National Computational Infrastructure for Lattice Gauge Theory

A DOE SciDAC Project

Lattice QCD Executive Committee

Richard Brower	Boston U.
Norman Christ	Columbia U.
Michael Creutz	BNL
Paul Mackenzie	FNAL
John Negele	MIT
Claudio Rebbi	Boston U.
Stephen Sharpe	U. of Washington
Robert Sugar (Chair)	UC Santa Barbara
Chip Watson	JLab

Collaboration and Partnerships

- Our collaboration includes nearly all senior lattice gauge theorists in the U.S.
  - Lattice gauge theorists, computer scientists, computer engineers
- Partnership with three DOE laboratories
  - Brookhaven
  - Fermilab
  - JLab
- Partnership with IBM
- Assistance from NSF PACI Centers
  - Pittsburgh Supercomputer Center
  - National Center for Supercomputer Applications

Participating Institutions

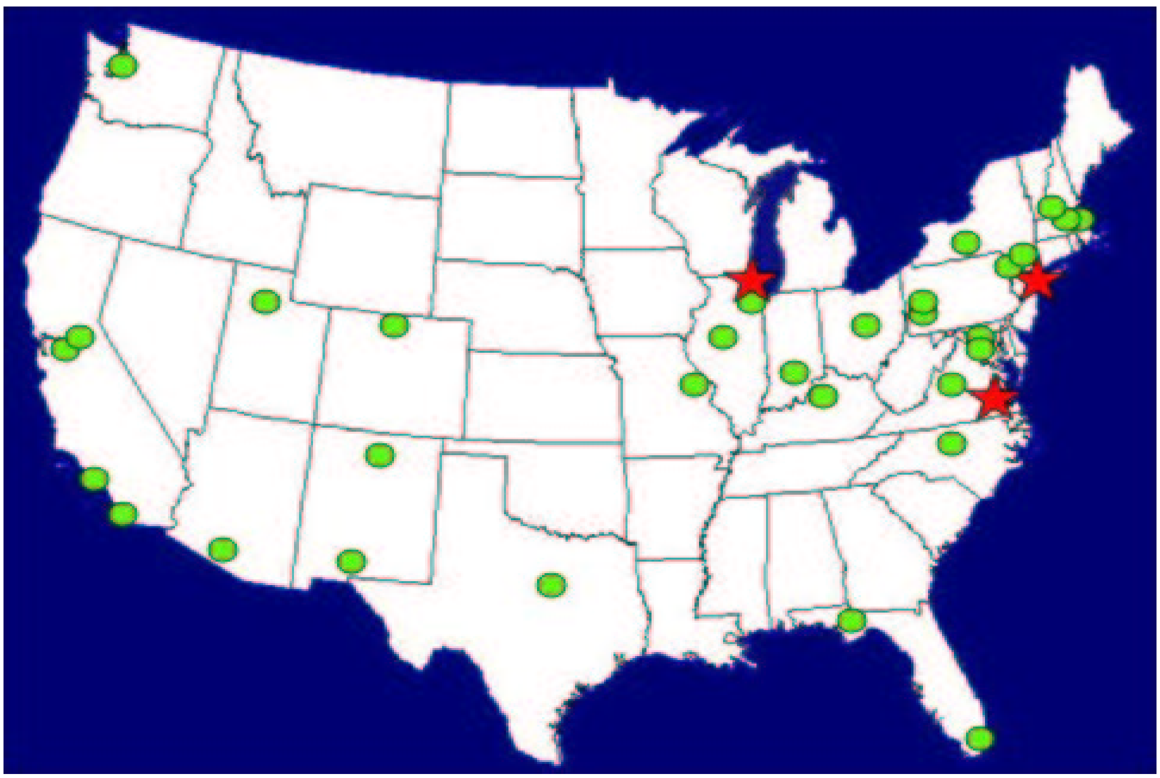

Objectives of High Energy and Nuclear Physicists

- Identify the fundamental building blocks of matter
- Determine the interactions among them

The Standard Model

- Quantum Chromodynamics (QCD)
- Electroweak Theory

Objectives of Lattice Gauge Theorists

- Understand the physical phenomena encompassed by QCD
- Make precision calculations of the predictions of QCD

Relationship to Experimental Programs

- **Major goals of the experimental programs in high energy and nuclear physics are to:**
  - **Verify the the Standard Model of High Energy Physics, or discover its limits**
  - **Determine the properties of hadronic matter under extreme conditions**
  - **Understand the structure of nucleons and other hadrons**
- **Lattice QCD calculations are essential to research in all these areas**
- **All major high energy and nuclear physics laboratories are impacted**
  - **Bates BNL Cornell FNAL JLAB SLAC**

Matching Computer Architectures to Scientific Problems

- **Strong tradition in lattice gauge theory**
- **Japanese Earth Simulator**
- **Argonne/LBL Blue Planet Proposal**
- **ORNL Cray X–1 Evaluation**

Special Purpose Hardware for QCD

- Simplifying features of lattice QCD calculations make building specially designed computers more cost effective than buying commercial ones
  - Uniform grids
  - Regular, predictable communications
  - Low memory requirements

- Approach pioneered by the Columbia group
  - Long series of successful machines
  - Gordon Bell prize for price/performance in 1998

- Approach adopted worldwide for Lattice QCD
  - Europe (APE)
  - United Kingdom (QCDOC)
  - Japan (CP-PACS/Hitachi)

Hardware Plans

- **Two hardware tracks:**
  - **QCD On a Chip (QCDOC)**
  - **Optimized Commodity Clusters**

- **Immediate goals**
  - **Price/performance of $1M per sustained Tflops in 2003**
  - **At least 10 Tflops sustained by 2004**

QCD Applications Program Interface

- **Unified programming environment**
  - **Enable very high efficiency on diverse multi–teraflops hardware**
  - **Allow physicists to focus on physics**
  - **Preserve investments in existing code**
  - **Produce portable code**

- **Three layer structure**
  - **Level 3: Highly optimized, computationally intensive subroutines**
  - **Level 2: Data parallel language to enable rapid production of efficient code**
  - **Level 1: Message passing and linear algebra routines**

Lattice Portal

- Provide U.S. Lattice Community access to all QCD data
  - Lattice archive with SQL-like storage and retrieval
  - File formats: XML metadata + Binary Data
  - Build on NERSC Gauge Connection Archive
  - Coordinate with ILDG on file formats, metadata, middleware
  - ORNL data is an example
- Parallel file transfers (bbftp, gridftp, jparss,...)
- Web based job submission and queuing standards
- Leverage PPDG SciDAC project
  - Build distributed U.S. infrastructure
  - Help create the International Lattice Data Grid

Current Status of QCD API

- Level 1 Message Passing (QMP)
  - Implemented in MPI, GM, QCDOC
  - In progress for Gigabit ethernet
- Level 1 Linear Algebra Routines (QLA)
  - Implemented in C and C++
  - Optimized SSE code for clusters
- Level 2 Data Parallel Language (QMP)
  - Implemented in C and C++
  - Optimization in progress
- Level 2 I/0 and Web Interface(QIO)
  - First release in Summer 2003

- **Level 3 Optimized Subroutines**
  - **Tested on QCDOC ASIC simulator**
  - **Being optimized in SSE for clusters**

- **Software Release and Documentation**
  - **`http://www.lqcd.org`**

Some Performance Highlights

- Results from gate–level simulator indicate QCDOC will achieve of order 50% of peak for optimized level 3 subroutines.

Estimate for Wilson Conjugate Gradient on $32^3 \times 64$ Lattice

Nodes	Sustained Teraflops	Percent of Peak
4096	2.1	52%
8192	4.2	51%
16384	8.1	49%
32768	15.6	48%

- Standard C code with MPI calls replace by QMP runs at 15% to 20% of peak on QCDOC simulator for very small local volumes

- Key subroutines run at approximately 1.5 Gflops on 2.0 GHz Pentium 4 with SSE instructions and QMP over GM

Conclusions

- SciDAC has brought the entire U.S. Lattice QCD Community together to develop the computational infrastructure it needs

- The software effort is creating a programming environment that will yield very flexible, high performance code

- We are in a position to build extremely powerful and cost effective computing platforms for the study of QCD

- The end product will be major progess in our understanding of the fundamental laws of nature