On the unconventional Hug integrator

Christophe Andrieu [*] and J.M. Sanz-Serna[†]

February 17, 2025

Abstract

Hug is a recently proposed iterative mapping used to design efficient updates in Markov chain Monte Carlo (MCMC) methods when sampling along manifolds is of interest. In this paper we show that Hug may be interpreted as a consistent discretization of a system of differential equations with a rather complicated structure. The proof of convergence of this discretization includes a number of unusual features we explore fully. We uncover an unexpected and, yet, undocumented property of the solutions of the underlying dynamical system that manifest itself by the existence of Hug trajectories that fail to cover the manifold of interest. This suggests caution when using the Hug update.

1 Introduction

Dynamical systems, and their integrators, play an important rôle in the design of modern and efficient Markov chain Monte Carlo (MCMC) algorithms that aim to sample from complex probability distributions [5]. Their ability to turn local information about the target density of interest, available in the form of gradients or higher-order derivatives, into efficient mappings adapted to the geometry of the problem at hand is the reason for this importance. HMC (Hybrid Monte Carlo) [3] is the best known illustration of this idea: here Hamilton's equations are used to follow the contours of an extended probability density of which the density of interest is a marginal. Together with volume preservation, this property explains the interest of HMC. Recently an alternative to HMC was proposed in [11] with the attractive feature that an update involves moving along contours of the probability density of interest itself, which should be contrasted with the corresponding, sometimes unintuitive, trajectories arising from solutions of Hamilton's equations.

The Hug update of [11] is based on the repeated composition of a "bounce" mapping (proposed by [12] and later developed in [2]). The bouncing mechanism makes it possible for the proposals generated by Hug to remain very close (hug) to the manifold of interest, which is remarkable since the algorithm is fully explicit, i.e. it does not resort to the numerical solution of the algebraic equations that define the manifold. In this

---

[*]School of Mathematics, University of Bristol, U.K.

[†]Departamento de Matemáticas, Universidad Carlos III de Madrid, Avenida de la Universidad 30, E-28911 Leganés (Madrid), Spain. E-mail: jmsanzserna@gmail.com

manuscript we show that Hug may be viewed as a consistent discretization of an underpinning dynamical system that we identify. Such dynamical systems provide an important step towards performance analysis of MCMC algorithms in high-dimensions, see [1] for example. By analysing the dynamical system, we have uncovered a phenomenon whereby Hug trajectories may oscillate in small neighbourhoods around the initial state of the system, depending on properties of the initial velocity. We investigate this phenomenon in detail on a simple model target distribution and provide a small numerical study indicating that conclusions of our precise analysis extend to more general scenarios. This suggests caution when using the Hug update.

The contents of this manuscript is as follows. Section 2 focuses on discrete-time algorithms. In Subsection 2.1 we present the Hug update in its original MCMC context where the aim is to maintain a single constraint while updating the state of the Markov chain. In Subsection 2.2 we present a generalisation of Hug's discrete dynamics [4] suitable for the exploration of manifolds resulting from multiple constraints. A number of properties of the algorithm are presented in Subsection 2.3. In Section 3 we identify the underpinning continuous-time dynamical system and explain the way in which Hug provides a convergent discretization of that system. In Subsection 3.1 we review results on linear projector differentials. Then, in Subsection 3.2 we present the continuous-time dynamical system and establish a number of properties that mirror those of the discrete time algorithm studied previously in Subsection 2.3. The convergence of Hug as a discretization of the continuous-time dynamical system is elucidated in Subsection 3.3. The analysis presents a number of nonstandard features; of particular interest from the numerical analysis point of view is that a *supraconvergence* phenomenon [6] takes place, whereby second-order convergence is obtained with first order consistency. In Subsection 3.4 we fully analyse a model system, showing that, unexpectedly, the trajectories generated by both the differential equations and Hug may fold back on themselves, thus precluding a good exploration of the manifold of interest. This unwelcome phenomenon depends on the initial condition and we obtain a full characterisation of the initial states for which the phenomenon occurs. In Section 4 we present simple numerical experiments suggesting that the phenomenon analysed in the model problem extends to more complex scenarios. Technical proofs are given in Section 5; these make extensive use the results by Golub and Pereyra [7] on differentiation of projections.

2   Hug algorithm

2.1   The algorithm

In this subsection we describe the Hug algorithm suggested in [11]. In the description we will make use of some properties that will be proved in Proposition 1 below.

Hug is a Markov Chain Monte Carlo method to obtain samples from a target unnormalized probability distribution π in $\mathbb{R}^n$, assumed to have a density $\exp(\ell(x))$ with respect to the Lebesgue measure. Hug introduces an auxiliary variable $v \in \mathbb{R}^n$ with a density $q(v|x)$ such that

$$q(v|x) = q(-v|x) \tag{1}$$

Algorithm 1 Hug

Require: Number of timesteps K; stepsize $\delta > 0$; current state x of the Markov chain; v-marginal density $q(\cdot|x)$.

      Initialize $k = 0$, $x_0 \leftarrow x$, draw velocity $v_0 \sim q(\cdot|, x_0)$.
      Timestepping:
  for $k = 0, \ldots, K - 1$, **do**
      Move to $x_{k+1/2} = x_k + (\delta/2)v_k$.
      Reflect: $v_{k+1} = R(x_{k+1/2})v_k$.
      Move to $x_{k+1} = x_{k+1/2} + (\delta/2)v_{k+1}$.
  end for
      Compute $\log(r) = \ell(x_K) - \ell(x_0) + \log q(v_K|x_K) - \log q(v_0|x_0)$.
      With a probability $\alpha = 1 \wedge r$, $x \leftarrow x_K$; else $x \leftarrow x$.
---

and generates a Markov Chain with invariant distribution $\pi(x)q(v|x)$. The marginal on x of this Markov chain has therefore $\pi(x)$ as an invariant distribution; one step of this marginal chain is described in Algorithm 1.

A salient feature of the algorithm is the use of velocity reflections. For $x \in \mathbb{R}^n$ with $\nabla\ell(x) \neq 0$, the $n \times n$ matrix $R(x)$ used in the reflections is

$$R(x) = I - \frac{2}{\|\nabla\ell(x)\|^2}\nabla\ell(x)\nabla\ell(x)^T, \tag{2}$$

so that, if w is a vector, $R(x)w$ is the result of reflecting w on the vector subspace of $\mathbb{R}^n$ tangent at x to the level hypersurface $\{y : \ell(y) = \ell(x)\}$. Here and elsewhere $\|\cdot\|$ denotes the standard Euclidean norm.

Figure 1 illustrates the computation of the iterate x_2 when $\delta = 0.1$, $x_0 = [1, 0]^T$, $v_0 = [1, 2]^T$ and the target is a bivariate Gaussian distribution with $\ell = -x_{(1)}^2 - 4x_{(2)}^2$ ($x_{(1)}$ and $x_{(2)}$ denote the scalar components of x). The figure suggests that the iterates x_k may be expected to be close to (or "hug") the level set $\{y : \ell(y) = \ell(x_0)\}$; this explains the name of the algorithm. The halfway points $x_{k+1/2}$ will not be too close to $\{y : \ell(y) = \ell(x_0)\}$, but this does not matter as they are just intermediate auxiliary values (like the internal stages of a Runge-Kutta integrator).

The typical choice of $q(v|x)$ satisfying (1) is given by an x-independent isotropic normal $v \sim \mathcal{N}(0, \sigma^2 I)$. Since the reflection matrix (2) is orthogonal, all iterates v_0, $\ldots$, v_K share a common Euclidean length, and therefore, for this simple choice of q, in the expression for the acceptance probability r the term

$$\log q(v_K|x_K) - \log q(v_0|x_0) = -\sigma^{-2}\left(\|v_K\|^2 - \|v_0\|^2\right)$$

vanishes. Furthermore, due to the "hugging" property, $\ell(x_K) - \ell(x_0)$ will typically be small so that the value of the acceptance probability r will be close to 1. At the same time, it could be expected that, once δ has been chosen, by taking the number K of timesteps sufficiently high, x_K will be far away from x_0. In conclusion, Hug may offer the possibility of generating proposals that are away from the current state of the chain and, at the same time, have high probability of being accepted.
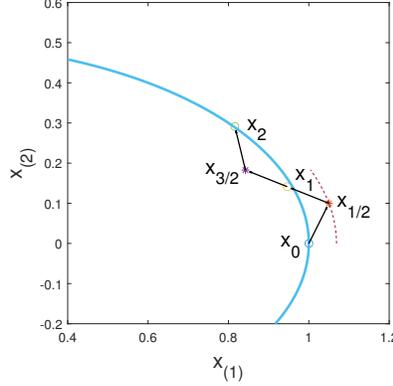
Figure 1: Computation of the iterate x_2 for a bivariate Gaussian. The plot depicts the plane of the variable $x = (x_{(1)}, x_{(2)}) \in \mathbb{R}^2$. The solid curve is the contour $\{y : \ell(y) = \ell(x_0)\}$. The dashed curve is (part of) the contour $\{y : \ell(y) = \ell(x_{1/2})\}$. The segments $[x_0, x_{1/2}], [x_{1/2}, x_1], [x_1, x_{3/2}], [x_{3/2}, x_2]$ share a common Euclidean length $(\delta/2)\|v_0\|$.

The x-moves in Hug will not change much the value of ℓ and in order to efficiently explore the whole state space, the reference [11] suggested to interleave steps of Hug with steps of a second Markov kernel, called Hop, that causes the state of the Markov chain to jump between different level sets of ℓ.

2.2   A generalization of the Hug timestepping mechanism

In what remains of the paper, we will consider that in the Hug time-stepping formulas ($x_0 \in \mathbb{R}^n$, $v_0 \in \mathbb{R}^n$)

$$
\begin{align}
x_{k+1/2} &= x_k + (\delta/2)v_k \tag{3} \\
v_{k+1} &= R(x_{k+1/2})v_k, \tag{4} \\
x_{k+1} &= x_{k+1/2} + (\delta/2)v_{k+1}, \tag{5}
\end{align}
$$

the reflection in (4) is taken with respect to the vector subspace of $\mathbb{R}^n$ tangent at $x_{k+1/2}$ to the level manifold $\{y : f(y) = f(x_{k+1/2})\}$ of a smooth function $f : \mathbb{R}^n \to \mathbb{R}^m$, $0 < m < n$. It is expected that now the iterates x_k will hug the level set $\{y : f(y) = f(x_0)\}$. To shorten the notation we shall hereafter set $\mathcal{M}(\eta) = \{y : f(y) = \eta\}$ ($\eta \in \mathbb{R}^m$). In the particular case of Algorithm 1, $m = 1$ and $f = \ell$. The present generalization has been considered in the thesis [4], where it is used to build algorithms to sample from *filamentary distributions*, i.e. distributions in $\mathbb{R}^n$ supported in the neighbourhood of a set $\mathcal{M}(\eta)$ for some fixed $\eta \in \mathbb{R}^m$.

We now present the expression for $R(x)$ that replaces (2) in the generalized scenario. For each $x \in \mathbb{R}^n$, $J(x)$ will denote the $m \times n$ Jacobian matrix of f evaluated at x. In order not to clutter the exposition with unwelcome details, we will assume

4

throughout that, for each x, $J(x)$ has full rank. (However this hypothesis may easily be weakened: for instance one might demand that $J(x)$ be full rank except at some isolated points or, more generally, that $J(x)$ be full rank for x in an open subdomain of $\mathbb{R}^n$.) The implicit function theorem shows that, for each x, $\mathcal{M}(f(x))$ is a smooth $n - m$ submanifold of $\mathbb{R}^n$. The vector space *tangent* at x to this manifold is given by

$$\mathcal{T}(x) = \{v \in \mathbb{R}^n : J(x)v = 0\}. \tag{6}$$

We shall denote by $\mathcal{N}(x) \subset \mathbb{R}^n$ the corresponding *normal* space, i.e. the m-dimensional vector subspace of $\mathbb{R}^n$ orthogonal to $\mathcal{T}$. For each x, $J(x)J(x)^T$ is invertible and

$$J(x)^+ = J(x)^T(J(x)J(x)^T)^{-1} \tag{7}$$

is the Moore-Penrose inverse of $J(x)$ (size $n \times m$). The $n \times n$ symmetric matrix

$$N(x) = J(x)^+ J(x) = J(x)^T(J(x)J(x)^T)^{-1}J(x) \tag{8}$$

represents the orthogonal projection of $\mathbb{R}^n$ onto $\mathcal{N}(x)$ and then

$$T(x) = I - N(x) \tag{9}$$

is the complementary projection onto $\mathcal{T}(x)$. In what follows we will use repeatedly the fact that, for any given x, each vector w may be decomposed as $T(x)w + N(x)w$ into a tangential component and a normal component.

After these preparations, we see that the expression for $R(x)$ to be used in (4) is $R(x) = I - 2N(x)$. For given x and w, there are different ways of computing $N(x)w$. One that suggests itself uses the decomposition $J(x)^T = QR$, where Q is $n \times m$ with orthonormal columns and R is $m \times m$ upper triangular with positive diagonal entries. The columns of Q are an orthonormal basis of $\mathcal{N}(x)$ and therefore $N(x)w = QQ^T w$. In the particular case $m = 1$, $f = \ell$, this amounts to finding the normalized gradient $(1/\|\nabla\ell(x)\|)\nabla\ell(x)$ as in (2).

2.3 Properties of the timestepping mechanism

Proposition 1 below summarizes some properties of (3)–(5). It requires some notation.

We denote by $\Psi_\delta : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n \times \mathbb{R}^n$ the map that carries out a single step of (3)–(5), i.e. $\Psi_\delta(x, v) = (x', v')$ if $v' = R(x + (\delta/2)v)v$, $x' = x + (\delta/2)v + (\delta/2)v'$.

The symbol $H(x)$ will denote the second derivative of f evaluated at x; this is the symmetric bilinear map $\mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^m$ such that

$$f(x + w) = f(x) + J(x)w + \frac{1}{2}H(x)[w, w] + o(\|w\|^2), \qquad w \to 0.$$

For a symmetric bilinear map $S : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^m$, the operator norm is defined as

$$\|S\| = \sup\{\|S[w_1, w_2]\| : \|w_1\| \le 1, \|w_2\| \le 1\}.$$

With this notation in place we have:

Proposition 1. *The following properties hold:*

1. *The map Ψ_δ is volume preserving.*

2. *The map Ψ_δ is time-reversible, i.e. for each x and v, $\Psi_\delta(x,v) = (x',v')$ if and only if $\Psi_\delta(x',-v') = (x,-v)$.*

Furthemore, let (3)–(5) hold for $k = 0,\ldots,K-1$. Then:

3. *For each $k = 0,\ldots,K-1$,*

$$T(x_{k+1/2})v_{k+1} = T(x_{k+1/2})v_k, \quad N(x_{k+1/2})v_{k+1} = -N(x_{k+1/2})v_k$$

   and

$$v_k + v_{k+1} = 2T(x_{k+1/2})v_k.$$

4. *The vectors $v_0,\ldots,v_K$ share a common Euclidean length $\|v_0\|$.*

5. *The vectors $x_{k+1/2}-x_k$, $k = 0,\ldots,K-1$, and $x_{k+1}-x_{k+1/2}$, $k = 0,\ldots,K-1$, share a common Euclidean length $(\delta/2)\|v_0\|$. The vectors $x_{k+3/2}-x_{k+1/2}$, $k = 0,\ldots,K-2$ share a common Euclidean length $\delta\|v_0\|$.*

6. *If, as x ranges in $\mathbb{R}^n$, $\|H(x)\|$ is upper bounded by β and $H(x)$ is γ-Lipschitz continuous, i.e. for each $x,y \in \mathbb{R}^n$,*

$$\|H(x) - H(y)\| \leq \gamma\|x-y\|,$$

   we have

$$\left\|f(x_K) - f(x_0)\right\| \leq \frac{\delta^2}{12}\|v_0\|^2\big(3\beta + \gamma(K-1)\delta\|v_0\|\big).$$

7. *If $H(x)$ is of the form $H(x)[w,w] = \|w\|^2\nu$, $\nu \in \mathbb{R}^m$, then $f(x_k) = f(x_0)$, $k = 0,\ldots,K$.*

For the particular case $f = \ell$, Properties 1. and 2. in tandem with standard results on Markov Chain Monte Carlo methods ensure that Algorithm 1 preserves the distribution with density $\propto \exp(\ell(x))$.

Property 6. appears, in a slightly different form, in [11] (for the case $f = \ell$) and in [4] (for the general case). It shows that indeed (3)–(5) produces points x_k that, for δ small, remain close to the manifold $\mathcal{M}(f(x_0))$. It is remarkable that for an explicit timestepping algorithm where no implicit equations are solved, the deviation $f(x_K) - f(x_0)$ grows at most *linearly* with $K\delta$, the accumulated time after taking K steps of length δ. It would be possible to formulate a localized version of Property 6., where rather than demanding that H is globally bounded and globally Lipschitz, those properties are only demanded in a domain that contains the iterates x_k.

Property 7. holds in particular if in the context of Algorithm 1 the target is an isotropic Gaussian. In that situation, if $q(v|x)$ is also an isotropic Gaussian, Algorithm 1 accepts all the proposals.

3 The ODE

Formulas (3)–(5) are clearly reminiscent of a discretization of a system of ODEs. In this section we shall prove that, in fact, (3)–(5) provide an integrator for a system that we will identify. The original reference [11] did not discuss the connection between Hug and differential equations.

3.1 The derivative $N'(x)$

The expression for the system of ODEs approximated by (3)–(5) includes the derivative $N'(x)$ of the projector $N(x)$ defined in (8). This derivative is studied next.

For fixed $x \in \mathbb{R}^n$, $N'(x)$ maps *linearly* each vector $w \in \mathbb{R}^n$ into an $n \times n$ matrix, that we will denote by $N'(x)[w]$, in such a way that $N(x + w) = N(x) + N'(x)[w] + o(\|w\|)$ as $w \to 0$. If $z \in \mathbb{R}^n$, the notation $N'(x)[w]z$ will be used, as expected, to refer to the n-vector obtained by multiplying the matrix $N'(x)[w]$ and z. It is important to note that for each fixed x, $N'(x)[w]z$ depends *bilinearly* on w and z.

According to [7], for fixed x and w,

$$N'(x)[w] = N'_\perp(x)[w] + N'_\parallel(x)[w], \tag{10}$$

where, the $n \times n$ matrices $N'_\perp(x)[w]$, $N'_\parallel(x)[w]$ are given by

$$N'_\perp(x)[w] = J(x)^+ H(x)[w, \cdot]T(x), \tag{11}$$
$$N'_\parallel(x)[w] = \left(N'_\perp(x)[w]\right)^T = T(x)\left(H(x)[w, \cdot]\right)^T\left(J(x)^+\right)^T. \tag{12}$$

Here, $J(x)^+$ and $T(x)$ were defined in (7) and (9) respectively and $H(x)[w, \cdot]$ is the $m \times n$ matrix that corresponds to the linear operator $z \in \mathbb{R}^n \mapsto H(x)[w, z] \in \mathbb{R}^m$ resulting from freezing at w the first argument of the bilinear operator $H(x)[\cdot, \cdot]$. In this way, the right hand-side of (11) is the product of an $n \times m$ matrix, an $m \times n$ matrix and $n \times n$ matrix.

The following result, that explains the subscripts $\perp$ and $\parallel$, gives properties of the matrices $N'_\perp(x)[w]$, $N'_\parallel(x)[w]$ that we will use repeatedly in the proofs of the results.

Lemma 2. *For each x and w:*

1. *The image subspace of $N'_\perp(x)[w]$ is contained in $\mathcal{N}(x)$.*

2. *The kernel of $N'_\perp(x)[w]$ contains $\mathcal{N}(x)$.*

3. *The image subspace of $N'_\parallel(x)[w]$ is contained in $\mathcal{T}(x)$.*

4. *The kernel of $N'_\parallel(x)[w]$ contains $\mathcal{T}(x)$.*

It is useful to present an illustration:

Example 3. *For the case $n = 2$, $m = 1$, $f(x) = -ax_{(1)}^2 - bx_{(2)}^2$, $a, b > 0$ (corresponding to a bivariate Gaussian in Algorithm 1), if $w = [w_{(1)}, w_{(2)}]^T$, one finds*

$$N'_\perp(x)[w] = ab\frac{w_{(2)}x_{(1)} - w_{(1)}x_{(2)}}{(a^2x_{(1)}^2 + b^2x_{(2)}^2)^2}\begin{bmatrix} ax_{(1)} \\ bx_{(2)} \end{bmatrix}\begin{bmatrix} -bx_{(2)} & ax_{(1)} \end{bmatrix}.$$

This is a rank one, 2×2 matrix with image spanned by the vector $[ax_{(1)}, bx_{(2)}]^T$, whose direction coincides with that of $\nabla f(x)$ (normal to the contour of f that contains x). The kernel is also spanned by ∇f. Note that the dependence on w is linear.

The matrix $N'_\parallel(x)[w]$ is obtained by transposition. Its image and kernel are orthogonal to $\nabla f(x)$.

3.2 The ODE being approximated

Algorithm (3)–(5) will be shown to approximate, in a sense to be specified later, the following system of ODEs in $\mathbb{R}^n \times \mathbb{R}^n$:[1]

$$\frac{d}{dt}x = T(x)v, \tag{13}$$

$$\frac{d}{dt}v = \Big(N'_\parallel(x)\big[(T(x) - N(x))v\big] - N'_\perp(x)\big[(T(x) - N(x))v\big] \Big)v. \tag{14}$$

In order to shorten the notation, we will introduce the symbols

$$v_\parallel = T(x)v, \qquad v_\perp = N(x)v$$

for the tangential and normal components of the velocity v. With this notation and taking into account the properties of the kernels of the matrices $N'_\parallel[v_\parallel - v_\perp]$ and $N'_\perp[v_\parallel - v_\perp]$ in Lemma 2, (14) may be written as

$$\frac{d}{dt}v = N'_\parallel(x)[v_\parallel - v_\perp]v_\perp - N'_\perp(x)[v_\parallel - v_\perp]v_\parallel. \tag{15}$$

Some properties of (13)–(14) are contained in the following result. Clearly 1., 2. and 3. here are exact counterparts of Items 1., 2. and 4. in Proposition 1. In addition Item 6. in that proposition is an approximate version of Item 4. below.

Theorem 4. *The system* (13)–(14) *has the following properties:*

1. *It preserves volume in $\mathbb{R}^n \times \mathbb{R}^n$.*

2. *It is time-reversible, i.e. it remains invariant after changing t into $-t$ and v into $-v$.*

3. *It conserves the Euclidean length of $v(t)$: $(d/dt)\|v(t)\|^2 = 0$.*

4. *$x(t)$ remains on $\mathcal{M}(f(x(0)))$: $(d/dt)f(x(t)) = 0$*

The next result gives expressions for $(d/d)v_\parallel$ and $(d/dt)v_\perp$. Note that Lemma 2 implies that in (16) the terms $-N'_\parallel(x)[v_\perp]v_\perp$ and $-N'_\perp(x)[v_\parallel]v_\parallel$ are respectively the tangential and normal components of $(d/dt)v_\parallel$. Similarly the tangential and normal components of $(d/dt)v_\perp$ are $N'_\parallel(x)[v_\parallel]v_\perp$ and $N'_\perp(x)[v_\perp]v_\parallel$.

---

[1]The —rather unexpected— expression for this system was derived by assuming the ansatz (18), where x and $v = v_\parallel + v_\perp$ are solutions of a system of ODEs to be determined, substituting the ansatz in (21)–(22) and imposing $\sigma_k = \mathcal{O}(\delta^2)$ and $\tau_k = \mathcal{O}(\delta^2)$ (cf. the proof of Theorem 6).

Proposition 5. *If $(x(t), v(t))$ is a solution of* (13)–(14)*, then:*

$$\frac{d}{dt}v_\| = -N_\|'(x)[v_\perp]v_\perp - N_\perp'(x)[v_\|]v_\|, \tag{16}$$

$$\frac{d}{dt}v_\perp = N_\|'(x)[v_\|]v_\perp + N_\perp'(x)[v_\perp]v_\|. \tag{17}$$

3.3 Relating the Algorithm and the ODE

Let us now study the relation between the Algorithm (3)–(5) and the system (13)–(14). The situation is much complicated by the way the algorithm treats the normal component of the velocities v_k. According to (20), in the limit $\delta \to 0$ with fixed x_0, $N(x_0)v_1 = -N(x_0)v_0$; therefore the v_k's cannot be seen as approximations to the values of a differentiable function at the steppoints $k\delta$. We will show presently that the values x_k, v_k generated by the algorithm approximate the values, $k = 0, \ldots, K$

$$X_k = x(k\delta), \qquad V_k = v_\|(k\delta) + (-1)^k v_\perp(k\delta), \tag{18}$$

where $(x(t), v(t))$ is a suitable solution of (13)–(14). The factor $(-1)^k$ makes it possible for the V_k's to mimic the reflections in the normal components of the v_k's.

It is convenient to rewrite the algorithm after elimination of the auxiliary halfway points $x_{k+1/2}$:

$$x_{k+1} = x_k + \delta T(x_k + (\delta/2)v_k)v_k, \tag{19}$$

$$v_{k+1} = \big(I - 2N(x_k + (\delta/2)v_k)\big)v_k. \tag{20}$$

The first step in the analysis is to show *consistency*, i.e. that when X_k, V_k are substituted into the equations (19)–(20), they originate small residuals

$$\sigma_{k+1} := X_{k+1} - X_k - \delta T(X_k + (\delta/2)V_k)V_k, \tag{21}$$

$$\tau_{k+1} := V_{k+1} - \big(I - 2N(X_k + (\delta/2)V_k)\big)V_k. \tag{22}$$

As with any other explicit timestepping algorithm, these residuals may be seen as truncation errors (also referred to as local errors), i.e. as the error of the algorithm at $t = (k+1)\delta$ if started from the exact values at $t = k\delta$. More precisely, if for some k, x_k and v_k happened to coincide with $x(k\delta)$ and $v_\|(k\delta) + (-1)^k v_\perp(k\delta)$ respectively, then the algorithm would deliver, after one step, values

$$x_{k+1} = X_{k+1} - \sigma_{k+1}, \qquad v_{k+1} = V_{k+1} - \tau_{k+1}.$$

The following result shows consistency of order one:

Theorem 6. (First-order consistency.) *As $\delta \to 0$ and $k \to \infty$ with $k\delta \to t$, the truncation errors σ_k, τ_k are $\mathcal{O}(\delta^2)$.*

Consistency is the key step to proving the next convergence result. It is remarkable that, with *first order* consistency, the algorithm is convergent of the *second order*. This phenomenon, known as *supraconvergence* is discussed in e.g. [6], where additional references are also supplied. Supraconvergence happens here because there is

δ	$\|x_1 - x(\delta)\|$	$\|x_2 - x(2\delta)\|$
1/16	4.23(-4)	4.87(-5)
1/32	1.15(-4)	6.56(-6)
1/64	3.00(-4)	8.50(-7)
1/128	7.62(-6)	1.08(-7)
1/256	1.93(-6)	1.36(-8)

Table 1: Errors in the x variable vs. δ after one or two timesteps.

substantial cancellation between the truncation errors at consecutive steps. A numerical illutration of this cancellation is provided in Table 1, where f is as in Figure 1, $x_0 = [\cos(1), (1/2)\sin(1)]^T$, $v_0 = [0, 1]^T$. We see that the errors after a single step (i.e. the truncation errors) decrease as $\mathcal{O}(\delta^2)$ in agreement with Theorem 6. However when a second step is taken with the same value of δ, the error decreases substantially. Errors after two steps exhibit a $\mathcal{O}(\delta^3)$ behaviour. These issues are discussed in detail in Section 5. For the time being we emphasize that this cancellation of the local error in consecutive steps is *not* the zigzagging of the midway approximation $x_{k+1/2}$ that may be seen in Figure 1.

Theorem 7. (Second-order convergence.) *Fix x_0, v_0 and a time interval $[0, T]$ and consider the solution $(x(t), v(t))$ of (13)–(14) with $x(0) = x_0$, $v(0) = v_0$ and the iterates (x_k, v_k) generated by (3)–(5). Then, as $\delta \to 0$,*

$$\max_{k\delta \leq T} \|x_k - x(k\delta)\| = \mathcal{O}(\delta^2),$$

and

$$\max_{k\delta \leq T} \|v_k - \big(v_\|(k\delta) + (-1)^k v_\perp(k\delta)\big)\| = \mathcal{O}(\delta^2).$$

The constants implicit in the $\mathcal{O}$ notation in the theorem vary with f, x_0, v_0 and T only.

Since the ODE preserves the value of f (Theorem 4, Property 4.), the $\mathcal{O}(\delta^2)$ deviations in the value of f for the x_k, proved in Property 6. of Proposition 1, match the second order convergence stated in Theorem 7.

3.4 Dynamics

The appearence of the system (13)–(14) is certainly unfriendly and it is not easy to guess the behaviour of its solutions. To get some insight, we illustrate its dynamics in the case of two examples.

Assume first that initially $v(0)$ is tangent, i.e. $v_\perp(0) = 0$. Then (17) shows that, for all t, $v_\perp(t) = 0$ and therefore $v(t) = v_\|(t)$. According to (16), the system reduces to:

$$\frac{d}{dt}x = v_\|,$$
$$\frac{d}{dt}v_\| = -N'_\perp(x)[v_\|]v_\|.$$

10

These are the equations of motion for a particle that is constrained to remain on the manifold $\mathcal{M}(f(x(0)))$ when there are no external forces (other than the normal force $-N'_\perp(x)[v_\parallel]v_\parallel$ exerted by the constraint). The preservation of $\|v_\parallel(t)\|^2$ is just the preservation of kinetic energy that takes place due to the absence of working forces.

As a second example we look at the case where $n = 2$, $m = 1$ and $f(x) = -ax^2_{(1)} - bx^2_{(2)}$, $a, b > 0$, as in Example 3. As pointed out there this would arise when sampling from a bivariate, anisotropic Gaussian target distribution. (Recall that for isotopic Gaussians all proposals are accepted.) There are four scalar equations in (13)–(14) and, using the first integrals $f(x)$ and $\|v\|^2$, the system is reducible to a two-dimensional one. For simplicity we restrict the attention to solutions where $f(x(0)) = -1$ (the general case may be retrieved from this by rescaling a and b). We may then parameterize the level set (ellipse) $\mathcal{M}(f(x(0)))$ in terms of an angular variable ϕ as follows:

$$x_{(1)} = a^{-1/2} \cos \phi, \qquad x_{(2)} = b^{-1/2} \sin \phi.$$

Then the unit vector tangent to the ellipse is

$$\mathbf{t}(\phi) = \mu^{-1/2}[-b^{1/2} \sin \phi, a^{1/2} \cos \phi]^T,$$

where

$$\mu = a \cos^2 \phi + b \sin^2 \phi, \tag{23}$$

and the corresponding unit normal vector is

$$\mathbf{n}(\phi) = \mu^{-1/2}[a^{1/2} \cos \phi, -b^{1/2} \sin \phi]^T.$$

We may parameterize $v_\parallel$ and $v_\perp$ in terms of scalars p and n:

$$v_\parallel = p\mathbf{t}, \qquad v_\perp = n\mathbf{n}.$$

In addition p and n are linked by the first integral $p^2 + n^2 = p(0)^2 + n(0)^2$ established in Theorem 4. Substitution of these expressions into (13), (15), leads, using the expressions for $N'_\parallel$ and $N'_\perp$ given in Example 3 and after considerable algebra, to the system

$$\frac{d}{dt}\phi = (ab)^{1/2}\mu^{-1/2}p, \tag{24}$$

$$\frac{d}{dt}p = \frac{1}{2}(ab)^{1/2}(a - b)\mu^{-3/2}\sin(2\phi)(c^2 - p^2), \tag{25}$$

where μ is the function of ϕ in (23) and $c^2 = p(0)^2 + n(0)^2$ is $\|v(0)\|^2$, $c \geq 0$. The system has to be considered only in the strip $|p| \leq c$ because $n(0)$ is real. The lines $p = \pm c$ that bound the strip are invariant.

In the isotropic case $a = b$, where the ellipses are circles, μ does not vary with ϕ, $p(t)$ remains constant and $\phi(t)$ varies linearly with t. In the anisotropic case $a \neq b$, it may be assumed that $b > a$, as the other case is reduced to this by interchanging the roles of the components $x_{(1)}$ and $x_{(2)}$ of x. The system has equilibria at $p = 0$, $\phi = j\pi/2$, j integer. One may easily prove that for j even the equilibria are centers and otherwise they are saddles. A sketch of the phase portrait may be seen in Figure 2,
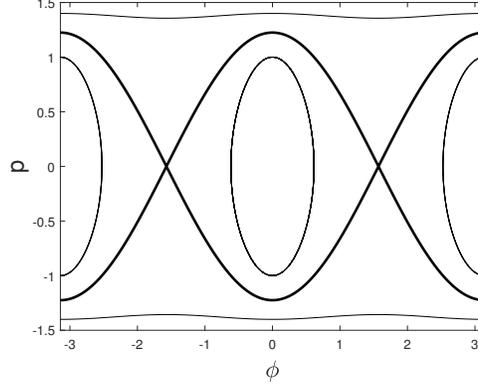
Figure 2: Phase plane of the system (24)–(25), when $a = 1$, $b = 4$, $c = \sqrt{2}$. The portrait is 2π-periodic in ϕ. The thicker lines correspond to the heteroclinic connections between saddle points. Outside the heteroclinic connections trajectories rotate; inside they librate.

where we note the similarity with the phase portrait of the standard pendulum equation. The saddle at $\phi = j\pi/2$, $p = 0$ (j odd) is connected to the neighbouring saddles at $\phi = (j \pm 2)\pi/2$, $p = 0$ via heteroclinic trajectories; these are the image in the (ϕ, p)-plane of solutions $(\phi(t), p(t))$ with

$$\lim_{t \to -\infty} (\phi(t), p(t)) = (j\pi/2, 0), \qquad \lim_{t \to \infty} (\phi(t), p(t)) = ((j + 2)\pi/2, 0)$$

or

$$\lim_{t \to -\infty} (\phi(t), p(t)) = (j\pi/2, 0), \qquad \lim_{t \to \infty} (\phi(t), p(t)) = ((j - 2)\pi/2, 0).$$

In the terminology of classical mechanics, trajectories at the top or the bottom of the figure (i.e. 'outside' the heteroclinic connections) *rotate*: in them p does not change sign as t varies and the angle ϕ increases monotonically (if $p > 0$) or decreases monotonically (if $p < 0$). The centers $\phi = j\pi/2$, $p = 0$ (j even) are surrounded by trajectories that *librate*: in them ϕ varies periodically around $j\pi/2$ between two values, say $\phi_{\max}, \phi_{\min}$, and p varies periodically around 0. In terms of the original system (13)–(14), in rotating trajectories, as $t \to \infty$, the point $x(t)$ describes again and again the ellipse $\mathcal{M}(f(x(0)))$ either clockwise (if $p(t) < 0$) or anticlockwise (if $p(t) > 0$). For librating trajectories, the solution $x(t)$ describes the arc of the ellipse parameterized by $\phi \in [\phi_{\max}, \phi_{\min}]$; when $\phi(t)$ has increased to $\phi_{\max}$ it starts decreasing and the trajectory in the x plane folds back on itself, see Figure 3. When $\phi(t)$ has decreased to $\phi_{\min}$ it starts increasing, causing the trajectory in the x plane to fold back on itself again. The process is repeated periodically. In the figure, the solution provided by the algorithm is seen to mimic the behaviour of the ODE (this has to be the case, if δ is sufficiently small, due to the convergence of the algorithm). Needless to say, this folding back phenomenon is completely unwelcome in the context
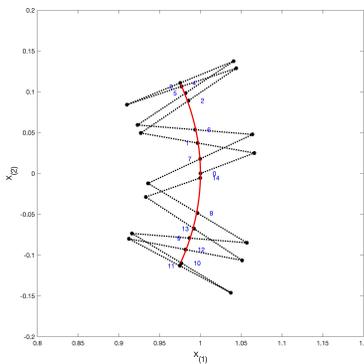
12

Figure 3: The continuous arc is the solution of the system (13)–(14) when $f(x) = -x_{(1)}^2 - 4x_{(2)}^2$, $x(0) = [1, 0]$, and $v(0) = [\sqrt{7/4}, 1/2]^T$ depicted in the x plane for $0 \leq t \leq 1.4$. (This corresponds to system (24)–(25) with $\phi(0) = 0$, $p(0) = 1/2$, $c = \sqrt{2}$.) The solution $x(t)$ first moves anticlockwise starting from $x(0)$ but then folds back on itself, moves clockwise and folds back once more. At the final time, $x(14)$ happens to be close to the initial location $x(0)$. Also depicted is the approximation provided by Algorithm (3)–(5), when $\delta = 0.1$ and $K = 14$. The numerical solution mimics the behaviour of the ODE; the numbers $0 - 14$ identify the stepnumber k of the iterates x_k.

of Algorithm 1, where one had the hope that, for suitable choices of K, the x_k would wholly explore $\mathcal{M}(f(x_0))$.

Note from Figure 2 that, generally speaking, the (desirable) rotating trajectories correspond to cases where $|p(0)|$ is large or, equivalently, $|n(0)| = \sqrt{c^2 - p(0)^2}$ is small relative to $c = \|v(0)\|$. This could have been expected because for $n(0) = 0$ the system just describes the inertial motion of a particle on the ellipse. In the scenario of Algorithm 1, $v(0)$ is a random variable with distribution $q(\cdot|x(0))$. For the standard choice where q is an $x(0)$-independent isotropic normal, $|n(0)|$ will not, on average, be small relative to $\|v(0)\|$ with the result that many trajectories will fold back on themselves. This suggests using distributions $q(v|x)$ where at each x vectors with larger tangential components are given more weight.

4   Numerical experiments and future work

As discussed in the introduction, this work was initially motivated by recent developments in the context of MCMC methods where the dynamics studied in the present paper underpins the design of updates of Markov chains aiming to explore, approximately, the contours of a probability density, or more generally a manifold. Analysis and observations in earlier sections point to potentially negative features of this dynamics. In the simple scenario of a two-dimensional ellipse we have provided a precise analysis of the dynamic, in particular characterising initial velocities leading to either
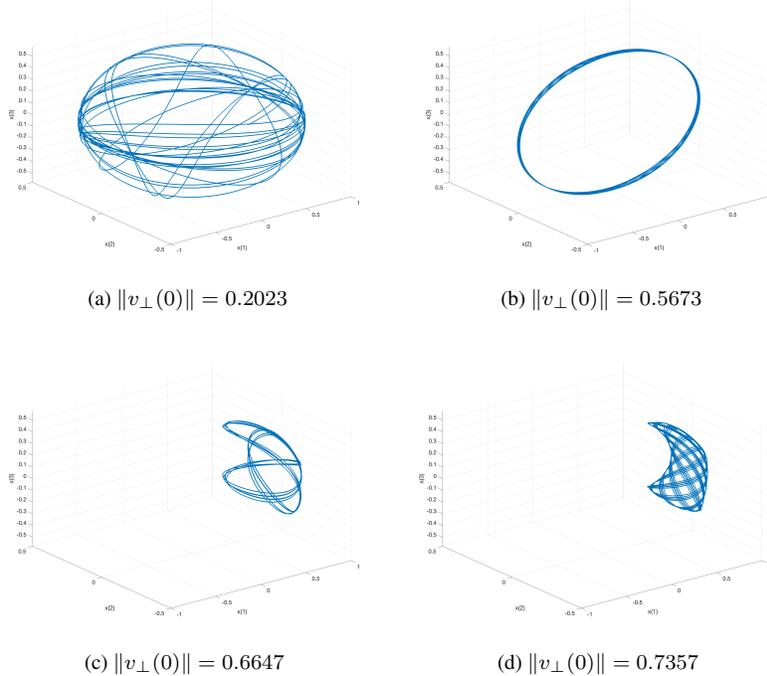
(a) $\|v_\perp(0)\| = 0.2023$

(b) $\|v_\perp(0)\| = 0.5673$

(c) $\|v_\perp(0)\| = 0.6647$

(d) $\|v_\perp(0)\| = 0.7357$

Figure 4: $K = 10,000$ integration steps of the dynamics constrained to a $3D$ ellipsoid $(n = 3)$, for four initial unit length velocities.

rotation or libration. However the two dimensional scenario is particular in that trajectories that fail to describe the full ellipse turn back on themselves exactly, resulting in a periodic behaviour. Natural questions are therefore what form this phenomenon takes in higher dimensions and whether it is even noticeable when used in the (random) proposal mechanism of an MCMC algorithm. Indeed, it is for example known that for a random vector $v \sim \mathcal{U}(\mathbb{S}^n)$ (where $\mathbb{S}^n$ is the unit hypersphere) and any fixed unit vector $u \in \mathbb{R}^n$ then $u^\top v$ tends to be statistically small in the sense that for $h \in [0, 1]$

$$\mathbb{P}(|u^\top v| \geq h) = \frac{\int_h^1 [1 - w^2]^{(n-2)/2} \mathrm{d}w}{\int_0^1 [1 - w^2]^{(n-2)/2} \mathrm{d}w},$$

confirming concentration on smaller values as n increases. This suggests that the negative phenomenon outlined earlier may perhaps naturally vanish in high dimensional scenarios. A full, and relevant, analytical investigation of these issues is beyond of the scope of the present manuscript and instead we present here a simple numerical study providing some insight into this phenomenon in the multivariate scenario.

We first explore the behaviour of the Hug integrator in the scenario where $n = 3$ and the constraint of interest is an ellipsoid defined by $f(x) = x^T A x$ with $A := \mathrm{diag}(1, 4, 3)$. In Fig. 4 we display $K = 10,000$ steps of the integrator for $\delta = 0.01$

14

and initial position $x(0) = [1, 0, 0]^T$ (therefore defining the level set of f) for four initial velocities $v(0)$ such that $\|v(0)\| = 1$. These are representative of what we have observed for numerous draws $v(0) \sim \mathcal{U}(\mathbb{S}_n)$: a large orthogonal component $\|v_\perp(0)\|$ results in trajectories that fail to explore a large area of the ellipsoid, due to a near oscillatory phenomenon. In fact we have found the size of the region visited to shrink for even larger values $\|v_\perp(0)\|$, but do not report the results here. For smaller values of K, more likely to be relevant in the context of Monte Carlo algorithms, we have similarly observed the influence of $\|v_\perp(0)\|$ on the size of the region explored by the dynamics.

In order to gain more insights into this phenomenon, we have run the following experiments. In the setup above we now consider $10,000$ realisations of the initial velocity $v^{(i)}(0) \overset{\text{iid}}{\sim} \mathcal{U}(\mathbb{S}_n)$ for $i = 1, \ldots, 10,000$ and report the scatter plot $i \mapsto \left(\|v_\perp^{(i)}(0)\|, \max_{0 \leq k \leq K} \|x^{(i)}(k) - x(0)\| \right)$, that is explore dependence of the largest Euclidean distance reached by the dynamic on $\|v_\perp(0)\|$. The results are reported in Fig. 5 for $K = 10, 100, 1000$ integration steps. This confirms our earlier observations, including the concentration of the whole dynamics around the initial position as $\|v_\perp(0)\|$ increases. We repeated this experiment for $n = 6$, $A = \text{diag}(1, 4, 3, 5, 1, 10)$ and $x(0) = (1, 0, 0, 0, 0, 0)^T$ and the results are reported in Fig. 6 with similar conclusions.

Our work suggests that MCMC based on Hug trajectories may benefit from the use of state-dependent and non-isotropic distributions for the velocity, limiting the likelihood of using velocities with dominant $v_\perp$ component; this however introduces a necessary tradeoff since the dynamics now does not leave the velocity distribution invariant, resulting in a higher rejection rate in the MH update [4]. More work is needed to understand this tradeoff.

5   Proofs

This section contains the proofs of the results.

5.1   Proof of Proposition 1

To prove 1., note that $(x, v) \mapsto \Psi_\delta(x, v) = (x', v')$ is the composition of three volume preserving maps, $(x, v) \mapsto (x + (\delta/2)v, v)$, $(x, v) \mapsto (x, R(x)v), (x, v) \mapsto (x + (\delta/2)v, v)$. Checking Property 2. is a trivial computation. Property 3. just restates (4). Property 4. follows from Property 3. Property 5. is a consequence of (3), (5) and Property 4.

With an obvious simplified notation, Taylor expansion around $x_{k+1/2}$ leads to
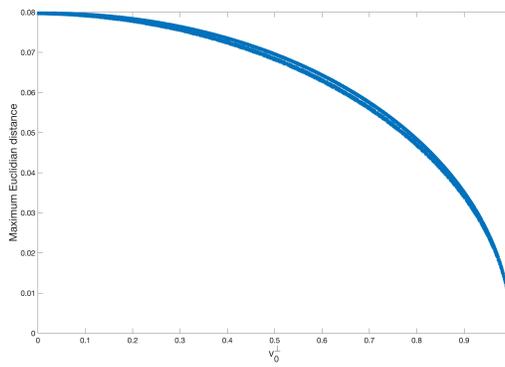
$$f_k = f_{k+1/2} - \frac{\delta}{2} J_{k+1/2} v_k + \frac{\delta^2}{4} H_k^+ [v_k, v_k],$$

$$f_{k+1} = f_{k+1/2} + \frac{\delta}{2} J_{k+1/2} v_{k+1} + \frac{\delta^2}{4} H_{k+1}^- [v_{k+1}, v_{k+1}],$$

(a) $K = 1000$
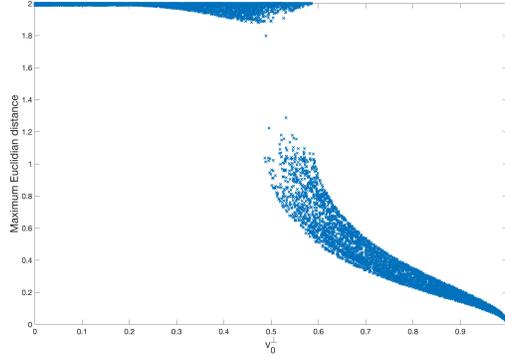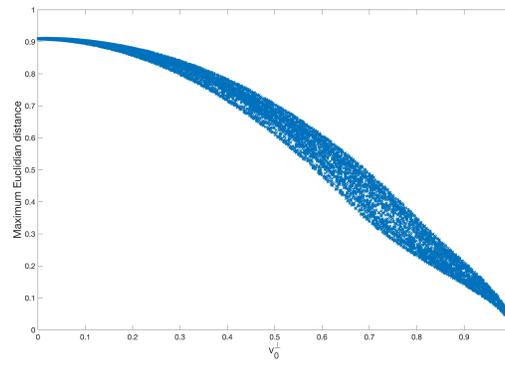
(b) $K = 100$
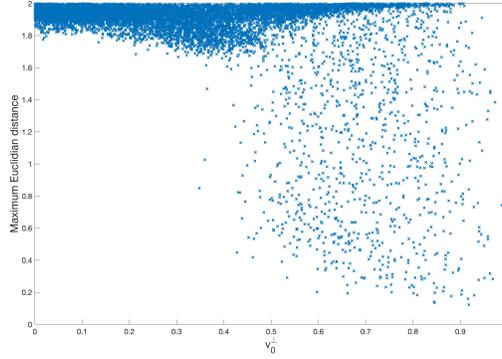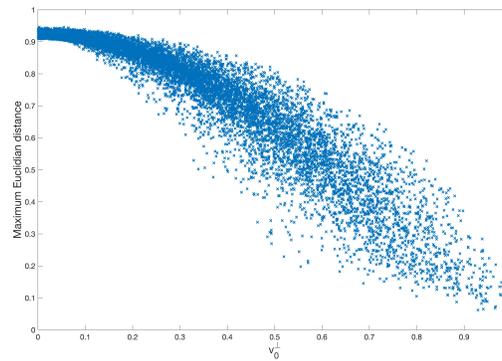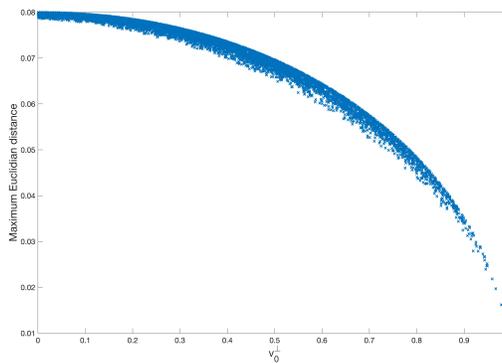

(c) $K = 10$

Figure 5: Scatter plots $i \mapsto \left(\|v_\perp^{(i)}(0)\|, \max_{0 \le k \le K} \|x^{(i)}(k) - x(0)\| \right)$ $(n = 3)$

(a) $K = 1000$

(b) $K = 100$

(c) $K = 10$

Figure 6: Scatter plots $i \mapsto \left(\|v_{\perp}^{(i)}(0)\|, \max_{0 \leq k \leq K} \|x^{(i)}(k) - x(0)\| \right)$ ($n = 3$)

17

with

$$H_k^+ = \int_0^1 H\big(x_{k+1/2} + \lambda(x_k - x_{k+1/2})\big)(1 - \lambda)d\lambda,$$

$$H_{k+1}^- = \int_0^1 H\big(x_{k+1/2} + \lambda(x_{k+1} - x_{k+1/2})\big)(1 - \lambda)d\lambda.$$

Subtracting, taking into account that $v_k + v_{k+1}$ is tangential by Property 3. and (6),

$$f_{k+1} - f_k = \frac{\delta^2}{4}\Big(H_{k+1}^-[v_{k+1}, v_{k+1}] - H_k^+[v_k, v_k]\Big) \tag{26}$$

and, summing by parts,

$$f(x_K) - f(x_0) = \frac{\delta^2}{4}\Big(H_K^-[v_K, v_K] - H_0^+[v_0, v_0] + \sum_{k=1}^{K-1}\big(H_k^- - H_k^+\big)[v_k, v_k]\Big).$$

By definition of the operator norm and Property 4., the norm of each of the terms $H_K^-[v_K, v_K]$ and $H_0^+[v_0, v_0]$ is bounded by

$$\beta \int_0^1 (1 - \lambda)d\lambda \, \|v_0\|^2 = \frac{\beta}{2}\|v_0\|^2.$$

On the other hand, $H_k^- - H_k^+$ equals

$$\int_0^1 \Big(H\big(x_{k-1/2} + \lambda(x_k - x_{k-1/2})\big) - H\big(x_{k+1/2} + \lambda(x_k - x_{k+1/2})\big)\Big)(1 - \lambda)d\lambda,$$

and, therefore, by the assumed Lispchitz continuity, and Property 5.

$$\|H_k^- - H_k^+\| \leq \int_0^1 \gamma(1 - \lambda)^2\|x_{k-1/2} - x_{k+1/2}\| \, d\lambda = \frac{\gamma}{3}\delta\|v_0\|.$$

After putting everything together, we get the bound in 6.

With the hypothesis in 7., Property 4. shows that the right hand-side of (26) vanishes, which concludes the proof. When $m = 1$ the manifolds $\mathcal{M}(\eta)$ are concentric hyperspheres, and the property $f(x_{k+1}) - f(x_k) = 0$ is apparent from the the geometry of the construction of x_{k+1} from x_k.

5.2 Proof of Lemma 2

1. From (11), the image of $N_\perp'(x)[w]$ is contained in the image of $J(x)^+$, which, according to (7) is contained in the image of $J(x)^T$, i.e. in the orthogonal of the kernel $\mathcal{T}(x)$ of $J(x)$ (see (6)), i.e. in $\mathcal{N}(x)$.

2. From (11), the kernel of $N_\perp'(x)[w]$ is contained in the kernel of $T(x)$, i.e. in $\mathcal{N}(x)$.

The other two claims are proved similarly.

5.3 Proof of Theorem 4

For Item 1. we have to prove that the vector field in (13)–(14) is divergence free. We first present some facts related to the computation of divergences. If $z \in \mathbb{R}^d \mapsto g(z) \in \mathbb{R}^d$ is a vector field, its divergence is of course the trace of the Jacobian

$$\text{div}_z(g(z)) = \sum_j \langle e_j, g'(z)e_j \rangle, \tag{27}$$

here the e_j are the unit coordinate vectors and $\langle \cdot, \cdot \rangle$ refers to the standard inner product. In the particular case of a linear vector field $z \mapsto Az$, this reduces to the trace of the matrix A:

$$\text{div}_z(Az) = \sum_j \langle e_j, Ae_j \rangle.$$

Since the trace is invariant by similarity, if Ω is an arbitrary invertible matrix,

$$\text{div}_z(Az) = \sum_j \langle e_j, Ae_j \rangle = \sum_j \langle e_j, \Omega^{-1}A\Omega e_j \rangle. \tag{28}$$

For a bilinear vector field $z \mapsto B(z, z)$, we observe that the Jacobian is formed adding the Jacobian with respect to the first argument of B (with the second frozen) to the Jacobian with respect to the second argument (with the first frozen). Therefore the divergence may be written as

$$\text{div}_z(B(z, z)) = \sum_j \langle e_j, B(e_j, z) \rangle + \sum_j \langle e_j, B(z, e_j) \rangle. \tag{29}$$

We are now ready to compute the divergence of the system (13)–(14). From (27) and (9),

$$\text{div}_x(T(x)v) = \sum_{j=1}^n \langle e_j, T'(x)[e_j]v \rangle = -\sum_{j=1}^n \langle e_j, N'(x)[e_j]v \rangle, \tag{30}$$

and, from (29), with a shortened notation,

$$\text{div}_v \left(\left(N'_\parallel(x)\big[(T(x) - N(x))v\big] - N'_\perp(x)\big[(T(x) - N(x))v\big] \right)v \right) =$$

$$\sum_{j=1}^n \langle e_j, \left(N'_\parallel\big[(T - N)e_j\big] - N'_\perp\big[(T - N)e_j\big] \right)v \rangle +$$

$$\sum_{j=1}^n \langle e_j, \left(N'_\parallel\big[(T - N)v\big] - N'_\perp\big[(T - N)v\big] \right)e_j \rangle. \tag{31}$$

The last summation vanishes because it is the trace of the matrix $N'_\parallel\big[(T - N)v\big] - N'_\perp\big[(T - N)v\big]$, which is skewsymmetric due to (12).[2] The properties of the images in

---

[2]In actual fact the traces of $N'_\parallel\big[(T - N)v\big]$ and $N'_\perp\big[(T - N)v\big]$ are both 0. According to Lemma 2, these matrices are nilpotent with $N'_\parallel\big[(T - N)v\big]^2 = 0$ and $N'_\perp\big[(T - N)v\big]^2 = 0$ which implies that all eigenvalues of $N'_\parallel\big[(T - N)v\big]$ and $N'_\perp\big[(T - N)v\big]$ vanish.

Lemma 2 show that the first summation in the right hand-side of (31) may be rewritten as

$$\sum_{j=1}^{n} \langle e_j, \Big((T-N)N_{\parallel}' \big[(T-N)e_j \big] + (T-N)N_{\perp}' \big[(T-N)e_j \big] \Big) v \rangle$$

or, resorting to (10),

$$\sum_{j=1}^{n} \langle e_j, (T-N)N' \big[(T-N)e_j \big] v \rangle.$$

Since $(T-N)^{-1} = T-N$, the formula (28) implies that this equals

$$\sum_{j=1}^{n} \langle e_j, N'[e_j]v \rangle,$$

which cancels (30). In this way volumen conservation has been established.

The time reversibility in Item 2. is clear: changing t into $-t$ and v into $-v$ reverses the sign of the left and right hand-sides of (13) and does not change the left and right hand-sides of (14) as the right hand-side is quadratic in v.

For Item 3., from (15)

$$\frac{1}{2} \frac{d}{dt} \langle v, v \rangle = \langle v, N_{\parallel}'(x)[v_{\parallel} - v_{\perp}]v_{\perp} \rangle - \langle v, N_{\perp}'(x)[v_{\parallel} - v_{\perp}]v_{\parallel} \rangle,$$

or, taking into account the properties of the images of $N_{\parallel}'(x)[v_{\parallel} - v_{\perp}]$ and $N_{\perp}'(x)[v_{\parallel} - v_{\perp}]$ in Lemma 2,

$$\frac{1}{2} \frac{d}{dt} \langle v, v \rangle = \langle v_{\parallel}, N_{\parallel}'(x)[v_{\parallel} - v_{\perp}]v_{\perp} \rangle - \langle v_{\perp}, N_{\perp}'(x)[v_{\parallel} - v_{\perp}]v_{\parallel} \rangle.$$

The right hand-side vanishes because the matrices involved are transposed of one another (see (12)).

Finally, for Item 4., (13) leads to

$$\frac{d}{dt} f(x) = J(x)v_{\parallel};$$

the right hand-side vanishes according to (6).

5.4   Proof of of Proposition 5

By using (9), (15), (13) and (10) successively, we may write

$$\begin{aligned} \frac{d}{dt}(Tv) &= \Big(\frac{d}{dt}T \Big) v + T \frac{d}{dt} v \\ &\quad - \Big(\frac{d}{dt}N \Big) v + TN_{\parallel}'[v_{\parallel} - v_{\perp}]v_{\perp} - TN_{\perp}'[v_{\parallel} - v_{\perp}]v_{\parallel} \\ &= -\big(N_{\parallel}'[v_{\parallel}] + N_{\perp}'[v_{\parallel}] \big)(v_{\parallel} + v_{\perp}) \\ &\quad\quad + TN_{\parallel}'[v_{\parallel} - v_{\perp}]v_{\perp} - TN_{\perp}'[v_{\parallel} - v_{\perp}]v_{\parallel}. \end{aligned}$$

Simplification with the help of Lemma 2 leads to (16)

Equation (17) is derived in a similar fashion.

5.5 Proof of Theorem 6

For σ:

$$\sigma_{k+1} = X_{k+1} - X_k - \delta T(X_k)V_k - \delta\big(T(X_k + (\delta/2)V_k) - T(X_k)\big)V_k$$

Now, note that

$$T(X_k + (\delta/2)V_k) - T(X_k) = \mathcal{O}(\delta),$$

and that, in view of (13),

$$X_{k+1} - X_k - \delta T(X_k)V_k = x((k+1)\delta) - x(k\delta) - \delta v_\|(k\delta) = \mathcal{O}(\delta^2).$$

We turn to τ. We write:

$$
\begin{aligned}
\tau_{k+1} &= V_{k+1} - \big(I - 2N(X_k)\big)V_k + 2\big(N(X_k + (\delta/2)V_k) - N(X_k)\big)V_k \\
&= V_{k+1} - T(X_k)V_k + N(X_k)V_k + 2\big(N(X_k + (\delta/2)V_k) - N(X_k)\big)V_k,
\end{aligned}
$$

and Taylor expanding

$$\tau_{k+1} = V_{k+1} - T(X_k)V_k + N(X_k)V_k + \delta N'(X_k)[V_k]V_k + \mathcal{O}(\delta^2),$$

or, from (18) (assuming k even) and (10), simplifying slightly the notation:

$$
\begin{aligned}
\tau_{k+1} &= v_\|((k+1)\delta) - v_\perp((k+1)\delta) - v_\|(k\delta) + v_\perp(k\delta) \\
&\quad + \delta N'_\|[v_\|(k\delta) + v_\perp(k\delta)](v_\|(k\delta) + v_\perp(k\delta)) \\
&\quad + \delta N'_\perp[v_\|(k\delta) + v_\perp(k\delta)](v_\|(k\delta) + v_\perp(k\delta)) + \mathcal{O}(\delta^2).
\end{aligned}
$$

We next recall (16) and (17) and Lemma 2 to get, simplifying further the notation:

$$
\begin{aligned}
\tau_{k+1} &= -\delta N'_\|[v_\perp]v_\perp - \delta N'_\perp[v_\|]v_\| - \delta N'_\|[v_\|]v_\perp - \delta N'_\perp[v_\perp]v_\| \\
&\quad + \delta N'_\|[v_\| + v_\perp]v_\perp + \delta N'_\perp[v_\| + v_\perp]v_\| + \mathcal{O}(\delta^2) \\
&= \mathcal{O}(\delta^2).
\end{aligned}
$$

For k odd the technique of proof is the same.

5.6 Proof of Theorem 7

Standard textbook proofs (see e.g. [9, Chapter 2.3]), going all the way back to Henrici [10], of the convergence of consistent one-step integrators assume that the map that advances the numerical solution over a single time-step is of the form $Id + \delta\Delta$, where the so-called increment function Δ is assumed to be Lipschitz in the neighbourhood of the ODE solution being approximated. Due to the flipping of the normal component of the velocity, that assumption does not hold for (3)–(5). There is a second difficulty. The truncation errors (21)–(22) are $\mathcal{O}(\delta^2)$ (first-order consistency) and the usual consistency+stability argument would lead to $\mathcal{O}(\delta)$ global error bounds (first order convergence), while Theorem 7 claims $\mathcal{O}(\delta^2)$ error bounds (second-order convergence).

In order to circumvent these difficulties, we consider the numerical integrator $\widehat{\mathcal{I}}$ such that one timestep of length 2δ with $\widehat{\mathcal{I}}$ is the result of taking successively two steps of length δ with (3)–(5). In symbols, the one step map $\widehat{\Psi}$ associated with $\mathcal{I}$ satisfies $\widehat{\Psi}_{2\delta} = \Psi_\delta \circ \Psi_\delta$. Clearly $\widehat{\Psi}_{2\delta}$ has the required $Id + (2\delta)\widehat{\Delta}$ structure, where $\widehat{\Delta}$ is differentiable, due to smoothness of f, and therefore locally Lispchitz. The consistency of (3)–(5) implies that the new integrator is also consistent (see [8, Theorem II.4.1]). In addition, the time-reversibility in Proposition 1 implies that its order of consistency has to be *even* (see e.g. [8, Section II.3, Section V.1]); simple examples (say $f(x) = \|x\|^2$) show that the order is exactly two. (An aside: Ψ_δ is consistent and time-reversible and yet is not consistent of even order. There is no contradiction, because the argument that shows even order for reversible integrators in [8, Section II.3, Section V.1] requires $\Psi_\delta = Id + \mathcal{O}(\delta)$ something that does not hold for (3)–(5).)

For $\widehat{\mathcal{I}}$, consistency of order two entails convergence of order two following the standard argument. The even numbered points (x_{2k}, v_{2k}) generated by (3)–(5) may be seen as coming from $\widehat{\mathcal{I}}$ and this proves

$$\max_{2k\delta \leq T} \left(\|x_{2k} - x(2k\delta)\| + \|v_{2k} - v(2k\delta)\| \right) = \mathcal{O}(\delta^2), \qquad \delta \to 0.$$

For an odd-numbered steppoint $(2k+1)\delta$, one may see the timestepping with (3)–(5) as being performed into two parts. First, one timesteps from $t = 0$ to $t = 2k\delta$ and then one performs a single timestep from $t = 2k\delta$ to $t = (2k+1)\delta$. The first part, as shown in the last display, introduces $\mathcal{O}(\delta^2)$ errors. The second part just introduces a single local error, which we know from Theorem 6 is of size $\mathcal{O}(\delta^2)$. This concludes the proof.

It may be of some interest to emphasize that since (3)–(5) is consistent of order one and $\widehat{\mathcal{I}}$ is consistent of order two, when timestepping with (3)–(5), the truncation error at the present step will almost be cancelled by the truncation error at the next step. (This is *not* the zigzagging of the midway points $x_{k+1/2}$ in Figure 1—that zigzagging cancels deviations of $x_{k+1/2}$ and $x_{k+3/2}$ normal to the level set of f.) The cancellation of truncation errors leading to supraconvergence is borne out in Table 1, where the columns give the truncation errors in x at the initial point (x_0, v_0) for (3)–(5) and for $\widehat{\mathcal{I}}$.

A similar phenomenon takes place implicitly in Property 6. of Proposition 1. In (26) we saw that the change in f over a single timestep is $\mathcal{O}(\delta^2)$ (which matches the first-order consistency of (3)–(5)). This may wrongly suggest that after $K = \mathcal{O}(1/\delta)$ steps the change in f would be $\mathcal{O}(\delta)$, rather than $\mathcal{O}(\delta^2)$ as proved in Proposition 1. Note that in (26), while H_{k+1}^- and H_k^+ differ by a $\mathcal{O}(\delta)$ amount, the same is not true of v_k and v_{k+1}, which differ by a $\mathcal{O}(1)$ amount, so that $f_{k+1} - f_k = \mathcal{O}(\delta^2)$ but

$f_{k+1} - f_k \neq \mathcal{O}(\delta^3)$. However, over two consecutive steps

$$
\begin{aligned}
f_{k+2} - f_k &= \frac{\delta^2}{4}\Big(H_{k+2}^-[v_{k+2}, v_{k+2}] - H_{[k+1}^+[v_{k+1}, v_{k+1}]\Big) \\
&\qquad + \frac{\delta^2}{4}\Big(H_{k+1}^-[v_{k+1}, v_{k+1}] - H_k^+[v_k, v_k]\Big) \\
&= \frac{\delta^2}{4}\Big(H_{k+2}^-[v_{k+2}, v_{k+2}] - H_k^+[v_k, v_k]\Big) \\
&\qquad + \frac{\delta^2}{4}(H_{k+1}^- - H_{k+1}^+)[v_{k+1}, v_{k+1}].
\end{aligned}
$$

Here v_{k+2} and v_k differ an by a $\mathcal{O}(\delta)$ amount, and the same is true for H_{k+2}^- and H_k^+ and for H_{k+1}^- and H_{k+1}^+. Therefore $f_{k+2} - f_k = \mathcal{O}(\delta^3)$ (which of course matches the fact that $\widehat{\mathcal{I}}$ is consistent of order two).

References

[1] A. Beskos, N. Pillai, G. Roberts, J.-M. Sanz-Serna, and A. Stuart. Optimal tuning of the hybrid Monte Carlo algorithm. *Bernoulli*, 19(5A):1501 – 1534, 2013.

[2] A. Bouchard-Côté, S. J. Vollmer, and A. Doucet. The bouncy particle sampler: A nonreversible rejection-free Markov chain Monte Carlo method. *Journal of the American Statistical Association*, 113(522):855–867, 2018.

[3] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Physics Letters B*, 195(2):216–222, 1987.

[4] M. C. Escudero. *Approximate Manifold Sampling*. PhD thesis, School of Mathematics, University of Bristol, February 2024.

[5] Y. Fang, J. M. Sanz-Serna, and R. D. Skeel. Compressible generalized hybrid Monte Carlo. *The Journal of Chemical Physics*, 140(17):174108, 05 2014.

[6] B. Garcia-Archilla and J. Sanz-Serna. A finite difference formula for the discretization of d^3/dx^3 on nonuniform grids. *Mathematics of Computation*, 57(195):239–257, 1991.

[7] G. H. Golub and V. Pereyra. The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate. *SIAM Journal on Numerical Analysis*, (2):413–432.

[8] E. Hairer, C. Lubich, and G. Wanner. *Geometric numerical integration*. Springer Series in Computational Mathematics. Springer-Verlag, Berlin, second edition, 2006.

[9] E. Hairer, S. Nørsett, and G. Wanner. *Solving Ordinary Differential Equations I: Nonstiff Problems*. Springer Series in Computational Mathematics. Springer Berlin Heidelberg, 1993.

[10] P. Henrici. *Discrete variable methods in ordinary differential equations*. New York: Wiley, 1962.

[11] M. Ludkin and C. Sherlock. Hug and hop: a discrete-time, nonreversible markov chain monte carlo algorithm. *Biometrika*, 110(2):301–318, 2023.

[12] E. A. Peters and G. de With. Rejection-free monte carlo sampling for general potentials. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 85(2):026703, 2012.