

A simplified approach for simulations of multidimensional compressible multicomponent flows: The grid-aligned ghost fluid method

Nikolaos Bempedelis, Yiannis Ventikos*

Department of Mechanical Engineering, University College London, WC1E 7JE, United Kingdom

Abstract

In the present work the authors present a simplified formulation for the extension of a ghost fluid method in multidimensional space. In the proposed method, the Riemann problems at the interface are formulated along the grid rather than in a normal to the interface direction. The information that is required to construct these Riemann problems is acquired “on-the-fly” from the adjacent to the interface cells. With respect to existing multidimensional ghost fluid formulations, the method is computationally less expensive, as the procedures of determining ghost fluid regions, extending, interpolating and extrapolating variables and computing geometrical quantities are avoided. More importantly, it is markedly simple with respect to its implementation. By introducing the proposed formulation in a well-established front tracking framework we perform an extensive validation of the method and demonstrate that despite its simplicity it yields highly accurate results while remaining free of oscillations.

Keywords: Front tracking, Compressible flows, Multimaterial flows, Ghost fluid method

1. Introduction

The concept of interface localization is commonly employed in computations of flow systems with separated phases. The most well-established techniques for the localization of the interface are the Volume of Fluid [1], Level Set [2, 3] and Front Tracking [4, 5, 6, 7] methods. In front tracking, the interface is represented with connected marker points which form a curve (in two dimensions) or a surface (in three dimensions). This means that two grids are present in the domain: a fixed Eulerian grid on which the flow governing equations are solved on and an overlaid Lagrangian grid which follows the motion of the interface. Thus, the topology of the flow (i.e. the regions occupied by each medium) is known at every instant.

Front tracking techniques may be categorized according to the way that the multicomponent domain is treated by the Eulerian solver. In the first category, the interface is given an artificial thickness of a few cells and the jumps of the discontinuous flow variables are smoothed out. This is the approach followed in one of the early implementations of front tracking [7]. In the second category, the jumps of variables across the interface are kept fully sharp. This is achieved by treating each subdomain (region occupied by each component) separately from its neighbors, in order to avoid the smearing and oscillations that would occur if the stencil included cells occupied by different materials (see the comprehensive review of Abgrall and Karni [8]). However, following the splitting of the domain in several subdomains, boundary conditions need to be defined at the borders of each subregion. The numerical implementation of this idea is carried out via the introduction of ghost cells within the domain.

Front tracking methods may be further categorized according to the way that the states of the ghost cells are defined. The two different approaches for the determination of what will be henceforth called inter-component boundary conditions are called dynamic and state-free tracking.

In *dynamic tracking*, the flow variables are stored at the nodes of the Lagrangian mesh as well. To be more precise, two sets of flow variables are stored at each Lagrangian node: one at each side of the interface. The interface position and states are updated by solving a Riemann problem along the interface. However, since no multidimensional Riemann solvers are available, dimensional splitting has to be applied. The equations are thus solved through successive one-dimensional steps; first in the normal to the interface direction, and then in the tangential one(s). The interface is propagated based on the velocity of the contact wave. The states of the interface are extrapolated to the nearby cells and are used as boundary conditions for each subregion. This is the approach followed in the work of Glimm and collaborators [4, 5, 6] and provides with highly accurate results [9, 10, 11, 12, 13, 14, 15]. However, storing the flow variables at the Lagrangian nodes along with the processes required for the solution of the flow equations at the interface result in a significant increase in the complexity and computational cost of the method.

*Corresponding author

Email address: y.ventikos@ucl.ac.uk (Yiannis Ventikos)

An alternative approach is *state-free tracking*, used in conjunction with an internal boundary capturing method, commonly referred to as a ghost fluid method. In state-free tracking there is no need to store the flow variables at the Lagrangian grid nodes. Instead, the information required to propagate the interface and define the states of the ghost cells is taken from the flow variables that are stored in the underlying Eulerian grid. Several techniques can be used to extract this information (values copied directly from the nearest cell or calculated via extrapolation/interpolation over the neighboring region). The advantage of this approach lies in the fact that it is straightforward and computationally less demanding than dynamic tracking. However, the two meshes are strongly coupled, which may be undesired in certain applications. The method that is presented in the current work falls under this category of front tracking methods. Other algorithms that are based on this idea are the methods of Hao and Prosperetti [16], Terashima and Tryggvason [17, 18], Bo et al. [19] and Wang et al. [20] (and its recent extension to two dimensions by Lu et al. [21]).

Over the past years, much work has been carried out in order to find a way to define the inter-component boundary conditions and interface propagation velocities in an accurate way using information from the Eulerian grid. The original ghost fluid method (GFM) was proposed in the seminal work of Fedkiw et al. [22]. Shortly after publication, the original method was extended in order to be able to handle deflagration and detonation discontinuities [23] and modified in order to be able to resolve interactions between liquid and gas media [24]. Liu and collaborators called attention to a significant shortcoming of these methods [25, 26, 27]; they were shown not to work consistently in the case of a strong shock impacting on a material interface. Several variations of the ghost fluid method have since been proposed, working on rectifying the deficiencies of the original versions, including the modified ghost fluid method (MGFM) of Liu et al. [26], the interface interaction method of Hu and Khoo [28] and the real ghost fluid method (rGFM) of Wang et al. [29]. All these methods however share a common idea: the solution of a Riemann problem at the interface, constructed from values stored at the neighboring cells of the Eulerian grid. The differences between each method lie in the way that the Riemann problem is constructed, solved and then used in the fixed grid. A review of several ghost fluid techniques was carried out by Sambasivan and Udaykumar [30]. More recently, a thorough study on the accuracy and errors introduced by each ghost fluid method was carried out by Xu and Liu [31].

The goal of the present work is to present a simplified framework for multidimensional simulations of compressible multicomponent flows. All aforementioned techniques for the definition of the inter-component boundary conditions are easy to implement in one dimension but require additional procedures when extended to higher dimensions. The proposed formulation retains the simplicity and cost of one-dimensional ghost fluid methods in higher dimensions, while at the same time manages to accurately resolve any wave-interface interaction. By implementing the aforementioned technique within the front tracking framework¹ developed by Glimm and his collaborators [4, 5, 6, 32, 33, 34, 35, 36, 19] we perform simulations that demonstrate the validity, robustness and accuracy of our method.

2. Methodology

2.1. Flow governing equations

For the purpose of demonstrating the qualities of the developed method, the fluids are assumed immiscible, while the effects of viscosity, thermal conductivity and surface tension are neglected. Furthermore, gases are considered non-condensable, while liquids are assumed to be pure (i.e. they cannot break down). Flows are thus modeled by the compressible Euler equations,

$$\frac{\partial \mathbf{q}}{\partial t} + \nabla \cdot \mathbf{f}(\mathbf{q}) = 0, \quad \mathbf{q}(\mathbf{x}, t) = \begin{pmatrix} \rho \\ \rho \mathbf{u} \\ E \end{pmatrix}, \quad \mathbf{f}(\mathbf{q}) = \begin{pmatrix} \rho \mathbf{u} \\ \rho \mathbf{u} \otimes \mathbf{u} + p \mathbf{I} \\ (E + p) \mathbf{u} \end{pmatrix}, \quad (1)$$

where ρ is the density, \mathbf{u} the velocity vector, p the pressure and E the total energy per unit volume. Total energy may be defined as the sum of the internal and kinetic energy of the fluid,

$$E = \rho e + \frac{1}{2} \rho |\mathbf{u}|^2, \quad (2)$$

where e is the specific internal energy of the fluid. The above system of equations is complete upon consideration of an equation of state which relates the pressure with the density and internal energy of the fluid,

$$p = p(\rho, e). \quad (3)$$

¹http://www.ams.sunysb.edu/~linli/FronTier++_Manual/index.html.

2.1.1. Equation of state

The implementation of an equation of state with the current approach is straightforward. As the flow topology is known at every instant, an appropriate expression may be used for each fluid. In the present method, both fluids (water and air) were modeled by the stiffened gas equation of state,

$$p + \gamma p_\infty = \rho(\gamma - 1)(e + e_\infty), \quad (4)$$

$$RT = \frac{(p + p_\infty)}{\rho} + (\gamma - 1) e_t \rho^{\gamma-1}, \quad (5)$$

which reduces to the ideal gas law in the case of air. As in the standard version of the stiffened gas law the internal energy and temperature are inaccurately calculated, these values are renormalized by introducing two extra parameters, e_∞ and e_t [37]. These parameters have been fitted to match the IAPWS-95 equation of state [38] and experimental data [39], as described with more detail in Hawker and Ventikos [14], and their values are given in Table 1,

Table 1: Thermodynamic parameters for the stiffened gas equation of state.

| Fluid | γ | $p_\infty(Pa)$ | $e_\infty(J/kg)$ | $R(J/Kg/K)$ | $e_t(J \cdot kg^{-\gamma} \cdot m^{3(\gamma-1)})$ |
|-------|----------|----------------|--------------------|-------------|---|
| Water | 4.4 | $6 \cdot 10^8$ | $7.456 \cdot 10^6$ | 1700 | $-1.8 \cdot 10^{-6}$ |
| Air | 1.4 | 0 | 0 | 287.058 | 0 |

where T is the temperature, γ is the adiabatic index, p_∞ is the stiffened pressure constant, R is the gas constant, e_∞ is the energy translation factor and e_t is the thermal energy factor.

In related works, water is modeled using either the Tait equation [40, 41] or the stiffened gas equation of state [42, 14, 43, 15]. In our work, the stiffened gas law was selected as it is a generalised equation of state and allows for the modeling of several different fluids. This is of importance when it comes to the implementation of the method; the numerical solver is formulated in the same way for both gas and liquid phases. Moreover, the Tait equation provides no means to calculate the temperature.

Recently, more complex equations of state (in equation or tabulated form) have been employed in order to investigate processes that occur in intense conditions. In Tully et al. [44], several equations of state were considered in the modeling of the shock induced collapse of cavities. They explored the effect of different thermodynamics over a large range of incident shock strengths (from 0.1 to 1000GPa), where warm dense matter regimes are involved. They found that more sophisticated equations of state provided with more relevant results for the density and temperature fields (especially as the shock strength increases) but that pressure and the overall trends of the collapse were captured accurately enough with simpler equations of state, such as the one used in the current work. In general, the methodology we propose in this paper can accommodate any equation of state, in formulaic (as in [45]) or look-up table format.

2.2. Interface handling and propagation

As dictated by the flow governing equations, a material interface gets advected by the local flow,

$$\frac{d\mathbf{x}(t)}{dt} = \mathbf{u}(\mathbf{x}, t), \quad (6)$$

where $\mathbf{x}(t)$ denotes the position of the interface in space and $\mathbf{u}(\mathbf{x}, t)$ the velocity field. In the numerical implementation of the method, the velocity at the interface points is interpolated from the values at the nearby cells of the fixed underlying grid. A second order Runge-Kutta method is used for the integration of the advection equation.

At this point we should note that tangential velocity is discontinuous across material interfaces for the Euler equations; this means that the velocity field is not well-defined on the interface. However, displacing an interface along itself (by any - physically meaningful or not - velocity field) has no effect on it². The reason for propagating the interface using the entire velocity vector rather than its normal-to-the-interface projection is related to the concept of the proposed grid-aligned formulation: one of its characteristics is that there is no need to compute geometrical quantities at the interface, such as the interface normal vector. As a result, we do not have access to this information.

The ‘‘locally grid-based’’ algorithm [36] is used for the handling and maintenance of the Lagrangian grid. The above algorithm essentially lets the front nodes freely advect in the domain, except for the case where two fronts (or two parts of the same front) interact. An Eulerian-grid-based reconstruction of the interface is carried out in such occurrences. Other algorithms may also be

²In numerical implementations, this introduces some error and increases the probability of mesh tangling. In practice however, tangling events are treated robustly by the numerical framework, whereas the advection error introduced by front tracking is very small (as shown in section 3.1).

employed, but provide with less robustness or accuracy as reported in [36]. Besides detection and resolution of tangled fronts, maintaining a well-shaped Lagrangian grid involves procedures such as node rearrangements and grid redistribution. In the simulations presented in this work, these are performed at a user-specified frequency (typically twenty time steps, unless otherwise stated). For more details regarding these algorithms the reader is referred to publications by the developers of the employed front tracking framework [5, 32, 33, 34, 35, 36].

In state-free tracking, the resolution of the Lagrangian mesh is related to that of the underlying Eulerian grid. The information for the propagation of the interface is extracted from the underlying grid, where the flow variables are stored. This effectively means that sub-grid level accuracy can not be achieved, despite having a sub-grid level representation of the interface. Therefore, maintaining a Lagrangian mesh with significantly higher resolution than the base grid is without value. Moreover, the flow solver on the underlying mesh is not aware of sub-grid level details in the representation of the interface.

Clearly, the interface is represented in more detail as the resolution of the Lagrangian grid increases, but wiggles are also more likely to emerge and possibly grow. Therefore, a balance between accuracy and stability (from an algorithmic point of view) needs to be found, especially for simulations in three dimensions where the resolution of any interactions between fronts is not easily taken care of. Keeping the resolution of the interface within a range of $\sim 1 - 2$ Lagrangian points per fixed grid cell is guaranteeing smooth behavior of the algorithm with few tangling incidents and at the same time a detailed enough representation of the interface.

2.3. Fixed grid solver

The system of equations (1) is solved in a fixed rectangular grid with equidistantly spaced nodes in each direction. The fluxes are calculated by a global Lax-Friedrichs flux vector splitting scheme with a characteristic-wise 5^{th} order WENO reconstruction [46]. The equations are integrated in time via a third-order Runge-Kutta method³. The time step is selected as the smallest between the time steps associated to each grid. The time step related to the Lagrangian grid is,

$$\Delta t_L = CFL \cdot SF \cdot \frac{\min(\Delta x, \Delta y, \Delta z)}{|\mathbf{u}|}, \quad (7)$$

where SF denotes a safety factor introduced in order to restrict tangling events and $|\mathbf{u}|$ is the local velocity vector magnitude. The timestep for the Eulerian solver is,

$$\Delta t_E = CFL \cdot \frac{\min(\Delta x, \Delta y, \Delta z)}{\lambda_{max}}, \quad (8)$$

where λ_{max} is the maximum wave speed in the domain. It should be noted that the time step related to the Lagrangian grid is expected to be much larger than that of the fixed grid as the speed of sound is not involved.

2.3.1. Inter-component boundary conditions

As the numerical scheme applies separately to each phase, boundary conditions within the domain are required, in order for the stencils to be complete across the interface. In our method, constant state inter-component boundary conditions are used, meaning that all cells that are required to complete the stencil have the same state.

The ghost cells are populated with the state of the interface, which is calculated using the interface interaction method (IIM) as it is both simple and robust. For a detailed evaluation of the performance of this method, the reader is referred to the original work by Hu and Khoo [28]. The states of the cells across the interface are computed using the method of characteristics,

$$u_I = u_l - \int_{p_l}^{p_I} \frac{dp}{\rho c}, \quad (9)$$

$$u_I = u_r + \int_{p_r}^{p_I} \frac{dp}{\rho c}, \quad (10)$$

where c denotes the speed of sound. The subscripts $()_l$, $()_r$ denote the cells on each side of the interface respectively, while the subscript $()_I$ stands for the state on the interface. By considering that the speed of sound is not expected to change significantly from the interface to the adjacent cells, the interface condition becomes,

$$u_I = \frac{\rho_l c_l u_l + \rho_r c_r u_r + p_l - p_r}{\rho_l c_l + \rho_r c_r}, \quad (11)$$

$$p_I = \frac{\rho_l c_l p_r + \rho_r c_r p_l + \rho_l c_l \rho_r c_r (u_l - u_r)}{\rho_l c_l + \rho_r c_r}. \quad (12)$$

³Simulations in sections 4.3 and 5 were performed with a second-order RK method in order to reduce the computational load; validation against the RK3 method was performed for the equivalent problems in two dimensions.

The density on the two sides of the interface may be calculated by performing an isentropic extrapolation from each side as follows,

$$\rho_{I,l} = \rho_l \left(\frac{p_I + p_{\infty,l}}{p_l + p_{\infty,l}} \right)^{1/\gamma_l}, \quad \rho_{I,r} = \rho_r \left(\frac{p_I + p_{\infty,r}}{p_r + p_{\infty,r}} \right)^{1/\gamma_r}, \quad (13)$$


where the subscripts $(\cdot)_{l,r}$ stand for the left and right side of the interface respectively.

2.3.2. Extension to higher dimensions

For the extension of the method to higher dimensions, the jump conditions should be applied in a direction that is normal to the interface. As a result, the Riemann problems should be constructed and solved in the normal to the interface direction. In the present work, a novel approach, referred to as the GA-IIM (grid-aligned interface interaction method), was developed. In this version, the stair-step representation of the flow topology (figure 1) in the fixed grid is fully adopted. This is an approach that relates to the work of Farhat et al. [47]. The jump conditions are applied on the interface following its mapping to the Eulerian grid (figure 1b), rather than in a normal to the interface direction. In a sense, the above way of coupling the flows is more appropriate for a pure-cell model. We note that the proposed grid-aligned formulation is not specific to the particular ghost fluid method and Riemann solver; all existing (known to the authors) methods can be formulated in such a manner.

At this point, we should make a note related to figure 1b. Eulerian gridlines represent cell edges, meaning that cells are filled in their entirety with a particular fluid at each instant. Evaluating which fluid occupies each cell (or more precisely its center) is a process that varies according to the technique used for the localization of the interface. For example, within a level set framework, this is specified by the sign of the level set. In front tracking, this information is provided by labels on the front points [5, 32, 35].

The information that is required to compute the interface condition is taken from the cells across the interface, while the information that is missing, namely the velocity component(s) in the tangential to the interface direction(s), is constantly extrapolated (i.e. directly copied) from the real fluid, as the jump of the tangential velocity is undefined in Euler's equations. Compared to the standard way of extending ghost fluid methods to higher dimensions, the proposed formulation is much simpler in its implementation and computationally less expensive, as we avoid interpolating or extrapolating any variables, determining ghost fluid regions, and computing any geometrical quantities. Despite its simplicity, the formulation gives accurate results, as will be shown in a following section.



(a) Topology described by the Lagrangian grid.

(b) Topology seen by the Eulerian grid.

Figure 1: Representation of the interface by the Lagrangian grid (left) and after mapping on the Eulerian grid (right). Eulerian gridlines represent cell edges.

2.3.3. Outline of the GA-IIM algorithm

In order to avoid oscillations and keep the jumps across the material interface sharp, each fluid is treated separately, by considering the interface as a boundary for each subdomain. The ideas that make up the GA-II method are the following:

- The single-phase stencils are completed by enforcing Dirichlet-type boundary conditions at the interface.
- The states of these conditions are the states of the fluids at the interface ($q_{BC,:} = q_{I,:}$).
- The conditions ($q_{I,:}$) are formulated and computed for the topology seen by the underlying grid (figure 1b).

Under the above assumptions, the extension of a ghost fluid method to higher dimensions can be carried out in the simplest possible - yet formally consistent - way. The problem is effectively reduced to expressions that are elementary and only involve the states of two adjacent cells.

As an example, consider the two-dimensional configuration of figure 1. In a horizontal-direction sweep that crosses the material interface (magnified part of figure 1b), we enforce fixed boundary conditions at the interface $\mathbf{q}_I = (\rho, p, u_n, u_t)_I$ for each fluid on the left and right side (subscripts n and t refer to the normal and tangential to the interface directions). Equations (11-13) are applied in the normal to the interface direction (using $(\rho, p, u)|_{I,r}$), yielding three out of the four components of the interface states $(\rho, p, u_n)_I$. The tangential to the interface velocity $u_t|_I$ simply takes its value from the adjacent cells v_l or v_r , for the fluids on the left and right side of the interface respectively. The same process is followed for interfaces in the vertical direction, with the difference that expressions (11-13) are computed using $(\rho, p, v)|_{I,l}$ and the tangential to the interface velocity $u_t|_I$ takes the values u_u or u_l (subscripts u and l now referring to the cells above and below the interface).

3. Initial assessment of the method's accuracy

The goal of the current section is (a): to evaluate the performance of front tracking against the level set method, (b): to validate the implementation of the IIM in a simple one-dimensional test configuration and finally (c): to assess the behavior of the grid-aligned formulation. This study is carried out using simple configurations in order to isolate the parameters under study: the linear advection of an interface and shock tube problems in one and two dimensions.

3.1. The linear advection of an interface

The slotted disk configuration [48] is a well-known test case to assess the capability of a method to properly follow the temporal evolution of an interface. In the current work, the rotation of a rigid body was considered, with the computational domain being a square of size $\mathbf{x} \in [0m, 100m]^2$. A disk of radius $R = 15m$ with a rectangular $[5m, 25m]$ slot on one side was centered at coordinates $(x, y) = (50m, 75m)$ (see figure 2). The body was rotated with a velocity field $\mathbf{u}(m/s)$, defined as:

$$\mathbf{u}(x, y) = \begin{cases} u(y) = 0.01(50 - y) \\ v(x) = 0.01(x - 50) \end{cases} \quad (14)$$

Following the completion of a full rotation, the interface was superposed on the initial field in order to allow for a comparison of the calculated solution with the initial condition (see figure 2). No deformation of the shape of the body is expected. The results are compared to those obtained by the level set method of Zuzio [49]. Two (underlying) meshes of equally sized elements were considered, with 32 and 64 elements in each direction. The Lagrangian grid parameters were set such that no redistribution of the mesh points takes place. The front tracking approach is shown to be significantly superior when compared to the results obtained by the level set method described in [49]; the interface is almost perfectly transported even when the grid resolution is low. More detailed studies comparing the performance and cost of front tracking, level set and volume of fluid methods may be found in [32, 33, 36].

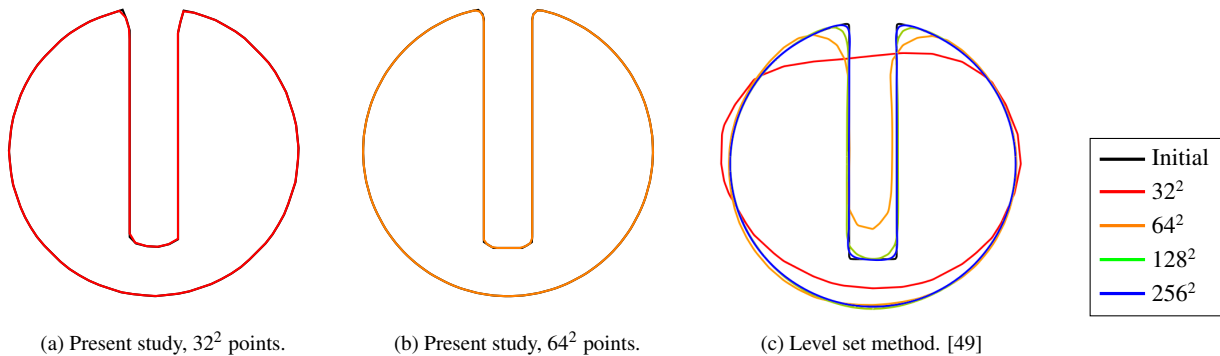


Figure 2: Zalesak disk test case. The initial interface is depicted in black, whereas the interface after one rotation is colored as per the legend on the right.

3.2. One-dimensional gas-dynamics and the effect of front tracking

The Sod shock tube problem is a well-known benchmark test case for compressible flow solvers. The computational domain is a one-dimensional ($x \in [-1m, 1m]$) tube, with a diaphragm placed at $x_c = 0m$. The states of the gas ($\gamma = 1.4$, $p_\infty = 0Pa$) on the left and right side of the interface are given:

$$(\rho, u, p) = \begin{cases} (1, 0, 1), & x \leq x_c \\ (0.125, 0, 0.1), & x > x_c \end{cases}. \quad (15)$$

At $t_0 = 0s$, the diaphragm is removed and the flow is allowed to evolve. This results in the propagation of three waves within the domain; a rarefaction wave, a contact discontinuity and a shock front. The aim of the test is to assess the ability of the solver to sharply capture the shock front and the contact discontinuity as well as to accurately reproduce the profile of the rarefaction wave. In the current study, the goal was to assess the behavior of front tracking when applied to the tracking of a discontinuity. For this reason, the results were compared against simulations that did not involve tracking of the discontinuity. Several grids were used in this study, with their characteristic quantities being presented in Table 2.

Table 2: Parameters of the grids used in the 1D shock tube simulations.

| Mesh | I | II | III | IV |
|----------------------|------|------|------|-------|
| No. elements | 50 | 100 | 200 | 400 |
| Grid spacing (m) | 0.04 | 0.02 | 0.01 | 0.005 |

The density, pressure and velocity profiles (obtained with meshes I and II) at $t = 0.4s$ are displayed in figure 3 along with the analytical solution to the problem. In the simulations with front tracking the contact discontinuity is kept sharp but conservation is not respected. This was expected as the ghost fluid method is not conservative in the region of the discontinuity. Clearly, across the tracked interface, the intercell fluxes do not cancel out. However, the method displays a consistent and convergent behavior as the grid gets refined, i.e. error is systematically reduced. This is outlined in figure 4, where the relative errors in mass, momentum and energy are presented. It is clear that the conservation error relates in a first-order rate with the spatial discretization (see also figure 5a).

The pure-cell definition of the Eulerian topology (schematically depicted in figure 1b) is another (secondary) source of error. This error is reflected in the sharp jumps (or oscillations) in figure 4. Clearly, this error is directly related to the spatial discretization and reduces with mesh refinement: the frequency and amplitude of the jumps matches the grid, i.e. their frequency doubles when the grid resolution is doubled, whereas their magnitude decreases.

The fact that the interface is propagated based on the underlying velocity field leads to a delayed response in its motion; in the early stages of the interaction all nearby cells have an almost zero velocity. The relative error in the position of the discontinuity (with respect to the analytical solution) is presented in figure 5b; it is evident that it is of the order of $O(\Delta x)$. This means that the dynamics of the discontinuity can be accurately predicted by using a sufficient spatial resolution.

3.3. The effect of the grid-aligned definition of the inter-component conditions

The explosion of a bubble in a square of size $x \in [0m, 2m]^2$ (see figure 6) was considered in order to assess the behavior of the GA-II method. The bubble has a diameter of $d = 0.2m$ and was placed at the center of the domain, which was discretized using 400 points in each direction (yielding a low resolution of 20 points per initial bubble radius, $R/\Delta = 20$). Both sides of the bubble were filled with the same fluid ($\gamma = 1.4$, $p_\infty = 0Pa$) of same density ($\rho = 1kg/m^3$) at different pressures $p_1 = 1Pa$, $p_2 = 100Pa$.

The solution along the $y = 1$ and $y = x$ lines (denoted with dashed curves in figure 6) is presented in figure 7. In that way we are able to evaluate any differences between the normal-to-the-interface and normal-to-the-grid definition of the boundary conditions, as the interface along the $y = 1$ line is locally aligned to the grid. The results are almost indistinguishable. We should take into consideration that the resolution in the $y = x$ direction is effectively $1/\sqrt{2} \approx 0.707$ times coarser, which is a quite significant factor. The differences in the velocity magnitude near the bubble center are not related to the grid-aligned ghost fluid formulation; they are artifacts of the numerical scheme that tries to capture the shock waves that converge radially and interact at the center of the bubble in a rectangular grid (see figure 8a). What is relevant to the developed method is the preservation of the circular shape of the interface. We hence introduce the concept of circularity, which is defined as the ratio between the surface area A and the perimeter P of our Lagrangian mesh,

$$Circ = \frac{4\pi A}{P^2}, \quad (16)$$

and takes the value of 1 for a perfect circle. The circularity of the contact discontinuity along time is plotted in figure 8b for various grid resolutions; overall, it is preserved with great accuracy.

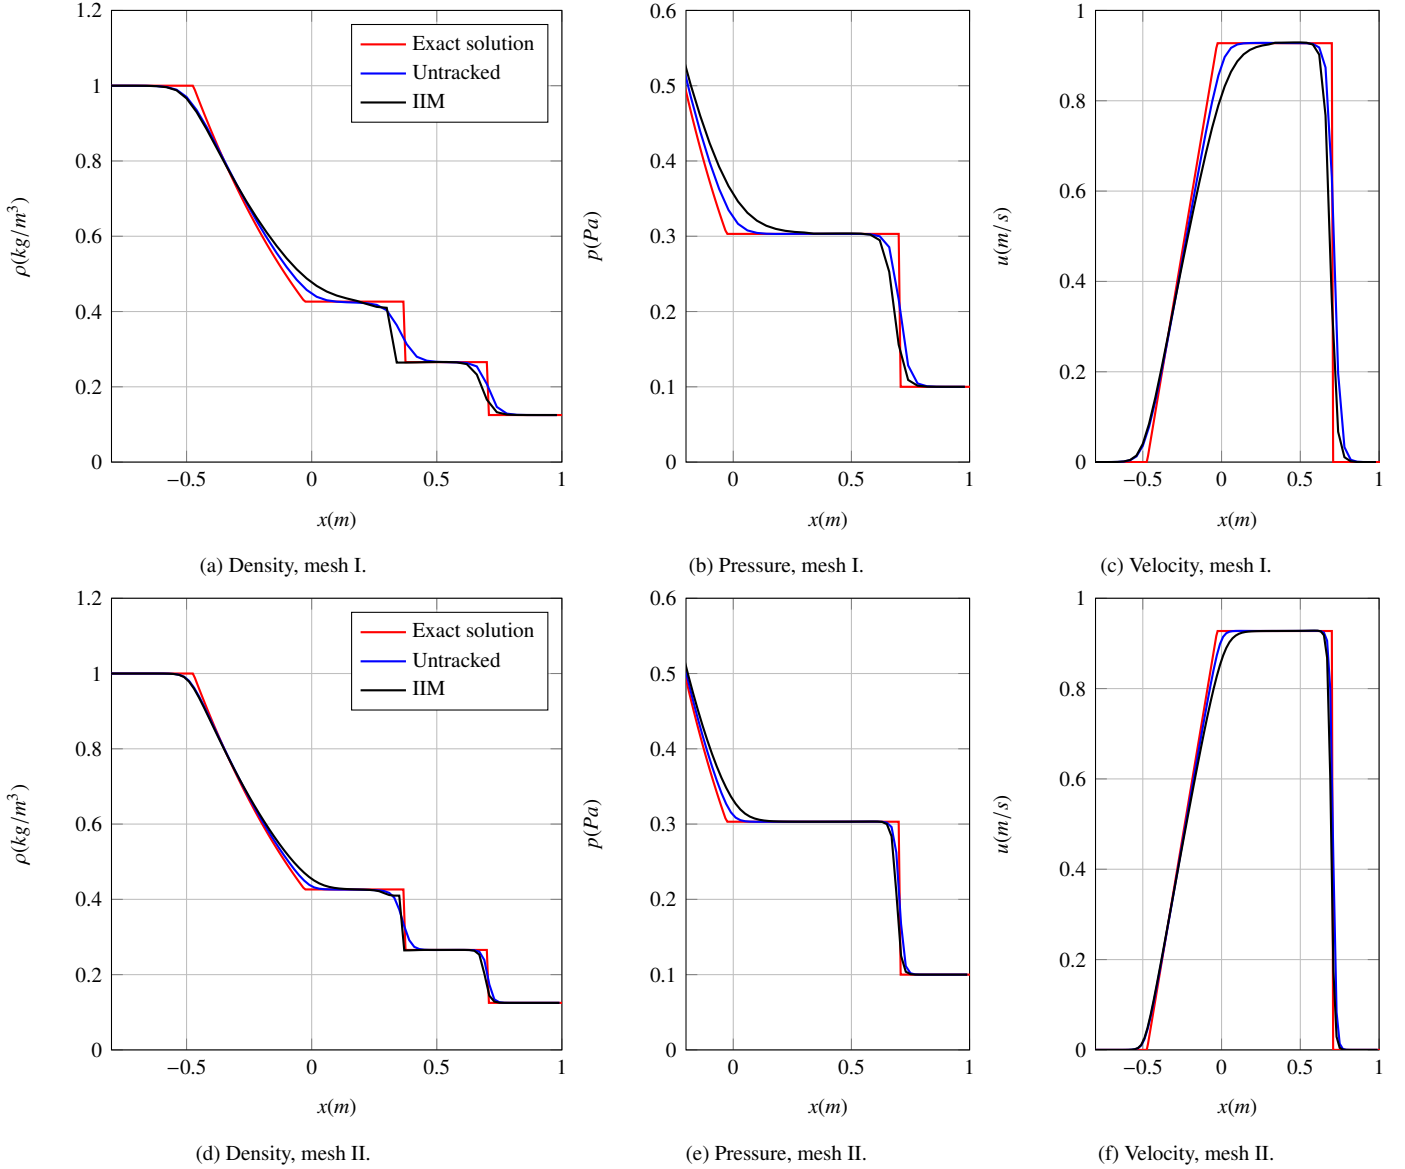


Figure 3: Density, pressure and velocity profiles for two different grids, $t = 0.4s$.

4. Validation of the method

The validation of the method continues with problems of shock- and pressure-driven collapse of bubbles. Four configurations of gradually increasing complexity are studied. In the first two cases we simulate the shock induced collapse of a cylindrical (2D) bubble, while in the latter two we study the shock induced and symmetric collapses of a spherical (3D) air bubble in water. As a final case for the validation of the method, we consider the high-speed breakup of a cylindrical water droplet in air.

4.1. 2D simulation of the shock induced collapse of a helium bubble in air

The problem that was considered first involves the passage of a $M = 1.22$ shockwave through a cylindrical helium bubble. The medium in which the shockwave propagates is air. This configuration is similar to the experiments of Haas and Sturtevant [50] which have been used by several researchers as a benchmark in order to validate their numerical methods [51, 52, 53, 17, 54, 43]. The simulation was carried out in two dimensions, in a rectangular domain of the following size: $x \in [0m, 0.0445m]$, $y \in [0m, 0.325m]$. The initial radius of the bubble was $r = 0.025m$ and its center was placed at $(x, y) = (0m, 0.15m)$. Only half of the bubble was modeled as the configuration is symmetric about the y axis. The shockwave that propagates in the y direction was initially placed at $y = 0.1m$. A reflecting boundary condition at $x = x_{max}$ modeled the shock tube wall, while non-reflecting conditions were set at

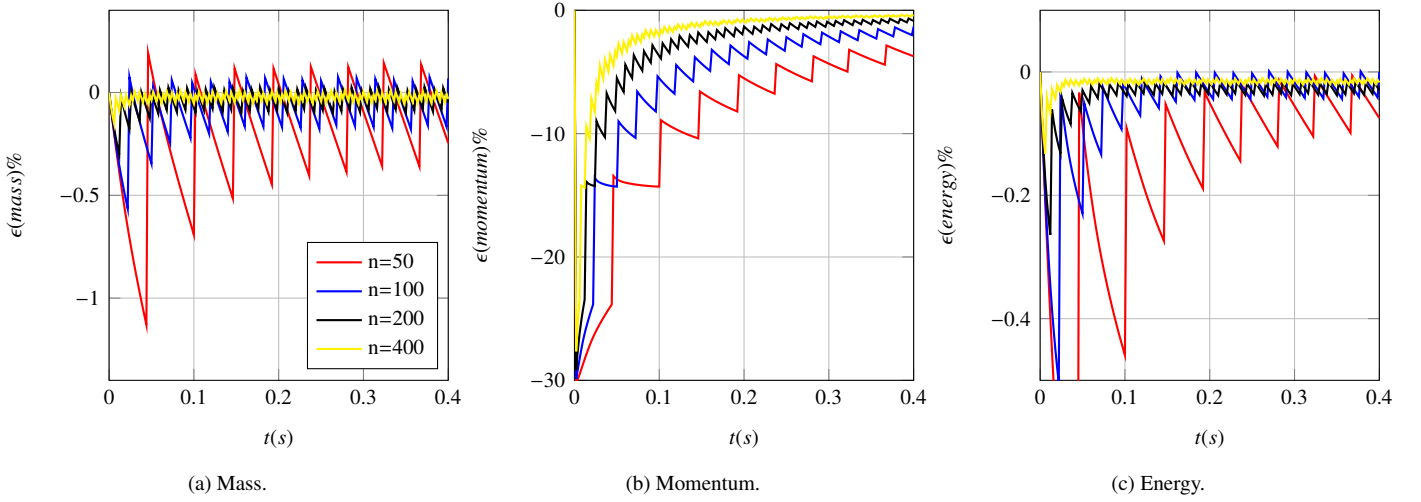


Figure 4: Relative mass, momentum and energy errors along time.

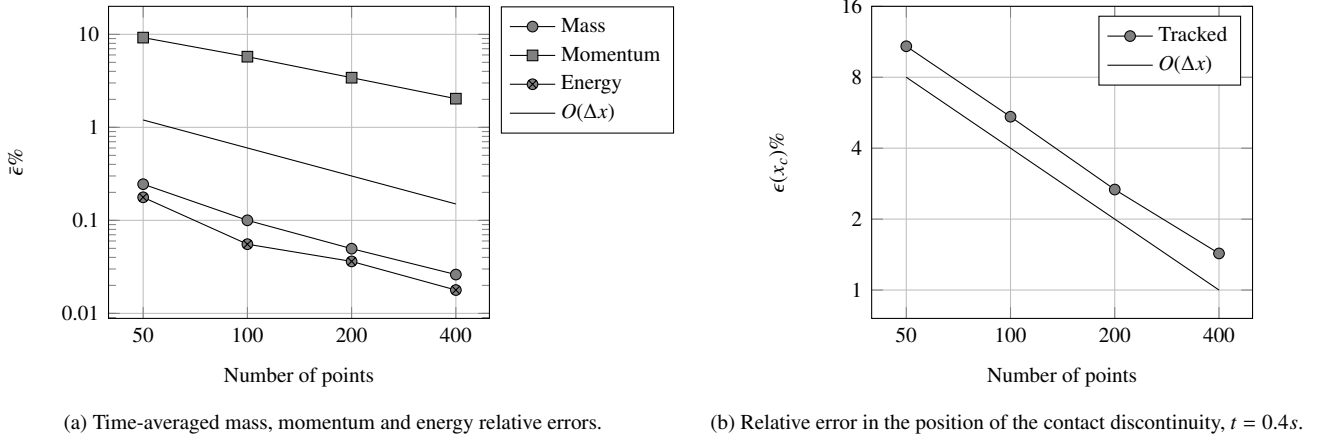


Figure 5: Relative errors for different grid resolutions.

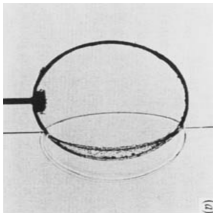
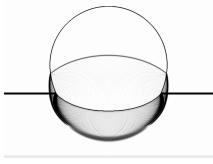
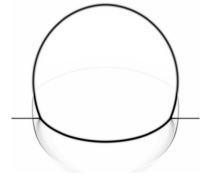
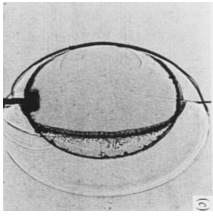



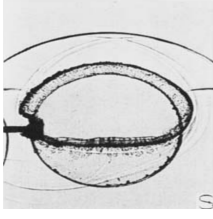

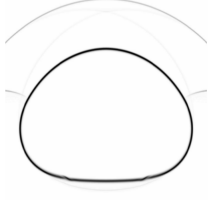
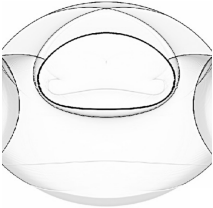
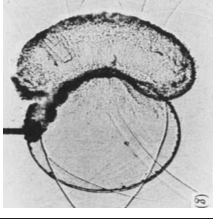
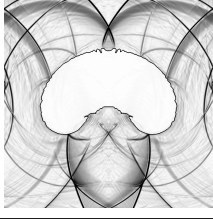
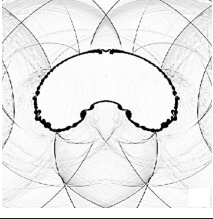
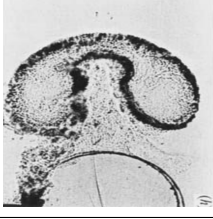
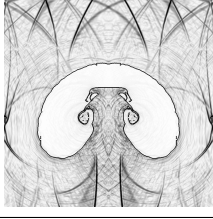
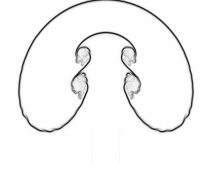
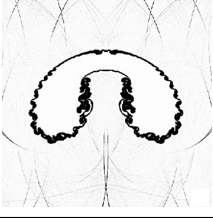

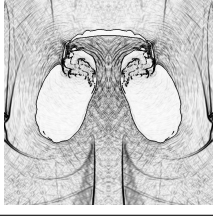

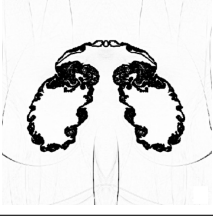
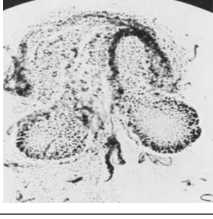
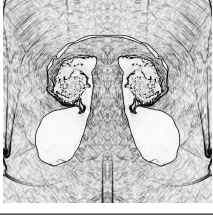

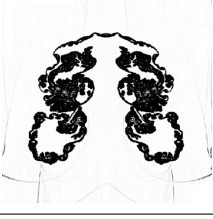
the downstream boundary of the domain ($y = y_{max}$). Constant post-shock conditions were applied at the upstream domain boundary ($y = y_{min}$). The bubble was filled with helium partly contaminated by air ($\gamma = 1.648$, $p_\infty = 0Pa$, $\rho = 0.235kg/m^3$) and the ambient medium was pure air ($\gamma = 1.4$, $p_\infty = 0Pa$, $\rho = 1.29kg/m^3$). The (unshocked) fluids were initially quiescent at atmospheric pressure $p = 101325Pa$.

A preliminary qualitative comparison between the results of our method and the experimental results of Haas and Sturtevant [50] as well as with the numerical studies of Shankar et al. [54] and Coralic and Colonius [43] is presented in Table 3. In our work, the domain was discretized using 200 points per initial bubble radius (ppbr). The mesh is sufficient to accurately resolve the flow, as will be shown in the analysis that follows. The resolution that was used in the works of Shankar et al. and Coralic and Colonius was 200ppbr and 500ppbr respectively. The scale and focal point of the plots at every instant are not the same. In general, the main features of the flow field and the interface are accurately reproduced by our method, although some differences in the orientation and integrity of the bubble can be observed (especially at later times) when comparing our results to those of Haas and Sturtevant [50]. However, these differences are attributed to the presence of the rod in the experiments (visible on the left side of the photographs), which disrupts the flow field. There is close agreement with the results from the other numerical works, except for some small scale features; this is attributed to the differences in the adopted flow models and numerical approaches. For a quantification of the results, the evolution of three distinct points on the interface of the bubble is presented in figure 9 for different grid resolutions. The results of our method are also compared against the results of Shankar et al. [54]; the two methods are in excellent agreement.

4.2. 2D simulation of the shock induced collapse of an air bubble in water

The interaction between a gas bubble and a shock wave in a liquid was the second configuration that was considered. A cylindrical air bubble of diameter $d = 0.001m$ was placed in a domain of size $x \in [0m, 0.001m]$, $y \in [0m, 0.002m]$, in mechanical

Table 3: Contours of the magnitude of the density gradient at various time instants after shock impact, 2D simulation of the $M = 1.22$ shock induced collapse of a $d = 0.05m$ helium bubble in air. Comparison against the experiment of Haas and Sturtevant [50] and the numerical works of Shankar et al. [54] and Coralic and Colonius [43].

| Time | Haas [50] | Present study | Shankar [54] | Coralic [43] |
|------------|---|---|--|---|
| $32\mu s$ |  |  |  | - |
| $62\mu s$ |  |  |  |  |
| $102\mu s$ |  |  |  |  |
| $245\mu s$ |  |  | - |  |
| $427\mu s$ |  |  |  |  |
| $674\mu s$ |  |  |  |  |
| $983\mu s$ |  |  |  |  |




Figure 6: High-pressure bubble explosion.

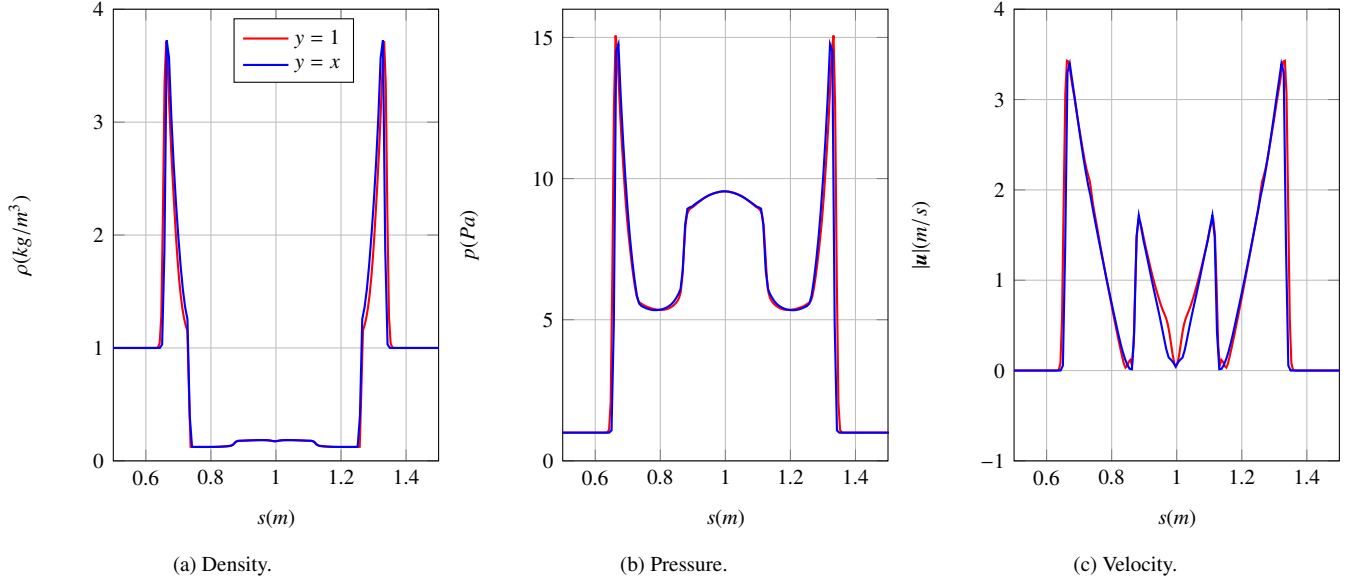




Figure 7: Density, pressure and velocity profiles, $t = 0.04s$.

equilibrium with the surrounding liquid, which was water. Only half of the bubble and the surrounding domain were simulated, due to symmetry. The bubble was centered at $(x, y) = (0m, 0.00075m)$. The density of water was set to $\rho_w = 993.89kg/m^3$ and that of air at $\rho_g = 1.204kg/m^3$. Both fluids were initially at rest under atmospheric pressure conditions. The thermodynamic parameters that were used to model the two fluids are presented in Table 1. Post-shock conditions were defined at the lower boundary ($p = 1GPa$, which corresponds to a $M \simeq 1.43$ shock wave in water), while non-reflecting conditions were set at the two remaining boundaries ($x = x_{max}$ and $y = y_{max}$). The main features of the collapse are presented in figure 10, in which the contours of the magnitude of the density gradient are depicted at four different instants after shock impact. For validation purposes, the figures are plotted alongside the results of Hawker and Ventikos [14] for grids of the same size. Note that the scale, focal points and color map ranges of the plots are not identical.

In figure 10a, $t = 230ns$ after the shockwave impacts the bubble, the interface is accelerated, whereas the shockwave gets transmitted in the gas and reflected as a strong rarefaction back in the liquid. The shockwave is curved at the side of the bubble. In figure 10b, the formation of the high speed jet may be observed. The gas in the bubble is compressed and heated after passage of several (reflected at the interface) shock waves. In the following figure (10c), at $t = 665ns$ after shock impact, the jet is about to impact the far bubble wall. A strong water hammer shockwave is generated upon impact. The transmitted shock wave that was propagating in the bubble has already reached the far bubble wall and gets transmitted in the water, and reflected back in the gas. Finally, in figure 10d the water hammer shockwave is seen traversing the bubble remains. The peak pressures during the collapse



(a) Velocity magnitude contour plot, $t = 0.04s$.



(b) Circularity of the contact discontinuity along time.

Figure 8: Circularity related parameters.





Figure 9: Evolution of the location of three points on the interface of the bubble along time after shock impact, 2D simulation of the $M = 1.22$ shock induced collapse of a $d = 0.05m$ helium bubble in air. Comparison for different grid resolutions (left) and against the work of Shankar et al. [54] (right).

are found not only in the strong water hammer wave, but also at the impact of the secondary microjets on the side walls of the bubble in the upstream direction and the interaction of all upstream moving waves. After the water hammer shockwave traverses the bubble remnants the strong interactions cease and the flow field is dominated by vortical motion. In general, all the features of the collapse identified in [14] are accurately captured by our method (see also an animation of this process, marked as Video 1, in Supplementary Material).

A quantification of the results is presented in Table 4. Four different grids were considered, with the resolution ranging from $50ppbr$ to $400ppbr$. Comparison is made against the work of Hawker and Ventikos [14], who used a resolution similar to the most refined mesh in our study. The time between the initial shock impact and the impact of the jet at the far bubble wall (first phase time) is used as a basic measure of convergence. However, the collapse process can not be characterized entirely by just one feature, and thus more measures are considered. The other measures to be considered are the jet speed at impact and the peak water-hammer shock pressure. It should be noted that the values in the table are approximate, as we only store the flow variables at a discrete frequency (every $5ns$), which might not coincide with the exact moments of jet or shockwave impact. All measures have reached a certain state of convergence for simulations with $100ppbr$ and above.

The behavior of the method at low resolutions is of particular significance when it comes to the realization of complex three-dimensional simulations. When looking at configurations such as a polydisperse bubble cloud, we would like to be confident that the dynamics of all bubbles are resolved with sufficient and similar accuracy. This is a strong point of the method: as can be seen in

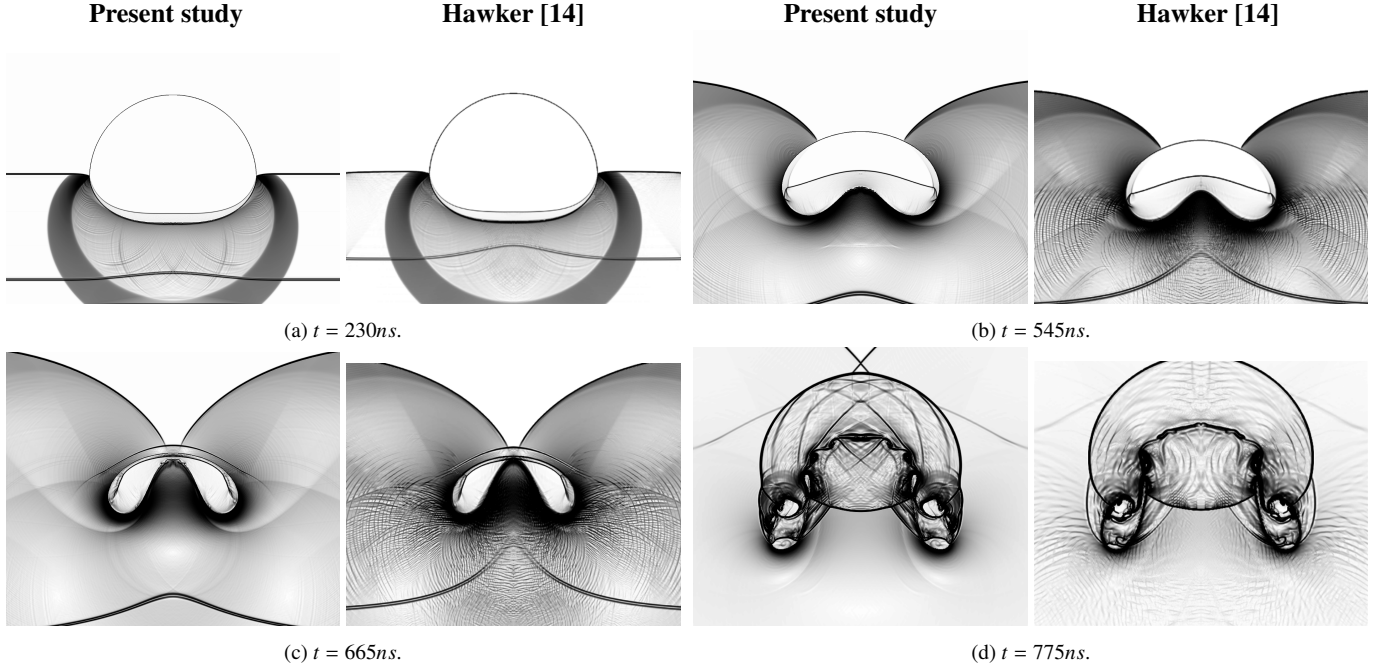


Figure 10: Contours depicting the magnitude of the density gradient at different time instants after shock impact, 2D simulation of the $1GPa$ shock induced collapse of a $d = 0.001m$ air bubble in water, $400ppbr$. Present method (left), Hawker and Ventikos [14] (right).

Table 4 the dynamics are accurately captured even with very low resolutions; the relative error for the first phase time between the coarsest and finest meshes is $< 4\%$.

Table 4: Characteristic collapse measures, 2D simulations of the $1GPa$ shock induced collapse of a $d = 0.001m$ air bubble in water. Results for different grid sizes are compared against the work of Hawker and Ventikos [14].

| | Present study | | | | Hawker [14] |
|--|---------------|-------|-------|-------|-------------|
| Resolution (<i>ppbr</i>) | 50 | 100 | 200 | 400 | 400 |
| First phase time (<i>ns</i>) | 705 | 690 | 685 | 680 | 675 |
| Jet speed at impact (<i>m/s</i>) | 2117 | 2128 | 2132 | 2131 | 2278 |
| Water-hammer shock pressure (<i>GPa</i>) | 3.037 | 3.343 | 3.422 | 3.480 | 3.00 |

4.3. 3D simulation of the shock induced collapse of an air bubble in water

The cylindrical inhomogeneity was subsequently replaced by a spherical bubble and the simulation was carried out in three dimensions. The resolution was limited compared to the two-dimensional study; a grid of 50 points per initial bubble radius was used. Snapshots of the collapse are presented in figure 11. The contours of the velocity magnitude are depicted on the two normal planes. Superposed on the velocity contours are the isolines of the magnitude of the density gradient. The bubble interface is depicted as well; one quarter of the sphere is transparent in order to allow for a view of the interior of the bubble. All the flow features (except for the micro jetting phenomenon due to the limited resolution) that were identified in the 2D simulation of the collapse are observed in the three-dimensional simulations as well. However, the collapse is faster as the focusing is now more intense [14]. As a result, the jet velocities and water-hammer shock pressures are higher. At $t = 660ns$ the bubble obtains its minimum volume. After that moment, it starts expanding, and continues being advected by the flow following a strongly vortical motion. As a result, the bubble remains get enclosed by the thin gas layer at the top ($t = 860ns$).

The characteristic measures of the collapse are presented in Table 5, compared against the work of Hawker and Ventikos [14]. For the sake of completeness, the three-dimensional results of our method are also compared to two-dimensional simulations with the same resolution. As discussed in the grid convergence study that was carried out in two dimensions, such a resolution is not sufficient to capture the dynamics of the collapse in a time-exact fashion. This is why the first phase time of our simulations is larger than what calculated by Hawker and Ventikos [14]. Regarding the other measures of validation and convergence, the jet speed at impact is predicted in excellent agreement with the reference work, while the water-hammer shock is predicted to be stronger.




Figure 11: 3D simulation of the 1GPa shock induced collapse of a $d = 0.001m$ air bubble in water, 50ppbr. The isolines that are superposed on the velocity magnitude contours depict the magnitude of the density gradient. One quarter of the bubble surface is transparent in order to allow for a view of the internal structures.

Table 5: Characteristic collapse measures, 2D and 3D simulations of the 1GPa shock induced collapse of a $d = 0.001\text{m}$ air bubble in water. Comparison against the work of Hawker and Ventikos [14].

| | 2D | 3D | Hawker, 2D [14] | Hawker, 3D [14] |
|--|-------|-------|-----------------|-----------------|
| Resolution (<i>ppbr</i>) | 50 | 50 | 400 | 100 |
| First phase time (<i>ns</i>) | 705 | 570 | 675 | 540 |
| Jet speed at impact (<i>m/s</i>) | 2117 | 2776 | 2278 | 2754 |
| Water-hammer shock pressure (<i>GPa</i>) | 3.037 | 5.464 | 3.00 | 4.04 |

The interface of the bubble at various time instants for both two-dimensional and three-dimensional simulations is depicted in figure 12. In the three-dimensional case only a slice of the bubble is shown in order to allow comparison with the two-dimensional results. Both simulations were carried out with a grid resolution of 50ppbr . In the first frame the interface is presented at $t = 500\text{ns}$ after the impact of the incident shock. The collapse in the case of the spherical bubble is shown to be much faster. In the second frame, the bubbles are depicted at the moment of impact of the main jet on the far wall. The jet in the case of the spherical bubble is much wider and the gas within the bubble is more compressed. In the frames that follow, the bubble interfaces are depicted at instants following the main jet impact. The patterns of the collapsing bubbles are similar. However, in the three-dimensional case, the contents of the bubble are more compressed ($t = 100\text{ns}$ after jet impact), and propagate at a greater speed ($t = 200\text{ns}$ after jet impact). It is interesting to mention that no three-dimensional instabilities are observed, as also reported by Hawker and Ventikos [14].




Figure 12: Bubble interfaces at different time instants after shock impact (SI) and main jet impact (JI), 2D (black) and 3D (red) simulations of the 1GPa shock induced collapse of a $d = 0.001\text{m}$ air bubble in water, 50ppbr . The thin (post-collapse) bubble regions are sub-Eulerian-grid structures; front tracking provides with sub-cell interface resolution. The underlying grid is omitted for reasons of plot clarity.

4.4. 3D simulation of the symmetric collapse of an air bubble in water

To further validate the presented method, the collapse of an air bubble in high-pressure water was considered next. The domain consisted of a cube of size $x \in [-0.003\text{m}, 0.003\text{m}]^3$ and the spherical bubble of radius $r = 0.001\text{m}$ was placed at its center. The density of the high-pressure ($p_w = 35758900\text{Pa}$) water was set to $\rho_w = 1011.4789\text{kg/m}^3$ while the gas within the bubble was initialized with a density and pressure equal to $\rho_g = 1.225\text{kg/m}^3$ and $p_g = 101300\text{Pa}$ respectively. Non-reflecting conditions were set in all of the domain's boundaries. Simulations were carried out with three different grid resolutions corresponding to 15, 25 and 40 points per initial bubble radius.

In figure 13 we compare the average radius and normalized volume of the bubble against the numerical solution to the Keller-Miksis equation [55] and the results of Johnsen and Colonius [42] who carried out two-dimensional axisymmetric simulations of the same configuration. The results from Johnsen and Colonius [42] are those predicted with their most refined mesh (where the bubble is discretized with 150 points along its initial radius). The Keller-Miksis equation is written for a bubble filled entirely with gas. In order to match the physics of the equation to our method, viscosity and surface tension are neglected. In any way, they are not expected to play a significant role due to the size of the bubble and the very strong pressure ratio ($p_w/p_g = 353$) that is driving the collapse. The above assumption can be easily confirmed by comparing the predictions of the model for both scenarios; the results are indistinguishable. The Keller-Miksis equation reads,

$$\left(1 - \frac{\dot{R}}{c_L}\right)R\ddot{R} + \frac{3}{2}\left(1 - \frac{\dot{R}}{3c_L}\right)\dot{R}^2 = \left(1 + \frac{\dot{R}}{c_L}\right)\frac{p_L(R, t) - p_\infty}{\rho_L} + \frac{R\dot{p}_L(R, t)}{\rho_L c_L}, \quad (17)$$

where $R(t)$ is the radius of the bubble, the dots indicate first and second time derivatives, c_L is the speed of sound through the liquid and $p_L(R, t)$ is the pressure on the liquid side of the bubble's interface. The gas inside the bubble is considered to be adiabatically

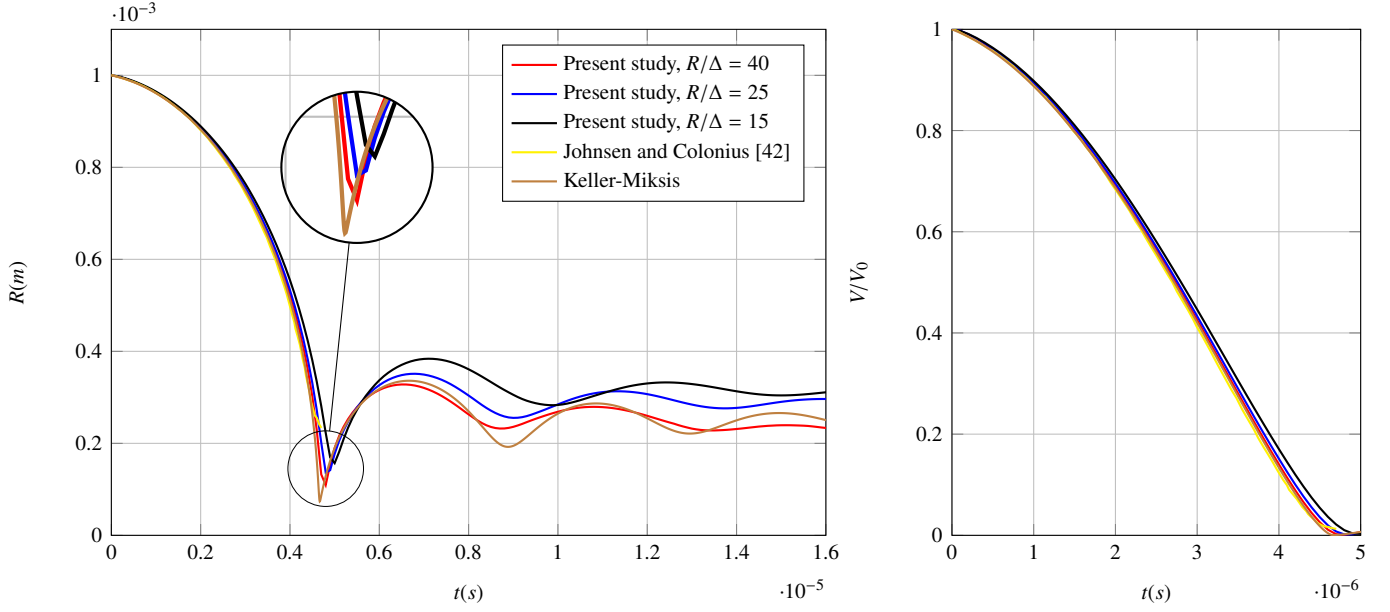


Figure 13: Evolution of the (left) average bubble radius and (right) normalized bubble volume along time for the symmetric collapse ($p_\infty/p_{g0} = 353$) of a $d = 0.002m$ air bubble in water. Comparison against the Keller-Miksis equation and the work of Johnsen and Colonius [42].

compressed, hence $p_B(t) = p_{g0} (R_0/R(t))^{3\gamma}$, where $R_0 = R(0)$ is the initial radius, p_{g0} is the initial pressure within the bubble and γ is the heat capacity ratio of air.

The results of our method fare particularly well (with respect to both convergence and size at collapse) compared to the work of Johnsen and Colonius [42], even though significantly lower resolutions are used. The numerical results are also in very good agreement with the predictions of the Keller-Miksis model [55]. We therefore conclude that the dynamics of the free-field Rayleigh collapse are accurately captured by our method.

Nonetheless, there are a few points that need to be discussed. When lower grid resolutions are employed, the collapse of the bubble at the final stages is slightly delayed. Overall, the subsequent rebounds and collapses of the bubble are captured by our method with success. However, in lower resolutions, the rebounds and oscillation periods are larger. This is partially attributed to the intensity of the preceding collapse phase; a smaller amount of energy is emitted in the form of waves, and is instead stored within the bubble as potential energy. It is also interesting to notice that strong damping of the oscillations is predicted in our simulations. It is typical to attribute such behavior to numerical diffusion or to the limited resolution of our simulations; from the late stages of the first collapse and onwards, the bubble is effectively discretized with a very small number of points. In this case however, it is the Keller-Miksis model that underestimates the damping of the oscillations, due to its thermodynamic assumptions (for example the assumption of constant pressure in the gas within the bubble).

An interesting phenomenon predicted by our simulations is the loss of sphericity of the interface. The sphericity Ψ is a measure of how closely the shape of an object approaches that of a perfect sphere and is defined as:

$$\Psi = \frac{\pi^{1/3} (6V)^{2/3}}{A}, \quad (18)$$

where V and A denote the volume and surface area of the object respectively. The sphericity of the bubble in our simulations is presented in figure 14. During the late stages of the first collapse the interface loses its sphericity and assumes a cross-like shape. The spherical shape is recovered at the first rebound, but following the second collapse we observe a complete loss of sphericity as the interface assumes a box-like shape. The interface of the bubble at these distinct stages is presented in figure 15, along with contours that represent the distance of a point from the origin of the axes.

This is a phenomenon that is well-known in the literature. Theoretical studies dating back to 1954 [56] suggest that bubbles undergoing volume changes do not always retain spherical symmetry. Sphericity loss was also reported in the work of Johnsen and Colonius [42]. In the much detailed work (at however different conditions) of Nagrath et al. [57], it was noted that during the violent collapse stage, the bubble deviated from spherical symmetry. Two types of instabilities were predicted: (a) Rayleigh-Taylor-like instabilities on the interface of the bubble and (b) shape instabilities. Such instabilities were also observed in the work of Kamran et al. [58].

An alternative and complementary interpretation as to the origins of the grid-shaped instabilities (figure 15) is proposed in this




Figure 14: Sphericity of the interface along time for the symmetric collapse ($p_{\infty}/p_{g0} = 353$) of a $d = 0.002m$ air bubble in water.

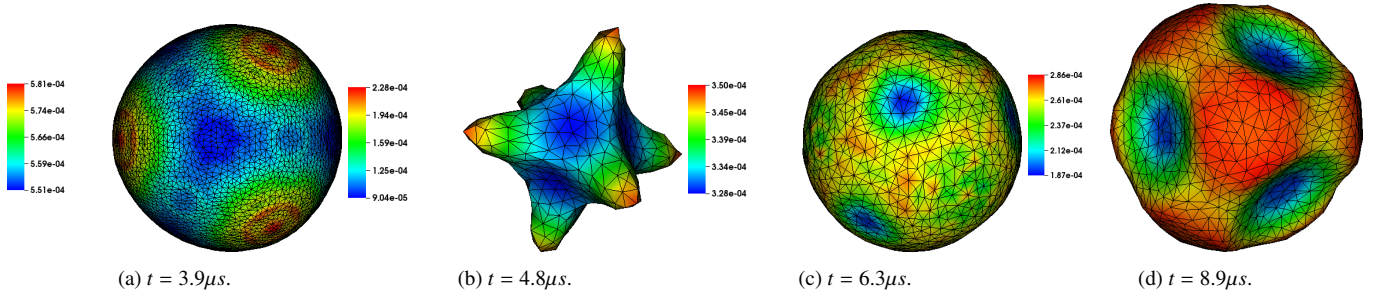


Figure 15: The interface of the bubble at various time instants, symmetric collapse ($p_{\infty}/p_{g0} = 353$) of a $d = 0.002m$ air bubble in water, 25ppbr.

work: we partly trace this to the numerical scheme (ours and others), similar to the discussion in section 3.3. Waves converge and diverge radially several times within the spherical volume and the numerical scheme is unable to resolve them accurately enough at the focal point, which by definition has a poor resolution of only a few cells (potentially only one). As a result, the scheme imparts on the (radial) waves its directionality properties. The perturbed waves then impact the interface and transmit the grid-shaped instabilities to the interface.

4.5. 2D simulation of the high-speed breakup of a water column in air

The shock-driven breakup of a cylindrical water droplet in air was the last problem to be considered for the validation of our method. This configuration is similar to the experiments of Igra and Takayama [59] which have been used by several researchers as a benchmark in order to validate their numerical methods [53, 60, 17, 61].

A cylindrical water droplet of diameter $d = 0.0048m$ was placed in a domain of size $x \in [-0.036m, 0.036m]$, $y \in [0m, 0.036m]$, in mechanical equilibrium with the surrounding air. Only half of the droplet and the surrounding domain were simulated, due to symmetry (about the x axis). The droplet was centered at $(x, y) = (0m, 0m)$. The density of water was set to $\rho_w = 1000kg/m^3$ and that of air at $\rho_g = 1.2kg/m^3$. Both fluids were initially at rest under atmospheric pressure conditions. The thermodynamic parameters that were used to model the two fluids are presented in Table 1. Post-shock conditions (corresponding to a $M = 1.47$ shock wave in air) were defined at the left boundary ($x = x_{min}$), while non-reflecting conditions were set at the two remaining boundaries ($x = x_{max}$ and $y = y_{max}$). The main features of the breakup are presented in figure 16, in which the contours of the magnitude of the density gradient are depicted at two different instants after shock impact. For validation purposes, the numerical Schlieren plots (obtained with a discretization corresponding to 50 points per initial droplet radius) are compared to the experimental results of Igra and Takayama [59]. Note that the values of time have been adjusted as discussed in Meng and Colonius [61].

For a quantification of the results, the drift of the cylinder (based on its most upstream point) is plotted in figure 17 for different grid resolutions. The results of our method are also compared against the experimental results of Igra and Takayama [59] and the numerical results of Chen [60].




Figure 16: Visualization of the high-speed droplet breakup process. (Left): Present study (50ppbr), numerical Schlieren contours, (right): Holographic interferograms [59]. Values of time have been adjusted as discussed in Meng and Colonius [61].




Figure 17: Evolution of the cylinder's drift (based on its most upstream point) along time after shock impact (without applying the time adjustment procedure described in [61]), 2D simulation of the high-speed breakup of a water column in air. Comparison against the experimental results of Igra and Takayama [59] and the numerical results of Chen [60].

5. Shock induced collapse of a triangular bubble array

Finally, the capabilities of the developed method are demonstrated by simulating a complex three-dimensional case. Drawing inspiration from the work of Betney et al. [15], the shock induced collapse of a triangular array of three bubbles was considered, where one bubble is smaller than the two others. In that work, bubbles were assumed to be cylindrical, and the simulations were thus carried out in two dimensions. In the present work, we study a similar configuration in three dimensions, i.e. considering spherical bubbles. Significant levels of energy focusing may be achieved through this arrangement. The large bubbles are the first to be impacted by the incident shock wave. At the same time, they shield the small central bubble. As a result, the small bubble collapses at a later stage, after being struck by the (much stronger) waves that are generated at the collapse of the large bubbles.

The radii of the large and small bubbles were set to $R_B = 0.0005m$ and $R_b = 0.0002m$ respectively (size ratio $R_B/R_b = 2.5$). The domain was discretized with 75 points per large bubble radius (30 points for the small one), resulting in a mesh consisting of 40,500,000 elements. The two large bubbles were centered at $(x, y, z) = (0.00055m, 0m, 0.00075m)$ and $(x, y, z) = (-0.00055m, 0m, 0.00075m)$, whereas the small one was placed at $(x, y, z) = (0m, 0m, 0.00125m)$. The properties and states of the bubbles and the surrounding medium were set as described in section 4.2.

The collapse process is depicted in figure 18 at different time instants after shock impact (an animation of this process, marked as Video 2, is also available in Supplementary Material). The collapse of the large bubbles is not symmetric: the lobes on the inner

side get compressed in delay compared to their free-field counterparts (figure 18a). This occurs as the incident shockwave weakens in the region between the bubbles, due to its interaction with the reflected rarefaction waves. The protection offered by the large bubbles to the small one can be seen in figure 18b. The central bubble only starts collapsing following the collapse of the main bubbles. Compression is more pronounced on its sides, resulting in a parachute-like shape, as seen in figure 18c. At the moment of lateral impact, strong peaks in pressure are generated. The large bubbles have assumed a horseshoe-shaped form. The gas within them is less compressed the farther we move away from the vicinity of the small bubble, where several strong wave interactions take place. Finally, the small bubble is pierced in half, and another pressure peak occurs (central bubble's water-hammer shock wave). At later times, vorticity dominates the motion of the bubble remnants (figure 18d).




Figure 18: Volume rendering of the magnitude of the density gradient at different time instants after shock impact, 3D simulation of the 1GPa shock induced collapse of a triangular bubble array configuration.

The characteristic measures for the collapse are presented in Table 6. These include the first phase time and water-hammer shock strength for the large bubbles, and the pressure observed at the moments of lateral and transverse collapse of the small bubble. In comparison with the spherical bubble collapse in a free field (section 4.3, Table 5), we observe that the first phase time of the large bubbles remains unaffected by the presence of the other bubbles (it is only slightly faster due to the increased resolution). Similar comments apply to the pressure of the water-hammer shock waves, as they are of equal strength. High pressures are found at the moment of lateral collapse of the small bubble. A peak in the pressure equal to ($p_{max} = 21.20GPa$) is found at the water-hammer shock of the small bubble. In general, significant levels of pressure amplification are achieved: four times greater than the water-hammer shock of a single bubble in a free-field, and more than twenty times the pressure of the incident shock.

The triangular array configuration was also simulated (keeping the same resolution) in two dimensions. The characteristic measures for the two-dimensional problem are also presented in Table 6. As mentioned in section 4.3, the collapse is a much faster process in three dimensions and the generated shock pressures are quite stronger. As a result, the focusing achieved in three dimensions is more intense.

Table 6: Characteristic collapse measures, 2D and 3D simulations of the 1GPa shock induced collapse of a triangular bubble array configuration. Reported times are measured after the impact of the incident shock on the large bubbles.

| | 3D | 2D |
|---|-------|------|
| Large bubble first phase time (<i>ns</i>) | 560 | 710 |
| Large bubble water-hammer shock pressure (<i>GPa</i>) | 5.31 | 2.89 |
| Central bubble lateral collapse time (<i>ns</i>) | 760 | 1030 |
| Lateral collapse pressure (<i>GPa</i>) | 14.65 | 5.13 |
| Central bubble transverse collapse time (<i>ns</i>) | 800 | 1070 |
| Central bubble water-hammer shock pressure (<i>GPa</i>) | 21.20 | 8.69 |

The mechanism that drives the collapse of the small bubble is the water-hammer shockwaves from the large bubbles. It thus makes sense to also measure the amplification relative to these waves instead of the incident one. This comparison shows that pressure amplification relative to the water-hammer shocks of the large bubbles is stronger in three dimensions.

It is important to note that our findings are complementary to those of Betney et al. [15]: a slightly different arrangement regarding bubble placement and size ratio is studied in that paper. We confirm the assertion that the exact positioning of bubbles plays an important role regarding the exact timing and magnitude of the various events and interactions, but also the order in which these events take place. For example, in the case considered in the work of Betney et al. [15], the peak overall pressure was found to occur at the moment of lateral impact, whereas in the configuration simulated in the current work it was found at the water hammer shocking of the smaller bubble. It is therefore further emphasized that the exact focusing mechanisms depend on the details of the bubble arrangement. This comment applies to both 2D and 3D configurations.

Further to this, the pressure peaks reported in the work of Betney et al. [15] were almost 40 times higher when compared to the incident shock strength. Our simulations show that in the full 3D (spherical bubble) configurations, this number can be even larger, a result which is of significant value in fields like cavitation damage and lithotripsy/histotripsy.

6. Conclusions

A novel way for the definition of the inter-component boundary conditions in sharp interface methods was proposed. The formulation is remarkably simple, but at the same time able to resolve any wave-interface interactions with accuracy that is close or equal to existing methods. By fully adopting the stair-step representation of an interface in a rectangular grid and using the interface states as boundary conditions the derived method is very simple; all preparatory procedures for the construction and population of the ghost cells are circumvented. More importantly, the formulation is markedly simple regarding its implementation. The proposed approach for the definition of inter-component boundary conditions in higher dimensions is general; it may be combined with any ghost fluid method and Riemann solver in order to approximate the interface state with the desired accuracy. Moreover, it can be combined with a level set technique for the localization of the interface, rendering the procedure of extending any single-phase solver to multiple phases trivial. However, thorough validation of such a method is a crucial step, due to the lower accuracy of level sets compared to front tracking (see section 3.1).

The behavior and accuracy of the method were analyzed using a series of standard test cases. Our method was shown to be capable of following the evolution of interfaces with great detail and to be consistent and convergent with grid refinement. Validation of the method continued by studying the collapse of bubbles and the breakup of droplets; the framework was shown to be performing very well in capturing the dynamics even at very low resolutions. The grid-aligned formulation for the ghost fluid method, which was the main point of this work, was shown to be capable of yielding accurate results without generating oscillations at the interface.

In the future, the developed methodology will be employed to carry out large scale simulations in three dimensions similar to the application studied in section 5. This will allow us to study the complex flow patterns in problems involving shockwave-interface interactions with great detail. Moreover, the grid-aligned formulation will be extended to include the effects of surface tension. We are confident that the grid-aligned formulation could accommodate a consistent implementation of capillary effects within a sharp-interface compressible flow framework. Finally, another path for future development would be considering adaptive mesh refinement (AMR) techniques, the implementation of which is significantly simplified by the proposed grid-aligned formulation.

Acknowledgments

This work was carried out in the framework of HAoS project, which has received funding from the European Union Horizon 2020 Research and Innovation programme, grant agreement No 675676. The authors would like to express their gratitude to Professors J. Glimm and X. Li of Stony Brook University for the use of their front tracking framework. Nikolaos Bempedelis kindly acknowledges the support of the Onassis Foundation.

References

- [1] C. W. Hirt, B. D. Nichols, Volume of fluid (VOF) method for the dynamics of free boundaries, *Journal of Computational Physics* 39 (1981) 201–225.
- [2] S. Osher, J. A. Sethian, Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations, *Journal of Computational Physics* 79 (1988) 12–49.
- [3] S. Osher, R. P. Fedkiw, Level set methods: an overview and some recent results, *Journal of Computational Physics* 169 (2001) 463–502.
- [4] J. Glimm, E. Isaacson, D. Marchesin, O. McBryan, Front tracking for hyperbolic systems, *Advances in Applied Mathematics* 2 (1981) 91–119.
- [5] J. Glimm, O. A. McBryan, A computational model for interfaces, *Advances in Applied Mathematics* 6 (1985) 422–435.
- [6] J. Glimm, J. Grove, B. Lindquist, O. A. McBryan, G. Tryggvason, The bifurcation of tracked scalar waves, *SIAM Journal on Scientific and Statistical Computing* 9 (1988) 61–79.
- [7] S. O. Unverdi, G. Tryggvason, A front-tracking method for viscous, incompressible, multi-fluid flows, *Journal of Computational Physics* 100 (1992) 25–37.
- [8] R. Abgrall, S. Karni, Computations of compressible multifluids, *Journal of Computational Physics* 169 (2001) 594–623.
- [9] J. W. Grove, R. Menikoff, The anomalous reflection of a shock wave at a fluid interface, *Journal of Fluid Mechanics* 219 (1990) 313–336.
- [10] K. Haller, Y. Ventikos, D. Poulikakos, P. Monkewitz, Computational study of high-speed liquid droplet impact, *Journal of Applied Physics* 92 (2002) 2821–2828.
- [11] X. Liu, E. George, W. Bo, J. Glimm, Turbulent mixing with physical mass diffusion, *Physical Review E* 73 (2006) 056301.
- [12] X. Liu, Y. Li, J. Glimm, X. Li, A front tracking algorithm for limited mass diffusion, *Journal of Computational Physics* 222 (2007) 644–653.
- [13] H. Lim, J. Iwerks, Y. Yu, J. Glimm, D. Sharp, Verification and validation of a method for the simulation of turbulent mixing, *Physica Scripta* 2010 (2010) 014014.
- [14] N. A. Hawker, Y. Ventikos, Interaction of a strong shockwave with a gas bubble in a liquid medium: a numerical study, *Journal of Fluid Mechanics* 701 (2012) 59–97.
- [15] M. R. Betney, N. A. Hawker, B. Tully, Y. Ventikos, Computational modelling of the interaction of shock waves with multiple gas bubbles in a liquid, *Physics of Fluids* 27 (2015) 036101.
- [16] Y. Hao, A. Prosperetti, A numerical method for three-dimensional gas-liquid flow computations, *Journal of Computational Physics* 196 (2004) 126–144.
- [17] H. Terashima, G. Tryggvason, A front-tracking/ghost-fluid method for fluid interfaces in compressible flows, *Journal of Computational Physics* 228 (2009) 4012–4037.
- [18] H. Terashima, G. Tryggvason, A front-tracking method with projected interface conditions for compressible multi-fluid flows, *Computers & Fluids* 39 (2010) 1804–1814.
- [19] W. Bo, X. Liu, J. Glimm, X. Li, A robust front tracking method: Verification and application to simulation of the primary breakup of a liquid jet, *SIAM Journal on Scientific Computing* 33 (2011) 1505–1524.
- [20] D. Wang, N. Zhao, O. Hu, J. Liu, A ghost fluid based front tracking method for multimedium compressible flows, *Acta Mathematica Scientia* 29 (2009) 1629–1646.
- [21] H. Lu, N. Zhao, D. Wang, A front tracking method for the simulation of compressible multimedium flows, *Communications in Computational Physics* 19 (2016) 124–142.
- [22] R. P. Fedkiw, T. Aslam, B. Merriman, S. Osher, A non-oscillatory Eulerian approach to interfaces in multimaterial flows (the ghost fluid method), *Journal of Computational Physics* 152 (1999) 457–492.
- [23] R. P. Fedkiw, T. Aslam, S. Xu, The ghost fluid method for deflagration and detonation discontinuities, *Journal of Computational Physics* 154 (1999) 393–427.
- [24] R. P. Fedkiw, Coupling an Eulerian fluid calculation to a Lagrangian solid calculation with the ghost fluid method, *Journal of Computational Physics* 175 (2002) 200–224.
- [25] T. Liu, B. Khoo, K. Yeo, The simulation of compressible multi-medium flow. I. A new methodology with test applications to 1D gas–gas and gas–water cases, *Computers & Fluids* 30 (2001) 291–314.
- [26] T. Liu, B. Khoo, K. Yeo, Ghost fluid method for strong shock impacting on material interface, *Journal of Computational Physics* 190 (2003) 651–681.
- [27] T. Liu, B. Khoo, C. Wang, The ghost fluid method for compressible gas–water simulation, *Journal of Computational Physics* 204 (2005) 193–221.
- [28] X. Y. Hu, B. C. Khoo, An interface interaction method for compressible multifluids, *Journal of Computational Physics* 198 (2004) 35–64.

- [29] C. Wang, T. Liu, B. Khoo, A real ghost fluid method for the simulation of multimediuim compressible flow, *SIAM Journal on Scientific Computing* 28 (2006) 278–302.
- [30] S. K. Sambasivan, H. Udaykumar, Ghost fluid method for strong shock interactions Part 1: Fluid-fluid interfaces, *AIAA Journal* 47 (2009) 2907–2922.
- [31] L. Xu, T. Liu, Accuracies and conservation errors of various ghost fluid methods for multi-medium Riemann problem, *Journal of Computational Physics* 230 (2011) 4975–4990.
- [32] J. Glimm, J. W. Grove, X. L. Li, K.-M. Shyue, Y. Zeng, Q. Zhang, Three-dimensional front tracking, *SIAM Journal on Scientific Computing* 19 (1998) 703–727.
- [33] J. Glimm, J. W. Grove, X. Li, N. Zhao, Simple front tracking, *Contemporary Mathematics* 238 (1999) 133–149.
- [34] J. Glimm, J. W. Grove, X. Li, D. C. Tan, Robust computational algorithms for dynamic interface tracking in three dimensions, *SIAM Journal on Scientific Computing* 21 (2000) 2240–2256.
- [35] J. Glimm, J. W. Grove, Y. Zhang, Interface tracking for axisymmetric flows, *SIAM Journal on Scientific Computing* 24 (2002) 208–236.
- [36] J. Du, B. Fix, J. Glimm, X. Jia, X. Li, Y. Li, L. Wu, A simple package for front tracking, *Journal of Computational Physics* 213 (2006) 613–628.
- [37] R. Menikoff, B. J. Plohr, The Riemann problem for fluid flow of real materials, *Reviews of Modern Physics* 61 (1989) 75–130.
- [38] W. Wagner, A. Pruß, The IAPWS formulation 1995 for the thermodynamic properties of ordinary water substance for general and scientific use, *Journal of Physical and Chemical Reference Data* 31 (2002) 387–535.
- [39] K. Nagayama, Y. Mori, K. Shimada, M. Nakahara, Shock Hugoniot compression curve for water up to 1 GPa by using a compressed gas gun, *Journal of Applied Physics* 91 (2002) 476–482.
- [40] G. J. Ball, B. P. Howell, T. G. Leighton, M. J. Schofield, Shock-induced collapse of a cylindrical air cavity in water: a Free-Lagrange simulation, *Shock Waves* 10 (2000) 265–276.
- [41] E. Lauer, X. Y. Hu, S. Hickel, N. A. Adams, Numerical investigation of collapsing cavity arrays, *Physics of Fluids* 24 (2012).
- [42] E. Johnsen, T. Colonius, Numerical simulations of non-spherical bubble collapse, *Journal of Fluid Mechanics* 629 (2009) 231.
- [43] V. Coralic, T. Colonius, Finite-volume WENO scheme for viscous compressible multicomponent flows, *Journal of Computational Physics* 274 (2014) 95–121.
- [44] B. Tully, N. Hawker, Y. Ventikos, Modeling asymmetric cavity collapse with plasma equations of state, *Physical Review E* 93 (2016) 1–6.
- [45] X. Hu, N. A. Adams, G. Iaccarino, On the HLLC Riemann solver for interface interaction in compressible multi-fluid flow, *Journal of Computational Physics* 228 (2009) 6572–6589.
- [46] C.-W. Shu, Essentially non-oscillatory and weighted essentially non-oscillatory schemes for hyperbolic conservation laws, in: *Advanced numerical approximation of nonlinear hyperbolic equations*, Springer, 1998, pp. 325–432.
- [47] C. Farhat, J.-F. Gerbeau, A. Rallu, FIVER: A finite volume method based on exact two-phase Riemann problems and sparse grids for multi-material flows with large density jumps, *Journal of Computational Physics* 231 (2012) 6360–6379.
- [48] S. T. Zalesak, Fully multidimensional flux-corrected transport algorithms for fluids, *Journal of Computational Physics* 31 (1979) 335–362.
- [49] D. Zuzio, Direct numerical simulation of two phase flows with adaptive mesh refinement, Ph.D. thesis, Ecole nationale supérieure de l’aéronautique et de l’espace, 2010.
- [50] J.-F. Haas, B. Sturtevant, Interaction of weak shock waves with cylindrical and spherical gas inhomogeneities, *Journal of Fluid Mechanics* 181 (1987) 41.
- [51] J. M. Picone, J. P. Boris, Vorticity generation by shock propagation through bubbles in a gas, *Journal of Fluid Mechanics* 189 (1988) 23–51.
- [52] J. J. Quirk, S. Karni, On the dynamics of a shock–bubble interaction, *Journal of Fluid Mechanics* 318 (1996) 129–163.
- [53] R. R. Nourgaliev, T.-N. Dinh, T. G. Theofanous, Adaptive characteristics-based matching for compressible multifluid dynamics, *Journal of Computational Physics* 213 (2006) 500–529.
- [54] S. K. Shankar, S. Kawai, S. K. Lele, Numerical simulation of multicomponent shock accelerated flows and mixing using localized artificial diffusivity method, *AIAA Paper* 2010-352 (2010) 2010.
- [55] J. B. Keller, M. Miksis, Bubble oscillations of large amplitude, *The Journal of the Acoustical Society of America* 68 (1980) 628–633.
- [56] M. Plesset, On the stability of fluid flows with spherical symmetry, *Journal of Applied Physics* 25 (1954) 96–98.
- [57] S. Nagrath, K. Jansen, R. T. Lahey, I. Akhatov, Hydrodynamic simulation of air bubble implosion using a level set approach, *Journal of Computational Physics* 215 (2006) 98–132.
- [58] K. Kamran, R. Rossi, E. Oñate, S. Idelsohn, A compressible Lagrangian framework for the simulation of the underwater implosion of large air bubbles, *Computer Methods in Applied Mechanics and Engineering* 255 (2013) 210–225.
- [59] D. Igra, K. Takayama, Investigation of aerodynamic breakup of a cylindrical water droplet, *Atomization and Sprays* 11 (2001).

- [60] H. Chen, Two-dimensional simulation of stripping breakup of a water droplet, *AIAA Journal* 46 (2008) 1135–1143.
- [61] J. Meng, T. Colonius, Numerical simulations of the early stages of high-speed droplet breakup, *Shock Waves* 25 (2015) 399–414.