CONF-820854--2

UCRL--90009

DE84 003193

PIECEWISE-PARABOLIC METHODS FOR ASTROPHYSICAL FLUID DYNAMICS

Paul R. Woodward

This paper was prepared for submittal to the
Proceedings of the NATO Advanced Workshop in
Astrophysical Radiation Hydrodynamics,
Munich, West Germany -- August 1982.

November 1, 1983

Lawrence
Livermore
National
Laboratory

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

MASTER

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

PIECEWISE-PARABOLIC METHODS FOR ASTROPHYSICAL FLUID DYNAMICS

Paul R. Woodward
Lawrence Livermore National Laboratory
Livermore, California 94550

ABSTRACT

A general description of some modern numerical techniques for the simulation of astrophysical fluid flow is presented. The methods are introduced with a thorough discussion of the especially simple case of advection. Attention is focussed on the piecewise-parabolic method (PPM). A description of the SLIC method for treating multifluid problems is also given. The discussion is illustrated by a number of advection and hydrodynamics test problems. Finally, a study of Kelvin-Helmholtz instability of supersonic jets using PPM with SLIC fluid interfaces is presented.

299

INTRODUCTION

There is a long history of the use of computer simulations in astrophysics because of the difficulty or impossibility of obtaining useful observational data. It should come as no surprise, for example, that nearly all our understanding of stellar structure and evolution has come from computer calculations. In other areas such as star formation, supernova explosions, formation of interstellar clouds, and the evolution of galaxies some observational data exists. However, computer simulations offer the potential for controlled experiments in which the various parameters of a theoretical model may be varied independently and unequivocal observations of the results may be made. Scant astronomical observational data often allows several theoretical interpretations. Computer simulations have been used to choose between such interpretations by means of checking for a rigorous consistency of theoretical arguments with known physical laws whose workings are too complex for the unaided mind to fathom.

For these reasons computer simulations have played an important role in astrophysics since the early 1960's, and there is every indication that their role will continue to grow along with the power of the computing machines they exploit. Consequently it is advisable for modern astrophysicists to develop some familiarity with the methods by which computer simulations are performed. Without such a familiarity the astrophysicist is forced to take all computational results at face value, and this can lead to misunderstandings. It is the purpose of this article to acquaint the uninitiated with some fundamental ideas behind a group of modern techniques for simulating fluid flow on computers. There is no intent to set down detailed equations, to give any more than an indication of the historical development of this subject, or to maintain any semblance of objectivity in choosing which techniques merit discussion. This editorial attitude will simplify the discussion so that the essential ideas of the various techniques which receive attention may be more clearly understood. The reader who desires a fuller account will be referred to a number of articles where he may find it.

THE RELATIONSHIP BETWEEN SIMPLE ADVECTION AND ASTROPHYSICAL FLUID DYNAMICS

Every informed astrophysicist should know that the foundation of fluid dynamics is the Boltzmann equation. Apart from the collision integral on the right-hand side, this equation describes simple advection, an incompressible motion of fluid in phase space. We write the collisionless Boltzmann equation in a two-dimensional phase space as

$$\frac{\partial f}{\partial t} + v \frac{\partial f}{\partial x} + g \frac{\partial f}{\partial v} = 0 \quad (1)$$

Here f is the density of the fluid in phase space, x is the distance coordinate, v the speed, and g the acceleration. When g is independent of v the motion in phase space is obviously incompressible. Inappropriate as it may seem, this equation with the last term dropped (i.e., $g = 0$) has been the conceptual basis for the design of numerical methods for treating the much more complicated, nonlinear equations of fluid dynamics. The success of this approach rests on the fact that the fluid equations are derived from the collisional Boltzmann equation, and hence at a fundamental level they involve simple advection. To see this, consider the equations for the one-dimensional, isentropic flow of a polytropic fluid. In Lagrangian coordinates moving with the fluid the governing equations are:

$$\frac{\partial x}{\partial t} = u \quad (2)$$

$$\frac{\partial v}{\partial t} - \frac{\partial u}{\partial m} = 0 \quad (3)$$

$$\frac{\partial u}{\partial t} + \frac{\partial p}{\partial m} = 0 \quad (4)$$

$$dm = \rho \, dx \quad (5)$$

$$v = 1/\rho \quad (6)$$

$$p = A p^\gamma \quad (7)$$

Here m is a mass coordinate related to the distance coordinate through the density ρ as in Eq. (5), A is a constant, and γ is the constant polytropic index. The Eulerian sound speed c and the Lagrangian sound speed C are defined by

$$c^2 = \gamma p v \quad (8)$$

$$C = \rho c \quad (9)$$

If we define two Riemann invariants R_+ and R_- by

$$R_{\pm} = u \pm \frac{2}{\gamma-1} c \quad (10)$$

it is a trivial matter to verify that these quantities satisfy the advection equations

$$\frac{\partial R_{\pm}}{\partial t} \pm C \frac{\partial R_{\pm}}{\partial m} = 0 \quad . \quad (11)$$

This implies that the two Riemann invariants are simply advected in opposite directions at the speed C in the Lagrangian coordinate system (R_+ goes to the right and R_- to the left). Thus the hydrodynamical problem has been reduced to simple advection.

Under more general circumstances fluid dynamics is not so simple. For large amplitude sound signals the two Riemann invariants interact and entropy is produced (A increases). In this case formulae like Eq. (10) for the Riemann invariants cannot be written; however, equations for the differentials $dR_{\pm}$ are still useful. Now A in Eq. (7) is no longer constant and we must augment our nonlinear system of equations with an energy equation

$$\frac{\partial E}{\partial t} + \frac{\partial u p}{\partial m} = 0 \quad . \quad (12)$$

Here E is the total energy per unit mass, which is related to the specific internal energy ϵ by

$$E = \epsilon + \frac{1}{2} u^2 \quad . \quad (13)$$

The equation of state which replaces Eq. (7) is now

$$p = (\gamma - 1) \rho \epsilon \quad . \quad (14)$$

The characteristic equations for changes in Riemann invariants now become

$$dR_{\pm} = du \pm \frac{dp}{C} = 0 \quad . \quad (15a)$$

Equations (15a) hold only along characteristic paths defined by

$$dm = \pm C dt \quad . \quad (15b)$$

In addition, a third characteristic equation

$$dR_0 = dp - c^2 dp = 0 \quad (15c)$$

applies along characteristic paths defined by $dm = 0$. Equations (15) hold only in smooth flow; if shocks or contact discontinuities are encountered on the characteristic paths, dR_+ and dR_0 are nonzero across these discontinuities. Even when no simple formulae for quantities R_+ can be written, the concept of advection of Riemann invariants along characteristic paths still proves useful.

Advection also enters the study of fluid dynamics more directly when Eulerian calculations are performed. In this article we will consider an Eulerian calculation to consist of two steps, a Lagrangian calculation followed by a remapping of the results of that calculation onto the Eulerian grid. The remap step of such a composite calculation requires the solution of the following nonlinear system of equations:

$$\frac{\partial \rho}{\partial t} + u \frac{\partial \rho}{\partial x} = 0, \quad (16)$$

$$\frac{\partial u}{\partial t} + \rho u \frac{\partial u}{\partial m} = 0, \quad (17)$$

$$\frac{\partial E}{\partial t} + \rho u \frac{\partial E}{\partial m} = 0. \quad (18)$$

We have written the last two equations in this set in an unusual manner which gives the best guide to the consistent numerical solution of the set of equations as a whole. Clearly, all three of these equations are just Eq. (1) with $g = 0$ and with different definitions of the density f , the coordinate x , and the advection speed v .

It should be noted that Eq. (1) has other uses in astrophysics. When g is determined from Poisson's equation, it becomes the equation of stellar dynamics. A similar equation with a source term on the right is the equation of radiative transfer. Equations like Eq. (1), often with source terms on the right, must be solved whenever a constituent of the fluid must be treated which has a long mean free path.

NUMERICAL TECHNIQUES FOR EXPLICIT CALCULATIONS OF SIMPLE ADVECTION

a. The PPM and PPB advection schemes

We will discuss two general approaches to solving the simple advection equation in one dimension

$$\frac{\partial f}{\partial t} + v \frac{\partial f}{\partial x} = 0 \quad (19)$$

For a fuller discussion of this subject and for historical references the reader is referred to van Leer (1977), Woodward and White (1983), and Colella and Woodward (1983). Of course, there is really only one way to solve Eq. (19) numerically; one takes a finite amount of data about $f(x)$, from it he infers an approximation to $f(x)$ everywhere, and then he uses this approximation to update his finite data in some straightforward way. This simple truth stands in relation to numerical analysis of partial differential equations as the following summary of algebra does to that subject: "You just let something be x and then solve for it" (E. J. Woodward 1959, private communication). Needless to say, many mathematicians have based careers on finding out just how the approximation to $f(x)$ everywhere should be constructed.

We will consider here only difference methods in conservation form. For these methods the finite data about $f(x)$ must in part consist of or be equivalent to the locations x_i of zone interfaces and the average values $\langle f \rangle_i$ of $f(x)$ within the intervals from x_i to x_{i+1} . Here the zone number i runs from 1 to some finite number N . For one method we will discuss, PPM, this is all the data provided for the calculation. For another method, PPB, additional moments of $f(x)$ within the zones must be provided. These moments will be written as $\langle f \chi^k \rangle$, where $k = 0, 1, 2$ and where χ is a zone-centered coordinate defined by

$$\chi = [x - \frac{1}{2} (x_L + x_R)] / \Delta x \quad (20)$$

Here x_L and x_R denote the locations of the left- and right-hand interfaces of the zone, that is $x_L = x_i$ and $x_R = x_{i+1}$ for zone i . The zone width is Δx , and therefore χ ranges from $-1/2$ to $1/2$ across the zone. The moments $\langle f \chi^k \rangle$ are defined by

$$\langle f \chi^k \rangle = \int_{-1/2}^{1/2} f(\chi) \chi^k d\chi \quad (21)$$

In the first approach to solving Eq. (19), here represented by the difference scheme PPM (the Piecewise-Parabolic Method; Woodward and Colella 1980, 1983; Colella and Woodward 1983), only the masses of the zones are used as data. These are given by Eq. (21) with $k = 0$. In order to approximate $f(x)$ everywhere PPM constructs a parabola to represent $f(\chi)$ within each zone:

$$f(\chi) = f_0 + f_1 \chi + f_2 \chi^2 \quad (22)$$

To compute the coefficients f_k , PPM first interpolates values f_L and f_R of f at the left- and right-hand interfaces of the zones. This is done by assuming that $f(x)$ is smooth, so that it can be approximated near x_L by the unique cubic curve which has the average values in the four zones nearest x_L which are prescribed by the data for the scheme. This cubic curve is then evaluated at x_L to give the interpolated density f_L . For the special case of uniform zone widths the formula for f_L , or $f_{i-1/2}$, has the following simple form:

$$f_{i-1/2} = \frac{7}{12} (\langle f \rangle_{i-1} + \langle f \rangle_i) - \frac{1}{12} (\langle f \rangle_{i-2} + \langle f \rangle_{i+1}) . \quad (23)$$

Once values for f_L and f_R have thus been obtained, the parabola in Eq. (22) is determined by demanding that it pass through f_L and f_R at x_L and x_R and that its integral over the zone should yield the prescribed zone mass $\langle f \rangle \Delta x$. The coefficients f_k must therefore be

$$f_0 = \frac{3}{2} \langle f \rangle - \frac{1}{4} (f_L + f_R) , \quad (24a)$$

$$f_1 = f_R - f_L , \quad (24b)$$

$$f_2 = 3 (f_L + f_R) - 6 \langle f \rangle . \quad (24c)$$

We will later describe important modifications to the above procedure which improve its performance near discontinuities in $f(x)$.

In the second approach, represented by the difference scheme PPB (the Piecewise-Parabolic Boltzmann scheme; van Leer 1977, Woodward and White 1983), the three moments $\langle f \rangle$, $\langle f\chi \rangle$, and $\langle f\chi^2 \rangle$ are used as data. This information is sufficient to uniquely determine the coefficients of the interpolation parabola in Eq. (22):

$$f_0 = \langle f \rangle - 15 \langle f\chi^2 \rangle_2 , \quad (25a)$$

$$f_1 = 12 \langle f\chi \rangle , \quad (25b)$$

$$f_2 = 180 \langle f\chi^2 \rangle_2 , \quad (25c)$$

$$\langle f\chi^2 \rangle_2 = \langle f\chi^2 \rangle - \langle f \rangle / 12 . \quad (25d)$$

When the advection velocity, v , in Eq. (19) is constant, the derivation of formulae to update the original data is straightforward, once the interpolation parabola have been determined at the original time $t = 0$. In this special case Eq. (19) has a trivial solution:

$$f(x,t) = f(x - vt, 0) \quad . \quad (26)$$

Thus the interpolation parabola for $f(x,0)$ also give $f(x,t)$, and this function need only be integrated over the zones to give the new data (see Fig. 1). For the PPB scheme care must be taken in evaluating the higher moment integrals in Eq. (21) to avoid confusing the x 's referring to time 0 and to time Δt . For details the reader is referred to Woodward and White (1983). To keep the logic in the program simple it is customary to demand that the timestep Δt be limited so that no zone interface moves out of the interval defined by the two zones adjacent to it.

Fig. 1 The PPM advection scheme moves a sine wave resolved with 3 zones a distance of $3/4$ zone widths to the right in a single timestep. (a) The zone-averaged values for sine wave (shown dotted) are used to construct a parabolic representation of the curve within each zone (solid lines). (b) These original interpolation parabola (dotted) have been translated $3/4$ zone widths to the right. Integrating these translated curves over the original zone intervals gives the new zone averages. These have been used to construct new interpolation parabola (solid lines), the final PPM representation of the translated sine wave.

When the advection velocity v depends upon x , as is usually the case in the remap step of a hydrodynamics calculation, interpolation parabolae are constructed directly for the zones at time Δt after they have translated, expanded, or contracted. The necessary data referring to these zones is provided by the output of the Lagrangian step of the calculation. Straightforward integrations over the interpolated structures again yield data referring to the original Eulerian mesh. As is indicated by the form of Eqs. (16), (17), and (18), the density should be interpolated as a function of a volume coordinate. The integrations over these density structures then yield masses advected from one zone to another. These should be used to define the advection velocities relative to the mass coordinate. The fluid velocity and specific total energy are then interpolated as functions of the mass coordinate, and integration of these structures over the advected masses gives the momenta and total energies advected from one zone to another.

b. Comparison of PPM and PPB on a Gaussian Advection Problem

The performance of the two schemes PPM and PPB is illustrated by the Gaussian advection tests shown in Figs. 2 and 3 and tabulated in Table I. The initial data for these runs is constructed using Eq. (21) with f set to f_0 , where

$$f_0(x) = \exp [(x - 1/2)^2 / 0.01125] \quad . \quad (27)$$

Thus $f_0(x)$ is a Gaussian of height 1 and standard deviation 0.075 centered at $x = 1/2$. Periodic boundary conditions are applied at $x = 0$ and $x = 1$. Using grids with 8, 16, 32, 64, and 128 zones, Eq. (19) has been solved with $v = 1$. Solutions were obtained with two values of the Courant number, σ , defined by

$$\sigma = v \Delta t / \Delta x \quad . \quad (28)$$

The values of σ used, 0.08 and 0.8, represent approximate worst and best cases, with the errors at $\sigma = 0.08$ as much as four times larger than those at $\sigma = 0.8$. The results of the PPM calculations are shown in Fig. 2, while those of the PPB calculations are shown in Fig. 3. In generating these figures the schemes themselves have been used to determine interpolation parabolae within the zones, so that $f(x)$ is defined everywhere for plotting. In addition integrated errors ϵ_0 , ϵ_1 , ϵ_{10} , and ϵ_{100} have been computed from the results at times 0, 1, 10, and 100 using the interpolated $f(x)$ obtained by the scheme:

$$\epsilon = \int_0^1 |f(x) - f_0(x)| dx / \int_0^1 f_0(x) dx . \quad (29)$$

These errors are tabulated in Table I.

A glance at Figs. 2 and 3 reveals a striking difference in accuracy between the PPM and PPB schemes. Apparently, it makes a great deal of difference which interpolation parabola is used to describe the internal structure of a zone. The PPM scheme requires grids of 16, 32, 64, or 128 zones if the Gaussian pulse is to be well represented after 0, 1, 10, or 100 transits of the grid, respectively. In contrast to this, the PPB scheme requires only 8, 8, 16, or 32 zones for the same purpose.

Thus the PPB scheme uses three times as much data per zone but needs only one fourth as many zones as PPM. One might have thought that the smoothness of the Gaussian would allow the higher moment data used by PPB to be effectively regenerated from the information provided by $\langle f \rangle$ in neighboring zones. The PPM results indicate that if such a reconstruction of the data is ever possible, it must require a much finer grid than is needed by PPB to adequately resolve the density structure. (This result holds true even when seventh-order curves are used in place of cubics to interpolate the interface values f_L and f_R in PPM.) Because PPB requires only 40% more time to update a zone than PPM (on the Cray-I PPB updates 1.2×10^6 zones in a second), it is therefore much more efficient to use PPB on the coarser grid.

A close look at the errors on the fine grids in Table I shows that both schemes are converging faster than we might have expected. Generally, if a polynomial of order n is used to describe the internal zone structure, the associated advection scheme is $n+1$ th order accurate. The interpolation polynomial cannot account for the $n+1$ th derivative of f within the zone, so the advected mass calculated at each interface contains an error proportional to this derivative, which is assumed to be of order 1, multiplied by Δx^{n+1} . However, a similar error is introduced in the advected mass at both interfaces of the zone, and these tend to cancel when the two advected masses are differenced to obtain the net change in the zone's mass. Thus the error in this net change is of order Δx^{n+2} . If we hold the Courant number σ fixed as Δx is reduced, then the number of timesteps needed to reach a specified time increases as $1/\Delta x$. This effect implies that the order of the error in computing to a fixed time, that is the order of the difference scheme, is one less than the order of the error incurred in a single timestep. Hence a scheme employing an n th-order polynomial to describe the internal structure of a zone should be $n+1$ th order accurate.

The errors in Table I for the fine grid calculations indicate that PPM converges nearly as fast as a fourth-order scheme and that PPB converges nearly as fast as a fifth-order scheme. The argument given above implies that both should be third-order schemes in a certain strict sense, but obviously as a practical matter both are more accurate than this. The accelerated convergence of PPM comes from its "higher-order spatial

Fig. 2a A Gaussian of height 1 and standard deviation 0.075 is advected by the PPM scheme at a Courant number of 0.08. Results for grids of 8, 16, 32, 64, and 128 zones are plotted against the exact solution at times 0, 1, 10, and 100. The parabolae used by PPM to interpolate within zones are used here for plotting. After the Gaussian has traversed the grid 100 times, only the finest grid gives adequate results.

accuracy". Because the density values at zone interfaces are interpolated using cubic curves, and because in the limit of small Courant numbers (or of σ near unity) these interface values alone give the advected masses, in this same limit the scheme must be fourth-order accurate. Thus the accelerated convergence of PPM is more noticeable in the runs with $\sigma = 0.08$, but it is still effective even when $\sigma = 0.8$.

The accelerated convergence of PPB is due to its exact conservation of the moments $\langle f_X \rangle$ and $\langle f_X^2 \rangle$. Consider the

Fig. 2b Same as Fig. 2a except that a Courant number of 0.8 was used. This larger timestep improves the results substantially, so that the 64-zone grid now gives good results at time 100.

initial replacement of the exact distribution $f_0(x)$ with the PPB representation $f(x)$. If we define the local error $\epsilon(x)$ by

$$\epsilon(x) = f(x) - f_0(x) , \quad (30)$$

then the preservation of the moments $\langle f_0 x^k \rangle$ by the PPB

Fig. 3a The same advection experiment shown in Fig. 2a is here carried out using the PPB scheme. The interpolation parabolae used by the scheme to represent the distributions within zones are used here for plotting. Note the quality of this representation on the 3-zone grid at times 0 and 1. The three finest grids give very good results at time 100. Note the low level of the negative densities generated on the coarsest grids, particularly in contrast to the PPM results in Fig. 2a.

representation guarantees that

$$\langle \epsilon \chi^k \rangle = 0, \quad \text{for } k = 0, 1, 2. \quad (31)$$

This very special property implies that any meaningful definition of global error which involves integrals over $\epsilon(x)$ rather than its absolute value will yield a sixth-order rather than a third-order error. For example, consider the following error definition:

Fig. 3b Same as Fig. 3a except that a Courant number of 0.8 was used. The results are improved by the use of this larger timestep, so that now only 16 zones are needed to produce good results at time 100.

$$\epsilon_{f0} = \int_0^1 \epsilon(x) f_0(x) dx / \int_0^1 [f_0(x)]^2 dx, \quad (32)$$

The numerator of this expression can be broken up into a sum of terms like

$$\begin{aligned} \Delta x \int_{-1/2}^{1/2} \epsilon(\chi) f_0(\chi) d\chi = & \Delta x \langle \epsilon \rangle f_0(0) + (\Delta x)^2 \langle \epsilon \chi \rangle \left. \frac{\partial f_0}{\partial x} \right|_{\chi=0} \\ & + \frac{1}{2} (\Delta x)^3 \langle \epsilon \chi^2 \rangle \left. \frac{\partial^2 f_0}{\partial x^2} \right|_{\chi=0} \\ & + \frac{1}{6} (\Delta x)^4 \langle \epsilon \chi^3 \rangle \left. \frac{\partial^3 f_0}{\partial x^3} \right|_{\chi=0} \end{aligned} \quad (33)$$

By virtue of Eq. (31), the first three terms on the right in Eq. (33) vanish. Therefore the numerator in Eq. (32) can be expressed as a sum of terms each of which is of order $(\Delta x)^4 \langle \epsilon \chi^3 \rangle$. Now $\epsilon(\chi)$ is of order $(\Delta x)^3$, and the number of terms is of order $1/\Delta x$, so the integrated error ϵ_{f0} is of order $(\Delta x)^6$. An advection timestep for a constant advection velocity consists of an exact "Lagrangian step" in which the zones all move over a constant distance followed by a remap which preserves three moments. If we regard the piecewise-parabolic representation in the zones which have shifted over a distance $v \Delta t$ as the "exact" distribution $f_0(x)$, then our argument above shows that ϵ_{f0} is increased by an amount of order $(\Delta x)^6$ when we use the exact moments of this $f_0(x)$ in the original Eulerian zones to construct a new piecewise-parabolic representation $f(x)$. Consequently, using the error ϵ_{f0} to measure the performance of the PPB scheme, we discover that this scheme is fifth-order accurate.

van Leer (1977) found another explanation for the unexpected accuracy of moment-conserving difference schemes. His work applied to the MUSCL scheme, which conserves only $\langle f \rangle$ and $\langle f \chi \rangle$, but similar arguments apply to PPB. The initial data constructed from $f_0(x)$ by finding the exact moments $\langle f_0 \chi^k \rangle$ can be broken into two parts, one which will be advected by PPB

with fifth-order accuracy and one which will be damped by a factor of order unity in each timestep. The second, rapidly damped part decreases in size as $(\Delta x)^3$. It therefore contributes an error of this size which, after several timesteps, does not continue to grow. This error makes the scheme third-order accurate, but after many timesteps this error is overwhelmed by the growing fifth-order error which arises from the advection of

TABLE I

Numerical Errors Generated in Gaussian Advection Tests

Scheme	Grid	ϵ_0	ϵ_1	ϵ_{10}	ϵ_{100}
PPM $\sigma = 0.08$	8	2.28e-1	1.04	9.82e-1	1.15
	16	1.92e-2	4.12e-1	9.00e-1	8.89e-1
	32	1.49e-3	5.87e-2	3.51e-1	6.54e-1
	64	1.17e-4	4.05e-3	3.96e-2	2.66e-1
	128	1.09e-5	2.76e-4	2.74e-3	2.65e-2
PPB $\sigma = 0.08$	8	2.54e-2	1.02e-1	3.50e-1	6.66e-1
	16	2.53e-3	6.99e-3	5.08e-2	2.14e-1
	32	3.00e-4	4.50e-4	2.37e-3	2.06e-2
	64	3.70e-5	4.74e-5	9.02e-5	7.66e-4
	128	4.62e-6	5.80e-6	6.36e-6	2.49e-5
PPM $\sigma = 0.8$	8	2.28e-1	5.28e-1	9.52e-1	1.13
	16	1.92e-2	1.29e-1	4.82e-1	9.57e-1
	32	1.49e-3	1.20e-2	9.67e-2	3.98e-1
	64	1.17e-4	8.59e-4	8.32e-3	7.13e-2
	128	1.09e-5	7.23e-5	7.14e-4	6.99e-3
PPB $\sigma = 0.8$	8	2.54e-2	3.67e-2	1.50e-1	4.30e-1
	16	2.53e-3	2.95e-3	1.11e-2	7.51e-2
	32	3.00e-4	3.26e-4	5.04e-4	3.77e-3
	64	3.70e-5	3.98e-5	4.20e-5	1.29e-4
	128	4.62e-6	4.96e-6	4.99e-6	6.33e-6

the first component of the initial data. This behavior can be clearly seen in the final entries in Table I for the PPB run with 128 zones at a Courant number of 0.8. These numbers show a relatively rapid initial growth of the error followed by a much slower phase of error accumulation.

In Fig. 4 the initial data for a pure sine wave on a grid of 16 uniform zones is decomposed into parts which are damped at different rates by the PPB scheme. In Fig. 4a the most slowly damped component is the one which closely resembles a sine wave. The other component shown is very rapidly damped. In order to make it visible, it has been amplified for plotting by the factor $100/(\Delta x)^3$. This component itself consists of two parts which are damped by PPB at different rates. These are shown in Fig. 4b. In fact, these three components of the sine wave are eigenfunctions of the PPB scheme; they are reduced in amplitude and shifted in phase by the scheme but their shapes are preserved. The larger one in Fig. 4b is damped by a factor of $7/8$ in each timestep at a Courant number of $1/2$, while the smaller one is damped only by a factor of $1/2$. The larger eigenfunction contributes an amount of order $(\Delta x)^3$ to the initial data but

Fig. 4 The initial representation of a pure sine wave by the PPB scheme on a 16-zone grid is here decomposed into parts which are damped at different rates by the scheme. On the left, the largest eigenfunction component, which closely resembles the sine wave, is shown. This eigenfunction is advected by PPB with fifth-order accuracy. The saw-toothed curve plotted with this eigenfunction is the difference between that eigenfunction and the sine wave. It has been multiplied by a factor $(100/(\Delta x)^3)$ to make it visible. This component, which is third-order in Δx , consists of two eigenfunctions of the PPB scheme. These are shown on the right. Both are damped strongly in a single timestep.

its contribution to the fundamental mode in the Fourier spectrum of this data is of order $(\Delta x)^6$. This is consistent with our previous discussion of the error measure ϵ_{f0} . These contributions for the smaller eigenfunction scale as $(\Delta x)^4$ and $(\Delta x)^8$.

PRACTICAL LIMITS TO THE ACCURACY OF A DIFFERENCE SCHEME FOR ADVECTION

From the above comparison of PPM and PPB one might be tempted to conclude that the best way to solve advection problems numerically would be to use a very high-order moment conserving method. This conclusion is false for a number of reasons. The most important of these reasons is that the advection problems which arise in computational physics are usually either multidimensional or nonlinear or both. For multidimensional problems, the amount of work for both the programmer and the computer rises rapidly as the accuracy increases beyond the second order. This can most easily be seen in two dimensions, with Cartesian coordinates x and y . Strang (1968) showed that a second-order, two-dimensional difference operator D_{xy} could be constructed from a symmetrized product $D_x D_y D_y D_x$ of second-order, one-dimensional difference operators. The second-order error generated by using $D_y D_x$ in place of $D_{xy} = D_x + D_y$ is cancelled by the error in using $D_x D_y$ subsequently. This error in $D_y D_x$ arises for nonlinear advection equations, such as the characteristic equations of isentropic gas dynamics [Eqs. (6)-(11)]. In that case the advection speeds are altered by D_x and D_y , so that $D_y D_x$ and $D_x D_y$ must give different results. For such nonlinear advection equations even third-order methods seem impractically complex.

For linear advection equations multidimensional calculations can be performed in a sequence of "passes" in which the fluid is permitted to move in only one dimension. However, if more than second-order accuracy is required these passes must be two dimensional in nature. The reason for this is illustrated in Fig. 5. If all gradients in the y direction are ignored in the x -pass, then the difference of the masses of the upper shaded triangles in the figure is being effectively equated to that of the lower shaded triangles. The error this introduces in the net change in the zone mass is proportional to

$$\Delta x (\Delta y)^2 \frac{\partial}{\partial x} \left(\frac{\partial \rho}{\partial y} \frac{\partial v_x}{\partial y} \right) .$$

Therefore a third-order difference scheme must account for this term. Even though an advection equation to be solved is nonlinear, so that strict third-order accuracy is unattainable by means of 1-D passes, it may still make sense to use a scheme as elaborate as PPB or PPM and to account for the above error term.

Fig. 5 The x-pass of an operator-split advection calculation in which a purely one-dimensional operator D_x is used. Ignoring gradients of ρ and v_x in the y-direction effectively causes the mass difference of the upper shaded triangles to be equated to that of the lower ones.

Such an application is stellar dynamics, in which the Boltzmann Eq. (1) must be solved. The equation is nonlinear by virtue of the coupling of the acceleration g to the phase space density through Poisson's equation. However, this coupling is weak, as only the velocity-integrated density enters Poisson's equation. Another reason why high-order advection techniques are appropriate in this application is that the high dimensionality of the problem means that only very coarse grids can be afforded. Therefore one is unlikely to find the calculation in the asymptotic regime where third- or fifth-order convergence is achieved. Actually, in such problems in 4 or 6 phase space dimensions one is lucky to resolve even the gross features of the problem adequately. Then schemes like PPB are especially attractive because of their good behavior when the flow is badly underresolved.

The usefulness of the PPB scheme in such stellar dynamical problems is illustrated by the gravitational two-stream instability problem discussed in detail by Woodward and White (1983). The time development of this flow is shown in Fig. 6. The problem begins with the imposition of a sinusoidal perturbation in the x-direction on the equilibrium density f_0 of fluid in phase space. This equilibrium consists of a uniform distribution of gravitating fluid in physical space. In velocity space the distribution in the v_y - and v_z -directions is a delta function at velocity zero. In the v_x -direction the distribution consists of the superposition of two Gaussians centered at $v_x = \pm 4$ and with standard deviations of unity. Each Gaussian contains a velocity-integrated density of $1/2$. The sinusoidal perturbation imposed on this equilibrium at time zero has a wavelength of 12

Figure 6a

RHO dt:4.69e-04 count:0.950: BoltzPPM 4/28/83- 5 blitzeq
15 contours: 6.980e-03 to 2.024e-01 n = 854 t = 5.00000e-01

RHO dt:4.69e-04 count:0.950: BoltzPPM 4/28/83- 5 blitzeq
15 contours: -1.250e-02 to 2.018e-01 n = 1782 t = 1.00000e-00

Figure 6b

RHO dt:4.69e-04 count:0.950 BoltzPPM 4/28/83- 5 blitzeb
 15 contours:-1.387e-02 to 2.918e-01 n = 2550 t = 1.50000e+00

RHO dt:4.69e-04 count:0.950 BoltzPPM 4/28/83- 5 blitzebq
 15 contours:-1.543e-02 to 2.017e-01 n = 3398 t = 2.00000e+00

Figure 6c

Figure Captions

Fig. 6 Contours of density in phase space for the gravitational two-stream instability problem described in the text. These results were obtained with the PPB scheme using a fine grid of 480x480 uniform zones. The flow begins with two Gaussian velocity distributions at $v = \pm 4$ and with unit standard deviation. A 5% sinusoidal perturbation of the velocity-integrated density is imposed. The flow is shown at time intervals of 1/2. The spiralling motion about the location of the initial density maximum causes ever thinner cords of phase space density to develop which eventually can no longer be resolved.

Fig. 7 PPB results at time 2.5 are displayed for two coarse grid calculations of the gravitational two-stream instability problem shown in Fig. 6. Grids of 30x30 (top) and 120x120 (bottom) uniform zones were used. As the grid is coarsened, fine-scale features are blended, but very little numerical noise is generated. The large-scale features near the center of the plots are adequately described even on the coarser grid. The PFE interpolation parabolae have been used to generate a 5x5 grid of density values per zone for plotting.

Figure 7

RHD dt: 9.50e-03 count: 0.950: BoltzPPM 5/11/83- 1 blitzad
15 contours: -2.126e-02 to 2.000e-01 n: 270 t: 2.51875e+00

RHD dt: 2.37e-03 count: 0.950: BoltzPPM 4/27/83- 3 blitzcar
15 contours: -1.106e-02 to 2.048e-01 n: 1050 t: 2.50119e+00

and an amplitude of 5%. The specification of a value for the gravitational constant G in Poisson's equation determines the relationship of these mass, length, and time units to the cgs system. We have set $G = \pi$ for this problem. The Jeans length for a collisional gas with a sound speed of unity would then be unity. Our perturbation, with a wavelength of 12 is therefore quite unstable.

In Fig. 6 the phase space density is represented at time intervals of $1/2$. Each snapshot displays 15 evenly spaced contours of phase space density with the lowest contour dotted. The calculation employed a very fine grid of 480×480 zones (only half of these were computed and the others were generated for plotting by symmetry). The standard deviations of the initial Gaussians were well resolved on this grid by six zones. The nature of the flow is a spiraling of most of the fluid about the origin and a translation of small parts of the fluid across the entire grid. Because the problem is periodic this translating envelope of fluid goes from one gravitationally bound clump to another, lingering between each pair. As the flow continues to wind up, ever thinner cords of phase space density are created, so that ultimately the flow must become unresolved. At this point PPB behaves relatively well. Instead of generating large amounts of numerical noise, as would be the case for most high-order methods, PPB averages over the individual thin structures to obtain a smoother structure. In this averaging process information is lost and entropy is created, but six moments of the local phase space distribution are nevertheless conserved. This averaging process is illustrated by the results in Fig. 7. These results were obtained with PPB using grids of 30×30 (Fig. 7a) and 120×120 zones (Fig. 7b). For plotting, the PPB interpolation parabola were used to generate data values on a 5×5 grid within each zone. This figure shows how detailed features are smeared out and blended as the grid is coarsened, but the important larger structures near the origin are not lost. In large part this blending constitutes a numerically accelerated representation of the physical relaxation process of phase mixing.

The gravitational two-stream instability problem just discussed is a good example of the type of problem for which PPB is well-suited. The advection equation is nonlinear, so that still higher order methods are probably unwarranted; nevertheless the continuous generation of smooth but fine-scaled structures demands a scheme which breaks down in a very graceful way. The PPB scheme is complicated, but the simplicity of Eq. (1) allows this complicated, moment conserving approach to be applied without undue programming difficulty. Problems of similar nature which are encountered in astrophysics and where the PPB approach would be useful are radiation transport and the treatment of nonthermal particles such as cosmic rays or neutrinos (in supernovae) which are loosely coupled to the thermal gas.

PRACTICAL LIMITS TO THE ACCURACY OF A DIFFERENCE SCHEME FOR HYDRODYNAMICS

For hydrodynamical simulations the practical limits placed upon a difference scheme are even stronger than for the case of advection. As noted earlier, the nonlinearity of the equations limits the accuracy to second order if the calculation is performed in a sequence of 1-D passes. However, if strong pressure waves are present the nonlinearity of the equations can effectively limit the accuracy to first order. This limitation is caused by discontinuities in the flow. Strong pressure waves can steepen into shocks, and these can interact to produce contact discontinuities. At these discontinuities the differential equations describing smooth flow are not obeyed, and the results of difference approximations to them are meaningless. The best one can expect from a difference scheme is a smeared out representation of a discontinuity which has a width of a zone (or a few zones) and which describes approximately the correct jumps in the flow variables across the discontinuity. Much labor can be expended to make the numerical representation of a discontinuity thin, but the width of the smeared out jump must scale linearly with Δx . Therefore, in so far as this width affects the overall accuracy of the calculation, that accuracy is limited to first order. In many flows of interest, the representation of discontinuities is crucial; obviously for these flows it would be a waste of time to build a difference scheme along the lines of PPB. It is just for such flows that the PPM hydrodynamics scheme was designed.

An excellent example of the first-order errors which can be caused by the numerical treatment of shocks is the simple problem of the propagation of a strong shock through a uniform gas described by a nonuniform grid. For simplicity, we will assume that the calculation is performed in Lagrangian coordinates and that the masses Δm_i of the $3N$ zones obey the following relations:

$$\begin{aligned} R &= 100^{1/N-1}, & \Delta m_1 &= (1-R^{-1})/(1-R^{-N}), \\ \Delta m_i &= \Delta m_{i-1}/R, & \text{for } i &\leq N, \\ \Delta m_i &= \Delta m_{2N+1-i}, & \text{for } N < i &\leq 2N, \\ \Delta m_i &= 1/N, & \text{for } N > 2N. \end{aligned}$$

We will assume that the left-hand boundary of the grid at $x = 0$ is a reflecting wall and that all gradients vanish at the right-hand boundary at $x = 3$. The initial conditions are uniform

flow with $u = -1$, $\rho = 1$, $p = 10^{-6}$, and the equation of state is $p = (2/3)\rho\epsilon$ (thus $\gamma = 5/3$). The solution to this flow problem consists of the motion of a shock from left to right across the grid at a velocity of $4/3$ with respect to the fluid and with a post-shock velocity of zero, density of 4, and pressure of $4/3$. The gas is brought to rest by this shock and all its kinetic energy is transformed into internal energy.

The numerical solution to this problem can be gravely in error when N is 20 or less. Solutions obtained with PPM in Lagrangian coordinates are displayed in Fig. 8. For these solutions, the values of N are 10, 20, 40, and 80. Zone-averaged values at time 2 are connected by straight lines in Fig. 8, and both the density and the total energy per unit mass are plotted. The exact solution is $\rho = 4$ for $m < 8/3$, $\rho = 1$ for $m > 8/3$, and $E = 1/2$ everywhere. All four calculations tend to give correct answers where the grid is uniform, but large errors in the post-shock entropy are generated where the grid is not uniform. Where the zones are gradually becoming smaller in the direction of shock propagation, not enough heat is generated in the shock, and the gas is overcompressed. Where the zones become larger, too much heat is generated in the shock, and the gas is undercompressed. At $m = 2$ the zones suddenly become smaller, and a large overcompression is produced. The large undercompression at $m = 0$ is caused by the reflecting wall there. Neither of these latter two errors is diminished when the grid is refined.

The peculiar behavior of the PPM scheme which is shown in Fig. 8 can be understood without a knowledge of the detailed workings of that difference scheme. It is only necessary to realize that Eqs. (3), (4), and (12) are differenced in conservation form. That is to say that at each interface L between two zones numerical approximations to the time-averaged fluxes $\bar{u}_L$, $\bar{p}_L$, and $(\bar{up})_L$ of specific volume, momentum, and total energy are computed. Thus for example

$$\bar{u}_L = \frac{1}{\Delta t} \int_0^{\Delta t} u_L(t) dt, \quad (34)$$

$$(\bar{up})_L = \frac{1}{\Delta t} \int_0^{\Delta t} u_L(t) p_L(t) dt. \quad (35)$$

If the differential equations (3), (4), and (12) are integrated over a rectangle in space-time with sides Δm and Δt , exact difference equations in conservation form are obtained:

Fig. 8 The propagation of a strong shock through a nonuniform grid using the PPM scheme in Lagrangian coordinates. The exact cancellation of large errors which normally guarantees that correct shock jumps are computed is vitiated by the nonuniformity of the grid, and incorrect post-shock entropies are therefore computed. Results for grids of 30, 60, 120, and 240 zones, as described in the text, are displayed.

Fig. 9 Time histories of the fluxes $\bar{u}_L$ and $(\bar{u}_p)_L$ (solid lines) and the fluxes $\bar{u}_R$ and $(\bar{u}_p)_R$ (dotted lines) as computed by the PPM scheme on a uniform Lagrangian grid for the strong shock in Fig. 8.

$$\begin{aligned}
 \langle v \rangle_{\text{new}} &= \langle v \rangle + (\bar{u}_L - \bar{u}_R) \Delta t / \Delta m, & (36) \\
 \langle u \rangle_{\text{new}} &= \langle u \rangle - (\bar{p}_R - \bar{p}_L) \Delta t / \Delta m, & (37) \\
 \langle E \rangle_{\text{new}} &= \langle E \rangle - [(\bar{u}p)_R - (\bar{u}p)_L] \Delta t / \Delta m. & (38)
 \end{aligned}$$

Here the subscripts L and R refer to the left- and right-hand interfaces of the zone, while the subscript "new" refers to the new time level. For a conservative difference scheme the numerical approximation consists entirely of a means of estimating the time-averaged fluxes $\bar{u}_L$, $\bar{p}_L$, and $(\bar{u}p)_L$. The advantage of conservative differencing is that Eqs. (36)-(38) are valid across discontinuities in the flow, when the differential equations break down.

If a difference scheme is dissipative, as it must be to convert kinetic to internal energy in a shock, and if it is conservative, then it must give precisely the correct jumps for a shock propagating through a uniform medium on a uniform computational mesh. The reason for this is that the translational symmetry of the problem guarantees that the same fluxes will be computed at both interfaces of a zone, but they will be computed a time interval $\Delta m / (\rho_0 v_s)$ later at one interface, say interface R. Here ρ_0 is the preshock density and v_s is the shock speed. Now to find the value of the jumps in v , u , and E for the zone as the shock passes we will integrate Eqs. (36)-(38) numerically using very small values for Δt . This procedure will give nearly continuous, smooth curves for the fluxes $\bar{u}_L$, $\bar{p}_L$, $\bar{p}_R$, $(\bar{u}p)_L$, and $(\bar{u}p)_R$ as functions of time. Such curves generated by the PPM scheme are shown in Fig. 9. These curves were constructed in the uniform section of the grid for the shock problem shown in Fig. 8. The dissipative character of the scheme guarantees that constant post-shock values are eventually attained. Integrating Eqs. (36)-(38) over the time interval 0 to t , shown in the figure and realizing that v , u , and E have their preshock values v_0 , u_0 , and E_0 at $t = 0$ and their post-shock values v_1 , u_1 , and E_1 at t_1 , we find

$$\begin{aligned}
 v_1 - v_0 &= - (u_1 - u_0) / \rho_0 v_s, & (39a) \\
 u_1 - u_0 &= (p_1 - p_0) / \rho_0 v_s, & (39b) \\
 E_1 - E_0 &= (u_1 p_1 - u_0 p_0) / \rho_0 v_s. & (39c)
 \end{aligned}$$

On the right we have replaced $\bar{u}(0)$ by u_0 , $\bar{u}(t_1)$ by u_1 , and similarly for p and (up) . Augmenting these equations with the equation of state gives a full set, and all the shock jumps and

the shock speed v_s may be determined. Because Eqs. (36)-(38) are exact, and because the detailed structure of the fluxes $\bar{u}(t)$, $\bar{p}(t)$, and $\bar{up}(t)$ in Fig. 9 cancelled out upon integration due to translational symmetry, Eq. (39) and the other equations determined in this way must be exact. This is an amazing fact; it implies that crude difference approximations to the derivatives in Eqs. (3), (4), and (12) may be used, and with a little care one can obtain the right answer under conditions where those differential equations themselves break down.

The above arguments depend critically upon the symmetry which guarantees that the curves $(\bar{up})_L(t)$ and $(\bar{up})_R(t)$ represented in Fig. 9 are identical except for a phase shift. This symmetry guarantees that the gross error which must be incorporated in the time integral of $(\bar{up})_L(t)$ for any difference scheme will be exactly cancelled by a compensating gross error in the integral of $(\bar{up})_R(t)$. Anything which causes this symmetry to be broken will upset this cancellation of errors, and the computed shock jumps will be wrong. This has happened in the runs shown in Fig. 8. In this case, the symmetry breaking at the reflecting wall at $m = 0$ is most easily grasped. At this wall $\bar{u}_L$ is always set to zero, the exact value. Thus $(\bar{up})_L$ is also zero and exact. However, $\bar{u}_R$ and $(\bar{up})_R$ for the zone next to the wall are very inexact, and all of this error must appear in the post-shock values of V and E , hence e , for this zone. As the grid is refined, the errors in $\bar{u}_R$ and $(\bar{up})_R$ are spread over shorter time intervals, but then so is the integration to obtain V_1 and E_1 . Thus the errors in V_1 and E_1 are large and independent of zone size. These errors do obviously depend upon the thickness of the numerical shock structure relative to the zone size, so this structure should be made as thin as possible. A glance at Fig. 8 shows that for PPM this structure is indeed close to the limit imposed by the width of a single zone.

The overheating at the reflecting wall in Fig. 8 can be understood because $\bar{u}_R$ must always be negative, while $\bar{p}_R$ must always be positive if the shock profile is monotone (i.e. there is no ringing). Then $(\bar{up})_R$ is always negative, while it should actually vanish. Consequently $(\bar{up})_R$ in Eq. (38) contributes a spurious positive energy to the first zone near the wall. The time integral of up must always be negative and should be roughly proportional to the numerical shock width, because up vanishes outside the numerical shock structure. Now, the dissipative mechanisms in the PPM difference scheme, which cause shocks to be spread out on the mesh, scale with the zone width. The result is a shock structure which is always about a zone wide. Therefore in the first region of the grid in Fig. 8, where the zones are decreasing in size as m increases, the time integrals of up at successive zone interfaces must also be decreasing in magnitude. Equation (38) then implies that the

computed post-shock energy should be too low, as indeed it is. Where the zones increase in size the sign of this effect is flipped, and the computed energy is too high. As the grid is refined the ratio of neighboring zone widths is reduced toward unity in a nearly linear fashion. This causes the noncancellation of errors in the time integrated energy fluxes at the different sides of each zone to be diminished in a nearly linear fashion. This brings about the roughly linear convergence of the solution which is shown in Fig. 8 away from the wall and the discontinuous jump in zone width at $m = 2$.

The above long discussion shows how a difference scheme which is formally second-order accurate may nevertheless converge only linearly to the solution of a problem involving discontinuities. The problem in Fig. 8 is not as artificial as it may seem. Zoning irregularities are unavoidable in many practical problems. An extreme case which appears frequently is the central region of a uniform grid in a cylindrical or spherical radial coordinate. Then factors of r or r^2 appear in the flux integrals and destroy the cancellation of errors. Error cancellation is also destroyed when two discontinuities collide. Then glitches like those near $m = 2$ in Fig. 8 can be generated. Finally, it should also be noted that PPM is a good difference scheme. Results of, say, the standard von Neumann-Richtmyer scheme on a problem like that in Fig. 8 are much worse (see for example Noh et al., 1979).

A somewhat extreme one-dimensional, hydrodynamical test problem which brings these considerations into better focus is the collision of two strong blast waves discussed in Woodward (1982) and exhaustively studied in Woodward and Colella (1983). A wave diagram for this problem is shown in Fig. 10. In this diagram the bunching together of many density contours marks the trajectory of a shock or contact discontinuity, while the spreading of density contours marks a rarefaction fan. At time zero the density is unity and the velocity vanishes everywhere on the unit spatial interval of the problem. For $x < 0.1$, $p = 1000$; for $x > 0.9$, $p = 100$; and for $0.1 < x < 0.9$, $p = 0.01$. Reflecting walls are located at $x = 0$ and $x = 1$. The gas has a gamma-law equation of state (Eq. 14) with $\gamma = 1.4$. The hot regions near the walls expand, driving strong shocks into the cold central gas. The hot gas is separated from the shells of dense, shocked gas by two contact discontinuities. At about $t = 0.028$ the shocks collide and cause a new contact discontinuity to be created. This contact discontinuity is soon strongly accelerated to the right by a rarefaction created by one of the strong shocks breaking out of the dense, previously shocked gas into the near vacuum next to the right-hand wall. Much of the wave interaction takes place in an extremely small region of space-time which is enlarged in the lower portion of Fig. 10.

RHD(x,1) d1=1.68e-85 PPKLR 6/1/83-1 wave1
68 contours: 1.528e-81 to 2.699e-81 level(1,1)/level(1,1) = 1.092

RHD(x,1) d1=1.68e-85 PPKLR 6/1/83-1 wave3
68 contours: 3.185e-81 to 2.747e-81 level(1,1)/level(1,1) = 1.078

Fig. 10. Contours of density on the space-time plane for the interacting blast wave problem of Woodward (1982) and Woodward and Colella (1983). Two strong shocks approach each other closely followed by contact discontinuities. Rarefaction fans reflected from the walls interact with these waves and with each other in a complicated fashion. The region of this interaction is enlarged in the lower part of the figure. The creation of a new contact discontinuity by the collision of the two shocks can be clearly seen in this enlargement.

The results in Fig. 10 may be regarded as exact. They have been generated on a nonuniform grid of 3096 zones with the smallest zone width set to $1/9600$ in the region of the shock collision. In addition, a five-zone section of the grid around each discontinuity is refined by an additional factor of 8, and the two principle contact discontinuities are treated as specially tracked interfaces between distinct fluids. The resulting solution is good to 1% everywhere, and is much better in most places. More detailed displays of this calculation are given in Woodward and Colella (1983). This interacting blast wave problem has been designed as a worst case to challenge Eulerian difference schemes to the utmost. The solution demands accurate treatment of the resolution of the jumps in the initial data into separate discontinuous waves and of the interaction of these waves with other discontinuous and strong continuous waves. A method which can solve this problem well should be able to handle just about anything which can arise in one-dimensional pure hydrodynamic flow.

PPM is such a scheme. The comparison with many other methods in Woodward and Colella (1983) makes this abundantly clear. However, the rate of convergence of PPM on this problem is linear. PPM contains many second-order features, many third-order features, and even a fourth-order feature, but for problems dominated by the interactions of discontinuous waves, as is this one, these high-order features serve mainly to keep the discontinuous waves sharp; they do not give high-order convergence. In fact, the main purpose for the use of parabolic interpolation in PPM is the sharp resolution of contact discontinuities. In Fig. 11 results of PPM solutions of the blast wave problem for uniform Eulerian grids of 200 and 1200 zones are presented at time 0.038. In this figure the computed zone-averaged densities are shown as dots while the "exact" solution is drawn as a solid line for comparison. These results show that the shocks and contact discontinuities are only about one zone wide. Nevertheless, the contact discontinuity formed by the shock collision contains 8 zones, and the jump at this discontinuity is not correctly computed on the 200-zone grid. A variety of solutions of lower quality are presented for this problem, along with quantitative measures of the errors and convergence rates in Woodward and Colella (1983). Here it suffices to say that PPM is the only one of six schemes studied in that article which converges as rapidly as linearly on this extremely difficult problem.

The results for this interacting blast wave problem are intended to discourage any overzealous pursuit of additional terms in Taylor expansions or additional moments to be conserved within zones. The intent is not to dampen the enthusiasm behind such pursuits but instead to redirect it in more profitable avenues, so far as hydrodynamical problems are concerned. One

such avenue has already been mentioned in passing; that is the technique of adaptive local mesh refinement. This technique has been quietly pursued by many people for many years, but recently it has been more systematically developed and more openly and systematically discussed by Gropp (1980), Berger (1982), and others. The simplest application of the technique is to refine the grid in 4 or 5 zones around each discontinuity in the flow. The grid should be refined in both space and time in order to keep the calculation both efficient and explicit. For PPM a mesh refinement factor of 8 is recommended, while for less powerful methods larger factors should be used. In the interacting blast wave problem discussed above, this procedure plus treating the problem as involving three distinct fluids brings the error for a 100-zone uniform Eulerian reference grid down below that of the unalloyed scheme using 1200 uniform Eulerian zones. Obviously this is a powerful technique, and it has the very attractive feature that it can be carried out in two dimensions quite easily.

Fig. 11 PPM results at time 0.038 for the interacting blast wave problem shown in Fig. 10. Results for uniform Eulerian grids of 200 to 1200 zones are plotted as dots against an extremely accurate solution (solid lines). Note especially the sharpness of the contact discontinuities. A linear convergence of the integrated error for this problem is observed, despite the formal second-order accuracy of the PPM scheme.

THE TREATMENT OF DISCONTINUITIES IN ADVECTION SCHEMES -- MONOTONICITY CONSTRAINTS AND CONTACT DISCONTINUITY STEEPENERS

We have seen that the most important limitation to the accuracy of hydrodynamical difference schemes is the necessity to represent discontinuities in the flow. One of the principal means of doing so may be discussed in the simpler context of the advection equation (19). This discussion will explain the treatment of flow discontinuities in the remap step of an Eulerian hydrodynamics calculation, their treatment in the Lagrangian step will be discussed later.

In Figs. 12 and 13 PPM and PPB advection calculations are presented which show the generation of oscillations near sharp jumps in the advected distribution. These oscillations always appear when high-order interpolation polynomials are used to interpolate between undersampled data. The oscillations are much less severe for the PPB scheme because the additional moment data provided for this scheme allows the smooth structure of the jumps to be nearly resolved on the finer grids. The format for Figs. 12 and 13 is similar to that of Figs. 2 and 3. Periodic boundary conditions are applied on the unit interval, the advection velocity is unity, and the results are displayed at times 0, 1, 10, and 100. The internal zone structures generated by the PPM and PPB schemes for grids of 8, 16, 32, 64, and 128 zones are plotted; and a Courant number of 0.08 was used so that these results represent a worst case. The advected function is

$$f_0(x) = 1 / \{1 + \exp [80(|x - 1/2| - 0.15)]\} . \quad (40)$$

The two jumps in this function are quite sharp; 95% of each jump is spread over a distance of only 0.092, and 50% is spread over only 0.028. On the finest grid these intervals are described by about 12 and 3.6 zones, respectively. This resolution proves insufficient to avoid oscillatory behavior for either the PPM or the PPB schemes, although the oscillations for the PPB scheme on the finest grid are very slight and appear only at the latest time shown in Fig. 13. Other tests reveal that with twice this resolution the PPB scheme will propagate these sharp features properly. However, if we are willing to give up the formal high-order accuracy of such a scheme, excellent results may be obtained at much less cost. The trick is to know when to give up formal accuracy and to what degree.

For the simple case of advection the source of oscillations near sharp jumps can be readily identified. The plots in Figs. 12 and 13 for time 0 reveal interpolated structures which are not monotone increasing (decreasing) even though they have been

constructed from monotone data. These over- and undershoots in the interpolated internal zone structures eventually give rise to over- and undershoots in the zone-averaged data. It is clear that if the interpolated function $f(x)$ is monotone, then regardless of which bins are used to construct new zone-averaged data (that is, regardless of the value of the Courant number) this new data must also be monotone. It is also clear that the first-order advection scheme which sets $f(x)$ constant within each zone must always preserve the monotonicity of its initial data. van Leer (1977) realized that an advection scheme may be made to preserve the monotonicity of its initial data if any non-monotone interpolated zone structures are flattened so that they become monotone. This flattening process can be viewed as a local

Fig. 12 Square wave advection tests using PPM without monotonicity constraints.

blending of the high-order scheme with the first-order scheme which uses constant zone structures. In order to treat the structure of each zone independently, the structures in neighboring zones are ignored when the flattening factor for the zone of interest is determined. This leads to van Leer's monotonicity constraint: no value interpolated within a zone shall lie outside the range defined by the zone averages for this zone and its two neighbors. If one wants to take the trouble, this constraint may be relaxed slightly to read as follows: the average density within the section of a zone to be advected into a neighboring zone and in the section which remains in the original zone for this timestep must lie within the range defined by the zone averages for this zone and its

Fig. 13 Square wave advection tests using PPB without monotonicity constraints.

Fig. 14a Interpolated and flattened zone structures appropriate for the representation of a sudden jump. This flattened structure will give monotone advection regardless of the timestep.

Fig. 14b Interpolated and flattened zone structures appropriate for monotone advection to the right at a Courant number of 0.5. The shaded region of the zone which is not advected into the zone on the right has the same average density as the zone on the left.

two neighbors. These monotonicity criteria are illustrated in Fig. 14. A similar monotonicity constraint for advection schemes was developed earlier by Boris and Book (1973), but that algorithm is not so easily understood and it has been much less successful when applied to other problems such as hydrodynamics.

An unfortunate characteristic of these monotonicity constraints is that they tend to alter a waveform dramatically once the decision is made to switch them on. This is illustrated by the Gaussian advection results in Fig. 15. These may be directly compared with the results in Fig. 2b, from which they differ only by the use of monotonicity constraints and contact discontinuity steepeners to be described shortly. The results in Fig. 15 are definitely an improvement over those in Fig. 2b, but the shape of the waveform has been completely lost in the process of keeping it monotone. One might well prefer the non-monotone results of the PPB scheme which are shown in Fig. 3b. In fact, the PPB results are so good that if any monotonicity constraints are to be applied to this scheme at all they must be applied with great care in order not to lose the advantage of this scheme's accuracy in smooth flow. In particular, the monotonicity constraints of van Leer (1977) and of Boris and Book (1973) cause zones near local extrema to be completely flattened. For the PPB scheme this is a disaster. The advantage of PPB is that its extra moment data allows the accurate treatment of very short wavelength Fourier modes. For these modes flattening zone structures near the frequent extrema results in gross wave distortion and damping.

For the PPB scheme it is worthwhile to expend some computer time to make sure that monotonicity constraints are really needed before they are applied. This procedure may be useful for other schemes as well, but for PPB it is essential if the constraints are to be applied at all. The entire procedure more than doubles the running time, so if the unconstrained scheme will do, it should be used. The monotonicity constraints described above modify the advection algorithm in regions where the second derivative of the advected function is large compared to the first derivative -- more precisely, when

$$\Delta x \left| \frac{\partial^2 f}{\partial x^2} \right| > \left| \frac{\partial f}{\partial x} \right| .$$

For schemes employing parabolic interpolation, large second derivatives of f do not necessarily pose any special problem. Therefore the monotonicity constraint may often be inappropriate, especially near extrema, where $\partial f / \partial x$ vanishes.

The key to modifying the monotonicity constraint is not to measure the monotonicity of the interpolated zone structure but instead to measure the ability of that structure when extrapolated into the neighboring zones to predict the correct zone-averaged values there. This will be a more direct measure of the extent to which the density distribution is undersampled by the grid than is the monotonicity criterion described above. In fact, both criteria will locate regions in which short wavelength Fourier modes play a large role, but the monotonicity criterion will also locate perfectly innocent extrema as well.

To measure the quality of the interpolation parabola given by Eqs. (22) and (25) (or by Eqs. (22)-(24) for PPM) we proceed as

Fig. 15 Gaussian advection tests using PPM with monotonicity constraints and contact discontinuity steepeners.

follows. First we check the extrapolation $f(x)$ into the zone on the left (subscript ZL). This extrapolation gives a guess $\langle f \rangle_{XZL}$ at the neighboring zone average $\langle f \rangle_{ZL}$. For a uniform mesh we have

$$\langle f \rangle_{XZL} = f_0 - f_1 + \frac{13}{12} f_2 \quad (41)$$

The quality Q_{ZL} of this extrapolation is then defined as

Fig. 16 Gaussian advection tests using PPB with monotonicity constraints under the control of a "quality" criterion. The Courant number is 0.8.

$$Q_{ZL} = \frac{|\langle f \rangle_{XZL} - \langle f \rangle_{ZL}|}{|\langle f \rangle_{XZL} - \langle f \rangle| + 0.005 (|\langle f \rangle_{ZL}| + |\langle f \rangle|)} \quad (42)$$

The extrapolated average $\langle f \rangle_{XZR}$ and its quality Q_{ZR} are then computed similarly for the zone on the right. The overall quality Q of the interpolated zone structure $f(x)$ is then taken to be

$$Q = \max \{Q_{ZL}, Q_{ZR}\} \quad (43)$$

If Q exceeds a threshold value of 0.4 (a more conservative figure is 0.2) in this zone or in either of its neighbors, then the interpolated zone structure may be modified according to the less restrictive of the two monotonicity constraints described above. Gaussian advection tests with PPB using this quality criterion and monotonicity constraint are shown in Fig. 16. These results are to be compared with those in Figures 3b and 15. The comparison with Fig. 15 shows that the quality criterion is smart enough to leave the results on the three fine grids essentially unchanged. Whether or not the coarse grid results are superior to those in Fig. 3b is a matter of taste and would depend upon the particular application as well. If the threshold for Q is raised above 0.4, only the 8-zone run will be affected by the monotonicity constraint. The threshold of 0.4 causes PPB to clip a pure sine wave on a grid of 3 zones or less, while a threshold of 0.2 clips the wave on grids of 4 zones or less.

Treatment of Discontinuities in PPM

The technique for treating discontinuities in PPM is most easily explained in terms of a specific example. Formulae for more general cases may be found in Colella and Woodward (1983). We will follow the scheme as it constructs a representation $f(x)$ from zone-averaged data $\langle f \rangle_i$ corresponding to an $f_0(x)$ which is abstracted from the square wave advection tests in Figs. 12 and 13:

$$f_0(x) = \tanh(x) \quad (44)$$

In Fig. 17 the six zones near the origin are shown, with $\Delta x = 3$ and the leftmost zone interface at $x = -10$. The dotted lines in Fig. 17a show the interpolation parabola constructed by the PPM scheme described earlier and used in the Gaussian advection tests in Fig. 2. The values of f at zone interfaces, f_L and f_R , have been generated from Eq. (23). These interface values are attained by the unique cubic curves which have the same average values as $f_0(x)$ in the four zones closest to each interface. The parabola within zones are constructed to have these interface values and the same zone-averaged values as $f_0(x)$. Note that the resulting interpolated $f(x)$ is not monotone. In Fig. 17b a construction

of zone-averaged slopes Δf_p is illustrated. The dotted lines show the zone averaged slopes for the parabolae which have the same zone-averaged values as $f_0(x)$ in each zone and its two neighbors. For a uniform grid Δf_p is given by

$$\Delta f_p = \frac{1}{2} (\langle f \rangle_{ZR} - \langle f \rangle_{ZL}) \quad (45)$$

The solid lines in Fig. 17b show monotonized slopes Δf_M . These are constructed by reducing Δf_p toward zero if necessary to keep the interpolated values implied by these slopes within the range of the three zone averages $\langle f \rangle_{ZL}$, $\langle f \rangle$, and $\langle f \rangle_{ZR}$.

The solid lines in Fig. 17a illustrate the interpolation parabolae implied by a provisional set of interface values computed by the full PPM scheme. Like the dotted lines in Fig. 17a these are also obtained by evaluating cubic polynomials at the zone interfaces. These interface values are determined by rewriting the interpolation formula in Eq. (23) in terms of $\langle f \rangle_{ZL}$, $\langle f \rangle$, Δf_{pZL} , and Δf_p and then substituting Δf_{MZL} and Δf_M for the two slopes:

$$f_L = \frac{1}{2} (\langle f \rangle_{ZL} + \langle f \rangle) - \frac{1}{6} (\Delta f_M - \Delta f_{MZL}) \quad (46)$$

Thus if the original slopes Δf_{pZL} and Δf_p are unchanged by the application of the monotonicity constraint, Eq. (46) gives the same result as Eq. (23). In our example only the slope for the zone nearest the origin is unaffected by the monotonicity constraint. The use of Eq. (46) gives values of f_L which are more reasonable than those obtained from Eq. (23). As the solid lines in Fig. 16a demonstrate, each value f_L lies between the neighboring zone averages $\langle f \rangle_{ZL}$ and $\langle f \rangle$. This monotonicity property is guaranteed by the use of monotonized slopes in obtaining f_L . A further improvement resulting from the use of Eq. (46) is the slight steepening of the implied internal structure for the zone nearest the origin. This is indeed a small effect, but such effects accumulate over thousands of timesteps and can eventually mean the difference between a contact discontinuity which is three zones wide and one which is five zones wide.

The full PPM algorithm includes a contact discontinuity detector and steepener which will cause the implied structures for the two zones between $x = -4$ and $x = 2$ to be steepened further. This part of the PPM algorithm is responsible for the very thin contact discontinuities in Fig. 11. It is also the cause of the very steep, incorrect structures in Fig. 15. The idea behind this part of the algorithm is to detect regions of

the flow in which the density jumps so sharply that these regions may be considered to be contact discontinuities (for hydrodynamics, we demand that the density jump not be accompanied by a pressure jump). When these regions are detected, the PPM

Fig. 17 (a) Interpolation parabolae for the unconstrained scheme used for the tests in Figs. (2) and (12) are shown as dots. The solid lines are generated from interface values constructed using monotone zone-averaged slopes. (b) Unconstrained (dotted) and monotone (solid) zone-averaged slopes. (c) Unsteepened (dotted) and steepened (solid) interpolation parabolae. (d) Unmonotonized, steepened parabolae (dotted) and monotone, steepened parabolae (solid).

algorithm causes them to be continually steepened, whether they are actually contact discontinuities or not. In doubtful cases some choice must be made, and it may be wrong, but we are careful to guarantee that nothing involved in the treatment of detected discontinuities will vitiate the formal second-order accuracy of the scheme in well-resolved, smooth flow. The behavior in the Gaussian advection tests in Fig. 15 results from the clipping of the peak of the Gaussian due to the monotonicity constraint. This introduces sudden jumps near the clipped peak which are detected as contact discontinuities and thus steepened still further. This behavior ceases when the grid is sufficiently fine. The point at which this transition occurs may be altered by adjusting the dimensionless constants in the contact discontinuity detector.

The action of the contact discontinuity steepener is illustrated in Fig. 17c. Here the dotted lines are the unsteepened interpolated structures which were drawn as solid lines in Fig. 17a. The solid lines in Fig. 17c are the steepened interpolated structures. Steepening occurs only for the two zones nearest the origin. For this case of uniform zone widths the discontinuity detection algorithm is as follows (more general formulae are given in Colella and Woodward 1983). First, an estimate of the second derivative of f is computed for each zone:

$$\Delta_2 f = \langle f \rangle_{ZR} - 2 \langle f \rangle + \langle f \rangle_{ZL} \quad . \quad (47)$$

Then a provisional steepness parameter S_0 is computed which compares the magnitudes of the first and third derivatives in each zone:

$$S_0 = [(\Delta_2 f)_{ZR} - (\Delta_2 f)_{ZL}] / [6 (\langle f \rangle_{ZL} - \langle f \rangle_{ZR})] \quad . \quad (48)$$

Note that S_0 is positive for sudden jumps in f , but it is negative for small plateaus in f . To locate these sudden jumps we must also demand that $\Delta_2 f$ change sign in crossing the region of the jump. Therefore we define a revised steepness parameter S_1 by

$$S_1 = S_0 \quad , \quad \text{if } (\Delta_2 f)_{ZL} (\Delta_2 f)_{ZR} < 0 \quad (49)$$

$$= 0 \quad , \quad \text{otherwise.}$$

We also demand that S_1 be zero unless some reasonable measure of the fractional change of f exceeds a small threshold of, say, 1%. This will avoid the steepening of tiny glitches of numerical origin. The final step in the detection procedure is to make

some quantitative decision as to the steepness which warrants special treatment as a discontinuity. The quantity S_1 is scaled from 0 to 1 via:

$$S_2 = \max \{0, \min \{1, 20 (S_1 - .05)\}\} . \quad (50)$$

The dimensionless constants in this formula have been chosen to guarantee that discontinuities one zone in width will be steepened fully (i.e., $S_2 \approx 1$ for these cases) regardless of their phase with respect to the zone boundaries. In our specific example in Fig. 17c, this criterion causes steepening to be applied only to the two zones nearest the origin. For the zone nearest the origin this steepening is substantial.

When the steepening parameter S_2 is nonzero, the previously interpolated interface values f_L and f_R are blended with values f_{SL} and f_{SR} which will generally give a steeper structure within the zone. The blending factor is just S_2 . The values f_{SL} and f_{SR} are obtained by extrapolating the monotone zone-averaged slopes in neighboring zones (see Fig. 17b) to the zone interfaces. Thus

$$f_{SL} = \langle f \rangle_{ZL} + \frac{1}{2} \Delta f_{MZL} , \quad (51a)$$

$$f_{SR} = \langle f \rangle_{ZR} - \frac{1}{2} \Delta f_{MZR} . \quad (51b)$$

Note that these interface values are formally second-order accurate, although they are intended for use in underresolved regions where order of accuracy is a concept of limited value.

The final step of the PPM construction is shown in Fig. 17d. This figure illustrates the application of a monotonicity constraint, first used by the PPM scheme in Woodward and Colella (1981). The zone averages of f , $\langle f \rangle$, and the interface values f_L and f_R determine parabolic internal structures for the zones of the form given in Eq. (22) through the relations in Eq. (24). These parabolas will be monotone so long as the absolute value of f_2 is less than that of f_1 . For the marginal cases where these absolute values are equal we have either

$$f_L = f_{ML} = 3 \langle f \rangle - 2f_R \quad (52a)$$

or

$$f_R = f_{MR} = 3 \langle f \rangle - 2f_L . \quad (52b)$$

When one of these marginal values is exceeded, either f_L or f_R is reset to that marginal value so that the implied interpolation parabola is monotone. Also, at extrema in $\langle f \rangle$ both f_L and f_R are set to $\langle f \rangle$. This flattening at extrema is done first. Then the difference of f across the zone is computed

$$\Delta f = f_R - f_L . \quad (53)$$

The criteria for resetting f_L and f_R can be written

$$f_L \rightarrow f_{ML} , \quad \text{if } \Delta f (f_L - f_{ML}) < 0 ; \quad (54a)$$

$$f_R \rightarrow f_{MR} , \quad \text{if } \Delta f (f_R - f_{MR}) > 0 . \quad (54b)$$

This monotonicity constraint has a dramatic effect on the steepened interpolation parabolae shown as solid lines in Fig. 17c and as dotted lines in Fig. 17d. The final result, the solid lines in Fig. 17d is a very sharp representation of the jump in $f_0(x)$. It is considerably sharper than the profile in Fig. 17a which is produced by the completely unconstrained PPM scheme. In fact the profile in Fig. 17d is much closer to $f_0(x)$ than is that in Fig. 17a. It is an empirical fact that the PPM interpolation procedure described above is sufficient to cause a square pulse 10 zones in width to be advected essentially without distortion. Unfortunately, as Fig. 15 demonstrates, other smoother pulses may be turned into square waves as they are advected if they are not well resolved. The philosophy here is that if such a pulse is not sufficiently resolved on the grid to be advected properly its shape may be greatly distorted but at least its position and height will be fairly well represented.

THE PPM METHOD FOR HYDRODYNAMICS

It is easiest, both conceptually and computationally, to perform hydrodynamic calculations with PPM if these are split into a sequence of one-dimensional calculations which in turn are split into Lagrangian hydrodynamics calculations followed by remap operations. Therefore we need only describe here the one-dimensional Lagrangian formulation of PPM hydrodynamics (for other formulations see Colella and Woodward 1983). The PPM hydrodynamics scheme is built upon the PPM advection scheme described above. The basic idea is to exploit the relationship between hydrodynamics and the advection of Riemann invariants which was discussed earlier. One would like to use the PPM advection scheme to update each family of Riemann invariants separately, but unfortunately hydrodynamics is not so simple. In general the Riemann invariants are defined only differentially through Eqs. (15), and they are not advected independently of one another. However, it is still true that the time dependence of the flow at any point in the fluid is determined by two separate domains of dependence bounded in space-time by the trajectory of that point and by two characteristic paths corresponding to pressure waves travelling toward the point from either side (see Fig. 18a). The information contained in each of these domains of dependence which affects the time evolution of the point in question still basically refers to only one of the two Riemann invariants. The approach in PPM is to let a nonlinear Riemann

solver handle any difficulties in sorting out what aspects of the information in each domain of dependence correspond to which Riemann invariant and to what extent the information in the two separate domains should interact. Such use of a nonlinear Riemann solver in hydrodynamic calculations was first introduced by Godunov (1959) and was first adapted for higher-order difference schemes by van Leer (1979).

A nonlinear Riemann solver is an essential ingredient in PPM. It computes the solution to Riemann's problem, that is the nonlinear interaction of two constant states of the fluid. The initial conditions for Riemann's problem are two constant states separated by a discontinuous jump. This sort of initial condition is found in the blast wave problem discussed earlier

Fig. 18a A space-time diagram in Lagrangian coordinates showing the domains of dependence for zone interface L . Paths of sound waves travelling to the right (left) are drawn as solid (dotted) lines. These trace out the domain of dependence $L+$ ($L-$). The characteristic equations (15a) hold along these paths if the flow is smooth. These paths have slopes equal to the local sound speed. This is approximated by PPM as $\langle C \rangle_{ZL}$ or $\langle C \rangle$ for the solid or dotted paths, respectively. The interaction of these sound waves produces the time evolution of u_L and p_L . The averages of these over the time interval of Δt shown here are the fluxes $\bar{u}_L$ and $\bar{p}_L$ used by PPM in the conservation laws in Eqs. (36) and (37).

and displayed in Fig. 10. The early evolution of that problem illustrates the character of the general solution to Riemann's problem (see also Fig. 18b). In general two nonlinear sound waves, either shocks or rarefactions, emerge from the initial jump, leaving a contact discontinuity behind there. The solution is self-similar, and in particular the state at the location of the initial jump is constant in time. A space-time diagram for such a Riemann problem is shown in Fig. 18b. The sound wave paths, or characteristic curves in Fig. 18b are straight lines, but they kink as they cross the shock and rarefaction waves. The algorithm for solving Riemann's problem involves an iteration, as will be described shortly.

We have already noted that the PPM hydrodynamics scheme is cast in strict conservation form, so that it will compute shock

Fig. 18b A space-time diagram showing the use of a Riemann solver in PPM to estimate the time-averaged fluxes $\bar{u}_L$, $\bar{p}_L$, and $(\bar{u}p)_L$ at zone interface L . A case similar to that occurring in the blast wave problem in Fig. 10 is shown. Constant states for the Riemann problem are obtained by averaging the interpolated initial data over the domains $L+$ and $L-$. The nonlinear interaction of these constant states produces a shock and a rarefaction wave and a contact discontinuity at interface L . The velocity and pressure $\bar{u}_L$ and $\bar{p}_L$ at the interface are constant in time and serve as estimates of the true time-averaged fluxes.

jumps correctly on uniform grids in otherwise smooth, well resolved flow. This means that PPM uses Eqs. (36)-(38) to update the zone-averaged specific volume V , flow velocity u , and specific total energy E . The remainder of the PPM scheme thus boils down to an algorithm for computing time-averaged fluxes $\bar{u}_L$, $\bar{p}_L$, and $(\bar{u}p)_L$ of zone volume, momentum, and total energy at the zone interfaces. This computation of the fluxes is illustrated in Fig. 18b. First, the interpolation algorithm from the PPM advection scheme described earlier is used to generate monotone internal zone structures for the variables V , u , and p . These variables are chosen for interpolation because of their close relationship to the variables appearing in the characteristic equations (15). The average Lagrangian sound speeds $\langle C \rangle_{ZL}$ and $\langle C \rangle$ in the zones adjacent to the interface L are then used to estimate the extents $\langle C \rangle_{ZL} \Delta t$ and $\langle C \rangle \Delta t$ of the two domains of dependence L^+ and L^- indicated in Fig. 18b. The interpolation parabolae are then used to compute average values $\langle V \rangle_{L^+}$, $\langle u \rangle_{L^+}$, and $\langle p \rangle_{L^+}$ of the variables within these spatial domains of dependence. These average values will be used together with the equation of state to determine the two constant states for a Riemann problem. The solution to this Riemann problem is illustrated in Fig. 18b. It yields constant values for u_L and p_L , which will serve as estimates for the true time-averages $\bar{u}_L$ and $\bar{p}_L$. The time average of the product up is then simply

$$(\bar{u}p)_L = \bar{u}_L \bar{p}_L \quad . \quad (55)$$

This procedure effectively replaces the time average of the sound wave interactions at the zone interface by the interaction of the spatial averages over the information carried by the sound waves. It is the job of the Riemann solver to correctly compute this nonlinear interaction and to sort out from the spatially averaged states just that information which the sound waves actually transport to the zone interface.

Before describing the technique used to solve Riemann's problem numerically, we first describe the means of treating a complicated equation of state. A more complicated treatment is described by Colella (1983) with examples of its use for strongly nonlinear flows in air (see also Colella and Glaz 1983). The simpler method here appears to work equally well. It is based upon a model equation of state suggested by J. LeBlanc:

$$p = p_{00} + (\gamma - 1) \rho e \quad . \quad (56)$$

Using this model equation of state we may derive formulae for the Eulerian and Lagrangian sound speeds c and C :

$$c^2 = (\gamma - 1) (\epsilon + pV) , \quad (57a)$$

$$c = \rho c . \quad (57b)$$

Here we have used the fact that an adiabatic change $d\epsilon$ in the specific internal energy is given by $-p dV$. The model equation of state (56) is meant to apply only to one zone and only for one timestep. The parameters p_{00} and γ are determined at the beginning of each timestep for each zone by reference to the true, presumably more complicated, equation of state. This true equation of state is asked to compute zone-averaged pressures $\langle p \rangle$ and sound speeds $\langle c \rangle$ from $\langle V \rangle$ and $\langle \epsilon \rangle = \langle E \rangle - \langle u \rangle^2/2$. Equations (56) and (57a) are then used to determine p_{00} and γ :

$$p_{00} = \langle p \rangle - \frac{\langle \epsilon \rangle \langle c \rangle^2}{\langle V \rangle (\langle \epsilon \rangle + \langle p \rangle \langle V \rangle)} . \quad (58a)$$

$$\gamma = 1 + [\langle c \rangle^2 / (\langle \epsilon \rangle + \langle p \rangle \langle V \rangle)] . \quad (58b)$$

Methods for solving Riemann's problem may be found in fluid dynamics text books such as Courant and Friedrichs (1948). We will only summarize a method here. Positive-definite nonlinear wave speeds $W_{\pm}$ for discontinuous waves travelling to the right (+) or to the left (-) may be found from the model equation of state and the jump conditions in Eqs. (39), with

$$W_{\pm} = \pm \rho_0 v_s . \quad (59)$$

These jump conditions also imply a useful relation for the jump in the internal energy:

$$\epsilon_1 - \epsilon_0 = -\frac{1}{2} (p_0 + p_1) (v_1 - v_0) . \quad (60)$$

If we approximate rarefaction waves by discontinuous jumps, the solution $\bar{u}_L$ and $\bar{p}_L$ to the Riemann problem illustrated in Fig. 18b satisfies

$$(\bar{u}_L - \langle u \rangle_{L\pm}) \pm (\bar{p}_L - \langle p \rangle_{L\pm}) / W_{\mp} = 0 . \quad (61)$$

The similarity of Eqs. (61) to the characteristic equations (15a) for smooth flow is striking. In fact, in the limit of small jumps between the states L^+ and L^- , the waves in the solution to Riemann's problem are weak and $W_{\mp}$ approaches $\langle C \rangle_{L^{\pm}}$. In this limit of weak waves Eqs. (61) and (15a) become identical.

The Riemann problem can be solved iteratively by beginning with the assumption of weak waves and replacing $W_{\mp}$ in Eq. (61) with $\langle C \rangle_{L^{\pm}}$. These sound speeds are computed from $\langle V \rangle_{L^{\pm}}$ and $\langle p \rangle_{L^{\pm}}$ by using the equation of state (56) together with Eqs. (57). We may then solve for a first guess $\bar{p}_{L1}$ at $\bar{p}_L$:

$$\bar{p}_{L1} = \langle p \rangle_{L^+} + [(\langle p \rangle_{L^-} - \langle p \rangle_{L^+}) - \langle C \rangle_{L^-} (\langle u \rangle_{L^-} - \langle u \rangle_{L^+})] \\ \times \langle C \rangle_{L^+} / (\langle C \rangle_{L^-} + \langle C \rangle_{L^+}) \quad (62)$$

Using the jump conditions (39a), (39b), and (60) together with the equation of state (56) and Eq. (59), we may derive a formula for the wave speeds $W_{\mp}$ in terms of the pressure $\bar{p}_L$:

$$W_{\mp}^2 = \frac{1}{2\langle v \rangle_{L^{\pm}}} [(\gamma_{L^{\pm}} - 1) (\langle p \rangle_{L^{\pm}} + \bar{p}_L) \\ + 2 (\bar{p}_L - p_{00L^{\pm}})] \quad (63)$$

Substituting our first guess $\bar{p}_{L1}$ for $\bar{p}_L$ in these equations gives first guesses $W_{\mp 1}$ at the wave speeds. These may now be inserted in Eq. (61) in place of $W_{\mp}$ to give two estimates, $\bar{u}_{L\pm 1}$, for $\bar{u}_L$. We will use Newton's method to find the value of $\bar{p}_L$ for which the difference between $\bar{u}_{L\pm}$ vanishes. Usually only a single iteration need be performed to achieve sufficient accuracy for our purposes. Using only the first guess for $\bar{p}_L$ is usually not sufficient.

The Newton iteration for $\bar{p}_L$ is illustrated in Fig. 19. It requires a formula for $\partial \bar{p}_L / \partial \bar{u}_{L\pm}$, as can be seen in the figure. This formula is obtained by differentiating the jump conditions for discontinuous waves:

$$\frac{\partial \bar{p}_L}{\partial \bar{u}_{L^{\pm}}} = \frac{\mp 4 \langle v \rangle_{L^{\pm}} W_{\mp}^3}{4 \langle v \rangle_{L^{\pm}} W_{\mp}^2 - (\gamma_{L^{\pm}} + 1) (\bar{p}_L - \langle p \rangle_{L^{\pm}})} \quad (64)$$

Inserting the first guesses $\bar{p}_{L1}$ and $\bar{u}_{L1}$ in this formula yields $(\partial \bar{p}_L / \partial \bar{u}_{L\pm})_1$. These derivatives give a second guess $\bar{p}_{L2}$ at $\bar{p}_L$, and the iteration can then either be continued or terminated as desired:

$$\bar{p}_{L2} = \bar{p}_{L1} + (\bar{u}_{L-1} - \bar{u}_{L+1}) \left[\frac{(\partial \bar{p}_L / \partial \bar{u}_{L-})_1 (\partial \bar{p}_L / \partial \bar{u}_{L+})_1}{(\partial \bar{p}_L / \partial \bar{u}_{L-})_1 - (\partial \bar{p}_L / \partial \bar{u}_{L+})_1} \right] \quad (65)$$

If the iteration is terminated at this point, we estimate $\bar{u}_L$ by

$$\bar{u}_{L2} = \bar{u}_{L+1} + \frac{(\bar{u}_{L-1} - \bar{u}_{L+1}) (\partial \bar{p}_L / \partial \bar{u}_{L-})_1}{(\partial \bar{p}_L / \partial \bar{u}_{L-})_1 - (\partial \bar{p}_L / \partial \bar{u}_{L+})_1} \quad (66)$$

The manner in which the PPM scheme treats a strong shock is illustrated in Fig. 20 from Woodward and Colella (1983). In this figure five zones near a strong shock are shown at various stages of the calculation. This calculation uses a uniform Eulerian

Fig. 19 The Newton iteration for the solution of Riemann's problem in the case involving two strong shocks. The curved lines through the initial states represent all the states which can be connected to the initial states through discontinuous waves.

grid and the shock moves three quarters of a zone width to the right in a timestep. The gas has a gamma law equation of state (Eq. (14)) with $\gamma = 1.4$. Ahead of the shock $u = 0$, $\rho = p = 1$, and behind the shock $u = 3.4056$, $\rho = 4.4091$, and $p = 16$. In Fig. 20a the initial zone-averaged velocities $\langle u \rangle$ are indicated by dotted lines, while the monotonized PPM interpolation parabolae are represented by the solid lines. The shock profile is very steep, with essentially only one zone describing the internal structure of the shock. Figure 20b illustrates the averaging over the domains of dependence for each zone interface in order to generate appropriate Riemann problems. The dotted lines again show the interpolation parabolae, while the solid lines indicate the averages in the domains of dependence. Figure 20c shows the Riemann problems generated in Fig. 20b as dotted lines. The solutions $\bar{u}_L$ to these Riemann problems determine the effective velocity gradients $\partial u / \partial m$ which compress the zones in the Lagrangian step of the calculation. These effective velocity gradients are shown by the solid lines in Fig. 20c. In computing $\bar{u}_L$ the Riemann solver is clearly doing its job of taking values of the Riemann invariant transported by this wave only from the "upstream" domains of dependence. In each case $\bar{u}_L$ is nearly equal to the "upstream" value $\langle u \rangle_{L+}$. This suggests that we might have computed this simple result much more cheaply. In this trivial instance that is indeed true, but the Riemann solver pays for itself when strong shocks interact with other continuous or discontinuous waves. Then reliable short cuts are very difficult to devise. In Fig. 20d the remap step is displayed. The monotonized interpolation parabolae are shown by the solid lines in the Lagrangian zones, which have been compressed and translated to the right. The dotted lines represent the average values of the velocity in the original Eulerian zones. Notice that the shock is once again confined essentially to a single zone.

The ability of the PPM scheme to compute such thin numerical shock structures without introducing post-shock oscillations is critically dependent upon the use of monotonicity constraints coupled with a nonlinear Riemann solver. Consider the case shown in Fig. 20. In order to avoid overcompressing the second zone from the left in that figure, the interface velocity computed at the right-hand side of that zone must be very nearly equal to the post-shock velocity. This is indeed the case in the calculation, because the monotonicity constraint has generated an interpolation parabola for that zone which is almost flat and also the Riemann solver has essentially chosen the "upstream" velocity from the two values presented to it. As important as avoiding overcompressing the second zone from the left in Fig. 20 is the need to begin compressing the fourth zone. Again, the monotonicity constraint keeps the average velocity large in the right-hand domain of dependence of the third zone, and the Riemann solver

Fig. 20 The propagation of a isolated strong shock on a uniform grid by the PPM scheme. (a) Original zone-averaged data (dotted lines) is used to generate monotone interpolation parabolae (solid lines) within the five zones shown. (b) By averaging over the interpolation parabolae (dotted lines) within the domains of dependence of the zone interfaces, constant states (solid lines) are obtained for Riemann problems. (c) Solution of the Riemann problems (dotted lines) yields interface velocities which imply compressional velocity gradients (solid lines) for the Lagrangian zones. (d) Interpolation parabolae (solid lines) are generated for the compressed Lagrangian zones which yield upon integration the new average velocities in the original Eulerian zones (dotted lines). The shock is again one zone wide after moving $3/4$ of a zone width to the right.

chooses this velocity to be assigned to the zone interface. This gives a relatively strong compressional velocity gradient for the fourth zone.

The results in Fig. 20 demonstrate that with the shock moving $3/4 \Delta x$ in a single timestep there is very little room for error. The third Eulerian zone is accelerated very nearly to the post-shock velocity in this time step; the slightest readjustment of the fluxes at the zone interfaces could result in an overshoot. One might well ask how other difference methods without monotonicity constraints and Riemann solvers accomplish this feat. Of course, the answer is that they do not accomplish it; but they may nevertheless propagate such strong shocks without oscillations. The means by which this is achieved is an artificial viscosity. The type of artificial viscosity which is most easily described is that proposed by Lapidus (1967). It amounts to adding to each interface flux $\bar{u}_L$, $\bar{p}_L$, and $(\bar{u}p)_L$ a small diffusive flux $\bar{u}_{DL}$, $\bar{p}_{DL}$, and $(\bar{u}p)_{DL}$:

$$\bar{u}_{DL} = C_{DL} (\langle v \rangle - \langle v \rangle_{ZL}) \quad , \quad (67a)$$

$$\bar{p}_{DL} = C_{DL} (\langle u \rangle_{ZL} - \langle u \rangle) \quad , \quad (67b)$$

$$(\bar{u}p)_{DL} = C_{DL} (\langle E \rangle_{ZL} - \langle E \rangle) \quad , \quad (67c)$$

where C_{DL} is a Lagrangian diffusion speed given by

$$C_{DL} = C_{D0} \min \{ \langle \rho \rangle_{ZL} , \langle \rho \rangle \} \times \max \{ \{ \langle u \rangle_{ZL} - \langle u \rangle \} , 0 \} \quad . \quad (67d)$$

Here C_{D0} is a constant of order unity. For the example in Fig. 20, the effect of the diffusive flux $\bar{u}_{DL}$ would be to increase $\bar{u}_L$. In this example the monotonicity constraints and the Riemann solver act in this same way. In effect they behave here like a very complicated prescription for computing C_{DL} which gives precisely the desired amount of diffusive flux to prevent oscillations near the shock. It is clear that no simple formula like Eq. (67) with a constant value of C_{D0} will match this performance. Therefore the value of C_{D0} must be chosen high enough to be safe in all situations. This leads to too much diffusion most of the time and much broader shock structures than PPM delivers.

It is generally true that one designs a difference scheme to give shock structures which are as narrow as possible. However,

one can overdo this. As is illustrated in Woodward and Colella (1983) and discussed further in Colella and Woodward (1983), difference schemes which produce shocks which are too narrow may be plagued by low level noise behind the shocks, particularly in the entropy distribution. These difficulties show up most strongly in two-dimensional calculations. For PPM this problem arises when a strong shock is nearly aligned with one of the coordinate directions (in 2-D) and when it is nearly stationary with respect to the grid. For such nearly stationary shocks PPM generates extremely narrow structures. Because the fundamental frequency of any numerical noise generated in such narrow shocks is just the frequency with which the shock crosses zone boundaries, nearly stationary shocks tend to generate noise with wavelengths of several zones. The PPM scheme is so accurate in smooth flow that this noise is not effectively damped. In order to eliminate this noise the shock structures must somehow be broadened slightly.

This problem and its solution are illustrated in Fig. 21. The dotted line in the figure displays the interpolated internal zone structures of the density computed by the PPM scheme described above. Twenty zones in the neighborhood of a very strong, nearly stationary shock are shown. The shock is extremely narrow, and the density oscillations behind it are partly due to spurious sound waves travelling away from the shock and partly due to spurious oscillations in the post-shock entropy. In an earlier section we argued that a difference scheme in conservation form must compute proper shock jumps. However, the argument of that section does not apply to the case in Fig. 21 because there is not sufficient dissipation in the PPM scheme to damp the relatively long wavelength oscillations behind the shock in a reasonable distance. The cure for this situation is to prevent their generation within the shock region by broadening the shock slightly. The oscillations are generated by a periodic expansion and contraction of the numerical shock structure as the shock crosses the zones. The shock is broadest when it is centered at the middle of a zone and narrowest when at a zone interface. The oscillations are drastically reduced when the shock is spread by additional dissipative mechanisms so that its shape varies only slightly with its phase relative to the zone boundaries. Such a PPM shock profile is shown by the solid line in Fig. 21.

The extra dissipation needed to give the solid line in Fig. 21 is added to PPM by flattening the interpolated zone structures in the shock region at the outset of both the Lagrangian step and the remap step. A more complicated and more effective dissipation involving jiggling of the computational grid near stationary shocks is discussed in Woodward and Colella (1983) and in Colella and Woodward (1983). This viscosity is "smart"; it recognizes

which shock structures may remain thin and which must be broadened. It is also effective; it accomplishes the greatest damping of post-shock noise with the least spreading of the shock structure. However, it is too complicated to be described here.

The dissipation mechanism we will describe broadens all shocks, whether they need it or not, to the point where post-shock oscillations are strongly damped in the worst cases. To accomplish this, the method relies on a measure of the shock steepness, S_{shk} . This measure must be sensitive to the pressure profile, which controls the generation of spurious sound waves behind the shock. The steepness of the pressure profile is measured by comparing narrow- and wide-based differences centered on a particular zone i . Thus we define a steepness parameter S_p by:

Fig. 21. Nearly stationary structures computed with PPM for a very strong shock on a uniform grid. The preshock (subscript 0) and post-shock (subscript 1) states of this gas with $\gamma = 5/3$ are: $p_0 = 1$, $p_1 = 4$; $p_0 = 7.5 \times 10^{-7}$, $p_1 = 0.75$; $u_0 = -0.899$, $u_1 = -0.149$. The dotted line shows the interpolated density structures in 20 zones near the shock for the PPM scheme without special shock broadening. The solid line shows these structures as computed by the PPM scheme with the shock broadening mechanism described in the text. This worst-case oscillation has been reduced in amplitude by a factor of 20 by slightly broadening the numerical shock structure as shown.

$$S_p = (\langle p \rangle_{i+1} - \langle p \rangle_{i-1}) / (\langle p \rangle_{i+2} - \langle p \rangle_{i-2}) \quad (68a)$$

This parameter is scaled from 0 to 1/2 via

$$S_{01} = \max \{0, \min \{0.5, 5 (S_p - 0.75)\}\} \quad (68b)$$

In well-resolved, smooth flow S_p will have a value very close to 1/2, so it is clear that S_{01} will be nonzero only near sudden jumps in p . The choice of the dimensionless constants in the definition of S_{01} indicates that we are aiming at broadened shock structures with values of S_p around 0.8 to 0.85, values at which PPM gives quite smooth post-shock regions in the worst of cases.

We wish to apply extra dissipation only near shocks. Therefore we multiply S_{01} by a parameter A_{shk} which is unity near shocks and vanishes elsewhere. We wish to ignore shocks with fractional pressure jumps dp_f which are less than 1/3, where

$$dp_f = (p_1 - p_0) / p_0 \quad (69)$$

Here the subscripts 0 and 1 refer to the pre- and post-shock states, respectively. Large pressure gradients can be set up without causing shocks if for example gravitational or centrifugal forces are at work. Therefore we examine instead the size of any compressional velocity jump relative to the sound speed. The shock jump conditions allow a threshold in dp_f to be related to one in the quantity dM^2 defined by

$$dM^2 = (u_1 - u_0)^2 / c_0^2 \quad (70)$$

For a gamma-law gas a threshold dp_{f0} corresponds to a threshold dM_0^2 given by

$$dM_0^2 = 2 (dp_{f0})^2 / \{ \gamma [2\gamma + (\gamma+1) dp_{f0}] \} \quad (71)$$

For each zone we estimate dM^2 by

$$\langle dM^2 \rangle = (\langle u \rangle_{ZR} - \langle u \rangle_{ZL})^2 / \min \{ \langle c^2 \rangle_{ZL}, \langle c^2 \rangle_{ZR} \} \quad (72)$$

Then the parameter A_{shk} is given by

$$A_{shk} = 1 \quad , \quad \text{if } \langle dM^2 \rangle > dM_0^2$$

$$\text{and} \quad \text{if } (\langle u \rangle_{ZL} - \langle u \rangle_{ZR}) > 0 \quad , \quad (73)$$

$$= 0 \quad , \quad \text{otherwise.}$$

A value of 1/3 for dp_{f0} in Eq. (71) is recommended. The shock steepness parameter is then

$$S_{shk} = A_{shk} S_{01} \quad . \quad (74)$$

To achieve sufficient dissipation it is necessary to take for S_{shk} not this value, but the maximum of this value as computed for the zone in question and for its neighbor on the higher pressure side.

Extra dissipation is achieved in the Lagrangian step by flattening the monotized interpolated zone structures of all variables by a factor $(1 - S_{shk})$. That is, the new structure is a blend of $(1 - S_{shk})$ times the monotized interpolated structure and S_{shk} times a totally flat structure. This has the effect of blending the PPM scheme with Godunov's first-order scheme near shocks which are too thin. In the remap step the same flattening procedure is used. For purely Lagrangian calculations best results are obtained if the value of S_{shk} is doubled. Some care must be taken when this technique is generalized to 2-D computations. One must be sure to sense the presence of shocks which compress mainly in the y-direction. Thus differences of the transverse component of velocity, u_y , must figure in the definition of $\langle dM^2 \rangle$ in Eq. (72). Also, a numerical approximation of the velocity divergence should replace the velocity difference in Eq. (73). Basically, the flattening factor for a zone should take the larger of the values which would be obtained for the zone in a 1-D x-sweep and a 1-D y-sweep. This will give much needed dissipation along the direction of a shock nearly aligned with the mesh in 2-D. An example of the consequences of omitting such dissipation is given in Fig. 8 of Woodward and Colella (1983).

MULTIFLUID CALCULATIONS WITH PPM

The Multifluid Lagrangian Step

When Eulerian hydrodynamics calculations are performed in a sequence of 1-D passes each of which consists of a Lagrangian step followed by a remap, an especially simple and powerful approach to multifluid hydrodynamics is easily employed. This approach was first used by DeBar in his KRAKEN code in the late 1960's and it is described in DeBar (1974). The idea behind DeBar's technique is to introduce a variable f_j for each fluid j in the problem. This variable is the fraction of the zone volume occupied by fluid j . In the Lagrangian step of the calculation the fluids are treated as well mixed within the zones, while in the remap step interfaces between the fluids are constructed within each zone. An essential feature of this method is its use of only the fractional volumes in the surrounding zones to construct the configuration for the fluids within each zone. This use of only local data makes the method both simple and robust. In particular, the method does not break down when a simply-connected region of fluid becomes multiply connected. Thus the method can follow complicated flows such as the breaking of water waves.

The Lagrangian step of a multifluid calculation in which the fluids are treated as well mixed can be performed in a number of ways. In the original KRAKEN code, the individual equations of state of the fluids were used to compute partial pressures which were summed to give the pressure of the zone. This pressure was used to accelerate the zones, and all fluids in the zone were assigned the same velocity. In KRAKEN an equation for the internal rather than the total energy was solved. This equation was solved separately for each fluid under the assumption that all fluids experienced the same fractional change in volume. Because of the differing compressibilities of the fluids, this last assumption led to inaccuracies in certain situations, and means of equilibrating the partial pressures within multifluid zones were introduced in other, later codes employing the same basic approach as KRAKEN.

We now describe a possible way of performing PPM Lagrangian hydrodynamics for a zone containing two fluids. If we assume that two fluids in a zone are always well mixed and in pressure equilibrium, then for small changes in zone volume this mixed fluid obeys an equation of state which can be derived from those of the individual fluids. If we call the two fluids A and B, with fractional volumes f_A and $f_B = 1 - f_A$, then pressure equilibrium implies that

$$p_A = p_B \quad , \quad (75a)$$

$$dp_A = dp_B \quad . \quad (75b)$$

For small adiabatic changes we must then have

$$c_A^2 dp_A = c_B^2 dp_B \quad . \quad (76a)$$

The constancy of the mass fractions of the fluids implies that for small changes

$$\frac{dp}{\rho} = \frac{dp_A}{\rho_A} + \frac{df_A}{f_A} \quad , \quad (76b)$$

where ρ , the density of the composite fluid, is given by

$$\rho = f_A \rho_A + f_B \rho_B \quad . \quad (77)$$

Differentiating this equation for ρ and using

$$df_B = - df_A \quad , \quad (78a)$$

we obtain

$$dp = \left(f_A + f_B \frac{c_A^2}{c_B^2} \right) dp_A + (\rho_A - \rho_B) df_A \quad . \quad (78b)$$

We now find

$$\frac{1}{f_B} \frac{df_A}{f_A} = \left(\frac{\rho_A c_A^2 - \rho_B c_B^2}{\rho_B c_B^2} \right) \frac{dp_A}{\rho_A} \quad . \quad (78c)$$

The definition of the composite sound speed, c , is

$$c^2 = \left. \frac{\partial p}{\partial \rho} \right|_S, \quad (79)$$

where the derivative is taken at constant entropy. Using the definition of the composite pressure, p ,

$$p = f_A p_A + f_B p_B, \quad (80)$$

we find

$$c^2 = f_A c_A^2 \left. \frac{\partial \rho_A}{\partial p} \right|_S + f_B c_B^2 \left. \frac{\partial \rho_B}{\partial p} \right|_S, \quad (81)$$

which leads to

$$c^2 = \frac{\rho_A \rho_B c_A^2 c_B^2}{[\rho_A \rho_B (f_A^2 c_B^2 + f_B^2 c_A^2) + f_A f_B (\rho_A^2 c_A^2 + \rho_B^2 c_B^2)]}. \quad (82)$$

Equations (80) and (82) allow both p and c^2 to be computed, and our model equation of state discussed earlier may be fitted to this data. Lagrangian PPM hydrodynamics may then be computed as for a single fluid, so long as the fractional volumes are adjusted after the computation to achieve precise pressure balance of the two fluids within the zone.

In practice, the method of treating fluids as well mixed within zones for the purpose of Lagrangian hydrodynamical calculations works very well. The reason for this must be that the main falsification introduced appears in the sound speed rather than in the pressure or velocity, which appear directly in the fluxes of specific volume, momentum, and total energy. The pressure and velocity are not greatly affected by assuming that the fluids are well mixed, because these variables are continuous across a fluid interface. In two dimensions, demanding that both fluids share a common velocity makes the treatment of slip along the fluid interface difficult. Any such slip must be smeared over a thin region about two zones wide.

The Multifluid Remap Step

In the remap step of a multifluid Eulerian calculation interfaces separating the various fluids within a zone must be

drawn. The SLIC method of Roh and Woodward (1977) is a particularly simple and effective means of drawing these interfaces. It is a simplification of the original method of DeBar (1974), and it is especially well behaved when many fluids are present in a single zone. For the special case of only two fluids in a zone SLIC is just about the simplest possible procedure for defining a fluid interface, and it is remarkable how well it works in practice. There are only five possible fluid configurations, and these are illustrated in Fig. 22.

The SLIC algorithm is specifically designed to work in 1-D passes using a minimum of data. The fluid interface within a zone is constructed using only the fractional volumes of the fluids in that zone and a knowledge of the presence or absence of each fluid in each of the two neighboring zones. Representing presence of the fluid by a 1 and absence by a 0, we may characterize this data in neighboring zones by an ordered pair of numbers for each fluid. The pair (1,0), for example, would mean that the fluid was present only in the neighboring zone on the left. SLIC chooses the horizontal fluid configuration on the left in Fig. 22 whenever these ordered pairs of numbers are the same for both fluids in the zone. If we call the fluids A and B, then the vertical configuration with A on the left will be chosen if we have either (1,0) for fluid A or (0,1) for fluid B. The remaining vertical configuration is the same except that the names of the fluids are interchanged. One of the two center-slab configurations on the right in Fig. 22 will be chosen if the ordered pairs for the two fluids are (0,0) and (1,1). These center-slab configurations occur only rarely in practice. They generally arise when there is a tendency to form spray at a badly distorted fluid interface. They will not be properly advected, and the presence of many such zones in a problem is an indication that the calculation should be terminated. Disregarding these zones, SLIC simply chooses to represent the fluid interface within a zone as horizontal or vertical. This gives a first-order accurate description of the fluid interface. This description is extremely robust; it makes it quite difficult for a fluid region to break up or to form thin structures. This

Fig. 22 The five permitted two-fluid configurations in SLIC for the x-pass of an "operator-split" calculation.

robustness is amply demonstrated by the 10-fluid problem presented in Noh and Woodward (1977).

After the lengthy discussion of advection algorithms at the beginning of this article it should seem natural to apply the PPM advection scheme to the treatment of the fluid interface. If this is done, it should be done only in regions where a single, isolated fluid interface is to be described over a large enough number of zone widths to make the high-order interpolations in PPM meaningful. A means for applying PPM to the problem of fluid interface definition is illustrated in Fig. 23. To construct the representation of the fluid interface which appears at the left in Fig. 23, the fractional volumes of the two fluids have been summed over five-zone strips. The sums over vertical strips have been chosen to define the fluid interface, because viewed from this direction the interface appears more horizontal. The figure illustrates how the fluid interface might appear if the PPM interpolation algorithm with monotonicity constraints were used to construct it from data summed over strips. In the central portion of Fig. 23 this form for the fluid interface has been used within each strip to construct interfaces within zones which match the fractional volumes prescribed there. This representation of the interface would be appropriate for the y-pass of the calculation, because the monotonicity constraint would act to damp strongly any short wavelength ripples on the fluid interface. Such damping serves to prevent the creation of spray, which cannot be adequately treated by the method unless it is well resolved on the computational mesh.

For an x-pass of the calculation the representation in the center of Fig. 23 is inappropriate. The portions of the fluid interface which have been flattened by the monotonicity constraint would now cause an undesirable proliferation of multi-fluid zones. In these zones the fluid interfaces are drawn nearly horizontal, while the SLIC algorithm would draw them vertical. A possible means of modifying the fluid interface to deal with this problem is illustrated at the right in Fig. 23. In constructing this part of the figure an additional monotonicity constraint has been applied. In each multifluid zone the usual monotonicity criterion has been applied in the y-direction. That is, the interface has been steepened if necessary to give values of the fractional volume in the x-volume coordinate which lies between those for the zones just above and below the zone in question. Note that this constraint is not applied to the data summed over strips but instead to that for individual zones.

The ideas illustrated in Fig. 23 are provisional and are now under development. They do not provide a tested means of improving the SLIC algorithm, but Fig. 23 should warn the ambitious reader that improving significantly upon SLIC is not as

easy as at first it seems. It also remains to be shown that improvements in drawing the fluid interface make sense in the absence of a better treatment of the hydrodynamics in the Lagrangian step and a better treatment of slip along the interface. A true treatment of fluid slip requires that the calculation not be split into separate x- and y-passes. However, significant improvements can be made by accounting for the shear term, formally third-order small, which was discussed earlier and illustrated in Fig. 5 (J. LeBlanc, private communication).

Fig. 23 A proposed piecewise-parabolic embellishment of the SLIC algorithm for defining a fluid interface. (a) Monotonized PPM interpolation parabolae are generated from fractional volume data summed in five-zone strips. The direction of summation is determined so that the fluid interface appears most nearly horizontal. (b) The interface shapes constructed in (a) are adjusted vertically to correspond to fractional volumes in individual zones. This representation is appropriate for advection in the direction of summation in strips (vertical in this example). (c) For advection in the other direction (horizontal in this example) the individual fluid interfaces are steepened if necessary to satisfy a monotonicity constraint in the direction of the strips (vertical). This will prevent excessive proliferation of multifluid zones. The block of 25 zones shown here would be used only to generate the fluid interface within the central zone. Generation of interfaces in other zones is indicated here to give an idea of the variety of situations which can arise.

SIMULATION OF UNSTABLE SUPERSONIC JETS WITH PPM

As an example of an astrophysical application of multifluid PPM hydrodynamics with SLIC tracking of fluid interfaces, we present here some results from a study of Kelvin-Helmholtz instabilities in supersonic jets. A detailed discussion of this work may be found in Woodward (1983). An earlier example of the use of SLIC imbedded in a different hydrodynamics code may be found in Woodward (1979, 1980). The motivation for the work on jet instability is the need to establish whether or not simple hydrodynamical models are appropriate for the study of jets emitted from the nuclei of active galaxies. A realistic simulation of a hydrodynamical jet must be three-dimensional; however, such calculations are beyond the capabilities of present computers. Consequently, two-dimensional simulations have been performed. In the very detailed work of Norman et al. (1982, 1983, see also this volume), cylindrical symmetry was assumed, and no catastrophic instabilities were observed. To complement this work, I have performed two-dimensional simulations of supersonic fluid sheets in Cartesian geometry. These simulations, unlike those in cylindrical geometry, permit the "jet" to meander. The meandering, or "firehose" instability is potentially more disruptive than those Kelvin-Helmholtz modes allowed in cylindrical geometry. Therefore it is important to see if this instability grows rapidly for popular hydrodynamic models of jets from active galactic nuclei. It is hoped that the behavior of fluid sheets in two dimensions will be a good indicator of that of roughly cylindrical jets in three dimensions.

It is well known that an isolated slip surface between two uniform fluids is stable when the motion of both fluids relative to a disturbance of the slip surface exceeds a critical Mach number slightly greater than unity. A similar result does not hold for a fluid sheet (or jet) because of the possibility of symbiotic interactions between perturbations at either side of the sheet. Such interactions, which result in a violent instability of the fluid sheet, are illustrated in Fig. 24. The calculation in Fig. 24 began with a sheet of gas with density 0.1, pressure 0.6, and x-velocity 6.32 (Mach 2) located between $y = -0.5$ and $y = 0.5$. This gaseous sheet was initially in pressure equilibrium with a surrounding ambient gas of density 1, pressure 0.6, sound speed 1, and x-velocity 0. Both gases obey a gamma-law equation of state with $\gamma = 5/3$. This system was perturbed by introducing a small vertical component to the velocity within the "jet" (we will use this term loosely from now on):

$$u_y = [0.05 \sin(2\pi x/3) + 0.025 \sin(4\pi x/3)] u_{x0}. \quad (83)$$

The subsequent flow, computed with PPM using a SLIC fluid interface, is shown at times 0.5, 1, 1.5, and 2 in Fig. 24. At each time 30 density contours are plotted with each contour level a fixed factor higher than the preceding one. The lowest and highest levels are given in the figure caption, and the lowest contour is dotted. The fluid boundary as well as the shocks in Fig. 24 are made evident by the bunching together of several density contours. Centered rarefaction fans are marked by regions in which several density contours appear to fan out from a common center. The calculation simulates only a small section of the jet, and periodic boundary conditions are applied at $x = 0$ and $x = 3$. The grid is very fine. There are 180×240 square zones in the region $0 < x < 3$ and $-2 < y < 2$ which is shown. Eighty additional zones above and 80 below this region grow steadily larger so that the y-boundaries are at ± 12 . This many additional y-grid lines proved to be unnecessary, because the jet was completely disrupted before sound signals could reach these distant boundaries. To demonstrate the correctness of the general flow pattern in Fig. 24, an identical calculation using precisely half the grid resolution in each direction is presented for comparison in Fig. 25. Note that even the amount of entrainment of ambient gas into the jet is the same in these two calculations. Only the fine details of the flow near the jet boundaries require the finer grid for a more accurate description.

Although the initial disturbance given in Eq. (83) has no y-dependence within the jet, the disturbance at time 0.5, in Fig. 24a has quite a different character. It tends to be similar along Mach lines rather than along vertical lines. Mach lines are the sound wave fronts which are generated in supersonic flow by signals emitted at specific points. The oblique features

within the jet in Fig. 24a are weak shocks aligned at the Mach angle in the frame of reference moving with the small kinks in the jet boundaries which generate these disturbances. The lack of y-dependence in the initial state was inconsistent with causality, and the flow has rearranged itself accordingly in short order. By time 1, in Fig. 24b, a characteristic and rather special structure has emerged. Nearly all the various weak shocks inside the jet have merged into two principal shocks of greater strength. These shocks, which form a zig-zag pattern inside the jet, are the means by which one jet boundary influences and destabilizes the other.

The flow pattern in Fig. 24b is fundamental to an understanding of supersonic jet instability. Beginning at the left in Fig. 24b, a shock is generated as the supersonic jet gas strikes a kink in the upper wall of the channel. This shock then propagates at roughly the Mach angle to the opposite wall of the channel. Here the shock strikes the wall and generates a dent in

Figure Captions

- Fig. 24 The fundamental odd mode of instability of a Mach 2 jet. The jet density is initially 0.1 and the ambient gas has a density of 1. Both are gamma-law gases with $\gamma = 5/3$. They are initially in pressure equilibrium, and the ambient sound speed is unity. Thirty density contours, each larger by a fixed factor, are plotted at times 0.5, 1, 1.5, and 2. A perturbation of the transverse velocity of the jet causes a characteristic zig-zag pattern of shocks within the jet to develop. Eventually some of the jet material is reversed in direction and ambient gas is entrained in the jet. By time 2.5 the jet is completely disrupted. The section of the grid which is shown is uniform, with 60 zones per jet width. The calculation was performed with PPM using a SLIC fluid interface.
- Fig. 25 Results of a calculation identical to that in Fig. 24 except for the use of precisely half the grid resolution in both x- and y-directions and half the number of timesteps. The general morphology of the flow in this calculation -- the amplitude of the jet oscillation, the locations of shock fronts, and the amount of entrainment of ambient gas into the jet -- agrees very well with the results for the finer grid in Fig. 24. Only fine details of the flow near the jet boundaries are not well represented here.

Fig. 24

AHQ, AHQ t=5.02e-01, 1.00e+00: PPMLR 9/27/82- 1 mach2eac, mach2eag
n= 292, 586; 30 contours: 3.87e-02 to 1.54e+00, 3.05e-02 to 2.14e+00

AHQ, AHQ t=1.50e+00, 2.00e+00: PPMLR 9/27/82- 1 mach2eak, mach2eao
n= 894, 1210; 30 contours: 2.75e-02 to 2.07e+00, 1.44e-02 to 2.20e+00

Fig. 25

RHO, RHO t=5.07e-01, 1.00e+00; PPMLR 9/26/82- 2 mach2dab, mach2dad
n= 154, 298; 30 contours: 3.78e-02 to 1.46e+00, 3.54e-02 to 1.91e+00

RHO, RHO t=1.50e+00, 2.01e+00; PPMLR 9/26/82- 2 mach2daf, mach2dah
n= 450, 600; 30 contours: 2.78e-02 to 1.93e+00, 1.50e-02 to 1.73e+00

it. The obliquity of the shock is such that the jet gas is turned downward here into the channel wall. This turning is aided by a small centered rarefaction fan which returns the post-shock pressure to the value in the adjacent ambient gas. The jet gas which has been turned into a collision course with the channel wall again causes a shock to form. When this shock reaches the opposite wall, the cycle is completed. The kink on each wall generates a shock which turns the flow into the kink on the opposite wall. When this turned gas strikes the wall, the pressure generated both sustains a shock and further excavates the kink in the wall. The effect of this interaction between the perturbations on the two channel walls is a rapid growth of the amplitude of the wobbling of the jet and of the strength of the imbedded shocks. By time 1.5, in Fig. 24c, some of the jet gas strikes the wall head on and is turned around to flow backwards along the jet boundary. In Fig. 24d, at time 2, the two vortices associated with this reverse flow are very large. Also the shocks inside the jet are nearly orthogonal to the flow. They therefore decelerate and heat the gas considerably. The propagation of the kinks in the jet is supersonic with respect to the ambient gas, and by time 1.5 fairly strong shocks are generated in that gas. In addition, chunks of ambient gas are entrained in the jet so that by time 2.5 the jet is completely disrupted.

Because the instability of the jet in Fig. 24 is driven by a reinforcement of the perturbation on one boundary of the jet by the perturbation on the other, it is natural to expect resonant effects. In fact, if the two perturbations are to reinforce each other the shocks set up on one side of the jet must strike the other side in proper phase with the perturbation there. Because the shocks must be inclined at nearly the Mach angle, there must therefore be a resonance condition involving the Mach number of the jet, the jet width, and the wavelength of the perturbation. These parameters are near resonance for the jet in Fig. 24. This particular resonant mode corresponds to a single zig-zag of shocks within the jet which are inclined at roughly the Mach angle. This is the fundamental odd mode, the resonant mode which has the longest wavelength, the fastest growth rate, and the most destructive effect upon the jet. It is a wobbling mode with alternating symmetrical regions of positive and negative velocity in the y -direction.

A whole series of modes can be constructed by drawing additional shock zig-zags spaced evenly in phase. Odd numbers of zig-zags give odd, wobbling modes, while even numbers of zig-zags give even, pinching modes. Even modes have even symmetry for the density distribution about the jet center line, and they are therefore the modes which appear when this line is treated as a reflecting boundary (as is the case when cylindrical symmetry is imposed). The fundamental even mode is shown in Fig. 26. The

initial, unperturbed jet for this calculation was identical to that used in the run in Fig. 24. However, in this case reflection symmetry about the jet center line at $y = 0$ was assumed. Also a sinusoidal perturbation of u_y within the jet was introduced with only half the wavelength of the fundamental odd mode:

$$u_y = 0.1 \quad y \quad u_{x0} \quad \sin(4\pi x/3) \quad . \quad (84)$$

The double zig-zag pattern which is characteristic of the fundamental even mode of jet oscillation is modified in Fig. 26 by the Mach reflections of the shocks at the jet center. For higher Mach numbers these become regular shock reflections, and a simple double zig-zag pattern of shocks within the jet appears. The grid used for the calculation in Fig. 26 is matched to that in Fig. 24. There are 90×120 square zones in the region $0 < x < 1.5$ and $0 < y < 2$, and 40 additional grid lines of increasing separation carry the computational domain out to $y = 5$. The flow is shown in Fig. 26 at time intervals of 0.5 beginning at $t = 0.5$. This mode rapidly causes breaking waves to form along the jet boundary and by time 2 a considerable amount of ambient gas has been entrained in the jet; however, this mode is still less disruptive than the fundamental odd mode in Fig. 24.

As can be seen at the earlier times displayed in Fig. 26, it is possible to draw a tighter shock zig-zig pattern within the jet at the Mach angle. Thus the fundamental even mode could be generated at higher frequencies; however, at higher frequencies its growth is less rapid and its effects are less disruptive of the jet. This is demonstrated in Fig. 27a. In this calculation the same even perturbation of the jet as in Eq. (84) but with half the wavelength was introduced. The flow is shown at time 1.5 in Fig. 27a, and the characteristic double zig-zig pattern of internal shocks is evident. During the evolution of this run the fundamental even mode shown in Fig. 27a competes with its first harmonic. As the instability grows and the effective jet width is reduced slightly by the entrainment of ambient gas, it becomes possible to fit the double zig-zag within the imposed wavelength of $3/4$ of the jet width. Then this fundamental even mode grows in strength to dominate the flow. To obtain the second even mode, with a quadruple zig-zag pattern of internal shocks, we must decrease the wavelength of the perturbation further. In the calculation shown at time 1.5 in Fig. 27b, a wavelength of 0.5 has been imposed:

$$u_y = 0.1 \quad y \quad u_{x0} \quad \sin(4\pi x) \quad . \quad (85)$$

Fig. 26

RHO, RHO $t=5.00e-01$, $1.00e+00$: PPMLR 10/14/83- 1 mach2z00, mach2z00
n= 314, 622: 30 contours: $3.83e-02$ to $1.78e+00$, $3.48e-02$ to $4.28e+00$

RHO, RHO $t=1.50e+00$, $2.00e+00$: PPMLR 10/14/83- 1 mach2z0c, mach2z0ad
n= 940, 1244: 30 contours: $2.77e-02$ to $4.64e+00$, $4.24e-02$ to $3.11e+00$

Figure Captions

Fig. 26 The fundamental even mode of instability of the Mach 2 jet in Fig. 24. This PPM computation uses a uniform grid within the region shown, with 60 zones per jet width as in Fig. 24. Density contours are shown at times 0.5, 1, 1.5, and 2. The growth of this resonant pinching mode is very rapid. The characteristic x-pattern of internal shocks appears by time 0.5, and the wave disturbances at the jet boundary break by time 1. Nevertheless, this mode is less disruptive than the "firehose" mode shown in Fig. 24. Note the Mach reflection of the internal shocks at the jet center line.

RHO. RHO 1.150e+00. 1.50e+00: PPNLR 10/14/83- 1 mach2xal. mach2val
n 1820. 1768: 30 contours: 2.43e-02 to 2.20e+00. 5.53e-02 to 2.76e+00

Fig. 27 The effects of short wavelength perturbations of even symmetry about the jet center line are indicated by these two PPM simulations. The unperturbed jets here are the same as in Fig. 24. In both simulations the grids are uniform in the region shown, with 120 zones per jet width, and both flows are shown at time 1.5. At the left the perturbation wavelength is 0.75 and at the right it is 0.5 jet widths. At the left a competition between the fundamental even mode of Fig. 26 and its first harmonic is eventually won by the former mode; at the right, the latter, harmonic mode is dominant.

The second even mode now appears clearly, but despite its rapid ruffling of the jet boundary it is less disruptive than the fundamental even mode shown in Fig. 27a. The relatively minor effect of this mode on the flow near the center of the jet can be gauged by the weakness of the shocks which it sets up there.

In Fig. 28 the second odd mode of jet instability can be seen. The calculation in Fig. 28 is identical to that in Fig. 24 except in two respects. First, the grid is somewhat coarser. There are 120x160 square zones in the region $0 < x < 3$ and $-2 < y < 2$. Twenty y-grid lines of increasing separation on either side of this region carry the computational domain out to $y = \pm 4$. The second difference between this calculation and that in Fig. 24 is of course the initial perturbation:

$$u_{y0} = 0.05 [\sin(2\pi x) + \sin(8\pi x/3)] u_{x0} \quad (86)$$

Initially the third resonant mode represented by the first term in u_{y0} grows to produce at time 1.25 the shock lattice structure shown in Fig. 28a. However, nonlinear coupling of this mode with the fourth one (the second term in u_{y0}) soon causes one of the zig-zag shocks to overwhelm the others, and the fundamental odd mode emerges at time 2.25 in Fig. 28b. By time 3 entrainment of ambient gas in the jet disrupts it considerably.

The experiment in mode-mode coupling represented in Fig. 28 demonstrates the dominance of the fundamental odd mode. In this case it has arisen as the difference frequency oscillation corresponding to the two shorter wavelength initial perturbations. That this behavior is not just a general tendency for ever longer wavelength oscillations to grow is demonstrated by two further experiments. The flow in Fig. 29 began with the same initial conditions as that in Fig. 28, except that the initial perturbations had longer wavelengths by a factor of 4:

$$u_{y0} = 0.05 [\sin(\pi x/2) + \sin(2\pi x/3)] u_{x0} \quad (87)$$

The grid has 360x120 square zones in the region $0 < x < 12$ and $-2 < y < 2$, with 10 additional y-grid lines of increasing separation expanding the computational domain to $y = \pm 2.5$. The flow is displayed at time 2, and the mode in Fig. 24 is clearly dominant. There is barely any hint of power at the difference frequency, corresponding to a wavelength of 12. In Fig. 30 a similar run is shown at time 3. Here the initial perturbation contained a wavelength of 12:

$$u_{y0} = [0.06 \sin(\pi x/6) + 0.02 \sin(\pi x/2)] u_{x0} \quad (88)$$

A somewhat coarser grid in x of 240 evenly spaced x-lines was used, while a somewhat finer grid in y of 160 evenly spaced y-lines was used from y = -2 to y = 2. On either side of this domain forty additional y-lines of increasing separation carry the grid to y = ±6. In the flow at time 3 the long wavelength initial perturbation has only modulated a pattern dominated by the shorter wavelength.

These simulations suggest that if a meandering supersonic jet has a principal wavelength discernable, this wavelength can be related to the other parameters of the jet by assuming that the

Fig. 28 The second resonant odd mode of instability of a Mach 2 jet. The same jet as in Fig. 24 has here been perturbed with disturbances with 3 and 4 wavelengths over the length of the grid. At first, the longer wavelength sets up an unstable resonant oscillation characterized by the triple zig-zag of shocks within the jet which is seen at the left (time 1.25). Later (time 2.25) at the right, nonlinear coupling between the initial disturbances has excited the dominant, fundamental odd mode.

fundamental odd resonant mode has been observed. This mode is so violently unstable that if it is not seen, the jet is unlikely to be a simple hydrodynamic jet of the type studied here. Instead, it may contain a stabilizing internal magnetic field. Alternatively it may be stabilized by a strong pressure gradient in the external ambient gas which gives the jet a strongly preferred direction of propagation, and resists meandering (rivers flowing down hills do not meander). Another possibility is that the jet Mach number is so large (of the order of 100) that it gets where it's going before the instability can develop. Remember that the resonant wavelength scales linearly with the jet Mach number when this is large (because of its relation to the Mach angle).

In the supersonic exhausts of jet engines the fundamental even mode is usually dominant for some time, but the fundamental odd mode, a helical mode for a cylindrical jet, eventually takes over (cf. Yu and Steiner 1983). The reason for the initial dominance of the even mode is an enormous even perturbation of the jet caused by its sudden emergence into a pressure reservoir to which it must adjust by either a lateral expansion or contraction. Norman et al. (1982) have suggested that hot spots of synchrotron radiation emitted from astrophysical jets may be due to

Fig. 29 The same initial jet as in Fig. 24 here has been perturbed by 5% sinusoidal oscillations of the transverse velocity with wavelengths of 3 and 4 jet widths. Because the shorter wavelength initial perturbation is near the fundamental resonant mode shown in Fig. 24, this mode grows so rapidly that it disrupts the jet before nonlinear mode-mode coupling can introduce any significant oscillation with a wavelength of 12 jet widths. Density contours are shown here at time 2, shortly before the jet is disrupted.

internal shocks of the type they observe in their calculations. If such hot spots correspond to the even modes in their calculations, perhaps the astrophysical jets, like their terrestrial counterparts, have suddenly emerged from a galaxy into an intergalactic pressure reservoir to which they adjust by exciting strong even modes of oscillation.

Fig. 30 The same initial jet as in Fig. 24 here has been perturbed by 6% and 2% sinusoidal oscillations in the transverse velocity with wavelengths of 12 and 4 jet widths. Despite its smaller initial amplitude, the shorter wavelength perturbation soon dominates the flow. Density contours are shown here at times 2 and 3. By time 3 the jet is nearly disrupted by a short wavelength oscillation. Although it is not immediately apparent, nonlinear effects have generated a large amount of the fundamental resonant mode with a wavelength of 3 jet widths.

The PPM hydrodynamics scheme with SLIC fluid interfaces is a powerful tool for investigating fluid flows. As the above investigation of jet instability demonstrates, nonlinear flow phenomena are easily and accurately computed which can then be understood in relatively simple terms. Although only one case may be computed at a time, the calculations are cheap enough that many can be afforded. This allows similarities and trends to be identified which may be used to predict the outcome of further simulations in order to test the correctness of their interpretation. From this process a phenomenological understanding emerges, which can be made quantitative as desired by performing additional simulations. This procedure allows theoretical work to be extended into regimes which are intractable for analytic methods.

The contrast between this numerical approach and analytic techniques is illustrated by the problem of jet instability. The numerical work sketched above and discussed more fully in Woodward (1983) should be compared with the analytic work of Ferrari et al. (1978) and of Hardee (1982, 1983). In that work long formulae involving special functions and graphs of growth rates are presented, but the role of resonant effects mediated by shocks at the Mach angle in defining the dominant mode in a given situation is obscured. Undoubtedly this is largely because such shocks, even when they are weak, are inherently nonlinear phenomena. Another reason that the Mach angle does not emerge from the linear analysis is that the modes studied for supersonic jets are not constrained to be causally meaningful. Flows which may actually be realized in meaningful experiments must be obtained by a superposition of a great many of the modes studied analytically. Thus although one may conclude from the analytic work that a supersonic jet is unstable, he may not learn what that jet will actually do in nature.

The numerical and analytic approaches can reinforce one another. The phenomenological understanding which emerges from a series of numerical simulations can be used to guide analytic work in profitable directions. As every student knows, it is helpful to know the answer before beginning the analysis of a problem. Numerical simulations, like physical experiments, make this possible. Consider the example of supersonic jet instability. In principle, we could find suitable trial perturbations for an analytic treatment by the following method. We could idealize the results of the simulations at the point where a principal mode has been established but where the shocks involved are still weak. This idealization would serve to remove unavoidable small numerical noise from the computed result and also to yield a perturbation simple enough to treat analytically. If growth rates for such perturbations could be computed analytically as functions of the jet parameters, the results would be most useful. As the power of computing machinery and

techniques continues to grow we may hope to see an accompanying growth of such interplay of the numerical and analytical approaches to theoretical physics.

Summary

Piecewise-parabolic methods for advection and for hydrodynamics have been discussed thoroughly here outside of the context of applied mathematics from which they have grown. This is not because other methods are without merit, but instead it is the result of the author's personal preference. Nevertheless, many of the fundamental issues discussed -- construction of interpolation functions, monotonicity constraints, contact discontinuity steepeners, conservation form, operator splitting, multifluid techniques, etc. -- are more universal than the limited context in which they have been discussed here. The author hopes that readers who do not intend to use piecewise-parabolic methods may still find much to interest them in these pages.

This work was performed under the auspices of the U. S. Department of Energy by the Lawrence Livermore National Laboratory under contract No. W-7405-ENG-48.

REFERENCES

- Berger, M.: 1982, "Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations," Thesis, Stanford University.
- Boris, J. P., and Book, D. L.: 1973, J. Comp. Phys. 40, pp. 202.
- Colella, P.: 1983, "Approximate Solution of the Riemann Problem for Real Gases," Lawrence Berkeley Laboratory Report LBL-14442, to appear in J. Comp. Phys.
- Colella, P., and Woodward, P. R.: 1983, "The Piecewise-Parabolic Method (PPM) for Gas Dynamical Simulations," J. Comp. Phys., in press.
- Courant, R., and Friedrichs, K. O.: 1948, Supersonic Flow and Shock Waves, New York: Interscience.
- DeBar, R.: 1974, "Fundamentals of the KRAKEN code," Lawrence Livermore National Laboratory Report UCIR-760.
- Ferrari, A., Trussoni, E., and Zaninetti, L.: 1978, Astron. Astrophys. 64, pp. 43.
- Godunov, S. K.: 1959, Nat. Sb. 47, pp. 271.

- Gropp, W. D.: 1980, SIAM J. Sci. Stat. Computing 1, pp. 191.
- Hardee, P. E.: 1982, Ap. J. 257, pp. 509.
- Hardee, P. E.: 1983, Ap. J. 269, pp. 94.
- Noh, W. F., Gee, M., and Kramer, G.: 1979, Lawrence Livermore National Laboratory Report UCID-18515.
- Noh, W. F., and Woodward, P. R.: 1977, Lecture Notes in Physics 59, pp. 330.
- Norman, M. L., Smarr, L., and Winkler, K.-H. A.: 1982, Astron. Astrophys. 113, pp. 285.
- Norman, M. L., Winkler, K.-H. A., and Smarr, L.: 1983, to appear in the proceedings of the conference on astrophysical jets, Torino, Italy, October, 1982.
- Strang, G.: 1968, SIAM J. Num. Anal 5, pp. 506.
- van Leer, B.: 1977, J. Comp. Phys. 23, pp. 276.
- van Leer, B.: 1979, J. Comp. Phys. 32, pp. 101.
- Woodward, P. R.: 1979, in The Large-Scale Characteristics of the Galaxy, ed. W. B. Burton, I.A.U. Symp. No. 84, pp. 159.
- Woodward, P. R.: 1980, in Early Solar System Processes and the Present Solar System, ed. D. Lal, Proc. Int. Sch. Phys. "Enrico Fermi," Course 73, pp. 1.
- Woodward, P. R.: 1982, in Parallel Computations, ed. G. Rodrigue, Academic Press, pp. 153.
- Woodward, P. R.: 1983, "Kelvin-Helmholtz Instability of Pressure-Confined, Supersonic Jets," Lawrence Livermore National Laboratory report, in preparation.
- Woodward, P. R., and Colella, P.: 1981, Lecture Notes in Physics 141, pp. 434.
- Woodward, P. R., and Colella, P.: 1983, J. Comp. Phys., in press.
- Woodward, P. R., and White, R. L.: 1983, "The Piecewise-Parabolic Boltzmann Scheme (PPB)," Lawrence Livermore National Laboratory report, in preparation.
- Yu, J. C., and Steiner, J. M.: 1983, Proc. A.I.A.A. Eighth Aero-acoustics Conf., Atlanta, April, 1983.

DISCLAIMER

This document was prepared as an account of work sponsored jointly by the U.S. Department of Energy and the Defense Advanced Research Projects Agency. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government thereof, and shall not be used for advertising or product endorsement purposes.