

Computational Science Laboratory
Technical Report CSL-TR-16 -6
July 26, 2016

Arash Sarshar, Paul Tranquilli, Brent
Pickering, Andrew McCall, Adrian Sandu and
Christopher J. Roy

“A Numerical Investigation of
Matrix-Free Implicit Time-Stepping
Methods for Large CFD Simulations”

Computational Science Laboratory
Computer Science Department
Virginia Polytechnic Institute and State University
Blacksburg, VA 24060
Phone: (540)-231-2193
Fax: (540)-231-6075
Email: sarshar@vt.edu
Web: <http://csl.cs.vt.edu>



Compute the Future



A Numerical Investigation of Matrix-Free Implicit Time-Stepping Methods for Large CFD Simulations

Arash Sarshar^{a,*}, Paul Tranquilli^a, Brent Pickering^b, Andrew McCall^b, Adrian Sandu^a,
Christopher J. Roy^b

^a*Department of Computer Science and Applications, VirginiaTech*

^b*Department of Aerospace and Ocean Engineering, VirginiaTech*

Abstract

This paper is concerned with the development and testing of advanced time-stepping methods suited for the integration of time-accurate, real-world applications of computational fluid dynamics (CFD). The performance of several time discretization methods is studied numerically with regards to computational efficiency, order of accuracy, and stability, as well as the ability to treat effectively stiff problems. We consider matrix-free implementations, a popular approach for time-stepping methods applied to large CFD applications due to its adherence to scalable matrix-vector operations and a small memory footprint. We compare explicit methods with matrix-free implementations of implicit, linearly-implicit, as well as Rosenbrock-Krylov methods. We show that Rosenbrock-Krylov methods are competitive with existing techniques excelling for a number of problem types and settings.

Keywords: Initial value problems, Computational Fluid Dynamics AMS 65L05, 65L07

1. Introduction

While flow problems are inherently unsteady, computer flow simulations have traditionally focused mainly on steady-state flow problems because they reduce the computational effort dramatically. Nevertheless, in many practical applications it is important to quantify the impact of unsteady flow phenomena on the forces and moments exerted on a body. These phenomena impact performance characteristics such as the lift and drag of a body, or the dynamic response of a control system. Historically, in aircraft design, these unsteady effects have required additional analyses to mitigate undesirable aeroelastic effects such as wing flutter and undesirable stall characteristics, among other issues [1–3]. Sometimes the unsteady effects are beneficial, e.g., when using leading edge extensions (LEX) to improve high angle of attack performance [4]. Studies of low Reynolds number unsteady flows have become much more relevant today with the development of micro

*Corresponding author

Email addresses: sarshar@vt.edu (Arash Sarshar), ptranq@vt.edu (Paul Tranquilli), bpickeri@vt.edu (Brent Pickering), mandrew9@vt.edu (Andrew McCall), sandu@cs.vt.edu (Adrian Sandu), cjroy@vt.edu (Christopher J. Roy)

Preprint submitted to Journal of

July 26, 2016

air vehicles (MAV) [5–11], with typical sizes as small as 15 cm. At these low Reynolds numbers viscous forces dominate the flow characteristics, leading to unsteady viscous effects such as laminar separation and von Kármán vortices, as commonly demonstrated in the flow over a cylinder [12].

The use of CFD allows for preliminary analyses of these designs to determine whether any undesirable unsteady effects will be present, before committing to the expensive development and testing of a physical system. However, a major limitation of unsteady flow analysis using CFD is the prohibitive amount of computational time required to simulate a time-accurate solution with the time integration schemes commonly used to solve the Navier-Stokes equations. This coupled set of nonlinear partial differential equations has to be solved iteratively to determine the solution for each time step. A fine mesh resolution typically required to capture the length and time scales of the flow. For making the computations feasible it is necessary to use a time discretization scheme that maximizes convergence speed without prohibitively restricting the time steps due to stability constraints. It is also important that the time integration scheme provides an accurate solution at every time step.

Explicit time integration methods have been used for time-accurate solutions of unsteady flow problems due to their low computational cost per step and moderate memory requirements. For example, in [13] a number of embedded high-order explicit Runge-Kutta methods with minimal memory storage have been developed for the compressible Navier-Stokes equations based on van der Houwen’s technique [14] for stage memory storage reduction. However, stability constraints restrict the maximum time steps that explicit methods can employ.

Implicit time-stepping methods have better stability properties than explicit methods, and therefore they can use very large time steps. However, the computational costs per step are also larger. The overall computational efficiency is given by the tradeoff between the computational cost per step and the total number of steps required to carry out the simulation.

The main cost of implicit methods is associated with solving a large system of nonlinear equations at each step [15]. Newton type methods for the solution of nonlinear systems are commonly used in the CFD literature in conjunction with preconditioned Krylov-based solvers for the inherent linear systems [16]. The popular Jacobian-free Newton Krylov (JFNK) methods employ finite difference approximations of the Jacobian-vector products required by Krylov solvers [17]. Studies of JFNK methods applied to solve Navier-Stokes equations [18] have shown that error tolerances of Krylov space solvers need to be carefully optimized for performance and accuracy.

There is considerable interest in developing numerical schemes that provide a suitable level of implicitness for time integration of stiff flow problems, such as to allow relatively large time steps while keeping the cost per time step comparable to that of explicit methods. In this paper we study the efficiency of several different matrix-free, both explicit and implicit, time integration methods applied to computational fluid dynamics problems of moderate to large dimensions. In addition to standard techniques, we also examine a new class of lightly-implicit time integration schemes, called Rosenbrock-Krylov methods which are particularly well suited to employ approximate Jacobian-vector products.

The remaining part of this paper is structured as follows. Section 2 reviews numerical methods accessible for the time integration of large systems of ordinary differential equa-

tions arising from flow problems. The numerical methods investigated here and their implementation are presented in Section 3. Section 4 applies these methods to a number of test problems and studies their effectiveness in terms of their numerical accuracy, stability, and computational efficiency in case of high dimensional problems. Conclusions and future work directions are discussed in Section 5.

2. Numerical time integration for CFD applications

Consider the autonomous initial value problem:

$$\frac{dy}{dt} = f(y), \quad y(t_0) = y_0, \quad t_0 \leq t \leq t_F, \quad y(t) \in \mathbb{R}^N, \quad f: \mathbb{R}^N \rightarrow \mathbb{R}^N. \quad (1)$$

In this paper equation (1) represents the system of ODEs resulting from the spatial semi-discretization of the Navier-Stokes equations for flow problems in the method-of-lines framework. The system is considered autonomous without loss of generality: any system can be written in autonomous form by appending the time variable to the solution vector. With only time derivatives remaining in equation (1), it is the choice of time-stepping method that determines the stability, accuracy, and efficiency of the numerical solution as the solution is propagated in time. This paper is concerned with the study of high-order implicit time marching schemes and their performance in large CFD applications.

We next review several important classes of numerical time integration algorithms.

2.1. Runge-Kutta methods

The historically well-known time integration schemes attributed to Runge and Kutta are well-studied [19, 20] and extensively utilized in flow applications [21, 22]. Let $y_n \approx y(t_n)$ be a numerical approximation of the solution of the system (1). An s -stage Runge-Kutta method advances the numerical solution to the next time step $t_{n+1} = t_n + h$ as follows:

$$k_i = f \left(y_n + h \sum_{j=1}^s a_{i,j} k_j \right), \quad i = 1, \dots, s; \quad (2a)$$

$$y_{n+1} = y_n + h \sum_{j=1}^s b_j k_j. \quad (2b)$$

The method coefficients

$$a = [a_{i,j}]_{1 \leq i, j \leq s} \quad b = [b_i]_{1 \leq i \leq s} \quad c = [c_i]_{1 \leq i \leq s},$$

are determined such that the method (2) has the desired accuracy and stability properties [23, II.1].

Explicit Runge-Kutta (ERK) methods are characterized by coefficients $a_{i,j} = 0$ for any $j \leq i$. This means that each stage value k_i (2a) depends only on previously stage vectors k_1, \dots, k_{i-1} . This leads to the convenient result that explicit Runge-Kutta methods need only one ODE right-hand-side function evaluation per stage, and no linear or nonlinear systems of equations are solved in the process. The stability requirements due to CFL conditions limit the step size h , and therefore impact the efficiency of the method.

Singly Diagonally Implicit Runge-Kutta methods (SDIRK) [24, IV.6] are characterized by coefficients $a_{i,j} = 0$ for any $j < i$, and $a_{i,i} = \gamma > 0$ for all stages $i = 1, \dots, s$. Solving for the stage vector k_i requires the solution of a nonlinear system of equations at each stage

$$F_i(k_i) = k_i - f(\xi_i + h\gamma k_i) = 0 \quad \text{for } i = 1, \dots, s, \quad (3)$$

which makes the computational cost per step significantly larger than for ERK. However, this also leads to improved stability properties and the ability to use much larger time steps. The nonlinear equation (3) is solved using Newton-type iterations:

$$\Delta k_i^{\{\ell\}} = -\left(\frac{\partial F_i}{\partial k_i}\right)^{-1} F_i(k_i^{\{\ell\}}), \quad k_i^{\{\ell+1\}} = k_i^{\{\ell\}} + \Delta k_i^{\{\ell\}}, \quad \ell = 0, 1, \dots \quad (4)$$

where

$$\frac{\partial F_i}{\partial k_i} = \mathbf{I}_N - h\gamma \mathbf{J}_n, \quad (5)$$

and \mathbf{J}_n is the Jacobian of the ODE right-hand-side function:

$$\mathbf{J}_n = \left. \frac{\partial f(y)}{\partial y} \right|_{y=y_n}. \quad (6)$$

The fact that $a_{i,i} = \gamma$ for all stages allows re-using the LU decomposition of (5) in the solution of linear systems appearing in equation (4) for all stage vectors $i = 1, \dots, s$.

2.2. Rosenbrock methods

Linearly implicit methods avoid the nonlinear systems (4) and solve only linear systems at each stage. One step of a Rosenbrock (ROS) method [24, IV.7] reads:

$$Y_i = y_n + h \sum_{j=1}^{i-1} a_{i,j} k_j, \quad k_i = f(Y_i) + h \mathbf{J}_n \sum_{j=1}^i \gamma_{i,j} k_j, \quad i = 1, \dots, s; \quad (7a)$$

$$y_{n+1} = y_n + h \sum_{i=1}^s b_i k_i. \quad (7b)$$

Therefore, the stage vectors k_i are found in succession by solving linear systems of the form:

$$(\mathbf{I}_N - h\gamma_{i,i} \mathbf{J}_n) k_i = f(Y_i) + h \mathbf{J}_n \sum_{j=1}^{i-1} \gamma_{i,j} k_j, \quad i = 1, \dots, s. \quad (8)$$

As with SDIRK methods, choosing $\gamma_{i,i} = \gamma > 0$ for $i = 1, \dots, s$ helps reduce the computational costs by allowing to reuse the same LU factorization (5) for all stages.

2.2.1. Rosenbrock-Wanner methods

Rosenbrock methods simplify the computational effort necessary to solve the stage vector equations by limiting the implicitness to linear terms containing the exact Jacobian right-hand-side products [25]. As a consequence, the accuracy of the method depends on the availability of the exact Jacobian. In many practical cases an exact

Jacobian is difficult to compute, however some approximation of the Jacobian may be available at reasonable computational cost. The Rosenbrock-Wanner (ROW) methods are Rosenbrock schemes that retain the order of accuracy for any matrix \mathbf{A}_n used in place of the exact Jacobian \mathbf{J}_n in (7). The preservation of accuracy is possible by imposing additional order conditions on the method coefficients [26]. Better approximations of the Jacobian $\mathbf{A}_n \approx \mathbf{J}_n$ will ensure better numerical stability. We note that while the formal definition for a ROW method is the same as in equation (7), the method coefficients are different due to the additional order conditions.

2.2.2. Rosenbrock-Krylov methods

The stage vectors in Rosenbrock-type methods are computed by solving the linear system of dimension N in equation (8). For large problems the solutions of these linear systems is best obtained via a Krylov-space iterative linear algebra solver such as GMRES [27].

Instead of using a Krylov-based iterative solver such as GMRES, Rosenbrock-Krylov (ROK) methods developed in [28] reformulate the method (7) using implicitness only in the Krylov subspace of dimension M constructed using modified Arnoldi iteration [27]

$$\mathcal{K}_M(\mathbf{J}_n, f(y_n)) = \text{range}\{\mathbf{V}_n\}, \quad \mathbf{V}_n \in \mathbb{R}^{N \times M}, \quad \mathbf{V}_n^T \mathbf{V}_n = \mathbf{I}_M, \quad \mathbf{V}_n^T \mathbf{J}_n \mathbf{V}_n = \mathbf{H}_n.$$

Here \mathbf{H} and \mathbf{V} are the upper Hessenberg and the orthogonal basis of the Krylov space, respectively, and are results of Arnoldi process.

A single time step of a ROK method is constructed as follows [28]:

$$F_i = f \left(y_n + \sum_{j=1}^{i-1} \alpha_{i,j} k_j \right), \quad (9a)$$

$$\phi_i = \mathbf{V}_n^T F_i, \quad (9b)$$

$$\lambda_i = (\mathbf{I}_M - h \gamma \mathbf{H}_n)^{-1} \left(h \phi_i + h \mathbf{H}_n \sum_{j=1}^{i-1} \gamma_{i,j} \lambda_j \right), \quad (9c)$$

$$k_i = \mathbf{V}_n \lambda_i + h (F_i - \mathbf{V}_n \phi_i), \quad (9d)$$

$$y_{n+1} = y_n + \sum_{i=1}^s b_i k_i. \quad (9e)$$

For $M \ll N$ the linear system (9c) is easily solvable using direct methods, and the stage vectors in full space can be recovered projecting the reduced space stage values back to full space [28]. The minimum dimension of the Krylov space is determined by the desired order of the numerical scheme. Readers interested in the derivation of the order conditions for this method may refer to [28]. In the extreme case $M = 0$ the method (9) reduces to an explicit Runge-Kutta method. In practice, however, we need the Krylov space to be large enough to capture some of the dominant eigenvalues of the full Jacobian, corresponding to fast-changing modes of (1), such as to alleviate the restrictions on step size of explicit methods imposed by the stiffness of the problem. An important question is whether the additional computational cost required by Arnoldi iteration can be compensated by the increases in step size due to the implicit nature of the method. The numerical results in Section 3 provide answers to this question.

2.3. Matrix-free implementations

As discussed in Sections 2.1 and 2.2, implicit time-stepping methods use the Jacobian matrix (6) to solve linear or nonlinear systems of equations at each step. The dimension of the state vector in (1) may become significantly large when a highly refined spatial discretization is required, whether to capture fine details of flow or when CFD is used in large data-driven applications such as climate research. Computation and storage of Jacobian matrices, even in sparse form, is not favorable in such scenarios.

In some cases the complexity of the spatial discretization scheme impedes construction of analytic Jacobian matrices. In CFD codes such as SENSEI-Lite [29], the use of complex upwind flux schemes including Roe’s and Van Leer’s flux scheme [30, V] in addition to MUSCL [31] reconstruction with flux limiters make the generation of an analytic Jacobian challenging. Furthermore, to allow for modularity as a research code, SENSEI-Lite allows these different flux schemes and limiters to be used interchangeably, dependent upon the test problem configuration. This interchangeability would necessitate the formulation of an analytic Jacobian for every possible configuration, eliminating the modularity of the code structure. The framework of matrix-free methods allows us to exploit the benefits of advanced time-stepping methods without forming the Jacobian matrix directly.

Krylov space iterative methods for solving the stage equations (4) or (8) rely on computing Jacobian-vector products. Instead of computing the Jacobian matrix and then multiplying, it is possible to approximate directly Jacobian-vector products using the finite-difference approximation of a directional derivative:

$$\mathbf{J}_n \cdot v = \left. \frac{\partial f(y)}{\partial y} \right|_{y_n} \cdot v \approx \frac{f(y_n + \varepsilon v) - f(y_n)}{\varepsilon}. \quad (10)$$

The optimum ε is chosen considering the trade-off between truncation and round-off errors [32]. Higher-order approximations are not favorable here as they require more right-hand-side function evaluations, which, considering the large dimensions of the problem, are costly to compute.

A more in-depth analysis of different strategies to compute Jacobian-vector products, and their effects on convergence and efficiency of implicit time integration methods, can be found in [33]. Of special interest is using exact Jacobian-vector products instead of finite difference approximations. Numerical experiments applied to discretizations of PDEs in [33] indicate that exact Jacobian-vector products provide more robustness in observed convergence orders, and increased runtime efficiency over methods using approximate products (10).

3. The software infrastructure for numerical investigations

3.1. Time-stepping schemes and their implementation

The time integration software used in the numerical experiments is MATLODE [34], a Matlab library for integration of ODE systems including implicit and explicit Runge-Kutta methods, Rosenbrock methods and Krylov based methods. The package also supports forward, adjoint, and tangent linear models, enabling sensitivity analysis applications. Aside from the methods available in MATLODE package, Matlab’s explicit time-stepping scheme based on Dormand and Prince [35] is also included in tests for comparison. Table 1 summarizes the methods used in numerical experiments and their properties.

Table 1: Overview of time stepping methods used in numerical experiments.

Method	Family	Stages	Order	Stability
ERK	Explicit Runge-Kutta	5	4	Conditionally stable
DOPRI5	Explicit Runge-Kutta	7	5	Conditionally stable
DOPRI853	Explicit Runge-Kutta	12	8	Conditionally stable
SDIRK	Implicit Runge-Kutta	5	4	L-stable
ROS4	Rosenbrock	4	4	L-stable
ROW	Rosenbrock-W	4	3	L-stable
ROK	Rosenbrock-Krylov	5	4	Conditionally stable
ODE45	Explicit Runge-Kutta	5	4	Conditionally stable

3.2. The fluid flow simulation code

SENSE-Lite, the CFD code employed in this paper, can solve both the Euler and Navier-Stokes equations for compressible flow [29]:

$$\frac{\partial \rho}{\partial t} + \frac{\partial \rho v_1}{\partial x_1} + \frac{\partial \rho v_2}{\partial x_2} = 0, \quad (11a)$$

$$\frac{\partial \rho v_1}{\partial t} + \frac{\partial \rho v_1^2}{\partial x_1} + \frac{\partial \rho v_1 v_2}{\partial x_2} + \frac{\partial p}{\partial x_1} = \frac{\partial \tau_{11}}{\partial x_1} + \frac{\partial \tau_{12}}{\partial x_2}, \quad (11b)$$

$$\frac{\partial \rho v_2}{\partial t} + \frac{\partial \rho v_1 v_2}{\partial x_1} + \frac{\partial \rho v_2^2}{\partial x_2} + \frac{\partial p}{\partial x_2} = \frac{\partial \tau_{12}}{\partial x_1} + \frac{\partial \tau_{22}}{\partial x_2}, \quad (11c)$$

$$\begin{aligned} & \frac{\partial E_t}{\partial t} + \frac{\partial \rho v_1 E_t}{\partial x_1} + \frac{\partial \rho v_2 E_t}{\partial x_2} + \frac{\partial v_1 p}{\partial x_1} + \frac{\partial v_2 p}{\partial x_2} = \\ & k \left(\frac{\partial q_1}{\partial x_1} + \frac{\partial q_2}{\partial x_2} \right) + \mu \left(\frac{\partial}{\partial x_1} (v_1 \tau_{11} + v_2 \tau_{12}) + \frac{\partial}{\partial x_2} (v_1 \tau_{12} + v_2 \tau_{22}) \right), \end{aligned} \quad (11d)$$

where

$$q_i = -k \frac{\partial T}{\partial x_i}, \quad (11e)$$

$$\tau_{i,j} = \mu \left(\frac{\partial v_j}{\partial x_i} + \frac{\partial v_i}{\partial x_j} \right) - \frac{2}{3} \mu (\nabla \cdot v) \delta_{i,j}. \quad (11f)$$

SENSEI-Lite uses a curvilinear, structured-grid, finite volume method. Second-order spatial accuracy is achieved through a standard total variation diminishing scheme consisting of MUSCL reconstruction and selectable flux limiters [31]. The code is written in C++ and MEX interfaces are used to call it from within time integrators implemented in Matlab. The primary function of the MEX code is to return the spatial residual for a given solution state; this is returned as a vector that is independent of any temporal information and can be used for building arbitrary time integration methods such as multi-stage ERK. This residual vector can optionally be returned as multiple vectors split according to the underlying equations, e.g., viscous and inviscid contributions from the Navier-Stokes equations as shown in (11); these vectors sum to the full residual and can also be used to integrate the equations independently. The Matlab client code is responsible for storing multiple solution state and update vectors as necessary, and can apply time-dependent source terms if desired.

4. Numerical results

4.1. Experimental setting

This section details the numerical experiments setup using MATLODE and SENSEI-Lite packages to study the performance of matrix-free time-stepping methods on unsteady flow problems. A reference solution is computed and stored by integrating the model using an explicit method with tight tolerances ($\sim 10^{-9}$) on errors. In each numerical experiment, the solution at the final integration time is compared against this reference solution, and the error is measured using the \mathcal{L}_2 norm:

$$Error = \|y_N - y_{ref}\|_2 \quad \text{where} \quad y = [\rho, \rho v, \rho u, E_t]^T$$

Where y_N and y_{ref} are numerical and reference state vectors at the final integration time respectively. All experiments use matrix-free time-stepping methods. For each test problem we evaluate the eigenvalues of the Jacobian matrix in equation (6) to get information about the dynamic modes of the state variable evolution. A wide spread of the eigenvalues indicates the existence of both slow and fast dynamics, in other words, of “stiff” dynamics. The largest eigenvalue gives an estimate of the largest stable step-size in explicit methods. Eigenvalues are computed using Matlab’s implementation of the implicitly restarted Arnoldi method that estimates the largest 1000 eigenvalues in magnitude for each test problem; this computation is also performed in matrix-free form.

4.2. Vortex-shedding cylinder test problem

The vortex-shedding cylinder test problem consists of a two-dimensional circular cylinder in a low subsonic flow (Mach 0.1) using a gas model for air at 5,000 ft altitude standard atmospheric conditions (278K, 84.31kPa). The viscosity for air is calculated based on local flow conditions using Sutherland’s law. The modeled is a free-stream flow, using far-field boundary conditions set at over 100 chord lengths away from the surface of the cylinder to minimize interactions with the boundary. The default cylinder diameter is $8 \times 10^{-5} m$, which yields a Reynolds number of approximately 200 and results in a steady and predictable two-dimensional shedding of alternating vortices behind the cylinder. At this low Reynolds number there are no sub-grid-scale turbulence effects, so all physically accurate spatial and temporal scales in the solution can be directly modeled; therefore, the CFD code can solve the laminar Navier-Stokes equations with no underlying turbulence model. For all test problems, the flux scheme used is Roe’s approximate Riemann solver and no flux limiter is employed in the MUSCL scheme as flux limiters were observed to be unnecessary and to generally reduce the accuracy of the FVM reconstruction in the continuous and smooth flow field around the cylinder. The parameters of this problem are modified to provide different tests, as summarized in Table 2. Care is taken so that the qualitative solution behavior (and appropriate Reynolds number) is maintained for all tests.

The first experiment is performed on the vortex-shedding cylinder test problem 1 with parameters given in Table 2. Figure 1 illustrates snapshots of the density component of the flow for cylinder test problem 1 at different times, showing the cyclic development of vortices behind the cylindrical object. This experiment uses fixed step sizes to study the temporal orders of convergence for each method, and the results are plotted in Figure 2. Numerical orders of convergence calculated for each method are reported in Table 3. One

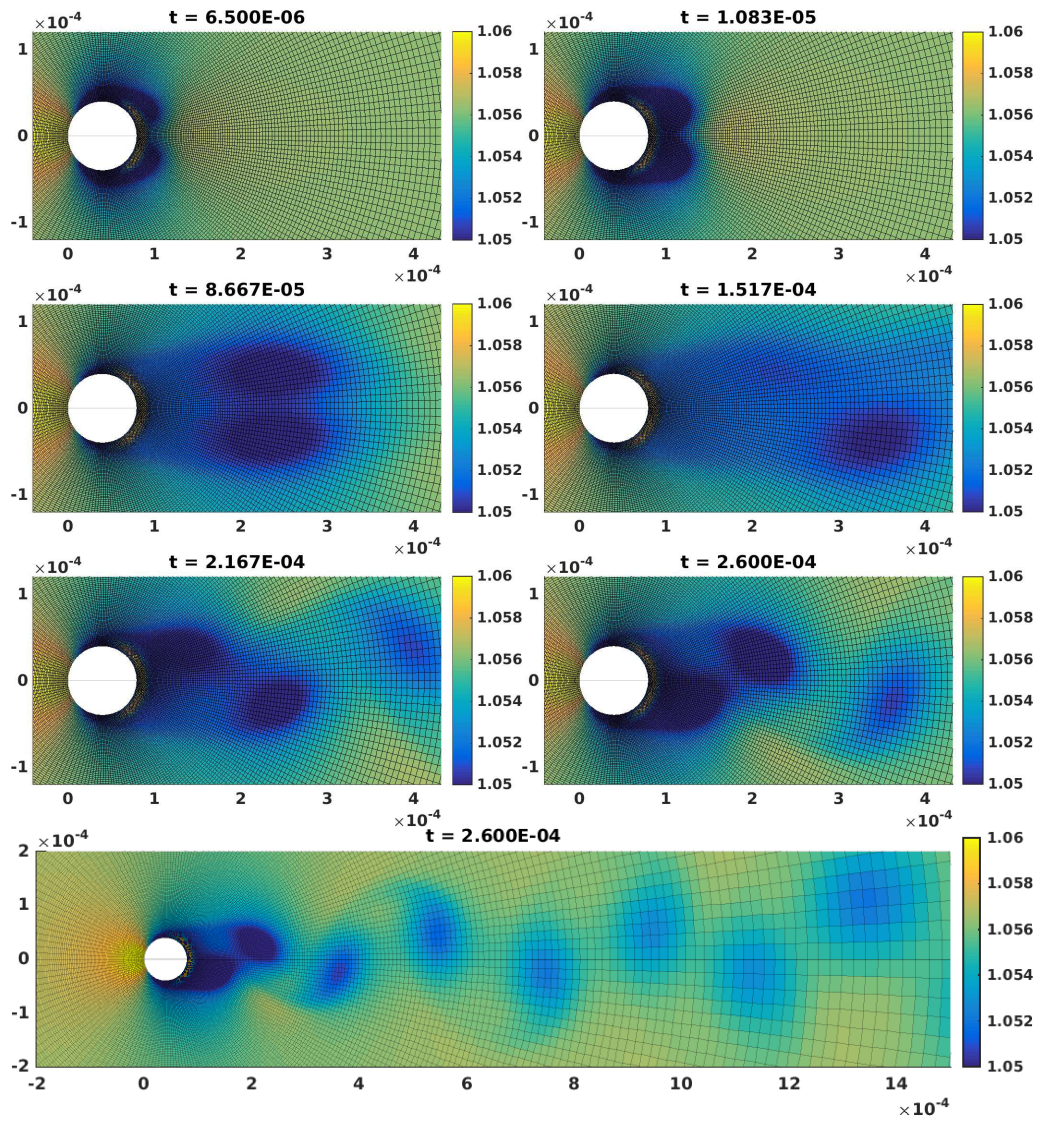


Figure 1: Snapshots of density component of the flow for vortex-shedding cylinder test problem 1 at different times.

Table 2: Parameters of the vortex-shedding cylinder test problems.

Parameter	Description	Test problem 1 value	Test problem 2 value
$\rho(kg/m^3)$	Reference fluid density	1.0565	1.0565
S	Sutherland's coefficient	1.45E-6	2.9E-6
$T(K)$	Temperature	278	278
$v(m/s)$	Reference velocity	340	340
$L(m)$	Diameter	8E-5	8E-5
Re	Reynold's number	165.90	82.95

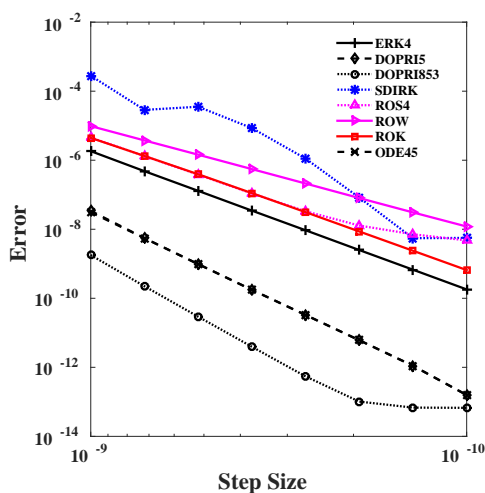


Figure 2: Convergence diagrams for the matrix-free methods of Table 1 using fixed step sizes over integration window $T = [0, 10^{-7}]$ second for cylinder test problem 1.

notable observation is the significantly lower numerical order of convergence for SDIRK method as a result of poor convergence of the Newton iteration, especially in the absence of preconditioners for the solution of linear systems. Readers interested in numerical experiments on a smaller problem that verify the theoretical order of convergence may consult the Appendix.

The cylinder test problem 2 with parameters given in Table 2, uses a different value for Sutherland's Law coefficient that translates into increased kinetic viscosity compared to vortex-shedding cylinder test problem 1. Figure 3 shows the numerical approximation of the first 1000 eigenvalues for the two test problems. As indicated in Figure 3 the more viscous test problem shows larger negative eigenvalues, therefore, we expect stricter stability bounds on the step sizes for this test problem. For both problems the cluster of eigenvalues with the largest negative real parts consists of only a limited number of modes. This fact becomes relevant when we consider ROK methods that use a reduced

Table 3: Numerical orders of convergence for various methods applied to the cylinder test problem 1.

Method	Numerical order	Theoretical order
ERK4	4.02	4
ERK5 (DOPRI5)	5.29	5
ERK5 (DOPRI853)	5.97	8
SDIRK	2.94	4
ROS4	3.11	4
ROW	2.96	3
ROK	3.85	4
ODE45	5.39	5

order Jacobian for implicit integration. Throughout the numerical experiments the ROK method employs a reduced space of dimension four, unless otherwise specified.

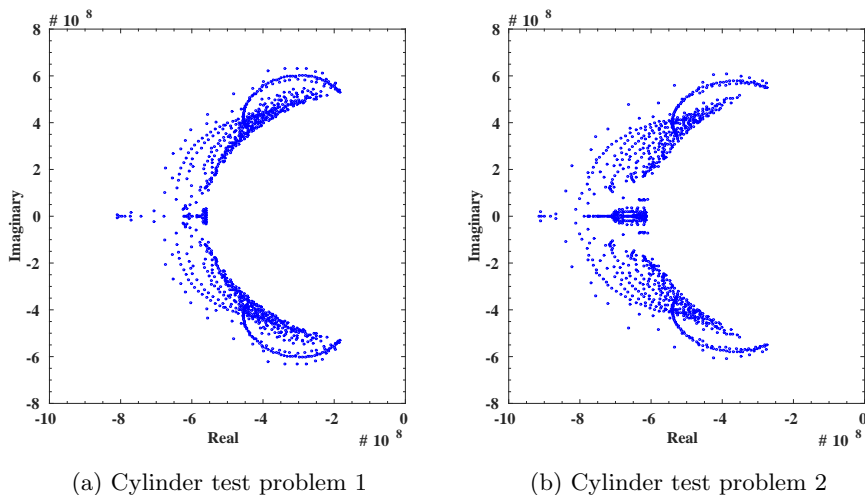
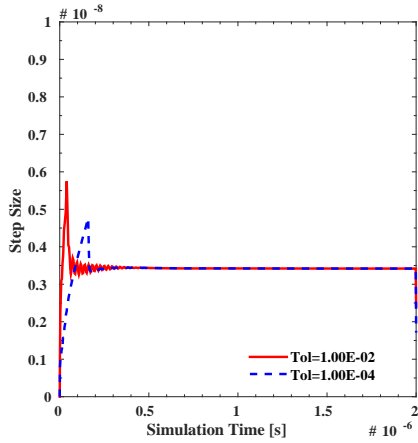
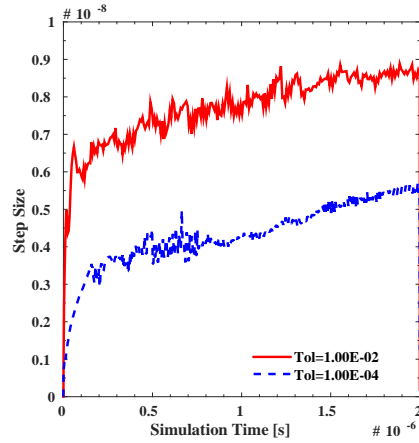


Figure 3: Eigenvalue distribution of the Jacobians of the right-hand-side functions for vortex-shedding cylinder test problems.

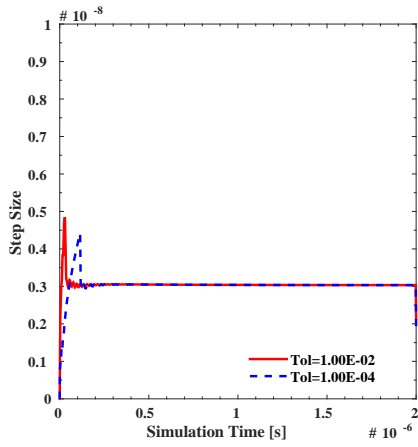
The stability of a method is related to its choice of timesteps in the adaptive time stepping framework. Figure 4 compares the step sizes of fourth order explicit Runge-Kutta method (ERK) to Rosenbrock-Krylov method of the same order for two different error tolerances. We can verify that the step sizes for the explicit method quickly reaches the upper bound set by the stability constraints regardless of the accuracy tolerance chosen for the method. On the other hand, by implicitly treating some of the stiff modes, the ROK method is able to achieve stable numerical integration for larger step sizes as is clear from Figures 4b and 4d. Furthermore, we notice that the step sizes scale well relative to the required accuracy.



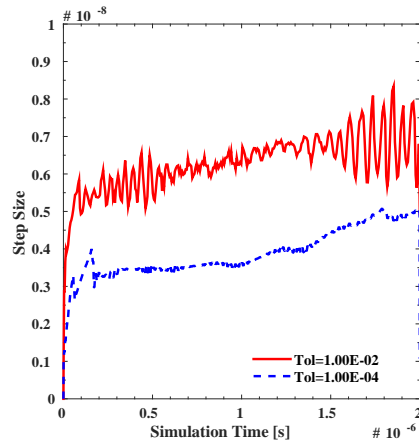
(a) ERK on cylinder test problem 1



(b) ROK on cylinder test problem 1



(c) ERK on cylinder test problem 2



(d) ROK on cylinder test problem 2

Figure 4: Adaptive time steps taken by the explicit Runge-Kutta method and by the ROK method for vortex shedding cylinder test problems.

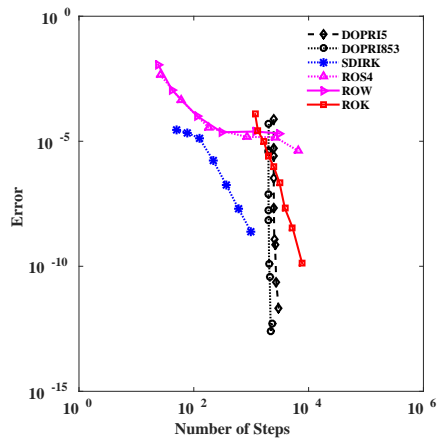
Figure 5 shows the Work-precision diagrams for the vortex-shedding cylinder test problems. Integration is performed over time window $T = [0, 2 \times 10^{-6}]$ seconds. These results lead to the following conclusions:

- As the adaptive time stepping method uses tighter tolerances, the integrator takes smaller steps leading to an increased number of total steps. This is the case for all integration methods presented here. However, the total number of steps for explicit methods does not change considerably for a wide range of tolerances, an effect observable in Figure 5 where the lines for explicit methods are nearly vertical. This is a result of the fact that for stiff problems the adaptive time steps are bounded by stability requirements rather than by accuracy constraints.
- Implicit methods are able to take fewer steps when the required solution tolerances are low, due to their improved stability properties. This is the case for ROS, ROW and SDIRK methods in Figure 5a and 5c. However, inspecting the runtime diagrams on Figures 5b and 5d reveals that the increased cost of these methods makes them considerably less efficient.
- The effect of stiffness of the problem can also be seen in the timing reported in Figures 5b and 5d. We notice that while the explicit methods take about the same amount of time to complete the integration for all choices of solution tolerance, the runtime scales better with tolerances for implicit methods.
- ROK is the most efficient method for the cylinder test problems for error tolerances below 10^{-6} . This is an indication that ROK is able to capture sufficiently many stiff components in its Krylov subspace, and that by treating them implicitly the method is able to take larger time steps.

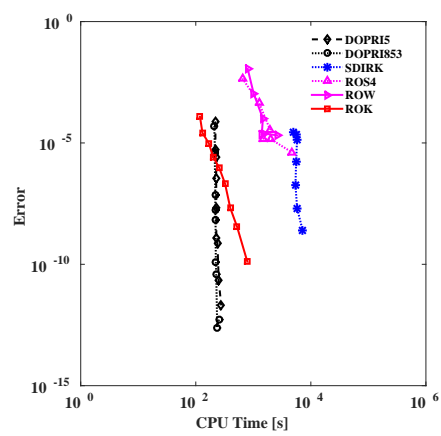
4.3. Flow over NACA0012 wing test problem

An analysis of unsteady flow over a NACA0012 airfoil is performed in addition to the cylinder test problem. Similar free-stream flow conditions are applied to this test problem, using atmospheric conditions at 5,000 ft. The Mach number is increased to 0.25; however, this flow is still within the subsonic flow regime. The chord length of this symmetric airfoil geometry is set to 0.001 meters, producing flow with a Reynolds number close to 5,000. Experimental analysis suggests that for the NACA0012 airfoil, a free-stream flow at this Reynolds number produces predominantly laminar flow over the airfoil and in its wake [10], allowing for turbulent effects to be neglected in this analysis. At an angle of attack of 15 degrees, vortices shed into the wake due to laminar separation of the flow over the upper surface of the airfoil generating an unsteady flow solution. As with the cylinder case, Roe's flux scheme is used with no flux limiter.

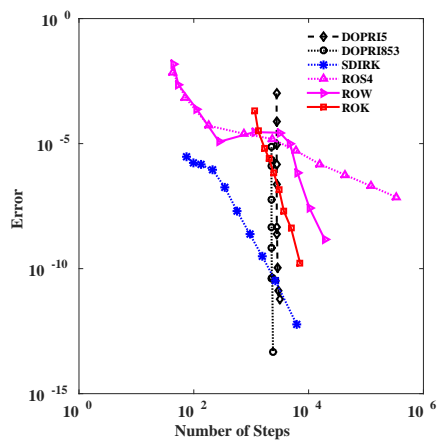
Figure 6 shows snapshots of the density field at different time moments. The distribution of Jacobian eigenvalues is shown in Figure 7, and indicates that a large number of fast eigenvalues are clustered together. The performance diagrams in Figure 8 indicate that the fastest methods here are Rosenbrock and ROW methods, until ROK and SDIRK become most efficient for errors below 10^{-4} . This test problem converges to a steady state solution, and this particular dynamics favors the implicit Rosenbrock methods, since the fast modes need only to be only damped out and not to be accurately solved.



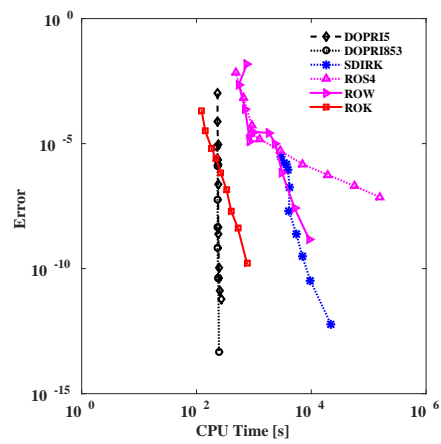
(a) Convergence diagram, cylinder test problem 1



(b) Work-precision diagram, cylinder test problem 1



(c) Convergence diagram, cylinder test problem 2



(d) Work-precision diagram, cylinder test problem 2

Figure 5: Relative performance of different integration methods applied to the cylinder test problems.

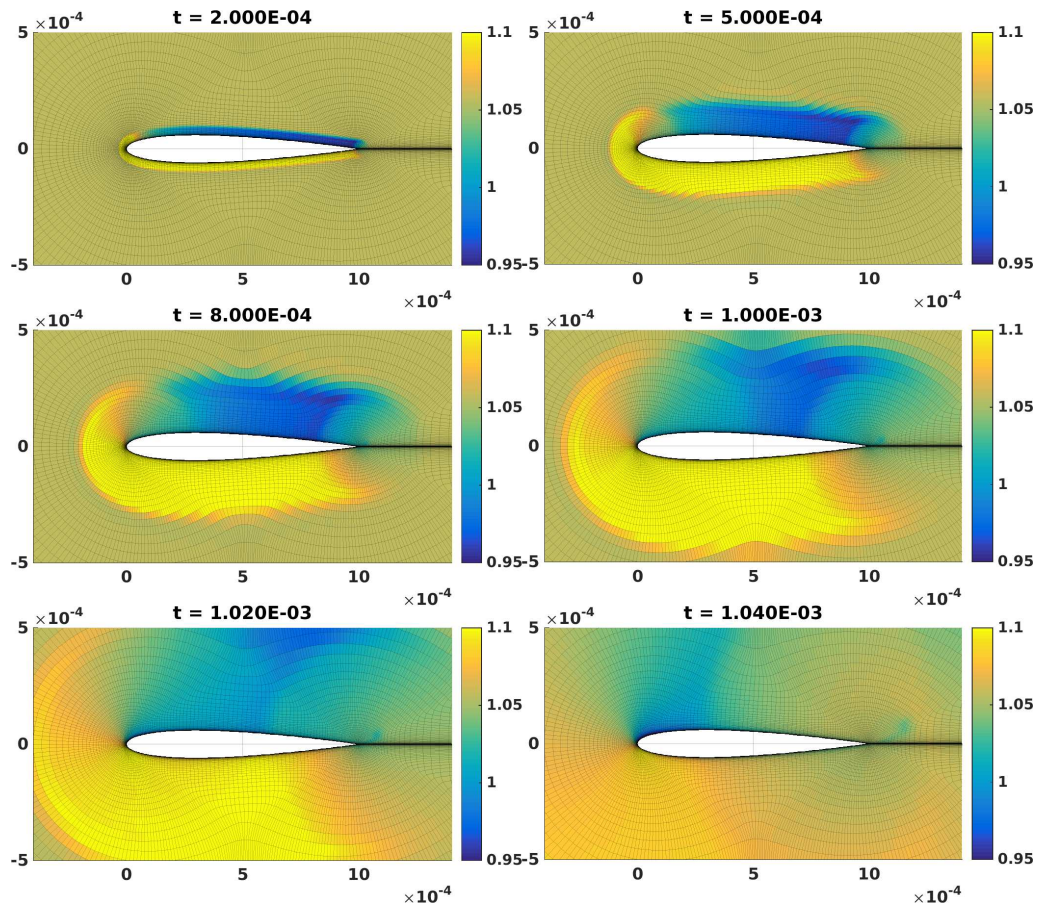


Figure 6: Snapshots of density component of the flow for NACA0012 at different times.

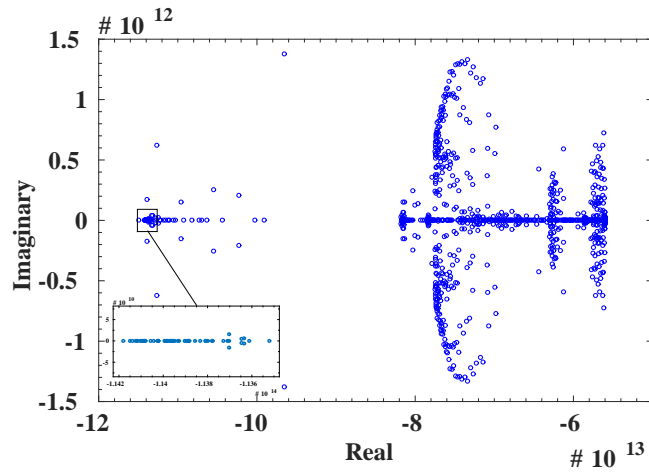


Figure 7: Eigenvalue distribution of the Jacobian for the NACA0012 test problem.

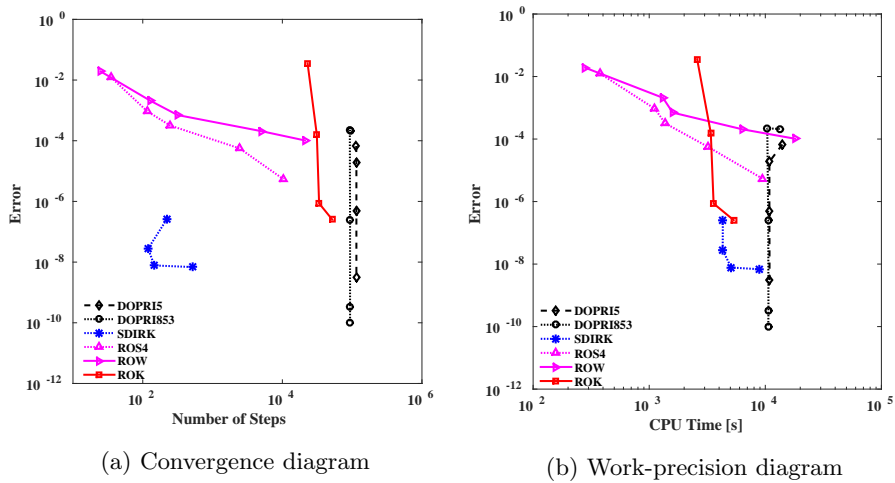


Figure 8: Relative performance of different integration methods applied to the NACA0012 test problem.

4.4. Vortex shedding cylinder with iso-thermal conditions test problem

The final test problem uses the vortex shedding cylinder problem 2, with iso-thermal conditions at cylinder boundaries for a wall temperature of $33^\circ K$. Among each class of explicit and implicit integrators we have selected the fastest representative methods. We have also chosen three variations of ROK method with 4, 8, and 12 Krylov basis vectors, respectively. The results in Figure 9 demonstrate that Krylov methods retain their computational superiority for a wide range of error tolerances. Furthermore, we observe that adding more basis vectors to the Krylov subspace requires extra cost (Figure 9b) and in turn increases step sizes slightly (Figure 9a), but ultimately the most efficient method is the one with minimum basis size, i.e. 4 vectors. Finally, it is worthwhile to point out that once the adaptive error controller of the step size reaches the GMRES tolerance, the error of JFNK methods does not decrease any further. This is seen in Figure 9 for the Rosenbrock-W method where the error curve flattens at an error level of about $5E-4$.

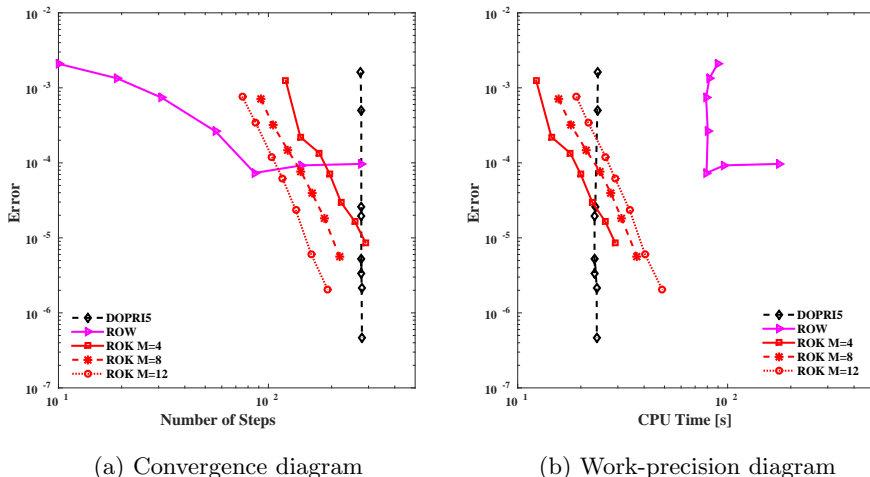


Figure 9: Relative performance of different integration methods applied to cooled cylinder test problem with $T_{\text{wall}} = 33^\circ K$.

5. Conclusions

This paper studies the performance of several types of high-order matrix-free time stepping methods when applied to solve unsteady flow problems. The methods under consideration are explicit Runge-Kutta, diagonally implicit Runge-Kutta and Rosenbrock schemes paired with iterative linear algebra solvers, and Rosenbrock-Krylov schemes. All implementations of implicit methods are matrix-free where the necessary Jacobian-vector products are approximated by finite differences; this is a typical setting for solving large-scale CFD applications.

Favorable properties of explicit methods include their easier implementation overhead and the low computational cost per step. As expected, our numerical experiments show

that the overall performance of explicit methods deteriorates quickly in the presence of mild stiffness, and for such problems they are not competitive with implicit methods.

Traditional implicit methods such as SDIRK and Rosenbrock can take large steps on stiff problems, as expected. However, their overall performance depends on how well the underlying linear systems are solved at each step. We observed that inaccurate linear solutions lead to loss of convergence, and that in absence of well-tuned linear algebra preconditioners the computational costs of matrix-free implicit methods are quite large.

Rosenbrock-Krylov methods are the most effective when applied to flow problems with a limited number of stiff modes. These methods are a suitable choice when exact full Jacobians are not available, as is the case in large CFD problems. The order conditions theory of Krylov-based methods accounts for the errors associated with linear algebra; indeed, our numerical experiments confirm that they show full order where other implicit methods suffer from reduced temporal convergence caused by inexact Jacobian approximations or by poorly converged linear system solutions. On the other hand, as the test problems become more stiff, Krylov-based methods lose their performance superiority due to the increased computational cost of creating a large Krylov basis required for numerical stability.

The numerical experiments in this study were performed on a 2D Navier-Stokes problem. Extension of the investigation to 3D unsteady flow problems is a future direction to pursue. Furthermore, finding automatic strategies for the cost-effective construction of Krylov basis is an important question for the advancement of Krylov-based family of methods.

6. Acknowledgments

This work was funded by the Air Force Office of Scientific Research (AFOSR) computational mathematics program via grant No. FA9550-12-1-0442 and by Computational Science Laboratory (CSL) in the Department of Computer Science at Virginia Tech.

7. Appendix: Order reduction with matrix-free methods

In this section we confirm the orders of convergence for the time-stepping methods used in this paper by applying them to Lorenz-96 test problem [36]. Proposed by Edward Lorenz, this model is often used to model chaotic behavior of atmospheric systems:

$$\frac{\partial u_i}{\partial t} = (u_{i+1} - u_{i-2}) u_{i-1} - u_i + F \quad \text{for } i = 1, \dots, 40.$$

The Jacobian of this system is a banded matrix that can be implemented in sparse format and used in the implicit methods of Section 2. Setting the external force factor $F = 8$ and using a range of fixed time steps to propagate the model forward, we observe nearly full order for all of the methods as reported in Table 4. One source of local error causing order reduction can be traced back to the truncation errors made by replacing the Jacobian-vector product in equation (10) with a first-order finite difference approximation. Poor convergence of the Newton's iteration as well as the Krylov-based solver for the linear system in the implicit methods are other sources of error contributing to inexact stage vectors and ultimately to order reduction. It is notable, however, that in

all the convergence tests the Rosenbrock-Krylov method retains its full order of convergence. Unlike fully and linearly implicit methods, the system solved in equation (9c) is formed using the Arnoldi iteration with an inexact Jacobian from the beginning and is solved using a direct method and therefore the issues arising from the convergence of the iterative solvers is avoided altogether. Interested readers may consult [28] for a detailed analysis of this phenomenon.

Table 4: Orders of convergence for methods applied to the Lorenz–96 problem.

Method	Numerical order	Theoretical order
ERK4	4.00	4
ERK5 (DOPRI5)	5.58	5
ERK5 (DOPRI853)	6.14	8
SDIRK	3.89	4
ROS4	3.91	4
ROW	2.94	3
ROK	3.85	4
ODE45	5.47	5

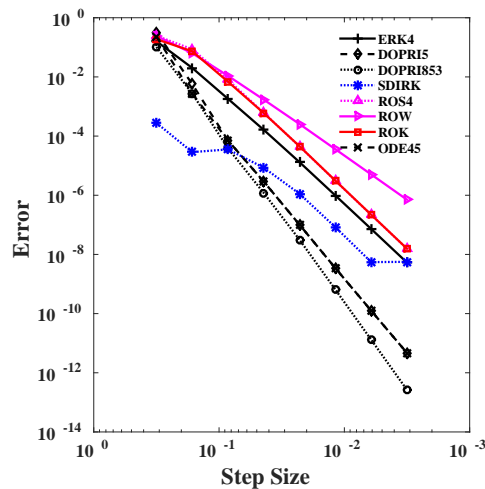


Figure 10: Convergence diagram for the Lorenz–96 test problem.

References

References

- [1] R. Yurkovich, Status of unsteady aerodynamic prediction for flutter of high-performance aircraft, *Journal of Aircraft* 40 (5) (2003) 832–842.

- [2] E. Lee-Rausch, J. Baitina, Wing flutter computations using an aerodynamic model based on the Navier–Stokes equations, *Journal of Aircraft* 33 (6) (1996) 1139–1147.
- [3] R. Niewoehner, J. Filbey, Computational study of F/A-18 E/F abrupt wing stall in the approach configuration, *Journal of Aircraft* 42 (6) (2005) 1516–1522.
- [4] Y. Hirato, M. Shen, S. Aggarwal, A. Gopalarathnam, J. Edwards, Initiation of leading-edge-vortex formation on finite wings in unsteady flow, 53rd AIAA Aerospace Sciences Meeting.
- [5] T. Nakata, H. Liu, Y. Tanaka, N. Nishihashi, X. Wang, A. Sato, Aerodynamics of a bio-inspired flexible flapping-wing micro air vehicle, *Bioinspiration and Biomimetics* 6 (4) (2011) 2–45.
- [6] Z. J. Wang, Unsteady forces and flows in low Reynolds number hovering flight: Two-dimensional computations vs robotic wing experiments, *The Journal of Experimental Biology* 207 (3) (2004) 449–460.
- [7] L.-J. Yang, F.-Y. Hsiao, W.-T. Tang, I.-C. Huang, 3D flapping trajectory of a Micro-Air-Vehicle and its application to unsteady flow simulation, *International Journal of Advanced Robotic Systems* 10 (2013) 264.
- [8] Flow Visualization in the Wake of the Flapping-Wing MAV 'DelFly II' in Forward Flight.
- [9] M. Sun, J. Tang, Unsteady aerodynamic force generation by a model fruit fly wing in flapping motion, *The Journal of Experimental Biology* 205 (1) (2002) 55–70.
- [10] R. F. Huang, C. L. Lin, Vortex shedding and shear-layer instability of wing at low-Reynolds numbers, *AIAA Journal* 33 (8) (1995) 1398–1403.
- [11] J. Choi, T. Colonius, D. Williams, Surging and plunging oscillations of an airfoil at low Reynolds number, *Journal of Fluid Mechanics* 763 (2015) 237–253.
- [12] T. von Kármán, *Aerodynamics*, McGraw-Hill, 1963.
- [13] C. A. Kennedy, M. H. Carpenter, R. M. Lewis, Low-storage, explicit Runge–Kutta schemes for the compressible Navier–Stokes equations, *Applied Numerical Mathematics* 35 (3) (2000) 177–219. doi:10.1016/S0168-9274(99)00141-5. URL <http://www.sciencedirect.com/science/article/pii/S0168927499001415>
- [14] P. Van der Houwen, Explicit Runge–Kutta formulas with increased stability boundaries, *Numerische Mathematik* 20 (2) (1972) 149–164.
- [15] H. Bijl, M. H. Carpenter, V. N. Vatsa, C. A. Kennedy, Implicit time integration schemes for the unsteady compressible Navier–Stokes equations: Laminar Flow, *Journal of Computational Physics* 179 (1) (2002) 313–329. doi:10.1006/jcph.2002.7059. URL <http://www.sciencedirect.com/science/article/pii/S0021999102970592>
- [16] G. Jothiprasad, D. J. Mavriplis, D. A. Caughey, Higher-order time integration schemes for the unsteady Navier–Stokes equations on unstructured meshes, *Journal of Computational Physics* 191 (2). doi:10.1016/S0021-9991(03)00330-9. URL <http://www.sciencedirect.com/science/article/pii/S0021999103003309>
- [17] P. Birken, G. Gassner, M. Haas, C.-D. Munz, Efficient time integration for discontinuous galerkin methods for the unsteady 3D Navier–Stokes equations, in: *European Congress on Computational Methods and Applied Sciences and Engineering (ECCOMAS 2012)*, number Eccomas, 2012.
- [18] N. Qin, D. K. Ludlow, S. T. Shaw, A matrix-free preconditioned Newton/GMRES method for unsteady Navier–Stokes solutions, *International Journal for Numerical Methods in Fluids* 33 (2) (2000) 223–248. doi:10.1002/(SICI)1097-0363(20000530)33:2<223::AID-FLD10>3.0.CO;2-V. URL [http://onlinelibrary.wiley.com/doi/10.1002/\(SICI\)1097-0363\(20000530\)33:2<223::AID-FLD10>3.0.CO;2-V/abstract](http://onlinelibrary.wiley.com/doi/10.1002/(SICI)1097-0363(20000530)33:2<223::AID-FLD10>3.0.CO;2-V/abstract)
- [19] J. Butcher, A history of Runge–Kutta methods, *Applied Numerical Mathematics* 20 (247–260).
- [20] W. H. Enright, D. J. Higham, B. Owren, P. W. Sharp, A survey of the explicit Runge–Kutta method (1994).
- [21] P. C. Jorgenson, R. C. Chima, Explicit Runge–Kutta method for unsteady rotor/stator interaction, *AIAA journal* 27 (6) (1989) 743–749.
- [22] C. A. Kennedy, M. H. Carpenter, R. M. Lewis, Low-storage, explicit Runge–Kutta schemes for the compressible Navier–Stokes equations, *Applied Numerical Mathematics* 35 (3) (2000) 177–219. doi:10.1016/S0168-9274(99)00141-5. URL <http://www.sciencedirect.com/science/article/pii/S0168927499001415>
- [23] E. Hairer, S. Norsett, G. Wanner, *Solving ordinary differential equations I: Nonstiff problems*, no. 8 in Springer Series in Computational Mathematics, Springer-Verlag Berlin Heidelberg, 1993.
- [24] E. Hairer, G. Wanner, *Solving ordinary differential equations II: Stiff and differential-algebraic problems*, 2nd Edition, no. 14 in Springer Series in Computational Mathematics, Springer-Verlag Berlin Heidelberg, 1996.
- [25] C. Lubich, A. Ostermann, Linearly implicit time discretization of nonlinear parabolic equations,

- IMA Journal on Numerical Mathematics 15 (1995) 555–583.
- [26] J. Rang, L. Angermann, New Rosenbrock W-methods of order 3 for partial differential algebraic equations of index 1, BIT Numerical Mathematics 45 (4) (2005) 761–787.
- [27] Y. Saad, Iterative methods for sparse linear systems, Siam, 2003.
- [28] P. Tranquilli, A. Sandu, Rosenbrock-Krylov methods for large systems of differential equations, SIAM Journal on Scientific Computing 36 (3) (2014) A1313–A1338.
URL <http://dx.doi.org/10.1137/130923336>
- [29] J. M. Derlaga, T. S. Phillips, C. J. Roy, V. Tech, SENSEI computational fluid dynamics code: a case study in modern fortran software development, in: 21st AIAA Computational Fluid Dynamics Conference, 2013, p. 2450.
- [30] C. B. Laney, Computational gasdynamics, Cambridge University Press, 1998.
- [31] B. van Leer, Towards the ultimate conservative difference scheme. v. a second-order sequel to Godunov’s method, Journal of Computational Physics 32 (1) (1979) 101–136.
- [32] D. Knoll, D. Keyes, Jacobian-free Newton–Krylov methods: a survey of approaches and applications, Journal of Computational Physics 193 (2004) 357 – 397. doi:<http://dx.doi.org/10.1016/j.jcp.2003.08.010>.
URL <http://www.sciencedirect.com/science/article/pii/S0021999103004340>
- [33] P. Tranquilli, S. R. Glandon, A. Sarshar, A. Sandu, Analytical jacobian-vector products for the matrix-free time integration of partial differential equations, Journal of Computational and Applied Mathematics (2016) –doi:<http://dx.doi.org/10.1016/j.cam.2016.05.002>.
URL <http://www.sciencedirect.com/science/article/pii/S0377042716302199>
- [34] A. Sandu, J. Linford, V. Rao, T. D’Augustine, Matlab integration package for stiff ODEs (2011).
URL <http://people.cs.vt.edu/~asandu/Software/Mat1ODE/mat1ode.html>
- [35] J. R. Dormand, M. A. Lockyer, N. E. McCorrigan, P. J. Prince, Global error estimation with Runge-Kutta triples, Computers & Mathematics with Applications 18 (9) (1989) 835–846.
- [36] E. N. Lorenz, Predictability: A problem partly solved, in: Proc. Seminar on predictability, Vol. 1, 1996, pp. 41–58.