CANADIAN APPLIED
MATHEMATICS QUARTERLY
Volume 14, Number 3, Fall 2006

STRUCTURE-PRESERVING AND EXPONENTIAL DISCRETIZATIONS OF INITIAL VALUE PROBLEMS

JOHN C. BOWMAN

ABSTRACT. Specialized integration algorithms for initial value problems, obtained by applying conventional explicit discretizations in a transformed space, are described. One example, conservative integration, is motivated by a theorem of Ge Zhong and Marsden [**17**] that establishes that in the absence of explicit time dependence, one must in practice choose between preserving symplecticity or conserving the Hamiltonian. Another example, exponential integration, is well suited to highly stiff ordinary differential equations. Fully Lagrangian methods for advection are shown to be a special case of exponential integration.

1 Introduction Many numerical libraries used in scientific computing offer a so-called "black-box" integrator that attempts to provide the user with a general solver for a generic ordinary differential equation. These solvers require little knowledge on part of the user as to the detailed dynamical behaviour of the system under consideration. This means that mathematical models are often discretized according to algorithms that have little to do with the original problem. So although the generic nature of these algorithms is often considered by users to be advantageous, Iserles [**26**] has pointed out that this feature is also their biggest disadvantage, as they cannot exploit the analytical structure of the equation. Instead, interest in the community is now shifting towards developing computational algorithms that are tailored to reflect known structural features of the problem under consideration (e.g., see [**13**], [**27**], [**34**], [**12**] and [**20**]). Such schemes typically end up being numerically more robust, exhibiting larger regions of stability and allowing the use of larger time steps.

Indeed, one can devise integration methods that respect desired properties of ordinary differential equations such as first integrals [**33, 28**] or

symplecticity [**30, 11, 31**]. For example, traditional numerical integration algorithms, which are polynomials in the time step, typically lead to systematic drifts of nonlinear first integrals. However, by applying traditional integrators in a transformed space, it is possible to obtain integrators that conserve first integrals to machine precision.

For the subclass of ordinary differential equations that are Hamiltonian, one can devise schemes that are symplectic; they conserve phase-space structure (including volume) by making the time step map a canonical transformation. In this work we revisit an interesting theorem of Ge Zhong and Marsden [**17**]: a symplectic map *with no explicit time-dependence* will conserve the actual value of the Hamiltonian only if it is the exact solution, up to a reparametrization of time. The implication is that unless one has the luxury of having an exact analytical solution to the problem, one must choose between having a conservative integrator or a symplectic one. The condition of no explicit time dependence was actually not mentioned by Ge and Marsden, but it left open the possibility of conservative variational symplectic integrators based on explicitly time-dependent symplectic maps, which were later developed for certain mechanics problems [**27**]. In this work we carefully restate and prove the theorem to make this condition explicit.

For time-independent symplectic maps and Hamiltonians, the Ge-Marsden theorem implies in practice that one must make a choice between developing a conservative integrator or a symplectic one; for some problems in classical mechanics, a conservative method can be shown to be superior (cf. Figure 1).

Conservative integrators are derived by transforming the dependent variable. Another class of integrators can be derived by transforming the independent variable and using the technique of operator splitting, which allows one to combine different numerical discretizations that are tailored to individual terms of an equation. For example, it can be applied to initial value problems when the temporal evolution can be expressed as a slowly varying perturbation of an exactly solvable differential equation. If the exactly solvable piece is stationary, so that its Green's function depends only on relative time differences, it is possible to discretize the perturbed problem with a scheme that is exact on the time scale of the solvable part. Examples of this method, exponential integrators and Lagrangian discretizations of advective equations, are briefly discussed in Section 3.

Other numerical discretizations that respect analytical structure include those that preserve the positive semi-definiteness of covariance matrices [**5, 4**] and unitarity (which underlies the conservation of the

trace of quantum-mechanical probability density matrices) [**32, 34**].

2 Conservative vs. symplectic integrators Given $\boldsymbol{f} : \mathbb{R}^{n+1} \to \mathbb{R}^n$, we consider the initial value problem obtained by evolving a vector $\boldsymbol{x} \in \mathbb{R}^n$ from the initial condition $\boldsymbol{x}(0) = \boldsymbol{x}_0$ according to

$$\frac{d\boldsymbol{x}}{dt} = \boldsymbol{f}(\boldsymbol{x}, t). \tag{1}$$

If $n = 2k$ and $\boldsymbol{x} = (\boldsymbol{q}, \boldsymbol{p})$ where $\boldsymbol{q}, \boldsymbol{p} \in \mathbb{R}^k$ satisfy

$$\frac{d\boldsymbol{q}}{dt} = \frac{\partial H}{\partial \boldsymbol{p}},$$
$$\frac{d\boldsymbol{p}}{dt} = -\frac{\partial H}{\partial \boldsymbol{q}},$$

for some function $H(\boldsymbol{q}, \boldsymbol{p}, t) : \mathbb{R}^{n+1} \to \mathbb{R}$, one says that (1) is Hamiltonian. Often, the Hamiltonian H has no explicit dependence on t. For such systems, Ge and Marsden noticed the following result (although they did not mention the condition of no explicit time-dependence).

Theorem 1 (Ge Zhong and Marsden 1988)**.** *A C^1 symplectic map M with no explicit time-dependence will conserve a C^1 time-independent Hamiltonian $H : \mathbb{R}^n \to \mathbb{R}$ $\iff$ M is identical to the exact evolution, up to a reparametrization of time.*

Proof. A C^1 symplectic scheme is a canonical map M corresponding to some approximate C^1 Hamiltonian $\widetilde{H}_\tau(\boldsymbol{x}, t) : \mathbb{R}^{n+1} \to \mathbb{R}$, where the label τ denotes the time step. If the mapping M does not depend explicitly on time, it can be generated by the approximate Hamiltonian $K(\boldsymbol{x}) = \widetilde{H}_\tau(\boldsymbol{x}, 0)$.

Suppose the symplectic map conserves the true Hamiltonian H:

$$0 = \frac{dH}{dt} = \frac{\partial H}{\partial q_i}\frac{dq_i}{dt} + \frac{\partial H}{\partial p_i}\frac{dp_i}{dt} + \frac{\partial H}{\partial t} = [H, K],$$

where

$$[H, K] = \frac{\partial H}{\partial q_i}\frac{\partial K}{\partial p_i} - \frac{\partial H}{\partial p_i}\frac{\partial K}{\partial q_i}.$$

The implicit function theorem then guarantees that in a neighbourhood of $\boldsymbol{x}_0 \in \mathbb{R}^n$ there exists a C^1 function $\phi : \mathbb{R} \to \mathbb{R}$ such that

$$H(\boldsymbol{x}) = \phi(K(\boldsymbol{x})) \quad \text{or} \quad K(\boldsymbol{x}) = \phi(H(\boldsymbol{x})) \iff [H, K] = 0.$$

Consequently, the trajectories in $\mathbb{R}^n$ generated by the Hamiltonians H and K coincide. □

While conservative variational symplectic integrators based on explicitly time-dependent symplectic maps have been proposed for certain mechanics problems [**27**] to circumvent the conditions of the Ge-Marsden theorem, in the time-independent case one must make a choice between symplecticity and conservation. In the 1990s, many papers on symplectic integration were to be found in the literature, but the available methods for enforcing conservation laws to machine precision typically required an *a posteriori* projection of a nonconservative solution back onto the energy manifold [**1, 25, 9**]. The need for a systematic explicit conservative method that evolves the system directly on the energy manifold, without arbitrarily removing one particularly degree of freedom (which is seldom practical [**18**]), led Shadwick, Bowman and Morrison [**33**] to develop a new class of algorithms known as conservative integrators. Traditional numerical discretizations of nonlinear initial value problems, based on polynomial functions of the time step, typically yield spurious secular drifts of nonlinear first integrals of motion (e.g., total energy), so that the numerical solution will *not* remain on the energy surface defined by the initial conditions. Shadwick, *et al.* [**33**] proposed a method for developing nontraditional explicit algorithms that exactly conserve nonlinear invariants to all orders in the time step (i.e., to machine precision).

For example, a three-mode truncation of the Fourier-transformed Euler equations for an inviscid 2D fluid

$$\begin{aligned}\frac{dx_1}{dt} &= f_1 = M_1 x_2 x_3,\\ \frac{dx_2}{dt} &= f_2 = M_2 x_3 x_1,\\ \frac{dx_3}{dt} &= f_3 = M_3 x_1 x_2,\end{aligned}$$

where $M_1 + M_2 + M_3 = 0$, conserves the energy $E = \frac{1}{2}\sum_k x_k^2$ since $\sum_k f_k x_k = 0$.

However, E is not conserved by conventional discretizations. For example, the Euler method $x_k(t+\tau) = x_k(t)+\tau f_k$ yields a monotonically

increasing new energy:

$$E(t+\tau) = \frac{1}{2}\sum_k \left[x_k^2 + 2\tau f_k x_k + \tau^2 S_k^2\right]$$
$$= E(t) + \frac{1}{2}\tau^2 \sum_k S_k^2.$$

Shadwick, *et al.* [**33**] determined a modification g_k of the original equations of motion that leads to *exact* energy conservation:

$$\frac{dx_k}{dt} = f_k + g_k.$$

On applying Euler's method to the system, one sees that the new energy evolves according to

$$E(t+\tau) = \frac{1}{2}\sum_k \left[x_k + \tau(f_k + g_k)\right]^2$$
$$= E(t) + \frac{1}{2}\sum_k \left[2\tau g_k x_k + \tau^2 (f_k + g_k)^2\right].$$

Energy conservation can be enforced by requiring that each term of the sum separately vanishes, thereby treating all modes on an equal footing. On solving for g_k, one then obtains the C-Euler discretization:

$$x_k(t+\tau) = \operatorname{sgn} x_k(t+\tau)\sqrt{x_k^2 + 2\tau f_k x_k}.$$

Here, a conventional Euler method can be used to determine the sign of $x_k(t+\tau)$. An important property of the C-Euler method is that it reduces to Euler's method as $\tau \to 0$:

$$x_k(t+\tau) = x_k\sqrt{1 + 2\tau\frac{f_k}{x_k}}$$
$$= x_k + \tau f_k + \mathcal{O}(\tau^2).$$

An easier way to derive the C-Euler scheme is to note that it is the result of applying the usual Euler algorithm to the transformed equation

$$\frac{dx_k^2}{dt} = 2 f_k x_k.$$

This suggests a general method for conservative integration: simply transform the dependent variables so that the conserved quantity can be expressed as a linear function of the transformed variables, apply a conventional integration method in the new space, and then transform back to the original space [**33**]. Indeed, the following trivial lemma [**28**] establishes that any multistage Runge-Kutta method will conserve such linear invariants:

Lemma 1. *Let $\boldsymbol{x}$ and $\boldsymbol{c}$ be vectors in $\mathbb{R}^n$. If $\boldsymbol{f} : \mathbb{R}^{n+1} \to \mathbb{R}^n$ has values orthogonal to $\boldsymbol{c}$, so that $I = \boldsymbol{c} \cdot \boldsymbol{x}$ is a linear invariant of*

$$\frac{d\boldsymbol{x}}{dt} = \boldsymbol{f}(\boldsymbol{x}, t).$$

Then each stage of the explicit s-stage discretization

$$\boldsymbol{x}_i = \boldsymbol{x}_0 + \tau \sum_{j=0}^{i-1} a_{ij} \boldsymbol{f}(\boldsymbol{x}_j, t + b_j \tau), \qquad i = 1, \ldots, s, \tag{2}$$

also conserves I, where τ is the time step and $a_{ij} \in \mathbb{R}$.

That is, higher-order conservative integration is made possible by finding a transformation $\boldsymbol{T} : \mathbb{R}^n \to \mathbb{R}^n$ such that the nonlinear invariants are linear functions of the new variable $\boldsymbol{\xi} = \boldsymbol{T}(\boldsymbol{x})$. Since the intermediate stage predictions appear in the final stage only through the function $\boldsymbol{f}$, only the final stage actually needs to be computed in the transformed space. The new value of $\boldsymbol{x}$ is then obtained by inverse transformation:

$$\boldsymbol{x}(t + \tau) = \boldsymbol{T}^{-1}(\boldsymbol{\xi}(t + \tau)).$$

Of course, the problem immediately arises that $\boldsymbol{T}$ may not be invertible. The following solutions are available:

1. Reduce the time step; for an analytic transformation $\boldsymbol{T}$, the inversion will be possible for a sufficiently small time step if the zeros of $\det(\boldsymbol{T}')$ are isolated (cf. Appendix A of [**28**]).
2. Switch to a traditional integrator for that time step.
3. Use an implicit backwards step (cf. Appendix A of [**33**]).

A simple second-order integrator, the predictor-corrector (PC) algorithm, can be obtained from (2) by setting $s = 2$, $b_0 = 0$, $a_{10} = b_1 = 1$, and $a_{20} = a_{21} = 1/2$. The conservative predictor-corrector (C-PC) obtained with the transformation $T(x) = x^2$ uses the conventional Euler

method as the predictor and also to determine the correct branch of T^{-1}:

$$\widetilde{x}_k = x_k + \tau f_k,$$
$$x_k(t+\tau) = \mathrm{sgn}(\widetilde{x}_k)\sqrt{x_k^2 + \tau(f_k x_k + \widetilde{f}_k \widetilde{x}_k)},$$

where $\widetilde{f}_k = f(\widetilde{x}_k)$.

For the one-dimensional autonomous case $dx/dt = f(x)$, the exact solution may be written as

$$x_\tau = x_0 + \tau f + \frac{\tau^2}{2} f' f + \frac{\tau^3}{6}(f'' f^2 + f'^2 f) + \mathcal{O}(\tau^4),$$

where all quantities on the right-hand side are evaluated at x_0. In comparison, when $T'(x_0) \neq 0$, C-PC yields the solution

$$x_\tau = x_0 + \tau f + \frac{\tau^2}{2} f' f + \frac{\tau^3}{4}\left(f'' f^2 + \frac{T'''}{3T'} f^3\right) + \mathcal{O}(\tau^4).$$

On setting $T(x) = x$, one sees that the C-PC solution reduces to the conventional PC prediction. Both the C-PC and PC algorithms are accurate to second order in τ; for $T(x) = x^2$, they agree through third order in τ.

In the singular case where $T'(x_0) = 0$, the conservative corrector reduces to

$$x_\tau = T^{-1}\left(T(x_0) + \frac{\tau}{2} T'(\tilde{x}) f(\tilde{x})\right);$$

if T and f are both analytic, the existence of a solution is then guaranteed as $\tau \to 0^+$ if the points at which T' vanishes are isolated.

Using the time step $\tau = 10^{-3}$, we see for the four-body classical mechanics problem in Figure 1 that conservative integration appears to yield more accurate results than the conventional second-order symplectic kinetic-potential energy (SKP) splitting

$$\tilde{p}_i = p_i - \frac{\tau}{2}\frac{\partial}{\partial q_i} V(q_1, q_2, \ldots, q_N),$$
$$q_i' = q_i + \tau \frac{\partial}{\partial \tilde{p}_i} K(\tilde{p}_1, \tilde{p}_2, \ldots, \tilde{p}_N),$$
$$p_i' = \tilde{p}_i - \frac{\tau}{2}\frac{\partial}{\partial q_i'} V(q_1', q_2', \ldots, q_N'),$$

which evolves the canonical variables (q_i, p_i) to (q_i', p_i') for $i = 1, \ldots, N$. This second-order scheme is similar to the one described by Ruth [**30**] and Forest and Ruth [**15**], with the roles of the coordinates and momenta interchanged. It conserves only an approximate Hamiltonian and can therefore lead to orbital drift.

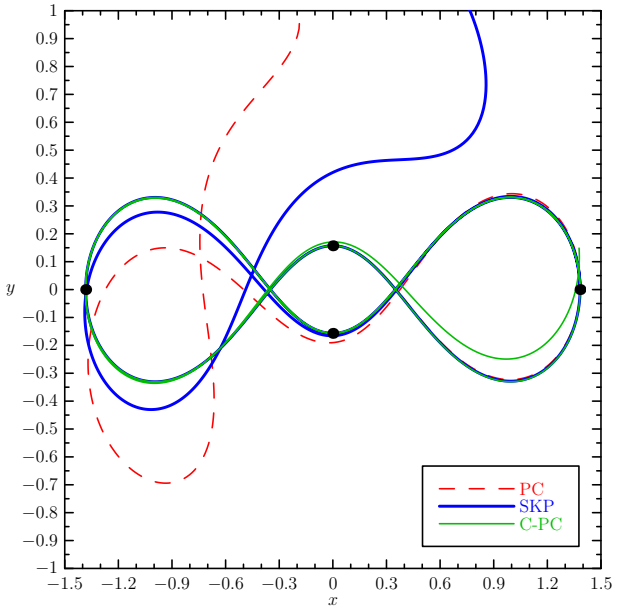

FIGURE 1: PC, symplectic SKP and C-PC solutions for a four-body choreography.

3 Exponential integrators Exponential integrators have a long history [**10, 29, 16, 35, 22, 2, 12, 20, 6**] and have been independently rediscovered many times. These methods respect the exact evolution of the system on a linear time scale and are therefore ideally suited to linearly stiff ODEs.

A typical stiff nonlinear initial value problem appears as

$$\frac{dx}{dt} + \eta\, x = f(x,t), \qquad x(0) = x_0. \tag{3}$$

where the nonlinear term f varies slowly in t compared with the value of the linear coefficient η:

$$\left|\frac{1}{f}\frac{\partial f}{\partial t}\right| \ll |\eta|\,.$$

The goal is to solve (3) on the linear time scale exactly, thereby avoiding the linear time-step restriction $\eta\tau \ll 1$. In the presence of the nonlinear term f, straightforward integrating factor and Rosenbrock methods [**19**] do not remove the explicit restriction on the linear time step τ, even if f is constant with respect to both arguments. In contrast, exponential integrators are designed to be the exact solution for this case.

Exponential integrators can be derived from the exact evolution equation

$$x(t_0+\tau) = P^{-1}(t_0+\tau)\left[x(t_0) + \int_{t_0}^{t_0+\tau} P(t) f(x,t)\,dt\right],$$

where $P(t) = e^{\eta(t-t_0)}$. On changing the independent variable from t to P, we find, noting $dt\,P = \eta^{-1} dP$, that

$$x(t_0+\tau) = P^{-1}(t_0+\tau)\left[x(t_0) + \eta^{-1}\int_1^{P(t_0+\tau)} f\,dP\right].$$

A rectangular approximation of the above integral then yields the exponential Euler algorithm (also known as the exponentially-fitted Euler method):

$$\begin{aligned} x_{i+1} &= P_\tau^{-1}\left[x_i + \eta^{-1}(P_\tau - 1) f_i\right] \\ &= P_\tau^{-1} x_i + \eta^{-1}(1 - P_\tau^{-1}) f_i, \end{aligned}$$

where $P_\tau = e^{\eta\tau}$ and τ is the time step. Note that the discretization is now with respect to $P(t)$ instead of t and that as $\eta\tau \to 0$, we recover the classical Euler scheme.

Similarly, a trapezoidal approximation of the integral in Equation (3) yields a second-order exponential predictor-corrector (E-PC) algorithm:

$$\begin{aligned} \widetilde{x}_{i+1} &= P_\tau^{-1} x_i + \eta^{-1}(1 - P_\tau^{-1}) f_i, \\ x_{i+1} &= P_\tau^{-1} x_i + \eta^{-1}(1 - P_\tau^{-1})\left(\frac{f_i + \widetilde{f}_{i+1}}{2}\right), \end{aligned}$$

where $\widetilde{f}_{i+1} = f(\widetilde{x}_{i+1})$.

The above derivation can be readily generalized by replacing the linear Green's function $e^{\eta(t-t')}$ above by any *stationary* Green's function $G(t-t')$. When the scalar variable x is replaced by a vector $\boldsymbol{x}$, the integrating factor becomes the matrix exponential $\boldsymbol{P}(t) = e^{\boldsymbol{\eta} t}$. It is interesting to note that Lagrangian discretizations of advection equations can also be written formally as exponential integrators. For the passive advection equation

$$\frac{\partial u}{\partial t} + v\frac{\partial}{\partial x}u = f(x,t,u), \qquad u(x,0) = u_0(x),$$

one sees that η represents the linear operator $v\partial/\partial x$ and $\mathcal{P}^{-1}u = e^{-vt\frac{\partial}{\partial x}}u$ corresponds to the Taylor series of $u(x-vt)$. A Lagrangian method has the advantage of not introducing any artificial numerical dissipation or interpolation of infinitesimal parcel values. A new fully Lagrangian method for passive scalar and self-advection dynamics that respects parcel rearrangement has been recently developed by Yassaei, *et al.* [**36**]; in the inviscid limit, it preserves the infinity of Casimir invariants associated with parcel rearrangement.

3.1 Numerical considerations There are two technical issues that must be addressed in the implementation of exponential integrators.

First, one must take care in evaluating the factor $(1-e^{-\eta\tau})$ to avoid catastrophic loss of precision when $|\eta\tau| < \log(1+\sqrt{\epsilon}) \approx \sqrt{\epsilon}$, where ϵ is the machine precision (the smallest positive number that when added to one in the machine representation yields a result larger than one). Fortunately, on many computers the function $f(x) = e^x - 1$ is a hardware instruction (to facilitate the accurate computation of hyperbolic trigonometric functions).

Second, if $\eta < 0$ and the nonlinear time is very long, the argument of the exponential in the inverse integrating factor $e^{-\eta\tau}$ can become so large in magnitude as to cause numerical overflow. This can be avoided simply by restricting the maximum size of the time step. (In the situation where $\eta > 0$, an underflow poses no problem if the exponential is approximated by zero.)

Coupled with a dynamically adjusted (adaptive) time step, exponential integrators have also been found to be very useful in single- and multiple-field Markovian statistical closure computations of forced-dissipative turbulence [**7, 8, 4, 23, 24**]. These algorithms can speed up such computations by many orders of magnitude, particularly in the relaxation of the triad interaction time in the dissipation range, where the linear time scale is short compared to the nonlinear time.

3.2 A charged particle in electromagnetic fields The motion of a charged particle in crossed electric and magnetic fields is an important problem in plasma physics. In magnetically confined fusion plasmas there are typically two widely separated frequencies associated with this problem, the high cyclotron frequency (typically 1 GHz) and the much lower frequency (typically 10 kHz) of oscillation of the electric field $\boldsymbol{E}$ and magnetic field $\boldsymbol{B}$. Because of this extreme separation of time scales, exponential integrators are well-suited to numerical integrations of $\boldsymbol{E}\times\boldsymbol{B}$ motion, described in cgs units by a vector equation for the particle velocity $\boldsymbol{v}$:

$$\frac{m}{q}\frac{d\boldsymbol{v}}{dt} = \frac{1}{c}\boldsymbol{v}\times\boldsymbol{B} + \boldsymbol{E}.$$

Here m and q are the particle mass and charge, respectively, and c is the speed of light.

This requires an efficient means of computing the matrix exponential $\exp(\boldsymbol{\Omega})$, where

$$\boldsymbol{\Omega} = -\frac{q}{mc}\,\tau \begin{pmatrix} 0 & B_z & -B_y \\ -B_z & 0 & B_x \\ B_y & -B_x & 0 \end{pmatrix}.$$

The computation of $\exp(\boldsymbol{\Omega})$ was optimized to 2 trigonometric functions, 1 division, 1 square root, and 35 additions or multiplications. The other necessary matrix factor, $\boldsymbol{\Omega}^{-1}[\exp(\boldsymbol{\Omega})-\mathbf{1}]$ requires care since $\boldsymbol{\Omega}$ is singular (but at least has distinct eigenvalues for $\boldsymbol{B} \neq \mathbf{0}$); it is evaluated as

$$\lim_{\lambda\to 0}[(\boldsymbol{\Omega}+\lambda\mathbf{1})^{-1}(e^{\boldsymbol{\Omega}}-\mathbf{1})].$$

In Figure 2 we plot the trajectory of an electron under the influence of the stationary magnetic field $\boldsymbol{B} = \hat{\boldsymbol{z}}$ and the time-varying electric field $\boldsymbol{E}(t) = \hat{\boldsymbol{x}}\exp(\cos(t))$, adopting the parameters $q = m = c = 1$. The initial velocity $\boldsymbol{v} = (1, 0, 1)$ was integrated forward 40 time steps of size 0.5; the trajectory was then computed from the velocity values using trapezoidal integration. We see that the E-PC algorithm performs much better than the conventional predictor-corrector (PC) algorithm, in comparison with the "exact" solution computed with a fifth-order Runge-Kutta integrator and a time step of 0.005.

For kinetic simulations of plasmas, exponential algorithms could potentially offer an alternative to gyrokinetic approximations [**14**] since they use the exact linear theory to eliminate the problem of widely-separated time-scales that gyrokinetic-averaging was developed to circumvent.

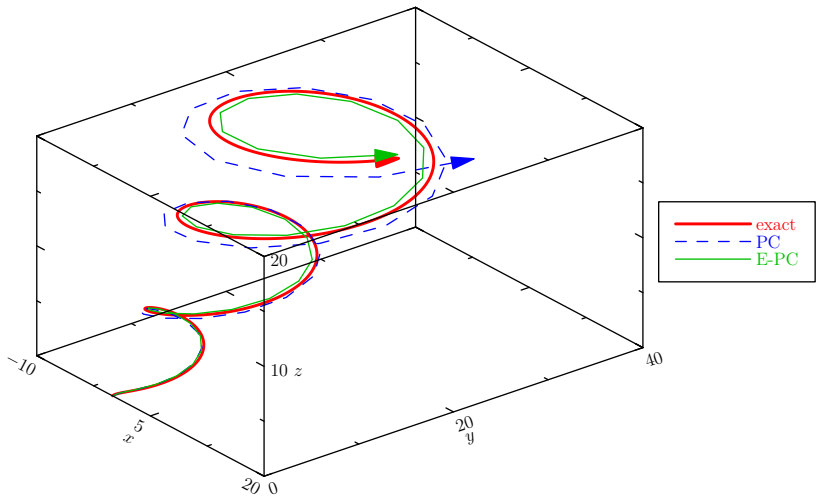

FIGURE 2: The motion of an electron under the Lorentz force as computed with the PC and E-PC integrators, in comparison with the exact solution.

3.3 Higher-order exponential integrators Higher-order exponential integrators have been discussed by [**21**], [**12**], [**20**] and [**6**]. They require the evaluation of the functions, for $n \in \mathbb{N}$,

$$\varphi_n(x) = x^{-n}\left(e^x - \sum_{k=0}^{n-1} \frac{x^k}{k!}\right).$$

Again, care must be exercised when evaluating $\varphi_1(x)$ and $\varphi_2(x)$ near 0. Accurate optimized double precision routines for evaluating these functions are available at `www.math.ualberta.ca/~bowman/phi.h`.

An adaptive exponential version of the highly efficient (3,2) embedded Bogacki-Shampine Runge-Kutta pair has been used to investigate a shell model of fluid turbulence [**6**]. Letting $z = -\eta\tau$, this four-stage method uses the coefficients $b_0 = 0$, $b_1 = 1/2$, $b_2 = 3/4$, $b_3 = 1$, and

$$a_{10} = \frac{1}{2}\varphi_1\left(\frac{1}{2}z\right),$$

$$a_{20} = \frac{3}{4}\varphi_1\left(\frac{3}{4}z\right) - a_{21}, \quad a_{21} = \frac{9}{8}\varphi_2\left(\frac{3}{4}z\right) + \frac{3}{8}\varphi_2\left(\frac{1}{2}z\right),$$

$$a_{30} = \frac{2}{3}\varphi_1(z) - a_{32}, \qquad a_{31} = \frac{1}{3}\varphi_1(z), \quad a_{32} = \frac{4}{3}\varphi_2(z) - \frac{2}{9}\varphi_1(z),$$
$$a_{40} = \varphi_1(z) - \frac{17}{12}\varphi_2(z), \quad a_{41} = \frac{1}{2}\varphi_2(z), \quad a_{42} = \frac{2}{3}\varphi_2(z), \quad a_{43} = \frac{1}{4}\varphi_2(z).$$

The value x_3 provides the desired solution, which can be shown, based on the work of [**20**], to have stiff order 3. Here *stiff order* means that the order is preserved even when η is a general unbounded linear operator. The value x_4 provides a second-order estimate that can be used to adjust the time step. Since $f(x_3)$ is just f at the initial x_0 for the next time step, no additional source evaluation is required to compute x_4. This is known as the first-same-as-last (FSAL) property. One notes that as $\eta\tau \to 0$, this exponential scheme reduces to the conventional adaptive [3,2] Bogacki-Shampine Runge-Kutta pair [**3**].

4 Conclusions Discretizations of ordinary differential equations that preserve physically relevant structure or known analytic properties are becoming of widespread interest. In this work we illustrated structure-preserving discretizations that can be obtained from conventional ones *via* transformation of the independent or dependent variables.

While traditional numerical discretizations of conservative systems generically yield artificial secular drifts of nonlinear invariants, explicit exactly conservative integrators obtained by discretizing the dependent variables in a space in which the conserved quantity is linear can outperform conventional (and in at least some cases, symplectic) integrators. Moreover, the transformation technique here is relevant to integrable and nonintegrable Hamiltonian systems and even to non-Hamiltonian systems such as force-dissipative turbulence. Likewise, exponential integrators are known to outperform conventional integrators when applied to linearly stiff systems. By combining these ideas, it is even possible to develop conservative exponential integrators. The underlying theme behind each of these methods is the goal of developing numerical discretizations that are aware of underlying analytical structure.

The author would like to acknowledge discussions with R. Sydora and thank the Natural Sciences and Engineering Research Council of Canada for financial support and Andrew Hammerlindl and Tom Prince for their contributions to `Asymptote` (`asymptote.sf.net`), the descriptive vector graphics language that was used to produce the figures in this work.

REFERENCES

1. A. Bayliss and E. Isaacson, *How to make your algorithm conservative,* Not. Amer. Math. Soc. **22** (1975), A594–A595.
2. G. Beylkin, J. M. Keiser and L. Vozovoi, *A new class of time discretization schemes for the solution of nonlinear PDEs,* J. Comp. Phys. **147** (1998), 362–387.
3. P. Bogacki and L. F. Shampine, *A 3(2) pair of Runge-Kutta formulas,* Appl. Math. Letters **2** (1989), 1–9.
4. J. C. Bowman and J. A. Krommes, *The realizable Markovian closure and realizable test-field model. II: Application to anisotropic drift-wave dynamics,* Phys. Plasmas **4** (1997), 3895–3909.
5. J. C. Bowman, J. A. Krommes and M. Ottaviani, *The realizable Markovian closure. I: General theory, with application to three-wave dynamics,* Phys. Fluids B **5** (1993), 3558–3589.
6. J. C. Bowman, C. R. Doering, B. Eckhardt, J. Davoudi, M. Roberts and J. Schumacher, *Links between dissipation, intermittency, and helicity in the goy model revisited,* Physica D **218** (2006), 1–10.
7. J. C. Bowman, *On inertial-range scaling laws,* J. Fluid Mech. **306** (1996), 167–181.
8. J. C. Bowman, *A wavenumber partitioning scheme for two-dimensional statistical closures,* J. Sci. Comput. **11** (1996), 343–372.
9. V. Brasey and Hairer, *Symmetrized half-explicit methods for constrained mechanical systems,* Appl. Numer. Math. **13** (1993), 23–31.
10. J. Certaine, *The solution of ordinary differential equations with large time constants,* Math. Meth. Dig. Comp. (1960), 129–132.
11. P. J. Channell and J. C. Scovel, *Symplectic integration of Hamiltonian systems,* Nonlinearity **3** (1990), 231–259.
12. S. Cox and P. Matthews, *Exponential time differencing for stiff systems,* J. Comp. Phys. **176** (2002), 430–455.
13. J. de Frutos and J. M. Sanz-Serna, *Erring and being conservative,* in: Numerical Analysis (1993) (D. F. Griffiths and G. A. Watson, eds.), Pitman Research Notes in Mathematics, Longmans Scientific and Technical, Harlow, 1994, 75–88.
14. D. H. E. Dubin, J. A. Krommes, C. Oberman and W. W. Lee, Phys. Fluids **26** (1983), 3524.
15. E. Forest and R. D. Ruth, *Fourth-order symplectic integration,* Physica D **43** (1990), 105–117.
16. A. Friedli, *Verallgemeinerte Runge-Kutta Verfahren zur Lösung steifer Differentialgleichungen,* Lecture Notes in Mathematics (1978), 214–219.
17. Ge Zhong and J. E. Marsden, *Lie-Poisson Hamilton-Jacobi theory and Lie-Poisson integrators,* Phys. Lett. A **133** (1988), 134–139.
18. C. W. Gear, *Maintaining solution invariants in the numerical solution of ODEs,* SIAM J. Sci. Stat. Comput. **7** (1986), 734–743.
19. E. Hairer, S. Nørsett and G. Wanner, *Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems*, Vol 14, Springer Series in Comput. Math., Springer, Berlin, 1991.
20. M. Hochbruck and A. Ostermann, *Explicit exponential Runge-Kutta methods for semilinear parabolic problems,* SIAM J. Numer. Anal. **43** (2005), 1069–1090.
21. M. Hochbruck, C. Lubich and H. Selhofer, *Exponential integrators for large systems of differential equations,* SIAM **19** (1997), 1552–1574.
22. M. Hochbruck, C. Lubich and H. Selfhofer, *Explicit integrators for large systems of differential equations,* SIAM J. Sci. Comput. **19** (1998), 1552–1574.

23. G. Hu, J. A. Krommes and J. C. Bowman, *Resistive drift-wave plasma turbulence and the realizable Markovian closure,* Phys. Lett. A **202** (1995), 117–125.
24. G. Hu, J. A. Krommes and J. C. Bowman, *Statistical theory of resistive drift-wave turbulence and transport,* Phys. Plasmas **4** (1997), 2116–2133.
25. E. Isaacson, *Integration schemes for long term calculation,* in: Advances in Computer Methods for Partial Differential Equations II (R. Vichnevetsky, ed.), I.M.A.C.S (AICA), 1977, 251–255.
26. A. Iserles, *Insight, not just numbers,* in: The 15th IMACS World Congress on Scientific Computation, Modelling and Applied Mathematics Vol **2** (A. Sydow, ed.), Wissenschaft and Technik Verlag, Berlin, 1997, 1–9.
27. C. Kane, J. E. Marsden and M. Ortiz, *Symplectic-energy-momentum preserving variational integrators,* J. Math. Phys. **40** (1999), 3353–3371.
28. O. Kotovych and J. C. Bowman, *An exactly conservative integrator for the n-body problem,* J. Phys. A.: Math. Gen. **35** (2002), 7849–7863.
29. S. Nøsett, *An A-stable modification of the Adams-Bashforth methods,* Lecture Notes in Mathematics **109** (1969), 214–219.
30. R. D. Ruth, *A canonical integration technique,* IEEE Trans. Nucl. Sci. **30** (1983), 2669–2671.
31. J. M. Sanz-Serna and M. P. Calvo, *Numerical Hamiltonian Problems*, Vol 7, Applied Mathematics and Mathematical Computation, Chapman and Hall, London, 1994.
32. B. A. Shadwick and W. F. Buell, *Unitary integration: A numerical technique preserving the structure of the quantum Liouville equation,* Phys. Rev. Lett. **79** (1997), 5189–5193.
33. B. A. Shadwick, J. C. Bowman and P. J. Morrison, *Exactly conservative integrators,* SIAM J. Appl. Math. **59** (1999), 1112–1133.
34. B. A. Shadwick, W. F. Buell and J. C. Bowman, *Structure preserving integration algorithms, in: Scientific Computing and Applications* Vol **7** (P. Minev, Y. S. Wong and Y. Lin, eds.), Advances in Computation: Theory and Practice, Nova Science Publishers, New York, 2001, 247–255.
35. J. Verwer, *On generalized linear multistep methods with zero-parasitic roots and an adaptive principal root,* Numer. Math. **27** (1977), 143–155.
36. M. A. Yassaei, J. C. Bowman and A. Basu, *A fully Lagrangian advection scheme,* submitted, J. Comp. Phys., 2007.

Department of Mathematical and Statistical Sciences,
University of Alberta, Edmonton, Alberta, Canada T6G 2G1
E-mail address: bowman@math.ualberta.ca