

Review Article

The Numerical Simulation of Two-Dimensional Fluid Flow with Strong Shocks

PAUL WOODWARD

Lawrence Livermore National Laboratory, University of California, Livermore, California 94550

AND

PHILLIP COLELLA

Lawrence Berkeley Laboratory, University of California, Berkeley, California 94720

Received July 15, 1982; revised August 25, 1983

Results of an extensive comparison of numerical methods for simulating hydrodynamics are presented and discussed. This study focuses on the simulation of fluid flows with strong shocks in two dimensions. By "strong shocks," we here refer to shocks in which there is substantial entropy production. For the case of shocks in air, we therefore refer to Mach numbers of three and greater. For flows containing such strong shocks we find that a careful treatment of flow discontinuities is of greatest importance in obtaining accurate numerical results. Three approaches to treating discontinuities in the flow are discussed—artificial viscosity, blending of low- and high-order-accurate fluxes, and the use of nonlinear solutions to Riemann's problem. The advantages and disadvantages of each approach are discussed and illustrated by computed results for three test problems. In this comparison we have focused our attention entirely upon the performance of schemes for differencing the hydrodynamic equations. We have regarded the nature of the grid upon which such differencing schemes are applied as an independent issue outside the scope of this work. Therefore we have restricted our study to the case of uniform, square computational zones in Cartesian coordinates. For simplicity we have further restricted our attention to two-dimensional difference schemes which are built out of symmetrized products of one-dimensional difference operators.

I. INTRODUCTION

Over the last 35 years, a great number of numerical schemes have been devised for the simulation of compressible gas dynamics on digital computers. A major difficulty which has enlivened research in this area is the problem of representing the shock and contact discontinuities which arise in these simulations. As early as 1950 a solution to this problem was proposed by von Neumann and Richtmyer [1]. With minor modifications this first solution to the problem is still in general use today. However,

certain disadvantages of von Neumann and Richtmyer's approach have led to continued efforts at more convenient, more accurate, and/or more elegant solutions. In this article we will discuss three different approaches to the representation of discontinuities which have evolved over the years. We will limit our consideration to approaches to "shock capturing," in which shocks are smeared out on the grid somewhat so that they may be treated in a relatively simple and convenient fashion. Our goal is to present a comprehensive comparison of these approaches, so that their relative merits can be accurately assessed.

The motivation for this work is the great need for accurate simulations of flows with strong shocks which exists in many fields of physics. Much experience indicates that the overall accuracy of such simulations is very closely related to the accuracy with which flow discontinuities are represented. Several algorithms have been proposed which perform well when applied to one-dimensional flow problems but which encounter major difficulties in two dimensions. Because two-dimensional calculations dominate present applications of gas dynamic simulation, we will focus on these more difficult simulations here.

Recently techniques for adapting a computational grid to aid in the resolution of flow discontinuities have received a great deal of attention. Although we will make some remarks about this subject, our main interest in this work is to compare the difference schemes which may be used on a particular grid and not the methods for choosing the grid itself. We regard the issues of choosing a difference scheme and of choosing a grid as largely independent, so that it makes sense to study them separately. We will also leave aside all issues relating to the treatment of the geometric source terms which arise when curvilinear coordinates are employed. Thus we will work exclusively in Cartesian coordinates, and we will use only uniform grids. In the same spirit we will generate all our two-dimensional difference schemes by forming symmetrized products of one-dimensional difference operators. All flows computed here will also employ an especially simple equation of state, namely, that of a gamma-law gas with gamma equal to 1.4. Even with all these restrictions, many computational difficulties remain, and there is a wide spread in the performance of the various methods considered below.

II. THREE MAJOR APPROACHES TO THE REPRESENTATION OF DISCONTINUITIES

a. *Artificial Viscosity*

This earliest and most commonly used approach to representing discontinuities was originally suggested by von Neumann and Richtmyer [1]. In the flow of real fluids we believe that there are no discontinuities. There are instead very thin regions of very steep gradients. If terms representing viscosity and heat conduction were included in the usual equations of hydrodynamics, the discontinuities in the solutions of these equations would no longer develop. However, we would have to resolve smooth structures on the very small distance scales characteristic of viscous

momentum transport and molecular heat conduction. Without changing the flow very much we can increase the physical coefficients of viscosity and heat conduction so that discontinuities are spread over distance scales which are negligible but still resolvable on a practical computational mesh. von Neumann and Richtmyer suggested that an artificial viscous pressure be used to smear shocks over a few computational cells, or zones. Later Lapidus [45] suggested that terms describing diffusion of mass, momentum, and energy be used for the same purpose. In both methods the terms added to the differential equations cause dissipation of kinetic energy into heat.

b. *Linear Hybridization*

In this approach the results of two difference schemes are blended together. A high-order difference scheme which is very accurate in smooth flow but badly behaved at discontinuities is blended with a low-order scheme. The low-order scheme should have sufficiently large dissipative truncation errors that it yields monotone profiles for flow discontinuities. In smooth flow, the high-order scheme is used, but near a discontinuity the low-order scheme is blended with it to an extent sufficient to guarantee monotone representations of the jumps at the discontinuity. We will refer to this criterion for determining the blending weight factor as a monotonicity constraint. In order to preserve the exact conservation of mass, momentum, and energy by the composite scheme when the blending factors vary over the grid, it is the fluxes of these conserved quantities at zone interfaces as computed by the two schemes which are blended. These combined fluxes are then differenced in order to update the mass, momentum, and energy of each zone.

This approach of linear hybridization is similar to the artificial viscosity approach in one respect. To a high-order method which oscillates near discontinuities some terms are added which are negligible in smooth flow and which are strongly dissipative near discontinuities. However, these terms are not motivated by analogy with a more realistic physical model for the flow than is given by the inviscid flow equations. Instead, the added terms are designed specifically to give the sharpest possible discontinuity profiles which are also monotone. They are able to perform this task better than the artificial viscosity terms; but under certain circumstances the monotonicity constraint proves not to be the physically appropriate condition, and then difficulties can arise.

c. *Godunov's Approach*

A third means of treating discontinuities, suggested by Godunov [2], is to introduce explicit nonlinearity into the difference method. In the two other approaches, difference schemes are derived from Taylor series expansions of the terms in the differential equations. This technique is fundamentally based upon the assumption that the solution is smooth. At a discontinuity this assumption is inappropriate; hence the need to force a well-behaved solution by introducing an

unphysically large viscosity or an unphysical constraint of monotonicity. In 1959 Godunov pioneered a new approach to this problem [2]. Instead of building up a full solution to the hydrodynamic equations by piecing together smooth, small-amplitude solutions, he built up a solution by piecing together discontinuous solutions. These discontinuous solutions closely approximate the smooth ones where those are appropriate, but they have the great additional advantage of approximating the true solution reasonably well even when that solution is not smooth.

Godunov made use of a nonlinear flow problem which is simple enough to permit exact solution—Riemann's shock tube problem. This simple solution describes the nonlinear flow which develops from a discontinuous jump separating two constant states. In general, the jump develops into two nonlinear waves, either shocks or rarefactions, with a contact discontinuity in between (cf. Ref. [3]). Godunov's approach was then to approximate a hydrodynamic flow by a large number of constant states, compute their interactions exactly, and average the results in a conservative fashion. This procedure leads to an accurate and very well-behaved treatment of shock discontinuities. Godunov's approach has also been extended to higher-order methods [4–9]. In all of these extensions, the Riemann solver is an essential element in allowing narrow discontinuities without unphysical oscillations.

A different way of using Riemann problems was developed by Glimm, Chorin, and others [10–15]. This method, the random choice method, represents the flow within each zone by the detailed solution to the Riemann problem sampled at a representative point within the zone. Diffusive errors from spatial averaging are avoided, but errors in the time dimension take their place. Unfortunately the very desirable properties of this method in one dimension do not persist in its multidimensional formulations to date (cf. [13]), so that we will not discuss the method further here.

III. ADVANTAGES AND DISADVANTAGES OF THE THREE APPROACHES

a. *Artificial Viscosity*

The major advantage of this approach is its simplicity. If a large enough artificial viscosity is added to the equations of hydrodynamics, their solutions will always be smooth and the design of effective difference methods can be done in a systematic, straightforward way. This procedure is especially successful when the artificial viscous effects are important on distance scales of a fixed length. When practical considerations force this length scale to vary with the size of the computational zones of a strongly nonuniform grid, large unphysical effects can result. This difficulty, which was first noticed in 1968 by Axelrod [16], has been discussed at length in Ref. [17] and will receive no further attention here.

The straightforward use of an artificial viscosity method on practical problems involves a trade-off of computation time and computer memory resources for program simplicity. Because shocks are smeared out over three or more zones by

these methods, fine grids are necessary if the grids are to be uniform to allow for the possibility of shock motion to any point in the problem domain. This necessity for fine grids is caused by the direct relationship between the accuracy of the computation and the thickness of the shocks in many problems of practical interest. The results we will show in a later section demonstrate that if a fully converged solution is desired, this straightforward approach is impractical. Even for the relatively simple 2-D test problems presented below full convergence requires 400,000 words of memory and rather large amounts of computer time.

A more practical approach toward the use of these methods has been suggested by Olinger, and has been developed by Gropp [18], Bolstad [19], and Berger and Olinger [20]. In this approach the mesh is locally refined in both space and time in the neighborhood of a discontinuity. Of course, this idea of local grid refinement is not entirely new, but Olinger and his collaborators have carried it out in a very general and systematic fashion. Local mesh refinement saves the computer memory required for a global fine grid at the cost of introducing substantial program complexity. The computer time needed to achieve a given level of accuracy can also be reduced, but savings here are not as great as one might expect. Considerable time is required to decide where and by how much the grid should be refined. To prevent sudden jumps in the solution at the locations of discontinuous changes in the mesh width, the regions of mesh refinement must be wider than one would suspect at first. Finally, the computation on the refined mesh regions must be specially organized for vector computers. This special organization involves both extra computer time and program complexity. A discussion of these issues can be found in [21]. A detailed assessment of the costs and benefits of the local mesh refinement approach is beyond the scope of this article. An investigation of the usefulness of this approach when combined with some of the difference schemes discussed in this article is presently under way.

An alternative approach is the use of a continuously adaptive mesh. In this approach the mesh is continuously deformed so that structures in the flow which have disparate characteristic length scales are all resolved most effectively. This sort of method is indispensable in problems which contain very thin features whose internal structures must be computed in detail. The method was pioneered in astrophysical and combustion physics simulations [22–29], in which very thin features whose internal structures are essential to the problem must be carefully tracked. In the work of Winkler [24–26] an additional equation relating the distribution of grid points to the flow solution is solved implicitly together with the flow equations. This is a very effective way to deal with the extreme grid distortions forced on the method in order to resolve very sharp features in the flow. Similar approaches have been developed by Dwyer *et al.* [27], and later by Miller and Miller [28], and Gelinis *et al.* [29]. In this last work the mesh equations are developed so that an estimate of the error of the computation is minimized.

All these continuously adaptive grid schemes have one major drawback—they demand an implicit treatment of the flow equations. This is caused by the very small zones which are used to resolve thin features. These zones would force a prohibitively small Courant limit on the time step for an explicit scheme. Consequently, these

methods are not to be recommended for use in problems for which only simple shocks and contact discontinuities need be well resolved. Such discontinuities have jumps which can be determined by conservation laws or other conditions which make no reference to the detailed internal structure of the jump. Therefore it is not necessary to force the computation to be implicit, and hence very slow, in order to obtain the correct jumps by resolving these structures.

Another disadvantage of the continuously adaptive grid schemes should be noted; namely, they become considerably more complicated in two dimensions. Most schemes for adapting a two-dimensional grid to a flow containing thin features involve nonorthogonal meshes (cf. [27] and [55]). This adds considerable complexity to the flow equations in the form of numerous terms involving the metric coefficients of the coordinate system defined by the grid. A detailed assessment of the general effectiveness of continuously adaptive grid schemes is beyond the scope of this article; however, it is clear that for a certain class of problems where the detailed internal structure of thin fronts must be computed they may prove to be very powerful.

b. *Linear Hybridization*

The linear hybridization approach is intermediate in complexity between the approaches of artificial viscosity and explicit nonlinearity. In this approach fluxes of the conserved quantities—mass, momentum, and energy—are computed at the zone interfaces using both a low-order- and a high-order-accurate difference scheme. A linear combination of these fluxes is then computed at each interface using weight factors which may be nonlinear functions of the local conditions of the flow. These blended fluxes are then differenced in order to update the conserved quantities within the zones. In general, the low-order scheme will be well-behaved at discontinuities but inaccurate in smooth flow. In contrast, the high-order scheme will generally oscillate at discontinuities but yield excellent results in smooth flow. The idea of linear hybridization is to combine the best features of both methods by combining their results in an intelligent way.

The process of linear hybridization requires about twice as much work as the simpler artificial viscosity method. The advantage of the method which offsets this additional work is an improved resolution of flow discontinuities when a uniform grid is used. The goal of the linear hybridization approach is to limit the width of discontinuities to the width of the sharply rising part of the shock profile of a high-order difference scheme, while eliminating the attendant post-shock oscillations by appropriate blending with a low-order scheme. The low-order scheme controls the high-order scheme by determining which of the features produced by that scheme should be believed and which should not. The composite method produces very sharp discontinuities under favorable circumstances, but it is clearly limited in a fundamental way by the resolving power of the low-order scheme. This limitation is most apparent when the interaction of discontinuities must be computed. The smearing of the interaction region by the low-order scheme can then cause the

composite method to fail to recognize the presence of a passively advected discontinuity such as a contact discontinuity. This effect can be seen in the test problems presented below.

A more important limitation of the linear hybridization approach is the difficulty in devising appropriate weight factors for the low-order and high-order fluxes in the absence of any additional information about the inherently nonlinear physical processes which operate near flow discontinuities. A linear hybridization of first- and second-order difference schemes of the von Neumann and Richtmyer type has been devised by DeBar [30], and has been in use since 1968. However, this method was unable to remove all oscillations at flow discontinuities without a rather large amount of smearing. Another early linear hybridization is described by Harten and Zwas in [31]. A blending algorithm was later devised by Boris and Book [32] which yields sharp discontinuities without oscillations.

Boris and Book suggested that the blending of low-order and high-order fluxes be controlled by a monotonicity constraint. They set the weight factor of the low-order flux to zero unless this would cause an extremum to be introduced which would not be computed using the low-order flux alone. In such a case the low-order flux is given just enough weight so that no new extremum is generated. For simple advection in one dimension, this monotonicity constraint that no new extrema should be generated is a direct consequence of the differential equation, and the results obtained using this constraint are excellent. However, problems arise when this constraint is generalized in a straightforward way to systems of nonlinear differential equations, such as the equations of hydrodynamics. Then no such monotonicity constraint is implied by the differential equations, and use of such a constraint can lead to difficulties. In particular, a smooth region with strong gradients can be turned into a sequence of discontinuous jumps, with the appearance of a staircase. This effect can be seen in all the results of the ETBFCT scheme of Boris [33], which are presented below in Section VI. This possibility that smooth flow may be represented by staircases does not generally hinder the convergence of a linear hybridization. As the grid is refined the staircases which are generated converge to the smooth solution by developing greater numbers of treads of lesser height.

An interesting case of oversteepening of wave structures can occur when linear hybridization is used in separate 1-D sweeps of a 2-D calculation. Then oblique shock fronts may be given a staircase appearance. This effect can be seen in the 2-D results of the ETBFCT scheme which are presented in Section VI. The effect is caused by forcing too narrow a description of the shocks which arise in the component 1-D sweeps. These shocks arise from 1-D projections of the true shock jump, and hence the velocity jumps across them will be in error. In a scheme with very low dissipation, noise can be introduced into the smooth post-shock flow if these incorrect shocks in the 1-D sweeps are too narrow. This effect has been pointed out in the context of the random choice method by Colella [13]. There is some indication of this effect in the ETBFCT results presented below.

The technique of "flux-corrected transport," or FCT, was first discussed by Boris and Book in terms of an antidiffusion operation to remove excessive diffusion from a

low-order difference scheme. The reformulation of Boris and Book's method in terms of a linear hybridization is due to Zalesak [34]. Harten [35] has also devised schemes based on a point of view similar to antidiffusion, which he calls artificial compression. Because his schemes have not yet been engineered for general use, we have not included them in our study. It is interesting to note that the MUSCL scheme of van Leer, which will be discussed in the following section, is a linear hybridization technique when it is reduced to the case of simple advection. In that case the monotonicity criterion used by van Leer to blend first- and second-order fluxes is based on the same philosophy as that of Boris and Book, and it gives nearly the same results (cf. [4]).

The linear hybridization technique of Boris and Book has been generalized and improved by Zalesak [34]. His new schemes, unlike all those discussed here, perform a 2-D computation for a zone in a single step. Therefore they are not included in this study. However, a detailed investigation of several new linear hybrid schemes of that type in one dimension is presently under way, and the results will be presented in a subsequent article [36].

c. Godunov's Approach

In this approach a narrow representation of flow discontinuities is made possible by building into the numerical method a knowledge of the propagation and interactions of nonlinear waves. This knowledge is built into the method in the form of a Riemann solver. The Riemann solver computes the nonlinear interaction of two constant states of the fluid, and tells us what nonlinear waves emerge from this interaction. Nonlinearity is included in this way, because the Riemann problem is the only nonlinear flow problem simple enough to permit a solution to be computed at every zone interface for each time step. Convergence of the method on a fine grid is based on the assumption that, aside from the jumps at flow discontinuities, the changes in the flow variables within the zones of such a grid are small. The fluxes of conserved quantities computed for an appropriately chosen discontinuous flow model using a Riemann solver agree with those valid for smooth flow up to and including terms of second order in the size of the jumps at the discontinuity. Of course, when the flow actually is discontinuous, the Riemann solver yields a more reliable result than a calculation based upon a smooth flow model.

In all the schemes of this class which will be discussed in this article, the Riemann solver is incorporated into the method in the same way. These schemes begin by determining a set of interpolation polynomials describing the distributions of the flow variables within the zones of the grid. The data used to construct these interpolation functions consist of zone-averaged values of flow variables related to conserved quantities. To update these zone averages, fluxes of the conserved quantities, averaged over the time step, are required at the zone interfaces. The constant states used as input to the Riemann solver are chosen to represent the domains of dependence for a given zone interface which are swept out during the time step by the characteristic curves of different families. The information contained in these domains of depen-

dence is determined using the interpolation polynomials describing the internal zone structures. The order of these polynomials determines the spatial order of accuracy of the resulting numerical method. This sequence of operations is discussed in more detail in Section Vc.

The major disadvantage of this approach is the complexity introduced into a difference scheme through a Riemann solver. The simplest such scheme, Godunov's first-order method, requires about twice as much computer time per zone per time step as the second-order MacCormack scheme with an artificial viscosity. This cost in program complexity and computational effort per zone is more than offset by the benefit of high accuracy when a higher-order scheme involving explicit nonlinearity is used. This high accuracy results from the narrow structures with which discontinuities are represented by such a scheme, and from the excellent representation of discontinuity interactions which the Riemann solver provides.

Although one form of program complexity is introduced via the Riemann solver, another form of complexity is made unnecessary by the narrow shock structures which the Riemann solver allows. These narrow structures permit the use of uniform grids for describing all discontinuities whose jump conditions can be built into a Riemann solver. Thus in many applications the need for greatly distorted adaptive grids, with all the complexities they involve, is eliminated.

A disadvantage of Godunov's approach is that the Riemann solver it requires becomes more complicated when the equation of state cannot be represented by a simple gamma law or an isothermal assumption. In this case a subiteration is required in the Riemann solver (cf. [37, 14]). However, it should be noted that correspondingly more work is required with the artificial viscosity approach. In that case one must smear a shock out over a larger number of zones, so that the physical processes such as molecular dissociation which cause the effective gamma to change are properly computed within the shock structure. With the linear hybridization method it is not clear how such a complex shock should be computed.

The complexity introduced into a numerical method by a Riemann solver is really needed only near discontinuities. It is therefore natural to hybridize such a method with a set of simple difference equations which is equivalent to the full nonlinear method up to terms of some order in the jumps of the state variables within a zone. This much faster calculation may be performed for the overwhelming majority of zones in which these jumps are small. Because of the overhead involved in sorting the zones into the two categories, this procedure yields no gain in speed for Godunov's method on a CRAY-1 computer. For the much more complicated PPM scheme, which will be described in Section Vc, a factor of two in speed may be gained in this way on a CRAY-1. The reasons for these disappointingly small speed-ups have been discussed in [21]. On a more favorably organized computer, a speed-up factor of about three should be possible for the PPM scheme. It should be noted that the speed of an artificial viscosity method using local mesh refinement in space and time to treat discontinuities would be affected by these same considerations on a CRAY-1 computer.

IV. THREE TEST PROBLEMS

a. *Two Interacting Blast Waves*

This one-dimensional test problem was introduced in [21] to illustrate the strong relationship between the accuracy of the overall flow solution and the thinness of discontinuities on the grid. It involves multiple interactions of strong shocks and rarefactions with each other and with contact discontinuities. This problem is




FIG. 1. Contours of density on the space-time plane for the interacting blast wave problem discussed in Section IVa. Sixty contours equally spaced in $\log \rho$ are shown. The region near the collision of the two strong shocks is enlarged in the lower part of the figure. The lowest density contour is dotted.

extremely difficult to solve on a uniform Eulerian grid, although it poses no particular difficulty for a Lagrangian calculation.

The initial condition consists of three constant states of a gamma-law gas, with $\gamma = 1.4$, which is at rest between reflecting walls separated by a distance of unity. The density is everywhere unity, while in the leftmost tenth of the volume the pressure is 1000, in the rightmost tenth it is 100, and in between it is 0.01. Two strong blast waves develop and collide, producing a new contact discontinuity. This evolution is quite complex. A wave diagram depicting the contours of density in space-time for this problem is shown in Fig. 1. In Fig. 1b the region of strong wave interactions near the collision of the two blast waves is shown in detail. Shocks and contact discontinuities are easily located in Fig. 1 by the concentration of many contour lines along their paths.

In an effort to obtain the most accurate possible solution for the evolution of this flow we have used a special version of the scheme PPMLR described in Section Vc. This special version of PPMLR treats the flow as containing three distinct fluids, so that the two contact discontinuities which result from the initial pressure jumps can be accurately tracked. During the initial phases of the flow evolution in which no wave interactions occur PPMLR is not used and the flow is computed exactly. During the subsequent evolution the PPMLR calculation is aided by refining the computational mesh by a factor of 8 in both space and time in the five zones nearest to each flow discontinuity. The reference mesh for this calculation contains 3096 zones. These zones have a width $1/2400$ in most of the computational domain. However, the mesh is finer near the location of the collision of the two blast waves in order to resolve the very narrow rarefaction wave which is formed there. Thus for $0.64 < x < 0.68$ and for $0.74 < x < 0.81$ we have $\Delta x = 1/4800$, while for $0.68 < x < 0.74$ we have $\Delta x = 1/9600$.

In Fig. 2a the flow is shown at time 0.01. On the right, the initial pressure jump has sent a strong shock into the cold gas and a strong rarefaction into the hot gas next to the right-hand wall. This rarefaction is in the process of reflecting from that wall, producing a region of nearly constant pressure and density next to the wall. These nearly constant values steadily decline in time due to the linear expansion of this gas away from the wall. Between the strong shock and the rarefaction near the right wall are two constant states separated by a strong contact discontinuity. At the left in the figure a similar structure has developed, but this structure has evolved further because of the higher initial pressure near the left-hand wall. The rarefaction produced by this initial high-pressure region has already reflected off the left-hand wall completely. A constant state with zero velocity is left next to this wall, and the reflected rarefaction is interacting with the strong contact discontinuity near $x = 0.3$. This interaction has produced kinks in the pressure and velocity profiles near $x = 0.2$ and a kink in the pressure profile at the contact discontinuity. The kinks near $x = 0.2$ mark the edge of a rarefaction reflected from the contact discontinuity. The edge of the rarefaction transmitted into the dense slab is about to overtake the strong shock near $x = 0.34$.

In Fig. 2b the flow is shown at time 0.016. On the right, the rarefaction is still in




FIG. 2. The interaction of two blast waves as computed by a special version of the PPMLR scheme with a nonuniform grid of 3096 zones. See discussion in Section IVa.

the process of reflecting from the right-hand wall. The reflected rarefaction is just beginning to emerge from the region of nearly constant pressure and density next to the wall, where it is interacting with the incident rarefaction. On the left, the rarefaction reflected from the contact discontinuity is just about to reach the left-hand wall. The transmitted rarefaction moving to the right has overtaken the strong shock and weakened it. This shock therefore generates less entropy and this has caused the sudden change in the density slope just to the left of the shock. This kink in the density profile marks the location of the mass element which just passed through the shock when the transmitted rarefaction overtook it.

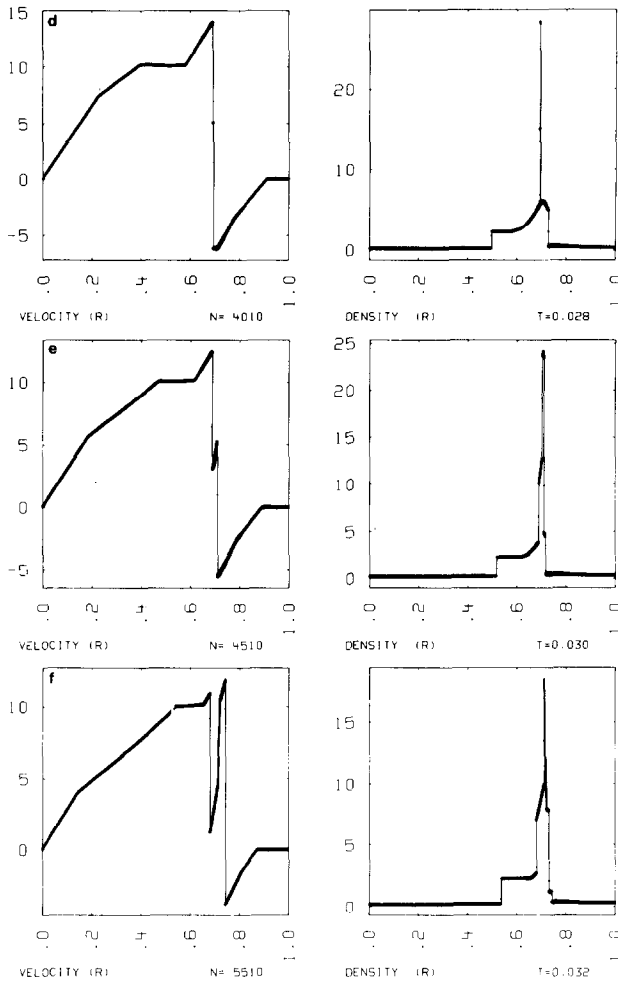


FIG. 2—Continued.

In Fig. 2c the flow is shown at time 0.026. On the right, the rarefaction has completely reflected from the right-hand wall and is beginning to interact with the contact discontinuity. A transmitted rarefaction has penetrated about one-third of the way into the dense slab. On the left the two rarefactions travelling in opposite directions in Fig. 2b have fully separated. Between them are two constant states separated by a contact discontinuity. The rarefaction on the left is in the process of reflecting from the left-hand wall. The reflected rarefaction is just emerging from the region of nearly constant pressure and density set up near the wall. The isolated rarefaction travelling to the right continues to weaken the shock, producing an entropy gradient in the shocked gas.




FIG. 2—Continued.

In Fig. 2d the flow is shown at time 0.028. The two blast waves are just now colliding. The peak density at this time does not reach the limiting value of 36, because the reflected shocks are significantly weaker than the incident ones. Although it is not apparent from the figure, the sharp density spike is indeed well resolved in the calculation, and the peak density shown is accurate to about a percent. In Fig. 2e the flow at time 0.030 is shown. Already by this time a rarefaction between the two reflected shocks has brought the peak density down to about 24. A great deal of structure is visible in the high-density regions. Starting at the left we have the reflected shock travelling to the left down a steep density and entropy gradient. This gradient continues behind the shock until a newly created contact discontinuity is reached at the location of the earlier shock collision. To the right of this contact discontinuity a rarefaction moving to the left can be discerned from the pressure profile and it can be seen in the wave diagram in Fig. 1b. Further to the right is the reflected shock moving to the right down a density gradient. Finally, we come to the contact discontinuity at the edge of the original right-hand dense slab.

In Fig. 2f the flow is shown at time 0.032. The reflected shock moving to the right has now moved into the hot gas near the right-hand wall. The interaction of this shock with the contact discontinuity has generated a very strong rarefaction wave moving to the left which is just now reaching the central contact discontinuity. In Fig. 2g, at time 0.034, this rarefaction is in the process of interacting with the central

contact discontinuity. The increase in the density directly behind the reflected shock moving to the left is caused by the entropy gradient in this gas. This reflected shock is just moving into a region of gas with constant entropy. Therefore the density gradient behind this shock will change sign in the near future. Note also that the rarefaction reflected from the left-hand wall is now beginning to penetrate the dense slab of cool gas.

In Fig. 2h the flow is shown at time 0.038. This is the time at which all our test runs in Section VIa are displayed. It should be clear from the above description of the flow evolution that this problem involves the multiple interactions of strong nonlinear continuous waves and discontinuities of various kinds. Because much of the important interaction takes place in a small volume, this problem is very difficult to compute on a uniform Eulerian mesh. It is therefore a good test for Eulerian difference schemes designed for flow problems involving strong shocks and their interactions in one dimension.

b. *A Mach 3 Wind Tunnel With a Step*

This two-dimensional test problem was introduced more than a decade ago in the paper by Emery [38] which compared several difference schemes. Since only the best of those schemes is of much interest nowadays, it is worthwhile to reuse the problem here to compare more modern difference methods. Results for this problem obtained with the original MUSCL code written by Woodward are presented by van Leer in [5]. Results obtained with a single-step Eulerian MUSCL are presented by Colella in [8]. Also results obtained with several methods are briefly presented in [7]. Here we will present much more comprehensive results. Obviously this problem has proven to be a useful test for a large number of methods and a large number of years.

The problem begins with uniform Mach 3 flow in a wind tunnel containing a step. The wind tunnel is 1 length unit wide and 3 length units long. The step is 0.2 length units high and is located 0.6 length units from the left-hand end of the tunnel. The tunnel is assumed to have an infinite width in the direction orthogonal to the plane of the computation (i.e., "slab symmetry" is assumed). At the left is a flow-in boundary condition, and at the right all gradients are assumed to vanish. The exit boundary condition has no effect on the flow, because the exit velocity is always supersonic. Initially the wind tunnel is filled with a gamma-law gas, with $\gamma = 1.4$, which everywhere has density 1.4, pressure 1.0, and velocity 3. Gas with this density, pressure, and velocity is continually fed in from the left-hand boundary. These admittedly artificial initial conditions make the problem very easy set up, a necessary feature of a useful test problem.

Along the walls of the tunnel reflecting boundary conditions are applied. The corner of the step is the center of a rarefaction fan and hence is a singular point of the flow. In the papers referenced above, nothing special was done at this singular point. Consequently the flows shown in those papers are seriously affected by large numerical errors generated just in the neighborhood of this singular point. These errors cause a boundary layer of about one zone in thickness to form just above the




FIG. 3. The time evolution of the Mach 3 wind tunnel problem discussed in Section IVb. The results were obtained with the PPMLR scheme using a uniform grid with $\Delta x = \Delta y = 1/80$. The contours of density are shown at time intervals of 0.5 up to time 3. At time 4, the contours of density, pressure, $A = p/\rho^{\gamma}$, v_y , v_x , and $(v_x/c) - 1$ are also plotted. In each plot 30 equally spaced contours are shown, with the contour for the lowest level or for any negative level drawn as a dotted line.

step in the wind tunnel. Shocks then interact with this boundary layer, and the qualitative nature of the flow in the tunnel is altered more or less dramatically, depending upon the difference scheme and the grid which is used. The sensitivity of various difference schemes to the treatment of the flow near the corner of the step can be assessed by comparing the results presented here with those in [7], for which no special boundary condition was applied at the corner.




FIG. 3—Continued.

In this paper we will attempt to minimize numerical errors generated at the corner of the step so that all the various schemes will tend to converge to the same flow in the limit of a very fine grid. To accomplish this we apply an additional boundary condition near the corner of the step. In the first row of zones above the step we will reset the first four zones starting just to the right of the corner of the step; in the row above we will reset the first two zones. In these zones we reset the density so that the entropy has the same value as in the zone just to the left and below the corner of the step. We also reset the magnitudes of the velocities, not their directions, so that the sum of enthalpy and kinetic energy per unit mass has the same value as in the same

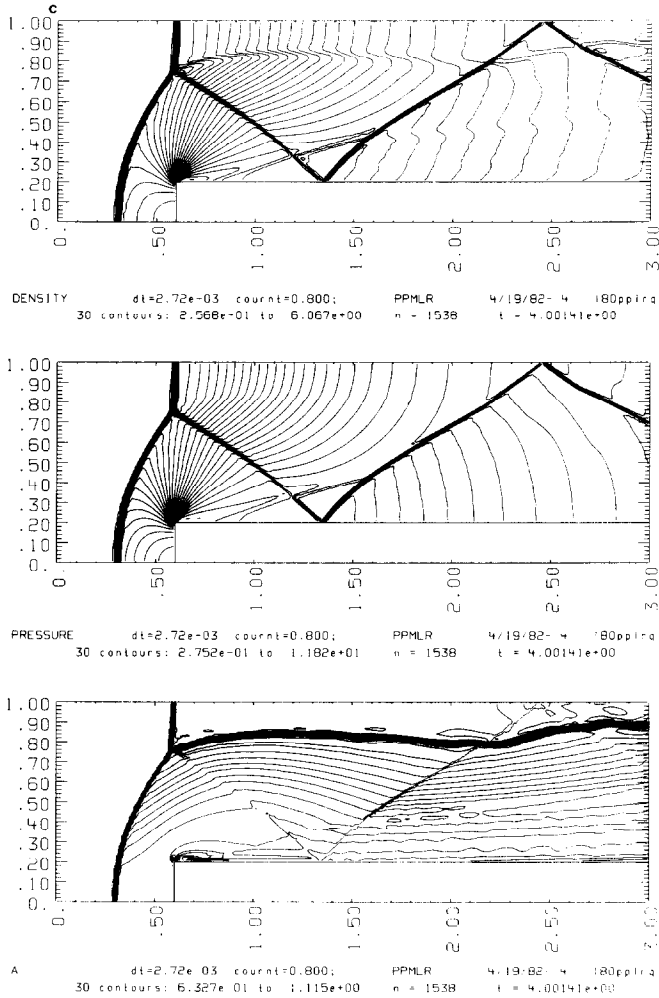


FIG. 3—Continued.

zone used to set the entropy. This condition is based on the assumption of a nearly steady flow in the region near the corner. It is clearly inappropriate at the very outset of the calculation. These conditions remove the grossest errors generated near the corner, but of course large errors in the flow direction there are bound to remain. These errors may be the cause of an overexpansion observed at the corner in all the runs, although similar effects occur in wind tunnel experiments of this type using real, viscous air.

The time evolution, up to time 4, of the density distribution in the wind tunnel is displayed in Fig. 3. The flow at time 4 is still unsteady. A steady flow develops by

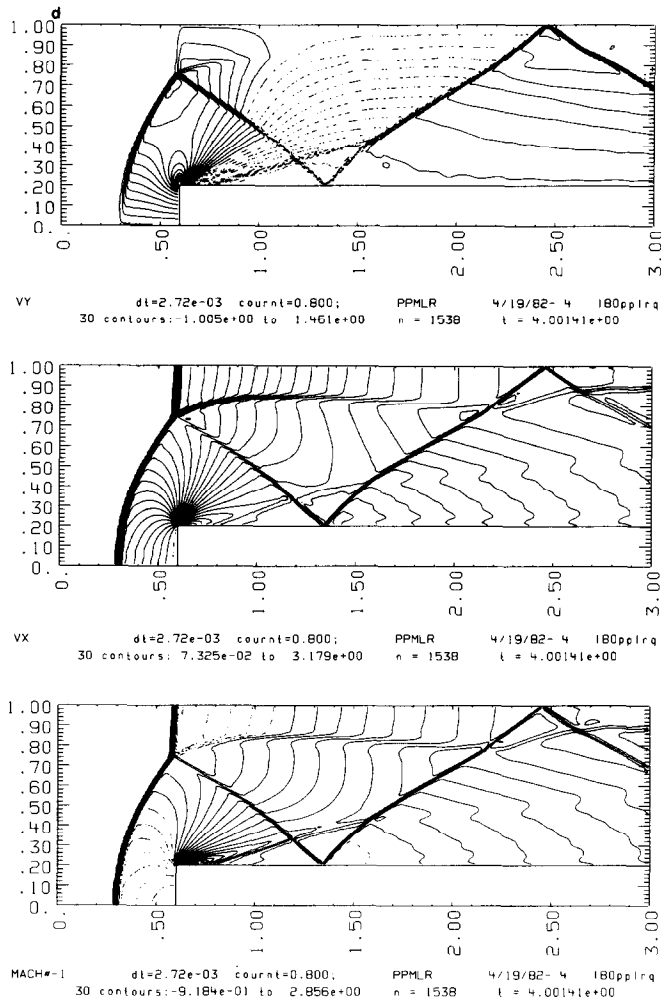


FIG. 3—Continued.

time 12 and is shown in Ref. [5]. Because this steady flow has very little structure we will focus here on the more interesting flow at early times. The results shown in Fig. 3 were obtained with our best scheme, PPM, using a very fine uniform grid of 80×240 zones. On this grid the PPM scheme appears to have “essentially” converged to its limit solution, so we may treat the results of the calculation as very nearly correct. Nevertheless, numerical effects may still remain, such as the overexpansion at the corner of the step and the weak oblique shock caused when the overexpanded flow finally strikes the upper surface of the step. These features do not appear to change in position or strength as the grid is refined, but this does not necessarily




FIG. 4. The self-similar flow resulting from the double Mach reflection of a Mach 10 shock from a wall is shown. These results were obtained with the PPMDE scheme using a uniform grid with $\Delta x = \Delta y = 1/120$. Contours are plotted for (a) density, (b) pressure, (c) $A = p/\rho^\gamma$, (d) v_y , (e) $v_x - 11.547$. In each plot 30 equally spaced contours are shown, with the contour for the lowest level or for any negative level drawn as a dotted line.




FIGURE 4 (continued)

imply that they are real. Another numerical effect is the mild Kelvin–Helmholtz instability (“water-wave” or “flag-waving” instability) of the contact discontinuity along the upper wall. This instability is visible at times 2.5 and 3. The instability is physical rather than numerical, but it is triggered by very small oscillations in the entropy behind the Mach shock. These oscillations are a numerical effect which is subsequently amplified by a physical instability.

c. Double Mach Reflection of a Strong Shock

This test problem was inspired by experimental and numerical studies of reflections of planar shocks in air from wedges [39–41, 37]. The flow can be set up experimentally by driving a shock down a tube which contains a wedge. At first the simple planar shock meets the walls of the tube at right angles, but when one wall begins to slope a complicated shock reflection occurs. A self-similar flow develops at this point which can be parameterized, for a given gamma-law gas, by the Mach number of the incident shock and the angle with which it encounters the reflecting wall of the

wedge. Our test problem involves a Mach 10 shock in air ($\gamma = 1.4$) which initially makes a 60° angle with a reflecting wall. The undisturbed air ahead of the shock has a density of 1.4 and a pressure of 1.

The reflecting wall lies along the bottom of the problem domain, beginning at $x = 1/6$. The shock makes a 60° angle with the x axis and extends to the top of the problem domain at $y = 1$. The short region from $x = 0$ to $x = 1/6$ along the bottom boundary at $y = 0$ is always assigned values for the initial post-shock flow. This boundary condition forces the reflected shock to be "attached" to the reflecting wall. The left-hand boundary is also assigned values for the initial post-shock flow, and at the right-hand boundary, at $x = 4$, all gradients are set to zero. The values along the top boundary are set to describe the exact motion of the initial Mach 10 shock. In setting up the problem this way we have attempted to make this test as easy as possible to run on a standard hydrodynamics code. Also, by tilting the incident shock rather than the reflecting wall we have avoided the complicated issue of how properly to model a boundary oblique to the mesh. Our intent here is to compare numerical methods for hydrodynamics, not methods for implementing boundary conditions.

The flow at time 0.2, computed by the PPM code on a fine grid with $\Delta x = \Delta y = 1/120$, is displayed in Fig. 4. Only the region between $x = 0$ and $x = 3$ is displayed in the figure, although the grid continues to $x = 4$. The parameters of this test problem correspond to double Mach reflection of the shock at the wall. Two Mach stems form, with two contact discontinuities. The second contact discontinuity is extremely weak and is more easily noticed by the velocity jump across it than by the density jump. The second Mach shock is rather weak, and it dies out entirely by the time it reaches the contact discontinuity from the first Mach reflection. This variation of the strength of the second Mach shock is very difficult to compute accurately. At the point where the first contact discontinuity approaches the reflecting wall the flow of the denser fluid is deflected by a pressure gradient built up in the region. The result is that a jet of the denser fluid is formed which shoots to the right along the wall. This feature of the flow thus bears a similarity to the classic problem of a "shaped charge" [42]. To compute the formation of this jet properly is extremely difficult. A further computational difficulty is presented by the region bounded by the second Mach shock, the curved reflected shock, and the reflecting wall. In this region there is very little vertical motion. Therefore numerical methods whose dissipation vanishes with zero flow velocity have a tendency to oscillate here.

The results shown in Fig. 4 are marred by a small numerical error. The problem was initialized with a pure shock jump which was given the narrowest possible description on the grid. Just such a precise shock jump was also set at the upper boundary at each time step. Because the PPM method, like any other method discussed here, must spread this shock out somewhat, a "starting error" is produced during the short time required for the numerical representation of the shock to reach a steady shock profile. This starting error takes the form of an overheated region of about two zones in width. It is visible only in the plot of the vertical component of velocity, v_y . It appears as a thin strip extending from the upper right-hand corner of each contour plot and reaching the bottom reflecting wall near $x = 2$. A similar

starting error mars all the computed results for this problem which are presented in Section VI.

During the course of this study, a number of other test problems were investigated. The three problems presented here involve most of the computational difficulties encountered in this broader set of tests. We include all three problems here, because it has been a common occurrence in the course of our work that a technique which works well on two of these problems will fail on the third. This behavior can be observed for a number of the difference schemes presented below.

V. REPRESENTATIVE NUMERICAL METHODS

a. *Artificial Viscosity*

The artificial viscosity approach is employed by so many difference schemes that it would be impossible to present them all here. Instead, we have chosen representatives from the two major classes of such schemes in general use. To represent the schemes based on fully zone-centered data we have chosen MacCormack's method [43], a variation on the Lax-Wendroff scheme [44]. Despite the stabilizing influence of dissipative truncation errors, MacCormack's scheme must be augmented by some additional viscosity in order to control numerical oscillations near shocks and contact discontinuities. We have added the artificial viscosity formulated by Lapidus [45].

We will represent the other major class of artificial viscosity methods by the scheme BBC. The class BBC represents makes use of the spatially staggered grid introduced by von Neumann and Richtmyer [1], and it also uses their artificial viscous pressure, q . It is really only the Lagrangian step of the BBC scheme which is of the von Neumann and Richtmyer type. This step is followed by a remap step from a Lagrangian grid back to the original Eulerian one. This remap is derived from the MUSCL algorithm of van Leer [4], adapted to the staggered grid. Thus the remap step of BBC is actually a linear hybrid of first- and second-order advection schemes. The Lagrangian step of the BBC scheme was devised by DeBar in 1968 for use in the KRAKEN code [30]. It uses a two-step formulation to do away with the time-staggered grid of von Neumann and Richtmyer. It simplifies their approach still further by centering velocity data at zone interfaces rather than at zone corners. The 2 - D Lagrangian step used in KRRAKEN was reformulated into component 1 - D sweeps for use in the original BBC code [46]. This method was later made second-order accurate by the addition of a MUSCL remap step modified and adapted to the staggered grid by Woodward. The present BBC scheme is described in detail in the Appendix.

b. *Linear Hybridization*

The linear hybridization approach will be represented by the single method ETBFCT of Boris [33] based on the ideas discussed by Boris and Book in [32]. This method is used currently in the code FAST2D, and results for a problem similar to

our double Mach reflection problem which were obtained with this method have been published recently [40]. A more thorough investigation of the antidiffusion approach than can be given here is now under way and will be published separately [36]. We have implemented the ETBFCT scheme in separate 1 – D sweeps in a sequence *xyyx*. This results in second-order accuracy for 2 – D computations (cf. [47]).

Boris recommends that ETBFCT be run at a Courant number less than one-half. Ikeda and Nakagawa have derived a slightly different and more restrictive limit on the Courant number for the SHASTA FCT algorithm [48]. They demonstrate that the method can generate small rarefaction shocks when their limit on the Courant number is exceeded. The results presented below were run at a Courant number of 0.4. The results are not significantly improved when the Courant number is lowered to 0.1. Any further reduction of the Courant number would be impractical.

c. Godunov's Approach

The approach of Godunov will be represented here by three difference schemes. The first, a single-step Eulerian version of Godunov's method, has been described by Godunov [49]. Simplified versions of Godunov's method are often used, as, for example, in the review of Sod [12]. Therefore the reader should note that the version used here employs the full nonlinear Riemann solver described by Godunov [49], modified for second-order convergence of the iterations as described by van Leer in [5] and simplified in its treatment of rarefactions as described in [13]. This version of Godunov's method gives the most accurate representation of shocks and their interactions.

A simple second-order scheme of this class is also studied here. It is a simplified version of the MUSCL scheme described by van Leer in [5], and it borrows some techniques from an advection scheme described by van Leer in [4]. It performs the Eulerian calculation in a single step, analogous to the MUSCL scheme described by Colella in [8]. The method described in [5] cannot handle very strong shocks, so our MUSCL has been revised to do so. Stephen Hancock [50] has independently devised a different means of handling strong shocks in MUSCL, and this method has been discussed by van Leer in [51]. Our simple MUSCL scheme is essentially the same as the scheme described in [8], except that slopes of the variable distributions within zones are obtained simply by differencing the average values of the variables in neighboring zones and applying van Leer's monotonicity constraint. We have included this scheme in our study to show what sort of performance may be expected from a relatively uncomplicated extension of Godunov's approach to second-order accuracy.

The most accurate scheme presented here involving a nonlinear Riemann solver is PPM, the piecewise-parabolic method. Previous versions of this scheme have been discussed in [7] and in [21], and the present version will be described in detail in a subsequent article [9]. Three key ingredients are responsible for the accuracy of the PPM scheme: (1) a Riemann solver to compute fully nonlinear wave interactions, (2) high-order interpolation techniques to properly describe the information present in the

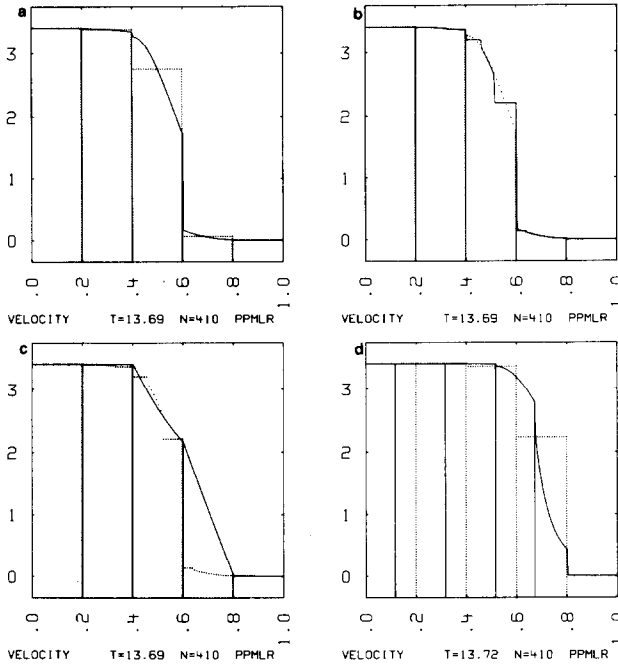


FIG. 5. A single time step in the computation of a strong shock using the PPM scheme. Velocities are displayed for 5 zones representing various stages in the computation. See the discussion in Section Vc of the text. (a) Interpolation parabolae (solid lines) are constructed from the original zone-averaged data (dotted lines). The shock is only one zone wide. (b) Data are obtained for Riemann problems which will give time-averaged fluxes at zone interfaces in a Lagrangian step of the calculation. The data (solid lines) are obtained by averaging the interpolation parabolae (dotted lines) over the domains of dependence of each zone interface. (c) The solution of the Riemann problems (represented by dotted lines) gives effective velocity gradients within the zones (solid lines) which compress the Lagrangian zones. (d) The compressed Lagrangian zones are shown by solid lines with internal structures inferred from zone-averaged data by the construction of new interpolation parabolae. New averages are computed using these parabolae for the original Eulerian zones (dotted lines). The shock has moved $\frac{3}{4}$ of a zone width and is again one zone wide.

domain of dependence of each zone interface, and (3) special monotonicity constraints and discontinuity detectors which help to keep discontinuities as sharp as possible without unphysical oscillations or unphysical steepening of the local gradients.

The concerted action of the above three key ingredients of PPM is best appreciated by looking at a particular example. In Fig. 5 the operation of the PPM scheme is shown in detail for the case of an isolated, fairly strong shock traveling to the right into a uniform gas at rest. The pressure and density ahead of the shock are both unity. The gas has a gamma-law equation of state with a gamma of 1.4. Behind the shock the pressure is 16, the density is 4.4091, and the velocity is 3.4056. The Eulerian zones are 0.2 length units wide, so that in the time step of 0.034056, the shock moves to the right three-quarters of a zone width.

The process of computing a single time step for five zones near the steady shock structure is displayed in four stages in Fig. 5. In Fig. 5a the dotted horizontal lines show the zone-averaged values of the velocity in the zones at the beginning of the time step. From this data the PPM scheme constructs interpolation parabolae which represent the distributions of the velocity within the individual zones with respect to a mass coordinate. These interpolation parabolae are shown as solid lines in Fig. 5a. In regions of smooth flow the resulting velocity distribution is continuous at zone interfaces, but in the case shown here discontinuities in the distribution have been introduced at some zone interfaces near the shock. These discontinuities are the result of two constraints: first that the interpolated parabolae within the zones must give back the original data when averaged over the zone masses, and second that they must also be monotone increasing or decreasing, with all values lying in the range defined by the zone-averaged velocities in the particular zone and in its two nearest neighbors. The representation of the shock on the numerical grid is too sharp to permit a continuous velocity profile under these constraints.

The version of the PPM scheme displayed in Fig. 5 is PPMLR, which performs an Eulerian calculation by breaking it up into a Lagrangian step in which pressure forces accelerate the zones and a subsequent remap step in which the effects of the fluid motion relative to a fixed grid are calculated. PPMLR moves the zone interfaces in the Lagrangian step, and hence computes zone compressions, by calculating approximate time-averaged interface velocities. This calculation is performed in three steps. First the domains traversed by sound waves reaching a zone interface during the time step are determined. For this purpose the sound waves are assumed to travel with the appropriate zone-averaged sound speeds. These domains of dependence are then represented by constant states of the gas which are obtained by averaging the interpolated spatial structure over the domains. The constant velocities so obtained and the spatial extents of the domains of dependence are shown by the solid lines in Fig. 5b. The dotted lines show the original interpolated velocity structures for comparison.

The purpose of averaging over the detailed structure contained in the domains of dependence of the interfaces is to allow the nonlinear hydrodynamic interaction of these domains to be computed by solving Riemann's shock tube problem. The results of this computation can be inferred from the solid lines in Fig. 5c. These display velocity distributions which are linear with respect to a mass coordinate and which join the time-averaged interface velocities computed from the Riemann problems. These velocity distributions may be regarded as the effective time-averaged distributions which compress the zones in the shock structure during the time step. The dotted lines in Fig. 5c show the distribution from Fig. 5b for comparison. For this case of an isolated shock, the effect of solving Riemann's problem has evidently been to equate the time-averaged interface velocity to the average velocity in the "upstream" domain of dependence. This choice of the interface velocity has kept the velocity gradient in the center zone in Fig. 5c from being too large and has therefore prevented an overcompression of that zone during the time step. Of course, the monotonicity constraint which caused the large initial velocity jump at the right-hand

interface has also played an important role in limiting the effective velocity gradient used to compress this zone.

The final remap step of PPMLR is illustrated in Fig. 5d. The solid lines in the figure show interpolation parabolae representing the velocity distributions within Lagrangian zones at the end of the time step. These velocity distributions are used to compute mass-weighted velocity averages within the original Eulerian zones. These averages are indicated by the dotted lines in the figure. Of course, mass-weighted averages are computed so that momentum is conserved in the process. Note that although the shock has moved three-quarters of a zone width in a single step, the representation of the shock remains sharp without introducing post-shock oscillations.

The PPMLR scheme illustrated in Fig. 5 has been used to compute the flow in the Mach 3 wind tunnel which is shown in Fig. 3. For this particular problem additional features of the method in treating shock discontinuities are important which do not come into play in the example in Fig. 5. When strong shocks travel slowly relative to the grid, as do all the shocks in the wind tunnel problem after time 2, the shock structures must be broadened to about 2 zone widths in the worst cases in order to prevent post-shock oscillations of about 1 to 2% in amplitude from developing. Such post-shock oscillations cannot be effectively damped behind the shocks, because the wavelengths of the oscillations are many zone widths in the case of slowly moving shocks. They must therefore be prevented by broadening the shock structure somewhat. This broadening is performed by two methods. First, the interpolated parabolic structures are flattened somewhat. Second, the grid to which the remapping is performed is jiggled slightly about its original position as the shock passes, so that the effective speed of the shock relative to the grid cannot vanish.

A similar broadening of the shock can be obtained by flattening the interpolated structures in shocks and by adding a small Lapidus artificial viscosity to the method, with a coefficient of 0.1—a coefficient a tenth the size of that used with MacCormack's scheme in the 2-D problems presented in this study. This simpler approach to shock broadening, however, gives a diffusive error term in smooth flow which is third-order small, while the other approach of jiggling the grid gives an error which is fifth-order small. These approaches are described in detail in [9], and they may be compared by studying the results of the double Mach reflection problem shown in Figs. 4 and 9f. These results were both obtained using a single-step Eulerian formulation of PPM, which we will refer to as PPMDE. In Fig. 4 the grid has been jiggled slightly near shocks, while in Fig. 9f a Lapidus viscosity with a coefficient of 0.1 has been used. The single-step Eulerian formulation of PPM is constructed on the same principles as the method with separate Lagrangian and remap steps, but is somewhat more complicated to describe. It is discussed in detail in [9].

Recently much attention has been given to approximate methods for solving Riemann's problem [51–54, 56]. In the schemes studied here we have used the approximate Riemann solver described in [13]. This approximates rarefactions by rarefaction shocks. We have used a Newton iteration with this method, so that more than two iterations are never required. For all but very strong shocks, a single iteration is sufficient. In the interacting blast wave problem, two iterations have been

used. In the single-step Eulerian codes, special provision for the spreading of a rarefaction fan must be made in the Riemann solver under certain special conditions (cf. [13]). The overall process is very fast on a CRAY-1 computer, and in the PPM scheme very little extra speed can be gained by simplifying this Riemann solver further. We believe, however, that approximate Riemann solvers like those devised by Harten and Lax [53] and by Roe [54] will be in much greater demand when larger systems of conservation laws, such as the equations of magnetohydrodynamics, are attacked.

VI. COMPUTED RESULTS FOR THE THREE TEST PROBLEMS

a. *Two Interacting Blast Waves*

This one-dimensional problem provides the most convenient basis for a thorough comparison of the various numerical methods. We have run the problem with each method using 9 different uniform grids, of 100, 150, 200, 300, 400, 600, 800, 1200, and 2400 zones. Each result at time 0.38 has been compared to a converged result obtained using a special version of PPMLR with a grid of 3096 zones, as described earlier in Section IVa. This special version tracks the two contact discontinuities which emerge from the initial pressure jumps by treating this as a flow containing three distinct fluids. This version of PPMLR also refines the grid automatically by a factor of 8 near all discontinuities in the flow and also near the two reflecting boundaries. This code yields a solution for the interacting blast wave problem which is as close to the exact solution as is necessary to accurately measure the errors produced by the difference schemes considered here. For the solution obtained with each scheme on each grid the fractional error, ε , was computed:

$$\varepsilon = \frac{1}{N} \sum_{i=1}^N |\langle \rho \rangle_i - \langle \rho \rangle_{i \text{ exact}}|.$$

Here N is the number of zones in the coarse grid, $\langle \rho \rangle_i$ is the average density in zone i of the coarse grid calculation, and $\langle \rho \rangle_{i \text{ exact}}$ is an appropriate average over values from the special 3096-zone run of PPMLR. For each run the number of seconds of CPU time on the CRAY-1 was measured. This time includes no input-output or plotting costs, only raw computation. It does include the computation of an equation for a transverse velocity component (which was always set to zero). Thus speeds for 2-D calculations are precisely half of the 1-D speeds. These 1-D speeds, and the integrated fractional errors, are tabulated in Table I.

To one familiar with only linearized error analysis, the results in Table I will come as a shock. In the last column a convergence rate is tabulated. It is derived from a one-sided difference approximation to the logarithmic derivative of the error with respect to the zone width. Of the schemes in the table, only Godunov's scheme is formally first-order accurate. All the others are formally second-order accurate.

TABLE I

Scheme	Courant No.	Zones	Time steps	CPU sec.	Fractional error	Convergence rate
MacCormack	0.8	200	400	0.278	0.261	0.32
		300	604	0.595	0.209	0.49
		400	806	1.05	0.181	0.53
		600	1216	2.33	0.143	0.55
		800	1624	4.11	0.114	0.66
		1200	2446	9.19	0.0862	0.72
BBC	0.8	200	332	0.456	0.196	0.43
		300	490	0.962	0.167	0.43
		400	652	1.69	0.145	0.44
		600	970	3.71	0.112	0.57
		800	1288	6.48	0.0920	0.65
		1200	1926	14.5	0.0717	0.64
ETBFCT	0.4	200	810	1.37	0.227	0.51
		300	1216	2.91	0.179	0.55
		400	1618	5.12	0.130	0.81
		600	2426	11.2	0.103	0.80
		800	3234	19.8	0.0787	0.72
		1200	4848	44.0	0.0606	0.76
Godunov	0.8	200	378	0.588	0.330	0.15
		300	568	1.27	0.293	0.24
		400	760	2.28	0.274	0.27
		600	1148	5.05	0.241	0.28
		800	1536	8.93	0.215	0.35
		1200	2320	20.1	0.182	0.40
MUSCL	0.8	200	408	1.36	0.162	0.46
		300	608	2.93	0.123	0.65
		400	808	5.17	0.0996	0.70
		600	1208	11.4	0.0723	0.77
		800	1610	20.1	0.0553	0.85
		1200	2412	44.9	0.0401	0.85
PPMDE	0.8	200	414	1.65	0.0990	0.61
		300	616	3.52	0.0639	1.05
		400	818	6.20	0.0467	1.08
		600	1222	13.7	0.0303	1.08
		800	1626	24.2	0.0218	1.10
		1200	2432	53.7	0.0144	1.07
PPMLR	0.8	200	310	1.63	0.0760	0.96
		300	454	3.44	0.0476	1.14
		400	598	5.99	0.0376	1.02
		600	888	13.1	0.0264	0.85
		800	1174	23.0	0.0205	0.87
		1200	1746	50.8	0.0133	0.98

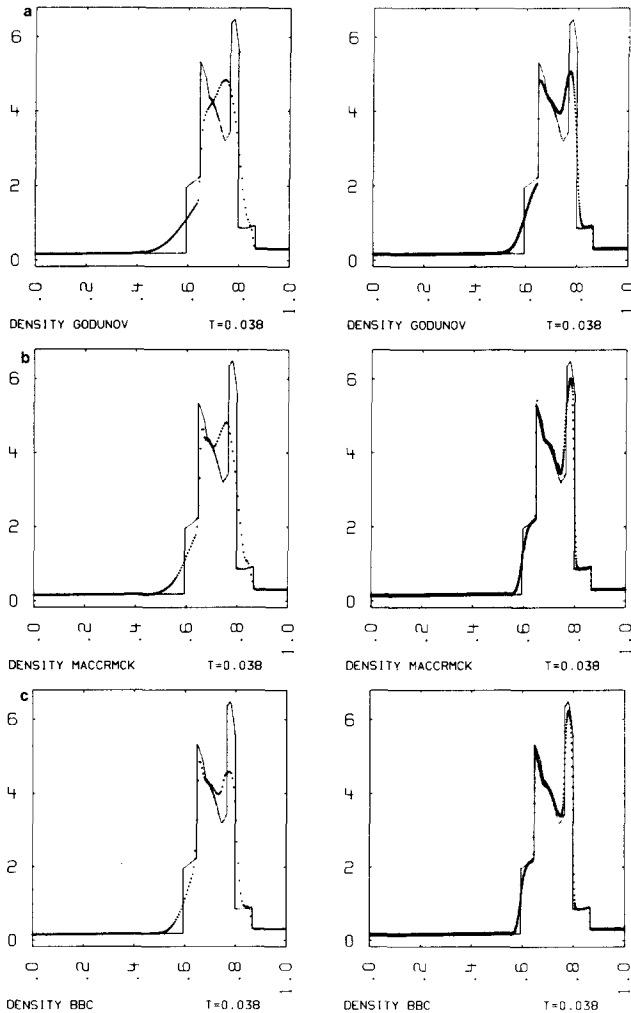


FIG. 6. Results for the interacting blast wave problem of Fig. 2 using 7 different numerical methods—(a) Godunov's method, (b) MacCormack's method, (c) BBC, (d) ETBFCT, (e) MUSCL, (f) PPMDE, and (g) PPMLR. Results are shown as dots for grids of 200 and 1200 zones. The solution obtained with the special version of PPMLR on a grid of 3096 zones is drawn for comparison in each figure as a solid line. Godunov's scheme gives unacceptable smearing of the three contact discontinuities while the shocks are only two zones wide. MacCormack's scheme gives broader shocks, with widths of three zones, while the contact discontinuities are sharper, but not sharp enough. BBC gives very similar performance to MacCormack's scheme on this problem. ETBFCT gives shocks only one zone wide, but it introduces a spurious contact discontinuity at the site of a discontinuity in the density slope. MUSCL performs very well on this problem, giving shocks one zone in width and contact discontinuities adequately resolved on the 1200-zone grid. The two formulations of PPM, PPMDE, and PPMLR give excellent and very similar results. Even on the 200-zone grid the two outer contact discontinuities are only one zone in width. Detailed comparisons of accuracy and CPU time for these runs can be made from the data in Table I.




FIG. 6—Continued.

However, the error from Godunov's method starts off by converging toward zero as $(\Delta x)^{0.15}$ and on the finest grids it is converging as $(\Delta x)^{0.4}$. This is very poor performance for a nominally first-order scheme. In fact, only the error of the PPM scheme converges to zero as fast as Δx . All the other schemes show less than linear convergence. Clearly, a linear error analysis based upon the assumption of smooth flow is entirely inappropriate to this sort of nonlinear problem.

Convergence is slow for this problem because the errors are mainly introduced at the discontinuities. In fact, our definition of the fractional error ε implies that in the limit of an extremely fine grid ε can converge to zero no faster than Δx . No matter how accurately we describe the smooth flow, large errors must occur at discontinuities which are spread over distances proportional to Δx . The computed results are displayed in Fig. 6 for the grids of 200 and 1200 zones. The solid lines show the special 3096-zone PPMLR results, while those of the various schemes are indicated as dots (one dot per zone). Both PPMDE and PPMLR have essentially converged on the 1200-zone grid. However, these are the only schemes to do so. The discontinuities which appear in the plots of Fig. 6 are, from left to right, the contact discontinuity which originally formed at $x = 0.1$, the shock which originally formed at $x = 0.9$, the contact discontinuity which formed when the two shocks collided, the contact discontinuity which originally formed at $x = 0.9$, and the shock which originally formed at $x = 0.1$. The hardest of these to represent are the contact discontinuities on the left and in the middle. Most of the error is generated in these features. The ETBFCT scheme also encounters difficulty at a discontinuity in the density slope. This discontinuity in slope marks the mass element which had just passed through the shock from the left-hand blast wave when the rarefaction wave reflected from the left-hand wall first reached that shock.

The results in Table I clearly indicate that the PPM scheme delivers the greatest accuracy at the least cost. Interpolating in the table we find that to match the accuracy of the 200-zone run with PPMLR the other schemes require more computer time by the following factors: 1.6 for PPMDE, 6.2 for MUSCL, 7.4 for BBC, 8.0 for MacCormack's scheme, and 13.5 for ETBFCT. The numbers of zones required are: 256 for PPMDE, 563 for MUSCL, 845 for ETBFCT, 1093 for BBC, and 1431 for MacCormack's scheme. Godunov's scheme cannot match the accuracy of PPMLR on the coarsest grid, even if it is given twelve times as many zones. Because the PPM schemes perform the entire nonlinear computation for all zones, the time comparisons given above would all grow by an additional factor of 2 or so if the PPM schemes were programmed to evaluate equivalent but much simpler difference equations in the smooth parts of the flow.

Due to the presence of strong contact discontinuities in this test problem it was necessary to augment the Lapidus viscosity in the calculations using MacCormack's scheme. In the Lapidus viscosity the diffusive flux at a zone interface, $i + \frac{1}{2}$, is multiplied by a coefficient proportional to the absolute value of the velocity difference of the neighboring zones, $|u_i - u_{i+1}|$. We have added to this a term $|c_i - c_{i+1}|/10$, where c_i is the speed of sound in zone i . This term damps the Gibbs oscillations which would otherwise result near contact discontinuities which move through the grid.

The timings for PPMLR and BBC are aided by the larger time steps which the splitting of these schemes into separate Lagrangian and remap steps allows, particularly for this problem. For these schemes the time step is limited by the maximum of the fluid velocity and the sound speed rather than by the maximum of their sum or difference. We should also point out that ours is not the fastest version of MUSCL conceivable. We have followed the approach of Ref. [8], which requires two solutions to Riemann's problem to update a zone, while the PPM methods require only a single Riemann solution. A MUSCL scheme can be formulated along the lines of PPM, and such a scheme would run a bit faster. We have presented results of the older formulation here in order to illustrate why the additional features of the PPM formulation are desirable. Finally, it should be noted that time requirements for multidimensional runs of the sort presented here would be spread over a much greater range for these schemes, because the computation time depends upon higher powers of $(1/\Delta x)$.

b. *A Mach 3 Wind Tunnel with a Step*

Results for all the representative difference schemes on the wind tunnel problem described in Section IVb are shown in Fig. 7. For each scheme the density distribution is shown at time 4 as computed on three different grids; $\Delta x = \Delta y = 1/20$, $\Delta x = \Delta y = 1/40$, and $\Delta x = \Delta y = 1/80$. Thirty equally spaced contours are shown in each plot, with the extreme contour levels given in the figure legend. If the extreme values in the computed zone averages are ρ_{\min} and ρ_{\max} , then the extreme contour levels are $(\rho_{\min} + \rho_{\max})/2 \pm (29/60)(\rho_{\max} - \rho_{\min})$. The density distribution is the most difficult to compute for this problem, because of the weak contact discontinuity caused by the Mach reflection of the bow shock at the upper wall. Results are shown for three grids to give a visual impression of the manner and rate of convergence of each scheme. For such 2-D problems, accuracy is a rather subjective quantity, and a full display of results on different grids allows each reader to judge for himself how the various schemes perform. All of these schemes are insensitive to the Courant number used, so long as it is small enough for stability. Most of the runs were made at a Courant number of 0.8. Runs with smaller Courant numbers required these time step reductions for stability.

It is natural to expect the least accuracy for this problem to result from using the only first-order scheme in our study, Godunov's scheme. A careful look at Fig. 7 shows this expectation to be fulfilled, but not by the large margin we might have anticipated. In some respects, such as the position of the shock reflection at the upper wall, the Godunov results on the coarsest grid are even superior to those of MacCormack's second-order-accurate scheme. In fact, the Godunov results are roughly comparable in accuracy to those of MacCormack's scheme on all three grids. Thus, even the rates of convergence of these methods are comparable. This example illustrates the limited usefulness of linearized error analysis for nonlinear problems of this type.

The surprisingly good performance of Godunov's scheme on this particular flow

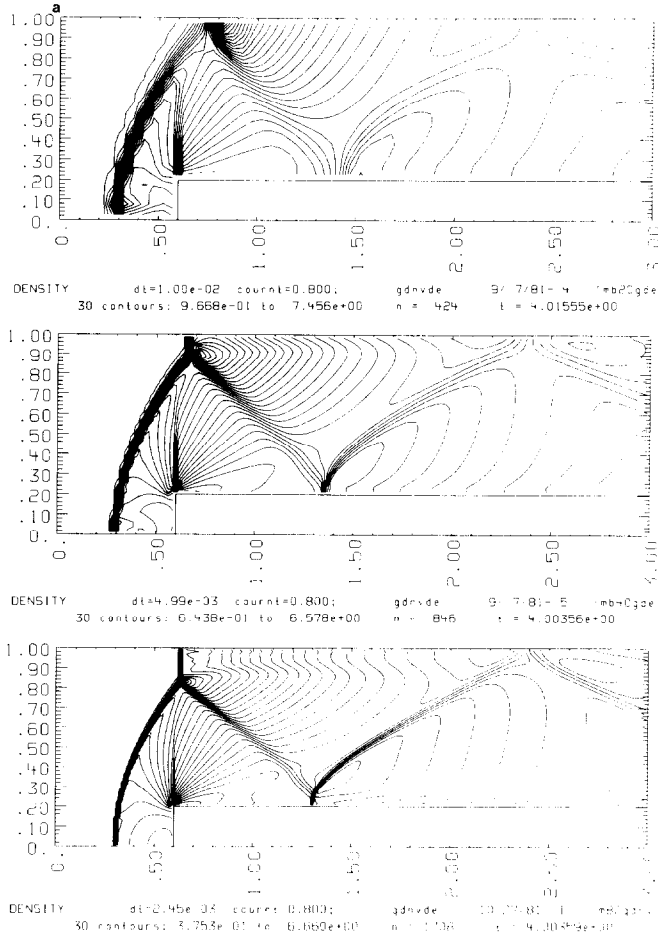


FIG. 7. Results at time 4 are plotted for the Mach 3 wind tunnel problem displayed in Fig. 3 and discussed in Section IVb. Density contours are shown for three grids, with $\Delta x = \Delta y = 1/20$ at the top, $\Delta x = \Delta y = 1/40$ in the middle, and $\Delta x = \Delta y = 1/80$ at the bottom.

FIG. 7a. Godunov's method. The general shape and position of the shocks is not correct, even on the finest grid. The Mach stem at the upper wall is absent on the coarsest grid and is too short and too far downstream on both finer grids. The shocks are very thin where they are strong, nearly stationary, and nearly aligned with the mesh. A numerical instability of $2 - D$, strong, stationary shocks is evident on all grids near the bottom wall and on the finest grid behind the Mach stem. This instability is more clearly illustrated in Fig. 8. The contact discontinuity emerging from the three-shock intersection is not smeared out as much as we might expect because this contact discontinuity is nearly aligned with the mesh. A strong numerical boundary layer along the step in the duct causes a second, spurious Mach reflection which is not diminished as the mesh is refined. Finally, unphysical rarefaction shocks are produced on all grids near the sonic line just above the corner of the step.




FIG. 7b. MacCormack's method. Despite the higher formal order of accuracy of this method, the general shape and position of the shocks it computes for this problem are roughly equivalent to those computed with Godunov's method. The Mach stem is absent on the coarsest grid and too short and too far downstream on the finest grid. The shocks have a fairly uniform width of 2 to 3 zones, and they are accompanied by mild pre- and post-shock oscillations. There is considerable smearing of the contact discontinuity generated at the Mach reflection near the upper wall. Also, a numerical boundary layer along the top of the step in the duct causes a spurious Mach reflection at the step, but this Mach stem is only two zones long. The numerical boundary layer may be identified downstream from this Mach stem by the sudden kinks in the density contours. These are most evident on the finest grid. Finally, the weak shock originating near the corner of the step in Fig. 3c and the contact discontinuity formed when this shock merges with that reflected from the step are not resolved by MacCormack's method, even on the finest grid.

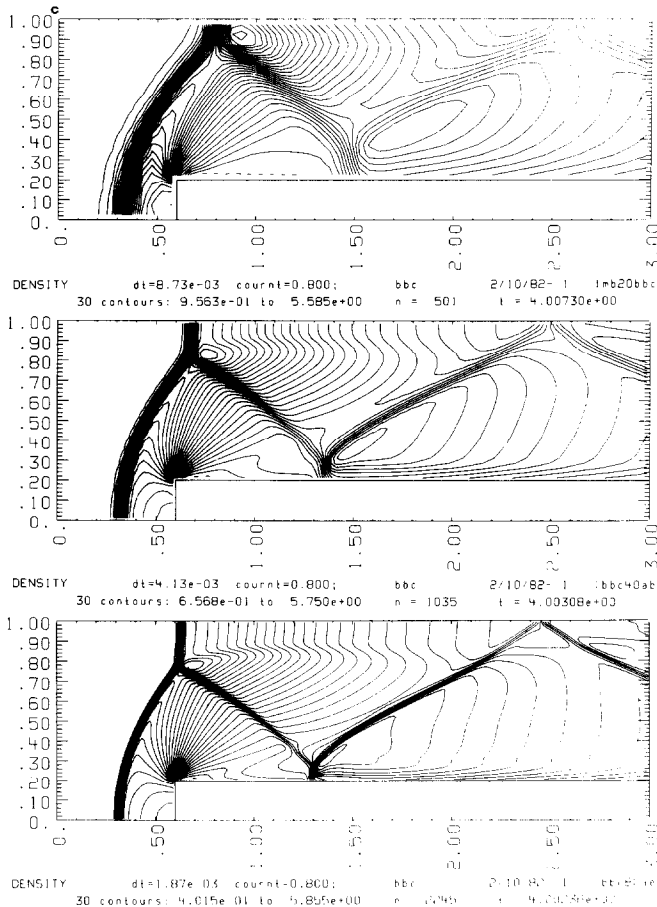


FIG. 7c. BBC. This method gives a more accurate representation of the general shape and position of the shocks than do either Godunov's or MacCormack's methods. The Mach stem is just visible on the coarsest grid, and it is almost correct on the finest grid. BBC gives roughly the result in this respect that MacCormack's method would on grids with $\sqrt{2}$ times as many zones in each direction. Therefore BBC is roughly a factor of 1.2 more efficient than MacCormack's method in solving this problem. The BBC shocks are about two zones wide. They are somewhat thinner than the MacCormack shocks, as is evident from the coarse grid results, and there are no pre- or post-shock oscillations. However, the shock has a small influence one zone upstream from the principal two-zone jump, as can be seen on the coarse grid. There is a bit less smearing of the contact discontinuity than with MacCormack's method because of the MUSCL remap step in BBC. However, the numerical boundary layer above the step is very strong, and spurious Mach reflections are produced. This is largely a result of the difficulty of applying an appropriate boundary condition at the corner of the step when a staggered grid is used.

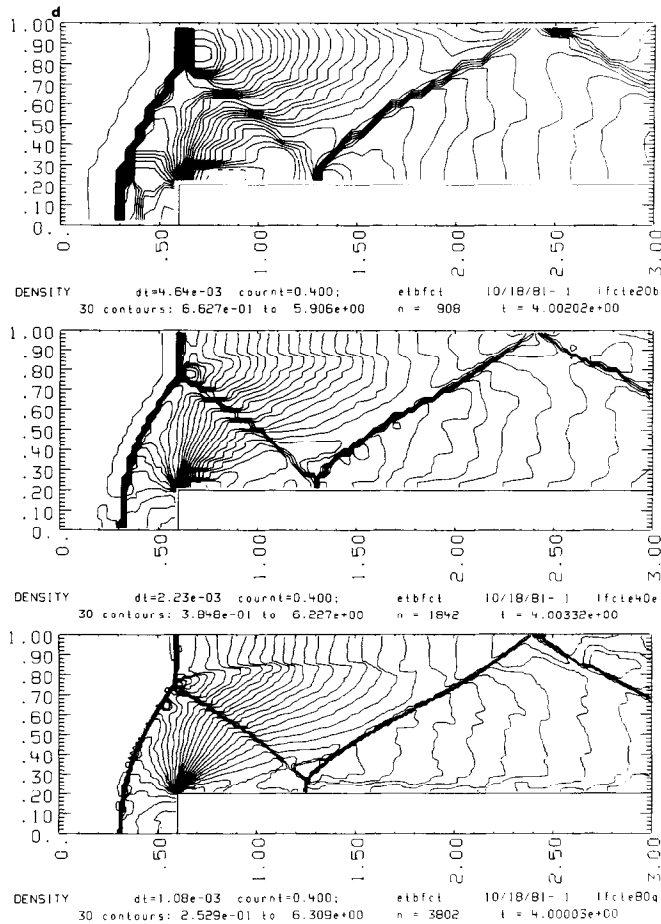


FIG. 7d. ETBFCT. Most notable is the staircase appearance of the contour lines. Possible causes for this are discussed in the text. Aside from this jitter, the accuracy of the general shock shape and position is extremely good. The Mach stem has the correct length and position on the middle grid. The shocks are extremely narrow except for broad preshock regions of substantial velocity and pressure undershoots which are delineated by low-level density contours upstream from the leftmost shocks. The contact discontinuities are very sharp when present. However, the contact discontinuity arising from the coalescence of the weak shock from the corner of the step and the shock reflected from the step appears in the wrong place on the finest grid, and, inexplicably, it appears on the coarse grid despite the absence of the weak shock which causes it. In addition, on the coarse grid the contact discontinuity near the upper wall does not appear, despite the length of the Mach stem. Finally, a numerical boundary layer along the top of the step in the duct produces spurious Mach reflections there which do not diminish as the mesh is refined.

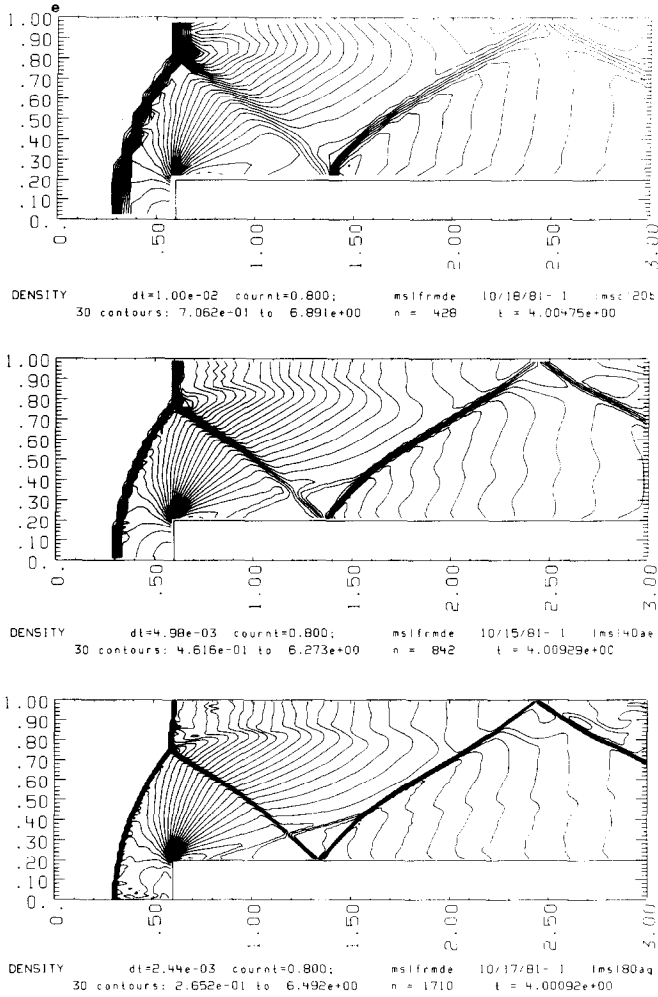


FIG. 7e. Single-step Eulerian MUSCL. The general position and shape of the shocks are as accurate as those computed by ETBFCT, while the jitter and spurious glitches produced by that scheme are absent. The shocks are extremely narrow, only a zone wide, without over- or undershoots. In fact, the shocks are too narrow, so that the numerical instability of strong, stationary shocks nearly aligned with the mesh develops near the bottom wall and behind the Mach stem. The entropy contour plot in Fig. 8 shows this instability more clearly. The contact discontinuities are spread somewhat, but both appear on the middle grid and the one near the upper wall can be seen on the coarse grid. Also note that the effects of the numerical boundary layer along the top of the step in the duct are very weak. The weak shock from the corner of the step is resolved on both finer grids.




FIG. 7f. Single-step Eulerian PPM, or PPMDE. The general position and shape of the shocks are slightly less accurate than those computed by MUSCL or ETBFCT, but essentially all numerical noise has been eliminated from the solution. The shocks are slightly broader than for the MUSCL calculation, but they are still essentially one zone wide. The length and position of the Mach stem are correct on the middle grid. The spreading of the contact discontinuities is very slight, and both are resolved on the middle grid. The effects of the numerical boundary layer along the top of the step in the duct are also very slight. Also, the weak shock from the corner of the step is resolved on both finer grids. The general accuracy of the computed flow is roughly equivalent to a BBC calculation with two times as many zones in each direction and to a calculation with MacCormack's scheme using $2\sqrt{2}$ times as many zones in each direction. Thus PPMDE is more efficient in solving this problem than BBC by roughly a factor 3.6, and relative to MacCormack's scheme this factor is roughly 4.4.




FIG. 7g. PPMLR, or PPM formulated as a Lagrangian step followed by a remap. These results are essentially equivalent to those computed with PPMDE. These calculations require 18% more computer time, so the costs of the calculations are roughly equivalent as well. The shocks are slightly broader than for PPMDE where they are strong and steady, that is, just ahead of the step and at the Mach stem, but they are thinner where they are weak and in motion, that is, at the regular reflections on the step and at the upper wall. The contact discontinuity near the upper wall is smoother than for PPMDE because this unstable slip line has suffered smaller numerical perturbations at earlier times in the calculation (see Fig. 3b). On the finer grids the effects of the numerical boundary layer along the top of the step are also slightly smaller than for PPMDE. Efficiency factors in relation to BBC and MacCormack's scheme are roughly 3.0 and 3.7, respectively.

problem is the result of its very narrow representation of nearly steady shocks. By time 4 nearly all the shocks in this problem are moving very slowly. Because the Riemann solver permits Godunov's method to represent a steady shock at a zone interface as a pure discontinuity, shocks which are nearly steady are very thin. When they are nearly at a zone interface they are represented by almost a pure jump, and otherwise one zone makes up the entire internal structure of the shock. This can be seen especially well on the coarse grid results. The plotting routine linearly interpolates between zone-averaged data. Thus when the shock is shown as a dark vertical bar, many contours have crowded in between two zone-averaged data values and the shock consists of nearly a pure jump. As we look upward along the bow shock starting at the bottom boundary, we can see the shock location shift from one zone interface to another, becoming periodically thicker and thinner as it does so. This behavior is most marked near the lower wall, because this part of the shock is strongest and moving most slowly.

The narrow steady shock structures produced by Godunov's method are responsible for its accuracy on this flow problem. However, these narrow structures also lead to difficulties. There is a tendency for noise to be emitted from the places where the shock relatively suddenly shifts over from one column of zones to another. This noise can be seen near the lower wall, where the flow velocity is so small that the noise is not damped effectively. The MUSCL scheme shows an even greater tendency to produce this sort of noise, because it has far less damping in regions of smooth flow. The noise is most noticeable on the finest grids.

To further elucidate the cause of this noise, the adiabatic constant, $A = p/\rho^\gamma$, is plotted in Fig. 8 for the MUSCL run on the finest grid. The fluctuations in the entropy generation at the shock at the bottom wall are fairly substantial. These are accompanied by fluctuations of about the same size in the velocity jump at the shock.




FIG. 8. Contours of the quantity $A = p/\rho^\gamma$, which is a function of the entropy, for the MUSCL computation of the wind tunnel using the fine grid $\Delta x = \Delta y = 1/80$. Errors introduced at an unphysical kink in the Mach shock near the upper wall are amplified by a physical instability, the Kelvin-Helmholtz instability. This numerical effect arises when a strong, narrow shock moves very slowly relative to the grid and is nearly aligned to it. This effect can be eliminated by increasing the dissipation of the numerical scheme near such a shock, so that the shock is broadened slightly.

At the shock near the upper wall, the source of these fluctuations at a sudden shift in the shock position from one column of zones to another is clearly visible. The slanted regions of alternately high and low entropy behind this Mach shock show how this sudden shift in the shock location has moved vertically along the Mach shock in the past. The periodicity of these entropy perturbations is perfectly matched to that of a Kelvin–Helmholtz instability (“water-wave” instability) of the contact discontinuity near the upper wall. The perturbations caused by the motion of the kink in the Mach shock have apparently been amplified by the physical instability of this slip line. In the results of both PPM schemes, PPMDE and PPMLR, the nearly steady shocks have been broadened slightly so that hardly any noise is produced behind them. As we have noted in Section Vc, this is accomplished by flattening of the internal zone structures for zones inside such shocks and, for PPMLR, by jiggling the grid slightly, and for PPMDE by the addition of a Lapidus artificial viscosity with a coefficient of 0.1.

Like MUSCL, the ETBFCT scheme produces very thin shock structures. From the above discussion we might expect them to create considerable noise in the smooth part of the flow. This appears to be the case. Particularly in the coarsest run, bunchings of the density contours can be traced to kinks in the shocks. However, the noise in the ETBFCT results has a different appearance than that produced by the MUSCL code. It is not limited to regions of small flow velocity in either the x or the y direction, but instead the noise is pervasive. It appears that either the noise is not sufficiently damped in the smooth flow or it is generated there as well. Especially on the coarser grids there is a marked tendency to oversteepen the gradients in the centered rarefaction fan, so that the density contours take on a staircase appearance.

Another feature of the thin shocks in the ETBFCT runs is not so apparent from the density plots shown in Fig. 7. A single density contour well in front of the main shock structure marks the extent of the preshock region which has been unphysically influenced by the shock. In this region substantial negative pressures are calculated which must be reset to floor values. The reason for this error is that the monotonicity constraint which determines the shock structure only guarantees monotone profiles for the variables ρ , ρu_x , ρu_y , and ρE , where E is the total energy per unit mass. Therefore the pressure profile need not be monotone. Neither MUSCL nor PPM has this difficulty. This may be because these schemes apply monotonicity constraints to the variables ρ , p , u_x , u_y instead. It should be noted that the computational costs for ETBFCT and MUSCL are practically identical. The MUSCL scheme therefore computes equally thin shock structures as ETBFCT at the same cost and without so much noise generation. We see in Fig. 7 that all noise errors are removed by the PPM schemes and the treatment of contact discontinuities is improved at an additional cost of about 20% in computer time.

The BBC scheme represents a combination of a Lagrangian step using an artificial viscosity and a remap step using a modification of the MUSCL linear hybridization of first- and second-order advection schemes. The artificial viscosity in the Lagrangian step spreads the shock enough to eliminate the problems associated with thin shocks which we have been discussing. Also the modified MUSCL remap helps

to keep the contact discontinuity sharp. In contrast Godunov's scheme and MacCormack's scheme spread the contact discontinuity more as it moves away from the three-shock intersection point. ETBFCT keeps this contact discontinuity very sharp, as does PPM. In the case of PPM this sharp contact discontinuity results from the use of parabolae to describe internal zone structures coupled with a special discontinuity detection algorithm. These features keep contact discontinuities as sharp as shocks in PPM runs.

The slowest feature of the computed flow to converge to its true value is the length and location of the Mach stem. This is directly related to the time at which the Mach stem first forms in the calculation. Because the formation of a Mach stem is a threshold phenomenon, this formation will occur too late if the shock is spread out over too wide a region. For Mach 3 wind tunnels with steps, our particular test problem is near the transition from a steady flow with a regular shock reflection at the upper wall (this occurs for steps of somewhat lesser height) and flows which involve Mach reflection. For our test problem a Mach stem forms at about time 1.5 and moves very slowly upstream along the upper wall. If the shocks are smeared out too much in the computation, the Mach stem forms late and does not reach its proper position at time 4.

Even on the finest grid MacCormack's scheme does not obtain the proper Mach stem location. Therefore despite the great speed with which this method can update a zone, it requires an impractically large number of zones, time steps, and computing time to compute the correct answer to this simple test problem. Godunov's scheme is not practical either for solving this problem. It computes the Mach stem location well, but smears the contact discontinuity over several zones. The Godunov results are also marred by bits of noise generated at kinks in the shocks and also by an unphysical rarefaction shock embedded in the centered rarefaction fan. This little rarefaction shock occurs where the Mach number of the flow is unity, so that one characteristic speed vanishes in the x -pass of the computation. The rarefaction shock is then allowed because of a subtle property of the scheme related to the fact that the error in the computed flux at a zone interface where the flow is sonic is one order smaller than everywhere else. The ETBFCT and MUSCL schemes compute the correct Mach stem location on the coarsest grid and with the least computer time. However, the thin shocks which make this possible cause underisable noise to enter the solution. Significantly less noise is generated by the MUSCL scheme, but if a smooth result is desired, the PPM schemes will produce the correct Mach stem location in the least time and on the coarsest grid.

It is interesting to note that two formulations of PPM, with a single Eulerian step and with a Lagrangian step followed by a remap, produce almost identical results. The two schemes also require almost identical amounts of time to update a zone. For multifluid problems the formulation with a separate Lagrangian step is much more convenient, while the single step Eulerian form is better suited to problems with complicated spatially dependent source terms. The two schemes can apparently be used interchangeably according to the nature of a particular application without concern about loss of accuracy or computing speed.




FIG. 9. Results at time 0.2 are plotted for the double Mach reflection problem displayed in Fig. 4 and discussed in Section IVc. Density contours are shown for three grids, with $\Delta x = \Delta y = 1/30$ at the top, $\Delta x = \Delta y = 1/60$ in the middle, and $\Delta x = \Delta y = 1/120$ at the bottom.

FIG. 9a. Godunov's method. The jet formed by the double Mach reflection is unresolved on even the finest grid. The strong shocks are about two zones wide with no under- or overshoots; however, the weak shock generated at the kink in the main reflected shock is extremely broad.

c. Double Mach Reflection of a Strong Shock

Results for the various difference schemes on the double Mach reflection problem described in Section IVc are shown in Fig. 9. Again, we show results on three grids — $\Delta x = \Delta y = 1/30$, $\Delta x = \Delta y = 1/60$, and $\Delta x = \Delta y = 1/120$ — to permit the reader to

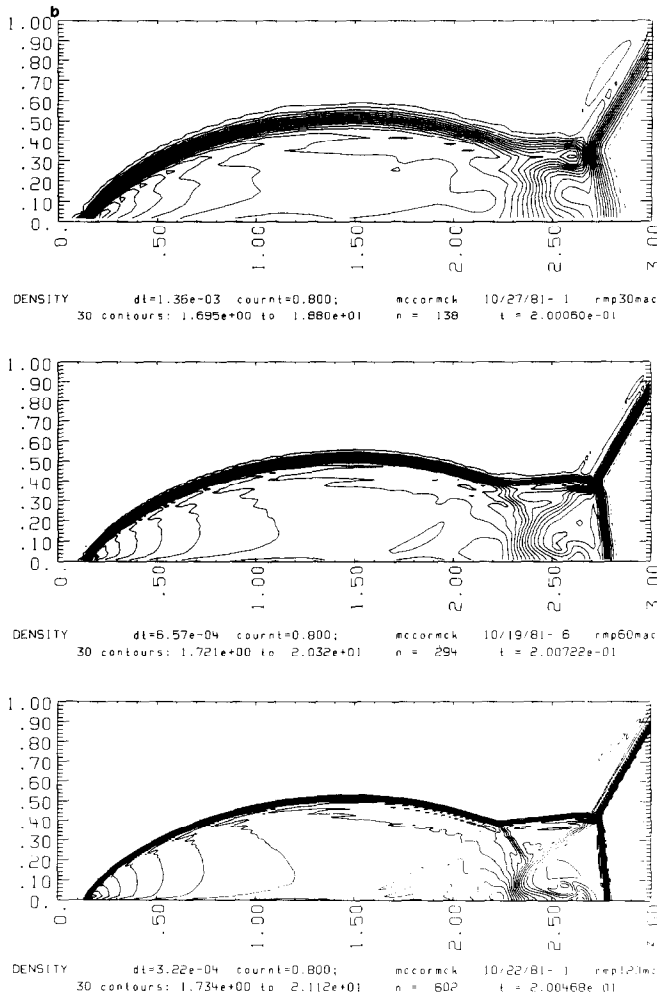


FIG. 9b. MacCormack's method. These results are far superior to those of Godunov's method. The jet produced by the double Mach reflection is described fairly well on the finest grid and adequately on the middle grid. The results on the coarse grid are poor. The shocks are fairly broad, 2 to 3 zones, with pre- and post-shock oscillations. The oscillations are particularly severe where the shock is moving slowly and is nearly aligned with the mesh. An unphysical structure near the main three-shock intersection is most noticeable on the coarse grid.

make his own assessment of the performance of these schemes. Because of the importance of contact discontinuities in this flow, density contour plots are shown in Fig. 9. The most complicated structures in the flow are in the rather small region of the double Mach reflection. The structure of the jet formed near the lower wall demands the most resolution from the difference scheme. Even on the finest grid this jet is less than five zones wide. The weakening of the second Mach stem as it approaches the




FIG. 9c. BBC. These results are roughly equivalent to those of MacCormack's method, except that the oscillations near the shocks and the unphysical structure near the main three-shock intersection have been eliminated. The weak shock generated at the kink in the main reflected shock is quite broad, due to the first, "linear" term in the artificial viscosity in Eq. (A12).

contact discontinuity is also difficult to resolve. Unlike the Mach 3 wind tunnel problem, this double Mach reflection contains both steady and unsteady structures. The curved reflected shock stretching across the problem domain is moving rapidly at its right-hand end and is not moving at all at its left-hand end. Therefore difference schemes which generate noise behind stationary shocks will oscillate toward the left-hand end of this shock.

As for the wind tunnel problem of the previous section, the least accurate results

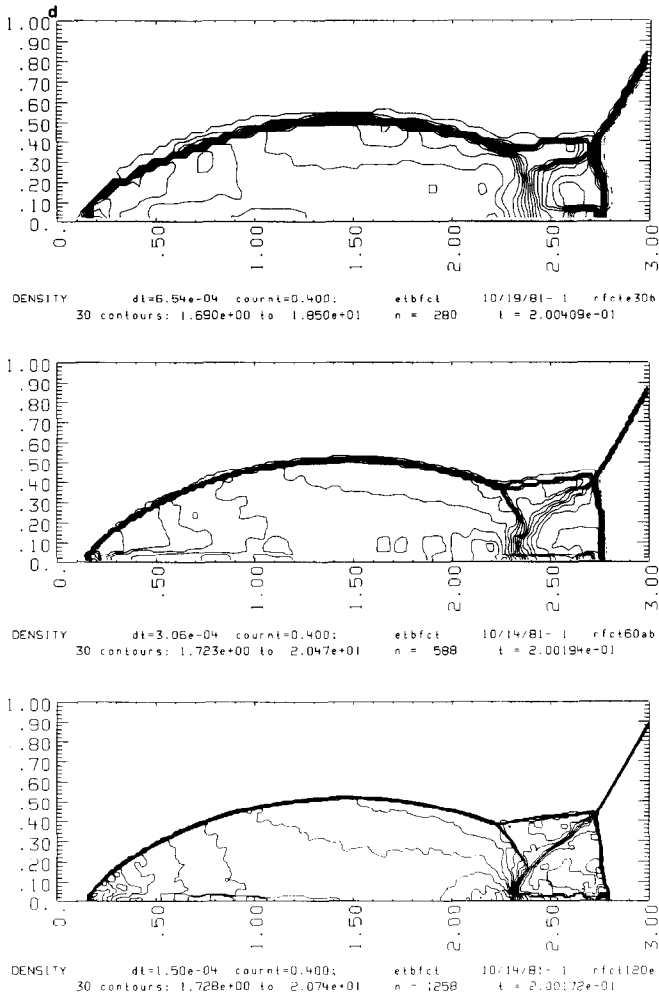


FIG. 9d. ETBFCT. Most noticeable is the staircase appearance of the contour lines. Possible causes of this are discussed in the text. The shape of the principal Mach stem appears to be incorrect on all grids. On the finest grid, this error may be caused by the incorrect description of the jet, a piece of which has broken off and struck the Mach shock. The contact discontinuity which turns to form this jet is considerably broadened on the middle grid and emerges from the three-shock intersection at the wrong angle on the coarse grid. Nevertheless, the weak shock at the second Mach reflection is quite well described.

are produced by Godunov's first-order scheme. However, for this double Mach reflection problem the Godunov scheme does not do as well in relation to MacCormack's scheme or BBC. The reason for this poorer performance is the unsteady nature of this problem. This causes the shock structures to be broader and the computation of their interactions less accurate. Also, the contact discontinuity is

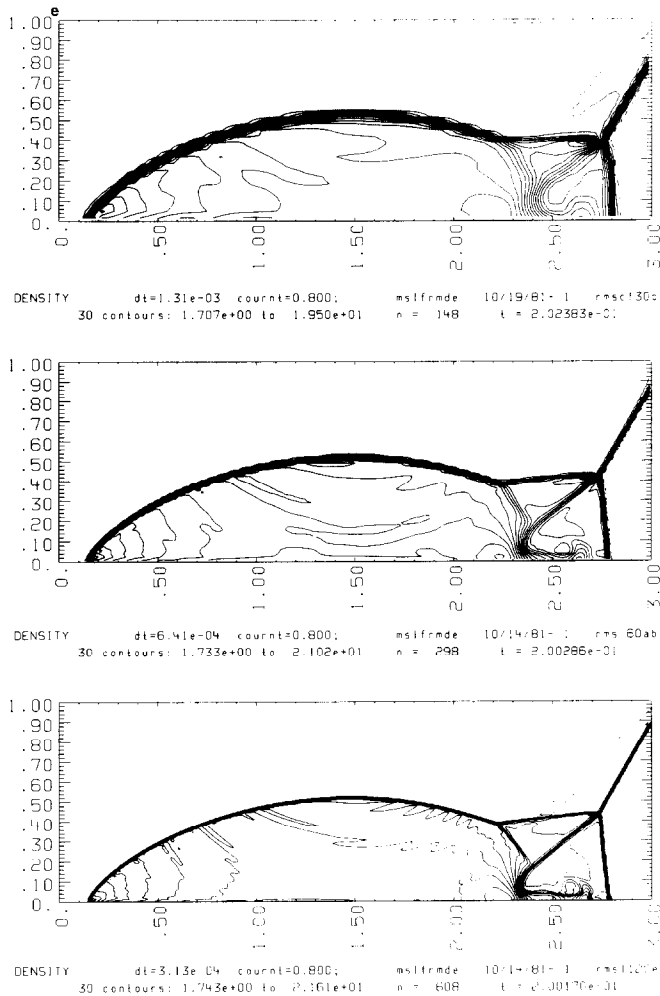


FIG. 9e. Single-step Eulerian MUSCL. Except for the presence of some numerical noise, these results are superior to those of MacCormack's method, BBC, and ETBFACT. The jet is resolved on the middle grid and can be clearly seen on the coarse grid as well. The weak shock at the second Mach reflection is also well resolved on all grids. The shocks are extremely thin, but this thinness has permitted a numerical instability to generate where the shocks move slowly and are nearly aligned with the mesh. A hint of this instability is also visible where the Mach stem strikes the lower boundary on the finer grids.

not aligned with one set of grid lines, so it is greatly diffused as it moves. Although MacCormack's scheme oscillates behind the nearly steady portion of the curved reflected shock, it produces overall results superior to Godunov's scheme. Its second-order accuracy permits a much better tracking of the moving contact discontinuity.

The performance of the BBC scheme is better on this problem relative to

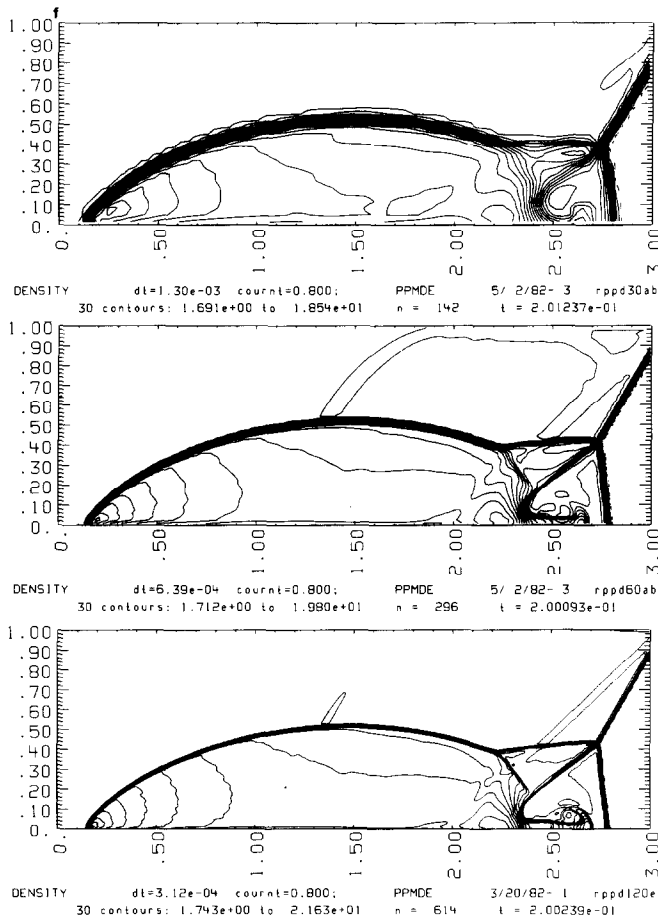


FIG. 9f. Single-step Eulerian PPM, or PPMDE. These results are an improvement over the MUSCL results in two respects—the numerical instability of slowly moving shocks nearly aligned with the mesh has been essentially eliminated and the contact discontinuity which bounds the jet is sharper. The results are marred by starting errors caused by the specification of pure jumps to describe the initial shock and to trace its motion along the upper boundary. These starting errors are diminished by the use of the more elaborate dissipation mechanism of the run shown in Fig. 4. However, starting errors are present to some degree in the results of all schemes. Comparison with the results for BBC and MacCormack's scheme indicates that PPMDE achieves roughly the same accuracy as these schemes when using less than half as many zones in each direction. Therefore for this problem PPMDE is more efficient than BBC and MacCormack's scheme by factors of more than 4.9 and 1.3, respectively.

MacCormack's scheme than it was on the wind tunnel problem. This should not surprise us much, because MacCormack's scheme was designed to compute steady aerodynamic flows, while BBC was intended for use on unsteady problems with strong shocks. The BBC scheme produces no oscillations behind the nearly steady reflected shock, but it gives a poorer description of the second Mach shock than

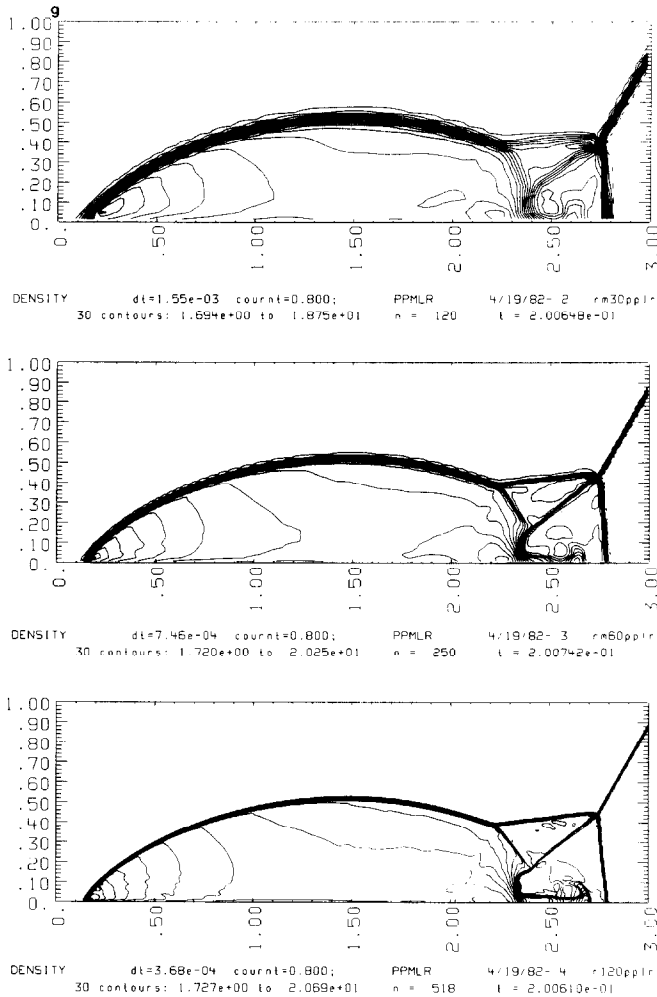


FIG. 9g. PPMLR, or PPM formulated as a Lagrangian step followed by a remap. These results improve on those of PPMDE in Fig. 9f by a reduction of the starting errors and a sharpening of the weak shock generated at the second Mach reflection. All structures are well resolved on the middle grid, and the results on the coarse grid are quite good. PPMLR gives superior performance to BBC and MacCormack's scheme when these schemes use twice as many zones in each direction. Hence it is more efficient for this problem than those schemes by factors of more than 4.4 and 1.2, respectively.

MacCormack's method. The poor description of the second Mach shock is caused in part by the "linear q " used in the Lagrangian step of BBC. Here we refer to the first term in Eq. A12 of the Appendix. This term vanishes linearly with the velocity jump across a zone rather than quadratically. This linear term in the artificial viscous pressure q was suggested for use in the BBC code by William Noh as a means of controlling post-shock oscillations.

A result on the coarsest grid obtained with our simplified MUSCL formulated as a Lagrangian step followed by a remap is shown in Fig. 10. This MUSCL run with a separate remap step shows no tendency to oscillate at the steady shock. In contrast the single-step Eulerian MUSCL runs generate noise all along the curved reflected shock. This comparison demonstrates the extra sensitivity of single-step Eulerian schemes to vanishing characteristic speeds. The extra accuracy of the single-step Eulerian approach in many circumstances is also illustrated by comparing these different MUSCL runs. In the PPM schemes we have paid a price for smooth results by broadening the shocks slightly. Nevertheless, this shock broadening seems to have had no adverse effects upon the accuracy of the calculation, even on the coarsest grid. The additional accuracy provided by the use of parabolae for internal zone structures and by contact discontinuity detection is especially noticeable in the PPM results on the fine grids. Here the contact discontinuity and the structure of the jet are particularly well represented. The extra cost in computer CPU time of the PPM runs over the MUSCL ones is 20%.

For PPM there is essentially no difference in accuracy or cost between the two formulations with and without a separate remap step. This is certainly not the case for MUSCL. This equivalence arises in PPM for two reasons. First, the PPM remap is formally third-order accurate. Second, the very thin stationary shocks which are a major advantage of the single-step Eulerian formulation have been broadened slightly in order to eliminate the noise which they would otherwise generate. For the PPM scheme one is therefore free to use the more convenient of the two formulations, depending upon the application at hand.

The results of the ETBFCT scheme for this double Mach reflection problem include some unusual features. First, there is pervasive noise in these results, as for the wind tunnel problem. Second, the shocks are steepened into staircase structures. Also these shock structures contain impossible internal states. The monotonicity constraint guarantees well-behaved shock structures for the variables ρ , ρu_x , ρu_y , and




FIG. 10. Results at time 0.2 for the double Mach reflection problem obtained using MUSCL formulated as a Lagrangian step followed by a remap. Density contours are shown for the coarse grid $\Delta x = \Delta y = 1/30$. The representation of the shock at the second Mach reflection and of the contact discontinuity is considerably worse than is obtained with the single-step Eulerian formulation of MUSCL.

ρE . However, the specific internal energy ε is given a very unreasonable structure. Inside the shock this variable attains twice its post-shock value. For simple gas dynamics problems this causes no difficulties, but if chemical reactions or radiation transport were important in the shock region, such an unreasonable structure for ε in the shock could cause problems.

A third unusual aspect of the ETBFCT results is their description of the jet which forms along the bottom wall. This jet does not develop the mushroom-shaped tip which is characteristic of a dense fluid spike plunging into a lighter fluid. Such structures commonly arise in studies of the Rayleigh–Taylor instability, for example. They are also computed by all the other difference schemes in this study except Godunov’s method. At least the PPM schemes show evidence of converging on the finest grids, so it seems likely that the ETBFCT jet structures are in error.

VII. CONCLUSIONS

In Section VIa we formulated an objective criterion for measuring the accuracy of the various difference schemes in this study in computing a blast wave interaction problem in one dimension. The result of that test was an ordering of the various schemes in terms of the accuracy achieved on a given uniform Eulerian grid. With the most accurate schemes listed first, that ordering was as follows:

1. PPM (both PPMLR and PPMDE)
2. MUSCL
3. ETBFCT
4. BBC
5. MacCormack’s scheme
6. Godunov’s scheme

If, instead of accuracy on a given grid, accuracy achieved in a given amount of CRAY-1 CPU time were used to make the list, then for our 1 – D test problem BBC and MacCormack’s scheme would move ahead of ETBFCT but otherwise this ordering would remain the same. This reordering is mainly due to the smaller time step which the ETBFCT scheme must use to ensure stability.

In the 2 – D test runs it is not so easy to define the qualities of the approximate solution which are most valuable. For example, would we prefer a smooth solution over a noisy one with certain features of the flow more accurately represented? Does the accurate representation of a jet in one part of the flow compensate for the presence of noise in another part? These are questions of individual taste, and in certain applications one feature of the solution will necessarily be more important than another. Despite the unavoidably subjective nature of any ordering of the schemes according to accuracy in two-dimensional problems, we conclude that roughly the same ordering listed above persists.

It is interesting to note that, at least for the second-order-accurate schemes in our

list, the ordering we have given with accuracy as the criterion is also an ordering according to general approach. Heading the list is Godunov's approach of building the nonlinear solution to Riemann's problem explicitly into the difference method. After Godunov's approach comes the approach of linear hybridization, or the blending of high-order- and low-order-accurate fluxes for use in the conservation laws. At the bottom of the list is the approach of smearing out discontinuities with an artificial viscosity. Although these approaches have yielded schemes of differing accuracy so far, we do not conclude that any one of them should be abandoned. Each has its own appeal. Efforts are presently under way to improve the artificial viscosity schemes by running them on adaptive grids, and better forms of the linear hybridization of fluxes are also being sought.

We suspect that all of this research may eventually lead to a conservation of difficulty principle. The schemes of the Godunov type achieve high accuracy by means of elaborate calculations of the fluxes at zone interfaces near flow discontinuities. This complicates and slows down such schemes substantially. The artificial viscosity methods perform their calculations at discontinuities very rapidly. However, to achieve the same accuracy they must refine the grid locally by at least a factor of two in each dimension and in time. Thus in 2-D at least 8 times as many zone updates must be performed. MacCormack's scheme with artificial viscosity is about 6 times faster than PPM. Therefore it requires only 35% more computer time to advance a shock on a doubly refined grid as PPM does on the original grid. At first glance, the approach of linear hybridization offers a mathematical trick, a monotonicity constraint on the flux, to get us out of this conservation of difficulty. However, when used in the form of the ETBFCT scheme our results show that this trick can lead us astray. We strongly suspect that modifications which make the flux blending more reliable in a linear hybridization will involve much of the same nonlinear information which is responsible for the success of PPM.

Our 2-D test results show that if we measure accuracy obtained per CPU second on a CRAY-1 computer we find a significantly smaller spread between the various schemes than if we measure accuracy per zone of a uniform Eulerian grid. This is of course related to the suspected conservation of difficulty principle. However, for many applications it is only the accuracy per zone which is important. In this case there is an enormous spread between the various schemes tested here. This measure of performance is appropriate for the many applications for which the computer time expended computing hydrodynamic effects is a small fraction of the overall computation. Such applications are coupled hydrodynamics and radiative transfer and the dynamics of a gas with a very complicated equation of state, which might result from a dynamic balance between many different chemical constituents.

Finally, we have seen that there is a rough equivalence of accuracy and computation time for the single-step and two-step Eulerian formulations of PPM: PPMDE and PPMLR. In the two-step formulation, a Lagrangian calculation is followed by a separate remap back to the original Eulerian grid. This equivalence of the two Eulerian schemes is not found for similar formulations of either MUSCL or Godunov's method. We conclude that this equivalence for PPM results from the

third-order accuracy of the PPM remap step. Therefore we advise that a remap be performed with one extra order of accuracy over the Lagrangian step if one desires roughly the same accuracy as can be obtained by working directly in Eulerian coordinates. Such two-step methods are very convenient for performing multifluid Eulerian computations.

Before closing it should be noted that the one-dimensional test problem presented here indicates a much wider spread in computational efficiency than is found in the two-dimensional tests. The one-dimensional problem thus argues against a conservation of difficulty principle. There are two reasons for this discrepancy. First, the largest errors in the 1-D problem arise at contact discontinuities. In two dimensions, contact discontinuities are often slip lines. When this slip is oblique to the mesh it gives rise to artificial compressional motions in the individual x - and y -passes of operator-split difference schemes such as the ones compared here. These compressional motions help to prevent rapid mass diffusion near the contact discontinuity. The second reason for the wide spread of computational efficiencies for the 1-D test problem is the extreme difficulty of that problem. It pushes the best of the schemes to their limits and the worst of the schemes well beyond theirs. A problem of this nature in two dimensions is presently completely out of the question, as convergence for the most accurate scheme considered here would require a grid of a million zones. The one-dimensional test problem is useful in showing the performance of the schemes under extreme conditions not soon to be encountered in practical two-dimensional calculations. Good performance on such an extreme problem should thus assure high-quality results under more ordinary circumstances.

APPENDIX—THE BBC SCHEME

Here we will describe the x -sweep of a 2-D BBC computation. A full 2-D step consists of an x -sweep, followed by two y -sweeps, and then by an x -sweep. All these four sweeps must use the same time step to achieve second-order accuracy (cf. [47]). We will denote zone-centered data by subscripts $i-1$, i , $i+1$ and interface centered data by subscripts $i-\frac{1}{2}$, $i+\frac{1}{2}$. The y -velocities are centered at the top and bottom interfaces of zones. We denote these interfaces by T_i and B_i , respectively. Time levels are denoted by superscripts n , $n+\frac{1}{4}$, $n+\frac{1}{2}$, $n+1$.

We begin with the Lagrangian step of the method. In terms of a mass coordinate m , defined by

$$dm = \rho dx, \quad (\text{A1})$$

we can write the Lagrangian equations of hydrodynamics in one dimension as follows:

$$\partial V / \partial t = \partial u / \partial m, \quad (\text{A2})$$

$$\partial u / \partial t = -\partial p / \partial m, \quad (\text{A3})$$

$$\partial E / \partial t = -\partial(up) / \partial m. \quad (\text{A4})$$

Here V is the specific volume, $1/\rho$, and E is the specific total energy. We have denoted u_x by u and will write v for u_y . Defining ε as the specific internal energy, we may write an internal energy equation familiar from thermodynamics:

$$\partial\varepsilon/\partial t = -p(\partial V/\partial t). \quad (\text{A5})$$

To perform a 1-D sweep of a 2-D calculation, we must remember that

$$E = \varepsilon + \frac{1}{2}(u^2 + v^2). \quad (\text{A6})$$

In this 1-D sweep the transverse velocity is passively advected:

$$\partial v/\partial t = 0. \quad (\text{A7})$$

These Lagrangian equations, specialized to the case of a gamma-law gas, are approximated by the following BBC difference equations, based on the data

$$\rho_i^n, E_i^n, u_{i-1/2}^n, v_{Bi}^n, \quad \text{for } i = 1, 2, \dots, N:$$

$$\varepsilon_i^n = E_i^n - \frac{1}{4}[(u_{i-1/2}^n)^2 + (u_{i+1/2}^n)^2 + (v_{Bi}^n)^2 + (v_{Ti}^n)^2], \quad (\text{A8})$$

$$p_i^n = (\gamma - 1) \rho_i^n \varepsilon_i^n, \quad (\text{A9})$$

$$\Delta m_i = \rho_i \Delta x_i = \rho_i(x_{i+1/2} - x_{i-1/2}), \quad (\text{A10})$$

$$q_i^n = \rho_i^n \{ [\max\{0, (u_{i-1/2}^n - u_{i+1/2}^n)\}]^2 + [\max\{0, (v_{Bi}^n - v_{Ti}^n)\}]^2 \}, \quad (\text{A11})$$

$$q_i^n = 0, \quad \text{when } \left[\frac{(u_{i+1/2}^n - u_{i-1/2}^n)}{\Delta x_i} + \frac{(v_{Ti}^n - v_{Bi}^n)}{\Delta y} \right] \geq 0,$$

otherwise

$$q_i^n = \frac{1}{3}(p_i^n q_i^n)^{1/2} + q_i^n, \quad (\text{A12})$$

$$\bar{\sigma}_i = \frac{\Delta t}{\Delta x_i} [(\gamma p_i + 16q_i)/\rho_i]^{1/2}, \quad (\text{A13})$$

$$u_{i-1/2}^{n+1/4} = u_{i-1/2}^n - \frac{\Delta t[(p_i^n + q_i^n) - (p_{i-1}^n + q_{i-1}^n)]}{2(\Delta m_{i-1} + \Delta m_i)}, \quad (\text{A14})$$

$$V_i^{n+1/2} = V_i^n + \frac{1}{2} \frac{\Delta t}{\Delta m_i} (u_{i+1/2}^{n+1/4} - u_{i-1/2}^{n+1/4}), \quad (\text{A15})$$

$$\varepsilon_i^{n+1/2} = \varepsilon_i^n - (p_i^n + q_i^n)(V_i^{n+1/2} - V_i^n), \quad (\text{A16})$$

$$p_i^{n+1/2} = (\gamma - 1) \varepsilon_i^{n+1/2} / V_i^{n+1/2}, \quad (\text{A17})$$

$$u_{i-1/2}^{n+1/2} = u_{i-1/2}^n - \frac{\Delta t[(p_i^{n+1/2} + q_i^n) - (p_{i-1}^{n+1/2} + q_{i-1}^n)]}{(\Delta m_{i-1} + \Delta m_i)}, \quad (\text{A18})$$

$$x_{i-1/2}^{n+1} = x_{i-1/2}^n + \Delta t u_{i-1/2}^{n+1/2}, \quad (\text{A19})$$

$$u_{i-1/2}^{n+1} = 2u_{i-1/2}^{n+1/2} - u_{i-1/2}^n, \quad (\text{A20})$$

$$V_i^{n+1} = V_i^n + \frac{\Delta t}{\Delta m_i} (u_{i+1/2}^{n+1/2} - u_{i-1/2}^{n+1/2}), \quad (\text{A21})$$

$$\sigma_i = \max\{\tilde{\sigma}_i, [(V_i^n - V_i^{n+1})/V_i^n]\}, \quad (\text{A22})$$

$$\varepsilon_i^{n+1} = \varepsilon_i^n - (p_i^{n+1/2} + q_i^n)(V_i^{n+1} - V_i^n), \quad (\text{A23})$$

$$v_{Bi}^{n+1} = v_{Bi}^n, \quad (\text{A24})$$

$$E_i^{n+1} = \varepsilon_i^{n+1} + \frac{1}{4}[(u_{i-1/2}^{n+1})^2 + (u_{i+1/2}^{n+1})^2 + (v_{Bi}^{n+1})^2 + (v_{Ti}^{n+1})^2]. \quad (\text{A25})$$

It is not immediately apparent, but William Noh has pointed out that this Lagrangian step is actually in conservation form. This follows from the fact that the BBC difference equations given above can be manipulated to yield a difference equation for a specific total energy, \tilde{E} , which is in a pure flux form, namely:

$$\begin{aligned} \tilde{E}_i^{n+1} = \tilde{E}_i^n - \frac{1}{2} \frac{\Delta t}{\Delta m_i} \{ & u_{i+1/2}^{n+1/2} [(p_{i+1}^{n+1/2} + q_{i+1}^n) + (p_i^{n+1/2} + q_i^n)] \\ & - u_{i-1/2}^{n+1/2} [(p_{i-1}^{n+1/2} + q_{i-1}^n) + (p_i^{n+1/2} + q_i^n)] \}, \end{aligned} \quad (\text{A26})$$

where

$$\tilde{E}_i = \varepsilon_i + \frac{1}{8 \Delta m_i} [(\Delta m_{i-1} + \Delta m_i) u_{i-1/2}^2 + (\Delta m_i + \Delta m_{i+1}) u_{i+1/2}^2] + \frac{1}{4} (v_{Bi}^2 + v_{Ti}^2). \quad (\text{A27})$$

This relation implies that the total energy, the sum over i of $\Delta m_i \tilde{E}_i$, is conserved by the scheme exactly if the boundary fluxes vanish. Because the formula for E_i in Eq. (A25) is the same as that for \tilde{E}_i in Eq. (A27), except for a redistribution among neighboring zones of a fixed amount of kinetic energy, this exact conservation of $\sum \Delta m_i \tilde{E}_i$ implies exact conservation of $\sum \Delta m_i E_i$ as well. The specific total energy used for the remap step of BBC is defined by Eq. (A25) because this simpler formula gives somewhat better results than Eq. (A27).

The remap step of BBC is based upon ideas from a MUSCL advection scheme described by van Leer in [4]. In each zone a linear distribution of density with respect to a volume coordinate is constructed having the zone-averaged value of ρ obtained from the Lagrangian calculation:

$$\rho_i^{n+1}(x) = \rho_i^{n+1} + \Delta \rho_i^{n+1} (x - x_i^{n+1}) / \Delta x_i^{n+1}, \quad (\text{A28})$$

for x within zone i . The slope $\Delta \rho_i^{n+1}$ is determined as follows:

$$\Delta \rho_i = s \max\{0, \min\{s \Delta \rho_{i-}, s \Delta \rho_{i+}\}\}, \quad (\text{A29})$$

where

$$\Delta\rho_{i\pm} = \pm \frac{2\Delta x_i}{(\Delta x_i + \Delta x_{i\pm 1})} (\rho_{i\pm 1} - \rho_i). \quad (\text{A30})$$

Here s is the sign of $\Delta\rho_{i+}$, and we have omitted superscripts $n + 1$ throughout. The advected mass δm at each zone interface is then determined by integrating the linear density distribution over the advected volume:

$$\begin{aligned} \delta m_{i-1/2} &= (x_{i-1/2}^{n+1} - x_{i-1/2}^n) [\rho_{i-1}^{n+1} + \frac{1}{2} \Delta\rho_{i-1}^{n+1} (x_{i-1/2}^n - x_{i-3/2}^n) / \Delta x_{i-1}] \\ &\quad \text{when } x_{i-1/2}^{n+1} > x_{i-1/2}^n, \text{ and} \\ \delta m_{i-1/2} &= (x_{i-1/2}^{n+1} - x_{i-1/2}^n) [\rho_i^{n+1} - \frac{1}{2} \Delta\rho_i^{n+1} (x_{i+1/2}^n - x_{i-1/2}^n) / \Delta x_i] \\ &\quad \text{when } x_{i-1/2}^{n+1} < x_{i-1/2}^n. \end{aligned} \quad (\text{A31})$$

The new density in Eulerian zone i is then

$$\rho_{iE}^{n+1} = [\Delta x_i^{n+1} \rho_i^{n+1} - (\delta m_{i+1/2} - \delta m_{i-1/2})] / \Delta x_i^n. \quad (\text{A32})$$

A similar remapping is performed for the specific total energy E , but the mass coordinate defined by Eq. (A1) plays the role of x in the above formulae. Also the advected mass $\delta m_{i-1/2}$ computed in Eq. (A31) plays the role of the advected volume.

For the momenta, the remap is also performed using a mass coordinate. The velocity $u_{i-1/2}$ is regarded as the x -momentum per unit mass averaged over a momentum box consisting of half the mass of each of zones $i - 1$ and i . The advected masses for such momentum boxes are then given by

$$\delta m_i = \frac{1}{2} (\delta m_{i-1/2} + \delta m_{i+1/2}). \quad (\text{A33})$$

The velocity v_{Bi} is regarded as the y -momentum per unit mass averaged over a momentum box consisting of half the mass of each of zone i and the zone below it. The advected masses for such momentum boxes are then given by simple averages of the advected masses $\delta m_{i-1/2}$ for zones in neighboring rows.

The formula for $\Delta\rho_i$ given above yields a second-order-accurate monotone advection scheme which is formally less accurate than the advection schemes discussed by van Leer in [4]. However, those more accurate schemes do not combine well with the Lagrangian step of BBC. In particular, at early times in the development of the blast wave problem shown in Fig. 1, BBC will generate impossibly high densities ranging from 8 to 10 depending upon which more accurate remapping scheme is employed. The remap step described here is formally less accurate but yields no unphysically large compressions.

ACKNOWLEDGMENTS

During the course of this work we have benefitted from stimulating discussions with many people. We would particularly like to acknowledge the enlightening discussions we have had with Drs. S. Zalesak and B. van Leer. This work was performed in part under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory and by the Lawrence Berkeley Laboratory under Contract W-7405-ENG-48. Partial support was also provided by the Office of Basic Energy Sciences of the U.S. Department of Energy.

REFERENCES

1. J. VON NEUMANN AND R. D. RICHTMYER, *J. Appl. Phys.* **21** (1950), 232.
2. S. K. GODUNOV, *Mat. Sb.* **47** (1959), 271.
3. R. COURANT AND K. O. FRIEDRICHS, "Supersonic Flow and Shock Waves," Interscience, New York, 1948.
4. B. VAN LEER, *J. Comput. Phys.* **23** (1977), 276.
5. B. VAN LEER, *J. Comput. Phys.* **32** (1979), 101.
6. B. VAN LEER AND P. R. WOODWARD, "Proceedings, TICOM Conference," Austin, Texas, 1979.
7. P. R. WOODWARD AND P. COLLELA, "Lecture Notes in Physics No. 141," p. 434, Springer-Verlag, New York/Berlin, 1981.
8. P. COLELLA, A direct-Eulerian MUSCL scheme for gas dynamics, *SIAM J. Sci. Statist. Comput.*, in press.
9. P. COLELLA AND P. R. WOODWARD, The piecewise parabolic method (PPM) for gas-dynamical simulations, *J. Comput. Phys.* **54** (1984).
10. J. GLIMM, *Comm. Pure Appl. Math.* **18** (1965), 697.
11. A. J. CHORIN, *J. Comput. Phys.* **22** (1976), 517.
12. G. A. SOD, *J. Comput. Phys.* **27** (1978), 1.
13. P. COLELLA, *SIAM J. Sci. Statist. Comput.* **3** (1982), 76.
14. A. J. CHORIN, *J. Comput. Phys.* **25** (1977), 253.
15. Z.-H. TENG, T.-P. LIU, AND A. J. CHORIN, *SIAM J. Appl. Math.* **42** (1982), 964.
16. T. AXELROD, private communication.
17. W. F. NOH, M. GEE, AND G. KRAMER, Lawrence Livermore National Laboratory Report UCID-18515, 1979.
18. W. D. GROPP, *SIAM J. Sci. Statist. Comput.* **1** (1980), 191.
19. J. BOLSTAD, "An Adaptive Finite Difference Method for Hyperbolic Systems in One Space Dimension," Ph. D. thesis, Stanford University, in preparation.
20. M. BERGER AND J. OLIGER, "Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations," submitted.
21. P. R. WOODWARD, Trade-offs in designing explicit hydrodynamics schemes for vector computers, in "Parallel Computation" (G. Rodrigue, Ed.), Academic Press, New York, 1982.
22. P. EGGLETON, *Mont. Not. Roy. Astron. Soc.* **151** (1971), 351.
23. J. I. CASTOR, C. G. DAVIS, AND D. K. DAVISON, "Dynamical Zoning within a Lagrangian Mesh by Use of DYN, a Stellar Pulsation Code," Los Alamos National Laboratory Report LA-66640, 1977.
24. K.-H. A. WINKLER, Technical Report MPI-PAE/Astro 90, Max-Planck-Inst. Phys. Astrophys., Garching, West Germany, 1976.
25. W. M. TSCHARNUTER AND K.-H. A. WINKLER, *Comput. Phys. Comm.* **18** (1979), 171.
26. K.-H. A. WINKLER AND M. J. NEWMAN, *Astrophys. J.* **238** (1980), 311.
27. H. A. DWYER, F. RAISZADEH, AND G. OTEY, "Lecture Notes in Physics No. 141," p. 170, Springer-Verlag, New York/Berlin, 1981.
28. K. MILLER AND R. MILLER, *SIAM J. Numer. Anal.*, in press.

29. R. J. GELINAS, S. K. DOSS, AND K. MILLER, *J. Comput. Phys.* **40** (1981), 202.
30. R. DEBAR, "Fundamentals of the KRAKEN Code," E. O. Lawrence Livermore National Laboratory Report UCIR-760, 1974.
31. A. HARTEN AND G. ZWAS, *J. Comput. Phys.* **6** (1972), 568.
32. J. P. BORIS AND D. L. BOOK, *J. Comput. Phys.* **11** (1973), 38.
33. J. P. BORIS, "Flux-Corrected Transport Modules for Solving Generalized Continuity Equations," N.R.L. Memorandum Report 3237, 1976.
34. S. T. ZALESK, *J. Comput. Phys.* **31** (1979), 335.
35. A. HARTEN, "The Method of Artificial Compression," AEC Research and Development Report C00-3077-50, New York University, 1974.
36. P. COLELLA, P. R. WOODWARD, AND S. ZALESK, in preparation.
37. P. COLELLA AND H. GLAZ, "Efficient Algorithms for the Solution of the Riemann Problem for Real Gases," Lawrence Berkeley Laboratory Report LBL-15776, 1983.
38. A. E. EMERY, *J. Comput. Phys.* **2** (1968), 306.
39. G. BEN-DOR AND I. I. GLASS, *J. Fluid Mech.* **92** (1979), 459.
40. D. BOOK, J. BORIS, A. KUHL, E. ORAN, M. PICONE, AND S. ZALESK, "Lecture Notes in Physics No. 141," p. 84, Springer-Verlag, New York/Berlin, 1981.
41. A. A. FURSENKO, V. P. GOLOVIZIN, V. A. KOMISSARUK, N. P. MENDE, AND A. I. ZHMAKIN, Leningrad Physical-Technical Institute Report No. 709, 1981.
42. G. BIRKHOFF, D. P. MACDOUGALL, E. M. PUGH, AND G. TAYLOR, *J. Appl. Phys.* **19** (1948), 563.
43. R. MACCORMACK, "Lecture Notes in Physics No. 8," Springer-Verlag, New York/Berlin, 1971.
44. R. D. RICHTMYER AND K. W. MORTON, "Difference Methods for Initial Value Problems," Interscience, New York, 1967.
45. A. LAPIDUS, *J. Comput. Phys.* **2** (1967), 154.
46. W. G. SUTCLIFFE, "BBC Hydrodynamics," E. O. Lawrence Livermore National Laboratory Report UCID-17013, 1973.
47. G. STRANG, *SIAM J. Numer. Anal.* **5** (1968), 506.
48. T. IKEDA AND T. NAKAGAWA, *Math. Comp.* **33** (1979), 1157.
49. S. K. GODUNOV, A. W. ZABRODYN, AND G. P. PROKOPOV, *Zh. Vycisl. Mat. Fiz.* **1** (1961), 1020.
50. S. HANCOCK, private communication.
51. B. VAN LEER, "On the Relation between the Upwind Differencing Schemes of Godunov, Engquist-Osher, and Roe," ICASE Report No. 81-11, 1981.
52. S. OSHER AND F. SOLOMON, *Math. Comp.* **38** (1982), 339.
53. A. HARTEN AND P. D. LAX, *SIAM J. Numer. Anal.* **18** (1981), 289.
54. P. L. ROE, *J. Comput. Phys.* **43** (1981), 357.
55. J. SALTZMAN, Courant Math. Res. D0E/ER/03077/174, 1982.
56. P. COLELLA, "Approximate Solution of the Riemann Problem for Real Gases," Lawrence Berkeley Laboratory Report LBL-14442.