Computational Infrastructure for Lattice Gauge Theory

John W. Negele

NSAC November 1, 2002

Scientific Objectives

Hardware Plans

Software Infrastructure

Deployment Plan

Management

Scientific Objectives

- Understand the physical phenomena arising from QCD

- Make precision calculations of the predictions of QCD

The next generation of Lattice QCD calculations will require computing facilities capable of sustaining tens of teraflops

Relevance to the DOE Mission

- Major goals of the DOE's experimental programs in high energy and nuclear physics:
  - Verify the Standard Model or discover its generalization
  - Determine the properties of hadronic matter under extreme conditions
  - Understand the structure and interactions of hadrons
- Lattice QCD calculations are essential to research in all these areas

Coupling with the Experimental Program

- Weak Decays of Strongly Interacting Particles
  - BaBar (SLAC)
  - Tevatron B-Meson Program (FNAL)
  - CLEO-c Program (Cornell)

- Quark–Gluon Plasma
  - RHIC (BNL)

- Structure and Interactions of Hadrons
  - Bates
  - BNL
  - FNAL
  - JLab
  - SLAC

Status of Lattice Calculations

- Some key quantities have been calculated to a few percent

- We know the resources required for accurate determination of a broad range of fundamental quantities

- Important progress is anticipated over the next five years
  - Major improvements in algorithms
  - Major increases in computing power

Some Lattice QCD Successes

- Qualitative features of QCD
  - Quark confinement
  - Chiral symmetry breaking
  - Quark–gluon plasma

- Examples of quantitative results
  - Temperature of the chiral symmetry restoration phase transition
  - Accurate determination of the hadron spectrum (quenched)
  - Determination of $\alpha_s(M_Z)$ and masses of the c and b quarks to a few percent.
  - Determination of B and D meson decay constants, and $\bar{B}B$ and $\bar{K}K$ mixing matrix elements to a precision of about 10–20%.

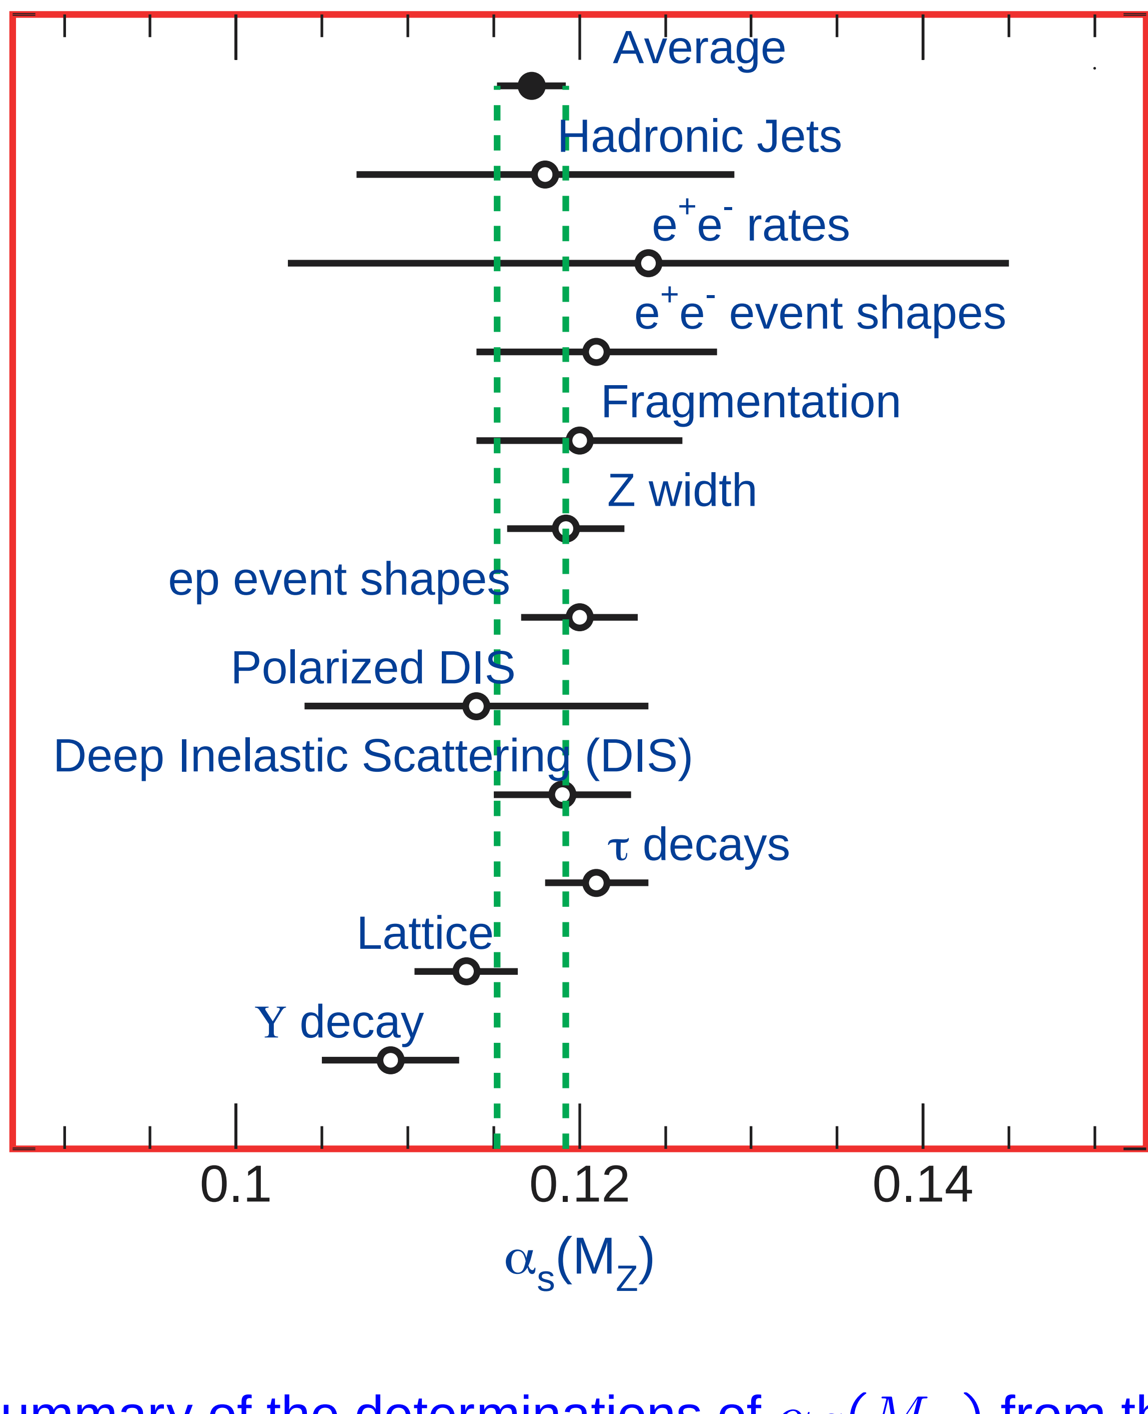

Summary of the determinations of $\alpha_S(M_Z)$ from the Particle Data Group

The Quark Gluon Plasma

- **At high temperatures and/or baryon densities one expects a phase transition or crossover from strongly interacting hadrons to a plasma of deconfined quarks and gluons.**

- **Lattice QCD calculations are required to calculate properties of strongly interacting matter in the vicinity of the transition.**
  - **Phase Diagram**
  - **Transition Temperature**
  - **Order of the Transition**
  - **Equation of State**
  - **Role of Instantons**

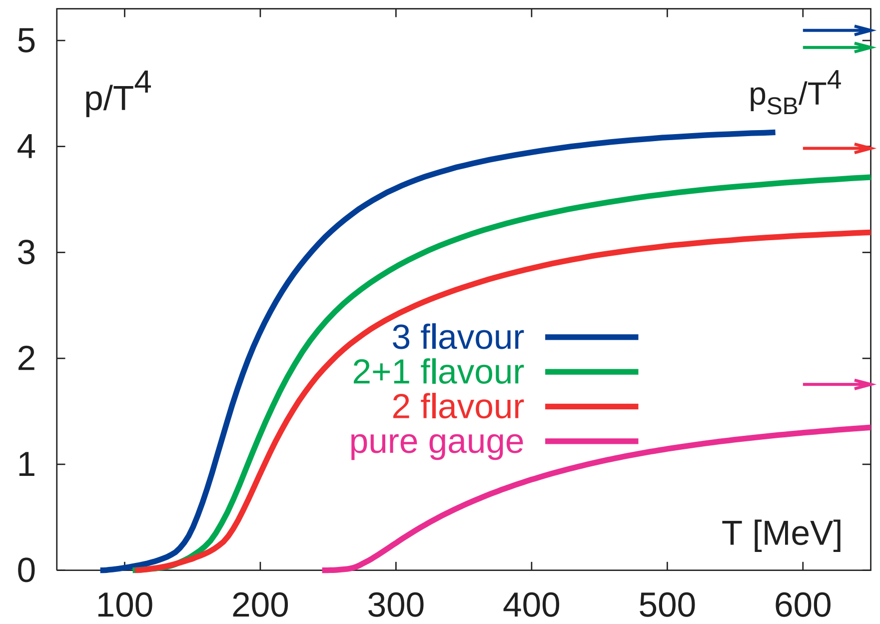

Equation of state for various flavors (F. Karsch hep-lat/0106019)

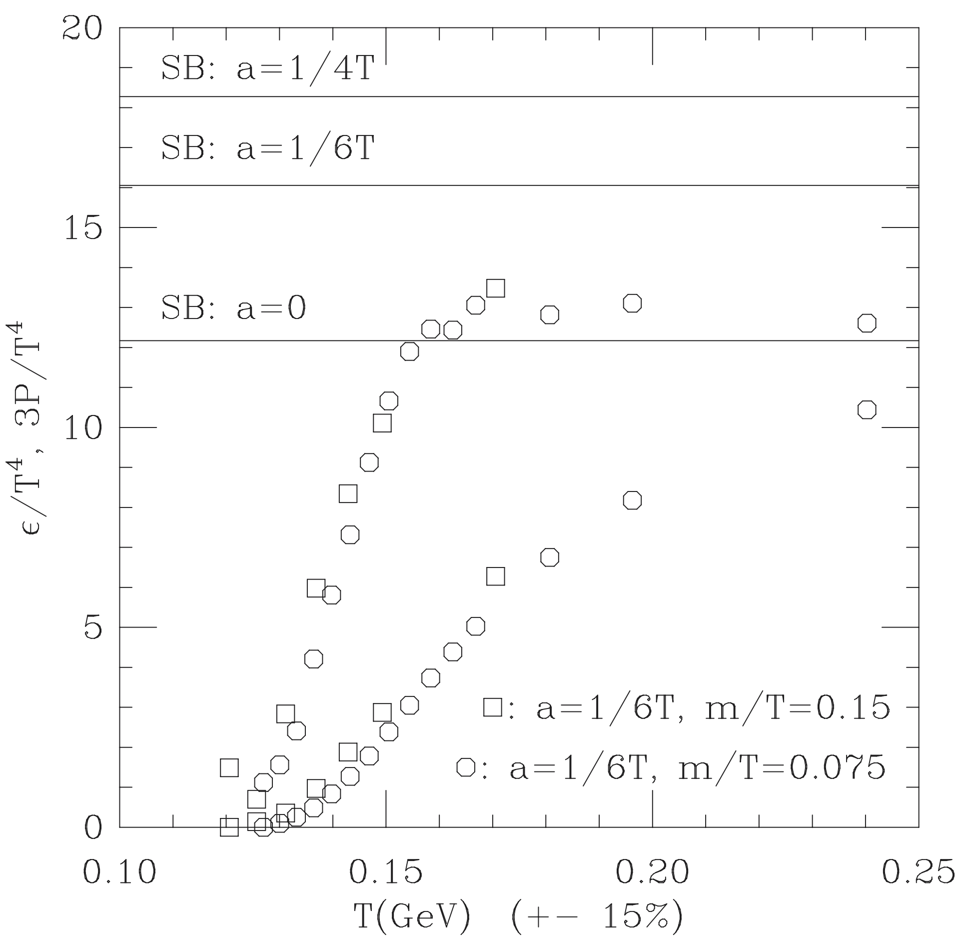

Pressure and energy density for 2 flavors (MILC)

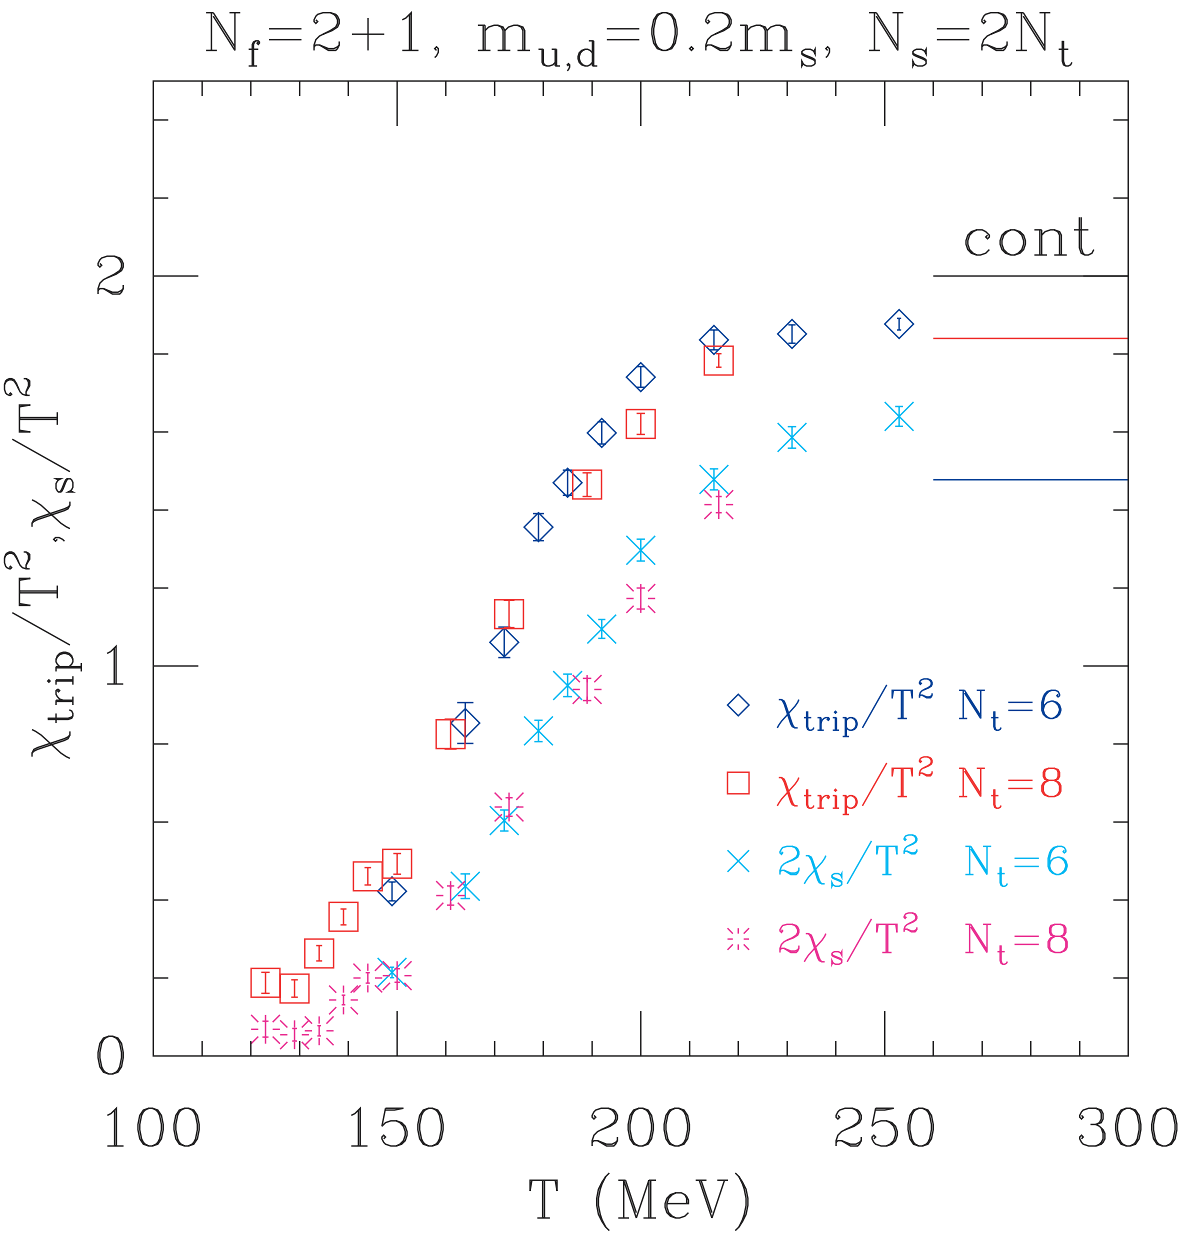

Triplet quark susceptibility, related to fluctuations in isospin density, and strange quark susceptibility, related to fluctuations in strangeness density (C. Benard *et al.* hep-lat/0209079)

Finite Baryon Density

Analytic continuation of the deconfinement line

de Forcrand, Philipsen

Critical line defined by $\left.\frac{\partial\chi}{\partial\beta}\right|_{\mu_c,\beta_c} = 0, \qquad \left.\frac{\partial^2\chi}{\partial\beta^2}\right|_{\mu_c,\beta_c} < 0.$

Implicit function theorem: $\chi(\beta,\mu)$ analytic $\Rightarrow \beta_c(\mu)$ analytic!

Symmetry: $\beta_c(\mu) = \beta_c(-\mu)$

$$\Rightarrow \boxed{\beta_c(\mu) = \sum_n c_n (a\mu)^{2n}}$$

Results for imaginary $\mu = i\mu_I$

$8^3 \times 4$, $N_f = 2$ staggered fermions, $am = 0.025$
($T_c(\mu = 0) \sim 170$ MeV, $a \sim 0.3$ fm)

imag. μ

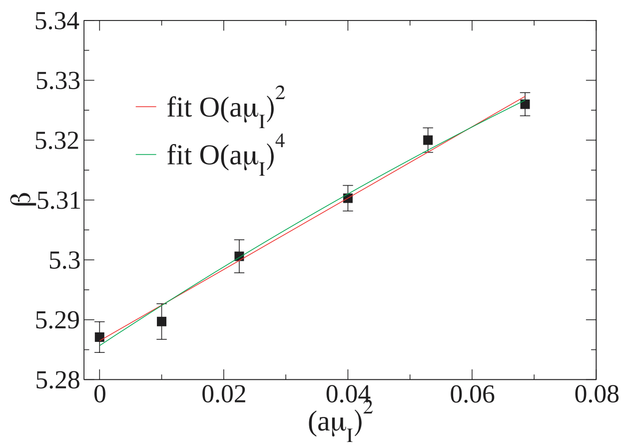

analytic cont.

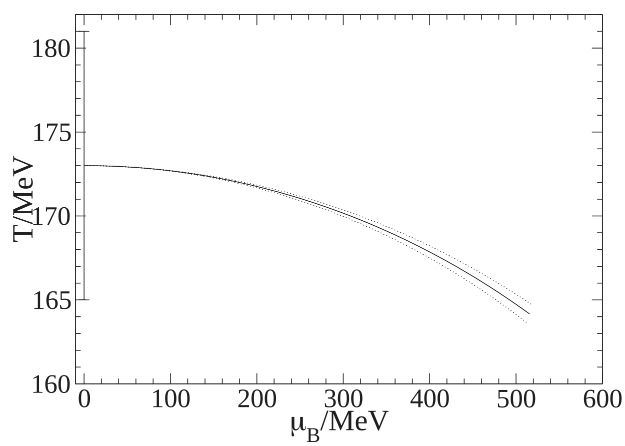

N.B.: systematics under control

order of p.t. from finite V study

Structure and Interactions of Hadrons

- Quark and gluon structure of the nucleon
  - Electromagnetic and strangeness form factors
  - Moments of light cone quark and gluon distributions
  - Moments of generalized parton distributions
- Spectroscopy
  - N^* spectra and transition form factors
  - Exotics
  - Glueballs
- Hadron-hadron interactions
  - Heavy-light meson and baryon interactions

Strangeness Form Factors from Parity-Violating Electron Scattering

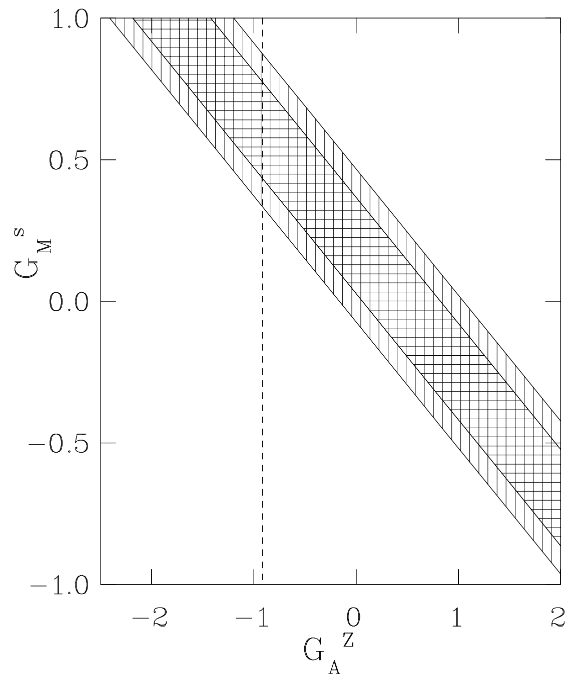

SAMPLE (Bates)

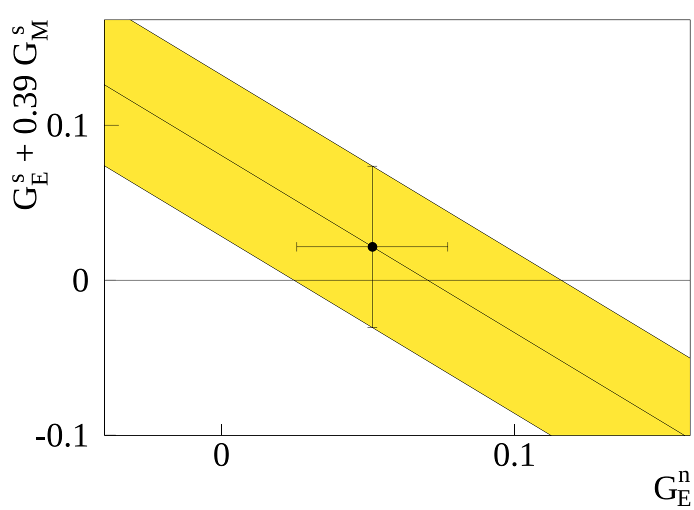

HAPPEX (Jlab)

Moments of Parton Distributions

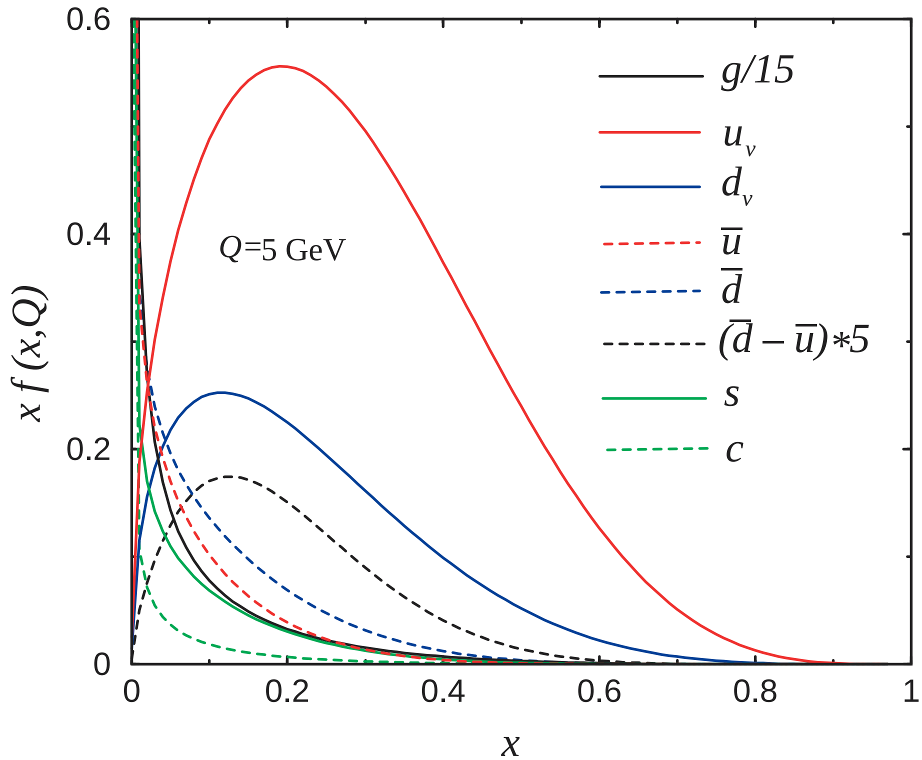

Parton distributions at 5 GeV

Leading Twist

$$\langle p|\bar{\psi}\gamma_\mu D_{\mu_1}\cdots D_{\mu_n}\psi|p\rangle \longrightarrow \langle x^n\rangle_q$$

$$\langle p|\bar{\psi}\gamma_5\gamma_\mu D_{\mu_1}\cdots D_{\mu_n}\psi|p\rangle \longrightarrow \langle x^n\rangle_{\Delta q}$$

$$\langle p|\bar{\psi}\gamma_5\sigma_{\mu\nu} D_{\mu_1}\cdots D_{\mu_n}\psi|p\rangle \longrightarrow \langle x^n\rangle_{\delta q}$$

Higher Twist

$$\langle p|\bar{\psi}\tilde{F}^{\mu\nu}\gamma_5\gamma_\mu\psi|p\rangle, \ldots$$

Higher Twist

$$\langle p'|\bar{\psi}\mathcal{O}D\cdots D|p\rangle$$

Chiral Extrapolation Requires Terascale Calculation

5% measurement at $m_\pi^2 = 0.05$ GeV^2

$m_\pi \sim 230 MeV$

$L \sim 4.3 fm.$

SESAM cost function

$$N_{\mathrm{OPS}} \sim 0.38[\frac{L}{4}]^{4.55}[\frac{0.8}{a}]^{7.25}[\frac{0.3}{m_\pi/m_\rho}]^{2.7}$$
$$\sim 8 \text{ Tflops-years}$$

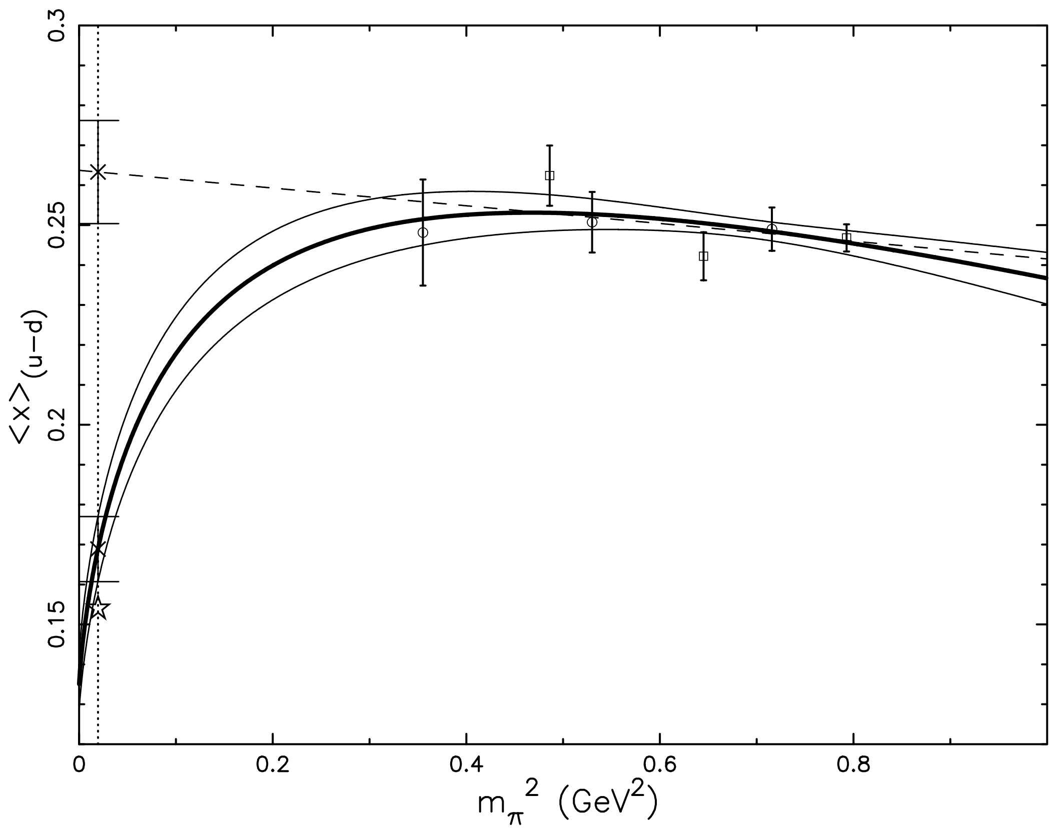

Quark momentum fraction in the nucleon

D. Dolgov *et al.* hep-lat/0201021

Analogous Extrapolation for Magnetic Moment

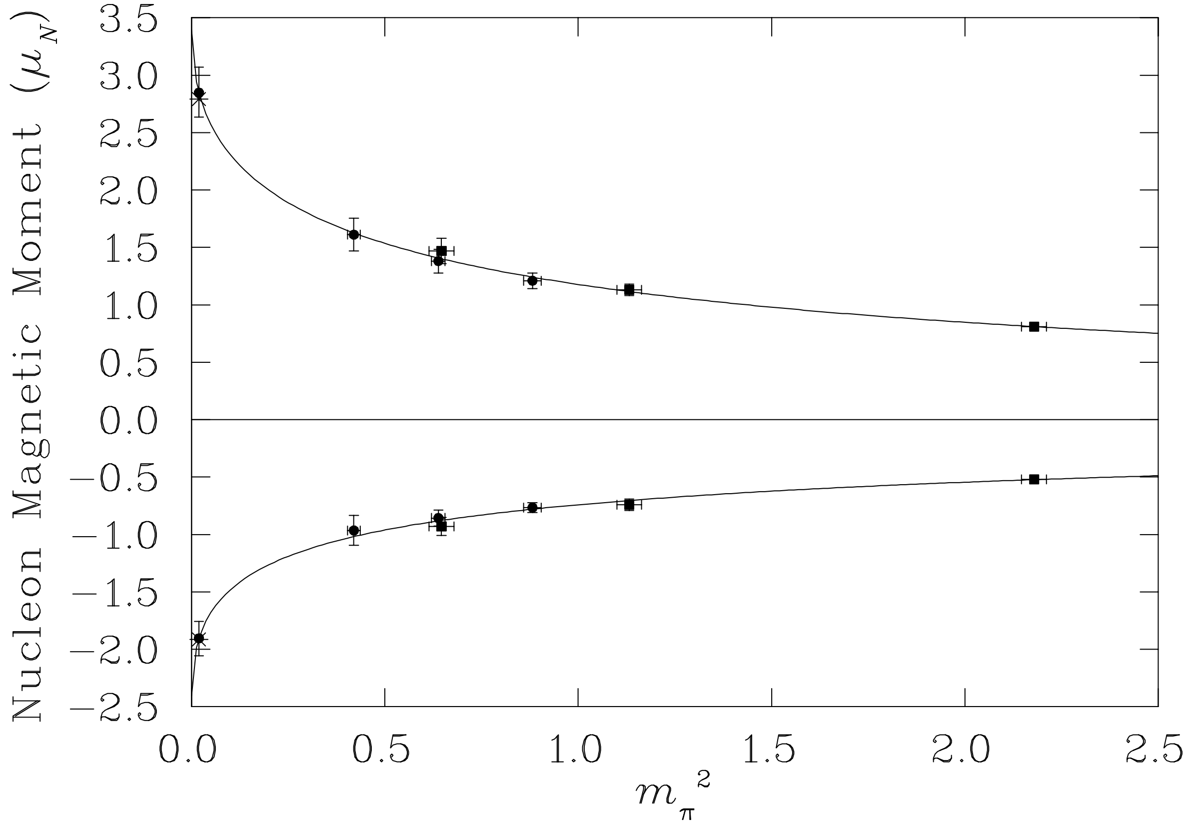

Nucleon magnetic moment

D. Leinweber *et al.* hep-lat/0103006

Nucleon Parity Partner

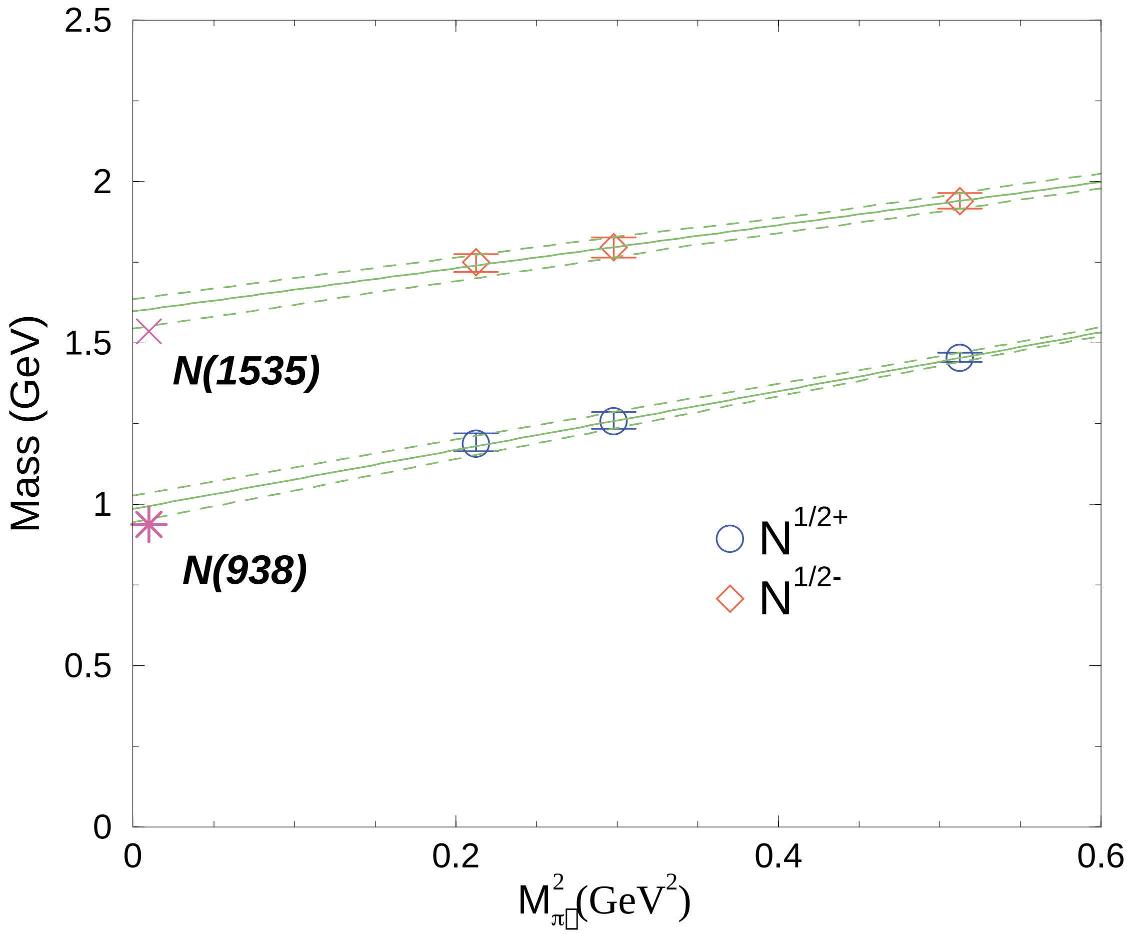

Quark mass extrapolation of nucleon and its parity partner

D. Richards *et al.* hep-lat/0011025

Hardware Plans

- Simplifying features of lattice QCD calculations make building specially designed computers far more cost-effective than buying commercial ones
  - Uniform grids
  - Regular, predictable communications
- Two hardware tracks:
  - QCD On a Chip (QCDOC)
  - Commodity Clusters
- Each track has its own strengths
- Each track may prove optimal for different aspects of our work
- The two-track approach positions us to exploit future technological advances, enables us to retain flexibility, and ensures robust national program

QCDOC

- The latest of the highly successful Columbia/Riken/BNL special purpose computers
    - The Columbia group has pioneered the design and construction of special purpose computers for QCD
    - The QCDSP won the Gordon Bell Prize in 1998 for price performance
- The QCDOC combines processor, networking, and memory on a single chip
- Partnership with IBM provides access to its technology for chip design and fabrication
- Targets price–performance of $1/Mflops for multi-teraflops machines in 2003

Proposed US QCDOC's

- 1.5 Tflops development machine at Columbia in 2002
  - Proposal now to HENP and MICS
  - $1.5M construction + $ 0.6M staff
  - National user facility
- 10^+ Tflops machine at BNL in 2004
  - Proposal to SciDAC in FY04

Commodity Clusters

- Market forces are producing rapid gains in processor and memory performance
  - Moore's Law $\Rightarrow$ 60% growth in performance per year
  - Pentium 4 currently provides exceptional performance for QCD
- Market for interconnects is growing
- Open Source System Software
  - Flexible programming environment
  - SciDAC Scalable Systems Software
- Targeted price–performance

	FY 2003	FY 2004	FY 2005	FY 2006
$/Mflops	2.0	1.2	0.9	0.7

Current and Planned Clusters

- **Prototype Clusters**

  99 64 Gflops Alpha cluster @ MIT

  99 48 Gflops Alpha cluster @ JLab

  01 80 node P3 cluster @ FNAL

- **SciDAC Clusters ($\sim$ \$2M $+$ Lab matching)**

  Myrinet

  8/02 48 node dual P4 cluster @ FNAL

  9/02 128 node single P4 cluster @ JLab

  2/03 128 node dual P4 cluster @ FNAL

  Gigabit Ethernet

  5/03 256 node dual P4 cluster @ JLab

  Next Generation Technology

  9-12/03 Clusters @ FNAL and JLab

- **Propose 10^{+} Tflops clusters at FNAL and JLab in FY05 and FY06**

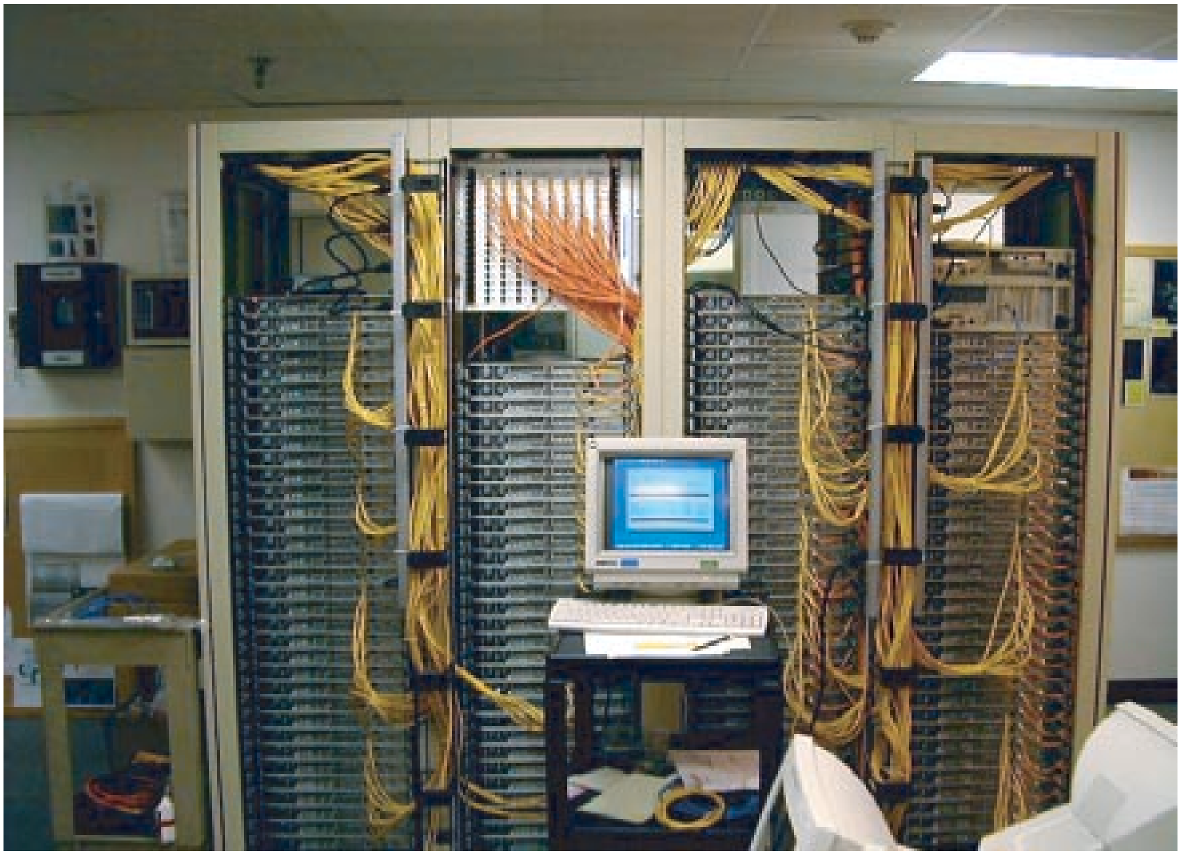

Jlab 128 node SciDAC Cluster

Projected performance and costs

	FY02	FY04	FY05	FY06
CPUs/Node	2	2	4	4
Gflops/Node	2.0	3.2	8.0	10.0
Nodes		192	384	512
Link Bandwidth (MB/s)		300 + 300	2× (400 + 400)	2 × (500 + 500)
Link Latency μsec	10	6	5	4
Performance (TFlops)		0.6	2.5	4.5
Hardware Cost ($M)		0.7	2.5	3.2
$/MFlops	2.0	1.2	0.9	0.7

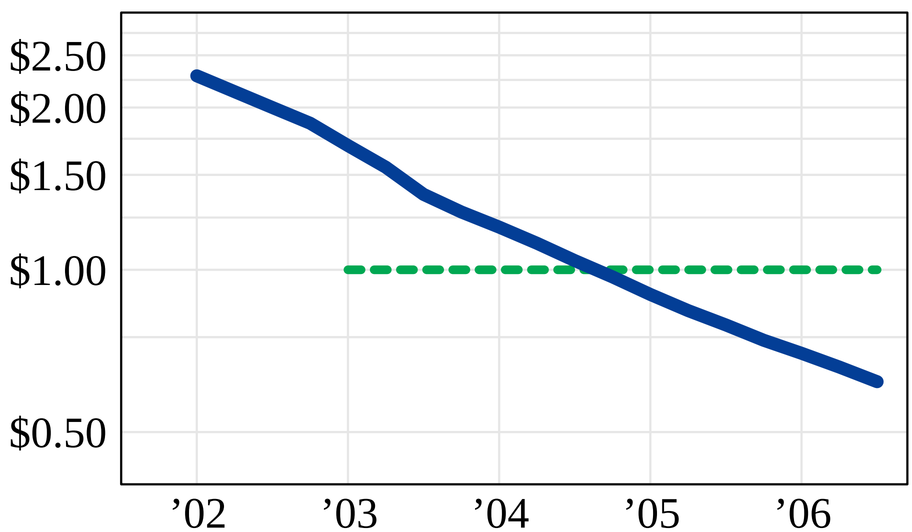

Price/performance estimates for clusters (blue line) and QCDOC (green line).

Hardware Deployment Plan

- Clusters
  - FY02–03: 2×0.5 Tflops (FNAL, JLab)
  - FY04: 2×0.5 Tflops (FNAL, JLab)
  - FY05–06: $2\times10^{+}$ Tflops (FNAL, JLab)
- QCDOC
  - FY03: 1.5 Tflops (Columbia)
  - FY04: 10^{+} Tflops (BNL)

SciDAC Software Infrastructure

- Create a programming environment that will enable very high performance on the QCDOC and clusters
- Physics code intended to run on all generations of all SciDAC machines
- Principal Components
  - QCD Application Programming Interface
  - Highly optimized linear algebra routines
  - Highly optimized communications
  - Optimization of computationally intensive subroutines
  - Porting and optimization of major community codes

Collaboration and Management Structure

- Our collaboration includes nearly all senior lattice gauge theorists in the US
  - Sixty-four senior scientists
  - Lattice gauge theorists, computer scientists, computer engineers

- Management Structure

Executive Committee

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

Scientific Program Committee

Peter Lepage	Cornell Univ
Robert Mahwinney	Columbia Univ
Colin Morningstar	Carnegie Mellon Univ
John Negele	MIT
Claudio Rebbi (Chair)	Boston Univ
Stephen Sharpe	Univ of Washington
Doug Toussaint	Univ of Arizona
Frank Wilczek	MIT

Oversight Committee

Steven Gottlieb (Chair)	Indiana Univ
Anna Hasenfratz	Univ of Colorado
Gregory Kilcup	Ohio State Univ
Julius Kuti	UC San Diego
Robert Pennington	NCSA
Ralph Roskies	Director, PSC
Terry Schalk	UC Santa Cruz

Software Coordinating Committee

Richard Brower (Chair)	Boston Univ
Carleton DeTar	Univ of Utah
Robert Edwards	JLab
Donald Holmgren	FNAL
Robert Mawhinney	Columbia Univ
Celso Mendes	Univ of Illinois
Chip Watson	JLab

Need to Move Quickly

- Major experiments that require terascale calculations for their interpretation have recently been completed, or will be completed within the next several years

- Theorists in Europe and Japan are moving rapidly to secure resources comparable to those we propose
  - UKQCD plans 5 Tflops in 2003
  - APE Collaboration will begin deploying multi-Tflops machines in 2003
  - DESY plans 20 Tflops (peak) APE NEXT in 2004
  - Japan ?

- We will propose three 10^{+} Tflops machines in 2004, 2005, and 2006

- We will propose steady level of funding to keep up with Moore’s Law

Conclusion

Lattice gauge theory was invented in the U.S., and U.S. theorists have traditionally been leaders of this field. If we are to play a significant role in the major advances expected in this area over the next five years, we must act now.

How Can NSAC Help?

We hope that NSAC will help our community by confirming that:

- Lattice QCD is an integral and essential part of contemporary nuclear physics research

- It is vital to act now to build the infrastructure needed to maintain a strong lattice QCD research program in the United States.