Large scale Domain Wall Fermion simulations on GPUs: Techniques and Properties

Chulwoo Jung, for RBC/UKQCD collaborations

David Murphy, Jiqun Tu, Robert Mawhinney, Christoph Lehner, Peter Boyle, ...

July 28, 2021

The RBC & UKQCD collaborations

UC Berkeley/LBNL
Aaron Meyer

BNL and BNL/RBRC
Yasumichi Aoki (KEK)
Peter Boyle (Edinburgh)
Taku Izubuchi
Yong-Chull Jang
Chulwoo Jung
Christopher Kelly
Meifeng Lin
Hiroshi Ohki
Shigemi Ohta (KEK)
Amarjit Soni

CERN
Andreas Jüttner (Southampton)

Columbia University
Norman Christ
Duo Guo
Yikai Huo
Yong-Chull Jang
Joseph Karpie
Bob Mawhinney
Ahmed Sheta
Bigeng Wang
Tianle Wang
Yidi Zhao

University of Connecticut
Tom Blum
Luchang Jin (RBRC)
Michael Riberdy
Masaaki Tomii

Edinburgh University
Matteo Di Carlo
Luigi Del Debbio
Felix Erben
Vera Gülpers
Tim Harris
Raoul Hodgson
Nelson Lachini
Michael Marshall
Fionn Ó hÓgáin
Antonin Portelli
James Richings
Azusa Yamaguchi
Andrew Z.N. Yong

KEK
Julien Frison

University of Liverpool
Nicolas Garron

Michigan State University
Dan Hoying

Milano Bicocca
Mattia Bruno

Peking University
Xu Feng

University of Regensburg
Davide Giusti
Christoph Lehner (BNL)

University of Siegen
Matthew Black
Oliver Witzel

University of Southampton
Nils Asmussen
Alessandro Barone
Jonathan Flynn
Ryan Hill
Rajnandini Mukherjee
Chris Sachrajda

University of Southern Denmark
Tobias Tsang

Stony Brook University
Jun-Sik Yoo
Sergey Syritsyn (RBRC)

RBC/UKQCD Domain Wall Fermion program

RBC/UKQCD has chosen DWF discretization of Dirac operators, which realizes 4d fermion as boundary modes of 5d fermion.

- Good chrial symmetry: Remnant symmetry breaking (residual mass) can be controlled separately from lattice spacing by increasing the extent of the 5th dimension (L_s) and the coupling between 4d slices ((z)Möbius, Möbius Accelerated DWF (MADWF)..)
- Protected zero mode: In contrast to Wilson fermions where the discretized Dirac operator can have poles near the valence mass (exceptional configurations), DWF formalism guarantees safety as long as valence mass is positive. Allowing simulation at physical point for moderate lattice spacing without the need for chrial extrapolation or link smearing for fermions, etc. $\rightarrow$ Focus on physical point. Avoid relying on ChPT.
- Focus on generating one longest possible Markov chain (vs. 'farming') : While evolving multiple chain can give practical advantages, this needs additional thermalization and potentially obscures autocorrelation, ergodicity issues. We are entering a regime where autocorrelation is becoming significant.

RBC/UKQCD 2+1f ensembles

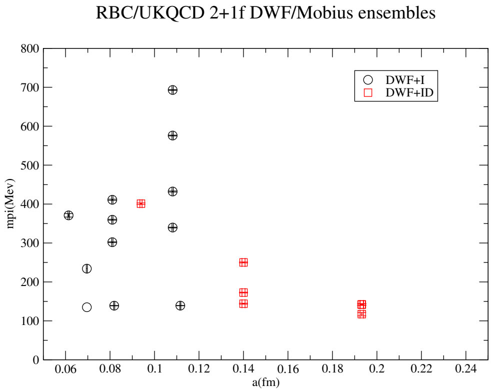

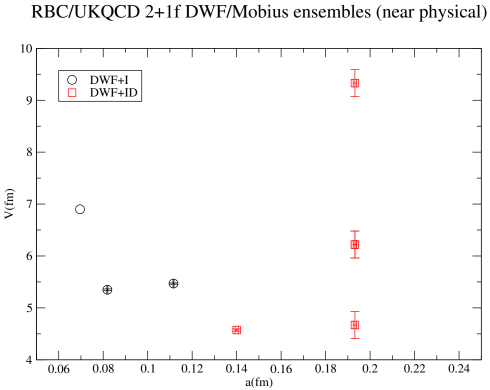

DWF+I: Iwasaki gauge action

DWF+ID: Iwasaki + Dislocation Suppressing Determinant Ratio(DSDR) : Supresses the chiral symmetry breaking on larger lattice spacing

96I on OLCF Summit

6 Nvidia V100 2 Power9 CPU, 512GB DDR4 + 96GB HBM2) × 4608 nodes
Mobius + Iwasak($\beta = 2.31$) i2+1 flavor physical , $96^3 \times 192 \times 12 (b + c = 2)$,
$am_l = 0.00054$, $am_h = 0.02144$, $a \sim$0.07fm, $L \sim$ 7fm
$16 \times 12^3 \times 12$ on $(1 \times 8 \times 8 \times 16 = 1024)$ ndoes × 6 GPUs
Evolution: CPS+ QUDA
Exact One Flavor Algorithm(EOFA) with Cayley Preconditioner: 1 flavor Mobius fermion action with

$$S_{pf} = \begin{pmatrix} 0 & \phi_1^\dagger \end{pmatrix} \left[I - k\Omega_-^T \frac{1}{H_T(m_1)} \Omega_- \right] \begin{pmatrix} 0 \\ \phi_1 \end{pmatrix} + \begin{pmatrix} \phi_2^\dagger & 0 \end{pmatrix} \left[I + k\Omega_+^T \frac{1}{H_T(m_2) - \Delta_+(m_1, m_2) P_+} \Omega_+ \right] \begin{pmatrix} \phi_2 \\ 0 \end{pmatrix},$$

Can be simulated with CG, avoids RHMC. Allow mixed precision, etc to reduce time on inversion.
Small memory footprint improve overall arithmetic intensity, especially significant on GPUs.
Multisplitting-preconditiond Conjugate Gradient(MSPCG, J. Tu, Wed 14:30 EDT)
Additive Schwarz 'done right' for Möbius CG Utilize Tensor core for 5D part
Low Overhead Transparent Multilevel Checkpoint/Restart (VeloC)

VeloC: Mid-MD Checkpointing capability

Motivation: Traditional checkpointing unit is 1 Hybrid Monte Carlo trajectory. After each trajectorory the gauge configuration (U_μ) and RNG state are stored.

- The job size and duration is increasing even with the rapid increase in computing power per node.
- Recent advances in evolution algorithm may help in reducing the autocorrelation time in MD units, but often needs more computation per trajectory.
- Longer runtime makes it more vulnerable against hardware errors
- Also useful in running small, but long jobs which often is problematic on LCFs

A portable, efficient and robust checkpointing with VeloC would allow for more resiliency against queue policy and/or machine failure.

Checkpointing strategy & Integration Stauts

VeloC interface implemented in CPS

- Registers necessary data (Gauge field (U_μ), gauge momenta(H_μ), Pseudofermions ($\psi(x)$)) to VeloC when they are created.
- Checkpoint with VeloC at the outermost MD integrator
- Recently changed to C++ interface. RNG checkpointing with serialization library in progress.
- Native (without VeloC) implementation also available in CPS. Currently being used for small volume RMHMC runs

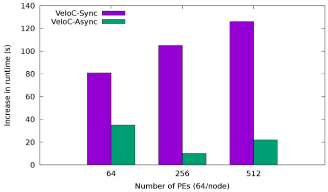

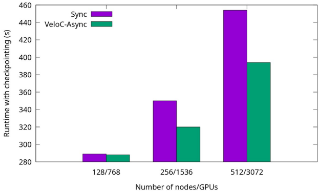

From B. Nicolae

96I evolution on Summit

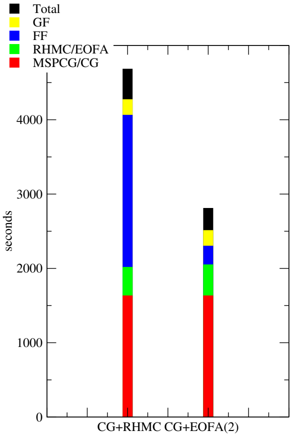

1024 nodes
CG+RHMC/CG+EOFA(2)

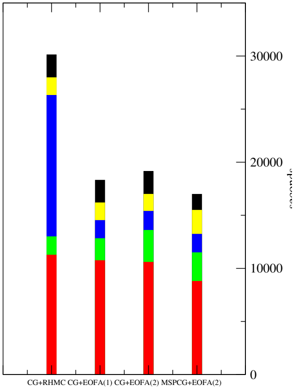

128 nodes
CG+RHMC/CG+EOFA(1)/
CG+EOFA(2)/MSPCG+EOFA(2)

While Multimass solver does not take much time, fermion force term calculation becomes time consuming. Arithmetic intensity low (no smearing).
EOFA with the Cayley preconditioner allows effective use of mixed precision solvers, efficient mass preconditioning, and reduce the number of pseudofermions significantly.

Measurement details

- Evec generation: Fine Grid Base (150 evecs, QUDA/CPS) + Coarse grid (5000 evecs, Grid/GPT)
  Coarse Grid with 4^4 block size achieves factor of $\sim$ 50 compresson
  Reduction in memory footprint also crucial: Multiple RHS projection in the blocked base, avoids explicit decompression
- Solver: QUDA MSPCG or Grid CG
  Half precision solver works well with single precision evecs with defect correction.
  Deflation reduce iteration count for light exact (1e-8) inversion by a factor of 20 $\sim$ 25.
- Wall source propagators for basic measurements ($m_\pi, m_K \cdots$): (48 sloppy (10^{-4}) + 12 exact (10^{-8})) propagators $\times$ 11 configurations so far
- Z3 Box propagators (for Ω): ($24^3, 32^3, 48^3$) Coulomb gauge-fixd box source with Z3 noise $\times$ 48 timeslices $\times \sim$ (30+30) configurations (separated by 20) Lattice rotation (ZT so far) employed to increase statstics. Planning to do XT, YT too

Residual mass, Z_A on 96I ensemble

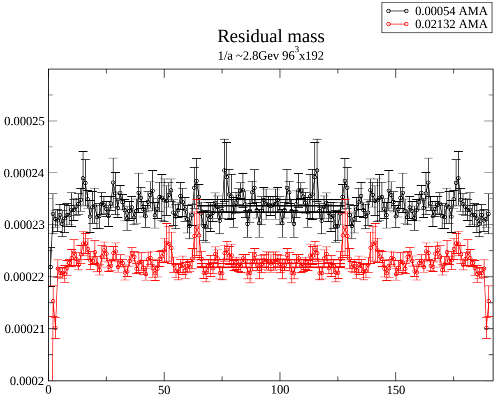

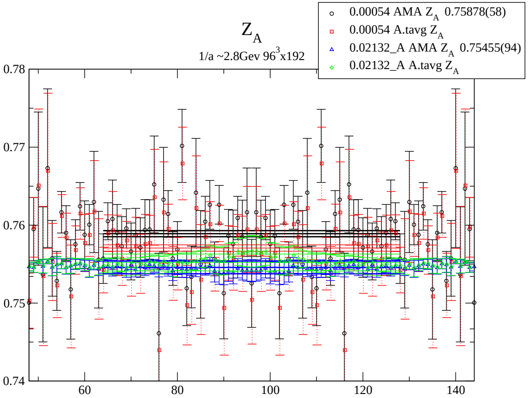

$$m_{res}a \sim 2 \times 10^{-4} \sim 0.6\text{Mev}$$

Omega mass on 96I ensemble

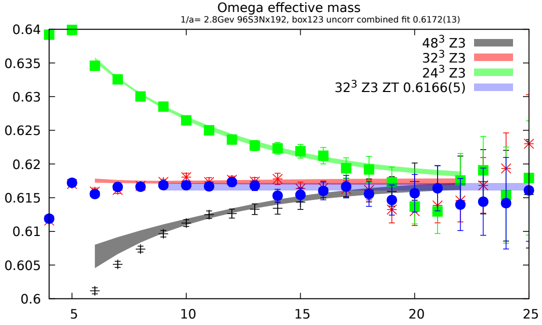

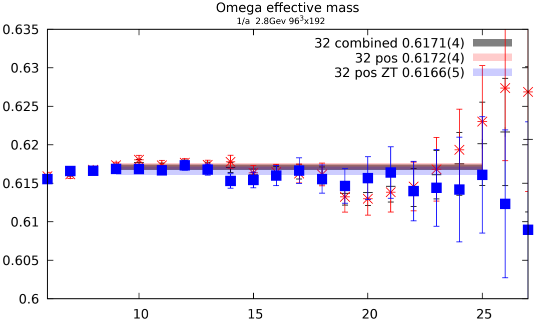

$$am_{\Omega} \sim 0.6171\ (4)$$

Wilson flow scale on 96I ensemble

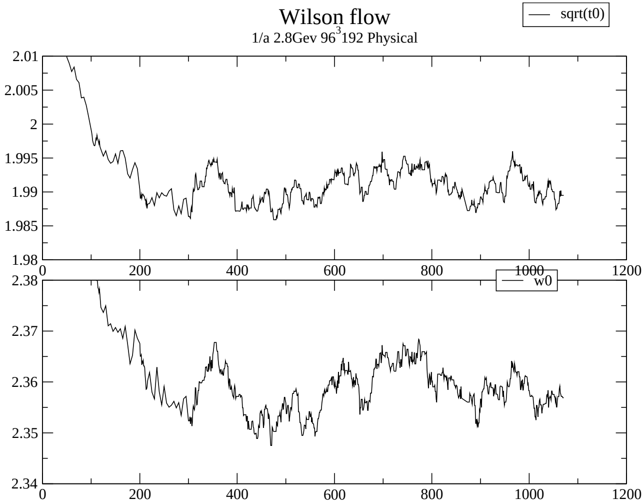

$\sqrt{t_0} = 1.9897(7)$
$w_0 = 2.356(2)$
$\tau_{int}(t_0, w_0) \sim 40\text{MD}$

Global topology on 96I ensemble

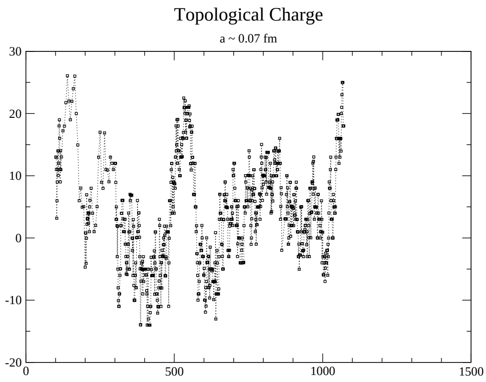

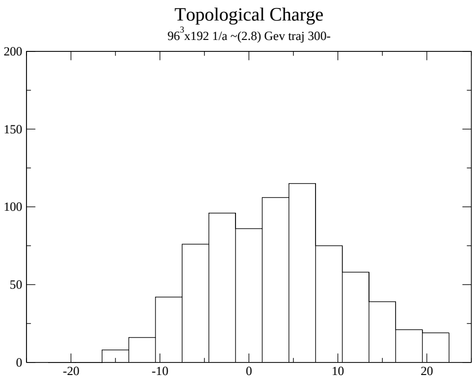

$$\tau_{int} \sim 50\text{MD}$$

Z_V on 96I ensemble

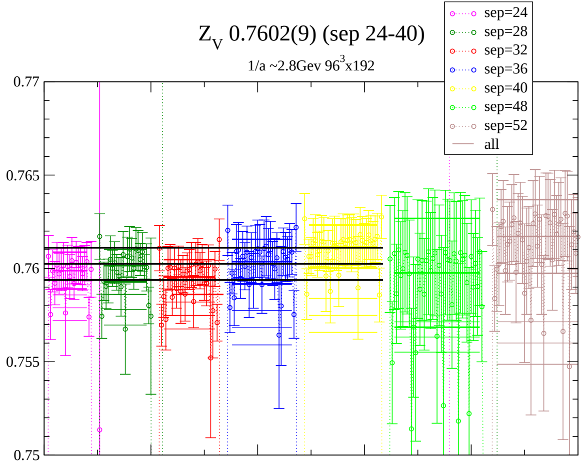

$$Z_V = \frac{<0|\sum_{\vec{x},t} \mathcal{V}_i^a(\vec{x},t) V_i^a(0.0)|0>}{<0|\sum_{\vec{x},t} V_i^a(\vec{x},t) V_i^a(0.0)|0>}$$

$$Z_V = 0.7602(9)\ \text{(preliminary)}$$

m_π, f_π on 96I ensemble

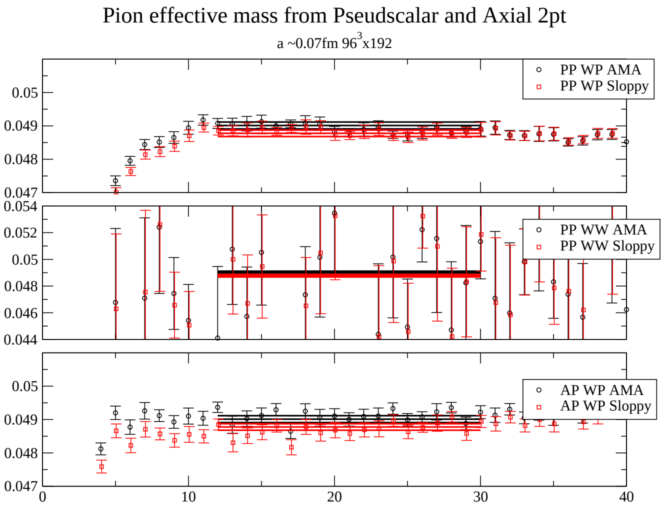

$$f_\pi = Z_V \sqrt{\frac{2}{m_\pi V} \frac{\mathcal{N}_{AP}^{WP2}}{\mathcal{N}_{PP}^{WW}}}$$

$$am_\pi = 0.04901(15),$$

$$af_\pi = 0.04817(14)$$

(Preliminary, stat. only)

m_K, f_K on 96I ensemble

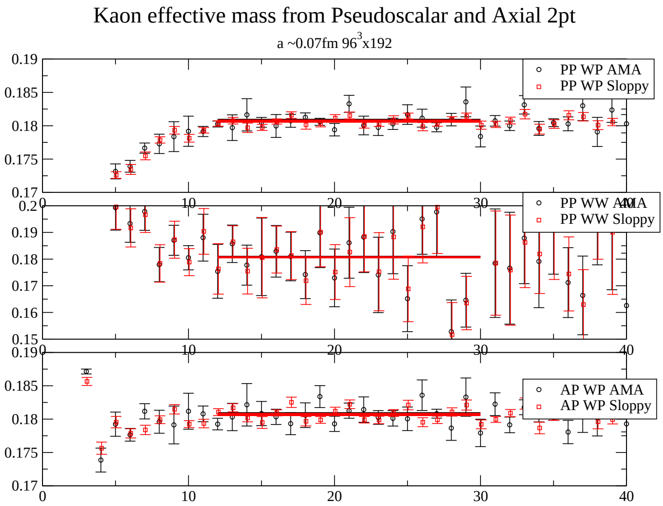

$am_K = 0.1808(1), af_K = 0.0575(1)$
(Preliminary, stat. only)

Comparison of other DWF+I physical point ensembles

	48I	64I	96I(Preliminary, stat only)
am_π	0.08049(13)	0.05903(13)	0.04901(15)
af_π	0.07580(8)	0.05550(10)	0.04817(14)
am_K	0.28853(14)	0.21531(17)	0.1808(1)
af_K	0.09040(9)	0.06653(10)	0.0575(1)
am_Ω	0.9702(10)	0.7181(7)	0.6171(4)
Z_V	0.71076(25)	0.74293(14)	0.7602(9)
$\sqrt{t_0}/a$	1.29659(28)	1.74496(62)	1.9897(7)
w_0/a	1.50125(94)	2.0495(15)	2.356(2)

Discussion & future plans

- $96^3 \times 192$ DWF+Iwasaki ensemble will provide RBC/UKQCD with a physical mass ensemble at the 3rd lattice spacing.
- AMA with efficiently generated corase-grid eigenvectors makes it possible to significantly reduce error despite relative small number of measurements. More eigenvectors already available for measurements. Errors already comparable with older ensembles.
- DWF ensemble generation on new and upcoming machines are likely to continue to be limited by the memory and internode bandwidth. While EOFA and MSPCG has helped mitigating these issues, evolution is more vulnerable compared to measurements, where various techniques (exact deflation, AMA, A2A, Split Grid...) are already devloped to mitigate, if not overcome, bandwidth issues. Continuing effort on optimizing for Exascale hardwares (DD-HMC (Boyle, Thur 22:45 EDT),..)
- Current lattice spacing was chosen to achieve reasonable topology tunneling. Going further requires a significant increase in the computing resource required to generate similar number of independent configurations $(cost) \sim a^{-10}$ (Critical slowing down).

While different observables may have different τ_{int}, it would be much desirable to ensure the ensemble is much longer than the (known) autocorrelation time. RBC/UKQCD is pursuing various approaches, many within the framework of US Exascale project (ECP)

Fourier Acceleration based approaches:

- Riemann Manifold Hybrid Monte Carlo (RMHMC): T. Nguyen, Thur. 14:15 EDT)
- Gauge fixed HMC(GFHMC): Include Gauge fixing term in the evolution (A. Sheta, Thur. 14:30 EDT)

ML based approaches

- LeapFrogLayers(L2HMC, S. Foreman, Thur. 14:45 EDT)
- Field Transformation HMC (L. Jin, Poster E3, X. Jin, Thur 13:45 EDT)
- Normalizing Flows (Foreman, Poster F4)

Thank you!