

Highly parallelisable simulations of time-dependent viscoplastic fluid flow with structured adaptive mesh refinement

Knut Sverdrup, Nikolaos Nikiforakis, and Ann Almgren

Citation: *Physics of Fluids* **30**, 093102 (2018); doi: 10.1063/1.5049202

View online: <https://doi.org/10.1063/1.5049202>

View Table of Contents: <http://aip.scitation.org/toc/phf/30/9>

Published by the [American Institute of Physics](#)

Articles you may be interested in

[Torque driven ferromagnetic swimmers](#)

Physics of Fluids **30**, 092001 (2018); 10.1063/1.5046360

[Study of the thrust–drag balance with a swimming robotic fish](#)

Physics of Fluids **30**, 091901 (2018); 10.1063/1.5043137

[Axisymmetric rim instability of water droplet impact on a super-hydrophobic surface](#)

Physics of Fluids **30**, 094101 (2018); 10.1063/1.5039558

[Analysis of flame curvature evolution in a turbulent premixed bluff body burner](#)

Physics of Fluids **30**, 095101 (2018); 10.1063/1.5044525

[The hydrodynamic behavior of a squirmer swimming in power-law fluid](#)

Physics of Fluids **30**, 083301 (2018); 10.1063/1.5045701

[Modeling polymer extrusion with varying die gap using Arbitrary Lagrangian Eulerian \(ALE\) method](#)

Physics of Fluids **30**, 093103 (2018); 10.1063/1.5045739

PHYSICS TODAY

WHITEPAPERS

ADVANCED LIGHT CURE ADHESIVES

Take a closer look at what these environmentally friendly adhesive systems can do

READ NOW

PRESENTED BY
 MASTERBOND
ADHESIVES | SEALANTS | COATINGS

Highly parallelisable simulations of time-dependent viscoplastic fluid flow with structured adaptive mesh refinement

Knut Sverdrup,^{1,a)} Nikolaos Nikiforakis,¹ and Ann Almgren²

¹*Cavendish Laboratory, University of Cambridge, Cambridge CB3 0HE, United Kingdom*

²*Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

(Received 20 July 2018; accepted 23 August 2018; published online 11 September 2018)

We present the extension of an efficient and highly parallelisable framework for incompressible fluid flow simulations to viscoplastic fluids. The system is governed by incompressible conservation of mass, the Cauchy momentum equation, and a generalised Newtonian constitutive law. In order to simulate a wide range of viscoplastic fluids, we employ the Herschel-Bulkley model for yield-stress fluids with nonlinear stress-strain dependency above the yield limit. We utilise Papanastasiou regularisation in our algorithm to deal with the singularity in apparent viscosity. The resulting system of partial differential equations is solved using the IAMR (Incompressible Adaptive Mesh Refinement) code, which uses second-order Godunov methodology for the advective terms and semi-implicit diffusion in the context of an approximate projection method to solve adaptively refined meshes. By augmenting the IAMR code with the ability to simulate regularised Herschel-Bulkley fluids, we obtain efficient numerical software for time-dependent viscoplastic flow in three dimensions, which can be used to investigate systems not considered previously due to computational expense. We validate results from simulations using this new capability against previously published data for Bingham plastics and power-law fluids in the two-dimensional lid-driven cavity. In doing so, we expand the range of Bingham and Reynolds numbers which have been considered in the benchmark tests. Moreover, extensions to time-dependent flow of Herschel-Bulkley fluids and three spatial dimensions offer new insights into the flow of viscoplastic fluids in this test case, and we provide missing benchmark results for these extensions. *Published by AIP Publishing.* <https://doi.org/10.1063/1.5049202>

I. INTRODUCTION

Viscoplastic fluids are non-Newtonian fluids which are characterised by a minimum induced stress necessary for flow to occur. For this reason, they are also commonly referred to as yield-stress fluids. When the imposed stress does not exceed the threshold value, the material is modelled as a rigid solid. In regions where the yield stress is exceeded, however, the material flows like a fluid. The ability of the material to support a stress under certain circumstances gives rise to phenomena such as non-flat surfaces at rest under gravity and the coexistence of yielded (flowing) and unyielded (rigid) regions within the fluid. The former can be demonstrated by distorting the surface of mayonnaise in a jar: gravity alone is not strong enough to surpass the yield stress and the surface remains in its distorted state. In addition to being fundamentally interesting from the perspectives of rheology, fluid mechanics, and mathematical modeling, yield stress fluids occur naturally and are paramount to the success of animals such as mudskippers¹ and snails.² Their importance in industries ranging from medicine^{3–6} to oil and gas exploration^{7–9} has led to extensive research contributions in the field.

Just as for Newtonian fluids, the pursuit of knowledge about viscoplastic fluids has relied heavily on computational methods in the last fifty years. Compared to the Newtonian

case, however, the numerical simulations are much more demanding in terms of processing time. This is largely due to the existence of a yield stress, since this results in a singularity at zero strain rate for the apparent viscosity of the fluid. Traditional methods for computing the solution are rendered useless for this case since infinite viscosities cannot be represented in the unyielded regions. There are now two main branches of algorithms designed to deal with this problem. The first utilises mathematical regularisation to approximate the viscosity function in the low-strain limit. In doing so, the rigid body approximation is effectively replaced by a fluid with very high viscosity in the unyielded regions. A regularisation parameter controls how close this approximation is to the idealised viscoplastic case. On the other hand, the problem can be reformulated in the framework of non-smooth optimisation theory and solved using augmented Lagrangians. Such methods can solve the unregularised problem without introducing any approximations but generally require more computational resources to find the solution.¹⁰

The regularisation approach was first explored by Bercovier and Engelman in 1980,¹¹ who utilised a simple yet efficient work-around by adding a small constant to the computed strain-rate so that even in the zero-strain limit the viscosity would remain finite. An alternative method proposed by Tanner and Milthorpe a few years later was the bi-viscosity model,¹² in which the viscoplastic fluid is characterised by a separate, large viscosity when the strain-rate is below a given threshold (see, e.g., Ref. 13). An exponential

^{a)}Electronic mail: ksk38@cam.ac.uk

regularisation factor was then introduced by Papanastasiou in 1987,¹⁴ and this method of regularisation is still widely used in modern codes which deal with viscoplasticity through regularisation. Important investigations based on Papanastasiou regularisation include those of Mitsoulis *et al.*^{15–18} and Syrakos *et al.*^{19–21}

Augmented Lagrangian methods were first applied to the viscoplastic flow problem in 1983 by Fortin and Glowinski,²² but the variational formulation on which it relies was derived by Duvaut and Lions in 1976, who studied existence, uniqueness, and regularity of solutions to the problem.²³ The augmented Lagrangian algorithm itself is due to the work of Hestenes.²⁴ Although the simulation of viscoplastic flow using this optimisation technique constitutes an important milestone in the history of its numerical treatments, the regularisation approach was much more popular due to the large discrepancy in computational resource requirements. Advances in convex optimisation over the next few decades led to the work of Saramito and Roquet,²⁵ where significant improvements were achieved in terms of convergence rates, and hence, computational efficiency. Other contributions²⁶ further confirmed the potential, and in recent years, state-of-the-art algorithms are actively being developed, notably by Treskatis *et al.*,²⁷ Saramito,²⁸ Bleyer,²⁹ and Dimakopoulos *et al.*³⁰ Due to the ability to solve the unregularised viscoplasticity problem, these methods have become increasingly popular since algorithmic advances have led to significant speed-up of their runtime. It is important to note, however, that these methods are still more costly than regularised approaches, so although they allow computation of exact locations of yield surfaces, they are unnecessarily expensive when this is less important, and a general understanding of the flow field is desirable. An example is within cement displacement complexities for engineering purposes,³¹ where the desirable insight is pressure distributions and the cement displacement efficiency.³²

Numerical treatment of viscoplastic flow problems has significantly improved over the last decades. Researchers interested in analysing such flows are, however, still limited by the computational cost associated with their solution. Notably, most research published in the field considers only two spatial dimensions and only steady-state solutions. Although open-source libraries such as FEniCS³³ and OpenFOAM³⁴ can be used to simulate regularised viscoplastic fluids in three dimensions, our contribution aims to provide a massively parallelisable tool for contemporary supercomputer architectures, with capabilities for structured AMR (Adaptive Mesh Refinement). In order to achieve this, we begin with the Incompressible Adaptive Mesh Refinement (IAMR)³⁵ code that uses a second-order accurate, approximate projection method to solve the incompressible Navier-Stokes equations. The code is built on the AMReX software framework for developing parallel block-structured AMR applications. We have implemented the Papanastasiou regularised Herschel-Bulkley (HB) model for the apparent viscosity function in IAMR so that it can be used to simulate generalised Newtonian fluids. As such, we are able to take advantage of modern supercomputer architectures in our simulation of viscoplastic fluids. For the first time, results are presented for the transient, three-dimensional lid-driven cavity problem for viscoplastic fluids. Notably, the

evolution of the three-dimensional yield surface from rest to steady-state is tracked, before the lid velocity is set to zero, allowing cessation of the viscoplastic fluid.

For further reading on developments in viscoplastic fluids, we refer the reader to the review papers by Barnes³⁶ and Balmforth *et al.*³⁷ Additionally, two excellent review papers on advances in numerical simulations of these fluids appeared in a recent special edition of *Rheologica Acta*.^{10,18} In Sec. II, we will introduce the governing partial differential equations (PDEs) and discuss relevant fluid rheologies. Section III is devoted to the numerical algorithm employed to simulate the fluid flow, including the regularisation strategy. Thorough validation is performed in Sec. IV, before we evaluate the code for Herschel-Bulkley fluids and three dimensions in Sec. V. Section VI concludes the article.

II. MATHEMATICAL FORMULATION

Our domain $\Omega \subset \mathbb{R}^d$ is in either two or three dimensions, i.e., $d \in \{2, 3\}$. For rank-2 tensors, we prescribe the scaled Frobenius norm

$$|\mathbf{A}| = \sqrt{\frac{1}{2} \text{tr}(\mathbf{A} \mathbf{A}^T)}, \quad (1)$$

as is customary in viscoplastic fluid mechanics. We take the gradient of a vector \mathbf{u} as the tensor with components $(\nabla \mathbf{u})_{ij} = \partial u_j / \partial x_i$, and the symmetric part of this gradient is given by $\mathcal{D}\mathbf{u} = \frac{1}{2}(\nabla + \nabla^T)\mathbf{u}$. The divergence of a tensor field is defined such that $(\nabla \cdot \mathbf{A})_j = \sum_{i=1}^d \partial A_{ij} / \partial x_i$. Variables are functions of position $\mathbf{x} \in \Omega$ and time $t \geq 0$.

A. Governing partial differential equations

We denote by $\rho \in \mathbb{R}$ the material density. The velocity field is introduced as $\mathbf{u}(\mathbf{x}, t) \in \mathbb{R}^d$, with components u , v , and w . The Cauchy stress tensor σ is defined as the sum of isotropic and deviatoric parts, $\sigma = -p\mathbf{I} + \tau$. Here, the pressure $p(\mathbf{x}, t) \in \mathbb{R}$ is multiplied by the identity tensor, while the deviatoric part of the stress tensor is denoted as $\tau(\mathbf{x}, t) \in \mathbb{R}_{\text{sym}}^{d \times d}$. The fluid motion is then governed by incompressible mass conservation and momentum balance through the Cauchy equation

$$\nabla \cdot \mathbf{u} = 0, \quad (2a)$$

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = \frac{1}{\rho}(-\nabla p + \nabla \cdot \tau + \mathbf{f}). \quad (2b)$$

Here, we have introduced \mathbf{f} to describe external body forces such as gravity acting on the fluid. Note that the mass conservation equation is simplified to (2a) due to incompressibility, i.e., constant density within a fluid parcel.

B. Rheology

We shall in the following restrict ourselves to relatively simple equations of state, of the form $\tau = \tau(\dot{\gamma})$, where the stress response is solely dependent on the rate-of-strain tensor $\dot{\gamma} = \mathcal{D}\mathbf{u} \in \mathbb{R}_{\text{sym}}^{n \times n}$. Although the Herschel-Bulkley model captures all aspects of the fluids considered, references for validation are only available for power-law and Bingham fluids. For this reason, we introduce these simpler rheological models first and illustrate how they can be combined.




FIG. 1. Herschel-Bulkley fluid: stress magnitude (left) and apparent viscosity (right) as functions of the magnitude of the rate-of-strain tensor $\dot{\gamma}$.

Newtonian flow is characterised by a dynamic coefficient of viscosity $\mu > 0$ which is independent of the strain, yielding a linear relationship between rate-of-strain and stress in the rheological equation

$$\tau = 2\mu\dot{\gamma}. \quad (3)$$

For fluids where the dependency of the stress on the rate-of-strain tensor is nonlinear, the apparent viscosity is a useful concept when considering rheological responses to strain. It is a generalisation of the constant viscosity for Newtonian flow, where we allow the viscosity to be a function of the magnitude of the rate-of-strain tensor. Denoting the apparent viscosity by η , we thus have

$$\eta = \frac{|\tau|}{2|\dot{\gamma}|}. \quad (4)$$

Many fluids are accurately modelled by a non-Newtonian behavior that captures shear-dependency through a smooth increase or decrease in apparent viscosity. Such fluids include pseudoplastics (shear-thinning, $\partial\eta/\partial|\dot{\gamma}| < 0$) and dilatants (shear-thickening, $\partial\eta/\partial|\dot{\gamma}| > 0$). A model which captures this behavior is the power-law fluid, with rheological equation

$$\tau = 2^n \mu |\dot{\gamma}|^{n-1} \dot{\gamma} \quad (5)$$

and apparent viscosity

$$\eta = 2^{n-1} \mu |\dot{\gamma}|^{n-1}. \quad (6)$$

A specific power-law fluid is characterised by its flow behavior index $n > 0$ (and μ , which is usually referred to as the consistency for power-law fluids and has units Pa s^n). From (5), it is immediately clear that the Newtonian case with $n = 1$ separates pseudoplastics ($n < 1$) from dilatants ($n > 1$).

Viscoplastic fluids have a stress threshold $\tau_0 > 0$ (the yield stress), below which they do not flow. We note that it is possible to describe elastic deformation for materials which do not flow, but we shall be considering constitutive equations which assume a rigid body approximation, i.e., zero strain rate below the yield stress.

The simplest type of viscoplastic fluid is the Bingham fluid,^{38,39} characterised by zero strain rate below the yield stress. In the yielded region, however, the stress depends linearly on the rate-of-strain magnitude, just like a Newtonian fluid. The stress-strain curve thus intercepts the $|\tau|$ -axis at the point $(0, \tau_0)$. As such, the Bingham fluid is a generalised Newtonian, with rheological equation

$$\begin{cases} \dot{\gamma} = 0 & \text{if } |\tau| \leq \tau_0 \\ \tau = 2\mu\dot{\gamma} + \frac{\tau_0}{|\dot{\gamma}|}\dot{\gamma} & \text{if } |\tau| > \tau_0 \end{cases} \quad (7)$$

and apparent viscosity

$$\eta = \mu + \frac{\tau_0}{2|\dot{\gamma}|}. \quad (8)$$

For Bingham plastics, we refer to μ as the plastic viscosity and note that η has a singularity for $\dot{\gamma} = 0$, as expected.

In many applications, it is desirable to capture both the yield stress of viscoplastic fluids and the power-law dependency occurring once the fluid starts flowing. A widely used rheological model for such fluids is due to the work of Herschel and Bulkley.⁴⁰ The Herschel-Bulkley fluid facilitates a very general description of non-Newtonian fluids, as it is a yield stress fluid with a nonlinear stress-strain dependency in the yielded region. As such, it can be thought of as a hybrid between Bingham plastics and power law fluids. The constitutive equation is

$$\begin{cases} \dot{\gamma} = 0 & \text{if } |\tau| \leq \tau_0 \\ \tau = 2^n \mu |\dot{\gamma}|^{n-1} \dot{\gamma} + \frac{\tau_0}{|\dot{\gamma}|} \dot{\gamma} & \text{if } |\tau| > \tau_0 \end{cases} \quad (9)$$

while the apparent viscosity is

$$\eta = 2^{n-1} \mu |\dot{\gamma}|^{n-1} + \frac{\tau_0}{2|\dot{\gamma}|}. \quad (10)$$

Bingham plastics are a special case of Herschel-Bulkley fluids in exactly the same manner as Newtonian fluids follow a specific power-law constitutive equation. Plots of the rheological characteristic for Herschel-Bulkley fluids are shown in Fig. 1, for various values of n . Note that the Bingham fluid is recovered for $n = 1$.

III. NUMERICAL ALGORITHM

A. Incompressible flow solver

An approximate projection method for solving the variable density incompressible Navier-Stokes equations on an adaptive mesh hierarchy has been implemented in the IAMR code. The algorithm was described for the constant viscosity case by Almgren *et al.*,³⁵ with extensions to low Mach number reacting flows with temperature-dependent viscosity provided by Pember *et al.*⁴¹ and Day and Bell,⁴² among others. In addition to solving the Navier-Stokes equations for velocity and pressure, the IAMR code allows for the (conservative or passive) advection of any number of scalar quantities. The implementation is such that the code can be run on architectures from single-core laptops through to massively parallel supercomputers.

In the approximate projection method as implemented in IAMR, an advection-diffusion step is used to advance the velocity in time; the solution is then (approximately) projected onto the space of divergence-free fields. The motivation for an approximate rather than exact projection is covered in detail in the literature.^{43–45}

Velocity is defined at cell centres at integer time steps and is denoted by $\mathbf{u}_{i,j,k}^m$. The pressure, on the other hand, is specified at cell corners and is staggered in time so that it is denoted as $p_{i+\frac{1}{2},j+\frac{1}{2},k+\frac{1}{2}}^{m+\frac{1}{2}}$. When no subscript is given (e.g., \mathbf{u}^m), we address the distribution of \mathbf{u} throughout the computational domain at time t^m . Considerations such as initialisation and boundary treatments are ignored at present, and we set external forces equal to zero. Instead, we focus on the steps required in order to evolve the solution a single time step Δt .

The first part of the advection-diffusion step is to extrapolate normal velocity components to each cell face at time $t^{m+\frac{1}{2}}$. This is done through a second-order Taylor series expansion. Consider the cell centre (i, j, k) , and suppose we want to extrapolate the first velocity component u to the face $(i+\frac{1}{2}, j, k)$. Then the truncated Taylor series is

$$u_{i+\frac{1}{2},j,k}^{m+\frac{1}{2}} \approx u_{i,j,k}^m + \frac{\Delta x}{2} \left(\frac{\partial u}{\partial x} \right)_{i,j,k}^m + \frac{\Delta t}{2} \left(\frac{\partial u}{\partial t} \right)_{i,j,k}^m. \quad (11)$$

Taking the first component of (2b), we have

$$\begin{aligned} \left(\frac{\partial u}{\partial t} \right)_{i,j,k}^m &= - \left(u \frac{\partial u}{\partial x} \right)_{i,j,k}^m - \left(v \frac{\partial u}{\partial y} \right)_{i,j,k}^m - \left(w \frac{\partial u}{\partial z} \right)_{i,j,k}^m \\ &\quad - \frac{1}{\rho} \left(\frac{\partial p}{\partial x} \right)_{i,j,k}^{m-\frac{1}{2}} + \frac{1}{\rho} \left(\frac{\partial \tau_{11}}{\partial x} + \frac{\partial \tau_{21}}{\partial y} + \frac{\partial \tau_{31}}{\partial z} \right)_{i,j,k}^m, \end{aligned} \quad (12)$$

which can be substituted into the last term of the Taylor expansion to avoid any dependence on temporal derivatives. In order to compute these extrapolated values, however, it is necessary to calculate discrete approximations of the spatial derivatives. This is done separately for each part. The first derivative of normal velocity is evaluated using a monotonicity-limited fourth-order slope approximation S as introduced by Colella in 1985.⁴⁶ Since the pressure field is defined at grid nodes, cell-centred approximations of the pressure gradient \mathbf{G} are computed by finite differences. Finally, transverse derivative terms such as $v\partial u/\partial y$ are evaluated by a separate extrapolation and subsequent upwinding procedure,

$$\left(v \frac{\partial u}{\partial y} \right)_{i,j,k}^m \approx \left(\widehat{v \frac{\partial u}{\partial y}} \right)_{i,j,k}^m. \quad (13)$$

With these discrete representations in place, the intermediate normal velocity at face $i+\frac{1}{2}$ can be approximated based on the cell to its left (i, j, k) as

$$\begin{aligned} u_{i+\frac{1}{2},j,k}^{m+\frac{1}{2},L} &= u_{i,j,k}^m + \left(\frac{\Delta x}{2} - u_{i,j,k}^m \frac{\Delta t}{2} \right) (Su)_{i,j,k}^m \\ &\quad - \frac{\Delta t}{2} \left(\widehat{v \frac{\partial u}{\partial y}} \right)_{i,j,k}^m - \frac{\Delta t}{2} \left(\widehat{w \frac{\partial u}{\partial z}} \right)_{i,j,k}^m \\ &\quad - \frac{\Delta t}{2\rho} \left(\left(-(\mathbf{G}_x p)_{i,j,k}^{m-\frac{1}{2}} + (\mathbf{D}_x \tau)_{i,j,k}^m \right) \right). \end{aligned} \quad (14a)$$

Similarly, an approximation based on the state to its right $(i+1, j, k)$ is

$$\begin{aligned} u_{i+\frac{1}{2},j,k}^{m+\frac{1}{2},R} &= u_{i+1,j,k}^m + \left(\frac{\Delta x}{2} + u_{i+1,j,k}^m \frac{\Delta t}{2} \right) (Su)_{i+1,j,k}^m \\ &\quad - \frac{\Delta t}{2} \left(\widehat{v \frac{\partial u}{\partial y}} \right)_{i+1,j,k}^m - \frac{\Delta t}{2} \left(\widehat{w \frac{\partial u}{\partial z}} \right)_{i+1,j,k}^m \\ &\quad - \frac{\Delta t}{2\rho} \left(-(\mathbf{G}_x p)_{i+1,j,k}^{m-\frac{1}{2}} + (\mathbf{D}_x \tau)_{i+1,j,k}^m \right). \end{aligned} \quad (14b)$$

These extrapolations need to be computed for tangential velocity components as well. At each face, we then have two extrapolated approximations for $\mathbf{u}^{m+\frac{1}{2}}$, one from each cell adjoining at the face. In order to pick the most accurate of these, we construct a face-based advection velocity \mathbf{u}^{adv} . The final, upwinded, extrapolated approximation at each face is denoted as $\tilde{\mathbf{u}}^{m+\frac{1}{2}}$.

With these extrapolated and upwinded approximations in place, (2b) is discretised in time to construct a new-time provisional velocity field, \mathbf{u}^* , without enforcing (2a), i.e., we define \mathbf{u}^* using

$$\begin{aligned} \frac{\mathbf{u}^* - \mathbf{u}^m}{\Delta t} &= - [\mathbf{u}^{\text{MAC}} \cdot \nabla \tilde{\mathbf{u}}]^{m+\frac{1}{2}} + \frac{1}{\rho} \left(-\mathbf{G} p^{m-\frac{1}{2}} + \frac{1}{2} (\nabla \cdot \tau(\mathbf{u}^m)) \right. \\ &\quad \left. + \nabla \cdot \tau(\mathbf{u}^*) + \mathbf{f}^m \right), \end{aligned} \quad (15)$$

where $\mathbf{G} p^{m-\frac{1}{2}}$ is a lagged approximation to the pressure gradient ∇p and the density ρ is a constant. IAMR is designed for more general flows with variable density; however, for the flows considered in this paper, the density is constant in space and time.

The explicit viscous term, $\tau(\mathbf{u}^m)$, is evaluated using only the velocity components at time t^m ; i.e., we define $\dot{\gamma}^m = \dot{\gamma}(\mathbf{u}^m)$ and $\eta^m = \eta(\dot{\gamma}^m)$ and write $\tau(\mathbf{u}^m) = 2\eta^m \dot{\gamma}^m$. The implicit term, $\tau(\mathbf{u}^*)$ as well as \mathbf{u}^* itself, is solved for with the same η^m , through a tensor solve of the form

$$\begin{aligned} \left(\mathbf{u}^* - \frac{\Delta t}{\rho} \nabla \cdot (\eta^m \dot{\gamma}^*) \right) &= \mathbf{u}^m - \Delta t [\mathbf{u}^{\text{MAC}} \cdot \nabla \tilde{\mathbf{u}}]^{m+\frac{1}{2}} \\ &\quad + \frac{\Delta t}{\rho} \left(-\mathbf{G} p^{m-\frac{1}{2}} + \nabla \cdot (\eta^m \dot{\gamma}^m) + \mathbf{f}^m \right). \end{aligned} \quad (16)$$

Note that all velocity components are solved for simultaneously.

The final part of the algorithm is the projection step and subsequent pressure update. The velocity field \mathbf{u}^* does not in general satisfy the divergence constraint as given by (2a). We solve

$$L_\rho \phi = D \left(\frac{1}{\Delta t} \mathbf{u}^* + \frac{1}{\rho} \mathbf{G} p^{m-\frac{1}{2}} \right), \quad (17)$$

where D is a discrete divergence and \mathbf{G} a discrete gradient. L_ρ is a second-order accurate approximation to $\nabla \cdot \frac{1}{\rho} \nabla$. The new-time velocity is then defined by

$$\mathbf{u}^{m+1} = \mathbf{u}^* - \Delta t \frac{1}{\rho} \mathbf{G} \phi \quad (18)$$

and the updated pressure by

$$p^{m+1/2} = \phi. \quad (19)$$

The resulting approximate projection satisfies the divergence constraint to second-order accuracy and ensures stability for the algorithm. Given suitable initial and boundary conditions, this algorithm can be used to evolve the system in time. For time-dependent flows approaching a steady-state solution, we consider the system steady when

$$\max \left| \frac{\mathbf{u}^{m+1} - \mathbf{u}^m}{\mathbf{u}^m} \right| < 10^{-5}. \quad (20)$$

For further details on the algorithm, see the original paper by Almgren *et al.*³⁵

B. Regularisation of viscosity

Equation (9) does not present any problems for unyielded regions ($\dot{\gamma} > 0$), but the apparent viscosity has a singularity when the strain rate approaches zero (and $|\tau| \rightarrow \tau_0$). Computational schemes such as those implemented in IAMR cannot be used in the presence of such singularities. Regularisation deals with the problem by replacing the ill-behaved apparent viscosity with a function that approximates the rheological behavior, but which stays bounded for arbitrarily small $\dot{\gamma}$. This is done by introducing an additional parameter ε to the apparent viscosity, which describes how big the effect of the regularisation is. A large value of ε allows for inexpensive computations even near unyielded flow, while the limit $\varepsilon \rightarrow 0$ recovers the unregularised description. Note that one must be careful to choose a small enough value for ε , in addition to fine enough mesh resolutions and convergence criterion. Otherwise, the mathematical approximation will not hold.

We employ the popular Papanastasiou regularisation,^{14,47} which utilises an exponential relaxation according to

$$\frac{1}{|\dot{\gamma}|} \rightarrow \frac{1 - e^{-|\dot{\gamma}|/\varepsilon}}{|\dot{\gamma}|}. \quad (21)$$

This is a good approximation for $|\dot{\gamma}|/\varepsilon \gg 1$, while in the small strain limit we have

$$\lim_{|\dot{\gamma}| \rightarrow 0} \frac{1 - e^{-|\dot{\gamma}|/\varepsilon}}{|\dot{\gamma}|} = \lim_{|\dot{\gamma}| \rightarrow 0} \frac{1}{|\dot{\gamma}|} \left(1 - \sum_{k=0}^{\infty} \frac{(-|\dot{\gamma}|/\varepsilon)^k}{k!} \right) = \frac{1}{\varepsilon} \quad (22)$$

so that it always remains bounded and recovers the unregularised model in the limit $\varepsilon \rightarrow 0$. Combination of (21) with the apparent viscosity as given in (9) gives the regularised viscosity as

$$\eta = \left(2^{n-1} \mu |\dot{\gamma}|^{n-1} + \frac{\tau_0}{2|\dot{\gamma}|} \right) (1 - e^{-|\dot{\gamma}|/\varepsilon}). \quad (23)$$

The effect of varying the regularisation parameter ε is shown in Fig. 2. In the Papanastasiou-regularised Herschel-Bulkley model, the singularity in apparent viscosity is replaced by the limiting value

$$\lim_{|\dot{\gamma}| \rightarrow 0} \eta = \frac{\tau_0}{2\varepsilon}. \quad (24)$$

C. Adaptive mesh refinement

For a vast range of flow problems, large computational domains are required even though we are only interested in small areas of the domain such as boundaries, interfaces, or areas where variables exhibit complex dynamic behavior. Refining the computational mesh in these regions is a well-known technique in order to save memory and processing power, and it can be applied with various degrees of sophistication. Within IAMR, and the AMReX framework in general, hierarchically structured AMR is employed, which allows for efficient local refinement by symmetrically splitting each cell into 2^d smaller ones, where d is the dimensionality of the problem. Furthermore, the dynamic refinement is spatio-temporal, meaning that we employ subcycling in the discrete time steps to avoid setting global time step restrictions based on the smallest cells. This is crucial for dynamic AMR to be used efficiently. Tagging criteria are utilised in order for the code to determine which cells should be split or merged at a given refinement step, and these can depend on the values of any directly simulated or derived variables, in addition to spatial position and time.

For details on the implementation of structured AMR in this code, we refer to the websites of AMReX (<https://amrex-codes.github.io/amrex/>) and IAMR (<https://amrex-codes.github.io/IAMR/>), where full documentation and source code are available.

D. Parallel implementation

AMReX is a mature, open-source software framework for building massively parallel block-structured AMR applications. AMReX contains extensive software support for explicit and implicit grid-based operations. Multigrid solvers, including those for tensor systems, are included for cell-based and node-based data. AMReX uses a hybrid MPI/OpenMP approach for parallelization; in this model individual grids are distributed to MPI ranks, and OpenMP is used to thread over logical tiles within the grids. Applications based on AMReX have demonstrated excellent strong and weak scaling up to hundreds of thousands of cores.⁴⁸⁻⁵¹




FIG. 2. Papanastasiou regularisation for the Bingham model: stress magnitude (left) and apparent viscosity (right) as functions of the magnitude of the rate-of-strain tensor $|\dot{\gamma}|$. For decreasing ε , we recover a closer approximation to the actual Bingham model. Note that logarithmic axes are employed in order to highlight the behavior in the low-strain limit.

IV. VALIDATION

In order to validate our code, we consider the two-dimensional lid-driven cavity test case. It is a historically significant and widely used benchmark problem for viscous flow simulations, and consequently a large amount of reference results exist in the literature. Our domain is a square of side length \mathcal{L} , i.e., $\Omega = [0, \mathcal{L}]^2$, and is filled with a fluid of constant density. Initially, the system is at rest. All walls except the top (the ‘‘lid’’) are held fixed so that the boundary conditions on these walls are $\mathbf{u} = (0, 0)$. At time $t = 0$, the lid is instantaneously prescribed the tangential velocity \mathcal{U} so that the relevant boundary condition for $t > 0$ is $\mathbf{u} = (\mathcal{U}, 0)$. Without any other external forces ($\mathbf{f} = 0$), this drives a recirculating flow in the cavity, which reaches a steady-state solution for non-turbulent regimes.

In order to obtain a dimensionless form of (2), we let $\hat{\mathbf{x}} = \mathbf{x}/\mathcal{L}$, $\hat{\mathbf{u}} = \mathbf{u}/\mathcal{U}$, and $\hat{t} = t/(\mathcal{L}/\mathcal{U})$. Additionally, we take \mathcal{U}/\mathcal{L} as a characteristic strain-rate, and let

$$\hat{\gamma} = \frac{\dot{\gamma}}{\frac{\mathcal{U}}{\mathcal{L}}}, \quad \hat{\eta} = \frac{\eta}{2^{n-1}\mu\left(\frac{\mathcal{U}}{\mathcal{L}}\right)^{n-1}}, \quad (25)$$

$$\hat{p} = \frac{p}{2^{n-1}\mu\left(\frac{\mathcal{U}}{\mathcal{L}}\right)^n}, \quad \hat{\tau} = \frac{\tau}{2^{n-1}\mu\left(\frac{\mathcal{U}}{\mathcal{L}}\right)^n}. \quad (26)$$

By substituting these dimensionless variables into (2), we find that the governing equations in dimensionless form are

$$\frac{\partial \hat{\mathbf{u}}}{\partial \hat{t}} + \hat{\mathbf{u}} \cdot \hat{\nabla} \hat{\mathbf{u}} = \frac{1}{Re} \left(-\hat{\nabla} \hat{p} + \hat{\nabla} \cdot \hat{\tau} \right), \quad (27a)$$

$$\hat{\nabla} \cdot \hat{\mathbf{u}} = 0, \quad (27b)$$

$$\hat{\tau} = 2\hat{\eta}\hat{\gamma}, \quad (27c)$$

$$\hat{\eta} = \left(|\hat{\gamma}|^n + \frac{1}{2}HB \right) \frac{1 - e^{-Pa|\hat{\gamma}|}}{|\hat{\gamma}|}, \quad (27d)$$

where we have introduced the dimensionless groups

$$Re = \frac{\rho \mathcal{L}^2}{2^{n-1}\mu\left(\frac{\mathcal{U}}{\mathcal{L}}\right)^n}, \quad HB = \frac{\tau_0}{2^{n-1}\mu\left(\frac{\mathcal{U}}{\mathcal{L}}\right)^n}, \quad Pa = \frac{\mathcal{U}}{\varepsilon \mathcal{L}}. \quad (28)$$

The Reynolds number Re is the ratio of inertial forces to (power-law) viscous ones, while the Herschel-Bulkley number HB quantifies the effect of yield stress versus power-law

viscosity. The Papanastasiou number Pa measures the degree of regularisation employed in the apparent viscosity, with higher Pa being a more accurate description of unregularised viscoplasticity. It is worth noting that the Reynolds number for Newtonian flow ($n = 1$) is $Re_N = \rho \mathcal{U} \mathcal{L} / \mu$.

When we have information about the dimensionless apparent viscosity $\hat{\eta}$ throughout our domain at a given time, we can compute the local Reynolds number field $Re_L = Re/\hat{\eta}$. This value, which is proportional to the inverse of the apparent viscosity, provides insight about which regions of the domain are dominated by the effects of viscoplasticity, and which ones have near-Newtonian behavior. Finally, for the special case $n = 1$, we introduce the Bingham number $Bi = \tau_0 \mathcal{L} / \mu \mathcal{U}$, to allow comparisons with articles written on the simulation of Bingham fluids.

In the remainder of this section, we use $\mathcal{L} = 1m$, $\rho = 1000 \text{ kg/m}^3$, $\mathcal{U} = 1 \text{ m/s}$, and $Pa = 400$. It is important to choose a high enough value of Pa when using regularised schemes, and our choice is based on the arguments by Syrakos *et al.*,¹⁹ who showed that this value results in highly accurate simulations. The cavity is discretised spatially with 256 cells in each dimension. The lid-driven cavity flow is then uniquely defined by the choices of μ , n , and τ_0 .

To the authors’ best knowledge, reference results for Herschel-Bulkley fluids have not been presented in the literature. There are, however, plenty of results for power-law fluids and Bingham plastics. Since these two characterise separate components of the Herschel-Bulkley model, we expect that validating each of them individually is a good test for our code. Evaluation of the full Herschel-Bulkley model, based on this validation, is given in Sec. V.

Power-law fluids have zero yield stress and are different from the Newtonian case when the flow behavior index is different from unity. They were studied by Bell and Surana in 1994,⁵² who presented results with $n = 1/2$ and $n = 3/2$. A third-order finite volume method with upwinding was then used to study the same problem by Neofytou in 2005.⁵³ The classical way to compare results from the lid-driven cavity test case is to drive the system to steady state and then look at the velocity profiles in each direction at slices perpendicular to the flow. This ensures that we can compare with results from codes that can only calculate steady-state solutions.

Figure 3 shows our results for power-law fluid velocity profiles and includes reference results from the

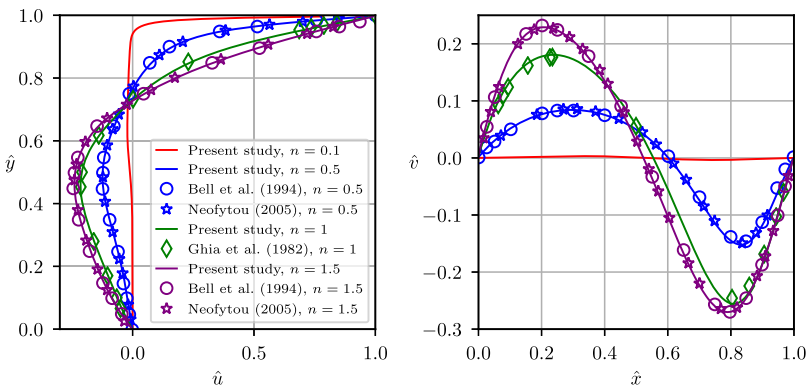


FIG. 3. Velocity profiles for power-law fluids with $\mu = 10 \text{ Pa s}^n$. Left: First component of velocity through the vertical slice $\hat{x} = 0.5$. Right: Second component of velocity through the horizontal slice $\hat{y} = 0.5$.

aforementioned papers. Additionally, we include the Newtonian case with comparisons from the classical paper by Ghia *et al.*⁵⁴ It is evident that decreasing n results in less lid-induced kinetic energy propagating further down in the domain, whereas increasing it has the opposite effect. Our results align very well with those found in the literature, especially those due to Neofytou, which have the highest accuracy. We have also included the velocity profile for $n = 0.1$, in order to illustrate the significant viscous resistance in the limit $n \rightarrow 0$.

We recall that Bingham fluids have $n = 1$ and non-zero yield stress τ_0 . Due to the importance of accurately capturing viscoplastic behavior, we have performed extra validation tests for this model. Figure 4 depicts velocity profiles at various Bingham numbers, with comparison to the studies of Neofytou⁵³ and Chai *et al.*⁵⁵ In the latter reference, a multiple-relaxation-time lattice Boltzmann model was employed to simulate non-Newtonian fluids. We see that increasing the yield-stress has an effect on the velocity profiles similar to that of lowering n for the power-law fluids, although the Bingham fluids have sharper transitions in the profile (especially noteworthy in the x -velocity). Our results are in excellent agreement with the relevant references, and we have also extended the range of Bingham numbers beyond those for which velocity profiles are available in the literature.

The second method we use for validation is locating the position of the main vortex centre within the cavity, at steady-state. A comprehensive study exploring this dependency was performed in 2014 by Syrakos *et al.*²⁰ We have performed simulations to steady-state for a large number of configurations of Re and Bi , and the resulting vortex locations are depicted in Fig. 5, alongside the results of Syrakos *et al.* We reiterate what they found in their paper: increasing the Bingham number moves the vortex upwards and to the right, while increasing the Reynolds number moves the vortex first toward the right, and then downwards and left to the centre. Our results agree very well with the reference results for the range of (Re, Bi) pairs covered in that study. Additionally, we have covered many more large Bingham numbers, obtaining results which follow the patterns to be expected. This illustrates the capability of our code to simulate fluids with very large yield stress. Please note that the number next to each point in the figure is the modified Reynolds number

$$Re^* = \frac{Re}{(Bi + 1)}. \quad (29)$$




FIG. 4. Velocity profiles for Bingham fluids with $Re = 100$. Left: First component of velocity through the vertical slice $\hat{x} = 0.5$. Right: Second component of velocity through the horizontal slice $\hat{y} = 0.5$.




FIG. 5. Vortex centre locations in the two-dimensional lid-driven cavity for Bingham fluids. Empty circles are results from Syrakos *et al.*,²⁰ while the remaining symbols are results from the present study. The number next to each point is the corresponding modified Reynolds number Re^* .

This is because, following Syrakos *et al.*,²¹ it provides a more natural measure of the transition to turbulence in the high- Bi region. For further discussions on choice of dimensionless parameters in viscoplastic fluid mechanics, we refer the reader to articles by de Souza Mendes⁵⁶ and Thompson and Soares.⁵⁷

The most important validation for viscoplastic fluids is the accurate determination of yield surfaces. Regions where the stress satisfies $|\tau| \leq \tau_0$ are known as unyielded regions. The interface which separates yielded and unyielded regions is called the yield surface. Since the yield surface separates two fundamentally different states of matter (viscous fluid flow and rigid solid behavior), locating it precisely is a valuable metric for evaluating numerical results. Finally, we would like to ensure that our time-marching scheme accurately captures the fluid movement. In order to validate both of these two attributes, we run a simulation starting from rest and evolving to the steady-state for the lid-driven cavity. Note that the criteria for the system having reached steady-state are that the maximum relative change in the velocity field from one iteration to the next is less than some tolerance (in our case 10^{-5}). At this point, we set the lid velocity to zero, and

let the system come to rest again through cessation. When the entire domain is an unyielded region, we stop the simulation. Although some authors have used time-stepping schemes to advance the system to steady-state,^{58,59} we are not aware of any results illustrating the yield surface development for instantaneous start from rest. On the other hand, the latter half of this numerical experiment, i.e., cessation of Bingham fluids

from the lid-driven steady-state, was studied numerically by Syrakos *et al.* in 2016.²¹ Comparisons with their yield surfaces and time measurements therefore serve as validation for our code.

Results with $Re = 1$ and $Bi = 10$ are depicted in Fig. 6. The yield surface is illustrated as the single black contour line $|\tau| = \tau_0$, while the heat map shows the local Reynolds number




FIG. 6. Yield surface and local Reynolds number for time-dependent flow in the 2D lid-driven cavity. The fluid is initially at rest, and gradually transitions to a steady state (a)–(e). At this point, the lid is stopped, allowing cessation in finite time (e)–(i). (a) $\hat{t} = 3.5 \times 10^{-6}$, (b) $\hat{t} = 1.05 \times 10^{-5}$, (c) $\hat{t} = 9.33 \times 10^{-5}$, (d) $\hat{t} = 9.540 \times 10^{-4}$, (e) $\hat{t} = 0.020555$, $\hat{t} = 0$, (f) $\hat{t} = 3.89 \times 10^{-5}$, (g) $\hat{t} = 6.480 \times 10^{-4}$, (h) $\hat{t} = 1.1664 \times 10^{-3}$, and (i) $\hat{t} = 1.2729 \times 10^{-3}$.

Re_L . This simple method for determining the yield surface can be improved upon for regularised constitutive equations, see, e.g., the work of Liu, Muller, and Denn.⁶⁴ Note that the scaling of the heat map is logarithmic (see top of Fig. 6).

For viscoplastic fluids, unyielded regions occur in two distinct manners. First, there are so-called stagnant zones, which are connected to no-slip walls and in which the fluid has zero velocity. On the other hand, unyielded zones with non-zero velocity are referred to as plug zones. These are not in contact with the walls and rotate in the interior of the cavity. In fact, these two types of unyielded zones can never be in contact with each other but must be separated by a layer of yielded fluid. This is because a plug zone connected to a stagnant zone would by association itself be stagnant and thus could not move.

The first five plots (a)–(e) show the transition from rest to steady-state. It is clear that the fluid yields immediately at a very thin layer near the lid since a large strain is applied here at $\hat{t} = 0$. Additionally, there are yielded regions which quickly propagate from the top corners and in toward the middle, eventually combining together and leaving an unyielded plug in the upper middle part of the cavity (a)–(c). This process occurs very quickly, in less than one percent of the time needed to reach steady-state but should in fact happen instantaneously, following the discussion above. The initial connection between the stagnant zone at the bottom and the rotating plug is a spurious artefact due to the regularisation approach. Since we regularise the apparent viscosity, small but non-zero velocities are permissible in unyielded regions.

The steady-state shape of the resulting plug, and the unyielded region at the bottom of the cavity (e), is characteristic for the chosen pair of Reynolds and Bingham numbers. Comparisons to results by Syrakos *et al.* show that the

steady-state yield surfaces are indistinguishable. At this point ($\hat{t} = 0.020\,552$), the lid is abruptly halted and held still. This removes the driving force for the flow from the system and leads to cessation of the flow. Since cessation occurs in finite time for viscoplastic fluids, the cessation time t_c , defined as the time until the entire cavity is in the unyielded state, is an important characteristic. In order to measure \hat{t}_c , we introduce a new time $\tilde{t} = \hat{t} - 0.020\,552$. We have chosen four times during flow cessation to compare with Syrakos *et al.*²¹ and the flow patterns are very similar to their published results. Additionally, our simulations yield $\hat{t}_c = 1.2844 \times 10^{-3}$, which is a relative difference of less than 0.9% compared to the previously published results. Note also the unphysical connection between stagnant and plug zones in (i), which is in agreement with the results and discussions of Syrakos *et al.*²¹ In order to rectify these spurious artefacts, one would need a very high Papanastasiou number or use an unregularised approach.⁶⁰

V. EVALUATION

A. Herschel-Bulkley fluids

The results in Sec. IV validate our code for fluids obeying power-law and Bingham rheologies, but there are no available comparisons in the literature for Herschel-Bulkley fluids, which exhibit power-law dependencies in addition to a non-zero yield stress. Consequently, simulations were performed for viscoplastic fluids with the flow index n equal to 0.5 and 2.0. In order to cover a wide range of flow behaviors, we also chose to vary the Herschel-Bulkley number HB and the modified Reynolds number Re^* , as given by (29). For each combination of n , HB , and Re^* , the system was advanced to




FIG. 7. Steady-state velocity profiles for Herschel-Bulkley fluids with $n = 0.5$ (top) and $n = 2.0$ (bottom), and various combinations of HB and Re^* , given in the legend in the top left plot. Note that for the dilatants with $HB = 100$, several lines almost overlap.

steady-state, in order to obtain velocity profiles through the centre of the domain and the final location of the main vortex. These results can be used as benchmark references for codes simulating Herschel-Bulkley fluids.

Figure 7 shows the steady-state velocity profiles for each different parameter combination. As expected, a small flow index enhances the effect of the yield stress, resulting in weak velocity variations. This leads to the largely overlapping x -velocity profiles in the upper left plot so that the y -velocity profiles in the upper right one are more useful comparisons. In the lower plots, however, where the fluid acts as a dilatant above the yield criterion, there is a trade-off between the effects of flow index and yield stress so that we obtain large velocity variations throughout the middle of the domain in both directions.

In Fig. 8, we show the steady-state locations of the main vortex centres for each Herschel-Bulkley simulation. Comparing the results with Fig. 5, we see the same effects of increasing HB as for Bi , i.e., vortex moving upwards and to the right. Similarly, increasing Re^* results in the vortex moving first right and then down toward the cavity centre. Both of these observations are as expected since the dimensionless groups capture similar aspects of the flow. Additionally, we note that decreasing the flow index n has a similar effect as increasing the yield stress, moving the vortex centre upwards and to the right. These additional data points serve as further benchmarking results for Herschel-Bulkley fluids.

B. Adaptive mesh refinement

The use of AMR improves the efficiency of our algorithm in two different ways, depending on whether we are interested in steady-state results or transient behavior. In the former case, we do not care about intermediate results and would ideally employ a large time step. Unfortunately, due to the semi-implicit viscous solves employed in IAMR, we are unable to reach steady-state for highly resolved simulations




FIG. 8. Steady-state vortex centres for Herschel-Bulkley fluids with different combinations of flow index n , Herschel Bulkley number HB , and modified Reynolds number Re^* . The number next to each point is the corresponding Re^* .

by iterating with large time steps. The large system requires longer runtimes per time step, and the time step decreases proportionally with the grid size. On the other hand, we can rapidly advance a low-resolution simulation to steady-state and then restart the simulation with one refinement level enabled throughout the entire domain. Due to the initial conditions, the finer resolution will converge after a much smaller amount of iterations. Subsequently, the simulation can be restarted again with a further level of refinement, and so on. In this manner, the computational expense associated with the time-dependent simulation to a highly resolved steady-state is greatly reduced. As an example, advancing the 2D case with $Re = 1$, $Bi = 10$ to steady-state on a 16-core desktop computer took 83 min and 13 s when using a resolution of 256^2 . By first advancing a low-resolution system with 64^2 cells to steady-state (33 s), and then refining the entire domain with two levels, the runtime for obtaining the same steady-state result is reduced to just under 15 min. Due to the built-in checkpoint functionality in AMReX, restarting simulations with different runtime options in this manner is straightforward.

In order to illustrate the usefulness of AMR in unsteady simulations, we refer to Fig. 9 (Multimedia view), which shows the full evolution of a Bingham fluid with $Re = 1000$ and $Bi = 1$ from rest to steady-state with AMR enabled. The base grid consists of 64 cells in each direction, and two levels of refinement are used to achieve high resolution near the yield surface. This is done by tagging cells where $|r| \leq 1.05\tau_0$ for refinement. Consequently, we only need to use large amounts of computational resources in some regions. Due to the block-structured AMR implementation, the grid distribution is highly parallel. Additionally, the temporal subcycling in time ensures that we only need to decrease the step size in refined cells—not globally.



FIG. 9. Illustration of spatio-temporal adaptive mesh refinement for tracking the development of the yield surface, with $Re = 1000$ and $Bi = 1$. The coarsest grid discretises with cell spacing $1/64$, but the two refinement layers ensure that the effective resolution near the yield surface is $1/256$. Multimedia view: <https://doi.org/10.1063/1.5049202.1>

C. Moving to three dimensions

Since the augmented IAMR code is efficient and parallelises well, it is possible to run unsteady simulations of fully three-dimensional viscoplastic flows. As an extension of the

lid-driven cavity for Bingham fluids in 2D, we consider a unit cube with stationary no-slip conditions on the floor and all four walls, while the lid (at $\hat{z} = 1$) drives the flow. We expect the results of such a simulation to be similar to the classical two-dimensional case, which means that we can compare the two

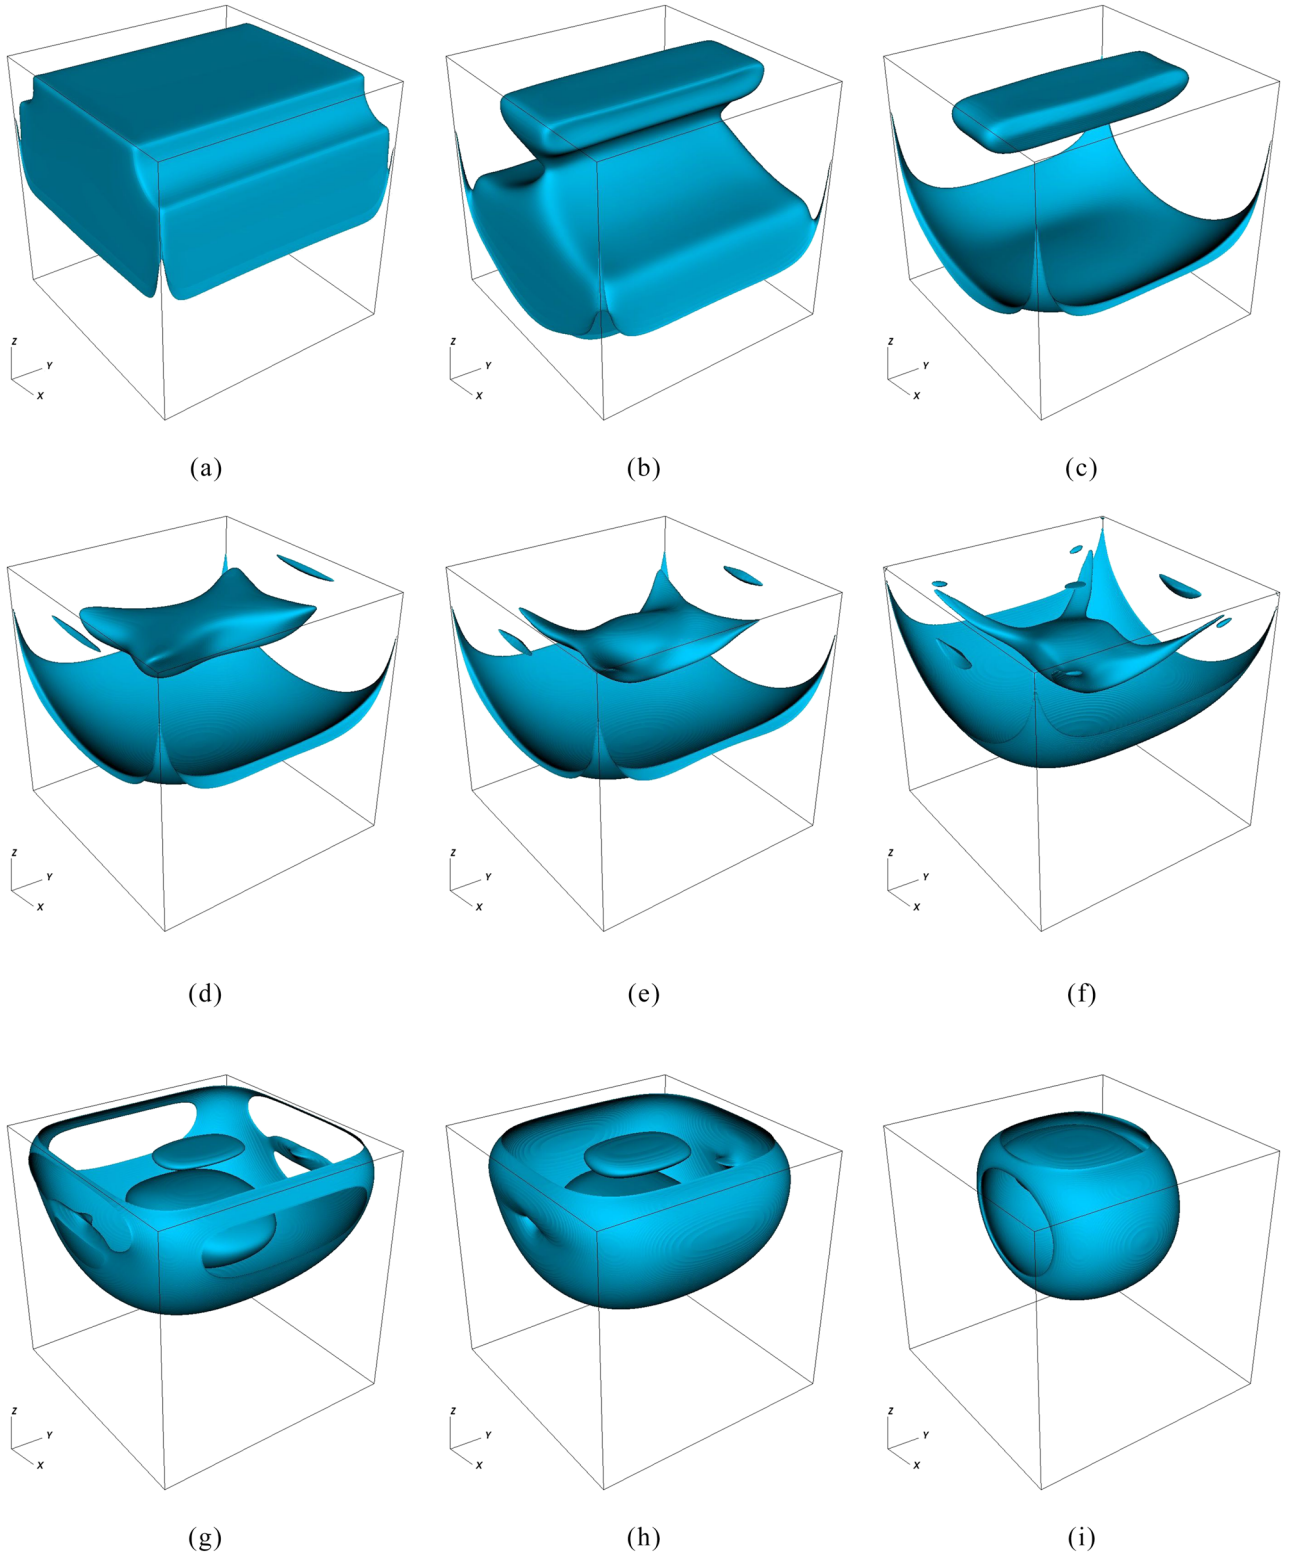


FIG. 10. Three-dimensional time-evolution of the lid-driven cavity flow, with $Re = 1$ and $Bi = 10$. The system advances from an instantaneous start from rest to steady-state, before the lid motion ceases and the fluid stops. The lid moves in the positive x -direction. The shapes of three-dimensional plug regions of this sort have never been identified for the lid-driven cavity before. (a) $\hat{t} = 1.0 \times 10^{-6}$, (b) $\hat{t} = 1.00 \times 10^{-5}$, (c) $\hat{t} = 6.82 \times 10^{-5}$, (d) $\hat{t} = 8.438 \times 10^{-4}$, (e) $\hat{t} = 1.4913 \times 10^{-2}$, $\hat{t} = 0$, (f) $\hat{t} = 3.97 \times 10^{-5}$, (g) $\hat{t} = 4.847 \times 10^{-4}$, (h) $\hat{t} = 7.706 \times 10^{-4}$, and (i) $\hat{t} = 1.0564 \times 10^{-3}$.

and expect similar flow characteristics. However, we would only get identical results by using periodic boundary condition on the walls which are parallel to the flow vortices. Since we apply no-slip conditions, there will be three-dimensional effects on the results, meaning that we can use them as separate reference results. The spatial domain is discretised using 256 cells in each direction, i.e., 16 777 216 cells in total. As before, the system is initially at rest. When steady-state is reached, we stop the motion of the lid and allow cessation to take place. Figure 10 illustrates how the yield surface evolves through time. Note that it is still computed as the contour $|\tau| = \tau_0$, but in 3D we can actually visualise it as a surface. The snapshot times in Fig. 10 are similar to those used in Fig. 6 but vary by small factors. Note also here the presence of numerical artefacts due to regularisation, demonstrated by connected stagnant and plug zones. Qualitatively, the resemblance with the 2D case is striking, as the flow clearly goes through the same steps to reach steady-state and cessation. Having said that, the third dimension allows a richer picture of the shape of the yield surface. Especially fascinating is the importance of the four vertical edges, which play a similar role as the 2D corners and lead to the yield surface stretching out in four directions from the centre. It would not be possible to

capture this type of yield surface shape without performing fully 3D fluid simulations. We emphasise that this is the first time fully three-dimensional yield surfaces have been presented for the lid-driven cavity case. Previously, the only 3D results are slices of rigid zones presented by Olshanskii⁶¹ (finite differences, 262 144 cells), and the heatmap slices presented in the bi-viscosity simulations performed by Elias *et al.*⁶² (819 468 elements).

In Fig. 11, steady-state yield surfaces are shown for various combinations of Re and Bi . Readers who are familiar with the two-dimensional test case (see, e.g., Ref. 20) will recognise the similarity between those results and slices through these three-dimensional regions. Note that the slices are taken at a slight angle in the xy -plane. These simulations illustrate the code's ability to solve the governing PDEs for time-dependent three-dimensional flows of non-Newtonian fluids following the Herschel-Bulkley model.

While the shapes of the three-dimensional yield surfaces shown in Figs. 10 and 11 provide interesting insight into the three-dimensional effects, they do not provide quantitative comparisons for other codes for three-dimensional viscoplastic fluid flows. Consequently, we show the corresponding steady-state velocity profiles in Fig. 12. The first




FIG. 11. Steady-state yield surface for Bingham fluids with various Reynolds and Bingham numbers in the three-dimensional lid-driven cavity. The lid moves in the positive x -direction. (a) $Re = 1$, $Bi = 1$, (b) $Re = 1$, $Bi = 10$, (c) $Re = 1$, $Bi = 100$, (d) $Re = 1000$, $Bi = 1$, (e) $Re = 1000$, $Bi = 10$, and (f) $Re = 1000$, $Bi = 100$.

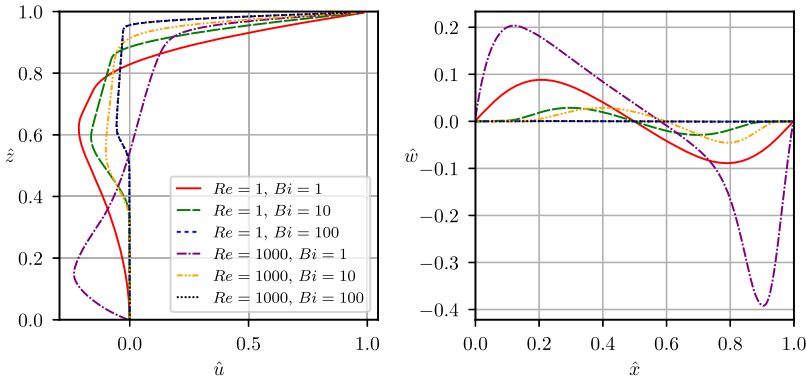


FIG. 12. Velocity profiles in the centre of the three-dimensional cavity, provided as reference solutions for 3D viscoplastic solvers. The effect of Reynolds and Bingham numbers is clear. Note that the two cases with $Bi = 100$ are indistinguishable.

velocity component is considered on the line $\hat{x} = 0.5$, $\hat{y} = 0.5$, while the x - and z -directions are swapped for the third component. As is evident, there are notable similarities between the two-dimensional simulations, as expected. These profiles can act as benchmark results for future yield stress simulation codes.

D. Parallel performance

Among the most impressive attributes of codes built within the AMReX framework, is the intrinsic scalability on state-of-the-art computer architectures. Due to major advances within parallel computing in the last few decades, accurate measurement of how well-suited algorithms are for parallel processing has received a lot of attention. Moreland and Oldfield provided an excellent overview of contemporary formal metrics for such measurements.⁶³

Suppose we are trying to solve a computational problem of size N . Then the speed-up of the parallel algorithm is given by the ratio of the parallel execution time to that of the best serial algorithm. The theoretically best possible speed-up occurs if the entire task can be divided into as many chunks of equal work as there are cores being used so that it is distributed among the available cores without any overhead. In this utopian case, we have speed-up equal to the processor core count. Although speed-up is a useful metric for analysing the parallel performance of an algorithm, it has one caveat. Since the serial runtime is necessary in order to compute them, it quickly becomes impractical for analysis of large-scale applications. Computing the serial runtime for very large N takes far too long. Scaling where the problem size is kept fixed, and the amount of allocated resources is increased,

is known as strong scaling. On the other hand, the problem size can be scaled so that it is proportional to the core amount. If one core is used to solve the problem with size N^* , one would then ideally expect the parallel runtime of a problem with size PN^* to remain constant when deployed on P cores. This type of analysis is called weak scaling. Finally, it may in some case be difficult to keep the problem size proportional to the job size so that weak scaling analysis becomes difficult to obtain. However, one can measure the amount of work done per time, referred to as the rate. The rate is computed as the ratio of problem size to parallel runtime and thus removes the weak scaling dependency between problem size and core amount. Sampling the rate across a selection of problem sizes and core amounts can thus help the simulation runner pick runtime parameters which yield a high computational efficiency.

Our main interest lies in ensuring that the scalability is unaffected by the extension to generalised Newtonian fluids. As such, we run a test-suite with varying problem sizes and processor cores for the 3D lid-driven cavity test when the domain is filled by a Bingham fluid with $Re = 1$ and $Bi = 10$. In order to avoid the problems associated with strong and weak scaling, we vary the number of cores and problem resolution but only use few cores for smaller problems and many cores for larger ones. Subsequently, we compute the rate for the cell updates, which gives proper insight into the resulting computational efficiency. In order to evaluate the parallel performance rather than the effect of the system state, we use a constant time step, advect the system by 100 time steps, and measure the average runtime per time step for the next ten.

Figure 13 illustrates the results of average runtimes and the rate for the scaling test. For a given problem size, the

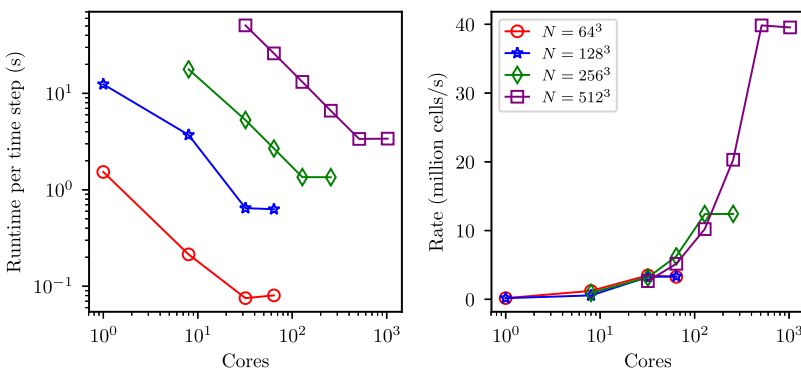


FIG. 13. Strong scaling results for a Bingham-fluid in the lid-driven cavity. The left plot shows the reduction in runtime per time step as the amount of processor cores is increased, for various problem sizes. The right plot shows how the corresponding computation rate grows. Note that both axes are logarithmic in the left plot, while only the primary axis employs logarithmic scaling in the right one.

reduction in runtime (and increase in rate) is basically optimal until a core limit is reached. At this point, the MPI communication overhead becomes too large for further parallelisation since the grid has already been broken down too far. However, as the problem gets bigger this core limit naturally increases so that the computational rate can continue increasing as long as larger problems are solved. When simulating at a given problem resolution, it is therefore best to choose the lowest amount of cores which maximise the rate. Note that these tests do not utilise OpenMP tiling and have not been optimised in terms of runtime parameters for parallelisation (such as maximum size for subgrids), and still the scalability is excellent.

VI. CONCLUSIONS

The publicly available software suite IAMR facilitates rapid simulations of incompressible fluids governed by the Navier-Stokes equations in a structured, adaptive framework with state-of-the-art parallelisation on contemporary super-computer architectures. By augmenting this code, originally designed for Newtonian flows, we allow it to deal with generalised Newtonian fluids using a regularised Herschel-Bulkley model. In doing so, we provide a code capable of simulating large-scale viscoplastic fluid systems efficiently enough to investigate three-dimensional, time-dependent systems. Thorough validation is performed through comparisons with results from peer-reviewed studies of the lid-driven cavity problem, and these benchmark results are expanded upon through a wider range of Reynolds and Bingham numbers. Finally, we have evaluated the ability of the code to deal with Herschel-Bulkley fluids and investigated three-dimensional yield surfaces in the lid-driven cavity. For both Herschel-Bulkley fluids and Bingham fluids in three dimensions, reference results are provided for the first time.

ACKNOWLEDGMENTS

K.S. would like to acknowledge the EPSRC Centre for Doctoral Training in Computational Methods for Materials Science for funding under Grant No. EP/L015552/1. Additionally, he acknowledges the funding and technical support from BP through the BP International Centre for Advanced Materials (BP-ICAM) which made this research possible.

- ¹S. S. Pegler and N. J. Balmforth, "Locomotion over a viscoplastic film," *J. Fluid Mech.* **727**, 1–29 (2013).
- ²M. W. Denny, "A quantitative model for the adhesive locomotion of the terrestrial slug, *ariolimax columbianus*," *J. Exp. Biol.* **91**, 195–217 (1981).
- ³A. J. Apostolidis and A. N. Beris, "Modeling of the blood rheology in steady-state shear flows," *J. Rheol.* **58**, 607–633 (2014).
- ⁴A. J. Apostolidis, M. J. Armstrong, and A. N. Beris, "Modeling of human blood rheology in transient shear flows," *J. Rheol.* **59**, 275–298 (2015).
- ⁵A. J. Apostolidis and A. N. Beris, "The effect of cholesterol and triglycerides on the steady state shear rheology of blood," *Rheol. Acta* **55**, 497–509 (2016).
- ⁶A. J. Apostolidis, A. P. Moyer, and A. N. Beris, "Non-Newtonian effects in simulations of coronary arterial blood flow," *J. Non-Newtonian Fluid Mech.* **233**, 155–165 (2016).
- ⁷S. Bittleston and D. Guillot, "Mud removal: Research improves traditional cementing guidelines," *Oilfield Rev.* **3**, 44–54 (1991).

- ⁸S. Taghavi, K. Alba, M. Moyers-Gonzalez, and I. Frigaard, "Incomplete fluid–fluid displacement of yield stress fluids in near-horizontal pipes: Experiments and theory," *J. Non-Newtonian Fluid Mech.* **167**, 59–74 (2012).
- ⁹I. A. Frigaard, K. G. Paso, and P. R. de Souza Mendes, "Bingham's model in the oil and gas industry," *Rheol. Acta* **56**, 259–282 (2017).
- ¹⁰P. Saramito and A. Wachs, "Progress in numerical simulation of yield stress fluid flows," *Rheol. Acta* **56**, 211–230 (2017).
- ¹¹M. Bercovier and M. Engelman, "A finite-element method for incompressible non-Newtonian flows," *J. Comput. Phys.* **36**, 313–326 (1980).
- ¹²R. Tanner and J. Milthorpe, "Numerical simulation of the flow of fluids with yield stress," in *Proceedings of Third International Conference on Numerical Methods in Laminar and Turbulent Flow*, edited by C. Taylor, C. Johnson, and I. Smith (Pineridge Press, Swansea, UK, 1983), pp. 680–690.
- ¹³A. Guha and S. Sengupta, "Analysis of von Kármán's swirling flow on a rotating disc in Bingham fluids," *Phys. Fluids* **28**, 013601 (2016).
- ¹⁴T. C. Papanastasiou, "Flows of materials with yield," *J. Rheol.* **31**, 385–404 (1987).
- ¹⁵E. Mitsoulis and T. Zisis, "Flow of Bingham plastics in a lid-driven square cavity," *J. Non-Newtonian Fluid Mech.* **101**, 173–180 (2001).
- ¹⁶T. Zisis and E. Mitsoulis, "Viscoplastic flow around a cylinder kept between parallel plates," *J. Non-Newtonian Fluid Mech.* **105**, 1–20 (2002).
- ¹⁷E. Mitsoulis, "On creeping drag flow of a viscoplastic fluid past a circular cylinder: Wall effects," *Chem. Eng. Sci.* **59**, 789–800 (2004).
- ¹⁸E. Mitsoulis and J. Tsamopoulos, "Numerical simulations of complex yield-stress fluid flows," *Rheol. Acta* **56**, 231–258 (2017).
- ¹⁹A. Syrakos, G. C. Georgiou, and A. N. Alexandrou, "Solution of the square lid-driven cavity flow of a Bingham plastic using the finite volume method," *J. Non-Newtonian Fluid Mech.* **195**, 19–31 (2013).
- ²⁰A. Syrakos, G. C. Georgiou, and A. N. Alexandrou, "Performance of the finite volume method in solving regularised Bingham flows: Inertia effects in the lid-driven cavity flow," *J. Non-Newtonian Fluid Mech.* **208**, 88–107 (2014).
- ²¹A. Syrakos, G. C. Georgiou, and A. N. Alexandrou, "Cessation of the lid-driven cavity flow of Newtonian and Bingham fluids," *Rheol. Acta* **55**, 51–66 (2016).
- ²²M. Fortin and R. Glowinski, "On decomposition-coordination methods using an augmented Lagrangian," in *Augmented Lagrangian Methods: Applications to the Numerical Solution of Boundary-Value Problems*, Studies in Mathematics and Its Applications, edited by M. Fortin and R. Glowinski (Elsevier, 1983), Chap. III, Vol. 15, pp. 97–146.
- ²³G. Duvaut and J. L. Lions, *Inequalities in Mechanics and Physics* (Springer, 1976).
- ²⁴M. R. Hestenes, "Multiplier and gradient methods," *J. Optim. Theory Appl.* **4**, 303–320 (1969).
- ²⁵P. Saramito and N. Roquet, "An adaptive finite element method for viscoplastic fluid flows in pipes," *Comput. Methods Appl. Mech. Eng.* **190**, 5391–5412 (2001).
- ²⁶D. Vola, L. Boscardin, and J. Latché, "Laminar unsteady flows of Bingham fluids: A numerical strategy and some benchmark results," *J. Comput. Phys.* **187**, 441–456 (2003).
- ²⁷T. Treskatis, M. A. Moyers-Gonzalez, and C. J. Price, "An accelerated dual proximal gradient method for applications in viscoplasticity," *J. Non-Newtonian Fluid Mech.* **238**, 115–130 (2016).
- ²⁸P. Saramito, "A damped Newton algorithm for computing viscoplastic fluid flows," *J. Non-Newtonian Fluid Mech.* **238**, 6–15 (2016).
- ²⁹J. Bleyer, "Advances in the simulation of viscoplastic fluid flows using interior-point methods," *Comput. Methods Appl. Mech. Eng.* **330**, 368–394 (2018).
- ³⁰Y. Dimakopoulos, G. Makrigiorgos, G. Georgiou, and J. Tsamopoulos, "The PAL (penalized augmented Lagrangian) method for computing viscoplastic flows: A new fast converging scheme," *J. Non-Newtonian Fluid Mech.* **256**, 23–41 (2018).
- ³¹S. Enayatpour and E. van Oort, "Advanced modeling of cement displacement complexities," in *SPE/IADC Drilling Conference and Exhibition* (Society of Petroleum Engineers, 2017).
- ³²P. Tardy, N. Flamant, E. Lac, A. Parry, C.utama, S. Almagro *et al.*, "New generation 3D simulator predicts realistic mud displacement in highly deviated and horizontal wells," in *SPE/IADC Drilling Conference and Exhibition* (Society of Petroleum Engineers, 2017).
- ³³M. Alnæs, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, and G. N. Wells, "The FEniCS project version 1.5," *Arch. Numer. Software* **3**, 9–23 (2015).

- ³⁴H. Jasak, A. Jemcov, Z. Tukovic *et al.*, “OpenFOAM: A C++ library for complex physics simulations,” in *International Workshop on Coupled Methods in Numerical Dynamics* (IUC, Dubrovnik, Croatia, 2007), Vol. 1000, pp. 1–20.
- ³⁵A. Almgren, J. Bell, P. Colella, L. Howell, and M. Welcome, “A conservative adaptive projection method for the variable density incompressible Navier–Stokes equations,” *J. Comput. Phys.* **142**, 1–46 (1998).
- ³⁶H. A. Barnes, “The yield stress—A review of $\pi\alpha\tau\alpha\ \rho\epsilon\iota$ —Everything flows?,” *J. Non-Newtonian Fluid Mech.* **81**, 133–178 (1999).
- ³⁷N. J. Balmforth, I. A. Frigaard, and G. Ovarlez, “Yielding to stress: Recent developments in viscoplastic fluid mechanics,” *Annu. Rev. Fluid Mech.* **46**, 121–146 (2014).
- ³⁸E. C. Bingham, “An investigation of the laws of plastic flow,” *Bull. Bur. Stand.* **13**, 309–353 (1916).
- ³⁹J. Oldroyd, “A rational formulation of the equations of plastic flow for a Bingham solid,” in *Mathematical Proceedings of the Cambridge Philosophical Society* (Cambridge University Press, 1947), Vol. 43, pp. 100–105.
- ⁴⁰W. H. Herschel and R. Bulkley, “Konsistenzmessungen von Gummi-Benzollösungen,” *Kolloid-Z.* **39**, 291–300 (1926).
- ⁴¹R. B. Pember, L. H. Howell, J. B. Bell, P. Colella, W. Y. Crutchfield, W. Fiveland, and J. Jessee, “An adaptive projection method for unsteady, low-mach number combustion,” *Combust. Sci. Technol.* **140**, 123–168 (1998).
- ⁴²M. S. Day and J. B. Bell, “Numerical simulation of laminar reacting flows with complex chemistry,” *Combust. Theory Modell.* **4**, 535–556 (2000).
- ⁴³M. Lai, J. Bell, and P. Colella, “A projection method for combustion in the zero Mach number limit,” in *11th Computational Fluid Dynamics Conference* (American Institute of Aeronautics and Astronautics, Inc., 1993), p. 3369.
- ⁴⁴A. S. Almgren, J. B. Bell, and W. G. Szymczak, “A numerical method for the incompressible Navier–Stokes equations based on an approximate projection,” *SIAM J. Sci. Comput.* **17**, 358–369 (1996).
- ⁴⁵D. F. Martin and P. Colella, “A cell-centered adaptive projection method for the incompressible Euler equations,” *J. Comput. Phys.* **163**, 271–312 (2000).
- ⁴⁶P. Colella, “A direct Eulerian MUSCL scheme for gas dynamics,” *SIAM J. Sci. Stat. Comput.* **6**, 104–117 (1985).
- ⁴⁷S. Mirzaagha, R. Pasquino, E. Iuliano, G. D’Avino, F. Zonfrilli, V. Guida, and N. Grizzuti, “The rising motion of spheres in structured fluids with yield stress,” *Phys. Fluids* **29**, 093101 (2017).
- ⁴⁸A. Almgren, V. Beckner, J. Bell, M. Day, L. Howell, C. Joggerst, M. Lijewski, A. Nonaka, M. Singer, and M. Zingale, “CASTRO: A new compressible astrophysical solver. I. Hydrodynamics and self-gravity,” *Astrophys. J.* **715**, 1221 (2010).
- ⁴⁹A. S. Almgren, J. B. Bell, M. J. Lijewski, Z. Lukić, and E. Van Andel, “Nyx: A massively parallel AMR code for computational cosmology,” *Astrophys. J.* **765**, 39 (2013).
- ⁵⁰A. Dubey, A. Almgren, J. Bell, M. Berzins, S. Brandt, G. Bryan, P. Colella, D. Graves, M. Lijewski, F. Löffler, B. O’Shea, E. Schnetter, B. Van Straalen, and K. Weide, “A survey of high level frameworks in block-structured adaptive mesh refinement packages,” *J. Parallel Distrib. Comput.* **74**, 3217–3227 (2014).
- ⁵¹S. Habib, A. Pope, H. Finkel, N. Frontiere, K. Heitmann, D. Daniel, P. Fasel, V. Morozov, G. Zagaris, T. Peterka, V. Vishwanath, Z. Lukić, S. Sehrish, and W.-k. Liao, “HACC: Simulating sky surveys on state-of-the-art supercomputing architectures,” *New Astron.* **42**, 49–65 (2016).
- ⁵²B. C. Bell and K. S. Surana, “ p -version least squares finite element formulation for two-dimensional, incompressible, non-Newtonian isothermal and non-isothermal fluid flow,” *Int. J. Numer. Methods Fluids* **18**, 127–162 (1994).
- ⁵³P. Neofytou, “A 3rd order upwind finite volume method for generalised Newtonian fluid flows,” *Adv. Eng. Software* **36**, 664–680 (2005).
- ⁵⁴U. Ghia, K. N. Ghia, and C. Shin, “High-Re solutions for incompressible flow using the Navier–Stokes equations and a multigrid method,” *J. Comput. Phys.* **48**, 387–411 (1982).
- ⁵⁵Z. Chai, B. Shi, Z. Guo, and F. Rong, “Multiple-relaxation-time lattice Boltzmann model for generalized Newtonian fluid flows,” *J. Non-Newtonian Fluid Mech.* **166**, 332–342 (2011).
- ⁵⁶P. R. de Souza Mendes, “Dimensionless non-Newtonian fluid mechanics,” *J. Non-Newtonian Fluid Mech.* **147**, 109–116 (2007).
- ⁵⁷R. L. Thompson and E. J. Soares, “Viscoplastic dimensionless numbers,” *J. Non-Newtonian Fluid Mech.* **238**, 57–64 (2016).
- ⁵⁸E. J. Dean, R. Glowinski, and G. Guidoboni, “On the numerical simulation of Bingham visco-plastic flow: Old and new results,” *J. Non-Newtonian Fluid Mech.* **142**, 36–62 (2007).
- ⁵⁹L. Muravleva, “Uzawa-like methods for numerical modeling of unsteady viscoplastic Bingham medium flows,” *Appl. Numer. Math.* **93**, 140–149 (2015).
- ⁶⁰Y. Dimakopoulos and J. Tsamopoulos, “Transient displacement of a viscoplastic material by air in straight and suddenly constricted tubes,” *J. Non-Newtonian Fluid Mech.* **112**, 43–75 (2003).
- ⁶¹M. A. Olshanskii, “Analysis of semi-staggered finite-difference method with application to Bingham flows,” *Comput. Methods Appl. Mech. Eng.* **198**, 975–985 (2009).
- ⁶²R. N. Elias, M. A. Martins, and A. L. Coutinho, “Parallel edge-based solution of viscoplastic flows with the SUPG/PSPG formulation,” *Comput. Mech.* **38**, 365–381 (2006).
- ⁶³K. Moreland and R. Oldfield, “Formal metrics for large-scale parallel performance,” in *International Conference on High Performance Computing* (Springer, 2015), pp. 488–496.
- ⁶⁴B. T. Liu, S. J. Muller, and M. M. Denn, “Interactions of two rigid spheres translating collinearly in creeping flow in a Bingham material,” *J. Non-Newtonian Fluid Mech.* **113**, 49–67 (2003).