



Algorithmic Challenges for Lattice Field Theory in the Multi-scale Era



SciDAC-3 PI Meeting
DOE July 24, 2013
Rich Brower Boston University



Narrative*

- Algorithms are the critical multi-disciplinary pursuit
 - Physics + Applied Math + Computer Science → Algorithms designers
- Need “Heroic” programming just to keep up
 - MB/Q, Titan, BlueWaters, Stampede → CUDA, JIT, **SUPER**
- Multi-grid QCD success story: why did it takes so long?
 - MG, DD,, **FastMATH**, Hyper, Qlua, FUEL
- Future: New theories, New lattices, Much more Multi-level
 - **Condensed Matter**, Graphene, Radial Quantization, FEM, etc.

Lattice Field Theory Coming of Age

Kenneth G. Wilson “Confinement of quarks”

Phys. Rev. D 10, 2445–2459 (1974)



K. Wilson: “Lecture at Lattice 1989 Capri”

“lattice gauge theory could also require a 10^8 increase in computer power AND spectacular algorithmic advances before useful interactions with experiment ...

- ab initio Chemistry
- 1. 1930+50 = 1980
- 2. 0.1 flops → 10 Mflops
- 3. Gaussian Basis functions

VS

- ab initio QCD
- 1. 1980 + 50 = 2030?*
- 2. 10 Mflops → 1000 Tflops
- 3. Clever Collective Variable?

Sustained Petaflops: 15 years ahead of schedule!

Expanding Physics Goals of Lattice Gauge Theory

- Ab initio QCD for **Nuclear/Astrophysics**

- I. Structure and excitation of Nuclei (QCD)

- II. Quark Gluon Plasma (QCD)

- Beyond the Standard Model for **High Energy Physics** :

- III. QCD for precision tests of Standard Model
(aka "Intensity Frontier")

- IV. Explore new Theories (not QCD) at TeV
(aka "Energy Frontier")

- Strongly coupled quantum systems for Novel Materials

Jefferson Lab

RHIC



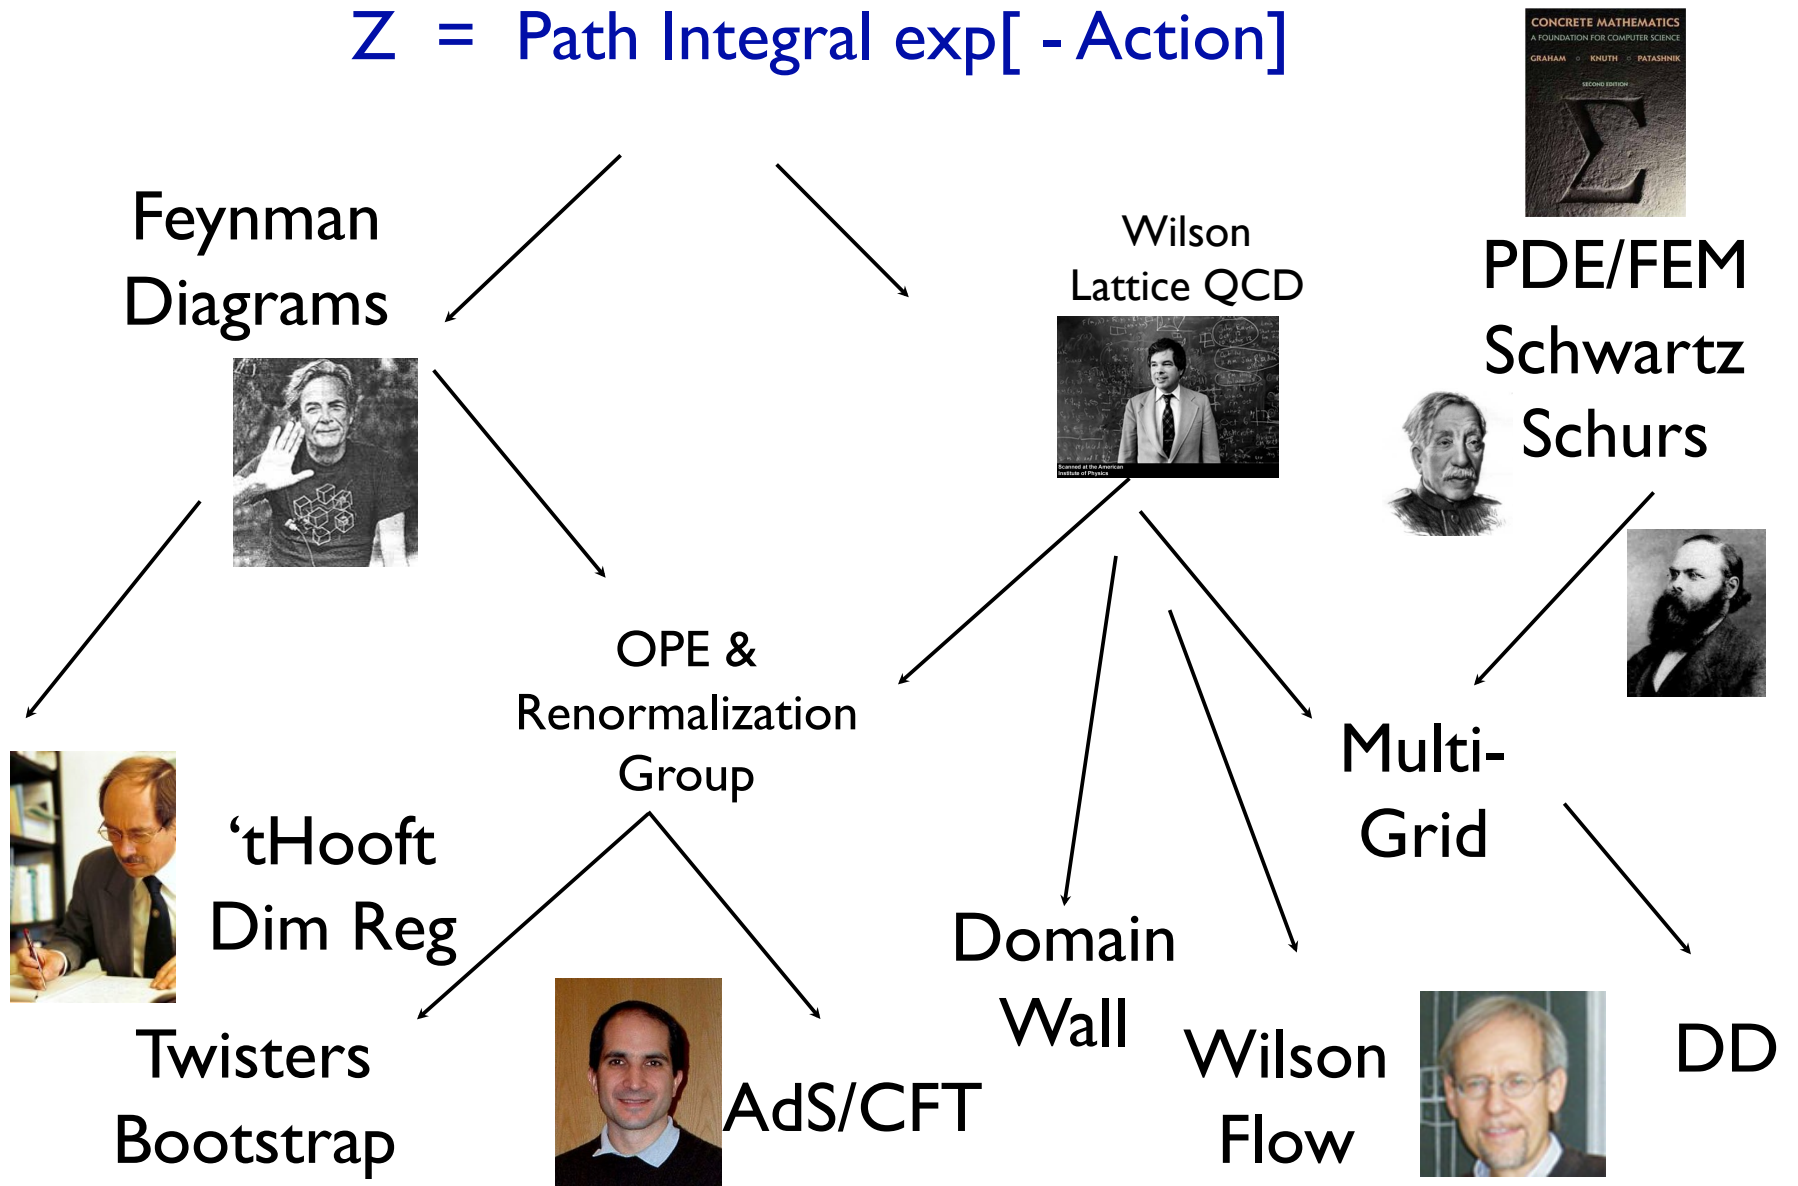
ATLAS
EXPERIMENT

Fermilab

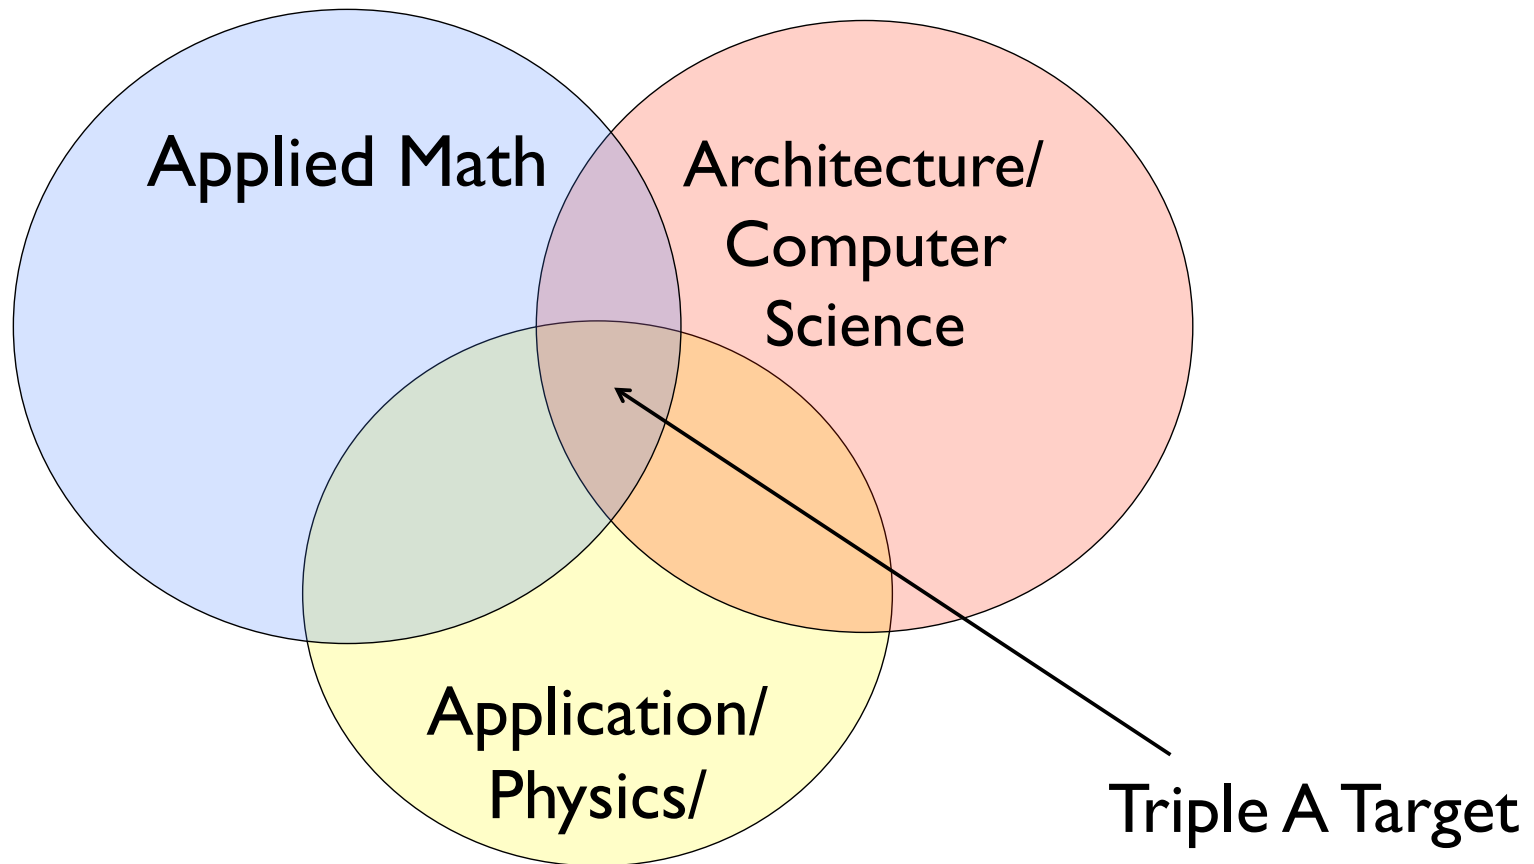


All prediction from Quantum Field Theory require "Algorithms"

$$Z = \text{Path Integral exp}[-\text{Action}]$$



Intersection of Algorithm Design



CHALLENGE: Need to **collaborate** between 3 disciplines
BUT also need **suitable abstraction/interfaces** so
each discipline can proceed semi-autonomously!

See posters for Details

BG-Q/IBM



Lattice QCD on the BGQ: Achieving 1 Pflops Production Jobs

Through advances in computer hardware and software, production lattice QCD jobs run by the BGQ and LUGQ Collaborations on BGQ simulations are reaching 1 PFlops. The largest computational costs in the solution of the Dirac equation for the fermion calculation is inverse preconditioning via Conjugate Gradient (CG) based solver for the fermion calculation. In reverse preconditioning via Conjugate Gradient (CG) based solver for the fermion calculation, the largest computational cost is the solution of the Dirac equation for the fermion calculation. In reverse preconditioning via Conjugate Gradient (CG) based solver for the fermion calculation, the largest computational cost is the solution of the Dirac equation for the fermion calculation.



Quantum Chromodynamics

Review Elementary Particles

QCD + Electroweak Interactions Produce Fermion Physics

Nonlinear Path-Block Reducing the Fermionic Effects of Quarks

Computers, Algorithms and Software

Production of LQCD software and lattice QCD computations to be performed with high performance across a variety of architectures, including embedded hardware, custom computing systems, and general purpose CPUs. The software must be able to execute on various hardware and be able to scale to high-end supercomputers.

Challenges

- Large data volume
- High precision
- High performance
- High accuracy
- High reliability
- High maintainability

Algorithm for Gauge Field Production

Algorithm for Measurement

Physics Results and Prospects

Major Developments Enabled by Physical Quark Masses

High Level Interface

```

            #include <HYPRE.h>
            int main() {
                // LQCD initialization
                // ...
            }
        
```

MG/fastMATH

Multigrid with HYPRE for Lattice QCD

Andrew Pochinsky
Massachusetts Institute of Technology
ap@MIT.EDU



Abstract

Lattice QCD simulations are a significant portion of computational time spent by scientific computing applications. The largest computational costs in the solution of the Dirac equation for the fermion calculation is inverse preconditioning via Conjugate Gradient (CG) based solver for the fermion calculation. In reverse preconditioning via Conjugate Gradient (CG) based solver for the fermion calculation, the largest computational cost is the solution of the Dirac equation for the fermion calculation.

Lattice QCD

Lattice QCD is one of the most demanding computational science fields. To continue to advance the field, lattice QCD simulations must move to larger and larger supercomputers. The computational cost of lattice QCD simulations is high and is expected to increase significantly in the future. The computational cost of lattice QCD simulations is high and is expected to increase significantly in the future.

Multigrid

Multigrid methods form a group of algorithms for solving systems of linear equations using discretization of the vector space into a hierarchy of subspaces. The most of multigrid is to reduce the convergence of a linear system to the solution of the original problem. This process is a linear transformation between the coarse and fine grids. The multigrid method is a linear transformation between the coarse and fine grids.

LQCD Software

HYPRE

The FAST Math Supercomputing Initiative (FAST) is a software library for high performance multigrid and iterative solvers for the solution of large linear systems of equations. The primary goal of the FAST Math is to provide users with an efficient and scalable multigrid and iterative solvers for the solution of large linear systems of equations. The primary goal of the FAST Math is to provide users with an efficient and scalable multigrid and iterative solvers for the solution of large linear systems of equations.

Linear System Interfaces

```

            #include <HYPRE.h>
            int main() {
                // HYPRE initialization
                // ...
            }
        
```

Primitive Solvers

A QUDA multigrid for solvers and preconditioners high level description is provided. The multigrid and preconditioners high level description is provided. The multigrid and preconditioners high level description is provided.

Solver Interface

A QUDA multigrid for solvers and preconditioners high level description is provided. The multigrid and preconditioners high level description is provided.

GPU/NVIDIA

QUADA: QCD in CUDA for Multi-GPU Lattice Field Theory

USQCD US Lattice Quantum Chromodynamics

USQCD US Lattice Quantum Chromodynamics

USQCD US Lattice Quantum Chromodynamics

Stage 1: Basic Linear Solvers

Maximize performance

Reduce on Card Memory Traffic

Stage 2: Scaling to Multi-GPUs

Communication Reduction between GPUs

Stage 3: Multi-scale Physics

Physics Aware Algorithm: Adaptive Multigrid

Future: Synthesis DD = MG to satisfy Arch + Physics

Titan/Intel-Phi/SUPER

Porting Lattice QCD Calculations to Novel Architectures

Ballint Joo, Frank Winter, Jefferson Lab, for the USQCD Collaboration

Computing Properties of Hadrons, Nuclei and Nuclear Matter from Quantum Chromodynamics

QDP-JIT and QUADA: Enabling Chroma on GPU Based Leadership Architectures

The Chroma software system is the standard workhorse of the gauge generation phase of LQCD calculations in Grid Nuclear Physics on leadership class systems such as QLCF/Titan, which feature GPU accelerated nodes. Chroma has long enjoyed accelerated solvers utilizing libraries such as QUDA. Gauge generation, however, requires the whole application to be accelerated, to avoid *Amal's Law* effects in parts of the code outside of QUDA.

Optimizing for Xeon Phi

We have been working in close collaboration with Intel Parallel Labs to develop high performance implementations of Lattice QCD kernels for Intel Xeon Phi. To achieve good performance requires careful attention to vectorization, cache-blocking, and mapping threads to blocks in a load balanced way. We wrote a code-generator to abstract vector instructions and to allow us to vary the spacing of prefetch instructions. The multi-node code uses an MPI Proxy to pick the optimal path between device nodes. The framework has been re-targeted to AFX on Intel Sandy Bridge and should be straightforward to apply to other multi-core, cache and vector based architectures such as BGQ.

Conclusions

QDP-JIT will allow effective exploitation of accelerated hardware resources, and forms the basis of our work in partnership with the SUPER SciDAC Institute to create a Domain Specific Compilation Framework for Lattice QCD. The lessons learned can be applied to other domain specific Frameworks using expression templates. Our work with Xeon Phi seeks to discover approaches for highly performant code on this architecture, targeting large scale Xeon Phi resources (e.g. Stampede) and to quantify the "Ning Gap" between optimized and "regular" code to help design future computers. Finally, the micro-benchmarks developed through our collaboration with Intel can be used to stress and evaluate proposed architectural changes on future version of Xeon Phi.

Lattice Gauge Theory HEP and NP posters

1. Lattice QCD on the BGQ: Achieving 1 PFlops Production Jobs

- IBM/Columbia (Bob Mawhinney)

2. Multigrid with Hyper for Lattice QCD

- FastMath/MIT (Rob Faulgot, Andrew Pochinsky)

3. Porting Lattice QCD Calculation to Novel Architectures

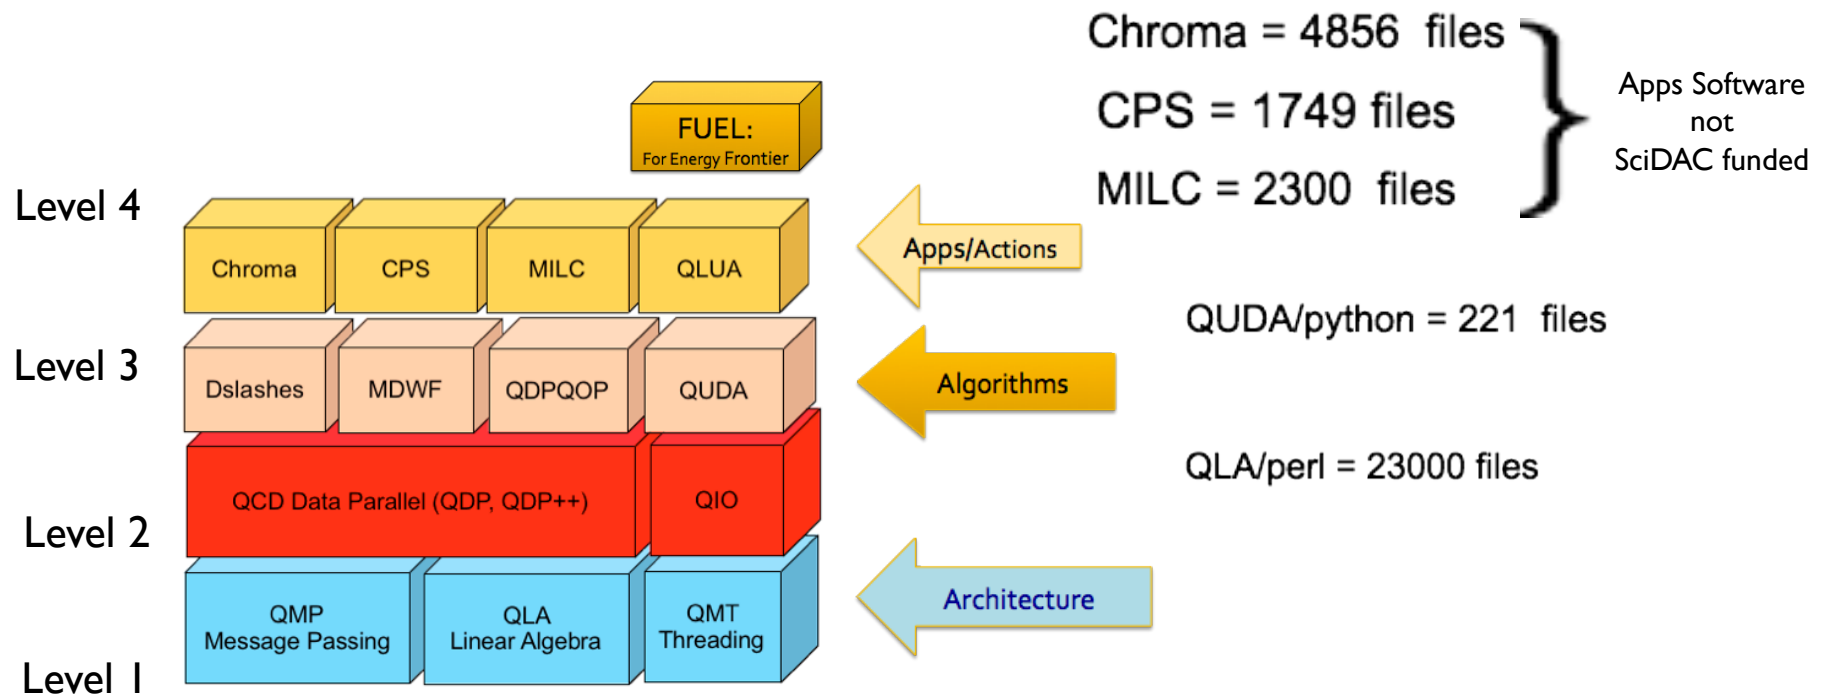
- SUPER/Jlab (Rob Fower/Balint Joo)

4. QUDA: QCD in CUDA for Multi-GPU Lattice Field Theory

- NVIDIA/BU (Mike Clark/Rich Brower)

USQCD Software Stack Stack

On line distribution: <http://usqcd.jlab.org/usqcd-software/>



The application codes
Chroma/CPS/MILC and a new QDP LUA
code base provide a rich set of tools.

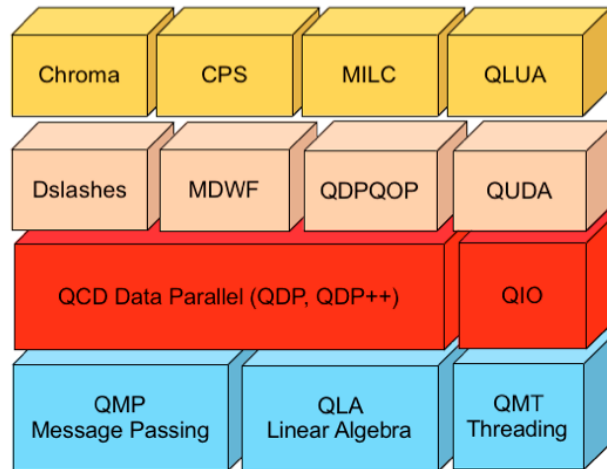

SciDAC LGT contributors

- ANL: James Osborn, Meifeng Lin, Heechang Na, (George T. Fleming)
- BNL: Frithjof Karsch, Chulwoo Jung, Hyung-Jin Kim, Yu Maezawa
- Columbia: Robert Mawhinney, Hantao Yin
- FNAL: James Simone, Alexei Strelchenko, Don Holmgren, Paul Mackenzie
- JLab: Robert Edwards, Balint Joo, Jie Chen, Frank Winter, Chip Watson
- W&M/UNC: Kostas Orginos, Andreas Stathopoulos, Rob Fowler (SUPER)
- LLNL: Pavlos Vranas, Chris Schroeder, Rob Faulgot (FASTmath)
- NVIDIA: Ron Babich, Mike Clark
- Arizona: Doug Toussaint, Alexei Bazavov
- Indiana/NCSA: Steve Gottlieb, Ran Zhou
- Utah: Carleton DeTar, Justin Foley
- BU: Richard Brower, Michael Cheng, Oliver Witzel
- MIT: Pochinsky Andrew, John Negele,
- Syracuse: Simon Catterall, (David Schaich in fall)
- Washington: Martin Savage, Saul Cohen
- Others: Peter Boyle, Jim Hetrick, Massimo Di Pierro, Patrick Dreher, et al
- “Team of Rivals” (apologies to contributors and projects *NOT* mentioned in 6 slides!)

Highest Priority is moving to 3 new architecture!



(May you live in Interesting Times!)



See posters for more details

BG-Q/IBM

Lattice QCD on the BGQ: Achieving 1 PFlops Production Jobs



Through advances in computer hardware and software, production lattice QCD jobs run by the RBC and UKQCD Collaborations on BGQ installations are reaching 1 PFlop. The largest computational costs in the solution of the Dirac equation for lattice QCD calculations are extensive multi-threading via OpenMP in the rest of our code base (lattice QCD) and Hines, Yu, Columbia has also been vital to achieving good performance. In addition, we employ a number of new Jefferson and Carnegie Mellon technologies to speed up our code further beyond the software improvements. With all of these techniques, we are able to simulate lattice QCD with physical pion masses in large volumes of size (5 fm)³. This has led to markedly reduced statistical and systematic errors in our results. With these advances as well as new theoretical ideas, we are now beginning production calculations for kaon decay involving disconnected quark diagrams, where the signal to noise ratios are much lower. These new calculations require constant on-demand model updates.



Quantum Chromodynamics

Quantum Chromodynamics (QCD) is the theory of the strong interaction, one of the four fundamental forces of nature. It describes the interactions between quarks and gluons, the constituents of matter and the force carriers, respectively. The theory is based on the principle of local gauge invariance under the SU(3) color group.

QCD is a non-Abelian gauge theory, meaning that the gluons themselves carry color charge and interact with each other. This leads to the phenomenon of asymptotic freedom, where the coupling constant becomes small at high energies, allowing for perturbative calculations. At low energies, the coupling becomes strong, and the theory is non-perturbative, requiring numerical methods like lattice QCD for study.

QCD is a non-Abelian gauge theory, meaning that the gluons themselves carry color charge and interact with each other. This leads to the phenomenon of asymptotic freedom, where the coupling constant becomes small at high energies, allowing for perturbative calculations. At low energies, the coupling becomes strong, and the theory is non-perturbative, requiring numerical methods like lattice QCD for study.

Computers, Algorithms and Software

The BG/Q supercomputer is a leading-edge system for high-performance computing. It features a unique architecture with a large number of processing elements (PEs) and a highly optimized interconnect. The system is designed for large-scale, parallel applications, making it ideal for lattice QCD simulations.

Our software stack is tailored for the BG/Q architecture, utilizing OpenMP for multi-threading and custom kernels for the Dirac equation. We have also implemented advanced memory management techniques to optimize the use of the system's large memory capacity.

Our software stack is tailored for the BG/Q architecture, utilizing OpenMP for multi-threading and custom kernels for the Dirac equation. We have also implemented advanced memory management techniques to optimize the use of the system's large memory capacity.

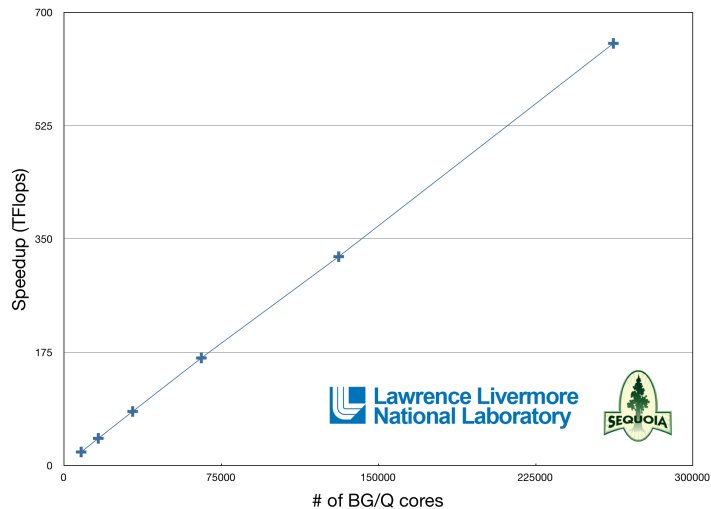
Physics Results and Prospects

Our recent results include the calculation of the pion and kaon masses, as well as the determination of the strong coupling constant. We have also performed calculations for the decay constants of the pion and kaon, which are crucial for understanding the properties of these mesons.

Our results are in excellent agreement with experimental data and other lattice QCD calculations. This demonstrates the power of the BG/Q supercomputer and our software stack in performing large-scale lattice QCD simulations.

Our results are in excellent agreement with experimental data and other lattice QCD calculations. This demonstrates the power of the BG/Q supercomputer and our software stack in performing large-scale lattice QCD simulations.

Weak Scaling for BAGEL DWF CG Inverter



Lawrence Livermore National Laboratory



Titan/Super/Intel-Phi

Porting Lattice QCD Calculations to Novel Architectures

Ballin Joo, Frank Winter, Jefferson Lab, for the USQCD Collaboration
Computing Properties of Hadrons, Nuclei and Nuclear Matter from Quantum Chromodynamics
<https://arxiv.org/abs/1508.04848>

QDP-JIT and QUDA: Enabling Chroma on GPU Based Leadership Architectures

The Chroma software system is the standard workhorse of the gauge generation phase of LQCD calculations in Cold Nuclear Physics on leadership class systems such as OLCF Titan, which feature GPU accelerated nodes. Chroma has long enjoyed accelerated solving utilizing libraries such as QUDA. Gauge generation, however, requires the whole application to be accelerated, to avoid *Amundin's Law effects* in parts of the code *outside* of QUDA. QDP-JIT is an implementation of the QDP++ layer on which Chroma is built. QDP++ expression templates are compiled into code generators which generate CUDA-PTX kernels for the expressions when first run, at which point grid and block dimensions are autotuned. QDP-JIT features a memory manager which can page data between host and device memories automatically and rearrange data to the most optimal layout (e.g. for coalesced access) without the need to instrument the large Chroma code with *#pragma* annotations.

Optimizing for Xeon Phi

We have been working in close collaboration with Intel Parallel Labs to develop high performance implementations of Lattice QCD kernels for Intel Xeon Phi. To achieve good performance requires careful attention to vectorization, cache-blocking, and mapping threads to blocks in a load balanced way. We wrote a code-generator to abstract vector intrinsics and to allow us to vary the spacing of prefetch instructions. The multiversion code uses an *AVX* on Intel Sandy Bridge and should be straightforward to apply to other multicores, cache and vector based architectures such as BGQ.

Wilson Dash Operator (single prec.)

Preconditioned CG (single prec.)

Wilson Dash Operator (double prec.)

Preconditioned CG (double prec.)

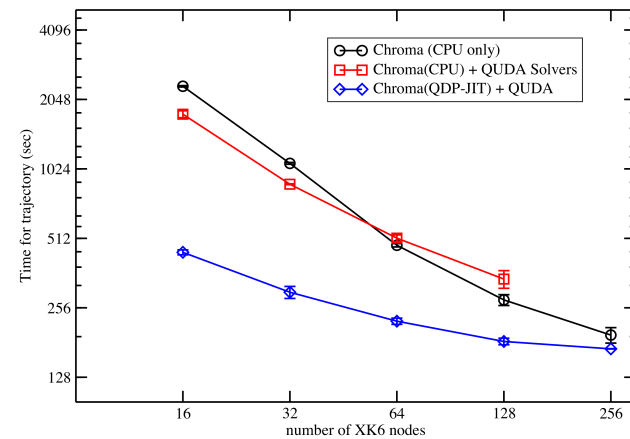
Performance of LQCD Kernels on Single Xeon Phi nodes compared with performance of QUDA on Titan K80 GPUs. The GPU performance are also available at <https://arxiv.org/abs/1508.04848>.

Performance of LQCD Kernels on Multiple Xeon Phi nodes of the Intel Eschelon Cluster using an MPI-communicative Proxy (F. Joo, D. Kabanov, K. Vaidyanathan, B. Sanyal, S. J. Kim, J. C. Franks, 2015, LQCD 1508.04848, 2015).

Conclusions

QDP-JIT will allow effective exploitation of accelerated leadership resources, and forms the basis of our work in partnership with the SUPER SciDAC Institute to create a Domain Specific Compilation Framework for lattice QCD. The lessons learned can be applied to other domain specific frameworks using expression templates. Our work with Xeon Phi seeks to discover approaches for highly performance code on this architecture, targeting large scale Xeon Phi resources (e.g. Stampede) to quantify the "Nirga Gap" between optimized and "regular" code to help design future compilers. Finally, the micro-benchmarks developed through our collaboration with Intel can be used to stress and evaluate proposed architectural changes on future version of Xeon Phi.

2 Flavor Wilson HMC (Gauge + 2 Flavor + Hasenbusch monomials), 32³x96 lattice



QUDA: (QCD in CUDA) Library

QUDA (QCD in CUDA) library started in 2008 with NVIDIA's CUDA implementation by Kip Barros and Mike Clark at Boston University. It has expanded to a broad base of USQCD SciDAC [1] software developers and is in wide use as the GPU backend for HEP and NP SciDAC application codes: Chroma, CPS, MILC, etc.

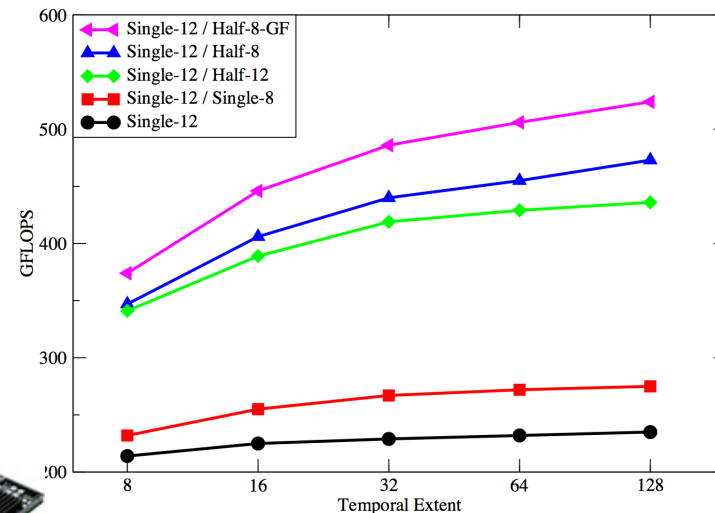
Provides:

- Various solvers for several discretizations,
- including multi-GPU support and domain-decomposed (Schwarz) preconditioners
 - Additional performance-critical routines needed for gauge-field generation

Maximize performance:

- Exploit physical symmetries
- Mixed-precision methods
- Autotuning for high performance on all CUDA-capable architectures
- Cache blocking

Kepler Wilson-Solver Performance



Kepler 20X with
2688 CUDA cores!

“QCD on CUDA” team – <http://lattice.github.com/quda>

- Ron Babich (NVIDIA)
- Kip Barros (LANL)
- Rich Brower (Boston University)
- Michael Cheng (Boston University)
- Mike Clark (NVIDIA)
- Justin Foley (University of Utah)
- Joel Giedt (Rensselaer Polytechnic)
- Steve Gottlieb (Indiana University)
- Bálint Joó (Jlab)
- Claudio Rebbi (Boston University)
- Guochun Shi (NCSA -> Google)
- Alexei Strelchenko (FNAL)
- Hyung-Jin Kim (BNL)
- Frank Winter (UoE -> Jlab)

The screenshot shows the GitHub interface for the `lattice/quda` repository. The top navigation bar includes a search bar, 'Explore', 'Gist', 'Blog', and 'Help'. The repository name 'lattice/quda' is displayed, along with 'Pull Request', 'Unwatch', 'Unstar', 'Fork', and '13' forks. The 'Issues' tab is selected, showing 42 open issues and 73 closed issues. The issues are sorted by 'Newest'. The left sidebar shows filters for 'Everyone's Issues' (42), 'Assigned to you' (10), 'Created by you' (26), and 'Mentioning you' (0). A 'Labels' section lists various categories: bug (4), clean-up (7), feature (19), optimization (14), and question (1). The main content area displays a list of issues, each with a title, labels, and the user who opened it. The issues listed are:

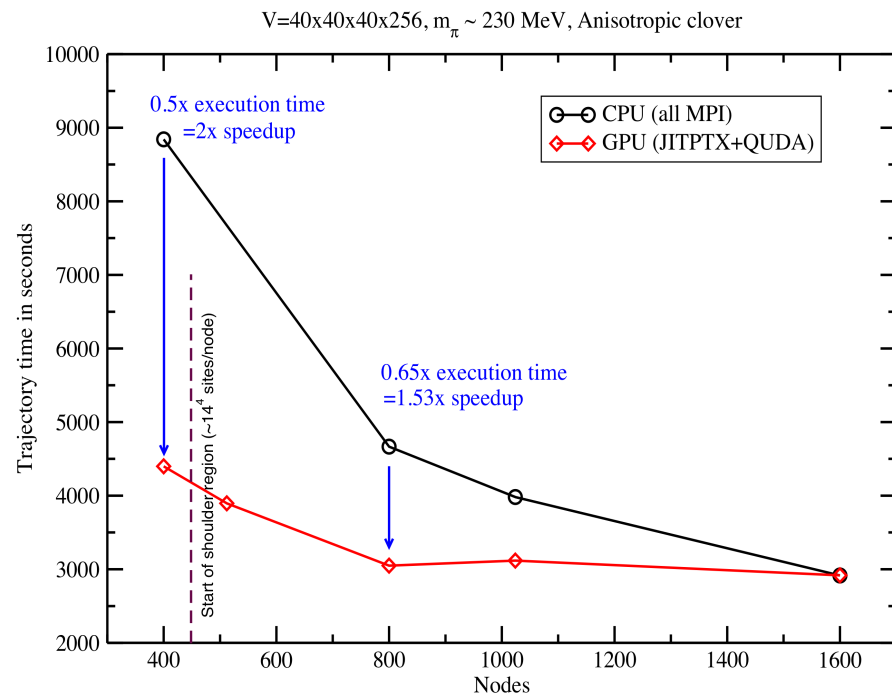
- #114: Investigate using only high precision for the solution vector in CG (feature, optimization)
- #113: Optimize multi-shift CG solver (optimization)
- #112: Implement I-BICGstab solver (feature, optimization)
- #111: Generalise QUDA's profiling utilities (feature, optimization)
- #107: Add support for loading / saving of spinor fields (feature)
- #105: Implement one-sided communication MPI back end (optimization)
- #104: Twisted mass CG solver has bad performance
- #103: Register optimization for each dslash kernel (optimization)

Gauge Generation using QDP-JIT/C on Titan*

“QDP-JIT .. implementation of the QDP++ layer on Chroma ...

QDP++ expression templates are compiled into code generators which generate **CUDA-PTX**

Grid and block dimensions are **autotuned**. QDP-JIT features a memory manager which can page data between host and device memories **automatically** and rearrange data to the most optimal layout (e.g. for coalesced access) without the need to instrument the large Chroma code with `#pragma annotations`” *



Gauge generation strong scaling on Titan: still worse then for CPU probably due to PCIe


*Quoted from Poster on “Porting Lattice QCD Calculations to Novel Architecture by Balint Joo and Frank Winter

Intel Xeon-Phi (stampede)



Xeon Phi 5110P

- 60 cores @ 1.053 GHz
 - connected by ring
 - 512Kb L2\$ / core
 - 32KB L1I\$ and 32KB L1D\$
 - in-order cores, 4 way SIMT
 - 512 bit wide Vector Engine
 - 16 way SP/8 way DP
 - can do multiply-add
- Peak DP Flops: 1.0108 TF
- Peak SP Flops: 2.0216 TF
- 8 GB GDDR (ECC)
 - 'top' shows ~6GB free when idle



Source: <http://software.intel.com/en-us/articles/intel-xeon-phi-coprocessor-codename-knights-corner>
<http://www.intel.com/content/www/us/en/processors/xeon/xeon-phi-detail.html>

Lagrangian for QCD

What so difficult about this!

$$S = \int d^4x \mathcal{L}$$

$$\mathcal{L}(x) = \frac{1}{4g^2} F_{\mu\nu}^{ab} F_{\mu\nu}^{ab} + \bar{\psi}_a \delta^{ab} \gamma_\mu (\partial_\nu + A_\mu^{ab}) \psi_b + m \bar{\psi} \psi$$

- **3x3** “Maxwell” matrix field & **2+** Dirac quarks
- **1** “color” charge g & “small” quark masses m .
- Sample quantum “probability”: $\text{Prob} \sim \exp[-S]$
- projectiles for HEP (Energy/Intensity) & Nuclear mater.


Wilson Dirac PDE on hypercubic Lattice

$$\bar{\Psi}_{ia}(x) \frac{1 - \overbrace{\gamma_{\mu}^{ij}}}{2} U_{\mu}^{ab}(x) \Psi_{jb}(x + \mu)$$

Dimension:

$$\mu = \hat{0}, \hat{1}, \hat{2}, \hat{3}$$

x_2 axis \uparrow





Color
a = 1,2,3

Spin
i = 1,2,3,4

x_1 axis \rightarrow

Put it on a Lattice



- Quarks and Gluon on Lattice :
 - $L^4 = 100 \times 100 \times 100 \times 100$ points
 - Discretize Dirac PDE as sparse $24 \times L^4$ by $24 \times L^4$ Matrix
 - Gauge Field $4 \times L^4$ 3×3 complex matrices
- Algorithms:
 - Krylov Solver for Quark PDE on $12 \times 3 \times 2 L^4$ unknowns for each
 - Semi-implicit Symplectic Hamiltonian Integrator:
 - Monte Carlo sample by Markov Process

Exact symmetries are powerful quantifier of errors that must vanish at zero lattice spacing

| | Classical Lagrangian/ PDE's | Lattice (i.e. Computer) | Quantum (i.e. Nature) |
|-----------------------------------|--------------------------------|----------------------------|--------------------------|
| Rotational(Lorentz) Invariance | ✓ | ✗ | ✓ |
| Gauge Invariance | ✓ | ✓ | ✓ |
| Scale Invariance | ✓ | ✗ | ✗* |
| Chiral Invariance | ✓ | ✓ | ✗* |

Result: QM spontaneously causes large (and unexpected) large scales.

* *Plus a small extra breaking due to mass of quarks.*

New Frontier: Higher resolution QCD

■ Lattice scales:

$$\begin{array}{ccccccc} a(\text{lattice}) & \ll & 1/M_{\text{proton}} & \ll & 1/m_{\pi} & \ll & L(\text{box}) \\ 0.06 \text{ fermi} & \ll & 0.2 \text{ fermi} & \ll 1 & .4 \text{ fermi} & \ll & 6.0 \text{ fermi} \\ & \underbrace{\hspace{10em}} & & \underbrace{\hspace{10em}} & & \underbrace{\hspace{10em}} & \\ & 3.3 & \times & 7 & \times & 4.25 & \simeq 100 \end{array}$$

■ Consequences:

- Increasing ill-conditioned Dirac operator
- Suffer from worse critical slowing down (CSD)
- $O(100^4)$ lattice volumes or more
- Single grid Krylov methods and homogenous Arch are optimal


Many more LGT mass scales to come

- quarks masses:
($u d s c b t = 2, 5, 100, 1300, 4190, 200000$ MeV)
- Electromagnetism (proton-nucleon splitting, $g-2$)
- Binding energy of nuclei (2.2 MeV for deuteron)
- TeV Strong Gauge BSM (near conformal) dynamics for composite Higgs
- ETC.


MG Scaling for Conformal Laplace

$$A\phi = b \implies \phi(x+h) - 2\phi(x) + \phi(x-h) + h^2 m^2 \phi(x) = b(x)$$

1 $h \rightarrow 2h$ Restriction $R = P^\dagger$



2 $h \rightarrow h$ Prologation P



(1) Blocking preserves the scale invariant const solutions (null state)

(2) Coarse operator is renormalized: $m \rightarrow 2m$ (in units $h = 1$)

QCD MG attempts in 1990's

See Thomas Kalkretuer
 hep-lat/9409008
 review on "MG Methods
 for Propagators in LGT".

Israel: Ben-Av, M. Harnatz,
 P.G. Lauwers & S.Solomon

Boston: Brower, Edwards,
 Rebbi & Vicari

Amsterdam: A. Hulsebos,
 J Smit J. C. Vick

Amsterdam: A. Hulsebos,
 J Smit J. C. Vick


| group | operator to be inverted | gauge field | lattice sizes |
|---|---|-------------|---------------|
| "Israel" [3, 13, and references therein] 1989-ongoing | $\not{D} + m$ | 2-d $U(1)$ | $\leq 256^2$ |
| | staggered fermions | 2-d $SU(2)$ | $\leq 256^2$ |
| | | 2-d $SU(3)$ | $\leq 128^2$ |
| "Amsterdam" [14, and references therein] 1990-1992 | $-\not{D}^2 + m^2$ | 2-d $SU(2)$ | $\leq 128^2$ |
| | staggered fermions staggered fermions and Wilson fermions | 2-d $SU(2)$ | $\leq 128^2$ |
| "Boston" [7, and references therein] 1990-1991 | $-\Delta + m^2$ | 2-d $U(1)$ | $\leq 64^2$ |
| | | 4-d $U(1)$ | $\leq 16^4$ |
| | | 2-d $SU(2)$ | $\leq 32^2$ |
| | $(\gamma_\mu + 1)D_\mu + m$ Wilson fermions | 2-d $U(1)$ | 64^2 |
| [29] 1990-1992 | $(\gamma_\mu + 1)D_\mu + m$ Wilson fermions | 2-d $U(1)$ | 64^2 |
| | | 4-d $SU(3)$ | 16^4 |
| "Hamburg" [21, 18, 22, 23, 1, 17, 19, 20, 2, 24] 1990-ongoing | $-\Delta + m^2$ | 2-d $SU(2)$ | $\leq 128^2$ |
| | | 4-d $SU(2)$ | $\leq 18^4$ |
| | $-\not{D}^2 + m^2$ | 2-d $SU(2)$ | $\leq 162^2$ |
| | staggered fermions | 4-d $SU(2)$ | $\leq 18^4$ |

Table 1: Overview of works on MG methods for propagators in lattice gauge theories.


QCD MG “failure” in 1990’s:



Universal of Critical Slowing down:




Gauss-Jacobi (Diamond), CG(circle),
3 level (square & star)



$\beta = 3$ (cross) 10 (plus) 100 (square)

Time = $F(\text{spontaneous length/quark compton length})$

Why Didn't It work?



Classical QCD (with zero mass quarks) has no scale. BUT spontaneous Conformal symmetry breaking magically gives the proton mass scale.

Success & Failures of MG attempts in 1990's : Why?

- Partial success: weak coupling “renormalization”
- Maintain Exact Gauge invariance
- Maintain exact γ_5 Hermiticity
- Local adaptive blocking: Projective MG
- Chiral Symmetry (density of small e.v.)
- Null vector: Atya- Singer Index Theorem



Boris Grigoryevich Galerkin ([Russian](#): Бори́с Григо́рьевич Гале́ркин, surname more accurately [romanized as Galyorkin](#); March 4 [[O.S.](#) February 20, 1871] 1871 – July 12, 1945),

$$H = \gamma_5 D = D^\dagger \gamma_5$$

$$\text{Prolongator } P \implies \text{Restrictor } R = P^\dagger \gamma_5$$

How do we get beyond the rough confinement barrier?

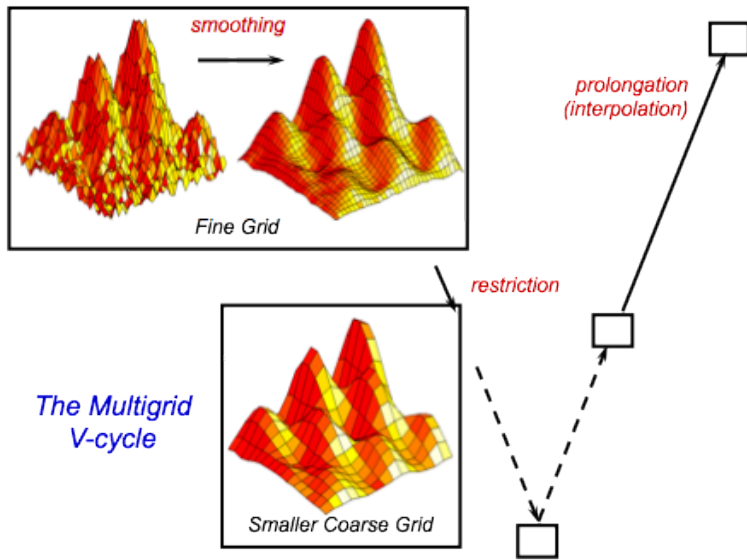
- 1990 Projective MG: First “partitioned” in Jacobi grid blocks. Second “project” near null block vectors
- Ok for weak fields (weak coupling Renormalization Group, ignores instantons for example.)
- 2005 David Keyes to BU with new “adaptive MG idea”. Brannick et al tried it for 2-d Dirac Eq. – slow algorithm but no critical slowing down!
- 2010 Practical QCD AMG: First “project” onto near null vectors (bad guys). Second “partition” into coarse grid.

First Success: Applied Math/Physics Collaboration Collaboration

Many different people (TOPS, QCD) and institutions involved in the collaboration



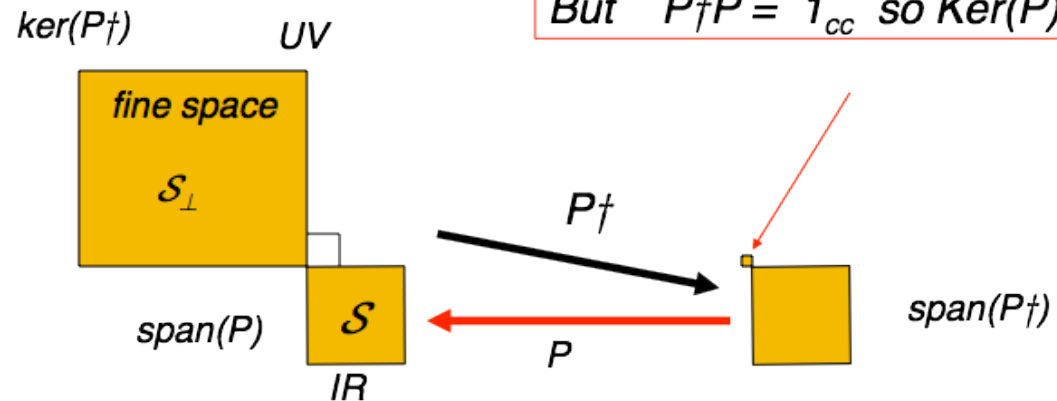
Adaptive Smooth Aggregation Algebraic Multigrid



Split the vector space into near null space \mathcal{S} and the complement \mathcal{S}_\perp

$$D: \mathcal{S} \simeq 0$$

But $P^\dagger P = 1_{cc}$ so $\text{Ker}(P) = 0$



(see Front cover of Strang's Undergraduate MIT math text!)

AMG on Wilson-clover Dirac Operator

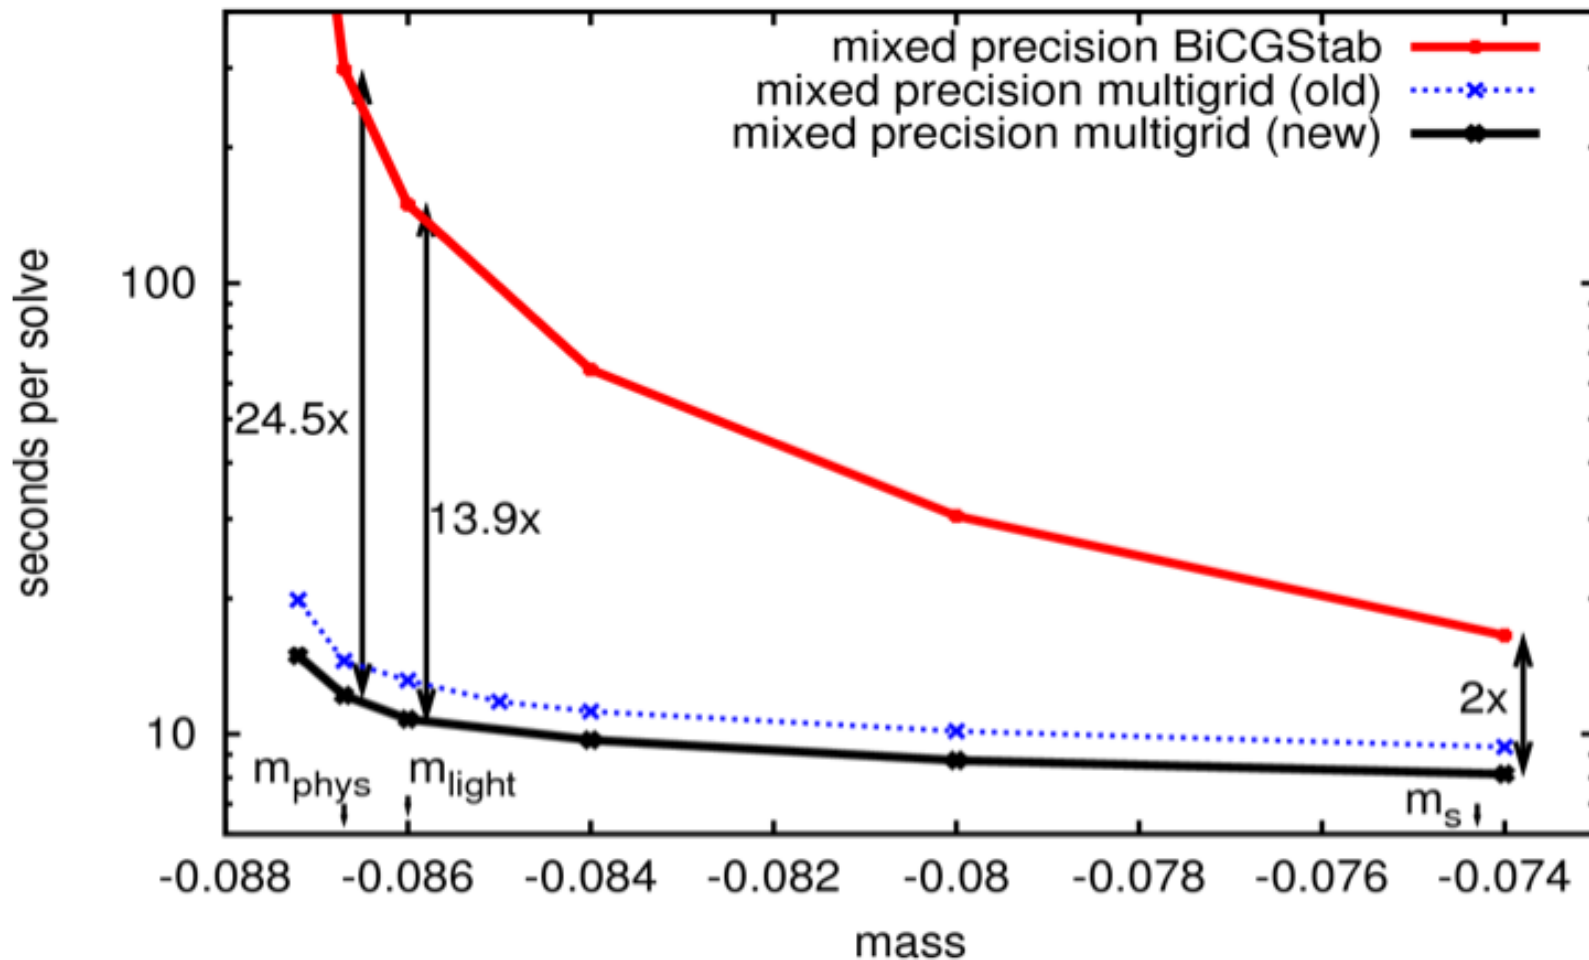
- Devil is in the details!
 - Rigorous MG proofs for normal equation ($D^\dagger D \psi = b$)
 - But would like to project D to avoid higher complexity.
 - Multigrid is recursive to multi-levels.
 - Must preserve Gauge invariance and γ_5 ($[\gamma_5, P] = 0$)
- First benchmarks for Wilson-Dirac Operator:
 - Asym $V=16^3 \times 64, 24^3 \times 64, 32^3 \times 96$ ($N_f = 2, 400\text{MeV pion}$)
 - $N_\nu = 20$ null vectors ψ^s_x with 4th order MR with subset refinement.
 - MG Blocks = $4^4 \times N_c \times 2$ and 3 level V MG cycle
 - pre and post-smoothing is done by 4 iteration GCR (later GMRES)
 - Extend to Red/Black preconditioning

James Osborn implement on BG/P in SciDAC-2 API

Future SciDAC-3 develop in HYPER framework/GPUs etc).

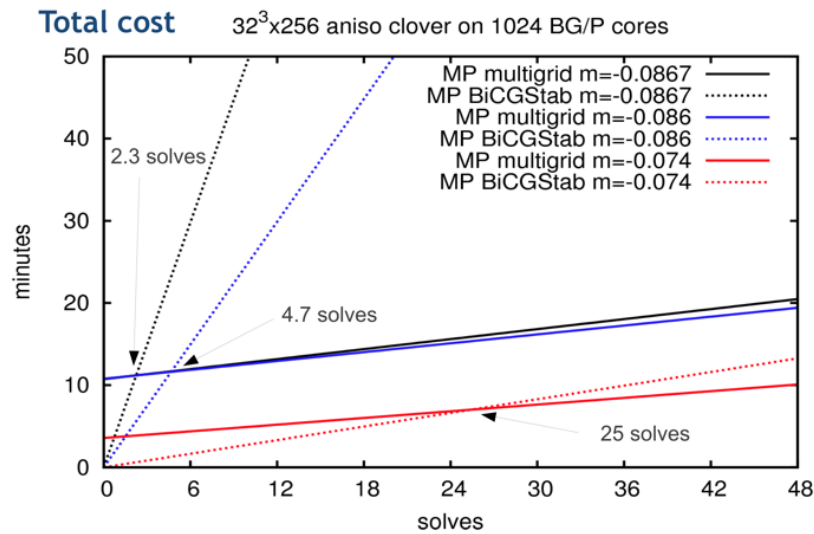
Adaptive Smooth Aggregation Algebraic Multigrid

$32^3 \times 256$ aniso clover on 1024 BG/P cores



"Adaptive multigrid algorithm for the lattice Wilson-Dirac operator" R. Babich, J. Brannick, R. C. Brower, M. A. Clark, T. Manteuffel, S. McCormick, J. C. Osborn, and C. Rebbi, PRL. (2010).

Good News/Bad News



James C. Osborn -- Calculating disconnected diagrams with multigrid -- INT, July 2011

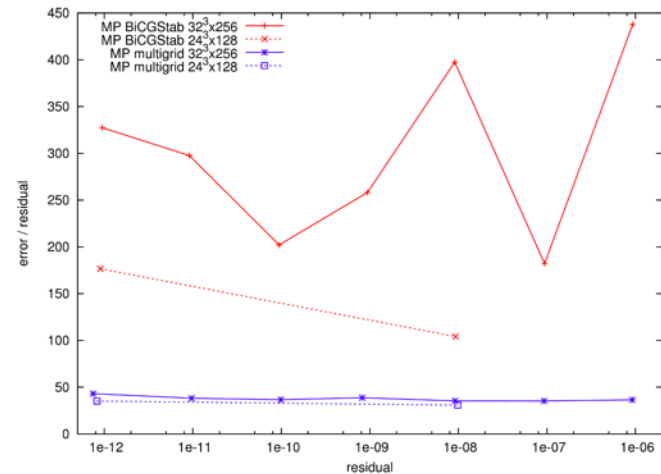
27

Actually MG error
 is smaller at fixed
 Residual

More Data: Should
 Archive MG
 projectors with
 lattice

Error vs residual

- Error:
 $e = x^* - x$
- Residual:
 $r = b - Ax$
 $= Ae$
- Residual not as
 sensitive to low
 modes



James C. Osborn -- Multigrid solver for Wilson clover fermions -- QCDA VI, Sep. 2010

18

Must put Best Algorithm on Best Hardware

- Problem: Wilson Clover for Light Quark is FASTER on the CPU than using the QUDA solver on GPUs!
- Solution put MG on GPU of course



- Cost in \$s reduced by a factor of at least

100+



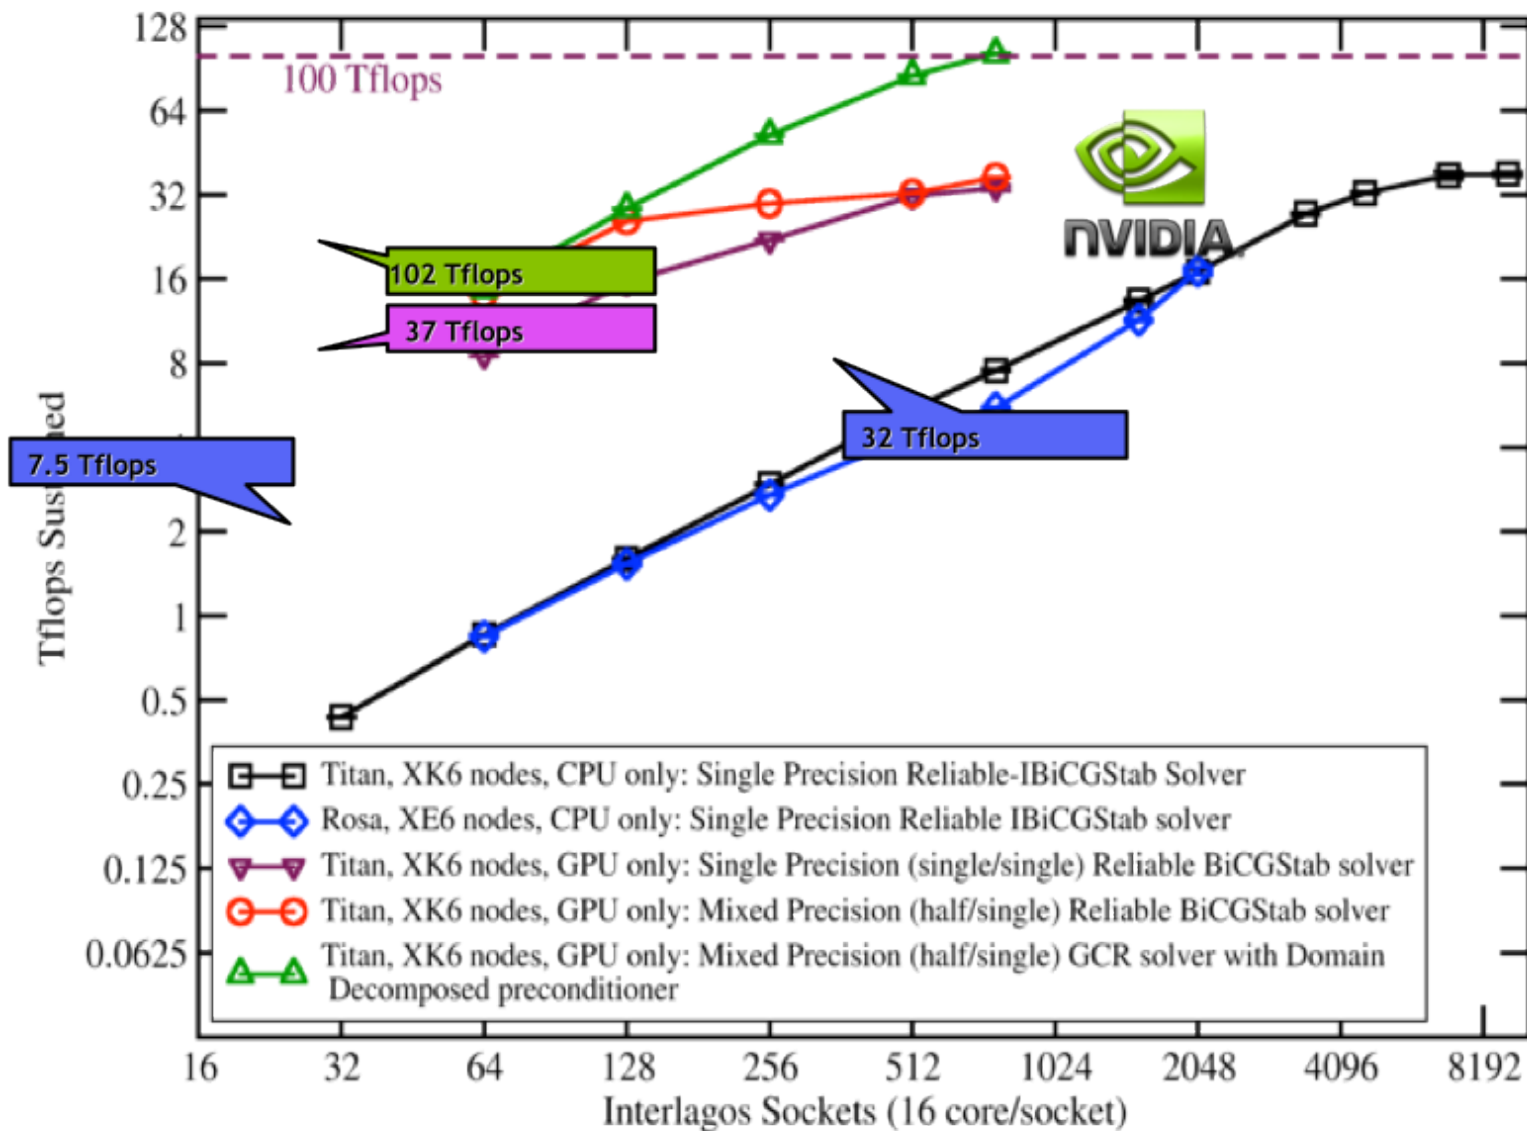
GPU $O(10+)$

MG $O(10+)$

(now with Mike Clark and Michael Cheng on NSF grant)

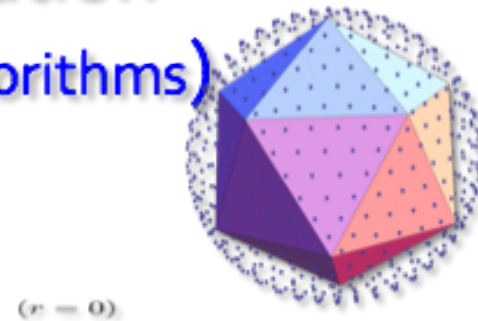

Communication Redcution: DD (Block Jacobi) on Titan

Strong Scaling: $48^3 \times 512$ Lattice (Weak Field), Chroma + QUDA





New Lattice Geometries on Curved Manifolds



- Conformal Radial Quantization
(finite element & graph theoretic algorithms)




- SUSY (Super Symmetry) & Graphene!



Graphene



A/B Bravais sub lattices
 Effective field 2+1 relativistic theory:
 4 copies of 2 component Dirac fields
 Phonons act like gauge fields.



Carbon sheet with Dirac fields: But lattice
 Couple to Coulomb potential and phonons
 Ideal for Lattice field theory, MG and GPU!

$$e_{eff}^2 = \frac{e^2}{\hbar v} \simeq 300 \times \frac{e^2}{\hbar c}$$

Graphene is 2+1 dimension Carbon sheet with Dirac fields: But lattice
 is real Hexagonal structure. Couple to coulomb potential and phones
 act like gauge fields! Ideal for Lattice field theory, MG and GPU!
 (Brower, Rebbi and Schaich)

$\mathbb{R}^4 \rightarrow \mathbb{R} \times \mathbb{S}^3$ Conformal Lattice

- Radial Quantization requires “spatial” spheres!
 - Goal is Conformal BSM Fix Points(or Scattering Length?)
 - Need **Finite Elements Method** to do 3d Ising on curves space!



Contents lists available

Physics

www.elsevier.com

Lattice radial quantization: 3D Ising

R.C. Brower^{a,*}, G.T. Fleming^b, H. Neuberger^{c,1}

^a Department of Physics, Boston University, Boston, MA 02215, USA

^b Department of Physics, Yale University, New Haven, CT 06520, USA

^c Department of Physics and Astronomy, Rutgers University, Piscataway, NJ 08855, USA

ARTICLE INFO

Article history:

Received 3 January 2013

Received in revised form 28 February 2013


Accepted 8 March 2013

Available online 15 March 2013

Editor: M. Cvetič


ABSTRACT

Lattice radial quantization Euclidean conformal field example, we employ a lattice dilations in the 3D Ising two descendants ($l = 1, 2$), from integer spacing for the lattice action will be required continuum limit.



$$a < \Delta r < L \quad \text{vs} \quad a < \Delta \log(r) < L$$

Finite Element Method Laplacian on a sphere



$$K_1(x, y) = [l_{12} - x - \frac{(l_{12} - x_0)y}{y_0}] / l_{12}$$

$$K_2(x, y) = [x - \frac{x_0 y}{y_0}] / l_{12}$$

$$K_0(x, y) = \frac{y}{y_0}$$


triangle on the
tangent plane

$$\text{with } \phi(x, y) = K_i(x, y)\phi_i \implies \int_{A_{012}} dx dy \partial_\mu \phi \partial_\mu \phi(x) =$$


$$\frac{1}{4A_{012}} [l_{02}^2 \phi_1^2 + l_{01}^2 \phi_2^2 + l_{12}^2 \phi_0^2] - \frac{1}{4A_{012}} [(l_{01}^2 + l_{02}^2 - l_{12}^2)\phi_1 \phi_2 + \text{cyclic}]$$

See *WEIGHTS OF LINKS AND PLAQUETTES IN A RANDOM LATTICE** N.H. CHRIST, R FRIEDBERG and T D. LEE Nuclear Physics (1982) quarks and gluons are more difficult!

Spectrum of Laplacian on a sphere



$s = 8$



$s = 512$



$$2.09439l + 2.09439l^2 - 5.75 * 10^{-6}l^3 - 2.95833 * 10^{-6}l^4$$

See posters for more Details

MG/fastMATH

Multigrid with HYPRE for Lattice QCD

Andrew Pochinsky
Massachusetts Institute of Technology
ap@mit.edu



Abstract

In lattice QCD calculations, a significant amount of computation time is spent in solving the Dirac equation. In this context, numerical solution of the Dirac equation is a major bottleneck. In this poster, we describe our approach to solving the Dirac equation using multigrid methods. We present a new multigrid algorithm for solving the Dirac equation using multigrid methods. We present a new multigrid algorithm for solving the Dirac equation using multigrid methods. We present a new multigrid algorithm for solving the Dirac equation using multigrid methods.

LQCD Software

Our LQCD software suite enables fast QCD computations to be performed with high performance across a variety of architectures, including hardware accelerators. Our software suite includes the following components: the multi-scale approach provides robust linear solvers algorithms with accelerated scaling compared to other multigrid methods. We present a new multigrid algorithm for solving the Dirac equation using multigrid methods. We present a new multigrid algorithm for solving the Dirac equation using multigrid methods.

Lattice QCD

Lattice QCD is one of the most demanding computational science fields. It is especially applicable to the study of strong interactions. LQCD requires high-precision floating-point arithmetic and simulation on hardware accelerating the parallel part. The algorithm for solving the discretized Dirac equation is a major bottleneck. The Dirac equation is solved by iterative methods. It is essential to have a high-performance solver for computing the Dirac equation. Presently several formulations of fermion actions are used in different areas of the LQCD program.

- Disaggregated hardware
- Mass-clone fermions
- Compact well fermions (several kinds)

Application of accelerated solver to lattice quantum field theories other than QCD is also of interest.

Multigrid

Multigrid methods form a group of algorithms for solving systems of linear equations using discretization of the exact operator into a hierarchy of approximations. They are especially useful for systems related to discretized linear differential equations and an array of techniques are used to produce accelerating multigrid based on coarse-grid methods. The most basic multigrid is the classical multigrid. The main idea of multigrid is to reduce the computational cost of solving the problem by using a sequence of coarser grids. The Dirac equation is solved by iterative methods. It is essential to have a high-performance solver for computing the Dirac equation. Presently several formulations of fermion actions are used in different areas of the LQCD program.

Linear System Interfaces

The FASTMath SciDAC Institute develops HYPRE [1], a software library of high performance preconditioners and solvers for the solution of large, sparse linear systems of equations. The primary goal of the HYPRE library is to provide users with advanced parallel preconditioners. The present project uses the library's parallel multigrid solvers for structured grids. The HYPRE's conceptual linear system interfaces are further abstracted via the HQL intermediate layer to better map into lattice QCD data types and operations.

Preconditioners

A QUDA interface for solvers and preconditioners high-level description is implemented. To simplify and streamline power needs to be balanced for optimal results. The design and the implementation are work in progress.

Solver catalog

It supports various ways to provide a way to combine preconditioners and solvers as well as a solver preconditioner, providing a level of abstraction to the user. Details of both the design and the implementation are work in progress.

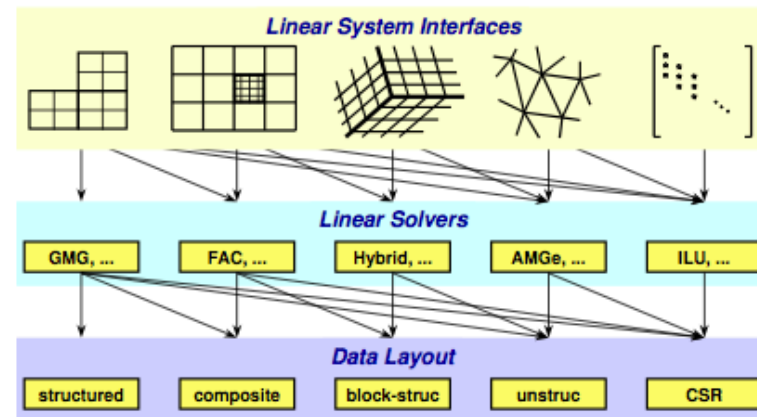
References

[1] <https://arxiv.org/abs/1708.05822>

Auto-tuning in use in QUDA and FUEL

HYPRE

The FASTMath SciDAC Institute develops HYPRE [1], a software library of high performance preconditioners and solvers for the solution of large, sparse linear systems of equations. The primary goal of the HYPRE library is to provide users with advanced parallel preconditioners. The present project uses the library's parallel multigrid solvers for structured grids. The HYPRE's conceptual linear system interfaces are further abstracted via the HQL intermediate layer to better map into lattice QCD data types and operations.



To support LQCD abstractions, two extensions are required to the HYPRE's core:

- support for more than three spatial dimensions
- support for complex numbers

New ideas for Rapid Prototyping using Lua Scripting language: Qlua (MIT) & FUEL (ANL)




Summary

- Success in Higher Resolution Physics and Heterogeneous computers are great for Lattice Gauge Theory Physics.
- Both require an increasingly sophisticated suit of multi-scale algorithms. Combine MG and DD to get both Fast convergence & Communication reduction: **BLTN algorithms?**
- Lattice Field Software will become increasingly intricate and expensive. Need new emphasis on tools adopted from or developed in collaboration with Applied Math and Computer Science.
- Rapid prototyping frameworks, auto-tuning, better compilers and restructuring of the API/Domain Specific Language.
- There is probably a magic bullet but current developments are beginning to suggest solutions.

Questions & Extra Slides

- BLTN (Better Late Than Never) Solver?

Physics+Math+Computing \Leftrightarrow Algorithm



New FUEL HMC framework (Framework for Unified Evolution of Lattices)

- High level layer focused on gauge configuration generation
 - motivation is to have flexible HMC framework to support wide range of beyond standard model theories
 - algorithmic abstraction: generation algorithm independent of gauge group, action, etc.
 - easy to write new high-level algorithms, tune parameters
 - serves as wrapper for efficient “level 3” routines
 - easy to plug in new routines
 - new routines can be written in any other language/framework
- Uses scripting language Lua
 - Small
 - Easy to port (ANSI C89)
 - Easy to use, yet powerful
 - Easy to embed and interface with libraries

Rich: Documentation is being addressed this spring/summer. Michael Chang and Mike Clark are nearly finished implementing the fine level of MG on the GPU. **James, Meifeng and I** are working on integrating Multigrid into Wilson colver evolution code **Kostas and Will 's Wilson isotropic HMC/Chroma trajectory**. Looks promising.

Andrew & Chris: Integration of HYPRE and Qlua is well underway. Rob Falgout, Christopher Schroeder and Andrew Pochinsky have completed an overall design of a HYPRE/USQCD interface (HQL) and begun its implementation. RF is largely finished extending HYPRE to handle more than 3 dimensions and fully expects to finish implementing complex numbers on schedule. CS and RF are making progress on the implementation of the HQL interface, and RF and AP are proceeding with the HQL-Qlua interface. AP is finishing extending Qlua to handle data types and procedures required to support HQL.

Some future directions

- Fermions PDEs are ubiquitous in Quantum Field theories and Nano materials. Lattice geometries and boundary conditions present new fun challenges.
 - Finite T lattice (e.g. 32 x speed up on $128^3 \times 96$ lattice*)
 - Old/New RG Geometric /Adaptive Hybrid MG
 - Monte Carlo Evolution of QCD: Sime implicit integrator
 - Graphene (again classical conformal LGT!)
 - Conformal Theories for LHC Higgsless models
 - Radial lattices for conform/string duals
 - Domain Wall/Overlap 5-d fore EXACT chirality

Other Dirac Operators (For quarks & Electrons)

- Multigrid & DD for Staggered and DW
- Eigenvector for Deflation and Disconnected Diagrams.
- Multi-scale Extension of Symplectic Integrators
- The Hexagonal Lattice for Dirac Electrons of Graphene
- Conformal Lattices on Spheres.
- Topological Defects and the Spin Connection of GR
- Anti-deSitter Space and Conformal Invariance for BSM
- and condensed matter
- etc.

Outline:

One: Keeping pace with current platforms!

BG/Q(IBM) ,Titan (NVIDIA), Stampede (Intel),...

Two: Exposing and Exploiting Multiscale Physics

Protons, Nuclei and beyond to Higgs

Three: Conforming Physics to Hardware.

“physics” and “architecture” is multi-scaled but not necessarily compatible.

Future: New lattices LHC Physics & Condensed Matter