# Lattice QCD

Paul Mackenzie

# Main message

- The Fermilab lattice theorists and Computing Division are principals in the USQCD Collaboration. USQCD has overseen the DoE's program to develop hardware and software infrastructure for lattice QCD in the US.
- We are in the middle of an 06-09 program under an OMB300 to install lattice computing hardware at Fermilab and JLab. John Kogut and Robin Staffin have asked us to formulate our plans for 2010 and beyond for possible inclusion in the DoE FY2010 budget.
- We would like Fermilab's enthusiastic support for:
  - making the forthcoming 08/09 hardware installation at Fermilab a big success,
  - support at HEPAP and with DoE/HEP leadership for a continued, successful project with Fermilab leadership.

# Bottom line

- Fermilab has made essential contributions to the success of the current lattice project: several FTEs, especially key people in leadership positions, and infrastructure for the support of several thousand nodes of computing hardware.
- Fermilab has been the linchpin of the DoE's efforts to deploy and utilize computing infrastructure for lattice QCD. It is absolutely vital that this leadership role continue, both if we are to make scientific progress, and if the US is to remain competitive in lattice QCD.

# Plan of talk:

- The USQCD Collaboration
- Current LQCD hardware project
- Physics achievements of current project
- Future physics goals and needs
- Plans for achieving the goals
- What we need from Fermilab

http://theory.fnal.gov/theorybreakout2007// contains

- 2007 DoE review of the LQCD project,
- USQCD whitepapers on plans for 2010-2014.

# USQCD Collaboration

Consists of the great majority of US lattice gauge theorists, 85 members.
Purpose: develop hardware and software infrastructure for the US lattice community.
(Physics projects are done by individual groups within USQCD.)

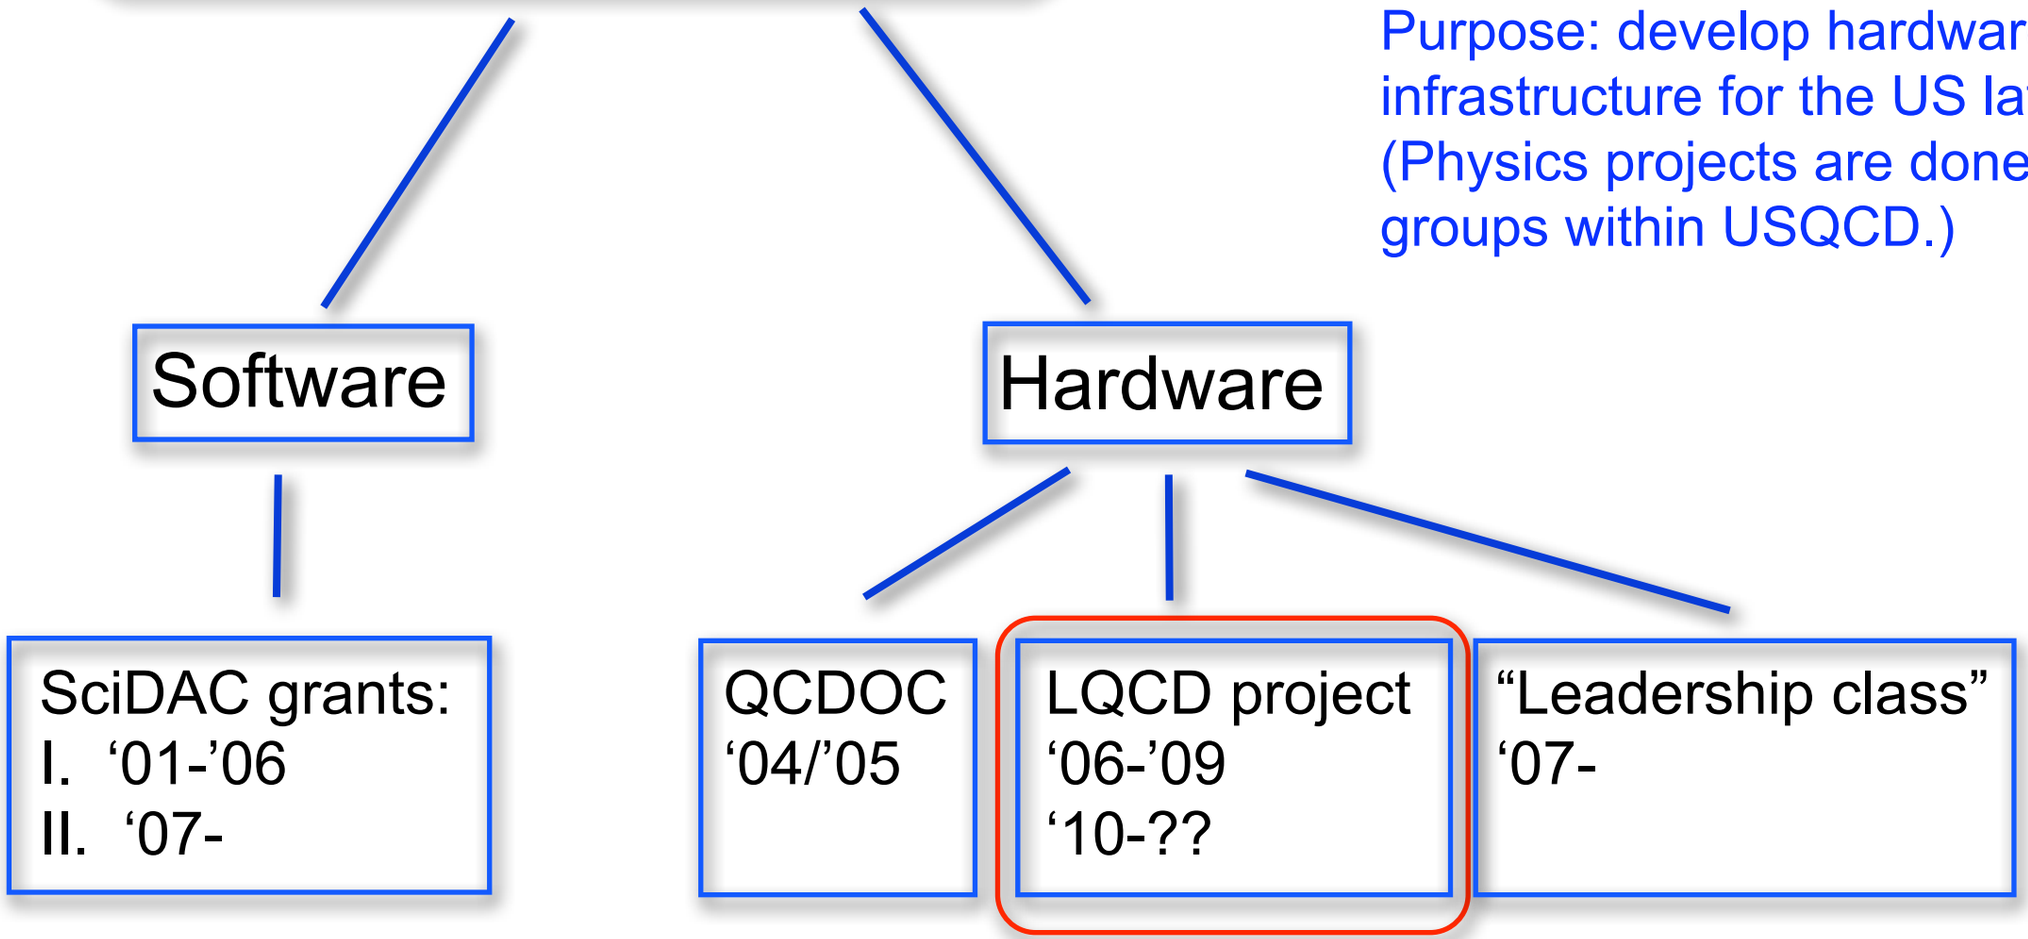


USQCD has grants for
- R&D for software and hardware through the SciDAC program.
- Hardware deployment from several sources, including the current LQCD project.

# USQCD Collaboration

# LQCD DoE Project

Federal project manager: John Kogut (HEP)
Federal project monitor: Sid Coon (NP)

Executive committee: Bob Sugar (chair), Rich Brower, Norman Christ, Mike Creutz, Paul Mackenzie, John Negele, Claudio Rebbi, David Richards, Steve Sharpe

Contract project manager: Bill Boroski
Associate Contract project manager: Bakul Banerjee
Original project manger and current technical lead, Don Holmgren

Scientific program committee: Andreas Kronfeld (chair)

Software committee: Rich Brower (chair)

Change control board: Bob Sugar (chair), Bill Boroski, Steve Gottlieb, Tom Schlagel, Vicky White, Roy Whitney

Fermilab employees (red)

The USQCD collaboration is funded through the LQCD project (as shown above), through SciDAC, and through base HEP and NP funds.

USQCD collaboration web page: http://www.usqcd.org

# USQCD timeline

USQCD Executive Committee formed.

First five-year SciDAC grant for lattice computing R&D.

Second five-year SciDAC grant for R&D.

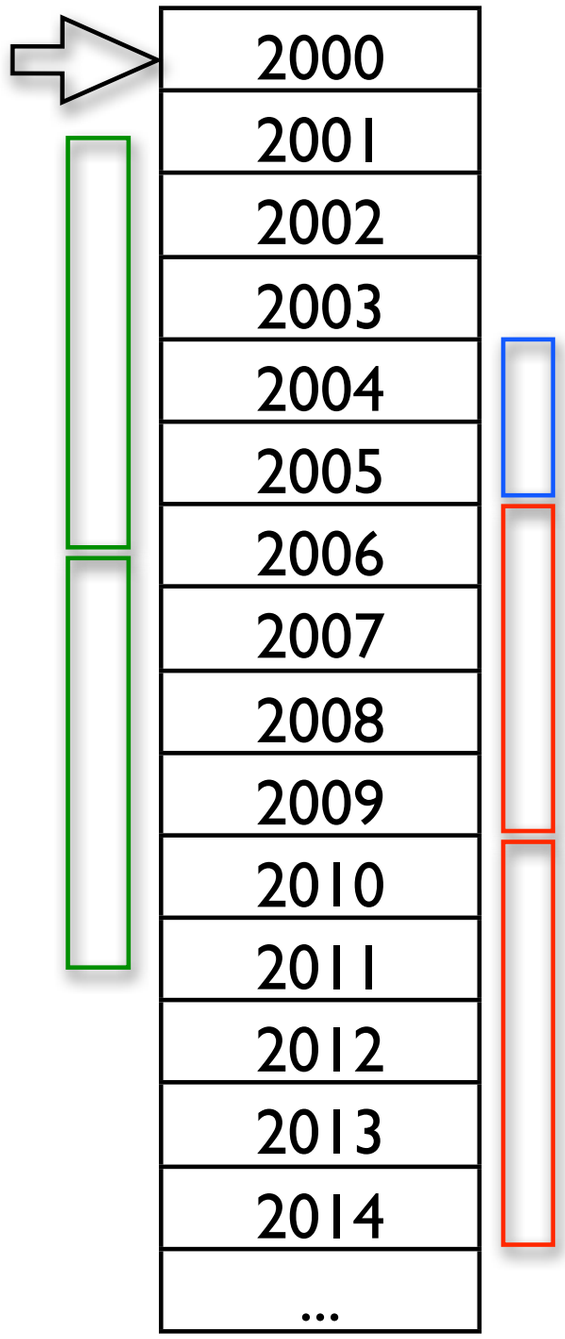


Construction of the QCDOC.

First cycle of continuous HEP and NP funding for hardware through LQCD project.

? We will ask the DoE for a continuation of the hardware project in 2010 and beyond.

# Current LQCD project

provides funds to

- Design and install new hardware for lattice QCD, such as the "Kaon" cluster at Fermilab.
- Operate existing lattice hardware: clusters at Fermilab and JLab, and the QCDOC at BNL (constructed before the official start of the current project).

*At Fermilab:*

Current lattice clusters are in the New Muon Lab.

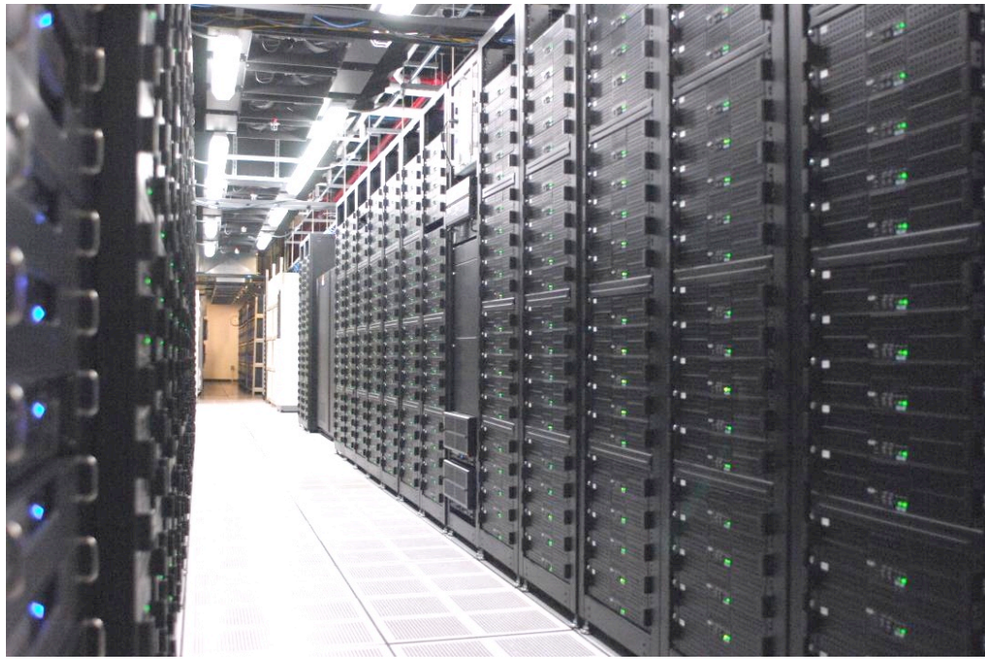

"Kaon" cluster, installed 2006.

# Total US teraflops for lattice QCD

17.5 TF at the end of the 08/09 installation.

Delivered teraflops using standard benchmarks of “asqtad” and “domain-wall” lattice codes, used in our reporting to the DoE.

(Multiply by 5 or 6 to compare peak teraflops used in advertising.)

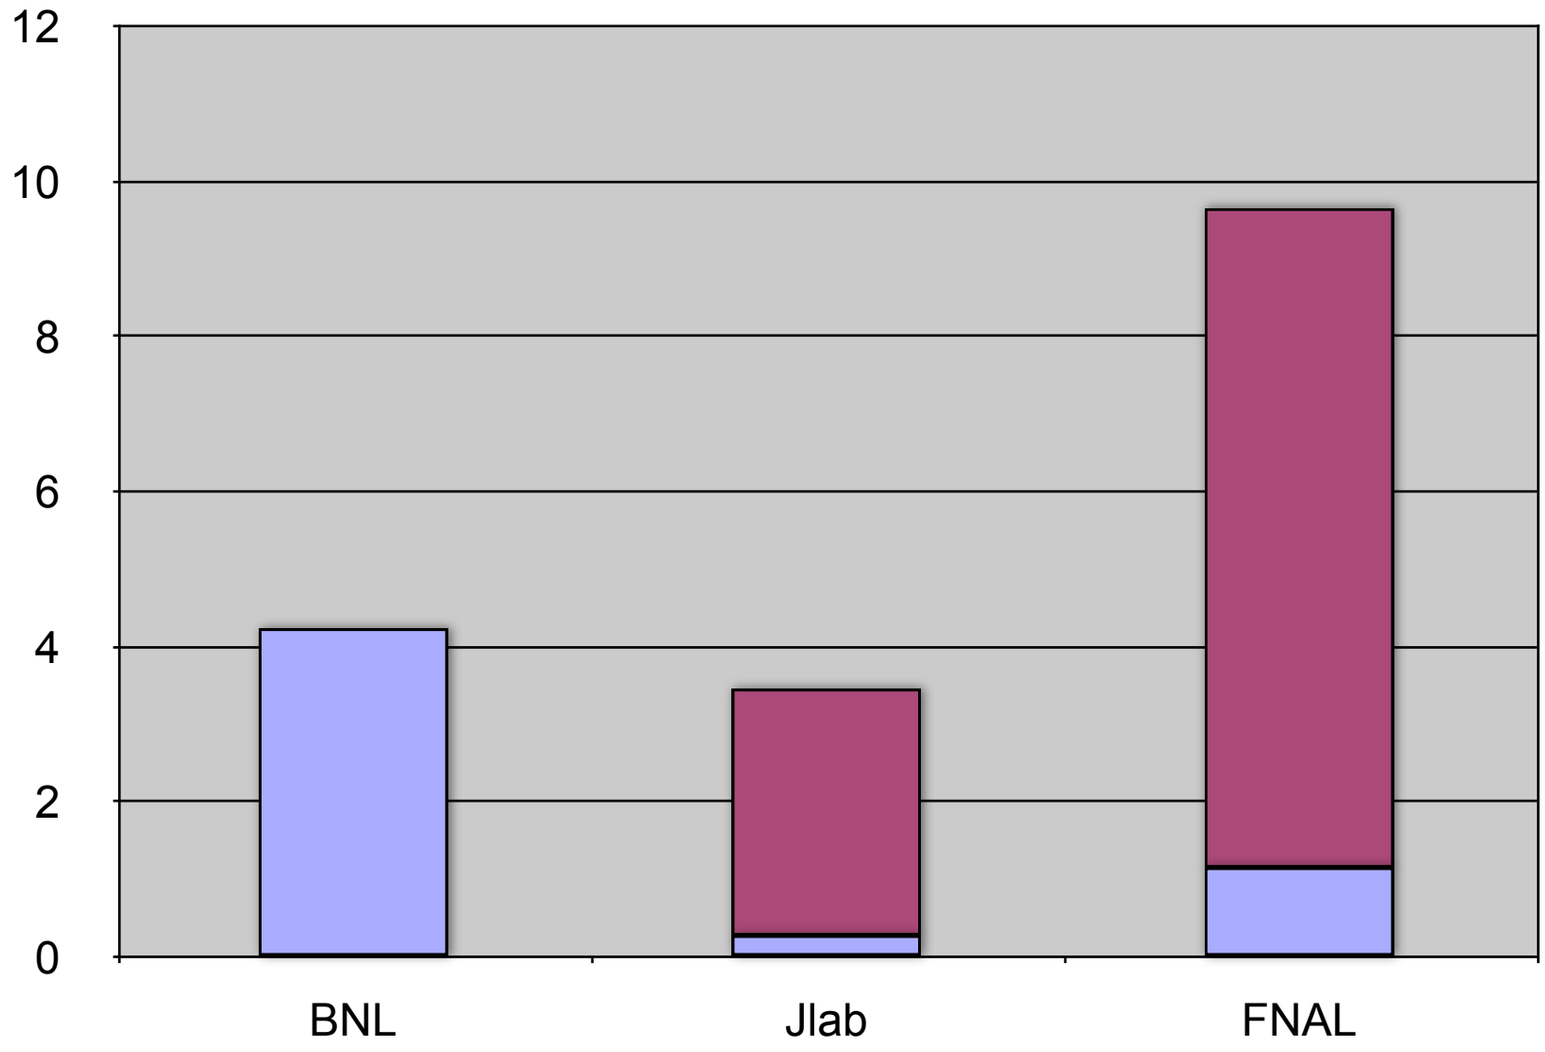


Maroon: constructed and operated through LQCD project,

Blue: operated through LQCD project, constructed with other DoE funding.

If the project does not continue in 2010- , funds to operate the computers must come from elsewhere.

# Structure of a lattice calculation

Few TB of data

Generate new gauge configurations on a QCDOC at BNL.
2 TF-years.

A single highly optimized program,
very long single tasks,
moderate I/O and data storage.

Transfer to Fermilab for analysis on clusters.
2 TF-years.

Large, heterogeneous analysis code base,
many small individual tasks,
heavy I/O and data storage.

Two comparable sized jobs with quite different hardware requirements.

# What this has cost:

## Total project scope and budget

- OMB Exhibit 300 IT investment funded to acquire and operate dedicated hardware at BNL, JLab, and FNAL for the study of quantum chromodynamics.
  - By the end of FY09, operate facilities with an aggregate capacity of 17.5 teraflop/s.
- Period of performance: FY06 through FY09.
- Budget: $9.2 million *(provided jointly by OHEP and ONP).*
  - *Covers project management, operations and maintenance of existing systems, and acquisition and deployment of new systems.*

| | FY06 | FY07 | FY08 | FY09 | Total |
|---|---|---|---|---|---|
| Hardware | $1,850 | $1,592 | $1,630 | $798 | $5,870 |
| Operations | $650 | $908 | $870 | $902 | $3,330 |
| Total | $2,500 | $2,500 | $2,500 | $1,700 | $9,200 |

*Shown in $K*

# Fermilab project funding profile

- Total Budget for LQCD OMB Exhibit 300 Computing Project:
  - Equipment: $5,870K
  - Operations and Project Management Support: $3,330K
  - Total: $9,200K
- Fermilab Allocation of Exhibit 300 Project Funds
  - Equipment: $3,913K *(67% of total equipment budget)*
  - Operations and Project Management Support: $1,520K *(46% of total ops budget)*
  - Total: $5,433K *(59% of total budget)*
  - FNAL budget allocation by FY ($K, including overheads.):

| | FY06 | FY07 | FY08 | FY09 | Total |
|---|---|---|---|---|---|
| Equipment | $1,548 | $25 | $1,585 | $755 | $3,913 |
| Operations and Proj. Mgmt. | $302 | $401 | $402 | $415 | $1,520 |
| Total | $1850 | $426 | $1,987 | $1,170 | $5,433 |

# Fermilab effort profile and in-kind contributions

- Level-of-Effort Profile, in FTE-yrs *(all funding sources shown)*

| Funding Source | FY06 | FY07 | FY08 | FY09 |
|---|---|---|---|---|
| LQCD OMB e300 | 2.1 | 2.5 | 2.5 | 2.5 |
| SciDAC | 2.8 | 2.8 | 2.8 | 2.8 |
| CD Base | 1.0 | 2.3 | 2.3 | 2.3 |
| Total | 5.9 | 7.6 | 7.6 | 7.6 |

- Additional Fermilab in-kind contributions
  - Computer facility upgrades to accommodate LQCD clusters
    - Lattice Computing Center at New Muon Lab (2/3 lattice use of $1.7M upgrade)
    - Grid Computing Center – Computer Room C (planned; 1/4 lattice use of $4.2M upgrade)
  - Electrical power to operate LQCD clusters
  - Infrastructure support
    - Networking, cyber security, mass storage, system disaster recovery, etc.

# DoE review of the LQCD project

May 14-15, 2007.
Review received July 9, 2007.
Maarten Golterman and John Kogut, co-chairs.

"The USQCD collaboration has been successful in providing an excellent large-scale computing infrastructure, through which a number of collaborations within the USQCD community has produced, and will produce first-rate scientific results. The infrastructure consists of the hardware installed at Brookhaven, Fermilab, and Jefferson Lab, as well as user-friendly software packages provided under the auspices of grants funded by the DOE SciDAC program. The organizational structure appears to be optimal in providing a platform for large-scale computing to the US lattice QCD community. In particular, the communication between scientists, technical staff and administrators associated with the project appears to be excellent, and highly efficient."

http://theory.fnal.gov/theorybreakout2007/LQCD_Project_2007_Review.pdf

# Cover letter from Robin Staffin and Dennis Kovar

July 5, 2007

Dear Dr. Boroski,

"...
You and the rest of the project management team were singled out for good planning, reporting, and hardware acquisition strategies and technical support. The scientific productivity and impact of LQCD were rated highly. Strong interactions with the experimental communities, involving flavor physics, nucleon structure and QCD thermodynamics were noted and encouraged to grow. The workshop which brought together members of the USQCD using LQCD clusters and the phenomenologists working on BaBar at SLAC was considered to be a model of possible future activities.
..."

Robin Staffin, Office of High Energy Physics
Dennis Kovar, Office of Nuclear Physics

http://lqcd.fnal.gov/directorsbriefing/LQCD_Project_2007_Review.pdf

# Statement of Bob Sugar

"Fermilab has been the linchpin of the DoE's efforts to deploy and utilize computing infrastructure for lattice QCD.  It is absolutely vital that this leadership role continue, both if we are to make scientific progress, and if the US is to remain competitive in lattice QCD."

MILC Collaboration
Chair of the USQCD Executive Committee
UC Santa Barbara
sugar@physics.ucsb.edu

# Statements from users: Peter Lepage

"Fermilab has been critically important to lattice QCD in recent years through its pioneering work on the cluster computers that are indispensable for the physics analysis part of lattice calculations --- that is, for the part that delivers the broad range of physics results that are directly useful to experimenters. Fermilab's continued leadership through 2010 and beyond will be essential if we are to continue to produce rich, high-quality (and highly relevant) results like those that have distinguished the US lattice QCD effort in recent years."

G. Peter Lepage
HPQCD Collaboration
Dean of Arts and Sciences, Cornell University
g.p.lepage@cornell.edu

# Statements from users: John Negele

Dear Paul,

As the LQCD Executive Committee plans for the next phase of the national LQCD Computing Project in 2010 and beyond, I want to be sure you and the Fermilab Directorate know how important I believe Fermilab's leadership has been for the success of the initial project, and how essential it will be for the next phase commencing in 2010.

Both in the precursor SciDAC lattice QCD infrastructure project and as the first project manager of the LQCD Computing Project, Don Holmgren played a seminal role in guiding and coordinating the entire national hardware effort at Jefferson Lab and BNL, as well as at Fermilab. Although my own work has addressed nuclear physics issues associated with the quark and gluon structure of the nucleon, it is absolutely clear that the quality of the national QCD resources that have been available to me and the level of support the DOE has provided to the whole field have been greatly influenced by Don's effort and leadership. Bill Boroski has seamlessly picked up where Don left off, and recently led our project through a highly successful DOE review. When Robin Staffin and Dennis Kovar write in their cover letter of the report to Bill Boroski "You and the rest of the project management team were singled out for good planning, reporting and hardware acquisition strategies and technical support", they speak for the whole community in praising Fermilab's leadership role in this project.

In looking to the future and preparing the case for DOE support of the next generation computer resources for lattice QCD, a key argument in convincing them of the outstanding leadership and technical excellence of the next phase of the project will be the strong continued commitment of the Fermilab management to this endeavor.

Best Regards,

John

John Negele
LHPC Collaboration
William A. Coolidge Professor of Physics, MIT
negele@mitlns.mit.edu

# Physics achievements

General physics goals of USQCD:

- Quark flavor physics
  - essential for B factories, CDF/D0, LHC-*b*
  - essential for high precision *K* and charm physics
- QCD thermodynamics
  - relevant to RHIC and to early universe
- Hadron spectroscopy and structure
  - relevant to JLab (also moments of pdfs=>Tevatron and LHC)

Lattice calculations seek to:

- Post-dict known experimental results
- Predict results to be confirmed by experiment
- Extract information unobtainable without the lattice

# Postdictions of known results

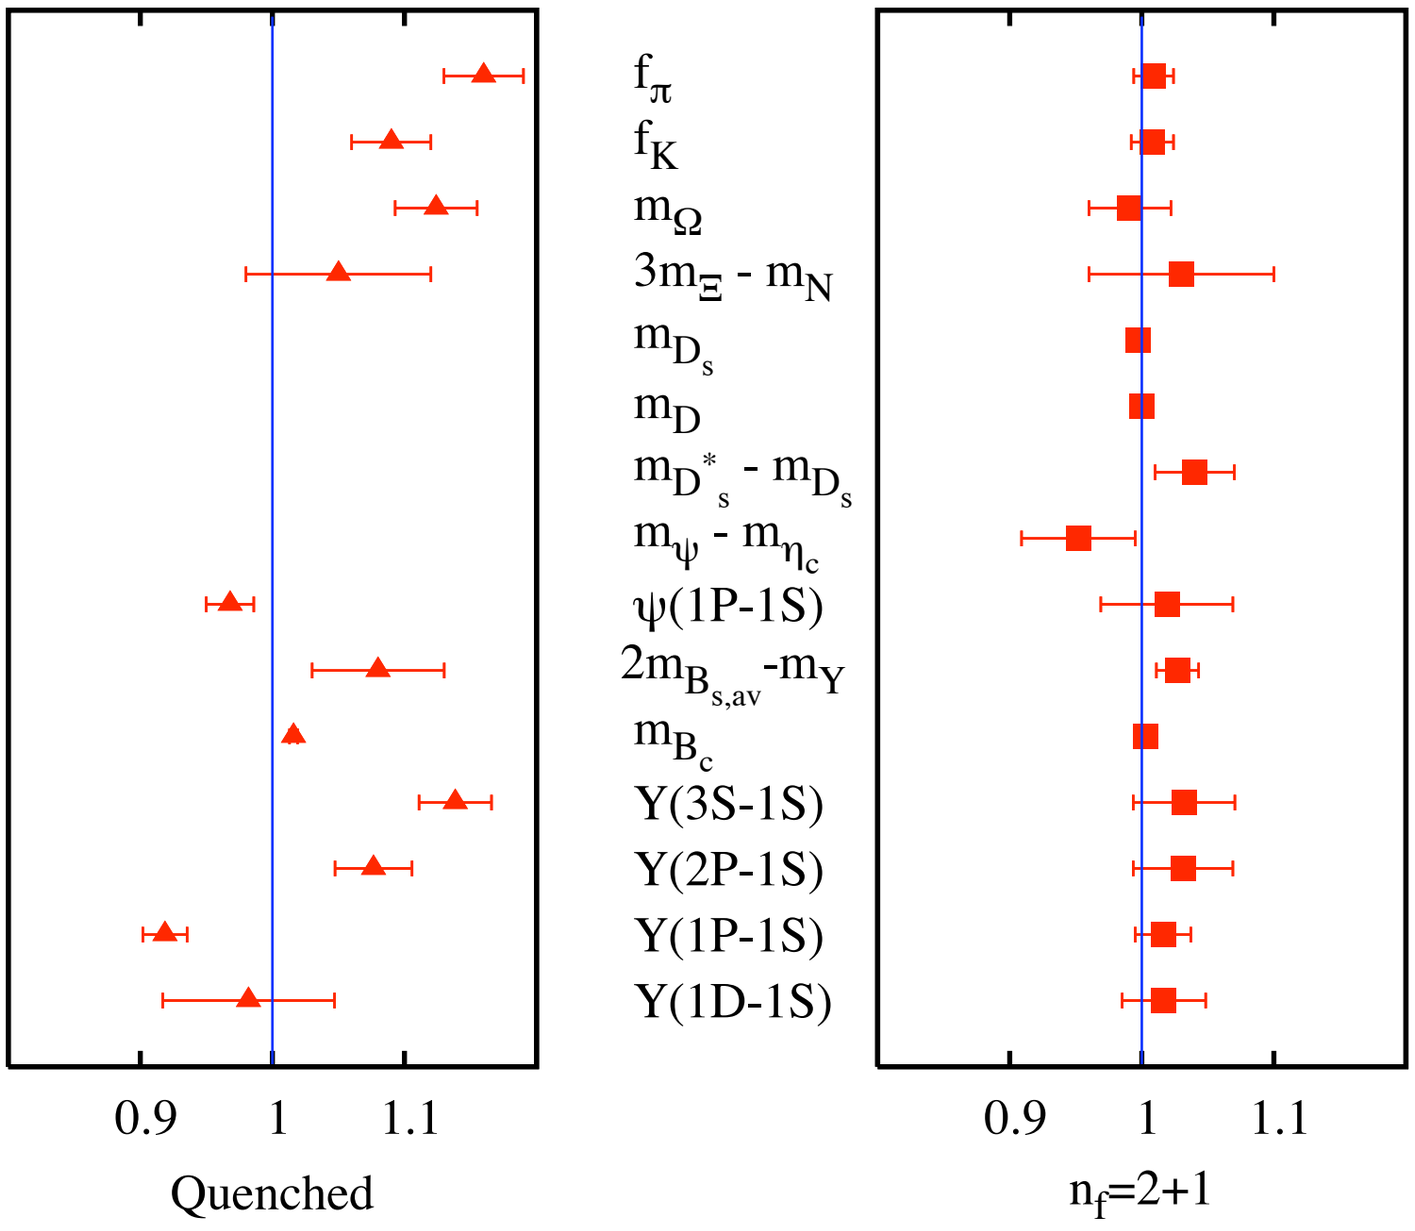


Old lattice calculations used the (computationally cheap) "quenched approximation" (left). This introduced errors of 10-15%, and no means of estimating them.

Modern "unquenched" calculations have removed these (right).

HPQCD, Fermilab, MILC
Phys. Rev. Lett. **92**:022001, 2004.

# Predictions: D Decay Constants

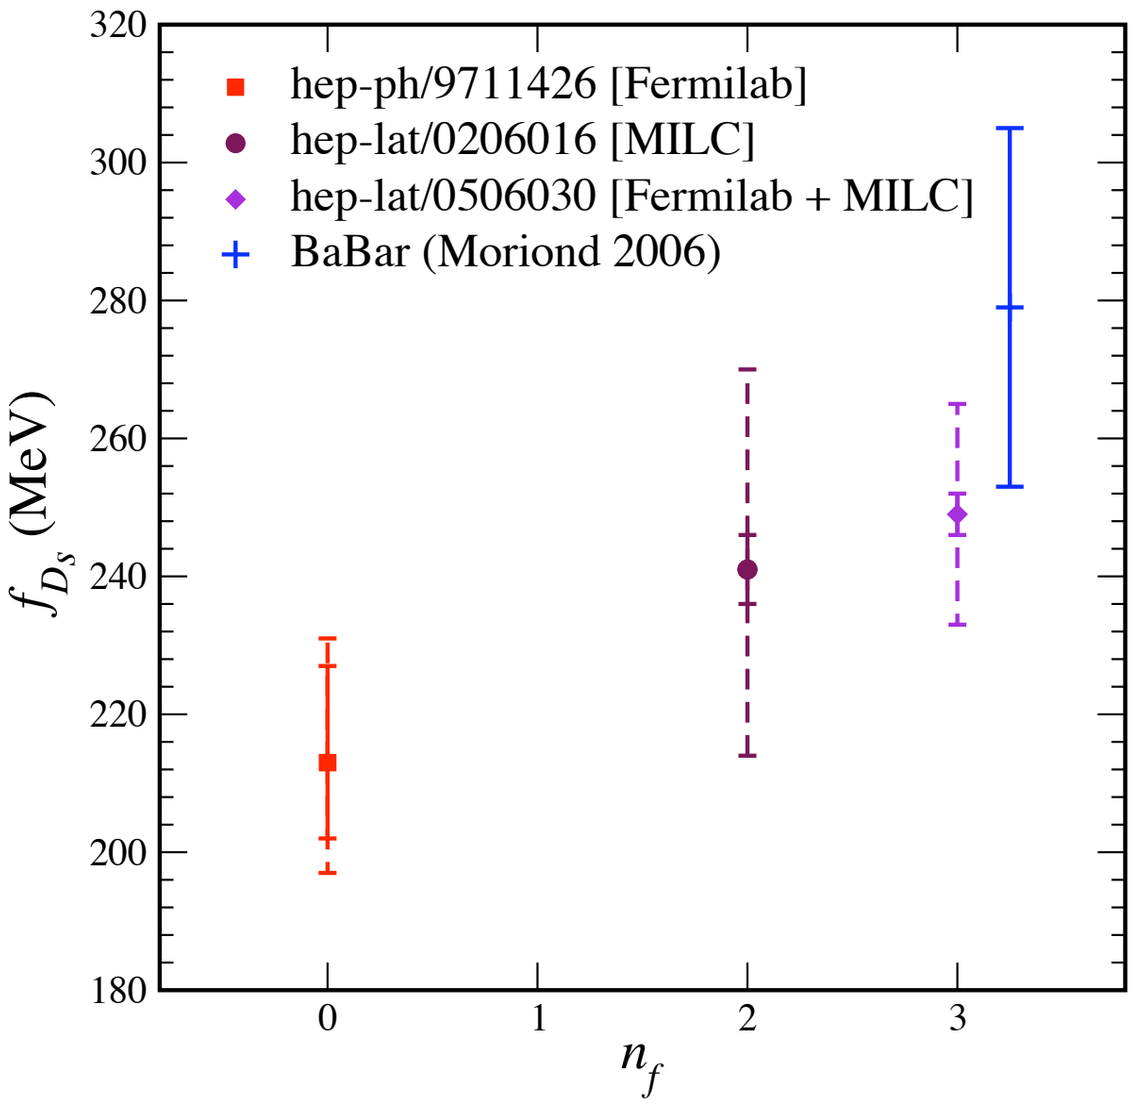


Phys. Rev. Lett. **95**: 122002, 2005.

Accuracies of D and Ds decay constants are being greatly improved by CLEO-c and the B factories, giving the lattice opportunities for predictions.

$$f_{D^+} = \begin{cases} 201 \pm 03 \pm 17 \text{ MeV [lattice]} \\ 223 \pm 17 \pm 03 \text{ MeV [CLEO]} \end{cases}$$

$$f_{D_s} = \begin{cases} 249 \pm 03 \pm 16 \text{ MeV [lattice]} \\ 279 \pm 17 \pm 20 \text{ MeV [BaBar]} \end{cases}$$

$$\frac{\sqrt{m_{D^+}} f_{D^+}}{\sqrt{m_{D_s}} f_{D_s}} = \begin{cases} 0.786 \pm 0.042 \text{ MeV [lattice]} \\ 0.779 \pm 0.093 \text{ MeV [expt]} \end{cases}$$

# Predictions: shape of D → Klν vs. $q^2$

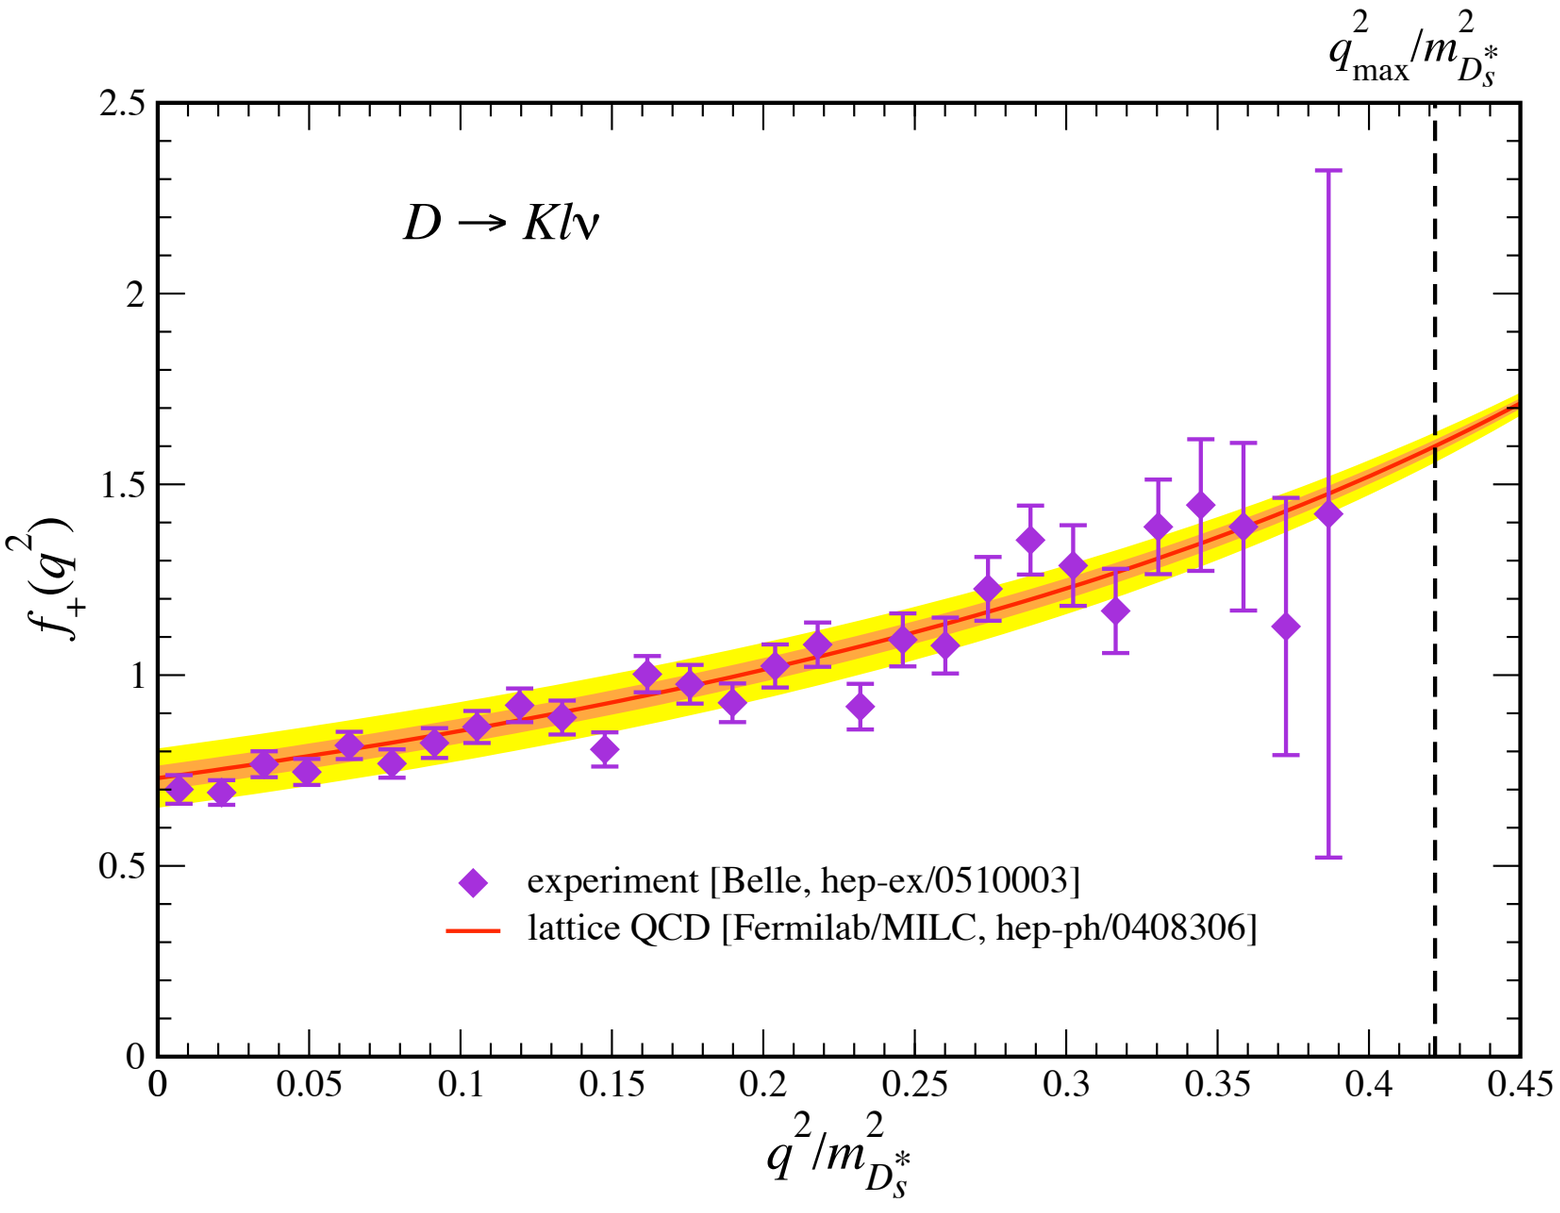


Shape of the semileptonic form factor predicted with lattice QCD (orange band) agrees with the shape subsequently measured by FOCUS, BaBar, and Belle (purple diamonds).

Phys. Rev. Lett. **94**: 011601, 2005.

# Predictions: $B_c$ Meson Mass

A rare hard-to-produce meson whose mass was unknown until recently.

$m_{B_c}^{n_f=0} = 6386 \pm 9 \pm 15 \pm 98$ MeV
[Phys. Lett. B 453, 289 (1999)]

$m_{B_c}^{2+1} = 6304 \pm 4 \pm 11^{+18}_{-0}$ MeV
[hep-lat/0411027 →
PRL **94** 172001 (2005)]

$m_{B_c}^{\text{expt}} = 6287 \pm 5$ MeV
[hep-ex/0505076 →
PRL **96** 082002 (2006)]

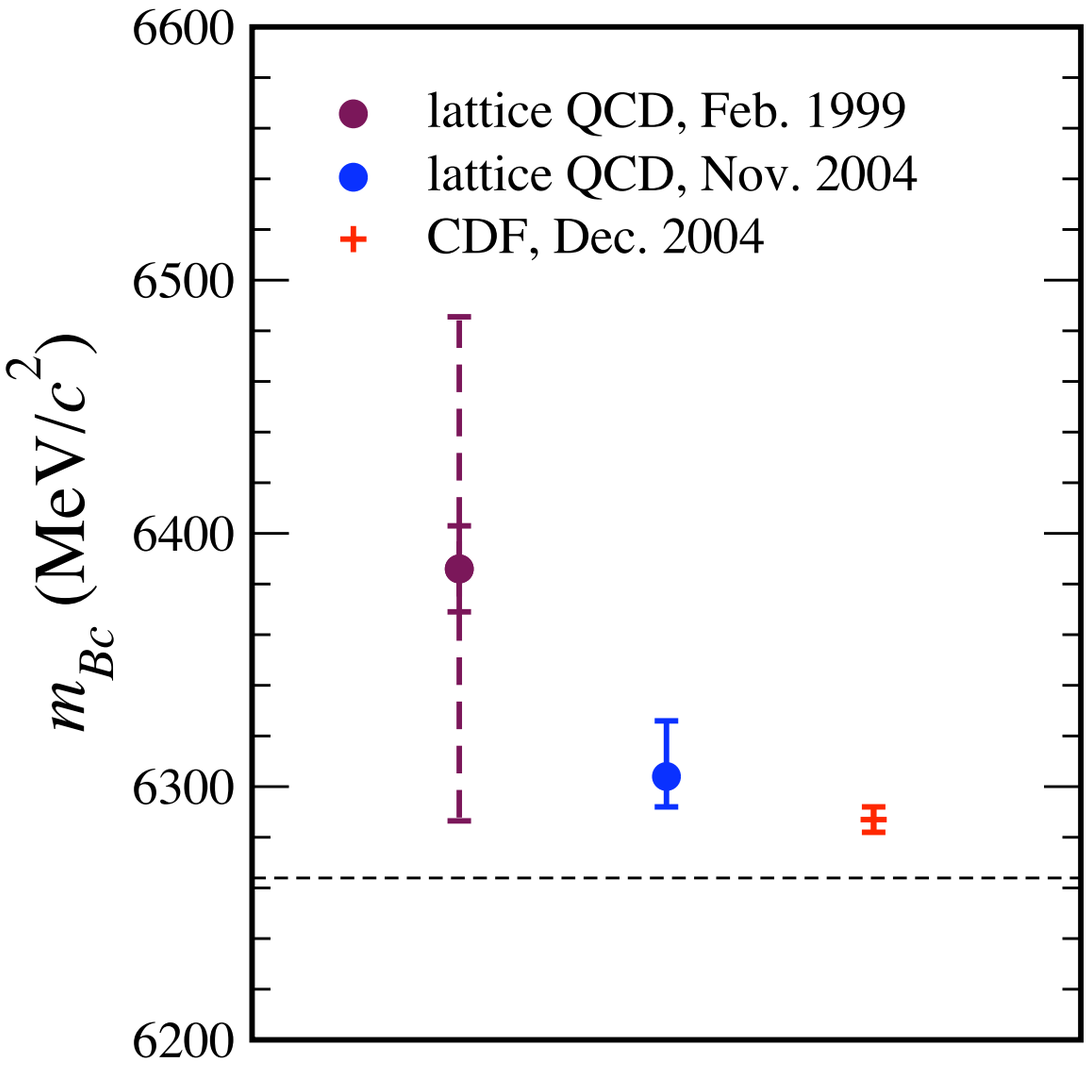


Phys. Rev. Lett. **94**: 172001, 2005.

# Extracting Standard Model parameters: $B_s\bar{B}_s$ mixing

Before Tevatron results:

***Impact on the Unitarity Triangle***

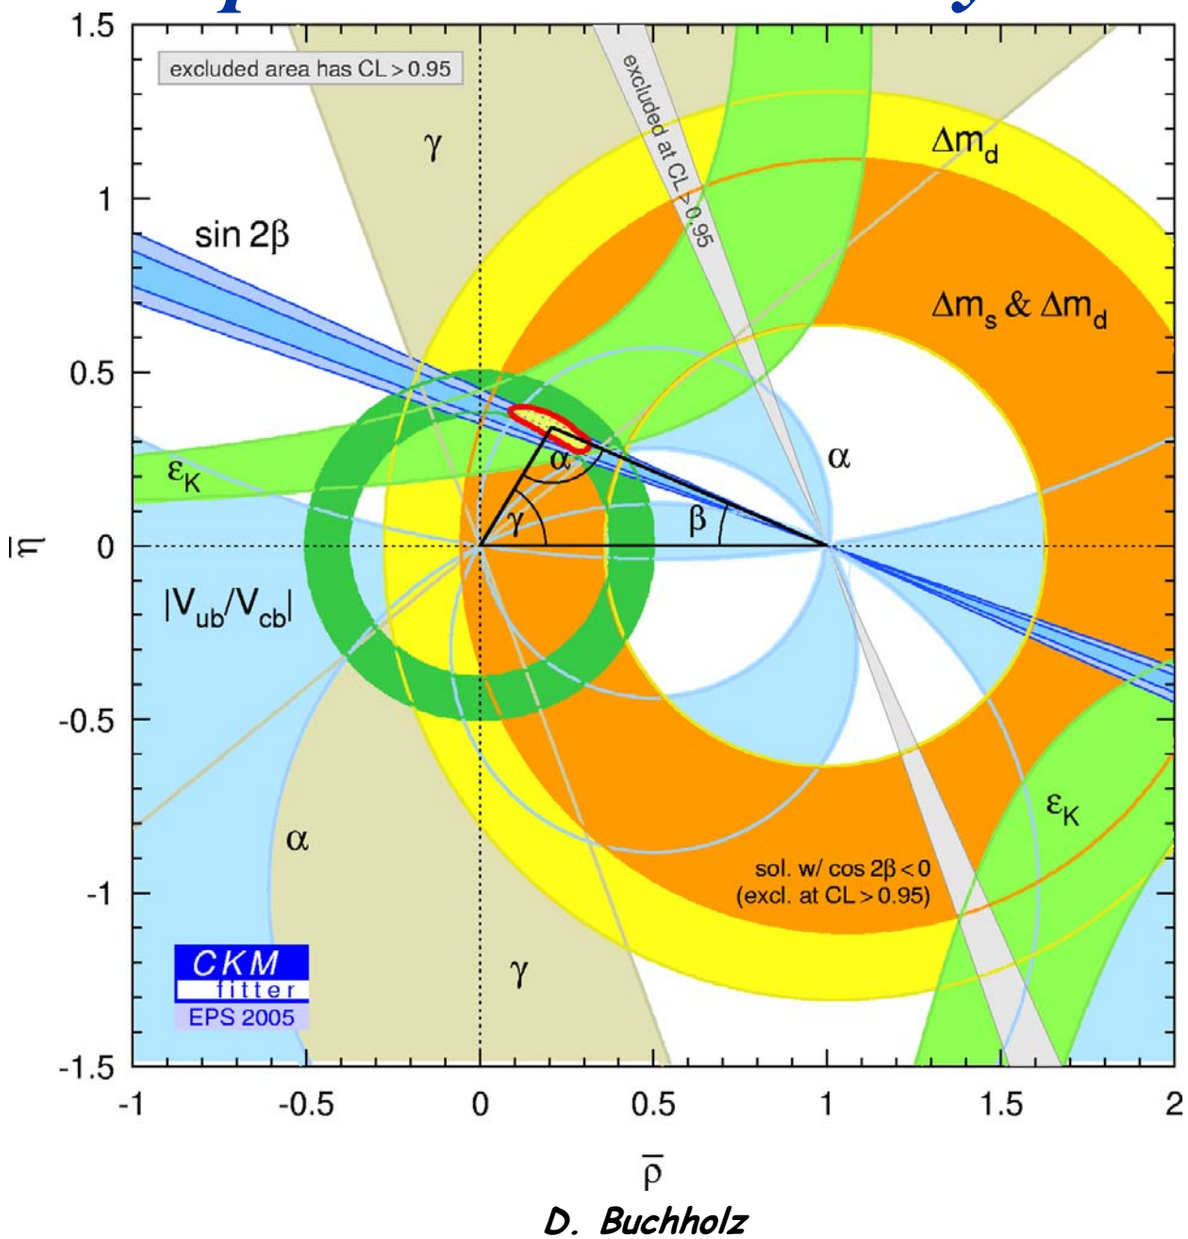


*D. Buchholz*

Bounds on the $\rho\eta$ plane from $B_s\bar{B}_s$ mixing can only be obtained with the aid of lattice QCD.

From Okamoto's Lattice 2005 review:

$$f_{B_s}/f_B\sqrt{\hat{B}_{B_s}/\hat{B}_B} = 1.210\left(^{+47}_{-35}\right)$$

$$\delta(|V_{td}|/|V_{ts}|) = 3-4\%$$

with forthcoming $\Delta M_{B_s}$

# $B_s\bar{B}_s$ Mixing

After Tevatron results:

## *Impact on the Unitarity Triangle*

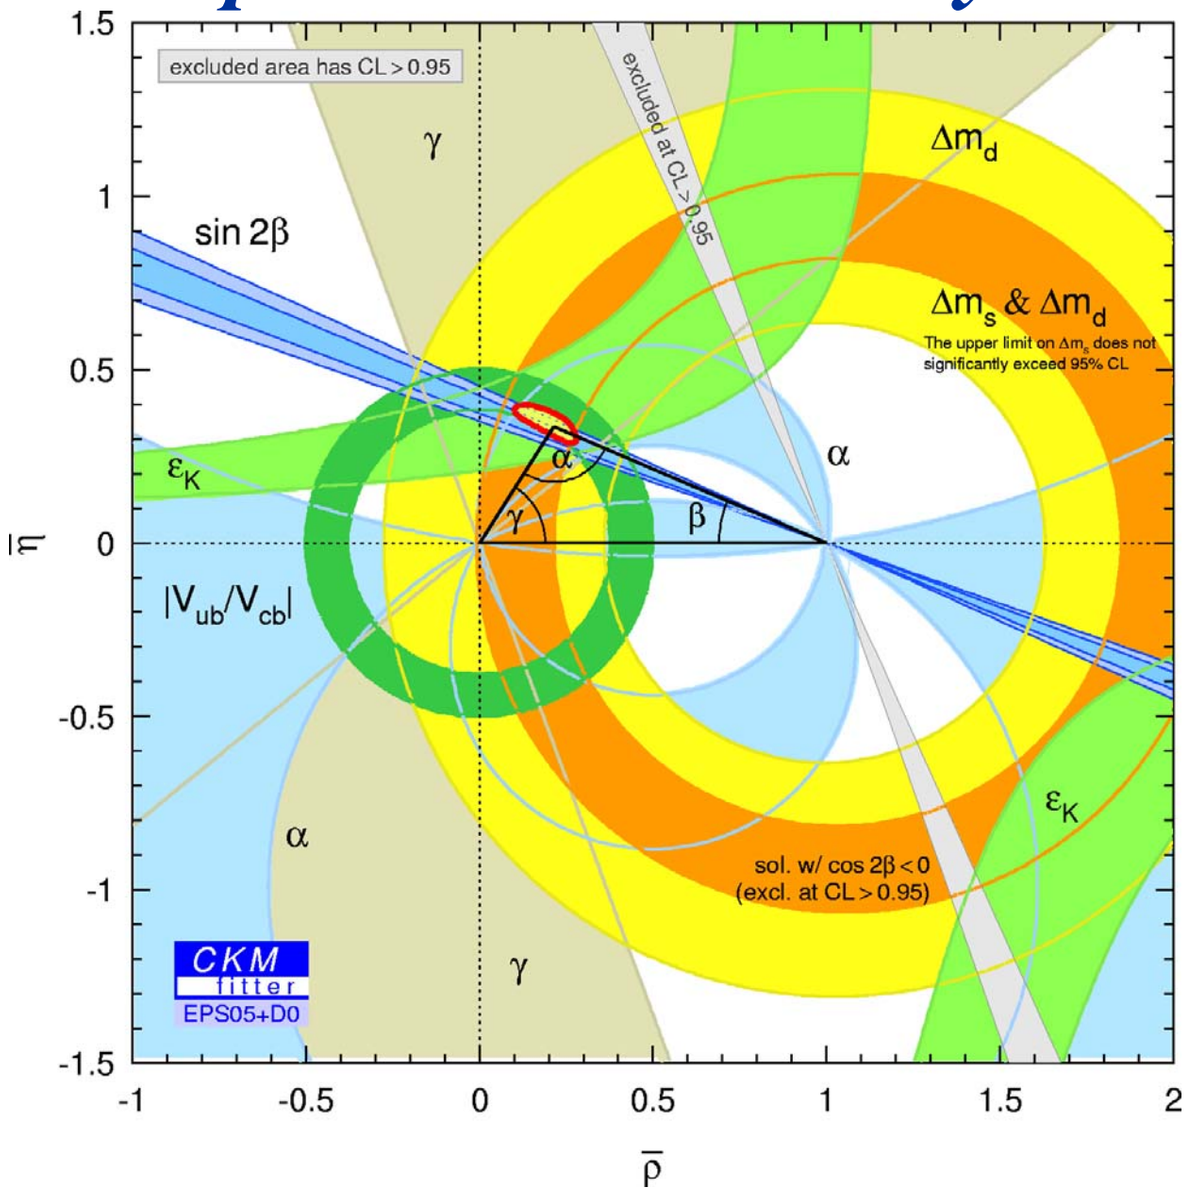


*D. Buchholz*

From Okamoto's Lattice 2005 review:

$$f_{B_s}/f_B\sqrt{\hat{B}_{B_s}/\hat{B}_B} = 1.210(^{+47}_{-35})$$

$$\delta(|V_{td}|/|V_{ts}|) = 3-4\%$$

Experimental uncertainties are 1%.
Only significant errors are now from lattice QCD.

# The job ahead

Lattice calculations are essential for extracting most quark masses and CKM matrix elements from experiment.
1% accuracy is needed for many quantities to complete the analysis of experiment. (4-5% is the standard of the present round of calculations of QCD amplitudes.) High precision QCD is needed.

$$\text{uncertainty} \propto \frac{1}{(\text{CPU time})^{\frac{1}{2+\ldots}}}$$

Finite lattice spacing, finite volume, chiral extrapolation, ...

$$\sim \frac{1}{(\text{CPU time})^{\frac{1}{7}}}$$

Statistics

$$\text{required CPU time} \propto \left(\frac{1}{\text{uncertainty}}\right)^{7}$$

Reducing the uncertainty by a factor of two by brute force requires a factor of 128 more computer time.

Orders of magnitude more computing power are coming.
This plus improved methods will be essential for delivering the necessary accuracy in lattice calculations.

# Beyond QCD:
## lattice gauge theory in the LHC/ILC era

- High precision lattice QCD is essential for completing the current round of HEP experiments, but it is also essential for developing well-established methods for nonperturbative field theory in preparation for potential new physics at the LHC and ILC.

- Of the four known interactions, only one (QED) is known to be completely perturbative. Any new physics that is discovered is likely to include nonperturbative sectors.

  - In SUSY theories, dynamical SUSY breaking provides a good explanation of the hierarchy.

  - In theories such as technicolor and topcolor, the role of the Higgs is played by pion-like bound states of new strong interactions.

- USQCD is encouraging pilot projects in strongly coupled BSM theories to prepare for nonperturbative phenomena found at the LHC and the ILC.

# The LQCD Project: 2010 and beyond

John Kogut has asked the USQCD Executive Committee for the case for continuation of the LQCD project in the years 2010 and beyond. We have prepared five whitepapers.

- Computational resources for lattice QCD, 2010-2014
- Fundamental parameters from future lattice calculations
- Nuclear Physics from Lattice QCD
- QCD Thermodynamics
- Lattice Field Theory in the LHC Era

For now, available at http://theory.fnal.gov/theorybreakout2007/.

# Current world lattice hardware

| Country | Sustained Teraflop/s |
|---|---|
| Germany | 10–15 |
| Italy | 5 |
| Japan | 14–18 |
| United Kingdom | 4–5 |
| Unites States | |
| LQCD Project | 9 |
| National Centers | 2 |
| US Total | 11 |

Table 3: Computing resources in sustained teraflop/s estimated to be available for the study of Lattice QCD in various countries, as of February, 2007.

(Tables in this section are from the white paper
http://theory.fnal.gov/theorybreakout2007/LatticeQCD2010-2014.pdf .

Available lattice computing power around the world will continue to expand exponentially in 2010 and beyond.

World lattice funding is organized and allocated on a national basis.

# What do we need to succeed?

- Orders of magnitude more computing power are required for 1% accuracy.
- We have two very large computing tasks with different requirements:
  - Gauge configuration generation
    - Single large, highly optimized program
    - Very long, single tasks. Low I/O
    - Suitable for inflexible, high-capability computers such as the QCDOC or the Blue Gene/P at Argonne.
  - Quark and hadron analysis.
    - Large, heterogeneous code base.
    - Many small individual tasks. Heavy I/O.
    - Dozens of large files, hundreds of small files per gauge configuration.
    - Ideal on flexible, high-capacity computers (such as Fermilab's clusters), at facilities with large mass storage and wide area networking. Significant clusters are also needed for development projects such as new physical theories, improved actions or methods.

# 1) DoE "Leadership Class" machines

will be suitable for the gauge configuration generation part of our computing requirements, and we will ask for as large a share of them as we can.

| | Peak teraflops | |
|---|---|---|
| | Oak Ridge XT4 | Argonne Blue Gene/P |
| 2007 | 119 | 114 |
| 2008 | 250 | 250-500 |
| 2009 | 1000 | 250-500 |

These high-capability machines will be difficult to use efficiently, and suitable only for the largest "capability" projects, such as configuration generation for lattice QCD. USQCD therefore hopes to get substantial allocations, perhaps 10% of the total allocations, based on our experience at supercomputer centers and our analysis of the competition. Benchmarks indicate that optimized lattice codes should run at 15% of peak speed on these types of machines. We therefore may get from these sources:

Comparable to the 18 teraflops expected from LQCD project resources in 2009; a good match.

| | Delivered teraflops for lattice QCD | |
|---|---|---|
| | Oak Ridge XT4 | Argonne Blue Gene/P |
| 2007 | 1.8 | 1.7 |
| 2008 | 3.75 | 3.75-7.5 |
| 2009 | 15 | 3.75-7.5 |

# 2) Flexible, high capacity computing clusters

- We will obtain as much as possible of our computing requirements from the "Leadership class" machines (right hand column). These are free to HEP (though not to the taxpayer!).

- A successful program will also require a continuation of the LQCD project with high-capacity machines (middle column) at a size comparable to the expected "leadership class" power, to analyze gauge configurations and for investigations of new physical theories, new actions, and new methods. These must be at facilities with high networking bandwidth and mass storage.

Combined program:

| Fiscal Year | Dedicated Hardware (Teraflop/s) | Leadership Class Machines (Teraflop/s) |
|---|---|---|
| 2010 | 34 | 33 |
| 2011 | 61 | 52 |
| 2012 | 100 | 82 |
| 2013 | 161 | 131 |
| 2014 | 256 | 208 |

Table 1: Expected throughput for dedicated hardware (column 2) and leadership class computers by fiscal year with a hardware budget of $3.0 million per year.

# Envisioned project

The cost for the Office of Science part of this project is

| Fiscal Year | Hardware Budget | Operations Budget | Total Budget |
|---|---|---|---|
| 2010 | 3.00 | 1.46 | 4.46 |
| 2011 | 3.00 | 1.52 | 4.52 |
| 2012 | 3.00 | 1.58 | 4.58 |
| 2013 | 3.00 | 1.64 | 4.64 |
| 2014 | 3.00 | 1.71 | 4.71 |
| Totals | 15.00 | 7.91 | 22.91 |

Table 2: Budgets for the proposed Project in millions of dollars.

This is an increase over the current project funding of $2.5 M/year, but one that we think is feasible.

Current funding for the LQCD is $2.0M/year from HEP, $0.5M/year from NP. We anticipate the bulk of the increase in funding needed coming from NP.

# Status of USQCD/DoE discussions

- We are preparing a proposal for configuration generation on the “leadership class” machines at Argonne and Oak Ridge.
- Bob Sugar has had a phone discussion with Robin Staffin on the continuation of the LQCD project in 2010 and beyond that was very friendly (though noncommittal).
- We are making plans through John Kogut to visit Germantown in the fall to brief HEP leadership on our plans for the future of the LQCD project.
- We are getting advice about whether to present our plans in other forums, such as HEPAP.

# Summary

- US lattice gauge theory
  - has accomplished great things in the current program,
  - has greater challenges ahead in the LHC era:
    - high precision lattice QCD for high precision *K, D, and B*, physics,
    - new strongly coupled physics at the LHC and ILC.
- Fermilab
  - has supplied the technical, management, and scientific leadership to bring about the current success.
  - The US lattice and HEP communities need Fermilab to continue in this role.
- We need Fermilab to
  - help us make sure that the 08/09 installation is a big success,
  - help us push for Fermilab's continuing leadership in lattice gauge theory with HEPAP and DoE/HEP leadership.