N96-18075

SOME OBSERVATIONS ON MESH REFINEMENT SCHEMES APPLIED TO SHOCK WAVE PHENOMENA*

James J. Quirk[†]
Institute for Computer Applications in Science and Engineering
NASA Langley Research Center, Hampton VA 23681-0001, USA.

SUMMARY

This workshop's double-wedge test problem is taken from one of a sequence of experiments which were performed in order to classify the various canonical interactions between a planar shock wave and a double wedge. Therefore to build up a reasonably broad picture of the performance of our mesh refinement algorithm we have simulated three of these experiments and not just the workshop case. Here, using the results from these simulations together with their experimental counterparts, we make some general observations concerning the development of mesh refinement schemes for shock wave phenomena.

INTRODUCTION

For problems governed by disparate physical scales, the potential savings to be gained from using local mesh refinement are often so large that any strategy will pay handsome dividends: a poor refinement scheme is better than none. Consequently the literature is littered with examples where some form of mesh refinement capability has been botched in a problem specific manner. Superficially the 'quick and dirty' approach appears attractive because the development costs are considerably less than those for a general scheme. In practice, however, the development costs of a general scheme can be recouped across a wide range of projects, and over time the cost/project becomes negligible. On the other hand, with the one-off approach the effective costs accumulate with each passing project and can become unexpectedly large over time. Moreover since one-off schemes rarely reach maturity, they tend to be needlessly expensive to run. Therefore, taken overall, we feel there is no merit in pursuing one-off refinement strategies.

Nevertheless, since an algorithm has to strike a balance between that which is desirable and that which is practicable, an element of 'horses for courses' remains even amongst general purpose mesh refinement schemes. Therefore a method, say, which was designed to provide the cheapest medium-accuracy solution to a steady flow problem might not be competitive when it comes to producing the most accurate solution to a time-dependent problem, and vice versa. Thus some care should be

---

PAGE __78__ INTENTIONALLY BLANK   **PRECEDING PAGE BLANK NOT FILMED**

taken in choosing the most appropriate form of mesh refinement before embarking on what might be an arduous exercise in software development. Given our research interests[10], in 1988 after much deliberation we plumped for a form of embedded mesh refinement first developed by Berger and co-workers[1, 2, 3].

Our formulation[10, 13] has now matured well beyond the development stage and so will not be described here. Instead we wish to engender some discussion as to the strengths and weaknesses of different refinement strategies as applied to investigations of shock wave phenomena. Our aim is not to promote one scheme over another, but to reveal some pitfalls which await the unwary. Therefore to place this discussion in the right context we will first present our numerical results for the workshop double-wedge test problem. This problem was inspired by a series of experiments performed by Takayama *et al.*[15] at Tohuku University to clarify the various types of reflection processes that can occur when a planar shock wave interacts with a double wedge. In view of this, it is worthwhile considering more than just the case chosen for the workshop and we will in fact present results from three different cases.

SHOCK DOUBLE-WEDGE INTERACTIONS

With reference to the schematic shown in Figure 1, we have simulated the interaction of a planar shock wave with three different double-wedge configurations (θ_1, θ_2). These configuration were chosen to match those in the experiments of Takayama *et al.*[15]. Given the instructions for the workshop, the flow was modelled using the two-dimensional Euler equations, taking the equation of state to be that of a perfect gas with ratio of specific heats (γ) set to 1.4 . The Mach number for the incident shock (M_S) was taken to be 2.16 giving a pressure ratio p_2/p_1 of 5.28.

The computational method used for our simulations is the same as in [11] i.e. a non-body-fitted grid was used in conjunction with a two-step finite-volume integration scheme. The effective resolution of the grid was equivalent to that of a uniform mesh of 2240 by 1280 cells. This was obtained using two levels of dynamic refinement, each by a factor of 4, on a uniform base grid of 140 by 80 cells.

Since this paper contains a number of interferograms and schlieren images whose sizes have been reduced solely to keep the length of this paper within acceptable limits we have also placed them on the World Wide Web at URL http://www.icase.edu/~jjq/flowviz/gal_dwedge.html so that they might be viewed at their original quality.

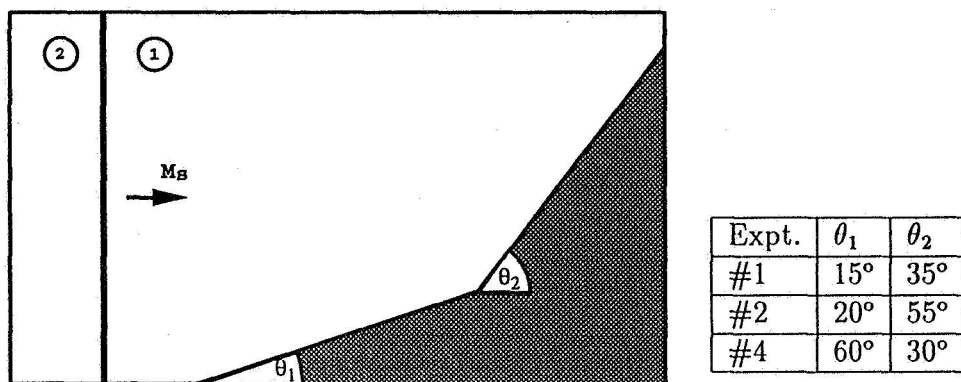| Expt.	θ_1	θ_2
#1	15°	35°
#2	20°	55°
#4	60°	30°

Figure 1: Double-wedge configurations used in the experiments of Takayama *et al.*[15] and our numerical simulations.

Experiment #1

A sequence of schlieren snapshots from our simulation of this interaction are shown in Figure 2. At early times, frames (a) and (b), there is single Mach reflection (SMR) of the incident wave from the first ramp. At intermediate times, frames (c) and (d), the Mach stem from this primary reflection interacts with the second wedge giving rise to a secondary reflection which is also of type SMR. At late times, frames (e) to (h), the secondary reflection interferes with the primary reflection. In Figure 3 a numerical interferogram is shown with its experimental counterpart. The two images are in good agreement, hence there is a reasonable quantitative agreement between simulation and experiment. Nevertheless, there are some clear discrepancies on the small scale. For example, in the experiment the base of the primary reflected shock has a small lambda foot due to its interaction with the boundary layer on the bottom wall of the shock tube (see bottom-left corner of image). This feature is missing in the numerical image since the simulation assumed that the flow was inviscid. In principle, adding viscous terms to the simulation is not difficult. However, a much finer grid would have to be used so as to resolve the relevant viscous scales. Thus the cost of the simulation would be increased dramatically and in this instance it is debatable whether the small improvements to be gained by adding physical viscosity would prove cost effective

Experiment #2

A sequence of schlieren snapshots from our simulation of this interaction are shown in Figure 4. At early times, frames (a) and (b), there is SMR of the incident wave from the first ramp as in Experiment #1. However, at intermediate times, frames (c) and (d), the reflection of the Mach stem is now complex Mach reflection (CMR) rather than SMR. At late times, frames (e) to (h), the secondary reflection again interferes with the primary reflection. In Figure 5 a numerical interferogram is shown with its experimental counterpart. The two images are in reasonable agreement, but the tie-up is noticeably poorer than in Experiment #1. Again the discrepancies are due to the lack of physical viscosity in the flow model. For example, in the experimental image there is a recirculation zone at the apex of the first ramp, and the base of the secondary reflected shock has a lambda foot due to its interaction with the boundary layer on the wedge. But these features cannot be reproduced by an inviscid simulation. Here the shock-boundary layer interactions are stronger than in Experiment #1 and have had quite a pronounced affect on the curvature with which both the primary and secondary reflected shocks run in to the wall. Consequently there would be some justification for switching to a viscous simulation for this experiment.

Experiment #4

A sequence of schlieren snapshots from our simulation of this interaction are shown in Figure 6. At early times, frames (b) to (c), the slope of the first wedge is sufficient that there is regular reflection (RR) and not SMR as in the other two experiments. At late times, frames (d) to (h), the incident shock diffracts around the convex corner formed by the two wedges. In Figure 7 a numerical interferogram is shown with its experimental counterpart. The two images are in good agreement except for those regions where viscous effects are expected to be important. Namely, the vortex core near the convex corner, and the foot of the reflected shock where it interacts with the boundary layer on the wall of the shock tube. This interaction affects the curvature of the reflected shock and would seem to account for the difference in the curvature of the fringes between the computational and experimental interferograms. However, the tie-up is sufficiently good that, as in Experiment #1, it is not clear that a viscous simulation would be worth the extra effort involved.
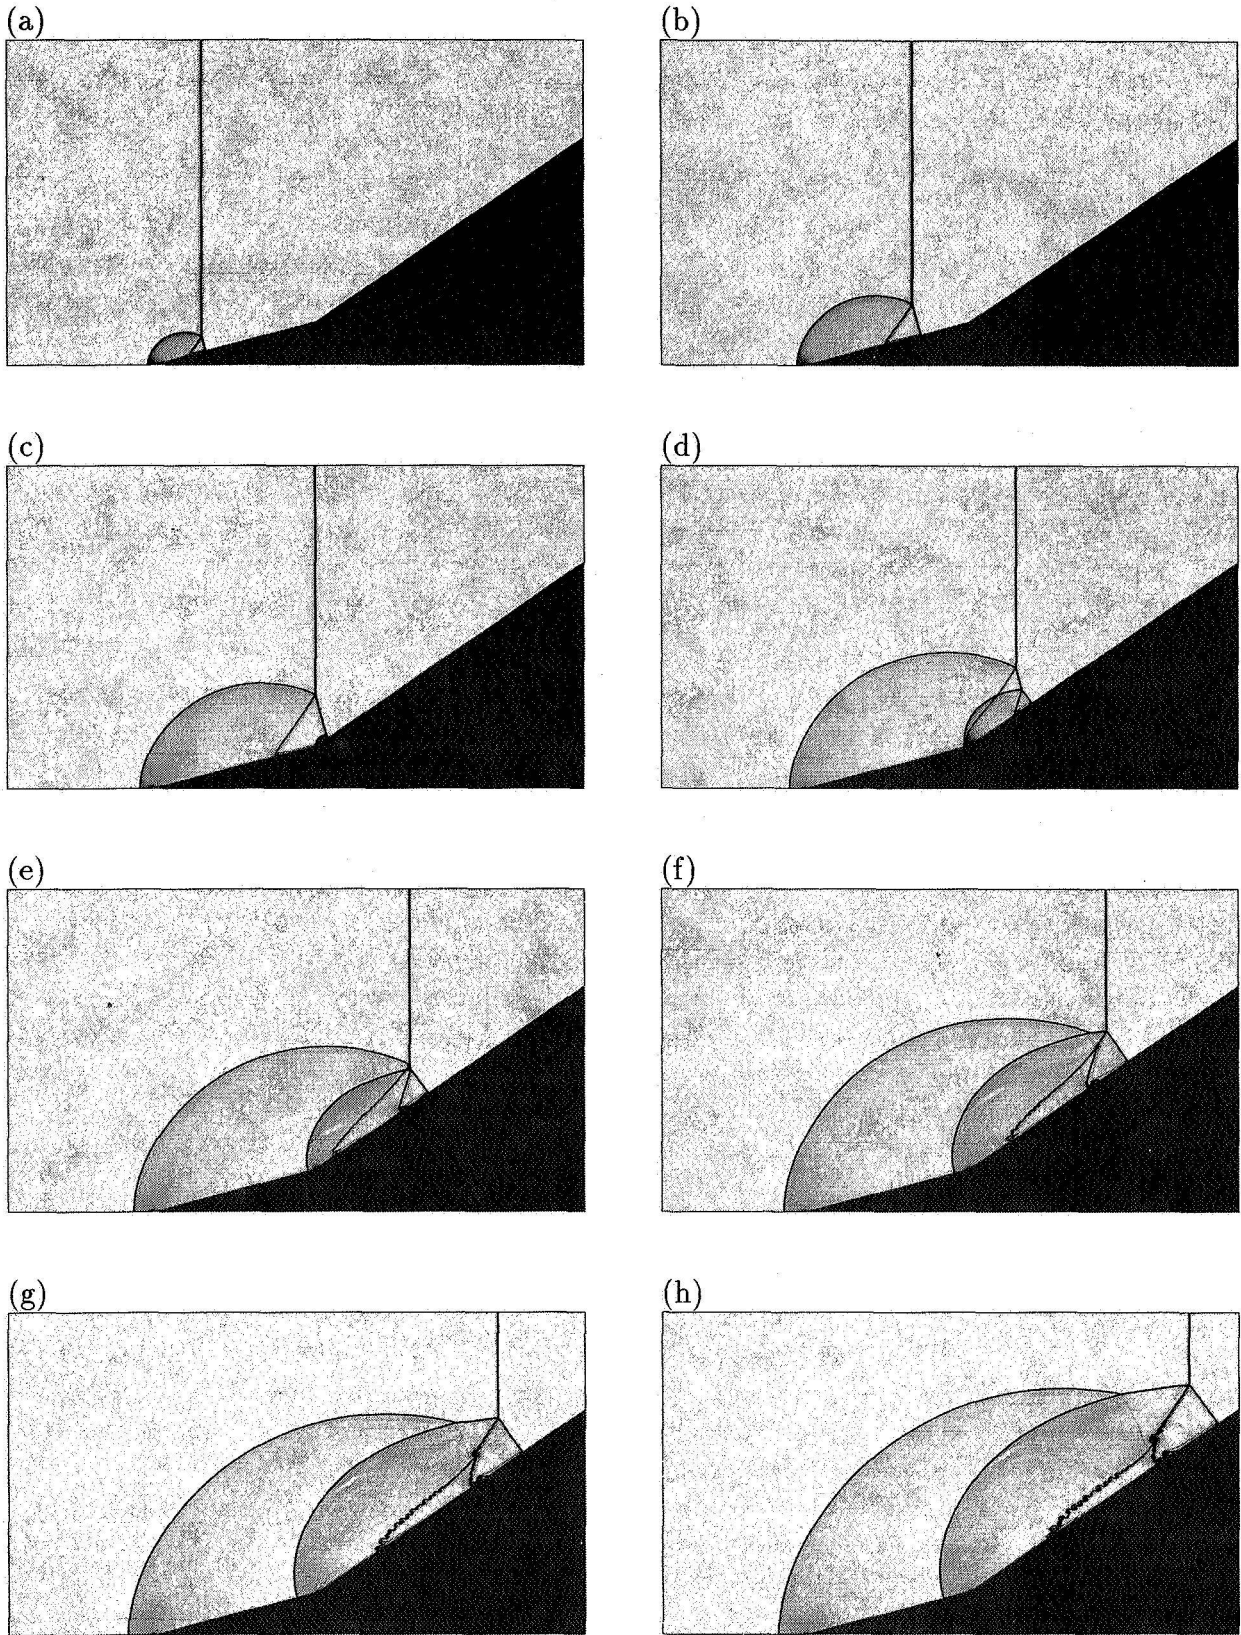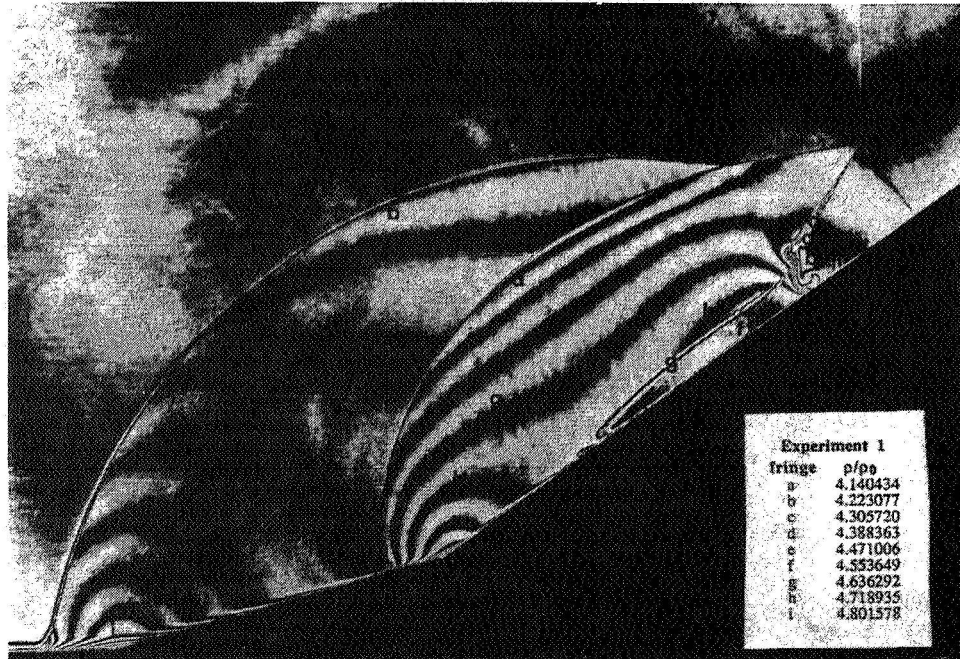
Figure 2: Sequence of schlieren snapshots from the simulation of Experiment #1.

(a) Experimental Interferogram, courtesy of Prof. Takayama
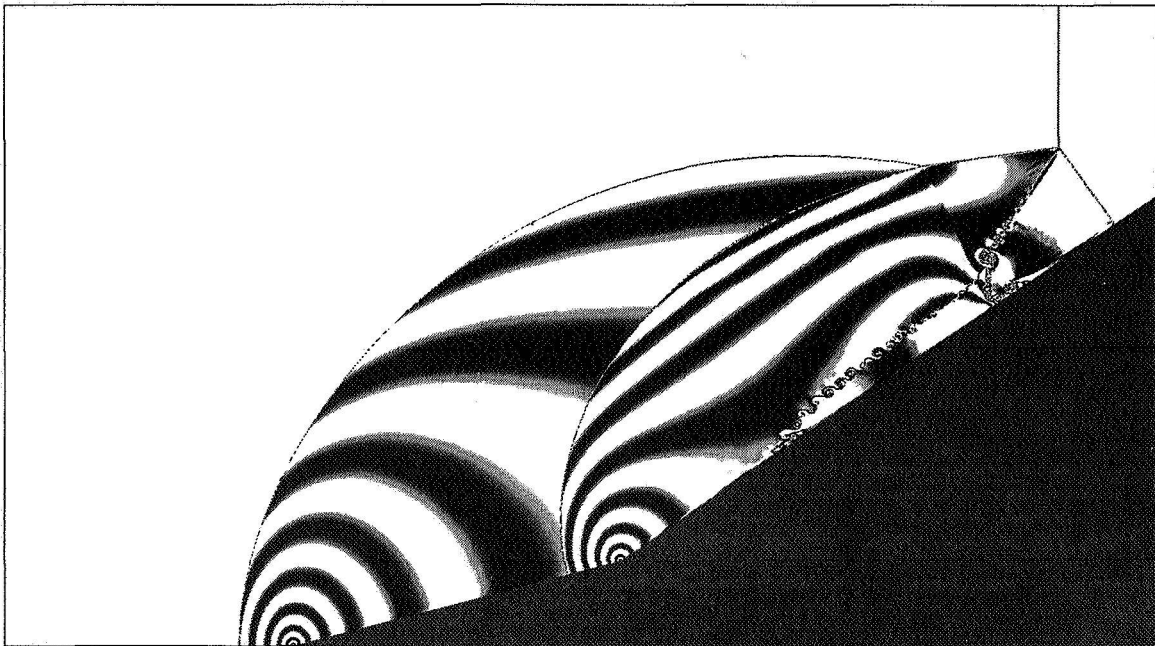

(b) Numerical Interferogram

Figure 3: Comparison between numerical and experimental interferograms for Experiment #1.

Figure 4: Sequence of schlieren snapshots from the simulation of Experiment #2.

(a) Experimental Interferogram, courtesy of Prof. Takayama

(b) Numerical Interferogram

Figure 5: Comparison between numerical and experimental interferograms for Experiment #2.
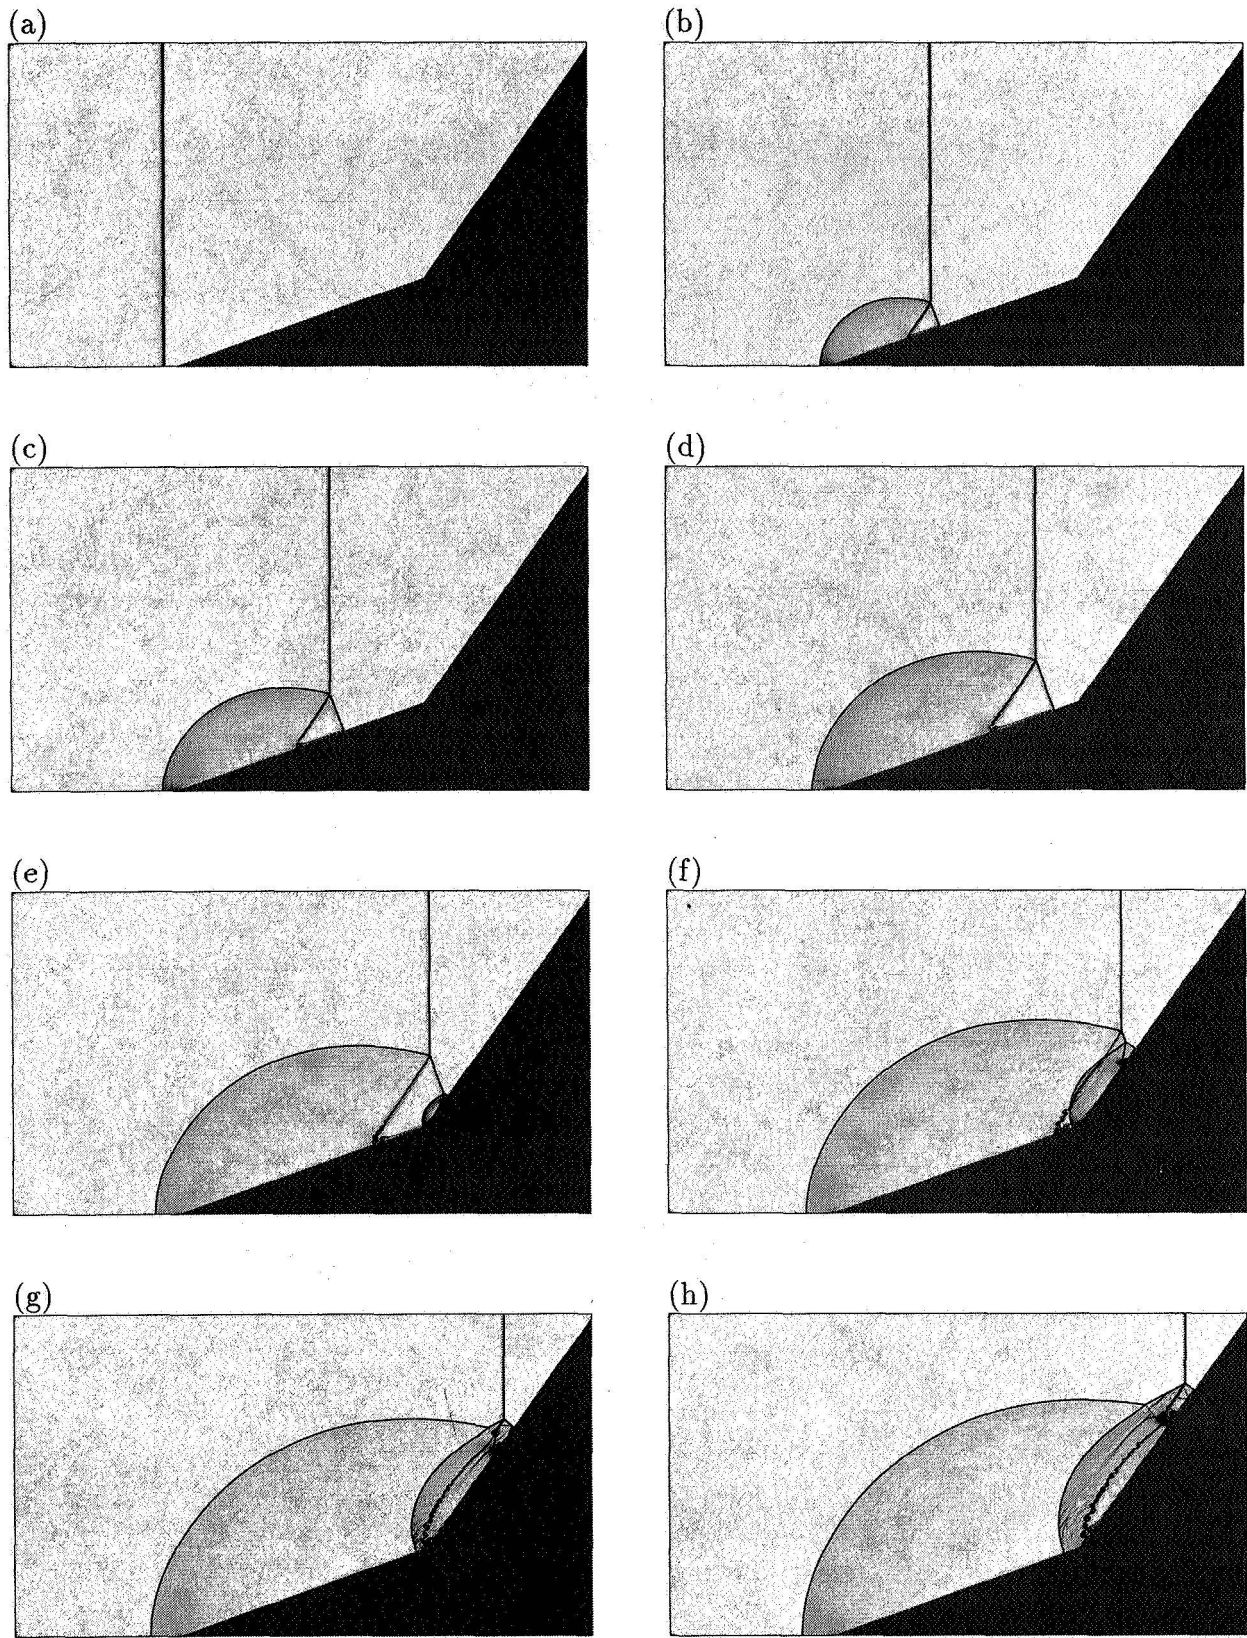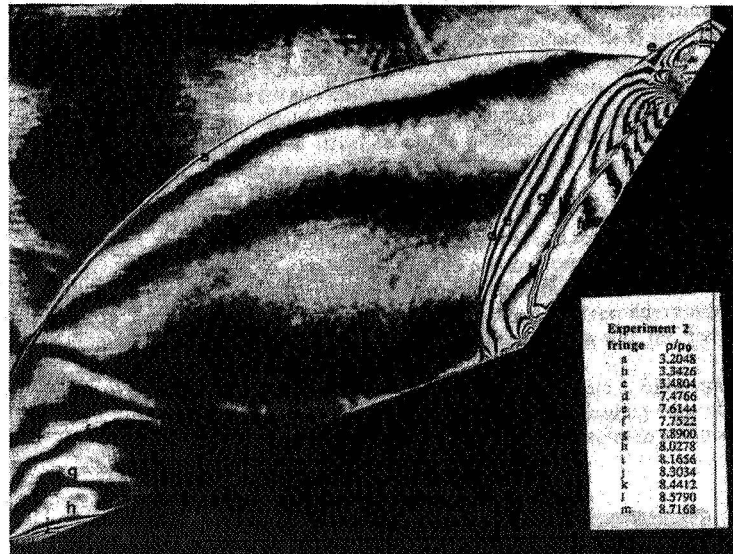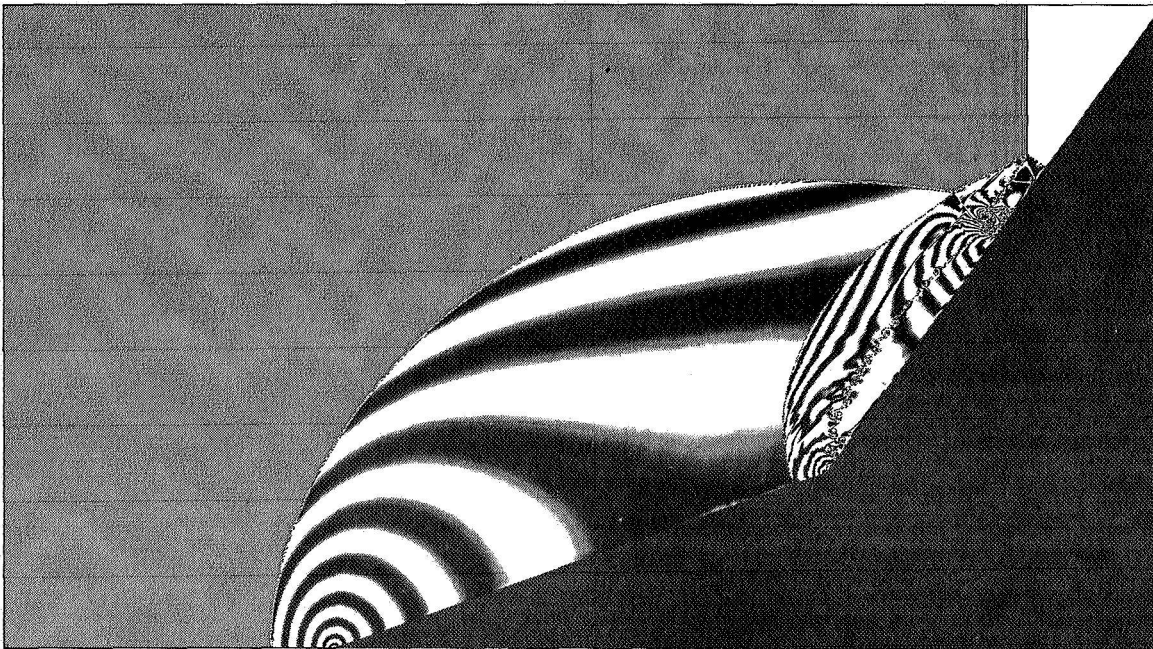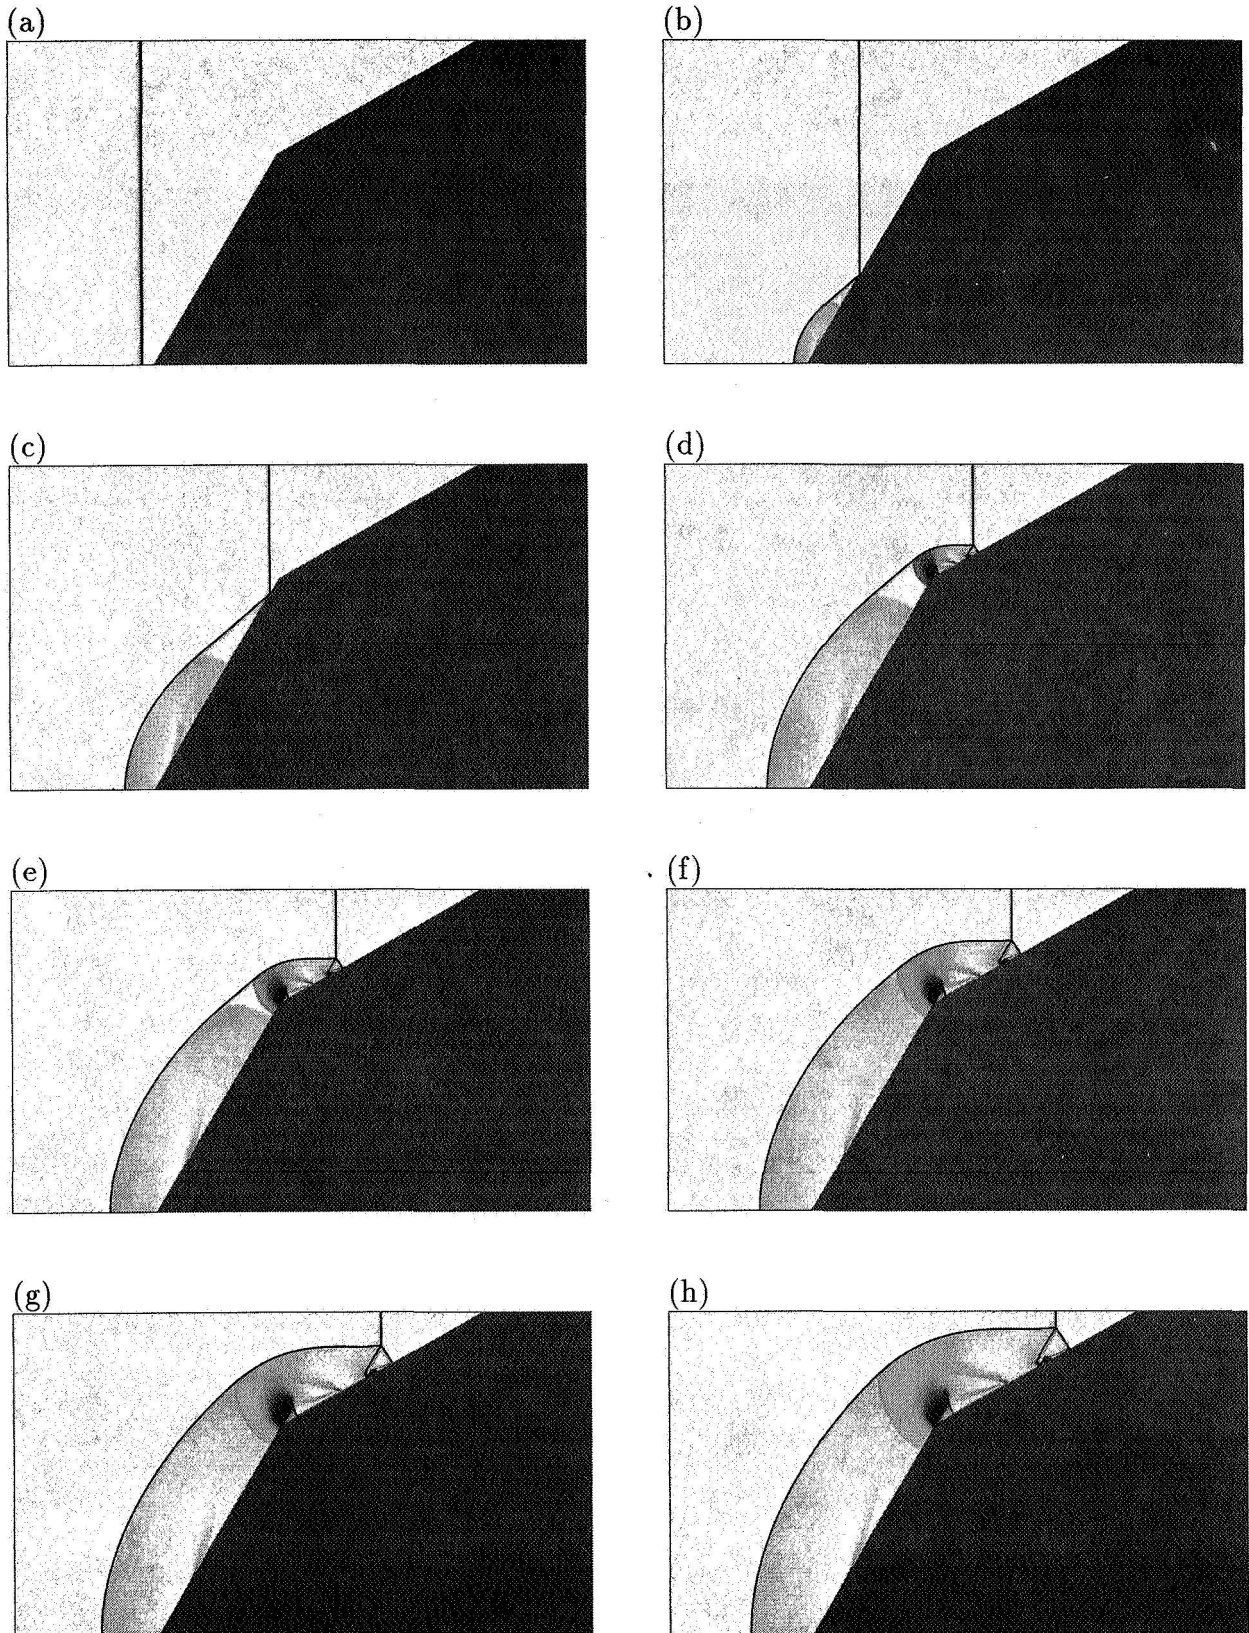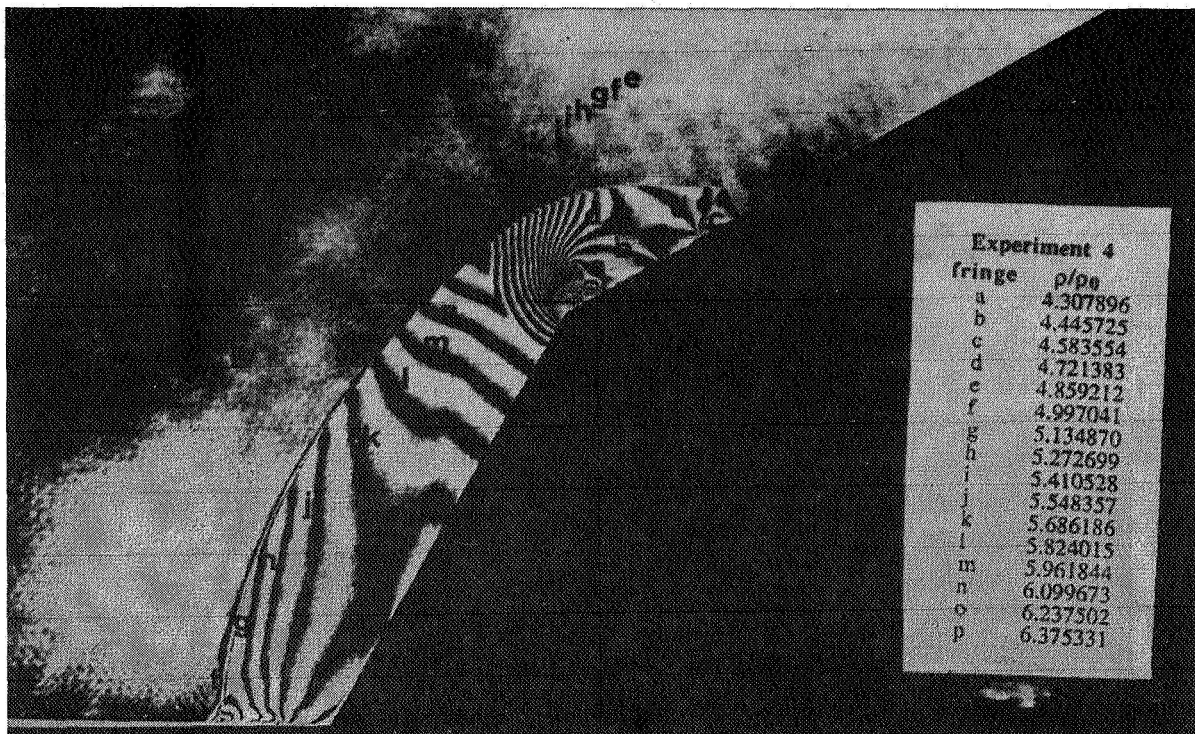
Figure 6: Sequence of schlieren snapshots from the simulation of Experiment #4.

(a) Experimental Interferogram, courtesy of Prof. Takayama
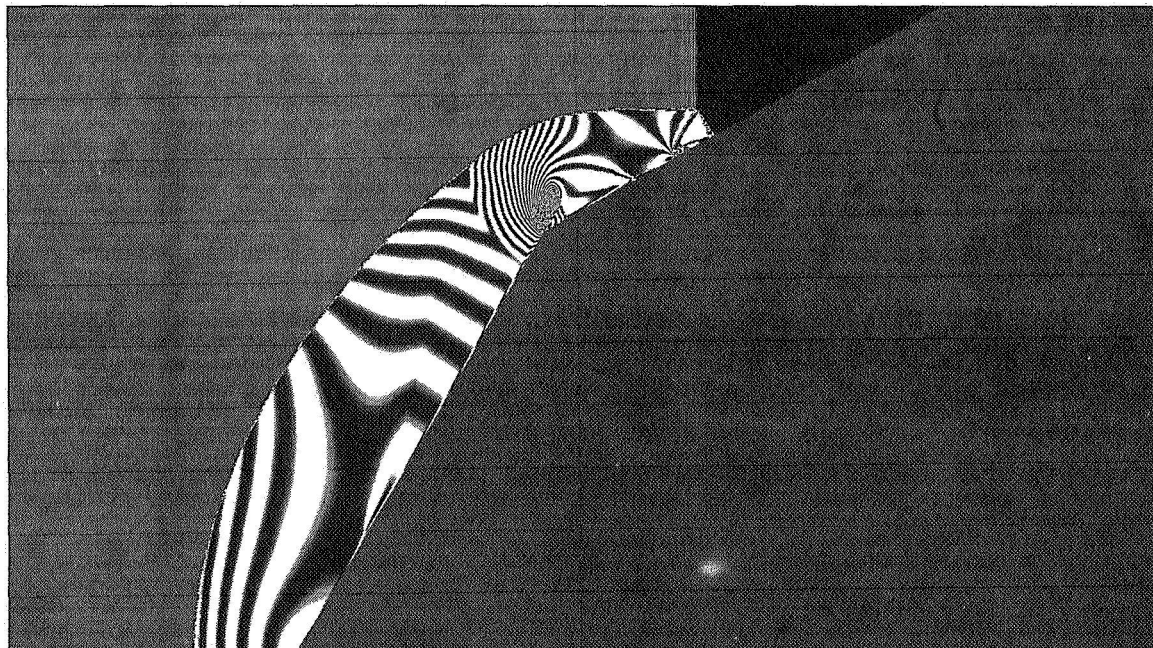

(b) Numerical Interferogram

Figure 7: Comparison between numerical and experimental interferograms for Experiment #4.

DISCUSSION

Here we restrict ourselves to making some specific observations about the development of mesh refinement methods for investigations of unsteady shock wave phenomena, and the reader who is unfamiliar with the basic techniques of mesh refinement is directed to[9, 17].

The majority of mesh refinement schemes give the impression of having been designed solely to minimize the number of grid cells that are required to compute a solution of a given resolution or accuracy. This design philosophy is presumably based on the notion that the effort required to integrate a discretized flow solution decreases as the number of grid cells decreases. But the following example demonstrates that the number of grid cells can have surprisingly little bearing on the cost of performing a time-dependent simulation and so this particular design philosophy is flawed.

Consider the propagation of a shock down a uniform mesh of N cells, each of width Δx. If a uniform time step is chosen such that the Courant number based on the speed of the shock is one (hence the shock traverses one cell per time step), it will take N integrations of N cells for the shock to pass through the domain i.e N^2 cell integrations. Now halve one cell in the grid such that there are $N-1$ cells of width Δx and two of width $\Delta x/2$. Again if a uniform time step is used to propagate the shock through this domain, without violating the CFL condition it will take $2N$ integrations of $N+1$ cells to propagate the shock through the domain i.e. $2N^2 + 2N$ integrations. Therefore although but a single cell has been added to the grid the cost of the simulation has more than doubled. Thus for time-dependent problems it is desirable to refine in time as well as space[10]. Here, using temporal refinement, the two small cells would be integrated $2N$ times and the other $N-1$ cells would be integrated N times as in the uniform mesh case i.e. a total of $N^2 + 3N$ integrations. Thus, for N reasonably large, the cost of the refinement becomes negligible. As an alternative to temporal refinement one could conceivably opt for an integration scheme which was stable for large Courant numbers, but for highly non-linear problems the loss in temporal accuracy would probably prove unacceptable.

A temporal refinement strategy is easily incorporated into hierarchical refinement schemes such as those based on quad-trees (e.g. [4]) or embedded patches (e.g. [3, 10]) since it is possible to avoid ever having to interpolate across discontinuities[10]. However, a temporal refinement strategy seems ill-suited to refinement schemes based on unstructured triangular meshes (as typified by[6]), at least when combined with a shock-capturing methodology, since one cannot avoid having to perform awkward non-linear interpolations at discontinuities. Such interpolations are unlikely to satisfy a shock-capturing scheme's unique smeared shock profile and so would result in spurious oscillations[10]. One convenient way around this difficulty would be to employ an integration scheme based on floating shock-fitting[7, 16] rather than shock-capturing. Then there would be no smeared discontinuities and the cause of the problem disappears. This strategy illustrates an important feature of the design of mesh refinement methods. It is often better to work around difficulties than to attempt to effect a cure. A refinement scheme contains many components and the best schemes seem to be those whose components work symbiotically.

Leaving aside the issue of temporal refinement, minimizing the number of grid cells will not automatically lead to an efficient method of refinement. Consider the case of an isolated discontinuity which runs oblique to the grid. It is clear that cellular quad-tree refinement (say [4]) is more efficient than embedded patch refinement (say [10]) in terms of the number of cells each method requires to

C-2

tile the discontinuity. However, it also has the larger storage overheads per mesh cell of the two associated data structures. For inert shock wave simulations which need only a small number of levels of refinement the storage overheads from quad-tree refinement are easily tolerated, but this might not be the case if the flow contained chemical reaction. Instead of a shock oblique to the grid consider a detonation wave which in addition to a shock front has some internal structure, albeit on a very fine scale, which must be resolved and cannot be captured. In this instance one would need a wide swathe of cells to cover the reaction zone which might be ten or more levels of refinement down in the quad tree because of the disparateness between the width of the reaction zone and the distance over which the detonation wave needs to be propagated. Therefore although the cells in the swathe are close to one another spatially they could lie far apart in the quad tree structure, which might impact on a parallel implementation of the scheme, and each cell would introduce a large overhead due to the accumulation of pointers down to its level in the data structure. Consequently embedded patch refinement might now prove to be more efficient because its storage overheads would be so much lower and it would better preserve the proximity of cells within the reaction zone.

Adaptive mesh refinement algorithms, unlike classical numerical methods, entail quite sophisticated software. Therefore arguments such as the one above must be tempered by the realization that specific implementation details can make or break an algorithm in terms of its practical performance. In particular the grid data structure needs to be well crafted. For example, the data storage needs to be flexible enough to cope with dynamic allocation and deallocation as local refinement is added and removed, and data accesses have to be efficient so as not to impact on performance. Now since it is all too easy to underestimate the level of commitment required to write, test and debug a pukka mesh refinement code, any newcomer would be well advised to take his or her own software skills in to account before choosing to code up any one particular method.

In fact the number of considerations that must be taken in to account before choosing a mesh refinement strategy are legion, even when one's needs are fairly specific. For example, our interests lie in investigating complex shock-wave phenomena, and given the results from the previous section it would appear that our refinement algorithm is well suited to our purposes. But suppose we were dissatisfied with the quality of our results for Experiment #2 (Figure 5) and wanted to perform a viscous simulation, would our scheme cope as well as in the inviscid case?

In the past the scheme has been used to perform viscous simulations of shock-boundary layer interactions[10], and so there is no reason to believe that it could not cope with a viscous simulation of Experiment #2. However, since viscous flow features tend to be anisotropic in nature, such a simulation would expose a weakness of our refinement scheme: it does not cope very well with anisotropic refinement. The method used[10] is basically limited to features like boundary layers which are affixed to solid surfaces. To refine a free shear layer which might happen to lie oblique to the mesh we would be forced to use isotropic refinement which would be needlessly expensive. This is an example where a change in the flow model can have a significant impact on the refinement efficiency, even though the application remains unchanged. Thus the correct choice of refinement strategy is never straightforward. To complicate matters even further, one cannot ignore the interplay between the method of refinement and the method of flow integration. For example, a triangular unstructured mesh has the geometric flexibility to allow for efficient anisotropic refinement but a certain amount of care must still be taken to generate meshes that are suitable for viscous simulations[8]. In general, depending on the application, one might wish to compromise the refinement efficiency so as to avoid compromising the accuracy of the flow integration (or vice versa). Of course the accuracy of a refinement scheme is, for the most part, ordained by the monitor functions which determine where refinement does or does not take place.

As is common practice we employ heuristic functions to determine where to refine, and coarsening takes place naturally by choosing not to refine and so involves no additional criteria[10]. For the present double wedge problems we used a combination of two monitor functions: density gradients were used to locate shocks and a local comparison between density and pressure gradients was used to locate contact discontinuities[13]. Now there are numerous reasons why this type of heuristic approach is unsatisfactory, not least of which is that it introduces tunable parameters and so increases the experience factor needed to operate a refinement scheme reliably. As Warren *et al.*[18] have shown, a poorly constructed heuristic monitor function can cause a mesh refinement scheme to home in on an incorrect solution. But this can happen with any refinement function, heuristic or not, which provides estimates for the local error without also providing estimates for how the local error affects the global error i.e. every refinement function in common use. To a large extent the mesh refinement community has been lulled into a false sense of security by the general experience that local errors are often benign. The test case discussed in [18] is a gentle reminder that small local errors can sometimes tip the balance and result in large global errors, but other more pathological examples are not difficult to find especially where chemical reaction is involved.

Figure 8 (a) shows a trace of the pressure behind the lead shock front of a one-dimensional detonation wave which exhibits a galloping instability[14]. By normal standards this computation would have been thought to be well resolved since 160 mesh points covered the so-called reaction half-length (giving some 256,000 cells over the time period shown) when contemporary simulations have ten or less points in the reaction half-length. However, when the simulation was repeated with the grid spacing halved, the dynamic behaviour of the detonation wave altered dramatically, see Figure 8 (b). At first glance one might assume that Figure 8 (b) came from the coarser computation since it looks more dissipative in that a two mode pulsation is decaying to a single mode pulsation. But in fact it is the extra dissipation in Figure 8 (a) that sustains a spurious two mode pulsation whereas the correct behaviour should be that of a two mode pulsation with a time-attractor limit cycle[14] i.e. Figure 8 (b). Interestingly the difference in behaviour arises not from an error in resolving the detonation shock front, but from a failure to resolve an innocuous part of the reaction zone which is smooth.
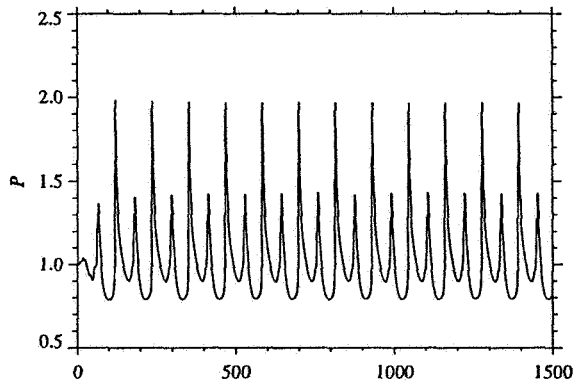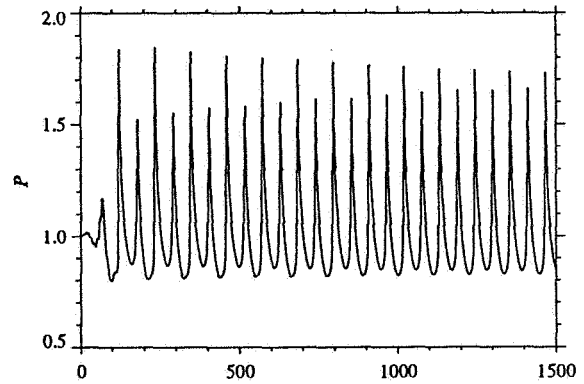
(a) 160 pts/$L_{\frac{1}{2}}$ (b) 320 pts/$L_{\frac{1}{2}}$

Figure 8: Variation in the computed pressure history trace for a galloping detonation wave when the mesh spacing is halved[14].

Clearly there is much room for improvement in the current crop of criteria used to control refinement. However, any attempts at devising rigorous mathematically based refinement criteria should not ignore certain practicalities. For example, in our simulations it can be necessary to adapt the grid tens of thousands of times[13] and so the method of determining where to refine must be reasonably cheap otherwise it would cripple the simulation. Also, the physical scales involved are so disparate that one cannot afford the luxury of periodically comparing the solution computed with refinement against that computed on a uniform mesh of the same high resolution, as is effectively done in[5], because this would require an unrealistic amount of storage.

For practical purposes the lack of a fool-proof refinement criteria does not undermine the usefulness of adaptive mesh refinement schemes for investigating shock wave phenomena, but it does complicate matters. Whenever we start to investigate a new problem we perform a sensitivity study to see how the computed results vary with, amongst other things, the effective resolution of the computational grid as controlled by our chosen refinement criteria. Thus we tool-up to a position where we think we can produce a reliable simulation. Note we would do more or less the same thing even if we were not employing mesh refinement, as past performance is no real guide as to how a numerical scheme will fair on a new problem.

For serious investigations the cost of tooling is generally spread over a parameter study and so is not excessive. The only drawback we find is that the results from sensitivity studies are rarely as conclusive as we would like. Many shock wave phenomena exhibit physical instabilities and so the notion of a grid converged solution is not always clear, or even appropriate since the flow model might preclude the possibility of having a sensible solution in the limit of the mesh spacing going to zero. For example, in [12] we presented results for the vortex sheet produced by a shock wave diffracting over a knife edge. These results show that an inviscid simulation can reproduce the correct behaviour and yet provide no limiting solution since the numerical dissipation which controls the fine scale structure of the vortex sheet, in the absence of physical viscosity, never bottoms out as the grid is refined. On the other hand, in some of our simulations of detonation phenomena it is clear that we are incapable of reaching a fully converged solution either because the physical scales are too disparate for our computing resources or the physical behaviour of the system is non-deterministic in that variations in discretization errors, no matter how small, lead to significant variations in dynamical behaviour.

Most CFD simulations are performed with the aim of producing quantitative answers to well understood problems, in which case the above vagaries are abhorrent. However, much of our work is performed in an attempt to fathom behaviour which is not known and simulations are used as a qualitative diagnostic and so a certain amount of subjectivity cannot be avoided. In short we use our mesh refinement algorithm to perform simulations which are more detailed than would otherwise be possible. Consequently we close this discussion without making any attempt to sell our scheme in terms of how efficiently it was able to compute the workshop double wedge problem. While this might be viewed as contrary we would argue that any results we could present would have little practical value: by comparison to our recent studies[13] the present simulations are so cheap as to be almost inconsequential. Moreover it should be appreciated that the cost of performing a time-dependent simulation can pale into insignificance when compared to the time taken to decipher the results, and to bandy performance figures would lose sight of the fact that our scheme has progressed well beyond the development stage and is used as an everyday tool.

CLOSING COMMENTS

At times adaptive mesh refinement appears to be more of an art than a science, therefore on a self-indulgent note we close with two quotations that sum up our thoughts on this branch of computational fluid dynamics.

The first quotation is taken from Shakespeare's *Twelfth Night* and explains why we feel there will always be a plethora of refinement schemes: "some [methods] are born great" in that they are so well suited to a particular class of problem they do not deserve to be replaced by some monolithic refinement scheme; "some [methods] achieve greatness" when they leapfrog the field by virtue of being able to exploit some new generation hardware feature and so methods tend to pass in and out of fashion; "and some [methods] have greatness thrust upon 'em" in that many fluids researchers cannot develop their own refinement code and must make do with whatever is available, so schemes that should be put to rest will not die by dint of their users.

Our second quotation is attributed to the son of the author Alexandre Dumas: "All generalizations are dangerous, even this one." In this paper we have tried to emphasize that context is all important where mesh refinement is concerned. Therefore whilst the subject is sorely in need of some formalism to guide us out of the present heuristic quagmire, there needs to be a realization that not all needs are the same. As we have shown, following rigorous criteria which are misplaced can prove disastrous. Therefore if some of our observations appear provocative it is only because we are attempting to correct an imbalance, as we see it, in current thinking. Grid convergence is a case in point. While rigorous Mathematical concepts of convergence are unambiguous, the practical concept of a grid converged solution to an unsteady problem, where the flow might be physically unstable, is hazy to say the least. And so common ground must be found between theoreticians and the practical exponents of mesh refinement before any real progress can be made in eliminating the heuristic elements from today's algorithms.

Acknowledgements

REFERENCES

[1] M. J. BERGER, *Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations*. Ph.D. thesis, Computer Science Dept., Stanford University (1982).

[2] M. J. BERGER AND J. OLIGER, *Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations*. J. Comput. Phys., **53**(1984), pp. 482–512.

[3] M. J. BERGER AND P. COLELLA, *Local Adaptive Mesh Refinement for Shock Hydrodynamics.* J. Comput. Phys., **82**(1989), pp. 67–84.

[4] Y.-L. CHIANG, B. VAN LEER AND K.G. POWELL, *Simulation of Unsteady Inviscid Flow on an Adaptively Refined Cartesian Grid.* AIAA Paper 92-0443, 1992.

[5] A. HARTEN, *Multiresolution Algorithms for the Numerical Solution of Hyperbolic Conservation Laws.* UCLA CAM Report 93-03, 1993.

[6] R. LÖHNER, K. MORGAN AND O. ZIENKIEWICZ, *Adaptive Grid Refinement for the Compressible Euler Equations.* In "Accuracy Estimates and Adaptivity for Finite Elements", Wiley, 1994.

[7] M. PARASCHIVOIU, J.-Y. TRÉPANIER, M. REGGIO AND R. CAMARERO, *A Conservative Dynamic Discontinuity Tracking Algorithm for the Euler Equations.* AIAA Paper 94-0081, 1994.

[8] S. PIRZADEH, *Unstructured Viscous Grid Generation by the Advancing-Layers Method.* AIAA J., **32**(1994), pp. 17–19.

[9] K. G. POWELL, P. L. ROE AND J. J. QUIRK, *Adaptive-mesh Algorithms for Computational Fluid Dynamics.* pp. 303–337, Algorithmic Trends in Computational Fluid Dynamics, Edited by M. Y. Hussaini, A. Kumar and M. D. Salas. Springer-Verlag, New York, 1993.

[10] J. J. QUIRK, *An Adaptive Grid Algorithm for Computational Shock Hydrodynamics.* Ph.D. thesis, College of Aeronautics, Cranfield Institute of technology (1991).

[11] J. J. QUIRK, *An Alternative to Unstructured Grids for Computing Gas Dynamic Flows Around Arbitrarily Complex Two-Dimensional Bodies.* Computers & Fluids, **23**(1994), pp. 125–142.

[12] J. J. QUIRK, *A Cartesian Grid Approach with Hierarchical Refinement for Compressible Flows.* pp. 200–209, Computational Fluid Dynamics '94,(Invited Lectures and Special Technological Sessions of the Second European Computational Fluid Dynamics Conference 5-8[th] September 1994, Stuttgart, Germany), Edited by S.Wagner et al., John Wiley and Sons Ltd., Chichester, 1994.

[13] J. J. QUIRK, *A Parallel, Adaptive Grid Algorithm for Computational Shock Hydrodynamics.* To appear in Applied Numerical Mathematics.

[14] M. SHORT AND J. J. QUIRK, *Evolution of a One-dimensional Pulsating Detonation Wave Driven by a Model 3-step Chain-branching Reaction.* Submitted.

[15] K. TAKAYAMA, O. ONODERA AND G. BEN-DOR, *Holographic Interferometric Study of Shock Transition over Wedges.* SPIE, **491**(1984), pp. 976-983.

[16] J. VAN ROSENDALE, *Floating Shock Fitting via Lagrangian Adaptive Meshes.* ICASE Report No. 94-89, 1994.

[17] J. F. THOMPSON AND N. P. WEATHERILL, *Aspects of Numerical Grid Generation: Current Science and Art.* AIAA Paper 93-3539, 1993.

[18] G. WARREN, W. K. ANDERSON, J. THOMAS AND S. KRIST, *Grid Convergence for Adaptive Methods.* AIAA Paper 91-1592, 1991.