

Why geometric numerical integration?

A. Iserles & G.R.W. Quispel

December 17, 2015

1 The purpose of GNI

Geometric numerical integration (GNI) emerged as a major thread in numerical mathematics some 25 years ago. Although it has had antecedents, in particular the concerted effort of the late Feng Kang and his group in Beijing to design structure-preserving methods, the importance of GNI has been recognised and its scope delineated only in the 1990s.

But we are racing ahead of ourselves. At the beginning, like always in mathematics, there is the definition and the rationale of GNI. The rationale is that all-too-often mathematicians concerned with differential equations split into three groups that have little in common. Firstly, there are the applied mathematicians, the model builders, who formulate differential equations to describe physical reality. Secondly, there are those pure mathematicians investigating differential equations and unravelling their qualitative features. Finally, the numerical analysts who flesh out the numbers and the graphics on the bones of mathematical formulation. Such groups tended to operate in mostly separate spheres and, in particular, this has been true with regards to computation. Discretisation methods were designed (with huge creativity and insight) to produce rapidly and robustly numerical solutions that can be relied to carry overall small error. Yet, such methods have often carried no guarantee whatsoever to respect qualitative features of the underlying system, the very same features that had been obtained with such effort by pure and applied mathematicians.

Qualitative features come basically in two flavours, the *dynamical* and the *geometric*. Dynamical features – sensitivity with respect to initial conditions and other parameters, as well as the asymptotic behaviour – have been recognised as important by numerical analysts for a long time, not least because they tend to impinge directly on accuracy. Thus, sensitivity with respect to initial conditions and perturbations comes under ‘conditioning’ and the recovery of correct asymptotics under ‘stability’, both subject to many decades of successful enquiry. Geometric attributes are invariants, constants of the flow. They are often formulated in the language of differential geometry (hence the name!) and mostly come in three varieties: *conservation laws*, e.g. Hamiltonian energy or angular momentum, which geometrically mean that the solution, rather than evolving in some large space \mathbb{R}^d , is restricted to a lower-dimensional manifold \mathcal{M} , *Lie point symmetries*, e.g. scaling invariance, which restrict the solution to the tangent bundle of some manifold, and quantities like *symplecticity* and *volume*, whose conservation corresponds to an evolution on the cotangent bundle of a man-

ifold. *The design and implementation of numerical methods that respect geometric invariants is the business of GNI.*

Since its emergence, GNI has become the new paradigm in numerical solution of ODEs, while making significant inroads into numerical PDEs. As often, yesterday's revolutionaries became the new establishment. This is an excellent moment to pause and take stock. Have all the major challenges been achieved, all peaks scaled, leaving just a tidying-up operation? Is there still any point to GNI as a separate activity or should it be considered as a victim of its own success and its practitioners depart to fields anew – including new areas of activity that have been fostered or enabled by GNI?

These are difficult questions and we claim no special authority to answer them in an emphatic fashion. Yet, these are questions which, we believe, must be addressed. This short article is an attempt to foster a discussion. We commence with a brief survey of the main themes of GNI *circa* 2015. This is followed by a review of recent and ongoing developments, as well as of some new research directions that have emerged from GNI but have acquired a life of their own.

2 The story so far

2.1 Symplectic integration

The early story of GNI is mostly the story of symplectic methods. A Hamiltonian system

$$\dot{\mathbf{p}} = -\frac{\partial H(\mathbf{p}, \mathbf{q})}{\partial \mathbf{q}}, \quad \dot{\mathbf{q}} = \frac{\partial H(\mathbf{p}, \mathbf{q})}{\partial \mathbf{p}}, \quad (2.1)$$

where $H : \mathbb{R}^{2d} \rightarrow \mathbb{R}$ is a *Hamiltonian energy*, plays a fundamental role in mechanics and is known to possess a long list of structural invariants, e.g. the conservation of the Hamiltonian energy. Yet, arguably its most important feature is the conservation of the *symplectic form* $\sum_{k=1}^d d\mathbf{p}_k \wedge d\mathbf{q}_k$ because symplecticity is equivalent to Hamiltonicity – in other words, every solution of a Hamiltonian system is a symplectic flow and every symplectic flow is Hamiltonian with respect to an appropriate Hamiltonian energy (Hairer, Lubich & Wanner 2006).

The solution of Hamiltonian problems using symplectic methods has a long history, beautifully reviewed in (Hairer, Lubich & Wanner 2003), but modern efforts can be traced to the work of Feng and his collaborators at the Chinese Academy of Sciences, who have used generating-function methods to solve Hamiltonian systems (Feng, Wu, Qin & Wang 1989). And then, virtually simultaneously, Lasagni (1988), Sanz-Serna (1988) and Suris (1988) proved that certain Runge–Kutta methods, including the well-known Gauss–Legendre methods, preserve symplecticity and they presented an easy criterion for the symplecticity of Runge–Kutta methods. GNI came of age!

Subsequent research into symplectic Runge–Kutta methods had branched out into a number of directions, each with its own important ramifications outside the Hamiltonian world:

- *Backward error analysis.* The idea of backward error analysis (AEB) can be traced to Wilkinson's research into linear algebra algorithms in the 1950ties.

Instead of asking “what is the numerical error for our problem”, Wilkinson asked “which nearby problem is solved *exactly* by our method?”. The difference between the original and the nearby problem can tell us a great deal about the nature of the error in a numerical algorithm.

A generalisation of AEB to the field of differential equations is fraught with difficulties. Perhaps the first successful attempt to analyse Hamiltonian ODEs in this setting was by Neishtadt (1984) and it was followed by many, too numerous to list: an excellent exposition (like for many things GNI) is the monograph of Hairer et al. (2006). The main technical tool is the B-series, an expansion of composite functions in terms of forests of rooted trees, originally pioneered by Butcher (1963). (We mention in passing that the Hopf algebra structure of this *Butcher group* has been recently exploited by mathematical physicists to understand the renormalisation group (Connes & Kreimer 1999) – as the authors write, “We regard Butcher’s work on the classification of numerical integration methods as an impressive example that concrete problem-oriented work can lead to far-reaching conceptual results”.) It is possible to prove that, subject to very generous conditions, the solution of a Hamiltonian problem by a symplectic method, implemented with constant step size, is exponentially near to the *exact* solution of a nearby Hamiltonian problem for an exponentially long time. This leads to considerably greater numerical precision, as well as to the conservation on average (in a strict ergodic sense) of Hamiltonian energy.

B-series fall short in a highly oscillatory and multiscale setting, encountered frequently in practical Hamiltonian systems. The alternative in the AEB context is an expansion into *modulated Fourier series* (Hairer & Lubich 2000).

- *Composition and splitting.*

Many Hamiltonians of interest can be partitioned into a sum of kinetic and potential energy, $H(\mathbf{p}, \mathbf{q}) = \mathbf{p}^\top M \mathbf{p} + V(\mathbf{q})$. It is often useful to take advantage of this in the design of symplectic methods. While conventional symplectic Runge–Kutta methods are implicit, hence expensive, *partitioned Runge–Kutta methods*, advancing separately in the ‘direction’ of kinetic and potential energy, can be explicit and are in general much cheaper. While perhaps the most important method, the Störmer–Verlet scheme (Hairer et al. 2003), has been known for many years, modern theory has led to an entire menagerie of composite and partitioned methods (Sanz-Serna & Calvo 1994).

Splitting methods¹ have been used in the numerical solution of PDEs since 1950s. Thus, given the equation $u_t = \mathcal{L}_1(u) + \mathcal{L}_2(u)$, where the \mathcal{L}_k s are (perhaps nonlinear) operators, the idea is to approximate the solution in the form

$$u(t+h) \approx e^{\alpha_1 h \mathcal{L}_1} e^{\beta_1 h \mathcal{L}_2} e^{\alpha_2 h \mathcal{L}_1} \dots e^{\alpha_s h \mathcal{L}_1} e^{\beta_s h \mathcal{L}_2} u(t), \quad (2.2)$$

where $v(t_0+h) =: e^{h \mathcal{L}_1} v(t_0)$ and $w(t_0+h) =: e^{h \mathcal{L}_2} w(t_0)$ are, formally, the solutions of $\dot{v} = \mathcal{L}_1(v)$ and $\dot{w} = \mathcal{L}_2(w)$ respectively, with suitable boundary conditions. The underlying assumption is that the solutions of the latter two

¹Occasionally known in the PDE literature as *alternate direction methods*.

equations are either available explicitly or are easy to approximate, while the original equation is more difficult.

A pride of place belongs to *palindromic compositions* of the form

$$e^{\alpha_1 h \mathcal{L}_1} e^{\beta_1 h \mathcal{L}_2} e^{\alpha_2 h \mathcal{L}_1} \dots e^{\alpha_q h \mathcal{L}_1} e^{\beta_q h \mathcal{L}_2} e^{\alpha_q h \mathcal{L}_1} \dots e^{\alpha_2 h \mathcal{L}_1} e^{\beta_1 h \mathcal{L}_2} e^{\alpha_1 h \mathcal{L}_1}, \quad (2.3)$$

invariant with respect to a reversal of the terms. They constitute a *time-symmetric map*, and this has a number of auspicious consequences. Firstly, they are always of an even order. Secondly – and this is crucial in the GNI context – they respect both structural invariants whose integrators are closed under composition, i.e. form a group (for example integrators preserving volume, symmetries, or first integrals), as well as invariants whose integrators are closed under symmetric composition, i.e. form a symmetric space (for example integrators that are self-adjoint, or preserve reversing symmetries). A basic example of (2.3) is the second-order *Strang composition*

$$e^{\frac{1}{2} h \mathcal{L}_1} e^{h \mathcal{L}_2} e^{\frac{1}{2} h \mathcal{L}_1} = e^{h(\mathcal{L}_1 + \mathcal{L}_2)} + \mathcal{O}(h^3).$$

Its order – and, for that matter, the order of any time-symmetric method – can be boosted by the *Yoshida device* (Yoshida 1990). Let Φ be a time-symmetric approximation to $e^{t\mathcal{L}}$ of order $2P$, say. Then

$$\Phi((1 + \alpha)h)\Phi(-(1 + 2\alpha)h)\Phi((1 + \alpha)h), \quad \text{where} \quad \alpha = \frac{2^{1/(2P+1)} - 1}{2 - 2^{1/(2P+1)}}$$

is also time symmetric and of order $2P+2$. Successive applications of the Yoshida device allow to increase arbitrarily the order of the Strang composition, while retaining its structure-preserving features. This is but a single example of the huge world of splitting and composition methods, reviewed in (McLachlan & Quispel 2002).

- *Exponential integrators.*

Many ‘difficult’ ODEs can be written in the form $\dot{\mathbf{y}} = A\mathbf{y} + \mathbf{b}(\mathbf{y})$ where the matrix A is ‘larger’ (in some sense) than $\mathbf{b}(\mathbf{y})$ – for example, A may be the Jacobian of an ODE (which may vary from step to step). Thus, it is to be expected that the ‘nastiness’ of the ODE under scrutiny – be it stiffness, Hamiltonicity or high oscillation – is somehow ‘hardwired’ into the matrix A . The exact solution of the ODE can be written in terms of the variation-of-constants formula,

$$\mathbf{y}(t + h) = e^{hA}\mathbf{y}(t) + \int_0^h e^{(h-\xi)A}\mathbf{b}(\mathbf{y}(t + \xi)) d\xi, \quad (2.4)$$

except that, of course, the right-hand side includes the unknown function \mathbf{y} . Given the availability of very effective methods to compute the matrix exponential, we can exploit this to construct *exponential integrators*, explicit methods that often exhibit favourable stability and structure-preservation features. The simplest example, the *exponential Euler* method, freezes \mathbf{y} within the integral in (2.4) at its known value at t , the outcome being the first-order method

$$\mathbf{y}_{n+1} = e^{hA}\mathbf{y}_n + A^{-1}(e^{hA} - I)\mathbf{b}(\mathbf{y}_n).$$

The order can be boosted by observing that (in a loose sense which can be made much more precise) the integral above is discretised by the Euler method, which is a one-stage explicit Runge–Kutta scheme, discretising it instead by multistage schemes of this kind leads to higher-order methods (Hochbruck & Ostermann 2010).

Many Hamiltonian systems of interest can be formulated as second-order systems of the form $\ddot{\mathbf{y}} + \Omega^2 \mathbf{y} = \mathbf{g}(\mathbf{y})$. Such systems feature prominently in the case of highly oscillatory mechanical systems, where Ω is positive definite and has some large eigenvalues. Variation of constants (2.4) now reads

$$\begin{aligned} \begin{bmatrix} \mathbf{y}(t+h) \\ \dot{\mathbf{y}}(t+h) \end{bmatrix} &= \begin{bmatrix} \cos(h\Omega) & \Omega^{-1} \sin(h\Omega) \\ -\Omega \sin(h\Omega) & \cos(h\Omega) \end{bmatrix} \begin{bmatrix} \mathbf{y}(t) \\ \dot{\mathbf{y}}(t) \end{bmatrix} \\ &+ \int_t^{t+h} \begin{bmatrix} \cos((h-\xi)\Omega) & \Omega^{-1} \sin((h-\xi)\Omega) \\ -\Omega \sin((h-\xi)\Omega) & \cos((h-\xi)\Omega) \end{bmatrix} \begin{bmatrix} \mathbf{0} \\ \mathbf{g}(\mathbf{y}(t+\xi)) \end{bmatrix} d\xi \end{aligned}$$

and we can use either standard exponential integrators or exponential integrators designed directly for second-order systems and using Runge–Kutta–Nyström methods on the nonlinear part (Wu, You & Wang 2013).

An important family of exponential integrators for second-order systems are *Gautschi-type methods*

$$\mathbf{y}_{n+1} - 2\mathbf{y}_n + \mathbf{y}_{n-1} = h^2 \Psi(h\Omega)(\mathbf{g}_n - \Omega^2 \mathbf{y}_n), \quad (2.5)$$

which are of second order. Here $\Psi(x) = 2(1 - \cos x)/x$ while, in Gautschi’s original method, $\mathbf{g}_n = \mathbf{g}(\mathbf{y}_n)$ (Hochbruck & Ostermann 2010). Unfortunately, this choice results in resonances and a better one is $\mathbf{g}_n = \mathbf{g}(\Phi(h\Omega)\mathbf{y}_n)$, where the *filter* Φ eliminates resonances: $\Phi(0) = I$ and $\Phi(k\pi) = 0$ for $k \in \mathbb{N}$. We refer to (Hochbruck & Ostermann 2010) for further discussion of such methods in the context of symplectic integration.

- *Variational integrators. Lagrangian formulation* recasts a large number of differential equations as minima of nonlinear functionals. Thus, for example, instead of the Hamiltonian problem $M\ddot{\mathbf{q}} + \nabla V(\mathbf{q}) = \mathbf{0}$, where the matrix M is positive definite, we may consider the equivalent variational formulation of minimizing the positive-definite nonlinear functional $L(\mathbf{q}, \dot{\mathbf{q}}) = \frac{1}{2}\dot{\mathbf{q}}^\top M \dot{\mathbf{q}} - V(\mathbf{q})$. With greater generality, Hamiltonian and Lagrangian formulations are connected via the familiar Euler–Lagrange equations and, given the functional L , the corresponding second-order system is

$$\frac{\partial L(\mathbf{q}, \dot{\mathbf{q}})}{\partial \mathbf{q}} - \frac{d}{dt} \left[\frac{\partial L(\mathbf{q}, \dot{\mathbf{q}})}{\partial \dot{\mathbf{q}}} \right] = \mathbf{0}.$$

The rationale of variational integrators parallels that of the *Ritz method* in the theory of finite elements. We first reformulate the Hamiltonian problem as a Lagrangian one, project it to a finite-dimensional space, solve it there and transform back. The original symplectic structure is replaced by a finite-dimensional symplectic structure, hence the approach is by design symplectic (Marsden & West 2001).

2.2 Lie-group methods

Let \mathcal{G} be a Lie group and \mathcal{M} a differentiable manifold. We say that $\Lambda : \mathcal{G} \times \mathcal{M} \rightarrow \mathcal{M}$ is a *group action* if

- a. $\Lambda(\iota, y) = y$ for all $y \in \mathcal{M}$ (where ι is the identity of \mathcal{G}) and
- b. $\Lambda(p, \Lambda(q, y)) = \Lambda(p \cdot q, y)$ for all $p, q \in \mathcal{G}$ and $y \in \mathcal{M}$.

If, in addition, for every $x, y \in \mathcal{M}$ there exists $p \in \mathcal{G}$ such that $y = \Lambda(p, x)$, the action is said to be transitive and \mathcal{M} is a *homogeneous space*, acted upon by \mathcal{G} .

Every Lie group acts upon itself, while the orthogonal group $O(n)$ acts on the $(n-1)$ -sphere by multiplication, $\Lambda(p, y) = py$. The orthogonal group also acts on the *isospectral manifold* of all symmetric matrices similar to a specific symmetric matrix by similarity, $\Lambda(p, y) = pyp^\top$. Given $1 \leq m \leq n$, the *Grassmann manifold* $\mathbb{G}(n, m)$ of all m -dimensional subspaces of \mathbb{R}^n is a homogeneous space acted upon by $SO(m) \times SO(n-m)$, where $SO(m)$ is the special orthogonal group – more precisely, $\mathbb{G}(n, m) = SO(n)/(SO(m) \times SO(n-m))$.

Faced with a differential equation evolving in a homogeneous space, we can identify its flow with a group action: Given an initial condition $y_0 \in \mathcal{M}$, instead of asking “what is the value of y at time $t > 0$ ” we might pose the equivalent question “what is the group action that takes the solution from y_0 to $y(t)$?”. This is often a considerably more helpful formulation because a group action can be further related to an *algebra action*. Let \mathfrak{g} be the Lie algebra corresponding to the matrix group \mathcal{G} , i.e. the tangent space at $\iota \in \mathcal{G}$, and denote by $\mathfrak{X}(\mathcal{M})$ the set of all Lipschitz vector fields over \mathcal{M} . Let $\lambda : \mathfrak{g} \rightarrow \mathfrak{X}(\mathcal{M})$ and $a : \mathbb{R}_+ \times \mathcal{M} \rightarrow \mathfrak{g}$ be both Lipschitz. In particular, we might consider

$$\lambda(a, y) = \left. \frac{d}{ds} \Lambda(\rho(s, y), y) \right|_{s=0},$$

where Λ is a group action and $\rho : \mathbb{R}_+ \rightarrow \mathcal{G}$, $\rho(s, y(s)) = \iota + a(s, y(s))s + \mathcal{O}(s^2)$ for small $|s|$. The equation $\dot{y} = \lambda(a(t, y), y)$, $y(0) = y_0 \in \mathcal{M}$ represents *algebra action* and its solution evolves in \mathcal{M} . Moreover,

$$y(t) = \Lambda(v(t), y_0) \quad \text{where} \quad \dot{v} = a(t, \Lambda(v, y_0))v, \quad v(0) = \iota \in \mathcal{G} \quad (2.6)$$

is a *Lie-group equation*. Instead of solving the original ODE on \mathcal{M} , it is possible to solve (2.6) and use the group action Λ to advance the solution to the next step: this is the organising principle of most *Lie-group methods* (Iserles, Munthe-Kaas, Nørsett & Zanna 2000). It works because a Lie-group equation can be solved in the underlying Lie algebra, which is a *linear space*. Consider an ODE² $\dot{y} = f(y)$, $y(0) \in \mathcal{M}$, such that $f : \mathcal{M} \rightarrow \mathfrak{X}$ – the solution $y(t)$ evolves on the manifold. While conventional numerical methods are highly unlikely to stay in \mathcal{M} , this is not the case for Lie-group methods. We can travel safely between \mathcal{M} and \mathcal{G} using a group action. The traffic between \mathcal{G} and \mathfrak{g} is slightly more complicated and we need to define a *trivialisation*, i.e. an invertible map taking smoothly a neighbourhood of $0 \in \mathfrak{g}$ to a neighbourhood of $\iota \in \mathcal{G}$ and taking zero to identity. The most ubiquitous example of trivialisation is the exponential map, which represents the solution of (2.6) as $v(t) = e^{\omega(t)}$, where ω is

²Or, for that matter, a PDE, except that formalities are somewhat more complicated.

the solution of the *dexpinv* equation

$$\dot{\omega} = \sum_{m=0}^{\infty} \frac{B_m}{m!} \text{ad}_{a(t, e^\omega)}^m \omega, \quad \omega(0) = 0 \in \mathfrak{g} \quad (2.7)$$

(Iserles et al. 2000). Here the B_m s are Bernoulli numbers, while ad_b^m is the *adjoint operator* in \mathfrak{g} ,

$$\text{ad}_b^0 c = c, \quad \text{ad}_b^m c = [b, \text{ad}_b^{m-1} c], \quad m \in \mathbb{N}, \quad b, c \in \mathfrak{g}.$$

Because \mathfrak{g} is closed under linear operations and commutation, solving (2.7) while respecting Lie-algebraic structure is straightforward. Mapping back, first to \mathcal{G} and finally to \mathcal{M} , we keep the numerical solution of $\dot{y} = f(t)$ on the manifold.

Particularly effective is the use of explicit Runge–Kutta methods for (2.7), the so-called Runge–Kutta–Munthe-Kaas (RKMK) methods (Munthe-Kaas 1998). To help us distinguish between conventional Runge–Kutta methods and RKMK, consider the three-stage, third-order method with the Butcher tableau³

$$\begin{array}{c|ccc} 0 & & & \\ \frac{1}{2} & \frac{1}{2} & & \\ 1 & -1 & 2 & \\ \hline & \frac{1}{6} & \frac{2}{3} & \frac{1}{6} \end{array}. \quad (2.8)$$

Applied to the ODE $\dot{y} = f(t, y)$, $y(t_n) = y_n \in \mathcal{M}$, evolving on the manifold $\mathcal{M} \subset \mathbb{R}^d$, it becomes

$$\begin{aligned} k_1 &= f(t_n, y_n), \\ k_2 &= f(t_{n+\frac{1}{2}}, y_n + \frac{1}{2}hk_1), \\ k_3 &= f(t_{n+1}, y_n - hk_1 + 2hk_2), \\ \Delta &= h(\frac{1}{6}k_1 + \frac{2}{3}k_2 + \frac{1}{6}k_3), \\ y_{n+1} &= y_n + \Delta. \end{aligned}$$

Since we operate in \mathbb{R}^d , there is absolutely no reason for y_{n+1} to live in \mathcal{M} . However, once we implement (2.8) at an algebra level (truncating first the *dexpinv* equation (2.7)),

$$\begin{aligned} k_1 &= a(t_n, \iota), \\ k_2 &= a(t_{n+\frac{1}{2}}, e^{hk_1/2}), \\ k_3 &= a(t_{n+1}, e^{-hk_1+2hk_2}), \\ \Delta &= h(\frac{1}{6}k_1 + \frac{2}{3}k_2 + \frac{1}{6}k_3), \\ \omega_{n+1} &= \Delta + \frac{1}{6}h[\Delta, k_1] \\ y_{n+1} &= \Lambda(e^{\omega_{n+1}}, y_n), \end{aligned}$$

³For traditional concepts such as Butcher tableaux, Runge-Kutta methods and B-series, the reader is referred to (Hairer, Nørsett & Wanner 1993).

the solution is guaranteed to stay in \mathcal{M} .

An important special case of a Lie-group equation is the linear ODE $\dot{v} = a(t)v$, where $a : \mathbb{R}_+ \rightarrow \mathfrak{g}$. Although RKMK works perfectly well in a linear case, special methods do even better. Perhaps the most important is the *Magnus expansion* (Magnus 1954), $v(t) = e^{\omega(t)}v(0)$, where

$$\begin{aligned} \omega(t) &= \int_0^t a(\xi) d\xi - \frac{1}{2} \int_0^t \int_0^{\xi_1} [a(\xi_2), a(\xi_1)] d\xi_2 d\xi_1 \\ &\quad + \frac{1}{4} \int_0^t \int_0^{\xi_1} \int_0^{\xi_2} [[a(\xi_3), a(\xi_2)], a(\xi_1)] d\xi_3 d\xi_2 d\xi_1 \\ &\quad + \frac{1}{12} \int_0^t \int_0^{\xi_1} \int_0^{\xi_2} [a(\xi_3), [a(\xi_2), a(\xi_1)]] d\xi_3 d\xi_2 d\xi_1 + \dots \end{aligned} \quad (2.9)$$

We refer to (Iserles & Nørsett 1999, Iserles et al. 2000, Blanes, Casas, Oteo & Ros 2009) for explicit means to derive expansion terms, efficient computation of multivariate integrals that arise in this context and many other implementation details. Magnus expansions are important in a number of settings when preservation of structure is not an issue, not least in the solution of linear stochastic ODEs (Lord, Malham & Wiese 2008).

There are alternative means to expand the solution of (2.7) in a linear case, not least the *Fer expansion*, that has found recently an important application in the computation of Sturm–Liouville spectra (Ramos & Iserles 2015).

Another approach to Lie-group equations uses *canonical coordinates of the second kind* (Owren & Marthinsen 2001).

2.3 Conservation of volume

An ODE $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$ is divergence-free if $\nabla \cdot \mathbf{f}(\mathbf{x}) = 0$. The flows of divergence-free ODEs are volume-preserving (VP). Volume is important to preserve, as it leads to KAM-tori, incompressibility, and, most importantly, is a crucial ingredient for ergodicity. Unlike symplecticity, however, phase space volume can generically *not* be preserved by Runge–Kutta methods, or even by their generalisations, B-series methods. This was proved independently in (Chartier & Murua 2007) and in (Iserles, Quispel & Tse 2007). Since B-series methods cannot preserve volume, we need to look to other methods.

There are essentially two known numerical integration methods that preserve phase space volume. The first volume-preserving method is based on splitting (Feng & Shang 1995). As an example, consider a 3D volume preserving vector field:

$$\begin{aligned} \dot{x} &= u(x, y, z) \\ \dot{y} &= v(x, y, z) \\ \dot{z} &= w(x, y, z) \end{aligned} \quad (2.10)$$

with

$$u_x + v_y + w_z = 0.$$

We split this 3D VP vector field into two 2D VP vector fields as follows

$$\begin{aligned}
\dot{x} &= u(x, y, z), & \dot{x} &= 0, \\
\dot{y} &= - \int u_x(x, y, z) dy, & \dot{y} &= v(x, y, z) + \int u_x(x, y, z) dy, \\
\dot{z} &= 0; & \dot{z} &= w(x, y, z).
\end{aligned} \tag{2.11}$$

The vector field on the left is divergence-free by construction, and since both vector fields add up to (2.1), it follows that the vector field on the right is also volume-preserving.

Having split the original vector field into 2D VP vector fields, we need to find VP integrators for each of these 2D VP vector fields. But that is easy, since 2D VP vector fields are essentially equivalent to 2D Hamiltonian vector fields (with the extra dimension ‘frozen’), and all symplectic methods (e.g. symplectic Runge–Kutta methods) are volume-preserving for Hamiltonian vector fields.

The above splitting method is easily generalised to n dimensions, where one splits into $n - 1$ 2D VP vector fields, and integrates each using a symplectic Runge–Kutta method.

An alternative VP integration method was discovered independently by Shang and by Quispel (Shang 1994, Quispel 1995). We again illustrate this method in 3D.

We will look for an integrator of the form

$$\begin{aligned}
x_1 &= g_1(x'_1, x_2, x_3) \\
x'_2 &= g_2(x'_1, x_2, x_3) \\
x'_3 &= g_3(x'_1, x'_2, x_3)
\end{aligned} \tag{2.12}$$

where (here and below) $x_i = x_i(nh)$, and $x'_i = x_i((n + 1)h)$. The reason the form (2.12) is convenient, is because any such map is VP iff

$$\frac{\partial x_1}{\partial x'_1} = \frac{\partial x'_2}{\partial x_2} \frac{\partial x'_3}{\partial x_3}. \tag{2.13}$$

To see how to construct a VP integrator of the form (2.12), consider as an example the ODE

$$\begin{aligned}
\dot{x}_1 &= x_2 + x_1^2 + x_3^3 \\
\dot{x}_2 &= x_3 + x_1 x_2 + x_1^4 \\
\dot{x}_3 &= x_1 - 3x_1 x_3 + x_2^5
\end{aligned} \tag{2.14}$$

It is easy to check that it is divergence-free.

Now consistency requires that any integrator for (2.14) should satisfy

$$\begin{aligned}
x'_1 &= x_1 + h(x_2 + x_1^2 + x_3^3) + \mathcal{O}(h^2) \\
x'_2 &= x_2 + h(x_3 + x_1 x_2 + x_1^4) + \mathcal{O}(h^2) \\
x'_3 &= x_3 + h(x_1 - 3x_1 x_3 + x_2^5) + \mathcal{O}(h^2)
\end{aligned} \tag{2.15}$$

and therefore

$$x_1 = x'_1 - h(x_2 + x_1'^2 + x_3^3) + \mathcal{O}(h^2) \quad (2.16)$$

$$x'_2 = x_2 + h(x_3 + x'_1 x_2 + x_1'^4) + \mathcal{O}(h^2) \quad (2.17)$$

$$x'_3 = x_3 + h(x'_1 - 3x'_1 x_3 + x_2'^5) + \mathcal{O}(h^2) \quad (2.18)$$

Since we are free to choose any consistent g_2 and g_3 in (2.12), provided g_1 satisfies (2.13), we choose the terms designated by $\mathcal{O}(h^2)$ in (2.15) and (2.16) to be identically zero. Equation (2.13) then yields

$$\frac{\partial x_1}{\partial x'_1} = (1 + hx'_1)(1 - 3hx'_1). \quad (2.19)$$

This can easily be integrated to give

$$x_1 = x'_1 - hx_1'^2 - h^2 x_1'^3 + k(x_2, x_3; h). \quad (2.20)$$

where the function k denotes an integration constant that we can choose appropriately. The simplest VP integrator satisfying both (2.14) and (2.20) is therefore:

$$\begin{aligned} x_1 &= x'_1 - h(x_2 + x_1'^2 + x_3^3) - h^2 x_1'^3 \\ x'_2 &= x_2 + h(x_3 + x'_1 x_2 + x_1'^4) \\ x'_3 &= x_3 + h(x'_1 - 3x'_1 x_3 + x_2'^5) \end{aligned} \quad (2.21)$$

A nice aspect of the integrator (2.21) (and (2.12)) is that it is essentially only implicit in one variable. Once x'_1 is computed from the first (implicit) equation, the other two equations are essentially explicit.

Of course the method just described also generalises to any divergence-free ODE in any dimension.

2.4 Preserving energy and other first integrals

As mentioned, Hamiltonian systems exhibit two important geometric properties simultaneously, they conserve both the symplectic form and the energy. A famous no-go theorem by Ge & Marsden (1988) has shown that it is generically impossible to construct a geometric integrator that preserves both properties at once. One therefore must choose which one of these two to preserve in any given application. Particularly in low dimensions and if the energy surface is compact, there are often advantages in preserving the energy.

An energy-preserving B-series method was discovered in (Quispel & McLaren 2008) cf. also (McLachlan, Quispel & Robidoux 1999).

For any ODE $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$, this so-called average vector field method is given by

$$\frac{\mathbf{x}' - \mathbf{x}}{h} = \int_0^1 \mathbf{f}(\xi \mathbf{x}' + (1 - \xi)\mathbf{x}) d\xi. \quad (2.22)$$

If the vector field \mathbf{f} is Hamiltonian, i.e. if there exists a Hamiltonian function $H(\mathbf{x})$ and a constant skew-symmetric matrix S such that $\mathbf{f}(\mathbf{x}) = S\nabla H(\mathbf{x})$, then it follows from (2.22) that energy is preserved, i.e. $H(\mathbf{x}') = H(\mathbf{x})$.

While the B-series method (2.22) is energy-preserving for any Hamiltonian H , it can be shown that no Runge–Kutta method is energy-preserving for all H . For a given *polynomial* H however, Runge–Kutta methods preserving that H do exist (Iavernaro & Trigiante 2009). This can be seen as follows.

Note that the integral in (2.22) is one-dimensional. This means that e.g. for cubic vector fields (and hence for quartic Hamiltonians) an equivalent method is obtained by replacing the integral in (2.22) using Simpson’s rule:

$$\int_0^1 g(\xi) \, d\xi \approx \frac{1}{6} [g(0) + 4g(\frac{1}{2}) + g(1)]. \quad (2.23)$$

yielding the Runge–Kutta method

$$\frac{\mathbf{x}' - \mathbf{x}}{h} = \frac{1}{6} \left[\mathbf{f}(\mathbf{x}) + 4\mathbf{f}\left(\frac{\mathbf{x} + \mathbf{x}'}{2}\right) + \mathbf{f}(\mathbf{x}') \right], \quad (2.24)$$

preserving all quartic Hamiltonians.

We note that (2.22) has second order accuracy. Higher order generalisations have been given in (Hairer 2010). We note that the average vector field method has also been applied to a slew of semi-discretised PDEs in (Celledoni, Grimm, McLachlan, McLaren, O’Neale, Owren & Quispel 2012).

While energy is one of the most important constants of the motion in applications, many other types of first integrals do occur. We note here that all B-series methods preserve all linear first integrals, and that all symplectic B-series methods preserve all quadratic first integrals. So, for example, the implicit midpoint rule

$$\frac{\mathbf{x}' - \mathbf{x}}{h} = \mathbf{f}\left(\frac{\mathbf{x} + \mathbf{x}'}{2}\right)$$

(which is symplectic) preserves all linear and quadratic first integrals. There are however many cases not covered by any of the above.

How does one preserve a cubic first integral that is not energy? And what about Hamiltonian systems whose symplectic structure is not constant? It turns out that generically, any ODE $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$ that preserves an integral $I(\mathbf{x})$, can be written in the form

$$\dot{\mathbf{x}} = S(\mathbf{x})\nabla I(\mathbf{x}), \quad (2.25)$$

where $S(\mathbf{x})$ is a skew-symmetric matrix⁴.

An integral-preserving discretisation of (2.25) is given by

$$\frac{\mathbf{x}' - \mathbf{x}}{h} = \bar{S}(\mathbf{x}, \mathbf{x}')\bar{\nabla}I(\mathbf{x}, \mathbf{x}'), \quad (2.26)$$

where $\bar{S}(\mathbf{x}, \mathbf{x}')$ is any consistent approximation to $S(\mathbf{x})$ (e.g. $\bar{S}(\mathbf{x}, \mathbf{x}') = S(\mathbf{x})$), and the *discrete gradient* $\bar{\nabla}I$ is defined by

$$(\mathbf{x}' - \mathbf{x}) \cdot \bar{\nabla}I(\mathbf{x}, \mathbf{x}') = I(\mathbf{x}') - I(\mathbf{x}) \quad (2.27)$$

⁴Note that in general $S(\mathbf{x})$ need not satisfy the so-called Jacobi identity.

and

$$\lim_{\mathbf{x}' \rightarrow \mathbf{x}} \bar{\nabla} I(\mathbf{x}, \mathbf{x}') = \nabla I(\mathbf{x}). \quad (2.28)$$

There are many different discrete gradients that satisfy (2.27) and (2.28). A particularly simple one is given by the Itoh–Abe discrete gradient, which for example in 3D reads

$$\bar{\nabla} I(\mathbf{x}, \mathbf{x}') = \begin{bmatrix} \frac{I(x'_1, x_2, x_3) - I(x_1, x_2, x_3)}{x'_1 - x_1} \\ \frac{I(x'_1, x'_2, x_3) - I(x'_1, x_2, x_3)}{x'_2 - x_2} \\ \frac{I(x'_1, x'_2, x'_3) - I(x'_1, x'_2, x_3)}{x'_3 - x_3} \end{bmatrix}. \quad (2.29)$$

Other examples of discrete gradients, as well as constructions of the skew-symmetric matrix $S(\mathbf{x})$ for a given vector field \mathbf{f} and integral I may be found in (McLachlan et al. 1999).

We note that the discrete gradient method can also be used for systems with any number of integrals. For example an ODE $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$ possessing two integrals $I(\mathbf{x})$ and $J(\mathbf{x})$ can be written

$$\dot{x}_i = S_{ijk}(\mathbf{x}) \frac{\partial I(\mathbf{x})}{\partial x_j} \frac{\partial J(\mathbf{x})}{\partial x_k}, \quad (2.30)$$

where the summation convention is assumed over repeated indices and $S(\mathbf{x})$ is a completely antisymmetric tensor. A discretisation of (2.30) which preserves both I and J is given by

$$\frac{x'_i - x_i}{h} = \bar{S}_{ijk}(\mathbf{x}, \mathbf{x}') \bar{\nabla} I(\mathbf{x}, \mathbf{x}') \Big|_j \bar{\nabla} J(\mathbf{x}, \mathbf{x}') \Big|_k \quad (2.31)$$

with \bar{S} any completely skew approximation of S and $\bar{\nabla} I$ and $\bar{\nabla} J$ discrete gradients as defined above.

3 Five recent stories of GNI

The purpose of this section is not to present a totality of recent research into GNI, a subject that would have called for a substantially longer paper. Instead, we wish to highlight a small number of developments with which the authors are familiar and which provide a flavour of the very wide range of issues on the current GNI agenda.

3.1 Highly oscillatory Hamiltonian systems

High oscillation occurs in many Hamiltonian systems. Sometimes, e.g. in the integration of equations of celestial mechanics, the source of the problem is that we wish to compute the solution across a very large number of periods and the oscillation is an artefact of the time scale in which the solution has physical relevance. In other cases oscillation is implicit in the multiscale structure of the underlying problem. A case in point are the (modified) *Fermi–Pasta–Ulam (FPU) equations*, describing a mechanical

system consisting of alternating stiff harmonic and soft nonlinear springs. The soft springs impart fast oscillation, while the hard springs generate slow transfer of energy across the system: good numerical integration must capture both!

A good point to start (which includes modified FPU as a special case) is the second-order ODE

$$\ddot{\mathbf{q}} + \Omega^2 \mathbf{q} = \mathbf{g}(\mathbf{q}), \quad t \geq 0, \quad \mathbf{q}(0) = \mathbf{u}_0, \quad \dot{\mathbf{q}}(0) = \mathbf{v}_0, \quad (3.1)$$

where $\mathbf{g}(\mathbf{q}) = -\nabla U(\mathbf{q})$ and

$$\Omega = \begin{bmatrix} O & O \\ O & \omega I \end{bmatrix}, \quad \omega \gg 1, \quad \mathbf{q} = \begin{bmatrix} \mathbf{q}_0 \\ \mathbf{q}_1 \end{bmatrix}, \quad \mathbf{q}_0 \in \mathbb{R}^{n_0}, \quad \mathbf{q}_1 \in \mathbb{R}^{n_1}.$$

An important aspect of systems of the form (3.1) is that the exact solution, in addition to preserving the total Hamiltonian energy

$$H(\mathbf{p}, \mathbf{q}) = \frac{1}{2}(\|\mathbf{p}_1\|^2 + \omega^2 \|\mathbf{q}_1\|^2) + \frac{1}{2}\|\mathbf{p}_0\|^2 + U(\mathbf{q}_0, \mathbf{q}_1), \quad (3.2)$$

where $\dot{\mathbf{q}} = \mathbf{p}$, also preserves the *oscillatory energy*

$$I(\mathbf{p}, \mathbf{q}) = \frac{1}{2}\|\mathbf{p}_1\|^2 + \frac{\omega^2}{2}\|\mathbf{q}_1\|^2 \quad (3.3)$$

for intervals of length $\mathcal{O}(\omega^N)$ for any $N \geq 1$. This has been proved using the *modulated Fourier expansions*

$$\mathbf{q}(t) = \sum_{m=-\infty}^{\infty} e^{im\omega t} \mathbf{z}_m(t).$$

The solution of (3.1) exhibits oscillations at frequency $\mathcal{O}(\omega)$ and this inhibits the efficiency of many symplectic methods, requiring step size of $\mathcal{O}(\omega^{-1})$, a situation akin to stiffness in more conventional ODEs. However, by their very structure, exponential integrators (and in particular Gautschi-type methods (2.5)) are particularly effective in integrating the linear part, which gives rise to high oscillation. The problem with Gautschi-type methods, though, might be the occurrence of resonances and we need to be careful to avoid them, both in the choice of the right filter (cf. the discussion in Subsection 2.1) and step size h .

Of course, one would like geometric numerical integrators applied to (3.1) to exhibit favourable preservation properties with respect to both total energy (3.2) and oscillatory energy (3.3). Applying modulated Fourier expansions to trigonometric and modified trigonometric integrators, this is indeed the case provided that the step size obeys the *non-resonance condition* with respect to the frequency ω ,

$$|\sin(\frac{1}{2}mh\omega)| \geq ch^{1/2}, \quad m = 1, \dots, N, \quad N \geq 2,$$

cf. Hairer & Lubich (2009).

All this has been generalised to systems with multiple frequencies, with the Hamiltonian function

$$H(\mathbf{p}, \mathbf{q}) = \overbrace{\frac{1}{2} \sum_{j=1}^s (\|\mathbf{p}_j\|^2 + \omega_j^2 \|\mathbf{q}_j\|^2)}^{\text{oscillatory}} + \overbrace{\frac{1}{2} \|\mathbf{p}_0\|^2 + U(\mathbf{q})}^{\text{slow}},$$

where

$$\mathbf{p} = \begin{bmatrix} \mathbf{p}_0 \\ \mathbf{p}_1 \\ \vdots \\ \mathbf{p}_s \end{bmatrix}, \quad \mathbf{q} = \begin{bmatrix} \mathbf{q}_0 \\ \mathbf{q}_1 \\ \vdots \\ \mathbf{q}_s \end{bmatrix}, \quad 0 < \min_{j=1,\dots,s} \omega_j, \quad 1 \ll \max_{j=1,\dots,s} \omega_j$$

for both the exact solution (Gauckler, Hairer & Lubich 2013) and for discretisations obtained using trigonometric and modified trigonometric integrators (Cohen, Gauckler, Hairer & Lubich 2015).

Further achievements and open problem in the challenging area of marrying symplectic integration and high oscillation are beautifully described in (Hairer 2014).

3.2 Kahan’s ‘unconventional’ method

A novel discretisation method for quadratic ODEs was introduced and studied in (Kahan & Li 1997). This new method discretised the vector field

$$\dot{x}_i = \sum_{j,k} a_{ijk} x_j x_k + \sum_j b_{ij} x_j + c_i \quad (3.4)$$

as follows,

$$\frac{x'_i - x_i}{h} = \sum_{j,k} a_{ijk} \left(\frac{x_j x'_k + x'_j x_k}{2} \right) + \sum_j b_{ij} \left(\frac{x_j + x'_j}{2} \right) + c_i. \quad (3.5)$$

Kahan called the method (3.5) ‘unconventional’, because it treats the quadratic terms different from the linear terms. He also noted some nice features of (3.5), e.g. that it often seemed to be able to integrate through singularities.

Properties of Kahan’s method:

1. *Kahan’s method is (the reduction of) a Runge–Kutta method.*

Celledoni, McLachlan, Owren & Quispel (2013) showed that (3.5) is the reduction to quadratic vector fields of the Runge–Kutta method

$$\frac{\mathbf{x}' - \mathbf{x}}{h} = 2\mathbf{f} \left(\frac{\mathbf{x} + \mathbf{x}'}{2} \right) - \frac{1}{2}\mathbf{f}(\mathbf{x}) - \frac{1}{2}\mathbf{f}(\mathbf{x}') \quad (3.6)$$

This explains *inter alia* why Kahan’s method preserves all linear first integrals.

2. *Kahan’s method preserves a modified energy and measure.*

For any Hamiltonian vector field of the form

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) = S\nabla H(\mathbf{x}), \quad (3.7)$$

with cubic Hamiltonian $H(\mathbf{x})$ and constant symplectic (or Poisson) structure S , Kahan’s method preserves a modified energy as well as a modified measure exactly (Celledoni et al. 2013).

The modified volume is

$$\frac{dx_1 \wedge \cdots \wedge dx_n}{\det\left(I - \frac{1}{2}hf'(\mathbf{x})\right)}, \quad (3.8)$$

while the modified energy is

$$\tilde{H}(\mathbf{x}) := H(\mathbf{x}) + \frac{1}{3}h\nabla H(\mathbf{x})^\top \left(I - \frac{1}{2}hf'(\mathbf{x})\right)^{-1} \mathbf{f}(\mathbf{x}). \quad (3.9)$$

3. *Kahan's method preserves the integrability of many integrable systems of quadratic ODEs.*

Beginning with the work of Hirota and Kimura, subsequently extended by Suris and collaborators (Petrera, Pfadler & Suris 2011), and by Quispel and collaborators (Celledoni et al. 2013, Celledoni, McLachlan, McLaren, Owren & Quispel 2014, van der Kamp, Kouloukas, Quispel, Tran & Vanhaecke 2014), it was shown that Kahan's method preserves the complete integrability of a surprisingly large number of quadratic ODEs.

Here we list some 2D vector fields whose integrability is preserved by Kahan's method:

- Quadratic Hamiltonian systems in 2D:

The 9-parameter family

$$\begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = \begin{bmatrix} bx^2 + 2cxy + dy^2 + fx + gy + i \\ -ax^2 - 2bxy - cy^2 - ex - fy - h \end{bmatrix}; \quad (3.10)$$

- Suslov systems in 2D:

The 9-parameter family

$$\begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = l(x, y) \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \nabla H(x, y), \quad (3.11)$$

where $l(x, y) = ax + by + c$; $H(x, y) = dx^2 + exy + fy^2 + gx + hy + i$;

- Reduced Nahm equations in 2D:

Octahedral symmetry:

$$\begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = \begin{bmatrix} 2x^2 - 12y^2 \\ -6x^2 - 4y^2 \end{bmatrix}; \quad (3.12)$$

Icosahedral symmetry:

$$\begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = \begin{bmatrix} 2x^2 - y^2 \\ -10xy + y^2 \end{bmatrix}. \quad (3.13)$$

The modified energy and measure for the Kahan discretisations of these 2D systems, as well as of many other (higher-dimensional) integrable quadratic vector fields are given in (Petrera et al. 2011, Celledoni et al. 2013, Celledoni et al. 2014).

Generalisations to higher degree polynomial equations using polarisation are presented in (Celledoni, McLachlan, McLaren, Owren & Quispel 2015).

3.3 Applications to celestial mechanics

GNI methods particularly come into their own when the integration time is large compared to typical periods of the system. Thus long-term integrations of e.g. solar-type systems and of particle accelerators typically need symplectic methods. In this subsection we focus on the former⁵.

One of the first symplectic integrations of the solar system was done in (Sussman & Wisdom 1992) where it was confirmed that the solar system has a positive Lyapunov exponent, and hence exhibits chaotic behaviour cf (Laskar 2003).

More recently these methods have been improved and extended (McLachlan 1995*b*, Duncan, Levison & Lee 1998, Laskar, Fienga, Gastineau & Manche 2011, Blanes, Casas, Farrés, Laskar, Makazaga & Murua 2013). Several symplectic integrators of high order were tested in (Farrés, Laskar, Blanes, Casas, Makazaga & Murua 2013), in order to determine the best splitting scheme for long-term studies of the solar system.

These various methods have resulted in the fact that numerical algorithms for solar system dynamics are now so accurate that they can be used to define the geologic time scales in terms of the initial conditions and parameters of solar system models (or vice versa).

3.4 Symmetric Zassenhaus splitting and the equations of quantum mechanics

Equations of quantum mechanics in the semiclassical regime represent a double challenge of structure conservation and high oscillation. A good starting point is the linear Schrödinger equation

$$\frac{\partial u}{\partial t} = i\varepsilon \frac{\partial^2 u}{\partial x^2} - i\varepsilon^{-1}V(x)u \quad (3.14)$$

(for simplicity we restrict our discourse to a single space dimension), given in $[-1, 1]$ with periodic boundary conditions. Here V is the potential energy of a quantum system, $|u(x, t)|^2$ is a position density of a particle and $0 < \varepsilon \ll 1$ represents the difference in mass between an electron and nuclei. It is imperative to preserve the unitarity of the solution operator (otherwise $|u(\cdot, t)|^2$ is no longer a probability function), but also deal with oscillation at a frequency of $\mathcal{O}(\varepsilon^{-1})$. A conventional approach advances the solution using a palindromic splitting (2.3), but this is suboptimal for a number of reasons. Firstly, the number of splittings increases exponentially with order. Secondly, error constants are exceedingly large. Thirdly, quantifying the quality of approximation in terms of the step-size h is misleading, because there are three small quantities at play: the step size h , N^{-1} where N is the number of degrees of freedom in space discretisation (typically either a spectral method or spectral collocation) and, finally, $\varepsilon > 0$ which, originating in physics rather than being a numerical artefact, is the most important. We henceforth let $N = \mathcal{O}(\varepsilon^{-1})$ (to resolve the high-frequency oscillations) and $h = \mathcal{O}(\varepsilon^\sigma)$ for some $\sigma > 0$ – obviously, the smaller σ , the larger the time step.

⁵A very readable early review of integrators for solar system dynamics is (Morbidelli 2002*b*), cf also (Morbidelli 2002*a*)

Bader, Iserles, Kropielnicka & Singh (2014) have recently proposed an alternative approach to the splitting of (3.14), of the form

$$e^{ih(\varepsilon\partial_x^2 - \varepsilon^{-1}V)} \approx e^{\mathcal{R}_0} e^{\mathcal{R}_1} \dots e^{\mathcal{R}_s} e^{\mathcal{T}_{s+1}} e^{\mathcal{R}_s} \dots e^{\mathcal{R}_1} e^{\mathcal{R}_0} \quad (3.15)$$

such that $\mathcal{R}_k = \mathcal{O}(\varepsilon^{\alpha_k})$, $\mathcal{T}_{s+1} = \mathcal{O}(\varepsilon^{\alpha_{s+1}})$, where $\alpha_0 \leq \alpha_1 < \alpha_2 < \alpha_3 < \dots$ – the *symmetric Zassenhaus splitting*. Here $\partial_x = \partial/\partial x$.

The splitting (3.15) is derived at the level of differential operators (i.e., prior to space discretisation), applying the symmetric Baker–Campbell–Hausdorff formula to elements in the free Lie algebra spanned by ∂_x^2 and V . For $\sigma = 1$, for example, this yields

$$\begin{aligned} \mathcal{R}_0 &= -\frac{1}{2}\tau\varepsilon^{-1}V = \mathcal{O}(1), \\ \mathcal{R}_1 &= \frac{1}{2}\tau\varepsilon\partial_x^2 = \mathcal{O}(1), \\ \mathcal{R}_2 &= \frac{1}{24}\tau^3\varepsilon^{-1}(\partial_x V)^2 + \frac{1}{12}\tau^3\varepsilon\{(\partial_x^2 V)\partial_x^2 + \partial_x^2[(\partial_x^2 V)\cdot]\} = \mathcal{O}(\varepsilon^2), \\ \mathcal{R}_3 &= -\frac{1}{120}\tau^5\varepsilon^{-1}(\partial_x^2 V)(\partial_x V)^2 - \frac{1}{24}\tau^3\varepsilon(\partial_x^4 V) + \frac{1}{240}\tau^5\varepsilon(7\{(\partial_x^2 V)^2\partial_x^2 \\ &\quad + \partial_x^2[(\partial_x^2 V)^2\cdot]\} + \{(\partial_x^3 V)(\partial_x V)\partial_x^2 + \partial_x^2[(\partial_x^3 V)(\partial_x V)\cdot]\}) \\ &\quad + \frac{1}{120}\tau^5\varepsilon^{-3}\{(\partial_x^4 V)\partial_x^4 + \partial_x^4[(\partial_x^4 V)\cdot]\} = \mathcal{O}(\varepsilon^4), \end{aligned}$$

where $\tau = ih$. Note that all the commutators, ubiquitous in the BCH formula, have disappeared: in general, the commutators in this free Lie algebra can be replaced by linear combinations of derivatives, with the remarkable property of *height reduction*: each commutator ‘kills’ one derivative, e.g.

$$[V, \partial_x^2] = -(\partial_x^2 V) - 2(\partial_x V)\partial_x, \quad [[V, \partial_x^2], \partial_x^2] = (\partial_x^4 V) + 4(\partial_x^3 V)\partial_x + 4(\partial_x^2 V)\partial_x^2.$$

Once we discretise with spectral collocation, \mathcal{R}_0 becomes a diagonal matrix and its exponential is trivial, while $e^{\mathcal{R}_1}\mathbf{v}$ can be computed in two FFTs for any vector \mathbf{v} because \mathcal{R}_1 is a Toeplitz circulant. Neither \mathcal{R}_2 nor \mathcal{R}_3 possess useful structure, except that they are *small!* Therefore we can approximate $e^{\mathcal{R}_k}\mathbf{v}$ using the Krylov–Arnoldi process in just 3 and 2 iterations for $k = 2$ and $k = 3$, respectively, to attain an error of $\mathcal{O}(\varepsilon^6)$ (Bader et al. 2014).

All this has been generalised to time-dependent potentials and is applicable to a wider range of quantum mechanics equations in the semiclassical regime.

4 Beyond GNI

Ideas in one area of mathematical endeavour often inspire work in another area. This is true not just because new mathematical research equips us with a range of innovative tools but because it provides insight that casts new light not just on the subject in question but elsewhere in the mathematical universe. GNI has thus contributed not just toward its own goal, better understanding of structure-preserving discretisation methods for differential equations, but in other, often unexpected, directions.

4.1 GNI meets abstract algebra

The traditional treatment of discretisation methods for differential equations was wholly analytic, using tools of functional analysis and approximation theory. (Lately, also tools from algebraic topology.) GNI has added an emphasis on geometry and this leads in a natural manner into concepts and tools from abstract algebra. As often in such mathematical dialogues, while GNI borrowed much of its conceptual background from abstract algebra, it has also contributed to the latter, not just with new applications but also new ideas.

- *B-series and beyond.* Consider numerical integration methods that associate to each vector field \mathbf{f} a map $\psi_h(\mathbf{f})$. A method ψ_h is called g -covariant⁶ if the following diagram commutes,

$$\begin{array}{ccc}
 \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) & \xrightarrow{\mathbf{x} = \mathbf{g}(\mathbf{y})} & \dot{\mathbf{y}} = \tilde{\mathbf{f}}(\mathbf{y}) \\
 \downarrow & & \downarrow \\
 \tilde{\mathbf{x}} = \psi_h(\mathbf{f})(\mathbf{x}) & \xrightarrow{\mathbf{x} = \mathbf{g}(\mathbf{y})} & \tilde{\mathbf{y}} = \psi_h(\tilde{\mathbf{f}})(\mathbf{y})
 \end{array}$$

It follows that if g is a symmetry of the vector field f and ψ is g -covariant, then ψ preserves the symmetry g . It seems that this concept of covariance for integration methods was first introduced in (McLachlan 1995a) and (McLachlan, Quispel & Turner 1998).

It is not hard to check that all B-series methods are covariant with respect to the group of affine transformations. A natural question to ask then, was “are B-series methods the only numerical integration methods that preserve the affine group?”. This question was open for many years, until it was answered in the negative by (Munthe-Kaas & Verdier 2015), who introduced a more general class of integration methods dubbed “aromatic Butcher series”, and showed that (under mild assumptions) this is the most general class of methods preserving affine covariance. Expansions of methods in this new class contain both rooted trees (as in B-series), as well as products of rooted trees and so-called k -loops (Iserles et al. 2007).

Whereas it may be said that to date the importance of aromatic B-series has been at the formal rather than at the constructive level, these methods may hold the promise of the construction of affine-covariant volume-preserving integrators.

- *Word expansions.* Classical B-series can be significantly generalised by expanding in *word series* (Murua & Sanz-Serna 2015). This introduced an overarching

⁶Also called equivariant.

framework for Taylor expansions, Fourier expansions, modulated Fourier expansions and splitting methods. We consider an ODE of the form

$$\dot{\mathbf{x}} = \sum_{a \in \mathcal{A}} \lambda_a(t) \mathbf{f}_a(\mathbf{x}), \quad \mathbf{x}(0) = \mathbf{x}_0, \quad (4.1)$$

where \mathcal{A} is a given *alphabet*. The solution of (4.1) can be formally expanded in the form

$$\mathbf{x}(t) = \sum_{n=0}^{\infty} \sum_{\mathbf{w} \in \mathcal{W}_n} \alpha_{\mathbf{w}}(t) \mathbf{f}_{\mathbf{w}}(\mathbf{x}_0),$$

where \mathcal{W}_n is the set of all words with n letters from \mathcal{A} . The coefficients $\alpha_{\mathbf{w}}$ and functions $\mathbf{f}_{\mathbf{w}}$ can be obtained recursively from the λ_a s and \mathbf{f}_a s in a manner similar to B-series. Needless to say, exactly like with B-series, word series can be interpreted using an algebra over rooted trees.

The concept of word series is fairly new in numerical mathematics but it exhibits an early promise to provide a powerful algebraic tool for the analysis of dynamical systems and their discretisation.

- *Extension of Magnus expansions.* Let \mathcal{W} be a *Rota–Baxter algebra*, a commutative unital algebra equipped with a linear map R such that

$$R(x)R(y) = R(R(x)y + xR(y) + \theta xy), \quad x, y \in \mathcal{W},$$

where θ is a parameter. The inverse ∂ of R obeys

$$\partial(xy) = \partial(x)y + x\partial(y) + \theta\partial(x)\partial(y)$$

and is hence a generalisation of a derivation operator: a neat example, with clear numerical implications, is the backward difference $\partial(x) = [x(t) - x(t - \theta)]/\theta$. Ebrahimi-Fard & Manchon (2009) generalised Magnus expansions to this and similar settings, e.g. dendriform algebras. Their work uses the approach in (Iserles & Nørsett 1999), representing individual ‘Magnus terms’ as rooted trees, but generalises it a great deal.

- *The algebra of the Zassenhaus splitting.* The success of the Zassenhaus splitting (3.15) rests upon two features. Firstly, the replacement of commutators by simpler, more tractable expressions and, secondly, height reduction of derivatives under commutation. Singh (2015) has derived an algebraic structure \mathfrak{J} which, encoding these two features, allows for a far-reaching generalisation of the Zassenhaus framework. The elements of \mathfrak{J} are operators of the form $\langle f \rangle_k = f \circ \partial_x^k + \partial_x^k \circ f$, where $k \in \mathbb{Z}_+$ and f resides in a suitable function space. \mathfrak{J} can be endowed with a Lie-algebraic structure and, while bearing similarities with the Weyl algebra and the Heisenberg group, is a new and intriguing algebraic concept.

4.2 Highly oscillatory quadrature

Magnus expansions (2.9) are particularly effective when the matrix $A(t)$ oscillates rapidly. This might seem paradoxical – we are all conditioned to expect high oscillation to be ‘difficult’ – but actually makes a great deal of sense. Standard numerical methods are based on Taylor expansions, hence on *differentiation*, and their error typically scales as a high derivative of the solution. Once a function oscillates rapidly, differentiation roughly corresponds to multiplying amplitude by frequency, high derivatives become large and so does the error. However, the Magnus expansion does not differentiate, it *integrates!* This has an opposite effect: the more we integrate, the smaller the amplitude and the series (2.9) converges more rapidly. Indeed, often it pays to render a linear system highly oscillatory by a change of variables, in a manner described in (Iserles 2002), and then solve it considerably faster and cheaper. Yet, once we contemplate the discretisation of (2.9) for a highly oscillatory matrix function $A(t)$, we soon come up another problem, usually considered difficult, if not insurmountable: computing multivariate integrals of highly oscillatory functions.

In a long list of methods for highly oscillatory quadrature (HOQ) *circa* 2002, ranging from the useless to the dubious, one method stood out: Levin (1982) proposed to calculate univariate integrals by converting the problem to an ODE and using collocation. This was the only effective method around, yet incompletely understood.

The demands of GNI gave the initial spur to the emergence in the last ten years to a broad swath of new methods for HOQ: Filon-type methods, which replace the *non-oscillatory* portion of the integrand by an interpolating polynomial (Iserles & Nørsett 2005), improved Levin-type methods (Olver 2006) and the method of numerical stationary phase of Huybrechs & Vandewalle (2006). The common characteristic of all these methods is that they are based on asymptotic expansions. This means that high oscillation is no longer the enemy – indeed, the faster the oscillation, the smaller the error!

Highly oscillatory integrals occur in numerous applications, from electromagnetic and acoustic scattering to fluid dynamics, quantum mechanics and beyond. Their role in GNI is minor. However, their modern numerical theory was originally motivated by a problem in GNI. This is typical to how scholarship progresses and it is only natural that HOQ has severed its GNI moorings and has become an independent area on its own.

4.3 Structured linear algebra

GNI computations often lead to specialised problems in numerical linear algebra. However, structure preservation has wider impact in linear algebraic computations. Often a matrix in an algebraic problem belongs to an algebraic structure, e.g. a specific Lie algebra or a symmetric space, and it is important to retain this in computation – the sobriquet “Geometric Numerical Algebra” might be appropriate! Moreover, as in GNI so in GNA, respecting structure often leads to better, more accurate and cheaper numerical methods. Finally, structured algebraic computation is often critical to GNI computations.

- Matrix factorization is the lifeblood of numerical algebra, the basis of the most

effective algorithms for the solution of linear systems, computation of eigenvalues and solution of least-squares problems. A major question in GNA is “Suppose that $A \in \mathcal{A}$, where \mathcal{A} is a set of matrices of given structure. Given a factorization $A = BC$ according to some set of rules, what can we say about the structure of B or C ?”. Mackey, Mackey & Tisseur (2005) addressed three such ‘factorization rules’: the *matrix square root*, $B = C$, the *matrix sign*, where the elements of B are ± 1 , and the *polar decomposition*, with unitary B and positive semidefinite C . They focussed on sets \mathcal{A} generated by a sesquilinear form $\langle \cdot, \cdot \rangle$. Such sets conveniently fit into two classes:

- (a) Automorphisms G , such that $\langle G\mathbf{x}, G\mathbf{y} \rangle = \langle \mathbf{x}, \mathbf{y} \rangle$, generate a *Lie group*;
- (b) Self-adjoint matrices S , such that $\langle S\mathbf{x}, \mathbf{y} \rangle = \langle \mathbf{x}, S\mathbf{y} \rangle$, generate a *Jordan algebra*; and
- (c) Skew-adjoint matrices H such that $\langle H\mathbf{x}, \mathbf{y} \rangle = -\langle \mathbf{x}, H\mathbf{y} \rangle$, generate a *Lie algebra*.

It is natural to expect that conservation of structure under factorization would depend on the nature of the underlying inner product. The surprising outcome of (Mackey et al. 2005) is that, for current purposes, it is sufficient to split sesquilinear forms into just two classes, unitary and orthosymmetric, each exhibiting similar behaviour.

- Many algebraic eigenvalue problems are structured, the simplest example being that the eigenvalues of a symmetric matrix are real and of a skew-symmetric are pure imaginary: all standard methods for the computation of eigenvalues respect this. However, many other problems might have more elaborate structure, and this is the case in particular for nonlinear eigenvalue problems. An important example, with significant applications in mechanics, is

$$(\lambda^2 M + \lambda G + K)\mathbf{x} = \mathbf{0}, \tag{4.2}$$

where both M and K are symmetric, while G is skew symmetric. The eigenvalues λ of (4.2) exhibit *Hamiltonian* pattern: if λ is in the spectrum then so are $-\lambda$, $\bar{\lambda}$ and $-\bar{\lambda}$.⁷ As often in numerical algebra, (4.2) is particularly relevant when the underlying matrices are large and sparse.

Numerical experiments demonstrate that standard methods for the computation of a quadratic eigenvalue problems may fail to retain the Hamiltonian structure of the spectrum but this can be obtained by bespoke algorithms, using a symplectic version of the familiar Lanczos algorithm, cf. (Benner, Fassbender & Stoll 2007/08).

This is just one example of the growing field of structured eigenvalue and inverse eigenvalue problems.

- The exponential from an algebra to a group: Recall Lie-group methods from Section 2.2: a critical step, e.g. in the RKMK methods, is the exponential map

⁷To connect this to the GNI narrative, such a pattern is displayed by matrices in the *symplectic Lie algebra* $\mathfrak{sp}(2n)$.

from a Lie algebra to a Lie group. Numerical analysis knows numerous effective ways to approximate the matrix exponential (Moler & Van Loan 2003), yet most of them fail to map a matrix from a Lie algebra to a Lie group! There is little point to expand intellectual and computational effort to preserve structure, only to abandon the latter in the ultimate step, and this explains the interest in the computation of the matrix exponential which is assured to map A in a Lie algebra to an element in the corresponding Lie group.

While early methods have used structure constants and, for maximal sparsity, Lie-algebraic bases given by space-root decomposition (Celledoni & Iserles 2001), the latest generation of algorithms is based upon *generalised polar decomposition* (Munthe-Kaas, Quispel & Zanna 2001).

Acknowledgments

This work has been supported by the Australian Research Council. The authors are grateful to David McLaren for assistance during the preparation of this paper, as well as to Philipp Bader, Robert McLachlan and Marcus Webb, whose comments helped to improve this paper.

References

- Bader, P., Iserles, A., Kropielnicka, K. & Singh, P. (2014), ‘Effective approximation for the semiclassical Schrödinger equation’, *Found. Comput. Math.* **14**(4), 689–720.
- Benner, P., Fassbender, H. & Stoll, M. (2007/08), ‘Solving large-scale quadratic eigenvalue problems with Hamiltonian eigenstructure using a structure-preserving Krylov subspace method’, *Electron. Trans. Numer. Anal.* **29**, 212–229.
- Blanes, S., Casas, F., Farrés, A., Laskar, J., Makazaga, J. & Murua, A. (2013), ‘New families of symplectic splitting methods for numerical integration in dynamical astronomy’, *Appl. Numer. Math.* **68**, 58–72.
- Blanes, S., Casas, F., Oteo, J. A. & Ros, J. (2009), ‘The Magnus expansion and some of its applications’, *Phys. Rep.* **470**(5-6), 151–238.
- Butcher, J. C. (1963), ‘Coefficients for the study of Runge-Kutta integration processes’, *J. Austral. Math. Soc.* **3**, 185–201.
- Celledoni, E. & Iserles, A. (2001), ‘Methods for the approximation of the matrix exponential in a Lie-algebraic setting’, *IMA J. Numer. Anal.* **21**(2), 463–488.
- Celledoni, E., Grimm, V., McLachlan, R. I., McLaren, D. I., O’Neale, D., Owren, B. & Quispel, G. R. W. (2012), ‘Preserving energy resp. dissipation in numerical PDEs using the “average vector field” method’, *J. Comput. Phys.* **231**(20), 6770–6789.
- Celledoni, E., McLachlan, R. I., McLaren, D. I., Owren, B. & Quispel, G. R. W. (2014), ‘Integrability properties of Kahan’s method’, *J. Phys. A* **47**(36), 365202, 20.

- Celledoni, E., McLachlan, R. I., McLaren, D. I., Owren, B. & Quispel, G. R. W. (2015), ‘Discretization of polynomial vector fields by polarization’, *Proc. R. Soc. A*. DOI: 10.1098/rspa.2015.0390, Published 16 Dec. 2015.
- Celledoni, E., McLachlan, R. I., Owren, B. & Quispel, G. R. W. (2013), ‘Geometric properties of Kahan’s method’, *J. Phys. A* **46**(2), 025201, 12.
- Chartier, P. & Murua, A. (2007), ‘Preserving first integrals and volume forms of additively split systems’, *IMA J. Numer. Anal.* **27**(2), 381–405.
- Cohen, D., Gauckler, L., Hairer, E. & Lubich, C. (2015), ‘Long-term analysis of numerical integrators for oscillatory Hamiltonian systems under minimal non-resonance conditions’, *BIT* **55**(3), 705–732.
- Connes, A. & Kreimer, D. (1999), ‘Lessons from quantum field theory: Hopf algebras and spacetime geometries’, *Lett. Math. Phys.* **48**(1), 85–96.
- Duncan, M., Levison, H. F. & Lee, M. H. (1998), ‘A multiple time step symplectic algorithm for integrating close encounters’, *Astron. J.* **116**, 2067–2077.
- Ebrahimi-Fard, K. & Manchon, D. (2009), ‘A Magnus- and Fer-type formula in dendriform algebras’, *Found. Comput. Math.* **9**(3), 295–316.
- Farrés, A., Laskar, J., Blanes, S., Casas, F., Makazaga, J. & Murua, A. (2013), ‘High precision symplectic integrators for the Solar System’, *Celestial Mech. Dynam. Astronom.* **116**(2), 141–174.
- Feng, K. & Shang, Z. J. (1995), ‘Volume-preserving algorithms for source-free dynamical systems’, *Numer. Math.* **71**(4), 451–463.
- Feng, K., Wu, H. M., Qin, M. Z. & Wang, D. L. (1989), ‘Construction of canonical difference schemes for Hamiltonian formalism via generating functions’, *J. Comput. Math.* **7**(1), 71–96.
- Gauckler, L., Hairer, E. & Lubich, C. (2013), ‘Energy separation in oscillatory Hamiltonian systems without any non-resonance condition’, *Comm. Math. Phys.* **321**(3), 803–815.
- Ge, Z. & Marsden, J. E. (1988), ‘Lie-Poisson Hamilton-Jacobi theory and Lie-Poisson integrators’, *Phys. Lett. A* **133**(3), 134–139.
- Hairer, E. (2010), ‘Energy-preserving variant of collocation methods’, *JNAIAM. J. Numer. Anal. Ind. Appl. Math.* **5**(1-2), 73–84.
- Hairer, E. (2014), Challenges in Geometric Numerical Integration, in ‘Trends in Contemporary Mathematics’, Springer Verlag, pp. 125–135.
- Hairer, E. & Lubich, C. (2000), ‘Long-time energy conservation of numerical methods for oscillatory differential equations’, *SIAM J. Numer. Anal.* **38**(2), 414–441 (electronic).

- Hairer, E. & Lubich, C. (2009), Oscillations over long times in numerical Hamiltonian systems, in ‘Highly oscillatory problems’, Vol. 366 of *London Math. Soc. Lecture Note Ser.*, Cambridge Univ. Press, Cambridge, pp. 1–24.
- Hairer, E., Lubich, C. & Wanner, G. (2003), ‘Geometric numerical integration illustrated by the Störmer-Verlet method’, *Acta Numer.* **12**, 399–450.
- Hairer, E., Lubich, C. & Wanner, G. (2006), *Geometric Numerical Integration*, Vol. 31 of *Springer Series in Computational Mathematics*, second edn, Springer-Verlag, Berlin.
- Hairer, E., Nørsett, S. P. & Wanner, G. (1993), *Solving ordinary differential equations. I*, Vol. 8 of *Springer Series in Computational Mathematics*, second edn, Springer-Verlag, Berlin.
- Hochbruck, M. & Ostermann, A. (2010), ‘Exponential integrators’, *Acta Numer.* **19**, 209–286.
- Huybrechs, D. & Vandewalle, S. (2006), ‘On the evaluation of highly oscillatory integrals by analytic continuation’, *SIAM J. Numer. Anal.* **44**(3), 1026–1048.
- Iavernaro, F. & Trigiante, D. (2009), ‘High-order symmetric schemes for the energy conservation of polynomial Hamiltonian problems’, *JNAIAM J. Numer. Anal. Ind. Appl. Math.* **4**(1-2), 87–101.
- Iserles, A. (2002), ‘On the global error of discretization methods for highly-oscillatory ordinary differential equations’, *BIT* **42**(3), 561–599.
- Iserles, A. & Nørsett, S. P. (1999), ‘On the solution of linear differential equations in Lie groups’, *R. Soc. Lond. Philos. Trans. Ser. A Math. Phys. Eng. Sci.* **357**(1754), 983–1019.
- Iserles, A. & Nørsett, S. P. (2005), ‘Efficient quadrature of highly oscillatory integrals using derivatives’, *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **461**(2057), 1383–1399.
- Iserles, A., Munthe-Kaas, H. Z., Nørsett, S. P. & Zanna, A. (2000), ‘Lie-group methods’, *Acta Numerica* **9**, 215–365.
- Iserles, A., Quispel, G. R. W. & Tse, P. S. P. (2007), ‘B-series methods cannot be volume-preserving’, *BIT* **47**(2), 351–378.
- Kahan, W. & Li, R.-C. (1997), ‘Unconventional schemes for a class of ordinary differential equations—with applications to the Korteweg-de Vries equation’, *J. Comput. Phys.* **134**(2), 316–331.
- Lasagni, F. M. (1988), ‘Canonical Runge-Kutta methods’, *Z. Angew. Math. Phys.* **39**(6), 952–953.
- Laskar, J. (2003), ‘Chaos in the solar system’, *Ann. Henri Poincaré* **4**(suppl. 2), S693–S705.

- Laskar, J., Fienga, A., Gastineau, M. & Manche, H. (2011), ‘La2010: a new orbital solution for the long-term motion of the Earth’, *Astronomy & Astrophysics*. DOI: 10.1051/0004-6361/201116836.
- Levin, D. (1982), ‘Procedures for computing one- and two-dimensional integrals of functions with rapid irregular oscillations’, *Math. Comp.* **38**(158), 531–538.
- Lord, G., Malham, S. J. A. & Wiese, A. (2008), ‘Efficient strong integrators for linear stochastic systems’, *SIAM J. Numer. Anal.* **46**(6), 2892–2919.
- Mackey, D. S., Mackey, N. & Tisseur, F. (2005), ‘Structured factorizations in scalar product spaces’, *SIAM J. Matrix Anal. Appl.* **27**(3), 821–850.
- Magnus, W. (1954), ‘On the exponential solution of differential equations for a linear operator’, *Comm. Pure Appl. Math.* **7**, 649–673.
- Marsden, J. E. & West, M. (2001), ‘Discrete mechanics and variational integrators’, *Acta Numer.* **10**, 357–514.
- McLachlan, R. (1995*a*), ‘Comment on: “Poisson schemes for Hamiltonian systems on Poisson manifolds” [Comput. Math. Appl. **27** (1994), no. 12, 7–16; MR1284126 (95d:65069)] by W. J. Zhu and M. Z. Qin’, *Comput. Math. Appl.* **29**(3), 1.
- McLachlan, R. I. (1995*b*), ‘Composition methods in the presence of small parameters’, *BIT* **35**(2), 258–268.
- McLachlan, R. I. & Quispel, G. R. W. (2002), ‘Splitting methods’, *Acta Numer.* **11**, 341–434.
- McLachlan, R. I., Quispel, G. R. W. & Robidoux, N. (1999), ‘Geometric integration using discrete gradients’, *R. Soc. Lond. Philos. Trans. Ser. A Math. Phys. Eng. Sci.* **357**(1754), 1021–1045.
- McLachlan, R. I., Quispel, G. R. W. & Turner, G. S. (1998), ‘Numerical integrators that preserve symmetries and reversing symmetries’, *SIAM J. Numer. Anal.* **35**(2), 586–599.
- Moler, C. & Van Loan, C. (2003), ‘Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later’, *SIAM Rev.* **45**(1), 3–49 (electronic).
- Morbidelli, A. (2002*a*), *Modern Celestial Mechanics*, Gordon & Breach, London.
- Morbidelli, A. (2002*b*), ‘Modern integrations of solar system dynamics’, *Annu. Rev. Earth Planet. Sci.* **30**, 89–112.
- Munthe-Kaas, H. (1998), ‘Runge-Kutta methods on Lie groups’, *BIT* **38**(1), 92–111.
- Munthe-Kaas, H. & Verdier, O. (2015), ‘Aromatic Butcher series’, *Found. Comput. Math.* To appear, DOI: DOI 10.1007/s10208-015-9245-0.
- Munthe-Kaas, H. Z., Quispel, G. R. W. & Zanna, A. (2001), ‘Generalized polar decompositions on Lie groups with involutive automorphisms’, *Found. Comput. Math.* **1**(3), 297–324.

- Murua, A. & Sanz-Serna, J. M. (2015), Word series for dynamical systems and their numerical integrators, Technical report, Universidad Carlos III de Madrid. arXiv: 1502.05528v2 [math.NA].
- Neishtadt, A. I. (1984), ‘The separation of motions in systems with rapidly rotating phase’, *Prikl. Mat. Mekh.* **48**(2), 197–204.
- Olver, S. (2006), ‘On the quadrature of multivariate highly oscillatory integrals over non-polytope domains’, *Numer. Math.* **103**(4), 643–665.
- Owren, B. & Marthinsen, A. (2001), ‘Integration methods based on canonical coordinates of the second kind’, *Numer. Math.* **87**(4), 763–790.
- Petrera, M., Pfadler, A. & Suris, Y. B. (2011), ‘On integrability of Hirota–Kimura type discretizations’, *Regul. Chaotic Dyn.* **16**(3-4), 245–289.
- Quispel, G. R. W. (1995), ‘Volume-preserving integrators’, *Phys. Lett. A* **206**(1-2), 26–30.
- Quispel, G. R. W. & McLaren, D. I. (2008), ‘A new class of energy-preserving numerical integration methods’, *J. Phys. A* **41**(4), 045206, 7.
- Ramos, A. G. C. P. & Iserles, A. (2015), ‘Numerical solution of Sturm–Liouville problems via Fer streamers’, *Numer. Math.* **131**(3), 541–565.
- Sanz-Serna, J. M. (1988), ‘Runge–Kutta schemes for Hamiltonian systems’, *BIT* **28**(4), 877–883.
- Sanz-Serna, J. M. & Calvo, M. P. (1994), *Numerical Hamiltonian Problems*, Vol. 7 of *Applied Mathematics and Mathematical Computation*, Chapman & Hall, London.
- Shang, Z. J. (1994), ‘Generating functions for volume-preserving mappings and Hamilton–Jacobi equations for source-free dynamical systems’, *Sci. China Ser. A* **37**(10), 1172–1188.
- Singh, P. (2015), Algebraic theory for higher-order methods in computational quantum mechanics, Technical report, DAMTP, University of Cambridge. arXiv: 1510.06896v1 [math.NA].
- Suris, Y. B. (1988), Preservation of symplectic structure in the numerical solution of Hamiltonian systems, in ‘Numerical solution of ordinary differential equations (Russian)’, Akad. Nauk SSSR, Inst. Prikl. Mat., Moscow, pp. 148–160, 232, 238–239.
- Sussman, G. J. & Wisdom, J. (1992), ‘Chaotic evolution of the solar system’, *Science* **257**(5066), 56–62.
- van der Kamp, P. H., Kouloukas, T. E., Quispel, G. R. W., Tran, D. T. & Vanhaecke, P. (2014), ‘Integrable and superintegrable systems associated with multi-sums of products’, *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **470**(2172), 20140481, 23.

Wu, X., You, X. & Wang, B. (2013), *Structure-Preserving Algorithms for Oscillatory Differential Equations*, Springer, Heidelberg; Science Press Beijing, Beijing.

Yoshida, H. (1990), 'Construction of higher order symplectic integrators', *Phys. Lett. A* **150**(5-7), 262–268.