

## Local coherence and deflation of the low quark modes in lattice QCD

---

**Martin Lüscher**

*CERN, Physics Department, TH Division,  
CH-1211 Geneva 23, Switzerland  
E-mail: [luscher@mail.cern.ch](mailto:luscher@mail.cern.ch)*

**ABSTRACT:** The spontaneous breaking of chiral symmetry in QCD is known to be linked to a non-zero density of eigenvalues of the massless Dirac operator near the origin. Numerical studies of two-flavour QCD now suggest that the low quark modes are locally coherent to a certain extent. As a consequence, the modes can be simultaneously deflated, using local projectors, with a total computational effort proportional to the lattice volume (rather than its square). Deflation has potentially many uses in lattice QCD. The technique is here worked out for the case of quark propagator calculations, where large speed-up factors and a flat scaling behaviour with respect to the quark mass are achieved.

**KEYWORDS:** Lattice QCD, Lattice Gauge Field Theories.

JHEP07(2007)081

---

## Contents

<b>1. Introduction</b>	<b>2</b>
<b>2. Spectral density and the <math>V^2</math>-problem</b>	<b>3</b>
2.1 Lattice parameters & field ensembles	3
2.2 Computation of the spectral density	3
2.3 Comparison with the Banks-Casher formula	4
<b>3. Inexact deflation</b>	<b>5</b>
3.1 Oblique projector algebra	5
3.2 Deflation of the Dirac equation	6
3.3 Deflation efficiency	7
<b>4. Domain-decomposed subspaces</b>	<b>7</b>
4.1 Block projection method	7
4.2 Building domain-decomposed subspaces from global fields	8
4.3 Deflation of the free-quark theory	8
<b>5. Local coherence and subspace generation</b>	<b>9</b>
5.1 Smoothness & local coherence	9
5.2 Numerical experiments	10
5.3 Subspace generation	10
5.4 Choice of parameters	11
<b>6. Deflation-accelerated solver for the Dirac equation</b>	<b>11</b>
6.1 Preconditioned Dirac equation	11
6.2 Krylov space solver and the deflation-relaxation interplay	12
6.3 Performance tests	12
6.4 Miscellaneous remarks	14
<b>7. Concluding remarks</b>	<b>15</b>
<b>A. Solution of the little Dirac equation</b>	<b>16</b>
A.1 Computation of the little Dirac operator	16
A.2 Even-odd preconditioning	16
A.3 Global-mode deflation	17
A.4 Solver performance	17
A.5 Using adapted precision	18

---

## 1. Introduction

The physical masses of the up and down quarks are much smaller than the typical low-energy hadronic scales such as the pion decay constant and the string tension. In numerical lattice QCD, the smallness of the quark masses still is a source of difficulty, for various reasons, but mainly because the available simulation techniques become inefficient close to the chiral limit.

It is not excluded, however, that many of the present limitations in lattice QCD can be overcome by “deflating QCD”, i.e. by treating the eigenmodes of the Dirac operator with small eigenvalues separately from the bulk of the quark modes. Deflation techniques are used in many areas of applied science and they are also an active research topic in numerical mathematics (see refs. [1, 2], for example, and references quoted there). In lattice QCD low-mode deflation was so far mainly used in connection with statistical error reduction methods [3]–[7] that now go under the headings of *low-mode averaging* and *all-to-all propagators*. Other applications of deflation methods in QCD include quark propagator computations in special situations, where only a small number of modes need to be deflated [8–10].

In the large-volume regime of QCD, the low-mode deflation methods proposed to date however tend to become useless in practice, because the number of eigenvalues of the Dirac operator below any fixed value, say 100 MeV, grows proportionally to the four-dimensional volume  $V$  of the lattice. The computational effort required for the calculation of the low quark modes and the deflation operations scales like  $V^2$  in this situation (or even a higher power of  $V$ ) and eventually offsets the benefits of low-mode deflation. As Banks and Casher [11] noted long ago, the average spectral density of the low quark modes is proportional to the quark condensate in the chiral limit. The  $V^2$ -problem is thus directly linked to the spontaneous breaking of chiral symmetry and is therefore present independently of the chosen lattice formulation of the theory.

At present little appears to be known about the space-time structure of the low quark modes, but a simple numerical inspection, reported in section 5, suggests that they are locally coherent to some extent. This property allows highly effective deflation subspaces to be built from only a few low modes, using block projectors. The numerical effort required for the preparation of the subspace and the deflation of the Dirac operator is then only of order  $V$  (rather than  $V^2$ ).

Before going into the details of the construction in sects. 4 and 5, the practical relevance of the  $V^2$ -problem is briefly discussed in section 2 and it is explained, in section 3, how to deflate the Dirac operator if the deflation subspace is not spanned by exact eigenmodes of the operator. The potential of the proposed deflation method is demonstrated in section 6, where a preconditioned solver for the lattice Dirac equation is described, whose efficiency decreases only slightly with the quark mass and which outperforms any solver previously used in lattice QCD by a large factor.

## 2. Spectral density and the $V^2$ -problem

### 2.1 Lattice parameters & field ensembles

All simulation results reported in this paper were obtained using the  $O(a)$ -improved Wilson formulation of lattice QCD [12, 13] with two flavours of mass-degenerate sea quarks. Only two lattices, of size  $48 \times 24^3$  and  $64 \times 32^3$ , were considered, both at the same inverse gauge coupling  $\beta = 5.3$ , sea-quark hopping parameter  $\kappa_{\text{sea}} = 0.13625$  and value  $c_{\text{sw}} = 1.90952$  [14] of the Sheikholeslami-Wohlert improvement term. At this point in parameter space, the lattice spacing  $a$  in physical units is estimated to be  $0.0784(10)$  fm [15], while the sea-quark mass is roughly equal to a quarter of the physical strange-quark mass  $m_s$ .


Representative ensembles of gauge-field configurations on these two lattices were generated by the authors of ref. [15] and were made available for the studies conducted here. The ensembles consist of 169 and 50 configurations, widely separated in simulation time so that the residual autocorrelations can, in most cases, be expected to be negligible. However, the discussion that follows is intended to be largely qualitative and the quoted errors and any systematic uncertainties will therefore be generously ignored.

### 2.2 Computation of the spectral density

In the Wilson theory, the spectrum of the (massive) lattice Dirac operator  $D$  is supported in an elliptic region in the complex plane and is thus not easily compared with the spectrum of the Dirac operator in the continuum theory and the Banks-Casher formula. This difficulty can be bypassed by considering the hermitian operator  $D^\dagger D$  instead of  $D$ , a choice which has other advantages as well. The computation of the low-lying eigenvalues of the operator, for example, becomes relatively straightforward. In this paper all eigenvalue and eigenmode calculations were performed using Chebyshev-accelerated subspace iterations (see appendix A of ref. [16]).

The spectral density of  $(D^\dagger D)^{1/2}$ , averaged over the ensemble of gauge-field configurations on the  $48 \times 24^3$  lattice, is shown in figure 1. Perhaps the most interesting feature of this distribution is that it is practically constant above the threshold region at the low end of the spectrum. The threshold of the density in infinite volume is, incidentally, expected to be at  $Z_A m_{\text{sea}}$  [16], where  $Z_A$  and  $m_{\text{sea}}$  denote the axial current renormalization constant and the bare current-quark mass of the sea quark ( $Z_A = 0.75(1)$  [17] and  $am_{\text{sea}} = 0.00761(7)$  [15] on the lattices considered here). As can be seen from the figure, this value appears to give a good indication on where the bulk of the spectrum in finite volume begins.

As discussed in ref. [16], the spectral density of  $(D^\dagger D)^{1/2}$  renormalizes multiplicatively, the renormalization factor  $Z_P$  being the same as the one of the pseudo-scalar density. For the specified lattice parameters, the conversion factor from the lattice to the  $\overline{\text{MS}}$  scheme of dimensional regularization at renormalization scale  $\mu = 2$  GeV was recently determined to be  $Z_P^{-1} = 1.84(3)$  [20]. The range of eigenvalues in figure 1 thus extends up to about 121 MeV after conversion to the  $\overline{\text{MS}}$  scheme, i.e. to a value approximately 25% larger than the physical mass of the strange quark [18]–[20].



**Figure 1:** Unrenormalized density  $\rho(\lambda)$  of the eigenvalues  $\lambda$  of  $(D^\dagger D)^{1/2}$  on the  $48 \times 24^3$  lattice, in units of ‘number of eigenvalues per MeV and  $\text{fm}^4$ ’. The lattice parameters are as specified in subsection 2.1 and the dotted vertical line indicates the theoretically expected position of the threshold of the density in infinite volume [16].

The spectral density on the  $64 \times 32^3$  lattice was also computed and turned out to be nearly the same as the one on the  $48 \times 24^3$  lattice. In particular, the average number of eigenmodes in the range covered by figure 1 increases from 29 on the smaller lattice to about 89 on the big lattice, which shows that the  $V^2$ -problem is not an academic one. The computation of the 32 lowest eigenvalues and associated eigenmodes of  $D^\dagger D$  on the  $48 \times 24^3$  lattice, for example, to a relative precision of  $10^{-3}$ , is in fact already a heavy task that requires the Dirac operator to be applied some  $2.5 \times 10^5$  times.

### 2.3 Comparison with the Banks-Casher formula

According to the Banks-Casher relation [11], the average number  $n(M)$  of eigenvalues of the massless Dirac operator of magnitude less than  $M$  is, in the continuum theory, given by

$$n(M) = \frac{2}{\pi} M \Sigma V + O(M^2), \tag{2.1}$$

where  $\Sigma$  denotes the  $u$ -quark condensate in the thermodynamic limit. This formula holds in any renormalization scheme, but  $\Sigma$  must refer to a definite normalization prescription. A recently quoted result in two-flavour QCD for the condensate in the  $\overline{\text{MS}}$  scheme is  $\Sigma = (251 \pm 13 \text{ MeV})^3$  [21]. Setting  $M = 100 \text{ MeV}$  for illustration, and assuming a  $2L \times L^3$  lattice, eq. (2.1) then yields the estimates  $n(M) = 21, 108$  and  $342$  for the average number of quark modes below  $M$  at  $L = 2, 3$  and  $4 \text{ fm}$ .

These figures are in a similar range as the numerically determined mode numbers reported in the previous subsection. A quantitative comparison must however take into account the exact physical sizes of the simulated lattices and the fact that the lattice Dirac

operator  $D$  includes the quark mass term. Inserting again the value of  $\Sigma$  quoted above and converting the lattice sizes to physical units (using  $a = 0.0784$  fm), the Banks-Casher formula then predicts the average number of eigenvalues in the range covered by figure 1 to be 20 and 63, respectively, on the  $48 \times 24^3$  and the  $64 \times 32^3$  lattice. These values are smaller than the actual numbers (29 and 89) of low modes on these lattices, but they are in the same ballpark and one should also not forget that there are systematic uncertainties in these calculations.

### 3. Inexact deflation

It should be quite clear at this point that good deflation methods in QCD should not assume the low eigenmodes of the Dirac operator to be accurately known. Eventually the only requirements are that the method is efficient and that the correctness of the final results is guaranteed. Inexact deflation was already discussed in ref. [10], for example, and will be driven to the extreme in this paper, partly following recent developments in the mathematical literature [1, 2].

#### 3.1 Oblique projector algebra

Deflation methods in QCD usually start from a set of quark fields,  $\phi_1(x), \dots, \phi_N(x)$ , which will here be assumed to be orthonormal but are otherwise left unspecified.<sup>1</sup> The orthogonal projector  $P$  to the space  $\mathcal{S}$  spanned by these fields (the *deflation subspace*) acts on a given quark field  $\psi(x)$  according to

$$P\psi(x) = \sum_{k=1}^N \phi_k(x) (\phi_k, \psi), \tag{3.1}$$

where the bracket  $(\chi, \psi)$  denotes the obvious scalar product in the linear space of all quark fields.

The restriction of the lattice Dirac operator  $D$  to the deflation subspace is referred to as the *little Dirac operator*. It is completely specified by the matrix

$$A_{kl} = (\phi_k, D\phi_l), \quad k, l = 1, \dots, N, \tag{3.2}$$

that represents its action on the basis fields. In the following, the little Dirac operator is assumed to be invertible, a requirement that will always be satisfied in practice. The linear operators

$$P_L\psi(x) = \psi(x) - \sum_{k,l=1}^N D\phi_k(x)(A^{-1})_{kl}(\phi_l, \psi), \tag{3.3}$$

$$P_R\psi(x) = \psi(x) - \sum_{k,l=1}^N \phi_k(x)(A^{-1})_{kl}(\phi_l, D\psi), \tag{3.4}$$

---

<sup>1</sup>The term *quark field* is reserved for lattice Dirac fields that carry a colour but no flavour index. The eigenmodes of  $D^\dagger D$  with small eigenvalues are referred to as the *low quark modes* or, somewhat abusively, as the *low modes of the Dirac operator*.

can then be defined, where the subscripts stand for “left” and “right” because  $P_L$  and  $P_R$  usually appear on the left and right of the Dirac operator. These operators are oblique projectors, i.e. they are not hermitian but satisfy

$$P_L^2 = P_L, \quad P_R^2 = P_R. \quad (3.5)$$

Other algebraic identities that follow directly from the definitions (3.1)–(3.4) are

$$P_L D = D P_R, \quad (3.6)$$

$$P P_L = P_R P = 0, \quad (3.7)$$

$$P_L(1 - P) = (1 - P)P_R = 1 - P. \quad (3.8)$$

In particular,  $P_L$  projects to the orthogonal complement of the deflation subspace.

### 3.2 Deflation of the Dirac equation

The inhomogeneous Dirac equation,

$$D\psi(x) = \eta(x), \quad (3.9)$$

may now be split into two independent equations by acting with the projectors  $P_L$  and  $1 - P_L$  from the left. The second equation can be solved immediately and the solution of the full system is then given by

$$\psi(x) = \chi(x) + \sum_{k,l=1}^N \phi_k(x)(A^{-1})_{kl}(\phi_l, \eta), \quad (3.10)$$

where  $\chi(x)$  must solve the deflated system

$$P_L D \chi(x) = P_L \eta(x) \quad (3.11)$$

subject to the constraint  $(1 - P_R)\chi(x) = 0$ . In view of the commutator property (3.6), this constraint is consistent with the deflated system and can be freely imposed. One may actually solve the deflated equation (3.11) without imposing any constraint and simply apply  $P_R$  to the calculated solution at the end of the computation.

The full quark propagator  $S(x, y)$  can be similarly split into two parts,

$$S(x, y) = P_R S(x, y) + \sum_{k,l=1}^N \phi_k(x)(A^{-1})_{kl} \phi_l(y)^\dagger, \quad (3.12)$$

the second term being the contribution along the deflation subspace while the first coincides with the Green function of the deflated system (3.11). In practice eq. (3.12) may be a starting point for the application of variance reduction methods such as those described in refs. [4, 7].

### 3.3 Deflation efficiency

Some insight into why deflation is potentially beneficial is obtained by noting that the deflated operator

$$\hat{D} = P_L D = P_L D(1 - P) \tag{3.13}$$

acts in the orthogonal complement  $\mathcal{S}^\perp$  of the deflation subspace. Moreover, a little algebra shows that  $\hat{D}$  is the Schur complement of  $D$  with respect to  $\mathcal{S}$  and that its inverse in  $\mathcal{S}^\perp$  is given by

$$\hat{D}^{-1} = (1 - P)D^{-1}(1 - P). \tag{3.14}$$

The condition number of the deflated system (3.11) is thus expected to be significantly smaller than the condition number of the full system if the low modes of the Dirac operator are sufficiently suppressed by the projector  $1 - P$ .

For any given normalized quark field  $\psi(x)$ , the deficit

$$\epsilon = \|(1 - P)\psi\|^2 \tag{3.15}$$

provides a practical measure of how well the field is approximated by the deflation subspace. Useful subspaces will have to be such that all low quark modes (in, say, the range considered in section 2) have small deficits  $\epsilon$ . However, contrary to what may be presumed, the construction of such subspaces does not require the low modes to be computed to any accuracy (see section 5).

## 4. Domain-decomposed subspaces

The deflation subspaces considered in the following are based on a division of the lattice into non-overlapping rectangular blocks of lattice points. Domain decompositions of this kind were previously introduced for the Schwarz preconditioning of the Dirac operator and the HMC algorithm [22, 23], but the subspaces constructed in this paper are not linked to the Schwarz preconditioning and can be used in many different ways.

### 4.1 Block projection method

Once the lattice is divided into blocks, local deflation subspaces may be defined by specifying  $N_s$  orthonormal quark fields  $\phi_l^\Lambda(x)$ ,  $l = 1, \dots, N_s$ , on each block  $\Lambda$ . The full deflation subspace is then spanned by the set of all these local subspaces and thus has dimension  $N = N_b N_s$ , where  $N_b$  denotes the number of blocks in the lattice. In particular, at fixed block size, the total number of basis fields scales proportionally to the lattice volume  $V$ .

Subspaces of this kind fit the general framework discussed in the previous section if the basis fields are relabelled by an index  $k$  running from 1 to  $N$ . The little Dirac operator, the deflation projectors and the deflated Dirac operator are thus defined as before. An obvious advantage of the construction is that the application of the projector  $P$  to a given quark field  $\psi(x)$ ,

$$P\psi(x) = \sum_{\Lambda} \sum_{l=1}^{N_s} \phi_l^\Lambda(x) (\phi_l^\Lambda, \psi), \tag{4.1}$$



requires a number of arithmetic operations proportional to the lattice volume times  $N_s$  (rather than  $N$ ). From the point of view of the operations count and the memory requirements, the subspace thus behaves as if it were spanned by only  $N_s$  fields. A notable exception to this rule is the little Dirac operator, which always acts in a space of dimension  $N$ .

### 4.2 Building domain-decomposed subspaces from global fields

In practice the block fields  $\phi_l^\Lambda(x)$ ,  $l = 1, \dots, N_s$ , will be obtained starting from a set of  $N_s$  globally defined quark fields  $\psi_l(x)$ . The procedure is very simple and begins by projecting the input fields to the blocks, i.e. by defining the fields

$$\psi_l^\Lambda(x) = \begin{cases} \psi_l(x) & \text{if } x \in \Lambda, \\ 0 & \text{otherwise,} \end{cases} \quad (4.2)$$

on each block  $\Lambda$ . The Gram-Schmidt orthonormalization process is then applied to these and the orthonormalized fields are taken to be the basis elements  $\phi_l^\Lambda(x)$ .

The subspace generated in this way contains the fields  $\psi_l(x)$ , but since the number of basis fields is multiplied by the number of blocks, the subspace tends to be much larger than the space spanned by the input fields.

### 4.3 Deflation of the free-quark theory

For illustration and in order to motivate the further developments, it is now helpful to briefly consider the case of the free-quark theory. As will become clear below, a good choice of the basis fields in this theory are the constant modes. Since the quark fields carry a Dirac and a colour index, one has  $N_s = 12$  orthonormal constant modes on each block.

If periodic or anti-periodic boundary conditions are imposed, the eigenmodes of the Dirac operator are plane waves of the form


$$\psi_p(x) = u_p e^{ipx}, \quad (4.3)$$

where  $u_p$  is a spinor that depends on the momentum  $p$  but not on the position  $x$ . Assuming an  $L^4$  lattice and a block division into blocks of size  $b^4$  (where  $L$  is an integer multiple of  $b$ ), a straightforward computation then shows that

$$\|(1 - P)\psi_p\|^2 = \epsilon_p \|\psi_p\|^2, \quad \epsilon_p = \frac{1}{12} p^2 (b^2 - a^2) + O(p^4 b^4). \quad (4.4)$$

The projection to the orthogonal complement of the specified deflation subspace thus suppresses the low-momentum modes by a factor proportional to  $p^2$  (see figure 2).

A second and perhaps more important observation is that the deflation efficiency does not depend on the lattice size. Even on very large lattices, all low modes with momenta  $p$  of magnitude up to some fraction of  $1/b$  are deflated with small deficits  $\epsilon_p$ . Figure 2 also illustrates the fact that high deflation efficiencies can be achieved by subspaces of fields that are only piecewise smooth, i.e. fields that are far from being approximate eigenmodes of the Dirac operator.



**Figure 2:** Approximation of a plane wave by a superposition of constant block modes. In the free-quark theory, piecewise constant deflation modes achieve high deflation efficiencies up to momenta  $p$  on the order of the inverse of the block size  $b$ .

## 5. Local coherence and subspace generation

The discussion in the previous section suggests that the  $V^2$ -problem can perhaps be solved using domain-decomposed deflation subspaces. However, no general prescription was given so far of how to choose the fields  $\psi_l(x)$ ,  $l = 1, \dots, N_s$ , from which these subspaces are built (cf. subsection 4.2). Such a prescription will now be developed, based on a property of the low quark modes referred to as local coherence.

### 5.1 Smoothness & local coherence

In the free-quark theory, the block projection method works out because the low-momentum modes are smooth on the scale of the block size  $b$ . The intuitive picture that goes along with this explanation is rather appealing but may be difficult to carry over to the full theory. In particular, the notion of smoothness ceases to have a well-defined meaning in presence of a non-trivial lattice gauge field.

A related concept which is better adapted to the situation in the full theory is local coherence. Loosely speaking, a set of quark fields is referred to as locally coherent if the fields are locally well approximated by a relatively small number of fields. When projected to the blocks of a block lattice, for example, such fields are contained in small subspaces of block fields, up to small deficits that depend on the block size and the dimension of the local subspaces.

It is quite clear that the block projection method can only work out if the low quark modes are locally coherent in this sense. Whether this is so appears to be difficult to tell on the basis of simple reasoning alone. The free-quark theory certainly provides little guidance at this point, because the physics of the low modes is completely different from the one in the full theory.

## 5.2 Numerical experiments

Local coherence is a property that can be investigated numerically in a straightforward manner. One begins with an accurate computation of the low-lying eigenvalues and associated eigenmodes of  $D^\dagger D$  and constructs a domain-decomposed subspace from an arbitrary subset of the calculated modes, following the lines of subsection 4.2. The question is then whether all other low modes are also contained in this subspace, up to small deficits  $\epsilon$ .

Several numerical experiments of this kind were performed in two-flavour QCD on the lattices specified in subsection 2.1. The results are quite impressive and unambiguously show that the low modes in this theory are locally coherent to a high degree. Moreover, the property appears to hold for every individual gauge-field configuration and not just on average.

If the  $64 \times 32^3$  lattice is divided into blocks of size  $4^4$ , for example, and if 12 eigenmodes out of 48 are selected for the construction of the domain-decomposed subspace, the remaining 36 modes turn out to lie in the subspace up to deficits  $\epsilon$  ranging from 0.03 to 0.06. The deficits increase with the block size, but become smaller if more modes are used for the subspace construction. On the  $48 \times 24^3$  lattice the situation is practically the same, i.e. similar deficits are obtained for a given block size and subspace dimension.

## 5.3 Subspace generation

As explained in subsection 4.2, the deflation subspaces constructed in this paper are obtained by restricting a set of quark fields  $\psi_l(x)$ ,  $l = 1, \dots, N_s$ , to the blocks of a block division of the lattice. The fields could be taken to be low eigenmodes of the Dirac operator, but it is far more economical to generate them by a relaxation process, starting from a set of random fields.

A relaxation method that can be used in this context is inverse iteration, where the fields are updated a number of times according to

$$\psi_l(x) \rightarrow \text{“}D^{-1}\text{”}\psi_l(x), \quad l = 1, \dots, N_s. \quad (5.1)$$

The inverse of the Dirac operator is put in quotes in this formula, because an accurate solution of the Dirac equation is not required. The application of a few cycles of the Schwarz alternating procedure [22], for example, actually already has the desired relaxation effect. Moreover, the procedure can be bootstrapped by using the current set of fields to deflate the Dirac equation and thus to accelerate the approximate solution of the equation in the next step (see section 6).

Inverse iteration rapidly depletes the components of the fields parallel to the high modes of the Dirac operator. After a few cycles, the fields then satisfy the bound

$$\|D\psi_l\| \leq M\|\psi_l\|, \quad l = 1, \dots, N_s. \quad (5.2)$$

for some value of  $M$  in the range of the low eigenvalues of  $(D^\dagger D)^{1/2}$ .

An important remark is now that such fields are, to a good approximation, linear combinations of the low quark modes. They are therefore locally coherent with these and consequently generate domain-decomposed subspaces that approximate the low modes up

to small deficits. Some further experimenting actually confirms this and also shows that the deflation efficiencies are not very different from those achieved by domain-decomposed subspaces built from exact low modes.

### 5.4 Choice of parameters

The deflation efficiency of the subspaces generated in this way depends on the block size, the dimension  $N_s$  of the local subspaces and on the number and quality of inverse iteration steps that were applied. Choosing small blocks and large numbers  $N_s$  of fields results in high deflation efficiencies but tends to increase the condition number of the little Dirac operator and thus the computer time required for the application of the oblique projectors  $P_L$  and  $P_R$ . Similarly, the beneficial effects of high numbers of fields and inverse iteration steps must be balanced against the effort spent for the subspace generation.

On the lattices specified in subsection 2.1, choosing blocks of size  $4^4$  and setting  $N_s = 20$  turns out to be a good compromise. Highly efficient deflation subspaces are obtained in this case if the relaxation procedure is stopped when the bound (5.2) is satisfied for a value of  $M$  in the  $\overline{\text{MS}}$  scheme equal to 100 MeV or so (cf. section 2). This level is reached after 11 inverse iteration steps and requires a computational effort equivalent to about 190 applications of the Dirac operator per field (if slightly less effective deflation subspaces are acceptable, one can do with 8 steps and 130 applications).

In general the parameters will have to be tuned empirically. If a deflated solver program like the one described in the next section is available, the deflation efficiency of a given subspace can be quickly determined by measuring the time required for the solution of the Dirac equation to a specified accuracy. Computations of the low quark modes are then again not required. The inverse iteration steps can, incidentally, be carried out at a valence quark mass different from the sea-quark mass. For reasons of efficiency, it is in fact recommended to set the bare mass in this process to a value close to (or even equal to) the critical mass.

## 6. Deflation-accelerated solver for the Dirac equation

Low-mode deflation is expected to be useful in several areas of lattice QCD, some of which [3]–[10] were already mentioned in section 1. The principal goal in this section is to show, in a concrete case, that the deflation subspaces constructed following the prescriptions given in the previous section are very effective and that they actually do provide a solution to the  $V^2$ -problem.

### 6.1 Preconditioned Dirac equation

Once the deflation subspace is generated, the deflated Dirac equation (3.11) can be solved straightforwardly using any of the well-known Krylov space algorithms (see ref. [24], for example). However, from the point of view of the execution time, such a solver may not perform too well, because the little system

$$\sum_{l=1}^N A_{kl} v_l = w_k, \quad k = 1, \dots, N, \quad (6.1)$$

must be solved, for one source vector  $w = (w_1, \dots, w_N)$ , each time the projector  $P_L$  is applied. As explained in appendix A, there are efficient algorithms to solve the little Dirac equation, but the computational effort remains non-negligible.

A better balance of deflation and other operations can be achieved by right-preconditioning the deflated equation. The solver discussed in the following includes the Schwarz preconditioner  $M_{\text{sap}}$  introduced in ref. [22], but a polynomial preconditioner or a fixed number of GMRES iterations [25, 26] may do just as well. In the case of the Schwarz preconditioner, the preconditioned deflated equation reads

$$P_L D M_{\text{sap}} \phi(x) = P_L \eta(x) \tag{6.2}$$

and the solution of eq. (3.11) is then given by  $\chi(x) = P_R M_{\text{sap}} \phi(x)$ . The important point to note is that the preconditioning reduces the iteration count of the Krylov space algorithm and thus the overhead generated by the deflation projector.

### 6.2 Krylov space solver and the deflation-relaxation interplay

Both the Schwarz preconditioner and the deflation projector involve approximate iterative procedures. The GCR algorithm is a recommended Krylov space solver in this situation, because it allows for inexact preconditioning without compromising the correctness of the solution (see ref. [24] for a general discussion and ref. [22] for a description of the algorithm in the context of lattice QCD).

An interesting feature of the GCR algorithm is that the Krylov space is extended, in each step, in a direction  $\xi(x) = M_{\text{sap}} \rho(x)$  where  $\rho(x)$  denotes the current residue. The latter satisfies  $P_L \rho(x) = \rho(x)$  by construction and is therefore orthogonal to the deflation subspace (cf. section 3). When acting on such a field, the alternating Schwarz procedure (which is basically a relaxation method) tends to be quite effective in producing an approximate solution of the Dirac equation  $D\xi(x) = \rho(x)$ . Low-mode deflation thus has the effect of improving the efficiency of the preconditioner.

Once  $\xi(x)$  is calculated, the minimal residue in the so extended Krylov space is determined by computing  $P_L D\xi(x)$  and by applying an orthogonalization process. There is thus an interplay between deflation and relaxation, where the low-mode and the high-mode components of the residue are reduced in alternation by the deflation projector and the Schwarz preconditioner.

### 6.3 Performance tests

The performance of the complete algorithm (DFL+SAP+GCR for short) was determined on the lattices specified in subsection 2.1, at the values of the (valence) quark mass that correspond to the hopping parameters  $\kappa_{\text{val}}$  listed in table 1. In this range of masses, the bare current-quark mass  $m_{\text{val}}$  decreases from the strange-quark mass  $m_s$  to approximately  $\frac{1}{6}m_s$  [15], where the condition number of the Dirac operator reaches a value of about 1900.

In order for the effects of low-mode deflation to be clearly seen, the performance measurements were extended to the even-odd preconditioned BiCGstab algorithm (EO+BiCGstab) [27, 28] and the Schwarz-preconditioned GCR algorithm without deflation (SAP+GCR) [22]. In all cases, the tests consisted in measuring the solver iteration

Lattice	$\kappa_{\text{val}}$	EO+BiCGstab		SAP+GCR		DFL+SAP+GCR	
		$N_{\text{BiCG}}$	$t$ [sec]	$N_{\text{GCR}}$	$t$ [sec]	$N_{\text{GCR}}$	$t$ [sec]
$48 \times 24^3$	0.13550	314	57	50	35	17	15
	0.13590	492	90	78	56	19	18
	0.13610	684	125	110	78	20	19
	0.13625	954	174	157	118	21	21
	0.13635	1269	231	227	170	22	22
$64 \times 32^3$	0.13550	323	72	52	45	17	20
	0.13590	520	115	83	71	20	23
	0.13610	748	165	120	103	21	26
	0.13625	1125	248	183	171	23	29
	0.13635	1663	366	294	267	25	32


**Table 1:** Average solver iteration numbers  $N_X$  and executing times  $t$ , using 24 and 64 processors, respectively, in the case of the  $48 \times 24^3$  and the  $64 \times 32^3$  lattice.

numbers and the computer time required for the solution of the full Dirac equation (3.9) to a precision where  $\|\eta - D\psi\| \leq 10^{-10}\|\eta\|$ . Timings were taken on a recent PC cluster with Infiniband network, using 12 and 32 double-processor nodes for the tests on the  $48 \times 24^3$  and the  $64 \times 32^3$  lattice respectively. Only highly optimized, parallel programs were used that include machine-specific enhancements such as those mentioned in ref. [22]. Quoted solver iteration numbers and timings are averages over 50 gauge-field configurations.

The algorithmic parameters were set to the same values on the  $48 \times 24^3$  and the  $64 \times 32^3$  lattice. In particular, the deflation subspaces were constructed by applying 11 inverse iteration steps to  $N_s = 20$  random fields and by projecting them to a division of the lattice into blocks of size  $4^4$ . In the case of the Schwarz preconditioner, the block size was taken to be  $8 \times 4^3$  and all other parameters were set to the standard values previously used in refs. [22, 23, 15]. A fairly small value,  $N_{\text{kv}} = 16$ , turned out to be a satisfactory choice for the maximal number  $N_{\text{kv}}$  of Krylov vectors that may be generated before the GCR algorithm is restarted (larger values, up to  $N_{\text{kv}} = 32$ , had to be used in the case of the SAP+GCR solver).

As is evident from the test results quoted in table 1, low-mode deflation significantly reduces both the solver iteration numbers and the time needed to solve the Dirac equation to a specified precision. Particularly impressive is the fact, illustrated in figure 3, that the deflated algorithm has a flat scaling behaviour with respect to the quark mass. Moreover, the solver iteration numbers on the two lattices are nearly the same, which is very much in line with the expectation that the efficiency of the domain-decomposed deflation subspaces is independent of the lattice volume and that they provide a solution to the  $V^2$ -problem.

Contrary to the solver iteration numbers, the timings quoted in last column of table 1 are sensitive to the time required for the application of the deflation projector  $P_L$  and thus to the average time needed for the solution of the little Dirac equation (see appendix A). The application of the projector actually consumed as much as 25% of the total time on



**Figure 3:** Average execution time  $t$  needed for the solution of the lattice Dirac equation on the  $64 \times 32^3$  lattice as a function of the bare valence quark mass  $m_{\text{val}}$  given in units of the lattice spacing  $a$ . The lattice, algorithm and test parameters are as specified in subsections 2.1 and 6.3. Dotted lines are drawn to guide the eye.

the small lattice and up to 30% on the big lattice.

#### 6.4 Miscellaneous remarks

- (i) *Partially quenched QCD.* In the tests reported in the previous subsection, the deflation subspace was generated only once per gauge-field configuration, i.e. the same subspace was used at all values of the valence-quark mass considered.
- (ii) *Deflation overhead.* The average time spent for the subspace generation was 150 and 184 seconds, respectively, on the  $48 \times 24^3$  and the  $64 \times 32^3$  lattice. These figures include the time needed for the computation of the little Dirac operator (3.3 and 3.9 seconds). The computational effort required for the preparatory work thus becomes quickly negligible if several quark propagators are to be computed.
- (iii) *Solver stability.* In the case of the deflated solver, the GCR iteration numbers  $N_{\text{GCR}}$  tend to be very stable. The iteration numbers observed in the tests actually deviated by at most 1 from their average values, except at the smallest quark mass on the  $64 \times 32^3$  lattice, where the maximal value of  $N_{\text{GCR}}$  ever seen was 27.
- (iv) *Acceleration of the HMC algorithm.* The HMC simulation algorithm [29] requires the lattice Dirac equation to be solved at regular intervals along the trajectories in

field space which lead from the current to the next configuration. Whether the use of low-mode deflation is profitable in this case depends on the quark mass and the precision requirements.

On the lattices specified in subsection 2.1, for example, an acceleration is achieved at hopping parameters  $\kappa_{\text{sea}} \geq 0.13625$ , if the relative solver tolerance is set to  $10^{-7}$  or less and if, say, 8 inverse iteration steps are used for the subspace generation. At these fairly small quark masses, the scaling behaviour of the HMC algorithm is then softened by nearly one power in the quark mass.

In practice much larger speed-up factors can conceivably be obtained by updating the deflation subspace along the trajectories in field space rather than generating the subspace from scratch each time the Dirac equation must be solved. Moreover, starting from the exact factorization

$$\det D = \det A \det \hat{D} \quad (6.3)$$

of the quark determinant, the HMC algorithm itself can perhaps be deflated too, in which case further accelerations and an improved stability of the algorithm will presumably be achieved.

## 7. Concluding remarks

An important qualitative result of this paper is the demonstration that the low quark modes can be simultaneously deflated using local subspaces of low dimension. Some further clarification (an analytic proof of the local coherence of the low modes, for example) would certainly be welcome, but the numerical studies conducted so far leave little doubt that the construction does indeed provide a solution to the  $V^2$ -problem.


Variance reduction methods, such as low-mode averaging [4] and all-to-all propagator techniques [6, 7], will probably be able to profit from these developments. The performance of the deflation-accelerated solver for the lattice Dirac equation discussed in the previous section is, in any case, quite impressive, particularly so at the smallest quark masses considered. In many cases the computational effort required for the calculation of hadronic correlation functions is thus significantly reduced.

The possible inclusion of deflation ideas in QCD simulation algorithms is an intriguing perspective, since this may allow simulations close to the physical values of the light-quark masses to be performed with an effort not very much larger than the one required at a sea-quark mass equal to, say, a fourth of the physical strange-quark mass.

## Acknowledgments

I am indebted to Leonardo Giusti and Peter Weisz for critical comments on a first version of the paper. The gauge-field configurations used for the numerical studies reported in this paper were generated by the authors of ref. [15]. All computations were performed on a dedicated PC cluster at CERN and on a CRAY XT3 at the Swiss National Supercomputing





**Figure 4:** Support of the function  $D\phi(x)$  (black points) if  $\phi(x)$  is supported on the grey block in the centre of the figure. In particular, the matrix elements (A.1) vanish unless the blocks  $\Lambda$  and  $\Lambda'$  coincide or are nearest neighbours.

Centre (CSCS). I am grateful to these institutions for providing the required computer resources.

## A. Solution of the little Dirac equation

In practice the dimension  $N$  of the domain-decomposed deflation subspaces introduced in this paper tends to be so large that an exact solution of the little Dirac equation (6.1) is not a viable option. The iterative solver proposed here is based on even-odd preconditioning, global-mode deflation and the GCR algorithm.

### A.1 Computation of the little Dirac operator

The block division of the lattice implies a decomposition of the little Dirac operator into  $N_s \times N_s$  block matrices  $B_{\Lambda\Lambda'}$ , whose matrix elements are given by

$$(B_{\Lambda\Lambda'})_{kl} = (\phi_k^\Lambda, D\phi_l^{\Lambda'}), \quad k, l = 1, \dots, N_s. \quad (\text{A.1})$$

Since the Wilson-Dirac operator has only nearest-neighbour hopping terms, most of these matrices vanish and a moment of thought reveals that the little Dirac operator actually couples nearest-neighbour blocks only (see figure 4).

The computation of the scalar products (A.1) is straightforward and requires a total effort proportional to the lattice volume times  $N_s^2$ . Note, however, that the operations count tends to increase rapidly if lattice Dirac operators with hopping terms extending over two or more links are considered.

### A.2 Even-odd preconditioning

The even-odd preconditioning familiar from the full lattice Dirac operator can also be applied to the little Dirac operator in its block form if there is an even number of blocks in each direction (which is here assumed to be the case). If the so-called symmetric preconditioning is chosen [18], the block matrices representing the preconditioned operator on even

blocks  $\Lambda, \Lambda'$  are given by

$$\hat{B}_{\Lambda\Lambda'} = \delta_{\Lambda\Lambda'} - \sum_{\Omega} (B_{\Lambda\Lambda})^{-1} B_{\Lambda\Omega} (B_{\Omega\Omega})^{-1} B_{\Omega\Lambda'}, \quad (\text{A.2})$$

where the sum extends over the common neighbours  $\Omega$  of the blocks  $\Lambda$  and  $\Lambda'$ .

The matrices (A.2) do not need to be stored in the memory of the computer, because the action of the preconditioned operator on a complex field can be computed in two steps, first hopping from the even to the odd blocks and then back to the even blocks. Some work can however be saved by storing the matrices  $(B_{\Lambda\Lambda})^{-1} B_{\Lambda\Lambda'}$  for all nearest-neighbour pairs  $\Lambda, \Lambda'$  of blocks.

On physically small blocks  $\Lambda$ , the diagonal block matrices  $B_{\Lambda\Lambda}$  tend to be safely invertible, but the program should check this and return to the original system if an ill-conditioned matrix is encountered (this never happened in the tests reported in this paper).

### A.3 Global-mode deflation

As explained in subsection 4.2, the basis fields  $\phi_l^\Lambda(x)$  on the blocks  $\Lambda$  are obtained starting from a set of global fields  $\psi_l(x)$ ,  $l = 1, \dots, N_s$ . The latter span a subspace in the generated deflation subspace which may be used to deflate the little Dirac operator. Actually only the components of the global fields on the even blocks are used to build this “little deflation subspace”, because the little Dirac equation is to be deflated in its even-odd preconditioned form.

The equation is deflated following the general procedures described in section 3. One simply has to replace the full Dirac operator by the even-odd preconditioned little Dirac operator and the quark fields by complex fields with  $N/2$  components. Note that the “little little Dirac operator” is an  $N_s \times N_s$  matrix that can be inverted to machine precision with a negligible effort.

Global-mode deflation is straightforward to implement and tends to reduce the condition number of the little system quite significantly (by about a factor 3 in the cases studied so far).

### A.4 Solver performance

Similarly to the full system, the deflated preconditioned little equation can be solved using the GCR algorithm. Tests of the complete solver were then performed using the same subspaces as in subsection 6.3. In particular, the number of GCR iterations  $N_{\text{GCR}}$  and the time  $t$  needed to solve the little equation to a relative precision of  $10^{-12}$  were determined and are quoted in columns 3 and 4 of table 2.

The dependence of these figures on the valence quark mass and the lattice volume is noticeable, but one can also see that the solver iteration numbers increase only relatively slowly towards the smaller quark masses. In practice all these variations are not too important, because the solution of the little system eventually consumes only a fraction of the time spent for the solution of the full system.

Lattice	$\kappa_{\text{val}}$	$N_{\text{GCR}}$	$t$ [sec]	$\overline{N}_{\text{GCR}}$
$48 \times 24^3$	0.13550	84	0.26	24
	0.13590	105	0.32	30
	0.13610	120	0.37	34
	0.13625	136	0.42	38
	0.13635	150	0.46	42
$64 \times 32^3$	0.13550	94	0.37	27
	0.13590	126	0.49	36
	0.13610	154	0.60	44
	0.13625	188	0.73	54
	0.13635	220	0.85	62

**Table 2:** GCR iteration numbers and time (using 24 and 64 processors, respectively, in the case of the  $48 \times 24^3$  and the  $64 \times 32^3$  lattice) needed for the solution of the little system.

### A.5 Using adapted precision

It is still worth including another improvement, however, which exploits the fact that the outer GCR algorithm (the one that solves the full system) is restarted from time to time, usually when the dimension of the generated Krylov space reaches the specified maximal value. Before each restart, the current residue is recomputed with high precision so that any inaccuracies which may have accumulated during the last cycle do not propagate to the next cycle.

For this reason it is permissible to solve the little Dirac equation to low precision inside the cycles of the outer algorithm. In the tests reported in subsection 6.3, for example, the required relative tolerances were set to  $10^{-6}$  and  $10^{-12}$ , respectively, inside and outside the cycles of the algorithm. The average solver iteration numbers are then practically reduced by a factor 2.

They can actually be reduced even further by adapting the precision as one proceeds from one Krylov vector to the next within a cycle. This is possible because the GCR algorithm operates directly on the minimal residuals in the generated Krylov spaces. Their magnitude decreases monotonically and need to be computed essentially only to a fixed decimal precision. The required precision for the solution of the little system can therefore be reduced in proportion to the norm of the quark fields on which the deflation projector  $P_L$  acts.

Once all these improvements are installed, the average iteration numbers  $\overline{N}_{\text{GCR}}$  required for the solution of the little system in the course of the cycles of the outer algorithm are reduced to the figures quoted in the last column of table 2. At the smallest quark mass on the  $64 \times 32^3$  lattice, for example, the time spent for the solution of the little system sums up to about 6 seconds, i.e. about 19% of the total time needed for the solution of the full system.

## References

- [1] J. Frank and C. Vuik, *On the construction of deflation-based preconditioners*, *SIAM J. Sci. Comput.* **23** (2001) 442.
- [2] R. Nabben and C. Vuik, *A comparison of deflation and the balancing preconditioner*, *SIAM J. Sci. Comput.* **27** (2006) 1742.
- [3] H. Neff, N. Eicker, T. Lippert, J.W. Negele and K. Schilling, *On the low fermionic eigenmode dominance in QCD on the lattice*, *Phys. Rev. D* **64** (2001) 114509 [[hep-lat/0106016](#)].
- [4] L. Giusti, P. Hernández, M. Laine, P. Weisz and H. Wittig, *Low-energy couplings of QCD from current correlators near the chiral limit*, *JHEP* **04** (2004) 013 [[hep-lat/0402002](#)].
- [5] T.A. DeGrand and S. Schaefer, *Improving meson two-point functions in lattice QCD*, *Comput. Phys. Commun.* **159** (2004) 185 [[hep-lat/0401011](#)].
- [6] SESAM collaboration, G.S. Bali et al., *Observation of string breaking in QCD*, *Phys. Rev. D* **71** (2005) 114513 [[hep-lat/0505012](#)].
- [7] J. Foley et al., *Practical all-to-all propagators for lattice QCD*, *Comput. Phys. Commun.* **172** (2005) 145 [[hep-lat/0505023](#)].
- [8] Ph. de Forcrand, *Progress on lattice QCD algorithms*, *Nucl. Phys. B (Proc. Suppl.)* **47** (1996) 228 [[hep-lat/9509082](#)].
- [9] R.B. Morgan and W. Wilcox, *Deflation of eigenvalues for GMRES in lattice QCD*, *Nucl. Phys.* **106** (Proc. Suppl.) (2002) 1067 [[hep-lat/0109009](#)].
- [10] L. Giusti, C. Hoelbling, M. Lüscher and H. Wittig, *Numerical techniques for lattice QCD in the  $\epsilon$ -regime*, *Comput. Phys. Commun.* **153** (2003) 31 [[hep-lat/0212012](#)].
- [11] T. Banks and A. Casher, *Chiral symmetry breaking in confining theories*, *Nucl. Phys. B* **169** (1980) 103.
- [12] B. Sheikholeslami and R. Wohlert, *Improved continuum limit lattice action for QCD with Wilson fermions*, *Nucl. Phys. B* **259** (1985) 572.
- [13] M. Lüscher, S. Sint, R. Sommer and P. Weisz, *Chiral symmetry and  $O(a)$  improvement in lattice QCD*, *Nucl. Phys. B* **478** (1996) 365 [[hep-lat/9605038](#)].
- [14] ALPHA collaboration, K. Jansen and R. Sommer,  *$O(a)$  improvement of lattice QCD with two flavors of Wilson quarks*, *Nucl. Phys. B* **530** (1998) 185 [*Erratum ibid.* **643** (2002) 517] [[hep-lat/9803017](#)].
- [15] L. Del Debbio, L. Giusti, M. Lüscher, R. Petronzio and N. Tantalo, *QCD with light Wilson quarks on fine lattices. I: first experiences and physics results*, *JHEP* **02** (2007) 056 [[hep-lat/0610059](#)]; *QCD with light Wilson quarks on fine lattices. II: DD-HMC simulations and data analysis*, *JHEP* **02** (2007) 082 [[hep-lat/0701009](#)].
- [16] L. Del Debbio, L. Giusti, M. Lüscher, R. Petronzio and N. Tantalo, *Stability of lattice QCD simulations and the thermodynamic limit*, *JHEP* **02** (2006) 011 [[hep-lat/0512021](#)];
- [17] M. Della Morte, R. Hoffmann, F. Knechtli, R. Sommer and U. Wolff, *Non-perturbative renormalization of the axial current with dynamical Wilson fermions*, *JHEP* **07** (2005) 007 [[hep-lat/0505026](#)].
- [18] JLQCD collaboration, S. Aoki et al., *Light hadron spectroscopy with two flavors of  $O(a)$ -improved dynamical quarks*, *Phys. Rev. D* **68** (2003) 054502 [[hep-lat/0212039](#)].

- [19] CP-PACS collaboration, A. Ali Khan et al., *Light hadron spectroscopy with two flavors of dynamical quarks on the lattice*, *Phys. Rev. D* **65** (2002) 054505 [Erratum *ibid.* **67** (2003) 059901] [[hep-lat/0105015](#)].
- [20] ALPHA collaboration, M. Della Morte et al., *Non-perturbative quark mass renormalization in two-flavor QCD*, *Nucl. Phys. B* **729** (2005) 117 [[hep-lat/0507035](#)].
- [21] JLQCD and TWQCD collaboration, H. Fukaya et al., *Two-flavor lattice QCD simulation in the epsilon-regime with exact chiral symmetry*, *Phys. Rev. Lett.* **98** (2007) 172001 [[hep-lat/0702003](#)]; *Two-flavor lattice QCD in the  $\epsilon$ -regime and chiral random matrix theory*, [arXiv:0705.3322](#).
- [22] M. Lüscher, *Solution of the Dirac equation in lattice QCD using a domain decomposition method*, *Comput. Phys. Commun.* **156** (2004) 209 [[hep-lat/0310048](#)].
- [23] M. Lüscher, *Schwarz-preconditioned HMC algorithm for two-flavour lattice QCD*, *Comput. Phys. Commun.* **165** (2005) 199 [[hep-lat/0409106](#)].
- [24] Y. Saad, *Iterative methods for sparse linear systems*, 2<sup>nd</sup> ed., SIAM, Philadelphia (2003); see also <http://www-users.cs.umn.edu/~saad/>.
- [25] H.A. van der Vorst and C. Vuik, *GMRESR: a family of nested GMRES methods*, *Num. Lin. Alg. Appl.* **1** (1994) 369.
- [26] C. Vuik, *New insights in GMRES-like methods with variable preconditioners*, *J. Comput. Appl. Math.* **61** (1995) 189.
- [27] H.A. van der Vorst, *BiCGSTAB: a fast and smoothly converging variant of BiCG for the solution of nonsymmetric linear systems*, *SIAM J. Sci. Stat. Comput.* **13** (1992) 631.
- [28] A. Frommer, V. Hannemann, B. Nockel, T. Lippert and K. Schilling, *Accelerating Wilson fermion matrix inversions by means of the stabilized biconjugate gradient algorithm*, *Int. J. Mod. Phys. C* **5** (1994) 1073 [[hep-lat/9404013](#)].
- [29] S. Duane, A.D. Kennedy, B.J. Pendleton and D. Roweth, *Hybrid Monte Carlo*, *Phys. Lett. B* **195** (1987) 216.