Computing the matrix sine and cosine simultaneously with a reduced number of products

Muaz Seydaoğlu [a,c,*], Philipp Bader [b], Sergio Blanes [c], Fernando Casas [d]

[a] *Faculty of Art and Science, Department of Mathematics, 49100 Mus, Turkey*
[b] *Departament de Matemàtiques, Universitat Jaume I, 12071 Castellón, Spain*
[c] *Instituto de Matemática Multidisciplinar, Universitat Politècnica de València, E-46022 Valencia, Spain*
[d] *IMAC and Departament de Matemàtiques, Universitat Jaume I, 12071 Castellón, Spain*

A R T I C L E   I N F O

A B S T R A C T

A new procedure is presented for computing the matrix cosine and sine simultaneously by means of Taylor polynomial approximations. These are factorized so as to reduce the number of matrix products involved. Two versions are developed to be used in single and double precision arithmetic. The resulting algorithms are more efficient than schemes based on Padé approximations for a wide range of norm matrices.

© 2021 IMACS. Published by Elsevier B.V. All rights reserved.

1. Introduction

Many dynamical systems are modeled by differential equations in which finding closed solutions is not possible and so one has to compute approximating solutions. These differential equations usually preserve some underlying geometric structure which reflects the qualitative nature of the phenomena they describe. It is then relevant that the approximations share with the exact solution of the differential equation these qualitative properties to render a right description. The design and analysis of numerical integrators preserving some of these geometric structures constitutes the realm of Geometric Numerical Integration, an active and interdisciplinary research area and the subject of intensive development during the last decades [6,11,16,19,21,24].

Exponential integrators can be considered as a class of geometric integrators tailored to stiff and oscillatory equations [7,8,14,15,17]. For large systems of equations these schemes usually require to compute the action of the exponential of a matrix on a vector [14,15]. However, for problems of moderate size it may be more appropriate to compute directly the exponential of the matrices involved.

When the problem is oscillatory, very often the formal solution involves both the sine and cosine of a matrix. Thus, for example, consider the Schrödinger equation in quantum mechanics,

$$i\frac{d\psi}{dt} = \mathcal{H}(t)\psi, \qquad \psi(t_0) = \psi_0,$$

---

* Corresponding author at: Faculty of Art and Science, Department of Mathematics, 49100 Mus, Turkey.
E-mail addresses: m.seydaoglu@alparslan.edu.tr (M. Seydaoğlu), bader@uji.es (P. Bader), serblaza@imm.upv.es (S. Blanes), casas@uji.es (F. Casas).

where $\mathcal{H}(t)$ is a Hermitian operator and ψ is a complex wave function. A usual procedure to get numerical approximations involves first a spatial discretization or working on a finite dimensional representation. In any event, one ends up with a matrix equation with a similar structure,

$$i\frac{du}{dt} = Au, \qquad u(t_0) = u_0 \in \mathbb{C}^N.$$

If A is a real and constant matrix, the unitary evolution operator is given by

$$U(t) = e^{-itA} = \cos(tA) - i\sin(tA). \tag{1}$$

There are different techniques to compute efficiently the exponential of a matrix [2,3,5,13,22,26–28]. However, using any of these general algorithms to approximate the unitary matrix e^{-itA} in (1) involves products of complex matrices making them computationally expensive. Alternatively, we propose an efficient procedure to compute the matrix sine and cosine that only involves a small number of products of real matrices. The algorithm is used in combination with squaring as

$$\cos(2A) = 2\cos^2(A) - I = I - 2\sin^2(A), \qquad \sin(2A) = 2\sin(A)\cos(A).$$

In this way, it only requires two products per squaring (instead of four products when considering the square of complex matrices), thus making the overall procedure more efficient.

There are other examples where the computation of the sine and cosine of a matrix can be of interest. For example, for wave equations given by the generic second order system

$$y'' + Ay = f(y, t),$$

with $y \in \mathbb{R}^N$, exponential integrators frequently require to solve separately the linear homogeneous problem

$$y'' + Ay = 0, \qquad y(0) = y_0, \quad y'(0) = y'_0. \tag{2}$$

Writing (2) as a first order system, the solution is given by

$$\begin{pmatrix} y(t) \\ y'(t) \end{pmatrix} = e^{tM} \begin{pmatrix} y_0 \\ y'_0 \end{pmatrix}, \qquad \text{with} \quad M = \begin{pmatrix} 0 & I \\ -A & 0 \end{pmatrix} \tag{3}$$

and

$$e^{tM} = \begin{pmatrix} \cos(t\sqrt{A}) & (\sqrt{A})^{-1}\sin(t\sqrt{A}) \\ -\sqrt{A}\sin(t\sqrt{A}) & \cos(t\sqrt{A}) \end{pmatrix} \equiv \begin{pmatrix} c(t^2 A) & s(t, A) \\ -As(t, A) & c(t^2 A) \end{pmatrix}. \tag{4}$$

Notice that the dimension of M is twice the dimension of A and so the cost of matrix-matrix multiplications grows, in general, by a factor of eight.

On the other hand, a closer look to the functions to be approximated clearly indicates that the algorithm used to evaluate the matrix sine and cosine for the unitary matrix (1) should not be used directly for this problem, since it requires computing first the square root of the matrix, $B = \sqrt{A}$, in addition to a multiplication and an inversion of this matrix. As a matter of fact, an efficient approximation to the exponential (4) was already presented in [4]. We propose in this case an improved algorithm based on a modification of the methods to compute the matrix sine and cosine with the goal of computing simultaneously the functions

$$c(t^2 A) \equiv \cos(\sqrt{t^2 A}) \qquad \text{and} \qquad s(t, A) \equiv (\sqrt{A})^{-1}\sin(\sqrt{t^2 A}).$$

For the double angle we will take into account that

$$c(4t^2 A) = 2c^2(t^2 A) - I, \qquad s(2t, A) = 2s(t, A)c(t^2 A),$$

thus requiring only two products per squaring. Notice that we do not use the property $\cos(2A) = I - 2\sin^2(A)$, since the function $\sin(A)$ is not computed in this case.

In summary, the purpose of this paper consists in developing algorithms that allow one to compute $\cos(A)$ and $\sin(A)$ or $c(t^2 A)$ and $s(t, A)$ simultaneously and providing full accuracy up to single or double precision with a reduced computational cost. Thus, in particular, we propose an algorithm that, with only six products, reproduces exactly the series defining $\cos(A)$ up to the power A^{16} and the series of $\sin(A)$ up to A^{17}. In this sense, the following definition will be helpful for the rest of the paper:

Definition 1. We say that a given function $f(A)$ satisfies $f(A) = \mathcal{O}(A^n)$ if it can be written as a convergent Taylor expansion, $f(A) = \sum_{k=n}^{\infty} c_k A^k$, for $\|A\| < \alpha$, with α a positive constant.

Thus, our algorithm allows one to write $\cos(A)$ with and error $\mathcal{O}(A^{18})$. The same procedure allows one, with one extra product (seven products in total), to approximate $\cos(A)$ and $\sin(A)$ with errors $\mathcal{O}(A^{26})$ and $\mathcal{O}(A^{25})$, respectively.

Although one can find in the literature several algorithms to compute $\cos(A)$ (see [26] and references therein), only few of them are designed to do so in a simultaneous way (see [1] and references therein). As our analysis shows and several numerical examples confirm, the technique we propose here outperform all of them.

2. The algorithms

The search of fast algorithms for evaluating matrix polynomials has received considerable interest in the literature [2,3, 18,20,23,25,29,30]. We next briefly summarize how to approximate the matrix sine and cosine functions by means of certain polynomials involving a reduced number of matrix products. This reduction essentially follows the same approach used in [9] to minimize the number of commutators appearing in different Lie-group integrators and was successfully adapted to the Taylor expansion of the exponential matrix in [2] and especially in [3].

Generally speaking, the strategy consists first in elaborating a recursive procedure to compute the polynomial approximating the matrix cosine with the minimum number of products and then using these same products to approximate the matrix sine as accurately as possible in the cheapest possible way.

Clearly, the most economic way to construct polynomials of degree 2^k is by applying the following sequence, which requires the evaluation of only k products. First we form the intermediate matrices

$$A_0 := I, \quad A_1 := A,$$
$$A_2 := z_{2,0}I + z_{2,1}A_1 + (x_1I + x_2A_1)(x_3I + x_4A_1),$$
$$A_4 := z_{4,0}I + z_{4,1}A_1 + z_{4,2}A_2 + (x_5I + x_6A_1 + x_7A_2)(x_8I + x_9A_1 + x_{10}A_2),$$
$$A_8 := \sum_{k=0}^{3} z_{8,2^{k-1}}A_{2^{k-1}} + (x_{11}I + \cdots + x_{14}A_4)(x_{15}I + \cdots + x_{18}A_4), \tag{5}$$
$$\vdots$$

and finally we take

$$P_{2^k} = A_{2^k},$$

where the coefficients $z_{i,j}, x_k \in \mathbb{R}$ have to be chosen in such a way that P_{2^k} coincides with the desired polynomial. Here the indices 2^k are chosen to indicate the highest attainable power, i.e., $A_{2^k} \in \mathbb{P}_{2^k}(A)$, where $\mathbb{P}_n(x)$ is the set of polynomials of degree $\leq n$ in the variable x. Of course, there are many redundancies in the coefficients since some of them can be absorbed by others. Thus, for instance, we may take just $A_2 = A_1A_1$, since any quadratic polynomial can be written in terms of A_0, A_1 and A_2. By the same token, we can take $A_4 = A_2(x_1A_1 + x_2A_2)$ because any polynomial of degree four can be expressed in terms of A_0, A_1, A_2 and A_4 for an appropriate choice of the coefficients x_1 and x_2. However, the problem is much more involved for higher degree polynomials.

According to the previous considerations, any polynomial of degree up to four can be computed with two products, whereas polynomials up to degree eight can be computed with only three products. This does not mean, however, that all such polynomials can be written with just three products. Thus, in particular, $P_7(A) = A^7$ requires at least four products where one of the intermediate results must be a polynomial of odd order (e.g. $A_2 = A^2$, $A_3 = AA_2$, $A_4 = A_2A_2$, $A_7 = A_4A_3 = P_7(A)$).

When a given polynomial cannot be reproduced by following the previous approach, new terms have to be incorporated. Thus, in particular

$$A_0 = I, \quad A_1 = A, \quad A_2 = A^2, \quad A_3 = AA_2$$
$$A_6 = B_{3,1} + B_{3,2}B_{3,3}, \qquad B_{3,i} = \sum_{k=0}^{3} x_{i,k}A_k \tag{6}$$
$$\vdots$$

and this generalizes the procedure.

We use this technique in the sequel to approximate first $\cos(A)$ and $\sin(A)$ simultaneously with the minimum number of products, and then we apply the same procedure to $c(t^2A)$ and $s(t, A)$.

2.1. Computing $\cos(A)$ and $\sin(A)$ simultaneously

Let us denote by

$$T_{2m}^c(A) = \sum_{k=0}^{m} \frac{(-1)^k (A^2)^k}{(2k)!}, \qquad T_{2m+1}^s(A) = A \sum_{k=0}^{m} \frac{(-1)^k (A^2)^k}{(2k+1)!}$$

the Taylor polynomial approximations of $\cos(A)$ and $\sin(A)$ up to order $2m$ and $2m+1$ in A, respectively, and by $T_{2m+1,\ell}^s$ with $\ell > 2m+1$, any polynomial of degree ℓ such that

$$T_{2m+1,\ell}^s(A) = T_{2m+1}^s(A) + \mathcal{O}(A^{2m+2}).$$

We next collect the best approximations we have obtained to the cosine and sine functions by polynomial functions when using an increasing number of products k.

$k = 3$ **products.** This constitutes a trivial problem, but it nevertheless illustrates the general procedure. With two products we can compute T_4^c as

$$A_2 = A^2, \qquad A_4 = A^4,$$
$$T_4^c(A) = I - \frac{1}{2!}A_2 + \frac{1}{4!}A_4, \tag{7}$$

and with one extra product we get

$$T_5^s(A) = A(I - \frac{1}{3!}A_2 + \frac{1}{5!}A_4). \tag{8}$$

$k = 4$ **products.** With three products we can compute T_8^c as

$$A_2 = A^2, \qquad A_4 = A_2^2, \qquad A_8 = A_4\left(-\frac{1}{6!}A_2 + \frac{1}{8!}A_4\right),$$
$$T_8^c(A) = I - \frac{1}{2!}A_2 + \frac{1}{4!}A_4 + A_8. \tag{9}$$

With one extra product we can approximate the matrix sine, but only up to order seven:

$$T_{7,9}^s(A) = A\left(I - \frac{1}{3!}A_2 + \frac{1}{5!}A_4 + \frac{6!}{7!}A_8\right). \tag{10}$$

According with the previous notation, $T_{7,9}^s(A) = T_7^s(A) + \mathcal{O}(A^9)$.

The order of approximation of the matrix sine can be increased by incorporating one extra product as follows:

$$A_8 = A_4\left(-\frac{1}{7!}A_2 + \frac{1}{9!}A_4\right),$$
$$T_9^s(A) = A\left(I - \frac{1}{3!}A_2 + \frac{1}{5!}A_4 + A_8\right). \tag{11}$$

In this way, $\sin(A) - T_9^s(A) = \mathcal{O}(A^{11})$.

$k = 6$ **products.** The following scheme allows one to express $T_{16}^c(A)$ with only four products:

$$A_2 = A^2, \qquad A_4 = A_2^2, \qquad A_8 = A_4(x_1 A_2 + x_2 A_4),$$
$$A_{16} = (x_3 A_4 + A_8)(x_4 I + x_5 A_2 + x_6 A_4 + x_7 A_8),$$
$$T_{16}^c(A) = I - \frac{1}{2}A_2 + x_8 A_4 + A_{16}. \tag{12}$$

In fact, we get two families of solutions depending on a free parameter that reproduce $T_{16}^c(A)$ in exact arithmetic. Since floating-point arithmetic has to be used in practice, we choose this free parameter x_1 so as to minimize the 1-norm of the vector $(x_1, \ldots, x_8)$ in order to avoid large round off errors. This results in

$$x_1 = \frac{7}{500}, \qquad\qquad\qquad x_2 = -\frac{7}{60000}, \qquad x_3 = \frac{1}{2500}(-1533 + 7\sqrt{36681}),$$

$$x_4 = -\frac{5(124581 + 391\sqrt{36681})}{10594584}, \quad x_5 = \frac{9775}{10594584}, \quad x_6 = -\frac{5(1001 + \sqrt{36681})}{508540032}, \tag{13}$$

$$x_7 = \frac{3125}{889945056}, \qquad\qquad\qquad x_8 = \frac{1549211 + 3246\sqrt{36681}}{63063000}.$$

Some of the coefficients are irrational numbers because they correspond to solutions of a nonlinear system of equations.

With two extra products we can approximate the matrix sine up to order $\mathcal{O}(A^{19})$ as follows:

$$C_{24} = (z_5 I + z_5 A_2 + z_6 A_4 + z_7 A_8 + z_8 T_{16}^c(A))A_8,$$
$$T_{17,25}^s(A) = A\left(z_0 I + z_1 A_2 + z_2 A_4 + z_3 A_8 + z_4 T_{16}^c(A) + C_{24}\right) \tag{14}$$

with

$$z_0 = \frac{8887}{4794}, \qquad z_1 = -\frac{1897}{3196}, \qquad z_2 = \frac{25259}{575280},$$

$$z_3 = -\frac{965093875}{9674368704}, \quad z_4 = -\frac{4093}{4794}, \qquad z_5 = \frac{25698275}{29023106112}, \tag{15}$$

$$z_6 = -\frac{3907675}{348277273344}, \quad z_7 = \frac{11865625}{3656911370112}, \quad z_8 = \frac{25}{308756448}.$$

In other words, (14)-(15) approximates the matrix sine up to a higher order than the matrix cosine.

$k = 7$ **products.** With five products we can compute T_{24}^c:

$$A_2 = A^2, \qquad A_4 = A_2^2, \qquad A_6 = A_4 A_2,$$
$$C_1 = a_{0,1} I + a_{1,1} A_2 + a_{2,1} A_4 + a_{3,1} A_6, \quad C_2 = a_{0,2} I + a_{1,2} A_2 + a_{2,2} A_4 + a_{3,2} A_6,$$
$$C_3 = a_{0,3} I + a_{1,3} A_2 + a_{2,3} A_4 + a_{3,3} A_6, \quad C_4 = a_{0,4} I + a_{1,4} A_2 + a_{2,4} A_4 + a_{3,4} A_6, \tag{16}$$
$$A_{12} = C_3 + C_4^2, \qquad A_{24} = (C_2 + A_{12})A_{12}$$
$$T_{24}^c(A) = C_1 + A_{24}.$$

The best solution we have obtained is:

$$
\begin{array}{ll}
a_{0,1} = 0, & a_{1,1} = 0, \\
a_{2,1} = 0.02264979811206039519, & a_{3,1} = -0.00013110924142135755, \\
a_{0,2} = 0.55751443809990408029, & a_{1,2} = -0.61577924683458386455, \\
a_{2,2} = 0.00747198841446687051, & a_{3,2} = -0.00003362444420476012, \\
a_{0,3} = 0.75936877868464999248, & a_{1,3} = -0.01560333979813817129, \\
a_{2,3} = 0.00010936989591908396, & a_{3,3} = -1.03893360877457159499 \cdot 10^{-6}, \\
a_{0,4} = 0, & a_{1,4} = -0.03964996874347447309, \\
a_{2,4} = 0.00015549007350382146 3, & a_{3,4} = -1.12673966307117002248 8 \cdot 10^{-6}.
\end{array}
\tag{17}
$$

Although we report here 20 digits for the coefficients, they can be in fact determined with arbitrary accuracy.

With two extra products we can approximate the matrix sine up to order $\mathcal{O}(A^{23})$ as follows:

$$C_{48} = (z_6 I + z_7 A_2 + z_8 A_4 + z_9 A_6 + z_{10} A_{12} + z_{11} T_{24}^c(A))T_{24}^c(A),$$
$$T_{23,49}^s(B) = A\left(z_0 I + z_1 A_2 + z_2 A_4 + z_3 A_6 + z_4 A_{12} + z_5 T_{24}^c(A) + C_{48}\right), \tag{18}$$

with

$$
\begin{array}{ll}
z_0 = 0.10090808375109885598, & z_1 = -0.07668753546445299316, \\
z_2 = 0.00084924846993243257, & z_3 = -0.00001220406904464391, \\
z_4 = 0.98499703159318860027, & z_5 = -0.84925233648155398756, \\
z_6 = 1, & z_7 = 0.00095544138280925799, \\
z_8 = 4.56337109377154270633 \cdot 10^{-6}, & z_9 = 2.73461259403000427141 \cdot 10^{-8}, \\
z_{10} = 0.00048550288474842477 & z_{11} = -4.15891109384923342531 \cdot 10^{-7}.
\end{array}
\tag{19}
$$

2.2. Computing $c(t^2 A)$ and $s(t, A)$ simultaneously

Given a real matrix A, let us denote by

$$P_m^c(t^2 A) = \sum_{k=0}^m \frac{(-1)^k (t^2 A)^k}{(2k)!}, \qquad P_m^s(t, A) = t\sum_{k=0}^m \frac{(-1)^k (t^2 A)^k}{(2k+1)!}$$

the Taylor expansions of the functions

$$c(t^2 A) = \cos(\sqrt{t^2 A}), \qquad \text{and} \qquad s(t, A) = (\sqrt{A})^{-1}\sin(\sqrt{t^2 A})$$

up to order m in A, respectively. Notice that they are approximations up to order $2m$ and $2m + 1$ in t to the respective functions. Analogously, we will denote by $P_{m,\ell}^s$, $\ell > m$, any polynomial of degree ℓ such that $P_{m,\ell}^s(A) = P_m^s(A) + \mathcal{O}(A^{m+1})$.

Next we show how the previous algorithms to approximate the sine and cosine functions can be adjusted to approximate $c(t^2 A)$ and $s(t, A)$. As before, we proceed according with the number of products involved.

$k = 3$ **products.** With two products we can compute $P_4^c(t^2 A)$:

$$B = t^2 A, \qquad B_2 = B^2, \qquad B_4 = B^2(-\frac{1}{6!}B + \frac{1}{8!}B_2),$$
$$P_4^c(B) = I - \frac{1}{2!}B + \frac{1}{4!}B_2 + B_4. \tag{20}$$

With the same number of products we can also evaluate $P_{3,4}^s(t, A)$,

$$P_{3,4}^s(t, A) = t\left(I - \frac{1}{3!}B + \frac{1}{5!}B_2 - \frac{6!}{7!}B_4\right), \tag{21}$$

whereas with one extra product we get

$$P_4^s(t, A) = t\left(I - \frac{1}{3!}B + \frac{1}{5!}B_2 + B_2\left(-\frac{1}{7!}B + \frac{1}{9!}B_2\right)\right). \tag{22}$$

$k = 4$ **products.** With three products we can compute $P_8^c(t^2 A)$:

$$B = t^2 A, \qquad B_2 = B^2, \qquad B_4 = B_2(x_1 B + x_2 B_2),$$
$$B_8 = (x_3 B_2 + B_4)(x_4 I + x_5 B + x_6 B_2 + x_7 B_4),$$
$$P_8^c(B) = y_0 I + y_1 B + y_2 B_2 + B_8, \tag{23}$$

whose coefficients are the same as those given in (13).

With one extra product we can approximate the matrix sine up to order eight as

$$C_{12} = (z_5 I + z_5 B + z_6 B_2 + z_7 B_4 + z_8 P_8^c(B))B_4,$$
$$P_{8,12}^s(t, A) = t\left(z_0 I + z_1 B + z_2 B_2 + z_3 B_4 + z_4 P_8^c(B) + C_{12}\right), \tag{24}$$

with the same values for the coefficients z_i as before.

$k = 5$ **products.** With four products we can compute $P_{12}^c(t^2 A)$:

$$B = t^2 A, \qquad B_2 = B^2, \qquad B_3 = B_2 B,$$
$$D_1 = a_{0,1} I + a_{1,1} B + a_{2,1} B_2 + a_{3,1} B_3, \quad D_2 = a_{0,2} I + a_{1,2} B + a_{2,2} B_2 + a_{3,2} B_3,$$
$$D_3 = a_{0,3} I + a_{1,3} B + a_{2,3} B_2 + a_{3,3} B_3, \quad D_4 = a_{0,4} I + a_{1,4} B + a_{2,4} B_2 + a_{3,4} B_3,$$
$$B_6 = D_3 + D_4^2$$
$$P_{12}^c(B) = D_1 + (D_2 + B_6)B_6, \tag{25}$$

with solution for the coefficients $a_{i,j}$ given in (17), whereas with one extra product we can approximate $P_{11}^s(t, A)$ as

$$C_{24} = (z_6 I + z_7 B + z_8 B_2 + z_9 B_3 + z_{10} B_6 + z_{11} P_{12}^c(A))P_{12}^c(B),$$
$$P_{11,24}^s(t, A) = t\left(z_0 I + z_1 B + z_2 B_2 + z_3 B_3 + z_4 B_6 + z_5 P_{12}^c(B) + C_{24}\right), \tag{26}$$

with the same coefficients as in (19).

2.3. Padé approximations

At this point it is useful to briefly review the schemes presented in [1] to compute the matrix sine and cosine simultaneously, since they will be compared in section 4 with our own procedure.

The methods presented in [1] are based on the identities

$$\cos(A) = \frac{e^{iA} + e^{-iA}}{2}, \qquad \sin(A) = \frac{e^{iA} - e^{-iA}}{2i},$$

and the use of Padé approximations of the exponential e^{iA}. For instance, taking a diagonal Padé of order eight for approximating e^{iA}, i.e. $r_4(iA) = [p_4(-iA)]^{-1} p_4(iA) = e^{iA} + \mathcal{O}(A^9)$ one gets

$$s_4 = \frac{A\left(I - \frac{11}{8}A^2 + \frac{37}{1176}A^4 - \frac{1}{70560}A^6\right)}{I + \frac{1}{28}A^2 + \frac{3}{3920}A^4 + \frac{1}{8}A^6 + \frac{1}{2822400}A^8},$$

$$c_4 = \frac{I - \frac{13}{28}A^2 + \frac{289}{11760}A^4 - \frac{19}{70560}A^6 + \frac{19}{2822400}A^8}{I + \frac{1}{28}A^2 + \frac{3}{3920}A^4 + \frac{1}{8}A^6 + \frac{1}{2822400}A^8},$$

with

$$s_4 = \sin(A) + \mathcal{O}(A^9), \qquad c_4 = \cos(A) + \mathcal{O}(A^{10}).$$

It is clear that s_4, c_4 can be computed simultaneously with 5 products (A^2, A^4, A^6, A^8, and the extra product for the numerator in s_4) and the computation of two inverse matrices. Since both denominators are the same, only one LU factorization is necessary. The totals cost is $(7 + \frac{1}{3})$ products. Notice that the same order (with very similar accuracy as we will see) is obtained with our novel approach at the cost of only 4 products (and a smaller number of matrices need to be stored).

3. Error analysis

Next we analyze how to bound the truncation errors of the previously considered Taylor polynomial approximations of order $2m$ and $2\tilde{m} + 1$ for the cosine and sine functions, respectively. They have the form

$$\cos(A) - T_{2m}^c = \sum_{k=m+1}^{\infty} \alpha_{2k} A^{2k}, \qquad 2m \in \{4, 8, 16, 24\}$$

$$\sin(A) - T_{2\tilde{m}+1,\ell}^s = \sum_{k=\tilde{m}+1}^{\infty} \tilde{\alpha}_{2k+1} A^{2k+1}, \qquad 2\tilde{m} + 1 \in \{5, 7, 17, 23\}. \tag{27}$$

On the other hand, the truncation errors of the approximations of the cosine and sine functions obtained by using Padé approximants for e^{iA} [1] can be written as

$$\cos(A) - c_m = \sum_{k=m+1}^{\infty} \gamma_{2k} A^{2k}, \qquad \sin(A) - s_m = \sum_{k=m}^{\infty} \hat{\gamma}_{2k+1} A^{2k+1}. \tag{28}$$

Clearly, the series (27) and (28) can be bounded in terms of $\|A\|$ as

$$\left\| \cos(A) - T_{2m}^c \right\| \le \sum_{k=m+1}^{\infty} |\alpha_{2k}| \, \|A\|^{2k}, \qquad \left\| \sin(A) - T_{2\tilde{m}+1,\ell}^s \right\| \le \sum_{k=\tilde{m}+1}^{\infty} |\tilde{\alpha}_{2k+1}| \, \|A\|^{2k+1}, \tag{29}$$

and

$$\left\| \cos(A) - c_m \right\| \le \sum_{k=m+1}^{\infty} |\gamma_{2k}| \, \|A\|^{2k}, \qquad \|\sin(A) - s_m\| \le \sum_{k=m}^{\infty} |\hat{\gamma}_{2k+1}| \, \|A\|^{2k+1}. \tag{30}$$

We denote by θ_{2m}^M the largest value of $\|A\|$ such that the bounds (29), (30) do not exceed a prescribed accuracy, u, for each method $M \equiv T_{2m}^c, T_{2\tilde{m}+1}^s, c_m, s_m$. To achieve maximum accuracy, we bound the previous forward absolute errors with the unit round off $u = 2^{-53}$, $u = 2^{-24}$ in double and single precision floating-point arithmetic, respectively. We have truncated the series of the corresponding functions after 150 terms to find θ_{2m}^M. The corresponding values for the new Taylor approximations of the cosine and sine functions are collected in Tables 1 and 2, together with the total number of matrix products corresponding to each procedure, Π_{2m}. For completeness, we have also included the values of θ_{2m}^M for the Padé approximations, as given in [1]. In this case we have also added to the total number of matrix products π_m the cost of evaluating two inverse matrices sharing the same LU factorization, i.e. $(2 + \frac{1}{3})$ products.

The comparison of the theoretical performance of the new Taylor polynomial approximations $T_{2m}^c, T_{2\tilde{m}+1}^s$ (with orders $\{4, 8, 16, 24\}$ and $\{5, 7, 17, 23\}$ respectively) and the Padé approximations c_m, s_m [1] (with orders $\{4, 8, 16, 24\}$) has been illustrated in Fig. 1: here we plot $\|A\|$ versus the number of matrix products required for each approximation of $\cos(A)$ and $\sin(A)$ simultaneously, both in double (left) and single (right) precision. From the figure the improvement achieved by the proposed Taylor polynomial approximations is apparent.

Remark 2. Notice that the previous error analysis is also valid for the approximations to the function $c(t^2 A)$ just by replacing $\|A\|$ by $t\sqrt{\|A\|}$ in the first expression of eq. (29). Since this result depends on both t and $\|A\|$, it is not appropriate to get error bounds simultaneously for both $s(t, A)$ and $As(t, A)$ of eq. (4). Nevertheless, in the frequent case in which A is a real and symmetric matrix, then e^{tM} in (4) is a symplectic matrix and $\|\sin(t\sqrt{A})\| \le 1$. In this case the values of $\tilde{\alpha}_{2k+1}$ appearing in the second expression of (29) are the same when considering the relative error for both functions $s(t, A)$ and $As(t, A)$.

Table 1

Number of matrix multiplications Π_{2m} and forward absolute error bounds θ_{2m}^M in double precision floating-point arithmetic, $u \leq 2^{-53}$, for the new Taylor algorithms T_{2m}^c, $T_{2\tilde{m}+1}^s$ and Padé approximations c_m, s_m [1]. The cost of the computation of two inverse matrices sharing the same LU factorization, i.e. $(2 + \frac{1}{3})$, has been included in the cost π_m for the Padé approximations.

$2\tilde{m}$	4	6	16	22
$2m$	4	8	16	24
$\theta_{2m}^{c_m}$	6.5633e-3	1.3959e-1	1.3879	3.7288
$\theta_{2m}^{s_m}$	2.4019e-3	1.1213e-1	1.3784	3.7287
π_m [1]	$\mathbf{5 + \frac{1}{3}}$	$\mathbf{7 + \frac{1}{3}}$	$\mathbf{10 + \frac{1}{3}}$	$\mathbf{12 + \frac{1}{3}}$
$\theta_{2m}^{T_{2m}^c}$	6.5633e-3	1.1495e-1	9.8108e-1	2.5675
$\theta_{2m}^{T_{2\tilde{m}+1}^s}$	1.777e-2	8.0438e-2	1.1184	1.97
Π_{2m}	**3**	**4**	**6**	**7**

Table 2

Same as Table 1, but now in single precision floating-point arithmetic.

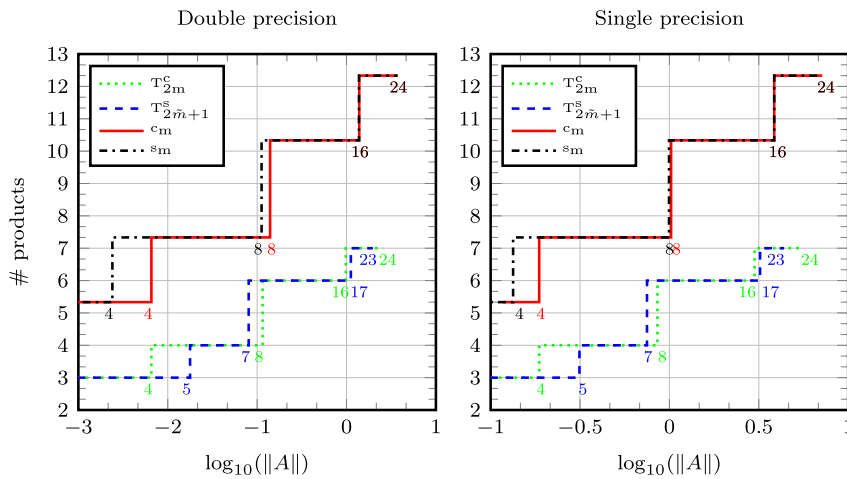| $2\tilde{m}$	4	6	16	22
$2m$	4	8	16	24
$\theta_{2m}^{c_m}$	1.8687e-1	1.0218	3.8571	7.1575
$\theta_{2m}^{s_m}$	1.3355e-1	9.9511e-1	3.8569	7.1575
π_m [1]	$\mathbf{5 + \frac{1}{3}}$	$\mathbf{7 + \frac{1}{3}}$	$\mathbf{10 + \frac{1}{3}}$	$\mathbf{12 + \frac{1}{3}}$
$\theta_{2m}^{T_{2m}^c}$	1.8709e-1	8.5756e-1	2.9935	5.5555
$\theta_{2m}^{T_{2\tilde{m}+1}^s}$	3.1386e-1	7.492e-1	3.2152	4.3819
Π_{2m}	**3**	**4**	**6**	**7**

Fig. 1. Orders and corresponding number of products of each method versus $\|A\|$ in double and single precision floating-point arithmetic.

4. Numerical experiments

In this section we collect the results we have obtained by measuring the performance of the new Taylor polynomials (denoted as 'cosmsinmT') and the Padé approximations (denoted as 'cosmsinmP') [1] to compute the matrix cosine and sine functions simultaneously. The platform for all numerical experiments is MATLAB R2013a and the matrix 1-norm has been used in the implementation of the algorithms. The experiments have been carried out for three sets of 2500 matrices of dimension $\leq N \times N$ each one, with $N = 16, 64$ and 1024. Each set of 2500 matrices is chosen as follows (the matrices are adjusted in order to have different norms):

- *Matrix gallery (blue)*: 52 test matrices have been chosen from the MATLAB gallery function [12]. 690 sampled matrices with different norms were tested. If the gallery does not return valid results for some dimension then the dimension is reduced. Furthermore, the dimension is reduced to 4×4 and 8×8 for matrices with norm larger than 10^4.
- *Random matrices (red)*: Using rand() and randn() functions in MATLAB to randomly generate matrices with entries drawn from different distributions. 400 matrices normally distributed, 500 matrices uniformly distributed in the interval
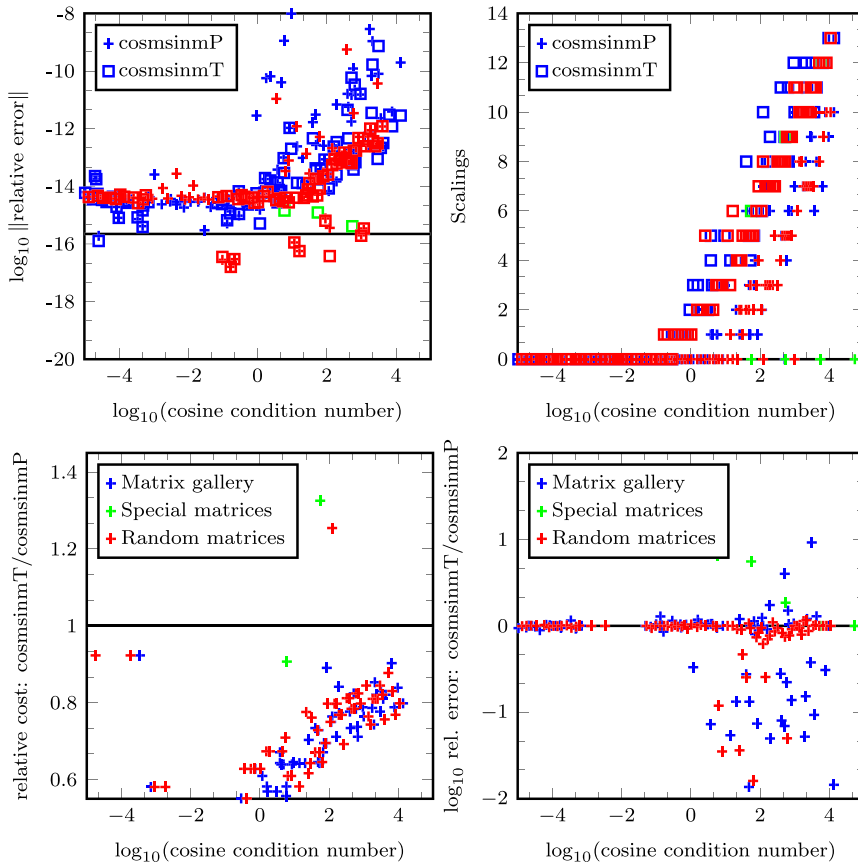
Fig. 2. Comparison of cosmsinmP and cosmsinmT by using results of the cosine function of the matrices of dimensions $\leq 16 \times 16$. Test Matrices are specified as Matrix gallery (blue), Random matrices (red) and Special matrices (green). Top left panel: relative errors of the methods. Top right panel: number of scalings used versus the condition number for each matrix. Bottom left panel: ratio of the costs. Bottom right panel: ratio of the relative errors. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

$(0, 1)$ and 501 matrices in the interval $(-0.5, 0.5)$. Additionally, using `spdiags()` and `rand()` functions in MATLAB to construct 400 triangular nilpotent matrices with random rank.

• *Special matrices (green)*: 9 matrices of the form

$$A = \begin{pmatrix} 1 & \lambda \\ 0 & -1 \end{pmatrix}, \tag{31}$$

where $\lambda = 1, 10, \ldots, 10^8$, possibly leading to *overscaling* (utilization of large value of the scaling parameter s).

The same test matrices as in Remark 5 of [3] have been generated and all matrices are adjusted to have 1-norms on the interval $(10^{-4}, 10^{4.1})$ in all numerical experiments. Notice that, given a matrix A, we first compute $\|A\|_1$ in order to decide, according to the results shown in Table 1 (for double precision) or Table 2 (for single precision), which method is used to approximate the functions as well as whether the scaling-squaring is required.

The condition numbers of each matrix function are computed by running the function `funm_condest1` from the Matrix Function Toolbox [12]. The reference solutions of the matrix cosine and sine have been calculated with *Mathematica* with 100 digits of precision. In all cases we have computed the relative error

$$\frac{\|F - f(A)\|_2}{\|f(A)\|_2},$$

where F is an approximated value of $f(A)$. In the following we show the results for double precision (similar results are obtained for single precision).

The results we have obtained for the 2500 matrices of dimension $\leq 16 \times 16$ are collected in Figs. 2 and 3. From the top left of these two figures we observe that in most cases the relative errors of both cosmsinmP and cosmsinmT methods produced in approximating the matrix cosine and sine functions change between $1.0e - 12$ and $1.0e - 15$, and they drop below machine accuracy for a few matrices, whereas the top right panels show that cosmsinmT method involves more
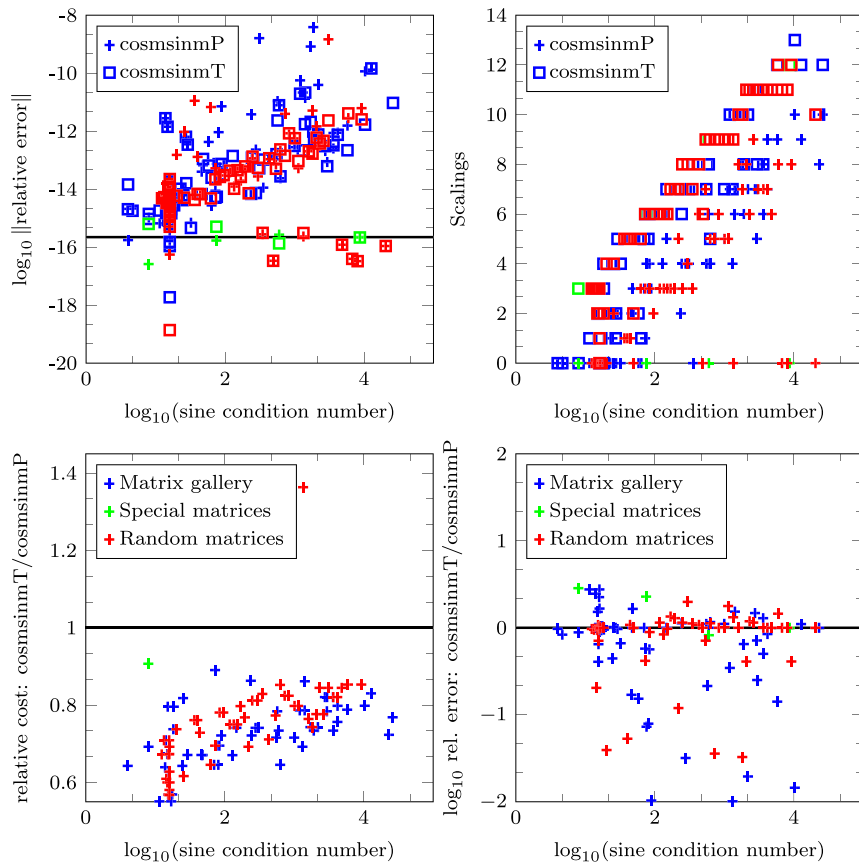
Fig. 3. Same as Fig. 2, but for the matrix sine function.

scalings (in particular, cosmsinmT leads to scaling for 1180 matrices, whereas cosmsinmP does so for 741 matrices). On the other hand, the bottom left panels show that in most cases the ratios of the cost cosmsinmT/cosmsinmP are below 1, with the new method cosmsinmT requiring fewer matrix products. Finally, it can be seen in the bottom right panels that the accuracy of both methods is in good agreement with the theoretical results we have obtained (here the values of the relative errors have been replaced by machine accuracy when they are smaller).

Figs. 4 and 5 show performance profiles [10] for the sets of matrices of dimension $\leq 64 \times 64$, and $\leq 1024 \times 1024$, respectively. A performance profile shows the probability that a measure is within a given factor of the best measure across all algorithms. Specifically, we show in Fig. 4 the relative errors (for both the cosine and the sine), the number of products and the computational time for matrices of dimension $\leq 64 \times 64$, whereas in Fig. 5 we only report the number of products and the computational times required. The abscissae are the digits of accuracy (for the relative error plots) and the interval $[1, 3]$ for the number of products and computation times.

From these figures it is clear that, in particular, cosmsinmT method is quite less expensive than cosmsinmP.

5. Conclusions

We have presented a new algorithm to compute the matrix cosine and sine. The algorithm contains several methods that are optimized for different values of the norm of the matrix and the desired accuracy, and can be combined with the scaling and squaring technique. Each of these methods is obtained by following a sequence in which each stage uses the results from all previous ones. An error analysis is also carried out and we have shown both theoretically as well as in the numerical experiments that the new algorithm is superior to other procedures from the literature that are based on Padé approximations to the matrix cosine and sine.

The new algorithm only involves matrix-matrix products and does not require to compute the inverse of matrices as is the case for Padé approximations. The cost of computing the inverse of a dense matrix can be taken as 4/3 the cost of the product of two dense matrices. However, for sparse matrices, the computational cost of the proposed algorithms grow almost linearly, whereas the cost of Padé approximations grows much faster because, in general, the inverse of a sparse matrix is a dense matrix. We thus expect a further increase in the efficiency of the new algorithms in this case.
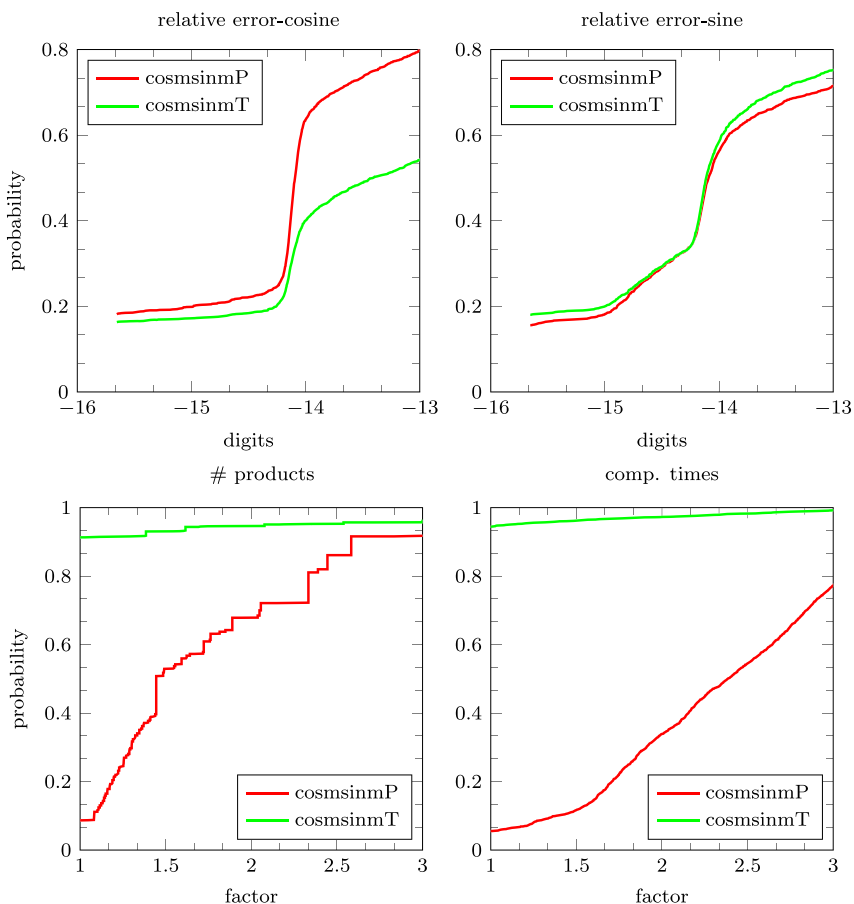
Fig. 4. Performance profiles of cosmsinmP and cosmsinmT by using results of the cosine and sine function of the matrices of dimension $\leq 64 \times 64$. Top panels: the percentage of matrices among the test set for which the relative logarithmic errors are lower than the horizontal axis. Bottom left panel: the percentage of matrices among the test set for which the method under consideration has a number of product lower than *factor*-times the smallest number of product across all algorithms. Bottom right panel: Restatement of bottom left panel in terms of the computation times averaged across 10 cycles.
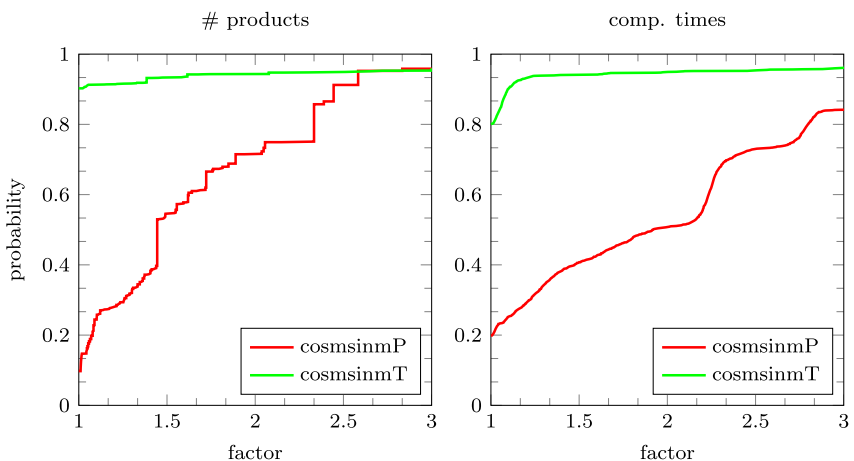

Fig. 5. Same as Fig. 4, bottom panels, for the set of matrices of dimension $\leq 1024 \times 1024$.

Acknowledgements

References

[1] A.H. Al-Mohy, N.J. Higham, S.D. Relton, New algorithms for computing the matrix sine and cosine separately or simultaneously, SIAM J. Sci. Comput. 37 (2015) A456–A487.

[2] P. Bader, S. Blanes, F. Casas, An improved algorithm to compute the exponential of a matrix, arXiv:1710.10989 [math.NA], 2017, preprint.

[3] P. Bader, S. Blanes, F. Casas, Computing the matrix exponential with an optimized Taylor polynomial approximation, Mathematics 7 (2019) 1174, https://doi.org/10.3390/math7121174.

[4] P. Bader, S. Blanes, E. Ponsoda, M. Seydaoğlu, Symplectic integrators for the matrix Hill's equation and its applications to engineering models, J. Comput. Appl. Math. 316 (2017) 47–59.

[5] P. Bader, S. Blanes, M. Seydaoğlu, The scaling, splitting and squaring method for the exponential of perturbed matrices, SIAM J. Matrix Anal. Appl. 36 (2015) 594–614.

[6] S. Blanes, F. Casas, A Concise Introduction to Geometric Numerical Integration, CRC Press, Boca Raton, FL, USA, 2016.

[7] S. Blanes, F. Casas, A. Murua, An efficient algorithm based on splitting for the time integration of the Schrödinger equation, J. Comput. Phys. 303 (2015) 396–412.

[8] S. Blanes, F. Casas, J.A. Oteo, J. Ros, The Magnus expansion and some of its applications, Phys. Rep. 470 (2009) 151–238.

[9] S. Blanes, F. Casas, J. Ros, High order optimized geometric integrators for linear differential equations, BIT Numer. Math. 42 (2002) 262–284.

[10] E.D. Dolan, J.J. More, Benchmarking optimization software with performance profiles, Math. Program. 91 (2002) 201–213.

[11] E. Hairer, C. Lubich, G. Wanner, Geometric Numerical Integration. Structure-Preserving Algorithms for Ordinary Differential Equations, 2nd ed, Springer, Berlin, 2006.

[12] N.J. Higham, Functions of Matrices: Theory and Computation, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2008.

[13] N.J. Higham, A.H. Al-Mohy, Computing matrix functions, Acta Numer. 19 (2010) 159–208.

[14] M. Hochbruck, C. Lubich, On Krylov subspace approximations to the matrix exponential operator, SIAM J. Numer. Anal. 34 (1997) 1911–1925.

[15] M. Hochbruck, A. Ostermann, Exponential integrators, Acta Numer. 19 (2010) 209–286.

[16] A. Iserles, A First Course in the Numerical Analysis of Differential Equations, 2nd ed., Cambridge University Press, 2008.

[17] A. Iserles, H.Z. Munthe-Kaas, S.P. Nørsett, A. Zanna, Lie group methods, Acta Numer. 9 (2000) 215–365.

[18] L. Lei, T. Nakamura, A fast algorithm for evaluating the matrix polynomial $I + A + \cdots + A^{N-1}$, IEEE Trans. Circuits Syst. I, Fundam. Theory Appl. 39 (1992) 299–300.

[19] B. Leimkuhler, S. Reich, Simulating Hamiltonian Dynamics, Cambridge University Press, 2004.

[20] W. Liang, R. Baer, C. Saravanan, Y. Shao, A.T. Bell, M. Head-Gordon, Fast methods for resumming matrix polynomials and Chebyshev matrix polynomials, J. Comput. Phys. 194 (2004) 575–587.

[21] C. Lubich, From Quantum to Classical Molecular Dynamics: Reduced Models and Numerical Analysis, European Mathematical Society, 2008.

[22] C.B. Moler, C.F. Van Loan, Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later, SIAM Rev. 45 (2003) 3–49.

[23] M.S. Paterson, L.J. Stockmeyer, On the number of nonscalar multiplications necessary to evaluate polynomials, SIAM J. Comput. 2 (1973) 60–66.

[24] J.M. Sanz-Serna, M.P. Calvo, Numerical Hamiltonian Problems, Chapman & Hall, London, 1994.

[25] J. Sastre, Efficient evaluation of matrix polynomials, Linear Algebra Appl. 539 (2018) 229–250.

[26] J. Sastre, J. Ibáñez, P. Alonso-Jordá, J. Peinado, E. Defez, Fast Taylor polynomial evaluation for the computation of the matrix cosine, J. Comput. Appl. Math. 354 (2019) 641–650.

[27] J. Sastre, J. Ibáñez, E. Defez, Boosting the computation of the matrix exponential, Appl. Math. Comput. 340 (2019) 206–220.

[28] R.B. Sidje, Expokit: a software package for computing matrix exponentials, ACM Trans. Math. Softw. 24 (1998) 130–156.

[29] C. Van Loan, A note on the evaluation of matrix polynomials, IEEE Trans. Autom. Control 24 (1979) 320–321.

[30] D. Westreich, Evaluating the matrix polynomial $I + A + \cdots + A^{N-1}$, IEEE Trans. Circuits Syst. 36 (1989) 162–164.