On the dynamics of a shock-bubble interaction

James J. Quirk[1]

Institute for Computer Applications in Science and Engineering

NASA Langley Research Center

Hampton, VA 23681, USA.

and

Smadar Karni[1,2]

Courant Institute of Mathematical Sciences

New York University

251 Mercer St.,New York, NY 10012

Abstract

We present a detailed numerical study of the interaction of a weak shock wave with an isolated cylindrical gas inhomogeneity. Such interactions have been studied experimentally in an attempt to elucidate the mechanisms whereby shock waves propagating through random media enhance mixing. Our study concentrates on the early phases of the interaction process which are dominated by repeated refractions and reflections of acoustic fronts at the bubble interface. Specifically, we have reproduced two of the experiments performed by Haas and Sturtevant: a $M_S = 1.22$ planar shock wave, moving through air, impinges on a cylindrical bubble which contains either helium or Refrigerant 22. These flows are modelled using the two-dimensional, compressible Euler equations for a two component fluid (air-helium or air-Refrigerant 22). Although simulations of shock wave phenomena are now fairly commonplace, they are mostly restricted to single component flows. Unfortunately, multi-component extensions of successful single component schemes often suffer from spurious oscillations which are generated at material interfaces. Here we avoid such problems by employing a novel, nonconservative shock-capturing scheme. In addition, we have utilized a sophisticated adaptive mesh refinement algorithm which enables extremely high resolution simulations to be performed relatively cheaply. Thus we have been able to reproduce numerically all the intricate mechanisms that were observed experimentally (*e.g.* transition from regular to irregular refraction, cusp formation and shock wave focusing, multi-shock and Mach shock structures, jet formation etc.), and we can now present an updated description for the dynamics of a shock-bubble interaction.

i

1 Introduction

In an extremely lucid paper, Haas and Sturtevant (1987) presented an experimental study of the interaction of weak shock waves with *isolated* inhomogeneities that took the form of either spherical or cylindrical bubbles. They argued that idealized experiments were necessary to shed light on the complex phenomenological behaviour whereby shock waves propagating through random media can alter the structure of fluid inhomogeneities. To this end, their experiments were a resounding success. A number of shadowgraphs were produced which provide much insight into several important mechanisms such as transition from regular to irregular refraction, folding processes, shock wave focusing, distortion of the bubble interface and vortex formation. However, given the nature of the experiments, certain subtleties of the interaction process were inevitably left unexplained. Note that such experiments are extremely difficult to control since they are conducted under imperfect conditions. For example, diffusion occurs across the membrane that forms the bubble interface and so the precise properties of the gas inhomogeneity are not known (Abd-El-Fattah & Henderson 1978a,b). Moreover, when the shock impinges on the bubble the membrane does not always rupture cleanly and it can interfere with the flow (Rupert 1992), as does the support structure needed to hold the bubble in place. There are also difficulties with the repeatability of the experiment due to variations in the incident shock strength (Haas & Sturtevant 1987), and problems with the interpretation of the flow visualization images due to unwanted three-dimensional effects (Wang & Widhopf 1990). Other problems arise in that certain quantities of interest cannot be measured directly either because of their intrinsic nature (*e.g.* vorticity) or because of practical limitations in the experimental setup.

Given this background, the purpose of the present study was to explore the extent to which a modern computational method could complement the experiments of Haas and Sturtevant (1987) in elucidating the basic mechanisms that govern the propagation of shocks through nonuniform gases. Additionally, it was thought that such a study could help bridge the gap between existing theories of shock reflection-refraction phenomena and experiment. For example, although Haas & Sturtevant were able to use the theory of geometrical acoustics to gain a good understanding of their experimental observations, this theory ignores wave nonlinearities and so it fails to account for all flow features. Similarly, while von Neumann theory (1963) is exact, within its assumptions, it cannot cope with regions of non-uniform flow and therefore it is not strictly applicable to shock interactions at curved interfaces (Ben-Dor & Takayama 1985). On the other hand, Whitham's theory (1957, 1958) and its extensions (Catherasoo & Sturtevant 1983, Schwendeman 1988) can cope with curved interfaces, but the theory is approximate and it does not take proper account of reflected waves. Moreover, this approach cannot provide any details for the flow structure behind the incident and refracted shock fronts. Hence the need for direct numerical simulations – simulations provide a controlled environment in which to isolate genuine flow phenomenology from experimental artifacts and they can provide global details of the flow dynamics to supplement the idealized, local descriptions provided by theory. Indeed, for the case of shock refraction at a planar interface, Henderson *et al.*(1991) have already demonstrated that careful simulations can be used to improve the classification of reflection-refraction phenomena which arose from experiment and analysis (Abd-El-Fattah & Henderson 1978a,b). Here we aim to shed more light on the refraction process at a curved interface.

Haas & Sturtevant's experiments have already inspired several numerical studies. For example,

both Picone & Boris (1988) and Yang *et al.*(1993, 1994) have performed computations aimed at determining the long-time evolution of the bubble inhomogeneities, while Löhner *et al.* (1987) have investigated the early-time dynamics of the interaction process. However, in these earlier studies the flow was modelled using a single gas rather than the exact binary system used by the experiment. This simplification, whilst expedient, inevitably reduced the accuracy of the results. Admittedly, for the cases studied here the errors so introduced are not catastrophic and, to some extent, can be tolerated. But these circumstances are fortuitous. Note that since some desired density jump must be imposed across the bubble interface, with a single gas component model the bubble cannot be in thermal equilibrium with its surroundings, as was the case with the experiments. Indeed, the error in temperature will be arbitrarily large, dependent on the density ratio (for the air-helium case studied here the temperature within the bubble would be 2.08 times too high). Now one of the motivations for studying shock-bubble interactions is to investigate mechanisms whereby air and fuel can be mixed efficiently in the short transit times available with supersonic combustion systems (Marble *et al.* 1987). Clearly, in such circumstances, gross errors in the temperature field could not be tolerated. In addition to the shortcomings of the flow model, these previous studies are under-resolved and are therefore prone to misinterpretation.

Our computational study avoids both of the above shortcomings. First, proper account is taken of the separate gas components; the flow is modelled by the compressible Euler equations for a two-component fluid (air & helium or air & Refrigerant 22 (R22) depending upon the experiment being simulated). Although this represents but a small generalization over the single component case, most popular shock-capturing schemes do not perform satisfactorily for multi-component flows in that they produce spurious oscillations at material interfaces (*e.g.* Abgrall 1988). Since such numerical artifacts can have a significant affect upon the evolution of a material interface, they are to be avoided. Here we employ a somewhat novel scheme to avoid this numerical difficulty (Karni 1994*a*). In essence, the scheme allows for a controlled conservation error so as to maintain the correct pressure equilibrium between different fluid components. While this relaxation of strict conservation runs against perceived wisdom in the design of numerical schemes for flows with shock waves (Lax 1954, 1972), it does produce excellent results. Second, we overcome the shortcoming of poor resolution by utilizing a sophisticated adaptive mesh refinement scheme (Quirk 1991). This scheme can reduce by several hundred-fold the cost of performing detailed simulations and so it allows for simulations that would otherwise prove to be prohibitively expensive.

As will be shown in this paper, our computational machinery provides a means of producing simulations which agree remarkably well with experiment. Since much of this machinery is general purpose and could be profitably used to investigate other quite different phenomena, we describe it in some detail (although the minutiae are necessarily skipped).

The organization for the rest of this paper is as follows. In the next section we present the compressible Euler equations for a two-component fluid and we demonstrate that schemes which are routinely applied in the single-component case do not work satisfactorily in this generalised case. In §3 we describe the major components of our numerical method for simulating multi-component flows, then in §4 we detail the computational set-up for the experiments that we have simulated. This is followed by four sections of results and discussion. First, in §5 we present a qualitative comparison against experiment concentrating on flow visualization. Then we present two quantitative comparisons with experiment, one section deals with the velocities of certain key

flow features, the other deals with the measurement of pressure traces at various locations along the axis of flow symmetry. These comparisons are followed by a discussion on the production of vorticity resulting from the passage of the shock through the bubble inhomogeneity. Although this discussion goes beyond the main purpose of this paper it is pertinent to several recent studies aimed at determining the long time evolution of the bubble. Finally, in §9 we close with some general remarks concerning our numerical study.

2 Multicomponent Flows

We model multicomponent flows using the compressible Euler equations augmented by a requisite number of species equations. For clarity, in this paper we focus on flows with only two components and so we employ just a single species equation; the extension of our discussion to several components follows straightforwardly.

In two-dimensions, using Cartesian coordinates (x, y), the governing equations may be written in conservation form

$$\mathbf{W}_t + \mathbf{F}(\mathbf{W})_x + \mathbf{G}(\mathbf{W})_y = 0$$

$$\mathbf{W} = \begin{pmatrix} \rho \\ \rho u \\ \rho v \\ E \\ \rho Y \end{pmatrix} ; \quad \mathbf{F}(\mathbf{W}) = \begin{pmatrix} \rho u \\ \rho u^2 + p \\ \rho u v \\ \rho u(E + p) \\ \rho u Y \end{pmatrix} ; \quad \mathbf{G}(\mathbf{W}) = \begin{pmatrix} \rho v \\ \rho v u \\ \rho v^2 + p \\ \rho v(E + p) \\ \rho v Y \end{pmatrix} . \tag{1}$$

Note that these equations are written in mixture form, ρ is the density of a binary mixture whose mass fractions are Y for component one and $1 - Y$ for component two. It is assumed that both fluid components are in pressure equilibrium and that they move with a single velocity whose components are u and v in the x and y directions, respectively. This assumption of no velocity slip is reasonable only if the density variation between components is moderate as is generally the case with two gases. Here, E is the total energy of the mixture per unit volume. Both fluid components are taken to be perfect gases, with ratios of specific heat $\gamma_1 = Cp_1/Cv_1$ and $\gamma_2 = Cp_2/Cv_2$. Therefore, the pressure, p, is given by

$$p = (\gamma(Y) - 1)(E - \frac{1}{2}\rho u^2 - \frac{1}{2}\rho v^2) \tag{2}$$

where the effective γ for the mixture depends on the species concentration, Y, and is found from standard thermodynamic reasoning to be

$$\gamma(Y) = \frac{\gamma_1 Cv_1 Y + \gamma_2 Cv_2(1 - Y)}{Cv_1 Y + Cv_2(1 - Y)} . \tag{3}$$

It is well known that solutions to (1) may develop discontinuous shock fronts, across which the governing equations are no longer valid in their differential form. Using Gauss's divergence theorem, equation (1) may be recast into an integral form which remains valid at a shock

$$\frac{\partial}{\partial t} \iint_D \mathbf{W} dx dy + \oint_{\partial D} \mathbf{F} dy - \mathbf{G} dx = 0 \tag{4}$$

and which can be used to deduce the Rankine-Hugoniot jump conditions (Courant & Friedrichs 1948). Numerically, (4) may be discretized to produce a so-called conservative shock-capturing scheme in which numerical approximations to the flux vectors $\mathbf{F}$ and $\mathbf{G}$ are used to evolve the field solution. Over recent years a whole new industry has grown up around the problem of how best to compute approximations to these numerical fluxes (Roe 1986). Irrespective of the flux formulation, however, a shock-capturing scheme necessarily results in a 'viscous' shock profile which is smeared rather than a perfect discontinuity (unless the discontinuity coincides with a cell interface). But it can be shown that a conservative discretization ensures that a numerically captured shock, although artificially smeared, has both the correct strength and speed; conversely, a nonconservative discretization may give physically inconsistent solutions (Lax 1954,1972; Hou & Le Floch 1991).

Given this fundamental property, a conservative formulation is almost universally accepted as the starting point for devising a shock-capturing scheme, and to date many successful schemes have been so developed for single-component flows. However, a major obstacle in extending conservative schemes to multi-component flows is ensuring that the correct pressure equilibrium is maintained between fluid components across a diffused material interface (*e.g.* Abgrall 1988; Colella *et al.* 1989; Larrouturou 1991; Ton *et al.* 1991; Karni 1994*a*; Bell *et al.* 1994). This numerical difficulty is illustrated in Figure 1. Even though each of the conserved variables might remain monotone across the smeared interface, the pressure that corresponds to the artificial intermediate state differs from the equilibrium pressure. Once generated such erroneous pressure fluctuations can propagate and contaminate the solution field. For example, Figure 2 (a) shows a snapshot from a one-dimensional, conservative computation of a shock-bubble interaction where the start data is identical to the air-helium case given in §4. Here the initial position of the bubble is marked by the vertical lines and the computed pressure field is shown some time after the shock has passed through the bubble and several reflections and refractions have taken place. Spurious pressure oscillations are clearly visible. For other sets of reasonable data, such oscillations get even larger. Now since material interfaces can be physically unstable, even slight numerical perturbations can trigger completely incorrect interfacial behaviour (Karni 1994*b*) and are therefore to be avoided. Note that in a reactive system such pressure perturbations could significantly alter the local release of chemical energy and so might compound the error.
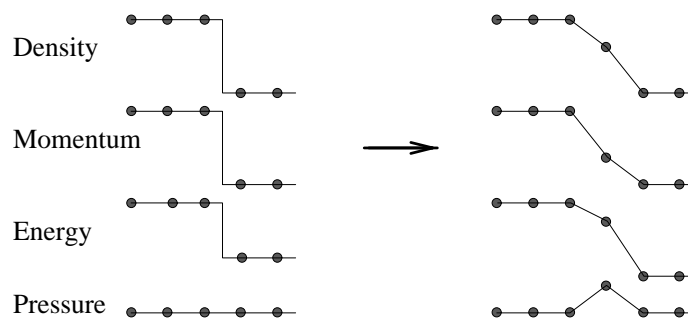

Figure 1: Pressure fluctuation at a material interface due to numerical diffusion.

Numerical problems with smeared interfaces can be avoided if fronts are fitted rather than
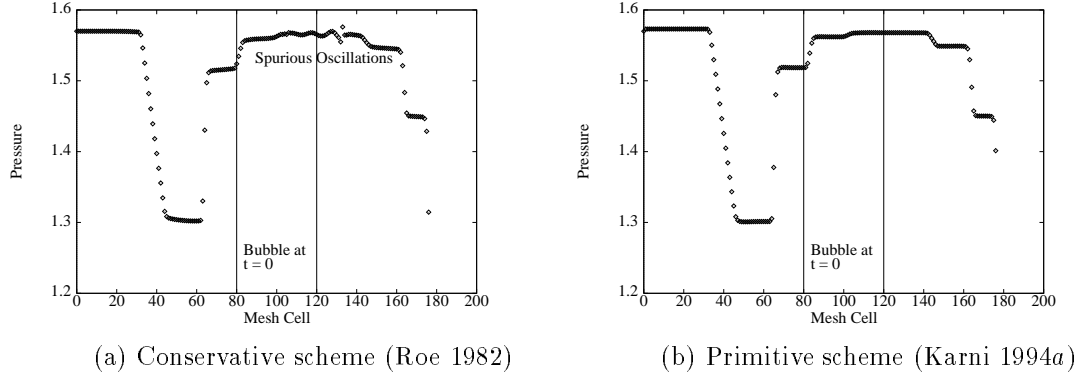
(a) Conservative scheme (Roe 1982)   (b) Primitive scheme (Karni 1994a)

Figure 2: Pressure profiles for a one dimensional 'shock-bubble' interaction.

captured (e.g. Grove & Meinkoff 1990), but front fitting introduces its own difficulties. Here we adopt an alternative approach and consider the governing equations in so-called primitive form. The Euler system (1) in primitive form is given by

$$\mathbf{U}_t + \mathbf{A}^p(\mathbf{U})\mathbf{U}_x + \mathbf{B}^p(\mathbf{U})\mathbf{U}_y = 0$$

$$
\mathbf{U} = \begin{pmatrix} \rho \\ u \\ v \\ p \\ Y \end{pmatrix} ; \quad
\mathbf{A}^p(\mathbf{U}) = \begin{pmatrix} u & \rho & 0 & 0 & 0 \\ 0 & u & 0 & \rho^{-1} & 0 \\ 0 & 0 & u & 0 & 0 \\ 0 & \gamma p & 0 & u & 0 \\ 0 & 0 & 0 & 0 & u \end{pmatrix} ; \quad
\mathbf{B}^p(\mathbf{U}) = \begin{pmatrix} v & 0 & \rho & 0 & 0 \\ 0 & v & 0 & 0 & 0 \\ 0 & 0 & v & \rho^{-1} & 0 \\ 0 & 0 & \gamma p & v & 0 \\ 0 & 0 & 0 & 0 & v \end{pmatrix} . \tag{5}
$$

To see the advantages of this formulation, consider a planar material interface aligned in the x direction with data such that $\frac{\partial}{\partial y} \equiv 0$. Across the interface, both the pressure, p, and the component of velocity normal to the interface, u, are constant. It follows that locally the primitive system reduces to three completely decoupled *linear* advection equations in ρ, v and Y and that both p and u remain constant. Thus, any reasonable discretization of (5) can be expected to produce solutions which are free of oscillations near the material interface, without introducing conservation errors.

Conservation errors, however, will occur near shocks and unless some measure is taken to control them, a primitive variable formulation will prove inadequate. Building on an idea first proposed by Zwas & Roseman (1973), Karni (1992) has developed a set of high-order correction terms which can be used to remove leading order conservation errors so as to produce a 'nearly' conservative, primitive variable scheme. This novel scheme rests on two observations: (i) Numerically captured shocks have 'viscous' profiles which are determined by the truncation error of the discretization scheme. (ii) A conservative discretization produces a consistent 'viscous' shock profile in the sense that a captured shock has both the correct strength and speed. In essence, the present primitive variable scheme employs correction terms so as to mimic the 'viscous' shock profile of a conservative scheme. In the next section, we outline the derivation of this scheme. But first, we demonstrate that for the one-dimensional shock-bubble problem it produces oscillation free solutions, cf. Figure 2 (a) and (b).

5

3 Numerical Method

We now describe the major components of our numerical method for investigating the dynamics of a shock-bubble interaction. These are: (i) the primitive variable discretization – this provides a sound basis for the integration of multicomponent flows; (ii) the parallel, adaptive mesh refinement implementation – this is essential in order to resolve intricate flow features, whilst maintaining low computational costs; (iii) graphical flow visualization – this facilitates the process of elucidating the phenomena under investigation.

3.1 A Non-Conservative Shock-Capturing Scheme

Following Strang (1968), we employ dimensional splitting to integrate the system (5) with the refinement that correction terms are applied to the right hand side (RHS) of each split equation so as to control conservation errors. Thus, we alternate between integrating

$$\mathbf{U}_t + \mathbf{A}^p(\mathbf{U})\mathbf{U}_x = \frac{\Delta t}{2}\mathbf{D_x} \qquad \text{and} \qquad \mathbf{U}_t + \mathbf{B}^p(\mathbf{U})\mathbf{U}_y = \frac{\Delta t}{2}\mathbf{D_y}. \tag{6}$$

The precise form of the correction terms, $\mathbf{D_X}$ and $\mathbf{D_y}$, depends upon the discretization for the left hand side (LHS) of each split equation. We shall now derive the correction terms, $\mathbf{D_X}$, assuming that the LHS of the x-sweep operator has been discretized using Roe's first-order upwind scheme (Roe 1982). In essence, this is done by comparing the x-sweep discretization for the primitive system (5) with the analogous x-sweep discretization for the conservative system (1).

If Roe's scheme is used to solve (1), the scheme is a first-order approximation to (1) but it is a second-order approximation to the *equivalent equation*

$$\mathbf{W}_t + \mathbf{F}(\mathbf{W})_x = \mathbf{W}_t + \mathbf{A}^c(\mathbf{W})\mathbf{W}_x = \frac{\Delta t}{2}\left(\frac{(|\mathbf{A}^c|\mathbf{W}_x)_x}{\lambda} - \mathbf{W}_{tt}\right) \tag{7}$$

where $\lambda = \Delta t/\Delta x$ is the ratio of the time step and mesh size used for the integration and $\mathbf{A}^c$ is the Jacobian matrix $\frac{\partial \mathbf{F}(\mathbf{W})}{\partial \mathbf{W}}$. The RHS of (7) constitutes the leading order terms in the truncation error of the scheme. To leading order, these dissipative terms determine the viscous path across the numerical shock transition. In this case, the numerical viscous path is physically consistent since it is produced by a conservative scheme and so it produces correct shock speeds and jumps.

Similarly, if Roe's upwind scheme is applied to solve equation (5), the scheme is a second-order approximation to the *equivalent equation*

$$\mathbf{U}_t + \mathbf{A}^p(\mathbf{U})\mathbf{U}_x = \frac{\Delta t}{2}\left(\frac{(|\mathbf{A}^p|\mathbf{U}_x)_x}{\lambda} - \mathbf{U}_{tt}\right). \tag{8}$$

In general, the two viscous forms (7) and (8) are different. The former, arising from a conservative discretization, yields shocks that satisfy the Rankine-Hugoniot conditions the latter does not. To enforce consistent shock profiles on the primitive solution, the difference between the two viscous expressions (appropriately transformed) has to be added to the RHS of the x-sweep operator for the primitive system to give (6) where

$$\mathbf{D_X} = \left\{\mathbf{T}\left(\frac{(|\mathbf{A}^c|\mathbf{W}_x)_x}{\lambda} - \mathbf{W}_{tt}\right) - \left(\frac{(|\mathbf{A}^p|\mathbf{U}_x)_x}{\lambda} - \mathbf{U}_{tt}\right)\right\} \tag{9}$$

and $\mathbf{T}$ is the conservative to primitive transformation matrix $\frac{\partial \mathbf{U}}{\partial \mathbf{W}}$.

If (6) is solved by Roe's upwind scheme, with its RHS (9) appropriately discretized, the solution procedure is conservative to the order of the numerical approximation. The correction terms may be written entirely in terms of the primitive variables

$$\mathbf{D_X} = \frac{\mathbf{T(T^{-1})}_x |\mathbf{A}^p| \mathbf{U}_x}{\lambda} - \mathbf{T(T^{-1})}_t \mathbf{U}_t. \tag{10}$$

Straightforward algebra shows that for the extended Euler system (5), the correction terms are given by

$$\mathbf{D_X} = \begin{pmatrix} 0 \\ \dfrac{1}{2\rho}\left(\dfrac{c_1 \rho_x u_x + \frac{1}{a^2}c_2 u_x p_x + c_4\left(\frac{\rho}{a}u_x^2 + \frac{1}{\rho a}p_x \rho_x\right)}{\lambda_x} - 4\rho_t u_t \right) \\ \dfrac{1}{2\rho}\left(\dfrac{\frac{1}{a^2}c_2 p_x v_x + 4|u|\rho_x v_x + c_4\frac{\rho}{a}u_x v_x}{\lambda_x} - 4\rho_t v_t \right) \\ \dfrac{\gamma-1}{2}\left(\dfrac{c_3 \rho u_x^2 + \frac{1}{a}c_4 p_x u_x + 2|u|\rho v_x^2}{\lambda_x} - 2\rho\left(u_t^2 + v_t^2\right) \right) \\ 0 \end{pmatrix} \tag{11}$$

where

$$
\begin{aligned}
c_1 &= |u-a| + 2|u| + |u+a|, & c_3 &= |u-a| + |u+a|, \\
c_2 &= |u-a| - 2|u| + |u+a|, & c_4 &= |u-a| - |u+a|
\end{aligned}
$$

and a is the sound speed.

The following observations can be made:

1. If (5) is used to replace time derivatives by space derivatives, all terms within $\mathbf{D_X}$ are scaled by either one or both of u_x and p_x, hence $\mathbf{D_X}$ vanishes near contact surfaces. Consequently, the correction terms, although derived for first-order upwinding, may also be used for second-order upwinding without degrading the latter's superior resolution of contact surfaces (Karni 1992, 1994a). Note that such schemes often reduce to first-order accuracy near shocks anyway, which is precisely where the correction terms come into play.

2. The correction terms are derived using asymptotic arguments based on the scheme truncation error. Conservation errors, while significantly reduced, are inherent to the method (Hou & Le Floch 1991, Karni 1992) and are not completely eliminated.

3. The correction terms depend on the ratio $\lambda = \dfrac{\Delta t}{\Delta x}$, and so some variation in their effect is to be expected with changes in Courant number ($wave\ speed * \lambda$). It is our experience that the correction terms work best at Courant numbers close to one, which is the upper bound on the size of time step for the integration process to be stable.

The correction terms for the y-sweep operator (6) may be similarly derived and are given by

$$
\mathbf{Dy} = \begin{pmatrix}
0 \\[2mm]
\dfrac{1}{2\rho} \left(\dfrac{\frac{1}{a^2}c_2 p_y u_y + 4|v|\rho_y u_y + c_4 \frac{\rho}{a} u_y v_y}{\lambda_y} - 4\rho_t u_t \right) \\[4mm]
\dfrac{1}{2\rho} \left(\dfrac{c_1 \rho_y v_y + \frac{1}{a^2}c_2 v_y p_y + c_4 \left(\frac{\rho}{a} v_y^2 + \frac{1}{\rho a} p_y \rho_y \right)}{\lambda_y} - 4\rho_t v_t \right) \\[4mm]
\dfrac{\gamma - 1}{2} \left(\dfrac{c_3 \rho v_y^2 + \frac{1}{a}c_4 p_y v_y + 2|v|\rho u_y^2}{\lambda_y} - 2\rho \left(u_t^2 + v_t^2 \right) \right) \\[4mm]
0
\end{pmatrix}
\qquad (12)
$$

where the coefficients c_1–c_4 are the same as those in (11) but with u replaced by v.

Given the derivation of the correction terms, our basic method of flow integration is as follows. The LHS for each split equation (6) is discretized using a second-order Roe scheme cast in *fluctuation-and-signal* form (Roe 1982). For the correction terms, temporal derivatives are replaced by spatial derivatives which are then centrally differenced, and pointwise values take the cell-centred values used by Roe's scheme. The correction terms contribute to cell-updates via a forward Euler time integration. Thus the x-sweep operator of our 'nearly' conservative primitive variable scheme takes the form

$$
\mathbf{U}_i^{n+1} = \mathcal{L}_x^{Roe}(\mathbf{U}_i^n) + \Delta t \left[\frac{\Delta t}{2} \mathbf{Dx}(\mathbf{U}_i^n) \right]
$$

where $\mathcal{L}_x^{Roe}$ is the standard Roe evolution operator. The y-sweep operator follows by analogy, and the two operators are alternated so as to arrive at a two-dimensional scheme (Strang 1968).

3.2   The AMR Algorithm – An Overview

The Adaptive Mesh Refinement (AMR) algorithm is a general purpose scheme for integrating systems of hyperbolic partial differential equations. It attempts to reduce the costs of integration by matching the local resolution of the computational grid to the local requirements of the solution being sought. For example, in simulations of gas dynamic flows, a fine mesh is generally used only in the vicinity of shock waves and other flow discontinuities, elsewhere a relatively coarse mesh is employed. Although the computational savings which accrue from local mesh refinement are totally problem dependent, they are often significant; savings of more than five hundred-fold have been obtained for simulations of detonation phenomena (Quirk & Hanebutte 1993). The foundations of the AMR algorithm lie with the work of Berger & Oliger (1984), but the derivative outlined here is due to Quirk (1991).

The AMR algorithm employs a hierarchical grid system. In the following, the term 'mesh' refers to a single topologically rectangular patch of cells and the term 'grid' refers to a collection of such patches. At the bottom of the hierarchy a set of coarse mesh patches delineates the computational domain. These patches form the grid G_0 and they are restricted such that there is continuity of grid lines between neighbouring patches. This domain may be refined locally by embedding finer mesh patches into the coarse grid G_0. These embedded patches form the next grid in the hierarchy, G_1. Each embedded patch is effectively formed by subdividing the coarse cells of the patches that it overlaps. The choice for the refinement ratio is arbitrary, but it must be the same for all the embedded patches. Thus, by construction, the grid G_1 also has continuity of grid lines. This process of adding grid tiers to effect local refinement may be repeated as often as desired, see Figure 3.

8

From stability considerations, many numerical schemes have a restriction on the size of time step that may be used to integrate a system of equations. The finer the mesh, the smaller the allowable time step. Consequently, the AMR algorithm refines in time as well as space. More but smaller time steps are taken on fine grids than on coarse grids in a fashion which ensures that the rate at which waves move relative to the mesh (the Courant number) is comparable for all grid levels. This avoids the undesirable situation where coarse grids are integrated at very small Courant numbers given the time step set by the finest grid's stability constraints: some schemes (*e.g.* Lax-Wendroff) give poor accuracy for small Courant numbers.

The field solution on each grid is retained even in regions of grid overlap and so all grid levels in the hierarchy coexist. The order of integration is always from coarse to fine since it is necessary to interpolate a coarse grid solution in both time and space to provide boundary conditions for its overlying fine grid. The various integrations at the different grid levels are recursively interleaved to minimize the span over which any temporal interpolation need take place. Periodically, for consistency purposes, it is necessary to project a fine grid solution on to its underlying coarse grid. Figure 4 shows the sequence of integration steps and back projections for a three level grid $\{G_0, G_1, G_2\}$ with refinement ratios of 2 and 4.

The integration of an individual grid is extremely simple in concept. Each mesh is surrounded by borders of dummy cells. Prior to integrating a grid, the dummy cells for every mesh patch in the grid are primed with data which is consistent with the various boundary conditions that have to be met. Each mesh patch is then integrated independently by an application dependent, black-box integrator that never actually sees a mesh boundary. Thus, in principle, any cell-centred scheme developed for a single topologically rectangular mesh can form the basis for the integration process.

In general it is necessary to adapt the computational grid to the changes in the evolving flow solution and so the grid structure is dynamic in nature. Monitor functions based on the local solution are used to determine automatically where refinement needs to take place so as to resolve small scale phenomena (Quirk 1991). For a simple example, Figure 5 shows several snapshots taken from the simulation of a shock wave diffracting around a corner. Each snapshot shows the outlines of the mesh patches which go to make the finest grid. This grid clearly conforms to the main features of the flow, namely the diffracted shock front and the vortex located at the apex of the corner (van Dyke 1982). Although the changes in grid structure shown here are dramatic, many adaptions have taken place between each frame. Note that while they might appear small, each mesh patch actually contains several hundred cells. A large number of small grid movements occurs because the adaption process dovetails with the integrations process, see Figure 4. Also note that the adaption always proceeds from fine to coarse so as to ensure that there is never a drop of more than one grid level at the edge of a fine grid to the underlying coarse grid. A grid adaption essentially produces a new set of mesh patches which must be primed with data from the old set of patches before the integration process can proceed. Where a new patch partially overlaps an old patch of the same grid level, for the region of overlap, data may be simply shovelled from the old patch to the new patch. In regions of no such overlap, the required field solution is found by interpolation from the underlying coarse grid solution.

In a typical application the finest grid will contain several hundred mesh patches. Thus, the mesh patch is a sufficiently fine unit of data for efficient parallelism. The parallel AMR algorithm (Quirk & Hanebutte 1993) is implemented using a Single Program Multiple Data (SPMD) model. Each
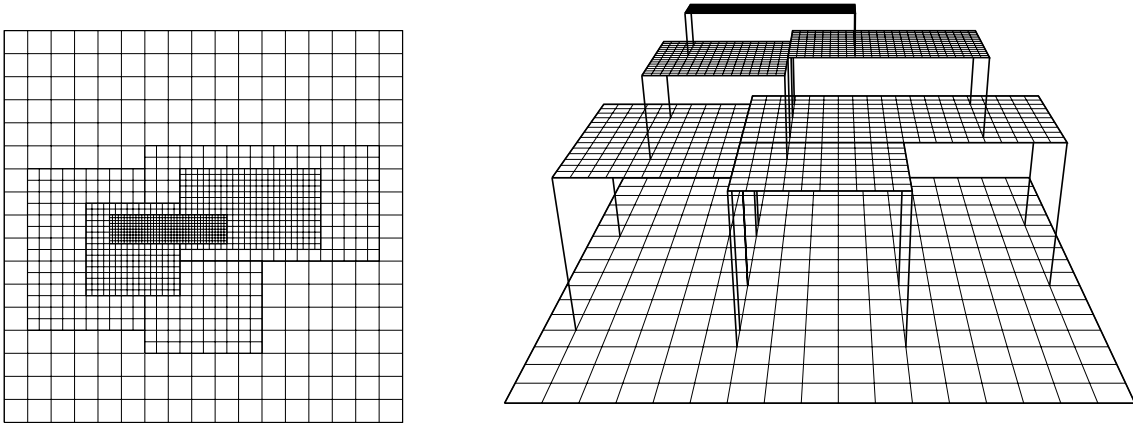
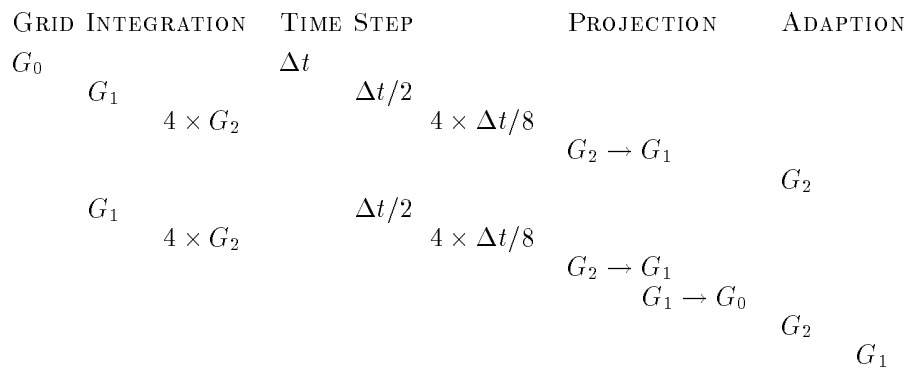Figure 3: The AMR algorithm employs a hierarchical grid system.

Grid Integration	Time Step	Projection	Adaption
G_0	Δt		
G_1	$\Delta t/2$		
$4 \times G_2$	$4 \times \Delta t/8$		
		$G_2 \rightarrow G_1$	
			G_2
G_1	$\Delta t/2$		
$4 \times G_2$	$4 \times \Delta t/8$		
		$G_2 \rightarrow G_1$	
		$G_1 \rightarrow G_0$	
			G_2
			G_1

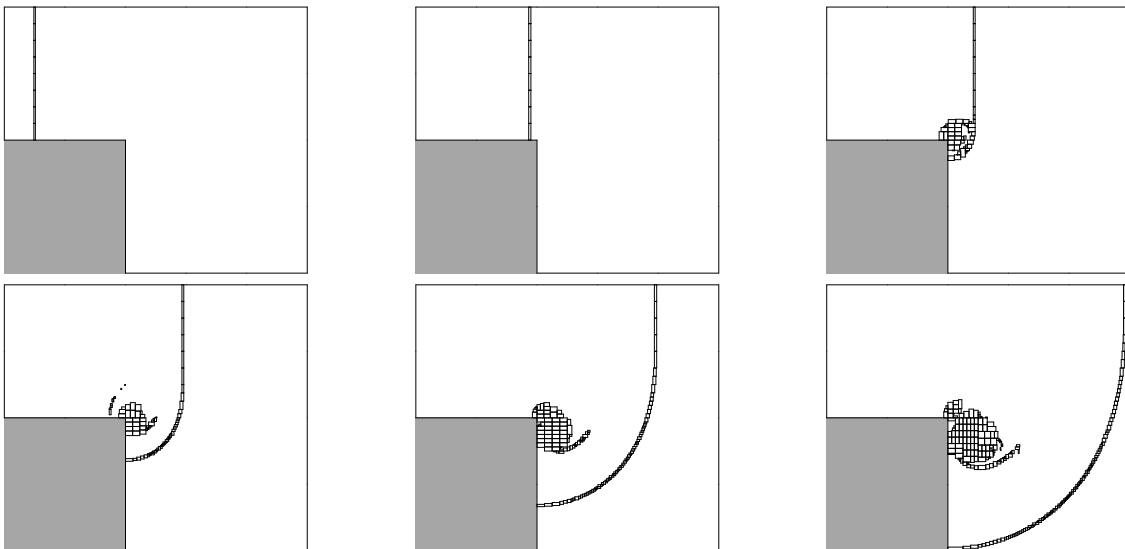Figure 4: Grid operations are recursively interleaved (to be read from top to bottom).

Figure 5: The AMR algorithm employs a dynamic grid system.

10

processing node executes the basic serial algorithm (Quirk 1991) in isolation from all other nodes, except that at a few key points messages are sent between the nodes to supply information that an individual node deems to be missing, that is off-processor. For example, during the integration of a grid, the only point at which a processor needs to know about other processors is during the priming of the dummy cells. Whereas in a serial computation all data fetches are from memory, for a parallel computation some are from memory and some necessitate receiving a message from another processor. Each time the grid adapts, the algorithm generates a schedule of tasks that have to be performed so as to prime correctly the dummy cells of a given grid. If running in parallel, this schedule is parsed to produce a schedule of those tasks that necessitate off-processor fetches. At which point, individual processors can exchange subsets of their fetch schedules, as appropriate, so that every node can construct a schedule of messages that it must send out at some later date. Thus, the priming process is carried out in two phases. First, all the local data fetches are performed as for the serial case. Second, each node sends out the data that has been requested of it. The node then waits for those data items it has requested. For each incoming message it can readily determine from its own schedules what to do with the off-processor data, and so the order in which messages arrive is unimportant. The adaption process and the back projection of the field solution between grid levels also necessitate sizeable amounts of communication, these are handled in a similar fashion to the priming of the dummy cells.

The problem of load balancing the AMR algorithm rests on determining the best distribution of the new patches amongst the processing nodes before the new field solution is interpolated from the old field solution. Currently, this is done using heuristic procedures (Quirk 1994b) which bear strong similarities to classical 'Bin Packing' algorithms (Graham 1969) with the added complication that they must account for the communication costs of data transfer between nodes.

3.3   Flow Visualization Images

In §5, in order to make a qualitative comparison between our numerics and experiment, we present a number of schlieren-type images (Liepmann & Roshko 1957). Such images are useful for identifying weak features which are often lost within contour plots. It should be appreciated that the sensitivity of our numerical images exceeds that which can be obtained experimentally and so there is little to be gained from trying to match exactly the experimental images as some other applications might warrant. Instead we have simply tried to elicit the maximum amount of information possible from our numerical results. Despite their simplicity, our schlieren-type images provide a very effective means of assimilating the various mechanisms that comprise shock refraction phenomena.

The plots shown in Figures 7 and 9 depict the magnitude of the gradient of the density field,

$$|\nabla \rho| = \sqrt{\left(\frac{\partial \rho}{\partial x}\right)^2 + \left(\frac{\partial \rho}{\partial y}\right)^2},$$

and hence they may be viewed as idealized schlieren images; the darker the image the larger the gradient. The density derivatives were computed using straightforward central-differencing, and the following nonlinear shading function, ϕ, was used to accentuate weak flow features,

$$\phi = \exp\left(-k\frac{|\nabla \rho|}{|\nabla \rho|_{max}}\right),$$

11

where k is a constant that took the value 600 for the light fluid and 120 for the heavy fluid. Using a 24 bit colour graphics system the grey shades outside the bubble were produced using the $< R, G, B >$ triplet $< 255\phi, 255\phi, 255\phi >$, and the shades within the bubble were produced using $< 204\phi, 204\phi, 255\phi >$.

We also present a number of realistically lit surface plots, see Figures 8,10 and 15, which are useful for determining the strengths of certain flow features. But lack of space prevents us from describing how these plots were produced.

4 Computational Set-up

For our investigation of the dynamics of a shock-bubble interaction we have reproduced numerically two of the experiments performed by Haas & Sturtevant (1987). Namely, the interactions of a $M_S = 1.22$ planar shock wave, moving through air, with a cylindrical bubble of either helium or Refrigerant R22 ($CHClF_2$). Whereas the helium bubble is lighter than the surrounding air and so acts as a divergent acoustic lens, the R22 bubble is heavier and therefore acts as a convergent acoustic lens. As will be seen in §5, these two cases lead to very different flow behaviour.

In the experiments the bubbles were produced by inflating a cylindrical former whose walls were made from a very thin membrane of nitrocellulose. Thus good control was exercised over the shape of the bubble and the resultant flows were almost two-dimensional, and so our computations which are two-dimensional can be expected to mimic the experiments fairly closely. Haas & Sturtevant produced three sets of results: (i) flow visualization in the form of spark shadowgraphs; (ii) velocities for certain key flow features; (iii) pressure traces measured at points downstream of the bubble along the axis of flow symmetry. We have produced similar sets of results from our simulations. However, it should be appreciated that the experimental results (shadowgraphs and velocities), unlike their computational counterparts, represent a compilation from a series of runs for each bubble case. Only a single spark shadowgraph could be taken from each run, and so the complete record was formed by repeating the experiment with different delay times to the exposure of the shadowgraph image. While this method produced excellent images, the accuracy of the velocity measurements necessarily suffered: since each measurement is derived from a sequence of images it is sensitive to the repeatability of the experiment. The general uncertainty in the velocity measurements is thought to be 11%, with the exception of a few instances for which it is thought to be as large as 30% (Haas & Sturtevant 1987).

A schematic of our computational set-up is shown in Figure 6. We have assumed that the flow field is symmetric about the axis of the shock tube and so only the top half of the flow field (ABCD) was computed. The following boundary conditions were applied to the flow domain: sides BC and DA were treated as solid walls using a standard reflecting boundary procedure (Quirk 1991); the inflow along side CD was specified using the exact flow conditions behind the incident shock wave; zeroth-order extrapolation was used along the side AB. Note that neither the upstream nor the downstream boundary treatment is critical since no physical waves reach these boundaries. Of more relevance are the so-called 'start-up' errors which are generated when a shock smears to its natural profile given an exact discontinuity as starting conditions (Hillier 1991). It is for this reason that the incident shock was placed some distance to the right of the bubble so that these errors, which manifest themselves as a pair of low frequency waves moving on the passive characteristics (Quirk

1991), would not have a chance to interfere with the shock-bubble interaction process.

All gas components were modelled as perfect gases; the appropriate values for the ratio of specific heats γ, the gas constant R, and the constant volume specific heat capacity C_V, used for the simulations are given in Table 1. The initial flow field was determined from standard shock relations given the strength of the incident shock wave ($M_S = 1.22$), taking the density and pressure of the quiescent flow ahead of the shock to be unity. The bubble was assumed to be in both thermal and mechanical equilibrium with the surrounding air, therefore its initial density was simply R_{air}/R_{bubble}. For the helium bubble case, it was assumed that the contamination of helium with air was 28% by mass as indicated by Haas and Sturtevant (1987). As can be seen from Table 1, this modifies the gas properties substantially. Given the experiences of Henderson *et al.* (1991), no attempt was made to model the effects of the membrane which was needed in the experiment to separate the two gas components. Therefore, ahead of the shock, each mesh cell was simply initialised with one of two states depending on whether its centre lay inside or outside of the bubble.
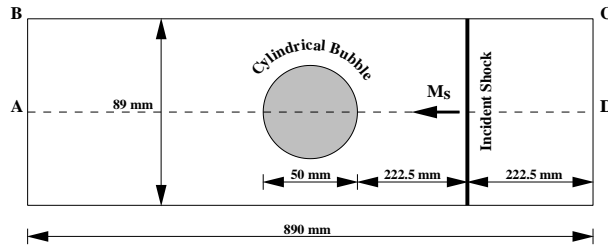

Figure 6: A schematic of the computational domain (not to scale).

Gas Component	γ	R kJ/kg K	C_V kJ/kg K
air	1.4	0.287	0.72
R22	1.249	0.091	0.365
He	1.67	2.08	3.11
He+28% air	1.648	1.578	2.44

Table 1: Gas properties for the simulations.

The computational domain was discretized using 20 coarse mesh patches each of which formed a square of 50 by 50 cells. Additionally, two levels of refinement were used, both with a refinement factor of 4, in order to resolve flow details. Thus the effective computational grid is equivalent to a uniform mesh of 16,000 by 800 cells with a spatial resolution of 0.056 mm. Both simulations were run as parallel computations on a small cluster of workstations (8 Sun Sparc10 Model 51s) and took two evenings each to complete. In this paper, we make no claims as to the computational efficiency of our numerical method, but it is sobering to consider that for the R22 bubble computation the equivalent uniform mesh calculation would require 3.26×10^{11} cell updates ($16 \times 1,592$ iterations on

a mesh 16,000 by 800 cells). For our flow solver, a single processor of a CRAY Y-MP might manage one cell update every $10\mu s$ in which case it would need 905 hours to run the simulation. Brute force computations on super computers do not represent a sensible option for investigations of shock wave phenomena.

5   Results and Discussion: Flow Visualization

In this section we present a number of flow visualization images which reveal certain subtleties of the shock-bubble interactions which were not apparent from either the experiment or previous numerical studies.

5.1   R22 Bubble – Convergent Case

Figure 7 shows a sequence of schlieren-type images from the simulation of the R22 bubble case, by way of comparison, the corresponding sequence of experimental images is also shown. Pleasingly, the simulation clearly reproduces all the salient features of the interaction. In order to bring out the quality of the simulation, and to show how it complements the experiment, we shall now describe this interaction in some detail. But first, so as to interpret correctly the images which follow, recall that the incident shock is moving from right to left and note that the original position of the bubble is marked by a light circle in the numerical images and by what looks like a dark circle with a T-shaped support in the experimental images.

Frame (a) of Figure 7 shows a view of the R22 bubble some $55\mu s$ after it is first hit by the incident shock wave from which it can be seen that the bubble has already undergone a slight deformation. What remains of the incident shock appears as two short vertical line segments near the top and bottom of the bubble. These segments are joined by a curved refracted shock which runs inside the bubble and a curved reflected shock which lies outside the bubble. A one-dimensional analysis for the precise moment the incident shock hits the bubble suggests that the reflected shock is 6.4 times weaker than the refracted shock. An appreciation of the relative strengths of these two waves can be gained from the surface plots for the density and pressure fields (Figure 8 (a)); the reflected wave is so weak it is hardly discernible. Note that the refracted shock lags behind the incident shock because the sound speed inside the bubble is lower than that outside the bubble. Haas and Sturtevant (1987) observed that the refracted shock is slightly thickened at its two endpoints, but no explanation was given as to why this was so. From the surface plots it is clear that the refracted shock is slightly weaker at its endpoints, both the pressure and density surfaces appear slightly chamfered. Thus the thickening is indicative of a compression system that matches the pressure jumps between the weak and strong parts of the refracted shock.

As time moves on, the difference in sound speeds between the bubble and the surrounding air becomes more apparent, and by $115\mu s$ (Figure 7 (b)) the refracted shock has folded such that two side limbs now run roughly normal to its central portion. The surface plots of the density and pressure fields for this time instant (Figure 8 (b)) reveal that each side limb varies markedly in its strength. In essence, for the flow inside the bubble, the air-R22 interface forms a concave ramp. Thus a series of compression waves are required to turn the flow through almost ninety degrees: each side limb is nearly horizontal and so the induced flow is vertical, but the induced flow behind

the central portion of the refracted shock is largely horizontal. Note that the two segments of the incident shock have started to diffract around the downstream half of the bubble, and that the bubble interface shows signs of incipient roll-ups where vorticity has been generated by the passage of the incident shock wave. Now since the flow model is inviscid, the development of these roll-ups will be controlled by vestigial numerical diffusion and so will depend upon the resolution of the computational grid. Nevertheless such roll-ups are qualitatively realistic and it is doubtful whether a viscous flow model would improve matters since a prohibitively fine mesh would be required to resolve the appropriate scales accurately.

By $135\mu s$ the system of compression waves which turns the flow around each of the two bends in the refracted shock has steepened and is clearly visible in the surface plots for the density and pressure fields (Figure 8 (c)). Thus the refracted wave does not extend beyond its junction with the side limbs as was suggested by Löhner *et al.* (1988). Whilst the thickening of the refracted wave shows up much more starkly in the experimental shadowgraphs than it does in the numerical schlieren images, it should be remembered that an experimental shadowgraph represents an integration of the curvature of the density field across the entire width of the shock tube facility used to perform the experiment. Consequently, any small three-dimensionality in the flow field will subtly alter the recorded image in ways that are not always easy to fathom. Here we believe the exaggerated thickening is one such experimental artifact. Because, referring to Figure 7 (c), within the upper of the two thickened limbs that appear in experimental image it is just possible to make out a line which matches the front shown by the numerical image.

Other artifacts of the experiment are much more obvious and so do not cause undue confusion. For example, it is clear from Figure 7 (c) that the bubble's support structure gave rise to a number of spurious waves. As did the walls of the shock tube, but we model reflections from the tube's walls and so these particular waves also appear in the numerical images. Looking beyond the present study, it would be interesting to perform a series of simulations to determine what influence such blockage effects have on the dynamics of interaction process.

Figure 7 (d) shows that by $187\mu s$ the refracted shock has almost been focused down to a point. The increase in peak pressure caused by this focusing is seen in the corresponding surface plots (Figure 8 (d)); at this time, the peak pressure is 2.1 times larger than the expected pressure behind an $M_S = 1.22$ shock wave. Outside the bubble, the top and bottom segments of the incident shock wave have now crossed, following their diffraction around the downstream half of the bubble, and two weak contact discontinuities are now visible. These contacts separate regions of fluid that have been induced into motion by either the diffracted part or the undisturbed part of the incident shock wave. The reflected shocks from the top and bottom walls of the shock tube have now started to pass through the bubble. Again these shocks lag behind their counterparts outside the bubble because of the difference in the sound speeds between the light and heavy fluids. The roll-ups along the bubble interface have become much more pronounced and are very prominent in the surface plot for the pressure field where they appear as tiny scallops (Figure 8 (d)). Note that the passage of the top and bottom reflected shocks through the corrugated bubble interface has given rise to a number of cylindrical acoustic waves which then recombine to form a shock in a manner reminiscent of Huygen's front reconstruction.

Once the refracted shock has been focussed it emerges from the downstream interface to become a transmitted wave which is cylindrical (Figure 7 (e)). The downstream interface of the bubble

necessarily aligns itself with the resultant velocity field which is almost radial and so it takes on a wedge-like shape. Note that the cylindrical transmitted wave is in the stages of catching up the two diffracted segments of the incident shock front. Although the agreement between experiment and computation is poor at this moment in time, it is worth remembering that each shadowgraph was produced from a separate experimental run. Therefore, the fact that we are generally able to match our numerical schlierens so closely to the shadowgraphs is testimony to the repeatability of the experiment. In this one instance, it would appear that the experimental run was relatively poor and that the gross features of the computation are correctly positioned.

If there is any criticism of the simulation, it should be directed at a few subtle shortcomings on the small scale. For example, the two-pronged feature emanating from the left-hand side of the bubble (Figure 7 (e) onwards), seems unduly exaggerated in our simulation. This feature is caused by a narrow jet of fluid which is shot forward during the focusing of the refracted wave. As yet, we cannot categorically state the cause of this exaggeration. It is probably due to the lack of real viscosity in our flow model. In the experiment viscosity causes the jet to spread thus reducing its range of influence. In the simulation, which is inviscid, any spreading of the jet is simply down to residual numerical diffusion. Given the resolution of our computation, this residual diffusion is very small and so the spreading of the jet will be underdone giving it an exaggerated range of influence. However, it is conceivable that the exaggeration is yet another obscure numerical failing of the type catalogued by Quirk (1994a).

By 342μs the bubble has moved appreciably from its original position and it has started to elongate (Figure 7 (g)). Inside the bubble there is a backward moving shock which was born from the internal reflection of the refracted shock from the downstream interface. In the numerical image a number of weaker waves are also apparent, these are caused by waves which pass through the bubble because of reflections from the walls of the shock tube and which subsequently lead to other internal reflections from the bubble interface. Outside the bubble, the transmitted wave has reflected from the walls of the shock tube. Interestingly, as can be seen from the surface plots for this time (Figure 7 (g)), spikes in the pressure and density fields still persist where the transmitted wave intersects the bubble interface. The apparent feathering of the transmitted shock is due to its passage over what is now a corrugated surface given the many roll-ups along the bubble interface.

The internally back-reflected shock wave eventually emerges from the upstream interface to become a backscattered wave (Figure 7 (h)). While the waves resulting from the reflection of the transmitted shock from the top and bottom walls of the shock tube in their turn start to pass through the bubble, further promoting the generation of vorticity along the interface. The bubble continues to elongate and by much later times it evolves into a large vortex pair (Figure 7 (h)). For these late times, when viscous effects might be expected to dominate proceedings, it is remarkable that an inviscid simulation gives such qualitatively good agreement with experiment.

Figure 7: Numerical schlieren images and experimental shadowgraphs (Haas & Sturtevant 1987) from the interaction of an $M_S = 1.22$ shock wave moving from right to left over an R22 cylindrical bubble. Times: (a) 55 μs, (b) 115 μs, (c) 135 μs, (d) 187 μs, (e) 247 μs.

Figure 7: (Contd.) Numerical schlieren images and experimental shadowgraphs (Haas & Sturtevant 1987) from the interaction of an $M_S = 1.22$ shock wave moving from right to left over an R22 cylindrical bubble. Times: (f) 318 μs, (g) 342 μs, (h) 417 μs, (i) 1020 μs. Experimental images ©Cambridge University Press 1987. Reprinted with permission of Cambridge University Press.

Figure 8: Surface plots of the density and pressure fields for the interaction of an $M_S = 1.22$ shock wave with an R22 cylindrical bubble. Times: (a) 55 μs, (b) 115 μs, (c) 135 μs.

Figure 8: (Contd.) Surface plots of the density and pressure fields for the interaction of an $M_S = 1.22$ shock wave with an R22 cylindrical bubble. Times: (d) 187 μs, (e) 247 μs, (g) 342 μs.

5.2   Helium Bubble – Divergent Case

Figure 9 shows a sequence of schlieren-type images from the simulation of the Helium bubble case, again the simulation reproduces all the features of the shock-bubble interaction process.

Figure 9 (a) shows a view of the helium bubble 32 μs after it is first hit by the incident shock wave. As before, there is a curved refracted shock which lies inside the bubble, however, since the helium has a higher sound speed than the surrounding air ($a_{\mathrm{air}}/a_{\mathrm{He}} = 0.35$), the refracted shock now moves ahead of the incident shock. Outside the bubble, the curved reflected wave is neither a simple shock nor a simple expansion wave. A one-dimensional analysis for the precise moment the incident shock hits the bubble suggests that the reflected wave should be a weak expansion (the density jump across this wave is 19% of the density jump between the undisturbed bubble and the surrounding air). Indeed, the surface plots for the pressure and density field confirm that this expectation is true near the axis of flow symmetry(Figure 10 (a)). However, away from this axis there is very little deformation of the bubble and the point of reflection acts as a solid surface giving rise to a reflected shock. Behind this shock there is an expansion system which accounts for the lower pressure to be found behind the rest of the reflected wave due to the collapse of the bubble.

The difference in sound speeds between the bubble and the surrounding air becomes more apparent by 52 μs (Figure 9 (b)) where the refracted shock has run well ahead of the incident wave. A four shock configuration has formed which Henderson *et al.* (1991) have termed twin regular reflection-refraction (TRR). A schematic for this shock configuration is shown in Figure 11. Given the relative positions of the four shocks no discernible contact discontinuity emanates from their intersection point as would be expected in the general case; although one does become visible by 72 μs (Figure 9 (d)). Around 62 μs (Figure 9 (c)) the refracted wave emerges from the left-hand side of the bubble to become the transmitted wave and the resultant internally reflected wave appears as two cusps. As can be seen from Figure 9 (d), this reflected wave is convergent and is being focused along the axis of the bubble but the local increase in pressure is quite small (Figure 10 (d)). By 82 μs (Figure 9 (e)) the internally reflected waves have crossed and are now diverging, here they appear as a small loop. The two branches of the transmitted shock have also now crossed. At 102 μs (Figure 9 (f)), along the axis of flow symmetry the side shock and the transmitted shock have almost merged. Meanwhile, both the original reflected wave and the transmitted shock have reflected from the walls of the shock tube. Interestingly, as can be seen from Figure 10 (f), such spurious reflections can lead to large increases in local pressure. Here the foot of the incident shock, where it meets the shock tube's walls, is reinforced substantially. This spike then proceeds to move away from the wall and eventually interacts with the bubble. At this time, what remains of the incident shock has just started to diffract around the downstream side of the bubble, and the internally reflected wave has emerged from the upstream interface as a weak back scattered wave. This has resulted in a very weak internally reflected wave, so weak in fact that it does not appear in the experimental images. As time moves on, the bubble becomes kidney shaped and spreads laterally in the process (Figure 10 (g)). This change in shape is driven by vorticity generated at the edge of the bubble due to the passage of the shock which induces a jet of air along the axis of flow symmetry. When this jet impinges on the air at the downstream edge of the bubble, which is less easily displaced than the lighter helium, it spreads laterally and the bubble forms a pair of distinct vortical structures (Figure 10 (i)).

21

Figure 9: Numerical schlieren images and experimental shadowgraphs (Haas & Sturtevant 1987) from the interaction of an $M_S = 1.22$ shock wave moving from right to left over a Helium cylindrical bubble. Times: (a) 32 μs, (b) 52 μs, (c) 62 μs, (d) 72 μs, (e) 82 μs.

Figure 9: (Contd.) Numerical schlieren images and experimental shadowgraphs (Haas & Sturtevant 1987) from the interaction of an $M_S = 1.22$ shock wave moving from right to left over a Helium cylindrical bubble. Times: (f) 102 μs, (g) 245 μs, (h) 427 μs, (i) 674 μs. Experimental images ©Cambridge University Press 1987. Reprinted with permission of Cambridge University Press.

Figure 10: Surface plots of the density and pressure fields for the interaction of an $M_S = 1.22$ shock wave with an He cylindrical bubble. Times: (a) 32 μs, (b) 52 μs, (c) 62 μs.

Figure 10: (Contd.) Surface plots of the density and pressure fields for the interaction of an $M_S = 1.22$ shock wave with an He cylindrical bubble. Times: (d) 72 μs, (f) 102 μs, (g) 245 μs.
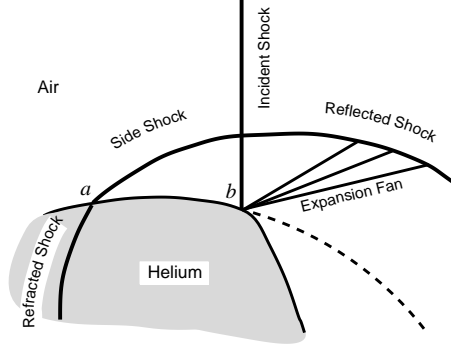
Figure 11: Schematic for twin regular reflection-refraction (TRR).

6   Results and Discussion: Velocities

The results from the previous section clearly indicate that our simulations are qualitatively correct, however, any serious numerical investigation should contain some form of validation exercise. Here, this included a quantitative check on the velocities of several prominent flow features. For each simulation, the positions of certain features were digitized from a sequence of schlieren-type images. Using these measurements, $x - t$ diagrams were then constructed so as to find the velocities. Whereas the experimentally measured velocities had an estimated uncertainty of 11%, here the uncertainty is much smaller. A shock might be smeared over 3 mesh cells, therefore given the resolution of the computational grid its location can be determined to within ± 0.17 mm. This equates to an uncertainty of less that 1% in the worst case velocity measurement. The uncertainty in velocity due to conservation errors is also small at less than 3%.

6.1   R22 Bubble – Convergent Case

The $x - t$ diagram for the shock interaction with the R22 cylindrical bubble is shown in Figure 12. Also shown in this figure is a schematic which identifies the various flow features that have been digitized. A comparison of our computed velocities with their experimentally measured counterparts (Haas & Sturtevant 1987) is given in Table 2. The agreement between the two sets of results lies well within the given 11% experimental error; the worst case (V_T) is just 5.8%. Note that we have ignored the large discrepancy for V_{di} since the experimental value appears to have been tabulated incorrectly; the experimental $x - t$ diagram indicates that V_{di} is close to 130 m/s which is in fair agreement with the computation. Overall, the general agreement between the two sets of velocities confirms the experimentalists' view that the contamination of R22 by air was so small (they estimated it at 3.4 % by mass) as to be negligible.

6.2   Helium Bubble – Divergent Case

The $x - t$ diagram for the shock interaction with the helium bubble is shown in Figure 13, and a comparison with experiment is made in Table 3. As with the R22 case, the two sets of results are in close agreement. However, the effects of air contamination are now significant. As detailed

26

in Section 4, we have assumed that the contamination of helium by air is 28% by mass. If no account is taken of contamination, the velocity results are very different even though the flow remains qualitatively similar. For example, the velocity V_R with 28% contamination is 943 m/s, alternatively, with zero contamination it is found to be 1073 m/s; an increase of 13.5%. The correction for contamination necessarily assumes that the air and helium are homogeneously mixed. Since this would not have been the case in the experiment, our correction can only be viewed as accounting for the gross affects of contamination.

Note that all the measured flow features move more or less at constant velocity; in Figure 13, the kink in the $x - t$ path of the incident shock front near $x = 40$ mm corresponds to the point at which the incident shock is engulfed by the curved transmitted shock wave.
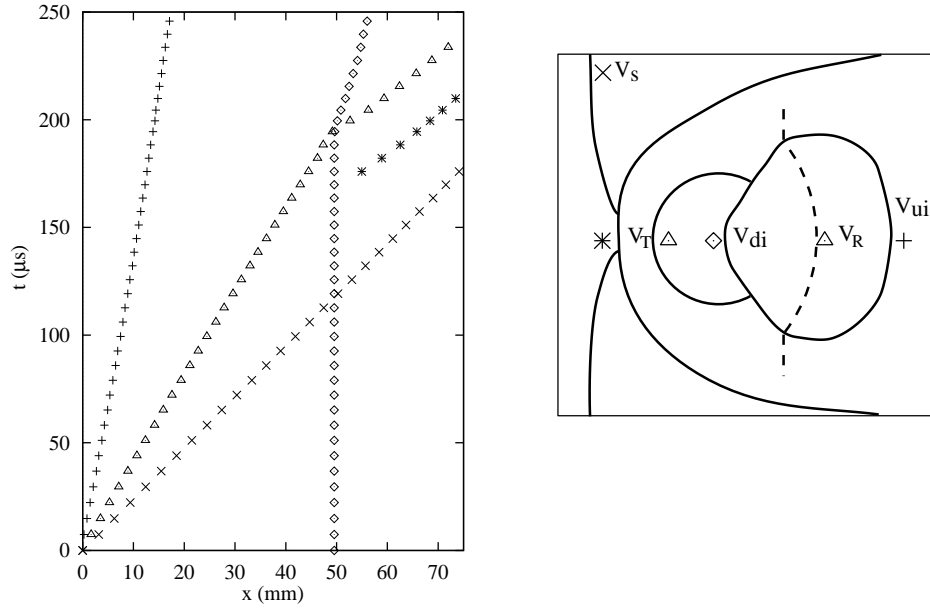

Figure 12: $x - t$ diagram for the R22 cylinder case with a schematic showing the points used to construct the diagram. Key: V_s − incident shock; V_R − refracted shock; V_T − transmitted shock; V_{ui}, V_{uf} − upstream edge of bubble, initial and final times; V_{di}, V_{df} − downstream edge of bubble, initial and final times.

Velocity	V_s	V_R	V_T	V_{ui}	V_{uf}	V_{di}	V_{df}
Computation	420	254	560	70	90	116	82
Experiment	415	240	540	73	90	78	78
% Discrepancy	+1.2	+5.8	+3.7	−4.1	+0.0	N/A	+5.1

Table 2: A comparison of the computed velocities for the R22 cylinder case with those measured experimentally by Haas and Sturtevant (1987); for key, see Figure 12.
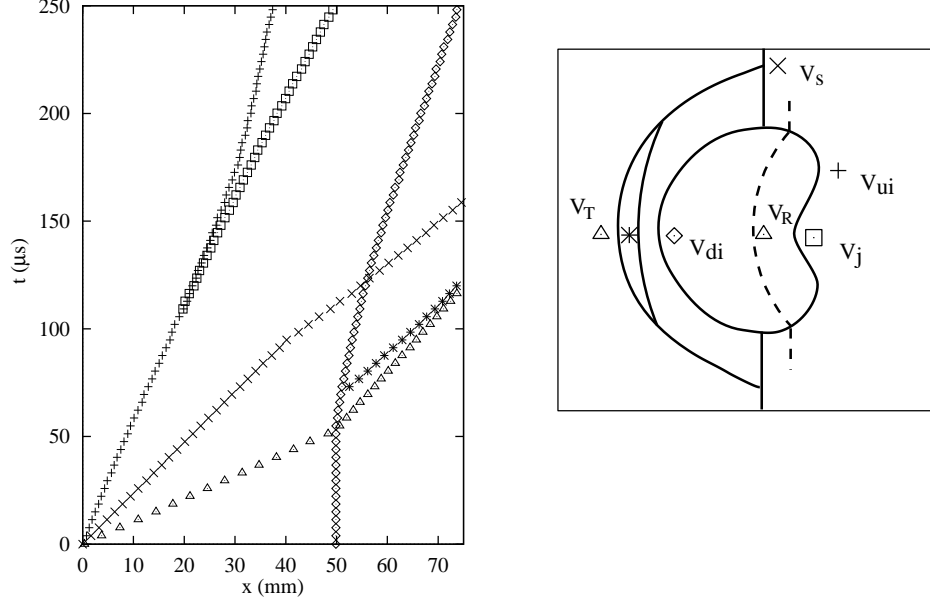
Figure 13: $x - t$ diagram for the He cylinder case with a schematic showing the points used to construct the diagram. Key: V_s – incident shock, V_R – refracted shock, V_T – transmitted shock, V_{ui} – upstream edge of bubble, V_{di} – downstream edge of bubble, V_j – air jet head.

Velocity	V_s	V_R	V_T	V_{ui}	V_{di}	V_j
Computation	422	943	377	178	146	227
Experiment	410	900	393	170	145	230
% Discrepancy	+2.9	+4.8	−4.1	+4.7	+0.7	−1.3

Table 3: A comparison of the computed velocities for the He cylinder case with those measured experimentally by Haas and Sturtevant (1987); for key, see Figure 13.

7 Results and Discussion: Pressure Traces

In addition to producing shadowgraphs, Haas and Sturtevant recorded pressure histories at several stations along the axis of flow symmetry so as to build up a more complete picture of the shock-bubble interaction process. For example, in the heavy bubble case, they noted that the diffracted wave generated a smooth pressure disturbance at a measuring station 3 mm downstream of the initial bubble position, and not a discontinuous disturbance as might be expected from a shock wave. In fact, as was shown in §5, the diffracted front barely constitutes a shock wave in the vicinity of the bubble interface: the surface plots in Figure 8 reveal that along the interface the pressure field ramps up gradually behind the diffracted wave and is not discontinuous, hence the smooth nature of the measured disturbance.

Although the experimental pressure traces only provide a local view of events and so are not as informative as the present pressure surfaces, it was hoped that they could be used to provide further

28

quantitative evidence as to the accuracy of the simulations. Unfortunately, the traces cannot be relied upon to provide an accurate benchmark since the measuring process was invasive. A pressure transducer was mounted on a movable endwall placed within the shock tube (Haas & Sturtevant 1987), thus the transducer actually measured the pressure disturbances for waves reflecting off the endwall and not the local pressure as desired. Consequently, the transducer would be expected to produce readings on the high side. This agrees with our findings. For example, the experiment gave the peak pressure in the heavy bubble case as 7.7 bar, but the simulation suggests that it is close to 5.1 bar. Also, the experiment indicated that the long time pressure, once all the disturbances have died away, to be about 2.2 bar. But this pressure should be close to the pressure behind a $M_S = 1.22$ shock wave which is only 1.56 bar (the simulation gave the long time pressure to be 1.6 bar). Here, the numerics provide a quantitative assessment of the errors introduced by the practicalities of the experimental setup.

Although we cannot make a useful comparison against experiment, for completeness we present the numerical pressure traces from the heavy bubble case, see Figure 14 below (cf. Figure 16 of Haas & Sturtevant, 1987)
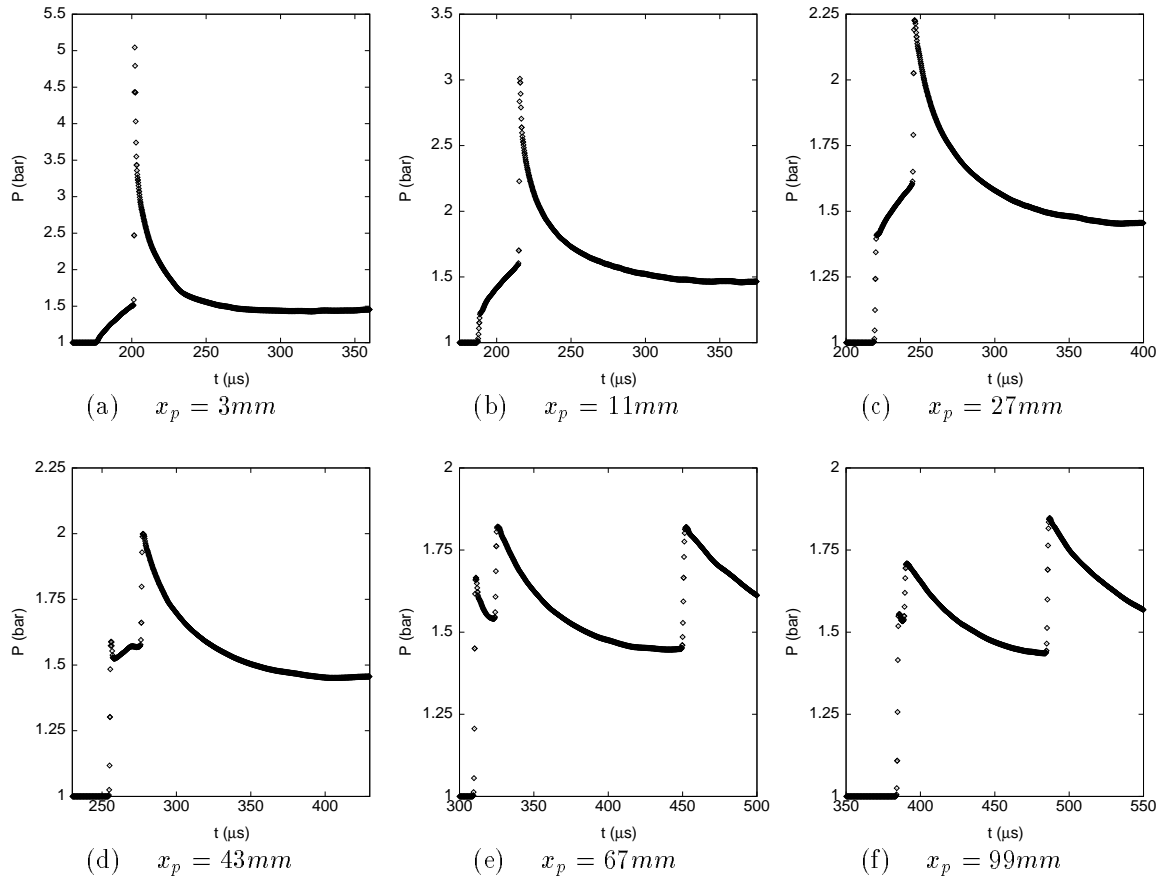

Figure 14: Pressure histories for several stations downstream of the R22 cylinder.

8   Results and Discussion: Vorticity Generation

Although it takes us beyond the main purpose of this paper, we can use our numerical results to make some observations which are pertinent to several recent studies on shock-induced mixing: the present two-dimensional, unsteady flow is analogous to a three-dimensional, steady flow that has been proposed as a mechanism to ensure efficient mixing of air and fuel in supersonic combustions systems (Marble *et al.* 1987; Drummond & Givi 1994). Essentially, vorticity which is impulsively generated by the passage of the shock through the bubble drives a mixing process which is reminiscent of the Richtmyer-Meshkov instability (Richtmyer 1960; Meshkov 1970; Rupert 1992). Thus the effectiveness of this type of mixing rests on the amount of vorticity generated during the early stages of a shock-bubble interaction, therefore it would be useful to have a means of predicting the vorticity produced from a given set of initial conditions.

The basic mechanism which produces the vorticity, $\boldsymbol{\omega}$, is not in doubt. Recall that the vorticity evolution equation, which is derived from the curl of the momentum equation, contains the baroclinic torque term

$$\nabla \times \left(\frac{1}{\rho} \nabla p \right).$$

This term may be recast so as to write the vorticity equation in the form

$$\frac{D\boldsymbol{\omega}}{Dt} = \quad \ldots \quad + \frac{1}{\rho^2} \nabla \rho \times \nabla p,$$

from which it can be seen that vorticity is produced whenever there is a misalignment in the gradients of the density and pressure fields (Shercliff 1977). In the case of a shock-bubble interaction, such a misalignment occurs because the propagating shock wave imposes a local pressure gradient which is largely independent of the local density gradient imposed by the bubble inhomogeneity.

Several authors have devised simple analytic expressions which serve to predict the amount of vorticity produced from an isolated shock-bubble interaction. Typically, this is done by making enough simplifying assumptions that the baroclinic torque can be integrated over the bubble for the duration of the interaction (*e.g.* Picone & Boris 1987 and Yang *et al.* 1994). Of the two referenced models, the one due to Yang *et al.* appears to provide the better predictions. For a range of interaction cases, it provides a vorticity prediction which is within 15% of that found by direct simulation, while the Picone & Boris model is sometimes out by more than a factor of 2. Yang *et al.* claim that the performance of their model stems from the fact that it retains the essential features of a shock-bubble interaction. But our simulations indicate that this claim is debatable. Consider the following observations on the vorticity produced by the helium bubble case.

Picone & Boris (1987) noted that the production of vorticity along the bubble interface appeared to track the fastest moving shock wave, the uncertainty arose from the low resolution of their computations. Our more detailed simulations show that this observation is essentially correct. Most of the vorticity is produced where the side shock intersects the bubble interface, point a on Figure 11. But a sizeable amount is also produced where the Mach stem crosses the interface, point b on Figure 11. These observations were gleaned from surface plots of the baroclinic torque term, for which there are just two localized spikes at the points a and b. However, over the leeward side of the bubble only the side shock produces any significant amount of vorticity: the Mach stem weakens as it diffracts around the bubble and it eventually becomes too weak to generate much vorticity, see

Figure 10 (f). Figure 15 shows two snapshots of the accumulated vorticity, cf. the surface plots of the pressure and density fields shown in Figure 10. Note that very little vorticity is generated due to plain shock curvature. Also note that the distribution of vorticity is not symmetric. More vorticity is deposited on the windward side of the bubble than on the leeward side.

The above observations undermine the assumptions upon which most vorticity prediction models are based. By way of example, take the Yang *et al.* (1994) model. Since vorticity is predominantly generated by the side shock, production effectively ceases once the transmitted wave emerges from the bubble. Note that the incident shock has only traversed half the bubble by the time this happens, see Figure 9 (c). But the model assumes that vorticity is generated continuously for the time it takes the incident shock to traverse the bubble (*i.e.* twice as long as is necessary), thus it might be expected to over predict the final amount of vorticity by a factor of 2. In part, this overprediction is counteracted by the fact that on the windward side of the bubble there are two major sources of vorticity production and not one as in the model. It is further assumed that the pressure gradient responsible for the baroclinic torque remains constant throughout the interaction and that it is equal to the jump across the incident shock. But this is not the case. Early on, the relevant pressure jump is significantly larger than that across the incident shock (see Figure 10 (a)), while at later times it is significantly smaller (see Figure 10 (f)). Other inaccuracies arise because it is assumed that this pressure gradient always acts horizontally. In fact, in the heavy bubble case it generally acts in a direction which is tangential to the bubble interface, see Figure 8 (b)–(d).

In view of the complex nature of a shock-bubble interaction, it would seem that any set of assumptions which are sufficient to yield a simple analytic expression for the vorticity are unlikely to be justifiable at all moments of the interaction. In practice, the individual errors from the separate assumption often partially cancel. Therefore, some models can provide reasonable vorticity predictions albeit in a slightly fortuitous manner.

Figure 15: Surface plots of the magnitude of the vorticity field for the interaction of an $M_S = 1.22$ shock wave with an He cylindrical bubble. Times: (a) 32 μs, (e) 102 μs.

9   Concluding Remarks

Our numerical results provide a comprehensive view of the complex phenomenological behaviour of shock-bubble interactions. For example, in the case of a weak shock interacting with a heavy bubble, they reveal the precise nature of the refracted shock front within the bubble and thereby identify for the first time the cause of the shock thickening that was observed experimentally. Previous computational efforts to reveal this internal structure were marred by both a lack of grid resolution and by numerical artifacts which interfered with the flow. Despite having only modest computing resources, we were able to obtain the necessary high resolution by using a sophisticated adaptive mesh refinement algorithm. Here, we estimate that this algorithm led to between a forty and fifty-fold decrease in computational effort. The actual computation time was further reduced by running each simulation as a parallel computation on a small cluster of workstations. Additionally, we took great care to avoid the numerical artifacts which generally plague multicomponent simulations. In essence, we elected to suffer a small, controlled conservation error so as to avoid spurious oscillations at material interfaces. Thus the present simulations are able to match their experimental counterparts both qualitatively and quantitatively, at least to the limits set by the physical model and by experimental uncertainties.

Simulations of complex time-dependent phenomena usually prove difficult to decipher, since they generate huge amounts of data. Therefore, in part, the clarity of the present numerical study rests with its careful use of computer graphics. For this paper we chose to present schlieren-type images of the computed flow since they enable us to compare our results directly against experiment. Such images are useful for identifying weak features which are often lost on contour plots and so they provide a very effective means for assimilating the various mechanisms that comprise shock refraction phenomena. We have also presented a number of realistically lit, surface plots which are invaluable for gauging the strengths of individual flow features. Here, for example, they are used to expose the weaknesses of current analytic models for predicting the amount of vorticity produced by a shock-bubble interaction.

Finally, while the present computational machinery is capable of producing excellent results, it is only fair to point out that it represents a considerable investment of effort. However, the machinery is general purpose and so the large development costs can be defrayed. Looking to the future, we hope to apply our machinery to other shock wave phenomena that are not yet fully understood, preferably as part of a combined theoretical-numerical-experimental study.

References

ABGRALL, R. 1988 Generalisation of the Roe scheme for the computation of mixture of perfect gases. *La Recherche Aérospatiale* **6**, 31-43.

ABD-EL-FATTAH, A.M. & HENDERSON, L.F. 1978*a* Shock waves at a fast slow gas interface. *J. Fluid Mech.* **86**, 15–32.

ABD-EL-FATTAH, A.M. & HENDERSON, L.F. 1978*b* Shock waves at a slow fast gas interface. *J. Fluid Mech.* **89**, 79–95.

BELL, J., BERGER, M., SALTZMANN, J. & WELCOME, M. 1994 Three-dimensional adaptive mesh refinement for hyperbolic conservation laws. *SIAM J. Sci. Comput.* **15**, 127–138.

BEN-DOR, G. & TAKAYAMA, K. 1985 Analytical prediction of the transition from Mach to regular reflection over cylindrical concave wedges. *J. Fluid Mech.* **158**, 365–380.

BERGER, M.J. & OLIGER, J. 1984 Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comp. Phys.* **53**, 482–512.

CATHERASOO, C.J. & STURTEVANT, B. 1983 Shock dynamics in nonuniform media. *J. Fluid Mech.* **127**, 539–561.

COLELLA P., HENDERSON L. & PUCKETT E.G. 1989 A numerical study of shock wave refractions at a gas interface. AIAA paper 89-1973-CP,*Proc. AIAA 9th CFD Conference, Buffalo, NY*, 426–439.

COURANT, R. & FRIEDRICHS, K.O. 1948 *Supersonic flow and shock waves*, Interscience, New York.

DRUMMOND, J.P. & GIVI, P. 1994 Suppression and enhancement of mixing in high-speed reacting flow fields. In *Combustion in high speed flows*, 191–229. Edited by Buckmaster, J., Jackson, T.L. & Kumar, A., Kluwer, Dordrecht.

GRAHAM, R.L. 1969 Bounds on certain multiprocessing anomalies. *SIAM J. Appl. Math.* **17**, 416–429.

GROVE, J.W. & MENIKOFF, R. 1990 Anomalous reflection of a shock wave at a fluid interface. *J. Fluid Mech.* **219**, 313–336.

HAAS J.-F. & STURTEVANT B. 1987 Interaction of weak shock waves with cylindrical and spherical gas inhomogeneities. *J. Fluid Mech.* **181**, 41–76.

HENDERSON, L.F., COLELLA P. & PUCKETT E.G. 1991 On the refraction of shock waves at a slow-fast gas interface. *J. Fluid Mech.* **224**, 1–27.

HILLIER, R. 1991 Computation of shock wave diffraction at a ninety degrees convex edge. *Shock Waves* **1**, 89–98.

HOU, T.Y. & LE FLOCH, PH. 1991 Why nonconservative schemes converge to wrong solutions: Error analysis. *Comm. Pure Appl. Math.* (submitted).

KARNI, S. 1992 Viscous shock profiles and primitive formulations. *SIAM J. Num. Anal.* **29**, 1592–1609.

KARNI, S. 1994*a* Multi-component flow calculations by a consistent primitive algorithm. *J. Comp. Phys.* **112**, 31–43.

KARNI, S. 1994*b* Interfacial instabilities in multicomponent fluids (in preparation).

LARROUTUROU, B. 1991 How to preserve the mass fractions positive when computing compressible multi-component flows. *J. Comp. Phys.* **95**, 59–84.

LAX, P.D. 1954 Weak solutions of nonlinear hyperbolic equations and their numerical computation. *Comm. Pure Appl. Math.* **7**, 159–163.

LAX, P.D. 1972 *Hyperbolic systems of conservation laws and the mathematical theory of shock waves.* SIAM Monograph series.

LIEPMANN, H.W. & ROSHKO, A. 1957 *Elements of Gasdynamics* John Wiley & Sons, Chichester, p. 157.

LÖHNER, R., PICONE, J.M. & BORIS, J.P. 1988 Wave structure produced by shock propagation through a dense bubble gas. In *Proc. 16th Intl. Symp. on Shock Tubes and Waves*, held at Aachen, Germany 1987. Edited by H. Gronig, Weinheim, Germany, 1988. 613–619.

MARBLE, F. E., HENDRICKS, G.J. & ZUKOSKI, E.E. 1987 Progress toward shock enhancement of supersonic combustion processes *AIAA Paper 87-1880*.

MESHKOV, Y.Y. 1970 Instability of a shock wave accelerated interface between two gases. *NASA translation TT F-13*, 074 R.F.

PICONE, J.M. & BORIS, J.P. 1988 Vorticity generation by shock propagation through bubbles in a gas. *J. Fluid Mech.* **189**, 23–51.

QUIRK, J.J. 1991 *PhD Thesis*, An adaptive mesh refinement algorithm for computational shock hydrodynamics. Cranfield Institute of Technology, U.K.

QUIRK, J.J. & HANEBUTTE 1993 A Parallel Adaptive Mesh Refinement algorithm. *NASA CR-191530, ICASE Report No. 93-63* (submitted to *J. Comput. Phys.*).

QUIRK, J.J. 1994*a* A Contribution to the Great Riemann Solver Debate. *Intl. J. Num. Meth. Fluids* **18**, 555–574.

QUIRK, J.J. 1994*b* Dynamic load balancing strategies for parallel adaptive mesh refinement (in preparation).

RICHTMYER, R.D. 1960 Taylor instability in shock acceleration of compressible fluids. *Comm. Pure Appl. Math.* **23**, 297–319.

ROE, P.L. 1982 Fluctuations and signals - A framework for numerical evolution problems. In *Numerical methods for fluid dynamics.* Edited by Morton, K.W. and Baines, M.J., 219–257, Academic Press, New York.

ROE, P.L. 1986 Characteristic-based schemes for the Euler Equations *Ann. Rev. Fluid Mech.* **18**, 337–65.

RUPERT, V. 1992 Shock-interface interaction: current research on the Richtmyer-Meshkov problem. In *Shock Waves*, Proceedings of the 18th Intl. Symp. on Shock Waves, held at Sendai, Japan 1991. Edited by Takayama, K., Springer-Verlag, Berlin, 1992. 83–94.

SCHWENDEMAN, D.W. 1988 Numerical shock propagation in non-uniform media *J. Fluid Mech.* **188**, 383–410.

SHERCLIFF, J.A. 1977 *Vector Fields*, Cambridge University Press, p. 276.

STRANG, G. 1968 On the construction and comparison of finite-difference schemes. *SIAM J. Num. Anal.* **5**, 506–517.

TON, V.T., KARAGOZIAN A.R., ENGQUIST B.E. & OSHER, S.J. 1991 482–512. paper WSS/C191-101, *Proc. Combustion Institute 1991 Fall Meeting, UCLA.*

VAN DYKE, M. 1982 *An album of fluid motion* The Parabolic Press, Stanford, California, p. 148.

VON NEUMANN, J. 1963 *Collected Works* vol. 6, Pergamon Press.

WANG, J.C.T. & WIDOPHF, G.F. 1990 Numerical simulation of blast flowfields using a high resolution TVD finite volume scheme. *Computers & Fluids* **18**, 103–137.

WHITHAM, G.B. 1957 A new approach to problems of shock dynamics. Part I Two dimensional problems. *J. Fluid Mech.* **2**, 145–171.

WHITHAM, G.B. 1958 On the propagation of shock waves through regions of non-uniform area or flow. *J. Fluid Mech.* **4** 337–360.

YANG J., KUBOTA T. & ZUKOSKI E.E. 1993 Applications of shock-induced mixing to supersonic combustion. *AIAA J.* **31**, 854–862.

YANG J., KUBOTA T. & ZUKOSKI E.E. 1994 A model for characterization of a vortex pair formed by shock passage over a light-gas inhomogeneity. *J. Fluid Mech.* **258**, 217–244.

ZWAS, G. & ROSEMAN, J. 1973 The effect of nonlinear transformations on the computation of weak solutions. *J. Comp. Phys.* **12**, 179–186