

QCD with dynamical Wilson fermions

Rajan Gupta

T-8, MS-B285, Los Alamos National Laboratory, Los Alamos, New Mexico 87545

Apoorva Patel

Theory Group, CERN, Genève 23, Switzerland

Clive F. Baillie

Physics Department, Caltech, Pasadena, California 91125

Gerald Guralnik and Gregory W. Kilcup

Physics Department, Brown University, Providence, Rhode Island 02912

Stephen R. Sharpe

Physics Department, University of Washington, Seattle, Washington 98195

(Received 27 March 1989)

We present results from a study of QCD with two flavors of Wilson fermions using the hybrid Monte Carlo algorithm, which incorporates the effects of fermion loops exactly. We evaluate the performance of the algorithm and its potential for large-scale computations. We argue that in the best case the algorithm slows down as $V^{5/4}m_q^{-13/4}$ at a fixed gauge coupling. We present improved algorithms for calculating the inverse and the determinant of the Wilson fermion operator. Results for the finite-temperature transition on 4×6^3 and 6×8^3 lattices are presented at $\beta=5.2-5.6$. We also give Wilson loop expectation values obtained on 8^4 lattices at $\beta=5.3$ for six values of κ . The data show evidence for screening in the $q\bar{q}$ potential. Lastly, on comparing Wilson and staggered-fermion results we find that $\beta=5.3$ is far from the scaling region.

I. INTRODUCTION

The challenge facing numerical calculations of QCD is to simulate the full theory on large lattices, with small lattice spacing, and with realistic quark masses. Progress towards this goal has been slow, mainly because of the difficulty of including dynamical fermions in the numerical algorithms. One avenue of progress has been the development of exact algorithms for doing the simulations. The hybrid Monte Carlo algorithm (HMCA) has recently emerged as the preferred exact algorithm.¹ In the near future simulations with this algorithm on lattices of sizes up to $\sim 16^4$ will be possible. It is the aim of this paper to do some of the groundwork necessary so as to choose the parameters of future simulations sensibly.

The HMCA (Ref. 2) can simulate multiples of two flavors of Wilson fermions, or multiples of four flavors of staggered fermions. To make the calculations as realistic as possible, therefore, we must use Wilson fermions. The two degenerate species of fermions in the simulations then represent the physical up and down quarks. Most of the previous work on the HMCA has considered the four-flavor staggered-fermion theory.³⁻⁵ This work has considered algorithmic issues (How does one tune the algorithm?) and the thorny question of the existence and the order of a phase transition in finite-temperature QCD. One of the purposes of the present study is to extend this work to the two-flavor Wilson fermion theory.

The results of this paper address the following issues.

(A) *Tuning and monitoring the algorithm.* This is a question at the core of determining the speed and reliability of any promising algorithm. We discuss this in Sec. II, and argue in Sec. III that the HMCA slows down proportional to $m_q^{-13/4}$, in the limit of small quark masses.

(B) *Efficient algorithms for the Wilson fermion matrix.* In Sec. IV we present details of an efficient matrix inversion algorithm for Wilson fermions. This algorithm is competitive with the best implementation of an incomplete lower upper (ILU) algorithm. We also discuss a simplification of the pseudofermion representation of the fermion determinant.

(C) *Comparison with approximate algorithms.* In order to substantiate the claim that the HMCA is competitive with approximate algorithms we provide a detailed comparison of our results with those from the hybrid and Langevin algorithms in Sec. V.

(D) *Critical parameters.* The first step towards choosing reasonable running parameters for large lattices is to determine (a) for a given coupling $\beta \equiv 6/g^2$, the critical value of the hopping parameter κ_c , where the pion mass vanishes; (b) for given β and hopping parameter κ (which together specify the lattice spacing a), the position of the finite-temperature crossover/transition. This is most easily expressed as $N_t^c = 1/(T_c a)$, where T_c is the transition temperature in physical units and N_t^c as defined need not be an integer. The first calculation gives the location of the chiral limit and the second provides a rough estimate of the scale (assuming that $T_c \approx 150$ MeV). With

these parameters in hand, and drawing upon the accumulated experience of pure gauge theory simulations, one can pick an appropriate lattice size to avoid severe finite-size effects in zero-temperature observables. Furthermore, one can estimate the interesting physical region where the u - and d -quark masses are small and the effects of vacuum polarization are substantial. Our results for T_c are described in Sec. VI, while detailed results for κ_c will be published elsewhere.⁶

(E) *Screening in the $q\bar{q}$ potential.* Section VII is devoted to illustrating the effects of dynamical quarks on the static $q\bar{q}$ potential. We make a careful study of the screening as a function of the quark mass, and present results showing the changeover for large Wilson loops from the area law of the pure gauge theory to the perimeter law for full QCD.

(F) *Comparison of staggered and Wilson fermion results.* In the scaling region the results obtained with these two different fermion schemes must come together, and only then can one claim to have meaningful predictions for the continuum theory. In Sec. VIII we show how far away we are from achieving this goal.

We end with some comments on the outlook for simulations of QCD with realistic parameters.

II. SYSTEMATIC TESTS OF THE HMCA WITH WILSON FERMIONS

Our implementation of the HMCA for QCD with Wilson fermions closely parallels that we used for staggered fermions.⁴ The Dirac operator for Wilson fermions is

$$\Delta[U]_{ij} = \delta_{ij} + \kappa \sum_{\mu} [(\gamma_{\mu} - r)U_{i,\mu}\delta_{i,j-\mu} - (\gamma_{\mu} + r)U_{i-\mu,\mu}^{\dagger}\delta_{i,j+\mu}]. \quad (2.1)$$

In the following we set $r=1$. The determinant of this operator is real ($\gamma_5\Delta^{\dagger}\gamma_5 = \Delta$) but not necessarily positive definite. In order to interpret the fermion determinant as a probability factor, we have to work with the operator $\Delta^{\dagger}\Delta$. Only when the power of $\det(\Delta^{\dagger}\Delta)$ is an integer, a path-integral representation for it can be given in terms of the pseudofermion fields ϕ :

$$\int D\phi D\phi^* \exp[-\phi^{\dagger}(\Delta^{\dagger}\Delta)^{-1}\phi] = \det(\Delta)\det(\Delta^{\dagger}). \quad (2.2)$$

Thus simulations in which Wilson fermions are represented by pseudofermions are restricted to multiples of two degenerate flavors. This subject is further discussed in Sec. VI.

We use the Φ version of the hybrid algorithm introduced by Gottlieb *et al.*⁷ It corresponds to evolving the system variables through the phase space in a fictitious (computer) time, with the Hamiltonian

$$H_{\Phi} = \frac{\alpha}{2} \text{Tr} \sum P_{i,\mu}^2 + \frac{\beta}{N} \text{Re Tr} \sum (1 - U_p) + \phi^{\dagger}(\Delta^{\dagger}\Delta)^{-1}\phi. \quad (2.3)$$

Here $U_{i,\mu}$ are the gauge link variables, $P_{i,\mu}$ are the momenta conjugate to them, and U_p is the 1×1 plaquette. The ϕ fields have no dynamics in the Φ algorithm, and

the simulated partition function is

$$Z = \int D\phi D\phi^* DU DP \exp(-H_{\Phi}). \quad (2.4)$$

The introduction of the conjugate momenta as Gaussian variables does not alter the QCD correlation functions defined in terms of U and ϕ . All expectation values are measured as simple time averages after the system has equilibrated.

The Φ algorithm is reversible and area preserving. Its basic steps are these: Given a configuration U , refresh the momenta P and generate the pseudofields $\phi \equiv \Delta^{\dagger}\eta$ with Gaussian random numbers η . Then evolve the system to U', P' using leapfrog discretization of the molecular-dynamics equations of motion for a trajectory consisting of n_{MD} steps each of size ϵ (our variant of leapfrog evolves the U at the first half step). Every trajectory repeats these steps. Because of the discretization, the Hamiltonian is not exactly conserved during the evolution along a trajectory. The HMCA removes this $O(n_{\text{MD}}\epsilon^3)$ error by accepting the new fields U', P' at the end of a trajectory with the probability⁸

$$P = \min(1, e^{-\delta H}), \quad \delta H = H(U', P', \phi) - H(U, P, \phi). \quad (2.5)$$

An important feature of the HMCA is that the Hamiltonian H used in the accept/reject step need not be the same as the Hamiltonian H_{Φ} used in the hybrid preprocessor.² This fact can be exploited to increase the efficiency of the algorithm. A convenient choice is to take H and H_{Φ} to have the same form but different values of the couplings α , β , and κ . We denote the couplings in H_{Φ} by a subscript MD. They are tunable parameters which only affect the acceptance rate, while the equilibrium distribution is governed by the couplings appearing in H .

One can gain ~ 10 – 20 % in the acceptance rate by a suitable choice of these adjustable couplings. The optimal choice depends on the details of the leapfrog algorithm.⁴ The parameter α can be tuned,⁹ but all our results have been obtained with the choice $\alpha_{\text{MD}} = \alpha = 1$. With our leapfrog scheme, we find that close to maximal acceptance can be obtained by selecting $\beta - \beta_{\text{MD}} \approx 0.01$, and $\kappa = \kappa_{\text{MD}}$. A little more improvement is possible by choosing κ_{MD} slightly smaller than κ , but we have not investigated this thoroughly. These rules of thumb are similar to those found in our studies of the staggered fermion HMCA. Table I shows some results of our search for an optimal value for β_{MD} .

In addition to the couplings appearing in the Hamiltonian, the HMCA has two more parameters characterizing the molecular-dynamics trajectory: the step size ϵ and the total number of steps n_{MD} before the global accept/reject. The length of the trajectory, $\epsilon \times n_{\text{MD}}$ is adjusted to obtain the best decorrelations. For $\epsilon \times n_{\text{MD}} \sim 1$, the acceptance is only weakly dependent on n_{MD} , and studies of the uncorrected hybrid algorithm applied to QCD (Ref. 10) and tests with the HMCA (Ref. 5) suggest that using $\epsilon \times n_{\text{MD}} \approx 0.7$ is optimal. These tests were done at couplings where the lattice correlation

TABLE I. Acceptance in the global Metropolis step on 8^4 lattices at $\beta=5.3$ as a function of β_{MD} . We show the data for three values of $\kappa=\kappa_{\text{MD}}$, and estimate that $5.29 < \beta_{\text{MD}} < 5.295$ is the best value. The convergence criterion R is defined in (2.6).

ϵ	κ	β_{MD}	$R \times 10^9$	Acceptance (%)
0.04	0.158	5.31	~ 2	54
0.04	0.158	5.29	~ 2	61
0.04	0.158	5.28	~ 2	49
0.0333	0.160	5.31	~ 2	58
0.0333	0.160	5.29	~ 2	68
0.0333	0.160	5.28	~ 2	51
0.04	0.162	5.295	~ 0.1	62
0.04	0.162	5.29	~ 0.1	62

length is approximately one lattice unit. In Sec. III we argue that $\epsilon \times n_{\text{MD}}$ has to be increased as $m_q \rightarrow 0$. Having fixed all the other parameters, the value of ϵ is chosen to maximize $\epsilon \times \text{acceptance}$, the rate of motion through the phase space.

The final attribute of the simulation is the accuracy required in the inversion $\chi = A^{-1}\phi$. This inversion, performed at each step of the trajectory using an iterative algorithm, consumes most of the CPU time. The accuracy criterion often used to test the convergence of the inversion is some variant of

$$\frac{|A\chi - \phi|^2}{|\chi|^2} < R. \quad (2.6)$$

An *ad hoc* choice for R does not guarantee that the final distribution of configurations will be correct. We need a criterion to judge when R is small enough.

For the HMCA one really has two choices to make: the value of R used to evaluate δH and the value used during the trajectory (R_{MD}). For the algorithm to remain exact, the inversions must be accurate when determining δH , and so a stringent criterion for R must be used. The evolution is reversible irrespective of the value of R_{MD} , if one uses a starting guess for χ independent of the values of χ at previous time steps. In practice, we use the previous value for χ as the input guess for the inversion. This means that our inversions must be accurate at each step to maintain detailed balance. Thus we use the same value for R during the evolution and in the evaluation of δH .

Two errors are introduced if R is not chosen small enough. First, δH is not evaluated correctly, and second, the evolution is not reversible. Either of these errors is sufficient, in general, to cause a failure of the identity¹¹

$$\langle e^{-\delta H} \rangle = 1. \quad (2.7)$$

Here the average is taken over all trajectories, whether or not the trajectory is accepted. We find that monitoring this identity allows one to choose a good value for R quite easily. We simply decrease R until the identity is satisfied within the error made in the evaluation of the average. The identity is visibly violated for low convergence criteria. For example, we find that $\langle e^{-\delta H} \rangle$ is 0.67(3), 0.93(2), and 1.00(1) for R values of 10^{-6} ,

2×10^{-7} , and 10^{-8} on a 4×6^3 lattice at $\beta=4.98$ and $m_q=0.025$ (confined phase).

The identity (2.7) holds only if the distribution of configurations has reached equilibrium. Once we have chosen a sufficiently small value for R , we can use the identity to test thermalization of our simulations.

The second criterion we use to choose R is to demand that the error made in evaluating δH be less than some bound. Though closely related to the first criterion, it turns out in practice to be more sensitive. It has the additional advantage that one need only use a few sample trajectories to evaluate the errors, whereas application of the first criterion requires taking an average over a large number of trajectories.

The tuning of parameters discussed above must result in δH lying mostly in the interval $[-1, 1]$ if it is to produce acceptances above 50%. This ~ 1 number, however, is a result of cancellation between three large numbers scaling as the lattice volume (each is $\sim 10^5$ for a 10^4 lattice). Consequently δH is very sensitive to the choice of R . By running a few trajectories with different values of R one can quickly select a safe value. We find that to compute δH to a fixed accuracy, it is necessary to decrease R as the quark mass is reduced. Our results have been obtained by demanding that δH be accurate to less than a percent. On 8^4 lattices, this required decreasing R from 10^{-9} for the heaviest quark masses to 10^{-12} for the lightest ones. To be more specific, in the above example with $R=10^{-6}$, 2×10^{-7} , and 10^{-8} , the values of δH for a particular trajectory were 4.2, 0.98, and 0.041, respectively. The correct answer was ≈ 0.076 , and the convergence criterion required to ensure an accuracy of $< 1\%$ was $R < 10^{-11}$.

From our tests we also deduce that effects of a poor accuracy criterion add coherently along the molecular-dynamics trajectory. Since the number of steps in a trajectory grows as $m_q \rightarrow 0$, one has no choice but to use a rather conservative value for R at small quark masses. It remains to be seen whether the high accuracy required in the calculation of δH proves to be a problem in working with large n_{MD} , or in implementing the HMCA on computers with limited precision.

III. THE COMPUTER TIME REQUIREMENTS OF THE HMCA

The analysis of the previous section leaves the efficiency of the HMCA to be judged by how ϵ depends on the variables m_q , β , and the lattice volume V for a fixed acceptance rate. (For Wilson fermions, the quark mass m_q is a function of the hopping parameter. We use it to give our arguments a general form.) The analytic arguments presented in Refs. 4 and 11, valid for both staggered and Wilson fermions, imply that, for fixed β and m_q , $\epsilon \propto V^{-1/4}$, making the HMCA a $V^{5/4}$ algorithm.

We find that the dependence of ϵ on m_q as $m_q \rightarrow 0$ is a much more important source of critical slowing down for the HMCA. The effective fermion action

$$S_F = \text{Tr} \ln(\Delta^\dagger \Delta), \quad (3.1)$$

is logarithmically singular as $m_q \rightarrow 0$. For finite m_q , how-

ever, one can appeal to analytic continuation and expand the change in the Hamiltonian during molecular-dynamics evolution in a Taylor series:

$$\delta H = \delta H_{\text{gauge}} + n_{\text{MD}} \sum_{n=3}^{\infty} C_n \epsilon^n S_F^{(n)}. \quad (3.2)$$

Here C_n are numerical coefficients, $S_F^{(n)}$ represents all the terms generated by differentiating the fermion action n times with respect to the gauge link matrices U and the factor of n_{MD} arises from integrating the result over the full trajectory. Differentiation of the logarithm produces negative powers of Δ , each successive differentiation adding one more. As $m_q \rightarrow 0$ the operator Δ^{-1} is expected to behave like m_q^{-1} making $S_F^{(n)}$ diverge like m_q^{-n} . The series expansion, therefore, makes sense provided ϵ goes to zero faster than m_q . For a trajectory of length $n_{\text{MD}} = O(\epsilon^{-1})$, the most divergent term as $m_q \rightarrow 0$ is the leading one, i.e., $n=3$. δH behaves like $\epsilon^2 m_q^{-3}$, requiring ϵ to decrease as $m_q^{3/2}$ in order to maintain constant acceptance. Our tests, predominantly with staggered fermions,¹² support this expectation.

The above argument has been made for the full effective fermion action (3.1). Our simulations, however, use a pseudofermion estimator of this action. Nevertheless, it is simple to see that the counting of powers of m_q goes through in the same way.

The efficiency of the HMCA in terms of CPU time requires the estimate of two extra factors.

(a) The number of iterations of the matrix inversion algorithm used to calculate $(\Delta^\dagger \Delta)^{-1} \phi$. This typically depends on the ratio of the largest to the smallest eigenvalue of the matrix Δ and grows as m_q^{-1} when $m_q \rightarrow 0$. (See the following section for more details.)

(b) The autocorrelation between the trajectories in terms of motion through the phase space. The decorrelation time is a function of the lattice correlation length ξ . For QCD, $\xi = \max(\xi_{\text{glue}}, m_\pi^{-1})$, where ξ_{glue} is the correlation length of gluonic degrees of freedom. Therefore $\xi = m_\pi^{-1} \propto m_q^{-1/2}$ for sufficiently light quarks. When the decorrelation time is large compared to the molecular-dynamics trajectory length, random-walk arguments suggest autocorrelations growing as ξ^2 . This is the situation if the trajectory length is held fixed as $m_q \rightarrow 0$. Alternatively, Duane¹³ has proposed use of trajectories of length $O(\xi)$. In this case the system evolves according to the guided walk of the hybrid algorithm, and the autocorrelation between trajectories becomes independent of ξ .

Combining all the factors, we find that when $n_{\text{MD}} = O(\epsilon^{-1})$ is held fixed, the HMCA degrades asymptotically as $V^{5/4} m_q^{-7/2}$ for fixed β . On the other hand, when $n_{\text{MD}} = O(\xi \epsilon^{-1})$ is allowed to grow, we have to go back and rederive the dependence of ϵ on m_q . We obtain $\epsilon \propto m_q^{7/4}$, and the total computer time becomes proportional to $V^{5/4} m_q^{-13/4}$. To illustrate the rapidly degrading performance, we only need to state that a twofold increase in the correlation length ξ , accompanied by a doubling of the lattice dimensions to maintain constant physics, necessitates an increase in the CPU effort by a factor of $2^{12} = 4096$ or $2^{11.5} = 2896$.

The improvement in the asymptotic behavior caused

by taking the trajectory length proportional to ξ is very small. To compare the two options we also have to take into account the relative size of the prefactors, which can only be determined empirically. We do foresee practical problems in using long trajectories on large lattices and at small quark masses, due to the growth of errors in δH mentioned in Sec. II. So far we have only used molecular-dynamics trajectories of length 0.7–1.0, and typical autocorrelation times increased from a few trajectories at the heaviest-quark masses to 50–100 trajectories at the lightest ones. More tests are clearly needed.

Our results differ from a more optimistic proposal of Toussaint.¹⁴ He suggests that for QCD, the decorrelation time (and hence the trajectory length) governing autocorrelation remains fixed as $m_q \rightarrow 0$; the slowing down at small quark mass is already accounted for by the decrease in ϵ . It is our view, however, that the decrease in ϵ is necessary to obtain accurate solutions to the molecular-dynamics equations. Even if we could integrate these equations exactly, we claim that the trajectory length would have to grow as $m_q^{-1/2}$ in order to produce decorrelated configurations.

We have painted a rather pessimistic picture, but to the extent that we can judge from our results, it is not far off the mark. Major improvements can be made by reducing the unknown proportionality constants that have entered the asymptotic analysis above. Some possibilities are as follows.

(a) Devising faster matrix inversion algorithms. The present best iterative algorithms for inverting the fermion matrix only reduce the prefactor, leaving the critical slowing down as $1/m_q$ unaffected. Alternatives designed to counter this, such as multigrid methods, have not yet become viable schemes for QCD in the present range of couplings.

(b) Finding a quicker and more efficient preprocessor than molecular dynamics, or a different discretization of the evolution equations that decreases the magnitude of the error δH allowing an increase in ϵ . Higher-order discretization schemes that change the leading term in δH from $n=3$ to $n=4$ have been considered,¹⁵ but they have not yet become practical for QCD.

IV. IMPROVED ALGORITHMS FOR CALCULATING THE INVERSE AND THE DETERMINANT OF THE DIRAC OPERATOR

The bottleneck in the hybrid Monte Carlo algorithm (HMCA), in approximate algorithms and in calculations of hadron masses and matrix elements is the inversion of the Dirac operator (2.1). In the following we explain the improvements we have made in calculating inverses. We also point out that it is possible to reformulate the determinant so that the pseudofermions live only on half the sites.

The Dirac operator can be written in even/odd block form as

$$\Delta = M \mathbf{1} + \kappa D = \begin{pmatrix} M \mathbf{1} & \kappa D_{eo} \\ \kappa D_{oe} & M \mathbf{1} \end{pmatrix}, \quad (4.1)$$

with $M = 1$ for Wilson fermions. Given a source ϕ , the

fermion propagator χ is obtained by solving the equation

$$\Delta\chi = \phi, \quad \chi = \begin{pmatrix} \chi_e \\ \chi_o \end{pmatrix}, \quad \phi = \begin{pmatrix} \phi_e \\ \phi_o \end{pmatrix}. \quad (4.2)$$

Multiplying both sides by $(M1 - \kappa D)$, the problem separates into

$$(M^2 1 - \kappa^2 D_{eo} D_{oe}) \chi_e = M \phi_e - \kappa D_{eo} \phi_o, \quad (4.3)$$

and

$$\chi_o = (\phi_o - \kappa D_{oe} \chi_e) / M. \quad (4.4)$$

The separation eliminates $O(\kappa)$ terms from the matrix to be inverted, and can be viewed as first-order preconditioning of the fermion matrix.¹⁶ Equation (4.3) can be solved using one's favorite matrix inversion algorithm—minimal residual, conjugate gradient, or anything else—and the full propagator, if needed, can be reconstructed at the end using (4.4). This trick substantially reduces the memory requirement of the code and roughly doubles the speed (in terms of iterations) at which the algorithm converges.

For this simple scheme to work, it is only necessary that the diagonal part of Δ be a multiple of the identity matrix, while the off-diagonal parts connecting even and odd sites contain all the interactions. This condition holds for the improved fermion action obtained using the $\sqrt{3}$ block transformation where the interaction terms connect sites at distance 1 and $\sqrt{3}$. In fact, Eqs. (4.3) and (4.4) were solved in the hadron spectrum calculation of Ref. 17.

For staggered fermions ($M = m$, $\kappa = \frac{1}{2}$, $r = 0$, $\gamma_\mu \rightarrow \eta_\mu$) even/odd decomposition is the standard procedure. Furthermore, the matrix in (4.3) is Hermitian and positive definite. In such cases the conjugate-gradient algorithm works well, and we have not found any tricks to improve upon it.

For heavy Wilson fermions the minimal residual algorithm converges in approximately the same number of iterations as the conjugate-gradient method, and it requires only half the number of arithmetic operators per iteration.¹⁸ For the generic equation $A\chi = \phi$ and the remainder $r = \phi - A\chi$, this algorithm iterates the equations

$$\alpha_n = (Ar_n, r_n) / |Ar_n|^2, \quad \chi_{n+1} = \chi_n + \omega \alpha_n r_n, \quad (4.5)$$

$$r_{n+1} = r_n - \omega \alpha_n Ar_n,$$

with the relaxation parameter ω set to 1. The algorithm works its way towards the final solution by minimizing the positive-definite quadratic form $|r_n|^2$ at each step. The minimization is attempted along the steepest-descent direction of the quadratic form $(r_n, A^{-1}r_n)$, which is in general complex. The difference in the two forms causes the iterates to get stuck when Ar becomes orthogonal to r . The failure is strongly dependent on the gauge configuration, the source vector ϕ , and the initial guess ϕ_0 . On general grounds, we expect the failure to depend on the overlap of ϕ with the eigenvectors corresponding to the small eigenvalues of A . Physically this situation arises at small quark masses.^{19,17} Our experience is that

for Wilson fermions (4.3) can be solved using the minimal residual algorithm for quark masses $m_q \geq m_s$, where m_s is the strange-quark mass.

Taking $\omega > 1$ overrelaxes the minimization step in (4.5). This overrelaxation is similar in spirit to the tuning of the diagonal element in the ILU decomposition of the matrix.¹⁸ Such overrelaxation is not helpful for the conjugate gradient or other similar methods where orthogonality of vectors at different iteration numbers has to be maintained. We find that choosing $\omega = 1.3$ decreases the number of iterations needed for convergence by $\approx 30\%$. Moreover, ω does not have to be tuned very sensitively to ensure the best convergence; a value in the range 1.15–1.50 provides reasonable improvement.

We have followed the idea of ‘‘polynomial preconditioning’’²⁰ further by converting (4.3) to

$$(M^4 1 - \kappa^4 D_{eo} D_{oe} D_{eo} D_{oe}) \chi_e = (M^2 1 + \kappa^2 D_{eo} D_{oe})(M \phi_e - \kappa D_{eo} \phi_o). \quad (4.6)$$

Again, $\omega \approx 1.3$ is optimal when solving this with the overrelaxed minimal residual algorithm. Solving (4.6) is $\approx 10\%$ slower than solving (4.3) in the region where both algorithms converge, but it has the advantage of being more stable. It shows no signs of failure even at $m_q \approx m_s/2$. At still smaller quark masses where this algorithm fails to converge, one can attempt further preconditioning. Otherwise it is always possible to revert back to the conjugate-gradient method applied to (4.3).

Quantitatively, for a single matrix inversion, the improved algorithm (even/odd split, minimal residual with overrelaxation) is about six times as fast in terms of CPU as a straightforward application of the conjugate-gradient method to (4.2). This factor applies, for example, to propagator calculations needed for evaluating hadron masses. In the molecular-dynamics evolution of the gauge configurations, however, we need to evaluate $X = (\Delta^\dagger \Delta)^{-1} \phi$. We achieve this in two steps:

$$\Delta^\dagger Y = \phi, \quad \Delta X = Y, \quad (4.7)$$

noting that $\Delta^\dagger(\kappa, r) = \Delta(-\kappa, -r)$. The first inversion enhances the proportion of small eigenvalue modes, so the second inversion is more time consuming. The conjugate-gradient method requires only one inversion to solve for X , and the final speed up compared to it is then about a factor of 2.5.

The ILU decomposition¹⁸ produces improvement comparable to our method. This is not surprising since both methods eliminate $O(\kappa)$ terms by preconditioning. The appearance of triangular matrices makes the ILU method inefficient to vectorize and parallelize. A high-order polynomial preconditioning, on the other hand, is simple to vectorize and parallelize. The big advantage for parallel machines (such as the Connection Machine) is that the ratio of multiplication by D (which requires nearest-neighbor communications only) to global dot products of vectors needed in the calculation of α (which are comparatively slow) is increased.

The algorithm we are proposing does nothing to alleviate the problem of critical slowing down; the number of iterations still increases as m_q^{-1} in the limit $m_q \rightarrow 0$. To

attack critical slowing we need to know something about the eigenvalue distribution of the matrix being inverted. Both Δ and Δ^\dagger have the same spectrum and so can be preconditioned using the same technique. Simple variants such as cycling ω through a set of values with iteration number, or adjusting the coefficients of the polynomial preconditioner²⁰ can be considered. We have not yet explored these variants. Fourier acceleration²¹ applied to (4.2) provides a factor-of-3 improvement, and should be tested for (4.6).

We now turn to a reformulation of the fermion determinant. Using the identity

$$\begin{aligned} \det(\Delta) &= \det \begin{pmatrix} \Delta_{ee} & \Delta_{eo} \\ \Delta_{oe} & \Delta_{oo} \end{pmatrix} \\ &= \det(\Delta_{oo}) \det(\Delta_{ee} - \Delta_{eo} \Delta_{oo}^{-1} \Delta_{oe}), \end{aligned} \quad (4.8)$$

we can rewrite the determinant as

$$\det(\Delta) = \det_e(\Delta_e) = \det_o(\Delta_o), \quad (4.9)$$

$$\Delta_e = M^2 1 - \kappa^2 D_{eo} D_{oe}, \quad \Delta_o = M^2 1 - \kappa^2 D_{oe} D_{eo}. \quad (4.10)$$

The subscript on \det denotes the subspace (even or odd sites) on which it is evaluated.

Physically the determinant contains contributions from all closed fermion vacuum-polarization loops. All such loops have even perimeter length on a hypercubic lattice. The determinant depends only on the trace of the fermion loop, which is independent of the choice of starting point. In a hopping-parameter expansion of $\det_e(\Delta_e)$, the hops are of length 2, whereas the hops in an expansion of $\det(\Delta)$ are of unit length. A loop of perimeter $2L$ can be covered in $2L$ hops of length 1 or in L hops of length 2. This is exactly what the result (4.9) expresses; Kramer's expansion has been simplified by restricting the starting point to be only on even (or odd) sites.

Clearly a similar expression can be written for $\det(\Delta^\dagger)$. Thus we can rewrite the pseudofermion functional integral (2.2) as

$$\det(\Delta) \det(\Delta^\dagger) = \int D\phi_e D\phi_e^* \exp[-\phi_e^\dagger (\Delta_e^\dagger \Delta_e)^{-1} \phi_e]. \quad (4.11)$$

It is straightforward to modify the HMCA to use the form (4.11). The advantage of simulating (4.11) is that the pseudofermion fields need to be defined only on half the lattice sites, thus reducing memory requirements. Though there is no reduction in the number of arithmetic operations per iteration, the phase space to be covered has been cut down by a factor of 2.

This formulation is very similar to that used for staggered fermions. The only difference is that in case of staggered fermions there is no need to take the absolute square of Δ_e . One can simply simulate the functional integral:

$$\det(\Delta) = \int D\phi_e D\phi_e^* \exp(-\phi_e^\dagger \Delta_e^{-1} \phi_e). \quad (4.12)$$

V. COMPARING THE HMCA WITH HYBRID AND LANGEVIN ALGORITHMS

In this section we address the following question: Is it more efficient to use an exact algorithm or to use an approximate algorithm and extrapolate to zero-step size? We compare the HMCA, hybrid, and Langevin algorithms using Wilson loops of size up to 3×3 as probes. The loops are measured on 6×8^3 lattices at $\beta=5.5$ and $\kappa=0.15$. This lattice is effectively at zero temperature for these parameters since, as shown in the next section, the finite-temperature phase transition for $N_t=6$ is at $\kappa \approx 0.1565$. We also compare our results to those of Fukugita, Oyanagi, and Ukawa,²² who use a second-order Langevin algorithm on a $9^3 \times 18$ lattice. Finally, we discuss the possibility of using an approximate algorithm at finite-step size at the end of the section.

All the algorithms spend approximately 90% of the time performing matrix inversions. These inversions can be carried out using the same optimal method in all cases. We, therefore, compare the efficiency of the algorithms by counting the number of matrix inversions required to achieve the same statistical accuracy for the Wilson loop data.

The results from our calculations are presented in Table II. For the HMCA we average over 400 trajectories of length 0.8, using $\epsilon=0.04$, which gives an acceptance rate of 70%. For the hybrid algorithm we also use 400 trajectories of length 0.8, with five values of ϵ : 0.40, 0.16, 0.08, 0.04, and 0.02. In both cases we use a bin size of 40 trajectories to estimate errors. The results of the first-order Langevin algorithm are extracted from runs of Langevin time $\tau=25$. We use two-step sizes, $\delta\tau=0.005$ and $\delta\tau=0.0025$, with measurements every $\tau=\frac{1}{8}$ and $\frac{1}{16}$, respectively. The data are grouped into bins of size $\tau=2.5$ to estimate errors.

For all the algorithms we find long thermalization times. We have not done a detailed autocorrelation study. Conservatively, we discard 400 trajectories for both the hybrid and the HMCA, and 25 time units for the Langevin runs, which we judge to be about twice the thermalization times. The HMCA results presented here differ by 2–3 standard deviations from the preliminary results presented in Ref. 12. The latter used shorter runs, consequently the lattices were not sufficiently thermalized.

For small enough ϵ ($\epsilon \leq 0.08$), the hybrid algorithm gives Wilson loop values larger than those from the HMCA. This is in accord with our implementation of the leapfrog scheme.⁴ On the contrary, making ϵ too large disorders the system. The $\epsilon=0.16$ results show that the disorder rapidly increases with the loop size. Since the approach to the $\epsilon=0$ values is not monotonic, we cannot use large ϵ results to extrapolate to zero step size. The data for $\epsilon < 0.08$ are close enough to the HMCA results that it may be possible to extrapolate them to $\epsilon=0$. The statistical errors on our data, however, do not allow us to test this possibility.

Upon running the HMCA with $\epsilon=0.08$, we discovered that the acceptance dramatically drops to only 15% compared to its value of 70% at $\epsilon=0.04$. The measured Wil-

son loops (up to 3×3) for the hybrid algorithm, however, have changed little between $\epsilon=0.08$ and $\epsilon \leq 0.04$. This discrepancy indicates that large fermion loops are making a significant contribution to δH . The finite-step size distortion is thus a function of the length scale of the observable, with longer-distance scale observables having larger distortions.

Since it is the long-distance observables that carry the physics of the theory, it is not obvious whether finite-step-size simulations are in the same universality class as QCD. Given this uncertainty, it is our conclusion that to be on the safe side hybrid simulations have to be run at values of ϵ where the HMCA has a high acceptance rate, and not be fooled by apparent agreement of short-distance observables. In such a case the HMCA is obviously comparable in efficiency to the hybrid algorithm.

For the first-order Langevin algorithm, the finite-step-size errors are very large and any extrapolation to zero-step size would be questionable. Thus one has to use a second-order formalism to be competitive with either the hybrid method or the HMCA. We quote the second-order algorithm results obtained by Fukugita, Oyanagi, and Okawa²² on $9^3 \times 18$ lattices in Table II(c). They use three values of the step size, $\delta\tau=0.02, 0.01, 0.005$, and

estimate the $\delta\tau=0$ values using linear extrapolation. The second-order data also show a strong dependence on the step size, e.g., the 3×3 Wilson loop expectation value changes by over 100% between $\delta\tau=0.02$ and 0.0. Therefore, we believe that Langevin simulations with at least three values of $\delta\tau$ are needed to reliably extrapolate to the $\delta\tau=0$ limit.

The statistical errors for the HMCA and the linearly extrapolated Langevin results are comparable. The agreement between the two sets of results is reasonable, given that the finite-size effects are not the same (the lattice volumes and fermion boundary conditions are different) and that the extrapolation involved in the Langevin data is substantial compared to the discrepancy. The data show the extent to which Langevin simulations underestimate the effect of dynamical fermions without the $\delta\tau \rightarrow 0$ extrapolation. The amount of extrapolation is increasing with the loop size and a distortion is likely to set in at some stage making a simple linear extrapolation insufficient.

To compare the efficiency of the second-order Langevin algorithm and the HMCA, we must take into account the different lattice volumes: the Langevin simulation used a lattice roughly four times as large. Since

TABLE II. (a) Comparison of the Wilson loop data with two flavors of dynamical Wilson fermions at $\kappa=0.15$ and $\beta=5.5$. The first five columns give results obtained using the hybrid algorithm and the last column lists the HMCA results. N_{inv} is the total number of matrix inversions required. (b) Same as (a) except that the first two columns contain results for the first-order Langevin algorithm. (c) Same as (a) except that the first three columns give results obtained by Fukugita, Oyanagi, and Okawa (Ref. 22) using a second-order Langevin algorithm and the fourth column is their linear extrapolation of the data to $\delta\tau=0$.

N_{inv}	(a)					HMCA 8000
	$\epsilon=0.40$ 800	$\epsilon=0.16$ 2000	$\epsilon=0.08$ 4000	$\epsilon=0.04$ 8000	$\epsilon=0.02$ 16 000	
1×1	0.4062(2)	0.5381(8)	0.5388(9)	0.5384(8)	0.5377(8)	0.5359(8)
1×2	0.1555(2)	0.3076(11)	0.3120(12)	0.3122(11)	0.3115(11)	0.3091(11)
1×3	0.0604(1)	0.1786(11)	0.1836(12)	0.1841(11)	0.1835(11)	0.1812(11)
2×2	0.0233(1)	0.1150(11)	0.1217(12)	0.1224(11)	0.1218(11)	0.1198(12)
2×3	0.0035(1)	0.0457(8)	0.0508(9)	0.0513(8)	0.0510(8)	0.0495(8)
3×3	0.0002(1)	0.0134(4)	0.0162(5)	0.0164(4)	0.0163(5)	0.0155(4)
(b)						
N_{inv}	$\delta\tau=0.005$	$\delta\tau=0.0025$	HMCA			
	5000	10 000	8000			
1×1	0.5029(9)	0.5231(8)	0.5359(8)			
1×2	0.2707(12)	0.2943(10)	0.3091(11)			
1×3	0.1476(11)	0.1680(11)	0.1812(11)			
2×2	0.0893(10)	0.1077(10)	0.1198(12)			
2×3	0.0312(8)	0.0420(7)	0.0495(8)			
3×3	0.0074(5)	0.0117(4)	0.0155(4)			
(c)						
N_{inv}	$\delta\tau=0.02$	$\delta\tau=0.01$	$\delta\tau=0.005$	$\rightarrow \delta\tau=0.0$	HMCA	
	2000	6000	8000	16 000	8000	
1×1	0.5137(9)	0.5257(7)	0.5330(9)	0.5389(10)	0.5359(8)	
1×2	0.2799(8)	0.2958(7)	0.3053(9)	0.3131(13)	0.3091(11)	
1×3	0.1548(8)	0.1689(7)	0.1779(8)	0.1848(11)	0.1812(11)	
2×2	0.0947(9)	0.1080(9)	0.1165(12)	0.1228(12)	0.1198(12)	
2×3	0.0339(5)	0.0419(5)	0.0474(7)	0.0511(8)	0.0495(8)	
3×3	0.0084(6)	0.0120(5)	0.0144(7)	0.0161(6)	0.0155(4)	

the smaller lattice is larger than the correlation length, one can view the larger lattice as four independent smaller lattices. The HMCA will then require only a fourth as many steps on the $9^3 \times 18$ lattice to reach the same statistical errors.

Combining the factors from the number of inversions (16000/8000) and the increase in statistics with lattice volume (4) with the $V^{-1/4}$ dependence of the HMCA step size ϵ (0.7), we conclude that the HMCA is about five times faster than the second-order Langevin algorithm on a $9^3 \times 18$ lattice. It might be objected that this crude analysis has been biased in favor of the HMCA by the assumption that the same inversion accuracy is required in the two schemes. Fukugita, Oyanagi and Ukawa²² claim that one can get away with a much less stringent convergence criterion R when using the Langevin algorithm, and therefore conclude that the Langevin algorithm is more efficient than the HMCA. From Table II(c) we note that Langevin simulations at finite $\delta\tau$ roughly correspond to an exact simulation with heavier dynamical fermions. So it is to be expected that a less stringent criterion for R can be used, but that as one reduces $\delta\tau$ the criterion will have to be tightened up. This effect should be added to the balance sheet on the side of the Langevin simulations.

The most conservative conclusion we draw from this exercise is that the HMCA is at least as efficient as the hybrid and Langevin algorithms on moderate-size lattices. The HMCA gains if an extrapolation to zero-step size is mandatory. (More so in calculations of the hadron spectrum and matrix elements, because the external quark propagators have to be calculated on only one set of background-field configurations for the HMCA.) However, it has not been resolved whether such a limiting procedure is really necessary, i.e., whether the configurations at finite $\delta\tau$ are in the same universality class as QCD. The fact that we find good agreement between the exact HMCA results, the hybrid-algorithm results with comparable ϵ , and linear extrapolation of the Langevin algorithm results is encouraging. On the other hand, we have to prove that such an agreement holds at all length scales. We need an unequivocal answer, but a numerical “proof” is hard to establish. In the best circumstances, universality will hold when the value of the limiting parameter (e.g., $\delta\tau$) is smaller than a certain nonzero bound. Then the effect of a nonzero-step size can be absorbed in the renormalization of the bare parameters of the theory (this includes the number of fermion flavors N_f). In such a happy situation we can forget about exact algorithms and extrapolations, and instead focus on showing that the dimensionless predictions of the theory, such as mass ratios, become constant as $g \rightarrow 0$ and $m_q \rightarrow 0$. On the other hand, if the bound is very sensitive to the value of m_q , then one has no recourse but to use an exact algorithm.

VI. FINITE-TEMPERATURE PHASE TRANSITION

We present results for the location of the finite-temperature phase transition for $N_f=4$ and 6 with β in the range 5.2–5.6. We have studied these transitions on $6^3 \times 4$ and $8^3 \times 6$ lattices. It may seem that these lattices

have too small a spatial volume to allow a study of finite-temperature effects. Our studies with staggered fermions suggest, however, that we will make only a small error in the location of the phase transition. What we cannot ascertain is the nature of the phase transition, i.e., whether there is a true transition (of either first or higher order) or simply a crossover. To answer this would require a finite-volume scaling analysis. Thus, when in the following we use the expression “phase transition,” we are simply identifying the region of parameter space where there is *at least* a rapid crossover and where there *may* be a phase transition.

We have studied the phase transition using hysteresis loops, varying κ at fixed β . The results are shown in Figs. 1–3. For $N_f=4$ most points are an average over 500 trajectories of length $\epsilon \times n_{MD} = 0.7$ –1.0. The exceptions are data at $\beta=5.3$ and $\kappa=0.150$ –0.156 (1000 trajectories each), $\kappa=0.157$, 0.158 (2000 trajectories each), and $\kappa=0.1575$ (5000 trajectories). The $N_f=6$ results are averaged over 200 trajectories for all points.

There is a rapid crossover visible at all values of β . Although the hysteresis loops do not show the behavior characteristic of a first-order transition, that may be an artifact of frequent tunneling caused by the small spatial volume. We do have some evidence for a two-phase system at $\kappa=0.1575$ and $\beta=5.3$, both from flip-flops in time histories and from histograms of the plaquette and the Wilson line.¹² However, this is only suggestive and confirmation will require a finite-volume analysis.

All the estimates for the transition couplings are compiled in Table III and shown in Fig. 4. We have included the results of Ukawa, who also uses the HMCA, at $\beta=4.5$ and 5.0 on 4×8^3 lattices.²³ The table also gives the known values of κ_c (Refs. 23 and 24) for two degenerate flavors of Wilson fermions.

With κ_c defined as the limit where the pion mass van-

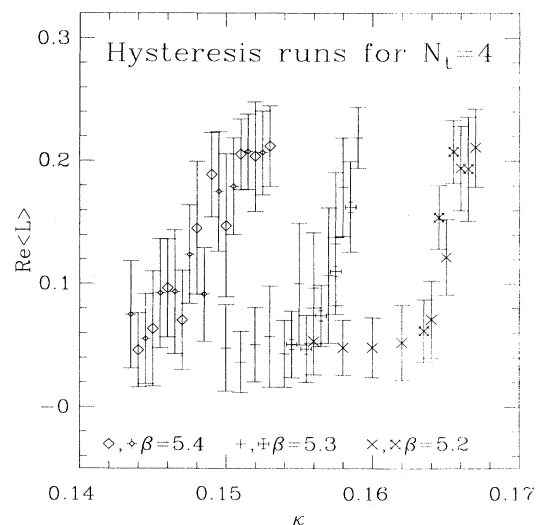


FIG. 1. The Wilson line as a function of κ during hysteresis runs on 4×6^3 lattices at $\beta=5.2, 5.3, 5.4$. Heating and cooling parts of the cycles are denoted by different symbols.

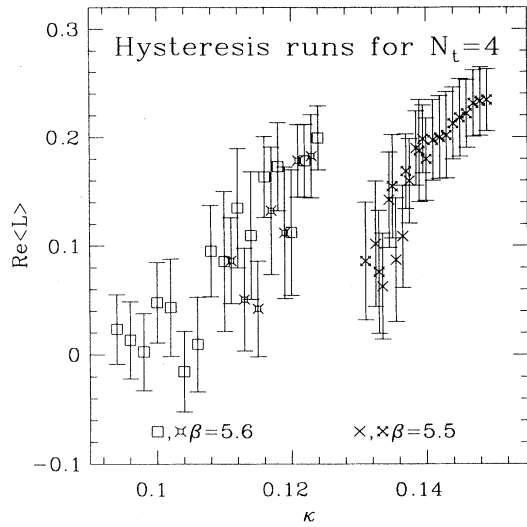


FIG. 2. Same as Fig. 1, but for $\beta=5.5, 5.6$.

ishes in the $T=0$ theory, the data for $N_t=4$ show that the phase transition line does not cross the κ_c line down to $\beta=4.5$. This feature was pointed out by Fukugita, Ohta, and Ukawa²⁵ and Ukawa,²³ and is in contrast with what has been observed for staggered fermions. Naively, in going from low to high temperatures, one might expect a “chiral-symmetry-restoration” transition along the κ_c line. According to this view, the finite-temperature transition line should cross the κ_c line. This is apparently what fails to happen for small N_t . Figure 4 shows that the situation is already considerably different for $N_t=6$. We argue in the following that the anomalous behavior at small N_t is an artifact of strong coupling and should

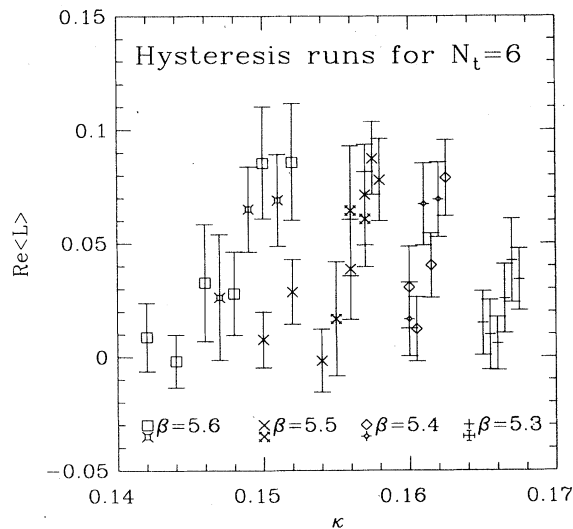


FIG. 3. Same as Fig. 1, but for $\beta=5.3, 5.4, 5.5, 5.6$ on 6×8^3 lattices.

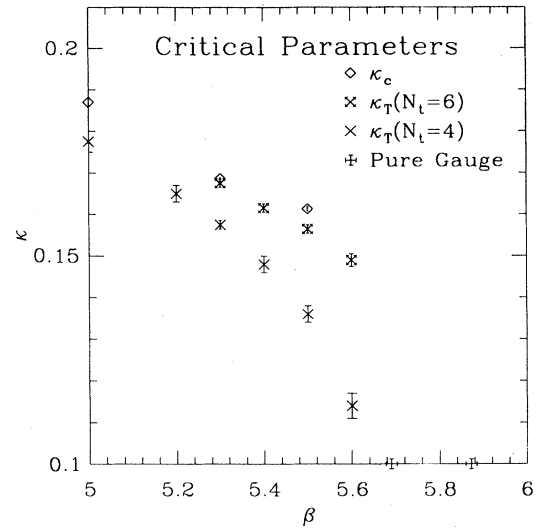


FIG. 4. The values of κ_T at the $N_t=4$ and 6 finite-temperature phase transition as a function of β . Also shown are the κ_c values corresponding to the chiral limit at $T=0$.

disappear as the lattice spacing goes to zero.

First we need a definition of κ_c which works at any temperature, since κ_c may vary with temperature in addition to depending on β . This is in contrast with staggered fermions for which the chiral limit is at $m_q=0$ for all lattice spacings and temperatures. Using the divergence of the nonsinglet axial-vector current, Iwasaki, Tsuboi, and Yoshić²⁶ define κ_c at finite temperature as the point where

$$\lim_{z \text{ large}} \frac{\left\langle \sum_{x,y,t} \nabla_z A_z(x,y,z,t) \pi(0) \right\rangle}{\left\langle \sum_{x,y,t} \pi(x,y,z,t) \pi(0) \right\rangle} \quad (6.1)$$

vanishes. In their first quenched calculation with this definition, they find κ_c to be almost independent of the temperature. Unfortunately, no results for the full QCD theory using this definition of κ_c exist yet.

TABLE III. Estimates for the location of the finite-temperature phase transition and κ_c for two flavors of Wilson fermions.

N_t	β	κ_T	κ_c
4	4.5	0.195–0.20	0.206(2)
4	5.0	0.175–0.18	0.187(1)
4	5.2	0.165(2)	
4	5.3	0.1575(10)	0.1686(3)
4	5.4	0.148(2)	
4	5.5	0.136(2)	0.1613(2)
4	5.6	0.114(3)	
6	5.3	≈ 0.1675	0.1686(3)
6	5.4	0.1615(10)	
6	5.5	0.1565(10)	0.1613(2)
6	5.6	0.1490(15)	

Second we need to understand how well different prescriptions for estimating κ_c coincide with each other. Wilson fermions have a clear interpretation in terms of spin and flavor degrees of freedom, and thus the construction of interpolating field operators to probe the physics is straightforward. This is achieved, however, at the cost of explicitly breaking chiral symmetry. As a consequence, the chiral limit is really only defined in the continuum limit. The customary definition of an effective chiral limit as the point where the pion mass vanishes defines a line $\kappa_c(\beta)$ in the $[\kappa, \beta]$ plane. But naive lattice operators do not satisfy chiral Ward identities on this line, and both perturbative and nonperturbative corrections have to be made before defining physical matrix elements. Even then the requisite identities are satisfied only up to corrections vanishing as powers of the lattice spacing. Therefore, different prescriptions for κ_c using different chiral properties can disagree by $O(a)$ terms.

This discussion is meant to emphasize that at strong coupling, where one is far from the continuum limit, there is no unique definition of κ_c . The only real avenue for progress is to do calculations for $N_f > 6$.

VII. SCREENING IN THE HEAVY $q\bar{q}$ POTENTIAL

Confinement in pure gauge QCD is visualized as the formation of a color-electric flux-tube between two isolated static quarks. The strength of this linear confinement is characterized by the string tension σ . A simple phenomenological spin-independent potential which is consistent with present calculations is $V = \alpha/r + \sigma r$. The picture changes when dynamical quarks are included. At short distances, the dynamical quarks essentially renormalize the gauge coupling, changing the value of α . At long distances, it becomes favorable for the string to break due to the creation of a $q\bar{q}$ pair from the vacuum. Linear confinement disappears, and the resulting configuration is a pair of mesons. This string-breaking scale is characterized by an effective quark mass m_{eff} which is a function of the quark mass m_q appearing in the action. For very heavy quarks $m_{\text{eff}} \approx m_q$, and one expects little deviation from the pure gauge potential for

separation $r < 2m_{\text{eff}}/\sigma$, except for a renormalization of the scale. For light quarks the rise in the potential is expected to cease at the length scale of confinement, about 1 fm.

Previous studies of screening of the $q\bar{q}$ potential^{22,27–29} have all used approximate algorithms. They have observed a flattening of the potential at large distances, while the major effect of dynamical fermions remains an overall shift in the gauge coupling. We have extended these calculations by using an exact algorithm, and by studying the behavior as a function of the quark mass for fixed β . Unfortunately, it is hard to make a direct comparison with the earlier calculations since that requires (a) extrapolation of data from approximate algorithms which typically underestimate the effect of dynamical quarks, and (b) a careful disentangling of the screening effect produced by dynamical quarks from that produced by a nonzero temperature.

We have measured up to 4×4 Wilson loops on 8^4 lattices at six values of κ : 0.156, 0.158, 0.160, 0.162, 0.165, 0.167. The gauge coupling was restricted to $\beta = 5.3$ so as to avoid substantial finite-temperature effects even at the smallest quark mass. At $\kappa = 0.167$, we encountered a severe critical slowing down of the algorithm caused by light modes of the Dirac operator. We could not measure the autocorrelation reliably, but it might be as large as a few hundred trajectories. This is an example of the shortcomings of present fermion algorithms at small quark masses as discussed in Sec. III.

Our results are presented in Table IV. The data is not sufficient to allow an extraction of the potential. Instead we study a simpler problem that exposes the screening effect of the dynamical quarks: the behavior of Wilson loops as a function of the quark mass. For every Wilson loop we have calculated in the presence of dynamical quarks, we find the matching β of the pure gauge theory that produces the same answer. The shift in the gauge coupling $\Delta\beta$ is tabulated as a function of the loop size and the quark mass in Table V. Pure gauge theory simulations needed for this matching have been performed on 8^4 lattices with β in the range [5.46, 5.7] in increments of 0.02 (0.025 for $\beta > 5.6$).

TABLE IV. Wilson loop expectation values for two flavors of Wilson fermions as a function of the loop size and κ . The data is for 8^4 lattices with antiperiodic boundary conditions for the Dirac operator in all four directions. N_{traj} is the total number of trajectories in the data sample.

	$\langle \text{Wilson loop} \rangle$ for $\beta = 5.3$ and $N_f = 2$					
κ	0.156	0.158	0.160	0.162	0.165	0.167
N_{traj}	900	1340	1020	1350	1710	2370
1×1	0.4847(3)	0.4909(2)	0.4957(2)	0.5016(4)	0.5152(6)	0.5299(6)
1×2	0.2472(3)	0.2548(3)	0.2606(3)	0.2679(5)	0.2852(5)	0.3047(4)
1×3	0.1274(3)	0.1337(2)	0.1386(2)	0.1451(4)	0.1604(4)	0.1785(4)
1×4	0.0658(2)	0.0703(2)	0.0740(2)	0.0788(3)	0.0905(3)	0.1050(4)
2×2	0.0712(2)	0.0770(2)	0.0812(2)	0.0870(4)	0.1014(4)	0.1195(4)
2×3	0.0214(1)	0.0245(1)	0.0267(1)	0.0299(2)	0.0385(3)	0.0507(3)
2×4	0.0065(1)	0.0079(1)	0.0089(1)	0.0105(2)	0.0149(2)	0.0219(2)
3×3	0.0040(1)	0.0051(1)	0.0058(1)	0.0070(1)	0.0105(1)	0.0169(2)
3×4	0.0007(1)	0.0011(1)	0.0013(1)	0.0016(1)	0.0031(1)	0.0060(1)
4×4	0.0000(1)	0.0001(1)	0.0002(1)	0.0004(1)	0.0007(1)	0.0020(1)

TABLE V. The shift in gauge coupling necessary to match the Wilson loop results for two flavors of dynamical quarks with those for the pure gauge theory, as a function of the loop size and κ .

$\Delta\beta$ from linear interpolation						
κ	0.156	0.158	0.160	0.162	0.165	0.167
1×1	0.155(1)	0.177(1)	0.196(1)	0.218(2)	0.266(2)	0.320(2)
1×2	0.160(1)	0.184(1)	0.203(1)	0.226(2)	0.277(1)	0.335(1)
1×3	0.162(1)	0.187(1)	0.206(1)	0.229(2)	0.282(1)	0.341(1)
1×4	0.162(1)	0.188(1)	0.208(1)	0.231(2)	0.284(1)	0.344(2)
2×2	0.165(1)	0.192(1)	0.212(1)	0.236(2)	0.289(1)	0.353(2)
2×3	0.166(1)	0.197(1)	0.216(1)	0.240(2)	0.295(2)	0.363(2)
2×4	0.165(2)	0.200(2)	0.218(2)	0.243(2)	0.298(2)	0.368(2)
3×3	0.171(2)	0.202(2)	0.222(2)	0.246(3)	0.301(2)	0.375(2)
3×4	0.167(8)	0.209(7)	0.224(4)	0.243(7)	0.306(4)	0.381(3)
4×4			0.209(27)	0.266(13)	0.308(13)	0.401(6)

For heavy quarks, the lowest-order hopping-parameter expansion predicts the coupling-constant renormalization to be

$$\Delta\beta(1 \times 1 \text{ loop}) = 48\kappa^4. \quad (7.1)$$

Since our observed values of $\Delta\beta(1 \times 1 \text{ loop})$ are much larger than the predictions of (7.1), it is clear that the quark masses we have used are sufficiently light to invalidate the lowest-order hopping-parameter expansion.

For large Wilson loops, the effect of string breaking should show up in our data as $\Delta\beta \rightarrow \infty$ when area $\rightarrow \infty$, with the rate of growth of $\Delta\beta$ with the loop size a function of the quark mass. Were the only effect of dynamical fermions a renormalization of the gauge coupling, $\Delta\beta$ would become a constant for loops beyond a certain size typified by the confinement scale. As the results in Table V and Fig. 5 show, for heavy quarks our data cannot rule out this latter possibility; a rise in $\Delta\beta$ with the loop size is barely discernible. But for light quarks (our largest value

of $\kappa=0.167$ roughly corresponds to the strange-quark mass) we do see the expected screening behavior— $\Delta\beta$ increases with the area of the loop.

The above results have implications for other physical observables such as the hadron spectrum. To see a departure from the behavior obtained in the quenched approximation, the potential felt by the quarks at the hadronic scale must be modified in shape and not just shifted by a renormalization of the gauge coupling. We see that this does not happen if the sea quarks are much heavier than the physical strange quark. Unfortunately, it is close to this mass that all the currently popular algorithms, approximate or exact, run out of steam.

What happens if we look at $\Delta\beta(\kappa_c)$? κ_c is expected to be determined by the behavior of the theory at the hadronic scale. At 1 loop the Λ parameter of QCD depends on the bare coupling as

$$\Lambda \sim \exp\left[-\frac{4\pi^2\beta}{33-3N_f}\right], \quad (7.2)$$

so in the asymptotic scaling region we expect

$$\Delta\beta(\kappa_c) = \left[\frac{2N_f}{33-2N_f}\right]\beta. \quad (7.3)$$

We have $\kappa_c=0.1686(3)$ from our hadron spectrum calculation.²⁴ It roughly agrees with the quenched simulation result, $\kappa_c=0.1692(1)$ at $\beta=5.7$ (Ref. 30). The matching value, $\Delta\beta(\kappa_c) \approx 0.4$, is well short of the answer 0.64 anticipated by (7.3) and tells us that we are far away from the asymptotic scaling region.

We make one final observation. Upon extrapolating our results to the chiral limit, we find $\Delta\beta(1 \times 1 \text{ loop}) \approx 0.38$ and $\Delta\beta(2 \times 2 \text{ loop}) \approx 0.42$. These values bracket $\Delta\beta(\kappa_c)$ and imply that the hadronic scale at present couplings is about one to two lattice spacings—a stark reminder of the coarseness of the lattice.

VIII. COMPARISON OF WILSON AND STAGGERED FERMION RESULTS

With the limited data available we can make a crude comparison between results for two flavors of Wilson and staggered fermions at $\beta=5.3$. For Wilson fermions we




FIG. 5. $\Delta\beta - \Delta\beta(1 \times 1 \text{ loop})$ vs loop area for the six values of κ at $\beta=5.3$. No error bars are plotted since the data are given in Table V.

use the quark mass

$$m_w = \ln \left[1 + \frac{1}{2} \left(\frac{1}{\kappa} - \frac{1}{\kappa_c} \right) \right]. \quad (8.1)$$

With $\kappa_c = 0.1686(3)$, the $N_f = 4$ transition for two flavors of Wilson fermions corresponds to $m_w \approx 0.76T_c$. On the other hand, the $N_f = 4$ transition for two flavors of staggered fermions is at staggered quark mass $m_S \approx 0.14T_c$ (Ref. 31). The quark mass renormalization is not identical for the two fermion schemes. In perturbation theory $m_w = (1 + 1.33/\beta + \dots)m_S$ (Ref. 32), while in the approximate scaling region ($\beta \geq 6.0$) quenched calculations typically give $m_w \approx 2m_S$ (Ref. 33). The difference observed here is, however, much larger and shows that $\beta = 5.3$ is not close to the scaling region.

A related piece of information is the value of the ρ mass extrapolated to the chiral limit for the two fermion schemes. In lattice units, $M_\rho \approx 1.3$ for staggered fermions,³⁴ while $M_\rho \approx 0.5$ for Wilson fermions.²⁴ Thus if we use M_ρ to set the scale, we find a much larger lattice spacing for staggered fermions than for Wilson fermions. The staggered-fermion scale is closer to that expected of quenched simulations at $\beta = 5.3$.

The scale of the theory is governed by the gauge coupling β and the vacuum-polarization effects of dynamical quarks. The above comparison shows that the use of staggered fermions underestimates the effect of dynamical quarks at $\beta = 5.3$. This is not unexpected, since the staggered-fermion flavor symmetry is exact only in the limit of vanishing lattice spacing. At finite lattice spacing one can roughly explain the difference by considering the effective number of staggered-fermion flavors to be less than its naive value and/or the effective quark mass to be greater than m_q . Both adjustments, decreasing the number of effective flavors and increasing the effective quark mass, reduce the effect of the quark loops.

What is perhaps surprising in the above analysis is the size of the difference between the two types of fermions. Based on quenched calculations, a reliable signature of the onset of scaling (to within 10% accuracy) is the agreement between Wilson and staggered-fermion results. Clearly we have a long way to go before seeing such scaling.

IX. CONCLUSIONS

We have analyzed the performance of the HMCA for Wilson fermions and find that it lives up to its billing of combining the best features of the molecular dynamics and the Langevin algorithms. Just as in the case of staggered fermions, removing the small-step-size errors of the hybrid preprocessor leads to little loss in performance; acceptance rates of $\approx 60-70\%$ are obtained in the global accept/reject step on 8^4 lattices without having to make the step size smaller than that used in the uncorrected hybrid algorithm. However, the HMCA is not a panacea as the $V^{5/4}m_q^{-13/4}$ scaling of the required computer time shows.

We have presented efficient algorithms for calculating the inverse and determinant of the Dirac operator for Wilson fermions. The inversion algorithm achieves comparable efficiency to the best ILU algorithm while being very simple to implement on vector and parallel machines. Progress remains to be made in reducing the critical slowing down as $m_q \rightarrow 0$.

The decision whether to use the HMCA or an approximate algorithm is not resolved by our study. In a particular example, the HMCA outperformed the Langevin algorithm, but it is crucial to determine the mass dependence of the latter to make a comparison useful for future research. An approximate algorithm will suffer from the same critical slowing down in matrix inversions. The possible advantage over the HMCA may come in the number of steps taken between statistically independent configurations. We doubt whether approximate algorithms can gain overwhelmingly over the HMCA dependence of $n_{\text{MD}} \propto V^{1/4}m^{-9/4}$. The crucial issue, however, is whether there is a need for the expensive extrapolation to zero-step size, or whether the approximate algorithm simulates a theory which lies in the same universality class as QCD.

Data for the location of the finite-temperature transition for $N_f = 4$ and 6 over the interesting region of β is given. The transition appears to be discontinuous, but in the absence of a finite-volume study we can make no definite conclusions.

Our data show the expected screening behavior in the $q\bar{q}$ potential. It also shows that the effect of dynamical quarks on the hadron spectrum will be hard to quantify unless simulations are done at quark masses comparable to or smaller than the strange-quark mass. Extrapolations from heavier masses towards the chiral region are unlikely to yield results quantitatively different from the quenched calculations.

Our study also highlights the large difference between results obtained using Wilson and staggered fermions at $\beta = 5.3$. This gauge coupling is too strong and lies well outside the scaling region. We have evidence that the staggered-fermion flavor symmetry is badly broken, and we believe that the Wilson fermion chiral symmetry is also badly broken. Realistic simulations at weaker couplings, smaller quark masses and bigger lattices will have to await superior methods and/or faster computers.

ACKNOWLEDGMENTS

This calculation was performed on the Cray X-MP's at the Los Alamos National Laboratory, the DD Division of CERN, the San Diego Supercomputer Center, the Pittsburgh Supercomputing Center, the National Magnetic Fusion Energy Computer Center at Livermore, and on Trace minisupercomputers at Multiflow Computer Inc. We thank these organizations for their support. We acknowledge useful conversations with Giorgio Parisi and Akira Ukawa on matrix inversion algorithms.

- ¹D. Weingarten, in *Field Theory on the Lattice*, proceedings of the International Symposium, Fermilab, Batavia, Illinois, 1988, edited by A. Kronfeld *et al.* [Nucl. Phys. B (Proc. Suppl.) (in press)].
- ²S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, Phys. Lett. B **195**, 216 (1987).
- ³R. Gupta, in *Field Theory on the Lattice*, proceedings of International Symposium, Seillac, France, 1987, edited by A. Billoire *et al.* [Nucl. Phys. B (Proc. Suppl.) **4**, 562 (1988)].
- ⁴R. Gupta, G. Kilcup, and S. Sharpe, Phys. Rev. D **38**, 1278 (1988); **38**, 1288 (1988).
- ⁵K. Bitar, A. D. Kennedy, R. Horsley, S. Meyer, and P. Rossi, Nucl. Phys. **B313**, 348 (1989); **B313**, 377 (1989).
- ⁶R. Gupta, G. Kilcup, A. Patel, and S. Sharpe (in preparation).
- ⁷S. A. Gottlieb, W. Liu, D. Toussaint, R. L. Renken, and R. L. Sugar, Phys. Rev. D **35**, 2531 (1987).
- ⁸N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, J. Chem. Phys. **21**, 1087 (1953).
- ⁹M. Creutz, H. Gausterer, and S. Sanielevici, Phys. Rev. D **39**, 689 (1989).
- ¹⁰S. Duane and J. Kogut, Nucl. Phys. **B275**, 398 (1986).
- ¹¹M. Creutz, Phys. Rev. D **38**, 1228 (1988).
- ¹²R. Gupta, in *Field Theory on the Lattice* (Ref. 1).
- ¹³S. Duane, Nucl. Phys. **B275**, 398 (1986).
- ¹⁴D. Toussaint, University of Arizona Report No. AZPH-TH/89-6 (unpublished).
- ¹⁵A. D. Kennedy, SCRI Report No. FSU-SCRI-88-132, 1988 (unpublished); J. B. Kogut, E. Kovacs, and D. K. Sinclair, Nucl. Phys. **B290**, 431 (1987).
- ¹⁶P. de Forcrand, H. Haraguchi, H. Hege, V. Linke, A. Nakamura, and I. Stamatescu, Phys. Rev. Lett. **58**, 2011 (1987).
- ¹⁷P. de Forcrand, R. Gupta, S. Güsken, K.-H. Mütter, A. Patel, K. Schilling, and R. Sommer, Phys. Lett. B **200**, 143 (1988).
- ¹⁸Y. Oyanagi, Comput. Phys. Commun. **42**, 333 (1986).
- ¹⁹P. Rossi, C. T. H. Davies, and G. P. Lepage, Nucl. Phys. **B297**, 287 (1988).
- ²⁰O. G. Johnson, C. A. Micchelli, and G. Paul, SIAM J. Num. Anal. **20**, 362 (1983).
- ²¹G. G. Batrouni, G. R. Katz, A. S. Kronfeld, G. P. Lepage, B. Svetitsky, and K. G. Wilson, Phys. Rev. D **32**, 2736 (1985); G. Katz, G. Batrouni, C. Davies, A. Kronfeld, G. Lepage, P. Rossi, B. Svetitsky, and K. Wilson, *ibid.* **37**, 1589 (1988).
- ²²M. Fukugita, Y. Oyanagi, and A. Ukawa, Phys. Rev. D **36**, 824 (1987).
- ²³A. Ukawa, in *Field Theory on the Lattice* (Ref. 1).
- ²⁴A. Patel, in *Field Theory on the Lattice* (Ref. 1).
- ²⁵M. Fukugita, S. Ohta, and A. Ukawa, Phys. Rev. Lett. **57**, 1974 (1986).
- ²⁶Y. Iwasaki, Y. Tsuboi, and T. Yoshié, University of Tsukuba Report No. UTHEP-190, 1988 (unpublished).
- ²⁷M. Campostrini, K. J. M. Moriarty, J. Potvin, and C. Rebbi, Phys. Lett. B **193**, 78 (1987).
- ²⁸P. de Forcrand, Nucl. Phys. **A461**, 361 (1987); P. de Forcrand, V. Linke, and I. O. Stamatescu, *ibid.* **B304**, 645 (1988); I. O. Stamatescu, in *Field theory on the Lattice* (Ref. 1).
- ²⁹M. Grady, D. K. Sinclair, and J. Kogut, Phys. Lett. B **200**, 149 (1988).
- ³⁰APE Collaboration, P. Bacilieri *et al.*, Nucl. Phys. B (to be published).
- ³¹S. A. Gottlieb, W. Liu, D. Toussaint, R. L. Renken, and R. L. Sugar, Phys. Rev. Lett. **59**, 1513 (1987).
- ³²H. Hamber and C. M. Wu, Phys. Lett. **133B**, 351 (1983).
- ³³R. Gupta, G. Guralnik, G. Kilcup, A. Patel, S. Sharpe, and T. Warnock, Phys. Rev. D **36**, 2813 (1987).
- ³⁴S. A. Gottlieb, W. Liu, D. Toussaint, R. L. Renken, and R. L. Sugar, Phys. Rev. D **38**, 2245 (1988).