BIT Numer Math (2009) 49: 177–197
DOI 10.1007/s10543-009-0214-3

BIT

Linearization-preserving self-adjoint and symplectic integrators

R.I. McLachlan · G.R.W. Quispel · P.S.P. Tse

Received: 1 May 2008 / Accepted: 20 January 2009 / Published online: 10 February 2009

Abstract This article is concerned with geometric integrators which are linearization-preserving, i.e. numerical integrators which preserve the exact linearization at every fixed point of an arbitrary system of ODEs. For a canonical Hamiltonian system, we propose a new symplectic and self-adjoint B-series method which is also linearization-preserving. In a similar fashion, we show that it is possible to construct a self-adjoint and linearization-preserving B-series method for an arbitrary system of ODEs. Some numerical experiments on Hamiltonian ODEs are presented to test the behaviour of both proposed methods.

Keywords Hamiltonian systems · Linearization-preserving methods · Symplectic integrators

Mathematics Subject Classification (2000) 37J15 · 65P10 · 65P30

Communicated by Christian Lubich.

This work was supported by the Australian Research Council and by the Marsden Fund of the Royal Society of New Zealand.

R.I. McLachlan
IFS, Massey University, Palmerston North, New Zealand
e-mail: R.McLachlan@massey.ac.nz

G.R.W. Quispel
Department of Mathematics, Centre of Excellence for Mathematics and Statistics of Complex Systems, La Trobe University, Melbourne, Victoria 3086, Australia
e-mail: R.Quispel@latrobe.edu.au

P.S.P. Tse (✉)
LSEC, Academy of Mathematics and Systems Science, Chinese Academy of Sciences, P.O. Box 2719, Beijing 100190, China
e-mail: ptse@lsec.cc.ac.cn

1 Introduction

Geometric integration is the numerical integration of systems of differential equations preserving exactly one or more of the structures existing in the differential equations. Such intrinsic[1] geometric structures can be found in many dynamical systems, where for example, many conservative and reversible Hamiltonian systems possess geometric properties such as the symplectic structure, the preservation of energy, phase-space volume, symmetries and reversing symmetries. The preservation of such geometric structures by numerical methods often leads to better long-time qualitative results in the phase space of the differential equations that other non-geometric numerical methods may not guarantee. More details on the theory of geometric integration and its many physical applications can be found in [4, 12, 16, 17].

Aside from the preservation of intrinsic geometric properties, we can also consider the dynamical behaviours of the numerical integrators (cf. [17]). The dynamics of a system in the neighbourhood of a fixed point is generally dominated by the eigenvalues λ of the linearized system at the fixed point. The time-h flow has a fixed point with eigenvalues $e^{h\lambda}$. However, the map generated by an integrator may have different eigenvalues, even if it does preserve the fixed point. For example, a Runge–Kutta method with stability function $R(z)$ (cf. [5]) maps the eigenvalue λ of the linearized system by $\lambda \mapsto R(h\lambda)$. For systems with bifurcations, such perturbations of the eigenvalues due to an integrator may affect the bifurcation parameter value or the bifurcation type observed in the numerical solutions.

An integrator may display one of the following dynamical behaviours with regards to bifurcations: the bifurcation type (e.g. Hopf) and the bifurcation parameter are both captured correctly; the bifurcation type is captured correctly but the bifurcation parameter is shifted; neither the bifurcation type nor the bifurcation parameter is captured correctly. For example, the self-adjoint Gauss methods capture most bifurcations (type and parameter) correctly. This is because at the bifurcation parameter value, the stability function of a Gauss method has unit modulus as the eigenvalues of the linearized system 'cross' the imaginary axis of the complex plane.

On the other hand, Gauss methods do *not* correctly capture the bifurcations associated with all eigenvalue resonances.[2] That is, for a fixed point with two eigenvalues λ_i and λ_j, Gauss methods do *not* correctly capture the bifurcation when $\lambda_i/\lambda_j \in \mathbb{Q}\backslash\{0, \pm 1\}$ (cf. [17, 21]). It is also conceivable, although not very common, that a discretization may cause a non-resonant bifurcation to become resonant or near-resonant. In order to preserve such bifurcations, we require that the integrator be linearization-preserving (defined in the next section). This property ensures that the integrator correctly captures all eigenvalues of the linearized system at fixed points, resulting in the correct preservation of the local dynamics before, during and after a bifurcation anywhere in the phase space.

In this article, we address the question of whether there exists a linearization-preserving numerical integrator which preserves other geometric properties common in

[1] By intrinsic we mean that the geometric property is coordinate-independent and the flow of such systems inherits these properties (cf. [17]).

[2] Please refer to [1] for the definition of eigenvalue resonances in Hamiltonian systems.

dynamical systems with structure. In particular, based on the construction of a pre-processed vector field, we propose two linearization-preserving geometric integrators which preserve the symplectic structure and any affine reversing symmetry group of an arbitrary Hamiltonian system.[3] The first linearization-preserving integrator is a symplectic and self-adjoint B-series method which can only be applied to canonical Hamiltonian ODEs. The second linearization-preserving integrator is a self-adjoint B-series method which can be applied to any ODE. At the end of this article, we present the implementation of these new integrators on two Hamiltonian ODEs: the Hénon–Heiles system and a one-parameter Hamiltonian system which undergoes a Hamiltonian Hopf bifurcation.

2 Linearization-preserving geometric integrators

Consider a system of ODEs $\dot{y} = f(y)$ and a set of initial conditions y_0. This system has the following Taylor expansion at a fixed point $\overline{y}$,

$$\dot{y} = f(y) = \partial f(\overline{y})(y - \overline{y}) + O(\|y - \overline{y}\|^2),$$

where $\partial f(\overline{y})$ denotes the Jacobian of f at $\overline{y}$. The Taylor expansion of the solution at $\overline{y}$ is

$$y(t) - \overline{y} = e^{t\partial f(\overline{y})}(y_0 - \overline{y}) + O(\|y_0 - \overline{y}\|^2). \tag{2.1}$$

Motivated by (2.1), we introduce the following definition.

Definition 2.1 Let an ODE and its initial condition be denoted by $\dot{y} = f(y)$ and y_0. We assume this ODE has fixed points. An integrator $\Phi_h : y_0 \mapsto y_1 \approx y(t_1)$ with step-size h is *linearization-preserving at a fixed point* $\overline{y}$ if it has the following Taylor expansion at $\overline{y}$,

$$y_1 - \overline{y} = e^{h\partial f(\overline{y})}(y_0 - \overline{y}) + O(\|y_0 - \overline{y}\|^2).$$

Furthermore, an integrator is said to be *linearization-preserving* if it is linearization-preserving at $\overline{y}$ for all fixed points $\overline{y}$ of $f(y)$.

To illustrate Definition 2.1, we consider an ODE which has a linear and a nonlinear part, $\dot{y} = Ly + N(y)$, with $\partial N(0) = 0$. In this case, a linear-nonlinear splitting method applied to this ODE is linearization-preserving at 0 but it is not linearization-preserving for all fixed points. In contrast, a B-series method (cf. [12]) which is linearization-preserving at 0 is linearization-preserving by linear covariance of B-series methods.

By Definition 2.1, it is natural to consider the exponential integrators[4] as the likely candidates for methods which are linearization-preserving. For example, consider the

[3] Recently, linearization-preserving integrators that also preserve a first integral have been constructed (cf. [18]).

[4] We use the term "exponential integrator" to describe a rather large class of numerical integrators in which the matrix exponential appears somewhere in the formulation of the method (cf. [19]).

order-two *exponentially fitted Euler method* (cf. [15, 19]),

$$y_1 = y_0 + h\varphi(h\partial f(y_0))f(y_0),$$

where h is the stepsize and the matrix function φ is given by

$$\varphi(z) = \frac{\exp(z) - 1}{z}. \tag{2.2}$$

To show that this method is linearization-preserving by Definition 2.1, we consider the following.

Let $\overline{y}$ be an arbitrary fixed point and let the initial condition y_0 be in the neighbourhood of $\overline{y}$, i.e. $y_0 = \overline{y} + \delta y_0$ for small δy_0. Then the Taylor expansion of $f(y_0)$ and its Jacobian $\partial f(y_0)$ at $\overline{y}$ are,

$$f(y_0) = \partial f(\overline{y})(y_0 - \overline{y}) + \frac{1}{2!}\partial^2 f(\overline{y})(y_0 - \overline{y}, y_0 - \overline{y}) + \cdots, \tag{2.3}$$

$$\partial f(y_0) = \partial f(\overline{y}) + \partial^2 f(\overline{y})(y_0 - \overline{y}) + \cdots. \tag{2.4}$$

By (2.4), (2.2) has the expansion at $\overline{y}$ (for $z = h\partial f(y_0)$),

$$\varphi(h\partial f(y_0)) = I + \frac{h}{2}\partial f(\overline{y}) + \frac{h^2}{3!}\partial f(\overline{y})^2 + \cdots + O(\|y_0 - \overline{y}\|). \tag{2.5}$$

Substitution of (2.3) and (2.5) into the exponentially fitted Euler method shows that the method is linearization-preserving at the fixed point $\overline{y}$. Since the exponentially fitted Euler method is a B-series method, it is linearization-preserving by linear covariance.

More generally, the exponentially fitted Euler method belongs to a class of exponential Runge–Kutta methods proposed by Hochbruck, Lubich and Selhofer in [15]. Such methods are B-series methods and they are linearization-preserving if their coefficients satisfy the order conditions for all *tall* rooted trees given in [15]. In the proposition below we give the condition under which a B-series method is linearization-preserving. Note that $\mathcal{T}$ denotes the set of all rooted trees (cf. [5]) excluding the empty tree $\emptyset$, and $\mathcal{TT} \subset \mathcal{T}$ denotes the set of all *tall* rooted trees from order one. That is,

$$\mathcal{TT} = \left\{ \bullet, \begin{array}{c}\bullet\\|\\\bullet\end{array}, \begin{array}{c}\bullet\\|\\\bullet\\|\\\bullet\end{array}, \ldots \right\}.$$

Proposition 2.1 *Let $B_f(a, y)$ with $a(\emptyset) = 1$ be a B-series arising from a B-series method applied to a vector field $f(y)$. Suppose the coefficients a of the B-series for all* tall *rooted trees $t \in \mathcal{TT}$ satisfy*

$$a(t) = \frac{1}{\gamma(t)}, \quad \forall t \in \mathcal{TT},$$

where γ is the density function (cf. [5]) for all rooted trees $t \in \mathcal{T}$. This B-series method is linearization-preserving.

Proof Consider Taylor expansions of the vector field $f(y)$ and its Jacobian $\partial f(y)$ at a fixed point $\overline{y}$, given by (2.3) and (2.4) with y_0 replaced by y. Only the elementary differentials of the *tall* trees in $\mathcal{TT}$ will contribute to the linear term in $y - \overline{y}$. All other elementary differentials contribute to higher order terms. Hence we can rewrite the B-series $B_f(a, y)$ (when applied to the expansion of $f(y)$ at $\overline{y}$) as,

$$B_f(a, y) - \overline{y} = y - \overline{y} + \sum_{t \in \mathcal{TT}} h^{r(t)} a(t) \mathcal{F}_f(t)(y - \overline{y}) + O(\|y - \overline{y}\|^2),$$

where $a(\emptyset) = 1$. Since all tall tree coefficients are exact,

$$\begin{aligned} B_f(a, y) - \overline{y} &= y - \overline{y} + \sum_{t \in \mathcal{TT}} \frac{h^{r(t)}}{r(t)!} \mathcal{F}_f(t)(y - \overline{y}) + O(\|y - \overline{y}\|^2) \\ &= \left[I + h\partial f(\overline{y}) + \frac{h^2}{2} \partial f(\overline{y})^2 + \cdots \right] (y - \overline{y}) + O(\|y - \overline{y}\|^2) \\ &= e^{h\partial f(\overline{y})}(y - \overline{y}) + O(\|y - \overline{y}\|^2). \end{aligned}$$

For every other fixed point of the vector field $f(y)$, the proposition holds by virtue of the covariance with respect to affine transformations by a B-series method with $a(\emptyset) = 1$ (cf. [21]). That is, a B-series method applied to the Taylor expansions of the same ODE at two different fixed points $\overline{y}_1, \overline{y}_2$ differ only by the same transformation. Hence, if a B-series method is linearization-preserving at one fixed point $\overline{y}$, it is linearization-preserving at every other fixed point of the ODE. □

Since a linearization-preserving B-series method guarantees that the coefficients for all tall rooted trees are exact, we have the following corollary.

Corollary 2.1 *A B-series method which is linearization-preserving is exact for all linear and for all affine systems.*

Thus, linearization-preservation delivers an enormous quantitative (as well as qualitative) advantage in integration of orbits that stay near (neutrally stable) fixed points for a long time, where the nonlinear terms in the vector field are small.

As mentioned previously, there exist exponential Runge–Kutta methods in [15] with B-series coefficients which satisfy Proposition 2.1. Such exponential integrators are linearization-preserving. However, not all exponential integrators are linearization-preserving. For example, the exponential general linear methods proposed by Minchev and Wright in [19] are exact for linear ODEs. But such methods use an inexact Jacobian, so they are not linearization-preserving at the origin or at any other fixed point of an arbitrary vector field.

In the following section, we present the construction of a fourth-order integrator, denoted Υ, which is symplectic, self-adjoint and linearization-preserving for any Hamiltonian system of ODEs. This new method Υ is a preprocessed vector field in-

tegrator[5] because given an original Hamiltonian vector field $f(y)$, we construct and integrate the preprocessed Hamiltonian vector field $\widetilde{F}(y)$ using a simple integrator, say Ψ. Formally, the numerical solution from $\Psi_{\widetilde{F}}(y)$ is equal to the more elaborate integrator Υ_f applied directly to $f(y)$. That is,

$$\Upsilon_f(y) = \Psi_{\widetilde{F}}(y).$$

Note that our method $\Upsilon_f(y)$ is different from the 'modifying' integrator proposed by Chartier, Hairer and Vilmart in [9]. In their paper, an arbitrary vector field $f(y)$ is modified.[6] Then the truncated version of the modified vector field is integrated using a low order geometric integrator. However, our preprocessed vector field $\widetilde{F}(y)$ is constructed to preserve the linearization defined in Definition 2.1 and it is not a truncation of another modified vector field. Moreover, our proposed method Υ_f is defined only for canonical Hamiltonian systems.

3 A symplectic, self-adjoint and linearization-preserving integrator $\Upsilon_f(y)$

In this section we first discuss the construction of the preprocessed vector field $\widetilde{F}(y)$ such that the preprocessed vector field integrator $\Upsilon_f(y)$ is exact for a linear Hamiltonian ODE $f(y)$. Then we extend the construction to nonlinear Hamiltonian ODEs so that the integrator $\Upsilon_f(y)$ becomes linearization-preserving. We shall defer a discussion of the geometric properties of method $\Upsilon_f(y)$ until Sect. 4.

3.1 Derivation of the method $\Upsilon_f(y)$ for linear ODEs

The aim in this section is to construct a preprocessed vector field integrator which is exact for linear Hamiltonian ODEs. For $y = (p, q) \in \mathbb{R}^{2d}$, let the linear Hamiltonian system be given by

$$\dot{y} = f(y) = J^{-1}\nabla H(y) = J^{-1}Sy = Ay, \qquad y(0) = y_0, \tag{3.1}$$

where S is a symmetric matrix and matrix $A = J^{-1}S$. The associated Hamiltonian function is quadratic,

$$H(y) = \frac{1}{2}y^T Sy,$$

and the exact solution of this linear Hamiltonian ODE is $y(t) = e^{tA}y_0$.

Because the original vector field $f(y)$ is Hamiltonian, we wish to construct a preprocessed vector field $\widetilde{F}(y)$ which is also Hamiltonian. Let $\widetilde{H}(y)$ denote the

[5] Other examples of preprocessed vector field integrators include generating function methods for Hamiltonian systems (cf. [6, 9, 11]).

[6] The modified vector field constructed by Chartier et al. in [9] is also different from the modified vector field in backward error analysis. The reader is referred to [9] for more detail.

preprocessed Hamiltonian for $\widetilde{F}(y)$. We let $\widetilde{H}(y)$ depend on the original quadratic Hamiltonian function $H(y)$ in the following way[7]

$$\widetilde{H}(y) = H(Y(y)), \quad \text{where } Y = y + h\Phi(hA)Ay,$$

and the matrix function Φ is yet to be determined. The Jacobian of vector Y is given by

$$\frac{\partial Y}{\partial y} = I + h\Phi(hA)A.$$

The preprocessed Hamiltonian $\widetilde{H}(y)$ leads to the construction of the preprocessed vector field $\widetilde{F}(y)$,

$$\dot{y} = \widetilde{F}(y) = J^{-1}\nabla\widetilde{H}(y) = J^{-1}\left(\frac{\partial Y}{\partial y}\right)^T \nabla_Y H(Y).$$

By the following relations between the matrices A and S in (3.1),

$$A = J^{-1}S, \quad \text{and} \quad A^T = -SJ^{-1},$$

and the assumption that Φ is an analytic matrix function with an expansion in powers of hA,

$$\Phi(hA) = a_0 + a_1 hA + a_2 h^2 A^2 + \cdots, \quad \text{where } a_0, a_1, \ldots \in \mathbb{R}, \tag{3.2}$$

we can show that

$$\Phi(hA)^T S = S\Phi(-hA).$$

Hence we can derive the preprocessed vector field $\widetilde{F}(y)$, which is linear in y,

$$\dot{y} = \widetilde{F}(y) = Fy,$$

where F is given by

$$F = A + hA^2\Phi(hA) - hA^2\Phi(-hA) - h^2A^3\Phi(-hA)\Phi(hA). \tag{3.3}$$

For our method to be self-adjoint (especially if we apply a B-series method to $\widetilde{F}(y)$), we require $\widetilde{F}(y,h)$ to have an expansion in even powers of h. That is, $\widetilde{F}(y,-h) = \widetilde{F}(y,h)$. This means the matrix function Φ must be an odd function. In this case we can further reduce the matrix F to,

$$F = A + 2hA^2\Phi(hA) + h^2A^3\Phi(hA)^2. \tag{3.4}$$

[7]The choice of $Y = \Phi(hA)y$ also yields an integrator which is exact for linear Hamiltonian ODEs. By following the derivation in Sect. 3.2, the integrator with $Y = \Phi(h\partial f(y))y$ is linearization-preserving, symplectic and self-adjoint for nonlinear Hamiltonian ODEs. But it is not a B-series method. We thank the referee for this comment.

Recall that a preprocessed vector field integrator $\Upsilon_f(y)$ can be thought of as the integration of the preprocessed vector field $\widetilde{F}(y)$ using a simple integrator $\Psi_{\widetilde{F}}(y)$ such that formally,

$$\Upsilon_f(y) = \Psi_{\widetilde{F}}(y).$$

Since the original vector field $f(y)$ and the preprocessed vector field $\widetilde{F}(y)$ are Hamiltonian, we choose to integrate $\widetilde{F}(y)$ using a symplectic method. To keep the following calculation simple, we use the order-two implicit midpoint rule, which also happens to be self-adjoint (cf. [12]). The implicit midpoint rule has the following stability function (cf. [5]),

$$R(z) = \frac{1+\frac{z}{2}}{1-\frac{z}{2}}, \quad \text{where } z \in \mathbb{C}.$$

In order to find the matrix function $\Phi(hA)$ such that our method is exact for the original linear ODE in (3.1), recall that the exact solution is $y_1 = e^{tA}y_0$. Hence,

$$R(hF) = \frac{I+\frac{hF}{2}}{I-\frac{hF}{2}} = e^{hA} \quad \Longrightarrow \quad F = -\frac{2}{h}\left(\frac{I-e^{hA}}{I+e^{hA}}\right).$$

Substituting F given by (3.4) into the above, this implies that the matrix coefficient $\Phi(hA)$ in the method must satisfy the following equation,

$$hA + 2\left(\frac{I-e^{hA}}{I+e^{hA}}\right) + 2h^2A^2\Phi(hA) + h^3A^3\Phi(hA)^2 = 0. \tag{3.5}$$

To solve for $\Phi(hA)$ in (3.5), first multiply (3.5) by $\frac{1}{hA}$ and rewrite the equation as,

$$(hA\Phi(hA) + I)^2 = -\frac{2}{hA}\left(\frac{I-e^{hA}}{I+e^{hA}}\right).$$

The solution $\Phi(hA)$ can be derived directly,

$$\Phi(hA) = \frac{1}{hA}\left(\sqrt{-\frac{2}{hA}\left(\frac{I-e^{hA}}{I+e^{hA}}\right)} - I\right). \tag{3.6}$$

Hence, the integration of the preprocessed vector field $\widetilde{F}(y) = Fy$ with F in (3.4) using the implicit midpoint rule, yields the exact solution of the original linear Hamiltonian ODE given by (3.1).

3.2 Derivation of the method $\Upsilon_f(y)$ for nonlinear ODEs

Now we consider a nonlinear canonical Hamiltonian ODE,

$$\dot{y} = f(y) = J^{-1}\nabla H(y),$$

where $H(y)$ is a Hamiltonian function of degree higher than two. To generalize the method from the previous section for this nonlinear ODE, while retaining the

linearization-preserving property, we require that the matrix function Φ be given as in (3.6), but that its argument contain the exact Jacobian $\partial f(y)$ from the nonlinear ODE, i.e.,

$$\Phi(h\partial f(y)) = \frac{1}{h\partial f(y)}\left(\sqrt{-\frac{2}{h\partial f(y)}\left(\frac{I - e^{h\partial f(y)}}{I + e^{h\partial f(y)}}\right)} - I\right). \tag{3.7}$$

The vector Y is now

$$Y = y + h\Phi(h\partial f(y))f(y), \tag{3.8}$$

and the Jacobian of $Y(y)$ is given by

$$\frac{\partial Y}{\partial y}(y) = I + h\Phi(h\partial f(y))\frac{\partial f}{\partial y}(y) + hD\Phi(h\partial f(y))(I \otimes f(y)). \tag{3.9}$$

For a $n \times n$ matrix Φ, the notation $D\Phi(h\partial f(y))$ represents the following $n \times n^2$ matrix, and $(I \otimes f(y))$ an $n^2 \times n$ matrix,

$$\begin{aligned} D\Phi(h\partial f(y)) &= \left[\begin{array}{cccc} \frac{\partial}{\partial y_1}\Phi(h\partial f(y)) & \frac{\partial}{\partial y_2}\Phi(h\partial f(y)) & \dots & \frac{\partial}{\partial y_n}\Phi(h\partial f(y)) \end{array}\right], \\ I \otimes f(y) &= \left[\begin{array}{ccccc} f(y) & 0_{n\times 1} & & \dots & \\ 0_{n\times 1} & f(y) & 0_{n\times 1} & & \\ & 0_{n\times 1} & \ddots & & \\ & & & & f(y) \end{array}\right]. \end{aligned} \tag{3.10}$$

In (3.10), $\frac{\partial}{\partial y_j}\Phi(h\partial f(y))$ refers to the component-wise partial derivatives of the elements of matrix Φ with respect to y_j. Hence $\frac{\partial}{\partial y_j}\Phi(h\partial f(y))$ is itself a $n \times n$ matrix.

Finally, the preprocessed Hamiltonian ODE with $\widetilde{H}(y) = H(Y(y))$ has the expression

$$\dot{y} = \widetilde{F}(y) = J^{-1}\nabla_y \widetilde{H}(y) = J^{-1}\left(\frac{\partial Y}{\partial y}(y)\right)^T \nabla_Y H(Y), \tag{3.11}$$

where Y is given by (3.8). This preprocessed vector field $\widetilde{F}(y)$ in (3.11) is integrated by the implicit midpoint rule,

$$Z = \frac{h}{2}\widetilde{F}(y_n + Z), \qquad y_{n+1} = y_n + 2Z. \tag{3.12}$$

4 Geometric properties of the method $\Upsilon_f(y)$

In this section we verify the properties of the method $\Upsilon_f(y)$ proposed in Sect. 3.2. For the rest of this section, the preprocessed vector field integrator $\Upsilon_f(y)$ is constructed in the following way.

Definition 4.1 Let the canonical Hamiltonian vector field with a (differentiable) Hamiltonian $H(y)$ be denoted by $f(y)$. The preprocessed vector field $\widetilde{F}(y)$ based on $f(y)$ is given by (3.11), where this expression depends on the matrix function $\Phi(h\partial f(y))$ in (3.7).

Furthermore, if $\Psi_{\widetilde{F}}(y)$ denotes the implicit midpoint rule which integrates $\widetilde{F}(y)$, then the entire integration process can be considered as a new integrator, $\Upsilon_f(y)$, acting on the original $f(y)$:

$$\Upsilon_f(y) = \Psi_{\widetilde{F}}(y).$$

In the proposition below, we make use of the fact that the preprocessed vector field method Υ_f is a B-series method $B_f(a, y)$ with $a(\emptyset) = 1$. This is because the preprocessed Hamiltonian given in Sect. 3.1 is an expansion in elementary Hamiltonians (cf. [12]). Alternately, that Υ_f is a B-series method is a direct consequence of results based on the substitution law $\star$ for B-series developed by Chartier, Hairer and Vilmart in [8].

Proposition 4.1 *The numerical method $\Upsilon_f(y)$, defined as in Definition 4.1 for a canonical Hamiltonian vector field $f(y)$, has the following properties*:

1. *it is a symplectic integrator,*
2. *it preserves all quadratic invariants associated with $f(y)$,*
3. *it is covariant with respect to any affine transformation,*
4. *it preserves any affine symmetry S of a dynamical system,*
5. *it is a self-adjoint integrator,*
6. *it preserves any affine reversing symmetry group of a dynamical system,*
7. *it preserves every fixed point $\overline{y}$ of $f(y)$,*
8. *it is linearization-preserving at every fixed point $\overline{y}$ of $f(y)$,*
9. *it is a fourth-order method.*

Proof We shall prove the statements in Proposition 4.1 by referring to existing results concerning B-series methods.

1. The preprocessed ODE $\dot{y} = \widetilde{F}(y)$ is a Hamiltonian system which has a preprocessed Hamiltonian $\widetilde{H}(y)$. Hence integrating $\widetilde{F}(y)$ by a symplectic method (e.g. implicit midpoint rule), yields a symplectic method. This can also be obtained directly from Corollary 4.23 in [8].
2. The preservation of all quadratic invariants comes from a result Chartier, Faou and Murua in [7].

3–4. From a result in [21], since $\Upsilon_f(y)$ is a B-series method.

5. Self-adjoint property comes from the preprocessed vector field $\widetilde{F}(y, h)$ being symmetric in h, i.e. $\widetilde{F}(y, h)$ is of even powers in h so that $\widetilde{F}(y, -h) = \widetilde{F}(y, h)$.
6. By virtue of being a self-adjoint B-series method (cf. [21]).
7. By being a B-series method with $a(\emptyset) = 1$.
8. By construction Υ_f is exact for all linear ODEs, hence it satisfies Proposition 2.1 and the result follows.

9. Since Υ_f is symplectic and linearization-preserving, it must satisfy the third-order symplecticity condition and the third-order condition for the tall tree. This implies Υ_f is of order three. Order four is by the symmetry of the method.[8]

□

More generally, with the same argument as given in 9. from the proof of Proposition 4.1, we have the following remark concerning the orders of linearization-preserving integrators.

Remark 4.1 All linearization-preserving B-series methods are (at least) order two. All symplectic and linearization-preserving B-series methods are (at least) order three. All self-adjoint, symplectic and linearization-preserving B-series methods are (at least) order four.

5 A self-adjoint and linearization-preserving integrator

The idea of preprocessed vector fields can also be used to construct linearization-preserving integrators which are self-adjoint. For this type of self-adjoint integrators, we are no longer restricted to constructing a preprocessed Hamiltonian vector field which is dependent on the Hamiltonian $H(y)$ from the original vector field. Hence the preprocessed vector field integrator can now be applied to an arbitrary vector field $f(y)$.

Let $f(y)$ denote an arbitrary vector field. Following the discussion in Sect. 3, we let the preprocessed vector field be

$$\widetilde{F}(y) = f(Y), \quad \text{where } Y = y + h\Phi(z)f(y) \text{ and } z = h\partial f(y). \tag{5.1}$$

If the (self-adjoint) implicit midpoint rule is used to integrate $\widetilde{F}(y)$, then a matrix coefficient function Φ for which the preprocessed vector field integrator is self-adjoint and linearization-preserving is

$$\Phi(z) = \frac{1}{z}\left(\frac{2}{z}\left(\frac{e^z - 1}{e^z + 1}\right) - 1\right).$$

Similarly, a simpler version of the linearization-preserving and self-adjoint integrator can be derived by considering the following preprocessed vector field which is integrated using the implicit midpoint rule,

$$\widetilde{F}(y) = \Phi(z)f(y), \quad \text{where } z = h\partial f(y).$$

In this case, the matrix function is

$$\Phi(z) = \frac{2}{z}\left(\frac{e^z - 1}{e^z + 1}\right).$$

[8] We thank the referee for his comments concerning this proof.

Such integrators are easier to implement than the integrator $\Upsilon_f(y)$ proposed in Sect. 3. However, the price we pay for this simpler construction is the loss of symplecticity.

6 Numerical experiments

In this section we implement the preprocessed vector field integrator proposed in Sect. 3.2. We also compare the numerical results from this new method against a standard geometric integrator (a symplectic Gauss method) and non-standard integrators (exponential general linear methods).

6.1 Implementation of the method $\Upsilon_f(y)$

Given a canonical Hamiltonian vector field $f(y)$, we wish to integrate it numerically using the method $\Upsilon_f(y)$ proposed in Sect. 3.2. Recall that $\Upsilon_f(y)$ is essentially the application of the implicit midpoint rule to the preprocessed vector field $\widetilde{F}(y)$ in (3.11). The most challenging part of the implementation[9] of $\Upsilon_f(y)$ often lies in the construction of the preprocessed vector field $\widetilde{F}(y)$, because it requires the matrices Φ in (3.7) and $D\Phi$ in (3.10).

The complexity of the two matrices depends on the Jacobian matrix $\partial f(y)$ of the original Hamiltonian ODE $\dot{y} = f(y)$. In particular, algebraic expressions for all elements in the matrix Φ are needed in order to partially differentiate them with respect to y for the construction of the matrix $D\Phi$. In the numerical experiments presented in the following sections, the algebraic expressions for the elements in matrix Φ are derived using *Sylvester's formula* for matrix functions (cf. [2]) using MAPLE.

Once the preprocessed vector field $\widetilde{F}(y)$ in (3.11) is constructed for the ODE $\dot{y} = f(y)$, we have to consider how to solve the implicit stage value Z in the implicit midpoint rule (3.12), before the numerical solution y_n can be updated. For our purposes, the fixed point iteration root-finding scheme is preferred over the full Newton iteration in order to avoid finding the full Jacobian of the preprocessed vector field $\widetilde{F}(y)$.

Suppose that the integrator $\Upsilon_f(y)$ has been successfully implemented for a Hamiltonian ODE $\dot{y} = f(y)$. We would like to check the error committed by the integrator throughout the integration period. The *global truncation error* committed by an integrator is approximated by taking the 2-norm of the difference, in double precision, between the numerical solution in quadruple precision and the exact solution of the Hamiltonian vector field $f(y)$. We also find the *error in the Hamiltonian* $H(y)$ from the numerical solution, by taking the difference, between the numerical Hamiltonian $H(y)$ and the Hamiltonian value $H(y_0)$ evaluated at the initial condition y_0 of each problem.

[9] Please refer to [21] for more details on the implementation of method $\Upsilon_f(y)$ and other numerical integrators in the comparison.

In each numerical experiment, we compare the performance of $\Upsilon_f(y)$, which is labelled *method*, against other numerical integration schemes. One numerical integrator in the comparison is the 2-stage symplectic Gauss method of order 4, *gauss4* (cf. [20]). Also included are two different exponential general linear methods (cf. [3, 19]) which use the inexact Jacobian. They are the explicit Lawson exponential method of order 4, *lawson4*, and the explicit Hochbruck–Ostermann exponential method of order 4, *hochost4*. Additionally, for the second Hamiltonian system, we implemented an explicit, 7-stage exponential Runge–Kutta method *exprk4*, from the class of exponential Runge–Kutta methods proposed by Hochbruck, Lubich and Selhofer in [15]. Even though the exponential Runge–Kutta method and the exponential general linear methods in the comparison are not geometric integrators for general ODEs, they are exact for linear ODEs $\dot{y} = Ay$ and, in the case of the exponential Runge–Kutta method *exprk4*, it is a linearization-preserving integrator (cf. [21]).

One important difference between our preprocessed vector field method Υ_f and the established exponential integrators is that exponential integrators such as *exprk4* are designed for long-time-step implementation, whereas our preprocessed integrator Υ_f has stepsize restrictions based on the matrix function $\Phi(z)$ in (3.6).[10] For example, the scalar case has a stepsize restriction of $z \neq i\pi$. However, this restriction is not serious for the numerical experiments presented in this paper since we are not attempting to integrate with stepsizes that are large compared to the frequencies of the system.

We now present an outline of the construction of method $\Upsilon_f(y)$ we have used when applied to a Hamiltonian vector field $f(y)$.

- Construct the matrix Φ in (3.7) using, e.g. MAPLE.
- Construct the matrix $D\Phi$ by differentiating the elements in matrix Φ.
- The expressions in matrices Φ and $D\Phi$ should be simplified as much as possible to reduce the number of operations required.
- Complex arithmetic is used when implementing the matrices Φ and $D\Phi$ due to the presence of square roots in the matrix elements. Since the numerical values of the elements in matrices Φ and $D\Phi$ should be real, we omit any imaginary values in Φ and $D\Phi$ caused by finite precision arithmetic.
- Construct the preprocessed vector field $\widetilde{F}(y)$ in (3.11).
- The implicit midpoint rule is applied to the preprocessed vector field $\widetilde{F}(y)$. To minimize cancellation effects in the numerical computation, we follow the discussion in [13] by rewriting the implicit midpoint rule as (3.12).
- Unless the stage value Z in (3.12) can be found directly from the equations, fixed point iteration is the root-finding method of choice for the implementation of the implicit midpoint rule applied to $\widetilde{F}(y)$.

All numerical experiments presented in the following sections are carried out using the arbitrary precision package ARPREC[11] in C++.

[10] We are grateful for the referee's comment regarding stepsize restrictions.

[11] The arbitrary precision package ARPREC, written by David Bailey, can be downloaded from http://crd.lbl.gov/~dhbailey/mpdist/.

6.2 The Hénon–Heiles system

In this section we numerically integrate the four-dimensional Hénon–Heiles ODE, where $y = (p_1, p_2, q_1, q_2) \in \mathbb{R}^4$. Numerical computation for this system was first done by Hénon and Heiles in [14] and the system describes planar particle motion in a two-dimensional potential. In particular, this Hamiltonian system is of interest because it contains non-trivial dynamics predicted by the KAM theory, and this dynamics can be accurately reproduced over long times by symplectic integrators.

The Hénon–Heiles system has the following Hamiltonian function,

$$H(y) = \frac{1}{2}(p_1^2 + p_2^2) + \frac{1}{2}(q_1^2 + q_2^2) + q_1^2 q_2 - \frac{1}{3}q_2^3,$$

where its four-dimensional vector field $f(y)$ and the Jacobian of $f(y)$ are given by,

$$f(y) = \begin{bmatrix} -q_1 - 2q_1q_2 \\ -q_2 - q_1^2 + q_2^2 \\ p_1 \\ p_2 \end{bmatrix}, \qquad \frac{\partial f}{\partial y}(y) = \begin{bmatrix} 0 & 0 & -1-2q_2 & -2q_1 \\ 0 & 0 & -2q_1 & -1+2q_2 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}.$$

The set of initial conditions used in this experiment is $y_0 = (0, 0, 0.8, -0.48)$, which is chosen to be near the fixed point $\overline{y} = (0, 0, \frac{\sqrt{3}}{2}, -\frac{1}{2})$ of the system.

The most complicated part in the implementation of our integrator *method* lies in the construction of the 4×4 matrix Φ in (3.7) and the 4×16 matrix $D\Phi$ in (3.10), which we derived using MAPLE. Because the Hamiltonian vector field $f(y)$ is nonlinear, the task of finding the non-zero matrices Φ and $D\Phi$ is not trivial.

Figure 1 displays the global truncation error in solution y against time $t = [0, 100]$. It is clear that the numerical solution from our *method* commits the smallest global truncation error over the integration period. We suspect that *method* performs so well in this case because we started the integration close to a fixed point $\overline{y}$, and so the linearization-preserving property contributes to the approximation of the exact solution.

Figure 2 shows that the numerical Hamiltonian values from the symplectic *gauss4* and *method* are bounded over time. However, *lawson4* and *hochost4*, which are not geometric methods, produce numerical Hamiltonian values that drift away from the initial value over time.

Furthermore, the self-adjoint and linearization-preserving integrator proposed in Sect. 5 also behaved similarly in a numerical comparison with the symplectic and self-adjoint *method* $\Upsilon_f(y)$. For the Hénon–Heiles system, the global truncation error in y and the error in the Hamiltonian $H(y)$ from the self-adjoint and linearization-preserving integrator are the same as those from *method* $\Upsilon_f(y)$. Hence we do not show the numerics from the self-adjoint integrator.

Before we leave the Hénon–Heiles system, we must remark that currently the cost of implementing *method* $\Upsilon_f(y)$ is prohibitively expensive due to the evaluation of matrices Φ and $D\Phi$ in (3.7) and (3.10) (see Table 1). Hence *method* $\Upsilon_f(y)$ cannot hope to compete against the other exponential integrators in the comparison. Even though it produces numerics which are qualitatively better than the other integrators in this comparison, its inefficiency is a drawback.

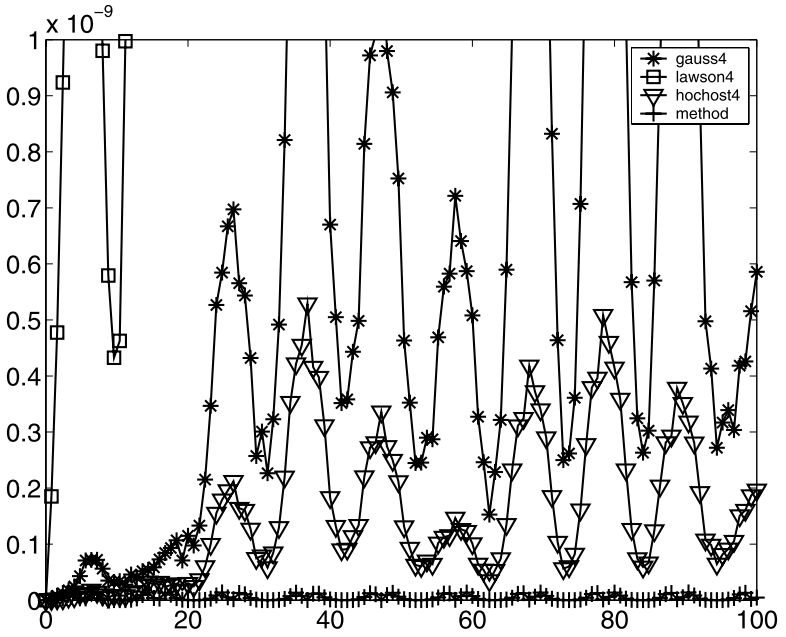

Fig. 1 The global truncation error in y against time $t = [0, 100]$ with stepsize $h = 0.01$, for the Hénon–Heiles system. The integrators in this comparison are *gauss4*, *lawson4*, *hochost4* and *method*

Table 1 The CPU times for the numerical integrators applied to the Hénon–Heiles system, with $y_0 = (0, 0, 0.8, -0.48)$ and stepsize $h = 0.001$ for $t = [0, 5]$

	gauss4	*lawson4*	*hochost4*	*method*
CPU time	46	8	14	5276

6.3 A Hamiltonian system with a Hamiltonian Hopf bifurcation

In this section we compare our *method* against the other integrators on a nonlinear, four-dimensional Hamiltonian ODE which has an extra parameter λ. This system goes through a Hamiltonian Hopf bifurcation (also known as a $1 : -1$ resonance) when $\lambda = 0$. The $1 : -1$ resonance in Hamiltonian system usually occurs when the quadratic part of the Hamiltonian is non-semisimple [22]. This is characterized by the eigenvalues of $\partial f(\overline{y})$ at a fixed point $\overline{y}$ as they collide along the imaginary axis of the complex plane, and either split into the complex plane or pass each other along the imaginary axis [10]. As mentioned previously, although the Gauss methods preserve correctly the parameter value at which this type of resonance bifurcation occurs, they may not preserve the eigenvalues of the linearized system before or after the bifurcation. In contrast, our *method* guarantees the correct eigenvalues for all parameter values.

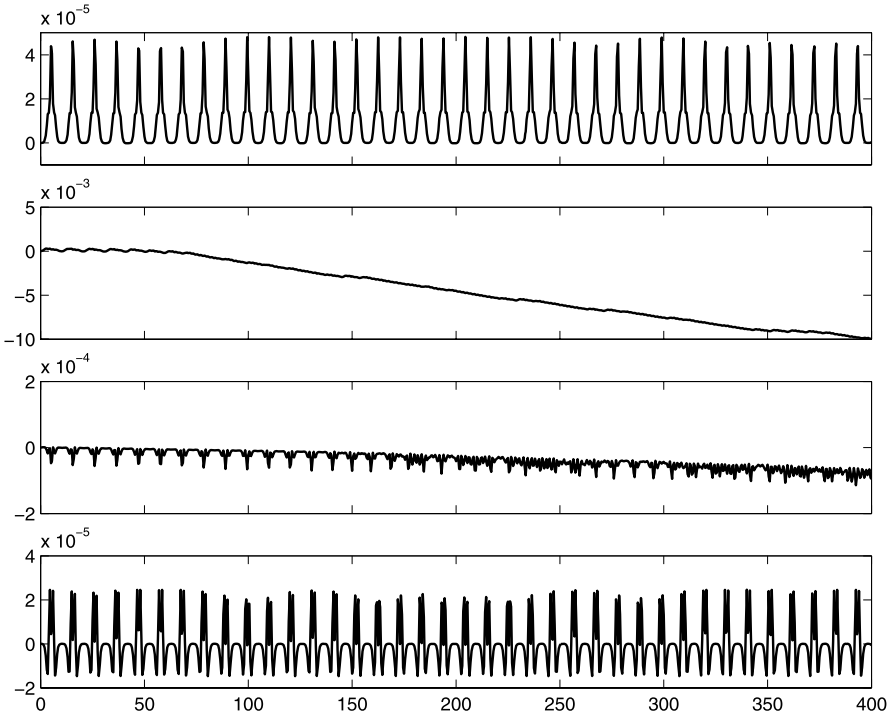

Fig. 2 The error in the Hamiltonian against time $t = [0, 400]$ with stepsize $h = 0.5$ for the Hénon–Heiles system. The integrators in this comparison are *gauss4* (*top*), *lawson4* (*second*), *hochost4* (*third*) and *method* (*bottom*)

We have constructed the following Hamiltonian used in this numerical experiment,

$$H(p_1, p_2, q_1, q_2) = \frac{1}{2}(p_1 q_2 - p_2 q_1 + \lambda p_2^2 - \lambda^2 q_1^2 - p_1^3),$$

where $y = (p_1, p_2, q_1, q_2)$. This Hamiltonian has one fixed point at the origin, $\overline{y} = (0, 0, 0, 0)$, and its vector field $f(y)$ and the corresponding Jacobian are given by,

$$f(y) = \begin{bmatrix} \frac{1}{2}p_2 + \lambda^2 q_1 \\ -\frac{1}{2}p_1 \\ \frac{1}{2}q_2 - \frac{3}{2}p_1^2 \\ -\frac{1}{2}q_1 + \lambda p_2 \end{bmatrix}, \qquad \frac{\partial f}{\partial y}(y) = \begin{bmatrix} 0 & \frac{1}{2} & \lambda^2 & 0 \\ -\frac{1}{2} & 0 & 0 & 0 \\ -3p_1 & 0 & 0 & \frac{1}{2} \\ 0 & \lambda & -\frac{1}{2} & 0 \end{bmatrix}. \tag{6.1}$$

At its fixed point at the origin, the four eigenvalues μ_i of $\partial f(0)$ are,

$$\mu_1 = \frac{\sqrt{-1 + 2\lambda\sqrt{-\lambda}}}{2}, \qquad \mu_2 = -\frac{\sqrt{-1 + 2\lambda\sqrt{-\lambda}}}{2},$$
$$\mu_3 = \frac{\sqrt{-1 - 2\lambda\sqrt{-\lambda}}}{2}, \qquad \mu_4 = -\frac{\sqrt{-1 - 2\lambda\sqrt{-\lambda}}}{2}.$$

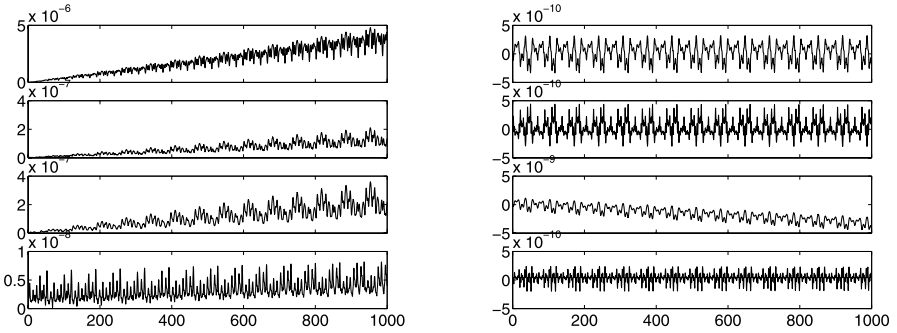

Fig. 3 The global truncation error in y (*left*), and the error in the Hamiltonian (*right*) against time t prior to the Hamiltonian Hopf bifurcation ($\lambda = -0.5$). The integrators included are *gauss4* (*top*), *hochost4* (*second*), *exprk4* (*third*) and *method* (*bottom*)

The Hamiltonian Hopf bifurcation occurs at $\lambda = 0$, where four distinct eigenvalues travel along the imaginary axis and two pairs collide at the Hamiltonian Hopf bifurcation at $\frac{1}{2}i$ and $-\frac{1}{2}i$ respectively. After the bifurcation, the eigenvalues split as a quadruplet into the complex plane.

We integrate this Hamiltonian system using *gauss4*, *hochost4* and *method* with the set of initial conditions $y_0 = (0.1, 0.1, 0.1, 0.141)$ and stepsize $h = 0.1$. In addition, we also include numerical results from the fourth-order, explicit 7-stage exponential Runge–Kutta method proposed by Hochbruck, Lubich and Selhofer in [15]. Although this method is not symplectic, it has a B-series and it is linearization-preserving at every fixed point of an arbitrary system (cf. [21]). We denote this method as *exprk4*.

In the following experiment, we pick three different parameter values λ in the numerical integration to represent the system before ($\lambda = -0.5$), during ($\lambda = 0$), and after ($\lambda = 0.5$) the Hamiltonian Hopf bifurcation. For $\lambda < 0$, the initial condition y_0 produces a quasi-periodic orbit that would eventually fill up the entire torus in the phase space. At the Hamiltonian Hopf bifurcation $\lambda = 0$, the eigenvalues $\mu_1, \ldots, \mu_4$ collide on the imaginary axis of the complex plane to produce a periodic orbit. For $\lambda > 0$ after the bifurcation, the origin changes stability from a centre to a saddle as the eigenvalues split into the complex plane. The solution produced by y_0 now moves rapidly away from the origin.

Figure 3 shows two graphs containing the global truncation error in solution y (left) and the error in the Hamiltonian $H(y)$ (right), from the four integrators for $t = [0, 1000]$. From the graph on the right hand side of Fig. 3, we see that the numerical Hamiltonian values are nearly preserved by the two symplectic integrators *gauss4* and *method*. Also integrator *hocost4* appears to produce numerical Hamiltonian values which are bounded for this stepsize h even though it is not symplectic (see Fig. 2). Furthermore, *exprk4* is clearly non-symplectic, which supports a conjecture in [21] that no exponential Runge–Kutta methods proposed in [15] can simultaneously be symplectic and linearization-preserving. The graph on the left hand side of Fig. 3 shows that *method* commits the smallest global truncation error in the quasi-periodic solution y, and its error accumulation rate is linear over time.

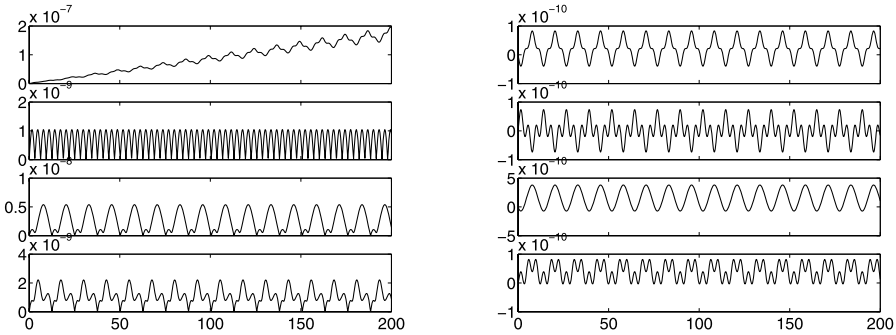

Fig. 4 The global truncation error in y (*left*), and the error in the Hamiltonian (*right*) against time t at the Hamiltonian Hopf bifurcation ($\lambda = 0$). The integrators included are *gauss4* (*top*), *hochost4* (*second*), *exprk4* (*third*) and *method* (*bottom*)

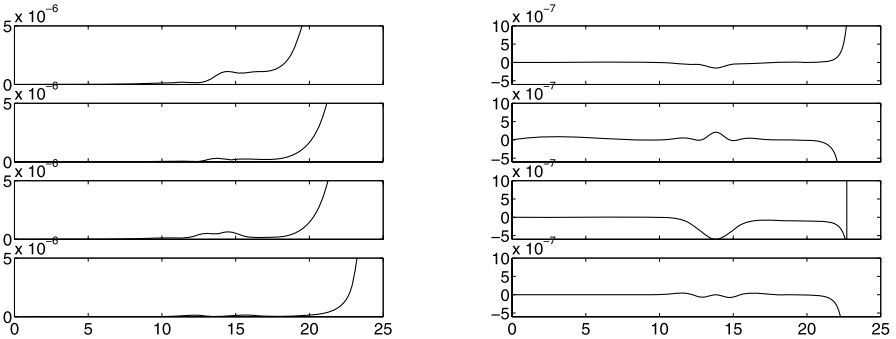

Fig. 5 The global truncation error in y (*left*), and the error in the Hamiltonian (*right*) against time t after the Hamiltonian Hopf bifurcation ($\lambda = 0.5$). The integrators included are *gauss4* (*top*), *hochost4* (*second*), *exprk4* (*third*) and *method* (*bottom*)

At the Hamiltonian Hopf bifurcation when parameter $\lambda = 0$, the left hand graph in Fig. 4 shows that *hochost4* and *method* do remarkably well with mutually comparable global truncation errors in the periodic solution y, which are smaller than those from the non-symplectic *exprk4* method and the symplectic *gauss4* method. Although we have only shown the selected interval $t = [0, 200]$ in Fig. 4, this behaviour persists over the entire integration period $t = [0, 1000]$. The right hand graph of Fig. 4 shows that the Hamiltonian values appear to be bounded against time t for all integrators, regardless of whether they are symplectic or not. This behaviour again persist for the entire integration period.

Finally, Fig. 5 displays the graphs of the global truncation error in y (left) and the error in the Hamiltonian $H(y)$ (right) for the system after the Hamiltonian Hopf bifurcation at $\lambda = 0.5$. The integration period in this case is much shorter because the origin is now of saddle-type, and the solutions close to the origin diverge rapidly away from the origin. In this case, we no longer expect symplectic integrators to produce solutions for which the Hamiltonian values are nearly preserved. Indeed, the right hand graph of Fig. 5 shows the error in the numerical Hamiltonian values

from all four integrators diverge roughly around time $t \approx 20$. With respect to the global truncation error in solution y, the left hand graph of Fig. 5 shows that *method* commits the smallest global truncation error in the comparison before the solution moves in the direction of the unstable manifold, away from the origin.

7 Conclusion

In this article we introduced a new type of geometric integrators that are linearization-preserving. Linearization-preserving integrators form one of the very few known natural classes of integrators defined for all vector fields. For B-series methods which are linearization-preserving, their coefficients for all tall rooted trees must be exact. Such methods may improve the accuracy of numerical solutions over traditional integrators, when the dominant error in the numerical solution is contributed by elementary differentials of the vector field corresponding to tall trees. Examples of such scenarios include the solution in the neighbourhood of a fixed point, the bifurcation of a system associated with eigenvalue resonances, and the approximation of an orbit in which sections of the trajectory is nearly straight. A more complicated example involves systems with heteroclinic cycles. Near there, orbits spend a lot of time near a periodic sequence of fixed points. It is difficult to preserve this type of behaviour accurately with other methods.

In this article, we have established (by construction) the existence of a fourth-order, linearization-preserving, self-adjoint and symplectic B-series method for Hamiltonian systems. Essentially, the implicit midpoint rule is applied to a preprocessed Hamiltonian vector field which is constructed from the original Hamiltonian vector field. A similar technique with the preprocessed vector field can also be used to construct a linearization-preserving B-series method which is self-adjoint.

For the experiments on Hamiltonian ODEs presented in this article, the new method $\Upsilon_f(y)$, which is the first method of its kind, typically committed a smaller global truncation error than those from the other integrators in the comparison. We suspect that the ability of method $\Upsilon_f(y)$ to be linearization-preserving, coupled with symplecticity and/or reversibility, contribute to this favourable error behaviour. However, we suspect that the linearization-preserving property alone does not necessarily lead to reduced global truncation error. This can be seen by comparing the numerical solution from the non-symplectic, 7-stage, exponential Runge–Kutta method *exprk4* proposed by Hochbruck et al. in [15], with the numerical solution from our *method*, implemented for the Hamiltonian system in Sect. 6.3.

It can also be observed that the non-symplectic exponential general linear method *hochost4* in this comparison display the same favourable behaviour in global truncation error as the symplectic *gauss4* and our *method*. As yet, the explanation of such favourable behaviour on Hamiltonian systems by certain exponential general linear methods remains an open problem.

While we are encouraged by the numerical results from our proposed method, one drawback with our method is that its construction depends on the third derivative of the Hamiltonian. Even with the advancement in automatic differentiation and symbolic computing that make numerical methods based on Jacobians and higher-order derivatives more practical, the involvement of the Hamiltonian third-derivative

makes our proposed method unavoidably complicated and expensive to implement. Similarly, although the proposed self-adjoint and linearization-preserving method requires only the second derivative of the Hamiltonian, it is no longer symplectic. Currently we are not aware of a symplectic and linearization-preserving integrator which only depends on the second derivative of the Hamiltonian.

Also, successful implementations of $\Upsilon_f(y)$ depend on whether matrices Φ and $D\Phi$ can be found. Careful consideration must be made before any approximation is introduced, or else desirable features like the preservation of linearization might be destroyed. The resolution of such difficulties currently remains an open problem.

More importantly however, the existence of such a preprocessed vector field integrator gives us hope that by employing similar techniques, new, high-order geometric integrators may be constructed to preserve simultaneously geometric properties not seen from previous constructions of numerical methods. Furthermore, we believe that linearization-preserving integrators have the potential to offer improved numerics for ODEs which feature certain types of dynamics, but issues concerning costs and efficiency must be addressed before their full potential can be realised.

Acknowledgement The authors thank the referees for their comments.

References

1. Arnold, V.I.: Mathematical Methods of Classical Mechanics. Springer, New York (1989)
2. Barnett, S.: Matrices: Methods and Applications. Clarendon Press, Oxford (1990)
3. Berland, H., Skaflestad, B., Wright, W.: EXPINT—A Matlab package for exponential integrators. Tech. Report 2005/4, Norwegian University of Science and Technology, Norway (2005)
4. Budd, C.J., Piggott, M.D.: Geometric integration and its applications. In: Handbook of Numerical Analysis, vol. XI, pp. 35–139. North-Holland, Amsterdam (2003)
5. Butcher, J.C.: Numerical Methods for Ordinary Differential Equations. Wiley, New York (2008)
6. Channel, P.J., Scovel, C.: Symplectic integration of Hamiltonian systems. Nonlinearity **3**, 231–259 (1990)
7. Chartier, P., Faou, E., Murua, A.: An algebraic approach to invariant preserving integrators: the case of quadratic and Hamiltonian invariants. Numer. Math. **103**, 575–590 (2006)
8. Chartier, P., Hairer, E., Vilmart, G.: A substitution law for B-series vector fields. Tech. Report, 2005/5498, INRIA, France (2005)
9. Chartier, P., Hairer, E., Vilmart, G.: Numerical integrators based on modified differential equations. Math. Comput. **76**, 1941–1953 (2007)
10. Dellnitz, M., Melbourne, I., Marsden, J.E.: Generic bifurcation of Hamiltonian vector fields with symmetry. Nonlinearity **5**, 979–996 (1992)
11. Feng, K.: Difference schemes for Hamiltonian formalism and symplectic geometry. J. Comput. Math. **4**, 279–289 (1986)
12. Hairer, E., Lubich, C., Wanner, G.: Geometric Numerical Integration. Structure-Preserving Algorithms for Ordinary Differential Equations. Springer, Berlin (2006)
13. Hairer, E., Wanner, G.: Solving Ordinary Differential Equations. II Stiff and Differential-Algebraic Problems. Springer, Berlin (1996)
14. Hénon, M., Heiles, C.: The applicability of the third integral of motion: Some numerical experiments. Astron. J. **69**, 73–79 (1964)
15. Hochbruck, M., Lubich, C., Selhofer, H.: Exponential integrators for large systems of differential equations. SIAM J. Sci. Comput. **19**, 1552–1574 (1998)
16. Leimkuhler, B., Reich, S.: Simulating Hamiltonian Dynamics. Cambridge University Press, Cambridge (2004)
17. McLachlan, R.I., Quispel, G.R.W.: Geometric integrators for ODEs. J. Phys. A **39**, 5251–5285 (2006)
18. McLaren, D.I., Quispel, G.R.W., Wright, W.M.: Private communication

19. Minchev, B.V., Wright, W.M.: A review of exponential integrators for first order semi-linear problems. Tech. Report, 2005/4, The Norwegian University of Science and Technology, Norway (2005)
20. Sanz-Serna, J.M., Calvo, M.P.: Numerical Hamiltonian Problems. Chapman & Hall, London (1994)
21. Tse, P.S.P.: Geometric Numerical Integration: On the Numerical Preservation of Multiple Geometric Properties for Ordinary Differential Equations. Ph.D. Thesis, Department of Mathematics, La Trobe University, Australia (2007)
22. van der Meer, J.C.: The Hamiltonian Hopf Bifurcation. Springer, Berlin (1985)