

Exploring the HMC trajectory-length dependence of autocorrelation times in lattice QCD

Harvey Meyer^{a,*}, Hubert Simma^a, Rainer Sommer^a, Michele Della Morte^b,
Oliver Witzel^b, Ulli Wolff^b

Alpha Collaboration

^a *Deutsches Elektronen Synchrotron DESY, Platanenallee 6, D-15738 Zeuthen, Germany*

^b *Humboldt-Universität, Institut für Physik, Newtonstrasse 15, D-12489 Berlin, Germany*

Received 13 June 2006; received in revised form 18 July 2006; accepted 23 August 2006

Available online 26 September 2006

Abstract

We study autocorrelation times of physical observables in lattice QCD as a function of the molecular dynamics trajectory length in the hybrid Monte-Carlo algorithm. In an interval of trajectory lengths where energy and reversibility violations can be kept under control, we find a variation of the integrated autocorrelation times by a factor of about two in the quantities of interest. Trajectories longer than conventionally used are found to be superior both in the $N_f = 0$ and $N_f = 2$ examples considered here. We also provide evidence that they lead to faster thermalization of systems with light quarks.

© 2006 Published by Elsevier B.V.

PACS: 12.38.Gc; 11.15.Ha; 05.10.Ln

Keywords: Wilson lattice QCD; Hybrid Monte-Carlo; Schrödinger functional; Trajectory length

1. Introduction

Hybrid Monte-Carlo (HMC) [1] is the most widely used algorithm to simulate non-Abelian gauge theories in the presence of $N_f > 0$ quark flavors. For Wilson-type quark actions it was found, in its most straightforward application, to become prohibitively expensive in the regime of phenomenological interest (with two very light quarks) [2]. Recent advances [3–6] however have led to a substantial lowering of the CPU time of simulations. These more sophisticated forms of HMC typically have a number of parameters that have to be tuned to guarantee an efficient simulation. One parameter that is however common to all HMC algorithms is the length τ of the trajectories of ‘molecular dynamics’ evolution, at the end of which an accept/reject step is performed to correct for any finite-step-size errors.

The algorithm efficiency’s dependence on τ has only been rarely and partially studied in the literature [7–10]. The reason for this is easy to understand. If we neglect the cost of evaluating the Hamiltonian, performing trajectories of length $\tau = 2$ represents the same computational cost as twice as many trajectories of length 1: indeed, as discussed in Appendix A, the average energy violations along a trajectory at fixed step-size are weakly dependent on the length within a reasonable range, and the reversibility violations increase rather slowly (see however [11]). Hence, with a high acceptance rate in both cases, to discriminate between these two running modes one has to determine the autocorrelation times of the observables of interest with some accuracy and reliability. This in turn requires very high statistics, which one is normally not ready to invest into parameter optimization. In this letter we study this question in various situations where high statistics can be achieved, and yet some realistic features of difficult simulations are present. A study of the $2d$ XY-model [12] in the early days of HMC followed a similar strategy, and a thorough analytic investigation of the free field theory case can be found in [13].

* Corresponding author.

E-mail address: harvey.meyer@desy.de (H. Meyer).

Table 1
The parameters of the physical systems considered

	$L^3 \times T$	BF	N_f	$1/d\tau$	β	κ	a/r_0 [Ref.]	aM_{PS}
A	8^4	A	0	30	6.086	–	0.16 [23]	–
	12^4	A	0	30	6.364	–	0.11 [23]	–
	16^4	A	0	30	6.57	–	0.08 [23]	–
B	$8^3 \times 32$	0	0	50	6.0	0.1338	0.19 [23]	0.38(2)
C	$24^3 \times 32$	0	2	32	5.3	0.1355	0.16 [25]	0.325(10)
D	12^4	$A \rightarrow 0$	2	32	5.6215	0.136665	0.11 [24]	–

A boundary field (BF) 0 means that it is vanishing, while A refers to ‘point A’ of Ref. [18].

Whether in the end our conclusions carry over to other forms of HMC algorithms will require direct testing. However we show that the dependence of the algorithm’s efficiency on τ can be significant, and a factor around two in computing time can easily be wasted in the case of an unfortunate choice of trajectory length.

2. Data

We employ the hybrid Monte-Carlo algorithm [1]. The Hamiltonian governing the evolution of the molecular dynamics is

$$H = \frac{1}{2} \sum_{\mu, x} \text{Tr} \{ \Pi_{\mu}(x) \Pi_{\mu}(x) \} + S[U], \quad (1)$$

where the momenta Π_{μ} are traceless Hermitian matrices. The gauge fields are thus updated according to $U'_{\mu}(x) = \exp[i \Pi_{\mu}(x) d\tau] U_{\mu}(x)$. In the $N_f = 2$ simulations, two pseudo-fermions are used to stochastically represent the determinant of the $O(a)$ improved Wilson fermions: the fermionic part of S results from even–odd [14] and mass-preconditioning [4] the Hermitian Dirac operator. The gauge part of S is the Wilson plaquette action.

When $N_f = 2$, we use the Sexton–Weingarten integration scheme [15], with a step-size four times larger for the fermionic forces; the quoted step-size $d\tau$ always denotes the largest one. A complete update cycle with trajectory length τ performs an integer number $\tau/d\tau$ of such steps followed by the Metropolis decision. The system has then advanced by τ molecular dynamics (MD) units, and in the rest of this paper all quantities referring to Monte Carlo time are given in MD units. If, for instance, successive measurements of observables are separated by M trajectories, an autocorrelation function $\Gamma(i)$ arises where i refers to *successive measurements*. From it we define the integrated autocorrelation time, that is directly relevant for the statistical error, in our MD units as

$$\tau_{\text{int}} = M\tau \left[\frac{1}{2} + \sum_{i \geq 1} \rho(i) \right], \quad \rho(i) = \Gamma(i) / \Gamma(0). \quad (2)$$

In numerical estimates the sum has to be truncated. If we specify a window W (in MD units) this amounts to the restriction $i \leq W/(M\tau)$. We refer the reader to [16] for the definition of $\Gamma(i)$, in particular for derived quantities, such as an effective mass. The error bars shown in plots of ρ are computed using Eq. (E.11) of Ref. [3].

We investigate the trajectory length dependence of autocorrelation times in three different systems (A, B, C: see Table 1). All simulations were done in 32-bit arithmetic, except for system C (64-bit). The system D will be used as a playground for thermalization. The lattice spacing is known in units of r_0 [19] at 1% level for the $N_f = 0$ runs and 5% for the $N_f = 2$ runs. All lattices have Schrödinger functional boundary conditions. Since we will be using somewhat longer trajectories than is usual, the question of reversibility violations arises. We check for this in the standard way (see e.g. [7]) by monitoring the Hamiltonian variation $\langle |\delta H|_{\leftrightarrow} \rangle$ under the following operations: a trajectory ‘forward’, reversing the sign of the momenta, and integrating back to the starting point.

The observables we will focus on are those which typically have autocorrelation times much larger than that of the plaquette. One observable (defined already in the pure gauge theory) is $dS/d\eta$; the inverse of its expectation value defines the Schrödinger functional renormalized coupling constant [18]. The angle η parametrizes the Abelian boundary field, the ‘point A’ of the fundamental domain corresponding to $\eta = 0$ where the boundary gauge field is proportional to $\lambda_3 + \sqrt{3}\lambda_8$ (λ_a is a Gell–Mann matrix). The others are fermionic: the correlators $f_A(x_0)$ and $f_P(x_0)$ correspond to the propagation of a quark and an antiquark from a boundary to a point in the bulk of the lattice where the axial current or pseudoscalar density annihilates them [17]. Finally we also consider f_1 , which is the amplitude for boundary-to-boundary propagation through the lattice. It serves to normalize the other correlators. In particular, $Z_A f_A(x_0) / \sqrt{f_1} \sim F_{\text{PS}} e^{-(x_0 - T/2)M_{\text{PS}}}$ holds far from the boundaries, where M_{PS} and F_{PS} correspond to the pseudoscalar mass and decay constant [17].

2.1. HMC in the pure gauge theory

We start by investigating autocorrelation times for gluonic quantities in the pure-gauge HMC (system A), where we have three different lattice spacings at constant physical parameters (see Table 2). The acceptance rates, as well as $\langle \delta H^2 \rangle$, are given (the step-size is the same in all these runs); we are clearly in the regime of high acceptance P_A where $\langle \delta H^2 \rangle \simeq 2\pi(1 - P_A)^2$ [28]. The dependence of $\langle \delta H^2 \rangle$ on τ is moderate. Therefore the dependence of P_A is even weaker, but note that the effect of lower acceptance at larger τ is automatically taken into account in the autocorrelation of observables. The reversibility violations as measured by $\langle |\delta H|_{\leftrightarrow} \rangle$ grow as $\sqrt{\tau}$ here, or even more slowly; the reader is referred to Appendix A for a more detailed discussion of $\langle \delta H^2 \rangle$ and $\langle |\delta H|_{\leftrightarrow} \rangle$. We now focus on the

Table 2

The autocorrelation times (in units of MD time) for $dS/d\eta$, the acceptance, the energy and reversibility violations for different choices of trajectory length. W is the window in MD time over which the autocorrelation time has been accumulated. From top to bottom: $L = 8, 12$ and 16 , while $La/r_0 = \text{const}$. Runs in single precision

A: pure gauge	$\tau = 1/2$	$\tau = 1$	$\tau = 2$	$\tau = 4$
$L = 8$ $\tau_{\text{int}}[dS/d\eta]$	6.10(35) $_{W=64}$	5.90(30) $_{W=36}$	3.14(10) $_{W=36}$	3.92(12) $_{W=44}$
$\tau_{\text{trunc}}[dS/d\eta]_{W=25}$	5.0(2)	4.2(1)	2.75(7)	3.46(9)
Acceptance [%]	97	97	97	96
$\langle \delta H^2 \rangle / (Ld\tau)^4$	0.789(6)	0.953(6)	1.221(8)	1.71(1)
$10^4 \cdot \langle \delta H _{\leftrightarrow} \rangle$	2.0	2.6	3.5	4.9
$L = 12$ $\tau_{\text{int}}[dS/d\eta]$	9.5(1.0) $_{W=86}$	6.4(5) $_{W=63}$	3.34(18) $_{W=34}$	4.36(24) $_{W=44}$
$\tau_{\text{trunc}}[dS/d\eta]_{W=25}$	5.8(3)	4.7(2)	3.1(1)	3.8(2)
Acceptance [%]	94	93	93	91
$\langle \delta H^2 \rangle / (Ld\tau)^4$	0.941(6)	1.134(8)	1.40(1)	1.92(3)
$10^4 \cdot \langle \delta H _{\leftrightarrow} \rangle$	5.1	6.4	8.8	12.2
$L = 16$ $\tau_{\text{int}}[dS/d\eta]$	6.5(1.0) $_{W=52}$	5.8(6) $_{W=52}$	4.2(4) $_{W=40}$	6.0(4) $_{W=68}$
$\tau_{\text{trunc}}[dS/d\eta]_{W=25}$	5.3(5)	4.2(5)	3.5(2)	4.8(3)
Acceptance [%]	89	88	86	84
$\langle \delta H^2 \rangle / (Ld\tau)^4$	1.08(1)	1.31(2)	1.58(2)	2.10(4)
$10^4 \cdot \langle \delta H _{\leftrightarrow} \rangle$	9.1	12	16	21

observable $dS/d\eta$ defined in [18]. We emphasize that, unlike the plaquette, this is an observable dominated by long-distance fluctuations and it typically has an autocorrelation time significantly larger than one.

The most direct way to evaluate the performance of different τ -choices for a given observable is to compare the corresponding normalized autocorrelation function, $\rho(t)$. This is done on Fig. 1, for trajectory lengths $1/2, 1, 2, 4$. There is a marked difference between, say, $\tau = 1/2$ and $\tau = 2$ in favor of the latter, at all MD-time separations. The choice $\tau = 4$ is slightly superior to $\tau = 2$ at short MD-time separation at the largest lattice spacing, while the opposite is true at the finest lattice spacing. Note *en passant* that at the shortest time separation, ρ is much smaller for $\tau = 2$ than for $\tau \leq 1$. In fact, the $\tau = 2$ autocorrelation function shows a significant non-monotonic behavior, as has been observed by us for hybrid overrelaxation algorithms, and also in a simple model studied in [13].

For a run of given MD-time length and measurement frequency, the error squared of the observable is proportional to $1/\tau_{\text{int}}$. When τ_{int} is large compared to the measurement frequency, it corresponds to the area under the curves shown in Fig. 1. In practice, a window W must be chosen where to stop the summation. This may be chosen in a self-consistent way which balances statistical against systematic errors. We use by default the prescription described in [16]. However, to rate relative performances of algorithms, we also compute here an integrated autocorrelation time τ_{trunc} where the summation is truncated at a fixed window $W = 25$. It turns out that typically 80% of the true τ_{int} has by then been accumulated; the uncertainty on τ_{trunc} is almost half of that on the full τ_{int} and the hierarchy between the different trajectory lengths is at any rate maintained (Table 2).

Fig. 2 then illustrates the dependence of the truncated autocorrelation time on the trajectory length τ . One sees that this quantity is minimized somewhere around $\tau = 2$, and this conclusion is largely independent of the lattice spacing in the range considered. Overall we see a variation by a factor two of τ_{trunc}

Table 3

The autocorrelations time (in units of MD time) of two primary observables (f_1 and f_P) and of the effective pseudoscalar mass and decay constant in the middle of the lattice, the acceptance, the energy and reversibility violations for different choices of trajectory length. Runs in single precision

B: quenched, $8^3 \times 32$	$\tau = 1/2$	$\tau = 2$	$\tau = 4$
$\tau_{\text{int}}[f_1]$	40(5) $_{W=270}$	24(4) $_{W=184}$	20(3) $_{W=172}$
$\tau_{\text{int}}[f_P]$	75(20) $_{W=245}$	44(8) $_{W=316}$	32(4) $_{W=256}$
$\tau_{\text{int}}[m_{\text{PS}}^{\text{eff}}(T/2)]$	25(4)	13.5(1.6)	12.1(1.2)
$\tau_{\text{int}}[f_{\text{PS}}^{\text{eff}}(T/2)]$	33(6)	19.0(2.6)	14.7(1.6)
Acceptance [%]	98	98	97
$\langle \delta H^2 \rangle / (L^3 T d\tau^4)$	0.827(5)	1.32(1)	1.87(3)
$10^4 \cdot \langle \delta H _{\leftrightarrow} \rangle$	4.9	9.9	11

for $1/2 \leq \tau \leq 4$. Generally speaking this is a substantial variation which translates directly into a corresponding speed-up of simulations whose cost is dominated by the HMC.

2.2. HMC in quenched QCD

As our next observables we now consider fermionic correlators. Although we are ultimately interested in dynamical, large-volume simulations (system C), we first investigate the autocorrelation times in system B, which, roughly speaking, is a quenched version of system C with in addition the spatial extent divided by 3. The computing time is now dominated by the measurements. The run-length is $t_{\text{run}} \simeq 32000$ in total (four independent lattices, called ‘replicas’, were simulated). We investigate the range $\tau = 1/2$ to $\tau = 4$, within which we see hardly any variation of the acceptance, and $\langle |\delta H|_{\leftrightarrow} \rangle$ scales roughly with $\sqrt{\tau}$: Table 3 shows this, as well as some relevant autocorrelation times.

The comparison of autocorrelation functions for f_1 and f_P is done on Fig. 3. The advantage of $\tau = 4$ and $\tau = 2$ over $\tau = 1/2$ is both substantial and statistically significant. This is reflected in a factor of about two in the integrated autocorrelation times. This gain is at least as marked in physical quantities such as the effective pseudoscalar mass and decay constant.

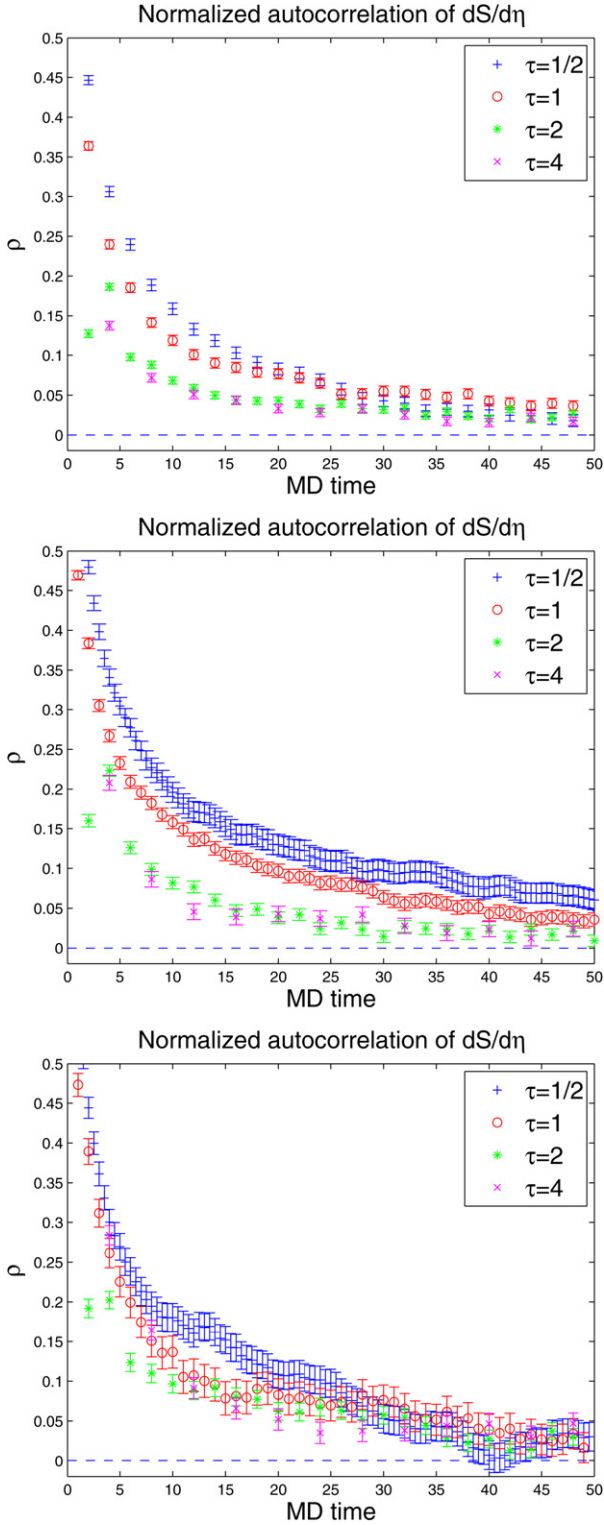


Fig. 1. Autocorrelation function of $dS/d\eta$ in a pure gauge HMC (system A), for different trajectory lengths (from top to bottom: $L = 8, 12$ and 16).

2.3. Mass-preconditioned HMC in $N_f = 2$ QCD

We now come to $N_f = 2$ dynamical simulations in a spatial volume of $(2 \text{ fm})^3$, and consider the same fermionic correlators as in the previous section. The quark mass is around m_s , the




Fig. 2. The truncated autocorrelation time of $dS/d\eta$ in the pure gauge system (A), as a function of the trajectory length τ and for different lattice spacings.

Table 4

The autocorrelations time (in units of MD time) of two primary observables (f_1 and f_P) and of the effective pseudoscalar mass and decay constant in the middle of the lattice, the acceptance, the energy and reversibility violations for different choices of trajectory length. The value of $\langle \delta H^2 \rangle / (L^3 T d\tau^4)$ marked by a * was obtained for one of the replicas; the other had one spike of $\delta H \simeq 2000$. Runs in double precision

C: $N_f = 2, \kappa = 0.1355$	$\tau = 1/2$	$\tau = 2$
$\tau_{\text{int}}[f_1]$	30(15) $_{W=145}$	24(8) $_{W=134}$
$\tau_{\text{trunc}}[f_1]_{W=50}$	23(5)	15(3)
$\tau_{\text{int}}[f_P]$	45(20) $_{W=185}$	26(10) $_{W=134}$
$\tau_{\text{trunc}}[f_P]_{W=50}$	28(7)	18(4)
$\tau_{\text{int}}[m_{\text{PS}}^{\text{eff}}(T/2)]$	9.5(2.5)	6.8(1.5)
$\tau_{\text{int}}[f_{\text{PS}}^{\text{eff}}(T/2)]$	14(4)	4.2(8)
Acceptance [%]	90	91
$\langle \delta H^2 \rangle / (L^3 T d\tau^4)$	0.147(4)*	0.164(6)
$10^4 \cdot \langle \delta H _{\leftrightarrow} \rangle$	1.0	2.5

strange quark mass, and we have two values of τ , $1/2$ and 2 . The total statistics in each case is about $t_{\text{run}} = 4000$ (2 replicas were simulated).

Note that the acceptance is practically the same in both runs. The autocorrelation functions of f_1 and f_P are compared on Fig. 4. Naturally the error bars are now larger; nonetheless a statistically significant advantage of the run at $\tau = 2$ is seen at MD-time separations up to about 30. This is confirmed when one looks at the truncated autocorrelation time with a window of $W = 50$. The autocorrelation function of f_1 is in fact surprisingly similar to that in the quenched simulation, Fig. 3. The resulting autocorrelation times are given in Table 4. In spite of large uncertainties, there is a significant reduction of $\tau_{\text{int}}[f_{\text{PS}}^{\text{eff}}(T/2)]$ when going from $\tau = 1/2$ to $\tau = 2$.

2.3.1. Thermalization with light quarks

Since we observe that autocorrelation times of long-distance observables are reduced by the use of longer trajectories, we also expect their thermalization to be accelerated by the latter. Consider the thermalization of system D (see Table 1). The exercise here is to

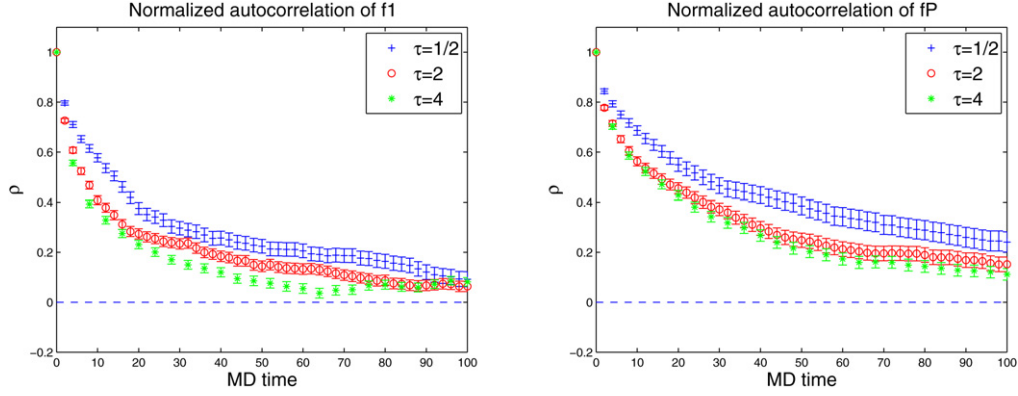


Fig. 3. Autocorrelation function of the correlators f_1 and f_P in quenched QCD (system B), for different trajectory lengths.

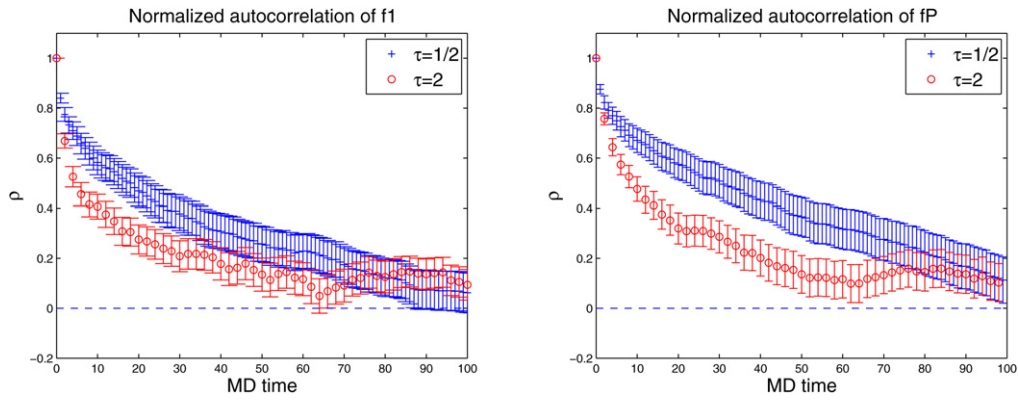


Fig. 4. Autocorrelation function of the correlators f_1 and f_P in $N_f = 2$ QCD (system C), for different trajectory lengths.

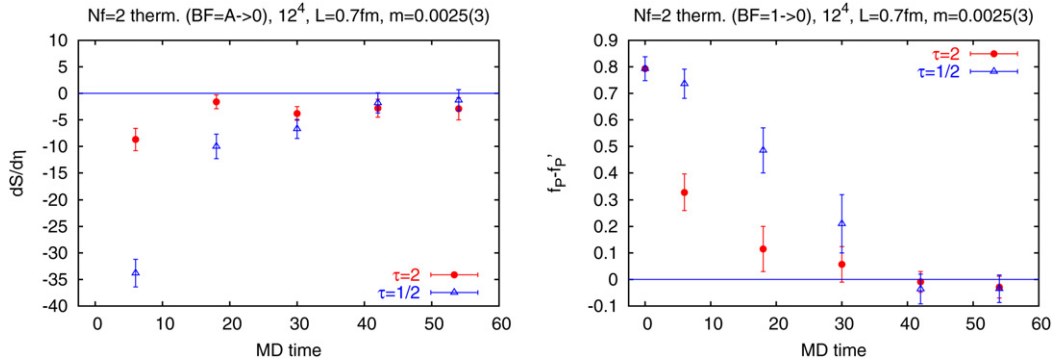


Fig. 5. Thermalization of two long-distance observables ($dS/d\eta$ and $f_P - f'_P$) having zero as expectation value. The starting configurations are taken from an ensemble where these expectation values are non-zero (system D).

1. start from the ensemble with non-zero boundary field ('point A' of [18]);
2. revert to homogeneous Dirichlet boundary conditions, and hence vanishing background field [18];
3. let the system rethermalize under the HMC evolution, for two different choices of trajectory length;
4. track how fast quantities whose expectation values vanish by symmetry in the homogeneous Dirichlet case relax to zero.

Fig. 5 shows the relaxation of the observables $dS/d\eta$ and $f_P - f'_P$. The latter is the asymmetry between the correlator

emanating from one boundary and the other. A data point at time t shows the value of the observable averaged over 16 independent replicas (corresponding to independent starting configurations), and averaged over the MD time interval $[t-6, t+6]$. The $dS/d\eta$ measurements were done after every trajectory and $f_P - f'_P$ every other trajectory.

The thermalization takes place faster with the choice of trajectory length $\tau = 2$, at least in the relatively early stages. In difficult simulations the time-consuming part is presumably the tail of the thermalization, but it would require very large statistics to demonstrate an algorithmic advantage in that regime.

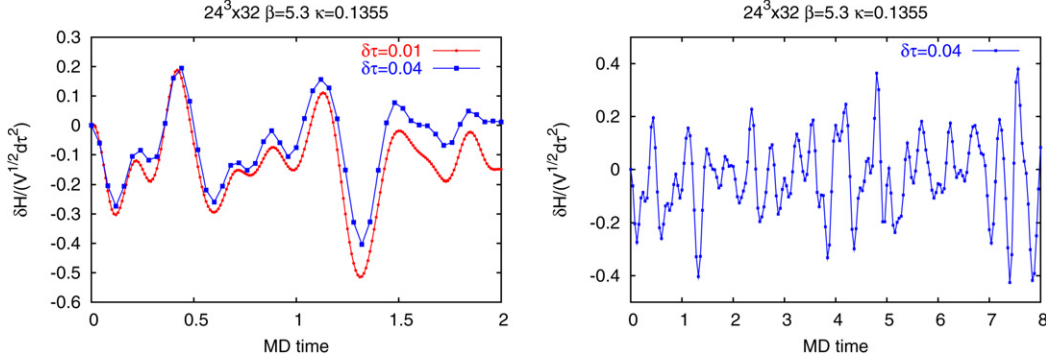


Fig. 6. Variation of the Hamiltonian along one MD trajectory on lattice C, in units of $\sqrt{V}d\tau^2$, where V is the number of lattice points.

3. Conclusion

We have investigated the dependence of autocorrelation times on the HMC trajectory length, focusing on long-distance observables, in a variety of different physical situations. We find this dependence to be substantial (a factor around 2) and statistically significant. The reduced correlation of successive measurements, done at fixed intervals of molecular dynamics time, is most clearly seen in the autocorrelation functions themselves.

The optimal choice of τ is observable dependent. However we have observed that trajectories longer than the ones conventionally used provide an advantage in computing standard physical quantities such as the pseudoscalar mass and decay constant. It will be interesting to see whether this conclusion also holds for HMC algorithms that use a different preconditioning of the pseudofermion action, for a different number of flavors, etc.

Naturally, other criteria are relevant in the final choice of trajectory length; the issues of stability and reversibility violations have been addressed in the appendix (see also Refs. [7,11,26] to name a few). We are currently performing a simulation at a quark mass of $m_s/2$ with $\tau = 2$ that shows good stability and controlled reversibility violations.

Acknowledgements

We thank DESY/NIC for computing resources on the APE machines and the computer team for support, in particular to run on the new apeNEXT systems. This project is part of ongoing algorithmic development within the SFB Transregio 9 “Computational Particle Physics” programme. R.S. thanks the Yukawa Institute for Theoretical Physics at Kyoto University; discussions during the YITP workshop *Actions and symmetries in lattice gauge theory* YITP-W-05-25 were useful to this work.

Appendix A. Stability and reversibility violations

In this appendix we discuss some important issues that one might worry about if one increases the length of MD trajectories. Most of what follows is not new but we find it useful to gather here the relevant points.

A.1. Energy violations along a trajectory

Fig. 6 shows the variation of the Hamiltonian along one MD trajectory on lattice C, for two different step sizes (here the leap-frog integrator was used with a ratio of 5 between the time-steps of the two pseudo-fermion forces [5]). The right plot illustrates the fact that the fluctuations of $\delta H(t)$ around zero do not grow very fast with MD time t .

Although a symplectic integrator such as leap-frog has discretization errors, there is a Hamiltonian which differs by $O(d\tau^2)$ terms from the MD Hamiltonian, and which is *exactly* conserved by the MD evolution [20,27]:

$$H_0 = H(t) + d\tau^2 h_1(t) + d\tau^4 h_2(t) + \dots \quad (\text{A.1})$$

Hence, for a given start configuration and momenta,

$$-\frac{\delta H(t)}{d\tau^2} = \delta h_1(t) + d\tau^2 \delta h_2(t) + \dots \quad (\text{A.2})$$

In words, the curve $\delta H(t)/d\tau^2$ is independent of $d\tau$, up to $O(d\tau^2)$ corrections. This is clearly what is seen on the left plot of Fig. 6. Note that since in equilibrium $\langle \delta H \rangle \simeq 1/2 \langle \delta H^2 \rangle = O(d\tau^4) > 0$ holds, we have $\langle \delta h_1(t) \rangle = 0$ and $\langle \delta h_2(t) \rangle < 0$ for any fixed t .

A.2. Scaling of energy violations

The scaling law $\langle \delta H^2 \rangle \sim V d\tau^4$ was proposed in [20,21]. Indeed, since δH is $O(d\tau^2)$ for one trajectory, δH^2 is $O(d\tau^4)$ and so is its average. As to the volume dependence, we confirm that $\langle \delta H^2 \rangle$ depends mainly on the number of lattice points V [22], while the dependence on the coupling β is very weak.¹ The values of $\langle \delta H^2 \rangle$ in Table 2 are consistent with $\langle \delta H^2 \rangle \propto V^{\eta(\tau)}$, where $\eta \in [1.07, 1.11]$.

The dependence of $\langle \delta H^2 \rangle$ on the trajectory length τ is very moderate but follows no obvious formula (see Table 2). Nevertheless, the trend is reasonably well described by τ^α , $0.3 < \alpha < 0.4$.

¹ Additional $\beta = 6.086$ simulations at $L = 12$ and 16 in system A with $\tau = 2$ give $\langle \delta H^2 \rangle / (L d\tau)^4 = 1.31(2)$ and $1.36(6)$, respectively.

A.3. Stability

Of course the statements made in the previous paragraph hold for exact arithmetics, and in practice it must be checked empirically at what trajectory length rounding errors introduce instabilities in the integration [26]. The other source of instability can be the occurrence of an exceptionally large force, which spoils the expansion in $d\tau$. In the simulation C (see Table 1), $|\delta H|$ was bounded by 1 in the $\tau = 2$ run, while one replica experienced one spike of $\delta H \simeq 2000$ in the $\tau = 1/2$ run (the simulation then continued normally).

A.4. Reversibility violations

Reversibility violations in our simulations come from rounding errors, and from the non-zero residuals of the inversions in the MD evolution.² These effects of course accumulate with the length of the trajectory. However Tables 2, 3 and 4 show that the increase of $\langle |\delta H_{\leftrightarrow}| \rangle$ with τ in simulations where $|\delta H|$ is bounded by one and is typically much smaller, is rather slow, roughly like $\sim \sqrt{\tau}$. Therefore it should not present a serious problem. In the trajectory of Fig. 6 with $d\tau = 0.04$, $10^4 \cdot \langle |\delta H_{\leftrightarrow}| \rangle$ is 0.06, 0.3, 1.4 and 1.4 for $\tau = 0.08, 0.8, 1.6$ and 8, respectively.

How a small reversibility violation influences the ensemble generated is not known (to us); it may lead to an incorrect sampling. The influence on expectation values of observables is even harder to pin down. The study [7] showed that even with $\langle |\delta H_{\leftrightarrow}| \rangle = 0.01$ one does not see an effect in a number of observables computed at the sub-percent level.

References

- [1] S. Duane, A.D. Kennedy, B.J. Pendleton, D. Roweth, Phys. Lett. B 195 (1987) 216.
- [2] A. Ukawa, CP-PACS Collaboration, JLQCD Collaboration, Nucl. Phys. Proc. 106 (Suppl.) (2002) 195; T. Lippert, *ibid.*, 193; H. Wittig, *ibid.*, 197.
- [3] M. Lüscher, Comput. Phys. Comm. 165 (2005) 199, hep-lat/0409106.
- [4] M. Hasenbusch, Phys. Lett. B 519 (2001) 177, hep-lat/0107019; M. Hasenbusch, K. Jansen, Nucl. Phys. B 659 (2003) 299, hep-lat/0211042.
- [5] A. Ali Khan, et al., QCDSF Collaboration, Phys. Lett. B 564 (2003) 235, hep-lat/0303026; C. Urbach, K. Jansen, A. Shindler, U. Wenger, Comput. Phys. Comm. 174 (2006) 87, hep-lat/0506011.
- [6] M.A. Clark, P. de Forcrand, A.D. Kennedy, PoS LAT2005 (2005) 115, hep-lat/0510004.
- [7] M. Della Morte, F. Knechtli, J. Rolf, R. Sommer, I. Wetzorke, U. Wolff, ALPHA Collaboration, Comput. Phys. Comm. 156 (2003) 62, hep-lat/0307008.
- [8] B. Gehrmann, U. Wolff, Nucl. Phys. Proc. 83 (Suppl.) (2000) 801, hep-lat/9908003; B. Gehrmann, Ph.D. Thesis, hep-lat/0207016.
- [9] A.D. Kennedy, B. Pendleton, Nucl. Phys. Proc. 20 (Suppl.) (1991) 118.
- [10] P.B. Mackenzie, Phys. Lett. B 226 (1989) 369.
- [11] C. Liu, A. Jaster, K. Jansen, Nucl. Phys. B 524 (1998) 603, hep-lat/9708017.
- [12] S. Gupta, Nucl. Phys. B 370 (1992) 741.
- [13] A.D. Kennedy, B. Pendleton, Nucl. Phys. B 607 (2001) 456, hep-lat/0008020.
- [14] K. Jansen, C. Liu, Comput. Phys. Comm. 99 (1997) 221, hep-lat/9603008; JLQCD Collaboration, S. Aoki, et al., Phys. Rev. D 65 (2002) 094507, hep-lat/0112051.
- [15] J.C. Sexton, D.H. Weingarten, Nucl. Phys. B 380 (1992) 665.
- [16] U. Wolff, ALPHA Collaboration, Comput. Phys. Comm. 156 (2004) 143, hep-lat/0306017.
- [17] M. Guagnelli, J. Heitger, R. Sommer, H. Wittig, ALPHA Collaboration, Nucl. Phys. B 560 (1999) 465, hep-lat/9903040.
- [18] M. Lüscher, R. Sommer, P. Weisz, U. Wolff, Nucl. Phys. B 413 (1994) 481, hep-lat/9309005.
- [19] R. Sommer, Nucl. Phys. B 411 (1994) 839, hep-lat/9310022.
- [20] R. Gupta, G.W. Kilcup, S.R. Sharpe, Phys. Rev. D 38 (1988) 1278.
- [21] M. Creutz, Phys. Rev. D 38 (1988) 1228.
- [22] S. Gupta, A. Irback, F. Karsch, B. Petersson, Phys. Lett. B 242 (1990) 437.
- [23] S. Necco, R. Sommer, Nucl. Phys. B 622 (2002) 328, hep-lat/0108008.
- [24] M. Della Morte, R. Frezzotti, J. Heitger, J. Rolf, R. Sommer, U. Wolff, ALPHA Collaboration, Nucl. Phys. B 713 (2005) 378, hep-lat/0411025.
- [25] M. Göckeler, R. Horsley, A.C. Irving, D. Pleiter, P.E.L. Rakow, G. Schierholz, H. Stuben, QCDSF Collaboration, hep-ph/0409312.
- [26] B. Joo, B. Pendleton, A.D. Kennedy, A.C. Irving, J.C. Sexton, S.M. Pickles, S.P. Booth, UKQCD Collaboration, Phys. Rev. D 62 (2000) 114501, hep-lat/0005023.
- [27] A.D. Kennedy, Lectures at Nara Workshop on Perspectives in Lattice QCD, 2005.
- [28] I. Montvay, G. Münster, Quantum Fields on a Lattice, Cambridge Univ. Press, Cambridge, UK, 1994.

² The solver starts from the vector zero in our implementation of the molecular dynamics.