

# HMC: avoiding rejections by not using leapfrog and some results on the acceptance rate

M.P. Calvo\*, D. Sanz-Alonso<sup>†</sup> and J.M. Sanz-Serna<sup>‡</sup>

December 11, 2019

## Abstract

We give numerical evidence that the standard leapfrog algorithm may not be the best integrator to use within the Hamiltonian Monte Carlo (HMC) method and its variants. When the dimensionality of the target distribution is high, the number of accepted proposals obtained with a given computational effort may be multiplied by a factor of three or more by switching to alternative integrators very similar to leapfrog. Such integrators have appeared in the numerical analysis and molecular dynamics literature rather than in the statistics literature. In addition, we provide some theoretical results on the acceptance rate of HMC.

**Keywords:** Hamiltonian Monte Carlo, numerical integrators, expected acceptance rate, expected energy error.

## 1 Introduction

In this paper we provide numerical evidence that the standard leapfrog algorithm may not be the best integrator to use within the Hamiltonian Monte Carlo (HMC) method and its variants. If the dimensionality of the target distribution is high, the number of accepted proposals obtained *with a given computational effort* may be multiplied by a factor of three or more by

---

\*Departamento de Matemática Aplicada e IMUVA, Facultad de Ciencias, Universidad de Valladolid, Spain. E-mail: mariapaz.calvo@uva.es

<sup>†</sup>Department of Statistics, University of Chicago, 5747 South Ellis Avenue, Chicago, Illinois 60637, US. E-mail: sanzalonso@uchicago.edu

<sup>‡</sup>Departamento de Matemáticas, Universidad Carlos III de Madrid, Avenida de la Universidad 30, E-28911 Leganés (Madrid), Spain. E-mail: jmsanzserna@gmail.com

switching to alternative, more accurate integrators very similar to leapfrog. Such integrators have appeared in the numerical analysis and molecular dynamics literature rather than in the statistics literature. In addition, we provide some theoretical results on the acceptance rate of HMC.

HMC, originally suggested by Duane et al. (1987), is a popular Markov chain Monte Carlo (MCMC) algorithm where proposals are obtained by numerically integrating a Hamiltonian system of differential equations (Neal, 2011; Sanz-Serna, 2014). This integration gives HMC its main potential advantage: the possibility of proposals that are far away from the current state of the chain and, at the same time, may be accepted with high probability. At present the leapfrog/Störmer/Verlet scheme is almost universally used to integrate the Hamiltonian dynamics and obtain a proposal that may be accepted or rejected. While rejections increase the correlation in any MCMC algorithm, they are particularly unwelcome in the case of HMC, because of the computational effort wasted in the numerical integration legs used to generate the rejected proposals (Campos and Sanz-Serna, 2015).

We emphasize that, when comparing different integrators, it is essential to take into account the computational effort. If the numerical integration of the Hamiltonian dynamics were exact, then all proposals would be accepted. By implication, for any reasonable (i.e. convergent) numerical integrator, it is possible to make the acceptance rate arbitrarily close to 100% by reducing the step-length  $\epsilon$  and increasing the number  $L$  of time-steps used to simulate the dynamics over a given time-interval  $0 \leq \tau \leq \tau_{end}$ . However such a reduction in  $\epsilon$  may not be advisable as it implies a corresponding increase in computational cost. In fact, *once an integrator has been chosen*, a very high acceptance rate signals that an inappropriately small value of  $\epsilon$  is being used and the computational cost of generating each proposal is too high; one would sample better at the same cost exploring the state space further by increasing the number  $N$  of steps in the Markov chain and solving the Hamilton equations less accurately in each proposal. On the other hand, *when choosing an integrator* one should prefer (all other things being equal) the method that maximizes the number of accepted proposals that may be obtained with a given computational cost.

Section 2 is introductory and contains a description of HMC and of the properties of the leapfrog algorithm. It also serves to recap a number of considerations that are essential to understand the remainder of the paper. Section 3 describes a two-parameter family of integrators that contains the leapfrog algorithm as a particular case. All members of the family may be implemented as easily as leapfrog and all use the same accept/reject strategy. It is not possible to identify analytically the algorithm of this family that

is best for the HMC method; in fact the acceptance rate depends on the so-called global errors of the integrators and global errors (as distinct from local errors) are extremely difficult to pin down analytically (Bou-Rabee and Sanz-Serna (2018) provide a short introduction to the basic concepts in numerical integration). For this reason we conduct numerical experiments to compare different integrators (Section 4). We employ three test problems: a simple multivariate Gaussian distribution, a Log-Gaussian Cox model and an example from molecular dynamics: the canonical distribution of an alkane molecule. The last distribution possesses many modes (corresponding to alternative equilibrium configurations of the molecule) and has been chosen in view of the complexity of the nonlinear forces. In all test problems the dimensionality  $d$  of the target distribution is high, ranging from 15 to 4096; in fact as  $d$  increases so does the computational cost of generating proposals and therefore it is more important to identify efficient integrators. The experiments show that several integrators of the family improve on leapfrog. Specially an integrator specifically designed by Blanes et al. (2014) to be used in the HMC context may provide, for a given computational effort, three times more accepted proposals than leapfrog.

In all experiments we found that the mean acceptance rate is (approximately) given by the formula (Gupta et al., 1990)  $a = 2\Phi(-\sqrt{\mu/2})$ , where  $\mu$  is the mean energy error and  $\Phi$  the standard normal cumulative distribution function. The high-dimensional asymptotic validity of this formula has been rigorously established by (Beskos et al., 2013) under the very restrictive assumption that the target is a product of independent identical copies of a given distribution. In Section 5 we prove (Theorem 2) that the formula also holds asymptotically in a general scenario for Gaussian distributions. In addition, we study (Theorem 1) the acceptance rate in the simplest model problem given by a univariate Gaussian target. The theorems have a clear implication for the construction of Hamiltonian integrators tailored to be used within HMC: reducing the expected energy error is the only way to improve the expected acceptance rate, which validates the approach used by Blanes et al. (2014).

Our contributions are summarized in the final Section 6.

We conclude this introduction with some caveats. A comparison of the merits of HMC with those of other sampling techniques or of the advantages of the various variants of HMC is completely outside our scope here; a recent contribution devoted to such comparisons is the paper by Radivojević and Akhmatkaya (2019), which also provides an extensive bibliography. Similarly, we are not concerned with discussing issues of HMC not directly related to inaccurate time-integrations. Some of those issues concern the ergodic-

ity of the method (Bou-Rabee and Sanz-Serna, 2017; Cances et al., 2007; Livingstone et al., 2019), the choice of the length  $\tau_{end}$  of the time-interval for simulating the dynamics (Hoffman and Gelman, 2014), the difficulties in leaving the neighbourhood of a probability mode in multimodal targets (Mangoubi et al., 2018) and the scalability of the method to large data-sets (Chen et al., 2014). Finally, while we limit the exposition to the simplest version of HMC, our study may be applied to other variants (Akhmatskaya et al., 2017; Horowitz, 1991; Radivojević et al., 2018) and even to molecular dynamics simulations with no accept/reject mechanism in place (Fernández-Pendás et al., 2016).

## 2 Hamiltonian Monte Carlo

We denote by  $\theta \in \mathbb{R}^d$  the random variable whose distribution we wish to sample and by  $\pi(\theta)$  the corresponding density. HMC is based on the introduction of an independent auxiliary variable  $p \in \mathbb{R}^d$  with density  $\mathcal{N}(p|0, M)$ . The joint density and the negative log-likelihood are then, respectively,  $\pi(\theta, p) = \pi(\theta)\mathcal{N}(p|0, M)$  and

$$H(\theta, p) = -\mathcal{L}(\theta) + \frac{1}{2}p^T M^{-1}p + \frac{1}{2} \log((2\pi)^d \det(M)),$$

where  $\mathcal{L}(\theta) = \log \pi(\theta)$ .<sup>1</sup> In a mechanical analogy, the function  $H$  is the Hamiltonian (total energy) of a system with potential energy  $-\mathcal{L}(\theta)$  and kinetic energy  $p^T M^{-1}p/2$ ;  $\theta$  are the positions,  $p$  the momenta and  $M$  is the mass matrix. Hamilton’s differential equations for the system are then given by

$$\frac{d}{d\tau}\theta = +\frac{\partial H}{\partial p} = M^{-1}p, \quad \frac{d}{d\tau}p = -\frac{\partial H}{\partial \theta} = \nabla_{\theta}\mathcal{L}(\theta), \quad (1)$$

where  $\tau$  represents time, and have the property that, for each  $\tau$ , their solution flow

$$(\theta(0), p(0)) \mapsto (\theta(\tau), p(\tau))$$

preserves the joint density,  $\pi(\theta(\tau), p(\tau)) = \pi(\theta(0), p(0))$ . HMC is a Markov chain Monte Carlo algorithm where at each step of the chain one integrates the system (1). Since, in practical applications, the exact solution flow cannot be obtained analytically, one has to resort to a numerical integration whose error is removed by an accept/reject mechanism. At present the

---

<sup>1</sup>The constant  $(1/2) \log((2\pi)^d \det(M))$  may be suppressed without altering the Hamiltonian dynamics or the acceptance probability.

leapfrog/Störmer/Verlet algorithm is the integrator of choice. Over a single time-step of length  $\epsilon > 0$ , the integrator reads

$$p \leftarrow p + \frac{\epsilon}{2} \nabla_{\theta} \mathcal{L}(\theta), \quad (2)$$

$$\theta \leftarrow \theta + \epsilon M^{-1} p, \quad (3)$$

$$p \leftarrow p + \frac{\epsilon}{2} \nabla_{\theta} \mathcal{L}(\theta). \quad (4)$$

In molecular dynamics, the substeps (2) and (4) are referred to as *kicks*; the momentum  $p$  undergoes an increment while the position  $\theta$  stands still. The substep (3) is called a *drift*; the system changes position without changing momentum.

The bulk of the computational effort of the leapfrog integrator typically lies in the evaluation of the gradient (score)  $\nabla_{\theta} \mathcal{L}(\theta)$ . While, in a single time-step, an evaluation is required by (2) and a second evaluation is required by (4), an integration leg consisting of  $L$  time-steps only demands  $L + 1$  evaluations, because the gradient in (4) is the same as the gradient in (2) at the next time-step. Thus the computational cost of the leapfrog integrator is essentially one gradient evaluation per time-step.

At each step of the Markov chain the Hamiltonian dynamics over an interval  $0 \leq \tau \leq \tau_{end} = L\epsilon$  are simulated by taking  $L$  time-steps of length  $\epsilon$  of the integrator starting from  $(\theta, p)$ , where  $\theta$  is the last sample and  $p$  is drawn from  $\mathcal{N}(0, M)$ . The numerical solution  $(\theta^*, p^*)$  at the end of the  $L$  time-steps is a (deterministic) proposal, which is accepted with probability

$$a = \min \left( 1, \exp \left( - [H(\theta^*, p^*) - H(\theta, p)] \right) \right). \quad (5)$$

It is hoped that, by choosing  $\tau_{end}$  appropriately,  $\theta^*$  will be far from  $\theta$  so as to reduce correlation and enhance the exploration of the state space.

In (5), the difference

$$\Delta H(\theta, p) = H(\theta^*, p^*) - H(\theta, p) \quad (6)$$

(recall that  $\theta^*$  and  $p^*$  are deterministic functions of  $(\theta, p)$ ) is the *increment* of the Hamiltonian function over the numerical integration leg. For the true solution, the value of  $H$  remains constant as  $\tau$  varies (conservation of energy) and therefore the *initial* energy  $H(\theta, p)$  coincides with the value of the Hamiltonian in the true dynamics at the *end* point  $\tau_{end}$ . Thus, (6) is also the *error* in  $H$  at  $\tau_{end}$  introduced by the numerical integrator. Since leapfrog is a *second order method*, in the limit where  $\epsilon \rightarrow 0$  and  $L \rightarrow \infty$  with  $L\epsilon = \tau_{end}$  fixed, it provides numerical approximations to the true values of

the Hamiltonian solutions that have  $\mathcal{O}(\epsilon^2)$  errors.<sup>2</sup> It follows that, for fixed  $(\theta, p)$ , (6) is also  $\mathcal{O}(\epsilon^2)$  in that limit. As a consequence, once  $\tau_{end}$  has been chosen to generate suitable proposals, the acceptance probability (5) may become arbitrarily close to 100% by decreasing  $\epsilon$  (but this will imply an increase of the computational cost of generating the proposal, as more time-steps have to be taken to span the interval  $[0, \tau_{end}]$ ). The sentence ‘‘HMC may generate proposals away from the current state that may be accepted with high probability’’ has been repeated again and again to promote HMC.

For given  $\tau_{end}$ , how should  $\epsilon$  be chosen? If  $\epsilon$  is too large, many rejections will take place and the chain will be highly correlated. On the other hand, a very high empirical acceptance rate signals that an inappropriately small value of  $\epsilon$  is being used and the computational cost of generating each proposal is too high; one would sample better at the same cost by increasing the number  $N$  of steps in the Markov chain and solving the Hamilton equations less accurately in each proposal. According to the analysis in Beskos et al. (2013), that only applies in a restrictive model scenario where the target is a product of independent copies of a fixed distribution, one should tune  $\epsilon$  so as to observe an acceptance rate of approximately 65%. The result is asymptotic as  $d \uparrow \infty$  and that reference recommends that in practice one aims at acceptance rates larger than that. This discussion will be retaken below, when commenting on Figures 5 and 6.

We emphasize that in an HMC context numerical integrations do not need to be very accurate. Even if the energy error in an integration leg is as large as, say, +0.5 (orders of magnitude larger than the accuracy usually sought when integrating numerically differential equations), then the proposal will be accepted with probability  $\exp(-0.5) \approx 0.61$ . In addition, negative energy errors always result in acceptance. Nevertheless it is important to recall that the Hamiltonian is an extensive quantity whose value is added when mechanical systems are juxtaposed. If the vector random variable  $\theta$  has  $d$  independent scalar components and  $M$  is diagonal, then the Hamiltonian for  $\theta$  is the sum of the Hamiltonians for the individual components  $(\theta_i, p_i)$  and the Hamiltonian system of differential equations for  $(\theta, p)$  is the uncoupled juxtaposition of the one-degree-of-freedom Hamiltonian equations for the  $(\theta_i, p_i)$ . In that model scenario, an acceptance probability of  $\exp(-0.5) \approx 0.61$  would require that each one-degree-of-freedom Hamiltonian system be integrated with an error of size  $0.5/d$  (on average over the components). Thus, as  $d$  increases the numerical integration has to become

---

<sup>2</sup>This and similar estimates to be found later require some smoothness in the Hamiltonian but we will not concern ourselves with such issues.

more accurate.

The leapfrog integrator has two geometric properties (Blanes and Casas, 2016; Sanz-Serna and Calvo, 2018): (1) it preserves the volume element,  $d\theta(\tau)dp(\tau) = d\theta(0)dp(0)$  and (2) it is time-reversible, i.e., if after  $L$  time-steps the integrator maps an initial condition  $(\theta, p)$  into  $(\theta^*, p^*)$ , then  $L$  steps from the initial condition  $(\theta^*, -p^*)$  lead to  $(\theta, -p)$ . These two properties are essential for the accept/reject mechanism based on the expression (5) to remove the bias introduced by numerical integration. For integrators that fail to be either volume-preserving or time-reversible, the expression of the acceptance probability is far more complicated (see e.g. Fang et al. (2014)); it is for this reason that explicit Runge-Kutta and other standard methods are not suitable to be used in HMC codes.

To conclude this section, we look at the model problem with univariate target  $\pi(\theta) = \mathcal{N}(\theta|0, \sigma^2)$ . If, in addition, the mass matrix is taken to be  $M = 1$ , then  $\pi(p) = \mathcal{N}(p|0, 1)$ . For this problem, it is proved by Blanes et al. (2014) that the expectation of the difference (6), for  $0 < \epsilon < 2\sigma$  and an arbitrary number  $L$  of time-steps, satisfies

$$0 \leq \mathbb{E}(\Delta H) \leq \frac{\epsilon^4/\sigma^4}{32(1 - \frac{\epsilon^2}{4\sigma^2})}, \quad (7)$$

if one assumes that the chain is at stationarity. With  $\epsilon = \sigma$ , the upper bound of the energy error takes the value  $1/24$  and halving  $\epsilon$  brings down the expected error to  $\leq 1/480$ . This proves that, for this model problem, the leapfrog integrator works well with step-lengths that are not small (as measured in the “natural” unit  $\sigma$ ).

The bound (7) requires that  $\epsilon/\sigma \in (0, 2)$ . In fact for  $\epsilon/\sigma \geq 2$  the leapfrog algorithm is *unstable* for the corresponding Hamiltonian equations given by the (scalar) harmonic oscillator  $d\theta/d\tau = p$ ,  $dp/d\tau = -\theta/\sigma^2$  (the frequency of the oscillator is the inverse of the standard deviation  $\sigma$ ). Recall that this means that, if for fixed  $\epsilon \geq 2\sigma$  more and more time-steps are taken, then the numerical solution grows unboundedly. One says that  $(0, 2)$  is the *stability interval* of the leapfrog algorithm.

As pointed out before, as  $\epsilon \downarrow 0$  with  $L\epsilon = \tau_{end}$ , the energy increment (6) is  $\mathcal{O}(\epsilon^2)$  at each fixed  $(\theta, p)$ ; since, according to (7), its expectation at stationarity is  $\mathcal{O}(\epsilon^4)$ , there is much cancellation between points  $(\theta, p)$  where (6) is positive and points where it is negative. Such a cancellation is a consequence of conservation of volume and time-reversibility in leapfrog integration and holds for arbitrary targets  $\pi(\theta)$  in  $\mathbb{R}^d$  (Beskos et al., 2013; Bou-Rabee and Sanz-Serna, 2018). In fact, for general volume-preserving,

time-reversible integrators of order of accuracy  $\nu$  and arbitrary targets, the expectation of  $\Delta H$  at stationarity is  $\mathcal{O}(\epsilon^{2\nu})$ .

### 3 A family of integrators

In this paper we work with a family of splitting integrators (Blanes et al., 2008), depending on two real parameters  $a$  and  $b$ , that have potential for replacing the standard leapfrog method within HMC algorithms. The use of integrators of this family in different applications has been studied in detail by Campos and Sanz-Serna (2017). The formulas for performing a single time-step are:

$$p \leftarrow p + (1/2 - b)\epsilon \nabla_{\theta} \mathcal{L}(\theta), \quad (8)$$

$$\theta \leftarrow \theta + a\epsilon M^{-1}p, \quad (9)$$

$$p \leftarrow p + b\epsilon \nabla_{\theta} \mathcal{L}(\theta), \quad (10)$$

$$\theta \leftarrow \theta + (1 - 2a)\epsilon M^{-1}p, \quad (11)$$

$$p \leftarrow p + b\epsilon \nabla_{\theta} \mathcal{L}(\theta), \quad (12)$$

$$\theta \leftarrow \theta + a\epsilon M^{-1}p, \quad (13)$$

$$p \leftarrow p + (1/2 - b)\epsilon \nabla_{\theta} \mathcal{L}(\theta). \quad (14)$$

We assume that  $a$  and  $b$  do not take the values 0 or 1/2; when they do some of the substeps are redundant. An integration based on this algorithm is just a sequence of kicks and drifts and therefore not very different from a standard leapfrog integration. Four evaluations of the gradient are required at a single time-step (8)–(14). However, since the last gradient at a time-step may be reused at the next time-step, an integration leg of  $L$  time-steps needs  $3L + 1$  gradient evaluations. Therefore, over a single time-step (8)–(14) is *three* times more expensive than the leapfrog (2)–(4). *To have fair comparisons, an HMC simulation based on the standard leapfrog (2)–(4) should be allowed three times as many time-steps as a simulation based on (8)–(14).* We will return to this issue later.

Since each individual kick and drift preserves volume, the integrator (8)–(14) is volume-preserving. In addition, its palindromic structure leads to its being time-reversible. As a consequence, the bias introduced by numerical integration errors may be suppressed by an accept/reject mechanism based on the same formula (5) that is used for leapfrog integrations. In summary, the implementation of HMC based on the new integrator is extremely similar to leapfrog-based implementations.

Campos and Sanz-Serna (2017) show that when the parameters do not satisfy the relation

$$a + b - 6ab = 0, \tag{15}$$

the stability interval of (8)–(14) is relatively short and as a consequence small values of  $\epsilon$  are required to cover the interval  $[0, \tau_{end}]$ . Therefore schemes for which (15) is violated cannot possibly compete with standard leapfrog. Accordingly, we hereafter assume that the relation holds and in this way we are left with a family that may be described in terms of the single parameter  $b$ . It is perhaps of some interest to point out that by imposing (15) we exclude the unique choice of  $a$  and  $b$  for which (8)–(14) is fourth-order accurate. That fourth-order accurate integrator has a very short stability interval and performs extremely poorly in the HMC context, as shown by Campos and Sanz-Serna (2017). All integrators considered below are then *second-order accurate*, just as standard leapfrog.

It is virtually impossible to get analytically accurate information on the energy increment (6) for given Hamiltonian system, integrator and value of  $\epsilon$ , in particular in our situation where we are interested in relatively *large* values of  $\epsilon$ ; therefore numerical experimentation seems to be essential to identify good values of  $b$ . The numerical experiments in the next section compare the following six values of  $b$ :

- $b = 1/3$ . In this case, it is easy to check that a single time-step of (8)–(14) yields the same result as three consecutive time-steps of standard leapfrog (2)–(4) with step-length  $\epsilon/3$ . Thus, in the experiments below, one may think that, rather than performing  $L$  time-steps with (8)–(14), each of length  $\epsilon$ , one is taking  $3L$  time-steps with length  $\epsilon/3$  with the standard leapfrog integrator. Hence, when comparing later an integration with a value  $b \neq 1/3$  with a second integration with  $b = 1/3$ , we are really comparing the results of the first integration with those standard leapfrog would deliver if it were run with a time-step  $\epsilon/3$ , so as to equalize computational costs. With  $b = 1/3$  the length of the stability interval of (8)–(14) is  $\eta = 6$  (of course three times the corresponding length for (2)–(4)). In what follows we shall refer to the method with  $b = 1/3$  as LF.
- $b = 0.35$  with stability interval of length  $\eta \approx 4.969$ .
- $b = 0.38111989033452$ . The rationale for this choice, suggested by Blanes et al. (2014) in the numerical analysis literature, will be summarized below. This method is called BlCaSa in Campos and Sanz-Serna (2017). The length of the stability interval is  $\eta \approx 4.662$ .

- $b = 0.391008574596575$ . This choice, due to Predescu et al. (2012) in the molecular dynamics literature without reference to HMC, results from imposing that, for the particular case where  $\nabla\mathcal{L}(\theta)$  is linear (so that  $\pi(\theta)$  is a Gaussian), the error in  $H(\theta, p)$  after a single time-step is  $\mathcal{O}(\epsilon^5)$  as  $\epsilon \downarrow 0$  (rather than  $\mathcal{O}(\epsilon^3)$  as it is the case for other values of  $b$ ). As a result, the energy increment (6) over an integration leg is  $\mathcal{O}(\epsilon^4)$  for Gaussian targets, which, according to the results in Beskos et al. (2013), implies an  $\mathcal{O}(\epsilon^8)$  bound for the expected energy error for such targets. This method is called PrEtAl in Campos and Sanz-Serna (2017) and has a stability interval of length  $\eta \approx 4.584$ .
- $b = 0.40$ , with  $\eta \approx 4.519$ .
- $b = 0.45$ , with  $\eta \approx 4.224$ .

Experiments not reported here show that values of  $b$  smaller than  $1/3$  or larger than  $0.45$  lead to integrators that are not competitive in the HMC context with the six listed above.

As it is the case for PrEtAl, the method BiCaSa was derived assuming a multivariate Gaussian model. However, while PrEtAl was derived so as to reduce the energy error in the limit  $\epsilon \downarrow 0$ , the derivation of BiCaSa takes into account the behaviour of the method over a finite  $\epsilon$ -interval in view of the fact that Hamiltonian simulations within HMC are not carried out with small values of the step-length. It was proved by Blanes et al. (2014) that for each integrator, assuming  $M = I$ ,<sup>3</sup> the expected energy increment at stationarity has, independently of the number  $L$  of time-steps at each integration leg, a bound

$$0 \leq \mathbb{E}(\Delta H) \leq \sum_{i=1}^d \rho(\epsilon/\sigma_i), \quad (16)$$

where  $d$  is the dimension of  $\theta$ , the  $\sigma_i$  are the marginal standard deviations (i.e. the square roots of the eigenvalues of the covariance matrix) of  $\theta$  and  $\rho$  is a function that changes with the integrator. For the standard leapfrog integrator and  $d = 1$ , we have seen this bound in (7). Blanes et al. (2014) determined  $b$  by minimizing

$$\rho_\infty = \max_{0 < \zeta < 3} \rho(\zeta);$$

---

<sup>3</sup>The result may be adapted to cover other choices of mass matrix (Blanes et al., 2014; Bou-Rabee and Sanz-Serna, 2018).

this optimizes the performance assuming that  $\epsilon$  is chosen  $\leq 3 \min_i \sigma_i$ , see Blanes et al. (2014) for further details.<sup>4</sup>

## 4 Numerical experiments

Appropriate choices of  $M$  may allow for proposal mechanisms tailored to the given target and to its possibly anisotropic covariance structure. From the point of view of the Hamiltonian dynamics,  $M$  may be chosen so as to reduce the spectrum of frequencies present in the system of differential equations, which facilitates the numerical integration, see e.g. Beskos et al. (2011). For simplicity we do not explore here special choices of  $M$  and all numerical experiments use the unit mass matrix  $M = I_d$ . In all of them we have randomized, following Neal (2011), the step-length at the beginning of each step of the Markov chain/integration leg. Once the “basic” step-length  $\epsilon$  and the number of time-steps  $L$  to complete a simulation of duration  $\tau_{end} = L\epsilon$  have been chosen, the integration to find the  $n$ -th proposal of the chain is performed with a step-length

$$\epsilon^{(n)} = (1 + u^{(n)})\epsilon, \quad u^{(n)} \sim \mathcal{U}(-0.05, 0.05).$$

As a result, the actual length  $L\epsilon^{(n)}$  of the integration interval will differ slightly from  $\tau_{end}$ . This randomization suppresses the artifacts that may appear for special choices of the step-length (Bou-Rabee and Sanz-Serna, 2017).

### 4.1 A Gaussian example

We have first considered a simple test problem, used by Blanes et al. (2014), with target density

$$\pi(\theta) \propto \exp\left(-\frac{1}{2} \sum_{j=1}^d j^2 \theta_j^2\right), \quad (17)$$

for  $d = 256$  or  $d = 1024$ . Assuming a diagonal covariance matrix is *not restrictive*, as the application of the numerical integrators commutes with the rotation of axes that diagonalizes the covariance matrix (see Bou-Rabee and Sanz-Serna (2018) for a detailed discussion). We report experiments

---

<sup>4</sup>For the standard leapfrog integrator the step-length  $\epsilon = \min_i \sigma_i$  is 50% smaller than the maximum allowed by stability (compare with Neal (1993, 2012)); the interval in the minimization here is three times as long because the integrators should be operated with step-lengths three times longer than (2)–(4).

with  $\tau_{end} = 5$ ; we also ran other values of  $\tau_{end}$  and the conclusions as to the merit of the various integrators do not differ from those drawn below.

As in this example the Hamiltonian system to be integrated is a set of  $d$  uncoupled linear oscillators with angular frequencies  $1, \dots, d$ , the condition for a stable integration is

$$d \times \epsilon \leq \eta,$$

where  $\eta$  is the length of the stability interval of the method. For all six integrators tested,  $\eta \geq 4$  and therefore for the step-length  $\epsilon_0 = 4/d$  all of them are stable. For the case  $d = 256$ , and  $\tau_{end} = 5$ ,  $\epsilon_0 = 4/d$ , we need  $L = \tau_{end}/\epsilon_0 = 320$  time-steps in each integration leg. Our experiments had  $L = 320, 360, 400, \dots, 960$  and  $\epsilon = 5/320, 5/360, \dots, 5/960$ , respectively; ( $\epsilon$  then varies between  $\approx 1.6 \times 10^{-2}$  and  $\approx 5.2 \times 10^{-3}$ ). For the methods with longer stability intervals, such as LF, we also used the coarser values  $\epsilon = 5/280, 5/240, 5/200$ , with (respectively)  $L = 280, 240, 200$ . Turning now to  $d = 1024$ , the choices  $\tau_{end} = 5$ ,  $\epsilon_0 = 4/d$ , require  $L = \tau_{end}/\epsilon_0 = 1280$  time-steps in each integration leg. Our experiments used  $L = 1280, 1440, 1600, \dots, 3840$ ,  $\epsilon = 5/1280, \dots, 5/3840$  ( $\epsilon$  ranges from  $\approx 3.5 \times 10^{-3}$  to  $\approx 1.3 \times 10^{-3}$ ), and in addition  $L = 1120, 960, 900$ ,  $\epsilon = 5/1120, 5/960, 5/900$  for the integrators with longer stability intervals.

For each  $b$  and  $\epsilon$ , we have generated a Markov chain with  $N = 5000$  samples, a number that is large enough for our purposes. The initial  $\theta^{(0)}$  is drawn from the target  $\pi(\theta)$ , and we have monitored the acceptance rate, the effective sample size (ESS) of the first component  $\theta_1$  of  $\theta$  (the component with largest variance/smallest precision), and the mean, across the 5000 integration legs/Markov chain steps, of the energy increments (6).

For the different methods we have plotted in Figures 1 ( $d = 256$ ) and 2 ( $d = 1024$ ), as functions of the step-size  $\epsilon$ , the mean of the changes in  $H$  (diamonds), the acceptance percentage (triangles), and the ESS as a percentage of the number  $N$  of samples in the Markov chain (squares). Note that we use a logarithmic scale for the energy increments (left vertical axis) and a linear scale for acceptance rate and ESS (right vertical axis).

We look first at the mean of  $\Delta H$  as a function of  $\epsilon$  for the different integrators. When comparing Figure 1 with Figure 2, note that both the vertical and horizontal scales are different;  $d = 1024$  uses smaller values of  $\epsilon$  and yet shows larger values of the mean energy error.

The differences between BICaSa, PrEtAl and  $b = 0.4$  are remarkable in both figures if we take into account that these three integrators have values of  $b$  very close to each other.<sup>5</sup> PrEtAl was designed by making  $\Delta H$  small

---

<sup>5</sup>When running BICaSa and PrEtAl it is important not to round the values of  $b$  given

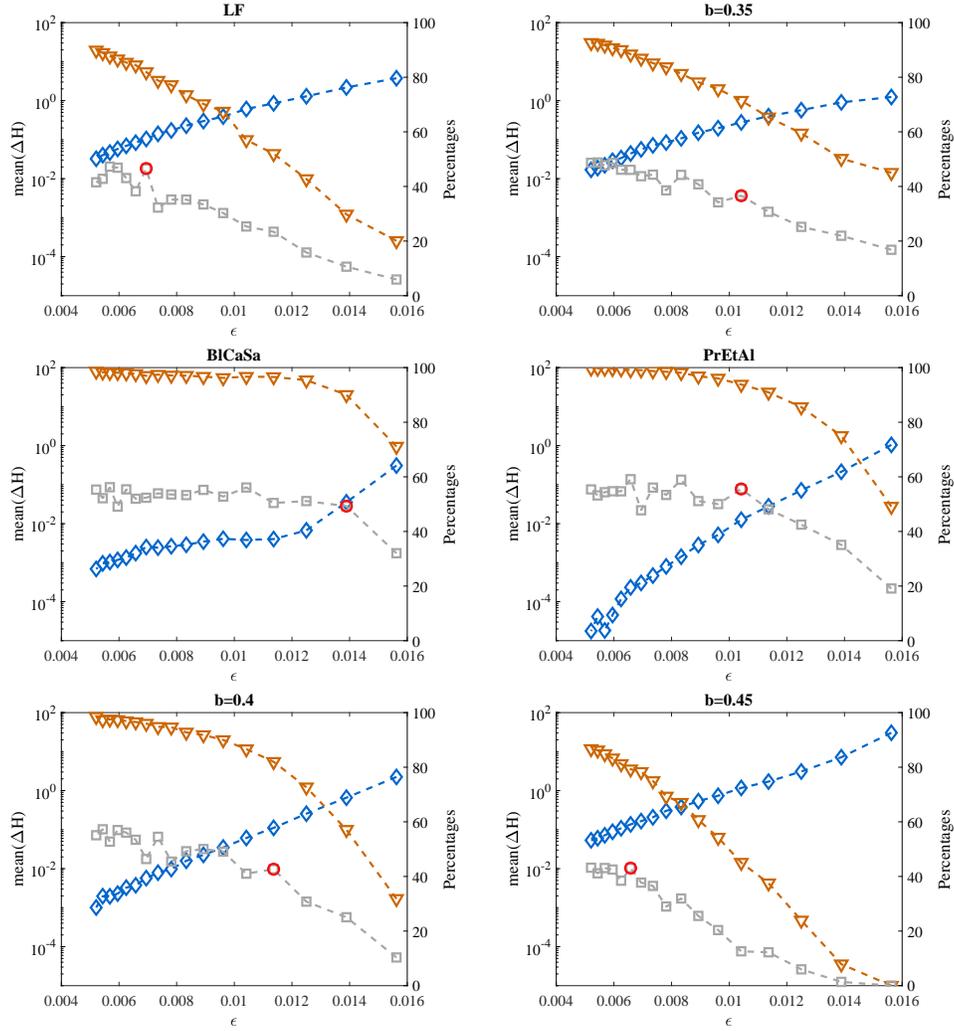


Figure 1: Gaussian model,  $d = 256$ . ESS percentage for  $\theta_1$  (squares), acceptance percentage (triangles) and the mean of  $\Delta H$  (diamonds) as functions of the step-length  $\epsilon$ , for different values of the integrator parameter  $b$ , when  $\tau_{end} = 5$ . The red circle highlights for each method the most efficient run as measured by the highest ESS per time-step.

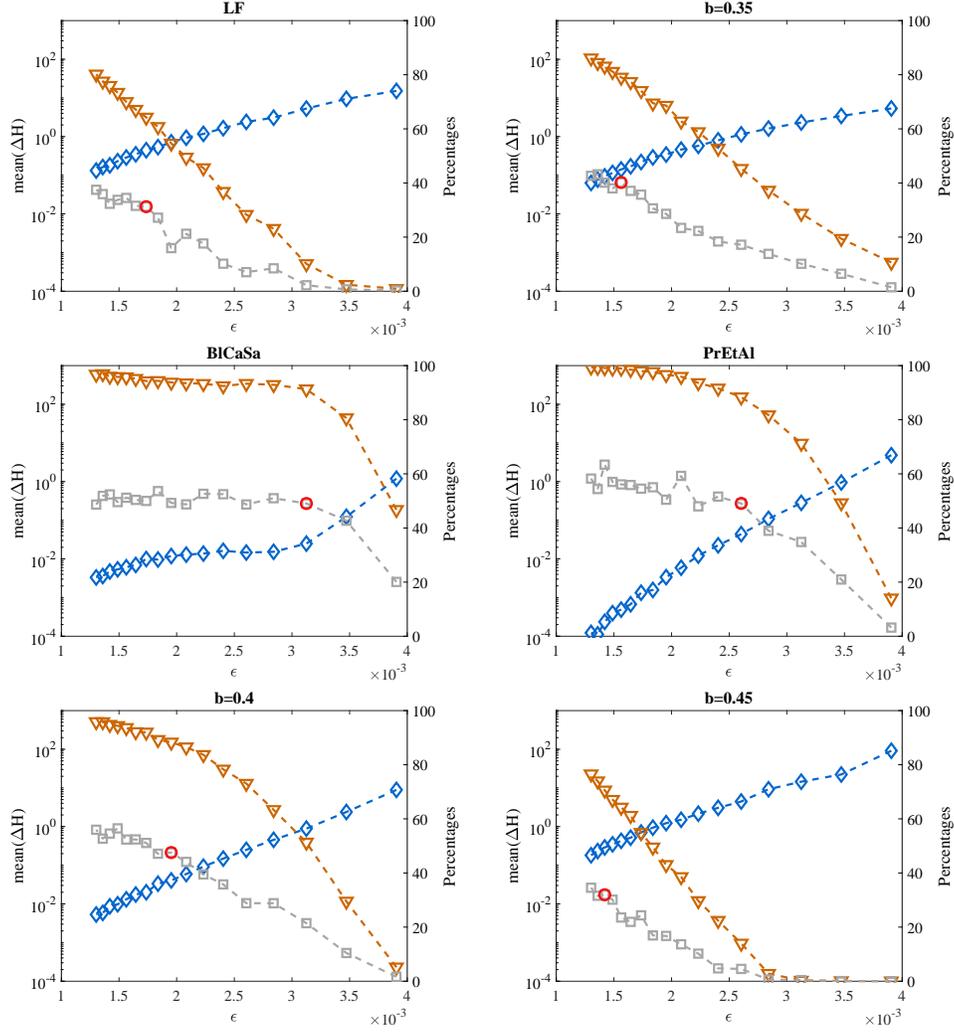


Figure 2: Model Gaussian,  $d = 1024$ . ESS percentage for  $\theta_1$  (squares), acceptance percentage (triangles) and the mean of  $\Delta H$  (diamonds) as functions of the step-length  $\epsilon$ , for different values of the integrator parameter  $b$ , when  $\tau_{end} = 5$ . The red circle highlights for each method the most efficient run as measured by the highest ESS per time-step.

in the limit  $\epsilon \downarrow 0$  for Gaussian targets and we see that indeed the graph of  $\Delta H$  as a function of  $\epsilon$  has a larger slope for this method than for the other five integrators. Even though, for  $\epsilon$  small, PrEtAl delivers very small energy errors, this does not seem to be particularly advantageous in HMC sampling, because those errors happen where the acceptance rate is close to 100%, i.e. not in the range of values of  $\epsilon$  that make sense on efficiency grounds. On the other hand, hand BICaSa was derived so as to get, for Gaussian targets, small energy errors over a range of values of  $\epsilon$ ; this is borne out in both figures where for this method the variation of  $\Delta H$  with  $\epsilon$  shows a kind of plateau, not present in any of the other five schemes.

In both figures, for each  $\epsilon$ , the values of the mean of  $\Delta H$  for each of the integrators  $b = 0.35$ , BICaSa, PrEtAl and  $b = 0.4$  are all below the corresponding value for  $b = 1/3$ . Since, for a given  $\epsilon$ , all integrators have the same computational effort, this implies that each of those four methods improves on LF as it gives less error for a given computational cost.

The mean acceptance as a function of  $\epsilon$  has a behaviour that mirrors that of the mean of  $\Delta H$ : smaller mean energy errors lead to larger mean acceptance rates as expected. For  $\epsilon$  small, PrEtAl has the highest acceptance rate, but again this is not advantageous for our sampling purposes. In Figure 1, BICaSa shows acceptance rates above  $\approx 70\%$  for the coarsest  $\epsilon \approx 1.6 \times 10^{-2}$ ; to get that level of acceptance with LF one needs to reduce  $\epsilon$  to  $\approx 1 \times 10^{-2}$ . The advantage of BICaSa over LF in that respect is even more marked in Figure 2.

To visualize the relation between the mean  $\Delta H$  and the mean acceptance rate, we have plotted in the right panel of Figure 3 the values corresponding to all our runs for  $d = 1024$  (i.e. to all six integrators and all values of  $\epsilon$ ). All points are almost exactly on a curve, so that to a given value of the mean energy error corresponds a well-defined mean acceptance rate, independently of the value of  $b$  and  $\epsilon$ . The left panel collects similarly all the runs with  $d = 256$ ; the points are again very approximately on the same curve as the points in the right panel, but the fit is now worse. This behaviour is related to a central limit theorem and will be discussed in Section 5. Figure 4 displays results corresponding to the six integrators,  $\tau_{end} = 5$ , and different step-lengths for the target (17) for the small values  $d = 2, 4$  where the central limit theorem cannot be invoked; as distinct from Figure 3 it does not appear to exist a single smooth curve that contains the results of the different integrators.

Returning to Figure 1 and Figure 2, we see that in both of them and 

---

 here and to compute accurately the value of  $a$  from (15).

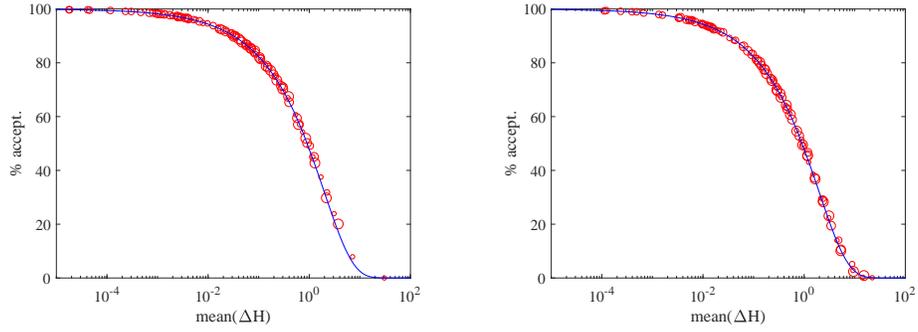


Figure 3: Gaussian model,  $d = 256$  (left),  $d = 1024$  (right). Acceptance percentage vs. mean energy error. Points come from all six integrators and all step-lengths. The diameter of the circles indicates the integrator used.

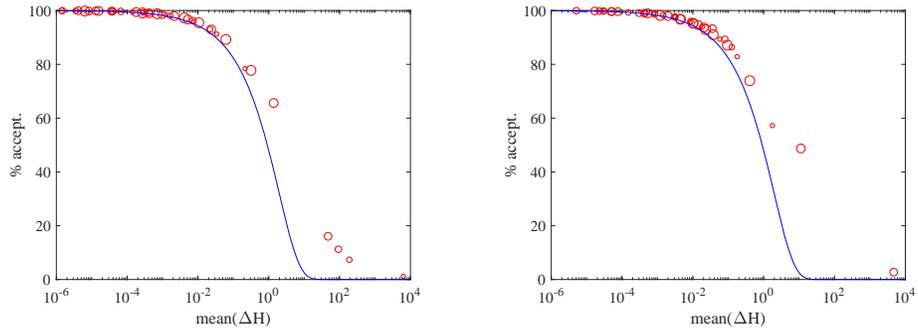


Figure 4: Gaussian model in low dimension,  $d = 2$  (left),  $d = 4$  (right). Acceptance percentage vs. mean energy error. Points come from all six integrators and different step-lengths. The diameter of the circles indicates the integrator used. Compare with Figure 3 and note that the horizontal scales are different.

for all six methods, the relative ESS is close to 50% when  $\epsilon$  is very small. This is no doubt due to the fact that, if we used exact integration of the Hamiltonian dynamics, we would see a relative ESS around 50%. We therefore should prefer integrators that reach that limiting value of  $\approx 50\%$  with the least computational work, i.e. with the largest  $\epsilon$ . In this respect the advantage of BiCaSa over LF is very noticeable. In each panel of both figures, we have marked with a circle the run with the highest ESS *per unit computational work*, i.e. the run where  $\text{ESS} \times \epsilon$  is highest. The location of the circles reinforces the idea that BiCaSa is the method that may operate with the largest  $\epsilon$ , i.e. with the least computational effort.

In order to better compare the efficiency of the different integration algorithms, we provide two additional figures, Figures 5 and 6, for  $d = 256$  and  $d = 1024$  respectively. For clarity, only the methods LF, BiCaSa and PrEtAl are considered. The horizontal axis in both panels of each figure represents values of  $\epsilon$ . In the left panel we show values of the acceptance percentage *multiplied* by  $\epsilon$  (so as to take into account that the computational effort is proportional to  $1/\epsilon$ ). The right panel shows values of the product of ESS and  $\epsilon$ .

For all step-lengths, the acceptance percentage per time-step, i.e. per unit computational work, is clearly smaller for LF than for the other two methods. Comparing PrEtAl and BiCaSa, we see that for the largest step-lengths (larger than  $8.0 \times 10^{-3}$  when  $d = 256$  and larger than  $2.5 \times 10^{-3}$  when  $d = 1024$ ), BiCaSa shows larger acceptance rates than PrEtAl, but when the step-length decreases the acceptance rates using PrEtAl are slightly higher. BiCaSa is the most efficient of the three methods, as can be seen by comparing the maxima of the three lines in each panel of both figures.

In the right panels a blue circle highlights the most efficient run for each method, which when  $d = 256$  corresponds to 360 time-steps per leg for BiCaSa (ESS = 2463, acceptance rate (AR) = 90.04%), 480 time-steps per leg for method PrEtAl (ESS = 2777, AR = 93.82%), and 720 time-steps per leg for LF (ESS = 2328, AR = 81.92%). When  $d = 1024$ , the most efficient runs are obtained when  $\epsilon = \tau_{end}/1600$  for method BiCaSa (ESS = 2452, AR = 91.30%), when  $\epsilon = \tau_{end}/1920$  for method PrEtAl (ESS = 2454, AR = 88.36%), and when  $\epsilon = \tau_{end}/2880$  for the LF integrator (ESS = 1562, AR = 64.24%). Note that the optimal acceptance rate is typically higher than the 65% resulting from the analysis of Beskos et al. (2013). There is no contradiction: while the analysis applies when  $d \uparrow \infty$  and  $\pi(\theta)$  is a product of identically distributed independent distributions, our figures have a fixed value of  $d$  and here  $\pi(\theta)$  is not of that simple form.

We conclude that in this model problem, for a given computational effort,

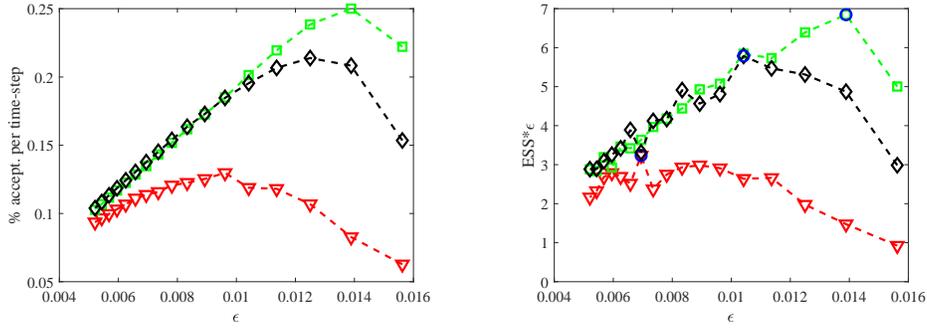


Figure 5: Gaussian model,  $d = 256$ . Comparison of the time integrators BICaSa (green squares), PrEtAl (black diamonds) and LF (red triangles). Acceptance percentage per time-step/per unit computational effort (left) and ESS per time-step/per unit computational effort (right) against the step-length  $\epsilon$ .

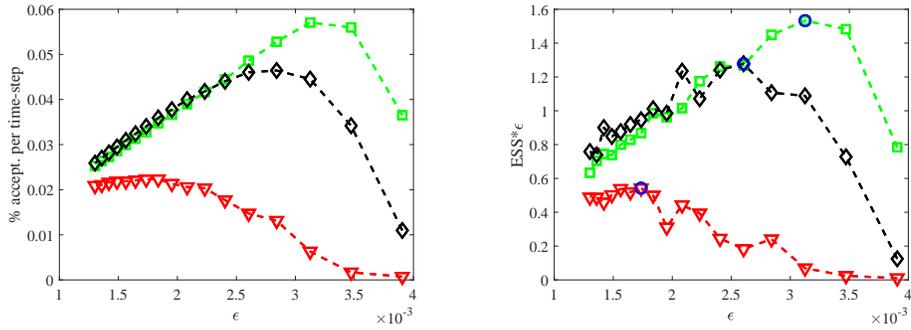


Figure 6: Gaussian model,  $d = 1024$ . Comparison of the time integrators BICaSa (green squares), PrEtAl (black diamonds) and LF (red triangles). Acceptance percentage per time-step/per unit computational effort (left) and ESS per time-step/per unit computational effort (right) against the step-length  $\epsilon$ .

BlCaSa generates roughly three times as many effective samples than LF. In addition the advantage of BlCaSa is more marked as the problem becomes more challenging.

The information obtained above on the performance of the different integrators when using ESS is essentially the same as when using the mean acceptance rate. Since ESS is a metric that depends on the specific observable chosen, the experiments below concentrate on the mean acceptance rate.

## 4.2 Log-Gaussian Cox model

The second test problem (Christensen et al., 2005; Girolami and Calderhead, 2011) is a well-known high dimensional example of inference in a log-Gaussian Cox process. A  $64 \times 64$  grid is considered in  $[0, 1] \times [0, 1]$  and the random variables  $\mathbf{X} = X_{i,j}$ ,  $1 \leq i, j \leq 64$ , represent the number of points in cell  $(i, j)$ . These random variables are conditionally independent and Poisson distributed with means  $m\Lambda_{i,j}$ , where  $m = 1/4096$  is the area of each cell and  $\Lambda_{i,j}$ ,  $1 \leq i, j \leq 64$ , are unobserved intensities, which are assumed to be given by  $\Lambda_{i,j} = \exp(Y_{i,j})$ , with  $\mathbf{Y} = Y_{i,j}$ ,  $1 \leq i, j \leq 64$ , multivariate Gaussian with mean  $\mu$  and covariance matrix  $\Sigma$  given by

$$\Sigma_{(i,j),(i',j')} = \sigma^2 \exp\left(-\frac{\sqrt{(i-i')^2 + (j-j')^2}}{64\beta}\right).$$

Given the data  $\mathbf{X} = \mathbf{x}$ , the target density of interest is

$$\pi(\mathbf{y}) \propto \prod_{i,j=1}^{64} \exp(x_{i,j}y_{i,j} - m \exp(y_{i,j})) \exp(-(\mathbf{y} - \mu\mathbf{1})^T \Sigma^{-1}(\mathbf{y} - \mu\mathbf{1})/2).$$

As in Christensen et al. (2005) we have fixed the parameters  $\beta = 1/33$ ,  $\sigma^2 = 1.91$  and  $\mu = \log 126 - \sigma^2/2 \approx 3.881$  and used three alternative starting values. The conclusions obtained as to the merits of the different integrators are very similar for the three alternative cases and, for this reason, we only report here results corresponding to one of them (choice (b) in Christensen et al. (2005)). The initial state is randomly constructed for each run in the following way. Given a random vector  $\Gamma$  with components  $\mathcal{N}(0, 1)$  the starting value  $\mathbf{y}$  is the solution of the nonlinear system  $\mathbf{y} = \mu\mathbf{1} + L\Gamma$ , where  $L$  is the Cholesky factor of the matrix  $(\Sigma^{-1} + \text{diag}(\mathbf{y}))^{-1}$  which, therefore, depends on  $\mathbf{y}$ . This nonlinear system has been solved by fixed-point iteration starting with  $\mathbf{y}^{(0)} = \mu\mathbf{1}$ , and the iteration has been stopped when  $\|\mathbf{y}^{(n+1)} - \mathbf{y}^{(n)}\| < 10^{-12}$  (which happens after 19 iterations).

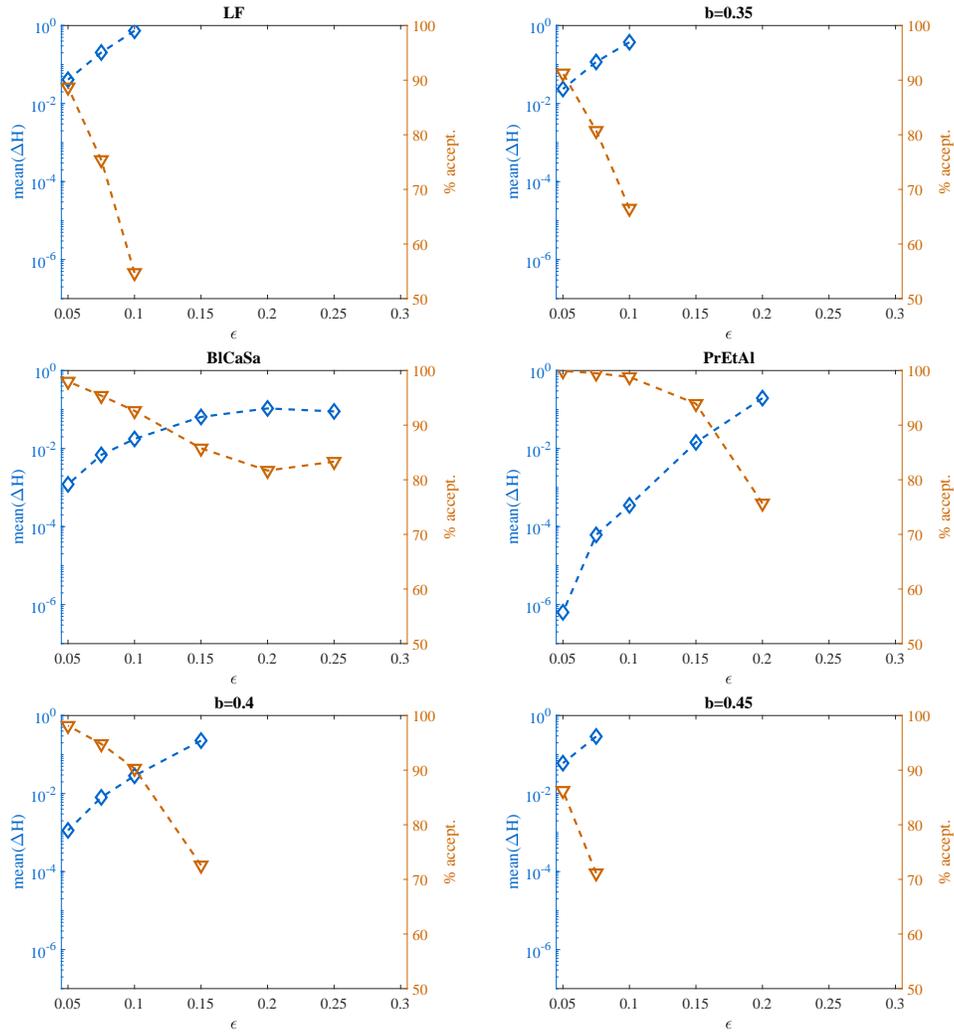


Figure 7: Log-Gaussian Cox problem. Acceptance percentage (triangles) and the mean of  $\Delta H$  (diamonds) as functions of the step-length  $\epsilon$ , for different values of parameter  $b$ , when  $\tau_{end} = 3$ .

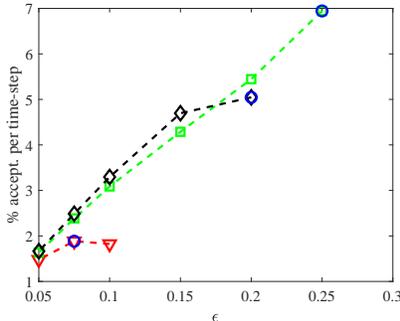


Figure 8: Log-Gaussian Cox problem. Acceptance percentage per time-step as a function of  $\epsilon$ , for PrEtAl (diamonds), BiCaSa (squares) and LF (triangles), when  $\tau_{end} = 3$ .

The length of each integration leg has been set equal to 3, and the “basic” step-lengths used are  $\epsilon = 0.3, 0.25, 0.20, 0.15, 0.1, 0.075, 0.05$ . After generating 1000 burn-in samples, we have run Markov chains with 5000 elements, a number that is sufficient for our purposes.

In Figure 7 we have represented the acceptance percentage and the mean of the increments in energy against the step-size  $\epsilon$  for the different values of the parameter  $b$ . We have omitted data corresponding to either acceptance percentages below 45 or values of the mean of  $\Delta H$  larger than 1.

Comparing the performance of the six methods, we see that for this problem, BiCaSa, in contrast with the other five integrators, operates well for  $\epsilon = 0.25$ . PrEtAl gives very high acceptance rates for small values of  $\epsilon$ , but this is not very relevant for the reason given in the preceding test example. The other four integrators, including LF, require very small step-lengths to generate a reasonable number of accepted proposals; their performance is rather poor.

The energy error in PrEtAl decays by more than five orders of magnitude when reducing  $\epsilon$  from 0.2 to 0.05; this is remarkable because the additional order of energy accuracy of this integrator as  $\epsilon \downarrow 0$ , as compared with the other five, only holds analytically for *Gaussian targets* (Predescu et al., 2012). Similarly BiCaSa shows here the same kind of plateau in the mean energy error we observed in the Gaussian model.

In Figure 8 we have plotted the acceptance percentage per step as a function of the step-length  $\epsilon$  for the two methods with the largest acceptance rates, PrEtAl (diamonds) and BiCaSa (squares) in addition to LF

(triangles). The blue circles in the figure highlight the most efficient run for each method. When the computational cost is taken into account, BICaSa is the most efficient of the methods being compared, by a factor of more than three with respect to LF.

### 4.3 Alkane molecule

The third test problem comes from molecular simulations and was chosen to consider multimodal targets very far away from Gaussian models. The target is the Boltzmann distribution  $\pi(\theta) \propto \exp(-V(\theta))$ , where the vector  $\theta$  collects the cartesian coordinates of the different atoms in the molecule and  $V(\theta)$  is the potential energy and, accordingly, in (8)–(14),  $\mathcal{L}(\theta)$  has to be replaced by  $-V(\theta)$ . We have considered the simulation of an alkane molecule described in detail in Cances et al. (2007) and later used by Campos and Sanz-Serna (2015), Campos and Sanz-Serna (2017) and Blanes et al. (2014). The expression for the potential energy is extremely complicated, as it contains contributions related to bond lengths, bond angles, dihedral angles, electrostatic forces and Van der Waals forces. In our numerical experiments molecules with 5, 9 or 15 carbon atoms have been used, leading to Hamiltonian systems with 15, 27 and 45 degrees of freedom (the hydrogen atoms are not simulated).

As Campos and Sanz-Serna (2017), we set the length of each integration leg as  $\tau_{end} = 0.48$ , with “basic” step-length  $\epsilon = 0.08, 0.06, 0.048, 0.04, 0.03$ , corresponding to 6, 8, 10, 12 and 16 steps per integration leg. In all runs, after generating 200 samples for burn-in, we consider Markov chains with 1000 samples; the results provided are obtained by averaging over 20 instances of the experiments. In all cases, the initial condition  $\theta^{(0)}$  for the chain is the equilibrium with minimum potential energy.

Figure 9 displays results for 15 carbon atoms. We have omitted data corresponding to either acceptance percentages below 45 or values of the mean of  $\Delta H$  larger than 1. It is clear that LF is far from being the best choice for this problem. BICaSa and PrEtAl have substantially larger acceptance rates than LF for each value of  $\epsilon$  where LF works. The results for 5 and 9 carbon atoms are not very different. With 5 atoms, LF,  $b = 0.35$ , BICaSa, and PrEtAl can operate with the largest value  $\epsilon = 0.08$ , but again, for all values of  $\epsilon$ , the acceptance rate of LF is well below that provided by BICaSa.

As in the preceding test problem, for both BICaSa and PrEtAl the variation of the mean  $\Delta H$  with  $\epsilon$  is as in the Gaussian model: the former method exhibits a plateau and the second a fast decay (a reduction of more than two orders of magnitude when halving  $\epsilon$  from 0.06 to 0.03). This is in spite of

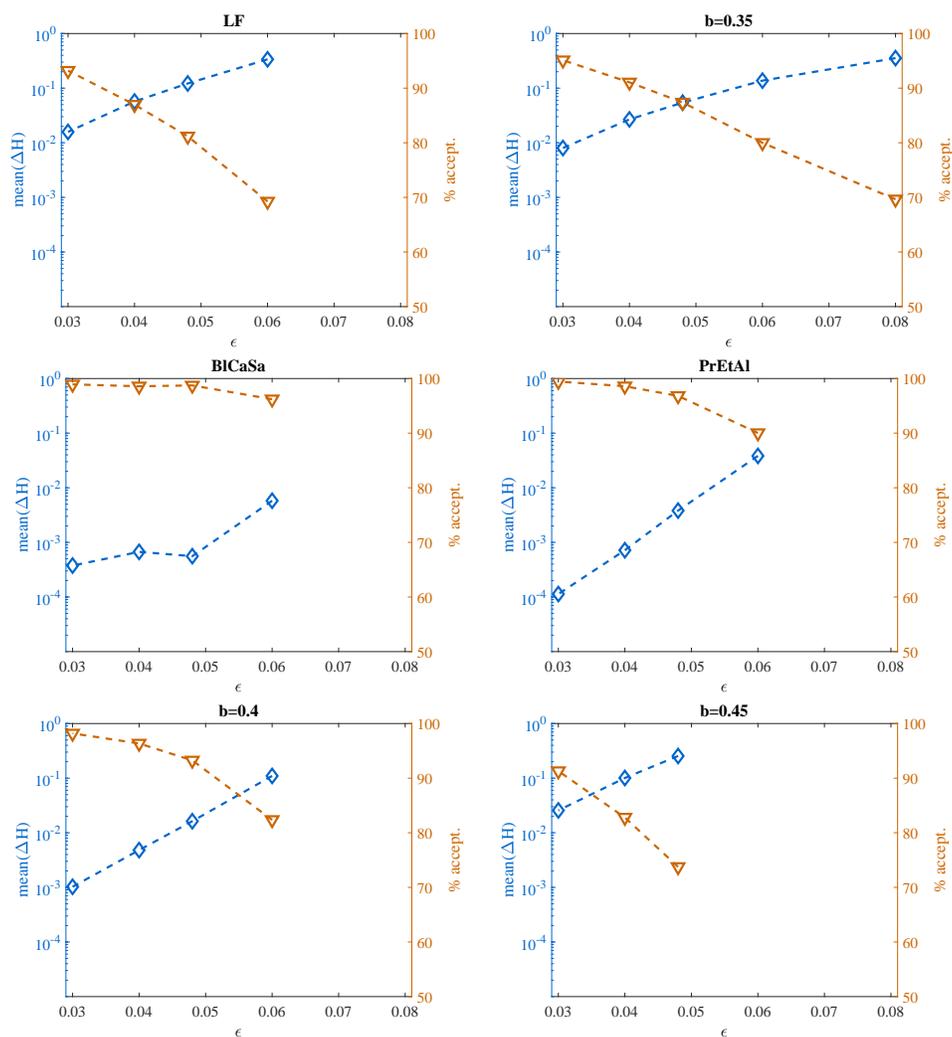


Figure 9: Alkane molecule with 15 carbon atoms. Acceptance percentage (triangles) and the mean of  $\Delta H$  (diamonds) as functions of the step-size  $\epsilon$ , for different values of parameter  $b$ , when  $\tau_{end} = 0.48$ .

the fact that both integrators were derived by considering only the Gaussian case.

## 5 Expected acceptance rate vs. expected energy error

In this section we investigate analytically some properties of the expected acceptance percentage and the expected energy increments.

### 5.1 A formula for the expected acceptance rate

The following result (Neal, 2011) is needed later.

**Lemma 1.** *Let the integration of the Hamiltonian dynamics (1) be carried out with a volume-preserving and time-reversible integrator, such that the transformation  $\Psi$  in the phase space that maps the initial condition  $(\theta, p)$  into the numerical solution  $(\theta^*, p^*)$  is continuously differentiable. Assume that  $(\theta, p) \sim \pi(\theta, p)$  (i.e. the chain is at stationarity) and that the energy increment (6) satisfies*

$$\mathbb{P}(\Delta H = 0) = 0. \tag{18}$$

*Then the expected acceptance rate is given by*

$$\mathbb{E}(a) = 2 \mathbb{P}(\Delta H < 0).$$

*Proof.* Since almost surely  $\Delta H \neq 0$ , we may write

$$\begin{aligned} \mathbb{E}(a) &= \int_{\mathbb{R}^d \times \mathbb{R}^d} \min(1, \exp(-\Delta H(\theta, p))) \exp(-H(\theta, p)) \, d\theta \, dp \\ &= \int_{\Delta H < 0} \exp(-H(\theta, p)) \, d\theta \, dp \\ &\quad + \int_{\Delta H > 0} \exp(-\Delta H(\theta, p)) \exp(-H(\theta, p)) \, d\theta \, dp, \end{aligned}$$

and it is enough to show that the last two integrals share a common value.

We denote by  $\widehat{\Psi}$  the composition  $S \circ \Psi$ , where  $S$  is the momentum flip  $S(\theta, p) = (\theta, -p)$  (so that  $\widehat{\Psi}(\theta, p) = (\theta^*, -p^*)$ ). Then, from the definition of  $\Delta H$ ,

$$\Delta H(\theta, p) + H(\theta, p) = H(\Psi(\theta, p))$$

and therefore

$$\begin{aligned}
\int_{\Delta H > 0} \exp(-\Delta H(\theta, p)) \exp(-H(\theta, p)) d\theta dp & \\
&= \int_{\Delta H > 0} \exp(-H(\Psi(\theta, p))) d\theta dp \\
&= \int_{\Delta H > 0} \exp(-H(\widehat{\Psi}(\theta, p))) d\theta dp,
\end{aligned}$$

where in the last equality we have taken into account that  $H \circ S = H$ . On the other hand,

$$\begin{aligned}
\int_{\Delta H < 0} \exp(-H(\theta, p)) d\theta dp &= \int_{\Delta H < 0} \exp(-H(\tilde{\theta}, \tilde{p})) d\tilde{\theta} d\tilde{p} \\
&= \int_{\Delta H > 0} \exp(-H(\widehat{\Psi}(\theta, p))) d\theta dp
\end{aligned}$$

Here the first equality is just a change in notation. The second uses the change of variables  $(\tilde{\theta}, \tilde{p}) = \widehat{\Psi}(\theta, p)$  which has unit Jacobian determinant because the integrator preserves volume. Pairs  $(\tilde{\theta}, \tilde{p})$  with  $\Delta H < 0$  correspond to pairs  $(\theta, p)$  with  $\Delta H > 0$  as a consequence of the time-reversibility of the integrator:

$$\begin{aligned}
\Delta H(\tilde{\theta}, \tilde{p}) &= H(\Psi(\tilde{\theta}, \tilde{p})) - H(\tilde{\theta}, \tilde{p}) = H(S(\theta, p)) - H(S(\Psi(\theta, p))) \\
&= H(\theta, p) - H(\Psi(\theta, p)) = -\Delta H(\theta, p). \quad \square
\end{aligned}$$

Later in the section we shall come across cases where, for simple targets and special choices of  $\epsilon$  and  $L$ , the assumption (18) does not hold. Those cases should be regarded as exceptional and without practical relevance. In addition, reviewing the proof, we see that the lemma could have been reformulated without (18): at stationarity and without counting the proposals with  $\Delta H(\theta, p) = 0$ , one out of two accepted steps comes from proposals with  $\Delta H(\theta, p) < 0$ .

## 5.2 The standard univariate Gaussian target

We now investigate in detail the model situation where  $\pi(\theta)$  is the standard univariate normal distribution and the mass matrix is  $M = 1$ , so that  $H = (1/2)p^2 + (1/2)\theta^2$ . The Hamiltonian system is the standard harmonic oscillator:  $d\theta/d\tau = p$ ,  $dp/d\tau = -\theta$ . For all time-reversible,

volume-preserving, one-step integrators of practical interest (including implicit Runge-Kutta and splitting integrators), assuming that  $\epsilon$  is in the stability interval, the transformation  $(\theta^*, p^*) = \Psi(\theta, p)$  associated with an integration leg has an expression (Blanes et al., 2014; Bou-Rabee and Sanz-Serna, 2018):

$$\theta^* = \cos(L\alpha_\epsilon)\theta + \chi_\epsilon \sin(L\alpha_\epsilon)p, \quad (19)$$

$$p^* = -\chi_\epsilon^{-1} \sin(L\alpha_\epsilon)\theta + \cos(L\alpha_\epsilon)p, \quad (20)$$

where  $\chi_\epsilon$  and  $\alpha_\epsilon$  are quantities that depend smoothly on the step-length  $\epsilon$  and change with the specific integrator, but are independent of the number  $L$  of time-steps. As  $L$  varies with fixed  $\epsilon$  and  $(\theta, p)$ , the end point  $(\theta^*, p^*)$  moves on an ellipse whose eccentricity is governed by  $\chi_\epsilon$ ; for a method of order of accuracy  $\nu$ ,  $\chi_\epsilon = 1 + \mathcal{O}(\epsilon^\nu)$  as  $\epsilon \downarrow 0$ . The angle  $\alpha_\epsilon$  is related to the speed of the rotation of the numerical solution around the origin as  $L$  increases; it behaves as  $\alpha_\epsilon = \epsilon(1 + \mathcal{O}(\epsilon^\nu))$ .

From (19)–(20) a short calculation reveals that the energy increment has the following expression

$$2\Delta H(\theta, p) = A\theta^2 + 2B\theta p + Cp^2, \quad (21)$$

with

$$\begin{aligned} A &= \sin^2(L\alpha_\epsilon)(\chi_\epsilon^{-2} - 1), \\ B &= \cos(L\alpha_\epsilon) \sin(L\alpha_\epsilon)(\chi_\epsilon - \chi_\epsilon^{-1}), \\ C &= \sin^2(L\alpha_\epsilon)(\chi_\epsilon^2 - 1). \end{aligned}$$

For future reference we note that

$$B^2 - AC = A + C. \quad (22)$$

Taking expectations at stationarity in (21),

$$2\mathbb{E}(\Delta H) = A + C, \quad (23)$$

i.e.

$$\mathbb{E}(\Delta H) = \sin^2(L\alpha_\epsilon)\rho(\epsilon)$$

with

$$\rho(\epsilon) = \frac{1}{2} (\chi_\epsilon^2 + \chi_\epsilon^{-2} - 2) = \frac{1}{2} \left( \chi_\epsilon - \frac{1}{\chi_\epsilon} \right)^2 \geq 0,$$

so that, regardless of the choice of  $L$ ,

$$0 \leq \mathbb{E}(\Delta H) \leq \rho(\epsilon). \quad (24)$$

As  $\epsilon \downarrow 0$ ,  $\rho(\epsilon) = \mathcal{O}(\epsilon^{2\nu})$ , and therefore  $\mathbb{E}(\Delta H) = \mathcal{O}(\epsilon^{2\nu})$ , where we remark that the exponent is doubled from the pointwise estimate  $\Delta H(\theta, p) = \mathcal{O}(\epsilon^\nu)$ , valid for  $\epsilon L = \tau_{end}$ . If  $L\alpha_\epsilon$  happens to be an integer multiple of  $\pi$  then  $A = B = C = 0$  and  $\Delta H \equiv 0$ ; in this case the transformation  $\Psi$  is either  $(\theta, p) \mapsto (\theta, p)$  or  $(\theta, p) \mapsto (-\theta, -p)$  with no energy error (but then the Markov chain is not ergodic, this is one of the reasons for randomizing  $\epsilon$ ). Also note that this gives an example where the hypothesis  $\mathbb{P}(\Delta H = 0) \neq 0$  in Lemma 1 does not hold.

The preceding material has been taken from Blanes et al. (2014). For the acceptance rate we have the following result that shows that  $\mathbb{E}(a)$  is of the form  $\varphi(\mathbb{E}(\Delta H))$  where  $\varphi$  is a monotonically decreasing function that does not depend on the integrator,  $\epsilon$  or  $L$ . Thus, even though the integrator BlCaSa was derived to minimize the expected energy error in the univariate standard normal target, this result shows that *it also maximizes the expected acceptance rate*.

**Theorem 1.** *When the target is the standard univariate Gaussian distribution, the mass matrix is set to 1 and the chain is at stationarity, the expected acceptance rate is given by*

$$\mathbb{E}(a) = 1 - \frac{2}{\pi} \arctan \sqrt{\frac{\mathbb{E}(\Delta H)}{2}},$$

*regardless of the (volume-preserving, time-reversible) integrator being used, the step-length  $\epsilon$  and the number  $L$  of time-steps in each integration leg.*

*Proof.* If in (21), either  $\sin(L\alpha_\epsilon) = 0$  or  $\chi_\epsilon = 1$ , then  $A = B = C = 0$ , which leads to  $\Delta H \equiv 0$ . In this case  $\mathbb{E}(a) = 1$  and  $\mathbb{E}(\Delta H) = 0$  and the result holds.

If  $\sin(L\alpha_\epsilon) \neq 0$  and  $\chi_\epsilon \neq 1$ , then  $A \neq 0$ ,  $B \neq 0$ ,  $C \neq 0$  and  $\Delta H$  vanishes on the straight lines

$$p = \frac{-B - \sqrt{B^2 - AC}}{C}\theta, \quad p = \frac{-B + \sqrt{B^2 - AC}}{C}\theta,$$

of the  $(p, \theta)$  plane with slopes  $m_1$  and  $m_2$  respectively. From the expressions for  $A$  and  $C$  we see that  $AC < 0$  and therefore  $m_1 m_2 = A/C < 0$ . For the sake of clarity, let us assume that  $C < 0$  and  $B > 0$ , which implies

$m_1 > 0$  and  $m_2 < 0$  (the proof in the other cases is similar). Choose angles  $\phi_1 \in (0, \pi/2)$ ,  $\phi_2 \in (\pi/2, \pi)$  with  $\tan(\phi_1) = m_1$ ,  $\tan(\phi_2) = m_2$ .

The energy increment  $\Delta H$  is then negative for points  $(\theta, p)$  with polar angles  $\phi \in (\phi_1, \phi_2) \cup (\phi_1 + \pi, \phi_2 + \pi)$ . Lemma 1 and the symmetry  $H(\theta, p) = H(-\theta, -p)$  yield

$$\begin{aligned} \mathbb{E}(a) &= 2 \int_{\phi \in (\phi_1, \phi_2) \cup (\phi_1 + \pi, \phi_2 + \pi)} \exp(- (1/2)(\theta^2 + p^2)) d\theta dp \\ &= 4 \int_{\phi \in (\phi_1, \phi_2)} \exp(- (1/2)(\theta^2 + p^2)) d\theta dp. \end{aligned}$$

The last integral is easily computed by changing to polar coordinates and then

$$\mathbb{E}(a) = \frac{2}{\pi}(\phi_2 - \phi_1).$$

We now recall the formula

$$\arctan(x) - \arctan(y) = \arctan\left(\frac{x - y}{1 + xy}\right)$$

(if the equality sign is understood modulo  $\pi$ , then the formula holds for all branches of the multivalued function  $\arctan(x)$ ) and, after some algebra find

$$\mathbb{E}(a) = \frac{2}{\pi} \arctan\left(\frac{2\sqrt{B^2 - AC}}{A + C}\right).$$

Taking into account (22) and (23), the last display may be rewritten in the form

$$\mathbb{E}(a) = \frac{2}{\pi} \arctan\left(\sqrt{\frac{2}{\mathbb{E}(\Delta H)}}\right); \quad (25)$$

the theorem follows from well-known properties of the function  $\arctan(x)$ .  $\square$

In the proof the slopes  $m_1$ ,  $m_2$  of the lines that separate energy growth from energy decay change with the integrator,  $\epsilon$  and  $L$ ; however the angle  $\phi_2 - \phi_1$  between the lines only depends on  $\mathbb{E}(\Delta H)$ .

Note that, as  $\mathbb{E}(\Delta H) \downarrow 0$ ,

$$\mathbb{E}(a) = 1 - \frac{\sqrt{2}}{\pi} \sqrt{\mathbb{E}(\Delta H)} + \mathcal{O}((\mathbb{E}(\Delta H))^{3/2}), \quad (26)$$

and, as a consequence, as  $\epsilon \downarrow 0$ ,  $L \uparrow \infty$  with  $L\epsilon = \tau_{end}$ ,

$$\mathbb{E}(a) = 1 - \mathcal{O}(\epsilon^\nu)$$

for an integrator of order of accuracy  $\nu$ . On the other hand, the formula (25) shows that, as  $\mathbb{E}(\Delta H) \uparrow \infty$ , the expected acceptance rate decays slowly, according to the estimate,

$$\mathbb{E}(a) \sim \frac{2\sqrt{2}}{\sqrt{\mathbb{E}(\Delta H)}}. \quad (27)$$

For instance, if  $\mathbb{E}(\Delta H) = 100$ , then (25) yields  $\mathbb{E}(a) \approx 0.089$ ; this should perhaps be compared with the data in Figure 4, where the target is low dimensional.

We conclude our analysis of the standard normal by presenting a result that will be required later. It is remarkable that the formulas in this lemma express moments of  $\Delta H$  as polynomials in  $\mathbb{E}(\Delta H)$  and that, furthermore, those polynomials do not change with the value of  $\epsilon$ ,  $L$  or with the choice of integrator.

**Lemma 2.** *For a univariate, standard Gaussian target with unit mass matrix and a time-reversible, volume preserving integrator, setting  $\mu = \mathbb{E}(\Delta H)$ , we have:*

$$\begin{aligned} \mathbb{E}((\Delta H)^2) &= 2\mu + 3\mu^2, \\ \mathbb{E}((\Delta H)^3) &= 18\mu^2 + 15\mu^3, \\ \mathbb{E}((\Delta H)^4) &= 36\mu^2 + 180\mu^3 + 105\mu^4. \end{aligned}$$

*Proof.* Raise (21) to the second (third or fourth) power. Compute expectations to write  $\mathbb{E}((\Delta H)^2)$  ( $\mathbb{E}((\Delta H)^3)$  or  $\mathbb{E}((\Delta H)^4)$ ) as a polynomial in  $A, B, C$ . Express  $B$  in terms of  $A$  and  $C$  by using (22) and a little patience. Finally use (23).  $\square$

As  $\mu \rightarrow 0$ ,  $\mathbb{E}((\Delta H)^2)/\mathbb{E}(\Delta H) \rightarrow 2$ , in agreement with Proposition 3.4 in Beskos et al. (2013). In addition, also  $\mathbb{E}((\Delta H)^4)/\mathbb{E}((\Delta H)^3) \rightarrow 2$ , a result that, while not contained in the paper by Beskos et al. (2013) may be proved with the tools used to prove that proposition.

### 5.3 Univariate Gaussian targets

For the univariate problem where  $\pi(\theta) = \mathcal{N}(0, \sigma^2)$  and  $M = 1$ , the Hamiltonian is  $(1/2)p^2 + (1/2)(\theta^2/\sigma^2)$  and the differential equations are  $d\theta/d\tau = p$ ,  $dp/d\tau = -\theta/\sigma^2$ . This problem may be reduced to the case  $\pi(\theta) = \mathcal{N}(0, 1)$ ,  $M = 1$  by scaling the variables, because for all integrators of practical interest the operations of rescaling and numerical integration commute. Specifically, the new variables are  $\theta/\sigma$ ,  $\tau/\sigma$  with  $p$  remaining as it was.

By changing variables we find that the bound (24) should be replaced by

$$0 \leq \mathbb{E}(\Delta H) \leq \rho(\epsilon/\sigma). \quad (28)$$

(The bound (16) is a direct consequence of this.) On the other hand, for arbitrary  $\sigma$ , Theorem 1 and Lemma 2 hold as they stand.

#### 5.4 Multivariate Gaussian targets

We now consider for each  $d = 1, 2, \dots$  a centered Gaussian target  $\pi_d(\theta)$ ,  $\theta \in \mathbb{R}^d$  with covariance matrix  $C_d$ . The  $C_d$  are assumed to be nonsingular, but otherwise they are allowed to be completely arbitrary. As pointed out before we may assume that the  $C_d$  have been diagonalized. We run HMC for each  $\pi_d(\theta)$  and, for simplicity, assume that the unit mass matrix  $I_d$  is used (but the result may be adapted to general mass matrices). Then  $\pi_d(\theta)$  is a product of univariate distributions  $\pi_{d,\ell}(\theta_\ell)$ ,  $\ell = 1, \dots, d$ , the Hamiltonian  $H_d$  is a sum of one-degree-of-freedom Hamiltonians  $H_{d,\ell}$  and the increment  $\Delta H_d(\theta, p)$  is a sum of increments  $\Delta H_{d,\ell}(\theta_\ell, p_\ell)$ . The integrator being used for  $\pi_d(\theta)$  is allowed to change with  $d$ ; it is only assumed that when applied to the standard normal univariate target takes the form (19)–(20). Similarly the step-length  $\epsilon$  and the length  $\tau_{end}$  of the integration interval are allowed to change with  $d$ .

Our next result shows that, even in this extremely general scenario, simple hypotheses on the expectations of the  $\Delta H_{d,\ell}$  ensure that, asymptotically, the variables  $\Delta H_d$  are normal with a variance that is twice as large as the mean (Neal, 2011).

**Theorem 2.** *In the general scenario described above, assume that*

$$M_d := \max_{1 \leq \ell \leq d} \mathbb{E}(\Delta H_{d,\ell}) \rightarrow 0 \quad (29)$$

and, for some  $\mu \in [0, \infty)$ ,

$$\mathbb{E}(\Delta H_d) = \sum_{\ell=1}^d \mathbb{E}(\Delta H_{d,\ell}) \rightarrow \mu. \quad (30)$$

Then, as  $d \uparrow \infty$ :

- The distributions of the random variables  $\Delta H_d$  at stationarity converge to the distribution  $\mathcal{N}(\mu, 2\mu)$ .

- The acceptance rates  $a_d$  for the targets  $\pi_d(\theta)$ ,  $d = 1, 2, \dots$ , satisfy

$$\mathbb{E}(a_d) \rightarrow 2\Phi(-\sqrt{\mu/2}).$$

*Proof.* The first item is a direct application of the central limit theorem. From Lemma 2, the variance of  $\Delta H_d$  is, with  $\mu_{d,\ell} := \mathbb{E}(\Delta H_{d,\ell})$ ,

$$s_d^2 = \sum_{\ell=1}^d (2\mu_{d,\ell} + 2\mu_{d,\ell}^2),$$

which tends to  $2\mu$  because, for any integer  $k > 1$ ,

$$\sum_{\ell=1}^d \mu_{d,\ell}^k \leq M_d^{k-1} \sum_{\ell=1}^d \mu_{d,\ell}. \quad (31)$$

We use the Lyapunov condition based on the centered fourth moment and consider the expression

$$\frac{1}{s_d^4} \sum_{\ell=1}^d \mathbb{E}((\Delta H_{d,\ell} - \mu_{d,\ell})^4).$$

We expand the fourth power and use Lemma 2 to write the expectation as a linear combination of  $\mu_{d,\ell}^2$ ,  $\mu_{d,\ell}^3$ ,  $\mu_{d,\ell}^4$  with coefficients that are independent of  $d$  and  $\ell$ . Then the sum in the display tends to 0 in view of (31) with  $k = 2, 3, 4$ .

For the second item, the boundedness of  $u \mapsto \min(1, \exp(-u))$  shows that

$$\mathbb{E}(a_d) \rightarrow \mathbb{E}(\min(1, \exp(-\zeta))),$$

where  $\zeta \sim \mathcal{N}(\mu, 2\mu)$ . The proof concludes by applying the lemma below.  $\square$

The following result, whose proof is an elementary computation and will be omitted, has been invoked in the proof of Theorem 2. The first equality should be compared with Lemma 1.

**Lemma 3.** *If  $\zeta \sim \mathcal{N}(\mu, 2\mu)$ , then*

$$\mathbb{E}\left(\min(1, \exp(-\zeta))\right) = 2\mathbb{P}(\zeta < 0) = 2\Phi(-\sqrt{\mu/2}).$$

As an application of Theorem 2, we review our findings in Figure 3.<sup>6</sup> In the scenario above, (17) plays the role of  $\pi_d(\theta)$ . We set  $\tau_{end} = 5$  and for each  $d$  choose a step-length  $\epsilon_d$  and one of the six integrators. According to (28),  $\mathbb{E}(\Delta H_{d,\ell}) \leq C\ell^4\epsilon_d^4$ , with  $C$  the largest of the constants corresponding to the different integrators tested. Assumption (29) will hold if  $\epsilon_d = o(d^{-1})$ . Furthermore, if  $\epsilon_d$  decays like  $d^{-5/4}$ , the sum in the left-hand side of (30) will be bounded above by a constant multiple of

$$\sum_{\ell=1}^d \ell^4 \epsilon_d^4 = \mathcal{O}(1),$$

and one may prove that (30) holds for some  $\mu < \infty$ . Therefore, for  $d$  large,  $\Delta H_d(\theta, p)$  will be asymptotically  $\mathcal{N}(\mu, 2\mu)$ . In fact the curve in Figures 3 and 4 is given by the equation  $a = 2\Phi(-\sqrt{\mu/2})$ . Of course if, as  $d$  varies,  $\epsilon$  decays faster than  $d^{-5/4}$  then the expectation of the acceptance rate would approach 100%; a slower decay would imply that the expectation would approach 0.

The equation  $a = 2\Phi(-\sqrt{\mu/2})$  is well known (Neal, 2011) and has been established rigorously by Beskos et al. (2013) in a restrictive scenario. There the target is assumed to be a product of  $d$  independent copies of *the same*, possibly multivariate, distribution  $\pi^*$ . The integrator and  $\tau_{end}$  are not allowed to vary with  $d$  and  $\epsilon_d$  has to decrease as  $d^{-1/4}$ . Some of these assumptions are far more restrictive than those we required above. On the other hand,  $\pi^*$  is *not* assumed to be Gaussian. Thus there are relevant differences between the analysis by Beskos et al. (2013) and the present study. However the underlying idea is the same in both treatments: one considers the energy error as a sum of many independent small contributions and for each contribution the variance is, approximately, twice as large as the mean. The role played by Proposition 3.4 in the paper by Beskos et al. (2013) is played here by Lemma 2, where we saw that the second moment  $\mathbb{E}((\Delta H)^2)$  (and by implication the variance) is approximately twice as large as the expectation  $\mathbb{E}(\Delta H)$  when the latter is small.

The function  $2\Phi(-\sqrt{\mu/2})$  decays exponentially as  $\mu \uparrow \infty$  (to be compared with (27)) and, as  $\mu \downarrow 0$  behaves as

$$1 - \frac{1}{2\sqrt{\pi}}\sqrt{\mu} + \mathcal{O}(\mu^{3/2})$$

(to be compared with (26)).

---

<sup>6</sup>The experiments for that figure had randomized step-lengths. Even though the effect of randomization may be easily analyzed, we prefer to ignore this issue in order to simplify the exposition.

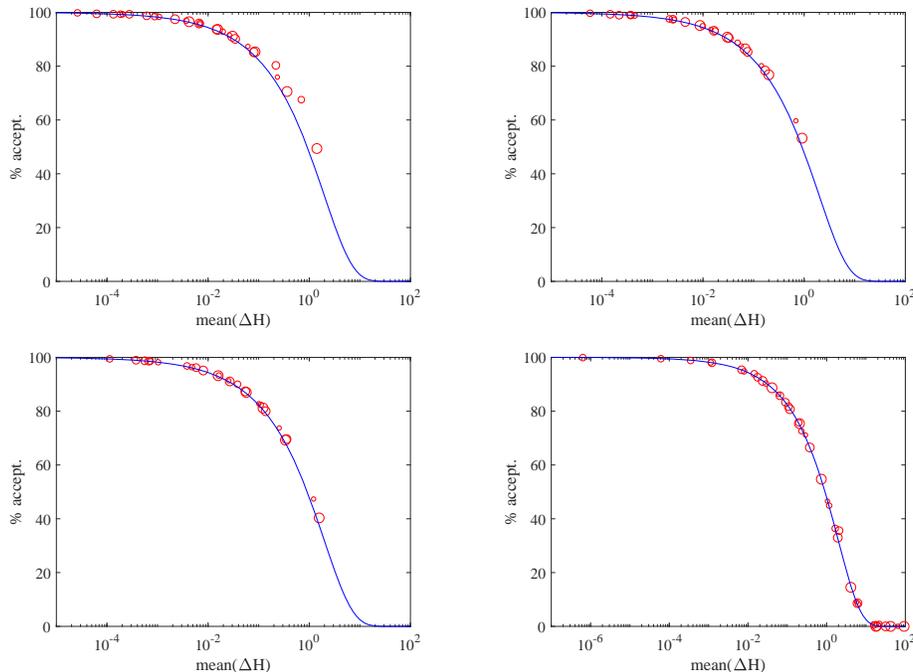


Figure 10: Alkane molecule with 5 carbon atoms (top left), 9 carbon atoms (top right) and 15 carbon atoms (bottom left), Log Gaussian Cox problem (bottom right). The dimensions  $d$  are 15, 27, 45 and 4096, respectively. Acceptance percentage vs. energy error. Points come from all six integrators and all step-lengths. The diameter of the circles indicates the integrator used. The continuous line corresponds to the case where  $\Delta H$  is  $\mathcal{N}(\mu, 2\mu)$ .

## 5.5 Other targets

Our final Figure 10 collects information from all six integrators and all step-lengths for the molecules with 5, 9 or 15 carbon atoms and the Log-Gaussian Cox problem. The points almost lie on the curve  $a = 2\Phi(-\sqrt{\mu/2})$ . Moreover, the fit gets better as the dimensionality increases. In other words, we experimentally observe here a behaviour that reproduces our theoretical analysis of the *Gaussian* case. This is in line with our earlier findings about the change of the average energy error as  $\epsilon$  varies. Thus the experiments in this paper seem to show that, for  $d$  large, properties of HMC for Gaussian targets also hold for general targets (or at least for a large class of relevant non-Gaussian targets).

## 6 Conclusions

Numerical experiments have proved that, for HMC sampling from probability distributions in  $\mathbb{R}^d$ ,  $d \gg 1$ , the integrator suggested by Blanes et al. (2014), while being as easy to implement as standard leapfrog, may deliver three times as many accepted samples as leapfrog with the same computational cost.

In addition we have given some theoretical results in connection with the acceptance rate of HMC at stationarity. We proved rigorously a central limit theorem behavior in a very general scenario of Gaussian distributions. We also provided a detailed analysis of the model problem given by the univariate normal distribution; this analysis implies that the the integrator in Blanes et al. (2014), while derived to minimize the expected energy error, also maximizes the expected acceptance rate.

**Acknowledgements.** M.P.C and J.M.S. were supported by project MTM2016-77660-P(AEI/FEDER, UE) MINECO (Spain). M.P.C. was also funded by project VA105G18 (Junta de Castilla y Leon, ES) cofinanced by FEDER funds. D.S-A. was supported by the NSF Grant DMS-1912818/1912802. J.M.S. would like to thank the Isaac Newton Institute for Mathematical Sciences (Cambridge) for support and hospitality during the programme “Geometry, compatibility and structure preservation in computational differential equations” when his work on this paper was completed.

## References

- E. Akhmetkaya, M. Fernández-Pendás, T. Radivojević, and J. M. Sanz-Serna. Adaptive splitting integrators for enhancing sampling efficiency of modified Hamiltonian Monte Carlo methods in molecular simulation. *Langmuir*, 33(42):11530–11542, 2017.
- A. Beskos, F. J. Pinski, J. M. Sanz-Serna, and A. M. Stuart. Hybrid Monte Carlo on Hilbert spaces. *Stochastic Processes and their Applications*, 121(10):2201–2230, 2011.
- A. Beskos, N. Pillai, G. O. Roberts, J. M. Sanz-Serna, and A. M. Stuart. Optimal tuning of the hybrid Monte Carlo algorithm. *Bernoulli*, 19(5A): 1501–1534, 2013.
- S. Blanes and F. Casas. *A Concise Introduction to Geometric Numerical Integration*. Chapman and Hall/CRC, 2016.

- S. Blanes, F. Casas, and A. Murua. Splitting and composition methods in the numerical integration of differential equations. *arXiv preprint arXiv:0812.0377*, 2008.
- S. Blanes, F. Casas, and J. M. Sanz-Serna. Numerical integrators for the hybrid Monte Carlo method. *SIAM Journal on Scientific Computing*, 36(4):A1556–A1580, 2014.
- N. Bou-Rabee and J. M. Sanz-Serna. Randomized Hamiltonian Monte Carlo. *The Annals of Applied Probability*, 27(4):2159–2194, 2017.
- N. Bou-Rabee and J. M. Sanz-Serna. Geometric integrators and the Hamiltonian Monte Carlo method. *Acta Numerica*, 27:113–206, 2018.
- C. M. Campos and J. M. Sanz-Serna. Extra chance generalized hybrid Monte Carlo. *Journal of Computational Physics*, 281:365–374, 2015.
- C. M. Campos and J. M. Sanz-Serna. Palindromic 3-stage splitting integrators, a roadmap. *Journal of Computational Physics*, 346:340–355, 2017.
- E. Cancès, F. Legoll, and G. Stoltz. Theoretical and numerical comparison of some sampling methods for molecular dynamics. *ESAIM: Mathematical Modelling and Numerical Analysis*, 41(2):351–389, 2007.
- T. Chen, E. Fox, and C. Guestrin. Stochastic gradient Hamiltonian Monte Carlo. In *International Conference on Machine Learning*, pages 1683–1691, 2014.
- O. F. Christensen, G. O. Roberts, and J. S. Rosenthal. Scaling limits for the transient phase of local Metropolis–Hastings algorithms. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):253–268, 2005.
- S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Physics letters B*, 195(2):216–222, 1987.
- Y. Fang, J. M. Sanz-Serna, and R. D. Skeel. Compressible generalized hybrid Monte Carlo. *The Journal of Chemical Physics*, 140(17):174108, 2014.
- M. Fernández-Pendás, E. Akhmetkaya, and J. M. Sanz-Serna. Adaptive multi-stage integrators for optimal energy conservation in molecular simulations. *Journal of Computational Physics*, 327:434–449, 2016.

- M. Girolami and B. Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.
- S. Gupta, A. Irbäck, F. Karsch, and B. Petersson. The acceptance probability in the hybrid Monte Carlo method. *Physics Letters B*, 242(3-4):437–443, 1990.
- M. D. Hoffman and A. Gelman. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014.
- A. M. Horowitz. A generalized guided Monte Carlo algorithm. *Physics Letters B*, 268(2):247–252, 1991.
- S. Livingstone, M. Betancourt, S. Byrne, and M. Girolami. On the geometric ergodicity of Hamiltonian Monte Carlo. *Bernoulli*, 25(4A):3109–3138, 2019.
- O. Mangoubi, N. S. Pillai, and A. Smith. Does Hamiltonian Monte Carlo mix faster than a random walk on multimodal densities? *arXiv preprint arXiv:1808.03230*, 2018.
- R. M. Neal. *Probabilistic inference using Markov chain Monte Carlo methods*. Department of Computer Science, University of Toronto Toronto, Ontario, Canada, 1993.
- R. M. Neal. MCMC using Hamiltonian dynamics. *Handbook of Markov chain Monte Carlo*, 2(11):2, 2011.
- R. M. Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- 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, et al. Computationally efficient molecular dynamics integrators with improved sampling accuracy. *Molecular Physics*, 110(9-10):967–983, 2012.
- T. Radivojević and E. Akhmatskaya. Modified Hamiltonian Monte Carlo for Bayesian Inference. *Statistics and Computing*, pages 1–28, 2019.
- T. Radivojević, M. Fernández-Pendás, J. M. Sanz-Serna, and E. Akhmatskaya. Multi-stage splitting integrators for sampling with modified Hamiltonian Monte Carlo methods. *Journal of Computational Physics*, 373:900–916, 2018.

- J. M. Sanz-Serna. Markov Chain Monte Carlo and Numerical Differential Equations. In *Current Challenges in Stability Issues for Numerical Differential Equations*, pages 39–88. Springer, 2014.
- J. M. Sanz-Serna and M. P. Calvo. *Numerical Hamiltonian Problems*. Dover Publications, 2018.