Lattice Gauge Theory Software & Algorithms*

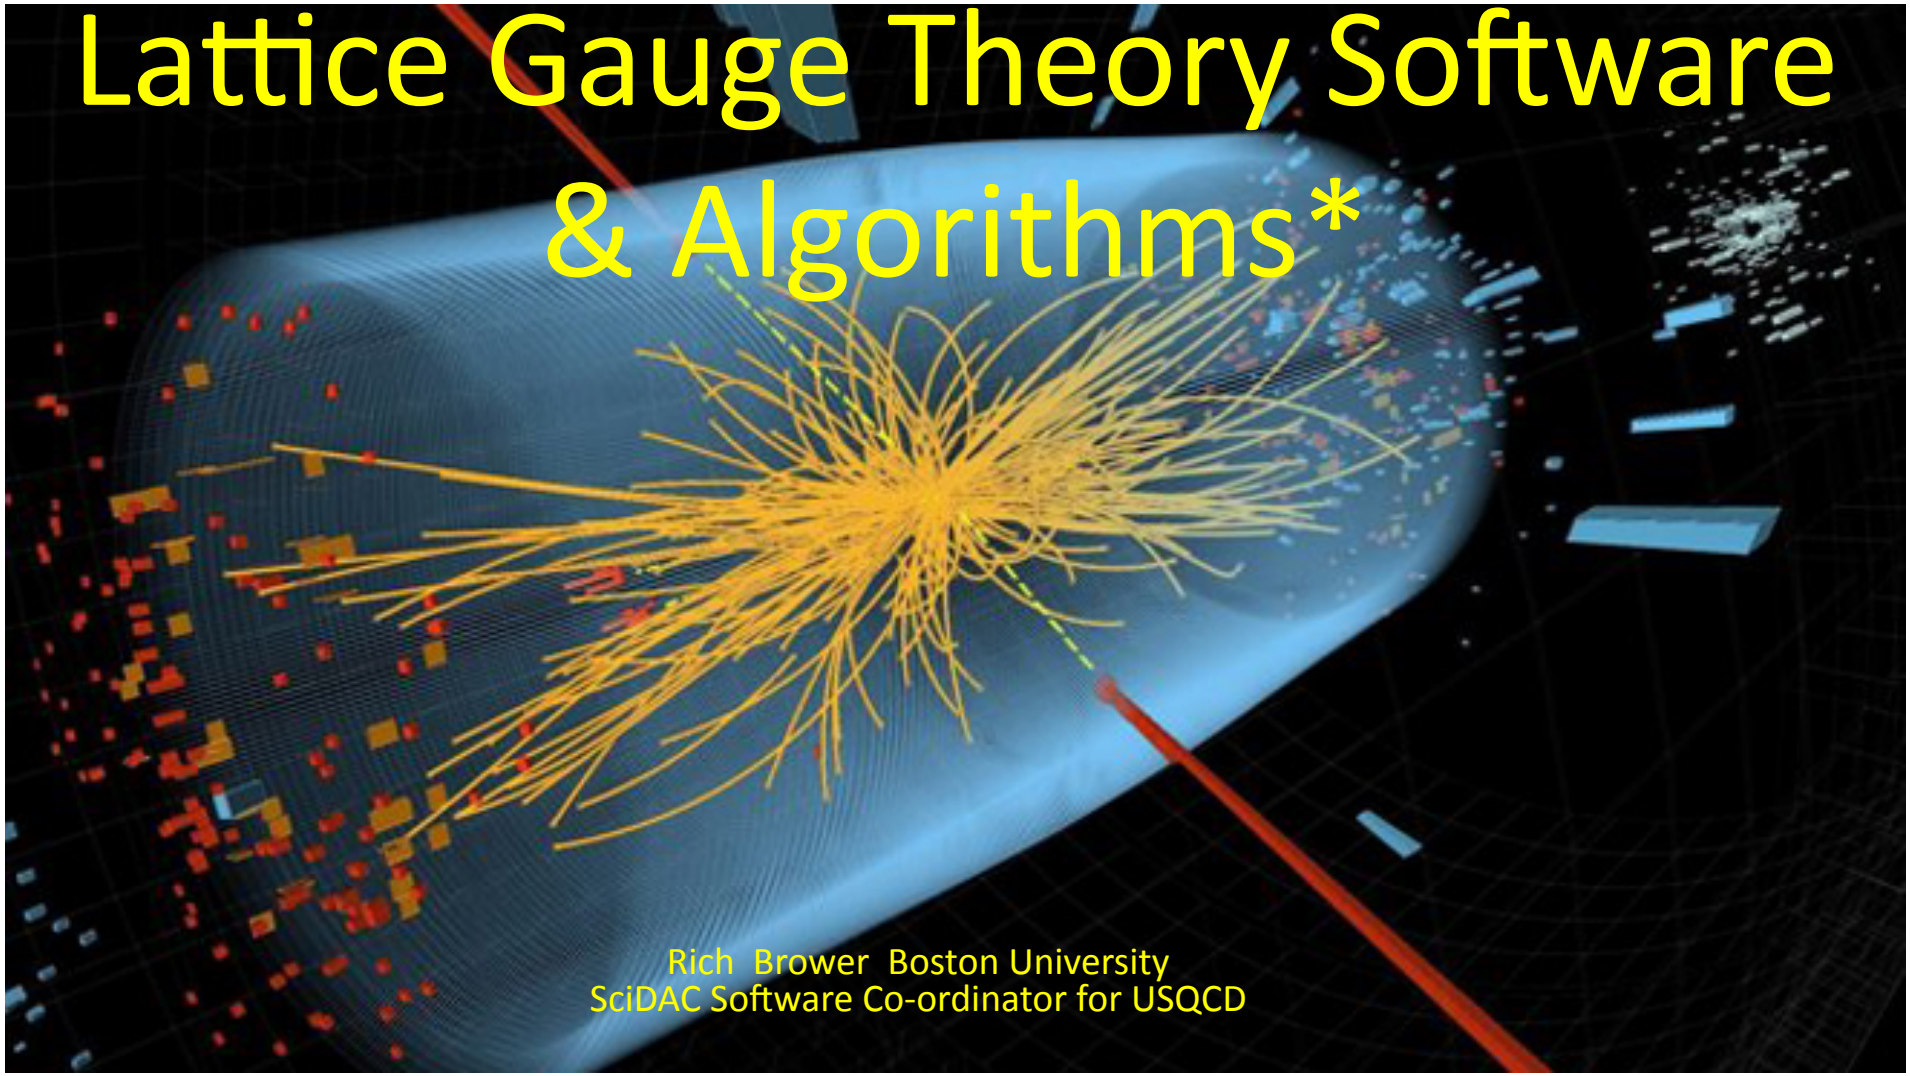

Rich Brower Boston University
SciDAC Software Co-ordinator for USQCD

Argonne Training Program on Extreme-Scale Computing
August 13, 2014

(*The Art of Exact Approximations!)

Outline

- Physics Architecture
- Hardware Architecture
- Software Architecture
  - Multigrid for QCD – 20+ year effort
  - GPUs -- back to the future

Finding the Overlap is difficult but very rewarding

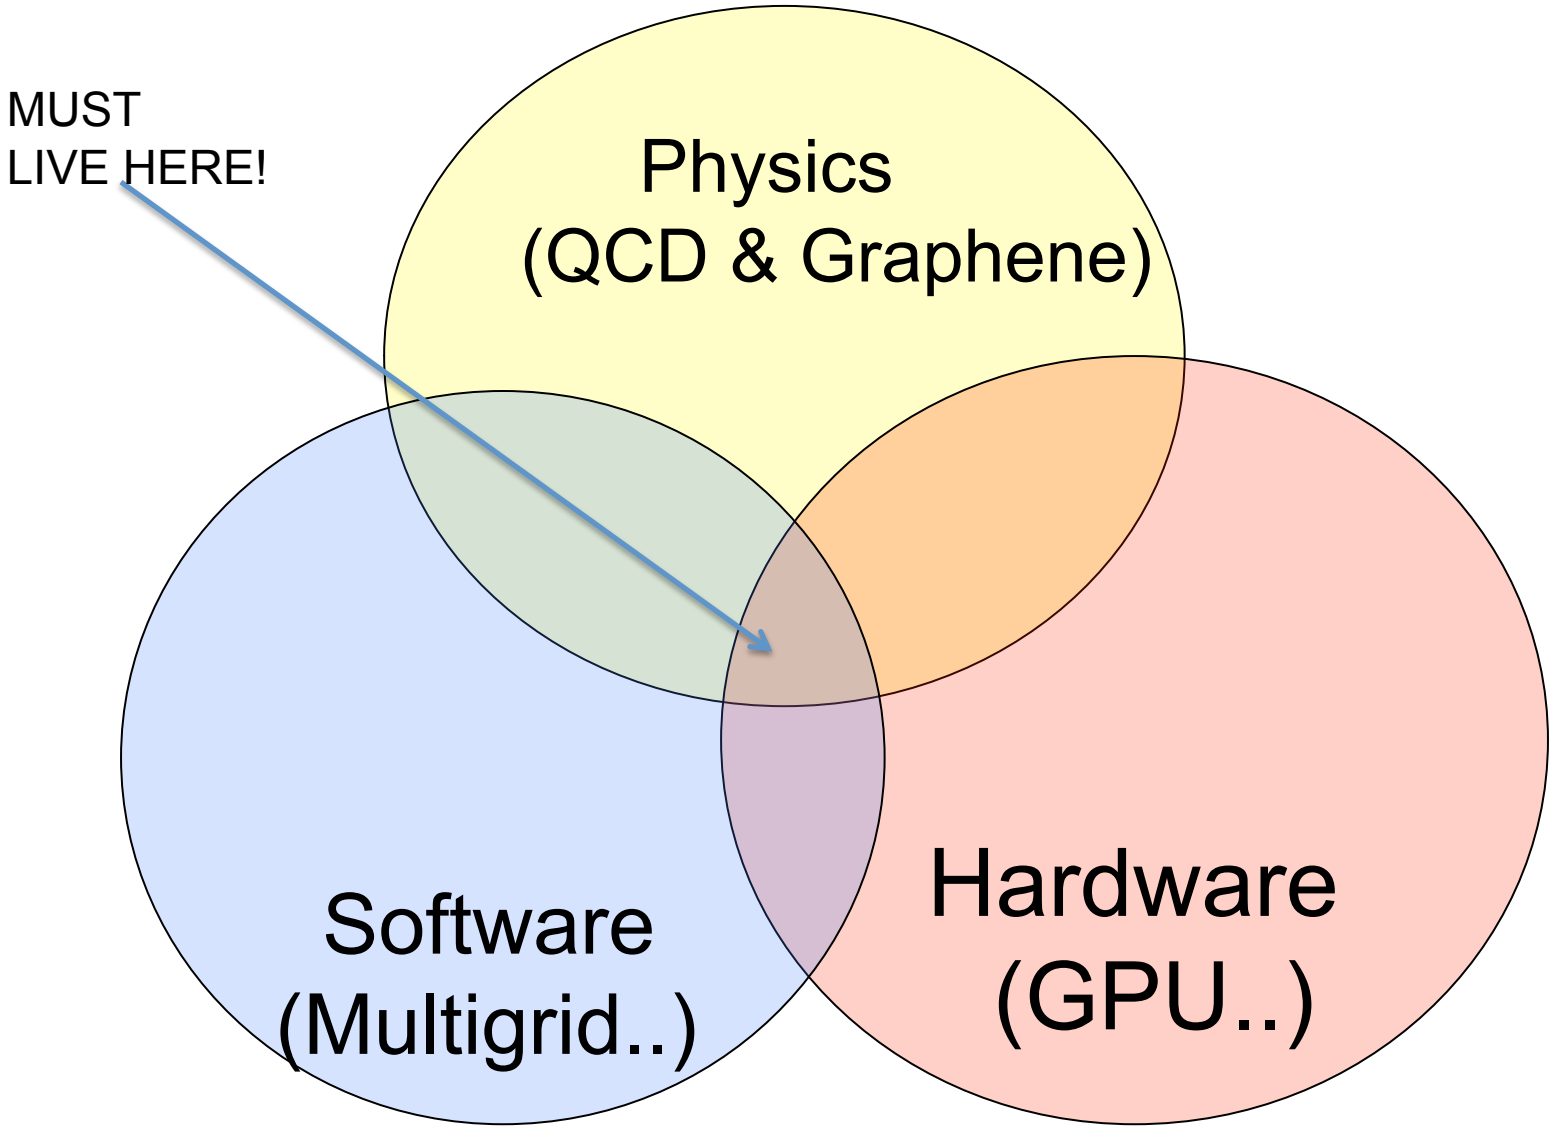

PHYSICS

All prediction from Quantum Field Theory require "Algorithms"

Z = Path Integral exp[- Action]

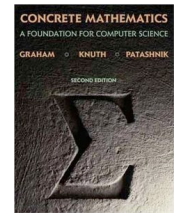

PDE/FEM
Schwartz
Schurs

Feynman
Diagrams

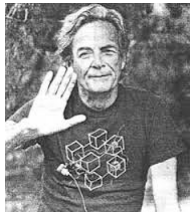

Wilson
Lattice QCD

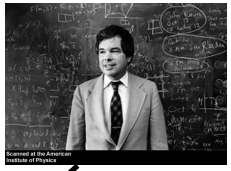

OPE &
Renormalization
Group

Multi-
Grid

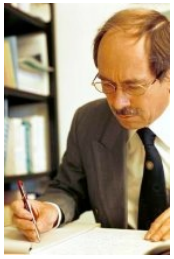

'tHooft
Dim Reg

Domain
Wall

Twisters
Bootstrap

AdS/CFT

Wilson
Flow

DD

Physic+Math+Computing <=> Algorithm

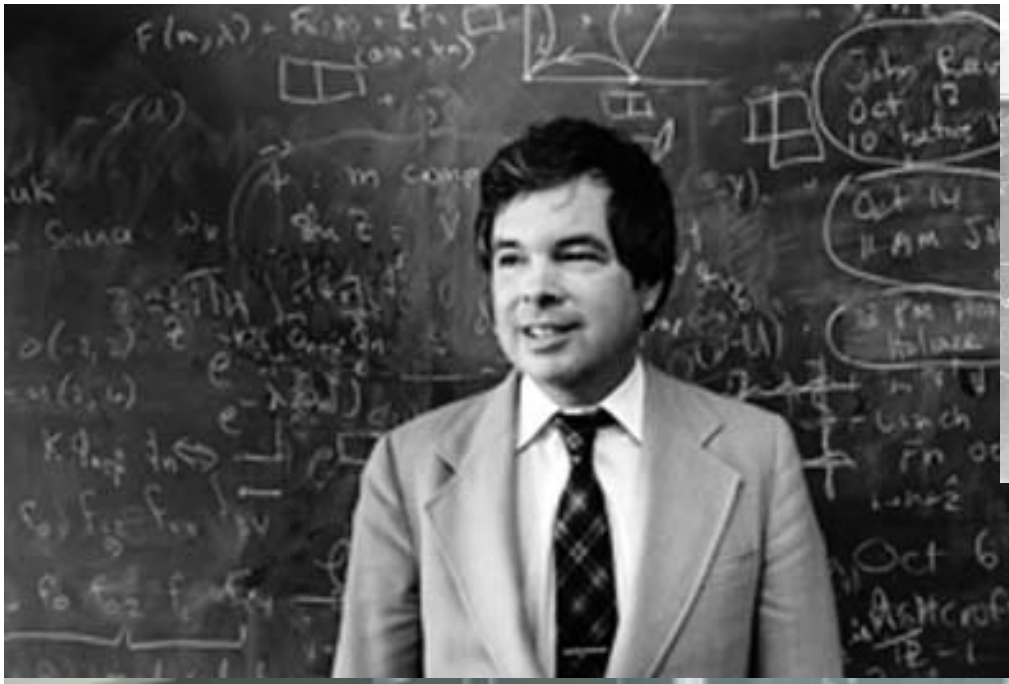

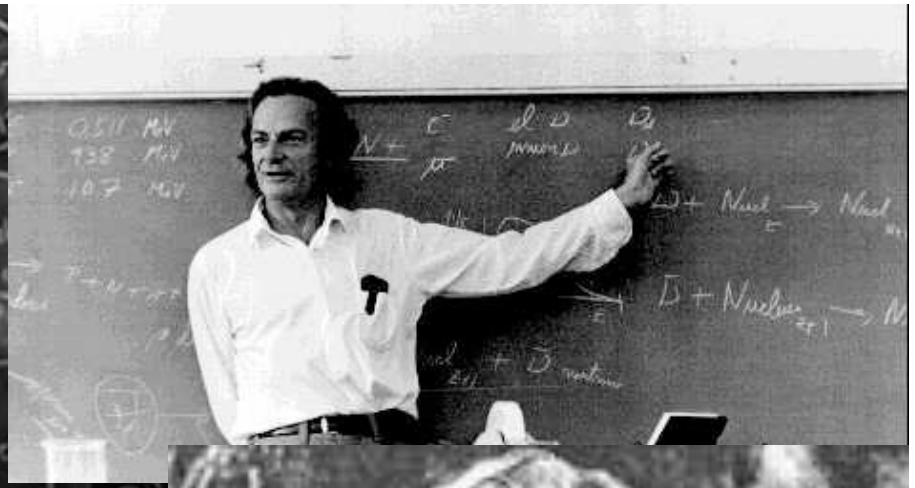

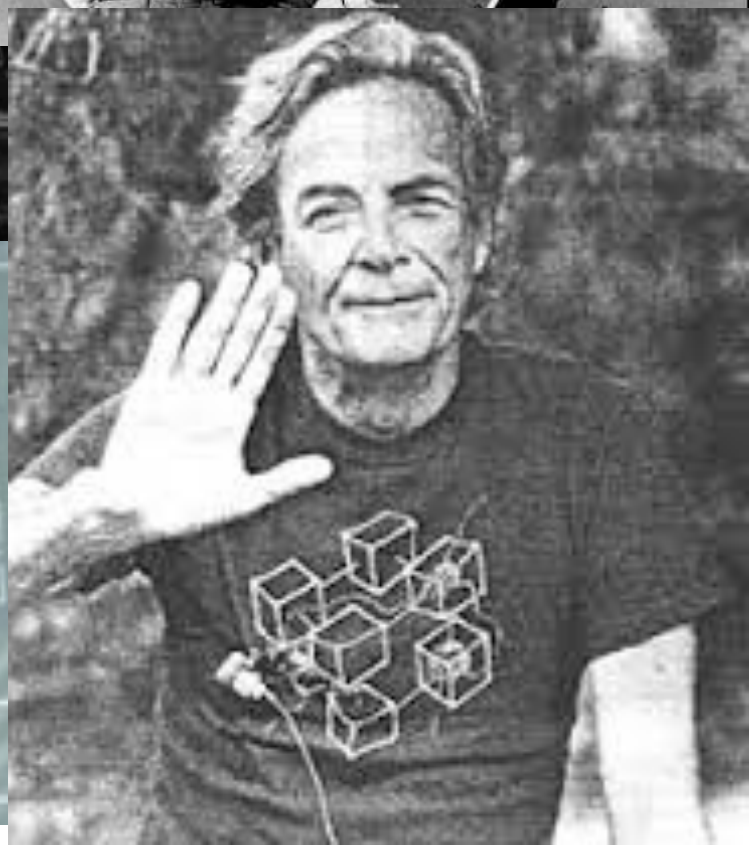

Kenneth G. Wilson "Confinement of quarks"

Phys. Rev. D 10, 2445–2459 (1974)

Abstract

A mechanism for total confinement of quarks, similar to that of Schwinger, is defined which requires the existence of Abelian or non-Abelian gauge fields. It is shown how to quantize a gauge field theory on a discrete lattice in Euclidean space-time, preserving exact gauge invariance and treating the gauge fields as angular variables (which makes a gauge-fixing term unnecessary). The lattice gauge theory has a computable strong-coupling limit; in this limit the binding mechanism applies and there are no free quarks. There is unfortunately no Lorentz (or Euclidean) invariance in the strong-coupling limit. The strong-coupling expansion involves sums over all quark paths and sums over all surfaces (on the lattice) joining quark paths. This structure is reminiscent of relativistic string models of hadrons.

Lattice Field Coming of Age

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?

*15 Years ahead of schedule!

QCD is Maxwell on SU(3) lattice

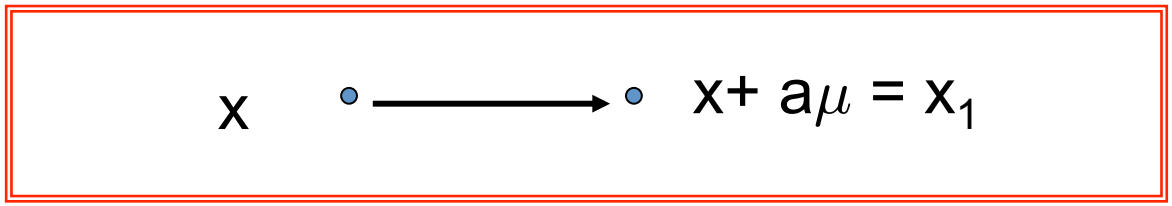

Finite difference: $$\frac{\partial\phi(x)}{\partial x_\mu} \to \Delta_\mu\phi(x) = \frac{\phi(x+a\mu)-\phi(x)}{a}$$

With Gauge field replace: $$\Delta_\mu\phi(x) = \frac{e^{i\int_{x+a\mu}^{x} A_\mu dx_\mu}\phi(x+a\mu)-\phi(x)}{a}$$

The new factor is covariant constant.

This is the Lattice Guage link: $U(x,x+a\mu) = e^{iaA_\mu(x)}$

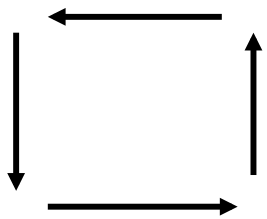

$$= e^{iaA_\mu(x)}e^{iaA_\mu(x)}e^{-iaA_\mu(x)}e^{-iaA_\mu(x)}$$

$$\simeq e^{ia^2(\partial_\nu A_\mu(x)-\partial_\mu A_\nu(x)-A_\mu(x+\nu)+i[A_\mu(x),A_\nu(x)])}$$

QCD on the Lattice: Base/Sparse vs Fiber/Dense

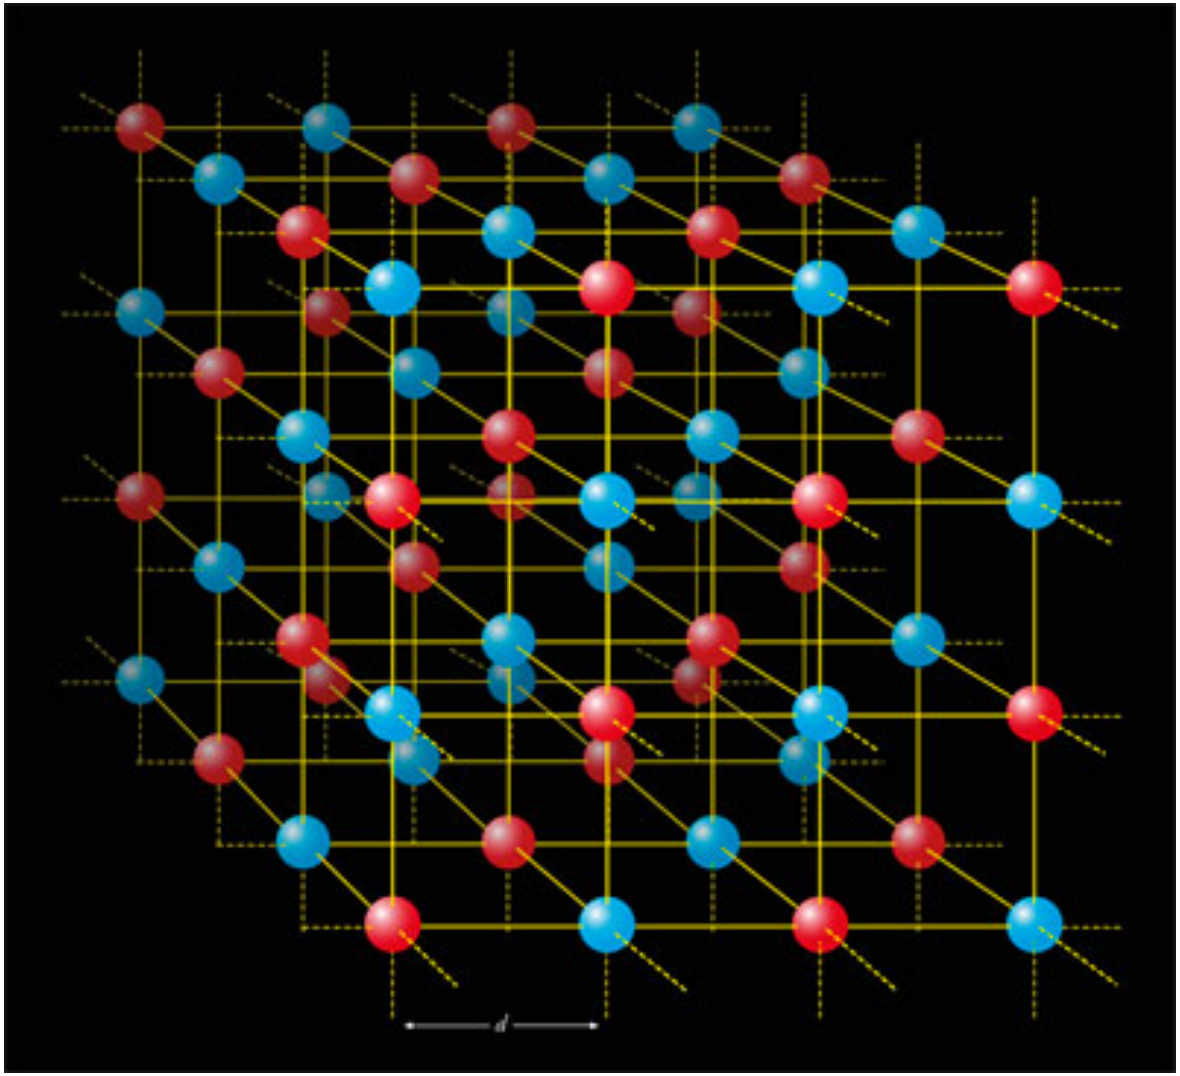

QCD Perfect scaling!

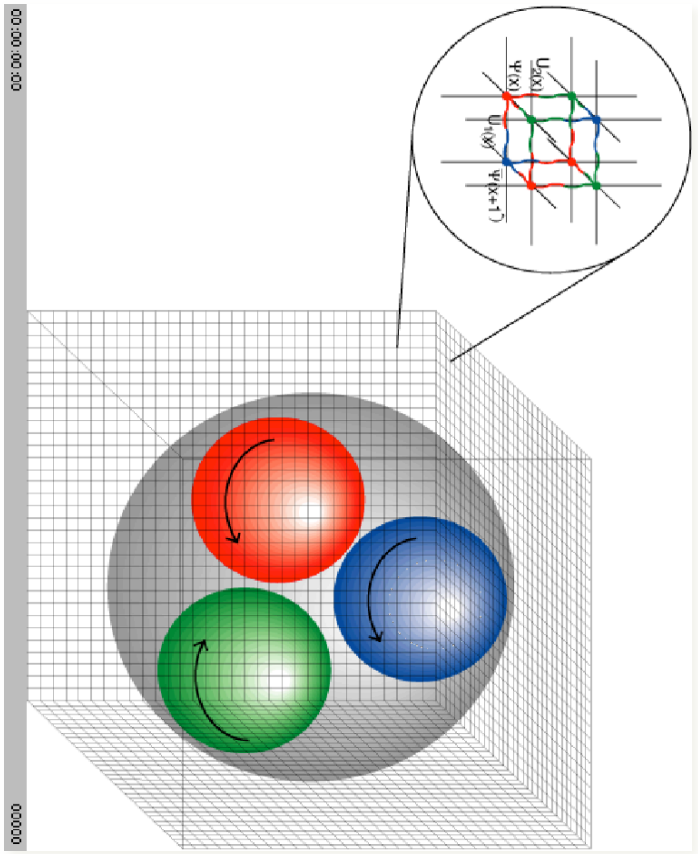

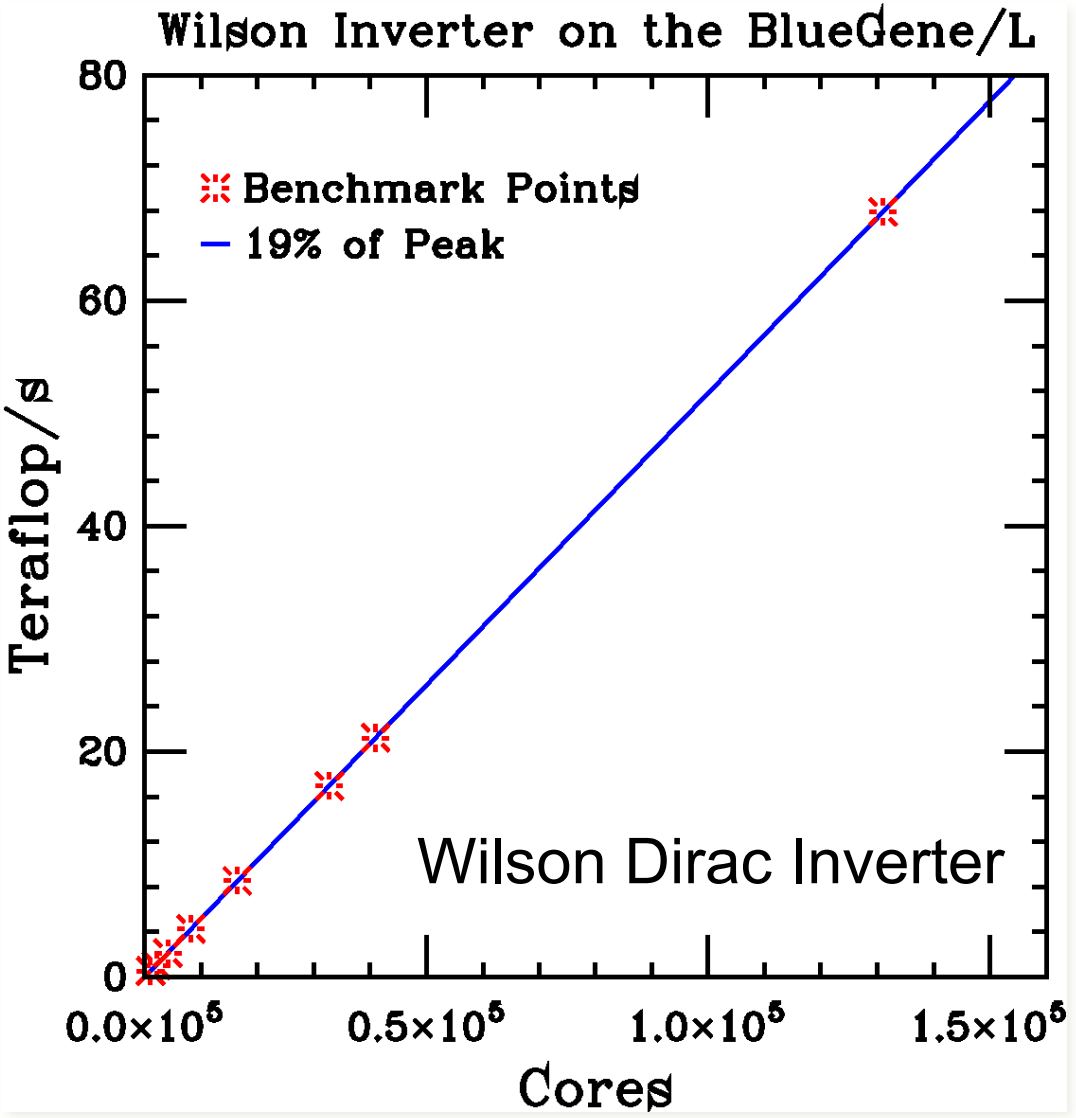

† LLNL BG/L weak scaling up to 131,072 cores: 2006 Gordon Bell award

by Vranas, Bhanot, Blumrich, Chen, Gara, Giampapa, Heidelberger, Salapura and Sexton

QCD: Fundamental Theory of Nuclear Force

Partition function

Anti-quark

quark

Gauge (Glue)

$$= \int d\bar{\Psi}(x) d\Psi(x) dA_\mu(x)$$ [Probability Density]

$$= \int d\bar{\Psi}(x) d\Psi(x) dA_\mu(x) \quad \exp[-\int d^4x \bar{\Psi} D \Psi - \frac{1}{g^2} \int d^4x F^2]$$

$$= \int dA_\mu(x) \quad Det[D] \quad \exp[-\frac{1}{g^2} \int d^4x F^2]$$

Maxwell (Curl)

Dirac Operator

Computational Approach, Numerical Methods

$$\mathrm{Prob}[U,\phi] = Z^{-1} e^{\beta Tr[U_{glue} + U_{glue}^{\dagger}] + \bar{\phi}(D_{quark}^{\dagger} D_{quark})^{-1}\phi}$$

- Monte Carlo importance sampling of gauge configurations:
  - Generate Quark Gluon background ensemble in Probability:
- Hybrid Monte Carlo:

  Molecular Dynamics Algorithm:
  - Multi-time step Hamiltonian evolution in "potential": - Log(Probability).
- Repeated solution of Dirac equation
  - (large sparse linear system) at each step

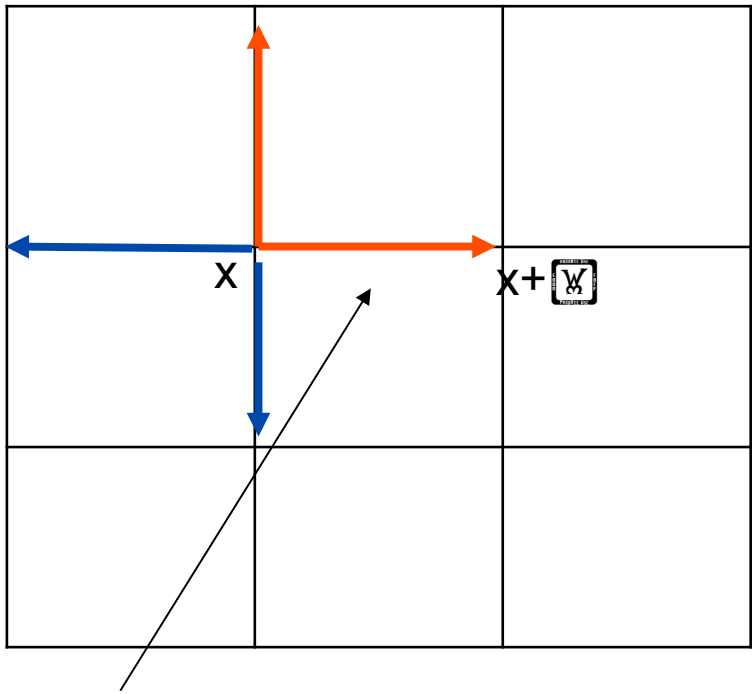

$$D_{quark} = m_q + \frac{1-\gamma_\mu}{2} U(x, x+\mu) + \frac{1+\gamma_\mu}{2} U(x+\mu, x)$$

QCD Lattice Measurement

$$\int dU_\mu d\bar{\Psi} d\Psi \; [\; \Psi(x)^3 \quad \bar{\Psi}(z)\gamma_\mu \Psi(x) \quad \bar{\Psi}(y)^3 \;] e^{-S} =$$

$$\int dU_\mu \; Det[D] \; [\; D^{-1}(x,y) D^{-1}(x,y) \; D^{-1}(x,z)\gamma_\mu D^{-1}(z,y) \;] e^{-S} +$$

$$\int dU_\mu \; Det[D] \; [\; D^{-1}(x,y) D^{-1}(x,y) D^{-1}(x,y) \; Tr[\gamma_\mu D^{-1}(z,z) \;] e^{-S}$$

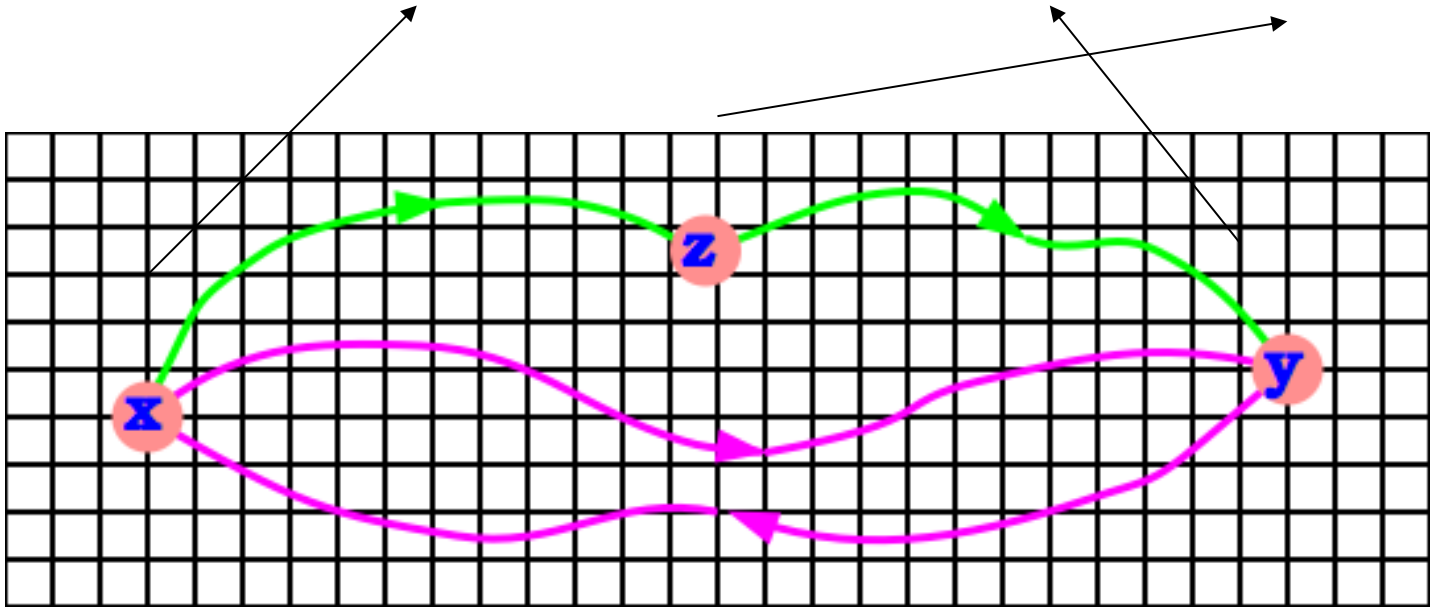

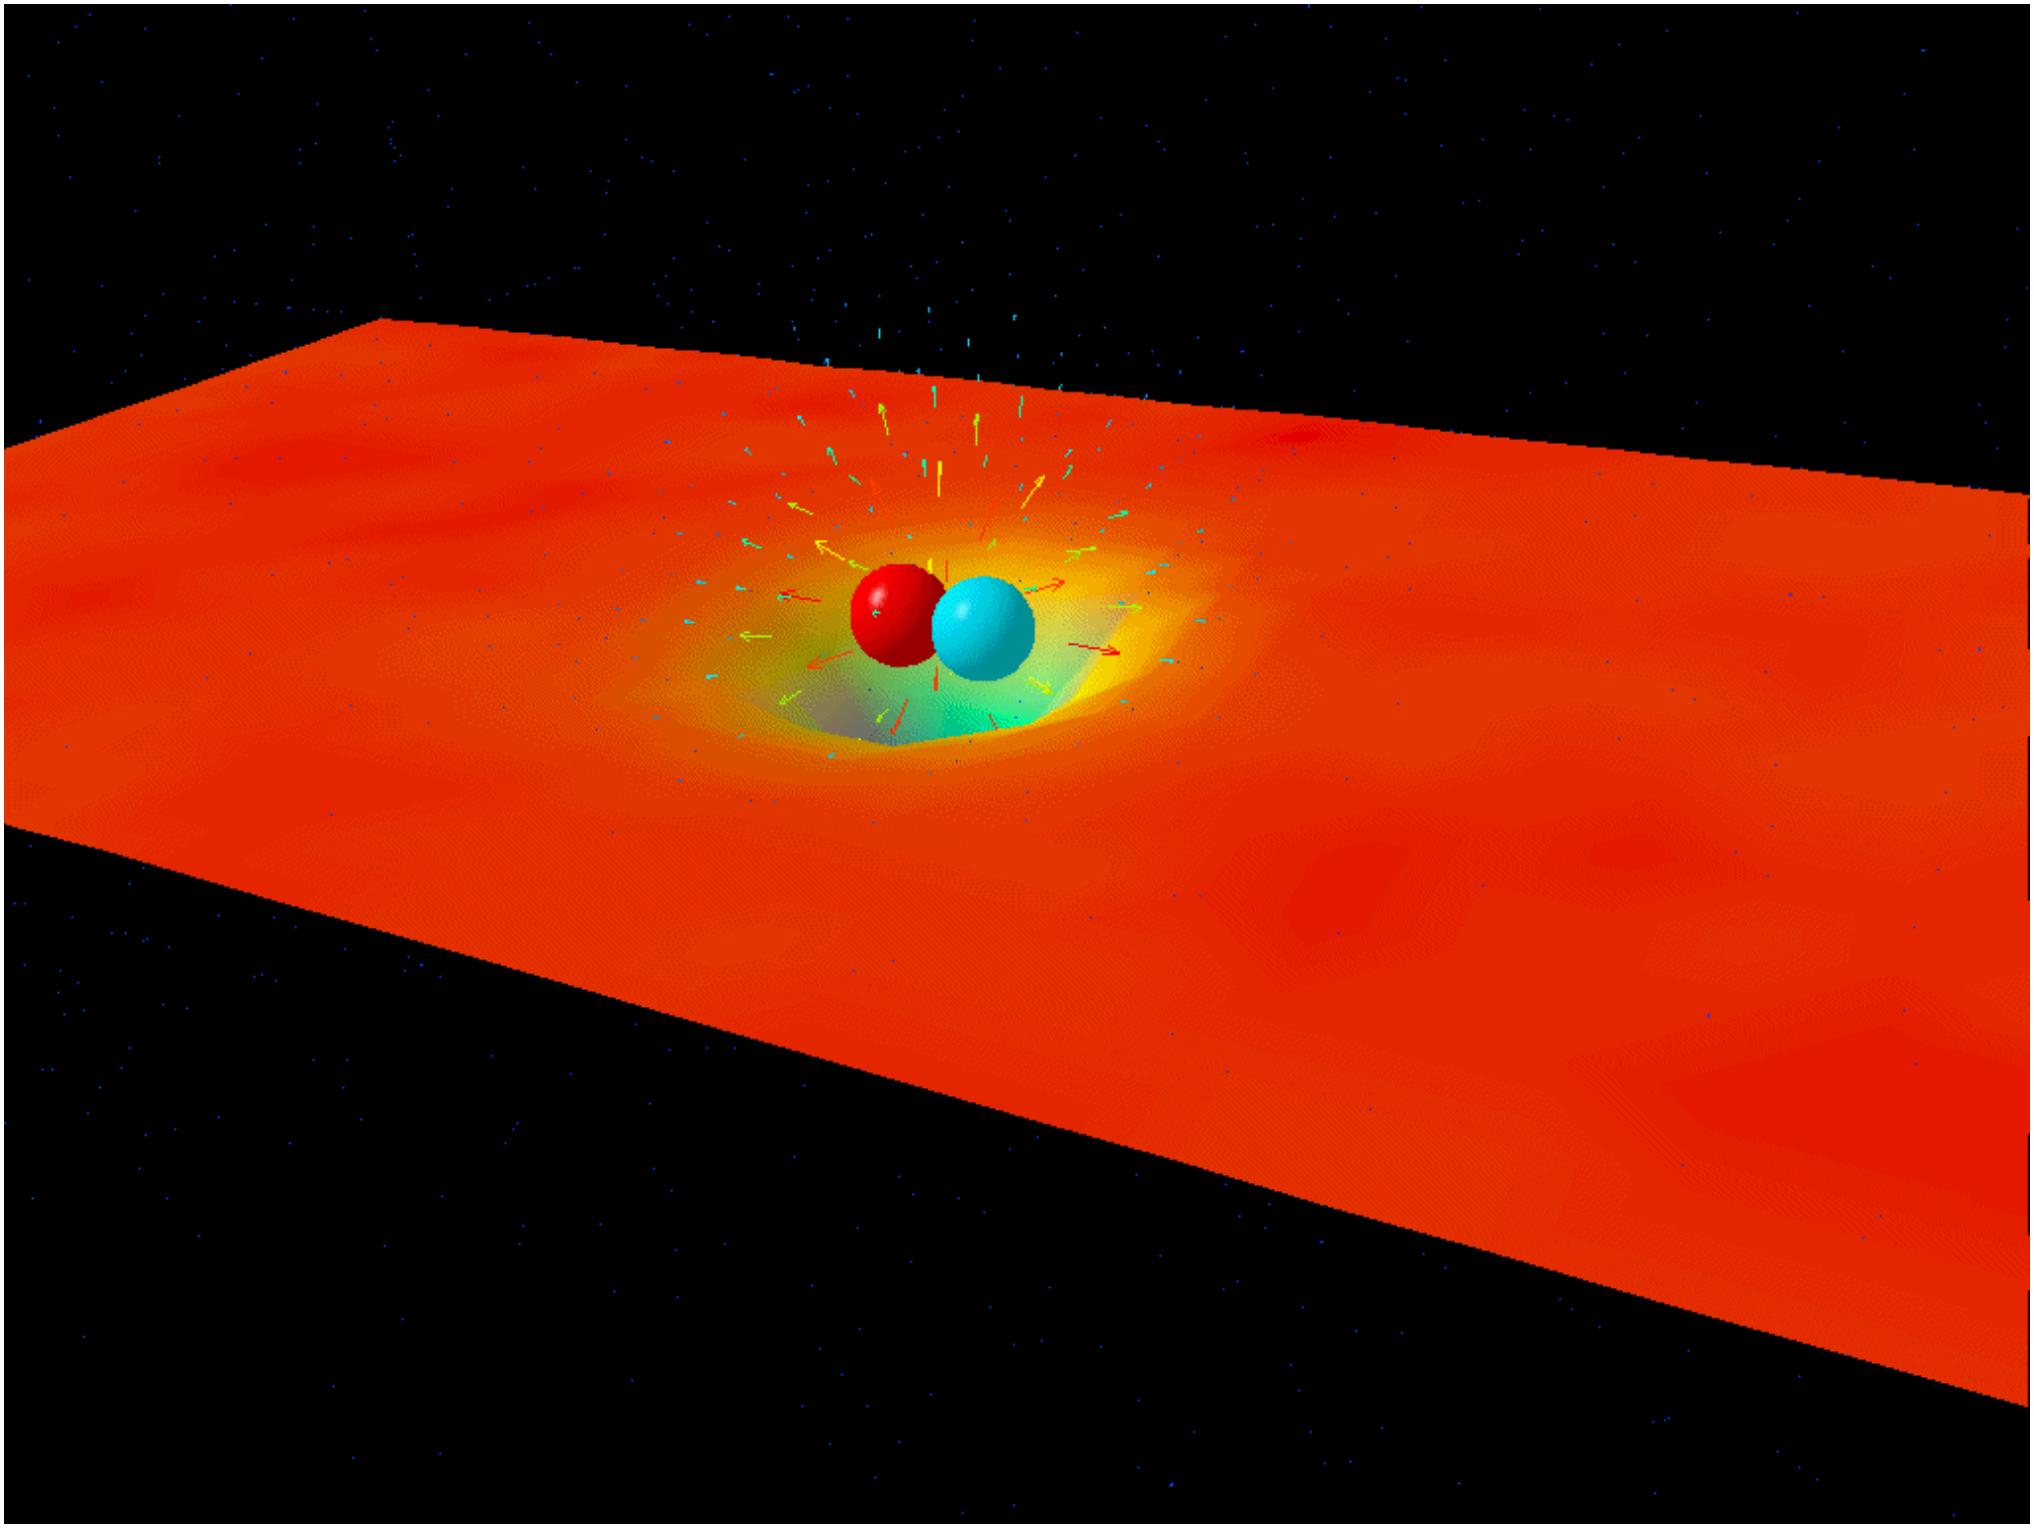

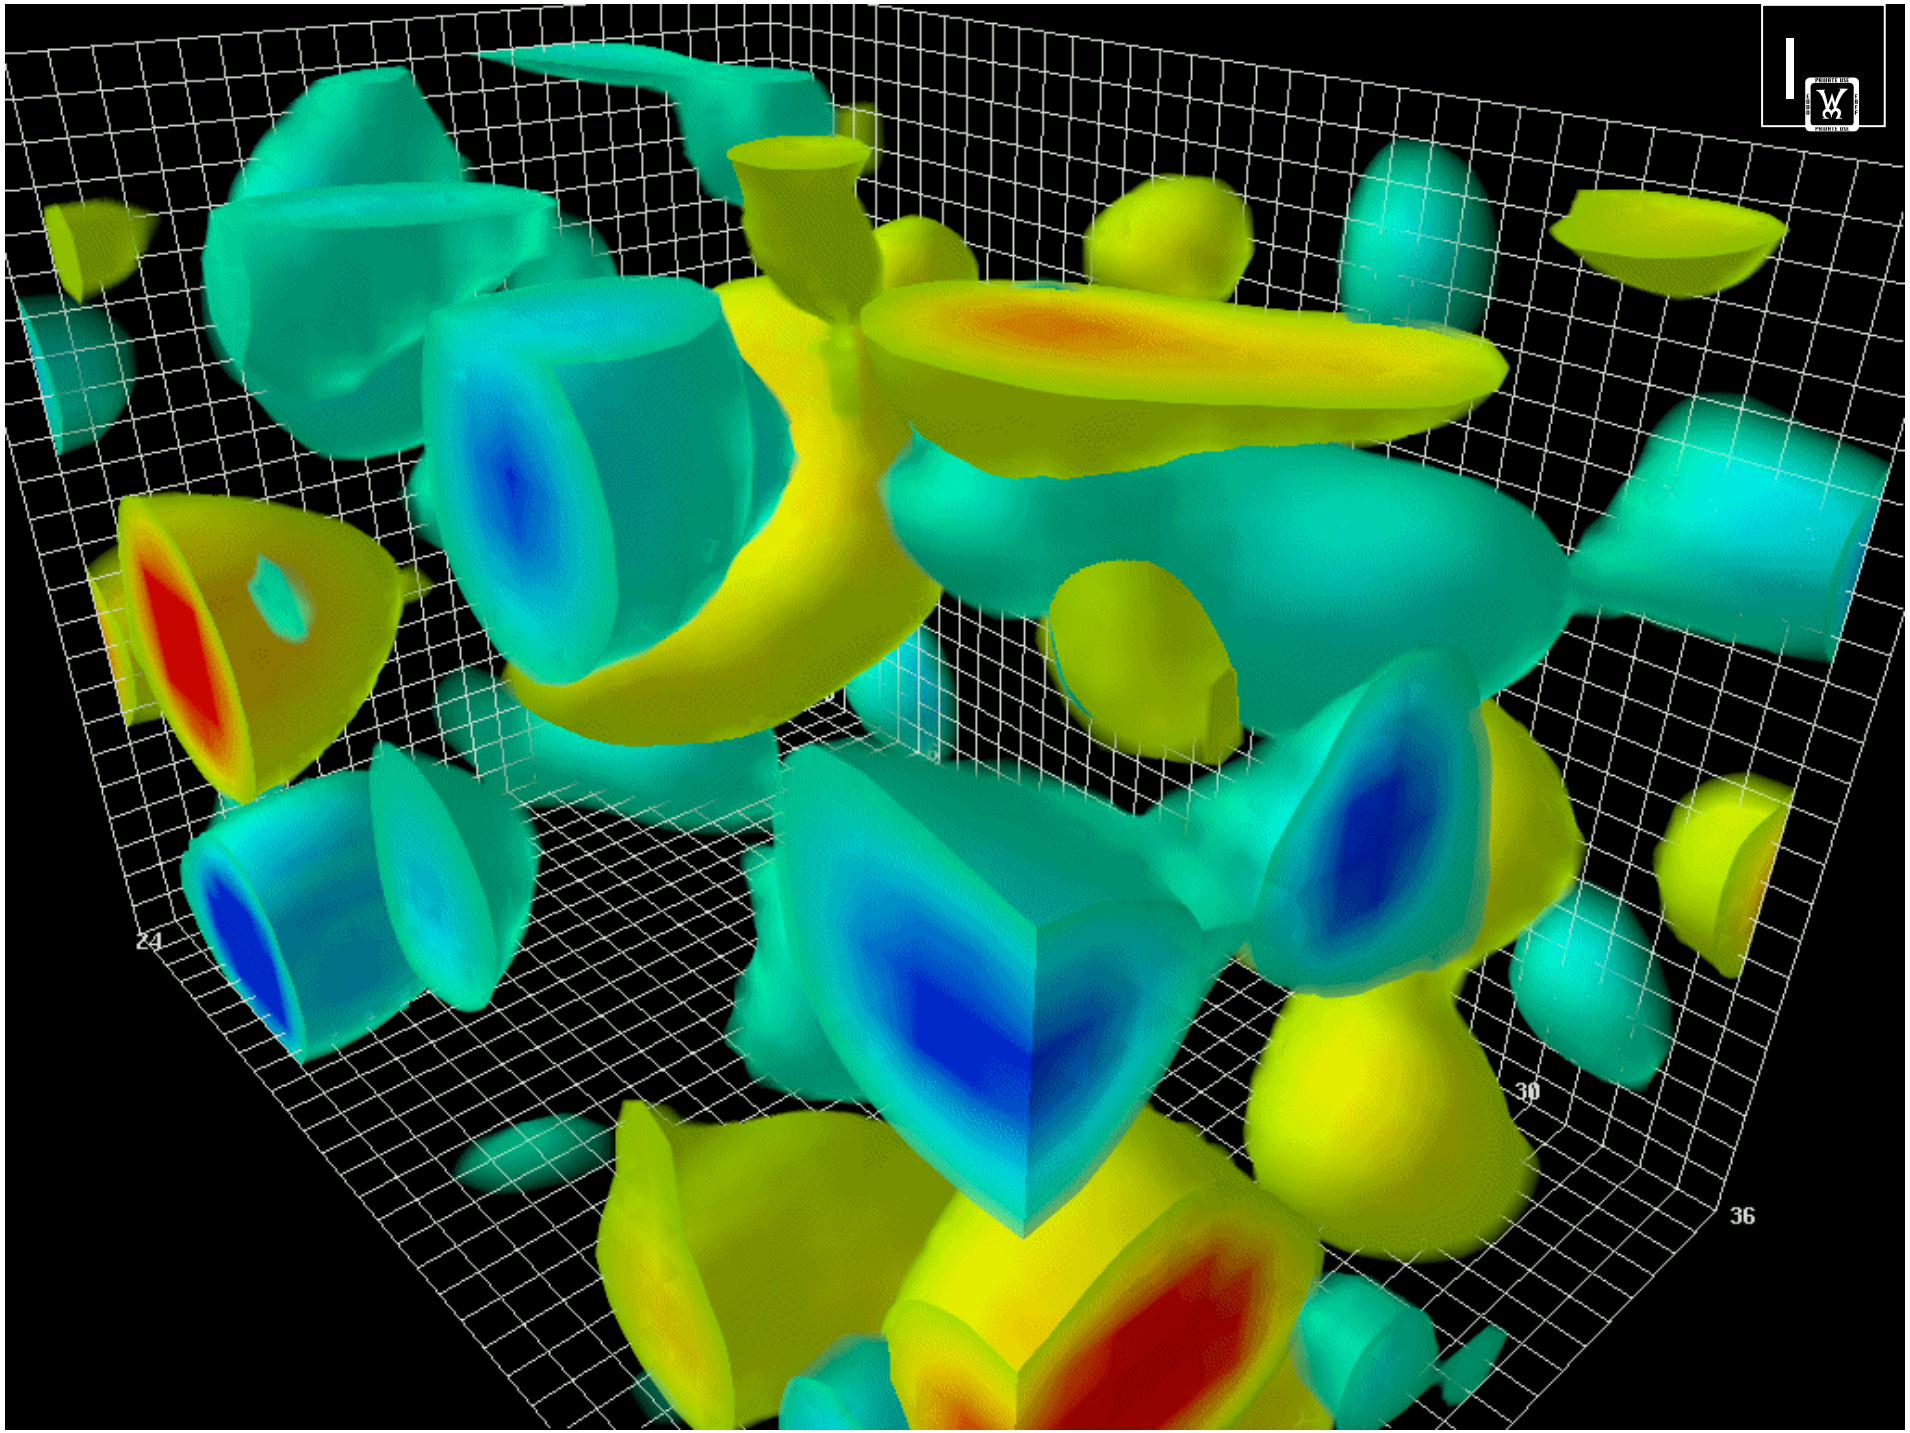
24
30
36

Lattice QCD – is extremely uniform

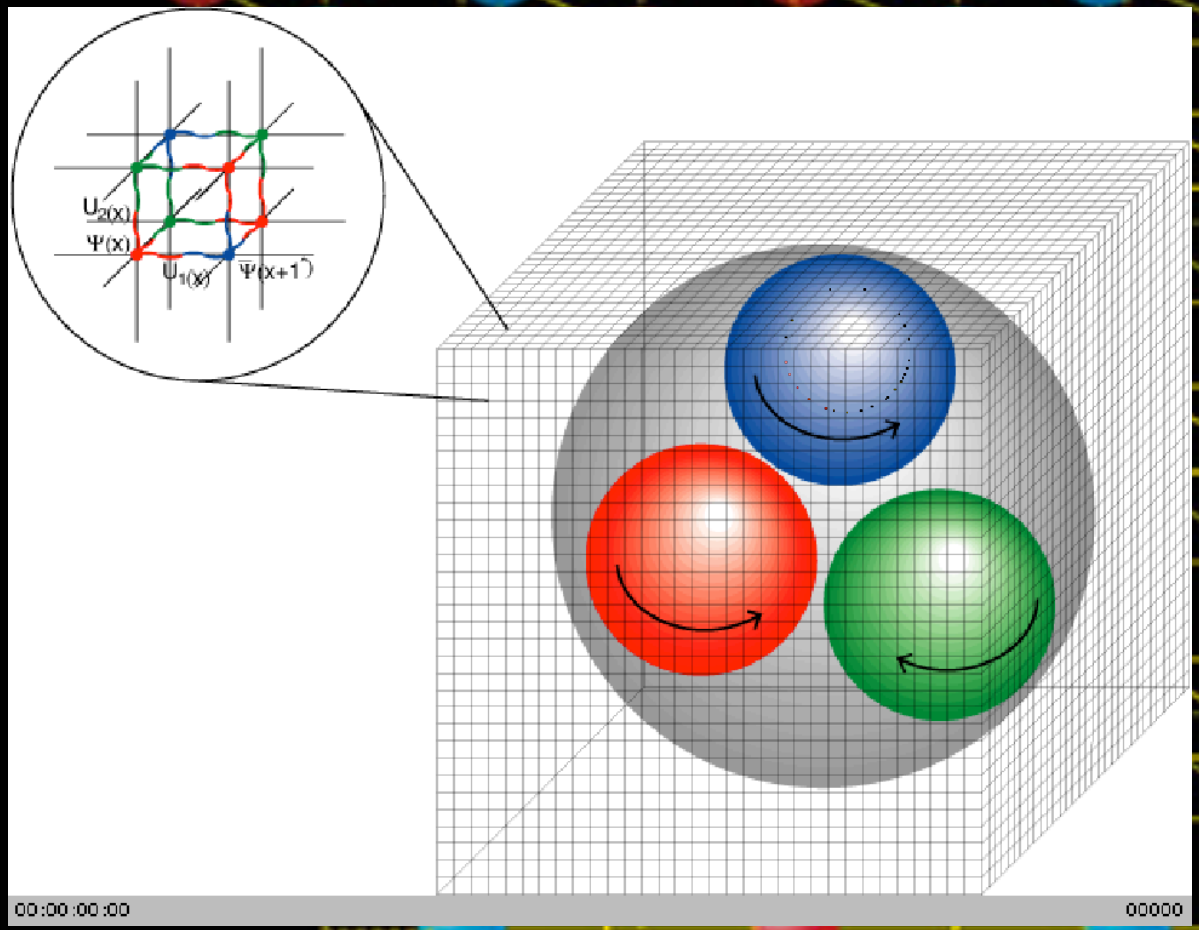

- Periodic or very simple boundary conditions
- SPMD: Identical sublattices per processor

Put PDE's on lattice

$$D\psi(x) = [m + \sum_{\mu} \gamma_\mu(\partial_\mu + iA_\mu(x))]\psi(x)$$

$$D\psi(x) = [m - \frac{1}{2a}\sum_{\pm\mu}(1 \pm \gamma_\mu)e^{\pm iA_\mu(x)}]\psi(x \pm \mu)$$

Improvements in next few years?

From a USQCD proposal 2011 (thanks!)

Quantity	CKM element	Present exp. err.	Present lat. err.	2014 lat. err.	2020 lat. err.
f_K/f_π	$\|V_{us}\|$	0.2%	0.6%	0.3%	0.1%
$f_+^{K\pi}(0)$	$\|V_{us}\|$	0.2%	0.5%	0.2%	0.1%
$D \to \pi\ell\nu$	$\|V_{cd}\|$	2.6%	10.5%	4%	1%
$D \to K\ell\nu$	$\|V_{cs}\|$	1.1%	2.5%	2%	$< 1\%$
$B \to D^{(*)}\ell\nu$	$\|V_{cb}\|$	1.8%	1.8%	0.8%	$< 0.5\%$
$B \to \pi\ell\nu$	$\|V_{ub}\|$	4.1%	8.7%	4%	2%
$B \to \tau\nu$	$\|V_{ub}\|$	21%	6.4%	2%	$< 1\%$
ξ	$\|V_{ts}/V_{td}\|$	1.0%	2.5%	1.5%	$< 1\%$
ΔM_s	$\|V_{ts}V_{tb}\|^2$	0.7%	10.5%	5%	3%

Hardware

Historical Perspective: First "commercial" QCD machine

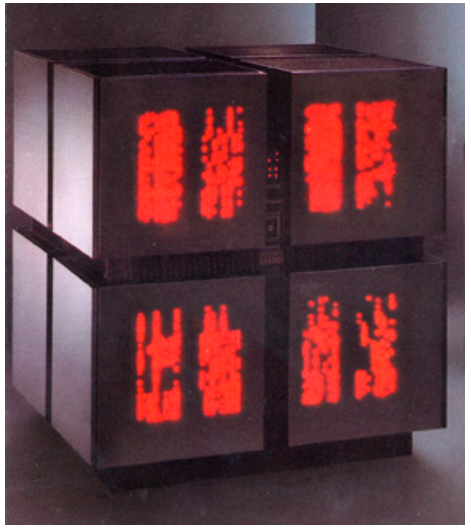

http://www.mission-base.com/tamiko/cm/cm-tshirt.html

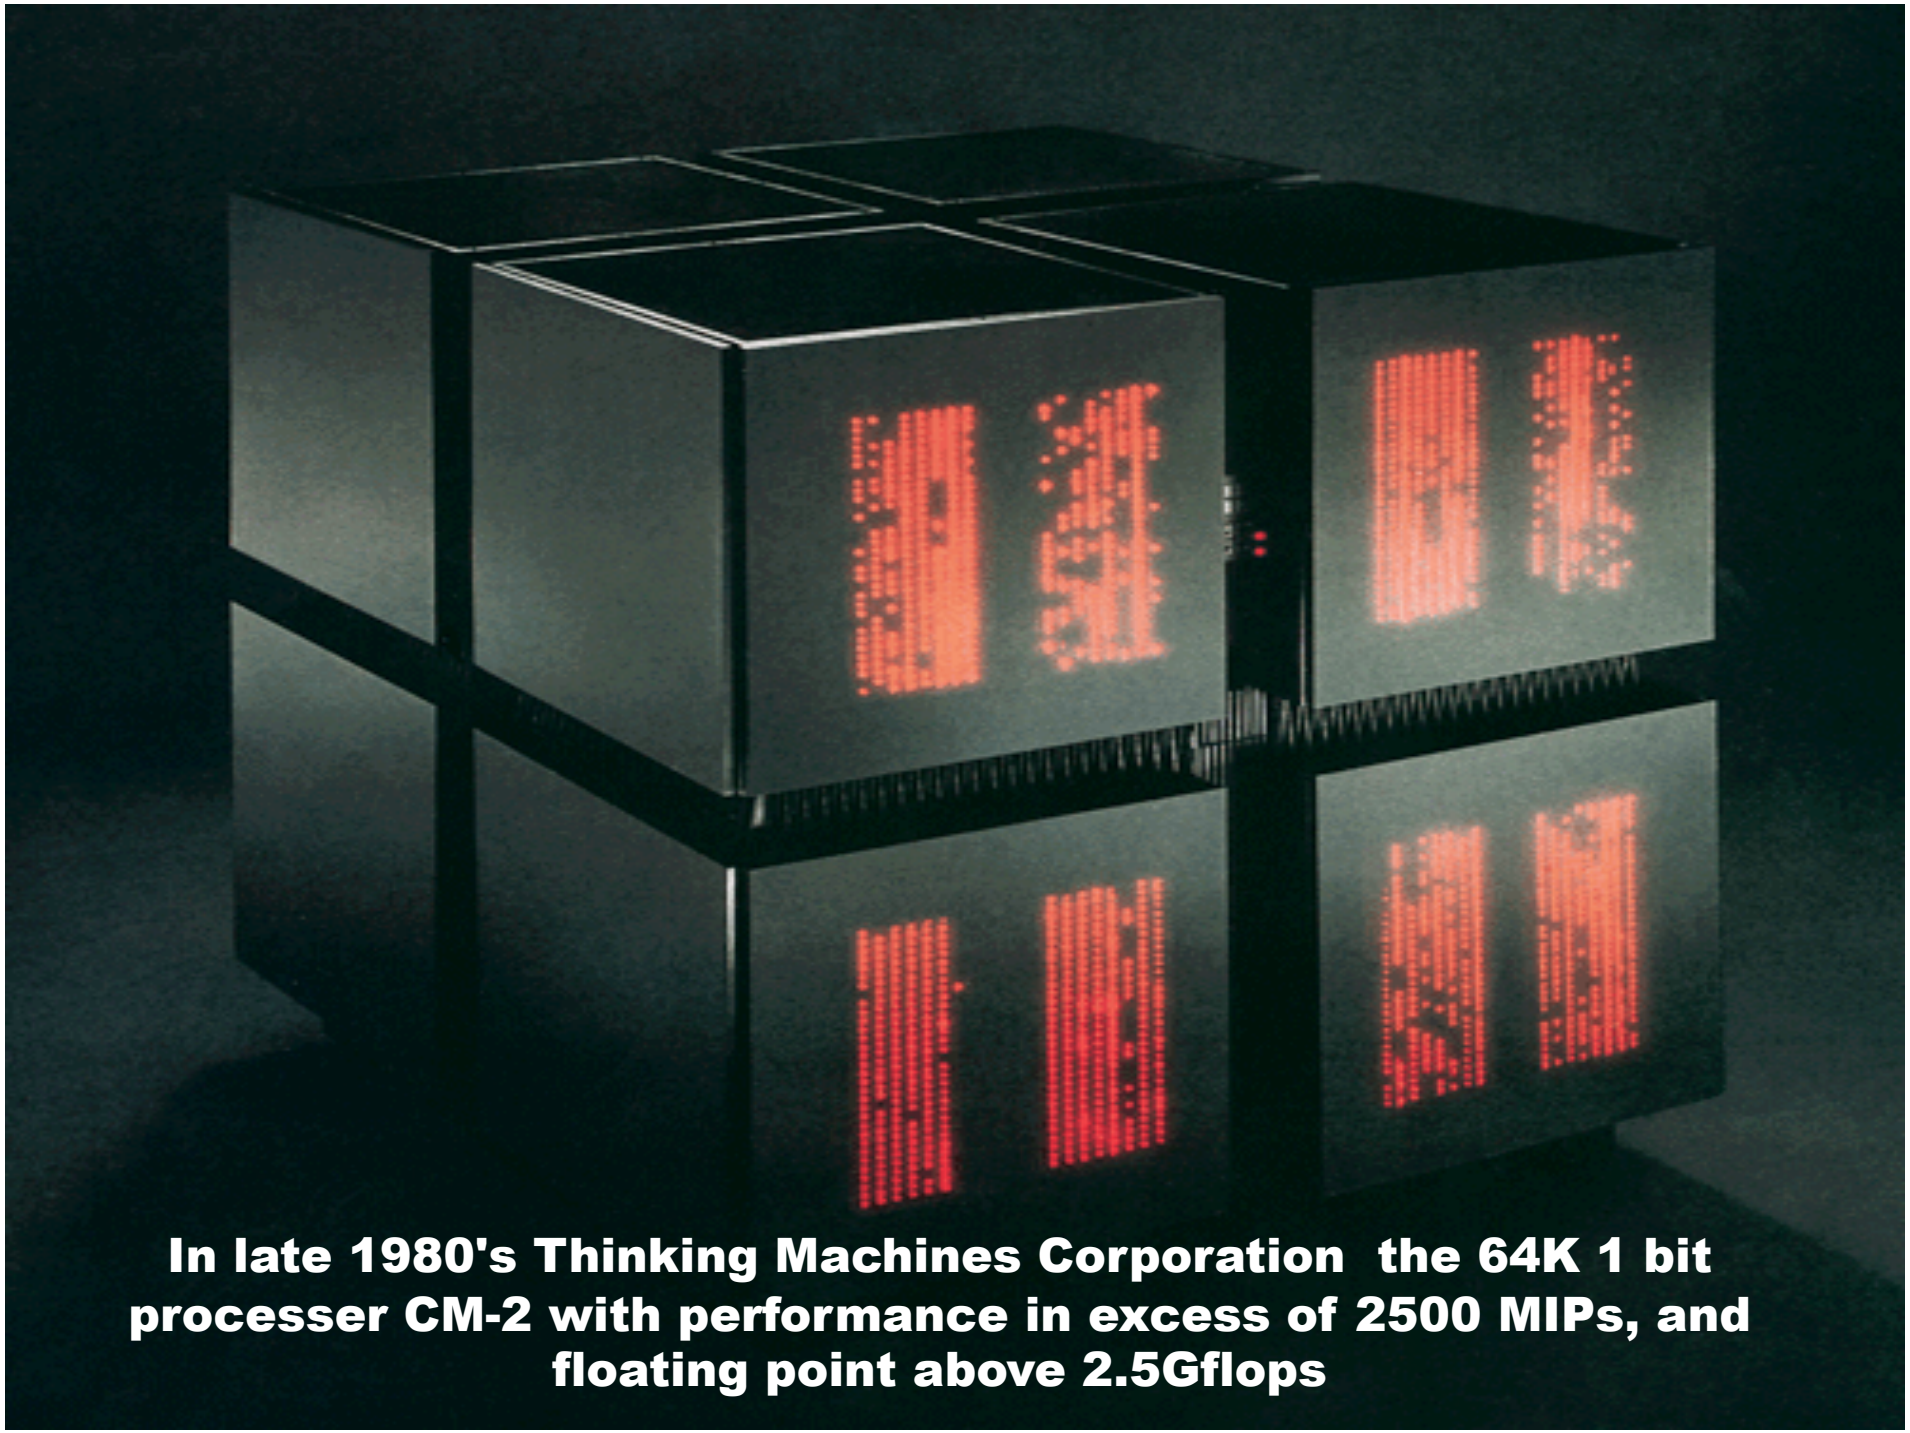

In late 1980's Thinking Machines Corporation the 64K 1 bit processer CM-2 with performance in excess of 2500 MIPs, and floating point above 2.5Gflops

Killed by Beowulf-clusters

BNL:QCDOC
JLab
FNAL
BG/L

First USQCD/Industry Collab: IBM

The BG/Q

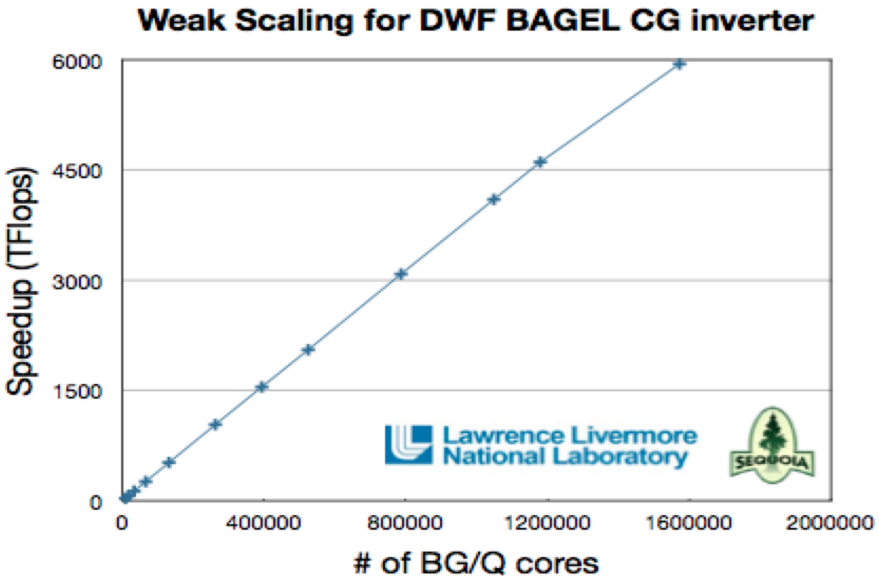

- The Columbia group participated in the design of the BG/Q. Under contract with IBM, they designed and implemented:
  - The interface between the processor core and the level-2 cache, and
  - The look-ahead algorithms used to prefetch data from level-2 cache and main memory, anticipating misses in the level-1 cache.
- Almost perfect weak scaling, 6 PF sustained, achieved on 96 racks of Sequoia (Columbia UK collaborator Peter Boyle).
- Our SciDAC-3 effort is working to port our libraries and codes. (Heechang Na poster.)

Indeed we are getting spectacular Computing Power
3 Multi-Petaflop/s Machines on Top 500

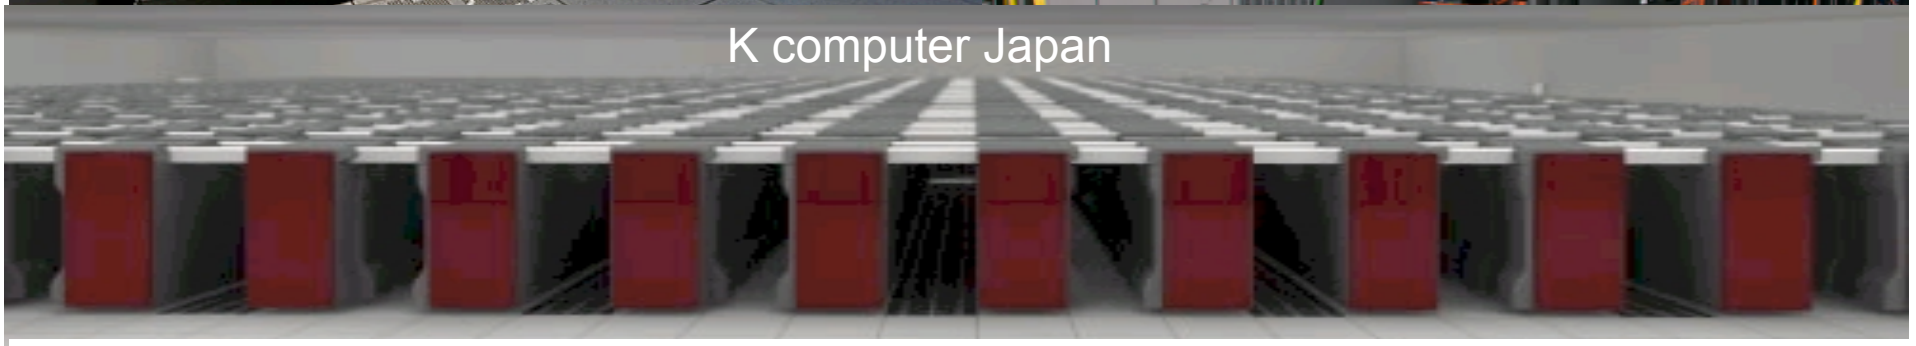

TOP 500: Cray Titan (United States, November 2012 - present)
IBM Sequoia Blue Gene/Q (United States, June 2012 – November 2012)
Fujitsu K computer (Japan, June 2011 – June 2012)
In 2013+ DOE BG/Q Mira, NSF Cray/Nvidia Blue Waters, Intel/MIC Stampede

Software/Algorithms

Machine with Programmers

http://en.wikipedia.org/wiki/ENIAC#Programming

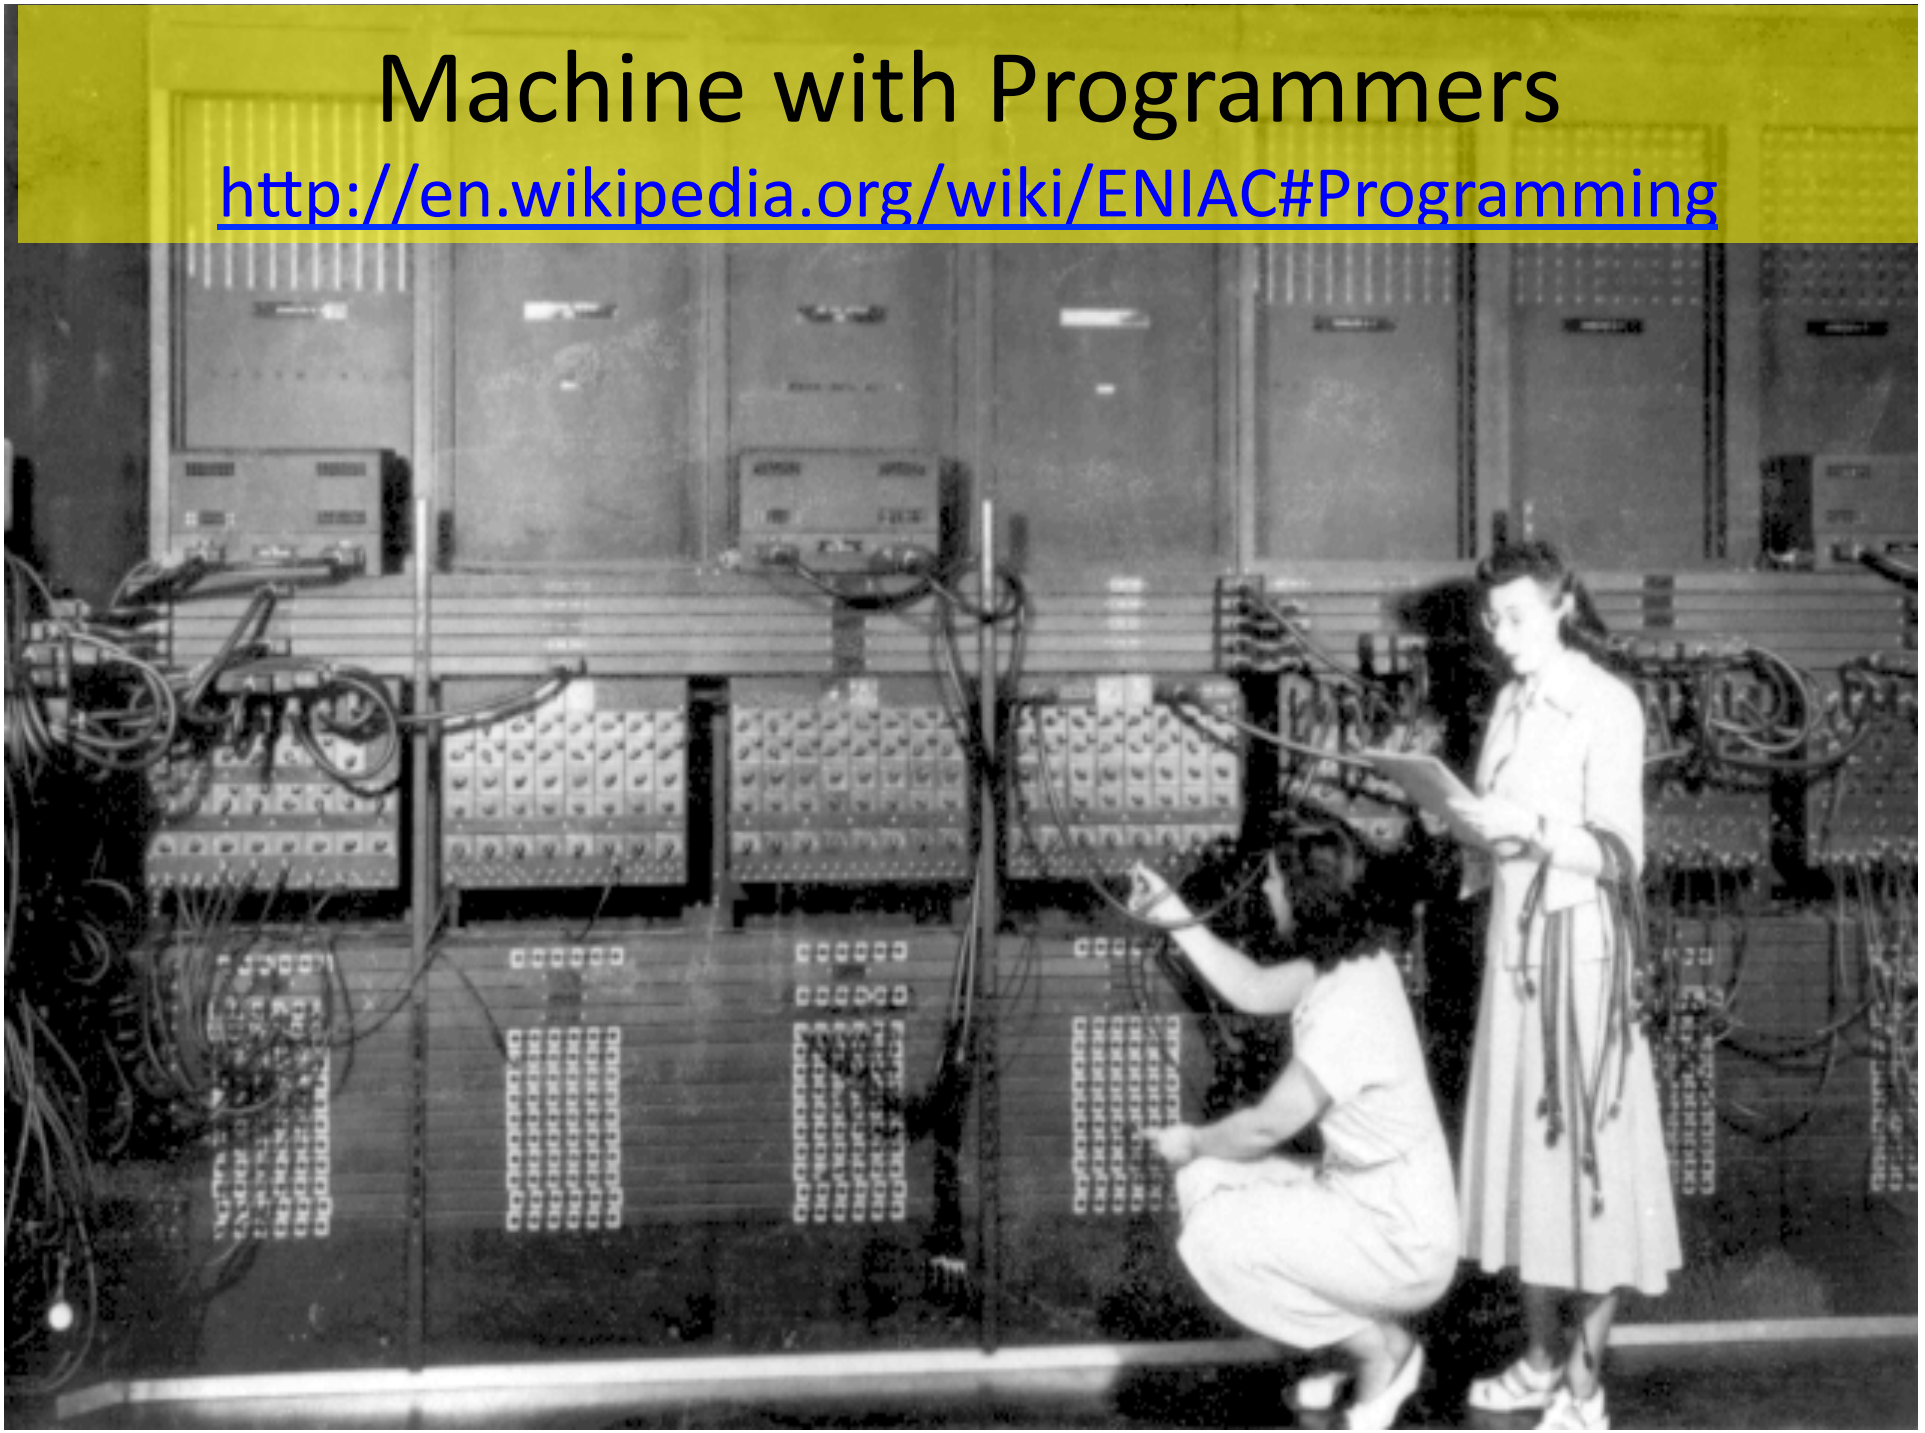

Who am I? Computational Physicists & Software Director for USQCD

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!)

USQCD Software Stack (SciDAC 2002-?)

(supports a "Team of Rivals")

http://usqcd.jlab.org/usqcd-software/

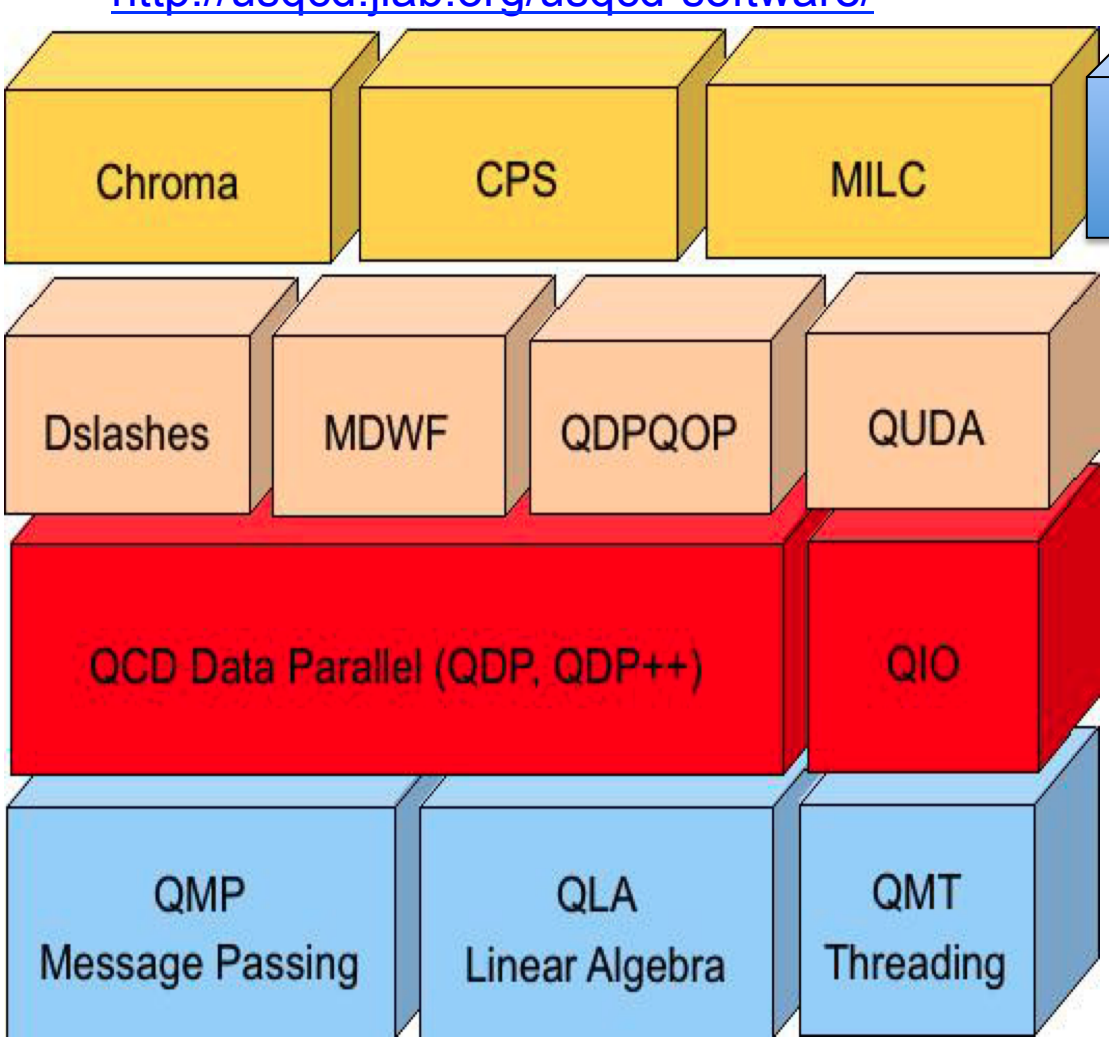

856 files

CPS = 1749 files

MILC = 2300 files

Algorithms

QLA/perl = 23000 files

QUDA/python = 221 files

Architecture

Highest Priority is moving to 3 new architecture!

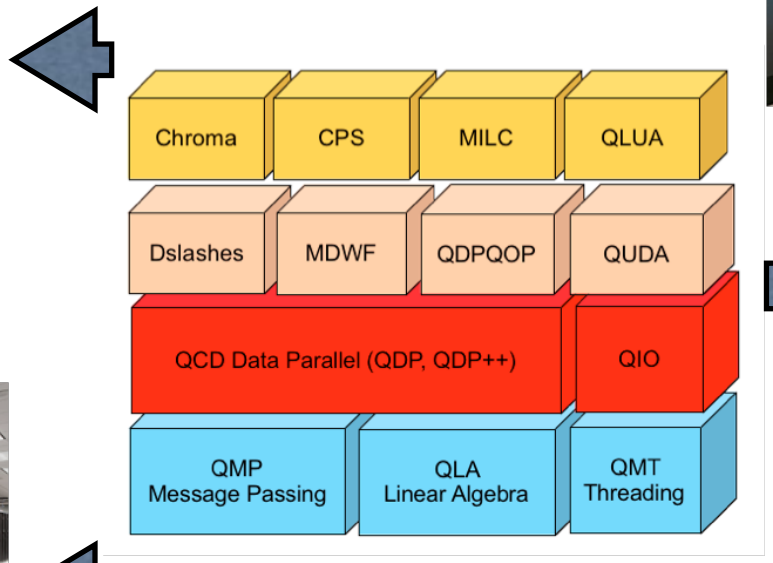

Argonne NATIONAL LABORATORY

IBM

OAK RIDGE National Laboratory

NVIDIA

TACC

intel

(May you live in Interesting Times!)

NCSA

CRAY X1

Algorithm must respond both to multi-scale physics and Heterogeneous Architrave

TWO CASE STUDIES

(1) Multigrid &
Domain Decomposition

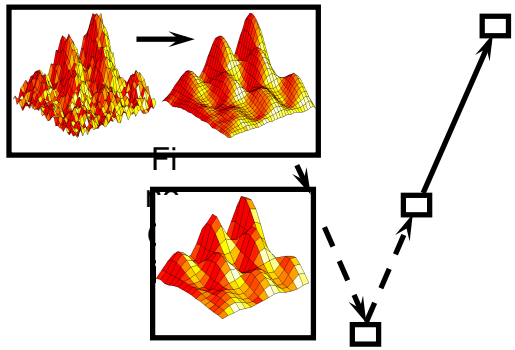

(2) GPU &
Lossless Data Compression

ALGORITHM ONE: MG

Lattice QCD is in a dramatic Turning Point

We now can resolve (at last) the pion scale of QCD!

$$a(lattice) \ll 1/M_{proton} \ll 1/m_\pi \ll L(box)$$

$$0.06 \text{ fermi} \ll 0.2 \text{ fermi} \ll 1.4 \text{ fermi} \ll 6.0 \text{ fermi}$$

$$\Longrightarrow \; L = O(100) \text{ or Minimum Lattice Volume } 100^4!$$

Just the beginning of multiscale lattice Field theory

quark masses: (udscbt = 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

At higher resolution
QCD Software is no longer so simple after all!

- We need to resolve many scale physics
  - quarks masses:
    (udscbt = 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
  - Origin of Higgs, Dark Matter, String Theory.

Part I: Scaling → MG

The Connection Machine

by

William Daniel Hillis
M.S., B.S., Massachusetts Institute of Technology

The first points out that computer science is missing many of the qualities that make the laws of physics so powerful: locality, symmetry, invariance of scale. This is why

Many fundamental equation have no scale (length or masses). Examples are the Poission, Maxwell, Massless Dirac and QCD equations. These appear to be ideal for naive MG rescaling algorithm. BUT quantum system such as QCD have magical scales emerging from the vacuum that give the proton/neutron (and David Keyes) their mass. This requires re-thinking Multigrid.

Multigrid: Case History in Algorithm Development

* "MG is always the Future": Carleton, JLab *2008*

** "Carleton, the future has arrived!": Rich, Oak Ridge *2013*

- History Lessons (1989-1992) *
  - Cause of early failure
- Modern Era (2008-2013)
  - 5 years to put into production the QCD MG Solver for Wilson-clover
- Future** (2013-2018)
  - Domain Wall & Staggered Solvers, HMC evolution, etc
  - Adaptation to heterogeneous architectures, etc.

PDE/Physics with NO scale!

$$\text{No change when} \quad x \to \lambda x, y \to \lambda y, z \to \lambda z$$

- Laplace:

$$-\nabla^2 \phi(x,y,z) = -[\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}]\phi = e^2 \delta(x)\delta(y)\delta(z)$$

- Solution:

$$\phi \to \lambda^{-1}\phi \implies \phi = \frac{e^2}{\sqrt{x^2+y^2+z^2}} = \frac{e^2}{r} \quad \text{up to a const}$$

- Same for Maxwell:

with c = 1 (space = time) e.g. (x,y,z,t)

MG Scaling for 1-d Laplace

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

- h ➔ 2h Restriction R = $P^{\dagger}$

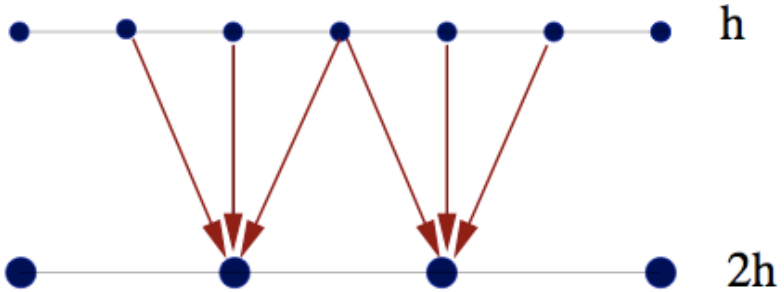

2 h ➔ h Prologation P

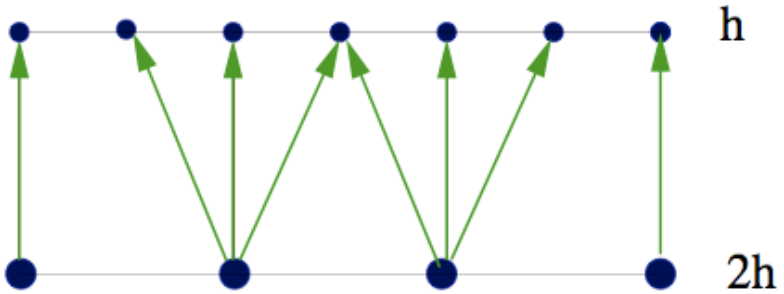

Preserve only scale invariant constant vector is scale invariant. In computer we set h = 1 so only scaling m ➔ 2 m (at m = 0 nothing change for course operator.)

Multi-grid V-Cycle

- L-grid correction scheme huge improvement
- Iterate until exact solve possible
- Interpolate back to fine grid
- *O(N) to O(N log N)* scaling

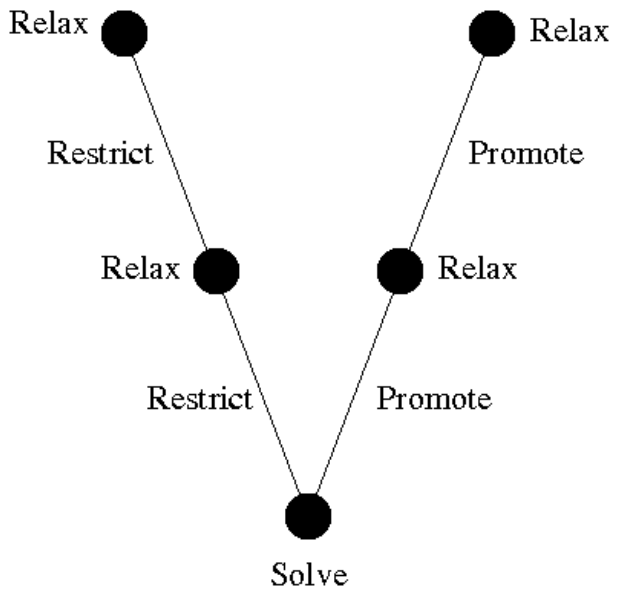

Dirac PDE.

- QCD PDE's (Maxwell+ massless Quark) has NO (length/mass) scale at all!
- Dirac:

$$\sum_{\mu=1,2,3,4} \gamma_\mu \frac{\partial}{\partial x_\mu}\psi + m\psi = \delta(x_1)\delta(x_2)\delta(x_3)\delta(x_4)$$

- Try same idea. Basically work but need to put on lattice and add interaction with field

QCD MG attempts in 1990's:

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

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

Boston: Brower, Edwards, Rebbi & Vicari

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

Hamburg: T. Kalkreuter, G. Mack & M. Speh

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

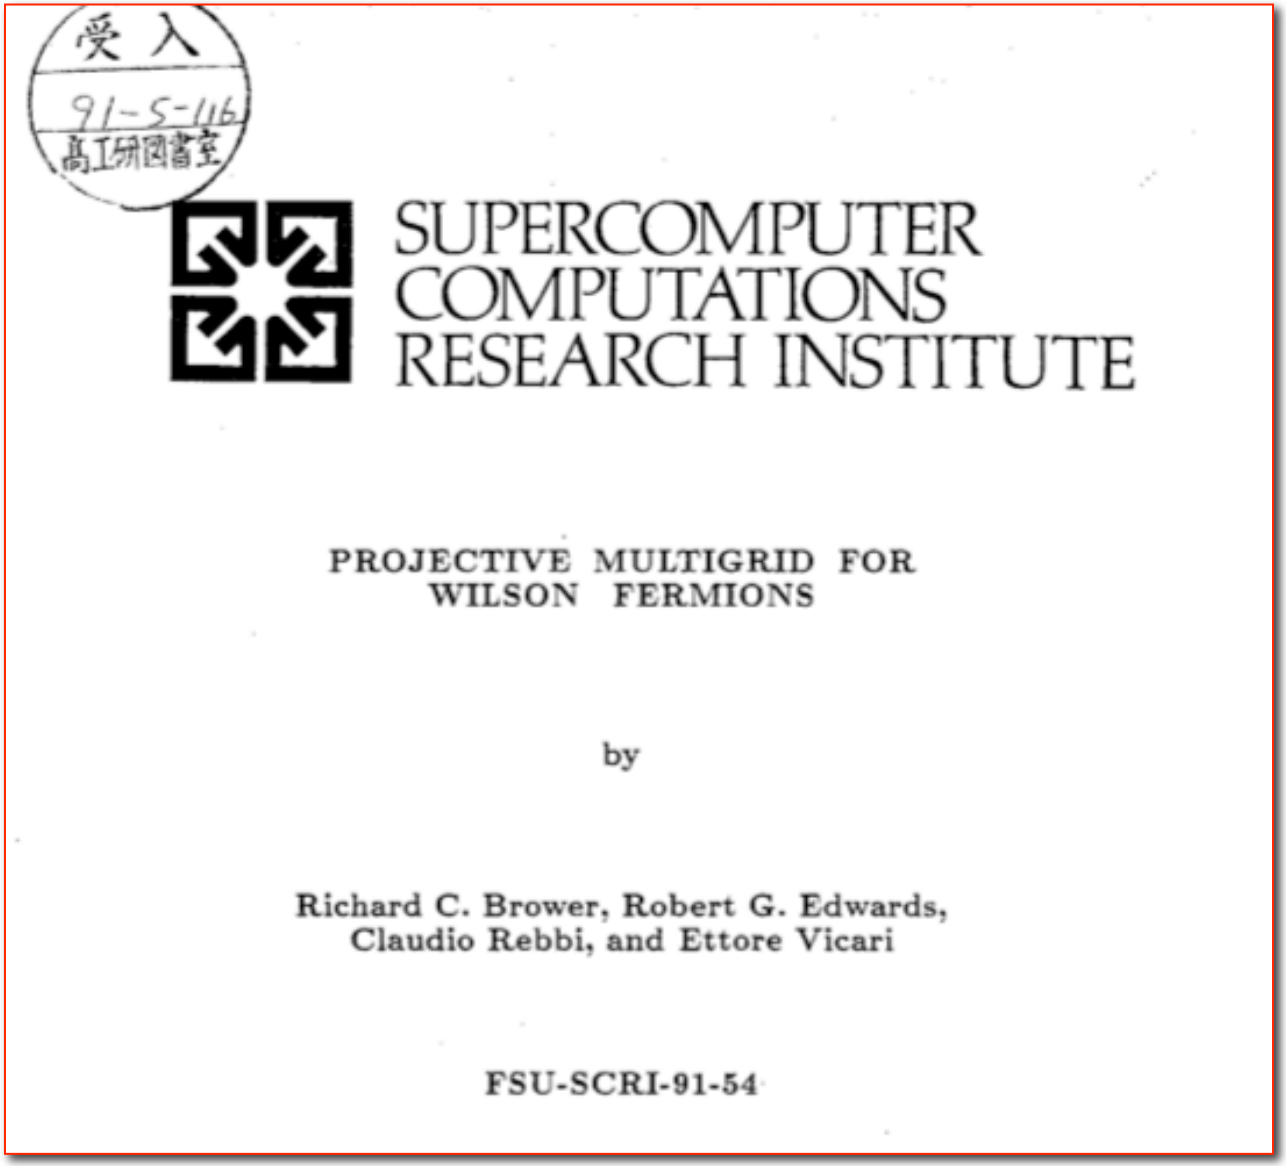

† R. C. Brower, R. Edwards, C.Rebbi,and E. Vicari,
"Projective multigrid for Wilson fermions", Nucl. Phys.B366 (1991) 689
(aka Spectral AMG, Tim Chatier, 2000)

2x2 Blocks for U(1) Dirac

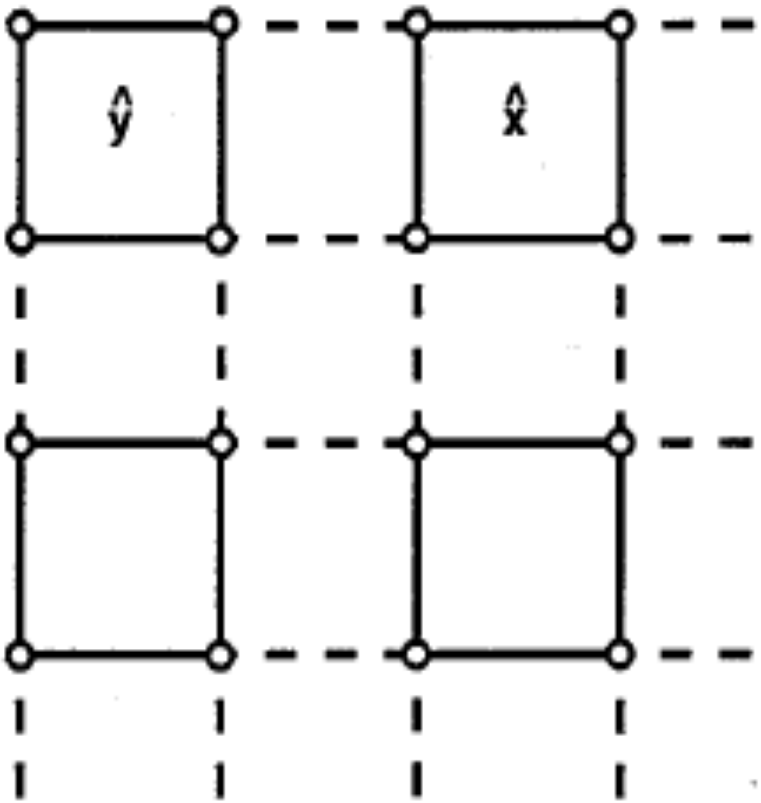

2-d Lattice,
$U_{\square}(x)$ on links $\square(x)$ on sites

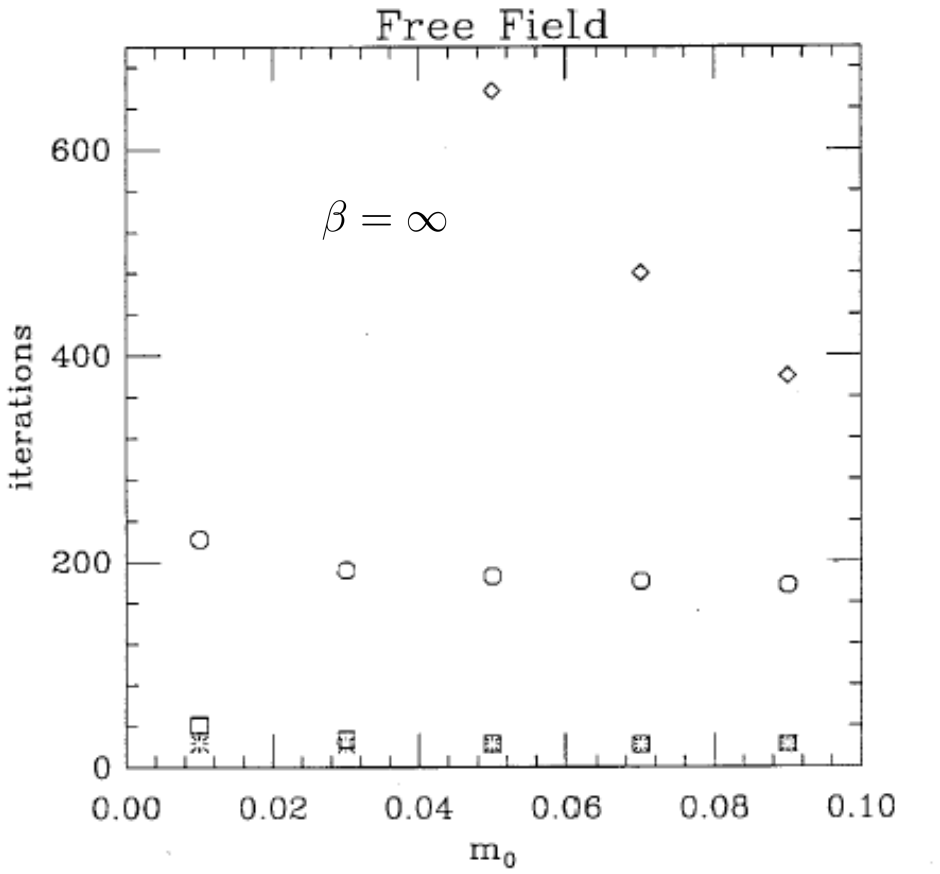

Gauss-Jacobi (Diamond), CG (circle),
V cycle (square), W cycle (star)

Universal critical slowing: τ = F(m l_σ)

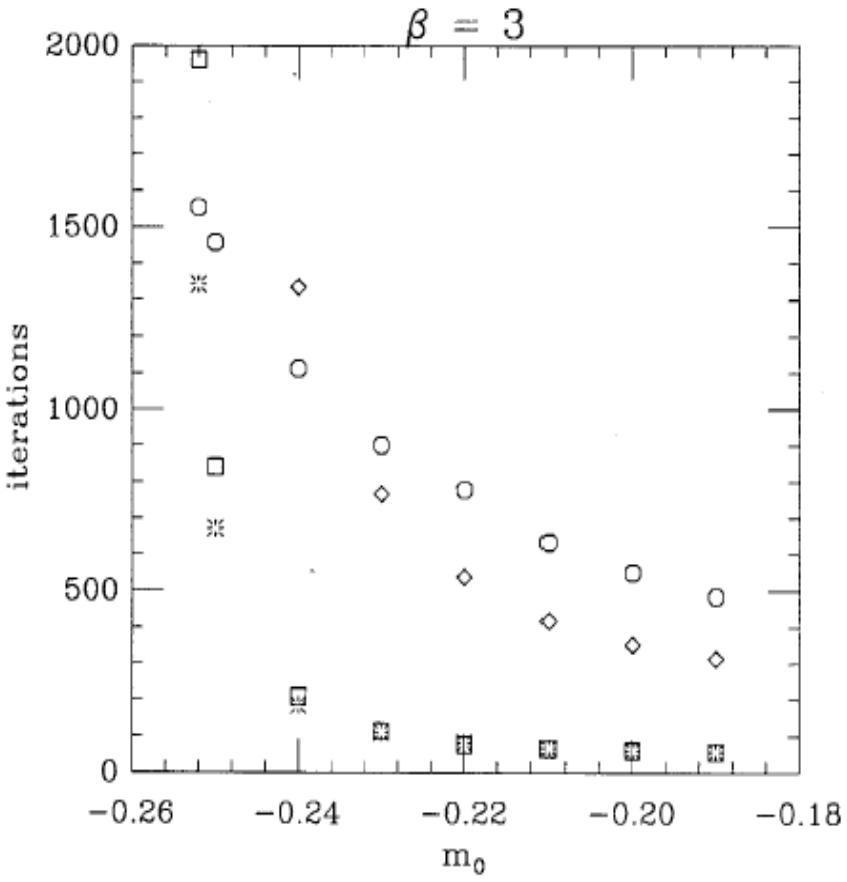

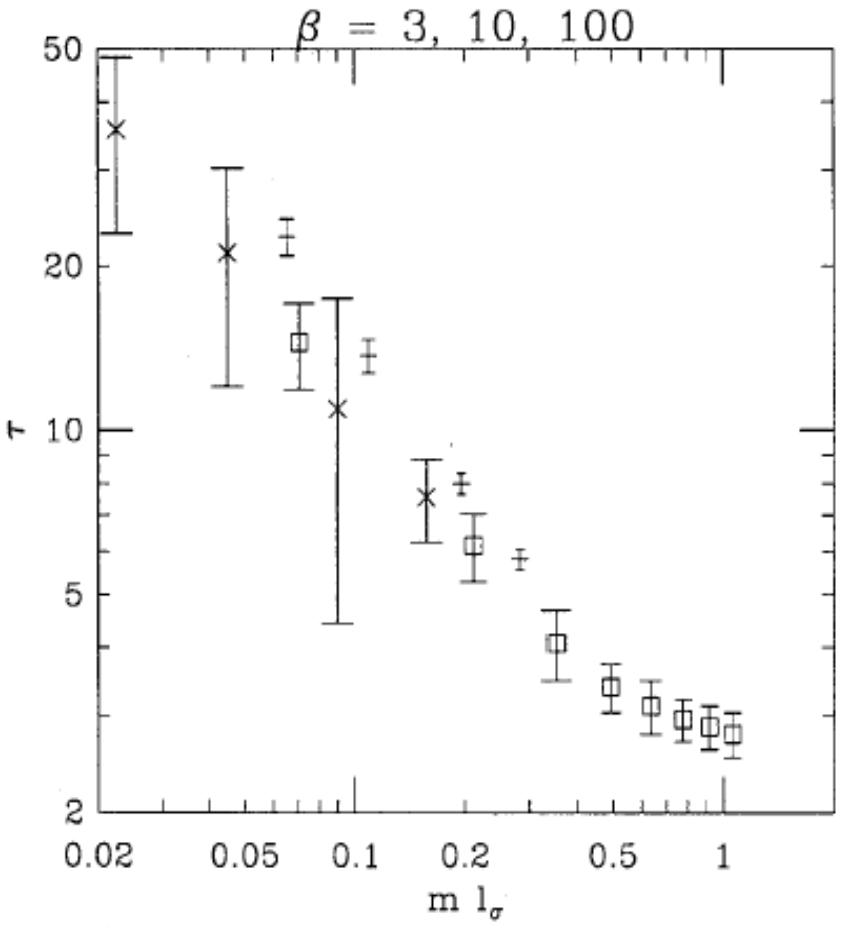

Gauss-Jacobi (Diamond), CG(circle), 3 level (square & star)

τ = 3 (cross) 10(plus) 100(square)

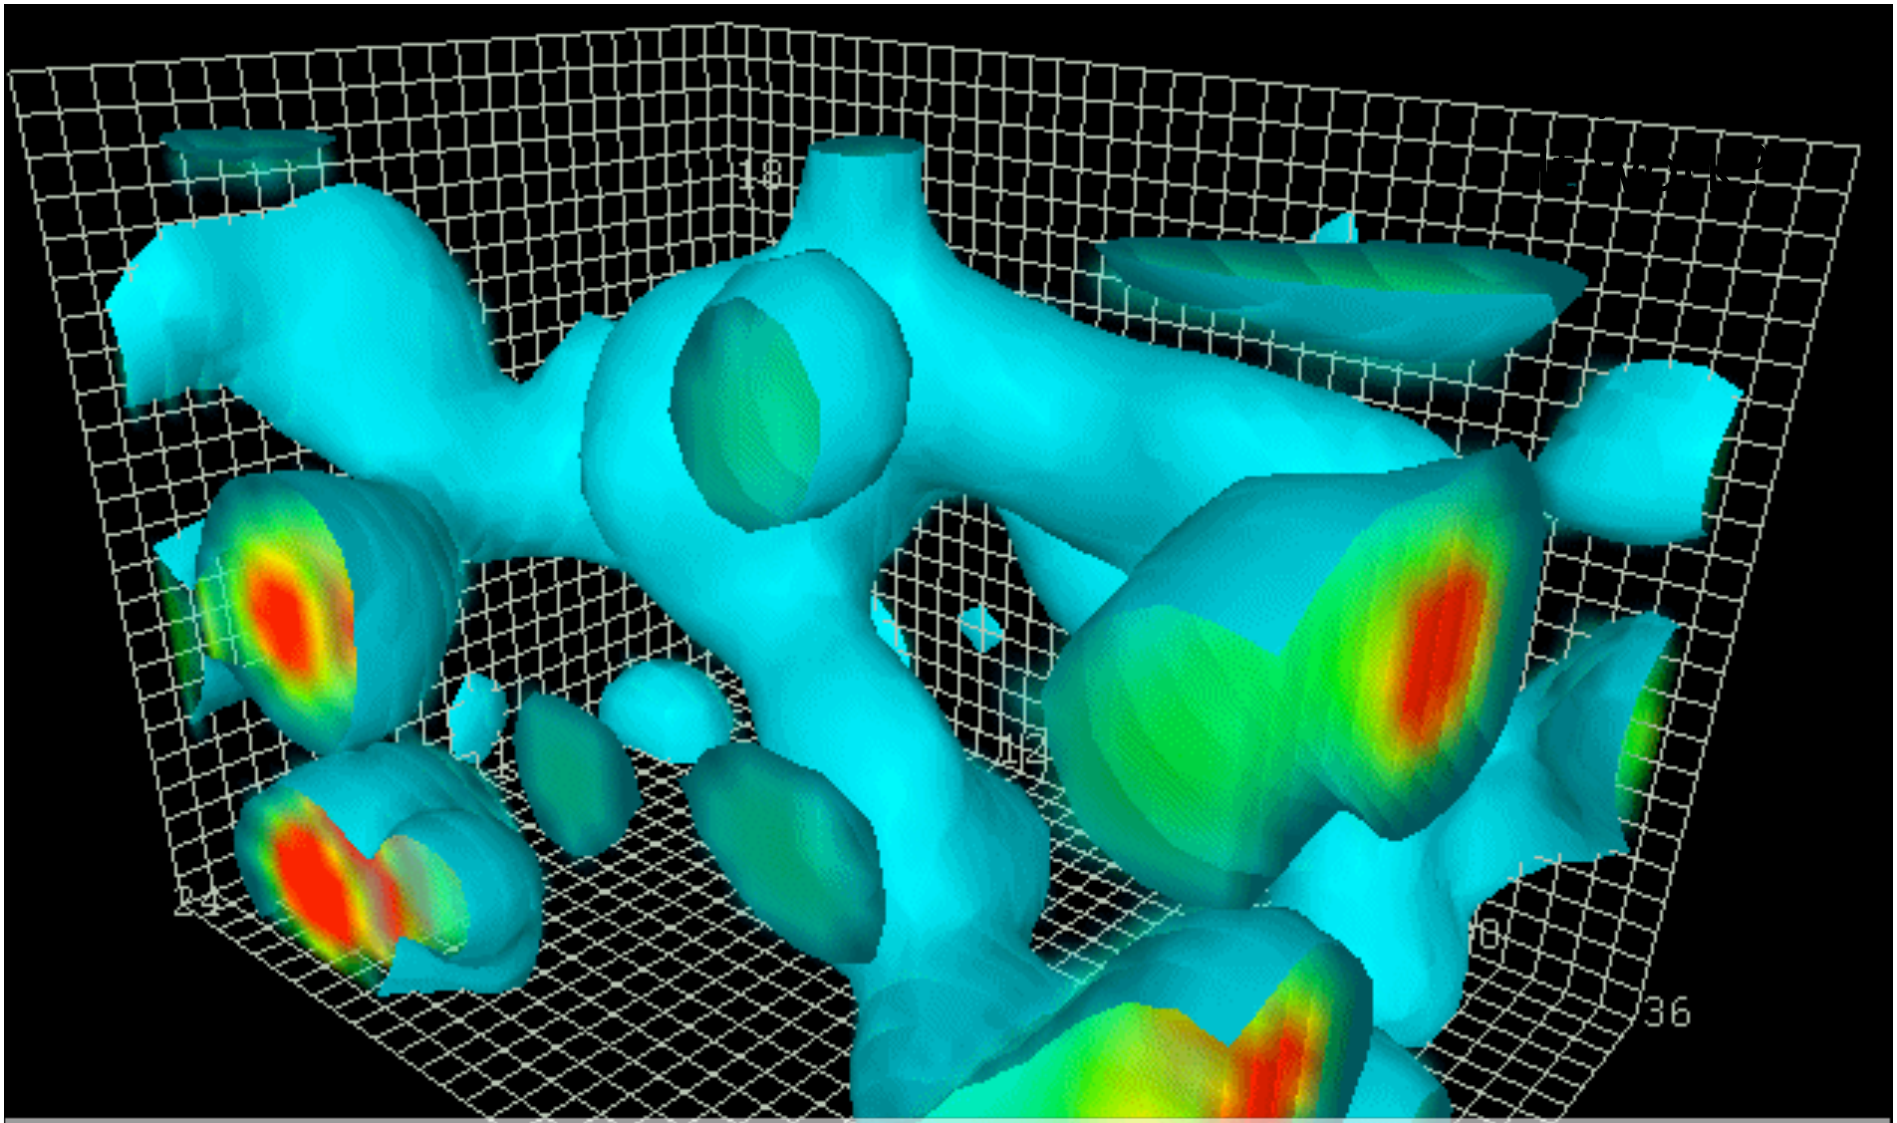

Classical QCD (with zero mass quarks) has no scale. BUT spontaneous Conformal symmetry breaking magically gives the proton mass scale.

QCD Multigrid Applied Math/Physics Collaboration

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

Critical NSF support!

- CU Boulder
  - Tom Manteuffel
  - Steve McCormick
  - Marian Brezina
  - John Ruge
  - James Brannick
  - Christian Ketelsen
  - Scott MacLachlan
- Lawrence Livermore
  - Rob Falgout
- Columbia /KAUST
  - David Keyes
- Boston University
  - Rich Brower
  - Claudio Rebbi
  - Mike Clark
  - James Osborn
  - Saul Cohen
- Penn State
  - James Brannick
  - Ludmil Zikatanov
- Tufts
  - Scott MacLachlan
- Argonne
  - James Osborn

Harvard/ NVIDIA
- Mike Clark

BU/NVIDIA
- Ron Babich

Seattle
- Saul Cohen

Adaptive Smooth Aggregation Algebraic Multigrid

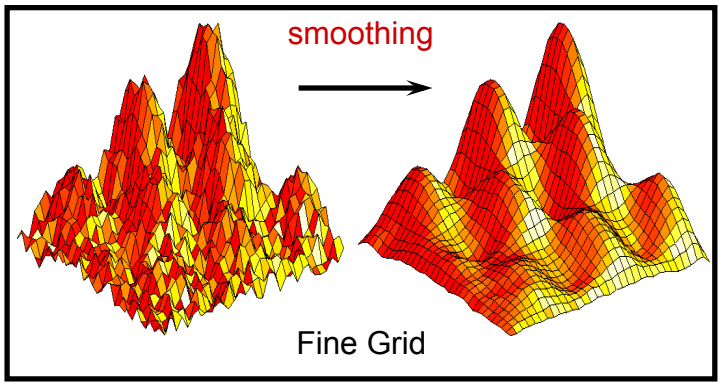

prolongation
(interpolation)

restriction

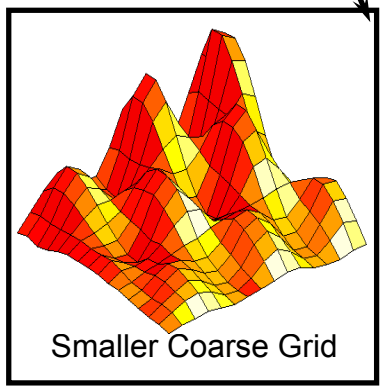

The Multigrid
V-cycle

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

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

Slow convergence of Dirac solver is due to small eigenvalues for vectors in near null subspace: $\mathcal{S}$.

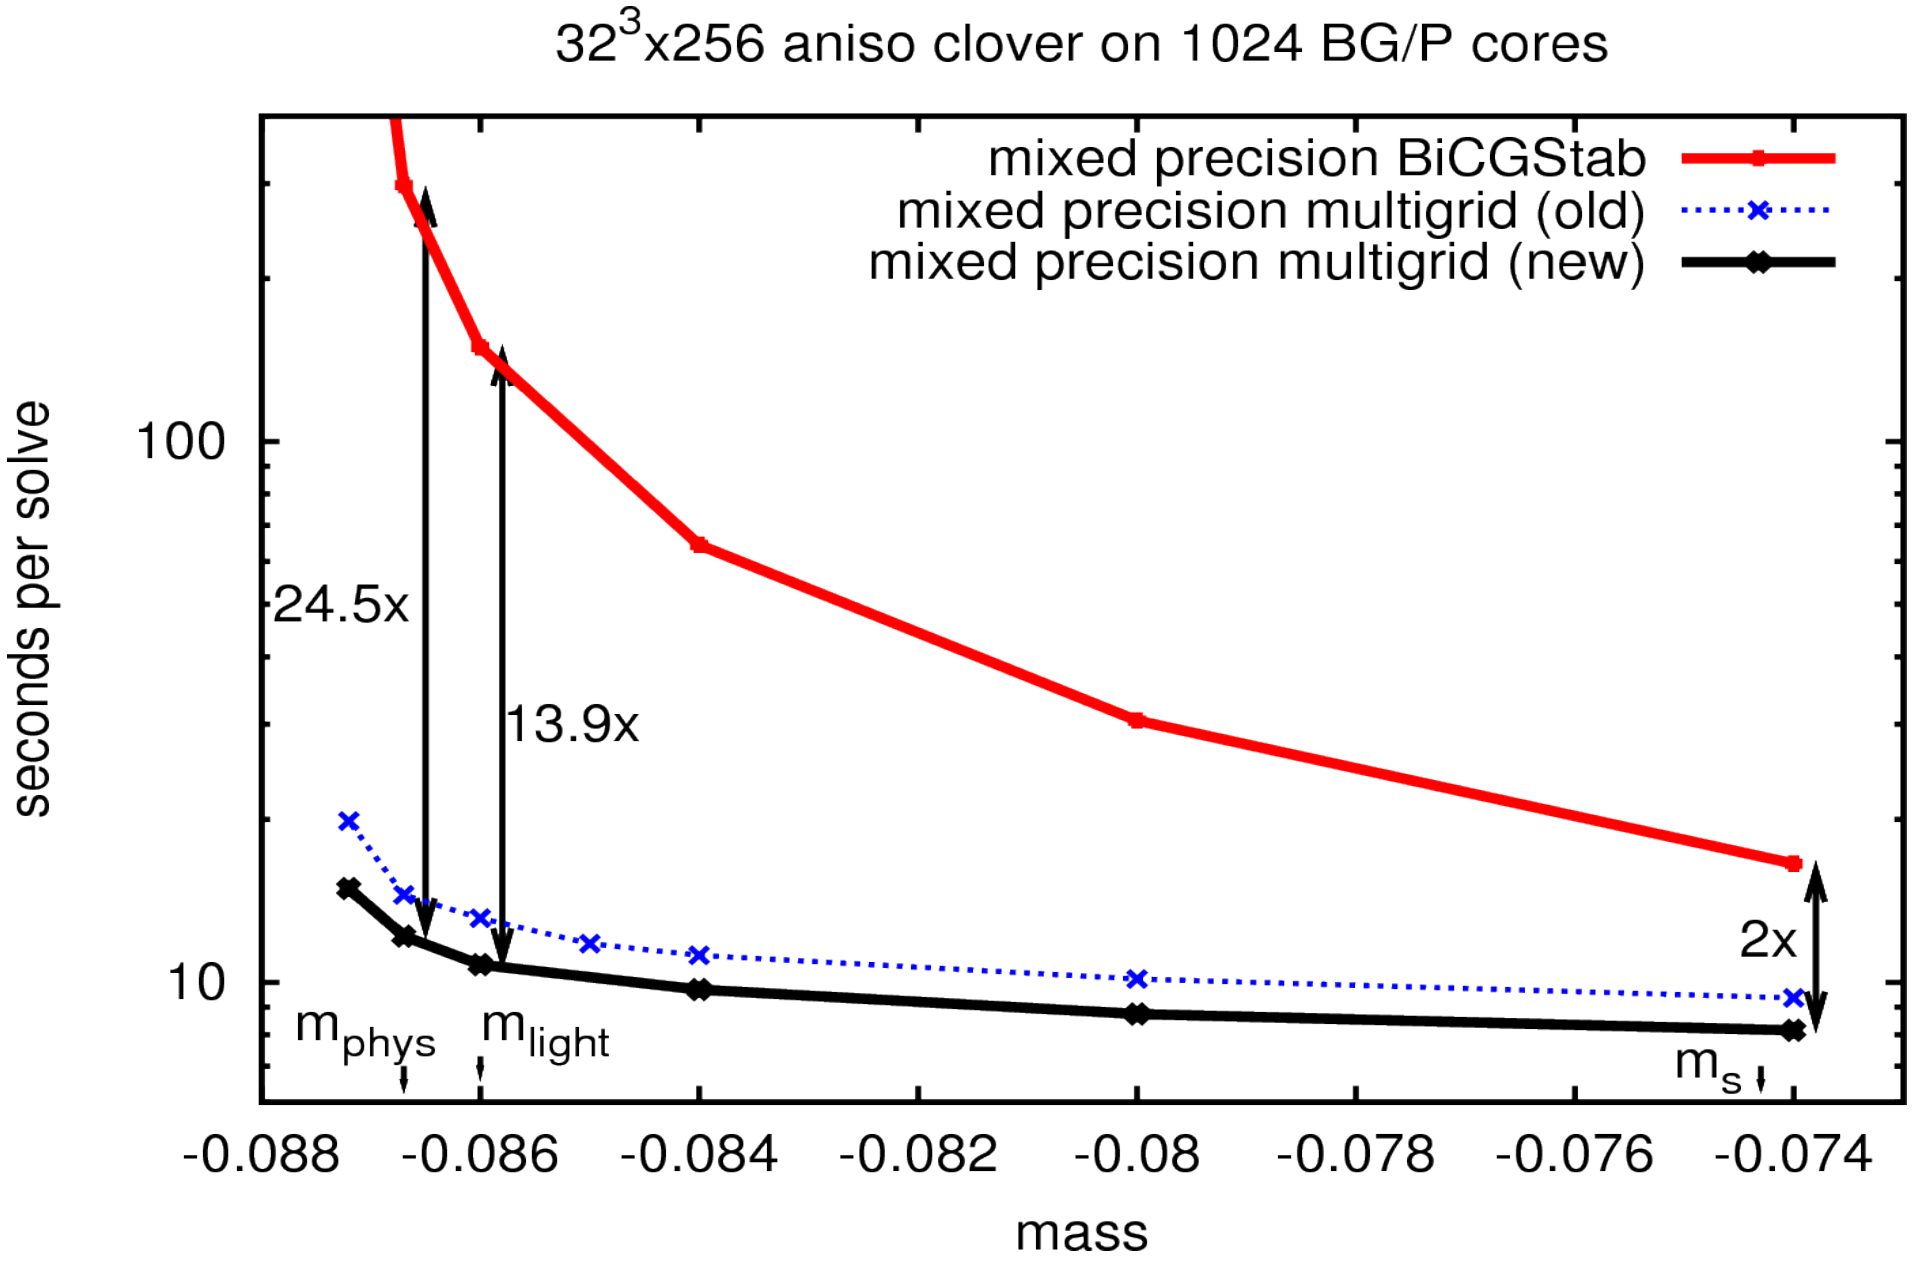

"Adaptive multigrid algorithm for the lattice Wilson-Dirac operator" R. Babich, J. Brannick, R. C. Brower, M. A. Clark, T. Manteuffel, S. McCormick, J. C. Osborn, and C. Rebbi, PRL. (2010).

Good News/Bad News

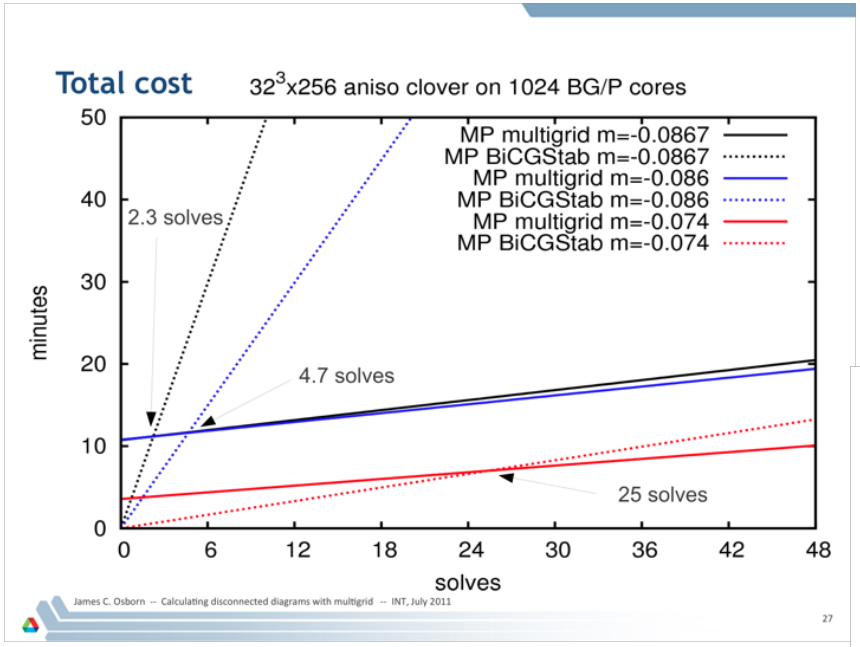

Save MG projectors with lattice ?

Actually MG error is smaller at fixed Residual

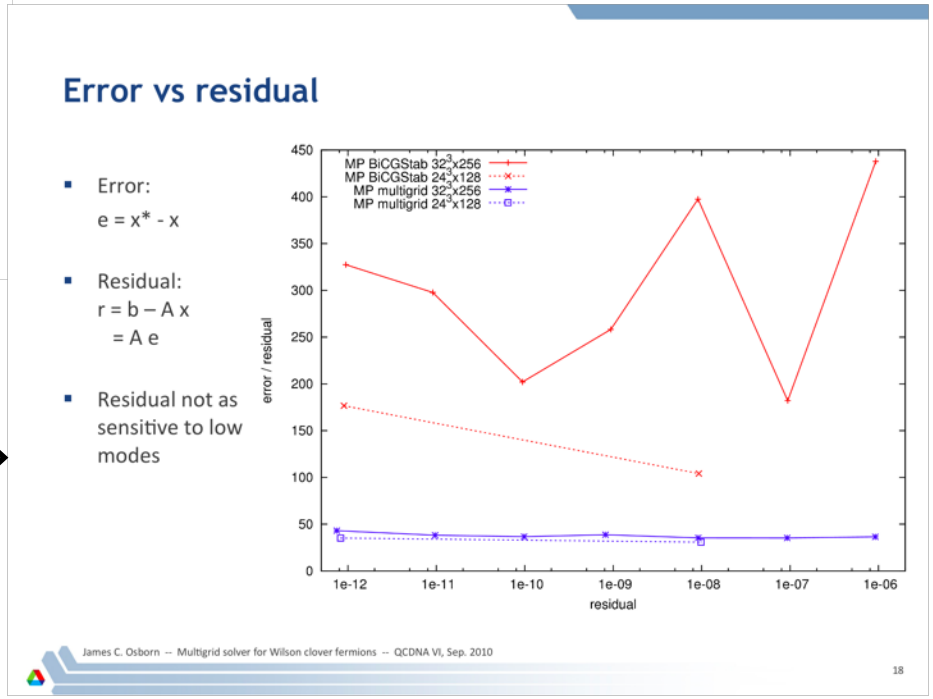

ALGORITHM TWO: GPU

Disruptive many-core Architectures

- 1/4 CM-2

Nvidia FERMI chip

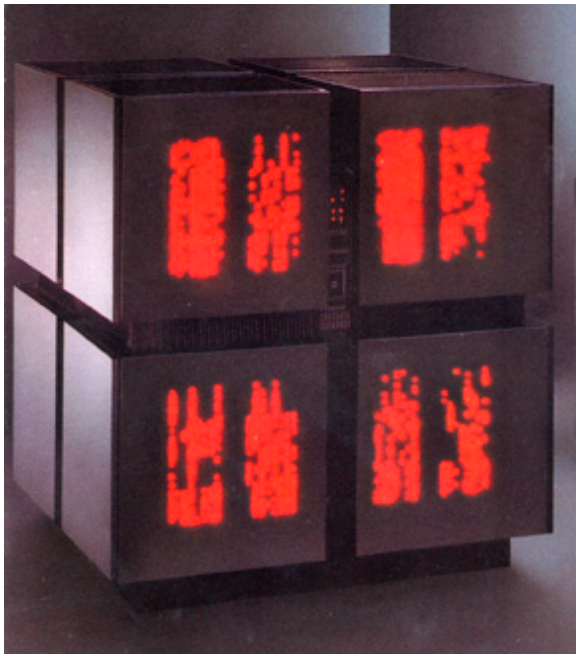

⇒

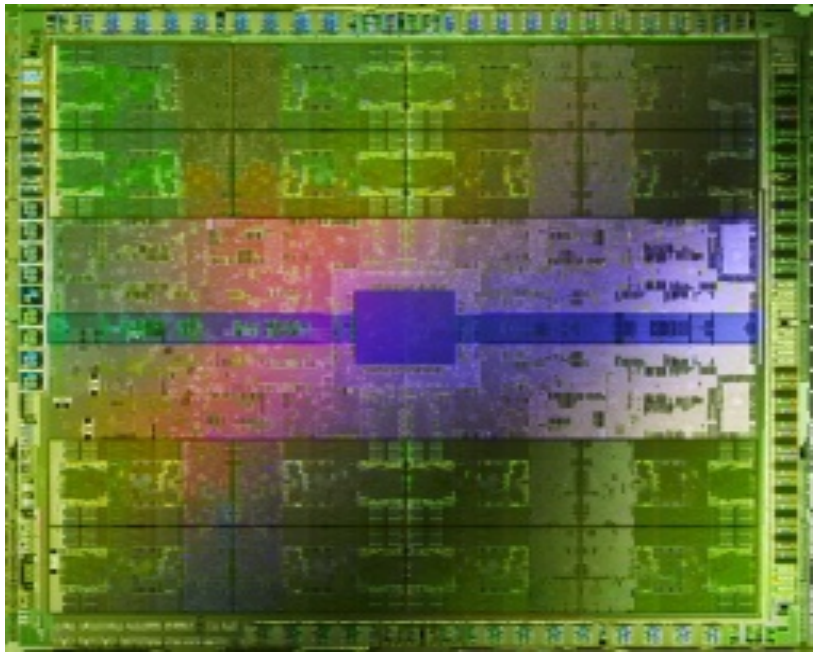

16 K bit serial PE. ⇒ 512 x 32 bit PE = 16 K bits

The pioneers of GPU computing in lattice QCD

2005-2006

Available online at www.sciencedirect.com
ScienceDirect
Computer Physics Communications 177 (2007) 631–639
Computer Physics Communications
www.elsevier.com/locate/cpc

Lattice QCD as a video game

Győző I. Egri [a], Zoltán Fodor [a,b,c,*], Christian Hoelbling [b], Sándor D. Katz [a,b], Dániel Nógrádi [b], Kálmán K. Szabó [b]

[a] *Institute for Theoretical Physics, Eötvös University, Budapest, Hungary*
[b] *Department of Physics, University of Wuppertal, Germany*
[c] *Department of Physics, University of California, San Diego, USA*

Received 2 February 2007; received in revised form 29 May 2007; accepted 7 June 2007

Available online 15 June 2007

Abstract

The speed, bandwidth and cost characteristics of today's PC graphics cards make them an attractive target as general purpose computational platforms. High performance can be achieved also for lattice simulations but the actual implementation can be cumbersome. This paper outlines the architecture and programming model of modern graphics cards for the lattice practitioner with the goal of exploiting these chips for Monte Carlo simulations. Sample code is also given.
© 2007 Elsevier B.V. All rights reserved.

"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 Institute)
- Steve Gottlieb (Indiana University)
- Bálint Joó (Jlab)
- Claudio Rebbi (Boston University)
- Guochun Shi (NCSA -> Google)
- Alexei Strelchenko

  (Cyprus Institute -> FNAL)
- Hyung-Jin Kim (BNL)
- Frank Winter (UoE -> Jlab)

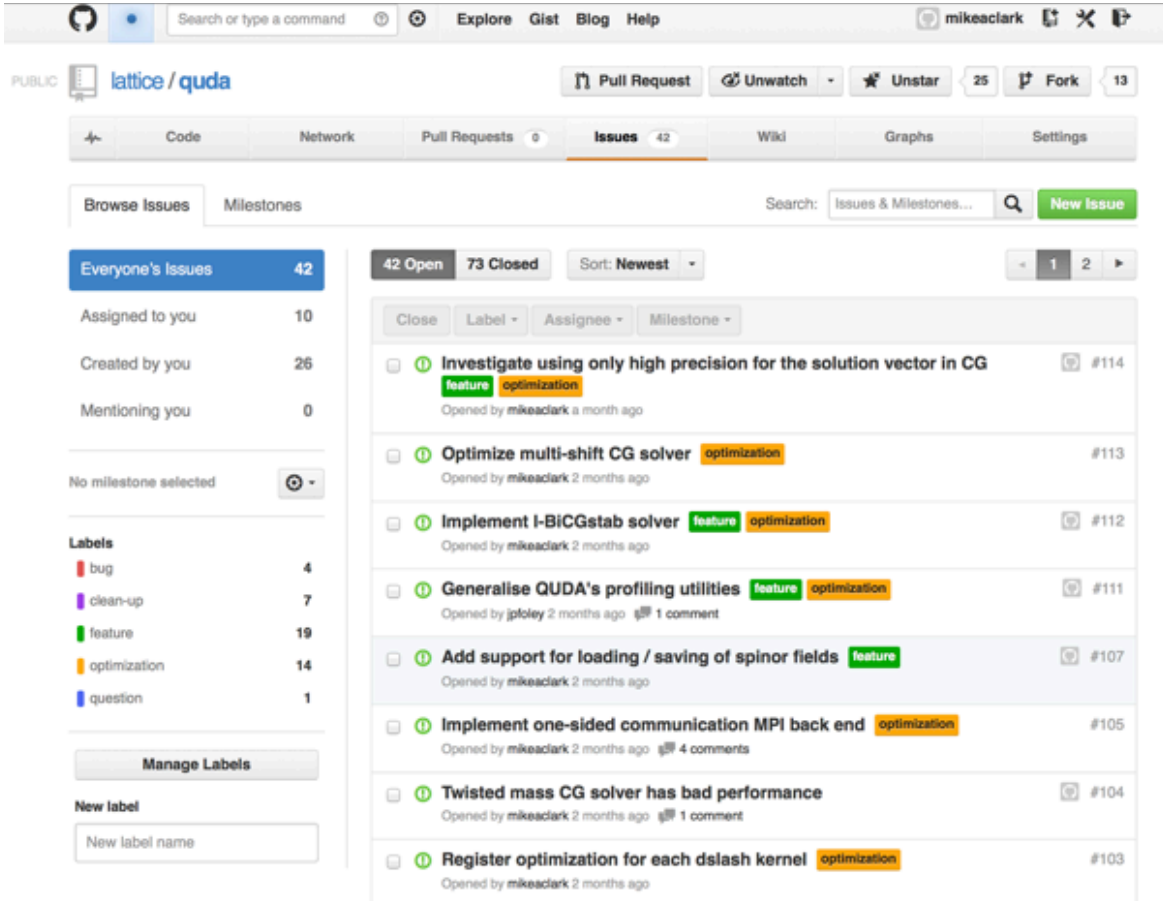

CUDA programming model

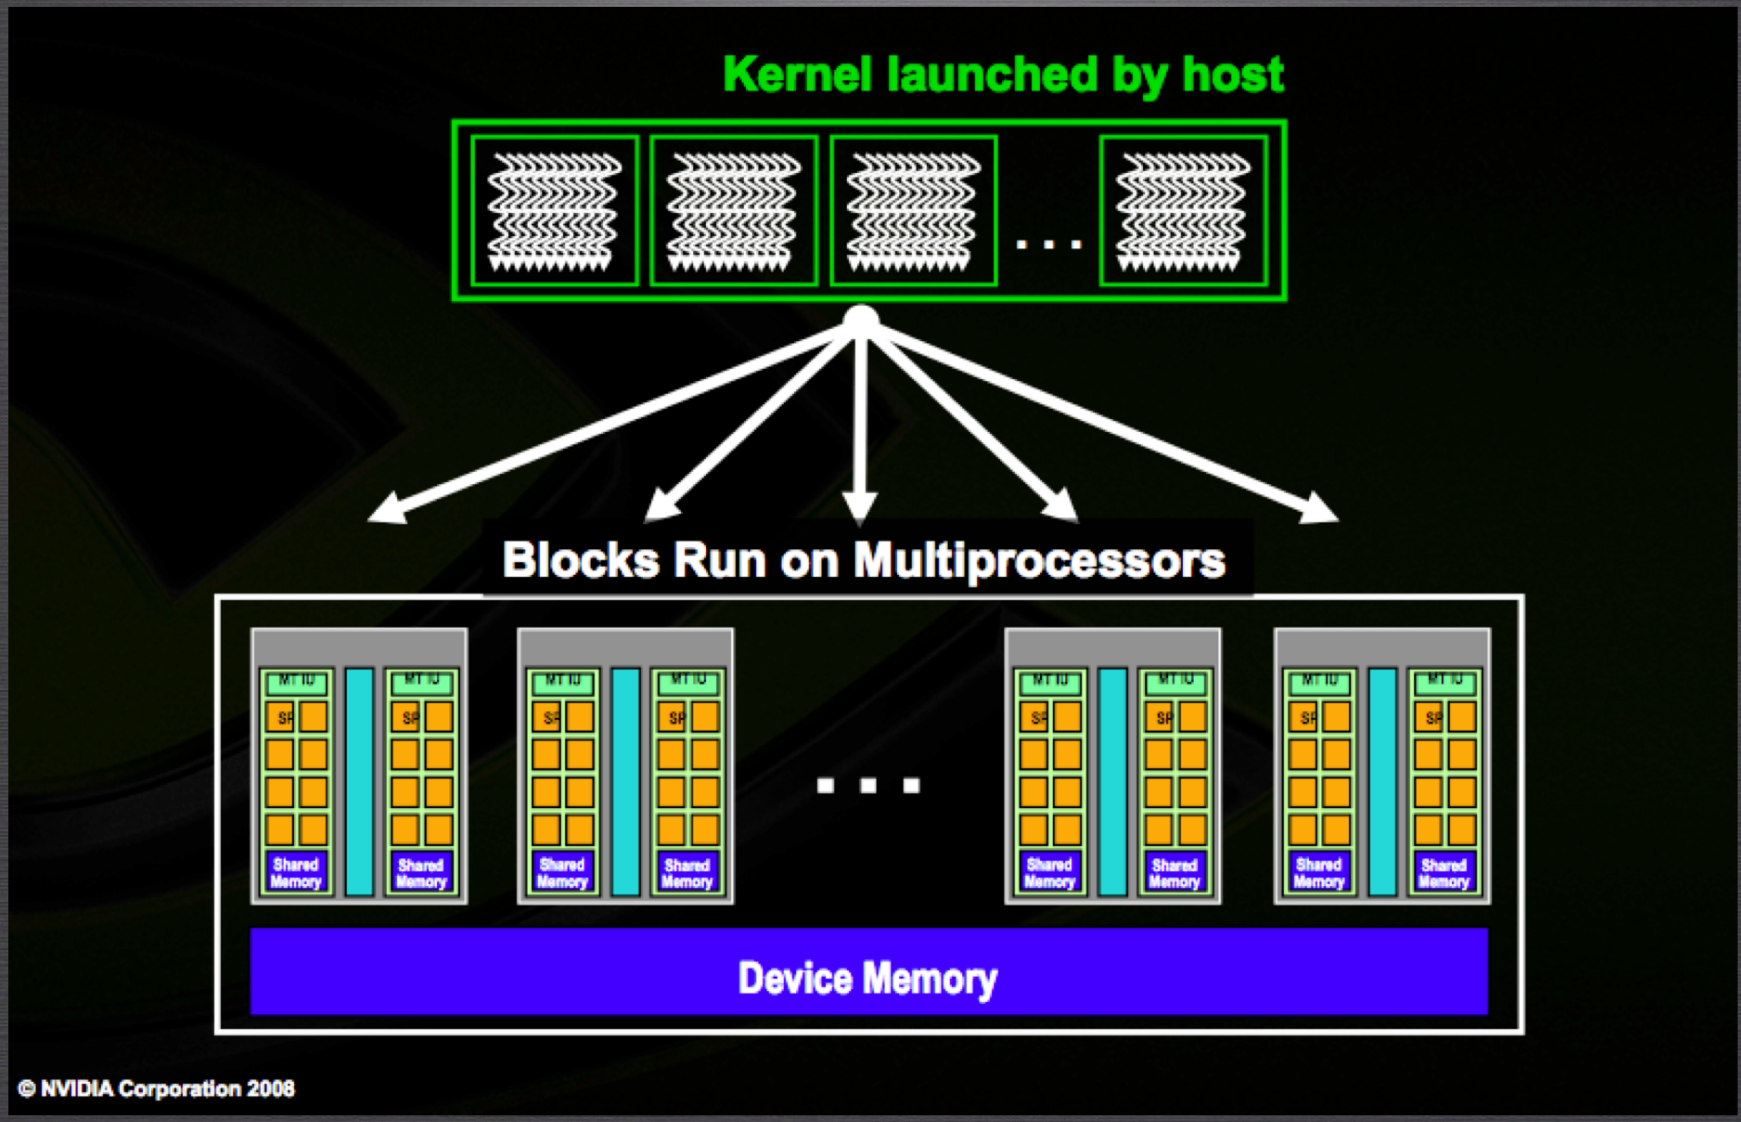

Serial on CPU and Data Parallel on GPU

Wilson matrix-vector: Data movement

- To carry out the Wilson matrix-vector product, per site, we must:
- Read one spinor per direction, for a total of 8 x 24 = 192 floats

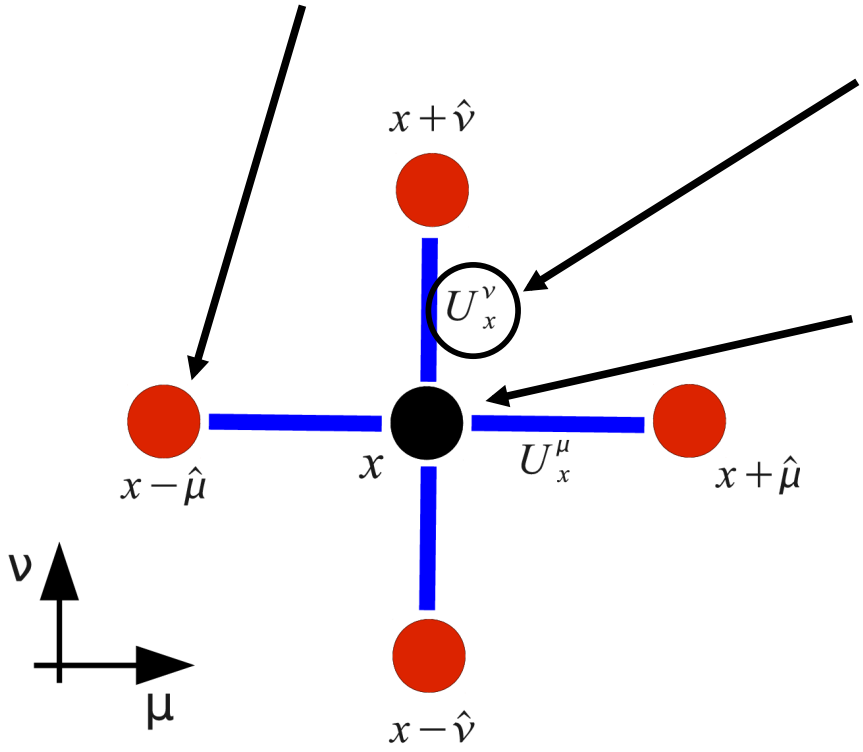

Read one gauge matrix per direction, for a total of
8 x 18 = 144 floats

Write one output spinor, consisting of 24 floats

Altogether, we must transfer 360 floats from memory, or **1440 bytes** in single precision.

Machine balance: Looking back

- In modern architectures, the main bottleneck is often bandwidth to memory, rather than raw floating point throughput.
- It wasn't always like this:

(Cartoon by John McCalpin, author of the STREAM benchmark)

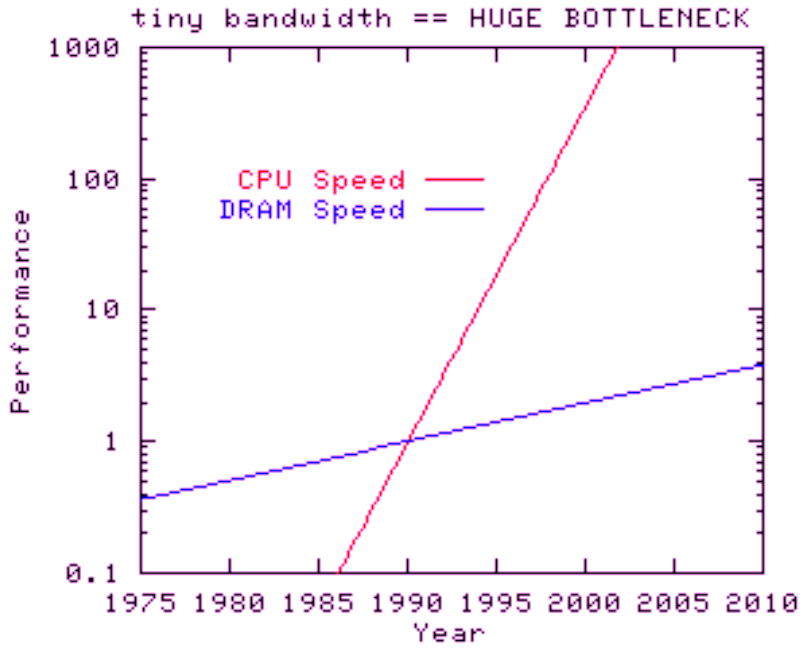

- There was a time when flops were much more precious, relative to bandwidth.

Tricks to reduce memory traffic

- Reconstruct SU(3) matrices from 8 or 12 real numbers on the fly, e.g.,

$$\begin{pmatrix} \mathbf{a} \\ \mathbf{b} \\ \mathbf{c} \end{pmatrix} = \begin{pmatrix} a_1 & a_2 & a_3 \\ b_1 & b_2 & b_3 \\ c_1 & c_2 & c_3 \end{pmatrix} \qquad \mathbf{c} = (\mathbf{a} \times \mathbf{b})^*$$

$$SU(2) : U = a_0\sigma_0 + \vec{a} \cdot \sigma$$

$$a_0^2 + \vec{a} \cdot \vec{a} = 1 \implies S_3 \text{ sphere}$$

Better still SU(3) has 8 parameters on S3 x S5

- Choose a gamma basis with γ_4 diagonal.

$$P^{\pm 4} = \begin{pmatrix} 1 & 0 & \pm 1 & 0 \\ 0 & 1 & 0 & \pm 1 \\ \pm 1 & 0 & 1 & 0 \\ 0 & \pm 1 & 0 & 1 \end{pmatrix} \Rightarrow \begin{cases} P^{+4} = \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \\ P^{-4} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix} \end{cases}$$

similarity transforms on D

- Fix to the temporal gauge (setting gauge links in the *t*-direction to the identity).

Mixed Precision Data Compression (communicate new info only!)

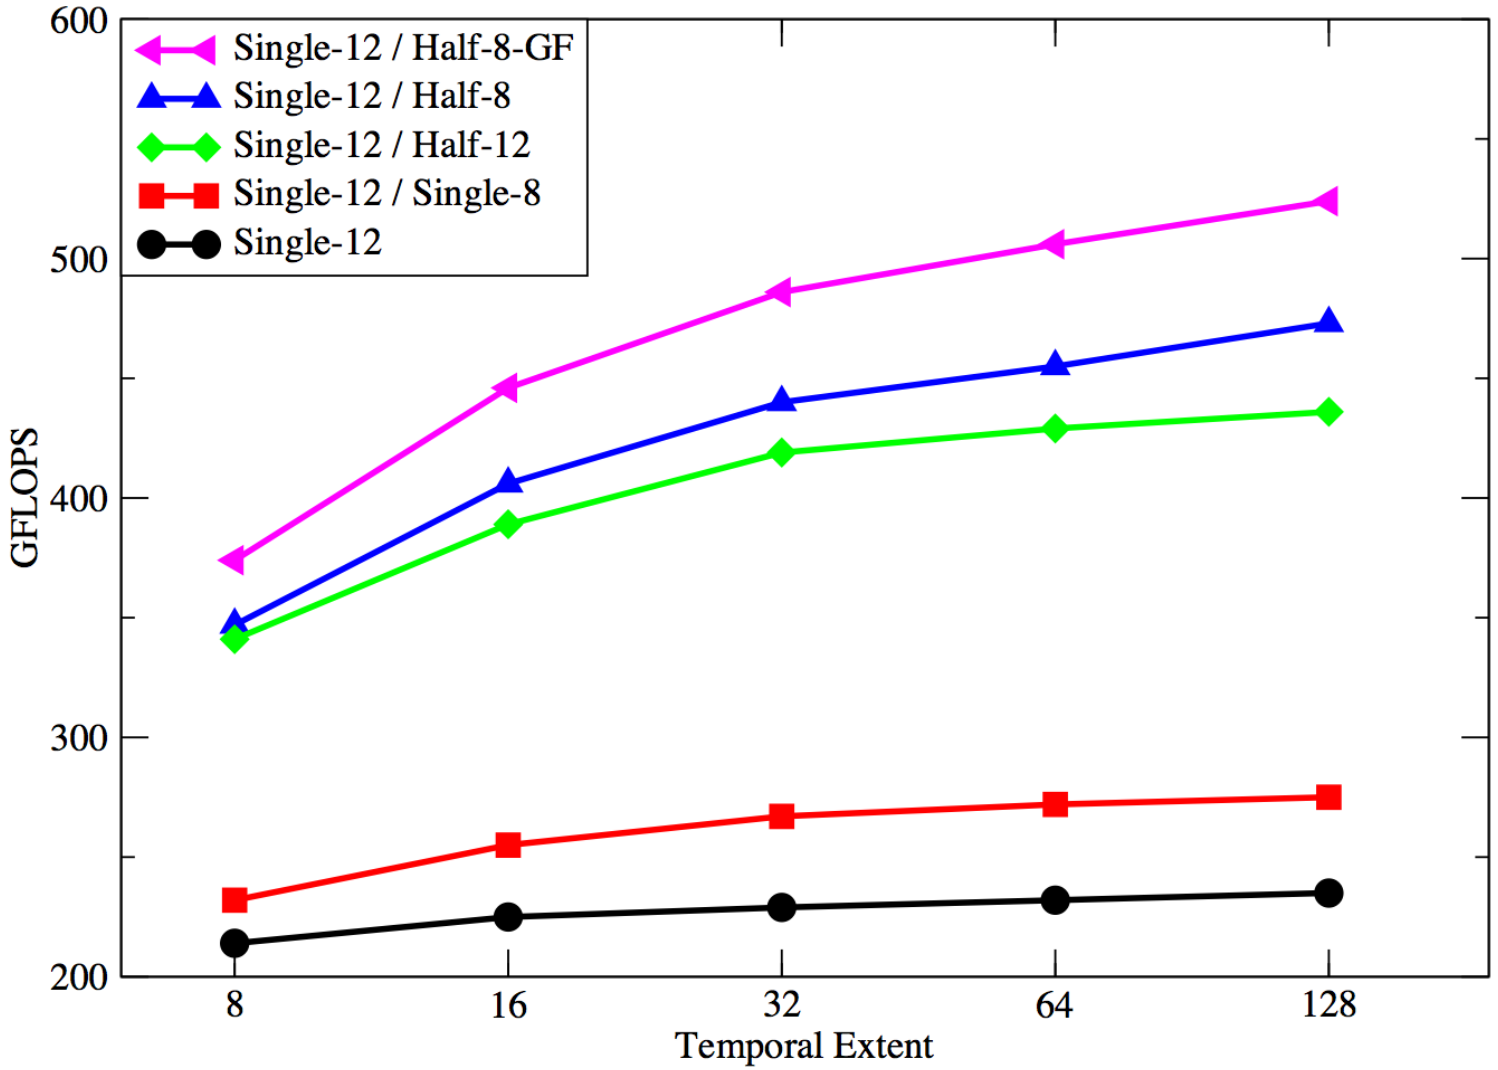

CUDA Programming

```
void saxpy_serial(int n, float a, float *x, float *y)
{
    for (int i = 0; i < n; ++i)
        y[i] = a*x[i] + y[i];
}
// Invoke serial SAXPY kernel
saxpy_serial(n, 2.0, x, y);
```

Standard C Code

```
__global__ void saxpy_parallel(int n, float a, float *x, float *y)
{
    int i = blockIdx.x*blockDim.x + threadIdx.x;
    if (i < n)  y[i] = a*x[i] + y[i];
}
// Invoke parallel SAXPY kernel with 256 threads/block
int nblocks = (n + 255) / 256;
saxpy_parallel<<<nblocks, 256>>>(n, 2.0, x, y);
```

Parallel C Code

Mixed precision with reliable updates

- Using a mixed-precision solver incorporating “reliable updates” (Clark et al., arXiv:0911.3191) with half precision greatly reduces time-to-solution while maintaining double precision accuracy.

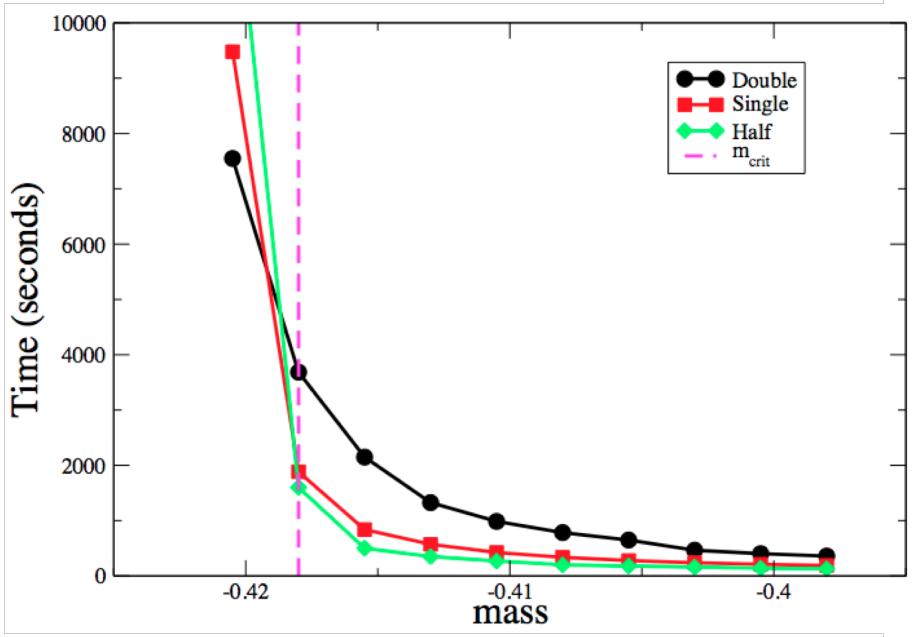

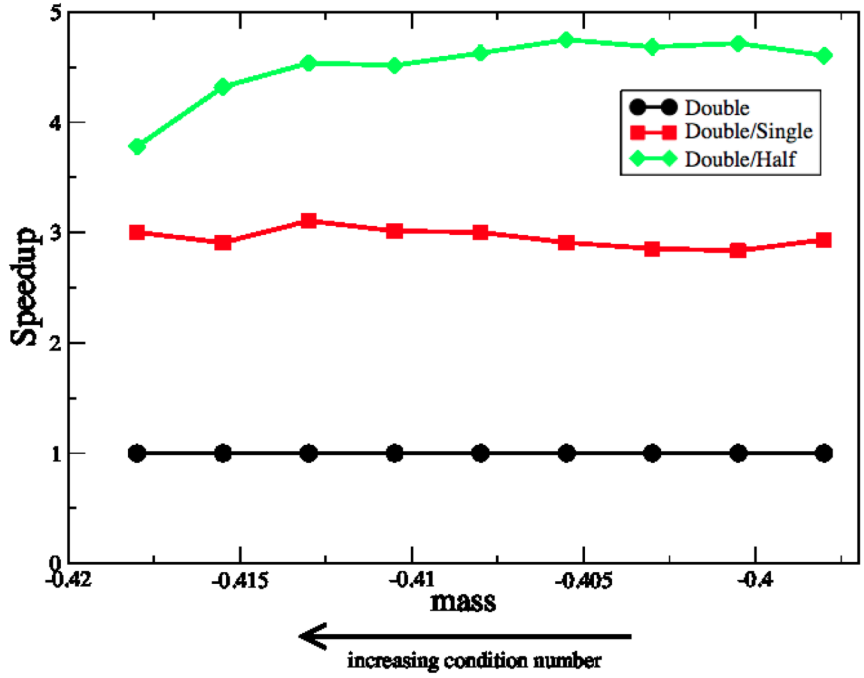

SYNTHESIS

MG+GPU Project (*)

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

O(100+)

GPU O(10+)

O(10+)

MG

Michael Cheng (BU) and Ron Babich/Mike Clark (NVIDIA)

(*) NSF PetaApps: “Multigrid QCD at the Petascale”
NSF SI2-SSI: “Scalable Hierarchical Algorithms for Extreme Computing”

Nov. 12, 2012 Titan supercomputer leads latest Top 500 list, newly-available Xeon Phi chip cracks top ten

MG on Accelerators

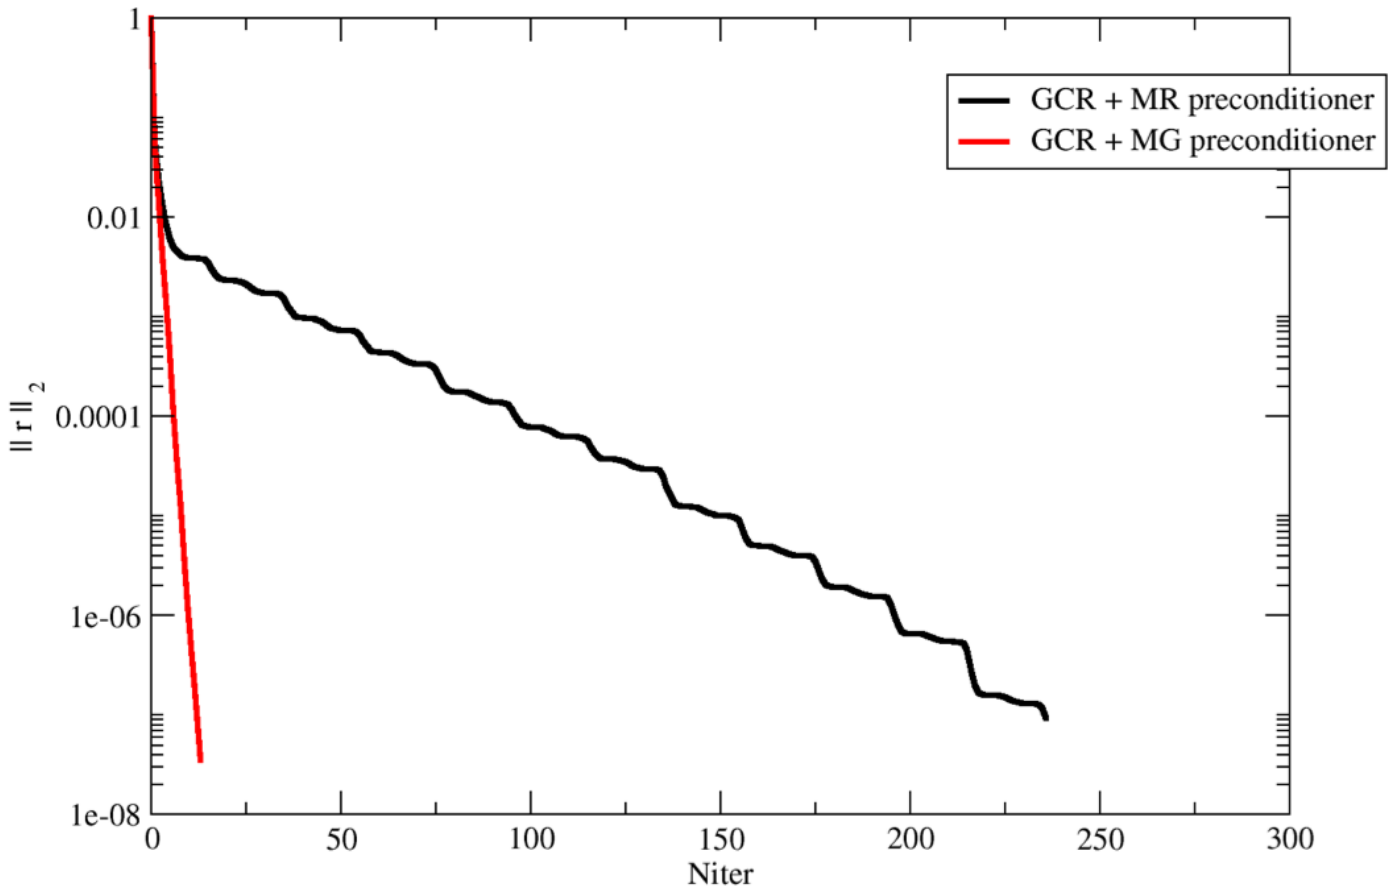

Domain Decomposition (Block Jacobi)

Communication-reducing Algorithms

- Reduce inter-node communication *and* synchronization
  - Inter-node communication comes from face exchange
  - Synchronization comes from global sums
- Utilize domain-decomposition techniques, e.g., Additive Schwarz

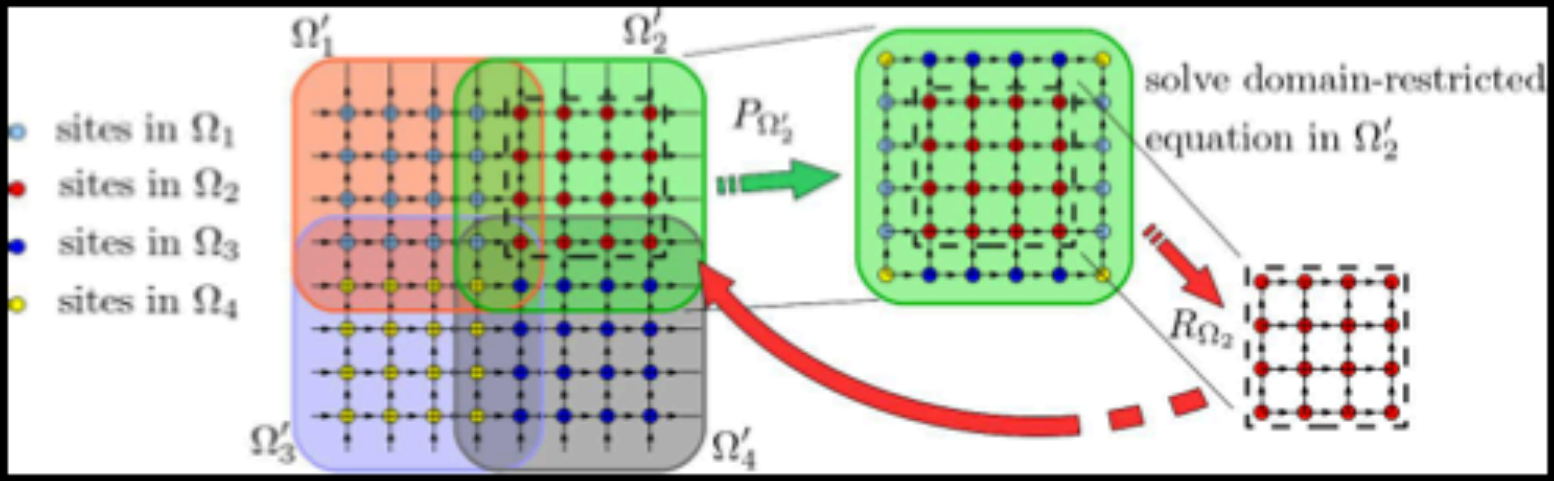

figure taken from Osaki

- Utilize s-step solvers to suppress global sums
  - Allows for the introduction of tower-kernels to decrease local comms

Domain Decomposition: Network Data Reduction

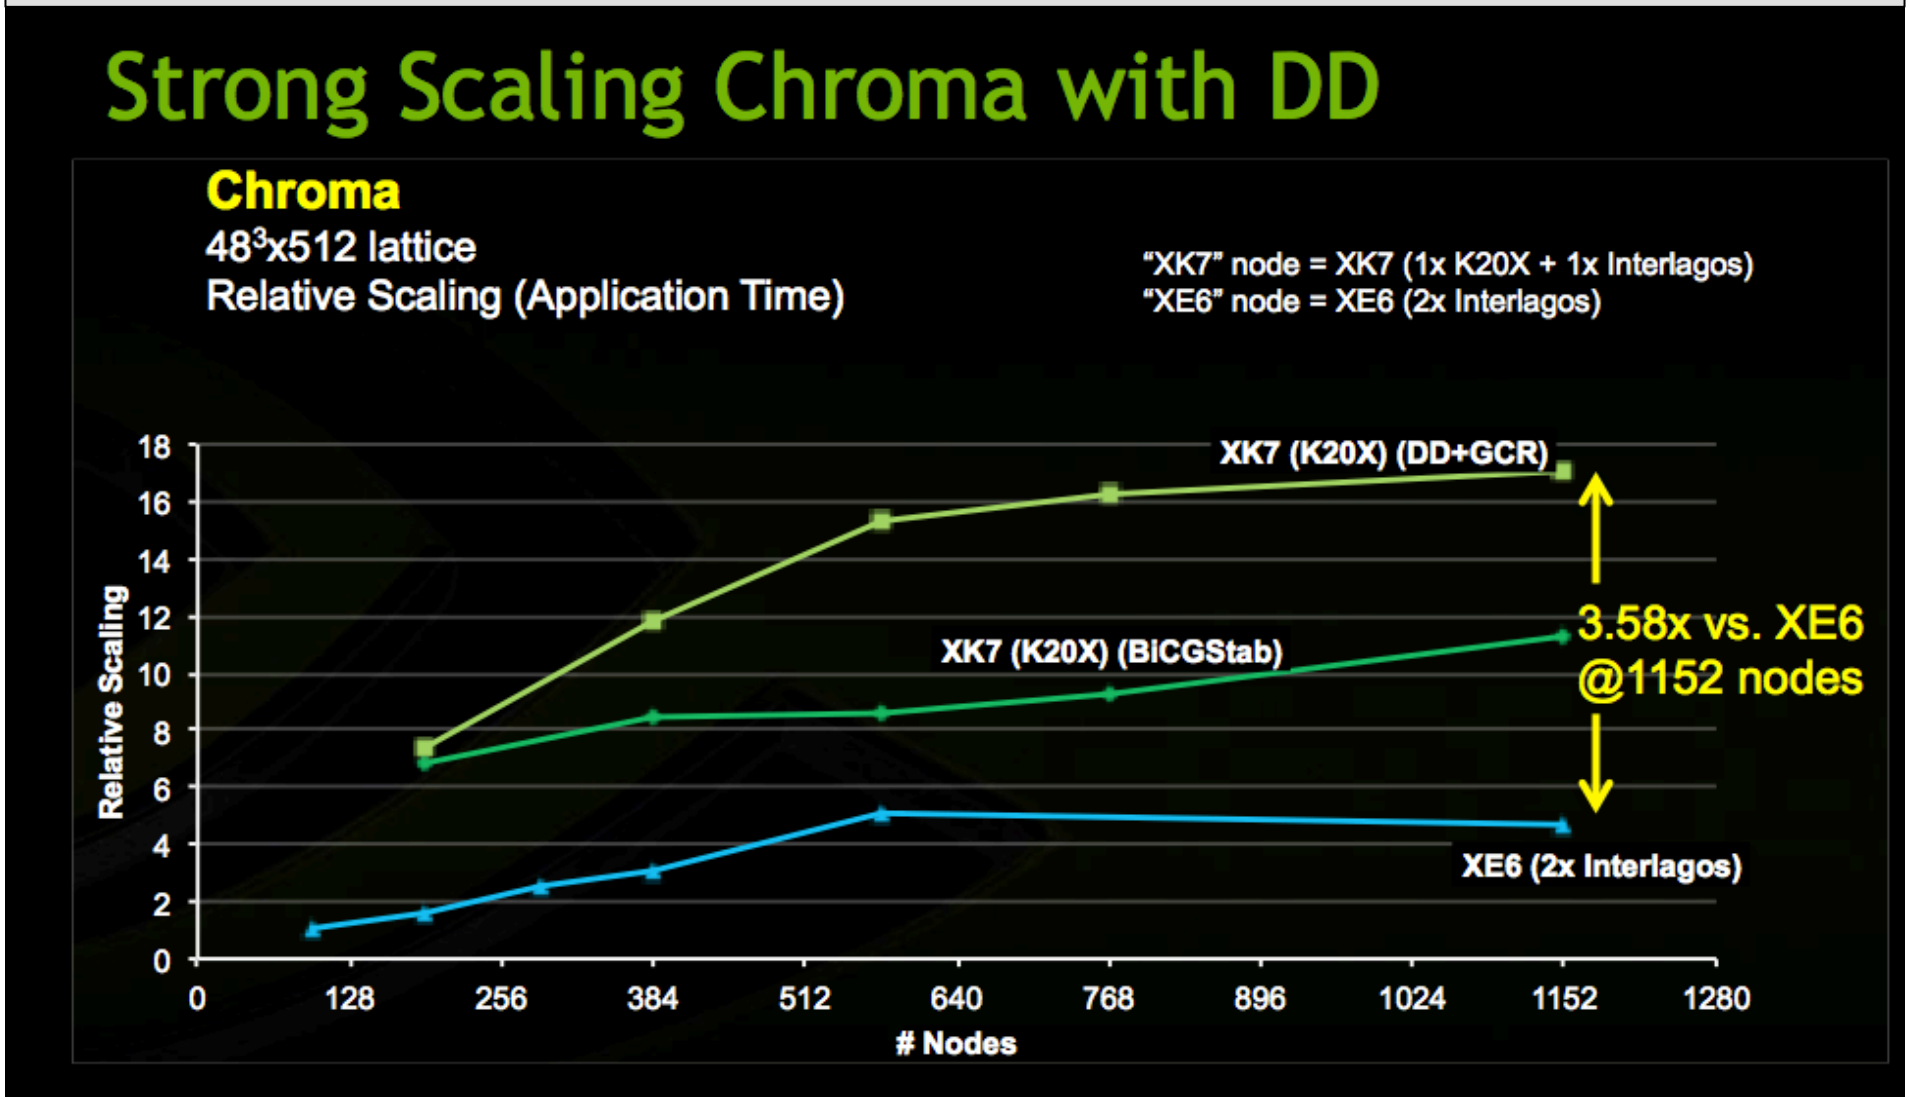

Third USQCD/Industry Collab: Intel

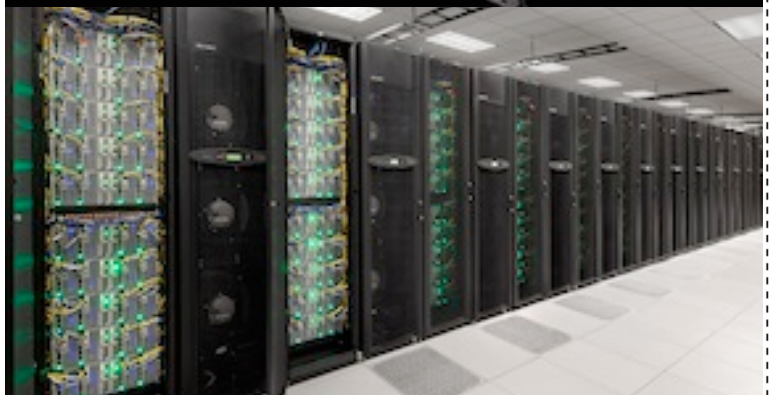

NSF Stampede with Intel Xeon Phi

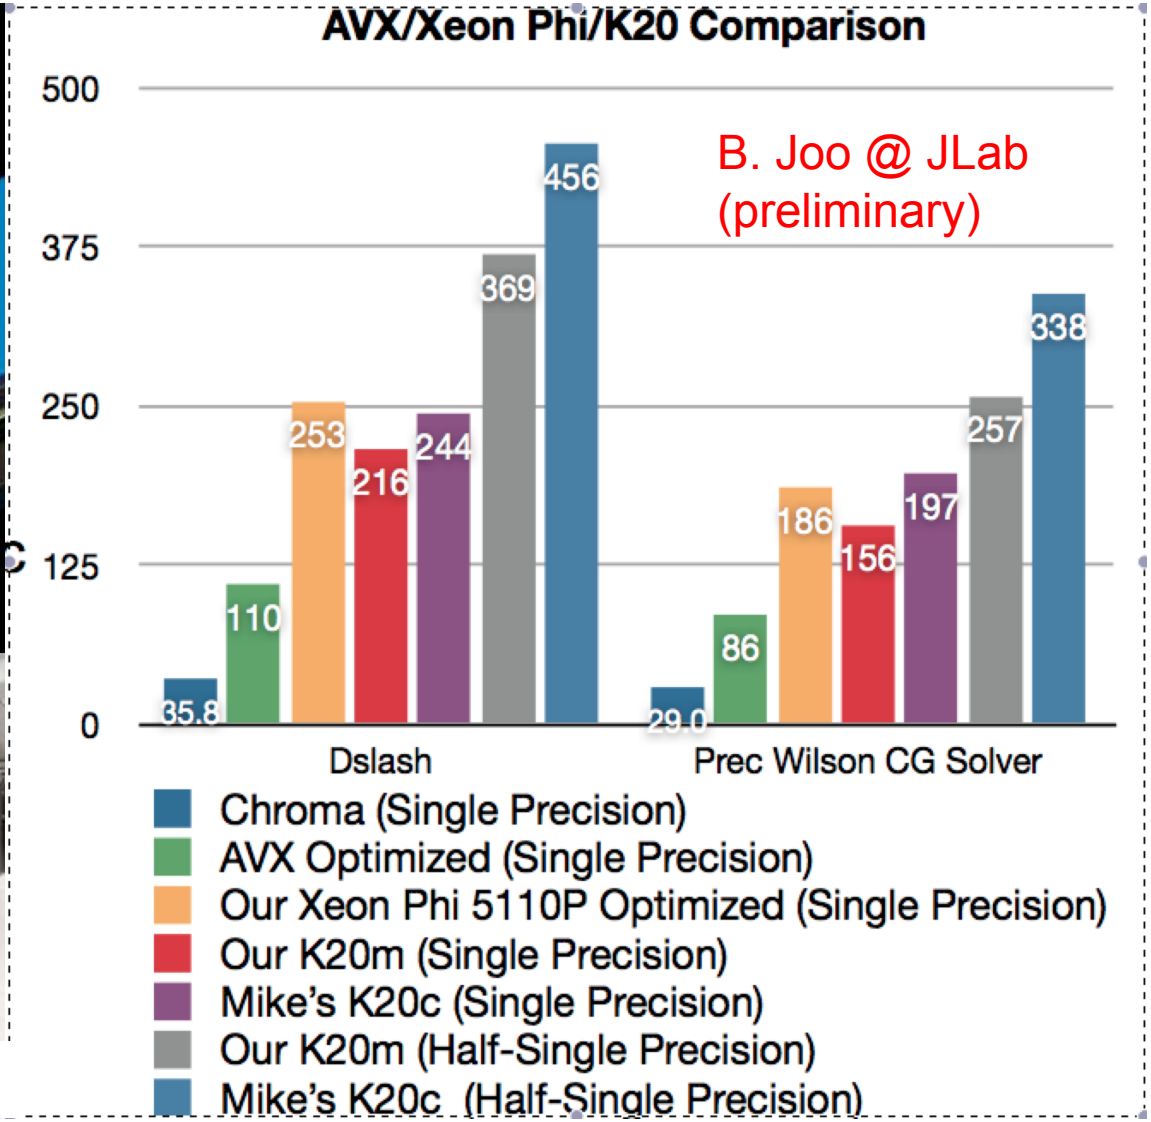

Future of Accelerator Nodes ?

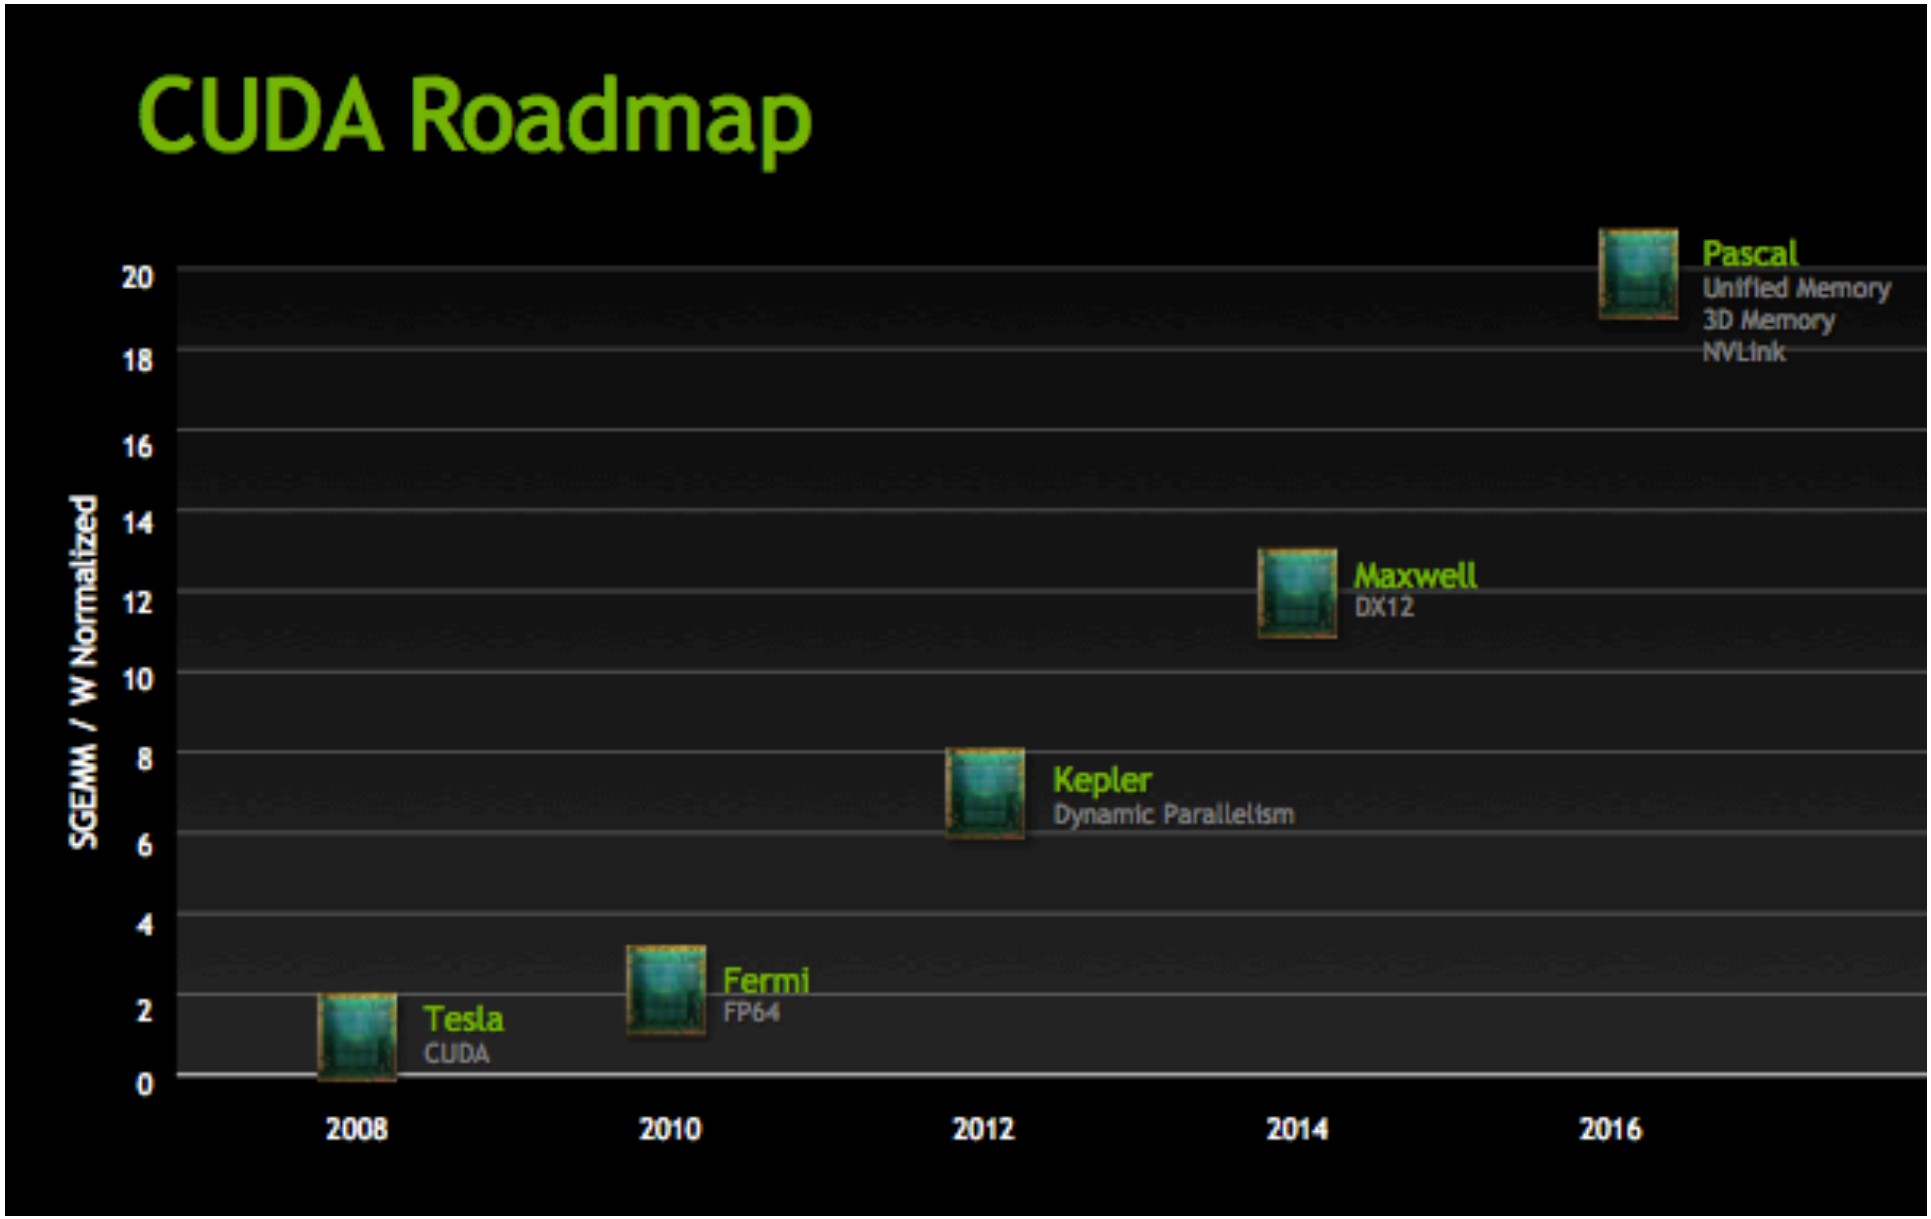

Future of Accelerators?

Introducing NVLINK and Stacked Memory

- NVLINK
  - GPU high speed interconnect
  - 80-200 GB/s
  - Planned support for POWER CPUs

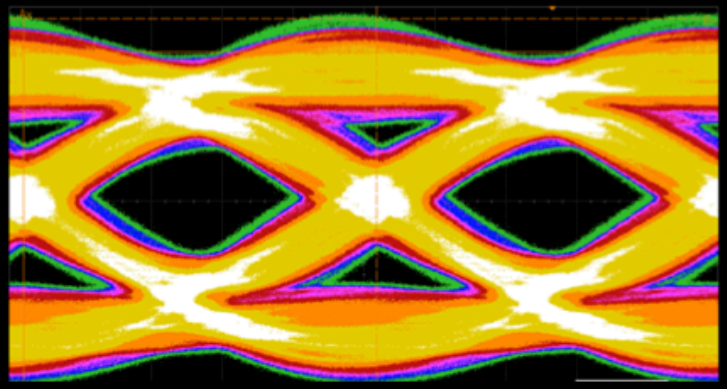

- Stacked Memory
  - 4x Higher Bandwidth (~1 TB/s)
  - 3x Larger Capacity
  - 4x More Energy Efficient per bit

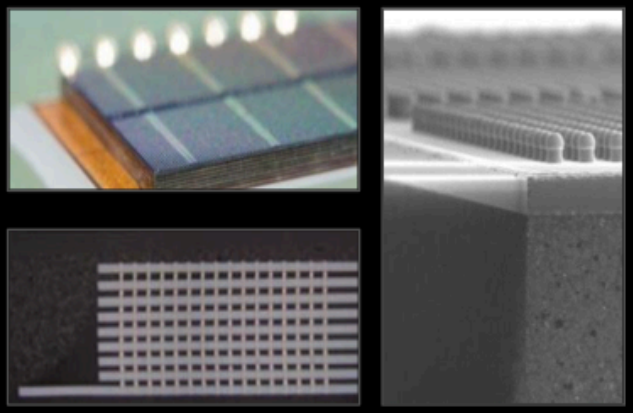

In 2015?

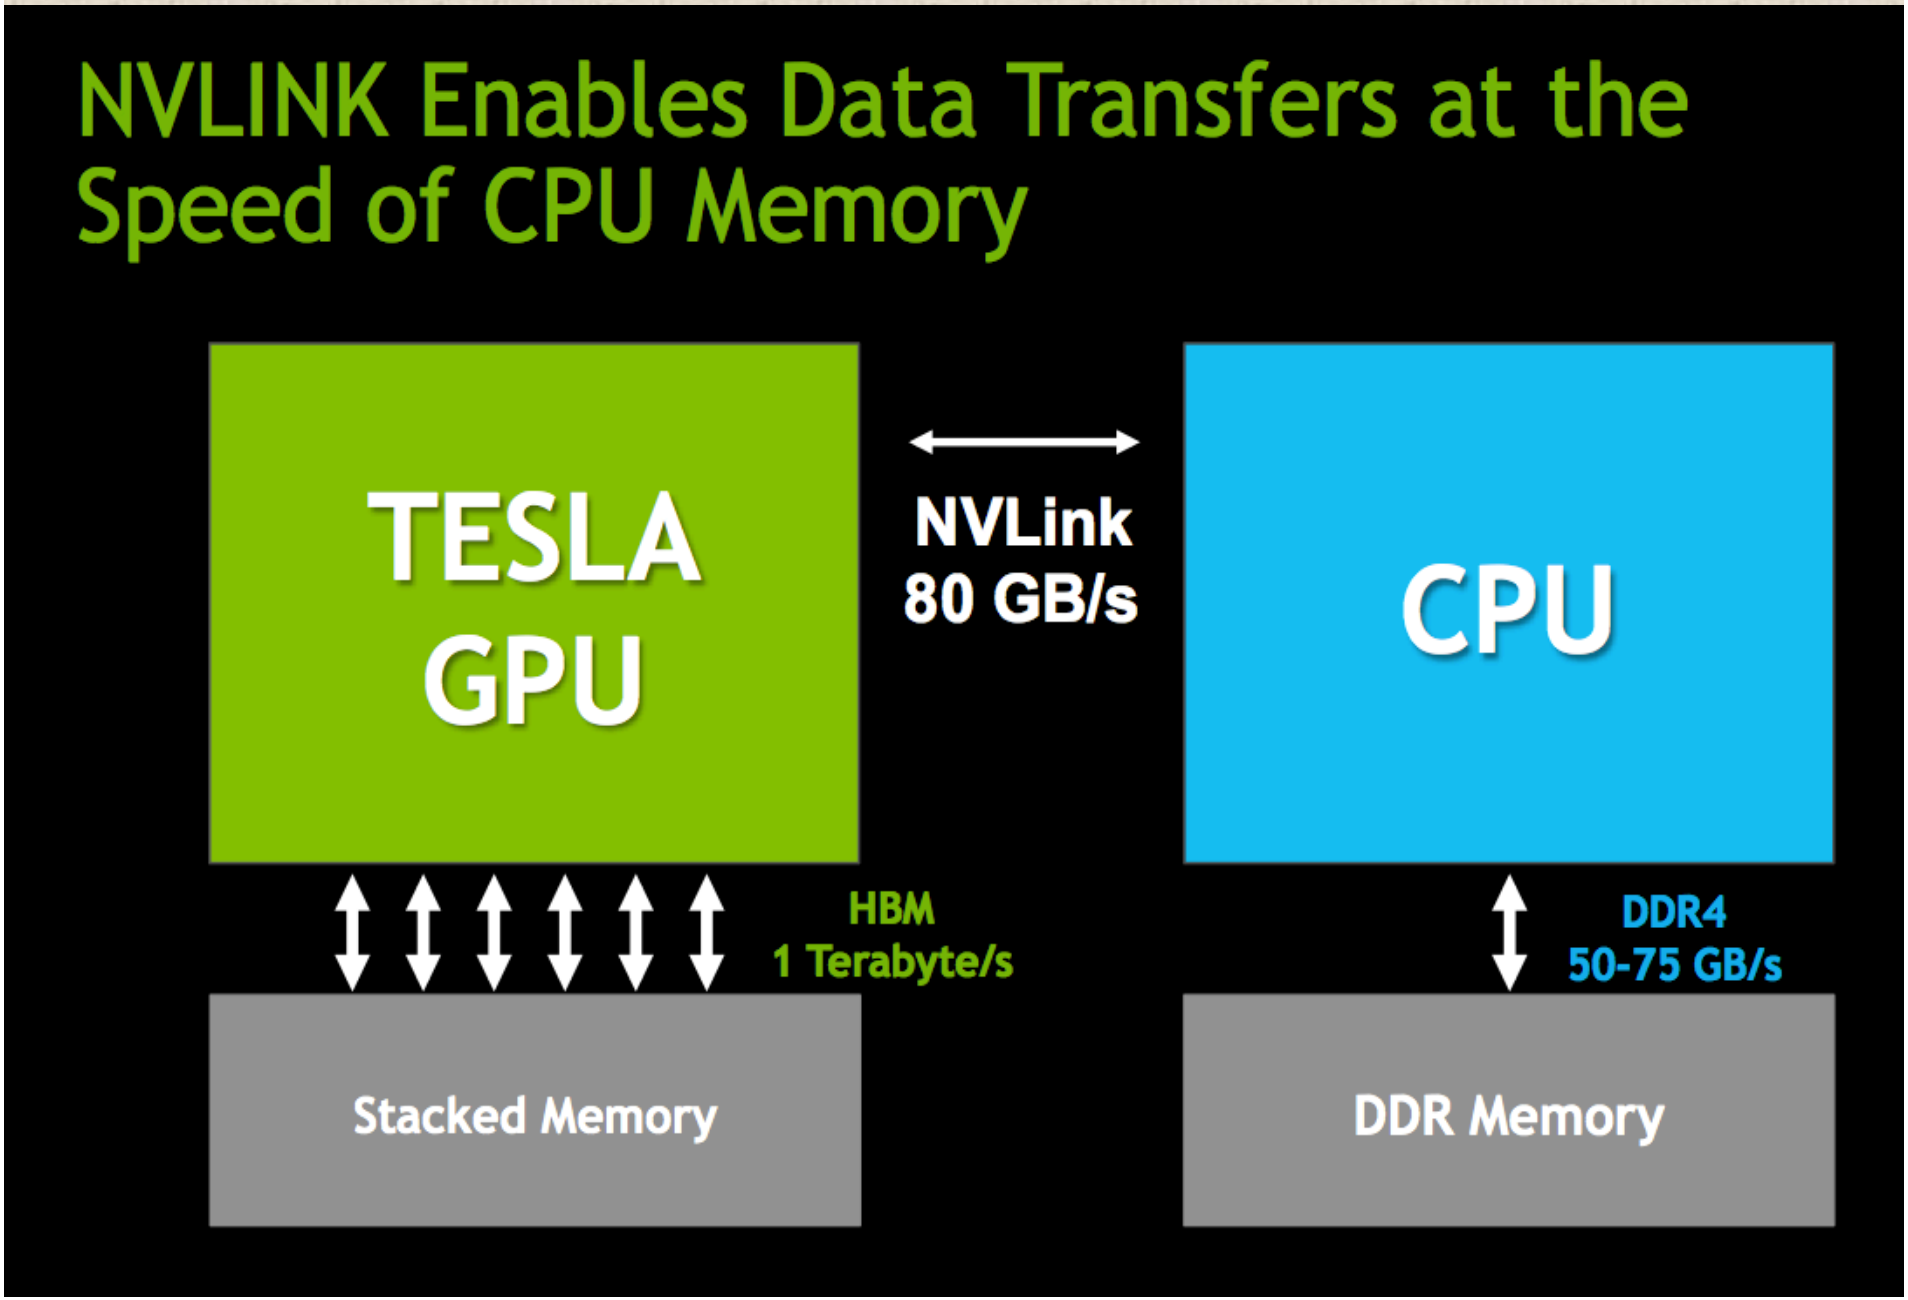
NVLINK Enables Data Transfers at the Speed of CPU Memory
TESLA GPU
NVLink 80 GB/s
CPU
HBM 1 Terabyte/s
DDR4 50-75 GB/s
Stacked Memory
DDR Memory

Lua (Moon in Portuguese): extensible scripting language to interface to High Performance codes and Applied Math tools

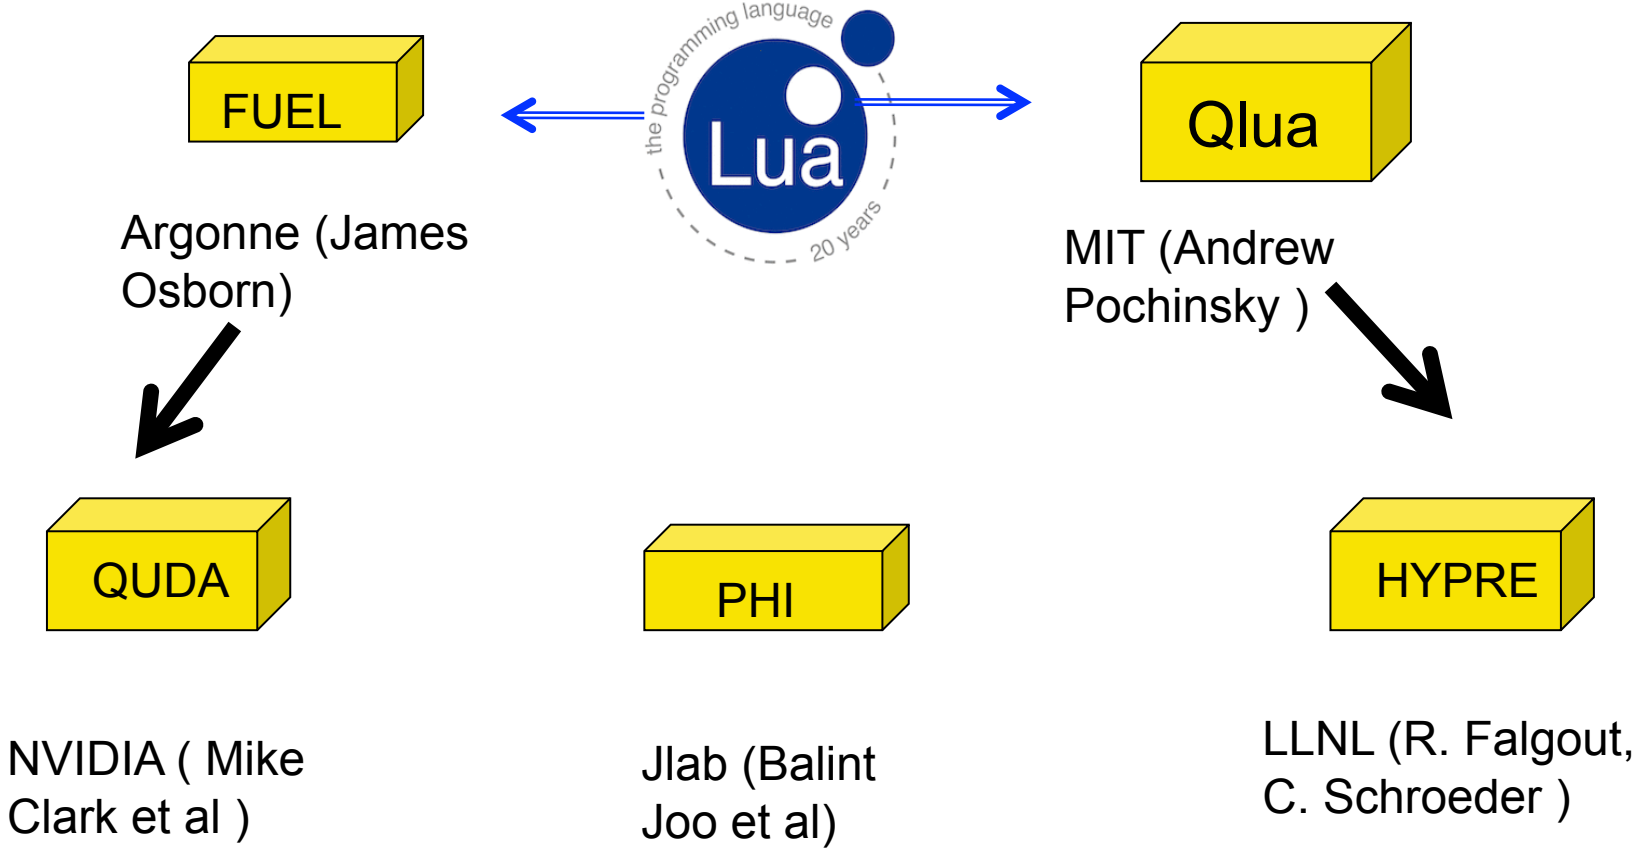

Many Hardware Targets:

New Lattice for Exponential Scales? Radial Quantization for CFT

- (near) Conformal Field Theories are important
  - BSM walking technicolor
  - AdS/CFT weak-strong duality
  - Model building

- Lattice difficulty: scales are (nearly) exponential.

- Hypercubic vs Radial Lattice

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

3-d Test of concept

$$\lim_{x\to 0} \mathcal{O}(x)|0\rangle = |\mathcal{O}\rangle$$

state-op corr

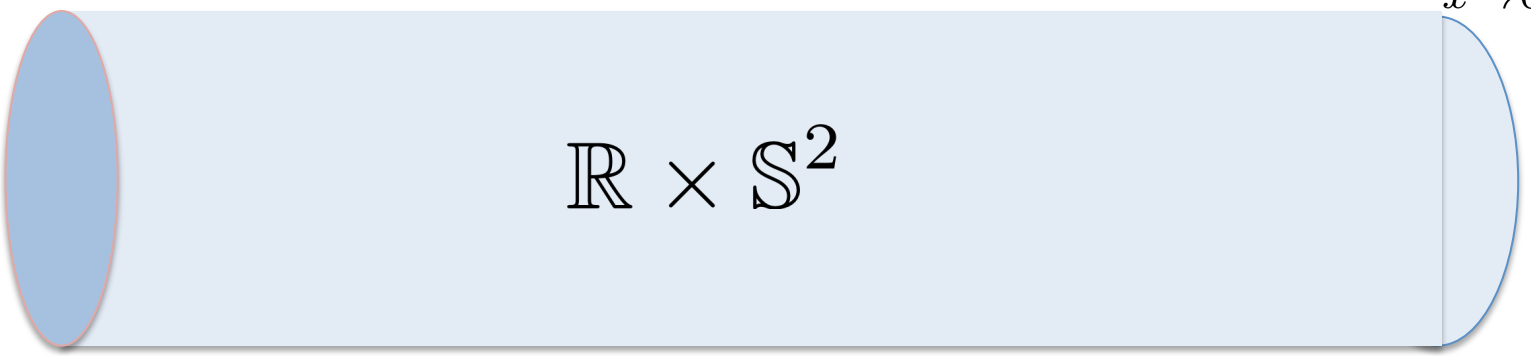

$(r = 0)$

$+\infty$ $\tau = \log r$ $-\infty$

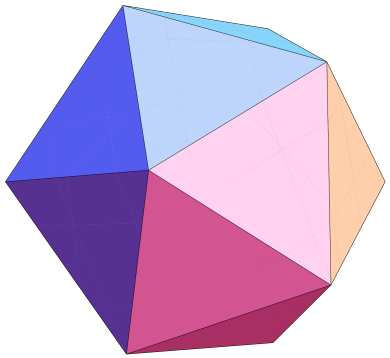

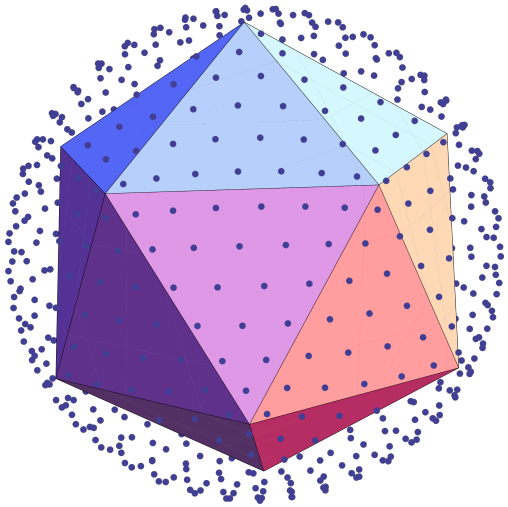

Graphene

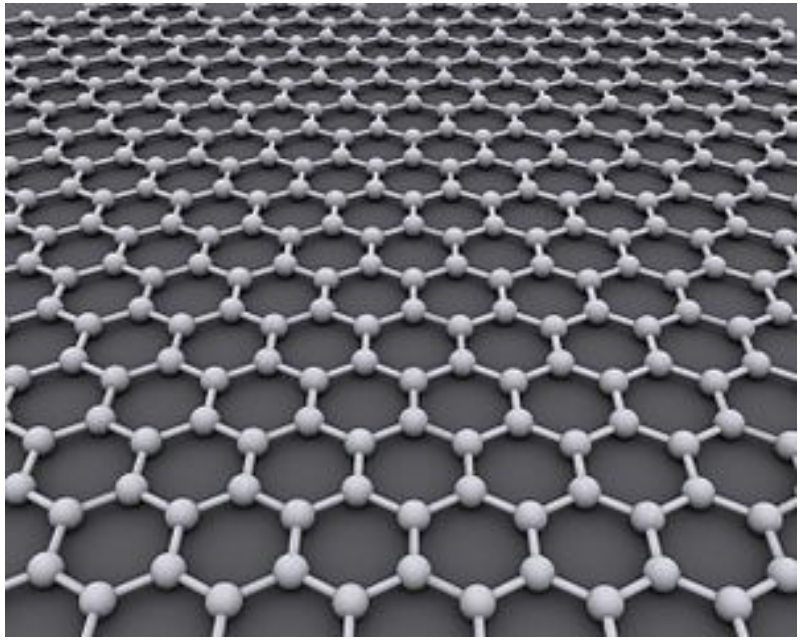

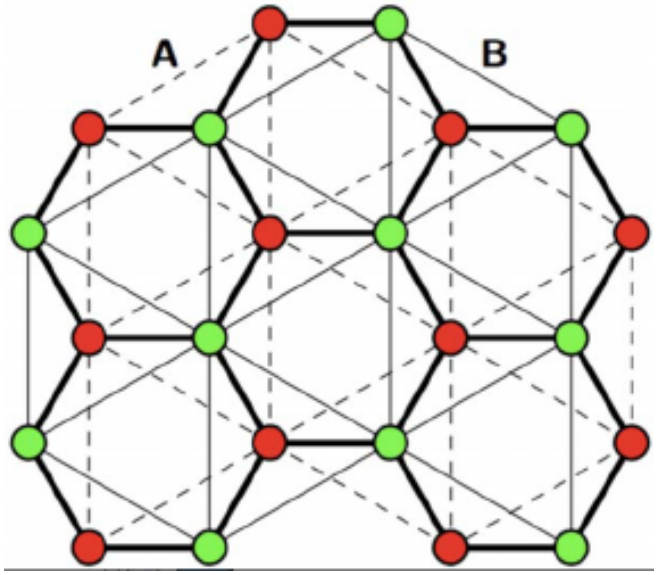

Graphene is 2+1 dimension Carbon sheet with Dirac fields: But lattice is real Hexagonal structure. Couple to coulomb potential and phones act like gauge fields! Ideal for Lattice field theory, MG and GPU! (Brower, Rebbi and Schaich)

Lattice GT (and hardware) – is NOT extremely uniform

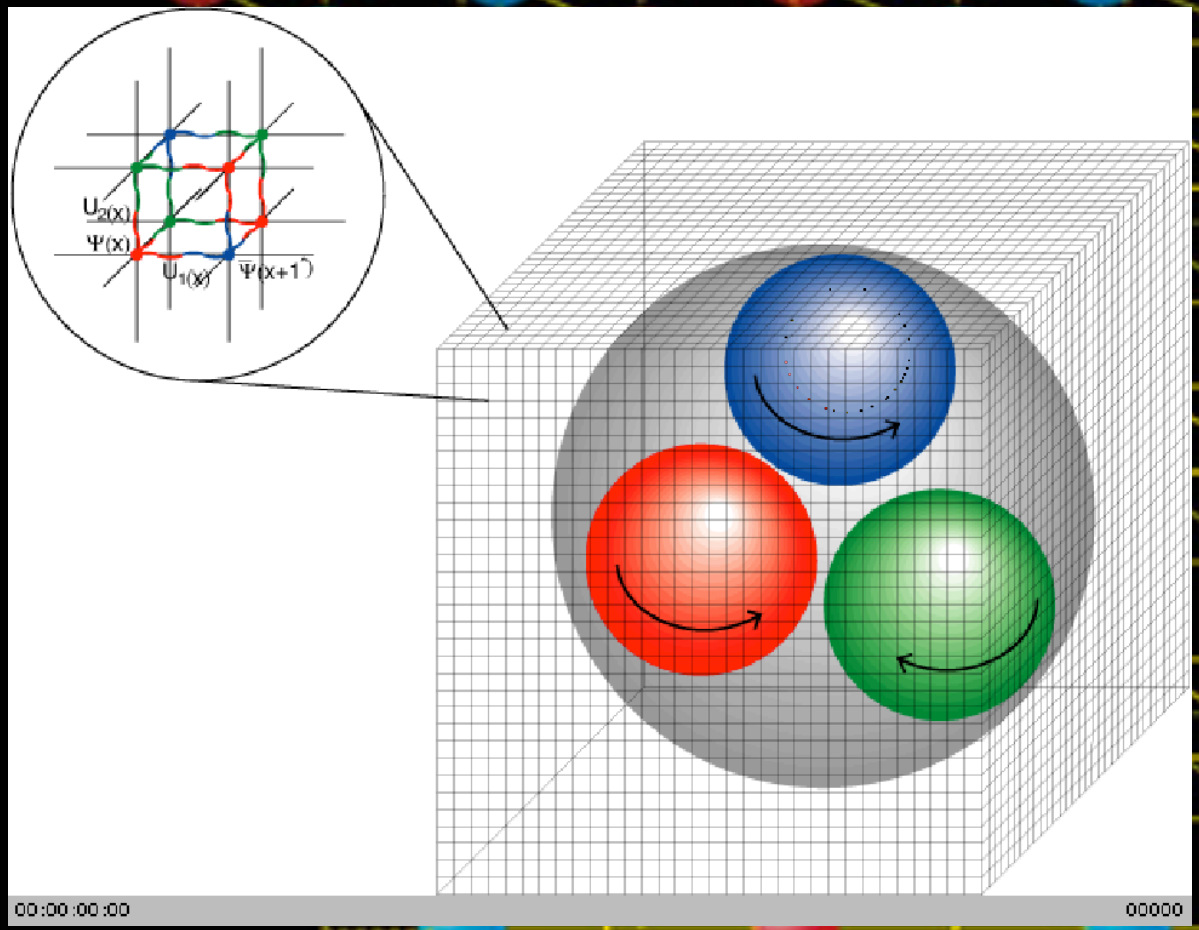

- Periodic or very simple boundary conditions NOT
- SPMD: Identical sublattices per processor NOT

Put PDE's on lattice

$$D\psi(x) = [m + \sum_{\mu} \gamma_\mu(\partial_\mu + iA_\mu(x))]\psi(x)$$

$$D\psi(x) = [m - \frac{1}{2a}\sum_{\pm\mu}(1 \pm \gamma_\mu)e^{\pm iA_\mu(x)}]\psi(x \pm \mu)$$

Q & A

- WE BADLY NEED A

"Formula Translation Language!"

It must include composable parallel toolbox that the compiler/operating system can optimize.

MPI++ and FFTW is not enough!