Geometric Integration Algorithms on Homogeneous Manifolds

Debra Lewis* and Peter J. Olver†

October 1, 2001

Abstract

   Given an ordinary differential equation on a homogeneous manifold, one can construct a "geometric integrator" by determining a compatible ordinary differential equation on the associated Lie group, using a Lie group integration scheme to construct a discrete time approximation of the solution curves in the group, and then mapping the discrete trajectories onto the homogeneous manifold using the group action. If the points of the manifold have continuous isotropy, a vector field on the manifold determines a continuous family of vector fields on the group, typically with distinct discretizations. If sufficient isotropy is present, an appropriate choice of vector field can yield improved capture of key features of the original system. In particular, if the algebra of the group is "full", then the order of accuracy of orbit capture (i.e. approximation of trajectories modulo time reparametrization) within a specified family of integration schemes can be increased by an appropriate choice of isotropy element. We illustrate the approach developed here with comparisons of several integration schemes for the reduced rigid body equations on the sphere.

1   Introduction.

Geometric integration techniques have become increasingly popular in the modern approach to numerical analysis. In the broad sense, "geometric integration" refers to numerical solution techniques for differential equations that preserve inherent geometric structures. Geometric integrators include symplectic and multisymplectic integrators that preserve the Hamiltonian or Poisson structure, [10, 12, 21, 35, 46, 9, 45, 2, 3, 29, 6], variational integrators that utilize the variational character of Lagrangian and canonical Hamiltonian systems, [4, 5, 31, 29], conservative integrators that preserve first integrals or conservation laws, [19, 20, 27, 42, 38, 39, 40, 25, 1], and symmetric integrators that preserve symmetries of the system, [7, 11, 17, 36]. A geometric integrator will track solutions over short time intervals as well as a standard scheme of the same order, e.g. a Runge–Kutta algorithm, while the extra expense required to construct and implement it will often be rewarded by significantly better performance in capturing the long term behavior and preserving geometric properties of the solutions.

In the more narrow sense, "geometric integration" refers to a family of Lie group integrators for ordinary differential equations and discretizations of partial differential equations. The rotation group plays a crucial role in many formulations of elasticity and plasticity and the advantages of exact rotations in numerical simulations of such materials, implemented via either the Rodriguez formula for the true exponential or the Cayley transform, have been amply demonstrated [41, 42, 40, 37]. Discretizations of specific dynamical systems on Lie groups that preserve not only the group structure, but additional geometric structures, have been used in the study of integrable systems; see, e.g. Moser and Veselov [31], Lewis and Simo [25], and McLaughlin and Scovel [28], for schemes preserving the group structure, the symplectic structure, and all point invariants of the generalized rigid body. Geometric integration schemes for general Lie groups and their associated bundles have been developed by Lewis and Simo [25, 26] and Munthe-Kaas, Iserles, Nørsett and their collaborators, [32, 15]. In particular, Munthe-Kaas [32] extended the classical Runge-Kutta algorithms to arbitrary Lie groups, creating a large, versatile family of geometric integrators. In general, if a Lie group G acts freely and transitively on a manifold M, then a differential equation on M uniquely determines a differential equation on G. Thus, replacing the original differential equation on M by the equivalent differential equation on G allows one to apply an appropriate Lie group integrators. The advantages of this approach over traditional integrators, e.g., a Runge–Kutta scheme, are discussed at length in the previously cited references.

If the action is transitive, but some points have continuous isotropy, then this construction is no longer unique — a given flow on M will correspond to continuous families of flows on the Lie group G. The purpose of this paper is to exploit this nonuniqueness in the design of geometric integration schemes that capture key features of the true dynamical system more accurately than other algorithms within a given class. We also show how to use isotropy to maintain the validity of conservation laws during the numerical approximation procedure. We illustrate our approach using flows on spheres, such as those arising in rigid body mechanics and micromagnetics, [24]. We have extensively tested several geometric integrators of various orders for the reduced free rigid body equations on the sphere. Further applications will be discussed elsewhere.

Our approach is reminiscent of the methods of Krupa, [18], who, in the compact group setting, applied an equivariant splitting of the dynamics into tangential and normal (isotropy) components to analyze stability and bifurcations of equilibria. The emphasis on isotropy and the utilization of any available freedom in generator selection is motivated in part by the first author's stability and bifurcation analyses of symmetric relative equilibria, [22, 23]. Our methods were originally inspired by the new equivariant approach to moving frames developed by the second author and Mark Fels, [13, 33]. In future work, we intend to apply the moving frame-based invariant numerical algorithms proposed in [34] to the systems under consideration in this paper.

2   The Basic Framework.

Let G be a Lie group acting transitively on a manifold M and let $\mathfrak{g}$ denote its Lie algebra of right-invariant vector fields on G. For simplicity, we shall assume that M is an embedded submanifold of a vector space $\mathcal{V}$. The tangent space $TM|_m$ at each $m \in M$ is then identified

with a subspace of $\mathcal{V}$. Given $m \in M$, let $G_m = \{g \cdot m = m\} \subset G$ denote the isotropy subgroup of $m \in M$ and $\mathfrak{g}_m \subset \mathfrak{g}$ its Lie subalgebra. In applications, G acts intransitively on the entire vector space, and M is a regular group orbit.

Let $\gamma \colon \mathfrak{g} \to \mathcal{X}$ denote the associated Lie algebra homomorphism[1] from the Lie algebra of G to the space of vector fields $\mathcal{X} = \mathcal{X}(M)$ on M. By transitivity, given a vector field $X \in \mathcal{X}$, we can construct a map $\psi \colon M \to \mathfrak{g}$ such that $X(m) = \gamma(\psi(m))(m)$. The map ψ is only prescribed up to isotropy. In other words,

$$X(m) = \gamma(\psi(m) + \zeta(m))(m), \tag{2.1}$$

for any map $\zeta \colon M \to \mathfrak{g}$ satisfying $\zeta(m) \in \mathfrak{g}_m$ for all $m \in M$.

A time-dependent vector field on M is defined as a smooth map $X \colon I \to \mathcal{X}$ from an interval $I \subset \mathbb{R}$ to the space of vector fields. According to (2.1), at each $m \in M$ and $t \in I$, we can write

$$X(m, t) = \gamma(\psi(m, t) + \zeta(m, t))(m), \tag{2.2}$$

where $\psi(m, t) \in \mathfrak{g}$, while $\zeta(m, t) \in \mathfrak{g}_m$ for all $m \in M$ and $t \in I$. We consider the associated non-autonomous flow $m(t) = \mathcal{F}_{t,t_0}(m)$ defined as the solution to the initial value problem

$$\frac{dm}{dt} = X(m, t), \qquad m(t_0) = m.$$

Since $X(m, t)$ is everywhere tangent to M, the flow $\mathcal{F}_{t,t_0}(m) \in M$ remains in M at all times (where defined), and our numerical approximations should reproduce this property as accurately as possible.

Let $\Delta t > 0$ denote the step size for the numerical algorithm, which may be fixed or variable. Our goal is to numerically approximate the time-dependent flow $\mathcal{F}_{t,t_0}$ at the subsequent time $t = t_0 + \Delta t$ by a map $\widetilde{\mathcal{F}}_{\Delta t, t_0} \colon M \to M$. The numerical algorithm will then approximate the flow by iterating this procedure[2] over the desired time interval. In accordance with standard procedures, we shall only analyze the local truncation error of the algorithm due to one iteration. Thus, without loss of generality, we may take the initial time $t_0 = 0$ from now on, and write

$$\mathcal{F}_t(m) = \mathcal{F}_{t,0}(m) = \exp(X(m, t)) \cdot m, \qquad m \in M, \tag{2.3}$$

for the time-dependent flow, and $\widetilde{\mathcal{F}}_{\Delta t} = \widetilde{\mathcal{F}}_{\Delta t, 0}$ for its numerical approximation. We shall further assume that the algorithm $\widetilde{\mathcal{F}}_{\Delta t}$ is obtained by evaluating a smooth time-dependent map $\widetilde{\mathcal{F}}_t \colon M \to M$ at time $t = \Delta t$. In our analysis, we shall ignore the machine-dependent effects of numerical approximation, e.g., round-off error, and assume that we can exactly compute the algorithmic flow to a specified order of truncation.

By transitivity of the group action, we can write our algorithm in the form

$$\widetilde{\mathcal{F}}_t(m) = g(m, t) \cdot m, \tag{2.4}$$

---

[1] Here is where right-invariance is required.

[2] For simplicity, we restrict our attention to single step algorithms.

where $g\colon M \times I \to G$ defines a time-dependent map from the manifold to the group. Assuming that $g(m,t)$ is sufficiently near the identity, we can regard

$$g(m,t) = \exp(\xi(m,t))$$

as arising from exponentiation of a Lie algebra-valued map $\xi : M \times I \to \mathfrak{g}$, and thus

$$\widetilde{\mathcal{F}}_t(m) = \exp(\xi(m,t)) \cdot m. \tag{2.5}$$

Appropriate maps ξ can be determined using standard families of algorithms (e.g. Runge–Kutta methods) with correction terms, if needed, constructed using the Lie bracket of the algebra $\mathfrak{g}$, as, for example, in [32].

In some situations, it is useful to generalize this framework by replacing exp with some form of "algorithmic exponential" Exp: $\mathfrak{g} \to G$. One important example is the Cayley transform for the orthogonal $O(n)$ and symplectic $\mathrm{Sp}(n)$ Lie groups, [44]. See [41, 42, 25, 26, 16] and the references therein for discussions of algorithmic exponentials, including applications of the Cayley transform to geometric integration. The computations in all cases are similar, and the reader will not lose much by assuming all exponentials are of standard type in this paper.

As we shall see, modification of the Lie algebra element used in the update, i.e. replacing (2.5) by

$$\widetilde{\mathcal{F}}_t(m) = \mathrm{Exp}(\xi(m,t) + \zeta(m,t)) \cdot m, \tag{2.6}$$

where the isotropy "improvement" $\zeta(m,t) \in \mathfrak{g}_m$ for all $m \in M$ and all $t \in I$, can change the qualitative features of the algorithmic trajectories, e.g. altering the apparent stability properties of equilibria or the conservation properties. Furthermore, in the presence of sufficient isotropy, given ξ, it is possible to design a map ζ such that the associated algorithmic update (2.6) captures the trajectories of the system to a higher order than is captured by the update determined by the original algorithm (2.5). In what follows, we shall consider families of algorithms in which the basic infinitesimal update' ξ is determined by some "standard" algorithm and specified choice of ψ, while the isotropy improvement term ζ is specified as a function of ξ and m.

Our primary goal is to exploit the ambiguity in the isotropy components of the generator and the infinitesimal update to increase the order of accuracy of a given algorithm. Therefore, we begin by formally defining the order of an approximation, based on the approximate and exact maps having the same Taylor expansion to order n.

Definition 1 Let $\mathcal{F}_t, \widetilde{\mathcal{F}}_t\colon M \to M$ be smooth time-dependent maps with $\mathcal{F}_0(m) = m = \widetilde{\mathcal{F}}_0(m)$. We shall say that $\widetilde{\mathcal{F}}_t(m)$ is an *order n approximation* of $\mathcal{F}_t(m)$ at $m \in M$ if

$$\lim_{t \to 0} \frac{\widetilde{\mathcal{F}}_t(m) - \mathcal{F}_t(m)}{t^n} = 0, \qquad \text{i.e.} \qquad \widetilde{\mathcal{F}}_t(m) - \mathcal{F}_t(m) = \mathcal{O}(t^{n+1}).$$

We shall say that $\widetilde{\mathcal{F}}_t$ is an *order n approximation* of $\mathcal{F}_t$ if $\widetilde{\mathcal{F}}_t(m)$ is an order n approximation of $\mathcal{F}_t(m)$ for all $m \in M$.

3   Algorithms on the Sphere.

As a warm-up to a general theory, we first treat the "simplest" case — the standard linear action of the rotation group $G = \mathrm{SO}(3)$ on $\mathbb{R}^3$. The orbits are spheres, and, without loss of generality, we concentrate on the unit sphere $M = S^2 = \{\,\|m\| = 1\,\} \subset \mathbb{R}^3$. We note that the isotropy subgroup $G_m \simeq \mathrm{SO}(2)$ for each $m \in S^2$ consists of the rotations around the axis through m.

To keep the presentation simple, we consider an autonomous vector field X on S^2 with flow $\mathcal{F}_t$. Let $\omega \colon S^2 \to \mathbb{R}^3$ denote the unique map satisfying

$$X(m) = \omega(m) \times m \qquad \text{and} \qquad \langle\, \omega(m), m \,\rangle = 0 \qquad \text{for all} \qquad m \in S^2.$$

Here we are making the usual identification of $\mathbb{R}^3 \simeq \mathfrak{so}(3)$ with the Lie algebra of the rotation group. We expand the flow of X in a power series

$$\begin{aligned}
\mathcal{F}_t(m) &= m + t\,\omega \times m + \tfrac{t^2}{2}\left[\,\omega \times (\omega \times m) + \dot\omega \times m\,\right] \\
&\quad + \tfrac{t^3}{6}\left[\,(-\|\omega\|^2\omega + \omega \times \dot\omega + \ddot\omega) \times m + 3\,\dot\omega \times (\omega \times m)\,\right] + \mathcal{O}(t^4).
\end{aligned} \tag{3.1}$$

Here $\dot\omega = \tfrac{d}{dt}\omega$ and $\ddot\omega = \tfrac{d^2}{d^2 t}\omega$ represent the time derivatives of the vector along the flow. We assume that we are not at an equilibrium point, and so

$$v \equiv \|\,\omega(m)\,\| \neq 0.$$

To facilitate the analysis, we introduce the orthonormal basis $\{m, \omega/v, m \times \omega/v\}$ and write

$$\omega^{(j)} = \frac{d^j}{dt^j}\,\omega = a_j\, m + b_j\,\frac{\omega}{v} + c_j\,\frac{\omega \times m}{v}\,, \qquad\qquad j = 0, 1, 2, \ldots.$$

Since $\langle\,\omega(m), m\,\rangle = 0$, we have $a_0 = a_1 = 0$; in addition, $b_0 = v$ and $c_0 = 0$. With respect to our orthonormal basis, the expansion (3.1) has the form

$$\mathcal{F}_t(m) = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + t \begin{pmatrix} 0 \\ 0 \\ -v \end{pmatrix} + \frac{t^2}{2} \begin{pmatrix} -v^2 \\ c_1 \\ -b_1 \end{pmatrix} + \frac{t^3}{6} \begin{pmatrix} -3\,b_1 v \\ c_2 \\ v^3 - b_2 \end{pmatrix} + \mathcal{O}(t^4). \tag{3.2}$$

On the other hand, an algorithmic update of the form

$$\widetilde{\mathcal{F}}_t(m) = \exp(\xi(m,t)) \cdot m, \qquad \text{where} \qquad \xi(m,t) = \sum_{j=1}^{\infty} \frac{t^j}{j!}\,\xi_j(m),$$

has the Taylor expansion

$$\begin{aligned}
\widetilde{\mathcal{F}}_t(m) &= m + t\,\xi_1 \times m + \tfrac{t^2}{2}[\,\xi_1 \times (\xi_1 \times m) + \xi_2 \times m\,] \\
&\quad + \tfrac{t^3}{6}[\,(-\|\xi_1\|^2\xi_1 + \tfrac{3}{2}\,\xi_1 \times \xi_2 + \xi_3) \times m + 3\,\xi_2 \times (\xi_1 \times m)\,] + \mathcal{O}(t^4).
\end{aligned} \tag{3.3}$$

(The first two terms in this expression are the same for both the matrix exponential and the Cayley transform, but differences appear at higher orders.) Note that $\xi(m)$ need not be tangent to M. We write the expansion (3.3) in terms of the basis used above, setting

$$\xi_j = \alpha_j\, m + \beta_j\,\frac{\omega}{v} + \gamma_j\,\frac{\omega \times m}{v}\,, \qquad j = 1, 2, \ldots, \qquad \|\xi_j\|^2 = \alpha_j^2 + \beta_j^2 + \gamma_j^2.$$

Then (3.3) takes the form

$$\widetilde{\mathcal{F}}_t(m) = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + t \begin{pmatrix} 0 \\ \gamma_1 \\ -\beta_1 \end{pmatrix} + \frac{t^2}{2} \begin{pmatrix} -\beta_1^2 - \gamma_1^2 \\ \alpha_1\beta_1 + \gamma_2 \\ \alpha_1\gamma_1 - \beta_2 \end{pmatrix}$$
$$+ \frac{t^3}{6} \begin{pmatrix} -3\,(\beta_1\beta_2 + \gamma_1\gamma_2) \\ \gamma_3 - \|\xi_1\|^2\gamma_1 + \frac{3}{2}(\alpha_1\beta_2 + \beta_1\alpha_2) \\ -\beta_3 + \|\xi_1\|^2\beta_1 + \frac{3}{2}(\alpha_1\gamma_2 + \gamma_1\alpha_2) \end{pmatrix} + \mathcal{O}(t^4).$$

(3.4)

Examining the terms in the discretization error $\dfrac{\widetilde{\mathcal{F}}_t(m) - \mathcal{F}_t(m)}{t}$ in order, we conclude that the algorithm is

a) consistent if and only if $\beta_1 = v$ and $\gamma_1 = 0$,

b) second order accurate if and only if, in addition, $\beta_2 = b_1$ and $\gamma_2 = c_1 - \alpha_1 v$,

c) third order accurate if and only if, in addition,

$$\beta_3 = b_2 + \tfrac{1}{2}\alpha_1^2 v + \tfrac{3}{2}c_1\alpha_1, \qquad \gamma_3 = c_2 - \tfrac{3}{2}(\,\alpha_2 v + \alpha_1 b_1\,).$$

In general, it can be seen that the method is order n accurate if and only if it is order $n - 1$ accurate and, in addition, the n–th order coefficients satisfy equations of the form

$$\beta_n = b_{n-1} + B_n(\,\xi_1, \ldots, \xi_{n-2}; \beta_{n-1}, \gamma_{n-1}; \omega, \ldots, \omega^{(n-2)}\,),$$
$$\gamma_n = c_{n-1} - \tfrac{n}{2}\,\alpha_{n-1}\beta_1 + C_n(\,\xi_1, \ldots, \xi_{n-2}; \beta_{n-1}, \gamma_{n-1}; \omega, \ldots, \omega^{(n-2)}\,),$$

for some functions B_n, C_n that depend upon the earlier terms in the expansion. Note that the isotropy coefficients α_j remain as free parameters in the matching of the two flows — although their values do affect the higher order coefficients in the expansion. In particular, it follows that if $\beta_1 \neq 0$, then α_{n-1} can be chosen so as to guarantee that the method is order n accurate.

In many situations, one is primarily interested in capturing the orbits, rather than actual trajectories. In other words, it may be sufficient to accurately describe where points in the phase space go, but not how long it takes them to get there. This freedom permits us to replace the true flow $\mathcal{F}_t$ with a flow with respect to some reparametrization $\tau = \tau(m, t)$ of the time. Both lows will have the same orbits, but move along them at different speeds. If an algorithm $\widetilde{\mathcal{F}}_t$ is an order $n - 1$ approximation of the true flow $\mathcal{F}_t$, then Corollary 4 states that the normal component of the approximation is accurate to order n. If the reparametrization has the form

$$\tau(m, t) = t + \frac{\theta(m)}{(n-1)!}\,t^{n-1}$$

for some function $\theta : M \to \mathbb{R}$, then

$$\frac{d^n}{dt^n}\left[\mathcal{F}_{\tau(m,t)}(m) - \mathcal{F}_t(m)\right]\Big|_{t=0} = \theta(m)X(m).$$

Thus an appropriate choice of θ can be used to eliminate the error in the direction of $X(m)$. It follows that an order $n - 1$ approximation to the flow $\mathcal{F}_t$ yields order n approximations to the

6

orbits of $\mathcal{F}_t$ if and only if the n–th order error is parallel to $X(m) = m \times \omega(m)$, i.e. if the n–th order error in the direction of $\omega(m)$ is zero.

Using the expansions (3.2) and (3.3), we see that a consistent algorithm of the form (2.6) captures orbits to second order if

$$\alpha_1 = \frac{c_1 - \gamma_2}{v}, \tag{3.5}$$

while a second order accurate algorithm captures orbits to third order if

$$\alpha_2 = \frac{1}{v} \left(\frac{2}{3}(c_2 - \gamma_3) - \frac{b_1}{v}(c_1 - \gamma_2) \right). \tag{3.6}$$

For the forward Euler method $\xi^{\mathrm{Eul}}(m, t) = t\,\omega(m)$; hence $\xi_1 = \omega(m)$ and $\xi_j = 0$, $j > 1$. In this case, (3.5) implies that $\alpha_1 = c_1/v$ yields a scheme that captures orbits to second order. For the backward Euler method $\xi^{\mathrm{BEu}}(m, t) = t\,\omega(\widetilde{\mathcal{F}}_t(m))$; it follows that $\xi_1 = \omega$ and $\xi_2 = 2\,\dot{\omega}$. Hence in this case $\alpha_1 = -c_1/v$ yields second order orbit capture. For the Heun, i.e. second order Runge–Kutta, method, with

$$\xi^{\mathrm{RK2}}(m, t) = \frac{t}{2}(\omega + \omega(\mathrm{Exp}(t\,\omega) \cdot m)) = t\,\omega + \frac{t^2}{2}\dot{\omega} + \frac{t^3}{4}\left(\ddot{\omega} - D\omega(m)(\dot{\omega} \times m)\right);$$

hence $\xi_1 = \omega$, $\xi_2 = \dot{\omega}$, and $\xi_3 = \frac{3}{2}(\ddot{\omega} - D\omega(m)(\dot{\omega} \times m))$. Since $\gamma_2 = c_1$, (3.6) implies that setting

$$\begin{aligned} \alpha_2 &= \frac{2}{3\,v}(c_2 - \gamma_3) = \frac{2}{3\,v}\left(c_2 - \frac{3}{2}\left(c_2 - \frac{1}{v}\langle\, D\omega(m)(\dot{\omega} \times m), m \times \omega \,\rangle \right) \right) \\ &= \frac{1}{v}\left\langle \frac{1}{v}D\omega(m)(\dot{\omega} \times m) - \frac{1}{3}\ddot{\omega}, m \times \omega \right\rangle \end{aligned}$$

yields an algorithm with third order orbit capture (but only second order overall accuracy).

The construction described above can be derived in a particularly simple and geometrically intuitive manner for the forward Euler method. If we assume that the algorithmic exponential is given by a rescaling of the matrix exponential (e.g. by the Cayley transform) and impose the restriction that the isotropy correction term be linear in t, then the update

$$m \quad \longmapsto \quad \widetilde{\mathcal{F}}_t^\sigma(m) = \exp(\tau(t, m)\,(\omega(m) + \sigma(m)m))$$

is a rigid rotation about an axis depending only on m. Hence the curve

$$\Gamma_\sigma(m, \epsilon) = \left\{ \widetilde{\mathcal{F}}_t^\sigma(m) : |t| \le \epsilon \right\}$$

is a segment of a circle in S^2. This segment will be the best circular approximation at m to the true orbit segment $\mathcal{O}(m, \epsilon) = \{ \mathcal{F}_t(m) : |t| \le \epsilon \}$ at m if and only if the geodesic curvatures of $\Gamma_\sigma(m, \epsilon)$ and $\mathcal{O}(m, \epsilon)$ agree at m. A straightforward calculation shows that these curvatures agree if and only if

$$\sigma(m) = \sigma_{\mathrm{i}}(m) := k_g(m)\,\|X(m)\|, \tag{3.7}$$

where

$$k_g(m) = \frac{\langle\, \ddot{m}, m \times \dot{m} \,\rangle}{\| m \|\,\| \dot{m} \|^3} \tag{3.8}$$

7

denotes the geodesic curvature of the orbit $\mathcal{O}(m, \epsilon)$ at m. Equation (3.8) follows immediately from the general formula

$$k_g(m) = \big\langle\, \alpha''(s), N(\alpha(s)) \times \alpha'(s) \,\big\rangle \tag{3.9}$$

for the geodesic curvature of a unit speed curve $\alpha(s)$ on a surface with unit normal $N(m)$, using the fact that $N(m) = m/\|\,m\,\|$ on a sphere centered at the origin. If $\mathcal{O}(m, \epsilon)$ is itself a segment of a circle, then $\Gamma_{\sigma_i}(m, \epsilon) = \mathcal{O}(m, \epsilon)$. Hence any torsion–free orbits, e.g. the separatrices of the reduced rigid body equations, are captured exactly by this version of the forward Euler method. Note that the choice $\sigma \equiv 0$, corresponding to using the generator orthogonal to m, yields segments of great circles.

To compare the orbit capture conditions derived by this geometric argument to those obtained using series expansions, we use the identities

$$\dot{m} = \omega \times m \qquad \text{and} \qquad \dddot{m} = \dot{\omega} \times m + \omega \times (\omega \times m),$$

the orthogonality condition $\langle\, \omega, m \,\rangle = 0$, and the hypothesis that $\|m\| = 1$ to express $\sigma_i(m)$ in the form

$$\sigma_i(m) = k_g(m)\,\|X(m)\| = \frac{\langle\, \ddot{m}, m \times \dot{m} \,\rangle}{\|\,\dot{m}\,\|^2} = \frac{\langle\, \dot{\omega}, m \times \omega \,\rangle}{v^2} = \frac{c_1}{v}.$$

Thus the condition $\alpha_1 = c_1/v$ for second order orbit capture in the forward Euler method is identical to the condition $\alpha_1 = \sigma_i(m)$, while the analogous condition for the implicit Euler method takes the form $\alpha_1 = -\sigma_i(m)$.

For higher order methods, the axis of rotation used in the update map $\widetilde{\mathcal{F}}_t^\sigma$ is typically time-dependent and hence the corresponding algorithmic trajectory segment typically is not circular (i.e. it has nonzero torsion). Hence the simple geometric argument used above cannot be applied. However, the strategy of curvature–matching can still be used. Since a smooth curve on a two dimensional manifold in $\mathbb{R}^3$ is determined up to a time reparametrization by its geodesic curvature, we can determine the conditions on the coefficients α_j, β_j, γ_j imposed by the restriction that the curvature of $\widetilde{\mathcal{F}}_t(m)$ match the curvature of $\mathcal{F}_t(m)$ to some order.

Lemma 2 *Let X be an everywhere nonzero vector field on a surface $S \subset \mathbb{R}^3$. Let $\mathcal{F}_t$ denote the flow of X, and suppose that $\widetilde{\mathcal{F}}_t$ is an order $n-1$ approximation to $\mathcal{F}_t$ for some $n \geq 2$. Then, $\widetilde{\mathcal{F}}_t$ captures the orbit of m under $\mathcal{F}_t$ to order n if and only if the geodesic curvature of the orbit of m under $\widetilde{\mathcal{F}}_t$ agrees with the geodesic curvature of the true orbit of m to order $n-2$.*

Proof: Let $N\colon S \to \mathbb{R}^3$ denote the normal vector field on S and let $\alpha(s)$ denote a parametrization of the true orbit through m by arc length. If we let

$$\mathcal{F}_t(m) = \sum_{j=0}^{\infty} \frac{t^j}{j!}\, m^{(j)} \qquad \text{and} \qquad k_g(\mathcal{F}_t(m)) = \sum_{j=0}^{\infty} \frac{t^j}{j!}\, k_g^{(j)}$$

denote the series expansions of the flow through m and the geodesic curvature (3.9) of the orbit at m, then

$$k_g^{(j)} = \frac{\big\langle\, m^{(j+2)}, N(m) \times X(m) \,\big\rangle}{\|X(m)\|^3} + \rho_j\left(m, \dot{m}, \ldots, m^{(j+1)} \right)$$

8

for some function ρ_j. An entirely analogous relation holds for the coefficients $\widetilde{m}^{(j)}$ and $\tilde{k}_g^{(j)}$ for the trajectory of $\widetilde{\mathcal{F}}_t$ through m. Thus if $\widetilde{\mathcal{F}}_t$ is an order $n-1$ approximation of $\mathcal{F}_t$, and hence $\widetilde{m}^{(j)} = m^{(j)}$ for $j = 0, \ldots, n-1$, then $\tilde{k}_g^{(j)} = k_g^{(j)}$ for $j = 0, \ldots, n-3$ and

$$\tilde{k}_g^{(n-2)} - k_g^{(n-2)} = \frac{\left\langle\, \widetilde{m}^{(n)} - m^{(n)}, N(m) \times X(m)\,\right\rangle}{||X(m)||^3}. \tag{3.10}$$

Thus the geodesic curvature $\tilde{k}_g(s)$ of the approximate orbit agrees with the geodesic curvature $k_g(s)$ of the true orbit to order $n-2$ if and only if the algorithm captures to order n the component of the orbit orthogonal to both the normal field N and the velocity field X, i.e. if and only if it captures the orbit to order n. ∎

4 The Reduced Free Rigid Body Equations.

We continue to consider the particular case of the rotation group $G = \mathrm{SO}(3)$ acting on the unit sphere $M = S^2$. We regard S^2 as a coadjoint orbit and use the symplectic structure

$$\Omega(m)(\xi \times m, \eta \times m) = -\left\langle\, m, \xi \times \eta\,\right\rangle$$

induced by the Lie–Poisson structure on $\mathfrak{so}(3)^* \simeq \mathbb{R}^3$. Let $\mathbb{I}$ be a positive definite symmetric matrix. The Hamiltonian function

$$H(m) = \tfrac{1}{2}\left\langle\, m, \mathbb{I}^{-1}m\,\right\rangle. \tag{4.1}$$

generates the vector field

$$X(m) = m \times \mathbb{I}^{-1}m \tag{4.2}$$

corresponding to the symplectic reduction of the free rigid body equations on $T^*\mathrm{SO}(3)$.

Euler's method for (4.2) is based on the family of maps

$$\mathcal{F}_t^\sigma(m) = \mathrm{Exp}(t(-\mathbb{P}_M \mathbb{I}^{-1}m + \sigma(m,t)\,m)) \cdot m,$$

where $\mathbb{P}_M$ denotes orthogonal projection onto the tangent space of $M = S^2$. The scalar function $\sigma\colon S^2 \times \mathbb{R} \to \mathbb{R}$ is arbitrary, and reflects the isotropy ambiguity in the flow. Formula (4.2) implies that

$$\begin{aligned}
\left\langle\, \ddot{m}, m \times \dot{m}\,\right\rangle &= \left\langle\, X(m) \times \mathbb{I}^{-1}m + m \times \mathbb{I}^{-1}X(m), m \times X(m)\,\right\rangle \\
&= \left\langle\, X(m), \mathbb{I}^{-1}X(m)\,\right\rangle - \left\langle\, m, \mathbb{I}^{-1}m\,\right\rangle ||X(m)||^2,
\end{aligned}$$

and so (3.7) yields

$$\sigma_{\mathrm{i}}(m) = \frac{\left\langle\, \ddot{m}, m \times \dot{m}\,\right\rangle}{||\dot{m}||^2} = \frac{\left\langle\, X(m), \mathbb{I}^{-1}X(m)\,\right\rangle}{||X(m)||^2} - \left\langle\, m, \mathbb{I}^{-1}m\,\right\rangle. \tag{4.3}$$

Note that if $I_1 \le I_2 \le I_3$ are the eigenvalues of of the inertia tensor $\mathbb{I}$, then

$$|\sigma_{\mathrm{i}}(m)| \le \frac{1}{I_1} - \frac{1}{I_3}$$

for all $m \in S^2$.

The reduced free rigid body is a conservative system on a two dimensional manifold; in the case of a triaxial body, when the eigenvalues $I_1 < I_2 < I_3$ are distinct, the level sets of the Hamiltonian (4.1) exactly determine the orbits of the system. Thus in this situation it is possible to specify orbit–preserving schemes without direct reference to the exact flow; it suffices to require that the scheme preserve the total energy to the desired order. If Exp is at least a second order approximation to the matrix exponential (e.g. if Exp is the Cayley transform), then the algorithmic flow satisfies

$$\mathcal{F}_t^\sigma(m) = m + t\,X(m) + \tfrac{t^2}{2}m_2 + \mathcal{O}(t^3),$$

where

$$m_2 = (\mathbb{I}^{-1}m + (\sigma - 2\,H(m)),m) \times X(m),$$

and hence

$$
\begin{aligned}
H(\mathcal{F}_t^\sigma(m)) - H(m) &= \tfrac{1}{2}\left\langle \mathcal{F}_t^\sigma(m) - m,\ \mathbb{I}^{-1}(\mathcal{F}_t^\sigma(m) + m)\right\rangle \\
&= t\left\langle X(m) + \tfrac{t}{2}m_2, \mathbb{I}^{-1}m + \tfrac{t}{2}\,X(m) + \tfrac{t^2}{4}m_2 \right\rangle + \mathcal{O}(t^3) \\
&= \tfrac{t^2}{2}\left(\left\langle X(m), \mathbb{I}^{-1}X(m)\right\rangle + (\sigma - 2\,H(m))\,||X(m)||^2\right) + \mathcal{O}(t^3).
\end{aligned}
$$

If m is an equilibrium, then $\mathcal{F}_t^\sigma(m) = m$ for any σ, while if $X(m) \neq 0$, then

$$H(\mathcal{F}_t^\sigma(m)) - H(m) = \tfrac{t^2}{2}(\sigma - \sigma_i(m))||X(m)||^2 + \mathcal{O}(t^3).$$

Thus in this situation we can derive the optimality of σ_i working directly from the energy constraint.

As the numerical results given below demonstrate, use of the isotropy correction σ_i yields an efficient, accurate version of the forward Euler scheme for the reduced free rigid body. While the scheme only captures orbits (and hence the energy) to second order for asymmetric bodies, the leading constants are quite small, yielding very good approximations even for very large time steps.

We tested the forward Euler method using three choices of σ:

1. $\sigma_b(m) = 2\,H(m) = \left\langle m, \mathbb{I}^{-1}m \right\rangle$, corresponding to the 'basic' generator $\omega_b(m) = \mathbb{I}^{-1}m$;

2. $\sigma_o \equiv 0$, corresponding to the orthogonal generator $\omega(m) = \mathbb{P}_M\mathbb{I}^{-1}m$;

3. σ_i given by the geodesic curvature improvement (4.3) of ω, yielding second order orbit capture.

As noted above, the separatrix is exactly captured if the infinitesimal updates ξ_o^{Eul} or ξ_i^{Eul}, which coincide on the separatrix, are used. On the other hand, when the 'basic' infinitesimal update ξ_b^{Eul} was used to integrate ten sample trajectories with initial conditions at random points on the separatrices of rigid bodies with randomly generated inertia tensors, the average errors over the integration interval $[0, 500]$ were: $9.72\ 10^{-2}$ for $\Delta t = 1$, $3.89\ 10^{-2}$ for $\Delta t = .1$, and $1.94\ 10^{-3}$ for $\Delta t = .01$. In table 3 we give some data generated using these choices of updates
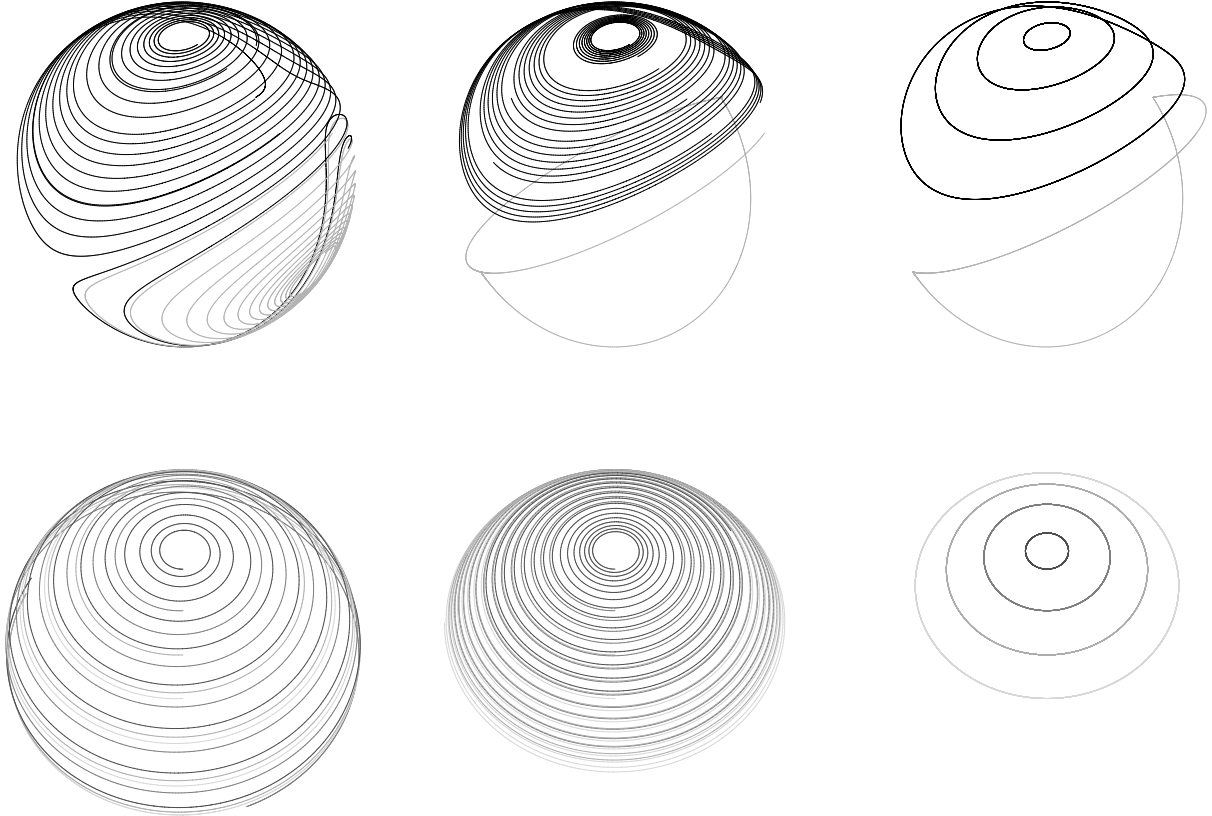
10

Figure 1: Sample trajectories computed over the interval $[0, 200]$ using the time step $\Delta t = 0.1$ and, left to right, ξ_b^{Eul}, ξ_o^{Eul}, and ξ_i^{Eul}. The upper row is computed using the inertia tensor of a triaxial rigid body, while the lower row is computed for an axisymmetric rigid body.

for ten randomly generated initial conditions and inertia tensors. We provide the errors in the energy for time steps $\Delta t = 10$, 1, .1, and .01. Finally, we consider an axisymmetric rigid body, with two randomly generated eigenvalues; recall that in this situation the choice ξ_i^{Eul}, combined with the use of the true exponential as the algorithmic exponential, yields the exact solution.

If the rigid body is axisymmetric, then all true trajectories consist either of equilibria (the 'poles' and the 'equator') or of steady rotations about the axis of symmetry. In this situation the forward Euler method with the infinitesimal update ξ_i^{Eul} associated to second order orbit approximation yields the true trajectories when the true exponential map is used as Exp. If the Cayley transform is used as the algorithmic exponential, then the orbits are captured exactly, but the algorithmic trajectories differ from the true trajectories by a time reparametrization. Note that the basic generator ξ_b^{Eul} and the orthogonal generator ξ_o^{Eul} yield only first order orbit approximations even in the axisymmetric case.

We used four different second order updates, corresponding to four different versions of the Heun method:

	$\xi_{\mathrm{n}}^{\mathrm{Eul}}$	$\xi_{\mathrm{o}}^{\mathrm{Eul}}$	$\xi_{\mathrm{i}}^{\mathrm{Eul}}$
Triaxial	$6.37\ 10^{-2}$	$2.23\ 10^{-2}$	$4.60\ 10^{-6}$
Axisymmetric	$2.46\ 10^{-1}$	$1.46\ 10^{-1}$	$7.38\ 10^{-14}$

Figure 2: Maximum energy error over the trajectories given in figure 1, with time step $\Delta t = .1$.

	Triaxial			Axisymmetric		
Δt	$\xi_{\mathrm{b}}^{\mathrm{Eul}}$	$\xi_{\mathrm{o}}^{\mathrm{Eul}}$	$\xi_{\mathrm{i}}^{\mathrm{Eul}}$	$\xi_{\mathrm{b}}^{\mathrm{Eul}}$	$\xi_{\mathrm{o}}^{\mathrm{Eul}}$	$\xi_{\mathrm{i}}^{\mathrm{Eul}}$
10	$9.04\ 10^{-2}$	$1.81\ 10^{-2}$	$4.56\ 10^{-3}$	$3.09\ 10^{-3}$	$3.70\ 10^{-3}$	0.00
1	$8.53\ 10^{-2}$	$8.54\ 10^{-3}$	$3.68\ 10^{-5}$	$1.55\ 10^{-3}$	$6.17\ 10^{-4}$	0.00
.1	$9.06\ 10^{-3}$	$1.44\ 10^{-3}$	$3.55\ 10^{-7}$	$2.33\ 10^{-4}$	$6.31\ 10^{-5}$	0.00
.01	$7.09\ 10^{-4}$	$1.62\ 10^{-4}$	$3.57\ 10^{-9}$	$2.43\ 10^{-5}$	$6.32\ 10^{-6}$	0.00

Figure 3: Average energy errors over ten sample runs with randomly generated initial conditions and inertia tensors, using versions of the forward Euler method and integrated over the interval $[0, 100]$.

- The standard Heun infinitesimal update with the 'basic' generator:

$$\xi_{\mathrm{b}}^{\mathrm{RK2}} = \tfrac{1}{2}\big(\omega_{\mathrm{b}}(m) - \omega_{\mathrm{b}}(\mathrm{cay}(\omega_{\mathrm{b}}(m)) \cdot m)\big)$$

- The standard Heun infinitesimal update with the orthogonal generator:

$$\xi_{\mathrm{o}}^{\mathrm{RK2}} = \tfrac{1}{2}\big(\omega(m) - \omega(\mathrm{cay}(\omega(m)) \cdot m)\big)$$

- The improved Heun infinitesimal update with the basic generator:

$$\xi_{\mathrm{ib}}^{\mathrm{RK2}}(m, \Delta t) := \xi_{\mathrm{b}}^{\mathrm{RK2}}(m, \Delta t) + \Delta t^3\, \sigma_{\mathrm{ib}}(m)\, m$$

where

$$\sigma_{\mathrm{ib}}(m) := \frac{\langle\, \mathbb{J}^n, u\,\rangle}{\langle\, \mathbb{J}^d, u\,\rangle}, \qquad \text{with} \qquad \begin{cases} u_j := (m_k m_\ell)^2 \\ J_j^n := -I_j(I_k + I_\ell)(I_k - I_\ell)^2 \\ J_j^d := 4\,I_1 I_2 I_3 I_j^2 (I_k - I_\ell)^2 \end{cases}$$

for any cyclic permutation (j, k, ℓ) of $(1, 2, 3)$.

- The improved Heun infinitesimal update with the orthogonal generator: $\xi_{\mathrm{io}}^{\mathrm{RK2}}$, where $\xi_{\mathrm{b}}^{\mathrm{RK2}}$ is replaced by $\xi_{\mathrm{o}}^{\mathrm{RK2}}$ and σ_{ib} replaced by an appropriate function σ_{io}. (The function σ_{io} is a rational function in m and the components of the inertia tensor, but is significantly more complicated than σ_{ib}.)

Δt	$\xi_{\text{b}}^{\text{RK2}}$	$\xi_{\text{o}}^{\text{RK2}}$	$\xi_{\text{ib}}^{\text{RK2}}$	$\xi_{\text{io}}^{\text{RK2}}$
	Triaxial			
10	7.42 10^{-2}	1.22 10^{-2}	6.11 10^{-2}	4.99 10^{-2}
1	4.36 10^{-3}	1.06 10^{-4}	9.61 10^{-4}	5.67 10^{-7}
.1	5.03 10^{-6}	1.10 10^{-7}	1.11 10^{-8}	2.14 10^{-11}
.01	5.01 10^{-9}	1.10 10^{-10}	2.94 10^{-13}	9.07 10^{-15}
	Axisymmetric			
10	2.70 10^{-3}	4.86 10^{-4}	6.34 10^{-3}	1.83 10^{-3}
1	6.99 10^{-5}	8.59 10^{-7}	4.07 10^{-6}	2.44 10^{-9}
.1	7.50 10^{-8}	8.64 10^{-10}	4.28 10^{-11}	2.72 10^{-14}
.01	7.50 10^{-11}	8.78 10^{-13}	0.00	3.81 10^{-14}

Figure 4: Average maximum energy errors over ten sample runs with randomly generated initial conditions and inertia tensors, integrated over the interval [0, 100] using versions of the Heun method.

As table 4 shows, the infinitesimal updates $\xi_{\text{b}}^{\text{RK2}}$ and $\xi_{\text{o}}^{\text{RK2}}$ yield algorithms that preserve the energy to third order in the time step. Hence the isotropy improvement functions σ_{ib} and σ_{io} are scaled by Δt^3, to yield a modification of the update, and hence the energy error, at fourth order. More generally, any choice of generator $\omega_{\text{gen}}(m) = \omega(m) + \sigma(m)\, m$ for some scalar function σ satisfying $(X \cdot \nabla)\sigma \equiv 0$ determines a Heun infinitesimal update $\xi_{\text{gen}}^{\text{RK2}}$ satisfying

$$H(\text{cay}(\xi_{\text{gen}}^{RK2}(m, \Delta t))) - H(m) = \mathcal{O}(\Delta t^4).$$

(The normal components of ω_{b} and ω are constant on the orbits of the rigid body flow.) The function σ_{ib} (respectively σ_{io}) was determined by symbolically computing the energy discretization for the infinitesimal update $\xi_{\text{b}}^{\text{RK2}}(m) + \sigma\, m$ (respectively $\xi_{\text{o}}^{\text{RK2}}(m) + \sigma\, m$) to fourth order with an unspecified parameter σ, then solving the resulting affine equation for σ. However, the general discretization error approach described in §3 would have yielded the same results. See §7 for a more detailed and general discussion of the use of conservation laws to determine update improvements.

If the body is axisymmetric, $\xi_{\text{ib}}^{\text{RK2}}$ preserves the energy to fifth order. As table 4 shows, some of these algorithms appear to have better *global* energy capture than the single step discretization energy error analysis (which we carried out symbolically using *Mathematica*) would suggest. Plots of the energy errors in the sample integrations, with randomly generated initial conditions and inertia tensors, show that the energy oscillates about a slow drift away from the correct value.

We consider six fourth order geometric methods. Four utilize a series expansion for the generator along a solution curve, while the other two use the RKMK4 algorithm of Munthe-Kaas, [32, 15], (with the Cayley transform as the algorithmic exponential). Using the Cayley transform, the map $\xi_{\text{b}}^{\text{sy4}}$ determined by the basic generator ω_{b} for the rigid body system on S^2

Δt	$\xi_{\rm b}^{\rm sy4}$	$\xi_{\rm o}^{\rm sy4}$	$\xi_{\rm ib}^{\rm sy4}$	$\xi_{\rm io}^{\rm sy4}$	$\xi_{\rm b}^{\rm RK4}$	$\xi_{\rm o}^{\rm RK4}$
	Triaxial					
10	$5.67\ 10^{-3}$	$3.37\ 10^{-2}$	$6.06\ 10^{-3}$	$1.90\ 10^{-2}$	$6.07\ 10^{-2}$	$3.59\ 10^{-3}$
1	$3.62\ 10^{-6}$	$2.72\ 10^{-7}$	$1.55\ 10^{-6}$	$1.68\ 10^{-7}$	$2.15\ 10^{-5}$	$1.46\ 10^{-7}$
.1	$3.56\ 10^{-10}$	$7.45\ 10^{-12}$	$3.56\ 10^{-10}$	$7.45\ 10^{-12}$	$2.97\ 10^{-10}$	$3.06\ 10^{-12}$
	Axisymmetric					
10	$1.10\ 10^{-3}$	$2.42\ 10^{-5}$	$1.10\ 10^{-3}$	$3.54\ 10^{-5}$	$4.81\ 10^{-4}$	$1.38\ 10^{-5}$
1	$3.24\ 10^{-8}$	$2.67\ 10^{-7}$	$3.24\ 10^{-8}$	$2.67\ 10^{-7}$	$8.17\ 10^{-8}$	$1.70\ 10^{-10}$
.1	$3.28\ 10^{-13}$	$2.67\ 10^{-10}$	$2.28\ 10^{-13}$	$2.67\ 10^{-10}$	$8.41\ 10^{-13}$	$4.85\ 10^{-15}$

Figure 5: Average maximum energy errors over ten sample runs with randomly generated initial conditions and inertia tensors, integrated over the interval [0, 100] using several fourth order methods.

is given by

$$\xi_{\rm b}^{\rm sy4}(m, \Delta t) = \sum_{j=1}^{4} \frac{\Delta t^j}{j!} \omega_{\rm b}^{(j-1)}(m) + \frac{\Delta t^3}{12}\left(\|\omega_{\rm b}(m)\|^2 \omega_{\rm b}(m) + \dot{\omega}_{\rm b}(m) \times \omega_{\rm b}(m)\right) +$$

$$+ \frac{\Delta t^4}{4!}\left(\|\omega_{\rm b}(m)\|^2 \dot{\omega}_{\rm b}(m) + \ddot{\omega}_{\rm b}(m) \times \omega_{\rm b}(m) + 2\langle \dot{\omega}_{\rm b}(m), \omega_{\rm b}(m)\rangle \omega_{\rm b}(m)\right),$$

where $\omega_{\rm b}^{(j)}(m) = \frac{\partial^j}{\partial t^j}\omega_{\rm b}(\mathcal{F}_t(m))|_{t=0}$. The corresponding algorithm for the rigid body using the orthogonal generator ω is

$$\xi_{\rm o}^{\rm sy4}(m, \Delta t) = \sum_{j=1}^{4} \frac{\Delta t^j}{j!} \omega^{(j-1)}(m) + \frac{\Delta t^3}{12}\left(\|\omega(m)\|^2 \omega(m) + \langle \ddot{\omega}(m), m\rangle m\right) +$$

$$+ \frac{\Delta t^4}{8}\langle \dot{\omega}(m), \omega(m)\rangle \omega(m).$$

The infinitesimal updates $\xi_{\rm b}^{\rm sy4}$ and $\xi_{\rm o}^{\rm sy4}$ can be modified by the addition of an appropriate multiple of the argument m to yield an additional order of energy, and hence orbit, capture. The scalar improvement functions, which are rational functions of the components of m and the inertia tensor, were determined by symbolic calculation. Note that the energy improvement term for the generator $\xi_{\rm o}^{\rm sy4}$ is identically zero if the body is axisymmetric; hence the results generated by $\xi_{\rm o}^{\rm sy4}$ and $\xi_{\rm io}^{\rm sy4}$ coincide in this case.

The last column of table 6 provides a comparison of the accuracy and efficiency of the representative updates tested here. The orbit (energy) accuracy of a given simulation was taken to be the maximum energy error over the total simulation (100 units of time), while the CPU time, in seconds was measured using the *Mathematica* `Timing` function. The same randomly generated initial conditions and inertia tensors were used in all of the numerical simulations.

The efficiency comparisons given in the last column of table 6 and in figure 7 suggest some interesting directions of future investigation. These comparisons show that, for the rigid body

14

Inf. update	Discret. error order	Energy error order	Relative CPU time	Rel. CPU within class	CPU/orbit acc. (ln/ln approx.)
ξ_b^{Eul}	1	1	1	1	1.22 x - 7.58
ξ_o^{Eul}	1	1	1.18	1.18	1.40 x - 10.8
ξ_i^{Eul}	1	2	1.45	1.45	.491 x - 8.24
ξ_b^{RK2}	2	3	2.10	1	.402 x - 5.72
ξ_o^{RK2}	2	3	2.42	1.15	.333 x - 5.78
ξ_{ib}^{RK2}	2	4	2.83	1.35	.248 x - 4.95
ξ_{io}^{RK2}	2	4	17.14	8.17	.235 x - 4.02
ξ_b^{sy4}	4	4	6.99	1.21	.267 x - 5.24
ξ_o^{sy4}	4	4	5.80	1	.233 x - 5.18
ξ_{ib}^{sy4}	4	5	8.15	1.40	.249 x - 4.97
ξ_{io}^{sy4}	4	5	11.24	1.94	.230 x - 4.70
ξ_b^{RK4}	4	4	12.08	2.08	.247 x - 4.22
ξ_o^{RK4}	4	4	13.19	2.27	.249 x - 4.91

Figure 6: Average convergence rates and CPU times. The first column gives the conventional discretization error order of convergence. The second gives the order of convergence of the energy error; note that this rate is the same for both the single–step and global errors. The third column gives the average CPU times relative to the infinitesimal update ξ_b^{Eul}. The fourth column gives the average CPU times relative to the fastest algorithm of that order of convergence of the conventional discretization error. The final column gives the affine approximations to the ln/ln plots of the average CPU times as functions of the average energy (and hence orbit) errors. All averages were taken over the same set of ten randomly generated initial conditions and inertia tensors.

system, the most dramatic gains in efficiency of orbit capture with the use of an improvement term occur for low order methods. This suggests that such improvements may be of particular value in situations, e.g. stiff systems, in which low order implicit methods are the methods of choice. Note that the improved Euler method is actually the most efficient method for low orbit accuracy simulations, while the (unimproved) symbolically computed fourth order update using the orthogonal generator is the most efficient of the methods tested for high accuracy simulations. The affine approximations of the CPU time/orbit accuracy graphs for these two methods intersect at a global orbit error of approximately $7.16\,10^{-6}$.

The significant differences in run times between the different fourth order methods illustrate some of the design issues involved in geometric integration. The RKMK4 methods, like their conventional counterparts, have the advantage of great versatility, but the algorithmic exponential evaluations and pullback corrections required at each stage result in relatively high computation costs. (Note that the Heun methods tested here do not require pullback corrections to achieve overall second order accuracy; the Heun infinitesimal updates are simply the averages of the generator evaluations at the current point and an Euler update of the current

Eul, b
Eul, o
Eul, i

RK2, b
RK2, o
RK2, ib

sy4, b
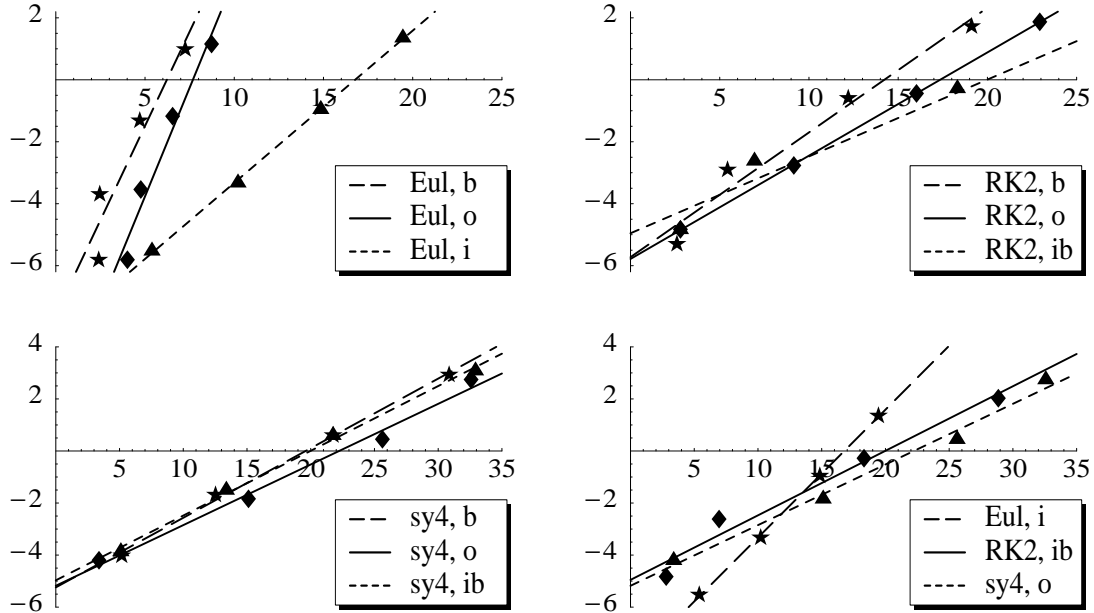sy4, o
sy4, ib

Eul, i
RK2, ib
sy4, o

Figure 7: CPU time as a function of orbit accuracy (ln/ln plots). The upper lefthand plot gives the data for the Euler methods, with infinitesimal updates $\xi_{\mathrm{b}}^{\mathrm{Eul}}$, $\xi_{\mathrm{o}}^{\mathrm{Eul}}$, and $\xi_{\mathrm{i}}^{\mathrm{Eul}}$; the upper righthand plot gives the data for three of the Heun updates: $\xi_{\mathrm{b}}^{\mathrm{RK2}}$, $\xi_{\mathrm{o}}^{\mathrm{RK2}}$, and $\xi_{\mathrm{ib}}^{\mathrm{RK2}}$; the lower lefthand plot gives the data for three of the symbolically derived fourth order updates: $\xi_{\mathrm{b}}^{\mathrm{sy4}}$, $\xi_{\mathrm{o}}^{\mathrm{sy4}}$, and $\xi_{\mathrm{ib}}^{\mathrm{sy4}}$. The lower righthand plot gives the data for the most efficient updates of each order (in overall accuracy): $\xi_{\mathrm{i}}^{\mathrm{Eul}}$, $\xi_{\mathrm{ib}}^{\mathrm{RK2}}$, and $\xi_{\mathrm{o}}^{\mathrm{sy4}}$.

point.) The infinitesimal updates ξ^{sy4} are significantly faster, with equal or superior rates of convergence, but have the disadvantage that these updates are specific to the rigid body system; analogous generator expansions and isotropy improvement terms must derived for any new dynamical system.

Our results do not suggest any clear reason to favor the basic generator over the orthogonal generator, or vice versa. In some algorithms, the basic generator yields the more efficient update; in others, the orthogonal generator gives the more efficient scheme. Note, for example, that the run times for $\xi_{\mathrm{io}}^{\mathrm{RK2}}$ are substantially longer than those for $\xi_{\mathrm{ib}}^{\mathrm{RK2}}$, while the symbolically computed fourth order update $\xi_{\mathrm{o}}^{\mathrm{sy4}}$ is both faster and more efficient than the corresponding update $\xi_{\mathrm{b}}^{\mathrm{sy4}}$ using the basic generator. Some of these differences are presumably due to suboptimal implementation of the schemes; the dramatic difference in speed between $\xi_{\mathrm{io}}^{\mathrm{RK2}}$ and $\xi_{\mathrm{ib}}^{\mathrm{RK2}}$ appears to be due to the relative complexity of σ_{io} in comparison to σ_{ib}. Algebraic manipulation of of the more complicated expressions appearing in some of the updates might lead to significant improvements in efficiency.

5   Higher Order Approximation of Orbits.

Using our experience with the rotation group acting on the sphere as a guide, we now turn to the general case of interest here, in which a Lie group G acts on $M \subset \mathcal{V}$. We assume, as before, that the exact flow $\mathcal{F}_t$ and its approximation $\widetilde{\mathcal{F}}_t$ can be expressed in the form

$$\mathcal{F}_t(m) = \text{Exp}(\psi(m,t)) \cdot m \qquad \text{and} \qquad \widetilde{\mathcal{F}}_t(m) = \text{Exp}(\xi(m,t)) \cdot m. \tag{5.1}$$

The maps $\psi, \xi \colon M \times I \to \mathfrak{g}$ are assumed to be sufficiently smooth, with

$$\psi(m,0) \equiv 0 \equiv \xi(m,0).$$

Exp can be taken to be standard exponentiation, or an algorithmic alternative. The following lemma is immediate from the definition of the tangent space.

Lemma 3   *Let M be an embedded submanifold of a vector space $\mathcal{V}$. Let $\varphi(s), \psi(s) \subset M$ be two continuous curves defined for $s \geq 0$, having a common endpoint $\varphi(0) = \psi(0) = m \in M$. If*

$$\lim_{s \to 0^+} \frac{\varphi(s) - \psi(s)}{s} = Z \tag{5.2}$$

exists, then $Z \in TM|_m$ is a tangent vector at m.

Note that we do not require either individual curve to be differentiable at m. One well-known example is to let

$$\varphi(s) = (\exp \sqrt{s}\, X)(\exp \sqrt{s}\, Y)m, \qquad \varphi(s) = (\exp \sqrt{s}\, Y)(\exp \sqrt{s}\, X)m,$$

where X, Y are two vector fields tangent to M. In this case, the right hand side of (5.2) yields the Lie bracket of the vector fields $Z = [\,X, Y\,]$, which is necessarily tangent to M.

Corollary 4   *If $\widetilde{\mathcal{F}}_t \colon M \to M$ is an order $n-1 \geq 0$ approximation of $\mathcal{F}_t \colon M \to M$, then*

$$\frac{d^n}{dt^n}\,[\,\widetilde{\mathcal{F}}_t(m) - \mathcal{F}_t(m)\,]\,\bigg|_{t=0} \in TM|_m$$

for all $m \in M$.

Proof: By the definition of an order $n-1$ approximation,

$$Z = \frac{d^n}{dt^n}\,[\,\widetilde{\mathcal{F}}_t(m) - \mathcal{F}_t(m)\,]\,\bigg|_{t=0} = n!\, \lim_{t \to 0} \frac{\widetilde{\mathcal{F}}_t(m) - \mathcal{F}_t(m)}{t^n}$$

exists at each m. Now, merely replace $s = t^n$ and use Lemma 3.   ∎

For a fixed $m \in M$, we can expand (5.1) as in powers of t, so

$$\psi(m,t) = t^\ell\, v_\ell + \mathcal{O}(t^{\ell+1}), \qquad \text{so that} \qquad \mathcal{F}_t(m) = m + t^\ell\, w_\ell + \mathcal{O}(t^{\ell+1})$$

for $v_\ell = v_\ell(m) \neq 0$. We call $\ell = \ell(m)$ the *leading order* and $w_\ell = w_\ell(m) \neq 0$ the *leading term* of $\mathcal{F}_t$ at m. Note that, by Lemma 3, $w_\ell \in TM|_m$. In most applications, including flows of autonomous vector fields X away from equilibrium points, $\ell \equiv 1$ and $v_1 = w_1 = X(m)$.

Let
$$\mathfrak{g}_m^{(n)} = \{ \, \zeta(t) = \zeta_1\, t + \cdots + \zeta_n\, t^n \mid \zeta_1, \ldots, \zeta_n \in \mathfrak{g}_m \, \} \simeq \mathfrak{g}_m^{\times n}$$
denote the space of $\mathfrak{g}_m$–valued polynomials of degree n without constant term. Given flows as in (5.1), let us define

$$K_m^{(n)} = \left\{ \, \zeta \in \mathfrak{g}_m^{(n)} \;\middle|\; \begin{array}{c} \mathrm{Exp}(\xi(m,t) + \zeta(t)) \cdot m \ \ \text{is an order } n-1 \\ \text{approximation of } \mathcal{F}_t(m) \text{ at } m \end{array} \right\}. \tag{5.3}$$

We define the map $\chi_m^{(n)} \colon K_m^{(n)} \to TM|_m$ by

$$\chi_m^{(n)}(\zeta) = \frac{d^n}{dt^n} \left[\, \mathrm{Exp}(\xi(m,t) + \zeta(t)) \cdot m - \mathcal{F}_t(m) \, \right] \bigg|_{t=0}. \tag{5.4}$$

Corollary 4 ensures that $\chi_m^{(n)}(\zeta)$ lies in $TM|_m$. Note that if Exp is at least a first order approximation of the true exponential map, then $\chi_m^{(1)}$ is constant.

Our main result on increasing the order of approximation follows. Under a certain range condition, (5.5), we show that one can, by adjusting the isotropy element and the time parametrization, increase the order of accuracy of the approximate flow. The range condition formalizes the notion of "sufficient isotropy" mentioned in the introduction.

Theorem 5 *Suppose $\widetilde{\mathcal{F}}_t(m) = \mathrm{Exp}(\xi(m,t)) \cdot m$ is an order $n-1$ approximation of the flow $\mathcal{F}_t(m) = \mathrm{Exp}(\psi(m,t)) \cdot m$. Assume further that the leading term of the flow $\mathcal{F}_t$ satisfies the range condition*
$$c\, w_\ell(m) \in \mathrm{Range}\ \chi_m^{(n)} \qquad \text{for some} \qquad c \in \mathbb{R}. \tag{5.5}$$

Then there exists

    a) a Lie algebra-valued polynomial $\sigma(m,t) \in K_m^{(n)}$, and

    b) an invertible time reparametrization $\tau \colon (-T, T) \to \mathbb{R}$ for some $T > 0$, satisfying

$$\tau(t) - t = \mathcal{O}(t^{n-\ell+1})$$

such that the map
$$\widetilde{\mathcal{F}}_t^{\,\sigma}(m) := \mathrm{Exp}(\xi(m,t) + \sigma(m,t)) \cdot m \tag{5.6}$$

is an order n approximation of the reparametrized flow $\mathcal{F}_{\tau(t)}$ at m.

Proof: If $\widetilde{\mathcal{F}}_t$ is already an order n approximation of $\mathcal{F}_t$, then there is nothing to show. By our hypothesis (5.5), there exists $\sigma \in K_m^{(n)}$ and $c \in \mathbb{R}$ such that $\chi_m^{(n)}(\sigma) = c\, w_\ell$. Thus for any choice of τ satisfying
$$\tau(m,t) = t + \frac{c\, t^{n-\ell+1}}{(n-\ell+1)!} + \mathcal{O}(t^{n-\ell+2}),$$

we have

$$\frac{d^n}{dt^n}\left[\widetilde{\mathcal{F}}_t^{\sigma}(m) - \mathcal{F}_{\tau(t)}(m)\right]\Bigg|_{t=0} = \chi_m^{(n)}(\sigma) - \tau^{(n-\ell+1)}(m,0)\, w_\ell = 0.$$

For such choices $\widetilde{\mathcal{F}}_t^{\sigma}(m)$ is an order n approximation of $\mathcal{F}_{\tau(t)}(m)$. ∎

We now consider the order conditions, up to order three, for any algorithm of the form (5.6). Let $\mathcal{F}_t$ denote the flow of the system $\dot m = \omega(m)$; then

$$\mathcal{F}_t(m) \;=\; \left(1 + t\,\omega + \frac{t^2}{2}\left(\omega^2 + \dot\omega\right) + \frac{t^3}{6}\left(\omega^3 + 3\,\dot\omega\omega + [\omega,\dot\omega] + \ddot\omega\right) + \mathcal{O}(t^4)\right)\cdot m.$$

Let $\widetilde{\mathcal{F}}_t : M \to M$ be an algorithmic update of the form

$$\widetilde{\mathcal{F}}_t(m) = \mathrm{Exp}(\xi(m,t))\cdot m,$$

where Exp is an n–th order approximation to the true exponential and

$$\xi(m,t) = \sum_{j=0}^{\infty} \frac{t^j}{j!}\,\xi_j(m).$$

Then

$$\widetilde{\mathcal{F}}_t(m) \;=\; \left(1 + t\,\xi_1 + \tfrac{t^2}{2}\left(\xi_1^2 + \xi_2\right) + \tfrac{t^3}{6}\left(\xi_1^3 + \tfrac{3}{2}(\xi_2\xi_1 + \xi_1\xi_2) + \xi_3\right) + \mathcal{O}(t^4)\right)\cdot m.$$

The algorithm $\widetilde{\mathcal{F}}_t$ is consistent if and only if $\xi_1 = \omega + \sigma_1$ for some $\sigma_1 \in \mathfrak{g}_m$. If the algorithm is consistent, then the local discretization error is

$$\frac{\widetilde{\mathcal{F}}_t(m) - \mathcal{F}_t(m)}{t} = \left(\frac{t}{2}(\xi_1\,\omega + \xi_2 - \dot\omega) + \mathcal{O}(t^3)\right)\cdot m.$$

Thus the algorithm is second order accurate if and only if

$$\xi_2 = \dot\omega + [\,\sigma_1,\omega\,] + \sigma_2$$

for some $\sigma_2 \in \mathfrak{g}_m$. If the second order condition is satisfied, then the local discretization error is

$$\frac{\widetilde{\mathcal{F}}_t(m) - \mathcal{F}_t(m)}{t} = \left(\frac{t^2}{6}\left(\xi_3 - \ddot\omega + \omega\,\sigma_1\omega - \tfrac{1}{2}\sigma_1\omega^2 + \tfrac{3}{2}\sigma_2\,\omega + \tfrac{3}{2}\sigma_1\,\dot\omega + \tfrac{1}{2}[\omega,\dot\omega] - \tfrac{1}{2}\sigma_1^2\,\omega\right) + \mathcal{O}(t^3)\right)\cdot m.$$

Thus the algorithm is third order accurate if and only if

$$\xi_3 = \ddot\omega + \tfrac{1}{2}(\,[\,\dot\omega,\omega\,] + +[\,\sigma_1,[\,\sigma_1,\omega\,]\,] + [\,\omega,[\,\omega,\sigma_1\,]\,] + 3\,[\,\omega,\sigma_2\,] + 3\,[\,\dot\omega,\sigma_1\,]\,) + \sigma_3$$

for some $\sigma_3 \in \mathfrak{g}_m$. Clearly, one can continue this process, obtaining general consistency and accuracy criteria for such algorithms.

6   Restrictions of Representations.

The most important case is when G acts via a linear representation $\rho\colon G \to \mathrm{GL}(\mathcal{V})$ on a vector space $\mathcal{V}$, and our submanifold $M \subset \mathcal{V}$ is a G-orbit. The simplest non-free case is the action of the rotation group $\mathrm{SO}(3)$ on $\mathbb{R}^3$ that formed the basis of our examples on the sphere $S^2 \subset \mathbb{R}^3$. In this section, we analyze our constructions in the general representation-theoretic context.

We use $\rho\colon \mathfrak{g} \to \mathfrak{X}(M)$ to denote the effect of the representation on the Lie algebra. The effect of the map (5.4) on a given Lie algebra-valued polynomial of the form

$$\zeta(t) = \zeta_{n-\ell} t^{n-\ell} + \cdots + \zeta_n t^n \in K_m^{(n)} \qquad \text{for} \qquad \zeta_{n-\ell}, \ldots, \zeta_n \in \mathfrak{g}_m$$

is

$$\chi_m^{(n)}(\zeta) = \rho(\zeta_{n-\ell}) z_\ell + z_n - w_n = \rho(\zeta_{n-\ell}) w_\ell + z_n - w_n.$$

Thus, the range condition (5.5) is satisfied if there exists $\zeta_{n-\ell} \in \mathfrak{g}_m$ and $c \in \mathbb{R}$ satisfying

$$\rho(\zeta_{n-\ell}) w_\ell + c\, w_\ell = w_n - z_n.$$

The latter condition is satisfied for *any* order $n-1$ approximation of $\mathcal{F}_t$ provided

$$\rho(\mathfrak{g}_m) \cdot w_\ell + \mathrm{span}\{\, w_\ell \,\} \;\supseteq\; TM|_m, \tag{6.1}$$

where $\rho(\mathfrak{g}_m) \cdot v = \{\, \rho(\zeta) v \,|\, \zeta \in \mathfrak{g}_m \,\}$. If $\mathcal{F}_t$ is the flow of an autonomous vector field X on M, then $w_1 = X(m) \neq 0$, unless m is an equilibrium. Hence in this case (6.1) takes the form

$$\mathfrak{g}_m \cdot X(m) + \mathrm{span}\{\, X(m) \,\} = TM|_m. \tag{6.2}$$

Since G acts transitively on M, (6.1) can be rephrased in terms of the Lie algebra as follows: Choose $\psi(m) \in \mathfrak{g}$ such that $\rho(\psi(m)) m = w_\ell$. Then (6.1) is satisfied if and only if

$$[\,\mathfrak{g}_m, \psi(m)\,] + \mathrm{span}\{\, \psi(m) \,\} + \mathfrak{g}_m = \mathfrak{g}, \tag{6.3}$$

where $[\,\mathfrak{g}_m, \psi(m)\,] = \{\, [\,\zeta, \psi(m)\,] \,|\, \zeta \in \mathfrak{g}_m \,\}$. Note that the latter condition does not depend on the choice of $\psi(m)$.

Let us introduce a class of subalgebras for which the range condition is immediate.

Definition 6   A subalgebra $\mathfrak{h} \subset \mathfrak{g}$ is called *full* if

$$[\,\mathfrak{h}, \eta\,] + \mathrm{span}\{\, \eta \,\} + \mathfrak{h} = \mathfrak{g} \qquad \text{for every} \qquad \eta \in \mathfrak{g} \setminus \mathfrak{h}. \tag{6.4}$$

It would be interesting to classify Lie algebra and subalgebra pairs that satisfy the full criterion (6.4).

Lemma 7   *If the isotropy subalgebra $\mathfrak{g}_m \subset \mathfrak{g}$ is a full subalgebra, then (6.3) follows, and hence the range condition (5.5)is automatically satisfied for all order $n-1$ approximations of the flow.*

Note that if $\mathfrak{h}$ is a full subalgebra, so is any conjugate subalgebra $\widetilde{\mathfrak{h}} = \operatorname{Ad} g \cdot \mathfrak{h}$. This implies that (6.3) only needs to be checked at a single point m in the orbit.

Example: Suppose $G = \operatorname{SO}(k)$ acts on $\mathcal{V} = \mathbb{R}^k$ via the usual representation. The $\operatorname{SO}(k)$ orbits are the spheres, and we concentrate on the unit sphere $M = S^k \subset \mathbb{R}^k$. The isotropy subgroup of any $m \neq 0$ is $G_m \approx \operatorname{SO}(k-1)$. By the preceding remark, to verify that the isotropy subalgebras are full, it suffices to work at the point $m = e_k = (0, \ldots, 0, 1)$, whose isotropy subgroup is

$$ H = G_m = \left\{ \left. \begin{pmatrix} U & 0 \\ 0 & 1 \end{pmatrix} \; \right| \; U \in \operatorname{SO}(k-1) \right\}. $$

Any $\eta \in \mathfrak{so}(k)$ can be written as $\eta = \zeta + z e_k^T - e_k z^T$ where $\zeta \in \mathfrak{h}$. If $\xi \in \mathfrak{h}$, then

$$ [\,\xi, \eta\,] = [\,\xi, \zeta\,] + (\xi\, z) e_k^T - e_k (\xi\, z)^T. $$

Since $so(k-1) \cdot x + \operatorname{span}\{\, x \,\} = \mathbb{R}^{k-1}$ for any nonzero $x \in \mathbb{R}^{k-1}$, the fullness condition (6.4) is satisfied. Thus, for the usual action of $\operatorname{SO}(k)$ on S^{k-1}, the range condition (5.5) holds for any vector field and any point $m \in S^{k-1}$.

Remark: If we consider the action $Q \cdot S = Q\, S\, Q^T$ of $\operatorname{SO}(k)$ on the manifold $\operatorname{Sym}(k)$ of $k \times k$ symmetric matrices, then (6.3) is satisfied for any matrix with $k - 1$ equal eigenvalues.

7   Conservation Laws.

A structure–preserving algorithm is one that exactly captures one or more features of the exact flow $\mathcal{F}_t$. In this section we consider the role of isotropy in designing algorithms that preserve integrals or conservation laws of the flow. For simplicity, we will assume that the conserved quantity depends only upon the coordinates of our trajectory, and not on any derivative.

Definition 8   A (vector-valued) *conservation law* or *first integral* of a flow $\mathcal{F}_t(m)$ is a smooth map $F \colon M \to W$ onto a vector space W such that

$$ F(\mathcal{F}_t(m)) = F(m) \qquad \text{for all} \qquad m \in M. $$

In general, a conservation law will only be preserved by our approximate flow up to a certain order. To make this precise,

Definition 9   The first integral $F \colon M \to W$ is said to be *conserved to order* $n - 1$ by the approximating flow $\widetilde{\mathcal{F}}_t(m) = \operatorname{Exp}(\xi(m, t)) \cdot m$ on a subset $K \subset M$ provided

$$ F(\widetilde{\mathcal{F}}_t(m)) - F(m) = \mathcal{O}(t^n), \qquad \text{for all} \qquad m \in K. \tag{7.1} $$

The Mather Division Theorem, [14], implies that satisfaction of (7.1) at a point m_0 is equivalent to the existence of $\epsilon > 0$ and a smooth map $\hat{F} \colon M \times (-\epsilon, \epsilon) \to W$ satisfying

$$ F(\widetilde{\mathcal{F}}_t(m)) - F(m) = t^n \hat{F}(m, t) $$

for all (m, t) in a neighborhood of $(m_0, 0)$.

The main result of this section is that, subject to a certain nondegeneracy hypothesis, if F is conserved to order $n-1$ by $\widetilde{\mathcal{F}}_t(m)$, then we can find a suitable isotropy correction $\sigma(m, t) \in \mathfrak{g}_m$ such that F is conserved exactly by the modified approximation

$$\breve{\mathcal{F}}_t(m) := \mathrm{Exp}(\xi(m, t) + t^{n-1}\,\sigma(m, t)) \cdot m. \tag{7.2}$$

Therefore, one can use the isotropy to retain at least the original order of overall approximation while conserving the first integral exactly. This requires, as above, that our algorithmic exponential approximate the true exponential to at least first order, namely

$$\left.\frac{d}{d\epsilon}\mathrm{Exp}(\epsilon\,\eta) \cdot m\right|_{\epsilon=0} = \gamma(\eta)(m) \qquad \text{for all} \qquad \eta \in \mathfrak{g}. \tag{7.3}$$

In contrast to the treatment given in Section 5, we do not require any explicit information about the Taylor expansion of the true flow. Rather, we use the Implicit Function Theorem to show that under appropriate nondegeneracy conditions an isotropy element can be found that yields a structure–preserving scheme.

Proposition 10 *Suppose that the first integral $F\colon M \to W$ is conserved to order $n-1$ by the approximate flow $\widetilde{\mathcal{F}}_t(m, 0)$, where*

$$\widetilde{\mathcal{F}}_t(m, \zeta) := \mathrm{Exp}(\xi(m, t) + t^{n-1}\,\zeta) \cdot m \qquad \text{for } \zeta \in \mathfrak{g}_m, \tag{7.4}$$

for all $m \in K$ in a compact set $K \subseteq M$. If the maps

$$\chi_m(\zeta) := \left.\frac{1}{n!}\frac{d^n}{dt^n}D(F \circ \widetilde{\mathcal{F}}_t)(m, 0)(0, \zeta)\right|_{t=0} \tag{7.5}$$

map $\mathfrak{g}_m$ onto W, then there exists $T > 0$ and an isotropy-valued map $\sigma\colon K \times (-T, T) \to \mathfrak{g}$ with $\sigma(m, t) \in \mathfrak{g}_m$ and

$$F(\widetilde{\mathcal{F}}_t(m, \sigma(m, t))) = F(m) \qquad \text{for all} \qquad (m, t) \in K \times (-T, T). \tag{7.6}$$

Proof: The linearization condition (7.3) on Exp implies that

$$\mathrm{Exp}(t\,\xi(m, t) + t^{n-1}\zeta) \cdot m - \mathrm{Exp}(t\,\xi(m, t)) \cdot m = \mathcal{O}(t^n)$$

for all $\zeta \in \mathfrak{g}_m$. Hence

$$F(\widetilde{\mathcal{F}}_t(m, \zeta)) - F(m) = (F(\widetilde{\mathcal{F}}_t(m, \zeta)) - F(\widetilde{\mathcal{F}}_t(m, 0))) - (F(\widetilde{\mathcal{F}}_t(m, 0)) - F(m)) = \mathcal{O}(t^n)$$

and the Mather Division Theorem implies the existence of smooth maps $\hat{F}_m : \mathfrak{g}_m \times (-\epsilon, \epsilon) \to W$ satisfying

$$F(\widetilde{\mathcal{F}}_t(m, \zeta)) - F(m) = t^n \hat{F}_m(\zeta, t) \qquad \text{and hence} \qquad D\hat{F}_m(0, 0)(\zeta, 0) = \chi_m(\zeta) \tag{7.7}$$

for every $m \in K$ and $\zeta \in \mathfrak{g}_m$.

The nondegeneracy hypothesis (7.5) on χ_m implies that there exists a subspace $V_m \subseteq \mathfrak{g}_m$ such that $\chi_m|_{V_m}$ is an isomorphism. It follows from (7.7) that for each $m \in K$ we can apply the Implicit Function Theorem to the restriction of $\hat{F}_m$ to $V_m \times (-\epsilon, \epsilon)$, finding $T_m > 0$ and a map $\sigma_m\colon (-T_m, T_m) \to V_m$ such that $F(\widetilde{\mathcal{F}}_t(m, \sigma_m(t))) = F(m)$. Taking $T := \min_{m \in K} T_m$ and setting $\sigma(m, t) := \sigma_m(t)$ yields the result. ∎

Corollary 11   *If there exists a compact set $K \subseteq M$ on which*

- $\mathrm{Exp}(t\,\eta) \cdot m - \exp(t\,\eta) \cdot m = \mathcal{O}(t^3)$ *for all $\eta \in \mathfrak{g}$*

- *the map $\widetilde{\mathcal{F}}_t$ given by (7.4) satisfies $\xi(m,0) = 0$ and $F(\widetilde{\mathcal{F}}_t(m,0)) - F(m) = \mathcal{O}(t^n)$*

- *the maps $\zeta \mapsto DF(m)\gamma\left([\,\xi_1(m),\zeta\,]\right)(m)$, where $\xi_1(m) := \frac{d}{dt}\xi(m,t)|_{t=0}$, map $\mathfrak{g}_m$ onto W*

for all $m \in K$, then the conclusions of Proposition 10 hold.

Proof: The Baker–Campbell–Hausdorff formula implies that

$$\exp(t\,\eta) \cdot m = \exp(t\,\eta)\exp(t^{n-1}\zeta) \cdot m = \exp\left(t\,\eta + t^{n-1}\zeta + \frac{t^n}{2}[\,\eta,\zeta\,] + \mathcal{O}(t^{n+1})\right) \cdot m,$$

and hence

$$\exp(t\,\eta + t^{n-1}\zeta) \cdot m = \exp\left(t\,\eta - \frac{t^n}{2}[\,\eta,\zeta\,]\right) \cdot m + \mathcal{O}(t^{n+1})$$

for all $\eta \in \mathfrak{g}$ and $\zeta \in \mathfrak{g}_m$. The Mather Division Theorem and the hypothesis $\xi(\cdot,0) = 0$ imply that there exist $\epsilon > 0$ and a smooth map $\hat{\xi} : K \times (-\epsilon,\epsilon) \to \mathfrak{g}$ such that $\xi(m,t) = t\,\hat{\xi}(m,t)$ for all $m \in K$ and $t \in (-\epsilon,\epsilon)$. Thus, using our hypothesis that Exp is at least a second order approximation to exp, we see that

$$\widetilde{\mathcal{F}}_t(m,\zeta) - \widetilde{\mathcal{F}}_t(m,0) = \exp(t\,\hat{\xi}(m,t) + t^{n-1}\zeta) \cdot m - \exp(t\,\hat{\xi}(m,t)) \cdot m + \mathcal{O}(t^{n+1})$$

$$= \exp\left(t\,\hat{\xi}(m,t) - \frac{t^n}{2}[\,\hat{\xi}(m,t),\zeta\,]\right) \cdot m - \exp(t\,\hat{\xi}(m,t)) \cdot m + \mathcal{O}(t^{n+1}).$$

Therefore,

$$\chi_m(\zeta) = \frac{1}{n!}\frac{d^n}{dt^n}D(F \circ \widetilde{\mathcal{F}}_t)(m,0)(0,\zeta)\bigg|_{t=0}$$

$$= \frac{1}{n!}\frac{d^n}{dt^n}\frac{d}{d\epsilon}F\left(\exp\left(t\,\hat{\xi}(m,t) - \frac{t^n}{2}[\,\hat{\xi}(m,t),\epsilon\,\zeta\,]\right) \cdot m + \mathcal{O}(t^{n+1})\right)\bigg|_{\epsilon=0}\bigg|_{t=0}$$

$$= -\tfrac{1}{2}DF(m)\gamma\left([\,\hat{\xi}(m,0),\zeta\,]\right)(m)$$

maps $\mathfrak{g}_m$ onto W. ∎

Remark: Given $z \in W^*$, set $F_z := F^* z$, i.e. $F_z(m) = z \cdot F(m)$ for all $m \in M$. Let $J : T^*M \to \mathfrak{g}^*$ denote the momentum map associated to the lifted action of G on M, i.e. $J(\alpha_m) \cdot \eta = \alpha_m \cdot \gamma(\eta)(m)$. Then

$$z \cdot DF(m)\gamma\left([\,\xi_1(m),\zeta\,]\right)(m) = J(\mathrm{d}F_z(m)) \cdot [\,\xi_1(m),\zeta\,] = \left(\mathrm{ad}^*_{\xi_1(m)}J(\mathrm{d}F_z(m))\right) \cdot \zeta$$

for all $m \in M$, $\zeta \in \mathfrak{g}$, and $z \in W^*$. Thus $\chi_m^* z = -\tfrac{1}{2}\mathrm{ad}^*_{\xi_1(m)}J(\mathrm{d}F_z(m))$ and the "ontoness" condition in Corollary 11 is equivalent to the condition that

$$\left(\mathrm{ad}^*_{\xi_1(m)}J(\mathrm{d}F_z(m))\right)\bigg|_{\mathfrak{g}_m} \neq 0$$

for any nonzero $z \in W^*$.

The condition (7.6) forms a system of equations for the isotropy corrections $\sigma(m,t)$ that exactly preserve the first integral. Using (7.7), we can rewrite this system as

$$\hat{F}_m(\zeta, t) = 0 \qquad \text{for} \qquad \zeta = \sigma(m,t).$$

In some situations, it may not be cost effective to solve this nonlinear equation. In such cases, we can replace the equation with its linearization

$$\chi_m(\zeta) + \hat{F}_m(0,0) = 0$$

at $t = 0$. If there exists $V_m \subset \mathfrak{g}_m$ such that $\chi_m|_{V_m}$ is an isomorphism, then

$$\sigma(m) = -(\chi_m|_{V_m})^{-1} \hat{F}(m,0,0)$$

determines an update $m \mapsto \hat{\mathcal{F}}_t(m, \sigma(m))$ that preserves F to order n.

The 'improved' algorithms of the free rigid body equations given in §4 can be regarded as implementations of this approach, given that the energy level sets coincide with the orbits of the system. (The only exception arises in the case of an axisymmetric body, for which the 'equator' consists of a circle of equilibria; however, all of the algorithms presented in §4 exactly capture the equilibria.)

8   Discussion.

In this paper, we have shown how to improve geometric integration algorithms through the application of isotropy subgroups. While the addition of isotropy terms does not affect the original dynamics, it can have nontrivial effects on the numerical approximation algorithms. We have seen that, under certain nondegeneracy hypotheses, the order of an algorithm can be increased, and conservative properties maintained by an isotropy modification of the naïve version of the original algorithm. In this introductory treatment, our results have been illustrated by algorithms on the sphere, particularly algorithms for reduced rigid body dynamics of interest in geometric mechanics. More substantial applications in micromagnetics can be found in [24]. While the examples treated so far indicate the desirability of such isotropy-enhanced geometric algorithms, the final verdict on the significance of our results in practical applications must await more substantial testing on complicated "real-world" problems.

Several further points are suggested by our results.

- If isotropy is used to improve a stable algorithm, will the resulting algorithm also be stable? For instance, can an order $n-1$ implicit method be combined with an explicit improvement to yield a stable order n method?

- If the group does not act transitively, the isotropy subalgebras may be non-conjugate and even of varying dimension. In such situations, it may not be possible to satisfy any condition to higher order than that of the original algorithm on the entire manifold, but conditions such as those described above can be used to determine a unique choice of

generator at points with additional isotropy. For example, if there is a submanifold N of M such that $\dim \mathfrak{g}_m > 0$ for $m \in N$ and $\dim \mathfrak{g}_m = 0$ for $m \in M \setminus N$, then the constructions described above can be used to specify the map ψ on N. The smoothness of such a modified ψ is not clear.

- Can the conservative and symmetry-preserving properties of the algorithms be further improved by application of the moving frame based numerical approximations introduced in [34]?

References

[1] Ablowitz, M. J., Herbst, B. M., and Schober, C., On the numerics of integrable discretizations, in: *Symmetries and integrability of difference equations (Estrel, PQ, 1994)*, CRM Proc. Lecture Notes 9, 1996, pp. 1–11.

[2] Ablowitz, M. J., Herbst, B. M., and Schober, C., On the numerical solution of the sine-Gordon equation. I. Integrable discretizations and homoclinic manifolds, *J. Comput. Phys.* **126** (1996), 299–314.

[3] Ablowitz, M. J., Herbst, B. M., and Schober, C., On the numerical solution of the sine-Gordon equation. II. Performance of numerical schemes, *J. Comput. Phys.* **131** (1997), 354–367.

[4] Bailey, C.D., Application of Hamilton's law of varying action, *AIAA J.* **13** (1975), 1154–1157.

[5] Baruch, M., and Riff, R., Hamilton's Principle, Hamilton's Law — 6^n correct formulations, *AIAA J.* **20** (1982), 687–692.

[6] Bridges, T.J., and Reich, S., Multi-symplectic spectral discretizations for the Zakharov-Kuznetsov and shallow water equations, *Phys. D* **152/153** (2001), 491–504.

[7] Budd, C.J., and Collins, C.B., Symmetry based numerical methods for partial differential equations, in: *Numerical analysis 1997*, D.F. Griffiths, D.J. Higham and G.A. Watson, eds., Pitman Res. Notes Math., vol. 380, Longman, Harlow, 1998, pp. 16–36.

[8] Budd, C.J., and Iserles, A., Geometric integration: numerical solution of differential equations on manifolds, *Phil. Trans. Roy. Soc. London A* **357** (1999), 945–956.

[9] Channell, P.J., and Scovel, C., Symplectic integration of Hamiltonian systems, *Nonlinearity* **3** (1990), 231–259.

[10] de Vogelaere, R., Methods of integration which preserve the contact transformation property of Hamiltonian equations, preprint, Dept. of Math., U. of Notre Dame, Report 4, 1956.

[11] Dorodnitsyn, V.A., Finite difference models entirely inheriting continuous symmetry of original differential equations, *Int. J. Mod. Phys. C* **5** (1994), 723–734.

[12] Feng, Kang, Difference schemes for Hamiltonian formalism and symplectic geometry, *J. Comp. Math.* **4** (1986), 279–289.

[13] Fels, M., and Olver, P.J., Moving coframes. II. Regularization and theoretical foundations, *Acta Appl. Math.* **55** (1999), 127–208.

[14] Golubitsky, M., and Guillemin, V., *Stable Mappings and their Singularities*, Springer–Verlag, New York, 1973.

[15] Iserles, A., Munthe–Kaas, H.Z., Nørsett, S.P., and Zanna, A., Lie group methods, *Acta Numerica* (2000), 215–365.

[16] Iserles, A., On Cayley-transform methods for the discretization of Lie-group equations, *Found. Comput. Math* **1** (2001), 129–160.

[17] Jaegers, P.J., *Lie group invariant finite difference schemes for the neutron diffusion equation*, Ph.D. Thesis, Los Alamos National Lab Report, LA–12791–T, 1994.

[18] Krupa, M., Bifurcations of relative equilibria, *SIAM J. Math. Anal.* **21** (1990), 1453–1486.

[19] Labudde, R.A., and Greenspan, D., Energy and momentum conserving methods of arbitrary order for the numerical integration of equations of motion. Part I, *Num. Math.* **25** (1976), 323–346.

[20] Labudde, R.A., and Greenspan, D., Energy and momentum conserving methods of arbitrary order for the numerical integration of equations of motion. Part II, *Num. Math.* **26** (1976), 1–16.

[21] Lasagni, F.M., Canonical Runge-Kutta methods, *ZAMP* **39** (1988), 952–953.

[22] Lewis, D., Lagrangian block diagonalization, *J. Dyn. Diff Eqn.* **4** (1992), 1–42.

[23] Lewis, D., Bifurcation of liquid drops, *Nonlinearity* **6** (1993), 491–522.

[24] Lewis, D., and Nigam, N., Geometric integration on spheres and some interesting applications, preprint, 2001.

[25] Lewis, D., and Simo, J.C., Conserving algorithms for the dynamics of Hamiltonian systems on Lie groups, *J. Nonlin. Sci.* **4** (1994), 253–299.

[26] Lewis, D., and Simo, J.C., Conserving algorithms for the N dimensional rigid body, *in: Proceedings of the Fields Institute workshop on Integration Algorithms for Classical Mechanics, Fields Institute Communications series* **10** (1996), 121–139.

[27] Marciniak, A., Energy conserving, arbitrary order numerical solutions of the N-body problem, *Num. Math.* **45** (1984), 207–218.

[28] McLachlan, R.I., and Scovel, J.C., Equivariant constrained symplectic integration, *J. Nonlin. Sci.* **5** (1995), 233–256.

[29] Marsden, J.E., Patrick, G.W., and Shkoller, S., Multisymplectic geometry, variational integrators, and nonlinear PDEs, *Comm. Math. Phys.* **199, no. 2** (1998), 351–395.

[30] McLachlan, R.I., and Quispel, G.R.W., Six lectures on the geometric integration of ODEs, *in: Foundations of Computational Mathematics*, R. DeVore, A. Iserles and E. Suli, eds., London Math. Soc. Lecture Note Series, vol. 284, Cambridge University Press, Cambridge, 2001, pp. 155–210.

[31] Moser, J., and Veselov, A.P., Discrete versions of some classical integrable systems and factorization of matrix polynomials, *Commun. Math. Phys.* **139** (1991), 217–243.

[32] Munthe–Kaas, H., Runge–Kutta methods on Lie groups, *BIT* **38** (1998), 92–111.

[33] Olver, P.J., Joint invariant signatures, *Found. Comput. Math.* **1** (2001), 3–67.

[34] Olver, P.J., Geometric foundations of numerical algorithms and symmetry, *Appl. Alg. Engin. Commun. Comput.* **11** (2001), 417–436.

[35] Sanz-Serna, J.M., Runge-Kutta schemes for Hamiltonian systems, *BIT* **28** (1988), 877–883.

[36] Shokin, Y.I., *The Method of Differential Approximation*, Springer–Verlag, New York, 1983.

[37] Simo, J.C., and Fox, D.D., On a stress resultant geometrically exact shell model. I. Formulation and optimal parametrization, *Comp. Meth. Appl. Mech Eng.* **72** (1989), 267–304.

[38] Simo, J.C., Rifai, M.S., and Fox, D.D., On a stress resultant geometrically exact shell model. Part VI: Conserving algorithms for nonlinear dynamics, *Int. J. Num. Meth. Eng.* **34** (1991), 117–164.

[39] Simo, J.C., and Tarnow, N., The discrete energy-momentum method. Conserving algorithms for nonlinear elastodynamics, *Z. Ang. Math. Phys.* **43** (1992), 757–792.

[40] Simo, J.C., Tarnow, N., and Wong, K., Exact energy-momentum conserving algorithms and symplectic schemes for nonlinear dynamics, *Comp. Meth. Appl. Mech Eng.* **100** (1992), 63–116.

[41] Simo, J.C., and Vu-Quoc, L., On the dynamics in space of rods undergoing large motions—a geometrically exact approach, *Comp. Meth. Appl. Mech Eng.* **66** (1988), 125–161.

[42] Simo, J.C., and Wong, K., Unconditionally stable algorithms for rigid body dynamics that exactly preserve energy and momentum, *Int. J. Num. Meth. Eng.* **31** (1991), 19–52.

[43] van Beckum, F.P.H., and van Groesen, E., Discretizations conserving energy and other constants of the motion, in: *Proc. ICIAM 87*, Paris, 1987, pp. 17–35.

[44] Weyl, H., *Classical Groups*, Princeton Univ. Press, Princeton, N.J., 1946.

[45] Wisdom, J., and Holman, M., Symplectic maps for the N-body problem, *Astro. J.* **102** (1991), 1528–1538.

[46] Zhong, G., and Marsden, J.M., Lie-Poisson Hamilton-Jacobi Theory and Lie-Poisson integrators, *Phys. Lett. A* **33**:3 (1988), 134–139.