

Adaptive Splitting Integrators for Enhancing Sampling Efficiency of Modified Hamiltonian Monte Carlo Methods in Molecular Simulation

Elena Akhmatskaya,^{*,†,‡} Mario Fernández-Pendás,[†] Tijana Radivojević,[†] and J.M. Sanz-Serna[¶]

BCAM - Basque Center for Applied Mathematics, IKERBASQUE, Basque Foundation for Science, and Universidad Carlos III de Madrid

E-mail: akhmatskaya@bcamath.org

Abstract

The modified Hamiltonian Monte Carlo (MHMC) methods, i.e., importance sampling methods that use modified Hamiltonians within a Hybrid Monte Carlo (HMC) framework, often outperform in sampling efficiency standard techniques such as molecular dynamics (MD) and HMC. The performance of MHMC may be enhanced further through the rational choice of the simulation parameters and by replacing the standard Verlet integrator with more sophisticated splitting algorithms. Unfortunately, it is not easy to identify the appropriate values of the parameters that appear in those algorithms. We propose a technique, that we call MAIA (Modified Adaptive Integration Approach), which, for a given simulation system and a given time step, automatically

*To whom correspondence should be addressed

[†]BCAM - Basque Center for Applied Mathematics, Alameda de Mazarredo 14, E-48009 Bilbao, Spain

[‡]IKERBASQUE, Basque Foundation for Science, María Díaz de Haro 3, E-48013 Bilbao, Spain

[¶]Departamento de Matemáticas, Universidad Carlos III de Madrid, Avenida de la Universidad 30, E-28911 Leganés (Madrid), Spain

selects the optimal integrator within a useful family of two-stage splitting formulas. Extended MAIA (or e-MAIA) is an enhanced version of MAIA, which additionally supplies a value of the method-specific parameter that, for the problem under consideration, keeps the momentum acceptance rate at a user-desired level. The MAIA and e-MAIA algorithms have been implemented, with no computational overhead during simulations, in MultiHMC-GROMACS, a modified version of the popular software package GROMACS. Tests performed on well-known molecular models demonstrate the superiority of the suggested approaches over a range of integrators (both standard and recently developed), as well as their capacity to improve the sampling efficiency of GSHMC, a noticeable method for molecular simulation in the MHMC family. GSHMC combined with e-MAIA shows a remarkably good performance when compared to MD and HMC coupled with the appropriate adaptive integrators.

Introduction

The role of numerical integrators in enhancing the performance of Hybrid/Hamiltonian Monte Carlo (HMC) has been a subject of active research in recent years.¹⁻⁴ It has been demonstrated that replacing the standard Verlet integrator with a splitting integrator specified by a suitable value of a parameter may significantly improve, for a range of step sizes, the conservation of the Hamiltonian and thus the acceptance rate of the proposals.^{1,2} Such integrators however possess shorter stability limits than the familiar Verlet algorithm.² In addition, the user is confronted with the problem of how best to choose the value of the parameter. The drawbacks of the use of splitting integrators more sophisticated than Verlet may be alleviated by resorting to the Adaptive Integration Approach, AIA.⁵ For a user-chosen time step, this approach automatically identifies an optimal, system-specific integrator, by using information on the highest frequencies of the harmonic interactions present in the system; this information is typically extracted from the input data intended for a molecular dynamics package. The term “optimal” refers to the fact that the selected integrator

minimizes, within a family of two-stage integrators, the expectation of the energy error for harmonic forces. When stability is an issue, AIA automatically chooses the Verlet integrator and, as the time step is reduced below the Verlet limit, AIA moves to more accurate integrators.

Another way to improve the performance of HMC is to introduce importance sampling as suggested in refs 6–11. Taking advantage of the fact that symplectic integrators preserve modified Hamiltonians more accurately than they preserve the true Hamiltonian, the authors proposed to sample with respect to modified/shadow Hamiltonians and to recover the desired distribution by reweighting. The resulting algorithms are capable of maintaining high acceptance rates and usually exhibit better efficiency than their predecessor HMC.^{11–14} Moreover, in many applications, the Verlet integrator is sufficient to provide an amount of accepted proposals adequate to generate good statistics, even with parameter settings for which HMC fails. Thus, the methods seem to be less sensitive than HMC to the choice of numerical integrator. However, it was shown that, for importance sampling HMC applied to high-dimensional statistical problems, replacing Verlet with optimized two-stage splitting integrators can improve the observed sampling efficiency by a factor of up to 4.¹¹ That reference, however, does not offer a recipe for the rational choice of the integrator for a given system and step size.

In this paper we propose an adaptive integration approach for molecular simulation, MAIA, which extends the ideas of AIA, in order to automatically select, for a given system and step size, the integrator with optimal conservation of the modified Hamiltonian, and therefore with the highest acceptance rate, in modified Hamiltonian Monte Carlo methods. Extended MAIA (or e-MAIA) offers the extra feature of a control on the stochasticity introduced in the momentum refreshment.

Modified Hamiltonian Monte Carlo Methods

The family of modified Hamiltonian Monte Carlo (MHMC) methods consists of HMC algorithms which, instead of sampling from the target canonical distribution

$$\pi(\mathbf{q}, \mathbf{p}) \propto \exp(-\beta H(\mathbf{q}, \mathbf{p})) \quad (1)$$

(where $\beta = 1/k_B T$ is the inverse temperature and k_B the Boltzmann constant) known up to a multiplicative constant, sample from an auxiliary importance canonical density

$$\tilde{\pi}(\mathbf{q}, \mathbf{p}) \propto \exp\left(-\beta \tilde{H}^{[k]}(\mathbf{q}, \mathbf{p})\right), \quad (2)$$

where H is a Hamiltonian and $\tilde{H}^{[k]}$ denotes a truncated modified Hamiltonian to be described later. The familiar notations \mathbf{q} and \mathbf{p} are used for positions and momenta, respectively. Such methods take advantage of two facts in order to enhance sampling efficiency of HMC. First, the closeness of $\tilde{H}^{[k]}$ to H makes it possible to implement an importance sampling approach and use samples of $\tilde{\pi}$ as a means toward computing expectations with respect to π . Second, the fact that the integrator preserves $\tilde{H}^{[k]}$ better than it preserves H leads to a more favorable value of the acceptance probability in the algorithms.

Symplectic integrators for the Hamiltonian dynamics with Hamiltonian function $H(\mathbf{q}, \mathbf{p})$, while not preserving the value of H exactly along the computed trajectory, do preserve exactly the value of a so-called modified Hamiltonian¹⁵⁻¹⁷

$$\tilde{H} = H + \Delta t H_2 + \Delta t^2 H_3 + \dots,$$

where Δt is the integration time step. For an integrator of order p , $\tilde{H} = H + \mathcal{O}(\Delta t^p)$, so that H_2, \dots, H_p vanish. In eq 2, $\tilde{H}^{[k]}$, $k > p$, is the truncation of \tilde{H} given by

$$\tilde{H}^{[k]} = H + \Delta t^p H_{p+1} + \dots + \Delta t^{k-1} H_k.$$

The expectation of the increment of H in an integration leg satisfies

$$\mathbb{E}_\pi[\Delta H] = \mathcal{O}(D\Delta t^{2p}), \quad (3)$$

(where D is the dimension), while

$$\mathbb{E}_{\tilde{\pi}}[\Delta \tilde{H}^{[k]}] = \mathcal{O}(D\Delta t^{2k}), \quad (4)$$

with $k > p$ (see ref 18 and section MAIA below) and therefore MHMC algorithms may benefit from high acceptance rates due to the better conservation of $\tilde{H}^{[k]}$.

The objective of a modified Hamiltonian Monte Carlo method is to sample from a distribution with probability density function

$$\pi(\mathbf{q}) \propto \exp(-\beta U(\mathbf{q})).$$

This is achieved indirectly, through sampling from the modified distribution (eq 2). In this paper we consider Hamiltonians of the form

$$H(\mathbf{q}, \mathbf{p}) = \frac{1}{2}\mathbf{p}^T M^{-1}\mathbf{p} + U(\mathbf{q}),$$

where M is a positive definite mass matrix and U is the potential, so that, under the target (eq 1), the variable \mathbf{q} has the marginal density $\propto \exp(-\beta U(\mathbf{q}))$.

Since in MHMC methods the samples are generated with respect to the modified or importance density, the computation of averages with respect to the target density after completion of the sampling procedure requires reweighting. If Ω_n , $n = 1, 2, \dots, N$, are the values of an observable along a sequence of states $(\mathbf{q}^n, \mathbf{p}^n)$ drawn from $\tilde{\pi}$ (eq 2), the averages

with respect to π (eq 1) are calculated as

$$\langle \Omega \rangle = \frac{\sum_{n=1}^N w_n \Omega_n}{\sum_{n=1}^N w_n},$$

where the importance weights are given by

$$w_n = \exp \left(-\beta (H(\mathbf{q}^n, \mathbf{p}^n) - \tilde{H}^{[k]}(\mathbf{q}^n, \mathbf{p}^n)) \right).$$

If the target density π and the importance density $\tilde{\pi}$ were not close, one would typically encounter high variability among weights, which would lead to large errors in the expectation, as many samples would not contribute significantly to the computation of $\langle \Omega \rangle$.

Let us now describe a generic MHMC algorithm. Given a sample (\mathbf{q}, \mathbf{p}) from the distribution $\tilde{\pi}$, the next sample $(\mathbf{q}^{\text{new}}, \mathbf{p}^{\text{new}})$ is defined as follows: (a) obtain the new momentum \mathbf{p}^* by applying a momentum update procedure that preserves the importance density $\tilde{\pi}$; (b) generate a proposal $(\mathbf{q}', \mathbf{p}')$ by simulating the Hamiltonian dynamics with Hamiltonian function H and with initial condition $(\mathbf{q}, \mathbf{p}^*)$ using a symplectic and reversible numerical integrator; (c) choose the next sample $(\mathbf{q}^{\text{new}}, \mathbf{p}^{\text{new}})$ to be $(\mathbf{q}', \mathbf{p}')$ (acceptance) with probability

$$\alpha = \min \left\{ 1, \frac{\tilde{\pi}(\mathbf{q}', \mathbf{p}')}{\tilde{\pi}(\mathbf{q}, \mathbf{p}^*)} \right\}. \quad (5)$$

Otherwise (rejection) set $(\mathbf{q}^{\text{new}}, \mathbf{p}^{\text{new}}) = (\mathbf{q}, -\mathbf{p}^*)$, i.e. carry out a momentum flip.

Since

$$\frac{\tilde{\pi}(\mathbf{q}', \mathbf{p}')}{\tilde{\pi}(\mathbf{q}, \mathbf{p}^*)} = \exp \left(-\beta \left(\tilde{H}^{[k]}(\mathbf{q}', \mathbf{p}') - \tilde{H}^{[k]}(\mathbf{q}, \mathbf{p}^*) \right) \right) = \exp \left(-\beta \Delta \tilde{H}^{[k]}(\mathbf{q}, \mathbf{p}^*) \right),$$

one may expect, in view of eqs 3 and 4, fewer rejections/momentum flips, and thus better sampling/more accurate dynamics when sampling with $\tilde{H}^{[k]}$ instead of H .^{8,19}

The first methods of the MHMC class were derived for atomistic simulations and dif-

ferred from each other in the ways of refreshing the momentum, computing modified Hamiltonians and integrating the Hamiltonian dynamics. For example, in (Separable) Shadow Hybrid Monte Carlo^{6,7} methods, a full momentum update is used, whereas in Targeted Shadow Hybrid Monte Carlo²⁰ and Generalized Shadow Hybrid Monte Carlo (GSHMC),⁸ suitable modifications of the partial momentum update of Horowitz²¹ are advocated in order to better mimic the dynamics and enhance sampling. More recent MHMC methods aim at specific applications, such as multiscale and mesoscale simulations (MTS-GSHMC and meso-GSHMC respectively)^{9,10} and computational statistics (Mix&Match Hamiltonian Monte Carlo).¹¹ As demonstrated in the original papers, for some particular problems, the use of MHMC methods resulted in a sampling efficiency several times higher than that observed with the conventional sampling techniques, such as MD, Monte Carlo (MC) and HMC. Further improvements can be achieved through the use of adaptive integration schemes as indicated in ref 5. The MAIA proposed in this work is an adaptive approach which can successfully replace the Verlet integrator in MHMC techniques with more powerful integration schemes.

In this paper we follow the momentum update procedure used in GSHMC,⁸ which is based on ideas from refs 21 and 22. We generate trial values

$$\begin{aligned}\mathbf{p}^{\text{trial}} &= \cos \varphi \mathbf{p} + \sin \varphi \mathbf{u} \\ \mathbf{u}^{\text{trial}} &= -\sin \varphi \mathbf{p} + \cos \varphi \mathbf{u}\end{aligned}\tag{6}$$

where $\varphi \in (0, \pi/2]$ is a parameter and the noise vector \mathbf{u} is drawn from the normal distribution $\mathcal{N}(0, \beta^{-1}M)$. A low value of φ will result in $\mathbf{p}^{\text{trial}}$ being close to \mathbf{p} and the behavior of the algorithm will be close to conventional MD. For φ near $\pi/2$, $\mathbf{p}^{\text{trial}}$ will be very different from \mathbf{p} . The proposed trial momentum $\mathbf{p}^{\text{trial}}$ is accepted ($\mathbf{p}^* = \mathbf{p}^{\text{trial}}$) with probability

$$\alpha_p = \min \left\{ 1, \frac{\hat{\pi}(\mathbf{q}, \mathbf{p}^{\text{trial}}, \mathbf{u}^{\text{trial}})}{\hat{\pi}(\mathbf{q}, \mathbf{p}, \mathbf{u})} \right\}\tag{7}$$

where $\hat{\pi}$ is the extended p.d.f.

$$\hat{\pi}(\mathbf{q}, \mathbf{p}, \mathbf{u}) \propto \exp(-\beta \hat{H}(\mathbf{q}, \mathbf{p}, \mathbf{u}))$$

corresponding to the extended Hamiltonian

$$\hat{H}(\mathbf{q}, \mathbf{p}, \mathbf{u}) = \tilde{H}^{[k]}(\mathbf{q}, \mathbf{p}) + \frac{1}{2} \mathbf{u}^T M^{-1} \mathbf{u}. \quad (8)$$

In case of rejection we set $\mathbf{p}^* = \mathbf{p}$.

In the present study, we use splitting methods to integrate the Hamiltonian dynamics. We use the symbols A and B to refer to the split systems

$$\dot{\mathbf{q}} = \mathbf{0}, \quad \dot{\mathbf{p}} = -\nabla U(\mathbf{q})$$

and

$$\dot{\mathbf{q}} = M^{-1} \mathbf{p}, \quad \dot{\mathbf{p}} = \mathbf{0},$$

respectively. These have solution flows

$$\phi_t^A(\mathbf{q}, \mathbf{p}) = (\mathbf{q}, \mathbf{p} - t \nabla U(\mathbf{q})) \quad (9)$$

and

$$\phi_t^B(\mathbf{q}, \mathbf{p}) = (\mathbf{q} + t M^{-1} \mathbf{p}, \mathbf{p}). \quad (10)$$

We limit our attention to the family of two-stage integrators where the map that advances the solution over a step of length Δt is defined as

$$\psi_{\Delta t} = \phi_{b\Delta t}^B \circ \phi_{\Delta t/2}^A \circ \phi_{(1-2b)\Delta t}^B \circ \phi_{\Delta t/2}^A \circ \phi_{b\Delta t}^B. \quad (11)$$

Here $0 < b < 1/2$ is a parameter that specifies the individual integrator within the family.

These integrators are symplectic, as compositions of flows of Hamiltonian systems, and reversible due to their palindromic structure. When $b = 1/4$, a step of length Δt of the integrator in eq 11 is equivalent to two successive steps of the velocity Verlet integrator each of length $\Delta t/2$.

When the error of the integrator is measured with respect to the modified Hamiltonian, the choice of parameter leading to the minimum error method is $b \approx 0.2306$;¹¹ we shall use the acronym M-ME (modified minimum error) to refer to this value. A minimization procedure parallel to the one used in ref 2 suggests the value $b \approx 0.2380$ ¹¹ that we shall call M-BCSS. Due to well-known properties of the Verlet method, the choice $b = 1/4$ yields the longest linear stability interval achievable in the family (eq 11). Therefore, generally speaking, one would select b close to $1/4$ if the value of Δt is so large that stability is the main concern and choose lower values of b , allowing for more accuracy in the conservation of the modified Hamiltonian, when the user is prepared to operate with smaller values of Δt .

In the next section we introduce algorithms which allow a rational choice of a value b in order to achieve the best accuracy of integration in a simulation of a given problem.

In our analysis we shall rely on the computationally efficient procedures for the calculation of modified Hamiltonians for the integrators of the family (eq 11) recently proposed in ref 11. The modified Hamiltonians are evaluated through numerical time derivatives $\nabla\dot{U}(\mathbf{q})$ of the gradient of the potential function U , which are computed from quantities available during the simulation (recalling that $-\nabla U(\mathbf{q})$ is equivalent to the force) using centered differences. Only two extra steps (one forward and another backward) are required. This makes feasible the use of those families of integrators in MHMC.

For the two-stage family (eq 11) of interest here the fourth-order modified Hamiltonian is given by¹¹

$$\tilde{H}^{[4]}(\mathbf{q}, \mathbf{p}) = \frac{1}{2}\mathbf{p}^T M^{-1}\mathbf{p} + U(\mathbf{q}) + \Delta t^2 \left(\lambda \mathbf{p}^T M^{-1} \nabla \dot{U}(\mathbf{q}) + \mu \nabla U(\mathbf{q})^T M^{-1} \nabla U(\mathbf{q}) \right), \quad (12)$$

with

$$\lambda = \frac{6b - 1}{24}, \quad \mu = \frac{6b^2 - 6b + 1}{12}.$$

Adaptive Algorithms

In this section we describe two adaptive algorithms: the Modified Adaptive Integration Approach (MAIA) and the extended MAIA (e-MAIA).

MAIA

MAIA is an algorithm which adapts the parameter b in eq 11 to the problem being solved and the value of Δt chosen by the user so as to maximize the expected acceptance rate α of the proposal $(\mathbf{q}', \mathbf{p}')$ in eq 5 or, equivalently, to minimize the expectation of the energy error function

$$\Delta \tilde{H}^{[4]}(\mathbf{q}, \mathbf{p}^*) = \tilde{H}^{[4]}(\mathbf{q}', \mathbf{p}') - \tilde{H}^{[4]}(\mathbf{q}, \mathbf{p}^*).$$

The analysis is based on a detailed study of the one-dimensional harmonic oscillator with potential $U(q) = (k/2) q^2$ ($k > 0$ a constant). For an integrator of the family in eq 11, the modified Hamiltonian in eq 12 takes the form

$$\tilde{H}^{[4]}(q, p) = \frac{1}{2} \frac{p^2}{M} + \frac{1}{2} k q^2 + \Delta t^2 \lambda \frac{k}{M^2} p^2 + \Delta t^2 \mu \frac{k^2}{M} q^2. \quad (13)$$

As shown in the section S1 in the Supporting Information, if $\omega = \sqrt{k/M}$ is the angular frequency of the harmonic oscillator and h denotes the nondimensional step size $h = \omega \Delta t$, then for the expected $\Delta \tilde{H}^{[4]}$ it holds

$$0 \leq \mathbb{E}[\Delta \tilde{H}^{[4]}] \leq \frac{1}{\beta} \rho(h, b), \quad (14)$$

where ρ is the function

$$\rho(h, b) = \frac{h^8 \left(b(12 + 4b(6b - 5)) + b(1 + 4b(3b - 2))h^2 - 2 \right)^2}{4(2 - bh^2)(4 + (2b - 1)h^2)(2 + b(2b - 1)h^2)(12 + (6b - 1)h^2)(6 + (1 + 6(b - 1)b)h^2)}. \quad (15)$$

Note that the expectation $\mathbb{E}[\Delta\tilde{H}^{[4]}]$ is taken with respect to the probability $\tilde{\pi}$ (eq 2) sampled by the algorithm. For a model consisting of D , possibly coupled, harmonic oscillators with angular frequencies ω_i , $i = 1, \dots, D$, the bound becomes

$$\mathbb{E}[\Delta\tilde{H}^{[4]}] \leq \frac{1}{\beta} \sum_{i=1}^D \rho(h_i, b),$$

with $h_i = \omega_i \Delta t$. Minimization of the right-hand side will therefore ensure optimal conservation of the modified Hamiltonian in the harmonic model.

In MAIA, given a physical problem which includes nonharmonic forces and a value of Δt , we estimate the fastest of the angular frequencies, $\tilde{\omega}$, of the two-body interactions⁵ and compute the nondimensional quantity

$$\tilde{h} = \sqrt{3}\tilde{\omega}\Delta t \quad (16)$$

($\sqrt{3}$ is a safety factor to be discussed presently). We then find the value of b that minimizes

$$\max_{0 < h < \tilde{h}} \rho(h, b). \quad (17)$$

Note that $(0, \tilde{h})$ is the shortest interval that contains all the values $h_i = \sqrt{3}\omega_i\Delta t$, where ω_i are the frequencies in the problem. In contrast to AIA,⁵ where the factor of $\sqrt{2}$ had to be used to avoid resonances of up to fourth order,²³ in MAIA, the factor $\sqrt{3}$, covering resonances of up to 5th order, was found to be appropriate.

The MAIA algorithm can be summarized as follows:

Given a physical system and a value of Δt , the MAIA algorithm determines the value of the parameter b to be used in eq 11 in the following way:

1. Find the periods or frequencies of all two-body interactions in the system. Determine the minimum period $\tilde{T} = 2\pi/\tilde{\omega}$, with the fastest frequency $\tilde{\omega}$, and compute the nondimensional quantity \tilde{h} in eq 16.
2. Check whether $\tilde{h} < 2\sqrt{2}$, which is the usual stability limit in molecular simulation for Verlet integrators.²⁴ If not, there is no value of b for which the scheme (eq 11) is stable for the attempted step size Δt and the integration is aborted.
3. Find the optimal value of the parameter b by minimizing eq 17 with the help of an optimization routine.

When Δt is “large” for the problem at hand, in the sense that stability is the primary concern, MAIA will choose $b = 1/4$, i.e. the Verlet integrator. Smaller values of Δt allow MAIA to reduce b and increase accuracy in the conservation of the modified Hamiltonian (see Figure 1). Figure 1 also shows the advantage of MAIA when compared with the older algorithm AIA,⁵ developed for the HMC method, which does not use modified Hamiltonians and samples with respect to the target canonical density.

e-MAIA

The overall performance of an MHMC method depends not only on the acceptance rate α of the proposal made after each MD integration leg (see eq 5), but also on the momentum update acceptance rate α_p in eq 7. The latter may play an important role in the quality of sampling,^{8,13} since α_p determines the frequency of the momenta resamplings. So far, we have looked for the integrator that maximizes α and our next objective is to find a way to control α_p .

As we did above, we build the analysis on the use of a harmonic oscillator model. For the scalar harmonic potential, the stationary marginal p.d.f.’s of the (stochastically independent




Figure 1: Parameter b for different integrators as a function of \tilde{h} (left) and bounds of the expected energy error measured with respect to the true, in solid lines, or modified Hamiltonian, in dashed lines (right). For more details, a log-log plot is also provided as an inset in the right panel. There are two lines for VV, as it may be used to sample from the true (HMC) or the importance density (GSHMC). AIA operates with respect to the true energy and MAIA with respect to its modified counterpart. Clearly the algorithms that operate with modified Hamiltonians possess smaller expected errors. This explains why, in general, VV GSHMC has higher acceptance rates than VV HMC and MAIA improves on AIA.

variables) p and u (see eq 6) are

$$\pi(p) \propto \exp\left(-\beta\left(\frac{1}{2}\frac{p^2}{M} + \Delta t^2 \lambda \frac{k}{M^2} p^2\right)\right), \quad \pi(u) \propto \exp\left(-\frac{\beta}{2}\frac{u^2}{M}\right), \quad (18)$$

respectively, and the extended Hamiltonian in eq 8 reads

$$\hat{H}(q, p, u) = \tilde{H}^{[4]}(q, p) + \frac{1}{2}\frac{u^2}{M},$$

with $\tilde{H}^{[4]}$ given in eq 13. As it was shown in ref 11, the difference in extended Hamiltonian satisfies

$$\begin{aligned} \Delta \hat{H} &= \hat{H}(q, p^{\text{trial}}, u^{\text{trial}}) - \hat{H}(q, p, u) \\ &= \Delta t^2 \lambda \left(\sin^2 \varphi \left(\frac{k}{M^2} u^2 - \frac{k}{M^2} p^2 \right) + 2 \cos \varphi \sin \varphi u \frac{k}{M^2} p \right), \end{aligned} \quad (19)$$

and from here it is found that

$$\mathbb{E}[\beta\Delta\hat{H}] = \Delta t^2 \beta \lambda \sin^2 \varphi \frac{\omega^2}{M} (\mathbb{E}[u^2] - \mathbb{E}[p^2]).$$

From eq 18 we have

$$\mathbb{E}[p^2] = \beta^{-1} M (1 + 2\Delta t^2 \lambda \omega^2)^{-1}, \quad \mathbb{E}[u^2] = \beta^{-1} M,$$

and then

$$\begin{aligned} \mathbb{E}[\beta\Delta\hat{H}] &= \Delta t^2 \lambda \sin^2 \varphi \omega^2 \left(1 - (1 + 2\Delta t^2 \lambda \omega^2)^{-1}\right) \\ &= \frac{2\Delta t^4 \lambda^2 \sin^2 \varphi \omega^4}{1 + 2\Delta t^2 \lambda \omega^2}. \end{aligned}$$

In terms of the dimensionless time step $h = \omega\Delta t$, one obtains

$$\mathbb{E}[\beta\Delta\hat{H}] = \frac{2h^4 \lambda^2 \sin^2 \varphi}{1 + 2h^2 \lambda}. \quad (20)$$

For the model consisting of D harmonic oscillators with angular frequencies ω_i , $i = 1, \dots, D$, the equivalent of eq 20 is

$$\mathbb{E}[\beta\Delta\hat{H}] = \sum_{i=1}^D \frac{2h_i^4 \lambda^2 \sin^2 \varphi}{1 + 2h_i^2 \lambda} \geq D \frac{2\bar{h}^4 \lambda^2 \sin^2 \varphi}{1 + 2\bar{h}^2 \lambda}, \quad (21)$$

where $h_i = \omega_i \Delta t$ are the dimensionless time steps and $\bar{h} = \bar{\omega} \Delta t$ with $\bar{\omega}$ equal to the slowest angular frequency among all the oscillators.

Using $\alpha_p \leq \exp(-\beta\Delta\hat{H})$ (see eq 7), from the inequality 21 and for a concrete choice of the angle φ_p , we can find the approximation

$$-\frac{\log \mathbb{E}[\alpha_p]}{D} \approx \frac{2\bar{h}^4 \lambda^2 \sin^2 \varphi_p}{1 + 2\bar{h}^2 \lambda}. \quad (22)$$

It has to be remarked that the fastest oscillation frequency features in the analyses of MAIA

and in its predecessor AIA⁵ but the *slowest* frequency is used in eq 22.

From eq 22, the expected acceptance rate in the momentum update may be controlled by three parameters: the parameter $\lambda = \lambda(b)$ that depends on the specific integrator being used, the parameter \bar{h} , which for a given problem is a function of Δt , and the angle φ . This fact motivates the algorithm that we call extended MAIA or e-MAIA. For a user-chosen Δt , e-MAIA first finds an integrator within the family of two-stage schemes that maintains the smallest expected modified energy error in the molecular dynamics part of the MHMC algorithm and then adjusts the value of φ so as to achieve a desired acceptance rate for the momentum update step. As explained above, the acceptance rates in the momentum update step depend on the choice of angle φ , whereas the MAIA analysis does not depend on φ . This means that, for some fixed values of φ and Δt , the integrator nominated by MAIA may not be favorable for maintaining an appropriate acceptance rate in the momenta. The goal of e-MAIA is to provide an adaptive choice of the angle φ to achieve a target, user-specified acceptance rate in the momentum update step while keeping the highest acceptance rate for positions.

While a high acceptance rate in the MD part has a positive effect on sampling with modified Hamiltonians, a too-frequent acceptance of momentum (close to 100 %) could lead to two undesired scenarios: (i) an accuracy deteriorating thermalization of the simulation, if the high acceptance rate is caused by a value of the angle φ very close to zero,^{13,19} or (ii) a disruption of the dynamical trajectories if the momenta are always resampled while φ is significantly bigger than zero.⁸ In the first scenario, the simulation will mimic an MD behavior in the NVE ensemble. The rationale for introducing e-MAIA is the possibility of simultaneously adapting the parameters b and φ to control both the acceptance probabilities α and α_p of the MD integration legs and the momentum updates.

The algorithm e-MAIA is as follows:

1. For a given physical problem, choose a time step Δt for the integration of the equations of motion, a target acceptance rate AR_p for the momentum update, and an initial value

φ_0 of the angle φ .

2. Find the slowest and the fastest angular frequencies in the harmonic interactions, $\bar{\omega}$ and $\tilde{\omega}$, respectively.
3. The integrator parameter b^* is obtained as in MAIA by optimization of the function ρ in eq 15. This choice of b^* guarantees the highest possible acceptance rate for harmonic interactions in the MD step.
4. The function which bounds the expected extended Hamiltonian error is given by (see eqs 21-22)

$$\tau(\bar{h}, b^*, \varphi) = \frac{2\bar{h}^4 \lambda^{*2} \sin^2 \varphi}{1 + 2\bar{h}^2 \lambda^*},$$

where λ^* is the value of λ when $b = b^*$ and

$$\bar{h} = \bar{\omega} \Delta t. \quad (23)$$

The angle φ^* is chosen as

$$\varphi^* = \arg \min_{\varphi \in (0, \pi/2]} \theta(\varphi), \quad (24)$$

with

$$\theta(\varphi) = \left| -\frac{\log(AR_p)}{D} - \tau(\bar{h}, b^*, \varphi) \right|.$$

5. If the selected φ^* is smaller than φ_0 , then either decrease the target AR_p and go to step 4 or, alternatively, define the function

$$\sigma(h, b, \varphi_0) = \rho(h, b) + \tau(h, b, \varphi_0) \quad (25)$$

and choose b^{**} that minimizes $\max_{0 < h < \bar{h}} \sigma(h, b, \varphi_0)$. (The fastest oscillation is used again for the momentum update part, since in this case we are constructing an upper bound of the expected energy error.)

We stress that for very small values of φ , an MHMC method loses its extra sampling abilities and behaves similarly to standard molecular dynamics. In e-MAIA this possibility is eliminated in step 5 of the algorithm in two optional ways. One way is to keep decreasing the target AR_p until φ^* rises above φ_0 . Another option is to optimize the joint bound function constructed for both expected errors, $\mathbb{E}[\beta\Delta\tilde{H}]$ and $\mathbb{E}[\beta\Delta\hat{H}]$. Though this sacrifices the position acceptance rates, the expected loss is small provided that $\varphi_0 \ll \pi/2$.

The reader should notice that, whereas MAIA in principle works for any method that samples with respect to modified Hamiltonians, e-MAIA only works for those MHMC methods which perform the momentum update step in the way described in this paper (eq 6).

Implementation

MultiHMC-GROMACS

MAIA and e-MAIA have been implemented in the MultiHMC-GROMACS^{5,25,26} software package. MultiHMC-GROMACS is a modified version of the popular molecular simulation code GROMACS^{27,28} and offers a set of recently developed algorithms that aim at improving the accuracy and sampling performance of the original GROMACS without sacrificing computational efficiency and parallel scaling properties. From the user’s point of view, there is no difference in setting up a simulation with MultiHMC-GROMACS and doing so with GROMACS, except for the need to specify a few additional parameters in the MultiHMC-GROMACS input file *.mdp*.

A detailed description of the MultiHMC-GROMACS package may be found elsewhere;^{5,25,26} here we briefly review those features that are relevant to the present study.

Hybrid Monte Carlo Methods

Hybrid Monte Carlo (HMC),²⁹ Generalized Hybrid Monte Carlo (GHMC)^{21,22} and Generalized Shadow Hybrid Monte Carlo (GSHMC)⁸ are all available in MultiHMC-GROMACS.

The implementation of GSHMC in GROMACS has been discussed in detail in refs 25 and 26. The other methods, i.e. HMC and GHMC, are implemented as special cases of GSHMC as described in ref 5.

Modified Hamiltonians

Two types of modified Hamiltonians are currently available in MultiHMC-GROMACS: the shadow Hamiltonian in a Lagrangian formulation⁸ for the Verlet/leapfrog integrator and the modified Hamiltonian (eq 12)¹¹ for two-stage integrators of the family eq 11.

Two-stage Integrators

The two-stage integrators of the family eq 11 have been implemented in MultiHMC-GROMACS as a concatenation of alternating updates of velocities and positions in the routine *do_md()* in *md.c*. The v-rescale thermostat³⁰ available in GROMACS has been adapted to work with the two-stage schemes.⁵ In addition, the two-stage integrators are optionally coupled with the SHAKE algorithm for simulation of constrained dynamics.⁵

At present MultiHMC-GROMACS may carry out simulations with the BCSS integrator,² the McLachlan's method,¹ their counterparts for modified Hamiltonians¹¹ and the velocity Verlet two-stage integrator. These integrators are given by specific choices of the parameter *b* in eq 11. In addition, the MultiHMC-GROMACS code implements the AIA integrator,⁵ where the algorithm tunes *b* to the particular simulation being carried out.

Implementing MAIA and e-MAIA

Similarly to AIA,⁵ MAIA and e-MAIA have been implemented in the GROMACS preprocessing module *grompp*. The preprocessing module is run only once before any simulation and, thus, does not introduce computational overheads in the simulation itself.

In the original GROMACS package, the module *grompp* reads the input files that contain the essential information about the simulated system, such as topology and structure, and

makes the necessary adaptation of that information for its use in the molecular dynamics module *mdrun*. The module *grompp* also checks the *.mdp* file for choices of input parameters for the simulation algorithms and, if necessary, generates warnings that allow users to reconsider the chosen setup. The *grompp* module finishes by producing the file *.tpr* to be used as input in the *mdrun* module for running molecular dynamics simulation.

In addition to these functionalities, a more advanced analysis of the harmonic interactions is included in *grompp* in MultiHMC-GROMACS. As has been explained in ref 5, the fastest harmonic interaction predetermines a maximal step size allowed for the stable numerical integration of the equations of motion. On the other hand, the slowest harmonic interactions are used in the e-MAIA algorithm to identify the best choice of the parameter φ . In MultiHMC-GROMACS, *grompp* searches for the periods corresponding to the fastest and slowest oscillations, \tilde{T} and \bar{T} respectively. The value \tilde{T} is used to define the upper limit of the dimensionless time step, $\tilde{h} = \sqrt{3}(2\pi/\tilde{T})\Delta t$, following the MAIA algorithm. The optimal value of the parameter b for a MAIA or e-MAIA integrator is then found as the argument that minimizes the maximum of ρ (eq 15) for the range of dimensionless time steps from zero to \tilde{h} . As in ref 5, the minimization is performed with a particle swarm optimization algorithm driven by a golden section search.³¹ The value \bar{T} is used to determine the angle φ , as explained in the e-MAIA algorithm.

Both b and φ are stored in the *input record* structure introduced by GROMACS for keeping all the input data during the whole simulation. Thus, b and φ can be accessed from every routine in the package.

The flowchart in Figure S1 summarizes MAIA and e-MAIA algorithms.

Numerical Experiments

In order to evaluate the efficiency of the proposed (e-)MAIA algorithm, we first compared its performance with that of several integration schemes which potentially can compete with

it. Then we estimated the performance of GSHMC combined with (e-)MAIA in comparison with other popular sampling methods. More precisely, e-MAIA was compared with fixed parameters integrators specifically derived for MHMC methods. The counterpart of BCSS² for modified Hamiltonians (M-BCSS) and the equivalent to the scheme of McLachlan¹ that minimizes the errors of modified Hamiltonians (M-ME) were included in the comparison. Both M-BCSS and M-ME have been recently derived¹¹ and implemented in MultiHMC-GROMACS. All three integrators were combined with the GSHMC method. In addition, e-MAIA was compared with integrators successfully used for molecular simulation in MD, HMC and GSHMC. The velocity Verlet and AIA combined with GSHMC were selected in this case. e-MAIA was compared with MAIA when both were implemented within GSHMC. GSHMC was compared with HMC and MD. For each tested sampling method the most efficient integrator was used: e-MAIA was chosen for GSHMC and AIA was employed in MD and HMC.

To provide a fair comparison, the following issues have been taken into account while producing the numerical results. To equalize the time spent on force calculations using Verlet and two-stage integrators, Verlet was always run with half a step size and twice the number of steps. Also, in the simulations with HMC and GSHMC, the number of Metropolis tests was kept constant regardless of the acceptance/rejection output. The computational overhead due to evaluation of modified Hamiltonians in GSHMC⁸ was taken into account by normalizing calculated integrated autocorrelation functions with respect to computational times. We notice that this overhead is, on average, of 1-2 % with respect to MD with the v-rescale thermostat or with respect to HMC, since both MD and HMC have the same computational cost. We also notice that the overheads of GSHMC tend to decrease when the trajectory lengths increase.

The tests were performed using two benchmark systems (see below), both run over a range of time steps Δt . The aim was to monitor the evolution of the parameters b and φ (eq 6) automatically chosen for each Δt in (e-)MAIA, and estimate their effect on the overall

sampling performance of GSHMC. In all plots in the Results, values of time steps correspond to two-stage integrators and assume twice smaller time steps for velocity Verlet.

Different lengths of MD trajectories L in GSHMC simulations were also tested. This parameter may play an important role in the sampling efficiency of GSHMC simulations when values chosen are either too small or too large.¹³ However, for the sake of clarity, in all tests presented in this work, the length of MD trajectories was fixed to 2000 steps when two-stage integrators were used and to 4000 otherwise. These values were found to be good choices for both GSHMC and HMC with different integration schemes and this is also confirmed by findings in ref 5. Also, as discussed above, for this trajectory length L the computational overheads of GSHMC w.r.t. MD are smaller than 1%.

With the obvious exception of e-MAIA, the angle used for the momentum refreshment (eq 6) was set to 0.2 for all tests unless stated otherwise.

Each individual test has been repeated 10 times and every result reported in this paper was obtained by averaging over the multiple runs to reduce statistical errors.

Benchmarks and Simulation Setup

The numerical experiments were performed using two benchmark systems: the coarse-grained VSTx1 toxin in a POPC bilayer³² and the atomistic 35-residue villin headpiece protein sub-domain.^{33,34} In the following we shall refer to these systems as toxin and villin, respectively.

Toxin is a coarse-grained system of 7810 particles. Four heavy particles were represented on average as one sphere.^{35,36} Coulomb and van der Waals interactions were solved using the shift algorithm. Both potential energies were shifted to 0 kJ mol⁻¹ at the radius of 1.2 nm. Periodic boundary conditions were considered in all directions. The target temperature was chosen to be 310 K and it was controlled in MD simulations by the v-rescale algorithm whereas no additional thermostat was required in HMC and GSHMC. No constraints were defined for this system. The total length of all simulations was 20 ns, which was sufficient for equilibration of the system for those choices of time steps that provided a stable integration.

The villin protein is a 9389 atoms system, composed of 389 atoms solvated with 3000 water molecules. Coulomb interactions were solved with the PME algorithm of order 6. van der Waals interactions were considered as in the toxin system but with a radius of 0.8 nm. Periodic boundary conditions were defined in all directions. As in the previous study,⁵ the bonds involving hydrogens were constrained using SHAKE/RATTLE. The use of these constraints does not affect the accuracy of simulations but allows for longer time steps which, due to the atomistic nature of the system, still are significantly shorter than in the coarse-grained toxin system. The temperature of 300 K was maintained using the v-rescale algorithm in MD and through the Metropolis tests in HMC and GSHMC.

Similarly to ref 5, an exhaustive study of the complete folding process of the villin protein is out of the scope of this work. Instead, the aim is to show the beneficial effect of (e-)MAIA on the accuracy and performance of the simulation of a constrained atomistic system. For our purposes, the length of 5 ns for each villin experiment was found sufficient. This simulation length was carefully chosen to make sure that the comparisons for simulations longer than those presented in this study are at least as favorable to the algorithms we suggest as we claim; they may be even more favorable as the simulations become longer. In the Supporting Information, in Figure S2 we show that the efficiency of GSHMC with e-MAIA, relative to MD, expressed in terms of radii of gyration, increases with simulation time.

Previously reported studies of the villin system have suggested that the use of a weak coupling thermostat and a barostat may lead to a better agreement with experiments.³⁷ Barostats are not considered, however, in this study, since the primary targets of the algorithms presented here are modified Hamiltonian Monte Carlo methods, which if no extra variations are introduced,^{8,26} sample in the NVT ensemble.

Results

Toxin

We start by measuring the acceptance rates of positions and momenta in the GSHMC simulations with different integration schemes. For the sake of clarity, we excluded from the plots the results for the MAIA algorithm, leaving only the data for e-MAIA. This makes sense because the position acceptance rates for MAIA and e-MAIA are always very similar (see step 3 of the e-MAIA algorithm), while e-MAIA has an obvious advantage over MAIA as far as the acceptance rates for momenta are concerned. We shall provide more details on this issue later.

The primary objective of the MAIA algorithm is to maximize the acceptance of position proposals in an MHMC method by minimizing the expected errors in modified Hamiltonians. Then, the first natural test for MAIA is to check whether the position acceptance rates observed in GSHMC simulations combined with MAIA are not below those observed with other two-stage integrators. In Figure 2, the effect of various integrators such as e-MAIA, the modified versions of BCSS² (M-BCSS) and ME¹ (M-ME),¹¹ the standard VV, and AIA⁵ on the acceptance rates in GSHMC simulations is investigated. The trends presented in the left plot are in good agreement with the theoretical prediction in Figure 1 (right panel). Indeed, the acceptance rates obtained with the modified adaptive approach e-MAIA, over the range of time steps considered, are never lower than the ones provided by the other integrators tested. For small time steps, all integrators, except AIA, guarantee high acceptance rates, but the situation changes as the time step increases and the shorter stability intervals of M-BCSS and M-ME result in acceptance rates well below those achieved with e-MAIA and VV. The low acceptance rates for AIA are not surprising, since this method was developed for sampling with respect to the true Hamiltonian and provides the lowest expected errors in Hamiltonian rather than in modified Hamiltonian. However, for the largest time step of 50 fs, the parameter b in AIA becomes equal to $1/4$ and thus AIA is equivalent to VV (see

Figure 1, left). The same applies to MAIA/e-MAIA for the longest time step, as can also be seen in Figure 1, left. This simply reflects the fact that the velocity Verlet integrator possesses the longest stability interval among the two-stage integrators and the adaptive methods AIA and MAIA select velocity Verlet when the time step goes beyond the stability limit of other two-stage integrators.




Figure 2: Toxin. Acceptance rates for positions (left) and momenta (right) observed in GSHMC simulations when using M-BCSS, M-ME, VV, AIA (all dashed lines) and e-MAIA (solid line). e-MAIA maintains the target AR_p of 90 % for each value of Δt (right).

The acceptance rates for momenta are shown in the right panel of Figure 2. For e-MAIA, we fixed the target acceptance AR_p to 90 % bearing in mind that too high (near 100 %) acceptance rates may degrade accuracy, whereas low acceptance rates normally reduce the sampling efficiency of GSHMC. With this target set, e-MAIA chose an appropriate value of φ for each time step being tested. The simulations with other integrators were run with the fixed value $\varphi = 0.2$. This was selected to achieve a good performance for the longest time steps. Obviously, with any integrator, the parameter φ can be adapted, by trial and error, to each simulation and time step, but we have to stress that, in practice, tuning blindly the value of φ is rather time-consuming and not necessarily results in the optimal choice of φ . That is why the ability of e-MAIA to optimize automatically such a choice is very welcome. As follows from Figure 2, right, for all tested time steps, e-MAIA maintained well the target AR_p by varying φ . The other integrators being combined with GSHMC led to very high,

unwelcome acceptance rates for most time steps tested.

We shall see next how the trends observed above for the acceptance rates impact the sampling efficiency of GSHMC. In the case of toxin, this efficiency was measured in terms of the integrated autocorrelation function IACF of the toxin drift d to the preferred interfacial location over the “convergence period”. The IACF is defined as

$$\text{IACF}^\Omega = \sum_{l=0}^{K'} \text{ACF}(l), \tag{26}$$

where $\text{ACF}(l)$, $l = 0, \dots, K' < K$ is the standard autocorrelation function for the time series Ω_k of K samples, $k = 1, \dots, K$.^{22,38} For GSHMC, the ACF’s are calculated taking into account the weights collected during simulations.¹¹ We notice that in all simulations performed the normalized weights are close to 1 due to small differences between modified and true Hamiltonians observed in the simulations as well as the choice of temperatures (common for molecular simulations of biological systems) leading to $\beta < 1$. This means that the metrics designed for weighted and nonweighted methods would not generate data that are too different. This, however, is not expected in a general case and is not common in statistical applications.³⁹ The IACF in eq 26 gives a quantitative measure of the time required, on average, to generate an uncorrelated sample. Low values of measured IACF imply low correlations between samples and thus more efficient sampling.

Figure 3, left, presents the IACFs (normalized with respect to computational time) obtained from GSHMC simulations using different integrators and time steps. Clearly, the simulations with e-MAIA provided the lowest values of IACFs and thus the best sampling for all choices of time step. All methods showed the good performance at $\Delta t = 40$ fs and, for this time step, the simulations with e-MAIA resulted in an efficiency (as measured by IACF) from 5 (vs M-BCSS, VV) to 9 (vs AIA) times higher than the simulations with other integration schemes. For the largest time step, $\Delta t = 50$ fs, the performance achieved using e-MAIA was 12 times better than in the simulations with M-BCSS and M-ME, but it did

not differ anymore from those observed in the simulations with VV and AIA, because for this long time step both AIA and e-MAIA chose velocity Verlet as integrator.




Figure 3: Toxin. Sampling efficiency of GSHMC combined with the integrators used in Figure 2. On the left, IACF of the drift, d , of the toxin to the preferred interfacial location evaluated as a function of Δt in GSHMC tests. On the right, the distribution of d observed in GSHMC simulations with various integrators using a time step of 30 fs. The solid black line (right) presents the “true” distribution produced with a ten times longer simulation (200 ns).

The right panel of Figure 3 compares the distributions of the distance d between the center of mass (c.o.m.) of the toxin and the c.o.m. of the bilayer, collected from simulations with $\Delta t = 30$ fs with different integrators, against the “true” distribution obtained from an MD simulation with velocity Verlet, over a time interval of length 200 ns, i.e., ten times longer. As for all tests in this section, the plots have results averaged over 10 repetitive runs. The curve corresponding to the simulation with e-MAIA (in red) shows the best match with the “true” distribution (in black).

The performances of e-MAIA and MAIA are compared in Figure 4. We chose the target AR_p in e-MAIA to be 90 % and the angle φ in MAIA to be equal to 1.1, which was the value found by e-MAIA for achieving the target $AR_p = 90$ % in GSHMC simulations at the smallest time step tested, $\Delta t = 20$ fs. Figure 4 reveals that, even though both e-MAIA and MAIA find the same integrator parameter b , leading to similar acceptance rates for positions, a good choice of the angle φ may visibly improve the sampling performance of GSHMC. The

improvement is by factors of 8 and 2 for $\Delta t = 40$ fs and $\Delta t = 50$ fs, respectively. The evolution, as the time step increases, of the optimal parameter φ as calculated by e-MAIA is also shown in Figure 4.




Figure 4: Toxin. e-MAIA (solid) vs MAIA (dashed). Acceptance rates for positions and momenta (left), IACFs (center) and the angle φ found by e-MAIA as a function of the time step (right) observed in GSHMC simulations. The angle φ used in MAIA was 1.1 and the target AR_p for e-MAIA was 90 %.

To finalize the numerical experiments on the toxin benchmark, we compared, using the normalized IACF metrics, the performance of three sampling methods, MD, HMC, and GSHMC. For each method, the best performing integrator was selected. Thus, GSHMC was combined with e-MAIA, based on the findings discussed above, whereas the AIA integrator was used for HMC and MD, according to the recommendations in ref 5. Figure 5, left, demonstrates the superiority of GSHMC over the other two methods, regardless the choice of time step. For the optimal choice of time step for this system, namely, $\Delta t = 40$ fs, the sampling efficiency of GSHMC is 4 times higher than that of HMC and 11 times better than that of MD. For the longest time step, $\Delta t = 50$ fs, the difference is even more dramatic and expressed in improvement factors of 17 and 30 over HMC and MD, respectively. Plotted in Figure 5, right, are the distributions of the distance d between the c.o.m. of the toxin and the c.o.m. of the bilayer produced by GSHMC, HMC, and MD simulations using a time step of 30 fs; they also confirm the better convergence of the GSHMC results to the “true” distribution.




Figure 5: Toxin. Sampling efficiency: GSHMC (e-MAIA) vs HMC (AIA) vs MD (AIA). The best integrator for each sampling method was employed. Sampling efficiency was measured by means of IACFs (left) and the distribution of the distance between the toxin and the membrane bilayer (right). The solid black line (right) presents the “true” distribution produced with a ten times longer simulation (200 ns).

Villin

As in the toxin case, we first inspected the acceptance rates for positions and momenta in GSHMC simulations with different integrators and found that the e-MAIA method worked as expected, i.e., provided the best position acceptance rates (Figure 6, left) and maintained the target momenta acceptance rate of 90 % (Figure 6, right) for all choices of time steps.




Figure 6: Villin. Acceptance rates for positions (left) and momenta (right) observed in GSHMC simulations when using M-BCSS, M-ME, VV, AIA (all dashed lines) and e-MAIA (solid line). e-MAIA maintains the target AR_p of 90 % for each value of Δt (right).

In contrast to the coarse-grained toxin benchmark, a quantitative analysis of the MAIA’s contribution to the GSHMC performance gain is not feasible with the atomistic villin benchmark. This is because such an analysis would require a long, computationally demanding series of simulations, for a range of time steps, integrators, and sampling methods. It is however possible to find evidence of the positive impact of MAIA on the sampling efficiency of GSHMC by using comparatively short simulation runs of 5 ns and metrics directly related to the quality of sampling.

One of such metrics is the radius of gyration (RG), which provides an estimation of the compactness of the desired structure, and is computed as

$$\text{RG} = \left(\frac{\sum_{i=1}^n \|r_i\|^2 M_i}{\sum_{i=1}^n M_i} \right)^{1/2},$$

where n is the number of atoms in the structure, r_i the distance between atom i and the center of mass of the structure, and M_i the mass of atom i . As in ref 37, we considered the experimental value of 0.94 nm as a target value and investigated the level of convergence to this value in short simulations when using different time steps, numerical integrators, and simulation methods.

Another metric used in this study relates to the positional root-mean-squared deviation (RMSD). The RMSD of a group of atoms in a molecule with respect to a reference structure can be calculated as

$$\text{RMSD} = \sqrt{\frac{1}{n} \sum_{i=1}^n \delta_i^2},$$

where δ_i is the distance between the positions of atom i in the two structures being compared.

Following the ideas from ref 37, we calculated the maximal RMSD of the α -carbon between any two visited structures in each simulation in order to judge the level of exploration of conformational space during the simulation.

In Figure 7 we plot, as functions of the time step, the radii of gyration and maximal RMSDs of the α -carbon calculated from the data collected in GSHMC simulations using

e-MAIA, M-BCSS, M-ME, VV and AIA integrators. Clearly, the simulations with e-MAIA (red solid line) produced the best approximations to the experimental data (left plot), the highest values of maximal RMSD (right plot) (implying better sampling) and the smallest performance degradation at the longest time steps.




Figure 7: Villin. Sampling efficiency of GSHMC combined with the integrators used in Figure 6: radius of gyration (left) and maximum RMSD of the α -carbon of the protein (right). The black solid line (left) represents the target experimental value of 0.94 nm.

The comparison of the results obtained using MAIA and e-MAIA in GSHMC simulations of villin confirmed the trends observed earlier in the toxin tests. Both methods achieved almost the same position acceptance rates, whereas the momenta acceptance rates were significantly higher in the simulations with e-MAIA (Figure 8, left). The latter was possible due to the automatic tuning of the parameter φ provided by e-MAIA for maintaining the target $AR_p = 90\%$ (Figure 8, right); its positive effect can be noticed in Figure 8, center.

Figure 9 compares the radii of gyration (left) and maximal RMSDs of the α -carbon (right) obtained from the simulations of villin using three different sampling methods, GSHMC, HMC, and MD. As in the toxin case, the best performing integrator was used for each sampler, i.e., e-MAIA was selected for GSHMC and AIA was combined with HMC and MD. For both metrics, GSHMC demonstrated the best results over the range of time steps. Its advantage over HMC and MD is most visible at longer time steps, when both HMC and MD lose accuracy and sampling efficiency.




Figure 8: Villin. e-MAIA (solid) vs MAIA (dashed). Acceptance rates for positions and momenta (left), radii of gyration (center) and the angle φ found by e-MAIA as a function of the time step (right) observed in GSHMC simulations. The angle φ used in MAIA was 0.9 and the target AR_p for e-MAIA was 90 %.




Figure 9: Villin. Sampling efficiency: GSHMC (e-MAIA) vs HMC (AIA) vs MD (AIA). The best integrator for each sampling method was employed. Sampling efficiency was measured through the radius of gyration (left) and the maximum RMSD of the α -carbon of the protein (right). The black solid line (left) represents the target experimental value of 0.94 nm.

Additionally, we have generated Ramachandran plots considering all residues of the protein except for glycine. In Figure 10 the Ramachandran plots, obtained for the largest time step $\Delta t = 6$ fs, are presented as two-dimensional joint distributions of φ and Ψ angles. Figure 10 confirms the advantages of GSHMC over other tested methods. Indeed, GSHMC combined with e-MAIA is the only method capable of sampling all regions including the less populated basins in the $\varphi, \Psi > 0$ region, which were out of reach for HMC and MD sampling.

Obviously, a deep atomistic study of the villin folding requires significantly longer runs than those presented here, as well as the incorporation of additional sampling techniques,

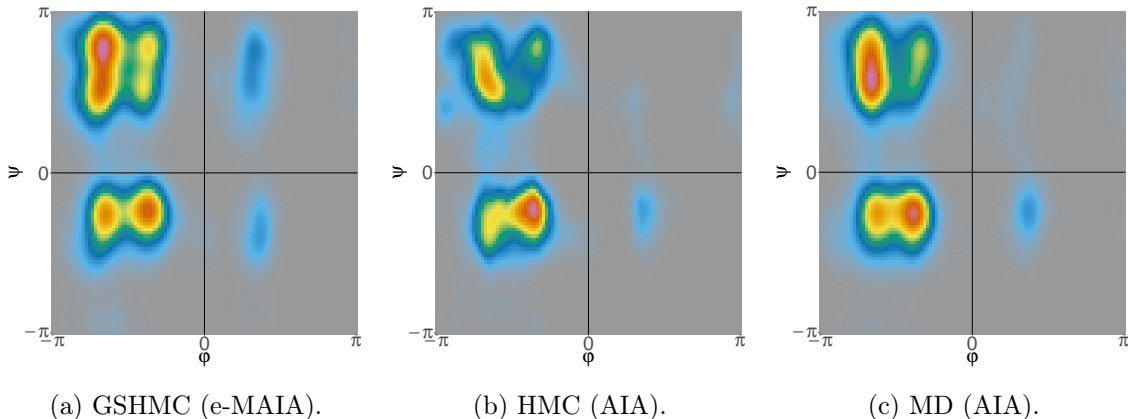


Figure 10: Villin. Sampling efficiency: GSHMC (e-MAIA) vs HMC (AIA) vs MD (AIA). Ramachandran plots for all residues of the protein except for glycine with φ torsion on the horizontal axis and Ψ on the vertical axis. The best integrator for each sampling method was employed. The time step was 6 fs, the largest in these tests.

such as, for example, parallel tempering, to the simulations. The latter can be implemented in a similar way, with similar cost for all three methodologies considered in our study. However, the simulations will clearly be more efficient if the underlying sampling method provides higher sampling efficiency, which is the case for GSHMC with e-MAIA.

Conclusions

We have proposed an adaptive approach for enhancing the accuracy and sampling efficiency of modified Hamiltonian Monte Carlo (MHMC) methods, with the aim of making them strong competitors of popular techniques such as molecular dynamics (MD) and hybrid Monte Carlo (HMC). Given a system to be simulated and a user-chosen time step, the new method, which we call Modified Adaptive Integration Approach or MAIA, identifies the two-stage numerical integrator which, when used in the Hamiltonian dynamics step of an MHMC method, provides the best conservation of the relevant modified Hamiltonian and thus the highest acceptance rate of the proposed trajectories. An enhanced variant of MAIA, e-MAIA, tailored to Generalized Shadow Hybrid Monte Carlo (GSHMC) methods, additionally supplies a value of the parameter φ that, for the problem under consideration, keeps

the momentum acceptance at a user-desired level. The MAIA algorithm has been implemented, with no computational overhead during simulations, in MultiHMC-GROMACS, a modified version of the popular software package GROMACS. The effect of the use of MAIA on the sampling efficiency of GSHMC has been demonstrated by using constrained atomistic and unconstrained coarse-grained benchmarks, and compared with the performance of other suitable integration schemes, including the popular velocity Verlet integrator. The tests revealed that the replacement in GSHMC of any fixed, two-stage integrator with e-MAIA leads systematically to improvements in sampling efficiency of up to an order of magnitude. The performance comparison of GSHMC, HMC, and MD combined with their best choices of numerical integrators (e-MAIA, AIA, AIA, respectively) confirmed the efficiency and robustness of the GSHMC-MAIA combination, whose advantages are especially noticeable when using the longest possible simulation time steps. For such cases, GSHMC, while maintaining good accuracy in simulation, provided a sampling efficiency (as measured with IACF) up to 30 times higher than the efficiency that may be achieved with MD.

Acknowledgements

The authors would like to thank the financial support from MTM2013-46553-C3-1-P funded by the Spanish Ministry of Economy and Competitiveness (MINECO). They have also received funding from the Project of MINECO with reference MTM2016-76329-R (AEI/FEDER, EU). This work has been possible thanks to the support of the computing infrastructure of the i2BASQUE academic network, and the technical and human support provided by IZO-SGI SGIker of UPV/EHU and European funding (ERDF and ESF). E.A. and T.R. thank support from Basque Government - ELKARTEK Programme, grant KK-2016/0002. M.F.P. would like to thank MINECO for funding through the fellowship BES-2014-068640. J.M.S.S. has been additionally supported by project MTM2016-77660-P(AEI/FEDER, UE) funded by MINECO (Spain). This research is also supported by the Basque Government through

the BERC 2014-2017 program and by MINECO: BCAM Severo Ochoa accreditation SEV-2013-0323.

References

- (1) McLachlan, R. I. On the Numerical Integration of Ordinary Differential Equations by Symmetric Composition Methods. *SIAM J. Scient. Comp.* **1995**, *16*, 151–168.
- (2) Blanes, S.; Casas, F.; Sanz-Serna, J. M. Numerical Integrators for the Hybrid Monte Carlo Method. *SIAM J. Scient. Comp.* **2014**, *36*, A1556–A1580.
- (3) Chao, W.-L.; Solomon, J.; Michels, D. L.; Sha, F. Exponential Integration for Hamiltonian Monte Carlo. International Conference on Machine Learning – ICML 2015. 2015; pp 1142–1151.
- (4) Campos, C. M.; Sanz-Serna, J. M. Palindromic 3-stage splitting integrators, a roadmap. *J. Comput. Phys.* **2017**, *346*, 340–355.
- (5) Fernández-Pendás, M.; Akhmatskaya, E.; Sanz-Serna, J. M. Adaptive multi-stage integrators for optimal energy conservation in molecular simulations. *J. Comput. Phys.* **2016**, *327*, 434–449.
- (6) Izaguirre, J. A.; Hampton, S. S. Shadow hybrid Monte Carlo: an efficient propagator in phase space of macromolecules. *J. Comput. Phys.* **2004**, *200*, 581–604.
- (7) Sweet, C. R.; Hampton, S. S.; Skeel, R. D.; Izaguirre, J. A. A separable shadow Hamiltonian hybrid Monte Carlo method. *J. Chem. Phys.* **2009**, *131*, 174106.
- (8) Akhmatskaya, E.; Reich, S. GSHMC: An efficient method for molecular simulation. *J. Comput. Phys.* **2008**, *227*, 4934–4954.
- (9) Escribano, B.; Akhmatskaya, E.; Reich, S.; Azpiroz, J. M. Multiple-time-stepping generalized hybrid Monte Carlo methods. *J. Comput. Phys.* **2015**, *280*, 1–20.

- (10) Akhmatskaya, E.; Reich, S. Meso-GSHMC: A stochastic algorithm for meso-scale constant temperature simulations. *Procedia Computer Science* **2011**, *4*, 1353–1362.
- (11) Radivojević, T. Enhancing Sampling in Computational Statistics Using Modified Hamiltonians. Ph.D. thesis, University of the Basque Country (UPV/EHU), Bilbao, 2016.
- (12) Akhmatskaya, E.; Reich, S. New Hybrid Monte Carlo Methods for Efficient Sampling: from Physics to Biology and Statistics (Selected Papers of the Joint International Conference of Supercomputing in Nuclear Applications and Monte Carlo: SNA + MC 2010). *Progress in nuclear science and technology* **2011**, *2*, 447–462.
- (13) Wee, C. L.; Sansom, M. S. P.; Reich, S.; Akhmatskaya, E. Improved Sampling for Simulations of Interfacial Membrane Proteins: Application of Generalized Shadow Hybrid Monte Carlo to a Peptide Toxin/Bilayer System. *J. Phys. Chem. B* **2008**, *112*, 5710–5717.
- (14) Escribano, B.; Lozano, A.; Radivojević, T.; Fernández-Pendás, M.; Carrasco, J.; Akhmatskaya, E. Enhancing sampling in atomistic simulations of solid-state materials for batteries: a focus on olivine NaFePO₄. *Theor. Chem. Acc.* **2017**, *136*, 43.
- (15) Sanz-Serna, J. M.; Calvo, M. P. *Numerical Hamiltonian problems*, 1st ed.; Applied Mathematics and Mathematical Computation 7; Chapman & Hall: London, 1994.
- (16) Leimkuhler, B. J.; Reich, S. *Simulating Hamiltonian dynamics*; Cambridge monographs on applied and computational mathematics; Cambridge Univ.: Cambridge, 2004.
- (17) Hairer, E.; Lubich, C.; Wanner, G. *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations; 2nd ed.*; Springer: Dordrecht, 2006.

- (18) Beskos, A.; Pillai, N.; Roberts, G.; Sanz-Serna, J. M.; Stuart, A. Optimal tuning of the hybrid Monte Carlo algorithm. *Bernoulli* **2013**, *19*, 1501–1534.
- (19) Akhmatskaya, E.; Bou-Rabee, N.; Reich, S. A comparison of generalized hybrid Monte Carlo methods with and without momentum flip. *J. Comput. Phys.* **2009**, *228*, 2256–2265.
- (20) Akhmatskaya, E.; Reich, S. In *New Algorithms for Macromolecular Simulation*; Leimkuhler, B., Chipot, C., Elber, R., Laaksonen, A., Mark, A., Schlick, T., Schütte, C., Skeel, R., Eds.; Springer Berlin Heidelberg: Berlin, Heidelberg, 2006; pp 141–153.
- (21) Horowitz, A. M. A generalized guided Monte Carlo algorithm. *Phys. Lett. B* **1991**, *268*, 247–252.
- (22) Kennedy, A.; Pendleton, B. Cost of the generalised hybrid Monte Carlo algorithm for free field theory. *Nucl. Phys. B* **2001**, *607*, 456–510.
- (23) Schlick, T.; Mandziuk, M.; Skeel, R. D.; Srinivas, K. Nonlinear Resonance Artifacts in Molecular Dynamics Simulations. *J. Comput. Phys.* **1998**, *140*, 1–29.
- (24) Mazur, A. K. Common Molecular Dynamics Algorithms Revisited: Accuracy and Optimal Time Steps of Störmer–Leapfrog Integrators. *J. Comput. Phys.* **1997**, *136*, 354–365.
- (25) Escribano, B.; Akhmatskaya, E.; Mujika, J. I. Combining stochastic and deterministic approaches within high efficiency molecular simulations. *Central European J. Math.* **2013**, *11*, 787–799.
- (26) Fernández-Pendás, M.; Escribano, B.; Radivojević, T.; Akhmatskaya, E. Constant pressure hybrid Monte Carlo simulations in GROMACS. *J. Mol. Model.* **2014**, *20*, 2487.
- (27) Berendsen, H.; van der Spoel, D.; van Drunen, R. GROMACS: A message-passing parallel molecular dynamics implementation. *Comp. Phys. Comm.* **1995**, *91*, 43–56.

- (28) Hess, B.; Kutzner, C.; van der Spoel, D.; Lindahl, E. GROMACS 4: Algorithms for Highly Efficient, Load-Balanced, and Scalable Molecular Simulation. *J. Chem. Theo. Comp.* **2008**, *4*, 435–447.
- (29) Duane, S.; Kennedy, A.; Pendleton, B. J.; Roweth, D. Hybrid Monte Carlo. *Phys. Lett. B* **1987**, *195*, 216–222.
- (30) Bussi, G.; Donadio, D.; Parrinello, M. Canonical sampling through velocity rescaling. *J. Chem. Phys.* **2007**, *126*, 014101.
- (31) Oh, S.; Hori, Y. Development of Golden Section Search Driven Particle Swarm Optimization and its Application. 2006 SICE-ICASE International Joint Conference. 2006; pp 2868–2873.
- (32) Jung, H. J.; Lee, J. Y.; Kim, S. H.; Eu, Y.-J.; Shin, S. Y.; Milesco, M.; Swartz, K. J.; Kim, J. I. Solution Structure and Lipid Membrane Partitioning of VSTx1, an Inhibitor of the KvAP Potassium Channel,. *Biochemistry* **2005**, *44*, 6015–6023.
- (33) Bazari, W. L.; Matsudaira, P.; Wallek, M.; Smeal, T.; Jakes, R.; Ahmed, Y. Villin sequence and peptide map identify six homologous domains. *Proceedings of the National Academy of Sciences* **1988**, *85*, 4986–4990.
- (34) McKnight, C. J.; Matsudaira, P. T.; Kim, P. S. NMR structure of the 35-residue villin headpiece subdomain. *Nature Structural & Molecular Biology* **1997**, *4*, 180–184.
- (35) Wallace, E. J.; Sansom, M. S. P. Carbon Nanotube/Detergent Interactions via Coarse-Grained Molecular Dynamics. *Nano Lett.* **2007**, *7*, 1923–1928.
- (36) Shih, A. Y.; Arkhipov, A.; Freddolino, P. L.; Schulten, K. Coarse Grained Protein-Lipid Model with Application to Lipoprotein Particles. *J. Phys. Chem. B* **2006**, *110*, 3674–3684.

- (37) van der Spoel, D.; Lindahl, E. Brute-Force Molecular Dynamics Simulations of Villin Headpiece: Comparison with NMR Parameters. *J. Phys. Chem. B* **2003**, *107*, 11178–11187.
- (38) Allen, M. P.; Tildesley, D. J. *Computer Simulation of Liquids*; Clarendon Press: New York, NY, USA, 1989.
- (39) Radivojević, T.; Akhmatkaya, E. Mix & Match Hamiltonian Monte Carlo. *ArXiv e-print 1706.04032* **2017**,

Graphical TOC Entry

