Hybrid Monte Carlo with non-uniform step size

Christian Holzgräfe,[a] Arnab Bhattacherjee,[b] and Anders Irbäck[c]
Department of Astronomy and Theoretical Physics, Lund University, Sölvegatan 14A, SE-223 62 Lund, Sweden

(Received 9 October 2013; accepted 7 January 2014; published online 24 January 2014)

The Hybrid Monte Carlo method offers a rigorous and potentially efficient approach to the simulation of dense systems, by combining numerical integration of Newton's equations of motion with a Metropolis accept-or-reject step. The Metropolis step corrects for sampling errors caused by the discretization of the equations of motion. The integration is usually performed using a uniform step size. Here, we present simulations of the Lennard-Jones system showing that the use of smaller time steps in the tails of each integration trajectory can reduce errors in energy. The acceptance rate is 10–15 percentage points higher in these runs, compared to simulations with the same trajectory length and the same number of integration steps but a uniform step size. We observe similar effects for the harmonic oscillator and a coarse-grained peptide model, indicating generality of the approach. © *2014 AIP Publishing LLC*. [http://dx.doi.org/10.1063/1.4862687]

I. INTRODUCTION

Molecular simulations often involve numerical integration of the classical equations of motion.[1] The Hybrid Monte Carlo (HMC) algorithm adds a Metropolis accept-or-reject step to the traditional molecular dynamics approach in order to remove sampling errors stemming from the discretization of the equations of motion.[2] While widely used in the area of lattice field theory, this algorithm has been less commonly used for molecular simulations, at least in part due to concerns that the rejection of computationally expensive trial states might mean a loss in efficiency.

In HMC, the acceptance rate is determined by the error in energy along the integration trajectories used to generate new trial states. It therefore depends strongly on the integration step size, $\delta\tau$.[3,4] The length of the trajectories, τ, is also important; for best decorrelation efficiency, τ must not be too small.[5–7] The need for sufficiently long trajectories brings into focus how errors propagate during the integration. In HMC, because the discretized equations of motion sample a shifted ensemble,[8,9] the simulated system evolves under non-stationary conditions. Hence, even after thermalization, each integration trajectory begins with a relaxation phase. This fact suggests that the usual choice of a constant integration step size need not be optimal. In this article, we test whether the introduction of a non-uniform integration step size can improve the performance of HMC for the harmonic oscillator and the Lennard-Jones system. We also briefly discuss results obtained for a C_α-C_β model for peptide helix formation.[10] For all three systems, we find that the use of a non-uniform step size can improve the acceptance rate, at unchanged average step size.

Many extensions and refinements of the original HMC update have been explored, with varying degrees of success. The approaches include finding suitable integrators more accurate than the standard leapfrog method,[11–17] introducing different time steps for fast and slow forces[18–20] or fast and slow Fourier modes,[8,21] and constructing shadow Hamiltonians that are better conserved by the discretized equations of motion.[22,23] Yet another approach is generalized HMC,[24–26] which regulates the extent to which momenta are randomized at the refreshment steps. The use of a non-uniform step size is a complementary method that can can be combined with any of these techniques, and whose potential benefits appear to have been overlooked so far.

II. METHODS

A. HMC

We consider a classical system with N particles defined by the Hamiltonian

$$H = \sum_{i=1}^{N} \frac{\mathbf{p}_i^2}{2m_i} + U(\mathbf{r}_1, \ldots, \mathbf{r}_N), \tag{1}$$

where U denotes potential energy. Based on numerical integration of the equations of motion, the HMC algorithm generates Boltzmann-weighted configurations $C = \{\mathbf{r}_1, \ldots, \mathbf{r}_N\}$, $p(C) \propto e^{-\beta U(C)}$ ($\beta = 1/kT$, k is Boltzmann's constant, T is temperature). Specifically, a HMC update consists of the following three steps:[2]

1. Draw Maxwell-Boltzmann distributed momenta $\mathbf{p}_i$ for all particles, $P(\mathbf{p}_i) \propto e^{-\beta \mathbf{p}_i^2/2m_i}$.
2. Evolve the system over a time τ using an integrator that is time-reversible and preserves phase-space volume, such as the leapfrog/Størmer/Verlet scheme

$$\begin{cases} \mathbf{p}_i(t+\delta\tau/2) = \mathbf{p}_i(t) - \delta\tau \nabla_i U(t)/2 \\ \mathbf{r}_i(t+\delta\tau) = \mathbf{r}_i(t) + \delta\tau \mathbf{p}_i(t+\delta\tau/2)/m_i \\ \mathbf{p}_i(t+\delta\tau) = \mathbf{p}_i(t+\delta\tau/2) - \delta\tau \nabla_i U(t+\delta\tau)/2 \end{cases} \tag{2}$$

which is iterated N_{MD} times ($\tau = N_{\mathrm{MD}} \times \delta\tau$).

[a]Electronic mail: christian.holzgraefe@thep.lu.se
[b]Present address: Department of Structural Biology, Weizmann Institute of Science, Rehovot 76100, Israel. Electronic mail: arnab.bhattacherjee@weizmann.ac.il
[c]Electronic mail: anders@thep.lu.se

 © 2014 AIP Publishing LLC

This article is copyrighted as indicated in the article. Reuse of AIP content is subject to the terms at: http://scitation.aip.org/termsconditions. Downloaded to IP: 130.235.189.230 On: Fri, 24 Jan 2014 14:59:42

3. Accept or reject the new configuration with acceptance probability $P_{\rm acc} = \min(1, {\rm e}^{-\beta\Delta H})$, where ΔH is the net energy change along the trajectory in step 2. With an exact integration, the acceptance rate would be one.

This update satisfies detailed balance with respect to the Boltzmann distribution independent of $\delta\tau$ and τ, which thus are tunable parameters. The leapfrog step (Eq. (2)) can be expressed as

$$T_{\rm LF}(\delta\tau) = T_p(\delta\tau/2)T_r(\delta\tau)T_p(\delta\tau/2), \tag{3}$$

where T_p and T_r are phase-space mappings corresponding to simple Euler steps in either momenta or positions. It is easy to verify that T_p and T_r are time-reversible $[T(\delta\tau)^{-1} = T(-\delta\tau)]$ and volume-preserving, from which it immediately follows that $T_{\rm LF}$ indeed has these required properties.

B. Local error

Leapfrog is a second-order integrator, so the local (one-step) error is of order $O(\delta\tau^3)$. Higher-order integrators that are time-reversible and volume-preserving can be constructed.[11–14] Another approach toward improved accuracy is to make a leapfrog-like ansatz and reduce rather than eliminate leading error terms.[15–17] The simplest ansatz of this kind, with five alternating T_p and T_r steps, can be written as

$$\begin{aligned} T_\xi(\delta\tau) = {} & T_p(\xi\delta\tau/2)T_r(\delta\tau/2) \\ & \times T_p((1-\xi)\delta\tau)T_r(\delta\tau/2)T_p(\xi\delta\tau/2). \end{aligned} \tag{4}$$

The symmetric construction ensures that T_ξ is time-reversible and volume-preserving. The remaining parameter ξ is determined so as to minimize the leading term of the one-step error in energy,[15,17] which is given by

$$\left[\frac{A(\xi)}{4}\sum_{i,j}\frac{(\nabla_i U\cdot\nabla_i)(\mathbf{p}_j\cdot\nabla_j)}{m_i m_j}U - \frac{B(\xi)}{12}\sum_{i,j,k}\frac{(\mathbf{p}_i\cdot\nabla_i)(\mathbf{p}_j\cdot\nabla_j)(\mathbf{p}_k\cdot\nabla_k)}{m_i m_j m_k}U\right]\times\delta\tau^3, \tag{5}$$

where ξ comes in through the coefficients $A(\xi) = 3\xi - 1 - \xi^2$ and $B(\xi) = (3\xi - 1)/2$. For quadratic potentials, the error can be made $O(\delta\tau^4)$ by selecting ξ so that $A(\xi) = 0$, which means $\xi = (3-\sqrt{5})/2 \approx 0.382$. This choice gives a higher-order harmonic (HOH) integrator,[17] with a small value for the other coefficient as well, $B \approx 0.073$. A previous and different criterion for determining ξ gives a similar result, namely, $\xi \approx 0.386$.[15] The local error of the leapfrog method is given by Eq. (5) with $A = B = 1$ $(\xi = 1)$.

The integrators that eliminate or reduce leading-order error terms are expected to be more accurate than leapfrog if the step size is small but not in general; the relative accuracies of the methods depend on the step size.[17] A strength of the leapfrog method is its ability to tolerate relatively large step sizes.

C. Global error for the harmonic oscillator

Knowledge of the local integration error is insufficient for understanding the dynamics of the HMC algorithm; one also needs to know how the error in energy propagates along a trajectory. This problem is a challenge to analyze in the general case, but can be examined in detail for the harmonic oscillator. Suppose a harmonic oscillator with energy $H = p^2/2 + x^2/2$ (unit mass and frequency) is integrated using leapfrog with time step $\delta\tau$ (Eq. (2)), starting from an initial state (x_0, p_0) drawn from the canonical distribution ${\rm e}^{-H}$ $(kT = 1)$. Let $\Delta H^{(n)} = H(x_n, p_n) - H(x_0, p_0)$ denote the energy change after n leapfrog steps. For $\delta\tau < 2$, one finds that $\Delta H^{(n)}$ averaged over the initial state is given by

$$\langle\Delta H^{(n)}\rangle = \frac{\delta\tau^4}{64(1-\delta\tau^2/4)}(1-\cos 2n\theta), \tag{6}$$

where $\theta = \arctan[\delta\tau\sqrt{1-\delta\tau^2/4}/(1-\delta\tau^2/2)]$. Throughout this $\delta\tau$ regime, the average energy change is an oscillating function of time $(t = n\times\delta\tau)$. The absence of energy drift can be attributed to the existence of a shadow Hamiltonian, $H_{\delta\tau} = p^2/2(1-\delta\tau^2/4) + x^2/2$, which is similar to H for small $\delta\tau$ and strictly conserved under the leapfrog evolution. For $\delta\tau > 2$, the integration becomes unstable, leading to an exponential growth of $\langle\Delta H^{(n)}\rangle$ with time.

Let us now modify the integration trajectories, by changing the first k time steps, $\delta\tau_i$, while leaving the rest equal to $\delta\tau$. An integration consisting of $n \geq k$ steps can then be described by the mapping $T_{\rm LF}(\delta\tau)^{n-k}A$, where $A = T_{\rm LF}(\delta\tau_k)\ldots T_{\rm LF}(\delta\tau_1)$. The first k steps, described by A, do not preserve the shadow energy $H_{\delta\tau}$, as $\delta\tau_i \neq \delta\tau$. Averaging over the initial state, one finds that

$$\langle H_{\delta\tau}^{(k)}\rangle = \frac{1}{2}\left(\frac{A_{px}^2 + A_{pp}^2}{1-\delta\tau^2/4} + A_{xx}^2 + A_{xp}^2\right), \tag{7}$$

where $H_{\delta\tau}^{(k)} = H_{\delta\tau}(x_k, p_k)$. After these k steps, $H_{\delta\tau}$ is conserved, so $\langle H_{\delta\tau}^{(n)}\rangle = \langle H_{\delta\tau}^{(k)}\rangle$ for $n > k$. One furthermore finds that for $n > k$ the average energy error, $\langle\Delta H^{(n)}\rangle$, oscillates with time about some level $\overline{\Delta H}$, as in Eq. (6). This level is, however, not the same as in Eq. (6), but depends on A and can be expressed as

$$\overline{\Delta H} = (1-\delta\tau^2/8)\langle H_{\delta\tau}^{(k)}\rangle - 1, \tag{8}$$

where $\langle H_{\delta\tau}^{(k)}\rangle$ is given by Eq. (7). Minimal $\overline{\Delta H}$ is thus achieved by selecting the initial time steps, and thereby A, so as to minimize the average shadow energy, as given by Eq. (7). For a uniform step size, which corresponds to setting A equal to the unit matrix, it can be readily verified that Eqs. (7) and (8) yield the same result as Eq. (6), namely, $\overline{\Delta H} = \delta\tau^4/64(1-\delta\tau^2/4)$.

It turns out to be easy to reduce $\overline{\Delta H}$ by changing a few initial steps. For $k = 1$, a reduction of $\overline{\Delta H}$ is observed for any $\delta\tau_1$ between 0 and $\delta\tau$. The maximal reduction is 17% and occurs for $\delta\tau_1 = \delta\tau\sqrt{2/3}$. For $k = 4$ and using $\delta\tau = 0.8$, we find numerically that the maximal reduction of $\overline{\Delta H}$ is 50%. Again, a reduced $\overline{\Delta H}$ is easy to achieve; a reduction of >40% is found for as much as 36% of all possible sequences with $0 < \delta\tau_i < 0.8$ $(i = 1, \ldots, 4)$. In Fig. 1(a), we show examples of

This article is copyrighted as indicated in the article. Reuse of AIP content is subject to the terms at: http://scitation.aip.org/termsconditions. Downloaded to IP: 130.235.189.230 On: Fri, 24 Jan 2014 14:59:42

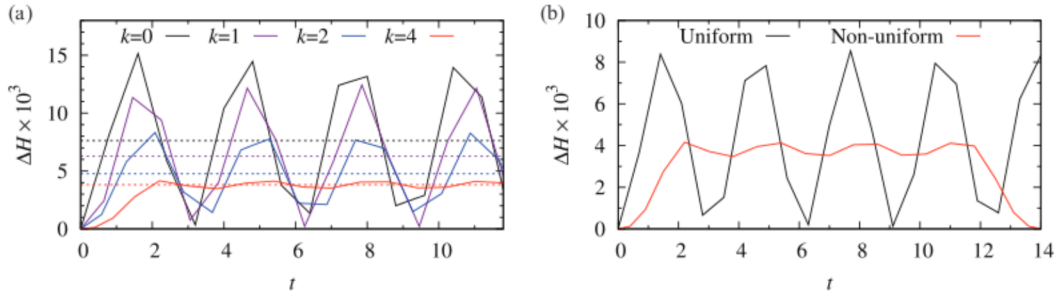

FIG. 1. Time evolution of the average energy error, $\langle \Delta H^{(n)} \rangle$, in leapfrog integration of the harmonic oscillator, $H = p^2/2 + x^2/2$, starting from initial states drawn from the canonical distribution e^{-H}. (a) Integration with step-size sequences of the form $(\delta\tau_1, \ldots, \delta\tau_k, \delta\tau, \delta\tau, \ldots)$ for $\delta\tau = 0.8$ and $k = 0$, 1, 2, and 4. Dashed horizontal lines represent $\overline{\Delta H}$, as given by Eq. (8). The initial step sizes are $\delta\tau_1 = \delta\tau\sqrt{2/3} \approx 0.65$ for $k = 1$, $(\delta\tau_1, \delta\tau_2) = (0.44, 0.62)$ for $k = 2$, and $(\delta\tau_1, \delta\tau_2, \delta\tau_2, \delta\tau_4) = (0.40, 0.50, 0.60, 0.70)$ for $k = 4$. (b) Integration with a uniform step size (0.70) and a symmetric non-uniform step-size sequence with the same number of steps ($N_{\mathrm{MD}} = 20$) and the same trajectory length ($\tau = 14$). The non-uniform sequence begins with (0.40, 0.50, 0.60, 0.70) and ends with the reverse four-step sequence. In the mid part, the step size is 0.80.

calculations with 0, 1, 2, and 4 modified initial steps. After the initial steps, the time step is the same in all four cases ($\delta\tau = 0.8$). The results show that the insertion of a few smaller steps in the beginning of the trajectories indeed can have large effects on the subsequent evolution of the average energy error.

This observation suggests that it might be useful to reduce the initial step sizes in HMC trajectories. To maintain time-reversibility, the reverse sequence of steps must then be added at the end of each trajectory. Through a broad set of calculations, we find that the unavoidable smaller final steps, fortunately, tend to further reduce the energy error, and often strongly so. Figure 1(b) shows an example of this. Here, trajectories with a symmetric non-uniform step-size sequence are compared to trajectories with the same number of steps, N_{MD}, and the same trajectory length, τ, but a uniform step size. With the non-uniform step-size sequence, there is a marked reduction in the average energy error, $\langle \Delta H^{(n)} \rangle$, toward the end of the trajectories. As a result, the final $\langle \Delta H^{(n)} \rangle$ is much smaller than it is with uniform step size, despite that N_{MD} and τ are the same in both cases.

For the harmonic oscillator, we thus find that the use of a non-uniform step size can reduce the average energy error along HMC trajectories, at no additional cost. In Sec. III, we show that this holds true for the Lennard-Jones system as well.

All the results presented above are for the $T_p(\delta\tau/2)T_r(\delta\tau)T_p(\delta\tau/2)$ form of the leapfrog integrator. Calculations on the same system, the harmonic oscillator, were also done with the alternative $T_r(\delta\tau/2)T_p(\delta\tau)T_r(\delta\tau/2)$ leapfrog scheme as well as with the above-mentioned HOH integrator.[17] For both these integrators, it was again easy to find symmetric non-uniform step-size sequences that brought a large reduction in the average energy error, compared to results obtained with uniform step size for the same N_{MD} and τ. Thus, the findings described above are not unique to the $T_p(\delta\tau/2)T_r(\delta\tau)T_p(\delta\tau/2)$ integrator. Nevertheless, in our Lennard-Jones study below, we focus on this particular integrator, for the sake of clarity.

D. Lennard-Jones system

We test the HMC method with non-uniform step size on a system of $N = 256$ particles with pairwise Lennard-Jones interactions, which has previously been studied using standard HMC.[27] In scaled units,[28] the Hamiltonian is given by

$$H = \sum_{i=1}^{N} \frac{\mathbf{p}_i^2}{2} + 4 \sum_{i<j} \left(\frac{1}{r_{ij}^{12}} - \frac{1}{r_{ij}^{6}} \right), \tag{9}$$

where $r_{ij} = |\mathbf{r}_i - \mathbf{r}_j|$ is the distance between particles i and j. The Lennard-Jones potential is truncated at $r_{ij} = 2.5$ and shifted to have it continuous. The particles reside in a periodic box of size L^3 with $L = 6.76$, which gives a density of $\rho = 0.83$. Two temperatures are used, $T = 0.5$ and $T = 1.0$. The time steps range from 0.004 to 0.04, where a value of 0.005 corresponds to 1.09×10^{-14} s for argon.[28]

All our simulations of this system are carried out using the leapfrog integrator of Eq. (2). For each set of simulation parameters studied, we generate 32 independent runs, starting from equilibrated initial conditions. The total number of elementary leapfrog steps in each run is kept the same for all parameter sets, namely, $N_{\mathrm{tot}} = 6.7 \times 10^6$, thus ensuring that the computational cost is approximately the same for all runs. For fair comparison, we moreover collect data as often as one can, that is, after each Metropolis question. This gives us $N_{\mathrm{tot}}/N_{\mathrm{MD}}$ measurements of a given property in each run. With this setup, relative computational efficiencies can be assessed by direct comparison of statistical errors. Throughout the paper, error bars and shaded bands represent one standard error of the mean.

III. RESULTS

We study the performance of HMC with non-uniform step size using the Lennard-Jones system as a testbed. Our calculations can be divided into two parts. First, we locate approximately optimal values of the step size $\delta\tau$ (Sec. III A) and the trajectory length τ (Sec. III B) for HMC with uniform step size. Second, keeping both τ and the number of integration steps per trajectory, N_{MD}, fixed at values found this way, we ask whether the acceptance rate can be improved by allowing for a non-uniform step size (Sec. III C).

A. Step-size dependence of standard HMC

We study the $\delta\tau$ dependence of standard HMC using a fixed trajectory length of $\tau = 0.24$. We vary $\delta\tau$ between

This article is copyrighted as indicated in the article. Reuse of AIP content is subject to the terms at: http://scitation.aip.org/termsconditions. Downloaded to IP: 130.235.189.230 On: Fri, 24 Jan 2014 14:59:42

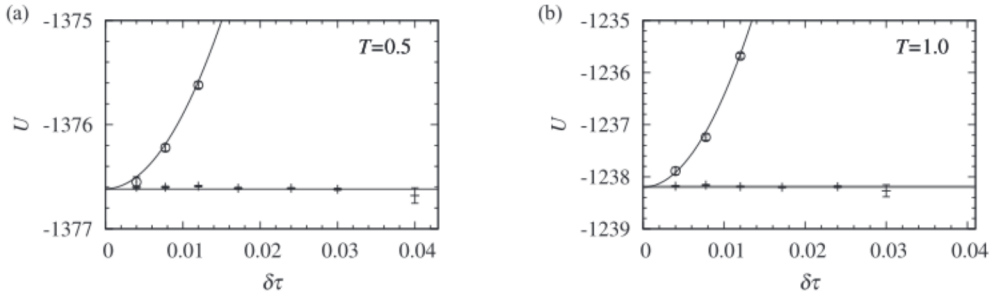

FIG. 2. Average potential energy for the Lennard-Jones system as obtained from HMC simulations with uniform step size (+) and from otherwise identical simulations with the Metropolis question left out (○), for different $\delta\tau$ and fixed $\tau = 0.24$, at (a) $T = 0.5$ and (b) $T = 1.0$. Curves are quadratic fits $U = a + b\,\delta\tau^2$, with a fixed, to the data obtained without the Metropolis step. Shaded horizontal bands represent results from ordinary MC runs.

0.004 and 0.04, so N_{MD} is in the range 6–60. For comparison, we also conduct otherwise identical simulations with the Metropolis step switched off. In addition, for reference, we carry out ordinary MC simulations with random single-particle translations as the elementary move.

Fig. 2 shows the average potential energy as obtained from these different simulations. The HMC data are independent of $\delta\tau$ and agree well with the MC results. Leaving out the Metropolis correction leads, by contrast, to noticeable deviations from the MC results. Since the local leapfrog error is cubic and the trajectory length is held fixed, the deviations should scale quadratically with $\delta\tau$, as is indeed observed in the simulations (Fig. 2).

Having verified that the HMC simulations correctly describe this physical property, we turn to the computational efficiency. Fig. 3 shows the acceptance rate and the statistical error on the potential energy, σ_U, as functions of $\delta\tau$. The acceptance rate depends strongly on $\delta\tau$. For systems with many degrees of freedom, the acceptance rate is expected to be well described by the relation $P_{\mathrm{acc}} = \mathrm{erfc}(\sqrt{\beta\langle\Delta H\rangle}/2)$,[4] where erfc(·) denotes the complementary error function and ΔH is the energy change over a trajectory. A control shows that the P_{acc} data in Fig. 3 to a good approximation obey this relation.

The computational efficiency depends not only on the acceptance rate. To assess the computational efficiency, we use the statistical error σ_U, as estimated from the variance of 32 single-run averages of U. Compared to the acceptance rate, σ_U is more difficult to determine, leading to noisier data. Nevertheless, the results show that σ_U exhibits a broad minimum, with close-to-minimal values in a region corresponding to acceptance rates between 0.5 and 0.9 (Fig. 3). Hence, precise fine-tuning of $\delta\tau$ is not required in order to achieve approximately optimal efficiency. Traditional molecular dynamics simulations typically use a step size of 0.005 or smaller, for which the deviations seen in Fig. 2 are modest. The inclusion of the Metropolis step makes it possible to increase the step size and thereby improve the sampling efficiency (Fig. 3).

B. Trajectory-length dependence of standard HMC

The above analysis shows that $\delta\tau = 0.02$ is a reasonable choice of step size at both temperatures studied. Next, we examine how the performance of standard HMC depends on the trajectory length τ, keeping $\delta\tau = 0.02$ fixed. Again, we verified that the results for the average potential energy agreed with MC data, independent of τ (see Fig. S1 of the supplementary material[29]).

The acceptance rate depends, as expected,[5] only weakly on τ (Fig. 4). At $T = 0.5$, P_{acc} decreases from 0.73 to 0.70 upon a 50-fold increase of τ, from 0.02 to 1.00. At $T = 1.0$, the corresponding drop in P_{acc} is from 0.49 to 0.43.

The statistical error σ_U shows, by contrast, a clear τ dependence (Fig. 4). An increased N_{MD} implies fewer measurements, since the total number of leapfrog steps is held constant. As a result, if the measurements were uncorrelated, one would expect σ_U to grow as $\sqrt{N_{\mathrm{MD}}}$. Our σ_U data at $T = 1.0$ show approximately this behavior for $N_{\mathrm{MD}} > 10$. In general, time correlations must, however, be taken into

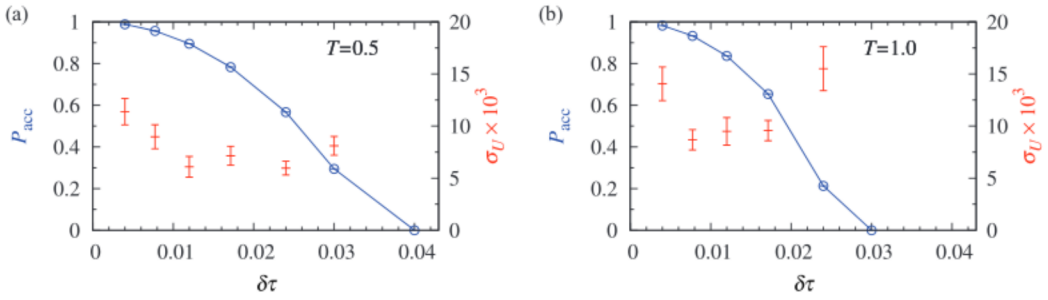

FIG. 3. The acceptance rate P_{acc} and the statistical error σ_U in HMC simulations of the Lennard-Jones system for fixed $\tau = 0.24$ and different uniform step sizes $\delta\tau$, at (a) $T = 0.5$ and (b) $T = 1.0$. Lines are drawn to guide the eye. Error bars on P_{acc} are small compared to the plot symbols. At the highest $\delta\tau$ studied, σ_U lies above the plotted range.

This article is copyrighted as indicated in the article. Reuse of AIP content is subject to the terms at: http://scitation.aip.org/termsconditions. Downloaded to IP: 130.235.189.230 On: Fri, 24 Jan 2014 14:59:42

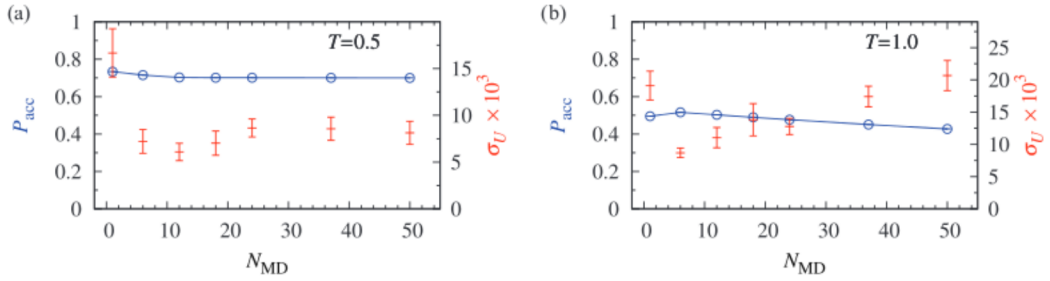

FIG. 4. The acceptance rate P_{acc} and the statistical error σ_U in HMC simulations of the Lennard-Jones system for fixed $\delta\tau = 0.02$ and different trajectory lengths $\tau = N_{MD} \times \delta\tau$, at (a) $T = 0.5$ and (b) $T = 1.0$. Lines are drawn to guide the eye. Errors bars on P_{acc} are small compared to the plot symbols.

account. To illustrate the strength of these correlations, we compute the normalized autocorrelation function for U, $\rho_U(t)$, for $N_{MD} = 1$ and 12 at both temperatures (see Fig. S2 of the supplementary material[29]). Table I shows the corresponding integrated autocorrelation times, $\tau_{int,U} = 1/2 \sum_{-\infty}^{\infty} \rho_U(t)$. A change from $N_{MD} = 1$ to $N_{MD} = 12$ is seen to yield a drastic drop in $\tau_{int,U}$. The gain in decorrelation efficiency is more than sufficient to balance the lower number of measurements, thus leading to a reduced σ_U for $N_{MD} = 12$ (Fig. 4). Knowing $\tau_{int,U}$, an alternative estimate of the statistical error on U can be calculated as $\tilde{\sigma}_U = [2\tau_{int,U}\mathrm{var}(U)/n]^{1/2}$, where $\mathrm{var}(U)$ is the variance of U and n is the number of measurements. Our two estimates $\tilde{\sigma}_U$ and σ_U agreed within 10%–20% in all cases. From Fig. 4, we conclude that, at both temperatures, N_{MD} should be chosen in the vicinity of $N_{MD} = 10$ for optimal efficiency.

C. Non-uniform step size

Having located approximately optimal parameters of standard HMC with uniform step size, we now ask whether the computational efficiency can be further improved by introducing a non-uniform step size. To this end, we pick four examples of low σ_U from Fig. 4, namely, $N_{MD} = 12$, 18 for $T = 0.5$ and $N_{MD} = 6$, 12 for $T = 1.0$. In each case, we seek to replace the original uniform step size by a step-size sequence that increases the acceptance rate, while leaving N_{MD} and τ unchanged. A higher acceptance rate should under these conditions guarantee enhanced computational efficiency.

In all four cases, such a step-size sequence proved easy to find by trial and error. The best sequences found are displayed in Table II. Most likely, even better sequences can be found by further fine-tuning. The acceptance rates achieved with the step-size sequences in Table II are consistently 10–15 percentage points higher than those observed in the corresponding simulations with uniform step size.

To see how the increase in acceptance rate arises, it is instructive to examine how the average error in energy evolves. Fig. 5 shows the evolution of $\langle \Delta H(t) \rangle$ along the integration trajectories in the two runs for $T = 0.5$ and $N_{MD} = 18$ in Table II. Two key properties are worth noting for the run with non-uniform step size. First, in the mid part of the trajectories, $\langle \Delta H(t) \rangle$ is smaller than it would have been without the smaller steps in the beginning. Note that the step size in this part (0.023) is larger than the step size used in the standard

TABLE I. Integrated autocorrelation times $\tau_{int,U}$, in units of the trajectory length τ, in HMC simulations of the Lennard-Jones system with a uniform step size of $\delta\tau = 0.02$.

T	N_{MD}	$\tau_{int,U}$
0.5	1	347 ± 3
0.5	12	4.21 ± 0.02
1.0	1	127.0 ± 0.4
1.0	12	3.28 ± 0.01

TABLE II. Acceptance rates obtained with uniform (P_{acc}^{old}) and non-uniform (P_{acc}^{new}) step size in HMC simulations of the Lennard-Jones system, for four combinations of T and N_{MD}. In the four standard HMC runs, the step size is $\delta\tau = 0.020$. The non-uniform step-size sequences $\delta\tau_1, \ldots, \delta\tau_k, \delta\tau_k, \ldots, \delta\tau_1$ ($N_{MD} = 2k$) are chosen so that the trajectory length is the same as in the corresponding simulation with uniform step size, that is $\tau = 2\sum_{j=1}^{k} \delta\tau_j = N_{MD} \times \delta\tau$.

T	N_{MD}	$\delta\tau_1, \ldots, \delta\tau_k$	P_{acc}^{old}	P_{acc}^{new}
0.5	12	0.014, 0.017, 0.019, 0.022, 0.024, 0.024	0.703	0.825
0.5	18	0.014, 0.015, 0.017, 0.020, 0.022, 0.023, 0.023, 0.023, 0.023	0.702	0.851
1.0	6	0.017, 0.020, 0.023	0.516	0.631
1.0	12	0.017, 0.019, 0.021, 0.021, 0.021, 0.021	0.503	0.617

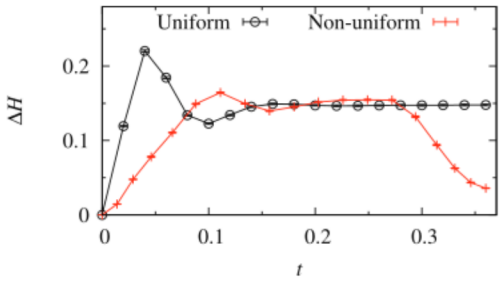

FIG. 5. Time evolution of the average energy error, $\langle \Delta H(t) \rangle$, along integration trajectories in HMC simulations of the Lennard-Jones system with uniform and non-uniform step size, respectively, at $T = 0.5$. Both runs use $N_{MD} = 18$ and $\tau = 0.36$. The non-uniform step-size sequence can be found in Table II. Lines are drawn to guide the eye. Error bars are small compared to the plot symbols.

This article is copyrighted as indicated in the article. Reuse of AIP content is subject to the terms at: http://scitation.aip.org/termsconditions. Downloaded to IP: 130.235.189.230 On: Fri, 24 Jan 2014 14:59:42

HMC run (0.020). A standard HMC run with the same step size (0.023) gives a significantly larger $\langle \Delta H(t) \rangle$ (see Fig. S3 of the supplementary material[29]). Second, the smaller steps at the end of the trajectories lead to a drop in $\langle \Delta H(t) \rangle$. Together, these two effects make the final $\langle \Delta H(t) \rangle$ much lower than it is in the simulation with uniform step size and identical τ and N_{MD}. The same two effects were observed for the harmonic oscillator (Fig. 1).

IV. DISCUSSION AND SUMMARY

In a HMC simulation, the system is kept in the desired equilibrium ensemble by the randomization of momenta and the Metropolis step. Each integration trajectory begins with a relaxation toward a slightly shifted ensemble. The shift represents a perturbation caused by discretization effects. The magnitude of the shift is not a simple function of the step size, but depends also on how the perturbation is switched on (Fig. 1(a)). This observation led us to introduce a non-uniform step size in our calculations, such that the steps are smaller in the tails of each integration trajectory. With this modification, we showed that the acceptance rate can be improved in simulations of both the harmonic oscillator and the Lennard-Jones system. This was accomplished while leaving the trajectory length and the number of integration steps unchanged.

We have also tested the HMC method with non-uniform step size on a third system, a C_α-C_β model for peptide helix formation[10] (see Text of the supplementary material[29]). Again, it was possible to increase the acceptance rate without changing τ or N_{MD} (see Table SI of the supplementary material[29]). For three quite different systems, we thus find that the performance of HMC can be improved by allowing for a non-uniform step size.

There is, to our knowledge, no existing procedure available for optimizing the step-size sequence along HMC trajectories, which would be useful to have. On the other hand, for all three systems we studied, very limited trial-and-error searches were sufficient to achieve improved performance compared to standard HMC with uniform step size. For the harmonic oscillator, we also performed a more systematic analysis, which supports the conclusion that step-size sequences with improved efficiency are easy to find.

In this article, we have mainly focused on the leapfrog integrator, which is of low (second) order but has the advantage of tolerating relatively large step sizes. A non-uniform step size can be used in HMC simulations based on other integrators as well. We briefly tested the HOH integrator[17] on the harmonic oscillator, and observed reductions in the average energy error comparable to those found with the leapfrog integrator.

To sum up, we have in this paper shown that the usual partitioning of HMC trajectories into steps of equal size need not be optimal. A change to non-uniform step size is easy to implement and can bring a noticeable reduction of energy errors at no extra cost.

ACKNOWLEDGMENTS

The simulations were performed on resources provided by the Swedish National Infrastructure for Computing (SNIC) at the LUNARC facility. A.B. acknowledges support from the Royal Physiographic Society in Lund.

[1] A. R. Leach, *Molecular Modelling: Principles and Applications*, 2nd ed. (Prentice Hall, 2001).
[2] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, Phys. Lett. B **195**, 216 (1987).
[3] M. Creutz, Phys. Rev. D **38**, 1228 (1988).
[4] S. Gupta, A. Irbäck, F. Karsch, and B. Petersson, Phys. Lett. B **242**, 437 (1990).
[5] S. Gupta, Nucl. Phys. B **370**, 741 (1992).
[6] A. D. Kennedy and B. Pendleton, Nucl. Phys. B **607**, 456 (2001).
[7] H. Meyer, H. Simma, R. Sommer, M. Della Morte, O. Witzel, and U. Wolff, Comput. Phys. Commun. **176**, 91 (2007).
[8] G. G. Batrouni, G. R. Katz, A. S. Kronfeld, G. P. Lepage, B. Svetitsky, and K. G. Wilson, Phys. Rev. D **32**, 2736 (1985).
[9] D. A. Sivak, J. D. Chodera, and G. E. Crooks, Phys. Rev. X **3**, 011007 (2013).
[10] A. Bhattacherjee and S. Wallin, Biophys. J. **102**, 569 (2012).
[11] M. Creutz and A. Gocksch, Phys. Rev. Lett. **63**, 9 (1989).
[12] M. Campostrini and P. Rossi, Nucl. Phys. B **329**, 753 (1990).
[13] E. Forest and R. D. Ruth, Physica D **43**, 105 (1990).
[14] H. Yoshida, Phys. Lett. A **150**, 262 (1990).
[15] I. P. Omelyan, I. M. Mryglod, and R. Folk, Phys. Rev. E **65**, 056706 (2002).
[16] T. Takaishi and P. de Forcrand, Phys. Rev. E **73**, 036706 (2006).
[17] C. Predescu, R. A. Lippert, M. P. Eastwood, D. Ierardi, H. Xu, M. Ø. Jensen, K. J. Bowers, J. Gullingsrud, C. A. Rendleman, R. O. Dror, and D. E. Shaw, Mol. Phys. **110**, 967 (2012).
[18] M. Tuckerman, B. J. Berne, and G. J. Martyna, J. Chem. Phys. **97**, 1990 (1992).
[19] J. C. Sexton and D. H. Weingarten, Nucl. Phys. B **380**, 665 (1992).
[20] S. J. Stuart, R. Zhou, and B. J. Berne, J. Chem. Phys. **105**, 1426 (1996).
[21] A. Irbäck, J. Chem. Phys. **101**, 1661 (1994).
[22] R. D. Engle, R. D. Skeel, and M. Drees, J. Comput. Phys. **206**, 432 (2005).
[23] C. R. Sweet, S. S. Hampton, R. D. Skeel, and J. A. Izaguirre, J. Chem. Phys. **131**, 174106 (2009).
[24] A. M. Horowitz, Phys. Lett. B **268**, 247 (1991).
[25] E. Akhmatskaya, N. Bou-Rabee, and S. Reich, J. Comput. Phys. **228**, 2256 (2009).
[26] J. A. Wagoner and V. S. Pande, J. Chem. Phys. **137**, 214105 (2012).
[27] B. Mehlig, D. W. Heermann, and B. M. Forrest, Phys. Rev. B **45**, 679 (1992).
[28] D. Frenkel and B. Smit, *Understanding Molecular Simulations: From Algorithms to Applications*, 2nd ed. (Academic Press, 2002).
[29] See supplementary material at http://dx.doi.org/10.1063/1.4862687 for three figures showing additional results from the simulations of the Lennard-Jones system, a table showing acceptance rates in the peptide simulations, and a description of the peptide model.

This article is copyrighted as indicated in the article. Reuse of AIP content is subject to the terms at: http://scitation.aip.org/termsconditions. Downloaded to IP: 130.235.189.230 On: Fri, 24 Jan 2014 14:59:42