

Multigrid algorithm for staggered lattice fermions

Richard C. Brower* and Evan Weinberg†

Boston University, Boston, Massachusetts 02215, USA

M. A. Clark‡

NVIDIA Corporation, Santa Clara, California 95050, USA

Alexei Strelchenko§

Scientific Computing Division, Fermilab, Batavia, Illinois 60510-5011, USA

(Received 7 March 2018; published 28 June 2018)

Critical slowing down in Krylov methods for the Dirac operator presents a major obstacle to further advances in lattice field theory as it approaches the continuum solution. Here we formulate a multigrid algorithm for the Kogut-Susskind (or staggered) fermion discretization which has proven difficult relative to Wilson multigrid due to its first-order anti-Hermitian structure. The solution is to introduce a novel spectral transformation by the Kähler-Dirac spin structure prior to the Galerkin projection. We present numerical results for the two-dimensional, two-flavor Schwinger model; however, the general formalism is agnostic to dimension and is directly applicable to four-dimensional lattice QCD.

DOI: [10.1103/PhysRevD.97.114513](https://doi.org/10.1103/PhysRevD.97.114513)**I. INTRODUCTION**

Increasingly powerful computers and better theoretical insights continue to improve the predictive power of lattice quantum field theories, most spectacularly for lattice quantum chromodynamics (LQCD) [1]. However, with larger lattice volumes and finer lattice spacing, exposing multiple scales, the lattice Dirac linear system becomes increasingly ill-conditioned threatening further progress. The cause is well known: as the fermion mass approaches zero, the Dirac operator becomes singular, due to the exact *chiral* symmetry of the Dirac equation at zero mass, causing *critical slowing down* [2]. The algorithmic solution to this problem for lattice QCD was recognized 25 years ago. The fine-grid representation for the linear solver should be coupled to multiple scales on coarser grids in the spirit of Wilson's real space renormalization group and implemented as a recursive multigrid (MG) preconditioner [3]. Early investigations in the 1990s introduced a gauge-invariant projective MG algorithm [4,5] with encouraging results for the Dirac operator in the presence of weak

(or smooth) background gauge fields near the continuum. However, in practice lattice sizes at that time were too small and the gauge fields were too rough to achieve useful improvements.

Not until the development of adaptive geometric MG methods [6,7] was a fully recursive MG algorithm found for the Wilson-Dirac discretization, which was able to transfer the strong background chromodynamics fields onto coarser scales and eliminate the ill conditioning of the Dirac kernel in the chiral limit. In spite of this achievement for the Wilson-Dirac and closely related twisted mass formulation [8,9], these are not the only important Dirac discretizations in common use in lattice field theory. Three other discretizations used extensively in high energy applications, which more faithfully represent chiral symmetry on the lattice, are referred to as the domain wall [10], overlap [11], and staggered [12] fermions. The application of adaptive geometric MG to these discretizations has proven to be more difficult, perhaps related to the improved lattice chiral symmetry. A two-level MG solver for domain wall fermions has been implemented [13,14] which shows some promise, and a non-Galerkin algorithm has been implemented for overlap fermions [15], but there has been no success at formulating an MG staggered algorithm. Moreover, since staggered lattice ensembles are now the largest available for LQCD, requiring $O(10^5)$ iterations for good convergence, improving staggered solvers is a critical issue. Here we introduce a novel solver with the Kähler-Dirac spin structure [16,17] that allows, at last, the construction of an effective multilevel adaptive geometric MG algorithm for staggered fermions.

*brower@bu.edu

†weinbe2@bu.edu

‡mclark@nvidia.com

§astrel@fnal.gov

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

The staggered fermion is a remarkable discretization [12,18] which closely resembles the continuum Dirac linear operator,

$$D^{ij}\psi_j(x) = [\gamma_\mu^{ij}(\partial_\mu - iA_\mu(x)) + m\delta^{ij}]\psi_j(x). \quad (1)$$

The lattice discretization replaces the derivative by a gauge-covariant central difference,

$$D_{x,y} = \frac{1}{2a} \sum_{\mu=1}^d \eta_\mu(x) [U(x, x+\mu)\delta_{x+\mu,y} - U^\dagger(x-\mu, x)\delta_{x-\mu,y}] + m\delta_{x,y}, \quad (2)$$

resulting in a sparse matrix operator on a hypercubic lattice with the background gauge fields $U(x, y)$ represented by highly oscillatory $SU(3)$ matrices on each link $\langle x, y \rangle$ of the lattice. The γ_μ matrices are replaced by a single staggered ± 1 sign: $\eta_\mu = (-1)^{\sum_{\nu < \mu} x^\nu}$. Similar staggered lattice realizations of Dirac fermions have proven valuable not only for lattice QCD investigations but also for a variety of physical systems such as graphene in condensed matter [19], supersymmetry [20], and strongly interacting conformal fixed points of possible interest for beyond the standard model (BSM) physics in the Higgs sector [21–29].

Unlike the Wilson and domain wall methods, the staggered discretization preserves the exact anti-Hermiticity of the continuum Dirac operator up to a real mass shift. In this sense it represents the most primitive (or even fundamental) discretization. It has no explicit spin matrices (γ_μ), so the Dirac spin structure only emerges in the continuum limit. Each 2^d lattice sub-block in four dimensions reassembles into four Dirac flavors (or tastes), the content of a single Kähler-Dirac fermion [30]. This is the structure that our MG algorithm exploits: dividing out the 2^d Kähler-Dirac spin structure transforms the spectrum into a near “circle” in the complex plane as illustrated in Fig. 3. The striking similarity of the resultant spectrum to the Wilson and overlap spectra is, we believe, essential to the success of our staggered MG algorithm.

In LQCD applications with staggered fermions, the system $D(U, m)_{ij}\psi_j = b_i$ is typically solved via Krylov methods on the Schur decomposed even/odd operator (or, equivalently, the red/black operator). Because the preconditioned operator is Hermitian positive definite, the system can be solved by the conjugate gradient (CG) algorithm. This method has proven robust, and there are some well established methods to fend off critical slowing down, such as EigCG [31] eigenvalue deflation or block Krylov solvers [32–35]. Block solvers do not remove critical slowing down, and deflation methods scale poorly with the volume in terms of the number of eigenvectors needed to remove critical slowing down. As explained in our earlier report [36], an adaptive geometric MG algorithm for the staggered

normal operator can easily be formulated which removes critical slowing down. However, this comes with a heavy overhead. A Galerkin coarsening of the normal equation introduces next-to-nearest neighbor (or corner) terms, resulting in a $2d + 2d(d-1)$ site coarse operator stencil, and in four dimensions increasing the off-diagonal terms from 8 to 32 terms. This becomes prohibitively expensive in terms of communication pressure in parallel strong scaling MG solvers [37–40].

The solution to this problem is to develop an MG algorithm directly on the staggered operator. In the interest of algorithm development, we consider a two-dimensional model system as opposed to the full four-dimensional QCD. The two-dimensional staggered fermion, coupled to an Abelian gauge theory, $U(x, x+\mu) = \exp[i\theta_\mu(x)]$ is the two-flavor Schwinger model in the continuum limit [41,42]. This is a fully nonperturbative quantum field theory which is an ideal analogue to four-dimensional QCD. Like QCD it exhibits confinement with a zero mass triplet of “pionlike” bound states in the chiral (zero mass) limit and instantons that present a topological mechanism which breaks chiral symmetry dynamically in the flavor singlet channel [43]. As such, this has proven to be a reliable test framework [6] prior to a full implementation for four-dimensional QCD. The reader is referred to an extensive literature to understand the physical features that guide our construction in two dimensions and the natural generalization to four dimensions.

The lattice Schwinger model has the action

$$S_{\text{lat}} = \bar{\chi}_x [D(U) + m_0]_{xy} \chi_y + \beta \sum_x U^{\text{plaq}}(x). \quad (3)$$

Introducing the lattice spacing a , the bare mass (m) and the gauge coupling (g) are given by dimensionless parameters, $m_0 = am$ and $\beta = 1/(a^2g^2)$, respectively. There are two important physical length scales determined by these parameters: (1) The fundamental gauge correlation length (or string length) measured by the Wilson loop area law is $l_\sigma = a\sqrt{2\beta}$. (2) The fundamental fermion length scale measured by the “pion” Compton wave length is $l_{M_\pi} = \frac{1}{M_\pi} \approx 0.5a(am)^{-2/3}\beta^{1/6}$ [42]. To approach the continuum both must be large relative to the lattice spacing. As an analogue to QCD, we should also approach the chiral regime with $l_{M_\pi}/l_\sigma \gg 1$. To control finite volume L^d and finite lattice spacing a errors, the four length scales should obey the constraint: $L \gg l_{M_\pi} \gg l_\sigma \gg a$.

This two-dimensional theory has been carefully selected because of its remarkable similarity to four-dimensional QCD in terms of both the underlying physics and the formal mathematical structure. Although at present our numerical tests are restricted to two dimensions, the entire formal structure is applicable to higher dimensions. The numerical analysis of a four-dimensional algorithm for

lattice QCD is under development in QUDA [44–46], an efficient GPU framework for LQCD applications. Results will be presented in a subsequent publication.

The organization of the paper is as follows. In Sec. II we give the mathematical framework of the staggered Dirac operator essential to our subsequent MG formulations. In Sec. III we consider a Galerkin projection of the original operator and explain why it fails as a MG preconditioner. We then contrast it with the coarse projection of our new Kähler-Dirac preconditioned operator. In Sec. IV we present in detail the construction of the staggered MG algorithm, followed by detailed numerical tests for the two-dimensional Schwinger model. In Sec. V we discuss some alternatives to our current implementation, which may be useful in the application of our staggered MG algorithm to four-dimensional LQCD and other staggered lattice simulations. For example, a method for exactly preserving complex conjugate eigenpairing and numerical tests thereof is presented in Sec. V and in Appendix A, respectively.

II. MATHEMATICAL PRELIMINARIES OF STAGGERED FERMIONS

The geometric structure of the staggered Dirac operator in Eq. (2) and its relationship to the low-lying eigenspectrum is central to our analysis. Indeed, the spectral properties are directly related to the failure of MG applied directly to the staggered operator. On the other hand, the geometric structure allows us to construct a new Kähler-Dirac preconditioned staggered operator with a dramatically changed spectrum suited to MG. To this end, we first present the free spectrum of the original staggered Dirac operator in comparison with our Kähler-Dirac preconditioned operator. Next, this is followed by a brief discussion of the impact on the Kähler-Dirac preconditioned spectrum when nonzero mass, the three-link Naik term, and interacting gauge fields are included.

A. Staggered Dirac operator

As noted in the Introduction, the staggered fermion is the lattice fermion discretization “closest” to continuum Dirac fermions. More concretely, many of its features are inherited directly from the discretization of the continuum action,

$$\mathcal{S} = \int d^d x \bar{\psi} [\gamma_\mu (\partial_\mu - iA_\mu) + m] \psi. \quad (4)$$

The naïve fermion discretization uses a central difference approximation for the first derivative, which causes the so-called “doubling” (or aliasing) problem [47]. In the continuum, a single naïve fermion gives 2^d Dirac fermions: 16 four-component spinors in four dimensions and 4 two-component spinors in two dimensions. The staggered construction reduces this multiplicity by spin diagonalizing

the Dirac structure, then dropping all but one of the $2^{d/2}$ copies. Explicitly, this spin diagonalization,

$$\begin{aligned} \bar{\psi}_{i,x} \gamma_\mu^{ij} [U(x, x + \mu) \psi_{j,x+\mu} - U(x, x - \mu) \psi_{j,x-\mu}] \\ \rightarrow \bar{\psi}_{i,x} \eta_\mu(x) [U(x, x + \mu) \psi_{i,x+\mu} - U(x, x - \mu) \psi_{i,x-\mu}], \end{aligned} \quad (5)$$

is achieved by the unitary field redefinition,

$$\psi_x \rightarrow \Omega_x \psi_x, \quad \bar{\psi}_x \rightarrow \bar{\psi}_x \Omega_x^\dagger, \quad (6)$$

with $\Omega_x = \gamma_1^{x_1} \gamma_2^{x_2} \dots \gamma_d^{x_d}$. Dropping all but one copy, with the replacement

$$\Omega_x^\dagger \gamma_\mu^{ij} \Omega_{x+\mu} = \eta_\mu(x) \delta_{ij} \rightarrow \eta_\mu(x) = \pm 1, \quad (7)$$

results in a partial solution of the *doubling* problem by reducing the fermion content to $2^{d/2}$ staggered Dirac fermions: 4 in four dimensions and 2 in two dimensions. It is convenient to write the staggered operator succinctly as

$$D\psi_x = \frac{1}{2} \sum_\mu \eta_\mu(x) [\mathbb{D}_\mu - \mathbb{D}_\mu^\dagger] \psi_x + m\psi_x, \quad (8)$$

in terms of the gauge covariant forward difference operator $\mathbb{D}_\mu \psi(x) \equiv U(x, x + \mu) \psi(x + \mu) - \psi(x)$ and its conjugate $\mathbb{D}_\mu^\dagger \psi(x) \equiv U^\dagger(x, x - \mu) \psi(x - \mu) - \psi(x)$.

The staggered operator has a few special properties not shared by other fermion discretizations. The staggered operator is anti-Hermitian up to a mass shift and is normal: $[D(U, m), D^\dagger(U, m)] = 0$, just like the continuum operator. This is in contrast to the Wilson discretization,

$$D_W(m) = \frac{1}{2} \sum_\mu [\gamma_\mu [\mathbb{D}_\mu - \mathbb{D}_\mu^\dagger] + r \mathbb{D}_\mu^\dagger \mathbb{D}_\mu] + m, \quad (9)$$

with its Hermitian second order Wilson (stabilization) term that decouples doublers but makes D_W non-normal in the interacting case. The Wilson term also explicitly breaks chiral symmetry. On the other hand, the staggered operator retains a single exact chiral symmetry in the interacting case,

$$\gamma_5 \rightarrow \epsilon(x) = \Omega_x^\dagger \gamma_5 \Omega_x = (-1)^{x_1 + x_2 + \dots + x_d}, \quad (10)$$

with $\epsilon(x)$ being the generator of the chiral symmetry. These good chiral properties give $\epsilon(x)D + D\epsilon(x) = 2m\epsilon(x) \rightarrow 0$ as $m \rightarrow 0$. The chiral projectors, $\frac{1}{2}(1 \pm \epsilon(x))$, partition the lattice into even and odd sublattices,

$$D\psi = b \rightarrow \begin{bmatrix} m & D_{eo} \\ D_{oe} & m \end{bmatrix} \begin{bmatrix} \psi_e \\ \psi_o \end{bmatrix} = \begin{bmatrix} b_e \\ b_o \end{bmatrix}. \quad (11)$$

Furthermore, D features an $\epsilon(x)$ Hermiticity, analogous to the γ_5 Hermiticity of the continuum Dirac operator.

The normal equations for the staggered operator are diagonal,

$$D^\dagger D\psi = \begin{bmatrix} m^2 - D_{eo}D_{oe} & 0 \\ 0 & m^2 - D_{oe}D_{eo} \end{bmatrix} \begin{bmatrix} \psi_e \\ \psi_o \end{bmatrix}. \quad (12)$$

The Schur-preconditioned system takes on a similar structure and is also Hermitian positive definite. For the free problem [i.e., unit gauge fields, $U(x, x + \mu) = 1$], there is an exact cancellation of all next-to-nearest neighbor “around-the-corner” terms in the normal operator. This is a result of the η_μ phases preserving a key property of the Dirac algebra when taking the product of η s around a plaquette,

$$\begin{aligned} \gamma_\mu \gamma_\nu (-\gamma_\mu) (-\gamma_\nu) &= -1 \\ \rightarrow \eta_\mu(x) \eta_\nu(x + \mu) \eta_\mu(x + \nu) \eta_\nu(x) &= -1. \end{aligned} \quad (13)$$

The result is a set of 2^d decoupled Laplace operators on a lattice with spacing $2a$ illustrated in Fig. 1. In this sense, the free staggered operator is truly the “square root” of the Laplace operator, similar to the continuum Dirac operator. We can immediately write down the eigenvalues of the free staggered operator, given by

$$\lambda(p, m) = m \pm i \sqrt{\sum_\mu \sin^2(p_\mu)}, \quad (14)$$




FIG. 1. The normal operator applied on an odd site “o.” All contributions to even sites “x” cancel due to D being normal. Links in black (solid) and red (dashed) correspond to ± 1 , respectively, due to the contributions of η_μ and the anti-Hermiticity of D . In the free field, it is clear that corner terms cancel.

where $p_\mu = 2n_\mu\pi/L$ for integers $n_\mu \in [-L/4 + 1, L/4]$ due to the shift-by-two translational invariance. The eigenvalues are imaginary (up to a real mass shift) and come in complex conjugate pairs. When an interacting gauge field is turned on, the “around-the-corner” terms no longer vanish, leaving the two decoupled components in Eq. (12). These next-to-nearest neighbor terms are the standard so-called clover term, resulting in an irrelevant, in the Wilsonian sense, spin-gauge interaction ($\sigma_{\mu\nu}F_{\mu\nu}$) in the continuum. The spectrum cannot be found analytically, but $\epsilon(x)$ Hermiticity symmetry ensures that the eigenvalues still appear in exact complex conjugate pairs.

B. Kähler-Dirac preconditioning

We now consider the spectral transformation which is essential to the staggered MG algorithm presented in Sec. III and tested numerically in Sec. IV. Here we will show that when the staggered operator is right preconditioned by the 2^d Kähler-Dirac blocks, the spectrum on the resultant $2a$ blocked lattice is dramatically different. In the free case, we prove that this transformation gives an exactly circular spectrum in the complex plane, similar to the overlap lattice Dirac discretization [11]. The inclusion of gauge fields and/or the three-link Naik term [48] are relatively small modifications of this basic circular structure.

The argument proceeds as follows. The staggered operator is composed of blocks containing 2^d sites, corresponding to 2^d degrees of freedom (d.o.f.) that in the continuum limit are recombined into a multiplet of Dirac fermions [49,50]. It is straightforward to see that the decomposition of the staggered operator in 2^d blocks of sites partitions the lattice, as illustrated in red in Fig. 2, into independent 2^d blocks B containing a plaquette of links. We will refer to these as Kähler-Dirac blocks. We also include the local mass term into this B block. The nearest-neighbor terms between the B blocks contribute to a block hopping term C , which is unitarily equivalent, up to the mass shift, to the block-local contributions in B . B and C only share sites at the corners of squares in two dimensions, cubes in three dimensions, and hypercubes in four dimensions. This is a dual decomposition: half of the links on the original lattice contribute to B , and half contribute to C , as represented in Fig. 2. We denote this partition between hopping terms within and across blocks as

$$D = \frac{1}{2} \eta_\mu [\mathbb{D}_\mu - \mathbb{D}_\mu^\dagger] + m = B + C. \quad (15)$$

We remark that we can interchange this dual description by shifting the coordinates $\vec{x}_i \rightarrow \vec{x}_i + \vec{1}$, where $\vec{1}$ is a vector of ones. We now construct the right-block-Jacobi or Kähler-Dirac preconditioned operator as

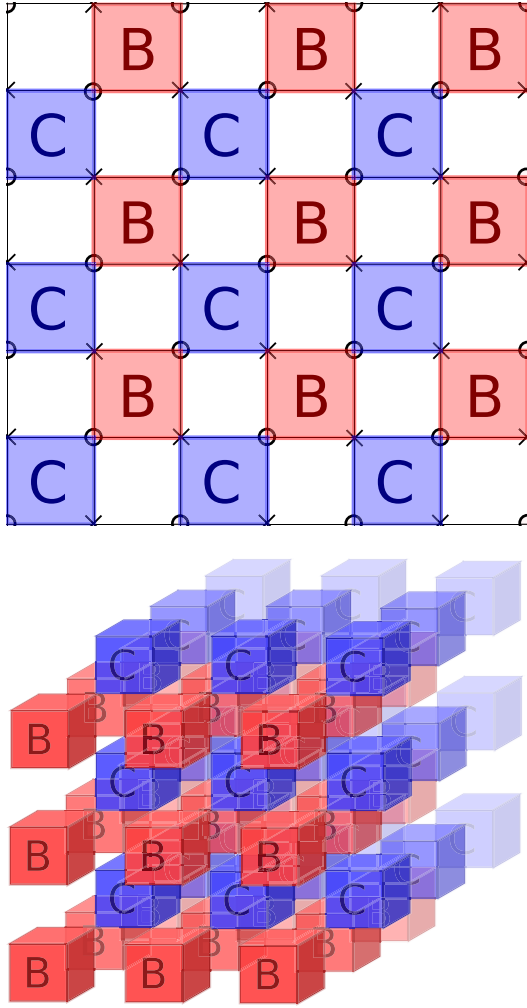


FIG. 2. On the top is the two-dimensional partition into the spin-tasted block B (in red) and the complementary connecting block C (in blue). On the bottom is the three-dimensional partition again into spin-tasted block B and the complementary connecting block C . In any dimension these partitions into B and C contain all N lattice sites of the periodic lattice and each share half of the dN links.

$$\boxed{A = DB^{-1} = \mathbb{I} + CB^{-1}} \quad (16)$$

This is a remarkably different operator with which we develop our MG algorithm.

To characterize these differences, we will first consider the free case. After rescaling and multiplying by $\epsilon(x)$, the generator of the exact staggered chiral symmetry, both terms are separately *Hermitian*, *traceless*, and *unitary*. More concretely, we define $\hat{B} = B\epsilon(x)/\sqrt{d+m^2}$ and $\hat{C} = C\epsilon(x)/\sqrt{d}$ and note

$$\hat{B}\hat{B}^\dagger = \hat{B}^2 = \mathbb{I}, \quad \hat{C}\hat{C}^\dagger = \hat{C}^2 = \mathbb{I}, \quad (17)$$

as a trivial consequence of the perfect cancellation of the corner terms for Eq. (13). These properties imply that \hat{B}

and \hat{C} have equal numbers of ± 1 eigenvalues, and further that the product $\hat{C}\hat{B}$ is a unitary matrix U . (The addition of a Naik term does not change \hat{B} , but it does contribute to \hat{C} .)

With this observation, our free Kähler-Dirac staggered operator A is given by

$$\begin{aligned} A &= DB^{-1} = \mathbb{I} + CB^{-1} \\ &= \mathbb{I} + \sqrt{\frac{d}{d+m^2}} \hat{C}\hat{B} = \mathbb{I} + \rho U. \end{aligned} \quad (18)$$

The eigenvalues of DB^{-1} lie on a circle centered at 1 as illustrated in the top panel of Fig. 3. The radius of the circle is $\rho = \sqrt{\frac{d}{d+m^2}}$. In the massless limit, the radius is exactly 1. This leads to an identical structure to the overlap operator,

$$D_{ov} = 1 + \gamma_5 \hat{\gamma}_5, \quad \hat{\gamma}_5 \equiv \text{sign}[\gamma_5 D_W(-M)], \quad (19)$$

under the mapping $(\gamma_5, \hat{\gamma}_5) \rightarrow (\hat{C}, \hat{B})$. Both \hat{C} and \hat{B} are algebraically similar to γ_5 and $\hat{\gamma}_5$, being Hermitian and unitary with an equal number of ± 1 eigenvalues. Adding a mass term to the overlap operator similarly rescales the unitary portion of the spectrum,

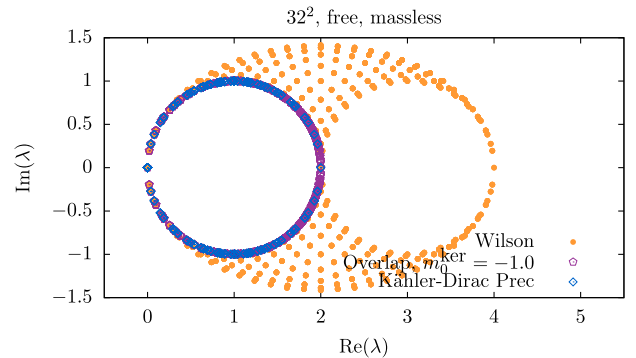



FIG. 3. The spectrum of the free, massless Kähler-Dirac preconditioned operator, with and without a Naik term on top, compared against the Wilson, overlap, and original staggered formulations on the bottom.

$$D_{ov} = \mathbb{I} + \frac{1-m}{1+m} \gamma_5 \hat{\gamma}_5, \quad (20)$$

introducing a mass gap. For comparison, in the bottom panel of Fig. 3, we show the free spectrum of the massless two-dimensional Wilson operator, the two-dimensional overlap operator, and our new two-dimensional Kähler-Dirac preconditioned operator. Very similar figures apply to four dimensions, except the Wilson spectrum now has four arcs in the positive real direction. Finally, we note that if we add a Naik term to the original staggered operator, the right preconditioning perturbs the unitarity of the spectrum but preserves the qualitative geometric features. A comparison of the Kähler-Dirac preconditioned operator to the original staggered operator is given in the top panel of Fig. 3. We compare the massless spectrum against all other fermion discretizations in Fig. 4.

The Kähler-Dirac operator no longer admits a simple “ γ_5 ” Hermiticity condition. However, it does obey a modified asymmetric $\gamma_5^{L/R}$ condition, which is essential for our discussions in Sec. III. The key observation is to note that we can change the “convention” that A is given by a right block preconditioning of D to a left block preconditioning via the transformation $\epsilon(x)A^\dagger\epsilon(x) = B^{-1}D$. We can rearrange this identity and note

$$\epsilon(x)B^{-1}A = A^\dagger\epsilon(x)B^{-1} \Rightarrow \gamma_5^L A = A^\dagger \gamma_5^L, \quad (21)$$

and likewise we can take advantage of the $\epsilon(x)$ Hermiticity of D to note

$$AB\epsilon = B\epsilon A^\dagger \Rightarrow A\gamma_5^R = \gamma_5^R A^\dagger. \quad (22)$$

Here we have defined $(\gamma_5^L, \gamma_5^R) = (\epsilon(x)B^{-1}, B\epsilon(x))$. This is a generalization of the idea of γ_5 Hermiticity: now, $\gamma_5^L \gamma_5^R = 1$ and $A^\dagger = \gamma_5^L A \gamma_5^R$. Also, just as is the case for the Wilson and staggered operators, these properties are enough to show

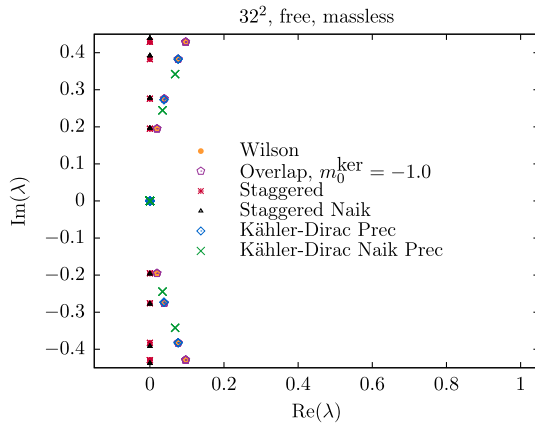


FIG. 4. The free field spectrum of the staggered, Kähler-Dirac preconditioned (with and without the Naik contribution), Wilson, and overlap operators.

that A features complex conjugate eigenpairs. Assume $A|\lambda^R\rangle = \lambda|\lambda^R\rangle$, where the superscript R denotes it is the right eigenvector corresponding to the eigenvalue λ . We can take the Hermitian conjugate of each side of the equation. Next, we can right multiply by γ_5^L and take advantage of $\gamma_5^{L/R}$ Hermiticity. This gives us $\langle\lambda^R|\gamma_5^L A = \langle\lambda^R|\gamma_5^L \lambda^*$; that is, A also has an eigenvalue λ^* with a left eigenvector $\langle\lambda^{*L}| \equiv \langle\lambda^R|\gamma_5^L$. This same exercise can be trivially repeated for left eigenvectors using γ_5^R to the same end.

C. Free spectrum after Kähler-Dirac preconditioning

For a detailed analysis of the spectrum, we introduce the flavor representation of the staggered operator [49–51], which is unitarily equivalent to a lattice Kähler-Dirac fermion in the free field [17,52]. Here each submatrix B is expressed in terms of the spin-taste gamma matrices which enumerate the components of a single continuum Kähler-Dirac fermion [30]. Its action is

$$S = b^d \sum_{X,\mu} \bar{q}(X) \left[\nabla_\mu (\gamma_\mu \otimes 1) - \frac{b}{2} \Delta_\mu (\gamma_5 \otimes \tau_\mu \tau_5) + m \right] q(X), \quad (23)$$

where $q(X)$ is the Kähler-Dirac field containing 2^d d.o.f., X is the Kähler-Dirac block index, $b = 2a = 2$ is the lattice spacing between Kähler-Dirac blocks, and the finite difference operators are defined as

$$(\nabla_\mu q)(X) = \frac{q(X + b\hat{\mu}) - q(X - b\hat{\mu})}{2b}, \quad (24)$$

$$(\Delta_\mu q)(X) = \frac{q(X + b\hat{\mu}) - 2q(X) + q(X - b\hat{\mu})}{b^2}. \quad (25)$$

In the language of staggered fermions, the γ_μ matrices generate the spin algebra, while the matrices $\tau_\mu = \gamma_\mu^\dagger$ generate the so-called taste algebra. It should be noted that if these lattice fermions are gauged on the lattice with twice the lattice spacing $b = 2a$, the resulting lattice theory of interacting Dirac-Kähler fermions [16,53] is no longer equivalent to the interacting staggered fermion and, of note, can generate a dynamical mass term. Likewise, on a continuum Riemann manifold, a Kähler-Dirac fermion admits a different gravitational gauging than Dirac fermions [54].

Our decomposition of $D = B + C$ is now partitioning Eq. (23) into local and nearest neighbor contributions. The local block B is given by

$$B = -B^\dagger \Leftrightarrow \sum_\mu \gamma_5 \otimes \tau_\mu \tau_5 \quad (26)$$

in the massless case. The inverse is given by

$$B^{-1} = -B^{-\dagger} \Leftrightarrow -\frac{1}{d} \left(\sum_\mu \gamma_5 \otimes \tau_\mu \tau_5 \right). \quad (27)$$

The transformation $D \rightarrow A = DB^{-1}$ gives the kernel

$$A = DB^{-1} = -\frac{1}{d} \sum_{\mu,\nu} \left[\nabla_{\mu}(\gamma_{\mu}\gamma_5 \otimes \tau_{\nu}\tau_5) + \frac{1}{2} \Delta_{\mu}(1 \otimes \tau_{\mu}\tau_{\nu}) \right]. \quad (28)$$

This operator can be explicitly diagonalized in arbitrary dimension by noting the Hermitian and anti-Hermitian projections of the operator commute, the Hermitian projection can be trivially diagonalized, and the imaginary part is prescribed by recalling the shifted unitary structure of the spectrum. This gives

$$\lambda(p_{\mu}) = 1 - \frac{1}{d} \sum_{\mu} \cos(p_{\mu}) \pm i \sqrt{1 - \left(\frac{1}{d} \sum_{\mu} \cos(p_{\mu}) \right)^2}. \quad (29)$$

This spectrum is visualized on the top panel of Fig. 3. The spectrum can be written as $1 - e^{i\theta}$ for where $\cos(\theta) = d^{-1} \sum_{\mu} \cos(p_{\mu}) \in [-1, 1]$. We note again that, up to a scaling, the low spectrum is similar to the Wilson, overlap, and staggered spectra.

1. Nonzero mass term

The spectrum undergoes a minor change when the original staggered operator is massive. The local block now becomes $\sum_{\mu}\gamma_5 \otimes \tau_{\mu}\tau_5 + m$, and the preconditioned spectrum becomes

$$\lambda(p_{\mu}) = 1 - \rho \left[\frac{\rho}{d} \sum_{\mu} \cos(p_{\mu}) \pm i \sqrt{1 - \left(\frac{\rho}{d} \sum_{\mu} \cos(p_{\mu}) \right)^2} \right], \quad (30)$$

where $S_n = \frac{1}{2} \sum_{\mu=x,y} \cos(np_{\mu})$ and $x = \frac{-1/48}{9/16}$, the ratio of the improvement coefficients in the Naik-improved action. The improved spectrum is shown on the left panel of Fig. 3, with the low modes emphasized in Fig. 4. We again make the critical observation that the spectrum is *qualitatively* similar to the original Kähler-Dirac spectrum, Wilson spectrum, and overlap spectrum.

3. Interacting staggered fermions in Kähler-Dirac form

We are ultimately interested in performing this right-block-Jacobi preconditioning on the interacting staggered operator, not the free operator. Procedurally, this is done by first gauging the staggered operator, and then performing the same unitary blocking transformation between the staggered

which parametrizes the arc of a circle $1 - \rho e^{i\theta}$ centered at $(1, 0)$ with radius $\rho = \sqrt{\frac{d}{d+m^2}}$. The arc is bounded to the range $\cos(\theta) = \rho d^{-1} \sum_{\mu} \cos(p_{\mu}) \in [-\rho, \rho]$, giving a gap $\lambda(0) = 1 - \rho^2 \pm i\sqrt{1 - \rho^2}$.

2. Naik term

Many modern LQCD simulations add a next-to-next-to-nearest neighbor improvement term known as a Naik [48] term. Two common realizations of this improvement, equivalent in the free-field limit, are AsqTad [55] and HISQ [56] fermions. The free operator [52] is given by

$$D(m)_{x,y}^{\text{Naik}} = -\frac{1}{2} \sum_{\mu} \eta_{\mu}(x) \left[\frac{9}{16} (\delta_{x,y-\hat{\mu}} - \delta_{x,y+\hat{\mu}}) \frac{1}{48} (\delta_{x,y-3\hat{\mu}} - \delta_{x,y+3\hat{\mu}}) \right] + m\delta_{x,y}. \quad (31)$$

The improved action admits the spectrum

$$\lambda_{\text{Naik}}(p, m) = m \pm i \sqrt{\sum_{\mu} \sin^2(p_{\mu}) \left[1 + \frac{1}{6} \sin^2(p_{\mu}) \right]^2}. \quad (32)$$

The effect of the Naik term on the Kähler-Dirac action is to modify the nearest neighbor term and add a next-to-nearest neighbor term. The 2^d Kähler-Dirac block $B_{\text{Naik}} = \frac{9}{16} B$ is *unchanged* up to a trivial rescaling. The new contributions are confined to C , which is no longer unitary: $\hat{C}_{\text{Naik}}^{\dagger} \hat{C}_{\text{Naik}} \neq \mathbb{I}$. Likewise, the spectrum is no longer a shifted unitary spectrum. Indeed, in two dimensions, the massless free spectrum is given by

$$\lambda(p_{\mu}) = 1 - [(1-x)S_1 + xS_2] \pm i \sqrt{1 - [(1-x)S_1 + xS_2]^2 - 2x[S_2 - S_1 + x(S_3 - 1)]}, \quad (33)$$

form and the Kähler-Dirac form. The local block no longer has a simple structure because of gauge links [17]. In two dimensions, the Kähler-Dirac block B attached to a unit corner at $2\vec{n}$ on the original staggered lattice is given by

$$\begin{pmatrix} m & 0 & -\frac{1}{2}U_x(2\vec{n}) & -\frac{1}{2}U_y(2\vec{n}) \\ 0 & m & -\frac{1}{2}U_y^{\dagger}(2\vec{n}+\hat{x}) & \frac{1}{2}U_x^{\dagger}(2\vec{n}+\hat{y}) \\ \frac{1}{2}U_x^{\dagger}(2\vec{n}) & \frac{1}{2}U_y(2\vec{n}+\hat{x}) & m & 0 \\ \frac{1}{2}U_y^{\dagger}(2\vec{n}) & -\frac{1}{2}U_x(2\vec{n}+\hat{y}) & 0 & m \end{pmatrix}. \quad (34)$$

Like the free case, the block B is still anti-Hermitian plus a massive shift. However, unlike the free case, the interacting

\hat{B} and \hat{C} are not unitary, and as such the product CB^{-1} does not have a unitary spectrum. Nonetheless, the spectrum is still approximately circular and centered at 1, as can be seen later in Fig. 9. This is a desirable property for matrix preconditioning in general [57,58], and it is essential for a successful MG algorithm. Importantly, this operator still maintains the $\gamma_5^{L/R}$ Hermiticity defined in Eqs. (21) and (22). This is true because the proofs of $\gamma_5^{L/R}$ Hermiticity solely depend on the $\epsilon(x)$ Hermiticity of D , which holds in the interacting case. By extension, the proofs of complex conjugate eigenpairing still hold. These comments carry over as appropriate when a N aik term is also included.

III. MULTIGRID COARSE OPERATOR

In forming the Galerkin projection of the staggered operator, we follow the methods of previous successful formulations of MG for the Wilson-Dirac discretization for LQCD [6]. Near-null vectors, or vectors which predominantly span the low-right eigenspace of the Wilson operator, are constructed by relaxing on the homogeneous equation with a random initial guess as is discussed in detail in Sec. IV. Later in this section we will also consider *exact low eigenvectors* programmatically as near-null vectors. The resulting near-null vectors are chirally doubled and block orthonormalized to construct the rows of the restrictor matrix R , which aggregates fine d.o.f. to a single site on the coarse lattice, and the prolongator matrix P , which maps coarse d.o.f. back to the fine lattice. Unless otherwise noted, $R = P^\dagger$. For the staggered operator, this implies that coarsening preserves the anti-Hermitian plus mass-shift structure. Block orthonormalization implies $P^\dagger P = \mathbb{I}$. The prologator and restrictor can be used to define the coarse operator,

$$\hat{D} = RDP. \quad (35)$$

The *hat* notation refers to an operator one level coarser than the “unhatted” operator.

We will begin by reviewing the Wilson formulation, largely to establish notation. We will extend this formulation to the staggered operator and show why this method fails to produce an effective recursive algorithm in this case. We will last repeat this formulation for the K ahler-Dirac preconditioned operator and show that, in contrast to the original staggered case, this method succeeds.

A. Review of Wilson-Dirac coarse operator

We begin by a basic restatement of the procedure for the adaptive geometric MG developed for the Wilson operator in QCD [5,6]. It is important to first note that the Wilson operator *does* obey γ_5 Hermiticity, that is, $\gamma_5 D_W \gamma_5 = D_W^\dagger$. The γ_5 Hermiticity is sufficient to prove that eigenvalues of D_W come in complex conjugate pairs as the limiting case of

$\gamma_5^L = \gamma_5^R$, as discussed in Sec. II B. Returning to MG, n_{vec}^1 near-null vectors are generated, where the “1” refers to coarsening the finest level, as discussed in Sec. IV. A key next step is chiral doubling: every near-null vector $|\psi_i\rangle$ is “doubled,” giving $\frac{1}{2}(1 \pm \gamma_5)|\psi_i\rangle$. For this reason, on the coarse operator, each coarse site has $2n_{\text{vec}}^1$ internal d.o.f., or, alternatively, a dense structure of n_{vec}^1 “coarse color” dof times two “chirality” dof. A successful implementation of Wilson MG critically depends on the preservation of chirality.

After performing a chiral doubling of the near-null vectors, we pack the doubled vectors into the prolongator

$$\tilde{P} = \left[\frac{1}{2}(1 + \gamma_5)|\psi_i\rangle, \frac{1}{2}(1 - \gamma_5)|\psi_i\rangle \right], \quad (36)$$

and again define $\tilde{R} = \tilde{P}^\dagger$. The *tilde convention* here is an indication that we have not (yet) block orthonormalized the $2n_{\text{vec}}^1$ vectors on each block. The chiral doubling implies $\gamma_5 \tilde{P} = \tilde{P} \sigma_3$, where $\sigma_3 = \text{diag}[1, \dots, 1, -1, \dots, -1]$ is a block Pauli matrix, or alternatively, the traditional σ_3 acting on the coarse chirality dof. It is easy to see that $\tilde{P}^\dagger D_W \tilde{P}$ is “ σ_3 ” Hermitian:

$$\sigma_3 \tilde{P}^\dagger D_W \tilde{P} \sigma_3 = \tilde{P}^\dagger \gamma_5 D_W \gamma_5 \tilde{P} = \tilde{P}^\dagger D_W^\dagger \tilde{P} = (\tilde{P}^\dagger D_W \tilde{P})^\dagger. \quad (37)$$

The essential property $\gamma_5 \tilde{P} = \tilde{P} \sigma_3$ is unchanged after we perform the last step, block orthonormalizing \tilde{P} to get P , because we performed our chiral doubling with a bona fide projector. The top chiral components and the bottom chiral components are already trivially orthonormal. This gives the final essential properties $\gamma_5 P = P \sigma_3$, and $\hat{D}_W = P^\dagger D_W P$ is σ_3 Hermitian. This methodology can be trivially extended to a recursive coarsening.

B. Failure of Galerkin projection of staggered operator

The prescription for (recursively) generating a coarse refinement of the Wilson operator D_W fails when naively translated to the staggered operator D with the only change being the replacement of γ_5 with $\epsilon(x)$, as noted by Eq. (10). While the iterative inversion of the even/odd preconditioned system exhibits critical slowing down, it does converge. However, this attempt at a Galerkin MG on the staggered operator D stalls completely at large volumes as illustrated in Fig. 5. We need to understand the cause of the failure of the Galerkin projection $\hat{D} = P^\dagger D P$ as a preconditioner. A MG algorithm may fail because the coarse operator does not accurately reproduce the low eigenspace of the fine operator, or because the coarse error “correction” is ineffective. We will study each of these properties for the staggered operator to attempt to understand the issue.




FIG. 5. The relative residual as a function of fine D applications for CG on the even/odd system compared with a failing MG solve on the staggered operator. The parameters of the MG solve are given in Table I.

As a spectral preconditioner, we expect the coarse operator to approximately preserve the low eigenmodes of the fine operator. In Fig. 6 we address this issue by comparing our failed staggered MG spectra with the successful Wilson MG spectra in Fig. 7.

First consider the staggered case. The center column (red squares) in Fig. 6 gives the positive imaginary component of the low-lying spectrum for the fine staggered operator. The spectrum is exactly paired with complex conjugate eigenvalues below the real axis due to $\epsilon(x)$ Hermiticity on the fine level and coarse levels. The other four columns give the spectrum for a recursively coarsened operator, constructing the prolongator/restrictor from exact low eigenvectors (left side) and near-null vectors (right side), where the near-null vectors are again generated as discussed in Sec. IV. Filled shapes correspond to the first coarse level. Hollow shapes correspond to the operator from a recursive coarsening. The horizontal black lines trace the low modes

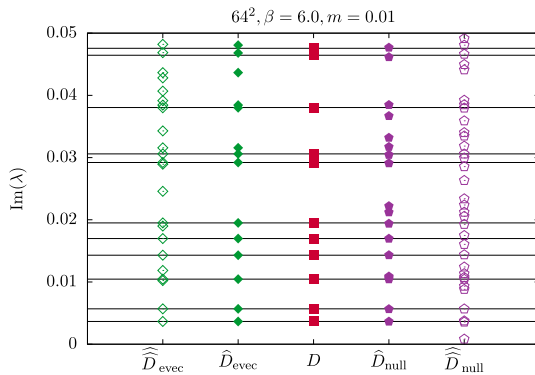


FIG. 6. The imaginary part of the spectrum of D for a representative configuration using a recursive symmetric Galerkin projection with both eigenvectors (left) and generated near-null vectors (right). Recursive coarsening introduces “spurious” low modes. The parameters of the MG coarsening are given in Table I.




FIG. 7. A representative spectrum of an interacting Wilson operator and its Galerkin coarsening. The MG coarsening aggregates fine d.o.f. over a 4^2 block into a coarse site using $n_{\text{vec}}^1 = 4$ near-null vectors.

of the fine operator across the coarsened operators. While these physical low modes are well preserved in all cases, there are many additional, *spurious* low eigenvalues in the coarse spectrum.

These spurious eigenvalues have a simple origin. Consider the normalized eigenvector of the coarse operator $\hat{D}|\hat{\lambda}\rangle = \hat{\lambda}|\hat{\lambda}\rangle$. We note that $\hat{\lambda} = \langle \hat{\lambda} | \hat{D} | \hat{\lambda} \rangle = \langle \hat{\lambda} | P^\dagger \rangle D (P | \hat{\lambda} \rangle)$. If we perform an eigendecomposition of $P|\hat{\lambda}\rangle = \sum_i c_i |\lambda\rangle$, we find that $\hat{\lambda} = \sum_i |c_i|^2 \lambda$, a consequence of D being normal. This implies $\hat{\lambda}$ is some linear combination of eigenvalues of D in the interval $[-\lambda_{\text{max}}, \lambda_{\text{max}}]$. A coarse eigenvalue can have an arbitrarily small imaginary part if, for example, $P|\hat{\lambda}\rangle$ is dominantly spanned by a pair of fine eigenvectors with complex conjugate eigenvalues. This eigenvector may have nothing to do with the low modes of D .

This would be less of an issue if higher modes were gapped along the real axis. This is true of the Wilson operator, as can be seen for a representative case in Fig. 7.¹ For the fine operator, whose eigenvalues are given by red squares, high modes are gapped along the *real* axis. For the coarse operator, whose eigenvalues are blue triangles, low modes are well preserved. Higher modes “collapse” toward the complex origin but are still well gapped along the *real* axis. This could be why MG on the Wilson operator does not break down, and may identically predict success for the Kähler-Dirac preconditioned operator.

Local collinearity versus the oblique projector. The Galerkin MG scheme involves two different projection operators:

- (i) The projection operator, $\mathbb{P} = PR$, from the fine space into a coarse subspace. Using the right eigenvectors as a basis for the fine vector space,

¹In the interacting case, the Wilson operator is no longer normal, and our convex hull proof breaks down. It appears that it is still sufficiently true, perhaps because the free Wilson operator is exactly normal. In a perturbative sense, the interacting Wilson operator is then “approximately” normal.

$V = \{|v_\lambda\rangle, 0 < |\lambda| \leq |\lambda_{\max}|\}$, our goal is for eigenvectors with small (near-null) eigenvalues, $|\lambda|/|\lambda_{\max}| < \epsilon$, to be approximately represented within the span of the coarse subspace, $\hat{V} = \mathbb{P}V$, in a least-squares sense.

(ii) The oblique (or Petrov-Galerkin) projector, $\mathbb{P}_{ob} = 1 - P(RDP)^{-1}RD$, that defines space of error components that are returned to the fine level with a complete solve in the coarse subspace. To not overburden the smoother this should at least not unduly amplify large eigenvectors, $|\lambda|/|\lambda_{\max}| > \epsilon$. Both are true projectors dividing the fine vector space V into disjoint subspaces, $\mathbb{P}(1 - \mathbb{P}) = 0$ and $\mathbb{P}_{ob}(1 - \mathbb{P}_{ob}) = 0$, though they do not define the same subspaces. The orthogonality, $\mathbb{P}_{ob}\mathbb{P} = 0$ is one-sided since $\mathbb{P}\mathbb{P}_{ob} \neq 0$.

Let us see how well the staggered MG handles these two requirements. In our construction, $R = P^\dagger$, so the coarse space projector is Hermitian. The statement of preserving the low eigenspace in the least-squares sense can be formulated as sufficiently minimizing

$$\|(1 - PR)\vec{v}_i\|_2, \quad (38)$$

for small eigenvalues of fine operator D . Since we generate our coarse space by geometric aggregation, this can be thought of as the *local collinearity* of near-null vectors with low eigenmodes. In the top left panel of Fig. 8, we see that starting with a block-orthonormalized basis either of near-null vectors or of low eigenvectors results in a good coverage of the low spectrum. This is typical of MG methods. At the bottom left this is extended to the next coarsest level with a similar result. This has important implications for eigenvector compression methods [59].

However, this is not sufficient for a successful coarse correction in an MG algorithm. The coarse correction should address the low modes of the fine operator without introducing large errors in the high mode subspace. The error after solving the coarse level is updated as $e \leftarrow e - P(RDP)^{-1}RDe = \mathbb{P}_{ob}e$. This is quantified by the magnitude of each eigenvector acted on by the Petrov-Galerkin or the so-called *oblique projector*,

$$64^2, \beta = 6.0, m = 0.01$$

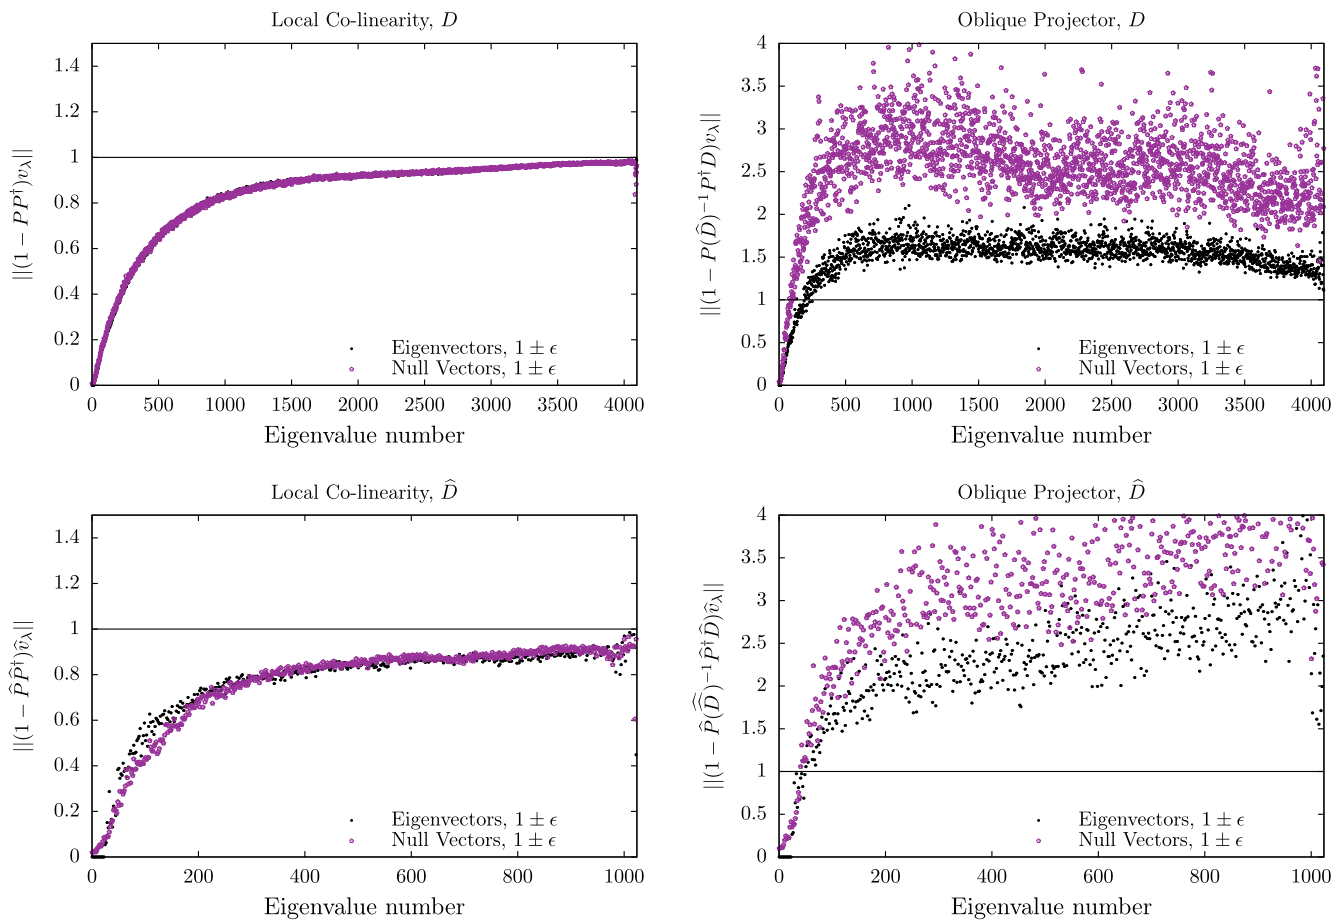


FIG. 8. On the left, a measurement of local collinearity: how well low near-null vectors can reconstruct eigenvectors. On the right, the contribution of eigenvector to the error after the coarse level solve. The top row of figures is the fine-to-first-coarse operator, and the bottom row is the next level from the coarse to doubly coarsened operator. The parameters of the MG coarsening are given in Table I. Both panels are sorted by increasing magnitude of the eigenvalues.

$$\|(1 - P(RDP)^{-1}RD)\vec{v}_\lambda\|_2. \quad (39)$$

The oblique projection of the coarse error ($\mathbb{P}e$) is zero: $\mathbb{P}_{ob}\mathbb{P} = [1 - P(RDP)^{-1}RD]PR = 0$. However, the oblique projection is not Hermitian so this does not imply the error in the orthogonal complement space ($e - PRe$) vanishes. This is illustrated on the right side of Fig. 8.

A magnitude less than or greater than one corresponds to a reduction or enhancement of the complementary error component, respectively. A successful coarse operator should strongly reduce the error component for low eigenmodes. In the context of MG, the enhancement from higher modes is addressed by the smoother. A larger enhancement requires a more expensive smoother; otherwise the solve stalls. In the top right panel of Fig. 8, we see that, for high modes, there is a large error enhancement. This is worse for a prolongator generated from near-null vectors than one generated from eigenvectors. In the lower right panel, we see the situation is even worse for a three-level algorithm. In all cases, an aggressive smoother is needed, increasingly so at coarser levels. This is why we

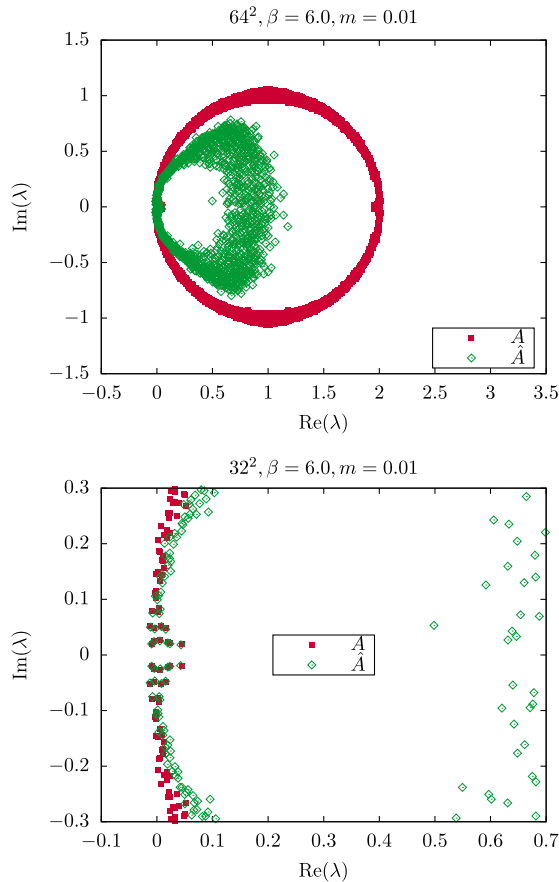


FIG. 9. On the top, a representative spectrum of the Kähler-Dirac preconditioned operator and its Galerkin projection from an interacting gauge field. On the bottom, a zoom in on the low spectrum. The parameters of the MG coarsening are given in Table I.

saw the MG algorithm fail. Now we turn to the same analysis for the Galerkin construction of the Kähler-Dirac preconditioned operator, which has in contrast minimal error enhancement, evident in Fig. 10.

C. Coarse Kähler-Dirac staggered operator: \hat{A}

We will coarsen the Kähler-Dirac preconditioned operator similar to how we coarsened the staggered operator, still using $\frac{1}{2}(1 \pm \epsilon(x))$ (unitarily rotated into the flavor basis) as a chiral projector on the near-null vectors. We will denote the method of coarsening the Kähler-Dirac preconditioned operator using $\frac{1}{2}(1 \pm \epsilon(x))$ as chiral projectors. Again, we will use $R = P^\dagger$. In Sec. V, we will discuss an *asymmetric* coarsening where $R \neq P^\dagger$. While not being of merit in two dimensions, it may be an interesting point of investigation in four-dimensional QCD. In this section we will consider the spectrum, collinearity, and oblique projector for a symmetric coarsening. Looking forward, in Sec. IV, we will demonstrate that symmetric coarsening produces a well behaved and robust recursive algorithm independent of the volume and the mass for physically relevant values of β .

As we described previously, the Wilson operator has a well behaved spectrum for MG as the high modes are well gapped along the real axis. This is also true for the Kähler-Dirac preconditioned operator in Fig. 9. As we discussed at the end of Sec. II B, the interacting spectrum is no longer a perfect circle in the complex plane. This does not undermine the qualitative benefits of the spectrum. Additionally, in the interacting case, a mass term still gaps the spectrum. In the bottom panel of Fig. 9, where we zoom in on the origin of the complex plane, we see that low modes are well preserved under our coarsening prescription, and there are no spurious modes near the complex origin. Eigenvalues of the coarse operator do not come in exact complex conjugate pairs, a consequence of using $\epsilon(x)$ as the chiral projector. This is inescapable because, in general, $\gamma_5^{L/R}$ does not define a good projector. The eigenvalues are *approximately* paired, which is consistent with a general preservation of the low spectrum. This may also be consistent with $\frac{1}{2}(1 \pm \epsilon(x))$ becoming equivalent to $\frac{1}{2}(1 \pm \gamma_5^{L/R})$, up to a unitary transformation, in the continuum limit, and as such preserving complex conjugate eigenpairs.

A careful study of the bottom panel of Fig. 9 shows that both the original operator and its coarsening feature eigenvalues with a negative real part, that is, lying in the left-half plane. We refer to these eigenvalues as *exceptional* eigenvalues, borrowing the language from Wilson-clover fermion literature [60]. The existence of modes in the left-half plane invalidate proofs which bound the convergence of Krylov solvers [61]. We will see in Sec. IV that, because error components in these exceptional modes are well solved by the coarse error correction, a recursive MG algorithm can successfully address this problem. As we will

$$64^2, \beta = 6.0, m = 0.01$$

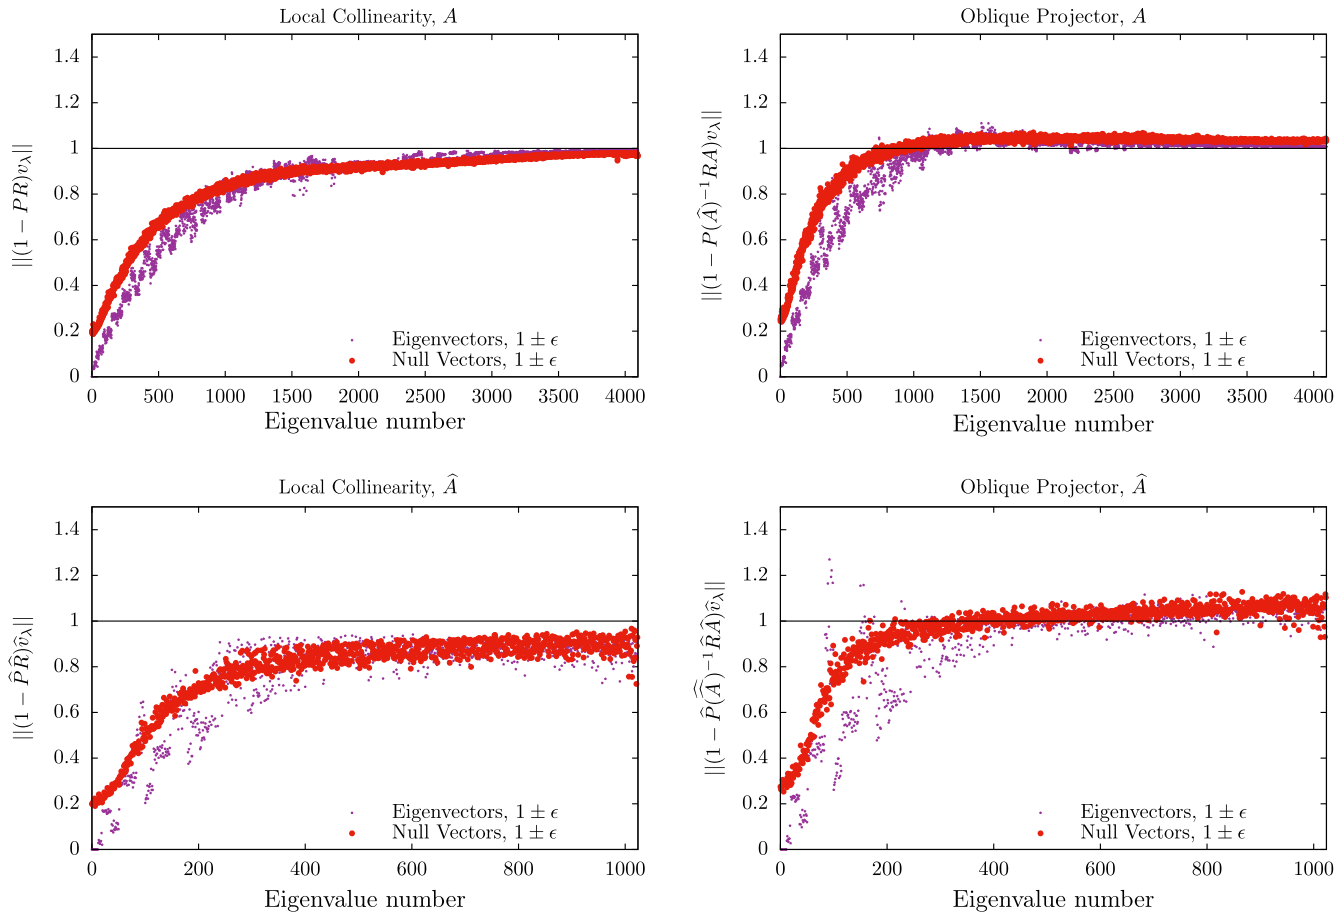


FIG. 10. On the left, a measurement of local collinearity, and on the right, the effect of the oblique projector on eigenvectors. The top row considers a representative Kähler-Dirac preconditioned operator; the bottom row considers its coarsening. Both panels are sorted by increasing magnitude of the eigenvalues. The parameters of the MG coarsening are given in Table I.

see in Sec. IV, this stabilizes the MG solve, independent of mass and volume, and is consistent with the success of MG for the Wilson operator beyond the critical mass.

Local collinearity versus the oblique projector. The overall failure of MG for the staggered operator stemmed from the large error enhancement to the high modes from the coarse correction. A predictor of success for MG on the Kähler-Dirac preconditioned operator would be a significant reduction of this enhancement. We would also still need to see strong local collinearity and a significant coarse error correction on low modes. In the left and right panels of Fig. 10, we consider the local collinearity and oblique projector, respectively, of the Kähler-Dirac preconditioned operator on a representative configuration. We explore using both near-null vectors and right eigenvectors to define the prolongator P and restrictor $R = P^\dagger$.

On the left, we see that local collinearity of low modes of the Kähler-Dirac operator is well maintained, similar to the original staggered operator. The benefit of coarsening the Kähler-Dirac preconditioned operator as opposed to the original staggered operator is most clearly noted by the

action of the oblique projector as displayed on the right panel of Fig. 10. The oblique projector reduces the error component on the fine level for roughly the lowest 15% of the spectrum. Above this threshold, the error component is enhanced, but only minimally.

IV. MG ALGORITHM PRESCRIPTION AND NUMERICAL RESULTS

The convergence rate of our new MG algorithm on the Kähler-Dirac preconditioned operator, illustrated in Fig. 11, is a dramatic improvement relative to the failed MG algorithm applied to the original staggered operator in Fig. 5. The *only* methodological difference is coarsening the Kähler-Dirac preconditioned operator instead of the original staggered operator. While starting from the Kähler-Dirac preconditioned operator was an essential step in formulating a MG algorithm, there are many other aspects that need to be tuned to make it robustly successful. The prescription for generating near-null vectors, as well as iterating on the coarser levels, is also important.

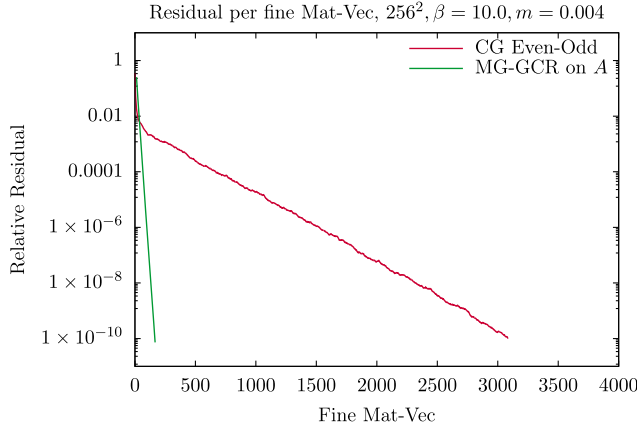


FIG. 11. The relative residual as a function of the number of fine operator applications for a representative $\beta = 10.0$, 256^2 configuration. The parameters of the MG solve are given in Table I.

With this in mind, we will begin by describing the complete prescription for defining our MG algorithm. Next, we will discuss the results of our tests of the MG algorithm. Importantly, we will see that as we scan in the quark mass, as shown in Fig. 12, our formulation has eliminated ill conditioning due to critical slowing down: unlike using CG on the even/odd preconditioned system, an MG solve takes a roughly constant number of outer iterations as the chiral limit is approached. Last, we will investigate the behavior of our MG algorithm as the continuum limit is approached, which is the most relevant regime for eliminating critical slowing down.

A. Definition of the MG algorithm

There are several parameters that need to be tuned when defining a MG algorithm. The parameters we choose are summarized in Table I. First, we consider a two-level




FIG. 12. The number of fine operations as a function of the bare mass for CG on the even/odd system, which exhibits critical slowing down, and for MG on the Kähler-Dirac preconditioned operator, where it is eliminated. Each data point is an average over 100 runs. Error bars are generally too small to be visible on the figures.

algorithm. We construct a right near-null vector ψ by relaxing on the homogeneous normal system $AA^\dagger\psi = 0$, using Gaussian distributed random vectors ψ_0 as the initial guess.² In practice, this is performed in multiple steps.

- (i) We convert the homogeneous system to the residual system $AA^\dagger\vec{e} = \vec{r} \equiv -AA^\dagger\psi_0$.
- (ii) We relax on the residual system using CG to a relative tolerance of 10^{-4} or a maximum number of 250 iterations.
- (iii) We reconstruct the near-null vector $\psi = \psi_0 + \vec{e}$, where \vec{e} is the result of relaxation.

This is performed n_{vec}^1 times, and then we *globally* orthonormalize the full set of near-null vectors. We subsequently chirally double the near-null vectors using $\frac{1}{2}(1 \pm \epsilon(x))$, and form the second-level operator $\hat{A} = P^\dagger A P$ from the block-orthonormalized chirally doubled null vectors. The coarse correction follows three steps. (1) Relax on the current residual, a process known as the presmoothen, (2) approximately solve the second-level system: $[RAP]Re = Rr$ (or, equivalently, approximately solve $\hat{A}\hat{e} = \hat{r}$), giving the prolonged error correction $e = P\hat{e}$, and (3) postsmooth on the error accumulated from steps (1) and (2). In step (2) we use a Krylov solver, and as such the MG preconditioner is not stationary. For this reason, we use the restarted generalized conjugate residual (GCR) [63] as a flexible outer solver, forming a K -cycle. We use a global MR for our presmoothen and postsmoothen. The specific details of these steps are given in Table I. In practice, we iterate on the even/odd preconditioned system on the fine level, with the prescription where we coarsen assuming the odd contributions are all zero, and we also ignore the odd contributions in the prolonged error. This technique proved successful for the Wilson operator [64].

A two-level algorithm does not fully eliminate critical slowing down, it just shifts it to the second level. We address this by generalizing to a *recursive* algorithm, where we perform a still coarser correction to the system in step (2) of the above description. We generate a third level, $\hat{\hat{A}}$, similar to how we generate the second level: we generate near-null vectors with $\hat{\hat{A}}\hat{\hat{e}}$, chirally double the near-null vectors using $\frac{1}{2}(1 \pm \sigma_3)$, and subsequently form a third level.

This clearly generalizes to still coarser levels. For our numerical experiments in Sec. IV, we only study a three-level algorithm. Unlike on the fine level, the Krylov solve we perform on the intermediate level is an iteration directly on $\hat{\hat{A}}$, as we found this was more stable in practice. We approximately solve the coarsest level via CG on the

²We remark that $A^\dagger A$ generally works just as well. We have also explored relaxing on A directly using BiCGstab and BiCGstab(l), $l = 6$ [62], which in practice works well at small volumes but degrades for larger volumes. The use of the normal operator may be why we can effectively capture exceptional eigenvalues.

TABLE I. The parameters we use for our K -cycle. For consistency, we use the same setup parameters throughout the procedures described in this paper. We remark that, for the setup, $A^\dagger A$ works just as well as AA^\dagger .

	Parameter	
Setup	Setup operator	Normal operator, AA^\dagger
	Setup solver	CG
	Max iterations	250
	Max residual tolerance per null vector	10^{-4}
	Number of null vectors, level 1 (n_{vec}^1)	8
	Size of aggregate block, level 1	4^2 (8^2 in the original lattice)
	Number of null vectors, level $l > 1$ (n_{vec}^l)	12
	Size of aggregate block, level $l > 1$	2^2
	Number of levels l_{max}	3
Solver, level 1	Operator	Schur preconditioned, $\mathbb{I} - A_{eo}A_{oe}$
	Restart length of GCR	32
	Relative residual tolerance	10^{-10}
	MR iterations for pre-,postsmooth	2
	MR relaxation parameter	0.85
Solver, level 2	Operator	\hat{A}
	Max iterations	16
	Restart length of GCR	8
	Relative residual tolerance	0.2
	MR iterations for pre-, postsmooth	2
	MR relaxation parameter	0.85
Solver, level 3	Operator	Normal operator, $\hat{A}\hat{A}^\dagger$
	Solver	CGNE
	Relative residual tolerance	0.2
	Maximum iterations	256

normal error (CGNE). Due to the exceptional eigenvalues which propagate to coarser levels, as noted in Fig. 9, numerical experiments with Krylov solvers acting on \hat{A} were in general not successful. This was either due to stability reasons [using BiCGstab(l) [62]] or due to cost (using GCR). We believe using the normal operator is of critical importance.

B. Results

A successful, recursive MG algorithm will shift critical slowing down to the coarsest level. In the context of the Schwinger model, and four-dimensional QCD, this means we want consistent convergence independent of mass and volume. We are also interested in the MG algorithm being successful in all physically interesting regimes. In the case of our target problems, this means we need to study the behavior with the bare coupling β . The continuum limit is taking $\beta \rightarrow \infty$ at constant physics, where the relevant region is $l_{M_x} > l_\sigma$. When β is too small, close to the cutoff scale, we are no longer studying relevant physics. A breakdown of MG for very small β is acceptable. The values of β studied, 3.0, 6.0, and 10.0, correspond to $l_\sigma \approx 2.4, 3.5,$ and $4.5,$ respectively. The lowest value of β is becoming rather unphysical.

1. Elimination of critical slowing down: Fine level

The indication of a successful two-level algorithm is the elimination of critical slowing down for the fine operator A , that is, constant iterations with respect to the mass and volume per each β . In Fig. 13, we present the number of applications of the fine operator A between the GCR algorithm and the MG preconditioner, which is proportional to the number of iterations for the outer GCR solve. In the left panel we consider the case of fixed physical $\beta = 6.0$ at varying volume. The number of A applications is roughly constant, independent of the volume and mass in the chiral limit. In the right panel, we consider our largest volume, 256^2 , fixed for three different values of β . We see that at $\beta = 10.0$ and 6.0 , critical slowing down has been essentially eliminated as a function of mass. At $\beta = 3.0$, where we are probing somewhat cutoff scale physics, the number of iterations appears not to be diverging with power law behavior, and as such critical slowing down has still been eliminated.

2. Elimination of critical slowing down: Intermediate level

A successful recursive algorithm eliminates critical slowing down at each level. Thus, we consider the average number of iterations for the second level coarse correction, averaged over each outer iteration. In the left panel of Fig. 14,

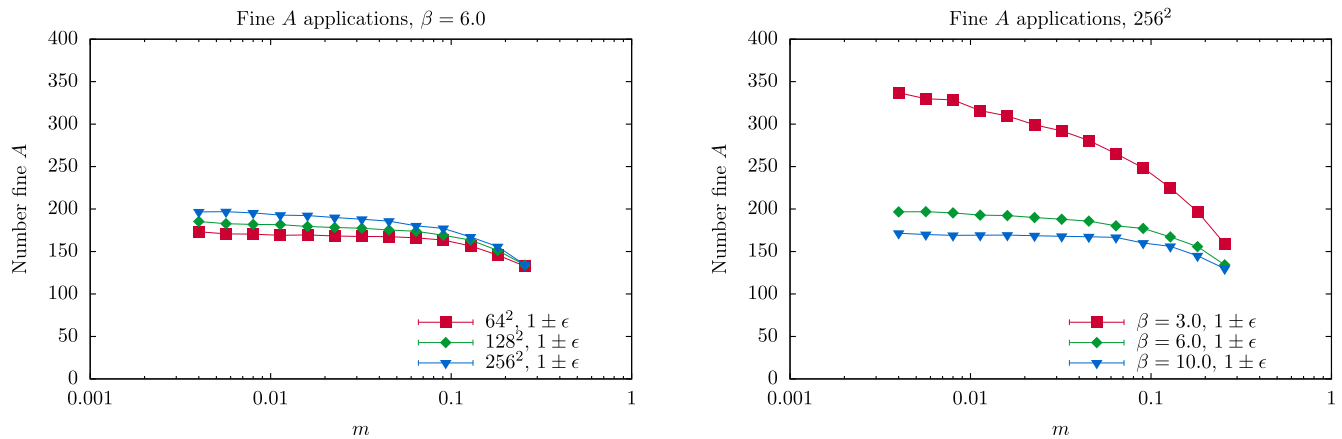


FIG. 13. The number of applications of the fine operator A per MG-preconditioned solve as a function of mass. On the left, we consider fixed β , and on the right, fixed volume. Each data point is an average over 100 runs. Error bars are generally too small to be visible on the figures.

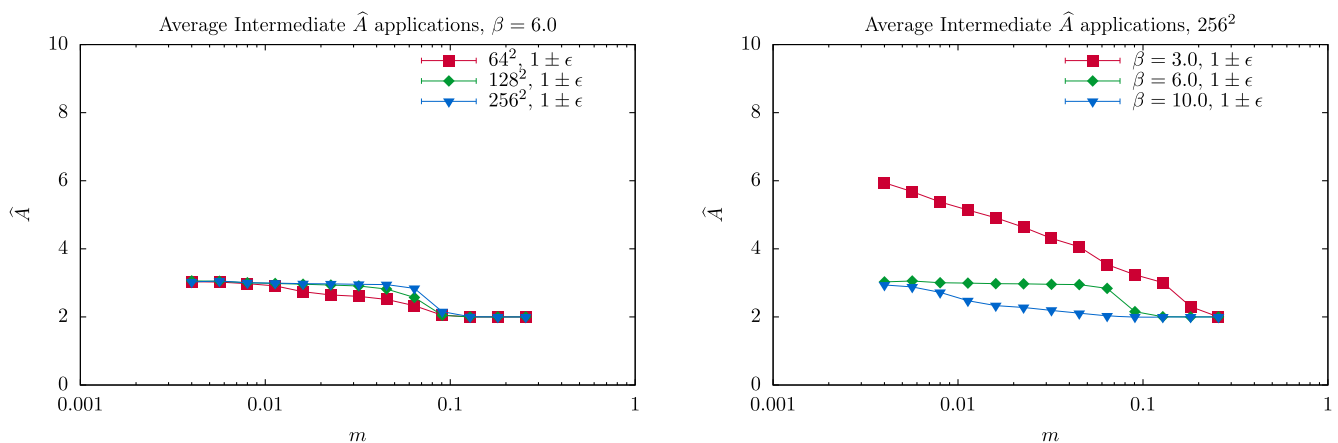


FIG. 14. The average number of iterations for the inner Krylov solve of \hat{A} as a function of mass. On the left, we consider fixed β , and on the right, fixed volume. Each data point is an average over 100 runs. Error bars are generally too small to be visible on the figures.

we consider the average number of intermediate iterations at fixed $\beta = 6.0$ at varying volume. In the chiral limit, the average number of iterations is essentially fixed at 3 independent of volume. In the right panel we consider a fixed volume of 256^2 for three different values of β . As with the fine level, we see an elimination of critical slowing down for each value of β .

We remark that in a highly optimized and tuned implementation, it is important that we use a K -cycle at the second level. In such an implementation, the maximum number of iterations on the coarsest level may be capped to some reasonable amount. This would cause the number of iterations on the intermediate level to increase. Since in a K -cycle the second level is solved to a fixed *residual* as opposed to a fixed number of iterations, the number of iterations at the *finest* level remains stable.

3. Critical slowing down: Coarsest level

The previous two paragraphs demonstrate an elimination of critical slowing down from finer levels. Thus, there

should be critical slowing down on the efficiently solvable coarsest level. In Fig. 15 we consider the average number of iterations for the coarsest solve via CGNE. In contrast to the previous two figures, these plots are on a log-log scale instead of a log-linear scale. In the left and right panels, we consider constant $\beta = 6.0$ and a constant volume of 256^2 , respectively. The number of iterations is divergent with power law behavior.³ Critical slowing down has been shifted to the coarsest level.

4. Comparison with a direct solve

In looking at the outermost level, the intermediate level, and the coarsest level in a three-level solve, we see that we have formulated an MG algorithm which shifts critical slowing down to the coarsest level. Furthermore, the solve is stable: in Fig. 11, we saw that an MG-GCR solve

³It will not exactly diverge for m extremely small due to finite volume cutoff effects.

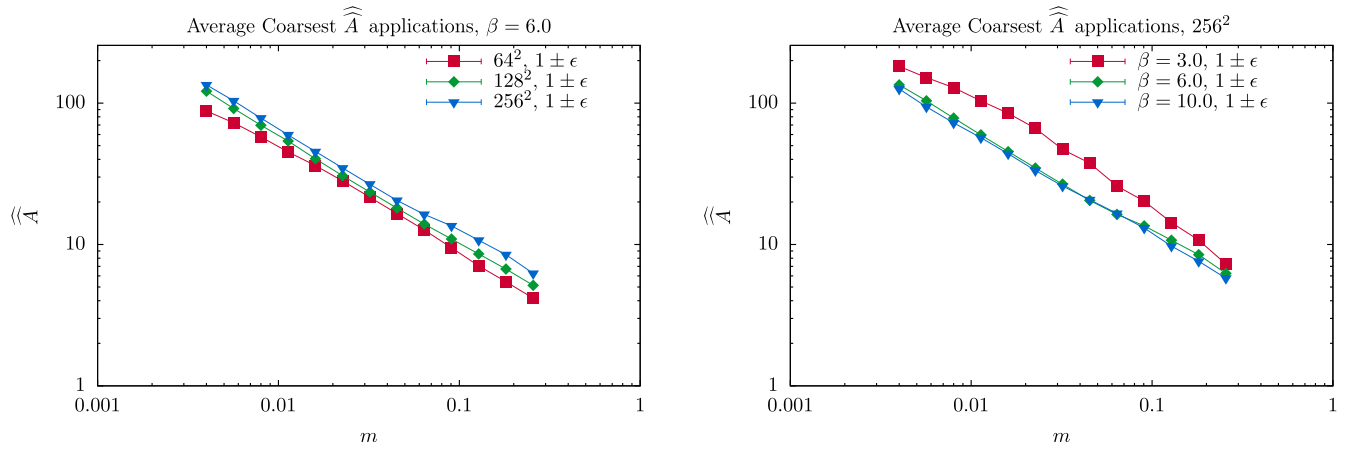


FIG. 15. The average number of iterations of CGNE on the coarsest level for, on the left, fixed β , and on the right, fixed volume. Note that this is a log-log plot. Each data point is an average over 100 runs. Error bars are generally too small to be visible on the figures.

converges smoothly at our most chiral point for $\beta = 10$. There are large reductions in the relative residual on each iteration. On the other hand, the traditional solve with CG on the even/odd operator, despite converging successfully, converges very slowly, an indication of critical slowing down.

This behavior persists independent of mass. In Fig. 12, we trace the number of iterations away from the chiral limit, seeing that it is roughly constant. Critical slowing down has been eliminated. On the other hand, the number of iterations for a solve with the even/odd operator diverges with mass with power law behavior. This is exactly the critical behavior that has been shifted to the coarsest level in Fig. 15. The benefit of our MG algorithm is drastic.

C. Continuum limit

It should be emphasized that our fixed prescription is effective in the most relevant regime: toward the continuum, where the lattice spacing vanishes relative to fixed physics, and in the chiral limit, where l_{M_x} diverges relative to l_σ .

For the two-dimensional Schwinger model, taking the continuum limit at constant physics corresponds to simultaneously doubling the length scale of the fine volume, halving the mass, and quadrupling β . In Table II, we

consider the use of MG while taking the continuum limit from two base configurations. First, we consider a base configuration of 64^2 at $m = 0.01$ and $\beta = 3.0$, where we have discussed earlier that an MG algorithm is successful. On two successive refinements toward the continuum limit, we see that there is a reduction in the number of outer applications of A and in the average number of iterations on the intermediate level. In tandem, the average number of iterations on the coarsest level increases: there is more critical slowing down to shift to the coarsest level, which is to be expected, toward the continuum limit. Our MG algorithm performs better as the continuum limit is taken. Next, we consider a base configuration of 64^2 at $m = 0.004$ and $\beta = 0.75$, an unphysically coarse configuration. In this case, an MG algorithm fails to converge. Again, on progressive refinements, the MG algorithm becomes convergent and becomes better behaved as the continuum limit is taken.

V. PRESERVING COMPLEX CONJUGATE PAIRS

A possible, if not necessary, generalization for four-dimensional QCD or other staggered fermion problems could be the exact preservation of complex conjugate pairs upon coarsening. Indeed it is possible to develop a

TABLE II. The effect of taking the lattice spacing to zero at constant physical box size and mass gap for two sequences of successive refinement. Cases where the maximum number of iterations is sometimes hit on the coarsest level are denoted with a dagger. All quantities are averaged over 100 runs.

L	m	β	Fine mat-vec	Intermediate average iterations	Coarsest average iterations
64	0.01	3.0	228.6(1.2)	3.45(4)	62.0(0.5)
128	0.005	12.0	159.4(0.4)	2.60(2)	95.7(0.6)
256	0.0025	48.0	147.4(0.5)	2.09(1)	205.0(0.6)
64	0.004	0.75	Max
128	0.002	3.0	290.5(1.8)	5.08(8)	206.2(1.5) [†]
256	0.001	12.0	189.9(0.4)	4.61(2)	249.6(0.2) [†]

prolongator P and a restrictor $R \neq P^\dagger$, abandoning chiral doubling with projectors, which preserves complex conjugate eigenpairs after coarsening the Kähler-Dirac preconditioned operator, or any operator satisfying $\gamma_5^{L/R}$ Hermiticity with $\gamma_5^L \gamma_5^R = \mathbb{I}$. The resulting formalism gives what we will call an *asymmetric* coarsening with $\sigma_1^{L/R}$ Hermiticity on the coarse level.

We consider a set of left and right vectors, $\langle \tilde{\psi}_i |$ and $|\psi_i\rangle$, respectively, which can generally be arbitrary and unequal. We perform a chiral doubling which gives

$$\tilde{R} = \begin{bmatrix} \langle \tilde{\psi}_i | \\ \langle \psi_i | \gamma_5^L \end{bmatrix}, \quad \tilde{P} = \begin{bmatrix} |\psi_i\rangle & \gamma_5^R |\tilde{\psi}_i\rangle \end{bmatrix}. \quad (40)$$

These prolongators and restrictors obey $\tilde{P}\sigma_1 = \gamma_5^R \tilde{R}^\dagger$ and $\sigma_1 \tilde{R} = \tilde{P}^\dagger \gamma_5^L$. This is sufficient to prove $\tilde{R}\tilde{A}\tilde{P}$ is σ_1 Hermitian. The next step is to *block bi-orthonormalize* R and P , enforcing $RP = \mathbb{I}$, by-products of which give us $\sigma_1^{L/R}$.

As a clarifying tangent, we will consider the case $\gamma_5^L = \gamma_5^R$ and $|\psi_i\rangle = |\tilde{\psi}_i\rangle$, that is, $\tilde{R} = \tilde{P}^\dagger$. This is true, for example, for the Kähler-Dirac preconditioned operator in the free field limit, or when considering the Wilson operator in general. The critical observation in this case is to recall that the process of (block) orthonormalization via a Gram-Schmidt is equivalent to a thin-QR decomposition. We define the block-dense matrix \mathbb{M} of block dimension (coarse dof) \times (coarse dof) as

$$\tilde{P}^\dagger \tilde{P} = \mathbb{M} = \Sigma^\dagger \Sigma, \quad (41)$$

where in the last step we have performed a Cholesky decomposition. We can rearrange Eq. (41) as

$$(\tilde{P}\Sigma^{-1})^\dagger (\tilde{P}\Sigma^{-1}) \equiv P^\dagger P = \mathbb{I}. \quad (42)$$

By definition, $P \equiv \tilde{P}\Sigma^{-1}$ is block orthonormal. With the definition $\sigma_1^\Sigma \equiv \Sigma\sigma_1\Sigma^{-1}$, we have $\gamma_5 P = P\sigma_1^\Sigma$, and $P^\dagger A P$ is σ_1^Σ Hermitian.

We return to the (block) bi-orthonormalization of R and P . The above procedure generalizes to a ‘‘thin-LU’’ decomposition. Equation (41) generalizes to

$$\tilde{R}\tilde{P} = \mathbb{M} = LU, \quad (43)$$

where in the last step we have performed an LU decomposition. We can rearrange Eq. (43) as

$$(L^{-1}\tilde{R})(\tilde{P}U^{-1}) \equiv RP = \mathbb{I}. \quad (44)$$

R and P are block bi-orthonormal. We can show $\hat{A} \equiv RAP$ admits a $\sigma_1^{L/R}$ Hermiticity condition via defining

$$\sigma_1^L = U^{-\dagger} \sigma_1 L, \quad \sigma_1^R = U \sigma_1 L^{-\dagger}, \quad (45)$$

and noting

$$\sigma_1^L R = P^\dagger \gamma_5^L, \quad P \sigma_1^R = \gamma_5^R R^\dagger. \quad (46)$$

The pair $\sigma_1^{L/R}$ obeys $\sigma_1^L \sigma_1^R = \mathbb{I}$, as can be verified by explicit calculation, requiring the critical and subtle observation that $\tilde{R}\tilde{P}$ is σ_1 Hermitian itself.

We emphasize that this construction is fully generic, whether or not $|\psi_i\rangle = |\tilde{\psi}_i\rangle$. We defer a discussion of numerical experiments with preserving complex conjugate eigenpairs to Appendix A. Our deference to an appendix reflects our observations that, in two dimensions, (recursively) preserving eigenpairing actually leads to a less effective, and sometimes unstable algorithm. This method, or a further development thereof, may bear some fruit in four dimensions.

We make the additional remark that we can now make the algorithmic choice to right-block-Jacobi precondition \hat{A} , analogous to the transformation we made to the staggered operator in the Kähler-Dirac form in the first place, and we continue to preserve complex conjugate eigenpairs if we coarsen again. Let us denote $\hat{A} = \mathcal{B} + \mathcal{C}$, where \mathcal{B} is the block-local contribution. The resulting right-block-preconditioned operator $\hat{A}\mathcal{B}^{-1}$ obeys a $\sigma_1^{rbj.L/R}$ Hermiticity condition with $\sigma_1^{rbj.L} = \sigma_1^L \mathcal{B}^{-1}$ and $\sigma_1^{rbj.R} = \mathcal{B}\sigma_1^R$. This recursive right-block-Jacobi preconditioning did not lead to an effective algorithm in two dimensions.

Exact preservation of eigenvectors. In the case of, for example, the Wilson operator, chiral doubling with $\frac{1}{2}(1 \pm \gamma_5)$ preserves complex conjugate eigenpairs. We can choose the vectors $|\psi_i\rangle \equiv |\tilde{\psi}_i\rangle$ to be right eigenvectors $|\lambda_i^{+,R}\rangle$ with eigenvalues λ_i^+ , where the $+$ denotes that the eigenvalue has a positive real part.

The coarse operator $P^\dagger D_W P$ exactly preserves the eigenvalue λ_i^+ , and $|\lambda_i^{+,R}\rangle$ is exactly preserved on the coarse subspace, that is, $PP^\dagger |\lambda_i^{+,R}\rangle = |\lambda_i^{+,R}\rangle$. However, even though chiral doubling guarantees the eigenvalue λ_i^- is also preserved by the coarse operator, it is *not* because $|\lambda_i^{-,R}\rangle$ is exactly preserved by the coarse subspace, that is, $PP^\dagger |\lambda_i^{-,R}\rangle \neq |\lambda_i^{-,R}\rangle$.

We can use *asymmetric* coarsening to preserve $|\lambda_i^{-,R}\rangle$. We can choose $|\psi_i\rangle = |\lambda_i^{+,R}\rangle$ and $\langle \tilde{\psi}_i | = \langle \lambda_i^{+,L} |$, then chirally double using Eq. (40), and subsequently block bi-orthonormalize the P and R . This operator preserves the eigenvalues λ_i^\pm , and additionally $PR |\lambda_i^{\pm,R}\rangle = |\lambda_i^{\pm,R}\rangle$ and $\langle \lambda_i^{\pm,L} | PR = \langle \lambda_i^{\pm,L} |$.

VI. CONCLUSION

The first successful MG algorithm in LQCD was constructed for the Wilson discretization of the Dirac

operator nearly a decade ago [6,7]. This advance relied on, at the time, the novel approach in LQCD to adaptively discover the near-null space and geometrically project onto coarse lattices. Remarkably, with the exception of the similar twisted-mass discretization, the basic method has not easily been generalized to two important methods: staggered and domain wall fermions, each of which feature improved chiral symmetry. A more fundamental understanding of MG methods in LQCD is clearly lacking. Here, we have taken a step toward this. For the staggered operator, we identified the spectral feature that was responsible for the failure of a straightforward generalization of Wilson MG and have overcome this problem by preconditioning by the Kähler-Dirac (spin-flavor) block structure. We demonstrate that this has a dramatic effect on the spectrum: in the singular, zero mass limit, the pure imaginary spectrum of the anti-Hermitian operator maps to a unitary circle of the form seen in the overlap operator.

The success of the resultant MG algorithm for this Kähler-Dirac preconditioned operator has been demonstrated numerically for the two-dimensional Schwinger model. Both the theoretical framework and the phenomenological features naturally generalize to the case of four-dimensional QCD. On this basis, we are optimistic that our staggered multigrid algorithm will have similar success in this application. Numerical tests for this conjecture are underway by extending the high performance MG framework of the QUDA library to coarsen staggered-like operators. These tests will be made on the largest available lattices to explore the scaling of the algorithm over a range similar to the two-dimensional tests presented here.

We have also made an effort to explore a range of projection methods that are capable of exactly preserving the complex conjugate pairs of eigenvalues present in the Kähler-Dirac preconditioned operator. We hope our

emphasis on spectral analysis and transformations will provide some flexibility in adapting our algorithm not only to four-dimensional QCD but also to similar Dirac discretizations found in BSM theories [2], supersymmetric Yang-Mills theory [20], and quantum critical behavior in condensed matter [19].

ACKNOWLEDGMENTS

We acknowledge fruitful conversations with Peter Boyle, Carleton DeTar, Arjun Gambhir, Martin Lüscher, James Osborn, Eloy Romero, Andreas Stathopoulos, and especially with Matthias Rottman and Andreas Frommer who have also explored the staggered problem and gave the insight that it is possible to exactly preserve complex conjugate eigenpairs for the coarsened operator. We also thank the Higgs Centre for Theoretical Physics for its hospitality at the 2016 QCDNA Workshop. R. C. B. and E. S. W. acknowledge support by DOE Grant No. DE-SC0015845 and by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. The work by A. S. was done for Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

APPENDIX: STUDIES OF PRESERVING COMPLEX CONJUGATE EIGENPAIRS

In Sec. V, we developed a formalism to exactly preserve complex conjugate eigenpairs for a coarsened Kähler-Dirac preconditioned operator. This used an *asymmetric* coarsening which gave a $\sigma_1^{L/R}$ on the coarse level. This formulation is largely successful; however, it can suffer from anomalously large real eigenvalues in the negative

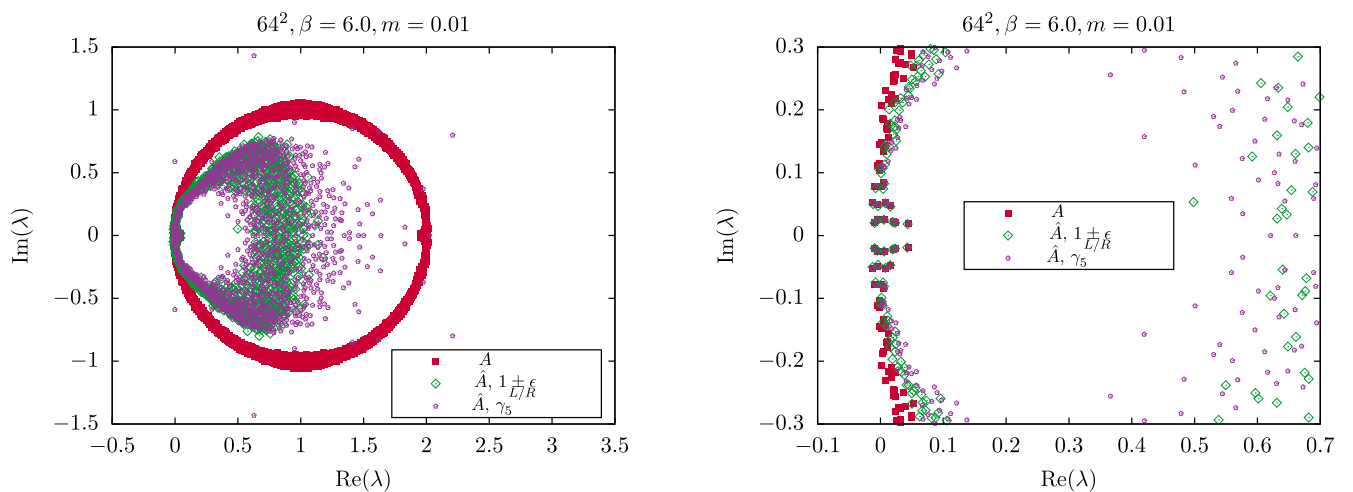


FIG. 16. On the left, the spectrum of the Kähler-Dirac preconditioned operator from an interacting gauge field, overlaid with the spectrum of the operator coarsened via naïve chiral doubling and complex eigenpair doubling. On the right, a zoom in on the low portion spectrum which shows eigenvalues in the left-half plane and further emphasized complex eigenpairing.

$$64^2, \beta = 6.0, m = 0.01$$

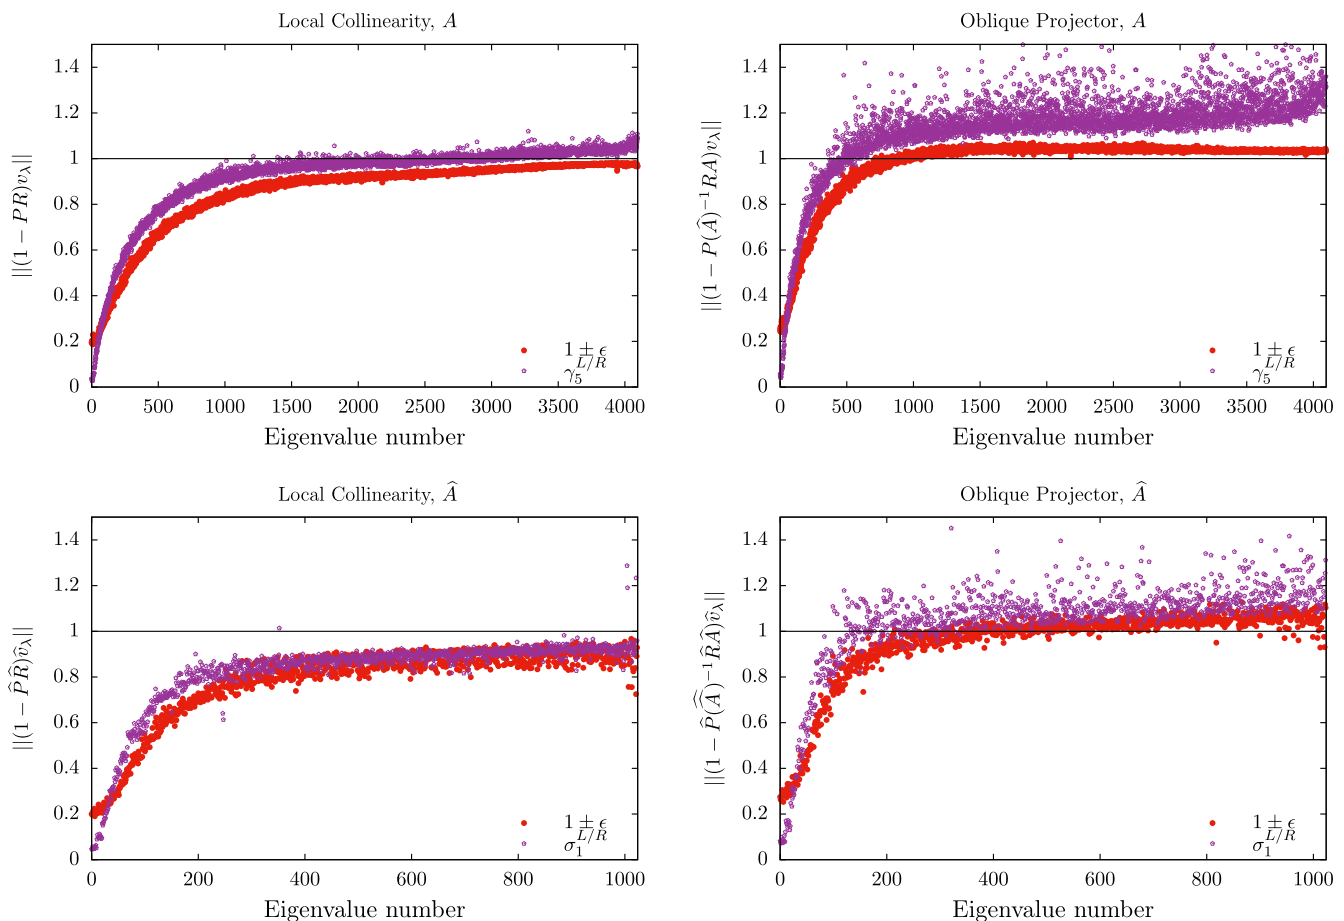


FIG. 17. On the left, a measurement of local collinearity: how well low near-null vectors can reconstruct higher eigenvectors. On the right, the effect of the oblique projector on eigenvectors. The top row considers a representative Kähler-Dirac preconditioned operator; the bottom row considers its coarsening. Both panels are sorted by increasing magnitude of the eigenvalues.

half plane, destabilizing the MG preconditioned solve, in cases where the *symmetric* coarsening proceeded without issues. If these stability issues can be addressed, it may lead to a better algorithm in two dimensions and four dimensions. As appropriate, this will be the topic of a future publication.

This appendix will follow the structure of Sec. III C, where we study the spectrum, local collinearity, and oblique projector of the asymmetrically coarsened operator *in the case where a recursive algorithm is successful*. We will then scan the iteration counts as a function of mass, similar to in Sec. IV B, and identify cases where the algorithm breaks down. Last, we will investigate one of these cases.

In Sec. III C we considered a representative spectrum of the Kähler-Dirac preconditioned operator and a symmetric coarsening. In the case of *asymmetric* coarsening, we again expect the low modes to be preserved well, but additionally come in complex conjugate pairs. This is exactly the case in

Fig. 16, where we overlay the spectrum of the asymmetric coarse operator. We also see a “feature” of $\sigma_1^{L/R}$ Hermiticity: there are pairs of purely real eigenvalues.

In the case of the Wilson or overlap operator, pairs of purely real eigenvalues have a physical interpretation. The smaller real eigenvalue corresponds to a physical chiral mode via the lattice index theorem [65], which thus needs to be well captured by an MG algorithm. The paired large real eigenvalue is merely a quirk of being on a finite lattice, and thus lives as an isolated large eigenvalue near the cutoff. On the other hand, the pairs of real eigenvalues for the coarsened Kähler-Dirac operator do not have an obvious physical intuition, just as the naïve staggered fermion operator does not trivially correspond to an index theorem [43]. These purely real eigenvalues are a symptom of unstable solves at larger volumes.

Returning to stable solves, we consider the local collinearity and the oblique projector under an asymmetric coarsening. These are overlaid on the data for a symmetric

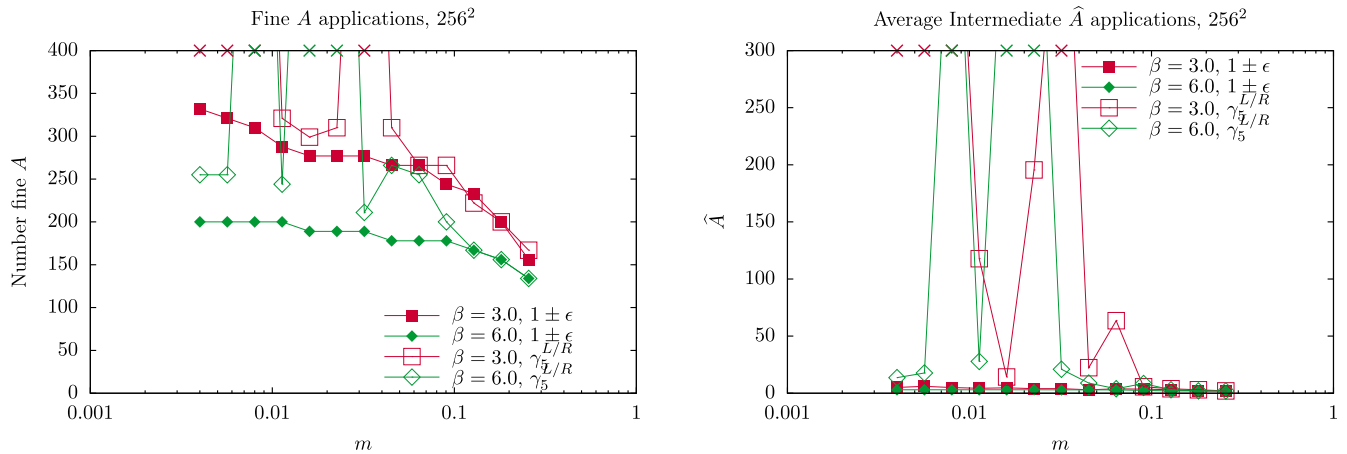


FIG. 18. On the left, the number of outer iterations, and on the right, the average number of iterations in the K -cycle on the intermediate level, as a function of mass at fixed β . Values marked with “x” indicate a failure to converge when an asymmetric coarsening is used. All data points are from a single configuration per β but are representative of a more general behavior.




FIG. 19. The spectrum of a representative coarse operator from a symmetric coarsening with the asymmetric coarsening overlaid. There are large purely real eigenvalues. Not included is a large negative eigenvalue $\lambda = -26.75$. Computing the fine operator spectrum was prohibitively expensive.

coarsening in Fig. 17. An asymmetric coarsening is roughly comparable to quality of a symmetric coarsening, indicative of a successful MG algorithm.⁴

As a next task, we consider MG preconditioned solves with the asymmetric coarsened operator. We will only present a subset of the cases considered in Sec. IV B and instead focus on the cases where the solve is unstable: large volumes. The number of fine operator applications and average intermediate applications are presented in Fig. 18.

⁴In general, the local collinearity is not bounded by 1 when $R \neq P^\dagger$. This is because $(1 - PR)$ is *not* a normal operator. Thus, for a normalized vector v , $v^\dagger(1 - R^\dagger P^\dagger)(1 - PR)v$ is not bounded by 1. This can be realized by the bi-orthonormal basis $p_1 = (1/2, 1/2, 1/2, 1/2)$, $p_2 = (1/2, 1/2, 1/2, -3/2)$, $r_1 = (1/2, 1/2, 1/2, 1/2)$, $r_2 = (1/2, -1/2, 1/2, -1/2)$.

In the cases where a data point is marked by a “x”, the solve failed. The failures are largely confined to smaller masses, but not with a discernible pattern; indeed, for $\beta = 6.0$, the lowest masses had stable solves.

We present the spectrum of the asymmetric coarsened operator, where an MG solve with an asymmetric coarsened operator fails, in the left panel of Fig. 19, where we see there are now large, real eigenvalues far in the right plane and also in the left plane. There is also a large negative real eigenvalue at approximately -26.75 . These pathological real eigenvalues are not part of the low subspace and are therefore not well captured by our MG algorithm. However, in the right panel, we see that the low spectrum is still well behaved. It is a point of future research to see if these anomalously large, real eigenvalues can be addressed.

- [1] K. G. Wilson, *Phys. Rev. D* **10**, 2445 (1974).
- [2] T. Blum, R. S. Van de Water, D. Holmgren, R. Brower, S. Catterall, N. Christ, A. Kronfeld, J. Kuti, P. Mackenzie, E. T. Neil, S. R. Sharpe, and R. Sugar, [arXiv:1310.6087](https://arxiv.org/abs/1310.6087).
- [3] K. G. Wilson, *Rev. Mod. Phys.* **47**, 773 (1975).
- [4] A. Hulsebos, J. Smit, and J. C. Vink, *Nucl. Phys. B, Proc. Suppl.* **20**, 94 (1991).
- [5] R. C. Brower, R. G. Edwards, C. Rebbi, and E. Vicari, *Nucl. Phys.* **B366**, 689 (1991).
- [6] J. Brannick, R. C. Brower, M. A. Clark, J. C. Osborn, and C. Rebbi, *Phys. Rev. Lett.* **100**, 041601 (2008).
- [7] R. Babich, J. Brannick, R. C. Brower, M. A. Clark, T. A. Manteuffel, S. F. McCormick, J. C. Osborn, and C. Rebbi, *Phys. Rev. Lett.* **105**, 201602 (2010).
- [8] R. Frezzotti, P. A. Grassi, S. Sint, and P. Weisz (Alpha Collaboration), *J. High Energy Phys.* **08** (2001) 058.
- [9] A. Frommer, K. Kahl, S. Krieg, B. Leder, and M. Rottmann, *SIAM J. Sci. Comput.* **36**, A1581 (2014).
- [10] D. B. Kaplan, *Phys. Lett. B* **288**, 342 (1992).
- [11] H. Neuberger, *Phys. Lett. B* **417**, 141 (1998).
- [12] J. Kogut and L. Susskind, *Phys. Rev. D* **11**, 395 (1975).
- [13] S. D. Cohen, R. C. Brower, M. A. Clark, and J. C. Osborn, *Proc. Sci. LATTICE2011* (**2011**) 030, [arXiv:1205.2933](https://arxiv.org/abs/1205.2933).
- [14] P. A. Boyle, [arXiv:1402.2585](https://arxiv.org/abs/1402.2585).
- [15] J. Brannick, A. Frommer, K. Kahl, B. Leder, M. Rottmann, and A. Strebel, *Numer. Math.* **132**, 463 (2016).
- [16] P. Becher and H. Joos, *Z. Phys. C* **15**, 343 (1982).
- [17] G. T. Bodwin and E. V. Kovács, *Phys. Rev. D* **38**, 1206 (1988).
- [18] N. Kawamoto and J. Smit, *Nucl. Phys.* **B192**, 100 (1981).
- [19] R. Brower, C. Rebbi, and D. Schaich, *Proc. Sci. LATTICE2011* (**2011**) 056, [arXiv:1204.5424](https://arxiv.org/abs/1204.5424).
- [20] S. Catterall, D. Schaich, P. H. Damgaard, T. DeGrand, and J. Giedt, *Phys. Rev. D* **90**, 065013 (2014).
- [21] T. Appelquist, G. T. Fleming, M. F. Lin, E. T. Neil, and D. A. Schaich, *Phys. Rev. D* **84**, 054501 (2011).
- [22] A. Hasenfratz, *Phys. Rev. Lett.* **108**, 061601 (2012).
- [23] Y. Aoki, T. Aoyama, M. Kurachi, T. Maskawa, K.-i. Nagai, H. Ohki, A. Shibata, K. Yamawaki, and T. Yamazaki, *Phys. Rev. D* **86**, 054506 (2012).
- [24] T. DeGrand, *Phys. Rev. D* **84**, 116901 (2011).
- [25] A. Cheng, A. Hasenfratz, Y. Liu, G. Petropoulos, and D. Schaich, *Phys. Rev. D* **90**, 014509 (2014).
- [26] A. Cheng, A. Hasenfratz, Y. Liu, G. Petropoulos, and D. Schaich, *J. High Energy Phys.* **05** (2014) 137.
- [27] M. Lombardo, K. Miura, T. J. N. da Silva, and E. Pallante, *J. High Energy Phys.* **12** (2014) 183.
- [28] Y. Aoki, T. Aoyama, M. Kurachi, T. Maskawa, K.-i. Nagai, H. Ohki, E. Rinaldi, A. Shibata, K. Yamawaki, and T. Yamazaki, *Phys. Rev. Lett.* **111**, 162001 (2013).
- [29] Y. Aoki, T. Aoyama, M. Kurachi, T. Maskawa, K. Miura, K. I. Nagai, H. Ohki, E. Rinaldi, A. Shibata, K. Yamawaki, and T. Yamazaki, *Phys. Rev. D* **89**, 111502 (2014).
- [30] v. E. Kähler, in *Forme differenziali e loro integrali*, edited by B. Segre (Springer, Berlin, 2011), pp. 160–258.
- [31] A. Stathopoulos and K. Orginos, *SIAM J. Sci. Comput.* **32**, 439 (2010).
- [32] M. R. Hestenes and E. Stiefel, *J. Res. Natl. Bur. Stand.* **49**, 409 (1952).
- [33] A. A. Nikishin and A. Y. Yeremin, *SIAM J. Matrix Anal. Appl.* **16**, 1135 (1995).
- [34] A. A. Dubrulle, *Electron. Trans. Numer. Anal.* **12**, 216 (2001).
- [35] M. A. Clark, A. Strelchenko, A. Vaquero, M. Wagner, and E. Weinberg, [arXiv:1710.09745](https://arxiv.org/abs/1710.09745).
- [36] E. S. Weinberg, R. C. Brower, M. Clark, and A. Strelchenko, *Proc. Sci., LATTICE2016* (**2017**) 273.
- [37] M. W. Benson, M.Sc. thesis, Lakehead University, 1973.
- [38] E. Chow, *Int. J. High Perform. Comput. Appl.* **15**, 56 (2001).
- [39] H. D. Sterck, U. M. Yang, and J. J. Heys, *SIAM J. Matrix Anal. Appl.* **27**, 1019 (2006).
- [40] E. Treister and I. Yavneh, *SIAM J. Sci. Comput.* **37**, A30 (2015).
- [41] J. Schwinger, *Phys. Rev.* **128**, 2425 (1962).
- [42] A. V. Smilga, *Phys. Rev. D* **55**, R443 (1997).
- [43] D. H. Adams, *Phys. Rev. Lett.* **104**, 141602 (2010).
- [44] M. A. Clark, R. Babich, K. Barros, R. C. Brower, and C. Rebbi, *Comput. Phys. Commun.* **181**, 1517 (2010).
- [45] R. Babich, M. A. Clark, B. Joo, G. Shi, R. C. Brower, and S. Gottlieb, [arXiv:1109.2935](https://arxiv.org/abs/1109.2935).
- [46] M. A. Clark, B. Joó, A. Strelchenko, M. Cheng, A. Gambhir, and R. Brower, in *Proceedings for the Supercomputing 2016 (SC'16) conference*, DOI: 10.1109/SC.2016.67.
- [47] H. B. Nielsen and M. Ninomiya, *Phys. Lett.* **105B**, 219 (1981).
- [48] S. Naik, *Nucl. Phys.* **B316**, 238 (1989).
- [49] F. Gliozzi, *Nucl. Phys.* **B204**, 419 (1982).
- [50] H. Kluberg-Stern, A. Morel, O. Napoloy, and B. Petersson, *Nucl. Phys.* **B220**, 447 (1983).
- [51] A. Duncan, R. Roskies, and H. Vaidya, *Phys. Lett.* **114B**, 439 (1982).
- [52] A. Bazavov, D. Toussaint, C. Bernard, J. Laiho, C. DeTar, L. Levkova, M. B. Oktay, S. Gottlieb, U. M. Heller, J. E. Hetrick, P. B. Mackenzie, R. Sugar, and R. S. Van de Water, *Rev. Mod. Phys.* **82**, 1349 (2010).
- [53] J. M. Rabin, *Nucl. Phys.* **B201**, 315 (1982).
- [54] T. Banks, Y. Dothan, and D. Horn, *Phys. Lett.* **117B**, 413 (1982).
- [55] K. Orginos, D. Toussaint, and R. L. Sugar, *Phys. Rev. D* **60**, 054503 (1999).
- [56] E. Follana, Q. Mason, C. Davies, K. Hornbostel, G. P. Lepage, J. Shigemitsu, H. Trotter, and K. Wong, *Phys. Rev. D* **75**, 054502 (2007).
- [57] D. Day and M. A. Heroux, *SIAM J. Sci. Comput.* **23**, 480 (2001).
- [58] M. Benzi and D. Bertaccini, *IMA J. Numer. Anal.* **28**, 598 (2007).
- [59] M. A. Clark, C. Jung, and C. Lehner, *EPJ Web Conf.* **175**, 14023 (2018).
- [60] T. A. DeGrand, A. Hasenfratz, and T. G. Kovacs, *Nucl. Phys. B, Proc. Suppl.* **73**, 506 (1999).
- [61] Y. Saad and M. H. Schultz, *SIAM J. Sci. Stat. Comput.* **7**, 856 (1986).
- [62] G. L. G. Sleijpen and D. R. Fokkema, *Electron. Trans. Numer. Anal.* **1**, 2000 (1993).
- [63] O. Axelsson, *Numer. Math.* **51**, 209 (1987).
- [64] J. C. Osborn, R. Babich, J. Brannick, R. C. Brower, M. A. Clark, S. D. Cohen, and C. Rebbi, *Proc. Sci. LATTICE2010* (**2010**) 037 [[arXiv:1011.2775](https://arxiv.org/abs/1011.2775)].
- [65] P. Hernandez, *Nucl. Phys.* **B536**, 345 (1998).