CERN-PH-TH/2012-161

Lattice QCD with open boundary conditions and twisted-mass reweighting

Martin Lüscher and Stefan Schaefer

CERN, Physics Department, 1211 Geneva 23, Switzerland

Abstract

Lattice QCD simulations at small lattice spacings and quark masses close to their physical values are technically challenging. In particular, the simulations can get trapped in the topological charge sectors of field space or may run into instabilities triggered by accidental near-zero modes of the lattice Dirac operator. As already noted in ref. [1], the first problem is bypassed if open boundary conditions are imposed in the time direction, while the second can potentially be overcome through twisted-mass determinant reweighting [2]. In this paper, we show that twisted-mass reweighting works out as expected in QCD with open boundary conditions and 2+1 flavours of O(a) improved Wilson quarks. Further algorithmic improvements are tested as well and a few physical quantities are computed for illustration.

1. Introduction

To be able to control the systematic errors in lattice QCD computations, simulations of lattices with spacing smaller than 0.05 fm and spatial extent of at least 4 fm have to be performed. Moreover, the quark masses should ideally be set to their physical values in these simulations.

An obstacle to progress along these lines is the well-established fact that all known simulation algorithms tend to get trapped in the sectors of gauge fields with fixed topological charge [3–7]. So far no remedy against this loss of ergodicity was found, but the problem can be bypassed by choosing open boundary conditions for the gauge field in the time direction [1]. The topological charge can then flow in and out of the lattice through its boundaries, while the physical states and the Hamiltonian are unchanged.

In this paper, the formulation of lattice QCD proposed by Wilson [8] is adopted, with counterterms added to cancel the leading effects in the lattice spacing a [9,10]. This version of the lattice theory has many desirable properties and is relatively easy to simulate. Chiral symmetry is however violated by effects of order a^2 and the spectrum of the Dirac operator is therefore not protected against accidental near-zero modes. Such modes can give rise to instabilities in simulations based on the HMC algorithm [11], which may, in the worst case, compromise the correctness of the simulations.

At fixed gauge coupling and quark masses, near-zero modes tend to be less frequent the larger the lattice volume V is, because the width of the distribution of the lowest eigenvalue of the Dirac operator decreases approximately like $V^{-1/2}$ [12–17]. While the shrinking of the width of the eigenvalue distribution has a stabilizing effect on large lattices, the twisted-mass determinant reweighting proposed in ref. [2] avoids the problem from the beginning through an intermediate infrared regularization of the quark determinant.

Encouraging first tests of this method were recently reported by Miao et al. [18]. Here we shall present the results of a more complete study that includes simulations of QCD with 2+1 flavours of quarks at a point in parameter space previously considered by the PACS-CS collaboration [19,20], where the quark masses are practically equal to their physical values. The simulations we have performed for these tests are also the first ones of full QCD with open boundary conditions. Moreover, an effort was made to improve the efficiency and robustness of the simulation algorithm by combining various known techniques (see sections 3 and 4). All simulations reported in this paper were performed using the publicly available `openQCD` program package [21].

2. Twisted-mass determinant reweighting

2.1 Lattice theory

As already mentioned, we consider lattice QCD with O(a)-improved Wilson quarks. The up and down quarks are assumed to be mass-degenerate and are referred to as the light quarks. There could be any number of heavier quarks, but the strange quark is the only one included in the simulations reported later.

The basic setup of the lattice theory and the notation employed are as in ref. [1]. In particular, open boundary conditions are imposed in the time direction and the

(four-dimensional hypercubic) lattice is assumed to have time-like extent T and spatial size L. For notational convenience, the lattice spacing is set to unity. As for the gauge action, we slightly depart from ref. [1] and replace the Wilson action by a more general expression, which includes the tree-level Symanzik-improved and the Iwasaki action (see appendix A).

2.2 Determinant regularization

Let D be the up quark lattice Dirac operator as it appears in the lattice action. The operator thus includes the (ordinary) mass term and the required $\mathrm{O}(a)$ counterterms. In ref. [2], two kinds of twisted-mass regularizations of the light quark determinant were proposed, which amount to replacing

$$\det\{D^\dagger D\} \to \det\{D^\dagger D + \mu^2\} \tag{2.1}$$

or

$$\det\{D^\dagger D\} \to \det\{(D^\dagger D + \mu^2)^2 (D^\dagger D + 2\mu^2)^{-1}\} \tag{2.2}$$

respectively. The twisted mass parameter $\mu > 0$ provides the desired infrared regularization and is usually set to a value on the order of the light quark mass.

With the regularization in place, the ensembles of gauge-field configurations generated in the simulations must be reweighted in order to obtain the correct expectation values of the observables of interest. The reweighting factors in the two cases are

$$W_1 = \det\{D^\dagger D (D^\dagger D + \mu^2)^{-1}\}, \tag{2.3}$$

$$W_2 = \det\{D^\dagger D (D^\dagger D + 2\mu^2)(D^\dagger D + \mu^2)^{-2}\}. \tag{2.4}$$

Both factors are ratios of quark determinants, which can be estimated stochastically with a modest computational effort (see subsect. 2.4).

Determinant reweighting usually becomes inefficient on large lattices, but as explained in ref. [2], the reweighting factors W_1 and W_2 are not expected to fluctuate wildly if μ is chosen appropriately. The second form, eq. (2.2), of the determinant regularization potentially fares better in this respect, because the contribution of the (very many) high modes of the Dirac operator to the reweighting factor is more strongly suppressed than in the case of the first form.

In our empirical studies, we found that W_2 in fact tends to fluctuate less than W_1, although the behaviour of the two factors is not qualitatively different and moreover

depends on the value of μ. For simplicity, and since this is the method used in the test runs reported later, we shall from now on focus on the regularization (2.2) of the quark determinant.

2.3 Even-odd preconditioned version

Twisted-mass determinant reweighting easily combines with the even-odd preconditioning of the Dirac operator. Let

$$D = \begin{pmatrix} D_{\mathrm{ee}} & D_{\mathrm{eo}} \\ D_{\mathrm{oe}} & D_{\mathrm{oo}} \end{pmatrix} \tag{2.5}$$

be the block decomposition of the Dirac operator with respect to an ordering of the lattice points x, where the even ones (those with even $x_0 + x_1 + x_2 + x_3$) come first. In practice the blocks on the diagonal are always invertible so that the even-odd preconditioned operator,

$$\hat{D} = D_{\mathrm{ee}} - D_{\mathrm{eo}}(D_{\mathrm{oo}})^{-1} D_{\mathrm{oe}}, \tag{2.6}$$

is well defined. Note that $\hat{D}$ acts on quark fields defined on the even lattice points.

When even-odd preconditioning is used, eqs. (2.2) and (2.4) get replaced by

$$\det\{D^\dagger D\} \to \det\{(D_{\mathrm{oo}})^2\} \det\{(\hat{D}^\dagger \hat{D} + \mu^2)^2 (\hat{D}^\dagger \hat{D} + 2\mu^2)^{-1}\}, \tag{2.7}$$

$$\hat{W}_2 = \det\{\hat{D}^\dagger \hat{D} (\hat{D}^\dagger \hat{D} + 2\mu^2)(\hat{D}^\dagger \hat{D} + \mu^2)^{-2}\}. \tag{2.8}$$

Note that the twisted mass term is added on the even sites of the lattice only. The regularizations (2.2) and (2.7) are therefore not the same.

2.4 Computation of the reweighting factor

An unbiased stochastic estimator for the reweighting factor W_2 is given by

$$W_{2,N} = \frac{1}{N} \sum_{k=1}^{N} \exp\{-\mu^4 (\eta_k, (D^\dagger D)^{-1} (D^\dagger D + 2\mu^2)^{-1} \eta_k)\}, \tag{2.9}$$

where $\eta_1, \ldots, \eta_N$ are random quark fields with normal distribution and the bracket $(\cdot\,,\cdot)$ denotes the obvious scalar product of such fields. The reweighting factor $\hat{W}_2$ can be similarly estimated by replacing D by $\hat{D}$ and by restricting the random fields to the even sites of the lattice.

In practice, the number N of random fields is usually taken to be in the range from, say, 12 to 48. The deviation $|W_{2,N} - W_2|$ is then often smaller than the statistical fluctuations of W_2, in which case the stochastic estimation of the reweighting factor does not lead to enhanced statistical errors. More accurate determinations of the reweighting factor (along the lines of ref. [22], for example) may however be required if observables sensitive to the low modes of the Dirac operator are considered and if there is an appreciable probability for the operator to have exceptionally small eigenvalues.

3. Frequency splitting of the quark determinant

The numerical integration of the molecular-dynamics equations can be a source of instability in the HMC algorithm even at relatively large quark masses. Empirically one knows that a frequency splitting of the quark determinant through mass shifts [23–25] or a domain decomposition of the Dirac operator [26] has a stabilizing effect and thus allows the equations to be integrated with larger step sizes than would otherwise be possible.

In this section, we describe a particular frequency-splitting scheme that naturally goes together with the twisted-mass determinant reweighting. The scheme has further merits and performed very well in all simulations reported later. For simplicity, we only discuss the case without even-odd preconditioning, but all results extend to the preconditioned quark determinants with the obvious modifications.

3.1 Factorization of the light-quark determinant

Let $\mu_0, \ldots, \mu_n$ be a set of twisted mass parameters satisfying

$$\mu_0 = \mu, \quad \mu_0 < \mu_1 < \ldots < \mu_n, \tag{3.1}$$

where μ is the regulator mass used for the determinant reweighting. The regularized light-quark determinant on the right of eq. (2.2) may then be written in the factorized form [23,24]

$$\det\left\{D^\dagger D + \mu_n^2\right\} \det\left\{\frac{D^\dagger D + \mu_0^2}{D^\dagger D + 2\mu_0^2}\right\} \prod_{k=0}^{n-1} \det\left\{\frac{D^\dagger D + \mu_k^2}{D^\dagger D + \mu_{k+1}^2}\right\}. \tag{3.2}$$

Each factor in this product of determinants may be represented through a pseudo-fermion functional integral. In total $n+2$ independent pseudo-fermion fields, $\tilde{\phi}_0$ and $\phi_0, \phi_1, \ldots, \phi_n$, with actions

$$\tilde{S}_{\mathrm{pf},0} = \big(\tilde{\phi}_0, (D^\dagger D + 2\mu_0^2)(D^\dagger D + \mu_0^2)^{-1}\tilde{\phi}_0\big), \tag{3.3}$$

$$S_{\mathrm{pf},k} = \big(\phi_k, (D^\dagger D + \mu_{k+1}^2)(D^\dagger D + \mu_k^2)^{-1}\phi_k\big), \quad k = 0, \ldots, n-1, \tag{3.4}$$

$$S_{\mathrm{pf},n} = \big(\phi_n, (D^\dagger D + \mu_n^2)^{-1}\phi_n\big), \tag{3.5}$$

need to be introduced for this representation.

When written as products over the eigenvalues of $D^\dagger D$, the determinants in eq. (3.2) are seen to be dominantly dependent on the eigenvalues in spectral intervals that are roughly delimited by the twisted masses $\mu_0, \ldots, \mu_n$. To some extent, at least, the factorization thus achieves a frequency splitting of the light-quark determinant. While such factorizations are known to stabilize the HMC algorithm, there is currently no solid theoretical understanding of why this is so. As a consequence, the choice of the twisted masses is, in practice, poorly guided and may require some fine-tuning [25].

In the course of our algorithmic studies, we found that frequency splittings where $\mu_n \simeq 1$ and

$$\mu_k \simeq 0.1 \times \mu_{k+1}, \quad k = 0, \ldots, n-1, \tag{3.6}$$

gave good results in all cases considered. Moreover, it is our experience that a fine-tuning of the masses is then not required. The rule (3.6) amounts to splitting the spectral range of the Dirac operator in equal segments on a log scale and is therefore referred to as "log-scale frequency splitting". Note that the rule implicitly fixes the number of twisted masses as a function of the reweighting mass $\mu = \mu_0$.

3.2 Strange quark determinant

While the physical strange quark is much heavier than the light quarks, the condition number of the strange-quark lattice Dirac operator D_s is not small in practice and a frequency splitting of strange-quark determinant thus seems advisable. Such splittings are naturally obtained when the RHMC algorithm [27,28] is employed for the strange quark.

The version of the RHMC algorithm used here essentially coincides with the one

described in sect. 2.6 of ref. [29]. The starting point is the factorization

$$\det D_s = W_s \det R^{-1} \tag{3.7}$$

of the strange-quark determinant, where

$$R = C \prod_{k=0}^{m-1} \frac{D_s^\dagger D_s + \omega_k^2}{D_s^\dagger D_s + \nu_k^2} \tag{3.8}$$

denotes the Zolotarev optimal rational approximation [30] of degree m of the operator $(D_s^\dagger D_s)^{-1/2}$ and

$$W_s = \det(D_s R) \tag{3.9}$$

the reweighting factor needed to correct for the approximation error. For a specified degree m and spectral range of $D_s^\dagger D_s$ in which the approximation R is to have the least possible error, the constant C and the twisted masses

$$\nu_0 < \omega_0 < \nu_1 < \ldots < \omega_{m-1} \tag{3.10}$$

are uniquely determined. The latter typically range over a few orders of magnitude and are about equally spaced on a log scale.

A further factorization of the strange-quark determinant is now obtained by breaking up the Zolotarev rational function (3.8) into two or more factors of the form

$$R_{l,j} = \prod_{k=l}^{j} \frac{D_s^\dagger D_s + \omega_k^2}{D_s^\dagger D_s + \nu_k^2}. \tag{3.11}$$

If $m = 12$, for example, a possible factorization is

$$\det R^{-1} = \text{constant} \times \det\{R_{0,4}^{-1}\} \det\{R_{5,8}^{-1}\} \det\{R_{9,11}^{-1}\}. \tag{3.12}$$

Each factor $\det\{R_{l,j}^{-1}\}$ is then represented through an integral over a pseudo-fermion field $\phi_{l,j}$ with action

$$S_{\mathrm{pf},l,j} = (\phi_{l,j}, R_{l,j}\phi_{l,j}). \tag{3.13}$$

In view of the strong ordering of the twisted masses ω_k and ν_k, a frequency splitting of the strange-quark determinant is achieved in this way, very much akin to the one of the light-quark determinant discussed in subsect. 3.1.

3.3 Molecular-dynamics forces

The molecular-dynamics force fields that derive from the actions $S_{\mathrm{pf},k}$ are given by

$$F_k(x,\mu)^a = -2(\mu_{k+1}^2 - \mu_k^2)\,\mathrm{Re}\,(\chi_k, \gamma_5 \partial_{x,\mu}^a D\psi_k), \quad k = 0,\ldots,n-1, \tag{3.14}$$

$$F_n(x,\mu)^a = -2\,\mathrm{Re}\,(\chi_n, \gamma_5 \partial_{x,\mu}^a D\psi_n), \tag{3.15}$$

where $\partial_{x,\mu}^a$ denotes the partial derivative with respect to the link variable $U(x,\mu)$ in direction of the SU(3) generator T^a and

$$\psi_k = (D + i\mu_k\gamma_5)^{-1}\gamma_5\phi_k, \qquad \chi_k = (D - i\mu_k\gamma_5)^{-1}\gamma_5\psi_k. \tag{3.16}$$

On physically large lattices, the force fields tend to be strongly ordered in magnitude. In particular, in the case of log-scale frequency splitting, the force F_k is about 10 times smaller than F_{k+1} and the force $\tilde{F}_0$ deriving from the action $\tilde{S}_{\mathrm{pf},0}$ is smaller than F_0 by roughly another order of magnitude. When integrating the molecular-dynamics equations, the pseudo-fermion forces can thus be integrated with different integration step sizes without compromising the accuracy of the integration [31].

Essentially the same comments apply to the forces deriving from the strange-quark pseudo-fermion actions (3.13). When the rational functions are expanded in partial fractions, the contributions of the fractions to the force are actually again given by eqs. (3.14),(3.16) except for a change in the proportionality factor and the fact that the Dirac operator D is replaced by D_s.

4. Integration of the molecular-dynamics equations

The numerical integration of the molecular-dynamics equations can be accelerated by using the improved integrators proposed by Omelyan, Mryglod and Folk [32] and a locally deflated solver for the lattice Dirac equation [33,34]. In the following subsections, we briefly describe the implementation of these improvements in our simulations.

4.1 Evolution equations

The molecular-dynamics equations

$$\partial_t \pi(x,\mu) = -T^a \partial_{x,\mu}^a S(U), \tag{4.1}$$

$$\partial_t U(x,\mu) = \pi(x,\mu)U(x,\mu), \tag{4.2}$$

evolve the gauge field $U(x,\mu)$ and its momentum $\pi(x,\mu) = \pi(x,\mu)^a T^a$ as a function of the molecular-dynamics time t. If integrated exactly, the evolution preserves the Hamilton function

$$H(\pi,U) = \tfrac{1}{2}(\pi,\pi) + S(U), \qquad (\pi,\pi) = \sum_{x,\mu} \pi(x,\mu)^a \pi(x,\mu)^a. \tag{4.3}$$

In these formulae, $S(U)$ stands for the total action, i.e. the sum of the gauge action, the pseudo-fermion actions and the terms proportional to $\ln\{D_{\rm oo}\}$ and $\ln\{(D_s)_{\rm oo}\}$, which need to be included if even-odd preconditioning is used (cf. subsect. 2.3; the dependence of the action on the pseudo-fermion fields is suppressed for simplicity).

4.2 Elementary integrators

The numerical integration schemes used in our simulations are based on the leapfrog integrator, the 2nd order Omelyan-Mryglod-Folk (OMF) integrator and a particular 4th order OMF scheme. All these integrators are sequences of the elementary update steps

$$\mathcal{I}_\pi(\epsilon):\ \pi \to \pi - \epsilon F, \tag{4.4}$$

$$\mathcal{I}_U(\epsilon):\ U \to \mathrm{e}^{\epsilon\pi} U, \tag{4.5}$$

where ϵ denotes the time step size and $F(x,\mu) = F(x,\mu)^a T^a$ the molecular-dynamics force integrated in the step. The leapfrog integrator, for example, amounts to applying the combination

$$\mathrm{LPFR}(\epsilon) = \mathcal{I}_\pi(\tfrac{1}{2}\epsilon)\mathcal{I}_U(\epsilon)\mathcal{I}_\pi(\tfrac{1}{2}\epsilon) \tag{4.6}$$

to the fields, while the 2nd order OMF integrator,

$$\mathrm{OMF}_2(\epsilon) = \mathcal{I}_\pi(\lambda\epsilon)\mathcal{I}_U(\tfrac{1}{2}\epsilon)\mathcal{I}_\pi((1-2\lambda)\epsilon)\mathcal{I}_U(\tfrac{1}{2}\epsilon)\mathcal{I}_\pi(\lambda\epsilon) \tag{4.7}$$

updates the gauge field in two steps and depends on a tunable parameter λ. In the case of the 4th order OMF integrator, $\mathrm{OMF}_4(\epsilon)$, there are 5 update steps, with step sizes given by eqs. (63) and (71) in ref. [32], and no tunable parameters.

4.3 Hierarchical integration

If the force F is set to the one deriving from the total action, the n-fold application of one of the elementary integrators integrates the molecular-dynamics equations from time 0 to some later time $\tau = n\epsilon$. In practice, a hierarchical integration scheme is used, where the different forces are integrated with different time step sizes [31].

Hierarchical integrators are constructed recursively in a way that is best explained by considering a simple case. Suppose the total action $S = S_0 + S_1$ is a sum of two terms that give rise to the forces F_0 and F_1. The construction then starts from an integrator of the form

$$\mathcal{J}_1(\tau) = \left\{ \mathrm{OMF}_2(\epsilon_1)|_{F\to F_1} \right\}^{n_1}, \qquad \epsilon_1 = \tau/n_1, \tag{4.8}$$

for the molecular-dynamics equations in which the total action is replaced by S_1. For a given integration time τ, one has the choice of the elementary integrator (OMF_2 in this case) at this level and of the number n_1 of times the latter is applied. Note that $\mathcal{J}_1(\tau)$ is just a sequence of update steps $\mathcal{I}_\pi(\epsilon)$ and $\mathcal{I}_U(\epsilon)$ with varying step sizes ϵ proportional to ϵ_1.

The force F_0 may now be included in the molecular-dynamics evolution by replacing all instances $\mathcal{I}_U(\epsilon)$ of the gauge-field update steps by an integrator like

$$\mathcal{J}_0(\epsilon) = \left\{ \mathrm{OMF}_4(\epsilon_0)|_{F\to F_0} \right\}^{n_0}, \qquad \epsilon_0 = \epsilon/n_0, \tag{4.9}$$

which integrates F_0 from the current time t to $t + \epsilon$. At this level, one again has the choice of the elementary integrator and the number of times it is applied. The integrator obtained in this way integrates F_0 and F_1 with average step sizes equal to $\tau/(10n_0n_1)$ and $\tau/(2n_1)$, respectively.

When the total action is a sum of n terms, a hierarchical integrator with n levels is required if the associated forces are to be integrated with different step sizes. The construction of the integrator always starts from the top level, where the smallest forces are integrated, and proceeds to the lower levels recursively by replacing the gauge-field update steps by a power of an elementary integrator. For a complete description of the integration scheme, the integration time τ, the list of elementary integrators, the numbers $n_0, n_1, \ldots$ and the forces integrated at each level must be specified.

4.4 Deflation acceleration

The frequency splitting of the quark determinant tends to give rise to a fairly large number of pseudo-fermion forces in the molecular-dynamics equations. When some

or all of these forces need to be computed, the twisted-mass lattice Dirac equation must be solved several times.

There is a range of algorithms that allow the lattice Dirac equation to be solved efficiently. In particular, the forces deriving from the rational-function actions (3.13) can be computed using a multi-shift conjugate-gradient (CG) algorithm [35]. The CG algorithm is also suitable for the solution of the Dirac equation at large quark masses, but at small and intermediate masses, the equation can be solved much more rapidly using the GCR solver, local deflation [33] and a preconditioner based on the Schwarz alternating procedure (SAP) [36].

Local deflation requires a deflation subspace to be generated before the integration of the molecular-dynamics equations starts and to be kept up-to-date in the course of the integration [34]. This overhead is however rapidly amortized along the molecular-dynamics trajectories if there are several pseudo-fermion forces, where the use of the deflated solver is highly profitable (as is the case at small quark masses).

4.5 Remark on solver tolerances

Whichever iterative procedure is used for the solution of the lattice Dirac equation, the algorithm is stopped as soon as the residue of the calculated approximate solution has decreased by some factor δ. It is tempting to relax the tolerance δ as one proceeds from the larger to the smaller forces, since the latter are small corrections to the total force and therefore need not be computed as accurately as the large forces [28].

This argumentation however ignores the fact that the deviations of the computed from the exact solutions depend on the condition number of the lattice Dirac operator. In the case of the fields (3.16), for example, simple norm estimates actually suggest that the relative error of the calculated fields ψ_k and χ_k scale proportionally to δ/μ_k and δ/μ_k^2, respectively. Such rigorous estimates tend to be too pessimistic, but they show that a loosening of the solver tolerances risks to compromise the accuracy (and thus the stability) of the numerical integration of the molecular-dynamics equations.

5. Algorithm stability and performance

One of the principal goals in this paper is to find out whether twisted-mass determinant reweighting works out on large lattices and at quark masses close to their physical values. The simulations reported below serve to study this question, but also provide a test of the simulation algorithm described in sections 3 and 4.

Table 1. Lattice parameters

Run	Lattice	S_{G}	β	c_{sw}	κ_u	κ_s
D_6	48×24^3	Wilson	5.3	1.90952	0.136350	–
E_8	64×32^3	Wilson	5.3	1.90952	0.136417	–
I_1	64×32^3	Iwasaki	1.9	1.71500	0.137740	0.136600
I_2	64×32^3	Iwasaki	1.9	1.71500	0.137796	0.136634

Table 2. Lattice spacing and meson masses

Run	a [fm]	L [fm]	m_π [MeV]	m_K [MeV]	$m_\pi L$	Reference
D_6	0.066	1.6	311	–	2.5	[39]
E_8	0.066	2.1	191	–	2.0	[39]
I_1	0.090	2.9	215	524	3.1	[19,20]
I_2	0.090	2.9	135	498	2.0	[20]

5.1 Lattice parameters

The parameters of the lattice theories we have simulated are listed in table 1. In the first two runs, D_6 and E_8, only the light quarks, with hopping parameter κ_u, were included, while the other two runs are simulations of 2+1 flavour QCD with strange-quark hopping parameter κ_s. The quoted values of the O(a) improvement coefficient c_{sw} were determined non-perturbatively in refs. [37] and [38], respectively, and the boundary improvement coefficients c_{G} and c_{F} were set to unity.

The basic physical parameters of the simulated lattices can be inferred from simulation results obtained in refs. [19,20,39] at nearby points in parameter space (see table 2). We quote these figures without errors and solely with the intention of giving a rough impression of the physical situation on the lattices we have considered. Table 2 shows that all lattices are at the edge of the large volume regime of QCD. From the point of view of the simulation stability, such lattices are particularly challenging and we therefore expect that stability can more easily be achieved when one proceeds to simulating larger and finer lattices (cf. sect. 1).

5.2 Algorithm parameters

In all runs reported here, even-odd preconditioning, the second kind of twisted-mass determinant reweighting and a twisted-mass determinant factorization with

Table 3. Simulation parameters

Run	τ	$\mu_0, \ldots, \mu_n$	Top level integrator*	$P_{\rm acc}$	$N_{\rm tr}$
D_6	2.0	$0.0045, 0.01, 0.1, 1.0$	$\{\mathrm{LPFR}\}^{10}$	0.94	1800
E_8	1.8	$0.0015, 0.01, 0.1, 1.0$	$\{\mathrm{OMF}_2\}^5$	0.93	896
I_1	1.2	$0.0020, 0.05, 0.5$	$\{\mathrm{OMF}_2\}^4$	0.88	1224
I_2	1.1	$0.0012, 0.05, 0.5$	$\{\mathrm{OMF}_2\}^6$	0.90	400

*The parameter λ of the OMF_2 integrator was set to 1/6 [32].

three or four masses μ_k were used (see table 3). The molecular-dynamics trajectory length τ was chosen to be about 2 on the finer lattices and 1 on the coarser ones. Experience suggests that shorter trajectory lengths would have a negative impact on the autocorrelation times of physical quantities [5]. The particular values of τ in table 3 were chosen so as to get high acceptance rates $P_{\rm acc}$ for the fields produced by the molecular-dynamics evolution. In the last column of table 3, the numbers $N_{\rm tr}$ of trajectories generated after thermalization are listed.

The choice of the rational approximation of the strange-quark determinant in the 2+1 flavour runs required some experimenting since the spectral range of $(\hat{D}_s^\dagger \hat{D}_s)^{1/2}$ is not known a priori. Eventually we settled on a Zolotarev rational function with 9 poles, a spectral approximation range $[0.03, 6.1]$ and the factorization

$$\det R^{-1} = \text{constant} \times \det\{R_{0,0}^{-1}\} \det\{R_{1,1}^{-1}\} \det\{R_{2,2}^{-1}\} \det\{R_{3,8}^{-1}\} \tag{5.1}$$

of the approximate strange-quark determinant (cf. subsect. 3.2). The approximation error is sufficiently small in this case to suppress the fluctuations of the reweighting factor $\hat{W}_s$ to a level of a few percent.

For the integration of the molecular-dynamics equations, a hierarchical integrator with 3 levels was used in all cases, $\{\mathrm{OMF}_4\}^1$ being the integrator on the first as well as on the second level. The top level integrators are listed in table 3. Only the force deriving from the gauge action is integrated at the lowest level and only the smallest forces (the ones deriving from $\tilde{S}_{\rm pf,0}$, $S_{\rm pf,0,0}$ and $S_{\rm pf,1,1}$) at the top level. Most forces are thus integrated at the first level.

For the solution of the lattice Dirac equation, the locally deflated SAP preconditioned GCR algorithm [33,36] was employed except at the largest twisted masses, where we used the CG and multi-shift CG [35] algorithms. The relative residues of the calculated solutions were required to be less than 10^{-10} in the force computations and at most 10^{-11} in the computation of the pseudo-fermion actions. With these

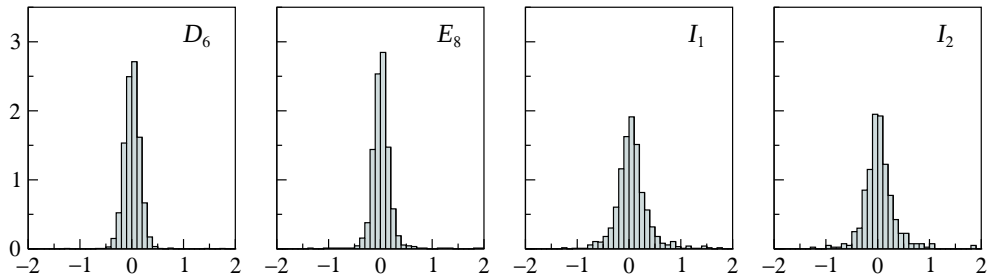

Fig. 1. Normalized distribution of the energy deficit ΔH at the end of the molecular-dynamics trajectories as measured in the runs $D_6, \ldots, I_2$.

tolerances, the reversibility of the numerical integration of the molecular-dynamics equations is guaranteed to a precision of about 10^{-9} in the link variables.

5.3 Integration instabilities

The molecular-dynamics evolution is probably chaotic and is in any case well known to be sensitive to integration inaccuracies. A manifestation of integration instabilities are large energy deficits ΔH at the end of the molecular-dynamics trajectories. These can be caused by accidental near-zero modes of the light-quark Dirac operator, but there exist further sources of instability as well (loose solver tolerances, for example, or coherent effects of the higher modes).

Twisted-mass determinant reweighting eliminates the first source of instability and is therefore expected to have a positive effect on the stability of the simulations. The energy deficits observed in our test runs are indeed well behaved (see fig. 1). In all cases, values of $|\Delta H|$ significantly larger than 2 are rare and occur with an estimated probability of at most a few permille. While such a high level of stability would be difficult to achieve without low-mode regularization of the quark determinant, the log-scale frequency splitting of the determinant no doubt has a stabilizing effect as well and perhaps also our choice of the molecular-dynamics integrator.

5.4 Reweighting efficiency

Stochastic estimates of the light- and strange-quark reweighting factors $\hat{W}_2$ and $\hat{W}_s$ can be obtained following the lines of subsect. 2.4. In all runs $D_6, \ldots, I_2$ we used 48 random fields for the estimation of $\hat{W}_2$ and a single field in the case of the strange-quark reweighting factor. The latter is actually nearly constant and little would be gained by calculating it more accurately.

For the reweighting to work out, the normalized reweighting factor should remain smaller than 2 or so, as otherwise the ensemble of fields generated in the simulation

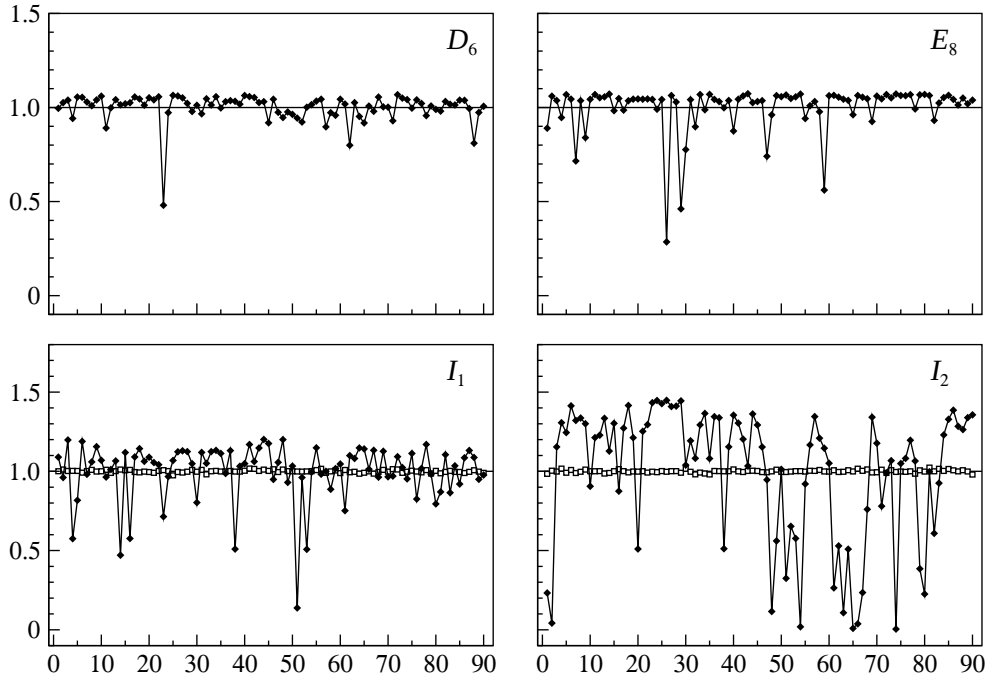

Fig. 2. Stochastic estimates of the reweighting factors $\hat{W}_2$ and $\hat{W}_2\hat{W}_s$ (diamonds; connecting lines are drawn to guide the eye) in the runs D_6, E_8 and I_1, I_2, respectively, plotted as a function of the configuration number. Configurations are separated by 8 trajectories except in run I_2, where they are separated by 4 trajectories. In the case of the 2+1 flavour runs, the strange-quark reweighting factor $\hat{W}_s$ is shown too (squares). All reweighting factors plotted in this figure are normalized such that their average is equal to 1.

is effectively reduced to the subset of configurations with the dominant weights. This condition is easily met in all simulations reported here (see fig. 2). Even in the most critical case, run I_2, the reweighting factor stays below 1.5 and only 15% of the gauge-field configurations have weight less than 0.5. On the simulated lattices and with the chosen values of the regulator mass μ_0, the efficiency of the simulations is thus not compromised by the determinant reweighting.

5.5 Low mode sampling

From the point of view of the sampling efficiency, observables that are sensitive to the low modes of the light-quark Dirac operator are a special case, because the series of measured values of such quantities tend to have large "spikes" at the points in simulation time where the Dirac operator happens to have near-zero modes. If the

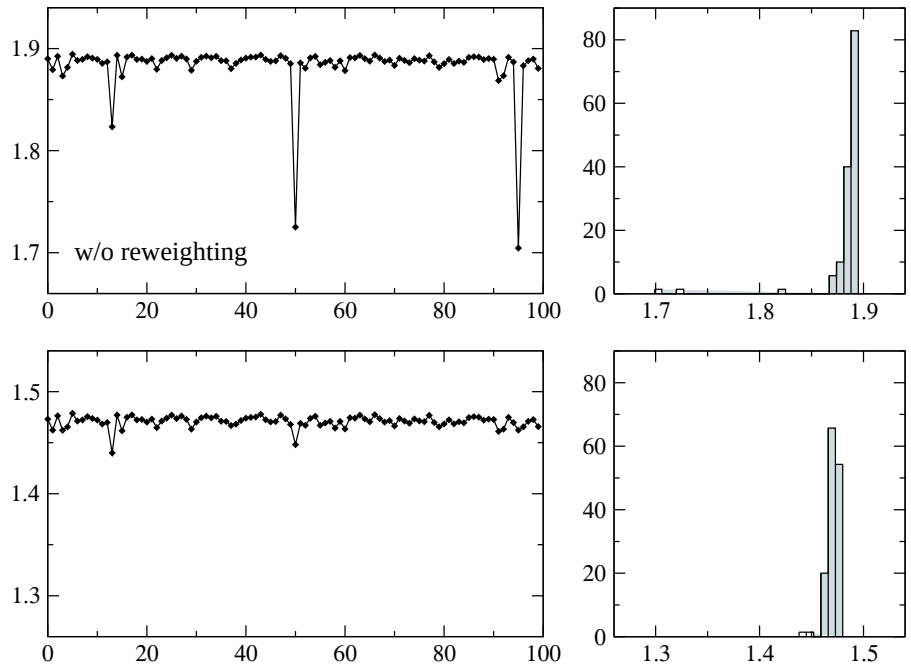

Fig. 3. Jackknife samples of the pion propagator $G_\pi(25,1)$, in lattice units and scaled by 10^3, as measured on the first 100 gauge-field configurations generated in run I_1 (configurations are separated by 8 trajectories). On the left the jackknife samples calculated with and without reweighting are plotted versus the omitted-configuration number. The normalized distributions of the samples are shown on the right.

data series is dominated by a few spikes, a reliable estimation of the expectation value of the observable and the associated statistical error is then practically excluded.

Exceptionally low eigenvalues of the Dirac operator can occur in simulations with twisted-mass reweighting too, but the spikes in the measurement data series now get suppressed by the reweighting factor. The product of the reweighting factor and the sum of the quark-line diagrams representing a hadronic correlation function in fact remains bounded when the Dirac operator becomes singular.

For illustration, consider the pion propagator

$$G_\pi(x_0, y_0) = -\sum_{\vec{x}} \langle (\bar{u}\gamma_5 d)(x)(\bar{d}\gamma_5 u)(y)\rangle \tag{5.2}$$

and its computation in run I_1 at $x_0 = 25$ and $y_0 = 1$ (see fig. 3). In order to reduce the statistical fluctuations, the propagator was evaluated using 10 random source fields at time y_0. As can be inferred from the series of the jackknife samples plotted

in the upper left diagram in fig. 3, the series of measured values has a few spikes in this run, which are up to 10 times larger than the median of the data.

Once the data are reweighted, the spikes however disappear, as theoretically anticipated, and the distribution of the jackknife samples is then entirely well behaved (lower row of diagrams in fig. 3). Note that the normalized reweighting factor assumes values much smaller than 1 more often than there are spikes in the propagator data series, because the propagator is sensitive to only those low modes of the Dirac operator whose eigenfunctions have significant support at time x_0 and y_0.

5.6 Simulation cost

Twisted-mass determinant reweighting adds very little to the total computational effort required for the simulations. Moreover, whether the first or the second form of twisted-mass reweighting is used makes nearly no difference, because the additional force term that must be included in the molecular dynamics evolution in the second case is tiny and can therefore be integrated with a large step size.

While the computer time required for the simulations depends on many parameters, the execution times measured in the runs reported in this paper may be of some interest. Using the `openQCD` program [21] and 12 nodes (96 cores) of a standard PC cluster†, the computer time required per unit of molecular-dynamics time in the D_6 run was 0.26 hours. For the runs E_8, I_1 and I_2, we used 32 nodes (256 cores) of the same machine, the execution times in these cases being 0.61, 0.93 and 1.74 hours per unit of molecular-dynamics time.

The true cost of a simulation however also depends on the autocorrelation times. As discussed in refs. [1,7], realistic lower bounds on the exponential autocorrelation times can be obtained by considering observables constructed using the Wilson flow [6]. The autocorrelation times estimated in this way turned out to be about 32, 20 and 15 in units of molecular-dynamics time in the D_6, E_8 and I_1 runs, respectively. These values are almost a factor 2 smaller than those determined at similar lattice spacings in the SU(3) gauge theory with open boundary conditions [1], but one should keep in mind that the estimates quoted here are based on much shorter data series and may need to be corrected once longer runs are performed.

† The machine used for the tests has 84 dual processor nodes with AMD Opteron 2352 (2.1 GHz quad-core) processors, 8 GB of DDR2-667 memory and DDR Infiniband interconnects.

6. Computation of physical quantities

When open boundary conditions are imposed, the QCD Hamiltonian and the space of physical states remain unchanged, but the presence of the boundaries at time 0 and T potentially complicates the analysis of the calculated correlation functions. Two cases of interest illustrating the issue are discussed below. We focus on run I_1 in this section, since the physical situation on this lattice is the one nearest to being representative of the large-volume regime of QCD.

6.1 Reference flow time

Extrapolations to the continuum limit require the physics on several lattices to be accurately matched. The matching can be based on a comparison of pseudo-scalar meson masses and decay constants, for example, but as explained in ref. [6], the finiteness of the Wilson flow in the continuum limit [6,40] may allow the lattices to be matched far more easily.

A quantity of interest in this context is the reference flow time t_0 implicitly defined through [6]

$$\big\{t^2\langle E(x)\rangle\big\}_{t=t_0} = 0.3, \qquad E(x) = \tfrac{1}{4}G^a_{\mu\nu}(x)G^a_{\mu\nu}(x), \tag{6.1}$$

where $G^a_{\mu\nu}(x)$ is a lattice expression for the gauge-field tensor at flow time t (see appendix B for the definition of the Wilson flow). On lattices with periodic boundary conditions, the expectation value $\langle E(x)\rangle$ is independent of x and coincides with its infinite-volume limit up to terms that vanish exponentially when lattice sizes T and L are taken to infinity. Note, however, that the asymptotic approach to the infinite volume limit can only be expected to set in at lattice sizes significantly larger than the smoothing range $\sqrt{8t}$ of the Wilson flow.

In the case of open boundary conditions, translation invariance in time is broken and $\langle E(x)\rangle$ consequently depends on x_0. Similarly to the finite-volume corrections on periodic lattices, the effects of the boundaries at time 0 and T decrease exponentially when one moves away from the boundaries. On large lattices and at flow times where the smoothing range of the Wilson flow is much smaller than the lattice sizes, $\langle E(x)\rangle$ is thus expected to be practically constant and equal to its infinite-volume value in a central range of x_0.

Up to statistical fluctuations of $1-2\%$, the simulation results for $\langle E(x)\rangle$ shown in fig. 4 are indeed consistent with the existence of a plateau in a broad range of x_0. In this calculation, a symmetric (clover) expression was employed for the gauge

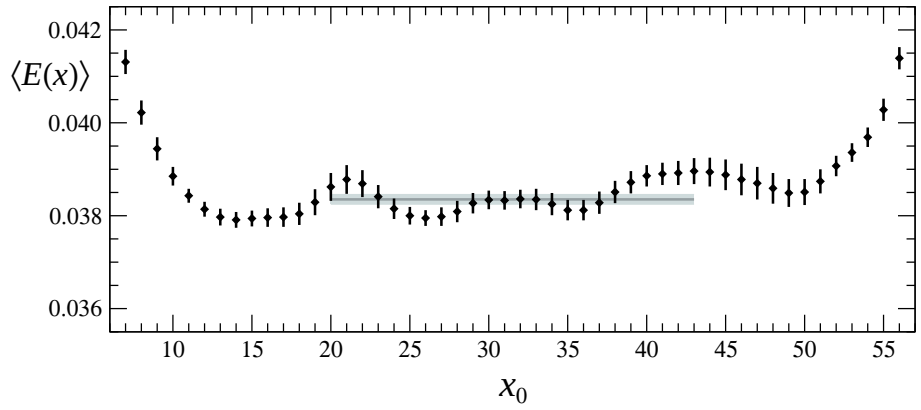

Fig. 4. Plot of $\langle E(x)\rangle$ at flow time $t = t_0$, calculated using 150 gauge-field configurations generated in run I_1. All quantities are given in lattice units. Statistical errors were estimated by the jackknife method after dividing the data into blocks of 6 consecutive measurements. The grey line with its error band was obtained through an uncorrelated least-squares fit of the data by a constant in the range $20 \leq x_0 \leq 43$.

field tensor $G^a_{\mu\nu}(x)$ and the expectation value of $E(x)$ was estimated by averaging

$$\overline{E}(x_0) = \frac{1}{4L^3} \sum_{\vec{x}} G^a_{\mu\nu}(x) G^a_{\mu\nu}(x) \tag{6.2}$$

over the gauge fields. The waves in the central region seen in fig. 4 can be explained by recalling that the Wilson flow suppresses the high-frequency modes of the gauge field. The calculated values of $\overline{E}(x_0)$ are therefore strongly correlated and thus tend to fluctuate coherently over distances in x_0 roughly equal to the smoothing range of the flow (which is about 5 lattice spacings in the case of fig. 4).

From the expectation values $\langle E(x)\rangle$ measured in the central region of the lattice, the result $t_0 = 2.792(10)$ is obtained for the reference flow time in run I_1. The associated smoothing range, $\sqrt{8t_0}$, is 4.7 lattice spacings and 0.43 fm in physical units. As a function of the flow time t, the behaviour of the dimensionless combination $t^2\langle E(x)\rangle$ in the range shown in fig. 5 is practically the same as in the pure gauge theory [6]. In particular, the combination rises nearly linearly beyond a smoothing range of 0.2 fm or so.

6.2 Pseudo-scalar meson masses

The pion mass can be extracted from the pion propagator $G_\pi(x_0, y_0)$ using standard methods (see subsect. 5.2 for the definition of the propagator). In the case shown in

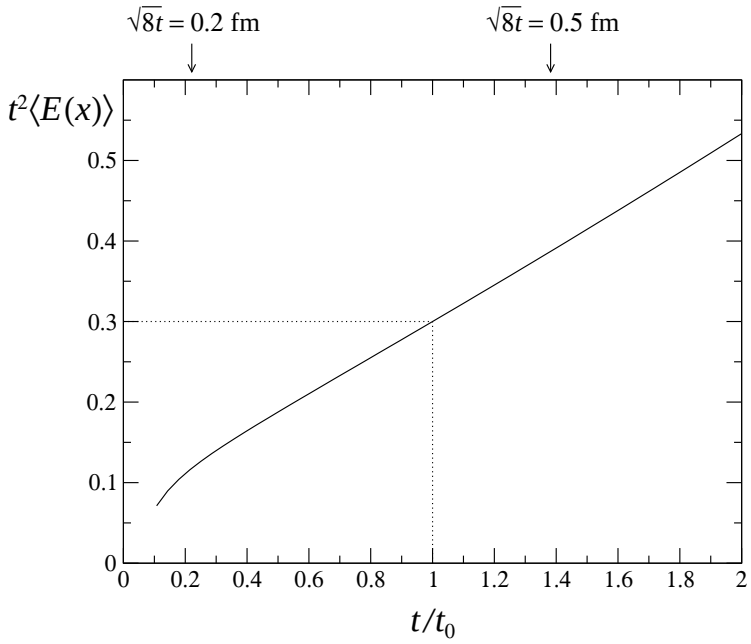

Fig. 5. Plot of $t^2\langle E(x)\rangle$ versus the flow time t in units of the reference scale t_0 as measured in the central region of the lattice in run I_1. Statistical errors are too small to be seen on the scale of the plot. The data at smoothing ranges $\sqrt{8t}$ less than 1.5 lattice spacings are strongly affected by lattice effects and are therefore not shown.

fig. 6, the source point is at $y_0 = 1$ and as a function of x_0 the propagator decreases roughly exponentially from there to the other end of the lattice. Apart from the fact that it falls off more rapidly, the kaon propagator behaves essentially in the same way. In the central region of the lattice, the effective masses determined from the propagators are constant within errors and a fit of the data yields $m_\pi = 0.0925(19)$ and $m_K = 0.2373(10)$ for the meson masses in lattice units (see fig. 7)†.

At small times x_0, the pseudo-scalar meson propagators plotted in fig. 6 receive contributions from higher-energy intermediate states as is the case on lattices with periodic boundary conditions. Deviations from a single-exponential curve are, however, also seen when x_0 approaches the boundary of the lattice at time T. Close to the chiral limit, and at distances from the boundary not smaller than 0.5 fm or so,

† In physical units, the calculated masses (203(4) and 520(2) MeV, respectively) are slightly smaller than the values quoted in table 2. The probability for the differences to be purely statistical is not completely negligible, but they could also derive from our interpolation of the results obtained in refs. [19,20] or from the presence of finite-volume effects in some of these data.

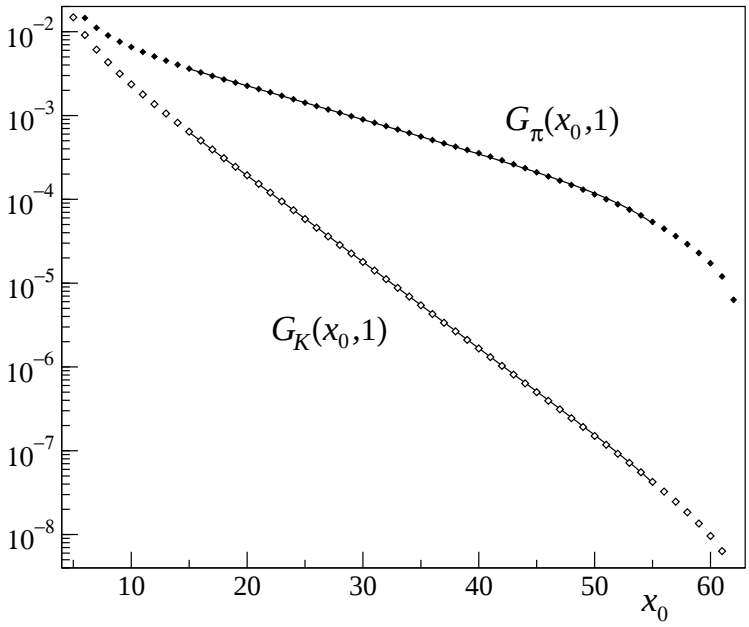

Fig. 6. Pion and kaon propagator in lattice units, calculated using an ensemble of 150 gauge-field configurations generated in run I_1. The lines are leading-order chiral perturbation theory fits to the data (eq. (6.4) and the corresponding expression for the kaon propagator).

these effects will be dominated by intermediate pseudo-scalar meson states and may therefore conceivably be described by chiral perturbation theory.

The boundary conditions to be used in chiral perturbation theory cannot be easily inferred from QCD. Dirichlet boundary conditions are however known to be the generic boundary conditions (i.e. those that do not require a fine-tuning or a particular symmetry pattern) in scalar field theories [41,42]. Since the flavour symmetry is preserved, the correct boundary conditions on the pion field π^a (where $a = 1, 2, 3$ is the isospin index) are thus likely to be

$$\pi^a(x)|_{x_0=0} = \pi^a(x)|_{x_0=T} = 0. \tag{6.3}$$

Note that these break the axial symmetries, as do the boundary conditions in the fundamental theory.

To leading order of chiral perturbation theory, the pion field satisfies the field equation $(-\Delta + m_\pi^2)\pi^a(x) = 0$. On-shell correlation functions of its zero-momentum component are therefore linear combinations of the exponentials $\exp\{\pm m_\pi x_0\}$. Tak-

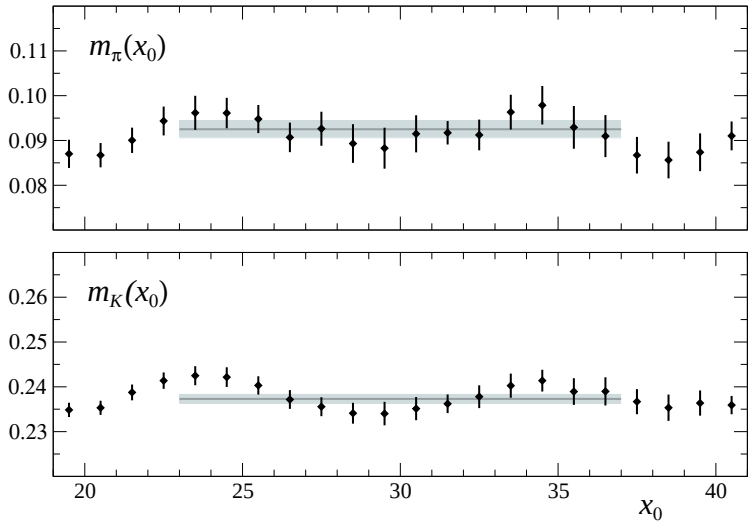

Fig. 7. Effective pion and kaon masses measured in the central region of the lattice in run I_1. The grey lines with their error bands were obtained through uncorrelated least-squares fits of the data.

ing the boundary conditions into account, the pion propagator plotted in fig. 6 is thus expected to be of the form

$$G_\pi(x_0, 1) \propto \sinh(m_\pi(T - x_0)) \tag{6.4}$$

in the central part of the lattice and at times x_0 close (but not too close) to T. Up to higher-order corrections, the same formula, with m_π replaced by m_K, should apply in the case of the kaon propagator.

The leading-order formulae actually fit the propagators quite well (curves in fig. 6). For the meson masses, the values quoted above were inserted in these fits and T was slightly adjusted from 63 to about 59 lattice spacings (in the chiral theory, T is an effective parameter whose value is dynamically determined by the properties of QCD near the boundaries). While further confirmation is clearly needed, the form of the meson propagators measured in run I_1 thus lends support to the conjecture that the chiral limit of QCD with open boundary conditions is described by the standard chiral effective theory with Dirichlet boundary conditions.

7. Concluding remarks

The use of open boundary conditions and twisted-mass determinant reweighting in numerical lattice QCD is profitable from the point of view stability, efficiency and conceptual clarity. While open boundary conditions slightly complicate the physics analysis of the calculated correlation functions, there are currently no practical alternative ways to avoid the well-known ergodicity problems related to the emergence of the topological charge sectors in the continuum limit. Twisted-mass determinant reweighting, on the other hand, ensures the absence of instabilities and sampling inefficiencies caused by accidental near-zero modes of the lattice Dirac operator when the Wilson formulation of the lattice theory is employed.

The simulation algorithm used in the runs reported in this paper combines twisted-mass determinant reweighting with a particular ("log-scale") Hasenbusch factorization [23,24] of the quark determinant and a hierarchical integrator for the molecular-dynamics equations based on some of the highly efficient integration rules proposed by Omelyan, Mryglod and Folk [32]. In all runs and with very little parameter tuning, an excellent stability and performance of the simulations could be achieved in this way.

There is every reason to expect that the situation will be essentially unchanged in this respect when larger and finer lattices than those considered here are simulated. The theoretical discussion in ref. [2] moreover suggests that the reweighting efficiency will depend only weakly on the lattice parameters if the second kind of twisted-mass determinant reweighting is used with an appropriate choice of the regulator mass.

Most simulations reported in this paper were performed on a dedicated PC cluster at CERN. We are grateful to the CERN management for funding this machine and to the CERN IT Department for technical support. Thanks also go to the John von Neumann Institute for Computing for computer time on a Blue Gene/P machine.

Appendix A. Gauge action

Let $\mathcal{S}_0$ and $\mathcal{S}_1$ be the sets of oriented 1×1 plaquette and 1×2 rectangular loops on the lattice (the time coordinate x_0 of the corners of all these loops must thus be in the range $0 \le x_0 \le T$). The gauge actions considered in this paper are of the form

$$S_{\mathrm{G}} = \frac{1}{g_0^2} \sum_{k=0}^{1} c_k \sum_{\mathcal{C} \in \mathcal{S}_k} w_k(\mathcal{C}) \, \mathrm{tr}\{1 - U(\mathcal{C})\}, \tag{A.1}$$

where $U(\mathcal{C})$ denotes the ordered product of the link variables $U(x,\mu)$ around $\mathcal{C}$ and $w_k(\mathcal{C})$ is a weight factor specified below. In order to ensure the correct normalization of the bare coupling g_0, the coefficients c_k must be such that

$$c_0 + 8c_1 = 1. \tag{A.2}$$

The Wilson plaquette, the tree-level Symanzik-improved [43] and the Iwasaki action [44] are obtained by setting $c_1 = 0$, $c_1 = -1/12$ and $c_1 = -0.331$, respectively. In all cases the standard convention $\beta = 6/g_0^2$ is used for the inverse coupling.

The weight factors $w_k(\mathcal{C})$ in eq. (A.1) are equal to 1 except for the space-like loops $\mathcal{C}$ on the boundaries of the lattice at time 0 and T, where

$$w_k(\mathcal{C}) = \tfrac{1}{2}c_{\rm G}. \tag{A.3}$$

As previously discussed in ref. [1], the coefficient $c_{\rm G}$ is required for O(a) improvement of correlation functions involving fields close to or at the boundaries of the lattice. In particular, setting $c_{\rm G} = 1$ ensures on-shell improvement at tree-level of perturbation theory.

Appendix B. Wilson flow

On lattices with periodic boundary conditions, the Wilson flow $V_t(x,\mu)$ of lattice gauge fields is defined by the equations

$$\partial_t V_t(x,\mu) = -g_0^2 \left\{\partial^a_{x,\mu} S_{\rm w}(V_t)\right\} T^a V_t(x,\mu), \qquad V_t(x,\mu)|_{t=0} = U(x,\mu), \tag{B.1}$$

where $S_{\rm w}$ denotes the Wilson plaquette action and the parameter $t \geq 0$ is referred to as the flow time (see ref. [6] for an introduction to the subject).

When open boundary conditions are imposed, the flow equation assumes a slightly different form,

$$\partial_t V_t(x,0) = -g_0^2 \big\{\partial^a_{x,0} S_{\rm G}(V_t)\big\} T^a V_t(x,0), \quad 0 \leq x_0 < T, \tag{B.2}$$

$$\partial_t V_t(x,k) = -\frac{g_0^2}{w(x_0)} \big\{\partial^a_{x,k} S_{\rm G}(V_t)\big\} T^a V_t(x,k), \quad 0 \leq x_0 \leq T, \quad k = 1,2,3, \tag{B.3}$$

where $S_{\rm G}$ denotes the Wilson action (A.1) with $c_{\rm G} = 1$. The weight factor

$$w(x_0) = \begin{cases} \frac{1}{2} & \text{if } x_0 = 0 \text{ or } x_0 = T, \\ 1 & \text{otherwise,} \end{cases} \tag{B.4}$$

is required to ensure the absence of $\mathrm{O}(a)$ lattice effects in expectation values of gauge invariant quantities at flow time $t > 0$.

Note that $\mathrm{O}(a)$ improvement is guaranteed if the flow is improved at tree-level of perturbation theory [40], a property which can be easily established by calculating the time-dependent propagator of the gauge field. The flow equation (B.2)–(B.4) can also be derived using an orbifold construction previously employed by Taniguchi [45] in a different context. $\mathrm{O}(a)$ lattice effects are then seen to be excluded by symmetry.

References

[1] M. Lüscher, S. Schaefer, *Lattice QCD without topology barriers*, JHEP 1107 (2011) 036

[2] M. Lüscher, F. Palombi, *Fluctuations and reweighting of the quark determinant on large lattices*, PoS (LATTICE 2008) 049

[3] L. Del Debbio, H. Panagopoulos, E. Vicari, *θ-dependence of SU(N) gauge theories*, JHEP 08 (2002) 044

[4] S. Schaefer, R. Sommer, F. Virotta, *Investigating the critical slowing down of QCD simulations*, PoS (LAT2009) 032

[5] S. Schaefer, R. Sommer, F. Virotta, *Critical slowing down and error analysis in lattice QCD simulations*, Nucl. Phys. B845 (2011) 93

[6] M. Lüscher, *Properties and uses of the Wilson flow in lattice QCD*, JHEP 1008 (2010) 071

[7] M. Lüscher, *Topology, the Wilson flow and the HMC algorithm*, PoS (Lattice 2010) 015

[8] K. G. Wilson, *Confinement of quarks*, Phys. Rev. D10 (1974) 2445

[9] B. Sheikholeslami, R. Wohlert, *Improved continuum limit lattice action for QCD with Wilson fermions*, Nucl. Phys. B259 (1985) 572

[10] M. Lüscher, S. Sint, R. Sommer, P. Weisz, *Chiral symmetry and O(a) improvement in lattice QCD*, Nucl. Phys. B478 (1996) 365

[11] S. Duane, A. D. Kennedy, B. J. Pendleton, D. Roweth, *Hybrid Monte Carlo*, Phys. Lett. B195 (1987) 216.

[12] L. Del Debbio, L. Giusti, M. Lüscher, R. Petronzio, N. Tantalo, *Stability of lattice QCD simulations and the thermodynamic limit*, JHEP 0602 (2006) 011

[13] L. Del Debbio, L. Giusti, M. Lüscher, R. Petronzio, N. Tantalo, *QCD with light Wilson quarks on fine lattices (II): DD-HMC simulations and data analysis*, JHEP 0702 (2007) 082

[14] P. H. Damgaard, K. Splittorff, J. J. M. Verbaarschot, *Microscopic spectrum of the Wilson Dirac operator*, Phys. Rev. Lett. 105 (2010) 162002

[15] G. Akemann, P. H. Damgaard, K. Splittorff, J. J. M. Verbaarschot, *Spectrum of the Wilson Dirac operator at finite lattice spacings*, Phys. Rev. D83 (2011) 085014

[16] K. Splittorff, J. J. M. Verbaarschot, *The Wilson Dirac spectrum for QCD with dynamical quarks*, Phys. Rev. D84 (2011) 065031

[17] M. Kieburg, K. Splittorff, J. J. M. Verbaarschot, *The realization of the Sharpe-Singleton scenario*, arXiv:1202.0620v1

[18] C. Miao, H. B. Meyer, H. Wittig, *Twisted-mass reweighting for O(a) improved Wilson fermions*, PoS (Lattice 2011) 041

[19] S. Aoki et al. (PACS-CS collab.), *2+1 flavor lattice QCD toward the physical point*, Phys. Rev. D79 (2009) 034503

[20] S. Aoki et al. (PACS-CS collab.), *Physical point simulation in 2+1 flavor lattice QCD*, Phys. Rev. D81 (2010) 074503

[21] `http://cern.ch/luscher/openQCD`

[22] A. Hasenfratz, R. Hoffmann, S. Schaefer, *Reweighting towards the chiral limit*, Phys. Rev. D78 (2008) 014515

[23] M. Hasenbusch, *Speeding up the Hybrid Monte Carlo algorithm for dynamical fermions*, Phys. Lett. B519 (2001) 177

[24] M. Hasenbusch, K. Jansen, *Speeding up lattice QCD simulations with clover-improved Wilson fermions*, Nucl. Phys. B659 (2003) 299

[25] C. Urbach, K. Jansen, A. Shindler, U. Wenger, *HMC algorithm with multiple time scale integration and mass preconditioning*, Comp. Phys. Commun. 174 (2006) 87

[26] M. Lüscher, *Schwarz-preconditioned HMC algorithm for two-flavor lattice QCD*, Comp. Phys. Commun. 165 (2005) 199

[27] I. Horvath, A. D. Kennedy, S. Sint, *A new exact method for dynamical fermion computations with non-local actions*, Nucl. Phys. (Proc. Suppl.) 73 (1999) 834

[28] M. A. Clark, A. D. Kennedy, *Accelerating dynamical fermion computations using the Rational Hybrid Monte Carlo (RHMC) algorithm with multiple pseudo-fermion fields*, Phys. Rev. Lett. 98 (2007) 051601

[29] M. Lüscher, *Computational strategies in lattice QCD*, in: *Modern perspectives in lattice QCD*, eds. L. Lellouch et al. (Oxford University Press, New York, 2011) [arXiv: 1002.4232]
[30] N. I. Achiezer, *Theory of approximation* (Dover Publications, New York, 1992)
[31] J. C. Sexton, D. H. Weingarten, *Hamiltonian evolution for the Hybrid Monte Carlo algorithm*, Nucl. Phys. B380 (1992) 665
[32] I. P. Omelyan, I. M. Mryglod, R. Folk, *Symplectic analytically integrable decomposition algorithms: classification, derivation, and application to molecular dynamics, quantum and celestial mechanics simulations*, Comp. Phys. Commun. 151 (2003) 272
[33] M. Lüscher, *Local coherence and deflation of the low quark modes in lattice QCD*, JHEP 0707 (2007) 081
[34] M. Lüscher, *Deflation acceleration of lattice QCD simulations*, JHEP 0712 (2007) 011
[35] B. Jegerlehner, *Krylov space solvers for shifted linear systems*, preprint IUHET-353 (1996) [arXiv: hep-lat/9612014]
[36] M. Lüscher, *Solution of the Dirac equation in lattice QCD using a domain decomposition method*, Comp. Phys. Commun. 156 (2004) 209
[37] K. Jansen, R. Sommer, *O(a) improvement of lattice QCD with two flavors of Wilson quarks*, Nucl. Phys. B530 (1998) 185 [E: *ibid.* B643 (2002) 517]
[38] S. Aoki et al. (CP-PACS collab.), *Nonperturbative O(a) improvement of the Wilson quark action with the RG-improved gauge action using the Schrödinger functional method*, Phys. Rev. D73 (2006) 034501
[39] P. Fritzsch et al. (ALPHA collab.), *The strange quark mass and Lambda parameter of two flavor QCD*, arXiv:1205.5380
[40] M. Lüscher, P. Weisz, *Perturbative analysis of the gradient flow in non-Abelian gauge theories*, JHEP 1102 (2011) 051
[41] K. Symanzik, *Schrödinger representation and Casimir effect in renormalizable quantum field theory*, Nucl. Phys. B190 [FS3] (1981) 1
[42] M. Lüscher, *The Schrödinger functional in lattice QCD with exact chiral symmetry*, JHEP05(2006)042
[43] P. Weisz, *Continuum limit improved lattice action for pure Yang-Mills theory (I)*, Nucl. Phys. B212 (1983) 1
[44] Y. Iwasaki, *Renormalization group analysis of lattice theories and improved lattice action. II – Four-dimensional non-Abelian SU(N) gauge model*, preprint UTHEP-118 (1983) and arXiv:1111.7054v1
[45] Y. Taniguchi, *Schrödinger functional formalism with Ginsparg–Wilson fermion*, JHEP 0512 (2005) 037;