AL

December 1994

HLRZ 94-76

WUB 94-41

Tuning Up an Oldtimer: Hybrid Monte Carlo with Wilson Fermions

K. Schilling^{a,c}, A. Frommer^b, V. Hannemann^b, Th. Lippert^a, and B. Nöckel^b

^a Fachbereich Physik, Bergische Universität, D-42097 Wuppertal, Germany

^b Fachbereich Mathematik, Bergische Universität, D-42097 Wuppertal, Germany

^c HLRZ c/o Forschungszentrum Jülich, D-52425 Jülich, Germany

Abstract

We show that the BiCGStab inversion algorithm helps to speed up by 50 % the computation of the fermionic force inside the Hybrid Monte Carlo (HMC) simulation of full QCD with Wilson fermions, in the chiral regime of small quark masses.

SCU 9506

CERN LIBRARIES, GENEVA

SCAN-9501346

Tuning Up an Oldtimer: Hybrid Monte Carlo with Wilson Fermions[†]

K. Schilling^{ac}, A. Frommer^b, V. Hannemann^b, Th. Lippert^a, and B. Nöckel^a

^aDepartment of Physics, University of Wuppertal, D-42097 Wuppertal, Germany

^bDepartment of Mathematics, University of Wuppertal, D-42097 Wuppertal, Germany

^cHLRZ c/o KFA, D52425 Jülich, Germany

We show that BiCGStab inversion algorithm helps to speed up by 50 % the computation of the fermionic force inside the Hybrid Monte Carlo (HMC) simulation of full QCD with Wilson fermions, in the chiral regime of small quark masses.

1. INTRODUCTION

The computational bottleneck of HMC is the evaluation of the molecular dynamics force, which requires the solution of the equations

$$M^\dagger M x = \phi \quad (1)$$

during each individual step along the molecular dynamics history of the field system, where the components of the vector ϕ are independent Gaussian random numbers. In the literature, the solution to the equation is normally reached by relaxation algorithms like minimal residual (MR) or conjugate gradient (CG).

At this point, the Wilson fermion discretization of QCD is particularly demanding, for three reasons: (a) the size of the Wilson-fermion matrix is four times larger than the corresponding Kogut-Susskind matrix, (b) the Wilson fermion matrix is a non-hermitian complex matrix, $M = 1 - \kappa D$ and (c) its non-diagonal part D is controlled by the hopping parameter κ , that needs tuning towards its critical value κ_c . This is the reason, why so little progress has been reached in $T = 0$ QCD with Wilson fermions.

At that stage one might take two attitudes to reach improvement: 1. Forget about HMC with the inversion problem altogether and try for alternatives[3]; or 2. launch another attempt to achieve more speed in the very inversion problem within HMC.

In this contribution we will take the second approach. Recently, a new and very promising relaxation method has been proposed by van der Vorst[2]. He designed BiCGStab to treat non-hermitian matrices directly. In this paper, we will demonstrate numerically, that this method is highly competitive in the chiral regime of lattice QCD[4].

We benchmark CGNE (which solves the 'normal equation' $M^\dagger M x = M^\dagger \phi$), MR and BiCGStab first on quenched gauge configurations at $\beta = 6.0$ on a 16^4 lattice. Next, we shall compare the efficiency of BiCGStab to CG and MR in a real HMC situation, at $\beta = 5.4$ on a 8^4 lattice.

For the details of BiCGStab, we refer the reader to ref.[2]. We just quote its net computational costs in table 1.

Method	M.M.	S.P.	AXPYs	St.
MR	1	2	2	3
CGNE	2	2	3	4
BiCG	2	2	5	6
CGS	2	2	7	7
BiCGStab	2	4	6	7

Table 1

Arithmetic and storage requirements: M.M = number of matrix multiplies, S.P = number of scalar products, St = vectorstorage needed.

[†]Work supported by the Deutsche Forschungsgemeinschaft under grant No. Schi 257/3-2, Schi 257/1-4, and by EC contract CHRX-CT92-0051.

Figure 1. CPU-time improvement factors on the 32-node CM5 as functions of κ normalized to $t_{BiCGStab}$ on a configuration which is representative for the ensemble.

2. RESULTS

2.1. Quenched Benchmark

We generated a sample of 10 decorrelated quenched $SU(3)$ gauge configurations, properly thermalized at $\beta = 6.0$. The propagators were computed on one source per configuration for the normal equation. For definiteness, we chose point-like sources, which are known to require more iterations than other sources. The inversion algorithms were implemented in data-parallel style on a 32-node CM5.

We chose the residual norm to be $r = 10^{-12}$ and we convinced ourselves that the resulting estimators from the different algorithms lie sufficiently close to each other, by looking at the difference between their outputs.

The observed improvement factor in terms of real computer time is shown in Fig. 1. Throughout the κ -region, we find BiCGStab to beat the other methods. These results hold in particular for the chiral regime, where, compared to MR, we gain a factor of 1.2 - 1.5 using BiCGStab.

Figure 2. $\log_{10} r$ as function of the number of iteration steps.

In Fig. 2, we display the iteration history for the three algorithms at $\kappa = 0.155$ on the same configuration as in Fig. 1. It shows that the convergence rate for the residual of MR equals that of CGNE. BiCGStab shows a convergence rate more than twice as fast as CGNE. This is related to the fact that BiCGStab is better conditioned than CGNE as it avoids squaring the condition number. Note that the residuals in BiCGStab exhibit the well-known fluctuations; they are not overly large, however.

We confirmed these observations by extension to our entire sample of ten statistically independent pure gauge configurations.

2.2. HMC Benchmark

We worked on a 8^4 lattice at $\beta = 5.4$ and performed HMC simulations at four values of κ approaching κ_c . At each κ -value, we performed 400 thermalization steps. In view of the expense in CPU-time, we only measured on ten configurations separated by 10 HMC sweeps each. Table 2 summarizes the HMC parameters used. We worked with single precision accuracy. Therefore, we expected larger numerical deviations in the ac-

κ	β	# of MD-steps	dt	r
0.158	5.4	25	0.02	10^{-8}
0.160	5.4	25	0.02	10^{-8}
0.161	5.4	25	0.02	10^{-8}
0.162	5.4	25	0.02	10^{-8}

Table 2

HMC parameters for both the full determinant and the determinant on an even subspace.

culated vectors of the investigated relaxation methods. We shall discuss these deviations and again we shall show that the final estimates x_f for the true solution x from the different methods are sufficiently close.

The κ -values chosen lie close to $\kappa_c = 0.1645$ as quoted by Gupta[5] for a 16^4 lattice. Taking this value as an approximation on our 8^4 system, $\kappa = 0.162$ would correspond to a quark mass $m_q \simeq 0.05$. An 8^4 lattice will of course exhibit severe finite size effects for such a small mass.

It has become standard to use the Wilson fermion determinant on an even subspace in order to accelerate the full QCD simulation. The determinant can be reformulated:

$$\det(M) = \det \begin{pmatrix} 1 & -\kappa D_{eo} \\ -\kappa D_{oe} & 1 \end{pmatrix}. \quad (2)$$

The computation of the block determinant leads to

$$\det(M) = \det(1 - \kappa^2 D_{eo} D_{oe}) = \det(M_e). \quad (3)$$

The simulation of the square of the determinant can be done on the even subspace sparing a factor of two in memory requirements and improving the condition of the numerical problem.

Using CG we have to multiply by the hermitian matrix $M_e^\dagger M_e$ during each iteration step. With MR and BiCGStab we again apply the two-step procedure as discussed before. We have employed the same HMC run parameters as in the simulation with the full determinant. In Fig. 3, we present the normalized iteration numbers. We find BiCGStab to converge again faster than CG and MR. The CPU-time improvement on the 32-node CM5 is nearly a factor of 1.5 compared to

Figure 3. Iteration numbers N in a HMC simulation as functions of κ normalized to $N_{BiCGStab}$ on a sample of ten configurations. The determinant is simulated on the even subspace.

both competing methods. As in the last section, we cross-checked the resulting difference norms and found the deviations being smaller than 10^{-6} .

REFERENCES

1. Duane, S., Kennedy, A. D., Pendleton, and B. J. Roweth, D., Phys. Lett. **B 195** 216 (1987).
2. van der Vorst, H.: BiCGStab: A Fast and Smoothly Converging Variant of Bi-CG for the Solution of Nonsymmetric Linear Systems, SIAM J. Sc. Stat. Comp. **13**, 631–644 (1992).
3. M. Lüscher, these proceedings.
4. A. Frommer *et al*, Int.J.Mod.Phys., *in print*.
5. Gupta, R. et al., Phys. Rev. D **44** 3272 (1991).