THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Stability of lattice QCD simulations and the thermodynamic limit

Stability of lattice QCD simulations and the thermodynamic limit

L. Del Debbio[1], L. Giusti[1,*], M. Lüscher[1], R. Petronzio[2], N. Tantalo[2,3]

[1] *CERN, Physics Department, TH Division, CH-1211 Geneva 23, Switzerland*

[2] *Università di Roma "Tor Vergata" and INFN sezione "Tor Vergata", Via della Ricerca Scientifica 1, I-00133 Rome, Italy*

[3] *Centro Enrico Fermi, Via Panisperna 89 A, I-00184 Rome, Italy*

Abstract

We study the spectral gap of the Wilson–Dirac operator in two-flavour lattice QCD as a function of the lattice spacing a, the space-time volume V and the current-quark mass m. It turns out that the median of the probability distribution of the gap scales proportionally to m and that its width is practically equal to $a/\sqrt{V}$. In particular, numerical simulations are safe from accidental zero modes in the large-volume regime of QCD.

1. Introduction

The formulation of lattice QCD proposed by Wilson long ago preserves many symmetries of the continuum theory exactly [1]. An infamous exception are the chiral symmetries, and although the symmetry-violating terms vanish proportionally to the lattice spacing (or its square if the theory is Symanzik-improved [2,3]), the presence of these terms complicates the lattice theory considerably. In particular, the fact that the massive Wilson–Dirac operator is not protected from arbitrarily small eigenvalues may lead to instabilities in numerical simulations.

Such instabilities were not observed, however, in recent simulations of the two-flavour Wilson theory on large lattices, not even at the smallest quark masses consid-

* On leave from Centre Physique Théorique, CNRS Luminy, F-13288 Marseille, France

ered [4–6]. It is tempting to attribute the absence of instabilities in these simulations to the use of a new simulation algorithm, but this explanation cannot be right, because the distribution of the spectral gap of the lattice Dirac operator (and thus the probability to find exceptionally small eigenvalues) is determined by the lattice action and the functional integral, and not by the method used to evaluate the latter.

Our aim in this paper is to clarify the situation by calculating the distribution of the gap on a set of lattices, using numerical simulations. In particular, we wish to determine, as explicitly as possible, the range in parameter space, where the Wilson theory can be simulated without running into instabilities.

2. Algorithmic stability and the spectral gap

In this section we introduce our notation and discuss the relevance of the distribution of the spectral gap for the stability of lattice QCD simulations. For any unexplained notation see ref. [3].

2.1 Lattice Dirac operator

The lattice theory is set up as usual on hypercubic lattices with spacing a. Periodic boundary conditions are imposed on all fields and in all directions, except for the quark fields, which are taken to be antiperiodic in time. Throughout this paper, we assume that there is a doublet of sea quarks with equal mass, although many results are likely to remain valid if the number of quark flavours is larger than two.

While the action of the gauge field will always be the Wilson plaquette action, we shall consider various lattice Dirac operators D_{m} (the subscript indicates that D_{m} includes the quark mass term). The Wilson–Dirac operator

$$D_{\mathrm{m}} = D_{\mathrm{w}} + m_0, \tag{2.1}$$

$$D_{\mathrm{w}} = \tfrac{1}{2} \left\{ \gamma_\mu \left(\nabla_\mu^* + \nabla_\mu \right) - a \nabla_\mu^* \nabla_\mu \right\}, \tag{2.2}$$

is the one we are primarily interested in, as well as its $\mathrm{O}(a)$-improved version [2,3]. In these equations, ∇_μ and ∇_μ^* denote the gauge-covariant forward and backward difference operators, and m_0 the bare quark mass.

Rather than the Dirac operator itself, we prefer to consider the hermitian operator

$$Q_{\mathrm{m}} = \gamma_5 D_{\mathrm{m}} \tag{2.3}$$

2

in the following. The determinants of these two operators are the same, but the fact that the spectrum of Q_m is real simplifies the discussion considerably. On a finite lattice, and for any specified gauge field, we then define the spectral gap

$$\mu = \min \left\{ |\lambda| \mid \lambda \text{ is an eigenvalue of } Q_m \right\}. \tag{2.4}$$

Evidently, the gap is a well-defined function of the gauge field and so is the spectral asymmetry

$$\eta = \tfrac{1}{2} \left\{ N_+ - N_- \right\}, \tag{2.5}$$

where $N_\pm$ are the numbers of positive and negative eigenvalues of Q_m.

In formulations of lattice QCD that preserve chiral symmetry via the Ginsparg–Wilson relation [7–11], the gap is bounded from below by the bare current-quark mass m and the asymmetry vanishes if $m > 0$ (see ref. [12], for example). However, these properties are not guaranteed in the Wilson theory, and it is possible, in the presence of some gauge-field configurations, that the gap is much smaller than m and that the asymmetry assumes a non-zero value.

2.2 Sources of instability

In the large-volume regime of QCD, and at large quark masses, the probability distribution $p(\mu)$ of the gap typically looks like the one shown in the upper plot in fig. 1. The subsets of gauge fields, where μ is far below the central value of the distribution, occur with such a small probability in this case that their contributions to the common physical observables can be safely neglected. Numerical simulations, using a preconditioned Hybrid Monte Carlo (HMC) algorithm [13], for example, will then normally run smoothly and produce a representative sample of field configurations as expected.

When the quark mass decreases, the centre of the gap distribution moves towards smaller values and eventually is no more than one or two standard deviations away from the origin (lower plot in fig. 1). At this point numerical simulations may run into instabilities for the following reasons:

(1) *Integration instabilities.* The HMC algorithm obtains the next field configuration by integrating the appropriate molecular-dynamics equations in field space, followed by an accept–reject step at the end of the integration. It is possible that the molecular-dynamics trajectories pass through field configurations where the gap μ is exceptionally small. The numerical integration becomes unstable in this case
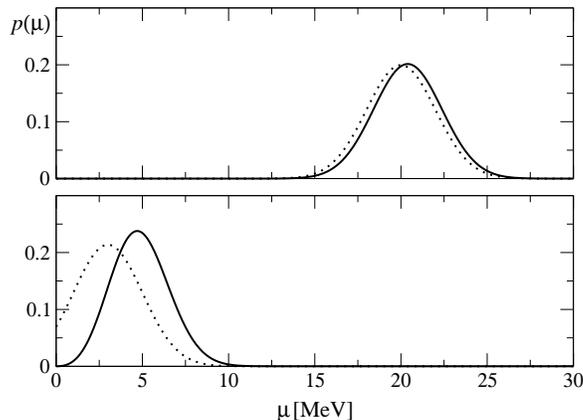
3

Fig. 1. Qualitative form of the normalized probability distribution $p(\mu)$ of the gap (full lines), at two values of the quark mass (upper and lower plot). The normalized weighted distributions proportional to $p(\mu)/\mu^2$ are also shown (dotted lines).

and liable to rounding errors. In particular, the reversibility of the integration is then no longer guaranteed, which invalidates the algorithm.

(2) *Ergodicity.* Along the molecular-dynamics trajectories, the spectral asymmetry η tends to be preserved. Changes in the asymmetry are in fact excluded in the limit where the integration step-size goes to zero. However, the sectors in field space, in which the asymmetry does not vanish, may become statistically relevant in the situation considered here. The HMC algorithm is then likely to give wrong results, because it may get stuck, for a very long time, in any one of these sectors.

(3) *Sampling inefficiencies.* Observables that are sensitive to the small eigenvalues of the Dirac operator, such as the pion propagator, may be poorly sampled by the simulation. The effect is illustrated by the dotted curves in fig. 1, which represent the distributions of the gap reweighted by the "observable" $1/\mu^2$. Expectation values of such quantities are difficult to compute reliably, since the subspace of fields where μ is very small is only rarely visited in the course of the simulation.

Instabilities of this kind can lead to underestimated statistical errors and incorrect results, and they may even suggest the presence of a phase transition when there is none. Since the QCD functional integral remains well defined, also in this difficult regime, improved simulation techniques may conceivably be developed, which do not suffer from any instabilities. Whether this is worth the effort is not obvious, however, because the theory may be strongly affected by lattice artefacts (resulting from a competition of mass and discretization effects) in the critical range of parameters.

4

3. Numerical studies

In the presence of any given gauge field, the spectral gap of the Dirac operator can be computed numerically, using a suitable iterative method (appendix A). Evidently, the histograms of the values calculated in the course of a numerical simulation approximate the gap distribution up to statistical errors. In this section, we report the results of such numerical studies and show that the data are well described by a few simple empirical laws.

3.1 Simulation parameters

The simulations listed in table 1 are part of an ongoing study of two-flavour QCD in the chiral regime [6]. Technically the project is based on the use of the Schwarz-preconditioned HMC simulation algorithm introduced in ref. [4], which allows the theory to be simulated in a range of lattices and quark masses that was, in practice, inaccessible so far.

Except for the last run in table 1, the unimproved Wilson theory was simulated, at inverse bare coupling β and hopping parameter $\kappa = (8 + 2am_0)^{-1}$. In run D_1 the coefficient c_{sw} of the Sheikholeslami–Wohlert improvement term [2,3] was set to the value determined by the ALPHA collaboration [14]. A number N_{cfg} of statistically decorrelated gauge-field configurations was generated in each case and later used for the calculation of the gap distributions.

Table 1. Simulation runs included in this study

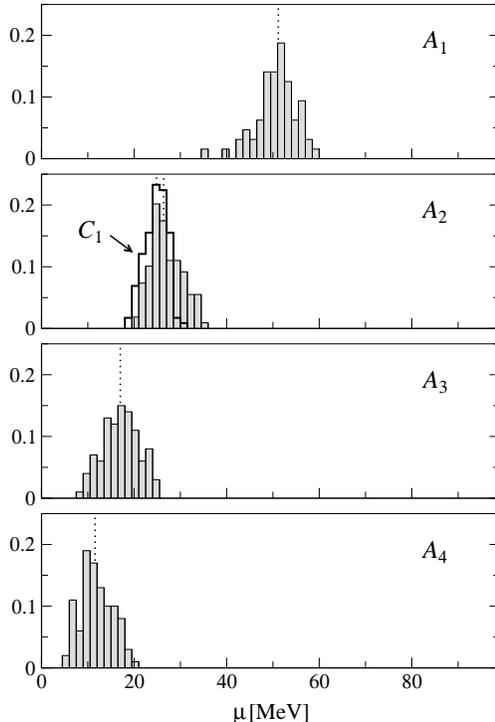| Run	Lattice	β	c_{sw}	κ	N_{cfg}	am_π
A_1	32×24^3	5.6	0	0.15750	64	0.2744(21)
A_2				0.15800	109	0.1969(16)
A_3				0.15825	100	0.1554(31)
A_4				0.15835	100	0.1204(44)
B_1	64×32^3	5.8	0	0.15410	100	0.1965(8)
B_2				0.15440	101	0.1481(11)
C_1	64×24^3	5.6	0	0.15800	116	0.1986(10)
D_1	48×24^3	5.3	1.90952	0.13550	104	0.3265(11)

5

Fig. 2. Normalized histograms of the spectral gap μ, as obtained in the runs A_1–A_4 and C_1. The bin size is 1.5 MeV and the dotted vertical lines indicate the position of the median $\bar{\mu}$ of the distributions.

Details of the simulations, the computation of the quark masses and other physical quantities will be given in a forthcoming publication [6]. Based on calculations of the Sommer reference scale r_0 [15], the lattice spacing in the Wilson theory is estimated to be about 0.08 and 0.06 fm at $\beta = 5.6$ and 5.8, respectively, while it is roughly 0.09 fm in the improved theory at $\beta = 5.3$. In physical units the spatial sizes of the lattices in table 1 are thus close to 2 fm in all cases.

As usual the conversion from lattice to physical units is ambiguous, but in this paper the latter serve for the purpose of illustration only. The pion masses quoted in the last column of table 1, for example, cover a range from 676 to about 294 MeV in the case of the runs A_1–A_4.

3.2 Mass-dependence of the gap distribution

The calculated distributions of the gap in the Wilson theory at $\beta = 5.6$ are plotted in fig. 2. A characteristic feature of these distributions is that their shape does not
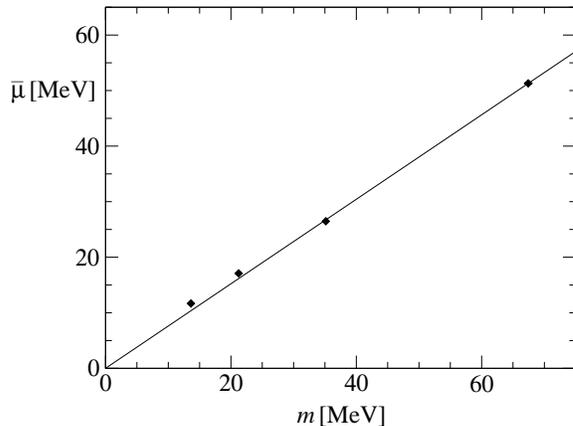
Fig. 3. Median $\bar{\mu}$ of the gap distribution on the 32×24^3 lattice (runs A_1–A_4) as a function of the bare current-quark mass m. The straight line from the point at the largest mass to the origin is drawn to guide the eye. Statistical errors are negligibly small on the scale of this plot.

show a strong dependence on the quark mass. Basically they are shifted to smaller values when the quark mass decreases. All these distributions are clearly separated from the origin, although the smallest mass may be quite close to where the unstable regime begins.

The distributions are roughly symmetric about their median (dotted lines in fig. 2), which is practically also the point where they assume their maximal value. As shown in fig. 3, the median scales approximately linearly with the current-quark mass. A small but significant effect is seen at the lighter quark masses, where the median is pushed to slightly larger values with respect to the straight scaling curve.

On the big lattices at $\beta = 5.8$, the situation is essentially the same, although here only two simulations have been completed so far (see fig. 4). In particular, the quark masses in these two runs are in a range where a nearly perfect scaling of the median as a function of the quark mass is again observed.

3.3 Statistical fluctuations of the gap

An important qualitative result of our empirical studies is that the fluctuations of the gap become smaller when the four-dimensional volume V of the lattice increases. The simulation C_1, for example, was performed at exactly the same coupling and quark mass as the simulation A_2, but on a larger lattice. As can be seen from fig. 2, the width of the gap distribution obtained in run C_1 is visibly reduced with respect to the one obtained in run A_2.
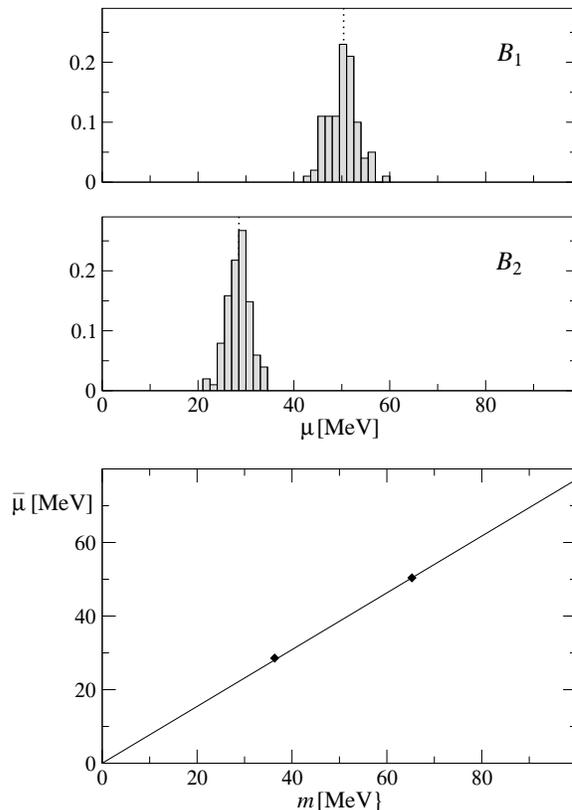
Fig. 4. Gap distributions and mass dependence of the median $\bar{\mu}$ on the 64×32^3 lattices (runs B_1 and B_2). The bin size is the same as in fig. 2.

A simple heuristic argumentation that may explain this effect goes as follows. Let U be a given lattice gauge field, μ the lowest eigenvalue of $|Q_{\mathrm{m}}|$ and ψ the associated eigenvector normalized to unity. Consider a small random fluctuation $U + \delta U$ of the gauge field. To first order, the change $\delta\mu$ in the gap is then given by

$$\delta\mu = a^4 \sum_x \psi(x)^\dagger \left[\delta Q_{\mathrm{m}} \psi\right](x). \tag{3.1}$$

At each point x, there are eight contributions to this sum, which are proportional to the random variation δU of the gauge field on the links attached there. Now if ψ extends over the whole lattice, i.e. if ψ it is not a localized mode, the size of these terms is proportional to a^3/V. On average this implies $\langle\delta\mu\rangle = 0$ and $\langle(\delta\mu)^2\rangle \propto a^2/V$.

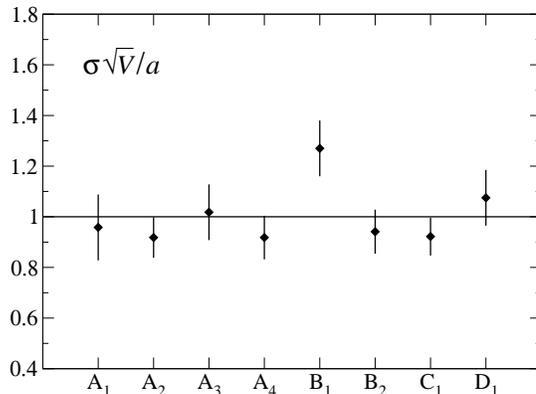While this kind of reasoning is extremely superficial, the width σ of the numerically

8

Fig. 5. Width σ of the gap distributions, given in units of $a/\sqrt{V}$, as obtained in runs A_1–D_1. Note that the combination $\sigma\sqrt{V}/a$ is dimensionless and can be computed directly from the lattice data, without intermediate conversion to physical units. The statistical errors were determined using the bootstrap method.

determined gap distributions appears to scale in the suggested way (see fig. 5) †. An interesting special case is that of the distributions obtained in the runs A_2 and C_1, where, as already mentioned, all parameters were the same except for the lattice volumes. The scaled widths $\sigma\sqrt{V}/a$, however, agree very well with each other. The approximate scaling law

$$\sigma \simeq \frac{a}{\sqrt{V}} \tag{3.2}$$

is actually consistent with all the data shown in fig. 5.

Somewhat surprisingly, the spectral gap thus seems to behave thermodynamically, like a free-energy density for example, with probability distribution

$$p(\mu) \propto \exp\left\{-\frac{V}{2a^2}\left(\mu - \langle\mu\rangle\right)^2\right\}. \tag{3.3}$$

The fact that σ is proportional to the lattice spacing suggests, on the other hand, that the observed widths result from short-distance fluctuations of the gauge field. These fluctuations scale more slowly with the volume than those expected in lattice theories

---

† The standard deviation of μ is subject to potentially large statistical uncertainties, because the tails of the gap distributions are poorly sampled. We therefore define the width of the distributions through $\sigma = \frac{1}{2}(v - u)$, where $[u, v]$ is the smallest range in μ, which contains more than 68.3% of the data.

that preserve chiral symmetry, where the eigenvalue distributions are universally computable, using chiral perturbation theory or random matrix theory.

4. Spectral gap in infinite volume

Analytical calculations of the gap distribution would evidently be very welcome at this point, but since chiral symmetry is violated in the Wilson theory, it is unclear how such calculations would proceed. The spectral density of the Dirac operator in infinite volume is somewhat more accessible and provides a useful reference for the situation on the lattices that can be simulated.

4.1 Spectral density

Let $\alpha_1, \alpha_2, \ldots$ be the eigenvalues of Q_{m}^2, ordered in ascending order and counting multiplicities. In the following we will be interested in the spectral density

$$\rho(\alpha) = \lim_{V \to \infty} \frac{1}{V} \sum_{k \geq 1} \langle \delta(\alpha - \alpha_k) \rangle \tag{4.1}$$

in infinite volume. It is, incidentally, possible to prove rigorously that the thermodynamic limit (4.1) exists, using a general argument based on a decomposition of the lattice into large blocks [16].

In the continuum theory, or if chiral symmetry is preserved by the lattice theory, the spectral density vanishes at $\alpha < m^2$. Moreover, for slightly larger values of α, the Banks–Casher relation [17] implies the asymptotic form

$$\rho(\alpha) \underset{\alpha > m^2}{=} \frac{\Sigma}{\pi \sqrt{\alpha - m^2}} + \mathrm{O}(1), \tag{4.2}$$

where Σ denotes the quark condensate. The spectral density thus has a well-defined threshold at $\alpha = m^2$ in these cases (see fig. 6).

In the following a *working hypothesis* is that there is a similar threshold $\bar{\alpha} > 0$ in the Wilson theory or its $\mathrm{O}(a)$-improved version (whichever is considered). We shall not need to know the shape of the spectral density in the vicinity of the threshold, but it must be guaranteed that $\rho(\alpha)$ vanishes if $\alpha < \bar{\alpha}$ and that the density is non-zero at or immediately above $\bar{\alpha}$. It is important to understand that $\bar{\alpha}$ is the point where the dense spectrum begins, while it is perfectly possible that the Wilson–Dirac
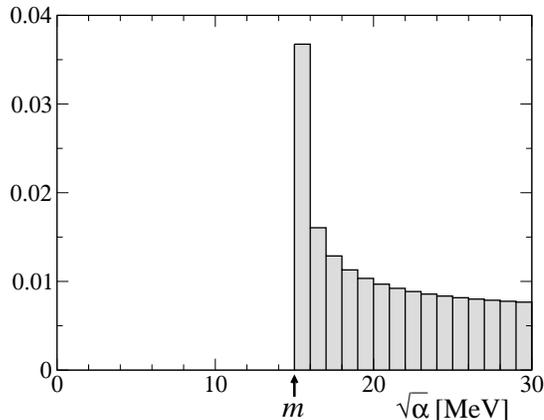
10

Fig. 6. Number of eigenvalues per bin and fm^4 at the low end of the spectral density in the continuum theory [cf. eq. (4.2)]. In this plot, $m = 15$ MeV and $\Sigma = (250\,\text{MeV})^3$ was assumed.

operator in finite volume has eigenvalues significantly smaller than $\sqrt{\overline{\alpha}}$. Any part of the spectrum with less than $O(V)$ modes per energy bin is in fact suppressed by the factor $1/V$ in the limit (4.1).

4.2 Adding valence quarks

To be able to relate the spectral density to correlation functions of local fields, we will need to consider partially quenched QCD, where $2N$ valence quarks are added to the theory [18,19]. The fermion action then becomes

$$S_{\text{F}} = a^4 \sum_x \left\{ \sum_{r=1}^{2N+2} \overline{\psi}_r(x) D_m \psi_r(x) + \sum_{k=1}^{N} |D_m \phi_k(x)|^2 \right\}, \qquad (4.3)$$

where $\overline{\psi}_r, \psi_r$ are the quark fields (2 sea quarks plus $2N$ valence quarks) and ϕ_k the pseudo-fermion fields that are required to cancel the valence-quark determinants. This defines a well-behaved euclidean field theory, whose renormalization can be expected to follow the usual rules. In particular, the obvious $U(2N + 2) \times U(N)$ flavour symmetry and some less obvious graded symmetries restrict the possible counterterms to the naively expected ones.

We now also introduce the bare scalar and pseudo-scalar densities

$$S_{rs} = \overline{\psi}_r \psi_s, \qquad P_{rs} = \overline{\psi}_r \gamma_5 \psi_s. \qquad (4.4)$$

These fields renormalize in the standard manner, i.e. for unequal flavours r, s, the

11

renormalized operators are $Z_\mathrm{S} S_{rs}$ and $Z_\mathrm{P} P_{rs}$, where the renormalization constants Z_S and Z_P are flavour-independent and can be taken to be the same as the ones in the theory without valence quarks.

4.3 Resolvent & moments of the spectral density

The resolvent

$$R(z) = \int_{\bar{\alpha}}^{\infty} \mathrm{d}\alpha \, \frac{\rho(\alpha)}{\alpha^2(z - \alpha)}, \tag{4.5}$$

is an analytic function of z with a cut along the real axis from $z = \bar{\alpha}$ to some value proportional to $1/a^2$. For later convenience, a factor α^{-2} is included in the integral to improve its convergence at large eigenvalues. Evidently, the spectral density is proportional to the discontinuity across the cut, and so can be uniquely recovered from the resolvent if the latter is known [20,21].

For $|z| < \bar{\alpha}$, the resolvent may be expanded in a convergent power series,

$$R(z) = \sum_{k=0}^{\infty} M_k z^k, \qquad M_k = -\int_{\bar{\alpha}}^{\infty} \mathrm{d}\alpha \, \frac{\rho(\alpha)}{\alpha^{k+3}}, \tag{4.6}$$

with coefficients M_k that can be written as

$$M_k = a^{4n-4} \sum_{x_1,\ldots,x_{n-1}} \langle P_{12}(x_1) P_{23}(x_2) \ldots P_{n1}(0) \rangle, \qquad n = 2k + 6. \tag{4.7}$$

This formula assumes that there are at least n quarks, but since each moment M_k may be considered separately, a sufficient number of valence quarks can always be added to the theory.

4.4 Renormalization

Equations (4.5)–(4.7) connect the spectral density to the basic field-theoretic correlation functions whose renormalization properties are well understood. The suggestion is then that the moments M_k get renormalized through multiplication by the factor $(Z_\mathrm{P})^n$. Since the sum over the coordinates $x_1, \ldots, x_{n-1}$ in eq. (4.7) includes the short-distance regions, it is, however, not totally obvious that the moments thus renormalized will indeed have a well-defined continuum limit.

In the continuum theory, and when inserted in correlation functions, the product $P_{12}(x_1) P_{23}(x_2)$, for example, has a short-distance expansion

$$P_{12}(x_1) P_{23}(x_2) \underset{x_1 \to x_2}{\sim} C(x_1 - x_2) S_{13}(x_2) + \ldots, \tag{4.8}$$

12

where the Wilson coefficient $C(x)$ diverges like $|x|^{-3}$ (up to logarithms). This singularity is integrable and does not give rise to any additional ultraviolet divergencies. As it turns out, all short-distance singularities of the correlation functions in eq. (4.7) are in fact integrable†. The ultraviolet divergencies of the moments are thus completely cancelled by the renormalization factor $(Z_P)^n$ and by the usual renormalization of the gauge coupling and the quark mass.

Recalling the expansion (4.6) of the resolvent $R(z)$ and its relation to the spectral density, it now follows that

$$\rho_R(\alpha) = Z_P^2 \rho(Z_P^2 \alpha) \tag{4.9}$$

has a universal continuum limit (once a particular renormalization condition for the pseudo-scalar density is adopted). The same must then also apply to the renormalized threshold

$$\bar{\alpha}_R = Z_P^{-2} \bar{\alpha}, \tag{4.10}$$

since the threshold is a property of the spectral density.

Without proof we mention in passing that the improved renormalized density in the $\mathrm{O}(a)$-improved theory is again given by eq. (4.9), provided Z_P is replaced by

$$Z_P \frac{1 + b_P a m_q}{1 + b_{PP} a m_q}, \tag{4.11}$$

where b_P and b_{PP} are improvement coefficients and m_q the additively renormalized bare quark mass (the notation is as in ref. [3]). The correction proportional to b_{PP} in this expression is in fact all that is needed to cancel the terms of order a that arise from the integrations over the short-distance singularities of the correlation function on the right of eq. (4.7).

4.5 Relation to the current-quark mass

The bare current-quark mass m is usually determined through vacuum-to-pion matrix elements of the isovector axial current and density. From the renormalization properties of the matrix elements it then follows that the renormalized quark mass

---

† When all coordinates $x_1, \ldots, x_{n-1}$ are scaled to zero, the operator product reduces to a coefficient function times the unit operator. The degree of divergence of the integral is $4 - n$ in this case, and convergence is thus guaranteed since $n \geq 6$.

13

is given by

$$m_{\mathrm{R}} = Z_{\mathrm{A}} Z_{\mathrm{P}}^{-1} m, \qquad (4.12)$$

where Z_{A} denotes the axial-current renormalization constant.

Once a definite renormalization condition for the isovector axial density is adopted, the threshold $\bar{\alpha}_{\mathrm{R}}$ and the quark mass m_{R} become physical quantities. In particular, in the continuum limit we have

$$\frac{\sqrt{\bar{\alpha}_{\mathrm{R}}}}{m_{\mathrm{R}}} = 1, \qquad (4.13)$$

not only when the limit is reached from a lattice theory that preserves chiral symmetry, because this ratio is dimensionless, unambiguously normalized and therefore independent of the regularization. In terms of the bare quantities this implies $\sqrt{\bar{\alpha}} = Z_{\mathrm{A}} m$, up to corrections of $\mathrm{O}(a)$ and independently of the normalization convention for the renormalized axial density [the renormalization constant Z_{P} cancels in the ratio (4.13)].

As already mentioned, the spectral gap in finite volume may not be related in any simple way to the threshold of the spectral density in the thermodynamic limit. It is nevertheless instructive to compare the median $\bar{\mu}$ of the gap distributions discussed in sect. 3 with the threshold $Z_{\mathrm{A}} m$.

Table 2. Comparison of the median $\bar{\mu}$ with $Z_{\mathrm{A}} m$

Run	Z_{A}	$\bar{\mu}/m$	$\bar{\mu} - Z_{\mathrm{A}} m$ [MeV]	$\langle \Delta \rangle$ [MeV]
A_1	$0.77(2)^a$	0.76(1)	$-0.5(5)$	2.8(1)
A_2		0.75(1)	$-0.6(4)$	3.2(1)
A_3		0.80(3)	0.6(5)	3.5(1)
A_4		0.85(3)	1.1(4)	3.9(1)
B_1	$0.78(2)^a$	0.77(1)	$-0.3(3)$	1.91(6)
B_2		0.79(1)	$-0.2(3)$	2.19(6)
C_1	$0.77(2)^a$	0.72(1)	$-1.7(3)$	2.18(7)
D_1	$0.75(1)^b$	0.68(1)	$-5.4(5)$	2.2(1)

[a] RI–MOM method [22]

[b] Schrödinger functional chiral Ward identity [23]

The numerically computed values of $\bar{\mu}/m$ quoted in table 2 actually agree quite well with the available estimates of Z_A. There are, however, significant differences in the last two rows of the table (runs C_1 and D_1), which underlines the fact that there is currently no solid theoretical understanding of the gap distributions in finite volume. On the other hand, the absolute deviation of the median from the threshold is, in most cases, smaller than the average splitting $\langle\Delta\rangle$ of the first four eigenvalues of $|Q_m|$ (see table 2; the figures quoted in the fourth column do not include the error on Z_A). In particular, the data are consistent with the working hypothesis on which our argumentation relied (cf. subsect. 4.1).

5. Conclusions

As explained in sect. 2, numerical simulations of the Wilson theory can be expected to be stable if the distribution of the spectral gap of the lattice Dirac operator is well separated from the origin. The range of stability may be defined through the inequality $\bar{\mu} \geq 3\sigma$, for example, where, as before, $\bar{\mu}$ and σ denote the median and width of the distribution. Using the empirical relations $\bar{\mu} \simeq Zm$ and $\sigma \simeq a/\sqrt{V}$, this bound becomes $m \geq 3a/Z\sqrt{V}$, which shows that the accessible range of quark masses depends on both the lattice spacing and the lattice size.

Another form of the stability bound is obtained by multiplication with the ratio $B = m_\pi^2/2m$, which is known to be practically independent of the quark mass [5,6]. For lattices of size $2L \times L^3$, this leads to the inequality

$$m_\pi L \geq \sqrt{3\sqrt{2}aB/Z}. \tag{5.1}$$

In particular, from our numerical studies of the (unimproved) Wilson theory, we deduce the stability bounds

$$m_\pi L \geq \begin{cases} 2.8 & \text{at } a = 0.08 \text{ fm,} \\ 2.3 & \text{at } a = 0.06 \text{ fm,} \end{cases} \tag{5.2}$$

where the mass dependence of B and Z was neglected for simplicity. Note that the pion mass in these formulae is the one in infinite volume, and not the possibly much larger mass computed on a lattice of size L.

It follows from these results that the range of stability includes all lattices where, say, $a \leq 0.1$ fm, $L \geq 2$ fm and $m_\pi L \geq 3$. Simulations of the Wilson theory on such

lattices, using the known simulation algorithms, are thus expected to be safe from the problems mentioned in sect. 2. The O(a)-improved theory is likely to behave in the same way, but so far this has only been checked on a single lattice with spacing $a = 0.09$ fm. If the median and width of the gap distribution are assumed to scale as in the unimproved theory, the stability bound deduced from the simulation of this lattice is $m_\pi L \geq 3.2$. Extensive studies of the improved theory will evidently be required to confirm this result, which is very much in line with the bounds (5.2).

Appendix A. Chebyshev accelerated subspace iteration

On a finite lattice, the few lowest eigenvalues of $A = Q_{\mathrm{m}}^2$ can be computed numerically by minimizing the associated Ritz functional, for example [24,25]. This method is relatively tolerant of rounding errors, which is an important advantage on computers that do not support double-precision arithmetic. Otherwise there is a choice of algorithms that can be significantly faster. Subspace iteration with Chebyshev acceleration and eigenvector locking is one of them, and it is our aim, in the following paragraphs, to describe this method in some detail (see ref. [26], for example).

A.1 Power method

A tight lower bound on the largest eigenvalue of A can be obtained by repeatedly applying the operator to a random quark field. Starting from some arbitrary field ψ with unit norm, the recursion

$$\chi = A\psi, \qquad \psi = \chi/\|\chi\|, \tag{A.1}$$

systematically enhances the upper spectral components of the field. The norm $\|\chi\|$ then provides an increasingly accurate estimate of the largest eigenvalue. In practice 20 iterations or so are usually sufficient for an accuracy better than 5%.

When A is replaced by the shifted operator $c - A$, where $c > \frac{1}{2}\|A\|$ is some fixed number, the power iteration converges to c minus the lowest eigenvalue of A. The
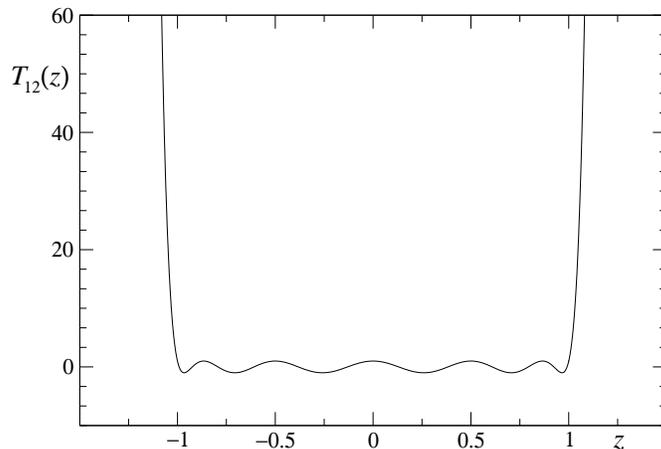
16

Fig. 7. Plot of the Chebyshev polynomial of order 12, showing the bounded oscil-
latory behaviour in the interval $[-1, 1]$ and the rapid increase of the polynomial away
from this range.

latter can thus be calculated in this way. However, since the low eigenvalues of A
are typically orders of magnitude smaller than c, accurate results are obtained only
after a very large number of iterations.

The power method is therefore not recommended for the computation of the low
eigenvalues of A, unless it is combined with an acceleration technique. In particular,
the shifted operator can be replaced by a Chebyshev polynomial and the iteration
may be extended to a subspace of quark fields. Both of these modifications lead to
significantly improved convergence rates.

A.2 Chebyshev polynomials

For $|z| \leq 1$ the Chebyshev polynomials $T_0(z), T_1(z), \dots$ are defined by

$$T_k(z) = \cos(k\theta), \qquad z = \cos\theta. \tag{A.2}$$

They oscillate between -1 and $+1$ in this range and rapidly increase or decrease to
$\pm\infty$ when $|z| > 1$, depending on whether k is even or odd (see fig. 7). Through the
linear transformation

$$z = \frac{2x - v - u}{v - u}, \tag{A.3}$$

the polynomials can be made to oscillate in an arbitrary interval $x \in [u, v]$ instead
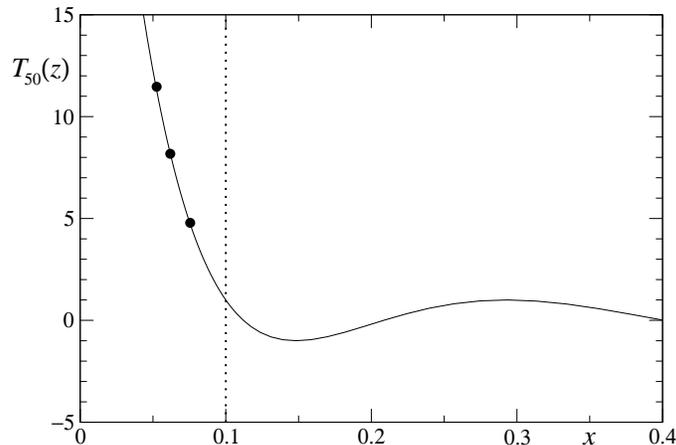of the standard range $[-1, 1]$.

17

Fig. 8. Plot of the Chebyshev polynomial of order 50, scaled to the interval $[u, v] = [0.1, 49]$. If the spectrum of A is contained in this range except for the few lowest eigenvalues (filled circles), the application of the polynomial to a given quark field enhances the field components along the subspace spanned by these modes.

We may now replace x by A and choose $[u, v]$ to contain the unwanted part of the spectrum of A. The other part is then strongly enhanced when the polynomial is applied to a given quark field (see fig. 8). Since the spectrum of the operator is not known beforehand, the interval bounds and the degree of the polynomial must be chosen adaptively in the course of the power iteration.

A simple mathematical fact that will be used in this context is summarized by

Lemma A.1. *For any given even degree k, and real numbers x, v, γ satisfying $x < v$ and $\gamma > 1$, there exists a unique value of the interval bound u such that $x < u < v$ and $T_k(z) = \gamma$ [where z is as in eq. (A.3)].*

Proof: Since γ is larger than 1, there is one and only one ω such that

$$\gamma = \cosh \omega, \qquad \omega > 0. \tag{A.4}$$

The inequality $x < u$, on the other hand, implies $z < -1$ and thus

$$T_k(z) = \cosh(k\nu), \qquad z = -\cosh \nu, \tag{A.5}$$

for some $\nu > 0$. We then conclude that the equation $T_k(z) = \gamma$ has a unique solution (respecting the specified constraints), which is given by $\nu = \omega/k$.

18

Having computed z as a function of k and γ, the interval bound

$$u = x + (v - x) \tanh^2 \left(\frac{\omega}{2k} \right) \tag{A.6}$$

is obtained from eq. (A.3). This shows that u is uniquely calculable, and it is now also easy to verify that the expression (A.6) has all the required properties.

A.3 Chebyshev accelerated subspace iteration

The algorithm described in this subsection computes increasingly accurate approximations to the n lowest eigenvalues of A and the associated eigenvectors. It operates on a d-dimensional subspace of quark fields, where d is usually taken to be quite a bit larger than n.

The computation proceeds iteratively, starting from a random set of orthonormal quark fields $\psi_1, \ldots, \psi_d$. In each cycle of the iteration, the fields are updated one by one so that at the end of the cycle the whole basis of fields is replaced by a new one. More precisely, a cycle consists of the following four steps:

1. Choose a Chebyshev polynomial $T_k(z)$ with even degree and an appropriate spectral range $[u, v]$, taking the current estimates $\alpha_1, \ldots, \alpha_d$ of the eigenvalues as input.

2. Update the basis vectors ψ_i, $i = 1, \ldots, d$, one after another through

$$\chi = T_k(z)\psi_i, \qquad z = \frac{2A - v - u}{v - u}, \tag{A.7}$$

$$\phi = \chi - \sum_{j=1}^{i-1} (\psi_j, \chi)\psi_j, \tag{A.8}$$

$$\psi_i = \phi / \|\phi\|. \tag{A.9}$$

The second and third equations here simply implement the modified Gram–Schmidt orthogonalization process. In particular, the new basis is guaranteed to be orthonormal.

3. Rotate the fields among themselves so as to diagonalize the operator A in the subspace spanned by them, i.e. so that

$$(\psi_i, A\psi_j) = \delta_{ij}\alpha_i \quad \text{for all} \quad i, j = 1, \ldots d, \tag{A.10}$$

$$\alpha_1 \leq \alpha_2 \leq \ldots \leq \alpha_d, \tag{A.11}$$

after the transformation.

4. Stop the algorithm if the approximate eigenvalues $\alpha_1, \ldots, \alpha_n$ and the associated eigenvectors satisfy the chosen convergence criterion.

This description is somewhat schematic and needs to be made more precise. It may not be obvious, for example, how to choose the polynomial in the first step, and there are various stopping criteria that may be applied.

A.4 Choice of the Chebyshev polynomial

Ideally the interval $[u, v]$ should contain the spectrum of A except for the d lowest eigenvalues. In particular, the upper limit v should be set to a value slightly larger than $\|A\|$. If this number is not already known, it may be calculated at the beginning of the subspace iteration using the ordinary power method.

To fix the lower bound u of the spectral interval and the degree k of the polynomial, a reasonable requirement is that

$$T_k(z)|_{x=\alpha_1} = \gamma^2, \qquad T_k(z)|_{x=\alpha_d} = \gamma, \tag{A.12}$$

where $\gamma > 1$ is the desired enhancement factor for the low modes (see fig. 8). In practice $\gamma = 3$ appears to be a sensible choice, but trying other values of γ may be worth while.

Recalling lemma A.1, it is clear that the conditions (A.12) determine both u and k. Explicitly, if we define ω and $\tilde{\omega}$ through

$$\cosh \omega = \gamma, \qquad \cosh \tilde{\omega} = \gamma^2, \tag{A.13}$$

the conditions become

$$u = \alpha_d + (v - \alpha_d) \tanh^2\left(\frac{\omega}{2k}\right), \tag{A.14}$$

$$\alpha_d + (v - \alpha_d) \tanh^2\left(\frac{\omega}{2k}\right) = \alpha_1 + (v - \alpha_1) \tanh^2\left(\frac{\tilde{\omega}}{2k}\right). \tag{A.15}$$

Normally k is such that $2k \gg \tilde{\omega}$, and the expansion at large k of eq. (A.15) then leads to

$$k = \frac{1}{2}\left\{\frac{(v - \alpha_1)\tilde{\omega}^2 - (v - \alpha_d)\omega^2}{\alpha_d - \alpha_1}\right\}^{1/2}, \tag{A.16}$$

while u is determined by eq. (A.14). Evidently k should be rounded to the closest even integer, and one may also wish to impose lower and upper limits on k at this point.

A.5 Stopping criterion

An upper bound on the deviation of the calculated eigenvalues α_i from the exact eigenvalues of A may be obtained by computing the residues

$$\rho_i = (A - \alpha_i)\psi_i \quad (i = 1, \ldots, n) \tag{A.17}$$

and the maximal eigenvalue ϵ^2 of the $n \times n$ residual matrix

$$R_{ij} = (\rho_i, \rho_j). \tag{A.18}$$

A well-known lemma then asserts that there are n orthonormal eigenvectors of A with eigenvalues $\hat{\alpha}_i$ such that $|\alpha_i - \hat{\alpha}_i| \leq \epsilon$ for all $i = 1, \ldots, n$.

This convergence criterion is safe but can be inefficient if the set $\{\alpha_1, \ldots, \alpha_n\}$ of approximate eigenvalues divides into well separated subsets of one or more eigenvalues. The error bounds obtained from the residual matrices associated to each subset of eigenvalues are then often quite a bit smaller than the bound obtained from the total residual matrix. Provided the subsets are indeed separated from one another, by a margin larger than the combined errors, these tighter bounds are completely safe too.

A.6 Eigenvector locking

A fairly obvious property of the subspace iteration is that the lower eigenvalues converge faster than the higher ones. The algorithm can thus be accelerated somewhat, by locking the subsets of eigenvalues and eigenvectors that have already converged. Locking means that these eigenvectors are not updated in the second step of each subspace iteration cycle and that, in the third step, the operator A is diagonalized in the complementary subspace only.

Another small acceleration is achieved by saving the last few eigenvectors, say $\psi_{d-r+1}, \ldots, \psi_d$, to some auxiliary fields before they are updated in the second step. In the third step the saved fields may then be included in the Ritz diagonalization, i.e. A is diagonalized in a subspace of dimension $d+r$, but only the first d eigenvectors are kept after the diagonalization.

A.7 Rounding errors

In order to avoid large rounding errors when the Chebyshev polynomials are applied to the quark fields, the Clenshaw recursion should be used [27]. Let ψ be a given

21

quark field and let us define

$$\chi_j = T_j(z)\psi, \qquad z = \frac{2A - v - a}{v - u}. \tag{A.19}$$

The computation then proceeds recursively according to

$$\chi_0 = \psi, \quad \chi_1 = z\psi, \tag{A.20}$$

$$\chi_{j+1} = 2z\chi_j - \chi_{j-1}, \qquad j = 1, 2, \ldots, \tag{A.21}$$

until the desired degree k is reached.

The use of 32-bit arithmetic in the subspace iteration does not lead to uncontrolled rounding errors as long as the degree k of the Chebyshev polynomial is not too large. In general the significance loss in the Clenshaw recursion grows proportionally to k, and degrees below 100 or 200 may therefore be safe. On large lattices, however, the lowest eigenvalues of A tend to be closely spaced and many orders of magnitude smaller than the maximal eigenvalue. Polynomials with significantly larger degrees will be required under these conditions and the use of double-precision arithmetic then becomes indispensable.

References

[1] K. G. Wilson, Phys. Rev. D10 (1974) 2445

[2] B. Sheikholeslami, R. Wohlert, Nucl. Phys. B259 (1985) 572

[3] M. Lüscher, S. Sint, R. Sommer, P. Weisz, Nucl. Phys. B478 (1996) 365

[4] M. Lüscher, Comput. Phys. Commun. 165 (2005) 199

[5] M. Lüscher, Lattice QCD with light Wilson quarks, in: Proceedings of the 23rd International Symposium on Lattice Field Theory, Dublin, 2005, eds. A. Irving, C. McNeile, C. Michael, PoS(LAT2005)002

[6] L. Del Debbio, L. Giusti, M. Lüscher, R. Petronzio, N. Tantalo, in preparation

[7] P. H. Ginsparg, K. G. Wilson, Phys. Rev. D25 (1982) 2649

[8] P. Hasenfratz, Nucl. Phys. B (Proc. Suppl.) 63 (1998) 53; Nucl. Phys. B525 (1998) 401

[9] P. Hasenfratz, V. Laliena, F. Niedermayer, Phys. Lett. B427 (1998) 125

[10] H. Neuberger, Phys. Lett. B417 (1998) 141; *ibid.* B427 (1998) 353

[11] M. Lüscher, Phys. Lett. B428 (1998) 342

[12] F. Niedermayer, Nucl. Phys. (Proc. Suppl.) 73 (1999) 105

[13] S. Duane, A. D. Kennedy, B. J. Pendleton, D. Roweth, Phys. Lett. B195 (1987) 216

[14] K. Jansen, R. Sommer (ALPHA collab.), Nucl. Phys. B530 (1998) 185 [E: *ibid.* B643 (2002) 517]

[15] R. Sommer, Nucl. Phys. B411 (1994) 839

[16] J. L. van Hemmen, J. Phys. A: Math. Gen. 15 (1982) 3891

[17] T. Banks, A. Casher, Nucl. Phys. B169 (1980) 103

[18] C. W. Bernard, M. F. L. Golterman, Phys. Rev. D49 (1994) 486

[19] S. R. Sharpe, N. Shoresh, Phys. Rev. D64 (2001) 114510

[20] J. C. Osborn, D. Toublan, J. J. M. Verbaarschot, Nucl. Phys. B540 (1999) 317

[21] P. H. Damgaard, J. C. Osborn, D. Toublan, J. J. M. Verbaarschot, Nucl. Phys. B547 (1999) 305

[22] D. Bećirević et al. (SPQ$_{CD}$R collab.), Non-perturbatively renormalised light quark masses from a lattice simulation with $N_f = 2$, hep-lat/0510014

[23] M. Della Morte et al. (ALPHA collab.), JHEP 0507 (2005) 007

[24] B. Bunk, K. Jansen, M. Lüscher, H. Simma, Conjugate gradient algorithm to compute the low-lying eigenvalues of the Dirac operator in lattice QCD, ALPHA collaboration notes (September 1994)

[25] T. Kalkreuter, H. Simma, Comput. Phys. Commun. 93 (1996) 33

[26] Y. Saad, Numerical methods for large eigenvalue problems (Manchester University Press, Manchester, 1992); this book is out of print but can be downloaded from `http://www-users.cs.umn.edu/~saad/books.html`

[27] W. H. Press, S. A. Teukolsky, W. T. Vetterling, B. P. Flannery, Numerical recipes in FORTRAN, 2nd ed. (Cambridge University Press, Cambridge, 1992)