The USQCD Collaboration

Overview

USQCD Collaboration

http://www.usqcd.org

Paul Mackenzie

Fermilab

mackenzie@fnal.gov

Lattice QCD Computational Science Workshop
Oak Ridge
April 29-30, 2013

The USQCD Collaboration

- Organizes computing hardware and software infrastructure for lattice gauge theory in the US.
- Represents almost all of the lattice gauge theorists in the US; ~ 150 people.
  - ~ 100 participating in physics proposals this year.
  - ~ 30 have served on the Executive of Scientific Program Committees.
- Physics calculations are done by smaller component collaborations within USQCD:
  - Fermilab, HotQCD, HPQCD, JLab, LHPC, LSD, MILC, NPLQCD, RBC, ...

Major areas of physics research

Lattice QCD calculations are essential to accomplishing the physics goals of high energy and nuclear physics. The USQCD physics program is driven by the experimental physics programs of the national labs and DoE- supported experiments.

- HEP Intensity frontier (Fermilab, SLAC, BNL, Belle-2, LHC-b)
  - Hadron spectrum; determining the parameters of the standard model: the CKM matrix, the quark masses, and the strong coupling constant, ..., and searching for inconsistencies due to beyond-the-standard-model physics.
- HEP Energy frontier (LHC)
  - Beyond the standard model: search for new particles and forces not yet discovered, ...
- NP Quark-gluon plasma in heavy ion collisions (RHIC)
  - De-confinement temperature; QCD plasma equation of state; transport coefficients (viscosity, ...)
- NP Hadronic and nuclear structure and interactions (JLab)
  - Resonance and exotics spectra, scattering lengths, and phases shifts; hadronic structure, ...

Science drivers: Fermilab and the LHC (HEP)

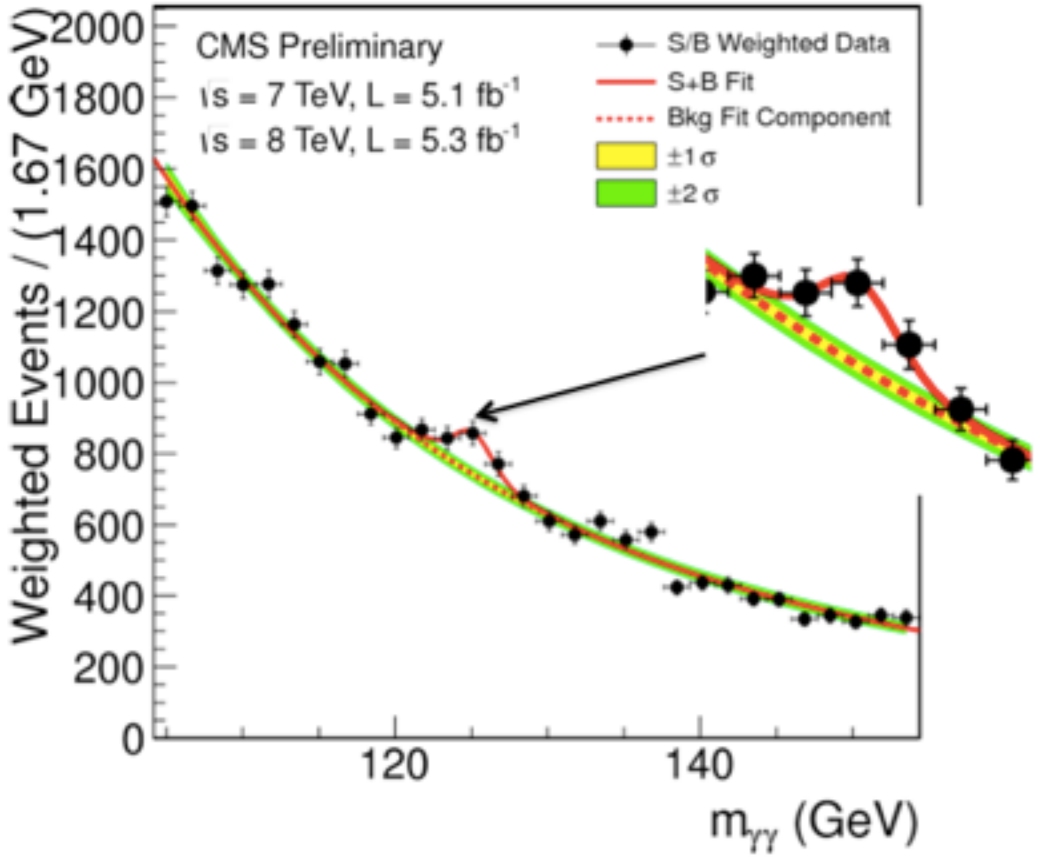

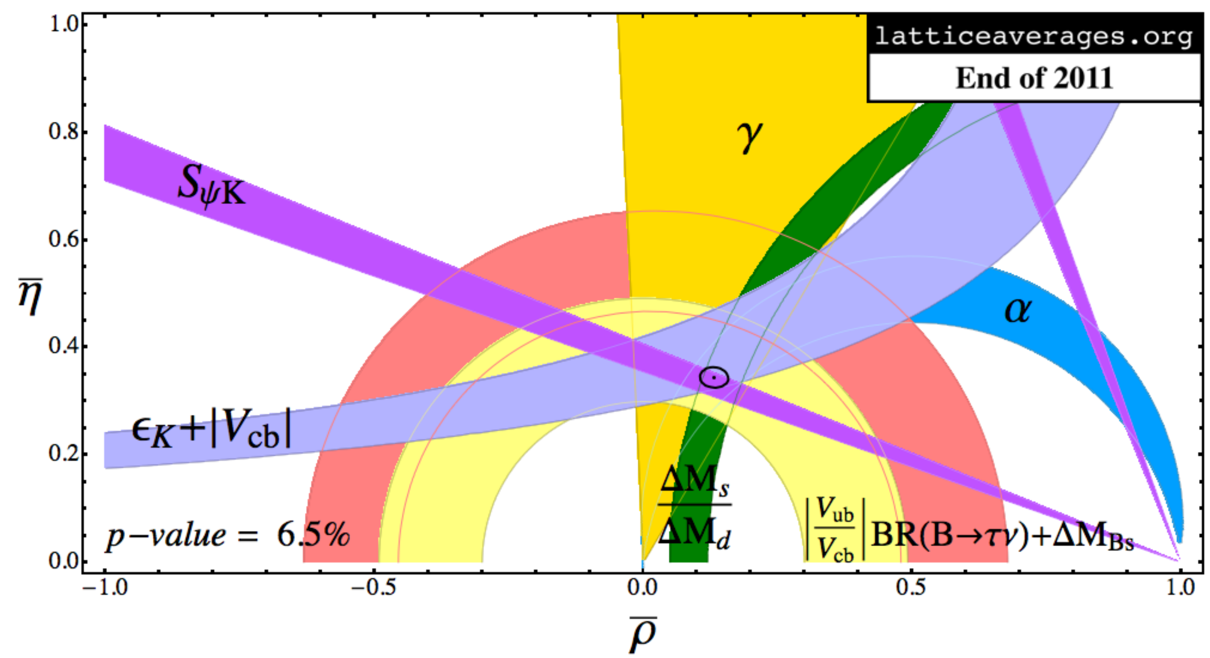

Laiho, Lunghi, and Van de Water

Fermilab director Pier Oddone: Lattice QCD calculations will make the data we obtain from quark factories (both electron-positron colliders as well as the Tevatron and LHC) far more useful in determining the fundamental parameters of the standard model and revealing any model inconsistencies indicative of new physics. For example, the existence of good lattice calculations allowed Fermilab's discovery of *BsBs*-bar mixing to make an important bound on the CP violating elements of the CKM matrix. Much more accurate calculations of this and other quantities are now needed to make full use of the data from the Fermilab's program

The Standard Model is *maddeningly* successful. It accounts for every particle physics experiment performed so far, sometimes to great precision (one part in a billion for the electron anomalous magnetic moment).

But, it contains obvious gaps and puzzles! Now that the “Higgs”, the last undiscovered element of the standard model seems to have been discovered, the search is on for physics “beyond-the-standard-model”. Why? It leaves many puzzles and unanswered questions.

- What is the relation between the ~20 seemingly random free parameters?
- Why is there more than one generation of quark?
- What is the relation between the three forces?
- How can gravity be incorporated?
- What is the dark matter?
- ...

The search for a more fundamental theory underlying the Standard Model is the central task of particle physics today.

Search for BSM physics proceeds on two frontiers.

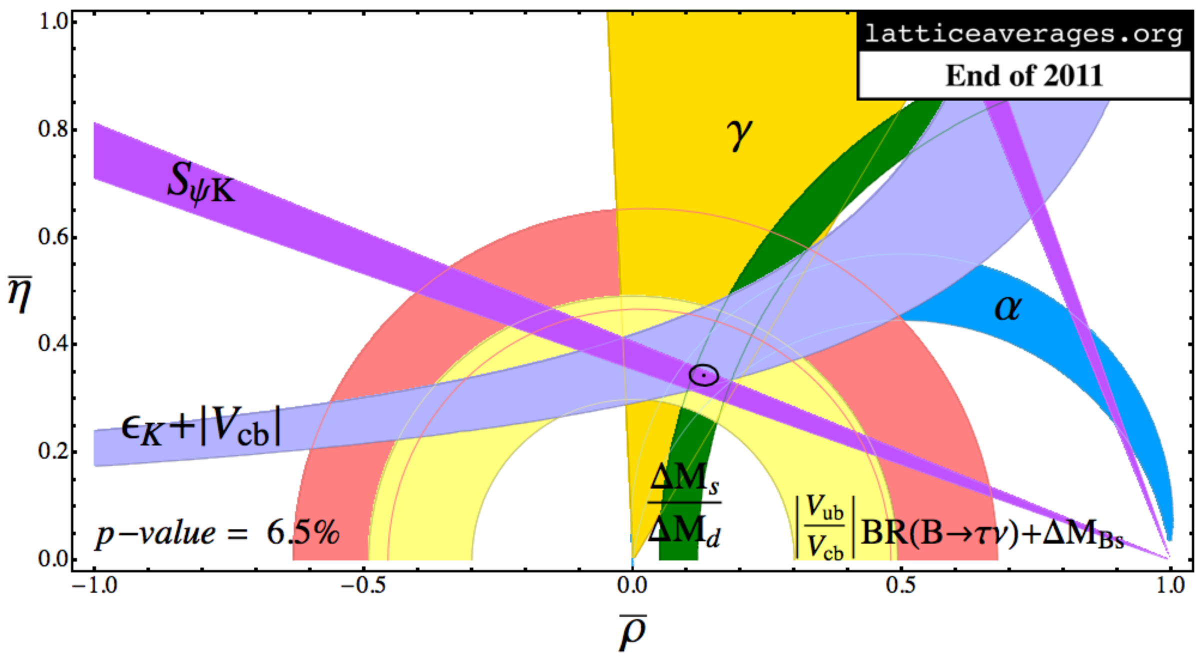

Intensity frontier:
Search for inconsistencies in determinations of standard model parameters due to BSM effects.

Detar and Izubuchi talks.

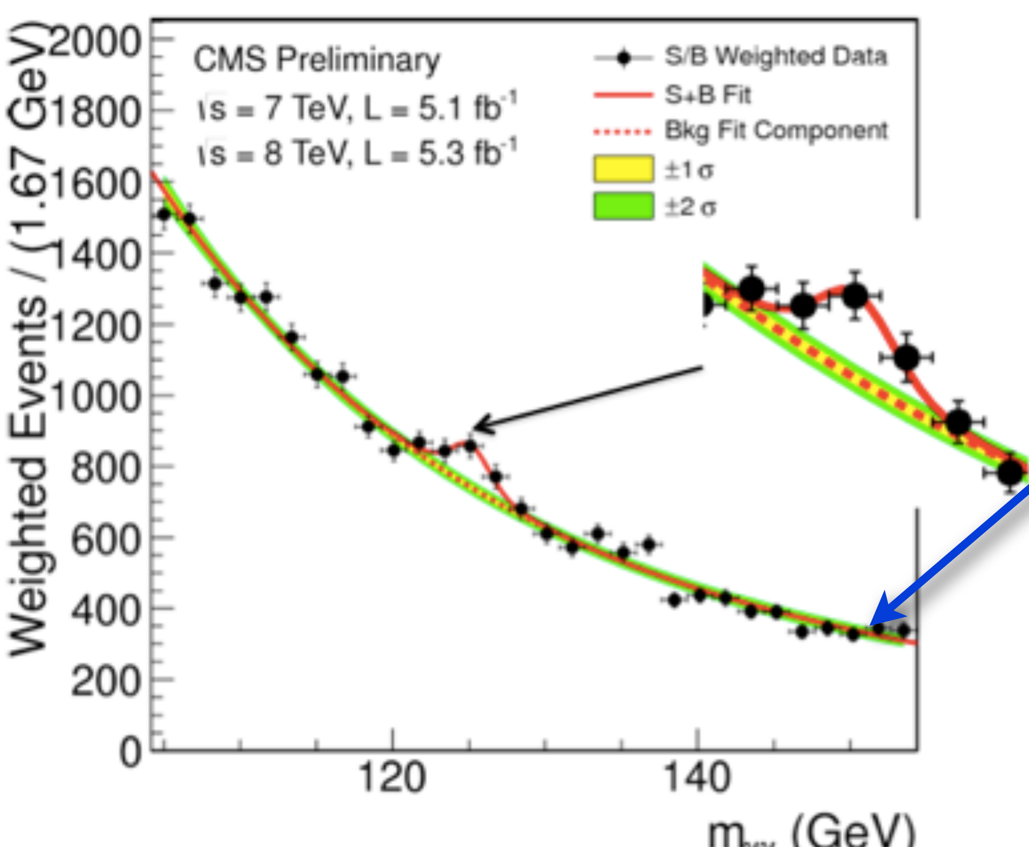

Energy frontier:
Experimental search for new bumps beyond the "Higgs".
Theoretical examination of whether the "Higgs" is something other than the standard model Higgs.

Kuti talk.

Science drivers: NP

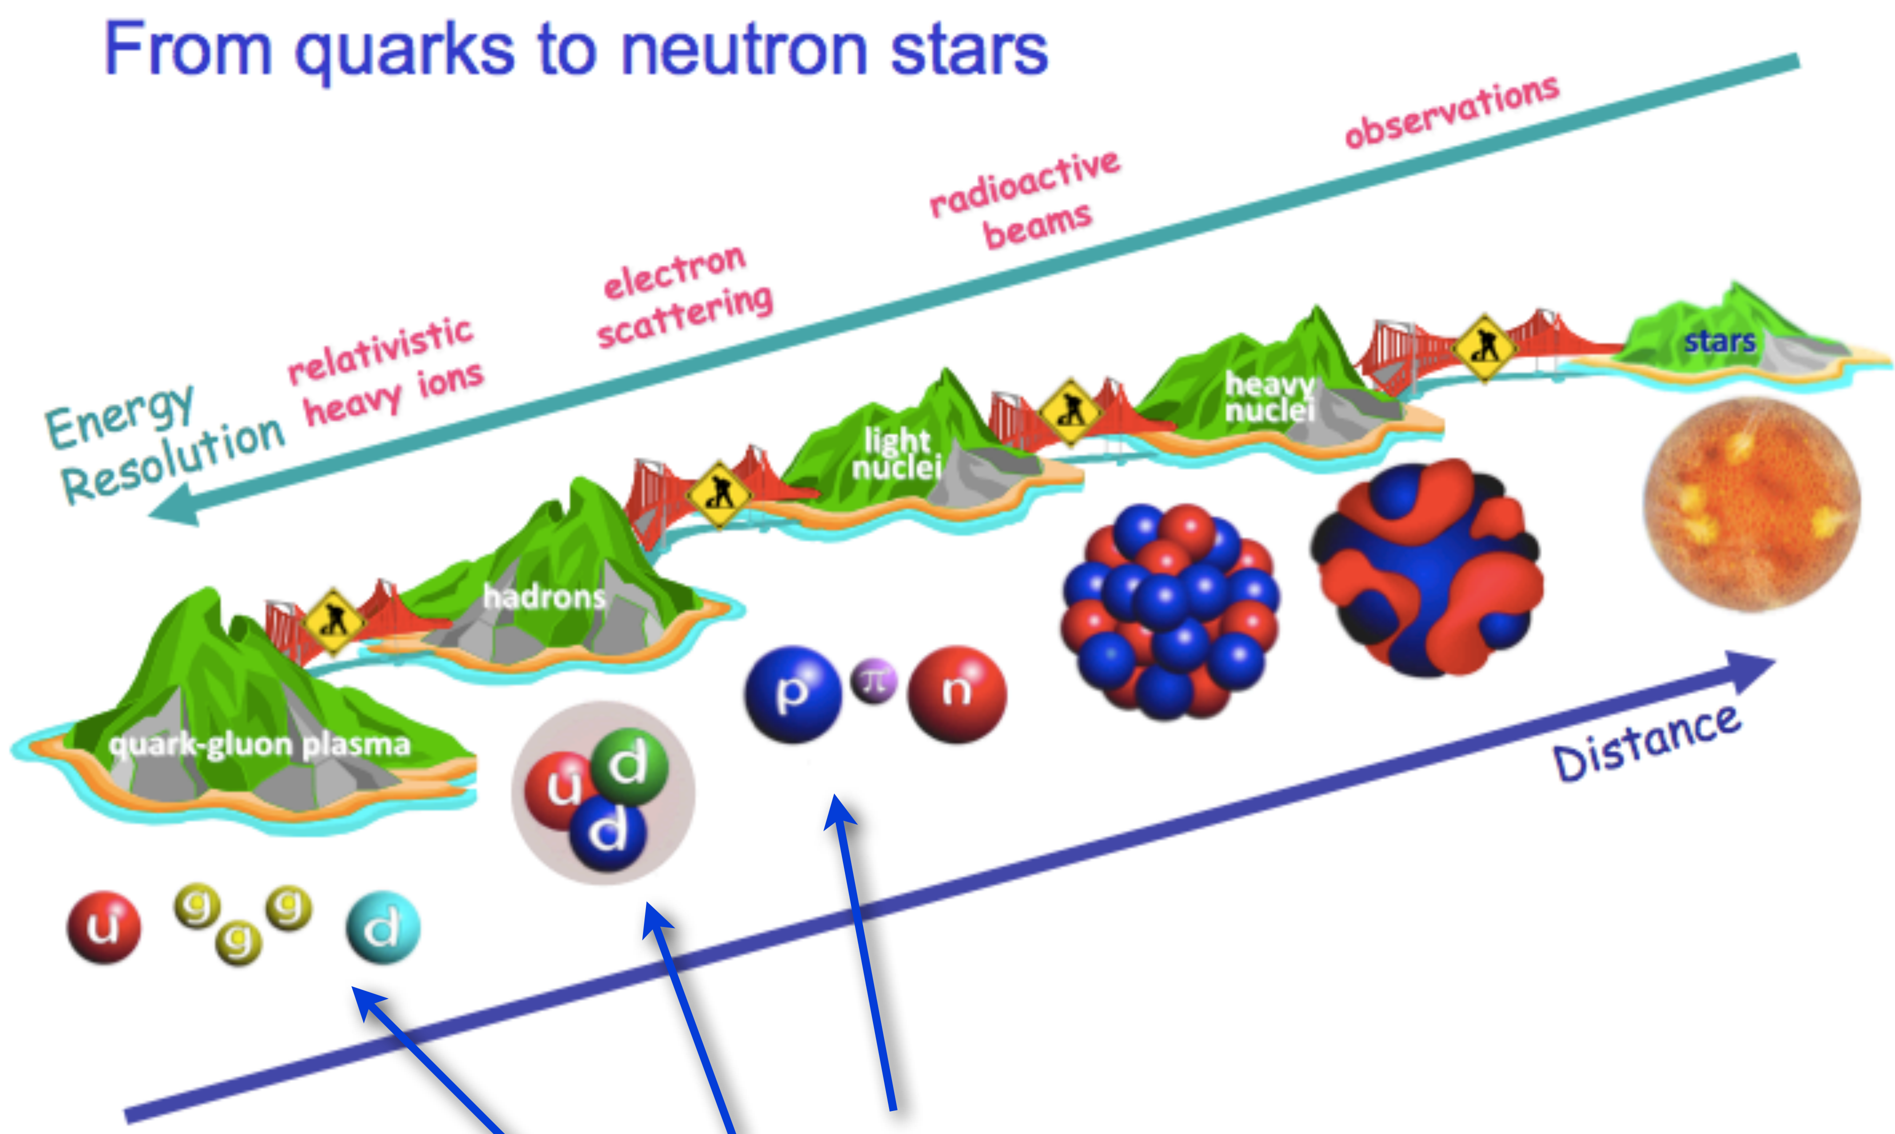

Witold Nazarewicz

The goal of lattice QCD is to derive the foundations of nuclear physics from first principles.

Orginos, Detmold talks.

Science drivers: Brookhaven National Lab (NP)

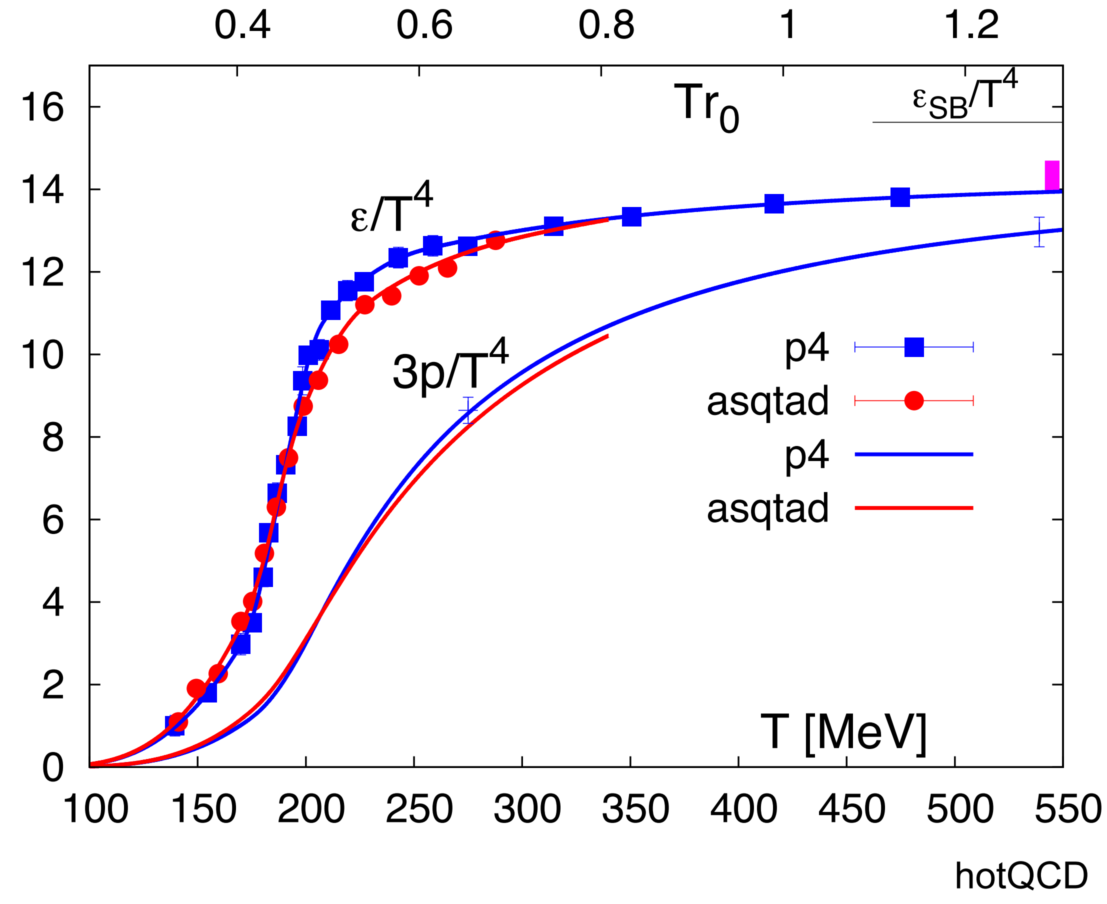

BNL Physics Department head Tom Ludlam: The lattice QCD calculations performed at BNL have direct relevance for the experimental program at RHIC, where an accurate determination of the equation of state of dense QCD matter with lattice gauge calculations is of central importance to the understanding of hydrodynamic properties from experimental data. In addition, we are counting on the USQCD research program to provide guidance in the search for a QCD critical point in heavy ion collisions, and an understanding of the properties of strongly interacting matter near this landmark point on the QCD phase diagram.

Petreczky, Mukherjee, Schmidt talks.

Science drivers: Jefferson Lab (NP)

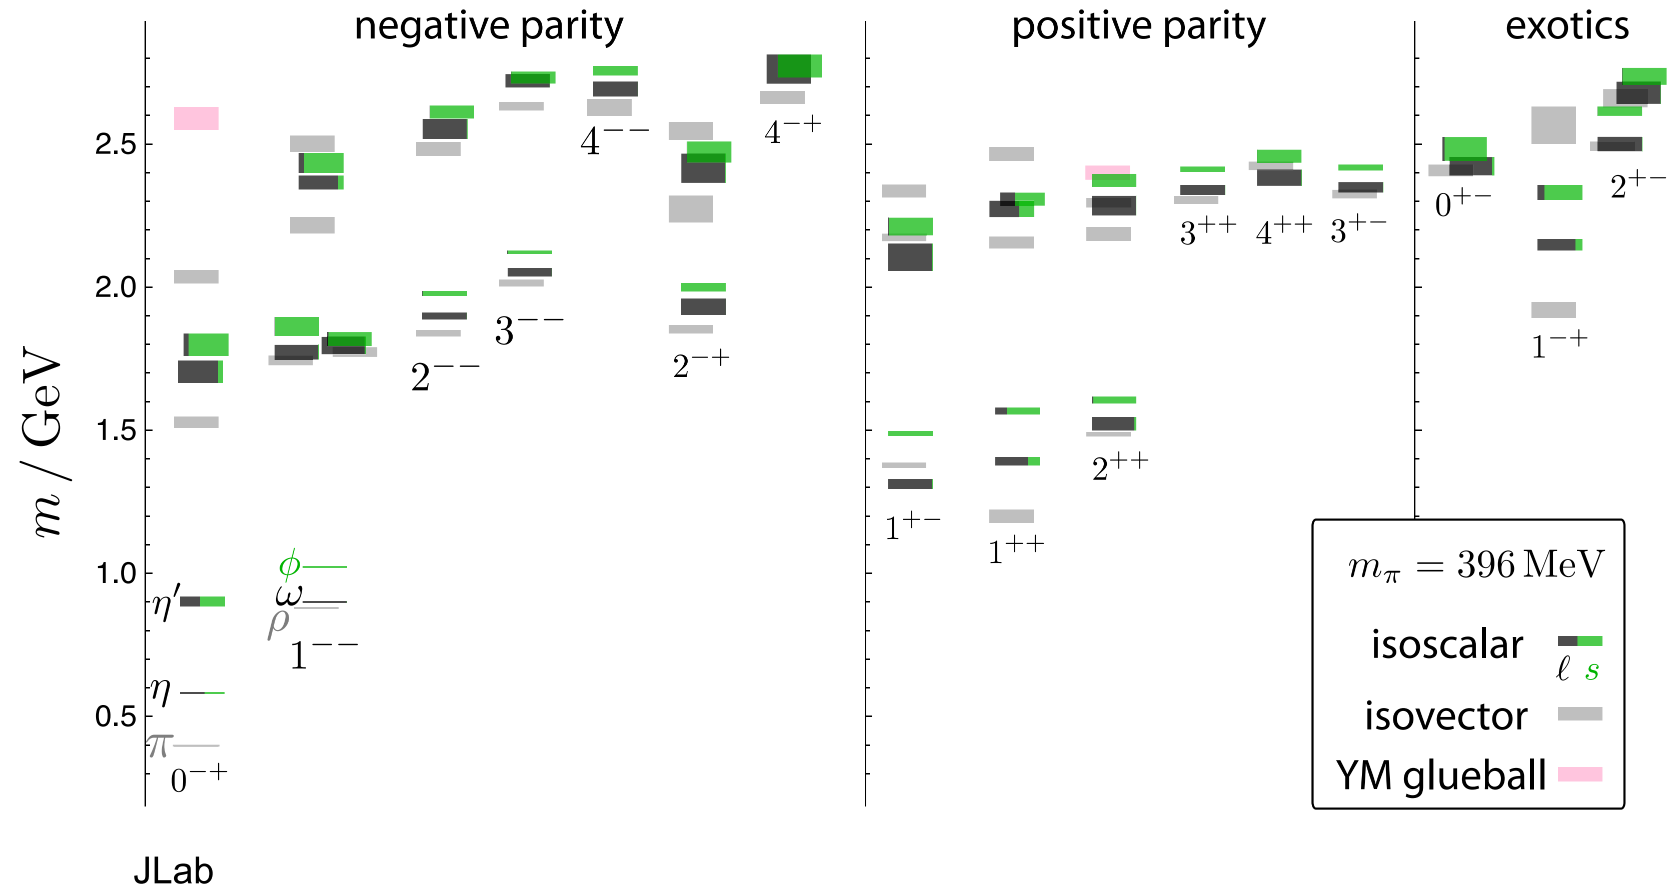

The structure and interactions of hadrons (nucleons and mesons).

JLab director Hugh Montgomery: The national efforts of the USQCD collaboration are key to the success of the lattice program at Jefferson Lab... A continued strong national program will ensure both the algorithmic developments, and the software infrastructure, to further exploit both frontier leadership-class and special-purpose computers, and thus provide the calculations that will capitalize on the DOE investment in the Jefferson Lab experimental program.

Edwards, Savage talks.

USQCD organization

On behalf of the US lattice community, USQCD oversees:

- SciDAC software grant (OHEP, ONP, OASCR).
  - Began 2001; ~$2.1 M/year currently.
  - Creates community libraries, optimized production programs, research on new approaches (GPUs are hot now), ...
- Community INCITE grants on ASCR Leadership Computing Facilities for capability computing.
- Design and deployment at national labs of cost-efficient capacity hardware funded by LQCD-ext Project (OHEP and ONP).
  - ~ $4 M/year. (Why? Coming later.)
  - Infiniband, GPU clusters.

USQCD Executive Committee

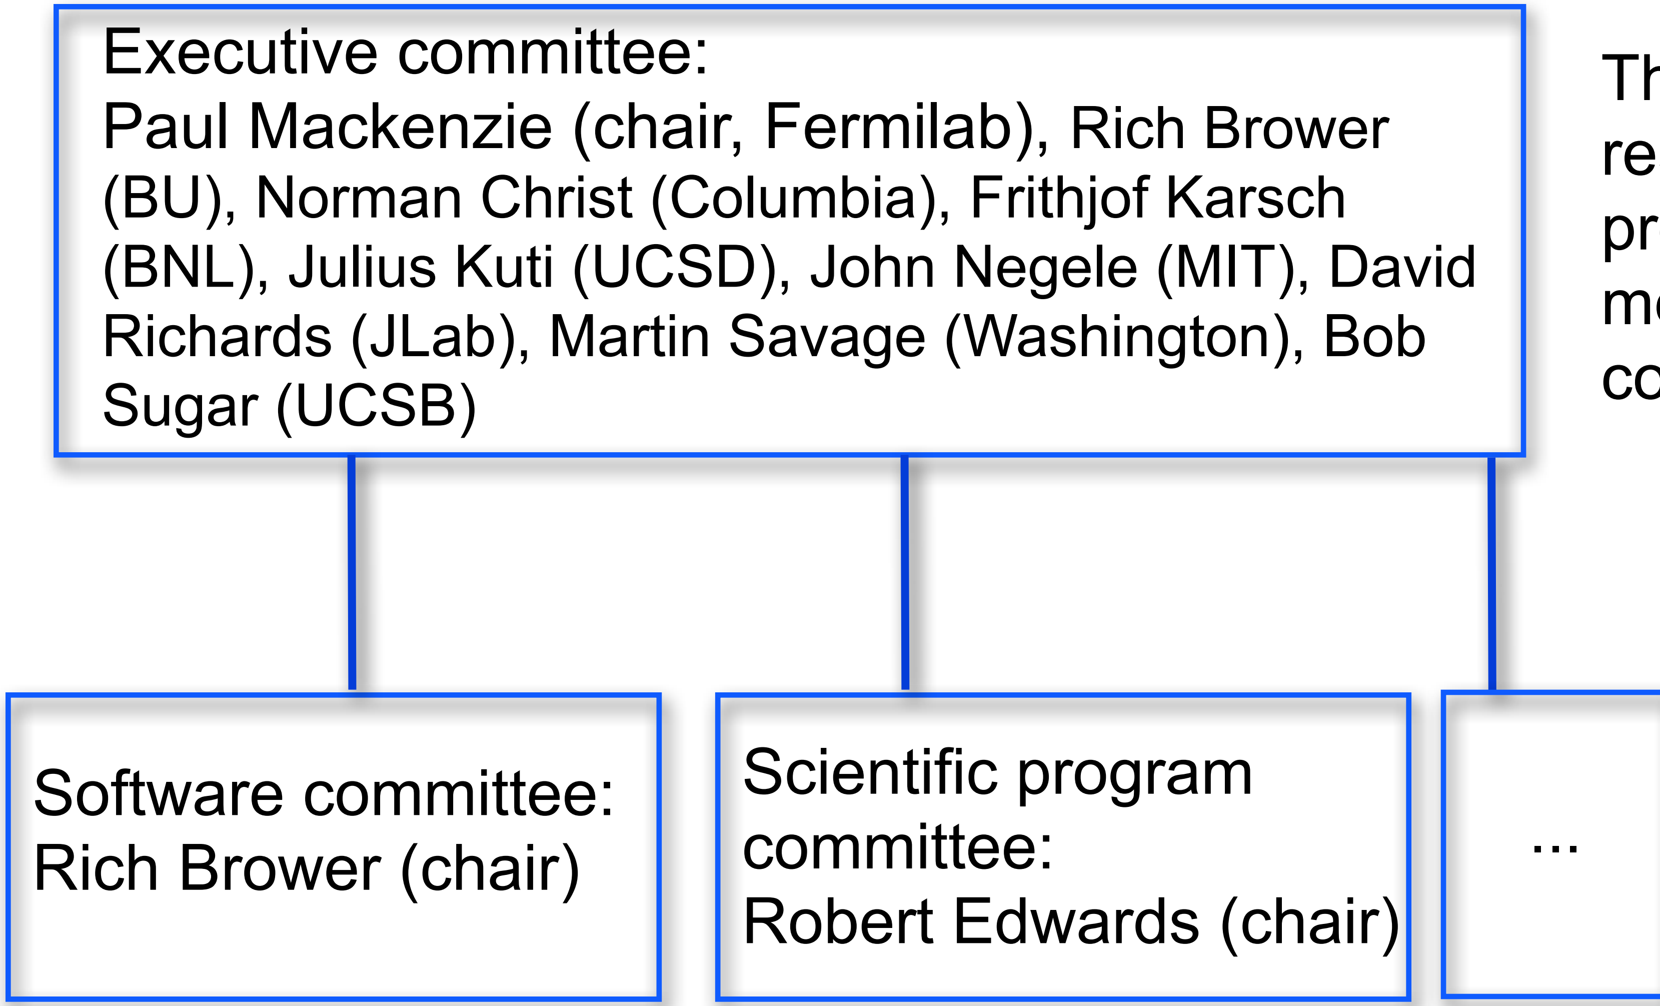

The Executive Committee is responsible for writing USQCD's proposals and for appointing the members of the other committees.

USQCD Software Committee

Software Committee:
Richard Brower (chair, BU), Carleton DeTar (Utah), Robert Edwards (JLab), Rob Fowler (UNC), Donald Holmgren (Fermilab), Robert Mawhinney (Columbia), Pavlos Vranas (LLNL), Chip Watson, (JLab).

- Organizes software work done under our SciDAC grant.
- Weekly conferences calls with 12-20 people, 40 people on mailing list.
- SciDAC grant pays for less than half of our software work.
  - $2.1 M/year, ~12 FTEs.
  - Much of the work of the software program is done by people on their regular salaries working to accomplish the goals of their physics collaborations.

USQCD timeline

- USQCD formed in 1999.
- SciDAC grants since 2001.
  - In SciDAC-3, grants from NP and HEP for about $1M each.
  - Essential for making effective use of Leadership Computing Facilities and our dedicated hardware, and for accomplishing our physics objectives.
- Capacity dedicated hardware grants from HEP and NP.
  - Installed at JLab, Fermilab, and BNL.
  - This year allocated 340 M cluster core-hours, 8 M GPU hours.
- INCITE grants since 2008.
  - This year, 140 M hours on Titan at OLCF;
  - 290 M hours on Mira and Intrepid at ALCF.

The lattice QCD computing problem

Lattice QCD approximates the continuum theory by defining the fields on a four dimensional space-time lattice.

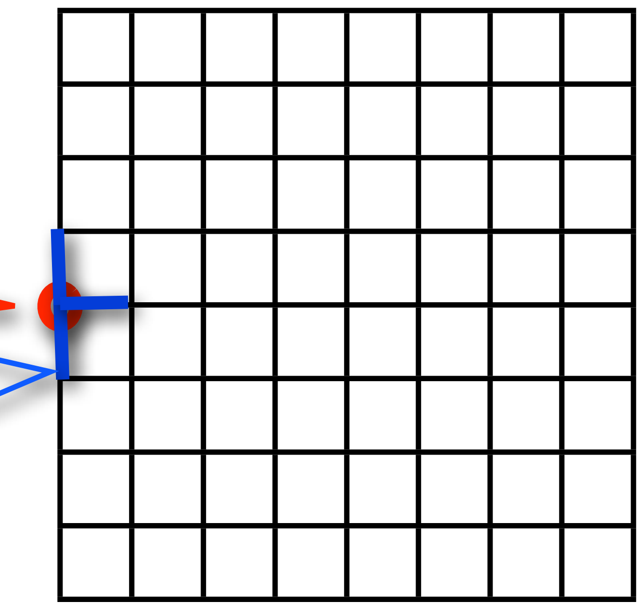

Quarks (*complex three vectors*) are defined on the sites of the lattice, and the gauge field gluons (*complex 3x3 matrices*) on the links.

Monte Carlo methods are used to generate a representative ensemble of gauge fields. Relaxation methods for sparse matrices are used to calculate the propagation of quarks through the gauge field.

Continuum quantum field theory is obtained in the zero lattice spacing limit. This limit is computationally very expensive.

Algorithms and methods

An ensemble of gauge configurations is created with Monte Carlo methods with symplectic Hamiltonian integration. A Markov chain of configurations is made, each one from the previous.

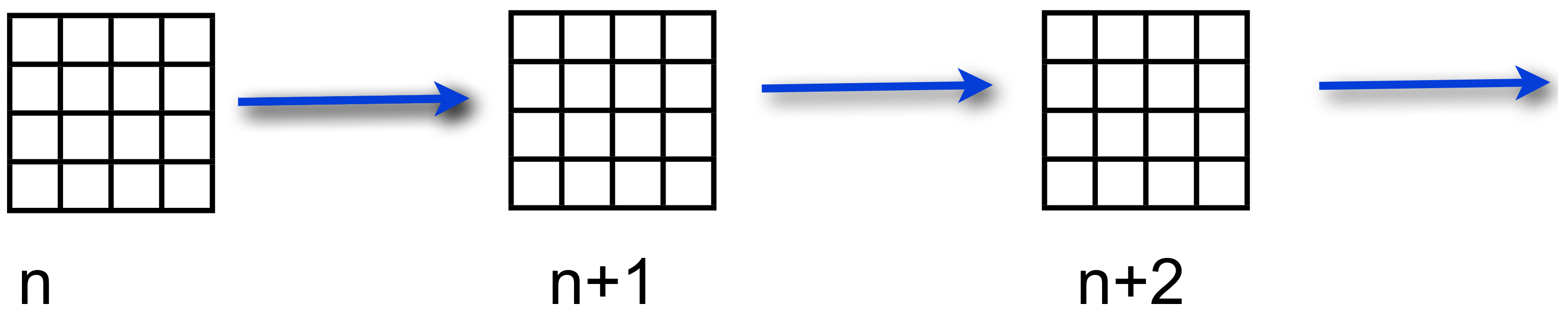

A capability computing job.

Once created, each configuration can be analyzed in parallel.

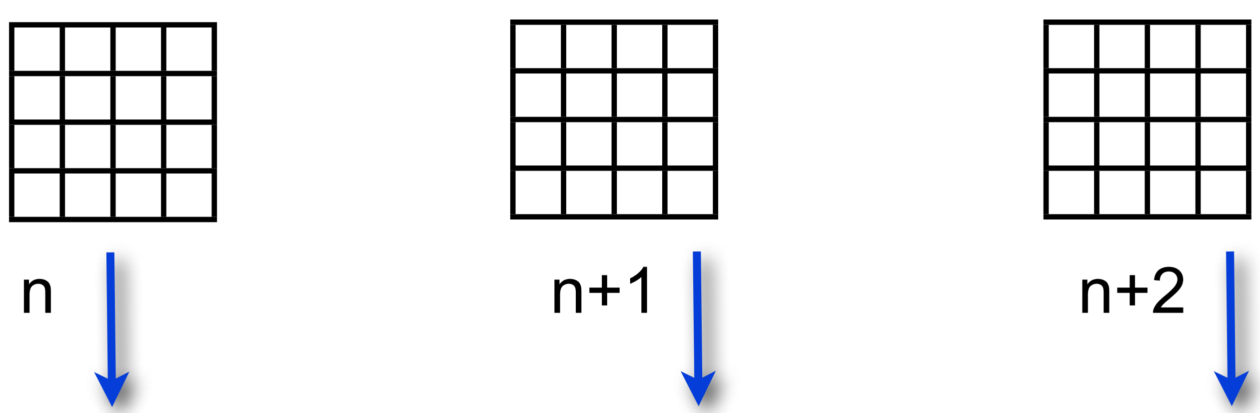

A capacity computing job.

The main numerical component of both jobs is solving a sparse matrix equation $Ax=b$,
with, for example, the bicongradstab algorithm.

Two main components of a typical lattice calculation

multi-TB
file sizes

Generate *O*(1,000) gauge configurations on a leadership facility or supercomputer center. Hundreds of millions of core-hours.

Transfer to labs for analysis on clusters. Larger CPU requirements.

Gauge configuration generation:
a single highly optimized program,
very long single tasks,
“moderate” I/O and data storage.

Hadron analysis.
Large, heterogeneous analysis code base,
10,000s of small, highly parallel tasks,
heavy I/O and data storage.

Two comparably sized jobs with quite different hardware requirements.

- Leadership class computing is essential for generating large ensembles of gauge configurations. This computing cannot be done any other way.
- We have an even greater need for flops analyzing these configurations.
    - Can often be done very efficiently in parallel on much smaller systems.

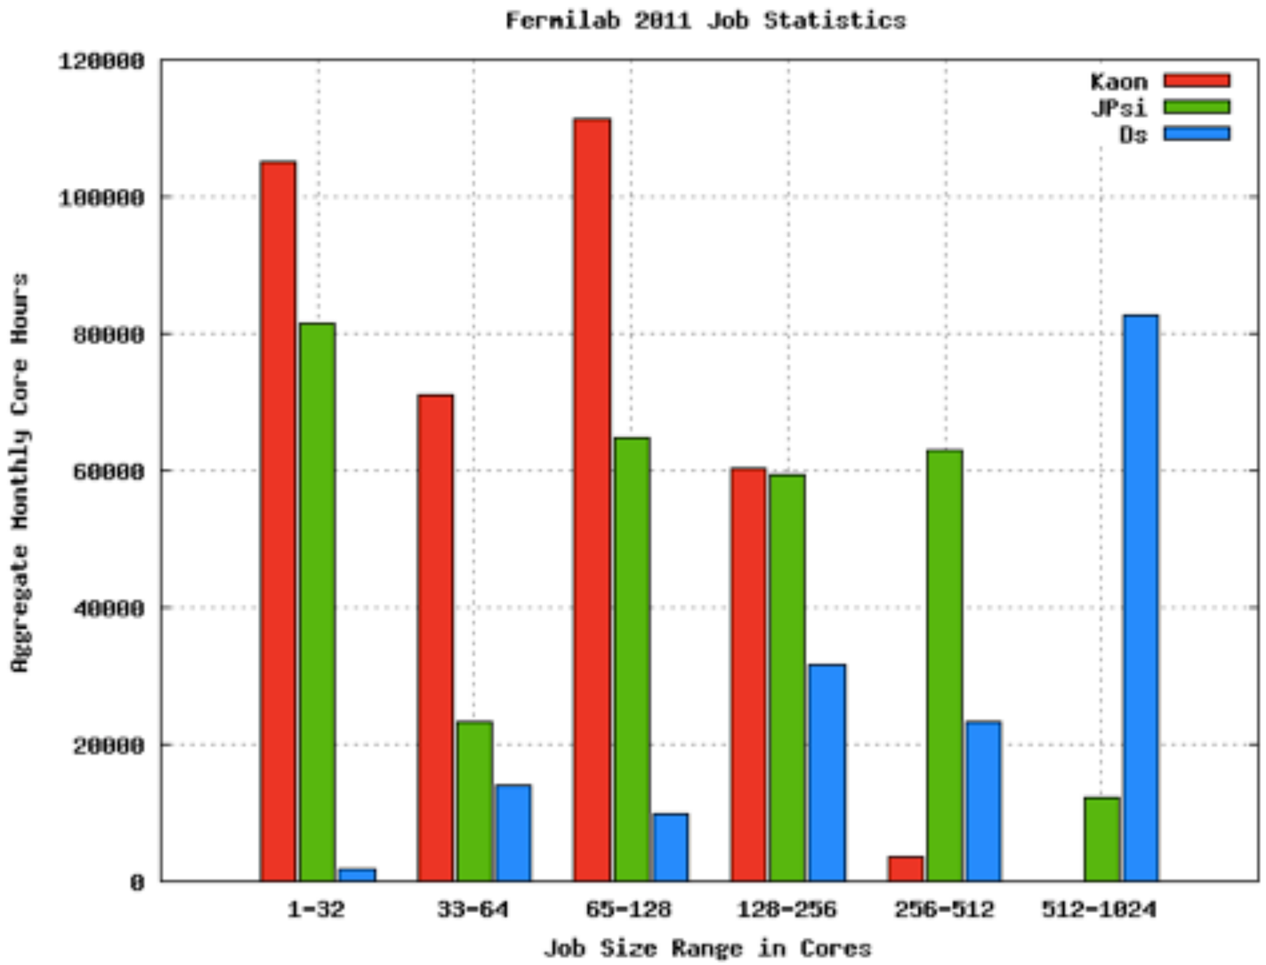

We have an approximately flat distribution of job-size needs from one-node jobs to multi-thousand node jobs on a log scale in job-size.

Fermilab Infiniband clusters.

SciDAC software program

- Organized by the the USQCD Software Committee.
- Essential to our program
  - for using hardware resources efficiently, both our INCITE resources and our LQCD-ext hardware,
  - for integrating new methodological developments,
  - for accomplishing our physics goals.
- Includes community libraries for QCD programming, called the QCD API, optimized high-level QCD codes and software packages, porting to new platforms, work with SciDAC centers and institutes and with computer scientists.

Brower, Clark, Gottlieb, Joo talks.

The QCD API

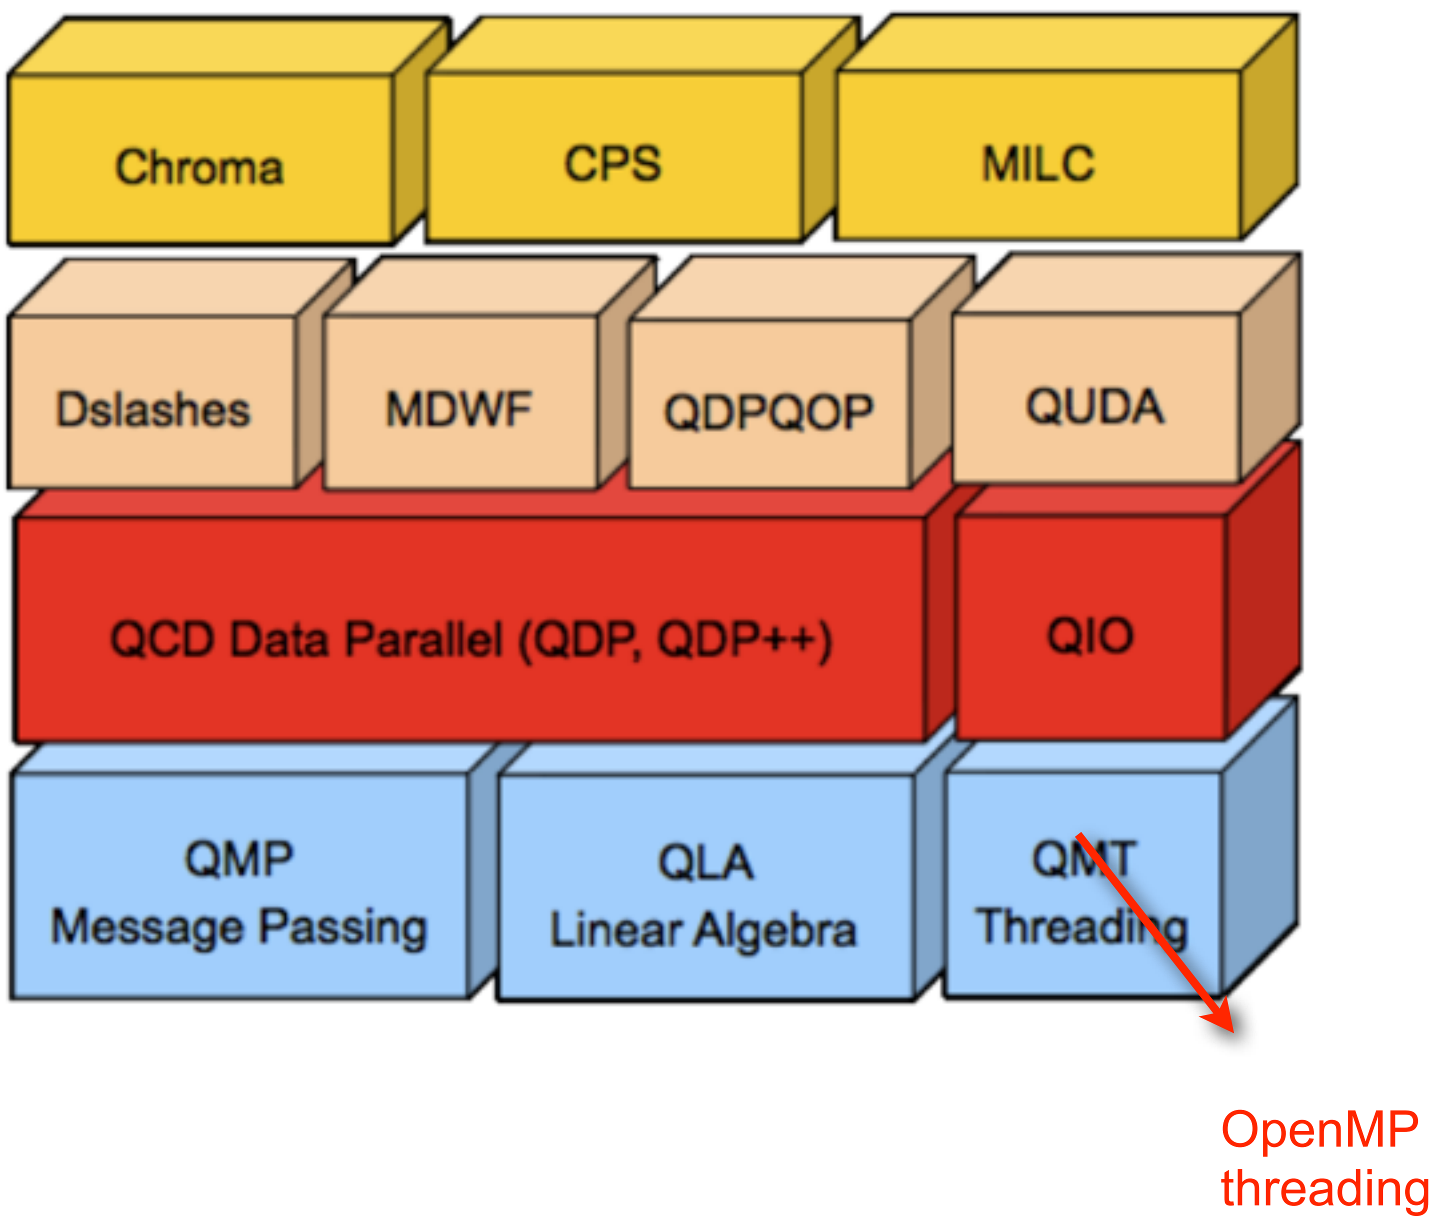

Basics created in SciDAC-1.

The QCD API

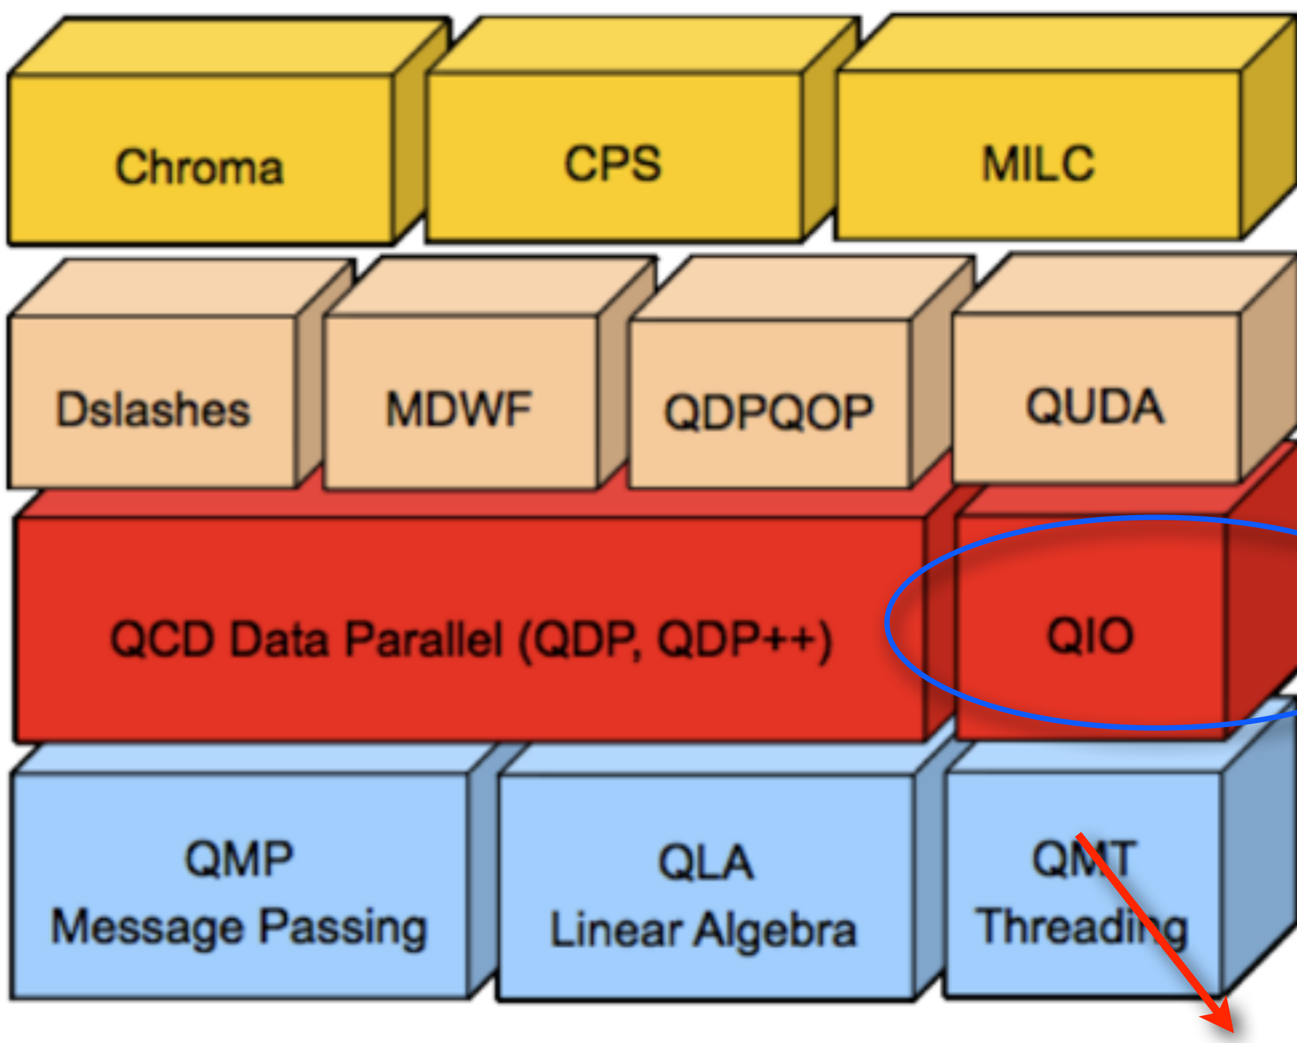

Basics created in SciDAC-1.

Added in SciDAC-2.

OpenMP
threading

The QCD API

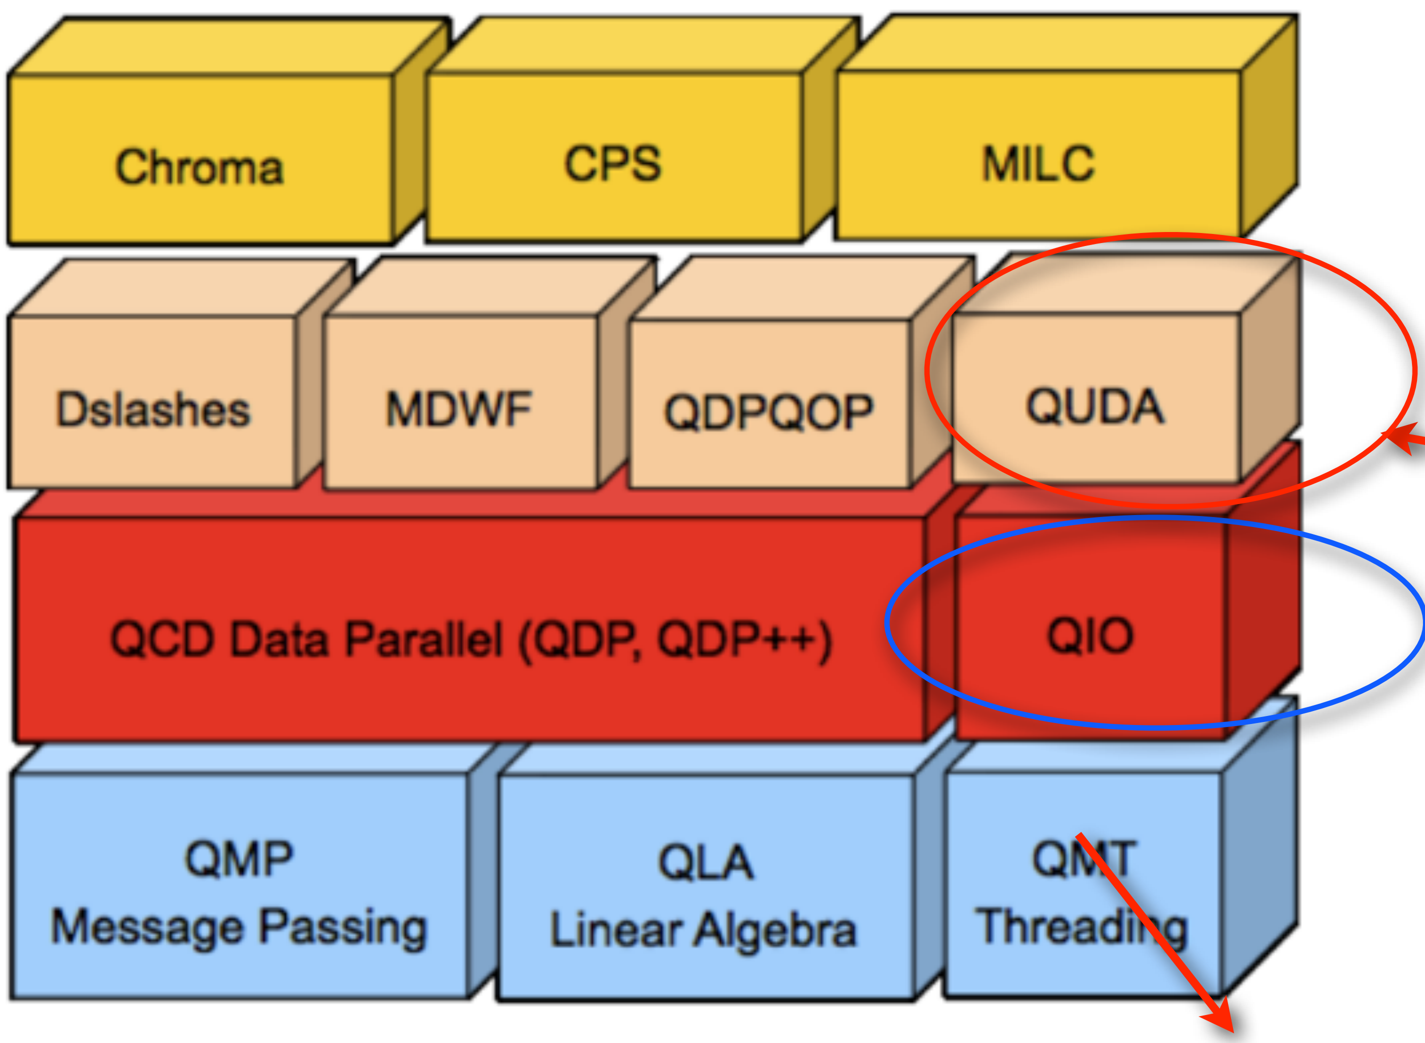

Basics created in SciDAC-1.

Added in SciDAC-2.

Active areas of development now.

OpenMP threading

Community software

E.g., Chroma and QDP++.

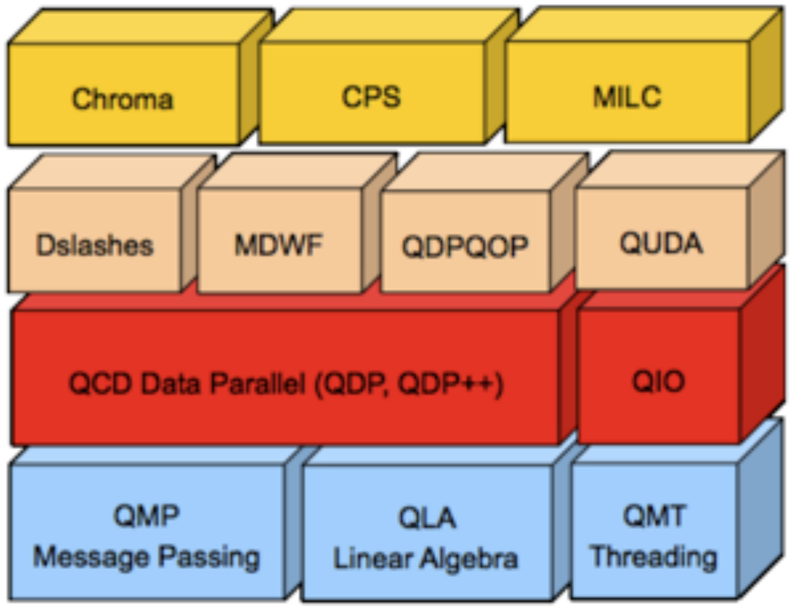

- Chroma was designed from bottom up in the USQCD era along with the QDP++ version of the QCD API.
- Level 2 (QDP++) Data parallel abstraction.
  - Hides architectural implementation and many optimizations.
  - Supports expressions & communications – close to pure math.
  - Eases rapid prototyping. Lowers entrance barrier for newcomers.
  - Use of expression templates in QDP++ hides loops over lattice site and internal space indices. Designed using modern software engineering techniques (design patterns, nightly test builds and regressions).
- Wide variety of highly optimized code available for various platforms.
- Nucl. Phys. Proc. Suppl. 140 (2005) 832, 290 citations.

Balint Joo's talk on an approach to updating QDP++ and Chroma for the many-core, heterogeous era.

The USQCD SciDAC program has enabled us to make optimal use of the hardware resources available.

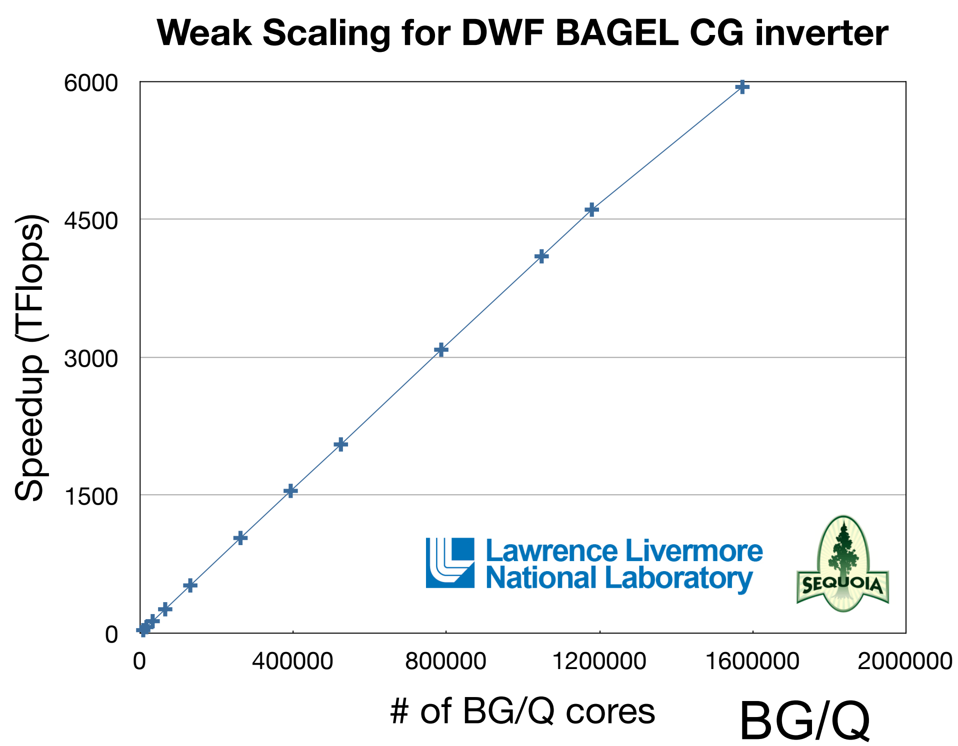

BG/Q

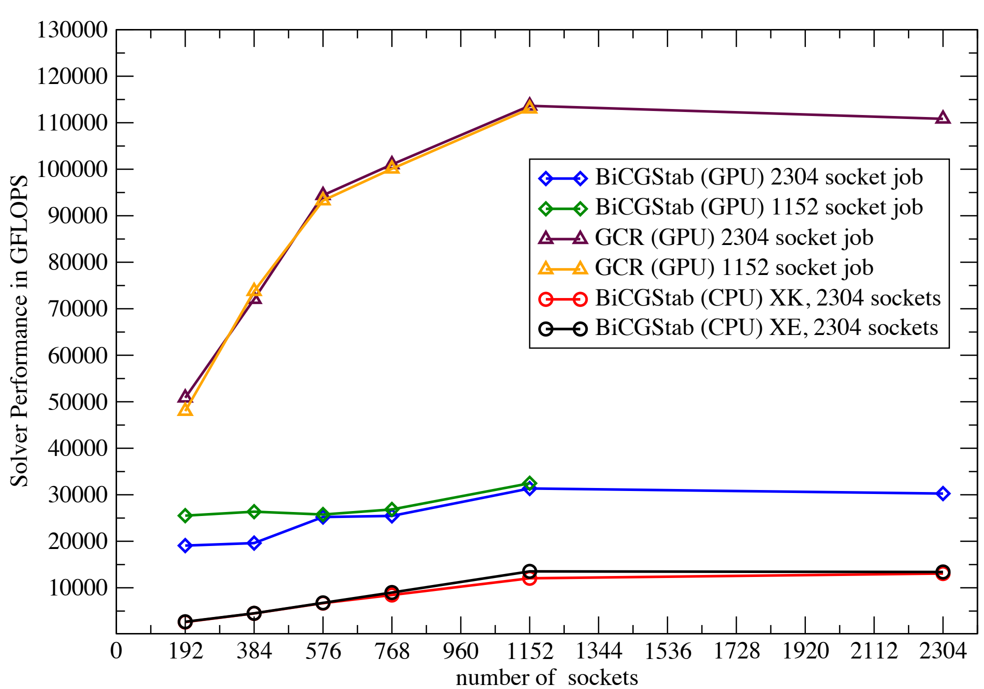

Cray/GPU

In 2008, Chulwoo Jung, James Osborn, and Andrew Pochinsky had highly optimized QCD software for the BG/P ready to go when it became available at the ALCF. Chulwoo's codes were able to identify a hardware error in the machine when it was delivered.

In 2013, high-performance codes have been created for Cray/GPU machines like Titan, the BG/Q, and capacity clusters.

The drive for more power. HEP: precision

High energy and nuclear physics have urgent, immediate needs for even more computing power.

In high energy physics, lattice calculations have been used to place important limits at a few per cent on the effects of beyond-the-standard-model physics on observed particle interactions.

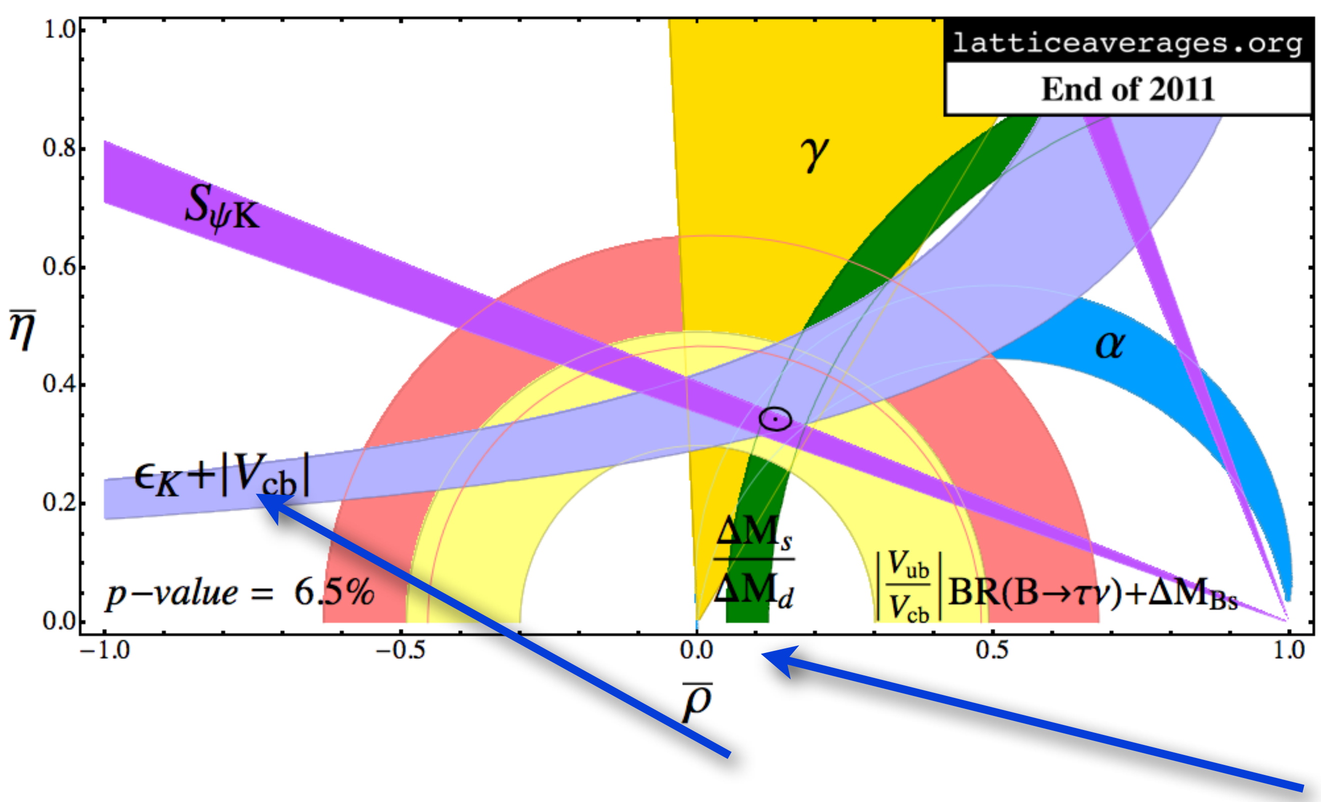

The standard model CP violating parameters rho and eta measured seven different ways.
New physics would lead to inconsistencies.

The experiments for *K*-anti*K* mixing and *B*-anti*B* mixing are known to an order of magnitude better, 0.5%. Commensurately accurate theory calculations are urgently needed.

The drive for more power. NP: volume

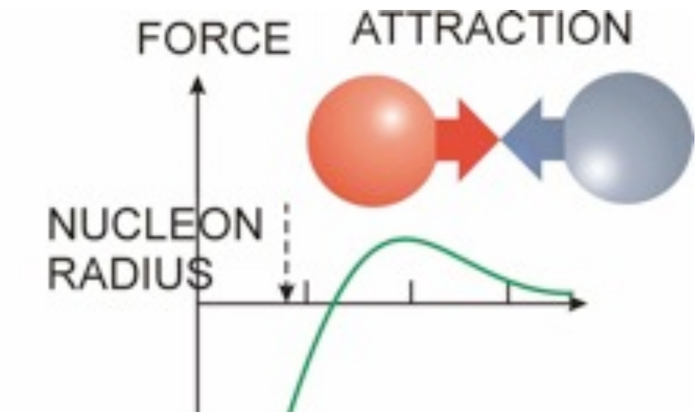

Nucleons and multi-nucleon states require much larger volumes than do the stable mesons which are the bread and butter of HEP calculations.

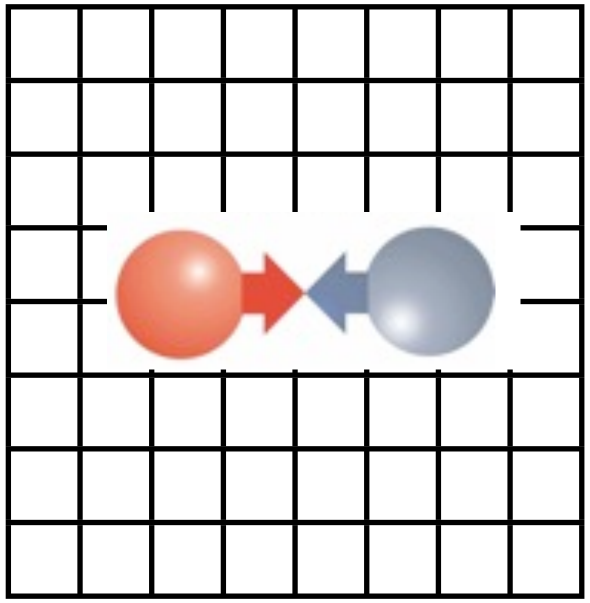

Current computers, even Titan, are not powerful enough to allow calculations at sufficiently large volumes for nuclear physics, while also achieving the physical quark masses and small lattice spacings needed for precise lattice calculations.

New strongly coupled theories may strain or break the uncertainty analysis that is well established for QCD.

Lattice gauge theories start from the fundamental equations. No model uncertainty.

- In QCD, uncertainties arise from
  - statistics
  - finite volume   Known coefficients
  - discretization
  - extrapolation to physical quark mass (in early calculations)

Asymptotic forms are convergent series with known functional forms

Coefficients that can be estimated within ranges from physical arguments.

⇒ Use Bayes' formula

In QCD, this is well understood and solid.

Better algorithms

- In the last 30 years, the development of better methods has contributed even more to lattice calculations than the factor of a billion raw machine speedup from the VAX 11/780 to the Blue Gene/Q.
  - Algorithms for generating gauge configurations have sped up by factors of 10-100 in the past ten years. (Clark and Kennedy, Phys.Rev.Lett. 98 (2007) 051601; hep-lat/0608015 , ...)
- Highest priority now is improving the quark solvers, a sparse-matrix problem.
  - JLab: perambulator methods for multi-propagator calculations.
  - Multi-grid methods, domain decomposition, All Mode Averaging, ...

Bigger machines

Lattice gauge theorists have been involved with the development of supercomputing from the beginning; our ability to program the largest current machines is enhanced by close relationships with vendors.

- Ken Wilson, inventor of lattice gauge theory, was an early proponent of scientific supercomputing.
  - In the 70s, he was programming array processors in assembly language to attack critical phenomena problems for which he won the Nobel Prize.
  - In the 80s, he pushed for establishment of the NSF supercomputing centers.
- After the introduction of Monte Carlo methods to lattice QCD in the early 80s, lattice gauge theorists worked to design machines aimed at lattice QCD
  - in academic efforts at Caltech (Cosmic Cube), Columbia, IBM (GF11, not a commercial project), Fermilab, ...
  - as part of the Thinking Machines project.

Coming hardware challenges in this decade

Memory per chip will grow by 100x.
Each core will communicate quickly only with nearby memory on chip.
~Six levels of (user-controlled?) cache.
Bandwidth per flop will drop sickeningly.

System attributes	2010	"2015"		"2018"	
System peak	2 Peta	200 Petaflop/sec		1 Exaflop/sec	
Power	6 MW	15 MW		20 MW	
System memory	0.3 PB	5 PB		32-64 PB	
Node performance	125 GF	0.5 TF	7 TF	1 TF	10 TF
Node memory BW	25 GB/s	0.1 TB/sec	1 TB/sec	0.4 TB/sec	4 TB/sec
Node concurrency	12	O(100)	O(1,000)	O(1,000)	O(10,000)
System size (nodes)	18,700	50,000	5,000	1,000,000	100,000
Total Node Interconnect BW	1.5 GB/s	20 GB/sec		200 GB/sec	
MTTI	days	O(1day)		O(1 day)	

from Dongarra and Beckman, via Thakur

Fewer cores/node (like BG/Q)

More cores/node (like GPUs, Titan)

Vision of the future from a couple years ago.
Two paths to the future were being thought through, based on current supercomputers.
We are heavily involved with both.

The QCDSP, the QCDOC, and the Blue Gene/L

- The Columbia group, led by Norman Christ, won the Gordon Bell prize for price/performance in 1998 for the QCDSP, a machine purpose-built for lattice QCD.
- It was succeeded by the QCDOC.
- A team led by Al Gara that had been part of these projects went to IBM and designed the closely related (and commercial product!) BG/L, which won the Gordon Bell prize for performance in 2005.
  - The system-on-a-chip design, tightly coupled standard processor and FP unit, torus network, and style of mechanical design (small easily replaced node cards) were modeled on the Columbia machines.

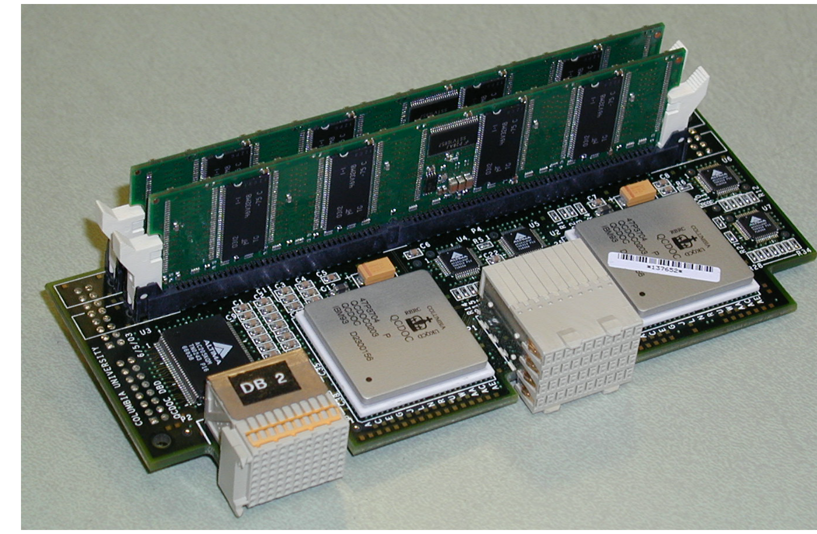

QCDOC compute card.

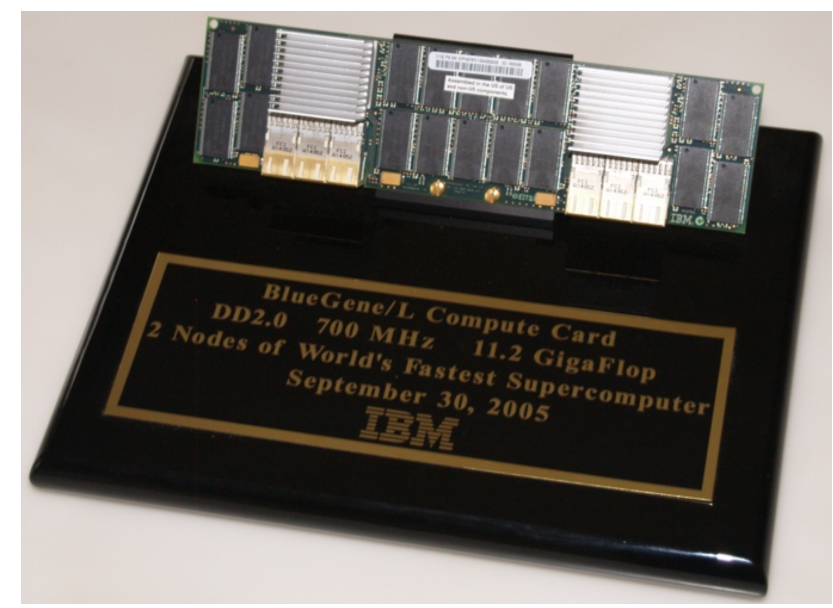

BG/L compute card.

The BG/Q

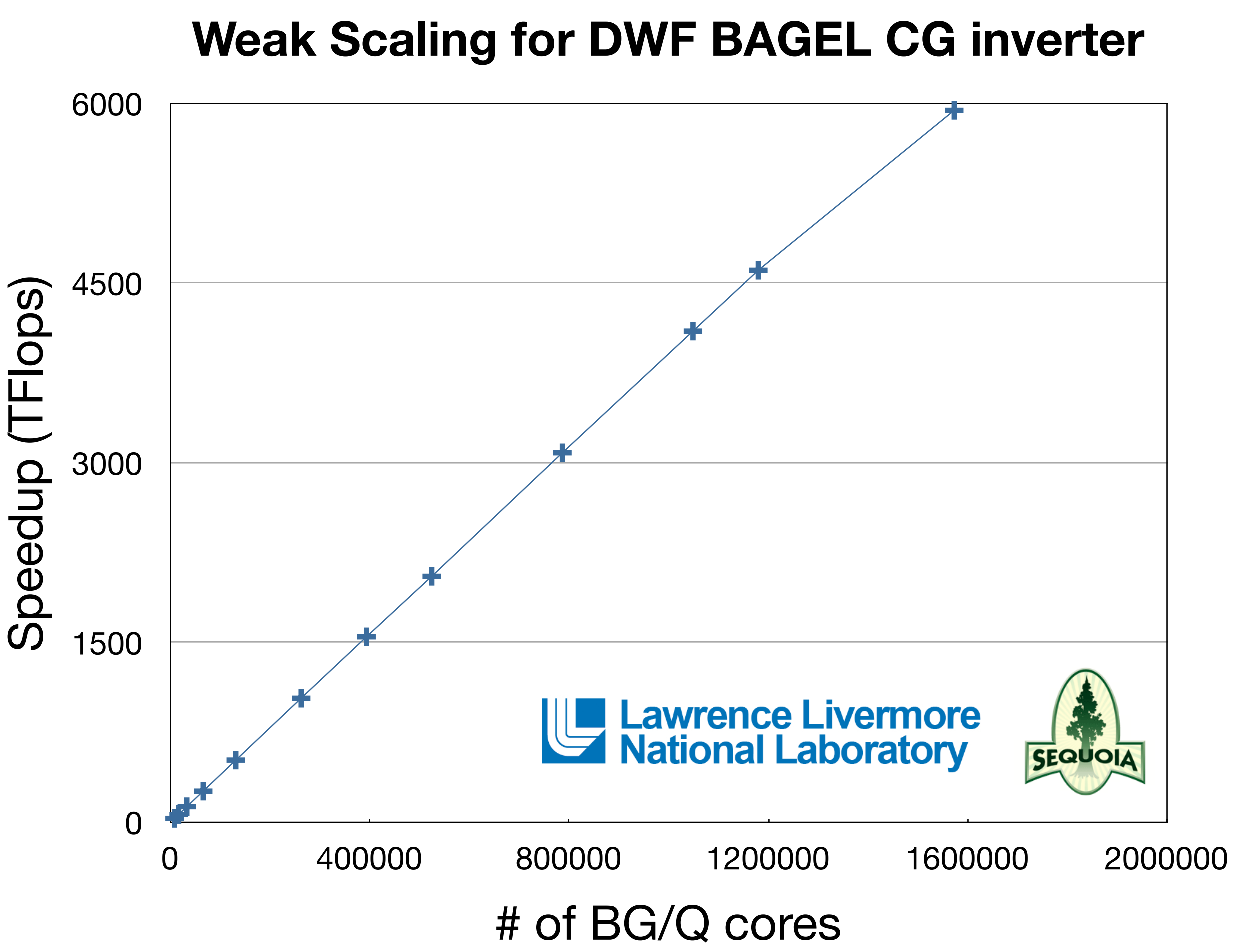

- The Columbia group participated in the design of the BG/Q. Under contract with IBM, they designed and implemented:
  - The interface between the processor core and the level-2 cache, and
  - The look-ahead algorithms used to prefetch data from level-2 cache and main memory, anticipating misses in the level-1 cache.
- Almost perfect weak scaling, 6 PF sustained, achieved on 96 racks of Sequoia (Columbia UK collaborator Peter Boyle).
- The cache management problem on current computers will become nightmarish on exascale computers. (6 levels of cache?)
  - These BG/Q prefetching methods may serve as an approach to the exascale cache problem.

GPUs

- Two USQCD members work directly for NVIDIA.
  - They work with academic collaborators to attain best performance on current machines,
  - Evaluates potential future architectures in terms of QCD (cache sizes, memory bandwidths, network bandwidth, latency sensitivity).
- Up to 10X peak price/performance vs. clusters for parts of code resident in GPU. But ... very low memory/core, bandwidth/core
  - Problem will grow worse on *all* computers throughout this decade. Lessons learned now are important.
- Optimization must reduce data movement, floating point not as important.
  - Easily reconstruct 8 or 12 of 18 SU(3) matrix components. Transfer only half.
  - Calculate desired double precision solution of $Ax=b$ in single or half precision, use double precision residual of result $r=b-Ax$ as a source to "polish" the result to double precision.

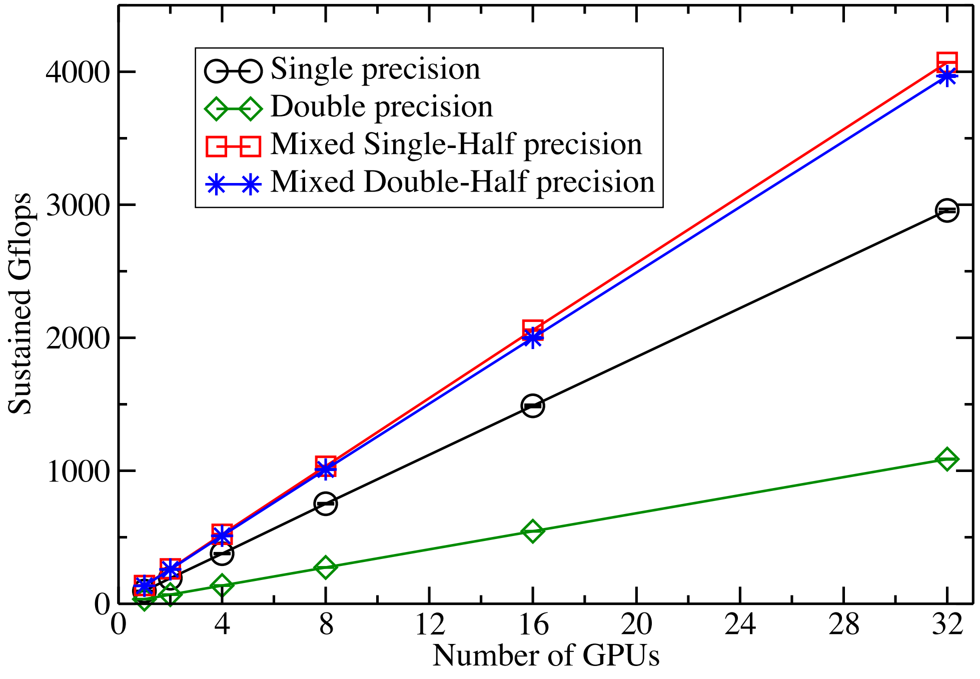

Many GPUs: Titan, ...

- Strong scaling to large numbers of GPUs brings even greater bandwidth challenges.
  - Ameliorated with further communications-minimizing algorithms.
  - E.g., Schwarz domain decomposition.

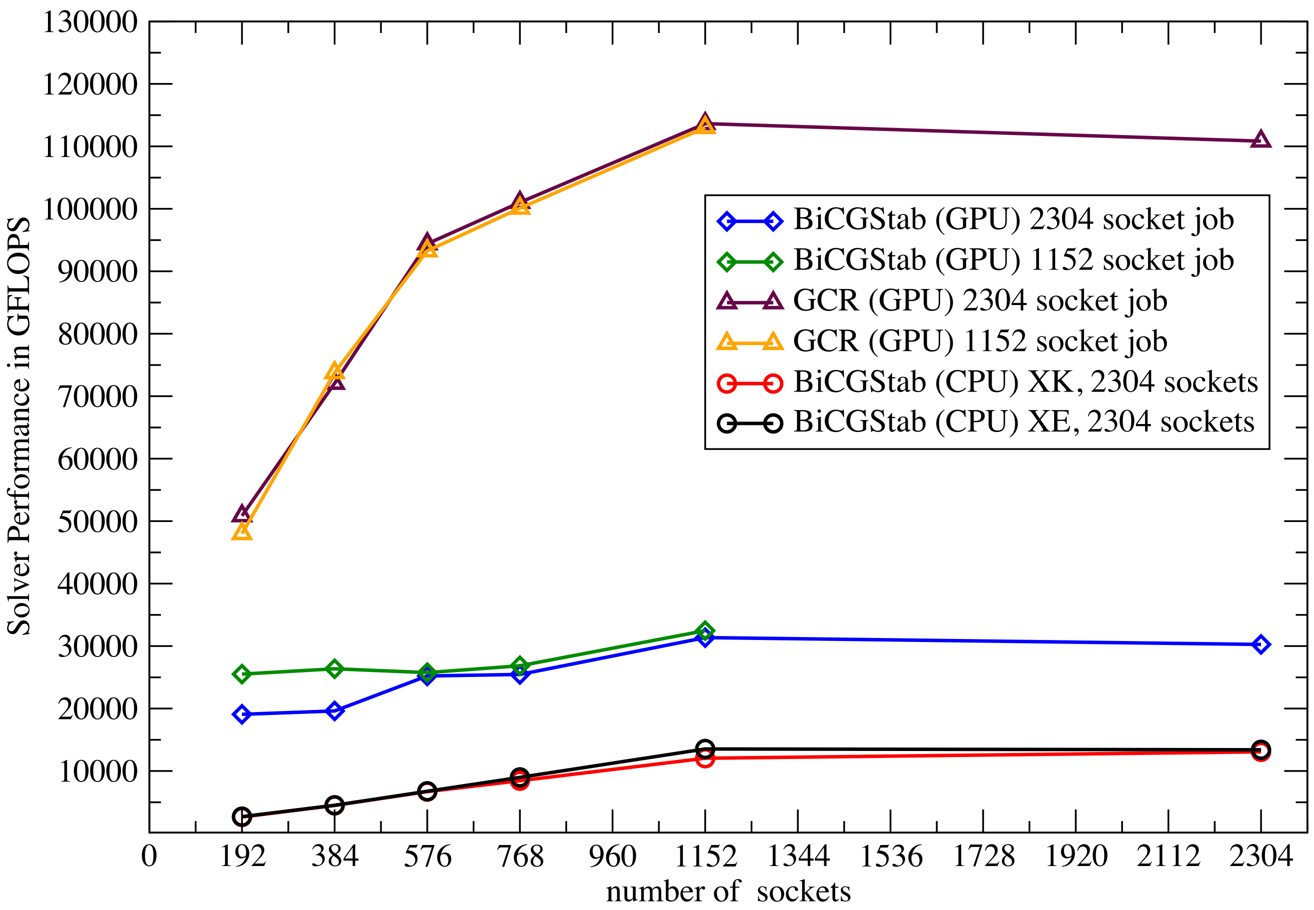

Clark, Gottlieb talks

Summary

- The experimental programs of high energy and nuclear physics have a critical need for numerical simulations of quantum chromodynamics to accomplish their programs.
- Lattice gauge theory calculations have been relentless drivers of scientific computing since the invention of lattice QCD and large-scale scientific computing in the ‘70s and ‘80s.
- Lattice calculations continue to urgently need vast amounts of scientific computing to accomplish their mission and we are pushing as hard as we can in every way we can to get it.

OLD

USQCD timeline

USQCD Executive Committee formed.

Software grants

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

SciDAC extension

SciDAC-3?

Absolutely essential for making effective use of Leadership Computing Facilities and our dedicated hardware, and for accomplishing our physics objectives.

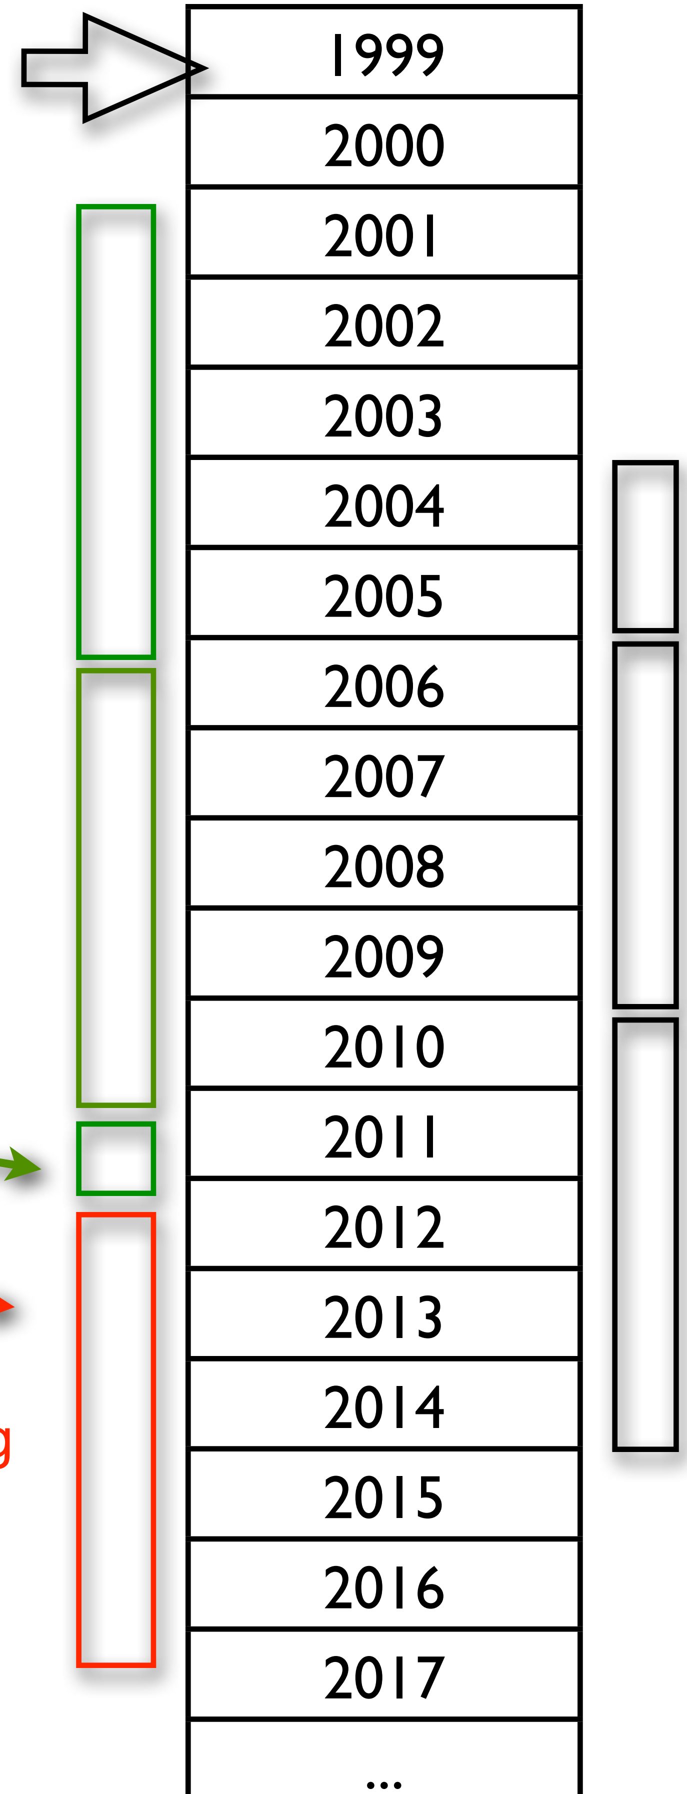

Hardware grants

Construction of the purpose-built QCDOC.

Funding from HEP and NP for hardware through LQCD and LQCD-ext projects.

Current hardware resources

- Last year, used 187 M core-hours at ALCF, 40 M core-hours at OLCF.
  - Expect about the same this year.
- The SPC is allocating on USQCD's dedicated hardware
  - 262.3 M Jpsi-core hours on clusters at JLAB and FNAL. (Jpsi core~2 BG/P cores.)
  - 4.2 M GPU-hours on GPU clusters at JLAB and FNAL.
- (Large resources at NERSC and the Teragrid are also used for lattice QCD, managed by individual member collaborations, not USQCD.)