QCD software and algorithms for the exascale

Peter Boyle (BNL)

- Grid software and the exascale
- SciDAC-5 project overview
- Project status
  - Multigrid
  - DDHMC
  - Critical slowing down

Brookhaven National Laboratory

Electronic and hardware trends: what can we say now about the future?

Location	System	Interconnect (GB/s) per node (X+R)	Floating point performance (GF/s) per node	Memory Bandwidth (GB/s) per node	Year	System peak (PF/s)	FP / Interconnect	FP / Memory	Memory / Interconnect
LLNL	BlueGene/L	2.1	5.6	5.5	2004	0.58	2.7	1.0	2.6
ANL	BlueGene/P	5.1	13.6	13.6	2008	0.56	2.7	1.0	2.7
ANL	BlueGene/Q	40	205	42.6	2012	20	5.1	4.8	1.1
ORNL	Titan	9.6	1445	250	2012	27	150.5	5.8	26.0
NERSC	Edison	32	460	100	2013	2	14.4	4.6	3.1
NERSC	Cori/KNL	32	3050	450	2016	28	95.3	6.8	14.1
ORNL	Summit	50	42000	5400	2018	194	840.0	7.8	108.0
RIKEN	Fugaku	70	3072	1024	2021	488	43.9	3.0	14.6
NERSC	Perlmutter/GPU	200	38800	6220	2022	58	194.0	6.2	31.1
ORNL	Frontier	200	181200	12800	2022	>1630	906.0	14.2	64.0

- All DOE Exascale computing is GPU accelerated
- Huge gains in floating point
  not matched by gains in memory (14x) and interconnect (300x)
- Machines increasingly better suited for dense matrices and machine learning
- Lots of diversity and difficulty:
  - Systems with AMD, Intel, Nvidia GPUs
  - Systems with CPU cores (+ HBM : Fujitsu, Intel SPR)
  - HIP, SYCL, CUDA and conventional programming
  - Host memory, GPU memory, DDR+HBM numa on CPU

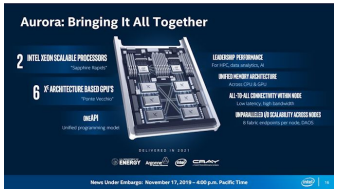

Forthcoming systems will increase floating point performance dramatically, but not interconnect.

- Lattice gauge theory algorithms for gauge field sampling *must* change to exploit.
- Lattice gauge theory correlation function calculations can run brilliantly

Brookhaven National Laboratory

Grid support & developments

- Grid has support for: HIP, SYCL, CUDA and conventional OpenMP loop acceleration
  - Portability key: if read one bit of Grid after this talk, make it: https://github.com/paboyle/Grid/blob/develop/Grid/threads/Accelerator.h
  - For loop macro captures loop bodies in (device) function objects
  - SIMD and SIMT with single source kernels
- Unified memory model OR distinct accelerator memory
  - Software managed device cache: O(1) cost lookup, O(1) true LRU eviction, evict-next/transient options
- Recent experimental additions:
  - RRRR / IIII data layout - relevant to Fugaku.
    Needs a project for A64FX support - Tilo, Nils?.
  - Padded cell & General stencil
    Optimise staggered smeared forces
    off axis stencil operators - HDCG style non-local coarsening
    (=complicated Covariant transport paths: c.f. Lehner, Wettig!)
- Broad use: GPT (Lehner), Hadrons (Portelli), Qlat (Jin), CPS (Jung), MILC (Detar)

Portability helps: RBC-UKQCD is running on

- (Europe) Leonard/Cineca, Lumi/CSC, Booster/Juelich
- (USA) Perlmutter/NERSC, Summit,Crusher, Frontier/ORNL
- (Small islands that are definitely not in Europe) Tursa/Edinburgh

Large effort by many people, notably C. Lehner, C. Jung, C. Kelly, A. Portelli

Performance

Perlmutter(Cray/Nvidia A100), Crusher (Cray/AMD MI250), Summit (IBM/Nvidia V100)

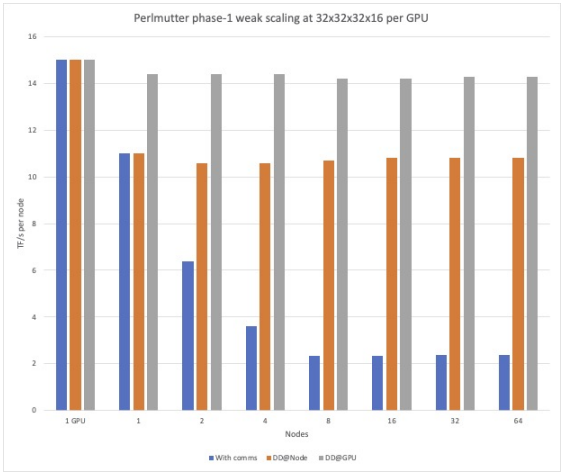

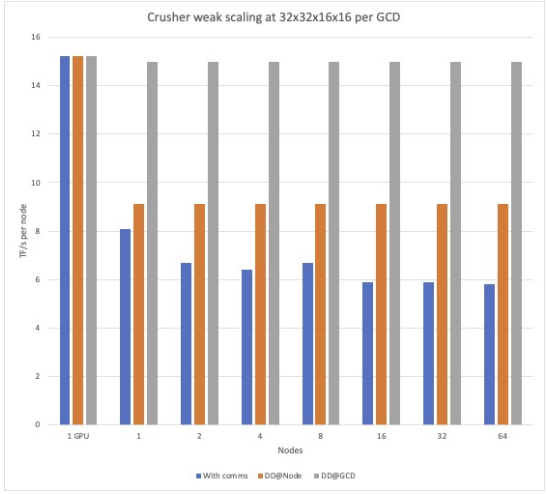

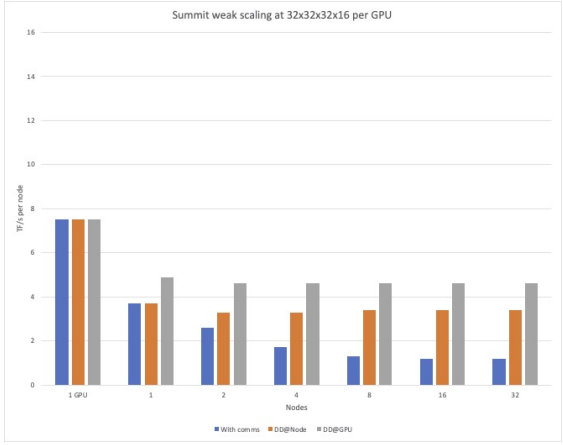

- Expect Aurora will have
  - Substantially more single node performance than Crusher
  - Same network technology as Crusher
  - $\Rightarrow$ Performance loss to communication will be a larger hit !!!

Brookhaven National Laboratory

Nvidia GPU performance

Provably unimprovable:

Dslash kernel: 39% FMA pipe, 80% L2, 78% memory; hard to improve by much

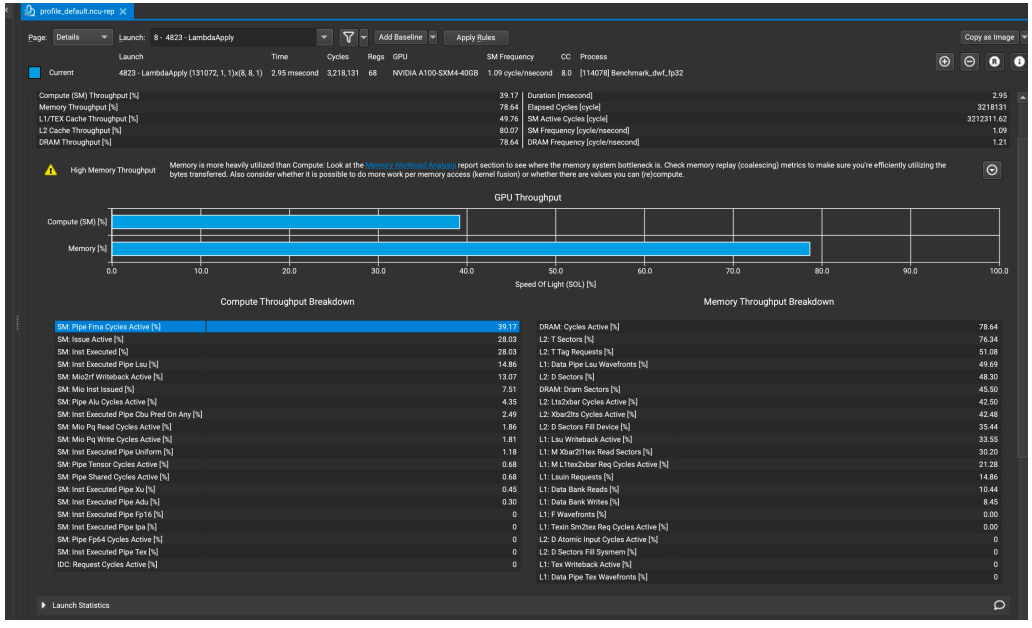

Brookhaven National Laboratory

Edinburgh Tursa / Juelich Booster / Cineca Leonardo

- 10TF/s per node *including* communication

32^4 comms and
Compute perfectly
Overlapped

Stopped using Nvlink
With GPU, use RDMA

Read coalesce kernels

750us -> 60us

6.6TF/s -> 9.9 TF/s

Grid benchmark
5.3 last Nov, now 9.9
With many improvements

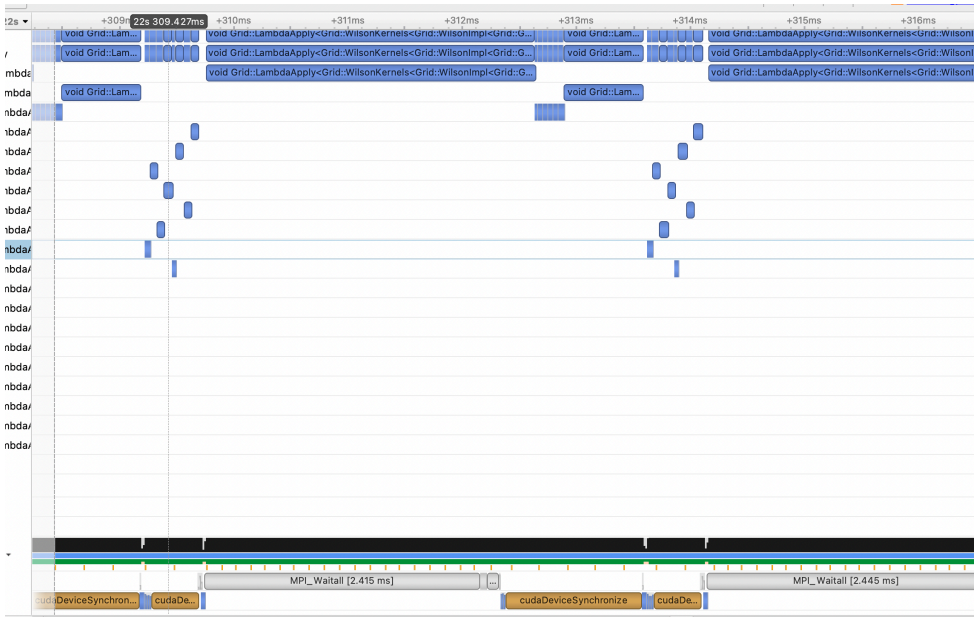

- Atos nodes, 4 x A100 with 4x Mellanox HDR
- Gives 185 GB/s bidirectional interconnect bandwidth
- Can 'prove' code is optimal: saturates memory bandwidth on every one of 22 kernels in sequence

Brookhaven National Laboratory

ORNL: Frontier/Crusher. CSC: LUMI

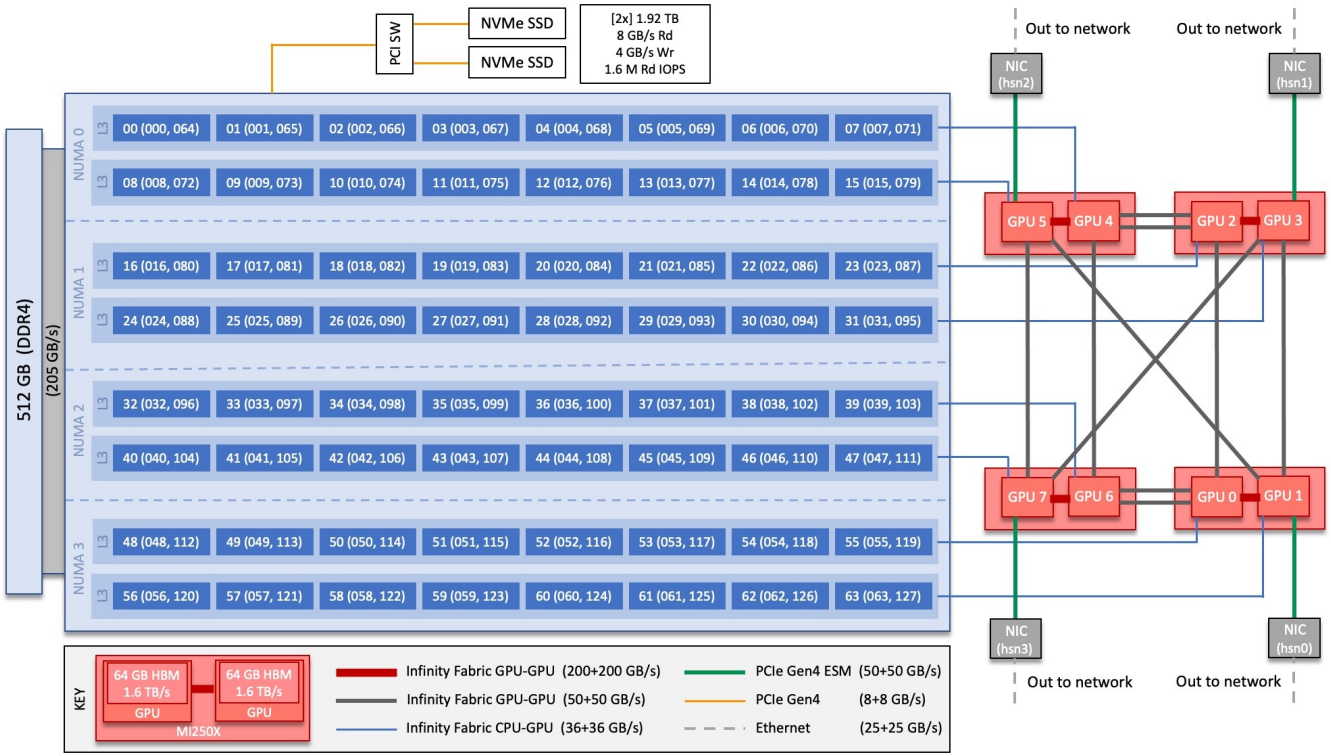

- 4x Slingshot 200 Gbit/s per node
- Mapping GPU's 0,1,2,3,7,6,5,4
- Then MPI ranks $\{000, 001, 010, 011, 100, 101, 110, 111\}$ are the vertices of a 2^3 cube in network

Brookhaven National Laboratory

Frontier/Crusher (ORNL), LUMI (CSC)

- 7.4TF/s per node including communication
- Cray/AMD 4x MI250 nodes, 4x Slingshot 11

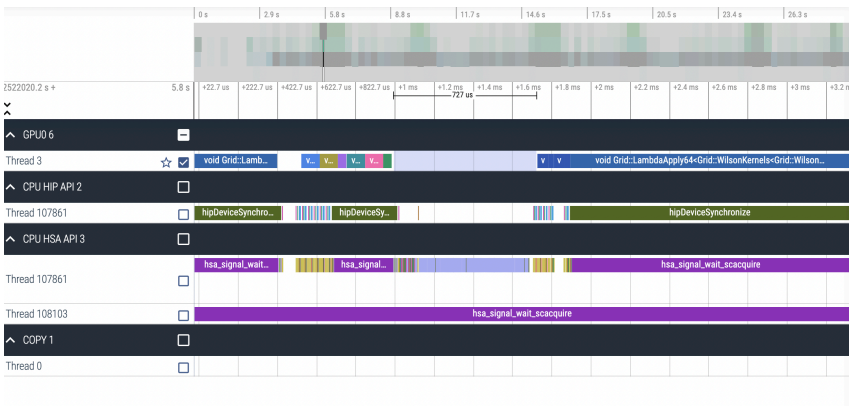

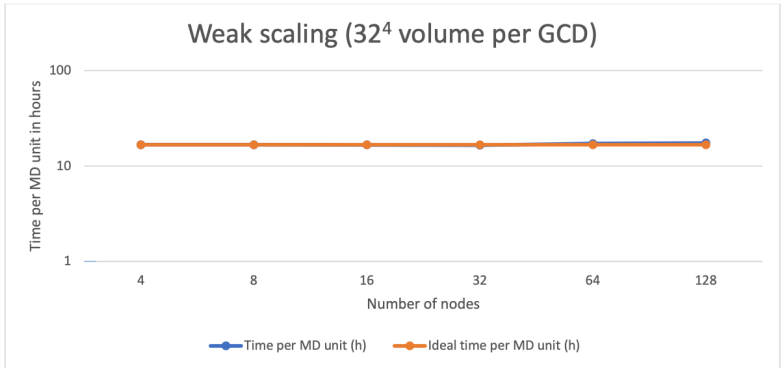

Nodes	Performance (TF / s)	Performance (TF / s / node)	Performance on LUMI-G (TF / s / node)	Time per MD unit (h)	Ideal Time per MD unit (h)
4	29.24	7.31	7.43	16.7	16.7
8	58.64	7.33	7.46	16.7	16.7
16	117.92	7.37	7.41	16.6	16.7
32	236.8	7.40	7.19	16.5	16.7
64	455.04	7.11	7.09	17.2	16.7
128	896	7.00	7.06	17.4	16.7

- After various fixes: constant performance
- 'tricks' to get MPI using DMA hardware between GPUs
- Infinity link intranode is relatively slow compared to NVlink
- Network limit is same as Tursa and performance within 30%
- No 'proof' it is optimal yet: might hope for further gains

Brookhaven National Laboratory

Algorithms for the future?

- Ratio of single GPU code to multinode code is 2:1 (Crusher)
- Ratio of single GPU code to multinode code is 7:1 (Perlmutter phase 1)
- Ratio on $500M Aurora may be similar to Perlmutter (but same network as Crusher)
- Need new algorithms $\Rightarrow$ SciDAC-5

SciDAC-5 project

Hybrid Monte Carlo

- Auxiliary Gaussian integral over conjugate momentum field $\int d\pi e^{\frac{-\pi^2}{2}}$
  Lives in Lie algbra; serves only to move U round the group Manifold

$$\int d\pi \int d\phi \int dU \quad e^{-\frac{\pi^2}{2}} e^{-S_G[U]} e^{-\phi^*(M^\dagger M)^{-1}\phi}$$

- Outer Metropolis Monte Carlo algorithm
  - Draw momenta
  - Draw pseudofermion as gaussian $\eta = M^{-1}\phi$
  - Metropolis acceptance step
- Metropolis proposal includes inner molecular dynamics at constant Hamiltonian:

$$H = \frac{\pi^2}{2} + S_G[U] + \phi^*(M^\dagger M)^{-1}\phi$$

- Drive as $\dot{U} = i\pi U$ derive HMC EOM from:

$$\dot{H} = 0 = \pi[\dot{\pi} + iU \cdot \nabla_U S_{TA}]$$

- Finite timestep performed in Lie algebra, keeps U on group manifold:

$$U' = e^{i\pi dt} U$$

- Force terms $\nabla_U S$ via product and chain rule; rules of matrix differentiation.
- Must invert $M^\dagger M$ at each timestep of evolution in MD force

$$\delta(M^\dagger M)^{-1} = -(M^\dagger M)^{-1}[(\delta M^\dagger)M + M(\delta M)](M^\dagger M)^{-1}$$

SciDAC-5 project

- "Multiscale acceleration: Powering future discoveries in High Energy Physics"
- 5 year project 2022/10 to 2027/10
- USQCD:
  - 3 Labs: ANL, BNL (lead), FNAL
  - 8 Universites: Columbia, BU, MSU, Illinois, UIUC, Utah
- SciDAC Fastmath :
  - LBNL, MIT, SUNY Buffalo
- Algorithmic research collab. between USQCD HET and SciDAC institutes (applied math/computer science)
- Funded by both DOE/HEP and DOE/ASCR
- 3 work packages:
  - WP1: Multigrid for Domain Wall and Staggered Fermions
  - WP2: Transformational sampling algorithms
  - WP3: Large domain decomposed HMC: minimise communication

Brookhaven National Laboratory

SciDAC-5 personnel

https://scidac5-fastmath.lbl.gov/
https://petsc.org/release/

SciDAC:

- MIT - Youssef Marzouk (FastMath Uncertainty Quantification)
- LBNL - Mark Adams (FastMath PETSc)
- SUNY Buffalo - Matt Knepley (PETSc), Joe Pusztay, Duncan Clayton

USQCD:

- ANL - James Osborne, Xiaoyong Jin
- BNL - Peter Boyle, Taku Izubuchi, Chulwoo Jung, Christopher Kelly, Nobu Matsumoto
- FNAL - Andreas Kronfeld, Jim Simone
- Boston University - Rich Brower
- Columbia - Norman Christ, Yikai Huo
- Indiana - Steve Gottlieb
- MSU - Alexei Bazavov
- UIUC - Aida El-Khadra
- Utah - Carleton Detar, David Clarke

Postdoc and PhD student positions are being filled

Brookhaven National Laboratory

Work package 1

- SciDAC: Adams, Boyle, Brower, Clayton, Gottlieb, Kelly, Knepley, Kronfeld, Pusztay
- Collaborators: Clark, Weinberg, Owen, others

PETSc Collaboration

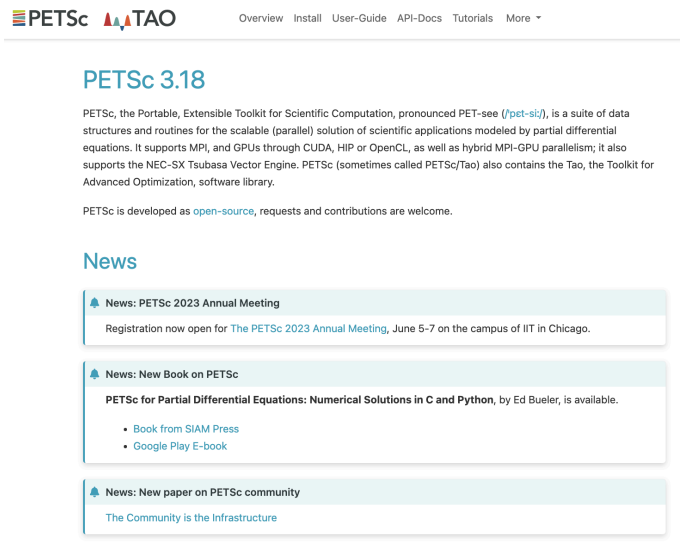

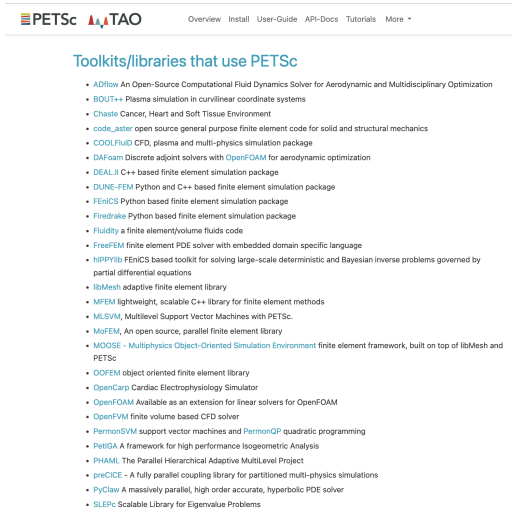

- Widely used world class PDE solver and multigrid HPC library in "C"
- Expert applied mathematicians
- Broad range of algorithms
- Aim: *expose our problems to PETSc team in a way we can work on them together*
- Status:
  - Completely documented Wilson operator
  - PETSc based implementation of Wilson (Pusztay, Knepley, Clayton)
  - free field Fourier verification
  - regresses to Grid, can load NERSC configurations (PB)
  - PETSc based implementation of DWF (PB)
  - Wilson multigrid under development

Brookhaven National Laboratory

SciDAC-5 multigrid & half plane condition

Algorithmic research: I bet we'll break every non-Hermitian Krylov solver in PETSc

- Spectrum of DWF presents problems to non-Hermitian solvers
- Why? Krylov space is the span of polynomials of matrix M. Let $|i\rangle$ be the set of right eigenvectors, $\mathscr{P}(x) = c_n x^n$ a polynomial

$$\begin{aligned} M|i\rangle &= \lambda_i |i\rangle \\ \eta &= \eta_i |i\rangle \\ \psi^{\text{Krylov}} &= \mathscr{P}(M)\eta = (c_n \lambda_i^n)\eta_i |i\rangle \\ \psi^{\text{True}} &= \frac{1}{\lambda_i}\eta_i |i\rangle \end{aligned}$$

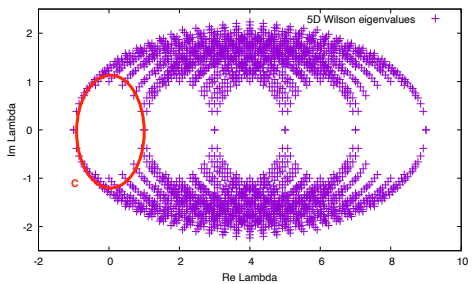

- There exists a contour C contained entirely within the (dense in large/infinite volume) spectrum such that

$$\oint_C \mathscr{P}^{\text{Krylov}}(z)dz = 0$$

$$\oint_C F^{\text{True}}(z)dz = \oint_C \frac{1}{z}dz = 2\pi i$$

- Thus the Krylov polynomal and true solution must differ *within the domain of the spectrum*
- Polynomial must differ from solution between discrete eigenvalues and low order smooth polynomial is inadequate
- Slow convergence, perhaps of order system size

Brookhaven National Laboratory

Domain wall: where to?

- arXiv:1205.2933, Cohen, Brower: real positive ($M^\dagger$ precondition)
  coarsen $M^\dagger M$ (2hop)
- arXiv:1402.2585, PB: real indefinite (RB-NE precondition)
  coarsen $M_{pc}^\dagger M_{pc}$ (4hop)
- arXiv:1611.06944, arXiv:2203.17119, PB, Yamaguchi: real indefinite (Γ_5 precondition)
  coarsen $\Gamma_5 M$ (1hop)
- arXiv:2004.07732, Weinberg et al: complex positive half plane ($M_{PV}^\dagger$ precondition),
  coarsen M (1 hop) and $M_{PV}^\dagger$ (1 hop)
  - 2D U(1) arXiv:2004.07732
  - 4D SU(3) arXiv:2203.17119

Possible plan:

- Directly coarsen 2-hop matrix $M_{PV}^\dagger M$
- PETSc opens up all sorts of algebraic and other multigrid options
- Good ideas arising from applied math community

Work package 2

- SciDAC: Izubuchi, Matsumoto, Marzouk, Christ, Jung, Boyle, Brower, Osborne, X. Jin
- Collaborators: Tomiya, L. Jin

Brookhaven National Laboratory

SciDAC-5 Sampling algorithms

- Generalised Fourier acceleration of HMC
  - Generalise Riemanian Manifold HMC ECP for new momentum distributions
    Jung, Christ + Marzouk arXiv:2112.04556, arXiv:1710.07036
- Trivialising maps and Field Transformation-HMC (Izubuchi, Matsumoto + Marzouk)

$$\int dU e^{-S[U]} = \int dV \left| \frac{dU}{dV} \right| e^{-S[U(V)]}$$

  - arXiv:2212.11387, (Nobu Matusmoto Lattice 2022), Matsumoto THIS WORKSHOP + Akio Tomiya, Luchang Jin, PB, Christoph Lehner, Chulwoo Jung
  - Effectively Luscher's Wilson flowed HMC *but with large flow timestep* arXiv:1009.5877
  - Currently in Qlat (Jin) quenched Wilson flow code
- General trivialising flows. Multiple Wilson loops (Matsumoto, Izubuchi, Tomiya, Jin)
  - Migrating plaquette flow HMC to Grid (PB); add fermions
    Generalise Grid smear HMC to include gauge action; Will add Jacobian
    Link smear in multiple steps, operating on disjoint subsets of links at each step
- UV smearing function $U(V)$ brings tunable Fourier acceleration with incomplete trivialisation

Work package 3 : DDHMC for odd flavours

- SciDAC: Bazavov, Boyle, Brower, D. Clarke, Kelly, Detar, El-Khadra, Kronfeld, Simone
- Collaborators: D. Bollweg, A. Yamaguchi

Brookhaven National Laboratory

Domain decomposition : arXiv:2203.17119

A matrix can be UDL factorised around its lower right block as follows,

$$\begin{pmatrix} D & C \\ B & A \end{pmatrix} = \begin{pmatrix} 1 & CA^{-1} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} S_\chi & 0 \\ 0 & A \end{pmatrix} \begin{pmatrix} 1 & 0 \\ A^{-1}B & 1 \end{pmatrix}, \tag{1}$$

where the Schur complement,

$$S = D - CA^{-1}B.$$

- Divide space time into "black" and "white" subsets.
- Reorder vectors as black and white vectors.
- Differential operators are matrices connecting black to black, white to white and cross terms

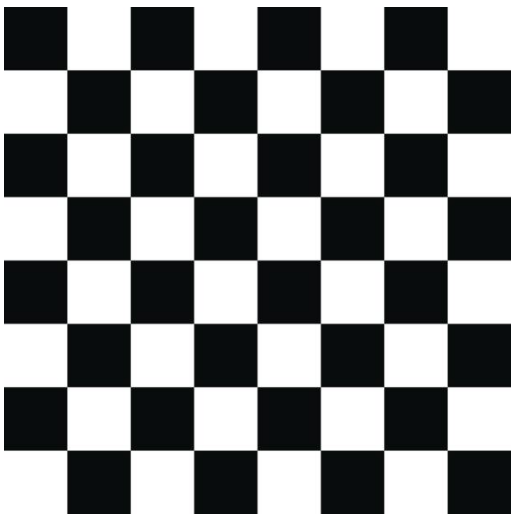

Schwarz-preconditioned HMC algorithm
for two-flavour lattice QCD

Martin Lüscher

CERN, Physics Department, TH Division
CH-1211 Geneva 23, Switzerland

Brookhaven National Laboratory

DDHMC refresher

Fermion operator may be factored:

$$\begin{pmatrix} D_\Omega & D_\partial \\ D_{\bar\partial} & D_{\bar\Omega} \end{pmatrix} = \begin{pmatrix} 1 & D_\partial D_{\bar\Omega}^{-1} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} D_\Omega - D_\partial D_{\bar\Omega}^{-1} D_{\bar\partial} & 0 \\ 0 & D_{\bar\Omega} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ D_{\bar\Omega}^{-1} D_{\bar\partial} & 1 \end{pmatrix}. \quad (2)$$

The factors L, M, and U are obvious and the determinant is:

$$\det D = \det D_\Omega \det D_{\bar\Omega} \det \left\{ 1 - D_\Omega^{-1} D_\partial D_{\bar\Omega}^{-1} D_{\bar\partial} \right\},$$

Luscher's approach:

- *small* domains 4^4 to 6^4
- HMC MD integrate gauge action and local determinants for each domain without communication
- Fits within L2 cache of a CPU core
- Small cell provides IR regulator for Dirichlet Dirac solves
- Exterior boundary gauge links are frozen (cross domain and in surface plane)

Boundary determinant

- Handling the Schur complement "boundary" determinant requires care

$$\chi = 1 - D_\Omega^{-1} D_\partial D_{\bar{\Omega}}^{-1} D_{\bar{\partial}}$$

- Luscher restricted to exterior boundary of Ω

$$R = \mathbb{P}_{\bar{\partial}} - \mathbb{P}_{\bar{\partial}} D_\Omega^{-1} D_\partial D_{\bar{\Omega}}^{-1} D_{\bar{\partial}}$$

- because in the right basis χ takes the form

$$\chi = \begin{pmatrix} 1-X & 0 \\ Y & 1 \end{pmatrix}$$

  so $\det \chi = \det R = \det(1-X)$

- For pseudofermion action $\phi_{\bar{\partial}}^\dagger (RR^\dagger)^{-1} \phi_{\bar{\partial}}$, 3D pseudofermion,

$$R^{-1} = \hat{\mathbb{P}}_{\bar{\partial}} - \hat{\mathbb{P}}_{\bar{\partial}} D^{-1} \hat{D}_{\bar{\partial}}$$

- $\delta R^{-1} = \mathbb{P}_{\bar{\partial}} D^{-1} \delta D D^{-1} D_{\bar{\partial}}$.
- Pauli-Villars (or Hasenbusch) requires

$$\phi_{\bar{\partial}}^\dagger P^\dagger R^{-\dagger} R^{-1} P \phi_{\bar{\partial}}.$$

  and $\delta R = \mathbb{P}_{\bar{\partial}} D_\Omega^{-1} (\delta D_\Omega) D_\Omega^{-1} D_\partial D_{\bar{\Omega}}^{-1} D_{\bar{\partial}} + \mathbb{P}_{\bar{\partial}} D_\Omega^{-1} D_\partial D_{\bar{\Omega}}^{-1} (\delta D_{\bar{\Omega}}) D_{\bar{\Omega}}^{-1} D_{\bar{\partial}}$.
- Doesn't work for odd flavours (nested solve needed)

What not to like about Lucher's approach

- Not friendly to odd flavours
- Luscher and collaborators don't use it - spikes in δH, too many inactive links
- Structured as lots of serial small volume evolutions: good for CPUs
- Need a GPU friendly approach: seek gain of communication avoidance, but not SMALL cell
- Make it as friendly as possible to GPUs: domain as big as a *multi-GPU exascale node*
  - Expect to lose IR bound on subdomain solves
  - Expect to gain fidelity and *large force suppression* with bigger inactive zones/more efficiency
  - Perturbative massless zero momentum two point function $\propto t^{-3}$

Two flavour determinant

$\tilde{D}$ is Dirichlet operator
Structure pseudofermion factors as:

$$\begin{aligned} \det P(D_l^\dagger D_l)^{-1} P^\dagger \quad &= \quad \det \tilde{D}_l (D_l^\dagger D_l)^{-1} \tilde{D}^\dagger \\ &\quad\quad \det \tilde{P}(\tilde{D}_l^\dagger \tilde{D}_l)^{-1} \tilde{P}^\dagger \\ &\quad\quad \det P(\tilde{P}^\dagger \tilde{P})^{-1} P^\dagger \end{aligned}$$

insert multiple Dirichlet BC Hasenbuch intermediate mass preconditioning.

$$\det \left\{ 1 - D_\Omega^{-1} D_\partial D_{\bar{\Omega}}^{-1} D_{\bar{\partial}} \right\} = \frac{\det D}{\det D_\Omega \det D_{\bar{\Omega}}},$$

2 flavour pseudofermion action:

$$S_{2f} = \phi^\dagger D_{\text{dirichlet}} (D^\dagger D)^{-1} D_{\text{dirichlet}}^\dagger \phi$$

Looks like a "Hasenbusch" operator ratio, but rather than mass preconditioning operators are identical everywhere except the boundary between domains.
Pseudofermion fills entire four volume, so different stochastic estimate of same determinant.

Force estimation

$$D_{\text{dirichlet}} = D + \Delta D$$

$$\begin{aligned} S_{2f} &= \phi^\dagger \phi \\ &+ \phi^\dagger \Delta D^\dagger D^{-1} \phi \\ &+ \phi^\dagger D^{-\dagger} \Delta D \phi \\ &+ \phi^\dagger \Delta D (D^\dagger D)^{-1} \Delta D \phi \end{aligned}$$

- Introduces surface to bulk pseudofermion coupling terms
- To be expect as "Y" term in

$$\chi = \begin{pmatrix} 1-X & 0 \\ Y & 1 \end{pmatrix}$$

Fractional powers: DD-RHMC

- What about the strange quark?
- Must be able to take fractional powers of determinant
- Alternate approach to boundary determinant allows to avoid nested inversion in Schur complement

RHMC boundary pseudofermion:

$$S_{1f}^{B} = \phi_1^{\dagger}(D_{\text{dirichlet}}^{\dagger}D_{\text{dirichlet}})^{\frac{1}{4}}(D^{\dagger}D)^{-\frac{1}{2}}(D_{\text{dirichlet}}^{\dagger}D_{\text{dirichlet}})^{\frac{1}{4}}\phi_1$$

RHMC local 1 flavor determinant ratio

$$S_{1f}^{L} = \phi_2^{\dagger}(P_{\text{dirichlet}}P_{\text{dirichlet}})^{\frac{1}{4}}(D_{\text{dirichlet}}^{\dagger}D_{\text{dirichlet}})^{-\frac{1}{2}}(P_{\text{dirichlet}}P_{\text{dirichlet}})^{\frac{1}{4}}\phi_2$$

RHMC boundary Pauli Villars

$$S_{1f}^{BP} = \phi_3^{\dagger}(P^{\dagger}P)^{\frac{1}{4}}(P_{\text{dirichlet}}P_{\text{dirichlet}})^{-\frac{1}{2}}(P^{\dagger}P)^{\frac{1}{4}}\phi_3$$

$16^3 \times 48$, $\beta = 2.13$, $m_{ud} = 0.01$ two flavour run

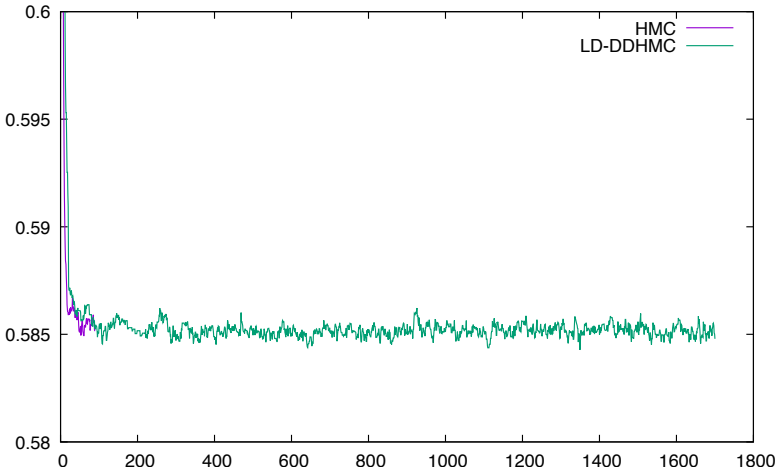

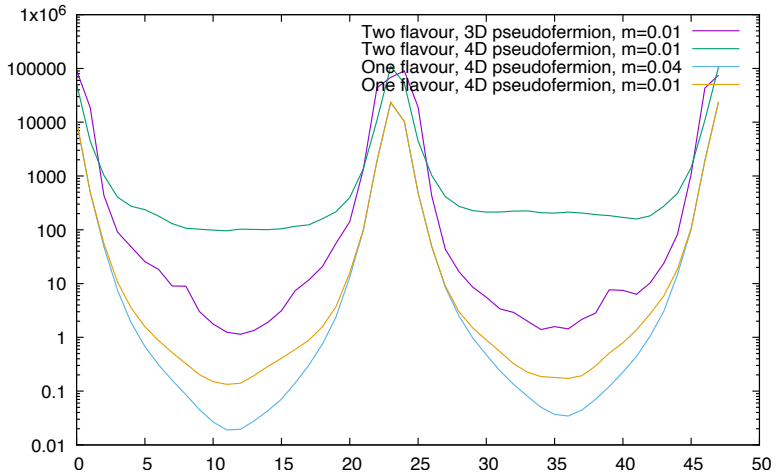

- 2 flavour 4D pseudofermion has larger force than 2f 3D pseudofermion
- 1+1 flavour 4D pseudofermion has smaller force than 2f 3D pseudofermion
- Odd flavour domain decomposition is now possible

Worth pursuing! arXiv:2203.17119

Brookhaven National Laboratory

Large volume run $\beta = 2.25$, $48^3 \times 96$, $m_{ud} = 0.00078$

- Close to physical, but $L_s = 12$ not $L_s = 24$
- After tuning, HMC trajectory on 16 Crusher nodes runs in 90mins
- For comparison: Summit - 128 nodes takes 1h
- Seek to tune DDHMC to match or exceed.
- Crusher has *good* communications: Expect any gain will be strictly limited

Nice surprise !

- Subdomain solves are MUCH faster than full solves
- Lower iteration count MULTIPLIES communicatoins gain
- Applies to all intermediate Hasenbusch factors in my DDHMC scheme

Cell	Iterations	Boundary	Conds
pppa	48x48x48x96	12351	antiperiodic
pppd	48x48x48x96	11864	open time
dddd	48x48x48x96	7349	open 4 dirs
pppd	48x48x48x48	10015	dirichlet
pppd	48x48x48x24	7416	dirichlet
ppdd	48x48x24x24	5150	dirichlet
pddd	48x24x24x24	4324	dirichlet
dddd	24x24x24x24	3692	dirichlet

Why so much? Look at difference in low mode eigenspectrum

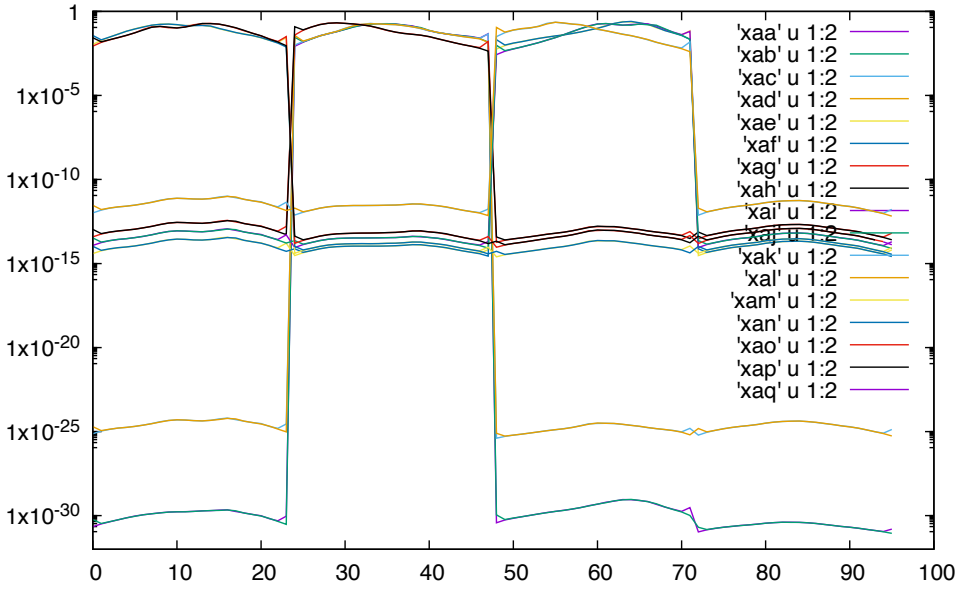

- Lowest lying eigenvectors of the (block diagonal) Dirichlet operator on four $48^3 \times 24$ cells in a $48^3 \times 96$ lattice.
- Modes are projected into $1 \pm \gamma_t$ parity
- components to look for edge effects at domain boundaries.

Brookhaven National Laboratory

Why so much?

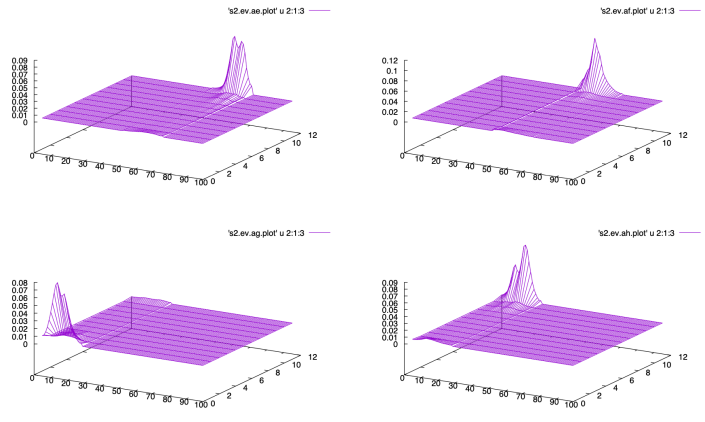

Lowest lying eigenvectors of the (block diagonal) Dirichlet operator on four $48^3 \times 24$ cells in a $48^3 \times 96$ lattice. XYZ slice norm is plotted vs 5th dimension coordinate and time plane coordinate. Modes are chiral in nature with small edge effects in the time direction. Mostly centred in the middle of the domains, and likely modes or instantons crossing domain boundaries are chopped with lifting of the eigenmode leaving only central objects.

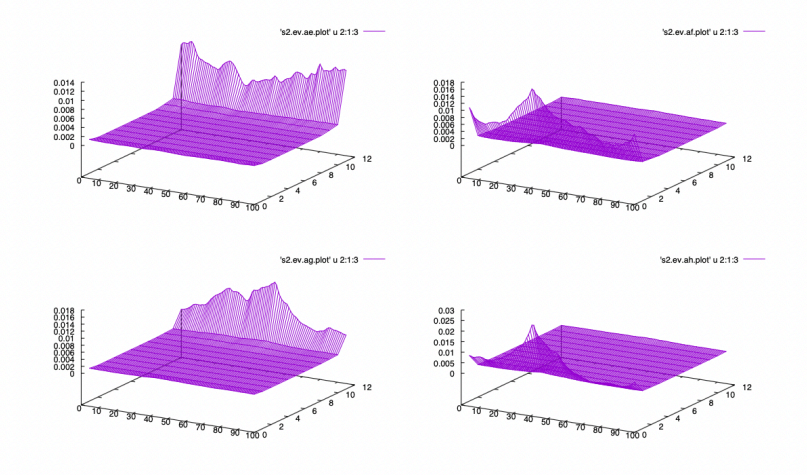

- Dirichlet BC on modes forces them to zero at boundary
- This forces (covariant) curvature and raises eigenvalue
- Enough to reduce solver iterations 4x
- Doesn't really support cute theories about cutting up instantons, global vs local topology (!)

Brookhaven National Laboratory

Why so much?

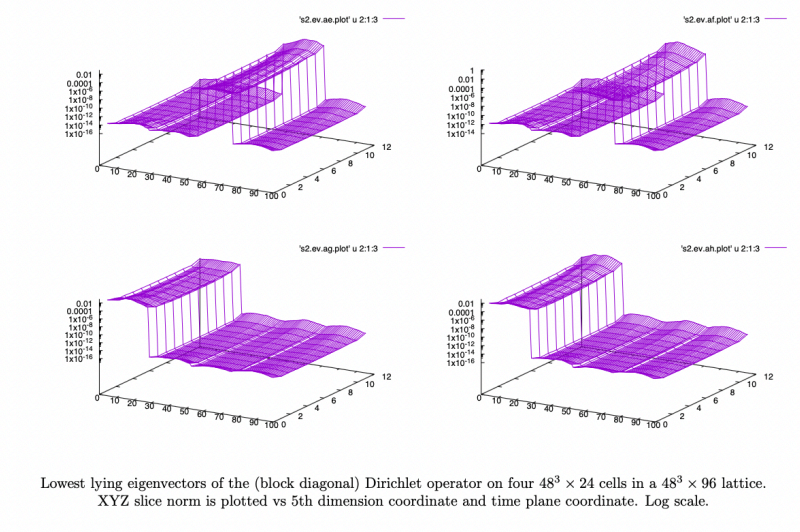

Lowest lying eigenvectors of the (block diagonal) Dirichlet operator on four $48^3 \times 24$ cells in a $48^3 \times 96$ lattice. XYZ slice norm is plotted vs 5th dimension coordinate and time plane coordinate. Log scale.

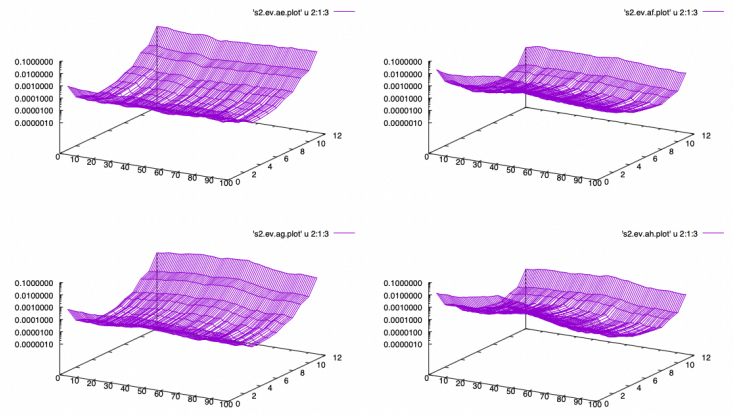

- Dirichlet BC on modes forces them to zero at boundary
- This forces (covariant) curvature and raises eigenvalue
- Enough to reduce solver iterations 4x
- Doesn't really support cute theories about cutting up instantons, global vs local topology (!)

Brookhaven National Laboratory

48^3 forces

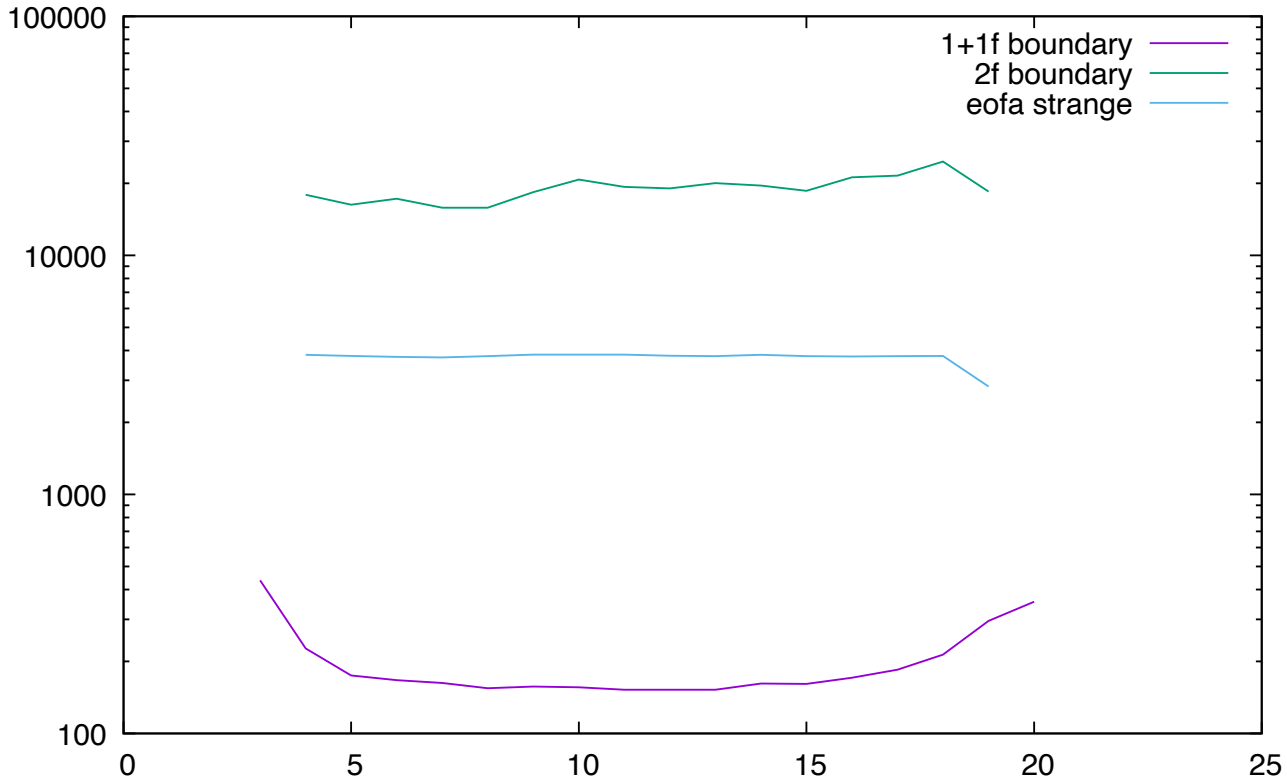

- Pattern of larger force with 2f boundary repeats
- Continue with rational 1+1 f boundary
- Introducing **partial** dirichlet BCs with further force reduction
  - Surface physical fields remain connected, 5d bulk is disconnected

Brookhaven National Laboratory

48^3 tuning status - Prelimibary Work In Progress

Work in progress:

- Do NOT yet have a wall clock gain or parity.
- Started to conserve Hamiltonian for 1f + 1f boundary determinant
- slightly slower time O(7000s) vs O(5500s)
  - I am recalculating an *expensive* Hamiltonian redundantly
  - Running Rational approximation in *double* precision and comparing to 2x faster mixed precision HMC
- Working on mixed precision multishift CG
- Working on tuning residuals pole by pole in rational approximation
- Working on right integrator timestep heirarchy
- Plan to look at multishift-multi-RHS (Keegan+De Forcrand arXiv:1808.01829)
- Two possible 2x speed ups, mixed precision more certain

Preliminary Sketch - this is NOT YET an equal acceptance comparison

algorithm	D_{full}	$D_{\text{dirichlet}}$
HMC	421648	0
DDHMC	161708	357316

Brookhaven National Laboratory

Summary

- Grid gives good cross platform single source performance
- CPU/GPU unified programming that is well optimised for both
- Engineering limits like bandwidths are hit
- Multiscale physics AND computer engineering BOTH dictate new algorithm research
- $\Rightarrow$ SciDAC-5 research directions
  - Multigrid for DWF (and staggered) are intrinsically harder than Wilson
  - Advanced multiscale aware gauge sampling
  - Domain decompostion to improve computational locality