A note on compressive limiting for two-material flows

J. W. Banks

Center for Applied Scientific Computing,
Lawrence Livermore National Laboratory,
Livermore, California, 94551

March 24, 2010

Abstract

In this short note we describe a simple extension to the multi-material shock capturing algorithm presented in [1] that can be used to maintain sharp material interfaces. The method takes the form of an artificial compression which is designed so that the material indicator jumps across only a few cells but which does not excite physical instabilities in the flow. Advantages of the approach include its simplicity and flexibility in that it provides a parameter that effectively determines captured interface thickness.

1 Introduction

One major criticism concerning the use of shock-capturing methods for multi-material flows is the excessive smearing of material interfaces which can result. This is particularly troublesome when the materials on either side of the interface have very different equations of state (EOS). For such cases a smeared interface can mean that exceedingly fine meshes must be used in order to accurately represent flow features. To make matters worse, typical shock capturing methods tend to smear linear discontinuities, such as material interfaces, continually. This means that for a given resolution the material interfaces become more and more diffuse as the simulation progresses [2].

One common approach to dealing with these issues is to reconstruct a material interface and treat the materials on either side independently. This is the approach taken for instance by the ghost-fluid method [3]. Typically the application of boundary conditions across reconstructed interfaces is a difficult process and can lead to lower order approximations near the interface than are used for the bulk of the flow. Furthermore, for instances when shocks and contacts coincide additional errors can be incurred.

An attractive middle ground is to use a capturing method that is constructed to maintain material interfaces at a constant width with a small number of cells. To that end we adopt the model used in [1]

$$\frac{\partial}{\partial t}\mathbf{u} + \frac{\partial}{\partial x_1}\mathbf{f}_1(\mathbf{u}) + \frac{\partial}{\partial x_2}\mathbf{f}_2(\mathbf{u}) = 0, \tag{1}$$

where

$$\mathbf{u} = \begin{bmatrix} \rho \\ \rho u_1 \\ \rho u_2 \\ E \\ \rho \phi_r \end{bmatrix}, \qquad \mathbf{f}_1(\mathbf{u}) = \begin{bmatrix} \rho u_1 \\ \rho u_1{}^2 + p \\ \rho u_1 u_2 \\ u_1(E+p) \\ \rho u_1 \phi_r \end{bmatrix}, \qquad \mathbf{f}_2(\mathbf{u}) = \begin{bmatrix} \rho u_2 \\ \rho u_1 u_2 \\ \rho u_2{}^2 + p \\ u_2(E+p) \\ \rho u_2 \phi_r \end{bmatrix}.$$

The variables have the meanings of ρ the density, u_1 the x_1-velocity, u_2 the x_2-velocity, E the total energy, p the pressure, and ϕ_r a material indicator. Initial conditions set ϕ_r to either 1 or 0 depending on the material and because the model does not allow mixing, exact solutions retain this separation for all time. The total energy is given by $E = \rho e + \frac{1}{2}\rho\left(u_1{}^2 + u_2{}^2\right)$ where $e = e(\rho, p, \phi_r)$ is the specific internal energy, prescribed by an EOS, which for the purposes of this paper this is taken to be the stiffened EOS

$$e(\rho, p, \phi_r) = \frac{p + \gamma(\phi_r)\pi(\phi_r)}{\rho(\gamma(\phi_r) - 1)}$$

where $\gamma(\phi_r)$ and $\pi(\phi_r)$ are the gas constant and stiffening parameter given as functions of the species indicator. For further details refer to [1, 4].

2   Artificial Compression

Beginning with [5], the virtues of artificial compression methods (ACM) were realized. One of the primary motivating factors behind the development of the ACM was the desire for numerical approximations with linear jumps to be able to achieve steady state profiles. Traditional shock capturing schemes continually smear such jumps and so it is generally not possible to time step to steady state solutions. With the ACM those jumps remain sharp and steady states can be achieved. However, the disadvantages associated with compressive schemes can be substantial. Firstly they are often only first order accurate even for smooth flows. Secondly they can produce erroneous flow features such as artificial steepening of smooth profiles which can lead to unrealistically high growth rates for physical instabilities. In general we agree that these are valid concerns, but the current situation is quite special in a number of ways. To begin with the species indicator only takes discrete values of 1 or 0 in the exact solution. This means that there is never a smooth transition that could potentially steepen. Also because we apply compression to the evolution equation for the species indicator only, instability growth requiring a coupling to the mass and momentum equations is not a concern. Finally, the L_1 convergence rate for non-compressive p^{th} order methods for linear jumps is typically $O(\Delta x^{p/(p+1)})$ but compressive schemes can converge as $O(\Delta x)$ [2]. For these reasons, carefully constructed first order compressive schemes can produce superior results when applied to advected linear discontinuities such as material interfaces.

The overall numerical approach to solving (1) follows the discussion in [1]. An overlapping grid approach is used to treat geometric complexities, adaptive mesh refinement is used to locally increase grid resolution where needed, and an energy correction is used to eliminate numerical errors associated with captured material interfaces. This energy correction is non-zero only in the vicinity of material interfaces and so the convergence to weak shock solutions is not altered. Because our goal is to maintain sharp material interfaces, the compressive scheme is applied only to the final equation in (1). As was shown in [2], a rather simple method that has the compressive character we desire uses two times the usual MinMod limiter, also known as the double MinMod limiter. The method from [1] uses the standard MinMod limiter and so the extension to use the double MinMod limiter for the last equation of (1) is rather trivial. In the interest of space, only the critical changes to the scheme in [1] are presented here and the reader is referred to [1] for further details. At some point in the algorithm the slope correction step is encountered to give high-resolution. As in [1] this is written

$$W_L = W_{i,j}^n + \frac{1}{2} R_{1,i,j}^n \left(I - \frac{\Delta t}{\Delta r_1} \max\{0, \Lambda_{1,i,j}^n\} \right) \alpha_{i,j}^n - \frac{\Delta t}{2\Delta r_2} R_{2,i,j}^n \Lambda_{2,i,j}^n \beta_{i,j}^n \,,$$

$$W_R = W_{i+1,j}^n - \frac{1}{2} R_{1,i+1,j}^n \left(I + \frac{\Delta t}{\Delta r_1} \min\{0, \Lambda_{1,i+1,j}^n\} \right) \alpha_{i+1,j}^n - \frac{\Delta t}{2\Delta r_2} R_{2,i+1,j}^n \Lambda_{2,i+1,j}^n \beta_{i+1,j}^n \,,$$

(2)

where W_L and W_R are the left and right states for an inter-cell Riemann problem, $R_k \Lambda_k R_k^{-1}$ is the eigen-decomposition of the flux Jacobian in the k-direction in index space, i and j are indices into a computational grid, n indicates the time step, Δr_k and Δt are spatial and temporal grid sizes, and α and β are given by

$$\alpha_{i,j}^n = \text{MinMod}\left(\left(R_1^{-1}\right)_{i,j}^n \left(W_{i,j}^n - W_{i-1,j}^n\right), \left(R_1^{-1}\right)_{i,j}^n \left(W_{i+1,j}^n - W_{i,j}^n\right) \right),$$

$$\beta_{i,j}^n = \text{MinMod}\left(\left(R_2^{-1}\right)_{i,j}^n \left(W_{i,j}^n - W_{i,j-1}^n\right), \left(R_2^{-1}\right)_{i,j}^n \left(W_{i,j+1}^n - W_{i,j}^n\right) \right).$$

(3)

Equations (2) and (3) are exactly equations (33) and (34) in [1]. In order to implement our compressive
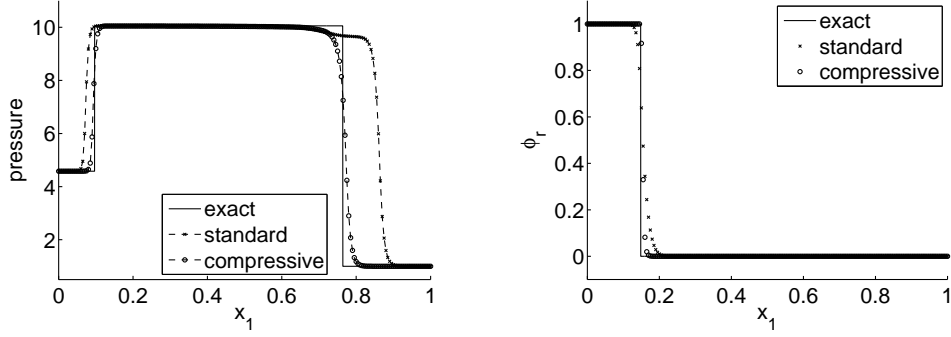
Figure 1: Pressure and species indicator ϕ_r at $t = 0.25$. The exact solution and numerical approximations with 200 cells using the standard scheme ($\kappa = 1$) as well as the compressive scheme ($\kappa = 2$) are shown.

scheme, equation (2) is altered slightly to

$$
W_L = W_{i,j}^n + \frac{1}{2} T R_{1,i,j}^n \left(I - \frac{\Delta t}{\Delta r_1} \max\{0, \Lambda_{1,i,j}^n\} \right) \alpha_{i,j}^n - \frac{\Delta t}{2\Delta r_2} T R_{2,i,j}^n \Lambda_{2,i,j}^n \beta_{i,j}^n \,,
$$

$$
W_R = W_{i+1,j}^n - \frac{1}{2} T R_{1,i+1,j}^n \left(I + \frac{\Delta t}{\Delta r_1} \min\{0, \Lambda_{1,i+1,j}^n\} \right) \alpha_{i+1,j}^n - \frac{\Delta t}{2\Delta r_2} T R_{2,i+1,j}^n \Lambda_{2,i+1,j}^n \beta_{i+1,j}^n \,,
$$

(4)

where

$$
T = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & \kappa \end{bmatrix}
$$

and κ is a parameter which dictates the treatment of the final equation of (1). For $\kappa = 0$ the first order scheme is used, for $\kappa = 1$ the standard MinMod scheme is recovered, and for $\kappa = 2$ the double MinMod scheme is achieved. Note that $\kappa = 2$ corresponds to the TVD limit in the Sweby diagram [6] and so great care must be exercised when using larger values. For $\kappa > 1$, the scheme is compressive and material interfaces will be maintained across a fixed number of cells. For instance, the choice of $\kappa = 2$ typically results in interfaces captured across two to four cells for all time. An important subtle distinction to note is that this formulation applies compressive limiting to the final equation, rather than applying the compressive limiter to the final component of the state vector. This seemingly small difference is in fact quite significant as will be illustrated via numerical example in section 3.

3   Numerical Examples

We present a few numerical examples to illustrate the effectiveness of the method when applied to two-material flows. The first is a 1-D flow with substantially different equations of state on either side of the material interface. The initial condition has uniform density $\rho = 1$ and uniform pressure $p = 1$. To the left of the origin a velocity of $u_1 = 1.5$ is prescribed and to the right $u_1 = -1.5$. For $|x_1| < 0.5$ the EOS roughly describes air with $\gamma = 1.4$ and $\pi = 0$. For $|x_1| \geq 0.5$ a very stiff gas is used with $\gamma = 4.4$ and $\pi = 50$. The flow is symmetric about the origin and the solution is computed only for $x_1 \in [0, 1]$ (the other half determined by symmetry if desired). The exact solution can be determined as the solution of three Riemann problems. Initially a right moving shock forms at the origin. This shock interacts with the material interface at $t \approx 0.2095$ producing a transmitted shock moving right, a reflected shock moving left, and a left moving material interface.

Computed numerical approximations at $t = 0.25$ using 200 uniformly spaced cells are shown in figure 1. The plot of ϕ_r shows the relative sharpness (≈ 3 cells) of the material interface when $\kappa = 2$ is used as
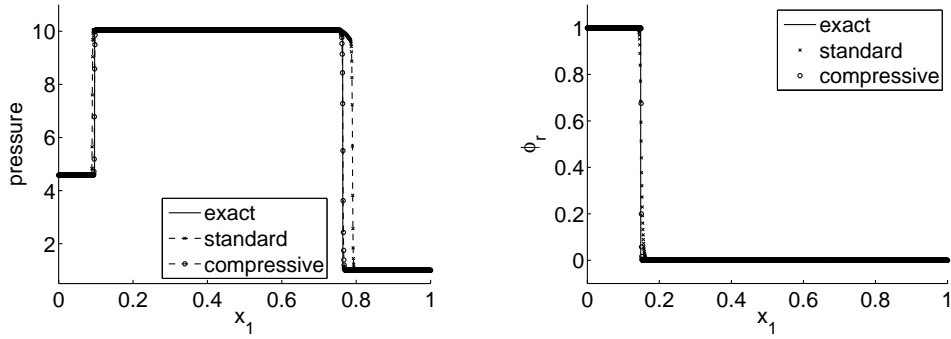
3

Figure 2: Pressure and species indicator ϕ_r at $t = 0.25$. The exact solution and numerical approximations with 1600 cells using the standard scheme ($\kappa = 1$) as well as the compressive scheme ($\kappa = 2$) are shown.
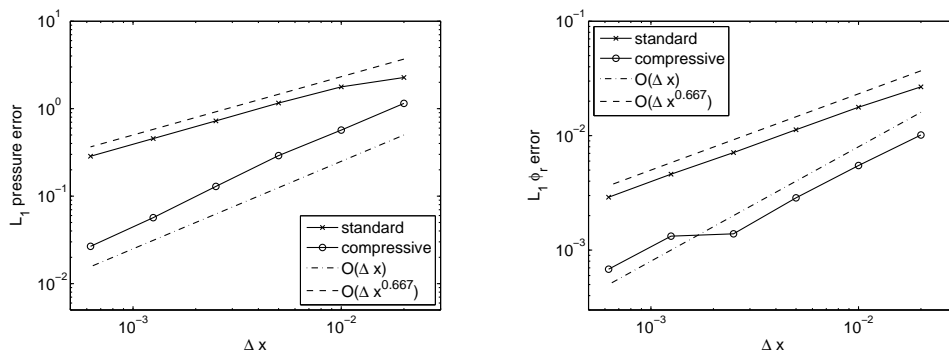

Figure 3: Convergence of the L_1 error for pressure and species indicator for the 1-D test problem. Errors are shown for both the compressive and standard schemes. Reference lines indicating $O(\Delta x)$ and $O(\Delta x^{2/3})$ are also included.

compared to the usual scheme with $\kappa = 1$. The pressure shows that the standard scheme has significantly misplaced the reflected and transmitted shocks, a result of overestimation of the sound speed in the smeared interface. Because the compressive scheme maintains a constant width interface, this type of error does not accumulate and the shocks are accurately located. Both methods will eventually converge to the exact solution if sufficiently small computational cells are used, but the rate for the standard method ($\kappa = 1$) is extremely slow because the captured material interface continually grows as a function of the number of time steps. Figure 2 shows the computed pressure and species indicator using 1600 cells. The compressive scheme still captures the interface with approximately 3 cells but the standard scheme now contains more than 20 cells. As a result, the shocks for the standard scheme are still significantly misplaced. Figure 3 shows a convergence study of the L_1 error for both pressure and species indicator over a range of resolutions corresponding to 50 through 1600 computational cells. In the figures, reference lines corresponding to the convergence rates $O(\Delta x)$ and $O(\Delta x^{2/3})$ are included. It is seen that the convergence rates for both pressure and species indicator correspond to $\approx O(\Delta x)$ for the compressive scheme while for the traditional scheme the convergence rates are $\approx O(\Delta x^{2/3})$.

Our second example is the impulsively driven cylinder presented in [1] and [7]. Here a cylinder of radius 0.8 containing air on the left ($\gamma = 1.4$, $\rho = 1.0$) and helium on the right ($\gamma = 1.67$, $\rho = 0.138$) is impulsively set into motion with speed 1. The gas has an initially uniform pressure $p = 1$. Further details of the setup, can be found in [1] and [7]. One significant concern when using compressive schemes is the excitation of physical instabilities such as the Kelvin–Helmholtz or Richtmyer–Meshkov instabilities. The flow generated by this test case exhibits both of these and the effects of our artificial compression can be
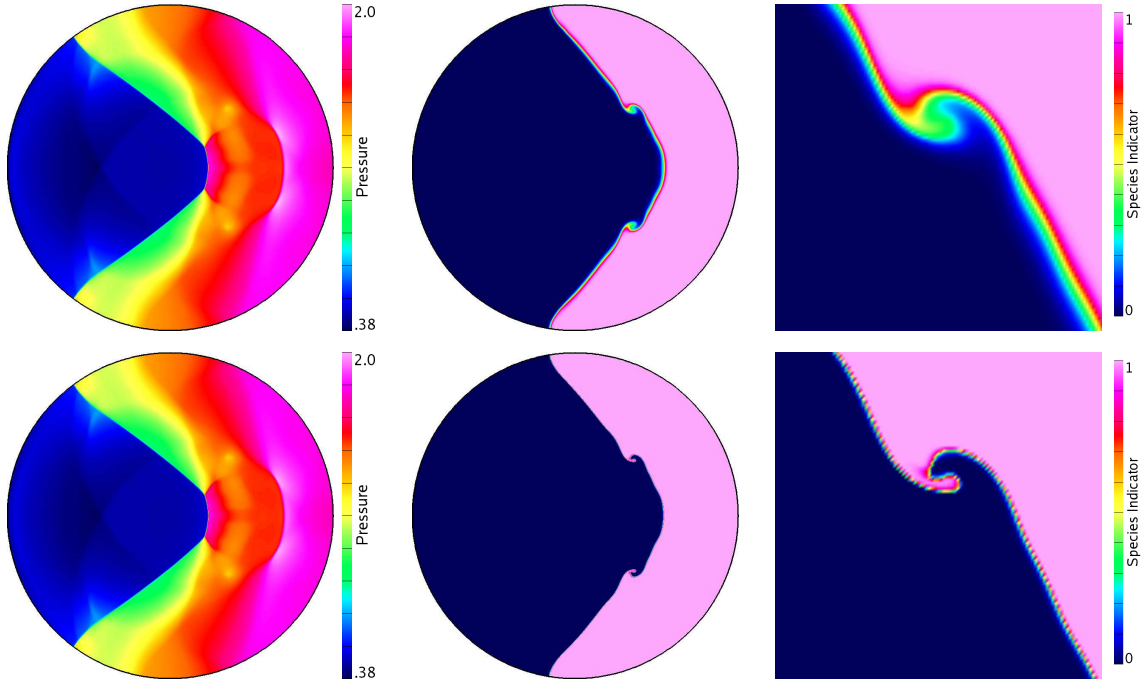
4

Figure 4: Pressure (left), species indicator (center), and a zoom of the species indicator (right) for impulsively driven cylinder at $t = 1.0$. On top are results from the standard MinMod limiter ($\kappa = 1$) and on bottom from the compressive scheme ($\kappa = 2$).

judged. Figure 4 shows the pressure and species indicator for both the standard MinMod scheme ($\kappa = 1$) and the compressive scheme ($\kappa = 2$) at the middle resolution presented in [1] with $\Delta x \approx 1/400$. Clearly the material interface is sharper for $\kappa = 2$ and occupies ≈ 3 cells, while for $\kappa = 1$ it is ≈ 10 cells wide. More to the point, there are no significant differences in the structure of the computed solution between the compressive and non-compressive schemes which indicates that instabilities have not been excited as a byproduct of sharp interface maintenance. To contrast this, figure 5 shows the pressure and species indicator obtained by a simulation which uses the double minmod limiter for the last component of α and β in equation (2) rather than the proposed scheme in equation (4). Clearly the results are unsatisfactory and illustrate both significant growth of instabilities as well as large amounts of numerical noise which have been produced through the use of a poorly constructed compressive scheme.

The simulations presented here use $\kappa = 1$ for the usual scheme and $\kappa = 2$ for the compressive scheme. However, we have found that for more complicated equations of state, such as the JWL EOS used for reactive flow modeling in [8], an intermediate value may be more appropriate. Typically $\kappa = 2$ is preferred, but this can cause convergence problems for the mixture EOS calculator. For such cases κ is reduced and the interface widens somewhat. For $\kappa > 1$ a steady finite width interface will be achieved and κ can be used to effectively determine the width, in cells, of the steady state interface.

4   Conclusions

In this brief note we have shown how artificially compressive schemes can be used to maintain sharp material interfaces in two-material flows. The typical detrimental effects of such schemes are shown not to arise here because of the special nature of material interfaces. A 1-D flow showed how artificially smeared interfaces can be detrimental to convergence rates and that an appropriate compressive scheme can mitigate this behavior. A 2-D example demonstrated how traditional schemes tend to mix material while compressive schemes are able to keep materials largely separate. At the same time, our compressive scheme was shown not alter
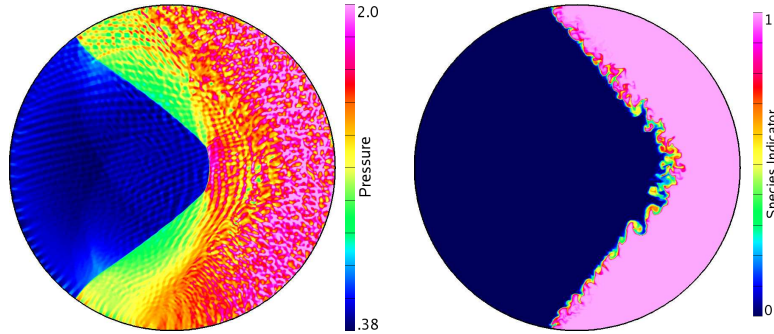
Figure 5: Pressure (left) and species indicator (center) at $t = 1.0$ computed using the double minmod limiter for the last component of α and β in equation (2) rather than the proposed scheme in equation (4).

the excitation of physical instabilities, or to produce erroneous numerical noise. One might be tempted to extend the compressive scheme to sharpen density jumps but extreme care should be exercised here. It is our experience that this can be done and density jumps will remain sharp, but physical instabilities, if present, quickly dominate the flow. Methods that retain a sharp capture of material interfaces for all quantities but which do not result in excessive instability growth are the subject of ongoing work.

Acknowledgements

References

[1] J. W. Banks, D. W. Schwendeman, A. K. Kapila, W. D. Henshaw, A high-resolution Godunov method for compressible multi-material flow on overlapping grids, J. Comput. Phys. 223 (2007) 262–297.

[2] J. W. Banks, T. Aslam, W. J. Rider, On sub-linear convergence for linearly degenerate waves in capturing schemes, J. Comput. Phys. 227 (14) (2008) 6985–7002.

[3] R. P. Fedkiw, T. Aslam, B. Merriman, S. Osher, A non-oscillatory Eulerian approach to interfaces in multimaterial flows (the ghost fluid method), J. Comput. Phys. 152 (1999) 457–492.

[4] R. Saurel, R. Abgrall, A simple method for compressible multifluid flows, SIAM J. Sci. Comput. 21 (3) (1999) 1115–1145.

[5] A. Harten, The artificial compression method for computation of shocks and contact discontinuities. I. single conservation laws, Commun. Pur. Appl. Math. 30 (5) (1977) 611–638.

[6] P. K. Sweby, High resolution schemes using flux limiter for hyperbolic conservation laws, SIAM J. Numer. Anal. 21 (1984) 995–1011.

[7] H. Zheng, C. Shu, Y. Chew, An object-oriented and quadrilateral-mesh based solution adaptive algorithm for compressible multi-fluid flows, J. Comput. Phys. 227 (14) (2008) 6895–6921.

[8] J. W. Banks, W. D. Henshaw, D. W. Schwendeman, A. K. Kapila, A study of detonation propagation and diffraction with compliant confinement, Combust. Theory and Modelling 12 (4) (2008) 769–808.