The performance of filtered leapfrog schemes in benchmark simulations

Article

Accepted Version

Williams Paul D. (Orcid ID: 0000-0002-9713-9820)

The Performance of Filtered Leapfrog Schemes in Benchmark Simulations

Paul D. Williams*

Department of Meteorology, University of Reading, Reading, United Kingdom

Jerry M. Straka and Katharine M. Kanak

Norman, Oklahoma

*Corresponding author contact details: Paul D. Williams, University of Reading, Meteorology Building, Whiteknights Road, Earley Gate, Reading, RG6 6ET, UK, Phone: +44 (0) 118 378 8424, Fax: +44 (0) 118 378 8316, E-mail: p.d.williams@reading.ac.uk

Keywords: leapfrog scheme, time filters, benchmark simulations, numerical weather prediction

Funding information: none

Abstract

The stabilisation of the leapfrog time-stepping scheme provided by the Robert–Asselin filter has enabled decades of atmospheric and oceanic research and weather and climate predictions. The unfortunate concomitant reduction from second-order accuracy to first-order accuracy inflicted by the filter has recently ushered in a new generation of leapfrog time filters that preserve the second-order accuracy, including the Robert–Asselin–Williams (RAW) filter and its variants. These modern filtered leapfrog schemes have previously been shown to improve numerical simulations made using both simple conceptual models and comprehensive general circulation models. However, their performance in standard benchmark experiments has not previously been assessed. Here we evaluate these filtered leapfrog schemes in four classic benchmark experiments: linear scalar advection; a nonlinear density current in the quasi-compressible equations; a nonlinear rising warm bubble in the fully compressible equations; and the linked behaviour of nonlinear twin tropical cyclones in the rotating shallow-water equations. For a given time-step size, the filtered leapfrog schemes are found to compare favourably with the third-order Runge–Kutta (RK3) scheme. They are also less computationally expensive than RK3, at roughly one-third to one-half the cost per time step. For a given computational expenditure, the filtered leapfrog schemes are found to produce smaller errors with respect to the analytical solution (where available) than RK3. Furthermore, the filtered leapfrog schemes are found to be numerically stable, even when the discretisation method splits the slow advection and diffusion modes from the fast acoustic and gravity-wave modes. Given that implementing filter upgrades requires only minimally invasive changes to an existing computer code, our results provide support for the continued use of filtered leapfrog schemes in atmosphere and ocean models.

1. Introduction

For most of the past five decades, the leapfrog time-stepping scheme, stabilised with a Robert–Asselin (RA) filter, has been widely used for weather prediction and climate simulation in atmosphere and ocean numerical models. The RA filter was introduced by Robert (1966) to reduce time splitting of solutions at odd and even time steps with the leapfrog scheme. Asselin (1972) showed that the filter effectively controls numerical instabilities by substantially damping the $2\Delta t$ computational mode. Unfortunately, the RA filter's non-zero damping of the lower-frequency physical modes reduces the formal time-stepping accuracy of the leapfrog scheme from second order to first order. (Hereafter, unless otherwise stated, order-of-accuracy refers to the effects of temporal truncation errors on the amplitude of linear oscillations.) Another way to control time splitting is to periodically restart the solution every few time steps with one of the forward schemes. While effective at temporarily removing the computational mode, this approach can result in undesirable computer code changes, and it also reduces the overall formal accuracy to first order.

Because methods for controlling the computational mode of the leapfrog scheme without reducing the overall O(2) accuracy have historically been lacking, many weather forecasting centres now use alternatives to the leapfrog scheme. Some of the most popular alternatives are discussed by Durran (1991) and include the O(2) and O(3) Adams–Bashforth schemes (e.g., Deardorff 1980; Durran 1991), the Crowley schemes (Smolarkiewicz 1982; Schlesinger 1985; Tremback et al. 1987; Wicker and Skamarock 2002), the O(2) and O(3) Runge–Kutta schemes (RK2, RK3; Wicker and Skamarock 1998; Wicker and Skamarock 2002), and variants of these (Shu 1988; Shu and Osher 1988; among others). Durran (1999) proposed combining the O(2) Adams–Bashforth and O(3) Adams–Moulton schemes to produce a predictor–corrector scheme. More recently, Wicker (2009) demonstrated the stability and use of this Adams–Bashforth–Moulton predictor–corrector scheme for integrations using split fast modes (sound and gravity waves) and slow modes (advection and turbulence), as proposed by Klemp and Wilhelmson (1978) and Skamarock and Klemp (1992), if the available computational memory is sufficient. Nevertheless, the leapfrog scheme remains

popular in climate models. The Robert–Asselin–Williams (RAW) filter is a generalisation of the RA filter with increased accuracy (Williams 2009, 2011, 2013). The RAW-filtered leapfrog schemes all possess O(2) phase errors, but the improved O(3)–O(7) amplitude errors may be beneficial in simulations whose overall numerical accuracy is limited by temporal differencing rather than spatial differencing.

Higher-order temporal and spatial differencing schemes have been shown, both theoretically and numerically, to profoundly affect the accuracy of simulated atmospheric and oceanic phenomena, including those on the storm-scale, mesoscale, synoptic scale, and planetary scale. Over the past few years, the fortunes of the leapfrog scheme have been revived by a series of proposed measures for damping the computational mode while preserving the O(2) accuracy of the physical mode. Williams (2009, 2011, 2013), Moustaoui et al. (2014), Li and Trenchea (2014), and Maurya et al. (2019) have each demonstrated how various time-filtered leapfrog schemes with O(3)–O(7) amplitude errors can be developed while retaining the O(2) phase error. These modified leapfrog schemes all have O(2) overall errors. For example, Williams (2009) developed the RAW filter to prevent the amplitude error reduction associated with the RA filter. Subsequently, Williams (2011) tested the RAW filter in semi-implicit systems, while Williams (2013) proposed and tested a family of schemes from combinations of O(2) and O(4) temporal filters. Using these filters allows for a family of leapfrog schemes with O(3), O(5), and O(7) amplitude errors to be formed, as predicted theoretically, depending on the values of three filter parameters. Finally, Babarsky and Sharpley (1997) developed time-centred advection schemes with up to fourth-order accuracy in space and time, but they employed the RA filter in their numerical tests, which likely degraded the order of the time accuracy in the numerical solutions.

It has been shown with three-dimensional prediction models that the RAW-filtered leapfrog scheme improves the skill of medium-term weather forecasts and ocean simulations (Amezcua et al. 2011; Young et al. 2014; Amezcua and Williams 2015). The RAW-filtered leapfrog scheme is now the default time-stepping scheme in the Taiwanese ocean model (Young et al. 2012). Ren and Leslie (2011) found good

performance of the RAW-filtered leapfrog scheme in an ice sheet model. Also, Moustaoui et al. (2014) demonstrated improved performance of a fourth-order, implicitly time-filtered leapfrog scheme, with an improved Courant–Friedrich–Lewy (CFL) criterion, compared to the RA-filtered leapfrog scheme and RK3 scheme for the nonlinear Lorenz (1963) system, a linear shallow-water problem, and twin tropical cyclones straddling the equator in a global shallow-water model. Li and Trenchea (2014) developed a higher-order RA-filtered leapfrog scheme, which achieved O(3) errors with specific smoothing parameters. Finally, Maurya et al. (2019) developed a family of stable high-order optimised hybrid RA-type time filters for leapfrog schemes, which strongly damp the computational modes without significantly damping the physical modes. The leapfrog schemes proposed by Maurya et al. (2019) have CFL criteria that are at least as permissive as the RK3 schemes.

Despite clear theoretical and numerical evidence that modern filtered leapfrog schemes deliver increased accuracy, what is hitherto lacking is a quantitative assessment of the performance of these schemes in standard linear and nonlinear multidimensional benchmark experiments. Previous studies of filtered leapfrog schemes have generally employed idealised nonlinear problems, as well as statistics from comprehensive general circulation models, with relatively little research activity in the middle of the model complexity hierarchy. For example, there is a gap in knowledge of how the schemes perform with scalar advection by idealised kinematic flow fields such as solid-body rotation. In addition, it is unknown how the schemes perform with more complex nonlinear flows that are not complicated by physics parameterisations of turbulence, friction, microphysics, and radiation. Problems of this sort include two-dimensional density currents (Straka et al. 1993), three-dimensional warm thermals (Carpenter et al. 1990; Tripoli 1992; Bryan and Fritsch 2002), and the linked behaviour of twin tropical cyclones in the two-dimensional shallow-water equations (Ferreira et al. 1996; Moustaoui et al. 2002; Moustaoui et al. 2014) and other vortex interactions such as binary vortices (Holland and Dietachmayer 1993). Finally, the performance of modern filtered leapfrog schemes has not previously been studied for geophysical systems solved with time-split integrations (Klemp and Wilhelmson 1978; Skamarock and Klemp 1992) involving split slow modes

(advection, turbulence, gravity waves, and physical processes) and fast modes (pressure gradients, divergence, sound waves, and in some cases gravity waves).

The present study aims to address the above omissions and provide support for the continued use of filtered leapfrog models, especially the RAW-filtered leapfrog schemes that are minimally invasive compared to the RA filter and require only trivial code changes. The study has various specific goals. First, we aim to show that the RAW-filtered leapfrog schemes – hereafter referred to as the Williams filtered leapfrog (WFL) schemes – can achieve accurate solutions without the excessively small time steps required by the RA-filtered leapfrog scheme. Second, we aim to show that the temporal truncation error of the WFL schemes and the RA-filtered leapfrog scheme for advection becomes increasingly important compared to diffusion tendencies (solved with an O(1) forward-in-time operator). This can be shown by examining pure advection problems, as well as using problems where both advection and diffusion tendencies are important and decreasing the eddy mixing coefficient. Third, we aim to show that the WFL schemes can be stably integrated in additive fast–slow mode-split models following methods used for leapfrog models. Fourth, we aim to show that the total energy and total mass are better conserved with the WFL schemes than the RA-filtered leapfrog scheme.

The layout of the present study is as follows. Section 2 summarizes the RA-filtered leapfrog scheme as well as the WFL schemes. The O(3) Runge–Kutta schemes (Shu 1988; Shu and Osher 1988; Wicker and Skamarock 2002) are summarised as well, as they are used for reference solutions when analytical solutions are lacking. The results of the numerical experiments are presented in Section 3. The advection of a cone in two-dimensional solid-body rotation is used to evaluate the effects of time integration schemes for a pure advection problem (Section 3a). The two-dimensional quasi-compressible equations are integrated for a dry density current problem (Section 3b). The three-dimensional fully compressible equations are integrated for a dry rising warm plume (Section 3c). The divergence-permitting two-dimensional shallow-water equations with full Coriolis force are integrated for the problem of the nonlinear interaction of twin tropical

cyclones (with opposite signs of rotation) straddling the equator (Section 3d). The results from each of the experiments are compared to analytical or nearly converged numerical solutions. The conclusions of this study, which show accurate and stable solutions for the family of WFL schemes, are summarised in Section 4.

2. Time integration schemes

This section will briefly review the finite-difference approximations that are associated with the various time-stepping schemes used in this study. Some of their properties are listed in Table 1.

a) The RA-filtered leapfrog scheme

The RA filter is effective in keeping the leapfrog scheme stable and is trivial to incorporate. The prediction of a dependent variable, A, in time can be written as:

$$\frac{\mathrm{d}A}{\mathrm{d}t} = f(A) \tag{2.1}$$

or in finite-difference form using the unfiltered leapfrog scheme as:

$$A_i^{n+1} = A_i^{n-1} + 2\,\Delta t\, f(A_i^n), \tag{2.2}$$

where superscript n is the time index and subscript i is a spatial index. The RA filter can be written as:

$$A_i^{n*} = A_i^n + \nu\left(A_i^{n-1*} - 2A_i^n + A_i^{n+1}\right) \tag{2.3}$$

and the RA-filtered leapfrog scheme becomes:

$$A_i^{n+1} = A_i^{n-1*} + 2\,\Delta t\, f(A_i^{n*}), \tag{2.4}$$

where asterisks denote filtered values. Use of this second-order filter leads to an amplitude-damped O(1) leapfrog scheme that significantly degrades physical modes (Schlesinger et al. 1983; Williams 2009). Only one evaluation of $f(A)$ is required per time step. The RA filter along with divergence damping (Skamarock and Klemp 1992) for integration of fluids problems with the leapfrog scheme provides computational

stability when incorporated in split fast- and slow-mode integrations. A smooth–de-smooth approach could be used in an attempt to recover the amplitude (Shuman 1957; Shapiro 1970; Perkey and Kreitzberg 1976).

b) The Williams filtered leapfrog (WFL) schemes

The RAW filter (Williams 2009, 2011, 2013) is a simple modification to the RA filter procedure to reduce the amplitude error. First, the prediction of a dependent variable, A, in time can be written in finite difference form as:

$$A_i^{n+1} = A_i^{n-1**} + 2\,\Delta t\,[\gamma f(A_i^{n*}) + (1-\gamma) f(A_i^n)]. \tag{2.5}$$

Here, the single and double asterisks denote values that have been filtered once and twice, respectively, using either the O(2) or O(4) filters defined below. A modification to the RA filter can be written using two O(2) temporal filters:

$$A_i^{n**} = A_i^{n*} + \frac{\nu\alpha}{2}\left(A_i^{n-1**} - 2A_i^{n*} + A_i^{n+1}\right) \tag{2.6}$$

and:

$$A_i^{n+1*} = A_i^{n+1} - \frac{\nu(1-\alpha)}{2}\left(A_i^{n-1**} - 2A_i^{n*} + A_i^{n+1}\right). \tag{2.7}$$

Importantly, this modification can be applied with very little additional coding and essentially no change in memory requirements or computational expense in existing models. Replacing the O(2) temporal filters with O(4) temporal filters (which invoke a longer time history) leads to:

$$A_i^{n**} = A_i^{n*} + \nu\alpha\left(A_i^{n-3**} - 4A_i^{n-2**} + 6A_i^{n-1**} - 4A_i^{n*} + A_i^{n+1}\right) \tag{2.8}$$

and:

$$A_i^{n+1*} = A_i^{n+1} - \nu(1-\alpha)\left(A_i^{n-3**} - 4A_i^{n-2**} + 6A_i^{n-1**} - 4A_i^{n*} + A_i^{n+1}\right). \tag{2.9}$$

The order of accuracy for amplitude errors of O(3) to O(7) is achieved by particular choices of α, γ, and ν and the choice of a O(2) or O(4) temporal filter, as shown in Table 1. The phase error is O(2) for all WFL schemes, making the overall formal accuracy of the WFL schemes O(2). Only one computation of $f(A)$ is

required per time step if $\gamma = 0$ or $\gamma = 1$, otherwise two computations of $f(A)$ are required. In practice, the O(2) temporal filter would be used for the first few time integration steps, until filtered values are available as far back as the n-3rd time step, when the O(4) temporal filter can be introduced and used for subsequent steps. Mode splitting with the WFL schemes can be accomplished exactly the same as it is with the RA-filtered leapfrog scheme (Skamarock and Klemp 1992).

c) The third-order Runge–Kutta schemes

The O(3) Runge–Kutta or RK3 scheme (Wicker and Skamarock 2002; called the "Linear-Case RK3" by Baldauf 2008, 2010) has become increasingly popular in atmospheric numerical modelling. It is somewhat similar to the O(3) Runge–Kutta Total Variation Diminishing (TVD) scheme (Shu 1988; Shu and Osher 1988) that is popular in some areas of computational fluid dynamics. The popularity of the strong-stability RK3 scheme owes in part to it being a forward-in-time scheme without a computational mode, having lower computational memory requirements, and having an improved stability regime based on the CFL criterion, such that the time step requirement is $c\Delta t/\Delta x < 1.73$ in one dimension (Wicker and Skamarock 2002) and $c\Delta t/\Delta x < 1.73/\sqrt{2}$ and $c\Delta t/\Delta x < 1.73/\sqrt{3}$ in two and three dimensions, respectively. The RK3 schemes have O(3) errors in time for linear problems, but O(2) errors for nonlinear problems (Baldauf 2008, 2010). These schemes require three functions to be computed per time step. Following Wicker and Skamarock (2002), the RK3 scheme is written as:

$$A_i^* = A_i^n + \frac{1}{3}\,\Delta t f(A_i^n) \tag{2.10}$$

$$A_i^{**} = A_i^n + \frac{1}{2}\,\Delta t f(A_i^*) \tag{2.11}$$

$$A_i^{n+1} = A_i^n + \Delta t f(A_i^{**}). \tag{2.12}$$

The O(3) Runge–Kutta TVD3 scheme (Shu and Osher 1988) is also a three-stage scheme and is given by:

$$A_i^* = A_i^n + \Delta t f(A_i^n) \tag{2.13}$$

$$A_i^{**} = \frac{3}{4}A_i^n + \frac{1}{4}A_i^* + \frac{1}{4}\,\Delta t f(A_i^*) \tag{2.14}$$

$$A_i^{n+1} = \frac{1}{3}A_i^n + \frac{2}{3}A_i^{**} + \frac{2}{3}\,\Delta t f(A_i^{**}). \tag{2.15}$$

It can be rewritten starting each stage at time *n* for mode splitting:

$$A_i^* = A_i^n + \Delta t f(A_i^n) \tag{2.16}$$

$$A_i^{**} = A_i^n + \frac{1}{4}\Delta t f(A_i^n) + \frac{1}{4}\,\Delta t f(A_i^*) \tag{2.17}$$

$$A_i^{n+1} = A_i^n + \frac{1}{6}\,\Delta t f(A_i^n) + \frac{1}{6}\,\Delta t f(A_i^*) + \frac{2}{3}\,\Delta t f(A_i^{**}). \tag{2.18}$$

Mode splitting of O(2) and O(3) Runge–Kutta schemes has been discussed by Wicker and Skamarock (1998, 2002) and Baldauf (2002, 2010). Both RK3 and TVD3 require divergence damping for stability when implemented with mode splitting. The RK3 scheme is used to produce the high-resolution reference solutions in this study because its order of accuracy is higher than that of the RA scheme. Specifically, RK3 is second-order accurate in time for nonlinear problems and third-order accurate in time for linear problems. Furthermore, RK3 is increasingly used in atmospheric science. With sufficient temporal and spatial resolution, RK3 is capable of providing nearly grid-converged solutions (as are all the schemes that are stable and consistent).

3. Results

Results from numerical simulations are shown in this section to assess the behaviour of the WFL schemes for two-dimensional linear and two- and three-dimensional nonlinear dry fluid flow problems. For each experiment, the numerical and physical details are summarised in Table 2. The WFL schemes that are tested include the one- and two-function O(3) and O(5) versions and the two-function O(7) version, made with combinations of two O(2) and two O(4) temporal filters, where the order achieved is the theoretical amplitude order of error for linear oscillations. Reference solutions are obtained using the O(1) RA-filtered leapfrog scheme and the RK3 scheme. The TVD3 schemes (Shu 1988; Shu and Osher 1988) were also used

to obtain reference solutions, but for the time steps employed in this study the RK3 and TVD3 solutions were very similar, and so the TVD3 solutions are not shown.

In the last three problems of Table 2, the forward–backward time-split scheme for velocity and pressure is implemented partly following Mesinger (1977) and Klemp and Wilhelmson (1978). First, the advection, diffusion, buoyancy, and filtering tendencies are computed using the long time step and stored. Next, the stored velocity tendencies are added linearly to the velocity fields on each small time step when the pressure gradient acceleration is integrated using the latest value of pressure. Then, the stored pressure tendencies are added to the pressure field on each small time step when the divergence term in the pressure equation is integrated using the latest velocity fields. This process is repeated for the required number of small steps, which is 12 for the leapfrog schemes and six for RK3.

a) Two-dimensional cone advection

The various time-stepping schemes are compared by using a series of two-dimensional numerical simulations of passive scalar advection. A cone of height 1 m and radius 8 m is linearly advected by an analytical wind field corresponding to solid-body rotation, where the cone centre is offset from the axis of rotation of the wind field. Equivalently, the system can be thought of as representing the advection of a passive tracer, whose non-dimensional concentration increases conically from zero outside a circle of radius 8 m to unity at the centre of the circle. This problem is known as rotational translation (Smolarkiewicz 1982; Schlesinger 1985; Tremback et al. 1987; Smolarkiewicz and Grabowski 1990; Blossey and Durran 2008). Other tests that were performed include linear translation, deformation, and a combination of rotational deformation and convergence. However, results from these tests are not shown for the sake of brevity, as they all behaved similarly to the rotational translation example (except for the deformation test, for which all schemes used herein performed poorly unless monotonic solutions were enforced by the spatial discretisation scheme or by an adjustment).

The analytical (exact) solution for the rotational translation of a cone is known, so that some exact error measures can be computed. The performance measures that will be used to assess the simulated cone height field include the domain-wide maximum and minimum scalar values from the numerical solution, the domain-wide root-mean-square error in the numerical solution relative to the exact solution, the dispersion and dissipation errors (which split up the domain-wide mean-square error into contributions from dispersion and dissipation; Takacs 1985), and the domain-wide L^{∞} error norm (i.e., the maximum absolute error in the numerical solution relative to the exact solution). The $|1-\Sigma A^2/\Sigma A_0^2|$ error measure will also be used (Tremback et al. 1987), where A is the scalar cone height field, A_0 is the exact solution, and both summations are taken over the entire simulation domain.

The domain extends from $x = 0$ to $x = L_x = 112$ m and from $y = 0$ to $y = L_y = 112$ m. The advective flow consists of clockwise rotation about the centre of the domain (i.e., $x = y = 56$ m). The initial condition is that the centre of the cone is located at $x = x_c = 84$ m and $y = y_c = 56$ m at $t = 0$. The analytical cone returns to its original location after each complete rotation period, which is 20π s. For the basic experiments, the grid spacing is $\Delta x = \Delta y = 2$ m and the time step is $\Delta t = 10\pi/628$ s. The effect of the time step is determined by using time steps of 2, 4, 8, 16, and 32 times smaller than the time step used for the basic experiments (i.e., $5\pi/628$, $2.5\pi/628$, $1.25\pi/628$, $0.625\pi/628$, and $0.3125\pi/628$ s). In addition, a smaller grid spacing of $\Delta x = \Delta y = 1$ m with time steps of 2 and 4 times smaller than the time step used in the basic experiments is also considered.

For the spatial differencing, the main experiments use the tenth-order constant grid flux scheme (Tremback et al. 1987). However, similar results were found for the tenth-order integrated flux scheme (not shown), although the results are slightly more accurate for the constant grid flux scheme. The tenth-order constant grid flux scheme collapses to the tenth-order advection scheme when the flow and grid spacing are constant (Tremback et al. 1987). The simulations shown and analysed here employ the constant grid flux scheme.

Using a high-order spatial discretisation helps ensure that the numerical errors are mostly due to temporal differencing rather than spatial differencing. Since the cone stays away from the domain boundaries, open boundaries are used without any problems (noting that analytical wind components are available outside the domain). No spatial diffusion or numerical filters are used in the simulations shown and analysed here. Tests were made with fourth-, sixth-, eighth-, and tenth-order spatial filters, including both non-monotonic formulations (Shapiro 1970; Purser 1987) and a monotonic formulation (Xue 2000). These results are not shown, as it was found that these filters had little effect on the general interpretation of the results, except (depending on the amount of smoothing) to produce slightly smoother solutions and better controlled Gibbs phenomena.

For both staggered and non-staggered grids, the values of the velocity components are needed at the edge of each grid zone for flux computations. These values may be obtained by averaging the velocities or by using the exact analytical velocities. The results presented herein are all produced using the averaging approach, which is common in atmosphere and ocean models. The averaging degrades the formal spatial accuracy to second order. It is logical that one should use spatial interpolation with an equivalent high-order scheme to match the spatial differencing scheme, in order to produce a formally high-order solution [personal communication, Bengt Fornberg, University of Colorado (2019) and Chi-Wang Shu, Brown University (2020)]. However, some tests (not shown) and some results shown in this paper lead to the conclusion that this approach does not always appreciably change the numerical solutions. This finding is probably due, at least in part, to the smooth (constant-in-time) flow fields considered with the cone rotation tests herein. Importantly, the interpolation order-of-accuracy effects appear to become more important with the cone problem using smaller time steps when smaller changes are resolved. Additionally, in cases with uneven and stretched grids (not shown), it has been found that it is important to use at least second-order spatial interpolation and not averages.

Figure 1 shows the simulated cone height field after exactly 12 complete rotations, made with the RA and RK3 temporal schemes and all the even-order constant grid flux schemes from second- to tenth-order, each using $\Delta x = \Delta y = 2$ m and $\Delta t = 0.625\pi/628$ s. Comparison with the analytical solution shows that the O(10) spatial differencing largely eliminates the spatial finite-difference truncation errors from the solution. As expected, the solutions for both RA and RK3 improve as the order of accuracy of the flux approximations is increased, with a cone height amplitude after 12 rotations (whose analytical value at any time is 1 m) of 0.77 and 0.81 m for the RA and RK3 eighth-order flux solutions and 0.82 and 0.88 m for the RA and RK3 tenth-order flux solutions. The RA and RK3 second-order flux simulations show little skill in the cone's location, shape, and amplitude retention after the first few rotations, and essentially no skill after eight rotations (not shown). There are still noticeable errors using the fourth- and sixth-order flux schemes. The cone shape appears only slightly distorted for the eighth- and tenth-order flux computations. There is no meaningful improvement found for twelfth- to sixteenth-order flux computations (not shown). From this set of experiments, it is concluded that O(10) flux computations limit the spatial differencing truncation errors sufficiently (with the temporal and spatial resolution employed for this problem) that the focus for all subsequent experiments (which will use O(10) flux computations) can be on the temporal differencing and time-step size.

A comparison of cone heights after exactly 12 complete rotations is shown in Figure 2 for the RA, RK3, and WFL (W03, WG3, W33, W43, W05, WG5, W55, and W77) schemes, each using $\Delta x = \Delta y = 2$ m. Included are examples produced using the RA scheme with five different time steps, ranging from $\Delta t = 10\pi/628$ s to $\Delta t = 0.625\pi/628$ s. The solutions for all the non-RA schemes are obtained using the largest RA time step ($\Delta t = 10\pi/628$ s). It is seen that as the time step gets smaller, the RA solutions converge to a nearly round cone. However, even with $\Delta t = 0.3125\pi/628$ s (not shown), the amplitude of the cone height with the RA scheme (0.86 m) is not as well preserved as it is with the RK3 and WFL schemes with $\Delta t = 10\pi/628$ s (all 0.87 m), despite their time step being 32 times as large. Moreover, the shape of the cone is well preserved with the RK3 and all the WFL schemes for all temporal resolutions considered. As Δt gets

very small, the RA cone shape only slowly approaches that produced by the RK3 and WFL schemes. In summary, minimal shape distortion is not attainable with RA unless the time step is reduced by a factor of at least 16, and even then, the amplitude preservation is still comparatively poor. On the other hand, even for the largest time steps, the WFL schemes produce only minor shape distortions that compare favourably with those produced by the RK3 scheme, with amplitude preservation that is as good as RK3, and all at one-third to two-thirds of the computational floating-point cost of RK3 for the same time-step size.

Various error measures for the cone height field after exactly 12 rotations are shown in Table 3 for the RA, RK3, W03, and W77 schemes. The W03 and W77 error measures are representative of the results from all the WFL schemes tested. The error measures for these two WFL schemes are very close, even with 64-bit computations. Note that in these cone advection simulations (but not in general) the L^∞ error norm is exactly equal to 1 m minus the maximum cone height value, because the maximum value and maximum error both occur at the grid point of the maximum value in the analytical solution. While all the WFL L^∞ error norms are generally acceptable, the best is obtained with the W77 scheme. The L^∞ error norms improve with the RA scheme as the time step is reduced. The RK3 L^∞ error norm is very similar to those from the WFL schemes. The root-mean-square errors and dispersion and dissipation errors all behave similarly to the L^∞ error norms for all schemes tested. The dispersion error exceeds the dissipation error for each scheme. For the RA solutions, the dispersion and dissipation errors are comparable in magnitude for the longest time step, but they grow an order of magnitude apart for the shortest time step. For the RK3 solution they are four orders of magnitude apart, and for the WFL solutions they are five orders of magnitude apart.

Figure 3 shows time series of the error measure $|1-\Sigma A^2/\Sigma A_0^2|$ for the cone height field up to 12 rotations. The RA solution with the longest time step is strongly damped in time, with the error measure reaching values of order unity after 12 rotations. Reducing the RA time step partly alleviates the damping, but even a time-step reduction by a factor of 16 does not prevent RA from having the largest errors of all the schemes studied. The RK3 solution has the next-largest errors and is also slightly damped, with the error measure

continuing to grow up to and beyond 12 cone rotations. Each of the WFL schemes performs better than RA and RK3, with significantly smaller errors and error growth rates. In particular, the W55, W77, and WG5 solutions show no detectable error growth between the first and 12th rotations, with the error measure after 12 rotations being an order of magnitude smaller than RK3. The values of $|1\text{-}\Sigma A^2/\Sigma A_0^2|$ for the pure (unfiltered) leapfrog scheme are also essentially unchanged for 1–12 and even 24–36 rotations (not shown). The error measures for W55, W77, and WG5 can therefore be considered essentially constant up to and beyond 12 cone revolutions using 32-bit computations, closely mimicking the behaviour of the pure leapfrog scheme. The best WFL scheme appears to be W33, in which the error measure actually decreases with time over the first ten rotations, reaching a value nearly three orders of magnitude smaller than the RK3 value.

Figure 4 shows additional simulations made with the RK3, RA, W03, and W77 schemes at higher spatial and temporal resolutions. Halving the grid length and time step (to preserve the Courant number) improves the solutions for all the schemes, as expected. Halving both the time step again (but now at fixed grid length) does not appreciably improve the RK3 or WFL simulations any further, but it does lead to further improvements in the RA simulations, which are still far from being converged. Note that all simulations have varying degrees of problems with the cone shape when coarser resolutions are used. These shape errors improve with finer spatial resolution and (to a lesser degree) with finer temporal resolution. The RA scheme is by far the worst performer, and the RK3 and WFL schemes are the best, even at the highest spatial and temporal resolutions analysed.

Since two-dimensional cone advection is the only problem in the present paper to have an analytical solution, we take the opportunity to explore the systematic relationships between time-step size, numerical error, and computational cost. We explore these relationships using simulations in which the spatial error is converged and the temporal error dominates. Figure 5 shows how the $|1\text{-}\Sigma A^2/\Sigma A_0^2|$ cone height error measure after 12 rotations varies with the time step. All simulations were run on the same computer using

the same compiler with the same compilation options. The RA scheme is clearly non-competitive, with errors that are orders of magnitude larger than any of the other schemes. The RA scheme also has poor (linear) convergence, as expected for a scheme with first-order formal accuracy. The RK3 scheme exhibits cubic scaling down to the smallest time step when double precision and tenth-order interpolation of the wind are used, as expected for a scheme with third-order formal accuracy (for linear systems; see Table 1). However, when the precision and order of interpolation of the wind are reduced, the RK3 error saturates before the smallest time step is reached, indicating that the computations can no longer resolve further error decreases, and that rounding errors from finite numerical precision now dominate the error measure. For the range of time steps over which temporal truncation errors dominate over rounding errors, the WFL schemes generally perform at least as well as RK3, with errors that are typically an order of magnitude (or more) smaller than RK3 at the largest time steps considered, even though their convergence is quadratic rather than cubic. Note that RK3 uses more computations than the WFL schemes – three times as many as WG3 and 1.5 times as many as W33 and W55 – and this increase in computations causes the error saturation to some degree.

Figure 5 gives useful insights into the empirical orders of accuracy of the schemes. However, a limitation of the figure is that the errors calculated for each time step are not controlled for computational expense. This limitation prevents us from using the figure to ascertain which schemes deliver the best accuracy for a given computational expenditure. To investigate this question, Figure 6 shows parametric plots of computational cost against error as the time-step size varies, allowing us to compare the computational costs of the schemes at given attained accuracies. The RA scheme costs many times more than any of the other schemes for any specified accuracy, making it a poor use of computational resource. For the range of time steps over which temporal truncation errors dominate over rounding errors, the second most expensive scheme is RK3, with the WFL schemes requiring the least computational effort to achieve the same accuracy. For example, for a specified tolerable error of 10^{-4}, the WFL schemes (W03, W33, WG3, W55, W05, WG5, W77) cost up to 70% less than RK3, making them the best use of computational resource.

To summarise, for the problem of pure linear advection of a scalar by a constant solid-body rotation flow, the RK3 and WFL schemes have comparable errors for any given time-step size. RK3 and WFL each significantly out-perform the O(1) RA scheme, probably because (for linear problems) their orders of accuracy are O(3) and O(2), respectively. Note, however, that RK3 is more computationally expensive than the WFL schemes for a given time-step size. When this factor is controlled for in the comparisons, we find that all the WFL schemes deliver better accuracy (i.e., smaller errors) than RK3 for a given computational expenditure. No particular WFL scheme notably out-performs any of the other WFL schemes, perhaps because these schemes are all inherently O(2). The RA scheme, with its O(1) overall temporal errors, exhibits the worst performance.

b) Two-dimensional density current

This section describes the use of the quasi-compressible fluid equations to simulate a nonlinear two-dimensional dry density current, with a constant O(2) eddy diffusion coefficient (K_m = 75 m^2 s^{-1}; Straka et al. 1993; Straka and Anderson 1993). The goals are to examine how the WFL schemes perform for a more complicated (but still idealised) flow and how stable the WFL schemes are with mode splitting. The initial conditions include a cold elliptical fluid region, with horizontal radius 4 km and vertical radius 2 km, containing a cosine-squared potential temperature anomaly of minimum value -15°C at the centre of the ellipse. The cold ellipse is released from rest, with its centre initially 3 km above a lower boundary and halfway across a horizontally periodic domain of horizontal extent 51.2 km. There is an upper boundary 6.4 km above the lower boundary. As time increases, the cold region descends to the lower boundary under gravity, and then spreads out symmetrically in both horizontal directions along the boundary as a density current.

When selecting boundary conditions to impose at the upper and lower boundaries in two-dimensional density current simulations, many previous studies (e.g., Straka et al. 1993) have chosen to set $du/dz = dw/dz = d\theta/dz = 0$ for the horizontal velocity component u, vertical velocity component w, and potential temperature θ. However, a few studies (e.g., Klemp and Wilhelmson 1978) have chosen to effectively set the vertical diffusion to zero, by applying $d^2u/dz^2 = d^2w/dz^2 = d^2\theta/dz^2 = 0$. To facilitate comparisons with the widest possible range of previous studies, here we choose to impose the former (Straka et al. 1993) boundary conditions at the upper and lower boundaries, because they seem to be more commonly used in the existing literature.

In contrast to the linear cone advection problem, there is no analytical solution for the nonlinear density current problem. Therefore, we must resort to using a high-resolution numerical simulation as the reference solution. Furthermore, a rigorous assessment of numerical convergence is more difficult with the staggered C-grid (as used here) than with the non-staggered A-grid used by Straka et al. (1993). Flow features, including their maximums and minimums, are found in different locations for different resolutions and schemes, meaning that the computation of errors against a high-resolution reference solution requires one of the solutions to be interpolated. We find that interpolating either solution gives essentially the same temporal and spatial derivative convergence rates as in Straka et al. (1993; their Figure 5) for the perturbation potential temperature simulated using the RA scheme, and therefore these results are not shown. Instead, perturbation potential temperature field differences are evaluated with respect to a high-order temporal scheme (here, RK3) with the same time step and spatial resolution as the test schemes. The maximums and minimums of perturbation potential temperature relative to the RK3 reference solution with high spatial resolution can also be compared. Finally, changes in the total energy and total mass in each simulation since $t = 0$ are analysed (Bryan and Fritsch 2002), as well as the domain-wide $\Sigma\theta'^2/\Sigma\theta'^2_0$ error measure for the perturbation potential temperature, where the subscript 0 here refers to $t = 0$ rather than to a reference solution (Tremback et al. 1987).

Figure 7 shows the high-resolution RK3 reference solution after 900 s, by which time the cold anomaly has descended to the lower boundary and begun spreading out horizontally. The time step for this reference solution is $\Delta t = 0.0078125$ s and the grid spacing is $\Delta x = \Delta z = 12.5$ m. Similar solutions with the same high spatial and temporal resolution using RA, WG3, and W77 (not shown) appear very similar to the RK3 reference solution. The maximum and minimum values differ between these four solutions by around 0.2% for the horizontal velocity, 0.4% for the vertical velocity and perturbation pressure, and less than 0.001% for the perturbation potential temperature. This level of agreement is expected for a nearly converged solution using a large eddy diffusion coefficient ($K_m = 75$ m^2 s^{-1}) and high resolution in both space and time. In addition, these solutions are all produced with the same spatial differencing routines and pressure solvers. In contrast, most of the solutions presented in Straka et al. (1993) were made using multiple models with various advection and pressure solvers.

Figure 8 shows the perturbation potential temperature fields obtained with the RK3, RA, and WFL schemes using $\Delta t = 0.0625$ s (although this time step was reduced by 25% to ensure stability for the W43 scheme) and $\Delta x = \Delta z = 100$ m. Unexpectedly, the various solutions are all very similar to each other, perhaps partly because we used half the time step used by Straka et al. (1993). (The present solutions are produced with tenth-order spatial differences, reducing the maximum stable Courant number to around 0.5383 times its value for second-order spatial differences; Straka and Anderson 1993.) More importantly, the diffusion term, which is integrated with a time-lagged forward-in-time scheme for stability in the RA and WFL solutions, has a significant impact on the solutions (see below). In most cases, similar perturbation potential temperature fields (with differences less than 0.2 K) can be obtained with $\Delta t = 0.4$ s, the exceptions being that RK3 can use $\Delta t > 1.0$ s, but W05, WG5, W43, and W77 all need to use $\Delta t < 0.125$ s.

Figure 9 shows the temporal evolution of the total (i.e., kinetic plus potential) energy and the total fluid mass in the RK3, RA, and W77 simulations. The other WFL simulations are omitted, because they behave nearly identically to W77. Although the analytical solution for this problem does not exist, we do know that

it must conserve the total energy and mass exactly, and so it is useful to test whether the numerical solutions also respect these conservation properties. In general, all the schemes perform quite well, conserving the initial energy and mass to within around 10^{-4}% over the first 900 s (as well as up to 3,600 s; not shown). However, when examining differences relative to RK3, W77 gains an order of magnitude less energy than RA loses. The total mass for all schemes oscillates spuriously with a period of around 40 s, as wave emission associated with the adjustment to the unbalanced initial condition challenges the numerical schemes. Superimposed on these high-frequency oscillations is a longer-term variation, in which (again relative to RK3) W77 loses an order of magnitude less mass than RA gains. We conclude that each numerical scheme tested delivers acceptable energy and mass conservation, with the WFL schemes generally performing much better than the RA scheme compared to the RK3 reference solution.

To explore the impacts of diffusion, additional simulations were run with the wind and perturbation pressure set to zero, leaving diffusion of potential temperature as the only process operating. Two different values of the eddy diffusion coefficient (K_m = 75 and 30 m^2 s^{-1}) were used. All ten schemes listed in Table 1 were tested, using Δt = 0.0625 s (reduced by 25% for W43) and $\Delta x = \Delta z$ =100 m, and the quantitative diagnostics that will be discussed here were similar across the schemes. These simulations are not shown graphically, but they allow the inference that diffusion is important in changing the potential temperature fields, with the values of $\Sigma\theta'^2/\Sigma\theta'^2_0$ at t = 900 s being around 0.86 (i.e., a 14% reduction since t = 0) and 0.94 (a 6% reduction) when K_m = 75 and 30 m^2 s^{-1}, respectively. However, the full simulations show that there are much larger changes when advection and diffusion are coupled together, with the values of $\Sigma\theta'^2/\Sigma\theta'^2_0$ at t = 900 s being around 0.34 (a 66% reduction) and 0.43 (a 57% reduction) when K_m = 75 and 30 m^2 s^{-1}, respectively. Whether diffusion acts alone or is coupled with advection, the impacts on potential temperature are less when K_m is reduced, as expected. When K_m is reduced from 75 m^2 s^{-1} to 30 m^2 s^{-1}, the ratio of the domain-mean absolute advection tendency to the domain-mean absolute diffusion tendency at t = 900 s approximately doubles for each of θ', u, and w, indicating that the advection becomes increasingly important when the diffusion is reduced.

In summary, the two-dimensional density current simulations are diffusion-limited by design (Straka et al. 1993). Therefore, the numerical solutions have similar errors regardless of the temporal integration scheme. The use of a collapsing flux stencil as the vertical boundaries are approached also could be a factor controlling or limiting the spatial and temporal accuracy, as could the O(1)-in-time and O(2)-in-space approximations for the diffusion and the O(2)-in-time and O(2)-in-space approximations for the pressure gradient and divergence. All the schemes tested conserve energy and mass to within around 10^{-4}%, but the WFL schemes are an order of magnitude closer to the high-resolution RK3 reference solution than the RA scheme is. Importantly, our results show that each of the WFL schemes appears to be stable when using split fast and slow modes. However, the WFL schemes that employ fourth-order temporal filters require smaller time steps to remain stable with mode splitting.

c) Three-dimensional rising warm plume

A nonlinear warm thermal current (i.e., rising warm bubble) with constant eddy diffusion coefficient (K_m = 2.0 to 7.5 m^2 s^{-1}) is used to examine the behaviour of the WFL schemes in a three-dimensional, fully compressible model. Unlike the density current experiments, the rising plume set-up avoids the complications of the main feature of interest interacting with an impermeable horizontal boundary (Carpenter et al. 1990; Tripoli 1992; Wicker and Skamarock 1998; Bryan and Fritsch 2002). The initial conditions include a warm sphere of radius 2 km, which contains a cosine-squared potential temperature anomaly of maximum value +2℃ at the centre of the sphere. The warm sphere is released from rest in the middle of a doubly periodic horizontal domain measuring 12 km by 12 km, with its centre initially 2 km above the lower boundary. As time increases to 1,000 s, the warm sphere deforms and develops a rotor as it ascends towards (but by 1,000 s does not reach) the upper boundary, which is at a height of 10 km. The evolving warm plume is axisymmetric about the vertical axis defined by $x = y = 6$ km, but the model equations being numerically integrated are fully three-dimensional.

Figure 10 shows the numerical solutions at 1,000 s for the RK3, RA, W03, WG3, W33, W05, WG5, W55, and W77 schemes. These solutions were made using an eddy diffusion coefficient of $K_m = 7.5$ m^2 s^{-1}, which is large enough to dampen the formation of shear instabilities at the top and sides of the plume (Grabowski and Clark 1991; Tripoli 1992). The solutions were made using $\Delta x = \Delta y = \Delta z = 100$ m and mostly using $\Delta t = 0.4$ s. However, a smaller time step of $\Delta t = 0.05$ s was found necessary to ensure stability for the WFL schemes that incorporate fourth-order temporal filters (W05, WG5, and W77). The axisymmetric radial overturning circulation, which results from stronger convection at the centre of the plume than the edges, and which thereby leads to the rotor formation, is clearly visible in Figure 10. In the absence of an analytical solution, a high-resolution RK3 simulation with $\Delta x = \Delta y = \Delta z = 25$ m and $\Delta t = 0.1$ s is used as a reference solution for comparison.

The rotors seen in Figure 10 are handled well (compared to the RK3 reference solution) by most of the schemes, and perhaps best by the RK3 and W33 test solutions. The maximum perturbation potential temperatures are systematically 6–12% higher in the RK3, RA, and WFL test solutions than the RK3 reference solution. Additionally, the maximum vertical velocities are 8–10% stronger for most of the test solutions (not shown). In Figure 10, the negative perturbation potential temperatures seen in the ambient fluid surrounding the plume are unrealistic artefacts of purely numerical origin. They occur largely due to spatial truncation errors, and they are around ten times smaller in amplitude in the RK3 reference solution.

Figure 11 shows simulations made using a smaller eddy diffusion coefficient ($K_m = 2.0$ m^2 s^{-1}) and finer numerical resolution ($\Delta x = \Delta y = \Delta z = 50$ m, $\Delta t = 0.2$ s), resulting in somewhat different dynamical behaviour at 1,000 s, which is described in detail by Grabowski and Clark (1991) and Tripoli (1992). The feedback associated with the dynamical instability can be seen in all the state variables for the high-resolution RK3 reference simulation ($\Delta x = \Delta y = \Delta z = 25$ m and $\Delta t = 0.1$ s) and in the perturbation potential temperature field for the test solutions made using the RK3 and RA schemes and the four WFL schemes

that incorporate second-order temporal filters (W03, WG3, W33, and W55). The use of a smaller eddy diffusion coefficient allows for advection to become increasingly important, and for shear instabilities to become more prominent, owing to sharper temperature and velocity gradients. For the perturbation potential temperature field, the test solutions that appear most similar to the high-resolution RK3 reference solution include the W55 scheme and perhaps the RK3 scheme, especially in the regions of developing instabilities on the leading edge of the plume. The maximum values of the perturbation potential temperature are 2–7% too high in all the test solutions compared to the reference solution. The positions of the structures on the leading edge of the plume in the high-resolution RK3 reference solution are similar to those in the W55 test solution, although the maximum perturbation potential temperature in the W55 solution is slightly higher. Reducing the time step of the test solutions to one-eighth of the value used above results in a substantial increase in the plume structure on its leading edge (not shown).

The results shown here for the rising warm plume indicate that the largest differences seem to be associated with the temporal order of accuracy of the schemes that are used. The differences using different time-stepping schemes become more apparent with higher spatial and temporal resolution when diffusion becomes increasingly less important and advection becomes increasingly more important. Here, diffusion is advanced with a first-order forward-in-time finite difference, whereas advection is integrated with second-order centred-in-time differences (except that they are first-order for RA). Further reduction of the eddy diffusion coefficient, or using a weak tenth-order spatial filter, along with smaller spatial resolutions, results in even more enhanced shear instabilities and even more differences (consistent with Grabowski and Clark 1991 and Tripoli 1992) when compared to the solutions shown here.

d) Two-dimensional twin tropical cyclones

The final set of benchmark experiments considers a large-scale atmospheric problem, in which the Coriolis force is critical. The problem is the linked behaviour of twin tropical cyclones (Moustaoui et al. 2002)

modelled using the two-dimensional rotating shallow-water equations cast on a β-plane (Ferreira et al. 1996; Moustaoui et al. 2002; Moustaoui et al. 2014). In the absence of the Coriolis force, if the two binary vortices are close enough together and strong enough, then they can interact via the Fujiwhara effect. In the presence of the Coriolis force, however, the twin cyclones can mutually and symmetrically affect each other, with the anticlockwise cyclone initialised north of the equator moving to the north-west and the clockwise cyclone initialised south of the equator moving to the south-west. This type of coupled cyclone motion results from a secondary asymmetric circulation, which is linked directly to the meridional gradient of planetary vorticity (i.e., the β-effect). Without the Coriolis force, or on an *f*-plane rather than a β-plane, this deviate motion does not occur.

The numerical model used here is based on rotating shallow-water equations that permit divergence and that include the full, un-approximated latitudinal variation of the Coriolis force (Anthes and Hoke 1975; Ferreira et al. 1996; Moustaoui et al. 2002; Moustaoui et al. 2014). The simulation domain is 9,600 km wide zonally and meridionally, centred on the equator. Vorticity centres with opposite signs are initialised at 7°N and 7°S (i.e., around 780 km either side of the equator), with rotational speeds decaying exponentially with increasing radius from their maximum value of 30 m s^{-1}. The simulations are run for four days (with tests in larger domains up to 12 days). A reference solution is made with the RK3 scheme, using temporal and spatial resolutions of $\Delta x = \Delta y = 20$ km and $\Delta t = 7.5$ s, whereas the test solutions are made using $\Delta x = \Delta y = 40$ km and $\Delta t = 15$ s. The smaller time step allows for stable simulations for all schemes.

The solutions are shown in Figure 12 and indicate that the cyclones migrate westward and move further apart over the first four days. Although there is a known weak instability with forward-type schemes (such as forward explicit Euler, RK3, TVD, Lax–Wendroff, etc.; Gadd 1980; Wicker and Skamarock 2002), no such instabilities are detected here in the high-resolution reference solution when advancing the Coriolis terms with the RK3 scheme, whether using a forward-in-time or forward–backward-in-time operator for

the fast (gravity wave) modes. The other schemes tested at a coarser resolution (RK3, RA, W03, WG3, W33, W05, WG5, W55, W77) each performed remarkably well compared to the reference solution after four days (96 h) of integration, in terms of the maxima and minima of the vertical vorticity and zonal and meridional winds. The extent to which the simulations agree is dependent on the coefficient of the sixth-order filter that is employed, with older studies using second-order filters but more recent studies using sixth-order filters. Longer simulations, up to 12 days, also produce remarkably good scheme-to-scheme solution appearance and errors (with respect to the maximum and minimum values of all fields), as well as good qualitative comparisons with the results of Moustaoui et al. (2014) at 50 and 100 h of simulation.

4. Summary and discussion

The stabilisation of the leapfrog time-stepping scheme provided by the RA filter has enabled large volumes of productive atmospheric and oceanic research and decades of successful weather and climate prediction outcomes. However, by spuriously damping the physical mode of the leapfrog scheme, our results show that the RA filter can unphysically distort dynamical fields and that it violates conservation of energy and mass more egregiously than the RK3 and WFL schemes. Only very small time steps reduce the distortion and non-conservation to acceptable levels. The distortion is particularly pronounced when the advection tendency plays an important role, such as in the linear two-dimensional cone rotation experiments. It is also particularly pronounced when weak eddy diffusion (either physical or numerical) is used in more complex nonlinear models, such as the density current and rising warm plume experiments, especially when centred-in-time differences are used for advection and forward-in-time differences are used (as required for stability) for diffusion. For example, we find that the O(2) RK3 and WFL solutions are somewhat better than the O(1) RA-filtered leapfrog solutions for the two-dimensional density current (with an eddy diffusion coefficient of $K_m = 75$ m^2 s^{-1}) and the three-dimensional rising warm plume (with $K_m = 7.5$ m^2 s^{-1}). When K_m is reduced by a factor of around two or more, which allows the advection to become more important

compared to the diffusion, we find that the problems in the RA-filtered leapfrog scheme worsen when compared to RK3 and WFL.

The results presented in this paper confirm that the WFL schemes that were formally shown to be linearly stable in Williams (2009, 2011, 2013) are each stable for more complex two- and three-dimensional systems of nonlinear equations for quasi-compressible and fully compressible fluid flows. Furthermore, the WFL schemes remain stable even when the discretisation method splits the slow advection and diffusion modes from the fast acoustic modes. The time-step size, spatial filtering, and divergence damping requirements of the WFL schemes that use second-order filters are generally similar to those of the RA scheme (although divergence damping is not necessarily needed for time-filtered schemes for the simulations in this paper). Importantly, the total energy errors and total mass errors are much smaller in magnitude (and consistent with the RK3 scheme) for all the WFL schemes than for the RA-filtered leapfrog scheme. The simplicity of the WFL schemes – being explicit rather than implicit – is attractive from the scalability point of view for future exa-scale weather and climate simulations.

As noted by Williams (2009, 2011, 2013), simple code modifications (to only a handful of lines) allow an O(1) amplitude error and O(1) overall error RA-filtered leapfrog code to be upgraded to an O(3) amplitude error and O(2) overall error WFL code. Such upgrades eliminate many of the problems of the RA-filtered leapfrog scheme and incur trivial additional computational cost and storage requirements, especially if the one-function WFL schemes are used. The two-function WFL schemes require more code modifications, but not prohibitively. The two-function WFL schemes may further improve simulations, especially for pure (or nearly pure) advection problems, depending on the order of accuracy of the WFL scheme used. Exploitation of variables while they are in the fastest layers of computer cache to the largest extent possible can improve computational efficiency more than expected. When large diffusion overwhelms any temporal truncation error improvement, such as in the present density current simulations, the two-function WFL schemes appear to provide no additional significant improvement beyond that provided by either the one-

function W03 or WG3 schemes. However, the two-function schemes do provide additional improvements when diffusion is not as important, as in the present rising warm plume experiments.

The one-function WFL schemes incur about one-third of the computational cost of the RK3 scheme per time step for the same time-step size, and yet they have a similar accuracy profile. However, the one-function WFL schemes (especially those using a fourth-order temporal filter) require a smaller maximum time step and use more memory. Nevertheless, even with the 73% longer time step permitted by the RK3 scheme compared to the pure leapfrog scheme (Durran 1991), the one-function WFL schemes may still be more cost-effective overall, given the reduced computational expense per time step. The two-function WFL schemes cost more than the one-function WFL schemes per time step, but they are still only about two-thirds of the cost of the RK3 scheme per time step for the same time-step size. Note also that in the leapfrog (RA and WFL) schemes with mode splitting, the small step sequence is looped through once per time step (with N small steps per time step). On the other hand, the RK3 scheme loops through the small step sequence three times per time step (with $N/3$, $N/2$, and N small steps). For example, when $N = 12$, there are a total of 22 small step calculations per time step for RK3, compared to only 12 for the leapfrog schemes. However, when $N = 12$ is used for the centred schemes, $N = 6$ can be used for RK3 (because it is a forward scheme), so there could be a total of only 11 small step calculations per time step.

When obtaining mode-split numerical solutions, the WFL schemes that use a fourth-order temporal filter were found to require reduced time steps to remain stable, compared to the WFL schemes that use a second-order temporal filter. This restriction was not found when the fourth-order temporal filter WFL schemes were used to obtain non-mode-split solutions, such as in the case of pure advection of a cone. This restriction is irrelevant for all WFL schemes when fast and slow modes were not split and integrated explicitly. Fortunately, it has been found that this problem can be alleviated in time-split simulations by integrating the velocity advection terms using the smaller time step, as has been done previously for the Crowley scheme (Tremback et al. 1987; Walko and Avissar 2008), making the integration of many scalar variables

efficient (because they do not need the reduced step) when the fourth-order temporal filter WFL schemes are used. In particular, the Walko and Avissar (2008) procedure permits stable and long integrations (longer than 4 days) as described with the twin tropical cyclones and other simulations not shown in this paper. The reasons for the reduced time step requirements of the fourth-order temporal filter are unknown and warrant further investigation.

The comparisons presented in this paper allow some provisional recommendations to be made regarding the most suitable WFL schemes for various purposes. For flows in which the diffusion is relatively small or zero, such as the advection-dominated flow above the planetary boundary layer in the free atmosphere, our linear advection and rising warm plume experiments suggest that the two-function W33 and W55 schemes deliver the best performance. In contrast, for diffusion-dominated flows, such as deep convection or the turbulent flow within the planetary boundary layer, our density current experiments suggest that the one-function W03 and WG3 schemes are sufficient. Our time-split experiments suggest that W05, WG5, W44, and W77 may be poorly suited to high-resolution non-hydrostatic simulations, because these schemes have limited stability with split modes.

One of the limitations of this study is that the more complex simulations are run for relatively short integration periods, compared to the typical usage of an atmosphere or ocean model. Consequently, it remains unclear if the WFL schemes are more suited to short-range numerical weather prediction forecasts or long-term climate projections. Longer simulations might produce more degradation of the RA-filtered leapfrog solution compared to the WFL and RK3 solutions, as seen in the two-dimensional cone advection simulations. Another limitation is that the results do not include comparisons to the filtered leapfrog schemes proposed by Moustoui et al. (2014) or Maurya et al. (2019), or to the forward temporal treatments used in the higher-order Crowley schemes (Tremback et al. 1987). The stability of the mode-split WFL schemes also remains to be formally shown. We call for future work to address these omissions.

Data availability statement

All source codes and model documentation are available from the authors.

References

Amezcua, J., Kalnay, E., & Williams, P. D. (2011) The effects of the RAW filter on the climatology and forecast skill of the SPEEDY model. *Monthly Weather Review*, **139**(2), 608–619. https://doi.org/10.1175/2010MWR3530.1

Amezcua, J., & Williams, P. D. (2015) The composite-tendency Robert–Asselin–Williams (RAW) filter in semi-implicit integrations. *Quarterly Journal of the Royal Meteorological Society*, **141**(688), 764–773. https://doi.org/10.1002/qj.2391

Anthes, R. A., & Hoke, J. E. (1975) The effect of horizontal divergence and the latitudinal variation of the Coriolis parameter on the drift of a model hurricane. *Monthly Weather Review*, **103**(9), 757–763. https://doi.org/10.1175/1520-0493(1975)103<0757:TEOHDA>2.0.CO;2

Asselin, R. (1972) Frequency filter for time integrations. *Monthly Weather Review*, **100**(6), 487–490. https://doi.org/10.1175/1520-0493(1972)100<0487:FFFTI>2.3.CO;2

Babarsky, R. J., & Sharpley, R. (1997) Expanded stability through higher temporal accuracy for time-centered advection schemes. *Monthly Weather Review*, **125**(6), 1277–1295. https://doi.org/10.1175/1520-0493(1997)125<1277:ESTHTA>2.0.CO;2

Baldauf, M. (2002) *Integration by time-splitting in mesoscale models. In: Doms G. Schättler U, editors. COSMO Newsletter No. 2;2002. P. 101-4.* COSMO Newsletter No. 2 p. 101-104. http://www.cosmo-model.org/content/model/documentation/newsLetters/newsLetter02/newsLetter_02.pdf

Baldauf, M. (2008) Stability analysis for linear discretisations of the advection equation with Runge–Kutta time integration. *Journal Computational Physics*, **227**(13), 6638–6659. https://doi.org/10.1016/j.jcp.2008.03.025

Baldauf, M. (2010) Linear stability analysis of Runge–Kutta-based partial time-splitting schemes for the Euler equations. *Monthly Weather Review*, **138**(12), 4475–4496. https://doi.org/10.1175/2010MWR3355.1

Blossey, P. N., & Durran, D. R. (2008) Selective monotonicity preservation in scalar advection. *Journal Computational Physics*, **227**(10), 5160–5183. https://doi.org/10.1016/j.jcp.2008.01.043

Bryan, G. H., & Fritsch, J. M. (2002) A benchmark simulation for moist nonhydrostatic numerical models. *Monthly Weather Review*, **130**(12), 2917–2928. https://doi.org/10.1175/1520-0493(2002)130<2917:ABSFMN>2.0.CO;2

Carpenter, R. L., Droegemeier, K. K., Woodward, P. R., & Hane, C. E. (1990) Application of the Piecewise Parabolic Method (PPM) to meteorological modeling. *Monthly Weather Review*, **118**(3), 586–612. https://doi.org/10.1175/1520-0493(1990)118<0586:AOTPPM>2.0.CO;2

Deardorff, J. W. (1980) Stratocumulus-capped mixed layers derived from a three-dimensional model. *Boundary-Layer Meteorology*, **18**(4), 495–527. https://doi.org/10.1007/BF00119502

Durran, D. R. (1991) The third-order Adams–Bashforth method: an attractive alternative to leapfrog time differencing. *Monthly Weather Review*, **119**(3), 702–720. https://doi.org/10.1175/1520-0493(1991)119<0702:TTOABM>2.0.CO;2

Durran, D. R. (1999) *Numerical Methods for Wave Equations in Geophysical Fluid Dynamics*. Springer-Verlag.

Ferreira, R. N., Schubert, W. H., & Hack, J. J. (1996) Dynamical aspects of twin tropical cyclones associated with the Madden–Julian Oscillation. *Journal of Atmospheric Sciences*, **53**(7), 929–945. https://doi.org/10.1175/1520-0469(1996)053<0929:DAOTTC>2.0.CO;2

Gadd, A. J. (1980) Two refinements of the split explicit integration scheme. *Quarterly Journal of the Royal Meteorological Society*, ***106*** (447), 215–220. https://doi.org/10.1002/qj.49710644715

Grabowski, W. W., & Clark, T. L. (1991) Cloud–environment interface instability: rising thermal calculations in two spatial dimensions. *Journal of Atmospheric Sciences*, **48**(4), 527–546. https://doi.org/10.1175/1520-0469(1991)048<0527:CIIRTC>2.0.CO;2

Holland, G.J., & Dietachmayer, G.S. (1993) On the interaction of tropical-cyclone-scale vortices. III: Continuous barotropic vortices. *Quarterly Journal of the Royal Meteorological Society*, **119**(514), 1381–1398. https://doi.org/10.1002/qj.49711951408

Klemp, J. B., & Wilhelmson, R. B. (1978) The simulation of three-dimensional convective storm dynamics. *Journal of Atmospheric Sciences*, **35**(6), 1070–1096. https://doi.org/10.1175/1520-0469(1978)035<1070:TSOTDC>2.0.CO;2

Li, Y., & Trenchea, C. (2014) A higher-order Robert–Asselin type time filter. *Journal Computational Physics*, **259**, 23–32. https://doi.org/10.1016/j.jcp.2013.11.022

Lorenz, E. N. (1963) Deterministic nonperiodic flow. *Journal of Atmospheric Sciences*, **20**(2), 130–141. https://doi.org/10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2

Maurya, P. K., Rajpoot, M. K., & Yadav, V. S. (2019) Higher-order optimized hybrid Robert-Asselin type time filters. *Journal Computational Physics*, **399**, 108941. https://doi.org/10.1016/j.jcp.2019.108941

Mesinger, F. (1977) Forward–backward scheme, and its use in a limited-area model. *Beiträge zur Physik der Atmosphäre (Contributions to Atmospheric Physics)*, **50**, 200–210.

Moustaoui, M., Teitelbaum, H., Basdevant, C., & Boughaleb, Y. (2002) Linked behavior of twin tropical cyclones. *Journal of Geophysical Research: Atmospheres*, **107**(D19), 4378. https://doi.org/10.1029/2000JD000066

Moustaoui, M., Mahalov, A., & Kostelich, E. J. (2014) A numerical method based on leapfrog and a fourth-order implicit time filter. *Monthly Weather Review*, **142**(7), 2545–2560. https://doi.org/10.1175/MWR-D-13-00073.1

Perkey, D.J., & Kreitzberg, C.W. (1976) A time-dependent lateral boundary scheme for limited-area primitive equation models. *Monthly Weather Review*, *104*, 744–755. https://doi.org/10.1175/1520-0493(1976)104<0744:ATDLBS>2.0.CO;2

Purser, R. J. (1987) The filtering of meteorological fields. *Journal of Applied Meteorology and Climatology*, **26**(12), 1764–1769. https://doi.org/10.1175/1520-0450(1987)026<1764:TFOMF>2.0.CO;2

Ren, D., & Leslie, L. M. (2011) Three positive feedback mechanisms for ice-sheet melting in a warming climate. *Journal of Glaciology*, **57**(206), 1057–1066. https://doi.org/10.3189/002214311798843250

Robert, A. J. (1966). The integration of a low order spectral form of the primitive meteorological equations. *Journal of the Meteorological Society of Japan. Ser. II*, **44**(5), 237–245. https://doi.org/10.2151/jmsj1965.44.5_237

Schlesinger, R. E. (1985) Effects of upstream-biased third-order space correction terms on multidimensional Crowley advection schemes. *Monthly Weather Review*, **113**(7), 1109–1130. https://doi.org/10.1175/1520-0493(1985)113<1109:EOUBTO>2.0.CO;2

Schlesinger, R. E., Uccellini, L. W., & Johnson, D. R. (1983) The effects of the Asselin time filter on numerical solutions to the linearized shallow-water wave equations. *Monthly Weather Review*, **111**(3), 455–467. https://doi.org/10.1175/1520-0493(1983)111<0455:TEOTAT>2.0.CO;2

Shapiro, R. (1970) Smoothing, filtering, and boundary effects. *Reviews of Geophysics*, **8**(2), 359–387. https://doi.org/10.1029/RG008i002p00359

Shu, C.-W. (1988) Total-variation-diminishing time discretizations. *Society for Industrial and Applied Mathematics. SIAM Journal on Scientific and Statistical Computing*, **9**(6), 12. https://doi.org/10.1137/0909073

Shu, C.-W., & Osher, S. (1988) Efficient implementation of essentially non-oscillatory shock-capturing schemes. *Journal of Computational Physics*, **77**(2), 439–471. https://doi.org/10.1016/0021-9991(88)90177-5

Shuman, F.G. (1957) Numerical methods in weather prediction: II. Smoothing and filtering. *Monthly Weather Review*, *85*(11), 357–361. https://doi.org/10.1175/1520-0493(1957)085<0357:NMIWPI>2.0.CO;2

Skamarock, W. C., & Klemp, J. B. (1992) The stability of time-split numerical methods for the hydrostatic and the nonhydrostatic elastic equations. *Monthly Weather Review*, **120**(9), 2109–2127. https://doi.org/10.1175/1520-0493(1992)120<2109:TSOTSN>2.0.CO;2

Smolarkiewicz, P. K. (1982) The multi-dimensional Crowley advection scheme. *Monthly Weather Review*, **110**(12), 1968–1983. https://doi.org/10.1175/1520-0493(1982)110<1968:TMDCAS>2.0.CO;2

Smolarkiewicz, P. K., & Grabowski, W. W. (1990) The multidimensional positive definite advection transport algorithm: Nonoscillatory option. *Journal Computational Physics*, **86**(2), 355–375. https://doi.org/10.1016/0021-9991(90)90105-A

Straka, J. M., Wilhelmson, R. B., Wicker, L. J., Anderson, J. R., & Droegemeier, K. K. (1993) Numerical solutions of a non-linear density current: A benchmark solution and comparisons. *International Journal for Numerical Methods in Fluids*, **17**(1), 1–22. https://doi.org/10.1002/fld.1650170103

Straka, J. M., & Anderson, J. R. (1993) Extension and application of a local, minimum aliasing method to multidimensional problems in limited-area domains. *Monthly Weather Review*, **121**(10), 2903–2918. https://doi.org/10.1175/1520-0493(1993)121<2903:EAAOAL>2.0.CO;2

Takacs, L. L. (1985) A two-step scheme for the advection equation with minimized dissipation and dispersion errors. *Monthly Weather Review*, **113**(6), 1050–1065. https://doi.org/10.1175/1520-0493(1985)113<1050:ATSSFT>2.0.CO;2

Tremback, C. J., Powell, J., Cotton, W. R., & Pielke, R. A. (1987) The forward–in-time upstream advection scheme: extension to higher orders. *Monthly Weather Review*, **115**(2), 540–555. https://doi.org/10.1175/1520-0493(1987)115<0540:TFTUAS>2.0.CO;2

Tripoli, G. J. (1992) A nonhydrostatic mesoscale model designed to simulate scale interaction. *Monthly Weather Review*, **120**(7), 1342–1359. https://doi.org/10.1175/1520-0493(1992)120<1342:ANMMDT>2.0.CO;2

Walko, R. L., & Avissar, R. (2008) The Ocean–Land–Atmosphere Model (OLAM) Part II: Formulation and Tests of the Nonhydrostatic Dynamic Core. *Monthly Weather Review*, **136**(11), 4045–4062. https://doi.org/10.1175/2008MWR2523.1

Wicker, L. J. (2009) A two-step Adams–Bashforth–Moulton split-explicit integrator for compressible atmospheric models. *Monthly Weather Review*, **137**(10), 3588–3595. https://doi.org/10.1175/2009MWR2838.1

Wicker, L. J., & Skamarock, W. C. (1998) A time-splitting scheme for the elastic equations incorporating second-order Runge–Kutta time differencing. *Monthly Weather Review*, **126**(7), 1992–1999. https://doi.org/10.1175/1520-0493(1998)126<1992:ATSSFT>2.0.CO;2

Wicker, L. J., & Skamarock, W. C. (2002) Time-splitting methods for elastic models using forward time schemes. *Monthly Weather Review*, **130**(8), 2088–2097. https://doi.org/10.1175/1520-0493(2002)130<2088:TSMFEM>2.0.CO;2

Williams, P. D. (2009) A proposed modification to the Robert–Asselin time filter. *Monthly Weather Review*, **137**(8), 2538–2546. https://doi.org/10.1175/2009MWR2724.1

Williams, P. D. (2011) The RAW filter: An improvement to the Robert–Asselin filter in semi-implicit integrations. *Monthly Weather Review*, **139**(6), 1996–2007. https://doi.org/10.1175/2010MWR3601.1

Williams, P. D. (2013) Achieving seventh-order amplitude accuracy in leapfrog integrations. *Monthly Weather Review*, **141**(9), 3037–3051. https://doi.org/10.1175/MWR-D-12-00303.1

Xue, M. (2000) High-order monotonic numerical diffusion and smoothing. *Monthly Weather Review*, **128**(8), 2853–2864. https://doi.org/10.1175/1520-0493(2000)128<2853:HOMNDA>2.0.CO;2

Young, C.-C., Liang, Y.-C., Tseng, Y.-H., & Chow, C.-H. (2014) Characteristics of the RAW-filtered leapfrog time-stepping scheme in the ocean general circulation model. *Monthly Weather Review*, **142**(1), 434–447. https://doi.org/10.1175/MWR-D-12-00333.1

Young, C.-C., Tseng, Y.-H., Shen, M.-L., Liang, Y.-C., Chen, M.-H., & Chien, C.-H. (2012) Software development of the TaIwan Multi-scale Community Ocean Model (TIMCOM). *Environmental Modelling & Software*, **38**, 214–219. https://doi.org/10.1016/j.envsoft.2012.05.017

Scheme	Amplitude Error	Phase Error	Filter Order	Function Evaluations	Filter Parameters (ν, α, γ)
RA (Robert–Asselin)	1	2	2	1	0.1, —, —
RK3 (Runge–Kutta)	3 (2)	3 (2)	—	3	—, —, —
W03 (Williams)	3	2	2	1	0.1, 0.5, 0
WG3 (Williams)	3	2	2	1	0.1, 0.5, 1
W33 (Williams)	3	2	2	2	0.1, 0.5, 0.5
W43 (Williams)	3	2	4	1	0.1, 0, 1
W05 (Williams)	5	2	4	1	0.1, 0.5, 0
WG5 (Williams)	5	2	4	1	0.1, 0.5, 1
W55 (Williams)	5	2	2	2	0.1, 0.5, (3-ν)/(4-ν)
W77 (Williams)	7	2	4	2	0.1, 0.5, (5-9ν)/[2(4-7ν)]

Table 1. Properties of the numerical time-stepping schemes used in this study, listing the order of accuracy of the amplitude and phase errors, the temporal filter order (where applicable), the number of function evaluations per time step, and the filter parameter values (where applicable). The amplitude and phase error orders are for linear oscillations, unless they appear in parentheses, in which case they are for nonlinear oscillations. Apart from RK3, these are all leapfrog-based schemes.

System	Domain and Resolution	Numerics	Initial Condition	System of Equations
Two-dimensional cone in constant solid-body rotation (Schlesinger 1985)	•L_x=L_y=112 m •T=12-36 rotations •Δx=Δy=2 m •Δt=10π/628 s •Sensitivity for Δx=Δy=1 m and Δt=(5, 2.5, 1.25, 0.625, 0.3125)π/628 s	•10th-order constant grid flux (tested 2nd-16th order; Tremback et al. 1987) •No filters	•Cone •Height=1 m •Radius=8 m •x_c=84 m •y_c=56 m •angular velocity ω=0.1 rad s^{-1}	•Scalar transport •Flux form, non-staggered A-grid
Two-dimensional density current (Straka et al. 1993)	•L_x=51.2 km •L_z=6.4 km •T=900 s •Δx=Δz=100 m •Δt=0.0625 s •Sensitivity for Δx=Δz=12.5-200 m and Δt=0.03125-2.0 s	•10th-order constant grid flux (tested 2nd-16th order; Tremback et al. 1987) •2nd-order fast modes (tested 2nd-10th order) •No filters •Divergence damping	•Ellipse •Cosine squared •x_c=25.6 km •x_r=4 km •z_c=3 km •z_r=2 km •ΔT_{min}=-15 °C	•Quasi-compressible (c_s=347 m s^{-1}) •Straka et al. (1993) •Split sound/advection and diffusion (forward–backward) •Flux form, staggered C-grid •Periodic in x •Free-slip top and bottom boundaries •Constant diffusion K_m=75 m^2 s^{-1} •Sensitivity to K_m=30 m^2 s^{-1}
Three-dimensional warm plume (Carpenter et al. 1990; Tripoli 1992; Wicker and Skamarock 1998; Bryan and Fritsch 2002)	•L_x=L_y=12 km •L_z=10 km •T=1000 s •Δx=Δy=Δz=100 m •Δt=0.4 s •Sensitivity for Δx=Δy=Δz=25-50 m Δt=0.03125-2.0 s	•10th-order constant grid flux (tested 2nd-16th order; Tremback et al. 1987) •2nd-order fast modes (tested 2nd-10th order) •No filters •Divergence damping	•Sphere •Cosine squared •x_c=y_c=6 km •x_r=y_r =2 km •z_c=2 km •z_r=2 km •$\Delta\theta_{max}$=+2 °C	•Fully-compressible •Similar to Bryan and Fritsch (2002) •Split sound/advection and diffusion (forward–backward) •Flux form, staggered C-grid •Periodic in x and y •Free-slip top and bottom boundaries •Constant diffusion K_m=7.5 m^2 s^{-1} •Sensitivity to K_m=2.0 m^2 s^{-1}
Two-dimensional twin tropical cyclones (Moustaoui et al. 2002; Moustaoui et al. 2014)	•L_x=L_y=9,600 km •T=12 days •Δx=Δy=40 km •Δt=15 days •Sensitivity for Δx=Δy=Δz=20 km Δt=7.5 days	•10th-order constant grid flux (tested 2nd-16th order; Tremback et al. 1987) •2nd-order fast modes (tested 2nd-10th order) •6th-order filter with $K_6\Delta t/\Delta x^6$=2.5×10^{-6} •Divergence damping	•Vorticity centres with opposite signs at 7°N and 7°S •Exponential decay from extremes •Maximum velocity V_0=30 m s^{-1}	•Rotating shallow-water equations •Anthes and Hoke (1975) •H_0 = 234.86 m •Explicit or split gravity wave/advection and diffusion (forward–backward) •Flux form, staggered C-grid •Divergence permitting •Radiative lateral boundaries •Normal velocity C_{gx}=58 m s^{-1}, C_{gy}=5 m s^{-1}

Table 2. Summary of the experimental designs for the benchmark simulations studied in this paper. L_x, L_y, and L_z are the domain dimensions; T is the simulation duration; Δx, Δy, and Δz are the numerical grid spacings; Δt is the time step; K_m is the eddy mixing coefficient; H_0 is the initial mean elevation of the fluid for the twin cyclone problem; c_s is the speed of sound; C_{gx} and C_{gy} are the normal boundary condition

gravity wave speeds; ΔT_{min} and $\Delta\theta_{max}$ are the temperature and potential temperature perturbations, respectively, of the specified initial conditions; and (x_c, y_c, z_c) and (x_r, y_r, z_r) are the centre and radius, respectively, of the specified feature.

Scheme	Maximum and Minimum Values (m)	L^∞ Error (m)	RMS Error (m)	Dissipation and Dispersion Errors (m^2)
RA ($\Delta t = 10\pi/628$ s)	0.476 -1.00×10^{-2}	0.524	2.49×10^{-2}	2.27×10^{-4} 3.93×10^{-4}
RA ($\Delta t = 2.5\pi/628$ s)	0.702 -1.28×10^{-2}	0.298	1.19×10^{-2}	4.35×10^{-5} 9.79×10^{-5}
RA ($\Delta t = 0.625\pi/628$ s)	0.823 -2.74×10^{-2}	0.176	7.04×10^{-3}	4.05×10^{-6} 4.55×10^{-5}
RK3 ($\Delta t = 10\pi/628$ s)	0.874 -4.62×10^{-2}	0.126	7.03×10^{-3}	6.76×10^{-9} 4.94×10^{-5}
W03 ($\Delta t = 10\pi/628$ s)	0.870 -6.17×10^{-2}	0.130	9.32×10^{-3}	1.61×10^{-9} 8.70×10^{-5}
W77 ($\Delta t = 10\pi/628$ s)	0.872 -6.63×10^{-2}	0.128	9.91×10^{-3}	9.51×10^{-10} 9.83×10^{-5}

Table 3. Values of various error measures for the cone height after 12 rotations, each rounded to three significant figures. All solutions use $\Delta x = \Delta y = 2$ m.

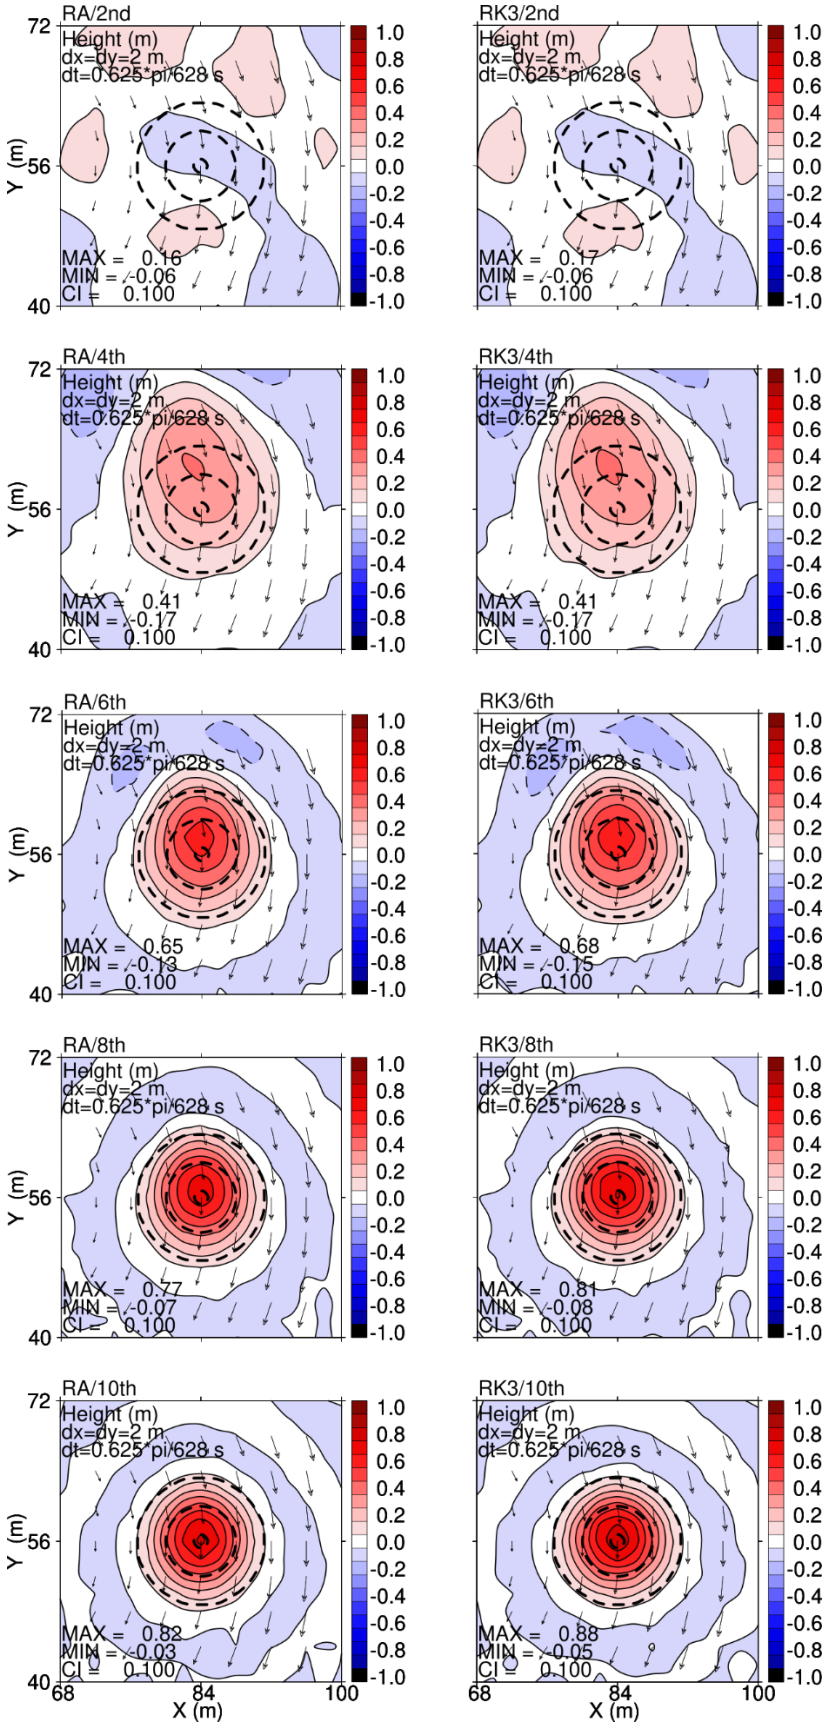

Figure 1. Cone height (m) after 12 rotations for RA with $\nu = 0.1$ (left column) and RK3 (right column) made using the O(2, 4, 6, 8, 10) constant grid flux schemes (top to bottom rows). Values of $\Delta x = \Delta y = 2$ m and $\Delta t = 0.625\pi/628$ s are used to minimize the time truncation error. The maximum value (MAX), minimum value (MIN), and contour interval (CI = 0.1 m) are stated on each plot. The thin, solid contours indicate positive and zero values, and the thin, dashed contours indicate negative values. The thick, dashed contours are the superimposed analytical solution at 0.1, 0.5, and 0.9 m. Wind vectors are plotted at every second gridpoint. Only the sub-domain from $x = 68$–100 m and $y = 40$–72 m is shown.

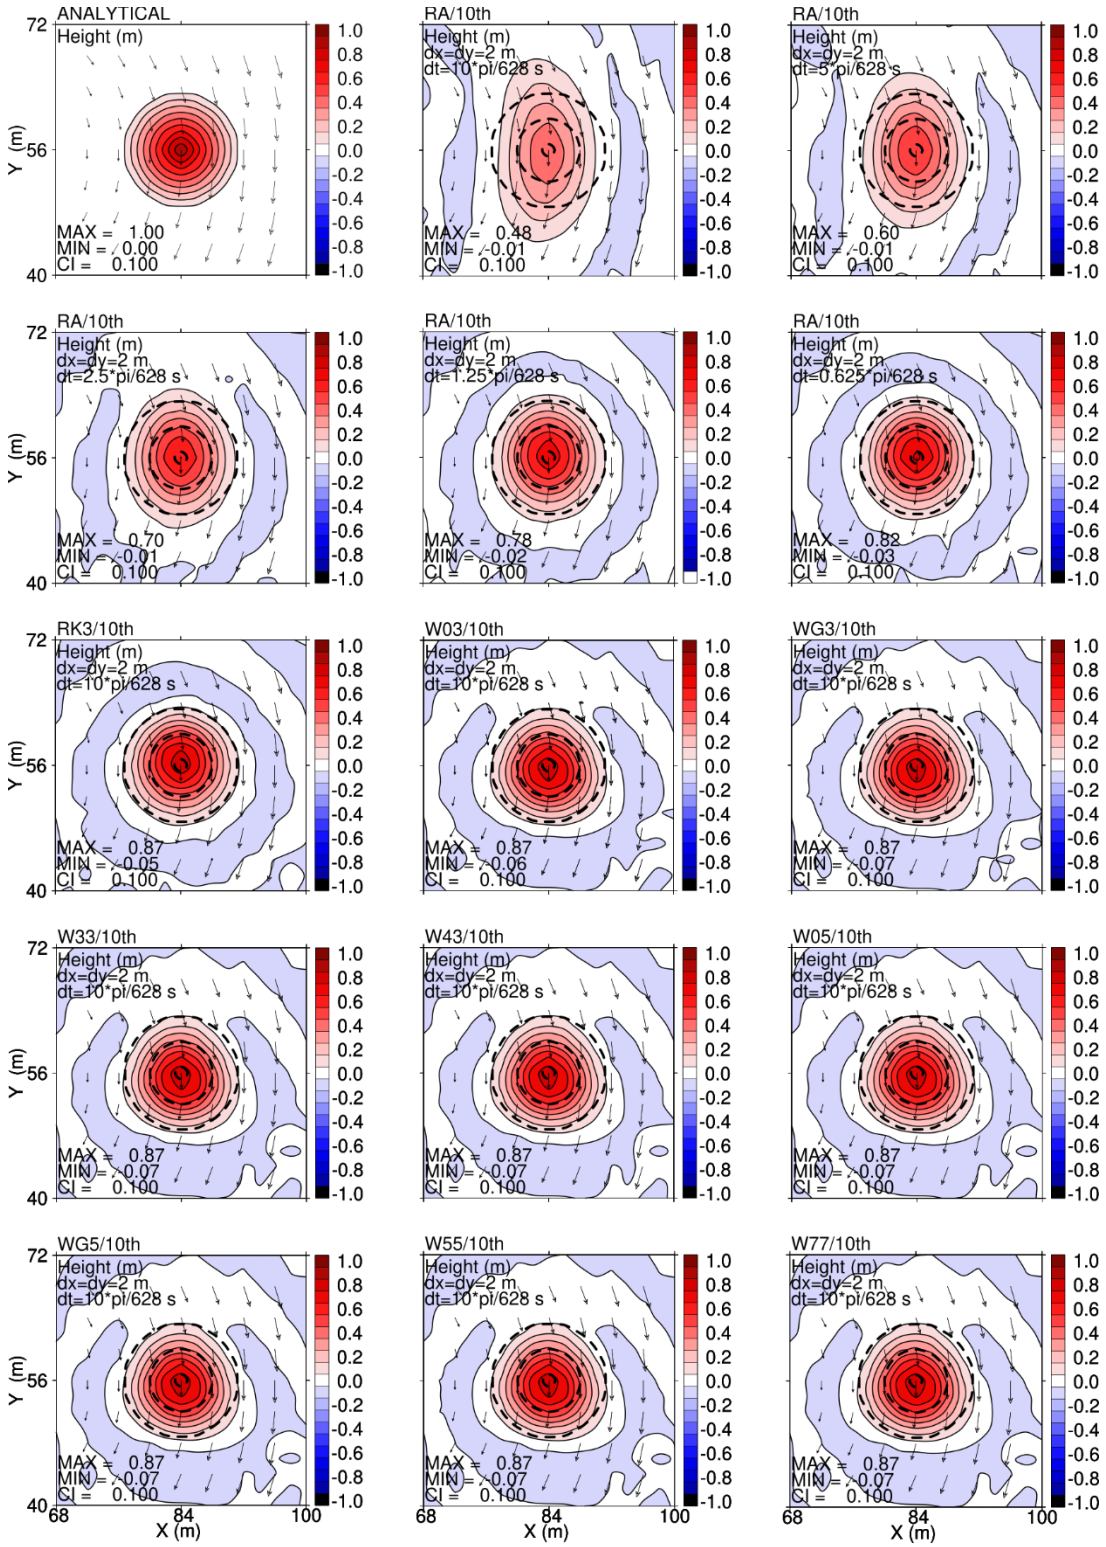

Figure 2. Cone height (m) after 12 rotations for all the schemes tested (RA, RK3, W03, WG3, W33, W43, W05, WG5, W55, W77). The RA solutions are made using Δt = 10π/628, 5π/628, 2.5π/628, 1.25π/628, and

$0.625\pi/628$ s. All other solutions are made using $\Delta t = 10\pi/628$ s. All solutions use $\Delta x = \Delta y = 2$ m. The maximum value (MAX), minimum value (MIN), and contour interval (CI = 0.1 m) are stated on each plot. The thin, solid contours indicate positive and zero values; there are no negative contours with the chosen contour interval. The thick, dashed contours are the superimposed analytical solution at 0.1, 0.5, and 0.9 m. Wind vectors are plotted at every second gridpoint. Only the sub-domain from x = 68–100 m and y = 40–72 m is shown.

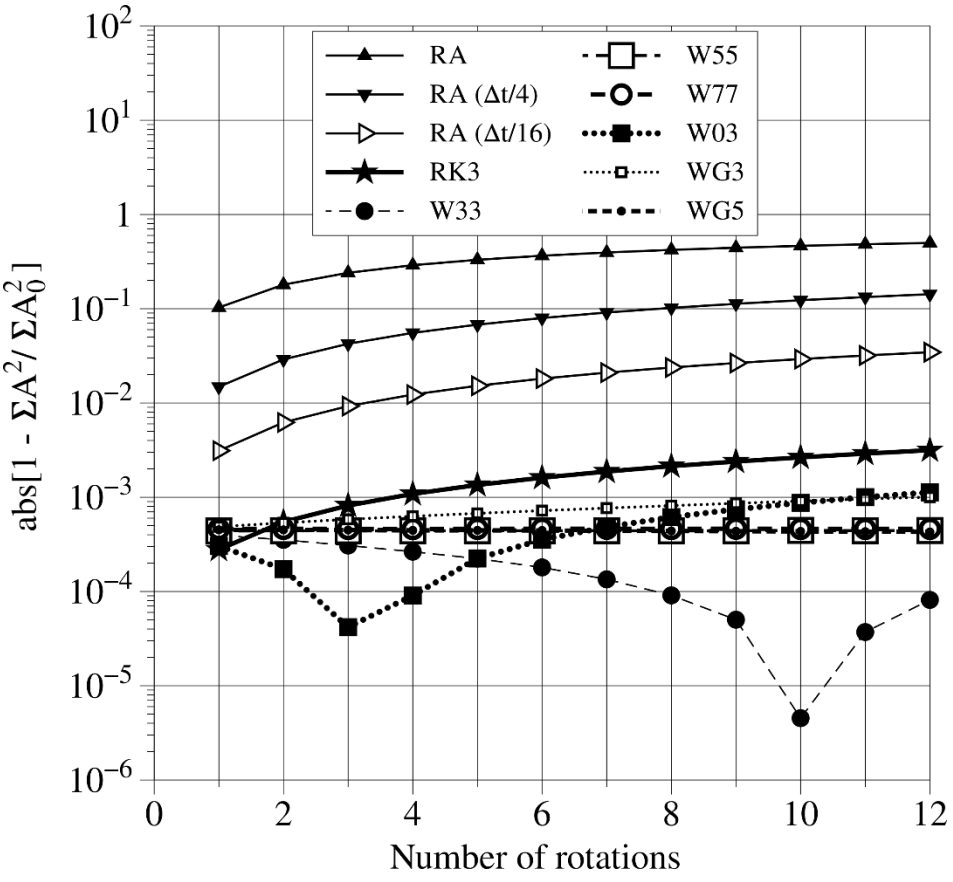

Figure 3. Time series of an error measure for the cone height field in the RA, RK3, W33, W55, W77, W03, WG3, and WG5 schemes. The error measures are evaluated once per complete rotation, from the start of the simulations up to 12 rotations. The error measure plotted is the absolute value of $1-\Sigma A^2/\Sigma A_0^2$, where A is the scalar cone height field, A_0 is the exact solution, and both summations are taken over the entire simulation domain. The RA solutions are made using $\Delta t = 10\pi/628$, $2.5\pi/628$ (labelled $\Delta t/4$), and $0.625\pi/628$ s (labelled $\Delta t/16$). All other solutions are made using $\Delta t = 10\pi/628$ s. All solutions use $\Delta x = \Delta y = 2$ m.

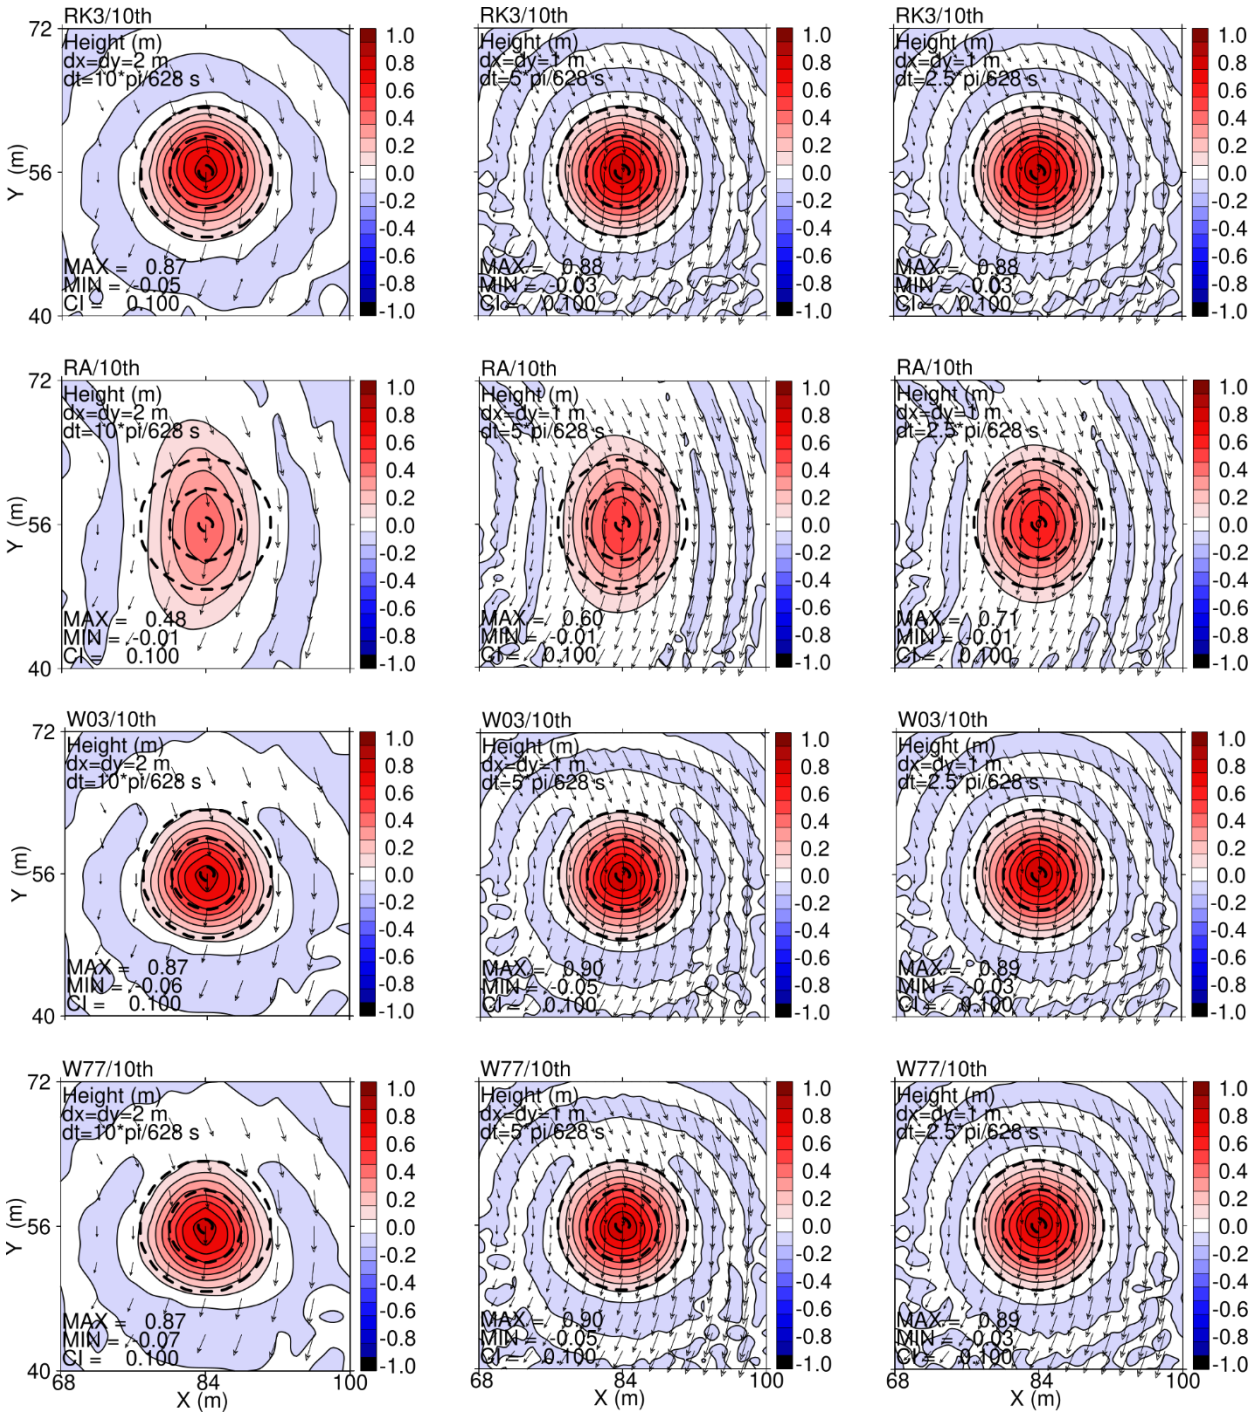

Figure 4. Cone height after 12 rotations for the RK3, RA, W03, and W77 schemes (top to bottom rows) with different time steps and spatial resolutions (left to right columns). The solutions were made using $\Delta t = 10\pi/628$ s and $\Delta x = \Delta y = 2$ m (left column), $\Delta t = 5\pi/628$ s and $\Delta x = \Delta y = 1$ m (middle column), and $\Delta t = 2.5\pi/628$ s and $\Delta x = \Delta y = 1$ m (right column). The maximum value (MAX), minimum value (MIN),

and contour interval (CI = 0.1 m) are stated on each plot. The thin, solid contours indicate positive and zero values; there are no negative contours with the chosen contour interval. The thick, dashed contours are the superimposed analytical solution at 0.1, 0.5, and 0.9 m. Wind vectors are plotted at every second gridpoint. Only the sub-domain from x = 68–100 m and y = 40–72 m is shown.

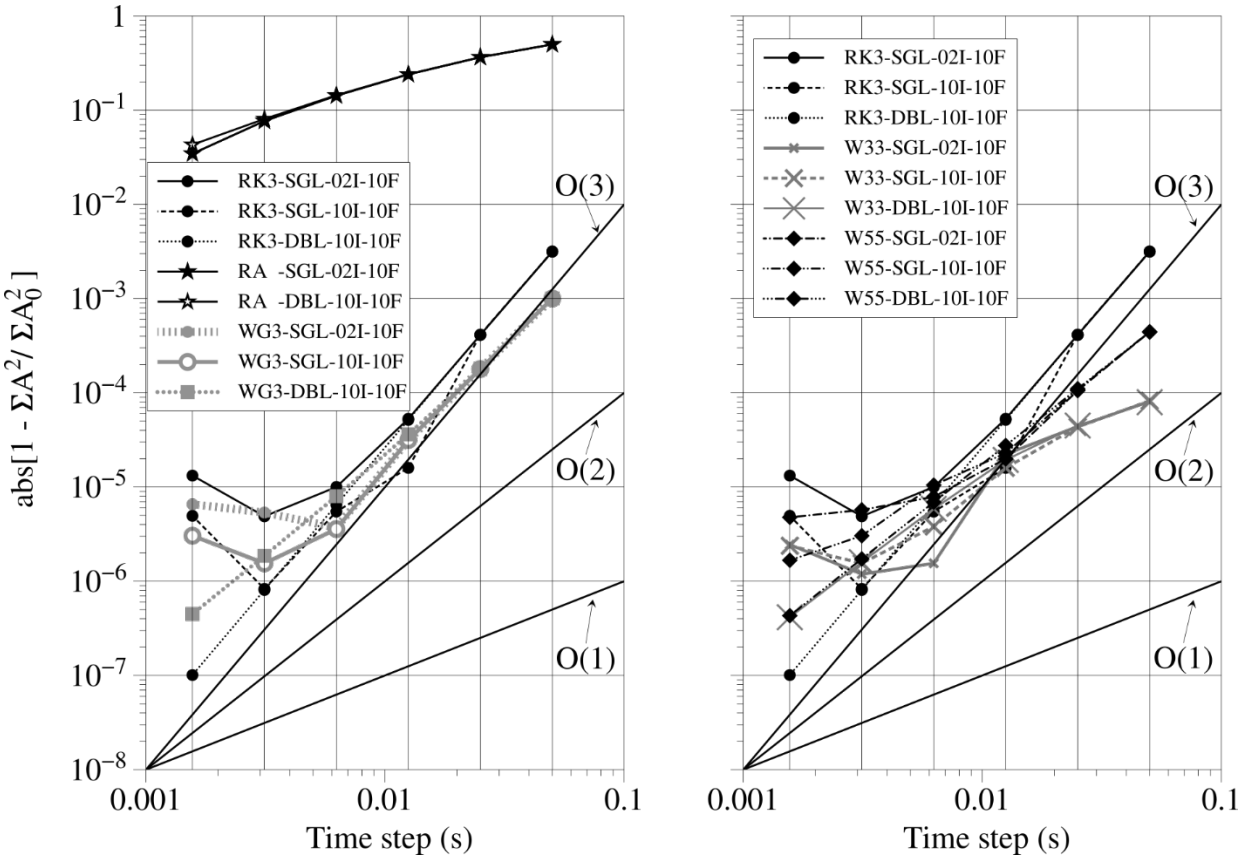

Figure 5. Variation of cone height error with time step after 12 rotations. The error measure plotted is the absolute value of $1-\Sigma A^2/\Sigma A_0^2$, where A is the scalar cone height field, A_0 is the exact solution, and both summations are taken over the entire simulation domain. The six time-step sizes used (Δt = 10π/628, 5π/628, 2.5π/628, 1.25π/628, 0.625π/628, and 0.3125π/628 s) are indicated with vertical lines. All solutions used $\Delta x = \Delta y$ = 2 m. Identical RK3 curves are shown on both panels for reference, with two additional schemes requiring a single function evaluation per time step shown on the left panel (RA and WG3) and two additional schemes requiring two function evaluations per time step shown on the right panel (W33 and W55). The simulations were run using either single-precision (32-bit) calculations (SGL) or double-precision (64-bit) calculations (DBL). The simulations used either 2nd-order (02I) or 10th-order (10I) interpolation of the wind. All simulations used the 10th-order constant grid flux scheme (10F). Both axes are logarithmic, and the reference lines labelled O(1), O(2), and O(3) indicate linear, quadratic, and cubic scaling.

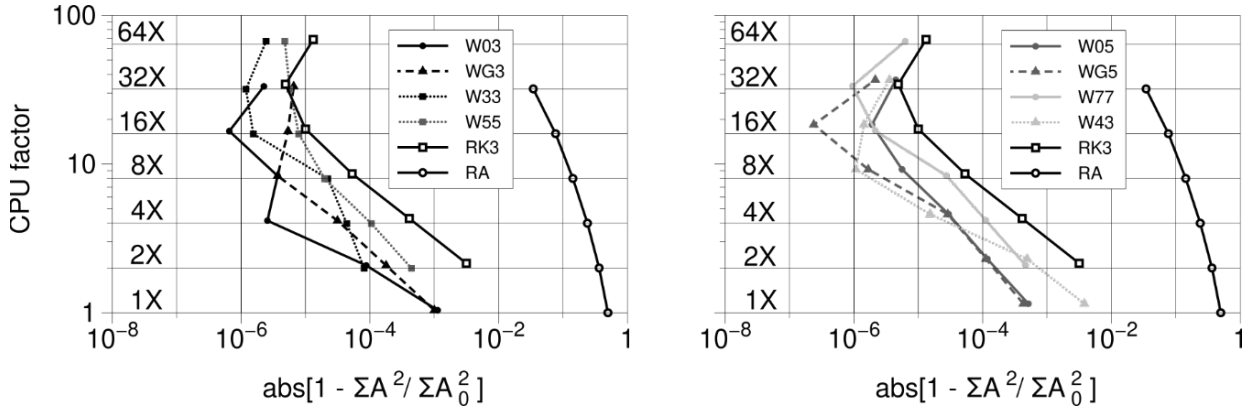

Figure 6. Variation of computational expense with cone height error for six different time-step sizes after 12 rotations. The cone height error measure, six time-step sizes, and horizontal resolution are the same as in Figure 5. In all cases plotted here, the numerical precision is single (32 bit), the wind interpolation is 2nd order, and the constant grid flux is 10th order. Identical RK3 and RA curves are shown on both panels for reference, with four additional schemes using second-order temporal filters shown on the left panel (W03, WG3, W33, and W55) and four additional schemes using fourth-order temporal filters shown on the right panel (W05, WG5, W77, and W43). The computational expense is measured by a Central Processing Unit (CPU) factor, which is defined as the total number of floating-point operations required to simulate 12 cone rotations (taking into account the relative costs of addition, subtraction, multiplication, and division) normalized by that of the RA scheme at the longest time step.

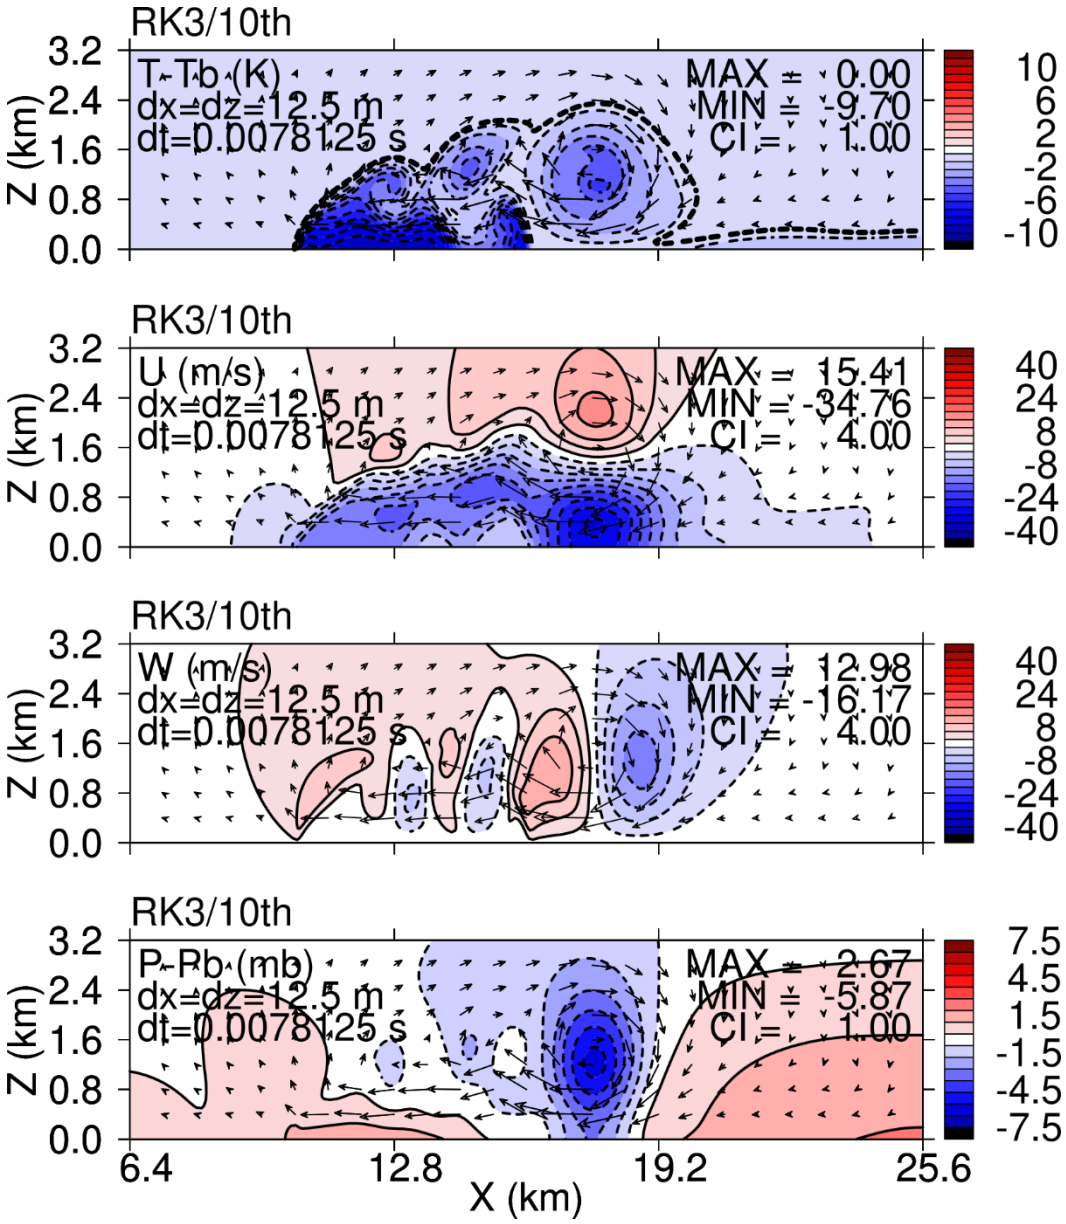

Figure 7. Reference solution for the two-dimensional density current after 900 s, obtained using the RK3 scheme with $\Delta t = 0.0078125$ s and $\Delta x = \Delta z = 12.5$ m. Plotted from top to bottom are the perturbation potential temperature (T-Tb in K), horizontal velocity (U in m s^{-1}), vertical velocity (W in m s^{-1}), and perturbation pressure (P-Pb in mb). The maximum value (MAX), minimum value (MIN), and contour interval (CI) are stated on each plot. The thin, solid contours indicate positive values, and the thin, dashed contours indicate negative values. There is no zero contour in the perturbation potential temperature plot, because there are no positive values and the maximum value is slightly less than zero. The zero contours

are omitted from the other three plots to keep them uncluttered. The thick, dashed contour is the perturbation potential temperature at -0.75 K to indicate the bulk location of the density current. Wind vectors are plotted every 800 m horizontally and 400 m vertically. Only part of the lower-left quadrant of the domain from x = 6.4–25.6 km and z = 0–3.2 km is shown.

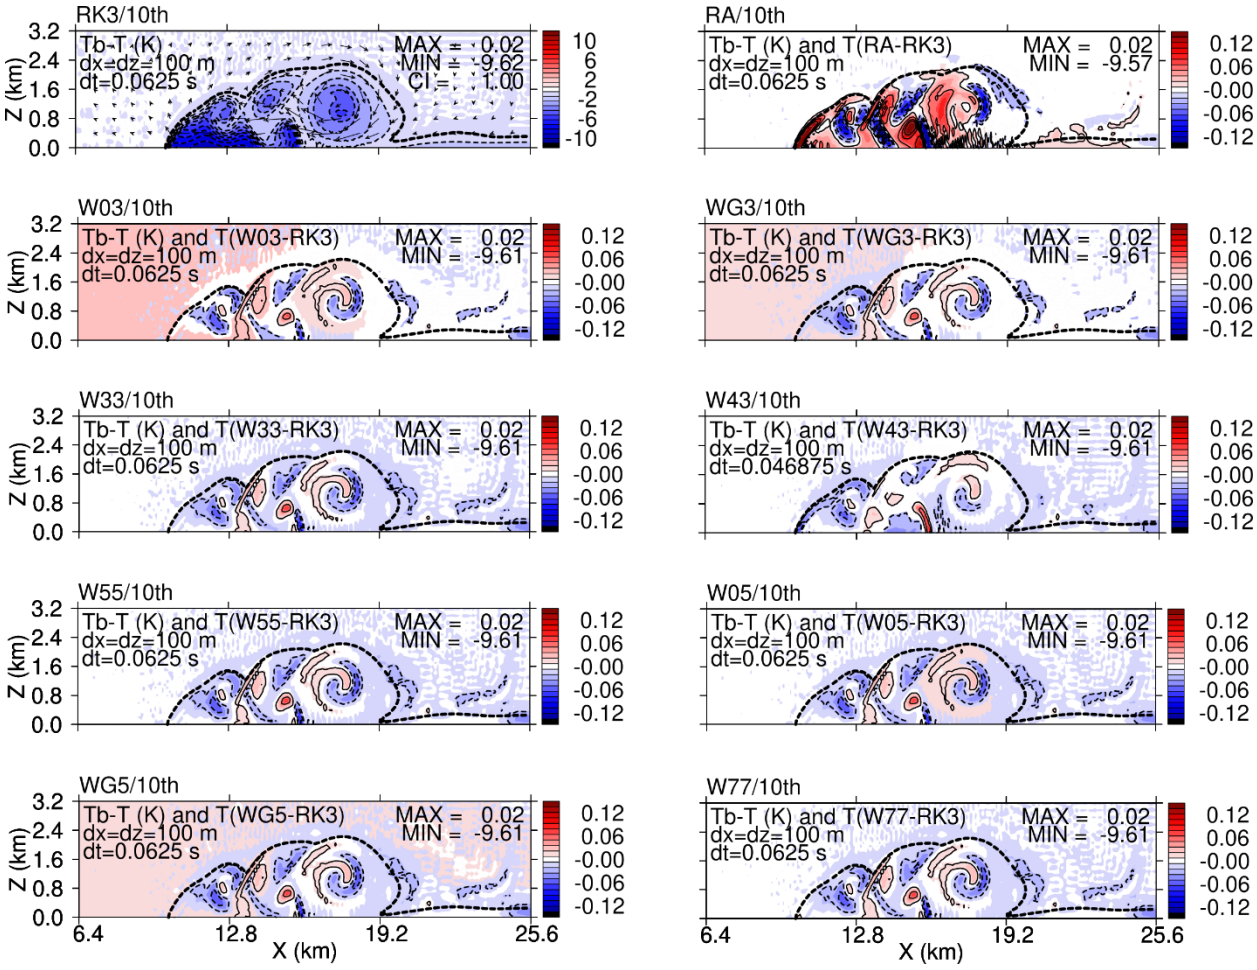

Figure 8. Perturbation potential temperature (K) for the two-dimensional density current after 900 s, obtained using the RK3, RA, and eight WFL schemes. All simulations use Δt = 0.0625 s (except that W43 uses three-quarters of this value) and $\Delta x = \Delta z$ = 100 m. The first plot shows RK3 and the subsequent plots show differences with respect to RK3. The maximum value (MAX) and minimum value (MIN) of the test simulations are stated on each plot. The thin, solid contours indicate positive values and the thin, dashed contours indicate negative values. The thick, dashed contours are the superimposed RK3 solution at -0.75 K to indicate the bulk location of the density current. Wind vectors in the RK3 simulation are plotted every 800 m horizontally and 400 m vertically. Only part of the lower-left quadrant of the domain from x = 6.4–25.6 km and z = 0–3.2 km is shown.

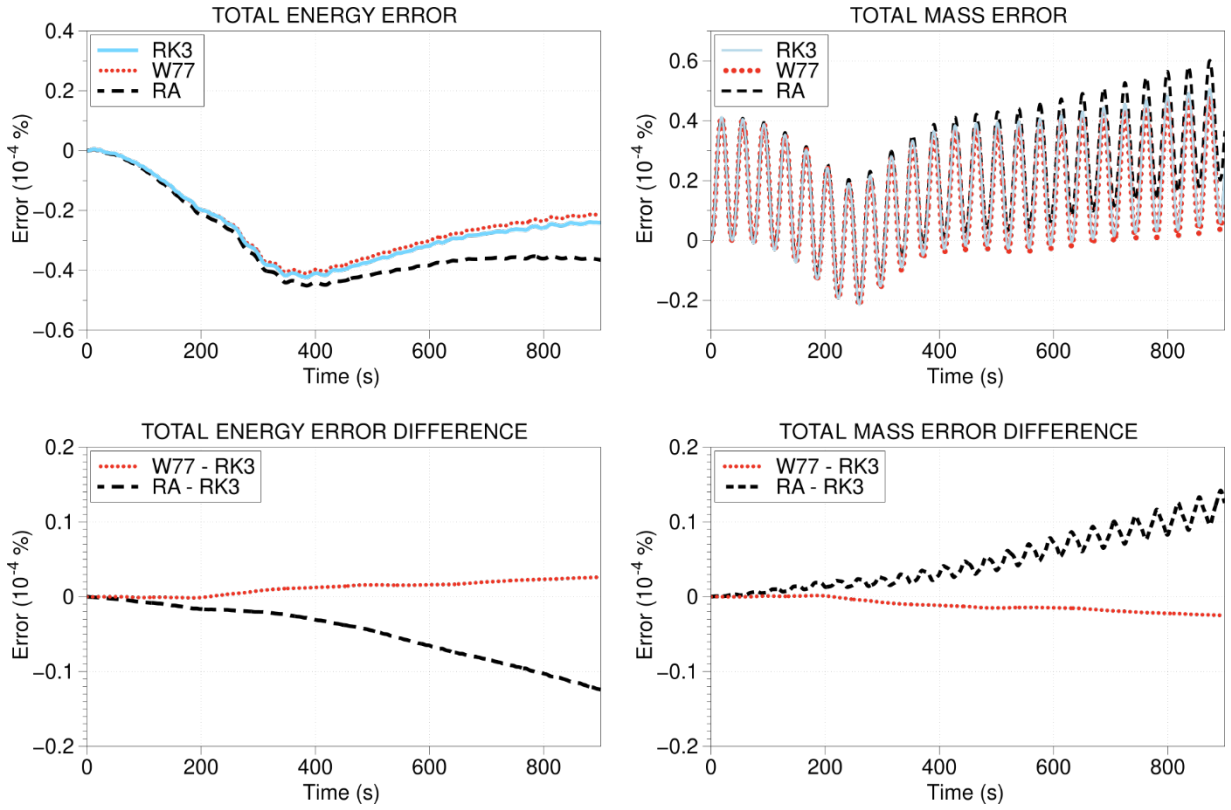

Figure 9. Time series of total energy (left column) and fluid mass (right column) for the first 900 s of the two-dimensional density current simulations, using the RK3, RA, and W77 schemes with $\Delta t = 0.0625$ s and $\Delta x = \Delta z = 100$ m. The upper row shows errors, which are defined as the percentage change in each simulation since $t = 0$. The lower row shows the corresponding error differences in RA and W77 with respect to RK3.

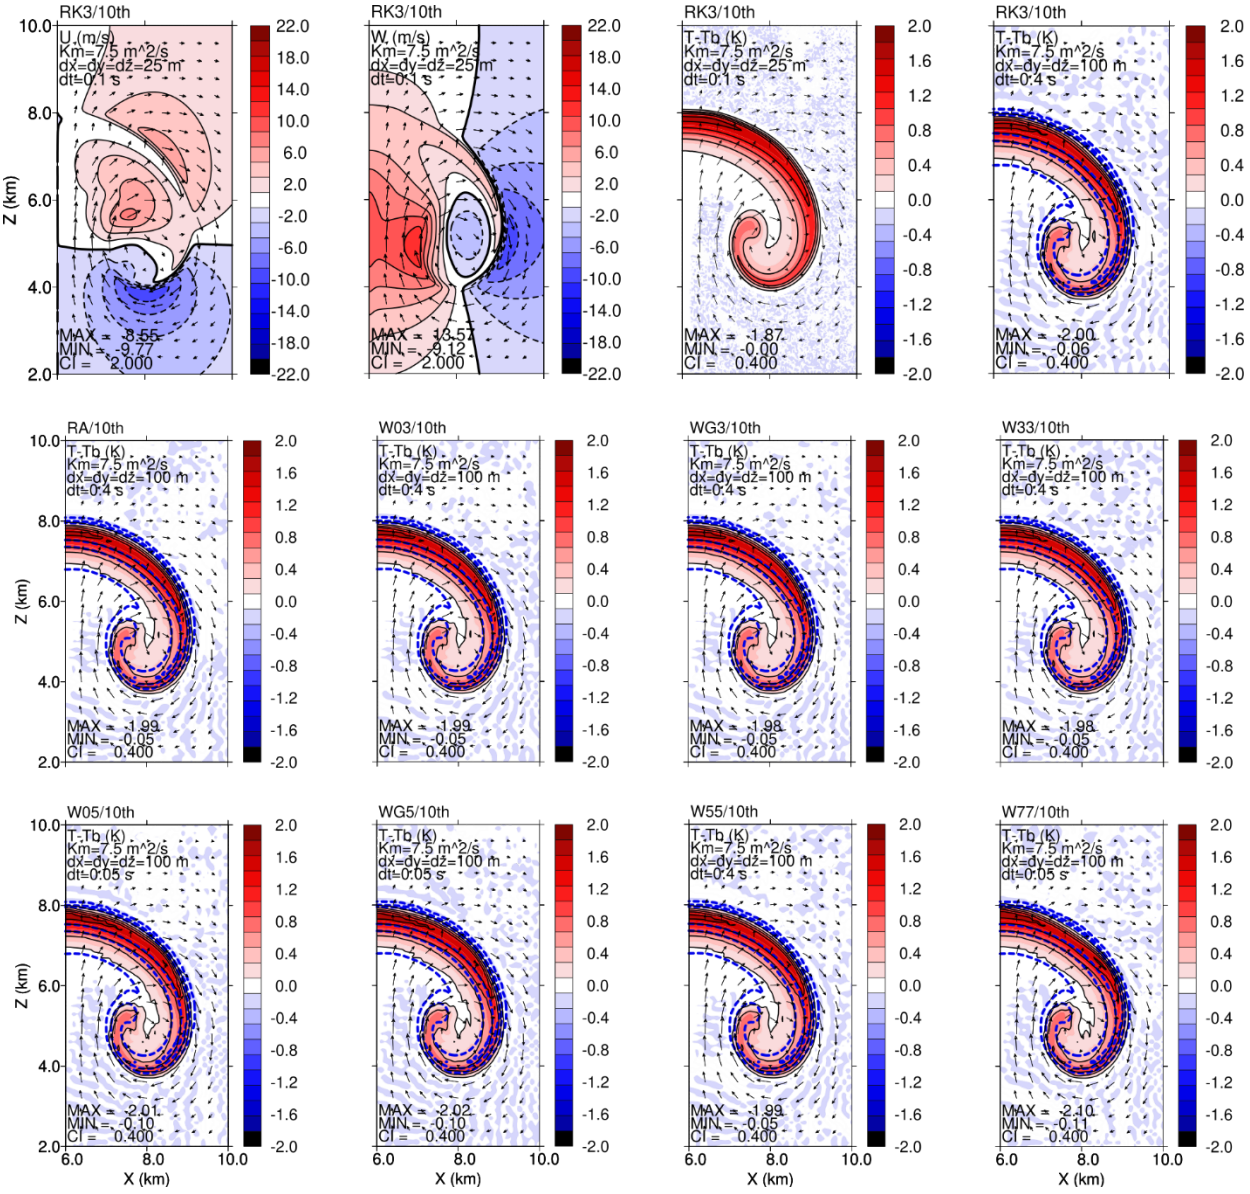

Figure 10. The three-dimensional rising warm plume after 1,000 s, simulated using the RK3, RA, W03, WG3, W33, W05, WG5, W55, and W77 schemes. The first three panels show the radial velocity (U in m s^{-1}), vertical velocity (W in m s^{-1}), and perturbation potential temperature (T-Tb in K) in the high-resolution RK3 reference solution, made using $\Delta x = \Delta y = \Delta z = 25$ m and $\Delta t = 0.1$ s. The following panels show the perturbation potential temperature (T-Tb in K) in the test solutions, made using $\Delta x = \Delta y = \Delta z =$ 100 m and $\Delta t = 0.4$ s (except that W05, WG5, and W77 use $\Delta t = 0.05$ s). The simulations use an eddy diffusion coefficient of $K_m = 7.5$ m^2 s^{-1}. The maximum value (MAX), minimum value (MIN), and contour interval (CI) are stated on each plot. The thin, solid contours indicate positive values, and the thin, dashed

contours indicate negative values. The thick, solid contours in the first two panels indicate the zero values, which are omitted from the subsequent panels to keep them uncluttered. The thick, dashed contours are the superimposed perturbation potential temperature in the high-resolution RK3 reference solution at 0.1, 0.6, and 1.1 K to indicate the bulk location of the plume. Wind vectors are plotted every 400 m horizontally and vertically. Only the sub-domain from x = 6–10 km and z = 2–10 km at y = 6 km is shown.

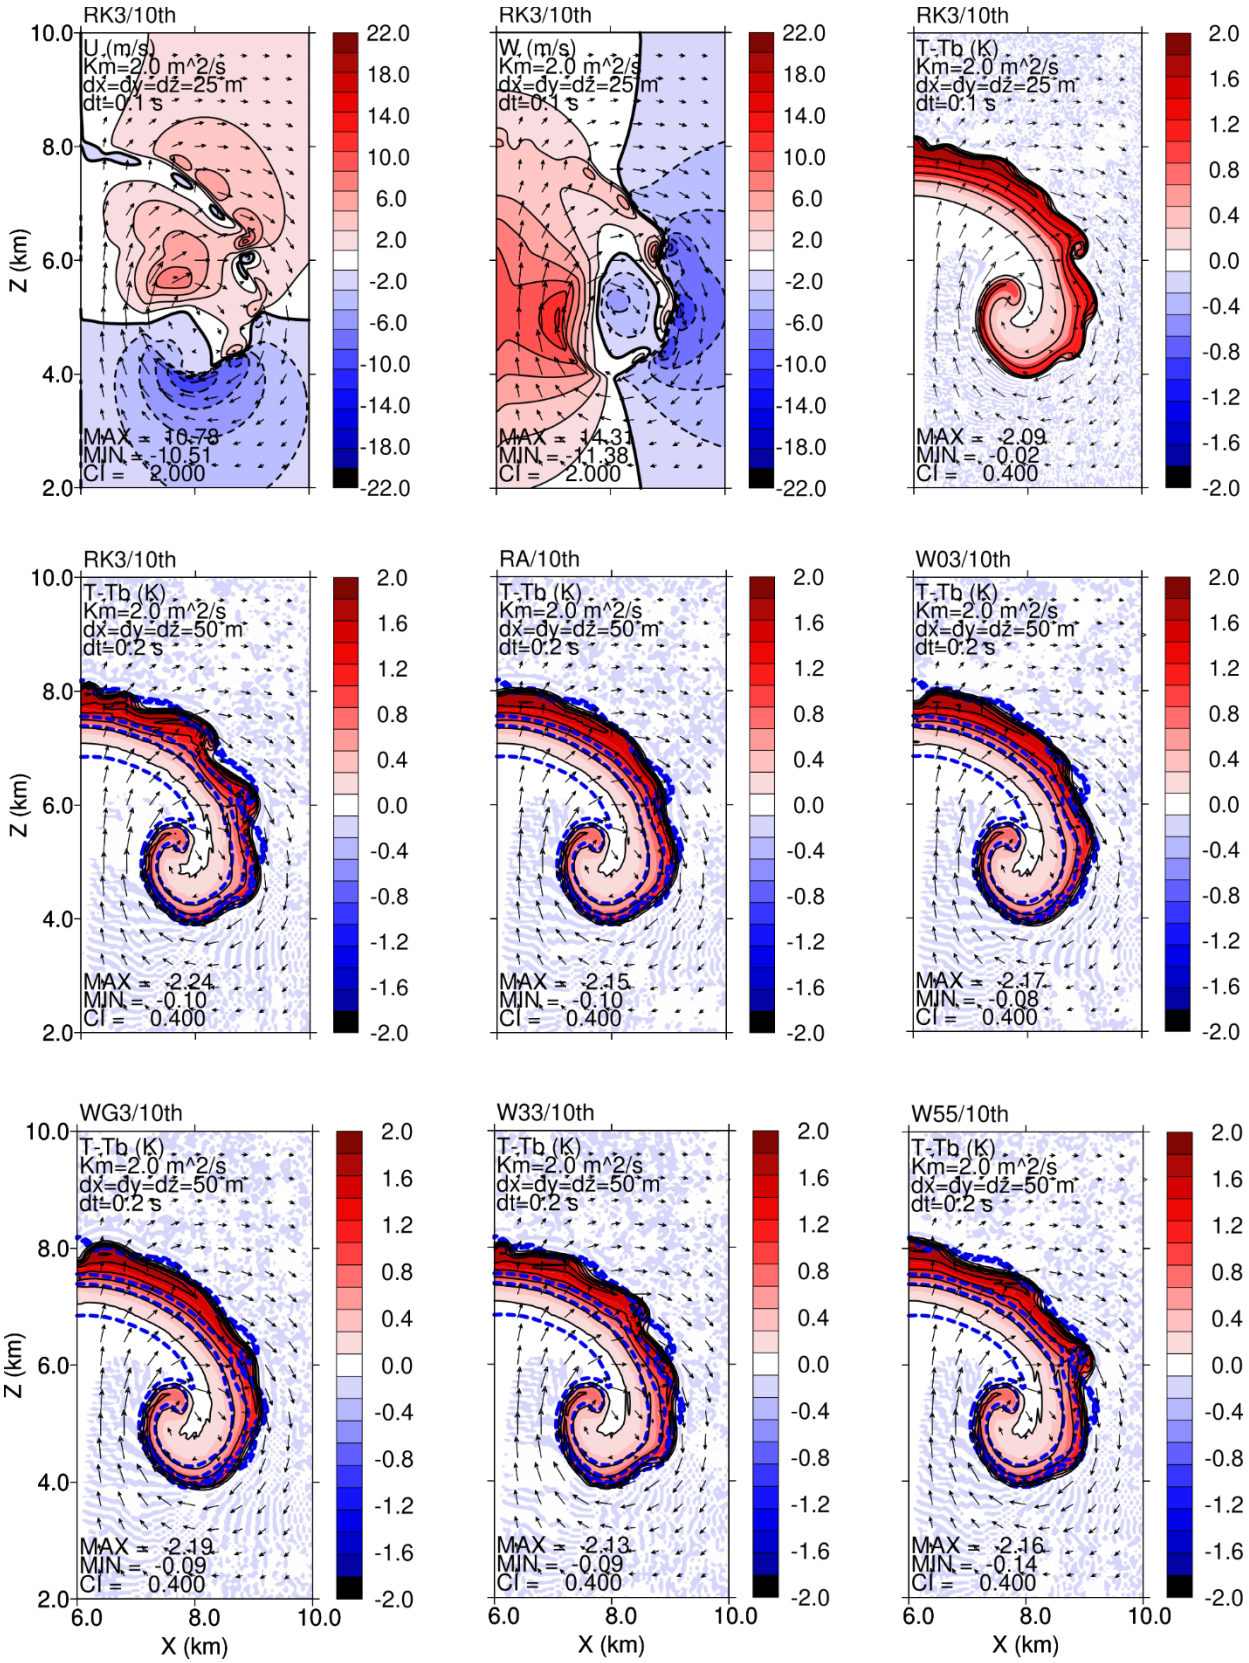

Figure 11. The three-dimensional rising warm plume after 1,000 s, simulated using the RK3, RA, W03, WG3, W33, and W55 schemes. The top row shows the radial velocity (U in m s^{-1}), vertical velocity (W in m s^{-1}) and perturbation potential temperature (T-Tb in K) in the high-resolution RK3 reference solution, made using $\Delta x = \Delta y = \Delta z = 25$ m and $\Delta t = 0.1$ s. The following panels show the perturbation potential temperature (T-Tb in K) in the test solutions, made using $\Delta x = \Delta y = \Delta z = 50$ m and $\Delta t = 0.2$ s. The simulations use an eddy diffusion coefficient of $K_m = 2.0$ m^2 s^{-1}. The maximum value (MAX), minimum value (MIN), and contour interval (CI) are stated on each plot. The thin, solid contours indicate positive values, and the thin, dashed contours indicate negative values. The thick, solid contours in the first two panels indicate the zero values, which are omitted from the subsequent panels to keep them uncluttered. The thick, dashed contours are the superimposed perturbation potential temperature in the high-resolution RK3 reference solution at 0.1, 0.6, and 1.1 K to indicate the bulk location of the plume. Wind vectors are plotted every 400 m horizontally and vertically. Only the sub-domain from $x = 6$–10 km and $z = 2$–10 km at $y = 6$ km is shown.

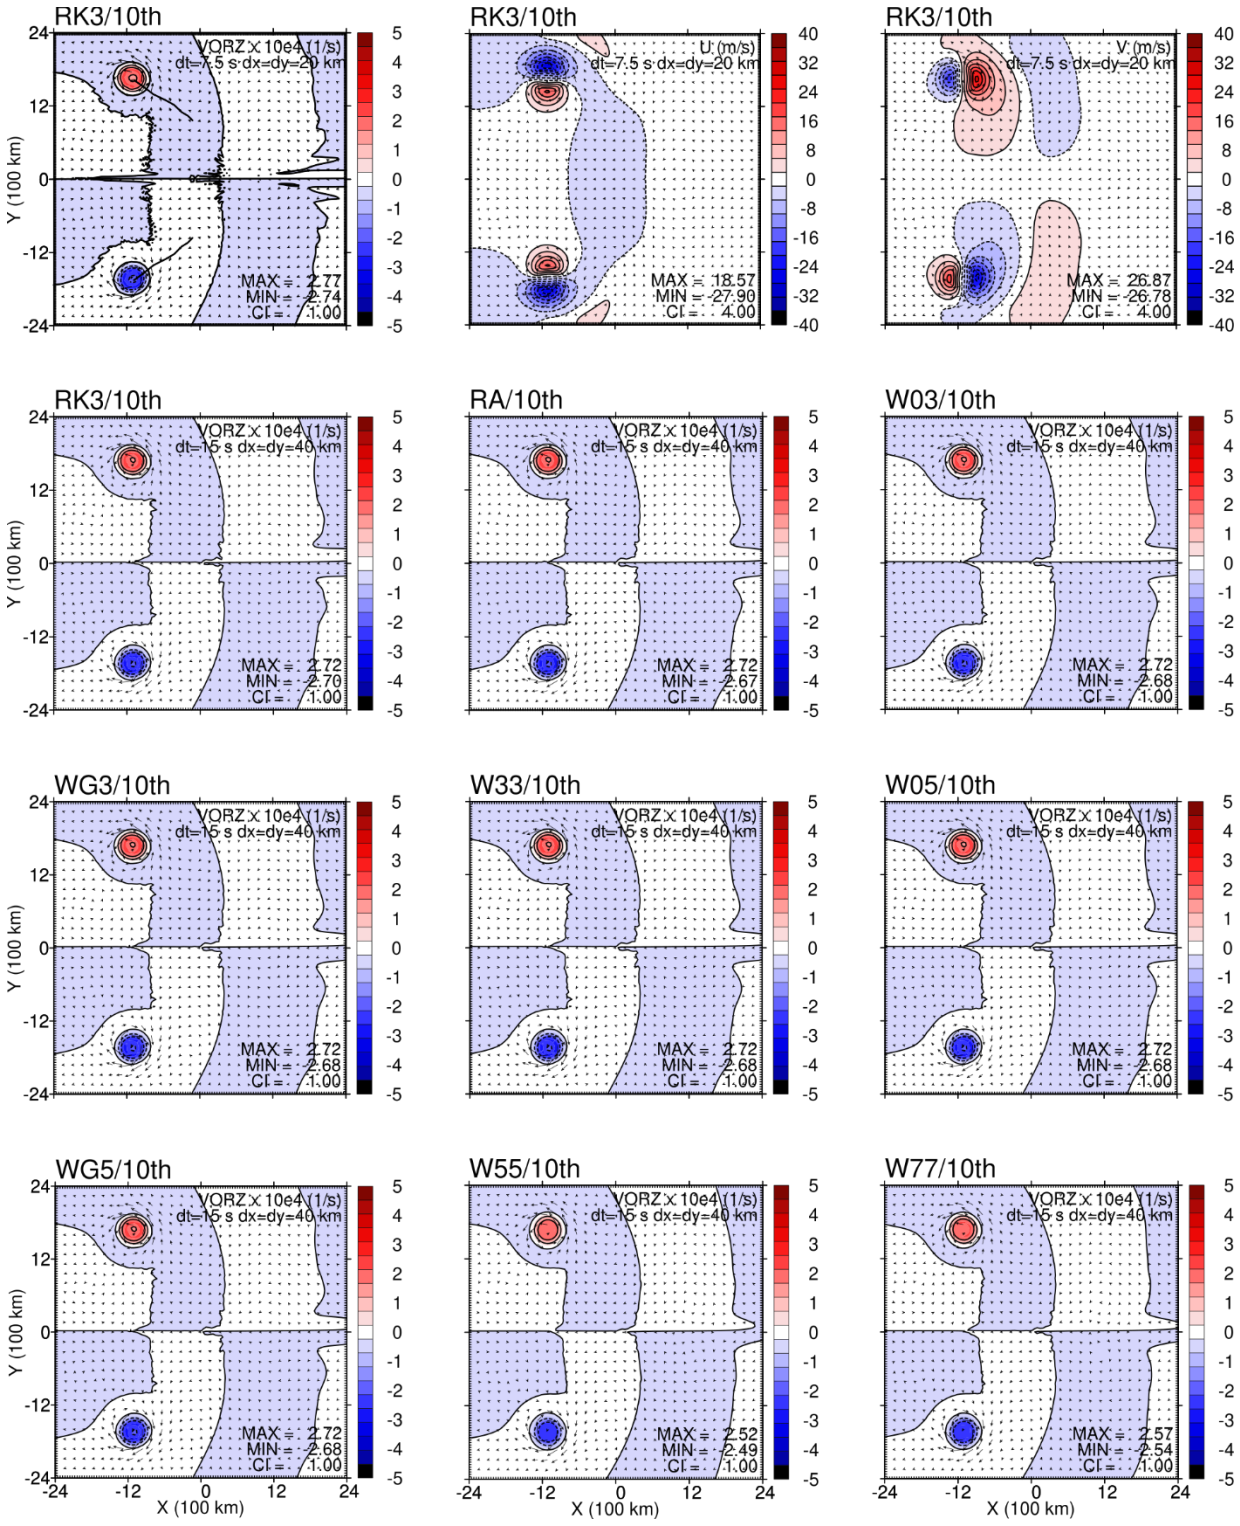

Figure 12. The two-dimensional twin tropical cyclones after four days, simulated using the RK3, RA, W03, WG3, W33, W05, WG5, W55, and W77 schemes. The top row shows the vertical vorticity (VORZ in 10^{-4} s^{-1}), zonal velocity (U in m s^{-1}), and meridional velocity (V in m s^{-1}) in the high-resolution RK3

reference solution, made using $\Delta x = \Delta y = 20$ km and $\Delta t = 7.5$ s. In the vertical vorticity plot for the reference solution, the solid lines indicate the cyclone tracks since $t = 0$. The following panels show the vertical vorticity (VORZ in 10^{-4} s^{-1}) in the test solutions, made using $\Delta x = \Delta y = 40$ km and $\Delta t = 15$ s. The maximum value (MAX), minimum value (MIN), and contour interval (CI) are stated on each plot. The solid contours indicate positive and zero values, and the dashed contours indicate negative values. The zero contours are omitted from the zonal and meridional velocity plots to keep them uncluttered. Wind vectors are plotted every 160 km zonally and meridionally. Only the sub-domain from $x = \pm 2{,}400$ km and $y = \pm 2{,}400$ km is shown.