

The PPM Compressible Gas Dynamics Scheme

Paul R. Woodward

University of Minnesota, LCSE

Feb. 19, 2005

1. Introduction.

The development of the PPM gas dynamics scheme grew out of earlier work in the mid 1970s with Bram van Leer on the MUSCL scheme [1-3]. The work of Godunov [4] inspired essential aspects of MUSCL. Features that I introduced into the MUSCL code – directional operator splitting and a Lagrangian step followed by a remap – were inspired by my work with the BBC code [6] at Livermore. These features of BBC were derived in turn from the work in the 1960s by DeBar on the KRAKEN code [5]. At Livermore, in the late 1970s, the use of moments of the internal cell distributions, which makes MUSCL resemble finite element schemes, was abandoned for reasons of compatibility with Livermore production codes. In order to recapture the accuracy of MUSCL while using only cell averages as the fundamental data, PPM, the Piecewise-Parabolic Method, was developed in collaboration with Phil Colella [7-10]. Over the last 20 years, PPM has evolved considerably in order to address shortcomings of the scheme as it was laid out in [9] and [10]. A complete description of the scheme has not appeared in print, but descriptions on the LCSE Web site associated with the PPMLib subroutine library exist [24]. PPM has also been extended to MHD [11-16] in collaboration with Wenlong Dai, applied extensively to the study of jets and supersonic shear layers [10,17-22] in collaboration with Karl-Heinz Winkler, Steve Hodson, Norm Zabusky, Jeffrey Pedelty, and Gene Bassett, used to simulate flow around stationary and moving objects [23] in collaboration with B. Kevin Edgar, applied extensively to convection and turbulence problems [25-46] in collaboration with David Porter, Annick Pouquet, Igor Sytine, and the Livermore ASCI turbulence team, and most recently extended to multifluid gas dynamics problems. Implicit versions of the method have also been worked on [47-52] in collaboration with Bruce Fryxell and, later, Wenlong Dai. A cell-based AMR version of PPM has been under development [see 53] in collaboration with Dennis Dingo. A complete description of the PPM algorithm at the time of this writing can also be found on the LCSE Web site [54]. Versions of the PPM gas dynamics scheme have become incorporated into 3 community codes aimed at astrophysics applications: FLASH [refs], ENZO [refs], and VH-1 [refs].

Because the several improvements and modifications to the 1984 or 1986 versions of PPM in the literature have not been published, we here lay out the present scheme for simple, single-fluid, ideal gas dynamics in some detail for those aspects that most directly affect simulation of turbulent flows. The aspects of the scheme that allow it to accurately capture and follow shocks will not be described here, for lack of space. For these aspects of the PPM algorithm, which are rather complex, the reader is referred to the Web documents [refs], since this treatment has changed since the scheme was last described completely in the literature.

2. Design Constraints.

Before launching into a systematic description of the PPM algorithm, it is worthwhile to first explain the goals and constraints that have influenced its design. These are:

1. Directional operator splitting.
2. Robustness for problems involving very strong shocks.
3. Contact discontinuity steepening.
4. Fundamental data in the form of cell averages only.
5. Minimal dissipation.
6. Numerical errors nevertheless dominated by dissipation, as opposed to dispersion.
7. Preservation of signals, if possible, even if their shapes are modified, so long as they travel at roughly the right speeds.
8. Minimal degradation of accuracy as the Courant number decreases toward 0.

The first of these design constraints was to guarantee very high processing speeds on all CPU designs. It has the side benefit of allowing the PPM scheme to employ complicated techniques to achieve high accuracy that would be impractical in a truly multi-dimensional implementation, especially in 3D. Accuracy is still limited in principle by the use of 1-D sweeps, but in practical problems errors produced by the use of 1-D sweeps have rarely, if ever, appeared to be important. The demand that the scheme be robust in the presence of very strong shocks has produced a general orientation that dissipation, in the right places, times, and amounts is good rather than bad. Early experience with Glimm's random choice scheme for multidimensional compressible gas dynamics in the 1970s illustrated the error of letting the formal dissipation of the scheme vanish in smooth flow when the simulated flow contains shocks. In low-speed turbulent flows, the need for dissipation in the numerical scheme is less obvious but no less real. Nevertheless, a design goal has always been to minimize the dissipation consistent with an accurate solution. The treatment of contact discontinuities involves a steepening algorithm that has an anti-diffusion effect. This contact discontinuity steepening is included in PPM in order to allow multifluid problems to be treated simply through the addition of passively advected constitutive properties, such as the constants of an analytic representation of the equation of state. This steepening of contact discontinuities and the action of monotonicity constraints together serve to preserve signals that can be moved across the grid accurately, even though this sometimes requires a falsification of the signal shape. Such preserved but falsified signals travel at the fluid velocity, since very short wavelength sound waves, with the exception of shocks, are difficult to propagate at the right speeds. In PPM we take the view that signals that cannot be propagated at the right speeds are best dissipated. This view has led to an intolerance in PPM for sound signals with very short wavelengths, such as, for example, 4 cell widths. The final design goal on the above list, that the accuracy of the scheme should be maintained in the limit of very many very small time steps, is achieved by increasing the formal order of the interpolation discussed below as the edge of a grid cell is approached.

The attitude that has been taken in the PPM design toward numerical dissipation is the most relevant of the above goals and constraints for the use of PPM in simulating turbulent flows. We will therefore describe in this paper the two portions of the PPM algorithm that most directly address numerical dissipation – the interpolation scheme used in PPM and the manner in which it is applied to constructing a picture of the structure of the flow within a grid cell. The remainder of the scheme that will not be described here – the approximate Riemann solver used in PPM, the construction of fluxes of conserved quantities at grid cell interfaces in Eulerian coordinates, the application of the conservation laws, and the application of a smart diffusion at shocks, and only at shocks – has much less impact on PPM’s numerical treatment of turbulence. So far as these other aspects of the scheme are concerned, it should suffice to point out that shocks are thin and obey the conservation laws that determine both their jumps and their propagation speeds and also to point out that these conservation laws assure that the numerical dissipation of turbulent kinetic energy by the PPM scheme results in the appropriate amount of heat being deposited locally, where the dissipation occurs.

3. PPM Interpolation.

The heart of any numerical scheme is its approach to interpolation. PPM utilizes the cell averages of various quantities in a fairly broad stencil surrounding a cell in 3D in order to construct a picture of the internal structure of the cell. Because the equations of hydrodynamics are coupled, the PPM interpolation of the relevant quantities is also coupled. However, to see how the interpolation process works, it is easiest to consider it first for the simple case of a single variable, which we will represent by the symbol a . We will represent the cell averages of a by $\langle a \rangle_i$. We write the value of a at the left- and right-hand interfaces of the cells as $a_{L,i}$ and $a_{R,i}$. We will determine the coefficients of a parabola $a(\tilde{x}) = a_0 + a_1\tilde{x} + a_2\tilde{x}^2$ representing the distribution of the variable a within the cell in terms of a cell-centered local coordinate $\tilde{x} = (x - x_M)/\Delta x$. Here x_M is the cell center and Δx is the cell width. For simplicity, we first assume a uniform grid. (In modern AMR codes, interpolation on non-uniform grids is unnecessary.) Our first task is to determine whether or not the function $a(x)$ is smooth in the region near our grid cell. To do this, we will essentially compare the first and third derivatives of the function in our cell. This process can be formulated in several different ways, all of which boil down to essentially the same criterion. We begin by determining the unique parabola that has the prescribed cell averages in our cell of interest and its two nearest neighbors.

If we write

$$\Delta a_{L,i} = \langle a \rangle_i - \langle a \rangle_{i-1}$$

then we have

$$a_1 = (\Delta a_L + \Delta a_R)/2; \quad a_2 = (\Delta a_R - \Delta a_L)/2$$

Here we have dropped the subscript i by writing Δa_R for $\Delta a_{L,i+1}$. Where possible, we will try to eliminate confusing subscripts by such devices. However, all our subscript indices will refer to cells, and

none to interfaces (no half-indices will be used). If no subscript index is given, the subscript i for our cell of interest will always be implied. In PPM, interpolated variables are discontinuous at interfaces, which makes the use of index subscripts such as $i + 1/2$ ambiguous. Thus, $a_{R,i}$, or simply a_R , is not equivalent to $a_{L,i+1}$.

We will consider a function to be smooth near our grid cell if the above unique parabola, when extrapolated to the next nearest neighbor cells, gives a reasonable approximation to the behavior of the function there. This will not be true, of course, if the third derivative of the function is sufficiently large. We will consider the function to be smooth in our cell of interest if the parabola in its neighbor on the left, when extrapolated into our cell's neighbor on the right gives an accurate estimate of that cell's average, and if also that cell's parabola gives a good estimate of the average in our cell's neighbor on the left. Writing the average of the extrapolated parabola in the neighbor cell as $\langle a_i \rangle_{i\pm 1}$ and in the next nearest neighbor cell as $\langle a_i \rangle_{i\pm 2}$, we find that

$$\langle a_{i-1} \rangle_{i+1} - \langle a \rangle_{i+1} = 2(a_{2,i-1} - a_2); \quad \langle a \rangle_{i-1} - \langle a_{i+1} \rangle_{i-1} = 2(a_2 - a_{2,i+1})$$

Here we have used the fact that, by construction of these parabolae, $\langle a_i \rangle_{i\pm 1} = \langle a \rangle_{i\pm 1}$. We denote the fractional error in this extrapolation by f_{err} , given by

$$f_{err} = \max \left\{ \left| (a_{2,i+1} - a_2) / (\Delta a_R + \text{sign}(\alpha, \Delta a_R)) \right|, \left| (a_2 - a_{2,i-1}) / (\Delta a_L + \text{sign}(\alpha, \Delta a_L)) \right| \right\}$$

Here sign is the Fortran sign transfer function, which applies the sign of its second argument to the absolute value of its first argument. We also denote a trivial value of the quantity a by α . Thus the term with the sign function is simply used to protect the divide operation in a Fortran program. Of course, we ignore this error estimate when the slope of the function is very small, since in that case it is of no consequence whether we get a “good” extrapolation or not. To get some perspective on this measure, f_{err} , of the smoothness of the function a on our grid, it is instructive to evaluate f_{err} for sine waves of wavelengths $n \Delta x$. For the case where $a(x) = \sin(2\pi x / (n \Delta x))$, we have $\langle a \rangle = s a_0$, where a_0 is the value at the center of the cell and $s = (2 / \Delta x) \sin(\Delta x / 2)$. The higher-order coefficients of the interpolation parabola are given by $a_1 = s \sin(\Delta x) \cos(2\pi x_M / (n \Delta x))$, where x_M is the coordinate of the cell center, and by $a_2 = s (\cos(\Delta x) - 1) a_0$. We therefore find that for values of n ranging from 4 to 15, our fractional error, equal simply to $1 - \cos(\Delta x)$, assumes the values: 1.00, 0.69, 0.50, 0.38, 0.29, 0.23, 0.19, 0.16, 0.13, 0.11, 0.10, and 0.09. If we assert that a sine wave over 14 cells is smooth and one over 10 cells is not, then we may conclude that it is reasonable to assert that a function is smooth if $f_{err} \leq 0.10$ and it is not smooth if $f_{err} \geq 0.20$. In between these two values, we may treat the function as partially smooth, by taking a linear combination of a smooth and an unsmooth interpolation parabola. The choice of which functions we will treat as smooth and which not comes from experience. Using only cell averages, it takes five cell values to determine, for example, whether one is really in a shock or not or, as we have just seen, whether the function is locally smooth or not. Under such conditions, it is simply unreasonable to believe that one can properly treat a sine wave with a wavelength of only 6 cells. There

are those in the community who believe this can be done, but PPM assumes this is impossible without further independent information provided upon which the numerical scheme can operate. Experience also shows that numerical noise, which can originate from a number of sources, tends to show up principally at wavelengths between 4 and 8 cells. Signals that appear at these wavelengths are thus quite possibly noise, and therefore deserve to be treated as if they might actually be noise. PPM takes the view that a little noise is not bad, but too much is intolerable. This view is motivated by much experience indicating that nearly total elimination of noise is usually accompanied by partial collateral elimination of the signal.

In PPM, we interpolate several different functions. Some of them, like the Riemann invariants, are defined only as differences (the Riemann invariants are defined by inexact differentials). Other functions, like the transverse velocities, are defined by cell averages. Some may have sharp transitions, associated with contact discontinuities, that deserve special treatment. Below, we describe the most general algorithm, for the special case of a uniform grid. This is the algorithm that PPM uses to interpolate a sub-grid structure for the Riemann invariant associated with the entropy, for which the differential form is given by

$$dA = d\rho - dp/c^2$$

From this most complicated interpolation algorithm, all other, simpler forms used in PPM can be derived as special cases. We will therefore describe the interpolation for the entropy Riemann invariant in detail, and later briefly note in which respects the other interpolations are degenerate cases of this one.

We begin the interpolation of dA by defining a variable A derived from it by determining the arbitrary constant of integration such that the cell average of A vanishes. Thus each cell has its own local scaling for A in which the cell average vanishes but the derivatives are correct. This local scaling is much like the local, cell-centered coordinate, \tilde{x} , that we use to define the interpolation parabola in a cell. There is a discontinuous jump in the integration constant when we go from one cell to its neighbor. This jump is

$$\Delta A_L = \Delta\rho_L - \Delta p_L / c_L^2 = \left(\langle \rho \rangle - \langle \rho \rangle_{i-1} \right) - \left(\langle p \rangle - \langle p \rangle_{i-1} \right) / c_L^2$$

For the gamma-law equation of state, we estimate c_L^2 via

$$c_L^2 = \frac{1}{2} \left(\langle c \rangle_{i-1}^2 + \langle c \rangle^2 \right) = \frac{1}{2} \left[\frac{\gamma \langle p \rangle_{i-1}}{\langle \rho \rangle_{i-1}} + \frac{\gamma \langle p \rangle}{\langle \rho \rangle} \right]$$

Here we see a common practice used in PPM, namely the evaluation of an estimate of the cell average of a function of primitive fluid state variables as the function of the cell averages. Experience shows that attempts to do better than this are usually unrewarded by compensating increases in simulation accuracy. In the specific case of the kinetic energy, PPM does do a better job, but the effect on the flow accuracy is only marginal in this case, and the cost is not insignificant.

Using the above definitions, we evaluate a fractional error estimate as follows:

$$f_{err} = \frac{1}{2} \frac{\left| \Delta A_{L, i+2} - 2 \Delta A_R + \Delta A_L \right| + \left| \Delta A_R - 2 \Delta A_L + \Delta A_{L, i-1} \right|}{\left| \Delta A_L \right| + \left| \Delta A_R \right| + d\alpha}$$

Here $d\alpha$ represents a trivial Riemann invariant difference that is used to protect the divide in the Fortran program. We note that this error estimate is very similar to but not exactly the same as the one discussed earlier. Using this error estimate, we construct a measure, Ω , varying between 0 and 1, of the roughness of the function near this grid cell as follows:

$$\Omega = \min \left\{ 1, \max \left\{ 0, 10(f_{err} - 0.1) \right\} \right\}$$

We now construct estimates A_{s1} and A_{m1} of the first-order parabola coefficient, which gives the change of the variable across the cell. The first, A_{s1} , applies to the parabola that has the 3 prescribed average values in this and the nearest neighbor cells. The second, A_{m1} , applies to this parabola after the application of a monotonicity constraint. These quantities will be used in constructing further quantities below. We have already seen that $A_{s1} = (\Delta A_L + \Delta A_R)/2$. Defining a sign variable, s , of absolute value unity to have the sign of A_{s1} , we may then compute A_{m1} in the following sequence of steps:

$$\begin{aligned} \tilde{A}_{m1} &= 2 \min \left\{ s \Delta A_L, s \Delta A_R \right\}, & \tilde{\tilde{A}}_{m1} &= s \max \left\{ 0, \min \left\{ s A_{s1}, \tilde{A}_{m1} \right\} \right\} \\ A_{m1} &= \Omega A_{s1} + (1 - \Omega) \tilde{\tilde{A}}_{m1} \end{aligned}$$

(The superiority of Fortran over mathematical notation for the description of numerical algorithms is here becoming apparent.) We now define two estimates, A_{sL} and A_{mL} , of the value at the left-hand cell interface. The first of these, A_{sL} , corresponds to an unconstrained interpolation polynomial, while A_{mL} corresponds to one that is constrained.

$$A_{sL} = \frac{1}{6} \left(A_{s1, i-1} - A_{s1} \right) - \frac{1}{2} \Delta A_L, \quad A_{mL} = \frac{1}{6} \left(A_{m1, i-1} - A_{m1} \right) - \frac{1}{2} \Delta A_L$$

We note that these interface value estimates involve the cell averages in two cells on either side of the interface. The unconstrained estimate is the value of the unique cubic curve that has the prescribed average values in these 4 cells. The constrained value is guaranteed to lie within the range defined by the cell averages adjacent to the interface. Of course, both these values refer to the choice of integration constant that gives a vanishing cell average to the right of this interface (in the cell of interest, i). We desire a linear combination of these two interface value estimates, which we will regard as a provisional estimate $A_{pL} = (1 - \Omega)A_{sL} + \Omega A_{mL}$. Rather than build this value A_{pL} by blending the unconstrained and constrained estimates, we find it more useful to build blended parabola coefficients, $A_{p1} = (1 - \Omega)A_{s1} + \Omega A_{m1}$, and to form A_{pL} from these. This procedure allows us to make use of the coefficients A_{p1} later in our contact discontinuity detection and steepening algorithm. We now proceed to modify our provisional interface values A_{pL} in cells where the function is not smooth. First, we set our interpolation function to a constant in cells where extrema occur, unless of course the function is smooth there. Thus, if $A_{pL} A_{pR} \geq 0$, we set new provisional values, indicated by a subscript q , as

follows: $A_{qL} = (1 - \Omega) A_{pL}$, $A_{qR} = (1 - \Omega) A_{pR}$. Otherwise we simply set $A_{qL} = A_{pL}$ and $A_{qR} = A_{pR}$. Now, in cells where the function is not smooth, we constrain the internal structure to be monotone. We must make sure in this process not to revise the cell structure further in the cells containing extrema. We therefore construct limiting values, A_{iL} and A_{iR} , which we must be careful to set to 0 in cells containing extrema. In the remaining cells, these values are: $A_{iL} = -2A_{pR}$ and $A_{iR} = -2A_{pL}$. This reflects the fact that a parabola having zero slope at one side of the cell will assume at the opposite interface the negative of twice this value, so long as the cell average is zero. If these limiting values are exceeded in cells where the function is not smooth, we must reset them to these limits. Therefore we construct new, but still provisional, values A_{rL} and A_{rR} (in Fortran, of course, no new names are required), which reflect this additional constraint:

$$\begin{aligned} \text{if } (A_{qR} - A_{qL})(A_{qL} - A_{iL}) < 0, \quad \text{then } A_{rL} &= (1 - \Omega)A_{qL} + \Omega A_{iL}, \quad \text{otherwise } A_{rL} = A_{qL} \\ \text{if } (A_{qR} - A_{qL})(A_{iR} - A_{qR}) < 0, \quad \text{then } A_{rR} &= (1 - \Omega)A_{qR} + \Omega A_{iR}, \quad \text{otherwise } A_{rR} = A_{qR} \end{aligned}$$

From these provisional interface values, we construct provisional interpolation parabola coefficients:

$$A_{r1} = A_{rR} - A_{rL}, \quad A_{r2} = 3(A_{rL} + A_{rR}), \quad A_{r0} = -A_{r2}/12$$

For interpolation of sound-wave Riemann invariant differences, for which we do no contact discontinuity steepening, the above values are our final results. However, for the entropy Riemann invariant differences, we continue as described below. We also note that for interpolation of a variable, such as the transverse velocity component, which is not expressed as a differential form, we may first construct differences of the cell averages and then proceed as outlined above.

We wish to detect cells that are inside sharp jumps in the entropy Riemann invariant that are associated with contact discontinuities. In such cells, it is inappropriate to try to fit smooth curves to determine the subgrid structure, since it is actually discontinuous (PPM is solving the Euler equations). For the purpose of computation, we will nevertheless require that the distribution inside the grid cell be a parabola; however, we may use the knowledge that the actual structure is discontinuous to construct an appropriate choice for this parabola. The idea here is quite simple. We build a test that detects cells that are within contact discontinuity structures. Then we obtain estimates of the interface values for such cells by extrapolating to the cell interfaces presumably smoother structures from outside the cell and hence, hopefully, from outside the discontinuity. Together with the prescribed cell averages, these more appropriate cell interface values allow us to build an improved, steeper parabola to describe the distribution within the cell. Obviously, the function in such cells is not smooth, so we apply monotonicity constraints to the new edge values and also to this improved parabola. This procedure works very well in practice. A similar procedure can be adapted for use with other numerical schemes that employ linear interpolation functions for the cell structures, but it is less effective. Apparently, the curvature provided by the use of parabolae allows the dimensionless constants that will be introduced below to make the steepening process effective for contact discontinuities without requiring that the steepening process be

applied for marginal cases, which can occasionally turn out not to actually be contact discontinuities. That is, using parabolae we are able to apply the steepening process only in unequivocal circumstances and thus to avoid steepening structures by mistake that should not actually be steepened.

We begin, as with our test for function smoothness, with a measurement of the size of the third derivative of the function relative to its first derivative. Reusing the symbol s for a different sign variable, we build a steepness measure S in a sequence of steps as follows:

$$S_1 = \left(A_{sR} - A_{sL} - A_{s1} \right) / \left(A_{s1} + 10 |s| d\alpha \right)$$

Here s has absolute value unity and has the same sign as A_{s1} in order to protect the divide without altering the sign of the overall expression. As before, $d\alpha$ is a trivial Riemann invariant difference. We must take care to realize that the numerator in this expression is not necessarily zero, since the edge values that appear there come from cubic curves rather than the parabolae that define the A_{s1} values. We must also take care to realize that because each cell has its own value of the integration constant for this Riemann invariant's differential form, we must set $A_{sR} = A_{sL, i+1} + \Delta A_R$. With those cautions, we see that this expression for S_1 is very similar to our measure of the lack of function smoothness, except that no absolute value signs appear in it. The first expression in the numerator for the variable difference across the cell involves cubic interpolation, which accounts for the third derivative of the function. From this we subtract A_{s1} , which removes the part that is due to the first 2 function derivatives. To obtain S_1 , we divide by an estimate of the first derivative. The lack of absolute value operations in this formula for S_1 is caused by our desire to detect sharp jumps in the function, and to reject sudden flat spots. Flat spots will correspond to negative values of S_1 . We now scale S_1 and limit it to the range from 0 to 1:

$$S_2 = \max \left\{ 0, \min \left\{ 1, 20 \left(S_1 - 0.05 \right) \right\} \right\}$$

To guard against applying our contact discontinuity steepening anywhere but at true sharp jumps in the entropy, we reset S_2 to zero where the second derivative of the function does not change sign and also where the amplitude of the jump is not sufficient to warrant this special treatment (we do not want to preserve, or worse, to amplify small numerical glitches). Thus:

$$\text{if } A_{s2, i-1} A_{s2, i+1} \geq 0, \quad \text{then } S_3 = 0, \quad \text{otherwise } S_3 = S_2$$

Here we use the earlier definition of the coefficients of the parabola that has the prescribed 3 cell averages, writing $A_{s2} = (\Delta A_R - \Delta A_L)/2$. To eliminate trivial jumps from consideration, we form S_4 :

$$\text{if } |\Delta A_L| + |\Delta A_R| < d\alpha, \quad \text{then } S_4 = 0, \quad \text{otherwise } S_4 = S_3$$

Before proceeding with the steepening algorithm, we make one further test to make certain that steepening is appropriate. This test makes sense for contact discontinuities, but for other variables, such as a fractional volume of a second fluid, it might not. Therefore this last test is optional, but we do use it for the entropy. The calling program to the interpolation routine provides two variable differences along

with the differences ΔA_L . If the jump is one of the type we seek to detect, the differences ΔB will be very small compared with the differences ΔD . Both these sets of differences are cell centered and have the same units. We enforce this demand by constructing our final steepness measure, S , as follows:

$$\text{if } \Delta D = 0 \text{ or } \Delta B / \Delta D \geq 0.1, \quad \text{then } S = 0, \quad \text{otherwise } S = S_4$$

In previous versions of PPM, the ΔB were pressure differences and the ΔD were estimates of $c^2 \Delta \rho$, while the density was the quantity being interpolated with potential contact discontinuity steepening. We still perform this additional test, although it is somewhat redundant when interpolating the entropy Riemann invariant, for which the differential form is just $\Delta A = \Delta \rho - \Delta p / c^2$.

With our contact discontinuity detection process complete, and all cells inside such structures marked by S , varying from 0 to 1, we are at last ready to begin the steepening operation. Much earlier, we computed constrained variable differences A_{m1} across the cells. In principle, these variable differences could be only partially constrained, but it is safe to assume that if our cell is in a contact discontinuity, the parameter Ω that measured the function roughness assumes the value unity in this and the neighboring cells. We define cell interface values, A_{cL} and A_{cR} , appropriate for a contact discontinuity structure by extrapolating the constrained slopes of the function in neighboring cells to the interfaces of our cell of interest: $A_{cL} = -\Delta A_L + A_{m1, i-1} / 2$ and $A_{cR} = \Delta A_R - A_{m1, i+1} / 2$. We then construct our nearly final estimates for the interface values as:

$$A_{rL} = (1 - S) A_{rL} + S A_{cL}, \quad A_{rR} = (1 - S) A_{rR} + S A_{cR}$$

We must once again apply the constraints on the implied parabola and then compute the coefficients of that parabola, just as before (however, hardly any cells are steepened, and hence this extra work occurs in a scalar loop for only these cells, which incurs hardly any additional computational cost). These operations will not be repeated here. They are just like those we performed to arrive at first A_{qL} and then A_{rL} beginning with A_{pL} .

It is worthwhile to note the ways in which the above interpolation algorithm addresses the previously listed design constraints for PPM. Some of the points are obvious, but others are less so. In particular, it is not obvious that dissipation errors will dominate dispersion errors, although this is in fact the case. The reason for this is the action of the monotonicity constraints on all short wavelength disturbances, which always fail our test for function smoothness. The monotonicity constraints are nonlinear, in that they alter the shape of a sine wave by introducing higher frequency components through effects such as the clipping of extrema. Their overall effect is strongly dissipative for very short wavelength signals, which of course are the ones for which the numerical scheme would otherwise introduce significant dispersion errors. This dominance of dissipation over dispersion error as a design goal may seem incompatible with the goal of minimal dissipation. Our task is to deliver as little dissipation as possible while still having dissipation overwhelm dispersion. PPM's use of parabolae rather than the more commonly used linear interpolation functions is meant to address this design goal. The use of cubic interpolation functions to help define these parabolae through interpolation of cell interface

values reduces dissipation errors still further, since not all parabolae provide equally good fits. However, even more importantly this use of cubic interpolation in PPM preserves the accuracy of the scheme in the limit of vanishing Courant number. Since, as is common practice, we will determine a single time step value for the entire grid based upon Courant number limitations for the single most demanding cell, the bulk of the cells we update will have very small Courant numbers. When AMR is added to the scheme, the Courant numbers used for the bulk of the cells are likely to become even smaller. Therefore it is very important for the accuracy of the calculation to hold up in this limit. The use of cubic interpolation of cell interface values is our way of addressing this issue in PPM. Over many years of applying PPM to a wide variety of flow problems, very little Courant number sensitivity has ever appeared.

The monotonicity constraints and contact discontinuity detection and steepening algorithms incorporated in PPM have important consequences for the propagation of barely resolved or effectively unresolved signals. Our design goal is to preserve these signals, so long as we can propagate them at roughly the proper speeds, rather than to destroy them through numerical diffusion processes. The monotonicity constraints and, to a far greater degree, the contact discontinuity steepening in PPM have this effect. They act upon passively advected signals, such as entropy variations, that can easily be propagated at the correct speeds. They apply an artificial compression along with their other effects, such as maintaining positivity where appropriate. This tends to maintain signal amplitude while altering signal shape when the signal is not adequately resolved. As the description of the contact discontinuity detection algorithm shows, we take great care to apply the steepening method only where appropriate, since some small signals, better described as numerical glitches of various types and causes, deserve to be dissipated.

4. Using the Interpolation Operators to Build a Subgrid-Scale Model for a Cell.

The above interpolation process is quite elaborate. Nevertheless, it is not the entire story of PPM interpolation. As we remarked earlier, the equations of gas dynamics are coupled, and therefore the PPM interpolations of the fluid state variables are also coupled. The goal is to come up with a complete and consistent picture of the internal structure of a grid cell – that is, to come up with a subgrid-scale model. This is not a turbulence model, but it *is* a subgrid-scale model. This model must satisfy, for consistency, certain primary constraints. First, the integrated mass, momenta, and total energy must be consistent with the prescribed cell averages of these quantities. This reflects the need for the numerical scheme to be in strict conservation form in order to correctly capture and propagate shocks. PPM regards the fundamental cell averages on which it operates to be the cell averages of density and pressure (both volume-weighted) and those of the 3 velocity components (all mass-weighted). From these, PPM defines the mass, momenta, and total energy of a cell to be: $\Delta m = \langle \rho \rangle \Delta x$, $\langle u_x \rangle \Delta m$, $\langle u_y \rangle \Delta m$, $\langle u_z \rangle \Delta m$, and $\frac{\langle p \rangle}{\gamma - 1} \Delta x + \frac{1}{2} (\langle u_x \rangle^2 + \langle u_y \rangle^2 + \langle u_z \rangle^2) \Delta m$. Here we do not bother to include the cell widths in the transverse directions, Δy and Δz , because they cancel out of all our equations for the x-pass. The last expression, for the cell's total energy, is misleading. At the beginning of the grid cell

update for a 1-D pass, we compute the total energy of the cell in just this way. However, once we have done this, we make a more accurate estimate of the cell's kinetic energy and then revise its thermal energy, keeping this total constant. We do this by using nearest neighbor cells to compute cell-centered slope estimates for all 3 velocity components. We then apply the standard monotonicity constraints to these slopes. All these computations are similar to the following steps, for the slope of u_y in the z direction:

$$\Delta_z u_y = \frac{1}{2} \left(\langle u_y \rangle_{k+1} - \langle u_y \rangle_{k-1} \right), \quad (\Delta_z u_y)_{\max} = 2 \min \left\{ s \left(\langle u_y \rangle_{k+1} - \langle u_y \rangle \right), s \left(\langle u_y \rangle - \langle u_y \rangle_{k-1} \right) \right\}$$

$$\Delta_{mz} u_y = s \max \left\{ 0, \min \left\{ s \Delta_z u_y, (\Delta_z u_y)_{\max} \right\} \right\}$$

Here s has absolute value unity and the same sign as $\Delta_z u_y$. We now estimate the cell's average kinetic energy to be given by

$$\begin{aligned} 2 \langle E_{kin} \rangle &= \langle u_x \rangle^2 + \frac{1}{12} \left[(\Delta_{mx} u_x)^2 + (\Delta_{my} u_x)^2 + (\Delta_{mz} u_x)^2 \right] \\ &+ \langle u_y \rangle^2 + \frac{1}{12} \left[(\Delta_{mx} u_y)^2 + (\Delta_{my} u_y)^2 + (\Delta_{mz} u_y)^2 \right] \\ &+ \langle u_z \rangle^2 + \frac{1}{12} \left[(\Delta_{mx} u_z)^2 + (\Delta_{my} u_z)^2 + (\Delta_{mz} u_z)^2 \right] \end{aligned}$$

This process captures the first correction to the average kinetic energy that arises from the internal structure of the velocity within the grid cell. The method would still be second-order accurate without this correction, but this correction proves to be of some marginal value (relative to its implementation cost) in flows involving strong shear layers. Terms of this sort make up part of some turbulence closure models, although one might argue that this is inappropriate, since they have nothing specifically to do with turbulence. In PPM we must compute estimates for cell averages of several other quantities, and including terms like these for all those computations would roughly double the cost of the scheme. This has been tried, and little noticeable benefit is delivered in practical problems. Therefore only these velocity terms are included in PPM. One can see why these terms are important as follows. The velocity slopes can be very large in a cell, because shear layers that are stretching due to the local flow or due to their own instability naturally become ever thinner. The contribution to the kinetic energy from the velocity slope terms can therefore be significant, and can thus have a significant effect upon the pressure (since the total energy is prescribed), and hence upon the dynamics. Nevertheless, this effect upon the dynamics is strongly localized, so that omitting the terms, as was done in PPM for many years, only tends to increase by roughly 50% the numerical friction in thin, strong (roughly sonic or supersonic) shear layers.

In order to propagate signals, we need to interpolate their subgrid structures. We therefore apply the previously described interpolation algorithm to the 5 Riemann invariants of the Euler equations for 3-D compressible flow. For the two transverse components of the velocity, u_y and u_z , which are advected passively with the fluid velocity in the x-pass, we form differences and then apply the

interpolation scheme described above. We do no contact discontinuity detection or steepening for these variables, even though they may in fact jump at such discontinuities. We have already discussed in detail the treatment for the entropy Riemann invariant. The sound wave Riemann invariants, which we denote by R_{\pm} , are defined by the following differential forms and their corresponding numerical approximations:

$$dR_{\pm} = du_x \pm \frac{dp}{C} \quad \Delta R_{\pm L} = \Delta u_{xL} \pm \frac{\Delta p_L}{C_L} = \langle u_x \rangle - \langle u_x \rangle_{i-1} \pm 2 \left(\langle p \rangle - \langle p \rangle_{i-1} \right) / \left(\langle C \rangle + \langle C \rangle_{i-1} \right)$$

Here we introduce the Lagrangian sound speed $C = \rho c$, and we approximate its cell average as $\langle C \rangle = \langle \rho \rangle \langle c \rangle = \sqrt{\gamma \langle p \rangle \langle \rho \rangle}$. The pressure, of course, is p , and only the x-component of velocity appears because we are describing the x-pass of the directionally split PPM algorithm. It is disturbances in R_+ that propagate to the right at speed $s_+ = u_x + c$ and disturbances in R_- that propagate to the left at speed $s_- = u_x - c$. We interpolate parabolae to describe the structure of these sound wave signals in the grid cells, using $\Delta R_{\pm L}$ in place of ΔA_L in the algorithm previously described, and of course performing no contact discontinuity detection or steepening. Performing interpolations for the Riemann invariants attempts to uncouple the gas dynamic equations as much as possible. However, we will not work in terms of these variables directly. Instead, we will immediately go about constructing interpolation parabolae for the primary fluid state variables – the density, pressure, and 3 components of velocity – in a manner that attempts to achieve as much consistency as possible with the interpolated structures of the Riemann invariant signals. To perform our hydrodynamical cell update, we will require internal cell structures for all these variables as well as the total energy. Any choice of structures for 5 variables will imply structures for all the others, but those implied structures may not satisfy reasonable constraints, such as monotonicity or positivity when appropriate. PPM attempts to achieve consistency between interpolated and implied structures by first interpolating constrained parabolae for the Riemann invariant signals, constructing the implied parabolae for the primary state variables, and then constraining those implied parabolae where appropriate. In a sense this is an attempt to “have it both ways,” which is of course impossible. However, experience has shown that the extra labor involved in this process delivers additional accuracy and robustness of a value worthy of its computational and programming cost.

From the interpolated parabolae for the sound wave Riemann invariants, we may compute the cell interface values of the pressure and x-velocity as follows:

$$p_L = \langle p \rangle + C_L (R_{+L} - R_{-L}) / 2, \quad p_R = \langle p \rangle + C_R (R_{+R} - R_{-R}) / 2$$

$$u_{xL} = \langle u_x \rangle + (R_{+L} + R_{-L}) / 2, \quad u_{xR} = \langle u_x \rangle + (R_{+R} + R_{-R}) / 2$$

We also compute a measure, Ω , of the lack of smoothness of the associated functions as the maximum of the values found for the two sound wave Riemann invariants separately during the process of their interpolation: $\Omega = \max \{ \Omega_{R_+}, \Omega_{R_-} \}$. We will use this measure Ω to control the application of constraints to our interpolation parabolae for p and u_x . We first apply monotonicity constraints,

controlled by Ω , to the cell interface pressures and x-velocities obtained above. When $\Omega = 1$, we demand that the interface values lie within the ranges defined by the averages in adjacent cells. We then, again when $\Omega = 1$, demand that the parabola defined by these constrained interface values and the cell averages are monotone. These constraints are essentially the same as those described earlier for the interpolation of the entropy Riemann invariant, and hence they will not be stated in detail here. Once the interpolation parabola for the pressure has been so defined, we may use the interpolated interface values for the entropy to determine interface values for the density. Using the maximum of the Ω measures for the pressure and for the entropy, we may then constrain the interface densities and ultimately the parabola that they and the cell averages define.

At the end of this lengthy process, we have interpolation parabola defined for the density, pressure, and all 3 velocity components. The formulae presented assume uniform cell sizes, although more general formulae are easily derived and were presented in the description of PPM in [9]. We can think of these interpolations as occurring in a cell number variable, so that they are valid, although potentially somewhat less accurate, even if the cell size is smoothly varying. Nevertheless, we must take care in interpreting these parabola. Some of the variables described by them, such as the velocity components, must be considered to vary according to mass fraction across each cell, since the (mass-weighted) cell average is associated directly with the conserved cell momentum. For the velocities, therefore, we take the interpolation variable \tilde{x} within the cell to have its origin at the center of mass and to describe fractions of the cell mass rather than of the cell volume. For the pressure or the density, whose volume integrals are directly associated with the mass and internal energy, we must interpret \tilde{x} as having its origin at the center of the cell in a volume coordinate and to describe volume fractions within the cell. Our use of the uniform grid formulae therefore reflects an assumption that all interpolated quantities are smoothly varying in a cell fraction variable, be it a volume or a mass fraction. We know this assumption to be false at contact discontinuities and slip surfaces (for some of the variables) and at shocks, but we have augmented our interpolation procedure to deal with these discontinuities. The interpolation is therefore valid, so long as we interpret it properly.

The process described above has enabled us to construct a subgrid-scale model for each cell. Since each fluid state variable varies as a simple parabola inside a cell, this subgrid-scale model cannot possibly describe subgrid-scale turbulent motions. To perform that function, the model would have to be augmented by the addition of one or more new state variables, such as a subgrid-scale turbulent kinetic energy variable, which could also be given an interpolation parabola within the cell. Nevertheless, our demand that the structure of a velocity component within a cell be no more complicated than a parabola, and our constraints on that parabola, provide a powerful mechanism to dissipate kinetic energy of small-scale motions unresolvable or marginally resolvable on our grid into heat. This is one of the important functions of turbulence closure models, especially when incorporated into difference schemes that lack any other method to accomplish this very necessary dissipation. In this respect, turbulence models can be used to stabilize otherwise nonlinearly unstable numerical schemes, a role that has nothing to do with turbulence itself and that confuses the purpose and function of a turbulence model. As we will see later,

with PPM we may compute turbulent flow directly from the Euler equations that govern it, and we accomplish the dissipation of small-scale motions through the truncation errors of our numerical scheme rather than from differencing an explicit viscous diffusion term. We take the view that the details of this dissipation are unimportant so long as kinetic energy turning up on the smallest possible scales in our computation is dissipated into heat with no unphysical side effects. We presume that in a far more expensive and careful simulation of all relevant physics, this same kinetic energy would be dissipated in any event, and the same amount of heat generated (since total energy is conserved), but the (hopefully unimportant) details on tiny length and time scales would differ. The assumption here is that the kinetic energy that appears on scales where the numerical dissipation comes into play has arrived there due to a physical process, such as a turbulent cascade, that is correctly simulated in our calculation based on the Euler equations because it has nothing to do with the molecular viscosity which has been omitted in deriving the Euler equations from the more general Navier-Stokes equations.

5. The PPM Approximate Riemann Solver, for the special case of weak waves.

We will update the cell averages in time by applying the laws of conservation of mass, momentum, and total energy. To do this, for a 1-D x-pass of the algorithm, we must compute time-averaged fluxes at the left- and right-hand cell interfaces of these 5 conserved quantities. This is done with the help of an approximate Riemann solver. A fundamental approximation that we make is that the solution to the Riemann problem for left and right states representing spatial averages over the appropriate domains of dependence is roughly equivalent to the average in time of the many separate Riemann problem solutions that each express the interaction of the specific Riemann invariant values arriving at the interface at that particular time. The Riemann solver is a nonlinear operator, and we assume that its output when operating on the averaged inputs is equivalent to the average of the outputs from the separate, time-varying inputs. This assumption makes the Riemann solver relatively simple and efficient. Our elaborate construction of subgrid-scale structures allows us to easily arrive at the spatial averages in the appropriate domains of dependence. In the linear regime, when we have simple advection of the Riemann invariants, we get the full advantage of or interpolation parabolae for the Riemann invariant signals. Experience shows that when variations in R_{\pm} are large, shocks rapidly develop, and these demand techniques only loosely, if at all, related to concepts of formal order of accuracy. These techniques are not discussed here, and the reader is referred to [54].

Here we will discuss how we handle the domain of dependence for the Riemann invariant R_{+} . The treatment for R_{-} is similar, and will not be described. We will treat the characteristic speed $s_{+} = u + c$ to be constant both in space and time within each cell for this single time step. Accounting for the space and time variation of s_{+} is demanded only if we desire formal third-order accuracy. We will focus our discussion on the left-hand interface of our cell of interest. We calculate a Courant number, σ_{+L} , which applies to one or the other of the adjacent cells, depending upon the signs of s_{+} in them:

$$\sigma_{+L} = s_{+, i-1} \Delta t / \Delta x, \quad \text{if } s_{+, i-1} \geq 0; \quad \text{otherwise} \quad \sigma_{+L} = -s_{+} \Delta t / \Delta x, \quad \text{if } s_{+} \leq 0$$

We note that the two possibilities listed above are not exhaustive. In the case where s_+ is directed away from the interface on both sides, we will have a centered rarefaction there. We will handle this case later. For the moment it will be sufficient to set σ_{+L} to zero in this rare event. The Courant condition demands that we control the time step Δt so that σ_{+L} does not exceed unity. We now estimate the spatial averages of the density and pressure, $\langle \rho \rangle_{+L}$ and $\langle p \rangle_{+L}$, in the domain of dependence via:

$$\begin{aligned} \langle \rho \rangle_{+L} &= \rho_{R,i-1} - \frac{\sigma_{+L}}{2} \left[\rho_{1,i-1} + \left(1 - \frac{2\sigma_{+L}}{3} \right) \rho_{2,i-1} \right], \quad \text{if } s_{+,i-1} \geq 0; \\ \langle \rho \rangle_{+L} &= \rho_L + \frac{\sigma_{+L}}{2} \left[\rho_1 - \left(1 - \frac{2\sigma_{+L}}{3} \right) \rho_2 \right], \quad \text{otherwise.} \end{aligned}$$

We obtain $\langle p \rangle_{+L}$ using similar formulae. From $\langle \rho \rangle_{+L}$ we may calculate a mass fraction Courant number, σ_{m+L} . When $s_{+,i-1} \geq 0$, we have $\sigma_{m+L} = \sigma_{+L} \langle \rho \rangle_{+L} / \langle \rho \rangle_{i-1}$, and otherwise we have $\sigma_{m+L} = \sigma_{+L} \langle \rho \rangle_{+L} / \langle \rho \rangle$. We can then obtain $\langle u_x \rangle_{+L}$ using the same formulae as for the density or pressure, but with σ_{+L} replaced by σ_{m+L} . Note that we choose the upwind domain of dependence in the case when R_+ Riemann invariant signals reach the interface from both adjacent cells. This can happen if the interface is inside a shock structure, in which case the R_+ signals from the cell on the right will become lost in the shock transition. Hence we take the signals from the cell on the left in this case, which will correspond to the proper post-shock state. For the R_- signals, we prefer signals from the cell on the right, if such signals can reach the interface from both adjacent cells. In the case when $s_{+,i-1} < 0$ and also $s_+ > 0$, the left-hand interface of our cell is located inside a centered rarefaction fan. We desire Riemann invariant information corresponding to a vanishing characteristic speed. We interpolate $\langle \rho \rangle_{+L}$, $\langle p \rangle_{+L}$, and $\langle u_x \rangle_{+L}$ linearly between the two edge values, as in: $\langle p \rangle_{+L} = f_{+L} p_{R,i-1} + (1 - f_{+L}) p_L$, where $f_{+L} = s_+ / (s_+ - s_{+,i-1})$.

A simple, linearized solution to the Riemann problem at the cell interface can be obtained from the demands that the time-averaged state at the interface have the same values of R_{\pm} as the domains of dependence which we have just determined. At this point we desire only the time-averaged pressure and x-velocity at the interface, which we denote by \bar{p}_L and \bar{u}_{xL} . Using the differential forms that define the Riemann invariants, we thus have: $\bar{u}_{xL} - \langle u_x \rangle_{\pm L} \pm (\bar{p}_L - \langle p \rangle_{\pm L}) / C_L = 0$. Therefore:

$$\begin{aligned} \bar{u}_{xL} &= \frac{1}{2} \left(\langle u_x \rangle_{+L} + \langle u_x \rangle_{-L} \right) + \frac{1}{C_L} \left(\langle p \rangle_{+L} - \langle p \rangle_{-L} \right) \\ \bar{p}_L &= \langle p \rangle_{-L} + C_L \left(\bar{u}_{xL} - \langle u_x \rangle_{-L} \right) \end{aligned}$$

Because this is a numerical scheme, and anything can happen, we subsequently make sure that \bar{p}_L is at least as big as some trivial, floor value. This simple, linear Riemann solver suffices for almost all the cell interfaces, but it is not good enough for interfaces that are in strong shock structures. For such cells, which are very small in number, we can do much better. The multidimensional test we apply to locate

such cells, how we revise the above Riemann solver, and how we compute and apply a smart dissipation in pathological cases are described elsewhere [54].

Up to this point, we have worked only with the time-averaged x-velocities and pressures at the cell interfaces. Now that we know the time-averaged interface velocities, we can trace back along the fluid streamlines from the interfaces into the adjoining cells in order to estimate what fraction of which cell crosses the interface during the time step. We may also estimate the time-averaged density at the interface by taking the spatial average at the beginning of the time step of this material that crosses the interface and compressing or expanding it *using the shock jump conditions* to the time-averaged pressure at the interface, which has just been computed. In the absence of shocks, this compression or expansion should be adiabatic, but if a strong shock is involved, the behavior could be extremely different. Using the shock jump conditions will give us an acceptable result in either case. We begin by computing the spatial averages, $\langle \rho \rangle_{0L}$, $\langle p \rangle_{0L}$, $\langle u_y \rangle_{0L}$, and $\langle u_z \rangle_{0L}$, in essentially the same way that we computed similar spatial averages earlier over the domains of dependence of the R_{\pm} Riemann invariant characteristics. Here we use the subscript 0 to denote the streamline characteristics. The time-averaged density at the cell interface is now:

$$\bar{\rho}_L = \langle \rho \rangle_{0L} \frac{(\gamma + 1) \bar{p}'_L + (\gamma - 1) \langle p \rangle_{0L}}{(\gamma - 1) \bar{p}'_L + (\gamma + 1) \langle p \rangle_{0L}}$$

6. Construction of the Time-Averaged Interface Fluxes and Application of the Conservation Laws.

Now that we have evaluated the time averages of the fluid state variables at the cell interfaces, it is a simple matter to construct from them the time-averaged fluxes of the conserved quantities. At each interface, we compute an advected volume, dx_L , an advected mass, dm_L , 3 advected momenta, $d\mu_{xL}$, $d\mu_{yL}$, $d\mu_{zL}$, and an advected total energy, dE_L . These are all signed quantities, with positive signs applying to advection to the right:

$$\begin{aligned} dx_L &= \bar{u}'_{xL} \Delta t, & dm_L &= \bar{\rho}_L dx_L, & d\mu_{xL} &= \bar{u}'_{xL} dm_L + \bar{p}'_L \Delta t, & d\mu_{yL} &= \bar{u}_{yL} dm_L \\ d\mu_{zL} &= \bar{u}_{zL} dm_L, & dE_L &= \frac{1}{2} \left[\left(\bar{u}'_{xL} \right)^2 + \left(\bar{u}'_{yL} \right)^2 + \left(\bar{u}'_{zL} \right)^2 \right] dm_L + \frac{\gamma}{\gamma - 1} \bar{p}'_L dx_L \end{aligned}$$

To guard against the production of negative densities, we also demand that dm_L not exceed advection of 95% of the upstream cell's mass. Our time step controls should prevent this situation, but this constraint provides additional code robustness.

We now apply the conservation laws for mass, momentum, and total energy in order to arrive at provisional new values for the cell averages at the end of the time step for this x-pass. We use the subscript N to denote the new time level and the superscript (1) to denote this first, provisional value. Despite our application of monotonicity constraints it is still possible in this Eulerian treatment for the new cell mass to be negative. We therefore demand that the new cell mass not fall below a floor value, given by $\rho_{\min} \Delta x$. The value of ρ_{\min} is of course problem dependent, and it must be supplied by

the user for any given run, along with trivial values for other positive definite quantities such as pressure and energy, p_{\min} and E_{\min} . To prevent problems of negative cell masses from occurring, we will compute a cell mass Courant number that we will use to help control the time step. This time step control is the Eulerian equivalent of the one used in Lagrangian calculations to prevent the tangling of grid cells.

$$\Delta m_N^{(1)} = \max \left\{ \left(\Delta m + dm_L - dm_R \right), \rho_{\min} \Delta x \right\}, \quad \langle \rho \rangle_N^{(1)} = \Delta m_N^{(1)} / \Delta x$$

The cell mass Courant number we use is $0.2 \left(\Delta m - \Delta m_N^{(1)} \right) / \Delta m_N^{(1)}$, and we use a similar cell volume Courant number equal to $0.2 \left(\bar{u}'_{xL} - \bar{u}'_{xR} \right) \Delta t / \left[\Delta x - \left(\bar{u}'_{xL} - \bar{u}'_{xR} \right) \Delta t \right]$. These Courant numbers assert that we do not want either the mass of the Eulerian cell or the volume of its Lagrangian counterpart to decrease by more than 80% in a single time step. We usually find that if the Courant number exceeds unity on any time step, even though we have taken measures to see that no disasters, such as square roots of negative numbers, will ensue, it is best to have the computation halve the time step and try again. This feature is built into our PPM code, including its massively parallel versions, and occurs automatically, although, gratefully, very rarely.

The provisional new cell velocities are now obtained from the momentum conservation laws:

$$\langle u_x \rangle_N^{(1)} = \left(\langle u_x \rangle \Delta m + d\mu_{xL} - d\mu_{xR} \right) / \Delta m_N^{(1)}$$

with similar equations for $\langle u_y \rangle_N^{(1)}$ and $\langle u_z \rangle_N^{(1)}$. Although it is optional, we usually choose to reset trivial cell interface and cell average velocities to zero. This prevents the generation of underflows, and it can allow us in some problems to avoid doing any computational labor in cells where nothing is happening yet. From these new velocities, we obtain the provisional new cell kinetic energy, $\langle E_{kin} \rangle_N^{(1)}$.

$$\langle E_{kin} \rangle_N^{(1)} = \frac{1}{2} \left[\left(\langle u_x \rangle_N^{(1)} \right)^2 + \left(\langle u_y \rangle_N^{(1)} \right)^2 + \left(\langle u_z \rangle_N^{(1)} \right)^2 \right]$$

Note that at this point we make only a simple estimate of the cell average of the kinetic energy. It would make the scheme a great deal more complex and expensive to attempt to determine the new internal cell velocity structures at this stage of the calculation. That will be done at the outset of the next 1-D pass, where the total energy in the cell will first be computed according to the above prescription, so that the total energy conservation law, stated below, will be obeyed.

$$\langle E \rangle_N^{(1)} = \left(\langle E \rangle \Delta m + dE_L - dE_R \right) / \Delta m_N^{(1)}, \quad \langle \varepsilon \rangle_N^{(1)} = \max \left\{ E_{\min}, \langle E \rangle_N^{(1)} - \langle E_{kin} \rangle_N^{(1)} \right\}$$

Here we use the symbol ε to represent the internal energy. We now compute the new provisional cell pressure from the gamma-law equation of state:

$$\langle p \rangle_N^{(1)} = (\gamma - 1) \langle \rho \rangle_N^{(1)} \langle \varepsilon \rangle_N^{(1)}$$

We now have a complete set of provisional new cell averages. These will actually be our final results for the overwhelming majority of the cells. However, near slowly moving strong shocks we still need to apply

our smart diffusion. This smart diffusion, applied solely in strong shocks under pathological conditions, is described elsewhere [54].

7. Summary.

We have given a complete description of the PPM scheme as it applies to flows containing no strong shocks. Due to the complexity of this method and the limitations of space, we have chosen a complete description of part of this scheme over a partial description of all of it. We have presented no simple examples of the accuracy the scheme delivers in practice, although such examples can be found in the references below. The reader can, however, go to the LCSE Web site at www.lcse.umn.edu, download a Windows application that runs PPM in simple situations, consult the user guide, and run test problems of various classic types for him- or herself. PPM achieves its robustness and accuracy through complexity, and this has its cost. From the above description, that cost might not be clear. On a difficult flow problem involving many strong shocks, on the average each cell update for a 1-D pass involves 921 flops and 249 vectorizable logical operations (which do not count as flops). This work is performed at a speed of 1255 Mflop/s on a 1.7 GHz Intel Pentium-M CPU working from its cache memory in a laptop machine. This laptop performance degrades to 954 Mflop/s for a full 128^3 3-D grid brick update, so that full 128^3 test runs can be performed overnight at a hotel. These speeds are for 32-bit arithmetic, which is all that PPM ever requires. On a 3.2 GHz Intel Pentium-4 CPU, these 32-bit rates become 1607 Mflop/s and 1273 Mflop/s, respectively, while the corresponding rates for 64-bit arithmetic, which is unnecessary but quoted for comparison purposes with other applications, are 1132 Mflop/s and 938 Mflop/s, respectively. PPM performance has scaled essentially linearly on every multiprocessing machine on which the code has ever been implemented, including ASCI machines at DoE labs with over 6000 CPUs [55,56,45].

8. References.

1. van Leer, B., "Towards the Ultimate Conservative Difference Scheme IV: A New Approach to Numerical Convection," *J. Comput. Phys.*, **23**, 276-299 (1977).
2. van Leer, B., "Towards the Ultimate Conservative Difference Scheme V: A Second-Order Sequel to Godunov's Method," *J. Comput. Phys.*, **32**, 101 (1979).
3. van Leer, B., and P. R. Woodward, Proc. Internat. Conf. Comput. Meth. Nonlinear Mech., edited by J. T. Oden (Amsterdam:North-Holland), p. 234.
4. Godunov, S. K., *Nat. Sb.*, **47**, 271 (1959).
5. DeBar, R., "Fundamentals of the KRAKEN Code," Lawrence Livermore National Laboratory Report UCIR-760 (1974).
6. Sutcliffe, W. G., "BBC Hydrodynamics," Lawrence Livermore National Laboratory Report UCID-17013 (1973).

7. Woodward, P. R., and P. Colella, "High-Resolution Difference Schemes for Compressible Gas Dynamics," *Lecture Notes in Phys.* **141**, 434 (1981).
8. Woodward, P. R., and P. Colella, "The Numerical Simulation of Two-Dimensional Fluid Flow with Strong Shocks," *J. Comput. Phys.* **54**, 115-173 (1984).
9. Colella, P., and P. R. Woodward, "The Piecewise-Parabolic Method (PPM) for Gas Dynamical Simulations," *J. Comput. Phys.* **54**, 174-201 (1984).
10. Woodward, P. R., "Numerical Methods for Astrophysicists," in *Astrophysical Radiation Hydrodynamics*, eds. K.-H. Winkler and M. L. Norman, Reidel, 1986, pp. 245-326.
11. Wenlong Dai and P. R. Woodward, "An Approximate Riemann Solver for Ideal Magnetohydrodynamics," *Journal Computational Physics* **111**, 354-372 (1994).
12. Wenlong Dai and P. R. Woodward, "Extension of the Piecewise-Parabolic Method to Multidimensional Ideal Magnetohydrodynamics," *Journal of Computational Physics* **115**, 485-514 (1994).
13. Wenlong Dai and P. R. Woodward, "Interactions between Magnetohydrodynamical Shocks and Denser Clouds," *Astrophysical Journal* **436**, 776-783 (1994).
14. W. Dai and P. R. Woodward, "A Simple Riemann Solver and High-Order Godunov Schemes for Hyperbolic Systems of Conservation Laws," *Journal of Computational Physics* **121**, 51 (1995).
15. Dai, Wenlong, and P. R. Woodward, "A Simple Finite Difference Scheme for Multidimensional Magnetohydrodynamical Equations," *Journal of Computational Physics* **142**, 331-369 (1998).
16. Dai, Wenlong, and P. R. Woodward, "On the Divergence-Free Condition and Conservation Laws in Numerical Simulations for Supersonic Magnetohydrodynamical Flows," *Astrophysical J.*, **493** (1998).
17. Woodward, P. R., "Simulation of the Kelvin-Helmholtz Instability of a Supersonic Slip Surface with the Piecewise-Parabolic Method (PPM)," in *Numerical Methods for the Euler Equations of Fluid Dynamics*, eds. Angrand, Dervieux, Desideri, and Glowinski, SIAM, 1985.
18. Woodward, P. R., "Simulations of Supersonic Jet Instability using the Piecewise-Parabolic Method (PPM)," in *High Speed Computing, Scientific Applications and Algorithm Design*, R. B. Wilhelmson, editor, Univ. of Illinois Press, 1988.
19. Winkler, K.-H., J. W. Chalmers, S. W. Hodson, P. R. Woodward, and N. J. Zabusky, "A Numerical Laboratory," *Physics Today*, Oct., 1987.
20. Pedelty, J. A., and P. R. Woodward, "Numerical Simulations of the Nonlinear Kink Modes in Linearly Stable Supersonic Slip Surfaces," *J. Fluid Mech.*, **225**, 101-120 (1991).
21. Bassett, G. M., and P. R. Woodward, "Numerical Simulation of Nonlinear Kink Instabilities of Supersonic Shear Layers," *Journal of Fluid Mechanics* **284**, 323-340 (Feb. 10, 1995).

22. Bassett, G. M., and P. R. Woodward, "Simulation of the Instability of Mach 2 and Mach 4 Gaseous Jets in 2 and 3 Dimensions," *Astrophysical Journal* **441**, (March 10, 1995).
23. Edgar, B. K., and P. R. Woodward, "Diffraction of a Shock Wave by a Wedge: Comparison of PPM Simulations with Experiment," AIAA Paper 91-0696, and *Video Journal of Engineering Research*, **3**, 25-33+cover illustration (1993).
24. Edgar, B. K., Woodward, P. R., and Anderson, S. E., PPM code library, with documentation and examples, available at www.lcse.umn.edu/PPMlib.
25. Porter, D. H., A. Pouquet, and P. R. Woodward, "Supersonic Homogeneous Turbulence," *Lecture Notes in Physics*, **392**, 105-125 (1991).
26. Porter, D. H., A. Pouquet, and P. R. Woodward, "A Numerical Study of Supersonic Turbulence," *Theoretical and Computational Fluid Dynamics*, **4**, 13-49 (1992).
27. Porter, D. H., A. Pouquet, and P. R. Woodward, "Three-Dimensional Supersonic Homogeneous Turbulence: A Numerical Study," *Phys. Rev. Lett.*, **68**, 3156-3159 (1992).
28. Porter, D. H., A. Pouquet, and P. R. Woodward, "Kolmogorov-Like Spectra in Decaying Three-Dimensional Supersonic Flows," *Phys. Fluids A*, **6**, 2133-2142 (1994).
29. Woodward, P. R., "Superfine Grids for Turbulent Flows," *IEEE Computational Science & Engineering*, Vol. **1**, No. 4, pp. 4-5+cover illustration, (December, 1994).
30. Woodward, P. R., D. H. Porter, B. K. Edgar, S. E. Anderson, and G. Bassett, "Parallel Computation of Turbulent Fluid Flow," Proc. Parallel CFD '94 Conference, Kyoto, Japan, May, 1994; published in *Computing Applications Mathematics*, Vol. **14**, no. 1, pp 97-105 (1995).
31. Mirin, A. A., R. H. Cohen, W. P. Dannevik, A. M. Dimits, D. E. Eliason, D. H. Porter, O. Schilling and P. R. Woodward, "Three-Dimensional Simulations of Compressible Turbulence on High-Performance Computing Systems," Eighth SIAM Conference on Parallel Processing for Scientific Computing, Minneapolis (1997), LLNL Report UCRL-JC-125949.
32. Schilling, O., R. H. Cohen, W. P. Dannevik, A. M. Dimits, D. E. Eliason, A. A. Mirin, D. H. Porter and P. R. Woodward, "Three-Dimensional High-Resolution Simulations of Compressible Rayleigh-Taylor Instability and Turbulent Mixing," Sixth Int'l. Workshop on the Physics of Compressible Turbulent Mixing, Marseille, France (1997), LLNL Report UCRL-JC-125308.
33. Cohen, R. H., W. P. Dannevik, A. M. Dimits, D. E. Eliason, A. A. Mirin, D. H. Porter, O. Schilling, and P. R. Woodward, "Three-dimensional high-resolution simulations of the Richtmyer-Meshkov mixing and shock-turbulence interaction," in Proc. Of the 6th International Workshop on the Physics of Compressible Turbulent Mixing, Marseille, France, 18-21 June, 1997, pp. 128-133. LLNL Report UCRL-JC-125309.

34. Cohen, R. H., A. M. Dimits, A. A. Mirin, W. P. Dannevik, R. G. Eastman, D. E. Eliason, O. Schilling, D. H. Porter and P. R. Woodward, “Three-Dimensional PPM Simulations of Re-shocked Richtmyer-Meshkov Instability,” VHS Video (1997), LLNL Report UCRL-MI-128783.
35. Porter, D. H., P. R. Woodward, and A. Pouquet, “Inertial Range Structures in Decaying Turbulent Flows,” *Physics of Fluids* **10**, 237-245 (1998).
36. Porter, D. H., A. Pouquet, and P. R. Woodward, “Intermittency in Compressible Flows,” Proc. Seventh European Turbulence Conference, St. Jean C. F., France, July, 1998, ed. Uriel Frisch.
37. Porter, D. H., A. Pouquet, I. Sytine, and P. R. Woodward, “Turbulence in Compressible Flows,” *Physica A* **263**, 263-270 (1999).
38. Sytine, I. V., D. H. Porter, P. R. Woodward, S. W. Hodson, and K.-H. Winkler 2000, “Convergence Tests for Piecewise Parabolic Method and Navier-Stokes Solutions for Homogeneous Compressible Turbulence,” *J. Comput. Phys.*, **158**, 225-238 (2000).
39. Cohen, R. H., W. P. Dannevik, A. M. Dimits, D. E. Eliason, A. A. Mirin, Y. Zhou, D. H. Porter, and P. R. Woodward, “Three-Dimensional Simulation of a Richtmyer-Meshkov Instability with a Two-Scale Initial Perturbation,” *Physics of Fluids*, **14**, 3692-3709 (2002); also available as LLNL Preprint UCRL-JC-144836 (2000).
40. Porter, D. H., Pouquet, A., and Woodward, P. R., “Measures of Intermittency in Driven Supersonic Flows,” *Physical Review E* **66**, 026301 (2002).
41. Porter, D. H., and P. R., Woodward, “3-D Simulations of Turbulent Compressible Convection,” *Astrophysical Journal Suppl.* **127**, 159-187 (2000); available at www.lcse.umn.edu/conv3d.
42. Porter, D. H., P. R. Woodward, and M. L. Jacobs, “Convection in Slab and Spheroidal Geometries,” Proc. 14th International Florida Workshop in Nonlinear Astronomy and Physics: Astrophysical Turbulence and Convection, Univ. of Florida, Feb., 1999; in *Annals of the New York Academy of Sciences* **898**, 1-20 (2000); available at www.lcse.umn.edu/convsph.
43. Woodward, P. R., D. H. Porter, I. Sytine, S. E. Anderson, A. A. Mirin, B. C. Curtis, R. H. Cohen, W. P. Dannevik, A. M. Dimits, D. E. Eliason, K.-H. Winkler, and S. W. Hodson, “Very High Resolution Simulations of Compressible, Turbulent Flows,” in *Computational Fluid Dynamics*, Proc. of the 4th UNAM Supercomputing Conference, Mexico City, June, 2000, edited by E. Ramos, G. Cisneros, R. Fernández-Flores, A. Santillan-González, World Scientific (2001); available at www.lcse.umn.edu/mexico.
44. Woodward, P. R., Porter, D. H., and Jacobs, M., “3-D Simulations of Turbulent, Compressible Stellar Convection,” Proc. 3-D Stellar Evolution Workshop, Univ. of Calif. Davis I.G.P.P., July, 2002; also available at www.lcse.umn.edu/3Dstars.
45. Woodward, P. R., Anderson, S. E., Porter, D. H., and Iyer, A., “Cluster Computing in the SHMOD Framework on the NSF TeraGrid,” LCSE internal report, April, 2004, available on the Web at www.lcse.umn.edu/turb2048.

46. LCSE movies from a variety of turbulent fluid flow simulations can be downloaded and/or viewed at www.lcse.umn.edu/MOVIES.
47. Fryxell, B., P. R. Woodward, P. Colella, and K.-H. Winkler, "An Implicit-Explicit Extension of the PPM Scheme for Lagrangian Hydrodynamics in One Dimension," *J. Comput. Phys.* **63**, 283 (1986).
48. Dai, Wenlong, and P. R. Woodward, "A High-Order Iterative Implicit-Explicit Hybrid Scheme for 2-D Hydrodynamics," in *Numerical Methods for Fluid Dynamics V*, pp. 377-383, eds. K. W. Morton and M. J. Baines, Oxford Science Publications (1995).
49. Dai, Wenlong, and P. R. Woodward, "Iterative Implementation of an Implicit-Explicit Hybrid Scheme for Hydrodynamics," *Journal of Computational Physics* **124**, 217-229 (1996).
50. Dai, Wenlong, and P. R. Woodward, "A Second-Order Iterative Implicit-Explicit Hybrid Scheme for Hyperbolic Systems of Conservation Laws," *Journal of Computational Physics* **128**, 181-196 (1996).
51. Dai, Wenlong, and P. R. Woodward, "Numerical Simulations for Radiation Hydrodynamics. I. Diffusion Limit," *Journal of Computational Physics* **141**, 182-207 (1998).
52. Dai, Wenlong, and P. R. Woodward, "Numerical Simulations for Radiation Hydrodynamics II. Transport Limit," *Journal of Computational Physics* **157**, 199-233 (2000).
53. Dingo, D., and P. R. Woodward, "A Parallel Cell-by-Cell AMR Method for the PPM Hydrodynamics Code," *International J. Modern Phys. C*, **14**, 1-24 (2003).
54. Woodward, P. R., 2005. "A Complete Description of the PPM Compressible Gas Dynamics Scheme," LCSE internal report available from the main LCSE page at www.lcse.umn.edu.
55. Woodward, P. R., and S. E. Anderson, "Portable Petaflop/s Programming: Applying Distributed Computing Methodology to the Grid Within a Single Machine Room," Proc. of the 8th IEEE International Conference on High Performance Distributed Computing, Redondo Beach, Calif., Aug., 1999; available at www.lcse.umn.edu/HPDC8.
56. Woodward, P. R., and D. H. Porter, 2005. "PPM Code Kernel Performance," LCSE internal report available at www.lcse.umn.edu.