INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS, VOL. 18, 555–574 (1994)

A CONTRIBUTION TO THE GREAT RIEMANN SOLVER DEBATE

JAMES J. QUIRK

Institute for Computer Applications in Science and Engineering, NASA Langley Research Center, Hampton, VA 23681, U.S.A.

SUMMARY

The aims of this paper are threefold: to increase the level of awareness within the shock-capturing community of the fact that many Godunov-type methods contain subtle flaws that can cause spurious solutions to be computed; to identify one mechanism that might thwart attempts to produce very-high-resolution simulations; and to proffer a simple strategy for overcoming the specific failings of individual Riemann solvers.

KEY WORDS Riemann solvers Shock waves Numerical artifacts

1. INTRODUCTION

Over recent years a plethora of shock-capturing schemes have been developed for the Euler equations of gas dynamics. During this period it has emerged that one of the more successful strategies for designing a shock-capturing scheme is to follow Godunov's lead[1] and utilize a classic initial value problem known as a Riemann problem.[2] Godunov assumed that a flow solution could be represented by a series of piecewise constant states. Thus the numerical representation closely approximates the true solution near discontinuities, and regions of smooth flow are reasonably well approximated by a series of step functions. He evolved this discretized flow solution by considering the non-linear interactions between its component states. Viewed in isolation, each pair of neighbouring states constitutes a Riemann problem, the solution to which may be found exactly.[3] The results from these separate Riemann problems may then be averaged so as to advance the flow solution through some time increment. Because it mimics much of the relevant physics, Godunov's scheme results in an accurate and well-behaved treatment of shock waves.

Although it provides the bedrock upon which most modern schemes are built, in its original form Godunov's method is of limited use. Firstly, the scheme proves to be highly dissipative and so requires an inordinately fine mesh to resolve complex shock-on-shock interactions. Secondly, since a Riemann problem has no closed form solution and can only be solved by some iterative method, Godunov's scheme is significantly more expensive than schemes which employ ordinary finite difference operators.

One of the first people to address this second shortcoming was Roe.[1] He argued that since the Riemann problems associated with Godunov's method arise from an approximation of the data, it might be sufficient to find only approximate solutions to these Riemann problems, provided that they still describe important, non-linear behaviour—his motivation being that approximate solutions can be computed much more cheaply than exact solutions. Thus the

CCC 0271–2091/94/060555–20 *Received 2 November 1992*
© 1994 by John Wiley & Sons, Ltd.

industry of designing approximate Riemann solvers was born.[4–7] Now, whilst Godunov-type schemes are often held up to be models of robustness, they can on occasions fail quite spectacularly. For example, when computing shock reflection problems, Roe's method can sometimes go awry by admitting solutions for which the Mach stem is inexplicably kinked. The existence of such failings partly explains why no consensus of opinion has been reached concerning the ideal Riemann solver. Whenever a new failing is unearthed, it adds fuel to the great Riemann solver debate: method X is better than method Y because of reasons A, B and C. It is our contention that for all the current crop of Riemann solvers at least one set of circumstances may be found for which any one solver is found wanting; some failings are just more visible than others.

In Section 2 we catalogue a number of situations in which anomalous behaviour is known to occur. This catalogue should serve to increase the general level of awareness within the shock-capturing community of the current limitations of Riemann solver technology. At present this awareness is not all that it should be: there are many instances in the literature where suspect numerical results are presented with either little or no adverse comment. We believe that one of the failings listed in our catalogue has hitherto gone unreported. In Section 3 we proffer a diagnosis of the mechanism which causes this new failing.

At this juncture it should be noted that any foibles that a specific Riemann solver might have may usually be controlled by the judicious use of a small amount of artificial dissipation. Indeed, it is worth pointing out that Colella and Woodward's PPM scheme,[8] which has proved itself to be more robust than most Godunov-type schemes, does in fact employ an elaborate artificial dissipation mechanism to supplement the dissipation provided via upwinding. As will be described in Section 4, we favour a strategy whereby the weaknesses of any one solver are overcome by combining it with one or more complementary solvers. The main advantages of this approach over that of adding artificial dissipation are twofold. Firstly, it does not degrade the resolution of the base Riemann solver; it is possible to control certain instabilities by changing the flavour of the dissipation mechanism rather than increasing the absolute level of dissipation. Secondly, it does not necessitate a host of tunable parameters and so this synergetic strategy does not negate the principal advantage of Godunov-type schemes over other shock-capturing methods. Of course, we are left with the difficulty of deciding when to use one Riemann solver in preference to another; however, we present a number of computations which suggest that this difficulty is not particularly bothersome.

Finally, in Section 5 we list the main conclusions that we have drawn from this work. Note that in this paper we do not address the first shortcoming of Godunov's method, namely its low resolution. Following van Leer,[9] it is assumed that a high-order extension to a first-order method can always be achieved by preprocessing the data supplied to the Riemann solver.

2. A CATALOGUE OF FAILINGS

We now present several instances where various Riemann solvers are known to give unreliable results. While most of these cases will be known to the aficionado, we believe that one of the cases which we are about to describe has hitherto gone unreported in the literature. Although our catalogue is not exhaustive, we hope that it might save some investigators from the harrowing experience of spending weeks or even months searching for coding errors that simply do not exist.

All the computational results shown in this section are for first-order schemes. Since such methods have low resolution, our calculations employed relatively fine meshes; for clarity, grids are drawn using every other fourth grid line.

2.1. *Expansion shocks*

By far the most widely investigated failing is that some Riemann solvers do not satisfy an 'entropy condition', such schemes can admit non-physical solutions such as expansion shocks. Osher[4] has found a general condition for a scheme to be entropy-satisfying when applied to scalar equations and he designates such schemes E-schemes. At present, however, any extension to a system of equations contains a large amount of empiricism and must therefore remain suspect. Indeed, Godunov's method is classified as an E-scheme but, as observed by Woodward and Colella,[10] it can give rise to nearly discontinuous expansion fans near sonic points. The density contours shown in Figure 1 illustrate this deficiency of Godunov's method quite clearly. These results are taken from the diffraction of a strong shock wave, $M_s = 5{\cdot}09$ with $\gamma = 1{\cdot}4$, around a 90° corner.

In its basic form Roe's scheme is another solver that admits expansion shocks; however, several fixes have been proffered which cure Roe's scheme of this particular affliction.[11–13] Such fixes are typical of the way in which Riemann solver deficiencies have been countered up to now. Whilst this strategy has proved reasonably successful, it has a number of drawbacks. Sometimes a fix uses a parameter which must be retuned between problems and hence one of the major advantages of Riemann-based schemes over, say, artificial dissipation schemes is lost. Alternatively, a scheme may require more than one fix and it may be unclear how the different fixes interact with one another.

2.2. *Negative internal energies*

Another situation in which some Riemann solvers are found wanting occurs whenever the dominant energy mode is kinetic rather than thermal. For such solvers the kinetic energy computed from a numerical approximation to the conservation laws of mass and momentum can exceed the total energy computed via an approximation to the conservation law of energy. Thus they can yield negative internal energies and hence negative pressures, which cause the scheme to fail. Einfeldt *et al.*[12] call any scheme which can be guaranteed not to yield negative

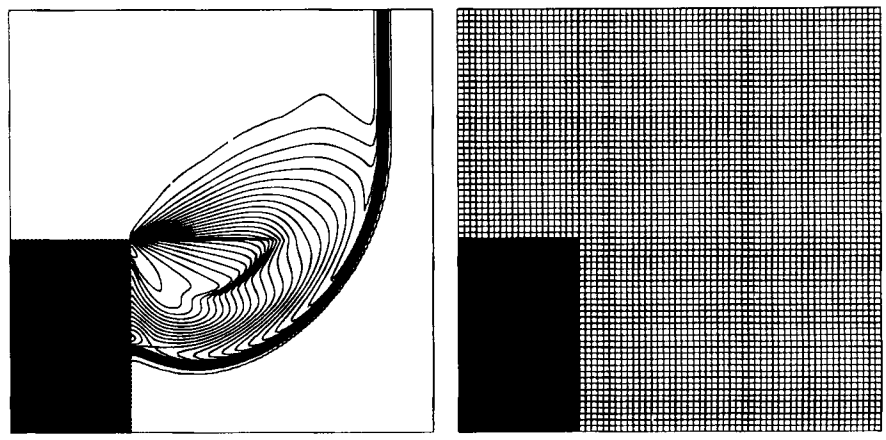

(a) (b)

Figure 1. A strong shock diffracting around a corner gives rise to an expansion shock: (a) density contours; (b) computational grid

pressures 'positively conservative'. They have shown that while Godunov's scheme is 'positively conservative', the reverse is true for any Godunov-type scheme based on a linearized Riemann solver. Indeed, the basic form of Roe's scheme is unable to cope with the test problem shown in Figure 1; the strength of the diffracting shock is sufficient to cause a negative pressure to be computed near the apex of the corner. Roe's scheme may be made 'positively conservative' by modifying its wave speeds; in essence, the scheme is made more dissipative by increasing the spread in velocity between the two acoustic waves.[12]

2.3. Slowly moving shocks

Since shock-capturing schemes do not resolve the internal structure of a shock wave, no physical significance can be attached to the discrete shock structure produced by a numerical scheme. Methods are built upon the premise that shock profiles are monotone; the precise structure comes about as a matter of course and is not preordained. Unfortunately, Roberts has shown that the nature of the shock structure produced by a particular scheme can have a large bearing on how well the scheme copes with slowly moving shock waves.[14] Godunov-type methods fare quite badly in this respect: as the shock moves relative to the mesh, the shock profile flexes, perturbing the supposedly passive characteristic fields as it does so.

Figure 2 shows a snapshot of the shock profile produced by Einfeldt's HLLE scheme[6] taken from the simulation of a shock wave which is moving slowly from left to right; the pre-shock state (density, velocity, pressure) is (1, −3·44, 1) and the post-shock state is (3·86, −0·81, 10·33). Note that for a Courant number of one it takes 50 time steps for this shock to traverse one mesh cell. The low-frequency perturbations observed in this figure are produced to a greater or lesser extent by any scheme which attempts to 'recognize' a shock wave. For fast moving shocks the post-shock noise will be of a much shorter wavelength than is the case here and will be effectively damped by the dissipation of the scheme. Roberts reports that Osher's scheme[15] does not produce low-frequency noise for slowly moving shocks, since it never connects two adjacent states by a shock, and he concludes that there may be advantages to using flux formulae that do not recognize the analytic shock jump conditions.

Another situation where the perturbation of a shock from its preferred profile results in perturbations on the passive characteristic fields occurs whenever a shock crosses a discontinuity in mesh spacing.[16] However, in this case sizeable perturbations may occur whatever the speed of the shock.

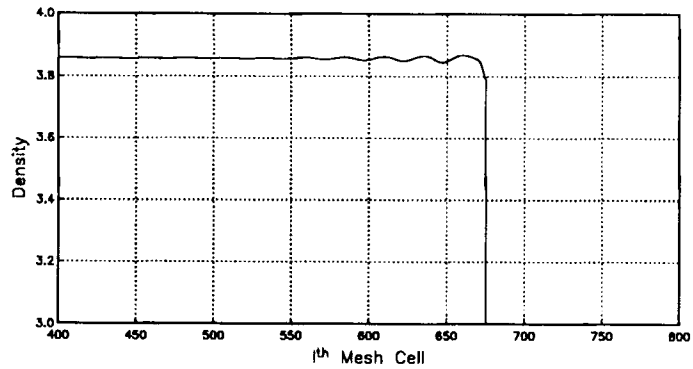

Figure 2. Low-frequency, post-shock oscillations occur for slowly moving shock waves

2.4. *The carbuncle phenomenon*

Several authors have by now reported a failing of Roe's scheme which has been dubbed the 'carbuncle phenomenon.'[17–19] For steady state, blunt body calculations Roe's scheme sometimes admits a spurious solution in which a protuberance grows ahead of the bow shock along the stagnation line. It appears that this effect is more pronounced the more closely the grid is aligned to the bow shock. Also, a carbuncle is more likely to appear for high-Mach-number flows than for low-Mach-number flows. Figure 3 shows such a spurious solution; here the freestream Mach number was taken to be 10. Note that along the stagnation line the bow shock is almost perfectly aligned with the grid. Consequently, parallel to the shock Roe's scheme will not add any dissipation via the contact and shear waves to counteract perturbations that appear through the acoustic waves; this appears to be a recurring theme whenever Roe's method fails. It is interesting to note that if Harten's entropy fix[13] is applied to the contact and shear waves, any shortcoming of Roe's scheme is invariably cured. However, there is no justification, either physical or mathematical, for applying this fix to these waves; it is just a convenient method for introducing an amount of artificial dissipation into the scheme.

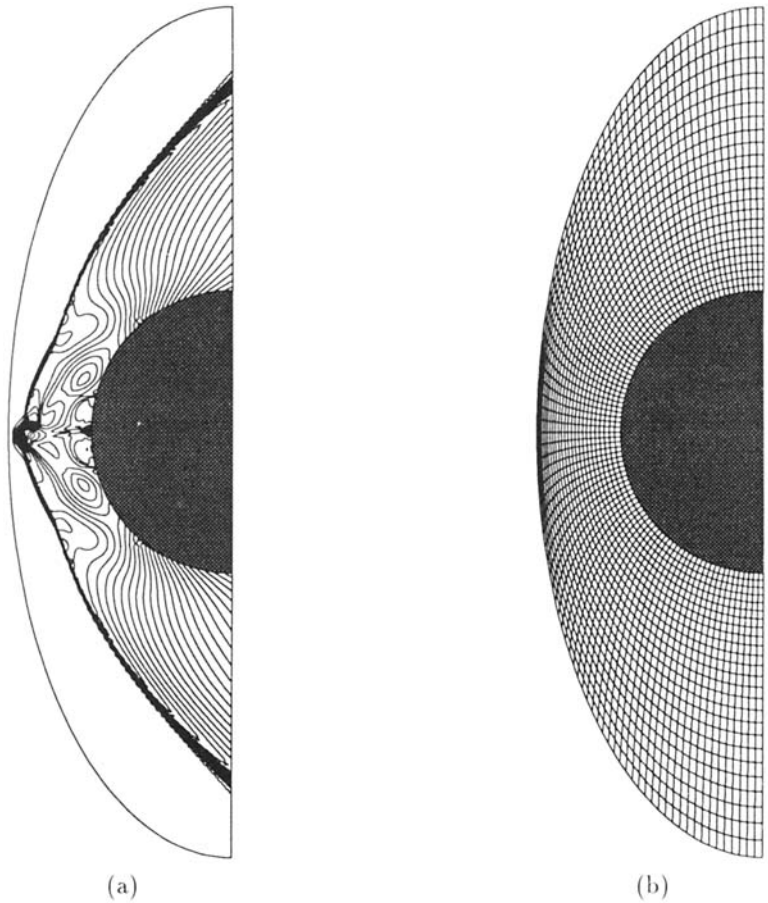

Figure 3. The *carbuncle phenomenon*: (a) density contours; (b) computational grid

2.5. *Kinked Mach stems*

During the course of developing a mesh adaptation scheme, we encountered a failing of Roe's scheme which is not dissimilar to the 'carbuncle phenomenon'.[16] When the reflection of a plane shock wave from a ramp lies in the double-Mach reflection regime, the principal Mach stem is sometimes inexplicably kinked. Figure 4 shows a snapshot of the pressure contours taken during the reflection of a plane shock, $M_s = 5{\cdot}5$ with $\gamma = 1{\cdot}4$, from a 30° ramp. The principal Mach stem is so severely kinked that it has given rise to a spurious triple point. Similarly strange results have been produced by Sawada[20] and Itoh *et al.*[21] As before, because of the way the Mach stem is aligned with the grid, there is probably insufficient dissipation added via the contact and shear waves to counteract perturbations that appear via the acoustic waves.

2.6. *Odd–Even decoupling*

By far the most insidious failing that we have come across has, we believe, gone unreported in the literature. During the course of producing very-high-resolution simulations, we have noticed a tendency for odd–even decoupling to occur along the length of planar shocks which are aligned with the mesh (for an example see Figure 11). Of the Riemann solvers that we have at our disposal, this failing afflicts an exact solver,[22] Roe's solver and Toro's linearized solver.[7] We emphasize the fact that this phenomenon only becomes apparent for very-high-resolution simulations which suffer some systematic perturbation. However, as will be shown below, the required perturbation can arise quite innocuously, so we suspect that this failing will prove fairly widespread once very-high-resolution simulations become commonplace owing to increases in computer power.

Now since we obtain our high grid resolution by means of a fairly complex mesh adaption scheme,[16] it seemed reasonable to suppose that this odd–even decoupling was attributable to some coding error, but an exhaustive search for such an error proved fruitless. Subsequently we have managed to reproduce this failing in a more controlled manner, as has a colleague using an independent code,[23] so we have little doubt that this tendency for odd–even decoupling to occur constitutes a genuine failing rather than being the manifestation of some deficiency of our code. That said, our adaptive mesh scheme clearly exasperates this failing. In the next section

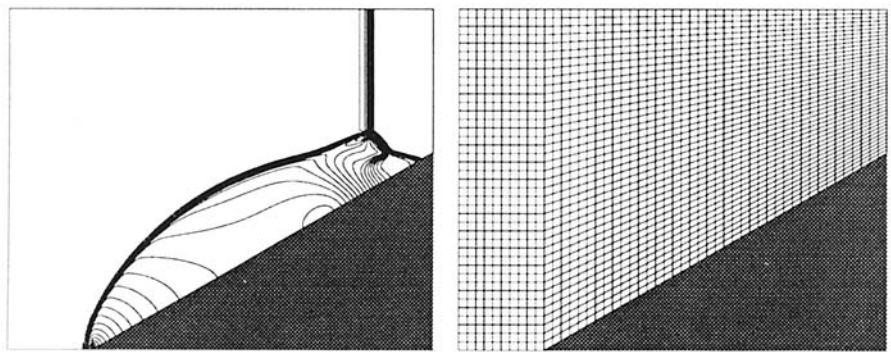

(a) (b)

Figure 4. The principal Mach stem arising from the reflection of a plane shock from a ramp is inexplicably kinked: (a) pressure contours; (b) computational grid

we shall present a possible diagnosis of the mechanism which causes this mode of failure; here we merely present the evidence that it exists.

Figure 5 shows several snapshots of the density contours from the simulation of a plane shock wave, $M_s = 6$ with $\gamma = 1{\cdot}4$, propagating down a duct. For this calculation we have used Roe's scheme together with a nominally uniform grid of 20 × 800 cells with unit spacing, the centreline of which is perturbed from that of a perfectly uniform mesh in the following manner:

$$Y_{i,\,jmid} = \begin{cases} Y_{jmid} + 10^{-6} & \text{for } i \text{ even,} \\ Y_{jmid} - 10^{-6} & \text{for } i \text{ odd.} \end{cases}$$

This perturbation to the grid centreline promotes odd–even decoupling along the length of the shock. Note that the shock has propagated some 15 channel widths before the decoupling first becomes apparent; see frame (b). For this point in the calculation Figure 6 shows a series of slices across the duct, for both the density and pressure fields, as one moves from the head to the foot of the shock. Interestingly, within the shock the decoupling of the pressure field is out of phase with the decoupling of the density field. As the shock continues to propagate down the duct, so the decoupling becomes progressively worse, until the shock breaks down completely.

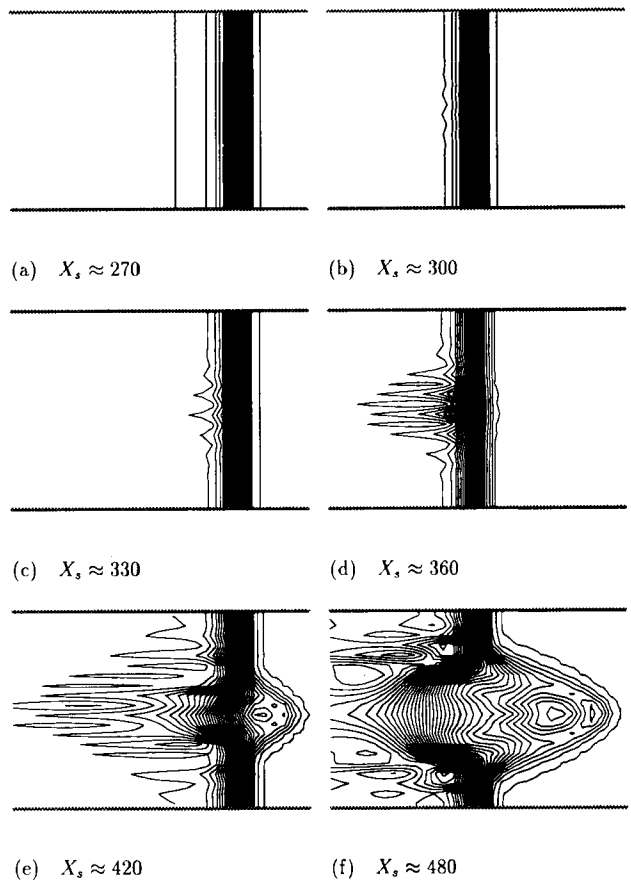

Figure 5. Odd–even decoupling occurs for a shock propagating down a duct

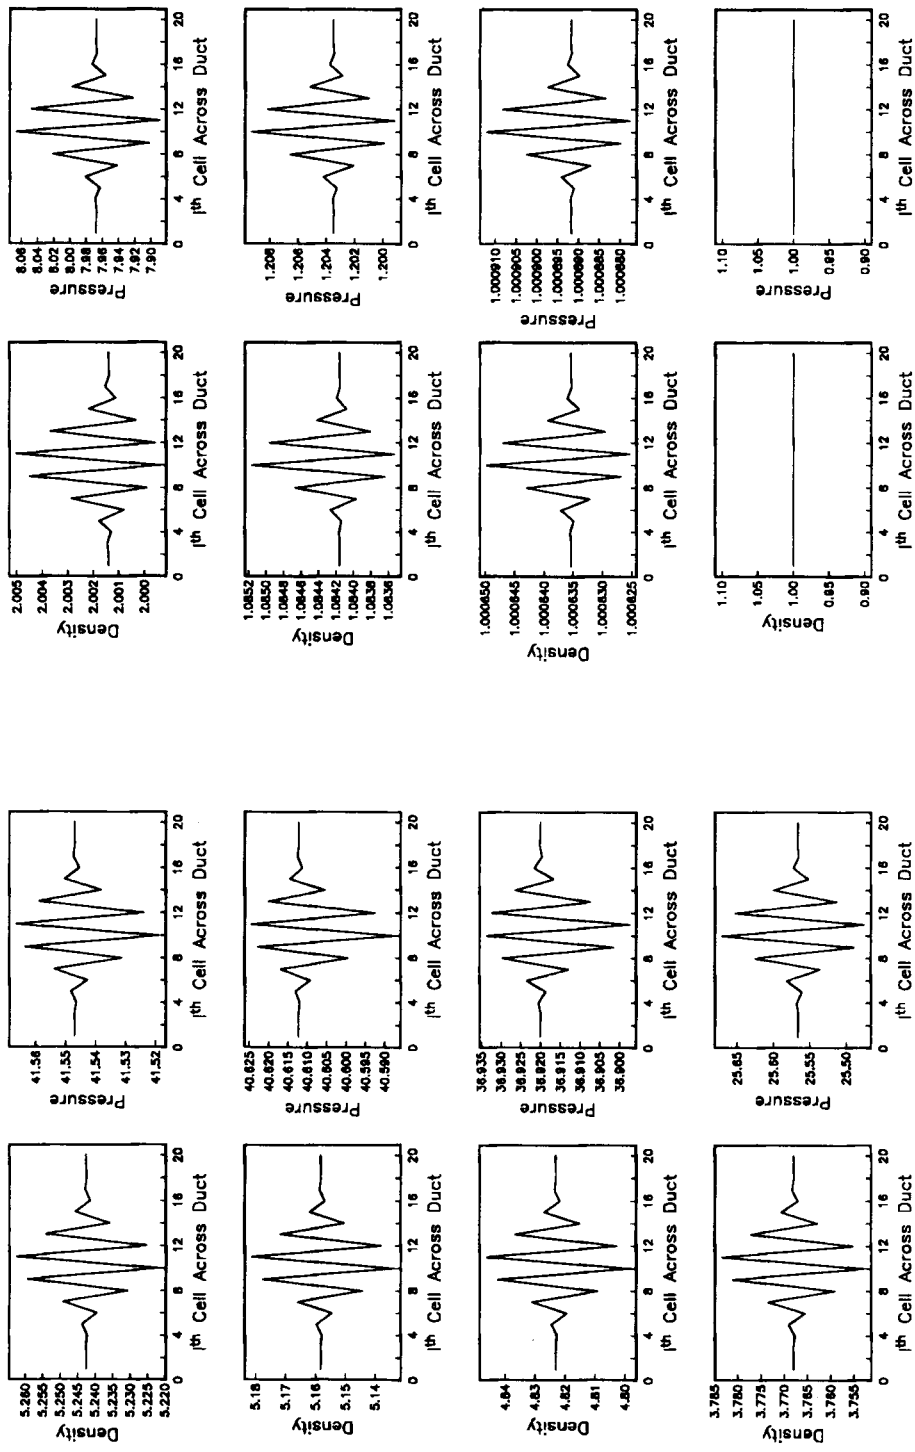

Figure 6. Transverse slices of the pressure and density fields for the shock shown in frame (b) of Figure 5

However, at no stage in the calculation does the code blow up in the sense that it generates a floating point exception; it simply goes astray.

3. ODD–EVEN DECOUPLING—A DIAGNOSIS?

As yet, the tools do not exist that would allow us to perform a rigorous non-linear stability analysis for some shock-capturing scheme applied to the Euler equations. However, it is possible to examine the way in which a scheme evolves certain sets of prescribed data so as to ascertain its likely stability characteristics. Since some Riemann solvers allow odd–even decoupling to develop along the length of a plane shock, it might prove fruitful to examine how different schemes evolve sawtooth-type initial data. To this end we consider the one-dimensional Euler equations with a passive component of shear velocity:

$$\frac{\partial}{\partial t}\begin{pmatrix}\rho\\ \rho u\\ \rho v\\ E\end{pmatrix}+\frac{\partial}{\partial y}\begin{pmatrix}\rho v\\ \rho uv\\ \rho v^2+p\\ (E+p)v\end{pmatrix}=0. \tag{1}$$

The quantities ρ, p, u, v and E are density, pressure, the passive shear component of velocity, the velocity in the y-direction and the total energy per unit volume respectively. For a perfect gas

$$E=\frac{p}{\gamma-1}+\tfrac{1}{2}\rho(u^2+v^2),$$

where γ is the ratio of specific heats. We assume that the computational mesh is uniform, with mesh spacing Δy, and that the discrete solution at time t^n is given by

$$\rho_j^n=\rho+\hat{\rho}^n, \qquad p_j^n=p+\hat{p}^n, \qquad u_j^n=u, \qquad v_j^n=0 \tag{2}$$

if j is even and by

$$\rho_j^n=\rho-\hat{\rho}^n, \qquad p_j^n=p-\hat{p}^n, \qquad u_j^n=u, \qquad v_j^n=0$$

if j is odd. Here $\hat{\rho}^n$ and $\hat{p}^n$ are the amplitudes of the sawtooth profiles for the density and pressure fields respectively. We shall consider two schemes which may be expressed in the form

$$\mathbf{W}_j^{n+1}=\mathbf{W}_j^n-\frac{\Delta t}{\Delta y}(\mathbf{G}_{j+1/2}^n-\mathbf{G}_{j-1/2}^n), \tag{3}$$

where $\mathbf{W}$ is the conserved variable vector $(\rho, \rho u, \rho v, E)^{\mathrm{T}}$ and $\mathbf{G}_{j+1/2}^n$ is a first-order flux function computed from the states $\mathbf{W}_j^n$ and $\mathbf{W}_{j+1}^n$.

3.1. Roe's scheme

Using Roe's scheme,[1] the interface flux for the system of equations (1) may be written as

$$\mathbf{G}_{j+1/2}^n=\tfrac{1}{2}(\mathbf{G}_j^n+\mathbf{G}_{j+1}^n)-\tfrac{1}{2}\sum_{k=1}^{k=4}\alpha_{j+1/2}^{(k)}|\lambda_{j+1/2}^{(k)}|\mathbf{e}_{j+1/2}^{(k)},$$

where the wave speeds $\{\lambda^{(k)}\}$, wave strengths $\{\alpha^{(k)}\}$ and eigenvectors $\{\mathbf{e}^{(k)}\}$ are given by

$$\lambda^{(1)} = \tilde{v} - \tilde{a}, \qquad \lambda^{(2)} = \tilde{v}, \qquad \lambda^{(3)} = \tilde{v}, \qquad \lambda^{(4)} = \tilde{v} + \tilde{a},$$

$$\alpha^{(1)} = \frac{\Delta p - \tilde{\rho}\tilde{a}\Delta v}{2\tilde{a}^2}, \qquad \alpha^{(2)} = \Delta\rho - \frac{\Delta p}{\tilde{a}^2}, \qquad \alpha^{(3)} = \rho\Delta u, \qquad \alpha^{(4)} = \frac{\Delta p + \tilde{\rho}\tilde{a}\Delta v}{2\tilde{a}^2},$$

$$\mathbf{e}^{(1)} = \begin{pmatrix} 1 \\ \tilde{u} \\ \tilde{v} - \tilde{a} \\ \tilde{h} - \tilde{a}\tilde{v} \end{pmatrix}, \qquad \mathbf{e}^{(2)} = \begin{pmatrix} 1 \\ \tilde{u} \\ \tilde{v} \\ \frac{1}{2}(\tilde{u}^2 + \tilde{v}^2) \end{pmatrix}, \qquad \mathbf{e}^{(3)} = \begin{pmatrix} 0 \\ 1 \\ 0 \\ \tilde{u} \end{pmatrix}, \qquad \mathbf{e}^{(4)} = \begin{pmatrix} 1 \\ \tilde{u} \\ \tilde{v} + \tilde{a} \\ \tilde{h} + \tilde{a}\tilde{v} \end{pmatrix}.$$

Here quantities written as $\tilde{(\cdot)}$ are the so-called Roe-averaged quantities, $\tilde{a}$ and $\tilde{h}$ are the Roe-averaged sound speed and total enthalpy respectively and $\Delta(\cdot)$ represents the forward difference operator $(\cdot)_{j+1} - (\cdot)_j$. Now for our chosen data

$$\tilde{(\cdot)}_{j+1/2} = \tilde{(\cdot)}_{j-1/2}, \qquad \Delta(\cdot)_{j+1/2} = -\Delta(\cdot)_{j-1/2}.$$

Therefore

$$\sum_{k=1}^{k=4} \alpha^{(k)}_{j+1/2} |\lambda^{(k)}_{j+1/2}| \mathbf{e}^{(k)}_{j+1/2} = -\sum_{k=1}^{k=4} \alpha^{(k)}_{j-1/2} |\lambda^{(k)}_{j-1/2}| \mathbf{e}^{(k)}_{j-1/2}.$$

Also $\mathbf{G}^n_{j-1} = \mathbf{G}^n_{j+1}$, so the evolution scheme (3) may be written as

$$\mathbf{W}^{n+1}_j = \mathbf{W}^n_j + \frac{\Delta t}{\Delta y} \sum_{k=1}^{k=4} \alpha^{(k)}_{j+1/2} |\lambda^{(k)}_{j+1/2}| \mathbf{e}^{(k)}_{j+1/2},$$

which can be simplified to

$$\mathbf{W}^{n+1}_j = \mathbf{W}^n_j + v_y \frac{\Delta p}{\tilde{a}^2} \begin{pmatrix} 1 \\ \tilde{u} \\ 0 \\ \tilde{h} \end{pmatrix}, \tag{4}$$

where v_y is the Courant number $\tilde{a}\Delta t/\Delta y$. Recognizing that $\mathbf{W}^n_j$ may be expressed as

$$\begin{pmatrix} \rho \\ \rho u \\ 0 \\ E \end{pmatrix} \pm \begin{pmatrix} \hat{\rho}^n \\ u\hat{\rho}^n \\ 0 \\ \hat{p}^n/(\gamma - 1) + (u^2/2)\hat{\rho}^n \end{pmatrix}$$

and that by definition

$$\tilde{a}^2 = (\gamma - 1)\left(\tilde{h} - \frac{\tilde{u}^2}{2} - \frac{\tilde{v}^2}{2}\right),$$

equation (4) may be manipulated to give

$$\hat{\rho}^{n+1} = \hat{\rho}^n - \frac{2v_y}{\tilde{a}^2}\hat{p}^n, \qquad \hat{p}^{n+1} = \hat{p}^n(1 - 2v_y).$$

From this it can be seen that the initial perturbation to the pressure field is damped provided that the CFL condition is met, i.e.

$$v_y \leq 1.$$

However, the form of the evolution for the density perturbation exposes a flaw in Roe's scheme; the density perturbation is fed directly from the pressure perturbation. Making the loose approximation that $\tilde{a}$ remains constant, for a one-off disturbance $(\hat{\rho}^0, \hat{p}^0)$ we have

$$\hat{\rho}^n = \hat{\rho}^0 - \frac{2v_y}{\tilde{a}^2}\hat{p}^0[1 + (1 - 2v_y) + (1 - 2v_y)^2 + \cdots + (1 - 2v_y)^{n-1}]$$

Thus

$$\hat{\rho}^\infty = \hat{\rho}^0 - \frac{\hat{p}^0}{\tilde{a}^2}. \tag{5}$$

Therefore if $\hat{p}^0$ is of opposite sign to $\hat{\rho}^0$, then for a one-off disturbance $\hat{\rho}$ will grow but remain bounded; but if the pressure field is continuously perturbed in a systematic manner, no matter how small the pressure perturbations, $\hat{\rho}$ will grow without bound, albeit slowly. For two-dimensional calculations, although we cannot prove it, we suspect that a strong shock wave moving normal to the y-direction provides this systematic perturbation.

Firstly, it is interesting to note that the failing reported in Section 2.6 is only observed for strong shocks. For a strong shock wave it seems reasonable to suppose that $|\hat{p}^0/\tilde{a}^2|$ is more likely to be larger than $|\hat{\rho}^0|$ than would be the case for weak shocks. Thus even if $\hat{\rho}^0$ and $\hat{p}^0$ are initially of the same sign, they need not remain so; see (5). Now consider frame (b) of Figure 5 and the associated profiles shown in Figure 6. Within the shock the odd–even decouplings of the pressure and density fields are indeed out of phase with one another, which is consistent with the observations made above. Such behaviour will cause the local sound speed to vary along the length of the shock and its profile will exhibit a sawtooth perturbation which is in phase with that of the pressure field. Consequently, the individual segments of the shock will be moving alternately faster and slower than the nominal shock speed. Such movements will exaggerate the sawtooth perturbation to the pressure field along the length of the shock, but diminish that for the density field. The increased pressure perturbations will then promote an increase in the density perturbations as detailed above, and so the whole process repeats itself.

Since there are two competing processes that affect the density perturbations, namely the relative movements of the shock and the decoupling along the length of the shock, we cannot categorically state that Roe's method is bound to break down. However, the weight of numerical evidence suggests that at least for strong shocks Roe's scheme will break down in the manner described here. Given our arguments, it should come as no surprise that Godunov's method also exhibits a tendency to allow odd–even decoupling to occur along the length of a strong shock wave. Since it is the sweep parallel to the shock that primarily causes the instability, the differences between using an exact Riemann solver and Roe's linearized solver for data that are nominally uniform should have little bearing on the growth of the instability.

Finally, before moving on, it should be noted that none of the popular entropy fixes which are applied to Roe's scheme cures this particular failing, except for the case where Harten's fix is applied to the shear and contact waves.* Simply altering the acoustic wave speeds can have

* To reiterate the comment made in Section 2.4, applying Harten's entropy fix to the linearly degenerate wave fields has no mathematical or physical justification, but is merely a convenient way in which to add an amount of artificial dissipation.

no affect; because of the symmetry of the data, both waves will be changed by the same amount and so the problem will persist. Also, moving to a high-order version of Roe's scheme will not improve matters, because the odd–even decoupling will cause a high-order flux function to drop to the first-order function.

3.2. *Einfeldt's HLLE scheme*

For Einfeldt's HLLE scheme[6] the interface flux function is given by

$$\mathbf{G}^n_{j+1/2} = \frac{b^{(+)}\mathbf{G}^n_{j+1} - b^{(-)}\mathbf{G}^n_j}{b^{(+)} - b^{(-)}} + \frac{b^{(+)}b^{(-)}}{b^{(+)} - b^{(-)}}(\mathbf{W}^n_{j+1} - \mathbf{W}^n_j),$$

where

$$b^{(+)} = \max(0,\ \lambda^{(4)}_{j+1/2},\ v_{j+1} + a_{j+1}), \qquad b^{(-)} = \min(0,\ \lambda^{(1)}_{j-1/2},\ v_j - a_j)$$

and $\lambda^{(1)}$ and $\lambda^{(4)}$ are the two acoustic wave speeds from Roe's method. Now for our chosen data $(\tilde{\ })_{j+1/2}$ is equal to $(\tilde{\ })_{j-1/2}$, therefore

$$b^{(+)}_{j+1/2} = -b^{(-)}_{j-1/2} \approx \tilde{a}, \qquad b^{(-)}_{j+1/2} = -b^{(+)}_{j-1/2} \approx -\tilde{a}.$$

Using these signal weightings, it may be found that

$$\hat{\rho}^{n+1} = (1 - 2v_y)\hat{\rho}^n, \qquad \hat{p}^{n+1} = (1 - 2v_y)\hat{p}^n,$$

where

$$v_y \approx \frac{\tilde{a}\Delta t}{\Delta y}.$$

From this it can be seen that both the density and pressure perturbations are damped provided that the CFL condition is satisfied. Just as important, however, is the fact that the pressure perturbation does not feed into the density perturbation, so we would not expect the HLLE solver to exhibit the odd–even decoupling that afflicts both Roe's scheme and Godunov's scheme; numerical experimentation confirms this expectation.

It is our contention that any scheme for which it can be shown that the perturbation to the pressure field feeds the perturbation to the density field will be afflicted by the odd–even decoupling shown in Figure 5. Thus it comes as no surprise to find that Toro's linearized Riemann solver[7] is afflicted by this failing, but Liou and Steffen's AUSM scheme[19] is not. The way is now open for some mathematician to perform a more rigorous analysis than we are able, so as to shed additional light on the mechanism which causes this particular failing.

4. AN ADAPTIVE RIEMANN SOLVER

Having exposed many of the weaknesses of Riemann solvers, we now present a simple strategy that we have found useful for improving the robustness of Godunov-type schemes. In essence, we select the precise flavour of upwinding to match the local flow data such that a particular Riemann solver is only employed in those situations where it is known to give reliable results. By recognizing the limitations of any one solver, it is possible to reap its advantages without suffering its attendant failings.

Our synergetic strategy has a number of attractions, not least of which is that some favoured solver need not be jettisoned simply because it occasionally fails. However, it does introduce the difficulty of how to decide when to use one Riemann solver in preference to another. It has been our experience that this added difficulty is not particularly bothersome, since we tend to combine a single high-resolution Riemann solver with just one or two other solvers that prove more reliable under conditions which are fairly well defined, and so a set of *ad hoc* switching functions suffice. For example, some of the worst failings of Riemann solvers occur in the vicinity of strong shock waves. To overcome such failings, we employ Einfeldt's HLLE scheme. Now it makes little sense to chop and change the choice of Riemann solver used along the length of a shock wave, since to do so would inevitably perturb a planar shock front. Hence we apply this particular Riemann solver throughout the immediate vicinity of a strong shock. Thus the HLLE switching function need only locate the position of a shock wave, but such functions already exist in the guise of mesh refinement monitor functions.

A simple test that identifies those cell interfaces which are in the vicinity of a strong shock is to check whether or not

$$\frac{|p_r - p_l|}{\min(p_l, p_r)} > \alpha, \tag{6}$$

where α is some threshold parameter which is problem-dependent and p_r and p_l refer to the pressures which act on the interface. If this condition is met, the two cells separated by the interface are flagged as lying within a strong shock. Thus when it comes to computing cell–interface fluxes, if the cells either side of an interface are both flagged as lying within a strong shock, the flux is computed using the HLLE solver. Note that since numerical shocks are invariably smeared over several mesh cells, it is worth locating shocks using a projection of the flow solution on a grid which is coarser than that used for the calculation. On such a grid a shock will appear much less smeared and so the left-hand side of the above switching function will be a fair indication of its strength. Once a set of cells have been flagged on this coarse mesh, the flags may be prolongated to the actual computational mesh so as to find those cells which lie in the vicinity of a shock wave.

Before proceeding further, several observations should be made. Firstly, although the HLLE solver is adjudged to be a low-resolution scheme, it does in fact resolve *isolated* shocks as well as an exact Riemann solver does. Consequently, using this robust solver in the vicinity of strong shock waves does not necessarily pollute a scheme's resolution, as would be the case if artificial dissipation were used to augment the dissipation provided via upwinding. For many inviscid calculations the amount of pollution proves to be negligible, and whilst some degradation would be expected for the case of a strong shock interacting with a boundary layer, it may well be unnecessary to employ the HLLE solver in such a situation because of the extra dissipation provided by the real viscous terms. Secondly, although the HLLE switching function requires a tunable parameter α, the retuning of this parameter is less involved than the retuning of an artificial dissipation mechanism; in general, it is far simpler to determine where extra dissipation should be added than it is to determine how much extra dissipation to add. For many problems, assuming that shocks are located as described above, a sensible threshold on the shock strength can be specified *a priori*. Lastly, it should be noted that our strategy of switching Riemann solvers may not prove suitable for those implicit schemes which require that the numerical flux function be differentiable.

Figure 7 shows how the HLLE solver may be used to correct the tendency of Roe's scheme

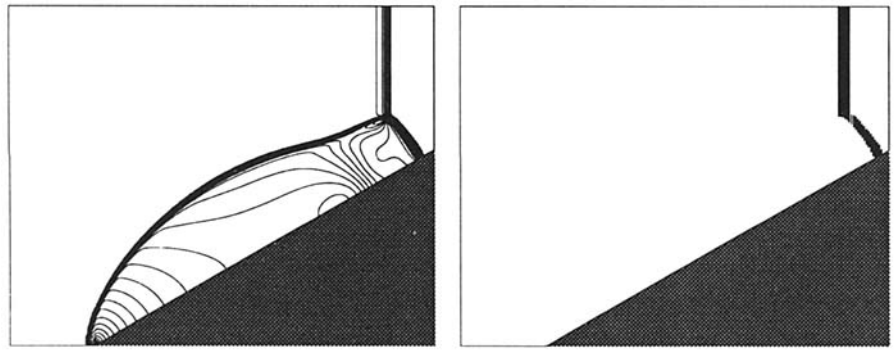

(a) (b)

Figure 7. The HLLE scheme can be used to circumvent the tendency of Roe's method to produce kinked Mach stems: (a) pressure contours; (b) HLLE switching function

to produce kinked Mach stems (cf. Figure 4). For this calculation the HLLE switching function was tuned such that it would only be activated by the incident shock and the principal Mach stem; α was simply set to half the strength of the incident shock. Note that apart from the region near the Mach stem, these new results are very similar to the old ones. This shows that the HLLE scheme has had no adverse affect on the resolution of Roe's scheme. Similarly, Figure 8 shows how the carbuncle phenomenon may be circumvented (cf. Figure 3). Here we have restricted the HLLE solver to cells near the stagnation line in order to demonstrate how localized the failing of Roe's scheme really is. In practice, however, we would advocate using the HLLE scheme along the whole length of the bow shock so as to maximize robustness without compromising resolution. Again a sensible value of α can be found *a priori* by using some large fraction of the shock strength along the stagnation line, which can be estimated, given the freestream Mach number, by assuming that the flow is locally one-dimensional. As shown in Figure 9, the HLLE solver may also be used to good effect to prevent Godunov's scheme from admitting expansion shocks (cf. Figure 1). Here we have employed the HLLE solver along the sonic line and in regions where the expansion waves are strong.

Having presented the gist of our strategy, we see little point in trying to sell a particular combination of solvers. Starting with some high-resolution Riemann solver, whose choice will inevitably be a matter of personal taste, the correct combination of solvers will depend both on that scheme's weaknesses and on the specific application in hand. In turn, the combination of Riemann solvers will dictate the choice of switching functions. Therefore we shall resist the temptation to recommend a specific course of action; instead, we present two simulations that show how an adaptive Riemann solver might be used to good effect. Briefly, both simulations were done using the two-dimensional analogue of the one-dimensional Euler equations given in Section 3. These equations were integrated using the two-step, finite volume scheme which is attributable to Hancock.[24] This scheme employs van Leer's MUSCL approach[9] to achieve a second-order extension to Godunov's method; hence different Riemann solvers may be slotted directly into the scheme so as to vary the flavour of the upwinding. Although the calculations were performed using an adaptive mesh algorithm,[16,25] the mesh refinement monitor function was such that the calculations employed a nominally uniform Cartesian mesh.

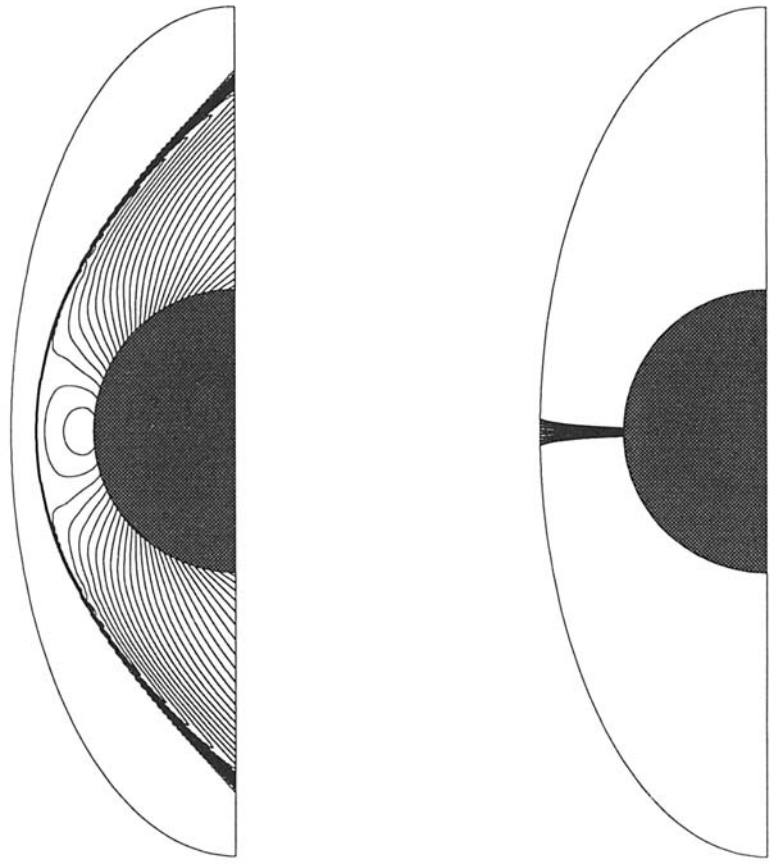

Figure 8. The HLLE scheme can be used to circumvent the *carbuncle phenomenon*: (a) density contours; (b) HLLE switching function

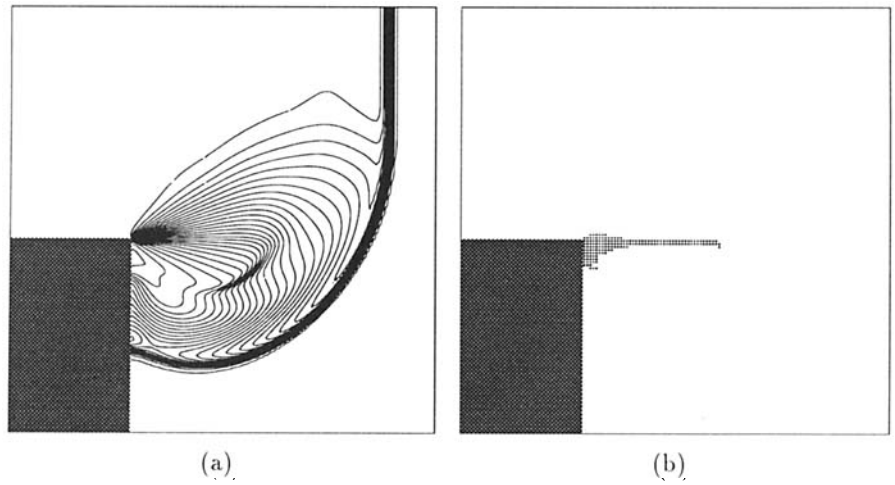

Figure 9. The HLLE scheme can be used to prevent Godunov's method from producing expansion shocks: (a) density contours; (b) HLLE switching function

Our first example concerns the simulation of a strong shock wave diffracting around a 90° corner; the shock Mach number and the ratio of specific heats are 5·09 and 1·4 respectively. We have computed this test problem using a combination of three different Riemann solvers: Toro's linearized Riemann solver was used to perform the MUSCL reconstruction step of Hancock's scheme as described by Quirk,[25] and the upwinding step was performed by adaptively selecting between Roe's solver and the HLLE solver. The parameter α used by the switching function (6) was set to unity so as to limit the HLLE solver to the incident and diffracted shock fronts and to a small region near the apex of the corner. Figure 10 shows a schlieren-type snapshot taken from this simulation. The different shades of grey depict the magnitude of the gradient of the density field; the darker the shade, the larger the magnitude; details of this shading procedure are given in Appendix A. Here it is not our intention to assess the accuracy of these results; the interested reader may do this using the experimental results of Bazhenova *et al.*[26] and the computational results of Hillier.[27] Instead, we wish to illustrate the fact that certain Riemann solver failings, if left unaddressed, can place an upper limit on the resolution of simulations that may be performed.

Consider the consequences of disabling the HLLE switching function so that Roe's solver alone is used for the upwinding stage of Hancock's scheme. The tendency of Roe's solver to allow odd–even decoupling to occur along a planar shock wave which is aligned with the grid will sooner or later cause this simulation to come to grief (see Figure 11), thus precluding the possibility of performing very detailed simulations. By way of comparing the resolution of these two sets of results, for Figure 10 there are 560 mesh cells from the apex of the corner to the point where the Mach stem meets the wall, while for Figure 11 there are only 120 cells. The question of whether or not our adaptive mesh algorithm contains some flaw which exasperates the odd–even decoupling is largely academic. The fact remains that Roe's solver is susceptible to this mode of failure whereas the HLLE solver is not. Whether the initial stimulus comes from a distorted mesh as in Section 2 or from some component of the mesh adaption scheme as seems likely here is immaterial.

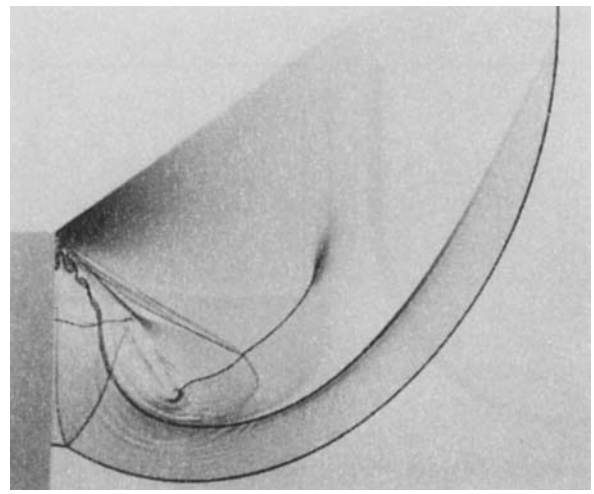

Figure 10. A schlieren- type snapshot from the diffraction of a strong shock wave around a 90° degree corner

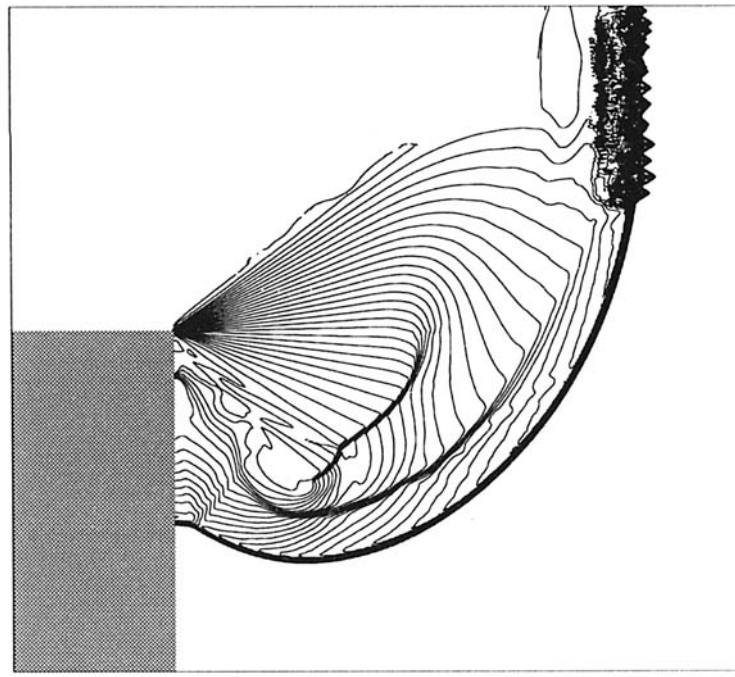

Figure 11. On its own, Roe's approximate Riemann solver cannot be used to reproduce the resolution of the simulation shown in Figure 10

So as not to leave the impression that the above shortcoming is somehow peculiar to Roe's method, we present a second set of results which are taken from the interaction of a planar shock wave with a half-diamond; the shock Mach number is 2·85, the ratio of specific heats is 1·4 and the angle at the apex of the diamond is 90°. As before, we have run this test problem using a combination of three different Riemann solvers, but this time we have substituted an exact Riemann solver[22] in place of Roe's linearized solver. Figure 12 shows a schlieren-type snapshot from this calculation; note that some 800 cells cover the width of the diamond, so this calculation is well resolved. Also, as an aside, we note that the quality of these results may be gauged by comparing them with the experimental results given by Glass *et al.*[28] Once again, if the HLLE switching function is disabled, the simulation is ruined by the odd–even decoupling that occurs along the length of the incident shock (see Figure 13). Note that this second calculation is of lower resolution than the first; only 400 cells cover the width of the diamond.

In this section we have attempted to show that the robustness of Godunov-type schemes may be improved by employing an adaptive Riemann solver, where the specific flavour of upwinding is altered to suit the local flow conditions. If used sensibly, this strategy can overcome most known failings of individual solvers. Despite our efforts, we recognize that the majority of shock-capturing practitioners will continue to use artificial dissipation as a band aid to fix a particular Riemann solver at the first signs of any failing, simply because it is expedient to do so. Whilst we find this approach disappointing, the principal aim of this paper is to emphasize the fact that most of the Riemann solvers that are in common use *must be augmented in some way* if they are to be used for the purpose of producing genuinely high-resolution simulations of shock hydrodynamic phenomena.

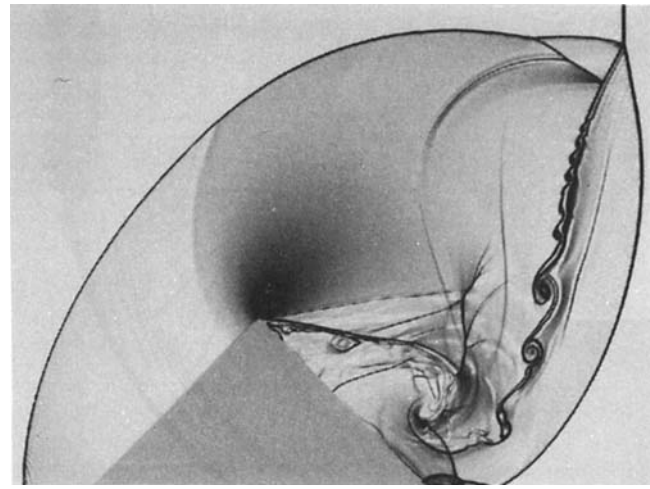

Figure 12. A schlieren-type snapshot from the interaction of a planar shock wave with a half-diamond

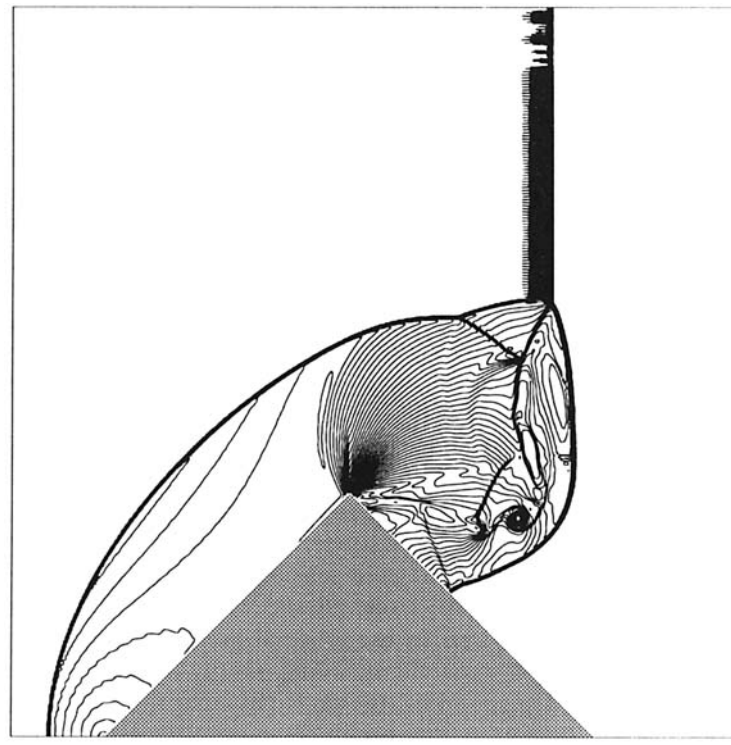

Figure 13. On its own, an exact Riemann solver cannot be used to reproduce the resolution of the simulation shown in Figure 12

5. CONCLUSIONS

Unless the dissipation provided via upwinding is augmented by some other mechanism, any Godunov-type scheme built upon a single Riemann solver will probably be flawed. For example, when subjected to some small but systematic form of perturbation, most popular Riemann solvers for the Euler equations, including the exact solver, cannot prevent odd–even decoupling occurring along the length of a strong shock wave which is aligned with the computational mesh. Thus far, this flaw has gone largely unnoticed simply because it is only exposed by very-high-resolution simulations. However, given that the required perturbations can arise quite innocuously, this mode of failure should prove fairly widespread once genuinely high-resolution simulations become commonplace owing to increases in computer power.

Although most flaws can be controlled by the judicious use of a small amount of artificial dissipation, to do so necessarily leads to a reduction in the resolution of the scheme. We favour an alternative approach whereby the failings of any one Riemann solver are circumvented by combining it with one or more complementary solvers. In essence, we advocate selecting the precise flavour of upwinding to suit the flow data. Admittedly, this synergistic strategy is not as aesthetically pleasing as having a single Riemann solver for all occasions, but we have shown that it can be made to work quite effectively; besides which, Riemann solvers are sometimes touted as being a solution-adaptive technique, so the concept of an adaptive Riemann solver is not that contrived.

Looking to the future, it is to be hoped that genuinely multidimensional Riemann solvers will overcome many of the shortcomings of today's dimensionally split schemes. However, given the way in which the present shortcomings have been stumbled across, these multidimensional schemes may themselves arrive complete with subtle failings with which to ensnare the unwary.

ACKNOWLEDGEMENTS

I am indebted to Dr. P. A. Jacobs for taking the time to reproduce the failing reported in Section 2.6 and I wish to thank Professor P. L. Roe for his comments concerning this work. This research was supported by the National Aeronautics and Space Administration under NASA Contract NAS1-18605 while the author was in residence at the Institute for Computer Applications in Science and Engineering (ICASE), NASA Langley Research Center, Hampton, VA 23681.

APPENDIX: SCHLIEREN-TYPE PLOTS

The plots shown in Figures 10 and 12 depict the magnitude of the gradient of the density field, namely

$$\sqrt{\left[\left(\frac{\partial \rho}{\partial x}\right)^2 + \left(\frac{\partial \rho}{\partial y}\right)^2\right]},$$

and hence they may be viewed as idealized schlieren images. So as to accentuate weak flow features, the following non-linear shading function has been used:

$$shade = \exp(-k\psi).$$

Here k is a constant and ψ is a weighting function given by

$$\frac{|\nabla\rho| - |\nabla\rho|_0}{|\nabla\rho|_1 - |\nabla\rho|_0},$$

where

$$|\nabla\rho|_0 = k_0|\nabla\rho|_{\max}, \qquad |\nabla\rho|_1 = k_1|\nabla\rho|_{\max},$$

k_0 and k_1 being constants. Note that *shade* is limited to values between 0 and 1, so for a 24-bit colour graphics system the grey level *shade* may be converted to an $\langle R, G, B\rangle$ triplet using

$$\langle 255 * shade,\ 255 * shade,\ 255 * shade\rangle.$$

For both figures the constants k, k_0 and k_1 were set to 5, 0·05 and $-0{\cdot}001$ respectively.

REFERENCES

1. P. L. Roe, 'Characteristic-based schemes for the Euler equations', *Ann. Rev. Fluid Mech.*, **18**, 337–365 (1986).
2. R. Courant and K. O. Friedrichs, *Supersonic Flow and Shock Waves*, Interscience, New York, 1948; reprinted as *Applied Mathematical Physics*, Vol. 21, Springer, New York, 1976.
3. M. Holt, *Numerical Methods in Fluid Dynamics*, Springer, Berlin, 2nd edn, 1984.
4. S. Osher, 'Riemann solvers, the entropy condition, and difference approximations', *SIAM J. Numer. Anal.*, **21**, 289–315 (1984).
5. M. Pandolfi, 'A contribution to numerical prediction of unsteady flows', *AIAA J.*, **22**, 602–610 (1984).
6. B. Einfeldt, 'On Godunov-type methods for gas dynamics', *SIAM J. Numer. Anal.*, **25**, 294–318 (1988).
7. E. F. Toro, 'A linearised Riemann solver for the time-dependent Euler equations of gas dynamics', *Proc. R. Soc. Lond. A*, **434**, 683–693 (1991).
8. P. Colella and P. R. Woodward, 'The piecewise parabolic method (PPM) for gas-dynamical simulations', *J. Comput. Phys.*, **54**, 174–201 (1984).
9. B. van Leer, 'Towards the ultimate conservative difference scheme, II. Monotonicity and conservation combined in a second-order scheme', *J. Comput. Phys.*, **14**, 361–376 (1974).
10. P. R. Woodward and P. Colella, 'The numerical simulation of two-dimensional fluid flow with strong shocks', *J. Comput. Phys.*, **54**, 115–173 (1984).
11. P. L. Roe, 'Some contributions to the modelling of discontinuous flows', in *Lectures in Applied Mathematics*, Vol. 22, American Mathematical Society, Providence, RI, 1985.
12. B. Einfeldt, C. D. Munz, P. L. Roe and B. Sjögreen, 'On Godunov-type methods near low densities', *J. Comput. Phys.*, **92**, 273–295 (1991).
13. H. C. Yee, R. F. Warming and A. Harten, 'Implicit total variation diminishing (TVD) schemes for steady-state calculations', *J. Comput. Phys.*, **57**, 327–360 (1985).
14. T. W. Roberts, 'The behaviour of flux difference splitting schemes near slowly moving shock waves', *J. Comput. Phys.*, **90**, 141–160 (1990).
15. S. Osher and F. Solomon, 'Upwind difference schemes for hyperbolic systems of conservation laws', *Math. Comput.*, **38**, 339–374 (1982).
16. J. J. Quirk, 'An adaptive grid algorithm for computational shock hydrodynamics', *Ph.D. Thesis*, College of Aeronautics, Cranfield Institute of Technology, 1991.
17. K. M. Peery and S. T. Imlay, 'Blunt-body flow simulations', *AIAA Paper 88-2904*, 1988.
18. H. C. Lin, 'Dissipation additions to flux-difference splitting', *AIAA Paper 91-1544-CP*, 1991.
19. M.-S. Liou and C. J. Steffen, Jr., 'A new flux splitting scheme', *NASA TM 104404*, 1991.
20. K. Sawada, 'A multi-dimensional extension of the preprocessing approach for cell centred finite volume scheme', *AIAA Paper 91-1536*, 1991.
21. K. Itoh, K. Takayama and G. Ben-Dor, 'Numerical simulation of the reflection of a planar shock wave over a double wedge', *Int. J. Numer. Methods Fluids*, **13**, 1153–1170 (1991).
22. E. F. Toro, 'A fast Riemann solver with constant covolume applied to the random choice method', *Int. J. Numer. Methods Fluids*, **9**, 1145–1164 (1989).
23. P. A. Jacobs, University of Queensland, personal communication, 1991.
24. G. D. van Albada, B. van Leer and W. W. Roberts, 'A comparative study of computational methods in cosmic gas dynamics', *Astron. Astrophys.*, **108**, 76–84 (1982).
25. J. J. Quirk, 'An alternative to unstructured grids for computing gas dynamic flows around arbitrarily complex two-dimensional bodies', *ICASE Rep. 92-7*, 1992; *Comput. Fluids*, **23** (1), 125–142 (1994).
26. S. K. Bazhenova, L. G. Gvozdeva and M. G. Nettleton, 'Unsteady interactions of shock waves', *Prog. Aerospace Sci.*, **21**, 249–331 (1984).
27. R. Hillier, 'Computation of shock wave diffraction at a ninety degrees convex edge', *Shock Waves*, **1**, 89–98 (1991).
28. I. I. Glass, J. Kaca, D. L. Zhang, H. M. Glaz, J. B. Bell, J. A. Trangenstein and J. P. Collins, 'Diffraction of planar shock waves over half-diamond and semicircular cylinders: an experimental and numerical comparison', *AIP Conf. Proc.*, 1989.