1. NAME OF INITIATIVE: Lattice Quantum ChromoDynamics Computing

List of major collaborating institutions (including non-US partners).

Laboratories: FNAL, BNL, TJNAF

Participants from these additional laboratories and universities: ANL, LANL, Washington U., Boston, NMSU, Duke, Columbia, Colorado, Utah, Kentucky, MIT, Pittsburgh, Illinois (Urbana), Florida International, Indiana, Florida St, University of the Pacific, Maryland, Ohio State, UC-Davis, UCSD, George Washington, Cornell, Carnegie Mellon, U. Washington, UCSB, Virginia, Yale, Baylor

2. SCIENTIFIC JUSTIFICATION:

Physics goals. How does it fit into the global physics goals for the entire field.

A major goal of the Department Of Energy (DOE) Office of Science is to identify the fundamental building blocks of matter and to understand how the forces among them give rise to the observed physical world. To this end, the Office of Science supports large experimental programs in high energy and nuclear physics investing approximately $750M per year in high-energy physics (HEP) and $400M per year in nuclear physics (NP). In addition, the National Science Foundation (NSF) provides approximately $70M per year in high-energy physics and $35M per year to nuclear physics. The purpose of the Information Technology (IT) investment described in this document is to provide the computational infrastructure needed to carry out research in theoretical physics in support of these experimental programs. Because of the direct relevance to weak decays of strongly interacting particles at BaBar (SLAC), the Tevatron B-Meson Program (Fermi National Accelerator Lab, FNAL), and the CLEO-c Program (Cornell) physics, roughly 60% of the investment in the high energy program will be directly impacted by these calculations. Because of the direct relevance of QCD to the hadron physics programs at CEBAF (Thomas Jefferson National Accelerator Facility, TJNAF) and RHIC (BNL) and the direct relevance of QCD to the quark-gluon plasma research at RHIC (Brookhaven National Lab, BNL), roughly 50% of the investment in nuclear physics is impacted by these calculations. There are two key goals: the first goal is to achieve results from theoretical calculations that are comparable to the experimental results thereby demonstrating an understanding of the science producing the experimental results. The second goal is to use the understanding of the science to provide guidance to the experiments, design next generation instrumentation and facilities and achieve scientific discoveries. To achieve these goals, high performance and cost-effective computational systems are required.

Important progress has been made towards understanding the fundamental laws of nature through the development of what is known as the Standard Model of High Energy Physics. The Standard Model forms the principal understanding for approximately half of nuclear physics. It provides fundamental theories of the strong, electromagnetic, and weak interactions, three of the four fundamental forces of nature. It has been successful in explaining a wealth of experiments conducted with particle accelerators and cosmic rays. However, knowledge of the Standard Model is incomplete because it has proven difficult to extract many of the predictions of quantum chromodynamics (QCD), the component that describes the strong forces of subatomic physics. The only means of doing so is through very large-scale numerical simulations within a framework known as lattice gauge theory. These simulations are necessary to solve fundamental problems in high energy and nuclear physics that are at the heart of the DOE's large experimental efforts in these fields. Major objectives of the experimental programs are to: 1) verify the Standard Model, or discover its limits, 2) determine the properties of strongly interacting matter under extreme high energy and density conditions, such as those that existed immediately after the "big bang" and are produced today in heavy-ion collision experiments, and 3) understand the structure of protons and neutrons and other strongly interacting particles. QCD simulations are essential to research in all of these areas. Computers sustaining tens of teraflop/s will be needed over the next several years if the calculations are to reach the level of accuracy required to enable the Office of Science to effectively capitalize on the investments it is making in current experiments.

3. VALIDATIONS FOR SCIENTIFIC JUSTIFICATION:

Examples of recommendations and supporting statements from the committees, panels, and the community at large.

Two federal advisory committees have endorsed the LQCD effort:

The High Energy Physics Advisory Panel (HEPAP) at its February 8-9, 2004 meeting agreed that the U. S. should maintain a world class program in lattice gauge theory in order to realize its investment in theoretical and experimental high energy physics. They agreed that the deployment of dedicated hardware is needed to match the accuracy of important experimental measurements, and recommended a long term deployment program.

Two Nuclear Science Advisory Committee (NSAC) reports endorse LQCD computational efforts. The *Report of the NSAC Subcommittee on Nuclear Theory* in its report dated November 2004 stated, "We urge the funding agencies to actively seek new resources, such as SciDAC or new initiative money to fund these facilities", and recommended support of this effort, even under flat budgets.

The need for this work is also reflected in the April 2002 NSAC Long Range Plan, "Opportunities for Nuclear Science":

> "Advances in computational physics and computer technology represent great opportunities… To exploit these opportunities, dedicated facilities must be developed with world-leading computational capabilities for nuclear physics research.

4. DESIRED SCHEDULE:

List major milestones (month & year) such as design complete, construction start, construction complete etc.

(I include near term work which is funded by SciDAC and FNAL program)

12/04 256-node Infiniband Cluster construction complete
03/05 256-node cluster expansion construction complete
06/06 1024-node cluster expansion construction complete

06/07 1024-node cluster expansion construction complete
06/08 1024-node cluster expansion construction complete

5. ROUGH ESTIMATE OF COST RANGES:

Whatever the best information available (e.g. $M +/-30~50%, $150~250M, etc.). Total cost range including non-DOE funding (if any other funding sources are assumed and if known, state from where and how much. Also indicate remaining R&D cost to go.

R&D cost/year: (FY05-FY08) M&S: $30K
FY05 (two 256-node clusters): $400K SciDAC funds (in hand)
$300K DOE Lattice Computing Initiative (LCI)
$150K FNAL in-kind FY05
SWF : $400K SciDAC funds (in hand)
< $100K FNAL program
FY06 (1024-node cluster): $2.2M +/- 20%, drawn from:
$255K SciDAC funds
$300K FNAL in-kind FY06
remainder DOE Lattice Computing Initiative
SWF: $430K SciDAC funds
remainder (~ $170K) FNAL + DOE LCI
FY07 (1024-node cluster): $2.2M +/- 20%
DOE LCI + FNAL in-kind
SWF: $800K FNAL + DOE LCI
FY08 (1024-node cluster): $1.6M +/- 20%
DOE LCI + FNAL in-kind
SWF: $1M FNAL + DOE LCI

6. DESIRED NEAR TERM R&D:

Major activities needed to be completed before start construction.

Design of lattice gauge computing clusters requires an ongoing R&D effort to evaluate emerging computing and networking hardware. Over the past 3 years the effort required has been about 3 man-months per year. R&D for the FY05 clusters consists of evaluation of a prototype Infiniband network fabric and computers built using PCI Express chipsets; both sets of hardware are in hand. R&D for FY05 will complete by 09/04.

During FY05 we will do R&D for the FY06 cluster. The effort again will be devoted to evaluating emerging computing hardware, assumed at this time to be dual core Intel-based commodity systems. Additional R&D effort will be required in the areas of distributed data storage, parallel file systems, and facility networking security. The latter is described in our recent OMB 300 submission and will require close cooperation among facility managers at FNAL, BNL, and TJNAF.

R&D efforts in FY06 and FY07 are similar to those described above.

7. BRIEF DESCRIPTION OF LABORATORY'S ANTICIPATED ROLE:
Expected unique capabilities to be provided by lab. Rough estimate of human resources from lab (#FTE in what type labor).
Fermilab will be one of three national labs, along with BNL and TJNAF, to house and operate the lattice gauge computing facilities described above. The BNL site will operate the U.S. QCDOC 5 TFlops supercomputer, starting in early FY05. FNAL and TJNAF already house and operate 200 GFlops clusters at each site, with expansions in the next 12 months at each site bringing capabilities to roughly 1 TFlops. During FY06-FY08, Fermilab will each year replace equipment which has reached 3 years of age with new 1024-node or larger clusters, with the goal of establishing a facility of roughly 10 TFlops capability by FY07 with annual "refreshes" to replace the oldest 33% of computing hardware.

By FY06, the three laboratories will be jointly responsible for the operation of the distributed US Lattice Gauge Computing facilities.

Fermilab's responsibilities will be:

- Cluster R&D, design, procurement, and installation of cluster hardware
- Operation of cluster hardware, including all user support
- Joint management responsibility with BNL and TJNAF for the US Lattice Gauge Computing "meta" facility
- Stewardship of lattice QCD data housed in Fermilab mass storage facilities.

Because of Fermilab's long experience in distributed computing, driven in the past by the very large computing requirements for supporting the experimental program, the lab is uniquely qualified to design, house, and operate the significant lattice gauge computing facilities described here. The field of lattice gauge computing also heavily depends on the reliable storage of and high quality (high bandwidth, low latency) access to the very significant quantities of data (gauge configurations, quark propagators) which are the end products of the computations. The laboratory faces similar but much larger requirements for its participation in the LHC computing efforts and thus is very well suited to fulfill the needs of lattice gauge computing.

To date the modest lattice gauge computing facility at Fermilab (currently 376 systems) requires approximately 1.5 FTE to operate, with another 2 FTE devoted to the development of software. The software work is part of a large collaborative effort funded by the SciDAC Lattice Gauge Computing project. As the clusters at Fermilab increase in size to an assumed 3000+ systems, the required labor to operate the facility will increase to at least 4 FTE. All labor associated with this initiative are Computing Professionals.