

PARALLEL-IN-TIME MAGNUS INTEGRATORS

B. T. KRULL AND M. L. MINION *

Abstract. Magnus integrators are a subset of geometric integration methods for the numerical solution of ordinary differential equations that conserve certain invariants in the numerical solution. This paper explores temporal parallelism of Magnus integrators, particularly in the context of nonlinear problems. The approach combines the concurrent computation of matrix commutators and exponentials within a time step with a pipelined iteration applied to multiple time steps in parallel. The accuracy and efficiency of time parallel Magnus methods up to order six are highlighted through numerical examples and demonstrate that significant parallel speedup is possible compared to serial methods.

Key words. Ordinary differential equations, nonlinear ordinary differential equations, Magnus expansions, isospectral flows, Lax pairs, parallel-in-time

AMS subject classifications. 34L30, 65L05, 65Y05

1. Introduction. The solution of ordinary differential equations (ODEs) is a well established field with applications across the spectrum of scientific disciplines. Numerical methods date back at least to Euler’s work in 1768 [5], and the accuracy, stability, and efficiency of various methods is well studied (see for example Refs. [8, 9]). In more recent years, the study of specialized numerical methods for ODEs that preserve certain mathematical properties of the numerical solution has seen increased interest. Examples include methods for Hamiltonian systems that numerically conserve invariants of the true dynamical system such as the energy or angular momentum. More generally, the true solution of an ODE poised in N -dimensional space may reside for all time on a manifold \mathcal{M} of dimension $d < N$, and the goal is to devise a method for which the numerical solution will also remain on \mathcal{M} . Such methods are referred to in general as *geometric integrators*. The interested reader is encouraged to consult Ref. [7] for a comprehensive introduction to the subject.

As a concrete example, consider the ODE given by

$$(1.1) \quad \frac{d}{dt}Y(t) = F(Y), \quad Y(0) = Y_0,$$

where Y and $F(Y)$ are both $N \times N$ matrices. Depending on the form of F , certain properties of the initial condition Y_0 may be preserved for all time, such as the determinant, orthogonality, idempotency, or the spectrum of eigenvalues. In general, standard numerical methods such as linear multistep or Runge-Kutta methods will not produce solutions that conserve such properties [16].

Magnus integrators are a subset of geometric integrators based on the expansion proposed by Wilhelm Magnus in 1954. The Magnus expansion is closely tied to the concept of Lie groups and algebras due to the presence of matrix commutators, also known as Lie brackets. The discussion concerning solutions to Eq. (1.1) can in fact be generalized to differential equations on Lie groups (see e.g. Ref. [12]), however in this work we strictly use matrix valued solutions. Magnus integrators can also be viewed as a type of exponential integrator (see e.g. Ref. [11] for a review) since they require that the matrix exponential be evaluated. The comprehensive review of the Magnus expansion and applications by Blanes, et. al. [1] provides

*Center for Computational Sciences and Engineering, Lawrence Berkeley National Laboratory, Berkeley, CA, 94720.

a description of several physical applications to which the Magnus expansion has been applied including nuclear, atomic, and molecular dynamics; nuclear magnetic resonance; quantum field theory and high energy physics; electromagnetism; optics; geometric control of mechanical systems; and the search for periodic orbits.

In general, numerical Magnus integrators are constructed by applying quadrature rules of suitable order to a truncation of the Magnus expansion. The review [12] presents various types of Magnus integrators up to sixth-order, and methods of order up to eight are considered in Ref. [2]. As the order increases, the number of commutator terms required in the Magnus expansion grows quickly, hence in these papers and others, a detailed discussion of how to minimize the number of commutators required for a given order is presented. So called *commutator-free* Magnus integrators (e.g. Refs. [3, 18]) have also been proposed that replace the need to compute matrix commutators with additional matrix exponentials. This can reduce the total computational cost of the method depending on the relative cost of computing commutators versus matrix exponentials. In this paper, an additional avenue for reducing the time to solution for Magnus integrators is investigated, namely parallelization in time.

The study of parallel numerical methods for ODEs dates back at least to Nievergelt in 1964 [17], and the field has seen an increase in activity in the last 15 years. (See Ref. [6] for a recent review.) The standard classification of parallel methods for ODEs includes parallelism across the method, across the problem, and across the time steps [4]. In this work we demonstrate the utility of parallelization across both the method and the time steps for Magnus integrators for both linear and nonlinear equations. Special attention is paid to schemes for solving nonlinear differential equations for isospectral flow problems, although the methodology that is described can be applied more generally.

The remainder of this paper is organized as follows. Section 2 presents the mathematical background behind the Magnus expansion and Magnus integrators followed by some specific Magnus integrators based on Gaussian collocation in section 3. The parallelization strategies for these integrators is presented in section 4. Numerical results comparing the efficiency of different parallel methods is presented in section 5 using the Toda lattice as the prototypical numerical example. The results demonstrate that significant parallel speedup over serial methods is possible and show that parallel higher-order methods are superior in terms of accuracy and time to solution compared to lower-order methods.

2. Mathematical Preliminaries. In this section, a review of the mathematics behind the construction of Magnus integrators for both linear and nonlinear problems is reviewed.

2.1. Matrix Calculus and Differential Equations. Given a constant matrix $A \in \mathbb{F}^{N \times N}$, where \mathbb{F} can be either \mathbb{R} or \mathbb{C} , the exponential of A is defined by the power series

$$(2.1) \quad e^A = \sum_{n=0}^{\infty} \frac{A^n}{n!}.$$

Given the time-dependent vector $y(t) \in \mathbb{F}^N$, the solution to the differential equation

$$(2.2) \quad y'(t) = Ay(t), \quad y(0) = y_0$$

is given by

$$(2.3) \quad y(t) = e^{At} y_0.$$

This is easily shown by differentiating the power series definition (2.1) of the exponential on the right hand side term-by-term to give

$$(2.4) \quad \frac{d}{dt}e^{At} = Ae^{At}.$$

Now consider the more general case of a time-dependent matrix $A(t)$. Again differentiating definition (2.1) term-by-term and using the product rule yields

$$(2.5) \quad \frac{d}{dt}e^{A(t)} = A'(t) + \frac{A(t)A'(t) + A'(t)A(t)}{2!} + \frac{A(t)'A(t)^2 + A(t)A'(t)A(t) + A'(t)A(t)^2}{3!} \dots$$

The right hand side of (2.5) can be rearranged to give

$$(2.6) \quad \frac{d}{dt}e^{A(t)} = \text{dexp}_{A(t)}(A'(t))e^{A(t)},$$

where the operator on the right is defined by

$$(2.7) \quad \text{dexp}_{A(t)}(A'(t)) = A'(t) + \frac{[A(t), A'(t)]}{2!} + \frac{[A(t), [A(t), A'(t)]]}{3!} \dots,$$

and brackets $[\cdot, \cdot]$ denote the matrix commutator $[X, Y] = XY - YX$.

Now consider the linear system of ODEs

$$(2.8) \quad y'(t) = A(t)y(t), \quad y(0) = y_0.$$

To find a solution, suppose that it can be written in the form

$$(2.9) \quad y(t) = e^{\Omega(t)}y_0$$

for some matrix $\Omega(t)$. Using Eq. (2.6), $\Omega(t)$ satisfies

$$(2.10) \quad \text{dexp}_{\Omega(t)}(\Omega'(t)) = A(t), \quad \Omega(0) = 0.$$

The same formal derivation can be applied to a nonlinear system of equations

$$(2.11) \quad y'(t) = A(y(t), t)y(t), \quad y(0) = y_0.$$

Again, if the solution to this equation is to take the form of Eq. (2.9), then $\Omega(t)$ satisfies

$$(2.12) \quad \text{dexp}_{\Omega(t)}(\Omega'(t)) = A(y(t), t), \quad \Omega(0) = 0.$$

In the next section, methods for finding $\Omega(t)$ are considered.

2.2. Magnus Expansion. In 1954, Magnus introduced an explicit expression for the solution of Eq. (2.8), which is reviewed here. The first step is the inversion of the operator $\text{dexp}_{\Omega(t)}$, which gives a differential equation for $\Omega(t)$

$$(2.13) \quad \Omega'(t) = \sum_{n=0}^{\infty} \frac{B_n}{n!} \text{ad}_{\Omega(t)}^n(A(t)), \quad \Omega(0) = 0,$$

where the B_n are the Bernoulli numbers and

$$(2.14) \quad \text{ad}_{\Omega}^k(A) = [\Omega, \text{ad}_{\Omega}^{k-1}A], \quad \text{ad}_{\Omega}(A) = \Omega A - A\Omega, \quad \text{ad}_{\Omega}^0(A) = A.$$

Next, a Picard-type iteration is applied to Eq. (2.13)

$$(2.15) \quad \Omega^{k+1}(t) = \int_0^t \sum_{n=0}^{\infty} \frac{B_n}{n!} \text{ad}_{\Omega^k}^n(A(t)).$$

Collecting terms in Eq. (2.15) yields an infinite series for $\Omega(t)$

$$(2.16) \quad \Omega(t) = \Omega^{(1)}(t) + \Omega^{(2)}(t) + \Omega^{(3)}(t) + \dots,$$

where

$$(2.17) \quad \Omega^{(1)}(t) = \int_0^t d\tau A(\tau)$$

$$(2.18) \quad \Omega^{(2)}(t) = \frac{1}{2} \int_0^t d\tau_2 \int_0^{\tau_2} d\tau_1 [A(\tau_2), A(\tau_1)]$$

$$(2.19) \quad \Omega^{(3)}(t) = \frac{1}{6} \int_0^t d\tau_3 \int_0^{\tau_3} d\tau_2 \int_0^{\tau_2} d\tau_1 [A(\tau_3), [A(\tau_2), A(\tau_1)]] + [[A(\tau_3), A(\tau_2)], A(\tau_1)]$$

$$(2.20) \quad \begin{aligned} \Omega^{(4)}(t) = \frac{1}{12} \int_0^t d\tau_4 \int_0^{\tau_4} d\tau_3 \int_0^{\tau_3} d\tau_2 \int_0^{\tau_2} d\tau_1 & [[[A(\tau_4), A(\tau_3)], A(\tau_2)], A(\tau_1)] \\ & + [A(\tau_4), [[A(\tau_3), A(\tau_2)], A(\tau_1)]] \\ & + [A(\tau_4), [A(\tau_3), [A(\tau_2), A(\tau_1)]]] \\ & + [A(\tau_3), [A(\tau_2), [A(\tau_1), A(\tau_4)]]]. \end{aligned}$$

Each subsequent term in the series contains an additional integration operator, commutators of one higher order, as well as an increasing number of commutator terms. The reader is referred to the original work of Magnus [14] or the extensive review [1] for further details. To summarize, the Magnus expansion gives an explicit formula for the solution of the linear equation given by Eq. (2.8) in the form of the exponential of a matrix defined by an infinite series given by Eq. (2.16).

2.3. The Magnus Expansion for Nonlinear Problems. The same formal procedure used in the last section to construct the solution to the linear problem Eq. (2.8) can also be applied to the nonlinear system Eq. (2.11). One can still represent the solution in terms of the exponential of the function $\Omega(t)$, and the only difference is that in the Magnus expansion terms given above in Eqs. (2.17)-(2.20), the terms $A(\tau)$ must be replaced with $A(y(\tau), \tau)$, which by the definition of the solution is $A(e^{\Omega(\tau)}y_0, \tau)$. Although this may appear at first a small change in notation, the implication is quite important. For the nonlinear problem, the Magnus expansion does not give an explicit formula for the function $\Omega(t)$ as in the linear case. Instead, the result is an equation for $\Omega(t)$ involving an infinite expansion of terms containing integrals of commutators dependent on $\Omega(t)$. The central insight of this paper is that this equation for $\Omega(t)$ can be solved efficiently by a fixed point iteration that is readily amenable to parallelization in the time direction.

2.4. Isospectral Flows. A special type of matrix differential equation for which the eigenvalues of the solution are independent of time is called *isospectral flow*. Problems of this form exist in application domains including electronic structure, wave dynamics, and linear algebra. Isospectral flow is often associated with the concept of a Lax pair: two matrices or operators dependent on time that satisfy the Lax equation

$$(2.21) \quad Y'(t) = [A(Y, t), Y], \quad Y(0) = Y_0,$$

where $Y(t), A(Y, t) \in \mathbb{F}^{N \times N}$.

It is straightforward to show that the solution to Eq. (2.21) can be written in the form of the transformation

$$(2.22) \quad Y(t) = (e^{\Omega(t)})Y_0(e^{\Omega(t)})^{-1},$$

where $\Omega(t)$ is defined by the Magnus expansion. Since the form of $Y(t)$ takes on a similarity transformation, the eigenvalues of $Y(t)$ do not change in time, hence the term *isospectral*. In the special case where A is Hermitian (or self-adjoint) and Y is skew-Hermitian (or skew-adjoint), the exponential $e^{\Omega(t)}$ is unitary, which reduces Eq. (2.22) to

$$(2.23) \quad Y(t) = (e^{\Omega(t)})Y_0(e^{\Omega(t)})^\dagger.$$

3. Numerical Methods Based on the Magnus Expansion. In this section, the process for constructing numerical methods for differential equations based on the Magnus expansion is discussed. In general, numerical methods are constructed by designing appropriate quadrature rules for the Magnus expansion truncated to a given order. The presentation here is focused on collocation type schemes based on Gaussian quadrature rules. As proved in Ref. [13], quadrature rules for the terms in the Magnus expansion based on s Gauss-Legendre quadrature nodes are sufficient for constructing a method of order $2s$. Here, methods of order two, four, and six are considered using both Gauss-Legendre and Gauss-Lobatto quadrature nodes. These methods correspond to quadrature rules applied to one, two, and four terms, respectively, in Eq. (2.16).

Considerable attention in the literature on Magnus integrators is devoted to designing methods requiring the minimum number of function evaluations and matrix commutators for a given order of accuracy [13, 1, 12]. In the context of time parallelization, the manner in which the cost of commutators and function evaluations are counted must reflect the fact that much of the work can be done in parallel, and the minimum parallel cost is not necessarily achieved by a direct parallelization of the serial method with the fewest number of commutators.

Methods for linear equations are discussed first, followed by a discussion of additional considerations for nonlinear problems in section 3.3.

3.1. Quadrature Rules for the Magnus Expansion. In this section, the specific types of quadrature rules used in the numerical methods are described. Quadrature rules based on Gauss-Lobatto or Gauss-Legendre quadrature rules using either two or three quadrature nodes are considered here. Table 1 lists the specific nodes used for each choice as well as the accompanying classical weights. For a method of a given order, each term in the truncated Magnus expansion must be approximated using the function values $A_m = A(t_m)$ (or $A_m = A(y(t_m), t_m)$ for nonlinear problems) at the quadrature nodes t_m corresponding to the quadrature nodes scaled to the time step

Name	Order	Nodes	$q^{(1)}$
Lob-2	2	0, 1	$\frac{1}{2}, \frac{1}{2}$
Lob-3	4	0, $\frac{1}{2}$, 1	$\frac{1}{6}, \frac{4}{6}, \frac{1}{6}$
Leg-3	6	$\frac{1}{2} - \frac{1}{2}\sqrt{\frac{3}{5}}, \frac{1}{2}, \frac{1}{2} + \frac{1}{2}\sqrt{\frac{3}{5}}$	$\frac{5}{18}, \frac{8}{18}, \frac{5}{18}$

TABLE 1

Quadrature nodes and weights for Gauss-Legendre and Gauss-Lobatto rules.

interval $[t_n, t_{n+1}]$. For the schemes described below, the same quadrature nodes are used at each term in the expansion in Eq. (2.16).

First consider the approximation to the first term of the expansion $\Omega^{(1)}(t)$ on the interval $[t_n, t_{n+1}]$ with $\Delta t = t_{n+1} - t_n$. Approximating the integral by Gaussian quadrature gives

$$(3.1) \quad \Omega^{(1)}(t_{n+1}) = \int_{t_n}^{t_{n+1}} A(t) dt \approx \Delta t \sum_{j=1}^M q_j^{(1)} A_j = \Omega_{n+1}^{(1)},$$

where M is the number of quadrature nodes. This is classical quadrature, and the well-known coefficients $q_j^{(1)}$ are given for completeness in Table 1.

In order to obtain a fourth-order method, the second term in the Magnus expansion must be included. Applying a quadrature rule to the double integral yields a general form for $\Omega_{n+1}^{(2)}$

$$(3.2) \quad \Omega_{n+1}^{(2)} = \Delta t^2 \sum_{i=1}^M \sum_{j=1}^M q_{i,j}^{(2)} [A_i, A_j].$$

The simplest approximation to $\Omega_{n+1}^{(2)}$ sufficient for fourth-order accuracy requires the calculation of only a single commutator term

$$(3.3) \quad \Omega_{n+1}^{(2)-1} = \Delta t^2 q^{(2)-1} [A_1, A_3],$$

with $q^{(2)-1} = 1/12$. The method denoted Lob-4-1 (where the 1 denotes one commutator term) uses this approximation. To compute $\Omega^{(2)}$ to the accuracy required for a sixth order method, three nodes can be used and it is necessary to compute three commutators

$$(3.4) \quad \Omega_{n+1}^{(2)-3} = q_1^{(2)-3} [A_1, A_2] + q_2^{(2)-3} [A_1, A_3] + q_3^{(2)-3} [A_2, A_3]$$

with the values

$$(3.5) \quad q_j^{(2)-3} = [-7.1721913818656e-2, -3.5860956909328e-2, -7.1721913818656e-2].$$

Despite the increased computational cost of two additional commutators, in a parallel implementation all three commutators can be computed simultaneously.

To achieve sixth-order accuracy, the first four terms of the Magnus expansion must be included. The sixth-order method denoted Leg-6 approximates the $\Omega^{(3)}$ term using three Gauss-Legendre nodes following the discussion in Ref. [12]. Specifically,

$$(3.6) \quad \begin{aligned} \Omega_{n+1}^{(3)} = \Delta t^3 (& [q_{1,1}^{(3)} A_1 + q_{1,2}^{(3)} A_2 + q_{1,3}^{(3)} A_3, [A_1, A_2]] + \\ & [q_{2,1}^{(3)} A_1 + q_{2,2}^{(3)} A_2 + q_{2,3}^{(3)} A_3, [A_1, A_3]] + \\ & [q_{3,1}^{(3)} A_1 + q_{3,2}^{(3)} A_2 + q_{3,3}^{(3)} A_3, [A_2, A_3]]). \end{aligned}$$

The values of the coefficients $q_{i,j}^{(3)}$ are the same as those in Ref. [12], and in matrix form are

$$(3.7) \quad q^{(3)} = \begin{bmatrix} 3.4538506760729e-3 & -5.5849500293944e-3 & -7.1281599059377e-3 \\ 1.6534391534391e-3 & 0.0 & -1.6534391534391e-3 \\ 7.1281599059377e-3 & 5.5849500293945e-3 & -3.4538506760729e-3 \end{bmatrix}.$$

Exploiting the linear property of commutators ($[A, X] + [A, Y] = [A, X + Y]$) allows one to combine terms that share the same inner single commutator and reduce the number of commutators from nine to three. Note that the single commutator terms, i.e. $[A_1, A_2]$, are computed during the formation of the $\Omega_{n+1}^{(2)}$ term and need not be computed again.

The fourth term in the Magnus expansion can be approximated using a low-order quadrature for a sixth-order method. Following the discussion in Ref. [2], the fourth term is approximated by

$$(3.8) \quad \Omega_{n+1}^{(4)} = \Delta t^4 q^{(4)} [B_0, [B_0, [B_0, B_1]]],$$

where $q^{(4)} = 1/60$ and

$$(3.9) \quad B_i = \Delta t \sum_{j=1}^3 q_j^{(1)} (t_j - 0.5)^i A_j$$

with $q_j^{(1)}$ given in the last row of Table 1.

In section 5, numerical examples are presented for five different Magnus integrators. Table 2 lists the specific discretization of each term included for a given method. The overall order of each method is determined either by the number of terms used in the expansion, or the order of the quadrature rules. For example, the methods Leg-2, Leg-4-3, and Leg-6 use the same quadrature nodes, but differ in the number of terms used in the expansion, while Lob-2 and Leg-2 use different nodes, but are both second-order because only one term in the expansion is used.

Name	Order	Nodes	Ω
Lob-2	2	Lob-2	$\Omega^{(1)}$
Leg-2	2	Leg-3	$\Omega^{(1)}$
Lob-4-1	4	Lob-3	$\Omega^{(1)} + \Omega^{(2)-1}$
Leg-4-3	4	Leg-3	$\Omega^{(1)} + \Omega^{(2)-3}$
Leg-6	6	Leg-3	$\Omega^{(1)} + \Omega^{(2)-3} + \Omega^{(3)} + \Omega^{(4)}$

TABLE 2
Description of the numerical schemes.

3.2. The Matrix Exponential and Solution Update. Once all of the quadrature approximations are applied and the value of Ω_{n+1} is computed, the solution can be updated by

$$(3.10) \quad y_{n+1} = e^{\Omega_{n+1}} y_n.$$

There are many approaches to computing the product of a matrix exponential and a vector of the form $e^A y$ [15], some of which explicitly compute the term e^A and some

which only approximate the product $e^A y$. The choice of method is problem dependent and does not affect the discussion of time parallelism of the methods. In the numerical examples presented here, the scaling-and-squaring method from [10] is used to form the matrix exponential explicitly.

3.3. Considerations for Nonlinear ODEs. Consider now problems of the form

$$(3.11) \quad y' = A(y, t)y$$

where non-linearity is introduced through the y -dependence of A . The terms in the Magnus expansion approximations introduced above now depend on the solution through A_m , and cannot simply be evaluated. The numerical solution at each quadrature node will be denoted y_m and is computed by

$$(3.12) \quad y_m = e^{\Omega_m} y_n,$$

where Ω_m is an approximation to the Magnus expansion on the interval $[t_n, t_m]$. The construction of Ω_m is discussed below.

A simple fixed-point Picard-type iteration is used to simultaneously solve for the values Ω_m and y_m . The iterative scheme is initialized by setting $y_m^{k=1} = y_n$ at each node m , where k denotes the iteration. The solution at each quadrature node is updated by

$$(3.13) \quad y_m^{k+1} = e^{\Omega_m^k} y_n.$$

Then Ω_m^{k+1} is computed using values $A(y_m^{k+1}, t_m)$ as described below.

To compute Ω_m^k , the process for constructing the quadrature rules in section 3.1 needs to be applied to each quadrature node. Evaluating the values $A(y_m^k, t_m)$ at each node t_m is straight-forward but needs to be performed each iteration. In all cases considered here, the same matrix commutators are used for each quadrature rule, so computing the commutators is also identical to the linear case. The significant difference is that a quadrature rule for each term in the Magnus expansion must be computed for each interval $[t_n, t_m]$ rather than just $[t_n, t_{n+1}]$ as in the linear case. The coefficients for each of the terms are included in Appendix A.

Once each Ω_m^k is computed, the matrix exponential can be computed for each node, and a new solution is obtained at each node by Eq. (3.13). The solution is considered converged when the maximum absolute value of the residual

$$(3.14) \quad R_m^k = e^{\Omega_m^k} y_n - y_m^k = y_m^{k+1} - y_m^k$$

is less than a predefined tolerance. In a traditional implementation of this iteration using Gauss-Legendre nodes, it is not necessary to compute the values Ω_{n+1}^k and y_{n+1} at the end of the time step during the iterations; however, when pipelining of iterations is employed as discussed in the next section, y_{n+1} is computed each iteration to update the initial condition for the next time step.

3.4. Considerations for Isospectral Flows. Consider now problems of the form of Eq. (2.21). The procedure for constructing Magnus integrators follows exactly that laid out in the previous section except that the solution is defined at quadrature nodes by

$$(3.15) \quad Y_m^{k+1} = e^{\Omega_m^k} Y_n e^{-\Omega_m^k},$$

and likewise for the computation of Y_{n+1}^{k+1} from Ω_{n+1}^k .

4. Parallelization in Time for Magnus Integrators. In this section, we investigate the theoretical computational cost of the Magnus integrators introduced in the previous section in both serial and parallel settings. We consider both parallelization across the method and parallelization across time. In the following discussion, it is assumed that arbitrarily many processors are available for a given problem and the cost of communication between processors is ignored. It is also assumed that the matrix exponential is formed explicitly as is done for the numerical results given in Section 5.

4.1. The Linear Case. First consider linear problems where $A(t)$ does not depend on the solution y . For each of the N time steps, the following tasks must be performed:

- L1. Evaluate $A_m = A(t_m)$ for each quadrature node t_m
- L2. Compute commutators necessary for each term in the truncated Magnus expansion
- L3. Apply quadrature rules to compute $\Omega(t_{n+1})$
- L4. Form the exponential of $\Omega(t_{n+1})$
- L5. Compute the solution y_{n+1} from y_n by matrix multiplication

Denote by C_A the computational cost of computing $A(t)$ for a given time. Then if M quadrature nodes are used, L1 has a computational cost of MC_A . Next let n^p denote the number of commutators required to compute the p th term in the Magnus expansion and C_C the cost of computing one commutator. Assuming that each commutator in the $p+1$ term can be formed with one additional commutator applied to a term from term p , the total number of commutators to compute is simply $n^1 + \dots + n^P$. Denoting this sum by N_C , the cost of L2 is $N_C C_C$. Task L3 requires only that a linear combination of the terms computed in L2 be computed. We can denote this cost by $N_C C_L$, where C_L is the cost of adding a term in the linear combination. Denote by C_E the cost of the matrix exponential, and hence the cost of L4 is C_E . Likewise denote by C_M the cost of multiplying the solution by a matrix which corresponds to the cost of task L5. Putting these together, the serial cost for N time steps is

$$(4.1) \quad \mathbf{C}_S = N(MC_A + N_C(C_C + C_L) + C_E + C_M).$$

Now consider the parallelization of the method for the linear problem across the step. In task L1, each function evaluation can be done concurrently, so that the cost is reduced from MC_A to C_A . For task L2, all the commutators of a given order can be computed concurrently, so that cost is reduced from $N_C C_C$ to PC_C . Task L3 can be done with cost $\log_2(N_C)C_L$, and the cost of task L4 and L5 remains the same.

Next consider the cost when both parallelization across the method and across time is employed. Given sufficient processors, tasks L1-L4 can all be computed on all time steps concurrently. Only task L5 must be done serially so that the total cost using both forms of parallelism becomes

$$(4.2) \quad \mathbf{C}_P = C_A + PC_C + \log_2(N_C)C_L + C_E + NC_M.$$

Clearly this is a significant reduction in computational cost. If the cost of computing the commutators and matrix exponential (L2 and L4) dominate the other terms, the theoretical parallel speedup approaches N .

An important point to make about this counting is that the cost per step when using parallelization across the step depends very little on the number of quadrature nodes or commutators used in each term since only the cost of task L3 depends on

these factors. Hence higher-order methods are only modestly more expensive per step than lower-order methods and there is less benefit from reducing the number of commutators required in each term of the Magnus expansion since multiple terms can be computed in parallel. Furthermore, for a given accuracy, higher-order methods will require fewer time steps (i.e. smaller N).

4.2. The Nonlinear Case. For nonlinear problems, the theoretical accounting of cost must be modified somewhat. Since we are using an iterative procedure to compute $\Omega(t_m)$, the following steps must be done for each iteration in each time step:

- N1. Evaluate $A(y_m^k, t_m)$ for each quadrature node t_m
- N2. Compute commutators necessary for each term in the truncated Magnus expansion
- N3. Apply quadrature rules to compute Ω_m^k at each quadrature node t_m
- N4. Form exponential of Ω_m^k at each quadrature node t_m .
- N5. Compute y_m^{k+1} at each quadrature node by matrix multiplication by Ω_m^k .

The main difference between these tasks and the linear case is that N2-N5 are done for each quadrature node instead of only once. For simplicity, the serial cost of these steps will be assumed to be M times that of the linear case. Hence, denoting by K_S the number of iterations required for each step in a serial implementation, the serial cost for the nonlinear Magnus method iteration becomes

$$(4.3) \quad \mathbf{C}_S = NK_S M(C_A + N_C(C_C + C_L) + C_E + C_M).$$

As will be shown below, the number of iterations required for convergence K_S depends on Δt in a nontrivial way.

If we allow parallelization across the method, the tasks above can all be computed concurrently at each quadrature node, and hence the cost of each iteration for the nonlinear method is essentially that of one step in the linear case using parallelization across the method.

$$(4.4) \quad \mathbf{C}_I = C_A + PC_C + \log_2(N_C)C_L + C_E + C_M.$$

The speedup across the method is bounded by M .

Now consider parallelization across the steps. The simplest way that this can be accomplished is to pipeline the iterations. We first divide the N time steps into blocks of size N_P with each step in a block assigned to a group of processors indexed by n_p . At each time step n_p for each iteration k , the initial condition is assigned the final value from iteration $k - 1$ of the time step $n_p - 1$. For each block, $N_P - 1$ pipelined iterations are required before the last processor has a consistent initial condition. After this initialization step, assume K_P additional iterations are needed for convergence on every time step in the block. As in the serial case K_P depends in general on the time step Δt and now also on N_P , increasing as N_P increases and decreasing as Δt decreases. Furthermore $K_P \geq K_S$.

The parallel cost on each block will be $(N_P - 1 + K_P)\mathbf{C}_I$ compared to $(N_P K_S)\mathbf{C}_I$ when no parallelization across time steps is applied (i.e. $N_P = 1$). The potential speedup from parallelization across time is then

$$(4.5) \quad \mathbf{S} = \frac{N_P K_S}{N_P - 1 + K_P} = \frac{K_S}{1 + (K_P - 1)/N_P}.$$

Clearly the speedup is bounded by the number of serial iterations required and the best speedup will occur when the quantity K_P/N_P remains small as N_P increases. This ratio will be investigated in the numerical examples.

5. Numerical Examples. In this section, the performance of the different Magnus integrators introduced in section 3 is examined in both serial and parallel settings. After a description of the test problem used, serial results demonstrating the relative accuracy and efficiency of the various methods are presented. Then, in section 5.4, some preliminary results on parallelization of the methods are included.

5.1. Test Case: The Periodic Toda Lattice. The numerical methods will be evaluated on the test problem of a d -particle periodic Toda lattice [19], a one-dimensional chain whose dynamics are governed by nonlinear nearest-neighbor interactions. The equations of motion are a Hamiltonian system for positions q_j and momenta p_j (assuming unit masses)

$$(5.1) \quad q_j' = p_j$$

$$(5.2) \quad p_j' = e^{-(q_j - q_{j-1})} - e^{-(q_{j+1} - q_j)}.$$

In order to cast the dynamics in terms of a Lax pair, one uses the Flaschka change of variables

$$(5.3) \quad \alpha_j = \frac{1}{2} e^{-(q_{j+1} - q_j)/2}$$

$$(5.4) \quad \beta_j = \frac{1}{2} p_j,$$

which leads to definitions of Y and A

$$(5.5) \quad Y = \begin{bmatrix} \beta_1 & \alpha_1 & 0 & \dots & \alpha_d \\ \alpha_1 & \beta_2 & \alpha_2 & \ddots & \vdots \\ 0 & \alpha_2 & \ddots & \ddots & 0 \\ \vdots & \ddots & \ddots & \ddots & \alpha_{d-1} \\ \alpha_d & \dots & 0 & \alpha_{d-1} & \beta_d \end{bmatrix}, A(Y) = \begin{bmatrix} 0 & -\alpha_1 & 0 & \dots & \alpha_d \\ \alpha_1 & 0 & -\alpha_2 & \ddots & \vdots \\ 0 & \alpha_2 & \ddots & \ddots & 0 \\ \vdots & \ddots & \ddots & \ddots & -\alpha_{d-1} \\ -\alpha_d & \dots & 0 & \alpha_{d-1} & 0 \end{bmatrix}.$$

The numerical example considered here is an 11-particle periodic Toda lattice taken from Ref. [20] with the initial conditions

$$(5.6) \quad q(0) = [0, \dots, 0]^T$$

$$p_j(0) = \begin{cases} 4, & 1 \leq j \leq 4, \\ 0, & j \geq 5. \end{cases}$$

The periodic Toda lattice is not asymptotically free and has considerably complicated motion. Figure 1 shows the position and momenta of the 11 particles up to $t_{final} = 10.0$.

5.2. Simulation parameters. The following numerical experiments are all performed on this 11-particle periodic Toda lattice with initial conditions as in Eq. (5.6) and a fixed $t_{final} = 10.0$. A Picard iteration tolerance of 10^{-12} is used on the maximum absolute value of the residual at the end of the timestep. The reference solution is taken as the Leg-6 method with $\Delta t = 2^{-15} t_{final}$ or 32768 steps. The reported error is defined as the matrix 2-norm of the absolute value of the difference between the solution in question and the reference solution.




FIG. 1. *Symmetric periodic 11-particle Toda lattice solutions with initial conditions as in Eq. (5.6).*

The methods have been implemented using the *libpfasst*¹ library. Communication between pipelined iterations is done in *libpfasst* using MPI, and the variable N_P corresponding to the number of parallel steps is also the number of MPI ranks. For parallelization across the method, OpenMP is used to parallelize the steps in the nonlinear iteration. All timing results were performed on a single compute node of NERSC’s Cray XC30 supercomputer, Edison, which contains 24 hardware cores and 64 GB of memory.

5.3. Serial Results. We first perform convergence tests to examine the error with respect to Δt for different Magnus methods. Figure 2 shows the behavior for each method summarized in Table 2. Each method displays the proper convergence rate for a range of Δt . For the second-order methods, note that the error for Leg-2 is significantly smaller than that of Lob-2, which implies the dominant error term for Lob-2 is due to the quadrature rule as opposed to the truncation of the Magnus expansion. The Leg-2 method requires more serial work in this case due to the use of three quadrature nodes instead of two. The difference between the fourth-order methods is less significant. Leg-4-3 is more accurate than Lob-4-1 (with a higher serial cost), but since the main difference between the two methods is how the second term in the Magnus expansion is treated, the difference between the two is smaller than for the second-order methods. Leg-6 is clearly more accurate than the other methods. Note that only about nine significant digits of accuracy is attainable for this problem using double precision due to the sensitivity of the solution to perturbations.

To better demonstrate the relative computational cost of each method, Figure 3 shows the total serial wall-time versus number of steps for the experiment above. As expected, Lob-2 is the method with the shortest time to solution. Lob-4-1, the simplest fourth-order method that can be constructed, is actually less expensive than the second-order Leg-2 scheme despite the fact that Lob-4-1 requires a matrix commutator. For the small problem size used here, the matrix commutators are relatively inexpensive, and the fact that matrix exponentials must be computed at three internal quadrature nodes for Leg-2 more than makes up for the lack of commutator terms. Leg-4-3 and Leg-6 are unsurprisingly the most expensive in serial.

Note that the cost of the methods displayed in Figure 3 does not grow exactly linearly with the number of time steps. This is due to the fact that the number of

¹*libpfasst* is available at <https://bitbucket.org/berkeleylab/libpfasst>.




FIG. 2. Error at $t_{final} = 10.0$ versus Δt for the Toda lattice test case.




FIG. 3. Total wall-time for the solution for the Toda lattice test case for fixed $t_{final} = 10.0$.

Picard iterations needed to converge to the tolerance depends on the time step Δt . Figure 4 shows the average number of iterations over all time steps for each method as a function of Δt . Note the higher-order methods require moderately fewer iterations than lower-order methods, and as the time step gets larger, the number of iterations required for convergence grows rapidly.

The three figures provided above demonstrate that it is not necessarily trivial to choose a method and time step that will provide a solution to a given accuracy with




FIG. 4. Average number of iterations as a function of Δt for serial methods on the Toda lattice test case.

the least computational effort. To illustrate this, Figure 5 shows the accuracy versus wall-clock time for the Toda lattice test. Reading from left to right for a given accuracy shows in increasing order, the methods with the fastest time-to-solution. While Lob-2 is by far the cheapest method, it only the fastest method for simulations where the error is $O(1)$. Lob-4-1 is the most efficient for error tolerances to about 10^{-6} , after which Leg-6 becomes the most efficient. For an error of about 10^{-6} , Leg-6 and Lob-4-1 are more than an order of magnitude more efficient than the second order methods.

5.4. Parallel in Time Results. In this section, the relative performance of parallel Magnus integrators is explored. We first consider the speedup due to parallelization over time steps by pipelining the Picard iterations. As discussed in Section 4, the theoretical speedup from pipelining is bounded by the number of serial iterations required for the method and depends on how the total number of parallel iterations required to reach convergence grows as the number of time parallel steps is increased. As in the serial case, the number of iterations required depends on the time step but now depends also on the number of parallel time steps in the pipeline. Figure 6 demonstrates this dependence for the Toda lattice test using method Leg-6 by plotting the average number of iterations required for convergence for different Δt and parallel steps, denoted by N_P . As in the serial case, the number of iterations required decreases with decreasing Δt and here it also increases for fixed Δt as N_P increases.

A second way to display the convergence behavior of the pipelined iteration is to plot the residual after each iteration for each time rank. Figure 7 shows this data for the Leg-6 method using a time step of $\Delta t = 10/128$ in the top panel and $\Delta t = 10/1024$ in the bottom panel for 16 parallel time steps. As discussed in section 4.2, the speedup achievable from pipelining depends on the ratio K_P/N_P . In this example it is clear that K_P/N_P is decreasing for large N_P as more pipelined time steps are used.

The nontrivial dependence of the parallel iterates on both Δt and N_P makes it difficult to predict which method with which parameters will minimize the time to




FIG. 5. Error versus time to solution for the serial Toda lattice test case. Each point on a line, read from top to bottom, represents a two-fold increase in the number of steps.




FIG. 6. Average number of pipelined iterations to reach residual of 10^{-12} for the parallel Toda lattice test case using the Leg-6 method.

solution for a given accuracy. As in the serial case, it is instructive to consider the accuracy versus wall-clock for different methods with different number of time-parallel steps. Figure 8 displays this information for each of the methods with increasing number of processors. Across a single method, e.g. Leg-6, there's an optimal value of parallelization due to the fact that the increased number of iterations required in the pipeline algorithm starts to cost more than it gains. It is also likely for this test

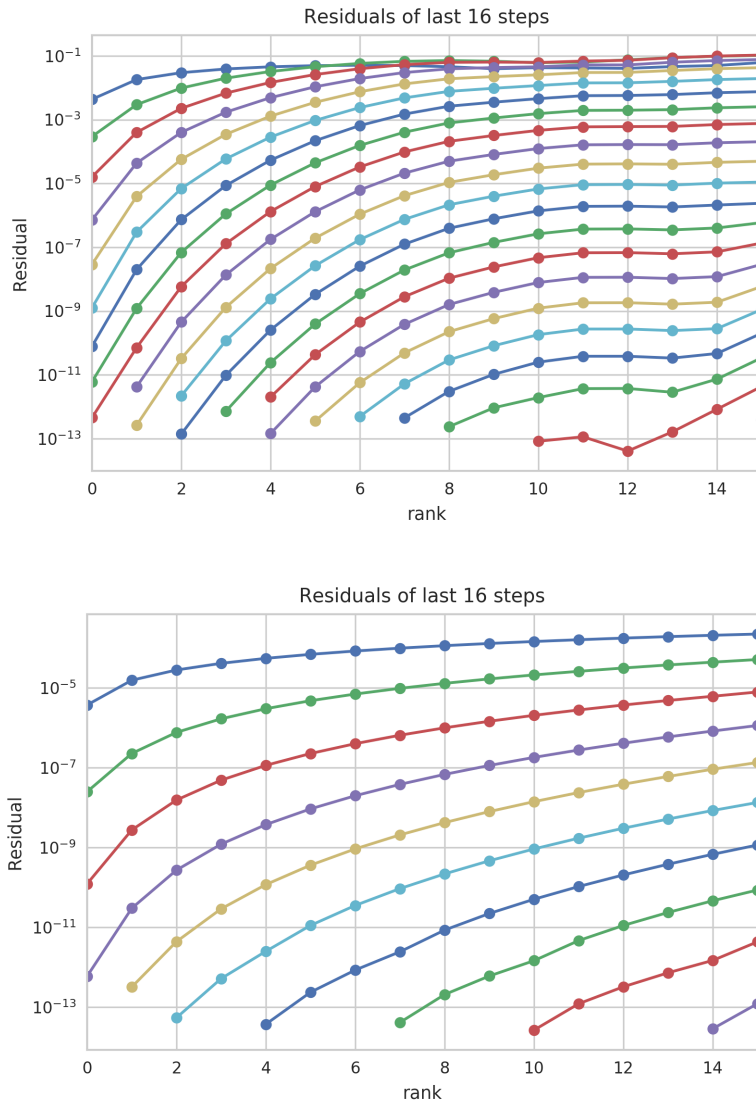


FIG. 7. Residual convergence for each processor in a 16 MPI task simulation using the second-order Legendre method Leg-6 with 128 steps (above) and 1024 steps (below). Every recurrence of a given color is another 6 iterations.

case that the small problem size implies that communication latency is not negligible. Nevertheless, it is again clear that higher-order parallel method gives a shorter time to solution than lower-order alternatives. The second order methods are only less expensive than the higher-order methods when no digits of accuracy are computed.

Finally, we present preliminary results using parallelization across the method and across time steps. Figure 9 shows the error versus time to solution for the serial implementation and time- and method-parallel methods. The particular configuration of $N_p = 8$ and 3 OpenMP threads uses the entirety of 1 hardware node on NERSC's

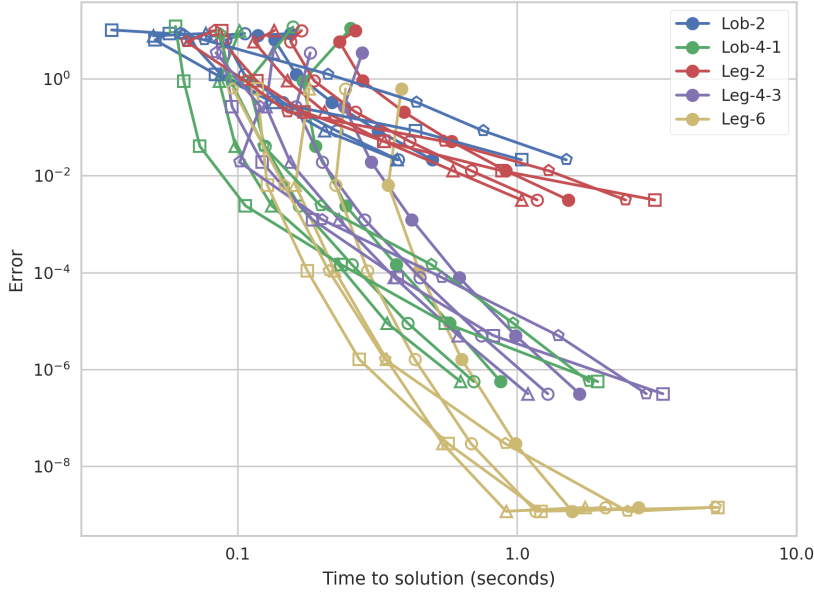


FIG. 8. Error versus time to solution for an 11-particle periodic Toda lattice example. Each color represents a different method and each marker represents a different number of time-parallel steps, $N_p = 1, 2, 4, 8, 16$. Closed circles represent the serial computation, open circles $N_p = 2$, triangles $N_p = 4$, squares $N_p = 8$, and pentagons $N_p = 16$.

Edison Cray XC30 supercomputer with no hyperthreading. The additional parallel across method provides an additional factor of at best two using simple OpenMP parallel do loops over nodes. Using both types of parallelism in this context gives up to a factor of 4.7 in the overall compute time compared to serial methods using 24 total processors. Note that compared to serial second-order methods, the parallel sixth-order method can compute a solution in less time with more than four orders of magnitude lower error.

6. Discussion. Much of the initial work on Magnus integrators is focused on linear problems where only a single evaluation of the Magnus expansion is required for each time step. In contrast, this paper explores the accuracy and cost of Magnus integrators applied to nonlinear problems. Nonlinear Magnus methods require solving an implicit equation involving the Magnus expansion to obtain the solution in each timestep, and the methods proposed here use a simple fixed point iteration for this solution that can be readily parallelized in time.

One conclusion presented here is that in both the linear and nonlinear case, straight-forward parallelization across the method is possible, leading to higher-order methods with only marginally higher computational cost than lower-order methods. This is complimentary to previous efforts in the literature where considerable effort is placed on reducing the complexity of each of the terms in regards to the number of commutators required.

The second level of parallelism described in this work, namely parallelization across the time steps through pipelined iterations, can further decrease the overall




FIG. 9. Error versus time to solution for an 11-particle periodic Toda lattice example. Solid markers indicate serial calculations and open markers indicate time- and method-parallel runs with 8 MPI tasks and 3 OpenMP threads.

time-to-solution for nonlinear problems. There is a non-trivial relationship between the number of pipelined time steps, the time step size, and the number of iterations required for convergence of the iteration, hence it is not easy to predict *a priori* which choice of parameters will lead to the shortest wall clock time given a desired level of accuracy. Nevertheless, for the test case considered here, in most situations, the parallel sixth-order methods require the least computation time and are far superior to serial second-order methods.

It is important to note that the results here are preliminary and were performed on a single test case. In general, the possible parallel speedup attainable for a given problem will depend on the sensitivity of the problem to perturbations, the relative cost of the operations such as computing commutators or the matrix exponential, and the ratio of computation to communication costs. In future work, the authors will use this parallel methodology to investigate real-time electronic dynamics, where the calculation of the right-hand side values is more expensive than both commutators and matrix exponentials.

7. Acknowledgments. The work here was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under contract number DE-AC02005CH11231. Part of the simulations were performed using resources of the National Energy Research Scientific Computing Center (NERSC), a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-

05CH11231.

Appendix A. Quadrature Rules for Intermediate Nodes. The necessary weights for implementing each of the methods Lob-4-1, Leg-2, Leg-4-3, and Leg-6 for nonlinear problems are presented here. For Lob-2, no additional rules are necessary.

For the computation of the single integrals in the first Magnus term, the necessary rules take the form

$$(A.1) \quad \Omega_m^{(1)} = \Delta t \sum_{j=1}^M q_{m,j}^{(1)} A_j,$$

and the coefficients $q_{m,j}^{(1)}$ correspond to the classical collocation schemes (see e.g. [9]) For Lob-3, only one additional rule is needed at the midpoint t_2 ,

$$(A.2) \quad q_{2,j}^{(1)} = \left[\frac{5}{24}, \frac{1}{3}, -\frac{1}{24} \right].$$

For Leg-3, rules for each node are required, and the $q_{m,j}^{(1)}$ are given by

$$(A.3) \quad q_{m,j}^{(1)} = \begin{bmatrix} \frac{5}{36} & \frac{2}{9} - \frac{\sqrt{15}}{15} & \frac{5}{36} - \frac{\sqrt{15}}{30} \\ \frac{5}{36} + \frac{\sqrt{15}}{24} & \frac{2}{9} & \frac{5}{36} - \frac{\sqrt{15}}{24} \\ \frac{5}{36} + \frac{\sqrt{15}}{30} & \frac{2}{9} + \frac{\sqrt{15}}{15} & \frac{5}{36} \end{bmatrix}.$$

For the second term $\Omega^{(2)}$ there are two versions described in the linear case by Eqs. (3.3) and (3.4). In the first case, only one additional coefficient is needed, namely $q_2^{(2)-1} = 1/48$. In the second case, we have

$$(A.4) \quad \Omega_m^{(2)-3} = q_{m,1}^{(2)-3} [A_1, A_2] + q_{m,2}^{(2)-3} [A_1, A_3] + q_{m,3}^{(2)-3} [A_2, A_3],$$

with the values of $q_{m,j}^{(2)-3}$ given by

$$(A.5) \quad q_{m,j}^{(2)-3} = \begin{bmatrix} -7.0825623244174e-4 & -3.5291589565775e-2 & -7.8891497044705e-2 \\ 2.0142743933468e-4 & 4.4826196136660e-3 & -1.8131905893999e-2 \\ -2.6081558162830e-6 & -5.6936734355286e-4 & -3.5152700676886e-2 \end{bmatrix}$$

The third term takes the form

$$(A.6) \quad \begin{aligned} \Omega_m^{(3)} = & [q_{m,1,1}^{(3)} A_1 + q_{m,1,2}^{(3)} A_2 + q_{m,1,3}^{(3)} A_3, [A_1, A_2]] + \\ & [q_{m,2,1}^{(3)} A_1 + q_{m,2,2}^{(3)} A_2 + q_{m,2,3}^{(3)} A_3, [A_1, A_3]] + \\ & [q_{m,3,1}^{(3)} A_1 + q_{m,3,2}^{(3)} A_2 + q_{m,3,3}^{(3)} A_3, [A_2, A_3]]. \end{aligned}$$

with

$$(A.7) \quad q_{1,i,j}^{(3)} = \begin{bmatrix} 1.4667828928181e-6 & -2.5468454487434e-6 & 7.1885579589404e-7 \\ -3.0653702506833e-7 & 6.9623363228690e-7 & -1.9684558120029e-7 \\ -2.2622163607144e-8 & -2.7279719400850e-9 & 8.5484354192049e-10 \end{bmatrix}$$

$$(A.8) \quad q_{2,i,j}^{(3)} = \begin{bmatrix} 1.0401143365317e-3 & -1.7143302808715e-3 & 1.9808827525182e-4 \\ -6.9105495969459e-5 & 2.9054016014502e-4 & -3.4658846939476e-5 \\ 9.2451884893203e-5 & 1.2595057164957e-5 & -2.4709074423914e-6 \end{bmatrix}$$

$$(A.9) \quad q_{3,i,j}^{(3)} = \begin{bmatrix} 4.1482959753609e-3 & -6.3874218931689e-3 & -3.5942319108173e-3 \\ 9.9737811032708e-4 & 1.2415302375576e-4 & -3.8059754231607e-4 \\ 3.7183849345731e-3 & 1.6935142950568e-3 & -1.0604085845381e-3 \end{bmatrix}.$$

Finally, for the fourth term, $\Omega_m^{(4)}$ is computed as in Eqs. (3.8) and (3.9), with $q^{(4)} = 1/60$ and $q_j^{(1)}$ in (3.9) replaced with $q_{m,j}^{(1)}$ from Eq. (A.3).

REFERENCES

- [1] S. BLANES, F. CASAS, J. A. OTEO, AND J. ROS, *The Magnus expansion and some of its applications*, Physics Reports, 470 (2009), pp. 151–238.
- [2] S. BLANES, F. CASAS, AND J. ROS, *Improved High Order Integrators Based On The Magnus Expansion*, Bit, 40 (2000), pp. 434–450.
- [3] S. BLANES AND P. C. MOAN, *Fourth- and sixth-order commutator-free Magnus integrators for linear and non-linear dynamical systems*, Applied Numerical Mathematics, 56 (2006), pp. 1519–1537.
- [4] K. BURRAGE, *Parallel methods for ODEs*, Advances in Computational Mathematics, 7 (1997), pp. 1–3.
- [5] L. EULER, *Institutiones calculi integralis*, no. v. 2, Acad. Imper. scientiarum, 1768.
- [6] M. J. GANDER, *50 years of time parallel time integration*, in Multiple Shooting and Time Domain Decomposition Methods, T. Carraro, M. Geiger, S. Körkel, and R. Rannacher, eds., Cham, 2015, Springer International Publishing, pp. 69–113.
- [7] E. HAIRER, C. LUBICH, AND G. WANNER, *Geometric numerical integration : structure-preserving algorithms for ordinary differential equations*, Springer, 2006.
- [8] E. HAIRER, S. P. NØRSETT, AND G. WANNER, *Solving ordinary differential equations I. nonstiff problems*, Mathematics and Computers in Simulation, 29 (1987), p. 447.
- [9] E. HAIRER AND G. WANNER, *Solving Ordinary Differential Equations II : Stiff and Differential-Algebraic Problems*, Springer Berlin Heidelberg, 1991.
- [10] N. J. HIGHAM, *The Scaling and Squaring Method for the Matrix Exponential Revisited*, SIAM Journal on Matrix Analysis and Applications, 26 (2005), pp. 1179–1193.
- [11] M. HOCHBRUCK AND A. OSTERMANN, *Exponential integrators*, Acta Numerica, 19 (2010), pp. 209–286.
- [12] A. ISERLES, H. Z. MUNTHE-KAAS, S. P. NØRSETT, AND A. ZANNA, *Lie-group methods*, Acta Numerica, (2000), pp. 215–365.
- [13] A. ISERLES AND S. P. NØRSETT, *On the solution of linear differential equations in Lie groups*, Philosophical Transactions: Mathematical, Physical and Engineering Sciences, 357 (1999), pp. 983–1019.
- [14] W. MAGNUS, *On the exponential solution of differential equations for a linear operator*, Communications on pure and applied . . . , VII (1954), pp. 649–673.
- [15] C. MOLER AND C. VAN LOAN, *Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later*, SIAM Review, 45 (2003), pp. 3–49.
- [16] H. MUNTHE-KAAS AND B. OWREN, *Computations in a free Lie algebra*, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 357 (1999), pp. 957–981.
- [17] J. NIEVERGELT, *Parallel methods for integrating ordinary differential equations*, Communications of the ACM, (1964), pp. 731–733.
- [18] M. THALHAMMER, *A fourth-order commutator-free exponential integrator for nonautonomous differential equations*, SIAM J. Numer. Anal., 44 (2006), pp. 851–864 (elect).
- [19] M. TODA, *Vibration of a chain with nonlinear interaction*, Journal of the Physical Society of Japan, 22 (1967), pp. 431–436.
- [20] A. ZANNA, *Numerical solution of isospectral flows*, PhD thesis, University of Cambridge, 1998.