

Toward a Renormalization-Group Preconditioned Conjugate Gradient for Domain Wall Fermions

Jonah Eick^{a,*} and Robert Mawhinney^a

^a*Department of Physics, Columbia University, New York, NY 10027, USA*

E-mail: je2492@columbia.edu, rdm10@columbia.edu

Recent advances such as multigrid and deflation have significantly accelerated Dirac operator solvers in lattice QCD. However, the substantial setup costs of these methods have impeded their application in the repeated Dirac inversions required for Hybrid Monte Carlo (HMC) ensemble generation. Building on earlier work in the RBC/UKQCD collaboration, which showed that renormalization-group (RG) blocked coarse lattices with $a^{-1} = 1$ GeV provide a good approximation of the low-mode structure of Mobius domain wall fermion (MDWF) Dirac operators on $a^{-1} = 2$ GeV lattices, we investigate how to exploit this correspondence in practice. Specifically, we study how filters and 5th-dimensional tuning on the coarse lattice can be used to control the influence of high modes in the coarse-fine mapping. Because RG blocking is computationally inexpensive, this approach may offer a way to incorporate multigrid-style acceleration directly into ensemble generation.

*The 42nd International Symposium on Lattice Field Theory (LATTICE2025)
Tata Institute of Fundamental Research (TIFR), Mumbai, India
November 2 - 8, 2025*

*Speaker

1. Introduction

Domain wall fermions (DWF) achieve approximate chiral symmetry on the lattice by supporting chiral modes on the 4D boundaries of a 5D bulk, but such fermion discretizations incur additional computational costs [1]. For QCD, where quark masses are non-zero, reaching an acceptably small residual chiral symmetry breaking – which decreases exponentially with the extent of the 5th dimension (L_s) at weak coupling – typically requires L_s values in the range of 10 to 20. Because this extra dimension increases the size of the Dirac operator, evaluating quark propagators becomes highly computationally demanding. To calculate the quark propagator ψ for a given source η and Dirac operator D , one must solve a large linear system of the form $D\psi = \eta$, which is typically done by applying iterative Krylov solvers such as the Conjugate Gradient (CG) algorithm to the normal equations:

$$D^\dagger D \psi = D^\dagger \eta. \quad (1)$$

For DWF operators, the RBC/UKQCD collaboration currently needs as many as 40,000 iterations of CG to accurately solve light quarks on large lattices [2].

Techniques such as deflation and multigrid [3, 4] are used to accelerate DWF Dirac operator solvers for computing propagators in the evaluation of observables, where there are many Dirac operator inversions per gauge configuration. However, the long setup times of these methods render them ineffective in the context of ensemble generation via the Hybrid Monte Carlo (HMC) algorithm. In the HMC, the gauge field is continually updated during the molecular dynamics integration, requiring repeated light-quark solves on a constantly evolving background where setup costs cannot be easily amortized. A long-standing objective has been to develop a method that accelerates the Dirac operator solver without incurring such prohibitive setup costs.

We report progress on a possible solution to this problem: the *Renormalization-Group Preconditioned Conjugate Gradient* (RGPCG). This method is based on the idea that lattice ensembles with corresponding actions at different lattice spacings can be tuned to lie along the same RG trajectory [5]. By RG-blocking a “fine” ($a^{-1} = 2$ GeV) gauge configuration, one is able to generate a computationally inexpensive “coarse” ($a^{-1} = 1$ GeV) lattice that preserves the low energy physics of the fine lattice. Previous studies [6, 7] demonstrate a robust correspondence between the near null spaces of the MDWF Dirac operators on these grids. This correlation suggests that the coarse Dirac operator can serve as an effective low-energy approximation of the fine operator. By restricting the residual to the coarse lattice, computing the low-mode correction where Dirac inversions are cheaper, and interpolating the result back, we can effectively precondition and accelerate the convergence of the fine CG solver. Because RG blocking is computationally inexpensive and allows the use of the *native* Dirac operator $D(U_c, m_c)$ on the coarse grid, this multi-level procedure circumvents the substantial setup costs associated with constructing Galerkin coarse operators in standard multigrid methods. However, a persistent challenge in this approach is interpolating coarse fields back to the fine lattice without coupling to unwanted high-energy modes. In the following sections, we detail this problem and explore potential remedies.

2. Renormalization-Group Blocking

The RBC/UKQCD collaboration previously generated a pair of 2+1 flavor MDWF ensembles via the HMC, with Iwasaki gauge actions and dislocation suppressing determinant ratio (DSDR) terms [5]. The fine and coarse ensembles have inverse lattice spacings $a^{-1} = 2$ GeV and $a^{-1} = 1$ GeV, and respective actions $S_f[U_f]$ and $S_c[U_c]$. Both ensembles have a physical spatial volume of approximately $(2.4 \text{ fm})^3$ and a pion mass tuned to $m_\pi \approx 300$ MeV. Low energy hadron masses and meson decay constants were calculated on both ensembles and were found to agree at the 2 – 7% level (see the first two columns of Table 1). This suggests that S_f and S_c lie, to good approximation, on a renormalization group trajectory.

	$\langle O \rangle_f$	$\langle O \rangle_c$	$\langle O \rangle_c^b$
size	$24^3 \times 64 \times 12$	$12^3 \times 64 \times 12$	$12^3 \times 64 \times 12$
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)

Table 1: Comparison of parameters and low-energy observables for the fine, coarse, and blocked coarse ensembles in [7].

In Ref. [5], the authors observed that one can move along the RG trajectory with a simple blocking procedure. Given a blocking kernel

$$G[U_c, U_f] = \prod_{x\mu} \delta(U_c(x, \mu) - g_b[U_f; x, \mu]), \quad (2)$$

which maps a fine gauge configuration U_f to a coarse gauge configuration U_c , we can generate a new *blocked* coarse lattice with action $S_c^b[U_c]$. The function $g_b[U_f; x, \mu]$ determines how gauge links are mapped to the coarse lattice. The effective action is obtained by integrating out the fine gauge links:

$$e^{-S_c^b[U_c]} \propto \int [dU_f] e^{-S_f[U_f]} G[U_c, U_f]. \quad (3)$$

In Refs. [5, 7] and our studies, we have employed an APE-like blocking kernel. As depicted in Figure 1, this kernel smears two fine links with the surrounding $1 \times 2 \times 1$ link staples, and projects the result onto a single $SU(3)$ coarse link. It is important to note that this blocking only takes place in the four physical spacetime dimensions. The 5th dimension is untouched, and its size on the coarse scale remains $L_S = 12$.




Figure 1: Single-step APE-like Blocking Kernel [5].

Figure 2: Mapping between S_f , S_c and S_c^b in [5].

The blocking kernel in Figure 1 contains an adjustable parameter α . It was previously found that for $\alpha = 0.688$, the effective action $S_c^b[U_c]$ produced from $S_f[U_f]$ closely matches $S_c[U_c]$. As shown in Table 1, the low energy hadron masses and meson decay constants calculated from the blocked coarse ensemble \mathcal{E}_c^b agree with those of the original coarse ensemble \mathcal{E}_c within 1 – 3%. Evaluating these observables required a choice of input quark mass, which is determined by demanding that the total light quark mass is the same on the coarse

and blocked coarse lattices: $[am_l + am_{\text{res}}]_{\text{coarse}} = [am_l + am_{\text{res}}]_{\text{blocked coarse}}$ [5]. These results establish a strong agreement between the HMC-generated coarse ID + MDWF ensemble and the RG-blocked coarse ensemble. It follows that APE blocking gives a computationally efficient way to produce a coarse lattice with roughly the same low energy physics as the target system, $S_f[U_f]$. Figure 2 shows a diagram relating the ensembles. This physical correspondence does not guarantee a correlation between the fine and (blocked) coarse MDWF. A microscopic correlation is essential for the RGPCG, and will be demonstrated in the next section.

3. Interpolation and Restriction

In the aforementioned work and our current studies, we use the standard Mobius fermion operator $D_{\text{MDWF}}(U, m, M_5)$ [8] with Mobius coefficients $b_5 = 2.5$ and $c_5 = 1.5$. The domain wall height $M_5 = 1.8$ is fixed by the DSDR term. To reduce numerical difficulty, we have used the strange quark for our input masses: $am_f = 0.019896$ and $am_c = 0.064150$. Mapping between fine and coarse fermion fields requires a choice of interpolation (I) and restriction (R) operators. In order to leverage the RG correspondence to solve the normal equations, the coarse Dirac operator must approximate the fine Dirac operator in the near null space. That is, the interpolation and restriction operators must satisfy

$$I(D^\dagger D)_c R \approx (D^\dagger D)_f \quad (4)$$

in the low mode space. Hermiticity requires that $R = I^\dagger$. The covariant interpolation operator is given by

$$I_{\text{cov}} = \prod_{\mu=0}^3 (1 + T_\mu^{\text{cov}}) P_{\text{inj}}, \quad T_\mu^{\text{cov}} \psi(x) = \frac{1}{2} \left[U_\mu(x) \psi(x + \hat{\mu}) + U_\mu^\dagger(x - \hat{\mu}) \psi(x - \hat{\mu}) \right]. \quad (5)$$

where P_{inj} injects the coarse lattice vector onto the fully even ($eeee$) sites of the fine lattice. The corresponding restriction operator is $R_{\text{cov}} = I_{\text{cov}}^\dagger = P_{\text{inj}}^\dagger \prod_{\mu=3}^0 (1 + T_\mu^{\text{cov}})$. This procedure is simplified in a smooth gauge, where parallel transport becomes unnecessary and the gauge links in T_μ^{cov} are replaced by the identity to define I_{fix} . Because the field only needs to be relatively smooth, fixing to an *approximate* Landau gauge avoids prohibitive costs and justifies using the simplified operators I_{fix} and $R_{\text{fix}} = I_{\text{fix}}^\dagger$. Note that if we use (approximate) Landau-gauge-fixing matrices $g(x)$ on the fine lattice, we must use the coarse restrictions of these matrices on the coarse lattice:

$$U_f \rightarrow g U_f g^\dagger, \quad \text{and} \quad U_c \rightarrow (P_{\text{inj}}^\dagger g P_{\text{inj}}) U_c (P_{\text{inj}}^\dagger g^\dagger P_{\text{inj}}). \quad (6)$$

An optimal set of operators will be *nearly* unitary: $RI \approx \mathbb{I}_c$ and $IR \approx \mathbb{I}_f$ (enforcing $R = I^\dagger$ makes exact unitarity impossible, creating a necessary tradeoff). To test this, we use the known coarse low modes $\{\phi_c\}$ to construct a low-lying vector $\psi_c^l = \alpha_i \phi_c^i$. Then we measure the overlap

$$(RI\psi, \psi) \equiv \frac{\langle RI\psi | \psi \rangle}{|RI\psi| |\psi|}, \quad (7)$$

finding $(R_{\text{cov}} I_{\text{cov}} \psi_c^l, \psi_c^l) \sim 0.96 - 0.98$ and $(R_{\text{fix}} I_{\text{fix}} \psi_c^l, \psi_c^l) \sim 0.94 - 0.95$. In the other direction, we use a low-lying fine vector ψ_f^l and find $(I_{\text{cov}} R_{\text{cov}} \psi_f^l, \psi_f^l) \sim 0.85 - 0.86$ and $(I_{\text{fix}} R_{\text{fix}} \psi_f^l, \psi_f^l) \sim 0.88 - 0.89$. Given that the gauge-fixed operators perform similarly (slightly worse for RI , slightly better for IR) and are simpler, we work on an approximately Landau-gauge-fixed lattice throughout this study and use $I = I_{\text{fix}}$, $R = R_{\text{fix}}$.

To verify Eq. (4), we need to compare the fine and coarse low-mode subspaces. In [7], the approximate inverse MDWF operators were constructed using the N lowest fine and coarse eigenvectors of the operators:

$$G_f = \sum_i^N \frac{|\phi_{f,i}\rangle \langle \phi_{f,i}|}{\lambda_{f,i}} \quad \text{and} \quad G_c = I \sum_i^N \frac{|\phi_{c,i}\rangle \langle \phi_{c,i}|}{\lambda_{c,i}} R, \quad (8)$$

respectively. Ideally, $G_c^{-1}G_f = \mathbb{I}_f$. The matrix elements of this product were computed using

$$\langle \phi_{f,m} | G_c G_f^{-1} | \phi_{f,n} \rangle = \sum_i^N \langle \phi_{f,m} | I | \phi_{c,i} \rangle \langle \phi_{c,i} | R | \phi_{f,n} \rangle \frac{\lambda_{f,n}}{\lambda_{c,i}}. \quad (9)$$

The matrix elements, for $N = 100$ and 1000 , are displayed in Figure 3. Despite a scaling factor of ~ 0.1 , the matrix is extremely diagonally dominant, confirming that $G_c^{-1}G_f \propto \mathbb{I}_f$. This is an essential result for our study, as it demonstrates that $(D^\dagger D)_c$ is a good approximation of $(D^\dagger D)_f$ in the near null space.




Figure 3: Matrix elements of $G_c G_f^{-1}$ for $N = 100$ (left) and $N = 1000$ (right) [7]. The diagonal structure resembles the identity matrix, but the magnitude is small. In the right figure the diagonal line is very faint.

4. Restarts with RG Corrections

Having established the correspondence between fine and coarse low mode spaces, we can apply the RG mapping to a CG solver for the normal equations. To test the usefulness of our method for a preconditioner, we first add a simple restart to the CG. In the restart, a partial solve is done on the coarse lattice and the result is interpolated and added back to the fine state before resuming the fine CG. To simplify the notation, we define $A_f \equiv (D^\dagger D)_f$ and $A_c \equiv (D^\dagger D)_c$. Note that in applying the CG to a system $Ax = b$, the residual (r) is related to the error (e) by $e = A^{-1}r$. The full procedure is

1. **Fine Solve:** Run CG on $A_f \psi_f = D_f^\dagger \eta$ for n_f iterations and compute residual $r_f = D_f^\dagger \eta - A_f \psi_f$.
2. **Coarse Correction:** Restrict to coarse lattice ($r_c = R r_f$) and run CG on $A_c \psi_c = r_c$ for n_c iterations to yield e_c .
3. **Restart:** Interpolate the correction ($e_f = I e_c$) and restart fine CG with $\psi'_f = \psi_f + b e_f$.

The scale factor is chosen to be $b = 9.0$, which accounts for the difference in the scale of the fine and coarse spectrums seen in Figure 3. A sample run with $n_f = 400$ and $n_c = 1000$ is displayed in Figure 4. Only the fine CG residual and iterations are shown; the coarse CG iterations are negligible in terms of runtime.

The simple restart results in a large ($\sim 10^4$) spike in the residual. Although the restarted residual eventually descends below the no-restart residual, the difference is negligible. The spike is caused by high-mode pollution introduced during the interpolation. Directly after the restart, the residual is

$$r'_f = D_f^\dagger \eta - A_f(\psi_f + b I e_c) = r_f - b A_f I e_c. \quad (10)$$

Even if e_c is highly accurate, any small high-mode fluctuations in $I e_c$ are drastically amplified by A_f . For the configuration studied here, the hermitian spectrum extends to $\lambda \lesssim 520$, while the lowest 1000 modes are




Figure 4: CG with one RG restart after 400 iterations.

bounded by $\lambda \lesssim 0.05$. Therefore, A_f amplifies small high-mode corrections by four orders of magnitude, explaining the spike. Such effects shift the residual enough that any genuine low-mode progress made on the coarse lattice is effectively wiped out. Another way to see this misalignment is in the error. The “true” error is $e_* = \psi_* - \psi_f$, where ψ_* is a 2000-iteration approximate solution and ψ_f the 400-iteration state. We find the overlap with the interpolated error is only $(e_*, Ie_c) \approx 0.65$. Thus the correction we add to the CG solution is not well aligned with the true solution. To avoid this problem and make the RG mapping beneficial for a preconditioner, we must filter out the UV noise responsible for the spike.

5. Filtering

We have explored various smoothing/filtering techniques. In this discussion we will focus on a Chebyshev lowpass filter. For a hermitian operator A with extremal eigenvalues λ_{\min} and λ_{\max} , the filter is given by

$$F_N(A) = \frac{g_0 c_0}{2} T_0(\tilde{A}) + \sum_{n=1}^{N-1} g_n c_n T_n(\tilde{A}), \quad \tilde{A} \equiv \frac{2A - (\lambda_{\max} + \lambda_{\min})}{\lambda_{\max} - \lambda_{\min}} \quad (11)$$

where T_n is the n -th Chebyshev polynomial, c_i are the Clenshaw coefficients, and g_i are determined from the Jackson kernel to suppress Gibbs oscillations [9]. Figure 5 shows the shape of this filter for different orders N as a function of the eigenspectrum. The filter aggressively damps the spectrum outside of the near null space. An effective filter F_N should give $F_N I A_c^{-1} R \approx A_f^{-1}$ in the near null space. To test this, we generate a fine low-mode vector ψ_l from 300 CG iterations on $A_f x = 0$, and compute $\psi'_l = F_N I A_c^{-1} R A_f \psi_l$. We expect $\psi'_l \approx \psi_l$ for an effective filter. The alignment of ψ'_l with ψ_l is described in Table 2. The results indicate that an $O(150)$

	No filter	$N = 150$	$N = 300$	$N = 450$	$N = 600$
$\langle \psi'_l \psi_l \rangle / (\psi'_l \psi_l)$	0.621	0.884	0.904	0.909	0.901

Table 2: Overlap of ψ_l and ψ'_l for different degrees of filtering.

or $O(300)$ filter is sufficient to protect against high mode coupling in the interpolation. Applying this filter to the restart algorithm, we update the iterative solution as $\psi'_f = \psi_f + b F_N e_f$ before restarting the CG. Figure 6 shows that a single restart with the Chebyshev filter (after 400 iterations) leads to a meaningful improvement in the residual. The $O(150)$ filter suppresses the restart spike, while the $O(300)$ filter makes the spike negligible,




Figure 5: Chebyshev filters on the eigenspectrum, for different orders (N).

as high-mode pollution has been smoothed out, and the residual converges in ≈ 300 fewer iterations. In terms of the error, we now find $(e_*, F_{300}Ie_c) \approx 0.95$, confirming that the filtered correction is much more aligned with the true solution.

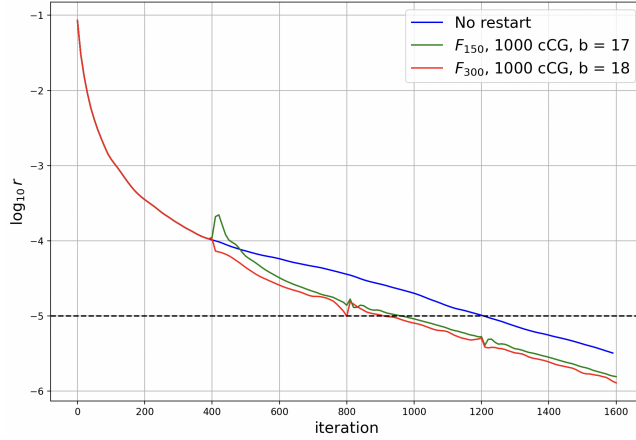


Figure 6: CG with filtered restarts. Blue: no restart. Green: one restart with the F_{150} filter. Red: one restart with F_{300} .

In terms of computational cost, the restart applied here roughly breaks even. Because evaluating a Chebyshev polynomial of degree N requires N matrix-vector multiplications, the filter F_N has a computational cost comparable to N applications of A_f . The 1000 coarse CG iterations amount to ≈ 60 fine CG iterations. Thus, in Figure 6, this extra cost is effectively equivalent to adding 360 additional CG iterations before the sharp dropoff in the F_{300} residual (red line). When F_{150} is used (green line), the residual reaches 10^{-5} in ≈ 220 fewer iterations than the no-restart baseline. However, the restart and filtering only takes ≈ 210 additional fine CG iterations, so the time to convergence is marginally improved. The small spikes at 800 and 1200 iterations are not subsequent restarts, but rather artifacts of the initial restart. This phenomenon is seemingly inherent to the CG and we hope to investigate it further.

Ultimately, a substantial speedup—which is the aim of an RG preconditioner—requires a change in the slope of the residual. Restarting the CG alone is not enough to cause such a speedup. Nevertheless, this

study shows that if we can more efficiently reach a similar degree of filtering, the RG procedure can be used to effectively precondition the normal equations.

6. The RG Preconditioner

In this section we describe how to incorporate our strategy in a preconditioner for the CG. First, although we have thus far viewed the coarse Dirac operator inversion as a coarse CG solve, it is also possible to use a spectral projection:

$$A_c^{-1} = \sum_{i=1}^N \frac{|\phi_{c,i}\rangle\langle\phi_{c,i}|}{\lambda_{c,i}}. \quad (12)$$

There are multiple tradeoffs between the coarse CG and the spectral projection. The latter requires calculating N coarse eigenvectors with the Lanczos algorithm before each Dirac solve, which introduces a setup cost. Furthermore, the inverse will only be supported on the coarse subspace spanned by the calculated eigenvectors. The coarse CG, on the other hand, requires no setup time and has support on the entire coarse eigenspace. However, the coarse CG introduces additional computational cost in each RGPCG iteration, and the coarse CG is a variable preconditioner, which may present further complications [10]. We plan to investigate this tradeoff in future work. Following [7], we consider the operator in Eq. (12). The resulting preconditioner is given by

$$M^{-1} = \mathbb{I}_f + bFI \sum_{i=1}^N \frac{|\phi_{c,i}\rangle\langle\phi_{c,i}|}{\lambda_{c,i}} RF. \quad (13)$$

It is essential that the preconditioner be symmetric, so the filter F is appended on the right side. Because this form of A_c^{-1} projects strictly into the subspace spanned by the N eigenvectors, we add the identity matrix \mathbb{I}_f to make M^{-1} surjective. The parameter b controls how strongly the preconditioner is incorporated. In [7], the preconditioner was tested with an exact filter using the eigenvectors of the fine operator: $F_* = \sum_i |\phi_{f,i}\rangle\langle\phi_{f,i}|$. Initial results for the RGPCG, with the preconditioner in Eq. (13), are shown in Figure 7.




Figure 7: Left: RGPCG with no filter. Right: RGPCG with the exact filter F_* . The blue lines are the CG with no preconditioner. Different scaling factors b are compared. From [7].

The RGPCG with no filter was found to give a small improvement, consistent with our earlier RG restart results. When the exact filter F_* is introduced, the RGPCG gives a $\sim 4x$ speedup. This method cannot be applied in ensemble generation, since the setup time required to obtain the exact fine eigenvectors defeats the purpose of the preconditioning. However, these results demonstrate the potential success of the RGPCG if an effective but less expensive filtering method is achieved.

7. Tuning the 5th Dimension

Another approach to securing the mapping between low mode subspaces is based on the 5D DWF structure. Because the 5th dimension is untouched in the RG blocking, it is likely that fine and coarse low modes have different 5th-dimensional shapes. Indeed, we have found that low-lying vectors on the coarse lattice have a slower fall-off in the 5th dimension, and the 5th dimensional profile is more jagged than on the fine lattice (see Figure 8). The mismatch is also evident in the residual mass: $am_{\text{res}} = 0.004522(12)$ on the fine lattice, while $am_{\text{res}} = 0.00847(21)$ on the RG blocked lattice. This effect may be relatively unnoticeable in the matrix elements of $G_c^{-1}G_f$, yet could introduce fluctuations in the coarse-lattice solution that appear as high mode noise after interpolation. We have addressed this problem by Wilson flowing the coarse lattice. In flow time t , the Wilson flow smooths the gauge field over a range $\sqrt{8t}$, which in turn reduces the propagation of light modes in the 5th dimension [11]. We have calculated the residual mass on \mathcal{E}_c^b as a function of flow time, finding $am_{\text{res}} = 0.00473$ at $t = 0.05$. Although the flow causes the residual masses to nearly match, the 5th-dimensional fall-off only changes modestly, as seen in Figure 8.

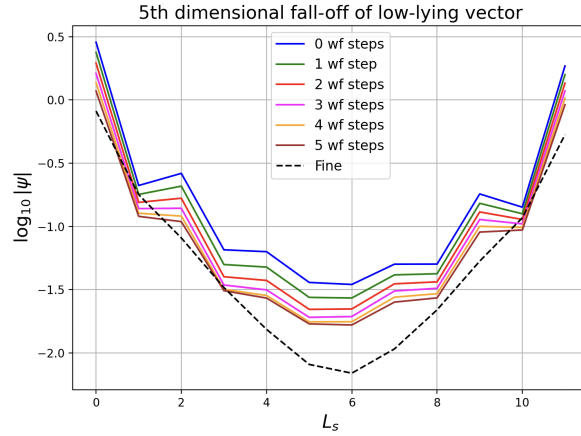


Figure 8: 5th dimensional fall-off of low-lying vectors. The solid colored lines show the fall-off on the coarse lattice after different numbers of flow steps ($\Delta t = 0.01$). The dashed black line is the fine lattice fall-off.

One could flow arbitrarily until the coarse fall-off reaches the fine fall-off, but the Wilson flow effectively changes the scale of the lattice. At $t = 0.03 - 0.04$, we have observed that the scale a^{-1} decreases by $\approx 5\%$. Thus, we can only flow a small amount in order to keep $a_c^{-1} \approx 1$ GeV. We plan to investigate different combinations of Wilson flow and RG blocking – such as flowing the fine gauge field before blocking it – in order to simultaneously tune the coarse lattice scale and the 5th-dimensional structure. Ultimately we care about the physical 4D modes on the walls, which gives us some freedom in how we treat the 5D bulk. Other methods of smoothing the 5th dimension – such as smearing over s -slices or filtering with the 5D transfer matrix – may be useful for suppressing high-mode pollution.

8. Conclusion

Using an APE-like blocking procedure, we can obtain an effective coarse action $S_c^b[U_c]$ that lies very close to the renormalization group trajectory of $S_f[U_f]$. A robust correspondence between the low mode spaces of the fine and coarse MDWF operators has been demonstrated. This differs from standard multigrid methods, wherein considerable time is spent constructing the coarse near null space. In our strategy, the RG provides the

near null space for free; the problem is how to *stay* in it. Interpolation introduces high-mode pollution, taking us out of the low mode space, which ruins the coarse-level correction to the CG solution. Conceptually, we understand how to handle this with Chebyshev filtering and other smoothing techniques, but such methods are computationally too expensive to be used in practice.

We believe that addressing the 5D structure of domain wall fermions is key to improving the mapping between low mode spaces. Some combination of smoothing and tuning the 5th dimension may suppress high-mode leakage in the 5D bulk. Another direction of interest is developing a cheaper low-mode filter on the fine lattice. We have seen that fine near-null eigenvectors provide a perfect filter, but they are too expensive. It may be possible to replicate the near null space of A_f using vectors in the near null space of A_c , which are much easier to compute. It remains to be seen whether interpolating and smoothing coarse low modes can produce approximate fine low modes at a significantly lower cost. If found true, this could make the RGPCG viable, and also be useful in deflation methods.

Acknowledgments

We would like to acknowledge earlier work by Jiqun Tu [5, 6] and Duo Guo [7], on which this project is based. We also thank our colleagues in the Columbia University Lattice QCD group for insightful discussions. The calculations in this work were done on the Genoa cluster at Brookhaven National Lab and made use of the Grid QCD Library. This research was supported in part by the U.S. Department of Energy, Office of Science, High Energy Physics under Award Number DE-SC0011941.

References

- [1] Y. Shamir, *Chiral Fermions from Lattice Boundaries*, *Nucl. Phys. B* **406** (1993) 90–106 [[hep-lat/9303005](#)].
- [2] T. Blum et al. (RBC/UKQCD Collaboration), *Domain wall QCD with physical quark masses*, *Phys. Rev. D* **93** (2016) 074505 [[arXiv:1411.7017](#)].
- [3] M. Lüscher, *Local coherence and deflation of the low quark modes in lattice QCD*, *JHEP* **07** (2007) 081.
- [4] P. A. Boyle, *Advances in algorithms for solvers and gauge generation*, *PoS LATTICE2023* (2024) 122.
- [5] J. Tu and R. Mawhinney, *A Numerical Study of Renormalization Group Transformations on Multiscale Lattices*, *EPJ Web Conf.* **175** (2018) 02006.
- [6] J. Tu, *Lattice QCD Simulations towards Strong and Weak Coupling Limits*, Ph.D. thesis, Columbia University (2020).
- [7] D. Guo, *Fermion Low Modes in Lattice QCD: Topology, the η' Mass and Algorithm Development*, Ph.D. thesis, Columbia University (2021).
- [8] R.C. Brower, H. Neff and K. Orginos, *The Möbius Domain Wall Fermion Algorithm*, *Comput. Phys. Commun.* **220** (2017) 1–19 [[arXiv:1206.5214](#)].
- [9] A. Weiße, G. Wellein, A. Alvermann and H. Fehske, *The kernel polynomial method*, *Rev. Mod. Phys.* **78** (2006) 275–306 [[arXiv:cond-mat/0504627](#)].
- [10] Y. Notay, *Flexible Conjugate Gradients*, *SIAM J. Sci. Comput.* **22** (2000) 1444–1460.
- [11] M. Lüscher, *Properties and uses of the Wilson flow in lattice QCD*, *J. High Energy Phys.* **08** (2010) 071.