Journal of Computational Physics 257 (2014) 594–626

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

A stabilized Runge–Kutta–Legendre method for explicit super-time-stepping of parabolic and mixed equations

Chad D. Meyer [a,*], Dinshaw S. Balsara [a], Tariq D. Aslam [b]

[a] Physics Department, Univ. of Notre Dame, 225 Nieuwland Science Hall, Notre Dame, IN 46556, USA
[b] WX-9 Group, Los Alamos National Laboratory, MS P952, Los Alamos, NM 87545, USA

ARTICLE INFO

Article history:
Received 8 February 2013
Received in revised form 12 July 2013
Accepted 8 August 2013
Available online 23 September 2013

Keywords:
Numerical methods
Super-time-stepping
Parabolic operators
PDEs

ABSTRACT

Parabolic partial differential equations appear in several physical problems, including problems that have a dominant hyperbolic part coupled to a sub-dominant parabolic component. Explicit methods for their solution are easy to implement but have very restrictive time step constraints. Implicit solution methods can be unconditionally stable but have the disadvantage of being computationally costly or difficult to implement. Super-time-stepping methods for treating parabolic terms in mixed type partial differential equations occupy an intermediate position. In such methods each superstep takes "s" explicit Runge–Kutta-like time-steps to advance the parabolic terms by a time-step that is s^2 times larger than a single explicit time-step. The expanded stability is usually obtained by mapping the short recursion relation of the explicit Runge–Kutta scheme to the recursion relation of some well-known, stable polynomial.

Prior work has built temporally first- and second-order accurate super-time-stepping methods around the recursion relation associated with Chebyshev polynomials. Since their stability is based on the boundedness of the Chebyshev polynomials, these methods have been called RKC1 and RKC2. In this work we build temporally first- and second-order accurate super-time-stepping methods around the recursion relation associated with Legendre polynomials. We call these methods RKL1 and RKL2. The RKL1 method is first-order accurate in time; the RKL2 method is second-order accurate in time. We verify that the newly-designed RKL1 and RKL2 schemes have a very desirable monotonicity preserving property for one-dimensional problems – a solution that is monotone at the beginning of a time step retains that property at the end of that time step. It is shown that RKL1 and RKL2 methods are stable for all values of the diffusion coefficient up to the maximum value. We call this a convex monotonicity preserving property and show by examples that it is very useful in parabolic problems with variable diffusion coefficients. This includes variable coefficient parabolic equations that might give rise to skew symmetric terms. The RKC1 and RKC2 schemes do not share this convex monotonicity preserving property. One-dimensional and two-dimensional von Neumann stability analyses of RKC1, RKC2, RKL1 and RKL2 are also presented, showing that the latter two have some advantages. The paper includes several details to facilitate implementation.

A detailed accuracy analysis is presented to show that the methods reach their design accuracies. A stringent set of test problems is also presented. To demonstrate the robustness and versatility of our methods, we show their successful operation on problems involving linear and non-linear heat conduction and viscosity, resistive magnetohydrodynamics, ambipolar diffusion dominated magnetohydrodynamics, level set methods and flux limited radiation diffusion. In a prior paper (Meyer, Balsara and Aslam 2012 [36]) we

* Corresponding author.

0021-9991/$ – see front matter
http://dx.doi.org/10.1016/j.jcp.2013.08.021

have also presented an extensive test-suite showing that the RKL2 method works robustly in the presence of shocks in an anisotropically conducting, magnetized plasma.

1. Introduction

A large number of important physical problems have their basis in parabolic partial differential equations (PDEs henceforth). Thermal and mass diffusion (Baehr and Stephan 2006 [8]), viscous fluid flow (Gassner, Lörcher and Munz 2007, 2008 [18,19]), viscous and resistive magnetohydrodynamics (Dumbser and Balsara 2010 [17]) are just a few examples. Furthermore, many subscale physical models employ non-linear parabolic or mixed hyperbolic–parabolic PDEs, such as Detonation Shock Dynamics (Aslam, Bdzil and Stewart 1996 [4]), flame propagation (Markstein 1965 [33]), Stefan problems (Rubenstein 1971 [44]) and thermal conduction (Balsara, Tilley and Howk 2008 [10]). Additionally, many mathematical models related to level set methods, image processing and computational geometry (Osher and Fedkiw 2002 [39], Sethian 1999 [45], and references therein) are rooted in non-linear parabolic PDEs.

While there are many physical problems related to non-linear parabolic PDEs, there is no consensus as to the numerical strategies to solve such problems. Explicit methods, while being quite simple and accurate to implement, suffer from an overly restrictive time-step to satisfy numerical stability. Implicit schemes are moderately efficient, but can be comparatively complicated to use in practice, and can sometimes lack robustness for sufficiently non-linear problems. When dealing with mixed hyperbolic–parabolic PDEs, often one utilizes methods suitable for each part of the equation through a Strang (Strang 1968 [47]) or operator splitting technique. The splitting schemes can be constructed up to second-order accuracy, but for sufficiently difficult problems can lead to poor solutions (LeVeque and Yee 1990 [32]).

A different approach, known as Runge–Kutta–Chebyshev (RKC henceforth), stabilized Runge–Kutta or super-time-stepping methods use "s" strategically devised explicit Runge–Kutta (RK) stages. See van Der Houwen and Sommeijer 1980 [50], van der Houwen 1996 [49], Verwer 1996 [52], Hairer and Wanner 1996 [22], Sommeijer, Shampine and Verwer 1998 [46], Alexiades, Amiez and Gremaud 1996 [3], O'Sullivan and Downes 2007 [38], Gurski and O'Sullivan 2011 [20], Gurski 2011 [21] for more information. The internal stages in the explicit RK method are intelligently chosen so that the overall time-step is stable up to $\sim s^2 \Delta t_{expl}$. Here Δt_{expl} is the maximum stable explicit time-step for the parabolic terms. These methods boast large efficiency gains compared to standard explicit schemes. While RKC1 and RKC2 schemes, as well as the ROCK schemes (Abdulle 2001 [1], 2002 [2], Medovikov 1998 [35], Lebedev 1994 [28], 2000 [29]) have been used with success, in this paper we show that there are improvements that can be made, especially when dealing with non-linear parabolic terms. We should also mention that all of the RK schemes described in this paragraph and also those introduced in this paper are designed to have short recursive sequences, which ensures that their memory requirements are reasonably modest.

When the diffusion coefficient is spatially varying or solution-dependent, the diffusion problems can produce steep fronts (Marshak 1958 [34], Reinicke and Meyer-ter-Vehn 1991 [42], Mousseau, Knoll and Rider 2000 [37]). For non-linear diffusion operators, it is difficult to prove stability of the RKC schemes. Often, the RKC schemes are augmented by an added damping to the underlying methods (Hundsdorfer and Verwer 2003 [25] and references therein), which serve to help stabilize the schemes for non-linear problems.

Here, we introduce the notion of a convexity property for these non-linear parabolic PDEs, where the underlying scheme can be shown to be stable for any diffusion coefficient less than or equal to the maximum diffusion coefficient. Our motivation for these new methods arises from a monotonicity analysis (LeVeque 1992 [31]) of the s stage Runge–Kutta scheme where we allow the diffusion coefficient to vary from zero to the maximum, and simultaneously enforce the monotonicity of the solution. For a scalar, purely parabolic operator, without source terms, the solution of the underlying PDEs should remain bounded by its initial data and, furthermore, remain monotone if it is initially monotone. In a thermal conduction problem, the temperature should not oscillate or form new extrema as time evolves.

These new methods are closely related to the earlier mentioned RKC schemes, but the convexity property adds an extra stability condition that is not present in those Super-time-stepping schemes. The new schemes presented in this paper are related to shifted Legendre polynomials, rather than shifted Chebyshev polynomials, and so we call these new schemes Runge–Kutta–Legendre (RKL henceforth). We refer to the formulation which is first-order accurate in time as RKL1 and the formulation which is second-order accurate in time as RKL2. The utility of the RKL2 scheme for thermal conduction problems, where the conduction coefficient is solution-dependent, was shown in Meyer, Balsara and Aslam (2012) [36]. While we have verified the convexity property on a case by case basis (up to $s = 64$), we hope mathematicians will generalize it for all s.

ROCK schemes [1,2,28,29,35] based on the RKC methodology are fourth-order accurate in time. We only consider RKL methods up to second-order. When coupled with a hyperbolic method, the temporal error due to operator splitting will be second-order in time, erasing any gains from a higher temporal order of accuracy in the parabolic solver. Higher order accurate RKL methods could be of interest for purely parabolic problems, or in an operator-unsplit arrangement, but that is outside of the scope of this work.

This paper is structured as follows. In Section 2 we motivate the method with simple examples and derive the RKL1 and RKL2 schemes. In Section 3, we introduce the convex monotone property for explicit parabolic methods. In Section 4,

we discuss the von Neumann stability of these methods. In Section 5, we discuss using non-linear operators with the RKL methods. In Section 6 we show the order of accuracy of the method. In Section 7, we demonstrate the usefulness of the RKL methods in a wide range of example problems. In Section 8, we compare the computational performance of the RKL methods to other, classical numerical methods. Section 9 contains our conclusions.

2. Description of the stabilized Runge–Kutta–Legendre method at first- and second-order

The explicit solution of parabolic equations suffers from the unfortunate restriction that the maximal explicit time-step is proportional to the square of the mesh size. Implicit schemes for parabolic equations may have no time-step restriction, but they require the inversion of large, sparse matrices, which can be computationally costly and difficult to parallelize. Super-time-stepping methods, such as RKC, have been developed to help alleviate this restrictive time-step constraint for parabolic equations while retaining the simplicity of an explicit method. Super-time-stepping methods are designed to solve the ODE system resulting from the discretization of a PDE. This takes the form

$$\frac{du}{dt} = \mathbf{M}u(t) \tag{1}$$

where $\mathbf{M}$ is a symmetric, constant coefficient matrix which represents the discretization of a parabolic operator. Such parabolic operators have non-positive, real eigenvalues. Later in this paper, we will relax this restriction on the parabolic operator $\mathbf{M}$, and thereby include operators with diffusion coefficients that are solution or position-dependent. Super-time-stepping methods fall into a category of stabilized Runge–Kutta methods for which additional stages are used to increase the stability of the method along the negative real axis, thereby increasing the maximum permissible time-step. This should be contrasted with traditional Runge–Kutta methods, where the additional stages are used to increase the accuracy of the method. The update of an s-stage system can be expressed in terms of the stability polynomial, also referred to as the amplification factor, R_s

$$u(t+\tau) = R_s(\tau\mathbf{M})u(t). \tag{2}$$

Stability is ensured if $|R_s(\tau\lambda)| \leqslant 1$ for all values of λ between 0 and the maximum (negative) eigenvalue of the matrix $\mathbf{M}$. The RKC methods are based on using shifted Chebyshev polynomials as the stability polynomial because the absolute value of these polynomials is bounded by unity over a large range of values. Additionally, they afford the greatest stability range along the negative real axis, thereby permitting one to take the largest possible stable superstep, τ. The desired temporal order of accuracy is achieved by comparing the stability polynomial in Eq. (2) to the analytic solution

$$u(t) = e^{t\mathbf{M}}u(0) \approx \left(1 + t\mathbf{M} + \frac{1}{2}(t\mathbf{M})^2 + \cdots\right)u(0). \tag{3}$$

We can achieve the desired order of accuracy by ensuring that the leading terms in the stability polynomial match exactly with the expansion of the exponential in the exact solution.

For a simple problem such as constant coefficient thermal conduction, the RKC methods have been shown to work well. It is not easy to show that more complicated problems will also behave well using the RKC methodology, but it is often suggested that they will. This, however, is not guaranteed, and Section 2.1 shows some examples where RKC methods show their deficiencies. Sections 2.2 and 2.3 catalog the first-order accurate RKL1 scheme and the second-order accurate RKL2 schemes, respectively. Section 2.4 describes how these schemes are implemented using Strang splitting for mixed problems which include hyperbolic and parabolic terms.

2.1. Motivation

Consider the simple example of thermal conduction between two different metals. The metals are initially at different temperatures and have different thermal conduction coefficients. When they are brought into thermal contact, heat will flow from hot to cold. We initialize the simulation with a 10 cm bar of aluminum (with conduction coefficient $\kappa = 2.37 \times 10^7$ erg/(cm s K)) at a temperature of 0 °C on the left, and a 10 cm bar of copper (with conduction coefficient $\kappa = 4.01 \times 10^7$ erg/(cm s K)) at a temperature of 100 °C on the right. This problem is discussed in depth in Section 6.1. Fig. 1(a) shows the results of this problem after 1 second of conduction, computed following the standard RKC2 methodology. The result is close to the analytical solution (shown as a solid line), but the method is not monotone and produces spurious oscillations within the profile. After many more time-steps the jumps shown in Fig. 1(a) will decrease in amplitude, but they could still be the site of unphysical shock formation if coupled with a hyperbolic operator. In addition, if a problem had solution-dependent coefficients, the oscillations could cause additional numerical errors or instability. These oscillations are due to the fact that although the RKC methods are stable and monotone for a constant conduction coefficient, they are not monotone for all conduction coefficients less than the maximum. When the conduction coefficient is spatially varying, the overall time-step must be limited by the most restrictive zone (i.e. the one with the highest conduction coefficient on a uniform mesh). We would therefore prefer a method which can retain a monotone solution for any conduction coefficient less than the maximum. Computationally, this is equivalent to wanting a monotone method for all

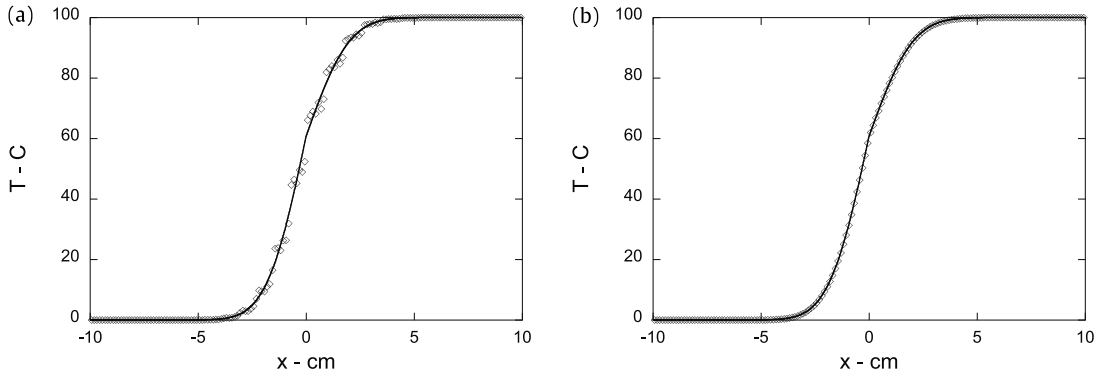

Fig. 1. (a) Shows the aluminum/copper heat conduction solution for the damped 7-stage RKC2 scheme, with damping coefficient $\varepsilon = 2/13$. (b) Shows the solution for the 7-stage RKL2 scheme.

time-steps shorter than the maximum. We refer to this property as the Convex Monotone Property (CMP), and discuss it at length in Section 3.

Fig. 1(b) shows the same result using the RKL2 scheme presented in this paper. The RKL2 method is described in Section 2.3 below. The solution, in this case, not only follows the analytic solution very closely compared to its RKC2 counterpart in Fig. 1(a), but is also monotone, with no spurious oscillations within the conduction profile. In Section 3, we show that the RKL methods have the CMP we desire. This result demonstrates that a method with the CMP for a parabolic problem with spatially varying coefficients will produce a much improved solution over a method which does not have the CMP.

Even if the solution is stable and monotone, we want to make sure that it does not have any other undesirable pathologies. To explore this, we can compute thermal conduction in a single metal but with a discontinuous initial temperature. Here, the left half of a 20 cm bar of aluminum is initially held at 0 °C, while the right half is held at 100 °C. Figs. 2(a)–2(d) show the results of this simulation for the RKC1, RKC2, RKL1 and RKL2 schemes, respectively. These tests were all run with 6 supersteps of 5 stages each, with the maximum stable time-step for each scheme. This ensures that each case will be both stable and monotone, but it also means that each case is run to a different final time. The most notable feature of the RKC schemes in Figs. 2(a) and 2(b) is the significant staircasing effect seen in the conduction profile. Although the result is monotone, these profiles show significant discontinuities even where a smooth solution is expected. This can be particularly problematic when the parabolic term is coupled with, for instance, the hydrodynamics equations, because a staircased pattern will be treated as a sequence of shocks by the hydrodynamics solver. Figs. 2(c) and 2(d) show the same problem with RKL1 and RKL2, respectively. The staircasing effect is much smaller for the RKL1 scheme and is virtually non-existent for the RKL2 scheme. Additionally, the solutions produced by the RKL schemes represent the expected thermal profile very accurately. In Section 4, we discuss how this comes about by appealing to the von Neumann stability analysis of the schemes.

2.2. RKL1: Runge–Kutta–Legendre method at first-order

The Runge–Kutta–Legendre (RKL) methods are based on using shifted Legendre polynomials as the stability polynomial of the scheme. Legendre polynomials, like Chebyshev polynomials, are bounded in magnitude by unity, and thus will also produce a stable scheme. The stability polynomial for the general s-stage RKL scheme can be written as

$$R_s(z) = a_s + b_s P_s(w_0 + w_1 z). \tag{4}$$

The parameter w_0 relates to the damping of the system, which we find is unnecessary for all RKL schemes, as discussed in Sections 4 and 5. Thus, we set $w_0 = 1$ for all RKL schemes. For a first-order scheme we choose $a_s = 0$. The value of the other parameters are set by comparison to the expansion in Eq. (3). For consistency at first-order, we have that $R_s(0) = 1$ and $R_s'(0) = 1$. This gives us $b_s = 1$ and $w_1 = 2/(s^2 + s)$. Therefore, for the s-stage RKL1 scheme we take

$$R_s(z) = P_s\left(1 + \frac{2}{s^2 + s} z\right). \tag{5}$$

Like the RKC scheme, we also require that every internal stage also has a corresponding stability polynomial which is a shifted Legendre polynomial, i.e. for the j-th stage of an s stage RKL1 scheme, we take

$$u^{(j)} = P_j\left(1 + \frac{2}{s^2 + s} z\right) u^n. \tag{6}$$

Thus, each stage in the scheme can be thought of as a first-order accurate approximation to the solution at time $t = \tau(j^2 + j)/(s^2 + s)$, where τ is the total time-step to be taken. Legendre polynomials obey a three point recursion relationship

$$(j) P_j(x) = (2j - 1) x P_{j-1}(x) - (j - 1) P_{j-2}(x) \tag{7}$$

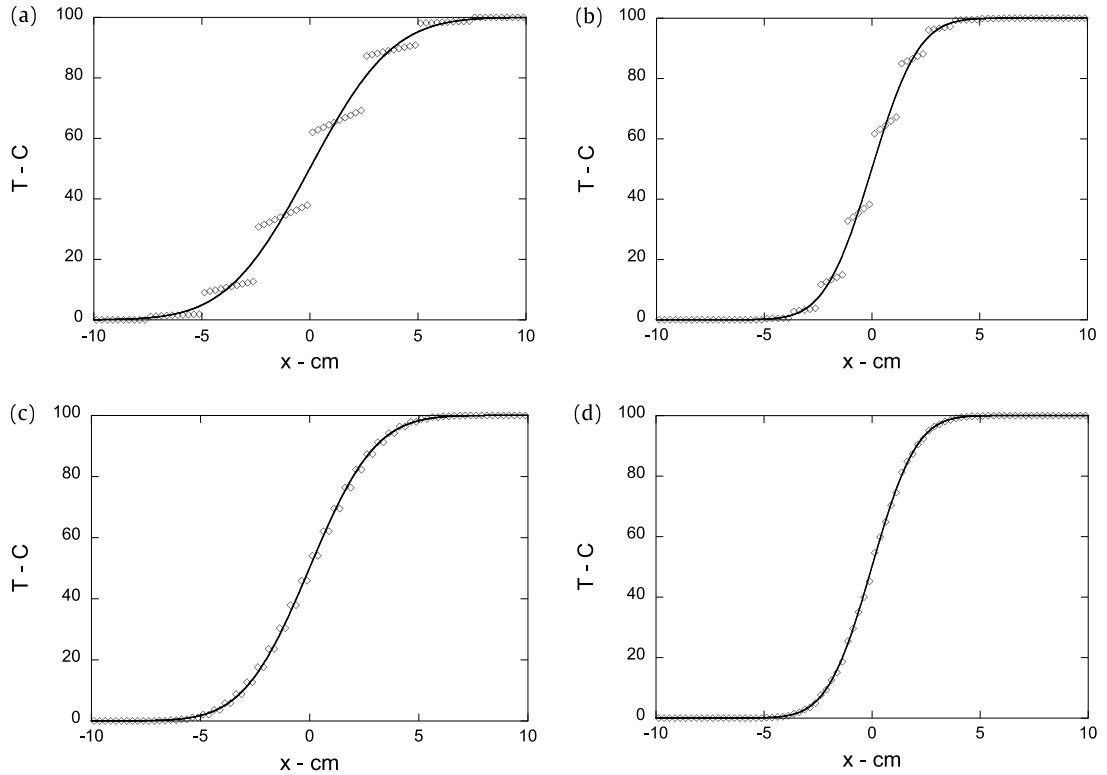

Fig. 2. Shows the heat conduction problem in aluminum with 5 stages and 6 supersteps for (a) RKC1 at $t = 4.67269$, (b) RKC2 at $t = 1.51382$, (c) RKL1 at $t = 2.89546$, (d) RKL2 at $t = 1.35121$.

and it is easily shown that the stability polynomials for RKL1 obey

$$P_j\left(1+\frac{2}{s^2+s}z\right) = \frac{2j-1}{j}\left(1+\frac{2}{s^2+s}z\right)P_{j-1}\left(1+\frac{2}{s^2+s}z\right) - \frac{j-1}{j}P_{j-2}\left(1+\frac{2}{s^2+s}z\right). \tag{8}$$

This recursion relation in Eq. (8) enables us to write down a Runge–Kutta scheme for operators **M** that have non-positive, real eigenvalues. Therefore, the s-stage RKL1 scheme is given by

$$\begin{aligned} &Y_0 = u(t_0)\\ &Y_1 = Y_0 + \tilde{\mu}_1\tau\mathbf{M}Y_0\\ &Y_j = \mu_j Y_{j-1} + \nu_j Y_{j-2} + \tilde{\mu}_j\tau\mathbf{M}Y_{j-1}; \quad 2 \leqslant j \leqslant s\\ &u(t+\tau) = Y_s \end{aligned} \tag{9}$$

with the parameters

$$\begin{aligned} &\mu_j = \frac{2j-1}{j}; \qquad \nu_j = \frac{1-j}{j}\\ &\tilde{\mu}_j = \frac{2j-1}{j}w_1 = \frac{2j-1}{j}\frac{2}{s^2+s}. \end{aligned} \tag{10}$$

The only thing that remains is the choice of the total time-step, τ, for the RKL1 scheme. Recall that Legendre Polynomials are bounded by unity in magnitude when their arguments are in the range $(-1, 1)$ so that the condition $-1 \leqslant 1 + 2/(s^2+s)z$ yields the time-step restriction. The explicit time-step is related to the spectral radius of the matrix by $\Delta t_{expl} = 2/\lambda_{\max}$. Therefore, we have that the RKL1 scheme has a maximum superstep of

$$\tau_{\max} = \frac{\Delta t_{expl}}{w_1} = \Delta t_{expl}\frac{s^2+s}{2}. \tag{11}$$

For the purpose of checking one's implementation of this method, we provide below an explicit instantiation of the method for $s = 4$. In this case, $w_1 = 1/10$, and the stability polynomials for the intermediate stages will be $P_j(1+\tau\mathbf{M}/10)$. The parameters at the various stages are given by

$$\mu_2 = \frac{3}{2}; \qquad \mu_3 = \frac{5}{3}; \qquad \mu_4 = \frac{7}{4}$$
$$\nu_2 = -\frac{1}{2}; \qquad \nu_3 = -\frac{2}{3}; \qquad \nu_4 = -\frac{3}{4}$$
$$\tilde{\mu}_1 = \frac{1}{10}; \qquad \tilde{\mu}_2 = \frac{3}{20}; \qquad \tilde{\mu}_3 = \frac{5}{30}; \qquad \tilde{\mu}_4 = \frac{7}{40} \tag{12}$$

and, the scheme is written as

$$\begin{aligned} Y_0 &= u(t_0) \\ Y_1 &= Y_0 + \frac{1}{10}\tau \mathbf{M} Y_0 \\ Y_2 &= \frac{3}{2}Y_1 - \frac{1}{2}Y_0 + \frac{3}{20}\tau \mathbf{M} Y_1 \\ Y_3 &= \frac{5}{3}Y_2 - \frac{2}{3}Y_1 + \frac{5}{30}\tau \mathbf{M} Y_2 \\ Y_4 &= \frac{7}{4}Y_3 - \frac{3}{4}Y_2 + \frac{7}{40}\tau \mathbf{M} Y_3 \\ u(t_0+\tau) &= Y_4. \end{aligned} \tag{13}$$

Here, Y_0 is the solution at the start of the time-step and Y_1, Y_2, Y_3 and Y_4 are the internal stages of the four-stage RKL1 scheme. It is possible to verify by back-substitution that each of the stages in Eq. (13) can be written as $Y_j = P_j(1 + w_1\tau\mathbf{M})Y_0$. Notice that the minimum storage required in computer memory is just three times the size of Y_0. This completes our description of the RKL1 method.

It is worth mentioning that there is an additional way of producing a first-order accurate RK method, which was employed in [3]. Rather than employing a recurrence relation like in Eq. (8), the stability polynomial in Eq. (2) can be factored

$$R_s(\tau\mathbf{M}) = \prod_{i=1}^{s}(1 + \Delta t_i \mathbf{M}). \tag{14}$$

This can be thought of as a series of forward-Euler time-steps. Although the overall stability polynomial is still the same, the intermediate stages are no longer guaranteed to be stable as some of the values of Δt_i will be longer than Δt_{expl}. For a linear operator, this will still produce a stable, first-order accurate solution at time $t + \tau$. When incorporating non-linear or asymmetric parabolic operators, however, this could result in an unstable scheme. The recursive procedure described in this paper ensures that each of the intermediate stages will be stable, consistent and accurate, which ensures a greater robustness in the scheme.

2.3. RKL2: Runge–Kutta–Legendre method at second-order

For a second-order accurate update, we can start with the same general stability polynomial as in Eq. (4), but we also require that the quadratic term is consistent with the expansion in Eq. (3). As we have done above in RKL1, we choose $w_0 = 1$, corresponding to the case without damping. Second-order accuracy requires that $R_s(0) = 1$, $R_s'(0) = 1$ and $R_s''(1) = 1$, where R_s is the stability polynomial. These requirements constrain the coefficients of the stability polynomial in Eq. (4) to be

$$a_s = 1 - b_s; \qquad b_s = \frac{P_s''(1)}{(P_s'(1))^2} = \frac{s^2+s-2}{2s(s+1)}; \qquad w_1 = \frac{P_s'(1)}{P_s''(1)} = \frac{4}{s^2+s-2}. \tag{15}$$

For $s < 2$, we have the freedom to choose $b_0 = b_1 = b_2 = 1/3$. Substituting this into the recursion relationship in Eq. (7) and rearranging terms, we have

$$\begin{aligned} a_j + b_j P_j(1 + w_1 z) = {} & \left(\frac{2j-1}{j}\right)\frac{b_j}{b_{j-1}}\left(a_{j-1} + b_{j-1}P_{j-1}(1+w_1 z)\right) \\ & + \left(\frac{1-j}{j}\right)\frac{b_j}{b_{j-2}}\left(a_{j-2} + b_{j-2}P_{j-2}(1+w_1 z)\right) \\ & + \left(\frac{2j-1}{j}\right)\frac{b_j}{b_{j-1}} w_1 z\left(a_{j-1} + b_{j-1}P_{j-1}(1+w_1 z)\right) \\ & + \left(1 - \left(\frac{2j-1}{j}\right)\frac{b_j}{b_{j-1}} - \left(\frac{1-j}{j}\right)\frac{b_j}{b_{j-2}}\right) - a_{j-1}\left(\frac{2j-1}{j}\right)\frac{b_j}{b_{j-1}} w_1 z. \end{aligned} \tag{16}$$

This can be written out as a numerical scheme

$$\begin{aligned} Y_0 &= u(t_0) \\ Y_1 &= Y_0 + \tilde{\mu}_1 \tau \mathbf{M} Y_0 \\ Y_j &= \mu_j Y_{j-1} + \nu_j Y_{j-2} + (1-\mu_j-\nu_j) Y_0 + \tilde{\mu}_j \tau \mathbf{M} Y_{j-1} + \tilde{\gamma}_j \tau \mathbf{M} Y_0; \quad 2 \leqslant j \leqslant s \\ u(t_0+\tau) &= Y_s \end{aligned} \tag{17}$$

with the parameters

$$\begin{aligned} \mu_j &= \frac{2j-1}{j}\frac{b_j}{b_{j-1}} = \frac{(2j-1)(j+2)(j-1)^2}{j(j-2)(j+1)^2} \\ \nu_j &= -\frac{j-1}{j}\frac{b_j}{b_{j-2}} = -\frac{(j-1)^3(j^2-4)}{j^3(j+1)(j-3)} \\ \tilde{\mu}_j &= \mu_j w_1 = \frac{(2j-1)(j+2)(j-1)^2}{j(j-2)(j+1)^2}\frac{4}{s^2+s-2}, \qquad \tilde{\mu}_1 = b_1 w_1 = \frac{4}{3(s^2+s-2)} \\ \tilde{\gamma}_j &= -a_{j-1}\tilde{\mu}_j = -\frac{(j-1)(j+2)(2j-1)(j^2-j+2)}{2j^2(j-2)(j+1)^2}. \end{aligned} \tag{18}$$

The maximum stable superstep for the RKL2 scheme is computed in the same way as in the previous subsection

$$\tau_{\max} = \frac{\Delta t_{expl}}{w_1} = \Delta t_{expl}\frac{s^2+s-2}{4}. \tag{19}$$

Notice that the requirement of second-order accuracy in time has reduced the timestep in Eq. (19) compared to the timestep limitation in Eq. (11). The j-th sub-stage in an s-stage scheme will be a second-order accurate approximation of the solution at time $t = \tau(j^2+j-2)/(s^2+s-2)$. An explicitly written-out instantiation of the RKL2 scheme is available for the $s=3$ and $s=4$ cases in [36], and we will not repeat it here. Notice from Eq. (17) that the minimum storage required in computer memory is four times the size of Y_0, though a comfortable implementation may use six times the size of Y_0.

2.4. Integrating these methods with a hydrodynamic solver

RKL methods can be used in conjunction with a pre-existing hydrodynamic solver to treat non-ideal parabolic terms such as viscosity and thermal conduction. In order to incorporate the methods together, we suggest combining the terms using Strang splitting [47] when second-order accuracy is desired. In this way, one would take a half-length step using RKL, then a full length timestep using the hyperbolic hydrodynamic solver, and finally a second half-length step using RKL. This constitutes a full update of one time-step. We choose the number of super-time-stepping stages we use by selecting the smallest odd value of s such that $\tau_{\max} > \Delta t_{hyp}/2$ for the method. For RKL1, this is

$$s = \left\lceil \frac{1}{2}\left(\sqrt{1+8\frac{\tau}{\Delta t_{parab}}} - 1\right)\right\rceil \tag{20}$$

and for RKL2

$$s = \left\lceil \frac{1}{2}\left(\sqrt{9+16\frac{\tau}{\Delta t_{parab}}} - 1\right)\right\rceil \tag{21}$$

where τ is the length of step we intend to take, i.e. $\Delta t_{hyp}/2$. Further implementation-related details have been provided in [36].

3. Demonstration of monotone stability

Although linear stability of a numerical method is a necessary condition, we have found that requiring a method to be monotone is very useful, as many physical systems have a monotone property. A numerical method is considered monotone if any value of the solution at the beginning of the time-step is increased, the solution at the end of the time-step can nowhere be decreased [31]. Here, we will explore the basic monotone stability of stabilized Runge–Kutta methods with "s" interior stages. The integer "s" can be comparable to the order of the RK scheme, or it can be much larger than the order of the scheme in order to promote greater time-step stability. Monotone stability is generally a more stringent requirement than linear stability. For the generic evolutionary PDE

$$\frac{du}{dt} = L(u), \tag{22}$$

where $L(u)$ contains spatial derivatives, we denote the numerical method that takes the solution from t to $t+\tau$ in zone "i" as

$$u_i^{t+\tau} = H\left(u^t; i\right). \tag{23}$$

Here, H can depend on a wide stencil of initial values of u^t. As the number of internal stages in an RK method is increased, the stencil becomes wider. A method is monotone if it can be shown that

$$\frac{\partial}{\partial u_i^t} H\left(u^t; j\right) \geqslant 0. \tag{24}$$

To make this concrete, we will explore the properties of RKL and RKC schemes as they relate to the canonical diffusion equation:

$$\frac{\partial u}{\partial t} - \alpha \frac{\partial^2 u}{\partial x^2} = 0. \tag{25}$$

Where we assume standard second-order central difference operator on a uniform spatial grid:

$$\frac{\partial^2 u}{\partial x^2} = \frac{u_{i+1} - 2u_i + u_{i-1}}{\Delta x^2}. \tag{26}$$

In the next few subsections we examine the role of monotone stability for various RK-based time-integration schemes.

3.1. The general 2-stage RK1 scheme

We introduce the methods with a simple example. The general 2-stage RK1 scheme is essentially comprised of two forward Euler explicit time steps, but each of these substeps may integrate a different portion of the overall time-step τ. With Δt_1 and Δt_2 denoting the time-steps of the first and second stages of a generic first-order 2-stage RK scheme, we obtain the following discrete updates:

$$\begin{aligned} \frac{u_i^1 - u_i^0}{\Delta t_1} &= \alpha \frac{u_{i+1}^0 - 2u_i^0 + u_{i-1}^0}{\Delta x^2} \\ \frac{u_i^2 - u_i^1}{\Delta t_2} &= \alpha \frac{u_{i+1}^1 - 2u_i^1 + u_{i-1}^1}{\Delta x^2}. \end{aligned} \tag{27}$$

Substituting the results from the first stage into the second stage yields the overall step to be:

$$u_i^2 = \delta_1\delta_2 u_{i-2}^0 + (\delta_1 + \delta_2 - 4\delta_1\delta_2)u_{i-1}^0 + (1 - 2\delta_1 - 2\delta_2 + 6\delta_1\delta_2)u_i^0 + (\delta_1 + \delta_2 - 4\delta_1\delta_2)u_{i+1}^0 + \delta_1\delta_2 u_{i+2}^0 \tag{28}$$

where $\delta_1 = \Delta t_1 \alpha/\Delta x^2$ and $\delta_2 = \Delta t_2 \alpha/\Delta x^2$ are scaled sub-steps. Note that the solution after 2 stages depends on five values of u at the initial time, and that this number will increase as the number of internal RK stages increases. To ensure monotone behavior of the solution, we simply need to make certain that all the coefficients in the above equation are non-negative. Notice that the coefficients are symmetric about u_i^0, due to the fact that the spatial operator is symmetric. Also note that the $\delta_1\delta_2$ term is non-negative. Thus there are in fact only two unique coefficients to check, namely $\delta_1 + \delta_2 - 4\delta_1\delta_2$ and $1 - 2\delta_1 - 2\delta_2 + 6\delta_1\delta_2$. By examining the minimum of these coefficients at any given pair of δ_1 and δ_2, we can map out the regions of monotone stability. This is shown in Fig. 3. Note that where $\delta_1 = \delta_2 = 1/2$, this corresponds to two equally sized forward Euler time-steps where we have the usual explicit time-step restriction $\Delta t = \frac{\Delta x^2}{2\alpha}$ for each step. But, importantly, there is a region where $\delta_1 \neq \delta_2$ that extends out significantly further, allowing for an increased overall time-step. In particular, the maximum value of $\delta_1 + \delta_2$ occurs at $\delta_1 = 1 - \frac{1}{\sqrt{2}}$, $\delta_2 = 1 + \frac{1}{\sqrt{2}}$. This particular scheme corresponds precisely to the undamped RKC1 scheme and yields a factor of two increase in overall time integration as compared to uniform stages of forward Euler. (Note that values of δ_1 and δ_2 can be reversed, but in doing so the first substep would be non-monotone.) It has been shown for undamped RKC1 that the overall timestep is increased by a factor equal to the number of stages of the RK, thus building an advantage for each additional stage utilized. Recall that the RKC1 scheme is derived from the recursion relation associated with shifted Chebyshev polynomials. As a result, the nature of the RK scheme, and the recursion relation used to derive it, set δ_1 and δ_2. In Section 3.2, we show that the RKL1 scheme, which is based on the recursion relations of shifted Legendre polynomials, yields a more salutary scheme for a broader class of non-constant coefficient diffusion problems.

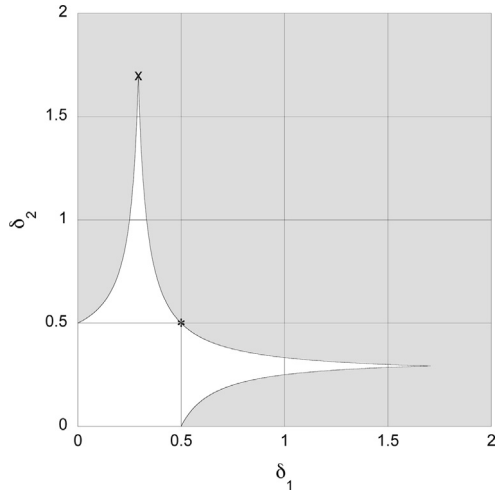

Fig. 3. The region of monotone stability for a 2-stage first-order Runge–Kutta method is the light area. The gray region indicates non-monotone behavior. The $*$ in the figure represents two equal forward Euler timesteps. The $\times$ indicates the undamped 2-stage RKC1 scheme.

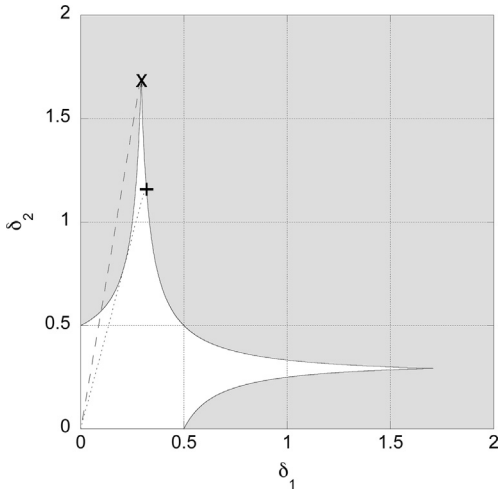

Fig. 4. Shows the lack of convex monotone property for the RKC1 scheme as indicated by the long dashed line traversing the gray region. The $+$ in the figure represents the 2-stage RKL1 scheme. The RKL1 scheme has the convex monotone property as indicated by the dotted line not crossing into the gray region.

3.2. The Convex Monotone Property (CMP)

In several physical problems, the diffusion coefficient can be spatially-varying or solution-dependent. In all such problems the largest value of the diffusion coefficient controls the overall time-step. However, the diffusion coefficient could take all values between zero and that maximal value. The time integration scheme must remain stable for that entire convex set of values.

The previous monotonicity analysis for 2-stage RK demonstrated that undamped RKL1 is the most efficient monotone scheme for the constant coefficient diffusion equation. However, to accommodate variable diffusion coefficients, we have devised an additional restriction, namely the convex monotone property (CMP). CMP ensures that the solution will not only be monotone for a single value of α, but will also maintain the monotonicity property for any diffusion coefficient between zero and α. For the 2-stage RK scheme, this is equivalent to having a clear linear path in (δ_1, δ_2) space to the origin that does not cross into the non-monotone region. The maximal value of $\delta_1 + \delta_2$ which satisfies this additional constraint is shown in Fig. 4. This occurs at $\delta_1 = \frac{1}{4}(3 - \sqrt{3})$, $\delta_2 = \frac{1}{4}(3 + \sqrt{3})$. This is precisely the 2-stage RKL1 scheme.

This graphical analysis from the previous subsection can be extended to second-order RK schemes and higher number of stages, but visualizing these is of course difficult for stages larger than 3. But, one can construct the difference scheme rather easily with symbolic computation, such as Mathematica. Then, one can construct the various polynomials consisting of δ_i, and check whether they satisfy the convex monotonicity constraints. Figs. 5(a) and 5(b) show the results from the previous example of 2-stage RKC1 and RKL1 polynomials respectively. The solid line shows the coefficient $1 - 2\delta_1 - 2\delta_2 + 6\delta_1\delta_2$, the

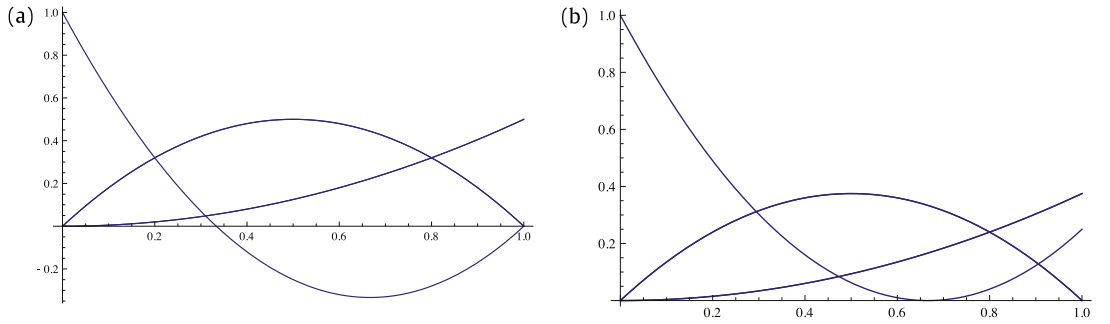

Fig. 5. (a) Shows the polynomials for a 2-stage undamped RKC1 scheme. Note the large range in which one of the polynomials is negative, indicating non-monotone behavior. (b) Shows that the 2-stage RKL1 scheme satisfies the convex monotone property.

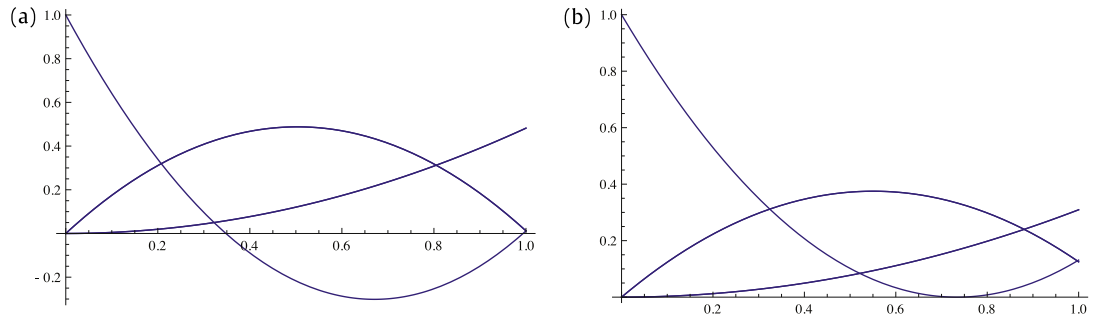

Fig. 6. (a) Shows the polynomials for a 2-stage damped RKC1 scheme, with $\varepsilon = 0.05$. Note the large range in which one of the polynomials is negative, indicating non-monotonic behavior. (b) Shows that the 2-stage damped RKC1 scheme with $\varepsilon = 2(\sqrt{6} - 2) \approx 0.898979$ satisfies the convex monotone property.

dashed line shows the coefficient $\delta_1 + \delta_2 - 4\delta_1\delta_2$ and the dotted line shows the coefficient $\delta_1\delta_2$ from Eq. (28). We see from Fig. 5(a) that for δ_1 and δ_2 along the dashed line in Fig. 4, one of the three coefficients becomes negative, showing that the undamped RKC1 scheme is not CMP. Fig. 5(b) shows that for δ_1 and δ_2 along the dotted line in Fig. 4, all the three coefficients are positive. That positivity is determined by the ratio of δ_1 and δ_2 which in turn is determined by the recursion relation of the shifted Chebyshev and Legendre polynomials.

Damping the RKC1 scheme helps to some degree, but one needs a rather large damping coefficient to maintain the CMP. Fig. 6 shows two variants of damping for the 2-stage RKC1 monotonicity polynomials. For $\varepsilon = 2/13$, we find that damped RKC1 still does not maintain the CMP. Not until $\varepsilon = 2(\sqrt{6} - 2) \approx 0.898979$ does the RKC1 scheme obtain CMP for $s = 2$ stages.

For a two stage RK scheme, the CMP is easy to trace out by hand and visualize on a two-dimensional plot. Fig. 5(b) has shown that the CMP for a two-stage RKL1 scheme is demonstrated by the positivity of three coefficients as the ratio of the sub-steps is kept constant and allowed to move on a dotted line in the two-dimensional plane shown in Fig. 4. As the number of stages increases, the CMP property becomes harder to verify. To take the example of the seven stage RKC1 and RKL1 schemes, one has to move along analogous straight lines in a seven-dimensional space that is determined by the seven sub-steps. The choice of geodesic line is determined by the recursion relation of the generating polynomial. For these seven stage RK schemes, CMP is ensured by the positivity of eight coefficients. Fig. 7(a) shows those eight coefficients for the 7-stage undamped RKC1 scheme, which is not CMP. Fig. 7(b) shows those eight coefficients for the damped RKC1 scheme with $\varepsilon = 2.000817$ that is required to maintain CMP. Fig. 7(c) shows the eight coefficients for the 7-stage RKL1 scheme that also displays CMP. The discussion in this paragraph has shown that it is very easy to post-facto verify the CMP property for a given super-time-stepping scheme by using a symbolic manipulation package. A general set of Mathematica scripts that demonstrate the CMP property for the RKC1, RKL1, RKC2 and RKL2 schemes has been provided in the supplementary sections of this paper. They can be run within a reasonable amount of time for lower values of "s" and we have verified the assertions made in this paper for all values "$s < 64$". While this is not a complete proof that the RKL schemes will all possess the CMP, we hope mathematicians can generalize the apparent CMP for RKL schemes.

The comparison between RKC1 and RKL1 schemes has been a very easy way to introduce the RKL family of super-time-stepping schemes. In practical problems, second-order accuracy in time is more desirable than first-order accuracy. In Fig. 8 we show the analogous comparison between an undamped RKC2 scheme, a damped RKC2 scheme with $\varepsilon = 2.000817$ and the RKL2 scheme. As in the previous paragraph, we have used 7-stage schemes. We verify once again that the undamped RKC2 scheme is not CMP, the damped RKC2 scheme requires substantial damping to become CMP and the RKL2 scheme is CMP. This suggests that second-order schemes share the same trends with respect to CMP as their first-order parents.

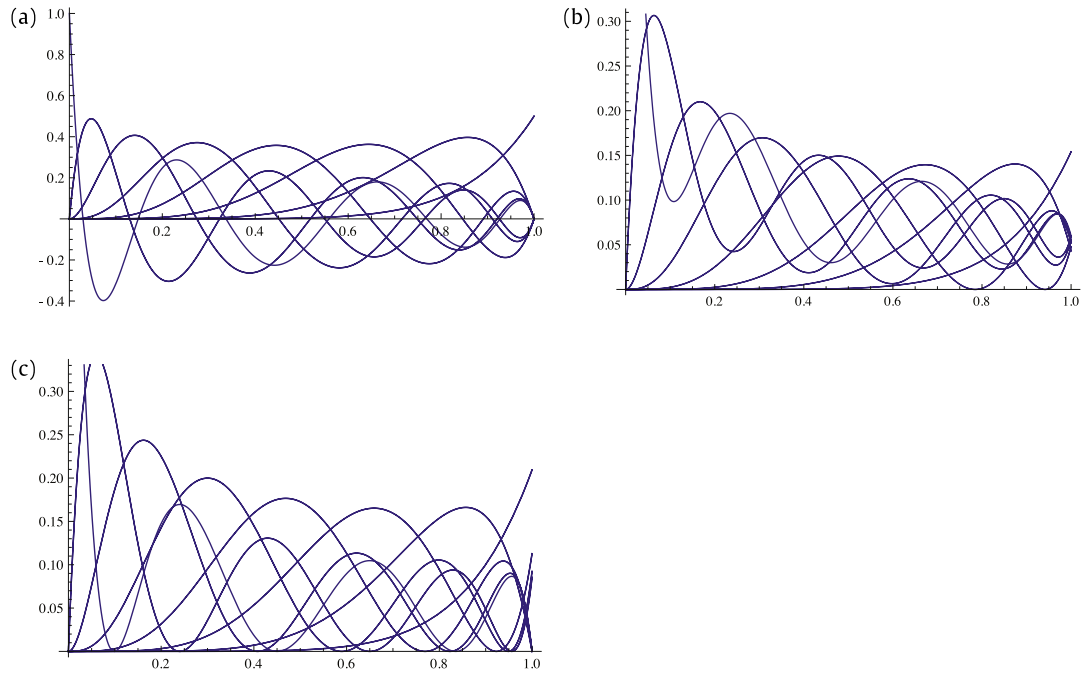

Fig. 7. (a) Shows the polynomials for a 7-stage undamped RKC1 scheme, the negative region indicates a lack of CMP. (b) Shows the polynomials for a 7-stage damped RKC1 scheme, with $\varepsilon = 2.000817$, to maintain CMP. (c) Shows that the 7-stage RKL1 scheme is also CMP.

There are two reasons why one would not want to use damping of RKC as a means by which to recover the convex monotonicity property. First, as the number of RK stages vary, so will the requisite damping parameter, and second, as the number of stages and thus damping parameter increase, the sufficiently damped RKC schemes that have the convex monotonicity property do not maintain $\tau \propto s^2$, making them less efficient than RKL schemes. The damped RKC schemes that maintain convex monotonicity appear to have $\tau \propto s^{3/2}$ as s becomes large. The RKL schemes maintain $\tau \propto s^2$ while satisfying the convex monotonicity property without any damping parameter.

4. Von Neumann stability analysis

4.1. One-dimensional stability analysis

In addition to the monotonicity analysis presented in Section 3, we consider in this section a full von Neumann stability analysis of the RKL schemes. We begin by expanding the discrete solution in terms of Fourier modes at time t and at time $t + \tau$.

$$u(x_i; t) = U_k^t e^{ikx_i}; \qquad u(x_i; t+\tau) = U_k^{t+\tau} e^{ikx_i}. \tag{29}$$

We apply the full s-stage update using the spatial discretization of the diffusion equation given in (26) as the operator. From this, we can derive the amplification factor $U_k^{t+\tau}/U_k^t$. In [36], we showed that for even values of s, the RKL2 scheme does not adequately damp the smallest wavelength modes on the mesh, and so we consider only odd values of s here.

It is very instructive to compare the RKL and RKC schemes with a second-order accurate semi-implicit method for some large, odd integer value of s. The TR-BDF2 method of Bank et al. (1985) [13] and Tyson et al. (2000) [48] is unconditionally stable for all time-steps and overcomes many of the deficiencies of the Crank–Nicholson scheme. In Fig. 9, we compare the amplification factors of the RKL and RKC schemes with the TR-BDF2 scheme. Each of the four plots shows the amplification factor for the respective scheme as a solid line, overlayed with the TR-BDF2 scheme as a dotted line and the analytic amplification factor as a dashed line. Figs. 9(a)–9(d) show the $s = 11$ RKC1, RKC2, RKL1 and RKL2 schemes, respectively. The first thing to note is that the RKC schemes in Figs. 9(a) and 9(b) have intermediate modes which are only slightly damped. In the RKL schemes in Figs. 9(c) and 9(d), those modes are all damped considerably. Fig. 9(c) shows that the smallest wavelength modes are not efficiently damped for the RKL1 scheme at its maximal time-step, which helps us understand the slight staircasing seen in Fig. 2(c). All the small wavelength modes of the RKL2 are effectively damped, as shown in Fig. 9(d). As a result, the scheme is free of staircasing as seen in Fig. 2(d). Thus, the staircasing seen in the RKC schemes can explained by the fact that some of the shorter wavelength modes are only slightly damped.

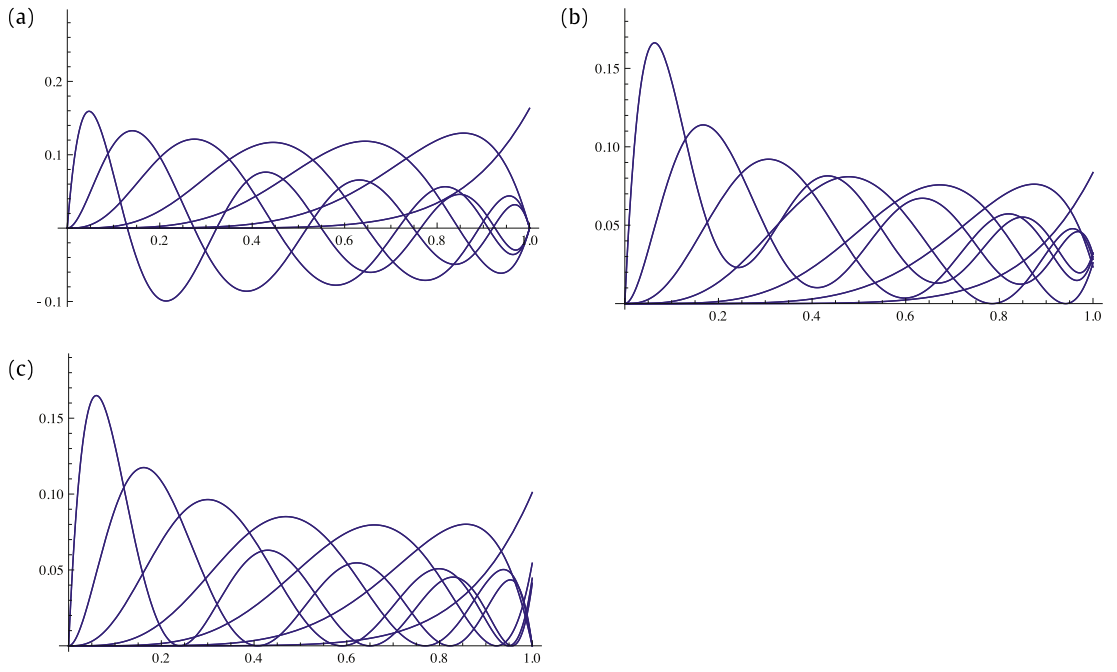

Fig. 8. (a) Shows the polynomials for a 7-stage undamped RKC2 scheme, the negative region indicates a lack of CMP. (b) Shows the polynomials for a 7-stage damped RKC2 scheme, with $\varepsilon = 2.000817$, to maintain CMP. (c) Shows that the 7-stage RKL2 scheme is also CMP.

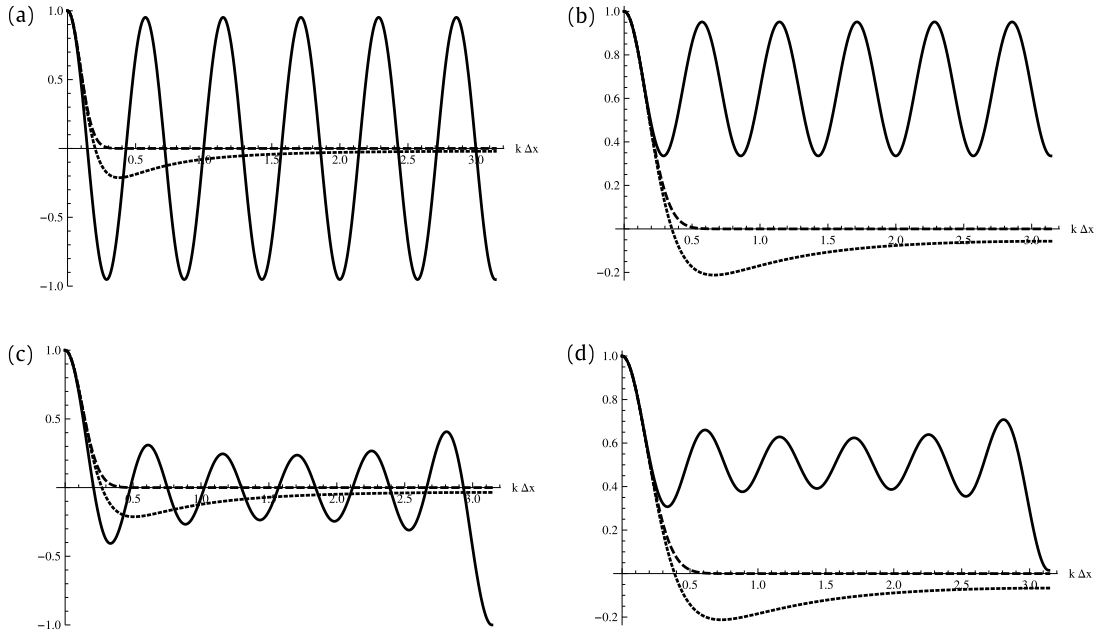

Fig. 9. (a)–(d) Shows the amplification factor for the $s = 11$ RKC1, RKC2, RKL1 and RKL2 schemes, respectively, as a solid line. The amplification factor for the TR-BDF2 scheme is overlaid as a dotted line and the analytic result as a dashed line.

4.2. Two-dimensional stability analysis

We also consider von Neumann stability for the two-dimensional diffusion equation. Using symbolic manipulation software on a two-dimensional Cartesian mesh, we were able to represent s-stage RKC and RKL schemes explicitly. Then, using a two-dimensional version of Eq. (29), we were able to obtain the amplification factors explicitly for values of $s \leqslant 11$. Figs. 10(a)–10(d) show a plot of the amplification factors of the $s = 11$ RKC1, RKC2, RKL1 and RKL2 schemes, respectively, as a function of $k_x \Delta x$ and $k_y \Delta y$. The plots are scaled from 1 to -1, with the white dashed lines representing the zero crossing. The first-order methods in Figs. 10(a) and 10(c) both have negative amplification factors for some values of wave-number,

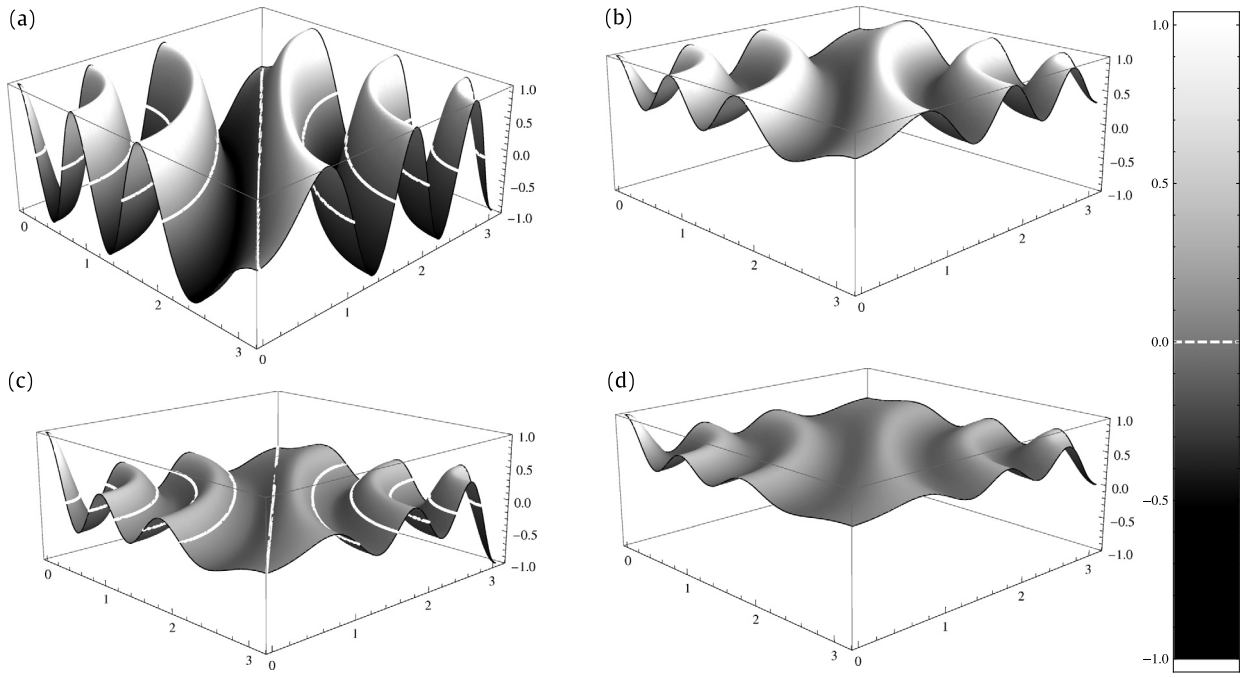

Fig. 10. Shows the amplification factor of the 2D diffusion equation for $s = 11$ stages. Image (a) shows this for RKC1, image (b) for RKC2, image (c) for RKL1 and image (d) for RKL2. The white dashed line marks the value of 0.

but the RKL1 method in Fig. 10(c) has amplification factors of a smaller magnitude than the RKC1 method in Fig. 10(a). Likewise, the amplification factor for the RKL2 method stays relatively close to 0.5 whereas for RKC2, they approach 1 at various points. This is in agreement with the one-dimensional analysis and shows that the RKL methods are also well suited for multi-dimensional problems.

5. Dealing with non-linear operators: stability in the complex plane

In this section, we loosen the restrictions we placed on the operator in Section 2, and consider operators which are non-linear and/or skew-symmetric. Recall the stability criterion given in Section 2: $|R_s(z)| \leqslant 1$. For parabolic equations with a constant coefficient, the only stability one needs is the one that pertains to negative definite values of λ, i.e. only the negative real axis of the complex plane matters. (Recall that λ is the eigenvalue of the parabolic operator.) However, parabolic operators with non-linearities or asymmetries can develop pulses that are effectively advected on a mesh [37]. In such situations the parabolic operator might develop a small advective part. This is easy to see explicitly for the example equation $\partial_t T = -\partial_x(\eta(T)\partial_x T)$. In this example, the heat conduction coefficient becomes temperature-dependent so that $\eta(T)$ has spatial variation in response to the spatial variation in T. In this case, one might want to retain some stability along the imaginary part of the complex plane. The form of the amplification factor remains the same, but z can now range over the complex plane instead of being restricted to the negative real axis.

To retain stability, the overall amplitude of the amplification factor should be bounded by unity in the complex plane. We, therefore, identify our domain of stability with the portion of the complex plane over which the amplitude of the amplification factor is less than unity. This is the reason that the RKC methods are based on damped Chebyshev polynomials. Undamped Chebyshev polynomials have points where the value is exactly ± 1, and at these points no asymmetry would be permitted because any foray into the complex plane would result in an unstable scheme. The Legendre polynomials do not have this deficiency; they only approach ± 1 at their boundaries. Therefore, we can safely leave the Legendre polynomials undamped and still retain stability in the complex plane. This domain of stability in the complex is plane is shown by the shaded domain in Fig. 11(a) for the $s = 5$ RKL2 scheme and in Fig. 11(b) for the $s = 5$ damped RKC2 scheme with $\varepsilon = 2/13$ as suggested by [52]. The dashed line in these two plots represents the limit placed on the time-step in each method. (It turns out that for odd values of s, the RKL2 and RKC2 schemes have a slightly larger stability limit than identified by the definition of the method. However, this is based on going outside the formal domain of the Legendre and Chebyshev polynomials, so the larger stability limit is never used in practice.) Such a study also enables us to display the advantages of RKL2 compared to RKC2. We see that the shaded region in Fig. 11(b) extends a little further along the negative real axis, showing that if the operator is purely parabolic then RKC2 will allow us to take slightly larger time steps. However, if the operator is non-linear and develops an advective character, i.e. a hyperbolic part, then the domain of stability for the RKL2 scheme covers a wider swath in the imaginary direction than the domain of stability for the RKC2 scheme. Thus RKL2

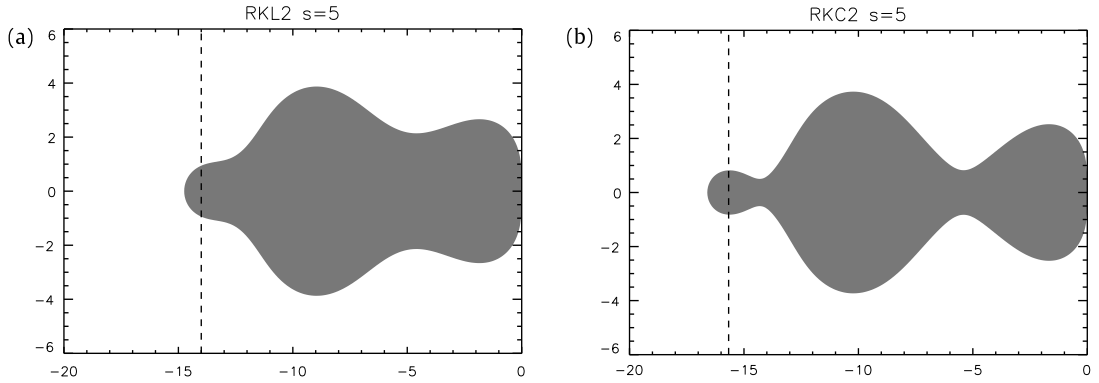

Fig. 11. (a) Shows the domain of stability in the complex plane for the $s = 5$ RKL2 scheme. (b) Shows the same for the damped $s = 5$ RKC2 scheme. The dotted line represents the maximum stable time step according to the respective method.

is much more robust than RKC2 when non-linear operators are used. The RKC2 scheme we showed in Fig. 11(b) uses the damping coefficient that was suggested by [25]. The RKL2 scheme does not need a damping coefficient, and none was used in this paper. These results hold equally well for the first-order methods, due to the nature of the underlying polynomials which make-up the methods. For a recent treatment of skew-symmetric operators within the context of RKC schemes, please see [21].

6. Accuracy analysis

6.1. One-dimensional heat conduction in aluminum–copper

In this sub-section, we investigate the properties of the RKC and RKL schemes using the example problem of heat conduction between dissimilar materials, as introduced in Section 2.1. Only the conservation of energy equation needs to be solved, as it is assumed that there is no motion in the materials, so mass and momentum are satisfied. Namely, we have:

$$\frac{\partial \varepsilon}{\partial t} + \frac{\partial}{\partial x}\left(\kappa \frac{\partial T}{\partial x}\right) = 0$$

where ε is energy, κ is the thermal conductivity and T is the temperature. For constant specific heats and density within each material, we have $\varepsilon = \rho c_p T$. This equation is solved on a uniform finite volume mesh with each of the RK schemes introduced earlier. The spatial scheme is conservative, so fluxes of $\kappa \frac{\partial T}{\partial x}$ need to be approximated at cell interfaces. The thermal gradient at the $i + 1/2$ interface is approximated by standard second-order central differences:

$$\left.\frac{\partial T}{\partial x}\right|_{i+1/2} = \frac{T_{i+1} - T_i}{\Delta x}.$$

For problems where material properties change across cells, we need to have an effective thermal conductivity at the interface. A commonly used approach (Patankar 1980 [41]) is to take the effective thermal conductivity to be the harmonic mean of the conductivities of the adjacent cells:

$$\kappa|_{i+1/2} = 2\frac{\kappa_i \kappa_{i+1}}{\kappa_i + \kappa_{i+1}}.$$

The spatial approximation is then:

$$\left.\frac{\partial}{\partial x}\left(\kappa \frac{\partial T}{\partial x}\right)\right|_i = \frac{F_{i+1/2} - F_{i-1/2}}{\Delta x},$$

where

$$F_{i+1/2} = 2\frac{\kappa_i \kappa_{i+1}}{\kappa_i + \kappa_{i+1}} \frac{T_{i+1} - T_i}{\Delta x}$$

and similarly for $F_{i-1/2}$.

Here, we are interested in the time dependent solution to the heat conduction problem, where a 10 cm aluminum rod at 0 °C is brought into contact with a 10 cm copper rod at 100 °C. The relevant material properties of several metals are given in Table 1 (Incropera and DeWitt 1990 [26] and Weast 1984 [53]).

Table 1
Thermophysical properties of select metals.

	ρ (g/cm^3)	c_p (erg/(g K))	κ (erg/(cm s K))
Aluminum	2.702	9.03×10^6	2.37×10^7
Copper	8.933	3.85×10^6	4.01×10^7
Titanium	4.500	5.22×10^6	2.19×10^6
Carbon (graphite)	2.267	7.16×10^6	1.50×10^7

Given that the solution is self-similar, one can construct the exact solution over an infinite domain. For a single material, this results in a simple error function, for 2 materials, one will have a piecewise set of error functions that yield a consistent energy flux across the material interface. The exact solution for the temperature field is:

$$T(x,t) = \begin{cases} T_{\text{int}} + (T_{\text{int}} - T_{l0}) \operatorname{erf}(\frac{x}{2\sqrt{t\alpha_l}}) & \text{for } x < 0 \\ T_{\text{int}} + (T_{r0} - T_{\text{int}}) \operatorname{erf}(\frac{x}{2\sqrt{t\alpha_r}}) & \text{for } x > 0 \end{cases}$$

where T_{l0} and T_{r0} are the initial temperature in the left and right materials α_l and α_r are the thermal diffusivities in the left and right materials, where $\alpha = \frac{\kappa}{\rho c_p}$. The interfacial temperature is given by:

$$T_{\text{int}} = \frac{\kappa_r T_{r0} \sqrt{\alpha_l} + \kappa_l T_{l0} \sqrt{\alpha_r}}{\kappa_r \sqrt{\alpha_l} + \kappa_l \sqrt{\alpha_r}}.$$

This exact solution is drawn as the solid line in Figs. 1 and 2, and is used in assessing the convergence properties of the various RK schemes.

Unlike in the mixed parabolic/hyperbolic problems, where the maximum stable time-step for the hyperbolic equation can be used to determine the number of stages of RK, here there is only a parabolic operator. So, to maintain uniformity in time integration we chose to take 5 stage RK at the coarsest grid of 80 computational cells, and chose the number of supersteps to an integrated final time just past one second. For subsequent resolutions, we doubled the number of supersteps, and then chose the number of stages needed to obtain a final time just past one second. Table 2 shows the errors for each of the RK schemes and the associated final integration time.

Several observations can be made. First, at low resolutions to moderate resolutions, the error in the RKL schemes is significantly lower than the RKC scheme of same order. Not until very high resolution, $N_x = \mathrm{O}(10^5)$, do the RKC schemes start to achieve asymptotic numerical convergence. The RKL schemes, on the other hand, achieve asymptotic numerical convergence much earlier, $N_x = \mathrm{O}(10^3)$. As demonstrated in Fig. 1, part of the reason is that the RKC schemes can demonstrate oscillatory solutions, leading to large errors at low to moderate resolutions. Furthermore, for non-linear problems, this may lead to a lack of robustness, whereas the RKL schemes have much better monotonicity properties. At high enough resolution, all four schemes achieve their design accuracy. It is worth noting that for numerical methods which take long overall time-steps, e.g. the RKL or RKC schemes, the numerical errors will be dominated by their temporal component. This is demonstrated in Table 2 by the fact that although each of the methods uses an identical (second-order accurate) spatial operator, the errors each follow different trends. Finally, we see that each of the schemes also shows a significant reduction in overall computational time as compared to standard 1 or 2-stage explicit schemes.

6.2. Viscous and conductive shock in one dimension

As a first test problem, we confirm the design convergence rate of the RKL schemes for the problem of a steady traveling viscous and heat conducting shock wave in Argon. With the inclusion of thermal conductivity and viscosity, we have the compressible Navier–Stokes equations. Appendix A describes the non-ideal MHD equations and the Navier–Stokes equations can be derived from them by setting the magnetic field to zero. If one assumes a steady propagating shock structure, the resulting Navier–Stokes equations can be transformed from a system of partial differential equations (PDEs) to a system of ordinary differential equations (ODEs). This set of ODEs can be solved to arbitrarily high accuracy via symbolic computation, such as Mathematica. Then this solution to the ODE system can then used as initial conditions for the PDE, and the shock structure will propagate at a constant speed. The Mathematica script to generate initial conditions is provided as a supplement to this paper. An exact solution to this problem was found in [16] for the special case of Prandtl number of 0.75, however we apply the method described above for a general shock.

Specifically, we examine a Mach $\sim$1.5 shock wave in Argon, with physical parameters shown in Table 3, and quiescent and post-shock states in Table 4. Fig. 12 shows the density, pressure, velocity and temperature profiles used as initial conditions. Table 5 shows some sample points in the shock structure corresponding to Fig. 12.

Computationally, the Strang split scheme as described in Section 2.4 is employed on a uniform finite volume grid. For the parabolic terms, a RKL1 or RKL2 time integration and central spatial differences is applied. For the hyperbolic terms, a predictor–corrector time integration scheme is used with a weighted essentially non-oscillatory (WENO) scheme based which uses a linearized Riemann solver, see Balsara (2004) [9], Balsara et al. (2009 [11], 2012 [12]).

Table 2
Errors associated with Al–Cu conduction problem for all the RKC and RKL schemes at first- and second-orders of accuracy.

N_x	N_{sts}	s	Average error	Maximum error	Final time
RKC1 with $\varepsilon = 0.05$					
80	2	5	2.77	23.2	1.298
160	4	7	1.69	17.6	1.271
320	8	9	8.78×10^{-1}	12.8	1.051
640	16	13	4.27×10^{-1}	6.38	1.096
1280	32	19	1.45×10^{-1}	2.40	1.171
2560	64	25	2.37×10^{-2}	3.97×10^{-1}	1.013
5120	128	37	7.42×10^{-3}	5.07×10^{-2}	1.110
10,240	256	51	3.61×10^{-3}	2.15×10^{-2}	1.054
20,480	512	71	1.77×10^{-3}	1.07×10^{-2}	1.022
40,960	1024	101	8.92×10^{-4}	5.37×10^{-3}	1.034
RKL1					
80	3	5	3.32×10^{-1}	2.47	1.206
160	6	7	1.28×10^{-1}	7.60×10^{-1}	1.126
320	12	11	6.75×10^{-2}	3.54×10^{-1}	1.327
640	24	15	3.17×10^{-2}	1.75×10^{-1}	1.206
1280	48	21	1.54×10^{-2}	8.74×10^{-2}	1.161
2560	96	29	7.46×10^{-3}	4.36×10^{-2}	1.093
5120	192	41	3.70×10^{-3}	2.18×10^{-2}	1.082
10,240	384	57	1.81×10^{-3}	1.09×10^{-2}	1.038
20,480	768	81	9.08×10^{-4}	5.44×10^{-3}	1.043
40,960	1536	113	4.47×10^{-4}	2.72×10^{-3}	1.012
RKC2 with $\varepsilon = 2/13$					
80	5	5	1.12	12.2	1.051
160	10	7	6.01×10^{-1}	6.37	1.051
360	20	11	2.53×10^{-1}	3.56	1.038
640	40	15	7.18×10^{-2}	9.22×10^{-1}	1.226
1280	80	21	6.75×10^{-3}	1.09×10^{-1}	1.204
2560	160	29	8.79×10^{-5}	1.33×10^{-3}	1.149
5120	320	39	7.05×10^{-6}	5.00×10^{-5}	1.040
10,240	640	55	1.76×10^{-6}	1.25×10^{-5}	1.034
20,480	1280	77	4.41×10^{-7}	3.15×10^{-6}	1.014
40,960	2560	109	1.10×10^{-7}	7.86×10^{-7}	1.016
RKL2					
80	6	5	1.73×10^{-1}	2.25	1.126
160	12	7	1.03×10^{-2}	1.16×10^{-1}	1.086
320	24	11	1.48×10^{-3}	9.36×10^{-3}	1.307
640	48	15	3.69×10^{-4}	2.42×10^{-3}	1.196
1280	96	21	9.20×10^{-5}	6.12×10^{-4}	1.156
2560	192	29	2.32×10^{-5}	1.58×10^{-4}	1.091
5120	384	41	5.79×10^{-6}	3.97×10^{-5}	1.081
10,240	768	57	1.46×10^{-6}	1.02×10^{-5}	1.038
20,480	1536	81	3.64×10^{-7}	2.53×10^{-6}	1.043
40,960	3072	113	9.17×10^{-8}	6.47×10^{-7}	1.011

At time $= 6.23175$ ns, we compare the computed solution with the exactly translated initial condition. See Fig. 13 for RKL1 and RKL2 computed solution of the density field and comparison with exact solution. Tables 6 and 7 evaluate the errors in the computed solution versus resolution. Note specifically that the first- and second-order design accuracy is observed for RKL1 and RKL2 respectively. Also note that for finer resolutions, the number of stages in the STS increases and subsequently saving computational time as compared to a standard explicit scheme.

6.3. Curvature flow in two dimensions

Another application of the RKL schemes is within the field of Detonation Shock Dynamics (DSD), see Bdzil and Stewart (2007) [15] for details of the underlying theory. In DSD, an asymptotic analysis of the reactive Euler equations reduces the system to a single intrinsic evolution equation relating the local curvature, κ, of the detonation shock to the local normal propagation speed of the detonation front, D_n. As shown in [4], one can utilize level set methods to treat the topologically complex front evolution of $D_n(\kappa)$ in multiple dimensions. The level set equation for propagating surfaces under the influence of their curvature were derived in Osher and Sethian (1988) [40]. The level set equation is

$$\frac{\partial \psi}{\partial t} + D_n |\nabla \psi| = 0$$

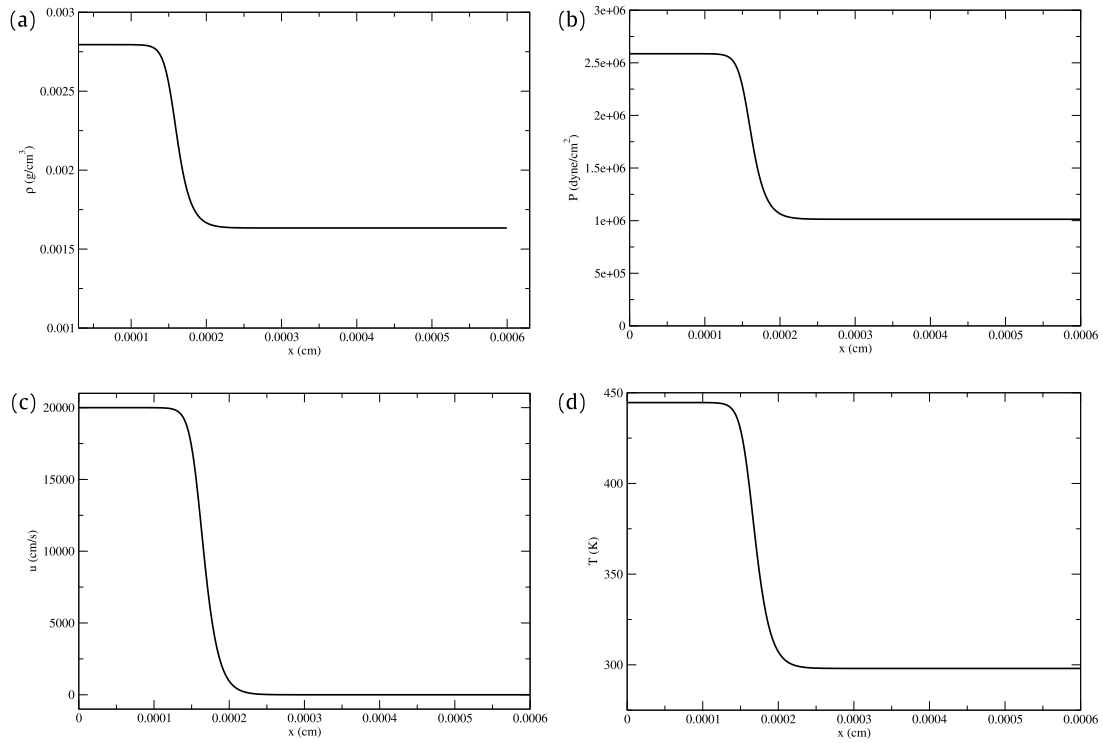

Fig. 12. Shows the density (a), pressure (b), velocity (c) and temperature (d) which serve as the initial conditions for the Argon shock problem.

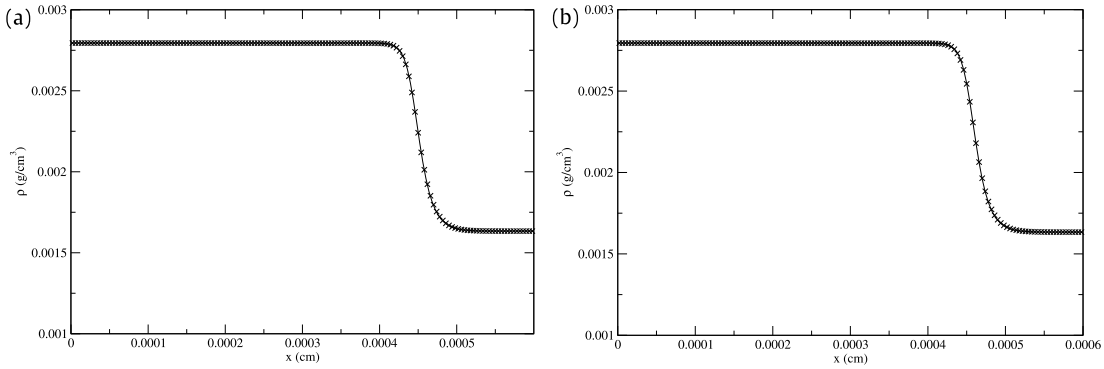

Fig. 13. (a) Shows the computed density from the RKL1 integration scheme for the Argon shock problem. (b) Shows the same for RKL2.

Table 3
Argon parameters used for Mach 1.5 shock wave.

R	2.08132×10^6 erg/(g K)
C_v	3.12181×10^6 erg/(g K)
γ	1.6667
κ	1800.82 erg/(cm s K)
μ	2.30741×10^{-4} g/(cm s)

Table 4
Quiescent and post-shock states for Mach 1.5 Argon shock.

	ρ (g/cm^3)	P (dyn/cm^2)	u (cm/s)	T (K)
Quiescent	1.63366×10^{-3}	1.01325×10^6	0	298
Post-shock	2.79472×10^{-3}	2.58615×10^6	2×10^4	444.606

Table 5
Sample initial conditions corresponding to Fig. 9.

x (cm)	ρ (g/cm^3)	P (dyn/cm^2)	u (cm/s)	T (C)
1.0125×10^{-4}	2.79324×10^{-3}	2.58433×10^{6}	1.99850×10^{4}	444.529
1.1025×10^{-4}	2.78947×10^{-3}	2.57970×10^{6}	1.99470×10^{4}	444.331
1.2075×10^{-4}	2.77213×10^{-3}	2.55836×10^{6}	1.97706×10^{4}	443.413
1.3125×10^{-4}	2.70386×10^{-3}	2.47401×10^{6}	1.90543×10^{4}	439.620
1.4025×10^{-4}	2.54079×10^{-3}	2.26997×10^{6}	1.71875×10^{4}	429.252
1.5075×10^{-4}	2.21912×10^{-3}	1.85400×10^{6}	1.27007×10^{4}	401.412
1.6125×10^{-4}	1.93207×10^{-3}	1.46064×10^{6}	7.43534×10^{3}	363.231
1.7025×10^{-4}	1.78699×10^{-3}	1.24969×10^{6}	4.13064×10^{3}	336.003
1.8075×10^{-4}	1.70178×10^{-3}	1.12039×10^{6}	1.92710×10^{3}	316.320
1.9125×10^{-4}	1.66357×10^{-3}	1.06077×10^{6}	8.65616×10^{2}	306.365
2.0025×10^{-4}	1.64839×10^{-3}	1.03675×10^{6}	4.30158×10^{2}	302.186

Table 6
Errors and convergence rates for the viscous and conductive shock in Argon: RKL1.

s	Zones	L_1 normed error	L_1 convergence	L_∞ error	L_∞ convergence
3	150	1.26×10^{-6}		7.06×10^{-6}	
3	300	3.42×10^{-7}	1.88	2.12×10^{-6}	1.74
5	600	9.86×10^{-8}	1.79	6.59×10^{-7}	1.69
7	1200	3.55×10^{-8}	1.47	2.70×10^{-7}	1.29
9	2400	1.51×10^{-8}	1.23	1.23×10^{-7}	1.13
11	4800	7.06×10^{-9}	1.10	5.86×10^{-8}	1.07
15	9600	3.46×10^{-9}	1.03	2.84×10^{-8}	1.05

Table 7
Errors and convergence rates for the viscous and conductive shock in Argon: RKL2.

s	Zones	L_1 normed error	L_1 convergence	L_∞ error	L_∞ convergence
3	150	1.51×10^{-6}		5.56×10^{-6}	
5	300	3.81×10^{-7}	1.99	1.42×10^{-6}	1.97
5	600	9.59×10^{-8}	1.99	3.58×10^{-7}	1.99
7	1200	2.40×10^{-8}	2.00	8.96×10^{-8}	2.00
9	2400	6.00×10^{-9}	2.00	2.25×10^{-8}	1.99
13	4800	1.50×10^{-9}	2.00	5.69×10^{-9}	1.98
19	9600	3.81×10^{-10}	1.98	1.50×10^{-9}	1.92

where

$$D_n(\kappa) = D_{CJ} + A(\kappa).$$

Here, D_{CJ} is a constant velocity associated with a planar detonation wave (the hyperbolic part of the level set equation). The $A(\kappa)$ term is a correction associated with a slowing of diverging waves (the parabolic part of the level set equation). The curvature is given in terms of the level set equation from

$$\kappa = \nabla \cdot \left(\frac{\nabla \psi}{|\nabla \psi|} \right).$$

Specifically for 2D, the curvature is given by

$$\kappa = \frac{(\frac{\partial \psi}{\partial y})^2 \frac{\partial^2 \psi}{\partial x^2} - 2 \frac{\partial \psi}{\partial x} \frac{\partial \psi}{\partial y} \frac{\partial^2 \psi}{\partial x \partial y} + (\frac{\partial \psi}{\partial x})^2 \frac{\partial^2 \psi}{\partial y^2}}{((\frac{\partial \psi}{\partial x})^2 + (\frac{\partial \psi}{\partial y})^2)^{3/2}}.$$

The $A(\kappa)$ term is generally non-linear, and for detonation propagation, the correction term must be bounded such that $D_n(\kappa) > 0$ (i.e. detonation shocks do not go backwards). A model $D_n(\kappa)$ which has this property is given by

$$D_n(\kappa) = \begin{cases} 1 - \frac{\alpha\kappa}{1+\alpha\kappa} & \text{for } \kappa \geqslant 0, \\ 1 - \alpha\kappa & \text{for } \kappa < 0. \end{cases}$$

For this particular non-linear $D_n(\kappa)$, one can integrate the evolution of a cylindrically expanding wave exactly. For an expanding cylindrical wave of radius r, the curvature is given by $\kappa = 1/r$, it is easy to show that the relationship between time, t, and radius, r, is given by

$$t = r + \alpha \ln r - r_0 - \alpha \ln r_0,$$

Table 8
The average error in burn time computed in the forward Euler Strang split curvature flow versus grid size. Also given is the number of times the parabolic operator, N_{par}, is evaluated per Strang split timestep.

$N_x = N_y$	N_{par}	Average error in t_b	Order of convergence
160	21	2.88×10^{-4}	
320	41	6.90×10^{-5}	2.06
640	82	1.70×10^{-5}	2.02
1280	163	4.21×10^{-6}	2.01
2560	326	1.05×10^{-6}	2.00
5120	652	2.62×10^{-7}	2.00

where r_0 is the initial radius at $t = 0$. This analytic test case will be utilized in the verification of the RKL2 scheme outlined next. To properly verify this test case, a burn table needs to be computed. The burn table, $t_b(x, y)$, is the arrival time field of the DSD shock wave. In a level set formulation, this occurs at a computational grid point when ψ changes sign. See [4] for details.

The resulting equation can be integrated via a Strang split approach. Note that both the hyperbolic and parabolic terms are quite non-linear. For the hyperbolic portion, with a Godunov based Hamilton–Jacobi solver (see Rouy and Tourin (1992) [43] for details) and second-order Heun's method for time integration in conjunction with second-order MinMod spatial interpolation, the integration requires a time-step such that in 2D:

$$\Delta t_{hyp} = CFL \frac{\min(\Delta x, \Delta y)}{\sqrt{2} D_{CJ}},$$

where *CFL* is taken as 0.9 in these tests. The explicit forward Euler parabolic time-step, with second-order central differences, is given in 2D by:

$$\Delta t_{expl} \leqslant \frac{\min(\Delta x^2, \Delta y^2)}{2\alpha}.$$

Unlike the MHD equations, the parabolic curvature term for this test case is computationally as costly as the hyperbolic operator. Thus, it is beneficial to reverse the order of the Strang split, such that a full hyperbolic operator of size Δt_{hyp} is taken, followed by a parabolic operator of size $2\Delta t_{hyp}$ and finally another hyperbolic operator of size Δt_{hyp}. In this case, for RKL2, s is chosen as the smallest odd value that satisfies:

$$s \geqslant \frac{1}{2}\left(\sqrt{9 + 32\frac{\Delta t_{hyp}}{\Delta t_{expl}}} - 1\right).$$

Next, results are catalogued for the RKL2 scheme as well as the traditional forward Euler explicit parabolic time integration. Note that even for forward Euler time integration, second-order overall convergence will take place, as in that case $\Delta t_{expl} \propto \Delta x^2$.

Specifically, we chose an initial condition that corresponds to a circle of initial radius, $r_0 = 1/10$ centered at $x = 1/9$ and $y = 1/12$. The initial level set function is set to the signed distance to this circle, namely

$$\psi = \sqrt{\left(x - \frac{1}{9}\right)^2 + \left(y - \frac{1}{12}\right)^2} - \frac{1}{10}.$$

For this test, we take $D_{CJ} = 1$ and $\alpha = 1/10$. As such, the analytic burn time is given as:

$$t_b = r - \frac{1}{10} + \alpha \ln r - \alpha \ln \frac{1}{10},$$

where $r = \sqrt{(x - \frac{1}{9})^2 + (y - \frac{1}{12})^2}$. The computational domain is given by $-1 < x < 1$ and $-1 < y < 1$. The equations are solved to $t = 1$, prior to the shock wave interacting with the computational boundary. The errors for the forward Euler scheme are given in Table 8, while the errors for the RKL2 scheme are given in Table 9. From the tables, it is clear that both algorithms yield second-order convergence, while the RKL2 scheme has roughly 30% lower error than forward Euler. Importantly, the number of parabolic evaluations is significantly reduced from the forward Euler case. At the finest resolution, the RKL2 scheme is a factor of 12.8 times faster than forward Euler.

6.4. Alfvén waves in three dimensions with resistive and ambipolar MHD

As a final demonstration of the accuracy of the RKL scheme, we present a decaying Alfvén wave in three dimensions. The equations of MHD with resistivity and ambipolar diffusion are described in Appendix A. The scheme is set-up with operator splitting as described in Section 2.4, and the resistivity and ambipolar diffusion operators are solved using the RKL2 scheme.

Table 9
The average error in burn time computed in the RKL2 Strang split curvature flow versus grid size. Also given is the number of times the parabolic operator, N_{par}, is evaluated per Strang split timestep.

$N_x = N_y$	N_{par}	Average error in t_b	Order of convergence
160	9	2.08×10^{-4}	
320	13	5.12×10^{-5}	2.02
640	19	1.28×10^{-5}	2.00
1280	27	3.14×10^{-6}	2.02
2560	37	7.83×10^{-7}	2.00
5120	51	1.96×10^{-7}	2.00

The equations for resistive and ambipolar MHD can be expanded in small perturbations about a constant state

$$\mathbf{u} = \mathbf{u}_0 + \mathbf{u}_1 e^{i(kx-\omega t)},$$

where only terms which are first-order in $\mathbf{u}_1$ are retained. For these fluids, the varying part of the conserved variables are given as

$$\mathbf{u}_1 = \{\rho_1, \rho_1 v_{x,1}, \rho_1 v_{y,1}, \rho_1 v_{z,1}, \varepsilon_1, B_{x,1}, B_{y,1}, B_{z,1}\}.$$

We select the Alfvén waves by setting the mean magnetic field and velocity values in the transverse directions to zero and choose the waves which correspond to perturbations of the y- and z-direction velocities and magnetic fields. Specifically, we take as a mean state

$$\rho_0 = 1, \qquad P_0 = 10, \qquad v_{x,0} = 1, \qquad v_{y,0} = v_{z,0} = 0, \qquad B_{x,0} = \sqrt{4\pi\rho_0}, \qquad B_{y,0} = B_{z,0} = 0.$$

For this mean state, the Alfvén wave has eigenvalues given by

$$\frac{\omega}{k} = v_{x,0} \pm \sqrt{\frac{B_{x,0}^2}{4\pi\rho} - \left(k\left(\frac{\eta}{8\pi} + \frac{B_{x,0}^2}{8\pi\Gamma\rho_{ion}\rho_{neutral}}\right)\right)^2} - ik\left(\frac{\eta}{8\pi} + \frac{B_{x,0}^2}{8\pi\Gamma\rho_{ion}\rho_{neutral}}\right).$$

We combine a y-component perturbation with a z-component perturbation 90° out of phase to obtain a circularly polarized Alfvén wave. Thus, the perturbations for the corresponding eigenvector of the torsional Alfvén wave are given by

$$v_{y,1} = \varepsilon\left(\pm\sqrt{\frac{B_{x,0}^2}{4\pi\rho} - \left(k\left(\frac{\eta}{8\pi} + \frac{B_{x,0}^2}{8\pi\Gamma\rho_{ion}\rho_{neutral}}\right)\right)^2} + ik\left(\frac{\eta}{8\pi} + \frac{B_{x,0}^2}{8\pi\Gamma\rho_{ion}\rho_{neutral}}\right)\right)$$

$$v_{z,1} = \varepsilon\left(\mp i\sqrt{\frac{B_{x,0}^2}{4\pi\rho} - \left(k\left(\frac{\eta}{8\pi} + \frac{B_{x,0}^2}{8\pi\Gamma\rho_{ion}\rho_{neutral}}\right)\right)^2} + k\left(\frac{\eta}{8\pi} + \frac{B_{x,0}^2}{8\pi\Gamma\rho_{ion}\rho_{neutral}}\right)\right)$$

$$B_{y,1} = \mp\varepsilon B_{x,0}$$

$$B_{z,1} = \pm i\varepsilon B_{x,0}$$

where $\eta = c/\sigma$ is the magnetic diffusivity, and we take a small value of $\varepsilon = 0.002$ for the amplitude of our perturbation. All of the other varying components are zero. The use of torsional Alfvén waves ensures that the magnetic pressure is uniform, at least in the ideal limit. In the linear analysis described above, there is no contribution to the energy equation, but both ambipolar diffusion and resistivity dissipate magnetic energy in the form of heat. Thus, in order to prevent pressure waves from disrupting the solution, we choose to hold the pressure fixed when computing the parabolic update in STS, ignoring the Poynting flux. This makes the update essentially isothermal, which does not affect the linear analysis.

We simulate these Alfvén waves on the unit cube with periodic boundary conditions. The direction of propagation is rotated to be along the diagonal of the cube, and the wavelength is set to be $\lambda = 2\pi/k = 1/\sqrt{3}$. In this way, the wave will be fully periodic in all directions. Specifically, we define a rotation matrix and coordinate transform between the numerical coordinates and the coordinates along the direction of the wave

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} \cos(\psi)\cos(\phi) - \cos(\theta)\sin(\psi)\sin(\phi) & \cos(\psi)\sin(\phi) + \cos(\theta)\sin(\psi)\cos(\phi) & \sin(\theta)\sin(\psi) \\ -\sin(\psi)\cos(\phi) - \cos(\theta)\cos(\psi)\sin(\phi) & -\sin(\psi)\sin(\phi) + \cos(\theta)\cos(\psi)\cos(\phi) & \sin(\theta)\cos(\psi) \\ \sin(\theta)\sin(\phi) & -\sin(\theta)\cos(\phi) & \cos(\theta) \end{pmatrix} \times \begin{pmatrix} x' \\ y' \\ z' \end{pmatrix}$$

where $\phi = \pi/2$, $\theta = \sin^{-1}(-\sqrt{2/3})$ and $\psi = \sin^{-1}((\sqrt{2}-\sqrt{6})/4)$, such that the x' direction points along the diagonal of the unprimed coordinate system. In order to initialize the magnetic field in a divergence-free manner, we also define a magnetic

Table 10
The L_2 errors and convergence rates for Alfvén waves in three dimensions with resistivity and ambipolar diffusion, as well as the number of stages, s, used in the RKL2 scheme when computing the solution. Six different accuracy analyses are shown. The left half of the table shows accuracy analyses for resistive simulations in the balanced case, the limiting case and the maximally diffusive case. The right half of the table shows accuracy analyses for ambipolar simulations in the same three cases. The stopping times for each of the simulations are also catalogued.

N_{zones}	s	L_2 error	L_2 order	s	L_2 error	L_2 order
Balanced cases	Resistive with $\eta = 0.5$ Stopping time = 0.866			Ambipolar diffusion with $\rho_{ion} = 10^{-5}$ Stopping time = 0.866		
24^3	3	1.17×10^{-4}		3	1.15×10^{-4}	
48^3	3	2.35×10^{-5}	2.31	3	2.33×10^{-5}	2.30
96^3	3	5.30×10^{-6}	2.15	3	5.28×10^{-6}	2.14
192^3	3	1.19×10^{-6}	2.16	3	1.19×10^{-6}	2.15
Limiting cases	Resistive with $\eta = 2.0$ Stopping time = 0.125			Ambipolar diffusion with $\rho_{ion} = 2.5 \times 10^{-6}$ Stopping time = 0.125		
24^3	3	4.27×10^{-4}		3	4.23×10^{-4}	
48^3	3	1.05×10^{-4}	2.02	3	1.04×10^{-4}	2.03
96^3	5	2.53×10^{-5}	2.06	5	2.49×10^{-5}	2.06
192^3	5	5.99×10^{-6}	2.08	5	5.90×10^{-6}	2.08
Maximally diffusive cases	Resistive with $\eta = 50.0$ Stopping time = 0.005			Ambipolar diffusion with $\rho_{ion} = 10^{-7}$ Stopping time = 0.005		
24^3	9	1.41×10^{-4}		7	9.86×10^{-5}	
48^3	13	2.53×10^{-5}	2.48	11	2.97×10^{-5}	1.73
96^3	17	4.88×10^{-6}	2.38	17	6.21×10^{-6}	2.26
192^3	25	3.19×10^{-7}	3.93	25	3.64×10^{-7}	4.09

vector potential

$$A'_x = 0, \qquad A'_y = \frac{1}{k}\left(\varepsilon B_{x,0} \cos\left(kx'\right)\right), \qquad A'_z = B_{x,0} y' + \frac{1}{k}\left(\varepsilon B_{x,0} \sin\left(kx'\right)\right).$$

For more information on setting up torsional Alfvén waves, please see [9] or [12].

For all the runs, we set the ion-neutral coupling constant, $\Gamma = 2.51327 \times 10^6$. We vary the values of η and ρ_{ion} in order to examine the functioning of the RKL2 scheme on these operators in different regimes. We examine 3 cases each for ambipolar diffusion and resistivity, for a total of 6 cases. In the first case, which we call the "balanced" case, we choose $\eta = 0.5$ or $\rho_{ion} = 10^{-5}$, such that the decay timescale of the wave is similar to the Alfvén crossing time. The stopping time is explicitly given in the accuracy analysis tables that follow. In the second case, we choose $\eta = 2.0$ or $\rho_{ion} = 2.5 \times 10^{-6}$, which is close to the limit for a propagating Alfvén wave, i.e. the real part of the eigenvalue approaches the background advection speed. We call this the "limiting" case. Finally, in the last case we choose $\eta = 50.0$ or $\rho_{ion} = 10^{-7}$, which is well into the regime where the Alfvén wave becomes non-propagating and the problem is dominated by the parabolic terms, so we call this the "maximally diffusive" case. Each case is run to a final time where the perturbations have decayed to about 75% of their original amplitudes. We have run each of these cases on meshes of increasing resolution from 24^3 to 192^3 zones in order to show the convergence of these schemes.

Table 10 presents the L_2 errors in the x-component of the magnetic field for each of the cases described above, as well as the maximum number of stages, s, used in the computation. In all three regimes, the RKL2 scheme performs well, showing second-order convergence for both ambipolar diffusion and resistive MHD. Even in the parabolically dominated cases, where the maximum hyperbolic time-step exceeds the explicit parabolic time-step by a factor of over 250, the scheme still fares well. This demonstrates the computational cost savings over explicit sub-cycling in time. For the 192^3 ambipolar diffusion simulation, 568 evaluations of the parabolic ambipolar diffusion operator would be required for a second-order accurate update with each hydrodynamic time-step. Taking into account Strang splitting, the RKL2 scheme achieves this with only 50 evaluations of the ambipolar diffusion operator, an 11.4× decrease in speed.

7. Results and applications

In this section, we expand on the results presented in [36]. In that paper, we showed that the RKL2 method was suitable for simulating anisotropic thermal conduction as occurs in magnetized plasmas. Here, we present three different applications which serve to demonstrate the wide-ranging applicability of the RKL methods to many different sorts of problems. In Section 7.1, we present results of a non-linear radiation-diffusion model. In Section 7.2, we present an example from detonation shock dynamics involving level sets. In Section 7.3, we present an example involving thermal conduction in a composite material. In Section 7.4, we present the Hartmann flow problem arising from viscous and resistive plasma flow.

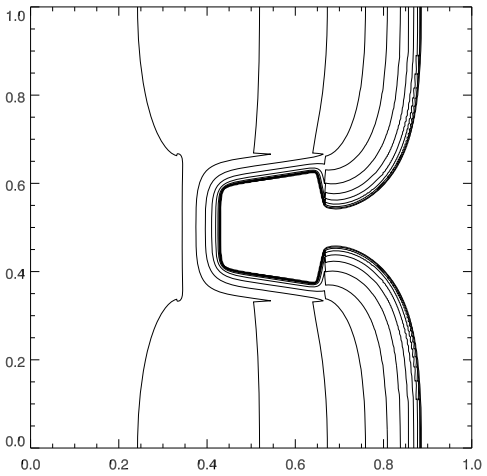

Fig. 14. A contour plot of the radiation-diffusion model problem at $t = 3$. There are 12 contours are evenly spaced between 0.1 and 1.2.

7.1. Radiation-diffusion model

The first application we present is the non-linear diffusion model presented in [37]. We follow [25] for the set-up and spatial discretization of the problem, using a mesh size of 400×400 zones and RKL2 super-time-stepping for the temporal discretization. We set the thermal conduction coefficient, $k = 0$ for this problem to maximize the steepness of the temperature gradient. Because this is a diffusion-reaction problem with no hyperbolic part, we are free to choose the number of stages and supersteps we will use in computing the solution. Using the spectral radius estimate presented in [25], we can see that if this problem were run to a final time of $t = 3.0$, nearly 1.2 million explicit, forward-Euler timesteps would be required, meaning 2.4 million evaluations of the update operator would be needed for second-order accuracy. Using RKL2, however, we can run this with only 150 supersteps with 174 stages each, for a total of 26,100 operator evaluations, a speedup of more than $90\times$.

Fig. 14 shows a contour plot of the simulation described above, with 12 evenly spaced contours from 0.1 to 1.2. We also ran the simulation with an explicit, second-order update, and the two schemes produce nearly identical results, indicating that the non-linear nature of the equations does not prevent the RKL2 scheme from significantly decreasing the computation required to get to the solution.

7.2. Surfaces propagating under curvature

In this example, we take the algorithm from Section 6.3 and apply it to detonation propagation in the "sandwich" geometry, see Hill and Aslam (2003) [23] and Aslam et al. (2009) [7] for details of this test. Essentially, a slab of high explosive is sandwiched between two slabs of inert material, and the detonation wave is observed to travel from one end to the other. The purpose of the test is to determine how well the inert material confines the detonation wave. For very heavy or stiff inert confiners, it is expected that the detonation will be confined better (and thus travel closer to its ideal speed, D_{CJ}) than if it is confined with a low-density compliant material. The theory behind this can be found in the discussions of Aslam and Bdzil (2002, 2006) [5,6], and [15].

The specific test problem stems from the experiment performed in [7]. Namely, an 8 mm thick slab of the plastic bonded explosive PBX 9502 (95% TATB and 5% Kel-F binder) was observed to detonate across a 127 mm length. Images of the propagating detonation wave were generated via proton radiography. Furthermore, shorting pins were placed along the edge of the high explosive charge to measure the velocity of the detonation wave as it swept across the slab of explosive. The $D_n(\kappa)$ curve was taken from Hill and Aslam (2010) [24], assuming a pressed density of 1.890 g/cm^3, initial temperature of 298.15 K and material lot HOL88H891-008. A plot of the $D_n(\kappa)$ curve is shown in Fig. 15. The relevant domain for this problem is $0 \text{ mm} < x < 127 \text{ mm}$ by $-4 \text{ mm} < y < 4 \text{ mm}$. At the top and bottom boundaries, a shock deflection angle of 35 degrees was imposed; see Aslam et al. (1996) [4] for discussion of boundary conditions with level sets and DSD. At the left and right boundaries, linear extrapolation was assumed. Initially, the wave is assumed to be flat, and enters the computational domain at $x = 0$. This corresponds to $\psi(x, y, t = 0) = x$. Since the wave is initially flat, it travels at the ideal D_{CJ} speed. But, as the angle boundary condition is enforced at the top and bottom of the domain, the curvature is locally increased, and as such the wave becomes curved near the boundaries. As the wave continues to propagate, it eventually becomes a steady propagating detonation. See Fig. 16 for the evolution of the shock shapes as well as the local detonation velocity attained as the wave travels over the domain.

Experimentally, a qualitatively curved wave was observed, see [7]. Also, from the experiment, over the last 80% of the explosive charge, a detonation velocity of 7.467 ± 0.029 mm/μs was observed. The DSD computation performed here yielded a velocity of 7.478 mm/μs (at the resolution of 1/32 mm, the DSD computation had already converged to 4–5 significant

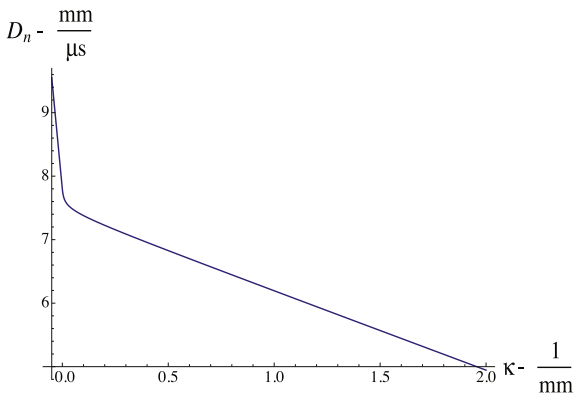

Fig. 15. The normal detonation shock speed curvature relation for the explosive PBX 9502.

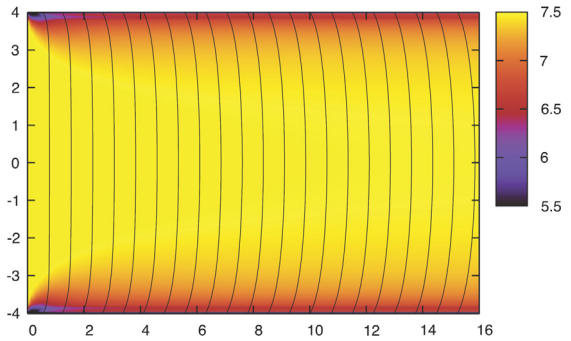

Fig. 16. Early time results of the DSD Sandwich test, with a spatial resolution of 1/32 mm. Contours indicate shock loci at time intervals of 0.1 microsec. The background color indicates the normal detonation speed as the wave passes over spatial point. Note that only a portion of the solution is shown, as the x-domain extends out to $x = 127$ mm.

figures). Importantly for the RKL2 scheme, with $s = 39$, a speedup of $10\times$ was observed over the typical forward Euler sub-cycling method.

7.3. Multi-material thermal conduction in two dimensions

There exists a plethora of physical applications where modeling multi-material thermal conduction is important. Here, we focus on one particular problem that arises in materials processing, namely the formation of the cermat (ceramic metallic) material titanium carbide, TiC. In particular, an active field of research is the self-propagating high-temperature synthesis of such materials; see Lee and Chung (1997) [30] for example. In their particular formulation, they combine small particles ($<$45 μm) of titanium and graphite in equal molar fractions along with aluminum particles in an argon atmosphere. These powders are thoroughly mixed and then compacted in a die. This mixture, if ignited at one, will support a self-propagating solid flame-like structure. Unlike traditional flames, there are no gaseous products, only $\mathrm{Ti} + \mathrm{C} \rightarrow \mathrm{TiC}$. The purpose of the aluminum is two-fold. First, aluminum is a good thermal conductor, and as such can aid in the propagation of the combustion wave. Secondly, aluminum has a relatively low melting temperature, and once melted it can serve as a solvent for the titanium and aid in the transport required to form the final product TiC. Simulating this entire process is beyond the scope of this work, but obtaining bulk thermal diffusivity of the mixture (prior to melting the Al) can be computed from modeling the constituents directly. Clearly, the material properties of the constituents are quite disparate (see Tables 1, 3 and 4), and this is where the convex monotone property of RKL can be helpful.

As an example, here we present the thermal conduction of a composite material presented in [30]. We examine the detailed thermal transport in a composite mixture of Ti, C, Al and Ar. The computational domain is 0.1 cm by 0.1 cm, filled with ∼40 μm sized particles of the constituents corresponding to volume fractions of roughly 34.2% Ti, 17.6% C, 28.2% Al and 20% Ar. See Fig. 17. This corresponds to a nearly stoichiometric mixture (1-to-1 molar ratio) of Ti and C, and is typical of the experiments performed in [30]. The procedure of filling this square domain was as follows. First, fill the box with the background Ar gas. Then replace ∼60% of the Ar with 40 μm circular Al particles placed randomly throughout the domain (they are allowed to overlap). Then, we add random 40 μm particles of Ti and C (in a 2-to-1 volumetric ratio) until we reach a porosity (volume fraction of Ar) of 20%. During the addition of the Ti and C particles, the underlying Al and Ar are displaced, resulting in the final volume fractions of 34.2% Ti, 17.6% C, 28.2% Al and 20% Ar.

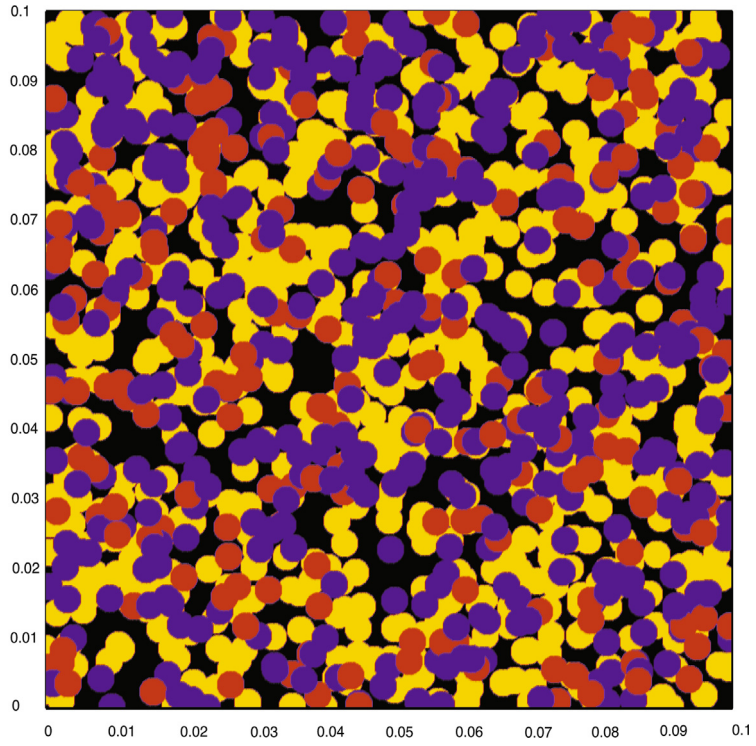

Fig. 17. Material distribution within the composite. Black indicates argon, yellow is aluminum, purple is titanium and red indicated carbon. (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

The initial conditions are relatively straightforward, namely we set the temperature, for $y < 0.05$ cm, to be 500 K, and for $y > 0.05$ cm to be 300 K. At 500 K, the Al is still in a solid state. Convection in the Ar is not an issue, and is treated as a stationary material; the dimensionless Rayleigh number for convection is well below unity. Along the $x = 0$ and $x = 0.1$ cm, we assume the walls are perfectly insulated, while the temperature along $y = 0$ and $y = 0.1$ cm are fixed at 300 K and 500 K respectively. Using 100 supersteps, with $s = 135$ yields the series of temperature fields shown in Fig. 18. This represents a speedup of $35\times$ as compared to the explicit forward Euler method. Note that the thermal profile is quite contorted during its evolution, as the connectivity of the solid materials is important for thermal diffusivity. For example, in this particular initial condition, the cold region 0.07 cm $< x <$ 0.09 cm, near $y = 0.06$ cm, is due to that particular region being completely surrounded by Ar. As such, very little energy flux is transported, and the region remains relatively cool. One can utilize such simulations to obtain bulk properties of complex composites, such as effective thermal diffusivities for continuum modeling.

7.4. Hartmann flow

In this sub-section, we describe a test problem arising from the plasma physics of viscous and conductive plasmas, which illustrates the interplay between flux freezing and diffusion. The Hartmann flow problem considers flow between two non-conducting plates moving at different velocities. The problem is described in detail in Jackson, 1975 [27].

We solve the problem on a non-uniform 2D mesh spanning $[0, 3.2]$ in the x-direction and $[0, 1]$ in the z-direction with 200×100 zones. Because the important dynamics in the problem are located near the top and bottom boundaries, we use a logarithmic grid near those boundaries. The first tenth of the physical domain is covered by 25 zones which increase in size exponentially. The middle 80% is treated by 50 zones uniformly spaced (with $\Delta z = 0.016$). Finally the last tenth of the region mirrors the first tenth to the boundary. The mesh near the lower boundary is determined by requiring the 25-th zone edge to be located at $y = 0.1$ and the 26-th zone edge to be located at $y = 0.116$. The first 25 zone boundaries are given by:

$$y_j = 0.002758291932\left(e^{0.144710717\,j} - 1\right).$$

The grid is uniform in the x-direction with $\Delta x = 0.016$. The density and magnetic field are fixed at the top and bottom boundaries. We impose an x-directional velocity of 0.5 at the top boundary and -0.5 at the bottom boundary. The viscosity in the problem ensures that these top and bottom x-velocities are persistently conveyed to the rest of the fluid so that the flow reaches steady state after a long time. We use periodic boundary conditions for the x-direction. The problem is initialized with a smooth linear velocity profile connecting the upper and lower boundaries: $v_x(z) = -0.5 + z$. Initially, the

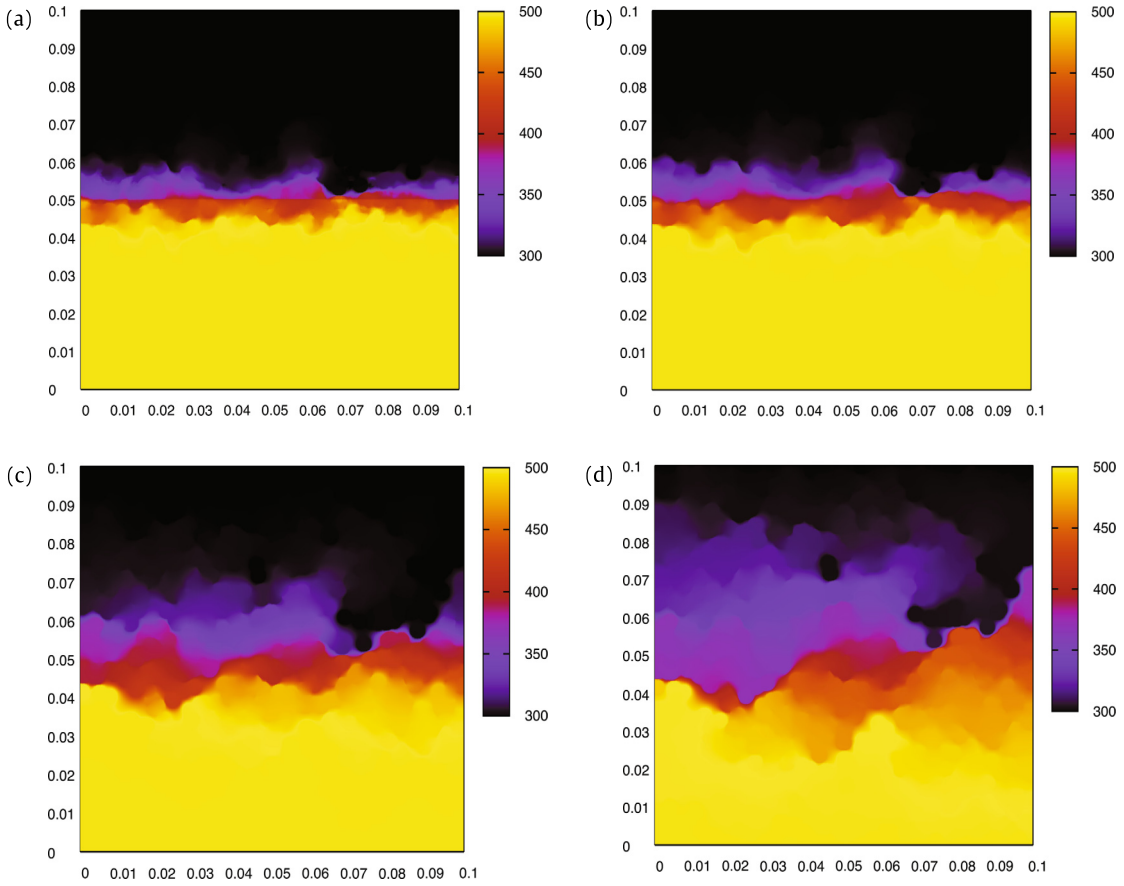

Fig. 18. Distribution of temperature for 4 different times during the evolution of thermal conduction in a composite. Times are roughly $t = 0.03$, 0.06, 0.25 and 1 millisecond.

magnetic field is uniform and constant in the z-direction. The density is also set to unity everywhere. In [27] there is an applied electric field in the y-direction, but we do not include that in present simulations.

The equations of viscous and resistive MHD are given in Appendix A. The operators for resistivity and viscosity were calculated using the RKL2 method for all the simulations in this sub-section. In [27] the problem is solved using an incompressible equation of state. However, we adopt the isothermal equation of state with an isothermal sound speed of unity. Using an isothermal set-up ensures that the viscous and Ohmic terms do not cause a runaway increase in the thermal energy. We run the problem to a final time of 20, which allows the flow to reach steady state. The steady state solution for velocity and magnetic field are given by [27]

$$v_x(z) = \frac{V_1}{\sinh(M)} \sinh\left(M\left(\frac{a-z}{a}\right)\right) + \frac{V_2}{\sinh(M)} \sinh\left(\frac{Mz}{a}\right)$$

$$B_x(z) = B_0\left(\frac{4\pi\sigma a^2}{c^2}\right)\left(\frac{V_2 - V_1}{2a}\right)\left(\frac{\cosh(\frac{M}{2}) - \cosh(\frac{M}{2} - \frac{Mz}{2})}{M\sinh(\frac{M}{2})}\right)$$

where V_1 and V_2 are the bottom and top velocities, a is the maximum extent in the z-direction of the domain σ is the electric conductivity, and M is the dimensionless Hartmann number, which is given by

$$M = \left(\frac{\sigma B_0^2 a^2}{\eta c^2}\right)^{1/2}.$$

By choosing viscosity and conductivity of $\eta = 1$ and $\sigma = 1$, the value of M becomes equal to B_0 in code units. The size of the boundary layer in this problem is approximately a/M. Therefore, as the magnetic field increases the size of the boundary layer decreases, and the region of constant magnetic field and near zero velocity becomes a greater part of the domain.

We have run this simulation with three values for the initial magnetic field B_0 given by 1, 10 and 100, in order to sample the different regimes of the problem. Fig. 19 compares the magnetic field and velocities for these simulations with the expected values given above. Figs. 19(a), 19(b) and 19(c) show the steady state x-component of the magnetic field

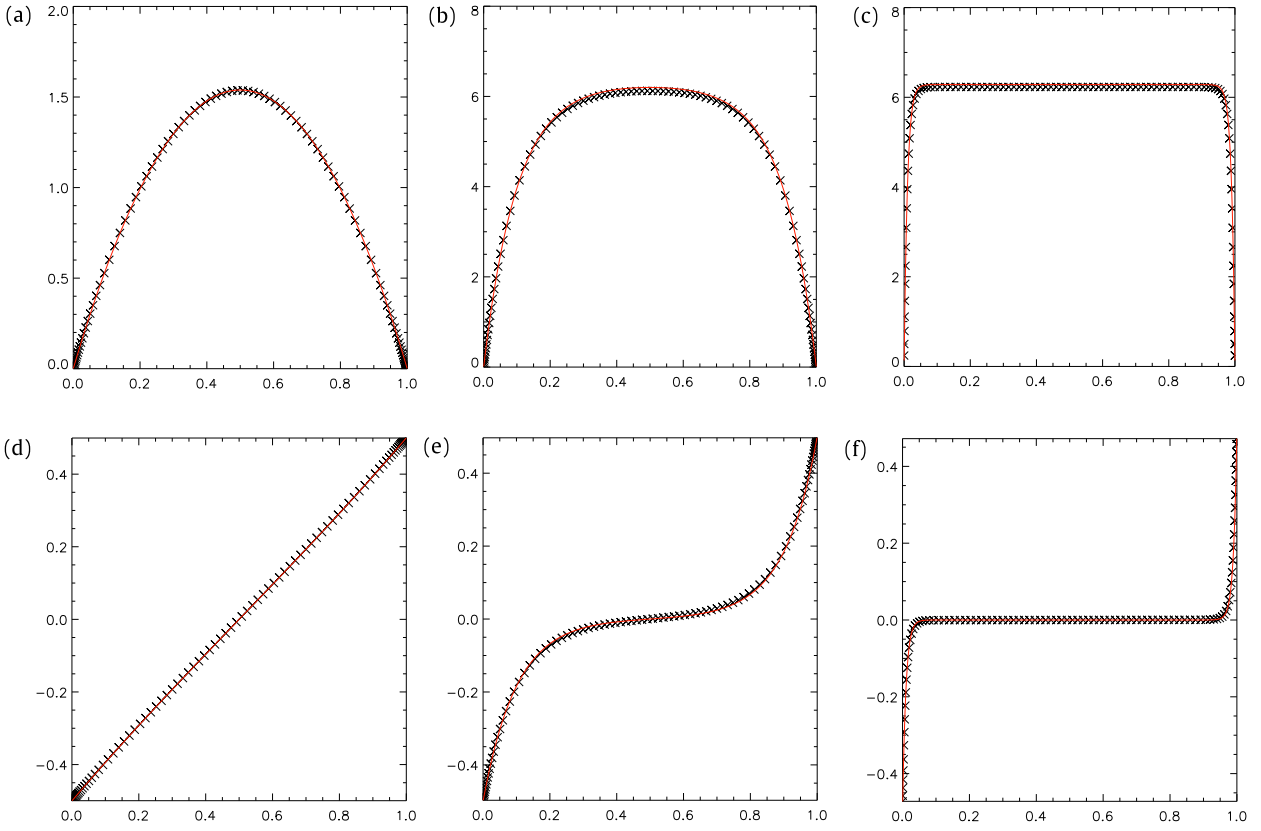

Fig. 19. Shows slices of B_x and v_x along the y-direction of the Hartmann problem when the simulations have reached steady-state. Simulation data is marked with crosses and the predicted value is given by the solid line. (a), (b) and (c) show the value of B_x for the case of $B_0 = 1$, 10 and 100 respectively. (d), (e) and (f) show value of v_x for the same.

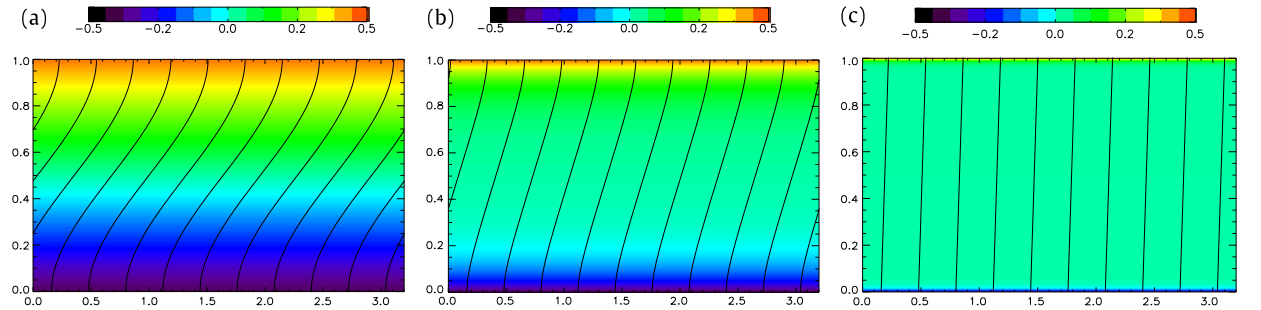

Fig. 20. Shows the full velocity profile of the Hartmann problem at steady state (background gradient) overlayed with the magnetic field lines for (a) $B_0 = 1$, (b) $B_0 = 10$, and (c) $B_0 = 100$.

overlaid with the expected profile along a vertical slice in the middle of the domain, and Figs. 19(d), 19(e) and 19(f) show the same for the x-component of the velocity. This is comparable to Figs. 10.2 and 10.3(a) in [27]. In each case presented, the numerically produced profile is shown to match the predicted profile quite closely. Fig. 20 shows the velocity overlaid with magnetic field lines. This can be compared with Fig. 10.3(b) in [27]. Fig. 20(a) shows the $B_0 = 1$ case, Fig. 20(b) shows the $B_0 = 10$ case, and Fig. 20(c) shows the $B_0 = 100$ case. The transition from uniform velocity field to the narrow boundary condition case is clearly seen.

For the $B_0 = 1$ case, the required s was found to be 121. For the $B_0 = 10$ case, the required s decreased to 81, and for the $B_0 = 100$ case, the required s was 29. This progression can be explained by realizing that the parabolic explicit time-step, Δt_{parab}, is the same for each of the cases, but the hyperbolic time-step, Δt_{hyp}, drops as the magnetic field increases, due to the larger magnetosonic speed.

8. Timing comparisons of RKL2 with other methods

In this section, we present timing comparisons between the RKL method presented in this paper and other existing methods. We compare the errors as well as the wall-clock time to solution in order to show the relative performance of the methods. Because the implicit treatment of a linear problem requires different solvers from the implicit treatment of a non-linear problem, we present two subsections below. In the first subsection, we present a problem with linear coefficients, and in the subsequent subsection, we present a non-linear problem with solution-dependent coefficients.

8.1. Timing analysis for a problem with linear coefficients

We first solve the heat equation in two dimensions, $\partial_t u = \partial_{xx} u + \partial_{yy} u$, on a Cartesian mesh. The domain of the problem is $[0, \pi]^2$ with Dirichlet boundary conditions; all boundaries were set to zero. We take as initial conditions a truncated sine series expansion of a square in the center of the domain, so that the solution for all time is given by

$$u(x,t) = \sum_{n=1}^{30}\sum_{m=1}^{30} \frac{4}{\pi^2} \frac{(\cos(m\tfrac{3}{8}\pi) - \cos(m\tfrac{5}{8}\pi))(\cos(n\tfrac{3}{8}\pi) - \cos(n\tfrac{5}{8}\pi))}{mn} e^{-(m^2+n^2)t} \sin(mx)\sin(ny).$$

Physically, these initial conditions are motivated by setting the solution to zero in the domain with the exception that we set the solution to unity in a central square of extent $[3\pi/8, 5\pi/8]^2$. The modes with $m > 30$ or $n > 30$ in the above sine series have been discarded because they decay out very rapidly in time. We find that the initial profile is sufficiently sharp to provide computational difficulty while still permitting accuracy analysis.

We compare four second-order accurate in time schemes:

(a) We use the RKL2 method presented in this paper. Because the target application of this method is to couple it with a hyperbolic part, we ran the code with $\Delta t \propto \Delta x$. Consequently, we had to increase the number of stages in the RKL2 scheme with increasing resolution. This is shown in Table 11. This choice also ensures that the temporal errors scale as the spatial errors.
(b) We use a time-explicit second-order accurate method. The timesteps in such a method clearly cannot keep up with the timesteps in the other two methods. As a result a second-order Runge–Kutta scheme based on Heun's method was used for explicit sub-cycling. We call this the RK2-Forward Euler method for want of a better name. Each timestep is inexpensive, however, as seen from Table 11, the number of timesteps increase as the square of the mesh size.
(c) We use the TR-BDF2 scheme [13,48], which represents a second-order accurate implicit method. This method requires two matrix inversions per time-step. The matrix inversions were done with pointwise multigrid using three Jacobi iterations on each upward and downward pass. Multiple multigrid V-cycles were used to reduce the error below the discretization error (but not to the level of round-off error).
(d) Because the Intel Math Kernel Library (IMKL) is a reference standard, we also ran the TR-BDF2 scheme with the direct sparse solver routines available in that library. We used the dpbsv routine for symmetric, positive-definite banded matrices. Although our multigrid solver is faster than this routine, we present this data because we realize that people would traditionally use IMKL for this problem.

We take as our stopping time, $t = 0.06777$ which corresponds to 125 time-steps on the coarsest mesh using the explicit method. We compute the TR-BDF2 inverses using a multigrid solver which we found worked much more efficiently than direct sparse matrix inversion. I.e. our multigrid solver proved to be faster than the best available sparse matrix inversion routine from the *Intel Math Kernel Library*.

Table 11 shows a comparison of the different schemes. The table shows the L_2 errors along with the wallclock computational time. First it is worth noting the level of speed-up the RKL2 scheme provides over the explicit sub-cycling. At a resolution of 1024^2, the RKL2 scheme is 23 times faster than the explicit RK scheme. This is in proportion to the total number of evaluations of the Laplacian operator taken by the two schemes. For the lower resolution cases, the multigrid TR-BDF2 is slower than even the explicit forward-Euler, but this is because the multigrid solver does not achieve its full efficiency on smaller meshes. However, it becomes more computationally efficient than the RK2-Forward-Euler method once the computational mesh has reached 512^2 zones. However, using a mesh of 1024^2 zones, RKL2 still outperforms the multigrid TR-BDF2 scheme by a factor of 5 while producing a solution with comparable numerical errors. The table shows that the numerical errors should remain comparable at higher resolutions.

Table 11 also shows that the RKL2 scheme has lower computational complexity on all the mesh sizes shown there. We also see that the computational complexity of the TR-BDF2 scheme increases more gradually with increasing resolution compared to the RKL2. We can always ask for the resolution at which both schemes will take comparable wall-clock times. Extrapolating the wall-clock time values, we would predict that the RKL2 and multigrid TR-BDF2 schemes should have the same computational cost on a mesh of approximately $30{,}000^2$ zones for this problem. We also see from Table 11 that the computational complexity of the direct solver from IMKL increases more adversely than multigrid with increasing mesh resolution. This is the expected result because multigrid scales as $N\log(N)$, whereas a direct solver scales approximately as $N^{1.3}$, where N is the number of mesh points.

Table 11
A comparison of the timing and errors of the RKL2, second-order Runge–Kutta, multigrid TR-BDF2 and direct sparse solver (IMKL) TR-BDF2 schemes. The number of time-steps (N_t) and stages (s) for STS is included as well as the L_2 error and the wall clock time.

Mesh	N_t	s	L_2 error	Time (s)	N_t	L_2 error	Time (s)
	RKL2				RK2 – forward Euler		
64^2	1	22	1.72×10^{-1}	0.002	125	5.35×10^{-4}	0.014
128^2	2	32	8.92×10^{-2}	0.022	500	1.39×10^{-4}	0.208
256^2	4	44	1.38×10^{-2}	0.358	2000	4.02×10^{-5}	3.34
512^2	8	63	2.83×10^{-3}	4.02	8000	1.42×10^{-5}	68.0
1024^2	16	88	8.77×10^{-5}	48.0	32,000	6.15×10^{-6}	1130
	TR-BDF2 multigrid				TR-BDF2 IMKL		
64^2	1		6.03×10^{-2}	0.063	1	6.04×10^{-2}	0.051
128^2	2		1.02×10^{-2}	0.546	2	1.02×10^{-2}	1.10
256^2	4		1.15×10^{-3}	4.07	4	1.15×10^{-3}	26.5
512^2	8		2.71×10^{-4}	36.1	8	2.71×10^{-4}	501
1024^2	16		6.66×10^{-5}	285	16	6.72×10^{-5}	14,500

The results of this sub-section show the clear advantage provided by the RKL methodology over the already existing explicit sub-cycling and implicit methods. The RKL2 method produces the expected speed-up over explicit sub-cycling. Additionally, for less computational cost and considerably easier implementation, the RKL2 method can deliver results with numerical errors which are comparable to those produced by the TR-BDF2 scheme.

8.2. Timing analysis for a problem with non-linear coefficients

In this subsection, we solve the non-linear porous medium equation, $u_t = \nabla^2 u^m$, in two dimensions with $m = 4$. This is equivalent to the heat equation with non-linear conduction coefficients equal to mu^{m-1}. This equation has a cylindrically-symmetric solution in two dimensions first derived by Barenblatt (1952) [14] which is given by:

$$u(r,t) = t^{-1/m} \max\left(C - \frac{m-1}{4m^2} r^2 t^{-1/m}, 0\right)^{\frac{1}{m-1}}$$

where $(x)_+$ means $\max(x, 0)$. The constant C must be positive and the solution is only valid for times $t > 0$. More information about the general properties of this equation and classes of solution can be found in [51], especially Chapter 4.

We compute this problem on a square domain spanning $[-0.5, 0.5]$ with Dirichlet boundary conditions set to 0. This solution has a sharp discontinuity which propagates outward at a finite speed. As time goes to 0, the solution becomes a delta-function. Therefore, we must initialize the problem at a slightly later time. We choose the initial time and constant C such that the initial profile goes to zero at $r_0 = 0.125$ and the average value of the profile over this region is 1. This gives us a starting time, $t_0 = 81/262{,}144 \approx 3.090 \times 10^{-4}$ and $C = 128/\sqrt{2} \approx 5.52 \times 10^{-3}$. This initial profile has a maximal value of 4/3. From this, we can derive the explicit time-step restriction for this problem, $\Delta t \leqslant \Delta x^2/4(4(4/3)^3)$. We run the solution to a final time of $t = 0.1208$, by which time the region containing a non-zero solution will have doubled in radius. For simplicity, we keep the time-step fixed throughout the simulation, however in an actual computation, the time-step could be allowed to increase as the maximum of the profile decreased.

As in the previous sub-section, we compare different methods for obtaining a solution to this problem. We use the RKL2 and RK2 methods as described in points (a) and (b) above for this problem as well. The implicit solution of the non-linear problem requires a different method than those mentioned above in points (c) and (d). For this case, we again use the TR-BDF2 scheme with a multigrid-accelerated Generalized Minimal Residual (GMRES) Newton–Krylov method. We find that the GMRES-multigrid method converges well for this problem under increasing resolution. We allowed the GMRES to make 5 restarts with a subspace of 30 basis vectors per restart. We cannot use a direct matrix inversion because the operator depends on the solution.

Table 12 compares the three schemes for obtaining a solution to this problem. The table shows wallclock computational time as well as L_2 errors for each of the schemes. Because the solution has a sharp discontinuity, O(1) errors are present in an approximately two-zone wide ring at the edge of the profile. To avoid contamination from these zones, the errors presented in the table are computed over a square centered on the origin with sides of length 0.25. Despite this, some decrement below second-order accuracy is still inevitable for a problem with such a steep discontinuity. As in the linear problem presented in the previous sub-section, the computational wallclock time for the STS method is significantly shorter than that of the RK2 explicit sub-cycling. At a resolution of 512^2 zones, the RKL2 scheme is 28 times faster than the time-explicit RK2 scheme. The GMRES method employed to solve the implicit problem presented in this sub-section was considerably slower than even the RK2 scheme at all resolutions. Each of the methods results in comparable errors in the solution profile. Note also that Table 12 shows that the multigrid GMRES method does not scale as well as the pointwise multigrid employed in the previous sub-section. In fact, the temporal scaling is similar to that of the RKL2 scheme.

The results in this sub-section demonstrate that the RKL2 scheme is capable of outperforming an implicit scheme even for a non-linear problem. With the twin advantages of lesser computational cost and smaller development effort, a solution

Table 12
A comparison of the timing and errors of the RKL2, second-order Runge–Kutta, and multigrid GMRES TR-BDF2 schemes. The number of time-steps (N_t) and stages (s) for STS is included as well as the L_2 error and the wall clock time. Note that the errors presented in this table are computed only using the central region of the solution.

Mesh	RKL2				RK2 – forward Euler			Multigrid GMRES		
	N_t	s	L_2 error	Time (s)	N_t	L_2 error	Time (s)	N_t	L_2 Error	Time (s)
64^2	30	50	1.37×10^{-3}	0.119	18,900	3.04×10^{-4}	1.85	30	5.79×10^{-4}	19.1
128^2	60	71	3.36×10^{-4}	1.36	76,500	7.48×10^{-5}	29.0	60	3.77×10^{-5}	195.4
256^2	120	100	6.94×10^{-5}	17.4	302,400	1.28×10^{-5}	455	120	2.44×10^{-5}	2210
512^2	240	142	1.93×10^{-5}	232	1,209,600	7.54×10^{-6}	6720	240	1.70×10^{-5}	23,961

of comparable accuracy can be produced using the RKL2 scheme. The results of this section show us that for linear and non-linear parabolic problems, the RKL schemes provide an easy, fast and reliable alternative to traditional sub-cycling or implicit methods.

9. Conclusions

We have presented a numerical scheme for solving parabolic operators in partial differential equation systems. The RKL methods are capable of computing explicit updates to parabolic problems with considerably higher computational efficiency than explicit forward-Euler methods. The results of the paper can be summarized as follows:

(1) The existing Runge–Kutta–Chebyshev methods are not monotone for even the simple case of one-dimensional thermal conduction between two different materials. We showed that the RKL method described in this paper is monotone and produces more favorable results for this test problem than the RKC methods.
(2) We introduced the Convex Monotone Property to explain this difference between the two methods. The RKL methods posses this property, whereas the RKC methods do not. With this property, the RKL methods will be monotone for a constant conduction coefficient less than the maximum, or for any time-step shorter than the maximum stable time-step. This is favorable for spatially-varying problems, and prevents spurious oscillations which could be very problematic for non-linear problems. This property is also useful for mixed hyperbolic and parabolic systems when it would be desirable to set the time-step from the hyperbolic solver.
(3) We showed that the RKL methods can be derived in both first- or second-order accurate formulations. We call these RKL1 and RKL2, respectively. They are based on the recursion relation for the shifted Legendre polynomials.
(4) The RKL1 and RKL2 schemes consistently perform at design accuracy for a wide range of linear and non-linear test problems in one, two and three dimensions.
(5) Explicit formulae for the numerical implementation of the RKL1 and RKL2 schemes are given in Section 3.
(6) We showed the RKL methods are suitable for use in a wide variety of applications, including viscosity and thermal conduction in hydrodynamics, resistivity and ambipolar diffusion in MHD, radiation-diffusion modeling, curvature flow in level set dynamics and complicated multi-material thermal conduction. In each of these cases a substantial decrease in computational cost was achieved when using RKL as compared to explicit forward-Euler without the complexity of implementation of an implicit method. This cost decrease did not adversely affect the quality of the result.
(7) We showed that the RKL methods are competitive with other traditional solution methods for parabolic problems. The wallclock times required for RKL2 is less than either traditional time-explicit RK sub-cycling or implicit methods using multigrid. This is true both for linear and non-linear problems.

Acknowledgements

DSB acknowledges support via NSF grants NSF-AST-0947765 and NSF-AST-1009091. DSB also acknowledges support via NASA grants from the Fermi program as well as NASA-NNX 12A088G. The majority of simulations were performed on a cluster at UND that is run by the Center for Research Computing. Computer support on NSF's XSEDE computing resources is also acknowledged.

The authors wish to thank Christopher Romick for providing the exact solution to the Argon shock structure. TDA was supported by the US Department of Energy. TDA would like to thank the John Bdzil (LANL), Joseph Powers (Notre Dame), Steve Son (Purdue) and Lori Groven (Purdue) for useful discussions during this work.

Appendix A

Many of the examples given here are drawn from CFD and MHD problems with non-ideal terms. Instead of cataloguing the equations in a piecemeal fashion, we provide a comprehensive set of equations for MHD with viscosity, thermal conduction, Ohmic dissipation and ambipolar diffusion. The hydrodynamic variant can be derived by setting the magnetic fields to zero. In [36] we provide a detailed discussion on thermal conduction in a tenuous plasma. As a result, thermal conduction in a tenuous plasma is only included here for the sake of completeness and without extensive discussion.

Let the fluid density be denoted by ρ, the pressure by "P", the thermal energy density by "e", the total energy density by ε, the velocity by $\mathbf{v}$ and the magnetic field by $\mathbf{B}$. We have the constitutive relation $e = P/(\gamma - 1)$ and the equation for the total energy density (in CGS units) $\varepsilon = e + \rho\mathbf{v}^2/2 + \mathbf{B}^2/(8\pi)$. The specific heat at constant volume is denoted by c_V and the ratio of specific heats is γ. In the next two paragraphs we describe the inclusion of non-ideal fluid terms (viscosity and thermal conduction) and non-ideal MHD terms (Ohmic dissipation and ambipolar diffusion).

The viscous stress π_{ij} is related to the viscous strain D_{ij} by

$$\pi_{ij} \equiv \mu D_{ij}; \qquad D_{ij} \equiv \frac{\partial v_i}{\partial x_j} + \frac{\partial v_j}{\partial x_i} - \frac{2}{3}(\nabla \cdot \mathbf{v})\delta_{ij} \tag{A.1}$$

where μ is the coefficient of shear viscosity and is determined experimentally. The energy flux associated with viscous heating is given by

$$F_{Viscous} = -\mathbf{v}\vec{\vec{\pi}}. \tag{A.2}$$

The thermal conduction in the classical limit is given by an energy flux of the form

$$\begin{aligned} \mathbf{F}_{class} &= -\kappa_0 T^{5/2}\mathbf{b}(\mathbf{b}\nabla T) \quad \text{(in the presence of magnetic field)} \\ \mathbf{F}_{class} &= -\kappa_0 T^{5/2}\nabla T \quad \text{(without a magnetic field)} \end{aligned} \tag{A.3}$$

where "κ_0" is a coefficient of thermal conduction, "T" is the temperature and "$\mathbf{b} \equiv \mathbf{B}/|\mathbf{B}|$" is the direction of the magnetic field. As the temperature gradient in a dilute plasma steepens, the electrons which carry the thermal conduction excite resonant waves in the plasma. As a result, the energy flux due to thermal conduction saturates. The saturated flux of thermal conduction is given by

$$\mathbf{F}_{sat} = -5\phi\rho c_s^3 \,\mathrm{sgn}(\mathbf{b}\nabla T)\mathbf{b} \tag{A.4}$$

where $\phi \approx 0.3$ is a dimensionless coefficient and c_s is the sound speed. The classical flux smoothly transitions over to a saturated flux as the temperature gradient increases, so that we have a net flux of thermal conduction given by

$$\mathbf{F}_{Conduction} = G(F_{class}/F_{sat})\mathbf{F}_{class} + \big[1 - G(F_{class}/F_{sat})\big]\mathbf{F}_{sat}. \tag{A.5}$$

The flux limiter, $G(x)$, arbitrates this smooth transition and is given by

$$G(x) = \max\big\{0, \min[1, 1 - 0.5x]\big\}. \tag{A.6}$$

This completes a basic description of the non-ideal terms for viscosity and heat conduction. As we see, the thermal conduction can become rather complicated and its good handling via the RKL2 super-time-stepping method is demonstrated in our prior paper. In this paper, we demonstrate how the viscous and thermal conduction operators work together by constructing a problem that has a semi-analytic solution via a system of coupled ODEs. To achieve this simplification, we have kept the thermal conduction and viscosity coefficients constant in the test problem that was designed here.

The magnetic field $\mathbf{B}$ evolves according to Faraday's law

$$\frac{\partial \mathbf{B}}{\partial t} + c\nabla \times \mathbf{E} = 0 \tag{A.7}$$

where "c" is the speed of light and $\mathbf{E}$ is the electric field. In the MHD approximation, the displacement current can be dropped from Maxwell's equations and the current density $\mathbf{J}$ is given by Ampere's law as

$$\mathbf{J} = \frac{c}{4\pi}\nabla \times \mathbf{B}. \tag{A.8}$$

In the MHD approximation, the electric field is related constitutively to the other flow variables as

$$\begin{aligned} &\mathbf{E} = \mathbf{E}_{Motional} + \mathbf{E}_{Ohmic} + \mathbf{E}_{Ambipolar} \quad \text{where } \mathbf{E}_{Motional} \equiv -\left(\frac{1}{c}\right)\mathbf{v} \times \mathbf{B}; \\ &\mathbf{E}_{Ohmic} \equiv \frac{c}{4\pi\sigma}\nabla \times \mathbf{B}; \qquad \mathbf{E}_{Ambipolar} \equiv -\left(\frac{1}{c}\right)\frac{1}{4\pi\Gamma\rho_{ion}\rho_{neutral}}\mathbf{B} \times \big[\mathbf{B} \times (\nabla \times \mathbf{B})\big]. \end{aligned} \tag{A.9}$$

Here, $\mathbf{E}_{Motional}$ is the motional electric field from ideal MHD and does not cause any dissipation; $\mathbf{E}_{Ohmic}$ is the Ohmic dissipation and can heat the plasma through resistive effects and $\mathbf{E}_{Ambipolar}$ is the electric field resulting from ion-neutral drift and heats the plasma via friction between the ions and neutrals. The electric conductivity of the plasma is given by σ; the ion and neutral densities are given by ρ_{ion} and $\rho_{neutral}$ and Γ is a measure of the collision rate between ions and neutrals. We do not include the Hall terms because the Hall operator is dispersive and, therefore, not within the purview of a scheme for treating (diffusive) parabolic equations. The Poynting flux, which mediates the flow of total electromagnetic energy density,

is given by

$$\mathbf{S} \equiv \frac{c}{4\pi}\mathbf{E}\times\mathbf{B} \equiv \mathbf{S}_{Motional} + \mathbf{S}_{Ohmic} + \mathbf{S}_{Ambipolar}$$
$$\text{where } \mathbf{S}_{Motional} \equiv \frac{c}{4\pi}\mathbf{E}_{Motional}\times\mathbf{B};\ \mathbf{S}_{Ohmic} \equiv \frac{c}{4\pi}\mathbf{E}_{Ohmic}\times\mathbf{B};\ \mathbf{S}_{Ambipolar} \equiv \frac{c}{4\pi}\mathbf{E}_{Ambipolar}\times\mathbf{B}. \tag{A.10}$$

We are now in a position to describe the non-ideal MHD equations. We begin by describing the primitive form, which is well-suited to analyzing the PDE system. The continuity equation is given by

$$\frac{D\rho}{Dt} = -\rho\nabla\mathbf{v}. \tag{A.11}$$

Here D/Dt is the material derivative. The momentum equation becomes

$$\rho\frac{D\mathbf{v}}{Dt} = -\nabla P + \frac{1}{c}(\mathbf{J}\times\mathbf{B}) + \overleftrightarrow{\nabla}\pi = -\nabla\left(P + \frac{\mathbf{B}^2}{8\pi}\right) + \frac{1}{4\pi}(\mathbf{B}\nabla)\mathbf{B} + \overleftrightarrow{\nabla}\pi \tag{A.12}$$

The equation for the evolution of the specific entropy is equivalent to the thermal energy equation and is given by

$$\rho T\frac{D}{Dt}\left[c_V \ln\left(\frac{P}{\rho^\gamma}\right)\right] = -\nabla\mathbf{F}^{cond} + \pi_{ij}\frac{\partial v_i}{\partial x_j} + (\mathbf{E}_{Ohmic} + \mathbf{E}_{Ambipolar})\mathbf{J}. \tag{A.13}$$

The above equation clearly shows that the non-ideal terms contribute exclusively to the generation of entropy. The magnetic field evolves according to Eq. (A.7) but can be explicitly written as

$$\frac{\partial\mathbf{B}}{\partial t} = \nabla\times(\mathbf{v}\times\mathbf{B}) - \frac{c^2}{4\pi}\nabla\times\left(\frac{1}{\sigma}\nabla\times\mathbf{B}\right) + \nabla\times\left\{\frac{1}{4\pi\Gamma\rho_{ion}\rho_{neutral}}\mathbf{B}\times\left[\mathbf{B}\times(\nabla\times\mathbf{B})\right]\right\}. \tag{A.14}$$

These were the equations that were used in the dispersion analysis. The same equations were also used for calculating the structure of the steady traveling shock with viscosity and heat conduction.

We are now in a position to assemble the MHD equations in a conservation form which makes the role of the ideal and non-ideal terms explicit. We have

$$\frac{\partial}{\partial t}\begin{pmatrix}\rho\\ \rho v_x\\ \rho v_y\\ \rho v_z\\ \varepsilon\\ B_x\\ B_y\\ B_z\end{pmatrix} + \frac{\partial}{\partial x}\begin{pmatrix}\rho v_x\\ \rho v_x^2 + P + \mathbf{B}^2/8\pi - B_x^2/4\pi\\ \rho v_x v_y - B_x B_y/4\pi\\ \rho v_x v_z - B_x B_z/4\pi\\ (\varepsilon + P + \mathbf{B}^2/8\pi)v_x - B_x(\mathbf{v}\cdot\mathbf{B})/4\pi\\ 0\\ (v_x B_y - v_y B_x)\\ -(v_z B_x - v_x B_z)\end{pmatrix}$$
$$+ \frac{\partial}{\partial y}\begin{pmatrix}\rho v_y\\ \rho v_x v_y - B_x B_y/4\pi\\ \rho v_y^2 + P + \mathbf{B}^2/8\pi - B_y^2/4\pi\\ \rho v_y v_z - B_y B_z/4\pi\\ (\varepsilon + P + \mathbf{B}^2/8\pi)v_y - B_y(\mathbf{v}\cdot\mathbf{B})/4\pi\\ -(v_x B_y - v_y B_x)\\ 0\\ (v_y B_z - v_z B_y)\end{pmatrix} + \frac{\partial}{\partial z}\begin{pmatrix}\rho v_z\\ \rho\ v_x v_z - B_x B_z/4\pi\\ \rho v_y v_z - B_y B_z/4\pi\\ \rho v_z^2 + P + \mathbf{B}^2/8\pi - B_z^2/4\pi\\ (\varepsilon + P + \mathbf{B}^2/8\pi)v_z - B_z(\mathbf{v}\cdot\mathbf{B})/4\pi\\ (v_z B_x - v_x B_z)\\ -(v_y B_z - v_z B_y)\\ 0\end{pmatrix}$$
$$= -\frac{\partial}{\partial x}F_{Non\text{-}ideal} - \frac{\partial}{\partial y}G_{Non\text{-}ideal} - \frac{\partial}{\partial z}H_{Non\text{-}ideal} \tag{A.15}$$

where the non-ideal fluxes in the x-, y- and z-directions can be written as:

$$F_{Non\text{-}ideal} \equiv \begin{pmatrix}0\\ -\pi_{xx}\\ -\pi_{xy}\\ -\pi_{xz}\\ F_{Viscous;x} + F_{Conduction;x} + S_{Ohmic;x} + S_{Ambipolar;x}\\ 0\\ -cE_{Ohmic;z} - cE_{Ambipolar;z}\\ cE_{Ohmic;y} + cE_{Ambipolar;y}\end{pmatrix};$$

$$G_{Non\text{-}ideal} \equiv \begin{pmatrix} 0 \\ -\pi_{xy} \\ -\pi_{yy} \\ -\pi_{yz} \\ F_{Viscous;y} + F_{Conduction;y} + S_{Ohmic;y} + S_{Ambipolar;y} \\ cE_{Ohmic;z} + cE_{Ambipolar;z} \\ 0 \\ -cE_{Ohmic;x} - cE_{Ambipolar;x} \end{pmatrix};$$

$$H_{Non\text{-}ideal} \equiv \begin{pmatrix} 0 \\ -\pi_{xz} \\ -\pi_{yz} \\ -\pi_{zz} \\ F_{Viscous;z} + F_{Conduction;z} + S_{Ohmic;z} + S_{Ambipolar;z} \\ -cE_{Ohmic;y} - cE_{Ambipolar;y} \\ cE_{Ohmic;x} + cE_{Ambipolar;x} \\ 0 \end{pmatrix}. \tag{A.16}$$

The separation between the ideal and non-ideal terms can be used to advantage in numerical implementations that rely on an operator split formulation.

References

[1] A Abdulle, Chebyshev methods based on orthogonal polynomials, Thesis No. 3266, Dept. Math., Univ. of Geneva, 2001.
[2] A Abdulle, Fourth order Chebyshev methods with recurrence relation, SIAM J. Sci. Comput. 23 (6) (2002) 2041–2054.
[3] V. Alexiades, G. Amiez, P.A. Gremaud, Super-time-stepping acceleration of explicit schemes for parabolic problems, Commun. Numer. Methods Eng. 12 (1) (1996) 31–42.
[4] T.D. Aslam, J.B. Bdzil, D.S. Stewart, Level set methods applied to modeling detonation shock dynamics, J. Comput. Phys. 126 (2) (1996) 390–409.
[5] T. Aslam, J. Bdzil, Numerical and theoretical investigations on detonation-inert confinement interactions, in: Twelfth International Detonation Symposium, 2002.
[6] T. Aslam, J. Bdzil, Numerical and theoretical investigations on detonation confinement sandwich tests, in: Thirteenth International Detonation Symposium, 2006.
[7] T. Aslam, S. Jackson, J. Morris, Proton radiography of PBX 9502 detonation shock dynamics sandwich test, in: Fifteenth APS Conference on Shock Compression of Condensed Matter, 2009.
[8] H. Baer, K. Stephan, Heat and Mass Transfer, Springer-Verlag, Berlin, Heidelberg, 2006.
[9] D.S. Balsara, Second order accurate schemes for magnetohydrodynamics with divergence-free reconstruction, Astrophys. J. Suppl. Ser. 151 (1) (2004) 149–184.
[10] D.S. Balsara, D.A. Tilley, J.C. Howk, Simulating anisotropic thermal conduction in supernova remnants—I. Numerical methods, Mon. Not. R. Astron. Soc. 386 (2) (2008) 627–641.
[11] D.S. Balsara, T. Rumpf, M. Dumbser, C.-D. Munz, Efficient, high accuracy ADER-WENO schemes for hydrodynamics and divergence-free magnetohydrodynamics, J. Comput. Phys. 228 (2009) 2480–2516.
[12] D.S. Balsara, C. Meyer, M. Dumbser, H. Du, Z. Xu, Efficient implementation of ADER schemes for Euler and magnetohydrodynamical flows on structured meshes – comparison with Runge–Kutta methods, J. Comput. Phys. 235 (2013) 934, http://dx.doi.org/10.1016/j.jcp.2012.04.051.
[13] R.E. Bank, W.M. Coughran Jr., W. Fichtner, E.H. Grosse, D.J. Rose, R.K. Smith, Transient simulation of silicon devices and circuits, IEEE Trans. Electron Devices ED-32 (1985) 1992.
[14] G.I. Barenblatt, On some unsteady motions of a liquid or a gas in a porous medium, Prikl. Mat. Meh. 16 (1) (1952) 67.
[15] J. Bdzil, D. Stewart, The dynamics of detonation in explosive systems, Annu. Rev. Fluid Mech. 39 (2007) 263–292.
[16] R. Becker, Stosswelle und Detonation, Z. Phys. 8 (1923) 321.
[17] M. Dumbser, D.S. Balsara, High-order unstructured one-step PNPM schemes for the viscous and resistive MHD equations, Comput. Model. Eng. Sci. 54 (3) (2010) 301–334.
[18] G. Gassner, F. Lörcher, C.-D. Munz, A contribution to the construction of diffusion fluxes for finite volume and discontinuous Galerkin schemes, J. Comput. Phys. 224 (2) (2007) 1049–1063.
[19] G. Gassner, F. Lörcher, C.-D. Munz, A discontinuous Galerkin scheme based on a space–time expansion II. Viscous flow equations in multi dimensions, J. Sci. Comput. 34 (3) (2008) 260–286.
[20] K. Gurski, S. O'Sullivan, A stability study of a new explicit numerical scheme for a system of differential equations with a large skew-symmetric component, SIAM J. Numer. Anal. 49 (1) (2011) 368–386.
[21] K. Gurski, A stability and cost study of explicit and dyadic time stepping for stiff nonsymmetric problems, in: AIP Conf. Proc., vol. 1368, 2011, pp. 239–242.
[22] E. Hairer, G. Wanner, Solving Ordinary Differential Equations II – Stiff and Differential Algebraic Problems, 2nd edition, Springer Ser. Comput. Math., vol. 14, 1996.
[23] L. Hill, T. Aslam, The LANL detonation-confinement test: prototype development and sample results, in: Twelfth APS Conference on Shock Compression of Condensed Matter, 2003.
[24] L. Hill, T. Aslam, Detonation shock dynamics calibration for PBX 9502 with temperature, density, and material lot variations, in: Fourteenth International Detonation Symposium, 2010.
[25] W. Hundsdorfer, J.G. Verwer, Numerical Solution of Time-Dependent Advection–Diffusion Reaction Equations, Springer-Verlag, 2003.
[26] F.P. Incropera, D.P. DeWitt, Introduction to Heat Transfer, John Wiley & Sons, New York, 1990.
[27] J.D. Jackson, Classical Electrodynamics, Wiley, New York, 1975, pp. 475–479.
[28] V.I. Lebedev, How to solve stiff systems of differential equations by explicit methods, Numer. Methods Appl. (1994) 45–80.
[29] V.I. Lebedev, Explicit difference schemes for solving stiff problems with a complex or separable spectrum, Comput. Math. Math. Phys. 40 (12) (2000) 1729–1740.

[30] W.-C. Lee, S.-L. Chung, Ignition phenomena and reaction mechanisms of the self-propagating high-temperature synthesis reaction in the titanium–carbon–aluminum system, J. Am. Ceram. Soc. 80 (1) (1997).
[31] R.J. LeVeque, Numerical Methods for Conservation Laws, Birkhäuser, 1992.
[32] R.J. LeVeque, H.C. Yee, A study of numerical methods for hyperbolic conservation laws with stiff source terms, J. Comput. Phys. 86 (1) (1990) 187–210.
[33] G.H. Markstein, Non-steady Flame Propagation, Pergamon, MacMillan, 1964.
[34] R.E. Marshak, Effect of radiation on shock wave behavior, Phys. Fluids (US) 1 (1958) 24.
[35] A.A. Medovikov, High order explicit methods for parabolic equations, BIT Numer. Math. 38 (2) (1998) 372–390.
[36] C.D. Meyer, D.S. Balsara, T.D. Aslam, A second-order accurate Super TimeStepping formulation for anisotropic thermal conduction, Mon. Not. R. Astron. Soc. 422 (2012) 2102–2115.
[37] V.A. Mousseau, D.A. Knoll, W.J. Rider, Physics-based preconditioning and the Newton–Krylov method for non-equilibrium radiation diffusion, J. Comput. Phys. 160 (2) (2000) 743–765.
[38] S. O'Sullivan, T.P. Downes, A three-dimensional numerical method for modelling weakly ionized plasmas, Mon. Not. R. Astron. Soc. 376 (4) (2007) 1648–1658.
[39] S. Osher, R. Fedkiw, Level Set Methods and Dynamic Implicit Surfaces, Appl. Math. Sci., Applied Mathematical Sciences, vol. 153, Springer, ISBN 978-0-387-95482-0, 2002.
[40] S. Osher, J. Sethian, Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton–Jacobi formulations, J. Comput. Phys. 79 (1) (1988) 12–49.
[41] S. Patankar, Numerical Heat Transfer and Fluid Flow, Taylor & Francis, 1980.
[42] P. Reinicke, J. Meyer-ter-Vehn, The point explosion with heat conduction, Phys. Fluids A, Fluid Dyn. 3 (1991) 1807.
[43] E. Rouy, A. Tourin, A viscosity solutions approach to shape-from-shading, SIAM J. Numer. Anal. 29 (3) (1992) 867–884.
[44] L.I. Rubenstein, The Stefan Problem, American Mathematical Society, 1971.
[45] J.A. Sethian, Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry Fluid Mechanics, Computer Vision, and Materials Science, vol. 3, Cambridge University Press, 1999.
[46] B.P. Sommeijer, L.F. Shampine, J.G. Verwer, RKC: an explicit solver for parabolic PDEs, J. Comput. Appl. Math. 88 (2) (1998) 315–326.
[47] G. Strang, On the construction and comparison of difference schemes, SIAM J. Numer. Anal. 5 (3) (1968) 506–517.
[48] R. Tyson, L.G. Stern, R.J. LeVeque, Fractional step methods applied to a chemotaxis model, J. Math. Biol. 41 (5) (2000) 455–475.
[49] P.J. van der Houwen, The development of Runge–Kutta methods for partial differential equations, Appl. Numer. Math. 20 (3) (1996) 261–272.
[50] P.J. van Der Houwen, B.P. Sommeijer, On the internal stability of explicit, m-stage Runge–Kutta methods for large m-values, Z. Angew. Math. Mech. 60 (1980) 479–485.
[51] J.L. Vázquez, The Porous Medium Equation: Mathematical Theory, Oxford University Press, New York, 2007.
[52] J.G. Verwer, Explicit Runge–Kutta methods for parabolic partial differential equations, Appl. Numer. Math. 22 (1) (1996) 359–379.
[53] R.C. Weast, Handbook of Chemistry and Physics, 52 edition, The Chem. Rubber Co., Florida, 1984, p. C475.