

Energy preserving methods on Riemannian manifolds

Elena Celledoni*, Sølve Eidnes*, Brynjulf Owren*, Torbjørn Ringholm*

May 19, 2018

Abstract

The energy preserving discrete gradient methods are generalized to finite-dimensional Riemannian manifolds by definition of a discrete approximation to the Riemannian gradient, a retraction, and a coordinate center function. The resulting schemes are intrinsic and do not depend on a particular choice of coordinates, nor on embedding of the manifold in a Euclidean space. Generalizations of well-known discrete gradient methods, such as the average vector field method and the Itoh–Abe method are obtained. It is shown how methods of higher order can be constructed via a collocation-like approach. Local and global error bounds are derived in terms of the Riemannian distance function and the Levi-Civita connection. Some numerical results on spin system problems are presented.

Keywords: Geometric integration, discrete gradients, Riemannian manifolds, numerical analysis.

Classification: 37K05, 53B99, 65L05, 82-08

1 Introduction

A first integral of an ordinary differential equation (ODE) is a scalar-valued function on the phase space of the ODE that is preserved along solutions. The potential benefit of using numerical methods that preserve one or more such invariants is well-documented, and several energy-preserving methods have been developed in recent years. Among these are the discrete gradient methods, which were introduced for use in Euclidean spaces in [1], see also [2]. These methods are based on the idea of expressing the ODE using a skew-symmetric operator and the gradient of the first integral, and then creating a discrete counterpart to this in such a way that the numerical scheme preserves the energy.

For manifolds in general, one can use the same schemes expressed in local coordinates. A drawback is that the numerical approximation will typically depend on the particular choice of coordinates and also on the strategy used for transition between coordinate charts. Another

*Department of Mathematical Sciences, Norwegian University of Science and Technology, N-7491 Trondheim
Elena Celledoni: elena.celledoni@ntnu.no; Sølve Eidnes: solve.eidnes@ntnu.no; Brynjulf Owren:
brynjulf.owren@ntnu.no; Torbjørn Ringholm: torbjorn.ringholm@ntnu.no

This work was supported by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 691070.

alternative is to use a global embedding of the manifold into a larger Euclidean space, but then it typically happens that the numerical solution deviates from the manifold. Even if the situation can be amended by using projection, it may not be desirable that the computed approximation depends on the particular embedding chosen. Crouch and Grossmann [3] and Munthe-Kaas [4, 5] introduced different ways of extending existing Runge–Kutta methods to a large class of differentiable manifolds. Both these approaches are generally classified as Lie group integrators, see [6] or the more recent [7] for a survey of this class of methods. They can also both be formulated abstractly by means of a post-Lie structure which consists of a Lie algebra with a flat connection of constant torsion, see e.g. [8]. In the present paper we shall state the methods in a slightly different context, using the notion of a Riemannian manifold. It is then natural to make use of the Levi-Civita connection, which in contrast to the post-Lie setting is torsion-free, and which in general has a non-zero curvature. For our purposes it is also an advantage that the Riemannian metric provides an intrinsic definition of the gradient. Taking an approach more in line with this, Leimkuhler and Patrick [9] considered mechanical systems on the cotangent bundle of a Riemannian manifold and succeeded in generalising the classical leap-frog scheme to a symplectic integrator on Riemannian manifolds.

Some classical numerical methods in Euclidean spaces preserve certain classes of invariants; for instance, symplectic Runge–Kutta methods preserve all quadratic invariants. This can be useful when there is a natural way of embedding a manifold into a linear space by using constraints that are expressed by means of such invariants. An example is the 2-sphere which can be embedded in \mathbb{R}^3 by adding the constraint that these vectors should have unit length. The classical midpoint rule will automatically ensure that the numerical approximations remain on the sphere as it preserves all quadratic invariants. In general, however, the invariants preserved by these methods are expressed in terms of coordinates. Hence the preservation property of the method may be lost under coordinate changes if the invariant is no longer quadratic. In [10], a generalization of the discrete gradient method to differential equations on Lie groups and a broad class of manifolds was presented. Here we develop this further by introducing a Riemannian structure that can be used to provide an intrinsic definition of the gradient as well as a means to measure numerical errors.

The structure of this paper is as follows: In section 2, we formulate the problem to be solved and introduce discrete Riemannian gradient methods, as well as presenting some particular examples with special attention to a generalization of the Itoh–Abe discrete gradient. We also briefly discuss the Euclidean setting as a special choice of manifold and show how the standard discrete gradient methods are recovered in this case. In the third section, we consider higher order energy preserving methods based on generalization of a collocation strategy introduced by Hairer [11] to Riemannian manifolds. We present some error analysis in section 4, and show numerical results in section 5, where the methods are applied to spin system problems.

2 Energy preservation on Riemannian manifolds

Consider an initial value problem on the finite-dimensional Riemannian manifold (M, g) ,

$$\dot{u} = F(u), \quad u(0) = u^0 \in M. \quad (2.1)$$

We denote by $\mathcal{F}(M)$ the space of smooth functions on M . The set of smooth vector fields and differential one-forms are denoted $\Gamma(TM)$ and $\Gamma(T^*M)$ respectively, and for the duality pairing between these two spaces we use the angle brackets $\langle \cdot, \cdot \rangle$.

A first integral associated to a vector field $F \in \Gamma(TM)$ is a function $H \in \mathcal{F}(M)$ such that $\langle dH, F \rangle$ vanishes identically on M . First integrals are preserved along solutions of (2.1),

$$\frac{d}{dt}H(u(t)) = \langle dH(u(t)), \dot{u}(t) \rangle = \langle dH(u(t)), F(u(t)) \rangle = 0.$$

2.1 Preliminaries

The fact that a vector field F has a first integral H is closely related to the existence of a tensor field $\Omega \in \Gamma(TM \otimes T^*M) =: \Gamma(\mathcal{T}_1^1 M)$, skew-symmetric with respect to the metric g , such that

$$F(u) = \Omega(u) \text{grad}H(u), \quad (2.2)$$

where $\text{grad}H \in \Gamma(TM)$ is the Riemannian gradient, the unique vector field satisfying $\langle dH, \cdot \rangle = g(\text{grad}H, \cdot)$. Any ODE (2.1) where F is of this form preserves H , since

$$\frac{d}{dt}H(u) = \langle dH(u), \dot{u} \rangle = \langle dH(u), \Omega \text{grad}H(u) \rangle = g(\text{grad}H(u), \Omega \text{grad}H(u)) = 0.$$

A converse result is detailed in the following proposition.

Proposition 1. *Any system (2.1) with a first integral H can be written with an F of the form (2.2). The skew tensor field Ω can be chosen so as to be bounded near every nondegenerate critical point of H .*

Proof. Similar to the proof of Proposition 2.1 in [2], we can write an explicit expression for a possible choice of Ω ,

$$\Omega y = \frac{g(\text{grad}H, y)F - g(F, y)\text{grad}H}{g(\text{grad}H, \text{grad}H)}. \quad (2.3)$$

Clearly, $g(y, \Omega y) = 0$ for all y . Since H is a first integral, $g(F, \text{grad}H) = \langle dH, F \rangle = 0$, so $\Omega \text{grad}H = F$. For a proof that Ω is bounded near nondegenerate critical points, see [2]. \square

In fact, such a tensor field Ω often arises naturally from a two-form ω through $\Omega y = \omega(\cdot, y)^\sharp$. A well-known example is when ω is a symplectic two-form. Note that Ω is not necessarily unique.

Retractions, viewed as maps from TM to M , will play an important role in the methods we discuss here. Their formal definition can be found e.g. in [12]:

Definition 1. Let ϕ be a smooth map $\phi : TM \rightarrow M$ and let ϕ_p denote the restriction of ϕ to T_pM , with 0_p being the zero-vector in T_pM . Then ϕ is a *retraction* if it satisfies the conditions

1. ϕ_p is defined in an open ball $B_{r_p}(0_p) \subset T_pM$ of radius r_p about 0_p ,
2. $\phi_p(x) = p$ if and only if $x = 0_p$,
3. $T\phi_p|_{0_p} = \text{Id}_{T_pM}$.

A generic example of a retraction on (M, g) is obtained via the Riemannian exponential, setting $\phi_p(x) = \exp_p(x)$, i.e. following along the geodesic emanating from p in the direction x .

2.2 The discrete Riemannian gradient method

We adapt the discrete gradients in Euclidean space to discrete Riemannian gradients (DRG) on (M, g) by means of a retraction map ϕ and a center point function c .

Definition 2. A discrete Riemannian gradient is a triple $(\overline{\text{grad}}, \phi, c)$ ¹ where

1. $c : M \times M \rightarrow M$ is a continuous map such that $c(u, u) = u$ for all $u \in M$,
2. $\overline{\text{grad}} : \mathcal{F}(M) \rightarrow \Gamma(c^*TM)$,
3. $\phi : TM \rightarrow M$ is a retraction,

such that for all $H \in \mathcal{F}(M)$, $u \in M$, $v \in M$, $c = c(u, v) \in M$,

$$H(v) - H(u) = g(\overline{\text{grad}}H(u, v), \phi_c^{-1}(v) - \phi_c^{-1}(u)), \quad (2.4)$$

$$\overline{\text{grad}}H(u, u) = \text{grad}H(u). \quad (2.5)$$

The DRG $\overline{\text{grad}}H$ is a continuous section of the pullback bundle c^*TM , meaning that $\pi \circ \overline{\text{grad}}H = c$, where $\pi : TM \rightarrow M$ is the natural projection. We also need to define an approximation to be used for the tensor field $\Omega \in \Gamma(\mathcal{T}_1^1 M)$. To this end we let $\overline{\Omega} \in \Gamma(c^*\mathcal{T}_1^1 M)$ be a continuous skew-symmetric tensor field such that

$$\overline{\Omega}(u, u) = \Omega(u) \quad \forall u \in M.$$

Inspired by [10, 13], we propose the scheme

$$u^{k+1} = \phi_{c^k}(W(u^k, u^{k+1})), \quad c^k = c(u^k, u^{k+1}) \quad (2.6)$$

$$W(u^k, u^{k+1}) = \phi_{c^k}^{-1}(u^k) + h\overline{\Omega}(u^k, u^{k+1})\overline{\text{grad}}H(u^k, u^{k+1}), \quad (2.7)$$

where h is the step size. The scheme (2.6)–(2.7) preserves the invariant H , since

$$\begin{aligned} H(u^{k+1}) - H(u^k) &= g(\overline{\text{grad}}H(u^k, u^{k+1}), \phi_{c^k}^{-1}(u^{k+1}) - \phi_{c^k}^{-1}(u^k)) \\ &= g(\overline{\text{grad}}H(u^k, u^{k+1}), h\overline{\Omega}(u^k, u^{k+1})\overline{\text{grad}}H(u^k, u^{k+1})) = 0. \end{aligned}$$

Here and in the following we adopt the shorthand notation $c = c(u, v)$ as long as it is obvious what the arguments of c are.

The Average Vector Field (AVF) method has been studied extensively in the literature; some early references are [14, 2, 15]. This is a discrete gradient method, and we propose a corresponding DRG satisfying (2.4)–(2.5) as follows:

$$\overline{\text{grad}}_{\text{AVF}}H(u, v) = \int_0^1 (T_{\gamma_\xi} \phi_c)^\top \text{grad}H(\phi_c(\gamma_\xi)) d\xi, \quad \gamma_\xi = (1 - \xi)\phi_c^{-1}(u) + \xi\phi_c^{-1}(v), \quad (2.8)$$

where $(T_x \phi_c)^\top : T_{\phi_c(x)}M \rightarrow T_x M$ is the unique operator satisfying

$$g((T_x \phi_c)^\top a, b) = g(a, T_x \phi_c b), \quad \forall x, b \in T_c M, \quad a \in T_{\phi_c(x)} M.$$

¹To avoid cluttered notation we will just write $\overline{\text{grad}}$ for the triple $(\overline{\text{grad}}, \phi, c)$ in the sequel.

Furthermore, we have the generalization of Gonzalez' midpoint discrete gradient [1],

$$\overline{\text{grad}}_{\text{MP}} H(u, v) = \text{grad}H(c(u, v)) + \frac{H(v) - H(u) - g(\text{grad}H(c(u, v)), \eta)}{g(\eta, \eta)} \eta \quad (2.9)$$

where $\eta = \phi_c^{-1}(v) - \phi_c^{-1}(u)$.

Note that both these DRGs involve the gradient of the first integral. This may be a disadvantage if H is non-smooth or if its gradient is expensive to compute. Also, the implicit nature of the schemes requires the solution of an n -dimensional nonlinear system of equations at each time step. An alternative is to consider the Itoh–Abe discrete gradient [16], also called the coordinate increment discrete gradient [2], which in certain cases requires only the solution of n decoupled scalar equations. We now present a generalization of the Itoh–Abe discrete gradient to finite-dimensional Riemannian manifolds.

2.3 Itoh–Abe discrete Riemannian gradient

Definition 3. For any tangent space $T_c M$ one can choose a basis $\{E_1, \dots, E_n\}$ composed of tangent vectors E_i , $i = 1, \dots, n$, orthonormal with respect to the Riemannian metric g . Then, given $u, v \in M$, there exists a unique $\{\alpha_i\}_{i=1}^n$ so that

$$\phi_c^{-1}(v) - \phi_c^{-1}(u) = \sum_{i=1}^n \alpha_i E_i.$$

The *Itoh–Abe DRG* of the first integral H is then given by

$$\overline{\text{grad}}_{\text{IA}} H(u, v) = \sum_{j=1}^n a_j E_j, \quad (2.10)$$

where

$$a_j = \begin{cases} \frac{H(w_j) - H(w_{j-1})}{\alpha_j} & \text{if } \alpha_j \neq 0, \\ g(\text{grad}H(w_{j-1}), T\phi_c(\eta_{j-1})E_j) & \text{if } \alpha_j = 0, \end{cases}$$

$$w_j = \phi_c(\eta_j), \quad \eta_j = \phi_c^{-1}(u) + \sum_{i=1}^j \alpha_i E_i.$$

We refer to [13] for proof that this is indeed a DRG satisfying (2.4)-(2.5).

2.4 Euclidean setting

Let $M = V$ be an \mathbb{R} -linear space, and let g be the Euclidean inner product, $g(x, y) = x^T y$. The operator Ω is a solution dependent skew-symmetric $n \times n$ matrix $\Omega(u)$. For any $u \in V$, we have $T_u V \cong V$. The retraction $\phi: V \rightarrow V$ is defined as $\phi_p(x) = p + x$, the Riemannian exponential on V , so that $\phi_c^{-1}(v) - \phi_c^{-1}(u) = v - u$. The gradient $\text{grad}H$ is an n -vector whose i th component is

$\frac{\partial H}{\partial u_i}$, and the definition of the discrete Riemannian gradient coincides with the standard discrete gradient, since (2.4) now reads

$$H(v) - H(u) = \overline{\text{grad}}(u, v)^T (v - u).$$

Furthermore, (2.6)-(2.7) simply becomes the discrete gradient method introduced in [1], given by the scheme

$$u^{k+1} - u^k = h \overline{\Omega}(u^k, u^{k+1}) \overline{\text{grad}}H(u^k, u^{k+1}), \quad (2.11)$$

where $\overline{\Omega}$ is a skew-symmetric matrix approximating Ω . Typical choices are $\overline{\Omega}(u^k, u^{k+1}) = \Omega(u^k)$, or $\overline{\Omega}(u^k, u^{k+1}) = \Omega((u^{k+1} + u^k)/2)$ if one seeks a symmetric method.

The DRGs (2.8) and (2.9) become the standard AVF and midpoint discrete gradients in this case. For the Itoh–Abe DRG, the practical choice for the orthogonal basis would be the set of unit vectors, $\{e_1, \dots, e_n\}$, so that $\alpha_i = v_i - u_i$, and we get (2.10) with

$$a_j = \begin{cases} \frac{H(w_j) - H(w_{j-1})}{v_j - u_j} & \text{if } u_j \neq v_j, \\ \frac{\partial H}{\partial u_j}(w_{j-1}) & \text{if } u_j = v_j, \end{cases}$$

$$w_j = \sum_{i=1}^j v_i e_i + \sum_{i=j+1}^n u_i e_i,$$

which is a reformulation of the Itoh–Abe discrete gradient as it is given in [16], [2] and the literature otherwise.

3 Methods of higher order

In the Euclidean setting, a strategy to obtain energy preserving methods of higher order was presented in [17] and later in [11], see also [18]. This technique is generalized to a Lie group setting in [10]. We will here formulate these methods in the context of Riemannian manifolds.

3.1 Energy-preserving collocation-like methods on Riemannian manifolds

Let c_1, \dots, c_s be distinct real numbers. Consider the Lagrange basis polynomials,

$$l_i(\xi) = \prod_{j=1, j \neq i}^s \frac{\xi - c_j}{c_i - c_j}, \quad \text{and let } b_i := \int_0^1 l_i(\xi) d\xi. \quad (3.1)$$

We assume that c_1, \dots, c_s are such that $b_i \neq 0$ for all i . A step of the energy-preserving collocation-like method, starting at $u^0 \in M$, is defined via a polynomial $\sigma : \mathbb{R} \rightarrow T_c M$ of degree s satisfying

$$\sigma(0) = \phi_c^{-1}(u^0), \quad (3.2)$$

$$\frac{d}{d\xi} \sigma(\xi h) \Big|_{\xi=c_j} = T_{U_j} \phi_c^{-1} \left(\Omega_j \text{grad}_j H \right), \quad U_j := \phi_c(\sigma(c_j h)) \quad (3.3)$$

$$u^1 := \phi_c(\sigma(h)), \quad (3.4)$$

where

$$\text{grad}_j H := \int_0^1 \frac{l_j(\xi)}{b_j} (T_{U_j} \phi_c^{-1})^\top (T_{\sigma(\xi h)} \phi_c)^\top \text{grad} H(\phi_c(\sigma(\xi h))) \, d\xi, \quad \text{and} \quad \Omega_j := \Omega(U_j).$$

Notice that with $s = 1$ and independently on the choice of c_1 , we reproduce the DRG method (2.6)-(2.7) with the AVF DRG (2.8).

Using Lagrange interpolation and (3.3), the derivative of $\sigma(\xi h)$ at every point ξh is

$$\frac{d}{d\xi} \sigma(\xi h) = \sum_{j=1}^s l_j(\xi) T_{U_j} \phi_c^{-1} \left(\Omega_j \text{grad}_j H \right), \quad (3.5)$$

from which by integrating we get

$$\sigma(\tau h) = \phi_c^{-1}(u_0) + h \sum_{j=1}^s \int_0^\tau l_j(\xi) \, d\xi T_{U_j} \phi_c^{-1} \left(\Omega_j \text{grad}_j H \right).$$

The defined method is energy preserving, which we see by using

$$\frac{d}{d\xi} (\phi_c(\sigma(\xi h))) = T_{\sigma(\xi h)} \phi_c \left(\frac{d}{d\xi} \sigma(\xi h) \right),$$

and (3.5) to get

$$\begin{aligned} H(u^1) - H(u^0) &= \int_0^1 g \left(\text{grad} H(\phi_c(\sigma(\xi h))), \frac{d}{d\xi} \phi_c(\sigma(\xi h)) \right) \, d\xi \\ &= \int_0^1 g \left(\text{grad} H(\phi_c(\sigma(\xi h))), T_{\sigma(\xi h)} \phi_c \left(\sum_{j=1}^s l_j(\xi) T_{U_j} \phi_c^{-1} \left(\Omega_j \text{grad}_j H \right) \right) \right) \, d\xi \\ &= \int_0^1 g \left((T_{\sigma(\xi h)} \phi_c)^\top \text{grad} H(\phi_c(\sigma(\xi h))), \sum_{j=1}^s l_j(\xi) T_{U_j} \phi_c^{-1} \left(\Omega_j \text{grad}_j H \right) \right) \, d\xi \\ &= \sum_{j=1}^s b_j g \left(\int_0^1 \frac{l_j(\xi)}{b_j} (T_{U_j} \phi_c^{-1})^\top (T_{\sigma(\xi h)} \phi_c)^\top \text{grad} H(\phi_c(\sigma(\xi h))) \, d\xi, \Omega_j \text{grad}_j H \right) \\ &= \sum_{j=1}^s b_j g \left(\text{grad}_j H, \Omega_j \text{grad}_j H \right) = 0, \end{aligned}$$

and hence repeated use of (3.2)-(3.4) ensures $H(u^k) = H(u^0)$ for all $k \in \mathbb{N}$.

3.2 Higher order extensions of the Itoh–Abe DRG method

From the Itoh–Abe DRG one can get a new DRG, also satisfying (2.4), by

$$\overline{\text{grad}}_{\text{SIA}} H(u, v) = \frac{1}{2} \left(\overline{\text{grad}}_{\text{IA}} H(u, v) + \overline{\text{grad}}_{\text{IA}} H(v, u) \right). \quad (3.6)$$

We call this the *symmetrized Itoh–Abe DRG*. Note that we need the base point c to be the same in the evaluation of $\overline{\text{grad}}_{\text{IA}} H(u, v)$ and $\overline{\text{grad}}_{\text{IA}} H(v, u)$. When $c(u, v) = c(v, u)$ and $\overline{\Omega}_{(u, v)} = \overline{\Omega}_{(v, u)}$, we get a symmetric DRG method (2.6)-(2.7), which is therefore of second order.

Alternatively, one can get a symmetric 2-stage method by a composition of the Itoh–Abe DRG method and its adjoint. Furthermore, one can get energy preserving methods of any order using a composition strategy. To ensure symmetry of an s -stage composition method, one needs $c_i(u, v) = c_{s+1-i}(v, u)$ for different center points c_i belonging to each stage and, similarly, $\overline{\Omega}_i(u, v) = \overline{\Omega}_{s+1-i}(v, u)$.

4 Error analysis

4.1 Local error

In this section, $\varphi_t(u)$ is the t -flow of the ODE vector field F . The most standard discrete gradient methods have a low or moderate order of convergence, and that is also the case for the DRG methods unless special care is taken in designing $\overline{\Omega}$ and $\overline{\text{grad}}H$. We shall not pursue this approach here, but refer to the collocation-like methods if high order of accuracy is required. We shall see, however, that the methods designed here are consistent and can be made symmetric. Analysis of the local error can be done in local coordinates, assuming that the step size is always chosen sufficiently small, so that within a fixed step, $u^k, u^{k+1}, c(u^k, u^{k+1})$ and the exact local solution $u(t_{k+1})$ all belong to the same given coordinate chart. From the definition (2.6)-(2.7) it follows immediately that the representation of $u^{k+1}(h)$ satisfies $u^{k+1}(0) = u^k$ and $\frac{d}{dh}u^{k+1}(0) = F(u^k)$. Then by equivalence of local coordinate norms and the Riemannian distance, we may conclude that the local error in DRG methods satisfies

$$d(u^{k+1}, \varphi_h(u^k)) \leq Ch^2.$$

Similar to what was also observed in [10], the DRG methods (2.6)-(2.7) are symmetric whenever $\overline{\text{grad}}H(u, v) = \overline{\text{grad}}H(v, u)$, $\overline{\Omega}(u, v) = \overline{\Omega}(v, u)$, and $c(u, v) = c(v, u)$ for all $u, v \in M$. In that case we obtain an error bound for the local error of the form $d(u^{k+1}, \varphi_h(u^k)) \leq Ch^3$.

The collocation-like methods of section 3 have associated nodes $\{c_i\}_{i=1}^s$ and weights $\{b_i\}_{i=1}^s$ defined by (3.1). The order of the local error depends on the accuracy of the underlying quadrature formula given by these nodes and weights. The following result is a simple consequence of Theorem 4.3 in [18].

Theorem 1. *Let ψ_h be the method defined by (3.2)-(3.4). The order of the local error is at least*

$$p = \min(r, 2r - 2s + 2)$$

where r is the largest integer such that $\sum_{i=1}^s b_i c_i^{q-1} = \frac{1}{q}$ for all $1 \leq q \leq r$. This means that there are positive constants C and h_0 such that

$$d(\psi_h(u), \varphi_h(u)) \leq Ch^{p+1} \quad \text{for } h < h_0, u \in M.$$

Proof. Choose h small enough such that the solution can be represented in the form $u(h\xi) = \phi_c(\gamma(\xi h))$, $\xi \in [0, 1]$, and consider the corresponding differential equation for γ in T_cM :

$$\frac{d}{dt}\gamma(t) = (\phi_c^* F)(\gamma(t)) = (T_{\gamma(t)}\phi_c)^{-1} \Omega \text{grad}H(\phi_c(\gamma(t))). \quad (4.1)$$

Notice that $(T_\gamma \phi_c)^{-1} = T_U \phi_c^{-1}$ where $U = \phi_c \circ \gamma$ and $T_{U(t)} \phi_c^{-1} : T_{U(t)} M \rightarrow T_c M$ for every t . We obtain

$$\frac{d}{dt} \gamma(t) = T_{U(t)} \phi_c^{-1} \Omega (T_{U(t)} \phi_c^{-1})^T (T_{\gamma(t)} \phi_c)^T \text{grad} H(\phi_c(\gamma(t))). \quad (4.2)$$

Considering the Hamiltonian $\tilde{H} : T_c M \rightarrow \mathbb{R}$, $\tilde{H}(\gamma) := \phi_c^* H(\gamma) = H \circ \phi_c(\gamma)$, we can then rewrite (4.1) in the form

$$\frac{d}{dt} \gamma(t) = \tilde{\Omega}(\gamma) \text{grad} \tilde{H}(\gamma), \quad \tilde{\Omega}(\gamma) := T_{U(t)} \phi_c^{-1} \Omega (T_{U(t)} \phi_c^{-1})^T, \quad (4.3)$$

where we have used that $\text{grad} \tilde{H} = T_{\gamma(t)} \phi_c^T \text{grad} H(\phi_c(\gamma(t)))$, which is now a gradient on the linear space $T_c M$ with respect to the metric inherited from M , $g|_c$. Locally in a neighborhood of c , (3.2)-(3.4) applied to (4.3) coincides with the methods of Cohen and Hairer, and therefore the order result [18, Thm 4.3] can be applied. Since the Riemannian distance $d(\cdot, \cdot)$ and any norm in local coordinates are equivalent, the result follows. \square

4.2 Global error

We prove the following result for the global error in DRG methods.

Theorem 2. *Let $u(t)$ be the exact solution to (2.1) where F is a complete vector field on a connected Riemannian manifold (M, g) with flow $u(t) = \varphi_t(u^0)$. Let ψ_h represent a numerical method $u^{k+1} = \psi_h(u^k)$ whose local error can be bounded for some $p \in \mathbb{N}$ as*

$$d(\psi_h(u), \varphi_h(u)) \leq Ch^{p+1} \quad \text{for all } u \in M.$$

Suppose there is a constant L such that

$$\|\nabla F\|_g \leq L,$$

where ∇ is the Levi-Civita connection and $\|\cdot\|_g$ is the operator norm with respect to the metric g . Then the global error is bounded as

$$d(u(kh), u^k) \leq \frac{C}{L} (e^{khL} - 1) h^p \quad \text{for all } k > 0.$$

Proof. Denoting the global error as $e^k := d(u(kh), u^k)$, the triangle inequality yields

$$e^{k+1} \leq d(\varphi_h(u(kh)), \varphi_h(u^k)) + d(\varphi_h(u^k), \psi_h(u^k)).$$

The first term is the error at nh propagated over one step, the second term is the local error. For the first term, we find via a Grönwall type inequality of [19],

$$d(\varphi_h(u(kh)), \varphi_h(u^k)) \leq e^{hL} d(u(kh), u^k) = e^{hL} e^k.$$

Using the local error estimate for the second term, we get the recursion

$$e^{k+1} \leq e^{hL} e^k + Ch^{p+1},$$

which yields

$$e^k \leq C \frac{e^{khL} - 1}{e^{hL} - 1} h^{p+1} \leq \frac{C}{L} (e^{khL} - 1) h^p.$$

□

Remark: Following Theorem 1.4 in [19], the condition that F is complete can be relaxed if $\varphi_t(u^0)$ and $\{u^k\}_{k \in \mathbb{N}}$ lie in a relatively compact submanifold N of M containing all the geodesics from u^k to $\varphi_{kh}(u^0)$. This is the case if, for instance, H has compact, geodesically convex sublevel sets, since both $\varphi_t(u^0)$ and $\{u^k\}_{k \in \mathbb{N}}$ are restricted to the level set $M_{H(u^0)} = \{p \in M \mid H(p) = H(u^0)\}$ and hence lie in the sublevel set $N_{H(u^0)} = \{p \in M \mid H(p) \leq H(u^0)\}$.

5 Examples and numerical results

We test our methods on two different variants of the classical spin system, whose solution evolves on the d -fold product of two-spheres, $(S^2)^d$,

$$\frac{ds_i}{dt} = s_i \times \frac{\partial H}{\partial s_i}, \quad s_i \in S^2, \quad i = 1, \dots, d, \quad H \in \mathcal{F}((S^2)^d). \quad (5.1)$$

The Riemannian metric g on $(S^2)^d$ restricts to the so-called round metric on each copy of the sphere. This metric coincides with the Euclidean inner product on the tangent planes of each of the spheres.

Geometric integrators for such systems are discussed widely in the literature, see e.g. [20, 21, 22, 23] and references therein. We study one or more bodies whose orientation is represented by a vector s_i of unit length in \mathbb{R}^3 , so that s_i lies on the manifold $M = S^2 = \{s \in \mathbb{R}^3 : \|s\| = 1\}$. Here and in what follows, $\|\cdot\|$ denotes the 2-norm. Starting with $d = 1$, our choice of retraction ϕ is given by its restriction to p ,

$$\phi_p(x) = \frac{p+x}{\|p+x\|}, \quad (5.2)$$

with the inverse

$$\phi_p^{-1}(u) = \frac{u}{p^T u} - p$$

defined when $p^T u > 0$. We note that $p^T x = 0$ for all $x \in T_p S^2$. The tangent map of the retraction and its inverse are given by

$$T_x \phi_p = \frac{1}{\|p+x\|} \left(I - \frac{(p+x) \otimes (p+x)}{\|p+x\|^2} \right), \quad T_u \phi_p^{-1} = \frac{1}{p^T u} \left(I - \frac{u \otimes p}{p^T u} \right), \quad (5.3)$$

where \otimes denotes the outer product² of the vectors. For $d > 1$, we use the retraction defined by $\Phi_p(x) = (\phi_{p_1}(x_1), \dots, \phi_{p_d}(x_d))$, where each $\phi_{p_i}(x_i)$ is given by (5.2).

²If x and y are in \mathbb{R}^3 , $x \otimes y$ is the matrix-matrix product of x taken as a 3×1 matrix and y taken as a 1×3 matrix.

5.1 Example 1: Perturbed spinning top

We consider first a nonlinear perturbation of a spinning top, see [22]. This is a spin system with one spin s . Given the inertia tensor $\mathbb{I} = \text{diag}(\mathbb{I}_1, \mathbb{I}_2, \mathbb{I}_3)$, and denoting by s^2 the component-wise square of s , we can define the Hamiltonian as

$$H(s) = \frac{1}{2} (\mathbb{I}^{-1} s)^\top \left(s + \frac{2}{3} s^2 \right).$$

The ODE system can be written in the form

$$\frac{ds}{dt} = \Omega(s) \text{grad} H(s), \quad \Omega(s) = \hat{s},$$

using the hat operator defined by $\hat{s}y = s \times y$. We approximate this system numerically, testing the scheme (2.6)-(2.7) with different discrete Riemannian gradients: the AVF (2.8), the midpoint (2.9), the Itoh–Abe (2.10) and its symmetrized version (3.6). For the three symmetric methods, we have chosen $c(s, \tilde{s}) = \frac{s+\tilde{s}}{\|s+\tilde{s}\|}$, so that $\phi_c^{-1}(\tilde{s}) = -\phi_c^{-1}(s)$. Using that $\text{grad} H(s) = \mathbb{I}^{-1}(s + s^2)$ and considering the transpose of $T_{\gamma_\xi} \phi_c$ from (5.3), the AVF DRG becomes

$$\begin{aligned} \overline{\text{grad}}_{\text{AVF}} H(s, \tilde{s}) &= \int_0^1 \frac{1}{\|l_\xi\|} \left(I - \frac{l_\xi \otimes l_\xi}{\|l_\xi\|^2} \right) \mathbb{I}^{-1} (\phi_c(\gamma_\xi) + \phi_c(\gamma_\xi)^2) d\xi \\ &= \int_0^1 \frac{1}{\|l_\xi\|} \left(\mathbb{I}^{-1} (\phi_c(\gamma_\xi) + \phi_c(\gamma_\xi)^2) - \phi_c(\gamma_\xi)^\top \mathbb{I}^{-1} (\phi_c(\gamma_\xi) + \phi_c(\gamma_\xi)^2) \phi_c(\gamma_\xi) \right) d\xi, \end{aligned}$$

with $\gamma_\xi = (1 - \xi)\phi_c^{-1}(s) + \xi\phi_c^{-1}(\tilde{s}) = (1 - 2\xi)\phi_c^{-1}(s)$ and $l_\xi = c + \gamma_\xi$. Similarly, the midpoint DRG becomes

$$\overline{\text{grad}}_{\text{MP}} H(s, \tilde{s}) = \frac{1}{\|s + \tilde{s}\|} \left(\mathbb{I}^{-1} \left(s + \tilde{s} + \frac{2}{3} (s^2 + s\tilde{s} + \tilde{s}^2) \right) + \frac{\frac{1}{2} \|s + \tilde{s}\|^2 - 2}{\|\tilde{s} - s\|^2} (H(\tilde{s}) - H(s)) (\tilde{s} - s) \right),$$

where we have used that $g(s, s) = s^\top s = 1$ for all $s \in S^2$. To obtain the basis of $T_c M$ for the definition of the Itoh–Abe DRG, we have used the singular-value decomposition. For the first order scheme, noting that $\phi_s^{-1}(s) = 0$, we choose $c(s, \tilde{s}) = s$, and get $\alpha_j = \phi_s(\tilde{s})^\top E_j$, for $j = 1, 2$. Then the DRG (2.10) can be written as

$$\overline{\text{grad}}_{\text{IA}} H(s, \tilde{s}) = \frac{H(\phi_s(\phi_s^{-1}(\tilde{s})^\top E_1 E_1)) - H(s)}{\phi_s^{-1}(\tilde{s})^\top E_1} E_1 + \frac{H(\tilde{s}) - H(\phi_s(\phi_s^{-1}(\tilde{s})^\top E_1 E_1))}{\phi_s^{-1}(\tilde{s})^\top E_2} E_2. \quad (5.4)$$

We solve the same problem using the 4th, 6th and 8th order variants of the collocation-like scheme (3.2)-(3.4). Choosing in the 4th order case the Gaussian nodes $c_{1,2} = \frac{1}{2} \mp \frac{\sqrt{3}}{6}$ as collocation points and setting $c(s, \tilde{s}) = s$, we get the nonlinear system

$$\begin{aligned} S_1 &= h \phi_{s_0} \left(\frac{1}{2} T_{S_1} \phi_{s_0}^{-1} (\Omega_1 \text{grad}_1 H) + \left(\frac{1}{2} - \frac{\sqrt{3}}{3} \right) T_{S_2} \phi_{s_0}^{-1} (\Omega_2 \text{grad}_2 H) \right), \\ S_2 &= h \phi_{s_0} \left(\left(\frac{1}{2} + \frac{\sqrt{3}}{3} \right) T_{S_1} \phi_{s_0}^{-1} (\Omega_1 \text{grad}_1 H) + \frac{1}{2} T_{S_2} \phi_{s_0}^{-1} (\Omega_2 \text{grad}_2 H) \right), \\ s_1 &= h \phi_{s_0} (T_{S_1} \phi_{s_0}^{-1} (\Omega_1 \text{grad}_1 H) + T_{S_2} \phi_{s_0}^{-1} (\Omega_2 \text{grad}_2 H)), \end{aligned}$$

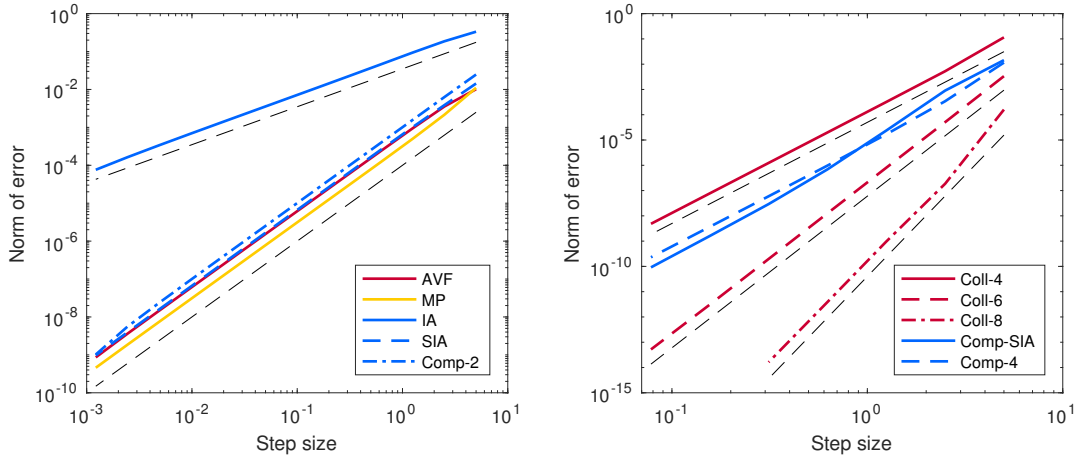


Figure 1: Error norm at $t = 10$ for the perturbed spinning top problem solved with different schemes, plotted with black, dashed reference lines of order 1, 2, 4, 6 and 8. Initial condition $s = (-1, -1, 1)/\sqrt{3}$ and $\mathbb{I} = \text{diag}(1, 2, 4)$. *Left*: The AVF, midpoint (MP), Itoh–Abe (IA) and symmetrized Itoh–Abe (SIA) DRGs and a 3-stage composition of the IA DRG scheme (Comp-2). *Right*: Collocation-type schemes of order 4, 6 and 8, a 3-stage composition of the SIA DRG scheme (Comp-SIA), and a 6-stage composition of the IA DRG scheme (Comp-4).

where

$$\sigma(\xi h) = \left((3 + 2\sqrt{3})\phi_{s_0}^{-1}(S_1) + (3 - 2\sqrt{3})\phi_{s_0}^{-1}(S_2) \right) \xi + \left(3(\sqrt{3} - 1)\phi_{s_0}^{-1}(S_2) - 3(1 + \sqrt{3})\phi_{s_0}^{-1}(S_1) \right) \xi^2$$

and we use the transposes of (5.3) and $\text{grad}H(s) = \mathbb{I}^{-1}(s + s^2)$ in the evaluation of $\text{grad}_1 H$ and $\text{grad}_2 H$. The 6th and 8th order schemes are derived in a similar manner, using the standard Gaussian nodes.

A second order scheme is derived by composing the Itoh–Abe DRG method with its adjoint, and a 4th order scheme is obtained by composing this method again with itself, as well as one by composition of the symmetrized Itoh–Abe DRG method with itself. In all stages of these composition methods, a symmetric $c(u, v)$ is used.

Plots confirming the order of all methods can be seen in Figure 1, where solutions using the different schemes are compared to a reference solution obtained using a very small step size. See the left hand panel of Figure 2 for numerical confirmation that our methods do indeed preserve the energy to machine precision, while the implicit midpoint method does not. In the right hand panel of Figure 2, the solution obtained by the Itoh–Abe DRG scheme with a step size $h = 1$ is plotted together with a solution obtained using the symmetrized Itoh–Abe DRG method with a much smaller time step. We observe, as expected for a method that conserves both the energy and the angular momentum, that the solution stays on the trajectories of the exact solution, although not necessarily at the right place on the trajectory at any given time.




Figure 2: *Left*: Energy error with increasing time for the AVF, midpoint (MP) and Itoh–Abe (IA) DRG methods, as well as the implicit midpoint (IMP) method, with step size $h = 1$, initial condition $s = (-1, -1, 1)/\sqrt{3}$ and $\mathbb{I} = \text{diag}(1, 2, 4)$. *Right*: Curves of constant energy on the sphere, found by our method with different starting values. The black solid line is the solution using the symmetrized Itoh–Abe DRG method with step size $h = 0.01$, while the red dots are the solutions obtained by the Itoh–Abe DRG method with step size $h = 1$.

5.2 Example 2: Heisenberg spin chain

We now consider the Heisenberg spin chain of micromagnetics. This problem is considered in [20, 23], where different geometric integrators are tested. Here, $s \in (S^2)^d$, and the Hamiltonian is

$$H(s) = \sum_{i=1}^d s_i^T s_{i-1}, \quad (5.5)$$

with $s_0 = s_d$ and $s_{d+1} = s_1$. The system (5.1) becomes, for this Hamiltonian,

$$\frac{ds_i}{dt} = \hat{s}_i (s_{i-1} + s_{i+1}), \quad i = 1, \dots, d,$$

and can be written in the block form

$$\frac{ds}{dt} = \Omega(s) \text{grad} H(s), \quad \text{where } \Omega(s) = \text{diag}(\hat{s}_1, \dots, \hat{s}_d). \quad (5.6)$$

For such a d -particle system, we may write the DRGs as

$$\overline{\text{grad}} H(s, \tilde{s}) = \left(\overline{\text{grad}}^1 H(s, \tilde{s}), \dots, \overline{\text{grad}}^d H(s, \tilde{s}) \right),$$

where we note that $\overline{\text{grad}}^i H(s, \tilde{s})$ is a discrete approximation to $\frac{\partial H}{\partial s_i}$. We thus get the AVF DRG

defined by

$$\begin{aligned}\overline{\text{grad}}_{\text{AVF}}^i H(s, \tilde{s}) &= \int_0^1 \frac{1}{\|l_{i,\xi}\|} \left(I - \frac{l_{i,\xi} \otimes l_{i,\xi}}{\|l_{i,\xi}\|^2} \right) (\phi_{c_{i-1}}(\gamma_{i-1,\xi}) + \phi_{c_{i+1}}(\gamma_{i+1,\xi})) \, d\xi \\ &= \int_0^1 \frac{1}{\|l_{i,\xi}\|} \left(\phi_{c_{i-1}}(\gamma_{i-1,\xi}) + \phi_{c_{i+1}}(\gamma_{i+1,\xi}) - l_{i,\xi}^\top (\phi_{c_{i-1}}(\gamma_{i-1,\xi}) + \phi_{c_{i+1}}(\gamma_{i+1,\xi})) l_{i,\xi} \right) \, d\xi,\end{aligned}$$

with $\gamma_{i,\xi} = (1 - 2\xi)\phi_{c_i}^{-1}(s_i)$ and $l_{i,\xi} = c_i + \gamma_{i,\xi}$. For the midpoint DRG we get

$$\overline{\text{grad}}_{\text{MP}}^i H(s, \tilde{s}) = c_{i-1} + c_{i+1} + \frac{H(\tilde{s}) - H(s) - (\text{grad}H(c(s, \tilde{s})))^\top \eta}{\eta^\top \eta} \eta_i,$$

where $\eta = (\eta_1, \dots, \eta_d)$ and $\eta_i = -2\phi_{c_i}^{-1}(s_i)$. In the numerical experiments, however, we have used a small modification of this,

$$\overline{\text{grad}}_{\text{MMP}}^i H(s, \tilde{s}) = c_{i-1} + c_{i+1} + \frac{\tilde{s}_i^\top \tilde{s}_{i-1} - s_i^\top s_{i-1} - (c_{i-1} + c_{i+1})^\top \eta_i}{\eta_i^\top \eta_i} \eta_i.$$

This DRG, which does indeed satisfy (2.4)-(2.5), leads to a more computationally efficient scheme than the original midpoint DRG. Each $\overline{\text{grad}}_{\text{IA}}^i H(s, \tilde{s})$ in the Itoh–Abe DRG is found as in the previous example, by (5.4). Higher order schemes are also derived in the same manner as before.

We test our schemes by comparing the numerical solutions with the exact solution

$$s_j(t) = (a \cos \theta_j + \tilde{a} \sin \theta_j) \cos \phi + \bar{a} \sin \phi, \quad \theta_j = jp - 2(1 - \cos p) \sin \phi,$$

for a choice of constants $\phi, p \in \mathbb{R}$ and orthogonal unit vectors $a, \tilde{a}, \bar{a} \in \mathbb{R}^3$, see [20]. Order plots for the methods are provided in Figure 3, using $d = 5$, $\phi = \pi/3$, $p = 2\pi/d$, $a = (1, 2, -1)/\sqrt{6}$, $\tilde{a} = (2, 1, 4)/\sqrt{21}$ and $\bar{a} = a \times \tilde{a}$. All schemes are shown to have the expected order.

6 Conclusions and further work

We have presented a general framework for constructing energy preserving numerical integrators on Riemannian manifolds. The main tool is to generalize the notion of discrete gradients as known from the literature. The new methods make use of an approximation to the Riemannian gradient coined the discrete Riemannian gradient, as well as a retraction map and a coordinate center function. An appealing feature of the new methods is that they do not depend on a particular choice of local coordinates or on an embedding of the manifold into a (larger) Euclidean space, but are of an intrinsic nature. Particular examples of discrete Riemannian gradient methods are given as generalizations of well-known schemes, such as the average vector field method, the midpoint discrete gradient method and the Itoh–Abe method. Extensions to higher order are proposed via a collocation-like method. We have analysed the local and global error behaviour of the methods, and they have been implemented and tested for certain spin systems where the phase space is $(S^2)^d$.

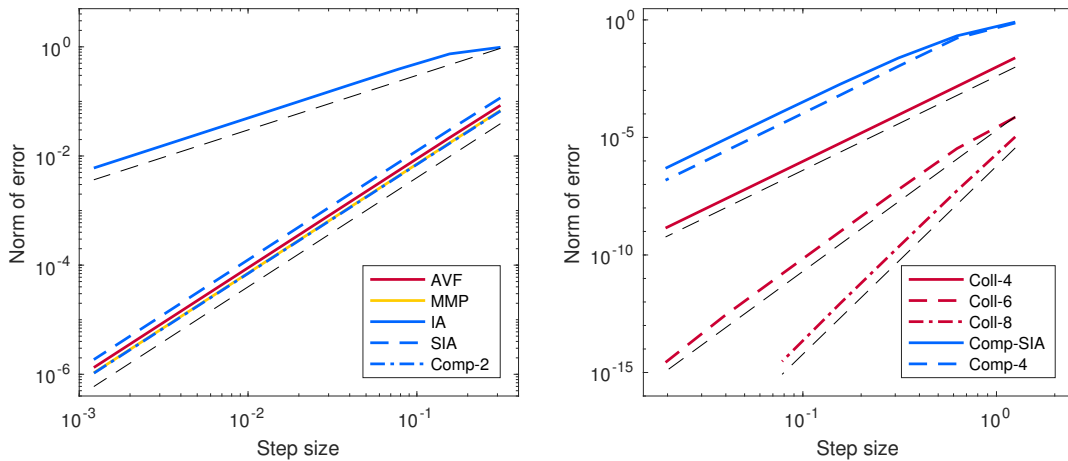


Figure 3: Error norm at $t = 10$ for the Heisenberg spin chain problem solved with different schemes, plotted with black, dashed reference lines of order 1, 2, 4, 6 and 8. *Left*: The AVF, modified midpoint (MMP), Itoh–Abe (IA) and symmetrized Itoh–Abe (SIA) DRGs and a 3-stage composition of the IA DRG scheme (Comp-2). *Right*: Collocation-type schemes of order 4, 6 and 8, a 3-stage composition of the SIA DRG scheme (Comp-SIA), and a 6-stage composition of the IA DRG scheme (Comp-4).

Possible directions for future research include a more detailed study of the stability and propagation of errors, taking into account particular features of the Riemannian manifold; for instance, it may be expected that the sectional curvature will play an important role. More examples should also be tried out, and we believe, inspired by [13], that there is a potential for making our implementations more efficient by tailoring them for the particular manifold, as well as the ODE problem considered.

References

- [1] O. Gonzalez, “Time integration and discrete Hamiltonian systems,” *J. Nonlinear Sci.*, vol. 6, no. 5, pp. 449–467, 1996.
- [2] R. I. McLachlan, G. R. W. Quispel, and N. Robidoux, “Geometric integration using discrete gradients,” *R. Soc. Lond. Philos. Trans. Ser. A Math. Phys. Eng. Sci.*, vol. 357, no. 1754, pp. 1021–1045, 1999.
- [3] P. E. Crouch and R. Grossman, “Numerical integration of ordinary differential equations on manifolds,” *Journal of Nonlinear Science*, vol. 3, no. 1, pp. 1–33, 1993.
- [4] H. Munthe-Kaas, “Lie–Butcher theory for Runge–Kutta methods,” *BIT Numerical Mathematics*, vol. 35, no. 4, pp. 572–587, 1995.
- [5] H. Munthe-Kaas, “Runge–Kutta methods on Lie groups,” *BIT Numerical Mathematics*, vol. 38, no. 1, pp. 92–111, 1998.

- [6] A. Iserles, H. Z. Munthe-Kaas, S. P. Nørsett, and A. Zanna, “Lie-group methods,” *Acta Numerica*, vol. 9, pp. 215–365, 2000.
- [7] E. Celledoni, H. Marthinsen, and B. Owren, “An introduction to Lie group integrators – basics, new developments and applications,” *Journal of Computational Physics*, vol. 257, pp. 1040–1061, 2014.
- [8] H. Z. Munthe-Kaas and A. Lundervold, “On post-Lie algebras, Lie–Butcher series and moving frames,” *Found. Comput. Math.*, vol. 13, no. 4, pp. 583–613, 2013.
- [9] B. Leimkuhler and G. W. Patrick, “A symplectic integrator for Riemannian manifolds,” *Journal of Nonlinear Science*, vol. 6, no. 4, pp. 367–384, 1996.
- [10] E. Celledoni and B. Owren, “Preserving first integrals with symmetric Lie group methods,” *Discrete Contin. Dyn. Syst.*, vol. 34, no. 3, pp. 977–990, 2014.
- [11] E. Hairer, “Energy-preserving variant of collocation methods,” *JNAIAM. J. Numer. Anal. Ind. Appl. Math.*, vol. 5, no. 1-2, pp. 73–84, 2010.
- [12] R. L. Adler, J.-P. Dedieu, J. Y. Margulies, M. Martens, and M. Shub, “Newton’s method on Riemannian manifolds and a geometric model for the human spine,” *IMA Journal of Numerical Analysis*, vol. 22, no. 3, pp. 359–390, 2002.
- [13] E. Celledoni, S. Eidnes, B. Owren, and T. Ringholm, “Dissipative schemes on Riemannian manifolds,” *arXiv preprint, arXiv:1804.08104*, 2018.
- [14] A. Harten, P. D. Lax, and B. van Leer, “On upstream differencing and Godunov-type schemes for hyperbolic conservation laws,” *SIAM Rev.*, vol. 25, no. 1, pp. 35–61, 1983.
- [15] G. Quispel and D. McLaren, “A new class of energy-preserving numerical integration methods,” *J. of Phys. A: Math. and Theor.*, vol. 41, no. 4, pp. 045206, 7, 2008.
- [16] T. Itoh and K. Abe, “Hamiltonian-conserving discrete canonical equations based on variational difference quotients,” *Journal of Computational Physics*, vol. 76, no. 1, pp. 85–102, 1988.
- [17] L. Brugnano, F. Iavernaro, and D. Trigiante, “Hamiltonian boundary value methods (energy preserving discrete line integral methods),” *J. Numer. Anal. Ind. Appl. Math*, vol. 5, no. 1, pp. 17–37, 2010.
- [18] D. Cohen and E. Hairer, “Linear energy-preserving integrators for Poisson systems,” *BIT Numerical Mathematics*, vol. 51, no. 1, pp. 91–101, 2011.
- [19] M. Kunzinger, H. Schichl, R. Steinbauer, and J. A. Vickers, “Global Gronwall estimates for integral curves on Riemannian manifolds,” *Rev. Mat. Complut.*, vol. 19, no. 1, pp. 133–137, 2006.
- [20] J. Frank, W. Huang, and B. Leimkuhler, “Geometric integrators for classical spin systems,” *Journal of Computational Physics*, vol. 133, no. 1, pp. 160–172, 1997.

- [21] D. Lewis and N. Nigam, “Geometric integration on spheres and some interesting applications,” *Journal of Computational and Applied Mathematics*, vol. 151, no. 1, pp. 141–170, 2003.
- [22] R. I. McLachlan, K. Modin, and O. Verdier, “Symplectic integrators for spin systems,” *Physical Review E*, vol. 89, no. 6, p. 061301, 2014.
- [23] R. McLachlan, K. Modin, and O. Verdier, “A minimal-variable symplectic integrator on spheres,” *Mathematics of Computation*, vol. 86, no. 307, pp. 2325–2344, 2017.