On order 5 symplectic explicit Runge-Kutta Nyström methods

LIN-YI CHOU †   linyi@scitec.auckland.ac.nz

P.W. SHARP   sharp@scitec.auckland.ac.nz

Department of Mathematics, University of Auckland, Private Bag 92019,
Auckland, NEW ZEALAND.

Abstract.   Order five symplectic explicit Runge-Kutta Nyström methods of five stages are known to exist. However, these methods do not have free parameters with which to minimise the principal error coefficients. By adding one derivative evaluation per step, to give either a six-stage non-FSAL family or a seven-stage FSAL family of methods, two free parameters become available for the minimisation. This raises the possibility of improving the efficiency of order five methods despite the extra cost of taking a step.

We perform a minimisation of the two families to obtain an optimal method and then compare its numerical performance with published methods of orders four to seven. These comparisons along with those based on the principal error coefficients show the new method is significantly more efficient than the five-stage, order five methods. The numerical comparisons also suggest the new methods can be more efficient than published methods of other orders.

1. Introduction

Time-dependent processes such as those arising in mechanics and chemistry where dissipation is insignificant can often be modelled by a Hamiltonian system of ordinary differential equations of the form

$$p' = -\frac{\partial H(p,q)}{\partial q}, \quad q' = \frac{\partial H(p,q)}{\partial p}, \tag{1}$$

where $' \equiv d/dt$ and $H : R^n \times R^n \mapsto R$ is sufficiently smooth.

A one step numerical method for (1) is called symplectic if it preserves the symplectic structure of the space of variables (p,q), thus reproducing the main qualitative property of the solution. Numerical experiments have shown symplectic methods can be more efficient than non-symplectic methods for long intervals of integration. In order to preserve the symplectic structure, a constant stepsize is used.

Symplectic Runge-Kutta methods are necessarily implicit. However, for separable Hamiltonians (those of the form $H(q,p) = T(p) + V(q)$), explicit Runge-Kutta Nyström (ERKN) methods can be symplectic if the coefficients of the method satisfy certain conditions.

---

†   The work was part of the requirements for the first author's M.Sc. thesis. The first author would like to thank the Department of Mathematics, University of Auckland for its financial support.

Okunbor and Skeel [6] investigated families of symplectic ERKN methods that use one, two and three stages. Calvo and Sanz-Serna in [1], [2] and [3] presented and tested a five-stage, order four method with minimised principal error coefficients. The method re-uses the last stage as the first stage of the next step, a property known as FSAL. This means only four derivations evaluations are required on each step, except for the first step. In [4], Calvo and Sanz-Serna presented a 13-stage, order seven FSAL method with minimised principal error coefficients. Okunbor and Skeel [7] investigated the existence of order five and six methods. They performed an extensive numerical search and found four individual five-stage, order five methods. They conjectured there were no six-stage methods of order six and through a numerical search obtained 16 individual seven-stage, order six methods.

For five-stage, order five methods, there are ten order conditions to satisfy and ten coefficients to satisfy these conditions. Hence, no coefficients are available to minimise the principal error coefficients. If one stage is added, and the order is kept at five, two coefficients become available, but this is at the expense of increasing by one the number of derivative evaluations required to take a step.

This trade-off between decreasing the error and increasing the cost for each step raises the interesting question of whether the introduction of the sixth stage will lead to a gain or loss in efficiency, where efficiency is measured by the number of derivative evaluations required to achieve a prescribed global error.

2.   The methods

2.1.   Definitions

When a Hamiltonian is separable, (1) can be written as

$$y'' = f(y), \tag{2}$$

where $f : R^n \mapsto R^n$.

The ERKN methods we consider use s-stages and generate order p approximations y_i and y_i' to $y(x_i)$ and $y'(x_i)$ respectively, $i = 1, 2, \ldots$, according to

$$y_i = y_{i-1} + hy_{i-1}' + h^2 \sum_{j=1}^{s} b_j f_j, \tag{3}$$

$$y_i' = y_{i-1}' + h \sum_{j=1}^{s} b_j' f_j, \tag{4}$$

where $h = x_i - x_{i-1}$ is constant, and

$$f_j = f(x_{i-1} + c_j h, y_{i-1} + c_j h y_{i-1}' + h^2 \sum_{k=1}^{j-1} a_{jk} f_k), \quad j = 1, \ldots, s.$$

The prime in b_j' denotes the derivative formula and not the derivative of b_j.

In all the methods we consider, the b_j are given by

$$b_j = (1 - c_j)b'_j, \quad j = 1, \ldots, s. \tag{5}$$

For an ERKN method to be symplectic, the a_{jk} must satisfy

$$a_{jk} = (c_j - c_k)b'_k, \quad k = 1, \ldots, j-1, \quad j = 2, \ldots, s. \tag{6}$$

Hence, once c_j, b_j, $j = 1, \ldots, s$ are known, the remaining coefficients of a symplectic ERKN are known.

For a method to be FSAL, its coefficients must satisfy $c_1 = 0$, $c_s = 1$ and

$$a_{sj} = b_j, \quad j = 1, \ldots, s-1. \tag{7}$$

If $c_1 = 0$ and $c_s = 1$ and the method is symplectic, (7) is automatically satisfied.

We decided against using the simplifying assumptions

$$\frac{c_j^2}{2} = \sum_{k=1}^{j-1} a_{jk}, \quad j = 2, \ldots, s,$$

because these led to a net reduction in the number of free parameters available for minimising the principal error coefficients.

In the remainder of this paper the term 'method' will mean 'symplectic method'.

2.2.   Six-stage, order five non-FSAL methods

There are 13 order conditions up to and including order five. However, three of these order conditions are dependent on the remainder (see, for example [7]). This means, since six-stage non-FSAL ERKN methods have 12 coefficients ($c_j, b_j, j = 1, \ldots, 6$), we can take two coefficients as free parameters and solve for the remaining coefficients. We found it convenient to take c_5 and c_6 as the free parameters.

Five of the 10 order conditions to satisfy are the quadrature conditions

$$\sum_{j=1}^{s} b_j c_j^k = (k+1)^{-1}, \quad k = 0, \ldots, 4, \tag{8}$$

where we take $0^0 = 1$ and $0^k = 0$, $k > 0$. Conditions (8) are easily solved for $b'_1, \ldots, b'_5$ in terms of the c and b'_6.

The next condition we solve is

$$\sum_{j=2}^{6} b'_j \sum_{k=1}^{j-1} a_{jk} = \frac{1}{6}. \tag{9}$$

When the expressions for $b'_1, \ldots, b'_5$ found from solving (8) are substituted into (9), the equation becomes a quadratic in b'_6. This is solved to give b'_6 and hence the other b' in terms of the c.

This leaves the following equations to satisfy, where summations are performed over i and j

$$b_i c_i a_{ij} = \frac{1}{8}, \quad b_i c_i^2 a_{ij} = \frac{1}{10}, \quad b_i c_i a_{ij} c_j = \frac{1}{30}, \quad b_i (a_{ij})^2 = \frac{1}{20}. \qquad (10)$$

When the expressions for the b' are substituted into the above equations, the equations become complicated functions of the c. Even with the help of MAPLE we were unable to solve these equations algebraically and resorted to a numerical approach.

We first attempted to use a non-interactive grid method consisting of an outer and inner grid. The outer grid was over c_5 and c_6 and was used to minimise the Euclidean norm of the order six principal error coefficients for the solution and derivative formulae. The inner grid was over c_1, c_2, c_3 and c_4 and was used to generate initial estimates of the solutions of (10). For each point on the inner grid, we used the package HYBRD (obtained from NETLIB) to solve (10). If a solution was found, we evaluated the norm of the principal error coefficients and updated our minimum and minimiser if required. As in [4], a solution was accepted only if $-1.5 \leq c_j \leq 1.5$, $j = 1, \ldots, 6$.

Although the above grid method produced tangible results after several days of CPU time, there was much duplication of effort because HYBRD would converge to the same solution from different points on the inner grid. In addition, the method had the air of brute force, something we found unappealing.

We then attempted an interactive grid search consisting of just the outer grid. By keeping a history stack of the previously converged solutions, we were able, after gaining some experience, to make rapid progress in the minimisation.

At intermediate points during the grid search, we tested the best ERKN method we had and compared its performance with published methods on a set of test problems. On a number of occasions we found the norm of the principal error coefficients significantly over-estimated the efficiency of the new method. On other occasions what we thought was a significant decrease in the norm led to little or no improvement in the efficiency.

We were unable to ascertain the reason for this sometimes poor correlation between the predicted and actual performance.

2.3.   Seven-stage, order five FSAL methods

The derivation and minimisation of seven-stage, order five FSAL methods is similar to that for the six-stage non-FSAL methods, the main difference being b_7' is now a free parameter in place of c_1. We used the interactive grid method and found through our testing at intermediate points of the search that methods from the FSAL family were more efficient than those from the non-FSAL family.

With a grid size of 0.01 for c_5 and c_6, the best method we found had $c_j, b_j,\ j = 1, 7$ as follows

$$
\begin{aligned}
c_1 &= 0.0000000000000000E + 00 & c_2 &= 0.2179621390175646E + 00 \\
c_3 &= 0.4424703708255242E + 00 & c_4 &= 0.1478460559438898E + 01 \\
c_5 &= 0.3400000000000000E + 00 & c_6 &= 0.7000000000000000E + 00 \\
c_7 &= 0.1000000000000000E + 01 \\
b'_1 &= 0.6281213570268329E - 01 & b'_2 &= 0.3788983131252575E + 00 \\
b'_3 &= 0.2754528515261340E + 00 & b'_4 &= -0.1585299574780513E - 02 \\
b'_5 &= -0.1785704038527618E + 00 & b'_6 &= 0.3479995834198831E + 00 \\
b'_7 &= 0.1149928196535844E + 00
\end{aligned}
$$

The remaining coefficients can be found using (5) and (6). The error in satisfying the order conditions is less than 3×10^{-16}.

The Euclidean norm of the order six principal error coefficients for the solution and derivative formula is 4.0×10^{-4} and 4.1×10^{-4} respectively. The corresponding norms for the four, five-stage order five methods given in [7] range from 1.7×10^{-2} to 2.4×10^{-2} for the solution formula and 3.3×10^{-2} to 4.3×10^{-2} for the derivative formula.
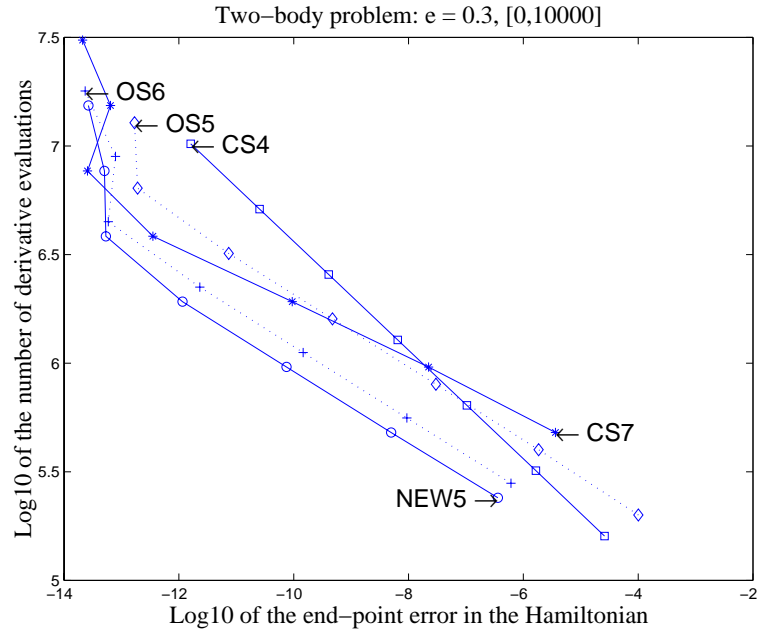

Figure 1. Base 10 log-log graph of the number of derivative evaluations against the norm of the end-point global error for the two-body problem with $e = 0.3$ and an integration interval of $[0, 10000]$.

An estimate of the relative efficiency of the new method and those in [7] can be obtained by assuming the global error varies as the fifth root of the Euclidean norms. The new method is predicted, after scaling by the number of derivative evaluations on each step, to be approximately twice as efficient as those in [7]. As stated at the end of the previous sub-section, decreasing the norm of the order six principal error coefficients does not always lead to a corresponding gain in efficiency. Hence, the statement that the new method is approximately twice as efficient as those in [7] should be taken as a general statement and not as one that holds for each and every problem.
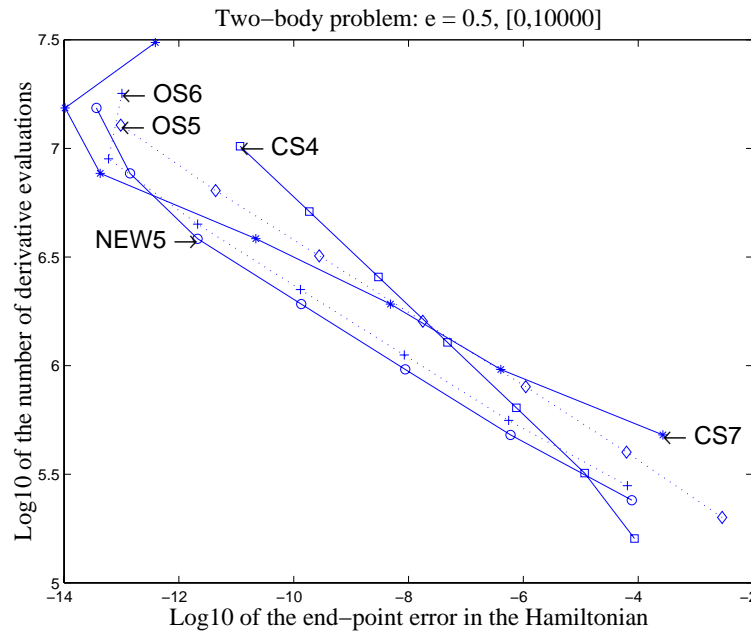

Figure 2. Base 10 log-log graph of the number of derivative evaluations against the norm of the end-point global error for the two-body problem with $e = 0.5$ and an integration interval of $[0, 10000]$.

3.   Numerical experiments

We performed extensive numerical comparisons of the new method, the order four method of [3], the order seven method of [4], and the 20 order five and six methods in [7]. Each method was tested on the simple pendulum problem, the Henon Heiles problem [5], and the two-body problem. For each method and problem we used stepsizes of 2^{-i}, $i = 2, \ldots, 7$ and several intervals of integration. For each integra-

tion we recorded the global error and the error in the Hamiltonian throughout the interval of integration. The testing was done in double precision Fortran 90.

Here we present a summary of the comparisons with an emphasis on the two-body problem because this problem is commonly used when comparing symplectic methods.

Figures 1, 2 and 3 give the log-log graphs of the number of derivative evaluations against the end-point error in the Hamiltonian for the two-body problem with eccentricities of 0.3, 0.5 and 0.7. The interval of integration is $[0, 10^4]$. Each figure contains the graphs for the new method, the methods of order four and seven in [3] and [4], and an order five and order six method from [7] (we chose the last method in Tables 1 and 2 of [7], the other methods in the tables led to similar conclusions). The five methods are denoted by NEW5, CS4, CS7, OS5 and OS6 respectively, where the digit is the order of the method.
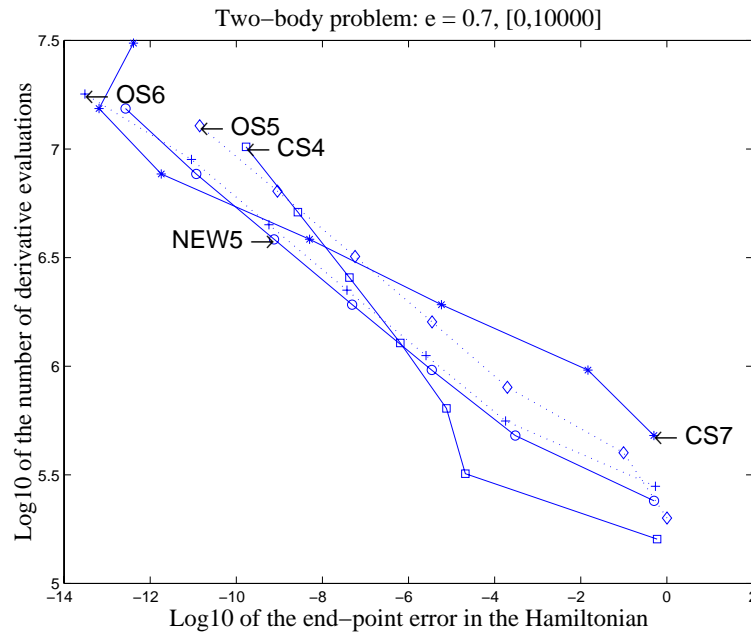

Figure 3. Base 10 log-log graph of the number of derivative evaluations against the norm of the end-point global error for the two-body problem with $e = 0.7$ and an integration interval of $[0, 10000]$.

We consider Figures 1 and 2 first. The new method is more efficient than OS5 which suggests it was advantageous to use an extra stage. The gain in efficiency is approximately a factor of two which agrees well with the gain predicted using the norm of the principal error coefficients. We also observe the new method is

more efficient than CS4 and CS7, and slightly more efficient than OS6. We verified using quadruple precision that the vertical placement of data points for small global errors was due to round-off errors in double precision.

The results in Figure 3 provide a caveat to the above conclusions. The new method has retained much of its efficiency relative to OS5 except at large global errors, but CS4 is more efficient than the new method at large global errors.

When the five methods are compared on the simple pendulum and Henon Heiles problem, the relative efficiency of the order five methods is similar to that on the two-body problem but OS6 is more efficient than the new method.

References

1. M.P. Calvo, J.M. Sanz-Serna, *Variable steps for symplectic integrators*, in Numerical Analysis, 1991, D.F. Griffiths and G.W. Watson, eds., Longmans Press, London, 1992, 32-48.

2. M.P. Calvo, J.M. Sanz-Serna, *Reasons for a failure. The integration of the two-body problem with a symplectic Runge-Kutta-Nyström code with stepchanging facilities*, in Applied Mathematics and Computational Reports, Report 1991/7, Universidad de Valladolid, Spain, 1991.

3. M. P. Calvo, J. M. Sanz-Serna, *The Development Of Variable-Step Symplectic Integrators With Application To The Two-Body Problem*, SIAM J. Sci. Comput., **14**, No.4 (1993), 936-952.

4. M. P. Calvo, J. M. Sanz-Serna, *High-Order Symplectic Runge-Kutta-Nyström Methods*, SIAM J. Sci. Comput., **14**, No.5 (1993), 1237-1252.

5. M. Hénon, C. Heiles, *The applicability of the third integral of motion: some numerical experiments*, Astron. J. **69** (1964), 73-79.

6. D. Okunbor, R.D. Skeel, *Explicit canonical methods for Hamiltonian systems*, Math. Comp., **59** (1992), 439-455.

7. D. Okunbor, R.D. Skeel, *Canonical Runge-Kutta-Nyström methods of order five and six*, Math. Comp., **51** (1994), 375-382.