Optimized Runge-Kutta Methods with Automatic Step Size Control for Compressible Computational Fluid Dynamics

Hendrik Ranocha¹ · Lisandro Dalcin² · Matteo Parsani² · David I. Ketcheson³

Received: 14 April 2021 / Revised: 21 July 2021 / Accepted: 23 July 2021 /

Published online: 22 November 2021

© The Author(s) 2021

Abstract

We develop error-control based time integration algorithms for compressible fluid dynamics (CFD) applications and show that they are efficient and robust in both the accuracy-limited and stability-limited regime. Focusing on discontinuous spectral element semidiscretizations, we design new controllers for existing methods and for some new embedded Runge-Kutta pairs. We demonstrate the importance of choosing adequate controller parameters and provide a means to obtain these in practice. We compare a wide range of error-control-based methods, along with the common approach in which step size control is based on the Courant-Friedrichs-Lewy (CFL) number. The optimized methods give improved performance and naturally adopt a step size close to the maximum stable CFL number at loose tolerances, while additionally providing control of the temporal error at tighter tolerances. The numerical examples include challenging industrial CFD applications.

Keywords Explicit Runge-Kutta methods · Step size control · Compressible Euler equations · Compressible Navier-Stokes equations · *hp*-adaptive spatial discretizations

Mathematics Subject Classification 65L06 · 65M20 · 65M70 · 76M10 · 76M22 · 76N99 · 35L50

✉ Hendrik Ranocha
mail@ranocha.de

Lisandro Dalcin
dalcinl@gmail.com

Matteo Parsani
matteo.parsani@kaust.edu.sa

David I. Ketcheson
david.ketcheson@kaust.edu.sa

¹ Applied Mathematics Münster, University of Münster, Münster, Germany

² Computer Electrical and Mathematical Science and Engineering Division (CEMSE), Extreme Computing Research Center (ECRC), King Abdullah University of Science and Technology (KAUST), Thuwal 23955-6900, Saudi Arabia

³ Computer Electrical and Mathematical Science and Engineering Division (CEMSE), King Abdullah University of Science and Technology (KAUST), Thuwal 23955-6900, Saudi Arabia

1 Introduction

Systems of hyperbolic conservation laws are used to model many areas of science and engineering, such as fluid dynamics, acoustics, and electrodynamics. In practical applications, these systems must often be solved numerically. Explicit Runge-Kutta schemes are the most commonly used time discretizations for hyperbolic partial differential equations (PDEs), because of their efficiency and parallel scalability [24, 46, 51]. Overall efficiency of the method also depends on choosing a time step that is as large as possible while still satisfying stability and accuracy requirements. Since stability requirements are frequently more restrictive in this setting, hyperbolic PDE practitioners often adapt the time step size based on a desired Courant-Friedrichs-Lewy (CFL) number. The CFL number involves the ratio of the maximum characteristic speed to the mesh spacing, which is essentially a proxy for the norm of the Jacobian. The optimal CFL number depends on the space and time discretizations chosen, and possibly on the problem; it is often determined by trial and error.

On the other hand, time integration research has long emphasized the efficiency of error-based step size control. Much effort has gone into the design of embedded Runge-Kutta pairs and step size controllers for this purpose. Compared to CFL-based control, error-based control has the advantage of not requiring a manually-tuned CFL number and allowing for control of the temporal error when necessary. CFL-based control has the advantage of (usually) yielding near-optimal efficiency once the appropriate CFL value has been found, as long as the calculation is indeed stability-limited. Error-based step size control for convection-dominated problems has been attempted previously; see e.g., [7, 80]. An ideal time integration algorithm would achieve the efficiency of the CFL-based controller in the stability-limited regime without the need for manually-tuned parameters, while automatically reducing the step size if error control becomes a more restrictive requirement. In this work, we develop such algorithms in the context of computational fluid dynamics (CFD).

Specifically, we focus on low-storage Runge-Kutta pairs (reviewed in Sect. 2) combined with PID step size controllers (reviewed in Sect. 3) and spectral element methods. Spectral element methods can be very efficient for large-scale computations [5, 27, 32, 34, 79]. Because stability is a challenging issue for these schemes, a lot of effort has been devoted to developing energy stable (linearly stable) [3, 54, 78], and entropy stable (nonlinearly stable) spatial discretizations [13, 15, 20, 22, 23, 67, 71, 72]. Stable fully-discrete schemes can be obtained from these semi-discretizations by using a slight modification of classical time integration schemes, based on the relaxation approach [41, 63, 65, 66].

In the paradigm of CFL-based error control, a common approach to time integrator design is to seek a large region of absolute stability (see e.g., [21] for a recent example of this approach in the context of CFD). For error-based control, a large region of absolute stability is again important (for both the main method and the embedded method). Additionally, when automatic error control is used with step sizes near the stability limit, the concept of step size control stability becomes crucial to the design of the controller. We demonstrate the importance of choosing good step size controllers in Sect. 4. There exists some previous work on developing error-based step size control techniques for convection-dominated problems, principally by Berzins et al. [7, 80].

We compare some existing Runge-Kutta pairs in Sect. 5, and develop optimized Runge-Kutta pairs for discontinuous spectral element semidiscretizations of hyperbolic conservation laws in Sect. 6. The spectral element methods applied for the numerical experiments are implemented in the *hp*-adaptive, unstructured, curvilinear grid solver SSDC [55]. SSDC is built on top of the Portable and Extensible Toolkit for Scientific computing (PETSc) [6], its mesh topology abstraction (DMPLEX) [44], and its scalable ODE/DAE solver library [1]. Further details on the spatial semidiscretizations can be found in [13, 14, 20, 56, 57, 67]. We perform numerical experiments using the novel schemes in Sect. 7, both for the compressible Euler and Navier-Stokes equations. Finally, we summarize and discuss our results in Sect. 8. We contributed our optimized methods to the freely available open source library DifferentialEquations.jl [62] written in Julia [8].

2 Runge-Kutta Methods and Adaptive Time Stepping

Using the method of lines, a spatial semidiscretization of a hyperbolic PDE yields an ordinary differential equation (ODE) system

$$\begin{cases} \frac{d}{dt}u(t) = f(t, u(t)), & t \in (0, T), \\ u(0) = u_0, \end{cases} \quad (1)$$

where $u : [0, T] \rightarrow \mathbb{R}^m$ and m is the number of degrees of freedom in the spatial discretization. An explicit first-same-as-last (FSAL) Runge-Kutta pair with s stages can be described by its Butcher tableau [12, 28]

$$\begin{array}{c|c} c & A \\ \hline & b^T \\ & \widehat{b}^T \end{array} \quad (2)$$

where $A \in \mathbb{R}^{s \times s}$ is strictly lower-triangular, $b, c \in \mathbb{R}^s$, and $\widehat{b} \in \mathbb{R}^{s+1}$. For (1), a step from $u^n \approx u(t_n)$ to $u^{n+1} \approx u(t_{n+1})$, where $t_{n+1} = t_n + \Delta t_n$, is given by

$$\begin{cases} y_i = u^n + \Delta t_n \sum_{j=1}^{i-1} a_{ij} f(t_n + c_j \Delta t_n, y_j), \\ u^{n+1} = u^n + \Delta t_n \sum_{i=1}^s b_i f(t_n + c_i \Delta t_n, y_i), & i \in \{1, \dots, s\}, \\ \widehat{u}^{n+1} = u^n + \Delta t_n \sum_{i=1}^s \widehat{b}_i f(t_n + c_i \Delta t_n, y_i) + \widehat{b}_{s+1} f(t_{n+1}, u^{n+1}). \end{cases} \quad (3)$$

Here, y_i are the stage values of the Runge-Kutta method and the difference $u - \widehat{u}$ is used to estimate the local truncation error. If $\widehat{b}_{s+1} = 0$, then (3) is an ordinary Runge-Kutta pair;

otherwise, it is referred to as an FSAL Runge-Kutta pair. The FSAL idea is to use the derivative of the new solution as an additional input for the error estimator [19]. If the step is accepted, this costs nothing since the value $f(t_{n+1}, u^{n+1})$ must be computed at the next step anyway. Usually, for a main method of order q , the embedded method is chosen to be of order $\hat{q} = q - 1$; i.e., the schemes are used in local extrapolation mode.

Remark 1 There are different notations for FSAL methods. A common alternative to our choice of using $A \in \mathbb{R}^{s \times s}$, $b, c \in \mathbb{R}^s$, and $\hat{b} \in \mathbb{R}^{s+1}$ is to embed the baseline s -stage Runge-Kutta method in a method with $s + 1$ stages and Butcher coefficients

$$\tilde{A} = \begin{pmatrix} A & 0 \\ b^T & 0 \end{pmatrix} \in \mathbb{R}^{(s+1) \times (s+1)}, \quad \tilde{b} = \begin{pmatrix} b \\ 0 \end{pmatrix} \in \mathbb{R}^{s+1}, \quad \tilde{c} = \begin{pmatrix} c \\ 1 \end{pmatrix} \in \mathbb{R}^{s+1}. \quad (4)$$

Then, the last row of $\tilde{A}$ is equal to $\tilde{b}$ and $\hat{b} \in \mathbb{R}^{s+1}$ can be defined as usual.

The common assumption $\sum_j a_{ij} = c_i$ is used throughout this article. For methods with error-based step size control, the initial step size is chosen using the algorithm described in [28, p. 169].

2.1 Low-Storage Methods

A typical Runge-Kutta implementation requires simultaneous storage of all of the stages and/or their derivatives. Each stage or derivative occupies m words; we refer to this amount of storage (sufficient for holding a copy of the solution on the spatial grid at one point in time) as a register. A low-storage Runge-Kutta method is one that can be implemented using only a few registers; herein we consider methods that require just three or four registers. Note that three registers are the fewest possible if one requires an error estimator and the ability to reject a step.

We consider the low-storage method classes (with and without the FSAL technique):

- $3S^*$: three-register methods that include an error estimate;
- $3S^*_{+}$: three-register methods that require a fourth register for the error estimate.

Let S_j denote a given storage register. The class $3S^*$ methods, introduced in [40], can be implemented using only three storage registers of size m if assignments of the form

$$S_j \leftarrow S_j + f(t, S_j) \quad (5)$$

can be made with only $m + o(m)$ memory. Otherwise, an additional register is required. The $3S^*$ method family is parameterized by the coefficients $c_i, \gamma_{1,i}, \gamma_{2,i}, \gamma_{3,i}, \beta_i, \delta_i$ and can be implemented as described in Algorithm 1. We will also use $3S^*$ or $3S^*_{+}$ to denote some strong stability preserving (SSP) Runge-Kutta methods that can be implemented using a slight modification of these algorithms, as described in [39].

The FSAL technique has been applied to low-storage methods in [53]; these schemes append an FSAL stage to $3S^*$ methods to get more coefficients for the embedded error estimator.

Algorithm 1 Minimum storage implementation of 3S* methods

```

 $S_1 \leftarrow u^n, S_2 \leftarrow 0, S_3 \leftarrow u^n$ 
for all  $i \in \{1, \dots, s\}$  do
     $S_2 \leftarrow S_2 + \delta_i S_1$ 
     $S_1 \leftarrow \gamma_{1,i} S_1 + \gamma_{2,i} S_2 + \gamma_{3,i} S_3 + \beta_i \Delta t_n f(t_n + c_i \Delta t_n, S_1)$ 
end for
 $u^{n+1} \leftarrow S_1$ 
 $\hat{u}^{n+1} \leftarrow (S_2 + \delta_s S_1 + \delta_{s+1} S_3) / \left( \sum_{i=1}^{s+1} \delta_i \right)$ 
if  $\hat{b}_{s+1} \neq 0$  then
     $\hat{u}^{n+1} \leftarrow \hat{u}^{n+1} + \hat{b}_{s+1} \Delta t_n f(t_{n+1}, u^{n+1})$ 
end if

```

All 3S*₊ methods use an additional storage location for the embedded error estimator. If the embedded method is not used, they reduce to 3S* methods without an embedded scheme. Their low-storage implementation is delineated in Algorithm 2.

Algorithm 2 Minimum storage implementation of 3S*₊ methods

```

 $S_1 \leftarrow u^n, S_2 \leftarrow 0, S_3 \leftarrow u^n, S_4 \leftarrow u^n$ 
for all  $i \in \{1, \dots, s\}$  do
     $S_2 \leftarrow S_2 + \delta_i S_1$ 
     $S_1 \leftarrow \gamma_{1,i} S_1 + \gamma_{2,i} S_2 + \gamma_{3,i} S_3 + \beta_i \Delta t_n f(t_n + c_i \Delta t_n, S_1)$ 
     $S_4 \leftarrow S_4 + \hat{b}_i \Delta t_n f(t_n + c_i \Delta t_n, S_1)$ 
end for
 $u^{n+1} \leftarrow S_1$ 
 $\hat{u}^{n+1} \leftarrow S_4$ 
if  $\hat{b}_{s+1} \neq 0$  then
     $\hat{u}^{n+1} \leftarrow \hat{u}^{n+1} + \hat{b}_{s+1} \Delta t_n f(t_{n+1}, u^{n+1})$ 
end if

```

2.2 Error-Based Step Size Control

We use step size controllers based on digital signal processing [25, 26, 73–75] implemented in PETSc [1, 6]. In particular, we use PID controllers that select a new time step using the formula¹

$$\Delta t_{n+1} = \epsilon_{n+1}^{\beta_1/k} \epsilon_n^{\beta_2/k} \epsilon_{n-1}^{\beta_3/k} \Delta t_n, \quad (6)$$

where q is the order of the main method, $\hat{q}$ is the order of the embedded method (usually $\hat{q} = q - 1$), $k = \min(q, \hat{q}) + 1$ (usually $k = \hat{q} + 1 = q$), β_i are the controller parameters, and

¹ Open source libraries such as PETSc [1] and DifferentialEquations.jl, [62] where we implemented a PID step size control, will usually limit the factor multiplying the time step size, e.g., using a limiter of the form $f(x) = 1 + \arctan(x - 1)$ [76].

Table 1 Classical step size controllers recommended in the literature

Controller	β_1	β_2	β_3
PI42	0.60	− 0.20	0.00
PI33	0.66	− 0.33	0.00
PI34	0.70	− 0.40	0.00

$$\epsilon_{n+1} = \frac{1}{w_{n+1}}, \quad w_{n+1} = \left(\frac{1}{m} \sum_{i=1}^m \left(\frac{u_i^{n+1} - \hat{u}_i^{n+1}}{\text{atol} + \text{rtol} \max\{|u_i^{n+1}|, |\hat{u}_i^{n+1}|\}} \right)^2 \right)^{1/2}, \quad (7)$$

where m is the number of degrees of freedom in u , and atol , rtol are the absolute and relative error tolerances, respectively. Some common controller parameters recommended in the literature are given in Table 1. Unless stated otherwise, we use equal absolute and relative error tolerances. The choice of the weighted/relative error estimate w_{n+1} is common in the literature [28, Equations (4.10) and (4.11)] and often the default choice in general purpose ODE software such as PETSc [1] or DifferentialEquations.jl [62]. This choice of w_{n+1} allows to decouple the time integration parameters from a possible spatial semidiscretization. In contrast to a quadrature-based approach, it weighs degrees of freedom of different refinement levels in the same way, which can be beneficial, since refined regions (of interest) are not weighed less than coarse regions (without interesting solution features).

If the factor multiplying the old time step Δt_n is too small or the solution is out of physical bounds, e.g., because of negative density/pressure in CFD, the step is rejected and retried with a smaller time step Δt_n . The default options used in all numerical experiments described in this work accept a step if the factor multiplying the step size is at least 0.9^2 . Otherwise, the step is rejected and retried with the step size predicted by the PID controller. If the solution is out of physical bounds, the step is rejected and retried with a time step reduced by a factor of four.

3 CFL- vs. Error-Based Step Size Control

The error-based step size control described above is efficient if the practical time step is limited by the constraint of accuracy. On the other hand, if the allowable time step is determined by stability, and an explicit time discretization is employed, then it is natural to use a step size of the form $\Delta t_n \propto 1/L_n$, where L_n is an approximation of the norm of the Jacobian of the ODE system.

In the time integration of hyperbolic PDEs, it is indeed often the case that the step size is limited in practice by stability rather than accuracy. Therefore it is common practice to use a step size control of the kind just described. For such systems, the norm of the Jacobian is proportional to $\max_i(\lambda_{\max}(u_i^n)/\Delta x_i)$, where Δx_i is a local measure of the mesh spacing (at grid point/cell/element i), and $\lambda_{\max}$ is the maximal (local) wave speed, related to the largest-magnitude eigenvalue of the flux Jacobian of the hyperbolic system. The step size control thus takes the form (referred to herein as a CFL-based control)

$$\Delta t_n = v \min_i \frac{\Delta x_i}{\lambda_{\max}(u_i^n)}, \quad (8)$$

Fig. 1 Unstructured grid used in the comparison of CFL- and error-based step size control

where ν is the desired CFL number. The appropriate choice of ν depends on the details of the space and time discretizations; it can be studied theoretically using linearization (see e.g., [49]) but is often determined experimentally. An additional complication is the question of how to define Δx . Even on uniform Cartesian grids and regular triangulations [48], multiple waves traveling in different directions make an optimal choice of ν difficult. This is even more a challenging question for unstructured grids.

A clear advantage of error-based control is the availability of an estimate of the temporal error. At first glance, error-based step size control seems inappropriate in the stability-limited regime, since the local error may not be very sensitive to small differences between stable and unstable step sizes, near the stable step size limit. A tight error tolerance that ensures stability at all steps might result in an excessively small step size. However, as described in [29, Section IV.2] and discussed below, it is possible to design error-based step size controllers that behave appropriately in the stability-limited regime.

Both classes of controllers require some user-determined parameters: ν and Δx for the CFL-based controller, and `atol` and `rtol` for the error-based controllers. In this section we show through an example that carefully-designed error-based controllers can achieve near-maximal efficiency in a way that is relatively insensitive to changes in the user parameters. In contrast, the efficiency of the CFL-based controller always bears a linear sensitivity to the parameters ν and Δx .

To demonstrate, we consider the two-dimensional advection equation with constant velocity $a = (1, 1)^T$ in the domain $[-5, 5]^2$ with periodic boundary conditions. An initial sinusoid of one wavelength in each direction is advected over the time interval $[0, 100]$. In space we apply the spectral collocation method of SSDC based on solution polynomials of degree $p = 4$ [55].

For the CFL-based controller, the ratio of the local mesh spacing and the maximal speed at a node i is estimated in this case as

$$\frac{\Delta x_i}{\lambda_{\max}(u_i^n)} = \sigma \frac{J_i}{\sum_{j=1}^d |(J\partial_x \xi^j)_i \cdot a|}, \quad (9)$$

where d is the number of spatial dimensions ($d = 2$ for this example), $a = (1, 1)^T$ is the constant advection velocity, J_i is the determinant of the grid Jacobian $\partial_x \xi$ at node i , $(J\partial_x \xi^j)_i$ is the contravariant basis vector in direction j at node i [45, Chapter 6], and σ is a normalizing factor depending on the solution polynomial degree, p , which is usually chosen such that a real stability interval of 2 corresponds to $\nu = 1$.

Two grids will be used: a regular, uniform grid with 8^2 elements and the curvilinear grid shown in Fig. 1. The results are summarized in Fig. 2 for the uniform grid and in Fig. 3 for the unstructured grid. We test three time discretizations: the popular fourth-order, five-stage, low-storage method CK4(3)5[2N] of [35]; the method KCL4(3)5[2R₊]C

Fig. 2 Performance of CFL- and error-based step size controls for a linear advection problem with $p = 4$ on a uniform grid. For CFL-based controllers, the maximal CFL number is always included. The error-based controller is a standard PI controller with $\beta_1 = 0.7$, $\beta_2 = -0.4$ and uses equal absolute and relative tolerances $\text{atol} = \text{rtol} = \text{tol}$

Fig. 3 Performance of CFL- and error-based step size controls for a linear advection problem with $p = 4$ on a nonuniform grid. For CFL-based controllers, the maximal CFL number is always included. The error-based controller is a standard PI controller with $\beta_1 = 0.7$, $\beta_2 = -0.4$ and uses equal absolute and relative tolerances $\text{atol} = \text{rtol} = \text{tol}$

of [37], which comes with an embedded third-order error estimator; and the SSP method SSP3(2)3[3S*₊] of [70], equipped with the embedded method of [18].

Remark 2 We use the same naming convention as [37], referring to an s -stage Runge-Kutta method of order q with embedded method of order $\hat{q}$ as NAME $q(\hat{q})_s$. Additional identifiers indicating low-storage requirements or other properties are appended, e.g., a subscript “F” for FSAL methods. The number of stages s denotes the effective number of RHS evaluations per step, which is one less than the number of stages for FSAL methods.

For low-storage methods, the required amount of memory based on certain assumptions is listed following the notation of [40]. In particular, nN methods need only n memory registers of size m if assignments of the form $S_j \leftarrow \alpha S_j + f(t, S_i)$ can be made without additional allocations. Similarly, nR methods use n memory registers and assignments of the form $S_j \leftarrow f(t, S_j)$; mS methods were described in Sect. 2.1. As described there, a subscript $+$ indicates methods that require an additional storage register if an embedded error estimator is used. Additional parts of the names of Runge-Kutta methods are usually taken directly from their sources. For example, the method $KCL4(3)5[2R_+]C$ is a fourth-order method with embedded third-order error estimator. It has five stages and requires two memory registers based on the nR assumption. If the embedded error estimator is used, it requires three memory registers. The C suffix is appended as suggested in [37] to indicate a particular design criterion (in this particular case, looking for a compromise between linear stability and accuracy).

The widely-used method $CK4(3)5[2N]$ has linear stability properties very similar to $KCL4(3)5[2R_+]C$ — both methods have the same maximum stable CFL number $\nu = 2.1$. Thus, they use the same number of RHS evaluations while yielding nearly the same errors. Thus, we omit the method $CK4(3)5[2N]$ in the following plots and use only the $KCL4(3)5[2R_+]C$ pair, for which it is possible to use error-based step size control. Impressively, the error-based controller manages, for a wide range of tolerances, to use almost exactly the same number of steps as the carefully tuned CFL-based controller. Over this range of tolerances, including $\text{tol} \in [10^{-5}, 10^{-3}]$, the step size is determined by stability. Hence, the number of RHS evaluations and the error are nearly independent of the tolerance in this regime. For tolerances larger than 10^{-3} , the final error increases for this long-time simulation while the number of RHS evaluations stays nearly the same. For tighter tolerances (below 10^{-6}), the error-based controller detects accuracy restrictions and increases the number of RHS evaluations. This also leads to a reduction of the final error until it plateaus again because of the dominant spatial error (at ca. $\text{tol} = 10^{-7}$). However, the number of RHS evaluations keeps increasing.

Using the same CFL number of $\nu = 2.1$ on the unstructured grid still results in a stable simulation. However, the CFL number can be doubled there without increasing the error significantly. Hence, the user has to tune this parameter carefully to get a stable and efficient simulation. In contrast, using error-based step size control we see behavior very similar to what was observed for the uniform grid. The same error tolerance can be used, resulting in the same optimal number of function evaluations determined manually for the CFL-based step size controller. This demonstrates the enhanced robustness properties of error-based step size control.

These examples suggest that error-based control is more robust to changes in the grid and less sensitive to the required user parameters. Similar results have been obtained using other Runge-Kutta schemes for this problem and for more challenging problems, some of which are presented later in this work. For practitioners whose primary interest is in applying the schemes to solve challenging scientific problems or developing spatial semidiscretizations, error-based time step controllers seem favorable, since the most important design choices have to be provided by the developers of the time integration schemes and the practitioners have to choose only the rather robust error tolerance of the solver.

4 Importance of Controller Parameters

Standard error-based controllers will often work acceptably in the asymptotic regime (i.e., the regime where the leading truncation error term strongly dominates all subsequent terms). However, as demonstrated in Sect. 3, applications involving convection-dominated problems are often constrained by stability, so that one may be working outside the asymptotic regime. In this case, the standard theory does not apply; instead, step size control stability has to be considered [30].

Following [29, Section IV.2], step size control stability can be explained using the linear model problem $\frac{d}{dt}u(t) = \lambda u(t)$. Given an explicit Runge-Kutta method with embedded error estimator and a PID controller (6) with parameters β_i , the update formulae become

$$\begin{cases} u^{n+1} = R(\Delta t_n \lambda) u^n, \\ e^{n+1} = E(\Delta t_n \lambda) u^n, \\ \Delta t_{n+1} = \text{PID}(\beta, \text{atol}, \text{rtol}, e^{n+1}, e^n, e^{n-1}), \end{cases} \quad (10)$$

where R is the stability polynomial of the main method, E is the difference of the stability polynomials of the embedded and the main method, and e is the (local) error estimate. By taking logarithms, this update formula can be reduced to a difference recursion with fixed points on the boundary of the stability region of the main method. To get a stable behavior, the spectral radius of the associated Jacobian has to be less than unity [29, Proposition IV.2.3]. For a PID controller (6), this Jacobian becomes [37]

$$J(z) = \begin{pmatrix} 1 & \frac{zR'(z)}{R(z)} & 0 & 0 & 0 & 0 \\ -\frac{\beta_1}{k} & 1 - \frac{\beta_1}{k} \frac{zE'(z)}{E(z)} & -\frac{\beta_2}{k} & -\frac{\beta_2}{k} \frac{zE'(z)}{E(z)} & -\frac{\beta_3}{k} & -\frac{\beta_3}{k} \frac{zE'(z)}{E(z)} \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}, \quad (11)$$

where k is the order of the error estimator; if $\hat{q} = q - 1$ is the order of the embedded method, then $k = q = \hat{q} + 1$. To get step size control stability, one can fix a controller such as the standard I controller and optimize the Runge-Kutta pair accordingly as demonstrated in [31]. The other possibility, pursued here, is to optimize the controller parameters for a given Runge-Kutta pair.

While one might hope that a controller designed to work well with one method will also work well with other methods, this is generally not the case. Rather, a controller should be designed for the given error estimator; cf. [4] for the case of linear multistep methods. To demonstrate this, we consider again the linear advection problem described in Sect. 3 with a uniform mesh. We will take the PI34 controller with $\beta_1 = 0.7$, $\beta_2 = -0.4$ [25], designed for use with the classical DP5(4)6_F method of [61], but use instead the BS5(4)7_F method of [10]. Note that both are 5(4) pairs designed with similar purposes in mind. Using a tolerance of $\text{tol} = 10^{-5}$, the integration requires 5 015 RHS evaluations and includes many rejected steps. Applying instead the optimized coefficient $\beta = (0.28, -0.23)$ derived later in this manuscript results in only 4 119 RHS evaluations and a nearly identical final error. A significant performance gain is obtained by applying appropriate controller parameters, cf. Table 2.

Table 2 Performance of different controllers for BS5(4)7_F: number of function evaluations (#FE) and rejected steps (#R) for the linear advection problem with uniform grid using polynomials of degree $p = 4$ as in Sect. 3

Scheme	β	tol	#FE	#R	Error
BS5(4)7 _F	(0.70, -0.40, 0.00)	10^{-5}	5 015	(132)	9.85E-07
	(0.28, -0.23, 0.00)	10^{-5}	4 119	(0)	9.79E-07

Fig. 4 Stability region scaled by the effective number of stages and spectral radius of the Jacobian (11) determining step size control stability for BS5(4)7_F. The standard PI34 controller is unstable near the negative real axis while the optimized PI controller is stable along the boundary of the stability region

The spectral radius of the Jacobian (11) determining step size control stability is plotted in Fig. 4. We see that the standard PI34 controller is unstable near the negative real axis while the optimized one is stable.

5 Comparison of Existing Methods

Here, we compare some general purpose methods and schemes designed for semidiscretizations of hyperbolic conservation laws. Since we are interested in error-based step size control, we consider only schemes with embedded error estimators. Hence, we consider the general purpose schemes

- BS3(2)3_F, third-order, four-stage FSAL method of [9],
- BS5(4)7_F, fifth-order, eight-stage FSAL method of [10],
- DP5(4)6_F, fifth-order, seven-stage FSAL method of [61],

the SSP schemes

- SSP3(2)3[3S*₊], third-order, three-stage SSP method of [70] with the embedded method of [18],

- SSP3(2)4[3S₊^{*}], third-order, four-stage SSP method of [47] with the embedded method of [18] which can be implemented efficiently in low-storage form as described in Appendix A,

and the low-storage methods optimized for hyperbolic conservation laws

- KCL3(2)4[2R₊]C, third-order, four-stage method of [37],
- KCL4(3)5[2R₊]C, fourth-order, five-stage method of [37],
- KCL4(3)5[3R₊]C, fourth-order, five-stage method of [37],
- KCL5(4)9[2R₊]S, fifth-order, nine-stage method of [37].

The results shown here, obtained with three commonly-used general purpose methods, are typical of what we have found in tests with a much wider range of methods. These results are sufficient to illustrate our main conclusions. We do not consider SSP4(3)10[3S₊^{*}] of [39] with the embedded method of [18] because step size control stability cannot be achieved for this method and any PID controller tested. The embedded method for SSP3(2)9[3S₊^{*}] proposed in [18] also does not lead to step size control stability. We have created a new embedded method with a stable optimized controller. However, it does not perform better than SSP3(2)4[3S₊^{*}], even with manually tuned CFL-based step size control.

In the following, we will use three representative test problems to compare the performance of these schemes. All test problems are semidiscretizations of the compressible Euler equations in d space dimensions

$$\partial_t u + \sum_{j=1}^d \partial_j f^j(u) = \sigma, \quad u(0) = u_0, \quad (12)$$

where the conserved variables $u = (\rho, \rho v^T, \rho e)^T$ are the density ρ , the momentum ρv , and the energy ρe . The flux for the spatial coordinate j is

$$f^j(u) = \begin{pmatrix} \rho v_j \\ (\rho v v_j + p \delta_{i,j})_i \\ (\rho e + p) v_j \end{pmatrix}, \quad (13)$$

where $p = \rho T = (\gamma - 1)(\rho e - \rho v^2/2)$ is the pressure, T is the temperature, and an ideal gas law with ratio of specific heats $\gamma = 7/5$ is assumed. The spatial semidiscretizations use entropy-dissipative nodal DG methods with polynomials of degree p on Legendre-Gauss-Lobatto nodes with upwind interface fluxes implemented in SSDC. We present detailed results for $p = 2$, which is a relevant choice in practical CFD applications. The results are similar for higher-order semidiscretizations such as polynomials of degree $p \in \{3, 4\}$, presented in the supplementary material in more detail.

5.1 Inviscid Taylor-Green Vortex

The inviscid Taylor-Green vortex in $d = 3$ space dimensions is a classical test case to study the stability of numerical methods [23]. The initial condition given by

$$\begin{cases} p(x) = 1 + \frac{1}{16} (\cos(2x_1) + \cos(2x_2)) (\cos(2x_3) + 2), \\ v_1(x) = \sin(x_1) \cos(x_2) \cos(x_3), \quad v_2(x) = -\cos(x_1) \sin(x_2) \cos(x_3), \\ v_3(x) = 0, \quad T(x) = 1, \end{cases} \quad (14)$$

with the Mach number $Ma = 0.1$ is evolved in the periodic domain $[-\pi, \pi]^3$. Unless stated otherwise, we use 8 elements per coordinate direction and the final time of $t = 20$. This test case is chosen as an example where the time step is mostly restricted by stability, the solution becomes turbulent, and a relatively low Mach number is used.

5.2 Isentropic Vortex

The isentropic vortex is a widely used benchmark problem [69] with the analytical solution. For the stationary case, the exact solution is given by

$$\begin{cases} T = T_\infty - \frac{(\gamma - 1)Ma_\infty^2 \beta^2}{8\gamma\pi^2} \exp(1 - r^2), \\ v_t = \frac{r\beta}{2\pi} \exp\left(\frac{1 - r^2}{2}\right), \quad \rho = T^{1/(\gamma-1)}, \end{cases} \quad (15)$$

where r is the distance from the axis of the vortex and v_t is the tangential velocity. The moving vortex solution is obtained by a uniform translation in the direction of the velocity vector field.

Herein, the simulation domain is a cube $[-5, 5]^3$ with periodic boundaries where the vortex rotates around the axis $(1, 1, 0)^T$, a direction not aligned with the grid. The parameters for this test are $\gamma = 1.4$, $Ma = 0.5$, $\beta = 5$ and $T_\infty = 1$. Unless stated otherwise, we use 8 elements per coordinate direction for optimizing controllers and 20 elements for examples with the final time of $t = 20$. This test case is chosen as an example where the time step can be restricted by accuracy for tight tolerances and because of the existence of an analytical solution.

5.3 Smooth Flow with Source Terms

The analytical solution

$$\rho(t, x) = \frac{3}{2} + \sin(\pi(x - t)), \quad v(t, x) = 1, \quad p(t, x) = 1 + A_p(1 + \sin(\omega_p t)) \quad (16)$$

is imposed as initial condition in the periodic domain $[-1, 1]$ and the source term

$$\sigma_{pe}(t, x) = \frac{A_p \omega_p}{\gamma - 1} \cos(\omega_p t) \quad (17)$$

is added to the right-hand side of the energy equation. The variation of the pressure with amplitude $A_p = 50$ and frequency $\omega_p = \pi/5$ results in a cyclic variation of the CFL restriction on the time step. Unless stated otherwise, we use 20 elements and the final time of

Table 3 Performance of general purpose schemes: number of function evaluations (#FE), rejected steps (#R), and L^2 error of the density for the inviscid Taylor-Green vortex (14), the isentropic vortex (15), and the flow with source term (16) using polynomials of degree $p = 2$

Scheme	β	tol	TGV (14)		Isent. vortex (15)			Source term (16)		
			#FE	#R	#FE	#R	Error	#FE	#R	Error
BS3(2)3 _F	(0.60, -0.20, 0.00)	10 ⁻⁵	5 256	(1)	1 476	(0)	5.77E-4	20 682	(4)	1.77E-3
BS5(4)7 _F	(0.28, -0.23, 0.00)	10 ⁻⁵	7 722	(1)	2 159	(0)	5.78E-4	30 356	(2)	1.77E-3
DP5(4)6 _F	(0.70, -0.40, 0.00)	10 ⁻⁵	8 014	(3)	2 217	(0)	5.78E-4	31 430	(4)	1.77E-3

$t = 20$. This test case is chosen to assess the ability of the schemes to adapt to varying time step restrictions and because of the existence of an analytical solution.

5.4 Optimization of Step Size Controllers

As explained in Sect. 4, the choice of appropriate step size controller parameters is important to obtain good performance when the schemes are run at the stability limit. Hence, we have optimized controller parameters for each scheme. In general, the optimal time step controller parameters for a given Runge-Kutta pair will depend somewhat on the problem under consideration. No single controller is optimal for all test cases, but for the experiments conducted in this work, good controllers are usually within ca. 5% of the optimal performance.

For the low-storage schemes of [37], we used the PI34 controller proposed originally with them. We also tested the PID controller using $\beta = (0.49, -0.34, 0.10)$ proposed in [36]. We also performed an optimization of controller parameters for each method, as follows.

We ran simulations of all three test cases described above and measured the performance of each scheme (in terms of the number of right-hand side evaluations). We used a brute-force search over the domain $\text{tol} \in [10^{-8}, 10^{-1}]$, sampling at each power of ten, and $\beta_1 \in [0.1, 1.0]$, $\beta_2 \in [-0.4, -0.05]$, $\beta_3 \in [0.0, 0.1]$ sampling at an interval of 0.01 in each parameter, and restricting a priori to parameter values yielding step size control stability for the given scheme (computed using NodePy [43]). The final time for these simulations was set to $t = 8$ for (14), $t = 4$ for (15), and $t = 20$ for (16) to make the brute-force optimization feasible. From the resulting data, consisting of thousands of runs with each method, an overall best choice of parameters was selected as in Sect. 6.1. Usually, this kind of min-max problem was approached by comparing the controllers minimizing the maximum, the median, or the 95% percentile of the RHS evaluations across all CFD simulations. Then, the final choice was made by human interaction taking into account step size control stability and design criteria for PID controllers.

Table 4 Performance of SSP schemes: number of function evaluations (#FE), rejected steps (#R), and L^2 error of the density for the inviscid Taylor-Green vortex (14), the isentropic vortex (15), and the flow with source term (16) using polynomials of degree $p = 2$

Scheme	β	τ_{01}	TGV (14)		Isent. vortex (15)			Source term (16)		
			#FE	#R	#FE	#R	Error	#FE	#R	Error
SSP3(2)3[3S* ₊]	(0.70, -0.37, 0.05)	10^{-5}	5 261	(3)	1 646	(0)	5.77E-4	20 681	(3)	1.77E-3
SSP3(2)4[3S* ₊]	(0.55, -0.27, 0.05)	10^{-5}	3 438	(1)	1 902	(0)	5.78E-4	20 638	(3)	1.77E-3

Table 5 Performance of low-storage schemes: number of function evaluations (#FE), rejected steps (#R), and L^2 error of the density for the inviscid Taylor-Green vortex (14), the isentropic vortex (15), and the flow with source term (16) using polynomials of degree $p = 2$

Scheme	β	τ_{01}	TGV (14)		Isent. vortex (15)			Source term (16)		
			#FE	#R	#FE	#R	Error	#FE	#R	Error
KCL3(2)4[2R ₊]C	(0.50, -0.35, 0.10)	10^{-5}	6 326	(1)	1 874	(0)	5.78E-4	24 886	(3)	1.77E-3
KCL4(3)5[2R ₊]C	(0.29, -0.24, 0.02)	10^{-5}	4 897	(65)	1 522	(0)	5.79E-4	23 902	(5)	1.77E-3
KCL4(3)5[3R ₊]C	(0.41, -0.28, 0.08)	10^{-5}	4 587	(42)	1 472	(0)	5.79E-4	23 817	(4)	1.77E-3
KCL5(4)9[2R ₊]S	(0.49, -0.34, 0.10)	10^{-5}	6 320	(3)	2 108	(0)	5.78E-4	30 374	(3)	1.77E-3

5.5 Results for Existing Schemes

For BS3(2)3_F, all of the controllers from Table 1 perform reasonably well, PI42 being slightly better than the others. In general, a wide range of controller parameters is acceptable for this scheme. As typified by the example in Sect. 4, standard controllers do not perform well for BS5(4)7_F. We found instead that $\beta = (0.28, -0.23, 0.00)$ is a reasonable choice for this scheme. For DP5(4)6_F, the PI34 controller (which was originally designed for it by Gustafsson [25]) performs reasonably well in our test cases and optimized controllers like $\beta = (0.61, -0.27, 0.01)$ do not perform significantly better.

Subsequently, we used the optimized controller parameters and ran full simulations (up to $t = 20$) for each method with a range of tolerances. Results are shown in Tables 3, 4, and 5, where polynomials of degree $p = 2$ have been used. There, we only show results for a tolerance $\tau_{01} = 10^{-5}$, since this choice is usually good for these small-scale test problems. Extended details are available in the supplementary material. For the inviscid Taylor-Green vortex, the time step is indeed restricted by stability for most tolerances, indicated by the approximately constant number of function evaluations, except for the very tight tolerance $\tau_{01} = 10^{-8}$ and some schemes. For the isentropic vortex (15), the step size is restricted by stability for tolerances $\gtrsim 10^{-6}$. For smaller tolerances, the number of function evaluations increases. However, this does not result in a significant change of the total error, which is determined mostly by the spatial semidiscretization. Finally, for the smooth flow with the source term (16), the step size is again restricted mostly by stability constraints.

5.5.1 General Purpose Methods

For very loose tolerances $\gtrsim 10^{-3}$, BS5(4)7_F and DP5(4)6_F result in a significant overhead caused by step rejections for some test cases. Otherwise, BS5(4)7_F performs better than DP5(4)6_F. Other fifth-order general purpose schemes like T5(4)6_F of [77] perform slightly better, usually yielding an improvement of ca. 5%. However, BS3(2)3_F is ca. 50% more efficient as long as the time step is restricted by stability.

These results do not change significantly if slightly higher-order semidiscretizations are employed in space (see supplementary material), up to polynomials of degree $p = 4$, resulting in fifth-order convergence in space. Hence, matching the order of accuracy in space and time is not strictly necessary if one is interested in fixed mesh sizes, especially in common CFD applications. This remains true even if the polynomial degree is increased to $p = 7$ for the test problems considered here. Then, the temporal error becomes significant and the error of the fully discrete method plateaus only at relatively tight tolerances such as 10^{-8} . Nevertheless, BS3(2)3_F is still the most efficient method for such high-order methods and tight tolerances.

5.5.2 SSP Methods

The popular method SSP3(2)3[3S*₊] can be equipped with the PI34 controller to give acceptable step size control performance; slightly better behavior can be achieved by choosing $\beta = (0.70, -0.37, 0.05)$. For loose and medium tolerances, this scheme performs similarly to BS3(2)3_F. BS3(2)3_F is significantly more efficient than SSP3(2)3[3S*₊] at tight tolerances.

SSP3(2)4[3S*₊] performs ca. 50% better than SSP3(2)3[3S*₊] or BS3(2)3_F at loose and medium tolerances for the inviscid Taylor-Green vortex using the optimized controller $\beta = (0.55, -0.27, 0.05)$. At loose tolerances, it is also ca. 15% more efficient than BS3(2)3_F for the isentropic vortex. However, the number of RHS evaluations increases drastically for tighter tolerances, making SSP3(2)4[3S*₊] less efficient than BS3(2)3_F for these parameters. The results for the flow with source term are similar but less pronounced. Hence, SSP3(2)4[3S*₊] can be more efficient than the best schemes so far but the embedded method does not seem to be reliable enough to make the choice of the tolerance as robust as for other schemes. Additionally, the choice of appropriate controller parameters can be crucial for SSP3(2)4[3S*₊], since some standard controllers do not perform well.

For higher polynomial degrees $p \in \{3, 4\}$, SSP3(2)4[3S*₊] is still a very interesting method that can even beat BS3(2)3_F. SSP3(2)3[3S*₊] is less efficient than SSP3(2)4[3S*₊] also for these higher polynomial degrees. For even higher polynomial degrees such as $p = 7$, the situation changes a bit since the temporal error becomes significant. While SSP3(2)4[3S*₊] is still the most efficient method so far for medium tolerances, it becomes less efficient than BS3(2)3_F for the vortex problems at a tolerance of 10^{-8} , since it is less optimized for accuracy than that general purpose method.

5.5.3 Low-Storage Methods

Some standard controllers like PI34 perform mostly acceptably well for KCL3(2)4[2R₊]C (based on the number of step rejections). Nevertheless, an optimized controller with parameter $\beta = (0.50, -0.35, 0.10)$ results in a few percent fewer function evaluations.

Fig. 5 Stability regions scaled by the effective number of stages of three representative methods (taking the FSAL property into account). The stability region of the main method is marked in gray and the boundary of the embedded method's stability region is drawn as black line

However, $\text{BS3}(2)3_F$ is up to 20% more efficient, in accordance with the real stability interval scaled by the effective number of stages, which is three for $\text{BS3}(2)3_F$ because of the FSAL property.

The PI34 controller does not perform well for the other low-storage schemes. For $\text{KCL4}(3)5[3R_+]C$ (but not for the other schemes), the PID controller with $\beta = (0.49, -0.34, 0.10)$ proposed in [36] performs much better. An optimized controller with parameter $\beta = (0.41, -0.28, 0.08)$ performs even slightly better, making this scheme more efficient than $\text{BS3}(2)3_F$ for the inviscid Taylor-Green vortex and slightly more efficient for the isentropic vortex. However, $\text{BS3}(2)3_F$ is still better for the other test case. $\text{KCL4}(3)5[2R_+]C$ and $\text{KCL5}(4)9[2R_+]S$ were more challenging for optimizing controller parameters and less efficient than $\text{KCL4}(3)5[3R_+]C$.

As for $p = 2$, $\text{KCL4}(3)5[3R_+]C$ is usually the most efficient existing low-storage method of [37] for $p \in \{3, 4\}$, which can also be more efficient than $\text{BS3}(2)3_F$ for the vortex problems. However, $\text{SSP3}(2)4[3S^*_+]$ is even more efficient there. Additionally, the sensitivity of the step size controller for the low-storage methods is bigger than for $\text{BS3}(2)3_F$. For $p = 7$, all of the low-storage methods considered here result in a non-negligible amount of step rejections for the inviscid Taylor-Green vortex. Nevertheless, the fourth-order accurate methods can be up to 15% more efficient than $\text{BS3}(2)3_F$ there. Nevertheless, $\text{SSP3}(2)4[3S^*_+]$ is still more efficient for this test problem. At tight tolerances such as 10^{-8} , $\text{KCL4}(3)5[3R_+]C$ is the most efficient method considered so far for the isentropic vortex and the smooth flow with source term.

5.5.4 Discussion

All of the general purpose schemes make use of the FSAL technique. Additionally, the stability region of the embedded scheme is always at least as big as the one of the main method. Although the 2R low-storage schemes were optimized for convection-dominated problems, they were outperformed for all test problems and at almost all tolerances by the general purpose method $\text{BS3}(2)3_F$. Possible reasons for this are that the 2R method coefficients are chosen subject to more stringent low-storage requirements, they do not exploit the FSAL technique, and they have embedded methods with a stability region that

is smaller than that of the main method in some areas. When the time step is restricted by stability, this can effectively reduce the allowable time step for the main method.

This last point is illustrated in Fig. 5, which shows the stability regions (scaled by the effective number of stages) of the main and embedded method for three pairs. We see that although the stability region of $BS3(2)3_F$ includes less of the real axis than that of $KCL4(3)5[3R_+]C$, the embedded method for $BS3(2)3_F$ extends further than that of $KCL4(3)5[3R_+]C$. The last stability region in the figure corresponds to a new method developed in the next section. Like $BS3(2)3_F$, it has the useful property that the stability region of the embedded method contains that of the main method.

In the results described above, the behavior of the methods for the inviscid Taylor-Green vortex was often slightly different from for the other test cases. This can partly be explained by the lower Mach number chosen for this example. Indeed, numerical experiments show that low Mach numbers put more stress on real axis stability than on the rest of the spectrum generated by linear advection. Hence, methods with stability regions that include more parts of the negative real axis but are not optimal for the linear advection spectrum can perform better for low Mach numbers; see also [17]. In this article, we focus on applications in CFD with medium to high Mach numbers. However, flows with small Mach numbers are usually computed using incompressible solvers and implicit time integration methods. Hence, we do not focus on this regime in this article. Nevertheless, we use the inviscid Taylor-Green vortex as test case to study the step size control stability on the negative real axis.

6 New Optimized Runge-Kutta Pairs

In the previous section we developed optimized step size controllers for existing Runge-Kutta pairs. Now we consider the optimization of Runge-Kutta pairs themselves (along with controllers). To do so, we begin with the $3S^*_+$ methods of [59], without embedded error estimator. Then, we design an embedded method and optimized controller parameters. The embedded method is optimized for step size control stability, good error metrics (see [37] and the supplementary material for the present work), to have a large stability region that includes that of the main method, and to have coefficients that are not too large. The resulting schemes given by double precision floating point numbers are optimized further using extended precision numbers in Julia [8] and the package Optim.jl [52], such that the order conditions are satisfied at least to quadruple precision. Coefficients of the new optimized methods are available in the accompanying repository [64] in full precision. Double precision coefficients are given in Appendix B. The stability region of a representative method is shown in Fig. 5c.

We also developed new pairs from scratch, based on the approach used in [59]. Specifically, we compute the Fourier footprint of the spectral element semidiscretization of the linear advection equation by varying the direction of the wave propagation velocity vector, the solution orientation, and the wave vector module and construct optimized stability polynomials using the algorithm described in [38]. Afterwards, low-storage Runge-Kutta schemes are constructed by minimizing their principal error constants, given their class, the number of stages s , the order of accuracy q , and the optimized stability polynomial as constraints. These optimizations are carried out using RK-Opt [42], based on the optimization toolbox of MATLAB. However, the resulting methods did not perform better than the pairs based on starting with methods from [59].

(a) Isentropic vortex (15), the chosen controller is within 3.1 % of the minimal #FE. **(b)** Flow with source term (16), the chosen controller is within 0.1 % of the minimal #FE.

Fig. 6 Performance of different controllers with $\beta_3 = 0$ for RK3(2)5_F[3S⁺₊] and two of the test problems with tolerance $\text{tol} = 10^{-5}$. The chosen controller for this scheme uses $\beta = (0.70, -0.23, 0.00)$ (marked with a black +). The number of function evaluations (#FE) is visualized only for those controllers that result in step size control stability along the whole boundary of the main method’s stability region

Table 6 Performance of optimized low-storage schemes: number of function evaluations (#FE), rejected steps (#R), and L^2 error of the density for the inviscid Taylor-Green vortex (14), the isentropic vortex (15), and the flow with source term (16) using polynomials of degree $p = 2$

Scheme	β	tol	TGV (14)		Isent. vortex (15)			Source term (16)		
			#FE	#R	#FE	#R	Error	#FE	#R	Error
RK3(2)5[3S ⁺ ₊]	(0.64, −0.31, 0.04)	10^{-5}	4 467	(2)	1 262	(0)	5.77E−4	19 692	(3)	1.77E−3
RK3(2)5 _F [3S ⁺ ₊]	(0.70, −0.23, 0.00)	10^{-5}	4 469	(1)	1 308	(0)	5.77E−4	19 690	(2)	1.77E−3
RK4(3)9[3S ⁺ ₊]	(0.25, −0.12, 0.00)	10^{-5}	4 205	(1)	1 478	(0)	5.79E−4	18 992	(1)	1.78E−3
RK4(3)9 _F [3S ⁺ ₊]	(0.38, −0.18, 0.01)	10^{-5}	4 207	(1)	1 371	(0)	5.80E−4	18 984	(1)	1.77E−3
RK5(4)10[3S ⁺ ₊]	(0.47, −0.20, 0.06)	10^{-5}	5 372	(2)	1 512	(0)	5.78E−4	24 107	(3)	1.78E−3
RK5(4)10 _F [3S ⁺ ₊]	(0.45, −0.13, 0.00)	10^{-5}	5 364	(1)	1 533	(0)	5.78E−4	24 112	(2)	1.77E−3

6.1 Optimization of Controller Parameters

The controller parameters for these new pairs are optimized using the same approach as described in Sect. 5.

Typical performance results of the optimization procedure are shown in Fig. 6 using RK3(2)5_F[3S⁺₊] as an example. For the isentropic vortex (15) with $\text{tol} = 10^{-5}$, the temporal accuracy starts to play a role and the controllers are not necessarily limited by stability. In this regime, controllers with larger β_1 and β_2 closer to zero perform better; they are more near to the simple deadbeat (I-) controller, which is in some sense optimal in the asymptotic regime. In contrast, the controllers operate near the stability boundary for the test case with source term (16). Here, controllers with more damping and smaller β_1 perform better. To find an acceptable controller for the scheme RK3(2)5_F[3S⁺₊], both kinds of problems have to be considered, seeking a compromise between efficiency in the asymptotic regime and near the stability boundary.

6.2 Results for New Schemes

Analogously to Tables 3, 4 and 5, results are summarized in Table 6 for the new optimized low-storage schemes with error control for $p = 2$; extended details and results for higher-order spatial semidiscretizations using solution polynomials of degree $p \in \{3, 4, 7\}$ are available in the supplementary material.

In general, the novel schemes are more efficient than all methods tested in Sect. 5. In particular, the novel third-order schemes are up to 18% more efficient than BS3(2)3_F, in accordance with the relative lengths of the real stability intervals. They are also up to 5% more efficient than KCL4(3)5[3R₊]C with the optimized PID controller for the inviscid Taylor-Green vortex and up to 13% more efficient for the isentropic vortex. Recall that BS3(2)3_F is more efficient than KCL4(3)5[3R₊]C for the other test problem. Only SSP3(2)4[3S₊^{*}] is more efficient than RK3(2)5_F[3S₊^{*}] for the inviscid Taylor-Green vortex, in accordance with the particularly large real stability interval, cf. Sect. 5.5. However, the new schemes are more efficient at realistic (medium to high) Mach numbers, for which they have been optimized.

The optimized fourth-order schemes can be even more efficient for the inviscid Taylor-Green vortex and the smooth flow with source term at medium tolerances, giving an improvement of a few percent. For the isentropic vortex, the third-order schemes are still up to 6% more efficient. The optimized fourth-order schemes are up to 25% more efficient than the best corresponding methods of [37] with the controller recommended there. However, the fourth-order accurate main methods of [59] make it particularly difficult to design good embedded methods and controllers. This can already be seen in the optimized controller coefficients, where the magnitudes of β_1 and β_2 differ less than for other optimized methods. While the controllers can be tuned to result in acceptable performance for these test problems, they do not necessarily lead to good performance for other setups.

The optimized fifth-order schemes are less efficient than the optimized third- and fourth-order schemes for these test cases unless the tolerance is very tight (so the spatial error dominates and the influence of the time integrator is negligible). These schemes (used with the controllers designed here) are much more efficient than the corresponding method of [37] (used with the controller prescribed there), by up to 25% for the source term problem, up to 35% for the isentropic vortex, and up to 18% for the inviscid Taylor-Green vortex.

For $p = 3$, the third-order accurate schemes are the most efficient ones of the optimized low-storage methods, except for the inviscid Taylor-Green vortex, where the fourth-order methods are up to 6% more efficient. For $p = 4$, the fourth-order accurate schemes are the most efficient new ones for these experiments, followed closely by the third-order accurate methods. For $p = 7$, the fourth-order accurate schemes are still the most efficient new ones; the third-order methods do not match the same small errors for tight tolerances and their temporal error dominates the spatial one. However, the fourth-order methods are difficult to control for loose tolerances, resulting in a significant number of step rejections. For sufficiently tight tolerances, the optimized fourth-order methods are more efficient than the fourth-order methods of [37].

In general, FSAL methods are often more efficient than non-FSAL schemes, especially at loose tolerances. Thus, we recommend to use the novel 3S₊^{*} (FSAL) methods for hyperbolic problems where the time step is restricted mostly by stability constraints.

Optimization of Runge-Kutta pairs for higher numbers of stages was not as successful. While we were able to obtain schemes with good theoretical properties, their performance did not show improvement compared to the schemes listed above. Improvements of the

Fig. 7 Convex hulls of the spectra of the spectral element semidiscretizations used for the optimization of stability polynomials as in [59]. The spectra are scaled such that $\min \operatorname{Re}(\lambda) = -1$

underlying optimization algorithms or the imposition of additional constraints might lead to better schemes in the future. However, the novel schemes developed in this work are already a significant improvement over the state of the art and perform well.

6.3 Further Optimizations

As shown in [59] and the numerical experiments above, the spatial error usually dominates the temporal error. Hence, it is interesting to optimize lower-order time integrators for higher-order spatial discretizations. Such an approach is presented in [58], but focused on first-order accurate time integrators. The results shown there demonstrate some speedup compared to the schemes presented in [59], but these come with a reduced accuracy.

Here, we choose third-order accurate Runge-Kutta methods and optimize them for a fifth-order spatial discretization. This resulted in a speedup compared to the optimized schemes described in the previous sections for some test cases; for other test cases, no speedup could be observed. This is in accordance with the general similarity of the scaled convex hulls of the spectra for different polynomial degrees $p = 1, 2, 3, 4$ shown in Fig. 7. Hence, we do not pursue this path of research further.

7 Additional Numerical Experiments and Comparisons

Hitherto, a careful selection of test cases was used to demonstrate issues and design criteria for explicit Runge-Kutta schemes applied to semidiscretizations of hyperbolic conservation laws. Next, more involved examples are used to demonstrate that the novel methods can be applied successfully to large-scale CFD problems including the compressible Navier-Stokes equations.

To be useful for engineering and applied problems in CFD, a CFL-based control must be automated as much as possible. Therefore, we use the approach described in Sect. 3 also for the viscous CFL number. The normalizing factor σ in (9) is chosen depending on the solution polynomial degree such that a method with a real stability interval of 2 is stable for the linear advection-diffusion equation on a uniform grid with a CFL factor $\nu = 1$. On top of that, a safety factor of 0.95 is applied, cf. [2].

Table 7 Number of function evaluations (#FE), rejected steps (#R), and L^2 error of the density for the viscous shock (19) using polynomials of degree p

Scheme	β	tol	$p = 2$			$p = 4$		
			#FE	#R	Error	#FE	#R	Error
BS3(2)3 _F	(0.60, −0.20, 0.00)	10 ^{−5}	615	(1)	3.76E−3	3 842	(3)	5.42E−5
SSP3(2)4[3S* ₊]	(0.55, −0.27, 0.05)	10 ^{−5}	450	(1)	3.76E−3	2 502	(3)	5.42E−5
KCL4(3)5[3R ₊]C	(0.41, −0.28, 0.08)	10 ^{−5}	552	(2)	3.76E−3	3 342	(29)	5.42E−5
RK3(2)5 _F [3S* ₊]	(0.70, −0.23, 0.00)	10 ^{−5}	533	(0)	3.76E−3	3 270	(2)	5.42E−5
RK4(3)9[3S* ₊]	(0.25, −0.12, 0.00)	10 ^{−5}	542	(0)	3.76E−3	3 403	(7)	5.47E−5
RK5(4)10 _F [3S* ₊]	(0.45, −0.13, 0.00)	10 ^{−5}	643	(0)	3.76E−3	3 931	(3)	5.42E−5

Except for the viscous shock described next, the other simulations conducted here start from a checkpoint of a developed solution and run on 8 nodes of Shaheen XC40 using 32 CPU cores each (one compute node of Shaheen XC40²). The general purpose and SSP methods are implemented using the explicit Runge-Kutta interface of PETSc. The other methods are implemented using their respective low-storage forms in PETSc.

7.1 Viscous Shock

The propagating viscous shock is a classical test problem for the compressible Navier-Stokes equations. The momentum $\mathcal{V}$ of the analytical solution satisfies the ODE

$$\alpha \mathcal{V} \frac{\partial \mathcal{V}}{\partial x} - (\mathcal{V} - 1)(\mathcal{V} - \mathcal{V}_f) = 0, \quad -\infty \leq x \leq +\infty. \quad (18)$$

The solution of this ODE can be written implicitly as

$$x - \frac{1}{2} \alpha \left(\log \left| (\mathcal{V}(x_1) - 1)(\mathcal{V}(x_1) - \mathcal{V}_f) \right| + \frac{1 + \mathcal{V}_f}{1 - \mathcal{V}_f} \log \left| \frac{\mathcal{V}(x_1) - 1}{\mathcal{V}(x_1) - \mathcal{V}_f} \right| \right) = 0, \quad (19)$$

where

$$\mathcal{V}_f \equiv \frac{\mathcal{U}_L}{\mathcal{U}_R}, \quad \alpha \equiv \frac{2\gamma}{\gamma + 1} \frac{\mu}{Pr \mathcal{M}}. \quad (20)$$

Here, $\mathcal{U}_{L/R}$ are the known velocities to the left and right of the shock at $\pm\infty$, $\mathcal{M}$ is the constant mass flow across the shock, $Pr = 3/4$ is the Prandtl number, and μ is the dynamic viscosity. The mass and total enthalpies are constant across the shock. Moreover, the momentum and energy equations become redundant.

² Shaheen XC40 is the petascale supercomputer hosted at KAUST, which features 6 174 dual socket compute nodes based on the 16 core Intel Haswell processors running at 2.3 GHz. Each node is equipped with 128 GB of DDR4 memory running at 2.3 GHz. Overall, the system has a total of 197 568 processor cores and 790 TB of aggregate memory.

Fig. 8 Solution polynomial degree distribution, computational domain and boundary mesh elements for the NASA junction experiment [55]; f is the fuselage length

For our tests, we compute $\mathcal{V}$ from (19) to machine precision using bisection. The moving shock solution is obtained by applying a uniform translation to the above solution. Initially, at $t = 0$, the shock is located at the center of the domain. We use the parameters $Ma = 2.5$, $Re = 10$, and $\gamma = 1.4$ in the domain given by $x \in [-0.5, 0.5]$ till the final time $t = 2$. The boundary conditions are prescribed by penalizing the numerical solution against the analytical solution, which is also used to prescribe the initial condition.

Some results for the most promising methods and optimized controllers are shown in Table 7; extended details are available in the Supplementary Material. The new scheme $RK3(2)5_F[3S^*_+]$ is ca. 18% more efficient than $BS3(2)3_F$ for relevant tolerances, in accordance with the relative real stability intervals. $SSP3(2)4[3S^*_+]$ is a very promising scheme for this kind of problem because of its improved stability properties around the negative real axis. In particular, $SSP3(2)4[3S^*_+]$ is ca. 50% more efficient than $BS3(2)3_F$ for relevant tolerances, also in accordance with the relative real stability intervals. Except for very tight tolerances and low solution polynomial degrees, the step size controllers detect the stability constraint accurately; the spatial error dominates and the error of the time integration schemes is negligible.

7.2 NASA Junction Flow

We consider the NASA junction flow problem as described in [55, Section 3.8]. The NASA junction flow test was designed to validate CFD for wing junction trailing edge separation and progression, and it is a collaborative effort between CFD computationalists and experimentalists [16]. Specifically, the NASA junction flow experiment is a series of wind tunnel tests conducted in the NASA Langley subsonic tunnel to collect validation data in the junction region of a wing-body configuration [68].

Here, we simulate the NASA junction flow with a wing based on the DLR-F6 geometry and a leading edge horn to mitigate the effect of the horseshoe vortex over the wing-fuselage junction [50]. A general view of the geometry is shown in Fig. 8b. The model crank chord is $\ell = 557.1$ mm, the wing span is 77.89ℓ , and the fuselage length is $f = 8.69\ell$. The wing leading edge horn meets the fuselage at $x_1 = 3.45\ell$, and the wing root trailing-edge is

Fig. 9 Q -criterion colored by the velocity magnitude of the NASA juncture flow

Table 8 Performance of different methods for error and CFL-based step size controllers: number of function evaluations (#FE), rejected steps (#R), and wall-clock time in seconds for the NASA juncture flow problem

Scheme	β	tol/ν	#FE	#R	Wall-clock time/s
BS3(2)3 _F	(0.60, -0.20, 0.00)	$\text{tol} = 10^{-8}$	1 786	1	586.49
		$\nu = 1.0$	3 598		1 171.60
SSP3(2)4[3S* ₊]	(0.55, -0.27, 0.05)	$\text{tol} = 10^{-8}$	1 172	1	376.18
		$\nu = 1.0$	2 340		747.45
RK3(2)5 _F [3S* ₊]	(0.70, -0.23, 0.00)	$\text{tol} = 10^{-8}$	1 527	1	496.58
		$\nu = 1.0$	3 056		991.35
RK4(3)9[3S* ₊]	(0.25, -0.12, 0.00)	$\text{tol} = 10^{-8}$	1 467	1	469.20
		$\nu = 1.0$	2 862		912.17
RK5(4)10 _F [3S* ₊]	(0.45, -0.13, 0.00)	$\text{tol} = 10^{-8}$	1 864	3	597.06
		$\nu = 1.0$	3 661		1 170.10

located at $x_1 = 5.31\ell$. In the wind tunnel, the model is mounted on a sting aligned with the fuselage axis. The sting is attached to a mast that emerges from the wind tunnel floor. The Reynolds number is $Re = 2.4 \times 10^6$ and the freestream Mach number is $Ma = 0.189$. The angle of attack is $\text{AoA} = -2.5^\circ$. We perform simulations in free air conditions, ignoring both the sting and the mast.

As shown in Fig. 8b, the grid is subdivided into three blocks, corresponding to three different approximation degrees, p , for the solution field. In particular, we use $p = 1$ in the far-field region, $p = 3$ in the region surrounding the model, and $p = 2$ elsewhere. In total, we use $\approx 6.762 \times 10^5$ hexahedral elements and $\approx 4.091 \times 10^7$ degrees of freedom (DOFs). We highlight that the boundary layer thickness over the fuselage for $x = 1\,000\text{--}2\,000$ mm is about 16 mm, while it is about 20 mm over the wing upstream of the separation bubble [33]. In the present simulation we use between eight and nine solution points in the boundary layer thickness δ_{99} . The mesh features a maximum aspect ratio of ca. 110. The grid is constructed using the commercial software Pointwise V18.3 released in September 2019; solid boundaries are described using a quadratic mesh.

Figure 9 shows the Q -criterion colored by the velocity magnitude of flow past the aircraft. The separation of the flow on the wing near the junction with the fuselage is visible.

Fig. 10 Overview of the Imperial Front Wing

A summary of the performance of the different methods is presented in Table 8. Here, the CFL adaptor with $\nu = 1.0$ tuned for linear advection-diffusion works for all Runge-Kutta methods. However, it was significantly less efficient than the error-based controller with a conservative tolerance of 10^{-8} ; the CFL controller used ca. 50% more RHS evaluations and wall-clock time. Thus, a tedious manual tuning to increase the CFL factor would be necessary to match the efficiency of the error-based controller which just works out of the box.

BS3(2)3_F is an efficient general purpose method for this CFD problem. Nevertheless, the optimized third- and fourth-order accurate methods are more efficient. Interestingly, SSP3(2)4[3S*₊] is again significantly more efficient, nearly 50% faster than BS3(2)3_F.

7.3 Viscous Flow Past a Formula 1 Front Wing

Here, we consider the flow past a Formula 1 front wing with a relatively complex geometry, supported by the availability of a CAD model and experimental results [60]. We refer to this test case as the Imperial Front Wing, originally based on the front wing and endplate design of the McLaren 17D race car [11]. The panel of Fig. 10 gives an overview of the Imperial Front Wing geometry. We denote by h the distance between the ground and the lowest part of the front wing endplate and by c the chord length of the main element. The position of the wing in the tunnel is further characterized by a pitch angle of 1.094° . Here we use $h/c = 0.36$ which can be considered as a relatively low front ride height, with high ground effect and hence higher loads on the wing. The corresponding Reynolds number is $Re = 2.2 \times 10^5$, based on the main element chord c of 250 mm and a free stream velocity U of 25 m/s. The Mach number is set to $Ma = 0.036$. This corresponds to a practically incompressible flow.

The computational domain is divided into 3.4×10^6 hexahedral elements with a maximum aspect ratio of ca. 250. Two different semidiscretizations with solution polynomials of degree $p = 1$ and $p = 2$ are used. The grid is constructed using the commercial software

Fig. 11 Time-averaged pressure coefficient, C_p , on the surface of the Imperial Front Wing

Table 9 Performance of different methods for error and CFL-based step size controllers: number of function evaluations (#FE), rejected steps (#R), and wall-clock time in seconds for the Imperial Front Wing using polynomials of degree $p = 2$

Scheme	β	tol/ν	#FE	#R	Wall-clock time/s
BS3(2)3 _F	(0.60, -0.20, 0.00)	$\text{tol} = 10^{-8}$	3 973	0	2 699.9
		$\nu = 1.0$	7 678		5 188.4
SSP3(2)4[3S* ₊]	(0.55, -0.27, 0.05)	$\text{tol} = 10^{-8}$	2 980	0	1 971.5
		$\nu = 1.0$	4 992		3 298.9
RK3(2)5 _F [3S* ₊]	(0.70, -0.23, 0.00)	$\text{tol} = 10^{-8}$	3 400	4	2 277.8
		$\nu = 1.0$	6 516		4 362.3
RK4(3)9[3S* ₊]	(0.25, -0.12, 0.00)	$\text{tol} = 10^{-8}$	3 195	1	2 115.2
		$\nu = 1.0$	6 111		4 046.1
RK5(4)10 _F [3S* ₊]	(0.45, -0.13, 0.00)	$\text{tol} = 10^{-8}$	4 094	3	2 714.1
		$\nu = 1.0$	7 811		5 186.8

Table 10 Performance of different methods for error and CFL-based step size controllers: number of function evaluations (#FE), rejected steps (#R), and wall-clock time in seconds for the Imperial Front Wing using polynomials of degree $p = 1$. SSP3(2)4[3S*₊] crashed for $\nu = 1.0$

Scheme	β	tol/ν	#FE	#R	Wall-clock time/s
BS3(2)3 _F	(0.60, -0.20, 0.00)	$\text{tol} = 10^{-8}$	1 411	0	356.75
		$\nu = 1.0$	1 579		398.55
SSP3(2)4[3S* ₊]	(0.55, -0.27, 0.05)	$\text{tol} = 10^{-8}$	1 352	0	331.97
		$\nu = 0.85$	1 304		321.54
RK3(2)5 _F [3S* ₊]	(0.70, -0.23, 0.00)	$\text{tol} = 10^{-8}$	1 127	6	280.19
		$\nu = 1.0$	1 341		333.21
RK4(3)9[3S* ₊]	(0.25, -0.12, 0.00)	$\text{tol} = 10^{-8}$	1 233	1	302.31
		$\nu = 1.0$	1 260		308.34
RK5(4)10 _F [3S* ₊]	(0.45, -0.13, 0.00)	$\text{tol} = 10^{-8}$	1 492	1	366.16
		$\nu = 1.0$	1 611		395.35

Pointwise V18.3 released in September 2019; solid boundaries are described using a quadratic mesh.

In Fig. 11, we present the contour plot of the time-averaged pressure coefficient on the surface of the front wing. The statistics have been obtained by averaging the solution for approximately five flow-through time units.

The performance characteristics of the different methods for $p = 2$ are summarized in Table 9. The results are in agreement with those obtained for the NASA juncture flow. The CFL adaptor with $\nu = 1.0$ works for all methods and is less efficient than error-based step size controllers. Again, BS3(2)3_F is an efficient general purpose scheme for this problem. Nevertheless, the optimized third- and fourth-order methods are more efficient and SSP3(2)4[3S*₊] is the most efficient scheme for this problem. The fifth-order method is less efficient than the other four schemes, as expected.

The results for $p = 1$ summarized in Table 10 are mostly similar to the ones presented before. In contrast to the results for $p = 2$, the CFL adaptor with $\nu = 1.0$ did not work for SSP3(2)4[3S*₊]; the simulation crashed when using $\nu = 1.0$ and manual tuning was necessary to get a working setup³. For $\nu = 0.85$, the CFL adaptor worked and was a few percent more efficient than the error-based controller. However, the latter did not require any manual tuning at all and worked robustly with default parameters for all Runge-Kutta methods. Here, SSP3(2)4[3S*₊] is less efficient than RK3(2)5_F[3S*₊] and RK4(3)9[3S*₊]. Otherwise, the results are similar to the ones obtained for the juncture flow and the setup using $p = 2$.

8 Summary and Conclusions

We studied explicit Runge-Kutta methods applied to dissipative spectral element semidiscretizations of hyperbolic conservation laws and CFD problems based on the compressible Euler and Navier-Stokes equations. In this context, we argued in Sect. 3 that error-based step size control can be advantageous compared to CFL-based approaches, since associated user-defined parameters are usually more robust and can be varied in rather large ranges without affecting accuracy or efficiency. Additionally, error-based step size control moves the burden of constructing critical parts of the controller from the developer of the spatial semidiscretization to the developer of the time integrator, easing the workflow for most researchers. The results for more complex test problems in Sect. 7 also support this conclusion.

In Sect. 4, we demonstrated that choosing good step size controller parameters is especially important if the time step is restricted by stability constraints, as is typical for many convection-dominated problems. We compared existing Runge-Kutta pairs in Sect. 5 and proposed an approach to optimize controller parameters for such methods. In general, the third-order method BS3(2)3_F of Bogacki and Shampine [9] performs well compared to both general purpose schemes and methods designed specifically for CFD applications. The strong-stability preserving method SSP3(2)4[3S*₊] of [47] with embedded method of [18] also performs well, but is more sensitive to the choice of error tolerance.

³ This behavior can be explained by the different shape of the stability region of SSP3(2)4[3S*₊] compared to the other methods, with a relatively larger real stability interval, cf. Table 1 in the supplementary material. Another measure of the size of the stability region might help but would not remove the necessity of manual tuning to get good performance for CFL-based controllers.

Table 11 Optimized PID controller parameters for some explicit Runge-Kutta methods with embedded error estimator

Method	Reference	β_1	β_2	β_3
BS3(2)3 _F	[9]	0.60	−0.20	0.00
BS5(4)7 _F	[10]	0.28	−0.23	0.00
DP5(4)6 _F	[61]	0.70	−0.40	0.00
T5(4)6 _F	[77]	0.57	−0.24	0.04
SSP3(2)3[3S* ₊]	[18, 70]	0.70	−0.37	0.05
SSP3(2)4[3S* ₊]	[18, 47]	0.55	−0.27	0.05
KCL3(2)4[2R ₊]C	[37]	0.50	−0.35	0.10
KCL4(3)5[3R ₊]C	[37]	0.41	−0.28	0.08
RK3(2)5[3S* ₊]	This article	0.64	−0.31	0.04
RK3(2)5 _F [3S* ₊]	This article	0.70	−0.23	0.00
RK4(3)9[3S* ₊]	This article	0.25	−0.12	0.00
RK4(3)9 _F [3S* ₊]	This article	0.38	−0.18	0.01
RK5(4)10[3S* ₊]	This article	0.47	−0.20	0.06
RK5(4)10 _F [3S* ₊]	This article	0.45	−0.13	0.00

In Sect. 6, we developed explicit low-storage Runge-Kutta pairs with optimized step size controllers. These novel schemes are more efficient than all of the existing schemes when applied to advection-dominated problems. We demonstrated their performance in several CFD applications with increasing complexity, including the compressible Euler and Navier-Stokes equations. We contributed our optimized methods to the freely available open source library `DifferentialEquations.jl` [62] written in Julia [8].

Although not demonstrated in this article, another advantage of error-based step size control becomes apparent for a cold startup of CFD problems, i.e., simulations around complex geometries that are initialized with a free stream flow. CFL-based approaches often need to adjust the CFL scaling at the beginning to cope with the initial transient period. In contrast, our error-based approach does not need special tuning and is robust in our experience.

A summary of existing and novel methods with optimized controller parameters is given in Table 11. Depending on whether dissipative/low-Mach effects dominate, SSP3(2)4[3S*₊] and RK3(2)5_F[3S*₊] are the most efficient schemes in our experience. Additionally, BS3(2)3_F is a surprisingly efficient general purpose method. It becomes increasingly complicated to design controllers that are stable and efficient across different applications for methods of higher order and/or with more stages. However, this is not necessarily a severe drawback, since third-order accurate methods like RK3(2)5_F[3S*₊], SSP3(2)4[3S*₊], and BS3(2)3_F are usually more efficient in CFD applications. As argued in Sect. 3, the error-based control has usually a relatively mild sensitivity with respect to the choice of the tolerance. In our experience, it is usually good to choose a relatively tight tolerance around 10^{-8} for applied CFD problems. Since the step size is almost always limited by stability, the tolerance does not matter that much, but a relatively tight tolerance helps for methods that are more difficult to control (e.g., those of higher order or with more stages).

The present work was influenced and partially motivated by the landmark work of Kennedy et al. [37], which focused on developing optimized Runge-Kutta methods for CFD. Herein, we focus on modern spectral element semidiscretizations that introduce dissipation at element interfaces, e.g., using upwind numerical fluxes. Hence, the stability regions of

our methods focus also on the negative real axis, whereas “imaginary axis stability is a high priority to the methods” designed in [37, p. 183]. Furthermore, we concentrate on the common case where the spatial error dominates and the step size is restricted by stability rather than temporal accuracy. Hence, we design Runge-Kutta pairs with large stability regions, for both the main and the embedded methods. In particular, the stability regions of our novel embedded schemes are larger than the ones of the corresponding main methods.

To our knowledge, this article is the first exploring the impact of controller parameters and step size control stability on the efficiency of explicit Runge-Kutta methods for CFD systematically. This provides important insights into the construction of new methods and augments best practices published before. In particular, we think the conventional wisdom that “coping with step size control instability is probably best accomplished by reducing step sizes” [37, p. 208] can be improved upon by instead optimizing the controller, since that results in a more robust and efficient scheme. As noted in [37, p. 208], “doing this optimization requires some caution because it is not sufficient in the design of a good controller for each of the eigensolutions to be damped. The time constants associated with these eigensolutions must not be too large or too small”. Herein, we proposed a way to conduct this optimization systematically and applied it to a wide range of schemes.

Of course, such an approach also comes with limitations, in particular if the main method is fixed such that only the embedded method and the controller can be designed freely. Some methods such as the fourth-order method used in this article constrain the range of embedded methods and controllers such that a good general purpose optimization is not necessarily successful. While the combined method can be efficient for certain problems, it is not necessarily similarly efficient for other problems, e.g., when going from inviscid to viscous flows. Other schemes such as the novel third- and fifth-order accurate optimized methods result in less stability restrictions, making the resulting methods and controllers efficient for a broad range of CFD problems. Thus, we would like to stress that designing a good time integration method should not only focus on the main method but consider the interaction of a main method, an error estimator, and a step size controller. Applying this principle to viscous flows will be a subject of future work.

Some previous work has focused on automated step size control for convection-dominated problems with the goal of achieving a temporal error that is of similar magnitude to that of the spatial error [7, 80]. The addition of such a control on top of the techniques employed here might lead to an even more efficient controller that is not adversely affected by excessively tight temporal error tolerance specification.

We expect the new methods developed in this article to perform well for convection-dominated flows in the subsonic regime; here, we tested them mainly with reference Mach numbers in the range 0.1–0.5. We have demonstrated their improved performance compared to some standard schemes also in other regimes, including viscous flows and the transonic/supersonic regime. For small Mach numbers, incompressible solvers with implicit time discretizations are usually applied. However, if compressible solvers should be used for low Mach numbers, methods could be optimized following the approach of this article.

Appendix A Efficient implementation of SSP3(2)4[3S*₊]

The Butcher coefficients of SSP3(2)4[3S*₊] are

$$\begin{array}{c|cccc}
 0 & & & & \\
 \hline
 1/2 & 1/2 & & & \\
 1 & 1/2 & 1/2 & & \\
 1/2 & 1/6 & 1/6 & 1/6 & \\
 \hline
 & 1/6 & 1/6 & 1/6 & 1/2 \\
 & 1/4 & 1/4 & 1/4 & 1/4
 \end{array} \quad (A1)$$

where spaces indicate zeros. Because of its low-storage structure, the method can be implemented efficiently and memory-friendly as

$$u \leftarrow u^n + \frac{1}{2} \Delta t_n f(t_n, u^n), \quad u \leftarrow u + \frac{1}{2} \Delta t_n f(t_n + \Delta t_n/2, u), \quad u \leftarrow u + \frac{1}{2} \Delta t_n f(t_n + \Delta t_n, u), \quad (A2a)$$

$$\hat{u} \leftarrow \frac{1}{3} u^n + \frac{2}{3} u, \quad u \leftarrow \frac{2}{3} u^n + \frac{1}{3} u, \quad u \leftarrow u + \frac{1}{2} \Delta t_n f(t_n + \Delta t_n/2, u), \quad \hat{u} \leftarrow \frac{1}{2} (\hat{u} + u). \quad (A2b)$$

At the end of one step (A2), u^{n+1} is stored in u and $\hat{u}^{n+1}$ is stored in $\hat{u}$. Usually, it is not important to know $\hat{u}^{n+1}$ but $\hat{u}^{n+1} - u^{n+1}$ to estimate the error; this difference can be obtained as $(\hat{u} - u)/2$ instead of $(\hat{u} + u)/2$ in the last assignment in (A2). If the low-storage assumption introduced in [40] can be applied, SSP3(2)4[3S*₊] can be implemented using only three memory locations for u^n , u and $\hat{u}$. Otherwise, an additional storage location is necessary to evaluate the right-hand side f . Note that the previous value u^n is already included in this count of memory locations.

Appendix B Coefficients of the Novel Runge-Kutta Pairs

The low-storage coefficients of the novel methods are listed in double precision in Tables B1, B2, B3, B4, B5 and B6. Full-precision results in electronic form are available in the accompanying repository [64]. We contributed our optimized methods to the freely available open source library DifferentialEquations.jl [62] written in Julia [8].

Table B1 Coefficients of the optimized, explicit, low-storage Runge-Kutta method RK3(2)5[3S*₊]

$\gamma_{1,i}$	$\gamma_{2,i}$	$\gamma_{3,i}$
+0.0000000000000000E + 00	+1.0000000000000000E + 00	+0.0000000000000000E + 00
+2.5876690703520788E − 01	+5.5284187451021605E − 01	+0.0000000000000000E + 00
−1.3243668739945030E − 01	+6.7318444003896738E − 01	+0.0000000000000000E + 00
+5.0556012314603993E − 02	+2.8031038045076351E − 01	+2.7525858134466369E − 01
+5.6705528079028777E − 01	+5.5215088735073936E − 01	−8.9505487092797853E − 01
δ_i	β_i	$\hat{b}_i$
+1.0000000000000000E + 00	+1.1479315633699007E − 01	+1.0463633713540937E − 01
+3.4076872093214550E − 01	+8.9335592952328596E − 02	+9.5204315749567586E − 02
+3.4143992805846252E − 01	+4.3558587173792318E − 01	+4.4824466455686685E − 01
+7.2293027328755899E − 01	+2.4735852952572862E − 01	+2.4490302954613102E − 01
+0.0000000000000000E + 00	+1.1292684944702953E − 01	+1.0701165301202518E − 01

Table B2 Coefficients of the optimized, explicit, low-storage Runge-Kutta method RK3(2)5_F[3S*₊]

$\gamma_{1,i}$	$\gamma_{2,i}$	$\gamma_{3,i}$
+0.0000000000000000E + 00	+1.0000000000000000E + 00	+0.0000000000000000E + 00
+2.5877719797257331E − 01	+5.5283549093013895E − 01	+0.0000000000000000E + 00
−1.3243803601407234E − 01	+6.7318716082030616E − 01	+0.0000000000000000E + 00
+5.0560339481908259E − 02	+2.8031039632976723E − 01	+2.7525632733046762E − 01
+5.6705320007393134E − 01	+5.5215254470206099E − 01	−8.9505261746740339E − 01
δ_i	β_i	$\hat{b}_i$
+1.0000000000000000E + 00	+1.1479359710235412E − 01	+9.4841667050357029E − 02
+3.4076558793345252E − 01	+8.9334428531133159E − 02	+1.7263713394303537E − 01
+3.4143826550033862E − 01	+4.3558710250086169E − 01	+3.9982431890843712E − 01
+7.2292753667879872E − 01	+2.4735761882014512E − 01	+1.7180168075801786E − 01
+0.0000000000000000E + 00	+1.1292725304550591E − 01	+5.8819144221557401E − 02

Table B3 Coefficients of the optimized, explicit, low-storage Runge-Kutta method RK4(3)9[3S*₊]

$\gamma_{1,i}$	$\gamma_{2,i}$	$\gamma_{3,i}$
+0.0000000000000000E + 00	+1.0000000000000000E + 00	+0.0000000000000000E + 00
-4.6556413012591804E + 00	+2.4992627526078262E + 00	+0.0000000000000000E + 00
-7.7202649248360644E - 01	+5.8668203654361373E - 01	+0.0000000000000000E + 00
-4.0244232134197242E + 00	+1.2051413654126708E + 00	+7.6210371111381703E - 01
-2.1296852467390187E - 02	+3.4747937967008691E - 01	-1.9811821590872183E - 01
-2.4350225192344701E + 00	+1.3213461401287232E + 00	-6.2289607063175667E - 01
+1.9856274809861678E - 02	+3.1196363243793707E - 01	-3.7522469934326264E - 01
-2.8107901128852841E - 01	+4.3514190558940874E - 01	-3.3554365390009466E - 01
+1.6894348958355357E - 01	+2.3596982994407883E - 01	-4.5609631107174843E - 02
δ_i	β_i	$\hat{b}_i$
+1.0000000000000000E + 00	+4.5037319691658841E - 02	+4.5506559279709452E - 02
+1.2629238543878065E + 00	+1.8592173220119687E - 01	+1.1759683104926386E - 01
+7.5749671775608729E - 01	+3.3297275092076306E - 02	+3.6582573305152133E - 02
+5.1635911581112226E - 01	-4.7842226210501985E - 03	-5.3115558343556296E - 03
-2.7463337920428273E - 02	+4.0558480626375678E - 03	+5.1782500127131271E - 03
-4.3826746539417710E - 01	+4.1850279996827944E - 01	+4.9546390221186826E - 01
+1.2735871036683928E + 00	-4.3818945074742778E - 03	-5.9993031327378659E - 03
-6.2947400454427949E - 01	+2.7128460973244426E - 02	+9.4050934345683165E - 02
+0.0000000000000000E + 00	+2.9522268113943101E - 01	+2.1693180876270352E - 01

Table B4 Coefficients of the optimized, explicit, low-storage Runge-Kutta method RK4(3)9_F[3S*₊]

$\gamma_{1,i}$	$\gamma_{2,i}$	$\gamma_{3,i}$
+0.0000000000000000E + 00	+1.0000000000000000E + 00	+0.0000000000000000E + 00
-4.6556414473350687E + 00	+2.4992627925744948E + 00	+0.0000000000000000E + 00
-7.7202650996458722E - 01	+5.8668203777188754E - 01	+0.0000000000000000E + 00
-4.0244366905198063E + 00	+1.2051460865230945E + 00	+7.6210066787213149E - 01
-2.1296762840185311E - 02	+3.4747937221867325E - 01	-1.9811825043394005E - 01
-2.4350225097901097E + 00	+1.3213460609651131E + 00	-6.2289592186990073E - 01
+1.9856272971319869E - 02	+3.1196364646941938E - 01	-3.7522483807759566E - 01
-2.8107911467910385E - 01	+4.3514195396843791E - 01	-3.3554383091351697E - 01
+1.6894341687548597E - 01	+2.3596981300287537E - 01	-4.5609550050311212E - 02
δ_i	β_i	$\hat{b}_i$
+1.0000000000000000E + 00	+4.5037326272637540E - 02	+2.4836759124515911E - 02
+1.2629238766481143E + 00	+1.8592173036998480E - 01	+1.8663277745621037E - 01
+7.5749671896859117E - 01	+3.3297296725697173E - 02	+5.6710807959369842E - 02
+5.1635894531407278E - 01	-4.7842041809589755E - 03	-3.4476954391492879E - 03
-2.7463274218026097E - 02	+4.0558359610313108E - 03	+3.6022450565166364E - 03
-4.3826731781279443E - 01	+4.1850277725960744E - 01	+4.5455706221450887E - 01
+1.2735872946026565E + 00	-4.3819019689193264E - 03	-2.4346652894276124E - 04
-6.2947402839274003E - 01	+2.7128437964460898E - 02	+6.6427553611035500E - 02
+0.0000000000000000E + 00	+2.9522270159645919E - 01	+1.6136970795235051E - 01

Table B5 Coefficients of the optimized, explicit, low-storage Runge-Kutta method RK5(4)10[3S*₊]

$\gamma_{1,i}$	$\gamma_{2,i}$	$\gamma_{3,i}$
+0.0000000000000000E + 00	+1.0000000000000000E + 00	+0.0000000000000000E + 00
+4.0436600785046961E - 01	+6.8714670697523461E - 01	+0.0000000000000000E + 00
-8.5034274642631846E - 01	+1.0930247604688987E + 00	+0.0000000000000000E + 00
-6.9508941670724198E + 00	+3.2259753823301613E + 00	-2.3934051593421395E + 00
+9.2387652253282782E - 01	+1.0411537008413965E + 00	-1.9028544220959867E + 00
-2.5631780399574042E + 00	+1.2928214888647027E + 00	-2.8200422105832073E + 00
+2.5457448699663476E - 01	+7.3914627692970059E - 01	-1.8326984641305650E + 00
+3.1258317338631691E - 01	+1.2391292570393000E - 01	-2.1990945107506979E - 01
-7.0071148005675854E - 01	+1.8427534793667669E - 01	-4.0824306603848765E - 01
+4.8396209709807264E - 01	+5.7127889426970779E - 02	-1.3776697911212080E - 01
δ_i	β_i	$\hat{b}_i$
+1.0000000000000000E + 00	-2.2801023055963646E - 03	+5.7345884846761938E - 02
-1.3317784091338497E - 01	+1.4073930208232305E - 02	+1.9714475180397338E - 02
+8.2604227852460299E - 01	+2.3326917941728226E - 01	+7.2152966056837173E - 02
+1.5137004305133324E + 00	+4.8082667004651816E - 02	+1.7396594898079398E - 01
-1.3058100631770482E + 00	+4.1190032211396227E - 01	+3.7036936004454879E - 01
+3.0366787893425076E + 00	-1.2914610713647529E - 01	-1.2155990390550650E - 01
-1.4494582670745926E + 00	+1.2207460110385798E - 01	+1.1803729454911216E - 01
+3.8343138733209576E + 00	+4.3578588031133875E - 02	+4.1556888233648698E - 02
+4.1222939719233249E + 00	+1.0250768752899050E - 01	+1.2278866279103799E - 01
+0.0000000000000000E + 00	+1.5593923403396062E - 01	+1.4562842322236844E - 01

Table B6 Coefficients of the optimized, explicit, low-storage Runge-Kutta method RK5(4)10_F[3S*₊]

$\gamma_{1,i}$	$\gamma_{2,i}$	$\gamma_{3,i}$
+0.0000000000000000E + 00	+1.0000000000000000E + 00	+0.0000000000000000E + 00
+4.0436601216857498E − 01	+6.8714670281614165E − 01	+0.0000000000000000E + 00
−8.5034272895758400E − 01	+1.0930247489147509E + 00	+0.0000000000000000E + 00
−6.9508941752621176E + 00	+3.2259753796071928E + 00	−2.3934051332441948E + 00
+9.2387651927310854E − 01	+1.0411537025101014E + 00	−1.9028544224217609E + 00
−2.5631780565098912E + 00	+1.2928214879121649E + 00	−2.8200422073999771E + 00
+2.5457448793652260E − 01	+7.3914627557881230E − 01	−1.8326984652773810E + 00
+3.1258317074119985E − 01	+1.2391292513718004E − 01	−2.1990944830846712E − 01
−7.0071144144405084E − 01	+1.8427534723701233E − 01	−4.0824306358478707E − 01
+4.8396210160238334E − 01	+5.7127889987965835E − 02	−1.3776697978802896E − 01
δ_i	β_i	$\hat{b}_i$
+1.0000000000000000E + 00	−2.2801003218369809E − 03	−2.0192554400120660E − 02
−1.3317784195088034E − 01	+1.4073931157901863E − 02	+2.7379034809591845E − 02
+8.2604228147502079E − 01	+2.3326917755084567E − 01	+3.0288186361459657E − 01
+1.5137004257557283E + 00	+4.8082667413538623E − 02	−3.6568438806222223E − 02
−1.3058100599350237E + 00	+4.1190032177069519E − 01	+3.9826647746767679E − 01
+3.0366788029241634E + 00	−1.2914610678077362E − 01	−5.7159594211406851E − 02
−1.4494582743988951E + 00	+1.2207460138487101E − 01	+9.8498551038485579E − 02
+3.8343138991763621E + 00	+4.3578585831744204E − 02	+6.6546015524560853E − 02
+4.1222937600129850E + 00	+1.0250768775680807E − 01	+9.0734795427481127E − 02
+0.0000000000000000E + 00	+1.5593923423620598E − 01	+8.4322893253308037E − 02

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s42967-021-00159-w>.

Acknowledgements Research reported in this publication was supported by the King Abdullah University of Science and Technology (KAUST). We are thankful for the computing resources of the Supercomputing Laboratory and the Extreme Computing Research Center at KAUST. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC 2044-390685587, Mathematics Münster: Dynamics-Geometry-Structure. Special thanks are extended to the McLaren F1 racing Team for providing data, CAD geometries, and setup of the Imperial Front Wing test case.

Funding Open Access funding enabled and organized by Projekt DEAL.

Compliance with Ethical Standards

Conflict of Interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not

permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Abhyankar, S., Brown, J., Constantinescu, E.M., Ghosh, D., Smith, B.F., Zhang, H.: PETSc/TS: a modern scalable ODE/DAE solver library (2018). [arXiv:1806.01437](https://arxiv.org/abs/1806.01437) [math.NA]
2. Al Jahdali, R., Boukharfane, R., Dalcin, L., Parsani, M.: Optimized explicit Runge-Kutta schemes for entropy stable discontinuous collocated methods applied to the Euler and Navier-Stokes equations. In: AIAA Scitech 2021 Forum, p. 0633 (2021). <https://doi.org/10.2514/6.2021-0633>
3. Almquist, M., Dunham, E.M.: Elastic wave propagation in anisotropic solids using energy-stable finite differences with weakly enforced boundary and interface conditions (2020). [arXiv:2003.12811](https://arxiv.org/abs/2003.12811) [math.NA]
4. Arévalo, C., Söderlind, G., Hadjimichael, Y., Fekete, I.: Local error estimation and step size control in adaptive linear multistep methods. *Numer. Algorithm* **86**, 537–563 (2021). <https://doi.org/10.1007/s11075-020-00900-1>
5. Baggag, A., Atkins, H., Keyes, D.: Parallel implementation of the discontinuous Galerkin method. Tech. Rep. NASA/CR-1999-209546, NASA, Institute for Computer Applications in Science and Engineering, NASA Langley Research Center, Hampton VA United States (1999)
6. Balay, S., Abhyankar, S., Adams, M.F., Brown, J., Brune, P., Buschelman, K., Dalcin, L., Dener, A., Eijkhout, V., Gropp, W.D., Kaushik, D., Knepley, M.G., May, D.A., McInnes, L.C., Mills, R.T., Munson, T., Rupp, K., Sanan, P., Smith, B.F., Zampini, S., Zhang, H., Zhang, H.: PETSc users manual. Tech. Rep. ANL-95/11—Revision 3.13, Argonne National Laboratory (2020)
7. Berzins, M.: Temporal error control for convection-dominated equations in two space dimensions. *SIAM J. Sci. Comput.* **16**(3), 558–580 (1995)
8. Bezanson, J., Edelman, A., Karpinski, S., Shah, V.B.: Julia: a fresh approach to numerical computing. *SIAM Rev.* **59**(1), 65–98 (2017). <https://doi.org/10.1137/141000671>. [arXiv:1411.1607](https://arxiv.org/abs/1411.1607) [cs.MS]
9. Bogacki, P., Shampine, L.F.: A 3(2) pair of Runge-Kutta formulas. *Appl. Math. Lett.* **2**(4), 321–325 (1989). [https://doi.org/10.1016/0893-9659\(89\)90079-7](https://doi.org/10.1016/0893-9659(89)90079-7)
10. Bogacki, P., Shampine, L.F.: An efficient Runge-Kutta (4,5) pair. *Comput. Math. Appl.* **32**(6), 15–28 (1996). [https://doi.org/10.1016/0898-1221\(96\)00141-1](https://doi.org/10.1016/0898-1221(96)00141-1)
11. Buscariolo, F.F., Hoessler, J., Moxey, D., Jassim, A., Gouder, K., Basley, J., Murai, Y., Assi, G.R.S., Sherwin, S.J.: Spectral/ hp element simulation of flow past a Formula One front wing: validation against experiments (2019). [http://arxiv.org/abs/1909.06701v1](https://arxiv.org/abs/1909.06701v1)
12. Butcher, J.C.: Numerical Methods for Ordinary Differential Equations. Wiley, Chichester (2016). <https://doi.org/10.1002/9781119121534>
13. Carpenter, M.H., Fisher, T.C., Nielsen, E.J., Frankel, S.H.: Entropy stable spectral collocation schemes for the Navier-Stokes equations: discontinuous interfaces. *SIAM J. Sci. Comput.* **36**(5), B835–B867 (2014). <https://doi.org/10.1137/130932193>
14. Carpenter, M.H., Parsani, M., Fisher, T.C., Nielsen, E.J.: Towards an entropy stable spectral element framework for computational fluid dynamics. In: 54th AIAA Aerospace Sciences Meeting. American Institute of Aeronautics and Astronautics (2016). <https://doi.org/10.2514/6.2016-1058>
15. Chan, J., Fernández, D.C.D.R., Carpenter, M.H.: Efficient entropy stable Gauss collocation methods. *SIAM J. Sci. Comput.* **41**(5), A2938–A2966 (2019). <https://doi.org/10.1137/18M1209234>
16. Christopher, L.R.: The NASA juncture flow test as a model for effective CFD/experimental collaboration. In: 2018 Applied Aerodynamics Conference. American Institute of Aeronautics and Astronautics (2018). <https://doi.org/10.2514/6.2018-3319>
17. Citro, V., Giannetti, F., Sierra, J.: Optimal explicit Runge-Kutta methods for compressible Navier-Stokes equations. *Appl. Numer. Math.* **152**, 511–526 (2020). <https://doi.org/10.1016/j.apnum.2019.11.005>
18. Conde, S., Fekete, I., Shadid, J.N.: Embedded error estimation and adaptive step-size control for optimal explicit strong stability preserving Runge-Kutta methods (2018). [arXiv:1806.08693](https://arxiv.org/abs/1806.08693) [math.NA]
19. Dormand, J.R., Prince, P.J.: A family of embedded Runge-Kutta formulae. *J. Comput. Appl. Math.* **6**(1), 19–26 (1980). [https://doi.org/10.1016/0771-050X\(80\)90013-3](https://doi.org/10.1016/0771-050X(80)90013-3)
20. Fernández, D.C.D.R., Carpenter, M.H., Dalcin, L., Zampini, S., Parsani, M.: Entropy stable h/p -nonconforming discretization with the summation-by-parts property for the compressible Euler and

- Navier-Stokes equations. *SN Partial Differ. Equ. Appl.* **1**(2), 1–54 (2020). <https://doi.org/10.1007/s42985-020-00009-z>
21. Figueroa, A., Jackiewicz, Z., Löhner, R.: Explicit two-step Runge-Kutta methods for computational fluid dynamics solvers. *Int. J. Numer. Methods Fluids* **93**(2), 429–444 (2021). <https://doi.org/10.1002/flid.4890>
  22. Fisher, T.C., Carpenter, M.H.: High-order entropy stable finite difference schemes for nonlinear conservation laws: finite domains. *J. Comput. Phys.* **252**, 518–557 (2013). <https://doi.org/10.1016/j.jcp.2013.06.014>
  23. Gassner, G.J., Winters, A.R., Kopriva, D.A.: Split form nodal discontinuous Galerkin schemes with summation-by-parts property for the compressible Euler equations. *J. Comput. Phys.* **327**, 39–66 (2016). <https://doi.org/10.1016/j.jcp.2016.09.013>
  24. Gottlieb, S., Ketcheson, D.I.: Time discretization techniques. In: Abgrall, R., Shu, C.-W. (eds) *Handbook of Numerical Analysis*, vol. 17, pp. 549–583. Elsevier (2016). <https://doi.org/10.1016/bs.hna.2016.08.001>
  25. Gustafsson, K.: Control theoretic techniques for stepsize selection in explicit Runge-Kutta methods. *ACM Trans. Math. Softw. (TOMS)* **17**(4), 533–554 (1991). <https://doi.org/10.1145/210232.210242>
  26. Gustafsson, K., Lundh, M., Söderlind, G.: A PI stepsize control for the numerical solution of ordinary differential equations. *BIT Numer. Math.* **28**(2), 270–287 (1988). <https://doi.org/10.1007/BF01934091>
  27. Hadri, B., Parsani, M., Hutchinson, M., Heinecke, A., Dalcin, L., Keyes, D.: Performance study of sustained petascale direct numerical simulation on Cray XC40 systems. *Concurrency Computat. Pract. Exper.* **32**(20), e5725 (2020). <https://doi.org/10.1002/cpe.5725>
  28. Hairer, E., Nørsett, S.P., Wanner, G.: *Solving Ordinary Differential Equations I: Nonstiff Problems*, Springer Series in Computational Mathematics, vol. 8. Springer-Verlag, Berlin/Heidelberg (2008). <https://doi.org/10.1007/978-3-540-78862-1>
  29. Hairer, E., Wanner, G.: *Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems*, Springer Series in Computational Mathematics, vol. 14. Springer-Verlag, Berlin/Heidelberg (2010). <https://doi.org/10.1007/978-3-642-05221-7>
  30. Hall, G., Higham, D.J.: Analysis of stepsize selection schemes for Runge-Kutta codes. *IMA J. Numer. Anal.* **8**(3), 305–310 (1988). <https://doi.org/10.1093/imanum/8.3.305>
  31. Higham, D.J., Hall, G.: Embedded Runge-Kutta formulae with stable equilibrium states. *J. Comput. Appl. Math.* **29**(1), 25–33 (1990). [https://doi.org/10.1016/0377-0427\(90\)90192-3](https://doi.org/10.1016/0377-0427(90)90192-3)
  32. Hutchinson, M., Heinecke, A., Pabst, H., Henry, G., Parsani, M., Keyes, D.: Efficiency of high order spectral element methods on petascale architectures. In: Kunkel, J., Balaji, P., Dongarra, J. (eds) *High Performance Computing. ISC High Performance 2016. Lecture Notes in Computer Science*, vol. 9697. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-41321-1_23
  33. Iyer, P.S., Malik, M.R.: Wall-modeled LES of the NASA juncture flow experiment. In: *AIAA Scitech 2020 Forum*, pp. 1–23 (2020). <https://doi.org/10.2514/6.2020-1307>
  34. Karniadakis, G.E., Sherwin, S.: *Spectral/hp Element Methods for Computational Fluid Dynamics*. Oxford University Press, Oxford (2013). <https://doi.org/10.1093/acprof:oso/9780198528692.001.0001>
  35. Kennedy, C.A., Carpenter, M.H.: Fourth order 2N-storage Runge-Kutta schemes. Technical Memorandum NASA-TM-109112, NASA, NASA Langley Research Center, Hampton VA 23681-0001, United States (1994)
  36. Kennedy, C.A., Carpenter, M.H.: Additive Runge-Kutta schemes for convection-diffusion-reaction equations. *Appl. Numer. Math.* **44**(1/2), 139–181 (2003). [https://doi.org/10.1016/S0168-9274\(02\)00138-1](https://doi.org/10.1016/S0168-9274(02)00138-1)
  37. Kennedy, C.A., Carpenter, M.H., Lewis, R.M.: Low-storage, explicit Runge-Kutta schemes for the compressible Navier-Stokes equations. *Appl. Numer. Math.* **35**(3), 177–219 (2000). [https://doi.org/10.1016/S0168-9274\(99\)00141-5](https://doi.org/10.1016/S0168-9274(99)00141-5)
  38. Ketcheson, D.I., Ahmadi, A.J.: Optimal stability polynomials for numerical integration of initial value problems. *Commun. Appl. Math. Comput. Sci.* **7**(2), 247–271 (2013). <https://doi.org/10.2140/camcos.2012.7.247>
  39. Ketcheson, D.I.: Highly efficient strong stability-preserving Runge-Kutta methods with low-storage implementations. *SIAM J. Sci. Comput.* **30**(4), 2113–2136 (2008). <https://doi.org/10.1137/07070485X>
  40. Ketcheson, D.I.: Runge-Kutta methods with minimum storage implementations. *J. Comput. Phys.* **229**(5), 1763–1773 (2010). <https://doi.org/10.1016/j.jcp.2009.11.006>
  41. Ketcheson, D.I.: Relaxation Runge-Kutta methods: conservation and stability for inner-product norms. *SIAM J. Numer. Anal.* **57**(6), 2850–2870 (2019). <https://doi.org/10.1137/19M1263662>. [arXiv:1905.09847](https://arxiv.org/abs/1905.09847) [math.NA]

42. Ketcheson, D.I., Parsani, M., Grant, Z.J., Ahmadi, A., Ranocha, H.: RK-Opt: a package for the design of numerical ODE solvers. *J. Open Source Softw.* **5**(54), 2514 (2020). <https://doi.org/10.21105/joss.02514>. <https://github.com/ketch/RK-Opt>
43. Ketcheson, D.I., Ranocha, H., Parsani, M., bin Waheed, U., Hadjimichael, Y.: NodePy: a package for the analysis of numerical ODE solvers. *J. Open Source Softw.* **5**(55), 2515 (2020). <https://doi.org/10.21105/joss.02515>. <https://github.com/ketch/nodepy>
44. Knepley, M.G., Karpeev, D.A.: Mesh algorithms for PDE with Sieve I: mesh distribution. *Sci. Program.* **17**(3), 215–230 (2009). <https://doi.org/10.3233/SPR-2009-0249>
45. Kopriva, D.A.: *Implementing Spectral Methods for Partial Differential Equations: Algorithms for Scientists and Engineers*. Springer Science & Business Media, New York (2009). <https://doi.org/10.1007/978-90-481-2261-5>
46. Kopriva, D.A., Jimenez, E.: An assessment of the efficiency of nodal discontinuous Galerkin spectral element methods. In: Ansoorge, R., Bijl, H., Meister, A., Sonar, T. (eds) *Recent Developments in the Numerics of Nonlinear Hyperbolic Conservation Laws*, pp. 223–235. Springer, Berlin/Heidelberg (2013). https://doi.org/10.1007/978-3-642-33221-0_13
47. Kraaijevanger, J.F.B.M.: Contractivity of Runge-Kutta methods. *BIT Numer. Math.* **31**(3), 482–528 (1991). <https://doi.org/10.1007/BF01933264>
48. Kubatko, E.J., Dawson, C., Westerink, J.J.: Time step restrictions for Runge-Kutta discontinuous Galerkin methods on triangular grids. *J. Comput. Phys.* **227**(23), 9697–9710 (2008)
49. Langseth, J.O., LeVeque, R.J.: A wave propagation method for three-dimensional hyperbolic conservation laws. *J. Comput. Phys.* **165**(1), 126–166 (2000)
50. Langtry, R.B., Kuntz, M., Menter, F.R.: Drag prediction of engine-airframe interference effects with CFX-5. *J. Aircr.* **42**(6), 1523–1529 (2005)
51. LeVeque, R.J.: *Finite Difference Methods for Ordinary and Partial Differential Equations: Steady-State and Time-Dependent Problems*. SIAM, Philadelphia, PA, USA (2007)
52. Mogensen, P.K., Riseth, A.N.: Optim: a mathematical optimization package for Julia. *J. Open Source Softw.* **3**(24), 615 (2018). <https://doi.org/10.21105/joss.00615>
53. Montijano, J.L., Rández, L., Ketcheson, D.I.: Low-storage FSAL embedded pairs of Runge-Kutta methods (2020) (In preparation)
54. O'Reilly, O., Lundquist, T., Dunham, E.M., Nordström, J.: Energy stable and high-order-accurate finite difference methods on staggered grids. *J. Comput. Phys.* **346**, 572–589 (2017)
55. Parsani, M., Boukharfane, R., Nolasco, I.R., Fernández, D.C.D.R., Zampini, S., Hadri, B., Dalcin, L.: High-order accurate entropy-stable discontinuous collocated Galerkin methods with the summation-by-parts property for compressible CFD frameworks: scalable SSDC algorithms and flow solver. *J. Comput. Phys.* **424**, 109844 (2021). <https://doi.org/10.1016/j.jcp.2020.109844>
56. Parsani, M., Carpenter, M.H., Nielsen, E.J.: Entropy stable discontinuous interfaces coupling for the three-dimensional compressible Navier-Stokes equations. *J. Comput. Phys.* **290**, 132–138 (2015). <https://doi.org/10.1016/j.jcp.2015.02.042>
57. Parsani, M., Carpenter, M.H., Nielsen, E.J.: Entropy stable wall boundary conditions for the three-dimensional compressible Navier-Stokes equations. *J. Comput. Phys.* **292**, 88–113 (2015). <https://doi.org/10.1016/j.jcp.2015.03.026>
58. Parsani, M., Ketcheson, D.I., Deconinck, W.: Optimized low-order explicit Runge-Kutta schemes for the high-order spectral difference method. In: *Proceedings of the 11th Finnish Mechanics Days*, pp. 49–54. University of Oulu, Department of Mechanical Engineering (2012). <http://hdl.handle.net/10754/333613>
59. Parsani, M., Ketcheson, D.I., Deconinck, W.: Optimized explicit Runge-Kutta schemes for the spectral difference method applied to wave propagation problems. *SIAM J. Sci. Comput.* **35**(2), A957–A986 (2013). <https://doi.org/10.1137/120885899>
60. Pegrum, J.: *Experimental study of the vortex system generated by a Formula 1 front wing*. Ph.D. thesis, Imperial College London (2007)
61. Prince, P.J., Dormand, J.R.: High order embedded Runge-Kutta formulae. *J. Comput. Appl. Math.* **7**(1), 67–75 (1981). [https://doi.org/10.1016/0771-050X\(81\)90010-3](https://doi.org/10.1016/0771-050X(81)90010-3)
62. Rackauckas, C., Nie, Q.: DifferentialEquations.jl—a performant and feature-rich ecosystem for solving differential equations in Julia. *J. Open Res. Softw.* **5**(1), 15 (2017). <https://doi.org/10.5334/jors.151>
63. Ranocha, H., Dalcin, L., Parsani, M.: Fully-discrete explicit locally entropy-stable schemes for the compressible Euler and Navier-Stokes equations. *Comput. Math. Appl.* **80**(5), 1343–1359 (2020). <https://doi.org/10.1016/j.camwa.2020.06.016>. [arXiv:2003.08831](https://arxiv.org/abs/2003.08831) [math.NA]
64. Ranocha, H., Dalcin, L., Parsani, M., Ketcheson, D.I.: Coefficients of optimized low-storage Runge-Kutta methods with automatic step size control for spectral element methods applied to compressible

- computational fluid dynamics. 2021. <https://github.com/ranocha/Optimized-RK-CFD>, <https://doi.org/10.5281/zenodo.4671927>
65. Ranocha, H., Lóczy, L., Ketcheson, D.I.: General relaxation methods for initial-value problems with application to multistep schemes. *Numer. Math.* **146**, 875–906 (2020). <https://doi.org/10.1007/s00211-020-01158-4>. [arXiv:2003.03012](https://arxiv.org/abs/2003.03012) [math.NA]
  66. Ranocha, H., Sayyari, M., Dalcin, L., Parsani, M., Ketcheson, D.I.: Relaxation Runge-Kutta methods: fully-discrete explicit entropy-stable schemes for the compressible Euler and Navier-Stokes equations. *SIAM J. Sci. Comput.* **42**(2), A612–A638 (2020). <https://doi.org/10.1137/19M1263480>. [arXiv:1905.09129](https://arxiv.org/abs/1905.09129) [math.NA]
  67. Rojas, D., Boukharfane, R., Dalcin, L., Fernández, D.C.D.R., Ranocha, H., Keyes, D.E., Parsani, M.: On the robustness and performance of entropy stable discontinuous collocation methods. *J. Comput. Phys.* **426**, 109891 (2021). <https://doi.org/10.1016/j.jcp.2020.109891>. [arXiv:1911.10966](https://arxiv.org/abs/1911.10966) [math.NA]
  68. Rumsey, C.L., Morrison, J.H.: Goals and status of the NASA juncture flow experiment, p. STO-MP-AVT-246. , NATO (2016)
  69. Shu, C.-W.: Essentially non-oscillatory and weighted essentially non-oscillatory schemes for hyperbolic conservation laws. Final Report NASA/CR-97-206253, NASA, Institute for Computer Applications in Science and Engineering, NASA Langley Research Center, Hampton VA United States (1997)
  70. Shu, C.-W., Osher, S.: Efficient implementation of essentially non-oscillatory shock-capturing schemes. *J. Comput. Phys.* **77**(2), 439–471 (1988). [https://doi.org/10.1016/0021-9991\(88\)90177-5](https://doi.org/10.1016/0021-9991(88)90177-5)
  71. Sjögreen, B., Yee, H.: High order entropy conservative central schemes for wide ranges of compressible gas dynamics and MHD flows. *J. Comput. Phys.* **364**, 153–185 (2018). <https://doi.org/10.1016/j.jcp.2018.02.003>
  72. Sjögreen, B., Yee, H.C., Kotov, D.: Skew-symmetric splitting and stability of high order central schemes. In: *Journal of Physics: Conference Series*, vol. 837, p. 012019. IOP Publishing (2017). <https://doi.org/10.1088/1742-6596/837/1/012019>
  73. Söderlind, G.: Automatic control and adaptive time-stepping. *Numer. Algorithms* **31**(1/2/3/4), 281–310 (2002). <https://doi.org/10.1023/A:1021160023092>
  74. Söderlind, G.: Digital filters in adaptive time-stepping. *ACM Trans. Math. Softw. (TOMS)* **29**(1), 1–26 (2003). <https://doi.org/10.1145/641876.641877>
  75. Söderlind, G.: Time-step selection algorithms: adaptivity, control, and signal processing. *Appl. Numer. Math.* **56**(3/4), 488–502 (2006). <https://doi.org/10.1016/j.apnum.2005.04.026>
  76. Söderlind, G., Wang, L.: Adaptive time-stepping and computational stability. *J. Comput. Appl. Math.* **185**(2), 225–243 (2006). <https://doi.org/10.1016/j.cam.2005.03.008>
  77. Tsitouras, C.: Runge-Kutta pairs of order 5(4) satisfying only the first column simplifying assumption. *Comput. Math. Appl.* **62**(2), 770–775 (2011). <https://doi.org/10.1016/j.camwa.2011.06.002>
  78. Vincent, P.E., Castonguay, P., Jameson, A.: A new class of high-order energy stable flux reconstruction schemes. *J. Sci. Comput.* **47**, 50–72 (2011). <https://doi.org/10.1007/s10915-010-9420-z>
  79. Vincent, P., Witherden, F., Vermeire, B., Park, J.S., Iyer, A.: Towards green aviation with Python at petascale. In: *SC'16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–11. IEEE Press (2016)
  80. Ware, J., Berzins, M.: Adaptive finite volume methods for time-dependent PDEs. In: *Modeling, Mesh Generation, and Adaptive Numerical Methods for Partial Differential Equations*, pp. 417–430. Springer (1995). https://doi.org/10.1007/978-1-4612-4248-2_20