Goals for our

"Experimental GPU cluster for Fundamental Physics"

Brower, Barba, Rebbi

And Bruce Boghosian

support by NSF grant number OCI-0946441

Project

- Build multi-GPU cluster to enable experimentation with algorithms, programming strategies, and system software etc.
  - Goal is peak performance on latest technology
  - Prototype for GPGPU at FERMI/Jlab & DOE SciDAC software
- First Target application:
  - Lattice Field Theory (QCD):
    - Brower /Rebbi/Clark/Babich et al
  - Computational Fluid Dynamics (CFD):
    - Barba/Boghosian/
- Train students in multi-core parallel programming
  - Using CUDA,OpenCL,Python,etc.

Harvard/BU
Cuda Center of Excellence

Tesla 1070 Nvidia Gift

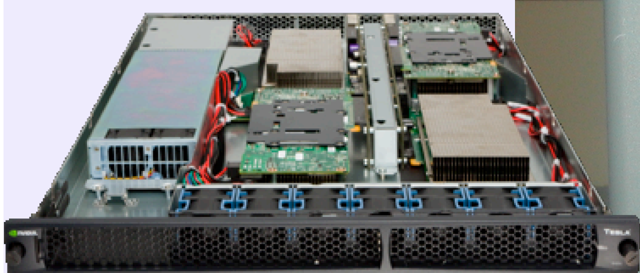

- *Disruptive QCD Technology! Graphic Processor Units*

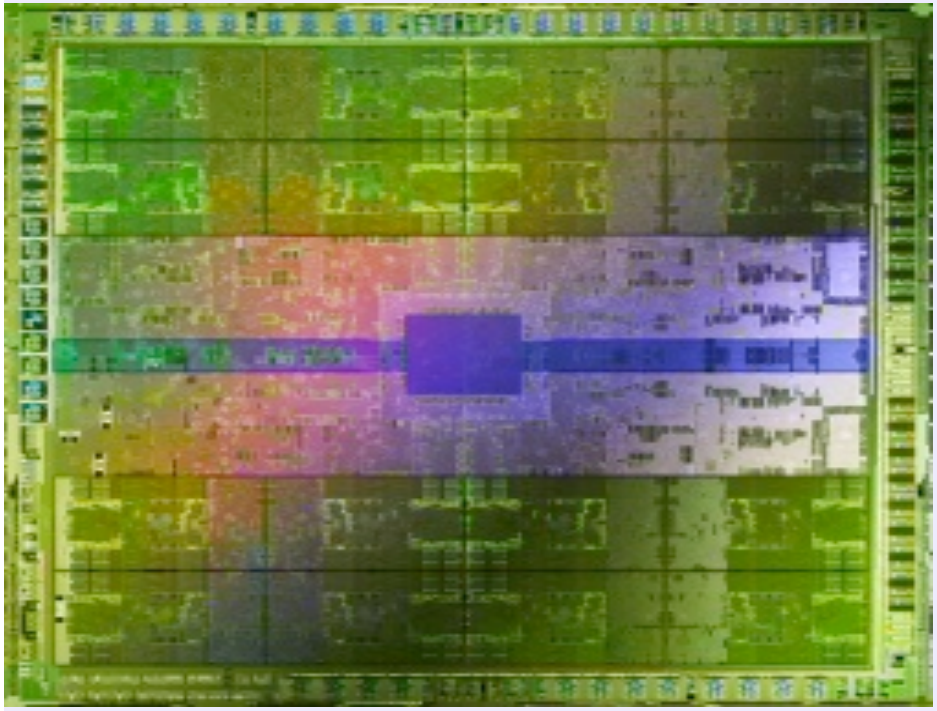

Nvidia's Fermi GPU:
512 cores x 32 bit FPU

4 GigaByte Mem, 1 Tflop peak

•GPU/BU code sent to Jefferson Lab for 3 million $ ARRA Cluster
⇒ 5 x performance @ 20% extra cost

Historical Perspective: First "commercial" QCD machine

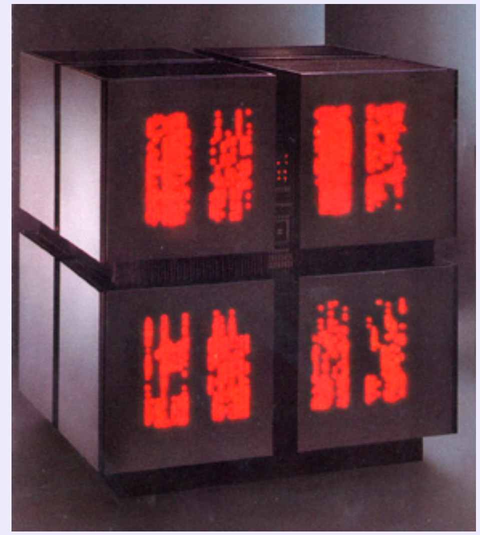

http://www.mission-base.com/tamiko/cm/cm-tshirt.html

First University installation in 1989*

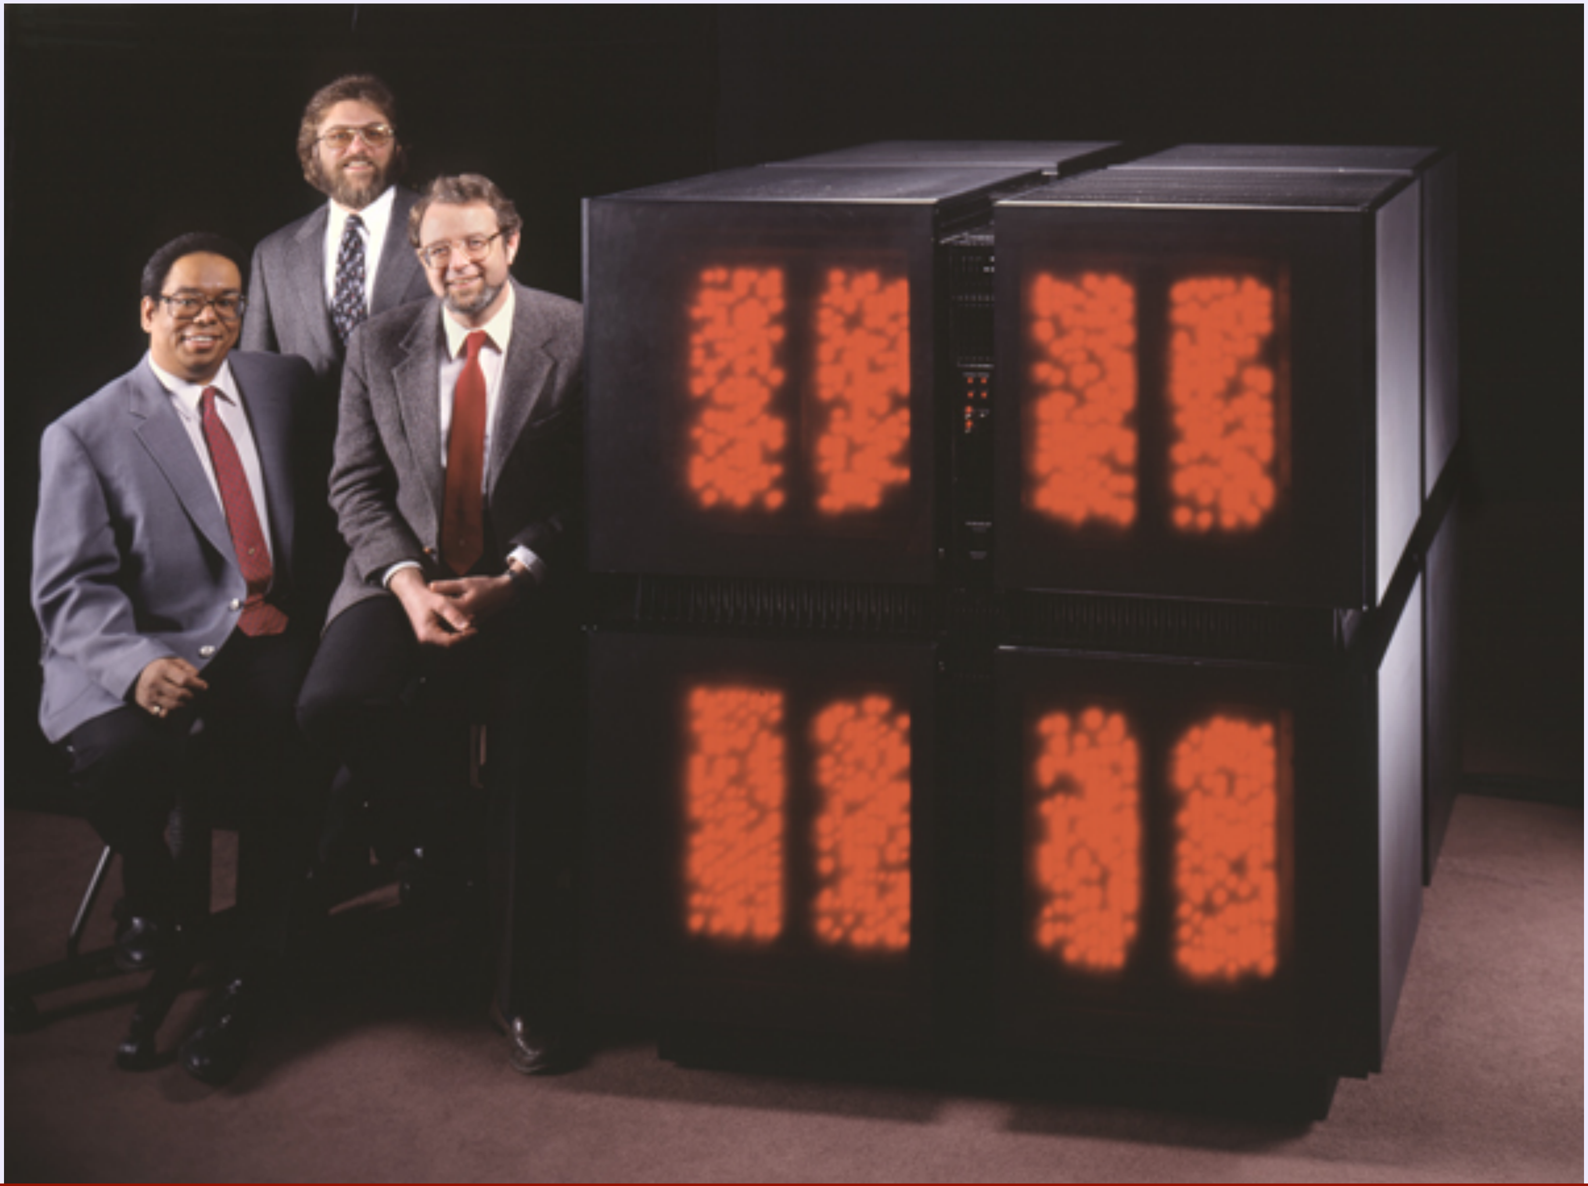

*(from left) Roscoe Giles, Glenn Bresnahan and Claudio Rebbi with CM-2

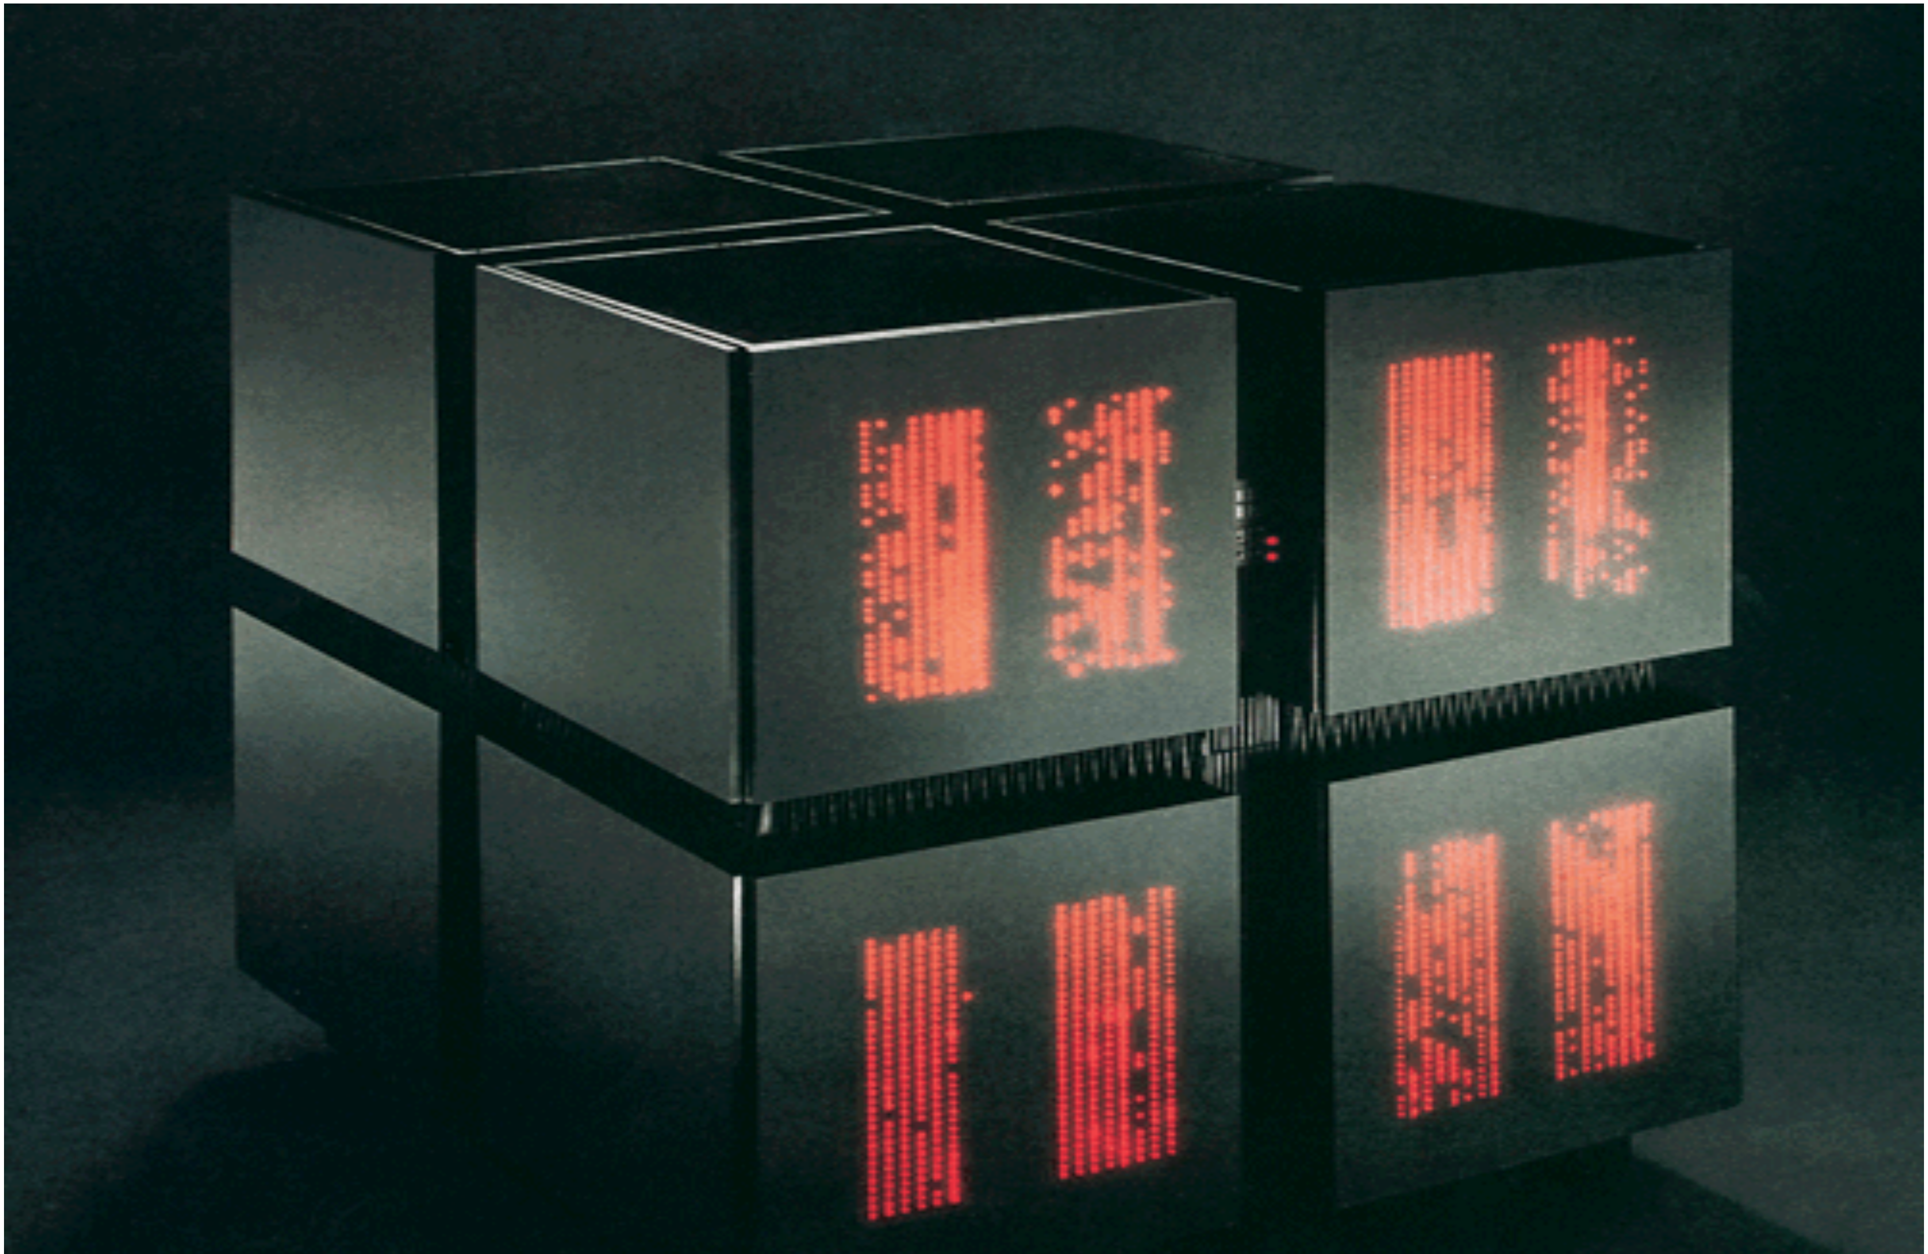

In late 1980's Thinking Machines Corporation the 64K 1 bit processer CM-2 with performance in excess of 2500 MIPs, and floating point above 2.5Gflops

Killed by Beowulf-clusters

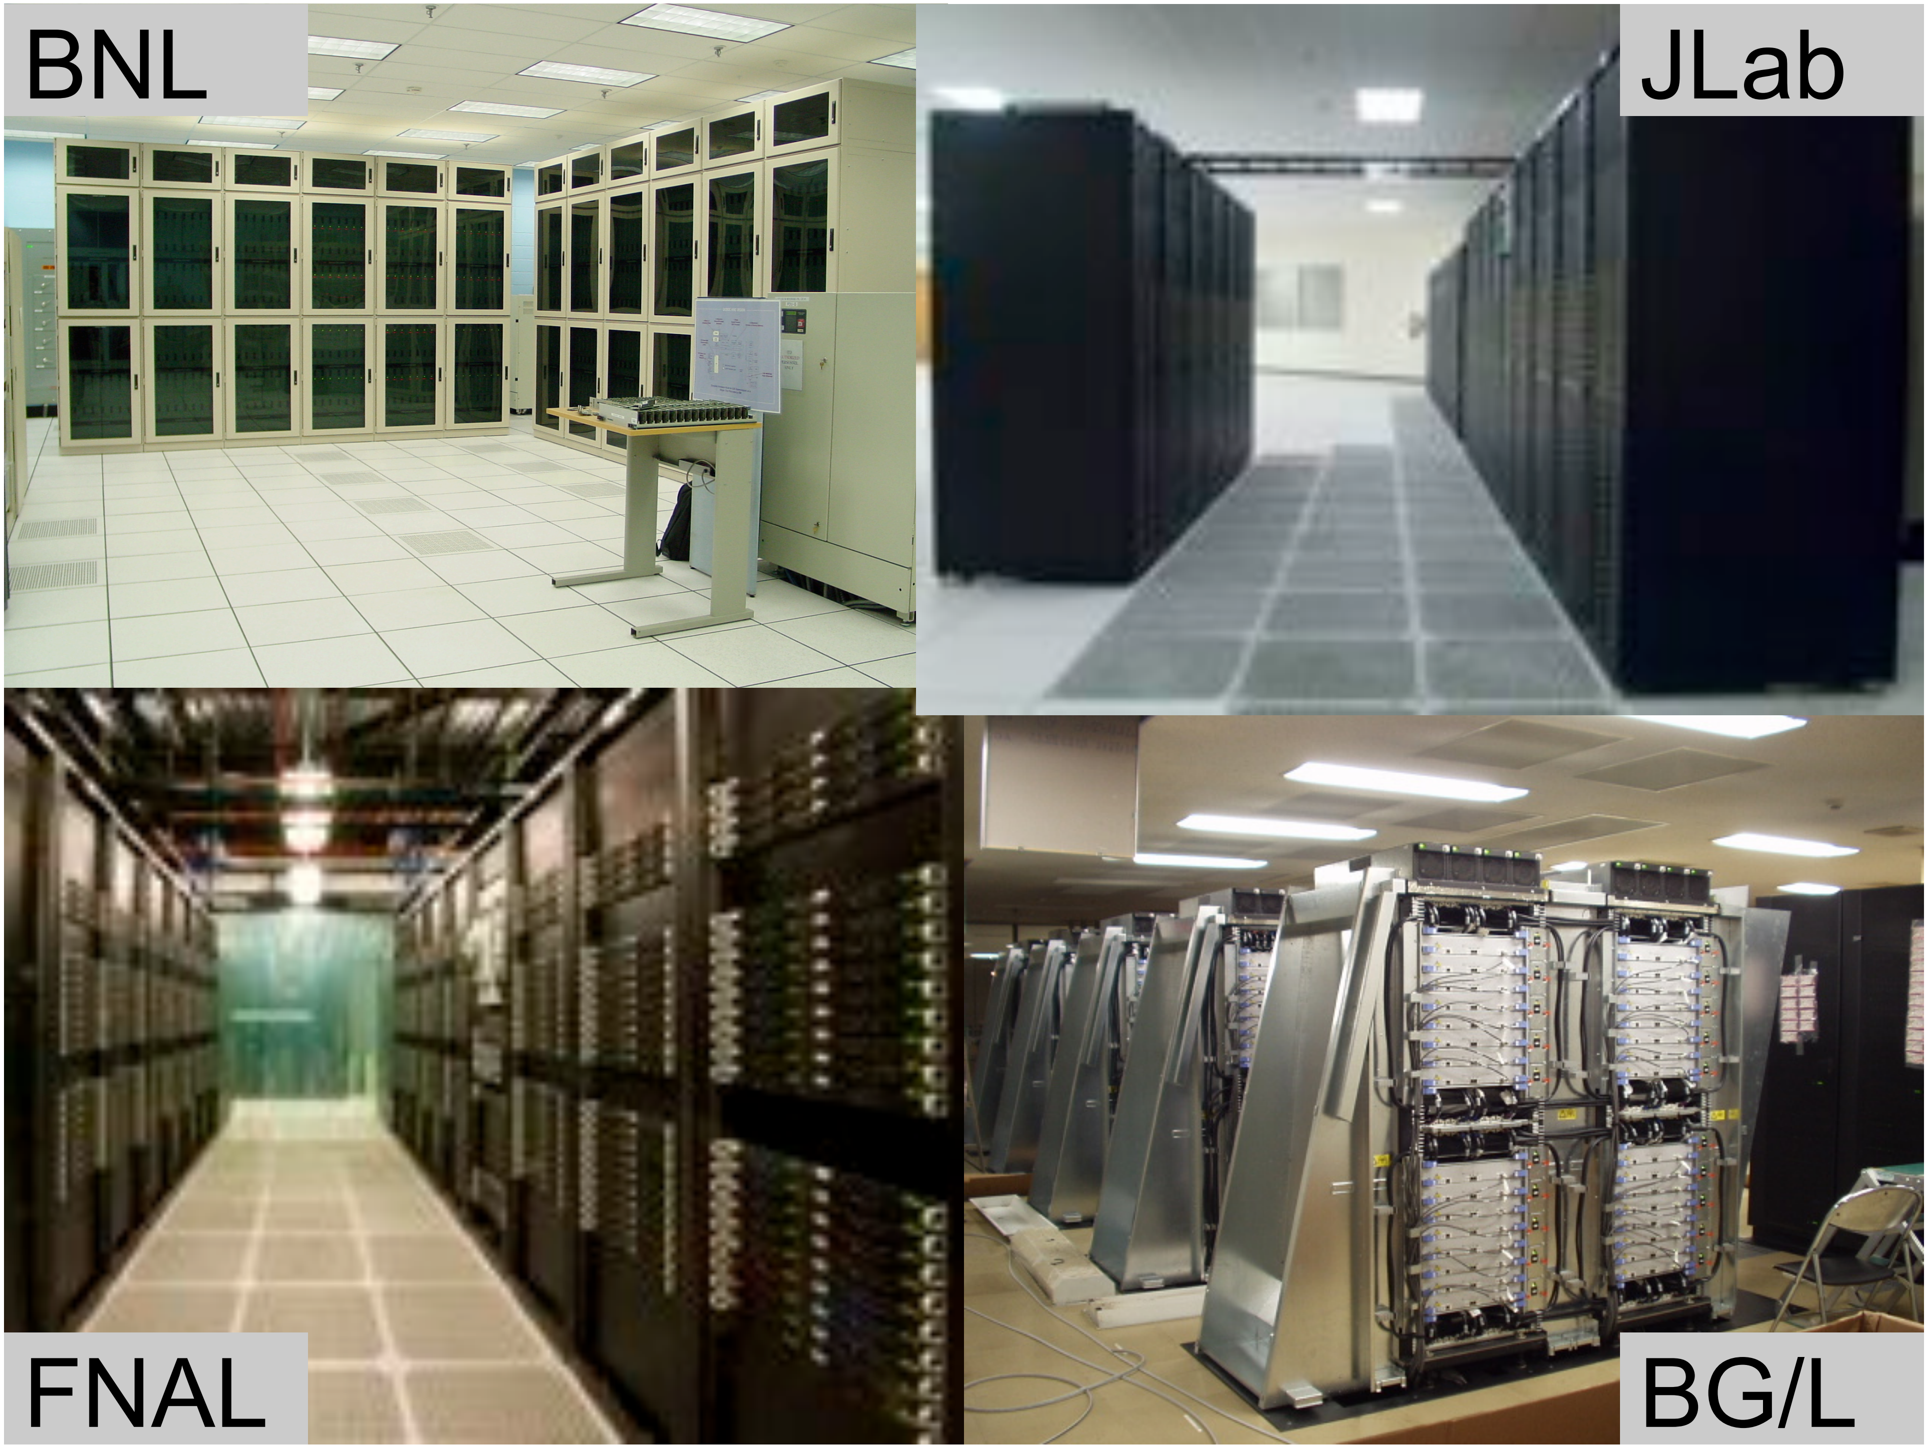
BNL
JLab
FNAL
BG/L

QCD Perfect scaling!

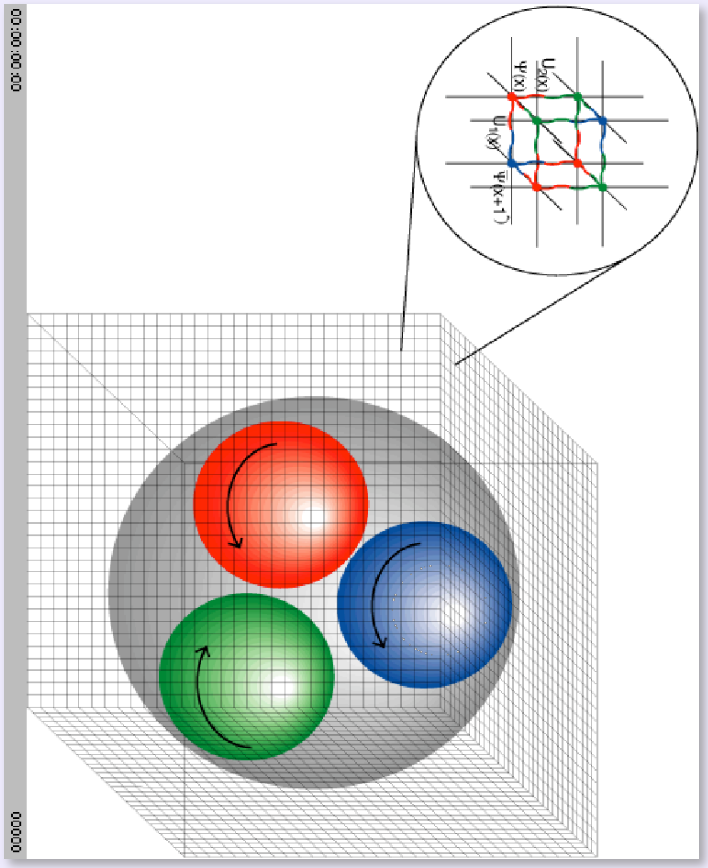

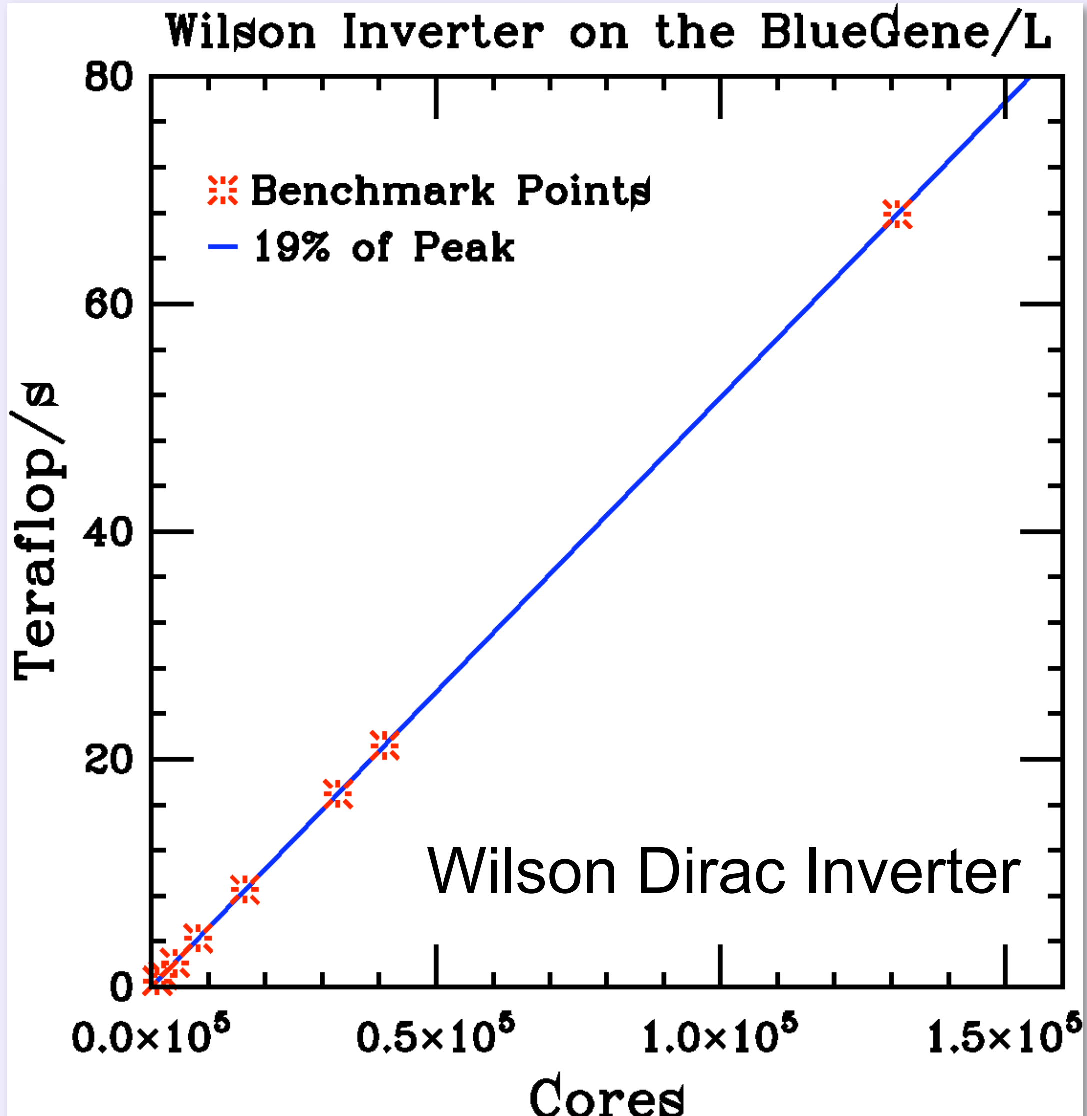

† LLNL BG/L weak scaling up to 131,072 cores: 2006 Gordon Bell award

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

BUT THIS IS NOT SUFFICIENT

K. Wilson (1989 Capri)

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

vs

ab initio Chemistry	ab initio QCD
1. 1930+50 = 1980	1. 1980 + 50 = 2030?*
2. 0.1 flops ➔ 10 Mflops	2. 10 Mflops ➔ 1000 Tflops
3. Gaussian Basis functions	3. Clever Collective Variable?

**Hopefully sooner but need \$1/Mflops ➔ \$1/Gflops!*

Disruptive many-core Architectures

- 1/4 CM-2

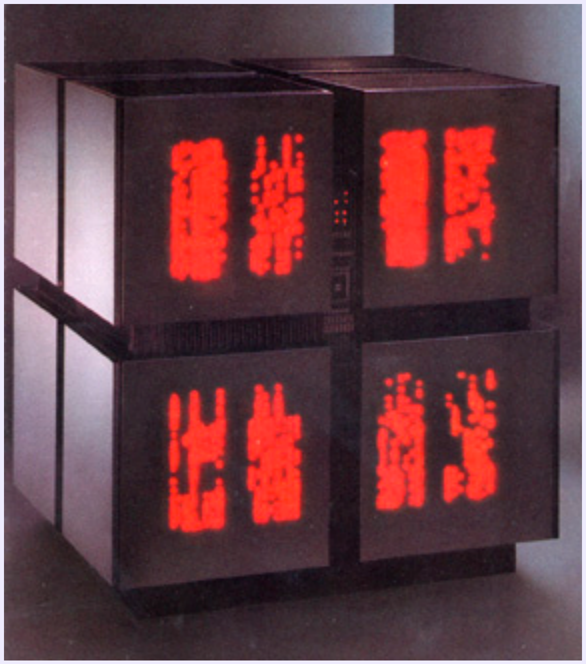

Nvidia FERMI chip

$\Rightarrow$

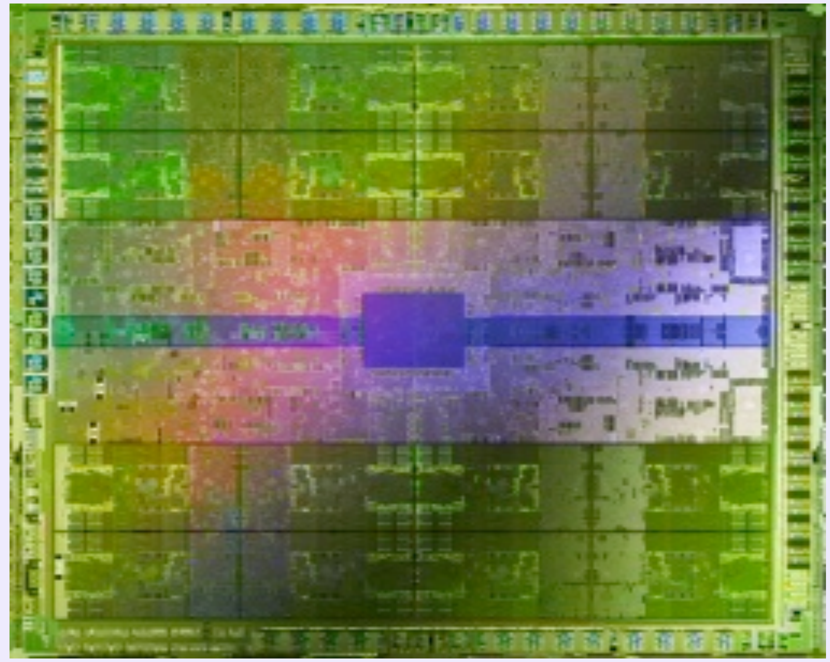

16 K bit serial PE. $\Rightarrow$ 512 × 32 bit PE = 16 K bits

Heterogeneous Challenge:

Physics & Algorithms & Hardware
ALL are becoming multi-scaled!

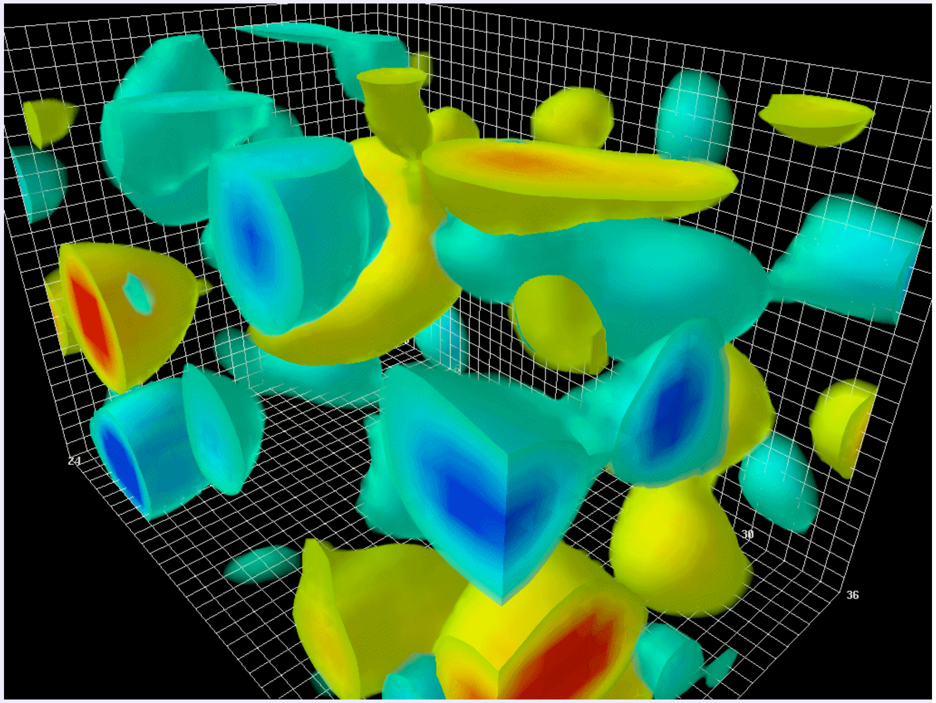

- Higher resolution lattices:
  - a(lattice) << $1/M_{proton}$ << $1/m_\pi$ << L (box)
  - 0.06 fermi << 0.2 fermi << 1.4 fermi << 6.0 fermi

3.3 x 7 x 4.25 $\simeq$ 100

QCD Physics

Multi –level Architecture

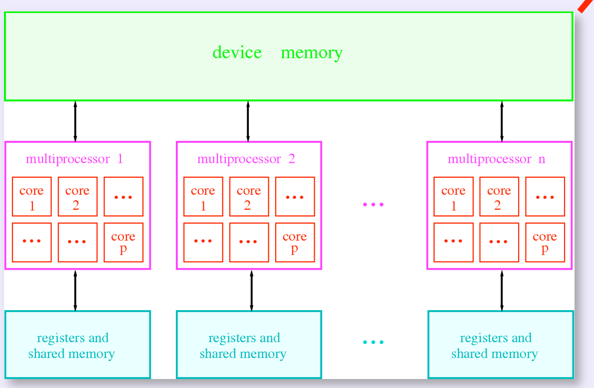

To reach beyond the Petaflops barrier we see this trend: BlueGene/Q, BlueWaters, Roadrunner, Qpace and of course GPGPU clusters!

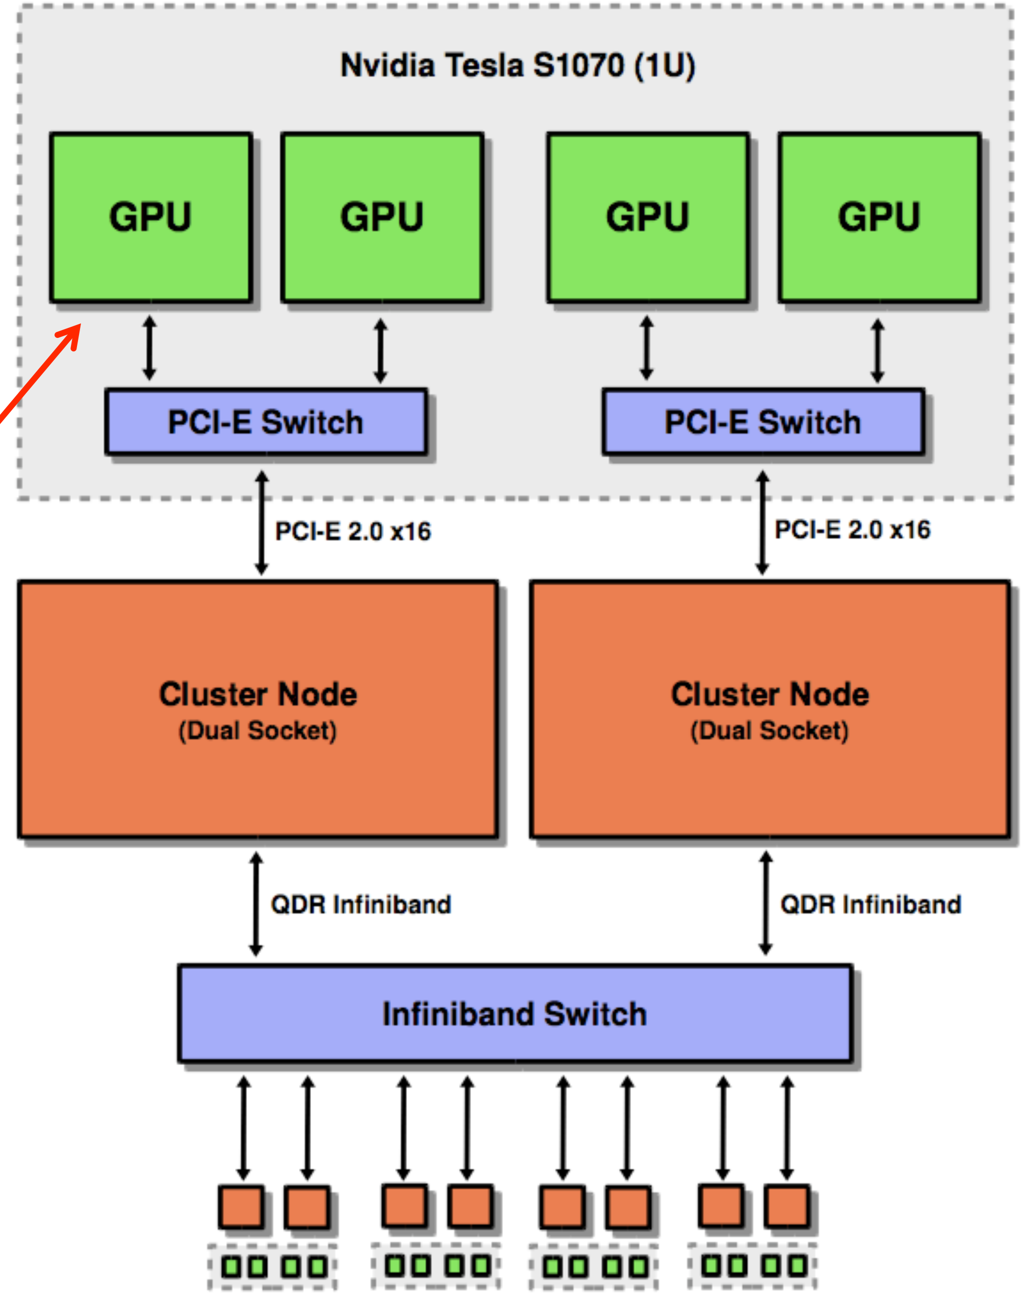

Multi-scale Algorithms

NSF PetaApps_Collaborative Research_BU /Colorado/Penn State Multi-grid/Schwartz & Domain Algorithms, 20 physics & applied math experts

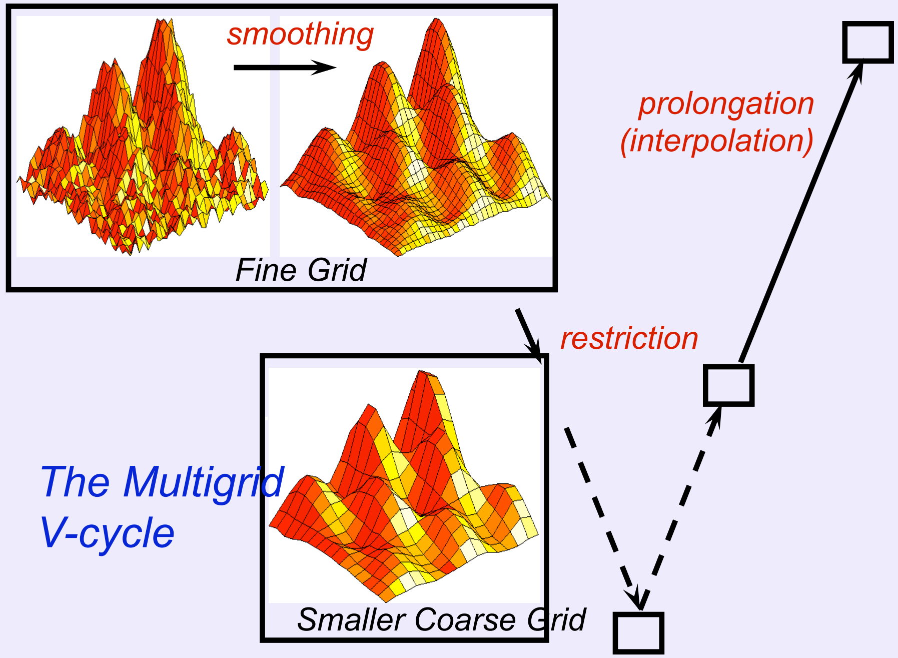

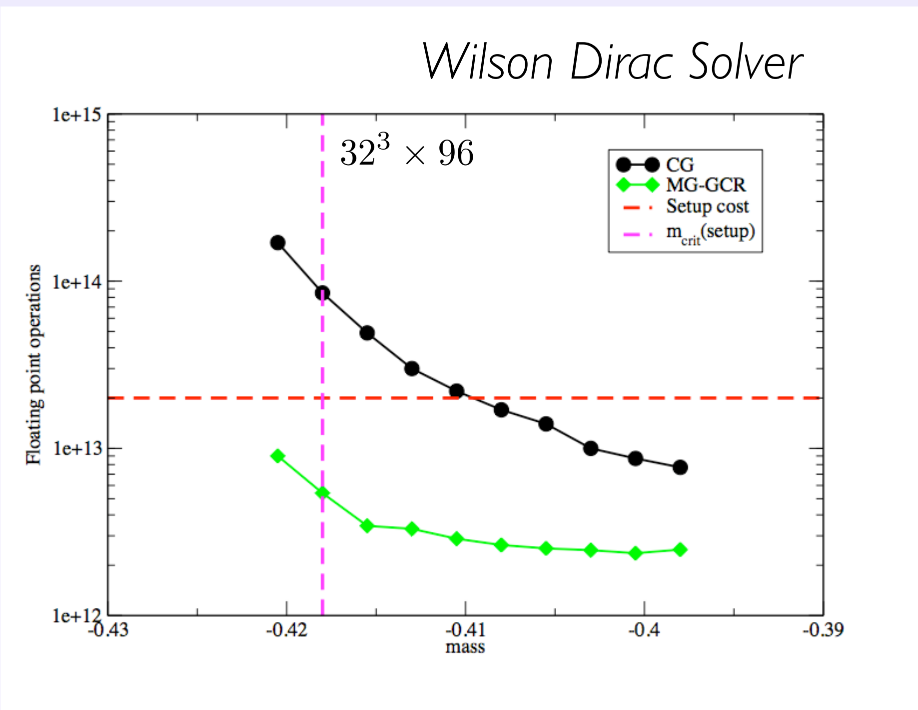

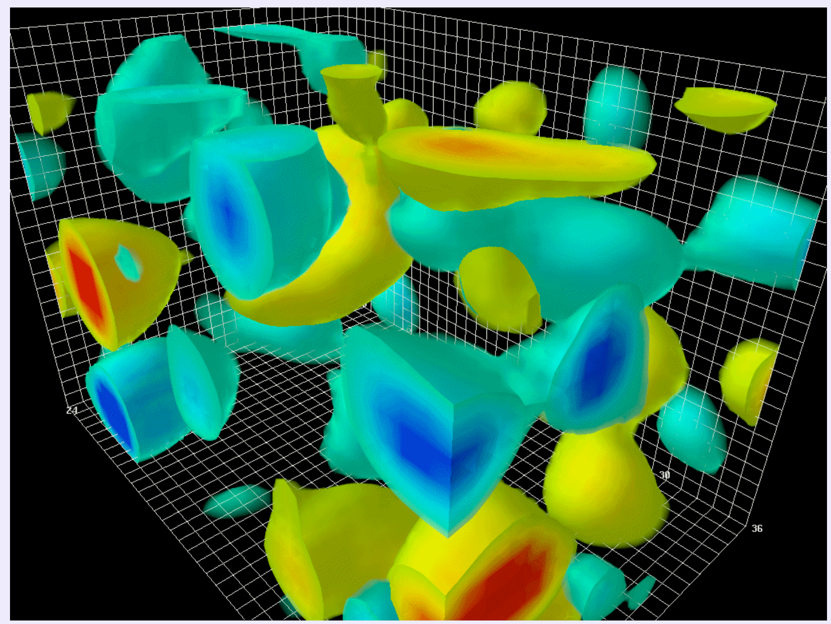

MG + GPU => 50 x speed up per $?

Broader impact

- Initial target applications in Lattice Field Theory and Computational Fluid Dynamics.
- But also to be a catalyst for local researchers in many fields to explore GPU architectures and share experience and methods.
- Access to faculty, post docs, graduate and undergraduate students in nano technology, chemistry, biological modeling, medical imaging, etc.
- Educational context and impact is crucial advantage of university based experimental GPGPU cluster!

FINI

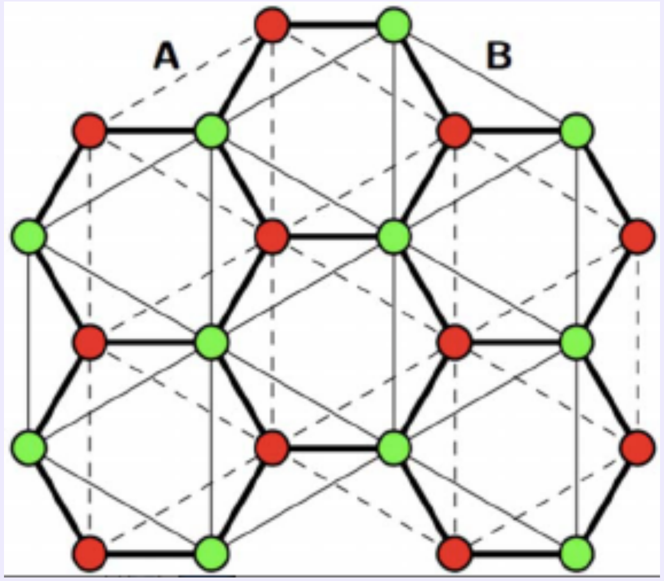

Launch workshop for the NSF-funded
Experimental GPU cluster for fundamental physics

organized by Lorena Barba, Richard Brower, Claudio Rebbi

Thursday, Nov. 12

9–9:10am Welcome & Introductions, Claudio Rebbi

9:10–9:30am **Richard Brower**, Boston University
"Potential impact of the 'Experimental GPU cluster for Fundamental Physics' NSF grant"

9:30–10:20am **Hanspeter Pfister**, Harvard University
"High-throughput science"

10:20–10:50am *BREAK*

10:50am–11:40pm **David Luebke**, Nvidia
"Graphics hardware & GPU computing: past, present, and future"

11:40–12:10pm **Richard Edgar**, Harvard University
"Diesel-powered supercomputing"

12:10–12:40pm **David Kaeli**, Northeastern University
"Many-core acceleration in biomedical applications"

12:40–2pm *LUNCH – sandwiches will be provided*

2–2:30pm **Michael Clark**, Harvard University
"GPU mixed-precision linear equation solver for lattice quantum chromodynamics, QCD"

2:30–3pm **Andreas Klöckner**, Brown University
"GPU metaprogramming using PyCUDA: methods and applications"

3–3:30 **Nicolas Pinto**, MIT
"Unlocking brain-inspired computer vision"

3:30–4pm *BREAK*

4–4:30pm **Bharat Sukhwani**, Boston University
"High-performance computing using GPUs: examples from computational biology"

4:30–5pm **Lorena Barba**, Boston University
"Toward GPU-accelerated meshfree flow simulation"

5–5:30pm **Tsuyoshi Hamada**, Nagasaki University
"42 TFlops N-body simulations on GPUs"

5:30pm Panel Discussion

Tutorial, Friday Nov. 13

9am–12pm "CPU/GPU programming with CUDA"

Performance Per MFLOP

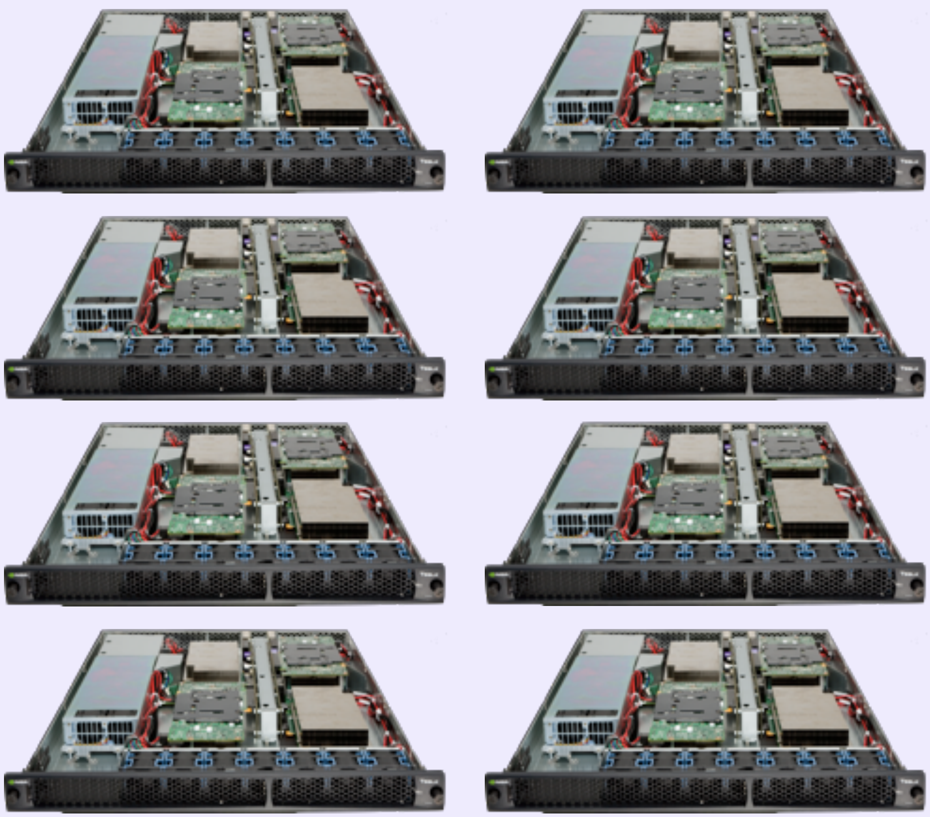

=

Performance Per Watt

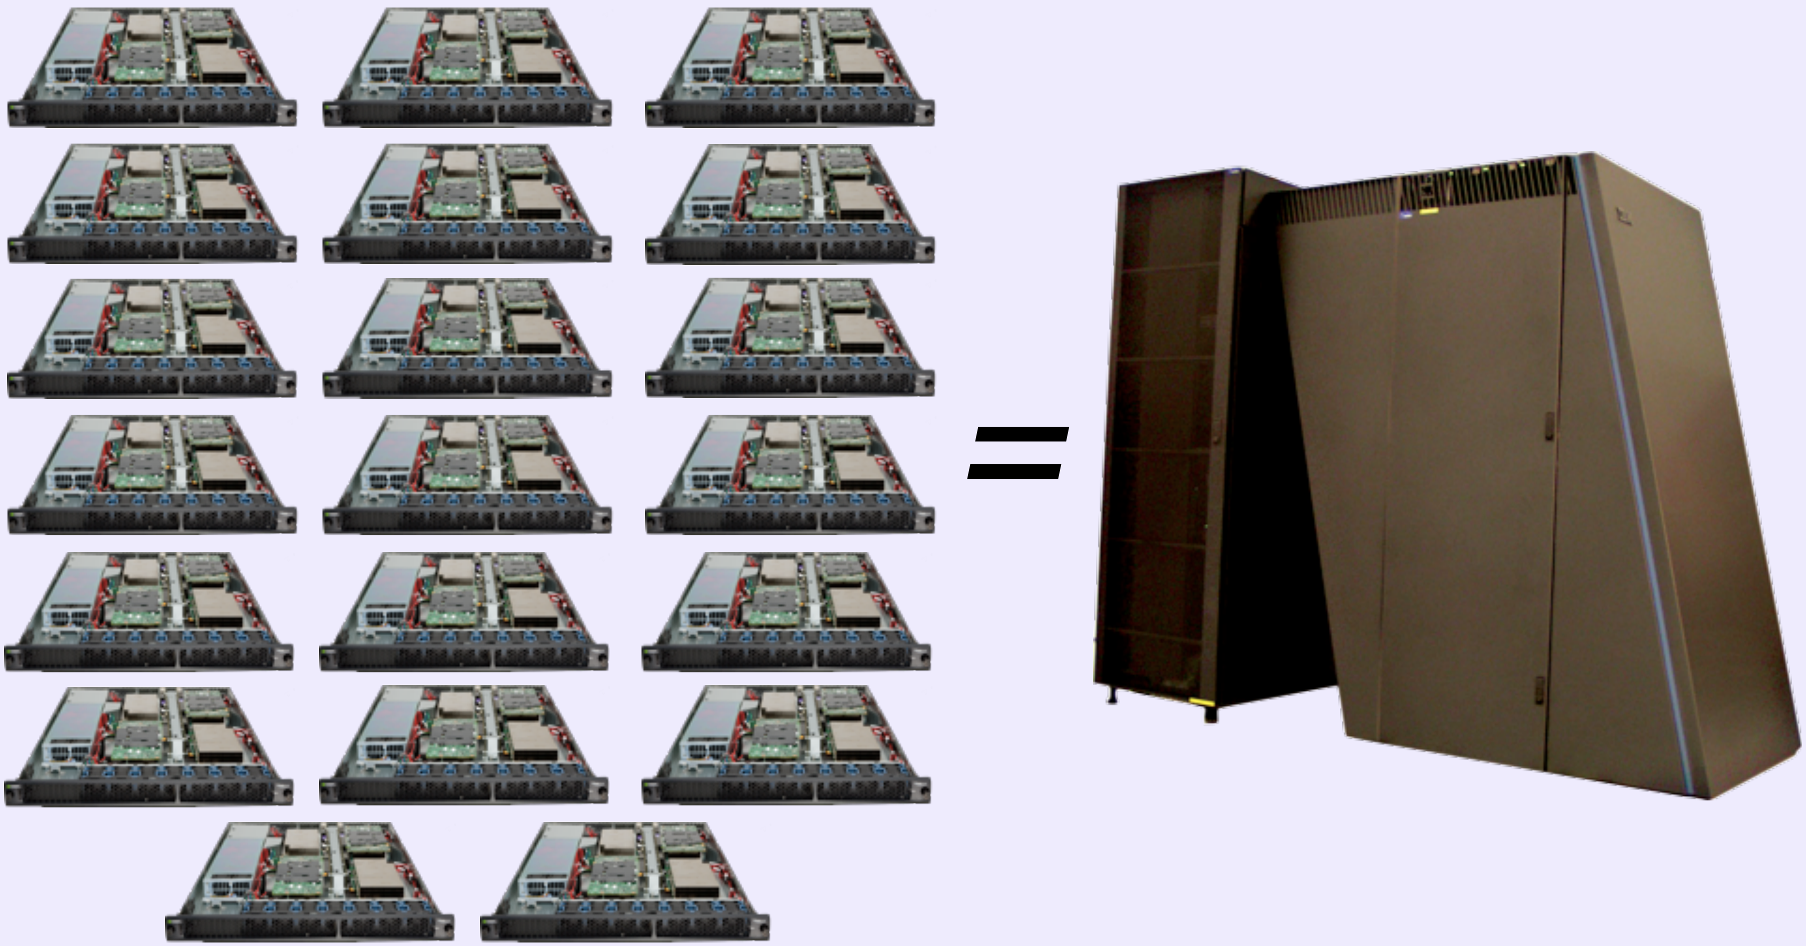

Performance Per $

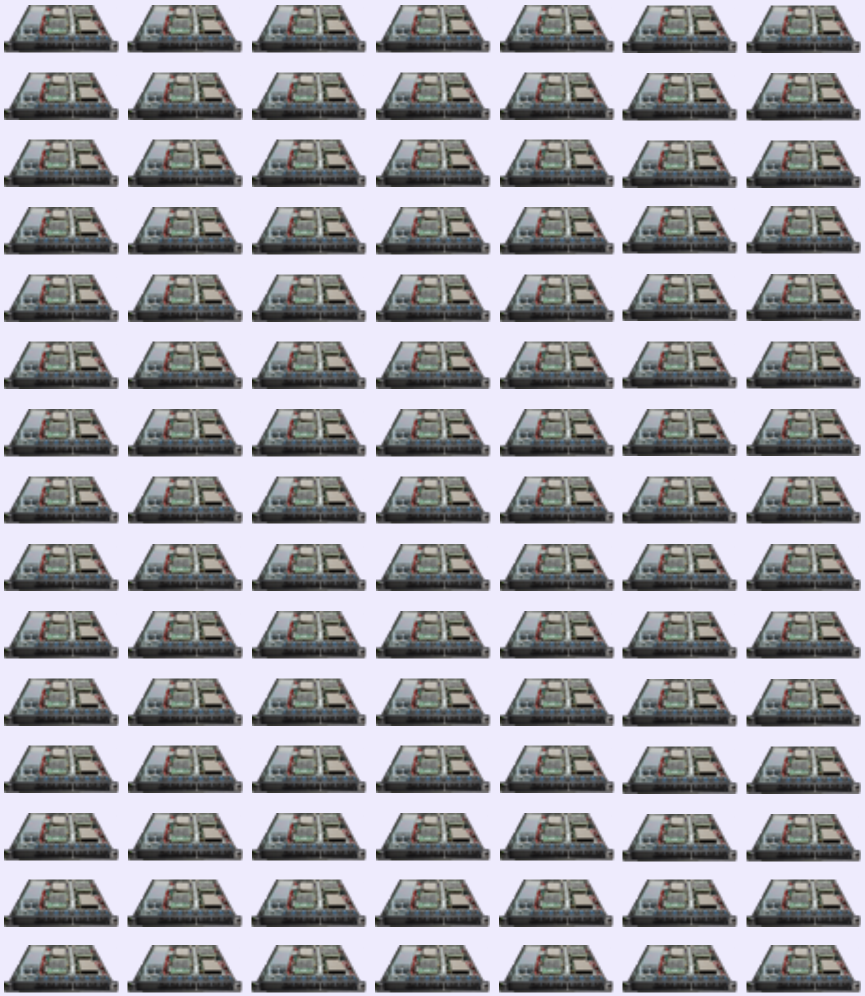

=

I. SciDAC API for Lattice QCD

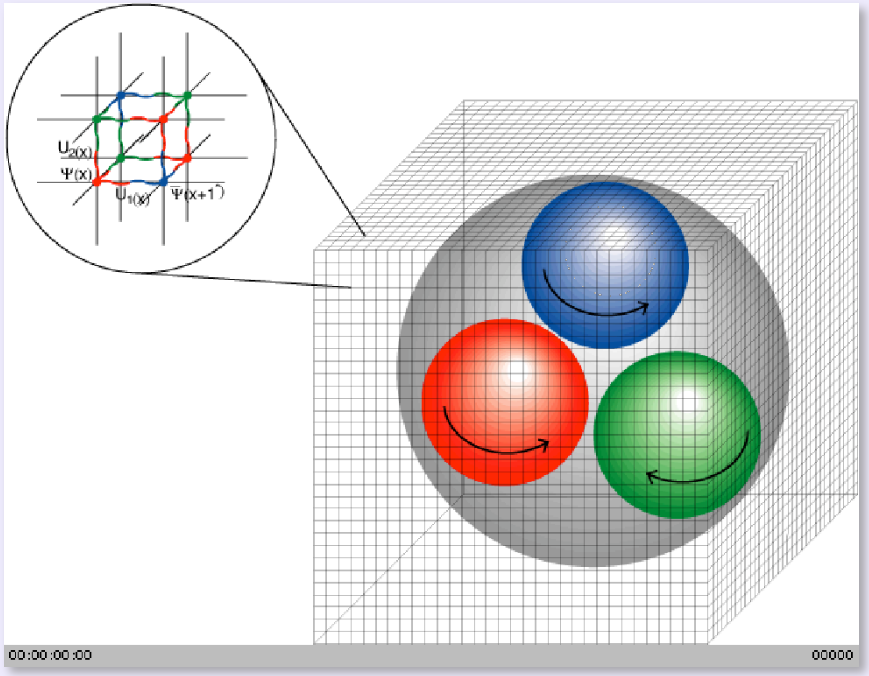

- (Perfect) Load Balancing: *Uniform periodic lattices & identical sublattices per processor.*
- (Complete) Latency hiding: *overlap computation / communications*
- Data Parallel: *operations on Small 3x3 Complex Matrices per link.*
- Critical kernels : *Dirac Solver, HMC forces, etc. ... 70%-90%*

Lattice Dirac operator:

$$[D\Psi]^i_\alpha(x) = \frac{1}{2a}\sum_\mu [U^{ij}_\mu(x)\gamma^{\alpha\beta}_\mu \Psi^j_\beta(x+\mu) - h.c] + \cdots$$

Lattice QCD was formulated in 1974: Kenneth G. Wilson, Confinement of Quarks. Phys.Rev.D10:2445-2459,1974.

Nvidia Tesla Quad S1070 1U System $8K

Processors	4 x Tesla T10P
Number of cores	960
Core clock	1.5 Hz
Performance	4 Teraflops
memory BW	16.0 GB
bandwidth	408 GB/sec
Memory I/0	2048 bit,800MHz
Form factor	1U (EIA 19" rack)
System I/O	2 PCIe x 16 Gen2
Typical power	700 W

Card	Cores	Bandwidth (GiB/s)	GFLOPS	Device Memory
8800 GTX	128	86.4	518.0	768 MB
Tesla C870	128	76.8	518.4	1.5 GB
GTX 280	240	141.7	933	1 GB
Tesla C1060	240	102	1000	4 GB

Spec for Fermi GPU!

GPU	G80	GT200	Fermi
Transistors	681 million	1.4 billion	3.0 billion
CUDA Cores	128	240	512
Double Precision Floating Point Capability	None	30 FMA ops / clock	256 FMA ops /clock
Single Precision Floating Point Capability	128 MAD ops/clock	240 MAD ops / clock	512 FMA ops /clock
Warp schedulers / SM	1	1	2
Special Function Units (SFUs) / SM	2	2	4
Shared Memory / SM	16 KB	16 KB	Configurable 48 KB or 16 KB
L1 Cache	None	None	Configurable 16 KB or 48 KB
L2 Cache	None	None	768 KB
ECC Memory Support	No	No	Yes
Concurrent Kernels	No	No	Up to 16

- Joining NVIDIA press conference was Oak Ridge National Laboratory who announced plans for a new supercomputer that will use NVIDIA GPUs based on the Fermi architecture.

Lack of Critical slowing down:

CG iteration count is insensitive to quark mass and lattice volume!

m_s (- 0.38922)

Lattice volumes

Mass:	$16^3 \times 64$	$24^3 \times 64$	$32^3 \times 96$
-.3980	40	40	41
-.4005	41	41	42
-.4030	42	42	43
-.4055	42	43	43
-.4080	43	44	45
-.4105	44	46	49
-.4130	45	49	52
-.4155	47	54	57
-.4180	50	62	89

Small increase is probably not significant?

Physical m^2_π

Chiral limit: $m^2_\pi = 0$

Nvidia’s Fermi GPU:
512 core x 32 floating = 16K bits