

## Variable stepsize commutator free Lie group integrators

Charles Curry · Brynjulf Owren

the date of receipt and acceptance should be inserted later

**Abstract** We introduce variable stepsize commutator free Lie group integrators, where the error control is achieved using embedded Runge-Kutta pairs. These are schemes for the integration of initial value problems posed on homogeneous spaces by means of Lie group actions. The focus is on commutator free methods, in which the approximation evolves by composing flows generated by Lie group exponentials. Such methods are encoded by a generalization of Butcher's Runge-Kutta tableaux, but it is known that more order conditions must be satisfied to obtain a scheme of a given order than are required for classical RK schemes. These extra considerations complicate the task of designing embedded pairs. Moreover, whilst the computational cost of RK schemes is typically dominated by function evaluations, in most situations the dominant cost of commutator free Lie group integrators comes from computing Lie group exponentials. We therefore give Butcher tableaux for several families of methods of order 3(2) and 4(3), designed with a view to minimizing the number of Lie group exponentials required at each time step, and briefly discuss practical error control mechanisms. The methods are then applied to a selection of examples illustrating the expected behaviour.

**Keywords** Lie group integrators, adaptive error control, geometric integration, commutator-free methods

**Acknowledgements** This work has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 691070, and from The Research Council of Norway (project 231632).

---

C. Curry · B. Owren  
Department of Mathematical Sciences, Norwegian University of Science and Technology (NTNU), 7491 Trondheim, Norway.  
Tel: +47 48201626  
E-mail: charles.curry@ntnu.no

## 1 Introduction

Off the shelf computer software for the numerical solution of ordinary differential equations usually comes with built-in error control and variable step size. Typically, such codes compute an estimate of the local truncation error in each step. This estimate is then compared to a user specified tolerance, `tol`. If its norm is smaller than `tol`, the step will be accepted and a step size to be used in the succeeding step is computed. If the estimate is larger than `tol`, the step will be rejected and attempted again with a new reduced step size. Several adjustments to this procedure can be made in order to improve the behaviour of the integrator. For instance one can seek to avoid an excessive number of rejected steps, or an oscillating behaviour of the sequence of step sizes.

For Runge-Kutta schemes, the most popular method for obtaining an error estimate is by using the method of embedded pairs, see for instance the monograph [9, p. 164–172] for details. For a given initial point  $y_n$  and a proposed step size  $h_n$ ,  $y_{n+1}$  is computed with a method of convergence order  $p$ , together with an auxiliary second approximation  $\hat{y}_{n+1}$  of convergence order  $\hat{p} \neq p$ . The resulting pair is usually labelled  $p(\hat{p})$  in the literature. Historically, it was customary to use  $\hat{p} > p$  and take  $y_{n+1} - \hat{y}_{n+1}$  as an estimate of the error in the lower order approximation  $y_{n+1}$ , a popular scheme of this form is the Runge-Kutta-Fehlberg 4(5) method. Yet, it seems unfortunate to not make use of the more accurate approximation, and it looks like the preferable choice these days is to use  $\hat{p} < p$  but adjust the step size changing formula in such a way that the global error behaves proportionally to the user given tolerance `tol`. This is done in several codes, such as the Dormand-Prince 5(4) method [6] implemented in the MATLAB solver `ode45` [17]. An important feature of embedded Runge-Kutta pairs is that they share the same internal stages. In this way, the cost of computing with the pair of method is not much higher than computing only the principal approximation  $y_{n+1}$  in each step. In the last few decades a subfield of numerical analysis called geometric integration has been established, see [8] for an exhaustive account. The main purpose is to develop numerical methods that preserve certain underlying geometric structures of the differential equation. Examples of such structures are symplecticity, volume, reversibility, and first integrals. In this paper we consider schemes which were designed for the situation where the differential equation has a natural formulation by means of a Lie group action on a smooth manifold, these schemes are called Lie group integrators, [3], and the particular type considered here are called commutator-free Lie group integrators [2]. Many of the aforementioned methods have primarily been implemented with constant step size.

Turning now to Lie group integrators, it is for our purpose useful to divide them into two categories. One consists of schemes which can be interpreted as applying a standard Runge-Kutta scheme to a local coordinate representation of the original differential equation, this category includes the Runge-Kutta-Munthe-Kaas methods. For such methods, error control and variable step

size can be implemented simply by applying the technique described above to the coordinate representation. The second category, however, does not have such a natural coordinate representation, the schemes are typically constructed by composing flows of simple vector fields. In this case, embedded pairs can only be derived by solving order conditions for Lie group integrators. For a general and modern exposition of order conditions of such integrators, see for example [12]. A detailed account of the order conditions for commutator-free Lie group integrators was given in [15]. The problem of deriving embedded pairs of efficient commutator-free schemes requires careful consideration. Not only is it important that the stages of the pair can be shared, in order for such schemes to be competitive to other Lie group integrators, they also need to reuse flow calculations to as large an extent as possible.

The aim of this paper is to show how one can derive embedded pairs of commutator-free Runge–Kutta methods with the smallest additional computational cost measured in terms of stages and flow calculations (exponentials) per step. Of course, it by no means clear that this way of measuring efficiency yields the smallest possible global error for a given computational cost. But we believe that our approach is a good starting point that illustrates in a clear way some of the challenges involved in constructing embedded pairs of commutator-free methods. The rest of this paper is organised as follows: in section 2 we begin by briefly describing the class of schemes we consider, and in particular we review the order conditions which were derived in [15]. We then give particular examples of embedded pairs of type 3(2) and 4(3). These schemes are constructed such that the pairs share common stages, and employ the well-known “First same as last” (FSAL) property known from classical Runge–Kutta methods. We also design the pair to reuse flow calculations to a largest possible extent. In the third section, we apply the CF43 method to three examples; first the Euler free rigid body, a standard test case for Lie group integrators. As a second example, we apply the method to the stiff Van der Pol equation, where we formulate the problem via the standard matrix-vector action by  $GL(2)$  on  $\mathbb{R}^2$ . Finally, we consider the heavy top, formulated via the coadjoint action of the Lie group  $SE(3)$  on the dual  $\mathfrak{se}(3)^*$  of its Lie algebra.

## 2 Commutator free methods

We consider the numerical solution of autonomous initial value problems on manifolds. For a given manifold  $M$  and vector field  $f$ , that is a section of the tangent bundle  $TM$ , a smooth curve  $y(t)$  is an integral curve of  $f$  if it satisfies

$$\dot{y} = f(y(t)), \quad t \in (a, b)$$

For any given starting location  $y(a) = p$ , the integral curve through  $p$  exists for sufficiently small time intervals. In the special case that  $M$  is a finite-dimensional vector space (which is identified with  $\mathbb{R}^n$ ), integral curves may be

approximated numerically using Runge-Kutta methods, i.e.

$$k_i = f \left( y_n + h \sum_{j=1}^s \alpha_i^j k_j \right), \quad i = 1, \dots, s$$

$$y_{n+1} = y_n + h \sum_{i=1}^s b^i k_i$$

It is not so apparent how to extend such methods to manifolds not possessing a linear structure. We begin by observing that the Euler scheme may be interpreted as a method which first ‘freezes’ the vector field  $f$  at the point  $p = y_n$  so that  $f_p(y) = f_p(y_n)$  for all  $y$ , and then flows a distance  $h$  along the integral curve of the frozen vector field  $f_p$ , i.e.

$$y_{n+1} = \exp(hf_p)y_n,$$

where  $\exp(tf_p) : y \mapsto y + tf_p(y)$  is the one-parameter group of diffeomorphisms of  $\mathbb{R}^n$  induced by the vector field  $f_p$  (i.e. the motion is along the integral curves of  $f_p$ , which are straight lines as  $f_p(y)$  is independent of  $y$ ). However, our notion of freezing the vector field at  $p$  relies on the identification of the tangent spaces  $T_p(M)$  and  $T_y(M)$  for any  $y$ . In the general setting, this requires a notion of parallel transport, which in turn requires a choice of linear connection. In general, if we freeze a vector field  $f$  by parallel transporting  $f_p(p) \in T_p(M)$  onto  $T_y(M)$ , the result is not independent of the path connecting  $p$  and  $y$  along which the transport is performed. For Lie group integrators, the connection is typically chosen to be flat so we can ignore this issue, but more generally assuming  $y$  is sufficiently close to  $p$ , the path could be taken to be the geodesic connecting  $y$  and  $p$ . In any case, the consequence is that  $\exp(tf_p).p$  will coincide with the geodesic  $\gamma_t$  in the direction of  $f(p)$  from  $p$ .

Two strategies emerge from the above comments when generalizing Runge-Kutta methods to manifolds. The first [2, 4, 16] consists of considering the  $k_i$  as vector fields frozen at different points, parallel transporting these to a base tangent space  $T_p(M)$ , taking a linear combination of the resulting tangent vectors and flowing along the geodesic in this direction. This strategy, when implemented using a Lie group action on a homogeneous space, results in the RKMK family of integrators, and allows the use of standard Runge-Kutta tableaux.

In some situations, it is advantageous to avoid the computation of commutators of vector fields typical of RKMK methods by taking compositions of exponentials of frozen vector fields. This leads to the commutator-free Runge-Kutta methods:

$$g_r = \exp\left(h \sum_k \alpha_{r,J}^k f_{g_k}\right) \cdots \exp\left(h \sum_k \alpha_{r,1}^k f_{g_k}\right)p, \quad r = 1, \dots, s$$

$$y_{n+1} = \exp\left(h \sum_k \beta_J^k f_{g_k}\right) \cdots \exp\left(h \sum_k \beta_1^k f_{g_k}\right)p$$

## 2.1 Lie group integrators

The above methods are in principle very general, but require the computation of geodesics and possibly parallel transport, which is typically impractical. The observation underlying Lie group integrators is that the geodesics  $\exp(hf_p)$  on a Lie group or reductive homogeneous space equipped with the canonical connection can be computed using the Lie group exponential. In practice, this typically means a matrix exponential, which can be approximated to machine accuracy with tolerable computational effort. Indeed, we suppose that the ODE we wish to solve may be written in the form

$$y' = (\lambda_* f(y))(y), \quad y(0) = p, \quad (2.1)$$

where  $f : M \rightarrow \mathfrak{g}$ , and  $\lambda_* : \mathfrak{g} \rightarrow \mathcal{X}(M)$  is the infinitesimal action arising from a group action  $\Lambda : G \times M \rightarrow M$  as

$$\lambda_*(u)(p) = \left. \frac{d}{dt} \right|_{t=0} \Lambda(\exp(tu), p)$$

The vector field  $\lambda_* f$  is then frozen straightforwardly at a point  $p$  by taking

$$\lambda_* f_p(y) = \left. \frac{d}{dt} \right|_{t=0} \Lambda(\exp(tf(p)), y)$$

The exponentials of these frozen vector fields are seen to obey

$$\exp(t\lambda_* f_p)y = \Lambda(\exp(tf(p)), y)$$

As a consequence, commutator-free RK Lie-group integrators take the form

$$\begin{aligned} g_r &= \exp\left(h \sum_k \alpha_{r,J}^k f_k\right) \cdots \exp\left(h \sum_k \alpha_{r,1}^k f_k\right) \\ f_r &= f(\Lambda(g_r, p)) \\ y_{n+1} &= \Lambda\left(\exp\left(h \sum_k \beta_J^k f_k\right) \cdots \exp\left(h \sum_k \beta_1^k f_k\right), p\right) \end{aligned}$$

As an example, consider the linear differential equation

$$y' = A(y)y,$$

where  $y \in GL(n)$  and  $A \in \mathfrak{gl}(n)$ . The anchor map  $\lambda_*$  associates  $A$  with a right-invariant vector field in the canonical manner; both the group action  $\Lambda$  and its associated infinitesimal action are given by matrix left-multiplication. The exponentials of frozen vector fields are simply matrix exponentials, i.e.

$$\exp(t\lambda_* f_p)y = \exp(tA(p))y.$$

Related examples are given in §3.

## 2.2 Order conditions

Let  $\hat{\varphi}_h : M \rightarrow M$  be the mapping corresponding to taking a single timestep of a given commutator-free RK method, and  $\varphi_h$  the mapping which flows for time  $h$  along the solution curve through the initial point. The method is said to be of order  $p$  if, for all  $C^\infty$ -functions  $\psi$  on  $M$ , we have

$$\psi(\varphi_h(p)) - \psi(\hat{\varphi}_h(p)) = \mathcal{O}(h^{p+1})$$

For a justification that this condition leads to global error estimates, see [5].

A commutator-free RK scheme is in general specified by the collection of coefficients  $\alpha_{r,j}^k, \beta_j^k$ . An order theory comprising systems of algebraic equations in the coefficients to be satisfied to attain a given order was derived in [15], analogous to the theory of order conditions of standard RK methods. We will summarize the results without proof here. First, define the coefficients

$$a_r^k = \sum_j \alpha_{r,j}^k, \quad b^k = \sum_j \beta_j^k, \quad c_k = \sum_i a_k^i$$

**Lemma 1** *A necessary condition for a commutator-free RK scheme with coefficients  $\alpha_{r,j}^k, \beta_j^k$  to have order  $p$  is that the associated  $a_r^k, b^k$  are the coefficients of a standard Runge-Kutta method of order  $p$  or greater.*

To attain a method of order 1 or 2, the above conditions are also sufficient. In particular, we can take  $J = 1$ , such that there is only one exponential computed at each stage. On the other hand, we must satisfy an additional condition to attain order 3:

**Lemma 2** *There are no order 3 methods employing exclusively  $J = 1$ . Given that  $J = 2$ , a sufficient condition to attain order 3 is that the coefficients  $a$  and  $b$  form a classical RK3 method, and in addition*

$$\sum_k \beta_1^k c_k + \frac{1}{2} \beta_2^k = \frac{1}{3},$$

We see that we only require two exponentials for computation of  $y_{n+1}$ ; one exponential suffices for the computation of  $g_r$  and  $f_r$ , i.e. we can take  $\alpha_{r,1}^k = a_r^k$ . The situation changes for methods of order 4:

**Lemma 3** *Given  $J = 2$ , the order 4 conditions are the classical conditions together with the non-classical order 3 condition and the following:*

$$\begin{aligned} \sum_k \beta_1^k c_k + \frac{1}{3} \beta_2^k &= \frac{1}{4} \\ \sum_k \beta_1^k c_k^2 + \frac{1}{3} \beta_2^k &= \frac{1}{6} \\ \sum_{j,k} \beta_1^k a_k^j c_j + \frac{1}{6} \sum_k \beta_2^k &= \frac{1}{12} \\ \sum_{i,j} b^i c_i \alpha_{j,1}^i c_j + \sum_{i,j,k} b^i a_i^j c_i \alpha_{j,2}^i c_j &= \frac{1}{12} \end{aligned}$$

To satisfy these conditions, it suffices to take two exponentials in one of the four intermediate stages, two in the final stage, and one exponential in other stages. We illustrate this by giving a sample tableau, displaying the coefficients of a method given in [15] which extends the classical RK4 method:

$$\begin{array}{c|cccc}
 0 & & & & \\
 \frac{1}{2} & \frac{1}{2} & & & \\
 \frac{1}{2} & 0 & \frac{1}{2} & & \\
 1 & \frac{1}{2} & 0 & 0 & \\
 \hline
 & -\frac{1}{2} & 0 & 1 & \\
 \hline
 & \frac{1}{4} & \frac{1}{6} & \frac{1}{6} & -\frac{1}{12} \\
 & -\frac{1}{12} & \frac{1}{6} & \frac{1}{6} & \frac{1}{4}
 \end{array}$$

This is a four-stage method, using two exponentials for the fourth stage, where the tableau displayed above is such that  $\alpha_{4,1} = (\frac{1}{2}, 0, 0)$ , and  $\alpha_{4,2} = (-\frac{1}{2}, 0, 1)$ .

To obtain a method of order 5, at least three exponentials are required for the final stage. There results a large nonlinear system of algebraic equations which has so far proved resistant to all attempts to attain a solution. Indeed, the construction of commutator-free methods of order 5 or higher remains an open problem.

### 2.3 Embedded pairs reusing exponentials

Classical Runge-Kutta schemes are typically implemented as an embedded pair, i.e. two different sets of coefficients  $b^i, \hat{b}^i$  are given such that the associated approximations  $y_n, \hat{y}_n$  are of different orders, typically  $\hat{p} = p-1$ . In this section, we extend this idea to commutator free RK methods, devising schemes with coefficients  $\beta_j^i, \hat{\beta}_j^i$  to create embedded pairs of orders differing by 1.

In practice, it is important to note that the dominant computational cost in the implementation of a commutator free method is likely to be the evaluation of Lie group (matrix) exponentials. Each *unique* horizontal row in the tableau generally requires the evaluation of one exponential. The observation of Owren [15] was that it is possible to design tableaux where certain rows coincide, such as the commutator free RK4 method above where the second and fourth rows are identical. This consideration looms large in the implementation of embedded pairs, where there is even greater opportunity for reuse of exponentials. A general CF3 scheme requires at least 3 exponentials, whilst a CF4 scheme requires 5. Our main achievement is the construction of CF32 and CF43 schemes using only one additional exponential and function evaluation compared to the constant stepsize scheme of the same order.

### 2.3.1 CF32

The order conditions to obtain a commutator free method of second order coincide with the classical RK2 conditions. As in the classical case the construction of a RK32 pair is always possible from an RK3 scheme, there is no difficulty in constructing a commutator free RK32 pair. In general, this will require one extra exponential and function evaluation. Suppose the tableau is constructed as follows

$$\begin{array}{c|ccc}
 0 & & & \\
 c_2 & a_{21} & & \\
 c_3 & a_{31} & a_{32} & \\
 \hline
 y & \beta_1^1 & \beta_1^2 & \beta_1^3 \\
 & \beta_2^1 & \beta_2^2 & \beta_2^3 \\
 \hline
 \hat{y} & a_{31} & a_{32} & 0,
 \end{array}
 \quad \text{or} \quad
 \begin{array}{c|ccc}
 0 & & & \\
 c_2 & a_{21} & & \\
 c_3 & a_{31} & a_{32} & \\
 \hline
 1 & \beta_1^1 & \beta_1^2 & \beta_1^3 \\
 & \beta_2^1 & \beta_2^2 & \beta_2^3 \\
 \hline
 y & \beta_1^1 & \beta_1^2 & \beta_1^3 \\
 & \beta_2^1 & \beta_2^2 & \beta_2^3 \\
 \hline
 \hat{y} & a_{31} & a_{32} & 0,
 \end{array}$$

implying a reuse of the third stage and hence saving an exponential. The two schemes above are identical, but the right hand tableau is of a more general type that allows  $\hat{y}$  to depend on the value of  $f(y)$ . Schemes of this type have the FSAL (first same as last) property, as the first function evaluation at the next step coincides with the function evaluation  $f(y)$  assuming the step is accepted. We will henceforth write all tableaux in this form, and abbreviate the row for  $y$  as FSAL.

The order 2 conditions become  $a_{31} + a_{32} = 1$  and  $a_{21}a_{32} = \frac{1}{2}$ , hence any 3-stage commutator free RK3 scheme obeying

$$a_{32} = \frac{1}{2a_{21}}, \quad a_{31} = \frac{2a_{21} - 1}{2a_{21}}$$

would admit an embedded pair without the need for an extra exponential. Unfortunately, such a scheme does not exist, as is shown readily using symbolic computation software. We therefore focus on constructing commutator free RK3 schemes reusing an exponential, a topic as yet unexplored in the literature. All of the results stated in the remainder of this section were proven using symbolic computation software.

Let  $\omega$  be a root of  $36z^2 + (9a - 30)z + 3a + 1$ . The general tableau of an RK32 scheme reusing the third stage in the second row of the fourth stage is given in Table 1.

Once  $a$  is chosen and a consistent choice of root for  $\omega(a)$  is made, there then follows a system of two linear equations for the four coefficients of  $\hat{y}$ , in general resulting in a two-parameter family once the rest of the tableau is set. In general, the values will be irrational, but we highlight that some choices of  $a$  give rational coefficients, for instance  $a = -\frac{1}{3}, \frac{1}{3}, \frac{2}{3}, \frac{7}{9}, \frac{101}{9}, \frac{34}{3}, \frac{35}{3}$  etc. In



0			
$a$	$a$		
$\frac{6\omega-1}{3}$	$\frac{6a\omega-3\omega-a}{3a}$	$\frac{\omega}{a}$	
1	$\frac{-12a\omega-18\omega-11a+11}{6}$	$\frac{18\omega(1-a)-11}{6a(3a+1)}$	$\frac{-12\omega+3a-10}{2(3a+1)}$
	$\frac{6a\omega-3\omega-a}{3a}$	$\frac{\omega}{a}$	
$y$	FSAL		
$\hat{y}$			

0		0	
$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
1	-1    2	$-\frac{1}{6}$	$-\frac{5}{12} \frac{1}{4}$
1	1 $-\frac{5}{4} \frac{1}{4}$	1	$-\frac{37}{12} \frac{9}{4} 2$
	-1    2    0		$-\frac{5}{12} \frac{1}{4}$
$y$	FSAL		
$\hat{y}$	0 $\frac{3}{4}$ 0 $\frac{1}{4}$	$\hat{y}$	0 $\frac{3}{4}$ 0 $\frac{1}{4}$

Table 1: In the first row  $a$  we give a general FSAL commutator-free RK-scheme of order 3(2) reusing the exponential of the third stage in the second row of the fourth stage.  $\omega$  is a root of the polynomial  $36z^2 + (9a - 30)z + 3a + 1$ . In the second row two concrete examples with rational coefficients are given.

0			
$a$	$a$		
$\frac{6\nu-1}{3}$	$\frac{6a\nu-3\nu+a}{3a}$	$\frac{\nu}{a}$	
1	$\frac{6a\nu-3\nu+a}{3a}$	$\frac{\nu}{a}$	
	$\frac{-12a\nu-6\nu+a+1}{6a}$	$\frac{-18a\nu-6\nu+1}{6a(3a-1)}$	$\frac{12\nu+3a-2}{2(3a-1)}$
$y$	FSAL		
$\hat{y}$			

Table 2: Commutator-free scheme of order 3(2) reusing the exponential of the third stage in the first row of the fourth stage.

general these can be found by setting the discriminant  $81a^2 - 972a + 756$  to be a square. We give two samples of rational tableaux from the case  $a = \frac{1}{3}$  in the second row of Table 1. We can also reuse the third stage in the first row of the fourth stage; indeed this may be preferable for some integrators as it permits the storage of the action of the exponential on  $y_0$  rather than the exponential itself. Let  $\nu$  be the root of  $36z^2 + (9a - 6)z - 3a + 1$ . The general tableau for such a scheme is found in Table 2. There are also one-parameter families of RK32 schemes reusing the second stage, either in the first or second row of  $y$ . Fix  $a$  and let  $\gamma$  be a root of  $4a(3a - 1)z^2 + 4(3a - 1)z + 3$  and  $\delta$  be a root of  $4az^2 + (12a - 2)z + 9a + 6$ . The respective tableaux are given in Table 3.

0				0			
$\frac{1}{3}$	$\frac{1}{3}$			$-\frac{1}{3}$	$-\frac{1}{3}$		
$c_3$	$a$	$(\frac{2}{3} - 2a)(2a\gamma + 1)$		$c_3$	$-\frac{2a\delta}{3} - 2a$	$a$	
1	$\frac{1}{3}$	0	0	1	$\frac{-6a\delta + 8a + 3}{6a}$	$\delta$	$-\frac{1}{2a}$
$y$	$(3a - 1)\gamma + \frac{2}{3}$	$-3a\gamma$	$\gamma$	$y$	$-\frac{1}{3}$	0	0
$\hat{y}$	FSAL			$\hat{y}$	FSAL		

Table 3: One parameter families of commutator-free schemes reusing the exponential of the second stage in the first row (left) or the second row (right)

The discriminants of the equations for  $\gamma$  and  $\delta$  are  $16(1 - 3a)$  and  $4(1 - 36a)$  respectively, which allows for easy generation of rational versions of the above schemes should this be desired, for instance  $a = 0$  works in both cases.

### 2.3.2 CF43

In the classical case, there are no RK43 pairs with only 4 stages, so it is essential to use a 5-stage FSAL scheme to attain RK43 at minimal cost. This is no longer true for commutator free methods, as it is possible to take a single CF4 method together with  $\beta, \hat{\beta}$  such that the underlying  $b = \hat{b}$ , but that only  $\beta$  fulfills all of the non-classical order four conditions. Nonetheless, we will preserve the greater generality afforded by the FSAL schemes in our presentation.

An order three method requires two exponentials in the final computation, so in the generic case a CF43 pair requires two extra exponentials compared to the CF4 case. It is possible to reduce the extra cost to one by reusing an exponential, but the full tableau must be generated with this in mind, as a generic CF4 scheme does not permit a CF43 pair which reuses an exponential. In general, it is possible to reuse two exponentials (but not more) in a CF43 tableau, one above the lines (in the  $\alpha$ s) and one below the lines (i.e., generating  $\hat{y}$ ). Once the reuse pattern has been specified, the top of the tableau is fixed, and a one parameter family of  $\hat{\beta}$  coefficients is typically admitted. Note that patterns involving reuse of the second stage do not give CF43 pairs. In contrast to the CF32 case, there are no rational CF43 pairs reusing the optimal number of exponentials. We give the exact form of one of the tableaux in Table 4, but typically print only floating point forms; more accurate descriptions of the coefficients are available from the authors on request.

0				
$c_2$	$p_1(\omega)$			
$c_3$	$p_2(\omega)$	$p_3(\omega)$		
1	$p_2(\omega)$	$p_3(\omega)$		
	$p_4(\omega)$	$p_5(\omega)$	$p_6(\omega)$	
1	$p_7(\omega)$	$p_8(\omega)$	$p_9(\omega)$	$\frac{\omega}{2}$
	$-\frac{1}{3}p_7(\omega)$	$p_{10}(\omega)$	$p_{11}(\omega)$	$-\frac{3\omega}{2}$
$y$	FSAL			
$\hat{y}$	$p_{12}(\omega)$	$p_{13}(\omega)$	$p_{14}(\omega)$	0 0
	one parameter family			

$$p_1(\omega) = \frac{1}{2}(7 - 288\omega^4 - 36\omega^3 + 48\omega^2 + 17\omega)$$

$$p_2(\omega) = \frac{1}{268}(-389 + 31824\omega^4 + 10962\omega^3 - 3651\omega^2 - 2027\omega)$$

$$p_3(\omega) = \frac{1}{268}(54 - 2880\omega^4 - 2520\omega^3 + 234\omega^2 + 553\omega)$$

$$p_4(\omega) = \frac{-51696\omega^4 - 13878\omega^3 + 7557\omega^2 + 2285\omega + 1244}{804}$$

$$p_5(\omega) = \frac{-521424\omega^4 - 323586\omega^3 + 61119\omega^2 + 61599\omega + 10976}{20100}$$

$$p_6(\omega) = \frac{-5328\omega^4 + 558\omega^3 + 93\omega^2 - 122\omega + 47}{300}$$

$$p_7(\omega) = \frac{1008\omega^4 - 1530\omega^3 + 501\omega^2 - 16\omega + 229}{536}$$

$$p_8(\omega) = \frac{541872\omega^4 + 76158\omega^3 - 84207\omega^2 - 19972\omega - 2703}{40200}$$

$$p_9(\omega) = \frac{-2304\omega^4 + 144\omega^3 + 174\omega^2 + 4\omega + 21}{150}$$

$$p_{10}(\omega) = \frac{256752\omega^4 + 67878\omega^3 - 170787\omega^2 - 10852\omega + 22877}{40200}$$

$$p_{11}(\omega) = \frac{-864\omega^4 - 396\omega^3 + 684\omega^2 + 264\omega + 11}{150}$$

$$p_{12}(\omega) = \frac{-51696\omega^4 - 13878\omega^3 + 7557\omega^2 + 2285\omega + 1244}{804}$$

$$p_{13}(\omega) = \frac{-521424\omega^4 - 323586\omega^3 + 61119\omega^2 + 61599\omega + 10976}{20100}$$

$$p_{14}(\omega) = \frac{-5328\omega^4 + 558\omega^3 + 93\omega^2 - 122\omega + 47}{300}$$

Table 4: A commutator-free pair of order 4(3). Here  $\omega$  is the unique real root of  $144z^5 + 90z^4 - 3z^3 - 13z^2 - 5z - 1$ . This is the general form of a 4(3)-pair reusing the second stage in the first part of the third stage, and reusing the second part of the third stage in the first part of the  $\hat{y}$  computation.

The same tableau is given below as a decimal approximation:

0					
4.785707347	4.785707347				
.8093268944	.7701000600	.03922683443			
	.7701000600	.03922683443	0		
1	.6195164818	.06934556872	-.4981889449		
	.4211354919	-.005776103764	-.1381183969	.2227590088	
1	-.1403784973	.006491728470	1.302163795	-.6682770264	
$y$	FSAL				
$\hat{y}$	.6195164818	0.06934556872	-.4981889449		
	-.158427746 $a$	0.008517658 $a$	$a$	-.618653791 $a$	-.231436127 $a$
	-0.075415454	-0.082788288		+.5828295568	+.3847010797

Another possible choice of reuse pattern leads to the next tableau, where the one parameter family for  $\beta_2$  has been set such that  $\beta_2^3 = 0$ .

0					
	.67104050				
	2.547687640	-1.355037274			
	2.547687640	-1.355037274	0		
	-.21944181	-0.0735967	.1003880		
1	.324015249	.15832891	-.21057643	.2282322824	
	-.108005081	.84426683	.44843513	-.6846968472	
$y$	FSAL				
	-.21944181	-0.0735967	.1003880	0	0
	.45603817	.93310478	0	-.2660264	0.06953371

Perhaps the optimal CF43 method in terms of reuse pattern is the following, in which the reuse always occurs in the first part of a split stage, and the method is a true 4-stage method not using FSAL. It is the only CF43 scheme with these properties:

0					
	1.351207192				
	0.5	0.097900176			
	0.5	0.097900176			
	7.900943678	2.989500877	-10.48834473		
$y$	.301574869	-0.054881885	.238291289	0.01501572796	
	-.1005249562	.1005249562	.5450471839	-0.04504718389	
$\hat{y}$	0.5	0.097900176	0	0	
	-.2989500877	-0.0522571042	.783338473	-0.03003145592	

### 3 Practical implementation

We give an outline of how the above methods are implemented in practice. First we show how an embedded pair allows for automatic step size control, following [9]. Indeed, suppose we have chosen an initial step size  $h$ , and obtain

approximate solutions  $y_1$  and  $\hat{y}_1$ . If  $y$  takes values in a normed space, we can consider  $y - \hat{y}$  as an estimate of the error, and aim to ensure

$$\|y_1 - \hat{y}_1\| < sc, \quad sc = Atol + \max(\|y_0\|, \|y_1\|) \cdot Rtol,$$

for some user-specified absolute and relative tolerances  $Atol, Rtol$ .

On the other hand, in many case  $y$  takes values in a manifold  $M$  for which there is no meaningful linear structure. For compact homogeneous manifolds such as  $S_n$  we can measure distances in the geodesic metric (e.g. in the Lie group setting the geodesics are those associated to the bi-invariant metric arising from the Killing form); on the other hand relative errors are only meaningful where scaling makes sense, so will generally be inappropriate. For non-compact Lie groups there is typically no one canonical metric, but it may be shown that the existence of plausible metrics is guaranteed, see [7, 11]; convergence results can typically be established generically [5].

It is also possible to measure distances by linearization in the Lie algebra, provided sufficient care is taken over possible ambiguities arising from the multivalued Lie group logarithm and non-trivial isotropies of the Lie group action. In the matrix Lie group setting, an error analysis was conducted in [18] by linearizing about the exact solution, in the case of  $SO(n)$  for instance this is equivalent to the geodesic distance. This suggests that we may have an appropriate notion of relative error where modelling perturbations from an initial point, as distances may be measured by linearization in the tangent space at the origin. For instance in the case of  $SO(2)$  it might be appropriate to allow larger errors as the rotation angle  $\theta$  strays further from zero. On the other hand, when the global topology becomes relevant the induced distance is inappropriate, as can be seen by considering  $S_2$  - if linearizing around the North pole, nearly adjacent points on the South pole will be considered further from each other than the North pole!

Suppose  $d$  is the chosen metric, and let  $err = d(y_1, \hat{y}_1)/sc$ . In general, for an embedded pair of order  $p(p-1)$ , the optimal step size is  $h_{opt} = h \cdot err^{-\frac{1}{p}}$ . In practice, it is usual to dampen the fluctuations in  $h$ , a typically procedure is to let

$$h_{new} = h \cdot \min(facmax, \max(facmin, fac \cdot err^{-\frac{1}{p}})),$$

for some  $fac < 1$  and appropriate  $facmin, facmax$ . The above discussion is in some sense only rigorous when the lower order approximation  $\hat{y}$  is used to continue the integration, as only then is  $d(y, \hat{y})$  a good measure of the local error. Nonetheless, practice has shown that it is usually better to continue with the higher order integrator, a procedure known as local extrapolation.

We have not addressed automatic selection of the initial step size. For problems in vector spaces, this is typically performed by constructing estimates of the derivatives of vector field  $F$ , using some combination of function evaluations and taking small step(s) with the Euler scheme, see [9]. Similar ideas may be employed for Lie group integrators, but require modification as derivatives cannot be approximated so simply using evaluations of vector fields at different points due to the difficulty of identifying nearby tangent spaces.

### 3.1 Free rigid body

We now give a simple illustrative example of Lie group integration. Euler's equations for the body angular momentum of a free rigid body in a reference frame parallel to the principal axes of inertia take the form

$$\frac{d\xi}{dt} = -mI^{-1}\hat{\xi}\cdot\xi, \quad (3.1)$$

where  $m$  is the body mass,  $I$  is the diagonal inertia tensor, and we use the standard convention of the hat-map

$$\xi = \begin{pmatrix} \xi_1 \\ \xi_2 \\ \xi_3 \end{pmatrix} \Rightarrow \hat{\xi} = \begin{pmatrix} 0 & -\xi_3 & \xi_2 \\ \xi_3 & 0 & -\xi_1 \\ -\xi_2 & \xi_1 & 0 \end{pmatrix}$$

The quantity  $|\xi|^2$  is readily shown to be conserved, i.e. equation (3.1) describes the evolution of  $\xi \in S^2$ . The equations allow for the immediate application of Lie group integrators, indeed the right hand side is of the form  $\lambda_* f(\xi)\xi$ , where the manifold  $M = S^2$ ,  $f : S^2 \rightarrow \mathfrak{so}(3)$  is the hat map, and the infinitesimal action is that of  $\mathfrak{so}(3)$  on  $S^2$  by matrix multiplication. The associated group action  $\Lambda$  is  $SO(3)$  acting by matrix multiplication. In this context, the commutator-free RK Lie-group integrator given by the coefficients  $\alpha, \beta$  is therefore

$$\begin{aligned} g_r &= \exp\left(h \sum_k \alpha_{r,J}^k f_k\right) \cdots \exp\left(h \sum_k \alpha_{r,1}^k f_k\right) \\ f_r &= \widehat{g_r \cdot \xi_n} \\ \xi_{n+1} &= \exp\left(h \sum_k \beta_J^k f_k\right) \cdots \exp\left(h \sum_k \beta_1^k f_k\right) \cdot \xi_n, \end{aligned}$$

where  $g_r \in SO(3)$ ,  $f_r \in \mathfrak{so}(3)$ , and  $\xi_n \in S^2$ . For instance, the scheme given in the bottom left table of 1 is implemented in the following MATLAB code:

```
function [y1,y1h] = RK32(f,y0,h)
F1 = h*f(y0);
g2 = expm(F1/3)*y0;
F2 = h*f(g2);
inter = expm(-F1+2*F2);
g3 = inter*y0;
F3 = h*f(g3);
g4 = inter*expm(F1-(5/4)*F2+(F3/4))*y0;
F4 = h*f(g4);
y1 = g4;
y1h = expm((3/4)*F2+(F4/4))*y0;
```

Figure 1 shows the results of implementing the above method on the free rigid body, alongside the CF43 pair (Table 4).




Fig. 1: The CF32 pair (bottom left, Table 1) and the CF43 pair (Table 4) applied to the free rigid body equations with  $I = (1, 2, 5)^T$  and random initial data on  $S^2$  on the time interval  $[0, 2]$ . The graph on the left shows the linear relationship between the prescribed error tolerance and the observed global error, whilst the right graph compares the computational cost (measured in number of exponentials) against global error.

### 3.2 The Van der Pol oscillator

We consider the following non-conservative oscillator with non-linear damping, formulated as a scalar second order differential equation

$$\ddot{x} - \mu(1 - x^2)\dot{x} + x = 0 \quad (3.2)$$

where  $\mu$  is a parameter that affects the stiffness of the system. One can rephrase this problem in the form (2.1) using the simple matrix-times-vector action of the Lie group  $GL(2)$  on  $\mathbb{R}^2 \setminus \{0\}$ ,

$$\frac{d}{dt} \begin{pmatrix} x \\ \dot{x} \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ -1 & \mu(1 - x^2) \end{pmatrix} \begin{pmatrix} x \\ \dot{x} \end{pmatrix} \quad (3.3)$$

The corresponding Lie group integrator can be interpreted as an exponential integrator in the sense defined for instance in [1]. It is not clear to us whether explicit Lie group integrators are generally well suited for stiff problems, but we believe it is still of interest to observe how the new embedded pair of commutator-free schemes behaves through the “needle” of the Van der Pol oscillator. We have implemented the CF32 scheme given in the bottom left tableau of Table 1 and applied it in a variable stepsize fashion to the Van der Pol oscillator (3.3) setting  $\mu = 60$  and  $y_0 = (1, 1)^T$  in all the experiments. In Figure 2 we show the two components of the solution computed by CF32 in the top graph, and note in particular the sharp downward spike (“needle”) for the second component in the approximate interval  $t \in [1.4, 1.56]$ . The relative and absolute tolerances were both set to  $10^{-3}$  in this experiment. The stepsizes chosen by the CF32 schemes are shown in the bottom graph of

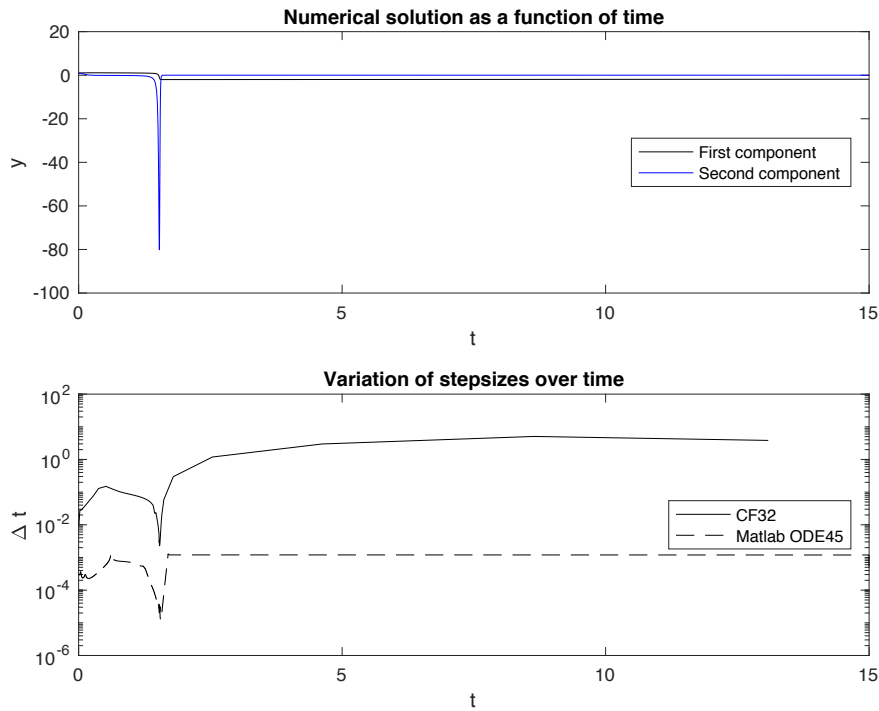


Fig. 2: The CF32 pair (bottom left, Table 1) applied to the Van der Pol equation with  $\mu = 60$  and  $y_0 = (1, 1)^T$  on the time interval  $[0, 15]$ . The figure shows the numerical solution (top) and the step size sequences selected by the method (bottom). The dashed line shows for comparison the step size sequence used by the Matlab solver ODE45. Both solvers used a tolerance  $TOL = 10^{-3}$ .

Figure 2, and one can see how the stepsizes are reduced through the needle. For comparison, we also show the stepsizes used by the builtin Matlab solver ODE45 which is based on the Dormand-Prince embedded Runge–Kutta pair [6]. We observe that the new CF32 solver behaves similarly to the Dormand-Prince scheme except that the former takes larger steps, this might be expected due to the fact that it computes (exact) matrix exponentials. For matrices with eigenvalues whose real parts tend to  $-\infty$ , such exact exponentials are bounded as opposed to their explicit Runge–Kutta counterparts which use polynomial approximations to the exponential map. In Figure 3 we visualise the difference between constant and variable stepsize for the problem (3.3). By running the variable stepsize code for a number of different choices for tolerances, we compute the global error at a fixed point  $t = 1.6$ , just after passing through the needle. The exact solution has been approximated in various different ways, one was to run various builtin Matlab solvers with strict tolerances. As a measure for the efficiency of the integrator, we have computed the number of exponential calculations used by the code to obtain a prescribed global error. The dashed line in Figure 3 shows the result for the constant stepsize version of the third order method in our embedded pair, and the solid



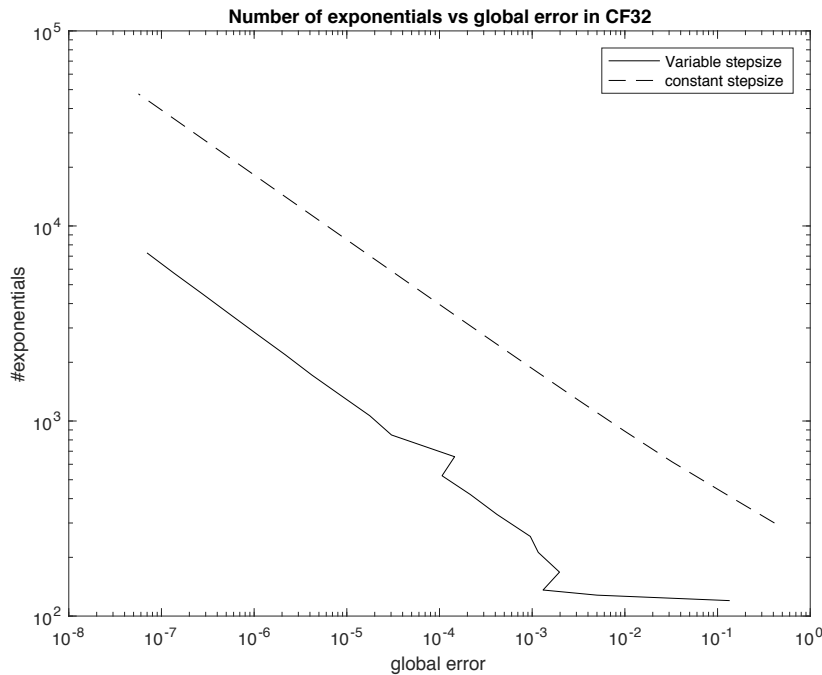


Fig. 3: Numerical integration of the Van der Pol equation with  $\mu = 60$ ,  $y_0 = (1, 1)^T$ . The horizontal axis is the global error at  $t = 1.6$ , and the vertical axis shows the number of exponentials that were computed. The solid line is the variable stepsize method, and the dashed line is the same third order method applied with constant stepsize.

line shows the result for the variable stepsize solver. The cost ratio between the constant and variable stepsize methods depend on the chosen global error, but for instance to obtain a global error of  $10^{-5}$  the constant stepsize integrator needs approximately 6.5 times as many exponential calculations as the variable stepsize method.

### 3.3 The heavy top

Mathematical models for the heavy top can be found in many text books, see for instance [10, 13]. Lie group integrators were applied to this problem in [14]. The heavy top is a rigid body, but because of the gravitational forces, it is not invariant under the action of  $SO(3)$  and the dynamics can therefore not be reduced to  $\mathfrak{so}(3)^*$  as the free rigid body. There is a smaller symmetry group  $S^1$  corresponding to rotation about the vertical axis and it turns out that the system can be formulated on  $\mathfrak{se}(3)^*$ , the dual of the Lie algebra of the special Euclidean group  $SE(3)$  consisting of translations and rotations in

Lie group	$SE(3) \cong SO(3) \ltimes \mathbb{R}^3$
Group product $SE(3)$	$(g, \mathbf{u}) \cdot (h, \mathbf{v}) = (g \cdot h, g \cdot \mathbf{v} + \mathbf{u})$
Inverse $SE(3)$	$(g, \mathbf{u})^{-1} = (g^{-1}, -g^{-1}\mathbf{u})$
Lie algebra	$\mathfrak{se}(3) \cong \mathfrak{so}(3) \ltimes \mathbb{R}^3$
Lie bracket $\mathfrak{se}(3)$	$[(\boldsymbol{\xi}, \mathbf{u}), (\boldsymbol{\eta}, \mathbf{v})] = (\boldsymbol{\xi} \times \boldsymbol{\eta}, \boldsymbol{\xi} \times \mathbf{v} - \boldsymbol{\eta} \times \mathbf{u})$
Manifold	$\mathfrak{se}(3)^* \cong \mathbb{R}^3 \times \mathbb{R}^3$
Coadjoint action by $SE(3)$ on $\mathfrak{se}(3)^*$	$(g, \mathbf{u}) \cdot (\boldsymbol{\mu}, \boldsymbol{\beta}) = \text{Ad}_{(g, \mathbf{u})}^*(\boldsymbol{\mu}, \boldsymbol{\beta})$ $= (g^T(\boldsymbol{\mu} - \mathbf{u} \times \boldsymbol{\beta}), g^T \boldsymbol{\beta})$
Infinitesimal generator of the action	$\lambda_*(\boldsymbol{\xi}, \mathbf{u})(\boldsymbol{\mu}, \boldsymbol{\beta}) = \text{ad}_{(\boldsymbol{\xi}, \mathbf{u})}^*(\boldsymbol{\mu}, \boldsymbol{\beta})$ $= (-\boldsymbol{\xi} \times \boldsymbol{\mu} - \mathbf{u} \times \boldsymbol{\beta}, -\boldsymbol{\xi} \times \boldsymbol{\beta})$
Exponential map	$\exp(t(\boldsymbol{\xi}, \mathbf{u})) = (\exp(t\boldsymbol{\xi}), \frac{\exp(t\boldsymbol{\xi}) - I}{t\boldsymbol{\xi}} \cdot t\mathbf{u})$

Table 5: The main properties of the group action for the heavy top equations

3-space. According to [10] the equations can be written in the form

$$\begin{aligned} \dot{\boldsymbol{\mu}} &= \boldsymbol{\mu} \times \mathbb{I}^{-1}\boldsymbol{\mu} + \boldsymbol{\beta} \times mg\boldsymbol{\chi} \\ \dot{\boldsymbol{\beta}} &= \boldsymbol{\beta} \times \mathbb{I}^{-1}\boldsymbol{\mu} \end{aligned} \quad (3.4)$$

Here  $\boldsymbol{\mu}$  is the body angular momentum and  $\boldsymbol{\beta}$  is the vertical direction as seen from the rotating body, more precisely  $\boldsymbol{\beta} = R^T \mathbf{e}_3$  where  $R$  is the attitude matrix of the top.  $\boldsymbol{\chi}$  is the unit vector in the direction from the fixed point to the center of mass of the heavy top,  $m$  is the mass,  $g$  the constant of gravity, and  $\mathbb{I}$  is the inertia tensor.

In this example, the coadjoint orbits are preserved, and for this reason it is natural to invoke Lie group integrators via the (right) coadjoint action of  $SE(3)$  on  $\mathfrak{se}(3)^*$ . It is convenient to take elements of both  $\mathfrak{se}(3)$  and  $\mathfrak{se}(3)^*$  to be vectors in  $\mathbb{R}^3 \times \mathbb{R}^3$ . Similarly, elements of  $SE(3)$  are represented as pairs  $(g, \mathbf{u})$  where  $g$  could be an orthogonal  $3 \times 3$ -matrix and  $\mathbf{u} \in \mathbb{R}^3$ .

The coadjoint action is the map  $A : SE(3) \times \mathfrak{se}(3)^* \rightarrow \mathfrak{se}(3)^*$

$$A((g, \mathbf{u}), (\boldsymbol{\mu}, \boldsymbol{\beta})) = \text{Ad}_{(g, \mathbf{u})}^*(\boldsymbol{\mu}, \boldsymbol{\beta}) = (g^T(\boldsymbol{\mu} - \mathbf{u} \times \boldsymbol{\beta}), g^T \boldsymbol{\beta})$$

whose infinitesimal generator is the map  $\lambda_* : \mathfrak{se}(3) \rightarrow \mathcal{X}(\mathfrak{se}(3)^*)$

$$\lambda_*(\boldsymbol{\xi}, \mathbf{u})(\boldsymbol{\mu}, \boldsymbol{\beta}) = \text{ad}_{(\boldsymbol{\xi}, \mathbf{u})}^*(\boldsymbol{\mu}, \boldsymbol{\beta}) = (-\boldsymbol{\xi} \times \boldsymbol{\mu} - \mathbf{u} \times \boldsymbol{\beta}, -\boldsymbol{\xi} \times \boldsymbol{\beta})$$

For the heavy top equations (3.4) we now have

$$(\dot{\boldsymbol{\mu}}, \dot{\boldsymbol{\beta}}) = \lambda_*(\mathbb{I}^{-1}\boldsymbol{\mu}, mg\boldsymbol{\chi})(\boldsymbol{\mu}, \boldsymbol{\beta})$$

so that with reference to (2.1) we have

$$f(\boldsymbol{\mu}, \boldsymbol{\beta}) = (\mathbb{I}^{-1}\boldsymbol{\mu}, mg\boldsymbol{\chi})$$

The key properties of the group action are summarised in Table 5. The result of applying a CF32 and a CF43 pair on the Kovalevskaya top (a particular instance of heavy top, notable for giving an integrable system) are detailed in Figure 4.




Fig. 4: The CF32 pair (bottom left, Table 1) and the CF43 pair (Table 4) applied to the Kovalevskaya heavy top on the time interval  $[0, 2]$ . The graph on the left shows the linear relationship between the prescribed error tolerance and the observed global error, whilst the right graph compares the computational cost (measured in number of exponentials) against global error.

## References

1. H. Berland, B. Owren, and B. Skaflestad. *B-series and order conditions for exponential integrators*. *SIAM J. Numer. Anal.*, 43(4):1715–1727, 2005.
2. E Celledoni, A. Marthinsen, and B. Owren. Commutator-free Lie group methods. *FUTURE GENERATION COMPUTER SYSTEMS*, 19(3):341–352, APR 2003.
3. E. Celledoni, H. Marthinsen, and B. Owren. An introduction to Lie group integrators—basics, new developments and applications. *J. Comput. Phys.*, 257(part B):1040–1061, 2014.
4. P. E. Crouch and R. Grossman. Numerical integration of ordinary differential equations on manifolds. *J. Nonlinear Sci.*, 3(1):1–33, 1993.
5. C. Curry and S. Schmeding. Convergence of Lie group integrators. arXiv:1807.11829, 2018.
6. J. R. Dormand and P. J. Prince. A family of embedded Runge-Kutta formulae. *J. Comput. Appl. Math.*, 6(1):19–26, 1980.
7. S. Faltinsen. Backward error analysis for Lie-group methods. *BIT*, 40(4):652–670, 2000.
8. E. Hairer, Ch. Lubich, and G. Wanner. *Geometric numerical integration*, volume 31 of *Springer Series in Computational Mathematics*. Springer, Heidelberg, 2010. Structure-preserving algorithms for ordinary differential equations, Reprint of the second (2006) edition.
9. E. Hairer, S. P. Nørsett, and G. Wanner. *Solving ordinary differential equations. I*, volume 8 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin, second edition, 1993. Nonstiff problems.
10. D. D. Holm. *Geometric mechanics. Part II. Rotating, translating and rolling*. Imperial College Press, London, second edition, 2011.
11. A. Iserles, H. Z. Munthe-Kaas, S. P. Nørsett, and A. Zanna. Lie-group methods. In *Acta numerica, 2000*, volume 9 of *Acta Numer.*, pages 215–365. Cambridge Univ. Press, Cambridge, 2000.
12. A. Lundervold and H. Z. Munthe-Kaas. On algebraic structures of numerical integration on vector spaces and manifolds. In *Faà di Bruno Hopf algebras, Dyson-Schwinger equations, and Lie-Butcher series*, volume 21 of *IRMA Lect. Math. Theor. Phys.*, pages 219–263. Eur. Math. Soc., Zürich, 2015.

13. J. E. Marsden and T. S. Ratiu. *Introduction to mechanics and symmetry*, volume 17 of *Texts in Applied Mathematics*. Springer-Verlag, New York, second edition, 1999. A basic exposition of classical mechanical systems.
14. A. Marthinsen, H. Z. Munthe-Kaas, and B. Owren. Simulation of ordinary differential equations on manifolds: some numerical experiments and verifications. *Modeling, Identification and Control*, 18(1):75–88, 1997.
15. B. Owren. Order conditions for commutator-free Lie group methods. *J. Phys. A*, 39(19):5585–5599, 2006.
16. B. Owren and A. Marthinsen. Runge-Kutta methods adapted to manifolds and based on rigid frames. *BIT*, 39(1):116–142, 1999.
17. L. F. Shampine and M. W. Reichelt. The MATLAB ODE suite. *SIAM J. Sci. Comput.*, 18(1):1–22, 1997. Dedicated to C. William Gear on the occasion of his 60th birthday.
18. J. Wensch. Extrapolation methods in Lie groups. *Numer. Math.*, 89(3):591–604, 2001.