

POSITIVITY PRESERVATION AND ADVECTION ALGORITHM TESTS FOR EDGE PLASMA TURBULENCE

J. L. PETERSON* AND G. W. HAMMETT†

Abstract. The unique conditions of edge tokamak plasmas motivate the employment of efficient numerical methods that can robustly handle steep temperature and density gradients. To that end, we compare an Arakawa finite difference technique with some recent high-resolution upwind methods that are designed to minimize non-monotonic overshoots while preserving the accuracy of solutions at smooth extrema. Versions of these algorithms are able to rigorously preserve positivity when that is physically relevant (such as for the advection of particle density or temperature). We explain in detail the use of these methods for the nonlinear Poisson bracket, an operator applicable to neutral fluid, gyrofluid, gyrokinetic, and general Hamiltonian simulations. For one-dimensional passive advection, the high-resolution upwind techniques maintain monotonicity and approach minimal levels of phase error and dissipation, especially at long-wavelengths. In a two-dimensional incompressible Navier-Stokes vortex merging problem we find that extrema-preserving methods can resolve details at lower resolution than the Arakawa technique, are less dissipative than traditional finite volume methods while still minimizing overshoots and ensuring positivity, and model nonlinear cascade behavior fairly well without additional subgrid damping.

Key words. numerical simulation, fluid, advection, positivity, high order upwind

AMS subject classifications. 76D05, 65M06, 65M08

1. Introduction. The unique environment of the edge region of fusion plasmas complicates the study of turbulence. Analytic approximations made to simplify tokamak core simulations do not necessarily hold true. Fluctuation levels may grow as large as the background distribution function, thereby limiting the applicability of perturbative formulations. Furthermore, steep temperature and density gradients and the general range of scales in edge turbulence strain computational resources. As the density and temperature can vary by two orders of magnitude or more over very short distances in the edge region, for instance, from the top of the pedestal to the scrape-off layer or in the vicinity of plasma blobs, the resolution requirements to accurately model the system can be challenging. These large gradients can lead to many of the same difficulties, such as unphysical overshoots and negative solutions, that motivated the development of modern shock-capturing algorithms.

In choosing from the virtual zoo of numerical algorithms the best possible method for edge turbulence simulation, one should consider many different properties of the computational scheme. First of all, it is important to consider the various conservation properties of the method. How well does the scheme preserve the density, momentum, energy, enstrophy and/or entropy of the system if physical dissipation is turned off? But one should also consider other important properties beyond integral conservation laws. For example, a second consideration is that one would like the system to satisfy certain physical constraints. Positivity-preserving algorithms can ensure that the temperature, density and distribution function solutions never drop below zero. This is particularly important in the presence of shocks, blobs and other steep-gradient features. Third, the accuracy of the employed method is important. High-order schemes can be advantageous, as they reduce the resolution required to accurately

*Lawrence Livermore National Laboratory, 7000 East Avenue, Livermore, CA 94551, USA (peterson76@llnl.gov)

†Princeton University, Princeton Plasma Physics Laboratory, P.O. Box 451, Princeton, NJ 08543, USA

model the physics of the system. Any reduction in resolution gets amplified according to the dimensions of the problem. A factor of 2 reduction in resolution leads to a factor of 2^d reduction in memory and a factor of $\approx 2^{d+1}$ reduction in computing time, which is a factor of ≈ 64 speed up for 5-dimensional gyrokinetics. (Though the reduction in the number of grid points is partially offset by the increased work per grid point needed for higher order methods, this can still be advantageous because many modern computers are limited by the bandwidth between memory and processor, so it is relatively efficient to do the extra calculations per memory access needed for higher order methods). Finally, when dealing with turbulence simulations, one must consider the ability of the algorithm to deal with cascades in k-space.

The problems with edge plasma turbulence simulation are not, it should be noted, confined to plasmas. The similarities between the equations governing drift-wave turbulence, two-dimensional neutral fluid turbulence, and general Hamiltonian systems, expand this problem to a much wider realm. While each physical problem has its own ideal algorithm, the knowledge on algorithm behavior gained in one field easily translates.

Naulin and Nielsen [21] compared a spectral method, a well-known Arakawa finite difference method [1], and a third order finite-volume WENO method [18, 26] due to Kurganov and Levy [14], for the two-dimensional high Reynolds number Navier-Stokes problem of vortex merging [24]. Naulin and Nielsen's conclusion was that the inherent numerical dissipation from the upwinding in the third order WENO method was so large that it was relatively inefficient for high Reynolds number flows. Their analysis focused on the conservation properties of the algorithms. Over time, the dissipation of the WENO method damped out fine features in the vorticity, ω , thereby reducing the total enstrophy $\Omega = \frac{1}{2} \int \omega^2 dA$ significantly faster than viscosity alone would at high resolution. Since the Arakawa finite difference method is designed mathematically to conserve Ω [1], Naulin and Nielsen concluded that the Arakawa method is preferable to the third order WENO scheme.

Our aim is to reexamine the vortex merger problem to test other aspects of numerical algorithms that may also be important for gyrokinetic and gyrofluid edge plasma simulations. In particular, while previous work focused on the conservation properties of algorithms, other properties such as positivity preservation and the modeling of cascades may be more important for edge turbulence applications. We will extend previous comparisons to include other modern finite volume methods with lower dissipation than a third order upwind method, including some recent high-order methods designed to reduce clipping of extrema by distinguishing smooth extrema from discontinuities. Section 2 lays out the vortex merger problem while Section 3 documents each algorithm and some of their unique properties. We present results for inviscid one-dimensional passive advection and for full two-dimensional nonlinear hydrodynamics in Section 4.

2. The Model System and General Problem Parameters. Due to the universality of nonlinear multi-dimensional advection, we take as our model system the same employed in Ref. [21], namely a two-dimensional neutral fluid vortex merger. We begin with the incompressible, unforced, 2D Navier-Stokes equations:

$$\frac{\partial \vec{v}}{\partial t} + (\vec{v} \cdot \nabla) \vec{v} = \nu \nabla^2 \vec{v} - \nabla p. \quad (2.1)$$

Here, $\vec{v} = v_x \hat{x} + v_y \hat{y}$ is the velocity vector of the fluid confined in the x-y plane, and ν is the kinematic viscosity. (The mass density has been normalized to unity.)

The scalar pressure, p , is governed by the incompressibility condition, $\nabla \cdot \vec{v} = 0$. One can apply the curl operator to express Eq. 2.1 in terms of the scalar vorticity, $\omega = (\nabla \times \vec{v}) \cdot \hat{z}$, and a stream function, $\psi : (\nabla \psi \times \hat{z} = \vec{v})$. The new system of equations is:

$$\frac{\partial \omega}{\partial t} = -[\omega, \psi] + \nu \nabla^2 \omega, \quad (2.2)$$

$$\nabla^2 \psi = -(\omega + \omega_{\text{offset}}). \quad (2.3)$$

A Poisson equation, Eq. 2.3, relates the vorticity and the stream function. The constant ω_{offset} will be discussed in greater detail below. The bracket term, $[\omega, \psi]$, is the standard Poisson bracket, given by

$$[\omega, \psi] = \vec{v} \cdot \nabla \omega = \frac{\partial \omega}{\partial x} \frac{\partial \psi}{\partial y} - \frac{\partial \omega}{\partial y} \frac{\partial \psi}{\partial x}. \quad (2.4)$$

The Poisson bracket term is our canonical nonlinear advection term. While this formulation is specifically for neutral fluids, this nonlinear term is ubiquitous, since many systems, including the gyrokinetic equations for magnetized plasma turbulence, are described by Hamiltonian dynamics that can be expressed in terms of Poisson brackets. For instance, in plasma simulations one is interested in the advection of plasma density n by the $\vec{E} \times \vec{B}$ drift velocity. In such cases $\vec{v}_{E \times B} \cdot \nabla n \propto [\Phi, n]$, where Φ is the electrostatic potential. As such, discretizations of Eq. 2.4 can be applied to plasma physics turbulence problems. All of the algorithms in Section 3 represent Eq. 2.4 differently.

Following the details outlined in [21], we solve Eqs. 2.2 and 2.3 on a doubly-periodic domain of dimensions 10×10 length units. The uniform grid spacing is $h = \Delta x = \Delta y = 10/N$, with N being the number of grid points in either the x or y directions, or $i, j \in [0, N - 1]$.

As our focus will be on the nonlinear Poisson bracket, Eq. 2.4, we will use an explicit “4-point stencil” for the viscous term on the right hand side of Eq. 2.2. Specifically we set

$$\nabla^2 \omega_{i,j} = \frac{1}{h^2} (\omega_{i-1,j} + \omega_{i+1,j} + \omega_{i,j-1} + \omega_{i,j+1} - 4\omega_{i,j}). \quad (2.5)$$

While the authors of Ref. [21] employ an implicit form for Eq. 2.5, there is no significant difference in solutions for the small time steps used in our simulations.

Additionally, to solve the Poisson equation, Eq. 2.3, we use a spectral technique based on the FFTW package [8] that takes advantage of the fact that in k-space the inversion of Eq. 2.3 is algebraic.

The problem is that of two merging Gaussian vortices. They are initially 3 length units apart, have unit maximum and are given in standard Gaussian form as:

$$\omega_l(\vec{r}_{i,j}) = e^{-|\vec{r}_{i,j} - \vec{r}_l|^2 / (2\sigma^2)}. \quad (2.6)$$

As to be expected, $|\vec{r}_{i,j}| = \sqrt{x_i^2 + y_j^2} = h\sqrt{i^2 + j^2}$ is a distance on the grid, \vec{r}_l is the location of the center of Gaussian number l , and σ is the standard deviation of the monopole. Following [21] we define our Gaussian monopoles such that $2\sigma^2 = (0.8)^2$ or $\sigma = 0.8/\sqrt{2} \approx 1.131$. The initial positions of the two peaks are $\vec{r}_1 = (3.5, 5.0)$ and $\vec{r}_2 = (6.5, 5.0)$. After summing these two Gaussians to give the initial vorticity $\omega(\vec{r}) = \omega_1(\vec{r}) + \omega_2(\vec{r})$, we calculate a small offset to the Poisson equation equivalent

to ensuring that the total net vorticity on the grid $\sum_{i,j}(\omega_{i,j} + \omega_{\text{offset}})$ has zero total circulation, so that the Poisson equation is invertible on a periodic domain. Specifically, this offset is written as:

$$\omega_{\text{offset}} = -\frac{h^2}{L^2} \sum_{i,j} \omega_{i,j}. \quad (2.7)$$

While in 2-D hydrodynamics the local vorticity $\omega(r)$ can have either sign, the equivalent Poisson equation in plasmas involves contributions from the ion and electron charge densities, both of which must individually be of only one sign, so positivity preservation of the ion and electron particle density is important.

Simulations are further parametrized by the Reynolds number, defined here as $Re = \sigma v_{max}^0 / (\sqrt{2}\nu)$, where v_{max}^0 is the maximum scalar velocity based on the initial conditions and is approximately 0.25. We examine high Reynolds number ($Re = 100,000$) and inviscid ($\nu = 0$) calculations. The time steps are variable during the course of the simulation and are designed to keep the Courant number at $CFL = v_{max}\Delta t/h = 0.1$, well below the stability limit.

2.1. Time-Advancement Algorithms and Positivity Constraints. To advance the system in time, we employ a third order strong stability preserving Runge-Kutta algorithm (SSP-RK3), as described in Ref. [11]. It is a multi-stage process that is slightly complicated by the coupled equations 2.2-2.3. Write Eq. 2.2 in the generic form

$$\frac{\partial \omega}{\partial t} = \xi(\omega, \psi). \quad (2.8)$$

The SSP-RK3 method evolves this equation in three stages for every total time step, Δt , with the total process a convex average of forward Euler steps. If each Euler step preserves certain properties (such as monotonicity, as the limiters of the various finite volume methods are designed to do), then the convex average of the final method will preserve these properties as well. Defining the vorticity and stream function at time step n as ω^n and ψ^n , then the SSP-RK3 method is

$$\omega^* = \omega^n + \Delta t \xi(\omega^n, \psi^n), \quad (2.9)$$

$$\omega' = \frac{3}{4}\omega^n + \frac{1}{4}\omega^* + \frac{1}{4}\Delta t \xi(\omega^*, \psi^*), \quad (2.10)$$

$$\omega^{n+1} = \frac{1}{3}\omega^n + \frac{2}{3}\omega' + \frac{2}{3}\Delta t \xi(\omega', \psi'). \quad (2.11)$$

It should be noted that each vorticity/stream function pair must satisfy the Poisson equation, Eq. 2.3, so this requires three inversions of Eq. 2.3 per time step.

As mentioned, we run the simulations with a CFL number of 0.1 to ensure that our comparison of spatial algorithms is not muddled by temporal errors. While our focus will be primarily on the different algorithms for evaluating the Poisson bracket, the choice of time-integration scheme in a plasma edge turbulence simulation is important as well.

The SSP-RK3 algorithm is one of a set of schemes designed to maintain the properties (such as monotonicity and positivity) of the simulation throughout the multi-stage integration process. There is a trade-off with these multi-stage algorithms. As the number of stages increases, one can get higher-order accuracy and/or increase

the maximum stable time step, but this of course increases the amount of work to invert the Poisson equation and evaluate the Poisson bracket for each stage.

SSP-RK algorithms are comprised of many Euler-like sub-stages, each of which makes a temporal advance of Δt_1 . In order for an SSP-RK algorithm to preserve monotonicity (and/or positivity, if relevant), it must satisfy a constraint on the maximum allowable sub-stage time advance [10]. Namely, if we define a CFL time scale by $\Delta t_{\text{CFL}} = h/v$, then to preserve monotonicity or positivity, each sub-stage of an SSP-RK method must satisfy a constraint of the form

$$\Delta t_{1,\text{max}} = \frac{\Delta t_{\text{CFL}}}{1 + \alpha}, \quad (2.12)$$

where the value of α depends on the spatial discretization algorithm and generally increases with the order of the spatial discretization. (Possible choices of α and its relation to monotonicity preservation are discussed in Ref. [27]. We will use their suggested value of $\alpha = 4$ for the cases in this paper. Values of α for positivity preservation are discussed below.) The length of each of these substeps is dependent upon the time-advancement algorithm. For instance, in the SSP-RK3 algorithm described in Eqs. 2.9-2.11, $\Delta t_1 = \Delta t$, the total time step. While the ratio of the total time step to a substep, $\Delta t/\Delta t_1$, is algorithmically dependent, it does define the maximum allowable total time step, $\Delta t_{\text{max}} = (\Delta t/\Delta t_1)\Delta t_{1,\text{max}}$. We can define an SSP-based relative efficiency for an algorithm as the ratio of this maximum allowable time step (relative to the Courant time Δt_{CFL}) to the number of stages in the Runge-Kutta scheme:

$$\begin{aligned} \epsilon_t &= \frac{\Delta t_{\text{max}}}{\Delta t_{\text{CFL}}} \frac{1}{N_{\text{stages}}} \\ &= \frac{\Delta t}{\Delta t_1} \frac{1}{N_{\text{stages}}(1 + \alpha)}. \end{aligned} \quad (2.13)$$

While detailed discussion of various optimal SSP algorithms can be found in Ref. [10], Table 2.1 highlights the efficiencies of a few methods. The three-stage RK method we will use, Eqs. 2.9-2.11, is actually more efficient than an optimal five-stage, fourth order method, because it has fewer stages. Even still, a recently discovered ten-stage SSP-RK4 algorithm [12] is nearly twice as efficient as the five-stage version, because it allows for a much larger time step. This analysis, however, only takes into consideration the work required for monotonicity and positivity. In situations where high accuracy is of greater concern, higher-order methods would be favorable.

In any case, we are focused primarily on spatial discretizations of the advection operator and so choose the simpler SSP-RK3 method with small enough time step as not to obscure the spatial discretization analysis. An estimate of the time stepping errors can be seen in Fig. 4.3, where the dissipation from the Arakawa scheme is entirely from time integration errors.

Time advancement can also alter the monotonicity and positivity properties of the system. Two of the Poisson bracket algorithms to be discussed, the XPPM and Suresh-Huynh algorithms, are monotonicity preserving, but they are not necessarily positivity preserving, a desirable property when simulating advection of physically positive quantities (examples include density, temperature, probabilities, or particle distribution functions). To understand the origin of this problem, note that these algorithms are monotonicity preserving (for constant advection in one dimension, if the time step is sufficiently small) in the sense that an initially monotonic region

TABLE 2.1

Stability-based efficiency of various time-advancement algorithms [10], measured by work required to integrate forward in time using the maximum possible time step to preserve stability, monotonicity and positivity. The parameter α depends on the spatial discretization algorithm, and the length of a substep is Δt_1 . While the fourth order 5-stage Runge-Kutta SSP method is more accurate than the lower order SSP-RK3 algorithm, it is slightly less efficient, as measured by Eq. 2.13. The low-storage ten-stage SSP-RK4 method [12] is nearly twice as efficient as the five-stage version of the same order, because it allows for a much larger time step.

Algorithm	Number of Stages N_{stages}	Time Step Size $\Delta t/\Delta t_1$	Relative Efficiency ϵ_t
SSP-RK3	3	1	$0.333/(1 + \alpha)$
5-stage SSP-RK4	5	1.508	$0.302/(1 + \alpha)$
Low-storage 10-stage SSP-RK4	10	6	$0.6/(1 + \alpha)$

remains monotonic after being moved by a time step. (In other words, the algorithms move existing extrema but do not create any new extrema). However, they do not necessarily prevent an existing extremum from being accentuated, and thus they do not necessarily ensure positivity of the solution.

Consider an example of how a monotonicity-preserving algorithm does not necessarily preserve positivity. These Poisson bracket algorithms are exact for smooth low-order solutions, which can in the vicinity of a minimum can be approximated as: $\omega(x, t) = C + (x - vt)^2$. Choose the cell i of width h to have its left edge at $x = 0$. During the time $0 < t < h/(2v)$, the location of the minimum moves from the edge of cell towards its center, so the cell-averaged value $\bar{\omega}_i(t)$ will drop. For certain negative values of C , it is possible that $\bar{\omega}_i$ begins positive but becomes negative. Of course, if $\omega(x)$ should be positive, then $C \geq 0$, and the initial values of $\bar{\omega}_i$ should be such that $\bar{\omega}_i$ remains positive. However, there exist certain initially positive values of $\bar{\omega}_i$ that could lead to negative values of $\bar{\omega}_i$ at later times.

It is possible to guarantee positivity of the cell-averaged solution by supplementing these algorithms with a simple additional limiter. A single forward Euler time step of Eq. 3.3 (neglecting the diffusion term for simplicity) can be written in the form

$$\bar{\omega}_i^{n+1} = \bar{\omega}_i^n - \sum_s \frac{\Delta t \hat{v}_{s,l}}{h} \hat{\omega}_{s,l} \quad (2.14)$$

where $\hat{\omega}_{s,l}$ is the value of ω interpolated (with limiters) to the face s of the cell l at time level n , and $\hat{v}_{s,l}$ is the velocity directed *outward* through face s . (For two-dimensional problems, $s = 1 \dots 4$, and $\hat{v}_{s,l} = -v_{i-1/2,j}$ for s corresponding to the left face of cell (i, j) .) To ensure positivity of the solution at time t_{n+1} , we must satisfy $\sum_s (\Delta t \hat{v}_{s,l}/h) \hat{\omega}_{s,l} \leq \bar{\omega}_i^n$. If we use a conservative upper bound on the left-hand side of this equation based on only the outward fluxes from this cell ($v_{s,l} \geq 0$), and if the edge interpolations are bounded by

$$0 \leq \hat{\omega}_{s,l} \leq (1 + \alpha) \bar{\omega}_i^n \quad (2.15)$$

for some constant α (to be discussed), then positivity is preserved if

$$\sum_s \Delta t \max(\hat{v}_{s,l}, 0)/h \leq 1/(1 + \alpha) \quad (2.16)$$

is satisfied for all cells. The maximum value over all cells of the left-hand side of this equation defines a generalized CFL parameter, so positivity is satisfied if $\text{CFL} \leq$

$1/(1 + \alpha)$. (Alternatively, for a given CFL parameter, one can limit the interpolated value of $\hat{\omega}_{s,l}$ to satisfy $\hat{\omega}_{s,l} < \bar{\omega}_l^n/\text{CFL}$ to ensure positivity.)

To preserve the order of accuracy of the algorithm, one chooses α large enough to avoid limiting any allowable, positive solutions that are sufficiently smooth. For the fourth order XPPM or SuHu-PPM, where solutions are of the form $\omega(x) \propto (x - x_0)^2$ near a zero, the maximum value for the ratio of the edge value to its cell average value, $\hat{\omega}_{s,l}/\bar{\omega}_l^n$, is 4. Therefore, setting $\alpha = 3$ will guarantee positivity. The value of $\alpha = 4$ used in the original SuHu algorithm will also ensure positivity without reducing the order of accuracy, if the edge interpolations are limited to satisfy Eq. 2.15 and the time step satisfies Eq. 2.16.

These results for the choice of α apply in the asymptotic limit as the grid is refined. One might consider more general polynomials near a minimum. As an example for XPPM or SuHu-PPM, which use cubic interpolating polynomials, one could have $\omega(x) \propto (x - x_0)^2(x - x_2)$ locally near a particular cell containing x_0 . A value of $\alpha \approx 3.38$ would be necessary to avoid limiting any case where the point x_2 (where ω changes sign and this local polynomial expansion breaks down) lies at least two cells away. However, as the grid is refined so that the point x_2 is many cells away, then $\alpha = 3$ becomes an acceptable limit to preserve positivity and the order of accuracy.

For the fifth order SuHu algorithm, where $\omega(x) \propto (x - x_0)^4$ is possible, one finds the maximum ratio of edge to cell-averaged values gives $\alpha \approx 6.43$, which would lead to the time step restriction, $\text{CFL} < 0.13$, somewhat stricter than the Suresh-Huynh original choice for just monotonicity, $\text{CFL} < 0.2$. However, if high order minima are rare and it is more likely to have distinct 2nd order minima, $\omega(x) \propto (x - x_0)^2(x - x_2)^2$, then the value of $\alpha \sim 3-4$ may be sufficient in most cases, allowing a larger time step with little loss of accuracy.

The limiter in Eq. 2.15 will preserve positivity if the SSP-RK time step is sufficiently small, with $\text{CFL} < 0.2$ for our parameters. As mentioned in Ref. [6], it is possible to couple finite volume methods with the Zalesak form of the Flux-Corrected Transport method (FCT) [28] to ensure global positivity for larger time steps.

The key idea of FCT is that it expresses a high-order flux as the sum of a low-order upwind flux (which is diffusive enough to guarantee positivity) and a correction flux, which has the form of an anti-diffusion. It then uses a two-step process, where the first step uses the low-order upwind flux to provisionally advance the solution, followed by a second step where it uses as much of the anti-diffusion flux as possible while still preserving positivity (or some other constraint if desired). This two-step process minimizes the amount of numerical diffusion that is needed. As has been pointed out [19], it is not necessary to apply the FCT algorithm on every sub-stage of a Runge-Kutta (RK) time-advancement algorithm. Instead, one could apply FCT only after a full (RK) time step, using the net high-order flux averaged over the Runge-Kutta sub-stages as the starting point for FCT. Furthermore, if an individual RK sub-stage generates negative values, the corrections from the other sub-stages are likely to make it positive again, so that it may be possible to skip the FCT procedure after most time steps. (Negative values during intermediate RK stages can be limited if necessary when calculating fluxes, which still satisfies the conservation law as long as only the fluxes are modified.)

3. The Poisson Bracket Algorithms. We compare six different algorithms for the spatial discretization of the Poisson bracket, Eq. 2.4. Like the authors of [21] we examine Arakawa finite differencing [1] and a third order WENO scheme [14, 18, 26]. Additionally, we compare the behavior of a variant of the Piecewise

Parabolic Method (PPM) with and without a limiter designed to improve the accuracy at smooth extrema [6, 7]. We also consider another WENO method meant to decrease the diffusion of solutions at extrema as laid out in Ref. [27]. While many of these high-resolution upwind algorithms have been tested for 1-D linear and shock problems, here we also compare them side by side with Arakawa on a 2-D nonlinear incompressible flow problem.

3.1. Arakawa Finite Differencing. A classic implementation of the Poisson bracket, Eq. 2.4, is the finite difference method proposed by Arakawa [1, 2]. It has the advantages of a straightforward, closed-form equation and of preserving discrete analogs of the exact linear and quadratic conservation properties of the Poisson bracket. That is, it guarantees conservation of the mean vorticity, which results from $\int d^2r [\psi, \omega] = 0$, and conservation of both quadratic invariants of 2-D hydrodynamics, the energy and the enstrophy Ω (the total squared vorticity, $\Omega = (1/2) \int d^2r |\omega|^2 \rightarrow (1/2)h^2 \sum_{i,j} |\omega_{i,j}|^2$), which result from the discrete equivalents of the properties $\int d^2r [\psi, \omega]\psi = 0$ and $\int d^2r [\psi, \omega]\omega = 0$. In Arakawa's method, the Poisson bracket is represented by:

$$\begin{aligned}
[\omega, \psi] = & -\frac{1}{12h^2} \times \{(\psi_{i,j-1} + \psi_{i+1,j-1} - \psi_{i,j+1} - \psi_{i+1,j+1})(\omega_{i+1,j} + \omega_{i,j}) \quad (3.1) \\
& - (\psi_{i-1,j-1} + \psi_{i,j-1} - \psi_{i-1,j+1} - \psi_{i,j+1})(\omega_{i-1,j} + \omega_{i,j}) \\
& + (\psi_{i+1,j} + \psi_{i+1,j+1} - \psi_{i-1,j} - \psi_{i-1,j+1})(\omega_{i,j+1} + \omega_{i,j}) \\
& - (\psi_{i+1,j-1} + \psi_{i+1,j} - \psi_{i-1,j-1} - \psi_{i-1,j})(\omega_{i,j-1} + \omega_{i,j}) \\
& + (\psi_{i+1,j} - \psi_{i,j+1})(\omega_{i+1,j+1} + \omega_{i,j}) \\
& - (\psi_{i,j-1} - \psi_{i-1,j})(\omega_{i-1,j-1} + \omega_{i,j}) \\
& + (\psi_{i,j+1} - \psi_{i-1,j})(\omega_{i-1,j+1} + \omega_{i,j}) \\
& - (\psi_{i+1,j} - \psi_{i,j-1})(\omega_{i+1,j-1} + \omega_{i,j})\}.
\end{aligned}$$

(The Arakawa algorithm can be considered as having a generalized finite-volume form, where the first four lines in the above equation correspond to the usual fluxes, between cell (i, j) and cells to its east, west, north, and south, but the next four lines represent some kind of corner fluxes, between the center cell and cells to the NE, SW, NW, and SE. In principle one could apply limiters to these fluxes to satisfy monotonicity constraints as one does with other finite-volume methods, but that would break the quadratic conservation properties of the Arakawa method.)

Note that one can consider a 1-D passive advection limit, where $\psi = vy$ and ω is independent of y , in which case the Arakawa formula applied to Eq. 2.2 reduces to:

$$\frac{\partial \omega_i}{\partial t} = -v \frac{\omega_{i+1} - \omega_{i-1}}{2\Delta x} + \nu \frac{\partial^2 \omega}{\partial x^2}, \quad (3.2)$$

which is just the standard form for centered second order finite differencing of the advection term. It is well known that this expression can lead to unphysical oscillations in the solution, particularly in regions of sharp gradients compared to the grid resolution, unless the diffusion term is sufficiently strong [16].

3.2. Finite Volume. While finite difference methods focus on the solution at points, finite volume techniques evolve the integral-averaged quantity within a grid cell. Specifically, we average Eq. 2.2 over a grid cell to obtain the following conservative

equation for the volume-averaged vorticity, $\bar{\omega}_{i,j} \equiv h^{-2} \int \int_{i,j} \omega dx dy$:

$$\frac{\partial \bar{\omega}_{i,j}}{\partial t} = \nu \nabla^2 \bar{\omega}_{i,j} - \frac{1}{h} \left(F_{i+1/2,j}^x - F_{i-1/2,j}^x + F_{i,j+1/2}^y - F_{i,j-1/2}^y \right). \quad (3.3)$$

The form of Eq. 3.3 is the same as Eq. 2.2, with the Poisson bracket term represented as a divergence of fluxes. Fundamentally we are evolving the same equation, but philosophically it is important to remember that these are volume-averaged vorticities, not necessarily the point-value vorticity at the grid center. The fluxes through the grid cell are defined in terms of the velocities and vorticities along grid edges:

$$F_{i+1/2,j}^x = v_{i+1/2,j}^x \times \omega_{i+1/2,j}^*, \quad (3.4)$$

$$F_{i,j+1/2}^y = v_{i,j+1/2}^y \times \omega_{i,j+1/2}^*, \quad (3.5)$$

The choice of ω^* depends upon the direction of the velocity. Finite-volume algorithms allow discontinuities at the grid boundaries, so along each grid edge there exist two values for the vorticity, depending upon which side of the interface one looks. Specifically, let us define the vorticity in grid cell (i, j) at the boundaries of that cell $(i \pm 1/2, j \pm 1/2)$ as

$$\omega_{i,j}^R = \omega_{i+1/2,j}, \quad (3.6)$$

$$\omega_{i,j}^L = \omega_{i-1/2,j}, \quad (3.7)$$

$$\omega_{i,j}^U = \omega_{i,j+1/2}, \quad (3.8)$$

$$\omega_{i,j}^D = \omega_{i,j-1/2}. \quad (3.9)$$

The heart of every finite volume method lies in the determination of these ‘‘right,’’ ‘‘left,’’ ‘‘up,’’ and ‘‘down’’ edge vorticities $\{\omega_{i,j}^R, \omega_{i,j}^L, \omega_{i,j}^U, \omega_{i,j}^D\}$. Once calculated, we choose ω^* to be the edge value of the vorticity in the cell from which a fluid element would originate (the ‘‘upwind’’ direction). Mathematically this is represented in the x-direction as:

$$\omega_{i+1/2,j}^* = \begin{cases} \omega_{i,j}^R & v_{i+1/2,j}^x > 0 \\ \omega_{i+1,j}^L & v_{i+1/2,j}^x < 0 \end{cases}, \quad (3.10)$$

$$\omega_{i-1/2,j}^* = \begin{cases} \omega_{i-1,j}^R & v_{i-1/2,j}^x > 0 \\ \omega_{i,j}^L & v_{i-1/2,j}^x < 0 \end{cases}. \quad (3.11)$$

The y-direction is symmetric:

$$\omega_{i,j+1/2}^* = \begin{cases} \omega_{i,j}^U & v_{i,j+1/2}^y > 0 \\ \omega_{i,j+1}^D & v_{i,j+1/2}^y < 0 \end{cases}, \quad (3.12)$$

$$\omega_{i,j-1/2}^* = \begin{cases} \omega_{i,j-1}^U & v_{i,j-1/2}^y > 0 \\ \omega_{i,j}^D & v_{i,j-1/2}^y < 0 \end{cases}. \quad (3.13)$$

Proceeding with any finite volume method involves the calculation of the vorticities $\{\omega_{i,j}^R, \omega_{i,j}^L, \omega_{i,j}^U, \omega_{i,j}^D\}$ and the edge velocities $\{v_{i \pm 1/2,j}^x, v_{i,j \pm 1/2}^y\}$. The velocity calculation is straightforwardly based upon the stream function. We use a centered

discretization for the velocities $\vec{v} = \nabla\psi \times \hat{z}$ and define $v_{i,j}^x = (\psi_{i,j+1} - \psi_{i,j-1})/2h$ and $v_{i,j}^y = -(\psi_{i+1,j} - \psi_{i-1,j})/2h$. We then interpolate linearly to the cell edges:

$$v_{i\pm 1/2,j}^x = \frac{1}{2} (v_{i,j}^x + v_{i\pm 1,j}^x), \quad (3.14)$$

$$v_{i,j\pm 1/2}^y = \frac{1}{2} (v_{i,j}^y + v_{i,j\pm 1}^y). \quad (3.15)$$

[Although we will be using higher order methods for the interpolations of ω to the cell faces, the final algorithm is formally only second order accurate in multiple dimensions because we are using simple mid-point evaluations of the face-averaged fluxes such as $F_{i+1/2,j}^x = \langle v^x \omega^* \rangle_{i+1/2,j} = (1/h) \int_{y_{j-1/2}}^{y_{j+1/2}} dy v^x(x_{i+1/2}, y) \omega^*(x_{i+1/2}, y) \approx v_{i+1/2,j}^x \omega_{i+1/2,j}^*$. The higher order aspects of the ω interpolations are still useful for reducing numerical dissipation. Extensions to full higher-order accuracy can be done, see for example Ref. [5].]

The edge vorticity calculation is unique to each finite volume method. However, all techniques involve two basic steps:

1. Interpolate from cell-averaged vorticities to cell boundaries. This is usually accomplished by fitting surrounding vorticity values to a high-order polynomial function. Call these first-stab estimates $\{\omega_{i,j}^{\dagger R}, \omega_{i,j}^{\dagger L}, \omega_{i,j}^{\dagger U}, \omega_{i,j}^{\dagger D}\}$.
2. Limit these initial interpolations, $\{\omega_{i,j}^{\dagger R}, \omega_{i,j}^{\dagger L}, \omega_{i,j}^{\dagger U}, \omega_{i,j}^{\dagger D}\}$, according to a set of pre-defined conditions. These initial estimates may (or may not) be altered during the process, but the result is the set $\{\omega_{i,j}^R, \omega_{i,j}^L, \omega_{i,j}^U, \omega_{i,j}^D\}$ to be used in Eqs. 3.10-3.13. Most often this is to prevent over- and under-shoots in the solution, but the exact mechanism varies with the algorithm.

Once one has determined the edge velocities $\{v_{i\pm 1/2,j}^x, v_{i,j\pm 1/2}^y\}$ from Eqs. 3.14 and 3.15 and $\{\omega_{i,j}^R, \omega_{i,j}^L, \omega_{i,j}^U, \omega_{i,j}^D\}$ from the method-specific interpolation/limiting algorithm, the procedure continues by using Eqs. 3.10-3.13 to determine the values of ω^* . Hence, one can use Eqs. 3.4 and 3.5 to determine the appropriate flux in Eq. 3.3. The coding to evolve the equation in time is identical for a finite volume method as for a finite difference equation, with the Poisson bracket term in Eq. 2.2 replaced by the flux terms in Eq. 3.3.

We now describe the algorithms for the interpolation and limiting of edge-valued vorticities in the x -direction. The y -direction is symmetric under the transformation: $\{R, L, i, j\} \rightarrow \{U, D, j, i\}$.

3.2.1. Third-Order WENO (WENO3). We employ the third order WENO method as suggested by Kurganov and Levy [14]. The interpolations are parabolic, while the limiters rigorously preserve monotonicity and reduce to piecewise constant values at extrema. First we determine if the current grid cell vorticity is either a maximum or a minimum. If so, we set the edge vorticity values to be equal to the cell-averaged value:

$$\begin{aligned} &\text{if } \text{sign}(\bar{\omega}_{i+1,j} - \bar{\omega}_{i,j}) \neq \text{sign}(\bar{\omega}_{i,j} - \bar{\omega}_{i-1,j}) \text{ then} \\ &\quad \omega_{i,j}^R = \omega_{i,j}^L = \bar{\omega}_{i,j}. \end{aligned} \quad (3.16)$$

Otherwise, we are not at an extremum and we proceed with our parabolic interpolations for $\{\omega_{i,j}^{\dagger R}, \omega_{i,j}^{\dagger L}\}$. Defining $\Delta_i^1 = (\bar{\omega}_{i+1,j} - \bar{\omega}_{i-1,j})$ and $\Delta_i^2 = (\bar{\omega}_{i+1,j} - 2\bar{\omega}_{i,j} + \bar{\omega}_{i-1,j})$

we have:

$$\omega_{i,j}^{\dagger R} = \bar{\omega}_{i,j} + \frac{1}{4}\Delta_i^1 + \frac{1}{12}\Delta_i^2, \quad (3.17)$$

$$\omega_{i,j}^{\dagger L} = \bar{\omega}_{i,j} - \frac{1}{4}\Delta_i^1 + \frac{1}{12}\Delta_i^2. \quad (3.18)$$

The limiters to ensure that rigorously no over- and under-shoots appear in the solution proceed thusly:

$$\begin{aligned} \text{if } \text{sign}(\bar{\omega}_{i,j} - \omega_{i,j}^{\dagger L}) \neq \text{sign}(\omega_{i,j}^{\dagger L} - \bar{\omega}_{i-1,j}) \text{ then} \\ \omega_{i,j}^L = \bar{\omega}_{i-1,j}, \end{aligned} \quad (3.19)$$

$$\omega_{i,j}^R = \bar{\omega}_{i-1,j} + \frac{5}{2}(\bar{\omega}_{i,j} - \bar{\omega}_{i-1,j}), \quad (3.20)$$

$$\begin{aligned} \text{else if } \text{sign}(\bar{\omega}_{i,j} - \omega_{i,j}^{\dagger R}) \neq \text{sign}(\omega_{i,j}^{\dagger R} - \bar{\omega}_{i+1,j}) \text{ then} \\ \omega_{i,j}^R = \bar{\omega}_{i+1,j}, \end{aligned} \quad (3.21)$$

$$\omega_{i,j}^L = \bar{\omega}_{i+1,j} + \frac{5}{2}(\bar{\omega}_{i,j} - \bar{\omega}_{i+1,j}). \quad (3.22)$$

If neither of these conditions is met, the limiter is not triggered and the initial estimates are used: $\omega_{i,j}^{R,L} = \omega_{i,j}^{\dagger R,L}$.

The limiters in the WENO3 scheme serve to ensure that edge vorticity values lie between the values of adjacent cells and rigorously preserve this monotonicity.

3.2.2. The Piecewise Parabolic Method (PPM). Originally developed by Colella and Woodward [7], the Piecewise Parabolic Method (PPM) is a finite volume method that seeks solutions comprised of local parabolas. Like the WENO3 method in Sec. 3.2.1, it becomes piecewise constant at extrema. In smooth regions away from extrema we use a fourth-order variant of PPM (PPM4) for the spatial discretizations, as described in Ref. [6]. As with all of the other spatial-differencing algorithms, we will use Runge-Kutta time advancement (instead of the original PPM time-advancement method), as also done in [5]. In our notation the initial interpolation becomes:

$$\omega_{i,j}^{\dagger R} = \frac{7}{12}(\bar{\omega}_{i,j} + \bar{\omega}_{i+1,j}) - \frac{1}{12}(\bar{\omega}_{i-1,j} + \bar{\omega}_{i+2,j}) \quad (3.23)$$

$$\omega_{i,j}^{\dagger L} = \frac{7}{12}(\bar{\omega}_{i,j} + \bar{\omega}_{i-1,j}) - \frac{1}{12}(\bar{\omega}_{i-2,j} + \bar{\omega}_{i+1,j}). \quad (3.24)$$

Next, PPM4 completes the interpolation by applying monotonicity constraints to eliminate solution oscillations. These limiters are triggered when monotonicity is violated. Explicitly these conditions are:

$$\begin{aligned} \text{if } (\bar{\omega}_{i,j} - \omega_{i,j}^{\dagger R})(\bar{\omega}_{i+1,j} - \omega_{i,j}^{\dagger R}) > 0 \text{ then} \\ \omega_{i,j}^{\dagger R} = (\bar{\omega}_{i,j} + \bar{\omega}_{i+1,j})/2 - (\delta\omega_{i+1}^* - \delta\omega_i^*)/6 \end{aligned} \quad (3.25)$$

$$\begin{aligned} \text{if } (\bar{\omega}_{i-1,j} - \omega_{i,j}^{\dagger L})(\bar{\omega}_{i,j} - \omega_{i,j}^{\dagger L}) > 0 \text{ then} \\ \omega_{i,j}^{\dagger L} = (\bar{\omega}_{i-1,j} + \bar{\omega}_{i,j})/2 - (\delta\omega_i^* - \delta\omega_{i-1}^*)/6, \end{aligned} \quad (3.26)$$

with the definitions:

$$\delta\omega_i^* = \begin{cases} \text{sign}(\Delta_i^1) \min(\frac{1}{2}|\Delta_i^1|, 2|\Delta_i^L|, 2|\Delta_i^R|) & \Delta_i^R \Delta_i^L > 0 \\ 0 & \text{otherwise} \end{cases}, \quad (3.27)$$

with $\Delta_i^1 = \bar{\omega}_{i+1,j} - \bar{\omega}_{i-1,j}$, $\Delta_i^L = \bar{\omega}_{i,j} - \bar{\omega}_{i-1,j}$ and $\Delta_i^R = \bar{\omega}_{i+1,j} - \bar{\omega}_{i,j}$. As mentioned, limiters only apply when monotonicity is violated in the edge interpolations. In smooth regions away from extrema, there will be no change to $\omega_{i,j}^{\dagger\{R,L\}}$ from Eqs. 3.25-3.26.

The next step in the PPM4 method is to apply additional limiters, which depend on whether or not the solution is a local extremum, defined not only by the averaged values, $\bar{\omega}$, but also by the initial interpolations, ω^\dagger . Specifically the solution is a local maximum or minimum if

$$\left(\omega_{i,j}^{\dagger R} - \bar{\omega}_{i,j}\right) \left(\bar{\omega}_{i,j} - \omega_{i,j}^{\dagger L}\right) \leq 0 \text{ or } \left(\bar{\omega}_{i-1,j} - \bar{\omega}_{i,j}\right) \left(\bar{\omega}_{i,j} - \bar{\omega}_{i+1,j}\right) \leq 0. \quad (3.28)$$

If the initial interpolation satisfies Eq. 3.28, we apply the piecewise constant limiters $\omega_{i,j}^R = \omega_{i,j}^L = \bar{\omega}_{i,j}$, like Eq. 3.16.

Away from extrema we apply further limiters that constrain the amount of variation in the local reconstructions, testing both edge interpolations, $\{\omega_{i,j}^{\dagger R}, \omega_{i,j}^{\dagger L}\}$:

$$\begin{aligned} \text{if } \left|\omega_{i,j}^{\dagger R,L} - \bar{\omega}_{i,j}\right| &\geq n \left|\bar{\omega}_{i,j} - \omega_{i,j}^{\dagger L,R}\right| \\ \text{then } \omega_{i,j}^{R,L} &= \bar{\omega}_{i,j} + n \left(\bar{\omega}_{i,j} - \omega_{i,j}^{\dagger L,R}\right) \\ \text{else } \omega_{i,j}^{R,L} &= \omega_{i,j}^{\dagger R,L}. \end{aligned} \quad (3.29)$$

We have used $n = 3$ here to avoid limiting the accuracy of monotonic cubic interpolations. It has been pointed out [19] that $n = 2$ can be used (as is done in the original Piecewise Parabolic Method) with only rare and slight loss of local accuracy (because locally quadratic behavior near flat regions is more generic than locally cubic behavior).

While more complex limiters for the parabolic interpolant exist [6], we follow Refs. [5] and [19] and choose the limiters described above.

3.2.3. Smooth Extrema Preserving Piecewise Parabolic Method (XPPM).

A major limitation of many finite volume methods is a reduction of accuracy at smooth extrema [6, 27]. As noted above, many algorithms resort to first-order (piecewise constant) solutions at local maxima and minima. One method of improving the performance of these methods is to apply a new limiter at extrema, to replace Eq. 3.28. A recently-developed limiter [6] uses estimates of the second derivative to relax the constraints on solutions at extrema, thereby allowing for solutions with higher accuracy.

The method is identical to that described in Sec. 3.2.2, with the exception of the limiters. Interpolation proceeds according to Eq. 3.23.

Instead of Eqs. 3.25-3.26, we use the following (which correct a typo in Eq. (19) of Ref. [6]):

$$\begin{aligned} \text{if } \left(\bar{\omega}_{i,j} - \omega_{i,j}^{\dagger R}\right) \left(\bar{\omega}_{i+1,j} - \omega_{i,j}^{\dagger R}\right) &> 0 \text{ then} \\ \omega_{i,j}^{\dagger R} &= \frac{1}{2} \left(\bar{\omega}_{i,j} + \bar{\omega}_{i+1,j}\right) - \frac{h^2}{6} (D^2\omega)_{i+1/2,j,\text{lim}} \end{aligned} \quad (3.30)$$

$$\begin{aligned} \text{if } \left(\bar{\omega}_{i-1,j} - \omega_{i,j}^{\dagger L}\right) \left(\bar{\omega}_{i,j} - \omega_{i,j}^{\dagger L}\right) &> 0 \text{ then} \\ \omega_{i,j}^{\dagger L} &= \frac{1}{2} \left(\bar{\omega}_{i-1,j} + \bar{\omega}_{i,j}\right) - \frac{h^2}{6} (D^2\omega)_{i-1/2,j,\text{lim}}. \end{aligned} \quad (3.31)$$

The value of $(D^2\omega)_{\{i,j\}\pm 1/2,\text{lim}}$ is determined from approximations to the second derivative:

$$\begin{aligned}(D^2\omega)_{i+1/2,j} &= \frac{3}{h^2} (\bar{\omega}_{i,j} - 2\omega_{i,j}^{\dagger R} + \bar{\omega}_{i+1,j}), \\(D^2\omega)_{i+1/2,j,L} &= \frac{1}{h^2} (\bar{\omega}_{i-1,j} - 2\bar{\omega}_{i,j} + \bar{\omega}_{i+1,j}), \\(D^2\omega)_{i+1/2,j,R} &= \frac{1}{h^2} (\bar{\omega}_{i,j} - 2\bar{\omega}_{i+1,j} + \bar{\omega}_{i+2,j}).\end{aligned}\quad (3.32)$$

If $(D^2\omega)_{i+1/2,j}$, $(D^2\omega)_{i+1/2,j\{L,R\}}$ all have the same sign, then

$$\begin{aligned}(D^2\omega)_{i+1/2,j,\text{lim}} &= s \min \left\{ C \left| (D^2\omega)_{i+1/2,j,L} \right|, C \left| (D^2\omega)_{i+1/2,j,R} \right|, \left| (D^2\omega)_{i+1/2,j} \right| \right\}; \\s &= \text{sign} \left((D^2\omega)_{i+1/2,j} \right).\end{aligned}\quad (3.33)$$

Otherwise, $(D^2\omega)_{i+1/2,j,\text{lim}} = 0$. Following Ref. [6] we select $C = 1.25$. To determine $(D^2\omega)_{i-1/2,j,\text{lim}}$, let $i \rightarrow i - 1$ and $\omega_{i,j}^{\dagger R} \rightarrow \omega_{i,j}^{\dagger L}$.

While modifications to limiters away from extrema exist [6], we follow Refs. [5] and [19] and focus only on algorithmic modifications at extrema as defined by Eq. 3.28. Recent work on further improvements to the XPPM approach can be found in Ref. [20].

Away from these points our implementation of the XPPM method continues as the PPM algorithm with Eq. 3.29. When Eq. 3.28 is satisfied, the smooth extrema limiter [6] initiates, beginning with several approximations to the second derivative:

$$\begin{aligned}(D^2\omega)_{i,j} &= \frac{6}{h^2} (\omega_{i,j}^{\dagger L} - 2\bar{\omega}_{i,j} + \omega_{i,j}^{\dagger R}), \\(D^2\omega)_{i,j,C} &= \frac{1}{h^2} (\bar{\omega}_{i-1,j} - 2\bar{\omega}_{i,j} + \bar{\omega}_{i+1,j}), \\(D^2\omega)_{i,j,L} &= \frac{1}{h^2} (\bar{\omega}_{i-2,j} - 2\bar{\omega}_{i-1,j} + \bar{\omega}_{i,j}), \\(D^2\omega)_{i,j,R} &= \frac{1}{h^2} (\bar{\omega}_{i,j} - 2\bar{\omega}_{i+1,j} + \bar{\omega}_{i+2,j}).\end{aligned}\quad (3.34)$$

If $(D^2\omega)_{i,j}$, $(D^2\omega)_{i,j\{C,L,R\}}$ all have the same sign, then

$$(D^2\omega)_{i,j,\text{lim}} = s \min \left\{ C \left| (D^2\omega)_{i,j,L} \right|, C \left| (D^2\omega)_{i,j,R} \right|, C \left| (D^2\omega)_{i,j,C} \right|, \left| (D^2\omega)_{i,j} \right| \right\}.\quad (3.35)$$

Otherwise, $(D^2\omega)_{i,j,\text{lim}} = 0$. The coefficient $s \equiv \text{sign} \left[(D^2\omega)_{i,j} \right]$, and C is a numerical constant. Following Ref. [6], we set $C = 1.25$. The final limiter is

$$\omega_{i,j}^{R,L} = \begin{cases} \bar{\omega}_{i,j} + \left(\omega_{i,j}^{\dagger R,L} - \bar{\omega}_{i,j} \right) \frac{(D^2\omega)_{i,j,\text{lim}}}{(D^2\omega)_{i,j}} & (D^2\omega)_{i,j} \neq 0 \\ \bar{\omega}_{i,j} & (D^2\omega)_{i,j} = 0 \end{cases}.\quad (3.36)$$

The net effect of this is that at smooth extrema, the limiters turn off and the high-order interpolation for $\omega_{i,j}^{\dagger R,L}$ is used, while at poorly resolved non-smooth extrema, it switches to approach a piecewise constant reconstruction.

3.2.4. Fifth Order Suresh-Huynh WENO (SuHu). Another method to ameliorate the loss of accuracy at extrema comes from Suresh and Huynh [27]. The technique uses geometrical arguments to maintain the order of the underlying interpolation method at smooth local extrema and guarantees global positivity for linear problems. (It should be noted that no such guarantee is made on nonlinear problems such as the Navier-Stokes problem here.) While it is possible to apply the SuHu limiter to any initial edge interpolation, we employ a fifth order quartic polynomial:

$$\omega_{i,j}^{\dagger R} = \frac{1}{60} (2\bar{\omega}_{i-2,j} - 13\bar{\omega}_{i-1,j} + 47\bar{\omega}_{i,j} + 27\bar{\omega}_{i+1,j} - 3\bar{\omega}_{i+2,j}), \quad (3.37)$$

$$\omega_{i,j}^{\dagger L} = \frac{1}{60} (2\bar{\omega}_{i+2,j} - 13\bar{\omega}_{i+1,j} + 47\bar{\omega}_{i,j} + 27\bar{\omega}_{i-1,j} - 3\bar{\omega}_{i-2,j}). \quad (3.38)$$

As shown by Eqs. 3.37 and 3.38, the polynomial interpolations are symmetric.

After interpolation, each edge value is limited by a function of the vorticities of the cell and its neighbors. Let us call that algorithm $l(\omega^\dagger, \omega_{-2}, \omega_{-1}, \omega_0, \omega_{+1}, \omega_{+2})$. Symmetry allows the same function to be used for each of the edge values $\{\omega_{i,j}^R, \omega_{i,j}^L, \omega_{i,j}^U, \omega_{i,j}^D\}$:

$$\omega_{i,j}^R = l\left(\omega_{i,j}^{\dagger R}, \bar{\omega}_{i-2,j}, \bar{\omega}_{i-1,j}, \bar{\omega}_{i,j}, \bar{\omega}_{i+1,j}, \bar{\omega}_{i+2,j}\right), \quad (3.39)$$

$$\omega_{i,j}^L = l\left(\omega_{i,j}^{\dagger L}, \bar{\omega}_{i+2,j}, \bar{\omega}_{i+1,j}, \bar{\omega}_{i,j}, \bar{\omega}_{i-1,j}, \bar{\omega}_{i-2,j}\right). \quad (3.40)$$

The function $\hat{\omega} = l(\omega^\dagger, \omega_{-2}, \omega_{-1}, \omega_0, \omega_{+1}, \omega_{+2})$ begins by calculating a ‘‘monotonicity-preserving’’ limit for ω , ω^{MP} :

$$\omega^{MP} = \omega_0 + \text{minmod}[(\omega_{+1} - \omega_0), \alpha(\omega_0 - \omega_{-1})]. \quad (3.41)$$

Higher values of α allow higher order accuracy in the interpolations but also reduce the allowable time step that will ensure monotonicity in an SSP Runge-Kutta algorithm. As suggested in Ref. [27], we set $\alpha = 4$. $\text{minmod}[x, y]$ is the median of x, y , and 0: $\text{minmod}[x, y] = (\text{sign}(x) + \text{sign}(y)) \times \min(|x|, |y|) / 2$. The multiple-argument $\text{minmod}[z_1, \dots, z_k]$ function returns the smallest argument if all are positive, the largest argument if all are negative and zero otherwise. Part of the efficiency of the SuHu limiter is a check to see if any limiters must be applied to the edge value. If we are within a certain tolerance ($\epsilon = 10^{-10}$), we do not have to apply any limiters to ω^\dagger :

$$\text{if } (\omega^\dagger - \omega_0) (\omega^\dagger - \omega^{MP}) < \epsilon \text{ then } \hat{\omega} = \omega^\dagger. \quad (3.42)$$

Otherwise we proceed with the limiting procedure, beginning with some second differences:

$$d_{-1} = \omega_{-2} + \omega_0 - 2\omega_{-1}, \quad (3.43)$$

$$d_0 = \omega_{-1} + \omega_{+1} - 2\omega_0, \quad (3.44)$$

$$d_1 = \omega_0 + \omega_{+2} - 2\omega_{+1}. \quad (3.45)$$

We then modify these differences:

$$d_+^{M4} = \text{minmod}[4d_0 - d_1, 4d_1 - d_0, d_0, d_1], \quad (3.46)$$

$$d_-^{M4} = \text{minmod}[4d_{-1} - d_0, 4d_0 - d_{-1}, d_0, d_{-1}]. \quad (3.47)$$

The limiting algorithm continues by calculating various geometric factors. The meaning and reasoning behind these are given in Ref. [27], but for simplicity, let us leave

these details and proceed with the algorithm.

$$\omega^{UL} = \omega_0 + \alpha(\omega_0 - \omega_{-1}), \quad (3.48)$$

$$\omega^{AVG} = \frac{1}{2}(\omega_0 + \omega_{+1}), \quad (3.49)$$

$$\omega^{MD} = \omega^{AVG} - \frac{1}{2}d_+^{M4}, \quad (3.50)$$

$$\omega^{LC} = \omega_0 + \frac{1}{2}(\omega_0 - \omega_{-1}) + \frac{4}{3}d_-^{M4}. \quad (3.51)$$

These values are used to calculate the minimum and maximum allowable values for ω^\dagger :

$$\omega_{\min} = \max[\min(\omega_0, \omega_{+1}, \omega^{MD}), \min(\omega_0, \omega^{UL}, \omega^{LC})], \quad (3.52)$$

$$\omega_{\max} = \min[\max(\omega_0, \omega_{+1}, \omega^{MD}), \max(\omega_0, \omega^{UL}, \omega^{LC})]. \quad (3.53)$$

Finally one modifies the edge value by taking the median of these two limits:

$$\hat{\omega} = \text{median}(\omega^\dagger, \omega_{\min}, \omega_{\max}). \quad (3.54)$$

Eqs. 3.41-3.54 make up the entire limiting function, $l(\omega^\dagger, \omega_{-2}, \omega_{-1}, \omega_0, \omega_{+1}, \omega_{+2})$, which can be applied to each edge interpolation, $\{\omega_{i,j}^{\dagger R}, \omega_{i,j}^{\dagger L}\}$, using Eqs. 3.39 and 3.40. Even though the SuHu method seems outwardly complex, the symmetries in both the interpolation and limiting phases of the algorithm allow for easy coding. The SuHu limiter maintains the order of the underlying algorithm even at extrema and guarantees monotonicity preservation for one-dimensional linear passive advection problems, but global constraints might not be enforced in the presence of nonlinearities. Positivity is also insured if a time-step restriction is satisfied.

3.2.5. Suresh-Huynh/Piecewise Parabolic Method Hybrid (SuHu-PPM).

To form a simple hybrid technique, we use the PPM interpolations as laid forth in Sec. 3.2.2 but the limiters proposed by Ref. [27] and detailed in Sec. 3.2.4. This is not out of the spirit of the SuHu method, as Ref. [27] states that the limiters therein can be applied to any type of interpolation, and here we apply them to the fourth order PPM interpolations. (This hybrid choice was motivated by results found in Sec. 4.) Explicitly, we use Eqs. 3.23-3.24 for the initial interpolations, but apply the limiters in Eqs. 3.39-3.40. (The value of α in Eq. 3.41 can be set to 3 when using the fourth order initial interpolations of PPM, which would allow a slightly larger time step, though we leave $\alpha = 4$ for the results in this paper.)

3.3. On Parallelization. Considering the recent advancements in high-powered computation, it is prudent to mention a few words on an algorithm's ability to parallelize efficiently over multiple processors and overall performance. Because higher order methods generally involve more FLOPS per memory access (and can reuse data in cache from adjacent cells in the higher-order interpolations), one will generally expect higher order methods to perform better on modern computer architectures where memory bandwidth is usually more of a bottleneck than the CPU.

Even though the test problem contained here was run on a single serial processor, all of the methods tested lend themselves to domain decomposition. The number of requisite ghost cells each processor needs depends upon the algorithm. The Arakawa method needs one layer to achieve its second-order accuracy, as cell (i, j) needs information from its immediate neighbors $(i \pm 1, j \pm 1)$. A higher-order Arakawa method,

such as the one discussed in Sec. 4.3, would require more information. For instance a fourth-order method needs $(i \pm 2, j \pm 2)$ [22] and, thus, two layers of ghost cells. The WENO3 method similarly only needs one layer of extra information per processor to calculate edge values $\{\omega_{i,j}^R, \omega_{i,j}^L, \omega_{i,j}^U, \omega_{i,j}^D\}$. However, in order to properly calculate the fluxes through the faces of cell (i, j) , we must know the edge vorticity values on both sides of the interfaces, or specifically $\{\omega_{i-1,j}^R, \omega_{i+1,j}^L, \omega_{i,j-1}^U, \omega_{i,j+1}^D\}$. This implies that each processor must first calculate all of the $\{\omega_{i,j}^R, \omega_{i,j}^L, \omega_{i,j}^U, \omega_{i,j}^D\}$ on its sub-grid using one ghost layer, then exchange information with its neighbors to be able to calculate the correct fluxes and advance the solution and then exchange the newly calculated vorticity values $\omega_{i,j}^{n+1}$ with its neighboring processors. Alternatively, one could subdivide the grid into domains with two ghost layers. This would allow each processor to accurately calculate all of the $\{\omega_{i,j}^R, \omega_{i,j}^L, \omega_{i,j}^U, \omega_{i,j}^D\}$ and $\{\omega_{i-1,j}^R, \omega_{i+1,j}^L, \omega_{i,j-1}^U, \omega_{i,j+1}^D\}$ on its sub-domain, thereby consolidating two message passing events into one. In most cases this will probably be better as it will reduce the total latency, but the choice of which method to use would depend on details of the specific computer architecture (such as the latency and bandwidth of the interprocessor communication relative to CPU speed), and thus the relative efficiencies of having a slightly larger domain on each processor and the speed at which processors can exchange information. Similarly, the PPM, XPPM and SuHu methods could all run with two ghost cells, but it will probably be more efficient to have three.

4. Results.

4.1. One-Dimensional Advection. By skipping the Poisson equation inversion, Eq. 2.3, with the initial conditions $v_x = 0$, $v_y = 1$ and without viscosity, $\nu = 0$, the equations of motion become simple advection. Eq. 2.1 reduces to a wave equation. Solutions are easily tractable, as the initial disturbance propagates with constant unit velocity. After a transit period ($t = 10m$, $m \in \mathbb{N}$), any distortions from the initial conditions are properties of the numerical algorithm only.

A standard test problem in one-dimensional advection is the transit of a Gaussian and a step-function pulse (see, for instance Ref. [16]). After initialization, we let these pulses transit through one period of motion and then compare one-dimensional slices of the solution to the initial conditions. Fig. 4.1 shows the overlaid initial (solid) and final (dotted) solutions for the six methods (all with $\text{CFL} = v\Delta t/\Delta x = 0.1$ so that time discretization errors are negligible). The Arakawa result displays a large phase error, as evidenced by the mismatch in peak locations. Furthermore, one clearly sees spurious oscillations that build up behind the pulses and propagate through the mesh. These kinds of errors could be particularly significant in cases where nonlinear interactions of coherent structures are important, so one would like to minimize phase errors between various components of the solution. This may require very high resolution with an Arakawa method and/or including relatively strong viscosity. Much of the total error in the Arakawa solution originates from the step function, which propagates through the mesh and adversely affects the smooth pulse. However, even without the step function, the solution would still oscillate unphysically negative (for a Gaussian with a Full Width at Half Maximum (FWHM) of 1, and 12.8 points per FWHM, the oscillations in the solution will reach -6% of the peak value after propagating a distance of 10). In many respects this is not unexpected, as in the one-dimensional limit the Arakawa discretization reduces to a second-order center-differenced method (see Eq. 3.2). The WENO3 and PPM4 solutions display significant diffusion near the Gaussian peak (PPM4 is slightly less

diffusive), but both show relatively small phase errors and asymmetries. However, the finite volume methods with smooth extrema preserving limiters reproduce the initial conditions quite well, with peak amplitude drops of $\sim 10^{-7}$. Additionally, no finite volume result has spurious oscillations; the solution stays above the initial minimum, preserving positivity.

As another test, we can initialize the vorticity to be a cosine wave in the direction of propagation with wave number k , $\omega_0(y) = \cos(ky)$. After the wave has moved a time $T = 2\pi/(kv)$, the numerical solution should again be a cosine wave, with a possible phase shift and modified amplitude: $\omega(y, T = 2\pi/(kv)) = a \cos(ky + \phi)$. Assuming no damping, the numerical cosine wave should satisfy $\langle \omega^2 \rangle = 1/2$. Let us characterize a numerical damping factor γ_n by

$$\langle \omega^2 \rangle (T) = \frac{1}{2} e^{-2\gamma_n T}. \quad (4.1)$$

An ideal algorithm would have small values for both γ_n and ϕ for all wave numbers $k\Delta y/\pi \in (0, 1)$. If one is primarily interested in large-scale physical features, low damping and phase errors at small k would be particularly desirable, even though in general, the phase and amplitude errors are not simple functions of only the wave number for most of these algorithms because of the nonlinearities in the limiters they use.

The phase error can be determined by the law of cosines, through the dot product of $\vec{\omega}_0 = \omega(y, t = 0)$ and $\vec{\omega}_T = \omega(y, t = T)$:

$$\cos(\phi) = \frac{\vec{\omega}_0 \cdot \vec{\omega}_T}{\omega_0 \omega_T}. \quad (4.2)$$

The phase can also be found less robustly by a non-linear least-squares fit of $\omega(y, T)$. We define the phase error E_ϕ as $(1 - \cos(\phi))/2$, which maps $\phi \in [0, \pi] \rightarrow E_\phi \in [0, 1]$. As shown in Fig. 4.2, the phase errors of the higher-order algorithms (XPPM, SuHu, SuHu-PPM) are all much lower than those of the WENO3, PPM4 and Arakawa methods across all wavelengths.

We display $\langle \omega^2 \rangle$ with calculated numerical damping rates, γ_n , of each method in Fig. 4.3. As the Arakawa spatial discretization method does not introduce any dissipation itself, the curve labelled Arakawa represents the small residual numerical dissipation from the third-order time-stepping method. While the WENO3 and PPM4 methods have relatively high rates of numerical dissipation, the extremum-preserving methods have significantly lower dissipation. XPPM drops to the minimum possible dissipation at long-wavelengths, because it reverts to an undamped centered method in well-resolved regions. The observation that SuHu is better than XPPM at higher wavenumber while XPPM is better at lower wavenumber motivated the choice of the SuHU-PPM hybrid, which uses the centered fourth order interpolation of the initial step of PPM in combination with the SuHu limiters. The SuHu-PPM method nearly overlaps the true solution as it approaches the minimum level of dissipation for long-wavelengths of $k\Delta y \lesssim \pi/3$.

To understand some of the asymptotic scalings for the numerical dissipation rate in Fig. 4.3, consider the following. If Eqs. 3.17-3.18 from the WENO3 method were used to interpolate the boundary values with no further limiters, then this would correspond to an upwind-biased third order method, which would have a numerical damping that asymptotically scales as $\gamma_n \sim |v|(\Delta x)^3 k^4$. As seen in damping rates of Fig. 4.3, the additional limiters in the WENO3 method give larger damping,

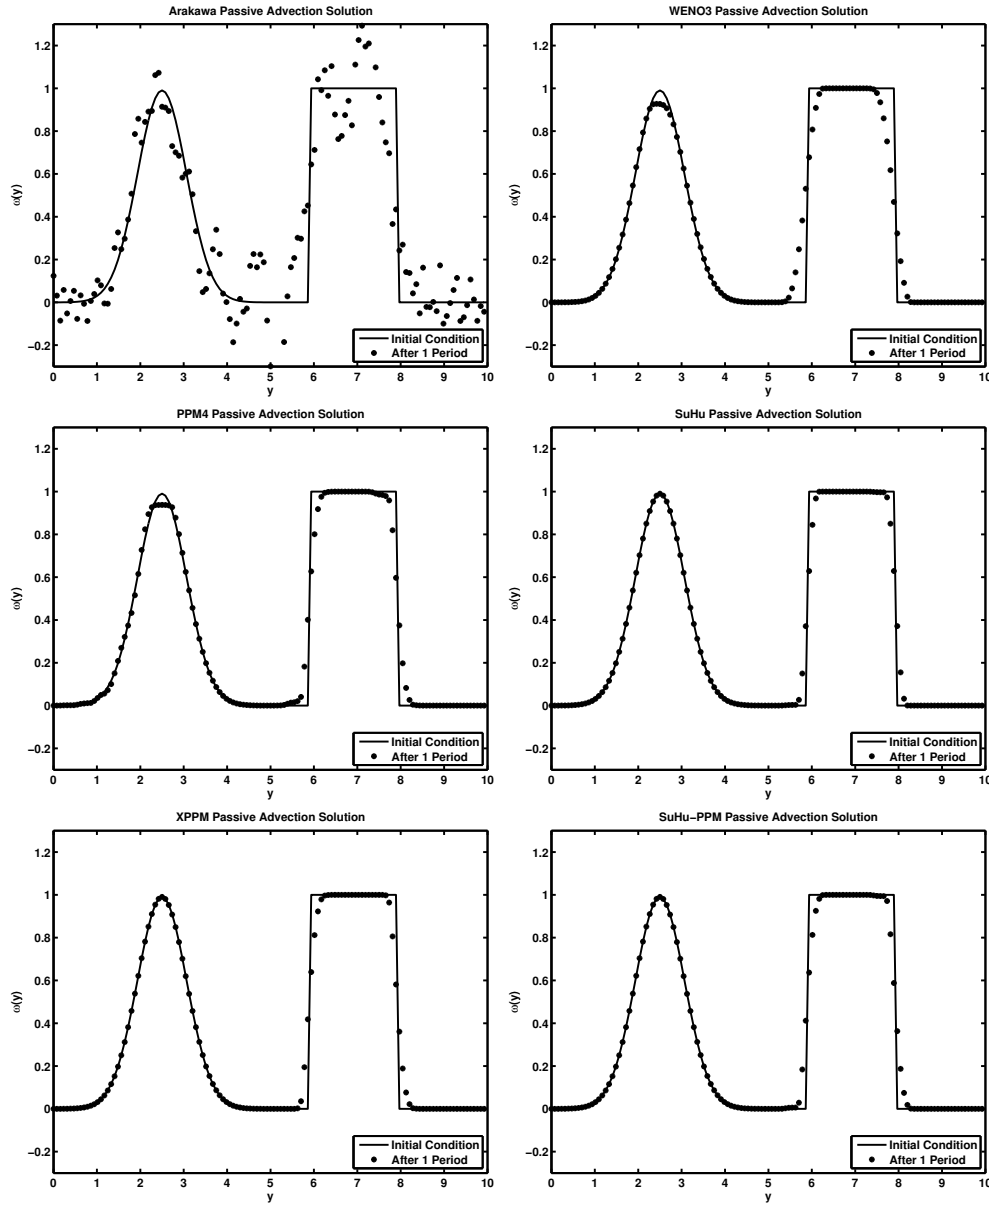


FIG. 4.1. Advection of a Gaussian and step-function pulse for one transit (dots) through a 1-D periodic domain, overlaid with the initial conditions (solid line), $\nu = 0$, $v_y = 1$. The Arakawa (top left) solution shows a large phase shift (the peak centers do not match) and significant spurious oscillations. The WENO3 method (top right) is significantly damped near the peak, due to first-order accuracy near extrema. While PPM4 (mid left) is less diffusive, neither peak exactly matches up with its initial condition and both display a slight asymmetry. As to be expected, XPPM (bottom left) only differs from PPM4 in its treatment of the smooth extremum. It and the other extrema-preserving methods, SuHu (mid right) and SuHu-PPM (bottom right), do a very good job at reproducing the correct Gaussian solution, with no negative solutions around the discontinuous pulse.

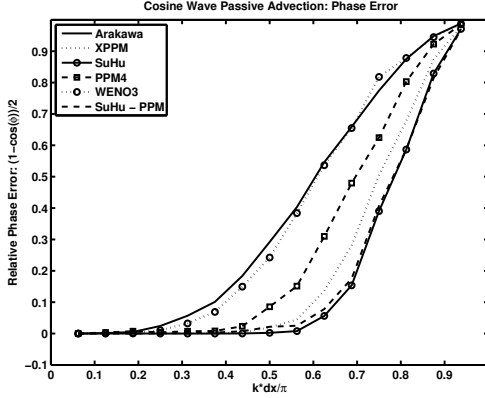


FIG. 4.2. Phase errors for advection of a cosine wave as a function of wave number. The higher-order methods have lower phase errors than WENO3 and Arakawa (which are comparable) across all wave numbers.

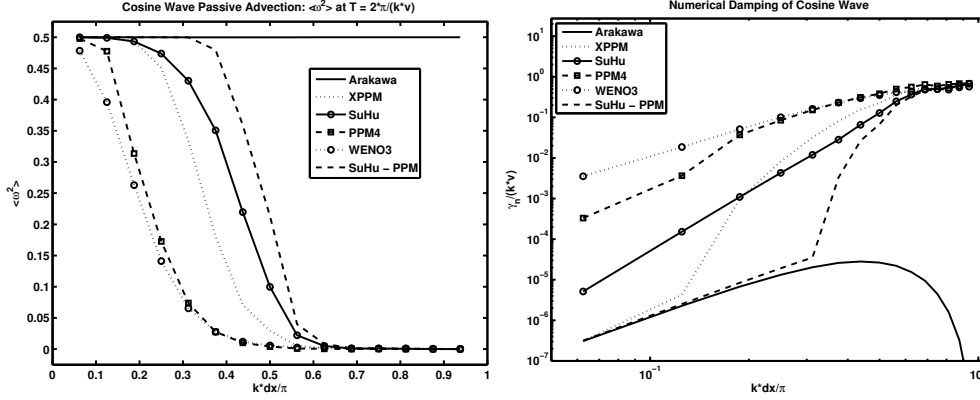


FIG. 4.3. $\langle \omega^2 \rangle$ as a function of wave number for the passive advection of a cosine wave for the different algorithms (left), along with associated numerical damping factors (right) as defined by $\langle \omega^2 \rangle(T) = \frac{1}{2} e^{-2\gamma_n T}$. As expected, the Arakawa method preserves $\langle \omega^2 \rangle = 1/2$ over all wavelengths and displays the minimum level of dissipation in the system, which is due entirely to the RK3 time advancement algorithm. The WENO3 method is the most diffusive, and the PPM4 solution is somewhat less damped. The extremum-preserving methods are closest to the true solution at low wave numbers and the SuHu-PPM hybrid technique approaches the minimum level of dissipation for $k\Delta y \lesssim \pi/3$.

$\gamma_n \propto k^3$, which can be understood as due to the fact that WENO3 reverts to a piecewise constant interpolation near extrema, i.e., simple first order upwind that locally introduces a numerical damping $\sim |v|\Delta x k^2$. Since the fraction of cells at extrema is proportional to $k\Delta x$, the resulting average numerical damping rate $\sim k^3$. Similar arguments apply to the regular PPM method, which is also dominated by clipping errors at extrema.

4.2. Two-Dimensional Vortex Merger. We now turn to the fully two dimensional nonlinear problem as described in Sec. 2. The two Gaussian monopoles begin to swirl around each other and merge together. During the process, sideband rings propagate outward across the mesh. Details on the vortex merging process can be




FIG. 4.4. Low ($N = 128$) and high-resolution ($N = 1024$) vorticity solutions for the vortex merger problem with high Reynolds number $Re = 10^5$ at $t = 100$, for the WENO3, Arakawa, SuHu and XPPM algorithms. The low-resolution solution is most accurate for SuHu and XPPM, as the merging peaks are still resolved. The low-resolution WENO3 solution is so diffusive that the two individual vortices are no longer visible, and the low-resolution Arakawa solution is overrun with numerical oscillations. At high resolution, fine banded features are missing from WENO3 while an altered color-scale shows that spurious oscillations still affect the Arakawa solution.

found in Ref. [24], but it should be noted that the vorticity should never drop below the floor as defined by the background. Any such solutions are solely artifacts of the numerical method.

We show in Fig. 4.4 $Re = 10^5$ vorticity solutions at $t = 100$ for WENO3 (top), Arakawa (upper-middle), SuHu (lower-middle) and XPPM (bottom) at two different resolutions, $N = 128$ (left) and $N = 1024$ (right). Beginning with the low-resolution case, we see that the WENO3 solution is extremely affected by numerical dissipation. The solution has washed out all fine features and has joined the two peaks into a single ellipsoid shape. Furthermore, the orientation of that ellipse does not agree with any of the other solutions. The low-resolution Arakawa solution is similarly plagued with numerical errors. A changed color scale shows that spurious oscillations $\sim 50\%$ of the converged solution obscure the result. However, both the SuHu and XPPM methods more faithfully reproduce the solution, even at low resolutions. Both of the merging peaks are resolved alongside the expanding banded arms. The $N = 1024$ extremum-preserving and Arakawa solutions show crisp details, such as the wake features seen trailing the vortexes, which are washed out by the WENO3 method. This is in agreement with the observations made by Naulin and Nielsen [21]. However, the high-resolution Arakawa solution still has excessive oscillations. This is seen not only in the scale differences of Fig. 4.4 (the maximum and minimum values of the Arakawa solution are larger than those of the other methods), but also in a one-dimensional slice through the domain center, as shown in Fig. 4.5.

With respect to overall levels of dissipation, we present in Fig. 4.6 the total enstrophy Ω at $t = 100$ (left plot) and at $t = 500$ (right plot), for different resolutions, $N = 1024, 512, 256, 128$. Plotted alongside is a curve marked “optimal sub-grid,” which corresponds to a hypothetical sub-grid model that leaves the spectrum unperturbed except for truncation at $k_{\max} = \pi/(\Delta x) = k_{\min}N/2$. The total enstrophy is calculated by summing the enstrophy spectrum Ω_k over all k up to k_{\max} from a well-resolved $N = 1024$ simulation using the fourth order Arakawa method described in Section 4.3. The optimal sub-grid curves are fairly flat, indicating that little of the enstrophy is above these cutoffs.

The solutions for both the WENO3 and Arakawa methods at $t = 100$ agree with those found in Ref. [21]. The numerical dissipation in WENO3 makes the enstrophy significantly lower than that calculated via the Arakawa method, which mathematically preserves the enstrophy. As such, the Arakawa solution only contains dissipation from the explicit viscous term and from the time-stepping algorithm, which is small. The extremum-preserving algorithms are the least dissipative of the finite volume schemes; however, even the PPM4 method presents a significant improvement over WENO3. (A correction to the finite volume methods that accounts for the fact that finite volume methods actually deal with cell-averaged vortices $\bar{\omega}_{i,j}$ and that $[\bar{\omega}_{i,j}]^2$ is slightly larger than $\overline{\omega_{i,j}^2}$ is very small and has no visual impact on this plot.) While Arakawa is closest to the correct total enstrophy when dissipation is not yet important (at early times) and converges most quickly at high resolution, it can have problems at low resolution at later times (as seen in the $t = 500$ plot) when dissipation should occur but the dissipation scale is not adequately resolved. Even at early times when Arakawa is closer to the correct total enstrophy because it has more enstrophy than the finite volume methods, the phases between the high k modes containing that extra enstrophy are not accurate. This leads to the large spurious oscillations seen in the low resolution results in Fig. 4.5.

In general, some level of dissipation is desirable for simulations of turbulent or

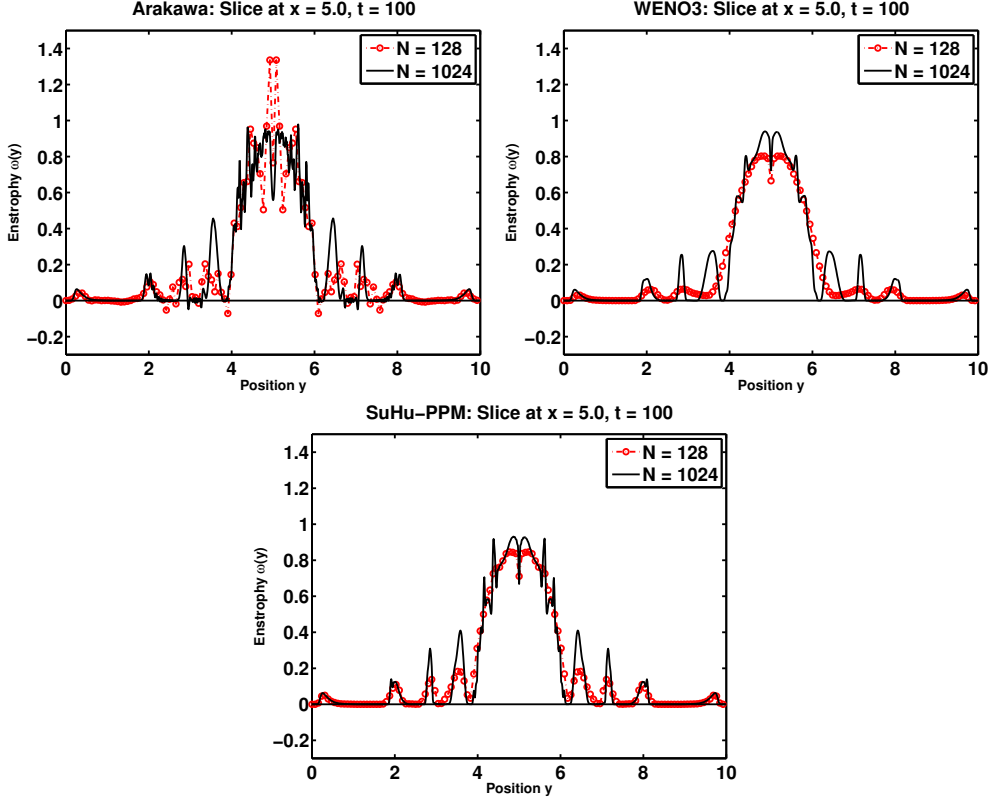


FIG. 4.5. One-dimensional slices of vorticity, $\omega(x = 5.0, y, t = 100)$ for the Arakawa (top left), WENO3 (top right) and SuHu-PPM (bottom) techniques. Both Arakawa resolutions show unphysical oscillations. The SuHu-PPM algorithm, like the other extremum-preserving techniques, more accurately locates some of the vorticity peaks at low resolution that are damped out by dissipation in WENO3, while still avoiding the negative overshoots of the Arakawa method.

other nonlinear cascades, as illustrated in another way by the enstrophy spectra Ω_k shown in Fig. 4.7. This plots $k\Omega(k)$ as a function of normalized wave-number k/k_{\min} , with $k_{\min} = 2\pi/L$, for different resolution simulations at late time, $t = 500$, for the same $\text{Re} = 10^5$ vortex merger problem as before. The left plot is for the Arakawa technique and the right plot is for the SuHu-PPM hybrid technique (XPPM and SuHu give similar results). The one-dimensional enstrophy spectrum Ω_k plotted here is the sum over all modes with $|\vec{k}|$ lying in the bin $k \pm \Delta k/2$, so that the total enstrophy is $\Omega = \sum_k \Omega_k$, where $k = k_{\min}, 2k_{\min}, \dots, \sqrt{2}(N/2)k_{\min}$.

As can be seen in Fig. 4.7, both methods converge to the proper answer at high resolution, and all methods work well for the energy-containing larger scales ($k/k_{\min} < 10$ in this case), but the Arakawa method at lower resolution has an artificial pileup of enstrophy at high k . Because of the conservation properties of Arakawa, the enstrophy is trapped in the system if there is no dissipation (or the dissipation scale is not sufficiently resolved), and instead of cascading down to unresolved smaller scales, the enstrophy is reflected back from the grid-scale, leading to a pileup. This pileup or bottleneck problem is of course well-known in turbulence simulations and is the motivation to add some kind of artificial dissipation or hyper-

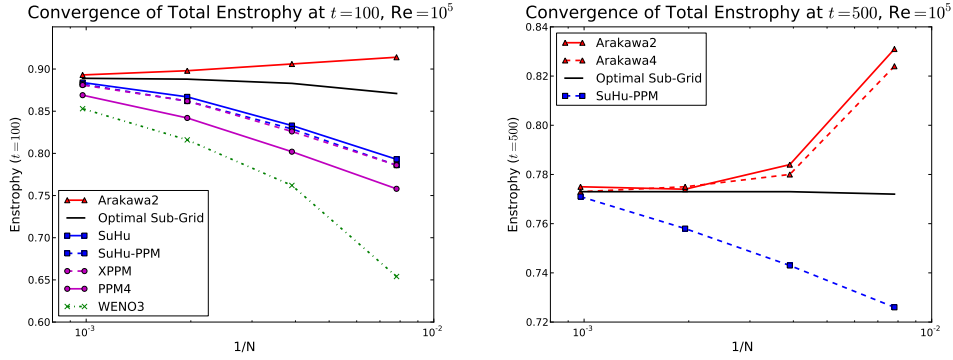


FIG. 4.6. Convergence of the total enstrophy at time $t = 100$ (left) and time $t = 500$ (right) vs. $1/N$, for grid sizes $N = 128, 256, 512, 1024$, for the high Reynolds number case. The $t = 100$ plot shows that WENO3 is the most dissipative, while the extremum-preserving methods are comparable to each other and significantly better than WENO3. Arakawa is closest to the optimal result at $t = 100$, but at later times when more dissipation should have occurred, the Arakawa method does not do as well for coarser resolutions where the dissipation scale is not adequately resolved. At the coarsest resolution the SuHu-PPM solution is closer than either Arakawa solution to that of the optimal sub-grid model.

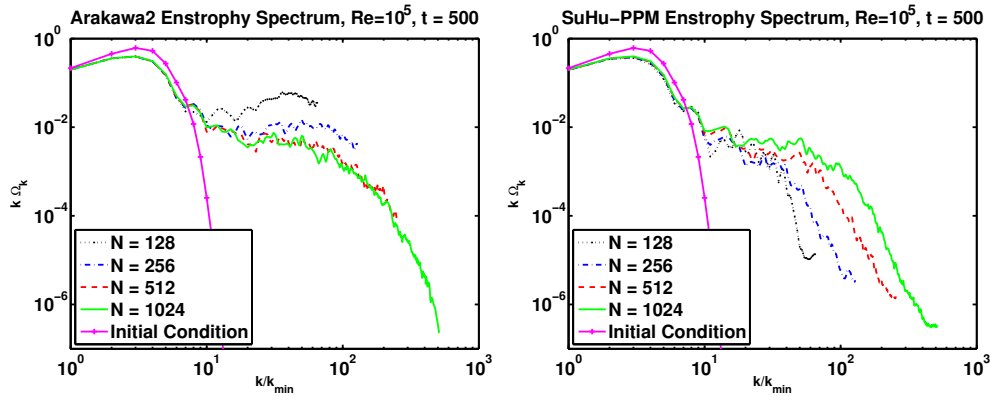


FIG. 4.7. Enstrophy spectra for high Reynolds number at different resolutions, for the Arakawa method (left) and the SuHu-PPM method (right). At lower resolutions the Arakawa method suffers from an artificial pileup of enstrophy at high k because it is not able to cascade to unresolved smaller scales where dissipation should occur. To avoid a pileup of small-scale enstrophy, sufficient damping needs to be included in the equations. The higher-order upwind methods, such as SuHu-PPM (right), are more robust on coarse grids, because they automatically provide some dissipation near the grid scale, which is needed to get the correct inertial range cascade behavior.

viscosity to such systems [3, 25]. The time it takes for nonlinearities to cascade part of the solution to dissipation scales is often very fast (for example, in fully-developed three-dimensional Kolmogorov turbulence, energy reaches the dissipation scale in a time of order the large-scale eddy turnover time, independent of Reynolds number), so dissipation quickly becomes significant in such cases.

In the lowest resolution $N = 128$ Arakawa case, the pileup spectrum for $k/k_{\min} > 20$ is roughly similar to what one might expect for thermalization of enstrophy, where there is no net cascade of enstrophy. (For 2-D thermal equilibrium, the energy spectrum is $E_k \propto k/(\alpha + \beta k^2)$, so the asymptotic result at high k would give $k\Omega_k \propto k^2$ for

positive β coefficient.) This is similar to recent observations of partial thermalization at high k for hyperviscosities $\propto \nabla^{2p}$ with excessively high p exponents [9]. Note that the standard expectations for 2-D turbulence [13, 23] with a net cascade would give $E_k \sim k^{-3}$ in the enstrophy inertial range, and thus $k\Omega_k \sim k^0$, though these results for fully developed turbulence are not necessarily expected to hold for the coherent vortex merger problem under study here.

The SuHu-PPM method (and other similar methods) is more robust at lower resolution, giving the qualitatively correct enstrophy inertial range, because it essentially introduces a kind of sub-grid model that provides some needed dissipation near the grid scale. While the numerical dissipation in these methods is the minimum necessary to ensure monotonicity (in one-dimensional linear advection), and acts qualitatively as a sub-grid model that improves robustness on coarse grids, it is not necessarily the optimal amount of damping for a sub-grid model purely from a turbulence modeling perspective (if small amounts of non-monotonic overshoots can be tolerated). For fluctuations near the grid scale, these methods revert to first order upwind differencing, which corresponds to a damping rate of order $|v|/\Delta x$, while eddy-viscosity concepts [15] indicate that a sub-grid model of the transfer from barely resolved scales to unresolved scales should introduce a damping rate near the grid scale of order the shearing rate, $|\nabla v|$. For a Kolmogorov spectrum $E(k) \sim k^{-5/3}$, where most of the energy is at large scales but most of the shearing is done by small scales, this means that upwind damping at the grid scale is too strong by a factor of approximately $\sqrt{k_{\max}^2 \int dk E(k)} / \sqrt{\int dk k^2 E(k)} \sim (k_{\max}/k_0)^{1/3}$, where k_0 is the wave number of the energy-containing eddies and $k_{\max} \sim 1/(\Delta x)$. While this damping is too strong at the grid scale at high resolution, where k_{\max}/k_0 is large, the numerical dissipation drops off relatively rapidly for $k/k_{\max} < 1/2$ to $1/3$ (see Fig. 4.3 and Fig. 4.7), like a hyperviscosity, so longer wavelength fluctuations are relatively unaffected.

For other types of flow, the numerical dissipation in these methods is not enough because it only damps modes with short scales in the direction of the flow, irrespective of the scales perpendicular to the flow direction. For example, consider passive advection $\partial\omega/\partial t + \nabla \cdot (\vec{v}\omega) = 0$ in a sheared flow $\vec{v} = \hat{x}sy$, with the initial condition $\omega(t=0) = \cos(k_0x)$. The exact solution should be $\omega = \cos(k_0x - k_0sty)$. If $k_0 \ll k_{\max}$ then the numerical dissipation will be negligible, even though the actual solution should cascade to unresolved scales when the y -component of the wave number $k_0st = k_{\max}$. This absence of damping will cause a recurrence error on time scales larger than $k_{\max}/(k_0s)$. For many turbulent systems, the gradients of the advected quantity will be randomly oriented with respect to the flow direction, so there will be sufficient damping near the grid scale. To further improve robustness one could consider adding a sub-grid model to introduce damping related to sheared flows.

It is also possible that a sub-grid model [15] (even a relatively simple Smagorinsky model) added to the Arakawa method could significantly improve its robustness with respect to the enstrophy spectral shape and the pileup problem, though such a sub-grid model by itself would not be enough to eliminate the phase errors that exist even in uniform advection or preserve exact positivity or monotonicity. One could further add limiters only to try to preserve positivity, which would introduce less dissipation than limiters that also try to preserve monotonicity. (These sub-grid models and limiters would of course break the conservation properties that the Arakawa method has by itself.) We leave exploration of these possibilities to future work.

4.3. A Higher-Order Finite Difference Solution. It is possible to construct a higher order Arakawa discretization, in the hopes of alleviating some of the issues seen with the lower-order method presented in Eq. 3.1. The lower order Arakawa discretization is a symmetric approximation to the Poisson bracket on a grid of size h :

$$[\omega, \psi]_h = [\omega, \psi]_{\text{True}} + C_2 h^2 + C_4 h^4 + \dots \quad (4.3)$$

In this case $[\omega, \psi]_{\text{True}}$ is the true value of the Poisson bracket, and the second term represents the error in the discretization, which is proportional to h^2 and scales with a constant factor C_2 . By applying the same algorithm but on a larger grid spacing, we can equally write:

$$[\omega, \psi]_{2h} = [\omega, \psi]_{\text{True}} + 4C_2 h^2 + 16C_4 h^4 + \dots \quad (4.4)$$

Equation 4.4 can be obtained from Eq. 3.1 by the transformation $h \Rightarrow 2h$ and $\{i, j\} \pm 1 \Rightarrow \{i, j\} \pm 2$. These can be combined algebraically to eliminate the second-order error term, thus giving us a new Poisson bracket discretization that has an error proportional to the fourth power of the grid spacing, much like many of the higher order finite volume methods. Explicitly this becomes:

$$\frac{4}{3} [\omega, \psi]_h - \frac{1}{3} [\omega, \psi]_{2h} = [\omega, \psi]_{\text{True}} - 4C_4 h^4 + \dots \quad (4.5)$$

The discretization given by Eq. 4.5 is a fourth order Arakawa finite difference approximation to the Poisson bracket with the same conservation properties as the lower order method detailed above. The grid stencil is now five points in any direction, the same as the higher order finite volume methods, with similar implications for parallelization.

Despite being a more accurate discretization of the Poisson bracket, this higher order method still traps enstrophy in the system and allows non-physical oscillations. Fig. 4.8, a recreation of the tests in Figs. 4.5 and 4.7 with the fourth order Arakawa method, shows similar results to the lower order Arakawa method. Since enstrophy is exactly conserved by the Arakawa methods, it cannot leave the system at the grid scale without explicit damping, and the power spectrum is altered if the dissipation scale is not adequately resolved. The higher order Arakawa solution still displays spurious oscillations, visibly violating positivity, even in the highest resolution simulation.

Higher order versions of the Arakawa Poisson bracket representation share the conservation properties of their lower order counterparts. Since positivity violation, spurious oscillation and turbulent spectrum alteration are all related to these conservation properties, it is not possible to eliminate these problems by merely employing a higher order solution.

5. Conclusions. We have compared a number of discretizations for the Poisson bracket nonlinearity for one-dimensional passive advection and for the vortex merger problem in two-dimensional incompressible Navier-Stokes. The ubiquitous nature of this nonlinearity allows for the application of these techniques to other equations, notably those used to describe edge gyrokinetic and gyrofluid plasmas.

In agreement with previous work [21], we find that the numerical dissipation in the third order WENO method [14, 18, 26] can be significant for some applications. The Arakawa method [1] is designed to be dissipation-free, and so of course does better in terms of conservation properties in cases where there should not be much

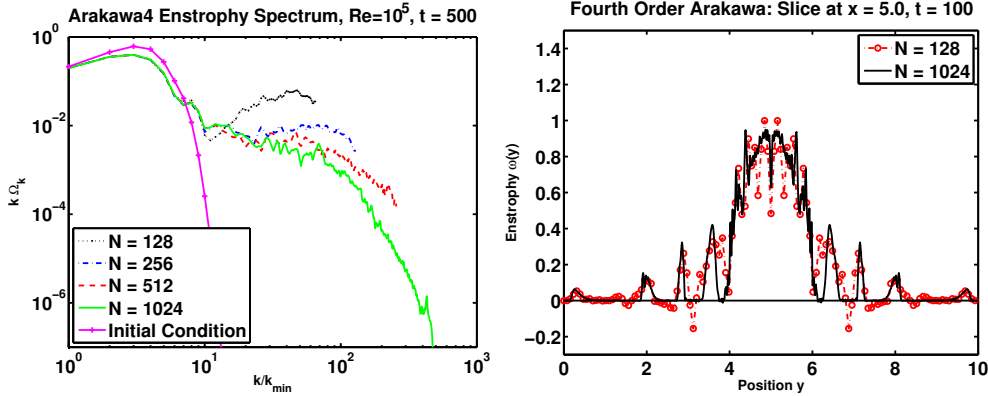


FIG. 4.8. Utilizing a higher order Arakawa method prevents neither spectrum pileup in low resolution cases (left) nor spurious oscillations (right).

dissipation. However, depending on the application at hand, there are other properties of interest beside dissipation rates, such as spectral shape, phase errors, spurious oscillations, and positivity. We have also extended the analysis to include a number of recent finite volume techniques that are both accurate to a higher order and seek to increase the accuracy of solutions at smooth extrema, leading to lower levels of numerical dissipation. These include the high-order WENO implementation of Suresh and Huynh [27], Colella and Sekora’s smooth extrema limiter [6] for the Piecewise Parabolic Method [7], and a hybrid technique that uses PPM4 interpolations and the Suresh and Huynh limiter. The extrema preserving techniques do extremely well in the passive advection problem, eliminating extraneous oscillations and reducing phase errors while keeping dissipation minimal. In the fully nonlinear problem they are less dissipative than the conventional finite volume methods, reduce unphysical oscillations, resolve details fairly well at low resolution and model cascade behavior. All methods of course converge to the right answer when the dissipation scale is well resolved, while these recent finite volume methods are fairly robust even at lower resolutions, because their numerical dissipation acts as an approximate sub-grid model (though not necessarily an optimal one).

Positivity is an important property to preserve for physical quantities such as density, temperature and distribution function in simulations of edge plasmas. While the finite-volume methods studied here do fairly well at preventing artificial overshoots and preserving positivity, because of the modifications to the limiters to preserve the accuracy of smooth extrema they may still produce small negative overshoots. However, it is possible to strictly preserve positivity with a small supplement to the limiters used in these algorithms, similar to the time step constraint they already use to enforce monotonicity. There are various ways the time step constraint might be further relaxed while maintaining positivity. One possibility, mentioned in Ref. [6], is to couple these algorithms with a Flux Corrected Transport algorithm [28].

Another class of techniques that may be worth exploring for these types of problems are Discontinuous Galerkin algorithms [4, 17, 29], as they combine some of the favorable features of finite volume and finite elements, and have relatively low phase errors while keeping a local stencil for efficient implementation.

Acknowledgments. We wish to thank J. Candy, P. Colella, D. Martin and J. Stone for their insight. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 and by the Princeton Plasma Physics Laboratory under Contract DE-AC02-09CH11466. This work was supported by the Center for the Study of Plasma Microturbulence, a DOE Scientific Discovery Through Advanced Computing project.

REFERENCES

- [1] A. ARAKAWA, *Computational design for long-term numerical integration of the equations of fluid motion: Two-dimensional incompressible flow. part i*, J. Comput. Phys., 1 (1966), pp. 119–143.
- [2] A. ARAKAWA, *Introduction to “Computational design for long-term numerical integration of the equations of fluid motion: Two-dimensional incompressible flow. part i”*, J. Comput. Phys., 135 (1997), pp. 101–102.
- [3] V. BORUE AND S. A. ORSZAG, *Forced Three-Dimensional homogeneous turbulence with hyperviscosity*, Europhys. Lett. EPL, 29 (1995), pp. 687–692.
- [4] B. COCKBURN AND C.-W. SHU, *The Runge-Kutta discontinuous Galerkin method for conservation laws v: Multidimensional systems*, J. Comp. Phys., 141 (1998), pp. 199–224.
- [5] P. COLELLA, M. R. DORR, J. A. F. HITTINGER, AND D. F. MARTIN, *High order finite-volume methods in mapped coordinates*, J. Comput. Phys., 230 (2011), pp. 2952–2976.
- [6] P. COLELLA AND M. D. SEKORA, *A limiter for PPM that preserves accuracy at smooth extrema*, J. Comput. Phys., 227 (2008), pp. 7069–7076.
- [7] P. COLELLA AND P. R. WOODWARD, *The piecewise parabolic method (PPM) for gas-dynamical simulations*, J. Comput. Phys., 54 (1984), pp. 174–201.
- [8] M. FRIGO AND S. G. JOHNSON, *The design and implementation of FFTW3*, Proc. IEEE, 93 (2005), pp. 216–231.
- [9] U. FRISCH, S. KURIEN, R. PANDIT, W. PAULS, S. S. RAY, A. WIRTH, AND J.-Z. ZHU, *Hyperviscosity, Galerkin truncation, and bottlenecks in turbulence*, Phys. Rev. Lett., 101 (2008), pp. 144501–144504.
- [10] S. GOTTLIEB, D. I. KETCHESON, AND C.-W. SHU, *High order strong stability preserving time discretizations*, SIAM J. Sci. Comput., 38 (2009), pp. 251–289.
- [11] S. GOTTLIEB, C.-W. SHU, AND E. TADMOR, *Strong stability preserving high-order time discretization methods*, SIAM Rev., 43 (2001), pp. 89–112.
- [12] D. I. KETCHESON, *Highly efficient strong stability-preserving runge-kutta methods with low-storage implementations*, SIAM J. Sci. Comput., 30 (2008), pp. 2113–2136.
- [13] R. H. KRAICHNAN, *Inertial ranges in Two-Dimensional turbulence*, Phys. Fluids, 10 (1967), pp. 1417–1423.
- [14] A. KURGANOV AND D. LEVY, *A Third-Order semidiscrete central scheme for conservation laws and Convection-Diffusion equations*, SIAM J. Sci. Comput., 22 (2000), pp. 1461–1488.
- [15] M. LESIEUR, *Turbulence in Fluids*, Springer, Dordrecht, The Netherlands, 2008.
- [16] R. J. LEVEQUE, *Finite volume methods for hyperbolic problems*, Cambridge University Press, Cambridge, U. K., 2002.
- [17] J.-G. LIU AND C.-W. SHU, *A High-Order discontinuous Galerkin method for 2D incompressible flows*, J. Comput. Phys., 160 (2000), pp. 577–596.
- [18] X.-D. LIU, S. OSHER, AND T. CHAN, *Weighted essentially non-oscillatory schemes*, J. Comput. Phys., 115 (1994), pp. 200–212.
- [19] D. MARTIN AND P. COLELLA, 2008. Private Communication.
- [20] P. MCCORQUODALE AND P. COLELLA, *A high-order finite-volume method for hyperbolic conservation laws on locally-refined grids*, Comm. Appl. Math. Comput. Sci., 6 (2011), pp. 1–25.
- [21] V. NAULIN AND A. H. NIELSEN, *Accuracy of spectral and finite difference schemes in 2D advection problems*, SIAM J. Sci. Comput., 25 (2003), pp. 104–126.
- [22] I. NAVON AND Z. ALPERSON, *Application of fourth-order finite differences to a baroclinic model of the atmosphere*, Meteor. Atmos. Phys., 27 (1978), pp. 1–19.
- [23] K. OHKITANI, *Wave number space dynamics of enstrophy cascade in a forced two-dimensional turbulence*, Phys. Fluids A: Fluid Dyn., 3 (1991), pp. 1598–1611.
- [24] J. J. RASMUSSEN, A. H. NIELSEN, AND V. NAULIN, *Dynamics of vortex interactions in Two-Dimensional flows*, Phys. Scr., T98 (2002), pp. 29–33.
- [25] B. D. SCOTT, *An implicit simulation of drift-wave turbulence in a sheared magnetic field*, J.

- Comput. Phys., 78 (1988), pp. 114–137.
- [26] C.-W. SHU, *Essentially Non-Oscillatory and weighted essentially Non-Oscillatory schemes for hyperbolic conservation laws*, Tech. Report TR-97-65, Inst. Comp. Appl. Sci. Eng. (ICASE), 1997.
 - [27] A. SURESH AND H. T. HUYNH, *Accurate Monotonicity-Preserving schemes with Runge-Kutta time stepping*, J. Comput. Phys., 136 (1997), pp. 83–99.
 - [28] S. T. ZALESAK, *Fully multidimensional flux-corrected transport algorithms for fluids*, J. Comput. Phys., 31 (1979), pp. 335–362.
 - [29] T. ZHOU, Y. LI, AND C.-W. SHU, *Numerical comparison of WENO finite volume and Runge-Kutta discontinuous Galerkin methods*, J. Sci. Comput., 16 (2001), pp. 145–171.