

Strawman on Dates and Deliverables for Year 1 of SciDAC-2

Richard C. Brower

August 24, 2006

Comments: Please help to complete and correct.

1. I want this list to be sparse and crisp, but everyone's name should be included at least once against a short term deadline.
2. I have listed the Software Committee Member who should *monitor* the progress at the Laboratories and Universities.
3. When deadlines are met new ones will be added.
4. To keep it unclustered **not** all collaborators in multi-lab projects are listed.
5. The FTE #'s are rounded to half integers, but check them nonetheless.
6. Dates/Deliverables will also be refined by presentation at Software Workshop on October 13-14, 2006. (Agenda will be added later in Sec. 4.)
7. Further descriptions on each task are provided below excerpts from the SciDAC-2 Proposal (<http://physics.bu.edu/~brower/scidac-2>)
 - Appendix on **Tasks and Milestones** in Sec 3.
 - Software Schedule and Gantt Chart in Sec 2

1 Dates/Deliverables: Sept 1, 2006 to Oct 31, 2007

- BNL:** (Monitor Bob Mawhinney) 2.5 FTE
 Level 3 RHMC Asqtad over MILC (Chulwoo Jung & Enno Schloz) Dec 1, 2006
 Level 3 Dirac Operator on BG/L & BG/P (Chulwoo Jung & Pavlos Vranas ?) Dec 1, 2006
 CPS system over QMP When ??
 QLA port to QCDOC When ??
 Common Run Time env. (Stratos, Chip Watson, Don Holmgren) March 1, 2007
 Support for ILDG (Who ??) When ???

- FNAL/IIT/Vanderbilt:** (Monitor Don Holmgren) 5.0 FTE
 QLA implementation for Opteron/Cray XT3 (Don Holmgren) Dec 1, 2006
 Evaluation of native QMP on Infiniband (Jim Kowalkowski) Dec 1, 2006
 Multi-core support collab with JLab (Who??) When ??
 Prototype Automated workflow & integrations w. LQCD (ITT with Xien-He Sun) Sept 1, 2007
 Implementation of Cluster Monitoring and Diagnostic Software (Vanderbilt with Ted Bapty) Sept 1, 2007
 Common Run Time env. (Stratos, Chip Watson, Don Holmgren) March 1, 2007
 Support for ILDG (Jim Simone)

- JLab:** (Monitors Robert Edwards/Chip Watson) 3 FTE
 Optimization of Chroma (QDP++) on BG/L and XT3 (Robert Edwards & Balint Joo) Jan 15, 2007
 Goals and standards for Level 4 code libraries (Balint?) March 1, 2007
 Algorithm development and optimization for on top of QDP++ On going
 Optimization for multi-core processors on the Intel/SSE3 architecture (Who ??) When ??
 Documentation and Training Workshop (Balint Joo) When??
 Code Testing framework (Who ??) When ??
 Data Analysis Library: Design and documentation (Who ??) When ??
 Common Run Time env. (Stratos, Chip Watson, Don Holmgren) March 1, 2007
 Support for ILDG (Balint Joo) When ??

- Boston University/MIT/DePaul University:** (Monitor Rich Brower) 3 FTE
 Implementation of QLA for BlueGene/L (Andrew Pochinsky & Joy Khoriaty) Jan 15, 2007
 Level 3 DW inverter for Mobius class actions (Andrew Pochinsky) May 1, 2007
 Double hummer enhanced gcc compiler for BlueGene/L (Andrew Pochinsky) Oct 1, 2006
 Implementation of QMP for BlueGene/L (James Osborn) Nov 1, 2006
 Implementation of QMP for BlueGene/P (James Osborn) Jan 15, 2007
 Multigrid for 4-d Wilson Fermions in 4-d (Mike Clark & TOPS) Jan 15, 2007
 Lattice Data Visualization — Toolkit working demo (Massimo DiPierro) Jan 15, 2007

- Arizona/Indiana/Utah:** (Monitor Carleton DeTar) 2 FTE
 Optimize Execution and Throughput for RMHC Asqtad on QCDOC (Dru Renner) March 1, 2007
 Convert MILC code over QLA library (James Osborn & Ludmila Levkova) March 1, 2007
 Optimize MILC for BG/P (James Osborn & Carleton DeTar) March 1, 2007
 Adapt QDP to MICL (Ludmila Levkova) When ??
 Task What ?? (Subhasish Basak) When ??

- North Carolina/Renci:** (Monitor Ying Zhang) 0.5 FTE
 PQDP profile of MILC and Chroma code: First pass (Rob Fowler & Ying Zhang) Jan 15, 2007
 Collaborations with DePaul and Renci on Visualization tools (Ying Zhang) on going

2 Software task schedule

(From SciDAC-2 Proposal)

The scheduling of software tasks is given in the Gantt chart of Fig. 1. The budget requests support for 15.7 FTE per year for the software effort. This is broken down among the three main divisions of software work as follows: 4.9 FTE for Machine Specific Software, 7.7 FTE to support Infrastructure for Application Code, and 3.1 FTE for Uniform Computing Environment. These resources are nearly doubled by the contributions of physicists and software engineers at the participating institutions with no direct SciDAC support. In Appendix 3 we briefly describe the tasks to be undertaken, the major milestones for the first two years, and the FTE assignments for each participating institution.

Figure 1: The schedule of software tasks: **In separate file.**

As indicated in the Gantt chart, our software project involves both near term milestones to provide time critical software components, and long term development tasks that are to be pursued on a continuous basis throughout the grant. Important software milestones to be achieved during the first year of the grant include: (i) Design of a multi-core library interface (QMC), and the evaluation of its performance relative to simply treating each core as a separate processor. If the decision is to go forward, then the software will be developed and integration with QLA and QDP will begin; (ii) Port of QMP to Infiniband and the BlueGene/L toroidal network; (iii) Optimization of essential components of QLA for the AMD Opteron and the IBM dual core PowerPC processor; (iv) Conversion of the basic modules of the MILC code to QLA/C, introduction of templates into CPS and their restructuring in Chroma; (v) Identification of physical attributes to be visualized, cataloging of relevant data sets and design of a prototype interface; (vi) Implementation of a basic common runtime environment; and (vii) Initial design work on the workflow software and the fault monitoring and mitigation system.

The tasks that we will pursue over the full five years of the grant are equally important. They include the continuous adaptation of the low level API (QMC, QMP, QLA) to new architectures; the integration of the higher level API (QDP) into the major, freely available application codes; the definition and expansion of the common QCD physics toolbox; the design of more sophisticated algorithms with critical components optimized; and the documentation, regression testing and distribution of software libraries. As the project proceeds critical evaluations of the cost/benefit of each task may substantially alter priorities and allocation of our FTE resources. The metrics for success are a continued growth in the number of application codes written in the ACD API, the performance of these codes, and the number of physicists who use the QCD API to obtain important research results. The growth in the user community is already very encouraging.

3 Tasks and Milestones of Participating Institutions

(From SciDAC-2 Proposal)

In this appendix we briefly describe the tasks that will be carried out by each of the collaborating institutions, the FTE budgeted for them, and indicate the major milestones for the first two years of the grant. More detailed descriptions of the tasks can be found in the work statements of the individual institutions that appear below.

BNL: BNL will continue to optimize software and implement new algorithms for the QCDOC. It will compile, install and test SciDAC software packages on this machine. BNL will continue the evolution of the Columbia Physics System (CPS) code. It will optimize the CPS for the BlueGene/L, and work on an implementation for the successor to the QCDOC. This work will continue throughout the project, although there will be greater emphasis on QCDOC software during the first three years, and on software for the QCDOC successor in the last two years. A total of 2.5 FTE is budgeted for this work.

FNAL: During the first year of the grant, FNAL will port SciDAC code from the Intel 32 bit to 64 bit environment, and will optimize the code for Opteron processors. During year one, it will also explore the approach for and determine the benefit of a native implementation of QMP over Infiniband, and if warranted, create the implementation. In collaboration with JLab and university researchers, FNAL will provide code to support multi-core processors. With computer scientists at Illinois Institute of Technology it will provide software for automated workflow, and with computer scientists at Vanderbilt it will create software to enhance the reliability of large systems. It will implement and/or deploy software to support the ILDG and other grid activities, and provide software support for the evaluation of new hardware. FNAL will work with JLab throughout the project to study commodity hardware for lattice QCD. During the first year of the grant, it will evaluate AMD Opteron processors and Pathscale Infinipath. Finally, it will assist in the overall project management. A total of 3.0 FTE per year is budgeted for these tasks.

JLab: In each year of this project JLab will carry out research aimed at improving algorithms and producing high performance code for the study of lattice QCD. During the first year of the project, JLab will focus on implementations and optimizations for multi-core processors and for the Intel/SSE3 architecture, and on support for data analysis activities. It will also expand the existing code testing framework, and provide enhanced user support in collaboration with other institutions via workshops, phone and email. JLab will work with FNAL throughout the project to study commodity hardware for lattice QCD. During the first year of the project, JLab will study the Intel dual core “Woodcrest” processor, and double data rate Infiniband fabrics. A total of 3.1 FTE per year is budgeted for these tasks.

Boston University: Boston University provides significant leadership for the project as a whole with Richard Brower serving as Software Coordinator and Claudio Rebbi as chair of the Scientific Project Committee. James Osborn of BU has special responsibility to develop the C implementation of QDP and work with collaborators at Arizona, Indiana and Utah to integrate it into the MILC code. He will also work closely with Andrew Pochinsky at MIT to optimize the QCD API for the BlueGene architecture. Brower and Rebbi are leading the physics side of the collaboration with TOPS to study multigrid methods for lattice QCD. A total of 0.97 FTE per year is budgeted for these tasks.

Columbia University: Columbia University will lead an international effort to design and prototype a specialized computer for QCD. During the first year, different design approaches will be studied, and a detailed report prepared describing the results of the study and proposing what is judged to be the best approach. During the second year, this approach will be pursued in greater detail, and a proposal will be submitted to the Executive Committee with a specific architecture, cost and schedule for design and construction. If this proposal is accepted, then the design and prototyping work will be pursued in subsequent years. A total of 1.0 FTE per year is budgeted for this project.

DePaul University: DePaul University will lead the design and development of a visualization tool for lattice QCD. Work will be done in collaboration with physicists involved in the project and with computer scientists at the University of North Carolina. The goals for the first year of the project are to identify and catalog the types of datasets to be visualized, identify appropriate smoothing and visualization algorithms, and develop a prototype interface. In subsequent years, plugins will be developed to read in the various types of datasets produced in lattice QCD simulations, and tools for manipulating the data in increasingly sophisticated ways will be created. A total of 1.08 FTE per year is budgeted for this effort.

University of Arizona, Indiana University and University of Utah: The MILC code is an integrated package of some 150,000 lines of scientific application code and a library of generic supporting codes, that is publicly available and widely used. Arizona, Indiana and Utah will work together to carry out a major overhaul of this code to exploit the advantages of the SciDAC software. During the first year of this effort, generic code that supports multiple science-specific applications will be converted to QLA/C to take advantage of its platform-specific optimizations. During the second year, key modules will be rewritten in QDP. Optimization and tuning of the RHMC algorithm, which is currently being incorporated into the code, will be carried out. The first production version of the algorithm will be made available by the end of year one of the grant. Production versions of the code optimized for the Cray XT3 and BlueGene/L will be completed during the first year of the grant, and multi-core and enhanced compiler improvements will be incorporated during the second year. As always, upgrades to the code will be made available to the lattice community as they are completed. Finally, improved documentation for the code will be produced and published on the web by the end of the second year of the grant. A total of 1.875 FTE per year is budgeted for this effort, divided approximately equally among the three universities.

Illinois Institute of Technology: Computer scientists at the Illinois Institute of Technology (IIT) will build a workflow management system for planning, capturing and executing LQCD analysis campaigns. This work will be done in collaboration with FNAL. During the first two years of the grant, a workflow system will be developed and integrated into the existing LQCD computing infrastructure, allowing users to describe their analysis campaign workflow through XML files or graphical interfaces, and submit them for execution. Next a scheduling system capable of interacting with the workflow system and the system performance monitor will be deployed. The final result will be an integrated workflow environment capable of handling multiple campaigns. A total of 1.083 FTE per year is budgeted for this project.

MIT: Andrew Pochinsky of MIT will lead an effort to optimize the QCD API for the BlueGene series of computers. During the first two years, the effort will focus on the BlueGene/L. The gcc compiler will be modified to make efficient use of the two arithmetic units on each processor. The QLA routines will be compiled with this modified compiler, and key routines will be hand optimized as required. A level-3 inverter for domain wall fermions will be written, and in collaboration with James Osborn of Boston University, an optimized version of QMP will be developed. This work will be aided by contract commitments made by IBM as part of the MIT purchase of a BlueGene/L. It is anticipated that in subsequent years these software developments will be extended to later models in the BlueGene line. A total of 0.925 FTE is budgeted for this effort.

University of North Carolina: Computer scientists at the University of North Carolina will develop a performance profiling library (PQDP) to analyze the performance of the MILC and Chroma codes during the first year of the project. During the second year, the PQDP will be validated by profiling the MILC code on a variety of HPC platforms, including the QCDOC, clusters, the BlueGene/L and the Cray XT3. In subsequent years the UNC SvPablo performance analysis toolkit will be extended to support analysis of C++ codes so that the PQDP can be used to study Chroma. Performance analysis will be carried out on both codes on a wide variety of HPC platforms, and a web-based performance database will be established. The goal is to optimize the performance of MILC and Chroma based on the collected performance data. Finally, UNC will work with computer scientists at DePaul on the visualization effort discussed above. A total of 0.5325 FTE per year has been budgeted for these tasks.

UCSB: As chair of the Lattice QCD Executive Committee Robert Sugar provides overall leadership and coordination of the project. UCSB will administer funds for travel not covered by grants to other participating institutions. These trips will include visits of collaboration members to participating institutions for joint work, and attendance at meetings directly related to the project. UCSB will also administer travel funds for Principal Investigators S. Sharpe and R. Sugar.

Vanderbilt University: Computer scientists at Vanderbilt will develop an automated fault monitoring and mitigation system for the large lattice QCD clusters being built at FNAL and JLab. This work will be done in collaboration with FNAL. During the first year, an integrated monitoring and control system will be designed using existing standards and tools. Also during the first year, a tool will be developed for definition of workflows, monitoring and mitigation actions, based on Vanderbilt's Generic Modeling Environment. This task will be closely coordinated with work at IIT. During the second year, model based generators will be developed to transform the designs into components and configurations for the runtime system. In subsequent years, refined versions of these tools will be developed. A total for 1.083 FTE per year is budgeted for this work.

4 Agenda for Software Workshop October 13-14, 2006