Part I

On the shearing box approximation for the compressible Euler equations

Aaron Wienkers

Trinity College
University of Cambridge

March 2016

Contents

Abstract

The shearing box approximation to the compressible unsteady Euler equations in a shearing, rotating, and stratified flow is investigated. A fundamental review is presented, detailing the development of a 2D finite volume method in a cartesian shearing box representing a circular accretion disc. An overall 2^{nd} order method is constructed using the HLLC approximate Riemann solver, a dimensionally split WAF flux-limited reconstruction scheme, and a symplectic source term integrator to conserve energy. The developed algorithm is validated against a comprehensive test suite modelled after that of Toro (1999) at each 1D, 2D, and source term construction step. The solver is finally employed to explore and further validate against a number of popular problems from accretion disc literature, looking both in the plane of the disc (r-ϕ) as well as at azimuthally invariant solutions in the r-z plane.

1 Introduction

Astrophysical discs are a common phenomenon in the universe — from galaxies and their resident active galactic nuclei, to x-ray binary stars, protoplanetary discs, and planetary rings. These discs are all characterised by rotationally-supported matter in a stable orbit around a central gravitating object. Accretion discs are an interesting subset of these discs because they support a radial inflow of mass and thus mass accretion onto a central object. Accretion discs form when an external source of in-falling matter has excess angular momentum and thus settles into a flattened disc (if the medium is collisional, e.g. gaseous). Since these discs are often dominated by centripetal support against gravity, they are differentially rotating which appears locally as a background shear flow. This shear flow contains a vast source of energy, which in non-rotating systems (such as plane Couette flow) would be fluid-dynamically unstable at these high Reynolds numbers. It may then be surprising at first glance that these shearing discs don't immediately become unstable, dissipate orbital energy, and infall onto the central object. Indeed this does not happen because the system is additionally restricted by angular momentum conservation. Thus in a rotating system, energy liberation must be accompanied with the transport of angular momentum radially outwards while mass accretes inwards. All accretion discs consequently require a mode of angular momentum transport out of the disc to continue supporting mass accretion.

For many years attention has been focused on searching for efficient modes of transporting angular momentum that could explain the deduced rate of mass accretion through observable discs. A number of purely hydrodynamical modes first suspected included simple viscous dissipation; however, the viscous timescale ($\propto R^2/\nu$) is far to slow to be relevant, often longer than the lifetime of the disc itself. Torques between gravitationally unstable regions, and more recently baroclinic instabilities have also been investigated as viable hydrodynamic modes of transport with more potential (Klahr and Bodenheimer, 2003; Marcus et al., 2015). Each of these modes of destabilising accretion discs has varying implications on the influence of stellar birth and evolution as well as the coagulation of dust into planetesimals. A complete understanding of the processes acting in these discs would then help elucidate the history of our solar system when applied to protoplanetary discs (Pringle, 1981).

These phenomena in accretion discs are able to be probed in a number of ways; however, upon exhausting perturbation theory, numerical models become an irreplaceable tool for studying the nonlinear dynamics of these accretion flows. To capture global disc properties and instability modes, a global (entire disc) simulation may be used. Alternatively, simulating just a subset of the disc is computationally cheaper and higher spatial resolutions can be used. However, it is impossible to deduce global disc properties of interest such as the mass accretion rate due to symmetries inherent to the local reference frame. Nonetheless, this paper will focus on numerics in the local approximation.

This paper will begin in §1 by developing the theory required to model accretion discs, which is needed to derive the shearing box approximation in light of its limitations. A short diversion will then be made to explicate the construction of the finite volume code in §2, beginning with the 1D formulation, extending to 2D, and finally adding a general source term solver. The Toro validation test results to this implementation will be discussed in §3, and the full shearing box algorithm developed here will be demonstrated in §4, starting first with an in-plane shearing sheet and then exploring known solutions in the r-ϕ plane shearing approximation.

1.1 Accretion Disc Dynamics

A global hydrodynamical model of an accretion disc may be realised by simply adding a gravitation term, $\nabla\Phi_g$, to the Navier-Stokes momentum equation, where the non-self-gravitational potential at radius r, and height off the midplane, z, is $\Phi_g = -GM/\sqrt{r^2+z^2}$. This assumption that self-gravitation between fluid elements is negligible is valid for low-mass discs and quantified by the Toomre Q stability criterion for a Keplerian disc,

$$Q = \frac{c_s\Omega}{\pi G\Sigma} > 1. \tag{1}$$

Here Σ is the local surface density of the thin disc, Ω is the angular velocity (orbital frequency), and $c_s = \sqrt{\gamma p/\rho}$ is the adiabatic sound speed. The addition of this radial gravitation term motivates a number of length and time scalings pertinent to astrophysical discs. Apart from the disc radius, R, a second length scale is motivated from a hydrostatic balance with the vertical component of gravity. This isothermal scale height is

$$H = \frac{c_s}{\Omega}. \tag{2}$$

Thus to a first approximation, the disc can be seen to flare radially, proportional to $r^{3/2}$. Relating these two length scales gives a non-dimensional disc aspect ratio, H/R. Three additional pertinent timescales also come out of the Navier-Stokes equations with a central gravitational potential: the dynamical, viscous, and thermal times. The relative magnitude of each will motivate and validate a number of important assumptions. The dynamical time, $t_d \sim \Omega^{-1}$, is the timescale on which important dynamic processes occur on, such as instability and turbulent transition. A viscous timescale, $t_v \sim R^2/\nu \sim \alpha^{-1} (R/H)^2 t_d$, sets the time for a fluid parcel to move radially through the disc solely due to viscous dissipation. Here, $\alpha < 1$ is the α-viscosity parameter of Shakura and Sunyaev (1973). Finally, the thermal timescale is locally set by the ratio of the thermal heat content to the viscous heating rate and is $t_t \sim \alpha^{-1} t_d$. Thus for thin $(H/R \ll 1)$ discs, $t_d < t_t \ll t_v$. It is argued based on the relative time-scales of this system that any long-term evolution can be ignored because the dynamical processes of interest are relatively quite fast. This means the evolution of the central mass can be ignored as well as viscous heating processes that may perturb the system from an initial equilibrium state. Viscosity may then be ignored altogether, and focus turned to the Euler system.

This simplified system of equations now supports two natural oscillation frequencies to small amplitude perturbations: the horizontal and vertical epicyclic frequencies, respectively Ω_r and Ω_z. For the Keplerian discs of interest in this paper, $\Omega_r = \Omega_z = \Omega$, and thus these epicycles (to first order) correspond to slightly eccentric and inclined orbits, respectively.

It should be acknowledged that some additional physics will not be included for the sake of simplicity, but which are known to be important in particular disc dynamics (Papaloizou and Lin, 1995, and references therein). Irradiation heating from the central object and radiative cooling dictate the viability of certain instabilities, in particular gravitational instability. Viscous heating also wages in this thermodynamic balance in regions of intense shear near the central object, warranting a general viscosity prescription. Most importantly in ionised discs, magnetic fields are necessary for realising the magneto-rotational instability known to locally destabilise ionised discs (Armitage, 2011).

1.2 Shearing Box Approximation

For studying *local* small-scale phenomena, it is numerically and often analytically beneficial to enter a local reference frame co-rotating and shearing with the disc. In local cartesian coordinates, $(x, y, z) = (r - R_0, R_0 (\phi - \phi_0 - \Omega_0 t), z)$, centred about a fiducial radius, R_0, and co-rotating with the orbital frequency, $\Omega_0 = \Omega(R_0)$, additional symmetries are introduced which can greatly simplify analysis (see Figure 1). Thus the equation of motion governing $\mathbf{u} = (u, v, w)$, in the rotating frame is

$$\rho \left(\frac{D\mathbf{u}}{Dt} + 2\Omega_0 \hat{\mathbf{z}} \times \mathbf{u} \right) = -\nabla p - \rho \nabla \Phi_t, \tag{3}$$

where $\Phi_t = \Phi_g + \Phi_c$ is the effective tidal potential — the sum of the centrifugal potential arising from the non-inertial reference frame and the gravitational potential. In this paper, the accretion disc will be assumed circular and in equilibrium, and thus will have a Keplerian rotation profile, $\Omega_k = \sqrt{GM/r^3}$ independent of ϕ, although this assumption will be challenged in Part II of this series.

The shearing box approximation is derived by expanding (3) in $\mathbf{x}/R_0 \ll 1$. The tidal potential in the local coordinates is then

$$\Phi_{\text{eff}} = -\frac{GM}{\sqrt{r^2 + z^2}} - \frac{1}{2}\Omega_0^2 r^2 = -q\Omega_0^2 x^2 + \frac{1}{2}\Omega_0^2 z^2 + \mathcal{O}\left(x^3 + z^3\right), \tag{4}$$

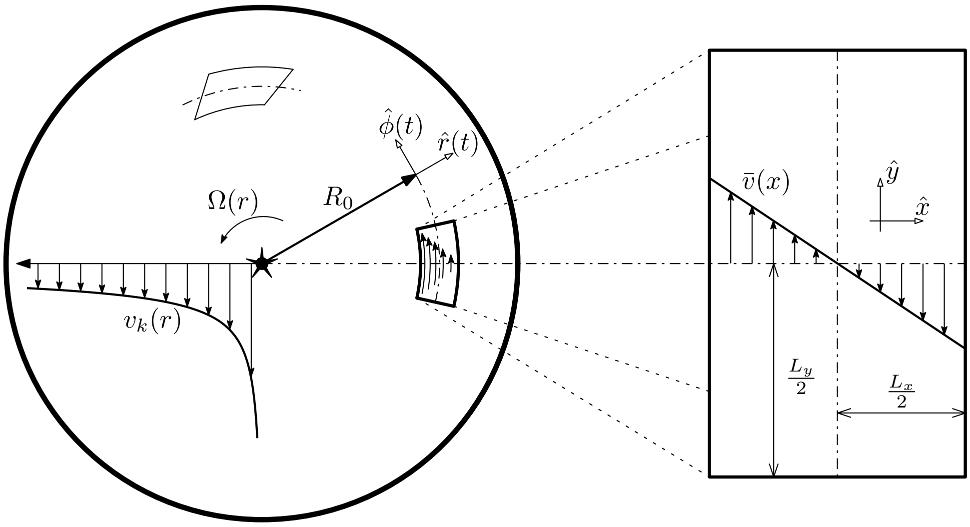

Figure 1: The global disc (left) and extracted local cartesian patch (right). The control volume is shown sheared and advected at a later time while the shearing box remains unchanged.

where the shear rate, $q \equiv -r\mathrm{d}\Omega/\mathrm{d}r$, is 3/2 for a Keplerian disc. This expansion can be seen as a projection removing the curvature of the solution by eliminating dependence on y. Using this expansion with (3), the local shearing box approximation in primitive variables can be written as

$$\frac{D\rho}{Dt} = -\rho\nabla\cdot\mathbf{u} \tag{5a}$$

$$\frac{D\mathbf{u}}{Dt} = -\frac{1}{\rho}\nabla p - 2\Omega_0\hat{\mathbf{z}}\times\mathbf{u} + 2q\Omega_0^2 x\hat{\mathbf{x}} - \Omega_0^2 z\hat{\mathbf{z}} \tag{5b}$$

$$\frac{Dp}{Dt} = -\gamma p\nabla\cdot\mathbf{u}, \tag{5c}$$

where $\frac{D}{Dt}(\cdot)$ is the total (Lagrangian) derivative. An ideal gas equation of state, $p = (\gamma-1)\,e$, was assumed to write equation (5c) in terms of pressure, providing closure. The ratio of specific heats, $\gamma = 7/5$, will be used for verification tests, though 5/3, is more appropriate for unionised accretion discs containing molecular hydrogen gas too cool to be rotationally excited. The shearing box admits an equilibrium velocity solution, $\bar{\mathbf{u}} = -q\Omega_0 x\hat{\mathbf{y}}$, which is exactly the local uniform Keplerian orbital motion by construction. Additionally, for an initially isothermal (i.e. $\gamma = 1$) disc, the equilibrium density and pressure profiles are proportional to $\exp\left(-z^2/2H^2\right)$.

Simple periodic boundary conditions can be applied in the $\hat{\mathbf{y}}$ direction. Writing the vector of primitive variables as $\mathbf{W} = (\rho, \mathbf{u}, p)^T$, this boundary condition is summarised as

$$\mathbf{W}(x, y, z, t) = \mathbf{W}(x, y \pm L_y, z, t) \tag{6}$$

for a domain of size (L_x, L_y). This boundary condition conflicts with the background Keplerian shear flow solution when applied in the $\hat{\mathbf{x}}$ direction. The $x = -L_x/2$ boundary is moving faster than at $x = +L_x/2$ by $v_{\text{boost}} = q\Omega_0 L_x$, and so the velocity must be Galilean boosted and shifted in the $\hat{\mathbf{y}}$ direction. This "shift-boost" periodic boundary emulates a differentially shifting tiled super-domain (see Figure 2), and is written explicitly as

$$\rho(x, y, z, t) = \rho(x \pm L_x, y \mp v_{\text{boost}}t, z, t) \tag{7a}$$

$$\mathbf{u}(x, y, z, t) = \mathbf{u}(x \pm L_x, y \mp v_{\text{boost}}t, z, t) \mp v_{\text{boost}}\hat{\mathbf{y}} \tag{7b}$$

$$p(x, y, z, t) = p(x \pm L_x, y \mp v_{\text{boost}}t, z, t). \tag{7c}$$

Since $t = 0$, the boundaries have sheared by $v_{\text{boost}}t$ relative to each other. Numerically, this necessitates periodically regridding the domain, which is realised by taking the modulus of the

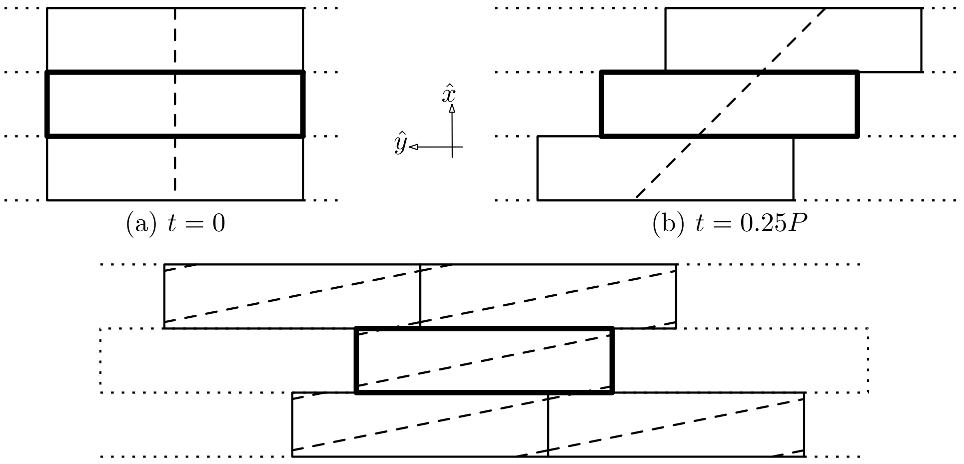

(a) $t = 0$

(b) $t = 0.25P$

(c) $t = 1.25P$

Figure 2: Tiling of the sheared super-domain motivating the shift-boost boundary condition. A tracer line (dotted) is laid at $t = 0$. The relative positions of the computational domain are shown at (b) $t = 0.25$ and (c) $t = 1.25$ regrid periods.

$v_{\text{boost}}t$ term with L_y in (7). Finally the boundaries in the vertical ($\hat{\mathbf{z}}$) direction must be chosen to approximate a vertically unbounded disc. Considering the vertical equilibrium profile, reflective boundary conditions are a good choice when enforced at a height $L_z/2 \gg H$ so as to minimise artificially constraining the fluid near the midplane.

It must be noted that the shift-boost boundary conditions physically do not conserve energy. This becomes apparent when expanding the boundary condition for the total energy density, $E = p/(\gamma - 1) + {}^1\!/\!_2\rho|\mathbf{u}|^2$. At either boundary, the Galilean boost gives each $E \to E \pm \rho v v_{\text{boost}} + \frac{1}{2}\rho v_{\text{boost}}^2$, where the second term does not cancel between the two boundaries. Intuitively, this is because the boundary conditions were designed to maintain the background shear and work must be done on the fluid to do so. This should not be seen as a limitation of the shearing sheet model, but rather elucidates the energy transfer from the global disc (Hawley et al., 1995). Although the total energy is not conserved, a new conserved quantity — the epicyclic energy — is found by manipulating the momentum equations (5b):

$$E_{\text{epi}} = \rho^2 \left(\delta u^2 + \frac{2}{2-q} \delta v^2 \right). \tag{8}$$

This epicyclic energy is a weighted perturbation energy for velocity fluctuations $\delta v = v_y + q\Omega_0 x$ and $\delta u = u$ from the equilibrium shear (Stone and Gardiner, 2010).

A number of symmetries arise as consequence for making the shearing sheet approximation with these boundary conditions. The system (5) is evidently (globally) horizontally invariant — the dynamics are independent of the (r, ϕ) coordinates considered to be the centre of the domain, since these were neglected at leading order in the expansions. This also implies the direction towards the central gravitating object is ambiguous. Therefore, this local model is limited from determining global disc properties such as angular momentum transport or the accretion rate (Regev and Umurhan, 2008). Finally, as apparent from the equilibrium profiles, the vertical symmetry about the plane $z = 0$ is retained from the global, inertial equations.

2 Numerical Method

The finite volume method is employed to numerically solve the physical model described by the system of equations (5) cast in conservative form:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot [\rho \mathbf{u}] = 0 \tag{9a}$$

$$\frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot [\rho \mathbf{u}\mathbf{u} + p\mathrm{I}] = \rho \Omega_0^2 (2qx\hat{\mathbf{x}} - z\hat{\mathbf{z}}) - 2\Omega_0 \hat{\mathbf{z}} \times \rho \mathbf{u} \tag{9b}$$

$$\frac{\partial E}{\partial t} + \nabla \cdot [(E+p)\mathbf{u}] = \Omega_0^2 \rho \mathbf{u} \cdot (2qx\hat{\mathbf{x}} - z\hat{\mathbf{z}}). \tag{9c}$$

The governing equations must be cast into this conservation form to converge to the weak solution upon discretisation. The classic example of this swindle is demonstrated by solution of Burger's equation in differential form, resulting in blatantly incorrect shock speeds. Reading off the variables in the partial time derivative defines the vector of conserved variables, $\mathbf{U} = (\rho, \rho\mathbf{u}, E)^T$. Then each term on the right of (9) constitutes the sourcing, $\mathbf{S}(\mathbf{U}, t)$. The initial part of this paper will focus on constructing and validating the solution of the homogeneous Euler set, where $\mathbf{S} = \mathbf{0}$.

2.1 Generalised Finite Volume Formulation

An ideal numerical scheme is designed to emulate the properties and behaviour of the physics it is being used to model. The foremost of these properties is being conservative, and so the following methods will by construction be conservative.

The vector of physically conserved continuous variables, $\mathbf{U}(x, y, z, t)$, is discretised by interpreting the value at the cell-centred point, $\mathbf{U}_{ijk}(t) = \mathbf{U}(i\Delta x, j\Delta y, k\Delta z, t)$, to be averaged over the the whole of the ijk^{th} finite control volume (or area in 2D). This integral average over cell volume V_{ijk} is

$$\begin{aligned} \mathbf{U}_{ijk}(t) &= \frac{1}{V_{ijk}} \int_{z_{k-\frac{1}{2}}}^{z_{k+\frac{1}{2}}} \int_{y_{j-\frac{1}{2}}}^{y_{j+\frac{1}{2}}} \int_{x_{i-\frac{1}{2}}}^{x_{i+\frac{1}{2}}} \mathbf{U}(x, y, z, t)\, \mathrm{d}x\, \mathrm{d}y\, \mathrm{d}z \\ &= \frac{1}{V_{ijk}} \iiint_{V_{ijk}} \mathbf{U}(\mathbf{x}, t)\, \mathrm{d}\mathbf{V}. \end{aligned} \tag{10}$$

The finite volume method is then founded on locally enforcing a general conservation law of the form

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbf{F}(\mathbf{U}) = \mathbf{S}(\mathbf{U}, t) \tag{11}$$

in each discretised control volume, V_{ijk}. Here $\mathbf{F}(\mathbf{U})$ is the *equation* flux function and $\mathbf{S}(\mathbf{U}, t)$ is the sourcing term, both mapping $\mathbb{R}^5 \to \mathbb{R}^5$ in the case of the 3D Euler system (9).

Integrating (11) over the finite volume, V_{ijk}, and using Liebniz's rule with the Gauss-Ostrogradsky theorem, gives an ODE describing the evolution of the finite volume as

$$\frac{\mathrm{d}}{\mathrm{d}t} \mathbf{U}_{ijk}(t) + \oiint_{\partial V_{ijk}} \mathbf{F}(\mathbf{U}) \cdot \mathrm{d}\mathbf{A} = \left(\mathbf{S}\left(\mathbf{U}, t\right)\right)_{ijk}. \tag{12}$$

The discretised finite volume source term is written here in the continuous form $(\mathbf{S}(\mathbf{U}, t))_{ijk} = \iiint_{V_{ijk}} \mathbf{S}(\mathbf{U}, t) \mathrm{d}\mathbf{V}$ to emphasise that the source term is volume-averaged. Often times the conservative set of equations define a flux function which is difficult to solve all at once. Thus, dimensional splitting algorithms help make the 3D update formula (12) computationally tractable.

2.2 Operator Splitting

Operator splitting is often employed to simplify the full 3D non-homogeneous conservation law into a series of 1D fractional step operators without compromising solver accuracy. Thus only a generalised 1D finite volume scheme needs to be designed.

Using the Method of Lines the general conservation equation (11) can be spatially semi-discretised as

$$\frac{\mathrm{d}\mathbf{U}}{\mathrm{d}t} = A_x\mathbf{U} + A_y\mathbf{U} + A_z\mathbf{U} + A_s\mathbf{U} \tag{13}$$

where A_x, A_y, and A_z are the discrete linear operators for some approximation to the partial derivatives. A_s is the source function in the special case that it can be written as a linear operator on $\mathbf{U}$, which is now a vector of vectors containing all spatial data reordered into a single column. It is trivial to solve this ODE giving $\mathbf{U}(t)$ in each discrete control volume:

$$\mathbf{U} = \mathbf{U}_0 \mathrm{e}^{t(A_x+A_y+A_z+A_s)} \tag{14}$$

Matrix exponentials are not in general commutative and so this exponent cannot be written as the product of the individual exponents. However, to make an overall 2^{nd} order scheme, (13) only needs to be solved at 2^{nd} order itself!

Now considering just a 2D method with a source term, there are two 2^{nd} order methods to split the exponent in equation (14):

$$\mathrm{e}^{t(A_x+A_y+A_s)} = \mathrm{e}^{\frac{1}{2}tA_x}\mathrm{e}^{\frac{1}{2}tA_y}\mathrm{e}^{tA_s}\mathrm{e}^{\frac{1}{2}tA_y}\mathrm{e}^{\frac{1}{2}tA_x} + \mathcal{O}\left(t^3\right) \tag{15}$$

$$= \frac{1}{4}\left[\mathrm{e}^{tA_x}\mathrm{e}^{tA_y}\mathrm{e}^{tA_s} + \mathrm{e}^{tA_x}\mathrm{e}^{tA_s}\mathrm{e}^{tA_y} + \mathrm{e}^{tA_y}\mathrm{e}^{tA_s}\mathrm{e}^{tA_x} + \mathrm{e}^{tA_s}\mathrm{e}^{tA_y}\mathrm{e}^{tA_x}\right] + \mathcal{O}\left(t^3\right) \tag{16}$$

The first of these operator splitting methods is Strang splitting, which is implemented as a half time step in x then y, a full time step to solve the source with A_s, and finally the second half time step in y then x. The second method, (16), is the parallel splitting scheme, which gains 2^{nd} order accuracy by averaging each of the 4 cyclic permutations of the linear operators. However, it will be shown that taking two half steps is not exactly equivalent to taking one full step, solely because of compounding truncation error. This compounding error often emerges as an increase in numerical diffusion with increasing number of time steps. Strang splitting decreases the effective time step by half and consequently increases the numerical diffusivity. Therefore, although parallel splitting is over twice as expensive, it is favoured over Strang splitting because it does not unnecessarily decrease the step size.

Now the 3D ODE given by equation (12) has been simplified to four 1D ODEs. Defining the solution operator in the $\hat{\mathbf{x}}$ direction, $\mathcal{L}_x^t$, which approximates (to within the constant of integration) the solution to

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{U}}{\partial x} = 0 \tag{17}$$

at time t for some initial condition $\mathbf{U}_0$, it is clear that, $\mathcal{L}_x^{\Delta t} = \exp(\Delta t A_x)$, by using the approximation to the discrete partial derivative operator, A_x. A first order solution to the 3D system, (11), is then given by $\mathbf{U}^{n+1} = \mathcal{L}_x^{\Delta t}\mathcal{L}_y^{\Delta t}\mathcal{L}_z^{\Delta t}\mathcal{L}_s^{\Delta t}\mathbf{U}^n$. The only task is to find each $\mathcal{L}^{\Delta t}$ to step $\mathbf{U}^n$ forward by Δt in a single dimension at a time.

Focussing on just the $\hat{\mathbf{x}}$ direction advection solution, $\mathcal{L}_x^{\Delta t}$, in 1D the closed surface integral term of equation (12) becomes a difference of fluxes evaluated at the cell edges, $x_{i\pm\frac{1}{2}}$. Integrating from time $n\Delta t$ to $(n+1)\Delta t$, gives the 1D conservative update formula,

$$\mathbf{U}_i^{n+1} = \mathbf{U}_i^n + \frac{\Delta t}{\Delta x}\left(\mathbf{f}_{i-\frac{1}{2}} - \mathbf{f}_{i+\frac{1}{2}}\right). \tag{18}$$

This allows the vector of conservative variables to be marched forward in time, given the time-averaged equation flux at each cell edge. This new flux is termed the *scheme*† flux, $\mathbf{f}_{i\pm\frac{1}{2}}$, which is the flux through the cell edges normal to the solution direction ($\hat{\mathbf{x}}$). It is clear now that this scheme is conservative, because the flux leaving one cell is identically the flux entering the adjacent cell.

Whereas the grid size Δx is specified by the user at run-time, the time step size in (18) must be adaptively chosen to help ensure numerical stability. The time step and cell size define a computational grid speed, $\Delta x/\Delta t$, which is the fastest speed that information such as waves,

---

[†] Not to be confused with the equation flux function, $\mathbf{F}(\mathbf{U})$.

characteristics, or flow velocities can be resolved to travel across the computational domain. The grid velocity must be greater than the speed of the physics of interest to have stability. It is convenient to define the Courant-Friedrich-Lewy (CFL) number,

$$\mathrm{CFL} = \frac{\Delta t}{\Delta x} u_{\max} \tag{19}$$

where $u_{\max} = \max_{\forall i}(|u| + a)$ over all cells, and a is the sound speed in a numerical context. The CFL number is then the fraction of a cell width that is acceptable for a disturbance to propagate in a single time step. $\mathrm{CFL} \leq 1$ is a necessary condition for any finite volume method to be stable and convergent, though a particular scheme may further restrict the stable CFL regime. In this paper, $\mathrm{CFL} = 0.9$ will be used, except during the first 5 time steps when $u_{\max}$ is not known with enough certainty, and so CFL will be reduced to 0.2.

The final 2nd order splitting scheme of choice for the 2D system with sourcing is

$$\mathbf{U}^{n+1} = \mathcal{L}_s^{\frac{1}{2}\Delta t} \frac{1}{2} \left(\mathcal{L}_x^{\Delta t} \mathcal{L}_y^{\Delta t} + \mathcal{L}_y^{\Delta t} \mathcal{L}_x^{\Delta t} \right) \mathcal{L}_s^{\frac{1}{2}\Delta t} \mathbf{U}^n. \tag{20}$$

The fractional steps are as follows: The source term is first stepped forward $\frac{1}{2}\Delta t$ using $\mathbf{U}^n$ as the initial condition. Then the x and y terms are stepped forward parallel in time, using the output of $\mathcal{L}_s^{1/2\Delta t}\mathbf{U}^n$ as the initial condition. Finally, this output is used again as the input for the final half-step of the source term to give $\mathbf{U}^{n+1}$. This scheme was chosen to minimise numerical diffusivity and artificial anisotropies in the advection terms while avoiding the 2x penalty incurred by using parallel splitting on the source term.

2.3 Advection Solver

With the full 3D advection step simplified into three uncoupled 1D steps, all that is left to do now is evaluate the scheme fluxes in the 1D conservative update (18) defining each solution operator, $\mathcal{L}$. Two upwind techniques for constructing 2nd order methods of finding these scheme fluxes will be investigated now, solving for the dimensionally split operators, $\mathcal{L}_x$ and $\mathcal{L}_y$. Centred schemes will not be explored here due to their excessive numerical diffusivity.

2.3.1 Riemann Problem Solution

The most common upwind methods for solving the compressible Euler equations are based on solving a Riemann problem at each intercell boundary arising from discretisation (see Figure 3a), to obtain the intercell scheme flux, $\mathbf{f}_{i\pm\frac{1}{2}}$. These are called Godunov Methods. The resulting Riemann problem is in a family of general Riemann problems which describe any initial value problem to a conservation equation with a piecewise initial condition consisting of a single discontinuity. Thus at each cell edge, a Riemann problem arises due to the domain discretisation (10) described by the initial condition,

$$\mathbf{U}(x, t=0) = \begin{cases} \mathbf{U}_L & x \leq 0 \\ \mathbf{U}_R & x > 0 \end{cases} \tag{21}$$

which evolves according to the conservative advection terms of the governing equations (left hand side of (9)). Given a solution of this Riemann problem, the original (1st order) Godunov upwind method evaluates the scheme flux using the state vector of the Riemann problem solution, $\check{\mathbf{U}}_{i+\frac{1}{2}}\left(\frac{x}{t}\right)$, at the cell edge (where $x = 0$ due to a translation by $x_{i+\frac{1}{2}}$), giving

$$\mathbf{f}_{i+\frac{1}{2}} = \mathbf{F}\left(\check{\mathbf{U}}_{i+\frac{1}{2}}(0)\right). \tag{22}$$

This is the Godunov intercell numerical flux. Thus the scheme flux to step equation (18) forward in time is evaluated using the particular Riemann solution state vector at the cell edge.

The Riemann problem solution for the Euler equations in general has no analytic solution, but can be solved either iteratively or by approximate means. The solution consists of left- and

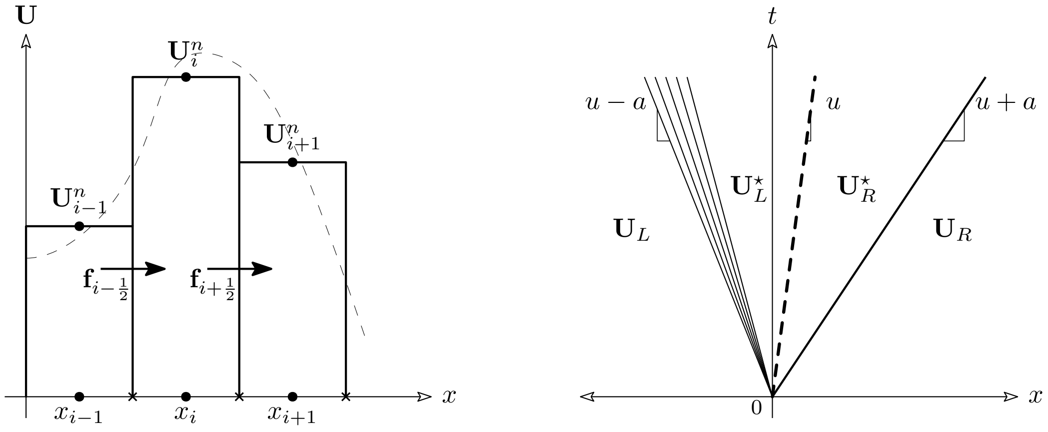

(a) Data discretisation and Riemann problem

(b) Riemann solution with characteristics

Figure 3: The problem discretisation onto the cell-centred points, •, shown in (a) gives rise to a Riemann problem at each cell-edged point, ×; and (b) the general Riemann solution centred about $x = 0$, with a left-moving rarefaction, central contact discontinuity, and right-moving shock.

right-travelling shock and rarefaction waves in different combinations, but always separated by a contact discontinuity. Then three characteristic waves (eigenvalues) in the solution, separate the four states shown in Figure 3b.

To update cell i, two Riemann problems must be solved — one each at $i \pm \frac{1}{2}$. First looking at the problem at edge $x_{i+\frac{1}{2}}$, $\mathbf{U}_L = \mathbf{U}_i$ and $\mathbf{U}_R = \mathbf{U}_{i+1}$ in Figure 3b. It is apparent that there is a similarity solution with similarity variable, x/t. This means that a complete solution can be found easily provided each of the characteristics as well as the starred states between them, because each characteristic separates constant states, except for the expansion fan (rarefaction) which requires special treatment.

An exact Riemann solution is explicated in Toro (1999); however, this is not practical because it often requires many iterations to find a sufficiently accurate solution. This becomes very computationally expensive when conducted in the lowest nested loop of the finite volume method. Approximate Riemann solutions must then be resorted to.

Adaptive Noniterative Riemann Solver (ANRS)

A cheap and robust approximate Riemann solver is Toro's Adaptive Noniterative Riemann Solver (1999) which will only gain utility in the HLLC approximate solver below to make a more accurate estimate of the sound speed. Only the pressure states need to be known to make an estimate of the characteristics of the Riemann problem and so just the pressure solution will be detailed.

ANRS improves upon three independent approximate solvers by discriminating between smooth parts of the solution (quantified by how nearly $p_L \sim p_R$) and parts where a more careful approximation must be used. A cheap first approximation to the starred states is determined with the Primitive Variable Riemann Solver (PVRS) of Toro (1991). The PVRS gives the starred pressure state (between the two characteristics) as

$$p^\star_{\text{PVRS}} = \frac{1}{2}(\rho_L + \rho_R) - \frac{1}{2}(u_R - u_L)\bar{\rho}\bar{a} \tag{23}$$

where $\bar{\rho}$ and $\bar{a}$ are just the average of the left and right density and sound speeds respectively. This starred state pressure is sufficient to determine how conditioned the local solution is. Defining $p_{\max}$ and $p_{\min}$ as the minimum and maximum pressures between the left and right states, and further defining $Q \equiv p_{\max}/p_{\min}$ gives a metric for how well-behaved the solution is.

For a sufficiently well-behaved solution, designated by $Q < Q_{\text{tol}} = 2$, the PVRS solution is a good approximation if $p_{\min} < p^\star_{\text{PVRS}} < p_{\max}$. In the case that Q is quite large, then a strong shock or rarefaction exists and care must be taken to find an accurate solution. Rarefactions are favoured when $p^\star_{\text{PVRS}} < p_{\min}$, calling utility of the Two Rarefaction Riemann Solver (TRRS) to determine the revised starred states. Assuming both non-linear waves are rarefactions, then

$$p^\star_{\text{TRRS}} = \left(\frac{a_L + a_R - \frac{\gamma-1}{2}(u_R - u_L)}{a_L/p_L^z + a_R/p_R^z} \right)^{1/z} \tag{24}$$

where $z = (\gamma - 1)/2\gamma$. Otherwise when $p^\star_{\text{PVRS}} > p_{\max}$, a shock exists and so a better approximation is given by the Two Shock Riemann Solver (TSRS). Assuming both waves are shock waves gives the updated starred pressure

$$p^\star_{\text{TSRS}} = \frac{g_L p_L + g_R p_R + u_L - u_R}{g_L + g_R} \tag{25}$$

where

$$g_K = \sqrt{2}\left(\rho_K((\gamma+1)p^\star_{\text{PVRS}} + (\gamma-1)p_K)\right)^{-1/2}$$

for $K = L, R$. Although the TSRS gives an implicit solution, (25) is made explicit by using the approximation to $p \sim p^\star_{\text{PVRS}}$ with acceptable accuracy reported in Toro (1991).

The above algorithm thus gives a better approximation to the solution of the Riemann problem than any of the three constituent solvers alone, since each excels at giving a cheap, accurate approximation in very specific regimes.

Harten - Lax - van Leer Contact (HLLC) Riemann Solver

The HLLC approximate Riemann solver is one of the most widely used approximate solvers, acting as a base method for many hybrid Riemann solvers. HLLC itself is a modified version of the original HLL (Harten-Lax-van Leer) solver (Harten et al., 1983). It was first used by Toro et al. (1994) to consider the centred contact (C) waves, and aptly modified to HLL(C).

A critical step in any Riemann problem solution is the accurate determination of the characteristics. The left and right wave speeds can be estimated using a variety of methods (Roe, Davies, etc.), but for consistency with Toro (1999), the ANRS will be used to estimate the pressure in the starred region. Accurately determining the wave speeds is particularly important for finite volume schemes which use each of the solution states (such as the Weighted Average Flux of §2.3.3), because errors in the wave speed estimates correspond to errors in the weights of the solutions states.

This starred pressure estimate is used to approximate the sound speed, $a^\star$, and finally the wave speeds (S_L, S_R, and S_C) of the left, right, and contact discontinuities respectively. For a compression wave, the characteristic is the rarefaction front, which travels at the sound speed, and so

$$S_L = u_L - a_L \tag{26a}$$

$$S_R = u_R - a_R; \tag{26b}$$

however, if the characteristic is a shock wave (e.g. $p^\star > p_{L/R}$), then an approximation of the true shock speed is

$$S_L = u_L - a_L \left(1 + \frac{\gamma+1}{2\gamma}\left(p^\star/p_L - 1\right)\right)^{1/2} \tag{27a}$$

$$S_R = u_R + a_R \left(1 + \frac{\gamma+1}{2\gamma}\left(p^\star/p_R - 1\right)\right)^{1/2}. \tag{27b}$$

This is an approximation only because $p^\star$ is approximate, which further motivates using a more accurate predictor Riemann solver such as ANRS. The speed of the contact surface is similarly estimated as

$$S_C = \frac{p_R - p_L + \rho_L u_L(S_L - u_L) - \rho_R u_R(S_R - u_R)}{\rho_L(S_L - u_L) - \rho_R(S_R - u_R)}. \tag{28}$$

With each of the characteristics approximated, the scheme fluxes can be computed based on the state occupying the cell edge, with an additional correction term for the interior starred states. Thus the fluxes to close equation (18) are

$$\mathbf{f}(\mathbf{U}_L, \mathbf{U}_R) = \begin{cases} \mathbf{f}_L = \mathbf{F}(\mathbf{U}_L) & 0 \leq S_L \\ \mathbf{f}_L^\star = \mathbf{F}(\mathbf{U}_L) + S_L(\mathbf{U}_L^\star - \mathbf{U}_L) & S_L \leq 0 \leq S_C \\ \mathbf{f}_R^\star = \mathbf{F}(\mathbf{U}_R) + S_R(\mathbf{U}_R^\star - \mathbf{U}_R) & S_C \leq 0 \leq S_R \\ \mathbf{f}_R = \mathbf{F}(\mathbf{U}_R) & S_R \leq 0 \end{cases} \tag{29}$$

where the left and right starred state vectors are given by

$$\mathbf{U}_K^\star = \left(\frac{S_K - u_K}{S_K - S_C}\right) \begin{pmatrix} \rho_K \\ \rho_K S_C \\ \rho_K v_K \\ E_K + (S_C - u_K)\left(S_C + \frac{p_K}{S_K - u_K}\right) \end{pmatrix} \tag{30}$$

when solving the dimensionally split equation in the $\hat{\mathbf{x}}$ direction. For solution in the $\hat{\mathbf{y}}$ direction, the contact wave speed (S_C) in the x-momentum relation should be replaced with u_K, and moved to the term in the y-momentum.

Following are two quite different techniques for constructing 2nd order schemes, each based on this HLLC approximate Riemann solution.

2.3.2 Limiters

The original 1st order Godunov method now must be modified to produce the desired 2nd order accuracy. An ideal 2nd order modified Godunov method is 2nd order everywhere the solution is smooth. These so-called “high resolution” schemes should be consistent, stable, and therefore convergent as well as being numerically inviscid and resolving discontinuities without artificial oscillations.

A metric for these oscillations inherent to 2nd order methods is the Total Variation. The Total Variation is defined as,

$$\mathrm{TV}\left(\mathbf{U}^n\right) = \sum_{i=-\infty}^{\infty} |\mathbf{U}_{i+1}^n - \mathbf{U}_i^n|, \tag{31}$$

such that any spurious oscillations introduced into $\mathbf{U}^n$, however small, will increase the total variation. If the total variation is monotonically decreasing in time, then it is assured there are no oscillations near discontinuities. This total variation diminishing (TVD) property is a necessary requirement to construct a high resolution 2nd order scheme; however, Godunov (1959) proved that the highest order monotone (and therefore TVD) scheme is only 1st order! A constraint on 2nd order solutions must then be imposed to regain TVD. Such limiting functions on either the scheme flux or the reconstructed slope can be specified to meet this requirement. Limiting selectively suppresses the technique used to gain 2nd order, returning the scheme to the 1st order monotone scheme where the solution becomes steep, while maintaining higher-order accuracy in smooth regions. This “steepness” is given by

$$\mathbf{r}_i = \frac{\Delta \mathbf{U}_{i,\mathrm{up}}}{\Delta \mathbf{U}_{i,\mathrm{down}}} \tag{32}$$

which are the upwind and downwind differences of the entire state vector $\mathbf{U}$.

Flux limiters, $\psi(r) \in \mathbb{R}$, in particular limit the fraction of the scheme flux allowed from the higher-order scheme:

$$\mathbf{f}_{i+\frac{1}{2}}^{\mathrm{TVD}} = \mathbf{f}_{i+\frac{1}{2}}^{\mathrm{LO}} + \psi_{i+\frac{1}{2}}\left(\mathbf{f}_{i+\frac{1}{2}}^{\mathrm{HI}} - \mathbf{f}_{i+\frac{1}{2}}^{\mathrm{LO}}\right). \tag{33}$$

It is apparent that the 1st order scheme is recovered if $\psi = 0$, or exactly the higher order scheme if $\psi = 1$, as in the definition of a general limiter. Sweby (1984) proved the requirements of the

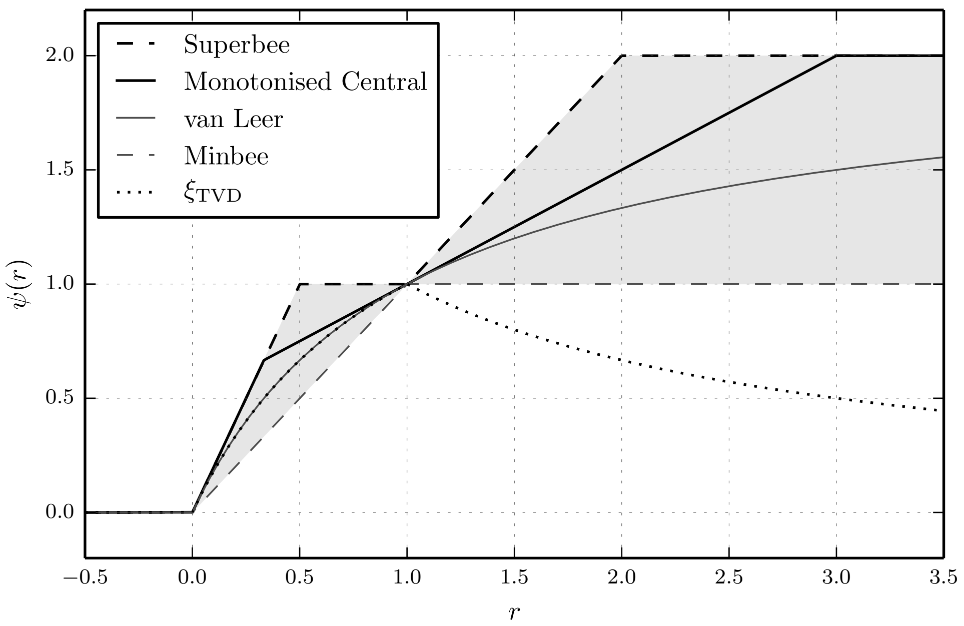

Figure 4: The Monotonised Central and van Leer flux limiter functions within the Sweby 2nd order TVD region (grey). The additional slope limiter restriction is also shown (dotted).

flux limiter function to create a monotone scheme. The Minbee and Superbee limiters outline the lower and upper limits respectively of this TVD region (in grey in Figure 4). Two popular limiters that will be used throughout this paper are the van Leer and Monotonised Central, defined as

$$\text{VAN LEER:}\ \psi(r) = \frac{2r}{1+r} \tag{34}$$

$$\text{MONOTONISED CENTRAL:}\ \psi(r) = \min(2r, \frac{1+r}{2}, 2) \tag{35}$$

where each is identically 0 for $r \leq 0$. These flux limiters are functions of a scalar steepness, and so just the density component of the steepness will be used here to represent the local flow condition.

In contract with flux limiters, slope limiters are used to make data-reconstructed higher order schemes monotone. They are analogous to flux limiters in switching from a higher order method to a monotone method; however, this is accomplished by limiting the slope of the reconstructed generalised Riemann problem. Each variable must now be limited independently, using the steepness for each respective variable, so that the slope limiter is a vector. There is an additional constraint on the Sweby flux limiter TVD region for slope limiters (see Figure 4). Slope limiting functions, $\xi(r)$, must additionally satisfy

$$\begin{aligned} \xi(r_i) \leq \xi_{\text{TVD}}(r_i) &= \min\left(\frac{2\beta_{i+\frac{1}{2}}r}{1+r}, \frac{2\beta_{i+\frac{1}{2}}}{1+r}\right) \\ &\approx \min\left(\frac{2r}{1+r}, \frac{2}{1+r}, 0\right), \end{aligned} \tag{36}$$

where the most conservative limit for $\beta_{i+\frac{1}{2}}$ is used. In general β is a function of the local Courant number, which would necessitate solving an additional Riemann problem. Applying this restriction to the flux limiters, it is apparent that the analogous slope limiters are simply

$$\xi_i = \min\left(\psi_i(r_i), \xi_{\text{TVD}}(r_i)\right). \tag{37}$$

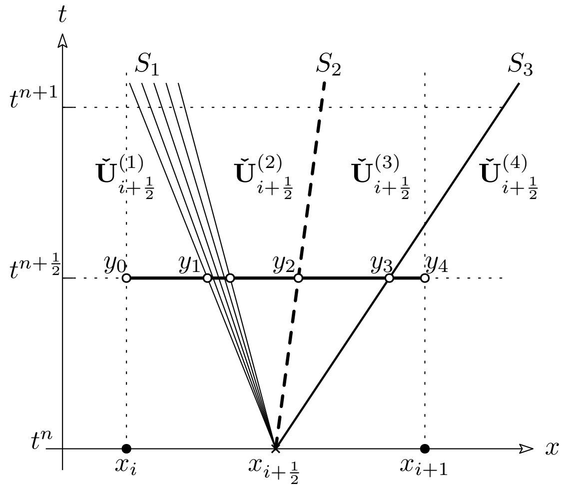

Figure 5: The WAF integration path (38) across a Riemann solution.

Similarly to flux limiters, in the limit that $\xi = 0$ the reconstructed Riemann problem is just the same as the original piecewise constant profile, and so the method is once again 1st order and monotone near discontinuities.

Care must be taken when the equilibrium solution consists of non-zero gradients in the conserved quantities, such as is the case with the background shear in the shearing box. The limiter still maintains monotonicity towards 0, and so even limits the equilibrium solution. This resulting TVD scheme is not a high resolution scheme anymore because it is not 2nd order at the equilibrium! This can be easily fixed by only limiting the perturbation quantities, such that the scheme becomes monotone relative to the equilibrium solution.

The van Leer limiter is chosen in light of these considerations and will be used as the default limiter throughout the rest of the investigation (Van Leer, 1974). On the lower side of the Sweby region, Minmod is the most diffusive high resolution limiter because it is the most conservative with resorting to the 1st order scheme. In contrast, while Superbee is the least diffusive it occasionally distorts certain smooth solution due to an overcompressive tendency. This refers to Superbee artificially steepening discontinuities by liberally allowing $\psi > 1$, corresponding to a negative weight on the 1st order method! Van Leer updated his limiter function (1977), designing a good multipurpose flux limiter, the Monotonised Central limiter. However, it is not in the slope limiter TVD region, and so van Leer's original limiter will be used so as to maintain consistency between the two finite volume schemes described below.

2.3.3 WAF TVD Scheme

One way to achieve a higher order Godunov method is to utilise more information from each of the Riemann solution states. The Weighted Average Flux (WAF) scheme does just this by weighting the solution state based on the occupancy in the current cell at $\Delta t/2$. In this way, WAF is a Generalised Godunov scheme, because it expands the interpretation of the Godunov flux in equation (22).

Given each of the four states of the HLLC Riemann solution, $\check{\mathbf{U}}_{i+\frac{1}{2}}(x,t)$ in an absolute coordinate system, Toro (1991, 1999) proposed using the entire Riemann solution by a weighted average of the flux at time $t^n + \Delta t/2$, between x_i and x_{i+1} (see Figure 5):

$$\mathbf{f}_{i+\frac{1}{2}} = \frac{1}{\Delta x}\int_{x_i}^{x_{i+1}} \mathbf{F}\left(\check{\mathbf{U}}_{i+\frac{1}{2}}\left(x, t^n + \Delta t/2\right)\right) \mathrm{d}x. \tag{38}$$

Assuming there are no rarefactions in the Riemann solution, such that each of the states, $\check{\mathbf{U}}^{(k)}_{i+\frac{1}{2}}$, are constants (which is the case in the HLLC solution), then (38) can be simplified as the

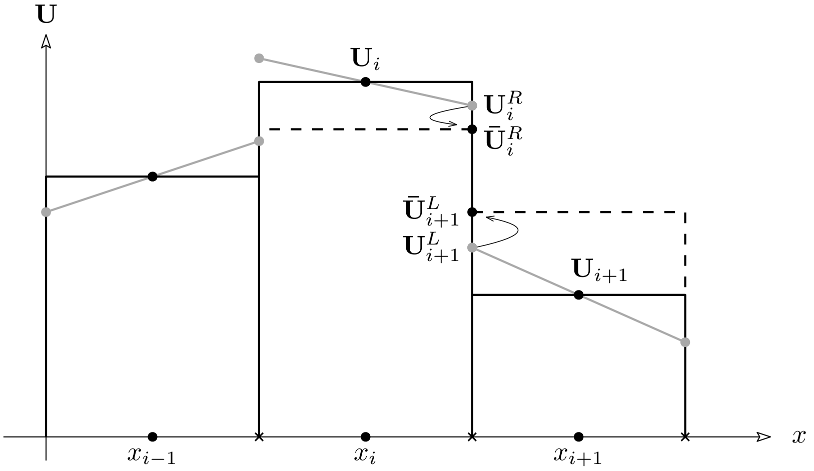

Figure 6: MUSCL-type reconstruction (grey), and MUSCL-Hancock evolution for edge $i + \frac{1}{2}$.

weighted sum,

$$\mathbf{f}_{i+\frac{1}{2}} = \sum_{k=1}^{N+1} \beta_k \mathbf{F}\left(\check{\mathbf{U}}_{i+\frac{1}{2}}^{(k)}\right). \tag{39}$$

$\beta_k = |y_k - y_{k-1}|/\Delta x$ are the weights for each of the states in Figure 5 and can be easily expressed in terms of the characteristics, S_k. Here, N is the number of characteristic lines in the Riemann solution (e.g. for HLLC, $N = 3$, but for a 1D system of equations, $N = 2$, as in the HLL Riemann solution).

The scheme flux (39) still must be modified to be TVD, so as to construct a high-resolution method. Applying the flux limiter between corresponding state fluxes of adjacent cell Riemann solutions, the WAF flux sum (39) is rewritten as

$$\mathbf{f}_{i+\frac{1}{2}} = \frac{1}{2}\left(F(\mathbf{U}_i) + F(\mathbf{U}_{i+1})\right) - \frac{1}{2}\sum_{k=1}^{N} \operatorname{sign}(S_k)\varphi_{i+\frac{1}{2}}^{(k)}\left(\mathbf{F}\left(\check{\mathbf{U}}_{i+\frac{1}{2}}^{(k+1)}\right) - \mathbf{F}\left(\check{\mathbf{U}}_{i+\frac{1}{2}}^{(k)}\right)\right). \tag{40}$$

Here, $\varphi_{i+\frac{1}{2}}^{(k)} = \varphi_{i+\frac{1}{2}}\left(r^{(k)}\right)$ is the WAF flux limiter function, which combines the weighted average and standard flux limiter, (34), by the relation,

$$\varphi(r^{(k)}) = 1 - (1 - |C_k|)\psi(r^{(k)}) \tag{41}$$

where C_k is the CFL number of the k^{th} Riemann characteristic, $C_k = S_k \Delta t/\Delta x$, and $r^{(k)}$ is the ratio of the jump in density across wave k of the upwind Riemann solution, $\check{\mathbf{U}}_{\text{up}}$, to the same characteristic wave in the local Riemann solution.

2.3.4 MUSCL-Hancock Scheme

An alternative technique for constructing a higher order method is to reconstruct the piecewise constant discrete data to more accurately define the initial left and right states of the Riemann problem. This is the idea behind MUSCL-type methods (Monotonic Upstream-centred Scheme for Conservation Laws). Specifically, the MUSCL-Hancock scheme interpolates the variables within each cell using a piecewise linear reconstruction defining a more representative Riemann problem at each cell edge (see Figure 6). So as to not bias one direction over the other, the slope is defined from a centred difference, $\boldsymbol{\Delta}_i = {}^1\!/\!{}_2(\mathbf{U}_{i+1} - \mathbf{U}_{i-1})$, and is a vector corresponding with the vector of conservatives.

As the MUSCL acronym implies, the scheme is monotone and therefore TVD. The TVD constraint is now a slope limiter applied during the reconstruction stage to mitigate spurious oscillations (as in the left edge of cell x_i in Figure 6). The limited slope, $\bar{\mathbf{\Delta}}_i = \xi_i \mathbf{\Delta}_i$, is instead used to determine the boundary-extrapolated values at the left and right edges:

$$\begin{aligned} \mathbf{U}_i^L &= \mathbf{U}_i^n - {}^1\!/\!_2 \bar{\mathbf{\Delta}}_i \\ \mathbf{U}_i^R &= \mathbf{U}_i^n + {}^1\!/\!_2 \bar{\mathbf{\Delta}}_i . \end{aligned} \tag{42}$$

A 2^{nd} order approximation to this generalised Riemann problem given by the piecewise linear left and right states is made by evolving the edge states forward by ${}^1\!/\!_2 \Delta t$,

$$\begin{aligned} \bar{\mathbf{U}}_i^L &= \mathbf{U}_i^L + \frac{1}{2} \frac{\Delta t}{\Delta x} \left(\mathbf{F}\left(\mathbf{U}_i^L\right) - \mathbf{F}\left(\mathbf{U}_i^R\right) \right) \\ \bar{\mathbf{U}}_i^R &= \mathbf{U}_i^R + \frac{1}{2} \frac{\Delta t}{\Delta x} \left(\mathbf{F}\left(\mathbf{U}_i^L\right) - \mathbf{F}\left(\mathbf{U}_i^R\right) \right) . \end{aligned} \tag{43}$$

This reduces the generalised Riemann problem into one soluble by the methods of §2.3.1, as illustrated by the dashed lines in Figure 6. Employing the HLLC approximate Riemann solver (29) between states $\bar{\mathbf{U}}_i^R$ and $\bar{\mathbf{U}}_{i+1}^L$ gives the scheme flux across edge $i + \frac{1}{2}$.

2.3.5 Well-Balanced Correction

A naïve implementation of the previous splitting method with either the WAF or MUSCL schemes produces significant numerical artefacts on the seemingly trivial problem of maintaining the vertical hydrostatic equilibrium solution (as pointed out in Slyz and Prendergast, 1999). This is because the hydrostatic pressure profile, which at cell interfaces manifests as a pressure jump, is interpreted by the Riemann solver as a propagating wave in the direction of the gradient of the potential. Even with 2^{nd} order parallel source splitting, the advection terms are not exactly balanced with the source term, as it is analytically in

$$\frac{\partial p}{\partial z} = -\rho \frac{\partial \Phi}{\partial z} . \tag{44}$$

So-called "well-balanced" schemes (Greenberg and Leroux, 1996) exactly satisfy a discrete form of this equilibrium; however, often assuming some underlying thermodynamic equilibrium to do so (LeVeque, 1997; Botta et al., 2004; Chandrashekar and Klingenberg, 2015). The method introduced in Käppeli and Mishra (2016) rather relies on a simple pressure reconstruction and assumes no thermodynamic equilibrium. A 2^{nd} order version of their technique is implemented here, which makes a slope-limited MUSCL-type reconstruction of the perturbation quantities, differing from the extrapolated pressure equilibrium in the neighbouring cell centres approximated as

$$\begin{aligned} p_i(x_{i-1}) &= p_i + \frac{\rho_{i-1} + \rho_i}{2} \frac{\Phi_i - \Phi_{i-1}}{\Delta z} \Delta z \\ p_i(x_{i+1}) &= p_i + \frac{\rho_i + \rho_{i+1}}{2} \frac{\Phi_{i+1} - \Phi_i}{\Delta z} \Delta z . \end{aligned} \tag{45}$$

The potential gradient needs to be calculated discretely so that this extrapolation satisfies the discrete pressure equilibrium defined as

$$\frac{p_{i+1} - p_i}{\Delta z} = -\frac{\rho_i + \rho_{i+1}}{2} \frac{\Phi_{i+1} - \Phi_i}{\Delta z} . \tag{46}$$

The MUSCL reconstruction step is instead approximated in primitive variables, with $\bar{\mathbf{\Delta}}_i$ being the centred difference of the perturbation quantities. The reconstruction (42) then gives the left and right edge states as

$$\begin{aligned} \mathbf{W}_i^L &= \mathbf{W}_i^n - {}^1\!/\!_2 \left(\bar{\mathbf{\Delta}}_i + \Delta_p^L \right) \\ \mathbf{W}_i^R &= \mathbf{W}_i^n + {}^1\!/\!_2 \left(\bar{\mathbf{\Delta}}_i + \Delta_p^R \right) \end{aligned} \tag{47}$$

where the pressure component of the limited slopes are supplemented with the equilibrium pressure difference at the cell edges,

$$\begin{aligned} \Delta_p^L &= -\rho_i \left(\Phi_i - \Phi_{i-1} \right) \\ \Delta_p^R &= -\rho_i \left(\Phi_{i+1} - \Phi_i \right) . \end{aligned} \tag{48}$$

Thus for an arbitrary thermodynamic and hydrostatic equilibrium, this correction maintains the steady state equilibrium to machine precision.

2.4 Source Solver

With the advection schemes detailed above, the final solution operator of the splitting method (20) that is left to be found is $\mathcal{L}_s$. This source operator approximates the solution to the initial value problem

$$\frac{\partial \mathbf{U}}{\partial t} = \mathbf{S}(\mathbf{U}, t). \tag{49}$$

Similarly to the advection step, the source solver must also be 2^{nd} order and conservative. An integration scheme is conservative if it conserves areas in phase space — a symplectic integrator — which is equivalent to requiring each eigenvalue of $\mathcal{L}$ for the particular integrand to be unity.

The source term of interest (the right side of (9)) is linear and so can be written as $\mathbf{S}(\mathbf{U}^n, n\Delta t) = A_s\mathbf{U}^n$, making it feasible to use an implicit integrator. It will become apparent that the best choice of an integrator for the shearing box source terms is Crank-Nicolson,

$$\left(I - \frac{1}{2}\Delta t A_s\right)\mathbf{U}^{n+1} = \left(I + \frac{1}{2}\Delta t A_s\right)\mathbf{U}^n. \tag{50}$$

Although Crank-Nicolson is generally not a symplectic integrator, with the shearing box source terms, the eigenvalues of $\mathcal{L}_s = \left(I - \frac{1}{2}\Delta t A_s\right)^{-1}\left(I + \frac{1}{2}\Delta t A_s\right)$ are all unity, and so the source solver is conservative. Additionally, Crank-Nicolson is 2^{nd} order accurate in time, and unconditionally stable. Thus the time step used to solve the source term can be exactly that given by the CFL condition from the advection time step. Otherwise for conditionally stable schemes with a restriction on Δt, the operator must be subcycled n_{sc} times, which is based on how much smaller the stable source time step is compared to the stable advection time step: $n_{\text{sc}} = \lceil \Delta t_x / \Delta t_s \rceil$. The source solution operator would then be replaced by

$$\mathcal{L}_s^{\frac{1}{2}\Delta t} \rightarrow \left(\mathcal{L}_s^{\frac{1}{2}\Delta t / n_{\text{sc}}}\right)^{n_{\text{sc}}}. \tag{51}$$

This will be necessary when solving for a general nonlinear source, where then an explicit ODE solver must be used.

This source integration is applicable for the unbalanced schemes; however, the well-balanced correction relies on the advection and source term truncation errors of the equilibrium profile exactly cancelling, which requires the time-marching update of (13) to be unsplit. The 2^{nd} order modified Runge-Kutta Strong Stability-Preserving scheme is then used to integrate $A_s\mathbf{U} + A_x\mathbf{U}$ forward in time together, so as to maintain the TVD property of the advection scheme. The conservative source integration must then be forfeited in exchange for a vertically well-balanced method.

2.5 Boundary Conditions

The boundary conditions are enforced by introducing fictitious ghost cells around the entirety of the domain (the number of which depends on the particular scheme). Then before each advection step, the ghost cells are updated to reflect the desired boundary conditions. This procedure is non-invasive to the the interior of the domain since each quantity is defined at cell-centre. By passively enforcing the boundary conditions in this way the finite volume method retains the desired conservative qualities.

Enforcing the periodic boundary (6) in $\hat{\mathbf{y}}$ is as straightforward as updating the top ghost cells with the values from the bottom (distance L_y away) and similarly with the bottom ghost cells. The domain in $\hat{\mathbf{y}}$ is now effectively of infinite extent and it is apparent that this boundary is conservative. Conservatively enforcing the shift-boost boundary (7) in $\hat{\mathbf{x}}$ is less trivial. This is because in general the shift, $v_{\text{boost}}t$, will not be an integer multiple of Δy needed to align with the opposite ghost cell, and so careful interpolation is necessary. First, the standard periodic boundary condition is applied (as above), with the boost from (7b) in conservative variables,

$$\rho v(x, y, z) \mapsto \rho v(x \pm L_x, y \mp v_{\mathrm{boost}} t) \mp \rho v_{\mathrm{boost}} \tag{52}$$

$$E(x, y, z) \mapsto E(x \pm L_x, y \mp v_{\mathrm{boost}} t) + \frac{1}{2} \rho v_{\mathrm{boost}}^2 \mp \rho v v_{\mathrm{boost}}. \tag{53}$$

Each row of ghost cells is then reconstructed with a 2^{nd} order slope-limited conservative reconstruction similar to that of MUSCL in §2.3.4. Finally, the values of each ghost cell are determined by integrating along the appropriately shifted fractions of the cells by using the definition of the finite volume discretisation (10), maintaining the conservative property of the finite volume method.

Reflective (no-flux) boundaries in $\hat{\mathbf{z}}$ can be realised by setting each ghost cell equal to the value of the nearest physical cell, except for the vertical velocity which is set to the negative of this value. However, with this prescription, the ghost cells will not be in hydrostatic equilibrium resulting in a continuous energy flux. The density and pressure must then be extrapolated into the ghost cells by fully reconstructing the discrete hydrostatic equilibrium with (46), and assuming a constant entropy in the ghost cells.

Therefore, because each of the advection solvers of §2.3, source solver outlined in §2.4, combined with the splitting method described in §2.2, as well as the boundary reconstruction above are 2^{nd} order, the numerical method outlined here is overall 2^{nd} order accurate in time and space. Additionally, because the advection step is conservative by construction, the source solver is symplectic, and the boundary conditions are all conservative, overall the algorithm is conservative.

2.6 Initial Conditions

A set of initial conditions for each particular problem or test completes the full specification of the governing initial boundary value problem. An initial force balance imposed by the pressure field is often desired, given a set of arbitrary initial perturbations on the background equilibrium solution. Ignoring for now the $\hat{\mathbf{z}}$ components and taking the divergence of the in-plane momentum equations (5b) gives the non-homogeneous Laplacian equation,

$$\nabla^2 p = \nabla \cdot (\mathbf{u} \cdot \nabla) \mathbf{u} + 2\rho\Omega_0 \left(\frac{\partial v}{\partial x} - \frac{\partial u}{\partial y} \right) + 2q\Omega_0^2 \tag{54}$$

specifying the pressure field. For the equilibrium shear flow, each of the right hand side terms cancel and the solution is trivially a constant pressure. This will not be the case in general for the test problems of §4, and requires solution of the $N_x \times N_y$ equations formed by discretising (54) using centred differences of p and $\mathbf{u}$. This set of equations is typically a well-behaved system which can be solved using relaxed Jacobi iterations. This solution is analogous when the $\hat{\mathbf{z}}$ components in (54) are included, however, the discretisation must match that of the *discrete* hydrostatic equilibrium (46).

3 Validation

The detailed algorithms and newly developed code are submitted to scrutiny via a series of 1D and 2D tests developed by Toro (1999) to comprehensively strain the working limits of these numerical methods as well as evince the strengths and weaknesses of each.

3.1 1D Tests

A series of five 1D tests define macroscopic Riemann problems producing varying magnitudes and combinations of shocks, rarefactions, and contact discontinuities. These are summarised in Table 1, including the simulation time and the location, x_c, of the initial discontinuity.

The results of a more difficult modification of the 1D tests are rather presented, which are conducted in 2D on an oblique plane. The initial discontinuity line is chosen to be located at $y = x$, which is the most unfavourable and diffusive direction for information to propagate in 2D. Each domain, $x, y \in \left[0, \sqrt{2}\right]$, is constructed with transmissive boundary conditions, and represents the solution on $N_x = N_y = 100$ cells. Each finite volume scheme tested employs the HLLC approximate Riemann solver and the van Leer limiter function. Variables are represented in code units which needn't be defined relative to an absolute at the moment, as long as dimensional consistency is ensured.

Test 1: Modified Sod A first and standard test is the Sod shock tube problem — the both experimentally and analytically well known test (Sod, 1978). The solution identically matches the Riemann problem shown in Figure 3b. It consists of a right-travelling contact wave (at $x = 0.55$ for the time shown in Figure 7), which separates a right shock wave (at $x = 0.75$) and a leftward sonic rarefaction wave (centred about $x = 0.3$).

Test 2: 123 Problem Secondly, to test the performance at near-vacuum conditions, a centred explosion will provide the desired large divergence. This is captured in the "123 Problem" which sets out two strong rarefaction waves symmetrically about a stationary contact line. Many methods find it difficult to capture the central $p^\star$ state of this test and fail spectacularly.

Test 3: Blast Wave Designed by Woodward and Colella (1984), the blast wave problem is a quite extreme test of the robustness of the implemented schemes. A large pressure difference (of 5 orders of magnitude) across the domain generates a strong rightward propagating shock wave with Mach number nearly 200. This quickly propagating shock is often hard for many methods to resolve, requiring many grid cells to represent the discontinuity. Additionally the shock wave generates a substantial amount of entropy which further strains some methods. A variant of this test will be used later in a 2D axisymmetric explosion.

Test 4: Collision of Two Shocks This carefully designed test of Toro (1999) produces two primary shock waves, whose collision results in two rightward propagating shock waves separated by a contact wave. This test is realised by creating a region of large negative divergence by setting up a high pressure region impinging on a low pressure state also converging toward the centre.

Test	Description	t_{stop}	x_c	ρ_L	u_L	p_L	ρ_R	u_R	p_R
1	Modified Sod	0.2	0.3	1.0	0.75	1.0	0.125	0.0	0.1
2	123 Problem	0.12	0.5	1.0	-2.0	0.4	1.0	2.0	0.4
3	Blast Wave	0.012	0.5	1.0	0.0	1000.0	1.0	0.0	0.01
4	Shock Collision	0.035	0.4	6.0	19.5975	460.894	6.0	-6.19633	46.095
5	Stationary Contact	0.012	0.8	1.0	-19.5975	1000.0	1.0	-19.5975	0.01

Table 1: Summary of the 1D test regimen.

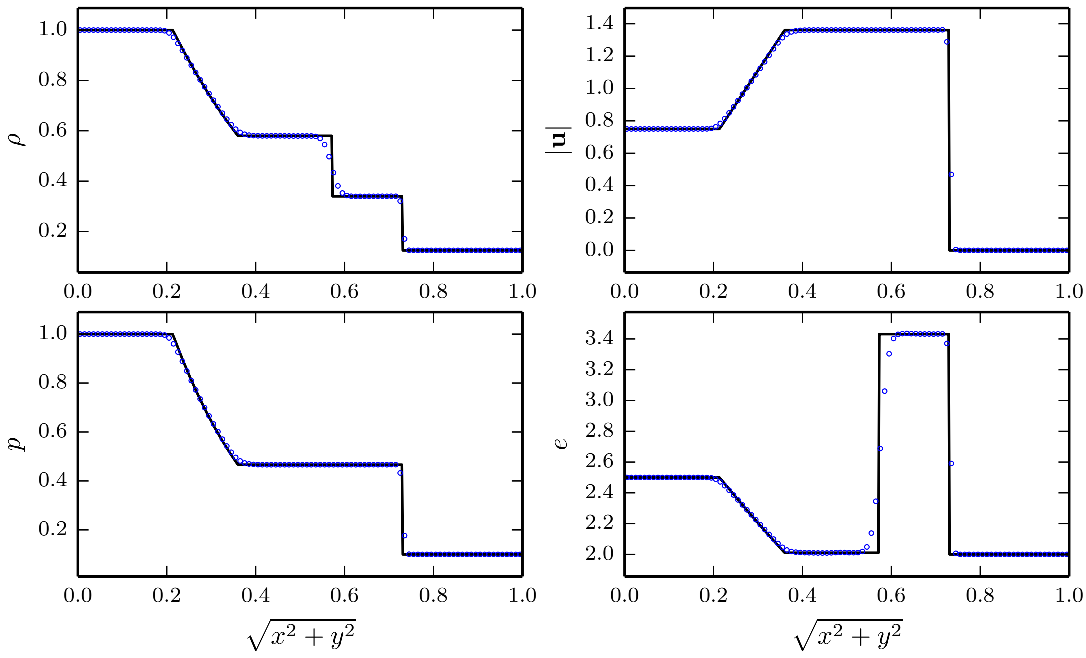

Figure 7: WAF TVD evolving Test 1 at $t = 0.2$, with the exact Riemann solution for reference.

Test 5: Stationary Contact Finally, as some methods find difficulty in capturing slowly-moving contact discontinuities, this final test was designed to produce just these conditions. A stationary contact line can be easily generated by embedding the blast wave conditions of Test 3 into a background flow moving at precisely the negative of the contact wave speed, $S_C = 19.5975$. This produces a contact discontinuity which remains at $x = x_c$, in addition to a left rarefaction and right shock wave.

3.1.1 WAF Performance

Of particular interest when analysing each test is the degree to which the exact solution is captured in smooth regions as well as the amount of numerical diffusion at each discontinuity, measured in the number of grid cells across the jump.

The Sod shock tube (Test 1) is very benign, and so the WAF TVD scheme performs quite well, nearly exactly resolving the smooth regions (Figure 7). This implementation also attains the same resolution across discontinuities as Toro's `NUMERICA` code. Convergence to the exact solutions are also demonstrated with increasing N.

The remaining test results are placed in Appendix A. Specifically, Test 5 fails due to WAF averaging an (albeit small) component of the large pressure state at the discontinuity. However, this is alleviated by using energy as the characteristic variable for limiting rather than density. Test 2 produces seemingly unacceptable disagreement in the internal energy plot; however, any errors in the centre state are amplified when dividing by the density of the near vacuum. Both Tests 4 and 5 contain spurious oscillations characteristic of the WAF TVD scheme with the van Leer limiter, and which are identically matched when employing `NUMERICA`. This non-monotone solution can be alleviated by using a more conservative flux limiter such as Minbee.

3.1.2 MUSCL Performance

MUSCL was found to be generally more dissipative than WAF, which is apparent when comparing the discontinuity resolution in Figure 8. MUSCL fails to evolve Test 2 due to the reconstructed profile of the energy or density dropping below 0, but this is easily alleviated by

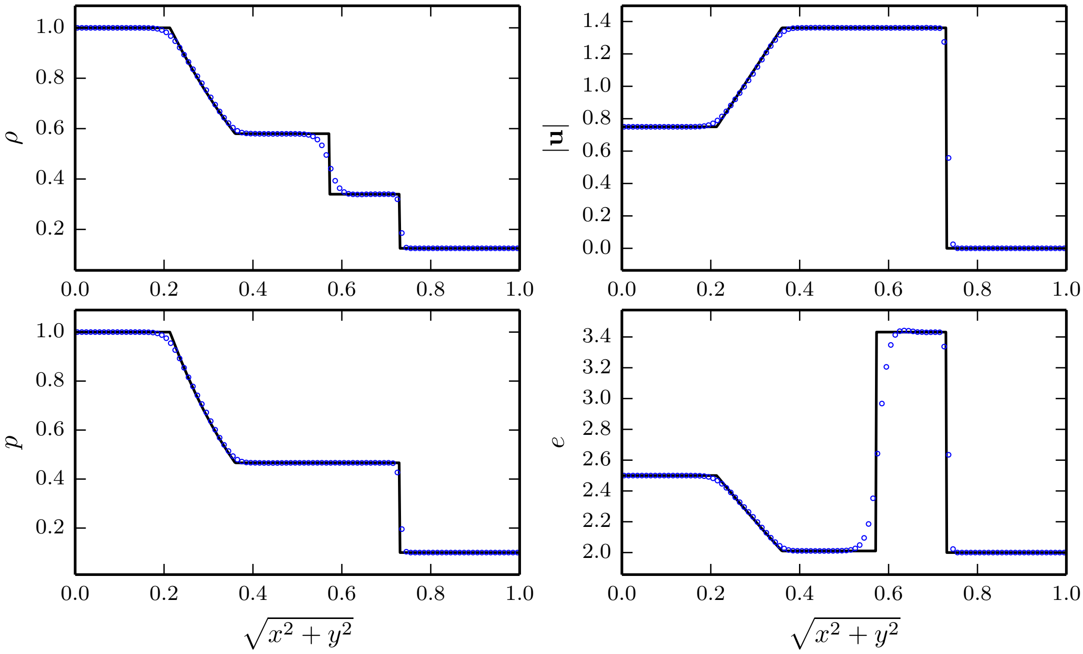

Figure 8: MUSCL TVD evolving Test 1 at $t = 0.2$, with the exact Riemann solution for reference.

enforcing a pressure and density floor during reconstruction. MUSCL also performs worse on Test 3, as measured by the density resolution of the contact wave. However, it does so with fewer oscillations as would be expected from a more diffusive scheme. Apart from the stationary contact line in Test 5, the WAF method tends to outperform MUSCL, and therefore will be the focus of the following 2D tests as well as the method of choice for further investigations.

3.2 2D Tests

The following 2D tests are designed to ensure there is minimal directional bias in the dimensional splitting method. This is accomplished by evolving an axisymmetric profile in 2D and comparing various radial slices. Ideally there would be no dispersion and the algorithm would produce truly axisymmetric results, agreeing exactly with the problem when formulated in 1D.

Unlike the 1D test problems which have an exact Riemann solution, there is no such simple exact solution with a geometric source term. A 1D axisymmetric model of the same problem using very high resolution (so as to be effectively converged to the true solution) must then be resorted to for verification. To solve the 1D radial axisymmetric problem, a geometric sourcing term must be introduced to the right side of the Euler equations,

$$\frac{\partial}{\partial t}\begin{pmatrix}\rho\\ \rho u_r\\ E\end{pmatrix}+\frac{\partial}{\partial r}\begin{pmatrix}\rho u_r\\ \rho u_r^2+p\\ (E+p)u_r\end{pmatrix}=-\frac{\alpha}{r}\begin{pmatrix}\rho u_r\\ \rho u_r^2\\ (E+p)u_r\end{pmatrix}, \tag{55}$$

where $\alpha = 0$ returns the 1D planar result, $\alpha = 1$ corresponds to 2D cylindrical symmetry, and $\alpha = 3$ is used for 3D spherical symmetry. This nonlinear source term can be solved using the 2 stage (2^{nd} order) Runge-Kutta integration scheme with appropriate subcycling of the split operator given by (51). This also happens to be an appropriate test to validate the source term implementation.

Explosion Test A 2D axisymmetric explosion is generated with similar initial conditions to the Sod shock tube problem, but in an initially static fluid. Fluid inside the explosion ($r < 0.4$)

is set to $\rho_{\text{in}} = 1.0$ and $p_{\text{in}} = 1.0$, while the outside initial state is $\rho_{\text{out}} = 0.125$ and $p_{\text{out}} = 0.1$. Again enforcing transmissive boundaries, and simulating until $t = 0.25$ allows time for the waves to propagate without reaching the boundaries. This results in a similar solution to the 1D Test 1 (Figure 7), except that each of the constant states now will be increasingly curved near $r = 0$ due to the geometric sourcing terms.

Weighted initial conditions based on the ratio of the area within r helps to mitigate the stair-step artefacts in the explosion boundary, which would propagate many frequencies independent of the numerical method. This initial condition is evolved on a centred 2D domain, $x, y \in [-1, 1]$, with 100×100 computational cells. The solution is compared against the 1D radial model (55) on $r \in [0, 1]$ using 4096 cells so as to ensure convergence.

Implosion Test An implosion test can similarly be set up by switching the initial states $\mathbf{W}_{\text{in}}$ and $\mathbf{W}_{\text{out}}$ from the above explosion test. These initial conditions are evolved on the same domain for $t = 0.15$. The geometry sourcing terms again cause increasing curvature with decreasing r. This is most evident in the shock body steepening, which eventually converges to $r = 0$ around $t = 0.22$ and forms a jet (or localised spike in pressure and density).

3.2.1 WAF Performance

The 2D tests will focus on the WAF TVD scheme since this will be the primary method for the shearing box results in §4; however, the MUSCL results are also included in Appendix A for reference.

The first thing to notice in Figure 9 is that the solution has the same number of points across the discontinuities as in the 1D Sod shock tube in Figure 7. It is similarly quite diffusive across the contact discontinuity (seen in the density and energy plots), and so is not due to the dimensional splitting method but rather the 1D advection scheme. The directional bias due to the splitting method can be quantified by measuring the standard deviation of the projection of the data onto the line $y = 0$ near the discontinuities. The largest standard deviation for WAF is measured to be $\sigma_{\max} = 0.022\Delta x$ whereas when using MUSCL is $\sigma_{\max} = 0.033\Delta x$. For both methods, the largest width is located at the contact discontinuity in the explosion test (Figure 9), and at the steepening shock front near $r = 0$ for the implosion (Figure 10).

This implementation is strained even further by evolving the implosion until collision at $r = 0$. Neither method deals with this shock collision very well at the low test resolution (Figures 28 and 29 in Appendix A) because the sourcing is singular there. Nonetheless the 2D axisymmetric simulations are shown to converge to the 1D case with decreasing Δx.

These tests have also built confidence that the source solution method described in §2.4 accurately implements the split source terms in 1D. Additionally, because the dimensional splitting has been shown to be convergent, it can be induced that the sourcing can similarly be extended to 2D, which will be exercised in the following section.

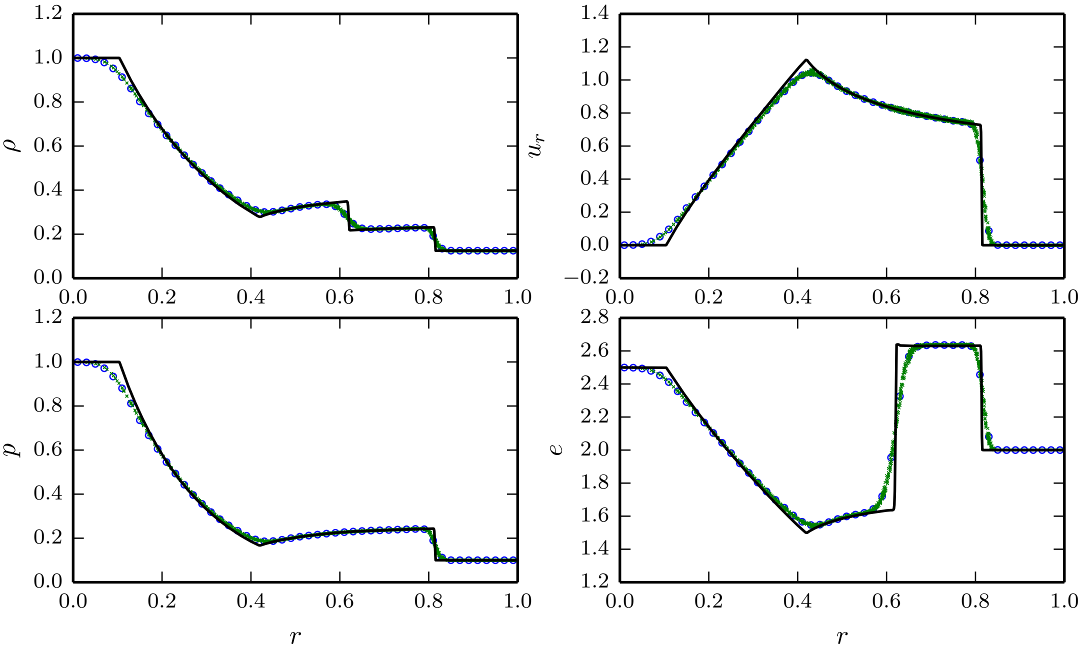

Figure 9: WAF TVD explosion test at $t = 0.25$, with the converged 1D solution for reference. Small ×'s mark the projected 2D values, and ∘'s indicate the average value in a bin of size Δx.

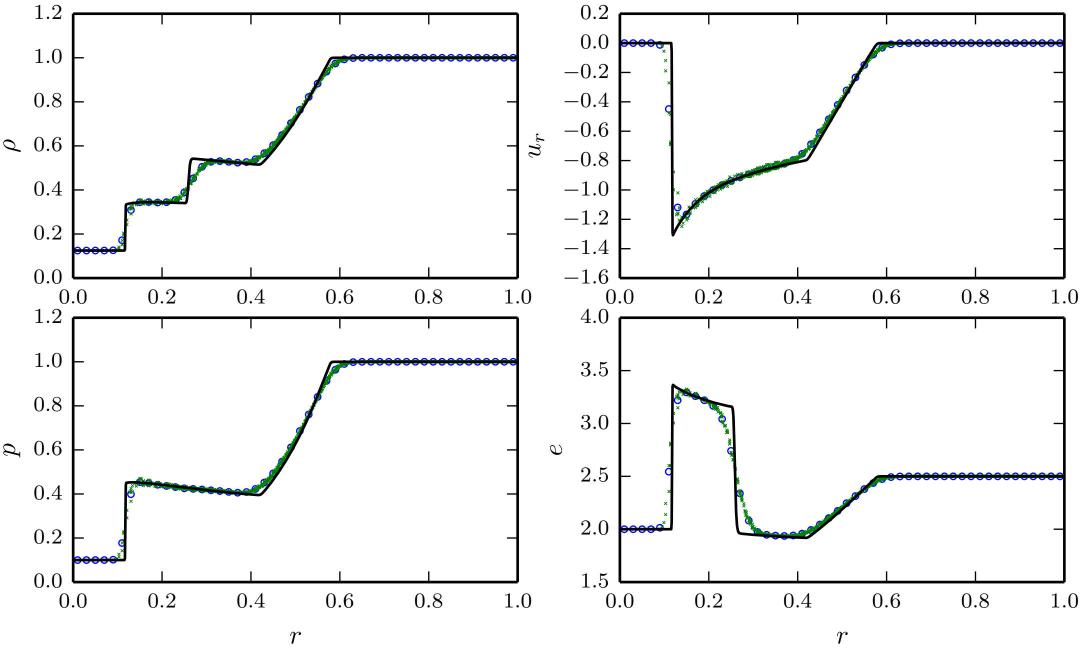

Figure 10: WAF TVD implosion test at $t = 0.15$, with the converged 1D solution for reference. Small ×'s mark the projected 2D values, and ∘'s indicate the average value in a bin of size Δx.

4 Results

The code implementation is now extended to solve the shearing box source terms as well as the conservative boundary conditions for x-y and x-z shearing sheets. A number of classic problems from the literature are explored in an effort to test these methods.

4.1 X-Y Shearing Box

Looking now just at the in-plane coordinates — that is just the $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ terms — the governing equations (9) written in the finite volume form evolving the vector of conservatives can be written explicitly as

$$\frac{\partial}{\partial t}\begin{pmatrix}\rho\\ \rho u\\ \rho v\\ E\end{pmatrix}+\frac{\partial}{\partial x}\begin{pmatrix}\rho u\\ \rho u^2+p\\ \rho uv\\ (E+p)u\end{pmatrix}+\frac{\partial}{\partial y}\begin{pmatrix}\rho v\\ \rho uv\\ \rho v^2+p\\ (E+p)v\end{pmatrix}=\begin{pmatrix}0\\ 2\rho v\Omega_0+2\rho q\Omega_0^2 x\\ -2\rho u\Omega_0\\ 2\rho uq\Omega_0^2 x\end{pmatrix}, \tag{56}$$

where it is clear from the form that the right hand side is the linear source, $\mathbf{S}(\mathbf{U})$. Periodic boundary conditions are applied in the $\hat{\mathbf{y}}$ direction, and shift-boost symmetry must be realised in the $\hat{\mathbf{x}}$ direction by applying the method described in §2.5.

4.1.1 Horizontal Epicycles

A first simple test is to track a uniform velocity perturbation to the background shear velocity. The solution comes from a linearisation of the global equations of motion, but is an exact solution in the shearing sheet. These are inertial (epicyclic) oscillations, termed because the restoring force is the Coriolis force, $\mathbf{F}_c = -2\rho\boldsymbol{\Omega}_0\times\mathbf{u}$. It is apparent that $\mathbf{F}_c\times\mathbf{u} = 0$, and so no work is done on the fluid particles; however, in discretisation this may not always be the case. The Coriolis force may then do work on a fluid particle, artificially increasing or decreasing the epicyclic energy. Although this implementation conserves the total energy by construction, the epicyclic energy (and the epicycle amplitude) in the inertial frame is not explicitly conserved to machine precision, but holds up to the truncation error of the source solver.

This test measures how well the epicyclic energy (8) is conserved by evolving epicycles for long times. The problem is set up following the test in Stone and Gardiner (2010), evolving an initial radial velocity perturbation $u_0 = 0.1c_s$ with $c_s = \Omega_0$ for a disk with aspect ratio unity. Three runs with increasing domain size ($L_x = L_y = 1, 10, 50$) and a 64×64 cell grid are meant to study the effect of an increased time step on the source integration method. Each was evolved for 100 orbital periods (equivalently epicyclic periods).

In each of the three runs the epicycles obtain a small error which is very weakly dependent on the domain size. The relative error in the velocity grows linearly at a rate of 10^{-5} per epicycle and is 90° out of phase. This corresponds to a relative error in the amplitude of $\sim 10^{-10}$ per epicycle, along with a small phase shift. Yet contrasting with Stone and Gardiner (2010), there is no visible phase error in the epicycles after 20 periods (Figure 11). Thus, although epicyclic motion is not conservative, the time for it to grow substantially is comparable with a typical viscous time scale, which exceeds the limiting assumptions of this shearing box model (5).

4.1.2 Shearing Internal Inertial Waves

Another dynamic which can arise as consequence of the rotating and shearing frame is inertial wave shearing. Goldreich and Tremaine (1978) originally investigated shearing inertial waves in the context of self-gravity, but the analysis is analogous with a point-mass potential. The solution is assumed to consist of transient perturbations to the radially unstratified shearing box equilibrium of the form $\delta\upsilon = \Re\left[\delta\hat{\upsilon}(t)\exp\left(\mathrm{i}k_x(t)x + \mathrm{i}k_y y\right)\right]$, for the eigenmode $\delta\hat{\upsilon}$ of some variable υ. By the process of the background shear, the radial wavenumber evolves in time as $k_x(t) = k_{x0} + q\Omega_0 k_y t$. Linearising the governing equations (5) for these small perturbations,

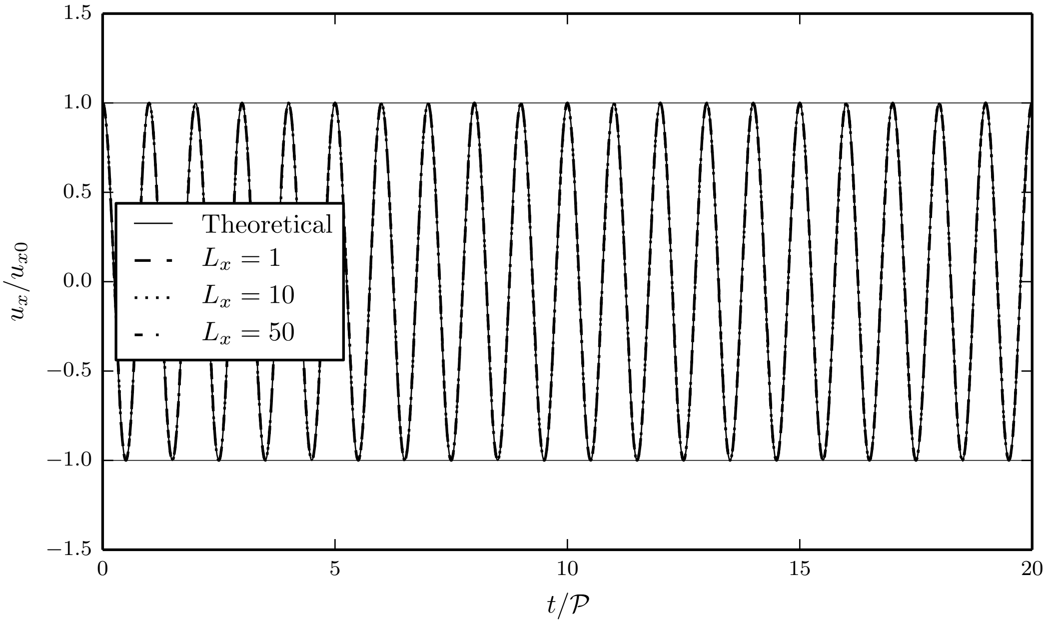

Figure 11: Epicycle velocity evolution for $L_x = L_y = 1$, 10, and 50.

$\delta\mathbf{W}$, admits four linearly independent solutions relating the perturbed variables — two vortical (incompressive[†]) solutions and two non-vortical (compressive) wave solutions.

This test will focus on the incompressive shearing wave solutions for short-wavelength perturbations which exhibit an initial growth or "wind-up" due to the shear before an asymptotic decay as $t \to \infty$. The solution is quite sensitive to numerical dissipation and dispersion which makes it a very good test of the implemented methods. Constraining the solutions to be vortical, $k_x\delta\hat{u} + k_y\delta\hat{v} = 0$, and solving with the linearised Euler equations requires a form of the initial conditions,

$$\delta\hat{v} = -\frac{k_x}{k_y}\delta\hat{u} \tag{57a}$$

$$\frac{\delta\hat{\rho}}{\rho_0} = \frac{\delta\hat{p}}{\gamma p_0} = -\frac{2\mathrm{i}}{c_s^2 k_y}(q-1)\Omega_0\delta\hat{u} \tag{57b}$$

where $k(t)^2 = k_x(t)^2 + k_y^2$ and $k_0^2 = k(0)^2$. Johnson and Gammie (2005) show that the perturbation kinetic energy, $\tilde{E}_k = {}^1\!/_2\rho|\delta\hat{\mathbf{u}}|^2$, for these vortical waves evolves as

$$\tilde{E}(t) = \frac{1}{2}\rho_0\delta\hat{u}_{x0}^2\frac{k_0^4}{k_y^2 k(t)^2}. \tag{58}$$

Thus the maximum amplification of the energy is $\mathcal{A} = 1 + (k_{x0}/k_y)^2$ which occurs at time $t = 4/(3\pi)\mathcal{P}$, where $\mathcal{P}$ is the orbital period.

Following the test of Stone and Gardiner (2010), the initial conditions are set as an equilibrium Keplerian background shear described by $\Omega_0 = 10^{-3}$, with $\rho_0 = 1$ and the unperturbed pressure determined by the isothermal sound speed, $c_s = 1.29 \cdot 10^{-3}$. The plane wave perturbation is characterised by $\mathbf{k} = (k_x, k_y) = 2\pi/L \cdot (-8, 2)$, with $\delta\hat{u}_{x0} = 10^{-4} \cdot c_s$, and the other perturbation variables satisfying (57). Thus, these parameters imply that a theoretical maximum

---

[†] Incompressive here does not refer to the compressibility of the Euler equations but rather to the absence of initial divergence.

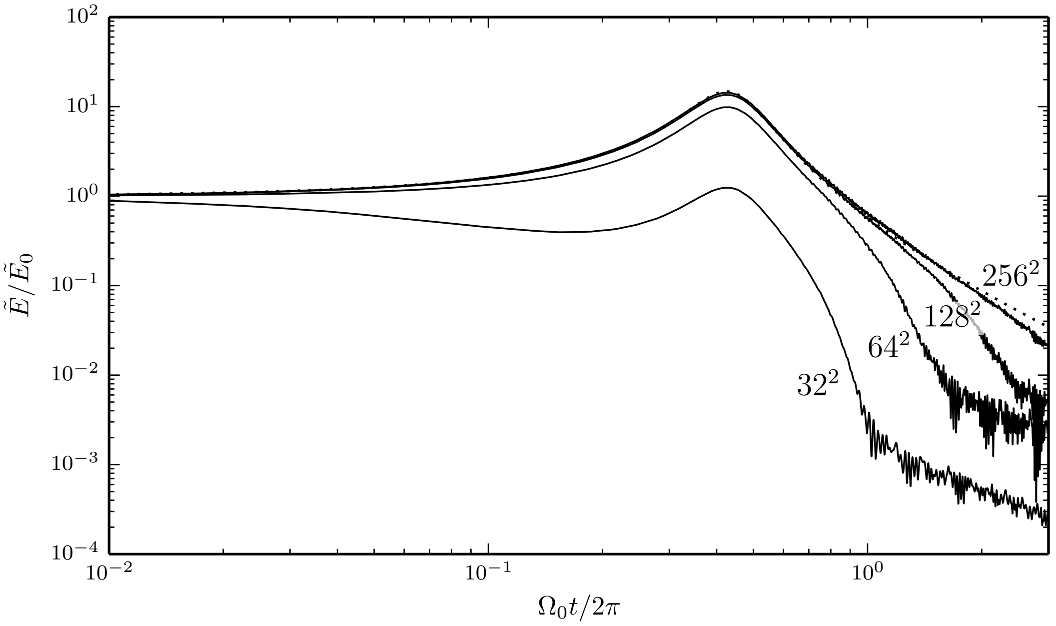

Figure 12: Perturbation kinetic energy evolution convergence study for vortical shearing waves.

amplification of $\mathcal{A} = 17$ should occur at $t = 4/(3\pi)\mathcal{P}$. These initial conditions are evolved in the shearing box domain with $L_x = L_y = 1$ using varying grid resolutions to determine convergence to the true solution.

A first attempt using the standard limiting method described in §2.3.2 results in a non-asymptotic decay after the initial wind-up, and an amplification factor of 15. This was determined to be due to the adverse effect of the limiter on the equilibrium shear profile. Even in the most benign case of uniform shear, the limiter on the $\hat{\mathbf{y}}$ momentum is not unity, and so it is never truly 2nd order. In fact, the limiter operates as designed, however, it ensures that each variable, including ρv is monotone, even if this artificially relaxes the background shear below the equilibrium. After modifying the limiter to only act on the perturbations from equilibrium, $\delta v = v - \bar{v}$, the monotone scheme regards the equilibrium state as zero. Very good convergence is then achieved even with 128^2 cells (Figure 12). Even when limiting only the perturbations, limiters can still affect convergence to a steady state solution, however in return for increased numerical stability. This is evidenced in the asymptotic decay dissipating faster than expected analytically. A final important note is that there is no evidence of aliasing, where the under-resolved trailing waves become so sheared that they are mistaken as leading waves and again periodically wind-up.

4.1.3 Vortex Energy Evolution

A final demonstration of the in-plane shearing sheet is with the evolution of a stable elliptic vortex. Moore and Saffman (1971) shows the required vorticity in a top hat vortex is

$$\omega_z = -\frac{q}{\chi}\frac{\chi + 1}{\chi - 1} \tag{59}$$

for the vortex aspect ratio, $\chi = \ell_y/\ell_x$, of a vortex with axis dimensions (ℓ_x, ℓ_y).

The perturbation kinetic energy, $\tilde{E}_k = {}^1\!/\!{}_2\rho|\delta\hat{\mathbf{u}}|^2$, and shear energy, $\tilde{E}_s = \rho\bar{v}\delta v$, are conserved in the anelastic approximation, as in Barranco and Marcus (2006), because the evolution of the equilibrium background states are restricted. However, because the implemented method here

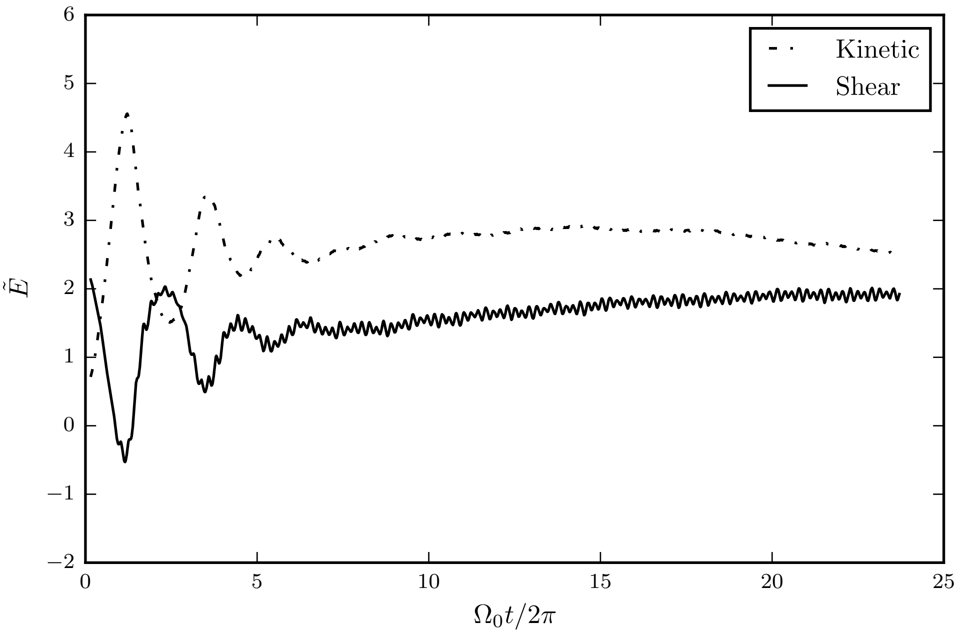

Figure 13: Kinetic and shear perturbation energy during evolution of a stable vortex.

solves the fully compressible Euler equations, there is no such conserved perturbation energy. Similarly to the epicyclic energy above, this energy is not conserved by construction like the total energy. Nonetheless, this problem can still give insight into the effects of numerical dissipation on physical systems.

A similar test is performed to that used by Barranco and Marcus (2006) in validating their pseudo-spectral anelastic hydrodynamics code. A vortex with dimensions $(\ell_x, \ell_y) = (0.25, 1)$ is evolved in a domain of size $(L_x, L_y) = (2, 4)$ with uniform density $\rho_0 = 1$. The Keplerian shear is described by $\Omega_0 = 1$ which requires $\omega_z = -0.625$ for stability. Finally, to ensure an initial discrete pressure equilibrium, the pressure field is solved by relaxed Jacobi iterations as described in §2.6 with pressure boundary conditions defined by $c_s = 1$.

After an initial transient, the perturbation energy relaxes to a steady state after nearly 100 orbital periods as the energy is dissipated as acoustic waves. The low frequency component of the initial transient seen in Figure 13 are the Kida oscillations (Kida, 1981), although they quickly dampen out due to numerical diffusion. These oscillations are much shorter lived than in a similar test of the nearly dissipationless pseudo-spectral methods used in Barranco and Marcus (2006). The high frequency oscillations apparent in the shear perturbation energy are exactly at the regridding frequency, and are due to the vortices in the neighbouring periodic boxes interacting each time they pass (see Figure 2).

4.2 X-Z Shearing Box

Now rather than looking in the midplane of the disc, a radial slice through the disc oriented normal to the midplane will be used as a new shearing sheet. Taking just the $\hat{\mathbf{x}}$ and $\hat{\mathbf{z}}$ terms from Equation (9) gives the azimuthally $(\hat{\mathbf{y}})$ invariant, local axisymmetric equations in the x-z shearing coordinates. The finite volume formulation, where now $\mathbf{U} = (\rho, \rho u, \rho w, E)^T$ is then

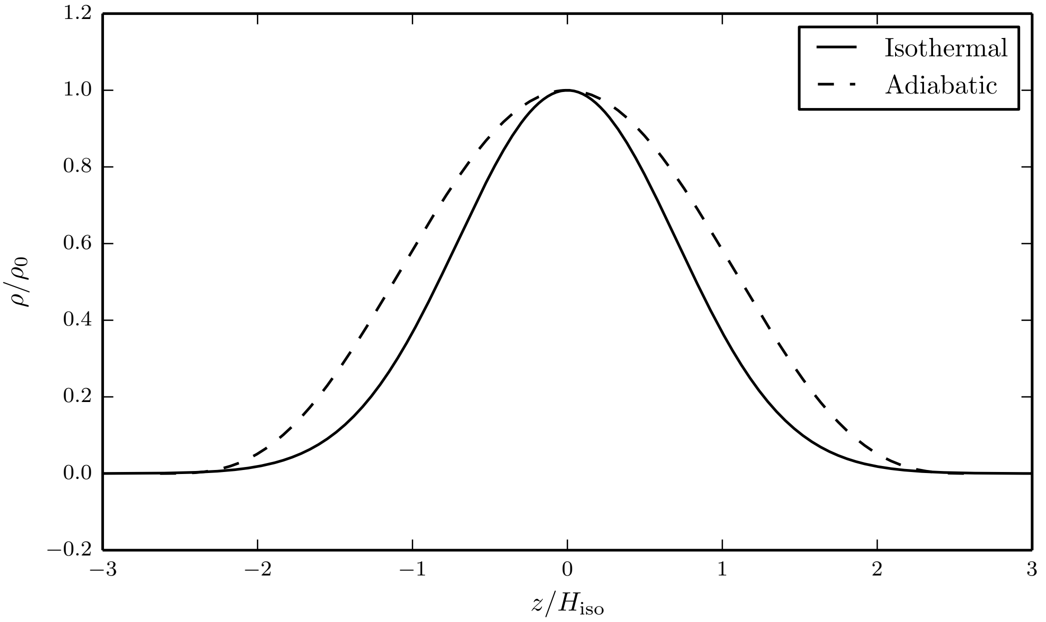

Figure 14: Density equilibrium profile corresponding to an isothermal and adiabatic thermodynamic equilibrium.

explicitly

$$\frac{\partial}{\partial t}\begin{pmatrix}\rho\\ \rho u\\ \rho w\\ E\end{pmatrix}+\frac{\partial}{\partial x}\begin{pmatrix}\rho u\\ \rho u^2+p\\ \rho uw\\ (E+p)u\end{pmatrix}+\frac{\partial}{\partial z}\begin{pmatrix}\rho w\\ \rho uw\\ \rho w^2+p\\ (E+p)w\end{pmatrix}=\begin{pmatrix}0\\ 0\\ -\rho\Omega_0^2 z\\ \rho\Omega_0^2(2qux-wz)\end{pmatrix}, \tag{60}$$

where the Coriolis terms in the momentum equations vanish completely. Additionally, by assuming a constant background shear, $\bar{v}=-q\Omega_0 x$, the effective potential gradient term $(2\rho q\Omega_0^2 x)$ in the $\hat{\mathbf{x}}$ momentum is balanced and cancels.

Periodic boundary conditions are similarly enforced in the $\hat{\mathbf{x}}$ direction. Assuming an initially vertical ($\hat{\mathbf{z}}$) isothermal equilibrium appropriate for an optically thin disc, the hydrostatic pressure profile is

$$p(z)=p_0\exp\left(-\frac{z^2}{2H^2}\right) \tag{61}$$

where H is the isothermal scale height. An initially adiabatic thermodynamic equilibrium is avoided, such that an artificial atmosphere need not be implemented in the vacuum beyond $\pm H_{\rm adiab}$, as in Figure 14). Finally, as described in §2.5, a pressure-extrapolated reflective boundary condition is enforced in the vertical direction at a height $z=\pm 4H$.

A typical test in shearing and vertically stratified domains involves inertia-gravity waves. Although inertia-gravity waves are axisymmetric, they cannot exist in this vertical 2D shearing sheet because the $\hat{\mathbf{y}}$ component of each wave mode is the source of the inertial restoring force. Inertia-gravity waves can exist in a so-called “2.5D” simulation where the azimuthal momentum is also evolved, but this will not be considered here.

4.2.1 Well-Balanced Equilibrium

A simple test of the well-balanced correction described in §2.3.5 is to evolve the isothermal hydrostatic equilibrium to ensure it is satisfied to machine precision. The *discrete* hydrostatic

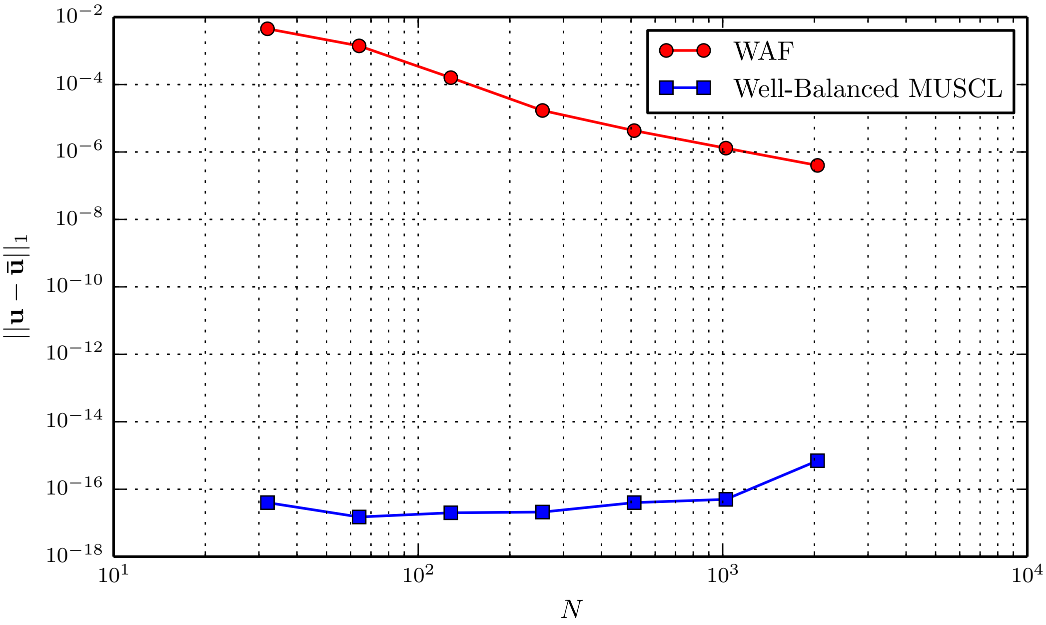

Figure 15: L_1-norm error in velocity for the well-balanced and unbalanced schemes.

equilibrium profile is evolved for 10 sound crossing times, H_{iso}/c_s, varying the vertical grid resolutions between $N_x = 32$ up to $N_x = 2048$. The 2^{nd} order well-balanced TVD MUSCL scheme is compared against the unbalanced WAF TVD scheme which was the focus of the previous tests. As evidenced in Figure 15, the well-balanced MUSCL method maintains hydrostatic equilibrium to machine precision. The unbalanced WAF scheme errors arise due to the operationally-split nature of the algorithm trying to balance the pressure gradient independently from the gravitational source. The largest discrepancy occurs at the midplane but remains on the order of the truncation error. This is evident in Figure 15 as is the 2^{nd} order accuracy of the WAF method. Each of the following tests will be conducted instead with this well-balanced MUSCL TVD method.

4.2.2 Vertical Epicycles

Similarly to the horizontal epicycles in the x-y shearing sheet, uniform velocity perturbations set off oscillations at the vertical epicyclic frequency (the orbital frequency for Keplerian discs). Vertical epicycles do not maintain the uniform velocity field as horizontal ones do because the body forces are not homogeneous anymore due to the pressure profile. The total vertical momentum will nonetheless remain sinusoidal. An additional difficulty in modelling the vertical epicycles is that the symplectic integration scheme used to conservatively integrate the source terms of the midplane shearing sheet cannot be used with the well-balanced MUSCL scheme of Käppeli and Mishra (2016), which requires explicit integration. Because this is integrating an energy component in the source term, the total energy of this system will inevitably not be conserved to machine precision.

The isothermal hydrostatic equilibrium for this test is defined by $c_s = 0.25\Omega_0$ and $\Omega_0 = 1$, such that the scale height is $H_{\mathrm{iso}} = 0.25$. Pressure-reconstructive reflecting boundaries are then enforced at $z = \pm 1$. A vertical velocity perturbation, $\delta w = 0.1 c_s$, finally is given to the quiescent fluid and evolved on a 128^2 cell grid for 100 epicycle periods.

The epicyclic amplitude is maintained to within a relative error of 10^{-6} over the course of 100 epicycles using the well-balanced MUSCL method, which is comparable to the unbalanced WAF scheme on the same test. In the WAF scheme, the truncation errors enter the epicyclic energy through the unbalanced portion of the advection and source terms, whereas in the well-balanced MUSCL scheme the error enters the epicyclic energy solely through the source term integration because the 2^{nd} order Runge-Kutta method is not symplectic.

4.2.3 Propagation of Waves through the Stratified Atmosphere

A final test on the robustness of this implementation is made by propagating plane waves vertically through the stratified disc. By a conservation argument, as the waves propagate into the rarefied atmosphere, the profile steepens eventually to the point of forming a shock front. There is no simple analytic solution for the waveform of these steepening waves; however, for isentropic waves, the velocity perturbation can be seen to amplify by

$$\mathcal{A} = \exp\left(\frac{z^2}{2H^2}\right) \tag{62}$$

which is just the inverse of the stratification scaling factor for an isothermal equilibrium disc.

Following the test in Käppeli and Mishra (2016), a discrete isothermal hydrostatic equilibrium is set up with $c_s = 0.25\Omega_0$ such that $H = 0.25$, and $\rho_0 = 1$. A wave generator is implemented at the disc midplane with vertical velocity perturbations $\propto \sin(4\pi\Omega_0 t)$, corresponding to 2 wavelengths per scale height since the sound crossing time is just Ω_0^{-1}. This configuration is evolved for 3 sound-crossing times with perturbations ranging between $|\delta v| = 10^{-5} c_s$ and $0.1 c_s$. 200^2 cells are initially used to see the performance at low resolution, but a converged reference profile is also computed in 1D with $N_y = 4096$ cells.

The expectations above are qualitatively met, where it is evident that perturbations of $|\delta v| = 0.05 c_s$ steepen sufficiently to form the saw-tooth shock fronts in Figure 17. At the same time, $|\delta v| = 0.01 c_s$ perturbations are still approximately sinusoidal (Figure 16). Still, if these waves were allowed to propagate into further rarefied fluid, they would eventually see a similar fate. The reason for the obvious asymmetry with larger amplitude perturbations is due to the perturbation velocity at $t = 0$ initially choosing a positive vertical velocity. This results in the leading shock on the right-most wave and a leading rarefaction (followed shortly by a shock) on the left-most wave. Compared to the converged reference profile, a small error in the wave speed appears with decreasing resolution. This is likely a limitation of the wave generator forcing at low resolutions rather than errors in the approximate Riemann solver.

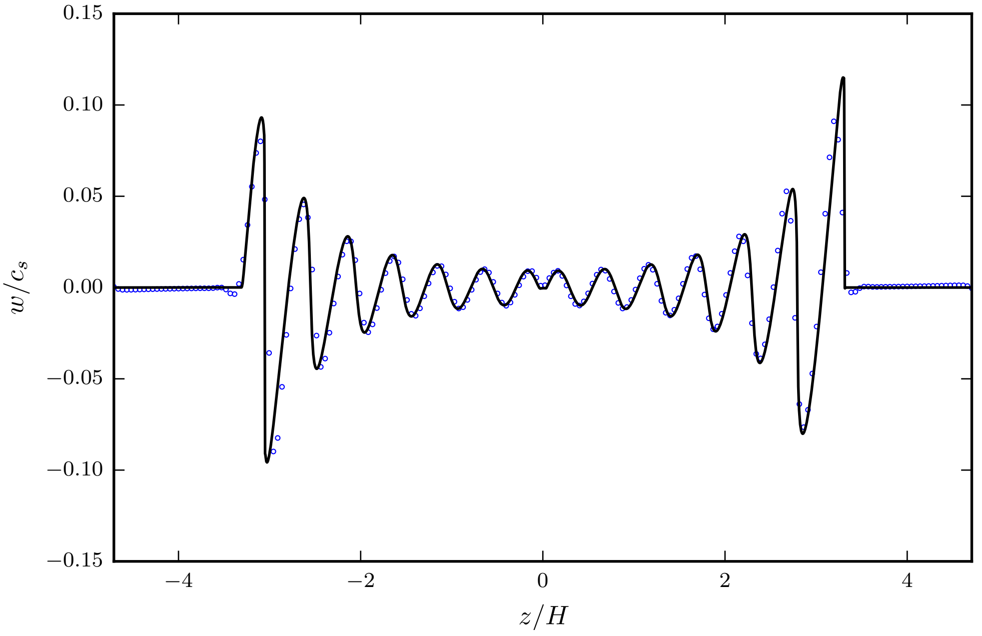

Figure 16: Perturbations of amplitude $|\delta v| = 0.01c_s$ remain sinusoidal to a good approximation. Numerical results with $N = 200$ cells along with the converged solution for reference.

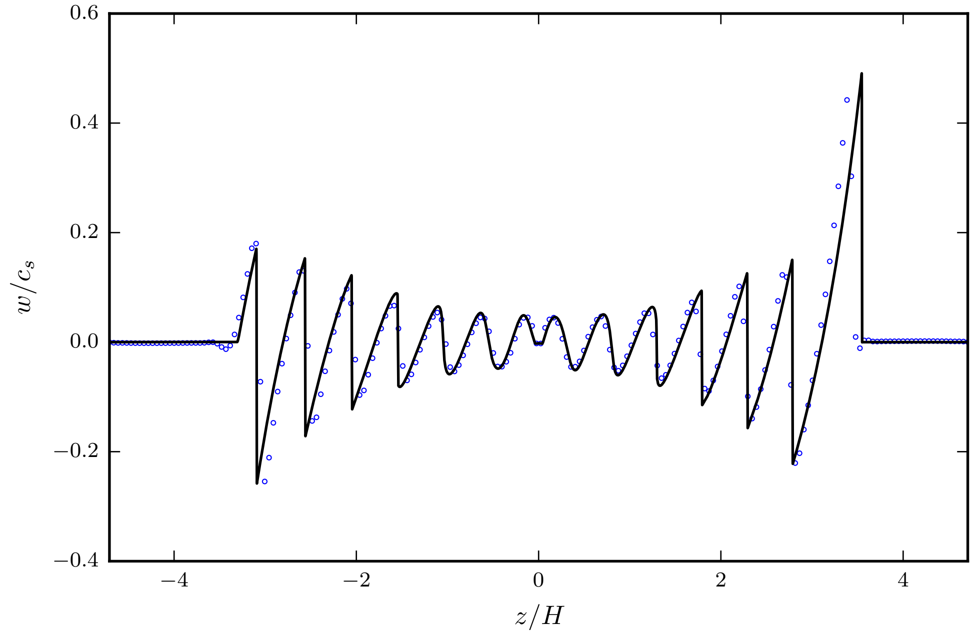

Figure 17: Perturbations of amplitude $|\delta v| = 0.05c_s$ form shocks in the rarefied disc atmosphere. Numerical results with $N = 200$ cells along with the converged solution for reference.

5 Conclusion

A comprehensive multidimensional test regimen has been presented for the compressible unsteady homogeneous Euler equations. Throughout the course of the presented 1D and 2D tests, the superiority of the 2^{nd} order flux-limited WAF scheme has become apparent over the numerically more diffusive MUSCL-Hancock method. The homogeneous governing equations were then supplemented with the shearing box source terms and used to explore various problems from the literature. The unbalanced WAF method was compared against the hydrostatically well-balanced and unsplit MUSCL method of Käppeli and Mishra (2016). In the 2D shearing box with a vertical hydrostatic balance, the well-balanced method succeeded at maintaining the discrete hydrostatic equilibrium to machine precision — a trivial analytic task in continuous space, but much more difficult numerically. This well-balanced scheme is necessary for sensitive stability analyses where localised perturbations on the order of the truncation error would artificially seed instabilities. Nonetheless, the unbalanced WAF method is still superior at capturing large perturbations from the hydrostatic balance as well as evolving the shearing sheet in the plane of the disc due to the ability of using a robust symplectic integrator which is conservative to machine precision. The total energy still was seen to be less relevant in the metric of the shearing box, and so future work will include constructing a numerical method to ensure the conserved epicyclic energy remains so (to machine precision) during discretisation — a more pertinent conserved quantity to the shearing box approximation. This requires stepping forward the $\hat{\mathbf{y}}$ perturbation momentum, $\rho(v + q\Omega_0 x)$, rather than ρv.

Part II of this series will extend the presented implementation to model general eccentric discs represented in the shearing box with time-varying non-orthogonal coordinates.

References

P. J. Armitage. Dynamics of Protoplanetary Disks. *Annual Review of Astronomy and Astrophysics*, 49(1):195–236, Sept. 2011.

J. A. Barranco and P. S. Marcus. A 3D spectral anelastic hydrodynamic code for shearing, stratified flows. *Journal of Computational Physics*, 219(1):21–46, Nov. 2006.

N. Botta, R. I. Klein, S. Langenberg, and S. Lützenkirchen. Well balanced finite volume methods for nearly hydrostatic flows. *Journal of Computational Physics*, 196(2):539–565, May 2004.

P. Chandrashekar and C. Klingenberg. A second order well-balanced finite volume scheme for Euler equations with gravity. *SIAM Journal on Scientific Computing*, 37(3):B382–B402, 2015.

S. K. Godunov. A difference method for numerical calculation of discontinuous solutions of the equations of hydrodynamics. *Matematicheskii Sbornik*, 1959.

P. Goldreich and S. Tremaine. The excitation and evolution of density waves. *The Astrophysical Journal*, 222:850–858, June 1978.

J. M. Greenberg and A. Y. Leroux. A well-balanced scheme for the numerical processing of source terms in hyperbolic equations. *SIAM Journal on Numerical Analysis*, 33(1):1–16, 1996.

A. Harten, P. D. Lax, and B. v. Leer. On Upstream Differencing and Godunov-Type Schemes for Hyperbolic Conservation Laws. *SIAM Review A*, 25(1):35–61, Jan. 1983.

J. F. Hawley, C. F. Gammie, and S. A. Balbus. Local Three-dimensional Magnetohydrodynamic Simulations of Accretion Disks. *The Astrophysical Journal*, 440:742, 1995.

B. M. Johnson and C. F. Gammie. Linear theory of thin, radially-stratified disks. *The Astrophysical Journal*, 626(2):978–990, 2005.

R. Käppeli and S. Mishra. A well-balanced finite volume scheme for the Euler equations with gravitation. *Astronomy and Astrophysics*, 587:A94, Feb. 2016.

S. Kida. Motion of an elliptic vortex in a uniform shear flow. *Physical Society of Japan*, 50(10): 3517–3520, 1981.

H. H. Klahr and P. Bodenheimer. Turbulence in Accretion Disks: Vorticity Generation and Angular Momentum Transport via the Global Baroclinic Instability. *The Astrophysical Journal*, 582(2):869–892, Jan. 2003.

R. J. LeVeque. Wave propagation algorithms for multidimensional hyperbolic systems. *Journal of Computational Physics*, 131(2):327–353, 1997.

P. S. Marcus, S. Pei, C.-H. Jiang, J. A. Barranco, P. Hassanzadeh, and D. Lecoanet. Zombie Vortex Instability. I. a Purely Hydrodynamic Instability to Resurrect the Dead Zones of Protoplanetary Disks. *The Astrophysical Journal*, 808(1), 2015.

D. W. Moore and P. G. Saffman. Structure of a Line Vortex in an Imposed Strain. In *Aircraft Wake Turbulence and Its Detection*, pages 339–354. Springer US, Boston, MA, 1971.

J. C. B. Papaloizou and D. N. C. Lin. Theory Of Accretion Disks I: Angular Momentum Transport Processes. *Annual Review of Astronomy and Astrophysics*, 33(1):505–540, 1995.

J. E. Pringle. Accretion Discs in Astrophysics. *Annual Review of Astronomy and Astrophysics*, 19(1):137–160, Sept. 1981.

O. Regev and O. M. Umurhan. On the viability of the shearing box approximation for numerical studies of MHD turbulence in accretion disks. *Astronomy and Astrophysics*, 481(1):21–32, Apr. 2008.

N. I. Shakura and R. A. Sunyaev. Black holes in binary systems. Observational appearance. *Astronomy and Astrophysics*, 24:337–355, 1973.

A. Slyz and K. H. Prendergast. Time-independent gravitational fields in the BGK scheme for hydrodynamics. *Astronomy and Astrophysics Supplement Series*, 139(1):199–217, Oct. 1999.

G. A. Sod. A survey of several finite difference methods for systems of nonlinear hyperbolic conservation laws. *Journal of Computational Physics*, 27(1):1–31, Apr. 1978.

J. M. Stone and T. A. Gardiner. Implementation Of The Shearing Box Approximation In Athena. *The Astrophysical Journal Supplement Series*, 189(1):142–155, July 2010.

P. K. Sweby. High Resolution Schemes Using Flux Limiters for Hyperbolic Conservation Laws. *SIAM Journal on Numerical Analysis*, 21(5):995–1011, Oct. 1984.

E. F. Toro. A linearized Riemann solver for the time-dependent Euler equations of gas dynamics. *Proceedings: Mathematical and Physical Sciences*, 434(1892):683–693, Sept. 1991.

E. F. Toro. *Riemann Solvers and Numerical Methods for Fluid Dynamics.* Springer, Berlin, Heidelberg, 2 edition, 1999.

E. F. Toro, M. Spruce, and W. Speares. Restoration of the contact surface in the HLL-Riemann solver. *Shock Waves*, 4(1):25–34, July 1994.

B. Van Leer. Towards the Ultimate Conservation Difference Scheme. II. Monotonicity and Conservation Combined in a Second-Order Scheme. *Journal of Computational Physics*, 14(4): 361–370, Mar. 1974.

P. Woodward and P. Colella. The numerical simulation of two-dimensional fluid flow with strong shocks. *Journal of Computational Physics*, 54(1):115–173, Apr. 1984.

A Appendix of Figures

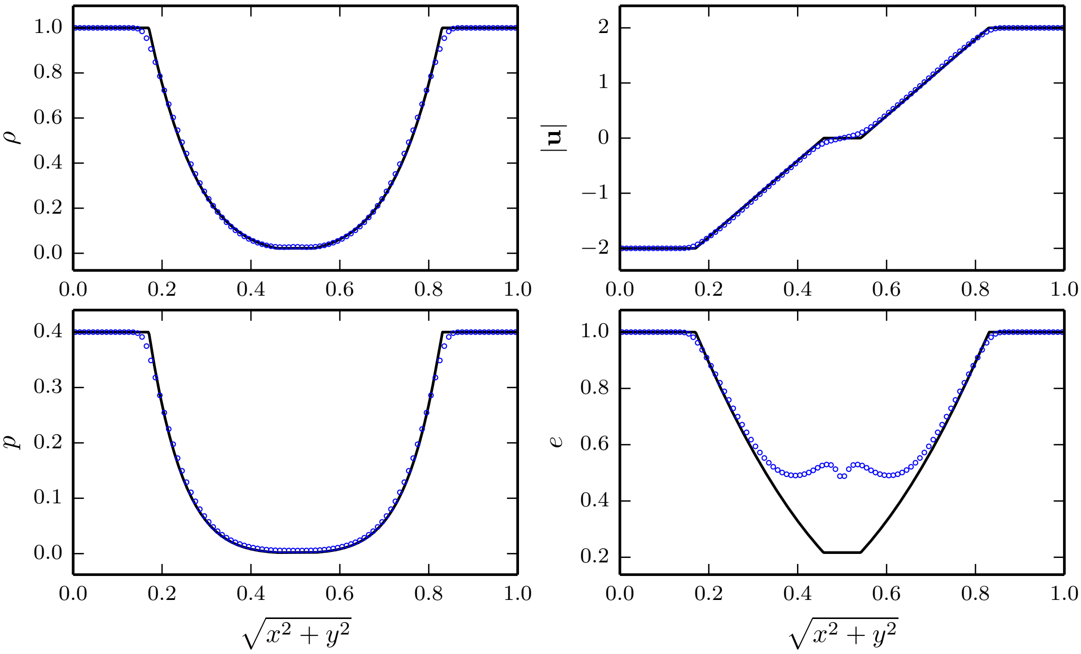

Figure 18: WAF TVD evolving Test 2 at $t = 0.12$, with the exact Riemann solution for reference.

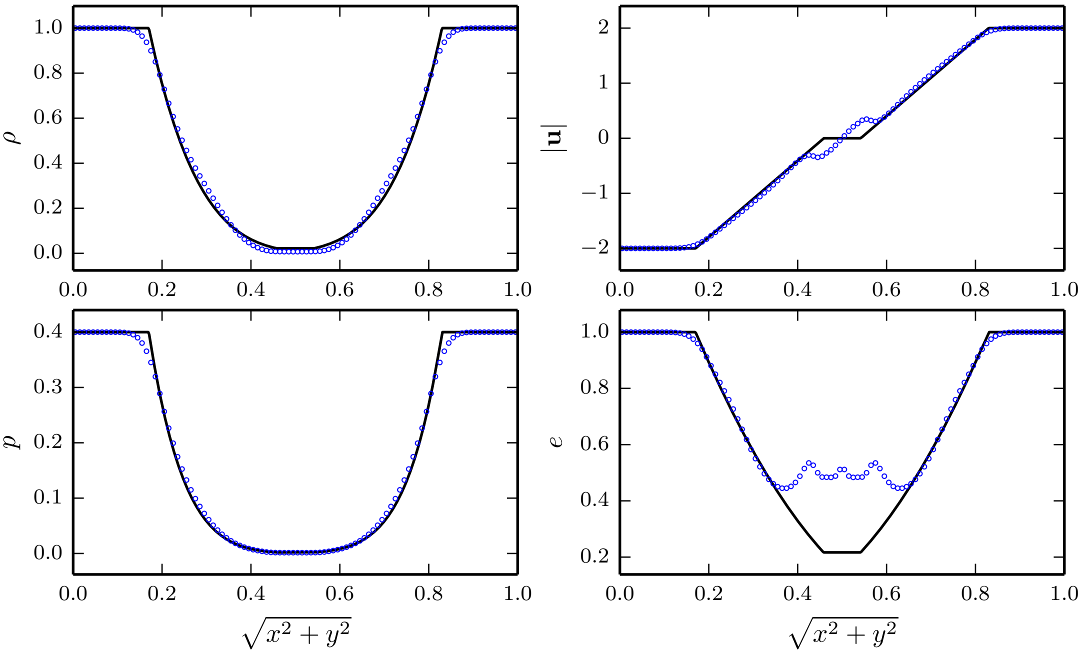

Figure 19: MUSCL TVD evolving Test 2 at $t = 0.12$, along with the exact Riemann solution.

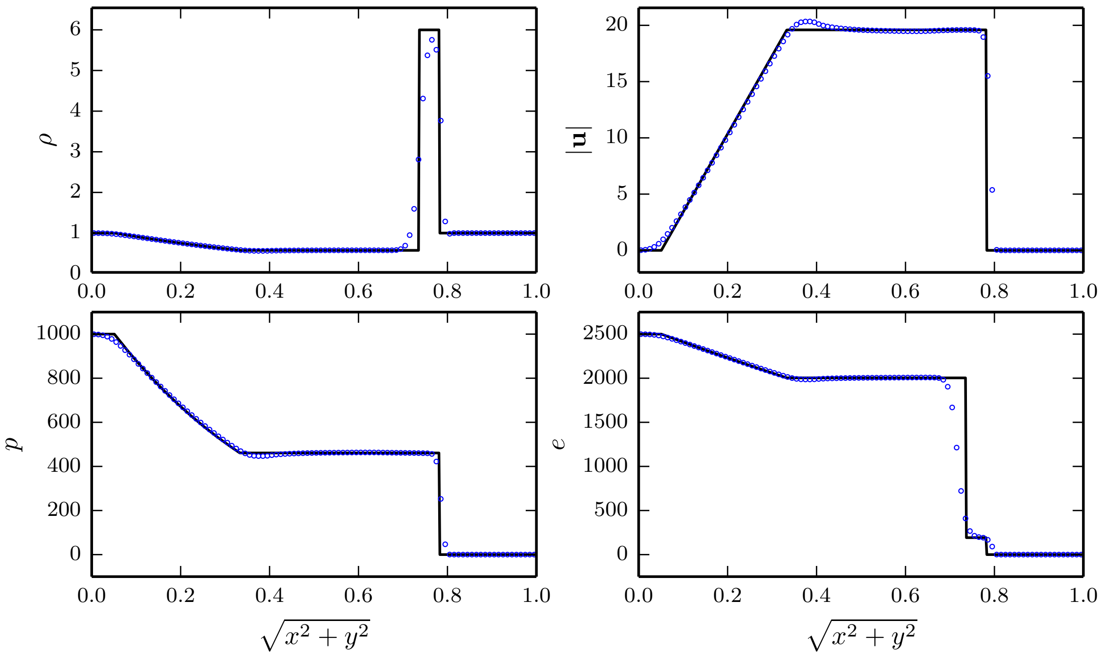

Figure 20: WAF TVD evolving Test 3 at $t = 0.012$, with the exact Riemann solution for reference.

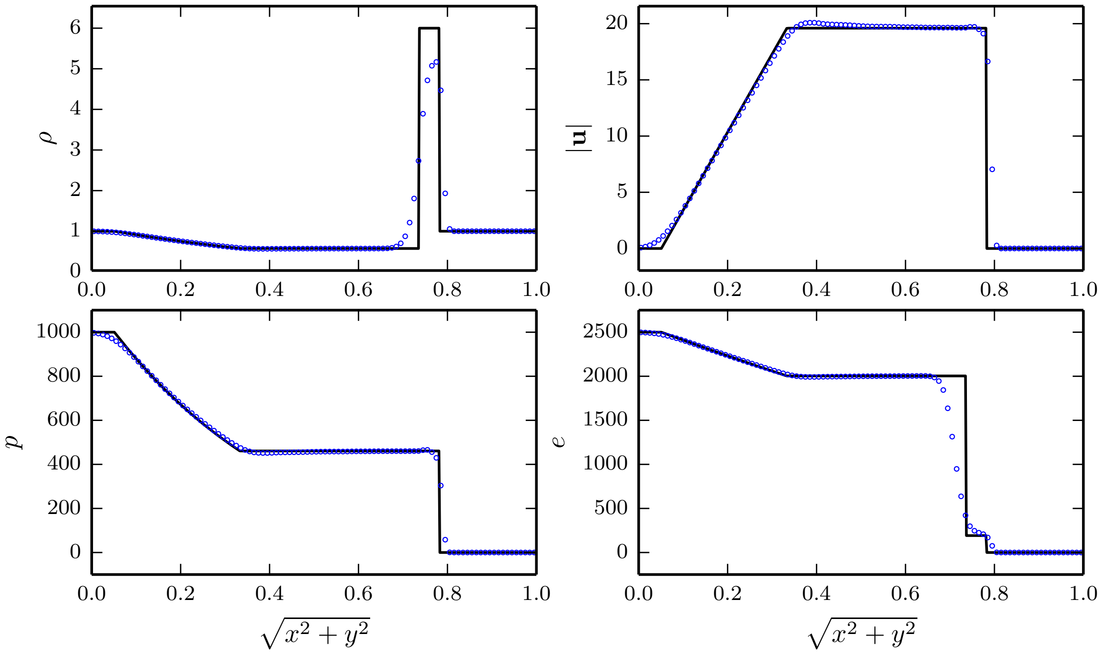

Figure 21: MUSCL TVD evolving Test 3 at $t = 0.012$, along with the exact Riemann solution.

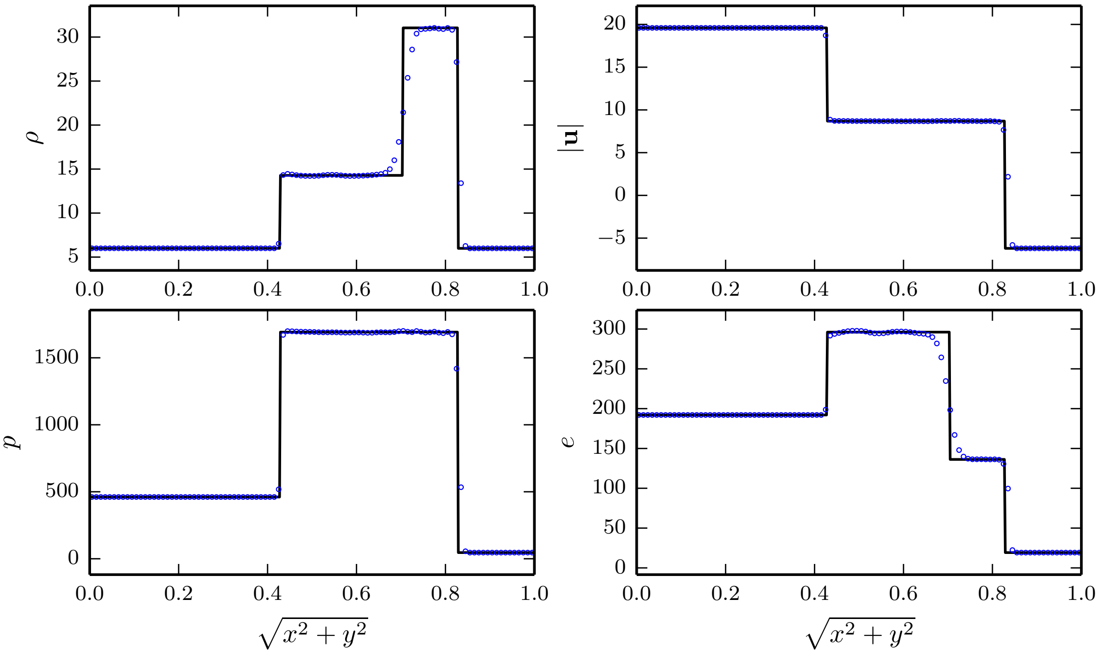

Figure 22: WAF TVD evolving Test 4 at $t = 0.035$, with the exact Riemann solution for reference.

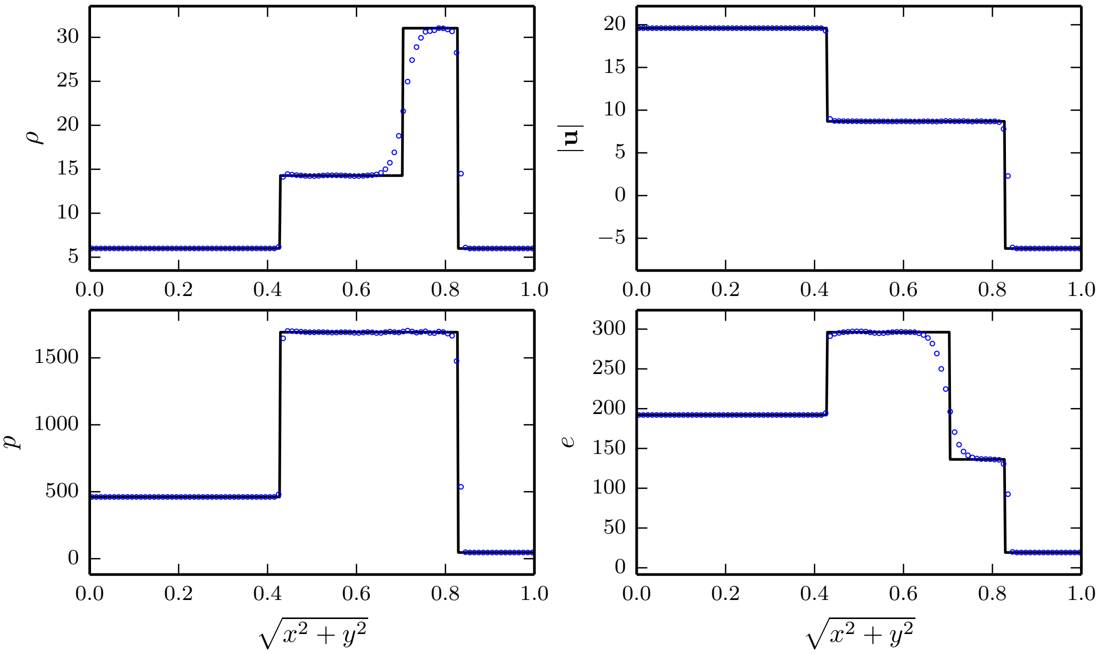

Figure 23: MUSCL TVD evolving Test 4 at $t = 0.035$, along with the exact Riemann solution.

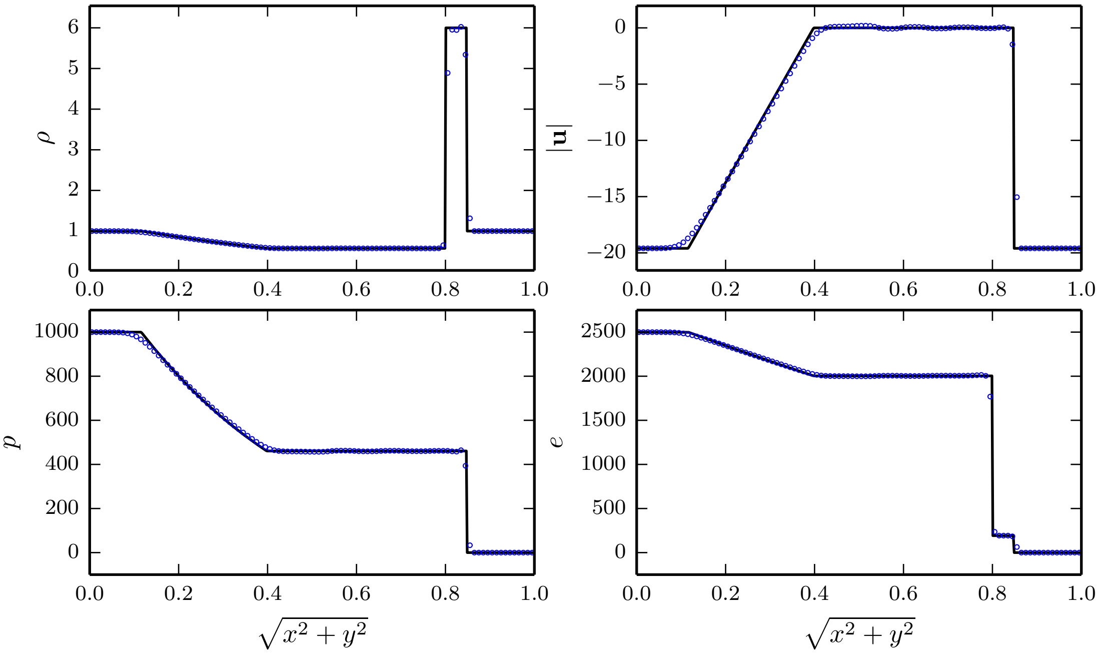

Figure 24: WAF TVD evolving Test 5 at $t = 0.012$, with the exact Riemann solution for reference.

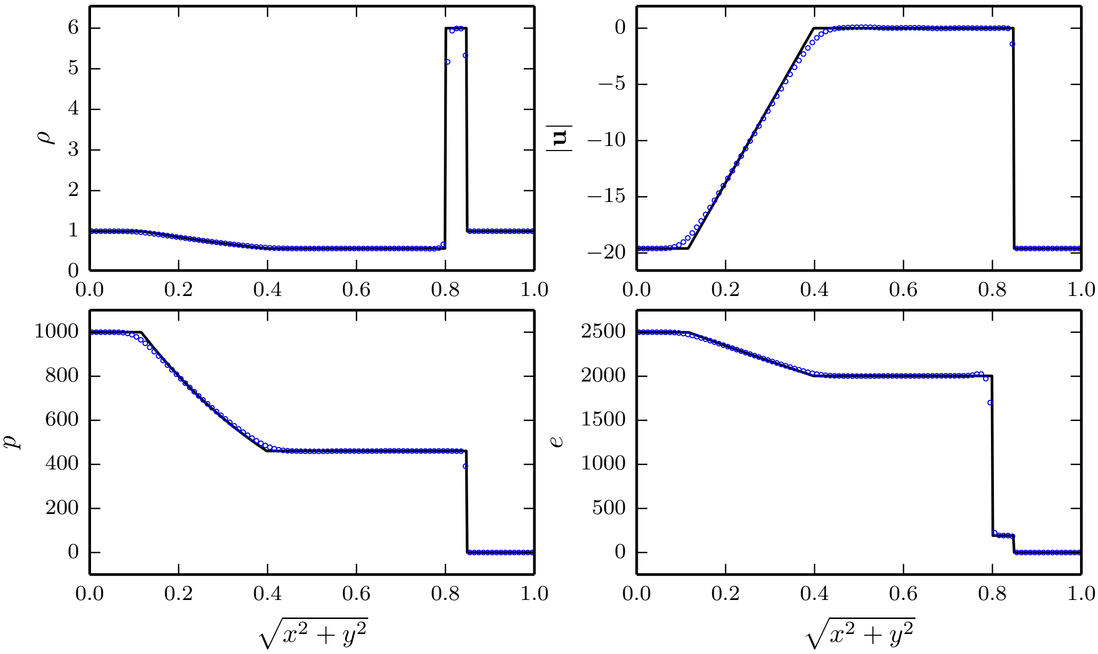

Figure 25: MUSCL TVD evolving Test 5 at $t = 0.012$, along with the exact Riemann solution.

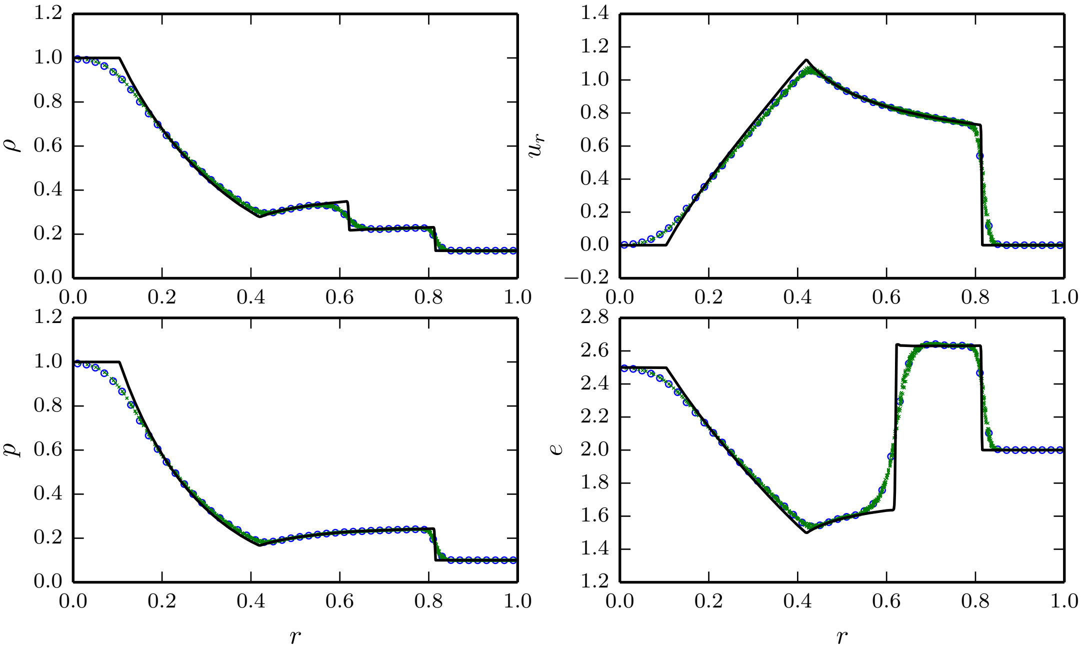

Figure 26: MUSCL TVD explosion test at $t = 0.25$, with the converged 1D solution for reference. Small ×'s mark the projected 2D values, and ○'s indicate the average value in a bin of size Δx.

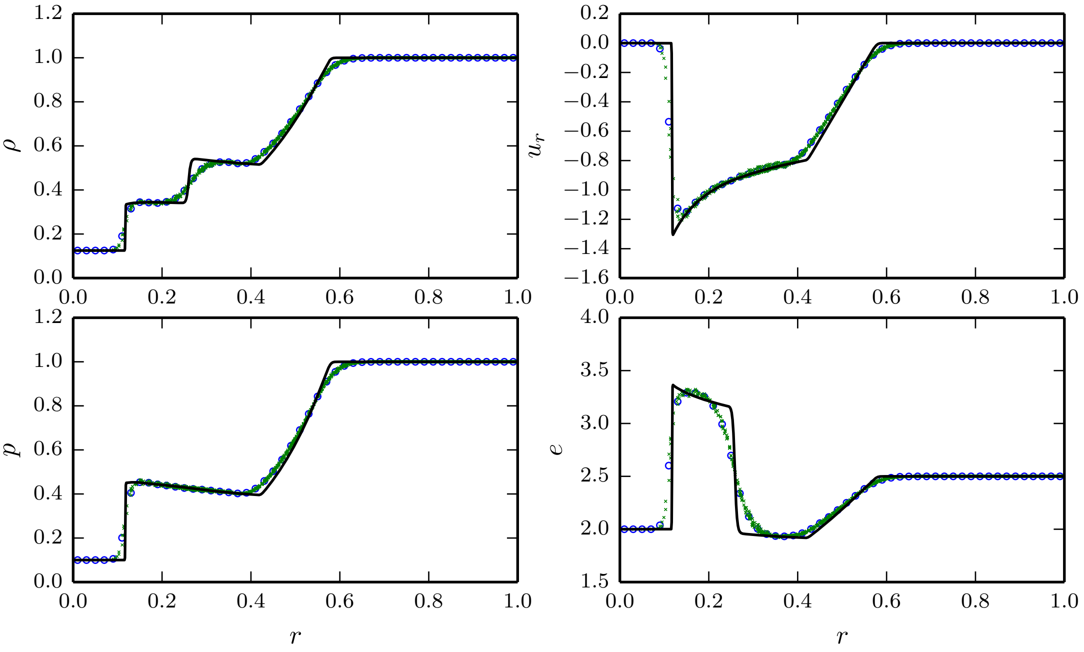

Figure 27: MUSCL TVD implosion test at $t = 0.25$, with the converged 1D solution for reference. Small ×'s mark the projected 2D values, and ○'s indicate the average value in a bin of size Δx.

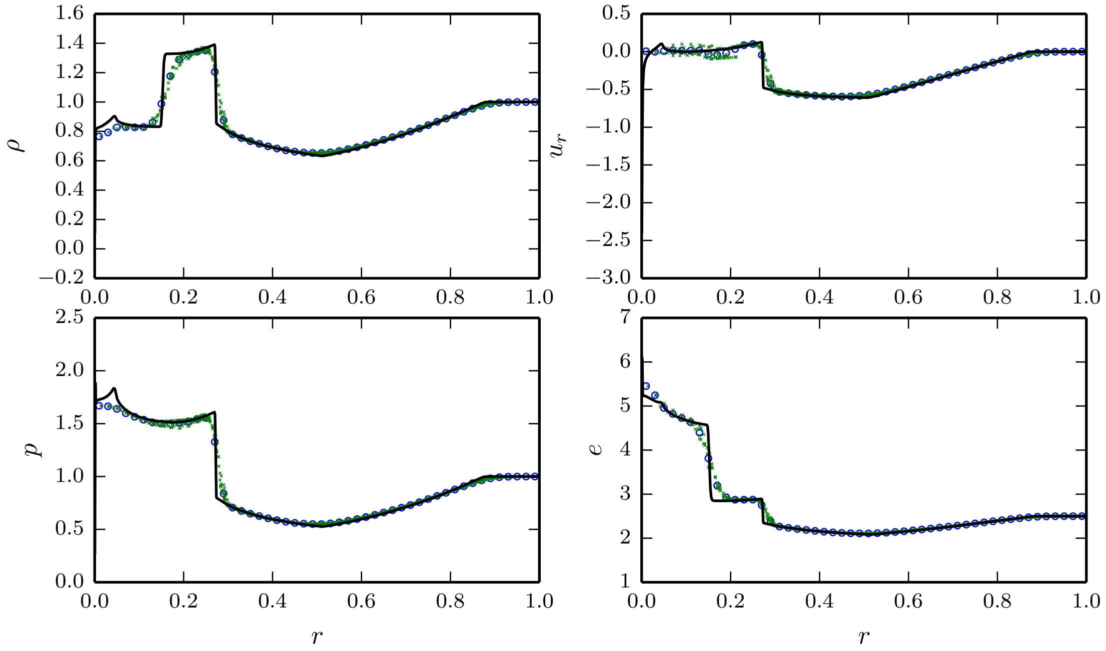

Figure 28: WAF TVD explosion test evolved past shock collision at $r = 0$, shown at $t = 0.4$, with the converged 1D solution for reference.

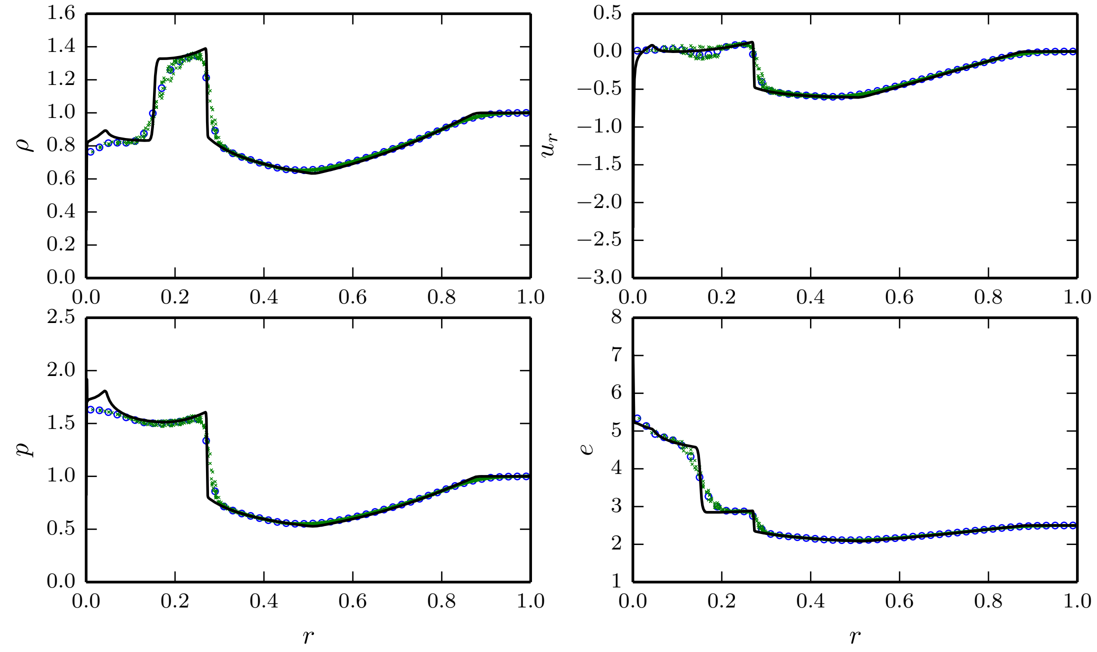

Figure 29: MUSCL TVD implosion test evolved past shock collision at $r = 0$, shown at $t = 0.4$, with the converged 1D solution for reference.