The accuracy of symplectic integrators

Robert I. McLachlan and Pau Atela

Campus Box 526, Program in Applied Mathematics, University of Colorado at Boulder, Boulder, CO 80309

We judge symplectic integrators by the accuracy with which they represent the Hamiltonian function. This accuracy is computed, compared and tested for several different methods. We develop new, highly accurate explicit fourth- and fifth-order methods valid when the Hamiltonian is separable with quadratic kinetic energy. For the near-integrable case, we confirm several of their properties expected from KAM theory; convergence of some of the characteristics of chaotic motions are also demonstrated. We point out cases in which long-time stability is intrinsically lost.

0. Introduction

The symplectic integration of Hamiltonian dynamical systems is by now an established technique. Ruth [1] has developed explicit methods for separable systems; his approach was extended to fourth order by Candy and Rozmus [2]. Channel and Scovel [3] and Feng et al. [4] have derived methods based on the Taylor series expansion of the time map of a general Hamiltonian. Feng [5] contains a survey of the Chinese program and an important generalisation which includes many known methods as special cases. In [6] Feng provides a discussion of the philosophy and history of symplectic integration.

For a summary of explicit symplectic integrators for separable Hamiltonians, see (2.2) and Table II.

Standard integrators do not generally preserve the Poincaré integral invariants of a Hamiltonian flow and cannot hope to capture the long-time dynamics of the system. Typically their numerical diffusion causes orbits to be attracted to elliptic orbits, or, coupled with forcing, creates completely unphysical attractors. These may look similar to the chaotic dynamics of some systems; however, the non-conservation of integral invariants presumably corrupts the long-time statistics of the flow. These failings have been discussed at length [2,3,5] and we will not consider non-symplectic methods further.

However, within the class of symplectic integrators (SIs) there is much variation in easy of use and in accuracy. It is true that the largest-scale structures in phase space, such as low-order resonances, tend to be the most persistent under the symplectic perturbations due to the discretisation. They will be captured by even a low-order SI with a fairly large time-step, as will the regular behaviour of near-integrable systems. But if quantitative results are required, or finer features must be resolved, or extremely slow scale phenomena are being studied, then the convergence of the dynamics of the discrete system to those of the continuous system must be studied. In this respect, not all symplectic integrators are created equal.

We wish to numerically integrate the flow of the Hamiltonian $H(\mathbf{q},\mathbf{p})$ with N degrees of freedom and time step k. We use the fact that any family of symplectic maps $\hat{\phi}(k,\mathbf{q},\mathbf{p})$ near the identity, depending smoothly on a parameter k, is locally the exact time-k map of a Hamiltonian flow, say that due to $\hat{H}(\mathbf{q},\mathbf{p},t)$. ($\hat{H}$ also depends on k). If, in addition, the map is close to the time-k map of the Hamiltonian $H(\mathbf{q},\mathbf{p})$, then $\hat{H}$ is close to H, and we shall identify $\| H-\hat{H} \|$ with the accuracy of the discrete map. (The choice of norm will be discussed later). This point of view is elaborated in §1.

Amongst methods of a given order we shall use this criterion to distinguish the "best". (Of course this can only be done in a general sense—for a specific problem there may always be specially-tailored methods). Identifying good methods is perhaps even more important in the symplectic than in the non-symplectic case, for SIs are usually used with fairly large time steps. In some cases this has led to higher-order methods being disadvantageous, because of their large error constants.

In the evolution of the real flow, the energy $H(\mathbf{q},\mathbf{p})$ is conserved. This will not be true in general for the flow $\hat\phi$ of $\hat H$. In fact, if H is conserved by $\hat\phi$ and if the original system has no other independent integrals, then the approximate and real flows are the same (up to a reparameterisation of time) (Ge Zhong [7]). We would thus have solved the equations of motion. So, one might consider minimising some norm of the *energy* truncation error $\nabla H\cdot\left(\phi(k,\mathbf{q},\mathbf{p})-\hat\phi(k,\mathbf{q},\mathbf{p})\right)$. In $N\geq 2$ degrees of freedom, an orbit of $\hat\phi$ can drift away from the original energy surface to a region of phase space with very different dynamics; one wishes to minimise this effect. This criterion does lead to good methods, but it has two drawbacks. It only considers the motion in one direction, that transverse to the energy surface; and it involves expanding the time-step to one higher order than needed for $\| H-\hat H\|$ (which is already complicated enough). Considering, say, the component of the truncation error not in the direction of the real flow would avoid the first problem but would be even more complicated. Because $\|H-\hat H\|$ must already be computed (see §2), finding the Hamiltonian error term is fairly simple; in addition we can hope that it captures more information about the global dynamics than the energy truncation error. The important point is to select a *well-defined*, relevant quantity for systematic comparisons.

In simple mechanical systems the Hamiltonian is often separable: $H(\mathbf{q},\mathbf{p})=T(\mathbf{p})+V(\mathbf{q})$. For this case Ruth [1] developed a class of explicit SI methods. (See Feng [5] for a wider class of systems for which explicit algorithms exist). These were extended to fourth order by Ruth in 1984 (unpublished), by Candy and Rozmus [2], and in an elegant formulation by Forest and Ruth [15]. In §2, we find the "best" methods of this class; we also deal with the important special case in which the kinetic energy is quadratic, i.e., $T(\mathbf{p})=\frac{1}{2}\mathbf{p}^T\mathbf{M}\mathbf{p}$. This leads to crucial simplifications that allow one to find much more accurate methods. There turns out to be a 4-stage, fourth-order method with is about 100 times more accurate than the Candy and Rozmus method; it is even more accurate than the fourth-order Gauss-Legendre (which is itself in some sense optimal for general Hamiltonians). We also develop an optimal 6-stage, 5th order method. These two methods are recommended for general use when H is separable with quadratic kinetic energy.

In §2.3 we consider a wider class of methods in which the derivatives of V higher than the first are also supplied, which leads to much simpler equations to be solved. In §3, we compute the error constants for three well-known symplectic integrators for general (non-separable) Hamiltonians. For these implicit methods, which must be solved by iteration, efficiency is strongly affected by the choice of predictor; so a method cannot be judged by its error constant alone.

There are three main numerical examples in §4. First the accuracy and efficiency of the "optimal" methods is tested in simple one and two degree-of-freedom problems. We also examine the convergence of some properties of a very fine chaotic web; this could not be studied at all without a symplectic integrator. Finally, we give an example of some intrinsic limitations of SIs in $N\geq 2$ degrees of freedom.

1. A framework for symplectic integration

We consider the class of symplectic integrators defined by smooth one-parameter families of maps close to the identity:

$$\mathbf{X}=\hat\phi(t,\mathbf{x}),\qquad \hat\phi(0,\mathbf{x})=\mathbf{x},$$

where $\mathbf{x}=(\mathbf{q},\mathbf{p})^T$ and $\mathbf{X}=(\mathbf{Q},\mathbf{P})^T$ are in $\mathbb{R}^{2N}$. With initial condition $\mathbf{x}_0$, the discrete orbit is given by the forward iterates $\hat\phi^{(n)}(k,\mathbf{x}_0)$, where k is the time-step. The map $\hat\phi$ is then generated by a time-dependent generating function of the second kind, say $K(\mathbf{q},\mathbf{P},t)$. (Arnol'd [8], pp. 242, 266–269). We will see now that an iterate $\hat\phi^{(n)}(k,\mathbf{x}_0)$ is the exact time-nk map of a $k-$periodic, time-dependent, time-discontinuous Hamiltonian $\hat H$. Think of K as generating a change of variables ($\hat\phi$) between two families of canonical coordinates systems, $\mathbf{x}(t)$ and $\mathbf{X}(t)$. We let $\mathbf{x}$ evolve with Hamiltonian identically 0, i.e. $\mathbf{x}(k) = \mathbf{x}(0)$.

$$
\begin{array}{ccccccc}
 & & & \overset{K(\mathbf{x},0)}{\hat\phi(0)=id} & & & \\
t=0 & & \mathbf{x}(0) & \longrightarrow & \mathbf{X}(0) & & \\
 & 0 & \downarrow & & \downarrow & \hat H & \\
t=k & & \mathbf{x}(k) & \longrightarrow & \mathbf{X}(k) & & \\
 & & & \underset{K(\mathbf{x},k)}{\hat\phi} & & &
\end{array}
$$

$\mathbf{X}$ is therefore evolving under the flow of the Hamiltonian

$$\hat H(\mathbf{X}, t) = 0 + \frac{\partial K}{\partial t} = \frac{\partial K}{\partial t}.$$

We take this as the definition of $\hat H$ for $0 \le t < k$. We let $\mathbf{X}$ evolve until $t = k$, and $\mathbf{X}(k)$ then becomes the initial condition for the next time-step. We extend the definition of $\hat H$ to $t \ge k$ by

$$\hat H(\mathbf{X}, nk + t) \equiv \hat H(\mathbf{X}, t) \qquad 0 \le t < k,\ n > 1.$$

The flows of H and $\hat H$ are sketched in Figure 1. The iterates $\hat\phi^{(n)}$ are now the *exact* time-nk map of the Hamiltonian $\hat H$.

Let $H(\mathbf{x})$ be the Hamiltonian we wish to integrate. We say the integrator is *consistent with the Hamiltonian H* if $\hat H(\mathbf{x}, 0) = H(\mathbf{x})$. Define the *Hamiltonian truncation error* as

$$E = H - \hat H.$$

This is sketched in Figure 1. In effect, the integrator calculates the exact flow of the nearby Hamiltonian $\hat H$. It is easy to check that if $\| E \| = \mathcal{O}(k^p)$ then the truncation error of the integrator is $\mathcal{O}(k^{p+1})$ and the method is of order p.

As we have two nearby Hamiltonians, H and $\hat H$, one would like to apply KAM theory to compare their flows; but in this case $\hat H$ has $\mathcal{O}(k^p)$ jump discontinuities at $t = nk$ so the theory does not apply. Instead one compares the nearby *maps* $\phi(k, \mathbf{x})$ and $\hat\phi(k, \mathbf{x})$ to establish which dynamics are preserved by the integrator. However, this formulation offers a parallel view and makes explicit the Hamiltonian truncation error which we would like to minimise.

This homotopy in k is not the only way to construct an $\hat H$. In fact, one can enforce any degree of smoothness required at $t = nk$, although this may increase $\frac{\partial \hat H}{\partial t}$. Such constructions proceed term by term in k, leading to questions of convergence; in contrast, our $\hat H$ clearly exists and is spatially smooth for all $\mathbf{x}$ and t. Similarly, one may construct an *autonomous* $\hat H$, matching $\mathbf{x}(k)$ term by term in k—see e.g. Auerbach and Friedman [18] for the one-degree-of-freedom case. If $\hat H$ converges it is a conserved quantity for the discrete system and describes those invariant tori which persist. However, the significance of this $\hat H$ in higher dimensions, or of considering the first term of a merely asymptotic series, is unclear. The secular energy drift seen in §4.3 for $N = 2$ indicates that the discrete system does not have a conserved quantity close to H and is really evolving in six, not four, dimensions.

2. Explicit methods for separable Hamiltonians

When the Hamiltonian is of the form $H(\mathbf{q}, \mathbf{p}) = T(\mathbf{p}) + V(\mathbf{q})$, Ruth [1] noticed that $K(\mathbf{q}_0, \mathbf{p}_1) = \mathbf{q}_0\mathbf{p}_1 + kH(\mathbf{q}_0, \mathbf{p}_1)$ is consistent with H and that the coordinate transformations it generates can be found explicitly. To extend the method to higher order, he applied several such transformations in turn:

$$(\mathbf{q}_0, \mathbf{p}_0) \overset{K_1(\mathbf{q}_0, \mathbf{p}_1)}{\longrightarrow} (\mathbf{q}_1, \mathbf{p}_1) \longrightarrow \cdots \overset{K_l(\mathbf{q}_{l-1}, \mathbf{p}_l)}{\longrightarrow} (\mathbf{q}_l, \mathbf{p}_l) \tag{2.1}$$

where

$$K_i = \mathbf{q}_{i-1}\mathbf{p}_i + k\big(a_i T(\mathbf{p}_i) + b_i V(\mathbf{q}_{i-1})\big),$$

a_i and b_i are parameters and k is the time step. $(\mathbf{q}_1, \mathbf{p}_1), \ldots, (\mathbf{q}_{l-1}, \mathbf{p}_{l-1})$ are intermediate stages as in standard Runge-Kutta methods. Recall that K_i is a generating function of the second kind which generates a change of coordinates according to

$$\mathbf{p}_{i-1} = \mathbf{p}_i + \frac{\partial K_i}{\partial \mathbf{q}_{i-1}}, \qquad \mathbf{q}_i = \mathbf{q}_{i-1} + \frac{\partial K_i}{\partial \mathbf{p}_i}.$$

(It may be checked that $l = a_1 = b_1 = 1$ generates a first-order symplectic map; applied to the pendulum, this gives the standard map). In the framework of §1, we have

$$\begin{array}{cccccc} & & & \overset{K(\mathbf{x},0)}{\underset{\hat{\phi}(0)=id}{}} & & \\ t=0 & & \mathbf{x}(0) & \longrightarrow & \mathbf{X}(0) & \\ & H_0 & \downarrow & & \downarrow & H \\ t=k & & \mathbf{x}(k) & \longrightarrow & \mathbf{X}(k) & \\ & & & \overset{\hat{\phi}}{\underset{K(\mathbf{x},k)}{}} & & \end{array}$$

with $\mathbf{x} = (\mathbf{q}_0\, \mathbf{p}_0)^T$, $\mathbf{X} = (\mathbf{q}_l\, \mathbf{p}_l)^T$, and K generating the transformation $\mathbf{x} \mapsto \mathbf{X}$. That is, integrating $\mathbf{x}$ under H_0 and transforming gives the exact flow of H; integrating 0 instead (so that $\mathbf{x}(0) = \mathbf{x}(k)$) and transforming gives the flow of $\hat{H}$. Since $H = H_0 + \frac{\partial K}{\partial t} = H_0 + \hat{H}$, we have $E = H_0$. The aim is to choose the a_i's and b_i's to make H_0 zero up to some order in k.

Let the force $\mathbf{F} = -V'$ and $\mathbf{P} = T'$ denote the gradients. The whole sequence of intermediate values is then given for $i = 1, \ldots, l$ by

$$\begin{aligned} \mathbf{p}_i &= \mathbf{p}_0 + k \sum_{m=1}^{i} b_m \mathbf{F}(\mathbf{q}_{m-1}) \\ \mathbf{q}_i &= \mathbf{q}_0 + k \sum_{m=1}^{i} a_m \mathbf{P}(\mathbf{p}_m) \end{aligned} \tag{2.2}$$

with transformed Hamiltonian

$$\begin{aligned} H_0(\mathbf{q}_0, \mathbf{p}_0) &= T(\mathbf{p}_l) + V(\mathbf{q}_l) - \sum_{i=1}^{l} \big(a_i T(\mathbf{p}_i) + b_i V(\mathbf{q}_{i-1})\big) \\ &\equiv \sum_{m=0}^{\infty} h_m(\mathbf{q}_0, \mathbf{p}_0) t^m \end{aligned}$$

where we have made a formal power series expansion of $\mathbf{q}_i$, $\mathbf{p}_i$, $T(\mathbf{p}_i)$, and $V(\mathbf{q}_i)$ about $(\mathbf{q}_0, \mathbf{p}_0)$ for $i = 1, \ldots, l$. For a method of order p we must choose $\mathbf{a}$ and $\mathbf{b}$ so that $h_m \equiv 0$ for $m = 0, \ldots, p-1$. We then call h_p the *principal error function.*

It is straightforward to check that h_m for $N = 1$ and for $N > 1$ are identical for $m \leq 5$. We therefore expanded h_m with $N = 1$ in Mathematica. Table I shows the number of equations which must be satisfied at each order to enforce $h_m \equiv 0$ (see also Sanz-Serna et al. [21], who have developed a general theory which allows one to calculate Table I for all m). For example, at third order H_0 has terms proportional to $\mathbf{P}_i\mathbf{P}_j\mathbf{P}_k\mathbf{F}''_{ijk}$, $\mathbf{F}_i\mathbf{F}'_{ij}\mathbf{P}_k\mathbf{P}'_{jk}$, and $\mathbf{F}_i\mathbf{F}_j\mathbf{F}_k\mathbf{P}''_{ijk}$. The combinatorial explosion as m increases is familiar from the theory of standard Runge-Kutta methods; but in that case there are also many more unknowns available. If the number of unknowns ($2l$) is greater than the number of equations (M) then the real solutions, if there are any, will lie on manifolds of dimension $2l - M$, assuming that all the equations are independent. We propose to use this freedom to minimise the principal error function and obtain more accurate methods.

Example: Second order methods.

Here $p = l = 2$ and the transformed Hamiltonian can be written (after some algebra)

$$\begin{aligned} H_0(\mathbf{q},\mathbf{p}) = {} & T(\mathbf{p})(1-a_1-a_2) + V(\mathbf{q})(1-b_1-b_2) \\ & + t\,[FP(2a_1b_2-1)] \\ & + t^2\left[F'P^2\tfrac{1}{2}(3a_1^2b_2-1) + F^2P'(3a_1b_2(1-\tfrac{1}{2}b_2)-1)\right] \\ & + \mathcal{O}(t^3) \end{aligned}$$

There is thus a one-parameter family of second-order methods given by

$$a_1 = 1-a_2, \qquad b_2 = \frac{1}{2(1-a_2)}, \qquad b_1 = 1-b_2.$$

The leapfrog and pseudo-leapfrog schemes [1,2] are the cases $a_2 = \frac{1}{2}$ and $a_2 = 0$, respectively. The principal error function h_2, the coefficient of t^2 in $E = H_0$, is

$$F'P^2\tfrac{1}{4}(1-3a_2) + F^2P'\frac{1-4a_2}{8(1-a_2)}.$$

As each function of $\mathbf{q}_0$ and $\mathbf{p}_0$ is an elementary differential of the same order $(p+1)$, we give them equal weight and minimise the sum of the squares of their coefficients. There are two minima, $a_2 = 1 \pm \frac{1}{2}\sqrt{2}$, and it is found that the smaller one has the smaller error. We may summarise this class of methods as follows:

Leapfrog:	$a_2 = \frac{1}{2}$,	$h_2 = F'P^2(-\frac{1}{8}) + F^2P'(-\frac{1}{4})$
Pseudo-leapfrog:	$a_2 = 0$,	$h_2 = F'P^2(\frac{1}{4}) + F^2P'(\frac{1}{8})(2.2)$
Optimal 2-stage:	$a_2 = 1-\frac{1}{2}\sqrt{2}$,	$h_2 = F'P^2\frac{1}{8}(3\sqrt{2}-4) - F^2P'\frac{1}{8}(3\sqrt{2}-4)$
		$\left(\frac{1}{8}(3\sqrt{2}-4) \sim 0.03\right)$

However, the leapfrog method has $b_1 = 0$ and so uses one instead of two force evaluations per time-step, usually the most expensive part of the computation; for a fair comparison we should therefore halve its step-size, giving an increase in accuracy of a factor of 4. These considerations give rise to the following

Definitions: Consider an explicit symplectic integrator of order p which is the exact time-k map of the Hamiltonian $\hat{H}$. The *Hamiltonian truncation error* is $E(t) = H - \hat{H}$. The *principal error function* is the coefficient of t^p in $H - \hat{H}$, which we write as

$$h_p = \sum_{j=1}^{J} f_j(\mathbf{a},\mathbf{b})g_j(\mathbf{F},\mathbf{P})|_{(\mathbf{q}_0,\mathbf{p}_0)}.$$

The *error constant* is

$$\Big(\sum_{j=1}^{J} f_j^2\Big)^{\frac{1}{2}}$$

If the method requires s force evaluations per time step, the *effective error constant* is

$$\left(\frac{s}{p}\right)^p \Big(\sum_{j=1}^{J} f_j^2\Big)^{\frac{1}{2}}$$

Amongst methods of the same class, we shall call the one with the smallest effective error constant *optimal.* Different norms of (f_j) would give different 'optimal' methods, but the Euclidean norm is easiest to minimise and gives good general-purpose methods.

The effective error constants for the above three methods are 0.070, 0.280, and 0.043, respectively. The optimal method shows a modest 39% improvement over simple leapfrog. In this case, the optimal method also minimised the energy truncation error.

We note in passing that in the literature, the name "leapfrog" is also applied to another, distinctly different method. Multi-level schemes of the form $\mathbf{x}_2 = \mathbf{x}_0 + 2kJ\nabla H(\mathbf{x}_1)$ are also called leapfrog and can be shown to be symplectic in the extended space $\mathbb{R}^{4N}$ with coordinates $(\mathbf{x}_0, \mathbf{x}_1)$ [4,16]. However, they are *not* symplectic in the usual sense. Orbits may lie on invariant tori in the extended space, but in general their projection to the physical phase space is not close to a real orbit. This class of methods seems to have little advantage over standard non-symplectic integrators.

We now apply our approach to higher-order methods.

2.1 Methods for separable Hamiltonians

Table I gives the number of equations which must be satisfied to get methods of different orders. For a third order method there are five equations. We therefore expect to need at least three stages. With $l = 3$, solutions lie on curves, and we find two such disjoint solution paths. On one path the minimum error constant is 0.267. The solution found by Ruth,

$$a_1 = \tfrac{2}{3}, \quad a_2 = -\tfrac{2}{3}, \quad a_3 = 1, \quad b_1 = \tfrac{7}{24}, \quad b_2 = \tfrac{3}{4}, \quad b_3 = -\tfrac{1}{24},$$

lies on the other path. The coefficients appearing in the principal error function h_3 are $-\frac{1}{54}$, $-\frac{1}{9}$, and $-\frac{5}{576}$, giving an error constant of 0.113. Although the solution curve can be found explicitly, this is not very instructive; instead we minimise the error constant numerically. After observing that the minimum possesses the symmetry $a_i = b_{4-i}$, it is easy to calculate the required solution: a_3 is the positive real root of $12a^4 - 24a^2 + 16a - 3 = 0$. The solution is given in Table II. The error constant is 0.058, about half that of the Ruth solution. (This solution also appears in Sanz-Serna [19], who *required* the symmetry; it is now seen to be the optimal third-order method.) Unfortunately neither solution curve intersects any of the coordinate planes $b_i = 0$.

For a fourth order method there are 8 equations to be solved. With 4 stages, Candy and Rozmus give two solutions, one of which is shown in Table II; the other is its symmetric partner ($\hat{a}_i = b_{5-i}$, $\hat{b}_i = a_{5-i}$). We calculate their error constant to be 0.311; the solution with $b_1 = 0$ has an effective error constant of 0.098. By Bézout's theorem, the maximum number of solutions of a system of polynomial equations is the total degree of the system, here $1^2 \cdot 2 \cdot 3^2 \cdot 4^3 = 1152$; but we have not carried out the global homotopy necessary to establish the actual number of real solutions [9]. However, we have not found any solutions in addition to the two given by Candy and Rozmus. There may exist better 5-stage schemes.

Yoshida [17] derives several sixth- and eighth-order methods by composing several steps of leap-frog. We find that his sixth-order methods ('fully symmetric', A, B, and C) have effective error constants of 5.39, 0.063, 1.42, and 1.35; it seems possible that his eight-stage method A is the optimal one.

2.2 Methods for quadratic kinetic energy

When $T(\mathbf{p}) = \frac{1}{2}\mathbf{p}^T\mathbf{M}\mathbf{p}$ is quadratic (the 'Runge-Kutta-Nyström' case) $\mathbf{P}''$ is identically zero. The third term of h_3 contains a factor $\mathbf{P}''$, and is automatically zero in this case. So with 4 stages, we have 7 equations in 8 unknowns (say $\mathbf{c}(\mathbf{a},\mathbf{b}) = \mathbf{0}$) and solutions lie on curves. We have not systematically traced the structure of these solutions in $\mathbb{R}^8$; instead we found roots by picking random initial points at which $\|\,\mathbf{c}\,\|_2$ was small and then minimising this norm. (The roots had too small a basin of attraction for Newton's method to be useful). We then numerically minimised the error constant starting from promising members of this set of solutions. Finally we solved $\mathbf{c}(\mathbf{a},\mathbf{b}) = \mathbf{0}$ in quadruple precision, starting from the approximately optimal solution. This gave the solution shown in Table II, with the remarkably small error constant 0.0025. Tests in §4.1 confirm the accuracy of this method. We recommend it for general use on this class of problems.

At higher order the gains are even more dramatic: h_4 now has 4 terms instead of 6, and h_5 has 5 instead of 10. We thus searched for 6-stage methods of order 5. In this case there are many solutions with equally small error constants located in different regions of parameter space. Amongst these we give in Table II the solution with the smallest $\max_i |b_i|$. The grounds for this heuristic are that truncated Taylor series tend to be more accurate close to the point of expansion. We wish to minimise the error due to the higher-order terms (which we have ignored until now) and which may contribute at the moderate values of the time step used in practice.

2.3 Methods using derivatives

If one composes s stages of the generating function (2.1) and expands about the point $(\mathbf{q}_0, \mathbf{p}_s)$ instead of $(\mathbf{q}_0, \mathbf{p}_0)$, the same terms appear in the transformed Hamiltonian as in §2.1, although with different coefficients. We may therefore generate the transformation $(\mathbf{q}_i, \mathbf{p}_i) \mapsto (\mathbf{q}_{i+1}, \mathbf{p}_{i+1})$ by the more general generating function with all these terms included. (In fact, Channel & Scovel [3] and Feng [4] derive 1-stage, arbitrary order methods by including sufficient terms in the generating function).

Here we include only those terms which allow $\mathbf{p}_i$ to be calculated explicitly, and take, for example,

$$K_i(\mathbf{q}_{i-1}, \mathbf{p}_i) = \mathbf{qp} + \left[k(a_i T + b_i V) + k^2 c_i \mathbf{F}.\mathbf{P} + k^3 d_i \mathbf{F}_j \mathbf{F}_k \mathbf{P}'_{jk} + k^4 e_i \mathbf{F}_j \mathbf{F}'_{jk} \mathbf{P}_l \mathbf{P}'_{lk}\right].$$

With $T = \frac{1}{2}\mathbf{p}^T \mathbf{M} \mathbf{p}$, this leads to the coordinate transformations

$$\begin{aligned} \mathbf{p}_i &= \left[I + k^2 c_i \mathbf{F}'\mathbf{M} + k^4 e_i (\mathbf{F}''\mathbf{F} + \mathbf{F}'^2)\mathbf{M}^2\right]^{-1} \left[\mathbf{p}_{i-1} + k b_i \mathbf{F} - k^3 d_i 2(\mathbf{F}\mathbf{F}'\mathbf{M})\mathbf{F}\right] \\ \mathbf{q}_i &= \mathbf{q}_{i-1} + k a_i \mathbf{M} \mathbf{p}_i + k^2 \mathbf{M}\mathbf{F} + k^4 e_i (\mathbf{M}^2 \mathbf{F}')\mathbf{F} \end{aligned}$$

where $\mathbf{F} = \mathbf{F}(\mathbf{q}_i)$, using matrix multiplication instead of tensor notation. We thus have 5 parameters per stage and can find methods with fewer stages than above (incidentally simplifying the algebra considerably). Unfortunately, because some of the terms must be excluded as they lead to implicit equations, not all the unknowns appear in every equation. This leads to less freedom in the solutions; the net result is that the methods have larger error constants than the number of parameters suggests. With $e_i \equiv 0$ (so that only first derivatives of $\mathbf{F}$ are required), the optimal two-stage, fourth-order method has an error constant of 0.018; allowing higher derivatives, the optimal two-stage, fourth-order method has an error constant of 0.004. Thus these methods are unlikely to be competitive with the "force-only" methods described above, except possibly in problems with a polynomial potential, in which case the derivative evaluations are extremely cheap.

3. Methods for general Hamiltonians

For general, non-separable Hamiltonians, no explicit methods are known. Here we apply our error criterion to compare three implicit methods.

Feng [4] and Channel and Scovel [3] recursively construct the generating function which generates the time-k map of the exact Hamiltonian with N degrees of freedom:

$$K(\mathbf{q}_0, \mathbf{p}_1, k) = \sum_{i=0}^{\infty} K_i(\mathbf{q}_0, \mathbf{p}_1) k^i.$$

For a method of order p, one includes terms of order up to k^p. Channel and Scovel give K_i, $i \leq 6$ explicitly, enabling methods of order 6 or less to be found. We call this the *Taylor series* method. It requires the first p derivatives of H and the formulae are lengthy. At each time-step, one solves N nonlinear equations for $\mathbf{p}_1$; $\mathbf{q}_1$ is then given explicitly.

We immediately have that

$$E = H - \hat{H} = \frac{\partial}{\partial t} \sum_{i=p+1}^{\infty} K_i(\mathbf{q}_0, \mathbf{p}_1) t^i = (p+1) K_{p+1}(\mathbf{q}_0, \mathbf{p}_0) t^p + \mathcal{O}(t^{p+1})$$

for $0 \leq t < k$. For example,

$$K_3 = \tfrac{1}{6}\left(H_{\mathbf{pp}} H_{\mathbf{q}}^2 + H_{\mathbf{pq}}^2 H_{\mathbf{p}} H_{\mathbf{q}} + H_{\mathbf{qq}} H_{\mathbf{p}}^2\right)$$

(cf. eqs. (2.2)), so the second-order method has an error constant of $\frac{\sqrt{3}}{2} \sim 0.87$, or $\frac{\sqrt{2}}{2} \sim 0.71$ if H is separable. (We shall include the separable case here for comparison, although these methods would not be used in practice in that case).

For higher orders we find the error constants listed in Table III. (Actually, we did the expansions for $N = 1$, and checked them for the separable case; as our error constant is only a heuristic, the extra algebra to find a small correction was judged not worthwhile). They are rather large, because this method completely drops terms of order k^{p+1} in the generating function. Other methods approximate these terms. However, the Taylor series method requires solving only N equations per time-step.

There is a class of symplectic integrators which have any order and only require evaluating $H_{\mathbf{p}}$ and $H_{\mathbf{q}}$—the right-hand-sides of the ODEs. These are the fully-implicit Gauss-Legendre Runge-Kutta methods discovered by Butcher [10] and shown to be symplectic by Sanz-Serna [11] and independently by Lasagni [12]. (Remarkably, these methods have optimal order and are also A-stable [13]). There are also singly-implicit symplectic methods, but these have no apparent computational advantage in the present non-stiff case.

For completeness we summarise the Gauss-Legendre methods. The s-stage method has order $2s$. Let the Hamiltonian system be $\dot{\mathbf{x}} = J\nabla H(\mathbf{x}) = \mathbf{f}(\mathbf{x})$. First one solves the $2Ns$ equations

$$\mathbf{g}_j = \mathbf{f}\big(\mathbf{x}_0 + k\sum_{i=1}^{s} a_{ij}\mathbf{g}_i\big), \qquad j = 1, \ldots s. \tag{3.1}$$

The integrated variables $\mathbf{x}_1 = \mathbf{x}(k)$ are given by

$$\mathbf{x}_1 = \mathbf{x}_0 + k\sum_{i=1}^{s} b_i\mathbf{g}_i.$$

The constants a_{ij} and b_i are given for $s = 1$, 2, and 3 in Table IV; Butcher [10] also gives the cases $s = 4$ and 5.

To find the principal error function we expand $\mathbf{x}_1$ in a Taylor series about $\mathbf{x}_0$. This is equal, to order k^p $(p = 2s)$, to the time-k flow of the Hamiltonian $H - H_0 t^p$. Comparing the terms of order k^{p+1} gives equations for $H_{0\mathbf{p}}$ and $H_{0\mathbf{q}}$, which are consistent if the method is symplectic. They are integrated to find the principal error function.

For $s = 1$, the familiar midpoint rule, we find

$$h_2 = \tfrac{1}{8}\left(H_{\mathbf{pp}}H_{\mathbf{q}}^2 - 2H_{\mathbf{pq}}H_{\mathbf{p}}H_{\mathbf{q}} + H_{\mathbf{qq}}H_{\mathbf{p}}^2\right)$$

giving an error constant of $\frac{\sqrt{6}}{8} \sim 0.31$ or $\frac{\sqrt{2}}{8} \sim 0.18$ in the separable case.

For $s = 2$, the principal error function is shown in Table V; the error constant is 0.032, 50 times smaller than that of the fourth-order Taylor series method. However, there are four times as many equations to solve, and the convergence of the straightforward iteration scheme

$$\begin{aligned}\mathbf{g}_1^{(l+1)} &= \mathbf{f}\big(\mathbf{x}_0 + k(a_{11}\mathbf{g}_1^{(l)} + a_{12}\mathbf{g}_2^{(l)})\big)\\ \mathbf{g}_2^{(l+1)} &= \mathbf{f}\big(\mathbf{x}_0 + k(a_{21}\mathbf{g}_1^{(l+1)} + a_{22}\mathbf{g}_2^{(l)})\big)\end{aligned}$$

is usually slower than simple iteration in the Taylor series method. In the latter case, an initial guess for $\mathbf{q}_1$ can be obtained from any explicit method such as Adams-Bashforth of the same order. For Gauss-Legendre methods, it is impractical to expand the solution of (3.1) in a Taylor series because one needs to estimate all the cross derivatives of $\mathbf{f}$, which is expensive when N is large. Instead we have simply stored some of the previous $\mathbf{g}$'s and extrapolated to get an initial value. The order of the extrapolation need not match the order of the integrator. With $\mathbf{g}_i^n = \mathbf{g}_i|_{t=nk}$, M-point polynomial extrapolation gives

$$\mathbf{g}_i^{n+1} = \sum_{j=1}^{M}(-1)^{j-1}\binom{M}{j}\mathbf{g}_i^{n+1-j}.$$

This eliminates about two iterations for the pendulum (see Table VI). Similar results are obtained on the Hénon-Heiles system, and with rational instead of polynomial extrapolation. Convergence can be further enhanced by using full Gauss-Seidel, i.e., reevaluating $\mathbf{f}$ after each element of $\mathbf{g}_i$ is updated, but this will usually not be cost-effective.

The faster convergence and improved predictors possible in the Taylor series method can be seen in Table VI for the pendulum. Because the number of iterations needed is $1+\mathcal{O}(k)$, for moderate k both methods appear in Figure 3 to be better than fourth order, and perform roughly equally. In the general case, the Taylor series method could be superior if the derivatives of H are fairly simple; otherwise Gauss-Legendre is preferred. The latter also extends simply to higher order.

The Yoshida methods may be applied to general Hamiltonians if leapfrog is replaced by the midpoint rule as the basic unit. The fourth-order version of this idea is also in Sanz-Serna and Abia [20]: Let $\phi(k)$ be the midpoint method with step-size k. Then $\phi(\beta k)\circ\phi((1-2\beta)k)\circ\phi(\beta k)$ is fourth order when $\beta=(2+2^{1/3}+2^{-1/3})/3$. We find an error constant of $(5543296+440624\cdot 2^{\frac{1}{3}}+3489430\cdot 2^{\frac{2}{3}})^{\frac{1}{2}}/6912\approx 0.659$ for this method.

Midpoint variants are due to Feng [4] and Stofer [21]. By expanding about the point $(\mathbf{x}_0+\mathbf{x}_1)/2$, Feng obtains a one-step, fourth order method requiring only three derivatives of H and instead of the four needed by the simple Taylor-series method.

$$\begin{aligned}\widetilde{H} &= H+\frac{k^2}{24}(H_{\mathbf{x}}JH_{\mathbf{xx}}JH_{\mathbf{x}})\\ &= H+\frac{k^2}{24}(H_{p_ip_j}H_{q_i}H_{q_j}-2H_{p_iq_j}H_{p_i}H_{q_j}+H_{q_iq_j}H_{p_i}H_{p_j})\cdot\\ \mathbf{x}_1 &= \mathbf{x}_0+kJ\widetilde{H}_{\mathbf{x}}\big((\mathbf{x}_0+\mathbf{x}_1)/2\big)\end{aligned}$$

Note that there are $2N$ implicit equations to be solved. Because of the centered expansion, the method is very accurate even though less information is used: the error constant is 0.086, less than three times as large as fourth-order Gauss-Legendre. In addition, one can use Adams-Bashforth as a predictor.

Stofer [21] reduces the number of derivatives to two with

$$\widetilde{H}=(1-2\beta)H(\mathbf{x})+\beta H\big(\mathbf{x}+\alpha kJH_{\mathbf{x}}(\mathbf{x})\big)+\beta H\big(\mathbf{x}-\alpha kJH_{\mathbf{x}}(\mathbf{x})\big)$$

where $\alpha^2\beta=1/24$. The error constant is $(9216+70(20\alpha^2-3))^{\frac{1}{2}}/1152$, which reaches a minimum of $\frac{1}{12}$ at $\alpha=\sqrt{0.15}$, although surprisingly it does not depend strongly on α. (This weak dependence was confirmed in numerical tests.) So this method seems to be preferable to that of Feng.

4. Examples

4.1 The pendulum

We first test the results of the previous sections by applying the methods to an integrable system: the one degree-of-freedom planar pendulum, with Hamiltonian

$$H(q,p)=\tfrac{1}{2}p^2-\cos q.$$

We took the initial condition $(2,0)$ on the upper homoclinic orbit and integrated forward to $t=5000$. One expects a method of order p to trace out both homoclinic orbits with an accuracy of $\mathcal{O}(k^p)$. This is shown in Figure 2, which also indicates that the energy error agrees with our "average-case" Hamiltonian error to within a factor of two. For example, the root-mean-square energy errors $(\overline{[H(q(nk),p(nk)-H(q(0),p(0))]^2}^{1/2})$ for the four fourth-order methods (optimal, Gauss-Legendre, Candy and Rozmus, and Taylor series) are in the ratio $1:1.6:54:175$ instead of the average Hamiltonian error $1:3.5:112:245$.

(Sometimes the performance is better than the average case: on a chaotic orbit in the 2 degree-of-freedom Hénon-Heiles system, errors for the optimal, Gauss-Legendre, and Candy and Rozmus methods were in the ratio $1:1.6:152$).

It is important to note that the error in the energy does not increase as $t \to \infty$, because the discrete orbits cannot cross the nearby preserved invariant tori. So in this case we do have the long-time stability often claimed for SIs.

Different order methods can only be compared in specific cases, because the error terms contain derivatives of different orders. Figure 3 shows that the optimal 4th order method is preferred for accuracies $< 10^{-2}$ ($k \sim 1$), and the 5th order method for accuracies $< 4 \times 10^{-6}$ ($k \sim 0.5$).

The orbital (action) error (proportional in this case to the energy error, except near $(\pi, 0)$ and $(0,0)$) will depend on the properties of the real system near the orbit. From KAM theory (here in its guise as the Moser Twist Theorem [14]), we know that invariant tori with sufficiently irrational rotation numbers are preserved by small $\mathcal{O}(\epsilon)$ symplectic perturbations. Here $\epsilon = \mathcal{O}(k^p)$. The destroyed tori are "replaced" by island chains interlaced by chaotic separatrix meshes in the perturbed system. These can give rise to larger errors in the discrete dynamics, as shown in Figure 4, where order L resonances magnify the energy error by more than an order of magnitude. The position of the resonances are also shifted from their unperturbed values, because although closed orbits are preserved, the speed along those orbits changes by $\mathcal{O}(\epsilon)$.

However, as $\epsilon \to 0$, the width of each resonant band is $\mathcal{O}(\epsilon^{L/2})$. For $L > 2$, this is subdominant to the already-present error of $\mathcal{O}(\epsilon)$, and we can expect to be able to ignore resonances between the time-step and the orbit if the time-step is sufficiently reduced. Thus, the resonances observed in Figure 4 are a nonlinear effect; linear theory cannot explain why some resonances are so much larger than others. The cases $L = 1$ and $L = 2$ indicate a kind of Nyquist sampling limit: we should not expect to reproduce dynamics with a time-scale shorter than twice the time-step.

4.2 Convergence of chaotic sets

Here we repeat the $1\frac{1}{2}$ degree-of-freedom example considered by Candy and Rozmus to investigate the convergence of the optimal fourth-order method. We use this method in the remainder of the paper. The Hamiltonian

$$H(q,p,t) = \tfrac{1}{2}p^2 + \tfrac{1}{2}q^2 + \epsilon\cos(q - 7t)$$

governs the motion of a linear oscillator perturbed by a resonant plane wave. See [2] and the references therein for a more extensive discussion of this problem. The averaged system has a separatrix mesh joining the hyperbolic points $(r\cos\theta, r\sin\theta)$ where $J_7(r) = \cos 7\theta = 0$. In the full problem, these separatrices split, forming a chaotic web, connected as $r \to \infty$.

A Poincaré section is defined by the period of the forcing: $T_j = j\frac{2\pi}{7}$, $j = 1,\ 2,\ \ldots$. We vary the number of time steps per forcing period, m, and plot (q,p) at $t = T_j$. As seen in Figure 5, even $m = 4$ roughly reproduces the orbit. However, capturing its finer-scale structure requires smaller time-steps. To study this, we integrated until the orbit dropped a level, took a window $[-2,2]\times[9,11]$, and fit a least-squares polynomial through all points lying within this window. This enables the mean path of the orbit to be subtracted out and the vertical scale expanded, as shown in Figure 6. For $m = 6$, the higher-order resonances inside the chaotic set are destroyed, and $m = 18$ ($k \sim 0.05$) is required to determine the shape of the set on this scale. By way of comparison, leapfrog requires $m = 500$ to reproduce the fourth-order $m = 18$ results. The tiny islands in the last frame of Figure 6 have period near 1000, yet are well approximated with only 18 time-steps per period.

Channel and Scovel [3] comment that the divergence of nearby orbits in a chaotic set destroys the accuracy of a *particular* orbit, but that the invariant tori bounding the chaotic set enable an SI to reproduce it. In fact, those bounding tori may be among the first to be destroyed by the perturbation due to the integrator. A convergence study is therefore essential.

The shape of the closure of an orbit (what we have been calling the chaotic set) is not its only property. Of greater importance are the long-time statistics of the behaviour of orbits. For example, in this system we might calculate statistics on a state variable defined as the separatrix level (see Fig. 7), or by the order in which neighbourhoods of hyperbolic points are visited, or by a partition of phase space along the asymptotic manifold of hyperbolic points. For simplicity, we defined an exit time as the first time for which

$r = \sqrt{q^2+p^2} < 8$ or $r > 17$, i.e., the time for the particle to move to a different separatrix level, in units of the forcing period $\frac{2\pi}{7}$. (These orbits almost always drop to the lowest level). For $\epsilon = 0.8$, the mean time is more than 50000, so we took $\epsilon = 2$. To obtain statistics, we took 225 initial conditions equally spaced in a $10^{-4} \times 10^{-4}$ square about $(0, 10.5939)$. Although the exit time for a *particular* initial condition will not converge until the entire orbit is tracked accurately, it is reasonable to hope that the cantori within the chaotic band, and thus the distribution of times for the whole ensemble will converge much more quickly. This is confirmed in Table VII: even $m = 8$ gives converged results (allowing for sampling error). Because of the extreme skewness of such distributions, determining the tail accurately (time$>$ 8000), and thus the mean, is much more difficult.

4.3 Several degrees of freedom

As we saw in §1, the discretisation of an N degree-of-freedom Hamiltonian evolves in the $N+1$ degree-of-freedom extended phase space $(q_1, \ldots, q_N, \xi, p_1, \ldots, p_N, p_\xi)$. (Generally one takes $\xi = t$ and $p_\xi = \hat{H}$, the value of the Hamiltonian function on the orbit). When $N \geq 2$, $N+1$-dimensional invariant tori of the discrete system do not separate regions of this space. Thus the error in the energy will not satisfy

$$\sup_{0 \leq t < \infty} |H(\mathbf{q}(t), \mathbf{p}(t)) - \hat{H}(\mathbf{q}(t), \mathbf{p}(t))| = \mathcal{O}(k^p) \tag{4.1}$$

as it does in the one degree-of-freedom case. For $1\frac{1}{2}$ degrees-of-freedom, the extended phase spaces of the discrete and the continuous systems are the same, and KAM tori separate it; so although the energy error can grow to be formally $\mathcal{O}(1)$, at least the discrete orbits remain confined to the correct region of phase space. But for $N \geq 2$ we lose even this property.

The problem is qualitatively worst when $N = 2$. Here KAM 2-tori of the real system separate a 3-dimensional energy level set. But the discrete system has 3 degrees of freedom, so KAM 3-tori do not separate the 5-dimensional energy level sets. The discrete orbits may therefore drift any distance across energy surfaces of the real system, thereby crossing invariant surfaces of the real flow, and possibly entering regions of phase space with quite different dynamics.

This drift is known as Arnol'd diffusion. If the real flow is almost integrable, or if the initial conditions are close enough to an elliptic periodic orbit, from the estimates of Nekhoroshev [8], Appendix 8) the drift is generically exponentially slow in the perturbation. We might therefore expect it to be negligible for practical purposes, and believe that this can account for some successful large-d.o.f. integrations (e.g., in [3]). It still remains to be seen what advantages symplectic integrators offer in fully chaotic, large-d.o.f. systems.

If one wishes to investigate chaotic orbits or isolated invariant sets the diffusion may be rapid. We investigated this phenomenon in the case of two coupled pendula, with Hamiltonian

$$H(q_1, q_2, p_1, p_2) = \tfrac{1}{2}(p_1^2 + p_2^2) - \cos q_1 - \cos q_2 - \epsilon \cos(q_1 - q_2)$$

The behaviour of the energy error is strikingly different depending on the initial conditions, although always $\mathcal{O}(k^p)$ after a *fixed* time (see Figure 8). With initial condition $(0, 2)$ near the elliptic orbit in which both pendula rotate anticlockwise in tandem, the energy error is quasi-periodic, indicating an orbit on an invariant torus of $\hat{H}$. With initial condition $(2, 2)$, on a chaotic orbit, it behaves like a random walk: even the typical $\mathcal{O}(\sqrt{t})$ growth is evident. As k is reduced, the rate of growth is reduced, but it can never be eliminated. Even up to $t = 1000$, errors are about 50 times larger than in the previous case.

The orbit quickly crosses to a different energy level of the real system and then into regions which are forbidden in the true dynamics. If there are qualitative errors like this, they can be delayed by reducing the time-step (but the same could be said of standard integrators). In any event, long-time stability in the sense of (4.1) is always lost.

In some $N = 2$ problems it may be possible to avoid this difficulty by a change of variables. For instance, if one coordinate is monotonic (e.g. $\dot{q}_2 > 0$) then the Hamiltonian may be rewritten in the reduced phase space (e.g. $\tilde{H}(q_1, p_1, t) = p_2(q_1, p_1, t)$, with time $t = q_2$). Now, by construction, the original energy is conserved exactly! Of course the value of the new energy (extending to $N = 2$ again) drifts as before, but the possibility of Arnol'd diffusion has been eliminated.

For $N > 2$ degrees of freedom, the real system also exhibits Arnol'd diffusion. However, the rate of diffusion, although small, depends on the dimension of the system, and would therefore be faster in the discrete than in the real system.

References

[1] Ruth, R.D., A canonical integration technique, *IEEE Trans. Nucl. Sci. NS*-**30** (1983), 2669–2671.

[2] Candy, J. and Rozmus, W., A symplectic Integration Algorithm for Separable Hamiltonian Functions, *J. Comp. Phys.* **92** (1991), 230–256.

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

[4] Feng Kang and Qin Meng-Zhao, The symplectic methods for the computation of Hamiltonian equations, *Numerical methods for partial differential equations, Proc. Shanghai* 1986, Lecture Notes in Mathematics **1297** 1–37, ed. Zhu You-lan and Gu Ben-yu, Springer, Berlin, 1987.

[5] Feng Kang and Qin Meng-Zhao, Hamiltonian Algorithms for Hamiltonian Systems and a Comparative Numerical Study, *Comp. Phys. Comm.*, to appear.

[6] Feng Kang, The Hamiltonian Algorithms for Hamiltonian Systems, *IMACS 1st Int. Conf. Comp. Phys.*

[7] Ge Zhong and Marsden, J.E., Lie-Poisson Hamilton-Jacobi theory and Lie-Poisson integrators, *Phys. Lett. A* **133**(1988), 134–139.

[8] Arnol'd, Mathematical Methods of Classical Mechanics (1978), Springer-Verlag, New York.

[9] Li, T.Y., Sauer, T., and Yorke, J.A., The random product homotopy and deficient polynomial systems, *Numer. Math.* **51** (1987), 481–500.

[10] Butcher, J.C., Implicit Runge-Kutta Processes, *Math. Comp.* (1964)

[11] Sanz-Serna, J.M., Runge-Kutta Schemes for Hamiltonian Systems, *BIT* **28** (1988), 877-883.

[12] Lasagni, F.M., Canonical Runge-Kutta methods, *ZAMP* **39**(11) (1988), 952–953.

[13] Burrage, K., and Butcher, J.C., Stability criteria for implicit Runge-Kutta methods, *SIAM J. Numer. Anal.* **16** (1979), 46–57.

[14] Moser, J. (1973). Stable and Random Motions in Dynamical Systems, Princeton University Press: Princeton.

[15] Forest, E., and Ruth, R., Fourth-order symplectic integration, *Physica D* **43** (1990), 105–117.

[16] Sanz-Serna, J.M., and Vadillo, F., Studies in numerical nonlinear instability, III: Augmented Hamiltonian systems, *SIAM J. Appl. Math* **47**(1) (1987), 92–108.

[17] Yoshida, H., Construction of higher order symplectic integrators, *Phys. Lett. A* **150** (1990), 262–269.

[18] Auerbach and Friedman, Long-time behaviour of numerically computed orbits: small and intermediate time-step analysis of one degree-of-freedom systems, *J. Comp. Phys* **93**(1) (1991), 189–224.

[19] Sanz-Serna, J.M., The numerical integration of Hamiltonian systems, lecture at Conf. Comp. Diff. Eqs., 1989.

[20] Sanz-Serna, J.M., and Abia, L., Order conditions for canonical Runge-Kutta schemes, *SIAM J. Numer. Anal.*, to appear; Partitioned Runge-Kutta methods for separable Hamiltonian problems, *Report 1990/8, Departamento de Matemática Aplicada y Computación, Universidad de Valladolid, Valladolid, Spain* (1990); Calvo, M.P. and Sanz-Serna, J.M., Order conditions for canonical Runge-Kutta-Nyström methods, *Report 1991/1*.

[21] Stofer, D.M., Some geometric and numerical methods for perturbed integrable systems, *Thesis*, ETH-Zürich (1988).

Table I Number of equations

Order	m	# terms in h_m		Total # equations	
		Separable	T quadratic	Separable	T quadratic
1	0	2	2	2	2
2	1	1	1	3	3
3	2	2	2	5	5
4	3	3	2	8	7
5	4	6	4	14	11
6	5	10	5	24	16

Table III Error constants of implicit methods

Order	# eqtns	General H	Separable	Quadratic	
Taylor series methods:					
2	N		0.87	0.71	0.71
3	N	1.15	0.87	0.85	
4	N	1.60	0.71	0.60	
5	N	2.22	0.80	0.66	
Gauss-Legendre methods:					
2	2N	0.31	0.18	0.18	
4	4N	0.032	0.011	0.009	
Feng midpoint:					
4	2N	0.086	0.033	0.031	
Stofer midpoint, $\alpha^2 = \frac{3}{20}$:					
4	2N	0.083	0.033	0.031	
Sanz-Serna (3 midpoint steps):					
4	6N	0.659	0.239	0.236	

Table IV Coefficients of Gauss-Legendre methods

$s = 1$:

$$a_{11} = \tfrac{1}{2},\ b_1 = 1$$

$s = 2$:

$$(a_{ij}) = \begin{pmatrix} \frac{1}{4} & \frac{1}{4} - \frac{\sqrt{3}}{6} \\ \frac{1}{4} + \frac{\sqrt{3}}{6} & \frac{1}{4} \end{pmatrix}$$

$$(b_i) = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \end{pmatrix}$$

$s = 3$:

$$(a_{ij}) = \begin{pmatrix} \frac{5}{36} & \frac{2}{9} - \frac{\sqrt{15}}{15} & \frac{5}{36} - \frac{\sqrt{15}}{30} \\ \frac{5}{36} + \frac{\sqrt{15}}{24} & \frac{2}{9} & \frac{5}{36} - \frac{\sqrt{15}}{24} \\ \frac{5}{36} + \frac{\sqrt{15}}{30} & \frac{2}{9} + \frac{\sqrt{15}}{15} & \frac{5}{36} \end{pmatrix}$$

$$(b_i) = \begin{pmatrix} \frac{5}{18} & \frac{4}{9} & \frac{5}{18} \end{pmatrix}$$

Table II Summary of Explicit Methods

Method	**a**	**b**	Error constant (& effective e.c.) General T	T quadratic
Order 2				
Leapfrog	$a_1 = a_2 = \frac{1}{2}$	$b_1 = 0,\ b_2 = 1$	0.280 (0.070)	
Pseudo-"	$a_1 = 1,\ a_2 = 0$	$b_1 = b_2 = \frac{1}{2}$	0.280	
Optimal	$a_1 = \frac{1}{\sqrt{2}} = 1 - a_2$	$b_1 = \frac{1}{\sqrt{2}} = 1 - b_2$	0.043	
Order 3				
Ruth [1]	$a_1 = \frac{2}{3},\ a_2 = -\frac{2}{3},\ a_3 = 1$	$b_1 = \frac{7}{24},\ b_2 = \frac{3}{4},\ b_3 = -\frac{1}{24}$	0.113	
Optimal	Let $z = -\left(\frac{2}{27} - \frac{1}{9\sqrt{3}}\right)^{1/3}$			
	$w = -\frac{2}{3} + \frac{1}{9z} + z$			
	$y = \frac{1}{4}(1 + w^2)$			
	$a_1 = \left(\frac{1}{9y} - \frac{w}{2} + \sqrt{y}\right)^{1/2} - \frac{1}{3\sqrt{y}} = 0.919661523017399857$			
	$a_2 = \frac{1}{4a_1} - \frac{a_1}{2},\ a_3 = 1 - a_1 - a_2$	$b_i = a_{4-i}$	0.058	
Order 4				
Candy [2],	$a_1 = a_4 = \frac{1}{6}(2 + 2^{1/3} + 2^{-1/3})$	$b_1 = 0,\ b_3 = (1 - 2^{2/3})^{-1}$		
Forest [15]	$a_2 = a_3 = \frac{1}{6}(1 - 2^{1/3} - 2^{-1/3})$	$b_2 = b_4 = (2 - 2^{1/3})^{-1}$	0.311 (0.098)	0.276 (0.087)
Optimal	$a_1 = 0.515352837431122936$	$b_1 = 0.134496199277431089$		
	$a_2 = -0.085782019412973646$	$b_2 = -0.224819803079420806$		
	$a_3 = 0.441583023616466524$	$b_3 = 0.756320000515668291$		
	$a_4 = 0.128846158365384185$	$b_4 = 0.334003603286321425$	—	0.0025
Order 5				
Optimal	$a_1 = 0.339839625839110000$	$b_1 = 0.119390029287567275$		
	$a_2 = -0.088601336903027329$	$b_2 = 0.698927370382475230$		
	$a_3 = 0.585856476825962118$	$b_3 = -0.171312358271600775$		
	$a_4 = -0.603039356536491888$	$b_4 = 0.401269502251353448$		
	$a_5 = 0.323580796554697639$	$b_5 = 0.010705081848235984$		
	$a_6 = 0.442363794219749458$	$b_6 = -0.058979625498031163$	—	0.0079

Table V Principal error functions of implicit methods
We take $N = 1$ and compute the coefficient h_4 of t^4 in $E = H - \hat{H}$ for three fourth-order methods.

	Taylor series	Gauss-Legendre	Kang mid-point
Term	$120h_4$	$864h_4$	$384h_4$
$H_{pppp}H_q^4$	1	1	1
$H_{ppp}H_q^3H_{qp}$	6	6	–4
$H_{pp}H_q^2H_{qp}^2$	7	–3	–8
$H_pH_qH_{qp}^3$	1	6	16
$H_{pp}H_q^3H_{qpp}$	4	–6	4
$H_pH_q^2H_{qp}H_{qpp}$	4	–6	4
$H_pH_q^3H_{qppp}$	1	–4	–4
$H_{pp}^2H_q^2H_{qq}$	8	3	8
$H_pH_{ppp}H_q^2H_{qq}$	9	–6	4
$H_pH_{pp}H_qH_{qp}H_{qq}$	29	–6	–16
$H_p^2H_{qp}^2H_{qq}$	7	–3	–8
$H_p^2H_qH_{qpp}H_{qq}$	7	12	–8
$H_p^2H_{pp}H_{qq}^2$	8	3	8
$H_pH_{pp}H_q^2H_{qqp}$	7	12	–8
$H_p^2H_qH_{qp}H_{qqp}$	4	–6	4
$H_p^3H_{qq}H_{qqp}$	4	–6	4
$H_p^2H_q^2H_{qqpp}$	1	6	6
$H_p^2H_{pp}H_qH_{qqq}$	9	–6	4
$H_p^3H_{qp}H_{qqq}$	6	–6	–4
$H_p^3H_qH_{qqqp}$	1	–4	–4
$H_p^4H_{qqqq}$	1	1	1

Table VI Convergence of implicit methods
We give the mean number of iterations for each method to converge, per time-step, with and without a predictor. The system is the pendulum with initial condition $(q, p) = (1, 1)$.

k	Gauss-Legendre		Taylor series	
	No predictor	5-pt extrapolation	No predictor	4th order Adams-Bashforth
$k = 0.25$	11.3	10.4	8.5	7.1
$k = 0.1$	8.8	7.2	6.2	4.6
$k = 0.01$	5.8	3.5	3.9	2.0

Table VII Exit time distribution

The exit times T for 225 initial conditions are divided into bins $1000i \leq T < 1000(i+1)$ $(0 \leq i \leq 11)$ and bin 12, $T \geq 12000$. We give the mean, $\overline{T}$, standard deviation s, and the number of orbits in each bin for different numbers of time-steps per period, m.

m	$\overline{T}$	s	Bin number												
			0	1	2	3	4	5	6	7	8	9	10	11	12
4	7185	14417	1	40	52	34	28	17	6	7	6	2	1	3	28
6	2285	3886	52	96	45	11	7	9	0	0	0	1	1	1	2
8	2133	2165	45	106	44	14	4	6	1	0	1	0	1	0	3
16	2760	5214	38	101	40	22	11	3	1	1	0	1	0	1	6
24	2346	4095	38	107	40	20	8	3	1	3	2	0	1	1	1

Figure 1 Real and approximate flows

ϕ is the flow of the Hamiltonian H; $\hat{\phi}$ is the flow of the approximate Hamiltonian $\hat{H}$.
$E = H - \hat{H}$ is $\mathcal{O}(k^p)$, k-periodic, and zero at $t = nk$.

Figure 2 Energy error

The pendulum is integrated up to $t = 5000$ with initial condition $(q, p) = (0, 2)$ and calculated the root-mean-square energy error over all time steps.

Figure 3 Relative efficiency of different methods

See caption for Fig. 2. Each implicit method requires two trigonometric function evaluations per iteration in this case; we have counted this as two force evaluations.

Figure 4 Effect of time-step resonances

Here we integrate the pendulum with the optimal fourth-order method for initial conditions $(0, p)$, $0 < p < 4$. The RMS energy error up to $t = 2000$ is shown.

Figure 5 $H = \frac{1}{2}(p^2 + q^2) - \epsilon \cos(q - 7t)$

Part of an orbit on the chaotic web formed by separatrix splitting (see Fig. 7). The initial condition is $(q_0, p_0) = (0., 10.5939)$ and there are $m = 4$ time-steps per forcing period. The Poincaré map is shown for 25000 periods. The orbit stays on the same two separatrix levels.

Figure 6 Convergence of a chaotic set

The part of the orbit outlined in Fig. 5 is magnified for four different numbers, m, of time steps per forcing period. $\triangle p$ is the difference between an iterate of the Poincaré map and the mean orbit (see text). The vertical scale around the mean orbit is expanded by $250\times$. The initial condition is as in Fig. 5 and 50000 forcing periods are shown, unless the orbit first drops to a lower level. The bottom picture shows some of the fine structure within the chaotic band.

Figure 7 Separatrix mesh

The separatrix mesh for the forced oscillator of §4.2, calculated from first-order averaging (from Candy and Rozmus [2]).

Figure 8 Energy error for coupled pendula

$H - h$ is the change in energy from that at the initial condition, $k = 0.5$. (a) Initial condition $(0, 2)$ (near elliptic periodic orbit); (b) initial condition $(2, 2)$ (chaotic orbit). Note the different vertical scales.