The Renormalization-Group Preconditioned Conjugate Gradient for Domain Wall Fermions

Jonah Eick

In collaboration with Robert Mawhinney

RBC & UKQCD Collaboration

43rd Symposium on Lattice Field Theory
University of Maryland, College Park

July 27, 2026

Columbia University in the City of New York

Introduction

Motivation

- Follow up to my talk at Lattice 2025 in Mumbai...
- The Dirac operator normal equations, $D^\dagger D\,\psi = D^\dagger\,\eta$, are solved with Krylov methods like the conjugate gradient (CG).
- Long setup times make multigrid and other deflation methods ineffective for HMC ensemble generation, where there is there is one light-quark solve per configuration (per time-step).
- The RBC/UKQCD collaboration needs up to 40,000 CG iterations per DWF solve for light quarks on large lattices... Expensive!
- The dream is for a "fast-set-up-time multigrid."

Renormalization Group Blocking

Initially, two HMC-generated ensembles ($N_f = 2+1$, ID + MDWF) $\mathcal{E}_f$ and $\mathcal{E}_c$

"Fine" ($\mathcal{E}_f$): $S_f[U_f]$, $24^3 \times 64 \times 12$, $\mathbf{a^{-1}} \approx \mathbf{2}$ GeV, $m_\pi = 300(3)$ MeV

"Coarse" ($\mathcal{E}_c$): $S_c[U_c]$ $12^3 \times 32 \times 12$, $\mathbf{a^{-1}} \approx \mathbf{1}$ GeV, $m_\pi = 307(5)$ MeV

Then produce a coarse ($a^{-1} \approx 1$ GeV) *RG-blocked* ensemble $\mathcal{E}_c^b$ with effective action $S_c^b[U_c]$:

$$e^{-S_b^c[U_c]} \propto \int [dU_f]\, e^{-S_f[U_f]}\, G[U_c, U_f]$$

Using a blocking kernel $\quad G[U_c, U_f] = \Pi_{x,\mu}\delta(U_c(x,\mu) - g_b[U_f; x, \mu])$

APE Blocking:

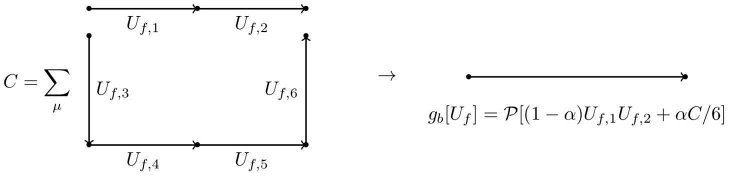

$\rightarrow$ This is a cheap procedure!

RG Correspondence

The low energy physics of S_f and S_c^b strongly correspond

Observables on HMC-generated $\mathcal{E}_f$ and $\mathcal{E}_c$ agree within $2-7\%$. Observables on RG-blocked $\mathcal{E}_c^b$ agree to 1% with $\mathcal{E}_c$.

$\mathcal{E}_f$ (HMC) $\downarrow$ $\mathcal{E}_c$ (HMC) $\downarrow$ $\mathcal{E}_c^b$ (RG-Blocked) $\downarrow$

	$\langle O\rangle_f$	$\langle O\rangle_c$	$\langle O\rangle_c^b$
size	$24^3\times64\times12$	$12^3\times64\times12$	$12^3\times64\times12$
β	1.943	1.633	-
am_l	0.000787	0.008521	0.007494
am_h	0.019896	0.065073	0.064150
a^{-1}(GeV)	2.001(18)	1.015(16)	1.010(16)
am_{res}	0.004522(12)	0.007439(86)	0.00847(21)
m_π(MeV)	300(3)	307(5)	308(8)
m_K(MeV)	491(5)	506(8)	507(11)
m_Ω(MeV)	1557(71)	1652(27)	1685(52)
f_π(MeV)	138(2)	147(2)	151(3)
f_K(MeV)	155(2)	166(3)	169(4)

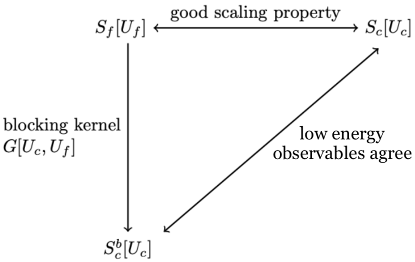

Operator Correspondence

Mapping Between Grids

Interpolation $\quad (I : \Psi_c \to \Psi_f)$

$$I_{\text{cov}} = \prod_{\mu=0}^{3} \left(1 + T_\mu^{\text{cov}}\right) P_{\text{inj}}, \quad T_\mu^{\text{cov}} \psi(x) = \frac{1}{2}\Big[U_\mu(x)\psi(x+\hat{\mu}) + U_\mu^\dagger(x-\hat{\mu})\psi(x-\hat{\mu})\Big].$$

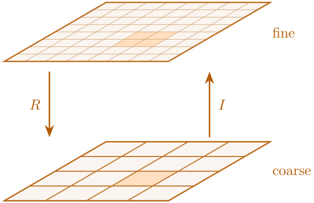

Restriction $\quad (R : \Psi_f \to \Psi_c)$

$$R_{\text{cov}} = I_{\text{cov}}^\dagger$$

Gauge fixing: we transform lattices to (*near*) Landau gauge.

- Approx. smoothness allows us to drop parallel transport in T_μ: $U_\mu \to 1$.
- This increases accuracy of fine–coarse mapping.

Low Mode Subspaces

For a successful preconditioner, we want $\boxed{I\,(D^\dagger D)_c^{-1}\,R \approx (D^\dagger D)_f^{-1}}$

$$(D^\dagger D)_f = \sum_i \lambda_{f,i}\,|\psi_{f,i}\rangle\,\langle\psi_{f,i}| \quad \text{and} \quad (D^\dagger D)_c = \sum_i \lambda_{c,i}\,|\psi_{c,i}\rangle\,\langle\psi_{c,i}|$$

Matrix elements: $\langle\psi_{f,m}|\,I(D^\dagger D)_c^{-1}\,R\,(D^\dagger D)_f\,|\psi_{f,n}\rangle$

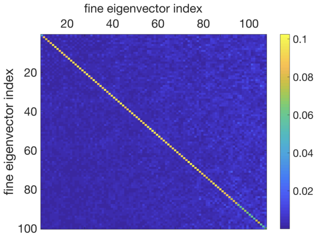

Diagonal structure $\implies I(D^\dagger D)_c^{-1}\,R\,(D^\dagger D)_f \approx 1$

Overlap: $|\langle I\psi_{c,i}, \psi_{f,i}\rangle|$

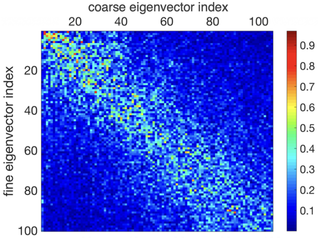

$I|\psi_{c,i}\rangle$ and $|\psi_{f,i}\rangle$ have large overlap.
Fine low modes have support on ∼20 coarse low modes.

$\implies$ $\mathbf{(D^\dagger D)_c}$ **is a good approximation of** $\mathbf{(D^\dagger D)_f}$ **in low-mode space!**

RG Preconditioner

Preconditioner: First Pass

The preconditioner takes the form

$$M^{-1} = \mathbb{I}_f + I\left(\sum_{i=1}^{N} \omega_i \frac{|\psi_{c,i}\rangle\langle\psi_{c,i}|}{\lambda_{c,i}}\right) R.$$

→ First term: identity $\mathbb{I}_f$ needed to make M^{-1} surjective

→ Second term: applied to residual r_f gives coarse approximation of the error,

$$\tilde{e}_f = I(D^\dagger D)_c^{-1} R r_f \approx (D^\dagger D)_f^{-1} r_f = e_f$$

→ ω_i is a tunable parameter: rescales per-mode

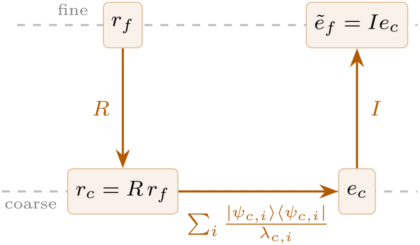

restrict → solve on coarse → interpolate

First Pass Result

RGPCG with M^{-1} using $N = 500$ coarse evecs and $\omega_i = 0.02$ for all i:

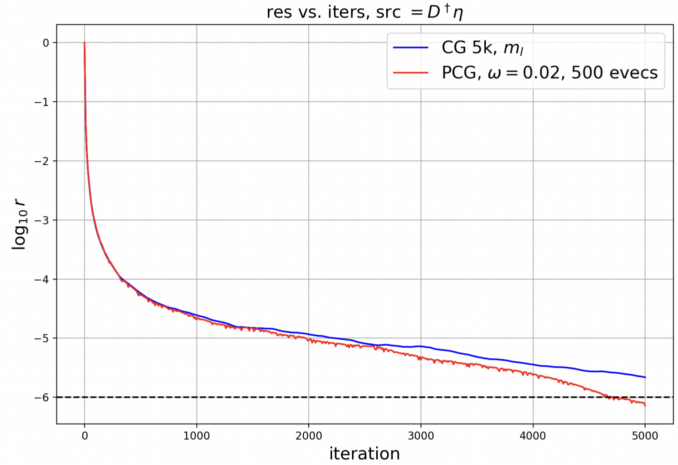

→ Only very small improvement

Noise and Filtering

The problem is *high-mode pollution* introduced in interpolation

Ex   True error after 400 iters: $e_* = x_* - x_{400}$ $(e_*, \tilde{e}_f) \approx 0.65$ $(e_*, F_{\text{cheby}}\tilde{e}_f) \approx 0.95$

Significant improvement after filtering with exact fine eigenvectors, $F_* = \sum_i |\psi_{f,i}\rangle\langle\psi_{f,i}|$, or a Chebyshev polynomial filter.

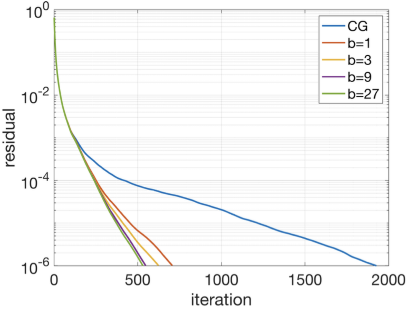

(RGPCG with exact filter F_*)

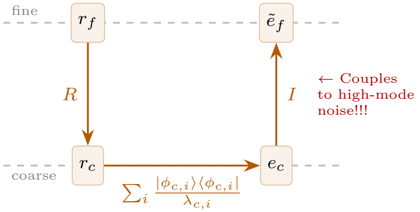

- These filtering methods are far too expensive! Need another approach ...

4D Chiral Improvement

5D and Spin Structure

We propose projecting the preconditioner onto 4D walls and physical LH/RH spin components

1. RG blocking is based on physical observables
2. Exponential localization of near-zero modes: $|\psi_{nz}(x)| \sim e^{-\alpha s}$
3. m_{res} is bigger on $\mathcal{E}_c^b$: D_c^{DW} couples to bulk and mixes chiralities more

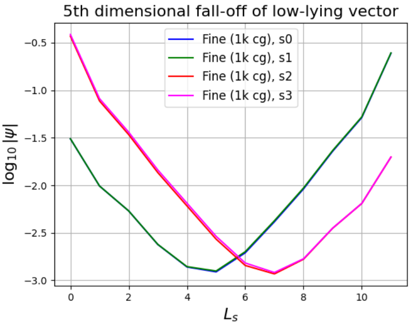

Improved Preconditioner

$$M_{4D}^{-1} = \mathbb{I}_f + \mathcal{P}_{4\to5}\, I_{4D} \left(\sum_{i=1}^{N} \omega_i \frac{|\psi_{c,i}^{4D}\rangle\langle\psi_{c,i}^{4D}|}{\lambda_{c,i}} \right) R_{4D}\, \mathcal{P}_{5\to4}$$

Where

$\mathcal{P}_{5\to4} = \frac{1}{2}(1-\gamma_5)\delta_{s,0} + \frac{1}{2}(1+\gamma_5)\delta_{s,L_s-1}$, $\mathcal{P}_{4\to5}$ injects into 5d vector

- Physically motivated
- Reduces operation cost by $1/L_s$
- Reduces coupling to high-mode pollution, while preserving most of the low-mode information

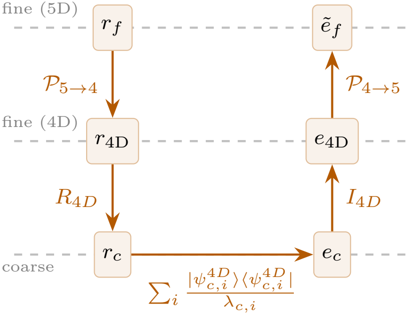

Tuning

1. Number of coarse low modes **(right)**: $N = 100$ is optimal
2. Tuning m by $\pm 10\%$ has negligible effect
3. Incorporating more s-slices (near walls): less improvement
4. ω_i **(right)**: hand-tuned per 5-10 modes *(red)* or computed from Rayleigh quotients *(blue)*:

$$\omega_{RQ,i} \propto |I\psi_{c,i}^{4D}|^2 / \langle I\psi_{c,i}^{4D}, (D^\dagger D)_f I\psi_{c,i}^{4D} \rangle$$

Note $\frac{\lambda_{20,f}}{\lambda_{1,f}} \sim 3$, so deflating first 10-20 modes can improve κ significantly

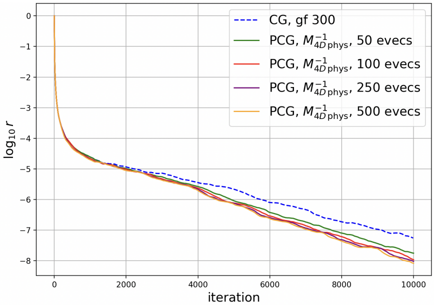

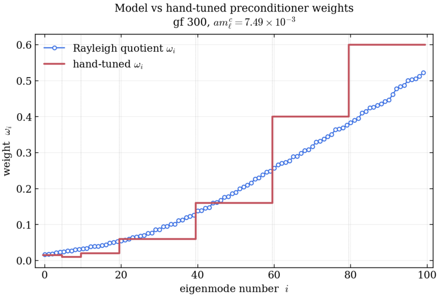

Current Results and Performance

Tuned RGPCG

RGPCG with M_{4D}^{-1}, 100 evecs, same ω_{RQ} – very similar on different configs

$\Longrightarrow$ No additional per-configuration tuning needed

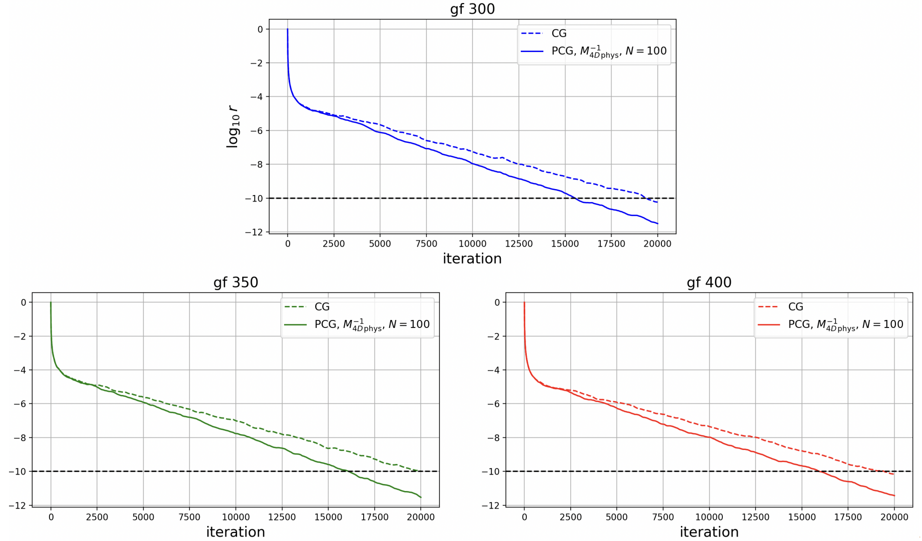

Performance I

Condition number:

$$\kappa = \frac{\lambda_{max}}{\lambda_{min}}$$

For CG, $\sqrt{\kappa}$ controls worst-case convergence (“iters to ϵ” $\lesssim \sqrt{\kappa}\,\ln(2/\epsilon)$)

Preconditioner	$\sqrt{\kappa}$	iters to $\lvert r\rvert \approx 10^{-10}$	Cost (flops/iter)
CG (no Prec.)	≈ 2947.1	≈ 19400	$\approx 3.4 \times 10^{10}$
M_{5D}^{-1}	≈ 2218.7	≈ 16550	$\approx 5.4 \times 10^{10}$
M_{4D}^{-1}	≈ 2123.1	≈ 15500	$\approx 3.9 \times 10^{10}$

Table: Convergence Analysis of PCG Strategies, $N = 100$

- $\frac{\sqrt{\kappa}_{4D}}{\sqrt{\kappa}_{CG}} \sim 0.72, \quad \frac{\text{Cost}_{4D}}{\text{Cost}_{CG}} \sim 1.15 \implies$ net improvement
- RGPCG reaches 10^{-10} precision in residual in 20% fewer iterations, 15% higher cost
- Slight overall reduction in convergence time

Performance II

At strange quark mass:

1. CG convergences much faster
2. PCG improvement is similar:

$$\frac{\sqrt{\kappa_{4D}}}{\sqrt{\kappa_{CG}}} \sim 0.80$$

(Compared to $\frac{\sqrt{\kappa_{4D}}}{\sqrt{\kappa_{CG}}} \sim 0.72$ for light quark, 5000+ iterations to 10^{-6} prec.)

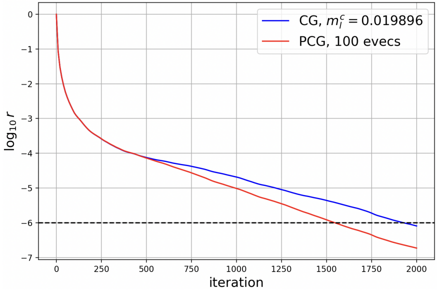

Plans to work with physical light quarks... (we hope for greater improvement due to bigger κ and more iters to ϵ)

Setup Cost

1. RG-blocking: C-shifts + $SU(3)$ projection (cheap), one tuning per HMC (negligible)
2. Coarse eigenvectors: $N = 100$ takes $\sim$ 1-2k outer CG iterations
3. ω_i: one tuning per HMC

Conclusion

Summary

- Using RG (APE blocking), we can generate $S_c^b[U_c]$ with low-energy physics matching that of $S_f[U_f]$

- Low-mode correspondence gives us a coarse near-null space with minimal setup time! Avoids cost of conventional Multigrid.

- Previously, the approach was hampered by high-mode pollution. Filtering remedies are prohibitively expensive.

- We find that restricting the RG preconditioner to the physical 4D components on the walls reduces high-energy noise while also reducing the cost by a factor of L_s

- Current results show RGPCG has $15 - 20\%$ speedup, (and we're looking for future improvements)

Outlook

1. Develop more systematic tuning of ω_i values
2. Try using 4d (Wilson) eigenvectors, rather than $\mathcal{P}_{5\to 4}\psi_c^{5D}$
3. In order to deflate the first 10-20 fine modes, try using an overcomplete spectrum of low modes at different m_l and m_{res}
4. Run at physical mass (and larger volumes)

Acknowledgments

Thank you for listening!

References:

1. J. Tu, *Lattice QCD Simulations towards Strong and Weak Coupling Limits*, Ph.D. thesis, Columbia University (2020). doi:10.7916/d8-74hj-ak60
2. D. Guo, *Fermion Low Modes in Lattice QCD: Topology, the η' Mass and Algorithm Development*, Ph.D. thesis, Columbia University (2021). doi:10.7916/d8-knph-2598.
3. J. Tu and R. D. Mawhinney, *A Numerical Study of Renormalization Group Transformations on Multiscale Lattices*, Columbia University (2017). doi:10.1051/epjconf/201817502006
4. T. Blum *et al.* (RBC Collaboration), *Quenched Lattice QCD with Domain Wall Fermions and the Chiral Limit*, Phys. Rev. D **69**, 074502 (2004), arXiv:hep-lat/0007038.