Multi-scale Algorithms for Lattice Gauge Theory

(Quantum ChromoDynamcs et al)

Richard C. Brower

(USQCD SciDAC Software Co-ord!)

Boston University

INT EXASCALE WORKSHOP

June 27, 2011 Seattle

QCD Multigrid Applied Math/Physics Collaboration

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

- CU Boulder
  - Tom Manteuffel
  - Steve McCormick
  - Marian Brezina
  - John Ruge
  - James Brannick
  - Christian Ketelsen
  - Scott MacLachlan
- Lawrence Livermore
  - Rob Falgout
- Columbia
  - 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 U

Mike Clark

- *Ron Babich*

- *INT Seattle*
  - *Saul Cohen*

Related Talks coming (I think)

- Rob Falgout: Multigrid Methods
- Andrew Pochinsky: QLUA & Software tools
- Pavlos Vranas: LQCD on the BG/P
- Kostas Orginos Eigenvalue Deflation
- Balint Joo: QUDA: Domain Decomposition on GPUs
- James Osborn (Next week) : MG software/Disco Diag
- And many more I sure!

Prototype for Exascale swimlanes

MG/Multiscale Narrative for QCD

- Past: 1990-2010
  The good olde days of " MPI Data Parallel code" for SPMD algorithms "easily" mapped on to uniform architecture of Beowulf clusters.

- Present 2010-2015
  Transition to high resolution multi-scale QCD and MPI network of threaded multi-core micros

- *Future: 2015-2020*
  Profound challenges to adapt efficient multi-scale algorithm to a complex heterogeneous architecture with deep memory hierarchies.

K. Wilson (Lattice 1989 Capri)

"One lesson is that 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*

1. *1930+50 = 1980*
2. *0.1 flops ➔ 10 Mflops*
3. *Gaussian Basis functions*

- *ab initio QCD*

1. *1980 + 50 = 2030?**
2. *10 Mflops ➔ 1000 Tflops*
3. *Clever Collective Variable?*

**Sustained on LGT*

Hopefully sooner but needs cost @ $10/Gflops!

Lattice QCD – extremely uniform

Dirac operator:

$$D\psi(x) = \sum_{\mu} \left(\partial_{\mu} + igA_{\mu}(x) \right) \psi(x)$$

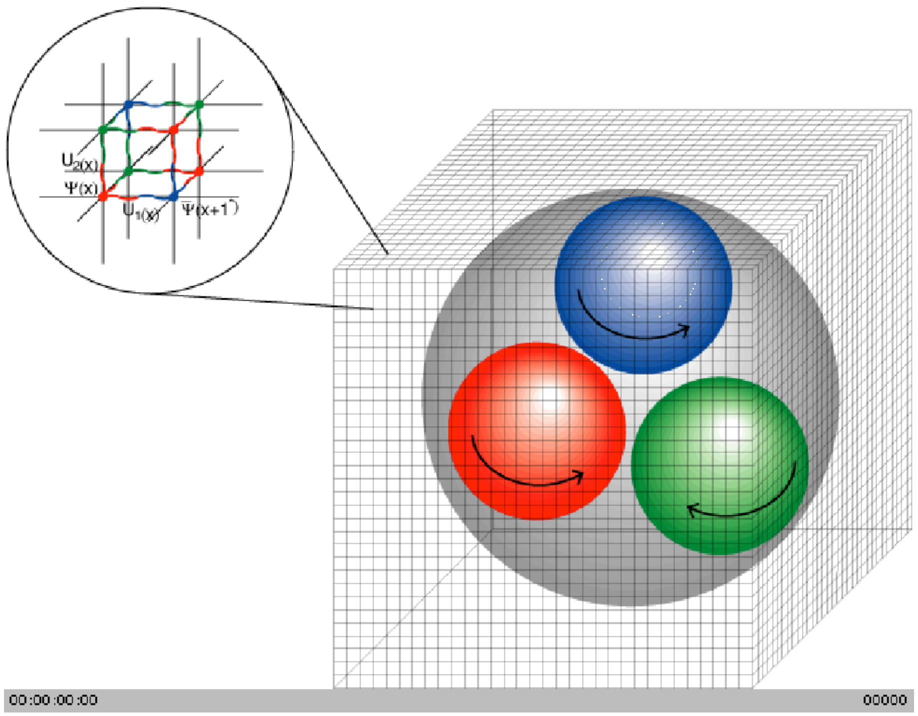

- Periodic or very simple boundary conditions
- SPMD: Identical sublattices per processor

Lattice Operator:

$$D\psi(x) = \frac{1}{2a} \sum_{\mu} \left[U(x)\psi(x+\hat{\mu}) - U^{\dagger}(x-\hat{\mu})\psi(x-\hat{\mu}) \right]$$

Lattice QCD code is simple/easy

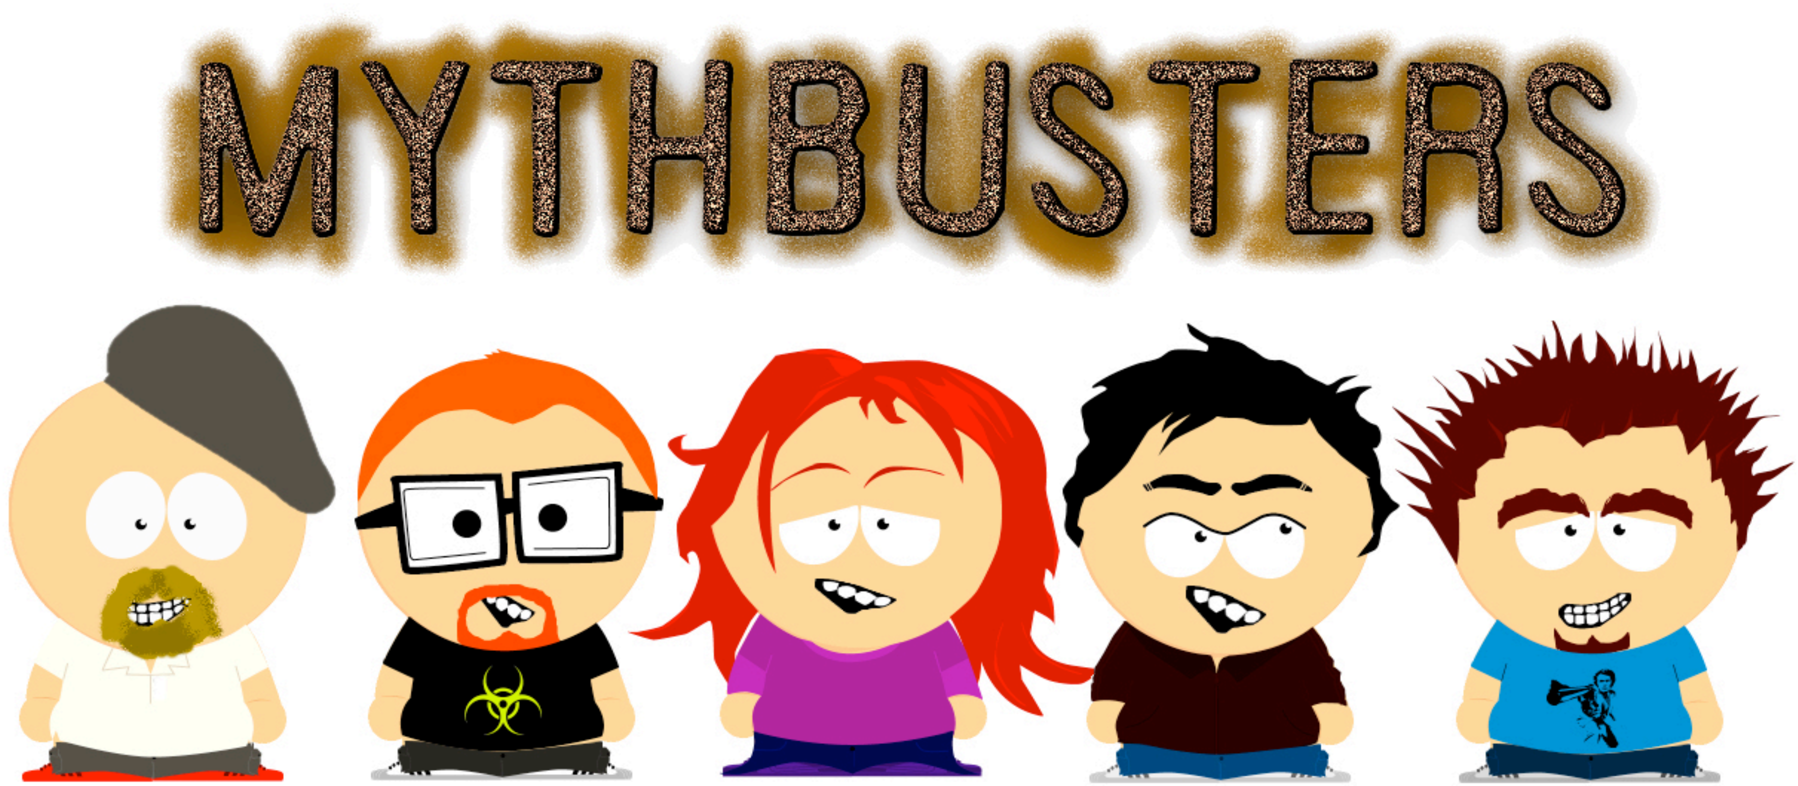

THIS MYTH IS BASED ON INDEQUATE RESOURCE

Accurate calculations MUST resolve multiples scales

AND

Exascale will require adaptation to heterogeneous Architecture

JOE BARBER 2008

Current SciDAC LGT Software Stack

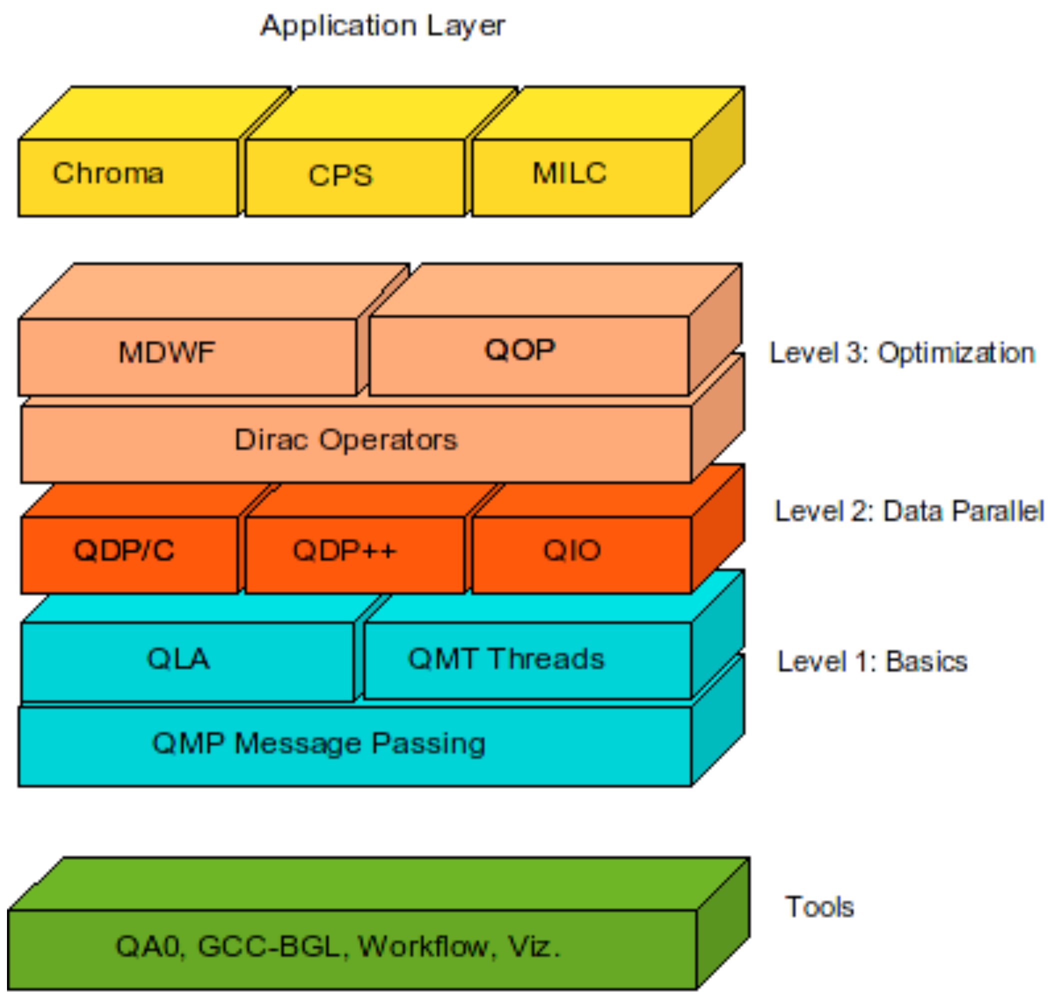

+ tools from collaborations with other SciDAC projects e.g. PERI

Chroma = 4856 files

CPS = 1749 files

MILC = 2300 files

QLA/perl = 23000 files

QUDA/python = 221 files

1990 Beowulf-clusters
Big Iron supers vanquished by the killer micro


BNL:QCDOC
JLab
FNAL
BG/L

2010: Power/Communication Wall kills Beowulf-clusters

Road to EXAFLOPS take back to the Future: Clusters of Many-core Vector Machines: Data control is the game.

Potential *Exascale* Architecture *Swimlanes*

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

Source: Andy White and Rick Stevens talk on “A decadal DOE plan for providing exascale applications and technologies for DOE mission needs” DOE ASCAC meeting, March 2010

IBM BG/Q++

CRAY Titan ++

exascale swimlanes

Physics:Higher resolution QCD

- Lattice scales:
  - 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

- Consequences:
  - Increasing ill-conditioned Dirac operator
  - Suffer from worse critical slowing down (CSD)
  - O(100^4) lattice volume
  - 1/4 Terabyte file for a single Dirac propagator

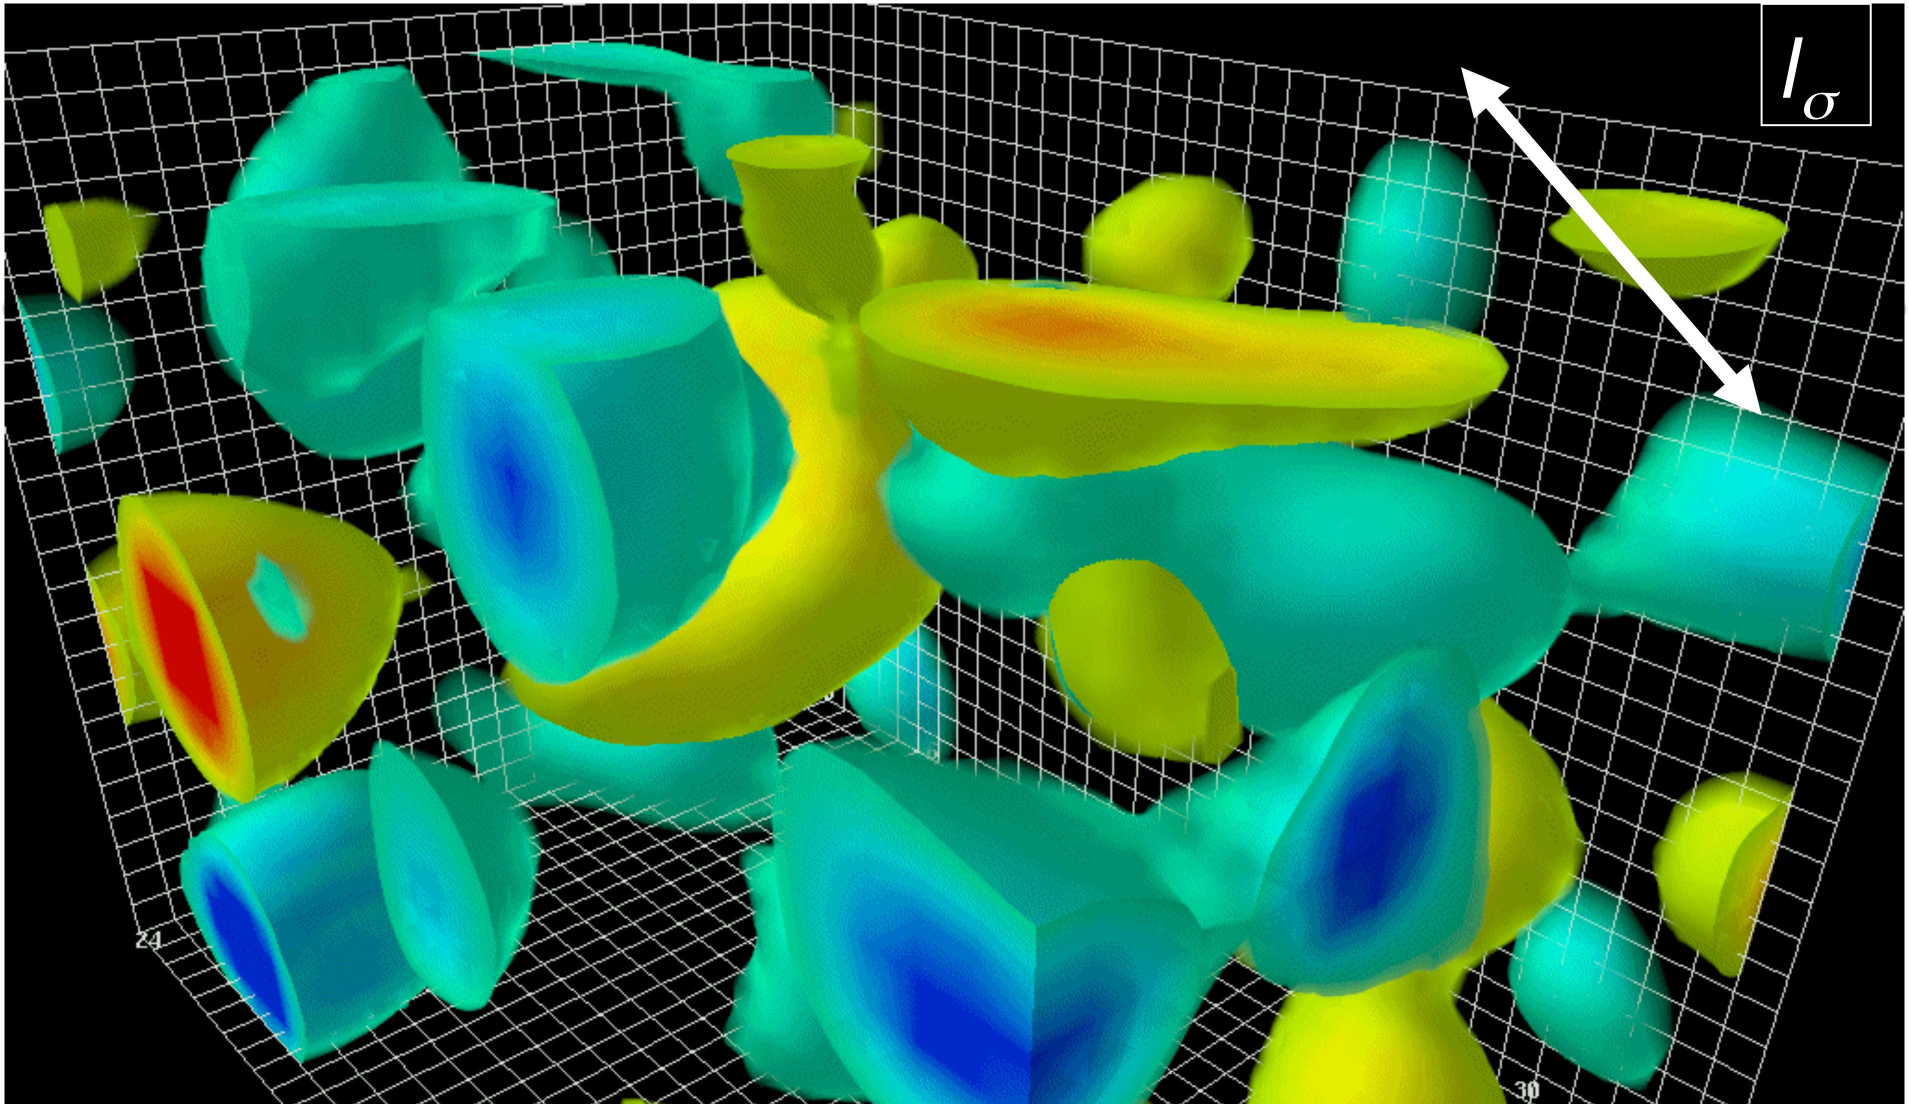

Classical QCD (with zero mass quarks) has no scale. BUT spontaneous Conformal symmetry breaking magically gives the proton mass scale and the "Higgs" gives quark masses!

Many more LGT mass scales to come

- 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
- ETC.

Multi-grid Free Laplace/Poisson

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

- Define the Prolongation op P

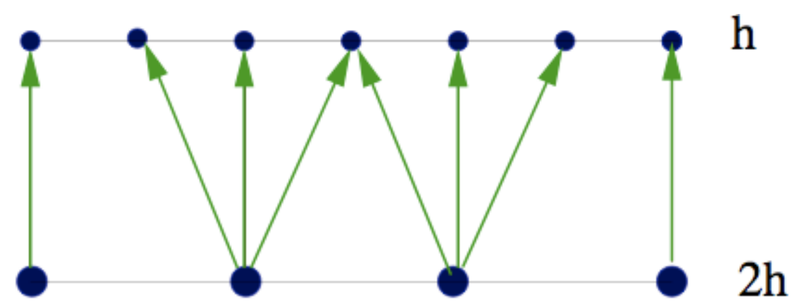

- Define the Restriction op R = $P^\dagger$

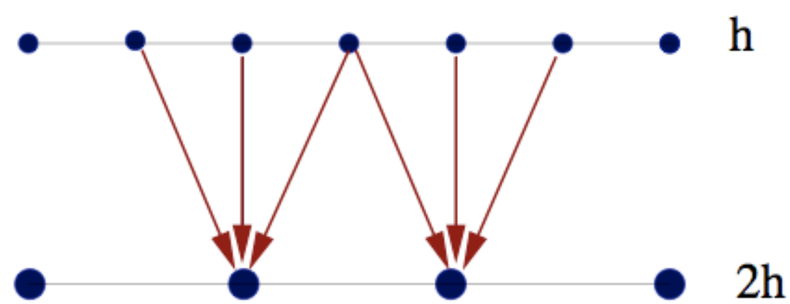

- Operator on **coarse** space

$$A_c = P^\dagger A P$$

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

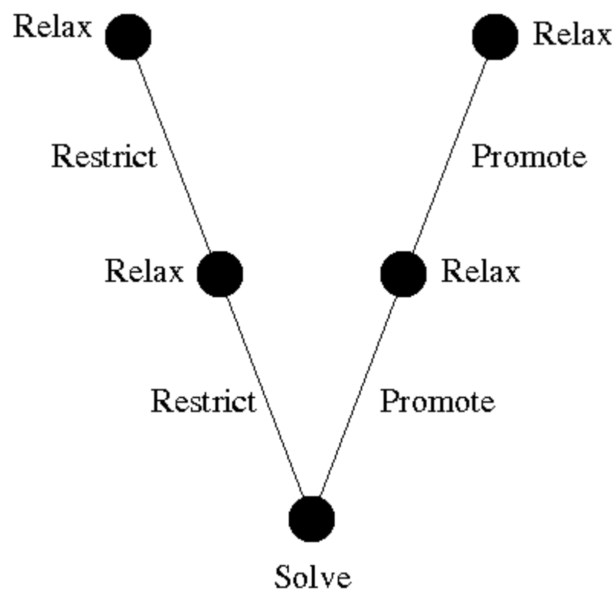

Dirac Eq. $m\psi(x) - \sum_{\mu=1}^{d} [\gamma_\mu \frac{\partial}{\partial_{x_\mu}} + iA_\mu(x)]\psi(x) = 0$

Naive Operator

Anti Hermitian: pure imaginary spectrum

$$D^{\text{naive}}_{x,x'} = m\,\delta_{x,x'} - \sum_\mu \frac{1}{2}\left[\gamma_\mu U_{x,\mu}\delta_{x+\hat{\mu},x'} - \gamma_\mu U^\dagger_{x,\mu}\delta_{x,x'+\hat{\mu}}\right]$$

- Straightforward discretization of continuum operator
- Results in 2^d doubler modes

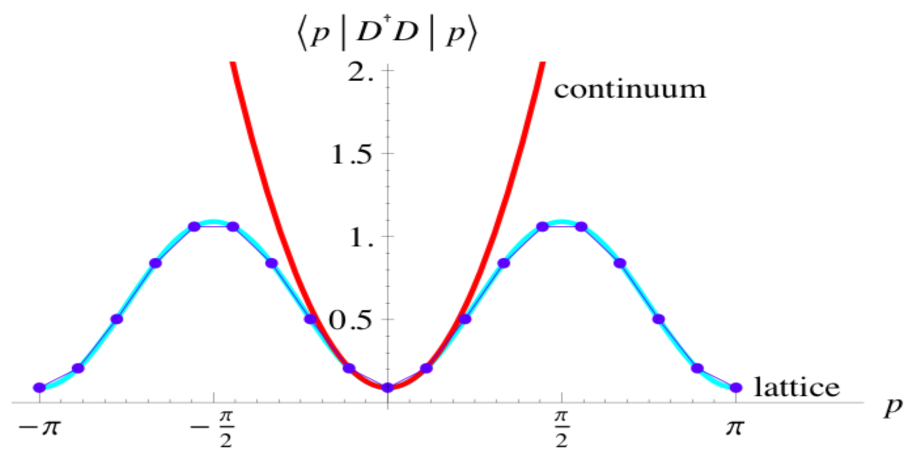

Naive central difference decouples even/odd site and leads to spurious low modes

Wilson Operator

Stabilize with added Laplace operator

$$D^{\text{Wilson}}_{x,x'} =$$

$$(m+d)\,\delta_{x,x'} - \sum_{\mu} \frac{1}{2}\left[(1+\gamma_\mu)U_{x,\mu}\delta_{x+\hat{\mu},x'} + (1-\gamma_\mu)U^{\dagger}_{x,\mu}\delta_{x,x'+\hat{\mu}}\right]$$

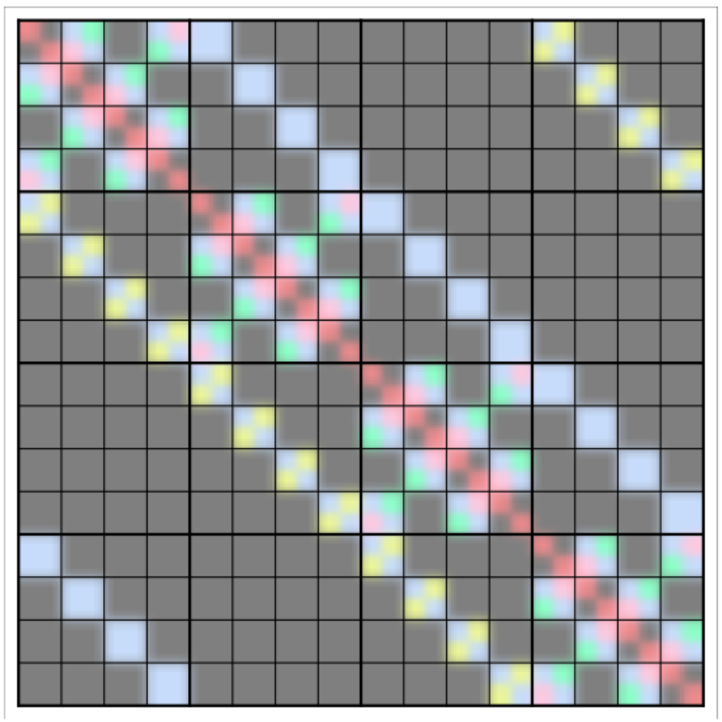

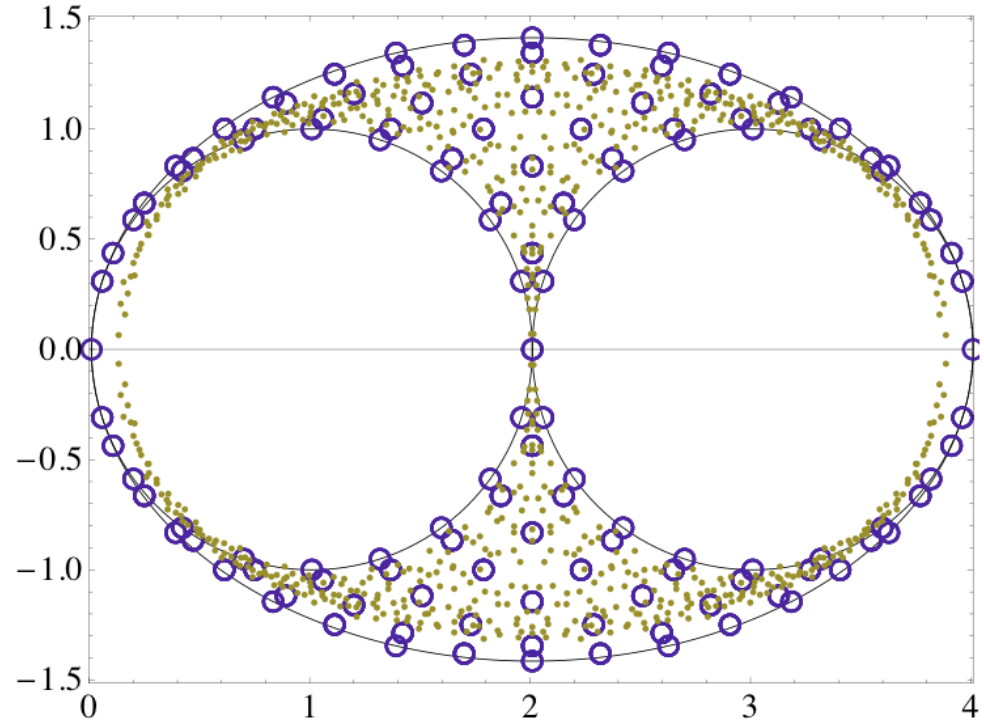

QCD MG attempts in 1990's:

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.*

SUPERCOMPUTER COMPUTATIONS RESEARCH INSTITUTE

PROJECTIVE MULTIGRID FOR WILSON FERMIONS

by

Richard C. Brower, Robert G. Edwards, Claudio Rebbi, and Ettore Vicari

FSU-SCRI-91-54

2x2 Blocks for 2-d U(1) Dirac
(Naive scaling of free problem)

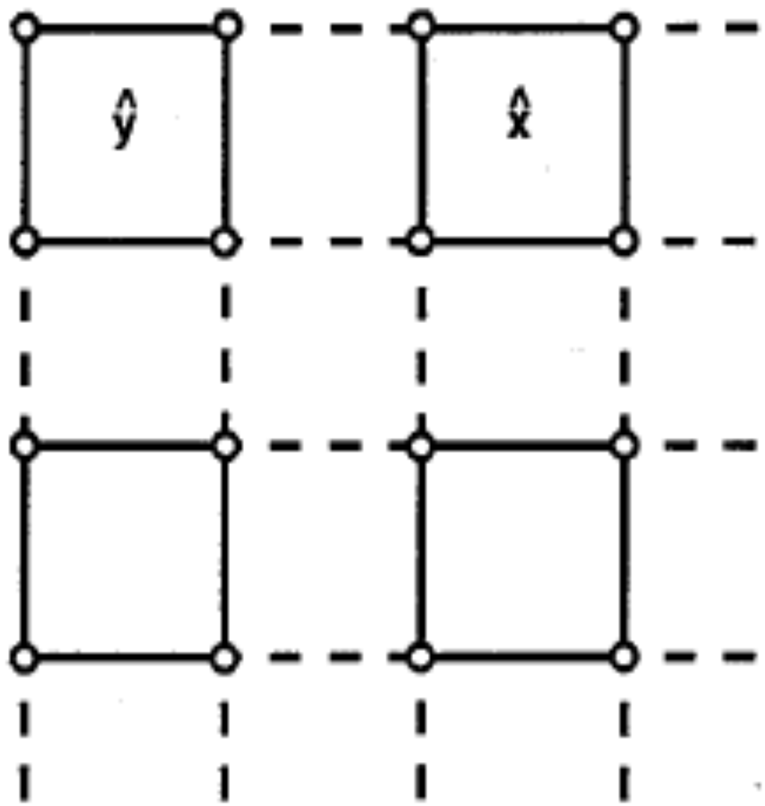

2-d Lattice, $U_\mu(x)$ on links $\Psi(x)$ on sites

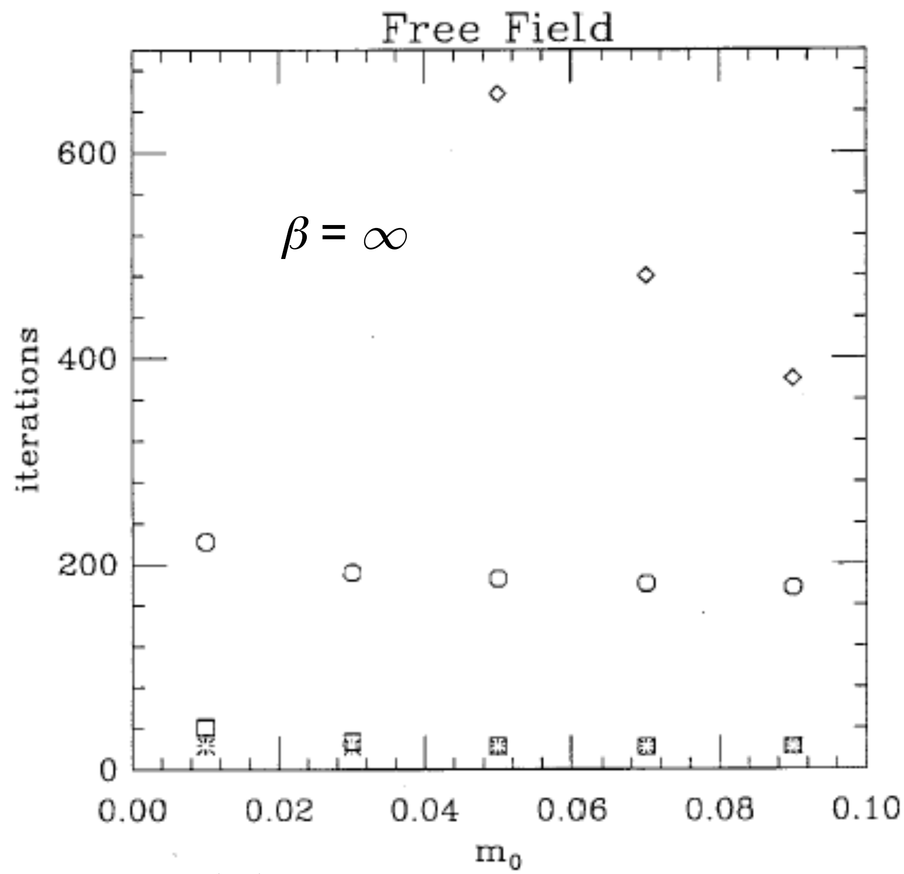

$$U_\mu(x) = 1$$

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

Universal critical slowing: $\tau = F(m\, l_\sigma)$

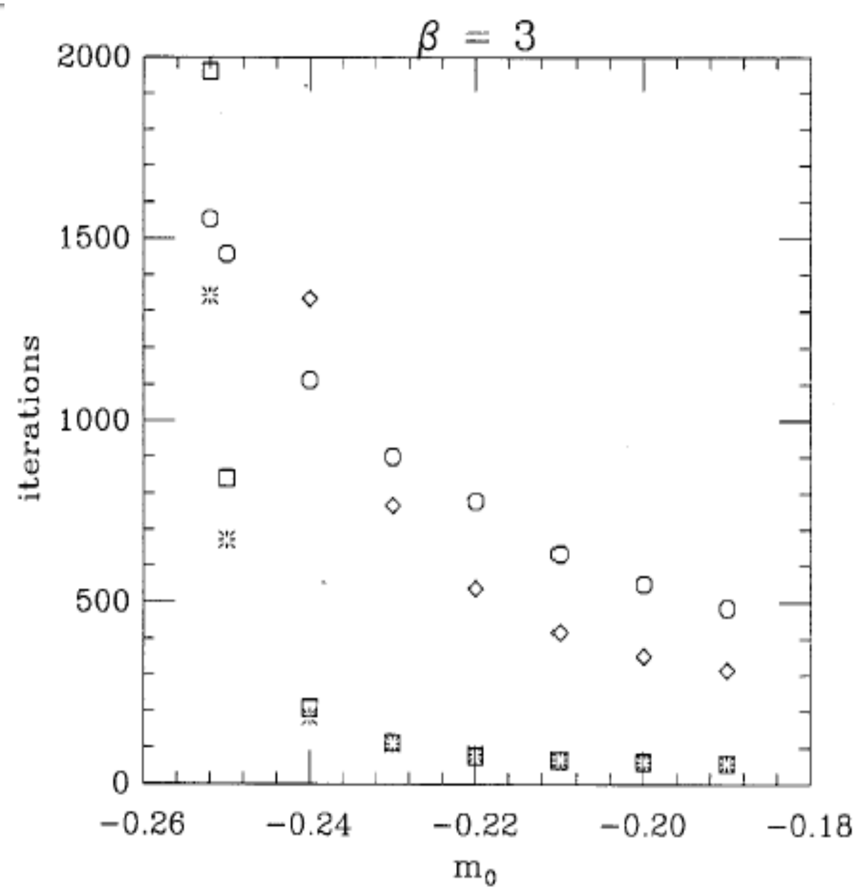

Gauss-Jacobi (Diamond), CG(circle), 3 level (square & star)

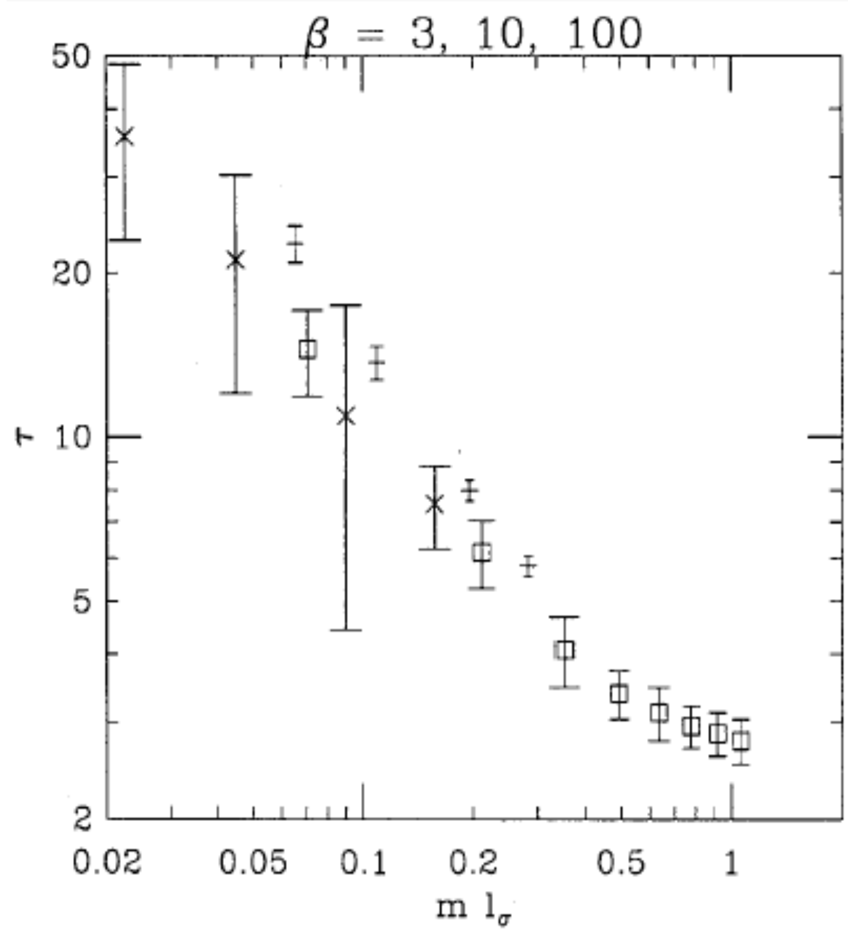

β = 3 (cross) 10(plus) 100(square)

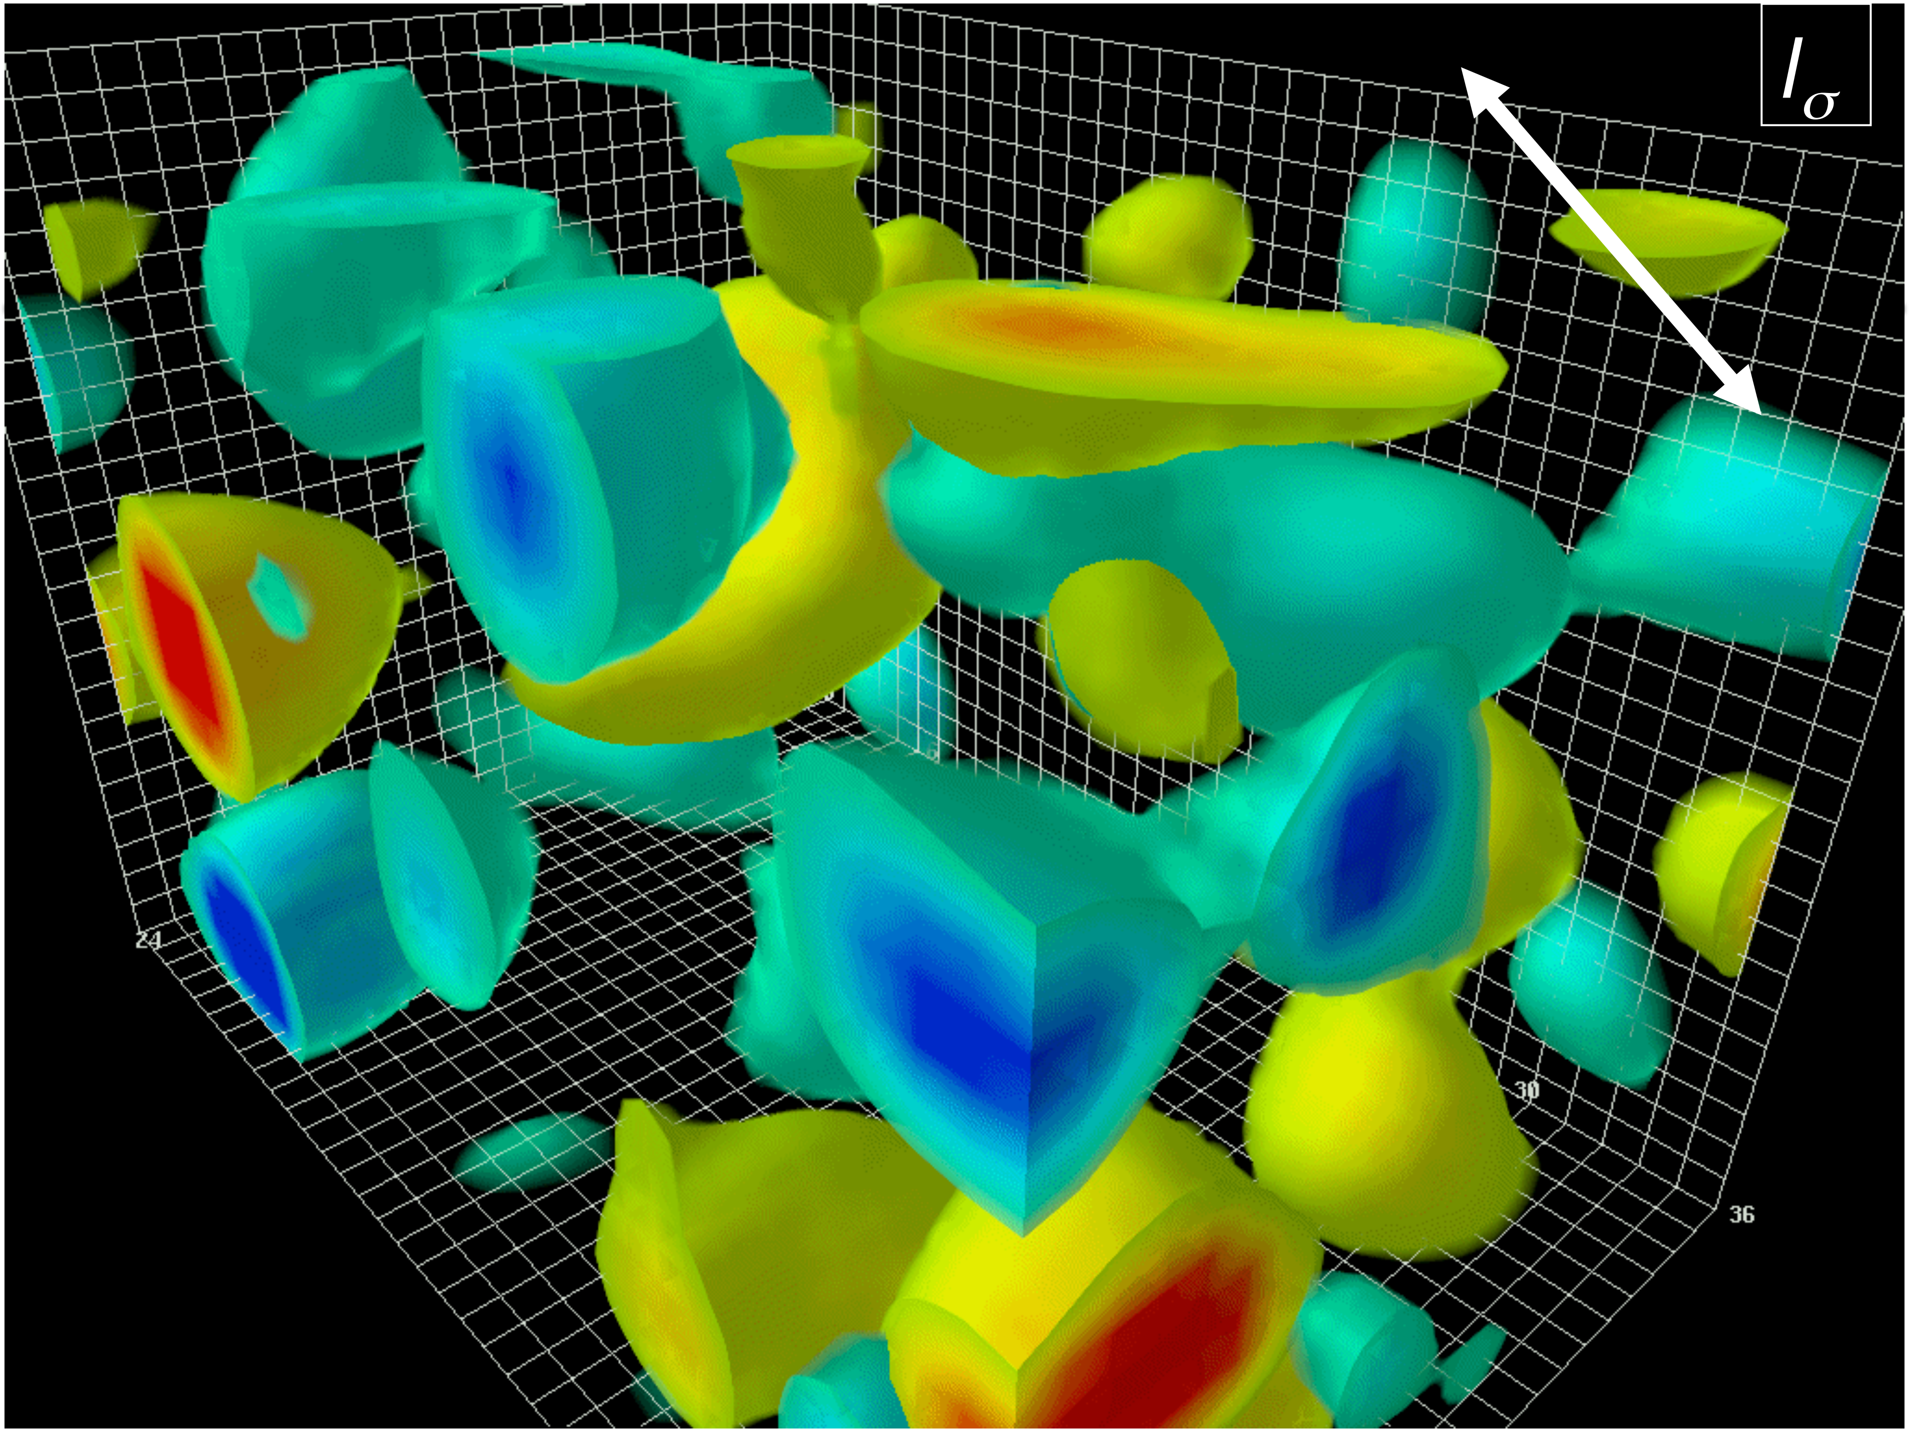

l_σ
24
30
36

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

- Maintain Gauge invariance
- Maintain γ_5 Hermiticity
- Local adaptive blocking: Projective MG
- Partial success (RG) weak coupling
- Learn from Failure

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

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

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!

- 2005 AMG: First "project" onto near null vectors (bad guys). Second "partition" into coarse grid.

Adaptive Smooth Aggregation Algebraic Multigrid

The Multigrid V-cycle

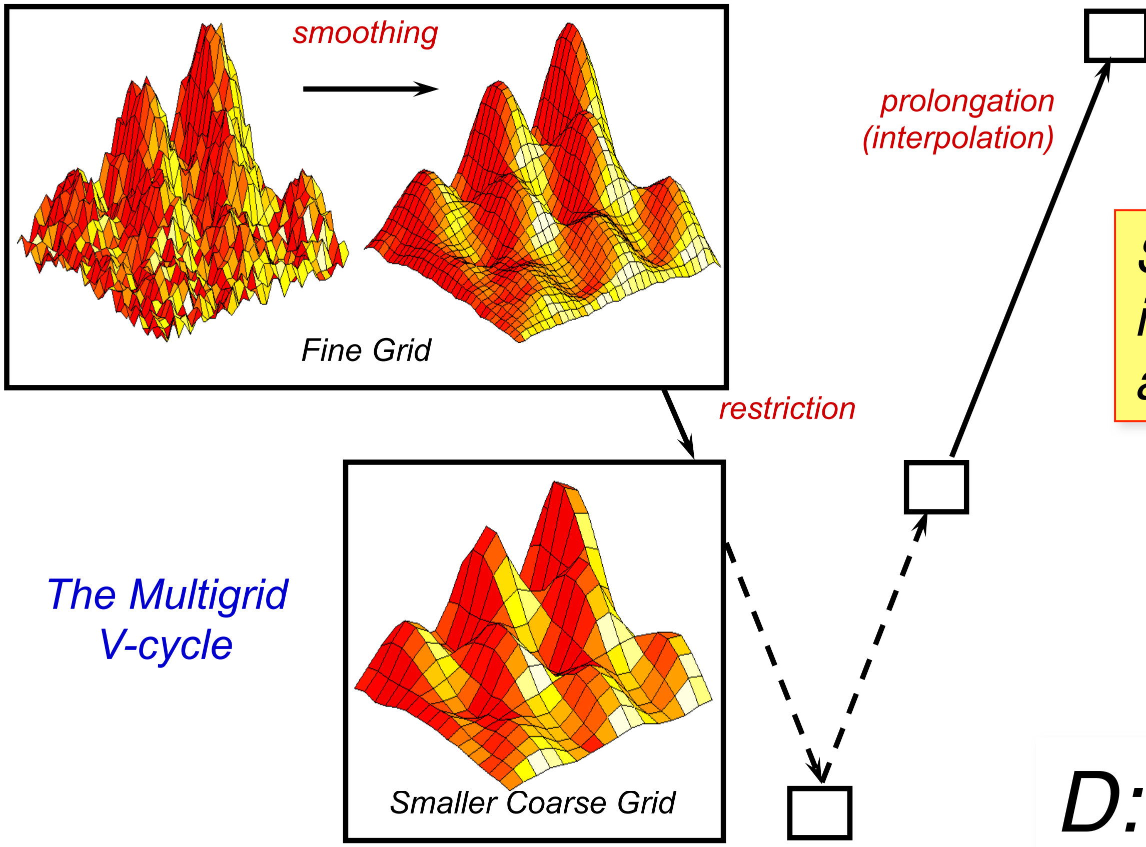

Spilt the vector space into near null space $\mathcal{S}$ and the complement $\mathcal{S}_\perp$

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

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

3 approaches to Near Null Space

1. "*Deflation*": N_ν *exact eigenvector projection*

$$
\begin{aligned}
\Pi &= 1 - |\psi_\lambda\rangle\langle\tilde{\psi}_\lambda| = 1 - D|\psi_\lambda\rangle\frac{1}{\lambda}\tilde{\psi}_\lambda| \\
&= 1 - D|\psi_\lambda\rangle\langle\tilde{\psi}_{\lambda'}|D|\psi_\lambda\rangle^{-1}\langle\tilde{\psi}_{\lambda'}| \\
&\to 1 - DP(RDP)^{-1}R.
\end{aligned}
$$

2. "*Inexact deflation*" *plus Schwarz DD (Lüscher)*

3. *Multi-grid uses coarse level for preconditioning*

Little Dirac: $$D_{cc} \equiv \hat{D} = P^\dagger D P$$

$$D_{Schur} = [D - DP\frac{1}{P^\dagger D P}P^\dagger D]_{ff}$$

Zero on coarse space

- *2 & 3 use the same splitting* $\mathcal{S}$ *and* $\mathcal{S}_\perp$

2-level Multigrid Cycle (simplified)

- Smooth: $x' = (1 - D)\, x + b \quad \Rightarrow r' = (1 - D)\, r$
- Project: $D_c = P^\dagger D P \quad \& \quad r_c = P^\dagger r$
- Solve: $A_c e_c = r \quad \Rightarrow \quad e_c = A^{-1}{}_c P^\dagger r$
- Prolongate $e = P$ e_c
- Update $x' = x + e \quad \Rightarrow r' = b - D(x + e)$

  $= [1 - D P (P^\dagger D P)^{-1} P^\dagger]\, r$

RESULT: D is "preconditioned" by $M = P (1/P^\dagger D P) P^\dagger$

$(1/M) D x = (1/M)\, b \quad \Rightarrow r' = [1 - D (1/M)]\, r$

Note since $P^\dagger r' = 0 \Rightarrow$ full (exact) deflation on $\mathcal{S}$

P†:fine → coarse (non-square matrix†)

(fine lattice vector space) *(coarse lattice vector space)*

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

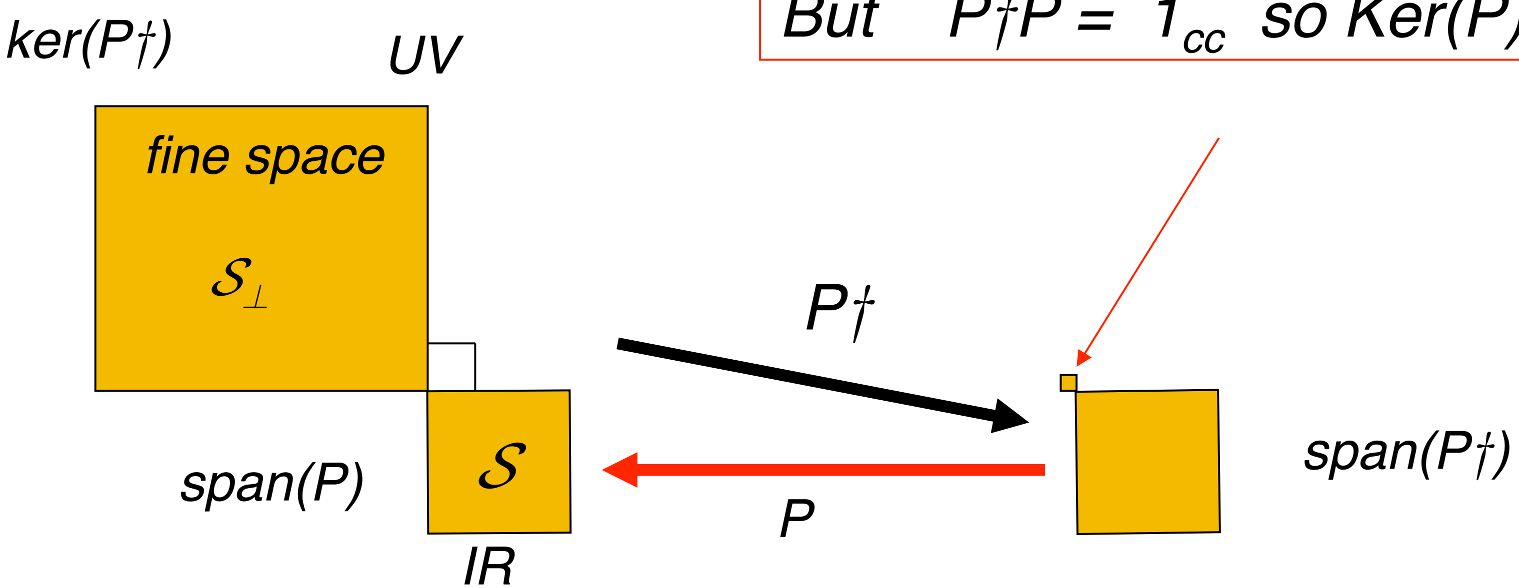

$\mathcal{S}$= span(P) = Image(P†)

rank(P) = rank(P†) =dim($\mathcal{S}$)= $N_\nu N_B = 2N_\nu L4/44$

† See Front cover of Gilbert Strang's undergraduate text !

Petrov-Galerkin oblique projection

$$D_{Schur} = [D - DP\frac{1}{P^\dagger DP}PD] = [1 - DP\frac{1}{P^\dagger DP}P]D$$

$$P^\dagger D_{Schur} = D_{Schur}P = 0$$

The coarse operator : $\hat{D} = D_{cc} = P^\dagger DP$

The projection op to coarse space : $PP^\dagger$

But the Schur complement use “oblique projects:

$$\Pi_L^\dagger = DP\frac{1}{P^\dagger DP}P^\dagger \quad , \quad \Pi_R = P\frac{1}{P^\dagger DP}P^\dagger D$$

$$D_{Schur} = D - \Pi_L^\dagger D\Pi_R = (1 - \Pi_L^\dagger)D(1 - \Pi_R)$$

General Problem: D ψ = b

- *"split" vector space into:*
  - *near null* $D\,\mathcal{S} \simeq 0$ & *Complement* $\mathcal{S}_\perp$
- *Schur decomposition (of course) does this!*
  - *Coarse = near null (IR) , Fine = complement (UV)*

Splitting is essential the idea of the Schur complement:

$$\begin{bmatrix} D_{cc} & D_{cf} \\ D_{fc} & D_{ff} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ D_{fc}D_{cc}^{-1} & 1 \end{bmatrix} \begin{bmatrix} D_{cc} & 0 \\ 0 & D_{ff} - D_{fc}D_{cc}^{-1}D_{cf} \end{bmatrix} \begin{bmatrix} 1 & D_{cc}^{-1}D_{cf} \\ 0 & 1 \end{bmatrix}$$

Implies $$D^{-1} = \begin{bmatrix} 1 & -D_{fc}D_{cc}^{-1} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} D_{cc}^{-1} & 0 \\ 0 & D_{Schur}^{-1} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -D_{cc}^{-1}D_{cf} & 1 \end{bmatrix}$$

with $D_{Schur} \equiv D_{fc}D_{cc}^{-1}D_{cf}$

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 block D to avoid higher complexity.
  - Multigrid is recursive to multi-levels.
  - Must preserves Gauge invariance and γ_5 ($[\gamma_5, P] = 0$)
- Current benchmarks for Wilson-Dirac Operator:
  - Asym $V = 16^3 \times 64$, $24^3 \times 64$, $32^3 \times 96$ ($N_f = 2$, 400MeV pion)
  - $N_\nu = 20$ null vectors ${\psi^s}_x$ 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 will give more details next week!).

Results 24^3x128

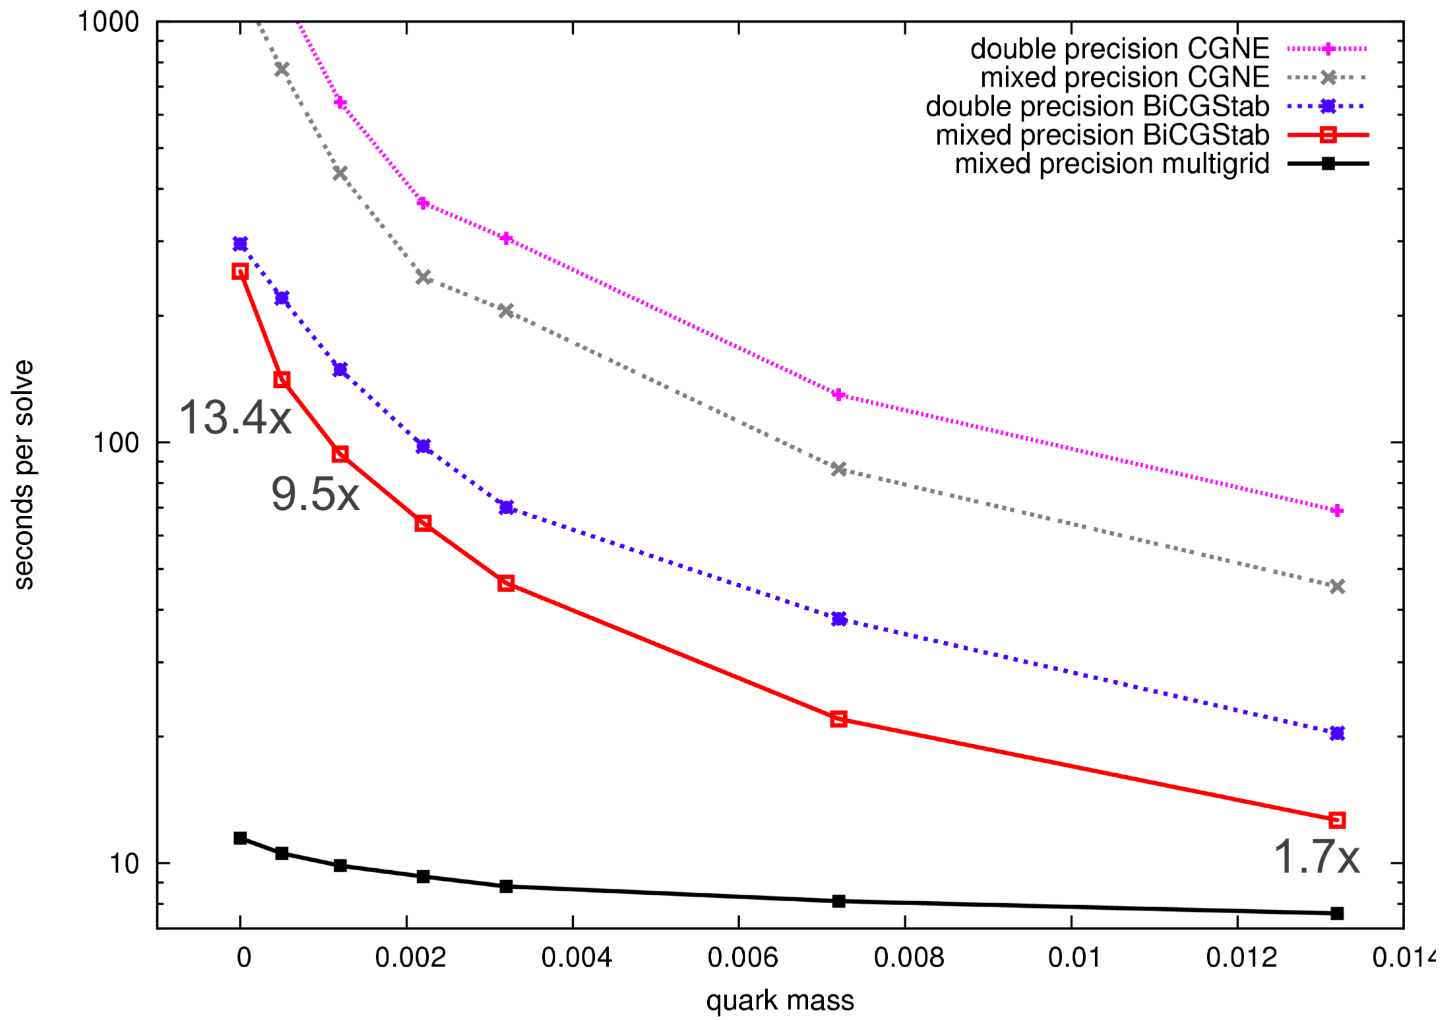

"Adaptive multigrid algorithm for the lattice Wilson-Dirac operator" R. Babich, J. Brannick, R. C. Brower, M. A. Clark, T. Manteuffel, S. McCormick, J. C. Osborn, and C. Rebbi, *Phys. Rev. Lett. 105, 201602 (2010).*

Lack of Critical slowing down:

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

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

m_s (- 0.38922)

Small increase is probably not significant?

physical $m2_\pi$

Chiral limit: $m2_\pi = 0$

"Adaptive Multigrid Algorithm for Lattice QCD" J. Brannick, R. C. Brower, M. A. Clark, J. C. Osborn, C. Rebbi, Phys. Rev. Lett. 100, 041601 (2008)

Results 32^3x256

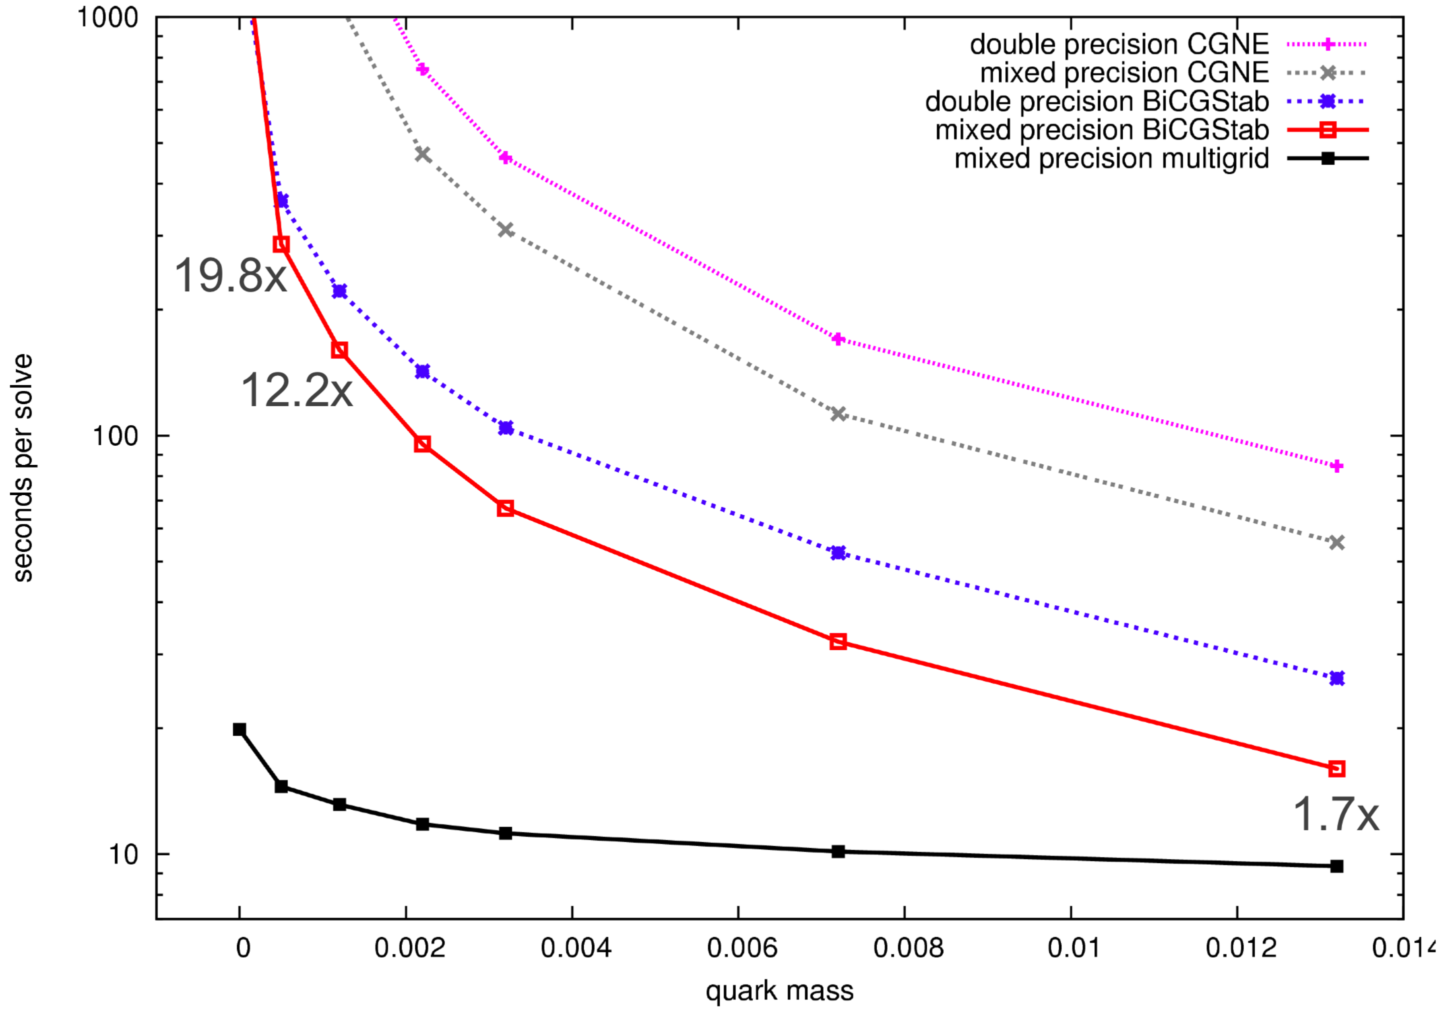

Surely there *is more tuning to be done with even greater speed up!*

2 level vs 3 level

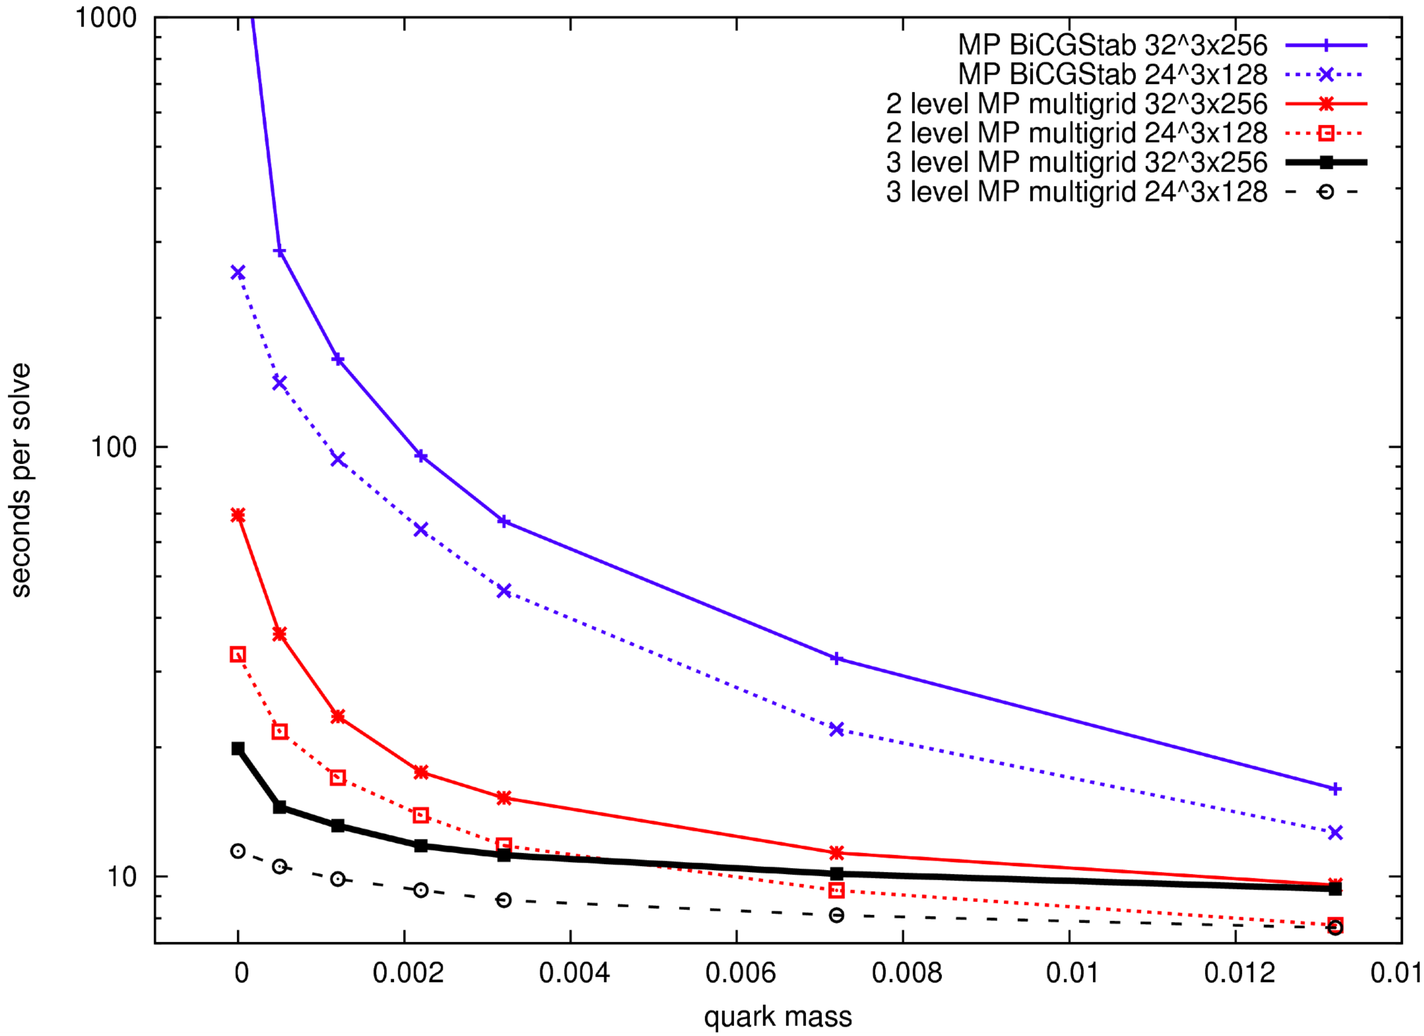

Total cost 32^3x256

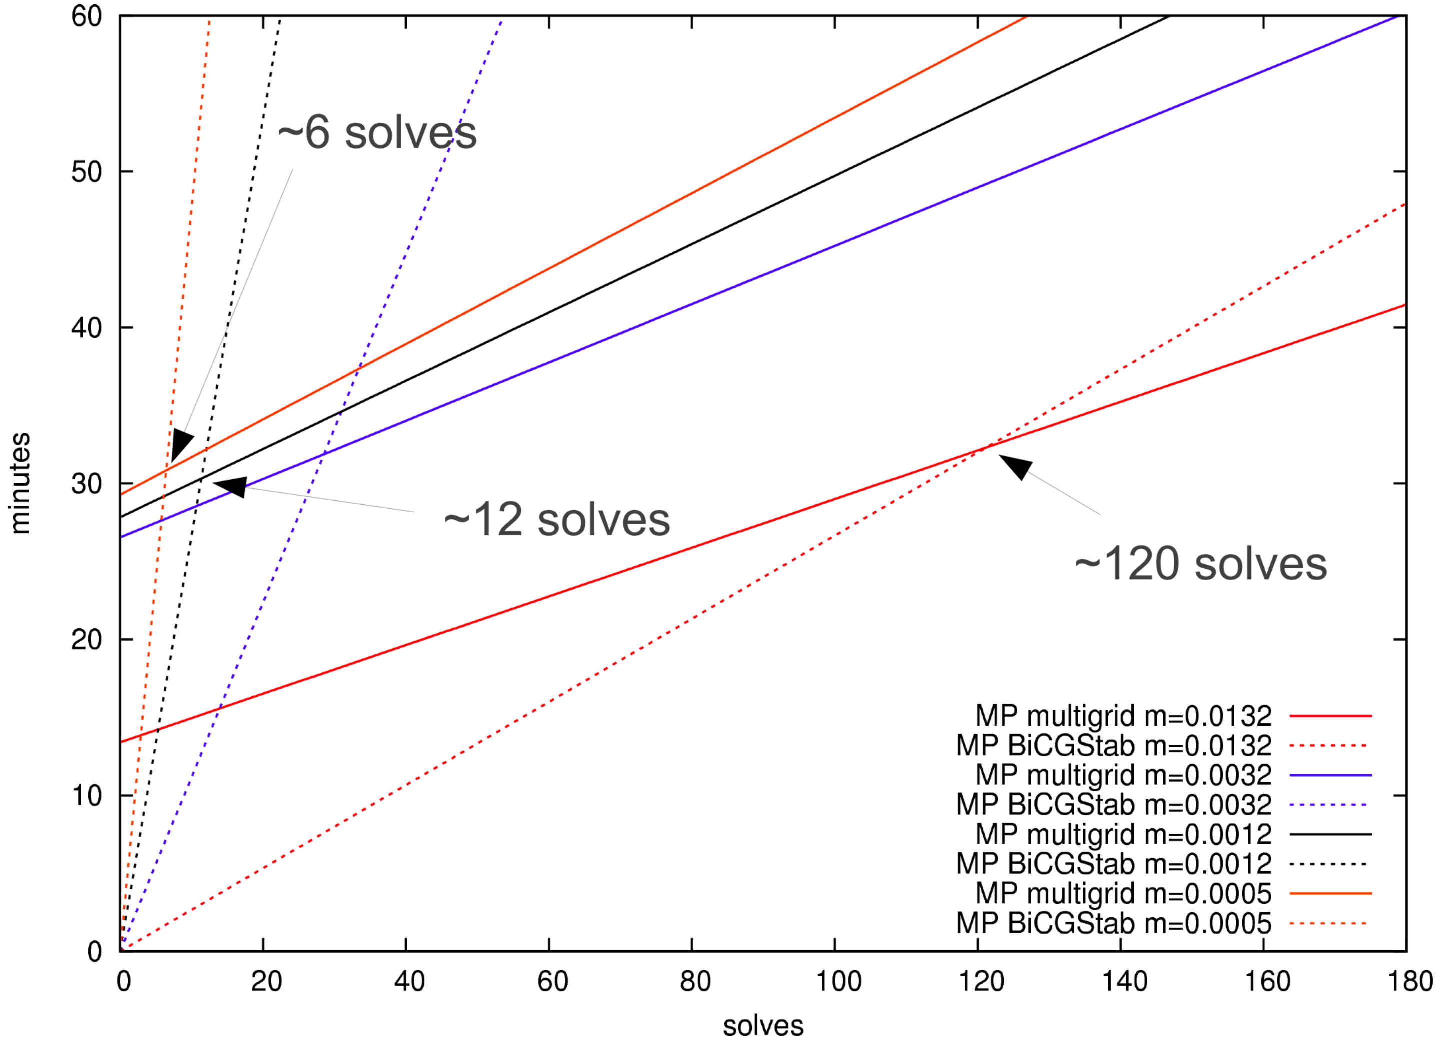

Error vs residual

- Error:

  e = x* - x

- Residual:

  r = b – A x

  = A e

- Residual not as sensitive to low modes

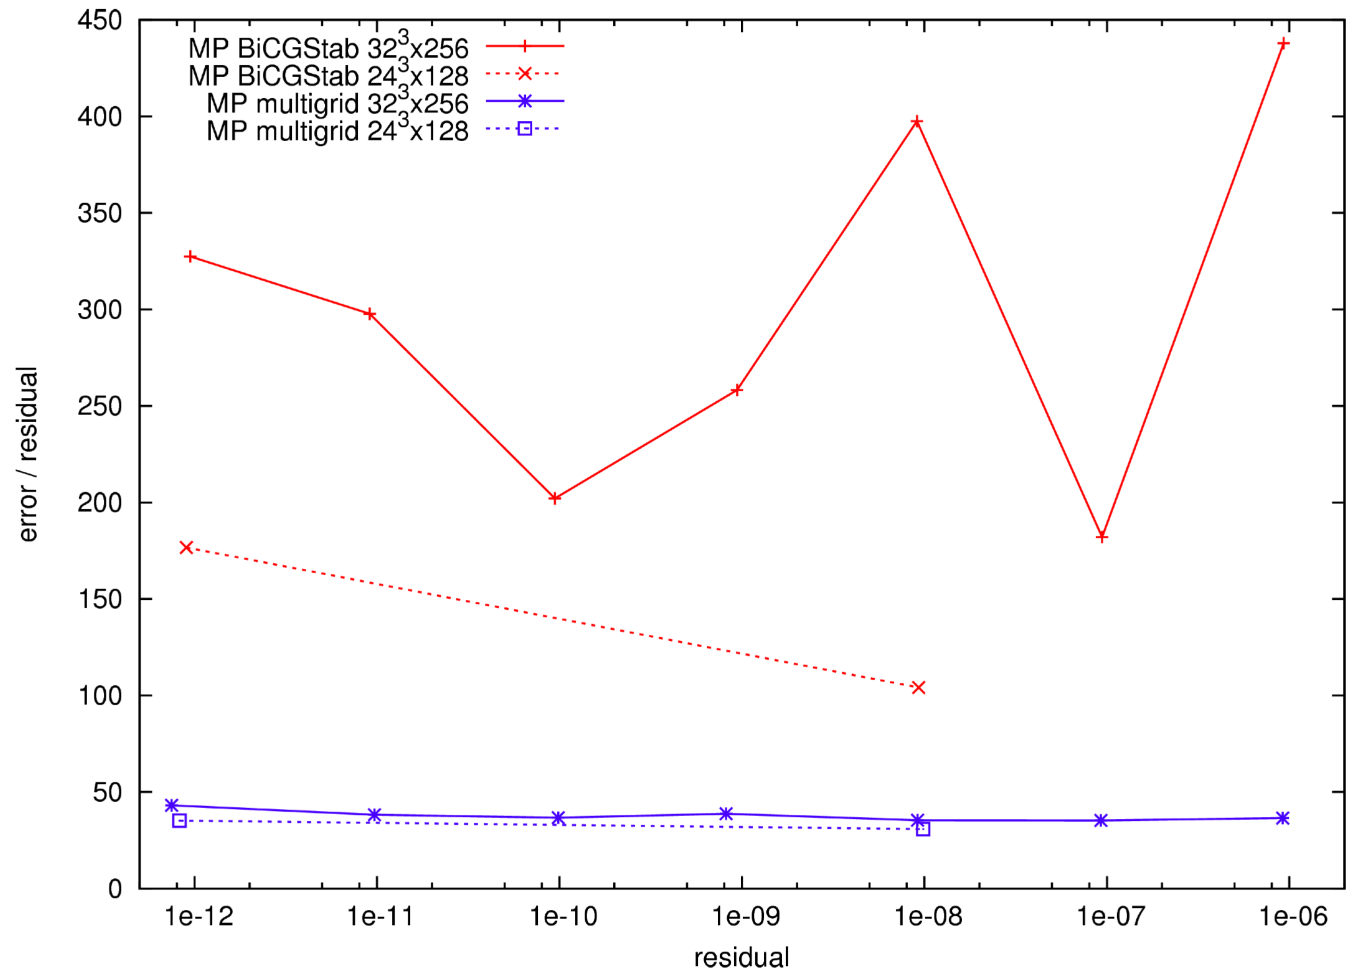

Speed up is even better at fixed error.

A Chiral Fermion

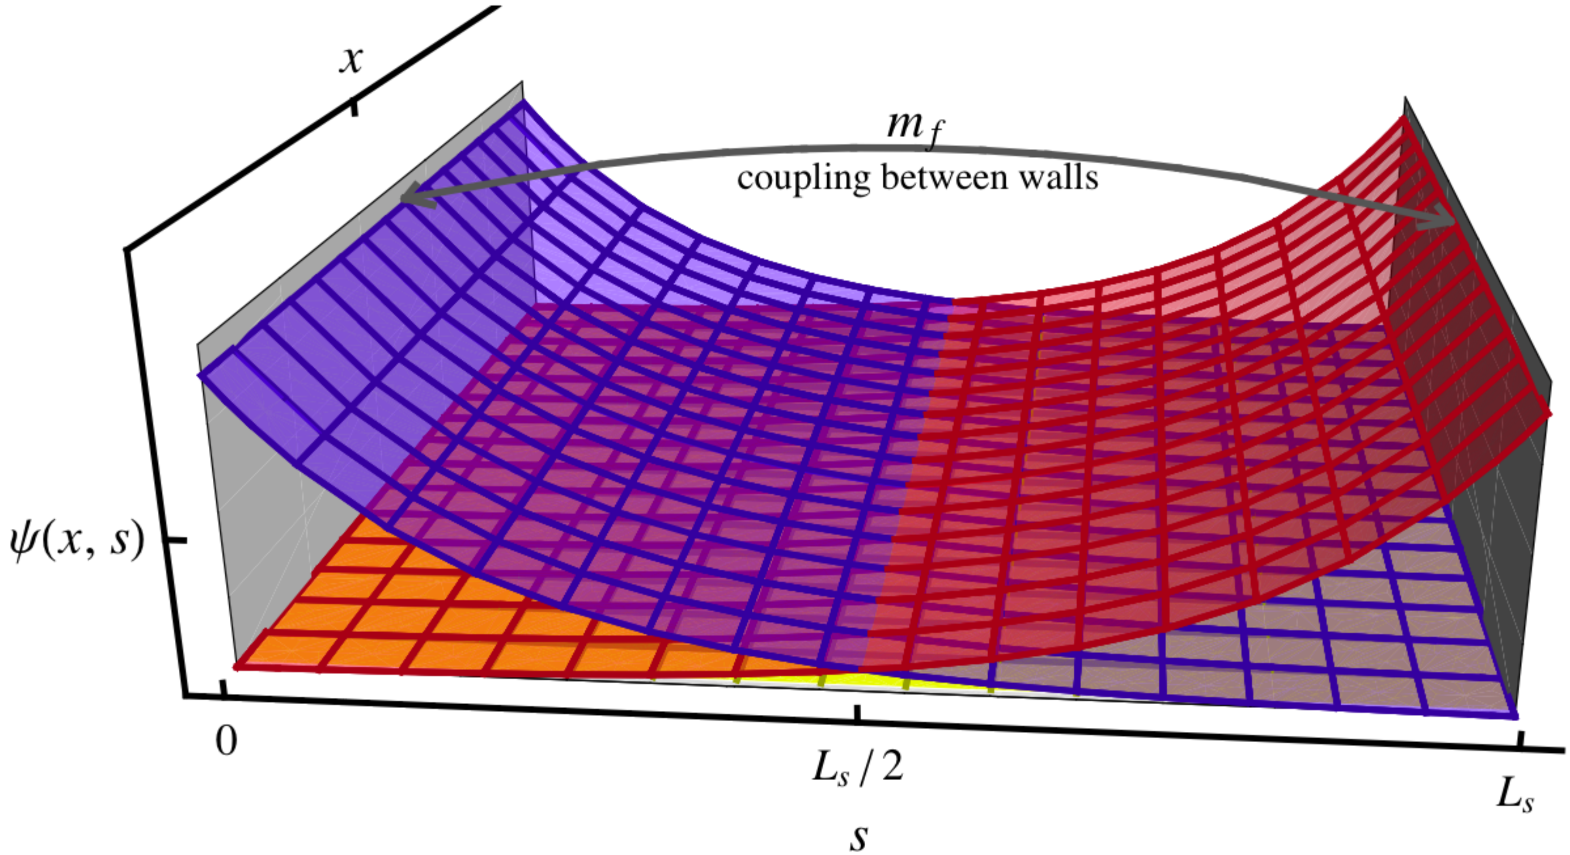

Domain Wall Multigrid Challenge

- 5-d Wilson with "negative mass" and Dirichlet B.C. in 5th "time" interval [0,Ls]
- New Gamma 5: $\Gamma_5 = \gamma_5 \mathcal{R}(s \leftrightarrow L_s - s)$
- D is now violently not normal: $D^\dagger D \neq D D^\dagger$
- Physical long wavelength mode are low singular (eigen)vectors NOT eigenvectors of D!

Domain-Wall Operator

$$D^{\mathrm{dwf}}_{x,s;x',s'} = \delta_{s,s'} D^{\parallel}_{x,x'} + \delta_{x,x'} D^{\perp}_{s,s'}$$

$$D^{\parallel}_{x,x'} =$$

$$(M_5 - d)\delta_{x,x'} + \frac{1}{2}\sum_{\mu}\left[(1-\gamma_\mu)U_{x,\mu}\delta_{x+\hat{\mu},x'} + (1+\gamma_\mu)U^{\dagger}_{x',\mu}\delta_{x-\hat{\mu},x'}\right]$$

$$D^{\perp}_{s,s'} = \frac{1}{2}\left[(1-\gamma_5)\delta_{s+1,s'} + (1+\gamma_5)\delta_{s-1,s'} - 2\delta_{s,s'}\right] - \frac{m}{2}\left[(1-\gamma_5)\delta_{s,L_s-1}\delta_{0,s'} + (1+\gamma_5)\delta_{s,0}\delta_{L_s-1,s'}\right]$$

Domain-Wall Operator

Spectrum

No zero e.v.

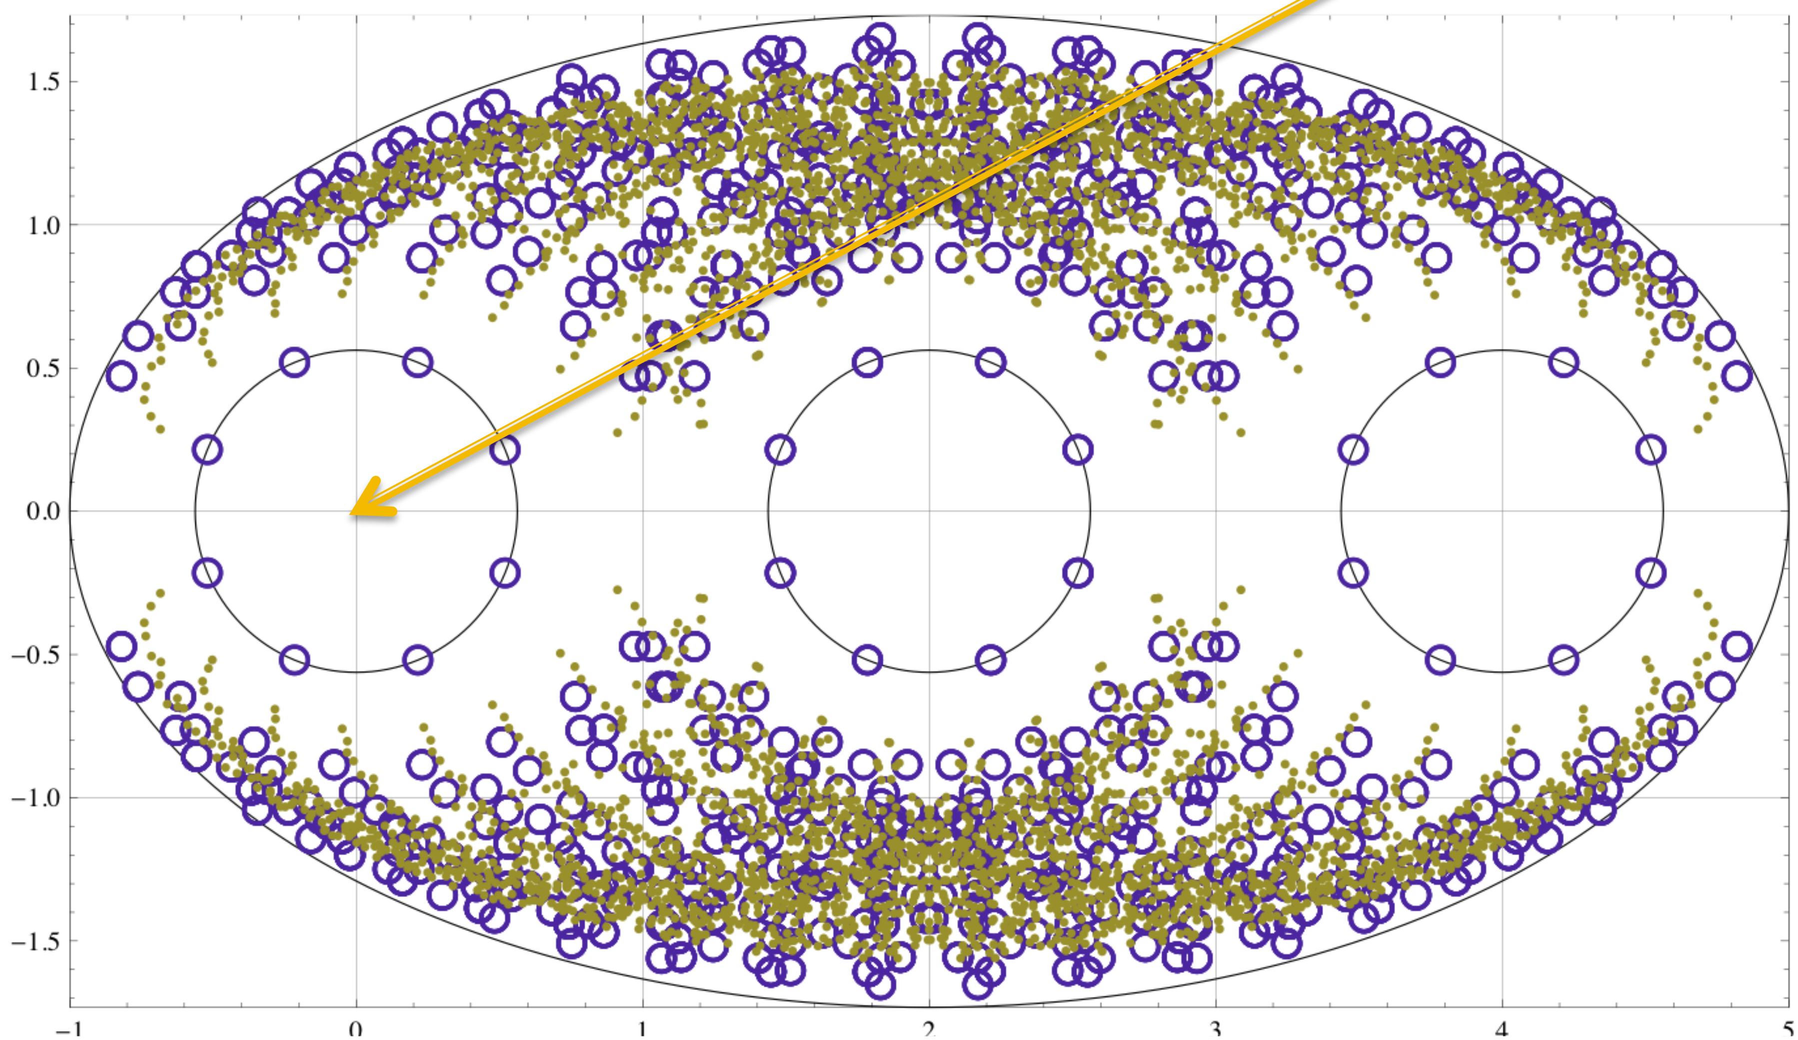

2 + 1 d but 4 + 1 d would have 5 eyes

Solution: Work with Normal Eq. and near zero singular vectors

$$\begin{aligned} D^{DW}_{s's}(M_5, m) &= D_{wil}(M_{cr})\delta_{s's} + [D^{(5)}_{s's}(m) + (M - M_{cr})\delta_{s's}] \\ &\equiv D_{s's}(M_{cr}) + \frac{1+\gamma_5}{2}\Delta + \frac{1-\gamma_5}{2}\Delta^\dagger \,. \end{aligned}$$

$$D^\dagger D B_n = \mu_n^2 B_n \quad \text{and} \quad D D^\dagger F_n = \mu_n^2 F_n$$

$$D B_n = \mu_n F_n \quad \text{and} \quad D^\dagger F_n = \mu_n B_n$$

See basic idea in "Supersymmetric Yang-Mills Theories from Domain Wall Fermions", David B. Kaplan, Martin Schmaltz , Chin.J.Phys. 38:543-550,2000

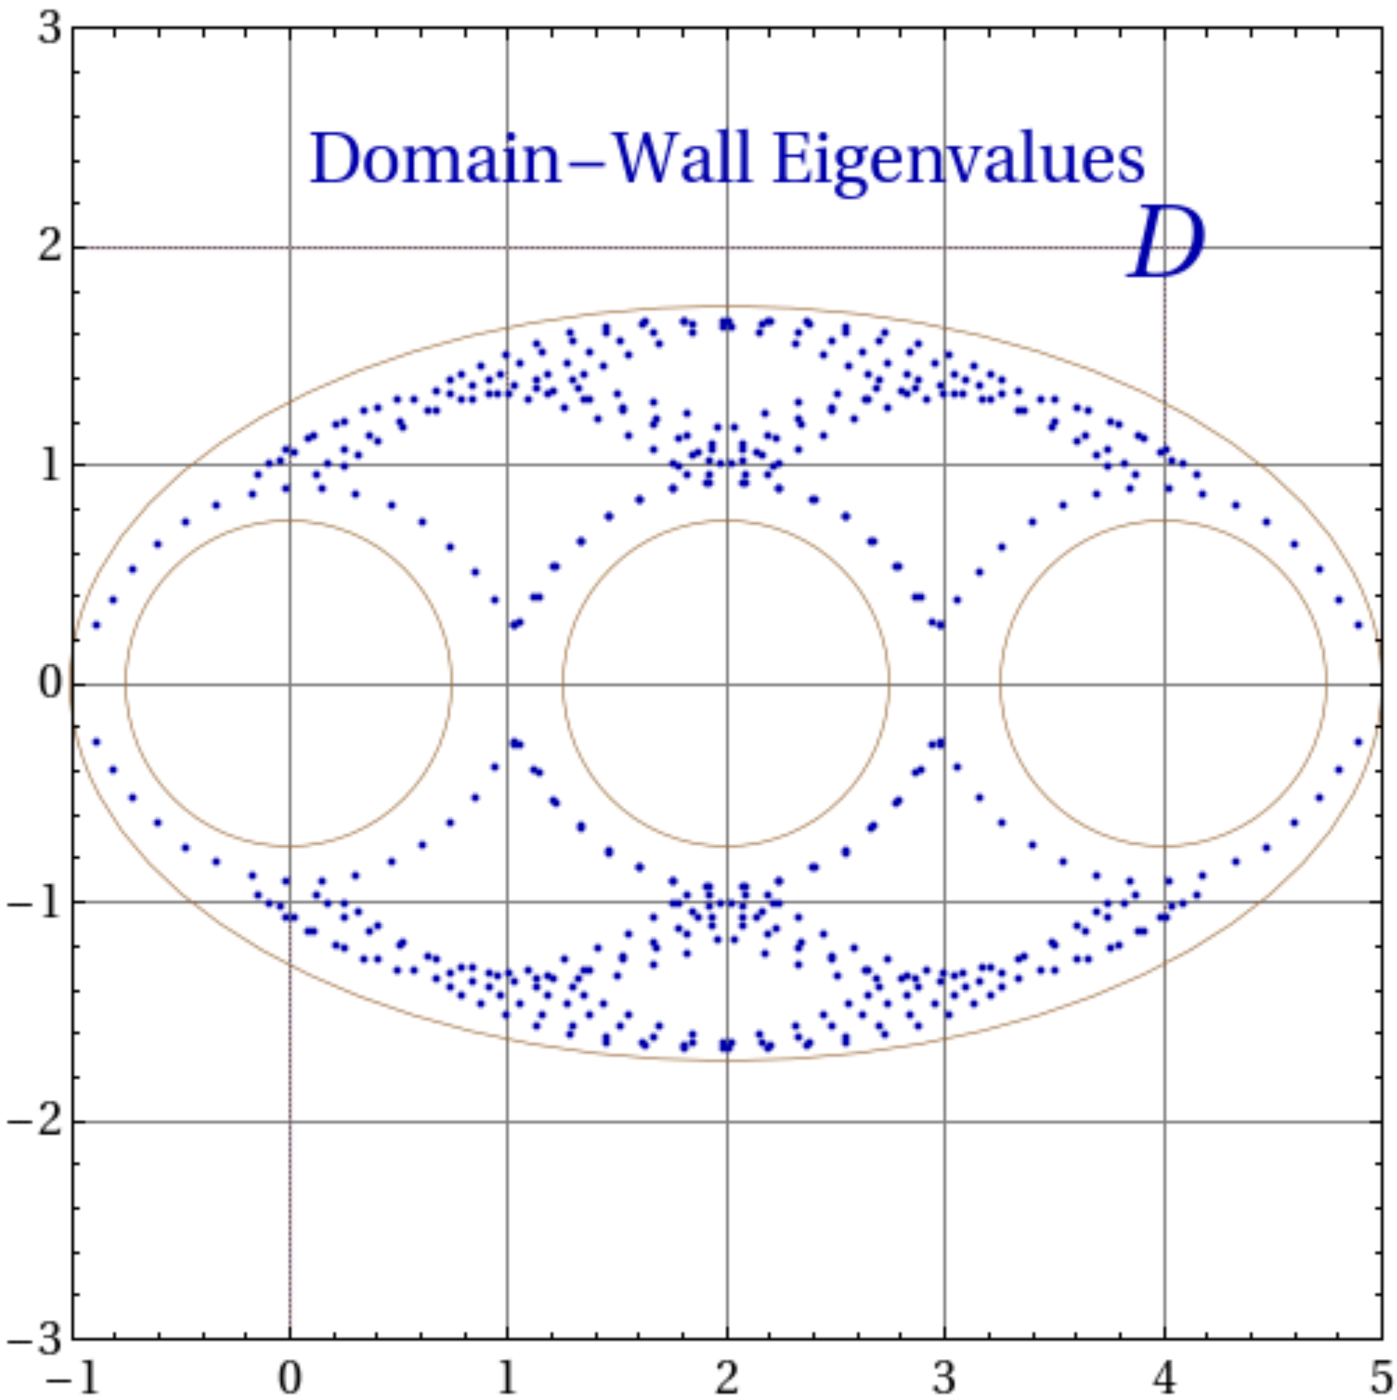

Non Normal
Non Hermitian
Non Pos. Def.

$$D = U\sqrt{D^\dagger D}$$

$$D^\dagger D = H^2$$

$$H = \Gamma_5 D$$

$$(\Gamma_5)^n D$$

Saul's Video!

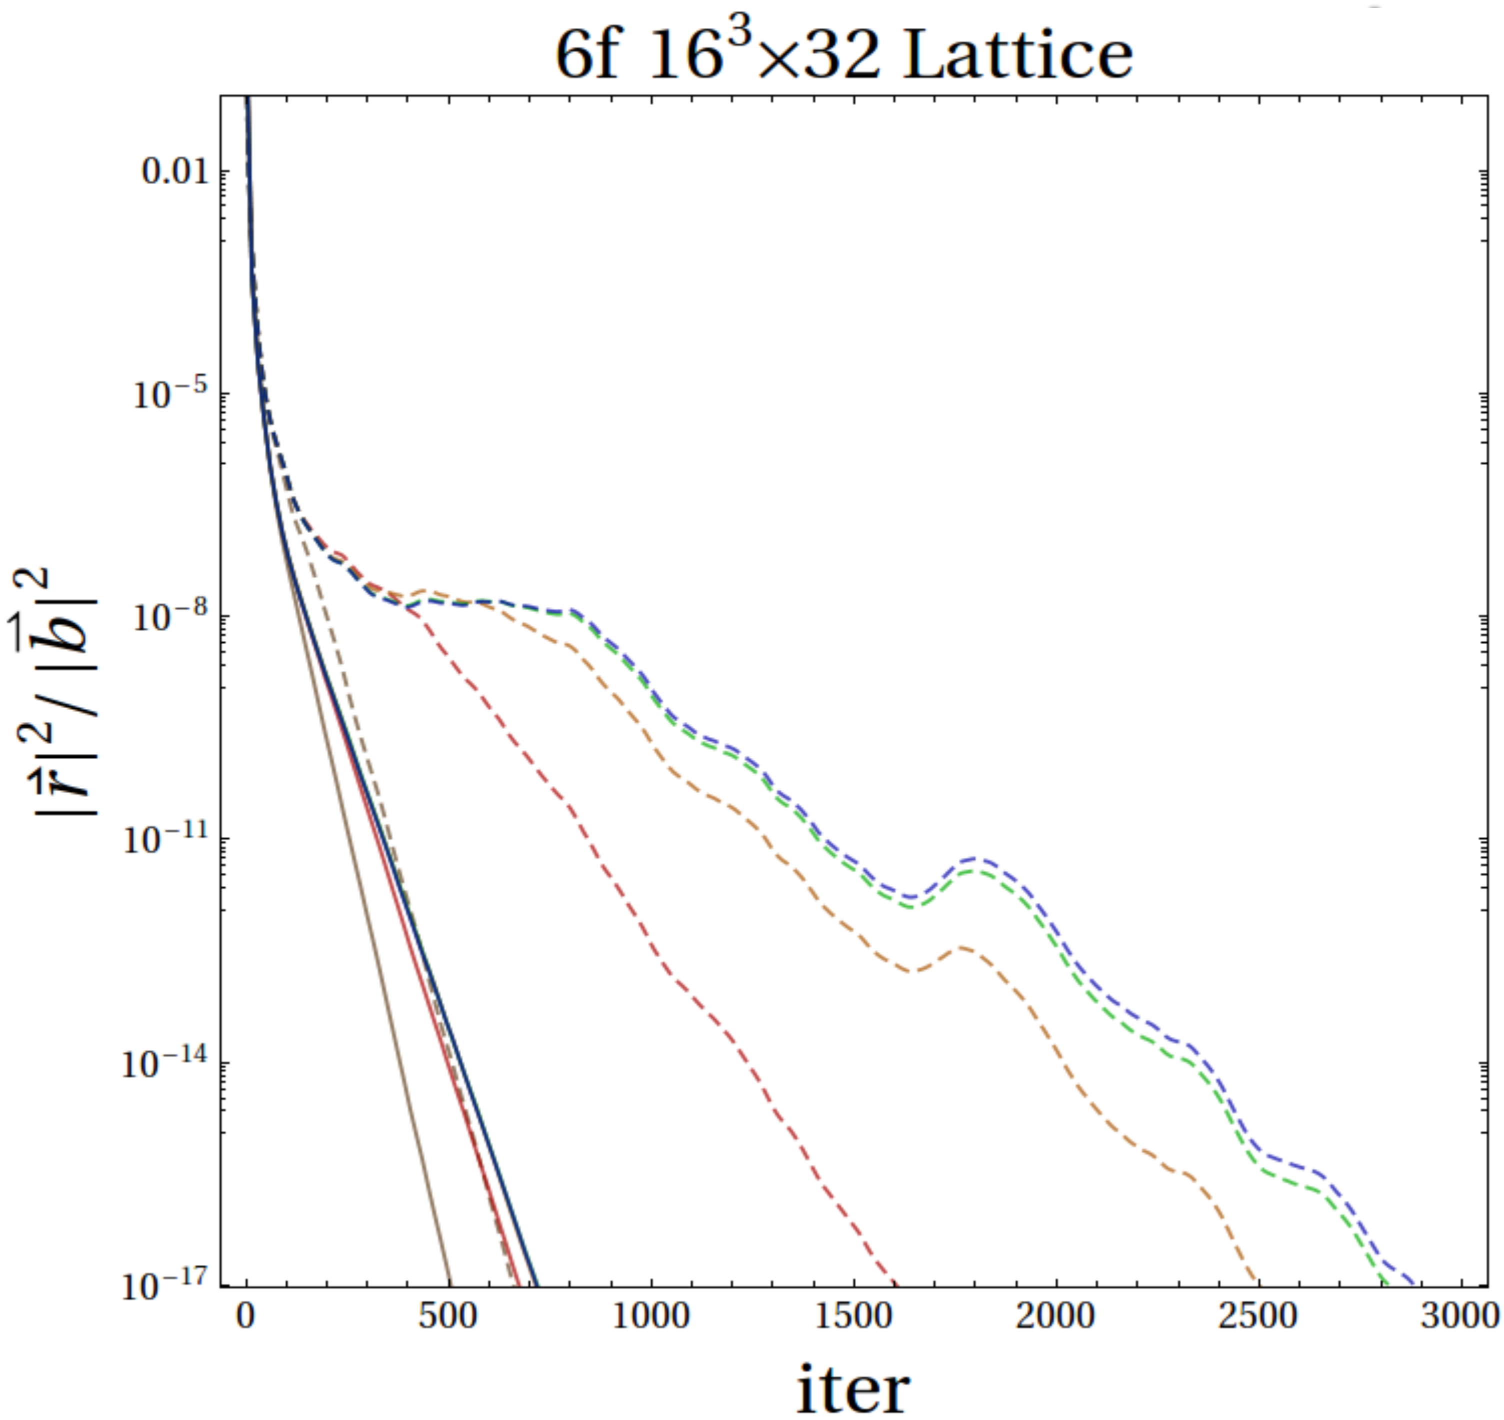

4^4xLs Blocks remove 5-d entirely

(Very Preliminary: More by Friday Saul?)

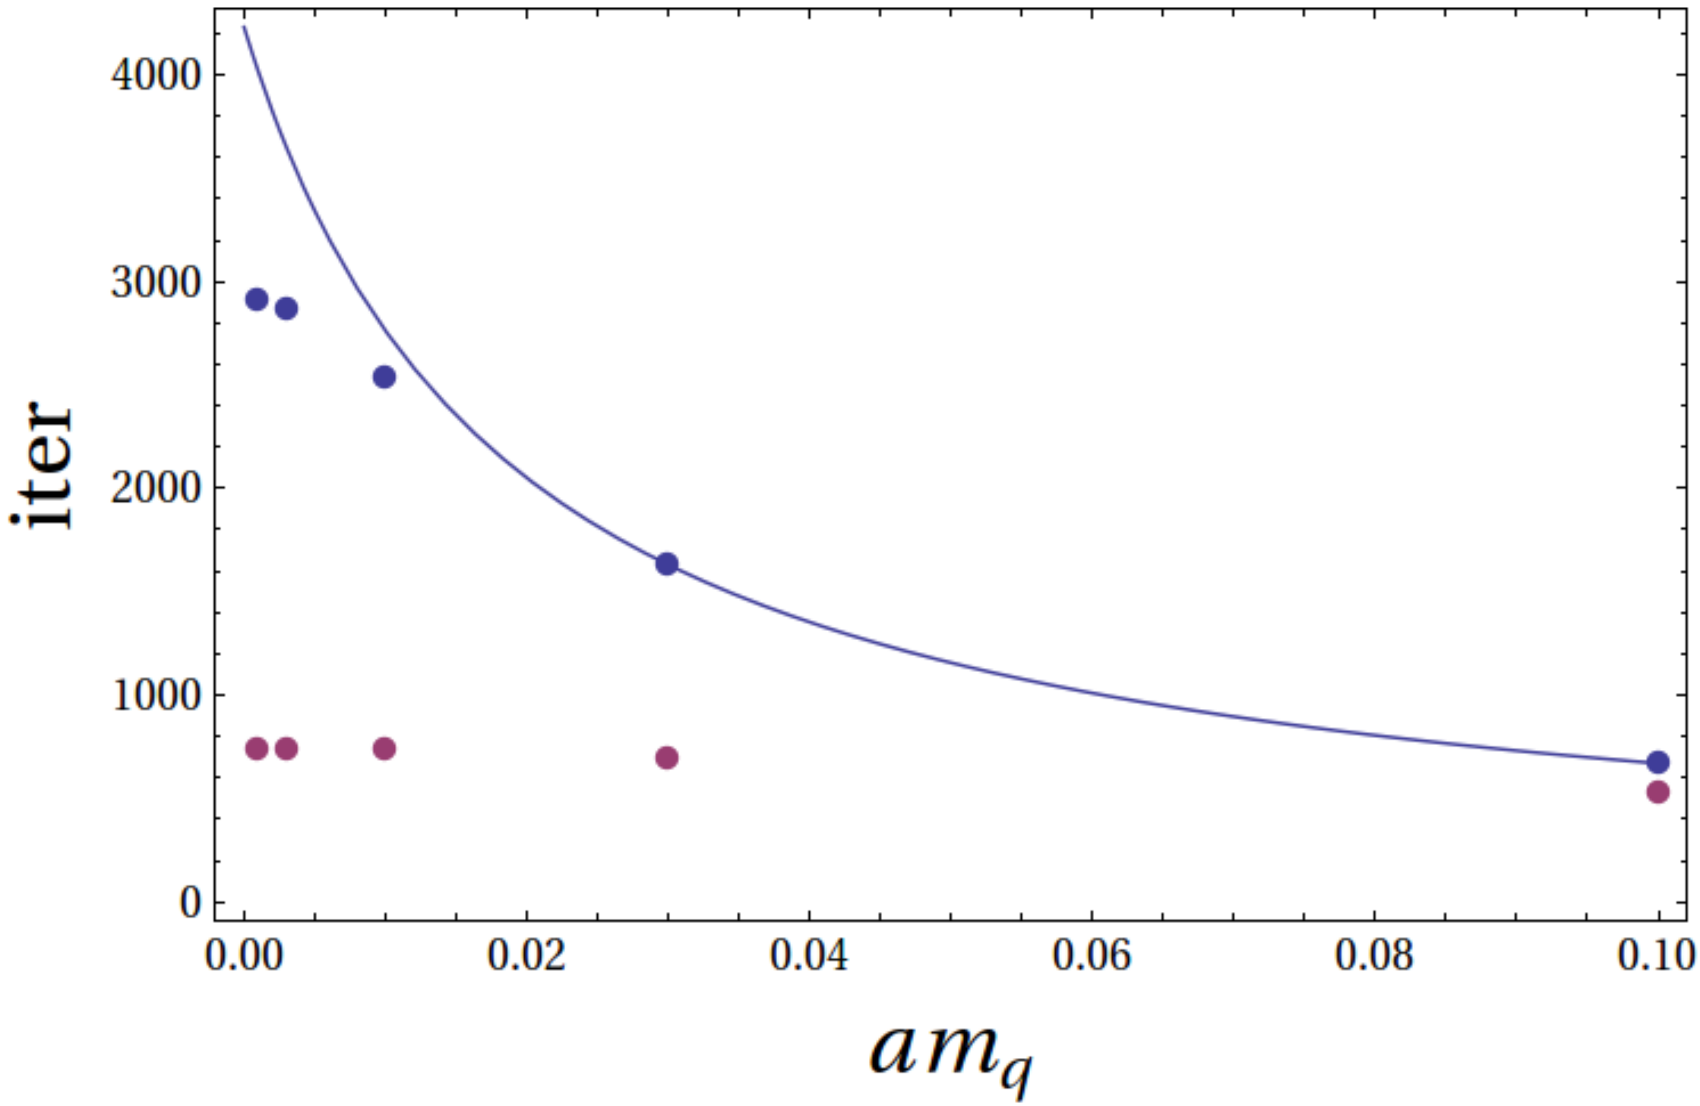

New Project: MG on GPU

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

100+

GPU O(10+) MG O(10+)

New apps:

1. Higgs-Nucleon coupling for Dark Matter detection
2. Beyond the Standard Model results…
3. Nuclear excitations and interactions etc.…
4. GPU-MG for graphene and solid modeling

Accomplishments & Goals

Multi-scale LGT algorithms

Multigrid Dirac solver:

- *Wilson clover solver in production, up to 20x speed up*
- *Multi-lattice support in QDP*
- *Domain Wall demo for Normal equations*

TODO:

- *Distribute MG for Wilson-clover as level 3 code*
- *MG for staggered inverter*
- *MG projection of DW operator and/or Overlap*
- *Investigate Domain Decomposition & data compression*
- *Reduce Data Movement for heterogeneous arch opt.*
- *Extend multi-scale to HMC with MG or DD (i.e. DDHMC)*
- *etc.*

Application to HMC: Lüscher's intermittently update of $\mathcal{S}$ subspace(0710.6417v1)

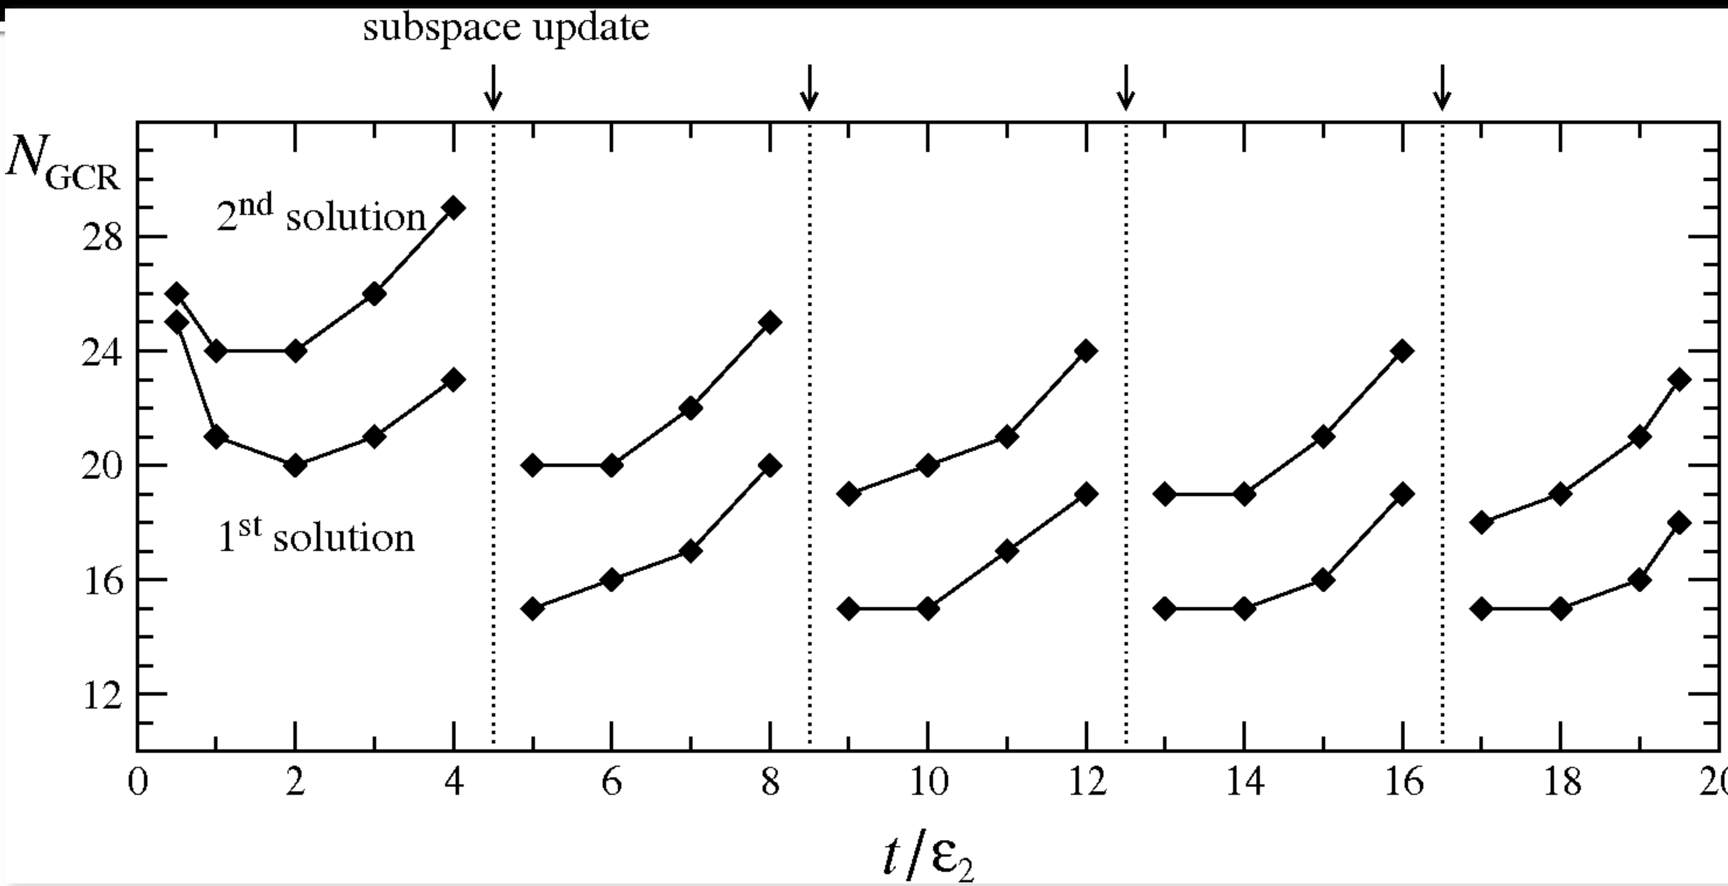

† *Combined with "Chronological Inverter" Brower, Ivaneko, Levi, Orginos*

Graphene

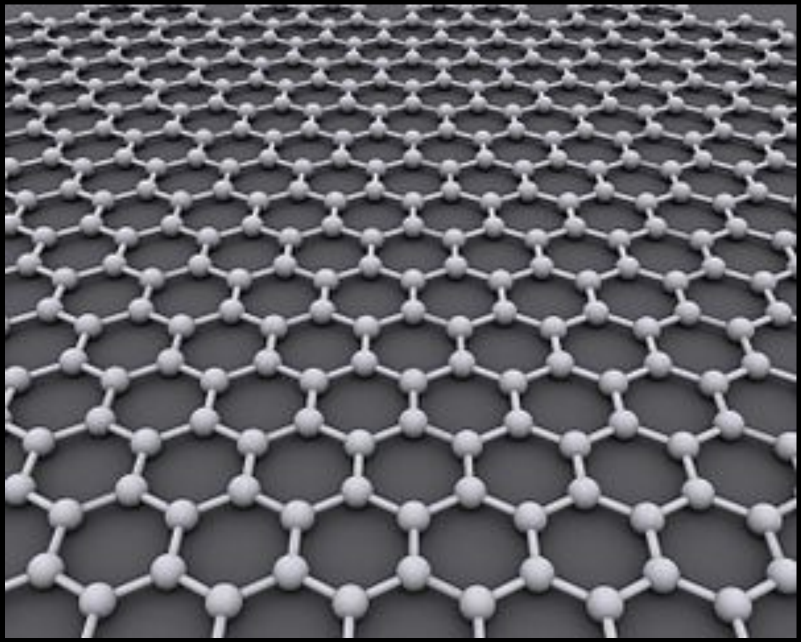

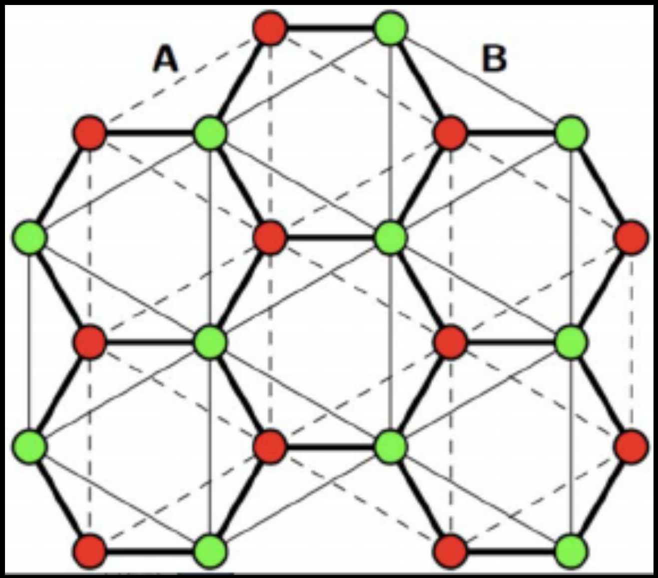

Graphene is 2+1 dimension Carbon sheet with Dirac fields: But lattice is real Hexagonal structure. Couple to coulomb potential and phones act like gauge fields! Ideal for Lattice field theory, MG and GPU! (Brower, Rebbi and Schaich)

LGT Algorithmic Space is large

- Multiple lattice for MG and DD
- New lattice geometries for SUSY and graphene
- New group representation for BSM strong dynamics
- Multi-step Symplectic Hamiltonian integrators for HMC
- Eigen solvers to partition UV vs IR effect
- New boundary condition and asymmetric lattices
- Multi-precision solvers and integrators with full precision results

Some References

- "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, arXiv:1005.3043v2 [hep-lat], Phys. Rev. Lett. 105, 201602 (2010).

- "Adaptive Multigrid Algorithm for Lattice QCD" J. Brannick, R. C. Brower, M. A. Clark, J. C. Osborn, C. Rebbi, arXiv:0707.4018 Phys. Rev. Lett. 100, 041601 (2008)

- "Multigrid solver for clover fermions", R. Babich, J. Brannick, R. Brower, M. Clark, S. Cohen, J. Os- born and C. Rebbi, Lattice 2010, Villasimius, Sardinia, Italy (2010).

- "The role of multigrid algorithms for LQCD", R. Babich, J. Brannick, R. C. Brower, M. A. Clark, S. D. Cohen, J. C. Osborn and C. Rebbi, PoS **LAT2009, 031 (2009) [arXiv:0912.2186 [hep-lat]].**

- "The removal of critical slowing down", J. Brannick, R. C. Brower, M. A. Clark, S. F. McCormick, T. A. Manteuffel, J. C. Osborn and C. Rebbi, PoS **LATTICE2008, 035 (2008) [arXiv:0811.4331 [hep-lat]].**

- “Algorithms for Lattice Field Theory at Extreme Scales”, Richard C. Brower, Ronald Babich, James Brannick, Michael A. Clark, Saul Cohen, Bálint Joó, Anthony Kennedy, James Osborn and Claudio Rebbi, SciDAC 2010, Chattanooga, TN (2010).
- ”Hadronic Physics using Lattice QCD and GPUs” B. Joo, R. Babich, R. C. Brower, M. A . Clark, J. Chen, J. Dudek, R. G. Edwards, M. J. Peardon, C. Rebbi, D. G. Richards, G. Shi, C. Thomas, W. Watson, USQCD Collaboration and Hadron Spectrum Collaboration, SciDAC 2010, Chattanooga, TN (2010).
- “Solving Lattice QCD systems of equations using mixed precision solvers on GPUs”, M. A. Clark, R. Babich, K. Barros, R. C. Brower and C. Rebbi, Comput. Phys. Commun. **181, 1517-1528 (2010) [arXiv:0911.3191 [hep-lat]].**
- “Blasting through lattice calculations using CUDA”, K. Barros, R. Babich, R. Brower, M. A. Clark and C. Rebbi, PoS **LATTICE2008, 045 (2008) [arXiv:0810.5365 [hep-lat].**
- “Moebius Algorithm for Domain Wall and GapDW Fermions”, R. Brower, R. Babich, K. Orginos,C. Rebbi, D. Schaich and P. Vranas, PoS **LATTICE2008, 034 (2008) [arXiv:0906.2813 [hep-lat]].**

FINI

- QUESTION?

Mapping Scales to Architecture

Future LGT Software Infrastructure will be more complex and more essential.
Multi-scale algorithms require collaboration with Applied Mathematicians and mapping tools/memory management collaboration with Software Engineers. Algorithms and Software must dance together (e.g. DD or mixed precision to reduce memory traffic)

New Physics problems: Nuclear scattering, Supersymmetry, LHC model need general gauge groups, propagator packaging. Better symplectic integrators for HMC. Asymmetric lattices/open boundary conditions etc. etc. (Graphene is a LGT with these already)

Libraries and "algorithmic middle ware" and high level interfaces to rapid prototype and deploy new algorithms. Should not try to repeat these software development in each code base worldwide. e.g . P and R (scatter/gather) ops, DSL High level tools like PETc ?

See QUDA (Balint Joo) and QLUA (Andrews Polchinsky) libraries for idea along theses lines!

Accomplishments & Goals (continued)

GPU Software

QUDA (QCD CUDA) [*http://lattice.githup.com/quda*](http://lattice.githup.com/quda)

- *Wilson*
- *Clover improved Wilson*
- *Twisted mass*
- *Improved staggered (astaq or HISQ)*
- *Mixed-precision implementation*

TODO (See slides by Babich, Gottlieb and Balint)

- *Finish Multi-GPU inverter for Domain Wall*
- *Put MG solver onto GPU (estimate O(100) cost reduction!)*
- *Extend QUDA to smeared sources*
- *Extend to do 2 and 3 point hadron correlators*
- *QUDA HMC and some strong scaling to O(100) GPU?*
- *New QDP implantation for GPU like clusters*
- *etc.*