Reprinted from

APPLIED NUMERICAL MATHEMATICS

Applied Numerical Mathematics 22 (1996) 217–223

Runge–Kutta type methods for orthogonal integration

Desmond J. Higham [1]

Department of Mathematics and Computer Science, University of Dundee, Dundee, DD1 4HN, Scotland, UK

ELSEVIER

ELSEVIER

Applied Numerical Mathematics 22 (1996) 217–223

APPLIED NUMERICAL MATHEMATICS

Runge–Kutta type methods for orthogonal integration

Desmond J. Higham [1]

Department of Mathematics and Computer Science, University of Dundee, Dundee, DD1 4HN, Scotland, UK

Abstract

A simple characterisation exists for the class of real-valued, autonomous, matrix ODEs where an orthogonal initial condition implies orthogonality of the solution for all time. Here we present first and second order numerical methods for which the property of orthogonality-preservation is always carried through to the discrete approximation. To our knowledge, these are the first methods that guarantee to preserve orthogonality, without the use of projection, whenever it is preserved by the flow. The methods are based on Gauss–Legendre Runge–Kutta formulas, which are known to preserve orthogonality on a restricted problem class. In addition, the new methods are linearly-implicit, requiring only the solution of one or two linear matrix systems (of the same dimension as the solution matrix) per step. Illustrative numerical tests are reported.

Keywords: Geometric integration; Implicit midpoint rule; ODEs on manifolds; Orthogonality; Structure preservation

1. Background

Suppose that $Y(t) \in \mathbb{R}^{m\times m}$ is a time-varying matrix that satisfies the autonomous initial value ordinary differential equation (ODE)

$$Y'(t) = G(Y(t)), \quad t > 0, \ Y(0) \text{ given}, \tag{1.1}$$

where $G : \mathbb{R}^{m\times m} \to \mathbb{R}^{m\times m}$ is smooth. (Note that the problem (1.1) could be re-written as a standard system of ODEs by stacking the columns of $Y(t)$ in a vector in $\mathbb{R}^{m^2}$. For our purposes, however, it is natural to regard the solution as a matrix.)

Recall that a real square matrix A is said to be orthogonal if $A^{\mathrm{T}}A = I$, where I is the identity matrix. Certain applications, including algorithms for computing Lyapunov exponents, continuous orthonormalisation and continuous singular value decomposition [2,7], give rise to Eq. (1.1) where

[1] The work of this author was supported by the Engineering and Physical Sciences Research Council of the UK under grants GR/H94634 and GR/K80228. Current address: Department of Mathematics, University of Strathclyde, Glasgow G1 1XH, Scotland.

PII S0168-9274(96)00033-5

$Y(0)$ is orthogonal and it is known that $Y(t)$ remains orthogonal for all $t > 0$. This prompts the following definition.

Definition 1.1. The problem (1.1) is said to preserve orthogonality if

$$Y(0)^{\mathrm{T}}Y(0) = I \implies Y(t)^{\mathrm{T}}Y(t) = I \text{ for all } t > 0.$$

The following classical result that characterises this property may be found, for example, in [5, Theorem 1].

Theorem 1.2. *The problem* (1.1) *preserves orthogonality if and only if* $G(Y) = F(Y)Y$ *for some function* $F : \mathbb{R}^{m\times m} \to \mathbb{R}^{m\times m}$ *that maps orthogonal matrices to skew-symmetric matrices; that is,* $Y^{\mathrm{T}}Y = I \Rightarrow F(Y)^{\mathrm{T}} = -F(Y)$.

Orthogonality-preservation is a qualitative property, and in line with other properties such as symplecticness, periodicity and reversibility (see, for example, [1,3,8]) it is of interest to identify numerical methods that share this property. This motivates the following definition.

Definition 1.3. Given $Y_0 = Y(0)$, suppose that a numerical method for (1.1) computes approximations $Y_n \approx Y(t_n)$, where $0 = t_0 < t_1 < t_2 < \cdots$. The method is said to preserve orthogonality if $Y_n^{\mathrm{T}}Y_n = I$ for all n whenever $Y(0)^{\mathrm{T}}Y(0) = I$ and the problem preserves orthogonality. In other words, from Theorem 1.2, the method preserves orthogonality if

$$\left.\begin{aligned} Y_0^{\mathrm{T}}Y_0 &= I \\ G(Y) &= F(Y)Y \\ Y^{\mathrm{T}}Y = I \Rightarrow F(Y)^{\mathrm{T}} &= -F(Y) \end{aligned}\right\} \implies Y_n^{\mathrm{T}}Y_n = I \text{ for all } n.$$

2. New methods

In this section we consider one-step discretisation methods for computing approximations $Y_n \approx Y(t_n)$ in (1.1). For convenience, we assume that a fixed step size $h > 0$ is used, so that $t_n = nh$. However, we point out that the property of orthogonality-preservation carries through directly to the case where local error control is applied in a variable time-stepping algorithm.

Recall that the implicit midpoint rule (IMR) for (1.1) is

$$Y_{n+1} = Y_n + hG\left(\tfrac{1}{2}(Y_n + Y_{n+1})\right).$$

It can also be written in more traditional Runge–Kutta form as

$$K_1 = G\left(Y_n + \tfrac{1}{2}hK_1\right), \tag{2.1}$$

$$Y_{n+1} = Y_n + hK_1. \tag{2.2}$$

The IMR is a one-stage, second order, implicit Runge–Kutta formula, and it is a member of the Gauss–Legendre class.

Now, using the splitting $G(Y) = F(Y)Y$, we consider the following linearly-implicit variation of the IMR:

$$Y_{n+1} = Y_n + hF(Y_n)\tfrac{1}{2}(Y_n + Y_{n+1}), \tag{2.3}$$

or, mimicking the traditional Runge–Kutta form,

$$K_1 = F(Y_n)\big(Y_n + \tfrac{1}{2}hK_1\big), \tag{2.4}$$

$$Y_{n+1} = Y_n + hK_1. \tag{2.5}$$

We see from (2.3), or from (2.4)–(2.5), that computing Y_{n+1} from Y_n requires the solution of a linear matrix equation; that is, m linear systems of dimension m, with the same coefficient matrix, $I - \frac{1}{2}hF(Y_n)$. We are interested in the case where $F(Y_n)$ is skew-symmetric and hence has purely imaginary eigenvalues. In this case $I - \frac{1}{2}hF(Y_n)$ is guaranteed to be nonsingular.

Theorem 2.1. *The method* (2.3) *is of order one and preserves orthogonality.*

Proof. Our proof is based on the ideas in [5, Theorem 2]. Assume that the ODE preserves orthogonality, and that $Y(0)^{\mathrm{T}}Y(0) = I$. Suppose, for induction, that $Y_n^{\mathrm{T}}Y_n = I$. Using (2.4)–(2.5), we have

$$Y_{n+1}^{\mathrm{T}}Y_{n+1} = Y_n^{\mathrm{T}}Y_n + h\big(Y_n^{\mathrm{T}}K_1 + K_1^{\mathrm{T}}Y_n\big) + h^2K_1^{\mathrm{T}}K_1.$$

Therefore, we need to show that

$$Y_n^{\mathrm{T}}K_1 + K_1^{\mathrm{T}}Y_n + hK_1^{\mathrm{T}}K_1 = 0. \tag{2.6}$$

Let $\Psi_1 = Y_n + \frac{1}{2}hK_1$. Then

$$Y_n^{\mathrm{T}}K_1 = \Psi_1^{\mathrm{T}}K_1 - \tfrac{1}{2}hK_1^{\mathrm{T}}K_1, \tag{2.7}$$

$$K_1^{\mathrm{T}}Y_n = K_1^{\mathrm{T}}\Psi_1 - \tfrac{1}{2}hK_1^{\mathrm{T}}K_1. \tag{2.8}$$

Now Y_n is orthogonal and so $F(Y_n)^{\mathrm{T}} = -F(Y_n)$. Hence,

$$\Psi_1^{\mathrm{T}}K_1 + K_1^{\mathrm{T}}\Psi_1 = \Psi_1^{\mathrm{T}}F(Y_n)\Psi_1 + \Psi_1^{\mathrm{T}}F(Y_n)^{\mathrm{T}}\Psi_1 = \Psi_1^{\mathrm{T}}\big(F(Y_n) + F(Y_n)^{\mathrm{T}}\big)\Psi_1 = 0. \tag{2.9}$$

Adding (2.7) to (2.8) and using (2.9) gives (2.6) as required.

We see from (2.4) that $K_1 = F(Y_n)Y_n + \mathrm{O}(h)$, and hence, in (2.5),

$$Y_{n+1} = Y_n + hG(Y_n) + \mathrm{O}\big(h^2\big).$$

This shows that the method matches the first two terms in the Taylor expansion of $Y(t_n + h)$, and hence is at least first order. It is clear from the numerical tests in the next section that the order is no higher than one. □

We note that the IMR does not preserve orthogonality in the sense of Definition 1.3. Following the proof in Theorem 2.1 it can be shown that orthogonality-preservation for the IMR holds if

$$F(\Psi_1)^{\mathrm{T}} = -F(\Psi_1), \tag{2.10}$$

where $\Psi_1 = Y_n + \frac{1}{2}hK_1$ in (2.1). The condition (2.10) follows if Ψ_1 is orthogonal, which requires $K_1^{\mathrm{T}}K_1 = 0$. However, this condition is not true in general. The numerical tests in the next section confirm that the IMR does not preserve orthogonality in this generality.

It is true, however, that the IMR preserves orthogonality when a further restriction is placed on F: if $G(Y) = F(Y)Y$ where F maps *all* matrices into skew-symmetric matrices, then (2.10) holds and hence $Y_n^{\mathrm{T}} Y_n = I$ for all n. In fact, Iserles and Zanna [5] give an elegant proof that this property is shared by all Runge–Kutta methods satisfying the "$M = 0$" condition, which is an algebraic constraint on the Runge–Kutta coefficients. In particular, all Gauss–Legendre methods have $M = 0$. (A slightly weaker version of this result was established by Dieci, Russell and van Vleck [2], and Wright [9] considered a related property.)

Theorem 2.1 shows that it is possible to construct an orthogonality-preserving method by adapting the IMR in a way that exploits the natural splitting $G(Y) = F(Y)Y$ arising from Theorem 1.2. However, this reduces the order of the method from two to one. Second order can be recovered by "bootstrapping" the argument that is supplied to F in (2.4) as follows (with $F_n = F(Y_n)$)

$$K_{1/2} = F_n\big(Y_n + \tfrac{1}{4}hK_{1/2}\big), \tag{2.11}$$

$$K_1 = F\big(Y_n + \tfrac{1}{2}hK_{1/2}\big)\big(Y_n + \tfrac{1}{2}hK_1\big), \tag{2.12}$$

$$Y_{n+1} = Y_n + hK_1. \tag{2.13}$$

Note that (2.11) and (2.12) represent linear matrix equations that must be solved for $K_{1/2}$ and K_1, respectively.

Theorem 2.2. *The method* (2.11)–(2.13) *is of order two and preserves orthogonality.*

Proof. The orthogonality-preservation can be established in two stages, both of which follow directly from the proof in Theorem 2.1. First, we see that $Y_n + \frac{1}{2}hK_{1/2}$ is orthogonal, since it is simply the result of applying method (2.4)–(2.5) over a step of length $h/2$. Then, since (2.12)–(2.13) has the same structure as (2.4)–(2.5) it follows that $Y_n + hK_1$ is orthogonal.

To find the order of the method, we first expand

$$K_{1/2} = \big(I - \tfrac{1}{4}hF_n\big)^{-1} F_n Y_n = F_n Y_n + \mathrm{O}(h),$$

and

$$F\big(Y_n + \tfrac{1}{2}hK_{1/2}\big) = F_n + F'[Y_n]\big(\tfrac{1}{2}hK_{1/2}\big) + \mathrm{O}\big(h^2\big) = F_n + \tfrac{1}{2}hF'[Y_n](F_nY_n) + \mathrm{O}\big(h^2\big).$$

(Here $F'[Y_n](X)$ denotes the Frechet derivative of F at the point Y_n, operating on X.) Hence,

$$\begin{aligned} K_1 &= \big(I - \tfrac{1}{2}hF_n + \mathrm{O}\big(h^2\big)\big)^{-1}\big(F_n + \tfrac{1}{2}hF'[Y_n](F_nY_n) + \mathrm{O}\big(h^2\big)\big)Y_n \\ &= \big(I + \tfrac{1}{2}hF_n + \mathrm{O}\big(h^2\big)\big)\big(F_n + \tfrac{1}{2}hF'[Y_n](F_nY_n) + \mathrm{O}\big(h^2\big)\big)Y_n \\ &= F_n + \tfrac{1}{2}h\big(F_n^2 + F'[Y_n](F_nY_n)\big)Y_n + \mathrm{O}\big(h^3\big). \end{aligned}$$

Finally, in (2.13),

$$Y_{n+1} = \big(I + hF_n + \tfrac{1}{2}h^2\big(F_n^2 + F'[Y_n](F_nY_n)\big)\big)Y_n + \mathrm{O}\big(h^3\big).$$

This agrees with the Taylor expansion of $Y(t_n + h)$ up to $\mathrm{O}(h^3)$, as required. It is clear from the numerical tests in the next section that the order is no higher than two. □

We emphasise that the new methods (2.4)–(2.5) and (2.11)–(2.13) are linearly-implicit in the sense that they only require either one or two linear matrix equations to be solved per step, respectively. The IMR, by contrast, requires a (generally) nonlinear system to be solved on each step.

The idea of adapting the Gauss–Legendre methods to exploit the splitting in Theorem 1.2 could be pursued to higher order. Given coefficients $\{a_{lj}\}_{l,j=1}^{\nu}$ and $\{b_l\}_{l=1}^{\nu}$ we could consider the Runge–Kutta based scheme

$$\Psi_l = Y_n + h\sum_{j=1}^{\nu} a_{lj}K_j, \tag{2.14}$$

$$K_l = F(U_l)\Psi_l, \tag{2.15}$$

$$Y_{n+1} = Y_n + h\sum_{l=1}^{\nu} b_l K_l. \tag{2.16}$$

Here U_l is some approximation to $Y(t_n + c_l h_n)$, where $c_l = \sum_{l=1}^{\nu} a_{lj}$. If the U_l are orthogonal then it follows from the proof of [5, Theorem 2] that this scheme preserves orthogonality when the Gauss–Legendre coefficients are used. However, aside from the difficulty of achieving high order with this approach, it should be noted that computing $\{\Psi_l\}_{l=1}^{\nu}$ involves the solution of ν coupled linear matrix equations—this can be written as m linear systems of dimension $m\nu$ that share the same coefficient matrix. In this case, a careful study would be needed to establish whether these customized methods are competitive with the more straightforward *projection* approach where discrete approximations from a standard time-stepping method are perturbed to nearby orthogonal matrices [2,4].

We also point out that (2.14)–(2.16) can be identified with *a single iteration* of the scheme proposed in [2, Theorem 4.1]. However, we emphasise that iterating to convergence, so that the Gauss–Legendre method is implemented, does not lead to orthogonality-preservation in the full generality of Definition 1.3.

3. Numerical tests

To illustrate the properties of the methods discussed in the last section, we present some numerical results. Two problems of the form (1.1) with $G(Y) = F(Y)Y$ and $m = 4$ were solved, using

Problem 1: $F(Y) = (Y\exp(Y) - (Y\exp(Y))^{\mathrm{T}})/2$,

Problem 2: $F(Y) = \big(Y\exp(Y) - (Y\exp(Y))^{\mathrm{T}}\big)/2 + (Y^{\mathrm{T}}Y - I)/10$.

Note that for Problem 1 we have $F(Y)^{\mathrm{T}} = -F(Y)$ for all Y, and hence the IMR preserves orthogonality. This is not the case for Problem 2, although $Y^{\mathrm{T}}Y = I \Rightarrow F(Y)^{\mathrm{T}} = -F(Y)$, as required in Theorem 1.2.

In implementing the IMR, we solved the nonlinear system at each step with the fixed point iteration

$$Y_{n+1}^{[v+1]} = Y_n + hG\big(\tfrac{1}{2}\big(Y_n + Y_{n+1}^{[v]}\big)\big), \quad v = 0, 1, 2, \ldots .$$

We used an Euler predictor, so $Y_{n+1}^{[0]} = Y_n + hG(Y_n)$, and the terminating criterion was

$$\big\|Y_{n+1}^{[v+1]} - Y_{n+1}^{[v]}\big\|_2 \leqslant 10^{-12}.$$

(Here $\|\cdot\|_2$ denotes the Euclidean norm.)

The problems were solved over $0 \leqslant t \leqslant 20$, using an arbitrary orthogonal $Y(0)$. (More precisely, we took $Y(0)$ to be the matrix `Q` from `[Q,R] = qr(magic(4))` in Matlab [6].) Each method was

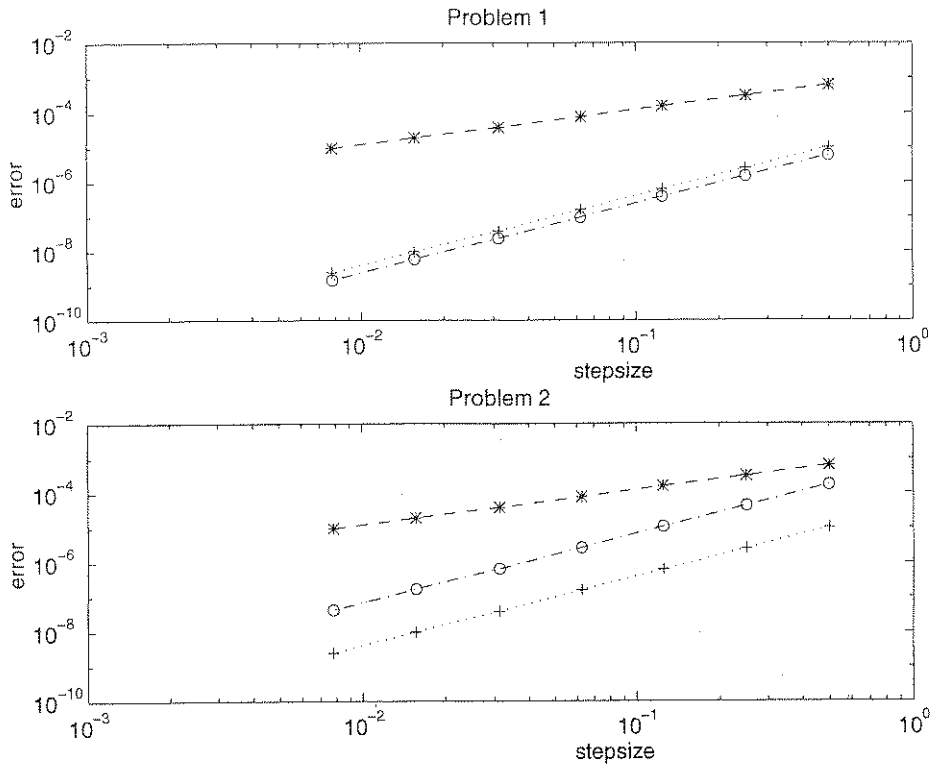

Fig. 1. Errors against step size on Problems 1 and 2. See text for details.

implemented with a fixed step size of $1/2$, $1/4$, $\ldots$, $1/256$. Letting Y_N^h denote the approximation $Y_N^h \approx Y(20)$ with a step size of h, the quantities

$$\mathrm{ge}^h := \left\| Y_N^h - Y_N^{h/2} \right\|_2,$$

which estimate the global error in Y_N^h, were recorded for each method. Fig. 1 records these values for Problems 1 (upper picture) and 2 (lower picture). The * symbol denotes the first order method (2.4)–(2.5), the + symbol denotes the second order method (2.11)–(2.13) and the o symbol denotes the IMR. In each case the slope of the line is consistent with the order of the method.

We also monitored the departure from orthogonality; that is,

$$\gamma^h := \min \left\{ \|E\|_2 \colon\ E \in \mathbb{R}^{m\times m},\ \left(Y_N^h + E\right)^{\mathrm{T}} \left(Y_N^h + E\right) = I \right\},$$

which can be computed via a polar decomposition of Y_N^h [4]. For the new methods on both problems, and for the IMR on Problem 1, γ^h was consistent with roundoff error, as expected. For the IMR on Problem 2, where orthogonality-preservation cannot be guaranteed, a plot of γ^h was almost indistinguishable from a plot of ge^h. Note that, since the true solution is orthogonal, the departure from orthogonality is trivially bounded by the global error. For the IMR on Problem 2 the trivial bound is sharp, which emphasises the lack of orthogonality in the numerical solution.

Acknowledgements

I thank David Griffiths and Paul Smith for discussions during the course of this work, and I thank Nick Higham for commenting on the original manuscript.

References

[1] B. Cano and J.M. Sanz-Serna, Error growth in the numerical integration of periodic orbits, with application to Hamiltonian and reversible systems, Technical Report 1, University of Valladolid (1995).
[2] L. Dieci, R.D. Russell and E.S. van Vleck, Unitary integrators and applications to continuous orthonormalization techniques, *SIAM J. Numer. Anal.* 31 (1994) 261–281.
[3] E. Hairer and D. Stoffer, Reversible long-term integration with variable step sizes, *SIAM J. Sci. Stat. Comput.*, to appear.
[4] D.J. Higham, Time-stepping and preserving orthonormality, Technical Report NA/161, University of Dundee, Scotland (1995).
[5] A. Iserles and A. Zanna, Qualitative numerical analysis of ordinary differential equations, Technical Report NA5, Department of Applied Mathematics and Theoretical Physics, University of Cambridge; also in: Lectures in Applied Mathematics (Amer. Math. Soc., Providence, RI, 1995).
[6] *MATLAB User's Guide* (The MathWorks, Natick, MA, 1992).
[7] V. Mehrmann and W. Rath, Numerical methods for the computation of analytic singular value decompositions, *Electron. Trans. Numer. Anal.* 1 (1993) 72–88.
[8] J.M. Sanz-Serna and M.P. Calvo, *Numerical Hamiltonian Problems* (Chapman and Hall, London, 1994).
[9] K. Wright, Numerical solution of differential equations for the analytic singular value decomposition, Technical Report Computing Science 405, University of Newcastle upon Tyne (1992).