Published by Institute of Physics Publishing for SISSA/ISAS
Received: *April 22, 2003*
Accepted: *May 20, 2003*

Lattice QCD and the Schwarz alternating procedure

Martin Lüscher

CERN, Theory Division
CH-1211 Geneva 23, Switzerland
E-mail: luscher@mail.cern.ch

Abstract: A numerical simulation algorithm for lattice QCD is described, in which the short- and long-distance effects of the sea quarks are treated separately. The algorithm can be regarded, to some extent, as an implementation at the quantum level of the classical Schwarz alternating procedure for the solution of elliptic partial differential equations. No numerical tests are reported here, but theoretical arguments suggest that the algorithm should work well also at small quark masses.

Keywords: Lattice Gauge Field Theories, Lattice QCD, Spontaneous Symmetry Breaking.

JHEP05(2003)052

Contents

1. Introduction

The simulation algorithms for (unquenched) lattice QCD that are currently in use rapidly become inefficient on large lattices and at small quark masses, where the effects of spontaneous chiral symmetry breaking set in. In the case of the HMC [1], the PHMC [2, 3] and the Multiboson [4] algorithms, the principal technical difficulty derives from the fact that the light quark masses have to be scaled proportionally to the square of the pion mass in the chiral limit. The lattice Dirac operator is then increasingly ill-conditioned, which affects all these algorithms in a similar and rather direct way, because they all start from a *global pseudo-fermion representation* of the quark determinant that involves an exact (or an accurate approximate) inversion of the Dirac operator.

In the present paper a simulation algorithm is proposed that exploits the underlying local structure of the theory and that may be expected to scale in a more favourable way in the chiral regime. The general strategy is closely related to the alternating procedure that was invented by the mathematician Schwarz in the 19th century to establish the existence of the solution of the Dirichlet problem

$$\Delta f(x)|_{x\in\Omega} = 0\,, \qquad f(x)|_{x\in\partial\Omega} = g(x)\,, \tag{1.1}$$

on arbitrary bounded domains Ω in the plane [5] (for an introduction to the subject in the context of discretized partial differential equations see ref. [6], for example). Very briefly this method obtains the solution iteratively by dividing Ω into a set of overlapping subdomains and by solving the Dirichlet problem on these, in each step of the iteration, with boundary values determined from the current approximation to the solution.

The design of simulation algorithms for lattice QCD that operate on overlapping blocks of lattice points is non-trivial, however, because the global correctness of the simulation must be guaranteed. In principle the problem can be solved using stochastic acceptance-rejection steps similar to those previously considered by Hasenbusch [7] (see also refs. [8]–[11]). The key question is then whether these correction steps can be implemented so that high acceptance rates are achieved, and a significant part of the present paper is therefore devoted to this issue.

2. General form of the algorithm

In this section the proposed algorithm is described in outline. Most technical details are deferred to the later sections and important improvements (such as preconditioning) are omitted in order to keep the presentation as simple as possible. Some algorithm theory, as summarized in appendix A, is nevertheless required to be able to understand the procedure.

2.1 Preliminaries

Although the algorithm is more generally applicable, only the standard Wilson formulation [12] of lattice QCD will be considered here (optionally including $\mathrm{O}(a)$ improvement [13, 14]) with a doublet of mass-degenerate quarks. The lattice spacing is set to unity for convenience and the SU(3) link variables are denoted by $U(x,\mu)$ as usual.

JHEP05(2003)052

After integration over the quark fields, the probability distribution of the gauge field reads

$$P[U] = \frac{1}{\mathcal{Z}}\,\mathrm{e}^{-S_{\mathrm{G}}}\det Q^2\,, \qquad Q \equiv \gamma_5(D_{\mathrm{w}} + m_0)\,, \tag{2.1}$$

where $\mathcal{Z}$ denotes the partition function, S_{G} the plaquette action, D_{w} the Wilson-Dirac operator and m_0 the bare quark mass.

2.2 Alternating procedure

Following the classical Schwarz procedure, the proposed algorithm visits rectangular blocks of lattice points according to some scheme and updates the link variables residing there. The blocks can have arbitrary sizes in principle, and their position may be chosen randomly, for example, so that all link variables are treated equally. However, as will become clear later, the algorithm is designed to perform particularly well if the blocks are small in physical units, i.e. if their edges are less than about 1 fm long.

On each block Λ that is visited in the course of this process, changes in the gauge field on Λ are proposed that satisfy detailed balance with respect to the distribution

$$P_{\Lambda}[U] = \frac{1}{\mathcal{Z}_{\Lambda}}\,\mathrm{e}^{-S_{\mathrm{G}}}\det(Q_{\Lambda} + Q_{\Lambda^*})^2\,. \tag{2.2}$$

A stochastic acceptance-rejection step then needs to be applied to correct for the difference between this distribution and the exact distribution (2.1). The operators Q_{Λ} and Q_{Λ^*} that appear here coincide with Q, except that they act on Dirac fields defined on the block Λ and its complement Λ^*, respectively, with Dirichlet boundary conditions. At the level of the fermion action, a complete decoupling of the quark fields inside and outside the block is thus achieved. In particular, the proposals for the link variables residing in Λ can be generated locally using a block version of the HMC algorithm, for example.

The stochastic acceptance-rejection step, on the other hand, involves an inversion of the Dirac operator Q on the full lattice. There is a fair amount of choice in the detailed implementation of this step, which can be exploited to reduce the influence on the acceptance probability of the link variables far away from the block. Moreover, the suggested procedure (which is explained in section 5) restricts the pseudo-fermion field that needs to be introduced at this stage to the boundary of the block. At least to some extent, the local character of the algorithm is thus preserved.

2.3 Hierarchical structure

In the form described above the algorithm requires a computational effort per block update that increases roughly proportionally to the lattice volume. By introducing a hierarchy of blocks such as the one shown in figure 1, one may, however, be able to do better than this. Proposed changes of the link variables on the smallest blocks are then obtained as before, and these are taken as proposed configurations on the next larger blocks, and so on, accepting or rejecting them so that detailed balance is satisfied with respect to the associated distributions $P_{\Lambda}[U]$. Finally a *simultaneous* global acceptance-rejection step is applied to the surviving configurations.

JHEP05(2003)052

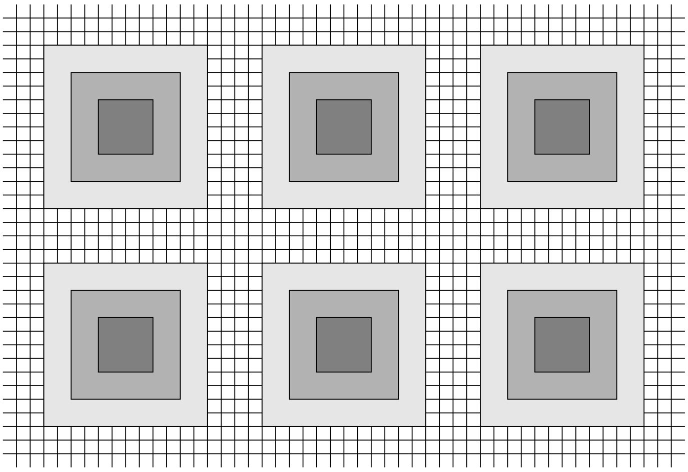

Figure 1: Proposed updates of the link variables on small blocks (dark grey squares) can be filtered through acceptance-rejection steps on increasingly larger blocks (light grey squares) before the global acceptance-rejection step is applied.

This procedure is logically correct since non-overlapping blocks are decoupled from each other and can be updated in parallel. On the other hand, it will only work out in practice if the configurations that have been accepted on the largest blocks have a high probability to be accepted on the whole lattice. For the chosen implementation of the acceptance-rejection steps, a theoretical argument will be given in section 5 that explains why this should be so (under certain conditions).

2.4 Low-mode reweighting

On physically small blocks Λ, the Dirac operator Q_Λ that acts on the quark fields in Λ is not expected to become ill-conditioned at small quark masses. The situation here is actually similar to the one encountered in studies of the Schrödinger functional, where the boundary conditions provide an infrared cutoff on both the gluon and the quark modes [15, 16]. As a consequence the block simulation algorithm should work well even if the bare quark mass m_0 is set to the critical mass m_c [17].

However, since the Wilson formulation of lattice QCD violates chiral symmetry, the Dirac operator on large lattices is not protected against the accidental presence of eigenvectors with eigenvalues much smaller than $m_0 - m_c$. When a gauge field configuration is proposed where this happens, it will only rarely pass the global acceptance-rejection step. Such configurations are therefore sampled with low statistics and this can easily lead to uncontrolled statistical fluctuations if a quantity is considered that is sensitive to the low modes of the Dirac operator (the propagator of the pseudo-scalar density, for example).

It may be possible to solve this problem by including the product

$$W[U] = \prod_{k=1}^{n} \lambda_k \tag{2.3}$$

of the first few eigenvalues $\lambda_1, \ldots, \lambda_n$ of Q^2 in the observables and the inverse factor $W[U]^{-1}$ in the global acceptance-rejection step. The computational overhead for this modification

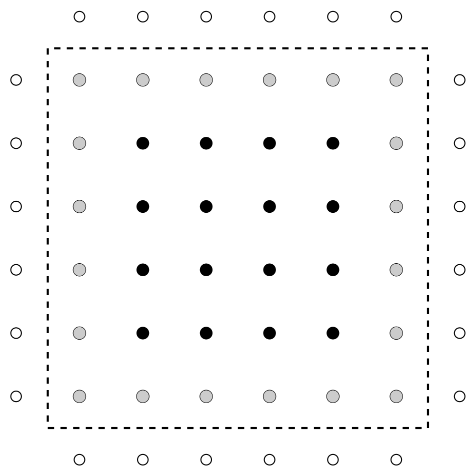

Figure 2: Two-dimensional slice of a 6^4 block (black and grey points inside the dashed square). The grey and the open points represent the interior and the exterior boundary points of the block respectively.

may not be negligible, but it should be noted in this connection that the eigenvalues do not need to be computed very accurately (as long as a definite procedure is used that obtains the eigenvalues independently of the previous gauge field configurations). Moreover the associated approximate eigenvectors can help to accelerate the inversion of the Dirac operator that is required in the acceptance–rejection step [18].

3. Domain decomposition

Let Λ be an arbitrary rectangular block of lattice points. Λ and its complement Λ^* (the set of points not in Λ) define a particular case of a domain decomposition of the lattice. In the following paragraphs the aim is to introduce some basic notation related to this decomposition, but the terminology is quite general and extends to the case where Λ is replaced by the union of a set of non-overlapping blocks.

3.1 Boundary points

On the lattice it is important to distinguish between interior and exterior boundary points (see figure 2). The boundary values for the classical Dirichlet problem on Λ, for example, should be specified on the set $\partial\Lambda$ of all exterior boundary points if the standard nearest-neighbour lattice laplacian is used, while the set of interior boundary points plays an analogous rôle from the point of view of the complementary domain Λ^* and is therefore denoted by $\partial\Lambda^*$.

A special feature of this geometry is that for each exterior boundary point there is a unique closest point on the interior boundary, referred to as its partner point. The converse is evidently not true, i.e. different exterior boundary points may have the same partner point.

JHEP05(2003)052

3.2 Decomposition of the Dirac operator

In position space the hermitean lattice Dirac operator is a sparse matrix that assumes the block-diagonal form

$$Q = \begin{pmatrix} Q_{\Lambda} & Q_{\partial\Lambda} \\ Q_{\partial\Lambda^*} & Q_{\Lambda^*} \end{pmatrix} \tag{3.1}$$

if the lattice points are ordered so that those in Λ come first. The operator Q_{Λ}, for example, acts on Dirac fields on Λ in the same way as Q, except that all terms involving the exterior boundary points are set to zero (which is equivalent to imposing Dirichlet boundary conditions on $\partial\Lambda$).

It is often convenient to consider Q_{Λ}, Q_{Λ^*}, $Q_{\partial\Lambda}$ and $Q_{\partial\Lambda^*}$ to be operators in the space of Dirac fields that are defined on the whole lattice rather than on Λ or Λ^* only. The embedding is done in the obvious way by padding with zeros. Equation (3.1) may then be written in the form

$$Q = Q_{\Lambda} + Q_{\Lambda^*} + Q_{\partial\Lambda} + Q_{\partial\Lambda^*}\,, \tag{3.2}$$

and an equivalent expression for the determinant in eq. (2.2) is $\det Q_{\Lambda}^2 \det Q_{\Lambda^*}^2$. This notation is perhaps slightly abusive but should not lead to any confusion since it is usually clear from the context which interpretation applies.

The operator $Q_{\partial\Lambda}$ and its hermitean conjugate $Q_{\partial\Lambda^*}$ connect the domains Λ and Λ^* to each other. Explicitly $Q_{\partial\Lambda}$ is given by

$$\begin{aligned} Q_{\partial\Lambda}\psi(x) \;=\; -\gamma_5\theta_{\Lambda}(x)\sum_{\mu=0}^{3}\Big\{&\tfrac{1}{2}(1-\gamma_\mu)\theta_{\Lambda^*}(x+\hat{\mu})U(x,\mu)\psi(x+\hat{\mu}) + \\ &+\tfrac{1}{2}(1+\gamma_\mu)\theta_{\Lambda^*}(x-\hat{\mu})U(x-\hat{\mu},\mu)^{-1}\psi(x-\hat{\mu})\Big\}, \end{aligned} \tag{3.3}$$

where θ_{Λ} and θ_{Λ^*} denote the characteristic functions of Λ and Λ^* respectively.[1] This operator thus transports the Dirac spinors on the exterior boundary points to the partner points on the interior boundary, while $Q_{\partial\Lambda^*}$ does the same in the opposite direction.

3.3 Dirichlet problem & space of boundary values

To be able to properly pose the Dirichlet boundary value problem for the lattice Dirac operator on Λ, the space of boundary values first needs to be specified. The situation is practically the same as in the case of the Schrödinger functional which was studied in ref. [16]. It is useful in this connection to introduce the projector-valued function

$$\theta_{\partial\Lambda}(x) = \sum_{\mu=0}^{3}\theta_{\Lambda^*}(x)\left\{\frac{1}{2}(1-\gamma_\mu)\theta_{\Lambda}(x-\hat{\mu}) + \frac{1}{2}(1+\gamma_\mu)\theta_{\Lambda}(x+\hat{\mu})\right\} \tag{3.4}$$

that will, in many respects, play the rôle of the characteristic function of the exterior boundary. The only non-vanishing terms on the right-hand side of eq. (3.4) are in fact

[1] The normalization conventions and all unexplained notations are as in ref. [14].

JHEP05(2003)052

those where x is in $\partial\Lambda$ and $x-\hat{\mu}$ or $x+\hat{\mu}$ its partner point, and the projectors $\frac{1}{2}(1\pm\gamma_\mu)$ are precisely such that the identity

$$Q_{\partial\Lambda} = Q_{\partial\Lambda}\theta_{\partial\Lambda} \tag{3.5}$$

holds.

The Dirichlet boundary value problem for the lattice Dirac operator is now to find a Dirac field $\psi(x)$ on $\Lambda\cup\partial\Lambda$ that satisfies

$$Q\psi(x) = 0 \quad \text{for all } x\in\Lambda\,, \tag{3.6}$$

$$\theta_{\partial\Lambda}(x)\psi(x) = \psi(x) = \eta(x) \quad \text{for all } x\in\partial\Lambda\,, \tag{3.7}$$

where $\eta = \theta_{\partial\Lambda}\eta$ is any prescribed field on $\partial\Lambda$. Recalling the decomposition (3.1) of the Dirac operator, it is clear that Q can be replaced by $Q_\Lambda + Q_{\partial\Lambda}$ in eq. (3.6). At all points in Λ the solution is then given by

$$\psi = -Q_\Lambda^{-1}Q_{\partial\Lambda}\eta\,, \tag{3.8}$$

where use has been made of the fact that $Q_{\partial\Lambda}$ moves the field η from the exterior to the interior boundary of the block. In particular, the solution is uniquely determined by the specified boundary values (if Q_Λ is invertible).

4. Block simulation algorithm

For the generation of proposed gauge field configurations on a given block Λ, the HMC and PHMC algorithms are both equally suitable. Some relevant details are given in this section for the case of the HMC algorithm. The correctness of the local form of this algorithm is easily shown by noting that it matches the general pattern outlined in subsection A.4.

4.1 Active and spectator link variables

As already mentioned in section 2, the proposed updates of the gauge field on Λ must satisfy detailed balance with respect to the distribution (2.2), which is now written in the form

$$P_\Lambda[U] = \frac{1}{\mathcal{Z}_\Lambda}\,\mathrm{e}^{-S_\mathrm{G}}\det Q_\Lambda^2 \det Q_{\Lambda^*}^2\,. \tag{4.1}$$

To achieve a complete decoupling from the surrounding lattice, only the link variables residing on a subset of links in Λ, the *active link variables*, should be changed. In particular, if the links shown in figure 3 are selected, the factor $\det Q_{\Lambda^*}$ is guaranteed to be independent of the active link variables, independently of whether the $\mathrm{O}(a)$-improved or the unimproved theory is considered.

At this point the HMC algorithm can be set up as usual by introducing a pseudo-fermion field $\phi(x)$ on Λ and the canonical momenta $\Pi(x,\mu)$ of the active link variables $U(x,\mu)$. The total action of the system is then taken to be

$$S = S_\mathrm{G} + \frac{1}{2}\,(\Pi,\Pi) + \left(\phi, Q_\Lambda^{-2}\phi\right)\,, \tag{4.2}$$

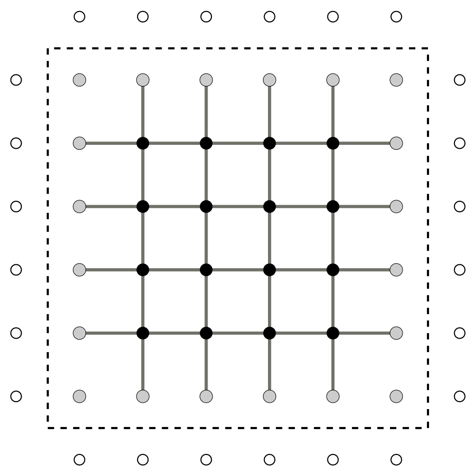

Figure 3: The block update algorithm only changes the link variables on a subset of links in the chosen block (solid lines). In particular, the link variables on the interior boundary and those connecting the interior to the exterior boundary are held fixed.

where the brackets $(\cdot,\cdot)$ denote the natural scalar products in the relevant spaces of fields. Note that the factor $\det Q_{\Lambda^*}$ can be ignored, since it does not depend on the active link variables. All inactive link variables actually play a spectator rôle in the following, which is rather similar to what happens in the case of the widely used simulation algorithms for pure gauge theories, where the link variables are updated one after the other.

4.2 Variable-speed molecular dynamics

Following the standard procedure, a new configuration of the active link variables is now generated by first choosing the pseudo-fermion field and the canonical momenta randomly with conditional probability proportional to e^{-S}. The momenta and the active link variables are then evolved from their initial values Π, U to some final values Π', U' by solving the molecular dynamics equations

$$\frac{\mathrm{d}}{\mathrm{d}\tau}\Pi(x,\mu) = -\gamma(x,\mu)F(x,\mu)\,, \tag{4.3}$$

$$\frac{\mathrm{d}}{\mathrm{d}\tau}U(x,\mu) = \gamma(x,\mu)\Pi(x,\mu)U(x,\mu)\,, \tag{4.4}$$

from $\tau = 0$ to $\tau = 1$. As usual the force $F(x,\mu)$ takes values in the Lie algebra of SU(3) and is determined by the requirement that

$$(\omega, F) = \lim_{\epsilon\to 0}\left\{\frac{1}{\epsilon}\left(S|_{U\to U_\epsilon} - S\right)\right\}, \qquad U_\epsilon(x,\mu) \equiv e^{\epsilon\omega(x,\mu)}U(x,\mu)\,, \tag{4.5}$$

for all variations $\omega(x,\mu)$ of the active link variables. It should be emphasized that the computation of the force is an operation that is local to the block Λ. In particular, only the Dirac operator Q_Λ needs to be inverted.

The action and the integration measure are both preserved in the course of the molecular dynamics evolution, for any choice of the (real) speed factor $\gamma(x,\mu)$. This factor is usually set to 1, but for reasons explained later, it may be advisable, in the present context, to take a smaller value on the links close to the boundary of Λ. The associated link variables are then slowed down and tend to deviate less from their initial values at the end of the evolution (stochastic equations where different field components are evolved at different speeds have previously been considered by Davies et al. [19]).

It is, incidentally, straightforward to incorporate this modification in the numerical integration scheme that is used to solve the molecular dynamics equations. As in the case of the standard HMC algorithm, an acceptance-rejection step is finally required to correct for the integration error and thus to obtain a transition probability that satisfies detailed balance with respect to the distribution $P_\Lambda[U]$.

5. Global acceptance-rejection step

The gauge field configurations that are generated by applying the local algorithm described above can be taken as proposed configurations for the theory on a larger block or on the whole lattice. Whether a proposed field is accepted or not is then decided by applying a stochastic acceptance-rejection step. A particular implementation of this step is now going to be discussed, first for the case of the transition from a given block Λ to the full lattice.

5.1 Acceptance probability

Following the general strategy summarized in subsection A.5, the probability distribution $P_\Lambda[U]$ is considered to be an approximation to the exact distribution $P[U]$. An auxiliary pseudo-fermion field χ is then added to the system with action

$$S_\chi = \big(\chi, (Q_\Lambda + Q_{\Lambda^*})Q^{-1}\varrho_\Lambda^2 Q^{-1}(Q_\Lambda + Q_{\Lambda^*})\chi\big)\,, \tag{5.1}$$

where a local shape function

$$\varrho_\Lambda(x) = \theta_{\partial\Lambda}(x) + \epsilon\,(1-\theta_{\partial\Lambda}(x))\,, \tag{5.2}$$

with parameter $\epsilon > 0$, has been included for reasons that will become clear in a moment. The simultaneous probability distribution of the enlarged theory,

$$\hat{P}[\chi, U] = \frac{1}{\hat{\mathcal{Z}}}\,\mathrm{e}^{-S_\chi} P_\Lambda[U]\,, \tag{5.3}$$

reproduces the correct distribution when the auxiliary field is integrated out. Note that ϱ_Λ is invertible and independent of the fields. The associated determinant is therefore a constant that cancels out once the normalization factors are taken into account.

Before the acceptance-rejection step can be applied, a pseudo-fermion field χ must be generated randomly with conditional probability $\hat{P}[\chi\,|\,U]$. This is easily achieved by setting

$$\chi = (Q_\Lambda + Q_{\Lambda^*})^{-1}\,Q\varrho_\Lambda^{-1}\eta\,, \tag{5.4}$$

where η is chosen with probability proportional to $\mathrm{e}^{-(\eta,\eta)}$. The proposed field U' generated by the block simulation algorithm is then accepted with probability

$$P_{\rm acc} = \min\big\{1, \mathrm{e}^{-\Delta S_\chi}\big\}\,, \qquad \Delta S_\chi \equiv S_\chi|_{U\to U'} - S_\chi\,, \tag{5.5}$$

which ensures that the gauge field is correctly updated with a transition probability that satisfies detailed balance with respect to the exact distribution $P[U]$.

5.2 Elimination of the volume terms

In the following lines it is shown that ΔS_χ only depends on the values of η on the exterior boundary $\partial\Lambda$, up to terms that are proportional to the parameter ϵ of the shape function ϱ_Λ and that can, therefore, be eliminated by taking the limit $\epsilon \to 0$.

The proof starts from the identities

$$\Delta S_\chi = (\zeta, \zeta) - (\eta, \eta) \,, \tag{5.6}$$

$$\begin{aligned} \zeta &\equiv \varrho_\Lambda Q'^{-1} \left(Q'_\Lambda + Q'_{\Lambda^*}\right) \left(Q_\Lambda + Q_{\Lambda^*}\right)^{-1} Q \varrho_\Lambda^{-1} \eta \\ &= \eta + \varrho_\Lambda Q'^{-1} \left(Q'_\Lambda - Q_\Lambda\right) \left(Q_\Lambda + Q_{\Lambda^*}\right)^{-1} \left(Q_{\partial\Lambda} + Q_{\partial\Lambda^*}\right) \varrho_\Lambda^{-1} \eta, \end{aligned} \tag{5.7}$$

where the last line has been obtained using the block decomposition (3.2) of the Dirac operator and the fact that Q_{Λ^*}, $Q_{\partial\Lambda}$ and $Q_{\partial\Lambda^*}$ are independent of the active link variables. The crucial observation is then that all terms proportional to $Q_{\partial\Lambda^*}\varrho_\Lambda^{-1}\eta$ vanish since $Q'_\Lambda - Q_\Lambda$ acts in the subspace of fermion fields supported on Λ. Recalling eqs. (3.5) and (5.2), this leads to

$$\zeta = \eta + \varrho_\Lambda Q'^{-1} \left(Q'_\Lambda - Q_\Lambda\right) Q_\Lambda^{-1} Q_{\partial\Lambda} \eta \tag{5.8}$$

and thus to an expression

$$\Delta S_\chi = 2\,\mathrm{Re}(\theta_{\partial\Lambda}\eta, \xi) + (\xi, \xi) + \mathrm{O}(\epsilon) \,, \tag{5.9}$$

$$\xi \equiv \theta_{\partial\Lambda} Q'^{-1} \left(Q'_\Lambda - Q_\Lambda\right) Q_\Lambda^{-1} Q_{\partial\Lambda} \eta \,, \tag{5.10}$$

of the type announced above.

At this point the parameter ϵ can safely be taken to zero because all terms that are inversely proportional to ϵ have disappeared. Moreover, since only the component $\theta_{\partial\Lambda}\eta$ of the field η enters the final expressions, the constraint

$$\eta = \theta_{\partial\Lambda}\eta \tag{5.11}$$

may now be imposed without loss. The field is thus restricted to the linear space of boundary values for the Dirichlet problem on Λ (cf. subsection 3.3).

5.3 Acceptance rate

Evidently the whole approach will fail in practice unless the acceptance probability (5.5) is reasonably large on average. This will obviously be the case if ΔS_χ is only rarely greater than 1, and the question thus arises whether there is any theoretical reason to expect this to be so.

Some insight into the problem can be gained by assuming, for a moment, that U' differs from U only on a single link l in Λ. In eq. (5.9) the first term then involves the propagation of the random field η from the exterior boundary $\partial\Lambda$ to the endpoints of l and a second propagation from there back to the boundary. The other term has a similar structure but involves altogether four quark propagators.

Quark propagators typically decay like d^{-3} in the short-distance regime and more rapidly at larger distances d. The contributions to ΔS_χ with two quark propagators are thus suppressed by a factor $d^{-3}d'^{-3}$ at least, where d and d' are the distances of the link l from the initial and final points on $\partial\Lambda$. The sum over all these points then still needs to be performed, but since η is a random field there are strong cancellations in this sum and on average the result is enhanced by a factor proportional to the number of boundary points only (rather than its square). An important suppression factor thus remains, particularly when the link l is a fair distance away from the boundary.

Reasonable acceptance rates can thus be expected if the proposed changes in the link variables on the links closer to the boundary of the block are not too large. By adjusting the speed factor $\gamma(x,\mu)$ in the molecular dynamics equations (4.3) and (4.4), this can be easily achieved. Note that there must also be important cancellations in the sum over the active link variables that is implicit in eq. (5.10), because they change in random directions in group space.

5.4 Numerical exercises

Some numerical confirmation of this qualitative argumentation is clearly desirable at this point. The decay of the quark propagator away from the boundary of the block, for example, can be checked by calculating the solution $\psi = -Q_\Lambda^{-1}Q_{\partial\Lambda}\eta$ of the Dirichlet problem on Λ. A representative result of such studies for the average of $|\psi(x)|^2$ over all boundary values and gauge fields is reported in figure 4.

Another check on the correctness of the theoretical argumentation is obtained by computing the average of the acceptance probability $P_{\rm acc}$ in the quenched approximation, where the proposed changes in the active link variables are generated using a link update algorithm. The results in this case are rather encouraging too and they demonstrate, in particular, the importance of the elimination of the volume terms.

In all these studies the dependence of the calculated quantities on the quark mass appears to be fairly weak, as long as the edges of the block are smaller than 1 fm or so. This is perhaps not totally surprising in view of what has been said above, but the observation suggests that the algorithm will remain effective also at small quark masses.

5.5 Hierarchical filter

As was briefly discussed in subsection 2.3, a hierarchical structure should be adopted on large lattices where the proposed gauge field configurations are filtered through a sequence of blocks of increasing size. The filter involves an acceptance–rejection step for each transition from a block Λ to the next larger block Ω. Equations (5.9)–(5.11) remain valid in this case if Q' is replaced by Q'_Ω. An important point to note here is that the distance of the links where the active link variables reside to the boundary of the blocks is increasing. The acceptance probability for the step from one block to the next is therefore expected to rise quickly.

JHEP05(2003)052

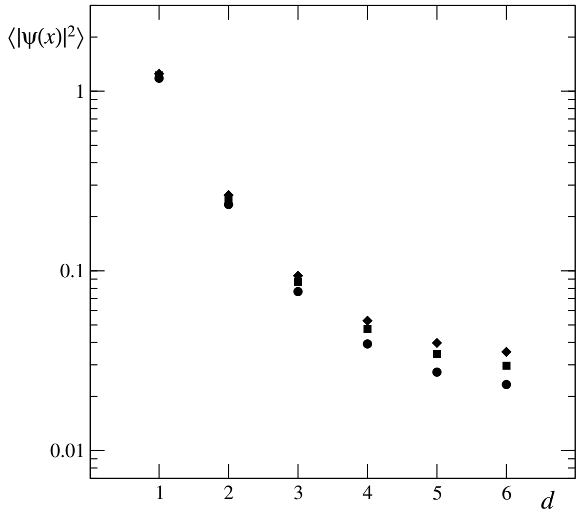

Figure 4: Average magnitude squared of the solution $\psi(x)$ of the Dirichlet problem (3.6), (3.7) on a 12^4 block as a function of the distance d from the exterior boundary of the block. The data shown are from a simulation of quenched QCD at a lattice spacing of about 0.1 fm. Diamonds, squares and circles correspond to values of the quark mass where the pion mass is known to be approximately equal to 320, 457 and 579 MeV respectively [20].

The proposed configurations that remain after the application of the hierarchical filter must finally be submitted to a global acceptance–rejection step. If several configurations are submitted simultaneously, as suggested in subsection 2.3, the acceptance probability is again given by eqs. (5.5) and (5.9)–(5.11), where Λ now stands for the union of the big blocks that contain the proposed configurations. The boundary $\partial\Lambda$ is then the union of the boundaries of these blocks and $\theta_{\partial\Lambda}$ the sum of the corresponding characteristic functions. In particular, there are terms in ΔS_χ that couple the random field η on the boundaries of different blocks.

6. Miscellaneous remarks

1. *Preconditioning.* As is generally the case in lattice QCD, the preconditioning of the Dirac operator can be expected to have a significant impact on the performance of the proposed algorithm. In particular, it is straightforward to implement even–odd preconditioning both in the block simulation algorithm and in the global acceptance–rejection step.

2. *Block simulation algorithm.* There is a range of algorithms that can be used to generate the proposed gauge field configurations on a given block. In particular, most improvements of the HMC and the PHMC algorithms that have been introduced over the years, such as the two-boson method [21, 22], should be beneficial here too.

   A more radical possibility is to generate the configurations through a standard link update algorithm, using an adapted gauge field action, and to apply a stochastic acceptance-rejection step to correct for the quark determinant [7]–[11]. Following the lines of subsection 5.2, the random pseudo-fermion field that needs to be introduced in the correction step can be reduced to the support of $Q'_{\Lambda} - Q_{\Lambda}$ in this case. This will no doubt increase the average acceptance probability, but whether the approach can compete with the block HMC algorithm, for example, remains to be seen.

3. *Parallel processing.* The proposed algorithm is well suited for parallel computers, because distant blocks can be updated independently of each other. Different distributions of the workload are conceivable, and it is not required that the smallest blocks be processed on single nodes, although in this case the communication overhead is minimized. It is, incidentally, advisable to translate the gauge field after each cycle rather than shifting the block positions. The program then always operates on the same blocks and their position can be chosen so as to fit the processor grid in the best possible way.

7. Conclusions

The problem to find efficient simulation algorithms for unquenched lattice QCD in the chiral regime has been around for many years. Whether the approach described in this paper provides a viable solution is not certain and can only be decided after extensive numerical tests have been performed.

It is quite clear, however, that significant progress in this area can only be made if the short- and long-distance effects of the sea quarks are treated differently. This insight is not new and has motivated the truncated determinant approximation of Duncan et al. [23, 24], for example, and a modification of the HMC algorithm with multiple molecular dynamics time scales [25]. To some extent, the two-boson method studied in refs. [21, 22] can also be put under this general heading.

The algorithm proposed here separates the short-distance effects from the long-distance ones by updating the gauge field on physically small blocks of lattice points that are visited sequentially. On the blocks the theory can be simulated even at vanishing quark masses, because the chosen boundary conditions imply an infrared cutoff on the spectrum of the Dirac operator. The long-distance effects of the sea quarks are then incorporated by a sequence of acceptance-rejection steps that lead from the small blocks to larger blocks and eventually to the full lattice.

Acknowledgments

I wish to thank Martin Hasenbusch and Ulli Wolff for correspondence on fermion simulation algorithms and Peter Weisz for a critical reading of the manuscript. The computer time for the numerical studies reported in section 5 has been kindly provided by DESY–Hamburg and by the Institute for Theoretical Physics at the University of Berne.

A. Transition probabilities

To prove the correctness of the algorithm proposed in this paper it suffices to write down the transition probabilities that are associated to the various algorithmic steps and to recall some basic facts from the general theory of numerical simulations. For completeness, the relevant theoretical results are summarized in this appendix.

A.1 Convergence theorem

To avoid any unnecessary complications, an abstract discrete system will be considered with a finite number of states s. The thermal equilibrium distribution of the states is denoted by $P(s)$ and it is assumed that $P(s) > 0$ for all s.

Numerical simulation algorithms for this system (as far as they follow the standard procedures) are defined by a transition probability $T(s \to s')$ that satisfies

1. $T(s \to s')$ is non-negative and such that $\sum_{s'} T(s \to s') = 1$ for all s.

2. The equilibrium distribution is preserved, i.e. for all states s' the equation $\sum_s P(s) \times T(s \to s') = P(s')$ holds.

3. The recurrence probability $T(s \to s)$ is non-zero for all s and any state can be reached from any other state in a finite number of transitions.

In practice the simulation starts from an initial state s_0 and generates a sequence of states $s_1, s_2, \ldots$ recursively with transition probability $T(s_k \to s_{k+1})$. The fundamental theorem of simulation theory then asserts that the states in the sequence will be asymptotically distributed according to the equilibrium distribution.

A.2 Composition rule & detailed balance

If T_1 and T_2 are any two transition probabilities that satisfy conditions 1–3, it is trivial to show that the same is true for the composed transition probability

$$T(s \to s') = \sum_r T_1(s \to r) T_2(r \to s') \,. \tag{A.1}$$

The associated algorithm generates an intermediate state r with probability $T_1(s \to r)$ and then the new state s' with probability $T_2(r \to s')$. Evidently any number of algorithms can be combined in this way.

Transition amplitudes that satisfy detailed balance,

$$P(s)T(s \to s') = P(s')T(s' \to s) \,, \tag{A.2}$$

are an important special case. If detailed balance holds, property 2 is an immediate consequence of property 1. However, the composed transition probability (A.1) in general does not satisfy detailed balance if T_1 and T_2 do, because the order of the factors in eq. (A.1) matters. It is possible to correct for this deficit by choosing the order randomly or by forming reversion-symmetric products.

JHEP05(2003)052

A.3 Acceptance-rejection method

Valid algorithms may often be obtained from transition amplitudes $T_0(s \to s')$ that satisfy detailed balance with respect to some approximate distribution $P_0(s)$ [26]. The idea is to first propose a new state s' according to this probability and to accept it with probability

$$P_{\mathrm{acc}}(s,s') = \min\left\{1, \frac{P_0(s)P(s')}{P(s)P_0(s')}\right\}. \tag{A.3}$$

Otherwise (i.e. when s' is rejected) the old state s is taken to be the new one.

The transition amplitude that is associated to this algorithm,

$$T(s \to s') = T_0(s \to s')P_{\mathrm{acc}}(s,s') + \delta_{ss'} - \delta_{ss'} \sum_r T_0(s \to r)P_{\mathrm{acc}}(s,r), \tag{A.4}$$

satisfies detailed balance and also conditions 1 and 3. It must be stressed, however, that for this statement to be true it is essential that $T_0(s \to s')$ satisfies detailed balance with respect to the distribution $P_0(s)$ and not only condition 2.

A.4 Using auxiliary stochastic variables

When the equilibrium probability distribution is too complicated to be treated directly, an equivalent but more accessible system may sometimes be found that involves auxiliary stochastic variables. In the case of the HMC algorithm, for example, the auxiliary variables are the pseudo-fermion field and the momenta of the link variables [1]. A simulation algorithm may then be set up for the enlarged system, and this leads to a correct algorithm for the original system under certain conditions.

In abstract terms the states of the enlarged system are labelled by pairs v, s, where v and s represent the auxiliary and the basic variables respectively. The equivalence to the original system is then expressed through the identity

$$P(s) = \sum_v \hat{P}(v,s) \tag{A.5}$$

in which $\hat{P}(v,s)$ denotes the equilibrium distribution of the enlarged system. In the following the conditional probability

$$\hat{P}(v|s) = \frac{\hat{P}(v,s)}{P(s)} \tag{A.6}$$

to find v given s will play an important rôle.

Now if $\hat{T}(v,s \to v',s')$ is any given transition probability for the enlarged system that satisfies conditions 1 and 2 (but not necessarily condition 3), an update algorithm for the original system is obtained as follows. Starting from an initial state s, the auxiliary variables v are first chosen randomly with conditional probability $\hat{P}(v|s)$. After that, a state v', s' is generated with probability $\hat{T}(v,s \to v',s')$ and s' is taken to be the new state of the original system. It is straightforward to show that the associated transition probability

$$T(s \to s') = \sum_{v,v'} \hat{P}(v|s)\hat{T}(v,s \to v',s') \tag{A.7}$$

satisfies conditions 1 and 2. Moreover, detailed balance holds if the transition probability $\hat{T}(v,s \to v',s')$ has this property relative to the equilibrium distribution of the enlarged system. Whether condition 3 is fulfilled depends on both $\hat{P}(v|s)$ and $\hat{T}(v,s \to v',s')$ and has to be verified in each case.

A.5 Stochastic acceptance-rejection method

The starting point here is again an algorithm with transition probability $T_0(s \to s')$ that satisfies detailed balance with respect to some approximate distribution $P_0(s)$. An enlarged system is then considered, exactly as above, but the auxiliary variables are now only used in the acceptance-rejection step where (A.3) is replaced by

$$P_{\rm acc}(s,s') = \sum_v \hat{P}(v|s) \min\left\{1, \frac{P_0(s)\hat{P}(v,s')}{\hat{P}(v,s)P_0(s')}\right\}. \tag{A.8}$$

This means that once a new state s' of the original system is proposed, a random choice of the auxiliary variables v first needs to be made, with conditional probability $\hat{P}(v|s)$, before it can be decided whether s' should be accepted or not.

The correctness of this method follows from the identity

$$\left[\frac{P(s)}{P_0(s)}\right] P_{\rm acc}(s,s') = \left[\frac{P(s')}{P_0(s')}\right] P_{\rm acc}(s',s) \tag{A.9}$$

that is easily derived from the definition (A.8) and the properties of the distributions $P_0(s)$ and $\hat{P}(v,s)$. In particular, the associated transition probability (A.4) satisfies detailed balance. It has recently been noted, however, that [11]

$$\begin{aligned} P_{\rm acc}(s,s') &\leq \min\left\{\sum_v \hat{P}(v|s), \sum_v \frac{P_0(s)\hat{P}(v,s')}{P(s)P_0(s')}\right\} \\ &= \min\left\{1, \frac{P_0(s)P(s')}{P(s)P_0(s')}\right\}, \end{aligned} \tag{A.10}$$

which shows that the stochastic acceptance-rejection method tends to be less efficient than the non-stochastic method.

References

[1] S. Duane, A.D. Kennedy, B.J. Pendleton and D. Roweth, *Hybrid Monte Carlo*, *Phys. Lett.* **B 195** (1987) 216.

[2] P. de Forcrand and T. Takaishi, *Fast fermion Monte Carlo*, *Nucl. Phys.* **53** *(Proc. Suppl.)* (1997) 968 [hep-lat/9608093].

[3] R. Frezzotti and K. Jansen, *A polynomial hybrid Monte Carlo algorithm*, *Phys. Lett.* **B 402** (1997) 328 [hep-lat/9702016];
The PHMC algorithm for simulations of dynamical fermions, I. Description and properties, *Nucl. Phys.* **B 555** (1999) 395 [hep-lat/9808011].

[4] M. Lüscher, *A new approach to the problem of dynamical quarks in numerical simulations of lattice QCD*, *Nucl. Phys.* **B 418** (1994) 637 [hep-lat/9311007].

JHEP05(2003)052

[5] H.A. Schwarz, *Gesammelte Mathematische Abhandlungen*, vol. 2, Springer Verlag, Berlin 1890.

[6] Y. Saad, *Iterative methods for sparse linear systems*, PWS Publishing Company, Boston 1996; see also `http://www-users.cs.umn.edu/~saad/`.

[7] M. Hasenbusch, *Speeding up finite step-size updating of full QCD on the lattice*, *Phys. Rev.* **D 59** (1999) 054505 [hep-lat/9807031].

[8] ALPHA collaboration, R. Frezzotti, M. Hasenbusch, U. Wolff, J. Heitger and K. Jansen, *Comparative benchmarks of full QCD algorithms*, *Comput. Phys. Commun.* **136** (2001) 1 [hep-lat/0009027].

[9] A. Hasenfratz and F. Knechtli, *Simulation of dynamical fermions with smeared links*, *Comput. Phys. Commun.* **148** (2002) 81 [hep-lat/0203010].

[10] A. Hasenfratz and A. Alexandru, *Evaluating the fermionic determinant of dynamical configurations*, *Phys. Rev.* **D 65** (2002) 114506 [hep-lat/0203026]; *Partial-global stochastic Metropolis update for dynamical smeared link fermions*, *Phys. Rev.* **D 66** (2002) 094502 [hep-lat/0207014].

[11] ALPHA collaboration, F. Knechtli and U. Wolff, *Dynamical fermions as a global correction*, hep-lat/0303001.

[12] K.G. Wilson, *Confinement of quarks*, *Phys. Rev.* **D 10** (1974) 2445.

[13] B. Sheikholeslami and R. Wohlert, *Improved continuum limit lattice action for QCD with Wilson fermions*, *Nucl. Phys.* **B 259** (1985) 572.

[14] 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].

[15] M. Lüscher, R. Narayanan, P. Weisz and U. Wolff, *The Schrödinger functional: a renormalizable probe for nonabelian gauge theories*, *Nucl. Phys.* **B 384** (1992) 168 [hep-lat/9207009].

[16] S. Sint, *On the Schrödinger functional in QCD*, *Nucl. Phys.* **B 421** (1994) 135 [hep-lat/9312079].

[17] 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 [hep-lat/9803017], erratum *ibid.* **B643** (2002) 517.

[18] L. Giusti, C. Hoelbling, M. Lüscher and H. Wittig, *Numerical techniques for lattice QCD in the ϵ-regime*, hep-lat/0212012, to appear in *Comput. Phys. Commun.*.

[19] C.T.H. Davies et al., *Fourier acceleration in lattice gauge theories, 3. Updating field configurations*, *Phys. Rev.* **D 41** (1990) 1953.

[20] CP-PACS collaboration, S. Aoki et al., *Light hadron spectrum and quark masses from quenched lattice QCD*, *Phys. Rev.* **D 67** (2003) 034503 [hep-lat/0206009].

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

[22] M. Hasenbusch and K. Jansen, *Speeding up lattice QCD simulations with clover-improved Wilson fermions*, hep-lat/0211042.

JHEP05(2003)052

[23] A. Duncan, E. Eichten and H. Thacker, *An efficient algorithm for QCD with light dynamical quarks*, *Phys. Rev.* **D 59** (1999) 014505 [hep-lat/9806020].

[24] A. Duncan, E. Eichten and J. Yoo, *Unquenched QCD with light quarks*, hep-lat/0209123.

[25] TrinLat collaboration, M. Peardon and J. Sexton, *Multiple molecular dynamics time-scales in Hybrid-Monte-Carlo fermion simulations*, hep-lat/0209037.

[26] N. Metropolis et al., *Equation-of-state calculations by fast computing machines*, *J. Chem. Phys.* **21** (1953) 1087.