

Simulation of shock-generated instabilities

R. M. Baltrusaitis^{a)}

*Thermonuclear Applications Group, Applied Theoretical and Computational Physics Division,
Los Alamos National Laboratory, Los Alamos, New Mexico 87545*

M. L. Gittings

*Thermonuclear Applications Group and Hydrodynamic Methods Group, Applied Theoretical
and Computational Physics Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545;
and Science Applications International Corporation, 10260 Campus Point Drive, San Diego,
California 92121*

R. P. Weaver

*Thermonuclear Applications Group, Applied Theoretical and Computational Physics Division,
Los Alamos National Laboratory, Los Alamos, New Mexico 87545*

R. F. Benjamin and J. M. Budzinski

*Hydrodynamic Group Experimentation Division, Los Alamos National Laboratory,
Los Alamos, New Mexico 87545*

(Received 14 March 1996; accepted 3 June 1996)

Direct 2-D numerical simulation of the fluid instability of a shock-accelerated thin gas layer shows flow patterns in agreement with experimental images. The Eulerian-based hydrodynamics code features Adaptive Mesh Refinement that allows the code to follow the vorticity generation and the complex flow resulting from the measured initial perturbations. These experiments and simulations are the first to address in quantitative detail the evolution of the Richtmyer–Meshkov instability in a thin fluid layer, and to show how interfluid mixing and vorticity production depend sensitively on initial perturbations in the layer. © 1996 American Institute of Physics. [S1070-6631(96)03509-X]

I. INTRODUCTION

Hydrodynamic instabilities, turbulence, and mixing constitute some of the most important unsolved problems in physics today. Examples range from astrophysical shocks and jetting¹ to capsule implosion in inertial confinement fusion.² Buoyancy-driven instabilities with small initial perturbations at a fluid interface progress from a phase described by linear theory into a phase of highly distorted flow denoted as “nonlinear growth.” The Rayleigh–Taylor³ (RT) and Richtmyer–Meshkov^{4,5} (RM) instabilities are such interfacial instabilities.⁶ RT instability occurs during constant acceleration of an unstable interface, and RM instability occurs during impulsive or shock acceleration. The linear phase of RT and RM instabilities is fairly well characterized, but the nonlinear instability growth is not well understood. Our goal is to examine the regime of highly nonlinear instability growth. This regime is characterized by the development of intense vorticity, complex and distorted flow patterns,⁷ and persistent features of the initial state.

We simulate the nonlinear instability growth observed in recent experiments involving shock-acceleration of a thin fluid layer.^{8–10} The experiments show that one of three flow patterns emerges from similar initial conditions and that the flows appear to be dominated by vortex dynamics. The experimental images show well-resolved flow patterns, but do not measure other flow characteristics. We report fluid simulations that accurately reproduce the experimentally observed flow patterns and thus enable better understanding of the flow by providing a description of vorticity production and transport. The vorticity distributions show clearly how

asymmetries in the initial interfacial perturbations account for the three patterns. The flow differences between these three different patterns suggest that interfluid mixing at the interfaces will proceed differently for each pattern, and this difference in mixing has implications for applications including inertial confinement fusion.

The direct numerical simulation of instability growth well into the nonlinear regime is becoming more feasible with recent advancements in computer software and hardware. Lagrangian techniques have traditionally been successful at following only the early-time development of these instabilities because the nonlinear growth of perturbations tends to tangle the computational mesh, especially in regions with high shear. Eulerian techniques have had more success but have been hampered by coarse zone-size resolution limited by available computer memory. We have developed an Adaptive Mesh Refinement (AMR) Eulerian code that optimizes the use of computer memory by dynamically refining the zone size of the mesh. This adaptation of the zoning allows the smallest zones (highest resolution) to follow interesting hydrodynamic features of the simulation such as shocks, vorticity, or density gradients. In regions of smaller physical gradients, the mesh recombines into larger zone sizes. For example, a uniform volume of material will be zoned at the coarsest level since all gradients are zero. As a shock wave moves into this volume, the region surrounding the strong pressure gradients of the shock will become highly resolved with fine zoning. After the pressure and density gradients have relaxed behind the shock front, the zone size will increase until it again reaches the coarsest level. This adaptation allows for an accurate direct simulation of instability growth by maintaining high effective resolution, while optimizing the use of computer memory and time.

^{a)}Corresponding author: Rose Mary Baltrusaitis. Electronic mail: rmbalt@lanl.gov

RT and RM instabilities develop when a perturbed interface is subjected to a normal acceleration. When the pressure gradient associated with the acceleration is opposed to the interfacial density gradient, instability occurs as the perturbations grow rapidly. For example, when one inverts a glass of water the RT instability is observed. At the water–air interface, the pressure gradient caused by gravity is pointed downward and the density gradient is directed upward (from the air to the water). Perturbations (i.e., ripples) at the water surface grow and the water spills out. When a fluid interface is accelerated by a shock wave rather than a constant acceleration, it is unstable and known as the “Richtmyer–Meshkov” instability. If pressure and density gradients are aligned in the case of constant acceleration, the interface is stable. However, aligned pressure and density gradients for impulsive or shock acceleration produce an unstable interface characterized by instability growth following phase inversion; i.e., when the shock moves from a dense fluid into a less dense one, the interfacial perturbations grow after reversing phase, with peaks in the original interface turning into valleys and valleys turning into peaks.

Richtmyer⁴ developed an impulse model that predicts a constant perturbation growth rate (following a brief transient) for initially small perturbations. Numerical simulations, beginning with the first Lagrangian simulations reported by Meyer and Blewett,¹¹ have generally agreed with Richtmyer’s prediction of constant growth rate. Most of these computations, as well as the impulse model, predict growth rates greater than observed in single-interface experiments such as those performed by Meshkov⁵ and Benjamin *et al.*¹² A notable exception is the recent computation using front-tracking techniques^{13,14} that shows agreement with experimental growth rates. These studies considered only the problem of a single, unstable interface, not the problem of RM instability at two nearby interfaces, which is the subject of our investigation.

Some applications involve the RM instability of nearby interfaces. For example, the thin shell of an ICF target is subjected to instabilities on the external, ablation/shell interface and on the internal, shell/fuel interface. These applications demand better understanding of RM instability occurring at nearby interfaces. Recent experiments^{8–10} with thin-layer geometry have generated an extensive set of highly-resolved images of time-dependent density distributions induced by the RM instability. Radiation and convergence effects are absent from these experiments which were designed to investigate only the underlying fluid dynamics. These experiments observed instability growth well into the nonlinear regime of a shocked, thin SF₆ gas layer embedded in air, and they found that such configurations produce flows more complex than single-interface experiments. The shock wave traverses both a light-to-heavy and a heavy-to-light interface, and the two interfaces can interact. There are many wave reverberations within the layer including transverse waves which affect the development of the flows observed in the experiments. Thus, this set of experimental observations shows the detailed phenomena of shock accelerating a thin fluid layer and provides a useful benchmark for code performance.

Although some qualitative theoretical work has been based on this data,¹⁵ no detailed quantitative studies have been published to date. We have applied our new computational technique to simulate these experimental results in order to gain better insight into these flows and to validate the code. Because the simulations successfully reproduce the experimental images, we can use the code to better understand and visualize the complex dynamic flows in these experiments. For example, experimentally inaccessible (but theoretically interesting) quantities such as vorticity can be computed and displayed by the code and their dynamic evolution examined. Studying these experiments provides a graphical and relatively controlled example of instability phenomena. The highly-resolved experimental data challenge both our physical understanding and our ability to simulate complex nonlinear hydrodynamics.

Three models have been used in the simulations with increasing degrees of sophistication: (1) pure materials with perturbation amplitudes and wavelengths typical of the data; (2) initially mixed materials with typical perturbation amplitudes and wavelengths; and (3) initial conditions as measured for each individual event. Because the dynamic results are sensitive to the precise initial conditions of the gas curtain, the only quantitative comparison to actual data is from simulations using the third model, but models one and two offer useful qualitative comparisons and suggestions for future work. For example, model one estimates the effect of using discontinuous interfaces rather than the diffusive interfaces of the present experiments.

Only a brief summary of the experimental techniques will be presented here. More detailed descriptions are available in Refs. 8, 9 and 10. The emphasis will be on our recent computational efforts to simulate the data.

II. THE EXPERIMENTS

The experimental discovery of three flow patterns produced by similar initial conditions in a thin layer occurred because of judicious choice of the initial conditions, efficacious application of the laser-sheet technique, and the acquisition of data from over 100 events that enabled the identification of correlations. The experimental techniques and results are summarized below.

Figure 1, taken from Ref. 8, shows a sketch of the test section of the shock tube. SF₆ gas flows vertically downward through a contoured nozzle into an air-filled volume. The nozzle produces a laminar jet of SF₆, and the jet has corrugated interfaces with the air on both sides. Make-up air injected into the shock tube minimizes the shear flow at the air–SF₆ interfaces. This laminar jet flow generates a “gas curtain” with a varicose cross section and diffusive boundaries. The dominant sinusoidal perturbation on this cross section has a wavelength of 6 mm and mean layer thickness of 3 mm. The peak SF₆ molar fraction across the curtain (that is, in the direction of shock propagation) varies between 40% and 60% on a shot-to-shot basis. The perturbations on each side of the layer range between 0 and 1 mm in amplitude. The precise amplitude is uncontrolled but carefully measured on each event about 100–200 μs before shock-wave impact with the SF₆ layer. These measurements of initial conditions




FIG. 1. The experimental setup. SF_6 flows vertically through the nozzle to form the initial conditions for the shock passage. The shock moves from left to right.

are essential for benchmarking codes because they enable computations to be initialized by measured, not idealized or estimated, conditions.

A Mach-1.2 shock-wave travels through the gas curtain, impulsively accelerating and compressing the two interfaces that form the thin layer. The initial perturbations on the interfaces grow and produce distinctive flow patterns. Two pulsed dye lasers illuminate a thin cross-section of the curtain. One laser produces the image of the initial conditions, taken 100–200 μs before shock impact, and the other images the dynamic condition at a preset time, typically 150–800 μs after shock impact. The light in each image is predominantly the Rayleigh-scattered light from SF_6 molecules, and the SF_6 molar fraction across the image is estimated using calibration scattering data from pure air and pure SF_6 . The background air (pre- or post-shock) is assumed to have a spatially-uniform density. Thus the laser-sheet images provide a mapping of SF_6 density, and this mapping is compared with computed density profiles.

Background light in the recorded image is primarily due to laser light scattered off windows, walls, the SF_6 nozzle, and other components of the experimental apparatus (see, for example, Fig. 5). Both background and calibration data are recorded periodically, interspersed with the actual shock-tube experiments. The laser intensity and spatial distribution vary by a few percent from shot to shot, and extraneous scattering sites (such as dust on windows) also change with time, leading to small experimental errors in the background subtraction and calibration of the data to which they are applied. The effect of these experimental errors on this work is partially addressed later in this paper.

Observations of over 100 experiments show that asymmetries in the initial conditions correlate with the three distinctive post-shock flow patterns. The range of upstream and downstream initial perturbations and the classification of the resultant dynamic patterns is shown in Fig. 2. Upstream mushrooms develop when the initial perturbations are predominantly on the upstream side of the layer, namely when

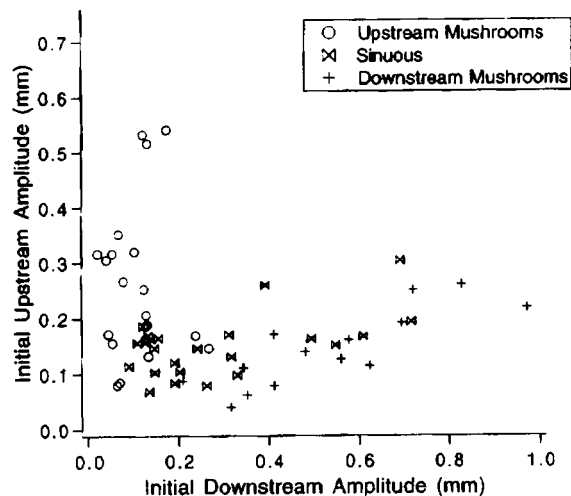


FIG. 2. The experimental data of initial amplitudes from various shots.

the initial upstream amplitude exceeds the initial downstream amplitude. Sinuous patterns develop when the perturbation amplitudes are nearly equal and when the initial downstream amplitude is slightly larger. The occurrence of downstream mushrooms correlates to predominant downstream perturbations. The images shown in Fig. 3(a) illustrate these three dynamic patterns in the data: sinuous, upstream mushrooms, and downstream mushrooms.

Experimental limitations preclude quantitative analysis beyond results reported here. Instability growth is largely two-dimensional as depicted in Fig. 7 of Ref. 10 but some images of downstream mushrooms showed diffusive regions that could be caused by some three-dimensionality of the flow. Also, measurement accuracy of shock strength for some events was limited by structure on the approaching shock front.

The simple physical interpretation of these flow patterns is based on vortex dynamics. Vorticity is produced at each interface by the interaction of the pressure gradient of the shock wave and the density gradient of the interface. The magnitude of this baroclinic vorticity production correlates with the magnitude of the perturbation amplitude. Thus the interface with the largest initial perturbation amplitude possesses the greatest post-shock vorticity, and this vorticity controls the subsequent flow. For example, the upstream mushroom is the signature of a vortex pair on the upstream side, which occurs when baroclinically generated vorticity is predominantly on the upstream interface because of the larger initial upstream perturbation. Although the basic physical mechanism is fairly simple, the flows comprise highly distorted interfaces and complex velocity distributions, which provide fluid simulations with a challenging benchmark problem.

III. SIMULATIONS

Our goals are to improve the understanding of RM instability of thin layers, and to validate and develop confidence in the AMR method and the underlying hydrodynam-

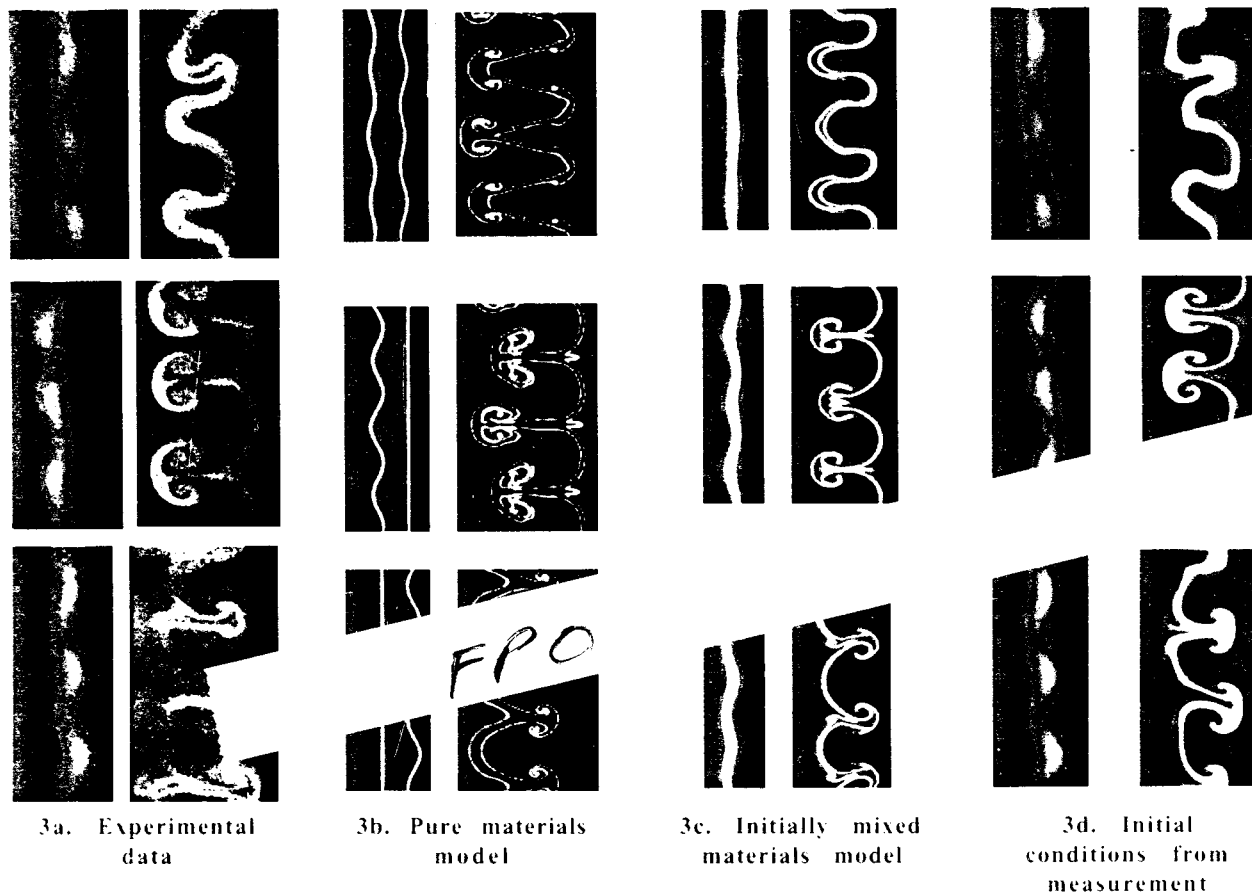


FIG. 3. The three generic dynamic flow patterns: sinuous (top), upstream (center), and downstream (bottom). Three perturbation wavelengths (18 mm) are shown. The shock moves from left to right. The density color palette goes from dark blue ($\rho=0$) to red ($\rho=0.006 \text{ g/cm}^3$). (a) The experimental data; the sinuous and upstream dynamic data is at 450 ms, while the downstream data is at 400 μs . (b) Pure materials. (c) Initially mixed materials. (d) Simulation with initial conditions as measured for each individual shot.

ics of the new code. The ability to simulate experimental data accurately in computationally-demanding physical regimes certainly adds to that confidence. We have modeled the flow patterns resulting from upstream, downstream, and

sinuous perturbations. Several models have been used to simulate these data with increasing levels of realism as described in Table I.

The physical conditions for the fluid properties were the

TABLE I. Summary of initial conditions for simulations.

Model of initial conditions	Attributes	Connection to data	Illustration
Pure materials Mach 1.2 shock	Sharp interfaces Maximum Atwood number Analytic description of amplitude and wavelength of interfaces	Qualitative only The three generic flow patterns are reproduced	Fig. 3(b)
Mixed materials Mach 1.2 shock	Diffuse interfaces Analytic description of amplitude and wavelength of interfaces	Qualitative but more representative of data	Fig. 3(c)
As measured for each shot	Minimal smoothing and conditioning Measured shock strength	Quantitative for each shot	Fig. 3(d)

same for all of the initial conditions models. The equations of state assumed that the air and SF₆ are gamma-law gases with gammas of 1.4 and 1.1, respectively. The ambient air pressure in the shock tube was set at 0.8×10^5 dynes/cm², corresponding to normal atmospheric pressure in Los Alamos. The temperature of the unshocked air and SF₆ was set to 300 K. The air density was set to 1.0 mg/cm³ and (pure) SF₆ density to 5.0 mg/cm³. The shock strength was set to Mach 1.2 for the analytic conditions.

A. RAGE code

The RAGE¹⁶ (for Radiation Adaptive Grid Eulerian) is a one-, two-, and three-dimensional multi-material Eulerian radiation-hydrodynamics code developed by SAIC (Science Applications International Corporation) and Los Alamos National Laboratory for use in solving a variety of high-deformation flow problems. This code is intended for general applications without tuning of algorithms or parameters. It avoids *ad hoc* models and attempts to rely on first-principles physics. Of particular importance to the computations in this paper is the fact that we are performing a direct numerical simulation: no explicit viscosity or sub-zonal turbulence modeling is incorporated in the code as yet. There does exist a small numerical viscosity (characteristic of all direct numerical fluid simulations), with a weak tensor viscosity added to improve rotational symmetry.

Among the features and assumptions of the RAGE code are the following:

- 1-D Cartesian, cylindrical, and spherical coordinate systems;
- 2-D Cartesian and cylindrical coordinate systems with square cells;
- 3-D Cartesian coordinate systems with cubical cells;
- adaptive mesh refinement redefinition every cycle on a cell-by-cell basis;
- all materials defined in the problem computationally present in each cell;
- pressure and temperature equilibrium for all materials in a cell;
- higher order piecewise linear Godunov numerical method;
- exact conservation of mass, momentum, and energy;
- radiation grey diffusion with nonequilibrium radiation and material temperatures.

For the purposes of this paper, the code was run in 2-D Cartesian geometry with radiation diffusion turned off. Otherwise, no special adjustments were made to simulate these experiments.

One of the key features of this code is the adaptation algorithm for the zoning. The adaptive algorithm allows cells to be subdivided or combined in each time-step cycle. When a zone is subdivided, the parent zone of the underlying mesh is preserved in background. When fine zones are combined, they always return to their original parent zone. Adjacent cells are not allowed to differ by more than a factor of two in linear dimension. At any given cycle, the computational grid can be structured to place finely subdivided cells exactly where needed and leave large cells in regions with slowly-varying physical quantities. (See Fig. 4 for an example of an




FIG. 4. The computational grid at $t=0$ for one of the pure materials initial conditions, illustrating the Adaptive Mesh Refinement (AMR) technique. The shock is moving in from the left. The interfaces and shock front are highly resolved. Zones behind the shock front have been recombined into parent cells.

adaptively-zoned grid.) The definition of where cells are needed varies from problem to problem. Most problems have had good success with adaptation algorithms that refine the cell size at material interfaces and at strong pressure and density gradients, although adapting on different variables is an option. Clearly the “art” of optimizing an AMR code is in the choice of the adaptation algorithm.

B. Results of simulations: Analytic initial conditions

The first model of the data uses generic initial conditions in the form of analytic functions describing the material distributions. We have chosen initial conditions of the SF₆ layer to be “typical” of the data without attempting to include the details of any individual experimental shot. The unperturbed thickness of the SF₆ layer was chosen to be 3 mm. The primary sinusoidal perturbation for the upstream and downstream cases had an amplitude of 0.5 mm, and for the sinusoidal case the amplitude on each side of the layer was 0.25 mm. The wavelength of the perturbation for all cases was 6 mm. A second perturbation mode, intended to provide some representative variation along the gas layer, was also imposed. This second mode had a wavelength of 12 mm and amplitude of 0.1 mm. The dimensions of the full computational grid were 24 mm parallel to the shock front (four perturbation wavelengths) and 1000 mm along the shock direction. The long dimension in the shock direction is necessary to prevent shocks from reaching the computational boundary and reflecting back to interact again with the SF₆ layer.

The largest (level-1) zone size in these problems was defined to be 1.8 mm by 1.8 mm. The AMR procedure allows cells to refine these dimensions by progressive factors of 2 down to a limit set by the user. We allowed refinement down to level-5 for these calculations, so that the finest zones were 0.1125 mm ($1.8 \text{ mm}/2^4$). At late times, approximately 50% of the total number of computational cells are at the highest resolution (level-5), although they occupy less than 1% of the area of the full computational grid. Test runs at level-6 showed no significant change in the computational details or growth rates. Each run at level-5 resolution re-

quired approximately 1 hour of single-processor Cray YMP time to reach a time of 800 μs after shock arrival at the SF_6 layer .

1. Pure material model

The first attempt to model these data used pure materials for the air and the SF_6 . That is, the initial SF_6 gas layer is pure with sharp SF_6 /air interfaces, which is different from the continuous density gradient in the actual experiment. Consequently, the vortex formation is faster and more violent in this computation than in the experimental data. Figure 3(b) shows the results of modeling initially pure materials. This faster growth of the instability can be seen best in the sinuous perturbation case, where the upstream mushrooms tend to form and roll-over at times when the experimental data show sinuous patterns. This difference should be expected because the pure-material simulations effectively have a larger Atwood number at the material interface which gives rise to a more rapidly growing instability. (The Atwood number is the difference divided by the sum of the densities across an interface.) This model gives qualitative agreement with the data, and suggests that future experiments with discontinuous interfaces can expect to observe sinuous patterns less often.

The effectiveness of the AMR zoning technique is shown in Fig. 4. The computational mesh for the region surrounding the SF_6 gas layer is shown just before the shock (moving toward the layer from the left) reaches the layer. The interface is highly resolved, as is the region around the shock front in the air to the left of the layer. The shocked air behind the front has already recombined fine zones into the coarser original parent mesh. This plot shows only the central 40 mm of the 1000 mm long computational grid. The computer memory and time savings with the AMR technique compared with that required to zone the entire problem at the finest level are clearly large.

2. Initially mixed material model

In this model the air and SF_6 are initially mixed, corresponding more closely to the experimental data. Other parameters of the model (analytic perturbation amplitudes and wavelengths) are identical to the pure-material model. The SF_6 mass fraction is assumed to be a Gaussian distribution along the shock direction, with the $1/e$ width equal to the perturbed interface width used in the pure material model. The mass fraction of SF_6 at the peak of the Gaussian is 0.9 (corresponding to a peak mole fraction of 0.6), as was typical of the experimentally measured data.⁸ This analytic model for the initial conditions, where the major change from the first model is the initial mixture of materials, presents a gentler density gradient to the incident shock, thereby producing less vorticity. The results with this model are shown in Fig. 3(c). The amplitude of the growth and the structure of the flow patterns are closer to the experimental data than for the pure-material case, especially for the sinuous case.

C. Results of simulations: Measured initial conditions

These simulations begin with the measured initial conditions and shock strengths for each event. They are the most demanding tests of the code, because they correspond closely to the experimental condition. Any difference between the simulated and the measured dynamic result that cannot be attributed to measurement errors or interpretation can only be attributed to inadequacies in the code. The success of these simulations is testimony to the accuracy of our computational technique.

The initial density profile requires minor conditioning before starting the simulation. For each event, the data from the laser shot preceding shock arrival is used to initialize the computational grid on a pixel-by-pixel basis. The effects of statistical fluctuations and steep, small-scale gradients in the data are minimized with a small amount of smoothing. Figure 5 shows the original experimental data image and a plot of the difference between corrected experimental densities and smoothed data for a slice through the curtain along the shock direction. These image files contain the SF_6 mole fraction after background subtraction and calibration information has been used to process the raw recorded image. Note the vertical band of unsubtracted background on the right side of the initial distribution, giving a spuriously long tail on the downstream side of the curtain. On the left side, the data goes slightly negative indicating an oversubtraction of background. Both of these effects must be truncated before initiating the simulation. The graph shows a single-pixel-wide slice horizontally through the image along the shock direction, before and after smoothing and truncation, that demonstrates the good fit to the region containing most of the SF_6 .

The simulation uses reflective boundary conditions so we limit the region of interest (ROI) of the initial profile by zero-derivative points. As shown in Fig. 5, the top and bottom of the ROI are located where the initial SF_6 profile is perpendicular to the ROI boundary, so that the reflective boundary conditions are approximately realized. The left and right ROI boundaries are parallel to the shock front. These ROI boundaries preserve at least three complete perturbation wavelengths (18 mm) so that the central flow patterns should be reliable even if ROI edges are slightly distorted. After conditioning and defining ROI boundaries, the measured initial condition is converted to density and air- SF_6 volume fraction, and then inserted into the calculational mesh. The shock strength for each event is determined from the experimentally-measured shock velocity determined by two pressure transducers upstream of the gas curtain.

Most of the computations were performed in the “center of mass” frame, in which the curtain is initially in motion and is brought to rest by the acceleration induced by the shock wave. This frame was used to minimize numerical advection, which effectively adds a diffusive component as the material transports through many cells of the calculational grid. Computations in the laboratory frame produced slightly more diffuse density distributions, with growth rates 5%–10% smaller than those computed in the center-of-mass frame.

Figure 3(d) shows the computed density distributions at

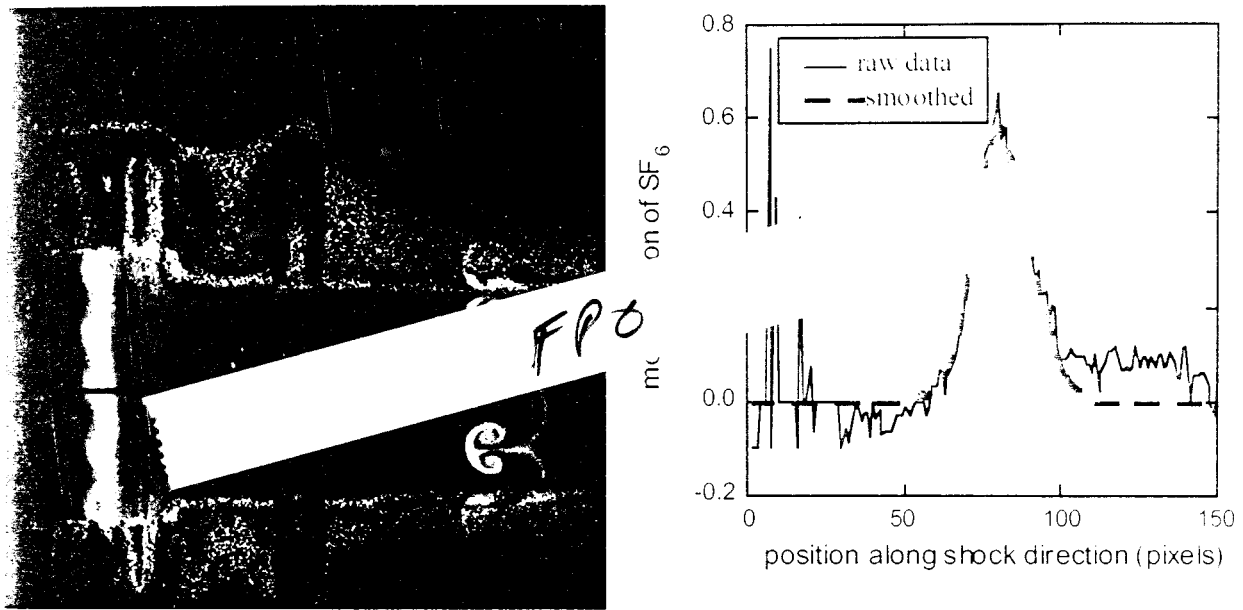


FIG. 5. An experimental data image and an example of the conditioning of the initial conditions for insertion into the code. The line on the image indicates the plotted pixels. See the text for discussion.

the same time and for the same initial conditions as the three experimental measurements shown in Fig. 3(a). The visual agreement between the data and the simulation is excellent, giving added confidence in the AMR technique as implemented in RAGE. In Fig. 6, we show quantitative one-dimensional comparisons between the data and the simulations shown in Figs. 3(a) and 3(d). The density distributions inside the boxes of both the data (left image) and simulation (right image) are projected onto the shock x -axis in the plots below the images and onto the orthogonal y -axis in the plots to the right of the images. These plots show the direct, quantitative comparison of mass flow in both directions between the data and simulation. Comparative results in Figs. 6(a)–6(e) show excellent agreement. The mass distributions in experiment and simulation are well matched. An example of an event not showing such excellent agreement is shown in Figs. 6(f)–6(g) to illustrate some potential difficulties with the experiment. This experimental image shows substantially more mass per wavelength in the time-evolved distribution compared with the initial conditions, although the mass of SF_6 within one period should remain approximately constant as long as the curtain remains two-dimensional. The simulation does not allow mass flow on the computational (reflective) boundaries and so must conserve SF_6 mass at all times. The mushrooms in the data image also appear to have a large diffuse halo not seen in the other events. The halo and mass change may indicate significant 3-D flows in this particular event.

Given the above success in reproducing the measurements, we have some confidence in using the simulation to examine details of flow dynamics not measured in these experiments. Figure 7 shows the early-time evolution of the pattern forming upstream mushrooms. A time sequence of

the density, pressure, and vorticity clearly shows the shock interaction with the perturbations in the density distribution. At $t=0$, the shock can be seen moving in from the left and just beginning to interact with the tail of the mass distribution. At $7 \mu\text{s}$, the shock is entering the main density distribution of the curtain and beginning to compress it, baroclinically generated vorticity patterns are beginning to form. At $26 \mu\text{s}$, the shock has just exited the compressed SF_6 layer. The vorticity pattern has formed predominantly on the upstream side by the predominantly upstream perturbations, and complex transverse structure in the pressure pattern has been formed by the shock interactions. Note the development of paired vortex structures (blue is negative, red is positive) predominantly on the upstream side of the SF_6 layer for this particular shot. The pressure plot shows the distortion of the shock front transmitted through the curtain, and the subsequent complicated set of signals as small rarefactions and shocks reverberate off density gradients behind the shock front. These signals continue to affect vorticity production until the transmitted shock is approximately one wavelength beyond the gas layer.

A time-sequence of the density and vorticity distributions until $600 \mu\text{s}$ for all three flow patterns is shown in Fig. 8. The fluid flow evolves into steadily more complicated and smaller scale features. As we pointed out earlier, no explicit viscosity is included in the code so that we cannot track the later flow development without adding energy dissipation and subzonal turbulence. Notice that the evolution of the vorticity keeps the two components of each vortex pair distinct and separate. The vorticity generated in the sinuous case is smaller in magnitude and is much more diffuse than the tight, well-defined vorticity patterns in the other two flow patterns.

IV. PARAMETER STUDIES

The experimental data and simulations show that the fluid flow patterns are strongly affected by rather subtle changes in the initial conditions. How well must we know the exact initial conditions to be able to predict the experimental outcome? If unmeasured aspects of the initial configuration significantly affect the flow, then our ability to predict the time evolution of that configuration will be limited.

For example, the disagreement between data and simulations shown in Figs. 6(f)–6(g) may be due to 3-D effects unmeasured in the initial configuration.

To assess sensitivity of results to experimental uncertainties in the initial conditions, we did several simulations using initial conditions distorted to the limit of experimental error bars (but still assuming that the experiment is 2-D). The events used for this analysis were carefully chosen to have

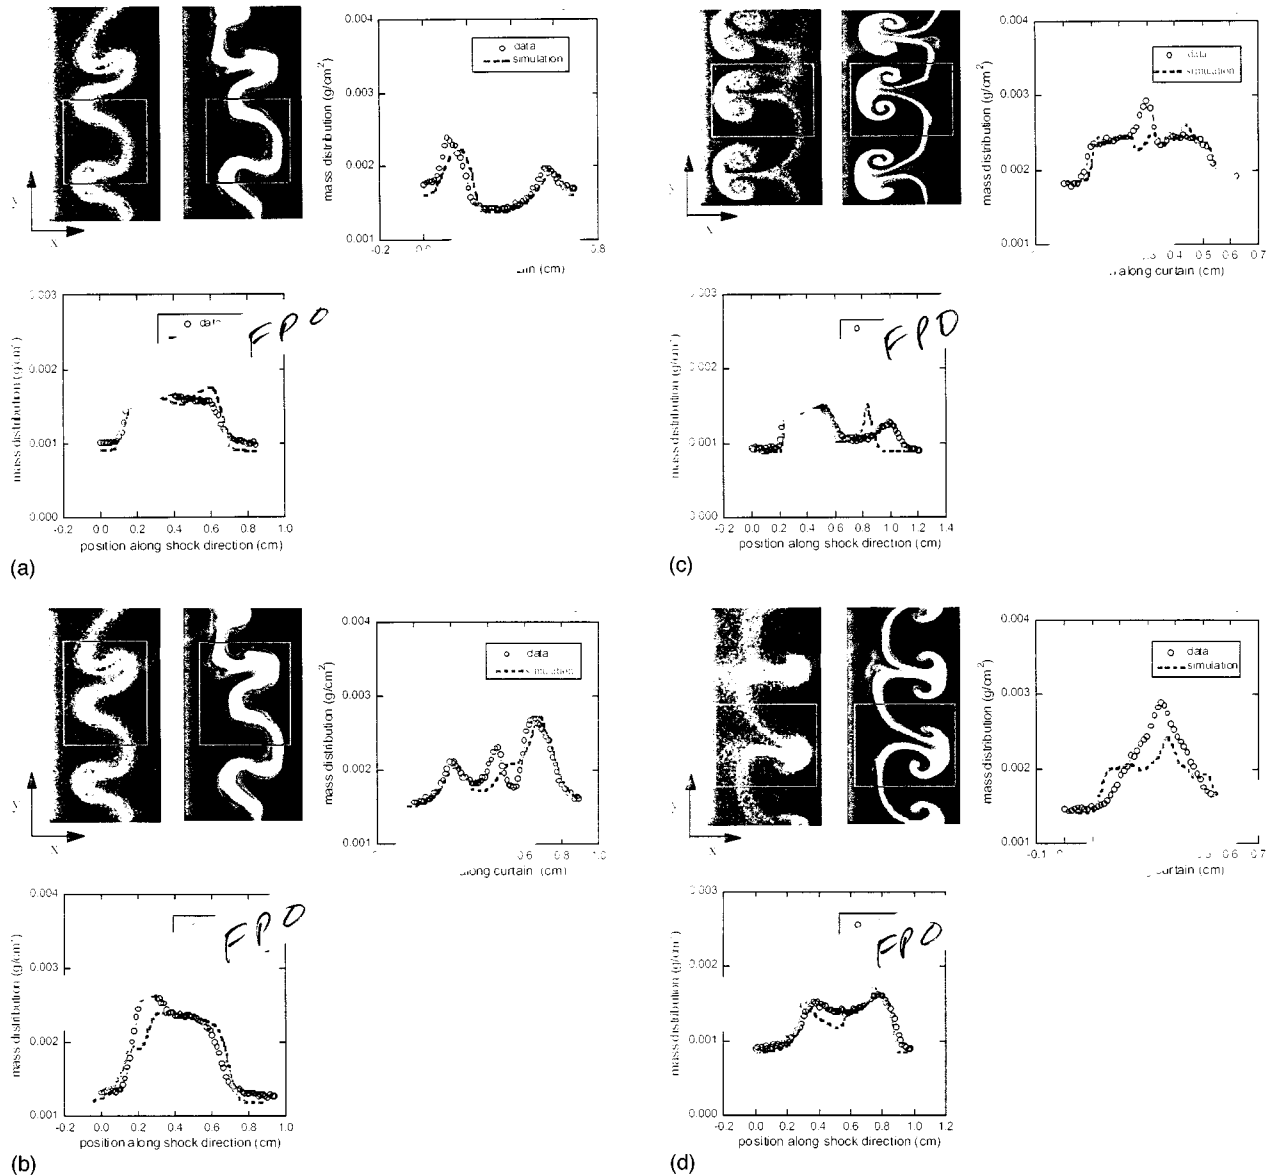








FIG. 6. (a) Shot 62. Agreement between data and simulation is excellent in both projections. (b) Shot 62. The three distinct peaks in the data in the y -projection show that the loops are aligned with the shock axis whereas in the simulation the center loop is tilted slightly. (c) Shot 197. The agreement on the mushroom cap is excellent. For the loops between mushrooms, the simulation does not flow as much mass downstream and does not diffuse as much as the data. In the y -projection, the central peak in the data is a result of the closed mushroom stem well-aligned with the shock axis. The stem in the simulation is more open and is also tilted slightly with respect to the shock axis. (d) Shot 193. The simulated growth and total mass are close to the data. As can be seen in the projection along the y -axis, the lower peak and broad wings show that the simulated stem is more open and the cap more rolled over than in the data. As with the upstream mushroom in (c), the loop between mushrooms is more diffuse in the data than in the simulation. (e) Shot 193, top mushroom. Agreement is similar to that for the bottom mushroom shown in (d). (f) Shot 190. The simulation growth in both projections is much less than in the data. The total mass in the window is also very different (the integral in either projection gives the total mass within the window). Note the large halo around the mushrooms in the data, especially when compared with the other downstream mushroom in (d) and (e). The halo and mass discrepancy may be an indication of 3-D growth in the data. See the text for discussion. (g) Shot 190, top mushroom. As in (f), the simulation growth is too small as is the total mass in the window. See the text for discussion.



(e)



(f)



(g)

FIG. 6. (Continued.)

good calibration data and background subtractions (see, for example, Fig. 5). The time-evolved, simulated density distributions showed little difference from the simulations using undistorted initial conditions. The computed flows are stable with respect to changes in the initial conditions within the experimental uncertainties. These sensitivity tests confirm the need for accurate measurements of initial conditions, and verify the adequate accuracy of the present experiments. However, our speculations about 3-D behavior to “explain” the discrepancies in Figs. 6(f)–6(g) also point out a loophole left to be closed in future experiments.

Three types of uncertainties were considered for these studies. Scaling the density by $\pm 5\%$ changed the growth rates by a few percent but otherwise did not affect the flows. Scaling the shock strength had similarly little effect on the results. Uncertainties in background subtraction and conditioning of the initial conditions were, however, more difficult to account for and seemed *a priori* to have a better chance of significantly modifying the calculations.

The first modification of the initial conditions was to add

long low-amplitude tails along the shock direction to the nominal measured pre-shock distribution. Physically, diffusion of SF_6 away from the curtain during the time before shock arrival might produce such long low-level tails, and a very slight shift in the background subtraction could hide their presence.

Figure 9(a) shows a comparison of three density slice plots from the original recorded data, the nominal initial distribution used in the simulation (showing the effects of smoothing), and the distribution modified to add tails. The simulation showed no appreciable effect due to the tails. Careful examination of the generation of vorticity shows the reason: the component of the density gradient orthogonal to the shock direction produces the initial vorticity, and tails only have small gradients parallel to the shock direction. The tails therefore do not affect the initial vortex generation that controls the evolution of the fluid flow.

A modification to the initial SF_6 distribution that directly affects the vorticity is one which changes the peak-to-valley ratio along the curtain (that is, parallel to the shock front). A

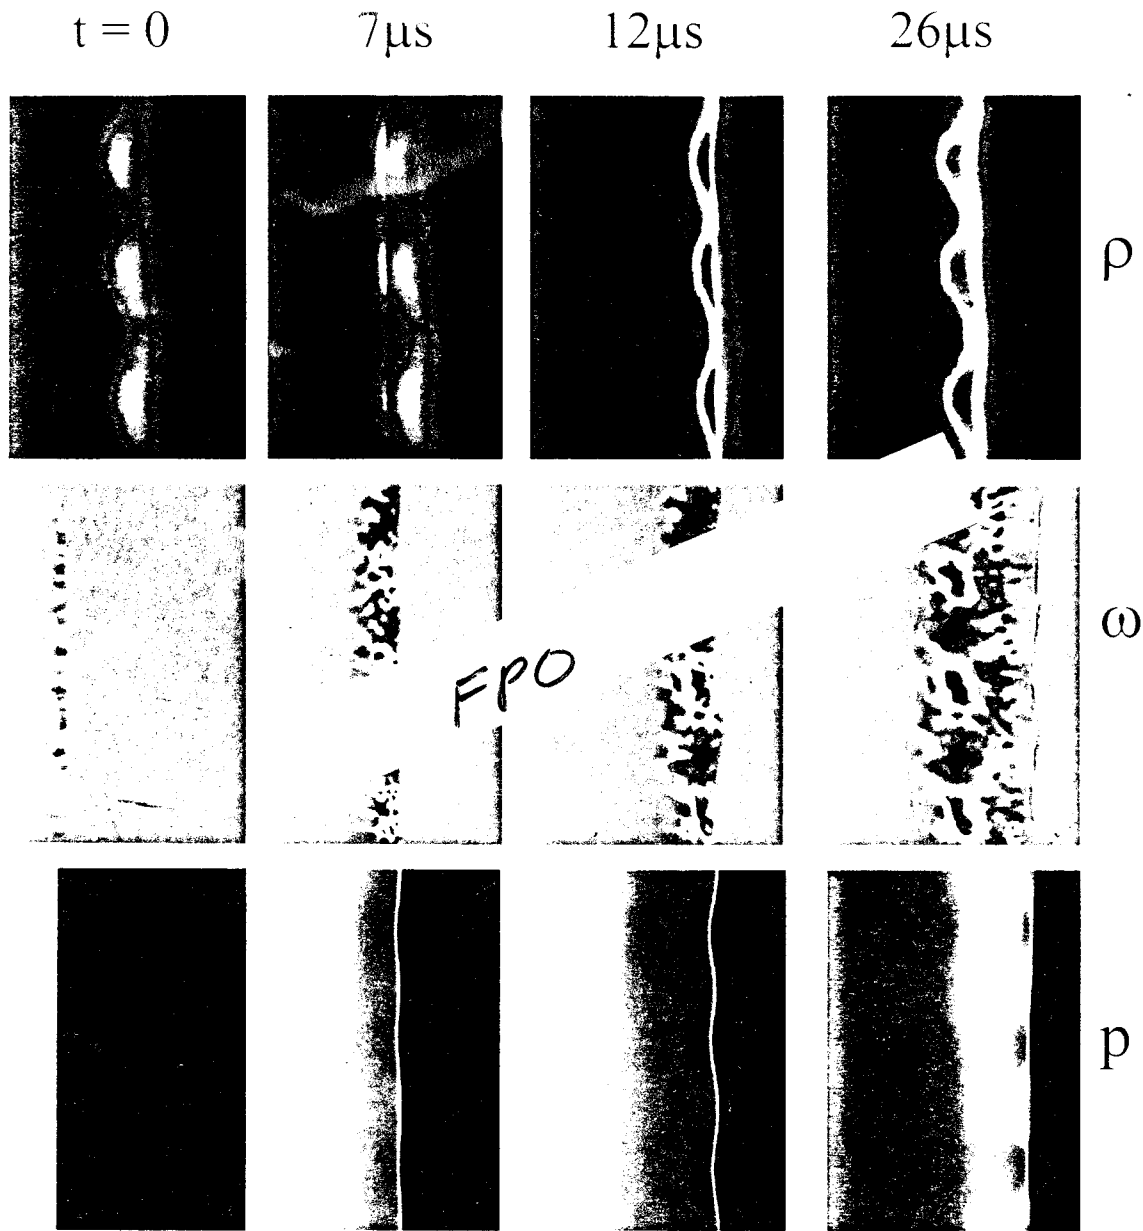


FIG. 7. Simulation of the early-time interaction of the Mach 1.2 shock with the SF_6 gas layer, for the simulation of shot 197 which results in an upstream mushroom. Density is shown in the top frames, vorticity in the center, and pressure in the bottom. The color palette goes from dark blue to red, corresponding to densities from 0 to 0.006 g/cm^3 , vorticity from $-50000/\text{s}$ to $+50000/\text{s}$, and the logarithm of the pressure from 0.7 to 1.5 bar.

small amplitude Gaussian in the shock direction was added to the measured initial conditions as above, but the amplitude of the Gaussian was modulated along the curtain to be between zero at the peak of the original distribution and maximum at the valleys. Figure 9(b) shows a slice-plot along the curtain comparing the density distribution in the original analysis and in this modification with the troughs filled in. This change in the data does affect the growth rates by 5%–10%, but is barely at the limits of a plausible experimental uncertainty.

Several other possibilities for experimental uncertainties exist but seem too speculative to address quantitatively without more information. For example, the laser measuring the

initial conditions of the curtain is fired around $100\text{--}200 \mu\text{s}$ before shock arrival. If the curtain is in motion or if put into motion by small pressure gradients in the shock tube, then not only will the initial density distribution be changed by the time of shock arrival but also the layer may have initial vorticity. Figure 10 shows the result of adding a slight overpressure or underpressure region in advance of the main shock on an initially-symmetric distribution. For simplicity, this calculation was run with analytic mass distributions and mixed materials. Nominal atmospheric pressure was 0.8 bars, and the extra region was set to 0.82 bars (top) and 0.78 bars (bottom). The main shock was 2.0 bars in both cases, close to a Mach 1.2 shock strength. Clearly, slight pressure




FIG. 8. Simulated time evolution of density and vorticity for the three analysed flow patterns. The color scale is the same as in Fig. 7. The shock can be seen moving in from the left in the $t=0$ density plots. In a time-sequence movie, the vorticity patterns rotate and weaken slightly.




FIG. 9. (a) Slice plots along the shock direction showing the effects of adding tails to the SF_6 distribution. (b) Slice plots along the SF_6 layer showing the amount of gas filling in the troughs of the layer.

gradients inside the shock tube (acoustic waves) can easily turn symmetric snake distributions into upstream mushrooms. Note that the initial condition producing the snake pattern has developed noticeable downstream perturbations by shock arrival, and the upstream mushroom case has clear upstream perturbations. The unperturbed case developed slight upstream mushrooms, and this simple model did not produce good downstream mushrooms.

Another significant uncertainty is the magnitude and effect of any three-dimensional nature of the initial distributions or subsequent flow. The experiment is designed to be mainly 2-D, but given the slightly unsteady flow creating the curtain and the intrinsically 3-D character of turbulence, the

2-D approximation must break down at some level. We speculate that the mass discrepancy and “halo” in the downstream mushroom shot 190 may be due to 3-D effects.

V. CONCLUSIONS

Vorticity production and the subsequent flow pattern of a shock-accelerated thin gas layer are very sensitive to perturbations in the initial mass distribution. Details of the compressed density profiles and the complex pressure patterns produced by shock interaction with the gas layer determine the structure and form of the nonlinear growth of the Richtmyer–Meshkov instability. Shock interactions with the gas layer continue well after the shock front has passed, in the form of multiple reflections and refractions of the initially-planar front. Vortex production is dominated by the initial encounter between the shock and density gradients in the gas layer (“baroclinic production of vorticity”). Initial density perturbations predominantly on one side of the gas layer produce vorticity predominantly on that side. The more symmetric configuration, with the perturbation amplitude roughly the same on both sides of the initial configuration, does not produce equal amplitude vortex pairs on both sides but rather produces a more diffuse, less intense vortex structure spanning the width of the gas layer. The vortex patterns persist with little change during the later nonlinear growth of the initial perturbations.

Our direct numerical simulation algorithm is doing a good job of reproducing the physics without any parameters adjusted for this particular problem. The simulation allows us to visualize and study the dynamics of the initial compression and the formation of vorticity in a graphical and controllable way.

The strengths of the AMR method as implemented in RAGE are demonstrated by our ability to resolve the fine features required to reproduce the physics accurately without




FIG. 10. A symmetric SF_6 layer is pulled upstream (top) or downstream (bottom) by a slight underpressure (top) or overpressure (bottom). The main shock hits at $t=0$.

stressing our computer resources. These features include the simulation of the collective effects of wave reverberation within the layer and the detailed generation of vorticity, both of which determine the type of dynamic structure to a large extent. We emphasize that these simulations are *quantitative* comparisons with very detailed, high-quality experimental data.

ACKNOWLEDGMENTS

We are pleased to acknowledge the contributions to the experiments by Dr. J. Jacobs of The University of Arizona. We also wish to thank John Grove and David Sharp for many fruitful discussions and comments. We especially thank Jas. Mercer-Smith for his unflagging support and encouragement. This work was performed under the auspices of the United States Department of Energy under Contract No. W-7405-ENG-36.

¹Recent examples can be found in A. Burrows, J. Hayes, and B. A. Fryxell, "On the nature of core-collapse supernova explosions," *Astrophys. J.* **450**, 830 (1995); J. M. Stone, J. Xu, and L. G. Mundy, "Formation of 'bullets' by hydrodynamical instabilities in stellar outflows," *Nature (London)* **377**, 315 (1995).

²Early work in the importance of Raleigh–Taylor instabilities in ICF implosions: R. E. Kidder, "Laser-driven compression of hollow shells: Power requirements and stability limitations," *Nucl. Fusion* **16**, 3 (1976).

³G. I. Taylor, "The instability of liquid surfaces when accelerated in a direction perpendicular to their planes. I," *Proc. Soc. London* **201**, 192 (1950).

⁴R. D. Richtmyer, "Taylor instability in shock acceleration of compressible fluids," *Commun. Pre Appl. Math.* **13**, 297 (1960).

⁵E. E. Meshkov, "Instability of the interface of two gases accelerated by a shock wave," *Fluid Dyn.* **4**, 101 (1969).

⁶Recent studies of these instabilities are documented in these workshop proceedings and others in the same series: P. F. Linden, D. L. Youngs, and S. B. Dalziel (editors), "Proceedings of the 4th International Workshop on the Physics of Compressible Turbulent Mixing," Cambridge, England (29 March–1 April 1993).

⁷J. W. Shaner, "Pattern formation by shock processes," *Physica D* **12**, 154 (1984).

⁸J. M. Budzinski, R. F. Benjamin, and J. W. Jacobs, "Influence of initial conditions on the flow patterns of a shock-accelerated thin fluid layer," *Phys. Fluids* **6**, 3510 (1994).

⁹J. W. Jacobs, D. L. Klein, D. G. Jenkins, and R. F. Benjamin, "Instability growth patterns of a shock-accelerated thin fluid layer," *Phys. Rev. Lett.* **70**, 583 (1993).

¹⁰J. W. Jacobs, D. G. Jenkins, D. L. Klein, and R. F. Benjamin, "Nonlinear growth of the shock-accelerated instability of a thin fluid layer," *J. Fluid Mech.* **295**, 23 (1995).

¹¹K. A. Meyer and P. J. Blewett, "Numerical investigation of the stability of a shock accelerated interface between two fluids," *Phys. Fluids* **15**, 753 (1972).

¹²R. F. Benjamin, D. Besnard, and J. Haas, "Shock and reshock of an unstable interface," LANL Report LA-UR 92-1185, Los Alamos National Laboratory (1993).

¹³J. W. Grove, R. Holmes, D. H. Sharp, Y. Yang, and Q. Zhang, "Quantitative theory of Richtmyer–Meshkov instability," *Phys. Rev. Lett.* **71**, 3473 (1993).

¹⁴R. L. Holmes, J. W. Grove, and D. H. Sharp, "Numerical investigation of Richtmyer–Meshkov instability using front tracking," *J. Fluid Mech.* **301**, 51 (1995).

¹⁵K. O. Mikaelian, "Rayleigh–Taylor and Richtmyer–Meshkov instabilities in finite-thickness fluid layers," *Phys. Fluids* **7**, 888 (1995).

¹⁶M. L. Gittings, "SAIC's Adaptive grid Eulerian Hydrocode," Defense Nuclear Agency Numerical Methods Symposium, 28–30 April 1992.

¹⁷N. Byrne, T. Betlach, and M. L. Gittings, "RAGE: A 2D Adaptive Grid Eulerian Nonequilibrium Radiation Code," Defense Nuclear Agency Numerical Methods Symposium, 28–30 April 1992.