

A projection-based, semi-implicit time-stepping approach for the Cahn-Hilliard Navier-Stokes equations on adaptive octree meshes

Makrand A. Khanwale^{a,b,1}, Kumar Saurabh^a, Masado Ishii^c, Hari Sundar^c,
James A. Rossmanith^b, Baskar Ganapathysubramanian^{a,*}

^a Department of Mechanical Engineering, Iowa State University, Iowa, 50011, USA

^b Department of Mathematics, Iowa State University, Iowa, 50011, USA

^c School of Computing, The University of Utah, Salt Lake City, UT, 84112, USA

ARTICLE INFO

Article history:

Received 3 May 2022

Received in revised form 14 September 2022

Accepted 19 December 2022

Available online 30 December 2022

Keywords:

Two-phase flows

Energy stable

Adaptive finite elements

Octrees

Variational multiscale approach

ABSTRACT

The Cahn-Hilliard Navier-Stokes (CHNS) system provides a computationally tractable model that can be used to effectively capture interfacial dynamics in two-phase fluid flows. In this work we present a semi-implicit, projection-based finite element framework for solving the CHNS system. We use a projection-based semi-implicit time discretization for the Navier-Stokes equation and a fully-implicit time discretization for the Cahn-Hilliard equation. We use a conforming continuous Galerkin (cG) finite element method in space equipped with a residual-based variational multiscale (RBVMS) formulation. Pressure is decoupled using a projection step, which results in two linear positive semi-definite systems for velocity and pressure, instead of the saddle point system of a pressure-stabilized method. All the linear systems are solved using an efficient and scalable algebraic multigrid (AMG) method. We deploy this approach on a massively parallel numerical implementation using parallel octree-based adaptive meshes. The overall approach allows the use of relatively large time steps with much faster time-to-solve than similar fully-implicit methods. We present comprehensive numerical experiments showing detailed comparisons with results from the literature for canonical cases, including the single bubble rise and Rayleigh-Taylor instability.

© 2022 Elsevier Inc. All rights reserved.

1. Introduction

Capturing the interfacial dynamics is crucial for gaining fundamental understanding of two-phase flows. In several applications like bio-microfluidics and advanced manufacturing such interface resolved simulations are critical for analysis and design [1–3]. More generally, the availability of robust and fast interface-resolved two-phase simulations can greatly enable the development of (data-driven) coarse-scale models that need not be interface-resolving. This is the primary motivation of the current work [4–7].

* Corresponding author.

E-mail addresses: khanwale@iastate.edu (M.A. Khanwale), maksbh@iastate.edu (K. Saurabh), masado@cs.utah.edu (M. Ishii), hari@cs.utah.edu (H. Sundar), rossmani@iastate.edu (J.A. Rossmanith), baskarg@iastate.edu (B. Ganapathysubramanian).

¹ Currently at Center for Turbulence Research, Department of Mechanical Engineering, Stanford University, CA, USA 94305.

We are particularly interested in the coupled Cahn-Hilliard Navier-Stokes equations as a means to capture interfacial dynamics. Cahn-Hilliard Navier-Stokes (CHNS) models have received increasing attention for their ability to capture interface resolved two-phase phenomena in various applications [8–12]. Such an approach provides a thermodynamically consistent description of the interfacial processes and evolution [13]. CHNS belongs to a family of models where a diffused field (also known as phase field) is used to track the interface. Such an approach circumvents modeling the jump discontinuities at the interface and allows for a diffuse transition between the physical properties from one phase to the other. Using Cahn-Hilliard equations to track the interface produces several advantages including mass conservation, thermodynamic consistency, and a free-energy-based description of surface tension with a well-established theory from non-equilibrium thermodynamics [14,15]. However, numerical schemes need to be carefully designed to allow the discrete numerical solutions of these CHNS models to inherit these continuous properties. In particular, CHNS models that use a volume averaged mixture velocity are attractive as they ensure a divergence-free mixture velocity, which, crucially, allows numerical schemes for such CHNS models to be designed by appealing to the extensive literature on incompressible single phase flows (see Volker [16]). Therefore, in this work, we continue to use a volume averaged mixture velocity for numerical development.²

We previously developed an energy-stable, fully-coupled, second order scheme in Khanwale et al. [19]. This scheme uses an equal order interpolation for pressure and velocity and uses a pressure stabilization based on the variational multi-scale method. While the approach of [19] is well-suited for systems affording naturally large timesteps (or where steady state solutions are desired), the coupled non-linear solves become expensive for turbulent systems, especially when small time steps are naturally required. Moreover, the analysis in [19] suggests that a fully-coupled pressure stabilized approach requires very careful design of preconditioners for efficient execution of large scale production simulations. To address these issues we seek to improve the results of [19] in three key aspects.

1. **Projection scheme for decoupled pressure solve:** We modify the fully coupled-scheme presented in Khanwale et al. [19] by decoupling the pressure equation from the velocity equations using a projection method; this is done in a manner that preserves second-order accuracy, energy-stability, and mass conservation. In conjunction with a block solution scheme, this results in elliptic (positive semi-definite) operators. Additionally, a projection variant of VMS method to bypass the discrete inf-sup condition is presented.
2. **Improved octree-based adaptive mesh:** We improve upon the octree based meshing framework presented in Khanwale et al. [19] through the use of a mesh-free kD tree construction. This new approach yields improved parallel performance [20,21] with a smaller memory footprint, and provides support for simulating on non-cubic domains via incomplete octrees.
3. **Scaling and speedup comparison:** We present a detailed scaling and performance analysis of the projection method on up to 7000 processors and compare it to the fully-coupled method presented in [19] to assess relative speedup.

1.1. Second-order energy-stable scheme projection scheme

We deploy our numerical method on adaptive meshes (see section 1.3), which can result in very restrictive time-steps for explicit schemes (like in Badalassi et al. [22]). This is in part due to the viscous and surface tension terms in Navier-Stokes, even if the Cahn-Hilliard equations are discretized using implicit or semi-implicit linearized methods [23–25]. Unlike other related works Badalassi et al. [22], Kim et al. [23], Yue et al. [26], we develop a numerical method for a thermodynamically consistent CH-NS model [27], which allows for unequal densities. This results in a variable coefficient pressure Poisson system that is carefully constructed such that we can use state-of-the-art Algebraic Multigrid solvers from PETSc.

There have been a number of attempts to use projection-based approaches to model CHNS models [17,24,28–33]. Shen and Yang [24,28,30] used a block-solver strategy with linearized time-schemes that reduce their discretization to a sequence of elliptic equations for the velocity and phase fields; however, in these papers the numerical methods were limited to first order provably energy-stable schemes. An upside of the first order nature of these schemes is the need to solve each block only once. However, the velocity used to advect the phase field in [24,28,30] is not guaranteed to be solenoidal, and needs to be modified to maintain energy-stability. A simplifying assumption made in [24,28,30] was to use a constant coefficient pressure Poisson equation.³ This is an elegant simplification, particularly for cases where the system behavior is not very sensitive to the density ratio.⁴ In this work we relax this assumption – with an eye on applications where the density ratio critically determines dynamics – and use the mixture density resulting in a variable coefficient pressure Poisson equation. We use a well-tuned Algebraic Multigrid (AMG) method to efficiently solve this system.

Zhu et al. [33] extended stabilized linearization of Cahn-Hilliard presented in [24,28,30] to a Scalar Auxiliary Variable (SAV) linearized Cahn-Hilliard time-scheme. The stabilized and SAV formulations for developing linear semi-implicit time schemes for the Cahn-Hilliard part of the CHNS models are very effective; however, they introduce tight time-step requirements. Instead of trying to linearize the Cahn-Hilliard equations, we use a fully implicit non-linear time-scheme similar

² CHNS models with mass averaged mixture velocity have also been explored in the literature [17,18].

³ The Poisson equation was made constant coefficient by replacing the variable specific density by the minimum of the specific densities of the two fluids.

⁴ As an example, the dynamics in the case of a rising bubble does not change once the density ratio is increased beyond a certain asymptotic value.

to Han and Wang [34] that allows us to take larger timesteps. This choice is a trade-off between time-step restrictions and complexity of each time step – both of which impact the total time-to-solve. Interestingly, we find that the fully-implicit non-linear scheme for the Cahn-Hilliard equations reduces to a one step Newton iteration for smaller timesteps, making it computationally equivalent to linear formulations.

Other related work: Han and Wang [34] used a block-iterative strategy with an energy-stable time scheme but with a non-linear scheme for Cahn-Hilliard. However, the time scheme in Han and Wang [34] is limited to the equal density case; and solenoidality of the velocity used for the advective phase field is not guaranteed. Guo et al. [17] recently reported a detailed analysis for a mass-averaged mixture velocity CHNS system using a fully coupled strategy using both projection and pressure coupled methods. While we do not include a rigorous proof of energy stability (we defer this to a subsequent publication), we numerically confirm energy-stability for a variety of canonical benchmarks.⁵

1.2. VMS-based stabilization for conforming Galerkin elements

Generally projection methods enforce solenoidality of the mixture velocity through a Helmholtz-Hodge decomposition step [35]. When we take the divergence of the Helmholtz-Hodge decomposition step we obtain the pressure Poisson equation used to update the pressure. When using continuous Galerkin finite elements and equal order polynomial basis functions for velocity and pressure, the discrete inf-sup condition [16] is not satisfied, which can lead to checkerboard instabilities for pressure when solving the pressure Poisson step. Therefore, a pressure stabilization approach is necessary even when using the projection approach. However, continuous Galerkin methods are attractive as they are comparatively easy to implement with adaptive octree based meshes [20,36]. To circumvent the inf-sup condition and stabilize the pressure Poisson equation we use a residual based VMS [37–39] approach. Moreover, the VMS approach allows for LES-type modeling of the flow physics as the mesh is coarsened away from the interface.

1.3. Scalable octree-based adaptive mesh generation

Adaptive spatial discretizations are popular in computational science [40–42] to improve efficiency and resolution quality. In several large-scale applications [20,43,44], the feasibility of high-fidelity simulations on modern supercomputers is primarily due to the localized resolution of adaptive discretizations. There are computational challenges to achieve spatial adaptivity in the distributed-memory setting, such as load-balancing, low-cost mesh generation, and mesh-partitioning. The scalability of the algorithms used for mesh generation and partitioning is crucial, especially when the problem requires frequent re-meshing. This is only possible when the overheads of solving a problem on an adaptive mesh are significantly lower than the savings achieved from the reduction of problem-size. Octrees are widely used in the community [20,44–46] due to their simplicity and their extreme parallel scalability.

In this work we employ an updated version of a parallel octree library, called DENDRO-KT, that provides high-quality space-adaptive resolution along with efficient mesh generation, partitioning, and traversal methods. The DENDRO-KT framework is a freely available open-source library. The DENDRO-KT library supports 2D, 3D, and 4D trees, and has been used to manage spacetime-adaptive discretizations for solving time-dependent PDEs [20]. A distinguishing feature of DENDRO-KT is the elimination of the element-to-nodal maps, instead giving access to nodal data at the finite elements using top-down and bottom-up traversals. The benefits of this are minimizing data footprint, avoiding indirect memory accesses, and avoiding computing complex topological maps for complex domains. Additionally, DENDRO-KT efficiently manages octrees over non-isotropic carved-out domains [21]. The methods of mesh-generation, partitioning, and traversal are detailed in section 4.

2. Governing equations

Consider a bounded domain $\Omega \subset \mathbb{R}^d$, for $d = 2, 3$ containing two immiscible fluids, and the time interval $[0, T]$. Let ρ_+ (η_+) and ρ_- (η_-) denote the specific density (viscosity) of the two phases. Let the phase field function, ϕ , be the variable that tracks the location of the phases and varies smoothly between $+1$ to -1 . We define a non-dimensional density, $\rho(\phi)$, such that

$$\rho(\phi) = \alpha\phi + \beta, \quad \text{where} \quad \alpha = \frac{\rho_+ - \rho_-}{2\rho_+} \quad \text{and} \quad \beta = \frac{\rho_+ + \rho_-}{2\rho_+}. \quad (1)$$

Note that our non-dimensional form uses the specific density of fluid 1 as the non-dimensionalizing density. Similarly, the non-dimensional viscosity is defined as

$$\eta(\phi) = \gamma\phi + \xi, \quad \text{where} \quad \gamma = \frac{\eta_+ - \eta_-}{2\eta_+} \quad \text{and} \quad \xi = \frac{\eta_+ + \eta_-}{2\eta_+}. \quad (2)$$

Let, $\mathbf{v}$ be the volume averaged mixture velocity, p the volume averaged pressure, ϕ the phase field (interface tracking variable), and μ the chemical potential. Then the governing equations in their non-dimensional form are as follows:

⁵ We note that the current method is constructed by decoupling pressure and velocity in the provably energy-stable method of Khanwale et al. [19].

$$\begin{aligned} \text{Momentum Eqns: } & \frac{\partial (\rho(\phi)v_i)}{\partial t} + \frac{\partial (\rho(\phi)v_i v_j)}{\partial x_j} + \frac{1}{Pe} \frac{\partial (J_j v_i)}{\partial x_j} + \frac{Cn}{We} \frac{\partial}{\partial x_j} \left(\frac{\partial \phi}{\partial x_i} \frac{\partial \phi}{\partial x_j} \right) \\ & + \frac{1}{We} \frac{\partial p}{\partial x_i} - \frac{1}{Re} \frac{\partial}{\partial x_j} \left(\eta(\phi) \frac{\partial v_i}{\partial x_j} \right) - \frac{\rho(\phi) \hat{g}_i}{Fr} = 0, \end{aligned} \quad (3)$$

$$\text{Thermo Consistency: } J_i = \frac{(\rho_- - \rho_+)}{2 \rho_+ Cn} m(\phi) \frac{\partial \mu}{\partial x_i}, \quad (4)$$

$$\text{Solenoidality: } \frac{\partial v_i}{\partial x_i} = 0, \quad (5)$$

$$\text{Continuity: } \frac{\partial \rho(\phi)}{\partial t} + \frac{\partial (\rho(\phi)v_i)}{\partial x_i} + \frac{1}{Pe} \frac{\partial J_i}{\partial x_i} = 0, \quad (6)$$

$$\text{Chemical Potential: } \mu = \psi'(\phi) - Cn^2 \frac{\partial}{\partial x_i} \left(\frac{\partial \phi}{\partial x_i} \right), \quad (7)$$

$$\text{Cahn-Hilliard Eqn: } \frac{\partial \phi}{\partial t} + \frac{\partial (v_i \phi)}{\partial x_i} - \frac{1}{Pe Cn} \frac{\partial}{\partial x_i} \left(m(\phi) \frac{\partial \mu}{\partial x_i} \right) = 0. \quad (8)$$

Non-dimensional mobility $m(\phi)$ is assumed to be a constant value of one. $\hat{\mathbf{g}}$ is a unit vector defined as $(0, -1, 0)$ denoting the direction of gravity. We use the polynomial form of the free energy density $\psi(\phi(\mathbf{x}))$ defined as follows:

$$\psi(\phi) = \frac{1}{4} (\phi^2 - 1)^2 \quad \text{with} \quad \psi'(\phi) = \phi^3 - \phi. \quad (9)$$

2.1. Details of non-dimensionalization

Let u_r and L_r denoting the reference velocity and length, m_r is the reference mobility, σ is the scaling interfacial tension, $v_r = \eta_+/\rho_+$, ε is the interface thickness, g is gravitational acceleration. Then relevant non-dimensional parameters are as follows: Peclet, $Pe = \frac{u_r L_r^2}{m_r \sigma}$; Reynolds, $Re = \frac{u_r L_r}{\nu_r}$; Weber, $We = \frac{\rho_r u_r^2 L_r}{\sigma}$; Cahn, $Cn = \frac{\varepsilon}{L_r}$; and Froude, $Fr = \frac{u_r^2}{g L_r}$, with u_r and L_r denoting the reference velocity and length, respectively. We use scaling relation suggested by [47] of $1/Pe = 3Cn^2$ to calculated appropriate Pe for a chosen Cn .⁶

2.2. Energy law

The system of equations eq. (3) – eq. (8) has a dissipative law given by:

$$\frac{dE_{\text{tot}}}{dt} = -\frac{1}{Re} \int_{\Omega} \frac{\eta(\phi)}{2} \|\nabla \mathbf{v}\|_F^2 d\mathbf{x} - \frac{Cn}{We} \int_{\Omega} m(\phi) \|\nabla \mu\|^2 d\mathbf{x}, \quad (10)$$

where the total energy is

$$E_{\text{tot}}(\mathbf{v}, \phi, t) = \int_{\Omega} \frac{1}{2} \rho(\phi) \|\mathbf{v}\|^2 d\mathbf{x} + \frac{1}{Cn We} \int_{\Omega} \left(\psi(\phi) + \frac{Cn^2}{2} \|\nabla \phi\|^2 + \frac{1}{Fr} \rho(\phi) y \right) d\mathbf{x}. \quad (11)$$

The norms used in the above expression are the Euclidean vector norm and the Frobenius matrix norm:

$$\|\mathbf{v}\|^2 := \sum_i |v_i|^2 \quad \text{and} \quad \|\nabla \mathbf{v}\|_F^2 := \sum_i \sum_j \left| \frac{\partial v_i}{\partial x_j} \right|^2. \quad (12)$$

This energy functional does not consider energy fluxes into/out of the boundaries. This scenario is representative of closed systems, systems with neutral contact line dynamics, or generally boundary conditions that do not contribute to energy.

3. Numerical method and its properties

We utilize a semi-implicit time-marching scheme for Navier Stokes (eq. (3)) with projection to decouple the pressure computation. Similar to Khanwale et al. [19] we use a fully implicit time marching scheme for the Cahn-Hilliard equations eqs. (7) to (8). This strategy delivers second order accuracy and energy-stability, which we verify using numerical experiments, for larger time-steps that are comparable to the fully implicit scheme in Khanwale et al. [19].

⁶ See Remark 1 in Khanwale et al. [19].

Let δt be a time-step and let $t^k := k\delta t$. We define the following time-averages:

$$\begin{aligned}\bar{\mathbf{v}}^k &:= \frac{\mathbf{u}^k + \mathbf{v}^{k+1}}{2}, & \bar{\mathbf{u}}^k &:= \frac{\mathbf{u}^k + \mathbf{u}^{k+1}}{2}, & \bar{\mathbf{v}}^k &:= \frac{3\mathbf{u}^k + \mathbf{u}^{k-1}}{2}, & \bar{p}^k &:= \frac{p^{k+1} + p^k}{2}, \\ \bar{\phi}^k &:= \frac{\phi^{k+1} + \phi^k}{2}, & \text{and } \bar{\mu}^k &:= \frac{\mu^{k+1} + \mu^k}{2},\end{aligned}\quad (13)$$

as well as the following function evaluations:

$$\begin{aligned}\tilde{\psi}^k &:= \psi(\bar{\phi}^k), & \tilde{\psi}'^k &:= \psi'(\bar{\phi}^k), & \tilde{\rho}^k &:= \rho(\bar{\phi}^k), & \rho^{k+1} &:= \rho(\phi^{k+1}), \\ \rho^k &:= \rho(\phi^k), & \text{and } \tilde{\eta}^k &:= \eta(\bar{\phi}^k).\end{aligned}\quad (14)$$

Using these temporal values, we define our time-discretized weak form of the Cahn-Hilliard Navier-Stokes (CNHS) equations.

Definition 1 (Variational form). Let $(\cdot, \cdot)$ be the standard L^2 inner product. We state the time-discrete variational problem as follows: find $\mathbf{v}^{k+1}(\mathbf{x}), \mathbf{u}^{k+1}(\mathbf{x}) \in \mathbf{H}_0^1(\Omega)$, $p^{k+1}(\mathbf{x}), \phi^{k+1}(\mathbf{x}), \mu^{k+1}(\mathbf{x}) \in H^1(\Omega)$ such that

$$\text{Cahn-Hilliard Eqn: } \left(q, \frac{\phi^{k+1} - \phi^k}{\delta t} \right) - \left(\frac{\partial q}{\partial x_i}, \tilde{u}_i^k \bar{\phi}^k \right) + \frac{1}{PeCn} \left(\frac{\partial q}{\partial x_i}, \frac{\partial \tilde{\mu}^k}{\partial x_i} \right) = 0, \quad (15)$$

$$\text{Chemical Potential: } - \left(q, \tilde{\mu}^k \right) + \left(q, \tilde{\psi}'^k \right) + Cn^2 \left(\frac{\partial q}{\partial x_i}, \frac{\partial \bar{\phi}^k}{\partial x_i} \right) = 0, \quad (16)$$

$$\begin{aligned}\text{Velocity prediction: } & \left(w_i, \tilde{\rho}^k \frac{v_i^{k+1} - u_i^k}{\delta t} \right) + \left(w_i, \tilde{\rho}^k \tilde{u}_j^k \frac{\partial \tilde{v}_i^k}{\partial x_j} \right) + \frac{1}{Pe} \left(w_i, \tilde{J}_j^k \frac{\partial \tilde{v}_i^k}{\partial x_j} \right) \\ & - \frac{Cn}{We} \left(\frac{\partial w_i}{\partial x_j}, \frac{\partial \bar{\phi}^k}{\partial x_i} \frac{\partial \bar{\phi}^k}{\partial x_j} \right) + \frac{1}{We} \left(w_i, \frac{\partial p^k}{\partial x_i} \right) \\ & + \frac{1}{Re} \left(\frac{\partial w_i}{\partial x_j}, \tilde{\eta}^k \frac{\partial \tilde{v}_i^k}{\partial x_j} \right) - \left(w_i, \frac{\tilde{\rho}^k \hat{g}_i}{Fr} \right) = 0,\end{aligned}\quad (17)$$

$$\text{Thermo Consistency: } \tilde{J}_i^k = \frac{(\rho_- - \rho_+)}{2\rho_+Cn} \frac{\partial \tilde{\mu}^k}{\partial x_i}, \quad (18)$$

$$\text{Projection: } \left(w_i, \tilde{\rho}^k \frac{u_i^{k+1} - v_i^{k+1}}{\delta t} \right) + \frac{1}{2We} \left(w_i, \frac{\partial (p^{k+1} - p^k)}{\partial x_i} \right) = 0, \quad (19)$$

$$\text{Solenoidality: } \left(q, \frac{\partial u_i^{k+1}}{\partial x_i} \right) = 0, \quad \left(q, \frac{\partial u_i^k}{\partial x_i} \right) = 0, \quad \left(q, \frac{\partial u_i^{k-1}}{\partial x_i} \right) = 0, \quad (20)$$

$$\text{Continuity: } \left(\frac{\rho^{k+1} - \rho^k}{\delta t}, q \right) + \left(\frac{\partial (\tilde{\rho}^k \tilde{u}_j^k)}{\partial x_j}, q \right) - \left(\frac{1}{Pe} \tilde{J}_j^k, \frac{\partial q}{\partial x_j} \right) = 0, \quad (21)$$

$\forall \mathbf{w} \in \mathbf{H}_0^1(\Omega)$, $\forall q \in H^1(\Omega)$, given $\mathbf{u}^k, \mathbf{u}^{k-1} \in \mathbf{H}_0^1(\Omega)$, and $\phi^k, \mu^k \in H^1(\Omega)$.

We redefine the pressure by absorbing the coefficient $1/We$ in its definition. Using the solenoidality of $\mathbf{u}^{k+1}$, eq. (19) is implemented in a two step strategy as follows.

Definition 2. Let $(\cdot, \cdot)$ be the standard L^2 inner product. The time-discretized projection variational problem can be stated as follows: find $\mathbf{u}^{k+1}(\mathbf{x}) \in \mathbf{H}_0^1(\Omega)$, $p^{k+1}(\mathbf{x}) \in H^1(\Omega)$ such that

$$\text{Pressure Poisson: } \left(\frac{\partial w_i}{\partial x_i}, \left(\frac{1}{\tilde{\rho}^k} \frac{\partial p^{k+1}}{\partial x_i} \right) \right) = - \frac{2}{\delta t} \left(w_i, \frac{\partial v_i^{k+1}}{\partial x_i} \right) + \left(\frac{\partial w_i}{\partial x_i}, \left(\frac{1}{\tilde{\rho}^k} \frac{\partial p^k}{\partial x_i} \right) \right), \quad (22)$$

$$\text{Velocity update: } \left(w_i, \tilde{\rho}^k u_i^{k+1} \right) + \frac{\delta t}{2} \left(w_i, \frac{\partial p^{k+1}}{\partial x_i} \right) = \left(w_i, \tilde{\rho}^k v_i^{k+1} \right) + \frac{\delta t}{2} \left(w_i, \frac{\partial p^k}{\partial x_i} \right), \quad (23)$$

$\forall \mathbf{w} \in \mathbf{H}_0^1(\Omega)$, given $\mathbf{v}^{n+1} \in \mathbf{H}_0^1(\Omega)$, and $p^k \in H^1(\Omega)$.

Here the pressure Poisson equation is derived by dividing eq. (19) by $\tilde{\rho}^k$ and then taking the divergence. Once the current pressure is known, we can update the velocity using eq. (23).

Remark 1. In Definition 1 and Definition 2 for the momentum equations and projection equations, the boundary terms in the variational form are zero because the velocity and the basis functions live in $\mathbf{H}_0^1(\Omega)$. Also we use the no flux boundary condition for ϕ and μ , which makes the boundary terms zero (i.e., $(q, \frac{\partial \tilde{\phi}^k}{\partial x_i} \hat{n}_i) = 0$ and $(q, \frac{\partial \mu^k}{\partial x_i} \hat{n}_i) = 0$). This choice of boundary conditions ensures that there are no boundary terms when the equations are weakened in the variational form. Other boundary conditions suitable for energy stability analysis using this particular energy functional include closed systems with neutral contact line dynamics. Additionally, a slight generalization of these boundary conditions to periodic conditions is also suitable for an energy-stability analysis as there are no boundary terms for this case. The numerical examples in the paper use these boundary conditions unless explicitly noted otherwise.

Remark 2. While $\phi \in [-1, 1]$ in the continuous equations, the discrete ϕ may violate these bounds. These bound violations may not change the dynamics of ϕ adversely, but they could lose the strict positivity of some quantities which depend on ϕ (e.g., mixture density $\rho(\phi)$ and viscosity $\eta(\phi)$). This effect is especially significant for high density and viscosity contrasts. We fix this issue by saturation scaling (i.e., we pull back the value of ϕ only for the calculation of density and viscosity). In particular, we define the following ϕ^* for the mixture density and viscosity calculations:

$$\phi^* := \begin{cases} \phi, & \text{if } |\phi| \leq 1, \\ \text{sign}(\phi), & \text{otherwise.} \end{cases} \quad (24)$$

Remark 3. It is important to note here that we are using a block iteration technique. Therefore, ϕ and μ are known when we are solving the momentum equations, and v_i is known when we are solving the advective Cahn-Hilliard equation. Also note that the mixture velocity used for advecting the phase field in eq. (15) is $\tilde{u}_i$, which is weakly solenoidal instead of $\tilde{v}_i$ which is not. We design the solve strategy such that this is always satisfied.

We now state a couple of important properties of the scheme.

Proposition 1 (Mass conservation). *The scheme of eq. (15) – eq. (16) with the following boundary conditions:*

$$\frac{\partial \tilde{\mu}}{\partial x_i} \hat{n}_i \Big|_{\partial \Omega} = 0, \quad \frac{\partial \tilde{\phi}}{\partial x_i} \hat{n}_i \Big|_{\partial \Omega} = 0, \quad \tilde{\mathbf{u}}^k \Big|_{\partial \Omega} = \mathbf{0}, \quad (25)$$

where $\hat{\mathbf{n}}$ is the outward normal to the boundary $\partial \Omega$, is globally mass conservative:

$$\int_{\Omega} \phi^{k+1} d\mathbf{x} = \int_{\Omega} \phi^k d\mathbf{x}. \quad (26)$$

Proof. The proposition can be easily proved using the test function in eq. (15) as 1, and subsequently using the divergence theorem along with the given boundary conditions prescribed by the function spaces. ■

We verify this claim numerically in sections 5.1 and 5.2. We recall the following result from Khanwale et al. [48].

Lemma 1 (Weak equivalence of forcing). *The forcing term due to Cahn-Hilliard in the momentum equation, eq. (17), with the test function $w_i = \delta t \tilde{v}_i^k$, becomes*

$$\frac{Cn}{We} \left(\frac{\partial}{\partial x_j} \left(\frac{\partial \tilde{\phi}^k}{\partial x_i} \frac{\partial \tilde{\phi}^k}{\partial x_j} \right), \delta t \tilde{v}_i^k \right) = \frac{\delta t}{We Cn} \left(\tilde{\phi}^k \frac{\partial \tilde{\mu}^k}{\partial x_i}, \tilde{v}_i^k \right) + \frac{\delta t Cn}{We} \left(\frac{\tilde{\mu}^k \tilde{\phi}^k}{Cn^2} - \frac{1}{2} \frac{\partial \tilde{\phi}^k}{\partial x_j} \frac{\partial \tilde{\phi}^k}{\partial x_j} - \frac{\tilde{\psi}}{Cn^2}, \frac{\partial \tilde{v}_i^k}{\partial x_i} \right), \quad (27)$$

$\forall \tilde{\phi}^k, \tilde{\mu}^k \in H^1(\Omega)$, and $\forall \tilde{\mathbf{v}}^k \in \mathbf{H}_0^1(\Omega)$, where $\mathbf{v}^k, \mathbf{v}^{k+1}, p^k, p^{k+1}, \phi^k, \phi^{k+1}, \mu^k, \mu^{k+1}$, satisfy eq. (17) – eq. (15).

Remark 4. In the literature, some variant of the forcing term $\tilde{\phi}^k \frac{\partial \tilde{\mu}^k}{\partial x_i}$, which is the first term on the right hand side of the expression above, is used effectively to construct semi-implicit schemes [17,24,28–34]. These forcing terms are equivalent to the thermodynamically consistent forcing term, $\frac{\partial}{\partial x_j} \left(\frac{\partial \tilde{\phi}^k}{\partial x_i} \frac{\partial \tilde{\phi}^k}{\partial x_j} \right)$, **only** if the test function is divergence free. The usual test functions used for a projection method with conforming Galerkin method are not divergence free. Thus, using a forcing term $\tilde{\phi}^k \frac{\partial \tilde{\mu}^k}{\partial x_i}$ results in methods that do not exhibit appropriate temporal convergence as conservation errors saturate. We show this effect numerically in panel (b) of Fig. 3 in section 5.1.




Fig. 1. Flowchart for the block iteration technique as described in section 3.

We solve the spatially discretized version of variational problems in Definition 1 and Definition 2 using a block iteration technique, i.e., we treat the velocity prediction equations (eq. (17)), pressure Poisson equation (eq. (22)), velocity update equations (eq. (23)), and the Cahn-Hilliard equations (eqs. (15) to (16)) as distinct sub-problems. Thus, three linear solvers of (1) velocity prediction, (2) pressure Poisson, and (3) velocity update are stacked together with a non-linear solver for Cahn-Hilliard equations inside the time loop. These solvers are each solved twice (two blocks) within every time step to preserve order of accuracy and self-consistency. See Fig. 1 for a flowchart of the approach. We emphasize that a block iterative approach allows us to make the coupling variables from one equation constant in the other during each respective linear/non-linear solve. For example, for the momentum equation, all the terms depending on ϕ (which is solved in the Cahn-Hilliard sub problem) are known. Similarly the mixture velocity used in the Cahn-Hilliard equation solve. We choose to perform two-block iterations as for the timesteps we choose for accuracy concerns, we achieve convergence in block iteration in two blocks.

We adopt a strategy here such that the mixture velocity used in the advection of the phase field (eq. (15)) is always weakly solenoidal. This means that the pressure Poisson and velocity update is performed two times as we run the block iteration twice. This is a robust strategy, especially for the cases of high density ratios where the pressure gradients between the two phases can be high. Another popular strategy, used for equal density ratios by Han and Wang [34], is to iterate between velocity prediction (eq. (17)) and Cahn-Hilliard equations (eqs. (15) to (16)) until a consistency tolerance is satisfied, and then project the velocity to calculate pressure and solenoidal velocity in the last block iteration. We do not use this strategy even though it is potentially cheaper (only requires pressure Poisson and velocity update to be performed once) due to the fact that we are not using divergence conforming elements unlike Han and Wang [34] and we simulate large density ratios.

3.1. Spatial discretization and the variational multiscale approach

We discretize the unknowns:

$$(\phi, \mu, \mathbf{v}, \mathbf{u}, p), \quad (28)$$

in space using conforming continuous Galerkin finite elements with piecewise polynomial approximations. However, approximating the velocity, $\mathbf{v}$, and the pressure, p , with the same polynomial order leads to numerical instabilities as this violates the discrete inf-sup condition or Ladyzhenskaya-Babuska-Brezzi condition (e.g., see Volker [16, page 31]). Even though the projection method decouples pressure and velocity, our use of equal interpolation order basis functions for pressure and velocity may lead to instabilities. In order to overcome this difficulty, we use the variational multi-scale (VMS) method [49], which adds stabilization terms to the pressure Poisson equation that transform the inf-sup stability condition to a coercivity statement [50]. Additionally, VMS provides a natural leeway into modeling high-Reynolds number flows [38] as it uses a projection based filter to decompose coarse and fine scales.

The VMS approach uses a direct-sum decomposition of the function spaces as follows. If $\mathbf{v} \in \mathbf{V}$, $p \in Q$, and $\phi \in Q$ then we decompose these spaces as:

$$\mathbf{V} = \mathbf{V}^c \oplus \mathbf{V}^f \quad \text{and} \quad Q = Q^c \oplus Q^f, \quad (29)$$

where $\mathbf{V}^c$ and Q^c are the finite dimensional cG(r) subspaces of $\mathbf{V}$ and Q , respectively, and the superscript f versions are the complements of the cG(r) subspaces in $\mathbf{V}$ and Q , respectively. We decompose the velocity as follows:

$$\mathbf{v} = \mathbf{v}^c + \mathbf{v}^f, \quad (30)$$

where the *coarse scale* components are $\mathbf{v}^c \in \mathbf{V}^c$, and the *fine scale* components are $\mathbf{v}^f \in \mathbf{V}^f$. We define a projection operator, $\mathcal{P}: \mathbf{V} \rightarrow \mathbf{V}^c$, such that $\mathbf{v}^c = \mathcal{P}\{\mathbf{v}\}$ and $\mathbf{v}^f = \mathbf{v} - \mathcal{P}\{\mathbf{v}\}$. Let

$$\tilde{v}_i^{c,k} = \frac{v_i^{c,k+1} + u_i^{c,k}}{2}, \quad \tilde{v}_i^{f,k} = \frac{v_i^{f,k+1} + u_i^{f,k}}{2}, \quad \text{and} \quad \tilde{v}_i^{c,k} + \tilde{v}_i^{f,k} = \tilde{v}_i^k. \quad (31)$$

Substituting this decomposition in the original variational form Definition 1 yields:

$$\begin{aligned} \text{Velocity prediction:} \quad & \left(w_i, \tilde{\rho}^k \frac{(v_i^{c,k+1} + v_i^{f,k+1}) - (u_i^{c,k} + u_i^{f,k})}{\delta t} \right) + \left(w_i, \frac{(\rho^{k+1} - \rho^k)}{\delta t} \tilde{v}_i^{f,k} \right) \\ & + \left(w_i, \tilde{\rho}^k \hat{u}_j^k \frac{\partial \tilde{v}_i^{c,k}}{\partial x_j} \right) + \left(w_i, \tilde{\rho}^k \hat{u}_j^k \frac{\partial \tilde{v}_i^{f,k}}{\partial x_j} \right) + \left(w_i, \frac{\partial (\tilde{\rho}^k \hat{u}_j^k)}{\partial x_j} \tilde{v}_i^{f,k} \right) \\ & + \frac{1}{Pe} \left(w_i, \hat{J}_j^k \frac{\partial \tilde{v}_i^{c,k}}{\partial x_j} \right) + \frac{1}{Pe} \left(w_i, \hat{J}_j^k \frac{\partial \tilde{v}_i^{f,k}}{\partial x_j} \right) + \frac{1}{Pe} \left(w_i, \frac{\partial \hat{J}_j^k}{\partial x_j} \tilde{v}_i^{f,k} \right) \\ & - \frac{Cn}{We} \left(\frac{\partial w_i}{\partial x_j}, \left(\frac{\partial \tilde{\phi}^{h,k}}{\partial x_i} \frac{\partial \tilde{\phi}^{h,k}}{\partial x_j} \right) \right) + \frac{1}{We} \left(w_i, \frac{\partial p^k}{\partial x_i} \right) \\ & + \frac{1}{Re} \left(\frac{\partial w_i}{\partial x_j}, \tilde{\eta}^k \frac{\partial (\tilde{v}_i^{c,k} + \tilde{v}_i^{f,k})}{\partial x_j} \right) - \left(w_i, \frac{\tilde{\rho}^k \hat{g}_i}{Fr} \right) = 0, \end{aligned} \quad (32)$$

$$\text{Pressure Poisson:} \quad \left(\frac{\partial w_i}{\partial x_i}, \left(\frac{1}{\tilde{\rho}^k} \frac{\partial p^{k+1}}{\partial x_i} \right) \right) = -\frac{2}{\delta t} \left(w_i, \frac{\partial (v_i^{c,k+1} + v_i^{f,k+1})}{\partial x_i} \right) + \left(\frac{\partial w_i}{\partial x_i}, \left(\frac{1}{\tilde{\rho}^k} \frac{\partial p^k}{\partial x_i} \right) \right), \quad (33)$$

$$\text{Velocity Update:} \quad \left(w_i, \tilde{\rho}^k u_i^{k+1} \right) + \frac{\delta t}{2} \left(w_i, \frac{\partial p^{k+1}}{\partial x_i} \right) = \left(w_i, \tilde{\rho}^k (v_i^{c,k+1} + v_i^{f,k+1}) \right) + \frac{\delta t}{2} \left(w_i, \frac{\partial p^k}{\partial x_i} \right), \quad (34)$$

where $\mathbf{w}, \mathbf{v}^{c,k+1} \in \mathcal{P}\mathbf{H}^r(\Omega)$, $\mathbf{v}^{f,k+1} \in (\mathcal{I} - \mathcal{P})\mathbf{H}^r(\Omega)$. Here $\mathcal{I}$ is the identity operator and $\mathcal{P}$ is the projection operator. Note that we add the following equation to the original momentum equation (eq. (17)) using mass conservation eq. (21):

$$\left(w_i, \frac{(\rho^{k+1} - \rho^k)}{\delta t} \tilde{v}_i^{f,k} \right) + \left(w_i, \frac{\partial (\tilde{\rho}^k \hat{u}_j^k)}{\partial x_j} \tilde{v}_i^{f,k} \right) + \frac{1}{Pe} \left(w_i, \frac{\partial \hat{J}_j^k}{\partial x_j} \tilde{v}_i^{f,k} \right) = 0. \quad (35)$$

This is important because the terms in blue in eq. (32) are not in conservative form, and would require higher regularity of $\tilde{v}_i^{f,k}$ to incorporate. (For interpretation of the references to color please refer to the web version of this article.) However, using eq. (35) along with integration-by-parts the terms in blue can be converted into conservative terms. Consequently, we now have extra terms:

$$\left(w_i, \tilde{\rho}^k \frac{(v_i^{f,k+1} - u_i^{f,k})}{\delta t} \right) + \left(w_i, \frac{(\rho^{k+1} - \rho^k)}{\delta t} \tilde{v}_i^{f,k} \right). \quad (36)$$

We can think of this addition as a second order discretization of $(w_i, \partial(\rho v_i^f)/\partial t)$. We assume, following [38], that the time derivative of fine scale momentum is zero.

We use the residual-based approximation proposed in Bazilevs et al. [38] for the fine-scale components, applied to a two-phase system [48], to close the equations:

$$\tilde{\rho}^k v_i^{f,k+1} = -\tau_m \mathcal{R}_m(\tilde{\rho}^k, \mathbf{v}^{k+1}, \mathbf{u}^k, \mathbf{u}^{k-1}, p^k) \quad (37)$$

with the following parameter values:

$$\tau_m = \left(\frac{4}{\Delta t^2} + v_i^c G_{ij} v_j^c + \frac{1}{\tilde{\rho}^k Pe} v_i^c G_{ij} J_j^c + C_l \left(\frac{\tilde{\eta}^k}{\tilde{\rho}^k Re} \right)^2 G_{ij} G_{ij} \right)^{-1/2}. \quad (38)$$

Here we set C_l and C_ϕ for all our simulations to 6 and the residuals are given by

$$\begin{aligned} \mathcal{R}_m(\tilde{\rho}^k, v_i^{c,k+1}, u_i^k, u_i^{k-1}, p^k) &= \tilde{\rho}^k \frac{v_i^{c,k+1} - u_i^k}{\delta t} + \tilde{\rho}^k \hat{u}_j^k \frac{\partial \tilde{v}_i^{c,k}}{\partial x_j} + \frac{1}{Pe} \hat{J}_j \frac{\partial \tilde{v}_i^{c,k}}{\partial x_j} \\ &\quad + \frac{Cn}{We} \frac{\partial}{\partial x_j} \left(\frac{\partial \phi^{h,k}}{\partial x_i} \frac{\partial \phi^{h,k}}{\partial x_j} \right) \\ &\quad + \frac{1}{We} \frac{\partial p^k}{\partial x_i} - \frac{1}{Re} \frac{\partial}{\partial x_j} \left(\tilde{\eta}^k \frac{\partial \tilde{v}_i^{c,k}}{\partial x_j} \right) - \frac{\tilde{\rho}^k \hat{g}}{Fr}. \end{aligned} \quad (39)$$

In the above expressions, we used the following notation:

$$\tilde{\rho}^k := \tilde{\rho}^k := \rho(\tilde{\phi}^{h,k}) \quad \text{and} \quad \eta^h := \eta(\tilde{\phi}^{h,k}). \quad (40)$$

Now, we replace the infinite-dimensional function spaces by their finite dimensional counterparts using conforming Galerkin finite elements where the trial and test functions are taken from the same spaces. Note that we only solve for the coarse-scale components. For notational simplicity we do not add the conventional superscript h denoting finite dimensional conforming Galerkin approximations. All the functions spaces in the definition below are finite dimensional. The resulting discrete variational formulation can then be defined as follows.

Definition 3 (VMS setting for velocity). Find $\mathbf{v}^{c,k+1} \in \mathcal{P}\mathbf{H}_0^r(\Omega)$ such that

$$\begin{aligned} \text{Velocity prediction:} \quad & \left(w_i, \tilde{\rho}^k \frac{v_i^{c,k+1} - u_i^k}{\delta t} \right) + \frac{1}{2} \left(w_i, \tilde{\rho}^k \hat{u}_j^k \frac{\partial v_i^{c,k+1}}{\partial x_j} \right) + \frac{1}{2} \left(w_i, \tilde{\rho}^k \hat{u}_j^k \frac{\partial u_i^k}{\partial x_j} \right) \\ & + \frac{1}{2} \left(\frac{\partial w_i}{\partial x_j}, \hat{u}_j^k \tau_m \mathcal{R}_m(\tilde{\rho}^k, v_i^{c,k+1}, u_i^k, u_i^{k-1}, p^k) \right) \\ & + \frac{1}{2Pe} \left(w_i, \hat{J}_j^k \frac{\partial v_i^{c,k+1}}{\partial x_j} \right) + \frac{1}{2Pe} \left(w_i, \hat{J}_j^k \frac{\partial u_i^k}{\partial x_j} \right) \\ & + \frac{1}{2Pe} \left(\frac{\partial w_i}{\partial x_j}, \hat{J}_j^k \frac{1}{\tilde{\rho}^k} \tau_m \mathcal{R}_m(\tilde{\rho}^k, v_i^{k+1}, u_i^k, u_i^{k-1}, p^k) \right) \\ & - \frac{Cn}{We} \left(\frac{\partial w_i}{\partial x_j}, \left(\frac{\partial \tilde{\phi}^{h,k}}{\partial x_i} \frac{\partial \tilde{\phi}^{h,k}}{\partial x_j} \right) \right) + \frac{1}{We} \left(w_i, \frac{\partial p^k}{\partial x_i} \right) \\ & + \frac{1}{Re} \left(\frac{\partial w_i}{\partial x_j}, \tilde{\eta}^k \frac{\partial (\tilde{v}_i^{c,k} + \tilde{v}_i^{f,k})}{\partial x_j} \right) - \left(w_i, \frac{\tilde{\rho}^k \hat{g}_i}{Fr} \right) = 0, \end{aligned} \quad (41)$$

$\forall \mathbf{w} \in \mathcal{P}\mathbf{H}_0^{r,h}(\Omega)$, given $\mathbf{u}^k, \mathbf{u}^{k-1} \in \mathcal{P}\mathbf{H}_0^{r,h}(\Omega)$ and $p^k, \phi^h, \mu^h \in \mathcal{P}H^{r,h}(\Omega)$.

Remark 5. Note that $\hat{u}_i^k$ is not decomposed here. The velocities at older times have their fine scale terms added to them (see the boxed term in eq. (43)) in the velocity update step in the previous time step and are completely known. Therefore, for ease of implementation in the code we do not need to decompose them. When we convert the fine scale velocities in the convective form again to get a complete velocity at time k (see term 3 and term 5), we get extra fine scale terms $-\left(w_i, \tilde{\rho}^k \left(u_i^{f,k}/dt\right)\right) + 1/2 \left(w_i, ((\rho^{k+1} - \rho^k)/\delta t) u_i^{f,k}\right)$. Our experimental results suggest that these terms are very small, and consequently we set them to zero. We emphasize here that eq. (41) is a linear equation. We deploy a semi-implicit time marching scheme that turns non-linear terms into a linear algebraic equation that must be solved at each time step. The viscous terms are considered to be implicit. This type of discretization improves efficiency as a full Newton iteration need not be solved while maintaining stability and robustness as viscous terms are still implicit. Note that this type of discretization is different from IMEX schemes [51], where the non-linear terms are completely explicit.

Now, let us define the variational VMS setting for the projection and update steps.

Definition 4 (VMS setting for pressure Poisson). Find $p^{k+1} \in H^r(\Omega)$ such that

$$\text{Pressure Poisson: } \left(\frac{\partial w_i}{\partial x_i}, \left(\frac{1}{\tilde{\rho}^k} \frac{\partial p^{k+1}}{\partial x_i} \right) \right) = -\frac{2}{\delta t} \left(w_i, \frac{\partial v_i^{c,k+1}}{\partial x_i} \right) - \frac{2}{\delta t} \left(\frac{\partial w_i}{\partial x_i}, \frac{1}{\tilde{\rho}^k} \tau_m \mathcal{R}_m \left(\tilde{\rho}^k, v_i^{c,k+1}, u_i^k, u_i^{k-1}, p^k \right) \right) + \left(\frac{\partial w_i}{\partial x_i}, \left(\frac{1}{\tilde{\rho}^k} \frac{\partial p^k}{\partial x_i} \right) \right), \quad (42)$$

$$\forall \mathbf{w} \in \mathcal{P}\mathbf{H}_0^{r,h}(\Omega) \text{ and } \forall q \in \mathcal{P}H^{r,h}(\Omega).$$

Remark 6. The boxed term in Definition 4 is very important for stability. If that term is not included we get checkerboard instabilities in pressure as we are using equal order interpolating basis functions for both velocity and pressure. The boxed term is a manifestation of the splitting of v_i^{k+1} into its coarse scale ($v_i^{c,k+1}$) and fine scale components ($v_i^{f,k+1}$) combined with the use of residual based VMS approximation in eq. (37).

As the pressure is calculated now, we can write the variational setting for the velocity update which transforms the non-solenoidal velocity $\mathbf{v}^{c,k+1}$ to a solenoidal velocity u^{k+1} at the current timestep.

Definition 5 (VMS setting for velocity update). Find $p^{k+1} \in H^r(\Omega)$ such that

$$\text{Velocity update: } \left(w_i, \tilde{\rho}^k u_i^{k+1} \right) + \frac{\delta t}{2} \left(w_i, \frac{\partial p^{k+1}}{\partial x_i} \right) = \left(w_i, \tilde{\rho}^k v_i^{c,k+1} \right) - \left(w_i, \tau_m \mathcal{R}_m \left(\tilde{\rho}^k, v_i^{c,k+1}, u_i^k, u_i^{k-1}, p^k \right) \right) + \frac{\delta t}{2} \left(w_i, \frac{\partial p^k}{\partial x_i} \right), \quad (43)$$

$$\forall \mathbf{w} \in \mathcal{P}\mathbf{H}_0^{r,h}(\Omega) \text{ given } \mathbf{v}^{c,k+1} \in \mathcal{P}\mathbf{H}_0^{r,h}(\Omega), \text{ and } p^{k+1}, p^k \in \mathcal{P}H^{r,h}(\Omega),$$

Here we have again used the splitting of v_i^{k+1} into its coarse scale ($v_i^{c,k+1}$) and fine scale components ($v_i^{f,k+1}$) combined with the use of residual based VMS approximation in eq. (37). We can now write the fully discrete Cahn-Hilliard variational form as,

Definition 6 (discrete Cahn-Hilliard equations). Find $\phi^h, \mu^h \in H^{r,h}(\Omega)$ such that

$$\text{Cahn-Hilliard: } \left(q, \frac{\phi^{h,k+1} - \phi^{h,k}}{\delta t} \right) - \left(\frac{\partial q}{\partial x_i}, \tilde{u}_i^{h,k} \tilde{\phi}^{h,k} \right) - \frac{1}{PeCn} \left(\frac{\partial q}{\partial x_i}, \frac{\partial (m^h \tilde{\mu}^{h,k})}{\partial x_i} \right) = 0, \quad (44)$$

$$\text{Potential: } - \left(q, \tilde{\mu}^{h,k} \right) + \left(q, \psi'(\tilde{\phi}^{h,k}) \right) + Cn^2 \left(\frac{\partial q}{\partial x_i}, \frac{\partial \tilde{\phi}^{h,k}}{\partial x_i} \right) = 0, \quad (45)$$

$$\forall \tilde{\mathbf{u}}^k \in \mathbf{H}_0^{r,h}(\Omega) \text{ and } \forall q \in \mathcal{P}H^{r,h}(\Omega).$$

In addition to assumptions made in eq. (36) and Remark 5, we reiterate the other assumptions made in getting these variational problems.

1. $\mathbf{v}^f = 0$ on the boundary $\partial\Omega$; similarly $\phi' = 0$ on $\partial\Omega$.
2. $\left(\frac{\partial w_i}{\partial x_j}, \tilde{\eta}^k \frac{\partial \tilde{v}_i^{f,k}}{\partial x_j}\right) = 0$ from the orthogonality condition of the projector. The projector utilizes the inner product that comes from the bilinear form of the viscous terms [38,52].
3. We use the pulled back ϕ^* (see Remark 2) for this calculation, which regularizes ϕ by smoothing out overshoots and undershoots. The pull back ensures $\phi^* \in H^r$, and that is projection on the mesh $\phi^{*,h} \in H^{r,h}$ as required for the cG formulation.

Remark 7. The above formulation is written for a generic order (r) for the interpolating polynomials (basis functions). Nevertheless, we restrict our attention to $r = 1$ for the numerical experiments in this paper.

Remark 8. Here we use an incremental pressure projection scheme. We show with detailed numerical experiments this scheme produces good results. See Guermond et al. [35] for an excellent review of pressure projection schemes in the weak setting, including some schemes which use rotational identities to achieve improvement. We note that the design and efficient HPC implementation of schemes incorporating rotational identities on octree based meshes is not straightforward. This is an active area of research in our groups.

4. Octree based domain decomposition

We use DENDRO-KT, a highly scalable parallel octree library, to generate, manage and traverse adaptive octree-based meshes in distributed memory. In this section, we summarize the core data-structures and algorithms used by DENDRO-KT. Additional details on these algorithms can be found in [20]. The DENDRO-KT interface is the next version of the DENDRO5 interface used in Khanwale et al. [19]. We present the improvements and critical changes in the octree framework compared to DENDRO5.

4.1. Distributed memory data structures

One of the distinguishing aspects of DENDRO-KT compared with other adaptive mesh refinement (AMR) frameworks is that it does not store any element-to-node maps. The overall memory footprint of the adaptive mesh is kept to a minimum, storing only the coordinates of the nodes. All other information is deduced on-the-fly during traversals of the kD -tree. Besides reducing the memory-footprint, this also avoids indirect memory accesses (via element-to-node maps) that are extremely inefficient on modern HPC architectures. In the distributed memory setting, this choice has implications for the inter-process and intra-process data structures.

A minimal data structure is maintained to send *owned* nodal data to kD neighboring processes: a list of the MPI processes owning neighboring partitions, and the number of, and list of local indices of, owned nodes to replicate to each neighboring process. We call this data structure a scatter-map, and is similar to the information that is typically used to exchange data to create *ghost* or *halo* regions. When ghost data is received from remote processes, it is aligned with a list of remotely-owned nodal coordinates, allowing ghost data to be intermingled, sorted, and traversed alongside locally-owned data without special handling.

Once the tree is constructed (described in the next section) all tree-related information is discarded, except the coordinates of the leaf octants and coordinates of the nodal points. Direct neighborhood access is not possible because of the lack of neighborhood maps, instead efficient traversal methods are provided (section 4.4) to enable FEM computations. Because of the lack of traditional maps used by unstructured mesh libraries, we refer to this approach as a *mesh-free* approach.

4.2. Octree construction and 2:1 balancing

DENDRO-KT refines an octant based on user-specified criteria proceeding in a top-down fashion. The user defines the refinement criteria by a function that takes the coordinates of the octant, and returns *true* or *false*. Since the refinement happens locally to the element, this step is embarrassingly parallel. In distributed-memory machines, the initial top-down tree construction enables an efficient partitioning of the domain across an arbitrary number of processes. All processes start at the root node (i.e., the cubic bounding box for the entire domain). We perform redundant computations on all processes to avoid communication during the refinement stage. Starting from the root node, all processes refine (similar to a sequential implementation) until the process produces at least $\mathcal{O}(p)$ octants requiring further refinement. The procedure ensures that upon partitioning across p processors, each processor gets at least one octant. Then using a space-filling-curve (SFC) based partitioning, we partition the octants across p partitions [53]. Once the algorithm completes this partitioning, we can restrict the refinement criterion to a processor's partition, which we can re-distribute to ensure load-balancing, with P , where $p < P$ partitions, according to the adaptive grid generated. We enforce a condition in our distributed octrees that no

two neighboring octants differ in size by more than a factor of two (2:1 balancing). This ratio makes subsequent operations simpler without affecting the adaptive properties. Our balancing algorithm uses a variant of TREESORT [53] with top-down and bottom-up traversal of octrees which is different from existing approaches [46,54,55]. As noted above, only the leaf octant coordinates and node coordinates (that lie on the constructed octree) are retained. Nodes that do not represent independent degrees of freedom, such as the so-called *hanging nodes* are also discarded.

4.3. Identification of unique, independent nodes

Before distributed nodal vectors can be defined and traversed, the nodes incident on the leaf octants must be enumerated and partitioned. We refer to the result of enumerating and partitioning the nodes as a *DA*, for *distributed array*, as it contains the coordinates of local and remotely accessed nodes, as well as the process-level scattermap needed for ghost exchange. It is straightforward to emit the coordinates of the nodes appearing at each element separately. However, most of these nodes are shared with neighboring elements, which may reside in other processes, and some of the nodes are dependent, or *hanging*, on the face of a larger neighboring element. The situation is additionally complicated by the possibility of a carved-out domain, which increases the number of ways that a node can be shared by several neighboring elements.

In this work the DA construction can be outlined as follows. Given the input of a distributed list of leaf octants, two intermediate graphs are produced, where edges represent dependencies from elements to nodes. The second graph is used to derive node ownership and the scattermap.

In more detail: For each leaf octant in the local mesh partition, the coordinates of candidate nodes are emitted. Anticipating the coordinates of potential hanging nodes on the element exterior, a set of *cancellation nodes* is also emitted, at one level deeper in the octree. Flags on the nodal coordinates indicate whether they are regular or cancellation nodes. Also, each node is emitted in a pairing with the emitting element, representing a dependency from the element to the node.

At this point there exists a distributed list of edges, ordered by element coordinates in the SFC. The next step is to identify hanging nodes using the cancellation nodes. The distributed list of edges is re-sorted and -partitioned, using nodal coordinate as the DISTREESORT key instead of the element. Thus all edges to a given node are stored contiguously.

- If a regular node is found at the same coordinate as a cancellation node, then it is hanging. There is a bijection between the nodes of a child face that are not shared by its parent, and the nodes of the parent face that are not shared with the child. Using this bijection, hanging nodes are mapped to non-hanging nodes, and the resulting edges are re-emitted. Effectively, the original (element, node) dependencies are corrected so that a hanging element is able to request the nodes needed for interpolation.
- If a regular node is found with no cancellation node at the same coordinate, the corresponding edge is emitted as-is.
- Edges of cancellation nodes are discarded.

After the nodal keys are updated, the second graph is realized by resorting the edges by the new nodal coordinates.

On each local partition of the graph, the set of elements is TREESORT-ed against the globally-known splitters of the mesh partition, producing a map from edges to process rank. Effectively the graph has been aggregated into a set of (process, node) dependencies. For each nodal coordinate, one of the connected process ranks is chosen to be the *owner* of that node. The other connected process ranks become *borrowers* of the node. The nodes are sent to the owners with a list of the borrowers, and to each of the borrowers with a tag specifying the owner. After this exchange, each process has received the coordinates of its owned nodes and remotely accessed nodes, along with the information to be stored in the scattermap.

4.4. Traversal-based mesh-free operations

DENDRO-KT supports both matrix and matrix-free computations. Since, we do not store any map data-structure, such an operation is achieved by performing top-down and bottom-up traversals of the tree. We briefly describe the key ideas of the approach here for performing matrix and vector assembly, as well as the inter-grid transfers needed AMR. Additional details on mesh-free traversals can be found in Ishii et al. [20].

Traversal-based vector assembly (see Fig. 2) The elemental vector couples elemental nodes in a global input grid vector with equivalent elemental nodes in a global output grid vector. Within a grid vector, the nodes pertaining to a particular element are generally not stored contiguously. If one were to read and write to the elemental nodes using an element-to-node map, the memory accesses would require indirection: $v_{glob}[map[e * npe + i]] = v_{loc}$. Not only do element-to-node maps cause indirect memory accesses; the maps become complicated to build if the octree is incomplete due to complex geometry. We take an alternative approach that obviates the need for element-to-node maps. Instead, through top-down and bottom-up traversals of the kD -tree, we ensure that elemental nodes are stored contiguously in a leaf, and there apply the elemental operations.

The idea of the top-down phase is to selectively copy nodes from coarser to finer levels until the leaf level, wherein the selected nodes are exactly the elemental nodes. Starting at the root of the tree, we have all the nodes in the grid vector. We create buckets for all child subtrees. Looping through the nodes, a node is copied into a bucket if the node is incident




Fig. 2. Illustration of top-down & bottom-up tree traversals for a 2D tree with quadratic element order. The leftmost figure depicts the unique shared nodes (nodes are color-coded based on level), as we perform top-down traversal nodes shared across children of the parent get duplicated for each bucket recursively, once leaf node is reached it might be missing elemental local nodes, which can be interpolated from immediate parent (see the rightmost figure). After elemental local node computations, bottom-up traversal performed while merging the nodes duplicated in the top-down traversal.

on the child subtree corresponding to that bucket. A node that is incident on multiple child subtrees will be duplicated. By recursing on each child subtree and its corresponding bucket of incident nodes, we eventually reach the leaf level.

Once the traversal reaches a leaf octant, the elemental nodes have been copied into a contiguous array. The elemental vector is computed directly, without the use of an element-to-node map. The result is stored in a contiguous output buffer the same size as the local elemental input vector.

After all child subtrees have been traversed, the bottom-up phase returns results from a finer to a coarser level. The parent subtree nodes are once again bucketed to child subtrees, but instead of the parent values being copied, the values of nodes from each child are accumulated into a parent output array. That is, for any node that is incident on multiple child subtrees, the values from all node instances are summed to a single value. The global vector is assembled after the bottom-up phase executes at the root of the octree.

Distributed memory is supported by two slight augmentations. Firstly, the top-down and bottom-up traversals operate on ghosted vectors. Therefore ghost exchanges are required before and after each local traversal. Secondly, the traversals are restricted to subtrees containing the owned octants. The list of owned octants is bucketed top-down, in conjunction with the bucketing of nodal points. A child subtree is traversed recursively only if one or more owned octants are bucketed to it. Note that because the traversal path is restricted by a list of existing octants, the traversal-based operation gracefully handles incomplete octrees without special treatment.

Traversal-based matrix assembly To implement matrix assembly, we have leveraged PETSc interface [56,57], that only requires a sequence of entries (id_{row} , id_{col} , val), and can be configured to add entries with duplicate indices [58]. Note that any other distributed sparse-matrix library can be supported in a similar fashion.

The remaining task is to associate the correct global node indices with the rows and columns of every elemental matrix. We use an octree traversal to accomplish this task. Similar to the traversal-based vector assembly, nodes are selectively copied from coarser to finer levels, recursively, until reaching the leaf, wherein the elemental nodes are contiguous. Note that integer node ids are copied instead of floating-point values from a grid vector. At the leaf, an entry of the matrix is emitted for every row and column of the elemental matrix, using the global row and column indices instead of the elemental ones. No bottom-up phase is required for assembly, as PETSc handles the merging of multi-instanced entries.

Traversal-based intergrid transfers An important aspect of AMR is the ability to effectively and efficiently transfer information from the current mesh to the newly refined/coarsened mesh—the intergrid transfer. DENDRO-KT supports efficient traversal-based intergrid transfers, and regions can be locally refined or coarsened in different regions of the mesh. For efficiency, we limit the local refinement or coarsening to only be by one level.⁷ In order to transfer information from mesh *A* (current) to mesh *B* (new), we simultaneously traverse both meshes in a synchronized fashion, i.e., we keep track of the points/data within a bucket for both meshes. Since *A* and *B* are (locally) off by one level, when we reach the leaf level on either *A* or *B*, three cases are possible,

- The buckets of *A* and *B* are at the same level and are leaves: we simply copy the finite-element values associated with this bucket (element),
- The bucket for *A* is at the leaf level, but the bucket for *B* is not: in this case, *B* is locally refined and we need to interpolate values from the element of *A* to its children. The elemental interpolation is done by evaluating the local shape functions at the locations of the child nodes.

⁷ Additional levels of coarsening or refinement will need multiple AMR steps, and are supported.

- The bucket for B is at the leaf level, but the bucket for A is not: in this case, B is locally coarsened and we need to project values from the elements of A to its parent.

Once the leaf-level transfer is done, the upward traversal is the same as for vector assembly.

5. Numerical experiments

5.1. 2D simulations: manufactured solutions

We use the method of manufactured solutions to assess the convergence properties of our method. We select an input “solution” which is solenoidal, and substitute it in the full set of governing equations. We then use the residual as a body force on the right-hand side of eqs. (17) to (15). We choose the following “solution” with appropriate body forcing terms:

$$\begin{aligned} \mathbf{v} &= \left(\pi \sin^2(\pi x_1) \sin(2\pi x_2) \sin(t), -\pi \sin(2\pi x_1) \sin^2(\pi x_2) \sin(t), 0 \right), \\ p &= \cos(\pi x_1) \sin(\pi x_2) \sin(t), \quad \phi = \mu = \cos(\pi x_1) \cos(\pi x_2) \sin(t). \end{aligned} \quad (46)$$

Our numerical experiments use the following non-dimensional parameters: $Re = 10$, $We = 1$, $Cn = 1.0$, $Pe = 3.0$, and $Fr = 1.0$. The density ratio is set to $\rho_-/\rho_+ = 0.85$.

For the first experiment we use a 2D uniform mesh with 512×512 bilinear elements (quads). Panel (a) of Fig. 3 shows the temporal convergence of the L^2 errors (numerical solution compared with the manufactured solution) calculated at $t = \pi$ to allow for one complete time period. For panel (a) we show the convergence for the forcing term of $\frac{Cn}{We} \frac{\partial}{\partial x_j} \left(\frac{\partial \phi}{\partial x_i} \frac{\partial \phi}{\partial x_j} \right)$. The figure shows the evolution of the error versus time-step on a log-log scale. The errors are decreasing with a slope of two for the phase-field function ϕ , thereby demonstrating second-order convergence. For velocity the slope on the log-log scale is initially about 2.0, but then tapers off for smaller time-steps; we expect this tapering off at smaller time-steps due to the fact that we used a fixed mesh and at some point the spatial errors dominate. As the scheme only uses pressure at previous timestep, the expected order for pressure here is 1 and in panel (a) we see the expected slope of 1 for pressure. In panel (b) we repeat the temporal convergence study in panel (a) with a forcing of $\frac{1}{CnWe} \phi \frac{\partial \mu}{\partial x_i}$. We observe that the velocity errors saturate at much larger timesteps compared to panel (a) and then they never decrease. Recall Remark 4 where we show that the forcing used in panel (a) and panel (b) are not equivalent unless the basis functions are divergence which is not the case for our conforming Galerkin finite element method. We use $\frac{Cn}{We} \frac{\partial}{\partial x_j} \left(\frac{\partial \phi}{\partial x_i} \frac{\partial \phi}{\partial x_j} \right)$ forcing for all the other cases in this work as it shows correct 2nd order convergence for velocity.

We next conduct a spatial convergence study. We fix the time step at $\delta t = 5 \times 10^{-4}$, and vary the spatial mesh resolution. Panel (c) of Fig. 3 shows the spatial convergence of L^2 errors (numerical solution compared with the manufactured solution) at $t = \pi$. We observe second order convergence for velocity, ϕ and pressure as expected for linear basis functions.

Panel (d) of Fig. 3 shows mass conservation for an intermediate resolution simulation with $\delta t = 5 \times 10^{-4}$ and 256×256 elements. We plot mass drift:

$$\int_{\Omega} \phi(\mathbf{x}, t) d\mathbf{x} - \int_{\Omega} \phi(\mathbf{x}, t = 0) d\mathbf{x}, \quad (47)$$

and expect this value to be close to zero as per the theoretical prediction of Proposition 1. We observe excellent mass conservation with fluctuations of the order of 10^{-7} , which is to be expected in double precision arithmetic. We detail the command-line arguments along with tolerances for iterative solvers used for the Algebraic Multigrid Method from PETSc in Appendix A.1.

5.2. 2D simulations: single rising bubble

To validate the framework, we consider benchmark cases for a single rising bubble in a quiescent water channel [59–61]. We set the Froude number ($Fr = u^2/(gD)$) to 1, which fixes the non-dimensional velocity scale to $u = \sqrt{gD}$, where g is the gravitational acceleration, and D is the diameter of the bubble. This scaling gives a Reynolds number of $\rho_c g^{1/2} D^{3/2} / \mu_c$, where ρ_c and μ_c are the specific density and specific viscosity of the continuous fluid, respectively. The Archimedes number, $Ar = \rho_c g^{1/2} D^{3/2} / \mu_c$, scales the diffusion term in the momentum equation. The Weber number here becomes $We = \rho_c g D^2 / \sigma$. We use the density of the continuous fluid to non-dimensionalize: $\rho_+ = 1$. The density and viscosity ratios are ρ_+/ρ_- and ν_+/ν_- , respectively. We present results for two standard benchmark cases.

Table 1 shows the parameters and the corresponding non-dimensional numbers for the two cases simulated in this work. The bubble is centered at $(1, 1)$, and since our scaling length scale is the bubble diameter, the bubble diameter for our simulations is 1. The domain is $[0, 2] \times [0, 4]$. Following the benchmark studies in the literature, we choose the top and bottom wall to have no slip boundary conditions and the side walls to have boundary conditions: $v_1 = 0$ (x -velocity) and $\frac{\partial v_2}{\partial x} = 0$ (y -velocity). We detail the command-line arguments along with tolerances for iterative solvers used for the




Fig. 3. Manufactured Solution Examples-chns-block-iteration-kt: (a) Temporal convergence of the numerical scheme for the case of manufactured solutions for conservative forcing; (b) Temporal convergence of the numerical scheme for the case of manufactured solutions for non-conservative forcing; (c) Spatial convergence of the numerical scheme for the case of manufactured solutions for conservative forcing; (d) the mass conservation for the case of manufactured solutions using 256×256 elements with time step of 5×10^{-4} . Here f denotes variables of interest which are velocities (v_1 and v_2), pressure p , phase field ϕ , and chemical potential μ . (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Table 1

Physical parameters and corresponding non-dimensional numbers for the 2D single rising drop benchmarks considered in section 5.2.

Test Case	ρ_c	ρ_b	μ_c	μ_b	ρ_+/ρ_-	ν_+/ν_-	g	σ	Ar	We	Fr
1	1000	100	10	1.0	10	10	0.98	24.5	35	10	1.0
2	1000	1.0	10	0.1	1000	100	0.98	1.96	35	125	1.0

Algebraic Multigrid Method from PETSc in Appendix A.2. We use a time step of 2.5×10^{-3} for test case 1 and 1×10^{-3} for test case 2. It is important to note that this is a substantially higher time step compared to fully linear block schemes presented in [24,28,30,31,33] and on par with fully implicit schemes such as [17,19,48] even though the velocity prediction, pressure Poisson, and velocity update steps are linear. This is the advantage of using a non-linear time-discretization for Cahn-Hilliard (CH). Our scheme is not limited by interfacial relaxation timescales. Interestingly, in practice, the non-linear iteration converges in a single Newton iteration, thus becoming computationally equivalent to a linear scheme. Finally, we notice that this scheme allows for a larger range of timesteps.

5.2.1. Test case 1

This test case considers the effect of higher surface tension, and consequently less deformation of the bubble as it rises. We compare the bubble shape in Fig. 4 with benchmark quantities presented in three previous studies [59–61]. We take $Cn = 1 \times 10^{-2}$ for this case. Panel (a) of Fig. 4 shows a shape comparison against benchmark studies in the literature for the case with Adaptive Mesh Refinement (AMR) vs Uniform Mesh. We see an excellent agreement in the shape of the bubble. Panel (b) of Fig. 4 shows a comparison of centroid locations with respect to time against benchmark studies in the




Fig. 4. 2D Single Rising Drop Test Case 1 (section 5.2.1). Shown in the panels are (a) comparisons of the computed bubble shape against results from the literature at non-dimensional time $T = 4.2$; (b) comparisons of the rise of the bubble centroid against results from the literature; (c) decay of the energy functional illustrating energy stability; and (d) total mass conservation (integral of total ϕ). (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

literature; again, we see an excellent agreement for both AMR and Uniform Mesh. We can see from the magnified inset besides of panel (a) of Fig. 4 that both the cases with AMR and uniform meshes show excellent agreement with benchmark studies.

We next check whether the numerical method follows the theoretical energy stability.⁸ We present the evolution of the energy functional defined in eq. (11) for test case 1. Panel (c) of Fig. 4 shows that the energy is decreasing in accordance

⁸ The energy function in eq. (11) should decrease as a function of time.

with the energy stability condition for both the cases of AMR and uniform meshes. The energy profile using the AMR mesh is nearly identical to the uniform mesh, suggesting that adaptive meshes can capture the energy dissipation correctly.

Finally, we check the mass conservation. Panel (d) shows the total mass of the system minus the initial mass. For the uniform mesh the change in the total mass is of the order of 10^{-8} , even after 1600 time steps. As expected, we see more drift for the AMR mesh compared to the uniform mesh. This is because the interpolation strategy used in adaptive mesh refinement to interpolate fields between adapted meshes is *not mass conserving*. Designing mass conserving interpolations for conforming Galerkin methods is an open question and out of the scope of the current work. Excellent mass conservation with the uniform mesh show that the numerical method delivers excellent mass conservation for long time horizons.

It is useful to compare the computational cost of using an AMR mesh versus using a uniform fine mesh. The uniform mesh has ~ 2 million elements. The AMR mesh on the other hand has a mesh count of about 54,000 elements. It is important to note that for this test case the We number is small and the shape of the bubble does not change much therefore, one does not require very small Cn to resolve the interfacial dynamics. Thus, in this case the resolution for interface is not very high and we do not anticipate a large efficiency gain.

5.2.2. Test case 2

This test case considers a lower surface tension resulting in high deformations of the bubble as it rises. As before, we compare the bubble shape in Fig. 5 with benchmark quantities presented in [59–61]. Panel (a) shows the shape comparison with benchmark studies in the literature. We see an excellent agreement in the shape of the bubble. All simulations (our results and benchmarks) exhibit a skirted bubble shape. We see an excellent match in the overall bubble shape with minor differences in the dynamics of its tails between the AMR and uniform meshes. Specifically, we see that the tails of the bubble in our case pinch-off to form satellite bubbles.⁹ The slight difference between the shape and location of the satellite drops between the AMR and uniform meshes is expected as there is some difference in the convection of the satellite drops due to lower resolution in the bulk. We performed this simulation with a $Cn = 0.005$ for both the AMR and uniform meshes. We can see in panel (a) of Fig. 5 that our simulation captures this filament pinch-off in the tails. The works of Aland and Voigt [60], Yuan et al. [61] did not observe these thin tails and pinch-offs, while Hysing et al. [59] described pinch-off of the tails and satellite bubbles.

Panel (b) of Fig. 5 compares the centroid location evolution with time. Again we see an excellent agreement with all three previous benchmark studies. We can see from the magnified inset in this panel that for both AMR and uniform meshes the plot approaches the benchmark studies and we see an almost exact overlap between the benchmark. Next, we report the evolution of the energy functional defined in eq. (11) for test case 2. Panel (c) of Fig. 5 shows the decay of the total energy functional in accordance with the energy stability condition for both AMR and uniform meshes. Finally, panel (d) of Fig. 5 shows the total mass of the system in comparison with the total initial mass of the system. We can see that for all spatial resolutions, the change in the total mass against the initial total mass is of the order of 1×10^{-8} , which illustrates that the numerical method satisfies mass conservation over long time horizons. As expected, for the AMR case just like in the test case 1 we see higher drift from mass conservation with AMR. As stated earlier, we attribute this to a non-mass conserving coarsening interpolation.

It is worth noticing the advantage of deploying AMR in this case. The uniform mesh has 8.4 million elements and takes 6 hours on 20 TACC Frontera node (120 node hours). The AMR mesh, on the other hand, has a maximum mesh count of about 0.1 million elements at peak refinement and takes 4 hours on 2 TACC Frontera nodes on (8 node hours). We get a much higher speed up for this test case compared to test case 1. High We results a lot in the deformation of bubble shape and therefore, requires very small Cn to resolve the interfacial dynamics. Therefore, the minimum resolution required in this case is significantly higher compared to the test case 1.

5.3. 2D simulations: Rayleigh-Taylor instability

We now demonstrate the performance of the numerical framework with large deformation in the interface and chaotic regimes (high Reynolds numbers) with a moderately high density ratio. While the bubble rise case in the previous subsection is an interplay between surface tension and buoyancy, buoyancy dominates the evolution of the Rayleigh-Taylor instability (RTI). Here the choice of non-dimensional numbers ensures that the surface tension effect is small (high Weber number). In contrast, other studies switch off the surface tension forcing terms in the momentum equations (see, e.g., [9,62–64]). We choose a finite but small surface tension because it ensures more stable filaments which are generated when the interfacial dynamics become more chaotic. This represents a more difficult case to simulate compared to the zero surface tension counterpart as in the chaotic regime of the RTI the thin filaments and satellite drops generated need to be resolved with a very high resolution. Additionally, one has to take smaller timesteps to resolve pinch offs during the breakup of these filaments and satellite drops.

The setup is as follows: the heavier fluid is on top of lighter fluid and the interface is perturbed. The heavier fluid on top penetrates the lighter fluid and buckles, which generates instabilities. This interface motion is challenging to track due to large changes in its topology. Additionally, the long term RTI generally encompasses turbulent conditions that require

⁹ Such instabilities require a low Cn number, as only a thin interface can capture the dynamics of the thin tails of the bubble.




Fig. 5. 2D Single Rising Drop Test Case 2 (section 5.2.2). Shown in the panels are (a) comparisons of the computed bubble shape against results from the literature at non-dimensional time $T = 4.2$; (b) comparisons of the rise of the bubble centroid against results from the literature; (c) decay of the energy functional illustrating energy stability; and (d) total mass conservation (integral of total ϕ). (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

resolving finer scales to capture the complete dynamics. We non-dimensionalize the problem by selecting the width of the channel as the characteristic length scale and the density of the heavier fluid as the characteristic specific density. Just as in the bubble rise case we use buoyancy-based scaling, setting the Froude number ($Fr = u^2/(gD)$) to 1, which fixes the non-dimensional velocity scale to be $u = \sqrt{gD}$, where g is the gravitational acceleration, and D is the width of the channel. Using this velocity to calculate the Reynolds number, we get $Re = \rho_L g^{1/2} D^{3/2} / \mu_L$, where ρ_L and μ_L are the specific density and specific viscosity of the light fluid, respectively. We set the Reynolds number to $Re = 3000$. These choices lead to a Weber number of $We = \rho_c g D^2 / \sigma$. To compare our results with previous studies, we simulate with the same initial

conditions as presented in Xie et al. [9], Khanwale et al. [19]. The We number is selected to be 100, so that the effect of surface tension is small on the evolution of the interface but still finite.

The Atwood number (At) is often used to parameterize the dependence on density ratio, with $At = (\rho_+ - \rho_-) / (\rho_+ + \rho_-)$. For the density ratio of 0.1, the Atwood number is $At = 0.82$. We chose specific density of the heavy fluid to non-dimensionalize, therefore $\rho_+ = 1.0$, and $\rho_- = 0.33$ for $At = 0.5$, while $\rho_- = 0.1$ for $At = 0.82$. ν_+/ν_- the viscosity ratio is set to 1. We use a no-slip boundary condition for velocity on all the walls along with no flux boundary conditions for ϕ and μ . The no-flux boundary condition for ϕ and μ inherently produces a 90 degree wetting angle (neutral contact angle) for both the fluids.

Weak surface tension reduces vortex roll-up in the simulations of immiscible systems, especially at lower At . Waddell et al. [65] experimentally showed different vortex roll-up amounts for miscible and immiscible systems. This difference is analogous to the difference between zero surface tension simulations and finite surface tension simulations. This effect is irrelevant to compare front locations against the literature (short time horizons). Nevertheless, it is crucial to accurately track the long time dynamics (as smaller filaments are more stable in the non-zero surface tension case). A finite surface tension would generate different long term mixing and breakup of the interface.

We run numerical experiment for a long term RTI with $Cn = 0.00125$ ¹⁰ with an adaptive mesh where the refinement varies from $8/2^{10}$ to $8/2^{14}$, for the Atwood number of $At = 0.82$. We use a variable timestep for this simulation. We use a timestep of 1.25×10^{-4} till non-dimensional time $t = 0.6$ as we have linear behavior of the interface. We then decrease the timestep to 5×10^{-5} till $t = 1.4$ as the interface starts accelerating, and then we further reduce it to 2.5×10^{-5} at $t = 1.4$ in anticipation of the shedding and breakup events near the interface. A carefully tuned algebraic multi-grid linear solver with successive over-relaxation is setup for the velocity, Cahn-Hilliard and pressure Poisson solvers. We detail the command-line arguments along with the tolerances used in Appendix A.3. Note that the resolution here is finer than the resolution used in [19] where we run the same case with an uniform mesh of 800×6400 elements.

Fig. 6 shows the snapshots of the interface shape along with the corresponding vorticity generated as it evolves in time for $At = 0.82$ and $Cn = 0.00125$. We observe the usual evolution of Rayleigh Taylor instability where the heavier fluid penetrates the light one, causing the lighter fluid to rise near the wall. The penetrating plume of the heavier fluid sheds small filaments at a non-dimensional time of around $t' = 1.9$. The penetrating plume is symmetric at early times, with symmetry breaking occurring at longer times due to increasingly chaotic behavior. The instability further proceeds to a periodic flapping. At longer times, the instability transitions to a chaotic mixing stage. Notice that there are strong vortex roll-ups generated near the thin filaments (e.g. see vorticity in panels (f), (g), (h)) which influence the onset and progression of chaotic behavior in RTI. As this is a finite surface tension case these filaments are more stable and require high resolution to resolve these fine scale vortical structures near the interface.

We observe from Fig. 6 that the development of the instability (at longer times) depends on the resolution of shedding filaments. Therefore, the long-time dynamics depend on the interface thickness that the Cn number controls. Khanwale et al. [19] presented analysis of how Cn numbers change the long term mixing behavior of RTI. They performed RTI simulations for $Cn = 0.005, 0.0025, 0.00125$ and showed that for $Cn = 0.00125$ the mixing in RTI is substantially different at long times. In the present work we choose the finest $Cn = 0.00125$ which requires a very high resolution mesh. However, in the present work we use an adaptive mesh to capture such a fine Cn interface representation. First we compare the results of short term with literature in panel (a) of Fig. 7 which shows the comparison of locations of the bottom and top front of the interface. We see an excellent agreement for short times with the literature. Note that for short times the effect of the finite surface tension is not noticeable and the front locations match very well with the literature. This is because for shorter times the dynamics are dominated by large scale flow structures. We also compare our interface fronts with Khanwale et al. [19] where RTI simulations were performed with a uniform mesh with finite surface tension (same We number as the present study) and we also see excellent match even for longer times. Panel (b) from Fig. 7 shows the decay of the energy functional with respect to time demonstrating energy stability for $At = 0.82$. Panel (c) from Fig. 7 shows the drift of mass which is expected to be as close to zero as possible. However, we see that for long term behavior the mass drift increases, we attribute this behavior to the interpolation operations during coarsening of the adaptive mesh refinement, this is similar to behavior in panel (d) of Fig. 4 and Fig. 5.

To quantify vortex dynamics at long times in RTI we perform vortex identification using Q-criterion [66]. Fig. 8 shows turbulent instability in RTI. Panel (a) shows the complex mixing structure of the interface in RTI in the chaotic regime, panel (b) of Fig. 8 shows the vorticity in the system after a very long time ($t' = t\sqrt{At} = 4.962$), panel (c) shows the Q-criterion at the same time. To identify vortices we threshold the Q-criterion between $Q = 1.9$ and $Q = 2$ which is shown in panel (d) of Fig. 8. It can be clearly seen that a myriad of small scale structures generated near the filaments due to shear instability, and are resolved in our simulation. Note that as this is a 2D case, therefore there is no vortex stretching, instead we have a redistribution of vorticity through the non-linear convection. It is well known [67] that there is a strong inverse cascade in the case of 2D turbulence and small scale structures combine to form large scale vortical structures. Therefore, it is even more important to resolve these near interface vortical structures which eventually feed in energy to the larger scales to correctly capture the large scale mixing in the system. This fine scale resolution has a profound influence on the higher-order statistics of Rayleigh-Taylor instability. We defer a detailed analysis of higher-order statistics to future work.

¹⁰ Note that this is a very small Cn number and requires a very fine mesh. This corresponds to the finest resolution selected in [19] for the same case.




Fig. 6. Rayleigh-Taylor instability in 2D (section 5.3): Dynamics of the interface as a function of time for $At = 0.82$ (density ratio of 0.1). In each panel the left plot illustrates the interface, and the right plot shows corresponding vorticity. Here normalized time $t' = t\sqrt{At}$.

To demonstrate the resolution of the adaptive mesh that we are using for this case we show a progressive zoomed insets of the mesh overlayed on top of ϕ in Fig. 9. Notice that we use an adaptive mesh where the refinement varies from element size of $8/2^{10}$ to $8/2^{14}$, with higher refinement near the interface. This resulted in a maximum problem size of about 0.8 million nodes, which ran on TACC Frontera with 16 nodes for 40 hours. The corresponding uniform mesh with a resolution of $8/2^{14}$ resulted in about 33.5 million dof which would be substantially more expensive to run.




Fig. 7. Rayleigh-Taylor instability 2D (section 5.3): (a) Comparison of positions of top and bottom front of the interface with literature; (b) decay of the energy functional illustrating Energy Stability for $At = 0.82$; (c) total mass conservation (integral of total ϕ - initial mass) for $At = 0.82$. Note the excellent mass conservation across $\sim 260,000$ time steps. As stated in the text, most of this deviation is attributed to the non-mass conserving interpolation schemes used to map solutions between two consecutive meshes.

5.4. 3D simulations: Rayleigh-Taylor instability

We now deploy our framework in 3D and simulate the Rayleigh-Taylor instability in 3D using adaptive octree meshes. The domain is set to be cuboid with dimensions $1 \times 8 \times 1$ in x_1, x_2, x_3 directions respectively. For the 3D simulations we choose the following initial condition for ϕ to describe the interface:

$$\phi(\mathbf{x}) = \tanh\left(\sqrt{2}\left[\frac{x_2 - l - g(\mathbf{x})}{Cn}\right]\right), \quad (48)$$

$$g(\mathbf{x}) = 0.05[\cos(2\pi x_1) + \cos(2\pi x_3)]. \quad (49)$$

Here l is the location in the vertical direction for the interface, which in this case is chosen to be $x_2 = 6$. Fig. 11(a) shows the initial condition. We use a $Cn = 0.0075$ and $At = 0.15$. For this lower At number simulation, the effect of non-zero surface tension is important. The non-dimensionalization follows the same logic as the 2D cases, with the Reynolds number set to $Re = 1000$, the Weber number ($We = \rho_c g D^2 / \sigma$) set to 1000, and viscosity ratio, ν_+ / ν_- , set to 1. For an Atwood number $At = 0.15$ the density ratio is $\rho_+ / \rho_- = 1.0 / 0.74$. Similar to the 2D cases, the boundary conditions are no-slip for velocity and no-flux for ϕ and μ on all the walls. This is different than the typical boundary condition of free slip on the side walls, which is used in the literature [10,62,68]. Another important thing to note is that generally zero surface tension is used, however in the current study as well as in Liang et al. [68] the surface tension is small but finite. We choose the no slip velocity on all walls to test the energy stability of the method, which requires the no-slip boundary conditions (see [19,30,48] for the setup of the proof).

We choose a time-step size of $\delta t = 0.001$. We refine near the interface to a level corresponding to element length of $8/2^{12}$, ensuring about six elements for resolving the diffuse interface, while the refinement away from the interface is $8/2^7$,




Fig. 8. Rayleigh-Taylor instability in 2D (section 5.3): turbulent dynamics in Rayleigh Taylor instability for $At = 0.82$ (density ratio of 0.1) at $t' = t\sqrt{At} = 4.962$. (a) Shows the interface; (b) shows vorticity (c) Q-criterion [66] (d) Q-criterion threshold between $Q = 1.9$ and $Q = 2$.

we also refine near the walls with a level of $8/2^9$. In Appendix A.3 we provide a detailed description of the preconditioners and linear solvers along with the command-line arguments that are used.

Fig. 10 shows the evolution of the interface along with the solution-adapted mesh. We color the mesh with the order parameter value (blue for the heavy fluid and white for the light fluid) to show the evolution of the system. It is seen that as the interface evolves it deforms and expands, causing the mesh density to gradually increase. This gradual growth helps with the efficiency of the simulation, since a uniform mesh for this case would be computationally expensive. At the point with the largest interfacial area, the simulation has around 52 million nodes, corresponding to 152 million degrees of freedom for the largest matrix assembly (of velocity prediction). The efficient and scalable implementation of the approach allows us to run this large scale simulation. We run the simulation till $t = 6.8$ on SDSC Expanse with 32 nodes (4096 processes). The simulation is subsequently run on TACC Frontera with 80 Cascade Lake nodes (4480 processes). We use this case to perform scalability analysis which is presented later in the scaling section (section 6).

Fig. 11 shows that the initial sinusoidal perturbation develops into penetrating plume of heavy fluid pushing down while the lighter fluid buckles and forms bubbles near the wall. The simulation maintains symmetry. As different parts of the interface move in opposite directions, Kelvin-Helmholtz instabilities cause the plumes to roll up, which in turn causes mushroom-like structures to develop (see Fig. 11(d)). However, the Kelvin-Helmholtz instability is arrested a bit due to finite




Fig. 9. Rayleigh-Taylor instability in 2D (section 5.3): Progressive zoomed in mesh for Rayleigh Taylor instability for $At = 0.82$ (density ratio of 0.1) at $t' = t\sqrt{At} = 3.006$. These plots are zoomed insets of the domain near the interfacial instabilities.




Fig. 10. Rayleigh-Taylor instability in 3D (section 5.4): Snapshots of the mesh at various time-points in the simulation for Rayleigh-Taylor instability for $At = 0.15$. The figures show the mesh on two slices to illustrate the refinement around the interface of two fluids represented by the gray iso-surface of $\phi = 0$. The phase field ϕ values color the mesh, where blue represents heavy fluid and white represents light fluid. Here $t(-)$ is the non-dimensional time. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

surface tension curbing breakup as seen in Fig. 11(f). This behavior is also presented in Liang et al. [68] who also used a CHNS model deployed using a Lattice Boltzman method. However, for the case of zero surface tension Jain et al. [10] show breakup at the same time $t = 7.8$. Another feature that is unique to our simulations compared to the simulations in the literature is the development of the bubbles inside the domain instead of at the walls. This is due to the no-slip conditions of velocity on the side walls instead of free-slip. Although these two types of spikes (upwards/downwards) begin to develop in a checkerboard pattern that preserves symmetry, their dynamics are different due to the velocity differential that the two fronts face.

The downward spikes undergo further deformation and we see the emergence of four protrusions from the mushroom structure (see Fig. 11(f)) caused by the shear generated between the fluids. Liang et al. [68] and Jain et al. [10] also report




Fig. 11. Rayleigh-Taylor instability in 3D (section 5.4): Zoomed in snapshots of the zero isosurface of ϕ which represents the interface at various time-points in the simulation for Rayleigh-Taylor instability for $At = 0.15$. Here $t(-)$ is the non-dimensional time.

Table 2

The level of refinement for two different meshes used to conduct the scaling studies.

	Refinement level			Mesh size
	Bulk	Wall	Interface	
Mesh M1	6	7	9	5 M
Mesh M2	8	10	13	30 M

four secondary filaments (but with breakup) in their simulations for the same At number, although their simulations were for zero surface tension (i.e., dynamics similar to miscible systems).

On the other hand, the mushroom structures from the upward spikes develop into long and thin circular films. The upward spikes develop circular films adjacent to the wall “bubbles” (i.e., structures near the wall that the heavier fluid generates as it is displaced by the lighter one) interact with these bubbles to merge and form larger structures (see Fig. 11(f)). While the central plumes have little-to-no interaction with the wall, the bubbles continue to rise and ultimately collide with the top wall.

We compare the front location of the downward spike with Liang et al. [68] in Fig. 12. We see a slight deviation from Liang et al. [68] which we attribute to the no-slip boundary conditions. Panel (a) from Fig. 13 shows the decay of the energy functional with respect to time demonstrating energy stability on the fully discrete 3D AMR meshes. Panel (b) from Fig. 13 shows the drift of mass which is expected to be as close to zero as possible. However, we see that for long term behavior the mass drift increases, we attribute this behavior to the interpolation operations during coarsening of the adaptive mesh refinement.

6. Scaling

In this section, we show preliminary scaling results of our solver. In order to conduct the scaling we choose the 3D Rayleigh Taylor instability problem (section 5.4). The domain is a cube with dimensions $8 \times 1 \times 1$. We chose two different meshes to conduct the scaling study with varying level of refinements. The overall mesh is characterized by three different regions of refinement: bulk/base level (refinement level in the bulk), wall level (refinement level at the wall), and interface refinement level (refinement level at the interface, indicated by $|\phi| \leq 0.95$). A refinement level of k , would mean a mesh resolution of $8/2^k$. Table 2 shows the different refinement levels. As the simulation progress, the region close to the interface, $|\phi| \leq 0.95$ is refined, while regions away from the interface are coarsened. We report the scaling behavior conducted over 5 time-steps.

Fig. 14 shows the strong scaling behavior of the individual solvers. The Cahn-Hilliard (CH), the velocity prediction, and the pressure Poisson (PP) updates were solved using PETSc’s Algebraic Multigrid (AMG), whereas the velocity update was solved using conjugate gradient with Additive Schwarz preconditioner (see Appendix A). We chose AMG for the first three updates on account of its $\mathcal{O}(N)$ complexity. Overall we see a good scaling behavior until the AMG setup cost begins to dominate. Previous studies have demonstrated the inability of AMG to scale to a very large number of processors [69,70] due to the associated communication and setup costs at such scales. This becomes particularly expensive in the case of adaptive meshes, as the AMG setup needs to be executed each time re-meshing is performed. We see that the velocity prediction suffers the most from AMG as it has the largest matrix size among the equations solved with AMG. We note that PETSc’s native AMG is degree-of-freedom (dof) agnostic, which makes the AMG cost grow with increasing dof. Geometric Multigrid (GMG) tends to eliminate such cost [70], and is particularly beneficial for hierarchical data structures like octrees.




Fig. 12. Normalized spike amplitude for the case of $Re = 1000$ and $We = 1000$ in the simulation of the 3D Rayleigh-Taylor instability (section 5.4).

Panels (a), (b), and (c) of Fig. 14 show the strong scaling behavior of the Cahn-Hilliard solve (Definition 6), velocity prediction step (Definition 3), and pressure Poisson step (Definition 5), respectively. In panels (a), (b), and (c) of Fig. 14 we see good scaling behavior for the first three points and then it starts to taper off as the communication and setup cost of PETSc's native AMG increases. This behavior is illustrated in the histogram of the breakdown of cost for each solver




Fig. 13. (a) Decay of the energy functional; (b) Mass conservation for the case of $Re = 3000$ and $We = 1000$ in the simulation of the 3D Rayleigh-Taylor instability (section 5.4).

in Fig. 17: we can see that the preconditioner (PC) setup cost increases with increasing processors for the Cahn-Hilliard solve, velocity prediction, and pressure Poisson; all of these updates are using AMG. On the other hand in panel (d) of Fig. 14 we can see that velocity update solve scales well, even beyond the first three points as it is using a simple conjugate gradient solver with an Additive Schwarz preconditioner, which has relatively low communication and setup costs. This point is highlighted in Fig. 17, which shows a histogram of the cost breakdown for the pressure Poisson solve with varying processor number.




Fig. 14. Strong scaling: Shown in the panels are the strong scaling behaviors on TACC Frontera for the (a) Cahn-Hilliard (CH), (b) velocity prediction (VP), (c) pressure Poisson (PP), and (d) velocity updates (VU).

Fig. 15 shows the scaling behavior of total time to solution as a function of the number of processors. We see a good scaling behavior until the AMG setup costs begin to dominate; this effect is shown in Fig. 16, in which the cost for each major component of the solution is tabulated as a function of the number of processors. A constant parallel cost would indicate ideal strong scaling. Overall, we see a constant growing cost with increase in processor with the AMG cost starting to dominate for $\geq 3K$ processors.

Fig. 17 shows the time taken by individual components for mesh M1 used in the scaling study. It can be clearly seen that each solver has different bottlenecks that would require different optimization strategies to overcome. For instance, the preconditioner (PC) setup and matrix assembly is dominant for the Cahn-Hilliard and velocity prediction solves, whereas vector assembly begins to equally dominate for the pressure Poisson (PP) solver. The matrix assembly becomes a bottleneck during the velocity update stage. With increasing number of processors, we can note that the percentage of time taken by AMG setup begins to dominate. This is a characteristic of AMG which has motivated the developers to look into GMG [70], especially at exascale. This is left as an avenue of future work. A more sophisticated roofline analysis will help to suggest more advanced optimization strategies like vectorization or cache blocking for each individual steps. However, we can clearly see a major advantage of using the current projection based approach. Unlike the previous analysis carried out in Khanwale et al. [19], where PC setup was the most dominant part and can take up to 80% of the total compute time,




Fig. 15. Strong scaling: shown in the panels are (a) the total time to solution as a function of the number of processors, and (b) the relative speed up as a function of the number of processors. These results were computed on the TACC Frontera supercomputer.




Fig. 16. Strong scaling parallel cost: (Run-time) $\times$ (number of cores) evaluated on the 3D Rayleigh Taylor problem on the Frontera supercomputer. A flat line would indicate ideal strong scaling. Here in the legend: CH-Solve corresponds to Cahn-Hilliard solve; VP-Solve corresponds to velocity prediction solve; PP-Solve corresponds to pressure Poisson Solve; VU-solve corresponds to velocity update solve.




Fig. 17. The panels show the percentage of time taken for matrix assembly, vector assembly, preconditioner (PC) setup, and other parts of the update as a function of the number of processors. The panels are organized as follows: (a) Cahn-Hilliard, (b) velocity prediction, (c) pressure Poisson, and (d) velocity update.

we see that we no longer rely on expensive ASM/LU¹¹ preconditioners to ensure the convergence of the solver. However, further optimization would require us to abandon the vendor optimized library, especially for improve performance during the matrix assembly and remeshing stages.

6.1. Performance comparison with the fully coupled method

The projection method presented in this paper in conjunction with the variational multiscale approach, allows us to construct operators which are elliptic in each block. This allows for a selection of cheaper preconditioners compared to the pressure stabilized approach in Khanwale et al. [19]. Additionally, the block projection approach also allows us to selectively provide implicit treatment and decreases cost, e.g. Cahn-Hilliard operators are fully implicit, while the Navier-Stokes operators are now semi-implicit, which linearizes the operator. To assess the performance gains of the projection based method in this paper compared to the fully-implicit method in Khanwale et al. [19], we perform a strong scaling analysis for both methods. For better comparison, the fully-coupled method from Khanwale et al. [19] is implemented in the current parallel framework using the same mesh (M2 from Table 2). We use the same case of 3D Rayleigh Taylor instability used for performing the rest of scaling analysis. The timestep size used in both the methods is also kept the same, so there is no artificial speedup due to small timestep for the projection method. Fig. 18 shows the strong scaling behavior

¹¹ Additive Schwarz preconditioners with LU were used locally for each block.




Fig. 18. Comparison of total solve time presented in this work in comparison to the fully coupled method of Khanwale et al. [19].

for both the fully coupled method from Khanwale et al. [19] and the projection method presented in the current work. We can observe that both methods scale well, however the projection method is consistently cheaper than the fully-coupled method. We observe a speedup of roughly 2.5 to 4 times with the block projection method.

7. Conclusions and future work

In this work we developed a projection based numerical framework for solving the Cahn-Hilliard Navier-Stokes (CHNS) model of two-phase flows. We developed a projection based second order in time numerical scheme to improve on the fully implicit scheme in Khanwale et al. [19]. The pressure projection method allowed us to decouple the pressure equation. The decoupling of pressure allowed us to use highly efficient and parallel Algebraic Multigrid Methods as all operators are close to elliptic. We developed a block iterative, hybrid semi-implicit-fully-implicit in time method. The new method requires Newton iteration only for Cahn-Hilliard resulting faster solution time at relatively larger time steps. We utilized a conforming Galerkin method for spatial discretization with variational multiscale approach (VMS) approach. VMS method allows us to stabilize the pressure projection equation as pressure and velocity is represented by equal order basis functions. We deployed this approach on the massively parallel octree based adaptive meshing framework based on KD-trees called DENDRO-KT. We verified numerically that the scheme is energy stable and mass conserving. We tested the framework for a comprehensive set of canonical cases for validation. We run the long time turbulent 2D Rayleigh-Taylor cases with adaptive meshing, resolving thin stable filaments with a very high resolution. Further, we performed a detailed scaling analysis of numerical framework and showed that the time required for preconditioning and Jacobian assembly is minimized compared to the method in Khanwale et al. [19].

We identify three main avenues of future work. First one entails coming up with efficient mass conserving interpolation strategies with conforming Galerkin methods. This is currently an open question. The second avenue is rigorously proving energy stability of the current numerical scheme. Rigorous energy stability of such second order projection schemes is tedious to prove. Finally, we notice that the algebraic multigrid methods become very expensive after certain number of processors are reached due to excessive communication. We are working on a tailor made multi-grid method for octree meshes which optimizes communication for extreme scalability of the iterative solvers.

CRedit authorship contribution statement

Makrand A. Khanwale: Conceptualization, Methodology, Software, Formal Analysis, Writing, Visualization. **Kumar Saurabh:** Software, Visualization, Writing. **Masado Ishii:** Software, Writing. **Hari Sundar:** Software, Writing, Analysis. **James A. Rossmanith:** Methodology, Analysis, Writing, Supervision. **Baskar Ganapathysubramanian:** Conceptualization, Methodology, Analysis, Writing, Resources, Supervision, Project Administration.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request. The source code that produced all simulations in this paper is publicly available on the “master” branch at <https://bitbucket.org/baskargroup/proteus/src/master/>.

Acknowledgements

The authors acknowledge XSEDE grant number TG-CTS110007 for computing time on TACC Stampede2 and SDSC Expanse. The authors also acknowledge computing allocation through Pathways award number PHY20033 on TACC Frontera, and Iowa State University computing resources through NSF 2018594. JAR was supported in part by NSF Grants DMS-1620128 and DMS-2012699. BG, KS, MAK were funded in part by NSF grant 1935255, 1855902. HS and MI were funded in part by NSF grant 2008772.

Appendix A. Details of solver selection for the numerical experiments

Solver settings presented below are broken up in 4 structures for 4 blocks we solve i.e., velocity prediction (see Definition 3), pressure Poisson (see Definition 4), velocity update (see Definition 5), and Cahn-Hilliard (CH) nonlinear solve (see Definition 6) respectively.

A.1. Manufactured solutions

For the cases presented in section 5.1 we use the flexible GMRES linear solver algebraic multigrid based preconditioning for pressure Poisson (`solver_options_pp`), CH (`solver_options_ch`), and a biCGstab based Krylov solver with Additive Schwarz based preconditioning for velocity prediction (`solver_options_momentum`). For better reproduction, the command line options we provide PETSC are given below which include some commands used for printing some norms as well. Here `ksp_atol` here is the absolute tolerance of the linear solver, `ksp_rtol` is the relative tolerance of the linear solver, and `snes_rtol`, `snes_atol` are relative and absolute tolerances of Newton iteration. PETSc AMG's inner smoothers are chosen to be a gmres (gmres) method with successive over relaxation preconditioning (sor).

```
solver_options_momentum = {
    ksp_atol = 1e-10
    ksp_rtol = 1e-10
    ksp_type = "bcgs"
    pc_type = "asm"
# monitor
    ksp_converged_reason = ""
};

solver_options_pp = {
    ksp_atol = 1e-10
    ksp_rtol = 1e-10
#multigrid
    ksp_type = "fgmres"
    pc_type = "gamg"
    pc_gamg_asm_use_agg = True
    mg_levels_ksp_type = "gmres"
    mg_levels_pc_type = "sor"
    mg_levels_ksp_max_it = 4
# monitor
    ksp_monitor = ""
    ksp_converged_reason = ""
};

solver_options_vupdate = {
    ksp_atol = 1e-10
    ksp_rtol = 1e-10
#ksp_monitor = ""
    ksp_converged_reason=""
    ksp_type = "cg"
    pc_type = "asm"
};

solver_options_ch = {
    snes_rtol = 1e-15
    snes_atol = 1e-15
    ksp_atol = 1e-12
    ksp_rtol = 1e-12
    ksp_stol = 1e-12
```

```
#multigrid
    ksp_type = "fgmres"
    pc_type = "gamg"
    pc_gamg_asm_use_agg = True
    mg_levels_ksp_type = "gmres"
    mg_levels_pc_type = "sor"
    mg_levels_ksp_max_it = 4
# monitor
    snes_monitor = ""
    snes_converged_reason = ""
    ksp_converged_reason = ""
};
```

A.2. Single bubble rise in 2D

In this case we use the flexible GMRES linear solver algebraic multigrid based preconditioning for velocity prediction (solver_options_momentum), pressure Poisson (solver_options_pp), and CH (solver_options_ch). A simple conjugate gradient method with no Additive Schwarz preconditioning is used for velocity update (solver_options_vu).

```
solver_options_momentum = {
    ksp_atol = 1e-8
    ksp_rtol = 1e-7

#multigrid
    ksp_type = "fgmres"
    pc_type = "gamg"
    pc_gamg_asm_use_agg = True
    mg_levels_ksp_type = "gmres"
    mg_levels_pc_type = "sor"
    mg_levels_ksp_max_it = 10
# monitor
    ksp_monitor = ""
    ksp_converged_reason = ""
};

solver_options_pp = {
    ksp_atol = 1e-8
    ksp_rtol = 1e-7
    ksp_stol = 1e-7

#multigrid
    ksp_type = "fgmres"
    pc_type = "gamg"
    pc_gamg_asm_use_agg = True
    mg_levels_ksp_type = "gmres"
    mg_levels_pc_type = "sor"
    mg_levels_ksp_max_it = 4
    pc_gamg_sym_graph = True

# monitor
    ksp_monitor = ""
    ksp_converged_reason = ""
};

solver_options_vupdate = {
    ksp_atol = 1e-8
    ksp_rtol = 1e-7
    ksp_stol = 1e-7
    ksp_converged_reason=""
    ksp_type = "cg"
    pc_type = "asm"
};

solver_options_ch = {
    snes_rtol = 1e-10
    snes_atol = 1e-10
    ksp_atol = 1e-8
    ksp_rtol = 1e-7
    ksp_stol = 1e-7

#multigrid
    ksp_type = "fgmres"
```

```

        pc_type = "gamg"
        pc_gamg_asm_use_agg = True
        mg_levels_ksp_type = "gmres"
        mg_levels_pc_type = "sor"
        mg_levels_ksp_max_it = 15

# monitor
    snes_monitor = ""
    snes_converged_reason = ""
    ksp_converged_reason = ""
};

```

A.3. Rayleigh-Taylor instability

In this case we use the flexible GMRES linear solver algebraic multigrid based preconditioning for velocity prediction (solver_options_momentum), pressure Poisson (solver_options_pp), and CH (solver_options_ch). A simple conjugate gradient method with no Additive Schwarz preconditioning is used for velocity update (solver_options_vu).

```

solver_options_momentum = {
    ksp_atol = 1e-8
    ksp_rtol = 1e-7

#multigrid
    ksp_type = "fgmres"
    pc_type = "gamg"
    pc_type = "none"
    pc_gamg_asm_use_agg = True
    mg_levels_ksp_type = "gmres"
    mg_levels_pc_type = "sor"
    mg_levels_ksp_max_it = 10

# monitor
    ksp_monitor = ""
    ksp_converged_reason = ""
};

solver_options_pp = {
    ksp_atol = 1e-8
    ksp_rtol = 1e-7
    ksp_stol = 1e-7

#multigrid
    ksp_type = "fgmres"
    pc_type = "gamg"
    pc_gamg_asm_use_agg = True
    mg_levels_ksp_type = "gmres"
    mg_levels_pc_type = "sor"
    mg_levels_ksp_max_it = 4

# monitor
    ksp_monitor = ""
    ksp_converged_reason = ""
};

solver_options_vupdate = {
    ksp_atol = 1e-8
    ksp_rtol = 1e-7
    ksp_converged_reason=""
    ksp_type = "cg"
    pc_type = "asm"
};

solver_options_ch = {
    snes_rtol = 1e-10
    snes_atol = 1e-10
    ksp_atol = 1e-8
    ksp_rtol = 1e-7

#multigrid
    ksp_type = "fgmres"
    pc_type = "gamg"
    pc_gamg_asm_use_agg = True

```

```

mg_levels_ksp_type = "gmres"
mg_levels_pc_type = "sor"
mg_levels_ksp_max_it = 15

# monitor
snes_monitor = ""
snes_converged_reason = ""
ksp_converged_reason = ""
};

```

This setup works very well with a relatively constant number of Krylov iterations as the number of processes are increased in the massively parallel setting. These same options are used for the simulation of Rayleigh-Taylor in 3D. The scaling results we present use the same setup of solvers except for the no of max ksp iterations on each multigrid level is increased to 30 (mg_levels_ksp_max_it).

References

- [1] F. Gibou, D. Hyde, R. Fedkiw, Sharp interface approaches and deep learning techniques for multiphase flows, *J. Comput. Phys.* 380 (2019) 442–463.
- [2] H. Xia, J. Lu, G. Tryggvason, Simulations of fused filament fabrication using a front tracking method, *Int. J. Heat Mass Transf.* 138 (2019) 1310–1319.
- [3] Y. Zhang, D. Han, Z. Dou, J.-C. Veilleux, G.H. Shi, D.S. Collins, P.P. Vlachos, A.M. Ardekani, The interface motion and hydrodynamic shear of the liquid slosh in syringes, *Pharm. Res.* 38 (2021) 257–275.
- [4] H. Ganti, P. Khare, Data-driven surrogate modeling of multiphase flows using machine learning techniques, *Comput. Fluids* 211 (2020) 104626.
- [5] S.L. Bruntton, B.R. Noack, P. Koumoutsakos, Machine learning for fluid mechanics, *Annu. Rev. Fluid Mech.* 52 (2020) 477–508.
- [6] K. Duraisamy, Perspectives on machine learning-augmented Reynolds-averaged and large eddy simulation models of turbulence, *Phys. Rev. Fluids* 6 (2021) 050504.
- [7] G.E. Karniadakis, I.G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics-informed machine learning, *Nat. Rev. Phys.* 3 (2021) 422–440.
- [8] Z. Guo, P. Lin, A thermodynamically consistent phase-field model for two-phase flows with thermocapillary effects, *J. Fluid Mech.* 766 (2015) 226–271.
- [9] Y. Xie, O. Wodo, B. Ganapathysubramanian, Incompressible two-phase flow: diffuse interface approach for large density ratios, grid resolution study, and 3D patterned substrate wetting problem, *Comput. Fluids* 141 (2015) 223–234, <https://doi.org/10.1016/j.compfluid.2016.04.011>.
- [10] S.S. Jain, A. Mani, P. Moin, A conservative diffuse-interface method for compressible two-phase flows, *J. Comput. Phys.* 418 (2020) 109606.
- [11] Z. Huang, G. Lin, A.M. Ardekani, Implementing contact angle boundary conditions for second-order phase-field models of wall-bounded multiphase flows, *arXiv preprint, arXiv:2103.07839*, 2021.
- [12] S. Mirjalili, A. Mani, Consistent, energy-conserving momentum transport for simulations of two-phase flows using the phase field equations, *J. Comput. Phys.* 426 (2021) 109918.
- [13] D.M. Anderson, G.B. McFadden, A.A. Wheeler, Diffuse-interface methods in fluid mechanics, *Annu. Rev. Fluid Mech.* 30 (1998) 139–165, <https://doi.org/10.1146/annurev.fluid.30.1.139>.
- [14] D. Jacqmin, An energy approach to the continuum surface tension method, in: 34th AIAA Aerospace Sciences Meeting & Exhibit, Vol. AIAA96, American Institute of Aeronautics and Astronautics, Reno, Nevada, USA, 1996, p. 0858.
- [15] D. Jacqmin, Contact-line dynamics of a diffuse fluid interface, *J. Fluid Mech.* 402 (2000) 57–88, <https://doi.org/10.1017/S0022112099006874>.
- [16] J. Volker, *Finite Element Methods for Incompressible Flow Problems*, Springer Series in Computational Mathematics Book, vol. 51, Springer, 2016.
- [17] Z. Guo, P. Lin, J. Lowengrub, S.M. Wise, Mass conservative and energy stable finite difference methods for the quasi-incompressible Navier–Stokes–Cahn–Hilliard system: primitive variable and projection-type schemes, *Comput. Methods Appl. Mech. Eng.* 326 (2017) 144–174, <https://doi.org/10.1016/j.cma.2017.08.011>.
- [18] M. Shokrpour Roudbari, G. Şimşek, E.H. van Brummelen, K.G. van der Zee, Diffuse-interface two-phase flow models with different densities: a new quasi-incompressible form and a linear energy-stable method, *Math. Models Methods Appl. Sci.* 28 (2018) 733–770, <https://doi.org/10.1142/S0218202518500197>.
- [19] M.A. Khanwale, K. Saurabh, M. Fernando, V.M. Calo, J.A. Rossmann, H. Sundar, B. Ganapathysubramanian, A fully-coupled framework for solving Cahn–Hilliard Navier–Stokes equations: second-order, energy-stable numerical methods on adaptive octree based meshes, *Comput. Phys. Commun.* 280 (2022) 108501, <https://doi.org/10.1016/j.cpc.2022.108501>, <https://linkinghub.elsevier.com/retrieve/pii/S001046552200220X>.
- [20] M. Ishii, M. Fernando, K. Saurabh, B. Khara, B. Ganapathysubramanian, H. Sundar, Solving PDEs in space-time: 4D tree-based adaptivity, mesh-free and matrix-free approaches, in: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2019, pp. 1–61.
- [21] K. Saurabh, M. Ishii, M. Fernando, B. Gao, K. Tan, M.-C. Hsu, A. Krishnamurthy, H. Sundar, B. Ganapathysubramanian, Scalable adaptive pde solvers in arbitrary domains, in: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2021, pp. 1–15.
- [22] V.E. Badalassi, H.D. Cenicer, S. Banerjee, Computation of multiphase systems with phase field models, *J. Comput. Phys.* 190 (2003) 371–397.
- [23] J. Kim, K. Kang, J. Lowengrub, Conservative multigrid methods for Cahn–Hilliard fluids, *J. Comput. Phys.* 193 (2004) 511–543.
- [24] J. Shen, X. Yang, A phase-field model and its numerical approximation for two-phase incompressible flows with different densities and viscosities, *SIAM J. Sci. Comput.* 32 (2010) 1159–1179, <https://doi.org/10.1137/09075860X>.
- [25] J. Shen, J. Xu, J. Yang, The scalar auxiliary variable (SAV) approach for gradient flows, *J. Comput. Phys.* 353 (2018) 407–416.
- [26] P. Yue, J.J. Feng, C. Liu, J. Shen, A diffuse-interface method for simulating two-phase flows of complex fluids, *J. Fluid Mech.* 515 (2004) 293–317.
- [27] H. Abels, H. Garcke, G. Grün, Thermodynamically consistent, frame indifferent diffuse interface models for incompressible two-phase flows with different densities, *Math. Models Methods Appl. Sci.* 22 (2012) 1150013.
- [28] J. Shen, X. Yang, Energy stable schemes for Cahn–Hilliard phase-field model of two-phase incompressible flows, *Chin. Ann. Math., Ser. B* 31 (2010) 743–758, <https://doi.org/10.1007/s11401-010-0599-y>.
- [29] S. Dong, On imposing dynamic contact-angle boundary conditions for wall-bounded liquid–gas flows, *Comput. Methods Appl. Mech. Eng.* 247–248 (2012) 179–200, <https://doi.org/10.1016/j.cma.2012.07.023>.
- [30] J. Shen, X. Yang, Decoupled, energy stable schemes for phase-field models of two-phase incompressible flows, *SIAM J. Numer. Anal.* 53 (2015) 279–296, <https://doi.org/10.1137/140971154>.
- [31] Y. Chen, J. Shen, Efficient, adaptive energy stable schemes for the incompressible Cahn–Hilliard Navier–Stokes phase-field models, *J. Comput. Phys.* 308 (2016) 40–56, <https://doi.org/10.1016/j.jcp.2015.12.006>.
- [32] S. Dong, Wall-bounded multiphase flows of N immiscible incompressible fluids: consistency and contact-angle boundary condition, <http://arxiv.org/abs/1610.10090>, arXiv:1610.10090, 2016.
- [33] G. Zhu, H. Chen, J. Yao, S. Sun, Efficient energy-stable schemes for the hydrodynamics coupled phase-field model, *Appl. Math. Model.* 70 (2019) 82–108.

- [34] D. Han, X. Wang, A second order in time, uniquely solvable, unconditionally stable numerical scheme for Cahn–Hilliard–Navier–Stokes equation, *J. Comput. Phys.* 290 (2015) 139–156, <https://doi.org/10.1016/j.jcp.2015.02.046>, arXiv:1407.7048.
- [35] J. Guermont, P. Mineev, J. Shen, An overview of projection methods for incompressible flows, *Comput. Methods Appl. Mech. Eng.* 195 (2006) 6011–6045, <https://doi.org/10.1016/j.cma.2005.10.010>, <https://linkinghub.elsevier.com/retrieve/pii/S0045782505004640>.
- [36] S. Xu, B. Gao, A. Lofquist, M. Fernando, M.-C. Hsu, H. Sundar, B. Ganapathysubramanian, An octree-based immersogeometric approach for modeling inertial migration of particles in channels, *Comput. Fluids* 214 (2021) 104764, <https://doi.org/10.1016/j.compfluid.2020.104764>, <https://linkinghub.elsevier.com/retrieve/pii/S0045793020303340>.
- [37] T.J. Hughes, L. Mazzei, K.E. Jansen, Large eddy simulation and the variational multiscale method, *Comput. Vis. Sci.* 3 (2000) 47–59.
- [38] Y. Bazilevs, V. Calo, J. Cottrell, T. Hughes, A. Reali, G. Scovazzi, Variational multiscale residual-based turbulence modeling for large eddy simulation of incompressible flows, *Comput. Methods Appl. Mech. Eng.* 197 (2007) 173–201.
- [39] N. Ahmed, T. Chacón Rebollo, V. John, S. Rubino, A review of variational multiscale methods for the simulation of turbulent incompressible flows, *Arch. Comput. Methods Eng.* 24 (2017) 115–164, <https://doi.org/10.1007/s11831-015-9161-0>.
- [40] T. Coupez, E. Hachem, Solution of high-Reynolds incompressible flow with stabilized finite element and adaptive anisotropic meshing, *Comput. Methods Appl. Mech. Eng.* 267 (2013) 65–85, <https://doi.org/10.1016/j.cma.2013.08.004>.
- [41] E. Hachem, S. Feghali, R. Codina, T. Coupez, Anisotropic adaptive meshing and monolithic Variational Multiscale method for fluid-structure interaction, *Comput. Struct.* 122 (2013) 88–100, <https://doi.org/10.1016/j.compstruc.2012.12.004>.
- [42] E. Hachem, M. Khaliloufi, J. Bruchon, R. Valette, Y. Mesri, Unified adaptive Variational MultiScale method for two phase compressible-incompressible flows, *Comput. Methods Appl. Mech. Eng.* 308 (2016) 238–255, <https://doi.org/10.1016/j.cma.2016.05.022>.
- [43] H. Bao, J. Bielak, O. Ghattas, L.F. Kallivokas, D.R. O'Hallaron, J.R. Shewchuk, J. Xu, Large-scale simulation of elastic wave propagation in heterogeneous media on parallel computers, *Comput. Methods Appl. Mech. Eng.* 152 (1998) 85–102.
- [44] M. Fernando, D. Neilsen, H. Lim, E. Hirschmann, H. Sundar, Massively parallel simulations of binary black hole intermediate-mass-ratio inspirals, *SIAM J. Sci. Comput.* 41 (2019) C97–C138, <https://doi.org/10.1137/18M1196972>.
- [45] H. Sundar, R.S. Sampath, G. Biros, Bottom-up construction and 2:1 balance refinement of linear octrees in parallel, *SIAM J. Sci. Comput.* 30 (2008) 2675–2708.
- [46] C. Burstedde, L.C. Wilcox, O. Ghattas, p4est: scalable algorithms for parallel adaptive mesh refinement on forests of octrees, *SIAM J. Sci. Comput.* 33 (2011) 1103–1133, <https://doi.org/10.1137/100791634>.
- [47] F. Magaletti, F. Picano, M. Chinappi, L. Marino, C.M. Casciola, The sharp-interface limit of the Cahn–Hilliard/Navier–Stokes model for binary fluids, *J. Fluid Mech.* 714 (2013) 95–126, <https://doi.org/10.1017/jfm.2012.461>.
- [48] M.A. Khanwale, A.D. Lofquist, H. Sundar, J.A. Rossmann, B. Ganapathysubramanian, Simulating two-phase flows with thermodynamically consistent energy stable Cahn–Hilliard Navier–Stokes equations on parallel adaptive octree based meshes, *J. Comput. Phys.* 419 (2020) 109674, <https://doi.org/10.1016/j.jcp.2020.109674>.
- [49] T.J. Hughes, G. Scovazzi, L.P. Franca, Multiscale and stabilized methods, in: *Encyclopedia of Computational Mechanics*, second edition, 2018, pp. 1–64.
- [50] T. Tezduyar, S. Mittal, S. Ray, R. Shih, Incompressible flow computations with stabilized bilinear and linear equal-order-interpolation velocity-pressure elements, *Comput. Methods Appl. Mech. Eng.* 95 (1992) 221–242.
- [51] U. Ascher, S. Ruuth, B. Wetton, Implicit-explicit methods for time-dependent partial differential equations, *SIAM J. Numer. Anal.* 32 (1995) 797–823.
- [52] T.J. Hughes, G. Sangalli, Variational multiscale analysis: the fine-scale Green's function, projection, optimization, localization, and stabilized methods, *SIAM J. Numer. Anal.* 45 (2007) 539–557.
- [53] M. Fernando, D. Duplyakin, H. Sundar, Machine and application aware partitioning for adaptive mesh refinement applications, in: *Proceedings of the 26th International Symposium on High-Performance Parallel and Distributed Computing, HPDC '17*, ACM, New York, NY, USA, 2017, pp. 231–242.
- [54] M. Bern, D. Eppstein, S.-H. Teng, Parallel construction of quadrees and quality triangulations, *Int. J. Comput. Geom. Appl.* 9 (1999) 517–532.
- [55] H. Sundar, R.S. Sampath, S.S. Adavani, C. Davatzikos, G. Biros, Low-constant parallel algorithms for finite element simulations using linear octrees, in: *SC'07: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis, ACM/IEEE*, 2007, pp. 1–12.
- [56] S. Balay, S. Abhyankar, M.F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W.D. Gropp, D. Karpeyev, D. Kaushik, M.G. Knepley, D.A. May, L.C. McInnes, R.T. Mills, T. Munson, K. Rupp, P. Sanan, B.F. Smith, S. Zampini, H. Zhang, H. Zhang, PETSc Web page, <https://www.mcs.anl.gov/petsc>, 2019.
- [57] S. Balay, W.D. Gropp, L.C. McInnes, B.F. Smith, Efficient management of parallelism in object oriented numerical software libraries, in: E. Arge, A.M. Bruaset, H.P. Langtangen (Eds.), *Modern Software Tools in Scientific Computing*, Birkhäuser Press, 1997, pp. 163–202.
- [58] S. Balay, S. Abhyankar, M.F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W.D. Gropp, D. Karpeyev, D. Kaushik, M.G. Knepley, D.A. May, L.C. McInnes, R.T. Mills, T. Munson, K. Rupp, P. Sanan, B.F. Smith, S. Zampini, H. Zhang, H. Zhang, PETSc Users Manual, Technical Report ANL-95/11 - Revision 3.11, Argonne National Laboratory, 2019.
- [59] S.-R. Hysing, S. Turek, D. Kuzmin, N. Parolini, E. Burman, S. Ganesan, L. Tobiska, Quantitative benchmark computations of two-dimensional bubble dynamics, *Int. J. Numer. Methods Fluids* 60 (2009) 1259–1288.
- [60] S. Aland, A. Voigt, Benchmark computations of diffuse interface models for two-dimensional bubble dynamics, *Int. J. Numer. Methods Fluids* 69 (2012) 747–761, <https://doi.org/10.1002/fld.2611>.
- [61] H.Z. Yuan, Z. Chen, C. Shu, Y. Wang, X.D. Niu, S. Shu, A free energy-based surface tension force model for simulation of multiphase flows by level-set method, *J. Comput. Phys.* 345 (2017) 404–426, <https://doi.org/10.1016/j.jcp.2017.05.020>.
- [62] G. Tryggvason, S.O. Unverdi, Computations of three-dimensional Rayleigh–Taylor instability, *Phys. Fluids A, Fluid Dyn.* 2 (1990) 656–659.
- [63] X. Li, B. Jin, J. Glimm, Numerical study for the three-dimensional Rayleigh–Taylor instability through the TVD/AC scheme and parallel computation, *J. Comput. Phys.* 126 (1996) 343–355.
- [64] J.-L. Guermont, L. Quartapelle, A projection FEM for variable density incompressible flows, *J. Comput. Phys.* 165 (2000) 167–188.
- [65] J.T. Waddell, C.E. Niederhaus, J.W. Jacobs, Experimental study of Rayleigh–Taylor instability: low Atwood number liquid systems with single-mode perturbations, *Phys. Fluids* 13 (2001) 1263–1273, <https://doi.org/10.1063/1.1359762>.
- [66] J. Hunt, A. Wray, P. Moin, Eddies, stream, and convergence zones in turbulent flows, *Center for turbulence research report CTR-S88*, pp. 193–208, 1988.
- [67] R.H. Kraichnan, Inertial ranges in two-dimensional turbulence, *Phys. Fluids* (1958–1988) 10 (1967) 1417–1423, <https://doi.org/10.1063/1.1762301>.
- [68] H. Liang, Q.X. Li, B.C. Shi, Z.H. Chai, Lattice Boltzmann simulation of three-dimensional Rayleigh–Taylor instability, *Phys. Rev. E* 93 (2016) 033113, <https://doi.org/10.1103/PhysRevE.93.033113>.
- [69] K. Saurabh, B. Gao, M. Fernando, S. Xu, M.A. Khanwale, B. Khara, M.-C. Hsu, A. Krishnamurthy, H. Sundar, B. Ganapathysubramanian, Industrial scale large eddy simulations with adaptive octree meshes using immersogeometric analysis, *Comput. Math. Appl.* 97 (2021) 28–44.
- [70] H. Sundar, G. Biros, C. Burstedde, J. Rudi, O. Ghattas, G. Stadler, Parallel geometric-algebraic multigrid on unstructured forests of octrees, in: *SC'12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, IEEE*, 2012, pp. 1–11.