Optimizing staggered multigrid for exascale performance

Venkitesh Ayyar,
Boston University

Lattice 2022, Bonn
August 11, 2022

In collaboration with Evan Weinberg, Kate Clark, Richard Brower

Solvers in lattice QCD

Use of solvers for $Ax = b$

→ Computing observables

→ Gauge field generation to compute change in action

In the physical region, need

$a \to 0$ and $am_q \to 0$

Critical slowing down

$A = \gamma_\mu D_\mu + m$

Small masses $\implies$ small eigenvalues

super-linear slowing of conventional solvers

Need 30,000 iterations with CG
for $192^3 \times 384$ lattice

EXASCALE COMPUTING PROJECT

BOSTON UNIVERSITY

Dealing with Critical Slowing down

Block Krylov solvers

Doesn't fundamentally address
Critical slowing down

Eigenvalue Deflation

Subtract out low eigenvectors

Doesn't scale well with matrix size

Multigrid

Coarsen operator

Apply solver on coarser grids

Smooth out low modes

Correct solution on finer levels using solution on coarser levels

Multigrid combined with Deflation at coarse levels works even better !

[1]Brannick, Brower, Clark, Osborn, Rebbi PRL 100, 041601 (2008)
[2]Brannick, Frommer, Kahl, Leder, Rottmann, Strebel: arxiv 1410.7170
[3]Whyte, Wilcox, Morgan arxiv : 1912.02868

General idea of Adaptive Multigrid

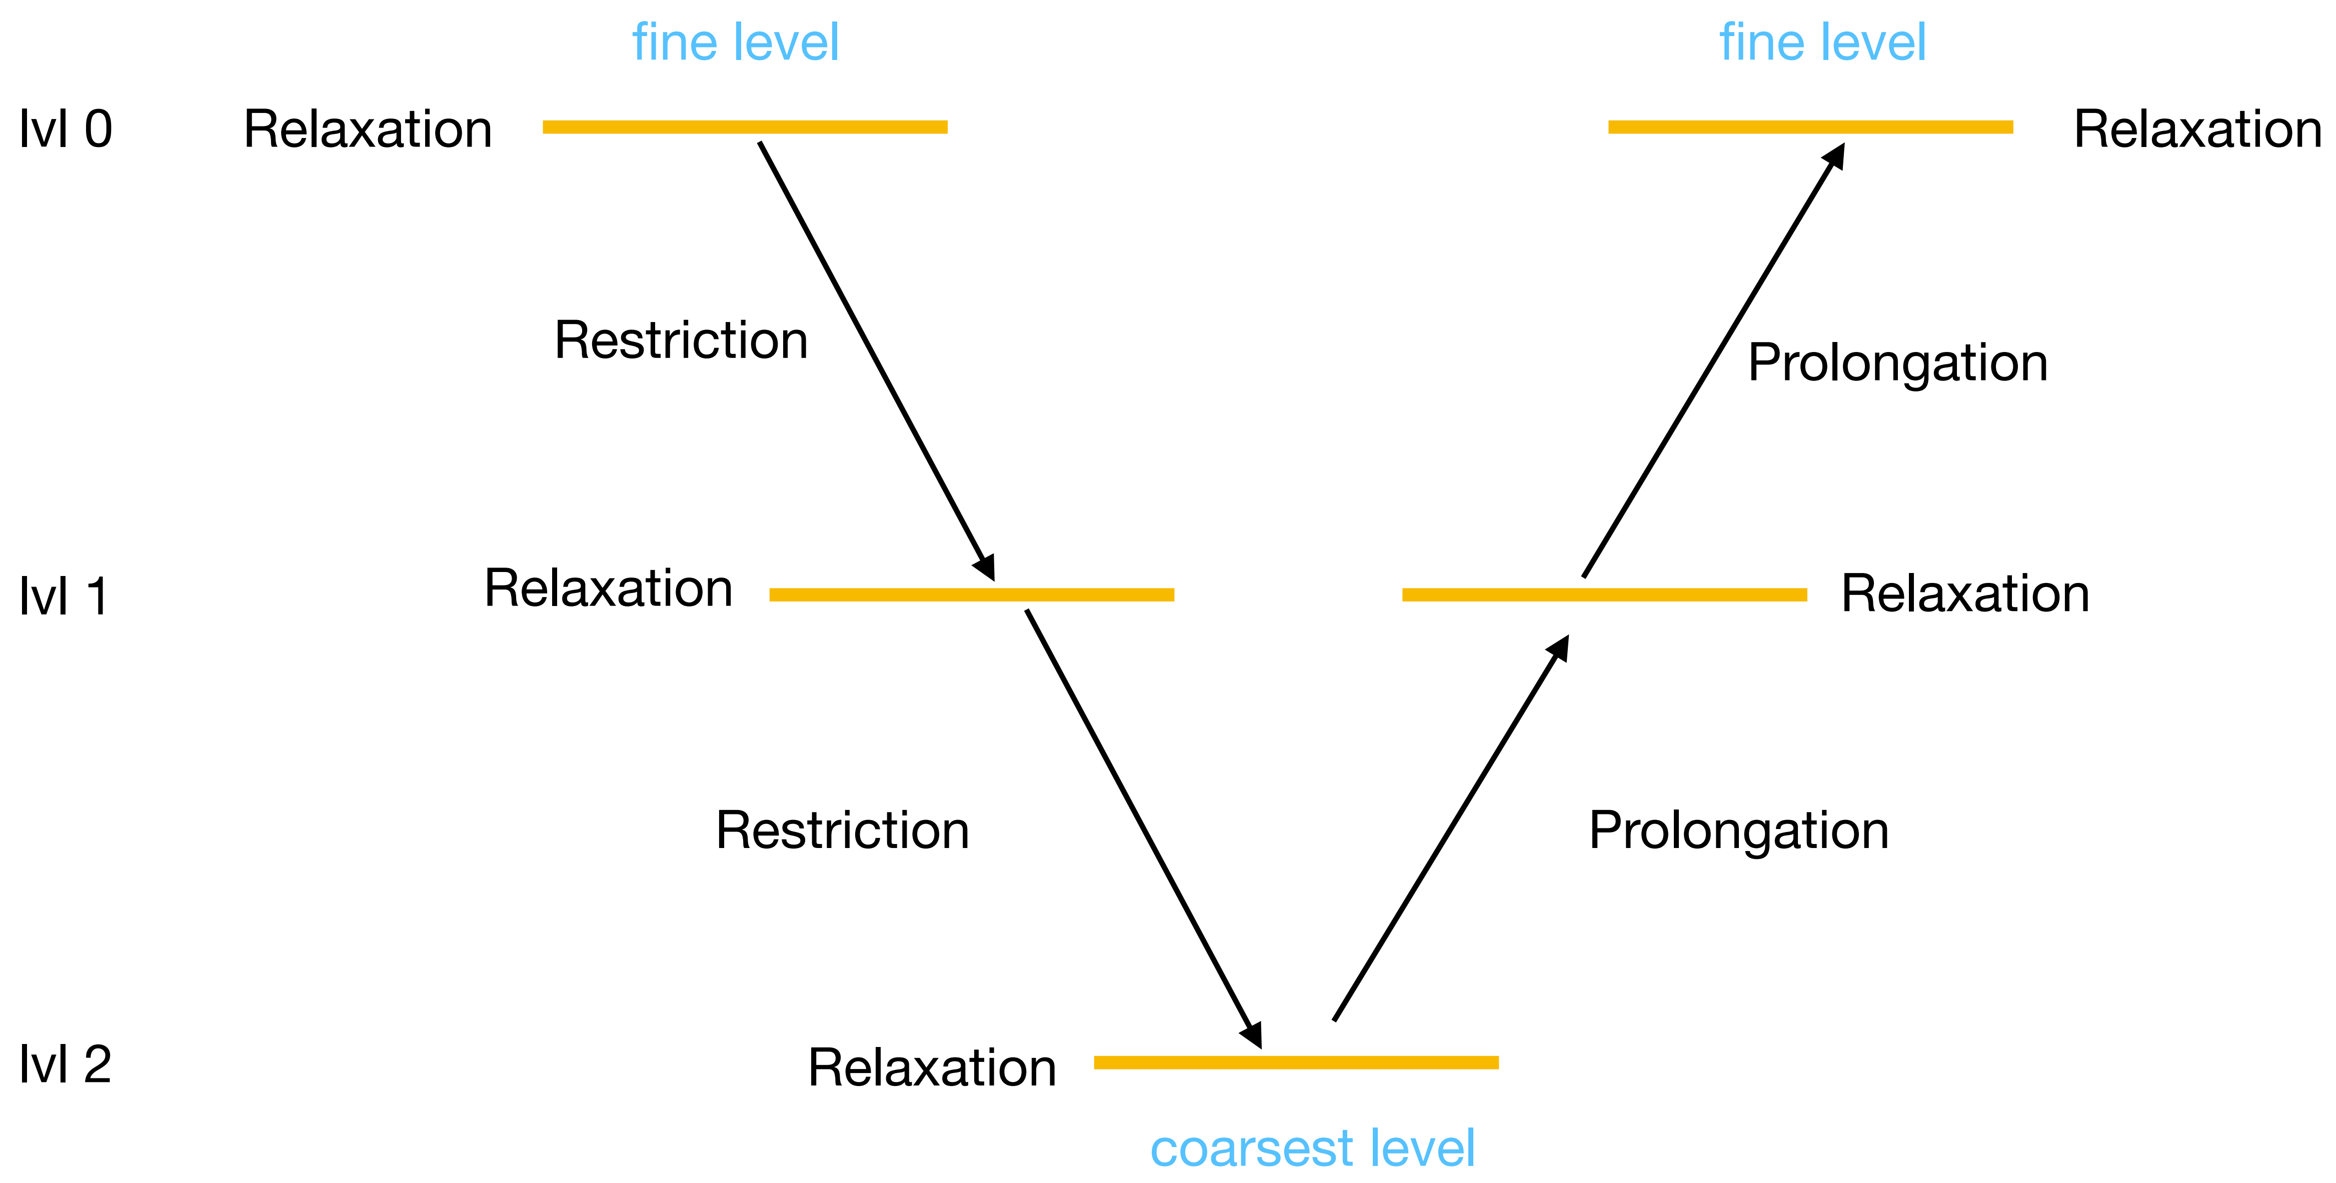

Relaxation	Solver A x = b iteratively
Restriction	Move to coarser grid
Prolongation	Move to finer grid

Adaptive multigrid procedure

Generate near-null vectors $Ax = 0$

Build restriction and prolongation operators from them

Apply solver on coarser lattices to smooth low modes

EXASCALE COMPUTING PROJECT

BOSTON UNIVERSITY

Multigrid Wilson vs Staggered

Figures from arXiv 1801.07823

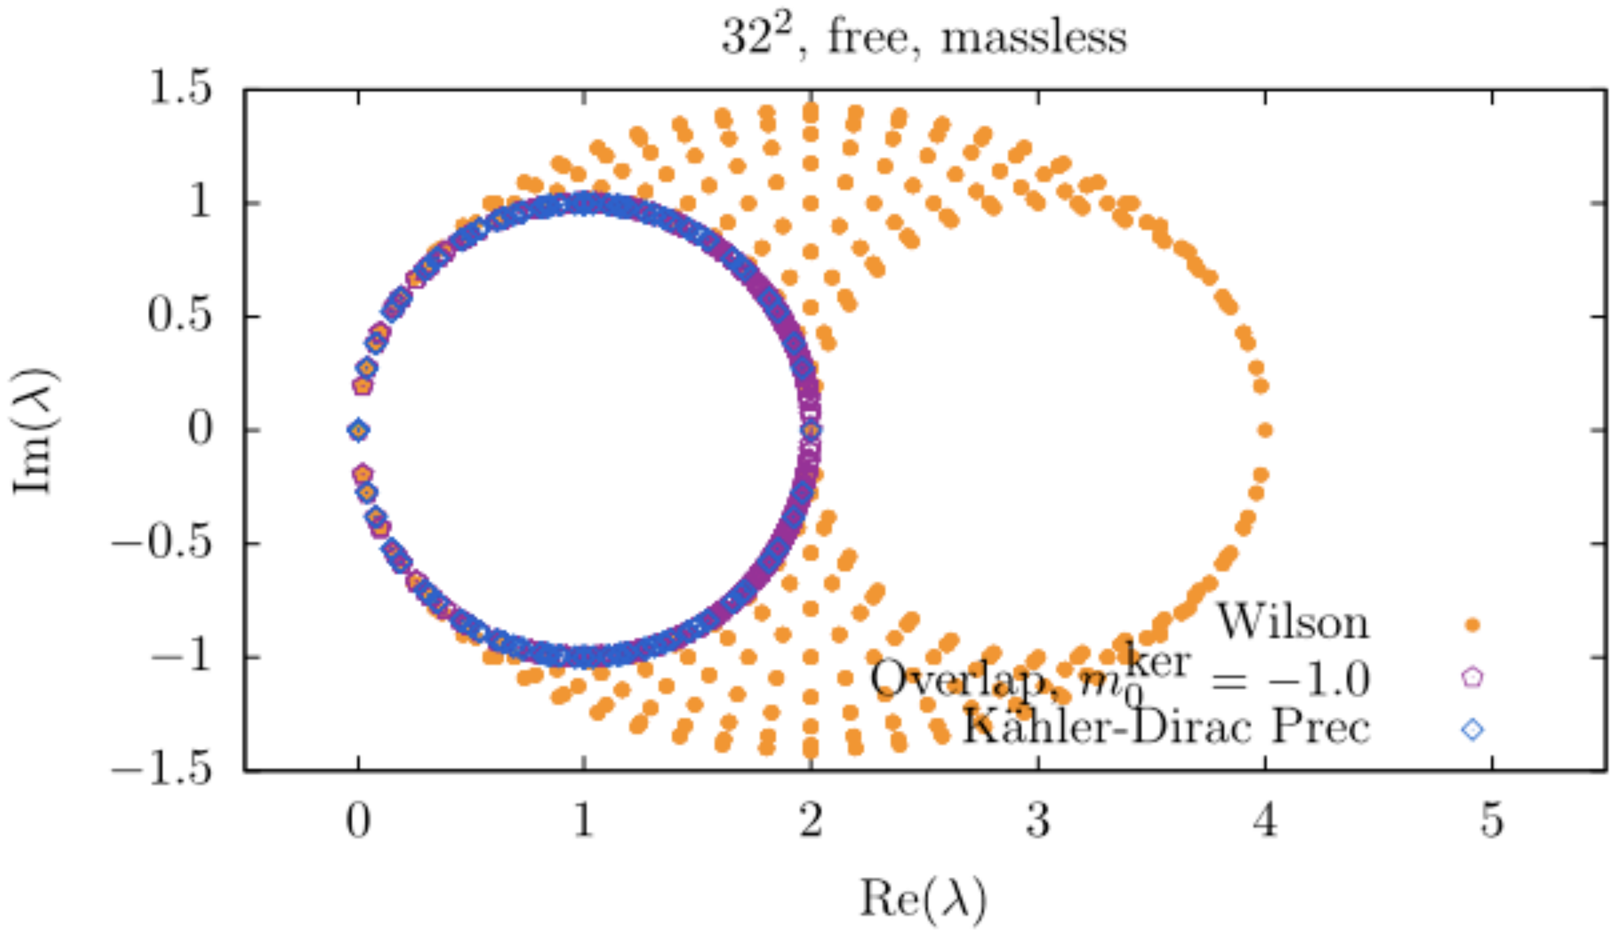

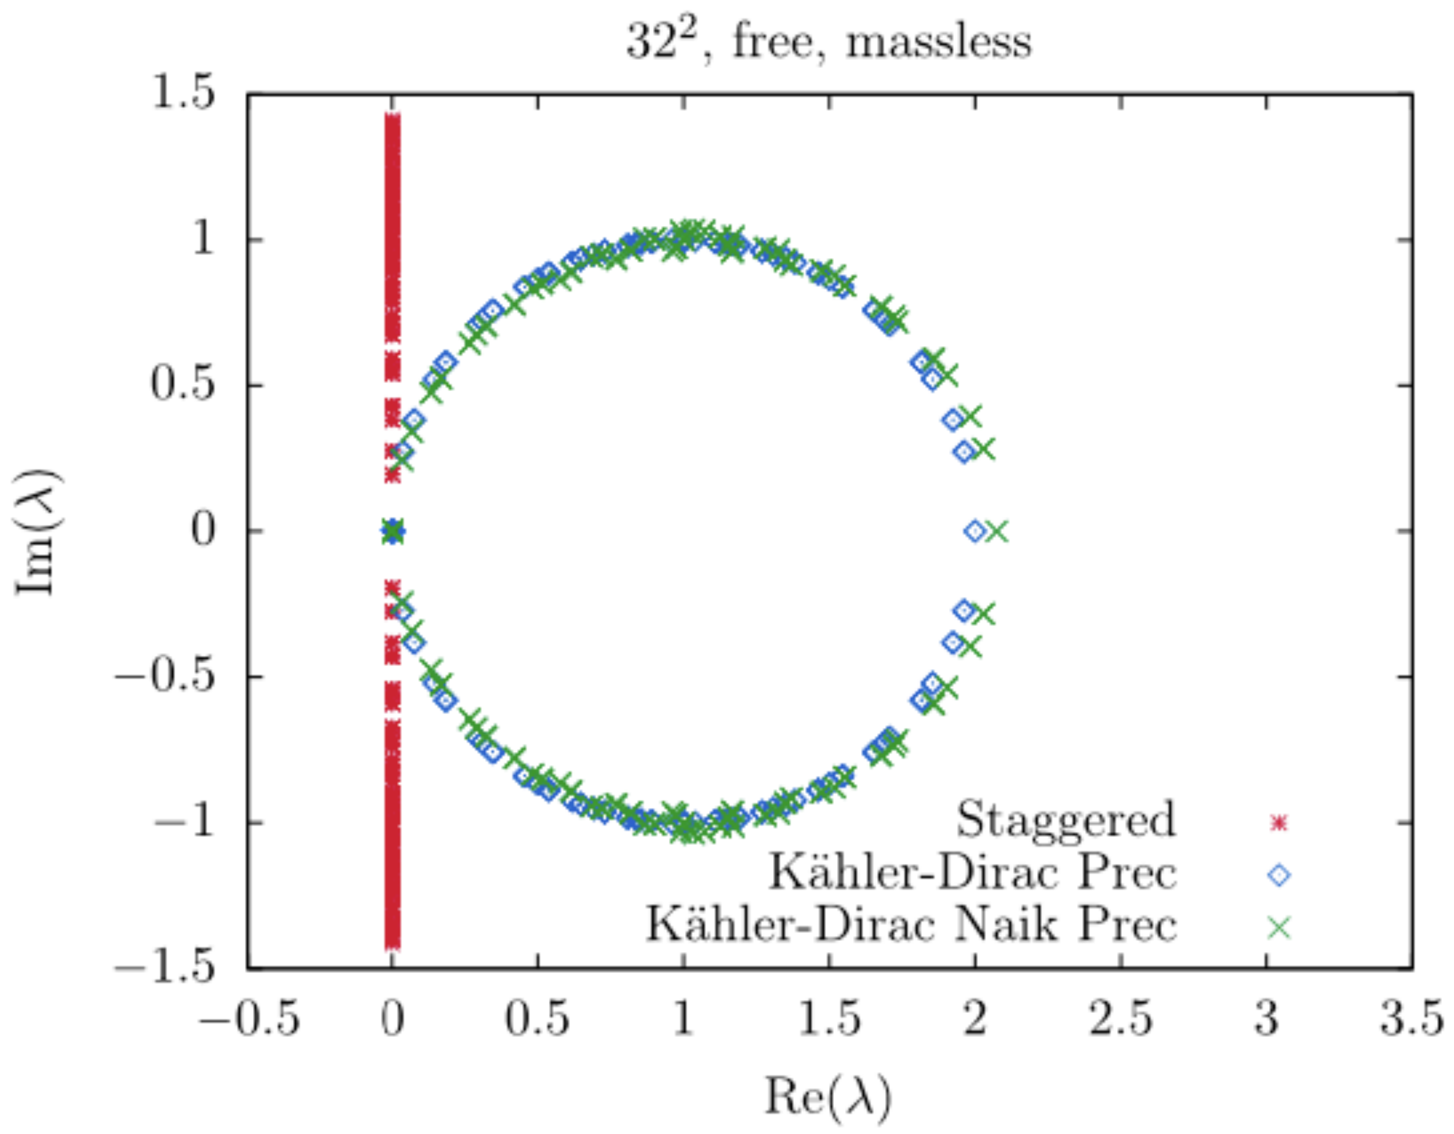

Wilson fermions : Brannick, Brower, Clark, Osborn, Rebbi PRL 100, 041601 (2008)

Circle-like spectrum for Wilson and Domain-wall operator

Standard Adaptive Multigrid implementation

Staggered fermions:

Straight line : No straight-forward implementation of Multigrid

Kahler-Dirac preconditioner : Spectrum becomes Wilson-like

Brower, Clark, Strelchenko, Weinberg, Phys Rev D 97, 2018

EXASCALE COMPUTING PROJECT

BOSTON UNIVERSITY

Staggered Multigrid

Kahler-Dirac preconditioner at first levels

Multiple coarse layers

Deflation on coarsest level

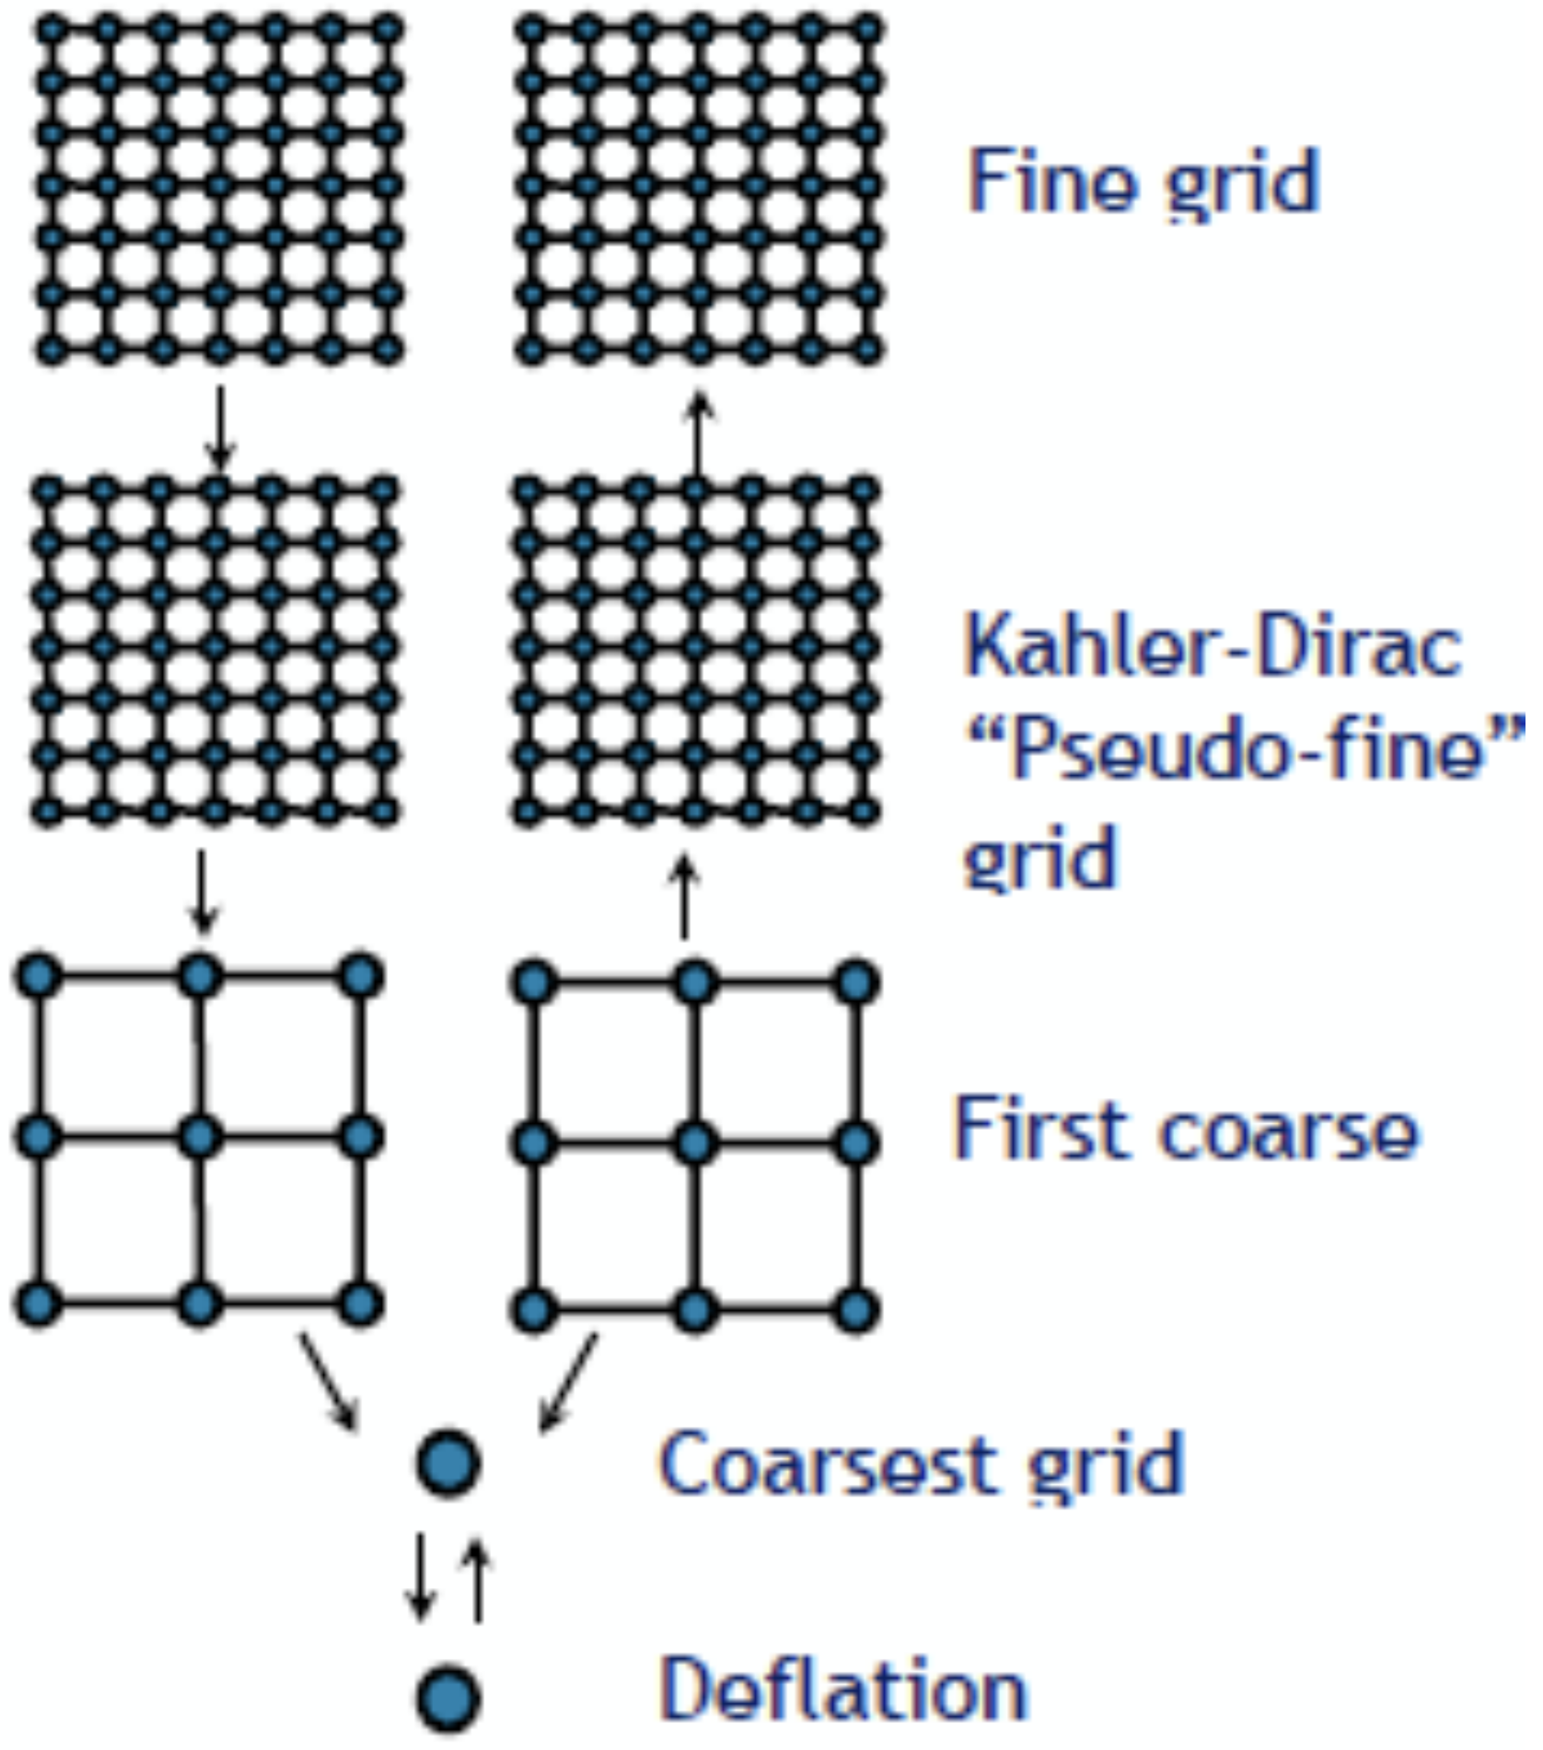

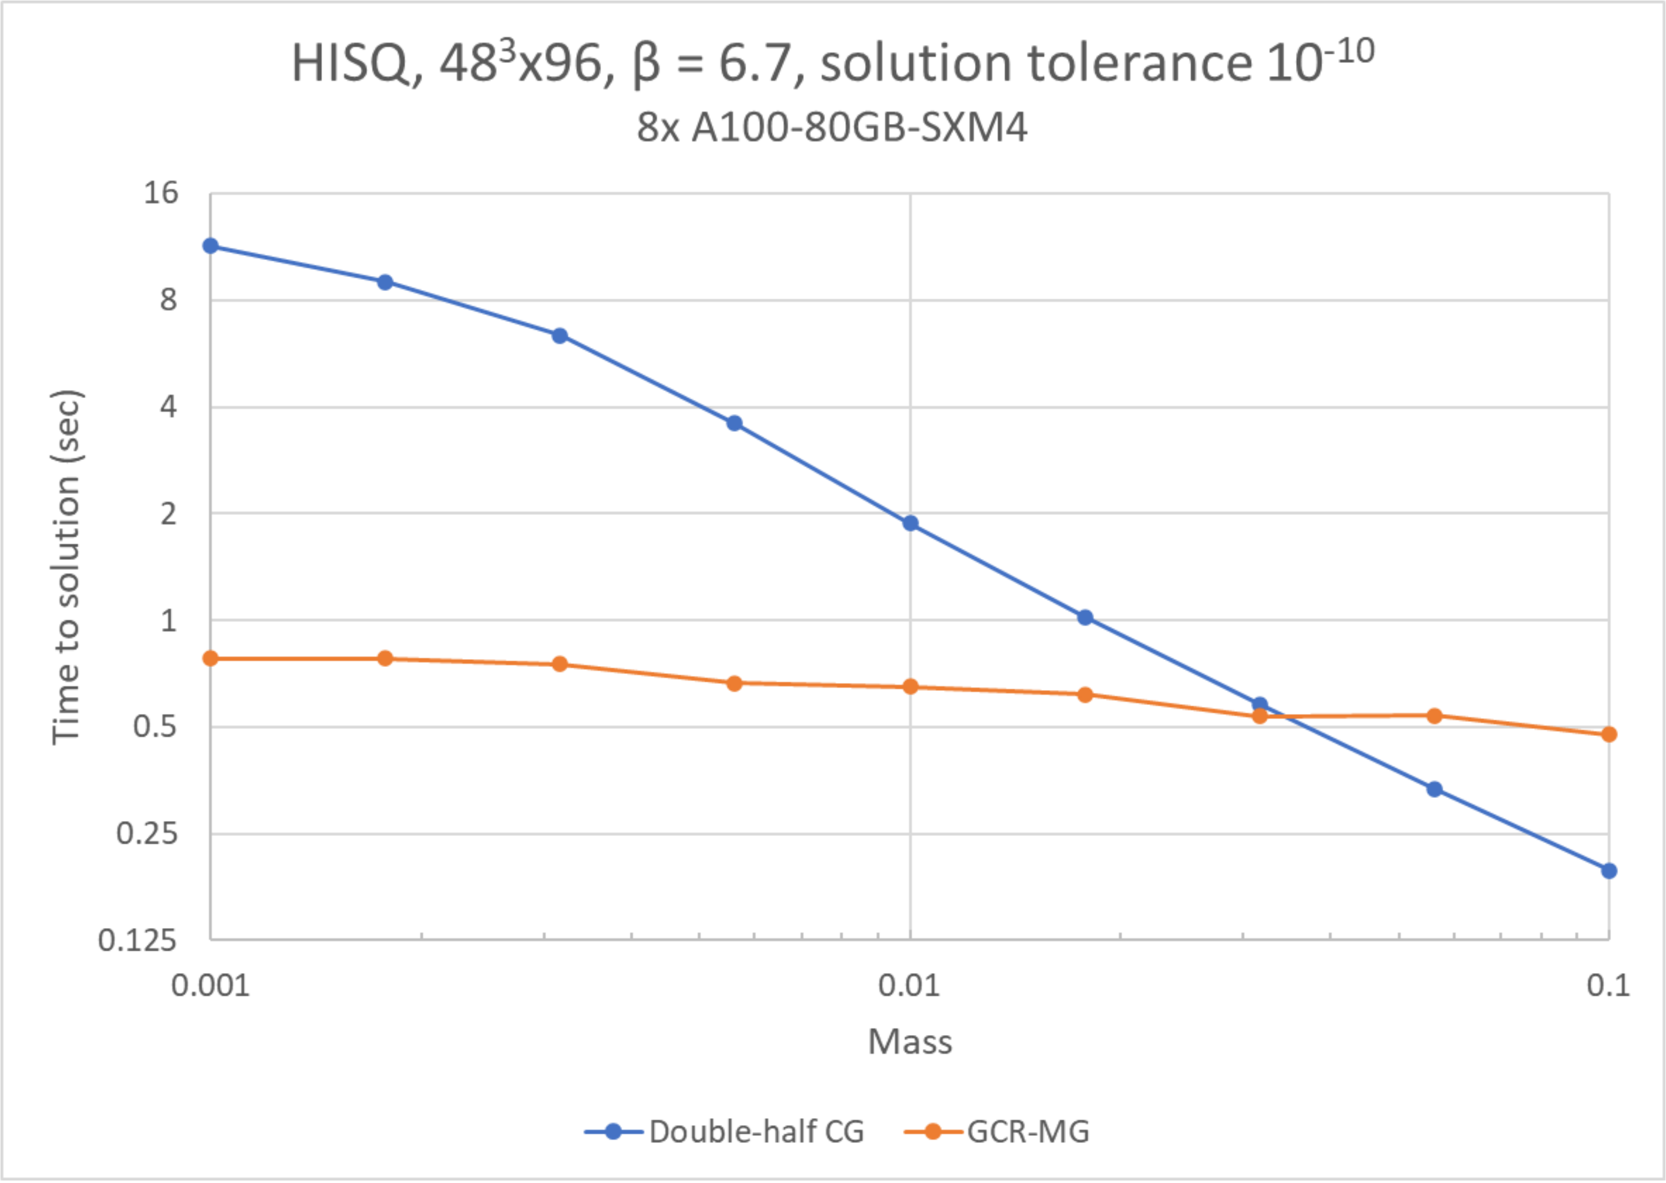

Multigrid CG comparison

Multigrid workflow details

General MILC workflow : 10 light masses

Pure CG workflow

All 10 light masses with multi-shift CG

Multigrid workflow

Peel of a few lighter masses (say 3) and use MG

Remaining 7 masses solved using multi-shift CG

Multigrid setup cost > CG setup

With enough # of solves, MG setup cost effect can be mitigated

Run details

Performance comparison

Multigrid vs **Conjugate Gradient**

KD-inverse +

HISQ operator

Pure Multi-shift CG

Machine	Lattice size	Nodes	GPUs
Summit (OLCF)	144^3 x 288	144	864

Summit Supercomputer at Oak Ridge national Lab

#4 on Top500 supercomputer list

NVIDIA GPUs and QUDA software

Future tests on Exascale machine **Frontier**

#1 on Top500 supercomputer list

Physical pion mass, lattice spacing a = 0.04 fm

We thank the MILC Collaboration for providing this configuration for our studies.

ECP EXASCALE COMPUTING PROJECT

BOSTON UNIVERSITY

Distribution of lattice blocks among GPUs

Summit
- 6 GPUs per node
- 18 nodes per rack

Multigrid with HISQ

Lattice	144, 144, 144, 288
Node geometry	6, 3, 6, 8
Local volume per GPU	24, 48, 24, 36
Nodes	144

Lattice dimensions : x, y, z, t

Blocking scheme

	Dimensions
Level 0	24, 48, 24, 36
Block 1	4, 6, 6, 6
Level 1	6, 8, 4, 6
Block 2	3, 2, 2, 3
Level 2	2, 4, 2, 2

Timing tests on Summit

Run type	MG-setup	Solve time per prop.	Total time (per 144 nodes)
Multi-grid	1573	329	1902
Multishift CG	-	768	768

Extending to 100 propagators per solve (projected timings)

Run type	MG-setup	Solve time	Total time (per 144 nodes)
Multi-grid	1573	32900	34473
Multishift CG	-	76800	76800

Solve time is for 10 light masses

Timings compared after scaling for same number of nodes

Next steps

Implementation

- Run with config on lattice $192^3 \times 384$ on Summit
- Build and run on Crusher at OLCF (lattice size $96^3 \times 192$) (32 nodes)
- Full run on Exascale machine **Frontier** with $192^3 \times 384$ lattice

Optimizations

- Investigating Chebyshev methods for near-null vectors[1]
- More aggressive blocking (perhaps 8,8,8,16)

Improvements

Non-telescoping

[1]Boyle, Yamaguchi arxiv: 2103.05034

EXASCALE COMPUTING PROJECT

BOSTON UNIVERSITY

Non-telescoping idea

Downsampling in Multigrid reduces
number of parallel operations

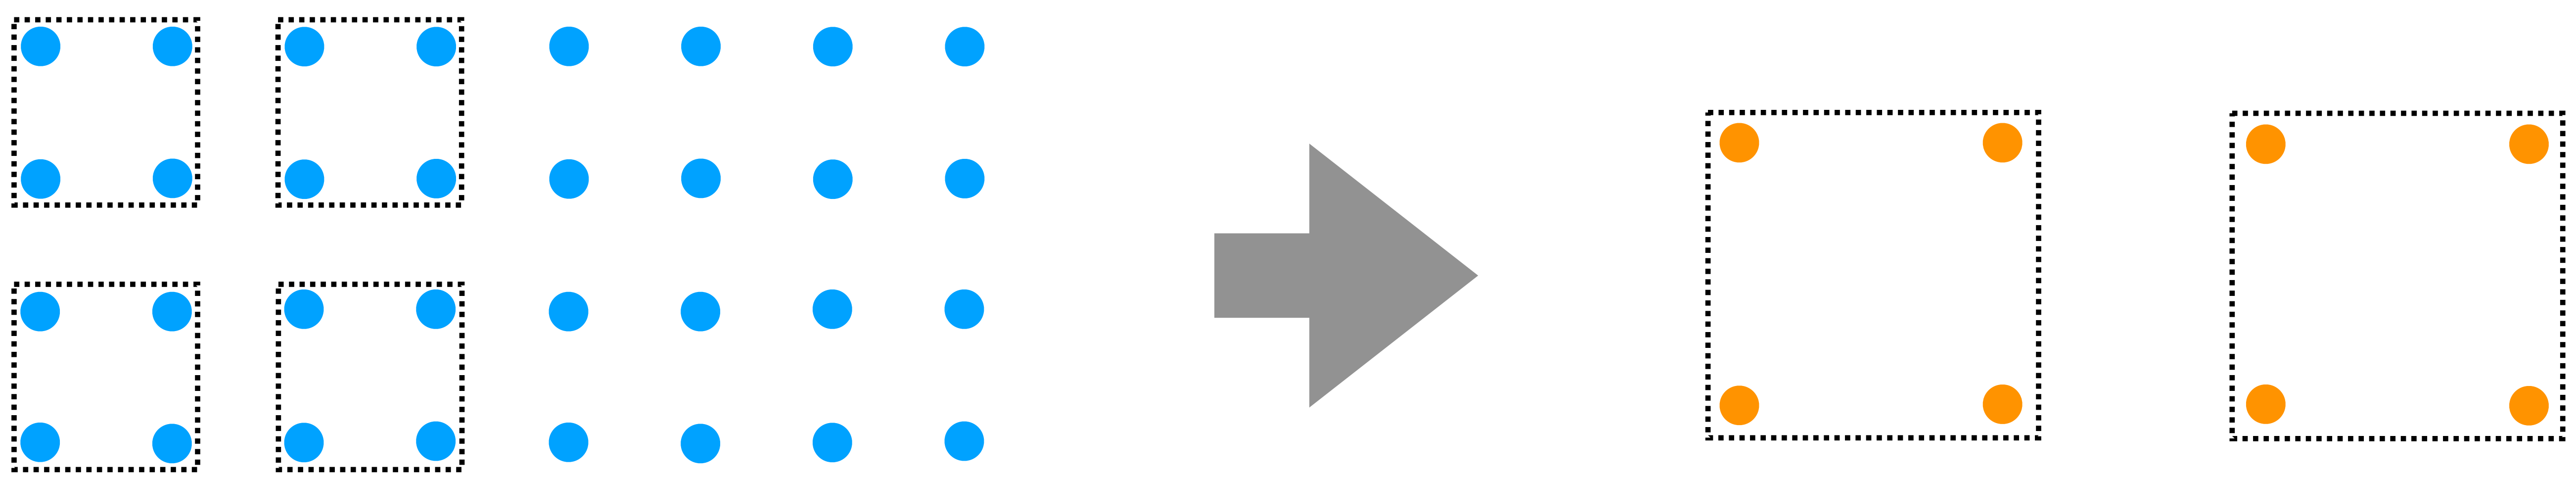

Can we do extra work to oversaturate GPUs ?

Benefit Multigrid ?

Faster convergence ?

Inherent choice of blocking scheme in Multigrid

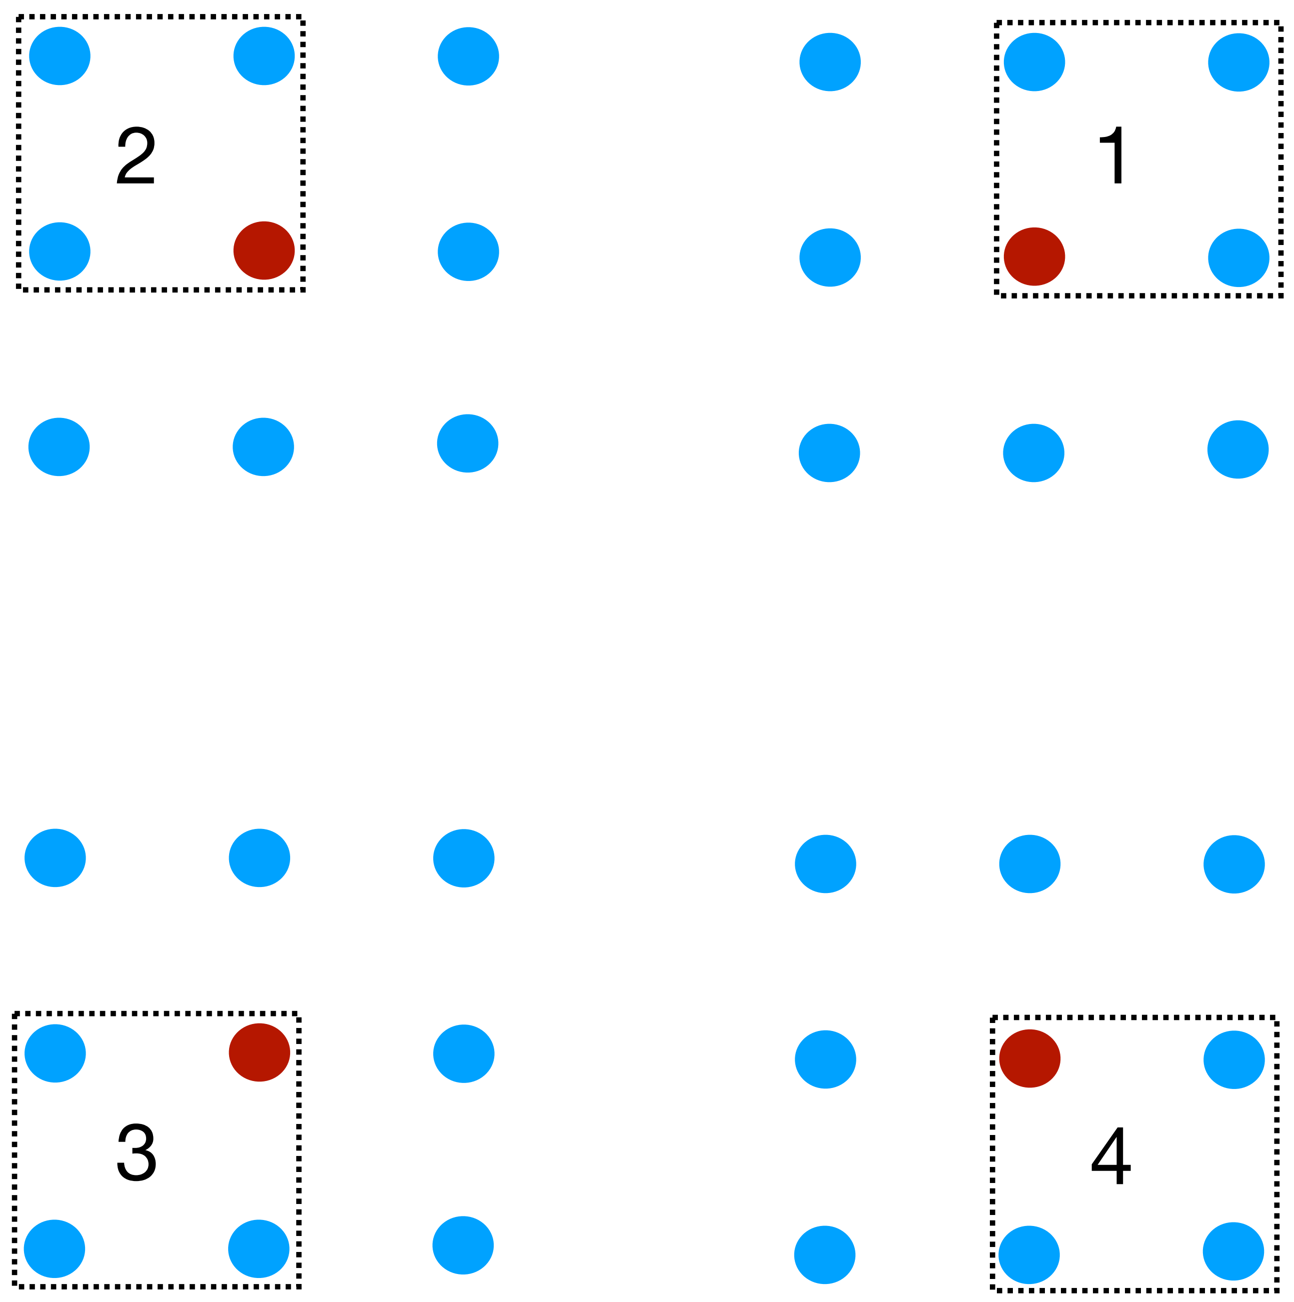

Can block with some or all ways

Utilize information from different copies

Capture correlations between farther sites

Non-telescoping idea

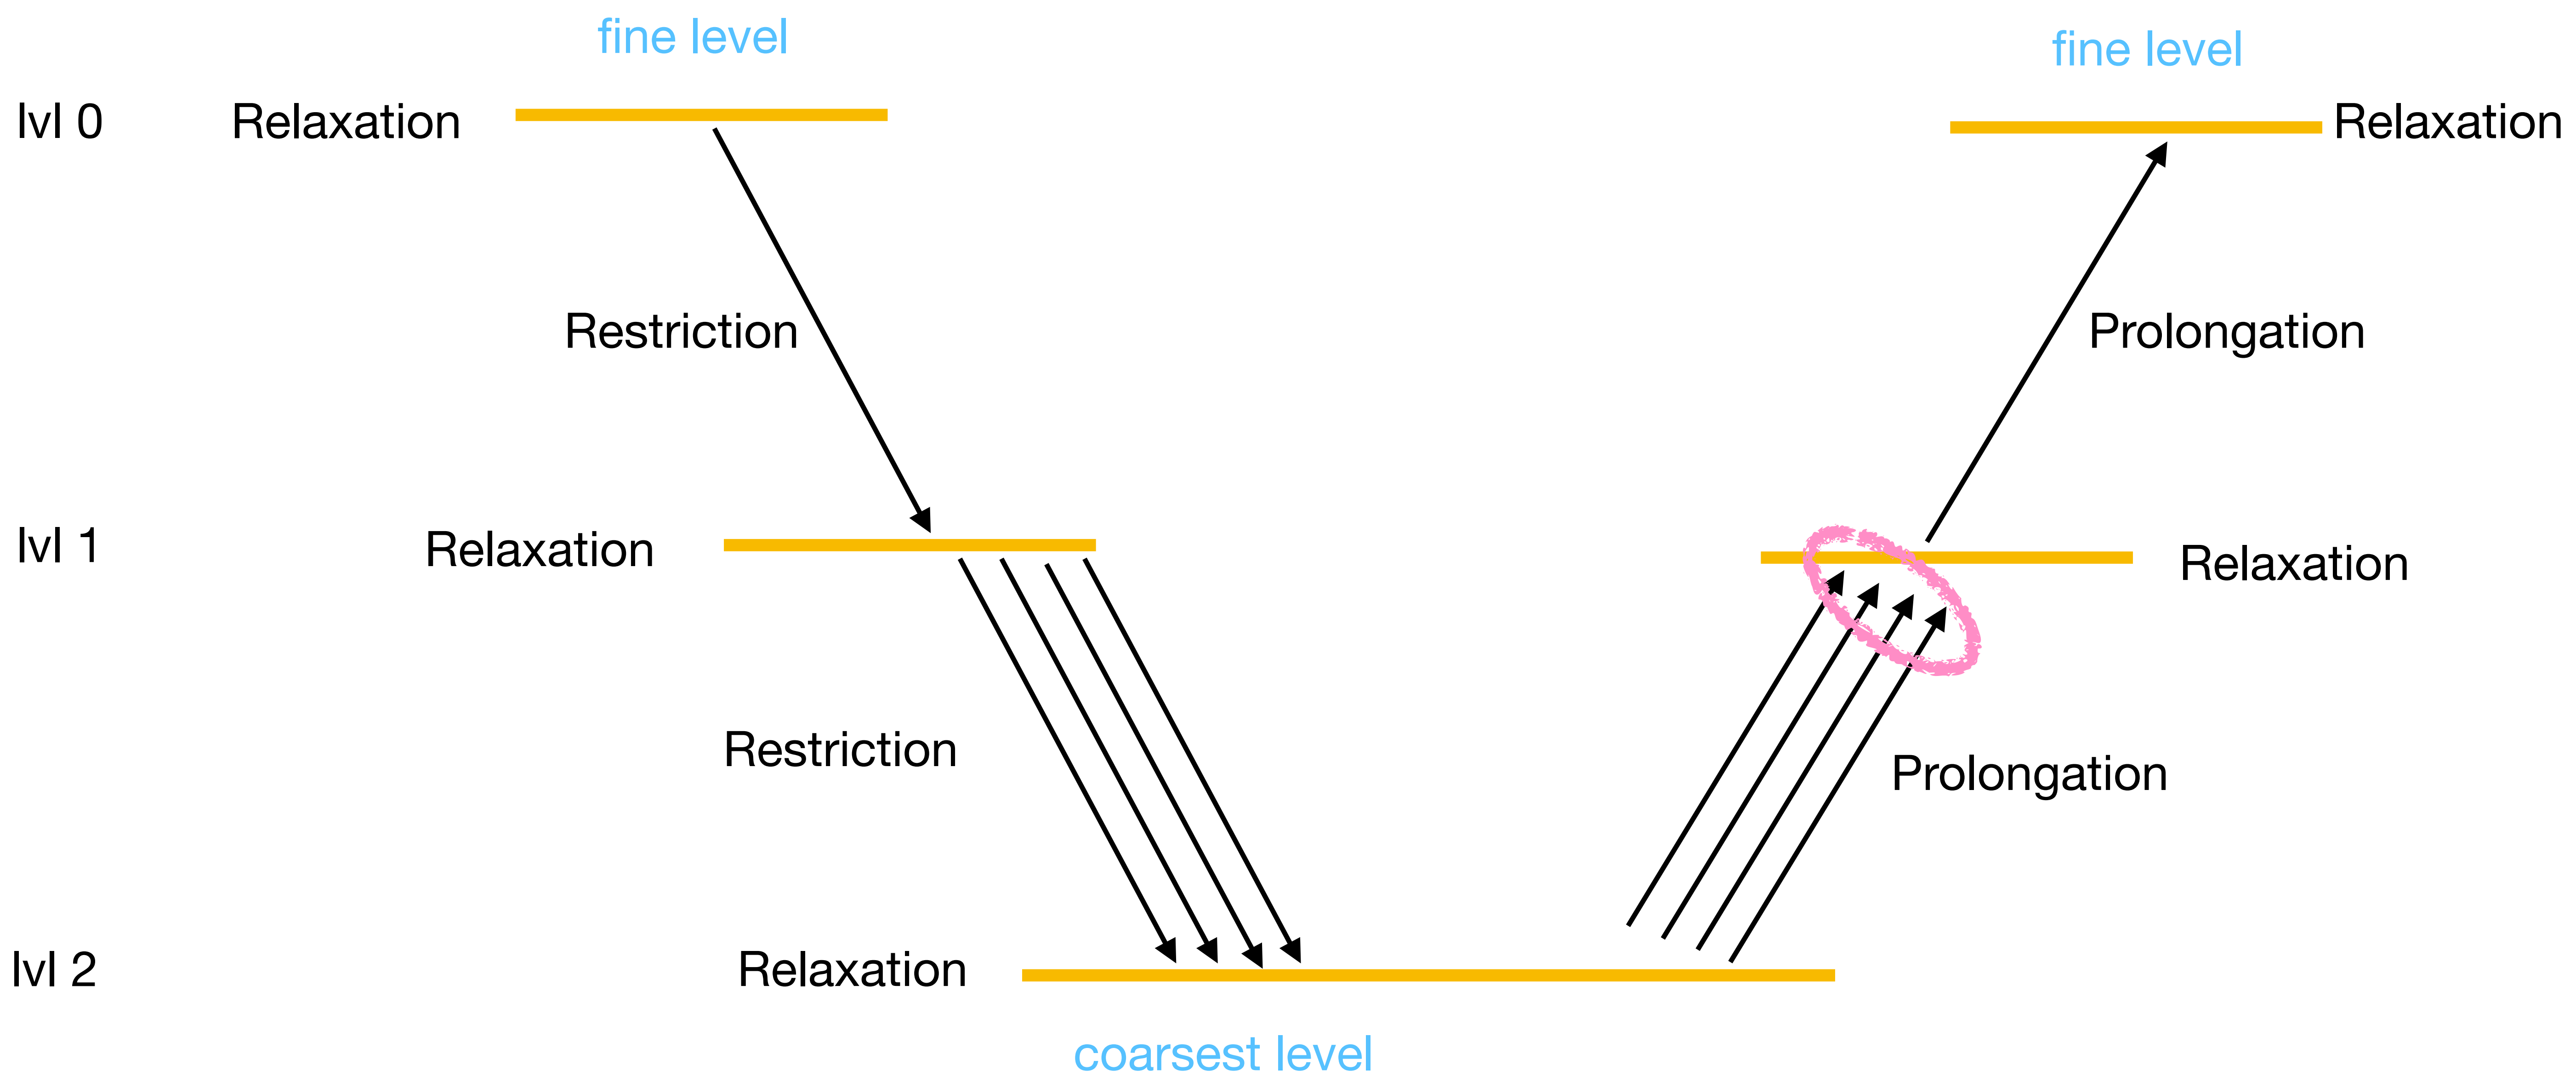

Restrict 4 copies to lowest level

Prolongate errors back and combine

eg: Minimal residual

Exploring with Laplace operator in 2D

EXASCALE COMPUTING PROJECT

BOSTON UNIVERSITY

Collaborators

Summit

supercomputer (Oak Ridge National Lab)

Richard Brower — BOSTON UNIVERSITY

MILC https://github.com/milc-qcd

Libraries

QUDA https://github.com/lattice/quda

https://github.com/lattice/quda/wiki/Multigrid-Solver

Evan Weinberg

Kate Clark — NVIDIA

Other talks at Lattice 2022

QUDA optimizations for LQCD

Kate Clark, 10:00 Thrs

Mathias Wagner, 16:30 Mon

Thank you

Backup slides

Multigrid details

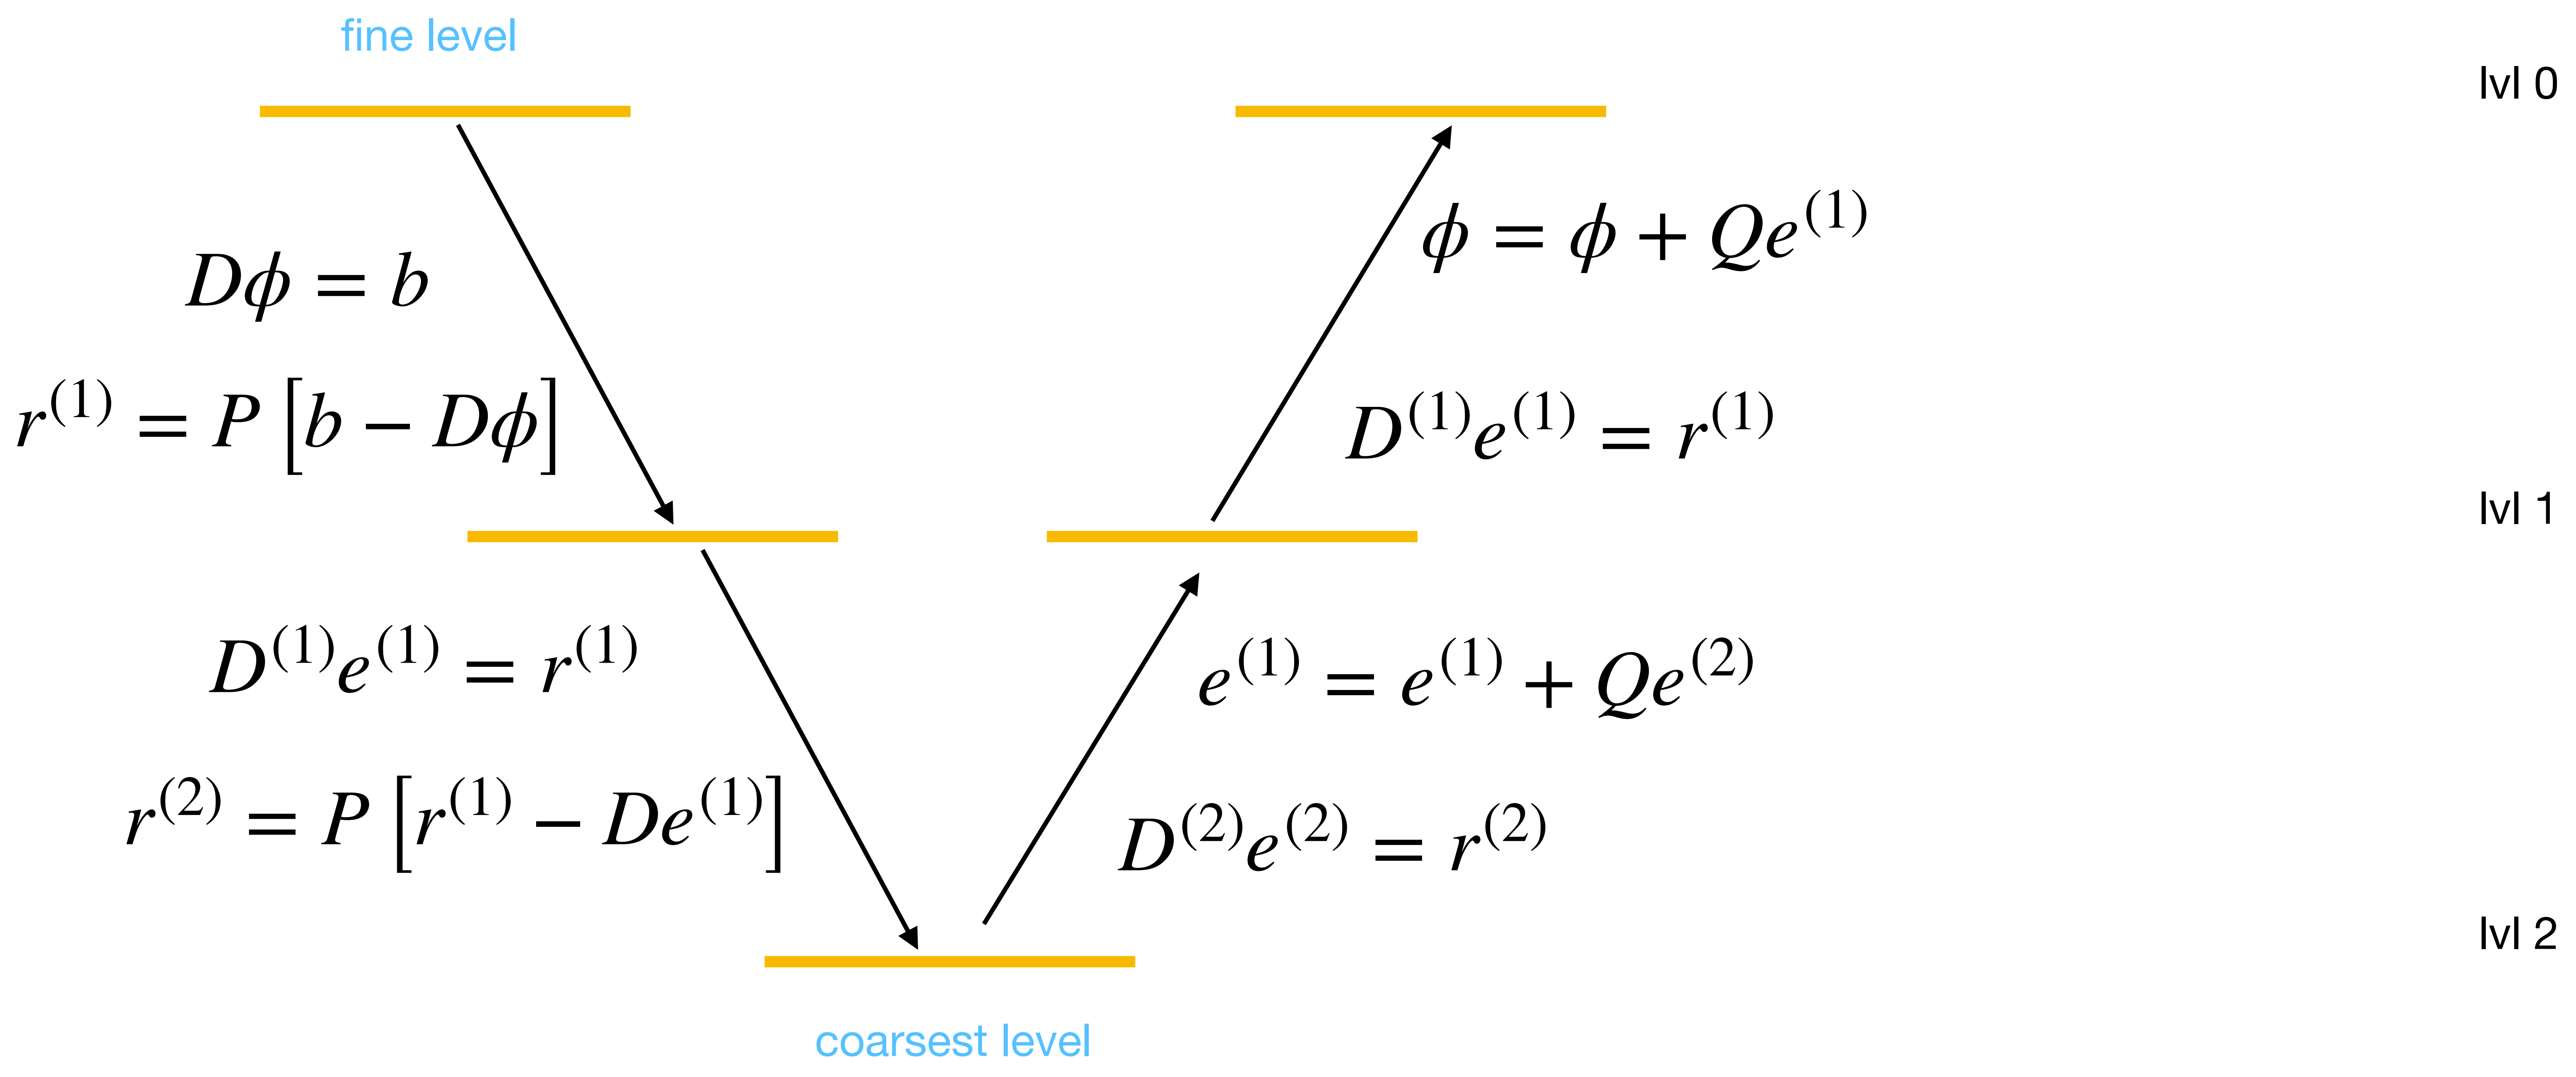

Exact solution

$$D\phi^* = b$$

Residue $r = b - D\phi$

Error $e = \phi - \phi^*$

$$Ae = r$$

P ≡ Restriction

Q ≡ Prolongation

Slide from Evan Weinberg

THEORY: KAHLER-DIRAC PRECONDITIONING

arXiv:1801.07823

- Key observation: the 2^d hypercube of degrees of freedom is equivalent to a Kahler-Dirac fermion (in the free field)
- Write the staggered operator as a dual-decomposition: $D_{stag} = B + C + m$
    - B: hopping terms within a 2^d block
    - C: hopping terms across blocks
- Perform a block-preconditioning by $(B+m)^{-1}$

$$D_{KD}(m) = (B+m)^{-1}[(B+m)+C]$$
$$= \mathbb{I} + (B+m)^{-1}C$$

Result: overlap-esque spectrum

- Perfect circle in free field
- "Fuzzy" when interactions are enabled

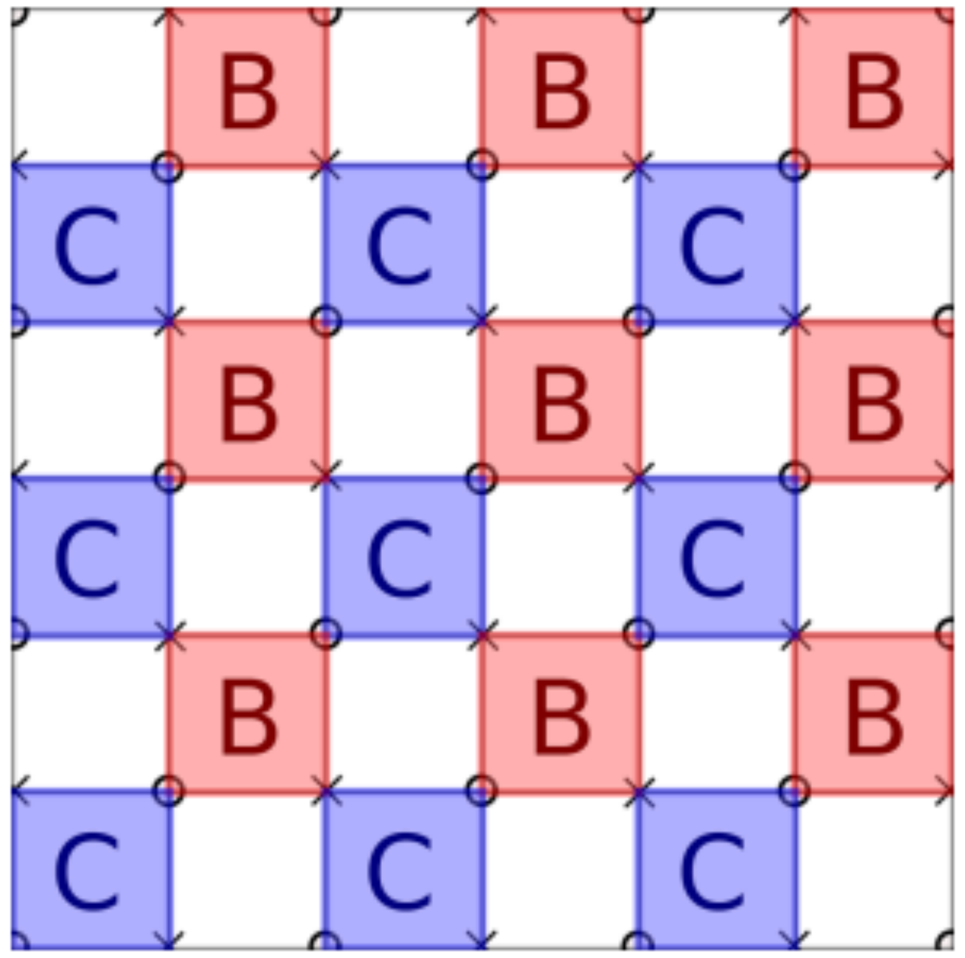

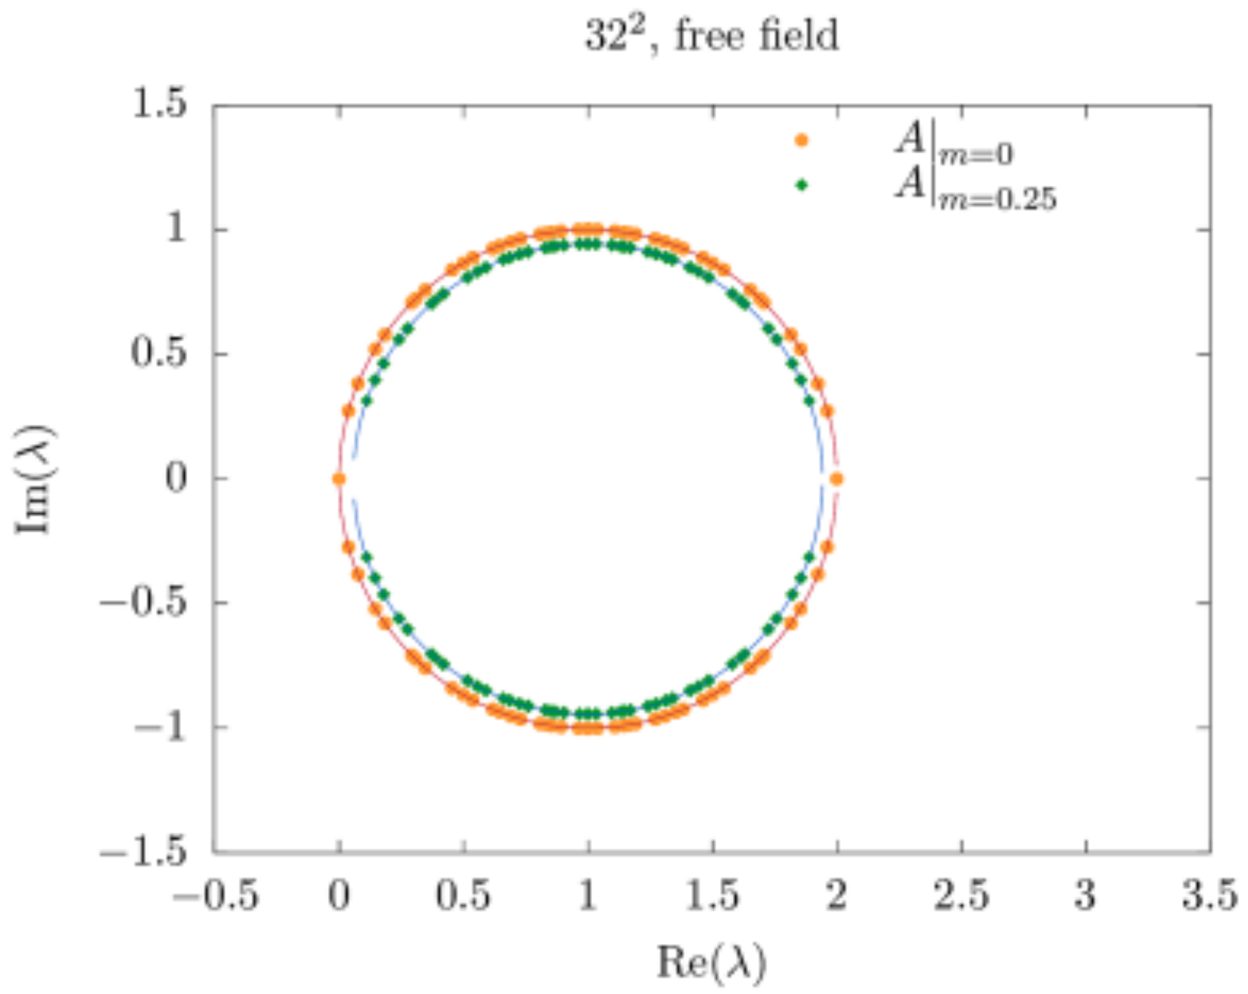

NVIDIA