FY13 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 1.0

Revision Date
August 20, 2012

PREPARED BY:
Don Holmgren, FNAL

CONCURRENCE:

August 20, 2012

William N. Boroski
LQCD-ext Contract Project Manager

Date

Lattice QCD Computing Project Extension (LQCD-ext)
Change Log: Alternatives Analysis for FY13 Procurement

Revision No.	Description	Effective Date
0.1	Document created from FY12 document.	July 22, 2012
0.2	Updates to all alternatives based on latest cost estimates	July 31, 2012
0.3	Expand scaling discussion and include metrics that take into account job sizes	August 6, 2012
0.4	Expand portfolio optimization discussion, add physics drivers	August 10, 2012
1.0	Final version after fixing typos and using “asqtad” consistently for staggered/asqtad/HISQ	August 20, 2012

Table of Contents

1 Introduction ... 1
2 FY13 Goals ... 1
3 Hardware Options ... 2
4 Alternatives ... 8
4.1 Alternative 1: A full rack of BG/Q deployed in Q1 2013 ... 8
4.2 Alternative 2: A half-rack of BG/Q deployed in Q1 2013, and a conventional cluster deployed late in FY13. ... 8
4.3 Alternative 3: A half-rack of BG/Q deployed in Q1 2013, and an accelerated cluster deployed late in FY13. ... 9
4.4 Alternative 4: A conventional cluster deployed in Q3 2013. ... 10
4.5 Alternative 5: A GPU-accelerated cluster deployed in Q3 2013. ... 10
4.6 Alternative 6: 50:50 (by budget) mixture of Conventional and GPU-Accelerated Clusters ... 10
4.6 Alternative 7: Status Quo (no additional deployment in FY13) ... 11
4.6 Other Alternatives ... 11
5 Discussion ... 11
6 Physics Drivers ... 17
7 Conclusion ... 18

1 Introduction

This document presents the FY13 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 ~$2.1M for capital procurements in FY2013. 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 FY07-FY12 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, after the acquisition of the first portion (66%) of the FY12 conventional cluster, but before the purchase of the FY12 GPU-accelerated cluster, these resources constitute an aggregate conventional cluster resource of about 50 teraflop/s sustained on LQCD benchmarks plus 660 GPUs deployed in 3 clusters. These 660 GPUs deliver approximately 90 effective teraflop/s based on the mixture of calculations run at JLab in the 2011-2012 allocation year; in GPU allocation units, they have a capacity of 4.7M GPU-hrs per year. The allocated project cost of operating these existing clusters in FY2013 is approximately $1.28M (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 FY13, 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 asqtad actions, with code for the clover action under development. Optimized multi-GPU codes are available for the asqtad, Wilson, clover, and twisted mass actions, with code for the DWF action under development.

2 FY13 Goals

The project baseline called for deployment in FY13 of conventional capacity totaling 44 TF. In FY11 and FY12, 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 FY13 goal was revised to increasing conventional capacity by between 15 and 22 TF, and accelerated capacity by between 4.6M and 6.9M "Fermi" GPU-hrs, equivalent to between 64 and 96 effective teraflops. Here, the effective GPU TFlops are based on FY11 allocations at JLab, and "Fermi" refers to the NVIDIA Fermi architecture. 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 FY13, the project will also decommission 8.4 teraflops of conventional cluster capacity (the JPsi cluster at Fermilab).

Sustained performance on conventional clusters is defined as the average of single precision DWF and improved staggered ("asqtad") actions 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 FY13 is to install these new resources by Sept 30, 2013 and release them to full production by Dec 1, 2013.

Beyond FY13, the objective is to take advantage of the improvements in technology implied by Moore's law, as well as the specific nature of LQCD calculations, to deploy a series of increasingly powerful resources for science.

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. An IBM BG/Q system, deployed at Brookhaven National Laboratory.
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

To date, benchmarks of DWF and asqtad 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 asqtad performance is 26 GF/node. Note that these figures correspond to 29% and 12.7% of peak, respectively.

Based on a cost estimate from BNL, the price of a full rack (1024 nodes) of BG/Q to the project will be $2.736M. This includes G&A, the first year maintenance cost and one additional year (FY14) of maintenance ($248.3K), plus software, front end and service nodes, and switches. The cost estimate for a half-rack is $1.432M, including the FY14 maintenance cost of $114.8K.

In the FY12 LQCD-ext cost estimate, there is an assumption that the baseline deployment plan will be used, which calls for all FY13 and FY14 hardware purchases to occur at Fermilab. Scenarios with half- or full racks of BG/Q would lead to reductions in labor expended at Fermilab in those years, since less hardware would be purchased at Fermilab in FY13 and therefore less hardware would be operated there in FY13 and FY14. For this analysis, the resulting FY13+FY14 Fermilab labor cost reductions are approximately $192K in the case of the purchase of a half-rack of BG/Q at BNL, and approximately $385K in the case of the purchase of a full-rack. The reductions were estimated by using the project's cost model, varying the number of procurements executed in FY13 at Fermilab and thereby varying the operated node counts in FY13 and FY14. For the purposes of life cycle cost calculations, for the third year of operations we reuse the FY14 Fermilab cost reduction figure of $75K for a half-rack, and $151K for a full rack. These labor cost reductions offset the $70K per year in the BG/Q cost estimate for deployment and operations labor at BNL.

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 three exponential fits to different subsets of the data.

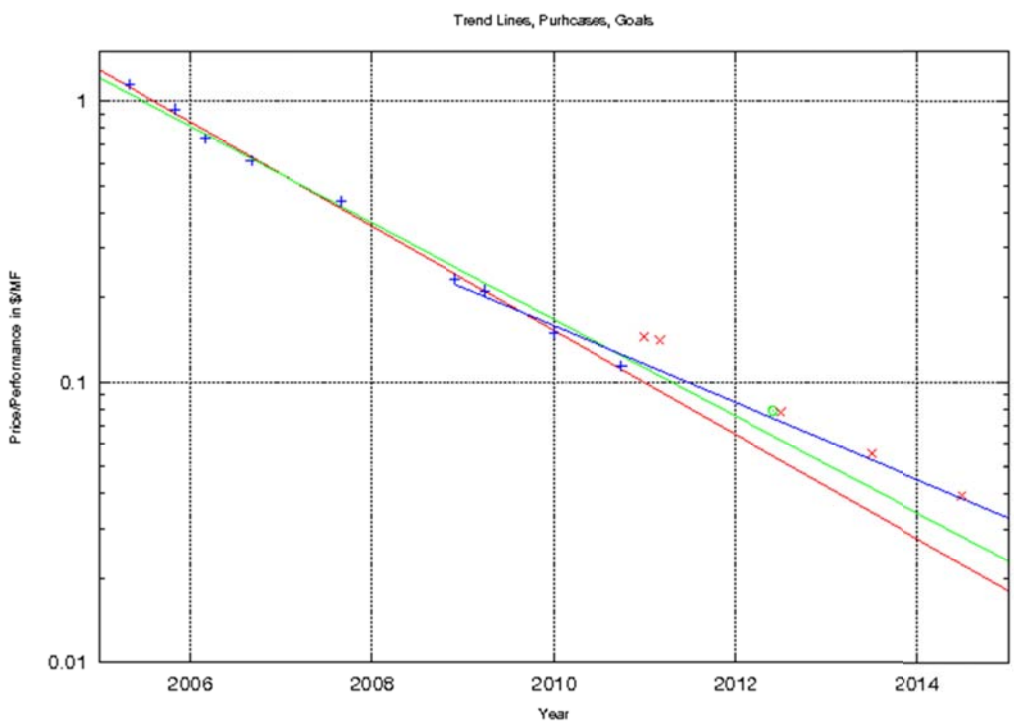

Here, the red line is the least-squares fit to the clusters purchased between 2005 and 2011, shown as blue "+" symbols. The red "X" symbols are goals used in the project plan. The green "o" is the JLab 12s cluster. The green line is the fit to the points from 2005 through 12s. The blue line is the fit to the points from late 2008 through 12s.

Using these fits, in mid-2013 the price/performance values corresponding to the three fit lines are $0.0343/MF, $0.0419/MF, and $0.0529/MF. An FY13 conventional cluster will likely have price/performance between the smallest and largest of these. For FY13 and FY14 purchases, quad-socket Intel “Sandy Bridge” and “Ivy Bridge systems will be available. The 12s cluster has dual “Sandy Bridge” sockets. The quad socket version will have lower cost per socket, and will have better scaling because more communication is local to each machine. Fewer Infiniband host channel adapters and switch ports are required (four processor sockets per HCA, rather than two per HCA on 12s). Since the equivalent job sizes will take half the node count on a quad-socket cluster compared to 12s, it is possible that performance can be maintained with modest Infiniband oversubscription.

Performance on conventional clusters varies between the various LQCD actions. On AMD-based systems, the difference between DWF and asqtad performance is relative small. For example, on the Ds cluster, DWF performance is 51.52 GF/node and asqtad performance is 50.55 GF/node. On Intel-based systems, the difference is considerably larger; on the 12s “Sandy Bridge” cluster at JLab, DWF performance is 56.5 GF/node, and asqtad performance is 43.74 GF/node. The $0.0419/MF middle trend line performance at mid-2013 uses the DWF:asqtad performance average. We can also estimate the DWF and asqtad price/performance values. For AMD-based systems, if performance follows Ds the estimates are $0.0415/MF for DWF and $0.0423/MF for asqtad. For Intel-based systems, if performance follows 12s the estimates are $0.0372/MF for DWF and $0.0480/MF for asqtad.

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 Dsg cluster at Fermilab, delivered in January 2013, has an effective throughput of 8.5 TF based on MILC f_π decay constant calculations. At a cost, including G&A, of $615K the price/performance of this cluster is $0.072/MF. There are a number of codes that achieve significantly higher throughput acceleration, such as those that do not require double precision, so $0.072/MF is a useful ceiling value. By mid-FY13 (14 months after the delivery of Dsg), assuming a similar 16-month halving time as the red trend line in the plot above, a GPU-accelerated cluster would achieve a ceiling of $0.040/MF for similar double precision calculations.

The following technology improvements support this increase in GPU-accelerated throughput for LQCD calculations:

- Such clusters would be based on systems and GPUs supporting PCIe Gen3, which provides higher bandwidth movement of data to and from the GPUs.
- The next generation of NVIDIA GPUs, Kepler, will certainly be available in time. In fact, the FY12 GPU cluster to be delivered to JLab in late FY12 or early FY13 will likely be based on Kepler. Kepler will have higher memory bandwidth and better floating point performance than the Fermi-architecture GPUs used on the Dsg cluster.
- NVIDIA software improvements, such as direct communications between GPUs in a node, and direct communications between GPUs and Infiniband interfaces, will improve the performance and scaling of parallel GPU codes for LQCD.
- Software improvements in the USQCD QUDA library will increase performance and the number of LQCD actions suitable for running on accelerated clusters.

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 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.

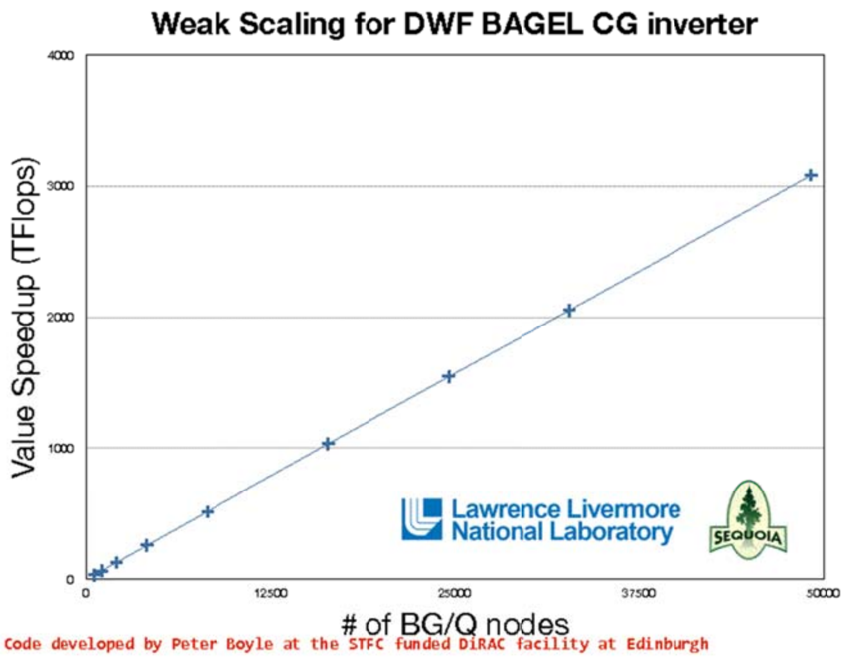

Because of longer network latencies, a conventional Infiniband cluster does not have flat weak scaling, but rather performance per core drops as job core counts increase (in weak scaling, the work accomplished per core is fixed as the number of cores are increased). The plot below shows weak scaling on the asqtad 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.

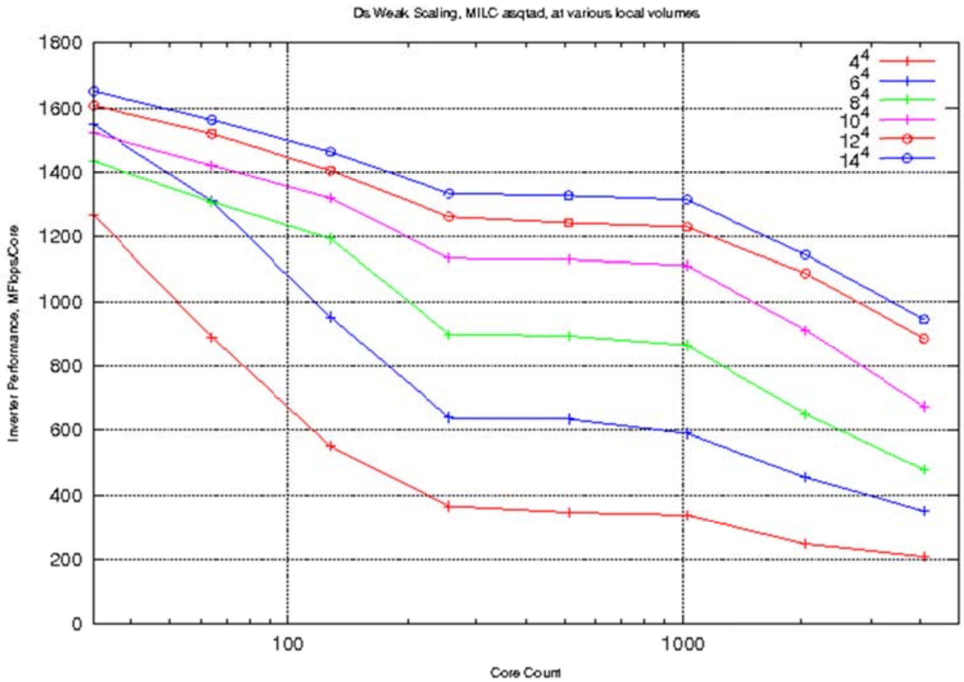

The USQCD cluster ratings use job sizes of 128. For local volumes of 14^4, performance per core of the asqtad inverter on jobs from 128 through 4096 cores on Ds is as follows:

- 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%.

For GPU-accelerated systems, weak scaling plots are not readily available, but the plots below show, respectively, strong scaling of the asqtad 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 asqtad 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.

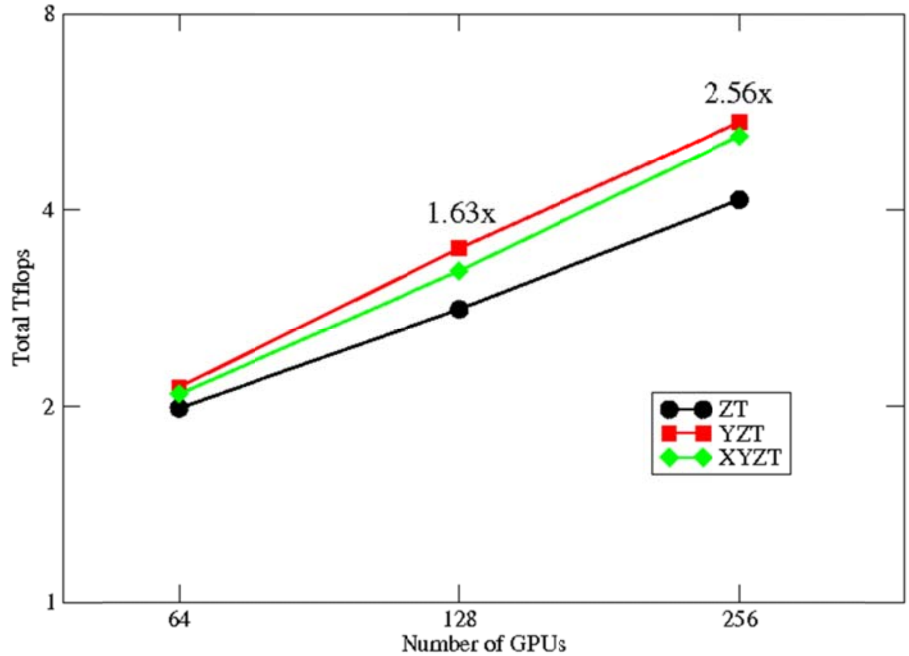

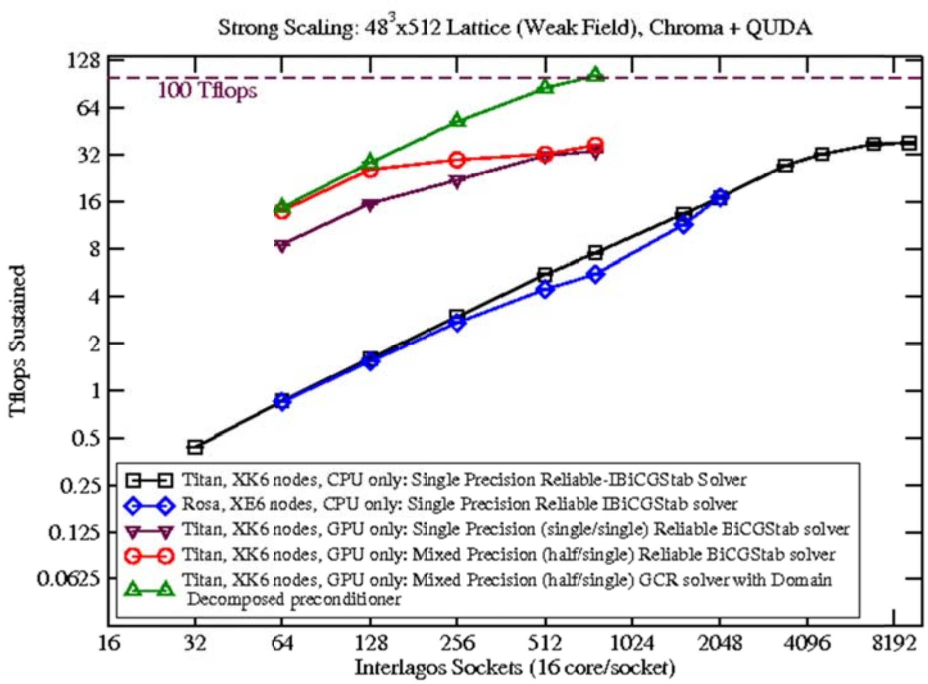

4 Alternatives

The following sections summarize the alternative technologies considered to achieve the some or all of the stated performance goals of this investment for FY13.

4.1 Alternative 1: A full rack of BG/Q deployed in Q1 2013

> *Deploy a full rack of BG/Q in the first calendar quarter of 2013 to sustain a total of at least 43.8 teraflop/s on the LQCD single precision benchmarks, for a total M&S cost of $2.73M.*

Based on the July 31, 2012 cost estimate, a full rack of BG/Q and necessary associated equipment would cost $2.736M, including G&A and maintenance contract costs of $248K total for the second year of operations (maintenance for year one is included in the purchase price). This system would have 1024 nodes and would achieve 42.75 GF/node, using the DWF:asqtad single precision performance average.

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

- Procure and install 43.8 TF in FY13 ($2.73M, including hardware maintenance)
- Operations at $70K per year ($210K for three years)
- Three-Year Lifecycle cost: $2.835M

Analysis: The three-year lifecycle costs in these alternatives only include the purchase costs of the hardware, two-year maintenance contracts or vendor warranties, and the incremental operations costs at the host laboratories over the three years. The costs do not include the power and space costs assumed to be contributed by the host laboratories.

The project hardware budget for FY13 is $2.09M. This alternative exceeds the available funds.

4.2 Alternative 2: A half-rack of BG/Q deployed in Q1 2013, and a conventional cluster deployed late in FY13.

> *Deploy a half-rack of BG/Q in the first calendar quarter of 2013 capable of sustaining at least 21.9 teraflop/s for a total M&S cost of $1.432M, and deploy a conventional cluster by the end of FY13 capable of sustaining at least 15.7 teraflop/s for a total M&S cost of $0.658M, for a total of 37.6 teraflop/s incremental capacity.*

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

- Procure and install a 21.9 TF BG/Q ($1.432M, including hardware maintenance)
- Operations of the BG/Q at $70K/year for a total of $0.210M
- Procure and install a 15.7 TF conventional cluster ($0.658M)
- Incremental procurement and operations costs for the conventional cluster for three years ($117K + $75K + $75K = $0.267M)
- Three-Year Lifecycle cost: $2.567M

Analysis: The hardware costs for this alternative are within the FY13 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 cluster deployment goal of 15 to 22 TF, and does not address the GPU-accelerated deployment goal of a new machine with annual capacity of 4.6M to 6.9M "Fermi" GPU-hrs.

For 128-core jobs, the price/performance of the BG/Q portion of the lifecycle cost is $0.075/MF, and that of the conventional cluster is $0.059/MF, with the overall price/performance $0.068/MF. For 2048-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.

For that fraction of the USQCD annual workload that requires large core counts per job and for which optimized BG/Q code is available, BG/Q hardware will have better price/performance than conventional cluster hardware.

4.3 Alternative 3: A half-rack of BG/Q deployed in Q1 2013, and an accelerated cluster deployed late in FY13.

Deploy a half-rack of BG/Q in the first calendar quarter of 2013 capable of sustaining at least 21.9 teraflop/s for a total M&S cost of $1.432M, and deploy a GPU-accelerated cluster by the end of FY13 of capacity at least 2.34M "Fermi" GPU-hrs/year, equivalent to 32.6 effective teraflop/s (16.5 effective teraflop/s for asqtad double precision) 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 ($1.432M, including hardware maintenance)
- Operations of the BG/Q at $70K/year for a total of $0.210M
- Procure and install a 2.34M GPU-hrs/yr accelerated cluster ($0.658M)
- Incremental procurement and operations costs for the accelerated cluster for three years ($117K + $75K + $75K = $0.267M)
- Three-Year Lifecycle cost: $2.567M

Analysis: The hardware costs for this alternative are within the FY13 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, but does meet the conventional cluster deployment goal.

The price/performance of the BG/Q portion of the lifecycle cost is $0.070/MF, and that of the accelerated cluster is $0.028/effective MF, with the overall price/performance $0.047/MF. The effective performance of 32.6 TF for the accelerated cluster uses the same mixture of GPU jobs sizes as seen on the JLab accelerated clusters in 2011. If we extrapolate instead 18 months from the MILC asqtad double precision analysis price/performance on the Dsg cluster, the projected double precision effective performance would be 16.5 TF. 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 $0.071/MF, with an overall price/performance including the BG/Q of $0.073/MF.

4.4 Alternative 4: A conventional cluster deployed in Q3 2013.

Deploy a conventional cluster by the end of June 2013 capable of sustaining at least 49.9 teraflop/s and costing a total of $2.09M.

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

- Procure 49.9 TF in FY13 ($2.09M)
- Incremental procurement and operations costs for three years ($234K + $151K +$151K = $0.536M total)
- Three-Year Lifecycle cost: $2.09M + $0.536M = $2.626M

Analysis: The hardware costs for this alternative are within the FY13 project budget. The overall price/performance for 128-core jobs is $0.053/MF. For 2048-core jobs, the price/performance of this alternative increases to $0.138/MF, assuming a 62% drop in performance per core relative to 128-core jobs.

This alternative exceeds the FY13 conventional cluster goal, and also exceeds the original baseline goal of 44 TF, but it does not address the GPU-accelerated cluster goal.

4.5 Alternative 5: A GPU-accelerated cluster deployed in Q3 2013.

Deploy a GPU-accelerated cluster by the end of June 2013 capable of delivering at least 7.44M "Fermi" GPU-hrs per year, equivalent to 104 effective TF (or 52.2 effective double precision TF), at cost of $2.09M.

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

- Procure a cluster with 930 "Fermi"-equivalent GPUs in FY13 ($2.09M)
- Incremental procurement and operations costs for three years ($234K + $151K +$151K = $0.536M total)
- Three-Year Lifecycle cost: $2.09M + $0.536M = $2.626M

Analysis: The hardware costs for this alternative are within the FY13 project budget. The overall price/performance is $0.025/effective MF, based on the performance of the mix of FY11 GPU allocated projects at JLab. Assuming a 60% drop in per GPU performance for large GPU-count jobs, the overall price/performance is $0.063/effective MF for 256-core problems.

4.6 Alternative 6: 50:50 (by budget) mixture of Conventional and GPU-Accelerated Clusters

Deploy a conventional and a GPU-accelerated cluster by the end of June 2013 capable of delivering respectively at least 24.9 TF, and at least 3.72M "Fermi" GPU-hrs per year, equivalent to 52 effective TF (or 26.1 double precision effective TF), at cost of $2.09M.

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

- Procure a 24.9 TF conventional cluster in FY13 ($1.045M)
- Procure a cluster with 465 "Fermi"-equivalent GPUs in FY13 ($1.045M)

- Incremental procurement and operations costs for three years ($234K + $151K +$151K = $0.536M total)
- Three-Year Lifecycle cost: $2.09M + $0.536M = $2.626M

Analysis: The hardware costs for this alternative are within the FY13 project budget. The 128-core job price/performance of the conventional cluster portion is $0.053/MF, and that of the GPU-accelerated cluster based on the performance of the mix of GY11 GPU allocated jobs at JLab is $0.025/effective MF, for an overall price/performance of $0.034/MF. For 2048-core jobs, the conventional cluster price/performance increases to $0.138/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 $0.063/effective MF for 256-GPU problems. The overall price/performance for such large jobs is $0.087/MF.

4.6 Alternative 7: Status Quo (no additional deployment in FY13)

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

The cost of this alternative is $1.28M in FY2013 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 new supercomputers that DOE and NSF will 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 untested.

- 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 of BG/Q). The full rack of BG/Q alternative will not be considered in this discussion, as the cost exceeds the project FY13 hardware budget.

Since a half-rack of BG/Q has a fixed price, all of the alternatives either direct the entire $2.09M FY14 hardware budget to Fermilab, or direct $1.432M to BNL for the BG/Q and $0.658M to

Fermilab. As covered at the FY12 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 FY13. 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. For calculations that require double precision and for which mixed precision algorithms nave not been implemented, the gain is much lower. At present, based on the ratio of allocation requests to available resources, conventional and accelerated resources are equally oversubscribed. The purchase of only GPU-accelerated hardware in FY13 would significantly over-provision this type of hardware, and the resulting USQCD hardware portfolio would not be balanced against demand.

Single GPU price/performance figures do not properly represent the wide range of acceleration observed across the various calculations and code bases. In contrast to conventional resources, where the difference in performance for asqtad and DWF actions does not exceed 30%, accelerated performance relative to conventional performance ranges over an order of magnitude. In the following discussions, we will use a range for GPU price/performance, the first calculated using extrapolations from the average effective performance observed on JLab 2011 GPU allocations, and the second calculated using extrapolations from the effective performance observed on the MILC asqtad allocation at Fermilab in 2012.

The table below shows price/performance for the various hardware options. For accelerated clusters, the two entries correspond to the JLab 2011 and MILC asqtad 2012 job types. 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. The large job estimates use degradation assumptions of 62% for clusters, 60% for GPUs, and no degradation for BG/Q. In this table, the BG/Q and conventional cluster performance estimates use the standard asqtad:DWF average. The BG/Q estimates use measured performance values (see Section 3 above), and the conventional cluster estimates use the middle price/performance trend line. The principle uncertainty for BG/Q performance is how well asqtad performance may improve from 12.7% of peak (developers estimate 15% to 20% of peak is achievable). The principal source of uncertainty for conventional cluster performance is the extrapolation for price/performance from the trend data; this analysis uses $0.0419/MF, and the achieved value will likely fall between $0.034/MF and $0.0529/MF.

Scenario	HW Cost	Ops Cost	Total Cost	Perf Small	P/P Small Jobs	Perf Large	P/P Large Jobs
				asqtad:dwf (TF)	asqtad:dwf $/MF	asqtad:dwf (TF)	asqtad:dwf $/MF
All BG/Q	$2.730	$0.210	$2.940	43.8	$0.067	43.8	$0.067
Note: Over Budget							
Half Rack BG/Q	$1.432	$0.210	$1.642	21.9	$0.075	21.9	$0.075
Conventional Cluster	$0.658	$0.267	$0.925	15.7	$0.059	6.0	$0.155
Overall	$2.090	$0.477	$2.567	37.6	$0.068	27.9	$0.092
Conventional Cluster	$2.090	$0.536	$2.626	49.9	$0.053	19.0	$0.138
Half Rack BG/Q	$1.432	$0.210	$1.642	21.9	$0.075	21.9	$0.075
Accelerated Cluster	$0.658	$0.267	$0.925	32.6	$0.028	13.0	$0.071
(GPU Double Precision)	$0.658	$0.267	$0.925	16.3	$0.057	6.5	$0.142
Overall	$2.090	$0.477	$2.567	54.5	$0.047	34.9	$0.073
Overall (GPU D.P.)	$2.090	$0.477	$2.567	38.2	$0.132	28.4	$0.217
GPU Cluster	$2.090	$0.536	$2.626	104.0	$0.025	41.6	$0.063
(GPU Double Precision)	$2.090	$0.536	$2.626	52.0	$0.051	20.8	$0.126
50% Conventional	$1.05	$0.27	$1.313	25.0	$0.053	9.5	$0.138
50% GPU	$1.05	$0.27	$1.313	52.0	$0.025	20.8	$0.063
(GPU Double Precision)	$1.05	$0.27	$1.313	26.0	$0.051	10.4	$0.126
Overall	$2.090	$0.536	$2.626	77.0	$0.034	30.3	$0.087
Overall (GPU D.P.)	$2.09	$0.54	$2.626	51.0	$0.052	19.9	$0.132

Considering small jobs only, the best overall price/performance figures, using the more conservative double precision GPU performance figures, are the conventional cluster, the GPU cluster, and the 50:50 conventional/GPU-accelerated cluster mix. Considering large jobs only, and again using the more conservative double precision GPU performance figures, the best overall price/performance figure is half-rack BG/Q plus conventional cluster scenario.

Optimizing the full USQCD portfolio of dedicated hardware should take into account the variations in performance across the hardware options of the LQCD actions of interest. As discussed in Section 3, parallel-GPU code for DWF action calculations is not currently available for GPU-accelerated clusters. Configuration generation has not yet been performed on GPU-accelerated clusters, and these clusters will exhibit strong scaling degradations. It is important to note that it is very probable that configuration generation will be performed on the GPU-accelerated resources at ORNL (Titan) and NCSA (Blue Waters) following the development of the necessary software; both of these resources will exhibit better strong scaling than GPU-accelerated clusters because of superior networking. Because of extensive software tuning efforts, BG/Q hardware has superior DWF performance compared to conventional clusters; because of the networking design, BG/Q systems also have much better strong scaling performance than either type of cluster.

The table below shows price/performance for the various hardware scenarios, but for the BG/Q resources, only DWF performance is considered. For the conventional cluster resources, the asqtad:dwf average is used, and for GPU-accelerated resources, the same range is used as before (double precision MILC asqtad performance, and performance based on the FY11 allocation mix at JLab). This shows the shift in price/performance under the assumption that the BG/Q might be used primarily for DWF calculations.

Scenario	HW Cost	Ops Cost	Total Cost	Perf Small	P/P Small Jobs	Perf Large	P/P Large Jobs
All BG/Q	$2.730	$0.210	$2.940	60.9	$0.048	60.9	$0.048
Note: Over Budget							
Half Rack BG/Q	$1.432	$0.210	$1.642	30.5	$0.054	30.5	$0.054
Conventional Cluster	$0.658	$0.267	$0.925	15.7	$0.059	6.0	$0.155
Overall	$2.090	$0.477	$2.567	46.2	$0.056	36.4	$0.070
Conventional Cluster	$2.090	$0.536	$2.626	49.9	$0.053	19.0	$0.139
Half Rack BG/Q	$1.432	$0.210	$1.642	30.5	$0.054	30.5	$0.054
Accelerated Cluster	$0.658	$0.267	$0.925	32.6	$0.028	13.0	$0.071
(GPU Double Precision)	$0.658	$0.267	$0.925	16.3	$0.057	6.5	$0.142
Overall	$2.090	$0.477	$2.567	63.1	$0.041	43.5	$0.059
Overall (GPU D.P.)	$2.090	$0.477	$2.567	46.8	$0.111	37.0	$0.196
GPU Cluster	$2.090	$0.536	$2.626	104.0	$0.025	41.6	$0.063
(GPU Double Precision)	$2.090	$0.536	$2.626	52.0	$0.051	20.8	$0.126
50% Conventional	$1.05	$0.27	$1.313	25.1	$0.052	9.5	$0.138
50% GPU	$1.05	$0.27	$1.313	52.0	$0.025	20.8	$0.063
(GPU Double Precision)	$1.05	$0.27	$1.313	26.0	$0.051	10.4	$0.126
Overall	$2.090	$0.536	$2.626	77.1	$0.034	30.3	$0.087
Overall (GPU D.P.)	$2.09	$0.54	$2.626	51.1	$0.051	19.9	$0.132

In the table above, considering small jobs only, the best overall price/performance figures, using the more conservative double precision GPU performance figures, are the half-rack BG/Q plus conventional cluster, the conventional cluster, the GPU cluster, and the 50:50 conventional:GPU cluster. Considering large jobs only, and again using the more conservative double precision GPU performance figures, the best overall price/performance figure is half-rack BG/Q plus conventional cluster scenario.

In communications from the USQCD Executive Committee and from the chair of the Scientific Program Committee, we have learned that in FY13 and later years, 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. Details regarding the anticipated projects are discussed in section 6 below.

To study the consequences of a mixture of large and small jobs types, we examine the portfolio of dedicated hardware under two scenarios: a half-rack of BG/Q and a conventional cluster, and a conventional cluster. The capacity of the hardware in these two resources will join with the existing Ds, 9q, 10q, and 12s clusters (considering conventional resources only). 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 asqtad-action jobs with 128-core jobs, and by the same 76% factor when comparing 4096-core asqtad-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. There are a pair of plots where the asqtad:DWF performance figure for BG/Q is used with 2048- and 4096-core jobs, and a pair of plots in which the DWF performance figure is used. 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.

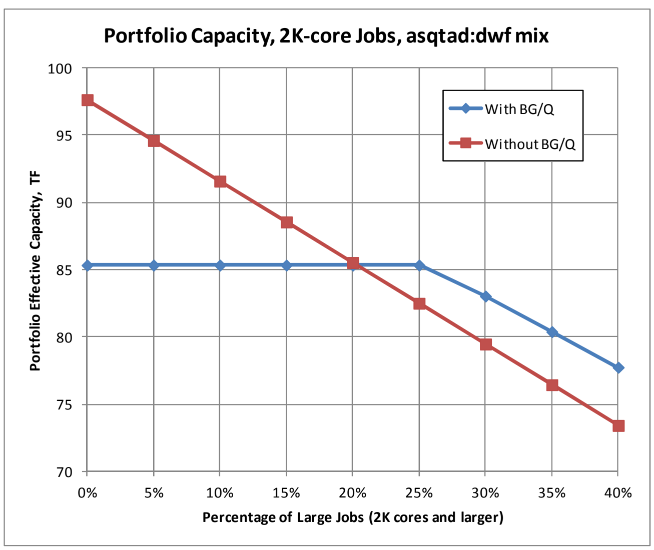

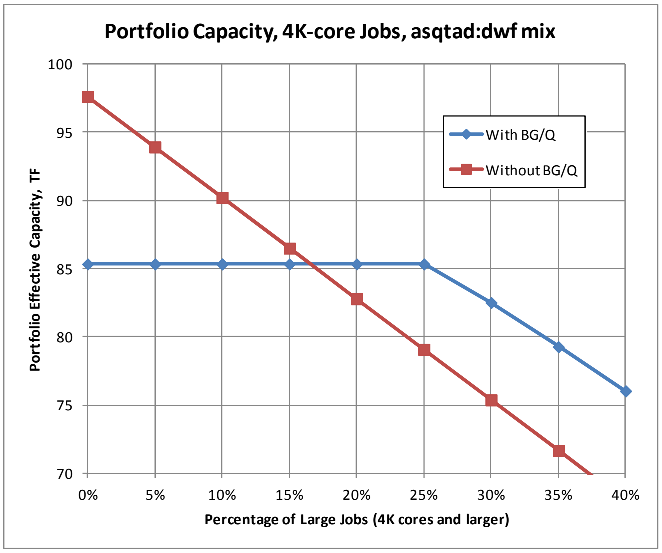
Portfolio Capacity, 4K-core Jobs, asqtad:dwf mix
Portfolio Effective Capacity, TF
100
95
90
85
80
75
70
0%
5%
10%
15%
20%
25%
30%
35%
40%
Percentage of Large Jobs (4K cores and larger)
With BG/Q
Without BG/Q

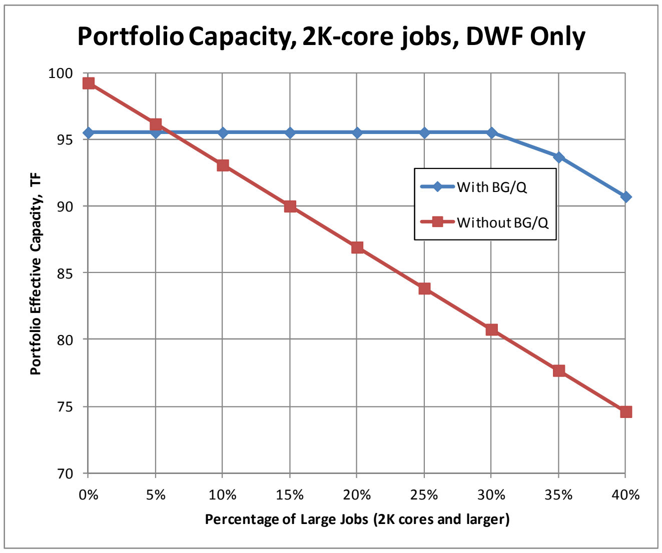
Portfolio Capacity, 2K-core jobs, DWF Only
Portfolio Effective Capacity, TF
100
95
90
85
80
75
70
0%
5%
10%
15%
20%
25%
30%
35%
40%
Percentage of Large Jobs (2K cores and larger)
With BG/Q
Without BG/Q

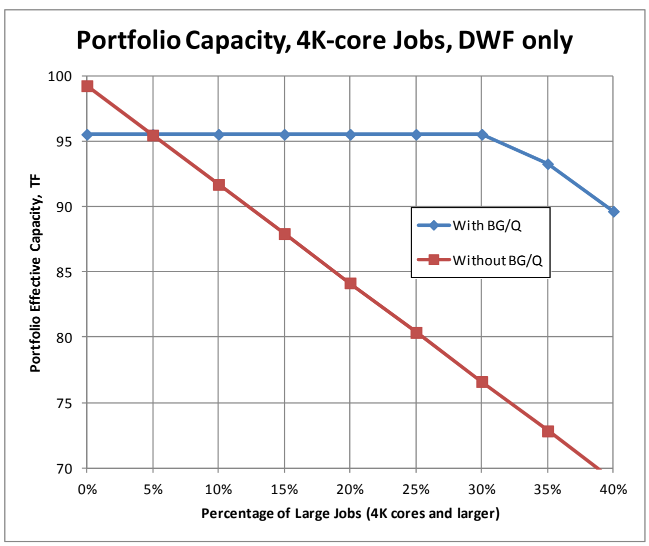
Portfolio Capacity, 4K-core Jobs, DWF only
Portfolio Effective Capacity, TF
100
95
90
85
80
75
70
0%
5%
10%
15%
20%
25%
30%
35%
40%
Percentage of Large Jobs (4K cores and larger)
With BG/Q
Without BG/Q

In all four cases in the plots above, for the conventional-only hardware ("Without BG/Q") as the large-job fraction increases, the effective capacity of the USQCD dedicated hardware portfolio continuously decreases. In the 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 (26% or 32%); 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 5% and 20%, then the effective capacity of the ensemble of USQCD dedicated hardware is larger for the BG/Q half-rack alternative. For the two cases that consider only DWF jobs on the BG/Q, the crossover in portfolio capacity occurs at approximately 5%. DWF allocations in recent years have exceeded this fraction. For the two cases that consider either asqtad or DWF large jobs for running on BG/Q, the crossover occurs at 20% or 16%, respectively, for 2048-core and 4096-core job sizes.

Because BG/Q hardware has such strong performance on the DWF action compared with conventional clusters, the purchase of a half-rack optimizes the overall computing capacity of the portfolio of USQCD dedicated hardware resources regardless of job size considerations under the supposition that as many DWF jobs as can be accommodated be run on that half-rack.

In the scenarios with only a FY13 GPU cluster or with a 50:50 mix of conventional and accelerated clusters, the new GPU hardware would be particularly suitable for MILC asqtad double precision calculations, since the NVIDIA "Kepler" hardware is expected to have relatively strong double precision performance compared with current "Fermi" GPUs, and because shifting MILC calculations preferentially to the GPUs would free the conventional cluster resources for DWF and other jobs. According to the figures in the tables above, even considering only small jobs, there is no clear price/performance advantage to either of these scenarios compared to one of the scenarios with a BG/Q half-rack. The half-rack scenarios have the additional advantages that the BG/Q hardware can absorb, if need be, work otherwise done on the conventional clusters involving any of the actions, and the BG/Q hardware has better strong scaling performance for large jobs. Scenarios involving GPU hardware have the disadvantage that current software does not support DWF production on the accelerated clusters.

6 Physics Drivers

The following paragraphs are drawn directly or paraphrased from communications with the USQCD Executive Committee, and with the Chair of the USQCD Scientific Program Committee.

USQCD has historically seen flat demand for computing capacity as a function of job size, with demand for roughly comparable numbers of core-hours in 1-node jobs, 10-node jobs, 100-node-jobs, and so forth. In the coming year, we expect approximately 20-25% of our demand in the job-size range of 2048 to 8192 cores. For jobs in this range, cluster performance is degraded by a factor of two to three, and consequently BlueGene hardware has a relative performance advantage. Because of the factor of ten increases in the scale of the resources at the DOE and NSF leadership centers in the coming year, the importance of USQCD dedicated hardware capacity at the upper end of its range will be larger than usual. 8192-core jobs will be discouraged on the ALCF BG/Q ("Mira"), and job sizes of less than 8 K cores are expected to be forbidden there.

There are several currently allocated projects that use job sizes in this range. An asqtad configuration generation project with lattice sizes 48^3x64 and 64^3x96 has been running this year at Fermilab. It ran with strong-scaling degradations of 2 to 2.5 relative to the rated performance of the Ds cluster, and it would have run more efficiently on a BG/Q. There are two currently allocated analysis projects for which we expect BG/Q performance to be similar to cluster performance: a DWF project (for which the most efficient code for BG/Q exists) on volumes of 48^3x96x24 and later 64^3x128x16, and a asqtad decay constant analysis project on volumes of 64^3*96 and later 96^3*192 (for which improvements in efficiency from the current 12.7% of peak are expected on the BG/Q via on-going software optimizations by members of the LQCD SciDAC-3 project).

The Executive Committee and the Scientific Program Committee believe that the demand for jobs in this "large job size" range will increase in the future. The future program will be formulated with the advice of the SPC, but these are some of the projects the Executive Committee foresees. A new isotropic clover configuration generation program has been approved and will start with lattice volumes of 32^3 and 48^3. Other planned medium-scale configuration projects include a domain-wall 32^3x64 ensemble with special boundary conditions chosen for study of the Delta I = 1/2 rule, and the zero-temperature 32^3x64 ensembles required for the analysis of the high-temperature QCD program. Planned large-scale analysis projects include correlation functions on high-temperature ensembles required to study di-lepton production in the quark-gluon plasma, clover fermion calculations of hadron structure, and creation of Dirac operator eigenvectors on domain-wall fermion ensembles. Many more such possible projects could be listed.

The USQCD Executive Committee believes that 25-30% of next year's program will run as well or better on a BG/Q as on clusters, and that the BG/Q would open possibilities at the large end of our job-size range that would not be feasible on clusters or GPUs.

7 Conclusion

Based on the assumption that at least 20% of USQCD production on the dedicated hardware resources in FY13 and following years will utilize large jobs (2048-cores or greater, corresponding to 25% or larger of a BG/Q half-rack), the alternatives that best optimize USQCD dedicated computing capacity are the BG/Q half-rack plus conventional cluster, and BG/Q half-rack plus GPU-accelerated cluster. This conclusion is further strengthened if allocated DWF projects are shifted to the BG/Q half-rack. We note that the budget available for the cluster ($0.658M) is too small to consider splitting into conventional and an accelerated cluster purchases. The decision of which of these two cluster types will be procured at Fermilab will occur by the start of calendar 2013, after additional information about USQCD requirements and available hardware performance and pricing are known.