

Accelerated overlap fermions

Waseem Kamleh,^{1,*} David H. Adams,^{1,2,†} Derek B. Leinweber,^{1,‡} and Anthony G. Williams^{1,§}

¹*CSSM Lattice Collaboration, Special Research Centre for the Subatomic Structure of Matter and Department of Physics and Mathematical Physics, University of Adelaide 5005, Australia*

²*Physics Department, Duke University, Durham, North Carolina 27708*

(Received 20 February 2002; published 9 July 2002)

Numerical evaluation of the overlap Dirac operator is difficult since it contains the sign function $\epsilon(H_w)$ of the Hermitian Wilson-Dirac operator H_w with a negative mass term. The problems are due to H_w having very small eigenvalues on the equilibrium background configurations generated in current day Monte Carlo simulations. Since these are a consequence of the lattice discretization and do not occur in the continuum version of the operator, we investigate in this paper to what extent the numerical evaluation of the overlap can be accelerated by making the Wilson-Dirac operator more continuum-like. Specifically, we study the effect of including the clover term in the Wilson-Dirac operator and smearing the link variables in the irrelevant terms. In doing so, we have obtained a factor of 2 speedup by moving from the Wilson action to a fat link irrelevant clover action as the overlap kernel.

DOI: 10.1103/PhysRevD.66.014501

PACS number(s): 11.15.Ha, 12.38.Gc

I. INTRODUCTION

The overlap formalism [1–4] leads, in the vector case, to a lattice formulation of QCD based on the overlap Dirac operator [5], given (in the massless case) by

$$D_o = \frac{1}{2a} [1 + \gamma_5 \epsilon(H_w)], \quad \epsilon(H) = \frac{H}{\sqrt{H^2}}, \quad (1)$$

(a = lattice spacing) where

$$H_w = \gamma_5 \left(D_w - \frac{m}{a} \right), \quad (2)$$

is a Hermitian operator constructed from the Wilson-Dirac operator D_w [6] with m being a tuning parameter.¹ The free field propagator of D_o has the correct continuum limit and is free of doublers when $0 < m < 2$. Because of its origin in the overlap formalism, D_o has good chiral properties [7]; this can also be seen from the fact that it satisfies [8] the Ginsparg-Wilson relation [9]

$$\gamma_5 D + D \gamma_5 = 2a D \gamma_5 D. \quad (3)$$

Lattice Dirac operators satisfying this relation have an exact, lattice-deformed chiral symmetry [10], can have exact zero modes with definite chirality [11], as well as the absence of mass renormalization and other promising theoretical properties [12–14].

The nice theoretical properties of the overlap Dirac operator come at a price: numerical evaluation of it via the poly-

nomial approximation is difficult due to the discontinuity at the origin of the matrix sign function $\epsilon(H)$. Practical methods have been developed in which $\epsilon(H)$ is approximated by a sum over poles $\epsilon_N(H)$, using either the so-called polar decomposition or the optimal rational polynomial approximation [15,16], both of which take the form

$$\epsilon_N(H) = H \left(c_0 + \sum_{k=1}^N \frac{c_k}{H^2 + d_k} \right). \quad (4)$$

The two approximations only differ in their choice of coefficients $\{c_0, c_k, d_k\}$, and both are evaluated (indirectly) using a multishift conjugate gradient (CG) matrix inverter [17] to calculate their action on a vector. This is an iterative procedure where each iteration requires one evaluation of the matrix operator H^2 acting on a vector (i.e. two evaluations of H), and the number of iterations required to reach a given solution precision is proportional to the condition number of H , $\kappa(H) = |\lambda_{\max}/\lambda_{\min}|$, which is the ratio of the largest eigenvalue of H to the smallest eigenvalue [16].

Triangle inequalities lead to an upper bound [18] given by $|\lambda_{\max}| \leq (8-m)/a$ for the operator H_w in Eq. (2). The lower bound $|\lambda_{\min}|$ can be zero though. The lattice gauge fields for which $\lambda_{\min} = 0$ form a subspace of measure zero in the space of all lattice gauge fields, so it is exceedingly unlikely that one would ever encounter them in a numerical simulation. However, our practical experience is that $|\lambda_{\max}| \leq 8$ while $|\lambda_{\min}|$ is often as small as 10^{-8} . This results in an unacceptably large value for the condition number $\kappa(H)$. There is a way to get around this problem though [19]. The typical spectrum of H_w is characterized by a handful of isolated low-lying eigenmodes, so one can project these out and deal with them explicitly. The condition number for the remaining part of the spectrum is then small enough that the approximation in Eq. (4) becomes feasible. In practical simulations, after projecting out the isolated low-lying modes, $\epsilon_N(H_w)$ takes roughly speaking $O(100-300)$ iterations to converge

*Electronic address: wkamleh@physics.adelaide.edu.au

†Electronic address: adams@phy.duke.edu

‡Electronic address: dleinweb@physics.adelaide.edu.au

§Electronic address: awilliam@physics.adelaide.edu.au

¹We are assuming that the Wilson parameter has been set to its canonical value $r=1$.

for $N \approx 14$, meaning that overlap fermions with the standard H_w are about $O(200-600)$ times more expensive than standard Wilson fermions.

Obviously it is desirable to improve upon this situation in order to make simulations with overlap fermions more feasible. In this paper we investigate ways to do this by modifying the operator H_w in the overlap formula in Eq. (1) so that its spectral properties are improved. The improvements we seek are twofold: (i) An upward shift in the magnitude of the low-lying eigenvalues of H_w so as to decrease the condition number, and (ii) a reduction in the *density* of low-lying eigenvalues, so as to make the projection method of Ref. [19] more efficient. Furthermore, our aim is to produce an implementation of the overlap formalism that will perform efficiently on large-scale parallel computing architectures. On such architectures, the cost of internode communication is typically high compared to the cost of intranode computation. We therefore demand that our candidate H be no less sparse than the Hermitian Wilson-Dirac operator, that is, possess at most nearest neighbor couplings.

II. FERMION ACTIONS

The continuum version $H_c = \gamma_5[\not{b} - (m/a)]$ of H_w has the lower bound $|\lambda_{\min}| \geq m/a$ since $H_c^2 = -\not{b}^2 + (m/a)^2 \geq (m/a)^2$. Hence the near zero values of the lowest eigenvalues of H_w on equilibrium backgrounds at currently accessible β are a result of the lattice discretization. Our aim is to shift the lowest eigenvalues away from zero by making H_w , or more specifically, the Wilson-Dirac operator D_w in H_w , more continuumlike. In the framework of nonperturbative improvement of lattice operators (see, e.g., [20]), $O(a)$ lattice artifacts in D_w are removed by adding the clover term of Ref. [21]. A simple heuristic argument for why this should be beneficial in the present situation is the following. We write the Wilson-Dirac operator as

$$D_w = \not{V} + \frac{a}{2} \Delta, \quad (5)$$

where the naive lattice Dirac operator \not{V} and lattice Laplace operator Δ are given by

$$\begin{aligned} a \not{V}_{x,x'} &\equiv a(\gamma^\mu \nabla_\mu)_{x,x'} \\ &= \frac{1}{2} \sum_\mu \{ \gamma_\mu [U_\mu(x) \delta_{x+e_\mu, x'} \\ &\quad - U_\mu^\dagger(x-e_\mu) \delta_{x-e_\mu, x'}] \}, \end{aligned} \quad (6)$$

$$\begin{aligned} a^2 \Delta_{x,x'} &= 8 \delta_{x,x'} - \sum_\mu [U_\mu(x) \delta_{x+e_\mu, x'} \\ &\quad + U_\mu^\dagger(x-e_\mu) \delta_{x-e_\mu, x'}]. \end{aligned} \quad (7)$$

The γ matrices are chosen to be Hermitian, so \not{V} is anti-Hermitian and Δ is Hermitian and positive. Define the operator C by the relation

$$\not{V}^2 = \nabla^2 + C \quad (8)$$

(where $\nabla^2 = \sum_\mu \nabla_\mu \nabla_\mu$). Then $C = \frac{1}{4} [\gamma_\mu, \gamma_\nu] [\nabla_\mu, \nabla_\nu]$ is $\sim O(a)$ and coincides with the usual clover term (with coefficient $c_{sw}=1$, the tree level value) up to an $O(a^2)$ term. Here and in the following $O(a^p)$ denotes a lattice operator whose leading term in a formal expansion in powers of the lattice spacing is $\sim a^p$. Now, setting the parameter m in H_w to its canonical value $m=1$ we have

$$\begin{aligned} H_w^2 &= \left(D_w - \frac{1}{a} \right)^* \left(D_w - \frac{1}{a} \right) \\ &= -\not{V}^2 - \Delta + \left(\frac{a}{2} \Delta \right)^2 + \frac{a}{2} [\Delta, \not{V}] + \frac{1}{a^2}. \end{aligned} \quad (9)$$

Straightforward calculations show that $\Delta + \not{V}^2 \sim O(a^2)$ and $[\Delta, \not{V}] \sim O(a)$; hence, by Eq. (8), we have $\not{V}^2 + \Delta = C + O(a^2)$. Hence we obtain the lower bound

$$H_w^2 \geq \frac{1}{a^2} - C - O(a^2) = \frac{1}{a^2} - O(a). \quad (10)$$

Thus the lower bound $1/a^2$ on the continuum version of H_w^2 is spoiled in the lattice case by an $O(a)$ term. If we now add C to Δ in Eq. (5), i.e. replace

$$D_w \rightarrow D_{cw} \equiv \not{V} + \frac{a}{2} (\Delta + C) \quad (11)$$

we find

$$\begin{aligned} H_{cw}^2 &= \left(D_{cw} - \frac{1}{a} \right)^* \left(D_{cw} - \frac{1}{a} \right) \\ &= -\not{V}^2 - (\Delta + C) + \left(\frac{a}{2} (\Delta + C) \right)^2 + \frac{a}{2} [\Delta + C, \not{V}] + \frac{1}{a^2} \\ &\geq \frac{1}{a^2} - O(a^2). \end{aligned} \quad (12)$$

Hence the $O(a)$ term ($-C$) in Eq. (10) has dropped out and the continuum lower bound $1/a^2$ is now only spoiled by an $O(a^2)$ term.

However, it is well-known that adding a clover term only improves the chiral properties of the Wilson-Dirac operator on smooth backgrounds, and that the localization of the real eigenvalues of D_{cw} is actually *worse* than for D_w on typical gauge backgrounds generated in Monte Carlo simulations [22–24]. This suggests that to further improve the chiral properties of D_{cw} we should consider smoothing the lattice gauge field. In Ref. [23], DeGrand *et al.* found that a significant improvement in the chiral properties can be achieved by applying an APE smearing procedure [25–28] to the link variables, leading to a fat link version of D_{cw} . (The idea of using fat links in fermion actions was first explored by the MIT group [29].) More recently, Zanotti *et al.* have shown in [30] that such improvement can be achieved by smearing only the link variables appearing in the irrelevant operators, i.e. in the Wilson and clover terms. This has the advantage of

preserving the short distance quark interactions. (The idea of using fat links in the irrelevant operators had been independently suggested previously in Ref. [18].)

Motivated by the preceding discussion, we compare the evaluation of the usual overlap Dirac operator with the operators obtained by replacing H_w in the overlap formula of Eq. (1) with the following variants (the lattice spacing has been set to 1 unless specified otherwise):

(i) *Hermitian Wilson-Dirac operator with clover term:*

$$H_{\text{cw}}(m, c_{\text{sw}}) = \gamma_5 \left[\nabla + \frac{1}{2} \left(\Delta - \frac{c_{\text{sw}}}{2} \sigma \cdot F \right) - m \right], \quad (13)$$

where

$$\sigma_{\mu\nu} = \frac{1}{2} [\gamma_\mu, \gamma_\nu], \quad F_{\mu\nu}(x) = \frac{1}{2} [C_{\mu\nu}(x) - C_{\mu\nu}^\dagger(x)], \quad (14)$$

$$C_{\mu\nu}(x) = \frac{1}{4} [U_{\mu\nu}(x) + U_{-\nu\mu}(x) + U_{\nu-\mu}(x) + U_{-\mu-\nu}(x)]. \quad (15)$$

(ii) *Fat link Hermitian Wilson-Dirac operator, both with and without clover term:*

$$H_{\text{fw}}(m, \alpha n_{\text{ape}}) = \gamma_5 \left(\nabla + \frac{1}{2} \Delta^{(\alpha n_{\text{ape}})} - m \right), \quad (16)$$

$$H_{\text{fcw}}(m, c_{\text{sw}}, \alpha n_{\text{ape}}) = \gamma_5 \left[\nabla + \frac{1}{2} \left(\Delta^{(\alpha n_{\text{ape}})} - \frac{c_{\text{sw}}}{2} \sigma \cdot F^{(\alpha n_{\text{ape}})} \right) - m \right], \quad (17)$$

where APE-smearing is carried out on the individual links in the irrelevant operators by making the replacement

$$U_\mu(x) \rightarrow U_\mu^{(\alpha)}(x) = \mathcal{P} \left((\alpha - 1) U_\mu(x) + \frac{\alpha}{6} \sum_{\pm\nu \neq \mu} U_\nu(x) U_\mu(x + a e_\nu) U_\nu^\dagger(x + a e_\mu) \right). \quad (18)$$

Here \mathcal{P} denotes projection of the RHS of Eq. (18) back to the SU(3) gauge group. That is, each link is modified by replacing it with a combination of itself and the surrounding staples to give a set of ‘‘fat links.’’ The means by which one projects back to SU(3) is not unique. We choose an SU(3) matrix $U_\mu^{(\alpha)}(x)$ such that the gauge invariant measure $\text{ReTr}(U_\mu^{(\alpha)}(x) X_\mu^\dagger(x))$ is maximal, where $X_\mu(x)$ is the smeared link before projection, that is, $U_\mu^{(\alpha)}(x) \equiv \mathcal{P} X_\mu(x)$. As the process of APE-smearing removes short-distance physics, it is preferable to only smear the irrelevant operators. Throughout this work ‘‘fat’’ means APE smearing of links in irrelevant terms only. Here α is the smearing fraction and n_{ape} is the number of smearing sweeps (18) we perform. As shown in [28], we can effectively reduce the two-dimensional parameter space (α, n_{ape}) to a one-dimensional space that depends solely on the product αn_{ape} , and this is reflected in the notation in Eqs. (16), (17).

Finally, as in [30], we can perform tadpole or mean-field improvement (MFI) [31] to bring our links closer to unity. This consists of updating each link with a division by the mean link, which is the fourth root of the average plaquette,

$$u_0 = \left\langle \frac{1}{3} \text{ReTr} U_{\mu\nu}(x) \right\rangle_{x, \mu < \nu}^{1/4}. \quad (19)$$

In the case of H_w and H_{fw} , mean-field improvement has little effect, entering in only as a single power in both cases. For H_w , mean field improvement effectively changes the value of m and renormalizes the Wilson parameter r . In the case of H_{fw} it has a similar effect but we have two mean link values, one for the untouched set of links and one for the smeared set. However, u_0 enters in as the fourth power in

front of the clover term, effectively raising c_{sw} towards its nonperturbative value. Hence our final two variants of H_w are the following.

(iii) *MFI clover Hermitian Wilson-Dirac operator, both with and without fat links:*

$$H_{\text{cl}}^{\text{mfi}}(m, c_{\text{sw}}) = \gamma_5 \left[\frac{1}{u_0} \nabla + \frac{1}{2} \left(\frac{1}{u_0} \Delta - \frac{c_{\text{sw}}}{2u_0^4} \sigma \cdot F \right) - m \right], \quad (20)$$

$$H_{\text{fcl}}^{\text{mfi}}(m, c_{\text{sw}}, \alpha n_{\text{ape}}) = \gamma_5 \left[\frac{1}{u_0} \nabla + \frac{1}{2} \left(\frac{1}{u_0^{\text{fl}}} \Delta^{(\alpha n_{\text{ape}})} - \frac{c_{\text{sw}}}{2(u_0^{\text{fl}})^4} \sigma \cdot F^{(\alpha n_{\text{ape}})} \right) - m \right], \quad (21)$$

where we have differentiated the mean link u_0 for the untouched links and u_0^{fl} for the fat links. We refer to the MFI fat clover action as the FLIC (fat-link irrelevant clover) action. The FLIC action was recently introduced and studied in Ref. [30]. If followed by a number (e.g. FLIC12) this denotes the number of APE-smearing sweeps (at $\alpha=0.7$) used in the action.

Before proceeding to the numerical results it is worth pointing out that the previous analytical results on the locality [32] and continuum limit of the axial anomaly [33–36] and index [37] of the overlap Dirac operator continue to hold when H_w is replaced by any of the variants given above in the overlap formula. In the case of the axial anomaly and index, this is essentially because the leading order term in the expansion of commutators of the covariant finite difference

operators in powers of the lattice spacing is unchanged, and the variants of H_w all coincide with H_w in the free field case. Regarding locality, the admissibility bounds of [18,32] on the plaquette variables get modified somewhat when the different variants of H_w are used. In light of the heuristic arguments above and our numerical results below, we expect that it should be possible to derive improved (i.e., less restrictive) bounds in these cases, although so far we have not been able to show this.

We also mention that more general variants of the overlap Dirac operator have been considered where one starts with an approximate solution to the Ginsparg-Wilson relation and then gets an exact solution by substituting into the overlap formula [38–42].² This has led to variants of the overlap action which are both easier to evaluate and more local than the original. However, it is not clear whether such general operators will have the good topological properties of the standard overlap Dirac operator, namely exact zero-modes with definite chirality in topologically nontrivial backgrounds, (cf. the counter-example of Chiu [43,44]). This is important in connection with the lattice implementation of the Witten-Veneziano formula for the η' mass with GW fermions [14], since for the argument to work it is essential that the would-be zero modes be exact zero modes.

III. SPECTRAL FLOW COMPARISON

In order to test the merits of each of our proposed actions, we first calculate the spectral flow of each of them to see if our reasoning regarding their low-lying spectra is valid. From the quadratic form of the lower bounds as a function of m , and based upon results given in Ref. [45], we expect there to be some peak value of m for which the gap around zero is the largest. We calculated the flow of the lowest 15 eigenvalues as a function of m for an ensemble of 10 mean-field improved Symanzik configurations at $\beta=4.38$ and size $8^3 \times 16$. The following flow graphs allow us to see the m value for the biggest gap, and also allow us to compare the different actions. As we are interested in the magnitude of the low-lying values rather than their sign, we plot $|\lambda|$ vs m .

We begin by examining the flow of the Wilson and clover action in Fig. 1. We see the Wilson spectrum is very poor, with a high density of very small eigenmodes and no gap away from zero. The addition of the clover term (at $c_{sw}=1$) provides some improvement, shifting the flow upwards and moving the peak values towards $m=1$ as expected. The presence of many small eigenmodes persists however, although their density is clearly reduced.

In Fig. 2 we examine the MFI clover and fat Wilson actions. Mean field improvement assists the basic clover action somewhat, spreading the spectrum upwards, although the lowest modes are not raised significantly. The mass value at which the low-lying density is minimized has moved signifi-

cantly away from $m=1.2$ to around $m=0.6$. As mentioned earlier, essentially all MFI does in this case is to change the value of c_{sw} to $1.0/u_0^3$, pushing it towards its nonperturbative value. Modifying the Wilson action by smearing the irrelevant operators provides a considerable improvement. While there are still some small modes present, their density has been greatly reduced, and the spectral flow now has a clear division between the isolated low-lying modes and the modes where the spectral density becomes high, which are well separated from zero. Smearing was performed with $\alpha=0.7$ and $n_{ape}=12$ smearing sweeps.

Results for the fat clover and FLIC12 actions are shown in Fig. 3. The spectral flow of the fat clover action clearly demonstrates the superiority of clover-improved actions. The gap around zero is enhanced again over the fat Wilson action, and the number of isolated low-lying modes is significantly reduced. As the fat links are already close to unity, the addition of mean field improvement only affects the fat clover flow slightly, raising the gap around zero a little and spreading the eigenvalues upwards slightly also. The low-lying density is again very good in this case and far superior to that of the Wilson action.

To confirm our results we choose the Wilson action as a “baseline” and compare it against the FLIC action (the best of the alternative actions) on a larger, finer lattice, $12^3 \times 24$ at $\beta=4.60$. This time we only use 4 smearing sweeps in the FLIC action since FLIC4 has less fattening and is the choice used in actual simulations [30]. We see that the Wilson action benefits significantly from the smaller lattice spacing, as there is now a visible separation from zero before the modes become dense. The FLIC action has the same characteristics as on the coarser lattice, but it now has a peak separation of the dense modes from zero of around 0.45. (See Fig. 4.)

Additionally, we tested the dependence of the FLIC action upon the amount of smearing done (see Fig. 5). As stated in [28], we only effectively need to vary the product αn_{ape} , so we fix α at 0.7 and vary n_{ape} between 0 and 12. We observe that the initial 4–6 sweeps have a significant effect, but beyond 6 sweeps the effect is marginal, with the low lying density remaining roughly constant and the eigenvalues being compressed very slightly downwards.

We note that there are other suggestions to accelerate the computation of the sign function by using improved gauge actions [46,47]. Our analysis is already based on improved gauge configurations and further improvements arising from the use of the fermion actions described herein are in addition to that of using improved gluon-field configurations (see Fig. 6). In particular, we use a 2-loop tadpole-improved Luscher-Wiese action. While explicit tests have not been performed using the FLIC action and Iwasaki glue, the “smoothing” effect of APE-smearing the irrelevant operators (which is responsible for the improved spectrum) does not depend on the details of glue that is used. Hence the improvements observed in the following are expected to hold for Iwasaki glue as well.

IV. RESULTS

Having obtained some understanding of the low-lying spectra of the various actions via the flow diagrams, we now

²Specifically, if D_{approx} is some approximate solution to the GW relation then $A=1-D_{approx}$ satisfies $A*A \approx 1$. An exact solution D to the GW relation, which is approximately equal to D_{approx} , is then obtained via the overlap formula by setting $D=1-(A/\sqrt{A*A})$.




FIG. 1. Spectral flow of the Wilson action (left) and the clover action (right) at $\beta=4.38$.




FIG. 2. Spectral flow of the MFI clover action (left) and the fat Wilson action (right) at $\beta=4.38$.




FIG. 3. Spectral flow of the fat clover action (left) and FLIC12 action (right) at $\beta=4.38$.




FIG. 4. Spectral flow of the fat Wilson action (left) and FLIC4 action (right) at $\beta=4.60$.




FIG. 5. Dependence of the FLIC spectrum at $\beta=4.60, m=1.35$ (left) and $\beta=4.38, m=1.45$ (right) on the number of APE smearing sweeps.

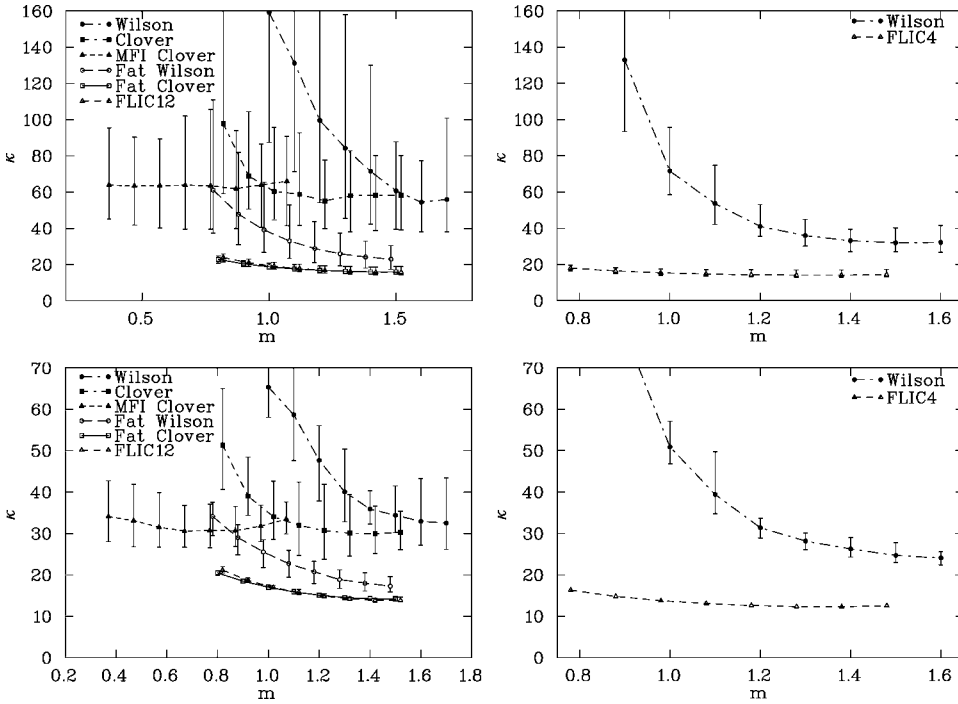


FIG. 6. Condition numbers of the various actions. (Top-left) Results for $\beta=4.38$ with 5 projected modes. (Bottom-left) Results for $\beta=4.38$ with 15 projected modes. (Top-right) Results for $\beta=4.60$ with 5 projected modes. (Bottom-right) Results for $\beta=4.60$ with 15 projected modes.

turn to quantitative comparisons. First we examine the condition number, κ , of the different actions as a function of m . We show below the condition numbers having projected out the lowest 5 eigenmodes and the lowest 15 eigenmodes on the 2 lattices that we used. The points are the mean condition numbers across the ensembles, and the error bars indicate the minimum and maximum condition numbers, giving an idea of the variation in κ . The smeared irrelevant-term actions here used 12 APE sweeps at $\alpha=0.7$ for the coarse lattice and 4 sweeps for the fine lattice. Some points are offset horizontally for clarity.

Two things are immediately noticeable. First, the smeared irrelevant-term actions are much better conditioned than the unsmeared actions, and second, the variation of κ between configurations is less. It should be noted that the variations (error bars) are displayed for all actions, but are smaller than the plot symbol at some points of the fat clover and FLIC lines. Projecting out an additional 10 eigenvalues has a significant effect on the unsmeared actions, but relatively little effect on the smeared actions due to reduction in the number of isolated low-lying values. In terms of condition number, the fat clover and FLIC actions are clearly and significantly superior to the other actions, with the FLIC action possessing a (slight) edge over the fat clover which arises from the mean field improvement.

As the clover term is quite fast to evaluate, we discard the fat Wilson as a candidate action at this point as it is the least well-conditioned of the smeared actions. Given the similarity between the clover-improved actions with and without mean-field improvement, we focus on the MFI clover and FLIC actions. We now compare in detail the performance for three actions: the Wilson, MFI clover and FLIC. To see how improving the condition number translates into a saving in CG iterations, we calculated the number of multi-CG iterations required to evaluate D_0 once across the ensemble for each of

these actions, using some typical simulation parameters. The number of modes we project out is determined by the spectral density at the highest mode. The spectral flow results show that in each of the actions tested, past a certain point the spectral density becomes very high and there is no computational benefit in projecting modes past this point. Therefore, we choose to stop projecting out modes when only a marginal improvement in the condition number is gained by projecting out more eigenmodes.

The Wilson and MFI clover are tested using the 14th order optimal rational polynomial (ORP) approximation [19]. The improved condition number of the FLIC actions allows us to use the 12th order polar decomposition, chosen to give a maximum deviation from $\epsilon(x)$ of less than 10^{-6} compared to the 3.1×10^{-5} of the 14th order ORP. The N th order polar decomposition is specified by

$$c_0=0, \quad c_k = \frac{1}{N \cos^2\left(\frac{\pi}{2N}\left(k - \frac{1}{2}\right)\right)},$$

$$d_k = \tan^2\left(\frac{\pi}{2N}\left(k - \frac{1}{2}\right)\right). \quad (22)$$

Low-lying modes are projected out where necessary. The sign function solution is calculated to a precision of 10^{-6} across the fine ensemble and the coarse ensemble used above. The value of m is chosen differently for each of the actions to optimize κ . Given the relative lack of improvement in using the MFI clover action compared to the Wilson, we discard it at this point and concentrate on comparing the Wilson and FLIC actions. As the results in Table I show, the FLIC action is by far the best in terms of convergence with a

TABLE I. Conjugate gradient (CG) iterations needed for a single evaluation of $\epsilon_N(x)$ using actual simulation parameters.

Action	β	Projections	Mean	Min	Max
Wilson	4.38	15	219	188	253
	4.60	15	202	190	212
MFI clover	4.38	15	200	178	240
FLIC12	4.38	10	92	89	100
FLIC6	4.38	10	90	86	101
FLIC4	4.60	15	109	106	112

reduction in iterations compared to the Wilson action of a factor of between 1.9 and 2.4.

However, what is not clear from this is how the savings in iterations translates into the most important quantity, a savings in compute time. Shifting from a standard Wilson action to a partially smeared action means that we now have two sets of gauge fields, the standard and smeared links. This doubles the number of vector-multiplications needed, and the standard spin-projection trick [48] is no longer applicable, providing an additional factor of 2 in both the multiplications needed and the communications needed. So, moving from the Wilson action to the FLIC action costs us a factor of 4 in vector multiplications and a factor of 2 in communications, plus the overhead for the clover term. On the other hand, evaluating the action of $\epsilon_N(x)$ on a vector costs $O(2N)$ vector multiplications in addition to the two evaluations of the kernel, H . While vector multiplications form a significant part of the cost of evaluating H , they are not the only part. There is a relatively high cost of communication compared to computation on the parallel architectures that we wish to use. It quickly becomes clear that the only real way to see how much of an improvement we have made is to do an actual calculation and compare the compute time needed.

To test the actual speedup, we choose to calculate the low-lying eigenmodes of $H_o^2 = D_o^\dagger D_o$ for the two different kernels, Wilson and FLIC. This calculation allows us to verify that both kernels give the appropriate spectral properties [16], and also allows us to calculate directly the relative compute time needed to evaluate D_o in each case. For the Wilson action we used the 14th order rational polynomial approximation in the region which it is bounded by unity ($0.025 < |x| < 1.918$) and where the maximal deviation from $\epsilon(x)$ is 3.1×10^{-5} . We used the mass parameter $m = 1.65$ and projected out 15 eigenvalues. For the FLIC action, we can take advantage of the improved condition number without reducing the accuracy of our approximation by using the polar decomposition at 12th order, which is sufficient to provide a maximal deviation of less than 3.1×10^{-5} . This saves us a (small) amount of computation. To optimize the condition number we choose to perform only 6 APE sweeps with the mass parameter set to $m = 1.45$ and projecting out 10 eigenvalues. To minimize the computation needed, we implement individual pole convergence testing in our multi-CG routine. The first pole is considered converged in the n th iteration according to the usual criterion based on the residue, $\|r_n\| < \delta$, where we chose $\delta = 10^{-8}$. The convergence criterion for the other poles is easily deduced by noting the

TABLE II. Actual compute time spent in the various parts of the algorithm.

Code portion	Wilson	FLIC6
1 kernel-vector evaluation (H)	0.022 sec	0.037 sec
1 multi-CG iteration (including H)	0.133 sec	0.154 sec
1 multi-CG iteration (excluding H)	0.089 sec	0.079 sec
1 overlap-vector evaluation	25.52 sec	13.67 sec

shifted polynomial structure of the residual, $r_n^i = P_n[H^2 + \sigma(i)]r_0 = \zeta_n^\sigma P_n(H^2)r_0 = \zeta_n^\sigma r_n$. Then the i th pole is considered converged if

$$\|r_n\| \zeta_n^{\sigma(i)} < 0.1 \times \delta, \quad (23)$$

where $\zeta_n^{\sigma(i)}$ is defined as in Eq. (2.44) of Ref. [17]. We have tested this convergence criterion by calculating individual residues and found it to be numerically very safe, and also to save significant amounts of computation. We consider the ten $8^3 \times 16, \beta = 4.38$ lattices. Computations are performed on 4 nodes of the Orion supercomputer (a Sun E420R cluster comprised of 40 nodes), with each node possessing 4 GB of RAM, 16 MB of L2-cache, and 4 UltraSPARC II 450 MHz processors and with nodes are connected by Myrinet networking. The lowest 6 eigenmodes of H_o^2 are calculated on each configuration using the Ritz functional method [49]. We measure the compute time spent in each of the different parts of the ‘‘inner-CG’’ calculation, with the following results.

The results show that using the FLIC action as the kernel in the overlap formalism provides a savings of a factor of 1.9 in actual compute-time spent in evaluating the overlap action (see Table II). This is easily understood by first observing that the time spent in the fermion matrix multiplication constitutes less than half of the compute time spent in the inner CG inversion. Secondly, we have only paid a factor of 2 in compute time moving from the Wilson action to the FLIC action, not the potential factor of 4. This is because the time spent in communication and performing the γ matrix algebra is not negligible when compared to the time spent in performing the gauge field multiplications. Finally, as the improved condition number of the FLIC kernel allows us to use the 12th order polar decomposition, we expend less effort per iteration in the CG component of the sign function evaluation. This is because the number of unconverged poles per iteration is reduced, as demonstrated in Table III.

These facts mean that the overall compute time per inner CG iteration increases by only 15% when moving to the FLIC kernel, and hence the savings of 55% in the total number of inner CG iterations needed translates into a savings in compute time. Thus we have shown that the FLIC action is numerically superior to the Wilson action as an overlap kernel. What has not been answered is what, if any, are the differences in physical properties of D_o using the different kernels. For example, overlap fermions are free of $O(a)$ errors irrespective of the choice of kernel, but in general may have different $O(a^2)$ errors. This will be addressed in future work.

TABLE III. Breakdown of the mean convergence for each of the poles.

Pole	Wilson	FLIC6	Pole	Wilson	FLIC6
1	188_{-21}^{+32}	85_{-6}^{+11}	8	55_{-3}^{+4}	19_{-1}^{+1}
2	188_{-21}^{+32}	82_{-4}^{+10}	9	39_{-2}^{+2}	14_{-1}^{+1}
3	188_{-21}^{+32}	65_{-4}^{+6}	10	28_{-2}^{+1}	10_{-0}^{+0}
4	188_{-21}^{+31}	50_{-2}^{+4}	11	19_{-1}^{+2}	7_{-1}^{+0}
5	161_{-13}^{+15}	39_{-2}^{+3}	12	14_{-1}^{+0}	4_{-0}^{+0}
6	116_{-8}^{+7}	31_{-2}^{+2}	13	9_{-1}^{+0}	
7	80_{-5}^{+5}	24_{-1}^{+2}	14	5_{-0}^{+1}	

V. CONCLUSION

Practical implementations of the overlap-Dirac operator use a sum over poles to approximate the matrix sign function. These approximations are evaluated using an iterative conjugate gradient routine. As each iteration requires about twice as much computational effort to evaluate as a single evaluation of H_w , reducing the number of iterations needed is the most direct way of reducing the expense of the overlap formalism. To succeed in this, we select an overlap kernel with an improved condition number motivated by analytic arguments. From the six candidate actions tested, the FLIC action has the best convergence properties, requiring fewer

low-lying projections than the Wilson action and providing a savings in iterations by about a factor of 2. This savings in iterations translates almost directly into a savings in computation time. We restate that only the irrelevant operators are smeared, and that minimal smearing is required, 6 sweeps at $\alpha=0.7$ for $\beta=4.38, a=0.165(2)$ or 4 sweeps at $\alpha=0.7$ for $\beta=4.60, a=0.125(2)$. As the FLIC action has only nearest neighbor couplings, it is well suited to calculations on highly parallel machines. We recognize that there will be some implementation dependence in our compute-time results, but believe that this dependence will be sufficiently small that all groups who wish to perform overlap calculations will benefit in moving from the Wilson to the FLIC kernel. As we have concluded that the FLIC action is a numerically superior kernel, we can proceed to investigate the dependence of the overlap action's physical properties on the kernel action.

ACKNOWLEDGMENTS

W.K. wishes to sincerely thank Urs Heller for his correspondence regarding the overlap formalism and also thanks Jianbo Zhang for a number of helpful discussions. We have also benefited from discussions with Robert Edwards, Urs Heller, and Herbert Neuberger. Additionally, we thank F. Bonnet for his contribution to the APE smearing code. The calculations were carried out on the Orion Supercomputer of the Australian National Computing Facility for Lattice Gauge Theory. This work was supported by the Australian Research Council.

-
- [1] R. Narayanan and H. Neuberger, Phys. Lett. B **302**, 62 (1993).
[2] R. Narayanan and H. Neuberger, Nucl. Phys. **B412**, 574 (1994).
[3] R. Narayanan and H. Neuberger, Phys. Rev. Lett. **71**, 3251 (1993).
[4] R. Narayanan and H. Neuberger, Nucl. Phys. **B443**, 305 (1995).
[5] H. Neuberger, Phys. Lett. B **417**, 141 (1998).
[6] K.G. Wilson, Phys. Rev. D **10**, 2445 (1974).
[7] H. Neuberger, Chin. J. Phys. (Taipei) **38**, 533 (2000).
[8] H. Neuberger, Phys. Lett. B **427**, 353 (1998).
[9] P.H. Ginsparg and K.G. Wilson, Phys. Rev. D **25**, 2649 (1982).
[10] M. Luscher, Phys. Lett. B **428**, 342 (1998).
[11] P. Hasenfratz, V. Laliena, and F. Niedermayer, Phys. Lett. B **427**, 125 (1998).
[12] P. Hasenfratz, Nucl. Phys. **B525**, 401 (1998).
[13] S. Chandrasekharan, Phys. Rev. D **60**, 074503 (1999).
[14] L. Giusti, G.C. Rossi, M. Testa, and G. Veneziano, Nucl. Phys. **B628**, 234 (2002).
[15] H. Neuberger, Phys. Rev. Lett. **81**, 4060 (1998).
[16] R.G. Edwards, U.M. Heller, and R. Narayanan, Nucl. Phys. **B540**, 457 (1999).
[17] B. Jegerlehner, hep-lat/9612014.
[18] H. Neuberger, Phys. Rev. D **61**, 085015 (2000).
[19] R.G. Edwards, U.M. Heller, and R. Narayanan, hep-lat/9905028.
[20] M. Luscher, hep-lat/9802029.
[21] B. Sheikholeslami and R. Wohlert, Nucl. Phys. **B259**, 572 (1985).
[22] C. Gattringer and I. Hip, Nucl. Phys. **B541**, 305 (1999).
[23] T. DeGrand, A. Hasenfratz, and T.G. Kovacs, Nucl. Phys. **B547**, 259 (1999).
[24] M. Stephenson, C. DeTar, T. DeGrand, and A. Hasenfratz, Phys. Rev. D **63**, 034501 (2001).
[25] M. Falcioni, M.L. Paciello, G. Parisi, and B. Taglienti, Nucl. Phys. **B251**, 624 (1985).
[26] APE Collaboration, M. Albanese *et al.*, Phys. Lett. B **192**, 163 (1987).
[27] F.D.R. Bonnet, P. Fitzhenry, D.B. Leinweber, M.R. Stanford, and A.G. Williams, Phys. Rev. D **62**, 094509 (2000).
[28] F. D. R. Bonnet, D. B. Leinweber, A. G. Williams, and J. M. Zanotti, Phys. Rev. D **65**, 114510 (2002).
[29] M.C. Chu, J.M. Grandy, S. Huang, and J.W. Negele, Phys. Rev. D **49**, 6039 (1994).
[30] J.M. Zanotti *et al.*, Phys. Rev. D **65**, 074507 (2002).
[31] G.P. Lepage and P.B. Mackenzie, Phys. Rev. D **48**, 2250 (1993).
[32] P. Hernandez, K. Jansen, and M. Luscher, Nucl. Phys. **B552**, 363 (1999).
[33] K. Fujikawa, Nucl. Phys. **B546**, 480 (1999).
[34] H. Suzuki, Prog. Theor. Phys. **102**, 141 (1999).
[35] D.H. Adams, hep-lat/9812003.
[36] Y. Kikukawa and A. Yamada, Phys. Lett. B **448**, 265 (1999).
[37] D.H. Adams, J. Math. Phys. **42**, 5522 (2001).

- [38] W. Bietenholz, I. Hip, and K. Schilling, Nucl. Phys. B (Proc. Suppl.) **106**, 829 (2002).
- [39] W. Bietenholz, Eur. Phys. J. C **6**, 537 (1999).
- [40] P. Hasenfratz, S. Hauswirth, K. Holland, T. Jorg, and F. Niedermayer, Nucl. Phys. B (Proc. Suppl.) **106**, 799 (2002).
- [41] P. Hasenfratz *et al.*, Int. J. Mod. Phys. C **12**, 691 (2001).
- [42] MILC Collaboration, T. DeGrand, Phys. Rev. D **63**, 034503 (2001).
- [43] T.-W. Chiu and T.-H. Hsieh, Phys. Rev. D **65**, 054508 (2002).
- [44] T.-W. Chiu, Phys. Lett. B **521**, 429 (2001).
- [45] R.G. Edwards, U.M. Heller, and R. Narayanan, Nucl. Phys. **B535**, 403 (1998).
- [46] R.G. Edwards and U.M. Heller, Phys. Rev. D **65**, 014505 (2002).
- [47] S.J. Dong *et al.*, Phys. Rev. D **65**, 054507 (2002).
- [48] M.G. Alford, T.R. Klassen, and G.P. Lepage, Nucl. Phys. **B496**, 377 (1997).
- [49] T. Kalkreuter and H. Simma, Comput. Phys. Commun. **93**, 33 (1996).