Multigrid Solvers in Bridge++

Issaku Kanamori (RIKEN)

based on works with Bridge++ Project members

T. Aoyama, K. Kanaya, H. Matsufuru, Y. Namekawa, H. Nemura and K. Nitadori

and

W.-L. Chen, K.-I. Ishikawa,

Sep. 25, 2024

German-Japanese seminar 2024 at Mainz

RIKEN
Center for

Introduction

Main bottle neck of the LQCD: solver to solve Dirac equation
both in HMC and measurements

As the quark mass *am* becomes smaller, takes more and more time
⇒Multigrid solver has a very mild mass dependence

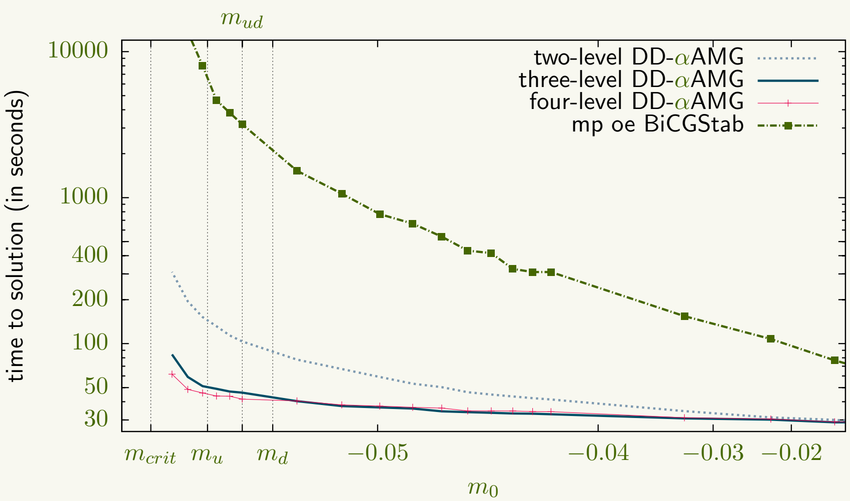

taken from M.Rottmann, Lattice 2015

64^4 lattice, 128 core

Outline

1. Multigrid Algorithm
2. Bridge++
3. Performance (Clover fermion)
4. Domainwall fermion
5. Summary and Outlook

Multigrid Algorithm

application to QCD: R. Babich et al. PRL 105 (2010) 201602
ML based: C. Lehner and T. Wettig, PRD 108 (2023) 034503
our implementation is based on DDαAMG: A. Frommer et al. , SIAM J. Sci. Comput. 36 (2014) A1581

Multigrid steps

- used as a preconditioner
- Restriction (R): $|x\rangle \to x(I,X) = \langle I(X)|x\rangle$
  fine grid (original lattice) $\to$ coarse grid
- Coarse grid solver: solve the coarse system
- Prolongation (P): $x(I,X) \to x(I,X)|I(X)\rangle$
  coarse grid $\to$ fine grid
- Smoother: improve the solution in the fine grid
- null space vector $|I\rangle$: adoptively improved
  $|I(X)\rangle = |I\rangle$ on domain X, otherwise 0

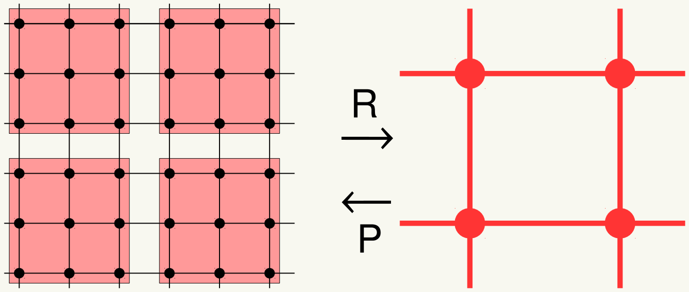

fine grid
12 dof/site

coarse grid
$2N_{\text{vec}}$ dof/site

Our implementation

I.K., K.-I. Ishikawa and H. Matsufuru, ICCSA2021 [2111.05012]

- 2-level multigrid, single precision
  - coarse solver: BiCGStab
  - post smoother: multiplicative Schwartz Alternating Procedure (SAP)
    inner: Minres
  - outer solver: Flexible BiCGStab
- setup: generate N_{vec} null space vectors
  initial SAP + 4 times adaptive MG preconditioner

Bridge++

first public version: 2009, the latest (v2.0.2): Feb. 2024

Present Project Members

Tatsumi Aoyama (U. of Tokyo),
Issaku Kanamori (RIKEN),
Kazuyuki Kanaya(U. of Tsukuba),
Hideo Matsufuru (KEK),
Yusuke Namekawa (Hiroshima U.),
Hidekatsu Nemura (Osaka U.),
Keigo Nitadori (RIKEN)

+ contributing to MG solver
  Ken-Ichi Ishikawa (Hiroshima U.),
  Wei-Lun Chen (SOKENDAI)

Features

- C++ object oriented framework
- Portable, easy to read, and extendable keeping reasonable performance
- Standard fermions, HMC, some measurements with test suite
- Extended to flexible data layout: version 2.0 (“alternative”)

Y.Akahoshi et al., J.Phys.Conf. Ser 2207 (2022)1, 012053

  - SIMD version for Intel AVX-512
  - GPU version with OpenACC (+ CUDA)
  - SIMD version for A64FX (Fugaku, QPACE4, etc.)

Bridge++ for Fugaku

- SIMD: 512 bits 8 doubles or 16 floats
  - `rrrr...iiii....` layout for complex numbers (not `riri...`)
  - 2-dim tiling in x-y directions

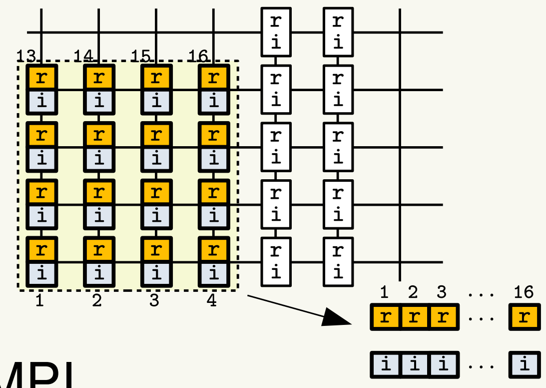

- Neighboring communication
  Persistent communication with Fujitsu extension of MPI
- can call QCD Wide SIMD library (QWS)

QWS `https://github.com/RIKEN-LQCD/qws` K.-I. Ishikawa et al., Comp. Phys. Comm. 282k (2023) 108510

- a solver library for Clover fermion
  mixed prec. nested BiCGStab with SAP preconditioner
- outcome of co-design for Fugaku
  102 PFlops with 147,456 nodes (i.e., almost full system)
- SIMD: `rrrr...iiii....` layout in 1-dim
- uses low level communication API (uTofu)

Target Machines (CPU/GPU)

	V100 (Volta)	H100 (Hopper)	A64FX
Number of SM	80	144	
FP64 Cores/GPU (CPU)	2560	9216	48
FP32 Cores/GPU (CPU)	5120	18432	48
Peak FP64 [TFlops]	7	26	3.07
w/ Tensor Cores		51	
Peak FP32 [TFlops]	14	51	6.14
w/ Tensor Cores		756	
Memory BW [TB/s]	0.90	2	1
PCIe [GB/s]	(Gen3) 32	(Gen5) 128	—
NVLinks [GB/s]	—	600	—
System	Cygnus 4× V100/node total 80 nodes	Pegasus 1× H100 /node total 120 nodes	Fugaku 1× A64FX/node total 158,976 nodes

GPU: PCIe edition, A64FX: Fugaku normal mode (2.0GHz)

Cygnus & Pegasus: Center for Computational Sciences, University of Tsukuba

Supercomputer “Fugaku”: RIKEN Center for Computational Science

A64FX: Fugaku

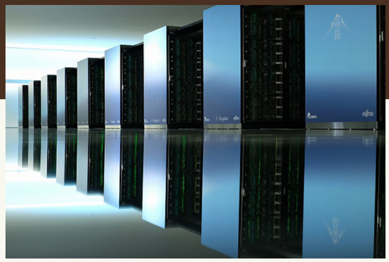

https://www.r-ccs.riken.jp/fugaku/about/

smoother: SAP from QWS library `https://github.com/RIKEN-LQCD/qws`

cf. Ishikawa, K.I., Matsufuru, PoS LATTICE2021 (2022) 278

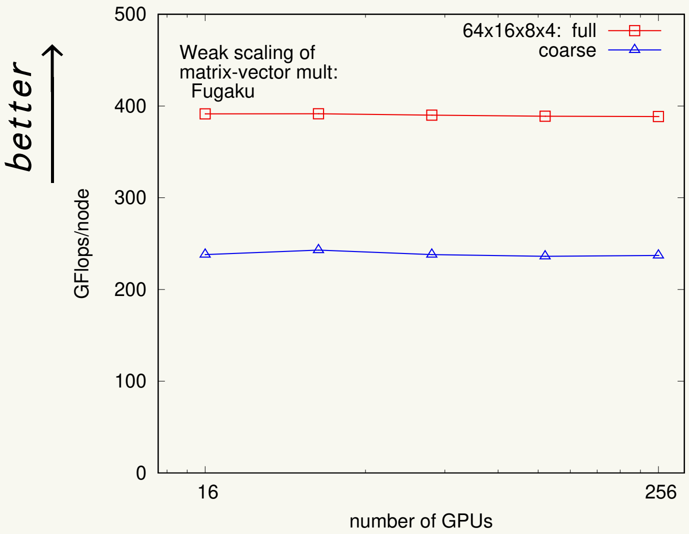

weak scalings of matrix vector multiplication

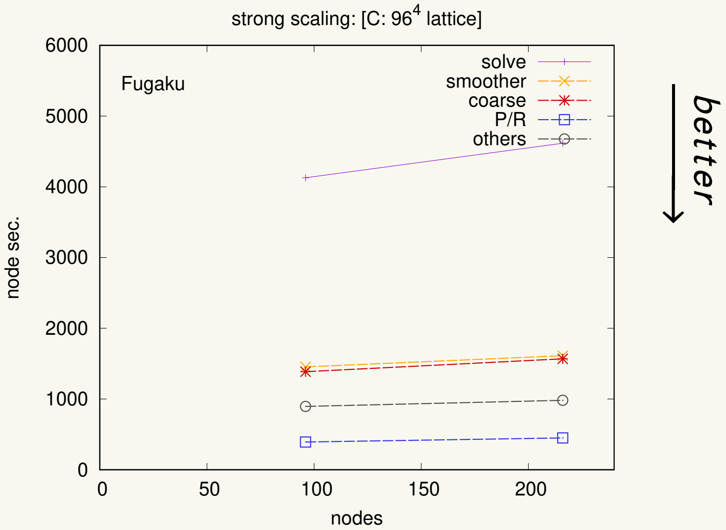

strong scaling of the MG solve configuration: 96^4 lattice, $M_\pi = 145$

K.-I.Ishikawa et al. [PACS] LATTICE2015, 075 (2016)

- solving time: 169 sec. (mixed prec. BiCGStab) $\Rightarrow$ 21.4 sec [216 nodes]
- good scaling
- QWS is really efficient: SAP $\sim$ +800 GFlops/node (i.e., $>$ 13%)

GPU Cluster: NVIDIA V100 (Cygnus)

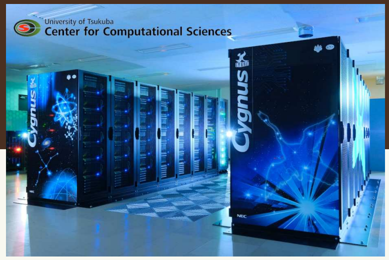

https://www.ccs.tsukuba.ac.jp/supercomputer/

some tunable parameters: not optimized

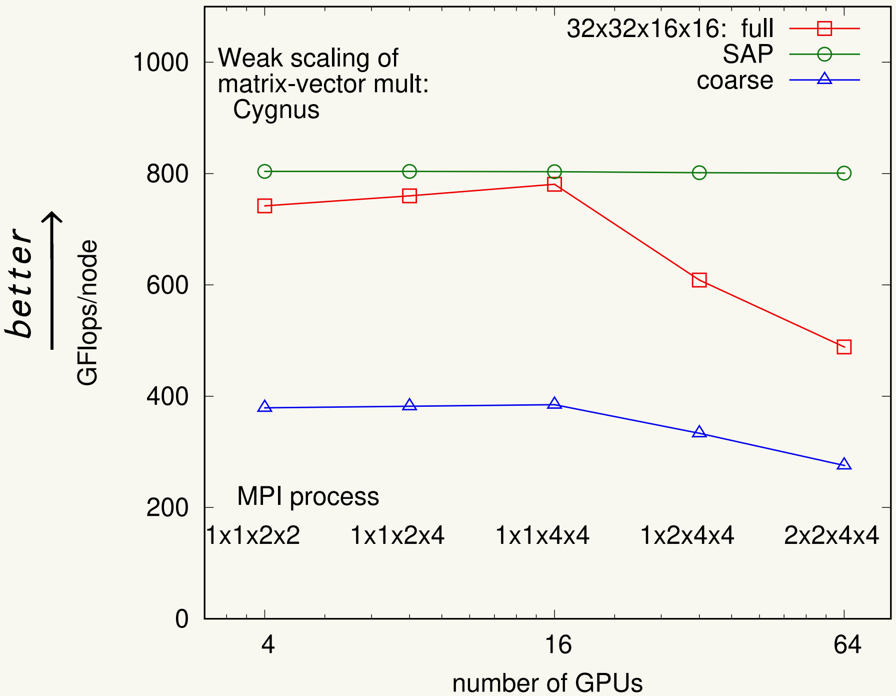

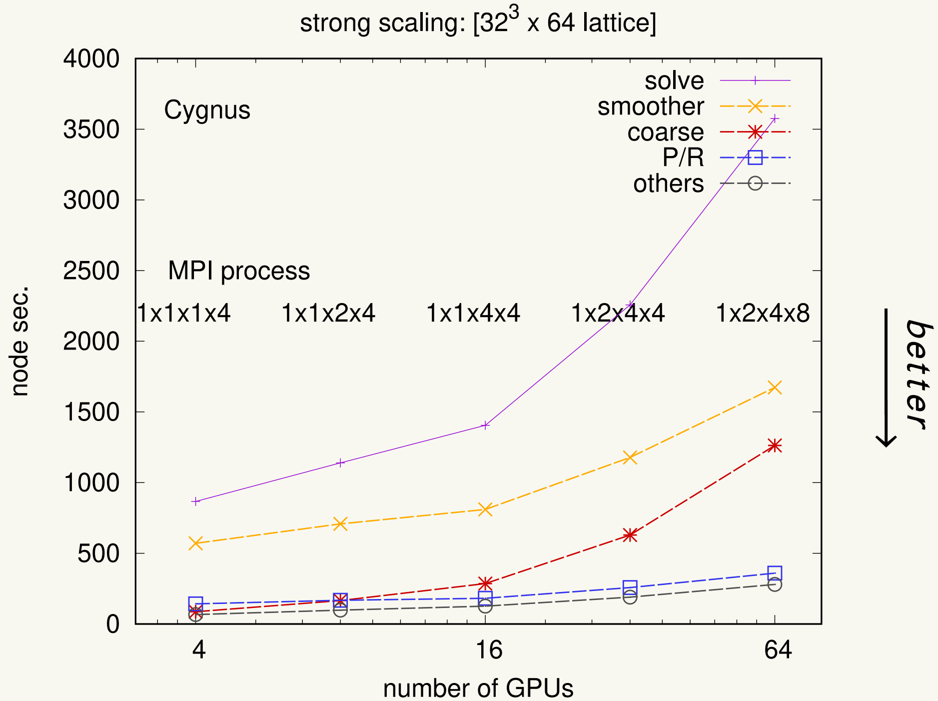

weak scalings of matrix vector multiplication [single prec.]
4 GPU/node

strong scaling of the MG solve
configuration:
$32^3 \times 64$ lat., $M_\pi = 156$ MeV
S.Aoki et al. [PACS-CS] Phys. Rev. D 79, 034503 (2009)

solving time: 393.3 sec. (mixed prec. BiCGStab) $\Rightarrow$ 216.8 sec. [4 GPU]

P/R: Prolongation and Restriction

GPU Cluster: NVIDIA H100 (Pegasus)

https://www.ccs.tsukuba.ac.jp/supercomputer/

some tunable parameters: not optimized

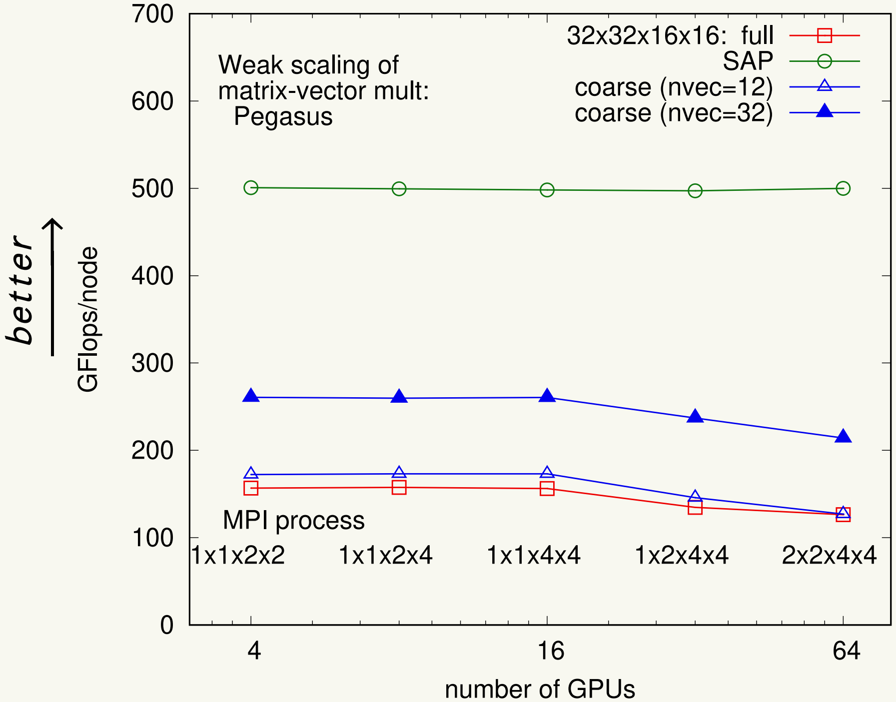

weak scalings of matrix vector multiplication [single prec.]

1 GPU/node

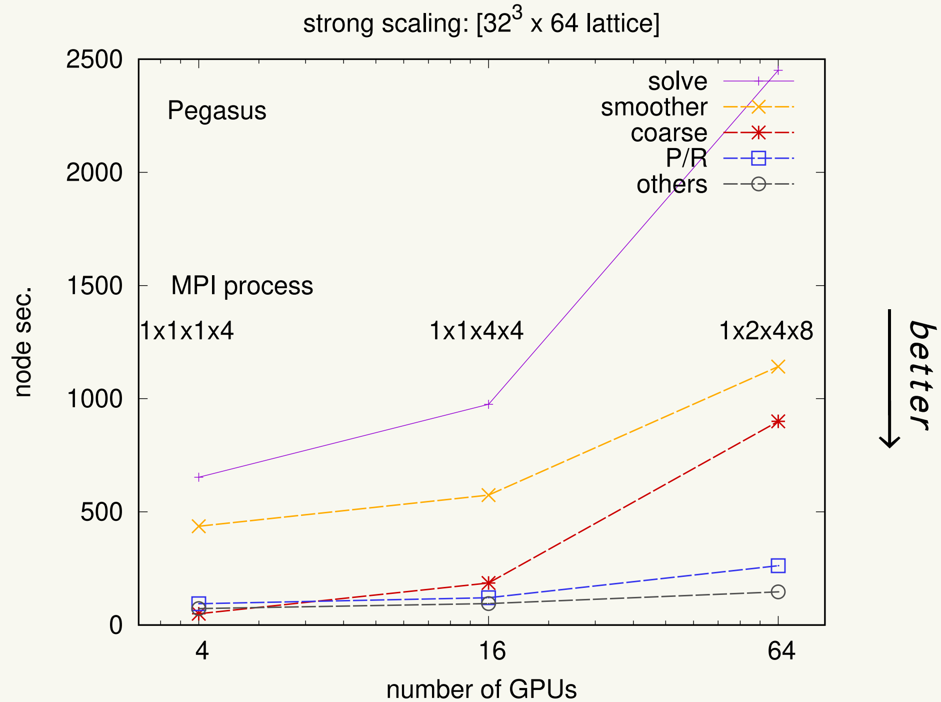

strong scaling of the MG solve

configuration:

$32^3 \times 64$ lat., $M_\pi = 156$ MeV

S.Aoki et al. [PACS-CS] Phys. Rev. D 79, 034503 (2009)

solving time: 484 sec. (mixed prec. BiCGStab) ⇒ 163 sec. [4 GPU]
solving time: 147 sec. (mixed prec. BiCGStab) ⇒ 38.3 sec. [64 GPU]

Some attempts but not drastic accelerations so far

S.D.Cohen et al., lattice 2011; P. Boyle, 1402.2085; R. Brower et al. PRD102 (2020) 094517; P. Boyle and A. Yamaguchi, 2103.05034; P. Boyle 2409.03904;

Real part of the spectrum is not positive definite

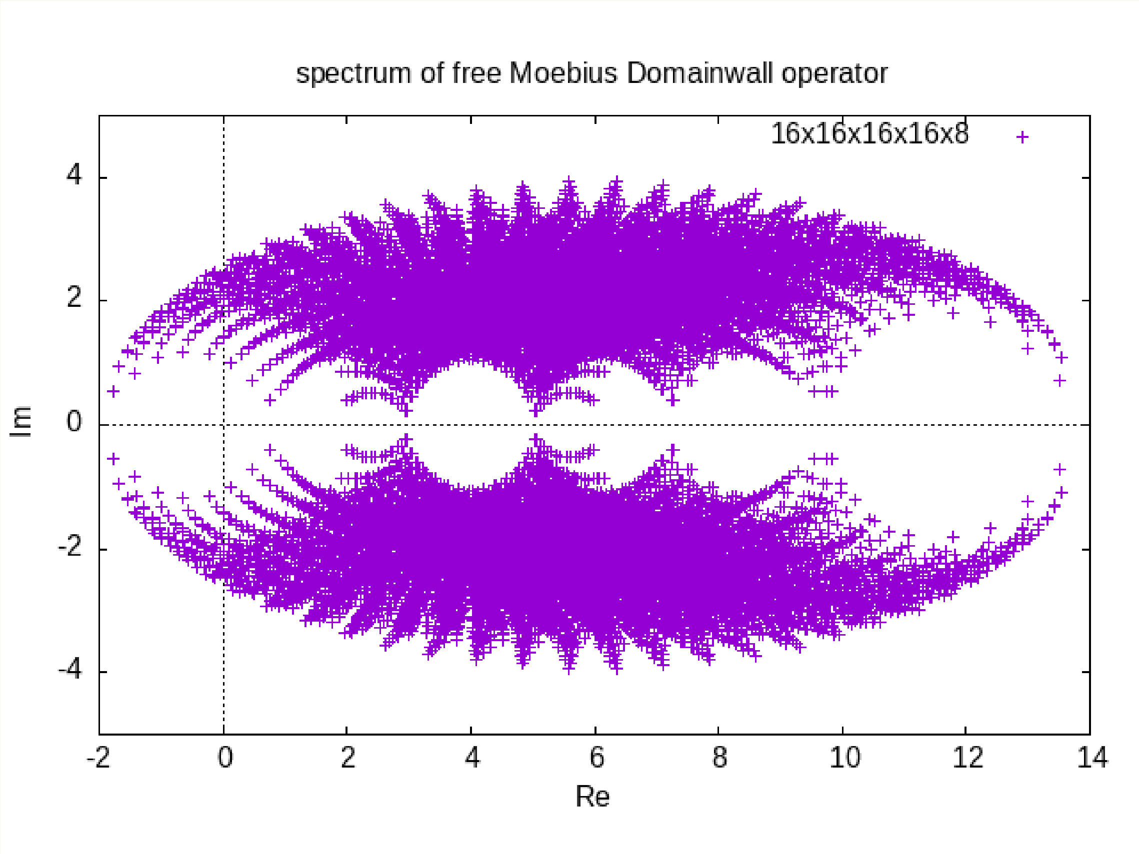

We usually use a CG solver with $D^\dagger D$

Towards Domainwall fermion

We know the physical modes are positive: $D_{\mathrm{PV}}^{-1}D$

$D_{\mathrm{PV}} = D(m = m_{\mathrm{PV}})$ with the Pauli-Villars mass m_{PV}

$\Rightarrow$ use D and $D^{\dagger}_{\mathrm{PV,coarse}}D_{\mathrm{coarse}}$ $D_{\mathrm{PV}}^{-1} \simeq D_{\mathrm{PV}}^{\dagger}$ as $1/(1 - ix) \simeq 1 + ix$ for small x

cf. R. Brower et al. PRD102 (2020)9, 094517

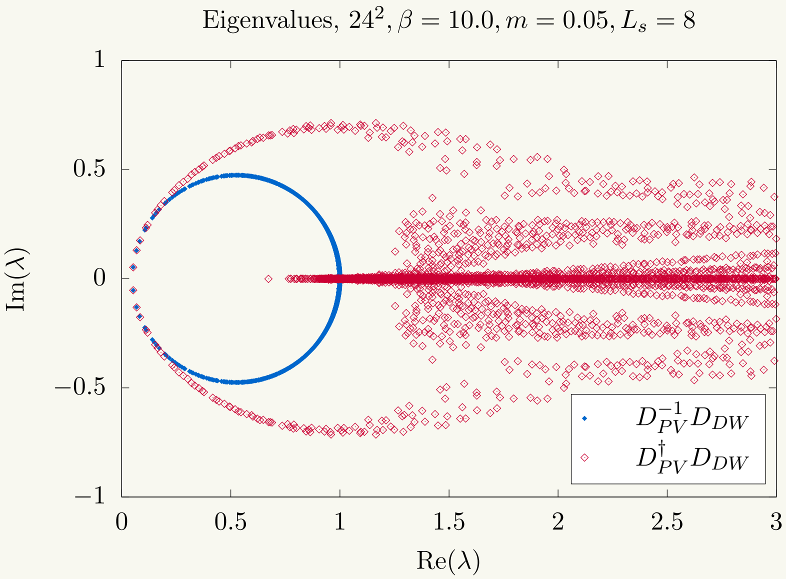

(from R.C.Brower et al.)

Towards Domainwall fermion

We know the physical modes are positive: $D_{\mathrm{PV}}^{-1}D$

$D_{\mathrm{PV}} = D(m = m_{\mathrm{PV}})$ with the Pauli-Villars mass m_{PV}

$\Rightarrow$ use D and $D_{\mathrm{PV,coarse}}^{\dagger} D_{\mathrm{coarse}}$ $\quad D_{\mathrm{PV}}^{-1} \simeq D_{\mathrm{PV}}^{\dagger}$ as $1/(1 - ix) \simeq 1 + ix$ for small x

cf. R. Brower et al. PRD102 (2020)9, 094517

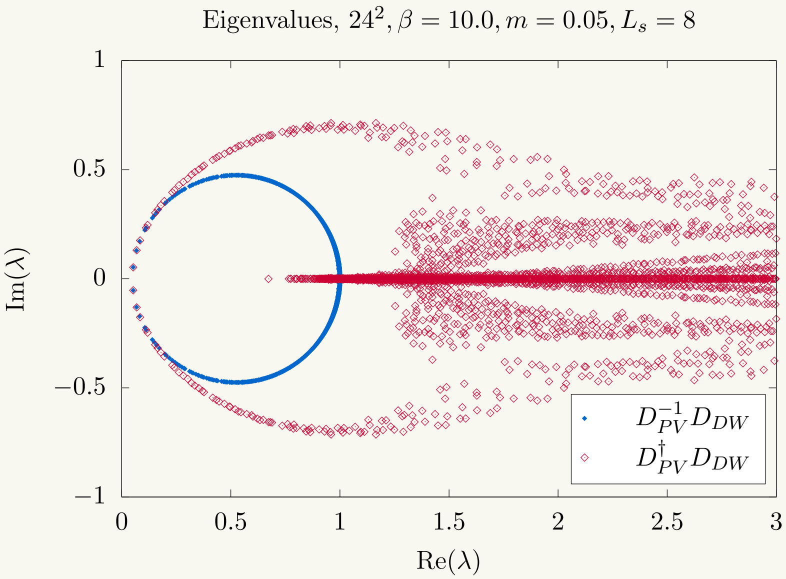

(from R.C.Brower et al.)

We tried a variant of this idea with $B_{\mathrm{PV}}^{\dagger} D$, but did not work

W.-L. Chen, I.K., H.Matsufuru, JPS meeting 2024.03

$$D = \begin{pmatrix} D_{EE} & D_{EO} \\ D_{OE} & D_{OO} \end{pmatrix} = \underbrace{\begin{pmatrix} D_{EE} & 0 \\ 0 & D_{OO} \end{pmatrix}}_{\equiv B} + \begin{pmatrix} 0 & D_{EO} \\ D_{OE} & 0 \end{pmatrix}$$

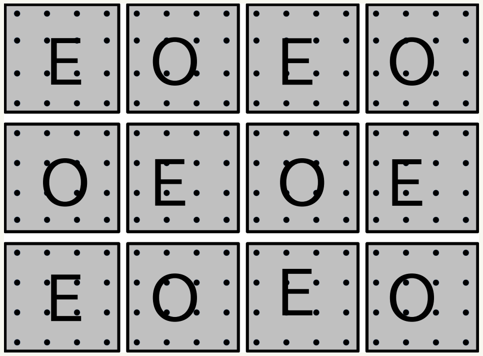

Spectrum?

configuration: JLQCD $48^3 \times 12$ 2+1 flavor finite-T conf. at physical point $T = 140$ MeV $< T_c$

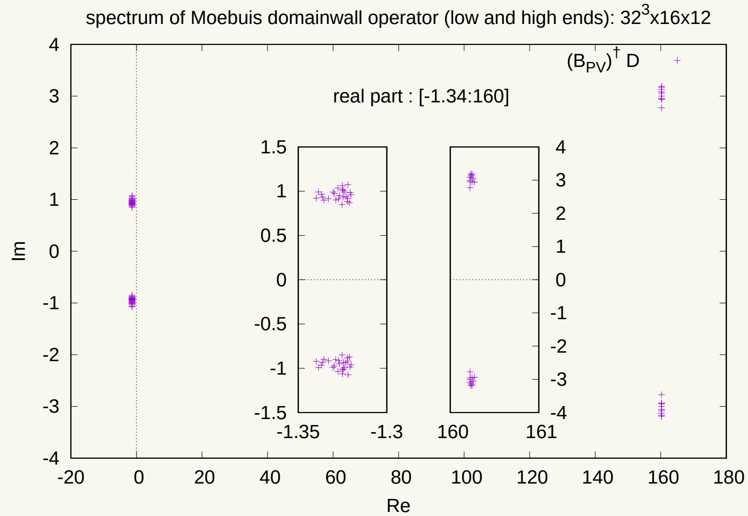

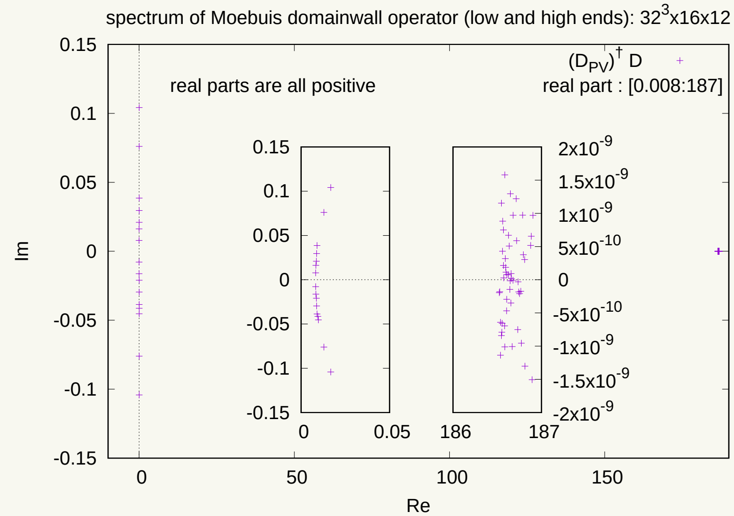

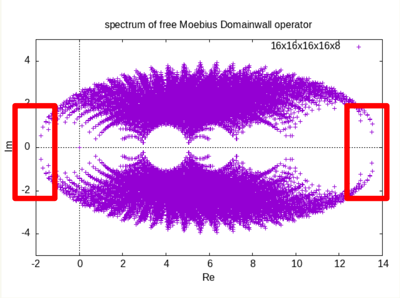

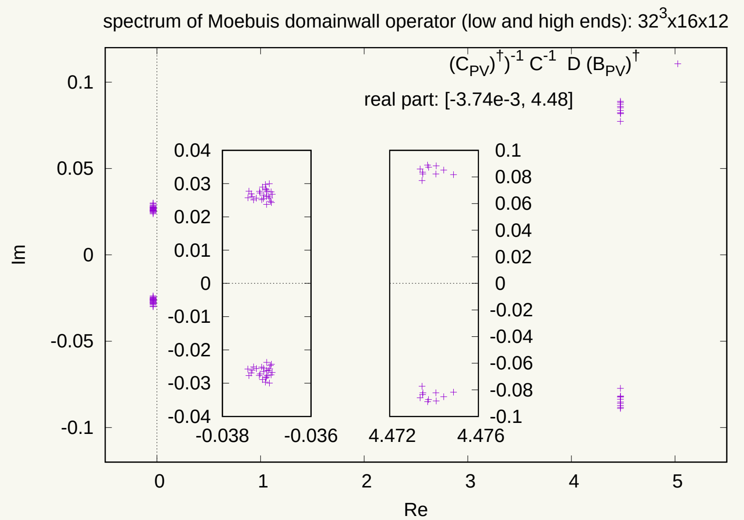

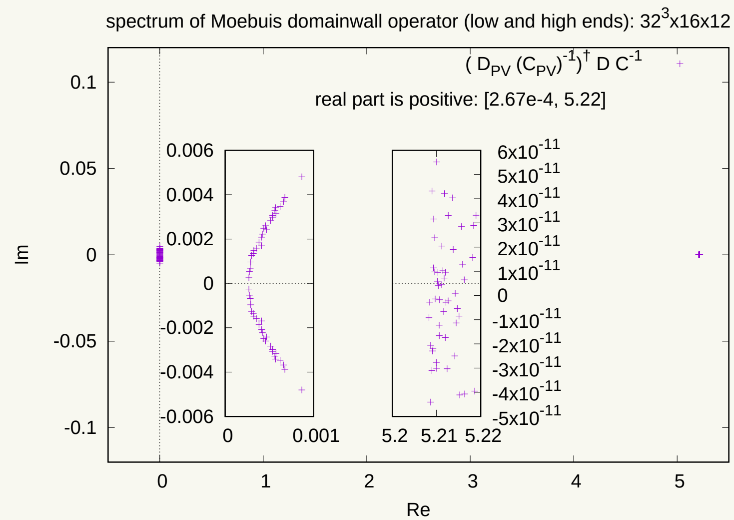

$$D = \begin{pmatrix} D_{ee} & D_{eo} \\ D_{oe} & D_{oo} \end{pmatrix} = \underbrace{\begin{pmatrix} D_{ee} & 0 \\ 0 & D_{oo} \end{pmatrix}}_{\equiv C} \begin{pmatrix} 1 & D_{ee}^{-1} D_{eo} \\ D_{oo}^{-1} D_{oe} & 1 \end{pmatrix} = \begin{pmatrix} 1 & D_{eo} D_{oo}^{-1} \\ D_{oe} D_{ee}^{-1} & 1 \end{pmatrix} C$$

spectrum with $B_{\rm PV}$ is not positive...

Summary and Outlook

- Bridge++ has an efficient implementation of multigrid solver for Wilson (Clover) fermion on Fugaku
- GPU version: implemented with OpenACC
- Domainwall fermion: on going

Outlook

- GPU: CUDA version
- Domainwall fermion
  - Investigation of the spectrum may help?
  - Block version may save us? P. Boyle

Acknowledgments

- Computational resource: Wisteria, Cygnus & Pegasus (through Multidisciplinary Cooperative Research Program in CCS, University of Tsukuba), Fugaku (RIKEN Center for Computational Science, ra000001)
- Grants: JSPS KAKENHI(20K03961, 19K03837)
- Configurations: Japan Lattice Data Grid, JLQCD collaboration