Suppressing dislocations in normalized hypercubic smearing

Thomas DeGrand

Department of Physics, University of Colorado, Boulder, CO 80309, USA

Yigal Shamir and Benjamin Svetitsky

Raymond and Beverly Sackler School of Physics and Astronomy, Tel Aviv University, 69978 Tel Aviv, Israel

Normalized hypercubic smearing improves the behavior of dynamical Wilson-clover fermions, but has the unwanted side effect that it can occasionally produce spikes in the fermion force. These spikes originate in the chain rule connecting the derivative with respect to the smeared links to the derivative with respect to the dynamical links, and are associated with the presence of dislocations in the dynamical gauge field. We propose and study an action designed to suppress these dislocations. We present evidence for improved performance of the hybrid Monte Carlo algorithm. A side benefit is improvement in the properties of valence chiral fermions.

I. INTRODUCTION

Smeared links are widespread in present-day lattice gauge theory simulations. A smeared link $V_{x,\mu} \in \mathrm{U}(N_c)$ is a parallel transporter from x to $x+\hat{\mu}$ that is constructed from the dynamical gauge field $U_{x,\mu} \in \mathrm{SU}(N_c)$ in the vicinity of the lattice sites x and $x+\hat{\mu}$. The replacement of the dynamical, or thin, links $U_{x,\mu}$ in the fermion action by smeared, or fat, links $V_{x,\mu}$ typically leads to a reduction of the discretization error. Intuitively, the fat links $V_{x,\mu}$ provide a smoother background for the fermions to propagate in, resulting in a more continuum-like behavior and thus smaller lattice artifacts.

In our own work on lattice gauge theory with fermions in higher representations we have been using normalized hypercubic (nHYP) smeared links in the Wilson-clover fermion action [1, 2]. We have indeed observed much reduced discretization effects. This allowed us to reach deeper into strong coupling, as well as to keep the clover term at its tree-level value $c_{SW} = 1$ [3, 4].

Our simulations with dynamical fermions were carried out with the standard Hybrid Monte Carlo (HMC) algorithm [5]. We use the familiar tools to accelerate the molecular dynamics (MD) integration: an additional heavy pseudofermion field as suggested by Hasenbusch [6]; multiple time scales for nested MD integration levels [7]; and a second-order Omelyan integrator [8].

This arsenal of techniques lends itself to many variations. As an example, one can try to improve the overall performance of the algorithm by using several intermediate Hasenbusch masses, each with its own pseudofermion action, at the same MD integration level. The idea is that, unlike the number of nested integration levels, which can only be changed discretely, the masses of the additional heavy pseudofermions can be tuned continuously, allowing for a more efficient optimization. This approach turned out to be successful for domain-wall fermions [9, 10].

We have experimented with these improvement schemes, but in many cases we have been stymied by continued low acceptance. Examination of our results suggests that the explanation of the problem lies in our smearing procedure. In the domain-wall simulations of the RBC and UKQCD collaborations [9, 10], smeared links have not been used, and the fermion force was obtained by directly differentiating the pseudofermion action with respect to the dynamical links $U_{x,\mu}$. In our simulations, on the other hand, we first differentiate the pseudofermion action with respect to the nHYP links $V_{x,\mu}$, and then apply the chain rule that is needed to convert the derivative with respect to the nHYP links into a derivative with respect to the dynamical (thin) links $U_{x,\mu}$. As will be clear below, this chain rule is the origin of our difficulties.

In this paper, we describe the problem and propose and study a remedy. Reduced acceptance is a result of spikes in the force; these appear in the chain-rule calculation because of dislocations in the dynamical gauge field, which produce large derivatives through the normalization step in the smearing. Our remedy is a new term in the gauge action, which suppresses these dislocations. This tames the fluctuations of the fermion force.

In Sec. II we present the evidence for the connection between the fat-to-thin chain rule and the acceptance of the HMC algorithm. In Sec. III we present the new term in the gauge action, designed to suppress dislocations. In Sec. IV we display the resulting improvement in performance of the HMC algorithm. We have also observed a side benefit—better behavior of chiral valence fermions, as we report in Sec. V. We conclude with a discussion in Sec. VI.

Of the smearing techniques in widespread use, Highly Improved Staggered Quarks, known as HISQ [11, 12], also make use of a normalization (i.e., reunitarization) step [see Eq. (1) below]. Our technique can be applied to them as well. Stout links [13] do not: The smeared

TABLE I: Ensembles used in this study and overall performance. The parameter γ is the coefficient of the new term in the gauge action, see Eqs. (10) and (14) below. β is the usual plaquette coupling, and κ the hopping parameter. n_1, n_2 and n_g are the number of steps of the three MD integration levels. The last column gives the acceptance rate. In all cases the lattice size is $L^3 \times T = 12^3 \times 24$, and the trajectory length is set to unity. The number of configurations in each ensemble ranges between 400 and 800.

γ	β	κ	n_1	n_2	n_g	acc.
0	9.6	0.1292	16	2	6	86%
0	9.65	0.129	16	2	6	92%
0.125	7.8	0.130	12	1	5	82%
0.25	5.6	0.130	12	1	5	92%

link is an analytic function of the thin links, and the problem we encounter with nHYP smearing is avoided. In favor of nHYP links, we recall their advantage, that the smearing range is small. Stout smearing, on the other hand, is usually repeated several times, giving rise to a less local fermion action. Local hypercubic geometry can be combined with the analytic stout recipe as HEX [14] or sHYP [2] smearing.

II. CHAIN RULE AND ACCEPTANCE RATE

Schematically, an nHYP link V is constructed as

$$V = \mathcal{P}(\Omega) \equiv \Omega Q^{-1/2} \ , \tag{1}$$

where[1]

$$Q = \Omega^\dagger \Omega \ , \tag{2}$$

and Ω is a weighted sum over paths (the precise definition will be given in the next section). Let S be a pseudofermion action that depends explicitly only on the fat links, and suppose that U is one of the thin links on which the weighted sum Ω depends. The thin-link force $\partial S/\partial U$ is related to the fat-link force $\partial S/\partial V$ via the "fat-to-thin" chain rule,

$$\frac{\partial S}{\partial U} = \sum_{i,j=1}^{N_c} \left(\frac{\partial S}{\partial V_{ij}} \frac{\partial V_{ij}}{\partial U} + \frac{\partial S}{\partial V_{ij}^*} \frac{\partial V_{ij}^*}{\partial U} \right) \ , \tag{3a}$$

$$\frac{\partial V}{\partial U} = \frac{\partial \Omega}{\partial U} Q^{-1/2} + \Omega \frac{\partial Q^{-1/2}}{\partial U} \ . \tag{3b}$$

Clearly, even if the fat-link force is well-behaved, small eigenvalues of Q can generally lead to large terms in the thin-link force. For Q to have an exceptionally small eigenvalue, the dynamical gauge field needs to be rough, or, loosely speaking, a dislocation should be present. At the same time, one should keep in mind that a locally rough gauge field does not always give rise to such exceptionally small eigenvalues. We will return to these considerations in more detail below.

The model we used for our tests is an SU(4) gauge theory with two Dirac fermions in the two-index antisymmetric (sextet) representation. The gauge action is the usual Wilson plaquette action

$$S_{\rm plaq} = \frac{\beta}{2N_c} \, {\rm Re\,tr} \sum_x \sum_{\mu \neq \nu} (1 - U_{x,\mu} U_{x+\hat{\mu},\nu} U^\dagger_{x+\hat{\nu},\mu} U^\dagger_{x,\nu}) \ , \tag{4}$$

with $N_c = 4$, to which we add a new piece, to be described below, with coefficient γ. As already mentioned we are using the Wilson-clover fermion action with nHYP links, with hopping parameter κ, and with $c_{SW} = 1$. We have one Hasenbusch mass $\mu = 0.2$ that effectively separates high- and low-momentum modes in the fermion determinant. The total pseudofermion action is thus

$$S_{\rm pf} = S_{\rm low} + S_{\rm high} + S_{\rm LU} \ , \tag{5}$$

$$S_{\rm low} = \phi_1^\dagger \frac{1}{M} (M M^\dagger + \mu^2) \frac{1}{M^\dagger} \phi_1 \ , \tag{6}$$

$$S_{\rm high} = \phi_2^\dagger \frac{1}{M^\dagger M + \mu^2} \phi_2 \ , \tag{7}$$

where ϕ_1 and ϕ_2 are two independent pseudofermion fields. $S_{\rm LU}$ is the additional pure-gauge term resulting from LU preconditioning, while M is the LU-preconditioned fermion matrix. The force due to $S_{\rm low}$, which is sensitive to the small eigenvalues of M, is integrated in the outer MD level of the Omelyan integrator with n_1 steps per trajectory. The next level, with n_2 steps, integrates the force due to $S_{\rm high}$, which is sensitive to the large eigenvalues of M, as well as the force coming from $S_{\rm LU}$. All these terms depend on the nHYP links. The force due to the new term, which we will introduce in the next section, is also integrated at this level. Finally, the force due to the Wilson plaquette action (4) is integrated in the innermost level with n_g steps.

We list the ensembles that we use for comparisons in Table I. This table also gives figures for the performance of the HMC algorithm before and after adding the new term, to be discussed below. In order to verify the similarity of the ensembles with and without the new term, we present some results for particle spectra and other physical quantities in Table II.

The crucial diagnostic information for two of our ensembles is presented in Table III. Let us begin with the data that pertain to the original action, that is, $\gamma = 0$. These data, which are maximum and average impulse before and after the chain rule, point to what has to be improved. The first section of three rows, labeled as "fat,high," provides information on the fat-link impulse resulting from $S_{\rm high}$. Next, the "fat,low" section gives

[1] In the numerical implementation we modify $Q = \Omega^\dagger \Omega + \zeta$, with $\zeta = 10^{-6}$, to avoid accidental crashes.

TABLE II: Physical properties of the ensembles. m_q is the quark mass as determined from the unimproved axial Ward identity [3]. r_1 is the larger Sommer scale [15, 16], while m_π and m_ρ are the pseudoscalar and vector meson masses respectively. All results are given in lattice units.

γ	β	m_q	r_1	$r_1 m_\pi^2/m_q$	$(m_\pi/m_\rho)^2$
0	9.6	0.0588(4)	3.04(6)	9.1(2)	0.43(1)
0	9.65	0.0471(4)	3.53(14)	12.2(5)	0.44(1)
0.125	7.8	0.0484(5)	3.12(5)	9.7(2)	0.40(1)
0.25	5.6	0.0575(3)	3.22(5)	10.7(2)	0.48(1)

TABLE III: Maximal and average impulse per trajectory before and after the chain rule. Results are shown for accepted and for rejected trajectories separately. We omit the standard deviation if it is less than 1%.

		$\gamma = 0,\ \beta = 9.6$		$\gamma = 1/4,\ \beta = 5.6$	
		accept	reject	accept	reject
fat,high	max	0.0783	0.0787	0.211	0.211
	avg	0.0312	0.0312	0.0847	0.0848
	max/avg	2.51	2.52	2.49	2.49
fat,low	max	0.093(1)	0.091(2)	0.119(2)	0.112(4)
	avg	0.0144	0.0144	0.0198	0.0198
	max/avg	6.4	6.3(1)	6.0(1)	5.7(2)
thin	max	0.40(2)	0.80(8)	0.87(2)	0.95(7)
	avg	0.0134	0.0134	0.0708	0.0708
	max/avg	30(1)	60(6)	12.3(3)	13(1)

information on the fat-link impulse from $S_{\rm low}$. Last, the "thin" section gives information on the thin-link impulse resulting from the total fat-link impulse after the application of the fat-to-thin chain rule.

In each section, the first row gives the maximal impulse per link in each MD trajectory, averaged over trajectories, separately for accepted and for rejected trajectories. The next row similarly gives the average impulse. The third row gives the ratio of mean maximal impulse to average impulse.

The main thing to notice about the fermions' fat-link impulses of the $\gamma = 0$ ensemble is that they are exactly the same for accepted and for rejected trajectories. In other words, there is absolutely no correlation between the fermions' fat-link impulses and the result of the Metropolis test of the trajectory.

By contrast, the thin-link impulses of the $\gamma = 0$ ensemble exhibit a clear distinction between accepted and rejected trajectories. For accepted trajectories the max/avg ratio is about 30, whereas for rejected trajectories it is twice as big. Histograms of the maximal thin impulse can be found in the upper row of Fig. 1. The difference in the mean value of the maximal impulse between the left and right panels is clearly visible. Also the shapes of the two distributions are quite different.

Our hypothesis is that when the bare coupling is strong enough, dislocations in the dynamical gauge field become abundant. Sometimes, such dislocations will give rise to exceptionally small eigenvalues of the matrix Q of Eq. (2). Through the fat-to-thin chain rule (3), the small eigenvalues of Q generate spikes in the thin-link impulse, which, in turn, results in a bigger probability for failing the Metropolis test at the end of the trajectory.

It is obvious that a large impulse in the final, thin-link force causes rejection of a trajectory. What is new is our observation that there is no large impulse in the initial, fat-link calculation. Evidently the problem lies in the chain rule. What contributes to the severity of this problem is that even a single spike for a single link at a single update step of the whole trajectory, if it is too big, has the potential of producing such a violation of MD energy conservation that will result in failing the Metropolis test. The question is whether we can do something about it.

III. DISLOCATION-SUPPRESSING ACTION FOR nHYP LINKS

In four dimensions, nHYP links $V_{x,\mu}$ are constructed from the dynamical gauge field $U_{x,\mu}$ via three successive smearing steps [1, 2]. Each step consists of first constructing a weighted sum over staples, which is then reunitarized. Explicitly,

$$\Omega_{x,\rho;\xi} = (1-\alpha_3)U_{x,\rho} + \frac{\alpha_3}{2}\left(U_{x,\xi}U_{x+\hat{\xi},\rho}U^\dagger_{x+\hat{\rho},\xi} + U^\dagger_{x-\hat{\xi},\xi}U_{x-\hat{\xi},\rho}U_{x-\hat{\xi}+\hat{\rho},\xi}\right) , \tag{8a}$$

$$\bar{V}_{x,\rho;\xi} = \mathcal{P}(\Omega_{x,\rho;\xi}) ,$$

$$\bar{\Omega}_{x,\mu;\nu} = (1-\alpha_2)U_{x,\mu} + \frac{\alpha_2}{4}\sum_{\substack{\rho\neq\mu,\nu\\ \xi\neq\mu,\nu,\rho}}\left(\bar{V}_{x,\rho;\xi}\bar{V}_{x+\hat{\rho},\mu;\xi}\bar{V}^\dagger_{x+\hat{\mu},\rho;\xi} + \bar{V}^\dagger_{x-\hat{\rho},\rho;\xi}\bar{V}_{x-\hat{\rho},\mu;\xi}\bar{V}_{x-\hat{\rho}+\hat{\mu},\rho;\xi}\right) , \tag{8b}$$

$$\tilde{V}_{x,\mu;\nu} = \mathcal{P}\left(\bar{\Omega}_{x,\mu;\nu}\right) ,$$

$$\tilde{\Omega}_{x,\mu} = (1-\alpha_1)U_{x,\mu} + \frac{\alpha_1}{6}\sum_{\nu\neq\mu}\left(\tilde{V}_{x,\nu;\mu}\tilde{V}_{x+\hat{\nu},\mu;\nu}\tilde{V}^\dagger_{x+\hat{\mu},\nu;\mu} + \tilde{V}^\dagger_{x-\hat{\nu},\nu;\mu}\tilde{V}_{x-\hat{\nu},\mu;\nu}\tilde{V}_{x-\hat{\nu}+\hat{\mu},\nu;\mu}\right) , \tag{8c}$$

$$V_{x,\mu} = \mathcal{P}\left(\tilde{\Omega}_{x,\mu}\right) .$$

The reunitarization operator $\mathcal{P}$ is defined in Eq. (1). Keeping track of this construction in reverse order, one can see that the staple sum extends into a different direction at each smearing step. The outcome is that a given fat link $V_{x,\mu}$ depends on a particular thin link $U_{y,\nu}$ if and only if there exists a hypercube to which both $V_{x,\mu}$ and $U_{y,\nu}$ belong.[2]

We are now ready to introduce the dislocation-suppressing action for nHYP smearing. This is done by adding to the pure-gauge action S_g a new term,

$$S_g = S_{\rm plaq} + S_{\rm NDS} , \tag{9}$$

where the new term is

$$S_{\rm NDS} = \frac{1}{2N_c}\sum_x {\rm tr}\left(\gamma_1\sum_\mu \tilde{Q}^{-1}_{x,\mu} + \gamma_2\sum_{\mu\neq\nu}\bar{Q}^{-1}_{x,\mu;\nu} + \gamma_3\sum_{\rho\neq\xi}Q^{-1}_{x,\rho;\xi}\right) . \tag{10}$$

The motivation for introducing the nHYP Dislocation Suppressing action, or NDS action for short, is clear. The chain rule can produce spikes in the thin-link force associated with small eigenvalues of $Q_{x,\rho;\xi}$, $\bar{Q}_{x,\mu;\nu}$ or $\tilde{Q}_{x,\mu}$. The NDS action is designed to suppress them, by creating a repulsive potential that is proportional to the sum of inverse eigenvalues of the Q matrices.

If we were to add the NDS action $S_{\rm NDS}$ to the usual plaquette action while holding β fixed, we would be pushed back into weaker coupling, and smaller lattice spacing. From the weak-coupling expansion $U_{x,\mu} = \exp(iaA_{x\mu})$ we obtain the bare coupling as

$$\frac{1}{g_0^2} = \frac{\beta}{2N_c} + \frac{1}{N_c}\left(\frac{\gamma_1\alpha_1}{3} + \gamma_2\alpha_2 + \gamma_3\alpha_3\right) . \tag{11}$$

The crucial question, which can only be addressed by performing numerical tests, is whether the NDS action can improve the performance of the HMC algorithm under the same physical conditions. This question will be studied in the next section.

In concluding this section we note that $S_{\rm NDS}$ is easily implemented in the existing code. Using the generic notation of Sec. II, first, $Q^{-1/2}$ is needed for the construction of the nHYP links, so one obtains $Q^{-1} = Q^{-1/2}Q^{-1/2}$ with basically no extra cost. Also, for the calculation of the force, we have

$$\frac{\partial}{\partial U}{\rm tr}\, Q^{-1} = 2\,{\rm tr}\left(Q^{-1/2}\frac{\partial Q^{-1/2}}{\partial U}\right) . \tag{12}$$

Once again, as can be seen from Eq. (3), both $Q^{-1/2}$ and $\partial Q^{-1/2}/\partial U$ have already been calculated, and so it is trivial to obtain their product.

IV. IMPROVEMENT OF MOLECULAR DYNAMICS UPDATE

In our numerical work we use the following values for the smearing parameters [1, 2]

$$(\alpha_1, \alpha_1, \alpha_1) = (0.75, 0.6, 0.3) . \tag{13}$$

Also, we have limited our numerical tests of $S_{\rm NDS}$ to the case

$$\gamma_1 = \gamma_2 = \gamma_3 = \gamma , \tag{14}$$

where the values of γ are shown in the first column of Table I. We have two ensembles without the NDS action, a $\beta = 9.65$ ensemble with a slightly weaker bare coupling

[2] Like the original thin links, the nHYP links $V_{x,\mu}$ reside in the fundamental representation. In our work on higher-representation fermions we first construct the nHYP links $V_{x,\mu}$, and then apply the appropriate group theoretic formulae to construct links in the desired representation from $V_{x,\mu}$. This also adds a step to the chain rule in calculating the MD force.

and a $\beta = 9.6$ ensemble with a slightly stronger bare coupling; one ensemble with $\gamma = 1/8$ and $\beta = 7.8$, and one with $\gamma = 1/4$ and $\beta = 5.6$. In trying to achieve the same physical conditions one can never do a perfect job at non-zero lattice spacing. Still, the physical properties listed in Table II show that all four ensembles exhibit reasonably similar physics.[3]

A. The thin-link force and acceptance

The first piece of evidence that the NDS action actually works is the performance figures shown in Table I. While maintaining the good acceptance rate intact, we have been able to reduce n_1 from 16 to 12, and n_2 from 2 to 1, which is a saving of more than 50% in the number of fermion inversions per trajectory.

In order to understand the origin of this improvement we look at the impulses of the ensemble with (γ, β)=(0.25, 5.6) and compare them to those of the (0, 9.6) ensemble, as shown in Table III. First, like the (0, 9.6) ensemble, also in the case of the (0.25, 5.6) ensemble the fat-link impulses are the same for accepted and for rejected trajectories. Once again, the Metropolis test is uncorrelated with the fat-link impulses produced by the pseudofermions.

In comparing the actual values of the impulses between the two ensembles we should keep in mind that the maximal and average impulses reflect the different numbers of update steps chosen in the two case. However, the different time increments cancel out in the max/avg ratios. Indeed, these ratios turn out to be equal in all four cases: they are basically the same for accepted and rejected trajectories, as well as for the (0, 9.6) and (0.25, 5.6) ensembles. This shows that the fermion sector did not bias the acceptance rate one way or another.

The main difference between the two ensembles is revealed in the thin-link impulse, in the bottom section of Table III. While in the case of the (0, 9.6) ensemble the max/avg ratios were 30 and 60 for accepted and for rejected trajectories respectively, in the case of the (0.25, 5.6) ensemble they are ~ 12 for both accepted and rejected trajectories. The NDS action has produced thin-link impulse ratios that are, first, smaller, and second, uncorrelated with the Metropolis test. These features are also seen in the histograms in the bottom panels of Fig. 1. Indeed, the two histograms of the (0.25, 5.6) ensemble have essentially the same shape.

The conclusion is that, while maintaining roughly the same physical conditions, the NDS action with $\gamma = 1/4$ successfully removes virtually all of the spikes of the thin-link impulse that resulted from the fat-to-thin chain rule. This is one of the main results of this paper.

A complementary observation is the following. Unlike the maximal thin-link impulse, the average thin-link impulse is the same for accepted and for rejected trajectories in both the (0, 9.6) and (0.25, 5.6) ensembles. However, this average value is more than 5 times bigger in the case of the (0.25, 5.6) ensemble than in the (0, 9.6) ensemble. The removal of the high-end tail of the distribution of the thin-link impulse by the NDS action has allowed us to increase the average impulse (by decreasing the numbers of steps), without harming the acceptance rate.

B. Integrator instability and "safety trajectories"

With $n_2 = 1$ in the $\gamma \neq 0$ ensembles of Table I, an interesting practical question is how aggressively can one reduce the number of steps of the outer level, n_1. For example, how would the acceptance rate change if n_1 is further decreased from 12 to 8?

We have carried out an exploratory study of this question on ensembles with parameter values that are similar to (but not necessarily identical with) those of the $\gamma \neq 0$ ensembles of Table I. Our main finding is that if we keep lowering n_1, at some point we will run into a situation where the HMC update experiences occasional, but long, sequences of rejections. The obvious first thought would have been that the fat-to-thin chain-rule spikes of the thin-link impulse are back. However, an examination of the pattern of impulses leads to a different picture. First, the long sequences of rejections are typically characterized by spikes in MD energy non-conservation as large as $\Delta S = O(100)$. Second, an examination of the MD time histories reveals that the occurrence of spikes of ΔS is virtually always correlated with (much smaller) spikes of both the maximal and the average fat-link force coming from $S_{\rm low}$.

The conclusion is that we are looking at a familiar integrator instability. The breakdown of the MD integration was nicely exemplified in the case of a free harmonic oscillator in Ref. [17]. If ω is the frequency of the oscillator, and $\delta\tau$ is the time increment of the (leapfrog) update, the breakdown occurs when the product $\omega\,\delta\tau$ exceeds a critical value that depends on the MD integration scheme.

In our simulations, the time increment $\delta\tau$ was held fixed. Of course, since we are dealing with an interacting field theory, many oscillators are present simultaneously, and their frequencies are changing with the MD evolution. In effect, there is therefore a maximal frequency ω_{max} that scales with λ_{min}^{-2}, where λ_{min}^2 is the smallest eigenvalue of $M^\dagger M$ (see Eq. (6)). We have looked at the low-lying spectrum of $M^\dagger M$ on some of our stuck streams and found that, indeed, the rise in the fat-link force of $S_{\rm low}$ is correlated with the occurrence of an exceptionally small eigenvalue. This, in turn, gives rise to an exceptionally large value of the product $\omega_{max}\,\delta\tau$, and,

[3] In the weak-coupling regime, it follows from Eqs. (11) and (13) that β must be shifted by -2.3γ in order to keep the bare coupling g_0 fixed. As can be seen from Tables I and II, the actual (γ, β) pairs that produce roughly equal physics involve a much bigger shift in β, showing that we are very far from the range of applicability of Eq. (11).

TABLE IV: Properties of the kernel operator [26]. The third column gives the average value of its 10 lowest eigenvalues. The last column gives the number of matrix multiplications needed for the construction of the overlap operator.

γ	β	$\bar{\lambda}$	$N_{\rm op} \times 10^5$
0	9.6	0.106(12)	3.0
0	9.65	0.163(11)	1.7
0.125	7.8	0.182(11)	1.7
0.25	5.6	0.322(15)	1.1

ultimately, to the onset of an integrator instability [17].

The alert reader would notice that the new problem is itself a sign of success. Indeed, it is the very smooth fat-link background, provided by the NDS action, which allows for the Wilson matrix to develop such small eigenvalues that are eventually capable of generating integrator instabilities.

Various solutions to this problem exist in the literature. First, obviously, the simplest solution is to avoid reducing the number of steps too much. The high acceptance rates of the $\gamma \neq 0$ ensembles reported in Table I suggest that, with $n_1 = 12$, we did not run into any integrator instabilities. This is confirmed by an examination of the histories of these runs. The (0.25,5.6) ensemble shows no ΔS spikes at all. Perhaps because of its smaller γ, the (0.125,7.8) ensemble has a few spikes, but none of them has generated a sequence of rejections.

One can do still better by adopting the strategy of Ref. [18]. According to this strategy, one uses a relatively small number of steps for most trajectories. Every once in a while, a larger number of steps is used for a "safety trajectory." The idea is that, in case the simulation has run into a sequence of rejections resulting from an integrator breakdown, that sequence will terminate at the next safety trajectory, where, thanks to its finer time increment, the trajectory will (very likely) be accepted.

We have found that, as long as integrator instabilities are rare, even a modest increase in n_1 is usually enough to eliminate all of them. As an example, the already noted high acceptance rates of the $\gamma \neq 0$ ensembles of Table I suggest that we might use $n_1 = 12$ only for the safety trajectories, while using a smaller number of steps, perhaps $n_1 = 8$, for most trajectories. The interval between two safety trajectories might be taken to be 5 or 10 trajectories.[4] The question of what is the optimal combination invites study, but it is clear that the insertion of safety trajectories is a very cheap cure for the instability problem.

[4] Reversibility of the MD update requires that the interval between two safety trajectories will be fixed beforehand.

V. IMPROVEMENT OF VALENCE CHIRAL FERMIONS

Chiral fermions—domain-wall fermions and overlap fermions—are widely used nowadays [19–21]. While domain-wall fermions are used both as dynamical [10, 22] and as valence fermions, overlap fermions are mostly used as valence fermions (see, however, Ref. [23]).

These chiral fermions are built from a kernel K, which is supercritical Wilson-like (hermitian) operator. Ideally, the kernel would have a spectral gap. In reality, there is never a clean gap. Instead the kernel operator has a mobility edge that is at $O(1)$ in lattice units, with a localized spectrum below the mobility edge and an extended spectrum above it [24]. The near-zero spectrum of localized eigenmodes is always undesirable. In the case of domain-wall fermions it is a dominant source for the residual mass, which is a measure of the imperfection of the chiral symmetry of the domain-wall system. In the case of overlap fermions, such near-zero eigenmodes need to be deflated during the construction of the overlap operator itself. When more of them are present, this makes the numerical construction more expensive and/or less accurate.

Since it is localized, a near-zero eigenmode of the kernel operator often owes its existence to a dislocation in the gauge field [25]. Now, the NDS action suppresses a certain family of dislocations, and so it is interesting to study whether it has any effect on the behavior of chiral fermions. As we will see, we indeed find a clear improvement.

We have used nHYP links to construct the kernel operator introduced in Ref. [26], and studied its properties on our set of ensembles. The third column of Table IV gives the average value $\bar{\lambda}$ of the 10 lowest kernel eigenvalues $|\lambda_i|$, $i = 1, \ldots, 10$. We see that $\bar{\lambda}$ grows with γ, and that, for $\gamma = 1/4$, it is significantly larger than in the other cases. This shows that the dislocation-suppressing effect of the NDS action also helps in reducing the number of near-zero eigenvalues of the kernel operator. This effect is also seen in Fig. 2, which shows histograms of the same 10 lowest kernel eigenvalues. Moreover, the depletion of the near-zero spectrum speeds up the numerical construction of the (valence) overlap operator. This can be seen in the last column of Table IV, which shows the average number of matrix multiplications by the kernel K that is needed for the construction process to converge to a given precision. For valence domain-wall fermions, we would correspondingly expect a reduction of the residual mass at a fixed size of the fifth dimension.

VI. DISCUSSION

The chain rule relating a force with respect to nHYP links to a force with respect to the dynamical links can give rise to relatively rare, but large, spikes of the total impulse, which, in turn, degrade the performance of the

HMC algorithm. In this paper we propose a new term in the gauge action aimed to suppress such spikes. We have shown that, for (approximately) fixed physical parameters, the new term indeed improves the performance of the HMC algorithm. As a side benefit, it improves the behavior of chiral valence fermions.

We have attributed the chain-rule spikes in the force to dislocations in the gauge field—a somewhat vague term. Indeed we have identified the dislocations, operationally, by the very fact some of the Q matrices (Eqs. (2) and (10)) have exceptionally small eigenvalues. This is analogous to what is customary when dealing with domain-wall or overlap fermions, where the presence of a dislocation is identified by the existence of a near-zero (localized) eigenvalue in the spectrum of the kernel operator. What is common to both cases is that one is not interested in the general roughness of the gauge field *per se*, but rather in concrete undesirable effects that this roughness can produce.

Over the years, a large body of work has been devoted to improving the performance of chiral fermions. While it is beyond the scope of this paper to review all this work, we would like to draw some useful lessons.[5] A common way to improve the behavior of chiral fermions is to use so-called improved gauge actions. Examples include various variants of the Symanzik action, the Iwasaki action, and the DBW2 action. Each of these actions consists of a sum over a few Wilson loops. The relative weight of each Wilson loop is fixed either by finding an approximate solution of a truncated renormalization group transformation, or by demanding the elimination of the leading discretization effects perturbatively. Generally speaking, these improved actions produce a smoother gauge field than the simple plaquette action (4). The decreased roughness of the gauge field typically gives rise to fewer near-zero modes in the kernel's spectrum [24, 27].

One method designed to suppress the near-zero eigenvalues of the chiral fermions' kernel K is the so-called Dislocation-Suppressing Determinant Ratio (DSDR) [10, 28]. The basic idea is that, if we were to add to the gauge action the term

$$-\log\det(K^\dagger K)\ , \tag{15}$$

this would produce a logarithmically divergent repulsive potential that entirely suppresses any exact zero modes of the kernel operator. In practice, using the term (15) also has undesirable effects, and so, instead, one replaces $K^\dagger K$ in the above expression by a certain rational polynomial of $K^\dagger K$.

The NDS action (10) is analogous to DSDR in that it targets those dislocations that are responsible for a specific undesirable feature. Now, such dislocations represent local, lattice-size structures in the dynamical gauge field that do not scale. Hence, there is no particular reason to fix the weight of the dedicated, NDS term relative to other terms in the gauge action, and it might be more natural to hold fixed the absolute coefficient of the NDS term while varying the coefficient of the plaquette action (or of any improved action, if one is being used). This is what is being done in effect in the domain-wall simulations of the RBC and UKQCD collaborations: a fixed DSDR term is used while the coefficient of the Iwasaki gauge action is being varied.

Such technical details need not obscure the basic fact that all of the various types of improvement usually play in concert, as was found in the context of domain-wall and overlap fermions [10, 28]. A new example is what we have found in this paper: The NDS action, designed specifically to remove chain-rule spikes in the force for nHYP smearing, also reduces the density of near-zero kernel eigenvalues for chiral fermions.

An alternative way to avoid the chain-rule spikes of the force is to reduce the values of the smearing parameters [see Eq. (13)]. It was observed in Ref. [29] that the matrix Q is positive definite if all the smearing parameters are smaller than 0.5. The dislocations will still be there, however, and their effect on chiral valence fermions will be undiminished. Moreover, weakening the smearing will destroy some of its benefit for approaching the continuum limit.

We mentioned HISQ fermions [11, 12] , which also incorporate the reunitarization step (1). Indeed the same correlation between small eigenvalues of the matrix Q and peaks of the MD force has been reported in this case.[6] It is likely that building a dislocation-suppressing action adapted for this type of smearing would similarly improve the performance of HISQ simulations.

Acknowledgments

We thank Anna Hasenfratz for useful discussions. This work was supported in part by the Israel Science Foundation under grant no. 449/13, and by the U. S. Department of Energy under grant DE-FG02-04ER41290. Our computer code is based on the publicly available package of the MILC collaboration [30]. The code for hypercubic smearing was adapted from a program written by A. Hasenfratz, R. Hoffmann and S. Schaefer [2].

[1] A. Hasenfratz and F. Knechtli, "Flavor symmetry and the static potential with hypercubic blocking," Phys. Rev. D **64**, 034504 (2001) [hep-lat/0103029].

[2] A. Hasenfratz, R. Hoffmann and S. Schaefer, "Hypercubic smeared links for dynamical fermions," JHEP **0705**, 029 (2007) [arXiv:hep-lat/0702028].

[3] T. DeGrand, Y. Shamir and B. Svetitsky, "Running coupling and mass anomalous dimension of

SU(3) gauge theory with two flavors of symmetric-representation fermions," Phys. Rev. D **82**, 054503 (2010) [arXiv:1006.0707 [hep-lat]].

[4] Y. Shamir, B. Svetitsky and E. Yurkovsky, "Improvement via hypercubic smearing in triplet and sextet QCD," Phys. Rev. D **83**, 097502 (2011) [arXiv:1012.2819].

[5] S. Duane and J. B. Kogut, "Hybrid Stochastic Differential Equations Applied to Quantum Chromodynamics," Phys. Rev. Lett. **55**, 2774 (1985); S. Duane and J. B. Kogut, "The Theory Of Hybrid Stochastic Algorithms," Nucl. Phys. B **275**, 398 (1986); S. A. Gottlieb, W. Liu, D. Toussaint, R. L. Renken and R. L. Sugar, "Hybrid Molecular Dynamics Algorithms for the Numerical Simulation of Quantum Chromodynamics," Phys. Rev. D **35**, 2531 (1987).

[6] M. Hasenbusch, "Speeding up the Hybrid-Monte-Carlo algorithm for dynamical fermions," Phys. Lett. B **519**, 177 (2001) [arXiv:hep-lat/0107019].

[7] C. Urbach, K. Jansen, A. Shindler and U. Wenger, "HMC algorithm with multiple time scale integration and mass preconditioning," Comput. Phys. Commun. **174**, 87 (2006) [arXiv:hep-lat/0506011].

[8] T. Takaishi and P. de Forcrand, "Testing and tuning new symplectic integrators for hybrid Monte Carlo algorithm in lattice QCD," Phys. Rev. E **73**, 036706 (2006) [arXiv:hep-lat/0505020].

[9] H. Yin and R. D. Mawhinney, "Improving DWF Simulations: the Force Gradient Integrator and the Móbius Accelerated DWF Solver," PoS LATTICE **2011**, 051 (2011) [arXiv:1111.5059 [hep-lat]].

[10] R. Arthur *et al.* [RBC and UKQCD Collaborations], "Domain Wall QCD with Near-Physical Pions," Phys. Rev. D **87**, 094514 (2013) [arXiv:1208.4412 [hep-lat]].

[11] A. Bazavov *et al.* [MILC Collaboration], "HISQ action in dynamical simulations," PoS LATTICE **2008**, 033 (2008) [arXiv:0903.0874 [hep-lat]]; "Progress on four flavor QCD with the HISQ action," PoS LAT **2009**, 123 (2009) [arXiv:0911.0869 [hep-lat]];

[12] A. Bazavov *et al.* [MILC Collaboration], "Scaling studies of QCD with the dynamical HISQ action," Phys. Rev. D **82**, 074501 (2010) [arXiv:1004.0342 [hep-lat]].

[13] C. Morningstar and M. J. Peardon, "Analytic smearing of SU(3) link variables in lattice QCD," Phys. Rev. D **69**, 054501 (2004) [hep-lat/0311018].

[14] S. Capitani, S. Dürr and C. Hoelbling, "Rationale for UV-filtered clover fermions," JHEP **0611**, 028 (2006) [hep-lat/0607006]; T. Kurth *et al.* [Budapest–Marseille–Wuppertal Collaboration], "Scaling study for 2 HEX smeared fermions: hadron and quark masses," PoS LATTICE **2010**, 232 (2010) [arXiv:1011.1780 [hep-lat]].

[15] R. Sommer, "A New way to set the energy scale in lattice gauge theories and its applications to the static force and alpha-s in SU(2) Yang-Mills theory," Nucl. Phys. B **411**, 839 (1994) [hep-lat/9310022].

[16] C. W. Bernard, T. Burch, K. Orginos, D. Toussaint, T. A. DeGrand, C. E. DeTar, S. A. Gottlieb and U. M. Heller *et al.*, "The Static quark potential in three flavor QCD," Phys. Rev. D **62**, 034503 (2000) [hep-lat/0002028].

[17] R. G. Edwards, I. Horvath and A. D. Kennedy, "Instabilities and nonreversibility of molecular dynamics trajectories," Nucl. Phys. B **484**, 375 (1997) [hep-lat/9606004]; B. Joo *et al.* [UKQCD Collaboration], "Instability in the molecular dynamics step of hybrid Monte Carlo in dynamical Fermion lattice QCD simulations," Phys. Rev. D **62**, 114501 (2000) [hep-lat/0005023].

[18] K. Jansen *et al.* [ALPHA Collaboration], "O(alpha) improvement of lattice QCD with two flavors of Wilson quarks," Nucl. Phys. B **530**, 185 (1998) [Erratum-ibid. B **643**, 517 (2002)] [hep-lat/9803017].

[19] D. B. Kaplan, "A Method for simulating chiral fermions on the lattice," Phys. Lett. B **288**, 342 (1992) [hep-lat/9206013].

[20] Y. Shamir, "Chiral fermions from lattice boundaries," Nucl. Phys. B **406**, 90 (1993) [hep-lat/9303005]; V. Furman and Y. Shamir, "Axial symmetries in lattice QCD with Kaplan fermions," Nucl. Phys. B **439**, 54 (1995) [hep-lat/9405004].

[21] H. Neuberger, "Exactly massless quarks on the lattice," Phys. Lett. B **417**, 141 (1998) [hep-lat/9707022].

[22] J. Noaki, S. Aoki, G. Cossu, H. Fukaya, S. Hashimoto and T. Kaneko, "Fine lattice simulations with chirally symmetric fermions," PoS LATTICE **2013**, 263 (2014); T. Kaneko *et al.* [The JLQCD Collaboration], "Large-scale simulations with chiral symmetry," PoS LATTICE **2013**, 125 (2013) [arXiv:1311.6941 [hep-lat]];

[23] S. Aoki, T. -W. Chiu, G. Cossu, X. Feng, H. Fukaya, S. Hashimoto, T. -H. Hsieh and T. Kaneko *et al.*, "Simulation of quantum chromodynamics on the lattice with exactly chiral lattice fermions," PTEP **2012**, 01A106 (2012); H. Fukaya *et al.* [JLQCD Collaboration], "Overlap/Domain-wall reweighting," PoS LATTICE **2013**, 127 (2013) [arXiv:1311.4646 [hep-lat]].

[24] M. Golterman and Y. Shamir, "Localization in lattice QCD," Phys. Rev. D **68**, 074501 (2003) [hep-lat/0306002]; M. Golterman, Y. Shamir and B. Svetitsky, "Mobility edge in lattice QCD," Phys. Rev. D **71**, 071502 (2005) [hep-lat/0407021]; "Localization properties of lattice fermions with plaquette and improved gauge actions," Phys. Rev. D **72**, 034501 (2005) [hep-lat/0503037].

[25] F. Berruto, R. Narayanan and H. Neuberger, "Exact local fermionic zero modes," Phys. Lett. B **489**, 243 (2000) [hep-lat/0006030].

[26] T. A. DeGrand [MILC Collaboration], "A Variant approach to the overlap action," Phys. Rev. D **63**, 034503 (2000) [hep-lat/0007046]; "Kaon B parameter in quenched QCD," Phys. Rev. D **69**, 014504 (2004) [hep-lat/0309026].

[27] D. J. Antonio *et al.* [RBC and UKQCD Collaborations], "Localization and chiral symmetry in three flavor domain wall QCD," Phys. Rev. D **77**, 014509 (2008) [arXiv:0705.2340 [hep-lat]].

[28] P. M. Vranas, *in Lattice Fermions and Structure of the Vacuum*, (Springer, New York, 1999); P. M. Vranas, "Gap Domain Wall Fermions," Phys. Rev. D **74**, 034512 (2006) [hep-lat/0606014]; H. Fukaya *et al.* [JLQCD Collaboration], "Lattice gauge action suppressing near-zero modes of H(W)," Phys. Rev. D **74**, 094505 (2006) [hep-lat/0607020]; D. Renfrew, T. Blum, N. Christ, R. Mawhinney and P. Vranas, "Controlling Residual Chiral Symmetry Breaking in Domain Wall Fermion Simulations," PoS LATTICE **2008**, 048 (2008) [arXiv:0902.2587 [hep-lat]].

[29] A. Cheng, A. Hasenfratz and D. Schaich, "Novel phase in SU(3) lattice gauge theory with 12 light fermions," Phys. Rev. D **85**, 094509 (2012) [arXiv:1111.2317 [hep-lat]].

[30] `http://www.physics.utah.edu/∼detar/milc/`

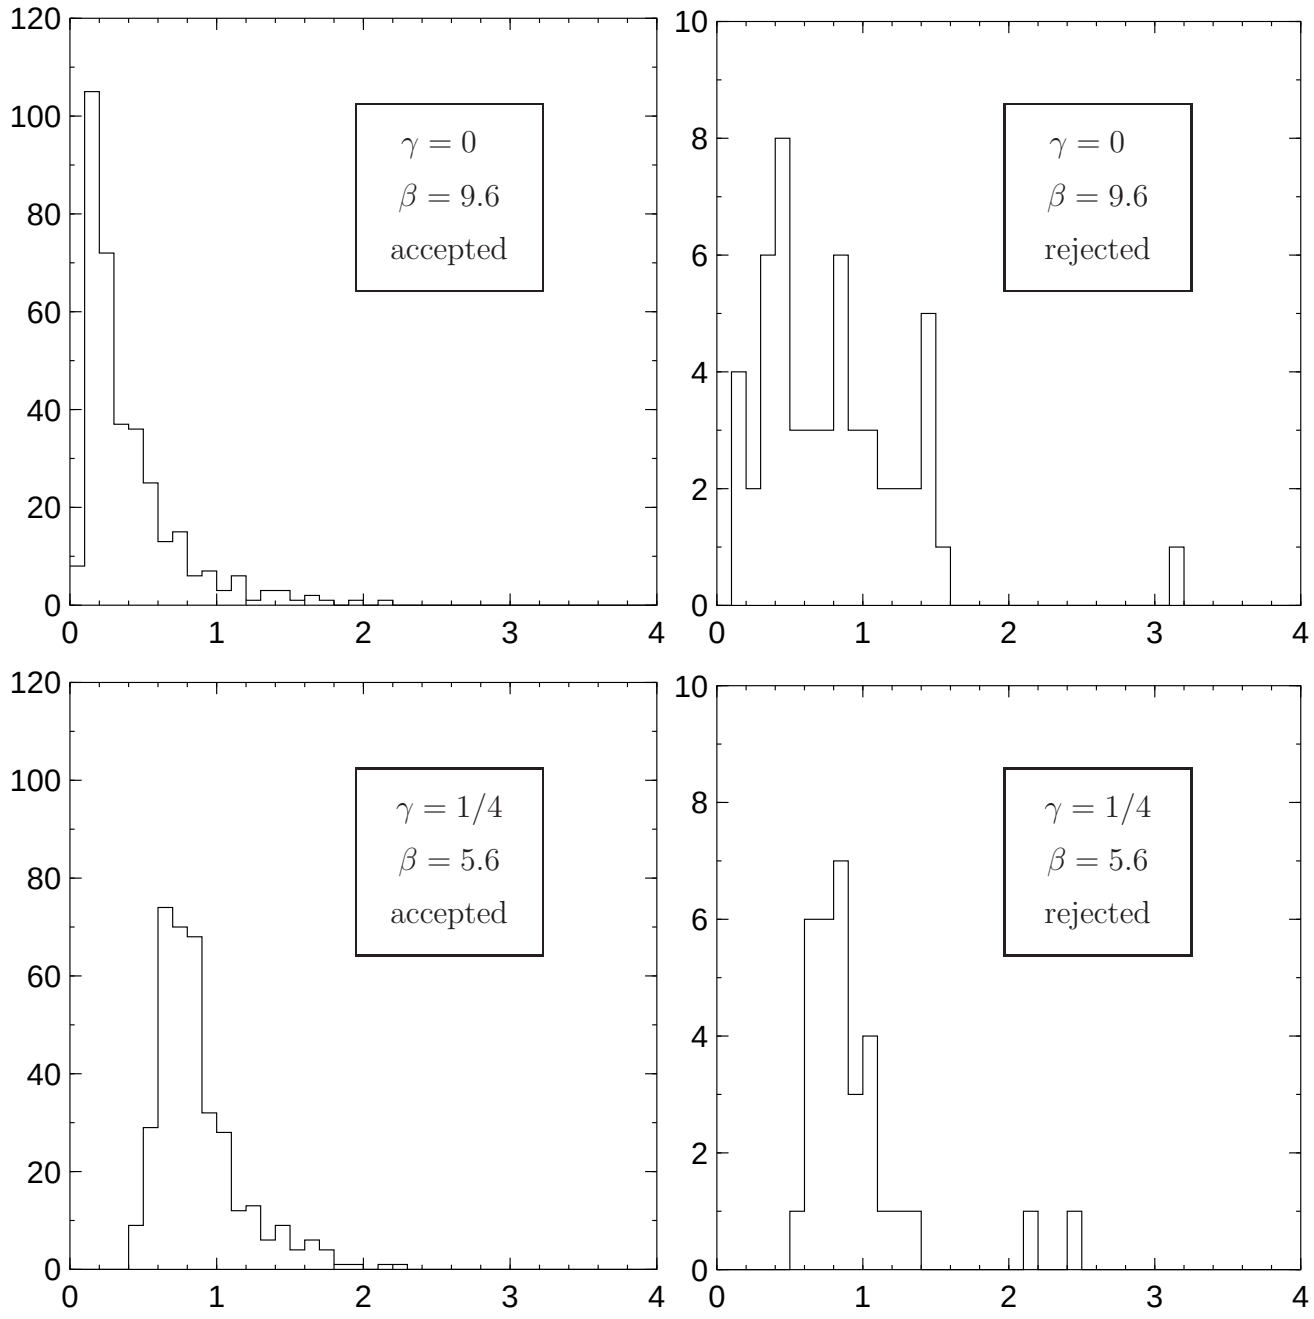

FIG. 1: Histograms of the maximal thin-link force (corresponding to the "thin" section of Table III), for 400 trajectories each. Notice the different vertical scales for accepted (left) and for rejected (right) trajectories.

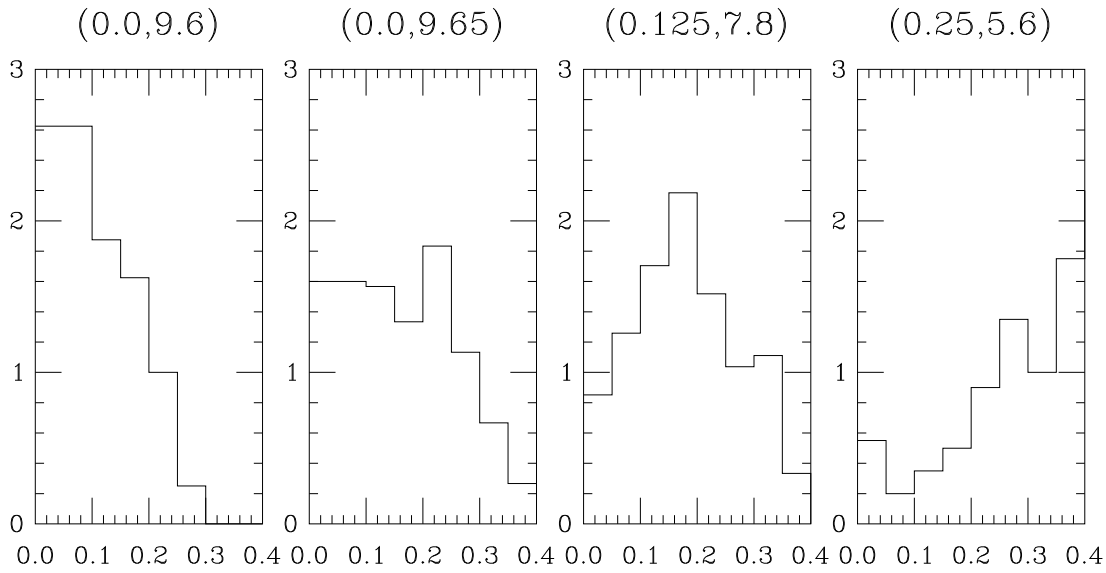

FIG. 2: Histograms of the lowest 10 eigenvalues of the kernel operator. The vertical axis is the average number of eigenvalues per bin, where the bin size is 0.05. For the (γ, β)=(0, 9.6) ensemble, the lowest 10 eigenvalues of all the configurations were within the shown interval, whereas for the other cases, some of these eigenvalues fall outside of this interval. Notice the depletion of near-zero eigenvalues as γ is increased.