Petascale Challenges in Lattice QCD

Steven Gottlieb, Indiana University

sg at indiana.edu

Outline

- Physics
- Plans
- Performance
- Petascale Challenges
- `http://physics.indiana.edu/~sg/milc.html`

Physics

- Quantum Chromodynamics (QCD): theory of the strong interaction
- One of four fundamental forces of Nature
- Perturbative Quantum Field Theory: works when coupling is weak
- Asymptotic Freedom implies QCD has a weak coupling at very high energy
- For QCD, need a nonperturbative method to deal with bound states
- Lattice QCD developed by K. Wilson in 1974, preserves essential symmetry called gauge invariance

Beyond Perturbation Theory

- Many phenomena of QCD require nonperturbative prowess
  - Confinement
  - Meson and Baryon Masses
  - Decay constants: f_π, f_K, f_D, etc.
  - Semileptonic form factors, *e.g.*, $D \to \pi l \nu$
  - Extraction of CKM matrix elements
  - Nucleon structure functions
  - Quark-gluon plasma
- Distinguishing new physics from SM physics
- Theories such as technicolor and other approaches to dynamical symmetry breaking

Relevant Experiments

QCD is an important force in Nature, and understanding it is essential to current experimental programs supported by NSF and DOE:

- Weak Matrix Elements: Fermilab (D0, CDF), SLAC (BaBar), Cornell (CLEO-c)
  - Decays of strongly interacting particles containing heavy quarks
- High Temperature QCD: Brookhaven (RHIC)
  - properties of strongly interacting matter under extreme conditions, such as existed in early universe
- Hadron Structure: Jefferson Lab (CEBAF), BNL (RHIC)
  - masses and internal structure of strongly interacting particles

Major Scientific Objectives

- Calculate weak interaction matrix elements of strongly interacting particles to the accuracy needed to make precise tests of the Standard Model.
- Determine the properties of strongly interacting matter at high temperatures and densities, such as those that existed immediately after the big bang.
- Calculate the masses of strongly interacting particles, and obtain a quantitative understanding of their internal structure and their interactions.
- Develop the tools needed to perform quantitative studies of strongly coupled theories that may be necessary to describe physical phenomena at shorter distance scales than have been explored to date.

Postdictions: 2003

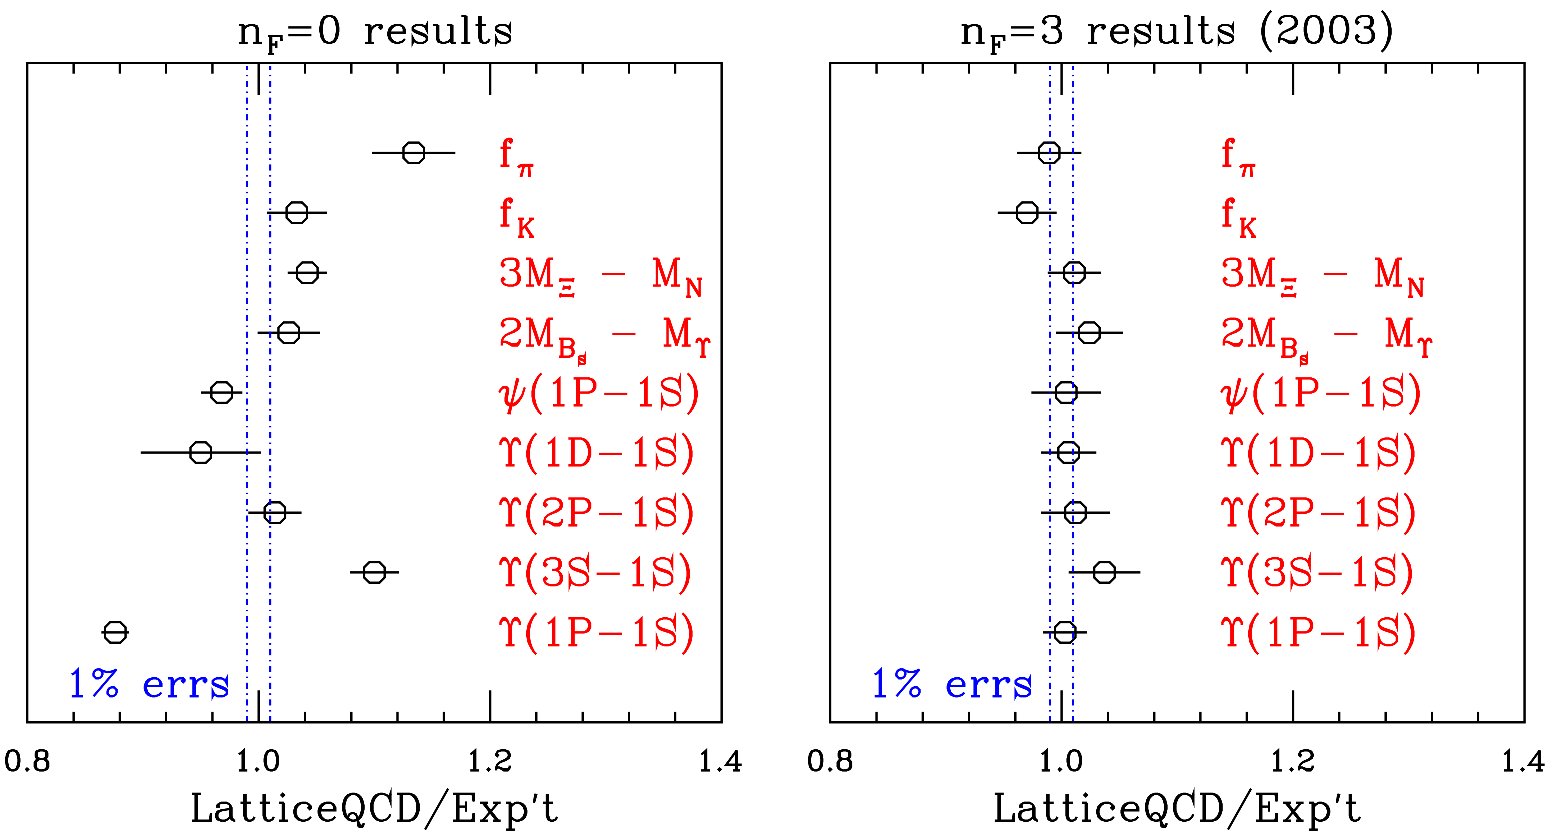

Predictions: 2005–6

Quantity	Lattice QCD	Experiment
$\mathbf{f_D}$	$\mathbf{201 \pm 3 \pm 17}$ MeV	$\mathbf{223 \pm 16 \pm 8}$ MeV
$\mathbf{f_{D_s}/f_D}$	$\mathbf{1.21 \pm 0.01 \pm 0.04}$	$\mathbf{1.25 \pm 0.12 \pm 0.10}$
$\mathbf{m_{Bc}}$	$\mathbf{6304 \pm 20}$ MeV	$\mathbf{6297 \pm 5}$ MeV
$\mathbf{f_B}$	$\mathbf{216 \pm 22}$ MeV	$\mathbf{229 \pm 36 \pm 34}$ MeV

D Meson Semileptonic Form Factor

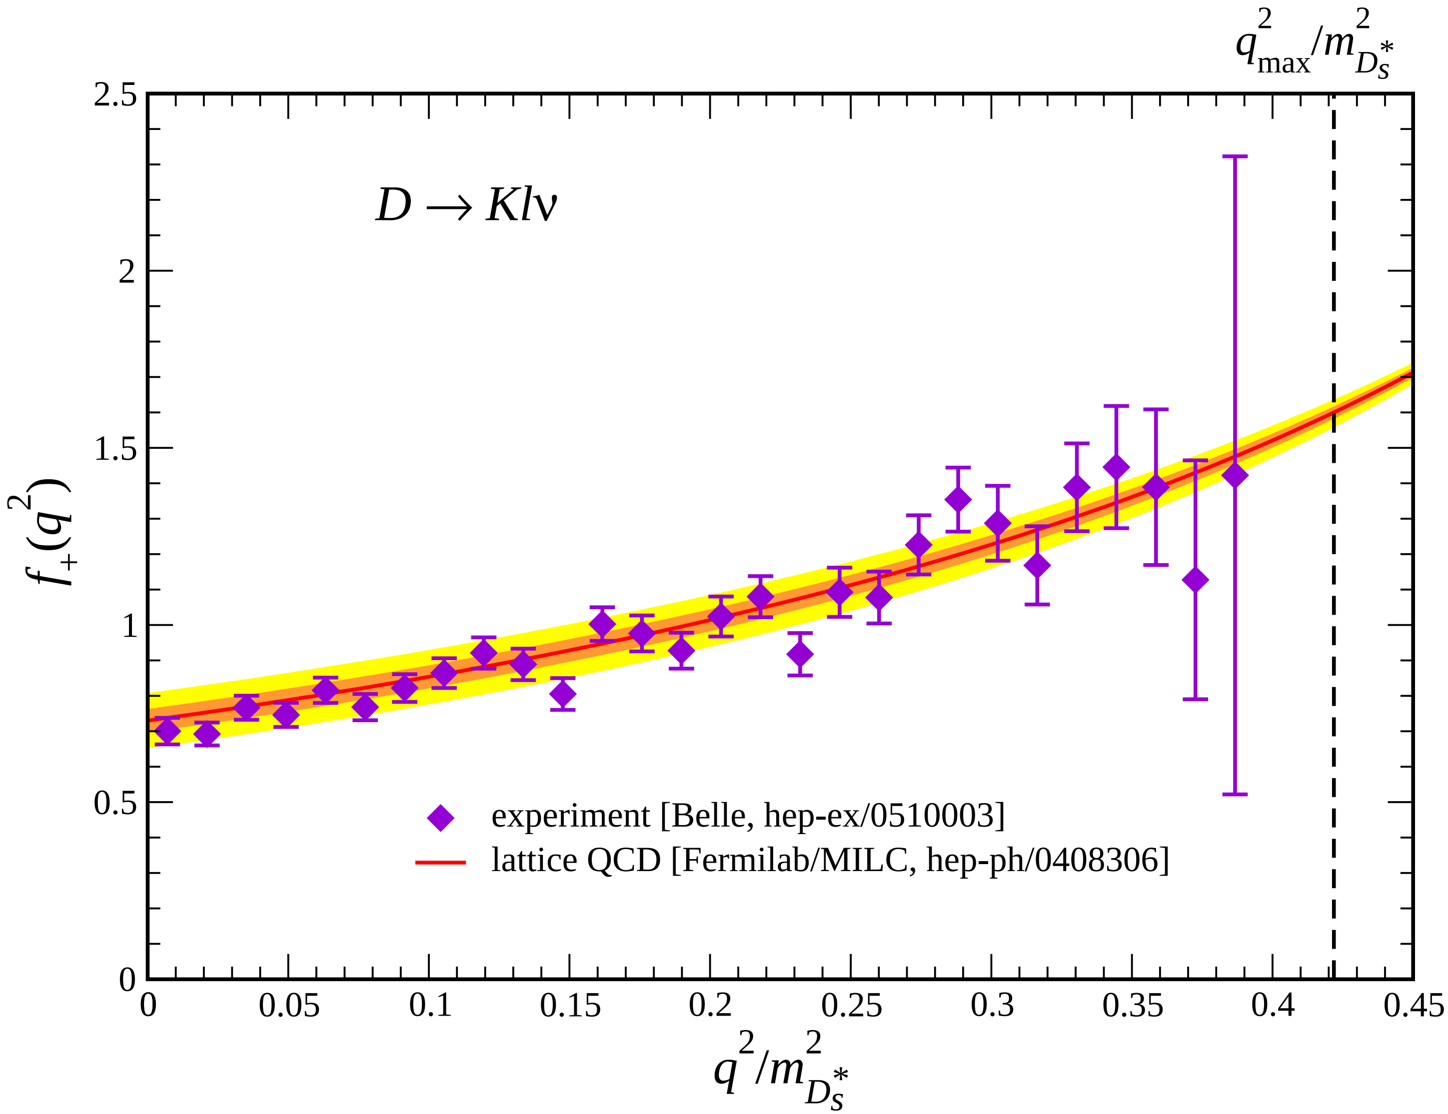

Control of Systematic Errors

To carry out a simulation we must select certain physical parameters:

- lattice spacing (a) or gauge coupling (β)
- grid size ($N_s^3 \times N_t$)
- sea quark masses ($m_{u,d}$, m_s)

To control systematic error we must

- take continuum limit
- take infinite volume limit
- extrapolate to light quark mass; can work at physical s quark mass

MILC Ensembles

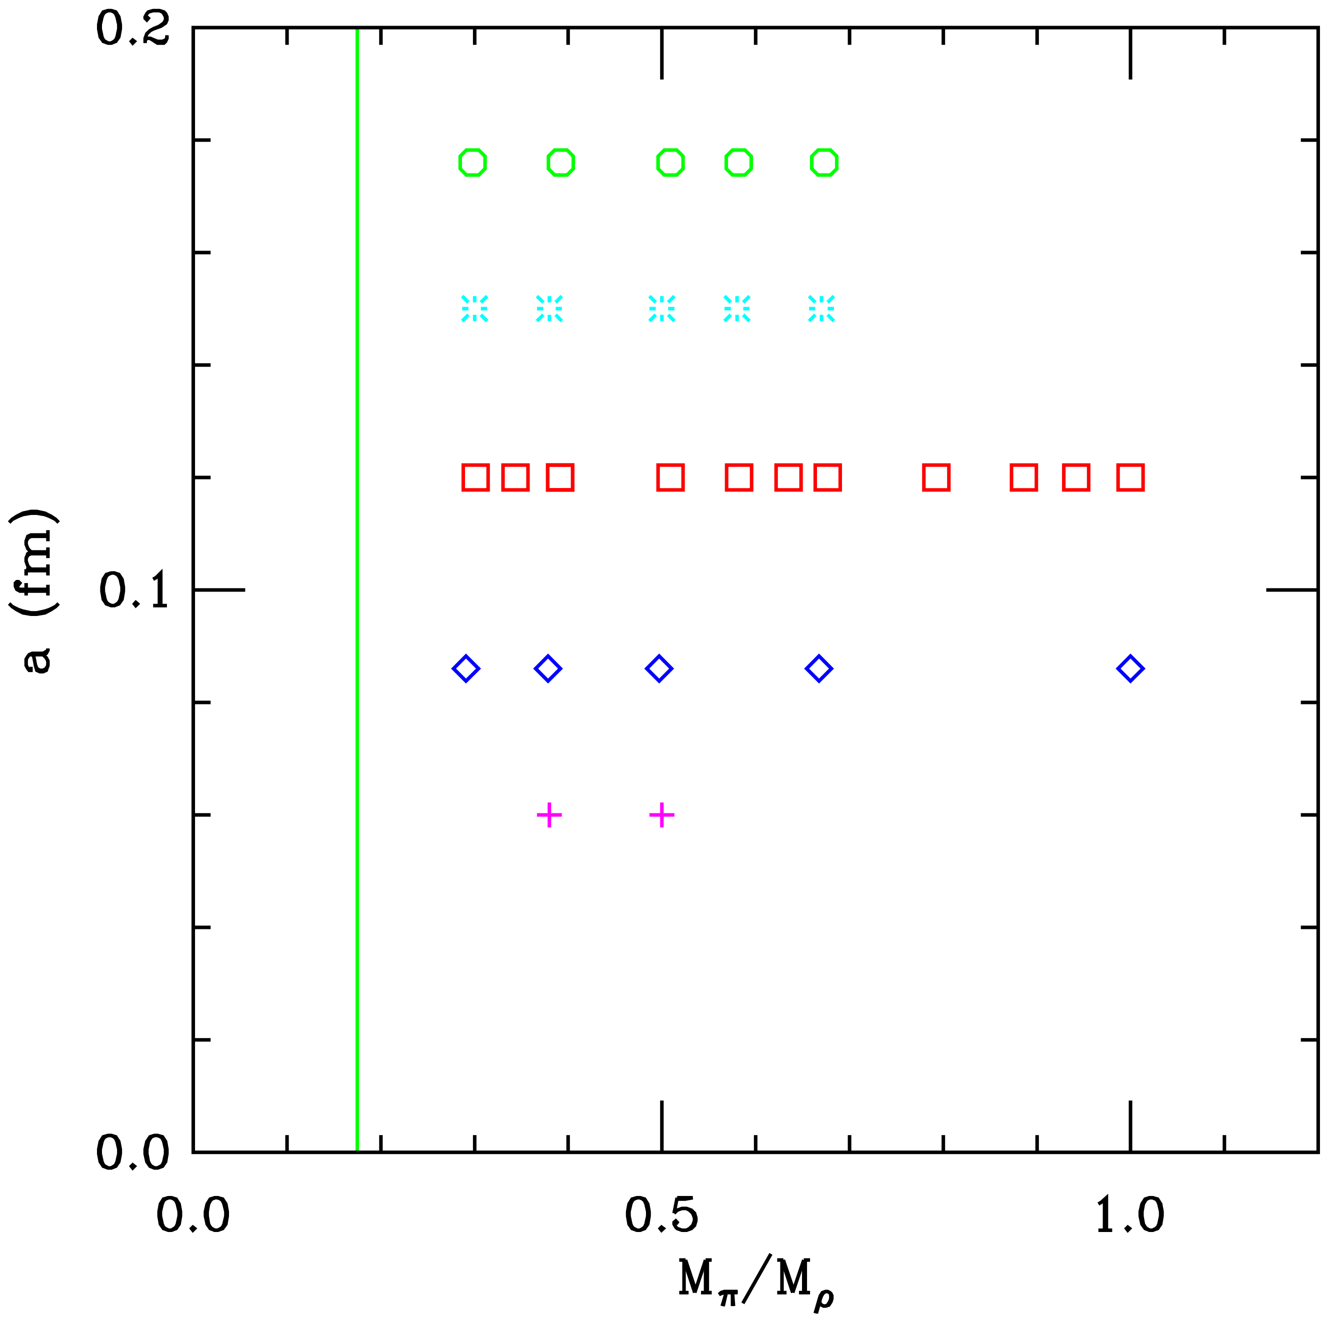

MILC Ensembles

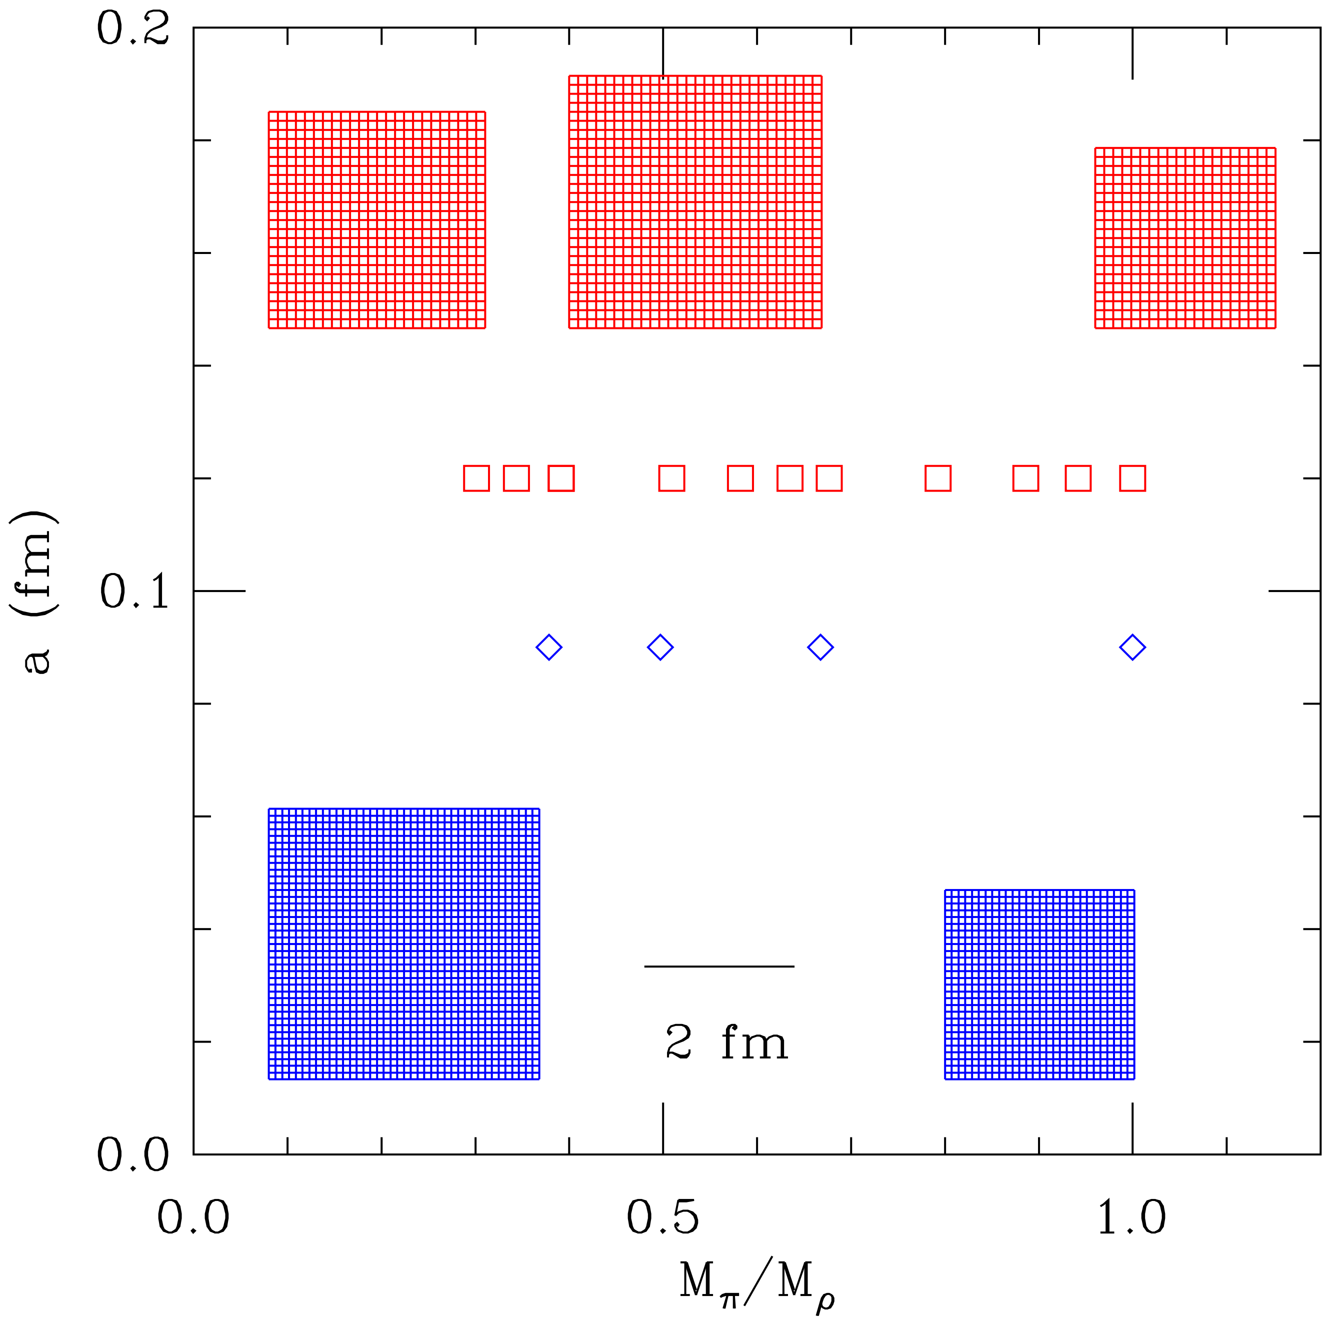

QCD on a Lattice: I

To define lattice QCD, continuum space-time is replaced by a 4d space-time grid.

Quarks described by 3-component complex vectors

Gluons described by 3×3 unitary matrices on the links connecting sites

Parallel transport required for comparing quark fields on different sites

Theory is easily vectorized or parallelized (domain decomposition)

With a regular grid communication patterns are predictable and mostly local. However, global sums are important.

Vector and matrix algebra are optimized via a library.

Assembly code via subroutine or inlining can be very useful.

QCD on a Lattice: II

Basic operations are done on the 3-component vectors and corresponding matrices. Access to memory is an issue:

- SP binary operation: 1 operation; 8 bytes input; 4 bytes output
- Matrix $\times$ vector: 60 ops (36 mult + 30 add); 96 bytes input; 24 bytes output
- 1.45 bytes input/flop and 0.36 bytes output/flop
- Modern CPUs are starved for memory bandwidth
  - cache access and size
  - does prefetching help?

Sparse matrix multiplies take much of the effort, but the operands are these small dense matrices

QCD on a Lattice: III

- Calculations must be carried out for several (small) lattice spacings in order to perform extrapolations to the continuum limit.
- It is too computationally expensive to perform simulations at the physical masses of the two lightest quarks (u and d). So, we work with a range of light quark masses and perform extrapolations to their physical values using chiral perturbation theory.
- It is necessary to increase the physical size of the box in which the simulations are performed as the light quark masses are decreased to avoid finite size errors. Thus, as the lattice spacing and quark masses are decreased, the number of lattice points must be increased.

Plans

- A recent algorithmic improvement called RHMC, based on rational function approximation is expected to reduce our requirements by about a factor of 4.
- It will still take several years to decrease the lattice spacing and to approach the physical light quark mass.
- Other methods for putting quarks on the lattice are more demanding.

a(fm)	$\mathrm{m_l/m_s}$	Lattice	Traj.	TF-Yr
0.06	0.2	$48^3 \times 144$	3,750	1.0
0.06	0.1	$60^3 \times 144$	4,500	2.0
0.06	0.05	$84^3 \times 144$	6,300	23.20
0.045	0.4	$56^3 \times 192$	4,000	0.6
0.045	0.2	$56^3 \times 192$	5,000	1.9
0.045	0.1	$80^3 \times 192$	6,000	13.7

Planned Runs

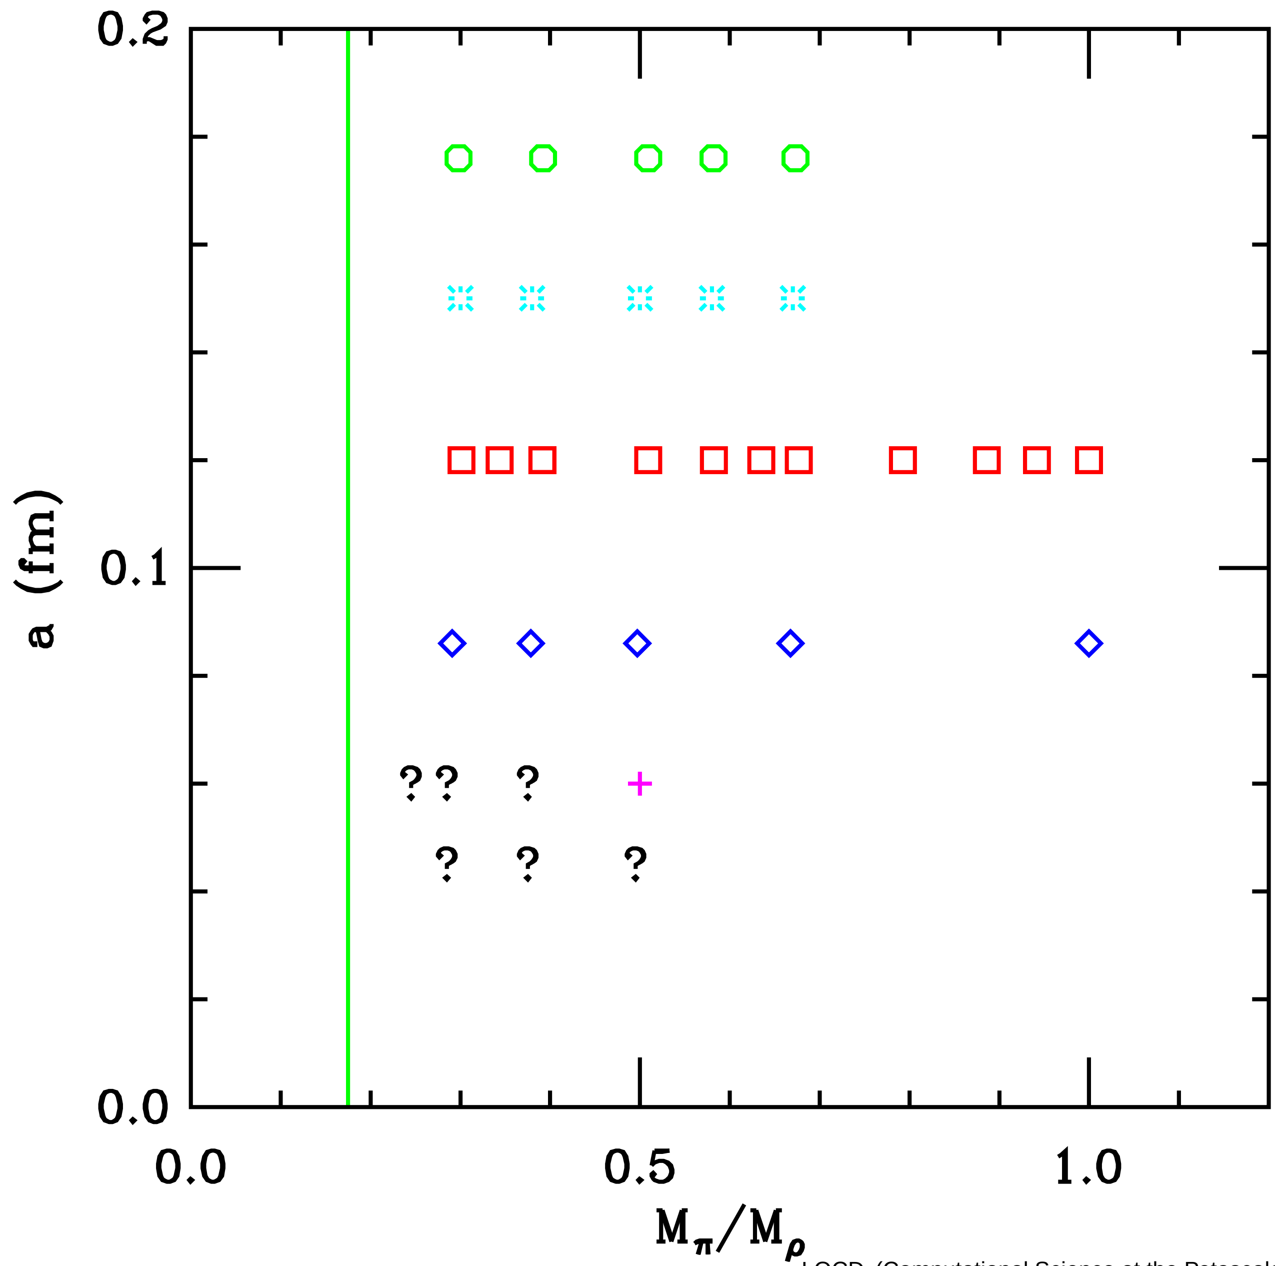

Performance I

- There are several major parts of the code:
  - Conjugate Gradient Solver (CG)
  - Fermion force contribution (FF)
  - Gauge force contribution (GF)
  - Calculation of “fat links” (Fat)
  - Calculation of “long links” (Long)
- Access to memory is a bottleneck, so single node performance is sensitive to the number of grid point. Thus, we like to do weak scaling benchmarks with L^4 sites per cpu.
- CG is the dominant part of code, but other parts are significant.

Performance II

Time distribution for a run on 2048 XT3 (BigBen) cpus using a $40^3 \times 96$ grid ($5 \times 10^2 \times 6$ per cpu) with $m_l = 0.1 m_s$:

Activity	time(s)	MF/cpu	per cent
CG	2987	530	58.5
FF	1125	579	22.0
GF	489	469	9.5
Fat	442	627	8.7
Long	24	340	<1
Input config.	41		<1
total above	5108		
unaccounted	104		1.9
wallclock	5212		

Performance vs CPUs

CG runs at 7.7 TF/s on 8192 cores for $112^3 \times 224$ grid!

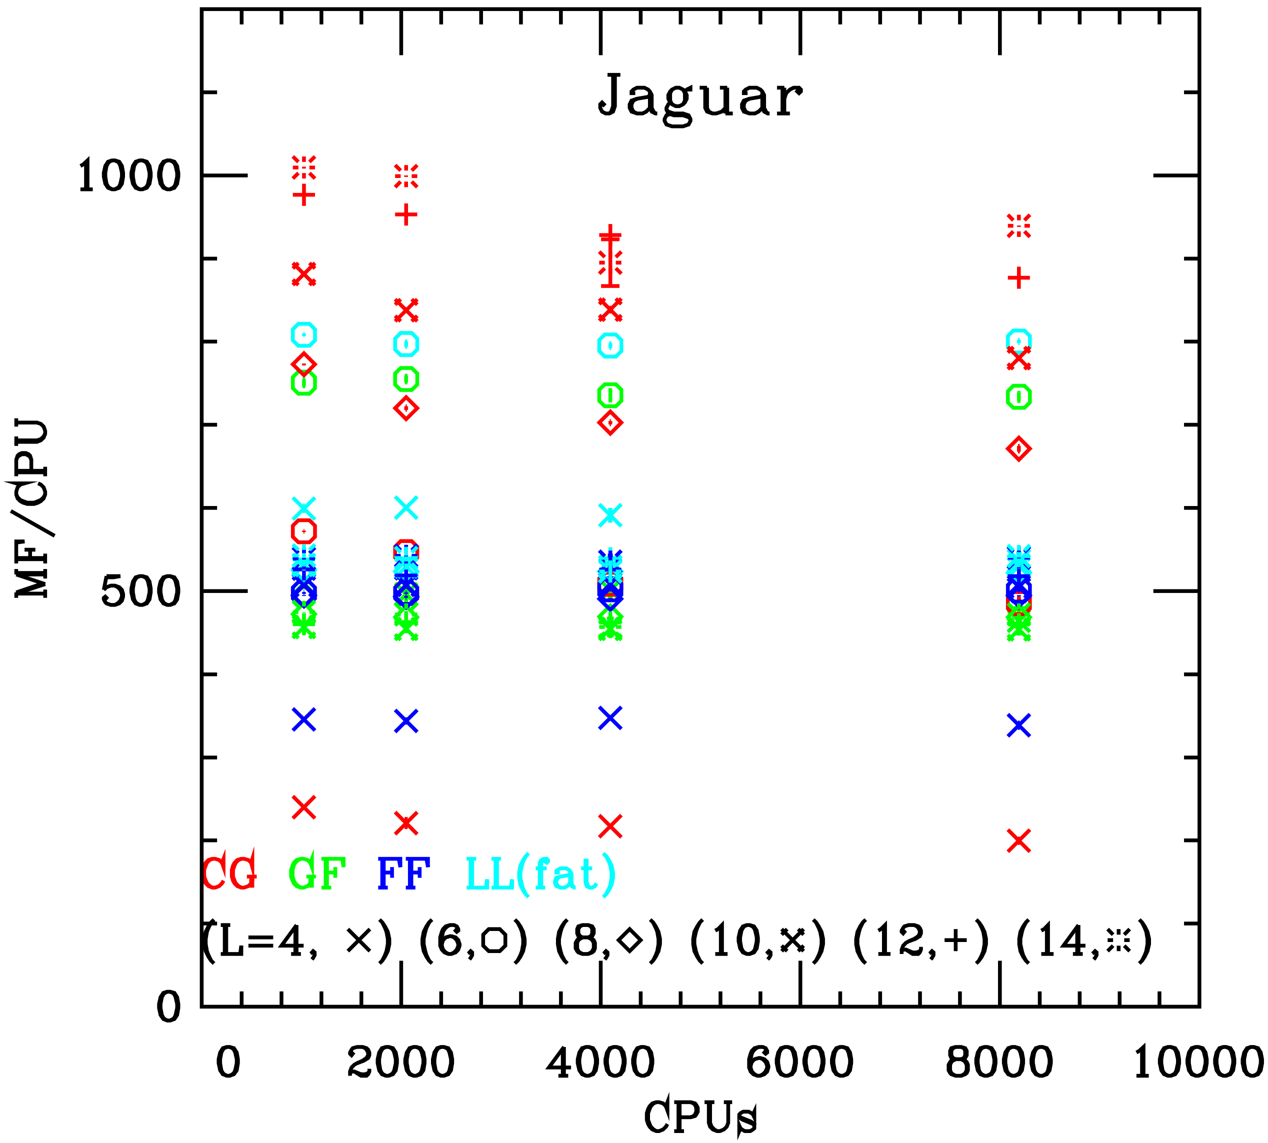

Performance vs L

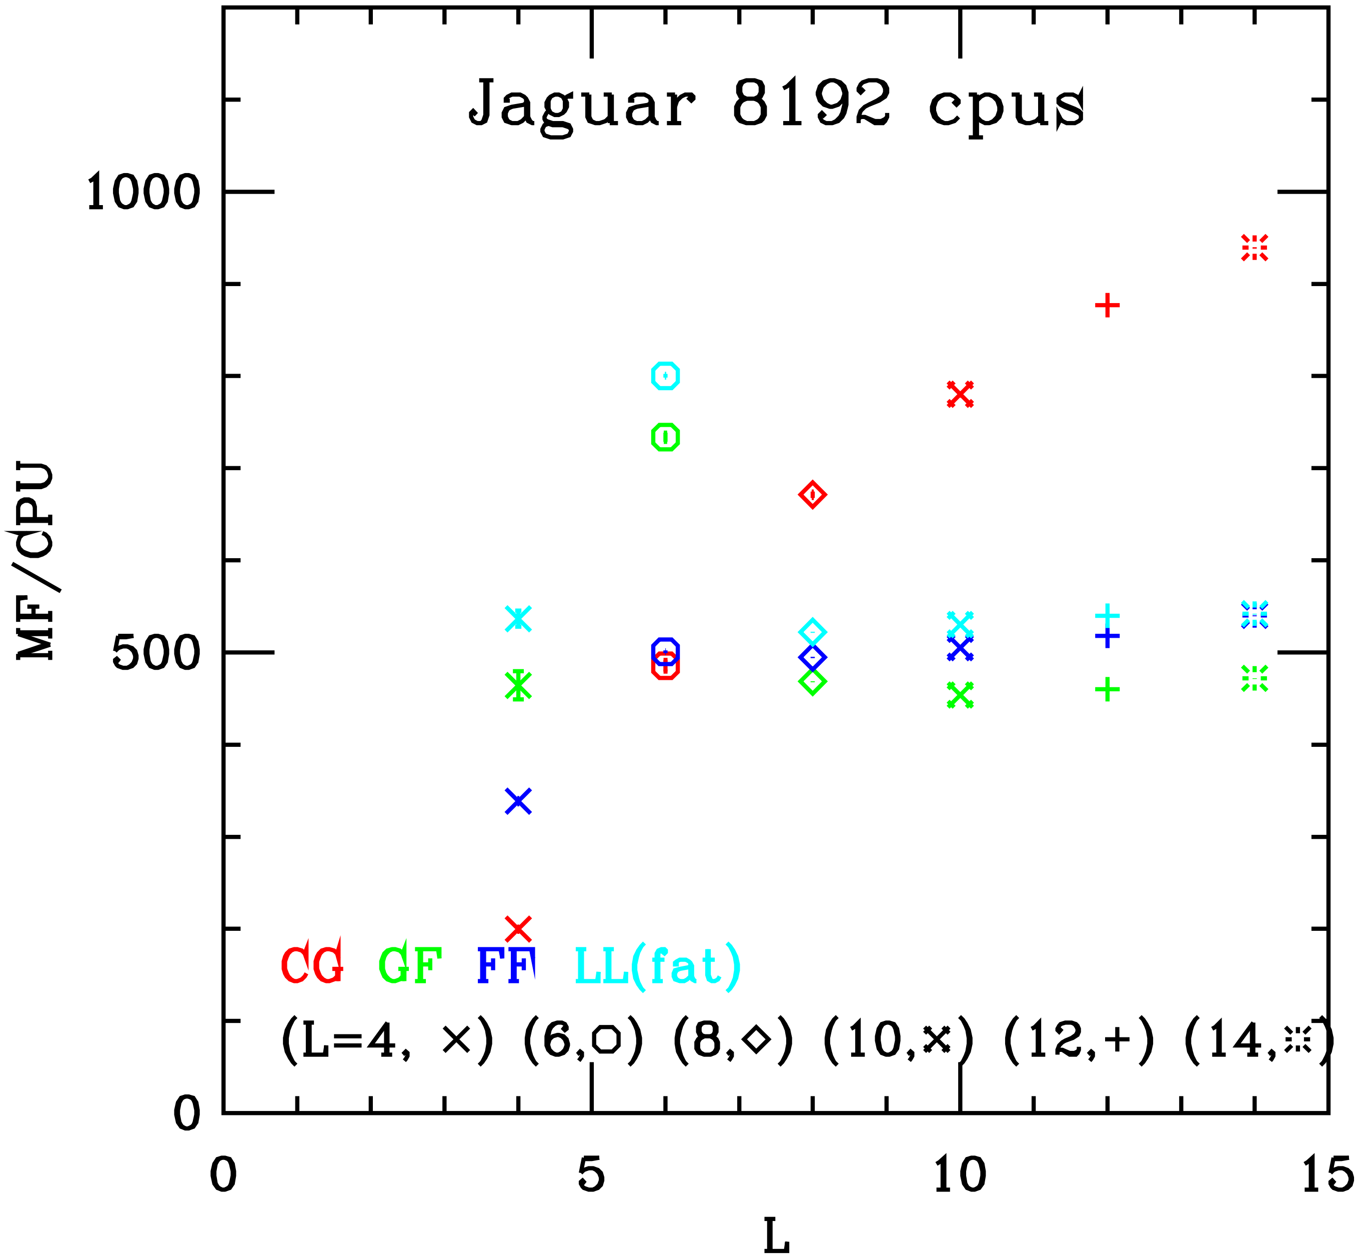

Performance Challenges

- Conjugate gradient is very important particularly for light quarks
- Using the new RHMC algorithm, the CG might not be as dominant
- Nearest neighbor communications are used for most messages, but global sum in CG can be a very important bottleneck, particularly for small L
- How many cpus will we need to use for petascale jobs?
- Can we use multiple cores more efficiently than just as separate MPI processes?
- Will network latency decrease as clock cycle does?