

Domain wall fermion QCD with the exact one flavor algorithm

C. Jung,¹ C. Kelly,² R. D. Mawhinney,² and D. J. Murphy²

¹*Department of Physics, Brookhaven National Laboratory, Upton, New York 11973, USA*

²*Department of Physics, Columbia University, New York, New York 10027, USA*



(Received 1 September 2017; published 13 March 2018)

Lattice QCD calculations including the effects of one or more nondegenerate sea quark flavors are conventionally performed using the rational hybrid Monte Carlo (RHMC) algorithm, which computes the square root of the determinant of $\mathcal{D}^\dagger\mathcal{D}$, where \mathcal{D} is the Dirac operator. The special case of two degenerate quark flavors with the same mass is described directly by the determinant of $\mathcal{D}^\dagger\mathcal{D}$ —in particular, no square root is necessary—enabling a variety of algorithmic developments, which have driven down the cost of simulating the light (up and down) quarks in the isospin-symmetric limit of equal masses. As a result, the relative cost of single quark flavors—such as the strange or charm—computed with RHMC has become more expensive. This problem is even more severe in the context of our measurements of the $\Delta I = 1/2$ $K \rightarrow \pi\pi$ matrix elements on lattice ensembles with G -parity boundary conditions, since G -parity is associated with a doubling of the number of quark flavors described by \mathcal{D} , and thus RHMC is needed for the isospin-symmetric light quarks as well. In this paper we report on our implementation of the exact one flavor algorithm (EOFA) introduced by the TWQCD Collaboration for simulations including single flavors of domain wall quarks. We have developed a new preconditioner for the EOFA Dirac equation, which both reduces the cost of solving the Dirac equation and allows us to reuse the bulk of our existing high-performance code. Coupling these improvements with careful tuning of our integrator, the time per accepted trajectory in the production of our $2 + 1$ flavor G -parity ensembles with physical pion and kaon masses has been decreased by a factor of 4.2.

DOI: [10.1103/PhysRevD.97.054503](https://doi.org/10.1103/PhysRevD.97.054503)

I. INTRODUCTION

Lattice QCD simulations are typically performed using variants of the hybrid Monte Carlo (HMC) algorithm, which includes the effects of dynamical sea quarks through the determinant of a fermion matrix evaluated by stochastically sampling a discretized QCD path integral. Conventional simulations choose the Hermitian fermion matrix $\mathcal{M} = \mathcal{D}^\dagger\mathcal{D}$ rather than the lattice Dirac operator $\mathcal{M} = \mathcal{D}$, since the latter, in general, has a complex spectrum, and is thus less amenable to standard numerical algorithms. While \mathcal{D} describes a single quark flavor, $\mathcal{D}^\dagger\mathcal{D}$ describes two degenerate quark flavors with the same mass. As a result the standard HMC algorithm naturally describes the light (up and down) quarks in the isospin-symmetric limit $m_u = m_d$ considered in most lattice calculations. Simulations including single quark flavors (such as the strange or charm) are typically performed by taking an overall square root of the determinant of $\mathcal{M} = \mathcal{D}^\dagger\mathcal{D}$,

leading to the rational hybrid Monte Carlo (RHMC) algorithm. While RHMC has found widespread usage in the lattice QCD community, RHMC calculations are typically more expensive than HMC calculations for the same input quark mass, in part because many of the techniques which have been developed to accelerate HMC simulations of degenerate quark flavor pairs are not applicable to RHMC.

A number of recent developments in the HMC algorithm used by the RBC/UKQCD Collaboration have driven down the cost of simulating degenerate pairs of isospin-symmetric quarks with the same mass. These developments include: extensive force tuning via Hasenbush mass preconditioning [1], the zMöbius domain wall fermion action [2], reduced L_s approximations to the light quark determinant [2], and the use of implicitly restarted, mixed-precision defect correction methods in the conjugate gradient algorithm [3]. In Table I we list timings for a recent large-scale calculation which utilizes these techniques. We now find that the single-flavor strange and charm quark determinants, which we simulate using the RHMC algorithm, are collectively the most expensive part of the calculation. To address this, we have turned to exploring TWQCD's recently proposed exact one flavor algorithm (EOFA), which allows for simulating single

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. 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.

TABLE I. Timings for one HMC trajectory of RBC/UKQCD’s $80^2 \times 96 \times 192 \times 32$ $N_f = 2 + 1 + 1$ ensemble with physical quark masses and $a^{-1} \approx 3$ GeV on a 12 288-node Blue Gene/Q partition [2].

Action Component	Timings	
Gauge	5970 s	12.0%
Light Quarks	19 600 s	39.4%
Strange and Charm Quarks	24 200 s	48.6%
Total	49 770 s	...

quark flavors without the need for RHMC [4]. Preliminary results have suggested that EOFA simulations can outperform RHMC simulations, both in terms of computer time and a reduced memory footprint, while producing exactly the same physics [5,6].

The RBC/UKQCD Collaboration’s ongoing efforts to probe direct CP -violation in $K \rightarrow \pi\pi$ decays provide a second motivation for exploring EOFA. The collaboration has recently reported the first calculation of the $\Delta I = 1/2$ $K \rightarrow \pi\pi$ decay amplitude with physical kinematics in Ref. [7], which, when combined with previous results for the $\Delta I = 3/2$ amplitude [8] determines the standard model CP -violating parameters ϵ and ϵ' entirely from first principles. An important ingredient in this calculation was the introduction of G -parity boundary conditions for the quark fields [7,9]: since the pion is G -parity odd, the pion momenta are quantized along G -parity directions as

$$p_\pi^i = \frac{(2n_i + 1)\pi}{L}, \quad n_i \in \mathbb{Z}, \quad (1)$$

allowing the ensemble parameters to be tuned such that the $K \rightarrow \pi\pi$ decay has both physical kinematics and the final pions in the ground state. Since the G -parity transformation $G = C e^{i\pi I_y}$ is the product of charge conjugation and a 180° isospin rotation about the y -axis—at the lattice boundary the light quark doublet transforms as $(u, d) \mapsto (\bar{d}, -\bar{u})$ —the G -parity Dirac operator inherently describes two quark flavors. The standard lattice technique for obtaining a Hermitian, positive-definite fermion matrix—by taking the square of the Dirac operator, $\mathcal{M} = \mathcal{D}^\dagger \mathcal{D}$ —results in a theory with four degenerate quark flavors on a G -parity ensemble, and a square root is required to reduce to a two-flavor simulation. Describing the light quark pair on a G -parity ensemble is a particularly attractive target for EOFA, since many of the techniques we use to accelerate the calculation of the light quark determinant for ensembles with periodic boundary conditions—including defect correction solvers, the forecasted force gradient integrator [10], and Hasenbusch mass preconditioning—are not applicable or of limited utility for RHMC simulations, but are expected to perform well in the context of EOFA. More generally, since there is no straightforward way to start the multishift conjugate gradient solver used for

RHMC with a nonzero initial guess, techniques which rely on forecasting or restarting the solver are not applicable.

In this work we discuss the RBC/UKQCD Collaboration’s implementation and tests of the exact one flavor algorithm, as well as the use of EOFA in generating gauge field configurations for our ongoing first-principles calculation of the ratio of standard model parameters ϵ'/ϵ from $\Delta I = 1/2$ $K \rightarrow \pi\pi$ decays with G -parity boundary conditions. We have independently implemented EOFA in the Columbia Physics System (CPS), BAGEL fermion sparse matrix library (BFM), and the Grid data parallel C++ QCD library (Grid), for Shamir and Möbius domain wall fermions, with periodic, anti-periodic, and G -parity boundary conditions. We will demonstrate in the following sections that a significant improvement over the RHMC algorithm in terms of wall clock time is indeed possible with EOFA after introducing a variety of preconditioning and tuning techniques. Early work in this direction was presented at the 34th International Symposium on Lattice Field Theory [6]; here we will elaborate on the details and discuss our first large-scale EOFA calculation.

II. THE EXACT ONE FLAVOR ALGORITHM

The exact one flavor algorithm was developed by the TWQCD Collaboration and used to enable efficient simulations of single quark flavors on GPU clusters, where memory usage is a significant constraint. In Ref. [11] the authors discuss their construction of a positive-definite pseudofermion action describing a single flavor of Wilson or domain wall quark, and elaborate on the details of this construction in Ref. [4]. The key is their observation that a ratio of determinants of domain wall fermion (DWF) Dirac operators can be factorized as

$$\det \left(\frac{\mathcal{D}(m_1)}{\mathcal{D}(m_2)} \right) = \frac{1}{\det(\mathcal{M}_L)} \cdot \frac{1}{\det(\mathcal{M}_R)}, \quad (2)$$

with \mathcal{M}_L and \mathcal{M}_R Hermitian and positive-definite. In a subsequent paper the authors benchmark EOFA against RHMC for $N_f = 1$ and $N_f = 2 + 1$ lattice QCD simulations, and demonstrate a number of advantages of the EOFA formalism [5]. These include substantial reductions in the pseudofermion force and in the memory footprint of the algorithm, since, in the context of EOFA, inversions of the Dirac operator can be performed using the ordinary conjugate gradient (CG) algorithm rather than the multishift CG used for RHMC. They ultimately find that they are able to generate HMC trajectories 15–20% faster using EOFA rather than RHMC after retuning their integration scheme to take advantage of these properties. More recently, TWQCD has used EOFA to generate $N_f = 2 + 1 + 1$ domain wall fermion ensembles with dynamical strange and charm quarks [12].

We note that the construction of the exact one flavor pseudofermion action has been detailed by TWQCD in

Ref. [4,11] and summarized in our own formalism in Ref. [6]. We will not repeat this discussion here, other than to give a brief overview and to introduce the notation used in this work. We write the 5D Möbius domain wall fermion (MDWF) operator \mathcal{D}_{DWF} in terms of the 4D Wilson Dirac operator D_W and 5D hopping matrix $L_{ss'}$ as

$$\begin{aligned} (\mathcal{D}_{\text{DWF}})_{xx',ss'} &= ((c+d)(D_W)_{xx'} + \delta_{xx'})\delta_{ss'} \\ &\quad + ((c-d)(D_W)_{xx'} - \delta_{xx'})L_{ss'} \\ (D_W)_{xx'} &= (4-M_5)\delta_{xx'} - \frac{1}{2} \sum_{\mu} [(1-\gamma_{\mu})U_{\mu}(x)\delta_{x+\hat{\mu},x'} \\ &\quad + (1+\gamma_{\mu})U_{\mu}^{\dagger}(x')\delta_{x-\hat{\mu},x'}] \\ L_{ss'} &= (L_+)_{ss'}P_+ + (L_-)_{ss'}P_- \end{aligned} \quad (3)$$

with

$$(L_+)_{ss'} = (L_-)_{s's} = \begin{cases} -m\delta_{L_s-1,s'}, & s=0 \\ \delta_{s-1,s'}, & 1 \leq s \leq L_s-1 \end{cases}. \quad (4)$$

Here x and s are spacetime indices in the 4D bulk and along the fifth dimension, respectively, with L_s denoting the total number of s sites, $P_{\pm} = (1 \pm \gamma_5)/2$ denoting the chiral projection operators, and $(R_5)_{ss'} \equiv \delta_{s,L_s-1-s'}$ denoting the operator which performs a reflection in the fifth dimension. We recover four-dimensional quark fields q and \bar{q} with definite chiralities from the five-dimensional quark fields ψ and $\bar{\psi}$ described by \mathcal{D}_{DWF} at the boundaries of the fifth dimension

$$\begin{aligned} q_R &= P_+\psi_{L_s-1} & q_L &= P_-\psi_0 \\ \bar{q}_R &= \bar{\psi}_{L_s-1}P_- & \bar{q}_L &= \bar{\psi}_0P_+ \end{aligned}. \quad (5)$$

Green's functions constructed from q and \bar{q} approximate continuum QCD arbitrarily well in the limit of vanishing lattice spacing and infinite 5D spacetime volume. The tunable parameters in Eq. (3) are the domain wall parameter M_5 , the bare quark mass m , and the Möbius scale $\alpha = 2c$; the parameter d is fixed at $d = 1/2$. DWF with the Shamir kernel is recovered from the more general Möbius operator in the limit $\alpha \rightarrow 1$. For more detail regarding our MDWF formalism we direct the reader to Ref. [13].

The construction of the exact one flavor action for domain wall fermions begins by factorizing the MDWF Dirac operator as [14]

$$\mathcal{D}_{\text{DWF}} = \mathcal{D}_{\text{EOFA}} \cdot \tilde{\mathcal{D}}, \quad (6)$$

with

$$\begin{aligned} (\mathcal{D}_{\text{EOFA}})_{xx',ss'} &\equiv (D_W)_{xx'}\delta_{ss'} + \delta_{xx'}(M_+)_{ss'}P_+ \\ &\quad + \delta_{xx'}(M_-)_{ss'}P_- \\ (\tilde{\mathcal{D}})_{ss'} &\equiv d(\delta_{ss'} - L_{ss'}) + c(\delta_{ss'} + L_{ss'}). \end{aligned} \quad (7)$$

The operator $\tilde{\mathcal{D}}$ relating \mathcal{D}_{DWF} and $\mathcal{D}_{\text{EOFA}}$ has no dependence on the gauge field, so we are free to replace \mathcal{D}_{DWF} with $\mathcal{D}_{\text{EOFA}}$ in Eq. (2) without modifying physical observables described by a properly normalized path integral. In fact, it can be shown analytically using the explicit form of $\tilde{\mathcal{D}}$ listed in Appendix A that

$$\det(\tilde{\mathcal{D}}) = ((c+d)L_s + m(c-d)L_s)^{12V}, \quad (8)$$

where $V = L^3T$ is the 4D spacetime volume. This substitution facilitates the construction of a proper action since the operator $\gamma_5 R_5 \mathcal{D}_{\text{EOFA}}$ is manifestly Hermitian for any choice of the Möbius scale α , whereas \mathcal{D}_{DWF} satisfies a less trivial γ_5 -Hermiticity condition when $\alpha \neq 1$ [15]. However, this comes at the cost of substantially more expensive inversions, since $\mathcal{D}_{\text{EOFA}}$ is dense in ss' whereas \mathcal{D}_{DWF} has a well-known tridiagonal block structure.

After introducing $\mathcal{D}_{\text{EOFA}}$, TWQCD's construction proceeds by applying the Schur identity

$$\begin{aligned} \det \left[\begin{pmatrix} A & B \\ C & D \end{pmatrix} \right] &= \det(A) \det(D - CA^{-1}B) \\ &= \det(D) \det(A - BD^{-1}C) \end{aligned} \quad (9)$$

to $\mathcal{D}_{\text{EOFA}}$, treated as a 2×2 block matrix in its spinor indices, and rearranging terms to arrive at the right-hand side of Eq. (2). Crucially, factors of $\gamma_5 R_5$ can be freely inserted under the determinant to replace $\mathcal{D}_{\text{EOFA}}$ with the Hermitian operator $H \equiv \gamma_5 R_5 \mathcal{D}_{\text{EOFA}}$, since $\det(\gamma_5) = \det(R_5) = 1$. The final form of the exact one flavor pseudofermion action is $S_{\text{EOFA}} = \phi^{\dagger} \mathcal{M}_{\text{EOFA}} \phi$, with

$$\begin{aligned} \mathcal{M}_{\text{EOFA}} &\equiv 1 - kP_-\Omega_-^{\dagger}[H(m_1)]^{-1}\Omega_-P_- \\ &\quad + kP_+\Omega_+^{\dagger}[H(m_2) - \Delta_+(m_1, m_2)P_+]^{-1}\Omega_+P_+. \end{aligned} \quad (10)$$

In Appendix A we collect explicit expressions for k , Ω_{\pm} , Δ_{\pm} , $\mathcal{D}_{\text{EOFA}}$, and $\tilde{\mathcal{D}}$ for Shamir and Möbius DWF, since, to the authors' knowledge, these expressions have not previously appeared in the literature. In Ref. [4] these operators are constructed recursively for the more general case of DWF with weights $\rho_s = c\omega_s + d$ and $\sigma_s = c\omega_s - d$ that are allowed to vary along the fifth dimension, subject to the constraint that ω_s is reflection-symmetric in s . Shamir and Möbius DWF are simpler, special cases with $\omega_s = 1$.

TABLE II. Summary of ensembles and input parameters used in this work. Here β is the gauge coupling, $L^3 \times T \times L_s$ is the lattice volume decomposed into the length of the spatial (L), temporal (T), and fifth (L_s) dimensions, and am_l and am_h are the bare, input light and heavy quark masses. On the 16I-G, 16ID-G, and 32ID-G ensembles G -parity boundary conditions are applied to the fermion fields at one or more of the spatial boundaries of the lattice; otherwise periodic boundary conditions are applied, and in all cases antiperiodic boundary conditions are used along the temporal direction.

Ensemble	Action	β	$L^3 \times T \times L_s$	Möbius Scale	G -Parity B.C.	am_l	am_h
16I	DWF + I	2.13	$16^3 \times 32 \times 16$	0.01	0.032
16I-G	DWF + I	2.13	$16^3 \times 32 \times 16$...	x	0.01	0.032
16ID-G	MDWF + ID	1.75	$16^3 \times 32 \times 8$	2.00	x, y, z	0.01	0.045
24ID	MDWF + ID	1.633	$24^3 \times 64 \times 24$	4.00	...	0.007 89	0.085
32ID-G	MDWF + ID	1.75	$32^3 \times 64 \times 12$	2.67	x, y, z	0.0001	0.045

III. SUMMARY OF ENSEMBLES USED IN THIS WORK

The properties of the lattices used in this work are summarized in Tables II and III. In all cases we use the Iwasaki gauge action (I) [16], and on some ensembles supplement this with the dislocation suppressing determinant ratio (DSDR) [17,18]; we abbreviate the combined action including both terms as “ID.” The additional DSDR term is designed to suppress the dislocations of the gauge field associated with tunneling between topological sectors, thereby reducing the degree of residual chiral symmetry breaking. For strong coupling simulations, where these dislocations occur frequently, the DSDR term reduces the costs associated with light quark masses while still maintaining good topological sampling. We simulate $N_f = 2 + 1$ quark flavors using domain wall fermions, with either the Shamir (DWF) [19,20] or Möbius (MDWF) [21–23] kernel. Finally, on ensembles marked “-G” we use G -parity boundary conditions in one or more of the spatial directions.

TABLE III. Summary of spatial volumes, lattice cutoffs, and pion masses in physical units for the ensembles used in this work. All values for the 16I and 32ID-G ensembles are from Refs. [7,24], respectively. On the 16I-G (16ID-G) ensemble we assume the lattice cutoff is the same as the 16I (32ID-G) ensemble since the same action and value of β has been used. The pion masses on the 16I-G and 16ID-G ensembles have been extracted using the fitted value of the lowest energy pion states from Table VIII and the continuum dispersion relation. Finally, the determination of the lattice scale for the 24ID ensemble was performed in Ref. [25], and the determination of the pion mass in Ref. [26].

Ensemble	L (fm)	a^{-1} (GeV)	m_π (MeV)	m_K (MeV)
16I	1.95(5)	1.62(4)	400(11)	527(13)
16I-G	1.95(5)	1.62(4)	388(14)	530(14)
16ID-G	2.29(1)	1.378(7)	429.8(5.0)	603.8(4.5)
24ID	4.82(19)	0.981(39)	137.1(5.5)	495(20)
32ID-G	4.57(2)	1.378(7)	143.1(2.0)	490.6(2.4)

The 16I ensemble was first generated and used to study light meson spectroscopy with domain wall fermions in Ref. [24]. The 16I-G ensemble is identical to the 16I ensemble except for the boundary conditions along the x -direction, which have been changed from periodic to G -parity. Likewise, the parameters of the 16ID-G ensemble have been chosen based on a series of $\beta = 1.75$ DSDR ensembles generated in Ref. [27], but have G -parity boundary conditions in all three spatial directions. Collectively, these three lattices are used as inexpensive, small-volume test ensembles with unphysical, heavy pion masses to perform cross-checks of the EOFA algorithm and its implementation in the BFM and CPS code libraries. The larger 24ID [26] and 32ID-G [28] ensembles have physical pion masses and are currently being generated as part of production RBC/UKQCD calculations.

IV. HYBRID MONTE CARLO WITH EOFA

In lattice QCD correlation functions are computed in terms of a discretized Euclidean path integral

$$\langle \mathcal{O}_1 \cdots \mathcal{O}_n \rangle = \frac{1}{\mathcal{Z}} \int \mathcal{D}U \left(\prod_f \mathcal{D}\psi_f \mathcal{D}\bar{\psi}_f \right) \times (\mathcal{O}_1[U] \cdots \mathcal{O}_n[U]) e^{-S[U, \bar{\psi}_f, \psi_f]}. \quad (11)$$

Here U is the gauge field, ψ_f is the quark field associated with flavor f , and $S[U, \bar{\psi}_f, \psi_f]$ is the action, which decomposes into a sum of contributions from the gauge field, fermions, and possibly other terms (e.g. the dislocation suppressing determinant ratio). To avoid having to deal with anticommuting Grassman variables in a computer, dynamical fermion flavors are integrated out and then reintroduced in terms of bosonic “pseudofermion” fields ϕ as

$$\begin{aligned} \frac{1}{\mathcal{Z}} \int \mathcal{D}\psi \mathcal{D}\bar{\psi} e^{-\bar{\psi} M \psi} &= \det(M) = \frac{1}{\det(M^{-1})} \\ &= \frac{1}{\mathcal{Z}} \int \mathcal{D}\phi \mathcal{D}\phi^\dagger e^{-\phi^\dagger M^{-1} \phi}, \end{aligned} \quad (12)$$

provided M is positive-definite. While pseudofermions can be represented straightforwardly in a computer, they come at the cost of applications of M^{-1} rather than M , which is not typically available in an explicit form. Even after discretization the integration in Eq. (11) is far too expensive to perform directly due to the enormous number of degrees of freedom on a typical lattice. Instead, Monte Carlo techniques are used to ergodically sample a sequence of representative configurations of the gauge field $\{U_i\}$, for which

$$\langle \mathcal{O}_1 \cdots \mathcal{O}_n \rangle \approx \frac{1}{N} \sum_{i=1}^N \mathcal{O}_1(U_i) \cdots \mathcal{O}_n(U_i). \quad (13)$$

The standard Monte Carlo technique used in modern lattice QCD calculations is known as the hybrid Monte Carlo (HMC) algorithm.

HMC generates a Markov chain of gauge field configurations $\{U_i\}$ by evolving a Hamiltonian system in unphysical molecular dynamics (MD) “time.” This Hamiltonian system is constructed by treating $U_\mu(x)$ as a generalized coordinate, introducing an $\mathfrak{su}(3)$ -valued conjugate momentum $\pi_\mu(x)$, and forming the standard Hamiltonian

$$H = \frac{1}{2} \pi^2 + S(U). \quad (14)$$

The associated equations of motion

$$\begin{cases} \partial_\tau U_\mu(x) = \pi_\mu(x) \\ U_\mu(x) \partial_\tau \pi_\mu(x) = -T^a \partial_{x,\mu}^a S(U) \end{cases} \quad (15)$$

can then be integrated using numerical integration techniques. The integration is performed over intervals of length $\Delta\tau$ —referred to as a single MD trajectory—as a sequence of N small steps $\delta\tau$, with $N = \Delta\tau/\delta\tau$. Finite precision integration errors are corrected stochastically with a Metropolis accept/reject step: after every N integration steps by $\delta\tau$ the total change in the Hamiltonian ΔH is computed, and the current gauge field $U'_\mu(x)$ is accepted as the next configuration in the Markov chain with probability

$$P_{\text{accept}} = \min(1, e^{-\Delta H}). \quad (16)$$

One can show that the resulting algorithm satisfies detailed balance provided the scheme used to numerically integrate Eq. (15) is reversible [29]. Ergodicity is achieved by performing a heatbath step each time the integration is restarted to pick a new conjugate momentum $\pi_\mu(x)$, and thus a new trajectory in the phase space $\{(U, \pi)\}$. HMC generates a sequence of gauge field configurations whose statistical independence is governed by the length of each MD trajectory, $\Delta\tau$. The number of MD trajectories separating statistically independent gauge field configurations is

typically determined *ex post facto* by examining the integrated autocorrelation times of representative physical observables.

The fermionic contribution to the Hamiltonian in Eq. (14) introduces a technical obstacle for the HMC algorithm since, generically, the lattice Dirac operator \mathcal{D} has a complex spectrum. Replacing \mathcal{D} with the Hermitian fermion matrix $M = \mathcal{D}^\dagger \mathcal{D}$ in Eq. (12) has a number of advantages. Most importantly, it allows M^{-1} to be applied to pseudofermion vectors using the conjugate gradient algorithm, and it allows for a straightforward pseudofermion heatbath step: at the beginning of each MD trajectory a random Gaussian vector η is drawn according to $P(\eta) \propto \exp(-\eta^\dagger \eta/2)$ and the initial pseudofermion field is seeded as $\phi = \mathcal{D}\eta$, ensuring that ϕ is correctly sampled as $P(\phi) \propto \exp(-\phi^\dagger M^{-1} \phi/2)$. However, the fermion matrix $M = \mathcal{D}^\dagger \mathcal{D}$ describes two degenerate quark flavors with the same mass. Single flavor simulations are typically performed by taking an overall square root of the fermion determinant,

$$\det(\mathcal{D}) = [\det(\mathcal{D}^\dagger \mathcal{D})]^{1/2}. \quad (17)$$

In the pseudofermion formalism applications of the operator $(\mathcal{D}^\dagger \mathcal{D})^{-1/2}$ are approximated by a matrix-valued function $f(\mathcal{D}^\dagger \mathcal{D})$,

$$\begin{aligned} [\det(\mathcal{D}^\dagger \mathcal{D})]^{1/2} &= \frac{1}{\mathcal{Z}} \int \mathcal{D}\phi \mathcal{D}\phi^\dagger e^{-\phi^\dagger (\mathcal{D}^\dagger \mathcal{D})^{-1/2} \phi} \\ &\simeq \frac{1}{\mathcal{Z}} \int \mathcal{D}\phi \mathcal{D}\phi^\dagger e^{-\phi^\dagger f(\mathcal{D}^\dagger \mathcal{D}) \phi} \end{aligned} \quad (18)$$

where $f(x)$ is a suitably constructed approximation to the inverse square root, valid over the spectral range of $\mathcal{D}^\dagger \mathcal{D}$. Variants of the HMC algorithm which construct f from different classes of functions have been proposed and used in the literature; the most common is the rational HMC (RHMC) algorithm [30], where

$$f(x) = \alpha_0 + \sum_{k=1}^N \frac{\alpha_k}{\beta_k + x} \quad (19)$$

is a rational function. While rational functions are in many ways a good choice—they are economical in the sense that the inverse square root can usually be well-approximated by a modest number of terms, and the multishift CG algorithm can be used to efficiently invert $(\mathcal{D}^\dagger \mathcal{D} + \beta_k)$ for all k simultaneously—the additional complexity of evaluating $f(\mathcal{D}^\dagger \mathcal{D})$ and the associated molecular dynamics pseudofermion force makes single flavor RHMC simulations significantly more costly than degenerate two flavor HMC simulations at the same bare quark mass. This additional cost can be largely attributed to the significant linear algebra overhead associated with multishift CG.

EOFA provides an alternative construction of a single-flavor pseudofermion action through Eq. (2): a ratio of fermion determinants can be factorized as a product of two determinants, each of which involves an operator which is Hermitian and positive-definite. This product can then be represented as a path integral over a bosonic pseudofermion field with a two-term action [Eq. (10)]

$$\det\left(\frac{\mathcal{D}_{\text{EOFA}}(m_1)}{\mathcal{D}_{\text{EOFA}}(m_2)}\right) = \frac{1}{\mathcal{Z}} \int \mathcal{D}\phi \mathcal{D}\phi^\dagger e^{-\phi^\dagger \mathcal{M}_{\text{EOFA}} \phi}, \quad (20)$$

leading to an algorithm which is “exact” in the sense that it avoids the numerical approximations required to implement the square root in RHMC [Eq. (18)] and related HMC variants. EOFA is also expected to be somewhat faster than RHMC, since there is no rational approximation entering into evaluations of the Hamiltonian or the pseudofermion force, eliminating the overhead associated with multishift CG. In the remainder of this section we elaborate on the details of the action, heatbath step, and pseudofermion force entering into the Hamiltonian equations of motion [Eq. (15)] for HMC with EOFA.

A. Action

The EOFA action [Eq. (10)] computes a ratio of determinants of $\mathcal{D}_{\text{EOFA}}$ upon integrating out the pseudofermion fields [Eq. (20)]. This ratio can be related to the conventional determinant ratio computed by the RHMC algorithm through Eqs. (6) and (8), leading to the relationship

$$\det\left(\frac{\mathcal{D}_{\text{DWF}}(m_1)}{\mathcal{D}_{\text{DWF}}(m_2)}\right) = \left(\frac{(c+d)^{L_s} + m_1(c-d)^{L_s}}{(c+d)^{L_s} + m_2(c-d)^{L_s}}\right)^{12V} \times \det\left(\frac{\mathcal{D}_{\text{EOFA}}(m_1)}{\mathcal{D}_{\text{EOFA}}(m_2)}\right). \quad (21)$$

We use this relationship as a test of the equivalence of RHMC and EOFA, as well as our implementation of the EOFA action, by stochastically computing the left side of Eq. (21) with the RHMC action

$$\mathcal{M}_{\text{RHMC}} = [\mathcal{D}_{\text{DWF}}^\dagger \mathcal{D}_{\text{DWF}}(m_2)]^{1/4} [\mathcal{D}_{\text{DWF}}^\dagger \mathcal{D}_{\text{DWF}}(m_1)]^{-1/2} \times [\mathcal{D}_{\text{DWF}}^\dagger \mathcal{D}_{\text{DWF}}(m_2)]^{1/4} \quad (22)$$

and the right side with the EOFA action [Eq. (10)] on the same gauge field configuration. Observing that we can, in general, rewrite a determinant as

$$\begin{aligned} \det(\mathcal{M}^{-1}) &= \frac{1}{\mathcal{Z}} \int \mathcal{D}\phi \mathcal{D}\phi^\dagger e^{-\phi^\dagger \mathcal{M} \phi} \\ &= \frac{1}{\mathcal{Z}} \int \mathcal{D}\phi \mathcal{D}\phi^\dagger e^{-\frac{1}{2}\phi^\dagger \Sigma^{-1} \phi} e^{\phi^\dagger (\frac{1}{2}\Sigma^{-1} - \mathcal{M}) \phi} \end{aligned} \quad (23)$$

suggests the following simple Monte Carlo integration scheme: we draw random pseudofermion vectors by

independently sampling the real and imaginary parts of each component from the standard normal distribution $\mathcal{N}(\mu = 0, \sigma = 1)$, and compute the expectation value

$$-\log \det(\mathcal{M}^{-1}) \approx \left\langle \phi_i^\dagger \left(\mathcal{M} - \frac{1}{2} \Sigma^{-1} \right) \phi_i \right\rangle, \quad (24)$$

where the average is computed using the jackknife resampling technique. This will accurately approximate the true log determinant for finite, realistically calculable values of N provided m_1 and m_2 are sufficiently close that the integrand is well-approximated by a Gaussian with unit variance. To address this latter systematic, we consider splitting Eq. (21) as a product of determinants

$$\begin{aligned} \det\left(\frac{\mathcal{D}(m_1)}{\mathcal{D}(m_2)}\right) &= \det\left(\frac{\mathcal{D}(m_1)}{\mathcal{D}(m'_1)}\right) \left[\prod_{i=1}^{N_m} \det\left(\frac{\mathcal{D}(m'_i)}{\mathcal{D}(m'_{i+1})}\right) \right] \\ &\times \det\left(\frac{\mathcal{D}(m_{N_m}')}{\mathcal{D}(m_2)}\right) \end{aligned} \quad (25)$$

with equally-spaced intermediate masses

$$m'_i = m_1 + \frac{m_2 - m_1}{N_m + 1} i, \quad i = 1, \dots, N_m, \quad (26)$$

and study the dependence of the result on N_m (this procedure is identical to the method introduced in Ref. [31] for computing quark mass reweighting factors). In the upper panel of Fig. 1 we plot the log determinants of $\mathcal{M}_{\text{RHMC}}^{-1}$ and $\mathcal{M}_{\text{EOFA}}^{-1}$ as a function of N_m , with $N = 10$ stochastic evaluations, computed using a single thermalized trajectory of the 16I, 16I-G, and 16ID-G ensembles. For the case of the 16ID-G ensemble, which uses the Möbius DWF fermion action, we also include the overall constant multiplying the right side of Eq. (21) so that in all cases we are computing the same determinant ratio of \mathcal{D}_{DWF} using either action.

We observe, as expected, that both formalisms agree for sufficiently large N_m . Likewise, we observe that at sufficiently small N_m the calculation generally becomes unreliable since we do not attempt to account for the systematic error associated with approximating the integrand of the path integral by a Gaussian with unit variance [i.e., setting $\Sigma = 1$ in Eqs. (23) and (24)]. In both cases “sufficiently” small or large N_m is controlled by the size of the splitting between m_1 and m_2 . We also observe that, for a given choice of N_m and N , both the statistical and systematic errors of the determinant ratio computed via EOFA are suppressed relative to the errors of the determinant ratio computed via RHMC. We argue that the observed error suppression can be explained by comparing the spectrum of $\mathcal{M}_{\text{RHMC}}$ to the spectrum of $\mathcal{M}_{\text{EOFA}}$, which we plot in the lower panels of Fig. 1 for a very small lattice volume (4^5)

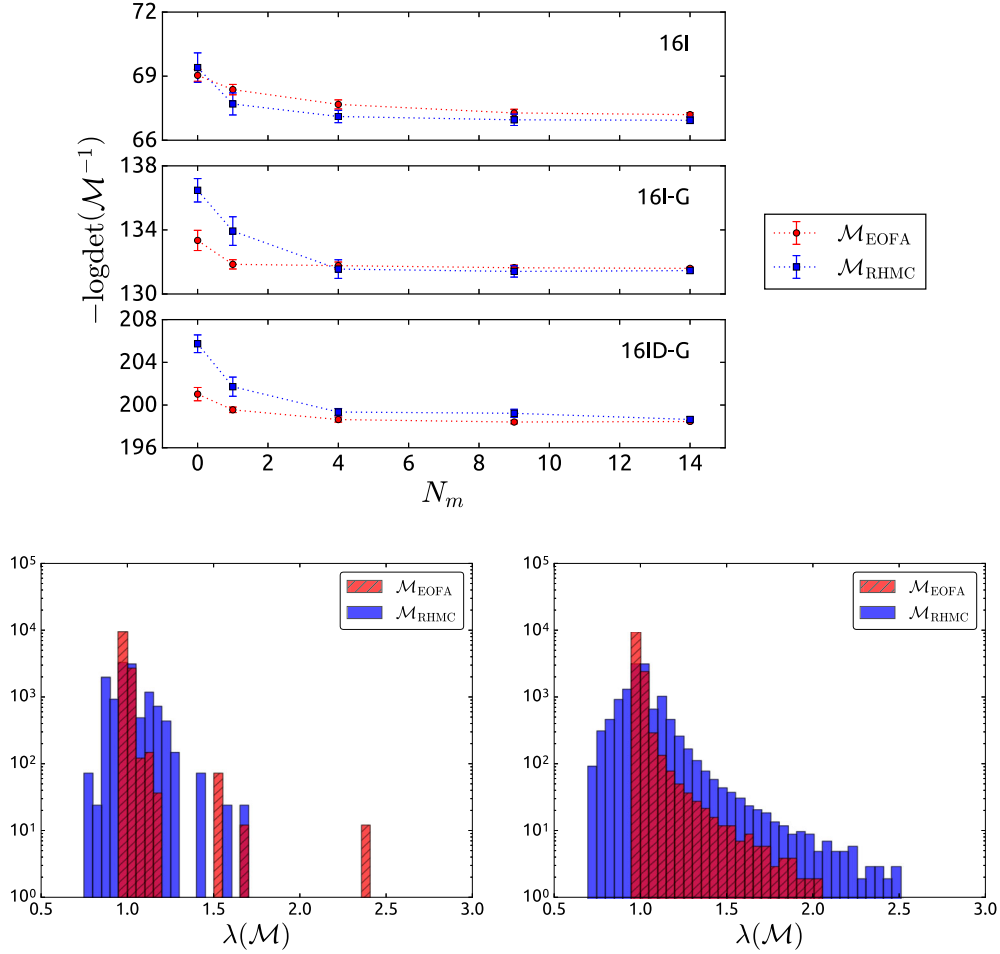


FIG. 1. Top: log determinants of the EOFA and RHMC actions as a function of the number of intermediate masses (N_m) used to compute Eq. (25), computed on a single, thermalized configuration of the 16I, 16I-G, and 16ID-G ensembles. We set (am_1, am_2) to $(0.032, 0.042)$, $(0.032, 0.042)$, and $(0.045, 0.055)$ on the 16I, 16I-G, and 16ID-G ensemble, respectively. We note that the error bars are purely statistical; for small N_m there is a large, unaccounted systematic error associated with setting $\Sigma = 1$ in Eqs. (23) and (24). Bottom: eigenvalue spectra of $\mathcal{M}_{\text{EOFA}}$ and $\mathcal{M}_{\text{RHMC}}$ on a 4^5 lattice with $am_1 = 0.01$, $am_2 = 1.0$, and $aM_5 = 1.8$. In the bottom left plot all of the gauge links are set to $U_\mu(x) = 1$ (i.e., the free field limit); in the bottom right plot each gauge link is set to an independent, random $SU(3)$ matrix.

where the complete spectrum can be computed directly. While both operators have similar condition numbers, we find that most of the spectrum of the EOFA action is concentrated into a small interval $[1, 1 + \Delta]$ with $\Delta \sim \mathcal{O}(0.1)$, leading to an action which is easier to estimate stochastically.

We propose that TWQCD's EOFA construction can be thought of as a kind of preconditioning which computes the same determinant ratio as RHMC but modifies the operator inside the determinant ($\mathcal{M}_{\text{RHMC}}$), mapping its spectrum onto a more compact interval. This suggests an additional application of the EOFA formalism: quark mass reweighting factors can be computed substantially more cheaply using the EOFA action than using the RHMC action, especially at light quark masses, even if the ensemble was generated using RHMC. This could be useful, for example, to include the dynamical effects of isospin breaking in

ensembles generated with isospin-symmetric up and down quarks.

B. Heatbath

At the beginning of each HMC trajectory we wish to draw a random pseudofermion field ϕ according to the distribution $P(\phi) \propto \exp(-\phi^\dagger \mathcal{M}_{\text{EOFA}} \phi)$. To do this, we first draw a random vector η by independently sampling the real and imaginary parts of each component from the normal distribution with $\mu = 0$ and $\sigma^2 = 1/2$, and then compute $\phi = \mathcal{M}_{\text{EOFA}}^{-1/2} \eta$. As with the RHMC algorithm we approximate the inverse square root by an appropriately constructed rational function, but we stress that in the context of EOFA this rational approximation enters only into the heatbath and is not necessary to compute the EOFA action itself or the associated pseudofermion force. Naively




FIG. 2. Left: relative error— ε , defined by Eq. (30)—in seeding the pseudofermion heatbath as a function of the number of poles in the rational approximation to the inverse square root (N_p), with $am_1 = 0.032$ set to the dynamical heavy quark mass, and a stopping residual of 10^{-10} for all CG inversions. Right: condition numbers of $\mathcal{M}_{\text{EFOA}}$ and $\mathcal{D}_{\text{DWF}}^\dagger \mathcal{D}_{\text{DWF}}$ as a function of the bare input quark mass (am_q); for $\mathcal{M}_{\text{EFOA}}$ this is the numerator mass ($am_1 = am_q$), while the denominator mass is fixed at $am_2 \equiv 1$. Both calculations were performed on a single, thermalized configuration of the 16I ensemble.

applying a rational approximation with the form of Eq. (19) to the operator $\mathcal{M}_{\text{EFOA}}$ results in

$$\begin{aligned} \mathcal{M}_{\text{EFOA}}^{-1/2} \simeq & \alpha_0 + \sum_{k=1}^{N_p} \alpha_l \left[\frac{1}{\gamma_l} - kP_- \Omega_-^\dagger [H(m_1)]^{-1} \Omega_- P_- \right. \\ & \left. + kP_+ \Omega_+^\dagger [H(m_2) - \Delta_+(m_1, m_2)P_+]^{-1} \Omega_+ P_+ \right]^{-1}, \end{aligned} \quad (27)$$

where we have defined $\gamma_l \equiv (1 + \beta_l)^{-1}$. In this form, the nested inversions required to seed the heatbath would make

$$\begin{aligned} \mathcal{M}_{\text{EFOA}}^{-1/2} \simeq & \alpha_0 + \sum_{k=1}^{N_p} \alpha_l \gamma_l \{ 1 + k\gamma_l P_- \Omega_-^\dagger [H(m_1) - \gamma_l \Delta_-(m_1, m_2)P_-]^{-1} \Omega_- P_- \\ & - k\gamma_l P_+ \Omega_+^\dagger [H(m_2) - \beta_l \gamma_l \Delta_+(m_1, m_2)P_+]^{-1} \Omega_+ P_+ \}. \end{aligned} \quad (29)$$

With this expression the EFOA heatbath step can be performed at the cost of $2N_p$ CG inversions using a rational approximation with N_p poles. Unlike the case of RHMC, multishift CG algorithms are not applicable to the EFOA heatbath since each of the $2N_p$ operators in Eq. (29) generates a different Krylov space. Furthermore, since the operators $\Delta_\pm P_\pm$ have a large number of zero modes and are therefore not invertible, there is no simple transformation by which this system can be recast into a form amenable to multishift CG.

In the left panel of Fig. 2 we test Eq. (29) on a single thermalized configuration of the 16I ensemble by computing the quantity

$$\varepsilon \equiv \frac{|\eta^\dagger \eta - \phi^\dagger \mathcal{M}_{\text{EFOA}} \phi|}{\eta^\dagger \eta} \quad (30)$$

after seeding the pseudofermion field ϕ with a random Gaussian vector η . In exact arithmetic $\varepsilon = 0$; in practice it measures the relative error in the heatbath step arising from the choice of CG stopping conditions and rational approximation to the inverse square root. We repeat this calculation, varying the number of poles in the rational approximation but keeping the stopping conditions fixed, and observe that ε reaches the limits of double-precision arithmetic even with a relatively modest number of poles compared to what is typically required to compute

EFOA prohibitively expensive. However, the Woodbury matrix identity

$$(A + BCD)^{-1} = A^{-1} - A^{-1}B(C^{-1} + DA^{-1}B)^{-1}DA^{-1} \quad (28)$$

and the cancellation between cross-terms involving products of the chiral projection operators can be used to manipulate this expression into the equivalent form

non-integer powers of $\mathcal{D}^\dagger \mathcal{D}$ accurately in the context of RHMC. In the right panel of Fig. 2 we demonstrate this explicitly by computing the condition numbers $\kappa = \lambda_{\max}/\lambda_{\min}$ of both operators as a function of the bare input quark mass. In Sec. VI we show how aggressive tuning of the rational approximation and stopping conditions, together with forecasting techniques for the initial CG guesses, can be combined to ameliorate the cost of the $2N_p$ inversions required to apply $\mathcal{M}_{\text{EOFA}}^{-1/2}$.

C. Pseudofermion force

The pseudofermion force

$$T^a \partial_{x,\mu}^a S_f(U) \equiv T^a \frac{d}{ds} S_f(e^{sT^a} U_\mu(x)) \Big|_{s=0} \quad (31)$$

measures the back-reaction of the pseudofermions on the HMC system [Eq. (15)] under an infinitesimal variation of the gauge field. In our notation $\{T^a\}$ is a basis for the Lie algebra $\mathfrak{su}(3)$, with the CPS normalization convention

$$\text{Tr}(T^a T^b) = -\frac{1}{2} \delta^{ab}. \quad (32)$$

The EOFA pseudofermion force can be worked out explicitly by differentiating the EOFA action [Eq. (10)] and applying the matrix identity

$$\partial_x M^{-1} = -M^{-1} (\partial_x M) M^{-1}, \quad (33)$$

resulting in

$$T^a \partial_{x,\mu}^a S(U) = k T^a (\gamma_5 R_5 \chi_1)^\dagger (\partial_{x,\mu}^a D_W) \chi_1 - k T^a (\gamma_5 R_5 \chi_2)^\dagger (\partial_{x,\mu}^a D_W) \chi_2, \quad (34)$$

with

$$\chi_1 \equiv [H(m_1)]^{-1} \Omega_- P_- \phi \quad (35)$$

and

$$\chi_2 \equiv [H(m_2) - \Delta_+(m_1, m_2) P_+]^{-1} \Omega_+ P_+ \phi. \quad (36)$$

Standard manipulations can be used to write a more general Dirac bilinear as

$$\begin{aligned} & a^\dagger (\partial_{x,\mu}^a D_W) b \\ &= -\frac{k}{2} \sum_{x,s,\mu} [a^\dagger(x,s) T^a U_\mu(x) (1 - \gamma_\mu) b(x + \hat{\mu}, s) \\ &\quad - a^\dagger(x + \hat{\mu}, s) U_\mu^\dagger(x) T^a (1 + \gamma_\mu) b(x)] \\ &= -\frac{k}{2} \sum_{x,s,\mu} [U_\mu(x) (\text{Tr}_{\text{spin}} [(1 + \gamma_\mu) a(x + \hat{\mu}, s) b^\dagger(x, s)] \\ &\quad + \text{Tr}_{\text{spin}} [(1 - \gamma_\mu) b(x + \hat{\mu}, s) a^\dagger(x, s)])], \end{aligned} \quad (37)$$

allowing Eq. (34) to be efficiently computed locally in terms of a trace over spinor indices, at the cost of the two inversions required to form χ_1 and χ_2 . Since Dirac bilinears of the form $a^\dagger (\partial_{x,\mu}^a D_W) b$ enter into the pseudofermion forces associated with many of the standard pseudofermion actions for Wilson and domain wall fermions—including the RHMC action—implementing the EOFA pseudofermion force requires little new code beyond what is required to implement the EOFA Hamiltonian.

V. SMALL VOLUME REPRODUCTION TESTS

To further test our implementation of EOFA we have reproduced the 16I (16I-G, 16ID-G) ensemble using EOFA for the strange quark (light quarks) in place of RHMC. For these tests we have made no serious effort to tune EOFA for performance; we have simply checked that replacing RHMC with EOFA, but leaving all other details of the simulation fixed, has no discernible impact on physical observables such as the average plaquette, topological susceptibility, and low energy spectrum.

A. Ensemble generation

The details of the integrator parameters and nesting are summarized in Tables IV and V, respectively. We use the abbreviations

$$\text{Quo}(m_1, m_2) \equiv \det \left(\frac{\mathcal{D}_{\text{DWF}}^\dagger \mathcal{D}_{\text{DWF}}(m_1)}{\mathcal{D}_{\text{DWF}}^\dagger \mathcal{D}_{\text{DWF}}(m_2)} \right) \quad (38)$$

and

TABLE IV. Basic integrator and HMC details for the generation of the 16I, 16I-G, and 16ID-G ensembles. We use nested Sexton-Weingarten integration schemes, detailed in Table V, with $\delta\tau$ the coarsest time step used to evolve the outermost level. We denote the CG stopping tolerances used for the force gradient forecasting, molecular dynamics, and Monte Carlo steps by r_{FG} , r_{MD} , and r_{MC} , respectively.

Ensemble	Integrator	$\delta\tau$	r_{FG}	r_{MD}	r_{MC}
16I	Force Gradient QPQPQ	0.1000	10^{-7}	10^{-8}	10^{-10}
16I-G	Omelyan ($\lambda = 0.2$)	0.2000	...	10^{-8}	10^{-10}
16ID-G	Force Gradient QPQPQ	0.1667	10^{-6}	10^{-7}	10^{-10}

TABLE V. Integrator layouts for the original RHMC runs. Here “Quo” is an abbreviation for the quotient action [Eq. (38)] and “RatQuo $_{1/n}$ ” is an abbreviation for the rational quotient action [Eq. (39)], with a rational function approximation used to apply $(\mathcal{D}^\dagger \mathcal{D})^{1/n}$ and its inverse. For the EOFA reproduction runs each instance of RatQuo $_{1/2}$ is replaced by an EOFA determinant with the same masses [Eq. (10)], while all other ensemble and integrator details are left fixed. The notation A:B for the update scheme denotes the number of steps of the next innermost integrator level (A) per step of the current level (B).

Ensemble	Level	Action	Update
16I	1	Quo(0.01, 0.2) + Quo(0.2, 1.0) + RatQuo $_{1/2}$ (0.032, 1.0)	4:1
	2	Gauge	1:1
16I-G	1	RatQuo $_{1/2}$ (0.01, 0.032)	1:1
	2	RatQuo $_{1/4}$ (0.032, 1.0) + RatQuo $_{1/4}$ (0.032, 1.0) + RatQuo $_{1/4}$ (0.032, 1.0)	8:1
	3	Gauge	1:1
16ID-G	1	RatQuo $_{1/2}$ (0.01, 0.05) + RatQuo $_{1/2}$ (0.05, 1.0) + RatQuo $_{1/4}$ (0.045, 1.0)	1:1
	2	DSDR	8:1
	3	Gauge	1:1

$$\text{RatQuo}_{1/n}(m_1, m_2) \equiv \left[\det \left(\frac{\mathcal{D}_{\text{DWF}}^\dagger \mathcal{D}_{\text{DWF}}(m_1)}{\mathcal{D}_{\text{DWF}}^\dagger \mathcal{D}_{\text{DWF}}(m_2)} \right) \right]^{1/n} \quad (39)$$

to denote the quotient and rational quotient actions, and on the EOFA reproduction ensembles replace each instance of RatQuo $_{1/2}(m_1, m_2)$ with the EOFA action [Eq. (10)] using the same mass parameters. The 16I and 16I-G EOFA reproduction runs were seeded with an ordered start—i.e., all gauge links were initially set to the unit matrix—and evolved for 1500 and 2500 MD time units, respectively. For the 16ID-G ensemble the last RHMC configuration (MD trajectory 908) was used to seed the start of the EOFA reproduction run, and then evolved for an additional 500 MD time units.

B. Basic observables

In Table VI we summarize results for the average plaquette $\langle P \rangle$, light quark and strange quark chiral $\langle \bar{\psi} \psi \rangle$ and pseudoscalar $\langle \bar{\psi} \gamma_5 \psi \rangle$ condensates, and the topological susceptibility $\chi_t \equiv \langle Q^2 \rangle / V$, computed on each ensemble; we observe statistically consistent results between the RHMC and EOFA ensembles in each case. Accompanying plots of

the time evolution can be found in Appendix C. The topological charge Q has been measured using the 5Li discretization introduced in Ref. [32] after cooling the gauge fields with 20 steps of APE smearing [33] using a smearing coefficient of 0.45. The ensemble averages were computed after binning over 50 (25) successive MD time units on the 16I and 16I-G (16ID-G) ensembles, where the bin size has been conservatively chosen based on the integrated auto-correlation times measured in Ref. [24] for the 16I ensemble and Ref. [27] for a series of $\beta = 1.75$ DSDR ensembles.

C. Low energy spectra

In Table VII we list results for the pion, kaon, Omega baryon, and residual masses, computed on the 16I ensemble. These calculations were performed using a measurement package previously introduced in Ref. [13], and based on the all-mode averaging (AMA) technique of Ref. [34]. Five *exact* light quark propagators were computed per trajectory using a deflated, mixed-precision CG solver with 600 low-mode deflation vectors and a tight stopping residual $r = 10^{-8}$, while *sloppy* propagators were computed for all time slices using a reduced stopping residual

TABLE VI. Average plaquettes, quark condensates, and topological susceptibilities (χ_t) computed on the 16I, 16I-G and 16ID-G lattices and their corresponding EOFA reproduction ensembles. The ensemble averages on the 16I (16I-G) lattices were computed using MD trajectories 500–1500 (500–2500) after binning over 50 successive MD time units. The ensemble averages on the 16ID-G lattices were computed using MD trajectories 500:900 for the RHMC ensemble, and MD trajectories 960:1360 for the EOFA ensemble, after binning over 25 successive MD time units. We do not compute χ_t on the 16ID-G ensemble since the short 400 MD time unit measurement runs are insufficient to adequately sample the topological charge, as evidenced by the time evolutions plotted in Appendix C.

Observable	16I		16I-G		16ID-G	
	RHMC	EOFA	RHMC	EOFA	RHMC	EOFA
$\langle P \rangle$	0.588 087(22)	0.588 106(26)	0.588 033(24)	0.588 039(16)	0.514 251(43)	0.514 200(48)
$\langle \bar{\psi}_l \psi_l \rangle$	0.001 697(5)	0.001 698(11)	0.001 715 1(72)	0.001 713 0(52)	0.005 543(11)	0.005 563(8)
$\langle \bar{\psi}_s \psi_s \rangle$	0.003 745 0(31)	0.003 743 5(74)	0.003 754 1(51)	0.003 752 9(34)	0.008 572 9(82)	0.008 589 5(69)
$\langle \bar{\psi}_l \gamma_5 \psi_l \rangle$	-0.000 015(14)	-0.000 012(19)	-0.000 003(15)	-0.000 006(12)	0.000 033(13)	-0.000 001(11)
$\langle \bar{\psi}_s \gamma_5 \psi_s \rangle$	-0.000 001(8)	-0.000 007(12)	-0.000 000 4(92)	-0.000 003 4(81)	0.000 017(10)	-0.000 002(8)
χ_t	$1.03(19) \times 10^{-5}$	$1.81(42) \times 10^{-5}$	$2.16(47) \times 10^{-5}$	$1.53(27) \times 10^{-5}$

TABLE VII. Low energy spectrum on the 16I ensemble computed from 100 independent measurements beginning with MD trajectory 500 and separated by 10 MD time units. Prior to fitting the correlation functions were binned over groups of 5 measurements. Corresponding effective mass plots can be found in Appendix C.

Observable	16I	
	RHMC	EOFA
am_π	0.2424(11)	0.2425(8)
am_K	0.3252(11)	0.3253(7)
am_Ω	1.003(15)	0.994(11)
$am_{\text{res}}^i(m_i)$	0.003 055 8(80)	0.003 052 3(78)

$r = 10^{-4}$. Strange quark propagators were computed with the tight residual $r = 10^{-8}$ for all time slices using ordinary CG with no deflation. AMA correlation functions were then computed by time-translational averaging of the sloppy propagators, using the available exact propagators to correct for bias. The light quark propagators were computed using Coulomb gauge-fixed wall (W) sources, with either local (L) or wall sinks; the strange quark propagators were computed using Coulomb gauge-fixed wall or Z_3 box (Z_3B) sources, and in both cases local sinks.

The pion and kaon masses were extracted by fitting to the asymptotic, large Euclidean time limit of the respective two-point correlation function,

$$\langle 0 | \bar{\mathcal{O}}(t) \mathcal{O}(0) | 0 \rangle \underset{t \rightarrow \infty}{\simeq} \frac{\langle 0 | \bar{\mathcal{O}}(t) | X \rangle \langle X | \mathcal{O}(0) | 0 \rangle}{2m_X V} (e^{-m_X t} \pm e^{-m_X(T-t)}), \quad (40)$$

where \mathcal{O} denotes the choice of interpolating operator, $X \in \{\pi, K\}$ is the ground state to which \mathcal{O} couples, and V and T are the spatial volume and temporal extent of the lattice, respectively. In particular, we performed simultaneous fits to the $\langle PP^{LW} \rangle$, $\langle PP^{WW} \rangle$, and $\langle AP^{LW} \rangle$ correlators, with $P(x) = \bar{\psi}(x) \gamma_5 \psi(x)$ and $A(x) = \bar{\psi}(x) \gamma_5 \gamma_4 \psi(x)$, and the first (second) superscript denotes the sink (source) type. The Omega baryon mass was extracted from the two-point correlation function

$$C_{\Omega\Omega}^{s_1 s_2}(t) = \sum_{i=1}^3 \sum_{\vec{x}} \langle 0 | \bar{\mathcal{O}}_{\Omega}^{s_1}(\vec{x}, t)_i \mathcal{O}_{\Omega}^{s_2}(0)_i | 0 \rangle \quad (41)$$

with the interpolating operator

$$\mathcal{O}_{\Omega}(x)_i = \varepsilon_{abc} (s_a^\dagger(x) C \gamma_i s_b(x)) s_c(x), \quad (42)$$

$s_1 = L$ and $s_2 \in \{W, Z_3B\}$. The correlators were then projected onto the positive parity component

$$P_+ C_{\Omega\Omega}^{s_1 s_2} = \frac{1}{4} \text{Tr} \left[\frac{1}{2} (1 + \gamma_4) C_{\Omega\Omega}^{s_1 s_2} \right] \quad (43)$$

and simultaneously fit to double exponential ansätze with common mass terms

$$C_{\Omega\Omega}^{s_1 s_2}(t) = (Z_1)_{\Omega\Omega}^{s_1 s_2} e^{-m_\Omega t} + (Z_2)_{\Omega\Omega}^{s_1 s_2} e^{-m_\Omega' t}. \quad (44)$$

Finally, the residual mass was determined by fitting the ratio

$$R(t) = \frac{\langle 0 | \sum_{\vec{x}} \bar{j}_{5q}^a(\vec{x}, t) | \pi \rangle}{\langle 0 | \sum_{\vec{x}} \bar{j}_5^a(\vec{x}, t) | \pi \rangle} \quad (45)$$

to a constant, where j_{5q}^a is the five-dimensional pseudoscalar density evaluated at the midpoint of the fifth dimension, and j_5^a is the physical pseudoscalar density constructed from the surface fields.

In addition, we have also measured the ground state pion energy, kaon mass, and residual mass on the 16I-G and 16ID-G ensembles. While the ground state of the kaon is at rest, the ground state of the pion has nonzero momentum $\vec{p}_{100} = (\pm\pi/L, 0, 0)$ on the 16I-G ensemble and $\vec{p}_{111} = (\pm\pi/L, \pm\pi/L, \pm\pi/L)$ on the 16ID-G ensemble due to the boundary conditions. These calculations make use of an extension of the AMA measurement package described above to G -parity ensembles; as discussed in Ref. [9], this requires the inclusion of additional diagrams that are generated by the mixing of quark flavors at the lattice boundaries through the G -parity operation. We measure on 51 configurations of the 16I-G ensemble, beginning with trajectory 500 and with a separation of 40 MD time units, and use sloppy and exact CG stopping tolerances of 10^{-4} and 10^{-10} , respectively, with a single exact solve per trajectory. We likewise measure on 21 configurations of the 16ID-G ensemble, beginning with trajectory 500 (960) for the RHMC (EOFA) ensemble and separated by 20 MD time units, and use the same AMA setup. We perform no additional binning for either ensemble since the separations between consecutive measurements are already comparable to the bin sizes used to compute the plaquette and quark condensates.

D. Pseudofermion forces on the 16I ensemble

TWQCD has observed that the average EOFA pseudofermion force is roughly half the size of the corresponding average RHMC pseudofermion force for a particular dynamical $N_f = 1$ QCD simulation with domain wall quarks performed in Ref. [4]. Following this observation, we examine the forces on the RHMC and EOFA variants of the 16I ensemble. We define a norm on the space of $\mathfrak{su}(3)$ -valued pseudofermion force matrices $F_\mu^a(x) \equiv \partial_{x,\mu}^a S(U)$ by

$$\|F_\mu(x)\| \equiv \left[\sum_a F_\mu^a(x) F_\mu^a(x) \right]^{1/2}, \quad (46)$$

and consider two measures of the force associated with a given configuration of the gauge field: the first is the RMS force




FIG. 3. Histograms of the RMS and maximum pseudofermion forces associated with force evaluations falling between trajectories 500 and 1500 of the 16I HMC evolutions. F_{RMS} and F_{max} are defined by Eqs. (47) and (48), respectively. Δt is the step size used to integrate the pseudofermion force contributions to the HMC evolution.

$$F_{\text{RMS}} \equiv \frac{1}{4V} \left[\sum_{x,\mu} \|F_{\mu}(x)\|^2 \right]^{1/2} \quad (47)$$

and the second is the maximum force

$$F_{\text{max}} \equiv \max_{x,\mu} \|F_{\mu}(x)\|, \quad (48)$$

in both cases taken over all lattice sites and link directions. While we expect Eq. (48) to be a more pertinent definition in the context of HMC simulations—we have empirically found that acceptance is controlled by the size of F_{max} —both F_{RMS} and F_{max} are, *a priori*, reasonable measures of the pseudofermion force.

In Fig. 3 we compare histograms of F_{RMS} and F_{max} between the RHMC and EOFA 16I ensembles. Each data point corresponds to a single evaluation of the pseudofermion force falling between MD trajectories 500 and 1500. We find that comparing the relative sizes of the RHMC and EOFA forces is highly dependent on whether one chooses F_{RMS} or F_{max} ; the mean EOFA F_{RMS} is roughly 30%

smaller than the mean RHMC F_{RMS} , but the distributions of F_{max} are nearly indistinguishable. This observation suggests that while the EOFA force distribution may have a smaller mean than the RHMC force distribution, the EOFA distribution also likely has longer tails, such that the largest forces have similar magnitudes. Since we expect the magnitude of the largest forces to correlate more strongly with the efficiency of the integrator than the magnitude of the average forces, as we have argued above, we interpret these results as suggesting that the optimal step size for an EOFA evolution should be similar to that of an RHMC simulation with the same mass parameters, even if the average force is somewhat smaller.

TWQCD has also observed a large hierarchy of scales in the pseudofermion forces associated with each of the two terms in Eq. (34); in Ref. [5] they find that the average force associated with the first term—involving the left-handed component of the pseudofermion field—is more than an order of magnitude smaller than the average force associated with the second term—involving the right-handed component—for two different dynamical QCD




FIG. 4. Histograms of the RMS and maximum pseudofermion forces associated with the left-handed and right-handed components of the pseudofermion field in Eq. (34).

simulations. They exploit this observation with a Sexton-Weingarten integration scheme, integrating the first term with a larger time step than the second, and find increased efficiency in their simulations. In Fig. 4 we compare histograms of the RMS and maximum left-handed and right-handed forces from 1000 thermalized configurations of the 16I EOFA ensemble. Our conclusions are analogous to the comparison between the EOFA and RHMC forces: if one considers F_{RMS} the left-handed force contribution is indeed substantially smaller than the right-handed force contribution, but if one instead considers F_{max} the force distributions are very similar in both magnitude and shape. Based on the latter observation we leave both terms in Eq. (34) on the same time step in our large-scale EOFA simulations.

We also note that in these small volume test runs we have not considered applying the Hasenbusch mass preconditioning technique [1] to the EOFA formalism. Introducing a set of Hasenbusch masses $\{m'_i\}_{i=1}^N$, with $m_1 < m'_i < m_2$, we can write the fermion determinant as

$$\det\left(\frac{\mathcal{D}(m_1)}{\mathcal{D}(m_2)}\right) = \det\left(\frac{\mathcal{D}(m_1)}{\mathcal{D}(m'_1)}\right) \left[\prod_{i=1}^{N-1} \det\left(\frac{\mathcal{D}(m'_i)}{\mathcal{D}(m'_{i+1})}\right) \right] \times \det\left(\frac{\mathcal{D}(m'_N)}{\mathcal{D}(m_2)}\right). \quad (49)$$

While the left-hand side can be simulated using a single pseudofermion field, the associated forces can be large if $m_1 \ll m_2$, requiring a small step size to maintain reasonable acceptance. The right-hand side, in contrast, involves $N + 1$ independent pseudofermion fields, but with possibly substantially reduced forces, allowing larger step sizes to be used. For light m_1 one typically observes that the gain from increasing the step size offsets the cost of simulating extra heavy flavors, leading to a more efficient simulation. In Sec. VII we demonstrate that Hasenbusch preconditioning allows for a substantial speed-up in the context of the 32ID-G ensemble. We also note that in addition to reducing the size of the pseudofermion forces, the Hasenbusch technique

preconditions the EOFA force in the sense that the size hierarchy between the left-handed and right-handed force contributions to a single determinant disappears in the limit $m'_i \rightarrow m'_{i+1}$. In practice, we find that the mass preconditioned simulation has comparable left-handed and right-handed force contributions even in the RMS sense.

E. Autocorrelations on the 16I-G ensemble

While EOFA and RHMC are formally equivalent in the sense of Eq. (21)—in the limit of infinite statistics they produce the same quark determinant, and hence the same ensemble averages for physical observables—it is in principle possible that the two algorithms approach this limit at different rates. Tables VI–VIII provide some evidence against an observable difference in our reproduction tests, since the error bars computed for the RHMC ensembles and the error bars computed for the EOFA ensembles are compatible. To more directly address this question we have also examined integrated autocorrelation times

$$\tau_{\text{int}}(\Delta_{\text{cut}}) = \frac{1}{2} + \sum_{\Delta=1}^{\Delta_{\text{cut}}} C(\Delta) \quad (50)$$

on the 16I-G ensembles [35] for the set of basic physical observables discussed in Sec. VB, where

$$C(\Delta) = \left\langle \frac{(X(t) - \bar{X})(X(t + \Delta) - \bar{X})}{\sigma_X^2} \right\rangle \quad (51)$$

is the autocorrelation function associated with the observable $X(t)$ at lag Δ . Our procedure for estimating errors on τ_{int} is identical to the procedure introduced in Ref. [27]: for each fixed Δ in Eq. (51) we bin the set of measurements $(X(t) - \bar{X})(X(t + \Delta) - \bar{X})$ over neighboring configurations and estimate the error on the mean $\langle \dots \rangle$ by bootstrap resampling.

In Fig. 5 we compare $2\tau_{\text{int}}$ —the separation between statistically independent molecular dynamics configurations—between the RHMC simulation and the EOFA simulation. Plots of the same observables as a function

of the molecular dynamics trajectory can be found in Appendix C. We note that the results for a given observable are consistent within the measured statistical errors, with the possible exception of the light quark pseudoscalar condensate, for which we measure a somewhat longer autocorrelation time in the RHMC simulation. We conclude that there is no evidence of a difference in autocorrelation times between the EOFA and RHMC algorithms in this work, but we also stress that this is an important issue and should continue to be addressed in future studies.

VI. OPTIMIZATION AND TUNING

In this section we discuss preconditioning and algorithmic techniques which reduce the cost of EOFA simulations. In some cases these are extensions of well-known lattice techniques to the EOFA formalism, while in other cases they are specific to EOFA. We illustrate these techniques using benchmark tests computed with the physical quark mass, Möbius DWF 24ID ensemble, and report timing results for code written in the Columbia Physics System (CPS) and running on 256-node or 512-node Blue Gene/Q partitions.

A. Inversions of $\mathcal{D}_{\text{EOFA}}$

Since the majority of the computational effort in an HMC simulation is associated with repeatedly inverting the Dirac operator, techniques to more efficiently apply the Dirac operator or to otherwise accelerate these inversions can have a dramatic impact on the overall efficiency of the integrator. To address the former, we make use of the BAGEL assembler generation library [36] to produce highly optimized kernels and fermion solvers for the Blue Gene/Q hardware. To address the latter, we make use of multiple preconditioning techniques, as well as a mixed precision defect correction CG solver.

The first preconditioning technique we apply—“even-odd” or “red-black” preconditioning—is well-known in the lattice QCD community. Lattice sites are labeled as even if $(x + y + z + t) \equiv 0 \pmod{2}$, or odd if $(x + y + z + t) \equiv 1 \pmod{2}$, inducing a 2×2 block structure on fermion operators

TABLE VIII. Low energy spectra on the 16I-G and 16ID-G ensembles computed from 51 and 21 measurements, respectively. On the 16I-G ensemble we also predict the ground state pion energy using the fitted am_π on the 16I ensemble and the continuum dispersion relation $aE_\pi^{\text{pred}} = \sqrt{(am_\pi)^2 + (a\pi/L)^2}$. Corresponding effective mass plots can be found in Appendix C.

Observable	16I-G		16ID-G	
	RHMC	EOFA	RHMC	EOFA
aE_π	0.3175(43)	0.3097(48)	0.4457(101)	0.4614(72)
aE_π^{pred}	0.311 97(4)	0.312 07(4)
am_K	0.3271(22)	0.3272(28)	0.4343(34)	0.4382(24)
$am'_{\text{res}}(m_l)$	0.003 140(90)	0.003 054(86)	0.009 19(14)	0.009 52(13)




FIG. 5. Integrated autocorrelation times for the average plaquette, for the light and strange quark chiral and pseudoscalar condensates, and for the square of the topological charge, measured using 2500 thermalized molecular dynamics trajectories of the RHMC (left) and EOFA (right) 16I-G ensembles.

$$M = \begin{pmatrix} M_{ee} & M_{eo} \\ M_{oe} & M_{oo} \end{pmatrix}. \quad (52)$$

Standard tricks can then be used to relate the linear system $M\psi = \phi$ to a better conditioned linear system involving only the odd sub-lattice; this preconditioned system is substantially cheaper to invert since the size of the problem has been halved. After inverting on the odd sub-lattice, the even component of ψ can also be recovered at modest cost, without ever needing to explicitly invert on the even sub-lattice. The details of this construction, and its extension to EOFA, are described in Appendix B.

The second preconditioning technique we apply—Cayley-form preconditioning—is unique to EOFA, and was introduced in Ref. [6]. The generic linear system one needs to solve in the context of EOFA has the form

$$(H(m_1) + \beta\Delta_{\pm}(m_2, m_3)P_{\pm})\psi = \phi, \quad (53)$$

where $H = \gamma_5 R_5 \mathcal{D}_{\text{EOFA}}$. For Möbius domain wall fermions $\mathcal{D}_{\text{EOFA}}$ is dense in ss' , and thus considerably more expensive to invert than \mathcal{D}_{DWF} , which has a tridiagonal ss' stencil, in terms of wall clock time. However, Eq. (6) suggests that Eq. (53) can be related to an equivalent system in terms of \mathcal{D}_{DWF} by using \tilde{D}^{-1} as a preconditioner. We elaborate on the mathematical details in Appendix B 2, and, in particular, demonstrate that $\Delta_{\pm}\tilde{D}$ has a relatively simple, rank-one form, allowing for substantially more efficient EOFA inversions—even when $\beta \neq 0$ —by working with the preconditioned system. This technique also has the advantage that it allows for EOFA simulations which re-use existing high-performance code for applying \mathcal{D}_{DWF} with little modification.

Finally, we use a restarted, mixed precision defect correction solver to perform the conjugate gradient inversions of the fully preconditioned EOFA system. For memory bandwidth-limited calculations—such as applying the Dirac operator—single precision computations can be performed at approximately half the cost of full double precision computations. In the defect correction approach to mixed precision CG, the following algorithm is used:

- (1) Solve the Dirac equation in single precision arithmetic using a reduced stopping tolerance (typically 10^{-4} or 10^{-5}).
- (2) Compute the current residual using the (single precision) solution in full double precision arithmetic.
- (3) If the desired final tolerance (typically 10^{-8} or smaller) has been reached, stop. Otherwise, return to step 1, using the residual vector computed in step 2 as the new CG source.

We observe that this algorithm outperforms straight double precision CG by approximately a factor of 2—as one would expect if the calculation is truly memory bandwidth-limited—provided the local lattice volume on each node is sufficiently large to avoid communications bottlenecks.

In Fig. 6 we plot the CG residual as a function of the wall clock running time of the inverter for a series of benchmark inversions of Eq. (53) on the 24ID ensemble. These benchmarks show the inverter performance as we sequentially introduce even-odd preconditioning, Cayley-form preconditioning, and finally, mixed precision CG. We also plot the time required to solve the family of linear systems

$$(\mathcal{D}_{\text{DWF}}^{\dagger} \mathcal{D}_{\text{DWF}} + \beta_k)\psi = \phi \quad (54)$$




FIG. 6. Wall clock time required to solve Eq. (53) to a stopping tolerance of 10^{-10} at the physical strange quark mass on the 24ID ensemble, as the preconditioning and algorithmic refinements discussed in the text are introduced sequentially. The dashed vertical line corresponds to the time required to apply $(\mathcal{D}_{\text{DWF}}^\dagger \mathcal{D}_{\text{DWF}})^{-1/2}$ by solving Eq. (54) using the high-performance implementations of even-odd preconditioned \mathcal{D}_{DWF} and multishift CG in the BAGEL library.

using multishift CG for the same set of poles $\{\beta_k\}$ used in the rational approximation to $x^{-1/2}$ in the RHMC evolution that generated the 24ID ensemble. This allows a baseline estimate of the cost of evaluating the EOFA Hamiltonian or pseudofermion force against the cost of evaluating the RHMC Hamiltonian or pseudofermion force at the same quark mass. We observe a factor of 3.9 speed-up for fully preconditioned EOFA over the even-odd preconditioned RHMC system. In both cases the underlying operator being inverted is \mathcal{D}_{DWF} ; the slower RHMC benchmark demonstrates the overhead associated with multishift CG relative to solving a single system with standard CG, both due to the inability to fully utilize mixed precision methods and due to the additional linear algebra required at each iteration.

B. Heatbath

Achieving the full performance improvement suggested by the inversion benchmarks in Sec. VI A is complicated by the form of the EOFA heatbath, which is expected to be more expensive than the RHMC heatbath, even with efficient EOFA code. Applying $\mathcal{M}_{\text{EOFA}}^{-1/2}$ [Eq. (29)] requires two independent CG inversions per pole used in the rational approximation to $x^{-1/2}$, since multishift CG is not applicable: we use two algorithmic techniques to reduce this cost. The first is a forecasting technique initially proposed by Brower *et al.* [37] in the context of more general HMC simulations, and later used successfully by TWQCD in the context of the EOFA heatbath [4]. The idea is the following: given a set of solutions $\{\psi_k\}_{k=1}^N$ to Eq. (53) for N different poles $\{\beta_k\}_{k=1}^N$, one can use the linear combination

$$\psi_{N+1} = \sum_{k=1}^N c_k \psi_k \quad (55)$$

minimizing the functional

$$\Phi[\psi] = \psi^\dagger (H + \beta_{N+1} \Delta_\pm P_\pm) \psi - \phi^\dagger \psi - \psi^\dagger \phi \quad (56)$$

as the initial CG guess for the next inversion with pole β_{N+1} . The coefficients c_k satisfy

$$\sum_{k=1}^N c_k \psi_k^\dagger (H + \beta_{N+1} \Delta_\pm P_\pm) \psi_k = \psi_l^\dagger \phi, \quad (57)$$

and can be computed explicitly using *e.g.* Gauss-Jordan elimination. Since Eq. (56) is the same functional minimized by the conjugate gradient algorithm itself, accurate initial guesses can be computed for modest N provided the $\{\beta_k\}_{k=1}^{N+1}$ are similar in magnitude. In Fig. 7 we test this forecasting technique using the 24ID ensemble and a rational approximation with 8 poles, and find that the iteration count required to solve Eq. (53) to a tolerance of 10^{-10} is more than halved for the last few poles.

The second technique we have used to accelerate the heatbath is motivated by observing that the coefficients entering into Eq. (29) span several orders of magnitude for a typical rational approximation to $x^{-1/2}$. We find typical values $k\alpha_l \gamma_l^2 / \alpha_0 \sim \mathcal{O}(10^{-3} - 10^{-5})$, suggesting that the inversions can be performed with reduced stopping tolerances relative to the desired accuracy of $\mathcal{M}_{\text{EOFA}}^{-1/2} \psi$, since the solution vectors are ultimately multiplied by small coefficients when the result is formed. We have explored the following simple optimization scheme to relax the stopping conditions for each pole:

- (1) Choose a desired tolerance for the heatbath, ε_{tol} , where ε is defined by Eq. (30).

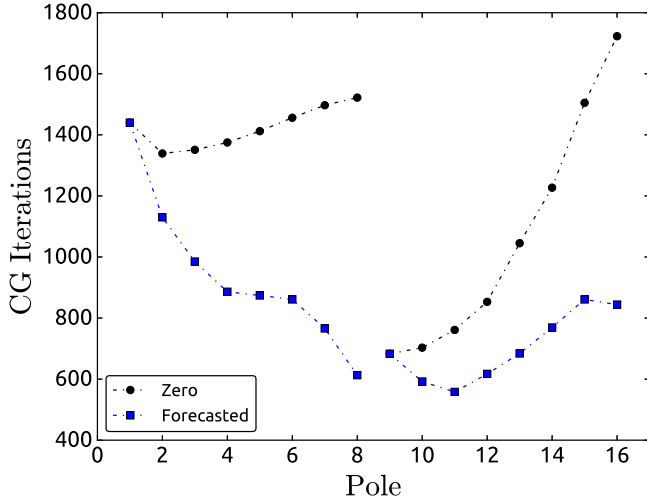


FIG. 7. CG iterations required to invert Eq. (53) for each of the 16 values of β entering into a rational approximation of $\mathcal{M}_{\text{EOFA}}^{-1/2}$ with 8 poles on the 24ID ensemble. The first 8 poles ($\beta = -\gamma_i$) are associated with the first (LH) term in Eq. (29), while the second 8 poles ($\beta = -\beta_i\gamma_i$) are associated with the second (RH) term. We find no improvement from using solutions to the LH system to forecast solutions to the RH system and vice-versa, since the Dirac operator being inverted in either case is evaluated with a different quark mass.

- (2) Choose one of the inversions required to compute $\mathcal{M}_{\text{EOFA}}^{-1/2}$ according to Eq. (29), and relax the stopping tolerance until the overall error in the heatbath ϵ reaches ϵ_{tol} .
- (3) Iterate over each inversion until all stopping conditions have been tuned.

We report results for the 24ID ensemble in Table IX. Using a rational approximation with 6 poles, and $\epsilon_{\text{tol}} = 10^{-10}$, we observe that the total heatbath time is more than halved while only slightly increasing the error. We have also checked that the final error and heatbath running time after tuning is insensitive to the exact order in which the stopping tolerances are tuned.

VII. LARGE-SCALE EOFA CALCULATIONS

In this section we turn to two ongoing ensemble generation calculations currently being performed by the

TABLE IX. The relative error (ϵ) and total running time for the EOFA heatbath on the 24ID ensemble before and after applying the tuning algorithm discussed in the text.

	ϵ	Total Heatbath Time
Untuned	1.52×10^{-11}	129.5 s
Tuned	7.79×10^{-11}	68.9 s

RBC/UKQCD Collaboration. The first is a strong-coupling $N_f = 2 + 1$ $24^3 \times 64 \times 24$ Iwasaki + DSDR lattice (24ID) intended for exploratory studies and calculations requiring high statistics [26]. The second (32ID) has been used for a first-principles calculation of the ratio of standard model CP -violation parameters ϵ'/ϵ from $\Delta I = 1/2$ $K \rightarrow \pi\pi$ decays in Ref. [7]. RBC/UKQCD is currently generating more gauge field configurations to reduce the statistical errors in the $\Delta I = 1/2$ decay amplitudes. Both ensembles have physical quark masses and large volumes, allowing for tests of the performance of EOFA in the context of state-of-the-art domain wall fermion calculations.

Tables X and XI summarize the details of the integrator parameters and nesting for these evolutions. The ensembles labeled RHMC correspond to the evolutions of Ref. [26] (24ID) and Ref. [7] (32ID-G). For the ensembles marked EOFA, we have changed the strange quark (light quark) action to EOFA for the 24ID (32ID-G) ensemble and returned the details of the integrator as described in the remainder of the section. For the 32ID-G ensemble—where, due to the G -parity flavor doubling, the EOFA action naturally describes the degenerate light quark pair—we have also switched from an Omelyan integrator to a force gradient integrator, and inserted additional Hasenbusch preconditioning determinants.

A. 24ID ensemble

We use the 24ID ensemble as a straightforward benchmark of RHMC against an equivalent EOFA simulation to describe a physical heavy quark flavor. Here this is the strange quark, but $N_f = 2 + 1 + 1$ simulations with dynamical strange and charm quarks are another obvious target of EOFA. We make no serious attempt to retune the

TABLE X. Basic integrator and HMC details for the generation of the 24ID and 32ID-G ensembles. We denote the coarsest time step used to evolve the outermost level by $\delta\tau$, and the CG stopping tolerances used for the force gradient forecasting, molecular dynamics, and Monte Carlo steps by r_{FG} , r_{MD} , and r_{MC} , respectively. We elaborate on the details of the integrator nesting in Table XI.

Ensemble	Integrator	$\delta\tau$	r_{FG}	r_{MD}	r_{MC}
24ID (RHMC)	Force Gradient QPQPQ	0.0833	10^{-5}	10^{-7}	10^{-10}
24ID (EOFA)	Force Gradient QPQPQ	0.0833	10^{-5}	10^{-7}	10^{-10}
32ID-G (RHMC)	Omelyan ($\lambda = 0.22$)	0.0625	...	10^{-7}	10^{-10}
32ID-G (EOFA)	Force Gradient QPQPQ	0.1667	10^{-5}	10^{-7}	10^{-10}

TABLE XI. Integrator layouts for the 24ID and 32ID-G ensembles. The notation A:B for the update scheme denotes the number of steps of the next innermost integrator level (A) per step of the current level (B).

Ensemble	Level	Action	Update
24ID (RHMC)	1	RatQuo _{1/2} (0.085, 1.0)	1:1
	2	Quo(0.00107, 0.00789) + Quo(0.00789, 0.0291) + Quo(0.0291, 0.095) + Quo(0.095, 0.3) + Quo(0.3, 0.548) + Quo(0.548, 1.0)	1:1
	3	Gauge + DSDR	1:1
24ID (EOFA)	1	EOFA(0.085, 1.0)	1:1
	2	Quo(0.00107, 0.00789) + Quo(0.00789, 0.0291) + Quo(0.0291, 0.095) + Quo(0.095, 0.3) + Quo(0.3, 0.548) + Quo(0.548, 1.0)	1:1
	3	Gauge + DSDR	1:1
32ID-G (RHMC)	1	RatQuo _{1/2} (0.0001, 0.007)	1:1
	2	RatQuo _{1/2} (0.007, 1.0) + RatQuo _{1/4} (0.045, 1.0)	1:2
	3	DSDR	1:2
	4	Gauge	1:1
32ID-G (EOFA)	1	EOFA(0.0001, 0.0058) + EOFA(0.0058, 0.0149) + EOFA(0.0149, 0.059) + EOFA(0.059, 0.177) + EOFA(0.177, 0.45) + EOFA(0.45, 1.0) + RatQuo _{1/4} (0.045, 1.0)	5:1
	2	DSDR	1:2
	3	Gauge	1:1

integrator after switching to EOFA beyond tuning the heatbath step with the following procedure:

- (1) Compute the largest and smallest eigenvalues of $\mathcal{M}_{\text{EOFA}}$ [Eq. (10)] for a few thermalized configurations of the gauge field, and use these measurements to inform the bounds of the rational approximations to $x^{-1/2}$ constructed via the Remez algorithm.
- (2) Add poles to the rational approximation, with all CG stopping tolerances set to r_{MC} , until $\varepsilon < r_{\text{MC}}$ [Eq. (30)] is reached.
- (3) With the rational approximation now fixed from step 2, tune the CG stopping tolerances corresponding to each pole, following the procedure described in Sec. VI B, and keeping $\varepsilon < r_{\text{MC}}$.

After tuning the heatbath, we then ran a single trajectory of the RHMC evolution and the EOFA evolution on a 256-node Blue Gene/Q partition. For the EOFA ensemble, we compare two schemes. The first (“dense”) is a straightforward implementation of Möbius DWF as proposed in Ref. [4]: we invert Eq. (53) directly, where $H = \gamma_5 R_5 \mathcal{D}_{\text{EOFA}}$ and the

other dense 5D operators appearing in the EOFA action are listed explicitly in Appendix A 2. We also do not apply the final step in our heatbath tuning procedure, leaving all CG stopping tolerances in the heatbath fixed at $r_{\text{MC}} = 10^{-10}$. In the second EOFA scheme (“preconditioned”) we fully tune the heatbath step and apply the Cayley-form preconditioning detailed in Appendix B 2 to inversions of Eq. (53). Timing breakdowns for the strange quark part of the evolution are reported in Table XII.

We observe that the dense EOFA formalism is actually somewhat slower than RHMC: the additional complexity of the EOFA heatbath, together with the more expensive inversions of the dense 5D operator $\mathcal{D}_{\text{EOFA}}$, negate the expected performance gains from the simpler forms of the Hamiltonian and force evaluations. We emphasize, however, that we have made no attempt to retune the integrator details to optimize for EOFA. After introducing Cayley-form preconditioning—so that we are inverting the tridiagonal operator \mathcal{D}_{DWF} rather than $\mathcal{D}_{\text{EOFA}}$ when we solve

TABLE XII. Strange quark timings for a single MD trajectory of the 24ID ensemble on a 256-node Blue Gene/Q partition. We compare RHMC to EOFA with (“preconditioned”) and without (“dense”) Cayley-form preconditioning.

Step	RHMC		EOFA (Dense)		EOFA (Preconditioned)	
	Time (s)	%	Time (s)	%	Time (s)	%
Heatbath	42.6	2.7	340.6	15.1	68.9	15.5
Force gradient integration (total)	1485.6	94.8	1840.6	81.8	355.9	80.1
Final Hamiltonian evaluation	39.4	2.5	68.8	3.1	19.8	4.4
Total	1567.6	...	2250.0	...	444.6	...
(Total RHMC)/Total	1.0	...	0.7	...	3.5	...

TABLE XIII. HMC details for the production ensemble generation run (1) of Ref. [7], as well as 13 tuning runs after switching to EOFA light quarks (2–14). We use the following notation: “O” denotes the Omelyan integrator, “FG” denotes the force gradient integrator, “ N_{traj} ” is the number of trajectories generated for the timing run, “acceptance” is the fraction of gauge field configurations which were accepted in the final Monte Carlo step, and “efficiency” is the ratio of the total job time per trajectory for the specified integration scheme to the total job time per trajectory of the scheme used in run 1. Entries in bold correspond to the original RHMC scheme (1) and the final, fully tuned EOFA scheme (12).

Run	Integrator Type	Light Hasenbusch Masses	$\Delta\tau$	r_{MD}	N_{traj}	Acceptance	Efficiency
	O	0.007	0.0625	10^{-8}	850	88%	...
2	O	...	0.0625	10^{-8}	10	40%	1.2
3	FG	0.043	0.0625	10^{-8}	10	100%	2.0
4	FG	0.018, 0.12	0.0625	10^{-8}	10	100%	1.8
5	FG	0.0118, 0.0412, 0.23	0.0625	10^{-8}	10	100%	1.7
6	FG	0.0075, 0.023, 0.11, 0.4	0.0625	10^{-8}	10	100%	1.7
7	FG	0.0058, 0.0149, 0.059, 0.177, 0.45	0.0625	10^{-8}	10	100%	1.5
8	FG	0.0103, 0.029, 0.12, 0.41	0.1000	10^{-6}	15	67%	4.0
9	FG	0.0103, 0.029, 0.12, 0.41	0.1000	10^{-7}	20	95%	3.0
10	FG	0.0103, 0.029, 0.12, 0.41	0.1667	10^{-7}	20	75%	4.5
11	FG	0.0058, 0.0149, 0.059, 0.177, 0.45	0.1000	10^{-6}	40	80%	3.0
12	FG	0.0058, 0.0149, 0.059, 0.177, 0.45	0.1667	10^{-7}	850	93%	4.2
13	FG	0.0058, 0.0149, 0.059, 0.177, 0.45	0.2000	10^{-7}	60	65%	4.5
14	FG	0.0058, 0.0149, 0.059, 0.177, 0.45	0.2000	10^{-8}	25	72%	3.9

Eq. (53)—we find that EOFA outperforms RHMC by a significant factor of 3.5.

B. 32ID-G ensemble

One particularly promising feature of EOFA in the context of G -parity ensembles is the potential for aggressive Hasenbusch mass preconditioning of the light quark determinant; this makes the 32ID-G ensemble a particularly interesting case study since the EOFA formalism is used to describe a physical mass light quark pair. In Ref. [7] the RBC/UKQCD Collaboration observed that mass preconditioning is not particularly effective for the RHMC light quark determinant, since each molecular dynamics step requires one multishift inversion of $\mathcal{D}^\dagger\mathcal{D}$ evaluated at the numerator quark mass and two multishift inversions of $\mathcal{D}^\dagger\mathcal{D}$ evaluated at the denominator quark mass. The latter two solves become prohibitively expensive if many intermediate masses are introduced, negating the expected gain from integrating the preconditioned pseudofermion forces with larger step sizes. The EOFA force, on the other hand, is no more expensive to evaluate than the force associated with the standard quotient action [Eq. (38)], so it is natural to expect better performance from Hasenbusch preconditioning.

In Table XIII we list details of the tuning runs we have used to explore potential schemes for evolving the 32ID-G ensemble with EOFA light quarks. We started by switching from an Omelyan integrator, for which the leading errors are $\mathcal{O}(\delta\tau^2)$, to a force gradient integrator, for which the leading errors are $\mathcal{O}(\delta\tau^4)$, and studied the effects of inserting mass preconditioning determinants one at a time

TABLE XIV. Measured spectral range of $\mathcal{M}_{\text{EOFA}}$, heatbath relative error (ϵ), and total time for the heatbath step (Δt_{HB}), using N_{LHSB} intermediate mass preconditioning steps and an order N_{poles} rational approximation to $x^{-1/2}$, with all CG stopping tolerances set to $r_{\text{MC}} = 10^{-10}$. Timings are reported for a 512-node Blue Gene/Q partition.

N_{LHSB}	Mass Ratio	λ_{min}	λ_{max}	N_{poles}	ϵ	Δt_{HB} (s)
0	0.0001/1.0	1.0	1150	11	6.91×10^{-11}	5263.2
1	0.0001/0.043	1.0	33.3	7	3.50×10^{-11}	2226.6
	0.043/1.0	1.0	22.8	7	6.82×10^{-12}	
2	0.0001/0.018	1.0	13.5	6	2.13×10^{-11}	2043.8
	0.018/0.12	1.0	6.4	5	1.11×10^{-11}	
	0.12/1.0	1.0	8.3	6	6.18×10^{-12}	
	0.0001/0.0118	1.0	8.9	6	6.08×10^{-12}	
3	0.0118/0.0412	1.0	3.3	4	4.09×10^{-11}	2307.8
	0.0412/0.23	1.0	5.5	5	1.29×10^{-11}	
	0.23/1.0	1.0	4.3	5	1.11×10^{-11}	
	0.0001/0.0075	1.0	5.9	5	9.63×10^{-12}	
	0.0075/0.023	1.0	2.8	4	3.55×10^{-11}	
4	0.023/0.11	1.0	4.6	5	1.00×10^{-11}	2080.7
	0.11/0.4	1.0	3.6	4	1.98×10^{-11}	
	0.4/1.0	1.0	2.5	4	2.34×10^{-11}	
	0.0001/0.0058	1.0	4.7	5	1.11×10^{-11}	
	0.0058/0.0149	1.0	2.3	4	1.64×10^{-11}	
5	0.0149/0.059	1.0	3.7	4	9.65×10^{-11}	2289.0
	0.059/0.177	1.0	3.0	4	4.14×10^{-11}	
	0.177/0.45	1.0	2.5	4	2.71×10^{-11}	
	0.45/1.0	1.0	2.2	4	1.64×10^{-11}	

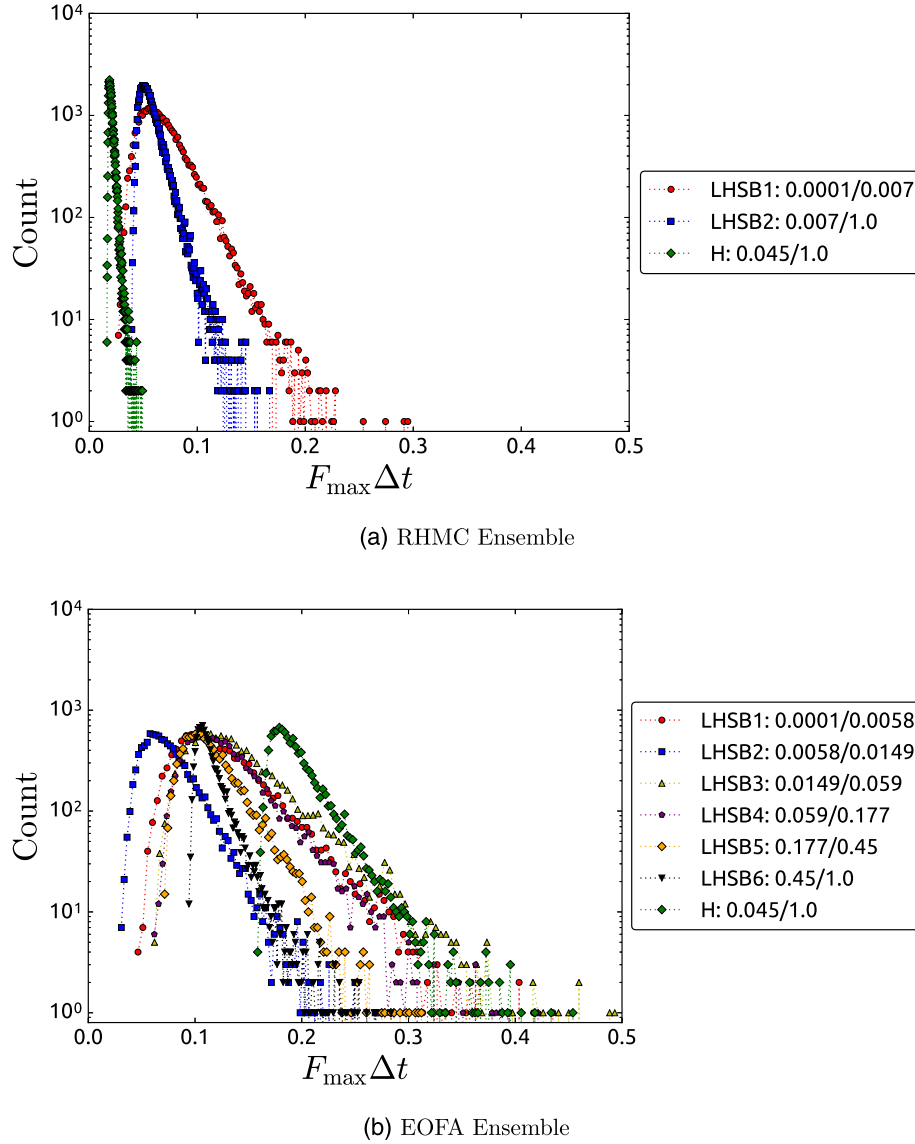


FIG. 8. Histograms of the maximum force, defined by Eq. (48), measured between trajectories 500 and 1350 on the 32ID-G RHMC ensemble and measured between trajectories 1350 and 2200 on the 32ID-G EOFA ensemble. We use the abbreviation “LHSB” in the legends to denote the various mass ratios entering into our mass preconditioning scheme for the light quark determinant, and “H” to denote the strange quark determinant.

(runs 1–7). We then identified two promising mass preconditioning schemes—one with four intermediate masses (runs 8–10), and the other with five intermediate masses (runs 11–14)—and continued tuning the step size, CG stopping conditions, and heatbath, to optimize the job time per trajectory and Monte Carlo acceptance. The initial RHMC scheme used in Ref. [7] corresponds to run 1, and the final EOFA scheme we have adopted for our continuing ensemble generation corresponds to run 12.

We find, in practice, that Hasenbusch mass preconditioning is extremely effective for the EOFA light quark determinant. In addition to reducing the size of the pseudofermion force, we also observe that the largest eigenvalue of the EOFA action, Eq. (10), decreases rapidly as $m_2 \rightarrow m_1$. As a

consequence, the heatbath is also less expensive with Hasenbusch preconditioning, since, as we increase the number of intermediate masses, we can simultaneously decrease the range and number of poles entering into the rational approximation used for each determinant. Table XIV summarizes the measured spectral range, the heatbath error, and the total heatbath cost for each of the runs 2–7. For this ensemble the first Hasenbusch mass reduces the cost of the heatbath by more than a factor of two, and subsequent Hasenbusch masses essentially leave the cost fixed.

For each of the runs 2–7 we generated ten trajectories, beginning from the same seed configuration, and analyzed the resulting distributions of F_{RMS} and F_{max} . In panel (a) of Fig. 8 we plot distributions of F_{max} from 850 trajectories of

<u>RHMC</u>	<u>EOFA</u>
<ul style="list-style-type: none"> • Omelyan integrator ($\delta\tau = 0.0625$) • One light quark Hasenbusch mass • Multishift CG with single precision \mathbb{D} but accumulating solution and search vectors in double precision, coupled with reliable update to correct residual • Even-odd preconditioning 	<ul style="list-style-type: none"> • Force gradient integrator ($\delta\tau = 0.1667$) • Five light quark Hasenbusch masses • Mixed precision defect correction CG • Even-odd preconditioning • Cayley-form preconditioning • Force gradient forecasting • Heatbath forecasting • Heatbath stopping tolerance tuning

FIG. 9. Comparison of optimizations used in the RHMC 32ID-G simulation to the optimizations used in the EOFA 32ID-G simulation.

the production RHMC ensemble generation calculation (run 1). Since we are using exactly the same RHMC action for the strange quark on the RHMC and EOFA ensembles, we tune by adjusting the number and magnitude of the intermediate light Hasenbusch masses such that the forces associated with each of the light quark determinants are comparable to the strange quark force. This allows us to simplify the integrator layout to a three-level scheme, with the light and strange quark determinants updated on the same level. We find that four intermediate Hasenbusch masses are sufficient to ensure that the strange quark force is dominant in the sense of F_{RMS} , and that five intermediate Hasenbusch masses are sufficient in the sense of F_{max} . Panel (b) shows an analogous force distribution for the latter mass preconditioning scheme.

In runs 8–10 we explore further tuning of a scheme with four light Hasenbusch masses, and in runs 11–14 we explore further tuning of a scheme with five light Hasenbusch masses. We note that the Monte Carlo acceptance is relatively poor in runs 8–10—as we argued in Sec. V, this is consistent with the view that the acceptance should be controlled by the largest integration errors accrued during the trajectory, which are proportional to F_{max} rather than F_{RMS} —and thus have abandoned this mass preconditioning scheme in favor of the scheme used in runs 11–14. We have then tuned the step size of the outermost integrator level ($\delta\tau$) and the CG stopping tolerance used in the molecular dynamics evolution (r_{MD}) to minimize the mean time required to generate an accepted gauge field configuration, resulting in the scheme of run 12. In addition, we have applied the heatbath tuning procedure described in Sec. VIB in all of the runs 8–14, allowing us to relax CG stopping tolerances for the

individual solves in the heatbath, while keeping the overall error bounded by $r_{\text{MC}} = 10^{-10}$. For the final scheme (12) this optimization further reduced the cost of the light quark heatbath from approximately 2300 s, as reported in Table XIV, to approximately 850 s after tuning.

Comparing the fully tuned EOFA scheme (12) to the original RHMC scheme (1) in Table XIII, we find that we are able to generate EOFA trajectories a factor of 4.2 times faster than RHMC trajectories, while maintaining a slightly higher acceptance rate of 93%. We emphasize, however, that this improved performance is only partially attributable to the simpler form of the EOFA Hamiltonian and force evaluations: we have also switched from an Omelyan integrator to a force gradient integrator, retuned the step sizes and integrator layout, and, in some cases, applied optimizations to the EOFA simulation that are not applicable to RHMC simulations (e.g. mixed precision CG). Figure 9 briefly summarizes the respective techniques used in the RHMC and EOFA evolution schemes. We have now adopted the EOFA scheme tested in run 12 for ensemble generation in our ongoing $\Delta I = 1/2 K \rightarrow \pi\pi$ calculation [38]. We expect the resulting performance gain to enable up to four times as many measurements in our current production run as we would have been able to generate using the initial RHMC evolution scheme, enabling a significantly more precise first-principles determination of the standard model ratio ϵ'/ϵ .

VIII. CONCLUSION

In this work we have explored the viability of the exact one flavor algorithm (EOFA) as an alternative to the rational Hybrid Monte Carlo (RHMC) algorithm in

molecular dynamics simulations of lattice QCD with domain wall fermions and periodic or G -parity boundary conditions. We have verified the formal equivalence of EOFA to RHMC through statistical tests of the EOFA action (Sec. IV), and checked, using a series of inexpensive, small volume ensembles with heavy pions, that physical observables such as the plaquette, quark condensates, topological susceptibility, and low energy spectrum are consistent between ensembles generated using EOFA and ensembles generated using RHMC (Sec. V). We have then discussed preconditioning and tuning techniques for EOFA simulations (Sec. VI and Appendix B), and finally, demonstrated that EOFA can substantially outperform RHMC for state-of-the-art lattice QCD simulations with large volumes and physical quark masses (Sec. VII). In particular, we find that we are able to generate gauge field configurations for the ongoing RBC/UKQCD calculation of the $\Delta I = 1/2$ $K \rightarrow \pi\pi$ decay amplitudes a factor of 4.2 times faster with EOFA. The keys to this dramatic speed-up are a novel preconditioning technique which relates inversions of the EOFA Dirac operator ($\mathcal{D}_{\text{EOFA}}$) to cheaper inversions of the standard domain wall fermion Dirac operator (\mathcal{D}_{DWF}), and the ability to apply mixed precision defect correction solvers and extensive Hasenbusch mass preconditioning in the context of EOFA.

Future work will explore further physics applications of EOFA. We intend to generate variants of the 24ID ensemble with nondegenerate up and down quark masses in the near future. These ensembles will enable exploratory studies of isospin breaking effects in the meson and baryon spectra, as well as in other precision lattice calculations such as the extraction of the CKM matrix element V_{us} from semi-leptonic kaon decays [39]. Other potential applications include domain wall QCD simulations with dynamical charm quarks in the sea, and simulations with light, $SU(3)$ -symmetric quarks. The latter simulations could be used, for example, to better constrain the strange quark dependence of our $SU(3)$ chiral perturbation theory studies [40], or to probe the location of the critical point separating the crossover and first-order phase transition regions in three-flavor domain wall QCD at finite temperature.

ACKNOWLEDGMENTS

The authors are grateful to Peter Boyle, Ting-Wai Chiu, and Norman Christ for helpful discussions in support of this work. Calculations were performed using the Blue Gene/Q computers of the RIKEN-BNL Research Center and Brookhaven National Lab. The software used includes the CPS QCD code (<https://github.com/RBC-UKQCD/CPS>) [41], supported in part by the USDOE SciDAC program, the BAGEL (<http://www2.ph.ed.ac.uk/paboyle/bagel/Bagel.html>) assembler kernel generator for high-performance optimized kernels and fermion solvers [36], and the Grid data parallel C++ QCD library (<https://github.com/paboyle/Grid>) [42]. R. D. M. and D. J. M. are

supported in part by U.S. DOE Grant No. DE-SC0011941. C. J. is supported in part by U.S. DOE Contract No. AC-02-98CH10886 (BNL). C. K. is supported by the Intel Corporation.

APPENDIX A: EOFA OPERATORS FOR SHAMIR AND MÖBIUS DWF

In this Appendix we list the operators which enter into $\mathcal{D}_{\text{EOFA}}$ [Eq. (6)] and the EOFA action [Eq. (10)]. The more general case of DWF with Zolotarev-type domain wall fermions is constructed implicitly in Ref. [4]; we explicitly list these operators for the more restrictive cases of Shamir and Möbius DWF used in our simulations. We use Θ_s to denote the discrete Heaviside theta function

$$\Theta_s = \begin{cases} 0, & s < 0 \\ 1, & s \geq 0 \end{cases} \quad (\text{A1})$$

and assume even L_s . The operators Ω_{\pm} and Δ_{\pm} are related by the identity

$$\Delta_{\pm} = k\Omega_{\pm}\Omega_{\pm}^{\dagger}, \quad (\text{A2})$$

and the Möbius operators reduce to the corresponding Shamir operators in the limit $c = d = 1/2$. We note that the dense Möbius expressions listed here are not used inside the inverter; we instead invert the preconditioned system discussed in Appendix B 2.

1. Shamir kernel

$$k = m_2 - m_1 \quad (\text{A3})$$

$$[\Omega_+]_{ss'} = \delta_{s,L_s-1}\delta_{s',0} \quad (\text{A4})$$

$$[\Omega_-]_{ss'} = \delta_{s,0}\delta_{s',0} \quad (\text{A5})$$

$$[\Delta_+(m_1, m_2)]_{ss'} = (m_2 - m_1)\delta_{s,L_s-1}\delta_{s',L_s-1} \quad (\text{A6})$$

$$[\Delta_-(m_1, m_2)]_{ss'} = (m_2 - m_1)\delta_{s,0}\delta_{s',0} \quad (\text{A7})$$

$$[M_+(m)]_{ss'} = \delta_{ss'} - \delta_{s,s'+1} + m\delta_{s,L_s-1}\delta_{s',0} \quad (\text{A8})$$

$$[M_-(m)]_{ss'} = \delta_{ss'} - \delta_{s,s'-1} + m\delta_{s,0}\delta_{s',L_s-1} \quad (\text{A9})$$

$$[\tilde{\mathcal{D}}(m)]_{ss'} = \delta_{ss'} \quad (\text{A10})$$

$$[\tilde{\mathcal{D}}(m)^{-1}]_{ss'} = \delta_{ss'} \quad (\text{A11})$$

2. Möbius kernel

$$k = \frac{2c(m_2 - m_1)(c + d)^{2L_s}}{[(c + d)^{L_s} + m_1(c - d)^{L_s}][(c + d)^{L_s} + m_2(c - d)^{L_s}]} \quad (\text{A12})$$

$$(\Omega_+)_{ss'} = (-1)^{s+1} \frac{(c - d)^{L_s - s - 1}}{(c + d)^{L_s - s}} \delta_{s',0} \quad (\text{A13})$$

$$(\Omega_-)_{ss'} = (-1)^s \frac{(c - d)^s}{(c + d)^{s+1}} \delta_{s',0} \quad (\text{A14})$$

$$[\Delta_+(m_1, m_2)]_{ss'} = \frac{(-1)^{s+s'} 2c(m_2 - m_1)(c + d)^{s+s'} (c - d)^{2(L_s - 1) - s - s'}}{[(c + d)^{L_s} + m_1(c - d)^{L_s}][(c + d)^{L_s} + m_2(c - d)^{L_s}]} \quad (\text{A15})$$

$$[\Delta_-(m_1, m_2)]_{ss'} = \frac{(-1)^{s+s'} 2c(m_2 - m_1)(c + d)^{2(L_s - 1) - s - s'} (c - d)^{s+s'}}{[(c + d)^{L_s} + m_1(c - d)^{L_s}][(c + d)^{L_s} + m_2(c - d)^{L_s}]} \quad (\text{A16})$$

$$\begin{aligned} [M_+(m)]_{ss'} &= \frac{(-1)^{s-s'} 2c(c + d)^{L_s - s + s' - 1} (c - d)^{s - s' - 1}}{(c + d)^{L_s} + m(c - d)^{L_s}} \Theta_{s-s'-1} \\ &+ \frac{(c + d)^{L_s - 1} - m(c - d)^{L_s - 1}}{(c + d)^{L_s} + m(c - d)^{L_s}} \delta_{ss'} \\ &+ \frac{(-1)^{s-s'+1} 2cm(c + d)^{s'-s-1} (c - d)^{L_s + s - s' - 1}}{(c + d)^{L_s} + m(c - d)^{L_s}} \Theta_{s'-s-1} \end{aligned} \quad (\text{A17})$$

$$\begin{aligned} [M_-(m)]_{ss'} &= \frac{(-1)^{s'-s+1} 2cm(c + d)^{s-s'-1} (c - d)^{L_s - s + s' - 1}}{(c + d)^{L_s} + m(c - d)^{L_s}} \Theta_{s-s'-1} \\ &+ \frac{(c + d)^{L_s - 1} - m(c - d)^{L_s - 1}}{(c + d)^{L_s} + m(c - d)^{L_s}} \delta_{ss'} \\ &+ \frac{(-1)^{s-s'} 2c(c + d)^{L_s + s - s' - 1} (c - d)^{s'-s-1}}{(c + d)^{L_s} + m(c - d)^{L_s}} \Theta_{s'-s-1} \end{aligned} \quad (\text{A18})$$

$$\begin{aligned} [\tilde{\mathcal{D}}(m)]_{ss'} &= (c + d)\delta_{ss'} + (c - d)P_+\delta_{s,s'+1} + (c - d)P_-\delta_{s,s'-1} \\ &- m(c - d)P_+\delta_{s,0}\delta_{s',L_s-1} - m(c - d)P_-\delta_{s,L_s-1}\delta_{s',0} \end{aligned} \quad (\text{A19})$$

$$\begin{aligned} [\tilde{\mathcal{D}}(m)^{-1}]_{ss'} &= \left[\frac{m(-1)^{s-s'+1}(c + d)^{s'-s-1}(c - d)^{L_s + s - s'}}{(c + d)^{L_s} + m(c - d)^{L_s}} + \frac{(-1)^{s-s'}(c - d)^{s-s'}}{(c + d)^{s-s'+1}} \Theta_{s-s'} \right] P_+ \\ &+ \left[\frac{m(-1)^{s'-s+1}(c + d)^{s-s'-1}(c - d)^{L_s + s' - s}}{(c + d)^{L_s} + m(c - d)^{L_s}} + \frac{(-1)^{s'-s}(c - d)^{s'-s}}{(c + d)^{s'-s+1}} \Theta_{s'-s} \right] P_- \end{aligned} \quad (\text{A20})$$

APPENDIX B: FOUR-DIMENSIONAL EVEN-ODD PRECONDITIONING

The inversions required to compute the exact one flavor Hamiltonian can be accelerated using a standard checkerboarding technique: we label lattice sites as “even” if $(x + y + z + t) \equiv 0 \pmod{2}$ or “odd” if $(x + y + z + t) \equiv 1 \pmod{2}$. This naturally induces a block structure in the Dirac operator \mathcal{D} , which can be LDU decomposed as

$$\underbrace{\begin{pmatrix} \mathcal{D}_{ee} & \mathcal{D}_{eo} \\ \mathcal{D}_{oe} & \mathcal{D}_{oo} \end{pmatrix}}_{\mathcal{D}} = \underbrace{\begin{pmatrix} 1 & 0 \\ \mathcal{D}_{oe}\mathcal{D}_{ee}^{-1} & 1 \end{pmatrix}}_L \underbrace{\begin{pmatrix} \mathcal{D}_{ee} & 0 \\ 0 & \mathcal{D}_{oo} - \mathcal{D}_{oe}\mathcal{D}_{ee}^{-1}\mathcal{D}_{eo} \end{pmatrix}}_D \underbrace{\begin{pmatrix} 1 & \mathcal{D}_{ee}^{-1}\mathcal{D}_{eo} \\ 0 & 1 \end{pmatrix}}_U. \quad (\text{B1})$$

Left-multiplying the linear system $\mathcal{D}\psi = \phi$ by

$$L^{-1} = \begin{pmatrix} 1 & 0 \\ -\mathcal{D}_{oe}\mathcal{D}_{ee}^{-1} & 1 \end{pmatrix} \quad (\text{B2})$$

results in the equivalent system

$$\begin{pmatrix} \mathcal{D}_{ee}\psi_e + \mathcal{D}_{eo}\psi_o \\ (\mathcal{D}_{oo} - \mathcal{D}_{oe}\mathcal{D}_{ee}^{-1}\mathcal{D}_{eo})\psi_o \end{pmatrix} = \begin{pmatrix} \phi_e \\ \phi_o - \mathcal{D}_{oe}\mathcal{D}_{ee}^{-1}\phi_e \end{pmatrix}, \quad (\text{B3})$$

leading to the following trick: assuming \mathcal{D}_{ee}^{-1} is available in an explicit form, it suffices to invert

$$(\mathcal{D}_{oo} - \mathcal{D}_{oe}\mathcal{D}_{ee}^{-1}\mathcal{D}_{eo})\psi_o = \tilde{\phi}_o, \quad (\text{B4})$$

with $\tilde{\phi}_o \equiv \phi_o - \mathcal{D}_{oe}\mathcal{D}_{ee}^{-1}\phi_e$. This system only involves the odd sublattice, and is thus substantially cheaper to invert than \mathcal{D} using an iterative algorithm like CG. The solution on the even sublattice can then be reconstructed for a trivial additional cost as

$$\psi_e = \mathcal{D}_{ee}^{-1}(\phi_e - \mathcal{D}_{eo}\psi_o). \quad (\text{B5})$$

This technique is already well understood in the context of RHMC with Shamir or Möbius DWF; in this Appendix we describe how to generalize the method to the exact one flavor algorithm.

In the context of EOFA, the generic linear system one needs to invert takes the form

$$(H(m_1) + \beta\Delta_{\pm}(m_2, m_3)P_{\pm})\psi = \phi. \quad (\text{B6})$$

We choose to multiply by an overall factor of $\gamma_5 R_5$, rewriting the system as

$$(\mathcal{D}_{\text{EOFA}}(m_1) + \beta\gamma_5 R_5 \Delta_{\pm}(m_2, m_3)P_{\pm})\psi = \gamma_5 R_5 \phi, \quad (\text{B7})$$

for the following reasons: first, we wish to re-use the existing high-performance implementation of the Wilson \mathcal{D} kernel in the BAGEL library without modification, and second, overall factors of $\gamma_5 R_5$ will cancel inside the inverter since we use CG applied to the normal equations and $(\gamma_5 R_5)^{\dagger}(\gamma_5 R_5) = 1$. Since $(\mathcal{D}_{\text{EOFA}})_{eo} = (\mathcal{D}_{\text{DWF}})_{eo}$ and $\Delta_{\pm} \propto \delta_{xx'}$ in the 4D bulk, only the operators coupling sites of the same parity need to be modified to implement even-odd preconditioned EOFA. We take somewhat different approaches for the Shamir and Möbius cases.

1. Shamir kernel

Recall that for the Shamir kernel $\mathcal{D}_{\text{DWF}} = \mathcal{D}_{\text{EOFA}}$, so the extension of an inverter for the even-odd preconditioned \mathcal{D}_{DWF} operator to instead solve Eq. (B7) is straightforward. With $\mathcal{D} = \mathcal{D}_{\text{DWF}}$, the same parity fermion matrix has the tridiagonal block structure

$$\begin{aligned} (\mathcal{D}_{\text{DWF}})_{ee} &= (\mathcal{D}_{\text{DWF}})_{oo} \\ &= \delta_{xx'} \{ (5 - M_5) \delta_{ss'} - P_+ \delta_{s,s'+1} - P_- \delta_{s,s'-1} \\ &\quad + m_1 P_+ \delta_{s,0} \delta_{s',L_s-1} + m_1 P_- \delta_{s,L_s-1} \delta_{s',0} \}. \end{aligned} \quad (\text{B8})$$

One can check by explicit calculation that the shift operators have the form

$$\begin{cases} \beta\gamma_5 R_5 \Delta_+(m_2, m_3) P_+ = \beta(m_3 - m_2) P_+ \delta_{xx'} \delta_{s,0} \delta_{s',L_s-1} \\ \beta\gamma_5 R_5 \Delta_-(m_2, m_3) P_- = -\beta(m_3 - m_2) P_- \delta_{xx'} \delta_{s,L_s-1} \delta_{s',0} \end{cases}, \quad (\text{B9})$$

so one can consider the operator appearing in Eq. (B7) as a slight generalization of Eq. (B8) to

$$\begin{aligned} \mathcal{D}_{ee} = \mathcal{D}_{oo} &= \delta_{xx'} \{ (5 - M_5) \delta_{ss'} - P_+ \delta_{s,s'+1} - P_- \delta_{s,s'-1} \\ &\quad + d_+ P_+ \delta_{s,0} \delta_{s',L_s-1} + d_- P_- \delta_{s,L_s-1} \delta_{s',0} \}, \end{aligned} \quad (\text{B10})$$

with

$$d_- = m_1 - \beta(m_3 - m_2) \delta_{i,-} \quad (\text{B11})$$

and

$$d_+ = m_1 + \beta(m_3 - m_2) \delta_{i,+}, \quad (\text{B12})$$

where the index i denotes the chirality of the shift operator. \mathcal{D}_{ee}^{-1} can be efficiently applied using the LDU decomposition of \mathcal{D}_{ee} , again as a slight generalization of the standard Shamir DWF case.

2. Möbius kernel and Cayley-form preconditioning

Using Eq. (7) we can write $\mathcal{D}_{\text{EOFA}}$ in the form

$$(\mathcal{D}_{\text{EOFA}})_{xx',ss'} = (D_W)_{xx'} \delta_{ss'} + \delta_{xx'} (\mathcal{D}^{\perp})_{ss'}. \quad (\text{B13})$$

The action of the operator appearing in Eq. (B7) on lattice sites of the same parity, then, is given by

$$\begin{aligned} \mathcal{D}_{ee} = \mathcal{D}_{oo} = & \delta_{xx'} \{ (4 - M_5) \delta_{ss'} + (M_+(m_1))_{ss'} P_+ \\ & + (M_-(m_1))_{ss'} P_- + \beta \gamma_5 R_5 (\Delta_{\pm}(m_2, m_3))_{ss'} P_{\pm} \}, \end{aligned} \quad (\text{B14})$$

with M_+ , M_- , and Δ_{\pm} as defined in Eqs. (A15)–(A18). The matrix elements of $\mathcal{D}_{ee}^{-1} = \mathcal{D}_{oo}^{-1}$ can be found by explicit numerical inversion as part of the setup cost; this is a trivial overhead since it suffices to invert only the ss' subblock of \mathcal{D}_{ee} . In this form, the exact factorization of the fermion determinant in Eq. (2) comes at the cost of dense $L_s \times L_s$ matrix operations. We argue that it is possible to do significantly better by introducing an additional preconditioning step.

We note that the system defined by Eq. (B7) can be more efficiently inverted for the case of Möbius DWF by using the operator $\tilde{\mathcal{D}}^{-1}$ as a right preconditioner, resulting in an equivalent system in terms of \mathcal{D}_{DWF} . For the special case $\beta = 0$ this is straightforward: observing that the relationship between $\mathcal{D}_{\text{EOFA}}$ and \mathcal{D}_{DWF} [Eq. (6)] can be used to manipulate

$$\mathcal{D}_{\text{EOFA}} \psi = \mathcal{D}_{\text{EOFA}} \cdot \tilde{\mathcal{D}} \cdot \underbrace{\tilde{\mathcal{D}}^{-1} \psi}_{\equiv \psi'} = \mathcal{D}_{\text{DWF}} \psi', \quad (\text{B15})$$

it suffices to solve $\mathcal{D}_{\text{DWF}} \psi' = \gamma_5 R_5 \phi$, from which $\psi = \tilde{\mathcal{D}} \psi'$ can be recovered at the cost of a single additional matrix multiplication. While we observe that $\mathcal{D}_{\text{DWF}}^{\dagger} \mathcal{D}_{\text{DWF}}$ has a slightly larger condition number than $\mathcal{D}_{\text{EOFA}}^{\dagger} \mathcal{D}_{\text{EOFA}}$, leading to a modest increase in the total number of CG iterations required to invert the system, \mathcal{D}_{DWF} also has a tridiagonal stencil in the fifth dimension, and can thus be applied in $\mathcal{O}(L_s)$ operations—unlike the $\mathcal{O}(L_s^2)$ operations required for the dense $\mathcal{D}_{\text{EOFA}}$ —leading to a substantial reduction in wall clock time for the inversion.

The $\beta \neq 0$ case is more involved, but can be treated in a similar manner. Right preconditioning Eq. (B7) with $\tilde{\mathcal{D}}^{-1}$ leads to

$$(\mathcal{D}_{\text{DWF}}(m_1) \pm \beta R_5 \Delta_{\pm}(m_2, m_3) \tilde{\mathcal{D}} P_{\pm}) \psi' = \gamma_5 R_5 \phi, \quad (\text{B16})$$

where we have used $\gamma_5 P_{\pm} = \pm P_{\pm}$. We define a new, preconditioned, shift operator $\tilde{\Delta}_{\pm}$ by

$$\tilde{\Delta}_{\pm}(m_1, m_2) \equiv R_5 \Delta_{\pm}(m_1, m_2) \tilde{\mathcal{D}} P_{\pm}, \quad (\text{B17})$$

and note that since $(\tilde{\Delta})_{eo} = (\tilde{\Delta})_{oe} = 0$, Eq. (B16) can be inverted efficiently even with $\beta \neq 0$ provided we can apply the operator $(\mathcal{D}_{\text{DWF}})_{ee} \pm \beta \tilde{\Delta}_{\pm}$ and its inverse in $\mathcal{O}(L_s)$ operations. This turns out to be possible after observing that $\tilde{\Delta}_{\pm}$ is rank-one, i.e., it can be written as a vector outer product

$$\tilde{\Delta}_{\pm} = u_{\pm} \otimes v_{\pm}. \quad (\text{B18})$$

To see this, we start by decomposing $\tilde{\mathcal{D}}$ into its chiral components— $\tilde{\mathcal{D}} = \tilde{\mathcal{D}}_+ P_+ + \tilde{\mathcal{D}}_- P_-$ —in terms of which we can also decompose

$$\tilde{\Delta}_{\pm} = R_5 \Delta_{\pm} \tilde{\mathcal{D}}_{\pm} P_{\pm}. \quad (\text{B19})$$

The 5D structure of these operators can be worked out by direct calculation, leading to Eq. (B18), with

$$\begin{cases} (u_+) _s = (-1)^s \frac{(c-d)^s}{(c+d)^{L_s+s+1}} ((c+d)^{L_s} + m_1(c-d)^{L_s}) \\ (v_+) _s = k \delta_{s, L_s-1} \\ (u_-) _s = (-1)^{s+1} \frac{(c-d)^{L_s-1-s}}{(c+d)^{2L_s-s}} ((c+d)^{L_s} + m_1(c-d)^{L_s}) \\ (v_-) _s = k \delta_{s,0} \end{cases}. \quad (\text{B20})$$

Matrix-vector products involving the preconditioned shift operator and a pseudofermion field can be computed from this decomposition as

$$\begin{cases} (\tilde{\Delta}_+ \psi) _s = k (u_+) _s P_+ \psi_{L_s-1} \\ (\tilde{\Delta}_+^{\dagger} \psi) _s = k \delta_{s, L_s-1} P_+ \left[\sum_{s'=0}^{L_s-1} (u_+) _{s'} \psi_{s'} \right] \end{cases} \quad (\text{B21})$$

and

$$\begin{cases} (\tilde{\Delta}_- \psi) _s = k (u_-) _s P_- \psi_0 \\ (\tilde{\Delta}_-^{\dagger} \psi) _s = k \delta_{s,0} P_- \left[\sum_{s'=0}^{L_s-1} (u_-) _{s'} \psi_{s'} \right]. \end{cases} \quad (\text{B22})$$

The inverses can be applied using the Sherman-Morrison formula:

$$\begin{aligned} & ((\mathcal{D}_{\text{DWF}})_{ee} \pm \beta (u_{\pm} \otimes v_{\pm}))^{-1} \\ &= (\mathcal{D}_{\text{DWF}})_{ee}^{-1} \mp \beta \frac{(\mathcal{D}_{\text{DWF}})_{ee}^{-1} (u_{\pm} \otimes v_{\pm}) (\mathcal{D}_{\text{DWF}})_{ee}^{-1}}{1 \pm \beta \langle v_{\pm}, (\mathcal{D}_{\text{DWF}})_{ee}^{-1} u_{\pm} \rangle}. \end{aligned} \quad (\text{B23})$$

In terms of

$$x_{\pm} \equiv (\mathcal{D}_{\text{DWF}})_{ee}^{-1} u_{\pm}, \quad (\text{B24})$$

which can be constructed numerically using the tridiagonal matrix algorithm [43], the necessary factors can be written as

$$\begin{cases} 1 + \beta \langle v_+, (\mathcal{D}_{\text{DWF}})_{ee}^{-1} u_+ \rangle = 1 + \beta k(x_+)_{L_s-1} \\ ((\mathcal{D}_{\text{DWF}})_{ee}^{-1} (u_{\pm} \otimes v_{\pm}) (\mathcal{D}_{\text{DWF}})_{ee}^{-1} \psi)_s = \frac{k(x_+)_s}{(c+d)^{L_s+m_1} (c-d)^{L_s}} P_+ \left[\sum_{s'=0}^{L_s-1} (c+d)^{s'} (c-d)^{L_s-1-s'} \psi_{s'} \right] \\ ((\mathcal{D}_{\text{DWF}})_{ee}^{-1} (u_+ \otimes v_+) (\mathcal{D}_{\text{DWF}})_{ee}^{-1} \psi)_s = \frac{k(c+d)^s (c-d)^{L_s-1-s}}{(c+d)^{L_s+m_1} (c-d)^{L_s}} P_+ \left[\sum_{s'=0}^{L_s-1} (x_+)_{s'} \psi_{s'} \right] \end{cases} \quad (\text{B25})$$

and

$$\begin{cases} 1 - \beta \langle v_-, (\mathcal{D}_{\text{DWF}})_{ee}^{-1} u_- \rangle = 1 - \beta k(x_-)_0 \\ ((\mathcal{D}_{\text{DWF}})_{ee}^{-1} (u_- \otimes v_-) (\mathcal{D}_{\text{DWF}})_{ee}^{-1} \psi)_s = \frac{k(x_-)_s}{(c+d)^{L_s+m_1} (c-d)^{L_s}} P_- \left[\sum_{s'=0}^{L_s-1} (c+d)^{L_s-1-s'} (c-d)^{s'} \psi_{s'} \right], \\ ((\mathcal{D}_{\text{DWF}})_{ee}^{-1} (u_- \otimes v_-) (\mathcal{D}_{\text{DWF}})_{ee}^{-1} \psi)_s = \frac{k(c+d)^{L_s-1-s} (c-d)^s}{(c+d)^{L_s+m_1} (c-d)^{L_s}} P_- \left[\sum_{s'=0}^{L_s-1} (x_-)_{s'} \psi_{s'} \right] \end{cases} \quad (\text{B26})$$

which allow Eq. (B23) to be applied to a pseudofermion vector in $\mathcal{O}(L_s)$ operations.

In Fig. 10 we benchmark representative even-odd preconditioned inversions of Eq. (B7) on the 24ID ensemble, with and without additional preconditioning by \tilde{D}^{-1} , at the physical strange quark mass. In addition to observing a substantial improvement in terms of wall clock time for the

inversion, we note that this preconditioning scheme also has the advantage that it requires little new code—assuming an existing high-performance implementation of \mathcal{D}_{DWF} —since $\mathcal{D}_{\text{EOFA}}$ is never applied directly in the preconditioned formalism.




FIG. 10. Comparison of wall clock inversion times for the two solves required to evaluate the EOFA Hamiltonian or pseudofermion force with and without Cayley-form preconditioning for the strange quark determinant on the 24ID ensemble. The dashed vertical lines show the corresponding total cost of the multishift inversions of $\mathcal{D}_{\text{DWF}}^\dagger \mathcal{D}_{\text{DWF}}$ needed to evaluate the RHC Hamiltonian or pseudofermion force on the same ensemble.

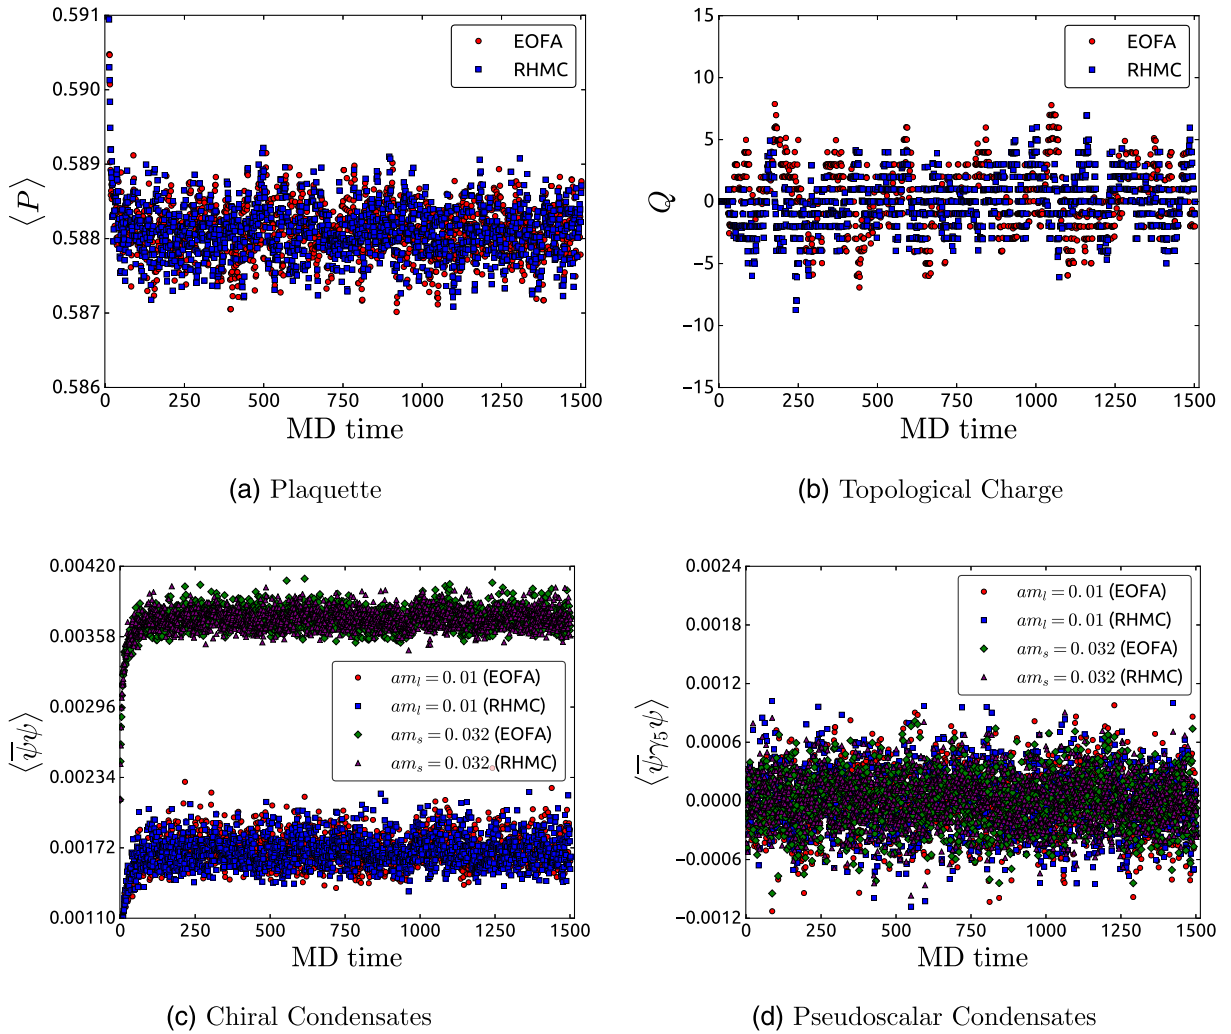
APPENDIX C: ADDITIONAL PLOTS FOR SMALL VOLUME REPRODUCTION TESTS
1. Evolution of the plaquette, quark condensates, and topological charge


FIG. 11. Molecular dynamics evolution of the average plaquette, topological charge, and quark condensates on the 16I ensembles.

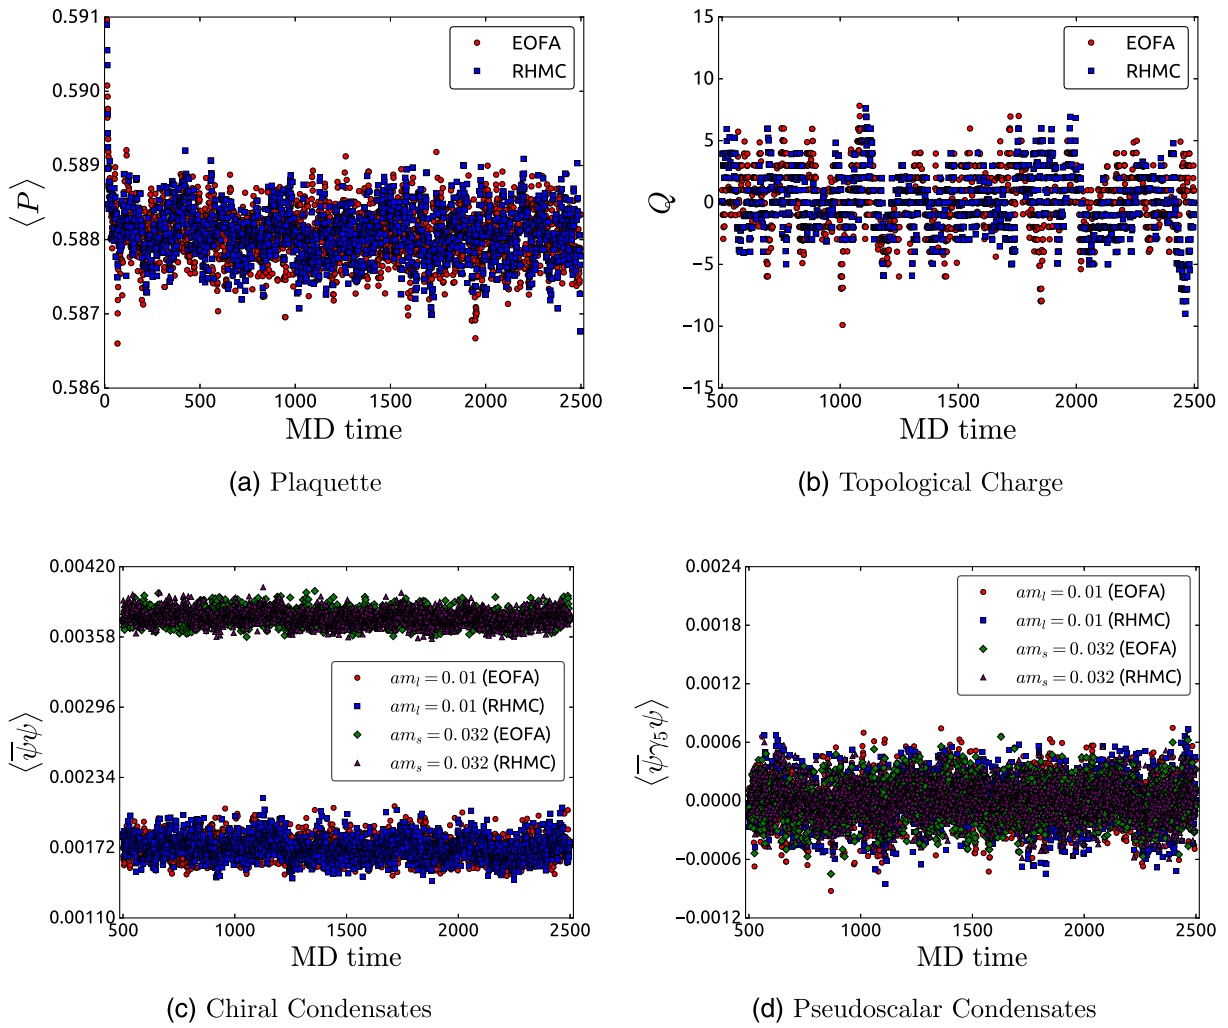


FIG. 12. Molecular dynamics evolution of the average plaquette, topological charge, and quark condensates on the 16I-G ensembles.

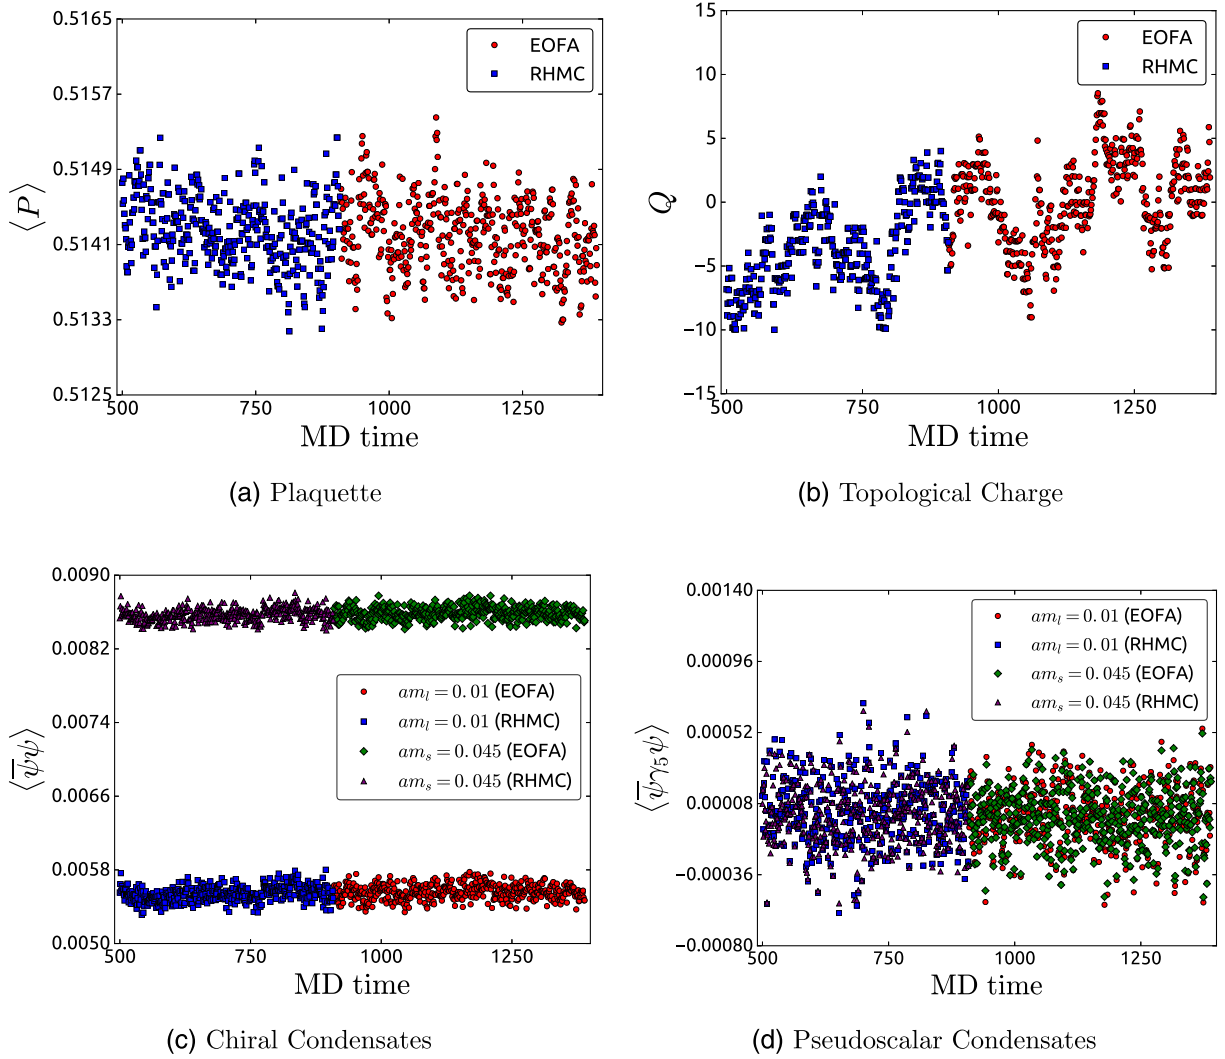


FIG. 13. Molecular dynamics evolution of the average plaquette, topological charge, and quark condensates on the 16ID-G ensembles.

2. Effective mass plots

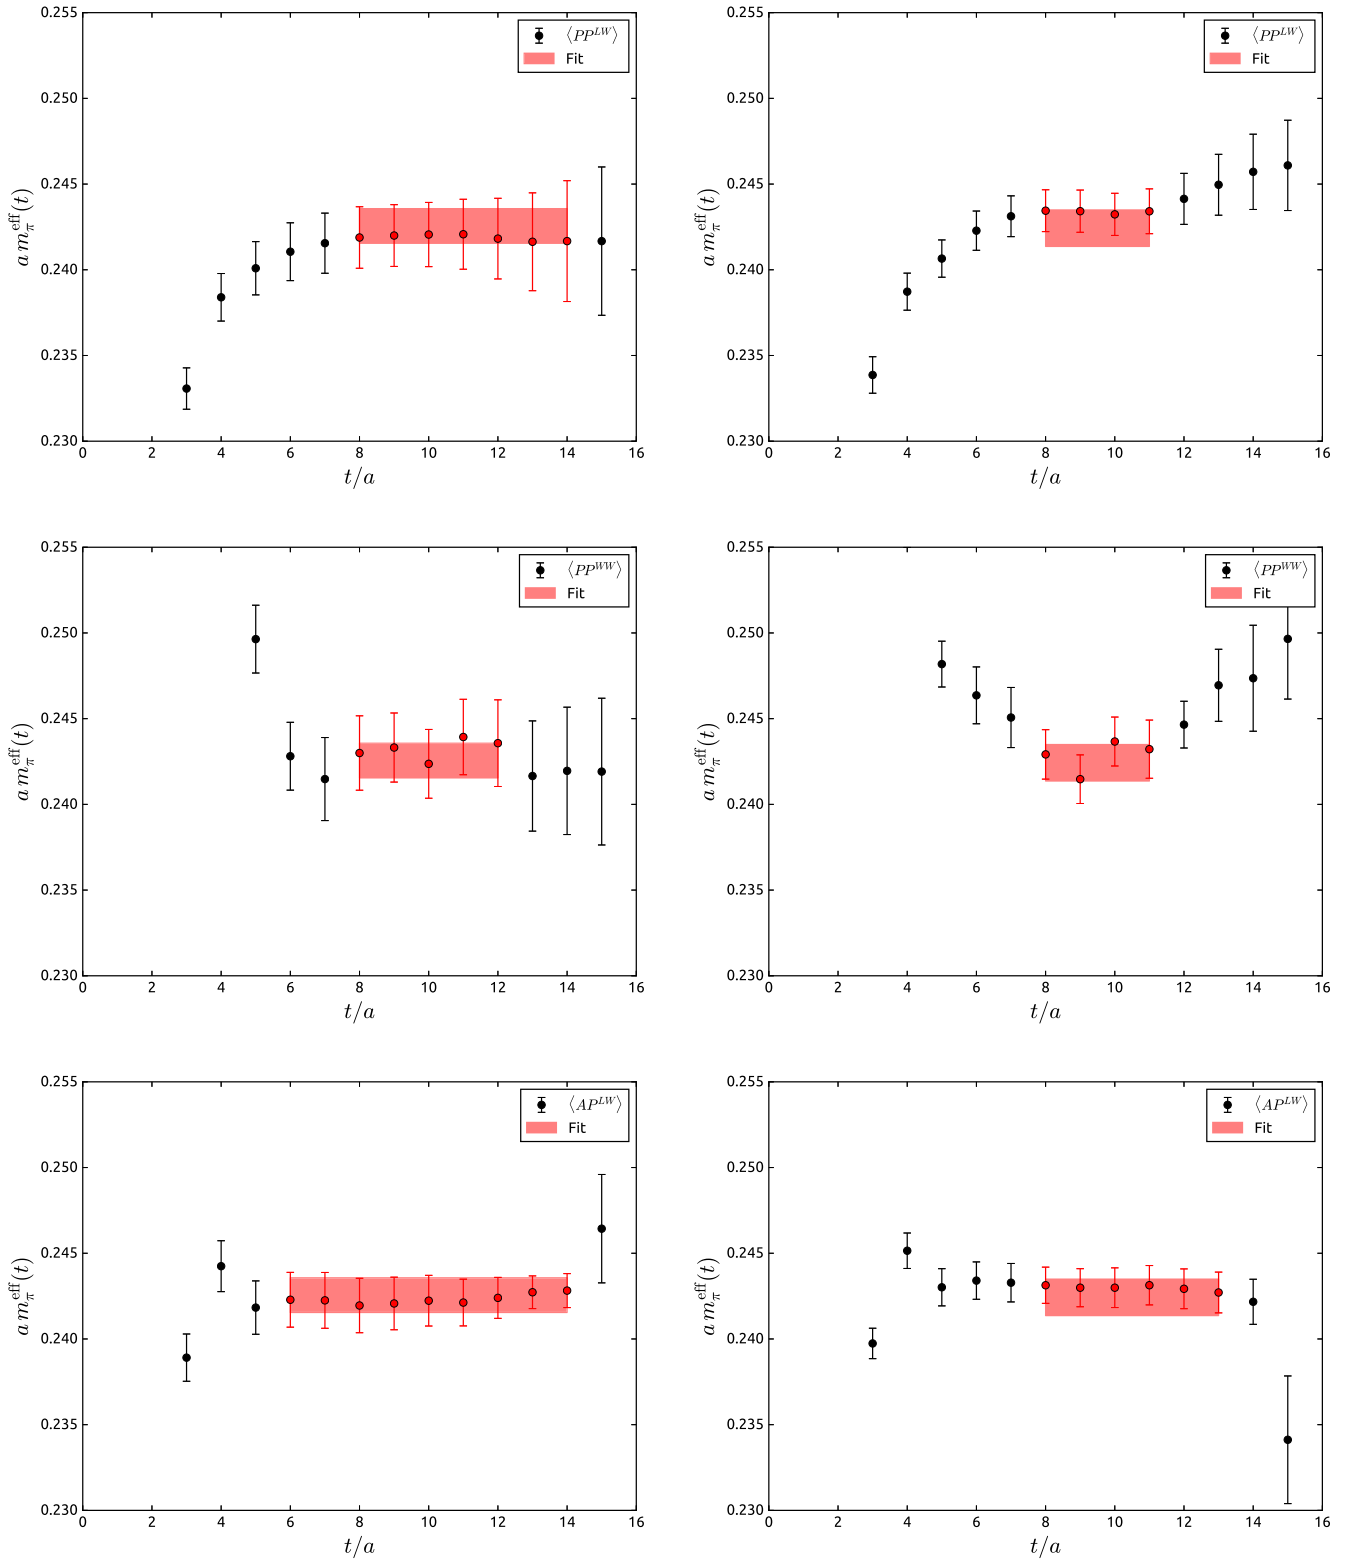


FIG. 14. Effective pion mass from a simultaneous fit to the $\langle PP^{LW} \rangle$ (top), $\langle PP^{WW} \rangle$ (middle), and $\langle AP^{LW} \rangle$ (bottom) correlation functions, as measured on the EOFA (left) and RHMC (right) 16I ensembles.

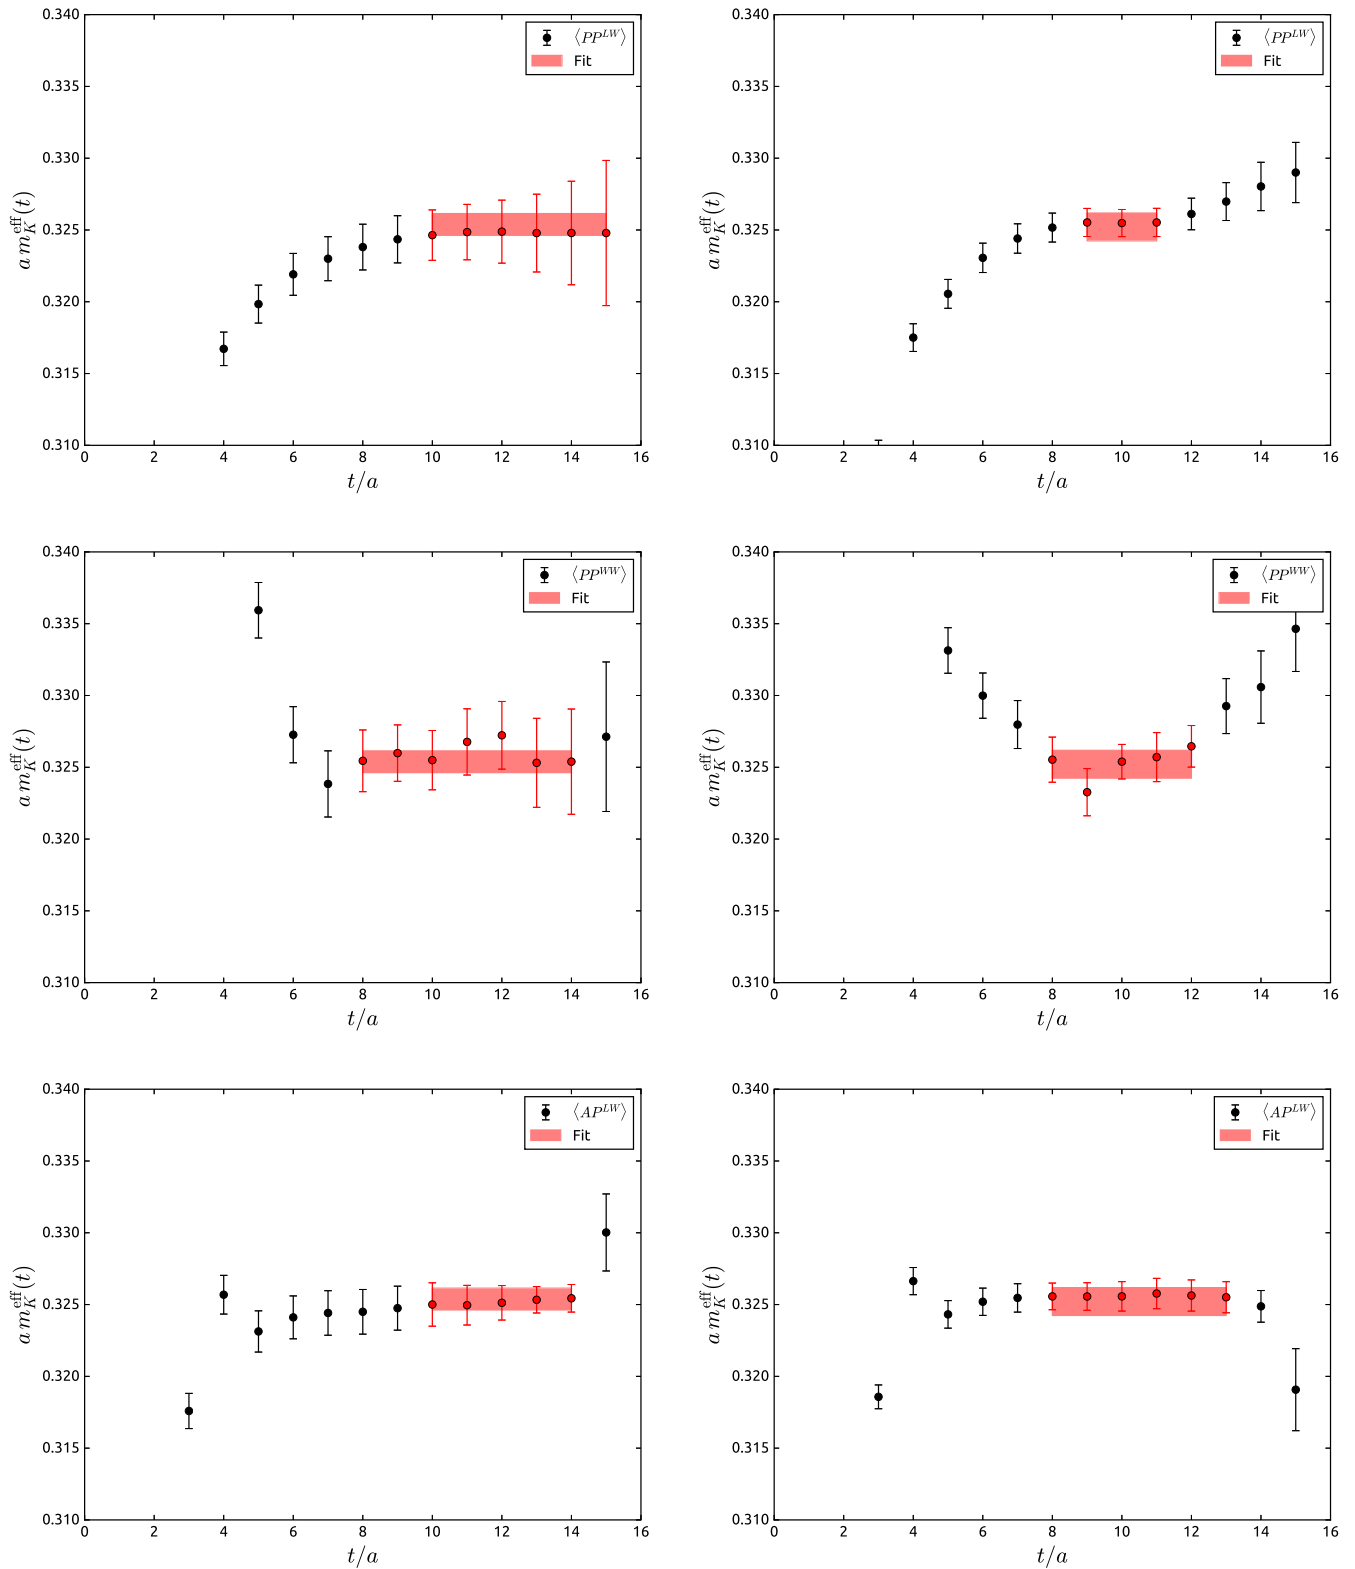


FIG. 15. Effective kaon mass from a simultaneous fit to the $\langle PP^{LW} \rangle$ (top), $\langle PP^{WW} \rangle$ (middle), and $\langle AP^{LW} \rangle$ (bottom) correlation functions, as measured on the EOFA (left) and RHMC (right) 161 ensembles.

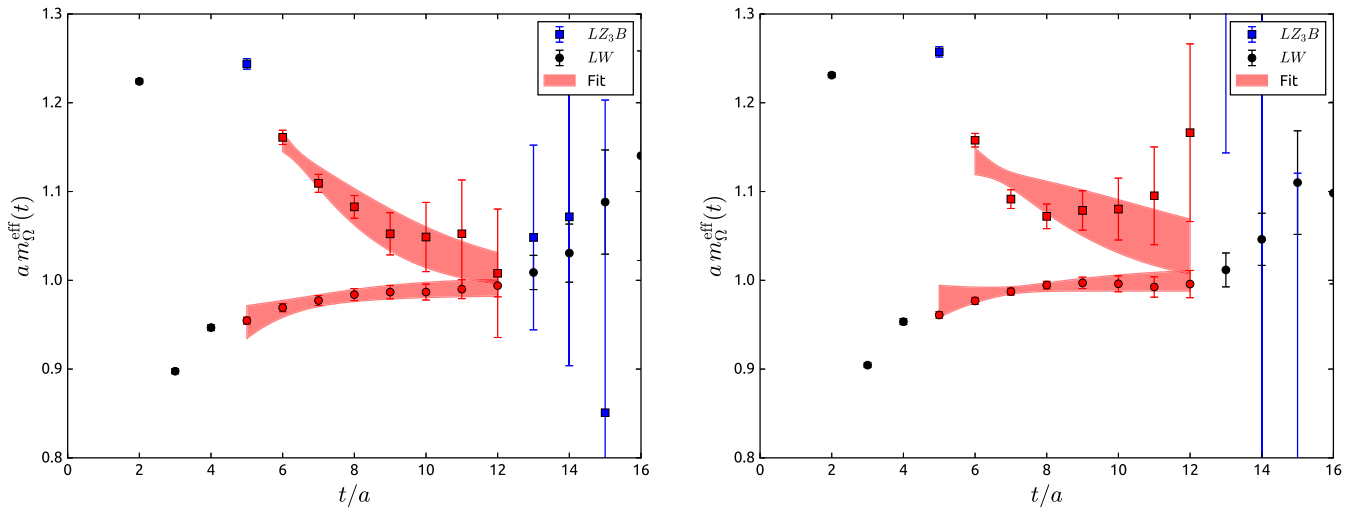


FIG. 16. Effective Ω baryon mass from a simultaneous two-state fit to wall and Z_3 noise sources, as measured on the EOFA (left) and RHMC (right) 16I ensembles.




FIG. 17. Effective $am'_{\text{res}}(m_l)$, as measured on the EOFA (left) and RHMC (right) 16I ensembles.

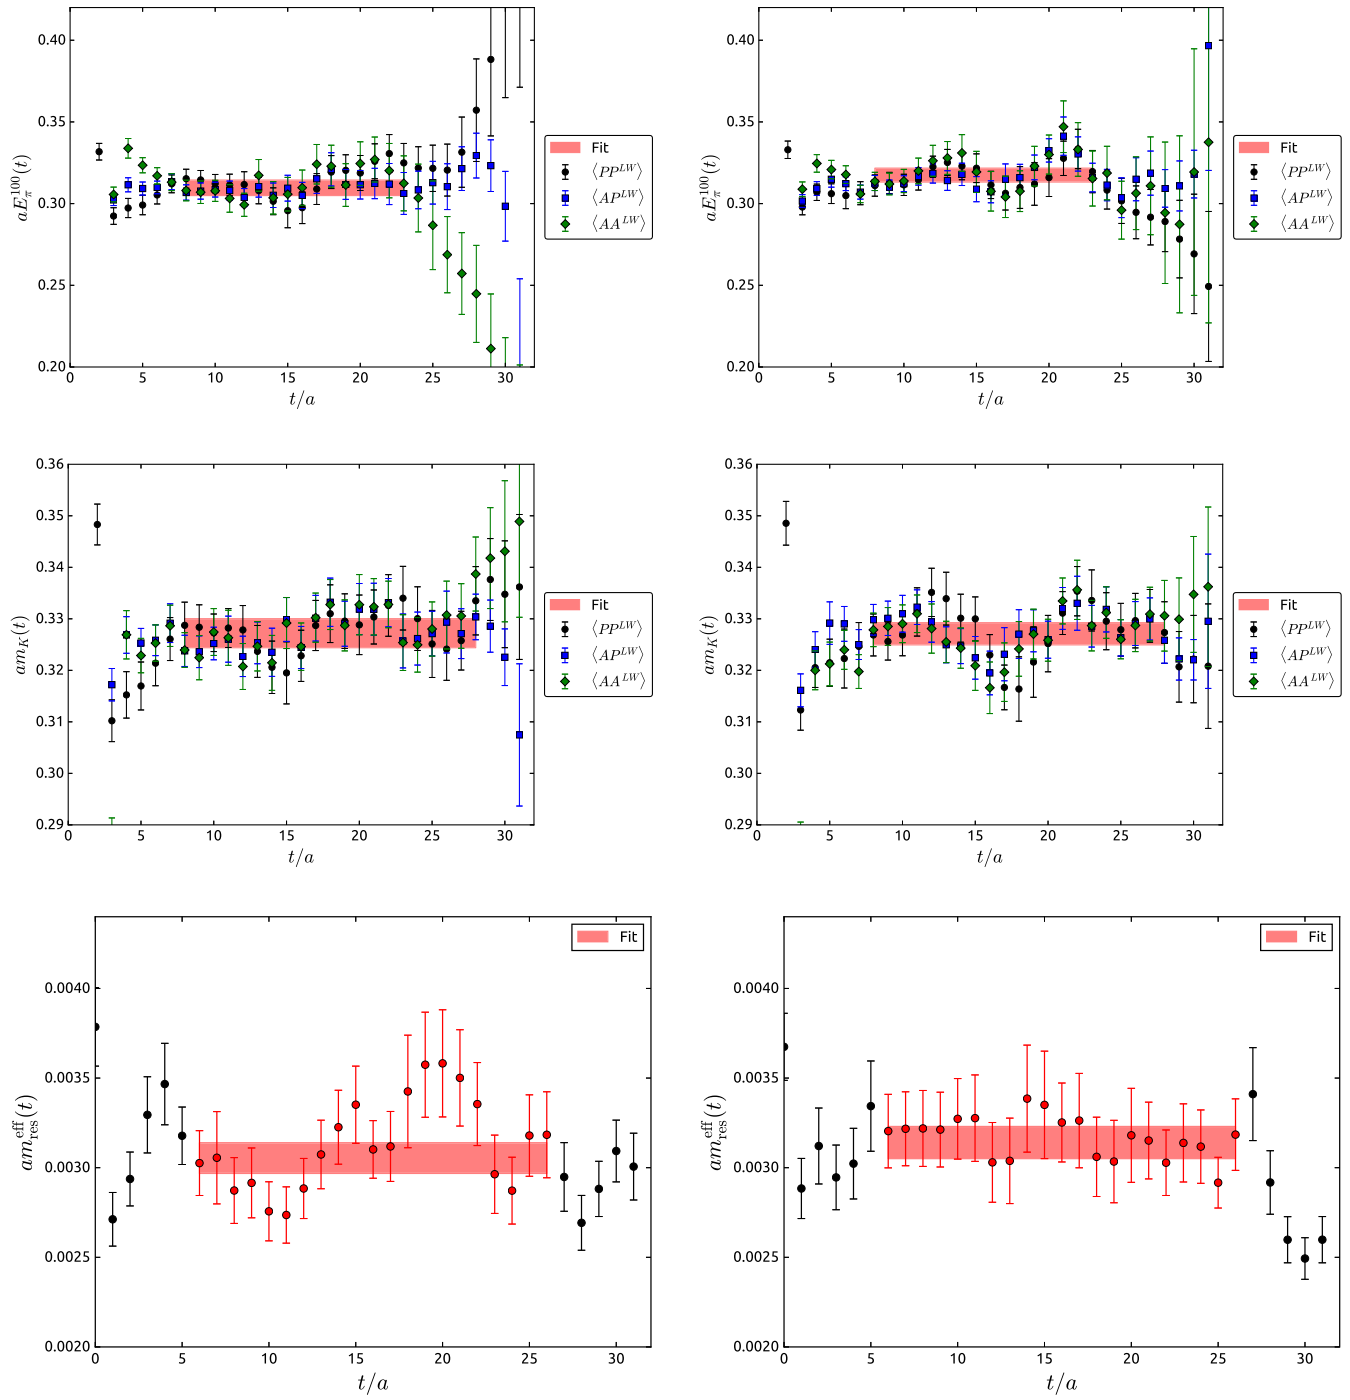


FIG. 18. Effective ground state pion energy (top), kaon mass (middle), and $am'_{\text{res}}(m_l)$ evaluated at the bare light quark mass, as measured on the EOFA (left) and RHMC (right) 16I-G ensembles.

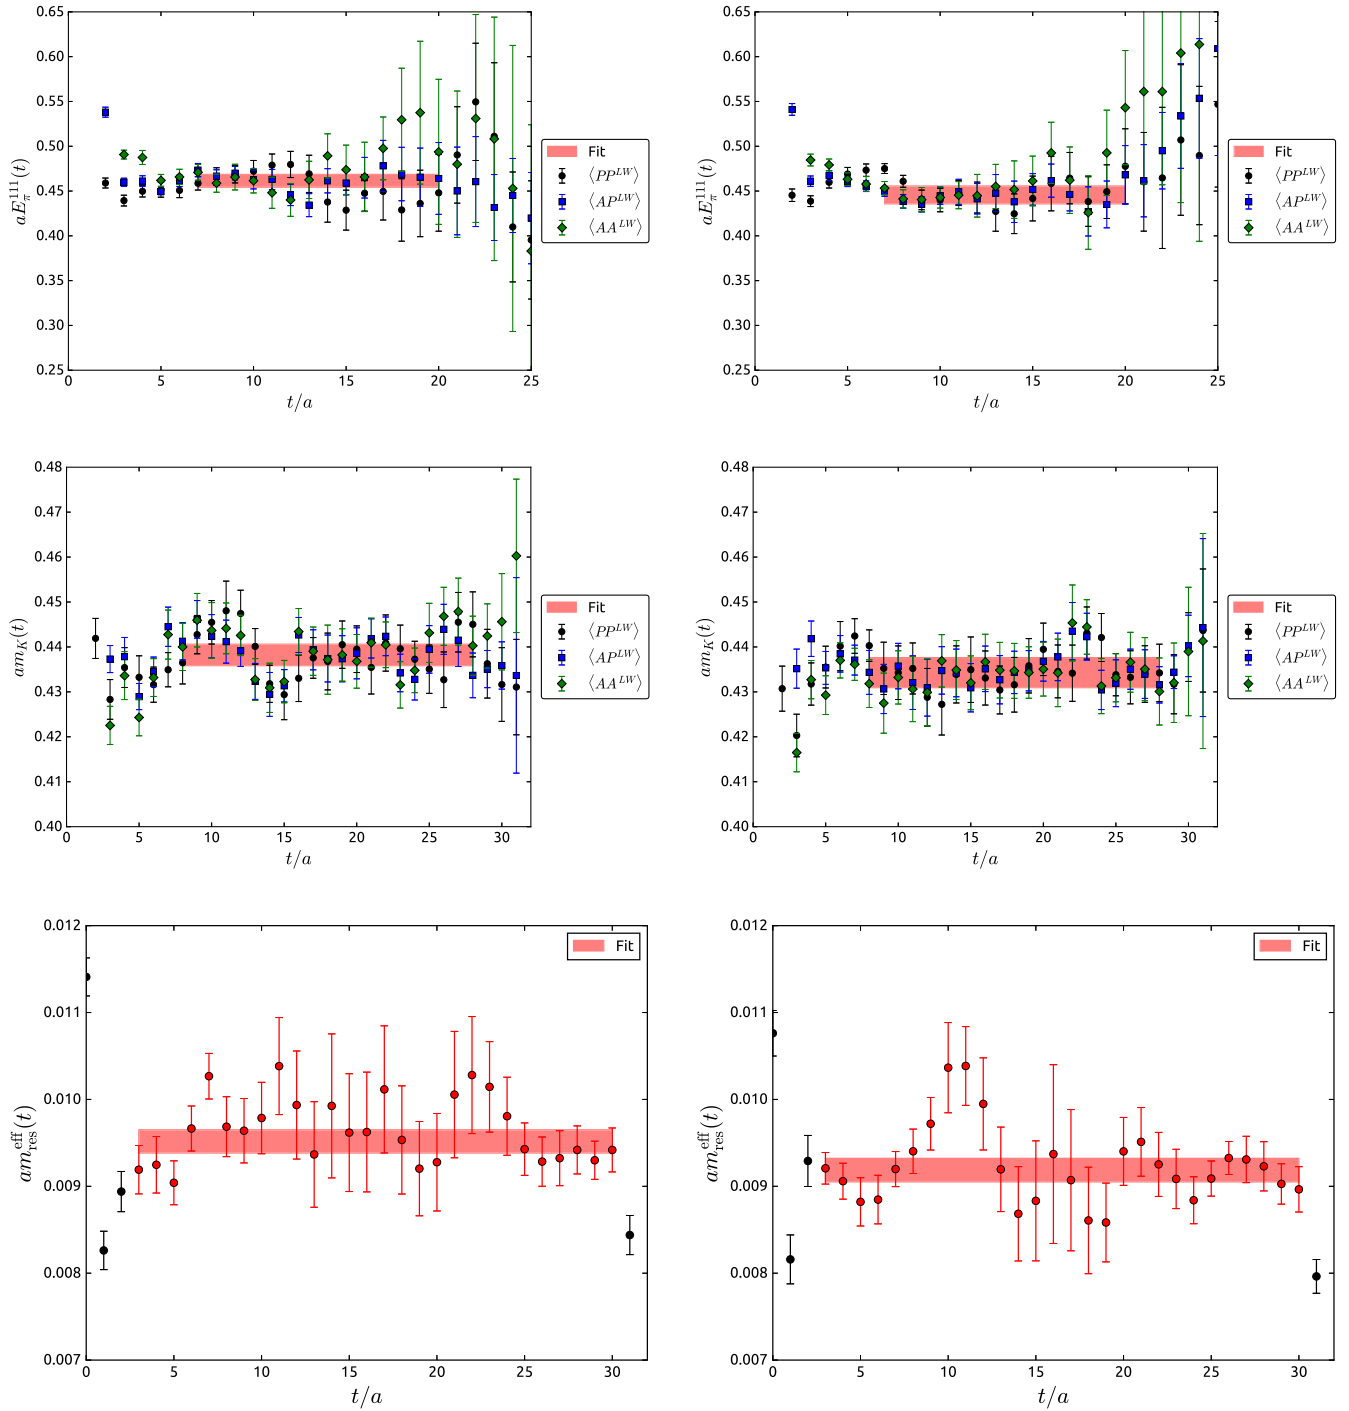


FIG. 19. Effective ground state pion energy (top), kaon mass (middle), and $am_{\text{res}}^{\text{eff}}$ evaluated at the bare light quark mass, as measured on the EOFA (left) and RHMC (right) 16ID-G ensembles.

- [1] M. Hasenbusch and K. Jansen, *Nucl. Phys.* **B659**, 299 (2003).
- [2] G. McGlynn, Ph.D. thesis, Columbia University, 2016.
- [3] We elaborate on the details of our defect correction solver in Sec. VI.
- [4] Y.-C. Chen and T.-W. Chiu (TWQCD Collaboration), *Phys. Lett. B* **738**, 55 (2014).
- [5] Y.-C. Chen and T.-W. Chiu (TWQCD Collaboration), *Proc. Sci.*, IWCSE2013 (2014) 059, [arXiv:1412.0819].
- [6] D.J. Murphy, *Proc. Sci.*, LAT2016 (2016) 272, [arXiv:1611.00298].
- [7] Z. Bai *et al.* (RBC, UKQCD Collaborations), *Phys. Rev. Lett.* **115**, 212001 (2015).
- [8] T. Blum *et al.*, *Phys. Rev. D* **91**, 074502 (2015).
- [9] C. Kelly, arXiv:1310.0434.
- [10] H. Yin and R. D. Mawhinney, *Proc. Sci.*, LATTICE2011 (2011) 051, [arXiv:1111.5059].
- [11] K. Ogawa, T.-W. Chiu and T.-H. Hsieh (TWQCD Collaboration), *Proc. Sci.*, LAT2009 (2009) 033, [arXiv:0911.5532].
- [12] Y.-C. Chen and T.-W. Chiu (TWQCD Collaboration), *Phys. Lett. B* **767**, 193 (2017).
- [13] T. Blum *et al.* (RBC, UKQCD Collaborations), *Phys. Rev. D* **93**, 074505 (2016).
- [14] We emphasize that the operator we call $\mathcal{D}_{\text{EofA}}$ is merely an alternative notation for TWQCD's rescaled Dirac operator D_T [4].
- [15] R. C. Brower, H. Neff, and K. Orginos, *Nucl. Phys. B, Proc. Suppl.* **140**, 686 (2005).
- [16] Y. Iwasaki and T. Yoshi, *Phys. Lett.* **143B**, 449 (1984).
- [17] P. M. Vranas, *Phys. Rev. D* **74**, 034512 (2006).
- [18] D. Renfrew, T. Blum, N. Christ, R. Mawhinney, and P. Vranas, *Proc. Sci.*, LATTICE2008 (2008) 048.
- [19] D. B. Kaplan, *Phys. Lett. B* **288**, 342 (1992).
- [20] Y. Shamir, *Nucl. Phys.* **B406**, 90 (1993).
- [21] R. Brower, H. Neff, and K. Orginos, *Nucl. Phys. B, Proc. Suppl.* **140**, 686 (2005).
- [22] R. Brower, H. Neff, and K. Orginos, *Nucl. Phys. B, Proc. Suppl.* **153**, 191 (2006).
- [23] R. C. Brower, H. Neff, and K. Orginos, *Comput. Phys. Commun.* **220**, 1 (2017).
- [24] C. Allton *et al.* (RBC, UKQCD Collaborations), *Phys. Rev. D* **76**, 014504 (2007).
- [25] P. A. Boyle *et al.*, *Phys. Rev. D* **93**, 054502 (2016).
- [26] R. Mawhinney, *Proc. Sci.*, LATTICE2016 (2016) 136.
- [27] R. Arthur *et al.* (RBC, UKQCD Collaborations), *Phys. Rev. D* **87**, 094514 (2013).
- [28] Z. Bai, T. Blum, P. A. Boyle, N. H. Christ, J. Frison, N. Garron, T. Izubuchi, C. Jung, C. Kelly, C. Lehner, R. D. Mawhinney, C. T. Sachrajda, A. Soni, and D. Zhang (RBC and UKQCD Collaborations), *Phys. Rev. Lett.* **115**, 212001 (2015).
- [29] H. J. Rothe, *Lattice Gauge Theories: An Introduction*, World Scientific Lecture Notes in Physics (World Scientific, Singapore, 2012).
- [30] M. A. Clark, *Proc. Sci.*, LAT2006 (2006) 004, [arXiv:hep-lat/0610048].
- [31] Y. Aoki *et al.* (RBC, UKQCD Collaborations), *Phys. Rev. D* **83**, 074508 (2011).
- [32] P. de Forcrand, M. G. Prez, and I.-O. Stamatescu, *Nucl. Phys.* **B499**, 409 (1997).
- [33] M. Albanese *et al.*, *Phys. Lett. B* **192**, 163 (1987).
- [34] T. Blum, T. Izubuchi, and E. Shintani, *Phys. Rev. D* **88**, 094503 (2013).
- [35] We focus on the 16I-G ensemble simply because significantly more thermalized molecular dynamics trajectories were produced for this ensemble than for the 16I or 16ID-G ensembles.
- [36] P. A. Boyle, *Comput. Phys. Commun.* **180**, 2739 (2009).
- [37] R. C. Brower, T. Ivanenko, A. R. Levi, and K. N. Orginos, *Nucl. Phys.* **B484**, 353 (1997).
- [38] C. Kelly, *Proc. Sci.*, LAT2016 (2016) 308.
- [39] P. A. Boyle *et al.* (RBC/UKQCD Collaborations), *J. High Energy Phys.* **06** (2015) 164.
- [40] R. D. Mawhinney and D. J. Murphy (RBC-UKQCD Collaborations), *Proc. Sci.*, LATTICE2015 (2016) 062.
- [41] C. Jung (RBC and UKQCD Collaborations), *Proc. Sci.*, LATTICE2013 (2014) 417.
- [42] P. Boyle, A. Yamaguchi, G. Cossu, and A. Portelli, arXiv:1512.03487.
- [43] B. N. Datta, *Numerical Linear Algebra and Applications*, 2nd ed. (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2010).