Nuclear Physics B280 [FS 18] (1987) 181–191
North-Holland, Amsterdam

MONTE CARLO SIMULATION OF FERMIONIC FIELDS*

Michael CREUTZ and Rajiv GAVAI**

Physics Department, Brookhaven National Laboratory, Upton, NY 11973, USA

Received 6 June 1986
(Revised 29 July 1986)

We present a simple algorithm for Monte Carlo simulation of field theories containing fermionic fields. It is based on a small step-size limit in a Metropolis et al. scheme. The approach has similarities with the Langevin evolution. A simple modification of the procedure permits an interpolation to the pseudofermionic algorithm.

1. Introduction

Monte Carlo simulation has become a prime tool in the study of nonperturbative phenomena in quantum field theory. The results for pure gauge fields have been spectacular, giving some of our first quantitative information on the solution of the SU(3) gauge theory of the strong interactions [1]. Despite the successes, progress has been considerably slower in attempts to extend these calculations to include fermionic fields. While numerical algorithms have in principle been found for simulating the dynamical quarks of the strong interactions, they are all extremely intensive in their demands for computer time [2–7]. This means that necessary checks for finite size effects, cutoff independence, and other sources of systematic error have been severely limited.

It is unlikely that the problem with fermions is merely one of not yet having found the optimal algorithm. The complexity of maintaining antisymmetry in the wave function of a filled Dirac sea requires detailed information on an entire spectrum of states. The extent to which this means that fermionic algorithms are inherently slower than bosonic procedures is still a matter of debate.

In this paper we wish to present a new algorithm for simulating fermionic fields. Our approach has close connections with the Langevin simulations advocated in

* This manuscript has been authored under contract number DE-AC02-76CHO0016 with the US Department of Energy.

** Permanent address after November 1 1986: Theory Group, Tata Inst. of Fundamental Research, Homi Bhabha Road, Bombay 400005, India.

ref. [2] as well as the pseudofermion approach of ref. [3]. We perform some simple tests of the algorithm with SU(3) lattice gauge theory on a small lattice.

As with most other promising techniques for dealing with fermionic fields [2–4, 6], we attempt to eliminate the need for a Monte Carlo simulation inside a Monte Carlo simulation by considering updating sweeps of the entire lattice volume but with only small changes in the value of any given gauge field. Whether this attempt can succeed depends on whether the size of the small changes can remain finite as the system volume is increased. This is a reasonable expectation as long as the correlation length is kept fixed, but it may be that the critical slowing down of a simulation with increasing correlation length is inherently more severe with fermions than with conventional bosonic simulations.

2. Preliminaries

We study a system of fermionic fields interacting with a bosonic field via a coupling term which is quadratic in the fermions. Thus the generic "partition function" we wish to simulate is

$$Z = \int \mathrm{d}\psi\, \mathrm{d}\bar{\psi}\, \mathrm{d}A \exp\bigl(-S_{\mathrm{G}}(A) - \bar{\psi} M(A) \psi\bigr). \tag{1}$$

Here A denotes the bosonic fields, which would be gauge fields in the case of the strong interactions, and ψ denotes the fermionic or quark fields. The quantity $S_{\mathrm{G}}(A)$ denotes the part of the action which depends on the gauge fields alone and $M(A)$ contains the kinetic, mass, and interaction terms for the fermions. The fermionic fields are taken to be anticommuting variables, integration over which is defined in the standard way to give

$$Z = \int \mathrm{d}A\, |M(A)| \exp\bigl(-S_{\mathrm{G}}(A)\bigr). \tag{2}$$

The problem with attempting a direct Monte Carlo simulation of the system in eq. (1) is that the fields ψ are not ordinary numbers, and the integrand cannot be interpreted as a probability. Eq. (2), on the other hand, is an ordinary integral, but the matrix M is extremely large (although quite sparse in practice) and it is impractical to calculate its determinant the large number of times necessary for a conventional Monte Carlo evaluation.

To proceed, we follow ref. [5] and rewrite (2) as

$$Z \propto \int \mathrm{d}\phi\, \mathrm{d}\phi^*\, \mathrm{d}A \exp\bigl(-S_{\mathrm{G}}(A) - \phi^* M^{-1}(A) \phi\bigr), \tag{3}$$

where ϕ is a commuting bosonic field. In order that the ϕ integral be well defined,

M should be a positive definite matrix. We will ensure this by considering M to be a square

$$M = (D + m)(D^\dagger + m), \tag{4}$$

where m is the fermion mass in lattice units. This is equivalent to considering a doubling of the fermionic species, one interacting with A via $D(A)$ and the other via $D^\dagger$. We will later discuss a possible way to remove this doubling.

To simulate this system, we wish to generate an ensemble of configurations of the fields ϕ and A with a probability distribution

$$P(A, \phi) \propto \exp\bigl(-S_G(A) - \phi^*(D^\dagger + m)^{-1}(D + m)^{-1}\phi\bigr). \tag{5}$$

The appearance of the inverse of the large matrix $D + m$ in the probability makes this difficult for direct attack, although the authors of ref. [5] observed that only one vector, $(D + m)^{-1}\phi$, is needed at any step and that this could be obtained via an iterative scheme such as the Gauss-Seidel or the conjugate gradient algorithm. The time involved for such an inversion will grow with the volume of the system and thus this gives an exact algorithm for simulating the fermionic system where the time per sweep of all variables will grow as the volume squared. The algorithms of refs [2–4, 6] attempt to avoid this severe growth at the expense of making a slow migration through phase space and only finding matrix elements of $(D + m)^{-1}$ once per full sweep of the gauge variables. Of course, the slower the flow through phase space, the more sweeps will be necessary to reach equilibrium.

3. The algorithm

Our algorithm consists of alternate sweeps through the ϕ and the A fields. The ϕ updating is particularly simple. We begin by generating a random vector χ with gaussian probability

$$P_\chi \propto \exp(-\chi^*\chi). \tag{6}$$

If we now take

$$\phi = (D + m)\chi, \tag{7}$$

this will be distributed with the desired probability

$$P_\phi \propto \exp\bigl(-\phi^*(D^\dagger + m)^{-1}(D + m)^{-1}\phi\bigr). \tag{8}$$

The construction of χ and ϕ is computationally fast because the individual components of χ are independent and because the matrix $D + m$ needed in eq. (7) is local, i.e. matrix elements are nonvanishing only for physically nearby sites, at

least for the usual models of interest. Thus we rapidly obtain a new ϕ field independent of its previous value. We note in passing that the Langevin implementation in ref. [2] does an equivalent construction in the use of gaussian random numbers to simulate the fermionic determinant.

Actually, as will become apparent, we will not need ϕ explicitly anywhere in the updating algorithm. Although we could eliminate this field and consider only χ, we find it simpler for discussion to consider the simple form of the coupled probability distribution in eq. (5).

In addition to the field χ, the updating procedure for the gauge field requires another quantity

$$\xi = (D^\dagger + m)^{-1}\chi = M^{-1}\phi . \tag{9}$$

This, unfortunately, is not so trivial to obtain. We use the conjugate gradient algorithm to find this quantity. The use of the conjugate gradient or similar inversion scheme is common also to the Langevin (ref. [2]) or microcanonical (ref. [4]) approaches.

We now come to the updating of the A field. What one would like to do is something like a Metropolis et al. [8] procedure where the acceptance of changes is governed by the action

$$S(A,\phi) = S_G(A) + \phi^*(D^\dagger + m)^{-1}(D+m)^{-1}\phi . \tag{10}$$

Unfortunately, every time A is changed, $D+m$ changes and its inverse on ϕ would need to be calculated again. It is this slow procedure we wish to avoid. As we will be working with small changes in A, consider the first derivative of eq. (10) with respect to A

$$\left.\frac{\partial S}{\partial A}\right|_\phi = \frac{\partial S_G}{\partial A} - 2\,\mathrm{Re}\left(\phi^*(D^\dagger + m)^{-1}(D+m)^{-1}\frac{\partial D}{\partial A}(D+m)^{-1}\phi\right)$$

$$= \frac{\partial S_G}{\partial A} - \xi^*\frac{\partial D}{\partial A}\chi - \chi^*\frac{\partial D^\dagger}{\partial A}\xi . \tag{11}$$

Note that if we consider the quantity

$$S_T(A,\chi,\xi) = S_G - \xi^*(D+m)\chi - \chi^*(D^\dagger + m)\xi , \tag{12}$$

then

$$\left.\frac{\partial S}{\partial A}\right|_\phi = \left.\frac{\partial S_T}{\partial A}\right|_{\chi,\xi} . \tag{13}$$

If we consider a small variation of A, to first order the change of S at constant ϕ

equals the change in S_T calculated at constant χ and ξ. As only changes in the action enter in the Metropolis [8] algorithm, if we update the A fields using the action S_T in such a procedure, this will be equivalent to using the correct action of eq. (10) up to terms of higher order in the changes. Our algorithm is to do exactly this, which is easily implemented efficiently because S_T is local.

This procedure makes errors which are of higher order in the size of the changes per sweep in the gauge field variables. A similar situation holds with the algorithms of refs. [2–4, 6]. One can, following ref. [2], envisage calculating the higher order terms in the effective equilibrium action, although it is not as straightforward here due to the probabilistic nature of the Metropolis algorithm. Furthermore, such formal calculations do not yield precise information on where still higher orders are negligible, making the algorithm a good approximation. Indeed, to have confidence that the errors induced due to these terms are small, one should study a desired measurable for a few values of this step size and extrapolate to the infinitesimal limit. Ref. [2] argued that for the Langevin algorithm a finite step-size represents a simulation with an effective action which differs from the initial one by terms which vanish with the step size. If this new action has the same continuum limit, then even these finite step simulations should give the same numerical results for physical observables. Nevertheless, an extrapolation to vanishing step is still necessary if one wishes to compare the results of different algorithms with a given set of parameters at a finite lattice spacing.

To the extent that the finite step errors are random, they may be thought of as an effective extra noise being added to the action. In a first order Langevin simulation, noise is used to simulate thermal fluctuations. Thus we might expect that a large part of the effect of these errors will be an effective increase of the system coupling. This increase appears in our simulation to be proportional to the Metropolis step-size.

4. A numerical test

We have tested the above procedure with SU(3) lattice gauge theory on a small system of 4^4 sites. For this analysis we used 8 flavors with a mass of 0.1 in lattice units; i.e. we used the action from ref. [4]. This choice of action permits comparisons with previous studies using other methods. A natural next step would be to use larger lattices and to reduce the number of flavors using known techniques.

As discussed above, we update the gauge fields with the Metropolis [8] procedure using the action of eq. (12). For a given link ℓ, the gauge field is characterized by a group element U_ℓ. To update this link, a trial U_ℓ' is obtained by multiplying U_ℓ by a matrix R from a set of 200 matrices. This set is formed by requiring that (i) for each matrix in the set its inverse also is in the set, and (ii) a given element of the set is obtained from $R = \exp(i\alpha \cdot \lambda)$, where λ_i, $i = 1, 8$, are the Gell-Mann matrices and α is a gaussian random vector chosen so that $\langle \alpha \rangle = 0$, and $\langle |\alpha|^2 \rangle = \rho^2$. The parameter

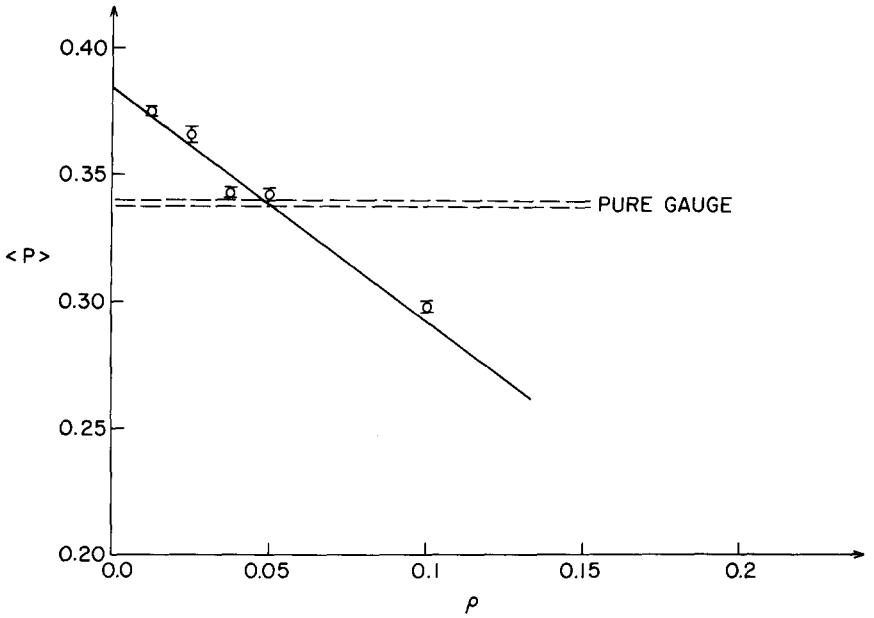

Fig. 1. The average plaquette as a function of the parameter ρ. The extrapolation to $\rho = 0$ should give the correct value for the plaquette in the interacting theory with fermions. The solid line is a linear fit.

ρ provides a way of controlling the size of the changes allowed in the gauge fields during a sweep of the lattice; indeed, ρ going to zero corresponds to taking a vanishing step size.

In fig. 1 we show the average plaquette

$$P = \left\langle \tfrac{1}{3}\mathrm{Re}\big(\mathrm{Tr}(U_p)\big)\right\rangle, \tag{14}$$

where U_p is the product of group elements around an elementary lattice square, as a function of the parameter ρ at $\beta = 4.5$. These points were obtained from 500 to 2000 iterations after equilibration, with the longer runs corresponding to the smaller values of ρ. The sweeps discarded for equilibration similarly varied from about 200 to 1000. The corridor marked by dashed lines denotes the value of the plaquette for the same value of β in the pure gauge theory. A linear extrapolation of our points to $\rho = 0$ gives a value $P = 0.383$, a value roughly consistent with that obtained in ref. [4]. In that paper the microcanonical algorithm was employed on an $8^3 \times 4$ lattice. We note that with 4 sites in the time direction the transition to a quark gluon plasma is expected at $\beta = 4.67 \pm 0.1$ [4]; consequently our simulation was done in the confining region. Expecting only finite size effects to differ, we are encouraged by the agreement of the two methods.

Note the rather steep ρ dependence in fig. 1. Indeed, for $\rho > 0.05$ the average plaquette is smaller than the pure gauge value, opposing the correct result that

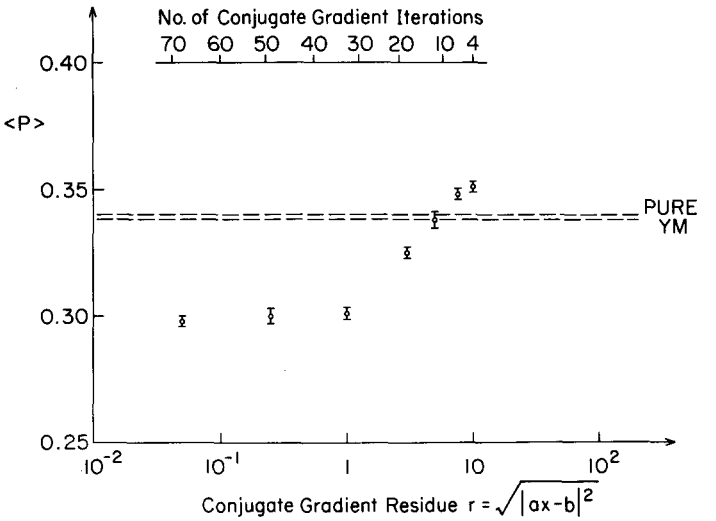

Fig. 2. The average plaquette at $\rho = 0.1$ as a function of the accuracy of the conjugate gradient inversion.

fermionic loops tend to order the system. We conclude that our procedure requires a rather small value of ρ. This may, however, be counteracted by the simplicity and speed of the procedure.

In addition to ρ, the algorithm has a dependence on the accuracy with which the conjugate gradient inversion for eq. (9) is carried out. The field ξ in that equation is obtained iteratively, and at the ith step the residue

$$r = \left\| (D^{\dagger} + m)\xi^{(i)} - \chi \right\| \tag{15}$$

determines the accuracy of the inversion. For the above analysis we required $r < 0.05$. To verify that this is adequate, we chose $\rho = 0.1$ and performed runs with several values of r. The results are displayed in fig. 2. The average plaquette is essentially independent of r for $r < 1.0$, showing that 0.05 was a conservative choice. Of course, the number of conjugate gradient iterations required decreases as r is increased, as seen from the top scale of fig. 2.

5. Comparisons with other approaches

The Metropolis algorithm in the limit of small step-size is quite close to the Langevin approach. In both cases small random changes are made to the field variables. A Metropolis program first considers small unbiased changes about the old field, and then, to maintain the desired peaking of the distribution towards lower action, rejects a fraction of those changes which go toward larger action. In

contrast, the Langevin approach always accepts the changes, but makes them in a direction which is biased towards lower action. In this approach the A field evolves in a "simulation time" coordinate via a stochastic differential equation of form

$$dA/d\tau = -\frac{\partial S}{\partial A} + \eta(\tau). \tag{16}$$

Here τ represents a noise term which simulates thermal fluctuations For more details see ref. [2]. Now note that the formula in eq. (13) gives the derivative of S with respect to A. If all fields χ and ξ are obtained from additional gaussian noise as in eqs. (6) and (9), we obtain essentially the algorithm in ref. [2]. For the purpose of numerical integration of eq. (16), the variable τ is made discrete with step size ε. On going from one τ step to the next, the noise term adds a random piece of magnitude $\varepsilon^{1/2}$ to each variable. This corresponds to the Metropolis step size in our algorithm; consequently, the parameter ε in ref. [2] is proportional to the square of our variable ρ.

The similarity of the approaches indicates that the errors due to a finite step size should be comparable. To directly make such a comparison, one should replot fig. 1 against a common definition of the step size. The simplest way to fix normalizations would be to use a number of simulation steps which characterizes the decorrelation of the system. As the step size is increased, this decorrelation time will decrease, but the finite step errors will increase. We conjecture that the behavior of the line in fig. 1 will be comparable for our algorithm and the Langevin approach, if the ordinate is taken as the inverse of this decorrelation time. As the momenta in a microcanonical simulation substitute for the noise in the Langevin approach, we expect a similar behavior for that method as well. Thus the choice between the algorithms amounts to which is easier to implement efficiently.

Our algorithm also has close connections with the pseudofermion method of ref. [3]. That approach attempts to find configurations of the field A with a Boltzmann weighting determined from the action

$$S_P = S_G + \mathrm{Tr}(\ln M). \tag{17}$$

For small changes in A, changes in this action are determined from the derivative

$$\frac{\partial S_P}{\partial A} = \frac{\partial S_G}{\partial A} + \mathrm{Tr}\left(M^{-1} \frac{\partial M}{\partial A} \right). \tag{18}$$

To estimate the needed matrix elements of M^{-1}, ref. [3] uses a Monte Carlo simulation to obtain commuting fields weighted by

$$P_\xi \propto \exp(-\xi^* M \xi). \tag{19}$$

Matrix elements of M^{-1} are given by expectation values of the ξ fields

$$(M^{-1})_{ij} = \langle \xi_j^* \xi_i \rangle . \tag{20}$$

If we now consider eq. (9) with χ obtained as in eq. (6), we see that our field ξ is constructed to have precisely the probability distribution in eq. (17). Indeed, the ξ fields in ref. [3] could just as well be obtained from a conjugate gradient inversion on gaussian random numbers as used here. Then our algorithm corresponds to using but a single configuration of fields to estimate M^{-1} via eq. (18). This would naively seem to have a substantial error, but as the step size for changes in A goes to zero, this error averages out.

Clearly our algorithm is an extreme case. One could interpolate between this and the algorithm of ref. [3] by averaging over some number N of ξ fields. This may also be thought of as considering N species of fermions each with its own ξ field, but then letting each species only contribute $1/N$ in the updating of the A field. As N increases to infinity, we obtain the pseudofermion algorithm of ref. [3]. In fig. 3 we show the step-size dependence of the average plaquette obtained for several values of N. Here we plot the plaquette versus the probability of accepting a trial change in

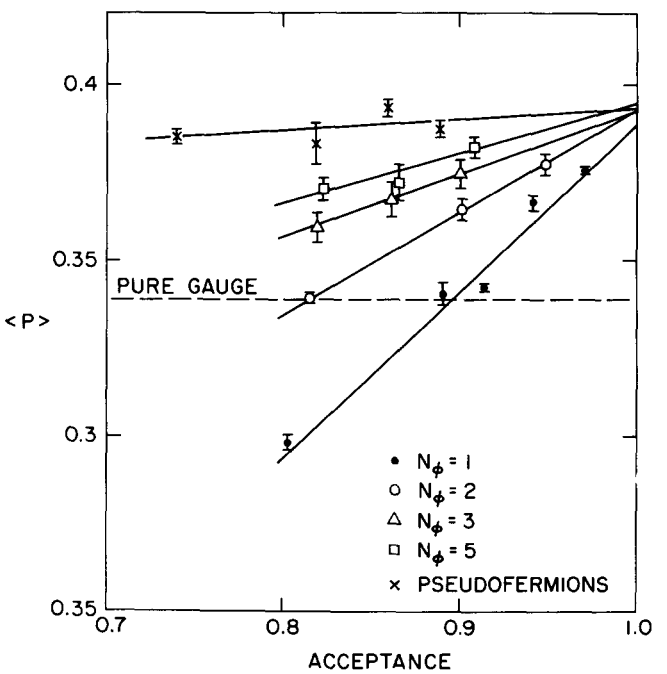

Fig. 3. The average plaquette as a function of the acceptance probability for a Metropolis hit. The various values of N_ϕ should all extrapolate to the same value at unit acceptance. The lines are linear fits.

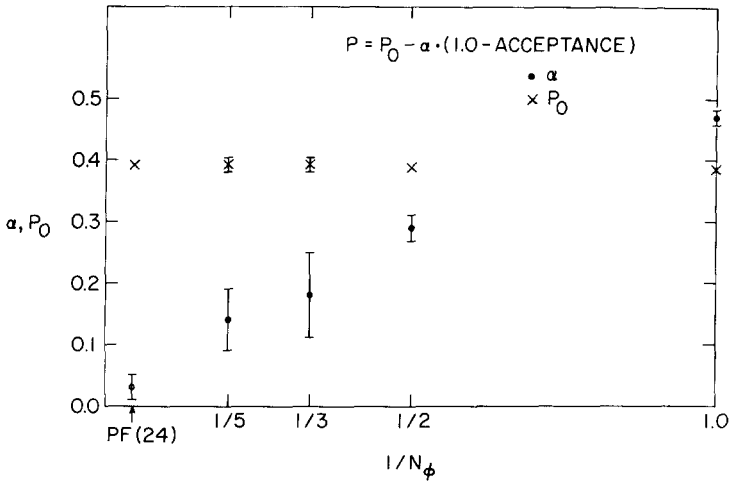

Fig. 4. The intercepts and slopes of the lines in fig. 3 as a function of the inverse of the number of ϕ fields.

the Metropolis updating. In this graph we also give points obtained from the pseudofermionic algorithm of ref. [3] where we used 24 pseudofermionic configurations to estimate the average in eq. (18). The solid lines in fig. 3 represent linear fits to the data. In fig. 4 we plot against $1/N$ the slopes and the intercepts at 100% acceptance for these fits. In this graph we also include a point for our pseudofermionic run, which is plotted at $1/N = \frac{1}{24}$, corresponding to the number of pseudofermionic configurations used in those runs. Note that all values of N are consistent with having a common intercept. Note also that the infinite N limit, which should represent the pseudofermionic algorithm, appears to have a small but nonvanishing slope.

This allowing of each species to contribute only fractionally in the updating of the A field may provide a scheme to reduce the effective number of fermion species overall. Naively, this can remove the extra doubling [9] introduced in eq. (4) as well as any inherent doubling in the basic formulation of the fermions. Such a possibility has been mentioned in the context of both the algorithms of refs. [2] and [3], where one simply puts a species reduction factor in front of the fermionic contribution while updating the A fields. There may, however, be some danger in this procedure because chiral symmetry breaking and anomalies suggest nonanalytic behavior as the number of fermionic species varies.

6. Conclusions

We have presented a particularly simple algorithm for the inclusion of fermionic fields in Monte Carlo simulations of quantum field theory. We replace the fermionic

action by scalar fields interacting with the inverse of the Dirac operator. These scalar fields are updated using gaussian random numbers. To update the gauge field, we use a small step-size limit of the Metropolis et al. [8] procedure and linearize the term involving the scalar fields. The procedure has similarities to the Langevin approach of ref. [2], but we feel that our method is simpler to implement. A variation of our approach permits an interpolation to the pseudofermionic algorithm of ref. [3]. A recent preprint [10] has presented a similar algorithm and suggested correcting for finite step errors by an appropriate acceptance or rejection of the new configuration at the end of each sweep of the gauge variables. Our algorithm is presumably also amenable to various acceleration techniques as extensively discussed in refs. [2] and [11].

References

[1] M. Creutz, L. Jacobs and G. Rebbi, Phys. Reports 95 (1983) 201
[2] G.G. Batrouni, G.R. Katz, A.S. Kronfeld, G.P.Lepage, B. Svetitsky and K.G. Wilson, Phys. Rev. D32 (1985) 2736
[3] F. Fucito, E. Marinari, G. Parisi and C. Rebbi, Nucl. Phys. B190 (1981) 369
[4] J. Polonyi and H.W. Wyld, Phys. Rev. Lett. 51 (1983) 2257;
J.B. Kogut, J. Polonyi, H.W. Wyld and D.K. Sinclair, Phys. Rev. Lett. 54 (1985) 1475
[5] D. Weingarten and D. Petcher, Phys. Lett. 99B (1981) 333
[6] A. Ukawa and M. Fukugita, Phys. Rev. Lett. 55 (1985) 1854
[7] G. Bhanot, U. Heller and I.O. Stamatescu, Phys. Lett. 129B (1983) 440;
Ph. De Forcrand and I.O. Stamatescu, Nucl. Phys. B261 (1985) 613;
M. Grady, Phys. Rev. D32 (1985) 1496;
A.D. Kennedy and J. Kuti, preprint NSF-ITP-85-21 (1985)
[8] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller and E. Teller, J. Chem. Phys. 21 (1953) 1087
[9] O. Martin and S. Otto, Phys. Rev. D31 (1983) 435
[10] R.T. Scalettar, D.J. Scalapino and R.L. Sugar, UC Santa Barbara preprint (1986)
[11] J. Goodman and A. Sokal, Phys. Rev. Lett. 56 (1986) 1015