LQCD Project Review Response

On May 24-25, 2005, a cost and schedule review of the Lattice QCD Computing Project was held at MIT, chaired by Dan Hitchcock of ASCR. The final report of the review committee was issued on June 27, 2005. On August 8, the project delivered its response to the review in a meeting at DOE in Germantown. This document contains written response of the project team. It is organized in the same sequence as the Hitchcock review report.

The text in italics following each response details actions taken since the project response was delivered in Germantown.

The significance and merit of the proposed initiative

Recommendation 1: In addition to exploiting existing opportunities, the group should facilitate exploratory studies in algorithms and comparative quantum field theory by allocating some time on the facility to this type of project. By comparative field theory, the committee means both variants of QCD (e.g., varying the number of colors and flavors, and quark representations, as well as quark masses) and also more radically different field theories (e.g., theories in different space-time dimensions, theories containing scalars, chiral gauge theories).

Response: This has been a long term scientific goal, and we will continue to allocate time for such studies. We have a very promising collaborative effort with the TOPS ISIC (David Keyes, adaptive multigrid) as part of our SciDAC work, and we will propose specific support for algorithm development in our upcoming SciDAC II proposal.

The FY 2006 allocations of project computing resources include a project titled "Investigations of twisted lattice supersymmetry" (Simon Catterall). This project is an example of an exploratory study of a comparative field theory of the type the Review Committee recommended we support. Allocations were also made to the projects "Improved Dynamical Chiral Fermion Algorithms" (Robert Edwards), and "All-to-all Propagators for Lattice Hadron Spectroscopy" (Jimmy Junge) which are exploratory studies of algorithms. The SciDAC-2 proposal for Lattice QCD Computing submitted in early March of 2006 includes algorithm development and exploration of comparative field theories in its work statement.

Recommendation 2: Visualization ought to be a powerful tool for understanding and finding surprises within the vast data set being generated. It also affords an opportunity to present the results to non-experts, including the interested public, in a memorable and attractive way. The team should develop a plan to incorporate specific visualization goals and approaches, as well as ensure sufficient visualization resources to make the approach feasible.

Response: Software development is not within the scope of this project. However, the SciDAC project plans to address this area. Further, as appropriate, acquisition plans will address visualization needs, providing the necessary hardware and (likely commercial) software infrastructure.

The SciDAC-II proposal for Lattice QCD Computing submitted in early March 2006 includes the development of visualization tools and techniques in its statement of work. This work will be centered at DePaul University and led by Massimo DiPierro, a lattice theorist who is an Assistant Professor in the School of Computer Science, Telecommunications, and Information Systems.

Recommendation 3: It is vital to the long-term health of the subject that young researchers get attracted into it. The team should consider ways in which this facility can be used to help the development of young researchers.

Response: We will continue to give high priority to proposals for computer time by young researchers. We will push to create new faculty positions and laboratory staff positions in our field, including joint appointments between the host laboratories and universities. A recent example was the appointment of Kostas Orginos, an outstanding young lattice gauge theorist, to a tenure track position by William and Mary/JLab. We plan to organize a series of summer schools in lattice gauge theory for graduate postdoctoral students. The Institute for Nuclear Physics in Seattle has agreed to host a summer school in 2007.

Of the 19 projects allocated project resources in FY06/FY07, 6 were submitted by researchers at the postdoctoral or assistant professor level: J. Dudek, J. Junge, J. Laiho, K. Orginos, J. Osborne, and P. Petreczky. Following discussions with the Lattice QCD Executive Committee the Theoretical Physics Program at the NSF indicated that it would create a set of five-year postdoctoral positions in lattice QCD. One award has been made, and a proposal for a second is under review.

In addition to the faculty position at William and Mary noted in our original response, another position was added with JLab support at Old Dominion University; both are in the field of nuclear theory, and both professors are using the LQCD facilities. Further, a bridge position at the University of New Hampshire, funded at 50% by JLab, was added in the field of LQCD.

The status of the technical design, including completeness of technical design and scope, feasibility and merit of technical approach and appropriateness and effectiveness of relevant R&D

Finding 1: The LQCD project presented a coherent four year plan for the acquisition and usage of computing resources for the LQCD community. The plan includes approximately equal investment in capability and capacity resources. The plan envisions adding additional capability resources over time and older resources would be utilized as capacity. The projected budgets and anticipated Moore's Law improvements in computational power should allow for the yearly acquisition of new clusters at about the same delivered performance on LQCD applications as the aggregate of existing computing resources.

Comment: In FY2006, the project begins with 5.8 Tflops of existing capacity and will add approximately 2.75 Tflops of new capacity. There are insufficient funds in any year of the project to add hardware matching the existing aggregate capacity; rather, roughly 25-30% additional capacity will be added each year.

The JLab "6N" cluster, which was released to production May 1, adds approximately 0.5 Tflops of capacity (0.3 Tflops funded by project funds, 0.2 Tflops funded by SciDAC and base funds). The Fermilab "Kaon" cluster, to be released to production by September 30, will add approximately 2.2 Tflops (1.9 Tflops funded by project funds, 0.3 Tflops funded by SciDAC and DOE supplemental funds). The total FY2006 incremental capacity will thus be approximately 2.7 Tflops, a fractional increase of 46%.

Recommendation 1: The committee recommends that the acquisition plan be modified to allow for a single joint acquisition, possibly every other year, alternating between the TJNAF and FNAL that would allow the delivery of resources to the program promptly in FY06 and beyond. The number of procurements should be reduced from eight to three or four.

Response: We agree that 8 procurements should be reduced to 3 or 4. Procurements will be a collaborative effort of the Project Manager and the Site Managers. In FY06, we propose that the cluster designed by the project be procured by FNAL. The project strongly feels that the cluster should be housed at FNAL because of their experience with Infiniband fabrics. The project also feels that it is critical that JLab gain experience with Infiniband, and recommends that JLab procure a 128-node cluster in FY05 with SciDAC and FY06 base funds, and extend this cluster to 256-nodes in FY06 with project funds; the resulting 400 Gflop cluster will meet the scientific needs of the approved DWF algorithm development and analysis of DWF quarks on asqtad lattices. In subsequent years the project will select the hardware (clusters vs. other supercomputers) and the location of the hardware in order to maximize the science according to the planned scientific program for the following year(s).

As agreed, this year JLab procured and brought online "6N", a 256-node cluster that has established Infiniband expertise at that site and which provided additional analysis capacity early in the year. Fermilab is procuring a large cluster to be brought online by the end of the fiscal year (500 dual Opteron nodes on project funds, and an additional 80 nodes on SciDAC and supplemental funds). The project plan to be presented at the May 25-26, 2006 review includes a single cluster procurement at JLab in FY07. At most a

single system procurement will occur in each of FY08 and FY09; if a mechanism can be found to combine the procurements, a combined FY08/FY09 purchase will occur in FY08.

Recommendation 2: If the FNAL construction schedule presented at the review, which delays the release of the computer there until September 2006, is accurate, the first computer delivered in FY 2006 should be put at TJNAF. If the revised FNAL schedule is accurate, which would enable the computer to be released to operation there in April 2006, the team should decide on the site for the computer based on where it can deliver the most science for the dollars invested.

Response: FNAL has committed to a schedule for the computer room refurbishment which will allow beneficial occupancy by April 2, 2006. We will follow the Program Manager's advice regarding the timing of the Federal Budget and will schedule the release of the RFP to first commit funds in fiscal Q2. Hardware delivery would therefore match the FNAL construction schedule. Further, we note that Intel roadmaps strongly favor delaying the procurement until mid-Q2.

As agreed at the Germantown meeting in August 2005, a large cluster would be procured and installed at Fermilab, likely based on emerging Intel hardware using fully buffered DIMM memory technology if such systems prove cost effective. This new memory technology promises increased memory bandwidth, which is critical to lattice QCD codes; however, systems will not be commercially available until June. The FNAL cluster will be released to production by the end of the fiscal year. The FNAL computer room refurbishment is scheduled to complete in mid-June, 2006. Procurement delays on this GPP project caused the later beneficial occupancy date. The RFP for the cluster was released at the beginning of the 3rd month of fiscal Q2 (March 3, 2006), with commitment of funds occurring in fiscal Q3 (May 19).

Recommendation 3: The cluster integration plan should be written down and an architectural diagram with hardware and software components clearly indicated. The plan should also include the software development and integration work items necessary to bring these resources into production. This plan should be presented to the LQCD scientific advisory board for review and approval.

Response: During the SciDAC project cluster designs were reviewed by the Oversight Committee, which included computing experts from outside LQCD. We will continue to follow this procedure and will also obtain the approval of the LQCD Executive Committee for each plan; this committee will have the responsibility of certifying that the plans fully meet the scientific requirements. The project plans will include the requested architectural diagrams as well as software development and integration details.

Integration plans for the JLab and Fermilab clusters were developed and included in the project WBS. The design of the JLab "6N" cluster, which deviated from the proposal

given by the project at the August 2005 meeting in Germantown (because of the use of dual core processors), was presented to the Change Control Board in December and approved by the CCB after the SciDAC prototype demonstrated better price/performance and good reliability. Design details of the Fermilab "Kaon" cluster were discussed at regular (weekly) meetings with the Chairman of the LQCD SciDAC Oversight Committee (Steve Gottlieb) and at biweekly project meetings which were attended by the Chairman of the Executive Committee.

Recommendation 4: The LQCD project plan should be expanded to identify dependencies on SciDAC and other projects for technology necessary for building the Metafacility. A clear set of Level 1 and/or Level 2 deliverables and milestones (e.g., single integrated login, single batch system, file and data sharing) for the Metafacility should be included in the plan. This will facilitate overall risk assessment and mitigation in the project.

Response: The WBS will be expanded appropriately, with Level 1 and/or Level 2 deliverables and milestones. We note that although the probability of occurrence of risks are low (SciDAC) to moderate (ILDG and other GRID developments), since the necessary software comes from projects external to this one, that the impact to the project deliverables of schedule slip are minimal and easily managed.

For FY06, only a limited Metafacility was planned, consisting of a common user runtime environment and utilities for transfer of file between the three laboratories. The single integrated login and batch system will not occur until at least FY08, and only if they improve the scientific productivity of the facilities. Metafacility-related milestones were included in the FY06 WBS (deployment of SciDAC libraries, deployment of common runtime environment, deployment of ILDG software). File transfer utilities and documentation were completed and deployed prior to the start of the project.

The feasibility and completeness of the proposed budget and schedule, including availability of manpower.

Recommendation 1: The project should consider alternative deployment strategies that result in fewer, larger systems over the same time period. This will reduce the required support effort to a feasible level within the project budget and associated subsidies.
An example of an alternative deployment strategy is to have single system delivery once a year, alternating between FNAL and TJNAF.
Because the facility work at FNAL was presented as being completed late in FY06, it appeared more effective to place a single larger system at TJNAF in early FY 06, and then a single larger system at FNAL in early FY 07. This would provide twice as much sustained computing between March 2006 to March 2007 as the schedule proposed by

the team. The team should use the amount of science delivered per dollar as the guiding principle for making system siting decisions.

Response: In FY06, a SciDAC Infiniband cluster at JLab similar to the FNAL FY05 cluster will be expanded, and a large FNAL cluster will be procured. In the subsequent years, 2 to 3 additional large procurements will occur, depending upon the timing of introductions of improved hardware to the market.

In FY06 at JLab an Infiniband cluster ("6N") was procured and released to production, providing approximately 0.5 Tflops of capacity (0.3 Tflops funded by the project, 0.2 Tflops funded by SciDAC). Fermilab is procuring a large cluster, to be integrated during fiscal Q4, which will provide an additional 2.2 Tflops of capacity. The project plans call for a single FY07 cluster to be procured and installed at JLab. In FY08 and FY09, at most one new system will procured per year.

Recommendation 2: The project should provide a cost benefit analysis for one site, two sites and three sites as part of the planning.

Response: The project will perform and include this analysis in the project plans.

A cost benefit discussion was prepared.

Recommendation 3: The cost projections for storage and consumables should be done to the same level as the costs for computational resources in order to ensure the user requirements are met in a balanced manner.

Response: We have gathered much additional information about the quantity and lifetime of data products and have modified the cost projections accordingly. The propagators discussed at the review, which take up most of the required storage, are intermediate data products that can be deleted 12-18 months after generation.

The LQCD collaboration uses an annual allocations process. Projects run for a year, from July 1 through June 30. Proposals for each year are due in February or March, and allocations are awarded in April or May. Each allocated project must provide disk and tape requirements. These are then used to make cost projections and plan budgets for the subsequent fiscal year.

Recommendation 4: The team should ensure wide impact of the valuable SciDAC-funded prototyping work with more timely publication of their results, both on the web site, but also in more widely shared publications and conferences. This effort should also seek out collaborations with other architectural and performance evaluation efforts.

Response: The project will increase the number of presentations and publications as recommended. We will also widen our collaborative efforts.

Norman Christ gave an invited talk on the QCDOC at the SciDAC 05 meeting titled "QCDOC: Project status and first results". It was published in the conference proceedings, J. Phys. Conf. Ser. 16, 129 (2005). Another article on the QCDOC, "Overview of the QCDSP and QCDOC computers," was published by the RBC/UKQCD groups, P. Boyle et al. in the IBM Research Journal 49, 351 (2005). There was a long list of papers presented at Lattice 05 reporting on initial research done with the QCDOC. Presentations related to SciDAC cluster hardware were presented at Lattice 05 and CHEP'06 (Mumbai, India). A BlueGene/L software workshop was held Jan 27 and 28 at Boston University (http://super.bu.edu/~brower/qcd-bgl/) which included several talks related to the SciDAC project.

Recommendation 5: The project team should reevaluate the principles used to determine which costs are included within the project to ensure an accurate presentation of the overall cost of the effort to DOE.

Response: The project has now produced effort and cost breakdowns showing the contributions from the laboratory base budgets, including power costs. We have also corrected the inconsistencies present in the earlier project plan of the lab/base effort breakdowns between the three labs.

Relevance of prototyping efforts outside the scope of the initiative and the status and plans for developing the required software for Lattice QCD computing.

Recommendation 1: The LQCD team should continue to monitor the market and benchmark the available options. The team should build on the existing collaboration of the participating labs in the prototyping effort to develop an integrated prototyping activity for LQCD. In the software area, the committee recommends the use of vendor-provided drivers to increase the communications performance (e.g., QMP over VAPI instead of MPI).

Response: In FY05/06, year 5 SciDAC funds will continue to be used for benchmarking available options. We will integrate this effort between FNAL and JLab, and as appropriate (e.g., storage hardware) with BNL. After FY06, SciDAC II funds, or base funds, will contribute to these activities, as they are of great interest and utility to the host laboratories. We will devote SciDAC-supported effort to implementation of the communications library (QMP) over VAPI; this work has begun.

During late FY05 and to date through FY06, the LQCD team has performed extensive hardware evaluations to support the JLab and FNAL FY06 cluster purchases. These evaluations included single- and dual-core single-socket Intel systems, dual-core dual-

socket Intel systems ("Dempsey" and "Woodcrest" processors), and dual-core dual-socket AMD systems. JLab and FNAL team members participated in these evaluations.

Efforts to implement a native version of QMP for Infiniband (QMP over VAPI) are in progress.

Recommendation 2: SciDAC support has been, and continues to be, absolutely essential for the success of the LQCD project. The leadership of the DOE in this area is recognized in the international lattice QCD community. Although this is somewhat outside the scope of this review, the committee recommends that the DOE consider continuing these efforts.

Response: The project enthusiastically concurs. Those project members involved in preparation of SciDAC II proposals will cite this recommendation.

The effectiveness of the proposed management structure.

Recommendation 1: Operations agreements with the sites over the lifetime of the project should be executed, which cover all contributions that are not included in the project cost, so that risks associated with escalation of operations costs can be reduced.

Response: We are preparing memoranda of understanding with the host laboratories. All such agreements will be referenced in our Project Execution Plan.

Operations agreements (Memoranda of Understanding) were executed between the project and each of the host laboratories.

Recommendation 2: The project should develop, and update on a yearly basis, a project-wide system deployment plan that optimizes the opportunity to deliver new science without artificial constraints on which programs can fund work at the three partner laboratories.

Response: The project has developed plans which are not constrained as noted in the recommendation.

The FY07 deployment plan will be reviewed May 25-26, 2006 at Fermilab by a DOE review panel. There are no constraints in this plan related to funding sources.

Recommendation 3: The strongly site-based management scheme reflects in part the history of forming this project. The laboratories should integrate their planning, prototyping and procurement activities. The approach to this should be in the revised PEP.

Response: Prototyping and procurement activities will be integrated and the approach described in the revised PEP.

The revised PEP does not explicitly discuss how the laboratories integrate their prototyping and procurement activities. However, in FY06 these activities have in fact been integrated. Specifically, FNAL and JLab jointly evaluated hardware for the JLab "6N" cluster. The application benchmarks used to evaluate hardware for both JLab and FNAL procurements were prepared by FNAL and JLab personnel, with JLab providing expertise on Domain Wall Fermion codes, and FNAL providing expertise on improved staggered action (asqtad) codes. Planning, prototyping, and procurement activities are discussed at the biweekly project conference calls which are attended by the site managers from JLab, BNL, and FNAL, the Chairman of the Executive Committee, and the Project Manager and Associate Project Manager.

Recommendation 4: Ensure the WBS is a tool for integrated planning as well as integrated reporting. The reporting should also document the actual physics output measured in terms of the allocations made by the Scientific Program Committee.

Recommendation 4.1: Incorporate schedules for integrated review of outyear plans into WBS to occur no later than June preceding beginning of FY.

Response: This will be done for FY07, FY08, and FY09.

The FY07 review will occur at FNAL on May 25-26, 2006.

Recommendation 4.2: Expand procurement processes to include all three sites and possible external experts, including evaluation of joint procurements.

Response: Procurements processes will include all three sites and when possible will be executed by a single site. Evaluation of proposed procurements will include outside experts (Oversight Committee).

No significant procurements are planned for BNL in FY06-FY09; rather, this laboratory will continue to operate the QCDOC, with small equipment purchases (approximately $20K) each year of storage and replacement parts. As discussed in the sections above, JLab and FNAL each made procurements in FY06. In FY07-FY09, at most one procurement per year will occur. In FY07, JLab will execute this procurement; evaluation of prototype hardware and planning will involve JLab and FNAL personnel.

Recommendation 4.3: Consider integrating technology tracking, hardware and software prototyping across all three sites.

Response: Cluster hardware and software prototyping will be integrated across JLab and FNAL. To the extent feasible, storage related hardware and software prototyping will be

integrated across all three sites; however, we note that at BNL the QCDOC currently only supports NFS-based disk systems. FNAL and JLab will use the same parallel filesystem (dCache or an alternative) and will integrate the related hardware and software efforts.

As noted above, no major procurements are planned at BNL in FY07-FY09. In FY06, JLab and FNAL jointly evaluated hardware for procurements. Starting with the JLab "6N" cluster and the Fermilab "Pion" and "Kaon" clusters, the systems at JLab and Fermilab are substantially compatible, and the labs cooperate closely on software evaluation (for example, the tracking and testing of Infiniband software). Prototyping activity for the FY07 acquisition will be coordinated at the project level and will occur at JLab and FNAL.

Recommendation 4.4: Since there are strong dependencies on some external efforts (SciDAC, ILDG, FNAL construction) schedule and contingency for these needs to be in WBS.

Response: We will include all such dependencies in the WBS and in the risk management strategy.

The WBS includes SciDAC, ILDG, and FNAL construction elements.

Recommendation 4.5: The laboratories should report the monthly progress of each laboratory in providing the capabilities and capacity agreed to by the Scientific Program Committee.

Response: The monthly reporting by the project will include this information as well as the actual monthly physics output (delivered flops by scientific project) at each site.

The monthly reporting to the program manager includes the delivered computing capacity at each lab (Tflops-yrs). The integrated node-hours delivered to the various scientific projects at the three labs are also tracked.

Recommendation 5: Consider moving metafacility operations to integrated project office.

Response: This will be done. In the organizational chart, the Metafacility Operations Manager (MFO) will report to the Project Manager; however, accounting will be done at JLab, where the MFO resides, similar to the site managers.

Metafacility operations have been moved to the integrated project office.

Recommendation 6: Charters for executive committee and Scientific Program Committee including how members are chosen should be produced and included in PEP.

Response: This will be done.

The PEP includes the charters and membership details of the Executive and Scientific Program Committees.

Recommendation 7: Try to integrate CCB and Scientific Program Committee review of change proposals.

Response: The Executive Committee will establish a permanent CCB for the duration of the project with at least three members. Members will be appointed by the Chair of the LQCD Executive Committee from the Executive Committee membership and from the user community. These members will not be a part of the participating Laboratories.

A permanent CCB was established, as detailed in the Project Execution Plan. The Scientific Program Committee is not involved in the review of change proposals.