Advances in Applied Mathematics and Mechanics
Adv. Appl. Math. Mech., Vol. **18**, No. 3, pp. 815-832

DOI: 10.4208/aamm.OA-2024-0226
June 2026

Should Exponential Integrators Be Used for Advection-Dominated Problems?

Lukas Einkemmer, Trung-Hau Hoang and Alexander Ostermann*

Department of Mathematics, University of Innsbruck, Austria

Received 16 October 2024; Accepted (in revised version) 28 May 2025

Abstract. In this paper, we consider the application of exponential integrators to problems that are advection dominated. In this context, we compare Leja and Krylov based methods to compute the action of exponential and related matrix functions. We set up a performance model by counting the different operations needed to implement the considered algorithms. This model assumes that the evaluation of the right-hand side is memory bound and allows us to evaluate performance in a hardware independent way. We find that exponential integrators, depending on the specific setting, either outperform or perform similarly to explicit Runge–Kutta schemes. We generally observe that Leja based methods outperform Krylov iterations in the problems considered. This is in particular true if computing inner products is expensive.

AMS subject classifications: 65L04, 65M22, 65F60

Key words: Exponential integrators, efficiency, advection-dominated problems.

1 Introduction

In the last few decades, exponential integrators have gained popularity as an option for solving complex systems of differential equations that exhibit stiffness (see, e.g., the review [21]). The main idea behind exponential integrators is as follows: the right-hand side of the problem is linearized, the linear component is integrated exactly and the remaining nonlinearity is discretized in an explicit way. In order to implement such schemes, approximations of the actions of the matrix exponential and related matrix functions are required. Doing this efficiently is important to obtain a competitive method.

In situations where the diffusion/advection coefficient is constant, the domain is a simple rectangle, and either periodic or homogeneous Dirichlet/Neumann boundary conditions are applied, the fast Fourier transform or related techniques can be employed

*Corresponding author.
Emails: lukas.einkemmer@uibk.ac.at (L. Einkemmer), trunghaugg@gmail.com (T. Hoang), alexander.ostermann@uibk.ac.at (A. Ostermann)

(see, e.g., [2,8,9,24]). If these assumptions do not hold, alternative methods must be used. Conventional approaches such as Padé approximations or diagonalization, however, are only practical when dealing with small systems. This paper focuses on two widely used classes of methods for approximating the action of a matrix functions on vectors in the context of large-scale systems. Specifically, we examine Krylov subspace methods (see, e.g., [19,29]) and Leja interpolation (see, e.g., [3,6]).

Exponential methods have been investigated extensively for diffusion-dominated problems (see, e.g., [20,21,27]). In this situation, it can be shown that they exhibit significantly improved performance compared to explicit methods and similar or in some cases even improved performance compared to implicit methods. Most commonly this comparison is performed under the assumption that no or no good preconditioner is available, although some ways to incorporate the solution of a simplified but similar problem into exponential integrators have been explored [7,11,15]. Significantly less research has been done in the context of advection-dominated problems (see [5,10,13,22]), the situation we consider in this paper. Here, we consider linear advection-diffusion problems with non-constant diffusion coefficients and a 2D compressible isothermal Navier–Stokes problem. Throughout this work, we aim to study exponential integrators as general-purpose methods, rather than focusing on solving specific problems where more efficient methods could be chosen.

To perform the comparisons, we evaluate the costs associated with the different methods. Our goal it to compare performance in a hardware independent way. To accomplish this, we develop a performance model (as is commonly done in computer science [25,30,31]) that counts operations such as matrix-vector multiplications, inner products, scalar multiplications, and performing linear combinations. This also allows us, for example, to investigate what effect expensive inner products (as is common in the case of modern distributed memory supercomputers) have on the overall cost.

The paper is structured as follows. We first recall exponential Rosenbrock methods in Section 2. In Section 3, we introduce the test problems for the numerical investigations. The performance model used to evaluate computational cost is explained in Section 4. In Sections 5 and 6, we present the numerical experiments for these problems. Finally, Section 7 provides the conclusions of the study.

2 Exponential Rosenbrock methods

In this paper, we compare the performance of explicit Runge–Kutta methods with exponential Rosenbrock schemes. While the former are well-established (see, e.g., the textbooks [17,18]), we recall two representations of the latter that will be used in this investigation. To present the schemes, we consider the autonomous problem

$$u'(t)=F(u(t)), \quad u(t_0)=u_0, \tag{2.1}$$

where F is a nonlinear function of u. Let u_k be the numerical solution, which approximates the exact solution $u(t_k)$ at time t_k. By linearizing Eq. (2.1) in each step at u_k, we