



Split Hamiltonian Monte Carlo revisited

Fernando Casas¹ · Jesús María Sanz-Serna² · Luke Shaw¹

Received: 15 July 2022 / Accepted: 6 September 2022
© The Author(s) 2022

Abstract

We study Hamiltonian Monte Carlo (HMC) samplers based on splitting the Hamiltonian H as $H_0(\theta, p) + U_1(\theta)$, where H_0 is quadratic and U_1 small. We show that, in general, such samplers suffer from stepsize stability restrictions similar to those of algorithms based on the standard leapfrog integrator. The restrictions may be circumvented by preconditioning the dynamics. Numerical experiments show that, when the $H_0(\theta, p) + U_1(\theta)$ splitting is combined with preconditioning, it is possible to construct samplers far more efficient than standard leapfrog HMC.

Keywords Markov chain Monte Carlo · Hamiltonian dynamics · Bayesian analysis · Splitting integrators

1 Introduction

In this paper we study Hamiltonian Monte Carlo (HMC) algorithms (Neal 2011) that are not based on the standard kinetic/potential splitting of the Hamiltonian.

The computational cost of HMC samplers mostly originates from the numerical integrations that have to be performed to get the proposals. If the target distribution has density proportional to $\exp(-U(\theta))$, $\theta \in \mathbb{R}^d$, the differential system to be integrated is given by the Hamilton's equations corresponding to the Hamiltonian function $H(\theta, p) = (1/2)p^T M^{-1}p + U(\theta)$, where $p \sim \mathcal{N}(0, M)$ is the auxiliary momentum variable and M is the symmetric, positive definite *mass matrix* chosen by the user. In a mechanical analogy, H is the (total) energy, while $\mathcal{T}(p) = (1/2)p^T M^{-1}p$ and $U(\theta)$ are respectively the kinetic and potential energies. The Störmer/leapfrog/Verlet integrator is the method of choice to carry out those integrations and is based on the idea of *splitting* (Blanes and Casas 2017), i.e. the evolution of (θ, p) under H is simulated by the separate evolutions

under $\mathcal{T}(p)$ and $U(\theta)$ (kinetic/potential splitting). However $H(\theta, p) = \mathcal{T}(p) + U(\theta)$ is not the only splitting that has been considered in the literature. In some applications one may write $H(\theta, p) = H_0(\theta, p) + U_1(\theta)$, with $H_0(\theta, p) = \mathcal{T}(p) + U_0(\theta)$, $U(\theta) = U_0(\theta) + U_1(\theta)$ and replace the evolution under H by the evolutions under H_0 and U_1 (Neal 2011). The paper (Shahbaba et al. 2014) investigated this possibility; two algorithms were formulated referred to there as “Leapfrog with a partial analytic solution” and “Nested leapfrog”. Both suggested algorithms were shown to outperform, in four logistic regression problems, HMC based on the standard leapfrog integrator.

In this article we reexamine $H(\theta, p) = H_0(\theta, p) + U_1(\theta)$ splittings, in particular in the case where the equations for H_0 can be integrated analytically (partial analytic solution) because $U_0(\theta)$ is a quadratic function (so that $\propto \exp(-U_0(\theta))$ is a Gaussian distribution). When U_1 is slowly varying, the splitting $H = H_0 + U_1$ is appealing because, to quote (Shahbaba et al. 2014), “only the slowly-varying part of the energy needs to be handled numerically and this can be done with a larger stepsize (and hence fewer steps) than would be necessary for a direct simulation of the dynamics”.

Our contributions are as follows:

1. In Section 3 we show, by means of a counterexample, that it is not necessarily true that, when H_0 is handled analytically and U_1 is small, the integration may be carried out with stepsizes substantially larger than those required by standard leapfrog. For integrators based on the $H_0 + U_1$

✉ Luke Shaw
shaw@uji.es

Fernando Casas
casas@uji.es

Jesús María Sanz-Serna
jmsanzserna@gmail.com

¹ Departament de Matemàtiques and IMAC, Universitat Jaume I, E-12071 Castellón, Spain

² Departamento de Matemáticas, Universidad Carlos III de Madrid, E-28911 Leganés, Spain

splitting, the stepsize may suffer from important stability restrictions, regardless of the size of U_1 .

2. In Section 4 we show that, by combining the $H_0 + U_1$ splitting with the idea of *preconditioning* the dynamics, that goes back at least to (Bennett 1975), it is possible to bypass the stepsize limitations mentioned in the preceding item.
3. We present an integrator (that we call RKR) for the $H_0 + U_1$ splitting that provides an alternative to the integrator tested in (Shahbaba et al. 2014) (that we call KRK).
4. Numerical experiments in the final Section 5, using the test problems in (Shahbaba et al. 2014), show that the advantages of moving from standard leapfrog HMC to the $H_0 + U_1$ splitting (without preconditioning) are much smaller than the advantages of using preconditioning while keeping the standard kinetic/potential splitting. The best performance is obtained when the $H_0 + U_1$ splitting is combined with the preconditioning of the dynamics. In particular the RKR integration technique with preconditioning decreases the computational cost by more than an order of magnitude in all test problems and all observables considered.

There are two appendices. In the first, we illustrate the use of the Bernstein-von Mises theorem (see e.g. section 10.2 in (van der Vaart 1998)) to justify the soundness of the $H_0 + U_1$ splitting. The second is devoted to presenting a methodology to discriminate between different integrators of the preconditioned dynamics for the $H_0 + U_1$ splitting; in particular we provide analyses that support the advantages of the RKR technique over its KRK counterpart observed in the experiments.

2 Preliminaries

2.1 Hamiltonian Monte Carlo

HMC is based on the observation that (Neal 2011; Sanz-Serna 2014), for each fixed $T > 0$, the exact solution map (flow) $(\theta(T), p(T)) = \varphi_T(\theta(0), p(0))$ of the Hamiltonian system of differential equations in \mathbb{R}^{2d}

$$\frac{d\theta}{dt} = \frac{\partial H}{\partial p} = M^{-1}p, \quad \frac{dp}{dt} = -\frac{\partial H}{\partial \theta} = -\nabla U(\theta), \quad (1)$$

exactly preserves the density $\propto \exp(-H(\theta, p)) = \exp(-T(p) - U(\theta))$ whose θ -marginal is the target $\propto \exp(-U(\theta))$, $\theta \in \mathbb{R}^d$. In HMC, (1) is integrated numerically over an interval $0 \leq t \leq T$ taking as initial condition the current state (θ, p) of the Markov chain; the numerical solution at $t = T$ provides the proposal (θ', p') that is accepted with probability

$$a = \min \left\{ 1, e^{-\left(H(\theta', p') - H(\theta, p)\right)} \right\}. \quad (2)$$

This formula for the acceptance probability assumes that the numerical integration has been carried out with an integrator that is both *symplectic* (or at least volume preserving) and *reversible*. The difference $H(\theta', p') - H(\theta, p)$ in (2) is the energy error in the integration; it would vanish leading to $a = 1$ if the integration were exact.

2.2 Splitting

Splitting is the most common approach to derive symplectic integrators for Hamiltonian systems (Blanes and Casas 2017; Sanz-Serna and Calvo 1994). The Hamiltonian H of the problem is decomposed in partial Hamiltonians as $H = H_1 + H_2$ in such a way that the Hamiltonian systems with Hamiltonian functions H_1 and H_2 may both be integrated in closed form. When Strang splitting is used, if $\varphi_t^{[H_1]}, \varphi_t^{[H_2]}$ denote the maps (flows) in \mathbb{R}^{2d} that advance the exact solution of the partial Hamiltonians over a time-interval of length t , the recipe

$$\psi_\epsilon = \varphi_{\epsilon/2}^{[H_1]} \circ \varphi_\epsilon^{[H_2]} \circ \varphi_{\epsilon/2}^{[H_1]}, \quad (3)$$

defines the map that advances the numerical solution a timestep of length $\epsilon > 0$. The numerical integration to get a proposal may then be carried out up to time $T = \epsilon L$ with the L -fold composition $\Psi_T = (\psi_\epsilon)^L$. Regardless of the choice of H_1 and H_2 , (3) is a symplectic, time reversible integrator of second order of accuracy (Bou-Rabee and Sanz-Serna 2018).

2.3 Kinetic/potential splitting

The splitting $H = H_1 + H_2$, $H_1 = T$, $H_2 = U$ gives rise, via (3), to the commonest integrator in HMC: the Störmer/leapfrog/velocity Verlet algorithm. The differential equations for the partial Hamiltonians T , U and the corresponding solution flows are

$$\begin{aligned} \frac{d}{dt} \begin{pmatrix} \theta \\ p \end{pmatrix} &= \begin{pmatrix} 0 \\ -\nabla U(\theta) \end{pmatrix} \implies \varphi_\epsilon^{[U]}(\theta, p) = (\theta, p - \epsilon \nabla U(\theta)), \\ \frac{d}{dt} \begin{pmatrix} \theta \\ p \end{pmatrix} &= \begin{pmatrix} M^{-1}p \\ 0 \end{pmatrix} \implies \varphi_\epsilon^{[T]}(\theta, p) = (\theta + \epsilon M^{-1}p, p). \end{aligned}$$

As a mnemonic, we shall use the word *kick* to refer to the map $\varphi_\epsilon^{[U]}(\theta, p)$ (the system is kicked so that the momentum p varies without changing θ). The word *drift* will refer to the map $\varphi_\epsilon^{[T]}(\theta, p)$ (θ drifts with constant velocity). Thus one timestep of the velocity Verlet algorithm reads (kick-drift-kick).

$$\psi_\epsilon^{[KDK]} = \varphi_{\epsilon/2}^{[U]} \circ \varphi_\epsilon^{[T]} \circ \varphi_{\epsilon/2}^{[U]}.$$

There is of course a position Verlet algorithm obtained by interchanging the roles of \mathcal{T} and U . One timestep is given by a sequence drift-kick-drift (DKD). Generally the velocity Verlet (KDK) version is preferred (see (Bou-Rabee and Sanz-Serna 2018) for a discussion) and we shall not be concerned hereafter with the position variant.

With any integrator of the Hamiltonian equations, the length $\epsilon L = T$ of the time interval for the integration to get a proposal has to be determined to ensure that the proposal is sufficiently far from the current step of the Markov chain, so that the correlation between successive samples is not too high and the phase space is well explored (Hoffman and Gelman 2014; Bou-Rabee and Sanz-Serna 2017). For fixed T , smaller stepsizes ϵ lead to fewer rejections but also to larger computational cost per integration and it is known that HMC is most efficient when the empirical acceptance rate is around approximately 65% (Beskos et al. 2013).

Algorithm 1 describes the computation to advance a single step of the Markov chain with HMC based on the velocity Verlet (KDK) integrator. In the absence of additional information, it is standard practice to choose $M = I$, the identity matrix. For later reference, we draw attention to the randomization of the timestep ϵ . As is well known, without such a randomization, HMC may not be ergodic (Neal 2011); this will happen for instance when the equations of motion (1) have periodic solutions and ϵL coincides with the period of the solution.

Algorithm 1 KDK Verlet

```

Input:  $\theta, U, H, M, \bar{\epsilon}, L$ 
1: Draw  $\xi \sim \mathcal{N}(0, M)$ ,  $\epsilon \sim \bar{\epsilon} \times \mathcal{U}_{[0.8, 1]}$  ▷ Randomise  $\epsilon$ 
2:  $\theta', p \leftarrow \theta, \xi$  ▷ Refresh momentum
3: for  $i = \{1, \dots, L\}$  do ▷ Do velocity Verlet integration
4:    $p \leftarrow p - \frac{\epsilon}{2} \nabla U(\theta')$ 
5:    $\theta' \leftarrow \theta' + \epsilon M^{-1} p$ 
6:    $p \leftarrow p - \frac{\epsilon}{2} \nabla U(\theta')$ 
7: end for
8:  $a \leftarrow \min \{1, \exp(H(\theta, \xi) - H(\theta', p))\}$ 
9: Draw  $\gamma \sim \mathcal{B}(a)$  ▷  $\gamma$  Bernoulli-distributed with mean  $a$ 
10:  $\theta \leftarrow \gamma \theta' + (1 - \gamma)\theta$  ▷ Accept proposal with probability  $a$ 
    
```

2.4 Alternative splittings of the Hamiltonian

Splitting $H(\theta, p)$ in its kinetic and potential parts as in Verlet is not the only meaningful possibility. In many applications, $U(\theta)$ may be written as $U_0(\theta) + U_1(\theta)$ in such a way that the equations of motion for the Hamiltonian function $H_0(\theta, p) = (1/2)p^T M^{-1} p + U_0(\theta)$ may be integrated in closed form and then one may split H as

$$H = H_0 + U_1, \tag{4}$$

as discussed in e.g. (Neal 2011; Shahbaba et al. 2014).

In this paper we focus on the important particular case where (see Section 5 and Appendix A)

$$U_0(\theta) = \frac{1}{2}(\theta - \theta^*)^T \mathcal{J}(\theta - \theta^*), \tag{5}$$

for some fixed $\theta^* \in \mathbb{R}^d$ and a constant symmetric, positive definite matrix \mathcal{J} . Restricting for the time being attention to the case where the mass matrix M is the identity (the only situation considered in (Shahbaba et al. 2014)), the equations of motion and solution flow for the Hamiltonian

$$H_0(\theta, p) = \frac{1}{2}p^T p + U_0(\theta) \tag{6}$$

are

$$\begin{aligned} \frac{d}{dt} \begin{pmatrix} \theta \\ p \end{pmatrix} &= \begin{pmatrix} 0 & I \\ -\mathcal{J} & 0 \end{pmatrix} \begin{pmatrix} \theta - \theta^* \\ p \end{pmatrix}, \\ \varphi_t^{[H_0]}(\theta, p) &= \exp \left(t \begin{pmatrix} 0 & I \\ -\mathcal{J} & 0 \end{pmatrix} \right) \begin{pmatrix} \theta - \theta^* \\ p \end{pmatrix} + \begin{pmatrix} \theta^* \\ 0 \end{pmatrix}. \end{aligned} \tag{7}$$

If we write $\mathcal{J} = Z^T D Z$, with Z orthogonal and D diagonal with positive diagonal elements, then the exponential map in Eq. (7) is

$$\begin{aligned} \exp \left(t \begin{pmatrix} 0 & I \\ -\mathcal{J} & 0 \end{pmatrix} \right) &= \begin{pmatrix} Z^T & 0 \\ 0 & Z^T \end{pmatrix} e^{t\Lambda} \begin{pmatrix} Z & 0 \\ 0 & Z \end{pmatrix}, \\ e^{t\Lambda} &= \begin{pmatrix} \cos(t\sqrt{D}) & D^{-1/2} \sin(t\sqrt{D}) \\ -D^{1/2} \sin(t\sqrt{D}) & \cos(t\sqrt{D}) \end{pmatrix}. \end{aligned} \tag{8}$$

In view of the expression for $\exp(t\Lambda)$, we will refer to the flow of H_0 as a *rotation*.

Choosing in (3) U_1 and H_0 for the roles of H_1 and H_2 (or viceversa) gives rise to the integrators

$$\begin{aligned} \psi_\epsilon^{[K RK]} &= \varphi_{\epsilon/2}^{[U_1]} \circ \varphi_\epsilon^{[H_0]} \circ \varphi_{\epsilon/2}^{[U_1]}, \\ \psi_\epsilon^{[R KR]} &= \varphi_{\epsilon/2}^{[H_0]} \circ \varphi_\epsilon^{[U_1]} \circ \varphi_{\epsilon/2}^{[H_0]}, \end{aligned} \tag{9}$$

where one advances the solution over a single timestep by using a kick-rotate-kick (KRK) or rotate-kick-rotate (RKR) pattern (of course the kicks are based on the potential function U_1). The HMC algorithm with the KRK map in (9) is shown in Algorithm 2, where the prefix *Uncond*, to be discussed later, indicates that the mass matrix being used is $M = I$. The algorithm for the RKR sequence in (9) is a slight reordering of a few lines of code and is not shown. Algorithm 2 (but not its RKR counterpart) was tested in (Shahbaba et al. 2014).¹

¹ It is perhaps of interest to mention that in Algorithm 2 the stepsize ϵ is randomized for the same reasons as in Algorithm 1. If only the stepsize used in the U_1 -kicks is randomized, while the stepsize in $\exp(\epsilon\Lambda)$ is kept constant, then one still risks losing ergodicity when ϵL coincides with

Algorithm 2 UncondKRK

Input: $\theta, Z, \Lambda, U, \mathcal{J}, H, \bar{\epsilon}, L$

- 1: Draw $\xi \sim \mathcal{N}(0, I)$, $\epsilon \sim \bar{\epsilon} \times \mathcal{U}_{[0.8, 1]}$
- 2: Compute $e^{\epsilon\Lambda}$
- 3: $\theta', p \leftarrow \theta, \xi$
- 4: **for** $i = \{1, \dots, L\}$ **do** ▷ Do KRK integration
- 5: $p \leftarrow p - \frac{\epsilon}{2} (\nabla U(\theta') - \mathcal{J}(\theta' - \theta^*))$
- 6: $\theta', p \leftarrow Z(\theta' - \theta^*), Zp$
- 7: $\theta', p \leftarrow e^{\epsilon\Lambda} \begin{pmatrix} \theta' \\ p \end{pmatrix}$
- 8: $\theta', p \leftarrow Z^T \theta' + \theta^*, Z^T p$
- 9: $p \leftarrow p - \frac{\epsilon}{2} (\nabla U(\theta') - \mathcal{J}(\theta' - \theta^*))$
- 10: **end for**
- 11: $a \leftarrow \min \{1, \exp(H(\theta, \xi) - H(\theta', p))\}$
- 12: Draw $\gamma \sim \mathcal{B}(a)$
- 13: $\theta \leftarrow \gamma\theta' + (1 - \gamma)\theta$

Since the numerical integration in Algorithm 2 would be exact if U_1 vanished (leading to acceptance of all proposals), the algorithm is appealing in cases where U_1 is “small” with respect to H_0 . In some applications, a decomposition $U = U_0 + U_1$ with small U_1 may suggest itself. For a “general” U one may always define U_0 by choosing θ^* to be one of the modes of the target $\propto \exp(-U(\theta))$ and \mathcal{J} the Hessian of U evaluated at θ^* ; in this case the success of the splitting hinges on how well U may be approximated by its second-order Taylor expansion U_0 around θ^* . In that setting, θ^* would typically have to be found numerically by minimizing U . Also Z and D would typically be derived by numerical approximation, thus leading to computational overheads for Algorithm 2 not present in Algorithm 1. However, as pointed out in (Shahbaba et al. 2014), the cost of computing θ^* , Z and D before the sampling begins is, for the test problems to be considered in this paper, negligible when compared with the cost of obtaining the samples.

2.5 Nesting

When a decomposition $U = U_0 + U_1$, with U_1 small, is available but the Hamiltonian system with Hamiltonian $H_0 = \mathcal{T} + U_0$ cannot be integrated in closed form, one may still construct schemes based on the recipe (3). One step of the integrator is defined as

$$\varphi_{\epsilon/2}^{[U_1]} \circ \left(\varphi_{\epsilon/2k}^{[U_0]} \circ \varphi_{\epsilon/k}^{[T]} \circ \varphi_{\epsilon/2k}^{[U_0]} \right)^k \circ \varphi_{\epsilon/2}^{[U_1]}, \quad (10)$$

where k is a suitably large integer. Here the (untractable) exact flow of H_0 is numerically approximated by KDK Verlet using k substeps of length ϵ/k . In this way, kicks with the small U_1 are performed with a stepsize $\epsilon/2$ and kicks with the large U_0 benefit from the smaller stepsize $\epsilon/(2k)$. This

one of the periods present in the solution. This prevents precalculation, prior to the randomization of ϵ , of the rotation matrix $\exp(\epsilon\Lambda)$.

idea has been successfully used in Bayesian applications in (Shahbaba et al. 2014), where it is called “nested Verlet”. The small U_1 is obtained summing over data points that contribute little to the loglikelihood and the contributions from the most significant data are included in U_0 .

Integrators similar to (10) have a long history in molecular dynamics, where they are known as multiple timestep algorithms (Tuckerman et al. 1992; Leimkuhler and Matthews 2015; Grubmüller et al. 1991).

3 Shortcomings of the unconditioned KRK and RKR samplers

As we observed above, Algorithm 2 is appealing when U_1 is a small perturbation of the quadratic Hamiltonian H_0 . In particular, one would expect that since the numerical integration in Algorithm 2 is exact when U_1 vanishes, then this algorithm may be operated with stepsizes ϵ chosen solely in terms of the size of U_1 , independently of H_0 . If that were the case one would expect that Algorithm 2 may work well with large ϵ in situations where Algorithm 1 requires ϵ small and therefore much computational effort. Unfortunately those expectations are not well founded, as we shall show next by means of an example.

We study the model Hamiltonian with $\theta, p \in \mathbb{R}^2$ given by

$$H(\theta, p) = H_0(\theta, p) + U_1(\theta),$$

$$H_0 = \frac{1}{2}p^T p + \frac{1}{2}\theta^T \begin{pmatrix} \sigma_1^{-2} & 0 \\ 0 & \sigma_2^{-2} \end{pmatrix} \theta, \quad U_1 = \frac{\kappa}{2}\theta^T \theta. \quad (11)$$

The model is restricted to \mathbb{R}^2 just for notational convenience; the extension to \mathbb{R}^d is straightforward. The quadratic Hamiltonian H_0 is rather general—any Hamiltonian system with quadratic Hamiltonian $(1/2)p^T M^{-1}p + (1/2)\theta^T W\theta$ may be brought with a change of variables to a system with Hamiltonian of the form $(1/2)p^T p + (1/2)\theta^T D\theta$, with M, W symmetric, positive definite matrices and D diagonal and positive definite (Blanes et al. 2014; Bou-Rabee and Sanz-Serna 2018). In (11), σ_1 and σ_2 are the standard deviations of the bivariate Gaussian distribution with density $\propto \exp(-U_0(\theta))$ (i.e. of the target in the unperturbed situation $U_1 = 0$). We choose the labels of the scalar components θ_1 and θ_2 of θ to ensure $\sigma_1 \leq \sigma_2$ so that, for the probability density $\propto \exp(-U_0(\theta))$, θ_1 is more constrained than θ_2 . In addition, we assume that κ is small with respect to

σ_1^{-2} and σ_2^{-2} , so that in (11) U_1 is a small perturbation of H_0 . The Hamiltonian equations of motion for θ_i , given by $\frac{d}{dt}\theta_i = p_i$, $\frac{d}{dt}p_i = -\omega_i^2\theta_i$, with $\omega_i = (\sigma_i^{-2} + \kappa)^{1/2} \approx \sigma_i^{-1}$, yield $\frac{d^2}{dt^2}\theta_i + \omega_i^2\theta_i = 0$. Thus the dynamics of θ_1 and θ_2 correspond to two uncoupled harmonic oscillators; the component θ_i , $i = 1, 2$, oscillates with an angular frequency ω_i (or with a period $2\pi/\omega_i$).

$$\begin{pmatrix} \theta_i \\ p_i \end{pmatrix} \leftarrow \begin{pmatrix} 1 & 0 \\ -\epsilon\kappa/2 & 1 \end{pmatrix} \begin{pmatrix} \cos(\epsilon/\sigma_i) & \sigma_i \sin(\epsilon/\sigma_i) \\ -\sigma_i^{-1} \sin(\epsilon/\sigma_i) & \cos(\epsilon/\sigma_i) \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -\epsilon\kappa/2 & 1 \end{pmatrix} \begin{pmatrix} \theta_i \\ p_i \end{pmatrix}$$

or

$$\begin{pmatrix} \theta_i \\ p_i \end{pmatrix} \leftarrow \begin{pmatrix} \cos(\epsilon/\sigma_i) - (\epsilon\sigma_i\kappa/2) \sin(\epsilon/\sigma_i) & \sigma_i \sin(\epsilon/\sigma_i) \\ (\epsilon^2\sigma_i\kappa^2/4) \sin(\epsilon/\sigma_i) - \sigma_i^{-1} \sin(\epsilon/\sigma_i) + \epsilon\kappa \cos(\epsilon/\sigma_i) & \cos(\epsilon/\sigma_i) - (\epsilon\sigma_i\kappa/2) \sin(\epsilon/\sigma_i) \end{pmatrix} \begin{pmatrix} \theta_i \\ p_i \end{pmatrix}.$$

We note, regardless of the integrator being used, the correlation between the proposal and the current state of the Markov chain will be large if the integration is carried out over a time interval $T = \epsilon L$ much smaller than the periods $2\pi/\omega_i$ of the harmonic oscillators (Neal 2011; Bou-Rabee and Sanz-Serna 2017). Since $2\pi/\omega_2$ is the longest of the two periods, L has then to be chosen

$$L \geq \frac{C}{\epsilon\omega_2} \approx \frac{C\sigma_2}{\epsilon}, \tag{12}$$

where C denotes a constant of moderate size. For instance, for the choice $C = \pi/2$, the proposal for θ_2 is uncorrelated at stationarity with the current state of the Markov chain as discussed in e.g. (Bou-Rabee and Sanz-Serna 2017).

For the KDK Verlet integrator, it is well known that, for stability reasons (Neal 2011; Bou-Rabee and Sanz-Serna 2018), the integration has to be operated with a stepsize $\epsilon < 2/\max(\omega_1, \omega_2)$, leading to a stability limit

$$\epsilon \approx 2\sigma_1; \tag{13}$$

integrations with larger ϵ will lead to extremely inaccurate numerical solutions. This stability restriction originates from θ_1 , the component with greater precision in the Gaussian distribution $\propto \exp(-U_0)$. Combining (13) with (12) we conclude that, for Verlet, the number of timesteps L has to be chosen larger than a moderate multiple of σ_2/σ_1 . Therefore when $\sigma_1 \ll \sigma_2$ the computational cost of the Verlet integrator will necessarily be very large. Note that the inefficiency arises when the sizes of σ_1 and σ_2 are widely different; the first sets an upper bound for the stepsize and the second a lower bound on the length ϵL of the integration interval. Small or large values of σ_1 and σ_2 are not dangerous *per se* if σ_2/σ_1 is moderate.

We now turn to the KRK integrator in (9). For the i -th scalar component of (θ, p) , a timestep of the KRK integrator reads

Stability is equivalent to $|\cos(\epsilon/\sigma_i) - (\epsilon\sigma_i\kappa/2) \sin(\epsilon/\sigma_i)| < 1$, which, for $\kappa > 0$, gives $2 \cot((\epsilon/(2\sigma_i))) > \epsilon\kappa\sigma_i$. From here it is easily seen that stability in the i -th component is lost for $\epsilon/\sigma_i \approx \pi$ for arbitrarily small $\kappa > 0$. Thus the KRK stability limit is

$$\epsilon \approx \pi\sigma_1. \tag{14}$$

While this is less restrictive than (13), we see that stability imposes an upper bound for ϵ in terms of σ_1 , just as for Verlet. From (12), the KRK integrator, just like Verlet, will have a large computational cost when $\sigma_1 \ll \sigma_2$. This is in spite of the fact that the integrator would be exact for $\kappa = 0$, regardless of the values of σ_1, σ_2 .

For the RKR integrator a similar analysis shows that the stability limit is also given by (14); therefore that integrator suffers from the same shortcomings as KRK.

We also note that, since as k increases the nested integrator (10) approximates the KRK integrator, the counterexample above may be used to show that the nested integrator has to be operated with a stepsize ϵ that is limited by the smallest standard deviations present in U_0 , as is the case for Verlet, KRK and RKR. For the stability of (10) and related multiple timestep techniques, the reader is referred to (García-Archilla et al. 1998) and its references. The nested integrator will not be considered further in this paper.

4 Preconditioning

As pointed out above, without additional information on the target, it is standard to set $M = I$. When $U = U_0 + U_1$, with U_0 as in (5), it is useful to consider a *preconditioned* Hamiltonian with $M = \mathcal{J}$:

$$\begin{aligned} H^{\{precond\}}(\theta, p) &= \frac{1}{2}p^T \mathcal{J}^{-1}p + U(\theta) \\ &= \frac{1}{2}p^T \mathcal{J}^{-1}p + \frac{1}{2}(\theta - \theta^*)^T \mathcal{J}(\theta - \theta^*) + U_1(\theta). \end{aligned} \tag{15}$$

Preconditioning is motivated by the observation that the equations of motion for the Hamiltonian

$$H_0^{[precond]}(\theta, p) = \frac{1}{2}p^T \mathcal{J}^{-1} p + \frac{1}{2}(\theta - \theta^*)^T \mathcal{J}(\theta - \theta^*),$$

given by $\frac{d}{dt}\theta = \mathcal{J}^{-1}p$, $\frac{d}{dt}p = -\mathcal{J}(\theta - \theta^*)$, yield $\frac{d^2}{dt^2}(\theta - \theta^*) + (\theta - \theta^*) = 0$. Thus we now have d uncoupled scalar harmonic oscillators (one for each scalar component $\theta_i - \theta_i^*$) sharing a *common* oscillation frequency $\omega = 1$.² This is to be compared with the situation for (6), where, as we have seen in the model (11), the frequencies are the reciprocals $1/\sigma_i$ of the standard deviations of the distribution $\propto \exp(-U_0(\theta))$. Since, as we saw in Section 3, it is the differences in size of the frequencies of the harmonic oscillators that cause the inefficiency of the integrators, choosing the mass matrix to ensure that all oscillators have the *same* frequency is of clear interest. We call *unconditioned* those Hamiltonians/integrators where the mass matrix is chosen as the identity agnostically without specializing it to the problem.

For reasons explained in (Beskos et al. 2011) it is better, when \mathcal{J} has widely different eigenvalues, to numerically integrate the preconditioned equations of motion after rewriting them with the variable $v = M^{-1}p = \mathcal{J}^{-1}p$ replacing p . The differential equations and solution flows of the sub-problems are then given by

$$\begin{aligned} \frac{d}{dt} \begin{pmatrix} \theta \\ v \end{pmatrix} &= \begin{pmatrix} 0 \\ -\mathcal{J}^{-1}\nabla_{\theta}U_1(\theta) \end{pmatrix} \\ \implies \varphi_t^{[U_1]}(\theta, v) &= \begin{pmatrix} \theta \\ v - t\mathcal{J}^{-1}\nabla_{\theta}U_1(\theta) \end{pmatrix}, \end{aligned}$$

and

$$\begin{aligned} \frac{d}{dt} \begin{pmatrix} \theta \\ v \end{pmatrix} &= \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix} \begin{pmatrix} (\theta - \theta^*) \\ v \end{pmatrix} \\ \implies \varphi_t^{[H_0^{[precond]}]}(\theta, v) &= \begin{pmatrix} \cos(t) & \sin(t) \\ -\sin(t) & \cos(t) \end{pmatrix} \begin{pmatrix} (\theta - \theta^*) \\ v \end{pmatrix} + \begin{pmatrix} \theta^* \\ 0 \end{pmatrix}. \end{aligned}$$

Since \mathcal{J} is a symmetric, positive definite matrix, it admits a Cholesky factorisation $\mathcal{J} = BB^T$. The inversion of \mathcal{J} in the kick may thus be performed efficiently using Cholesky-based solvers from standard linear algebra libraries. It also means it is easy to draw from the distribution of $v \sim B^{-T}\mathcal{N}(0, I)$.

Composing the exact maps $\varphi_{\epsilon}^{[-]}$ using Strang’s recipe (3) then gives a numerical one-step map $\psi_{\epsilon}^{[+]}$ in either an RKR or

KRK form. The preconditioned KRK (PrecondKRK) algorithm is shown in Algorithm 3; the RKR version is similar and will not be given.

Algorithm 3 PrecondKRK

Input: $\theta, B^{-T}, \bar{\epsilon}, U, \mathcal{J}, \theta^*, L, H$
 1: Draw $\xi \sim B^{-T}\mathcal{N}(0, I), \epsilon \sim \bar{\epsilon} \times \mathcal{U}_{[0.8,1]}$
 2: $\theta', v \leftarrow \theta, \xi$
 3: **for** $i = \{1, \dots, L\}$ **do**
 4: $v \leftarrow v - \frac{\xi}{2}(\mathcal{J}^{-1}\nabla_{\theta}U(\theta') - (\theta' - \theta^*))$
 5: $\theta' \leftarrow (\theta' - \theta^*)$
 6: $\theta', v \leftarrow \theta' \cos(\epsilon) + v \sin(\epsilon), v \cos(\epsilon) - \theta' \sin(\epsilon)$
 7: $\theta' \leftarrow \theta' + \theta^*$
 8: $v \leftarrow v - \frac{\xi}{2}(\mathcal{J}^{-1}\nabla_{\theta}U(\theta') - (\theta' - \theta^*))$
 9: **end for**
 10: $a \leftarrow \min\{1, \exp(H(\theta, \xi) - H(\theta', v))\}$
 11: Draw $\gamma \sim \mathcal{B}(a)$
 12: $\theta \leftarrow \gamma\theta' + (1 - \gamma)\theta$

Of course it is also possible to use the KDK Verlet Algorithm 1 with preconditioning ($M = \mathcal{J}$) (and v replacing p). The resulting algorithm may be seen in Algorithm 4.

Algorithm 4 PrecondVerlet

Input: $\theta, B^{-T}, \bar{\epsilon}, U, \mathcal{J}, \theta^*, L, H$
 1: Draw $\xi \sim B^{-T}\mathcal{N}(0, I), \epsilon \sim \bar{\epsilon} \times \mathcal{U}_{[0.8,1]}$
 2: $\theta', v \leftarrow \theta, \xi$
 3: **for** $i = \{1, \dots, L\}$ **do**
 4: $v \leftarrow v - \frac{\xi}{2}\mathcal{J}^{-1}\nabla_{\theta}U(\theta')$
 5: $\theta' \leftarrow \theta' + \epsilon v$
 6: $v \leftarrow v - \frac{\xi}{2}\mathcal{J}^{-1}\nabla_{\theta}U(\theta')$
 7: **end for**
 8: $a \leftarrow \min\{1, \exp(H(\theta, \xi) - H(\theta', v))\}$
 9: Draw $\gamma \sim \mathcal{B}(a)$
 10: $\theta \leftarrow \gamma\theta' + (1 - \gamma)\theta$

Applying these algorithms to the model problem (11), an analysis parallel to that carried out in Section 3 shows that the decorrelation condition (12) becomes, *independently of* σ_1 and σ_2

$$L \gtrsim C/\epsilon$$

and the stability limits in (13) and (14) are now replaced, *also independently of the values of* σ_1 and σ_2 , by

$$\epsilon \approx 2, \quad \epsilon \approx \pi,$$

for Algorithm 4 and Algorithm 3 respectively. The stability limit for the PrecondRKR algorithm coincides with that of the PrecondKRK method. (See also Appendix B.)

The idea of preconditioning is extremely old; to our best knowledge it goes back to (Bennett 1975). The algorithm in (Girolami and Calderhead 2011) may be regarded as a

² The fact that the frequency is of moderate size is irrelevant; the value of the frequency may be arbitrarily varied by rescaling t . What is important is that all frequencies coincide.

θ -dependent preconditioning. For preconditioning in infinite dimensional problems see (Beskos et al. 2011).

5 Numerical results

In this section we test the following algorithms:

- Unconditioned Verlet: Algorithm 1 with $M = I$.
- Unconditioned KRK: Algorithm 2.
- Preconditioned Verlet: Algorithm 4.
- Preconditioned KRK: Algorithm 3.
- Preconditioned RKR: similar to Algorithm 3 using a rotate-kick-rotate pattern instead of kick-rotate-kick.

The first two algorithms were compared in (Shahbaba et al. 2014) and in fact we shall use the exact same logistic regression test problems used in that reference. If x are the prediction variables and $y \in \{0, 1\}$, the likelihood for the test problems is $(\tilde{x} = [1, x^T]^T, \theta = [\alpha, \beta^T]^T)$

$$\mathcal{L}(\theta; x, y) = \prod_{i=1}^n \left(1 + \exp(-\theta^T \tilde{x}_i)\right)^{-y} \left(1 + \exp(\theta^T \tilde{x}_i)\right)^{y-1}. \tag{16}$$

For the preconditioned integrators, we set U_0 as in (5) with θ^* given by the maximum a posteriori (MAP) estimation and \mathcal{J} the Hessian at θ^* .

For the two unconditioned integrators, we run the values of L and ϵ chosen in (Shahbaba et al. 2014) (this choice is labelled as A in the tables). Since in many cases the autocorrelation for the unconditioned methods is extremely large with those parameter values (see Fig. 1), we also present results for these methods with a principled choice of T and ϵ (labelled as B in the tables). We take $T = \epsilon L = \pi / (2\omega_{\min})$, where ω_{\min} is the minimum eigenvalue of \sqrt{D} given in Eq. (8). In the case where the perturbation U_1 is absent, this choice of T would decorrelate the least constrained component of θ . We then set ϵ as large as possible to ensure an acceptance rate above 65% (Beskos et al. 2013)— the stepsizes in the choice B are slightly smaller than the values used in (Shahbaba et al. 2014), and the durations T are, for every dataset, larger. We are able thus to attain greater decorrelation, although at greater cost. For the preconditioned methods, we set $T = \pi/2$, since this gives samples with 0 correlation in the case $U_1 = 0$, and then set the timestep ϵ as large as possible whilst ensuring the acceptance rate is above 65%.

In every experiment we start the chain from the (numerically calculated) MAP estimate θ^* of θ and acquire $N_s = 5 \times 10^4$ samples. The autocorrelation times reported are calculated using the emcee function `integrated_time`

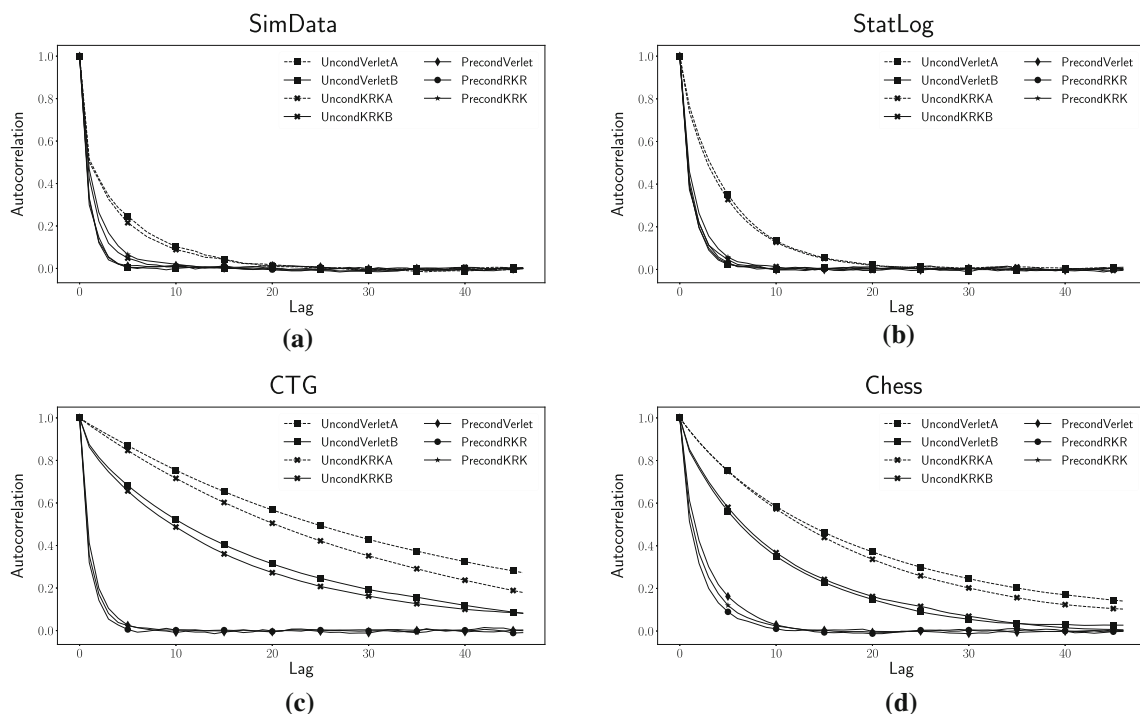


Fig. 1 Autocorrelation function plots for the slowest moving component associated to the IAC τ_{\max} for each dataset. For the unconditioned methods, we show the principled choice B (solid line) and the choice A from (Shahbaba et al. 2014) (dotted). The values of ϵ and T are as given in the tables

with the default value $c = 5$ (Foreman-Mackey et al. 2013). We also estimated autocorrelation times using alternative methods (Geyer 1992; Neal 1993; Sokal 1997; Thompson 2010); the results obtained do not differ significantly from those reported in the tables.

Finally, note that values of $\bar{\epsilon}$ quoted in the tables are the maximum timestep that the algorithms operate with, since the randomisation follows $\epsilon \sim \bar{\epsilon} \times \mathcal{U}_{[0.8,1]}$. All code is available from the github repository <https://github.com/lshaw8317/SplitHMCRvisited>.

5.1 Simulated data

We generate simulated data according to the same procedure and parameter values described in (Shahbaba et al. 2014). The first step is to generate $x \sim \mathcal{N}(0, \sigma^2)$ with $\sigma^2 = \text{diag} \{ \sigma_j^2 : j = 1 \dots, d - 1 \}$, where

$$\sigma_j^2 = \begin{cases} 25 & j \leq 5 \\ 1 & 5 < j \leq 10 \\ 0.04 & j > 10 \end{cases}.$$

Then, we generate the true parameters $\hat{\theta} = [\alpha, \beta^T]^T$ with $\alpha \sim \mathcal{N}(0, \gamma^2)$ and the vector $\beta \in \mathbb{R}^{d-1}$ with independent components following $\beta_j \sim \mathcal{N}(0, \gamma^2)$, $j = 1, \dots, d - 1$, with $\gamma^2 = 1$. Augmenting the data $\tilde{x}_i = [1, x_i^T]^T$, from a given sample x_i , y_i is then generated as a Bernoulli random variable $y_i \sim \mathcal{B}((1 + \exp(-\hat{\theta}^T \tilde{x}_i))^{-1})$. In concreteness, a simulated data set $\{x_i, y_i\}_{i=1}^n$ with $n = 10^4$ samples is generated, $x_i \in \mathbb{R}^{d-1}$ with $d - 1 = 100$. The sampled parameters $\theta \in \mathbb{R}^d$ are assumed to have a prior $\mathcal{N}(0, \Sigma)$ with $\Sigma = \text{diag} \{ 25 : j = 1 \dots, d \}$.

Results are given in Table 1. The second column gives the number L of timesteps per proposal and the third the computational time s (in milliseconds) required to generate a single sample. The next columns give, for three observables, the products $\tau \times s$, with τ the integrated autocorrelation (IAC) time. These products measure the computational time to gen-

erate one independent sample. The notation τ_ℓ refers to the observable $f(\theta) = \log(\mathcal{L}(\theta; x, y))$ where \mathcal{L} is the likelihood in (16), and τ_{θ^2} refers to $f(\theta) = \theta^T \theta$. The degree of correlation measured by τ_ℓ is important in optimising the cost-accuracy ratio of predictions of y , while τ_{θ^2} is relevant to estimating parameters of the distribution of θ (Andrieu et al. 2003; Gelman et al. 2015). Following (Shahbaba et al. 2014), we also examine the maximum IAC over all the Cartesian components of θ , since we set the time T in order to decorrelate the slowest-moving/least constrained component. Finally the last column provides the observed rate of acceptance.

Comparing the values of $\tau \times s$ in the first four rows of the table shows the advantage, emphasized in (Shahbaba et al. 2014), of the $H_0 + U_1$ (4) over the kinetic/potential splitting: Unconditioned KRK operates with smaller values of L than Unconditioned Verlet and the values of $\tau \times s$ are smaller for Unconditioned KRK than for unconditioned Verlet. However when comparing the results for Unconditioned Verlet A or B with those for Preconditioned Verlet, it is apparent that the advantage of using the Hessian \mathcal{J} to split $U = U_0 + U_1$ with $M = I$ is much smaller than the advantage of using \mathcal{J} to precondition the integration while keeping the kinetic/potential splitting.

The best performance is observed for the Preconditioned KRK and RKR algorithms that avail themselves of the Hessian both to precondition and to use rotation instead of drift. Preconditioned RKR is clearly better than its KRK counterpart (see Appendix B). For this problem, as shown in Appendix A, U_1 is in fact small and therefore the restrictions of the stepsize for the KRK integration are due to the stability reasons outlined in Section 3. In fact, for the unconditioned algorithms, the stepsize $\bar{\epsilon}_{KRK} = 0.03$ is not substantially larger than $\bar{\epsilon}_{Verlet} = 0.015$, in agreement with the analysis presented in that section.

The need to use large values of L in the unconditioned integration stems, as discussed above, from the coexistence of large differences between the frequencies of the harmonic

Table 1 SimData: For methods labelled A, parameters from (Shahbaba et al. 2014): $T = 0.3$, $\bar{\epsilon}_{Verlet} = 0.015$, $\bar{\epsilon}_{UKRK} = 0.03$. For the unconditioned methods labelled B, $T = \pi/2\omega_{\min} = 0.6$, and

$\bar{\epsilon}_{Verlet} = 0.015$, $\bar{\epsilon}_{UKRK} = 0.03$. For the preconditioned methods, $T = \pi/2$, and $\bar{\epsilon}_{Verlet} = T/3 \approx 0.52$; the other preconditioned methods operate with $\bar{\epsilon}_{Precon} = T \approx 1.57$

| | L | s [ms] | $\tau_\ell \times s$ | $\tau_{\theta^2} \times s$ | $\tau_{\max} \times s$ | AP |
|----------------|-----|----------|-----------------------|----------------------------|------------------------|------|
| UncondVerlet A | 20 | 4.70 | $3.5 \times s = 16.5$ | $11.4 \times s = 53.6$ | $7.0 \times s = 32.9$ | 0.69 |
| UncondVerlet B | 40 | 8.49 | $3.7 \times s = 31.4$ | $2.6 \times s = 22.1$ | $2.0 \times s = 17.0$ | 0.68 |
| UncondKRK A | 10 | 3.04 | $3.4 \times s = 10.3$ | $11.1 \times s = 33.8$ | $6.6 \times s = 20.1$ | 0.76 |
| UncondKRK B | 20 | 5.27 | $3.9 \times s = 20.5$ | $3.3 \times s = 17.4$ | $3.0 \times s = 15.8$ | 0.69 |
| PrecondVerlet | 3 | 1.60 | $2.5 \times s = 4.0$ | $2.3 \times s = 3.7$ | $2.3 \times s = 3.7$ | 0.79 |
| PrecondKRK | 1 | 1.22 | $2.8 \times s = 3.4$ | $3.4 \times s = 4.2$ | $3.5 \times s = 4.3$ | 0.75 |
| PrecondRKR | 1 | 0.99 | $1.6 \times s = 1.6$ | $2.1 \times s = 2.1$ | $2.1 \times s = 2.1$ | 0.87 |

Table 2 StatLog: For methods labelled A, parameters are from (Shahbaba et al. 2014): $T = 1.6$, $\bar{\epsilon}_{Verlet} = 0.08$, $\bar{\epsilon}_{UKRK} = 0.114$. For the unconditioned methods labelled B, $T = \pi/2\omega_{\min} = 3.26$, and $\bar{\epsilon}_{Verlet} = 0.08$, $\bar{\epsilon}_{UKRK} = 0.114$. For the preconditioned methods, $T = \pi/2$, and $\bar{\epsilon}_{Verlet} = T/3$; the other preconditioned methods operate with $\bar{\epsilon}_{Precon} = T/2$

| | L | s [ms] | $\tau_\ell \times s$ | $\tau_{\theta^2} \times s$ | $\tau_{\max} \times s$ | AP |
|----------------|-----|----------|-----------------------|----------------------------|------------------------|------|
| UncondVerlet A | 20 | 1.99 | $5.5 \times s = 11.0$ | $5.8 \times s = 11.6$ | $9.8 \times s = 19.5$ | 0.69 |
| UncondVerlet B | 40 | 3.34 | $7.6 \times s = 25.4$ | $2.5 \times s = 8.3$ | $2.6 \times s = 8.7$ | 0.64 |
| UncondKRK A | 14 | 1.73 | $6.2 \times s = 10.7$ | $5.7 \times s = 9.9$ | $9.5 \times s = 16.5$ | 0.72 |
| UncondKRK B | 28 | 2.79 | $8.7 \times s = 24.3$ | $2.9 \times s = 8.1$ | $2.9 \times s = 8.1$ | 0.65 |
| PrecondVerlet | 3 | 0.64 | $2.5 \times s = 1.6$ | $2.6 \times s = 1.7$ | $2.7 \times s = 1.7$ | 0.88 |
| PrecondKRK | 2 | 0.60 | $2.9 \times s = 1.7$ | $3.2 \times s = 1.9$ | $3.3 \times s = 2.0$ | 0.88 |
| PrecondRKR | 2 | 0.53 | $2.3 \times s = 1.2$ | $2.5 \times s = 1.3$ | $2.7 \times s = 1.4$ | 0.94 |

Table 3 CTG: For runs labelled A, parameters are from (Shahbaba et al. 2014): $T = 1.6$, $\bar{\epsilon}_{Verlet} = 0.08$, $\bar{\epsilon}_{UKRK} = 0.123$. For the unconditioned runs labelled B, $T = \pi/2\omega_{\min} = 7.85$, and $\bar{\epsilon}_{Verlet} = 0.08$, $\bar{\epsilon}_{UKRK} = 0.118$. For the preconditioned methods, $T = \pi/2$, and $\bar{\epsilon} = T/2$

| | L | s [ms] | $\tau_\ell \times s$ | $\tau_{\theta^2} \times s$ | $\tau_{\max} \times s$ | AP |
|----------------|-----|----------|-----------------------|----------------------------|-------------------------|------|
| UncondVerlet A | 20 | 1.06 | $5.9 \times s = 6.2$ | $20.1 \times s = 21.2$ | $80.3 \times s = 84.8$ | 0.69 |
| UncondVerlet B | 98 | 4.28 | $6.1 \times s = 26.1$ | $5.1 \times s = 21.8$ | $36.0 \times s = 154.2$ | 0.64 |
| UncondKRK A | 13 | 0.95 | $6.5 \times s = 6.2$ | $17.9 \times s = 17.0$ | $53.0 \times s = 50.3$ | 0.77 |
| UncondKRK B | 66 | 3.89 | $6.1 \times s = 23.7$ | $5.1 \times s = 19.8$ | $37.2 \times s = 144.6$ | 0.65 |
| PrecondVerlet | 2 | 0.36 | $2.6 \times s = 0.9$ | $2.1 \times s = 0.7$ | $2.6 \times s = 0.9$ | 0.76 |
| PrecondKRK | 2 | 0.41 | $1.8 \times s = 0.7$ | $1.8 \times s = 0.7$ | $2.4 \times s = 1.0$ | 0.90 |
| PrecondRKR | 2 | 0.35 | $1.9 \times s = 0.7$ | $1.7 \times s = 0.6$ | $2.1 \times s = 0.7$ | 0.93 |

Table 4 Chess: For runs labelled A, parameters are from (Shahbaba et al. 2014): $T = 1.8$, $\bar{\epsilon}_{Verlet} = 0.09$, $\bar{\epsilon}_{UKRK} = 0.2$. For runs labelled B, $T = \pi/2\omega_{\min} = 5.71$, and $\bar{\epsilon}_{Verlet} = 0.087$, $\bar{\epsilon}_{UKRK} = 0.142$. For the preconditioned methods, $T = \pi/2$, and $\bar{\epsilon} = T/2$

| | L | s [ms] | $\tau_\ell \times s$ | $\tau_{\theta^2} \times s$ | $\tau_{\max} \times s$ | AP |
|----------------|-----|----------|------------------------|----------------------------|------------------------|------|
| UncondVerlet A | 20 | 1.52 | $12.2 \times s = 18.5$ | $18.9 \times s = 28.6$ | $42.3 \times s = 64.1$ | 0.62 |
| UncondVerlet B | 65 | 4.20 | $3.6 \times s = 15.1$ | $1.5 \times s = 6.3$ | $19.9 \times s = 83.6$ | 0.68 |
| UncondKRK A | 9 | 0.90 | $13.3 \times s = 11.9$ | $21.3 \times s = 19.1$ | $37.7 \times s = 33.8$ | 0.72 |
| UncondKRK B | 40 | 3.31 | $4.1 \times s = 13.6$ | $1.9 \times s = 6.3$ | $22.1 \times s = 73.1$ | 0.64 |
| PrecondVerlet | 2 | 0.46 | $2.6 \times s = 1.2$ | $3.1 \times s = 1.4$ | $5.2 \times s = 2.4$ | 0.63 |
| PrecondKRK | 2 | 0.50 | $1.6 \times s = 0.8$ | $2.5 \times s = 1.2$ | $4.6 \times s = 2.3$ | 0.81 |
| PrecondRKR | 2 | 0.44 | $1.6 \times s = 0.7$ | $2.2 \times s = 1.0$ | $3.8 \times s = 1.7$ | 0.85 |

oscillators. In this problem the minimum and maximum frequencies are $\omega_{\min} = 2.6$, $\omega_{\max} = 105.0$.

5.2 Real data

The three real datasets considered in (Shahbaba et al. 2014), StatLog, CTG and Chess, are also examined, see Tables 2–4. For the StatLog and CTG datasets with the unconditioned Hamiltonian, KRK does not really provide an improvement on Verlet. In all three datasets, the preconditioned integrators clearly outperform the unconditioned counterparts. Of the three preconditioned algorithms Verlet is the worst and RKR the best.

StatLog

Here, $n = 4435$, $d - 1 = 36$. The frequencies are $\omega_{\min} = 0.5$, $\omega_{\max} = 22.8$.

CTG

Here, $n = 2126$, $d - 1 = 21$. The frequencies are $\omega_{\min} = 0.2$, $\omega_{\max} = 23.9$.

Chess

Here, $n = 3196$, $d - 1 = 36$. The frequencies are $\omega_{\min} = 0.3$, $\omega_{\max} = 22.3$.

Acknowledgements This work has been supported by Ministerio de Ciencia e Innovación (Spain) through project PID2019-104927GB-C21, MCIN/AEI/10.13039/501100011033, ERDF (“A way of making Europe”).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indi-

cate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: Bernstein-von Mises theorem

From the Bernstein-von Mises theorem (see e.g. section 10.2 in (van der Vaart 1998)), as the size of the dataset n increases unboundedly, the posterior distribution $\pi(\theta | x, y)$ becomes dominated by the likelihood and is asymptotically Gaussian; more precisely $\mathcal{N}(\hat{\theta}, n^{-1}\mathcal{I}_F(\hat{\theta})^{-1})$, where $\hat{\theta}$ represents the true value and \mathcal{I}_F denotes the Fisher information matrix. This observation shows that, at least for n large, approximating the potential $\propto \exp(-U(\theta))$ by a Gaussian $\propto \exp(-U_0(\theta))$ with mean θ^* as in Section 5.1 is meaningful.

An illustration of the Bernstein-von Mises theorem is provided in Figure 2 that corresponds to the simulated data problem described in Section 5.1. As the number of data points increases from $2^7 = 128$ to $2^{14} = 32,768$, the scaled values ω_j/\sqrt{n} where ω_j^2 are the eigenvalues of the numerically calculated Hessian $\mathcal{J}(\theta^*)$ that we use in U_0 converge to the square roots of the eigenvalues of the Monte Carlo estimation of the Fisher information matrix $\mathcal{I}_F(\hat{\theta})$ calculated using the true parameter values and the randomly generated x_i .

A further illustration is provided in Fig. 3 where again the number of data points increases from $2^7 = 128$ to $2^{14} = 32,768$. The following parameter values are used:

- The preconditioned algorithms, where solutions of the Hamiltonian H_0 are periodic with period 2π , have $T = \pi/2$. For the KRK and RKR splittings we take two timesteps per proposal, i.e. $L = 2$ and for the preconditioned Verlet, $L = 3$.
- For the unconditioned algorithms we set $T = (\pi/2)/\omega_{\min}$. i.e. a quarter of the largest period present in the solutions of H_0 . Both Verlet and UncondKRK are operated with $L = 30$ timesteps per proposal.

Note that since, as n varies the value of L for each algorithm remains constant, the number of evaluations of ∇U_1 (for methods with rotations) or ∇U (for the Verlet integrator) remains constant. The figure shows that, as n increases, the acceptance rate for the methods Unconditioned KRK, Preconditioned KRK, and Preconditioned RKR based on the splitting (4) approaches 100%. These methods are exact when $U_1 = 0$ and $\exp(-U)$ coincides with the Gaussian $\exp(-U_0)$ and therefore have smaller energy errors/larger

acceptance rates as n increases. On the other hand the integrators based on the kinetic/potential splitting are not exact when the potential U is quadratic, and, correspondingly, we see that the acceptance rate does not approach 100% as $n \uparrow \infty$.

Appendix B: Integrating the preconditioned Hamilton equations

KRK and RKR are two possible reversible, symplectic integrators for the equations of motion corresponding to the preconditioned Hamiltonian (15), but many others are of course possible. In this Appendix we present a methodology to choose between different integrators. The material parallels an approach suggested in (Blanes et al. 2014) to choose between integrators for the kinetic/potential splitting; an approach that has been followed by a number of authors (see (Blanes et al. 2021) for an extensive list of references). The methodology is based on using a Gaussian model distribution to discriminate between alternative algorithms, but, as shown in (Calvo et al. 2021), is very successful in predicting which algorithms will perform well for general distributions.

To study the preconditioned $H_0 + U_1$ splitting, we select the model one-dimensional problem

$$H(\theta, p) = \frac{1}{2}(p^2 + \theta^2) + \frac{1}{2}\kappa\theta^2. \tag{17}$$

We assume that $\kappa > -1$ so that the potential energy $(1/2)\theta^2 + (\kappa/2)\theta^2$ is positive definite. The application of one step (of length ϵ) of an integrator for this problem in all practical contexts takes the linear form

$$\begin{pmatrix} \theta_{n+1} \\ p_{n+1} \end{pmatrix} = M_{\kappa,\epsilon} \begin{pmatrix} \theta_n \\ p_n \end{pmatrix}, \quad M_{\kappa,\epsilon} = \begin{bmatrix} A_{\kappa,\epsilon} & B_{\kappa,\epsilon} \\ C_{\kappa,\epsilon} & D_{\kappa,\epsilon} \end{bmatrix}. \tag{18}$$

From Eq. (18), it is clear that an integration leg of length $T = \epsilon L$ (with initial condition $\theta(0) = \theta_0, p(0) = p_0$) is given by

$$\begin{pmatrix} \theta_L \\ p_L \end{pmatrix} = M_{\kappa,\epsilon}^L \begin{pmatrix} \theta_0 \\ p_0 \end{pmatrix}.$$

We now apply two restrictions to the integration matrix in Eq. (18). Reversibility imposes that $A_{\kappa,\epsilon} = D_{\kappa,\epsilon}$; symplecticity (in one dimension equivalent to volume-preservation) implies that (Blanes et al. 2014)

$$\det(M_{\kappa,\epsilon}) = A_{\kappa,\epsilon}^2 - B_{\kappa,\epsilon}C_{\kappa,\epsilon} = 1. \tag{19}$$

The eigenvalues of the matrix $M_{\kappa,\epsilon}$ are then

$$\lambda = A_{\kappa,\epsilon} \pm \sqrt{A_{\kappa,\epsilon}^2 - 1},$$

Fig. 2 The ordered rotation frequencies ($\omega_j = \sqrt{\lambda_j}$, with λ_j an eigenvalue of \mathcal{J}) scaled by \sqrt{n} as the number of data points n increases converge to the square roots of the eigenvalues of the (estimated) Fisher information matrix

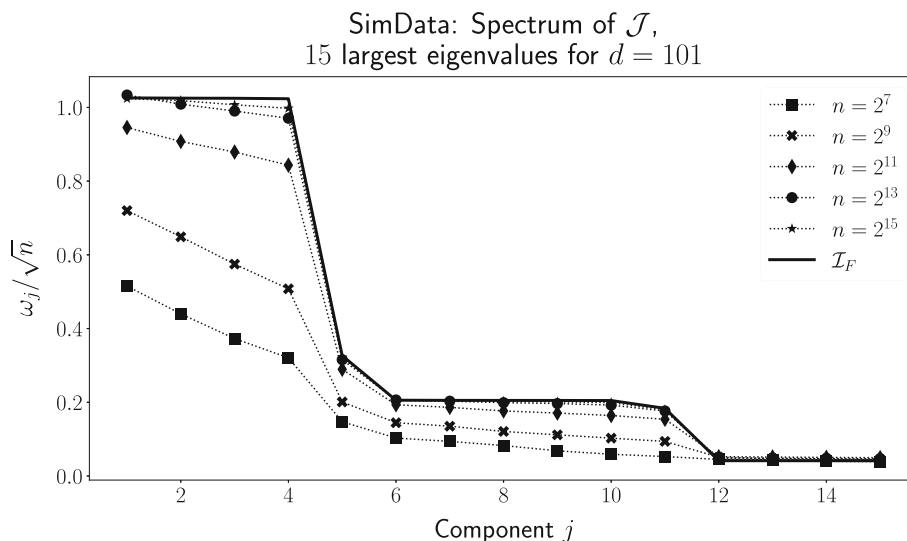



Fig. 3 As n varies the algorithms are run with a fixed number L of timesteps per proposal. For methods using rotations the acceptance rate approaches 100% as $n \uparrow \infty$



which shows that there are three cases:

1. $|A_{\kappa,\epsilon}| > 1$. For one of the eigenvalues, $|\lambda| > 1$ and so the integration is unstable.
2. $|A_{\kappa,\epsilon}| < 1$. The integration is stable as both eigenvalues have magnitude 1.
3. $|A_{\kappa,\epsilon}| = 1$. The symplectic condition Eq. (19) necessarily implies $B_{\kappa,\epsilon}C_{\kappa,\epsilon} = 0$, which gives two sub-cases:
 - (a) $|B_{\kappa,\epsilon}| + |C_{\kappa,\epsilon}| = 0$. The matrix $M_{\kappa,\epsilon} = \pm I$ and the integration is stable.
 - (b) $|B_{\kappa,\epsilon}| + |C_{\kappa,\epsilon}| \neq 0$. Then, if $B_{\kappa,\epsilon} \neq 0$,

$$M_{\kappa,\epsilon}^L = \begin{bmatrix} A^L & LA^{L-1}B \\ 0 & A^L \end{bmatrix}$$

and the integration is (weakly) unstable. Similarly there is weak instability if instead $C_{\kappa,\epsilon} \neq 0$

Thus for stable integration, one may find $\eta_{\kappa,\epsilon}$ such that $A_{\kappa,\epsilon} = \cos(\eta_{\kappa,\epsilon}) \in [-1, 1]$; in addition we define $\chi_{\kappa,\epsilon} = B_{\kappa,\epsilon} / \sin(\eta_{\kappa,\epsilon})$ for $\sin(\eta_{\kappa,\epsilon}) \neq 0$ and let $\chi_{\kappa,\epsilon}$ be arbitrary if $\sin(\eta_{\kappa,\epsilon}) = 0$. In this way, for the model problem, all stable, symplectic integrations have a propagation matrix of the form

$$M_{\kappa,\epsilon} = \begin{bmatrix} \cos(\eta_{\kappa,\epsilon}) & \chi_{\kappa,\epsilon} \sin(\eta_{\kappa,\epsilon}) \\ -\chi_{\kappa,\epsilon}^{-1} \sin(\eta_{\kappa,\epsilon}) & \cos(\eta_{\kappa,\epsilon}) \end{bmatrix}. \tag{20}$$

We now state a lemma analogue of Proposition 4.3 in (Blanes et al. 2014).

Lemma 1 Denote $A = \cos(L\eta_{\kappa,\epsilon})$, $B = \chi_{\kappa,\epsilon} \sin(L\eta_{\kappa,\epsilon})$, $C = -\chi_{\kappa,\epsilon}^{-1} \sin(L\eta_{\kappa,\epsilon})$. Given the initial conditions θ_0, p_0 , and integrating the dynamics of the Hamiltonian of the model problem Eq. (17) using the integrator in Eq. (20) for L steps to give new values of θ_L, p_L , the energy error may be expressed as:

$$\Delta \equiv H(\theta_L, p_L) - H(\theta_0, p_0)$$

$$= \frac{1}{2} (C + (1 + \kappa)B) (C\theta_0^2 + 2A\theta_0 p_0 + Bp_0^2). \tag{21}$$

Proof Applying the symplectic condition Eq. (19), the energy error $\Delta \equiv H(\theta_L, p_L) - H(\theta_0, p_0)$ then follows

$$\begin{aligned} 2\Delta &= p_L^2 + (1 + \kappa)\theta_L^2 - p_0^2 - (1 + \kappa)\theta_0^2 \\ &= (C\theta_0 + Ap_0)^2 + (1 + \kappa)(Bp_0 + A\theta_0)^2 \\ &\quad - p_0^2 - (1 + \kappa)\theta_0^2 \\ &= (C^2 + (A^2 - 1)(1 + \kappa))\theta_0^2 \\ &\quad + 2A(C + (1 + \kappa)B)\theta_0 p_0 + (A^2 - 1 + (1 + \kappa)B^2)p_0^2 \\ &= (C + (1 + \kappa)B) (C\theta_0^2 + 2A\theta_0 p_0 + Bp_0^2). \end{aligned}$$

□

Theorem 1 *With the notation of the lemma, assume that the initial conditions $\theta_0 \sim \mathcal{N}(0, 1/(1 + \kappa))$, $p_0 \sim \mathcal{N}(0, 1)$ are (independently) distributed according to their stationary distributions corresponding to the Hamiltonian for the model problem Eq. (17). Then the expected energy error follows*

$$\mathbb{E}[\Delta] = \sin^2(L\eta_{\kappa,\epsilon})\rho(\epsilon, \kappa) \leq \rho(\epsilon, \kappa),$$

where ρ is given by

$$\begin{aligned} \rho(\epsilon, \kappa) &= \frac{1}{2} \left(\sqrt{1 + \kappa} \chi_{\kappa,\epsilon} - \frac{1}{\sqrt{1 + \kappa} \chi_{\kappa,\epsilon}} \right)^2 \\ &= \frac{(C_{\kappa,\epsilon} + (1 + \kappa)B_{\kappa,\epsilon})^2}{2(1 + \kappa)(1 - A_{\kappa,\epsilon}^2)}. \end{aligned} \tag{22}$$

Proof Since $\mathbb{E}[\theta_0 p_0] = \mathbb{E}[\theta_0]\mathbb{E}[p_0] = 0$ and $\mathbb{E}[p_0^2] = 1$, $\mathbb{E}[\theta_0^2] = 1/(1 + \kappa)$, the expectation of Eq. (21) is

$$\begin{aligned} \mathbb{E}[\Delta] &= \frac{1}{2} (C + (1 + \kappa)B) \left(\frac{C}{1 + \kappa} + B \right) \\ &= \frac{1}{2} \left(\frac{C}{\sqrt{1 + \kappa}} + \sqrt{1 + \kappa}B \right)^2. \end{aligned}$$

Substituting the expressions for B, C from the definitions in the theorem above into the last display and dropping the subscripts to give $\chi = \chi_{\kappa,\epsilon}$, $s = \sin(L\eta_{\kappa,\epsilon})$ and $c = \cos(L\eta_{\kappa,\epsilon})$ gives

$$\begin{aligned} \mathbb{E}[\Delta] &= \frac{1}{2} s^2 \left(\frac{1}{(1 + \kappa)\chi^2} + (1 + \kappa)\chi^2 - 2 \right) \\ &= \sin^2(L\eta_{\kappa,\epsilon})\rho(\epsilon, \kappa). \end{aligned}$$

□

Since $\eta_{\kappa,\epsilon}$ and $\chi_{\kappa,\epsilon}$ depend on the integrator Eq. (20), the ρ function also depends on the integrator. Note that ρ does not change with L . For the model problem, integrators with smaller ρ lead to smaller averaged energy errors at stationarity of the chain and therefore to smaller empirical rejection rates. By diagonalization it is easily shown as in (Blanes et al. 2014) that the same is true for all Gaussian targets $\propto \exp(-U(\theta))$, $U = U_0 + U_1$. This suggests that, all other things being equal, integrators with smaller ρ should be preferred (see a full discussion in (Calvo et al. 2021)).

KRK vs. RKR

For the KRK integration, we find (similarly to Section 3) that a stable integration requires $-1 < \cos(\epsilon) - \epsilon\kappa \sin(\epsilon)/2 < 1$, so that, the stability limit for $\kappa > 0$ is

$$\epsilon < \frac{2 \cot(\epsilon/2)}{\kappa}. \tag{23}$$

Note that $\epsilon < \pi$ for any value of $\kappa > 0$. For $-1 < \kappa < 0$, the stability limit is $\epsilon < \pi$. Application of the formula Eq. (22) gives the ρ function of the integrator as

$$\rho^{[KRK]}(\epsilon, \kappa) = \frac{\kappa^2 \csc(\epsilon)(-4\epsilon \cos(\epsilon) + (4 + \kappa\epsilon^2) \sin(\epsilon))^2}{8(1 + \kappa)(4\kappa\epsilon \cos(\epsilon) + (4 - \kappa^2\epsilon^2) \sin(\epsilon))}.$$

For $\kappa = 0$, ρ vanishes as expected because then the integration is exact.

Similarly, for RKR, stable integration requires $-1 < \cos(\epsilon) - \epsilon\kappa \sin(\epsilon)/2 < 1$, so that, the stability limit for $\kappa > 0$ of RKR is the same we found in (23). Again for $-1 < \kappa < 0$, the stability limit is $\epsilon < \pi$, as for KRK.

Application of the formula Eq. (22) gives the ρ function of the RKR integrator as

$$\rho^{[RKR]}(\epsilon, \kappa) = \frac{\kappa^2 \csc(\epsilon)(\kappa\epsilon \cos(\epsilon) + 2 \sin(\epsilon) - (2 + \kappa)\epsilon)^2}{2(1 + \kappa)(4\kappa\epsilon \cos(\epsilon) + (4 - \kappa^2\epsilon^2) \sin(\epsilon))}.$$

The following result implies that for all Gaussian problems, at stationarity, RKR always leads to smaller energy errors/higher acceptance rates than KRK.

Theorem 2 *For each choice of $\kappa > -1$, $\kappa \neq 0$, and $\epsilon > 0$ leading to a stable KRK or RKR integration*

$$\rho^{[RKR]}(\epsilon, \kappa) < \rho^{[KRK]}(\epsilon, \kappa).$$

Proof From the expressions for ρ given above, we have to show that

$$\begin{aligned} &4(\kappa\epsilon \cos(\epsilon) + 2 \sin(\epsilon) - (2 + \kappa)\epsilon)^2 \\ &< (-4\epsilon \cos(\epsilon) + (4 + \kappa\epsilon^2) \sin(\epsilon))^2. \end{aligned}$$

□

It is therefore sufficient to show that

$$2(\kappa\epsilon \cos(\epsilon) + 2 \sin(\epsilon) - (2 + \kappa)\epsilon) < -4\epsilon \cos(\epsilon) + (4 + \kappa\epsilon^2) \sin(\epsilon) \tag{24}$$

and

$$4\epsilon \cos(\epsilon) - (4 + \kappa\epsilon^2) \sin(\epsilon) < 2(\kappa\epsilon \cos(\epsilon) + 2 \sin(\epsilon) - (2 + \kappa)\epsilon). \tag{25}$$

The inequality (24) may be rearranged as

$$2\kappa\epsilon(\cos(\epsilon) - 1) < 4\epsilon(1 - \cos(\epsilon)) + \kappa\epsilon^2 \sin(\epsilon),$$

or

$$-4\kappa\epsilon \sin^2(\epsilon/2) < 8\epsilon \sin^2(\epsilon/2) + 2\kappa\epsilon^2 \sin(\epsilon/2) \cos(\epsilon/2).$$

Since for stable runs $\epsilon < \pi$, so that $\sin^2(\epsilon/2) > 0$, the last display is equivalent to

$$-4\kappa < 8 + 2\kappa\epsilon \cot(\epsilon/2)$$

or

$$-4 < \kappa(\epsilon \cot(\epsilon/2) + 2).$$

For $0 < \epsilon < \pi$, $\epsilon \cot(\epsilon/2) + 2$ takes values between 4 and 2 and therefore the last inequality certainly holds for each $\kappa > -1$.

The inequality (25) may be similarly rearranged as

$$\kappa\epsilon(2 - \epsilon \cot(\epsilon/2)) < 4(2 - \epsilon \cot(\epsilon/2)) \cot(\epsilon/2).$$

Since, for $0 < \epsilon < \pi$, $2 - \epsilon \cot(\epsilon/2) > 0$, we conclude that (25) is equivalent to

$$\kappa\epsilon < 4 \cot(\epsilon/2),$$

a relation that, according to (23), holds for stable integrations with $\kappa > 0$ and is trivially satisfied for $\kappa < 0$, $\epsilon \in (0, \pi)$. \square

Multistage splittings

In addition to the Strang formula (3) one may consider more sophisticated schemes

$$\psi_\epsilon = \varphi_{a_1\epsilon}^{[H_1]} \circ \varphi_{b_1\epsilon}^{[H_2]} \circ \varphi_{a_2\epsilon}^{[H_1]} \circ \dots \circ \varphi_{a_{m-1}\epsilon}^{[H_1]} \circ \varphi_{b_{m-1}\epsilon}^{[H_2]} \circ \varphi_{a_m\epsilon}^{[H_1]} \tag{26}$$

where $\sum a_j = \sum b_j = 1$. These integrators are always symplectic and in addition are time reversible if they are palindromic i.e. $a_i = a_{m-i+1}$, $i = 1, \dots, m$, $b_i = b_{m-i}$, $i = 1, \dots, m - 1$. In the case of the kinetic/potential splitting of

H , integrators of the form (26) when used for HMC sampling may provide very large improvements on leapfrog/Verlet (see (Calvo et al. 2021; Blanes et al. 2021) and their references). For the preconditioned $H_0 + U_1$ splitting in this paper, we have investigated extensively the existence of formulas of the format (26) that improve on the RKR integrator based on the Strang recipe (3). We proceeded in a way parallel to that followed in (Blanes et al. 2014). For fixed m , $m = 3$ or $m = 4$, and a suitable range of values of κ and ϵ , we choose the values of a_i and b_i so as to minimize the function ρ in Theorem 1, thus minimizing the expected energy error at stationarity in the integration of the model problem. The outcome of our investigation was that, while we succeeded in finding formulas that improve on the Preconditioned RKR integrator, the improvements were minor and did not warrant the replacement of RKR by more sophisticated formulas.

References

Andrieu, C., De Freitas, N., Doucet, A., Jordan, M.I.: An Introduction to MCMC for Machine Learning. *Mach. Learn.* **50**(1), 5–43 (2003)

Bennett, C.H.: Mass Tensor Molecular Dynamics. *J. Comput. Phys.* **19**(3), 267–279 (1975)

Beskos, A., Pillai, N., Roberts, G., Sanz-Serna, J.M., Stuart, A.: Optimal Tuning of the Hybrid Monte Carlo Algorithm. *Bernoulli* **19**(5A), 1501–1534 (2013)

Beskos, A., Pinski, F.J., Sanz-Serna, J.M., Stuart, A.M.: Hybrid Monte Carlo on Hilbert spaces. *Stoch. Process. Appl.* **121**(10), 2201–2230 (2011)

Blanes, S., Calvo, M.P., Casas, F., Sanz-Serna, J.M.: Symmetrically Processed Splitting Integrators for Enhanced Hamiltonian Monte Carlo sampling. *SIAM J. Sci. Comput.* **43**(5), A3357–A3371 (2021)

Blanes, S., Casas, F.: *A concise introduction to geometric numerical integration*. CRC Press, Boca Raton, Florida (2017)

Blanes, S., Casas, F., Sanz-Serna, J.M.: Numerical Integrators for the Hybrid Monte Carlo Method. *SIAM J. Sci. Comput.* **36**(4), A1556–A1580 (2014)

Bou-Rabee, N., Sanz-Serna, J.M.: Randomized Hamiltonian Monte Carlo. *Ann. Appl. Probab.* **27**(4), 2159–2194 (2017)

Bou-Rabee, N., Sanz-Serna, J.M.: Geometric Integrators and the Hamiltonian Monte Carlo Method. *Acta Numer.* **27**, 113–206 (2018)

Calvo, M.P., Sanz-Alonso, D., Sanz-Serna, J.M.: HMC: Reducing the Number of Rejections by not Using Leapfrog and some Results on the Acceptance Rate. *J. Comput. Phys.* **437**, 110333 (2021)

Foreman-Mackey, D., Hogg, D.W., Lang, D., Goodman, J.: emcee: the MCMC Hammer. *Publ. Astron. Soc. Pac.* **125**(925), 306–312 (2013). <https://doi.org/10.1086/670067>

García-Archilla, B., Sanz-Serna, J.M., Skeel, R.D.: Long-time-step Methods for Oscillatory Differential Equations. *SIAM J. Sci. Comput.* **20**(3), 930–963 (1998)

Gelman, A., Carlin, J.B., Stern, H.S., and Rubin, D.B.: *Bayesian Data Analysis* (3rd ed.). Chapman and Hall/CRC (2015)

Geyer, C.J.: Practical Markov Chain Monte Carlo. *Stat. Sci.* **7**(4), 473–483 (1992)

Girolami, M., Calderhead, B.: Riemann Manifold Langevin and Hamiltonian Monte Carlo Methods. *J. R. Stat. Soc.: Ser. B (Statistical Methodology)* **73**(2), 123–214 (2011)

- Grubmüller, H., Heller, H., Windemuth, A., Schulten, K.: Generalized Verlet Algorithm for Efficient Molecular Dynamics Simulations with Long-range Interactions. *Mol. Simul.* **6**(1–3), 121–142 (1991)
- Hoffman, M.D., Gelman, A.: The No-U-Turn sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* **15**(1), 1593–1623 (2014)
- Leimkuhler, B., Matthews, C.: *Molecular Dynamics*. Springer International Publishing, Cham (2015)
- Neal, R.M.: Probabilistic Inference Using Markov Chain Monte Carlo Methods. Department of Computer Science. University of Toronto Toronto, ON, Canada (1993)
- Neal, R.M.: MCMC Using Hamiltonian Dynamics, In *Handbook of Markov Chain Monte Carlo*, eds. Brooks, S., Gelman, A., Jones, G.L., and Meng, X.L., 139–188. Chapman and Hall/CRC (2011)
- Sanz-Serna, J.M.: Markov Chain Monte Carlo and Numerical Differential Equations. In: Dieci, L., Guglielmi, N. (eds.) *Current Challenges in Stability Issues for Numerical Differential Equations*, pp. 39–88. Springer International Publishing, Cham (2014)
- Sanz-Serna, J.M., Calvo, M.P.: *Numerical Hamiltonian Problems*. Chapman and Hall, London (1994)
- Shahbaba, B., Lan, S., Johnson, W.O., Neal, R.M.: Split Hamiltonian Monte Carlo. *Stat. Comput.* **24**(3), 339–349 (2014)
- Sokal, A.: Monte Carlo Methods in Statistical Mechanics: Foundations and New Algorithms. In: DeWitt-Morette, C., Cartier, P., Folacci, A. (eds.) *Functional Integration: Basics and Applications*, pp. 131–192. Springer, Boston, MA (1997)
- Thompson, M.B.: A Comparison of Methods for Computing Autocorrelation Time. arXiv preprint [arXiv:1011.0175](https://arxiv.org/abs/1011.0175) (2010)
- Tuckerman, M., Berne, B.J., Martyna, G.J.: Reversible Multiple Time Scale Molecular Dynamics. *J. Chem. Phys.* **97**(3), 1990–2001 (1992)
- van der Vaart, A.: *Asymptotic Statistics*. Cambridge University Press, Cambridge, UK (1998)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.