

COVER SHEET

NOTE: This coversheet is intended for you to list your article title and author(s) name only—this page will not appear on the CD-ROM.

Title: *Three-Dimensional Parallel Adaptive Mesh Refinement Simulations of Shock-Driven Turbulent Mixing in Plane and Converging Geometries* for Proceedings of the 21st International Conference on Parallel Computational Fluid Dynamics (ParCFD 2009)

Authors: Manuel Lombardini, Ralf Deiterding

PAPER DEADLINE: *****NOVEMBER 25, 2009****

PAPER LENGTH: ****16 PAGES (Maximum)****

ABSTRACT

This paper presents the use of a dynamically adaptive mesh refinement strategy for the simulations of shock-driven turbulent mixing. Large-eddy simulations are necessary due to the high Reynolds number turbulent regime. In this approach, the large scales are simulated directly and small scales at which the viscous dissipation occurs are modeled. A low-numerical centered finite-difference scheme is used in turbulent flow regions while a shock-capturing method is employed to capture shocks. Three-dimensional parallel simulations of the Richtmyer-Meshkov instability performed in plane and converging geometries are described.

FLOW CONFIGURATION

Acceleration-induced mixing of fluids, commonly observed in natural phenomena such as supernova collapse or in technologies involving supersonic combustion, originates from the initial mis-alignment of the pressure gradient at existing shocks with the local density gradients. The underlying fluid instability acting in this problem is known as the Richtmyer-Meshkov instability (RMI), and was first studied numerically by Richtmyer [1] and confirmed experimentally by Meshkov [2]. The RMI fundamentally occurs when a perturbed interface separating two fluids of different densities is accelerated impulsively by a shock wave depositing baroclinic vorticity at the interface. It is often thought of as an impulsive version of the Rayleigh-Taylor instability. The mixing between the fluids is enhanced when the mixing layer is processed by additional compressible waves such as shock waves and expansion fans.

Description of the Problem

The reference problem consists of a thin air/SF₆ plane interface in a light-to-heavy configuration, initially impacted by a planar incident shock of Mach number M_I and then reshocked after reflection of the transmitted shock off the end-wall of the shocktube [3]. Besides, a canonical simulation of the RMI in a 90° wedge has

Manuel Lombardini, Graduate Aeronautical Laboratories, California Institute of Technology, Pasadena, California 91125, U.S.A.

Ralf Deiterding, Oak Ridge National Laboratory, P.O. Box 2008 MS6367, Oak Ridge, Tennessee 37831, U.S.A.

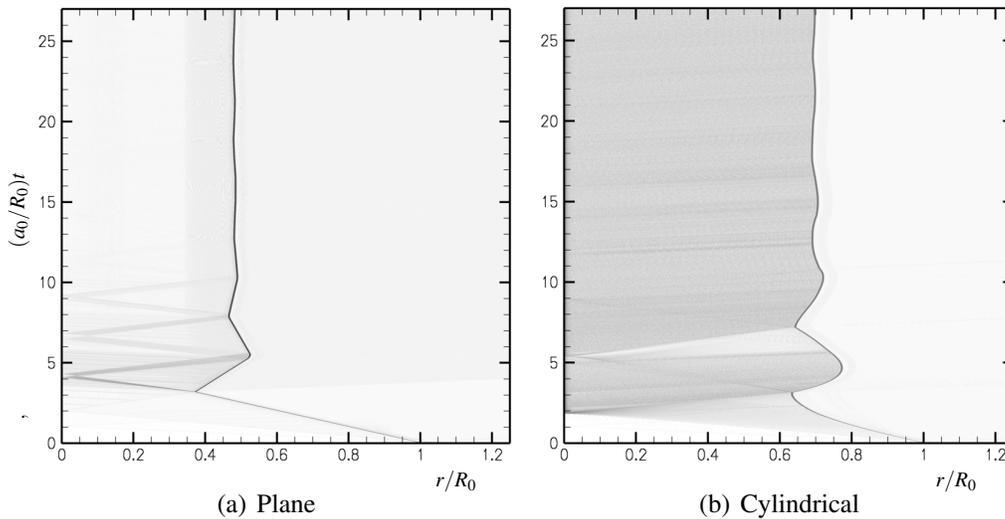


Figure 1. Wave diagrams for light-to-heavy air/SF₆ $M_I = 1.2$ -shock interactions.

been set up in parallel to focusing shocktube experiments currently conducted by the group of Prof. Dimotakis at GALCIT [4]. An imploding cylindrical shock impacts a perturbed, cylindrically-shaped density interface that separates light air (outside) from heavy SF₆ (inside). The transmitted shock converges down the wedge, reflects off the z -axis, and reshocks the interface, initiating a strong turbulent mixing between light and heavy fluids, similar to the plane case.

A study of the wave diagrams describing the unperturbed shock-contact interaction with reshock shows differences between the two geometries, as presented in Figure 1. In plane geometry, the heavy-to-light reshock refraction produces a reflected expansion wave traveling towards the end-wall and reverberating between the mixing layer and the end-wall. The successive echoing expansion waves of decreasing intensity deposit baroclinic vorticity over the course of their passage across the mixing layer from head to tail. In converging geometry, a shock traveling towards the end-wall is formed after the reflected expansion fan is generated by the first reshock interaction. After reflection off the end-wall, this shock re-interacts a second time with the layer. This sequence is repeated in a seemingly self-similar way, generating vorticity at the layer impulsively.

Initial and Boundary Conditions

The initial contact discontinuity deformation is the result of small-wavelength disturbances superposed to a smaller-amplitude long-wavelength symmetry-breaking perturbation in the directions transverse to the main flow (x -direction of propagation of the plane shock or r -direction of propagation of the converging shock). In plane geometry, the y - and z -disturbances wavelengths are comparable. However, because of the geometric contraction of the flow, azimuthal and axial disturbances at the cylindrical interface are initially chosen such that they are of comparable wavelengths at the reshock time.

The plane shocktube is constituted of a reflecting end-wall, and periodic sidewalls in the transverse directions of the shocktube center line. In the converging simulation, reflecting boundaries form both wedge walls and the wedge axis extends in the periodic z -direction. Thanks to the ghost fluid method [5], an inner cylindrical

wall of radius $r = r_{in}$ is added to regularize the wedge corner, where flow quantities can become singular. Assuming that shock-wave/boundary-layer interactions do not play a dominant role in the growth of the mixing zone, slip boundary conditions are applied at all reflecting walls. Consistent inflow boundary conditions are prescribed at the open end of the shocktube/wedge.

COMPUTATIONAL PROCEDURE

The flow presented in the previous section exercises both the large-eddy simulations (LES) and the shock-capturing features of the solver with dynamically adaptive meshes.

Large-Eddy Simulations

The reshock process produces a large dynamical range of turbulent scales, necessitating the use of LES. The stretched-vortex subgrid-scale (SGS) model of Misra & Pullin [6], extended to compressible flows by Kosovic *et al.* [7] and subgrid scalar transport by Pullin [8], is based on an explicit structural modeling of small-scale dynamics. The model utilizes stretching vortices as the essential subgrid element in the closure of Favre-filtered Navier-Stokes equations by providing the subgrid momentum stress τ_{ij} , the subgrid turbulent temperature flux q_i^T , and the mixture fraction flux q_i^Ψ . The model parameters required are determined entirely locally without the traditional use of spatial averaging. This model also enables a computational paradigm for multi-scale LES that extends estimates of some turbulent statistics from the resolved cutoff-scale to the Kolmogorov and Batchelor scales.

The resolution requirements imposed by the flow physics vary greatly both spatially and temporally for these simulations. For example, different key features such as shock waves of different strengths and turbulent mixing regions (as seen on Figures 5 and 4) need more resolution than other smoother regions of the flow. This is provided presently through LES within the AMROC framework [9] and based on the structured adaptive mesh refinement algorithm (AMR) of Berger & Oliger [10]. Discrete conservation of mass, momentum, and energy is accomplished by using a flux-based conservative finite-difference approach [11].

Hybrid Shock-Capturing/Low-Numerical Dissipation Scheme

The numerical method is formulated for Cartesian uniform grids, and is applied to each subgrid of the mesh hierarchy. It is an extension of the hybrid method by Hill & Pullin [12] to structured AMR meshes. A weighted, essentially non-oscillatory (WENO) scheme is used to capture discontinuities (such as shock waves or fine/coarse mesh interfaces) but switches to a low-numerical dissipation, explicit, tuned center-difference scheme (TCD) in the smooth or turbulent regions, optimal for the functioning of explicit LES such as the SGS stretched-vortex method. To ensure discrete numerical stability of the inviscid terms (momentum, scalar and energy convection terms), the centered discretization are written in a stable, energy preserving (skew-symmetric) formulation adapted to compressible flows [13].

For the subgrid activity to be correctly computed, thereby assuring the quality of the LES, the use of WENO is restrained to regions containing shock waves only. Switching between WENO and TCD has been optimized using a detection criterion based on Lax's entropy condition [14].

Adaptive Mesh Refinement

The core idea of the AMR method [10] is to cluster regions that require refinement into non-overlapping rectangular subgrids (patches) $G_{l,m}$ which define the domain of an entire level $l = 0, \dots, L$ by $G_l := \bigcup_{m=1}^{M_l} G_{l,m}$. The refinement process proceeds recursively and constructs a hierarchy of subgrids successively contained within the next coarser level domain. In contrast to cell-based mesh adaptation techniques the patch-based approach does not require a special coarsening operations; subgrids are simply removed from the hierarchy. The coarsest possible resolution is thereby restricted to the level 0 grid. Usually, it is assumed that the spatial and temporal discretization on a particular level are refined by the same factor, thereby ensuring the validity of a basically unaltered CFL-type stability condition on all levels of the hierarchy.

The numerical update is applied on each level by calling a single-grid routine implementing the finite volume scheme in a loop over all the subgrids $G_{l,m}$. The regularity of the input data allows a straightforward implementation of the scheme and further permits optimizations to take advantage of high-level caches, pipelining, etc. New refinement grids are initialized by interpolating the vector of conservative quantities from the next coarser level. However, data in cells already refined is copied directly from the previous refinement patches. *Ghost* or *halo* cells around each patch are used to decouple the subgrids computationally. Ghost cells outside of the root domain G_0 are used to implement physical boundary conditions. Ghost cells in G_l have a unique interior cell analogue and are set by copying the data value from the patch where the interior cell is contained (synchronization). For $l > 0$, internal boundaries can also be used. If recursive time step refinement is employed, ghost cells at the internal refinement boundaries on the level l are set by time-space interpolation from the two previously calculated time steps of level $l - 1$. Otherwise, spatial interpolation from the level $l - 1$ is sufficient.

Parallelization Strategy

The parallelization strategy implemented in AMROC is a rigorous domain decomposition approach. On a parallel machine with P identical nodes the root domain G_0 is split into P non-overlapping portions by $G_0 = \bigcup_{p=1}^P G_0^p$. The key idea now is that all higher level domains are required to follow the decomposition of the root level, i.e. $G_l^p := G_l \cap G_0^p$. The work on an arbitrary subdomain $\Omega \subset G_0$ is heuristically estimated with the expression $\mathcal{W}(\Omega) = \sum_{l=0}^{l_{\max}} [\mathcal{N}_l(G_l \cap \Omega) \prod_{\kappa=0}^l r_{\kappa}]$. Herein, $\mathcal{N}_l(\cdot)$ denotes the total number of finite volume cells on level l in a given domain and the product accounts for the recursive time step refinement. Note that minor workload imbalances due to the hybridization of WENO with the TCD scheme are currently neglected. Decompositions G_0^p with similar workload are found at runtime as the hierarchy evolves with a partitioning algorithm based on a generalization of Hilbert's space-filling curve [15]. The space-filling curve defines an ordered sequence on the cells of the root level that can easily be split in load-balanced portions. As such curves are constructed recursively, they are locality-preserving and lead to moderate data redistribution and parallel synchronization costs.

As a scalability test for AMROC we consider a 3-level simulation of a confined explosion problem using only the WENO scheme. The time per base level iteration of the most expensive operations is displayed in Figure 2(a). As it can be expected with the chosen parallelization approach, linear speed-up is achieved for the block-based numerical update of the finite volume scheme ("Integration"), while the scaling

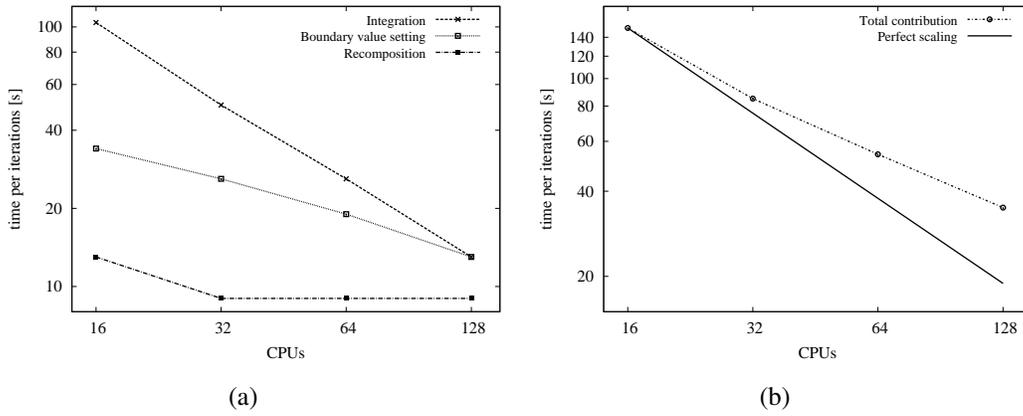


Figure 2. CPU scalability test: 3D adaptive computation with 3 grid levels (coarse grid of $30 \times 30 \times 30$ cells, uniformly refined grid of $120 \times 120 \times 120$ cells (1.73 Mcells)) of a circular shock-wave expanding in an enclosed box. The reflected shocks interact in a complex manner. (a) represents the breakdown of computational costs in second per iteration, (b) the total time per iteration.

of the operations “Boundary value setting” and “Recomposition” of the refinement hierarchy, that involve parallel communication, is less optimal. We have found the parallel performance of the overall algorithm (cf. Figure 2(b)) suitable for effective parallel computations on several hundred processors.

COMPUTATIONAL RESULTS

To illustrate the need of AMR for the RMI, consider our larger converging cylindrical simulation ($M_I = 2.0$ incident cylindrical shock as it impacts the cylindrical interface). The simulation is conducted in a wedge-like geometry of angle $\pi/2$. The computational domain, in cylindrical coordinates (r, θ, z) is:

$$r_{in} < r < r_{out}, \quad 0 < \theta < \pi/2, \quad 0 < z < R_0 \quad (1)$$

The inner cylindrical reflecting wall of radius $r_{in} = 0.04R_0$ and the outer radial distance $r_{out} = 0.94\pi/2R_0$, where R_0 is the location of the initial shock-contact interaction. There are 7 initial perturbation modes in the z -direction and 6 in the θ -direction.

Flow History

Figure 3 depicts the three-dimensional evolution of mixing zone throughout the simulation through the levels of mass fraction (blue for light portions of fluid, red for heavy ones, white for equal mass fraction of light and heavy fluids). In Figure 4, two-dimensional Schlieren images in azimuthal and axial views are arranged together to show the density gradients across the mixing layer and the compressible waves traveling within the domain. From these figures, we recall the different stages in the mixing layer growth. Subfigures (a) show the state of the flow following the initial light-to-heavy shock refraction. A transmitted shock is produced, followed by the interface accelerating towards the center. Growing mushroom-like structures characteristic of baroclinic instabilities are visible in the azimuthal and axial directions. The imploding shock reflects off the axis of the wedge and reshock the contracting

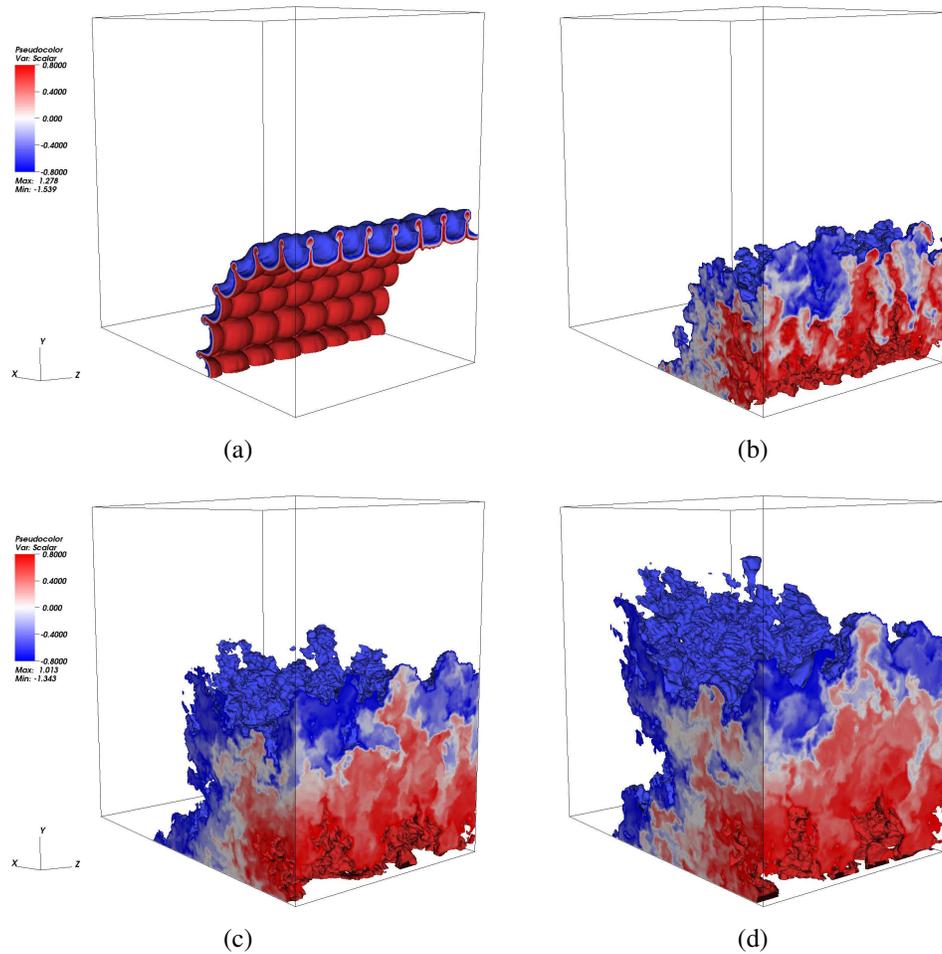


Figure 3. Light-to-heavy air/SF₆ $M_I = 2.0$ -converging cylindrical RMI: levels of mixed fluid across the mixing zone (white for regions of 50% mass fraction) (a) after the first imploding shock interaction (flow contraction); (b) after the first exploding reshock (flow expansion); (c) at about than twice the reshock time and (d) at late time once most compressible waves have existed the computational domain.

interface, inverting its motion. The energy deposited by the reshock is such that the growth of the mixing layer is further increased. As shown in Subfigures (b), as the fluids begin to interpenetrate, the spectrum of physical scales involved in the flow dynamics broadens: large scales become larger while smaller scales appear. Subfigures (c) and (d) describe late stages in the growth of the mixing region. As can be inferred from Subfigure 4(c), the density gradients are reduced and the flow can be considered quasi-incompressible as only weakly compressible waves remain in the wedge. Portions of heavy fluid penetrating into light fluid (also known as spikes) and light fluid into heavy fluid (or bubbles) can still be distinguished. Ultimately (from Subfigures (d) and at later times), the flow becomes fully turbulent (see also spectral analysis in Subsection 5.2). The mixing zone extends in the axial and radial directions, with no recollection of the orderly azimuthal and axial modes initially present, and the largest scales produced become comparable to the domain size.

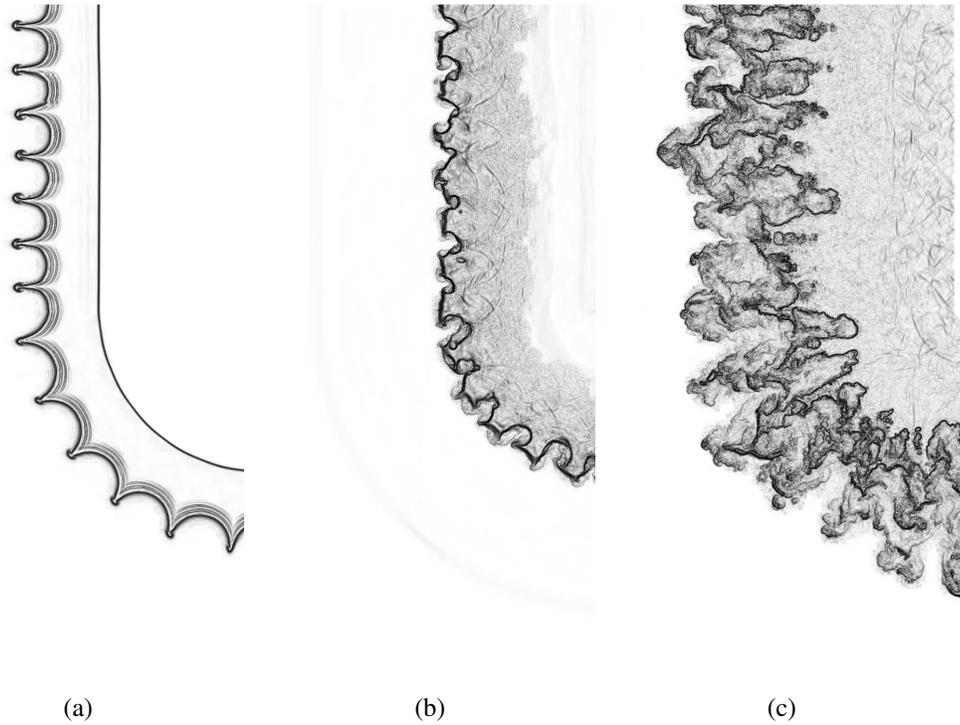


Figure 4. Light-to-heavy air/SF₆ $M_I = 2.0$ -converging cylindrical RMI: azimuthal view (plane slice $z = 0$) and axial view (plane slice $x = 0$) of Schlieren density gradient photographs showing shock waves and density stratifications (a) after the first imploding shock interaction (flow contraction); (b) after the first exploding reshock (flow expansion) and (c) at about twice the reshock time.

Computational Details

In the displayed computation, the domain is discretized with $95 \times 95 \times 64$ cells on the base grid with three additional levels of refinement based on the local density gradient. The refinement ratio between each level is equal to two for all levels and directions, and the subgrid cutoff scale is set to that of the finest mesh. As the flow evolves, the distribution of the AMR hierarchy to different processors is adjusted dynamically to balance the work and all parallel data structures are automatically rearranged. The simulation was performed using 32 AMD Opteron 2.5 GHz-quad-processor nodes (16 GB memory each) and consumed about 70,000 h CPU time. The cell count varies from a minimum of approximately 10 million cells in the early times of the simulation, when the mixing zone has not yet radially expanded, to a peak of around 140 million at late times. AMR reduces the computational expenses compared to the equivalent finest unigrid $760 \times 760 \times 512$ problem that would have used approximately up to 3 times more storage and taken more than 3 times longer to complete. In Figure 5, the mesh adaptation and distribution to the 128 available processors is displayed for two exemplary snapshots in time. The graphics in the left column show the mixing zone (white for regions of 50% mass fraction) on background planes that represent the domains of different mesh refinement in different gray shades (the 4 levels are displayed), the right graphics display the domains of different processors by color.

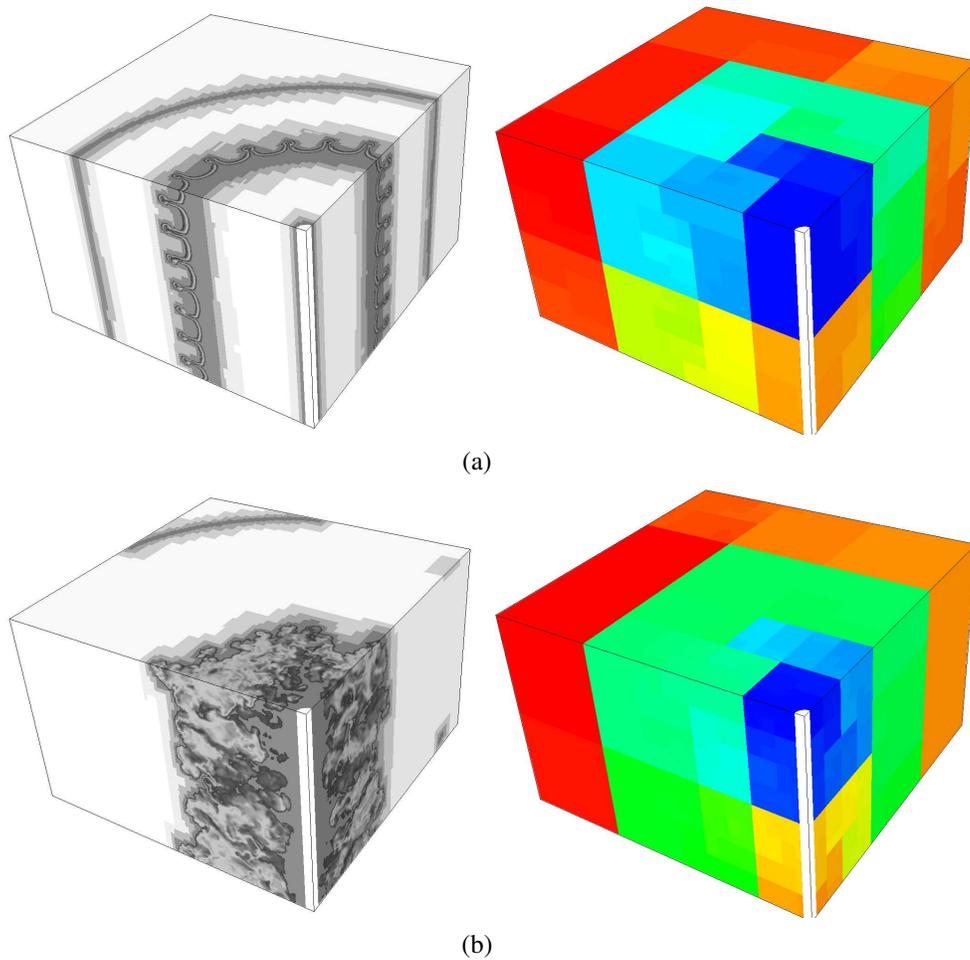


Figure 5. Mesh adaptation and processor distribution for the light-to-heavy air/SF₆ $M_I = 2.0$ -converging cylindrical RMI: (a) after the first imploding shock interaction (flow contraction) and (b) after the first exploding reshock (flow expansion).

DATA ANALYSIS

A detailed quantitative analysis of results can however be obtained from post-processing on saved parallel data files that are read in on smaller CPU count. Owing to the symmetry of the flow for each geometry, the data (at a given time) are scrutinized over surfaces normal to the main flow direction. In plane (resp. cylindrical) geometry, we consider plane (resp. cylindrical) surfaces normal to the direction of the incident shock and determined by their axial position x (resp. radius r).

Surface-Averaged Statistics

The investigation includes space-time histories of instantaneous plane/cylindrical surface-averages $\langle \cdot \rangle$ of diverse quantities Q , taken parallel/concentrically to the main shock. Typically, we need to evaluate surface-averages of about a hundred base quantities ($\langle \rho \rangle$, $\langle \rho^2 \rangle$, etc.) across the entire domain. For each time considered, the post-processing required around 1,000 h CPU for the largest simulations performed (cf. Subsection 3.3). From the base surface-averages, we define, for example in the

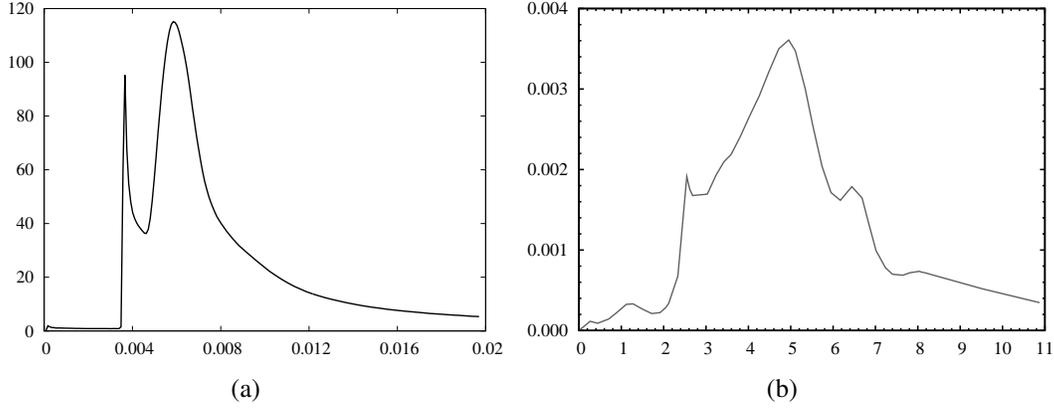


Figure 6. Light-to-heavy air/SF₆ $M_I = 1.2$ -RMI: resolved+subgrid volume-averaged TKE (in arbitrary units) vs. time (in ms) for the (a) plane and (b) cylindrical geometries.

cylindrical geometry, the perturbations with respect to the surface-average $\langle Q \rangle$ or the density-weighted surface-average \tilde{Q} :

$$Q'(r, \theta, z, t) = Q(\mathbf{x}, t) - \langle Q \rangle(r, t), \quad (2)$$

$$Q''(r, \theta, z, t) = Q(\mathbf{x}, t) - \tilde{Q}(r, t), \quad (3)$$

where

$$\tilde{Q}(r, t) = \frac{\langle \rho Q \rangle}{\langle \rho \rangle}, \quad (4)$$

with ρ total density field. From the perturbation field Q'' defined at a surface of radius r at time t , a density-weighted variance of Q follows:

$$\text{Var}_\rho(Q)(r, t) = \tilde{Q}^2 - \tilde{Q}^2 = \frac{\langle \rho Q^2 \rangle}{\langle \rho \rangle} - \frac{\langle \rho Q \rangle^2}{\langle \rho \rangle^2}. \quad (5)$$

Among important turbulent quantities, we focus on the turbulent kinetic energy (TKE), that is a result of a resolved-scale contribution $\langle \mathbf{K} \rangle$ and a subgrid (modeled) counterpart $\langle \mathbf{k} \rangle$:

$$\langle \mathbf{K} \rangle = \frac{1}{2} \text{Var}_\rho(u_i u_i), \quad (6)$$

$$\langle \mathbf{k} \rangle = \frac{1}{2} \frac{\langle \tau_{ii} \rangle}{\langle \rho \rangle}. \quad (7)$$

The resolved+subgrid TKE is then summed over the direction of anisotropy of the problem, leading to a measure of the volume-averaged TKE at a given time, whose evolution is shown in Figure 6. This figure summarizes the main events of the flow. First, the total amount of TKE deposited by the initial shock, visible as a small bump near $t = 0$, as well as that owing to the first reshock. The first reshock is followed by a steep decay of TKE and a second interaction of comparable energy to that of the reshock but with slower increase. The second interaction is clearly indicative of the first expansion wave interacting with the mixing layer (cf. Subsection 2.1). In

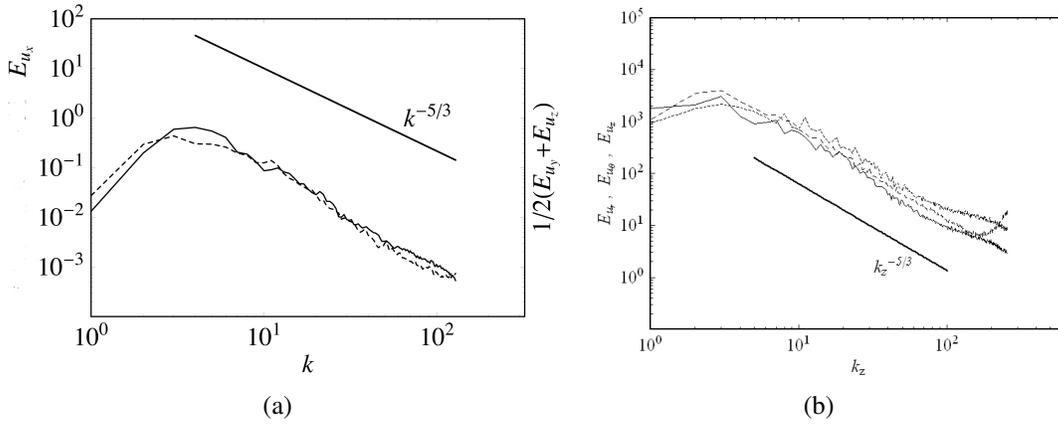


Figure 7. Light-to-heavy air/SF₆ $M_I = 1.2$ -RMI: late-time power spectra computed over a plane or cylindrical slice across the mixing zone center. For the plane geometry (a), k -power spectrum of the velocity component u_x (anisotropic contribution, solid line) and combined power spectrum of the components u_y, u_z (isotropic contribution, dashed line), with $k = \sqrt{k_y^2 + k_z^2}$ (all computed wavenumbers shown and $k_{max} = 128$). For the cylindrical geometry (b), k_z -power spectra of the velocity components u_r (anisotropic contribution, solid line), u_θ (small-dashed line), and u_z (long-dashed line), with $k = k_z$ (all computed wavenumbers shown and $k_{z,max} = 256$).

the converging flow, the shock forming behind the reflected expansion do not seem to have a significant influence on the TKE evolution. After about twice the reshock time, the TKE starts to slowly decay, since none of the secondary reverberations deposit enough energy across the mixing zone to sustain the turbulent activity.

Power Spectra

Diverse spectra, including velocity components, density, and scalar spectra, as well as other quantities such as Taylor and Kolmogorov microscales can be evaluated within the turbulent mixing zone. Consider first the plane geometry. Periodic boundary conditions in the (y, z) -cross-section of the shocktube allow for two isotropic directions within the mixing layer and the calculation of radial power spectra at a given instant in time. For an arbitrary field $f(x, y, z, t)$, we define the instantaneous radial power spectrum of f on the plane located at the center of the mixing zone $x = x_c(t)$, at the instant t by:

$$E_f(k; x_c(t)) = \frac{k}{2} \int_0^{2\pi} \left| \hat{f}(k, \beta; x_c(t)) \right|^2 d\beta, \quad (8)$$

where $\hat{f}(\cdot, \cdot; x_c(t))$ is the (y, z) two-dimensional Fourier transform of $f(x_c(t), \cdot, \cdot, t)$ in polar wavenumber space with k and β the radial and angular wavenumbers, related to the Cartesian wavenumbers $k_y = k \cos(\beta)$ and $k_z = k \sin(\beta)$. In the converging geometry, the spectra are defined over a surface of radius $r = r_c(t)$ slicing through the mixing zone, and associated to the wavenumber k_z , the z -axis being the only periodic direction of the flow. The Fourier coefficients are then averaged over the azimuthal direction of the flow. Results for both geometry are represented in Figure 7 from data across the mixing zone at late-time during the decay of TKE (many reshock times).

The inertial subrange with Kolmogorov scaling approaching $k^{-5/3}$ seems reached for each velocity components, in both geometries. Except for very large scales which are aware of the anisotropy of the main flow, the velocity spectrum of each component are comparable in the inertial subrange, which suggest that the turbulent mixing has reached some isotropy. Some anisotropy is detectable at the smallest scales and it is still to be determined whether or not it is an effect of the subgrid modeling.

CONCLUSIONS

We highlighted the use of AMR for shock-driven flows containing a strong direction of anisotropy (the axis of the straight shocktube or the radial direction of the wedge) and leading to late-time turbulent mixing. Surface-averaged statistics in the directions transverse to the direction of the base flow showed similarities and differences between the plane and converging geometries. Future work will consist in performing comparable simulations when the incident shock is traveling from the heavy fluid to the light one, in both geometries.

REFERENCES

1. R. D. Richtmyer, "Taylor instability in shock acceleration of compressible fluids," *Commun. Pure Appl. Math.*, 13, 297-319 (1960).
2. E. E. Meshkov, "Instability of the interface of two gases accelerated by a shock wave," *Soviet Fluid Dyn.*, 4(5), 101-108 (1969).
3. M. Vetter and B. Sturtevant, "Experiments on the Richtmyer-Meshkov instability of an air/SF₆ interface," *Shock Waves*, 4, 247-252 (1995).
4. C. Bond, D. J. Hill, D. I. Meiron, and P. E. Dimotakis, "Shock focusing in a planar convergent geometry: experiment and simulation," *article in press in J. Fluid Mech.*
5. R. P. Fedkiw, T. Aslam, B. Merriman, and S. Osher, "A non-oscillatory Eulerian approach to interfaces in multimaterial flows (the ghost fluid method)," *J. Comput. Phys.*, 152(2), 457-492 (1999).
6. A. Misra and D. I. Pullin, "A vortex-based model for large-eddy simulation," *Phys. Fluids*, 9(8), 2443-2454 (1997).
7. B. Kosovic, D. I. Pullin, and R. Samtaney "Subgrid-scale modeling for large-eddy simulations of compressible turbulence," *Phys. Fluids*, 14(4), 1511-1522 (2000).
8. D. I. Pullin, "A vortex-based model for the subgrid flux of a passive scalar," *Phys. Fluids*, 12(9), 2311-2319 (2000).
9. R. Deiterding, "Construction and application of an AMR algorithm for distributed memory computers," In Adaptive Mesh Refinement – Theory and Applications, T. Plewa, T. Linde, and V. G. Weirs, Eds., vol 41 of Lecture Notes in Computational Science and Engineering, 361-372 (2005).
10. M. J. Berger and J. Olinger, "Adaptive mesh refinement for hyperbolic partial-differential equations," *J. Comput. Phys.*, 53, 484-512 (1984).
11. D. J. Hill, C. Pantano, and D. I. Pullin, "Large-eddy simulation and multiscale modeling of a Richtmyer-Meshkov instability with reshock," *J. Fluid Mech.*, 557, 29-61 (2006).
12. D. J. Hill and D. I. Pullin, "Hybrid tuned center-difference-WENO method for large-eddy simulation in the presence of strong shocks," *J. Comput. Phys.*, 194, 435-450 (2004).
13. A. E. Honein and P. Moin, "Higher entropy conservation and numerical stability of compressible turbulence simulations," *J. Comput. Phys.*, 201, 531-545 (2004).
14. M. Lombardini, "Richtmyer-Meshkov instability in converging geometries," PhD thesis , Caltech (2008).
15. M. Parashar and J. C. Browne, "On partitioning dynamic adaptive grid hierarchies," in *Proc. of the 29th Annual Hawaii Int. Conf. on System Sciences* (1996).