Journal of Non-Newtonian Fluid Mechanics 256 (2018) 23–41

Contents lists available at ScienceDirect

Journal of Non-Newtonian Fluid Mechanics

journal homepage: www.elsevier.com/locate/jnnfm

The PAL (Penalized Augmented Lagrangian) method for computing viscoplastic flows: A new fast converging scheme

Y. Dimakopoulos[a,*], G. Makrigiorgos[a], G.C. Georgiou[b], J. Tsamopoulos[a]

[a] Department of Chemical Engineering, University of Patras, Patras 26504, Greece
[b] Department of Mathematics and Statistics, University of Cyprus, Nicosia 1678, Cyprus

ARTICLE INFO

Keywords:
Viscoplastic fluids
Papanastasiou regularization
Augmented Lagrangian method
Penalty method
Bubble rise
Lid-driven cavity
Filament stretching

ABSTRACT

Computation of viscoplastic fluid flows has always been a challenging task. Viscoplastic models are intrinsically discontinuous at the yielded-unyielded interface, which leads to numerical difficulties, because of the singularity in the Jacobian matrix of the resulting discretized equations. For this reason, several modeling or numerical approaches have been proposed, the most popular being the Papanastasiou regularization (PR) and the Augmented Lagrangian (AL) methods, respectively. Recently, studies on AL methods have focused on developing accelerated algorithms, since the required computational cost of using AL is extremely high. In the present work, a fast converging and efficient algorithm is proposed for tracking the yield surface and predicting the flow field of viscoplastic fluids accurately. The numerical procedure is the Penalized Augmented Lagrangian (PAL) method, which is based on a monolithic Newton solver for AL, where the governing equations of the Lagrange-multiplier tensor for both the rate-of-strain projection and the extra-stress tensors are penalized. To test the efficiency of our algorithm, five benchmark flow-problems with fixed, free and moving boundaries are studied. First, the problem of the steady rise of a bubble in a viscoplastic medium is addressed validating the new algorithm with the findings by Dimakopoulos et al. (2013). Then the entrance flow in a rectangular channel is solved, where a primary unyielded region is found around the centerline in the developed part of the flow and secondary unyielded regions near the entrance. In addition, the lid-driven cavity problem is solved, which is an often used test for various numerical algorithms and the results are compared to relevant studies for viscoplastic fluids such as those of Syrakos et al. (2013, 2014) and Treskatis et al. (2016). Furthermore, the developed flow in a square duct is examined, similarly to Saramito (2016). Finally, the transient filament stretching of a shear-thinning, yield stress fluid is examined, and the results are compared to those by Balmforth et al. (2010). In all cases, either steady or transient, the algorithm captures the yield surfaces correctly, while maintaining a low computational cost, because the convergence of the method requires only a few (i.e. 5–30) Newton iterations. Based on these extensive tests, PAL is found to be superior combining accuracy and speed to all existing solution methods for viscoplastic fluids.

1. Introduction

1.1. Yield stress fluids and computational challenges

Numerous engineering applications and natural phenomena involve materials characterized as yield–stress materials. From a rheological standpoint, they are also called viscoplastic fluids and are distinguished from other non–Newtonian fluids, because of the existence of a stress threshold, namely the yield–stress, which plays a crucial role in their flow. The dual identity of viscoplastic fluids lies in the fact that they behave as solids if the measure of the applied stress does not exceed the yield–stress value, and flow as liquids if the yield–stress is surpassed. Although the existence of a real yield-stress in viscoplastic fluids has been debated, in most applications it is regarded as a very useful concept for describing their behavior. Modeling yield-stress materials is one of the elusive tasks in rheology due to the complexity and variety of their structure. A recent review on viscoplastic fluids discussing the aforementioned issues is the work of Balmforth et al. [5]. First, Bingham [8] attempted to model viscoplastic materials. In a simplified form, he assumed that the rate of deformation is zero when the material is in the unyielded state, while the material flows as a Newtonian fluid once it yields. A modification of the Bingham model, to include shear thinning effects, led to the development of the Herschel-Bulkley (HB) model [33]. The general three-dimensional form of this constitutive law

* Corresponding author.
E-mail address: dimako@chemeng.upatras.gr (Y. Dimakopoulos).

https://doi.org/10.1016/j.jnnfm.2018.03.009
Received 19 October 2017; Received in revised form 2 March 2018; Accepted 11 March 2018
Available online 14 March 2018
0377-0257/

can be expressed as:

$$\max\left(0,\ \frac{|\underline{\tilde{\tau}}| - \tilde{\tau}_y}{\tilde{k}}\right)^{1/n} \frac{\underline{\tilde{\tau}}}{|\underline{\tilde{\tau}}|} = \underline{\tilde{\dot{\gamma}}} \tag{1}$$

Henceforth, the tilde $\sim$ over a variable denotes its dimensional form. In the above equations, $\underline{\tilde{\tau}}$ is the extra stress tensor and $\underline{\tilde{\dot{\gamma}}}$ is the rate-of-strain tensor, which is defined as $\underline{\tilde{\dot{\gamma}}} = \tilde{\nabla}\underline{\tilde{v}} + (\tilde{\nabla}\underline{\tilde{v}})^T$. The magnitudes of the aforementioned tensors are given by:

$$|\underline{\tilde{\tau}}| \equiv \left[\frac{1}{2}\underline{\tilde{\tau}} : \underline{\tilde{\tau}}\right]^{1/2},\ |\underline{\tilde{\dot{\gamma}}}| \equiv \left[\frac{1}{2}\underline{\tilde{\dot{\gamma}}} : \underline{\tilde{\dot{\gamma}}}\right]^{1/2} \tag{2}$$

The model generally includes three parameters. They are the consistency index $\tilde{k}$ (a generalized plastic viscosity), the flow index n, which controls the shear–thinning or thickening behavior of the material and, of course, the yield stress, $\tilde{\tau}_y$. When $n = 1$, the consistency index becomes the plastic Bingham viscosity $\tilde{\eta}_p$. In addition, the fluid is characterized by its density $\tilde{\rho}$ and interfacial tension with air $\tilde{\gamma}$.

One crucial limitation of the model is that for stresses lower than the yield stress, the solid region is assumed to behave as a rigid (inelastic) material. This introduces ambiguity, because in the unyielded domain the extra stress remains undetermined and Eq. (1) reduces to:

$$\underline{\tilde{\dot{\gamma}}} = 0 \tag{3}$$

while at the same time, the velocity is subjected to the momentum balance throughout. Indeed, the rate of deformation should be zero, since the material is assumed to be a rigid solid. However, the equation above has infinite solutions regarding the stress field. Any arbitrary combination of shear and normal stresses that may develop in the unyielded region by the flow can satisfy Eq. (3), as long as the second invariant of the stress tensor for this combination is lower than the yield stress, providing an infinity of acceptable solutions. Thus, the problem is ill-posed. Unfortunately, this is not the only drawback of this model. The existence of the max term in Eq. (1) or the related discontinuous viscosity introduces a numerical singularity in the Jacobian matrix, if a Newton iterative scheme is adopted, when solving this non-regularized formulation, due to the discontinuity of the derivative of the max function .This is also stressed in the work of Saramito [52] and Saramito and Wachs [55]. One way to overcome these shortcomings is to determine the flow field in the yielded region only along with the yield surface, as accomplished by Beris et al. [7] and Smyrnaios et al. [56]. However, this approach can be used only when the physics of the problem allow us to have a good estimation of the location of the yield surface before solving the problem.

The analysis in the present work is centered in the Herschel–Bulkley model, which does not account for the elasticity of the material. The more general approach is to include the elastic deformation in the unyielded region of these materials, which are then called elasto-viscoplastic. Such constitutive models have recently appeared in the literature, Saramito [53,54], and avoid one of the discrepancies of any purely viscoplastic model. Typically, in the unyielded region, the material is considered to be a Neo-Hookean elastic solid, [24,25] and thus, the stress there may be determined. Nevertheless, multiplicity of the solutions is not removed, even in this case, as first shown by Cheddadi et al. [12].

Predicting the yield surface has been extensively investigated in the field of computational fluid mechanics. The vast majority of papers involving flows of viscoplastic fluids, present numerical algorithms, which are based on one of two approaches. The first approach is the regularization models, where the discontinuous viscosity function is smoothed in order to asymptotically match the original model, as discussed in the review by Mitsoulis and Tsamopoulos [37]. Arguably the most commonly used approximation among them is Papanastasiou's regularization (PR, Papanastasiou [43]). It approximates the fluid rheology by a generalized Newtonian one, with a viscosity continuously changing in the yielded and unyielded regions, according to Eq. (4):

$$\underline{\tilde{\tau}} = \tilde{\eta}_{pap}(\underline{\tilde{\dot{\gamma}}})\ \underline{\tilde{\dot{\gamma}}},\quad \tilde{\eta}_{pap}(\underline{\tilde{\dot{\gamma}}}) = \tilde{k}\,|\underline{\tilde{\dot{\gamma}}}|^{n-1} + \tilde{\tau}_y\left(\frac{1 - \exp(-\widetilde{N}|\underline{\tilde{\dot{\gamma}}}|)}{|\underline{\tilde{\dot{\gamma}}}|}\right) \tag{4}$$

It is a popular method, because it is easy to implement and introduces just one regularization parameter ($\widetilde{N}$) in order to smooth the discontinuous viscosity function, via the exponential function. Relevant to the present study and exclusively based on the Papanastasiou regularization is the work of Damianou and Georgiou [13]. They examined the Poiseuille flow of a Herschel–Bulkley fluid in a rectangular duct, emphasizing wall-slip and assuming that it follows a law similar to HB. From a computational perspective, they regularized both the constitutive model and the slip equations according to Damianou et al. [14] and obtained accurate results when the regularization parameters were sufficiently large. Transient simulations of viscoplastic fluids with the Papanastasiou regularization have been reported by Dimakopoulos and Tsamopoulos [17] when air is displacing it from a tube and by Papaioannou et al. [42] when it is injected in a tube. In addition to the finite element method used in the previously mentioned publications, Syrakos et al. [57–60] have used the finite volume method with the Papanastasiou regularization for the steady and transient lid–driven cavity flow and the flow inside an extrusion damper. Finally, Burgos et al. [10] have compared the predictions of the Papanastasiou model as well as the biviscosity model to the original HB model in the determination of yield surfaces in an antiplane shear flow. It is no wonder why the Papanastasiou approximation is so popular and employed in various viscoplastic flows, regardless of geometry and type of flow. It is simple to implement in any solver (e.g. FD, FEM or FVM) and it follows physical intuition. Of course, it comes with drawbacks, as it will be seen later.

The second approach, which is the focus of the present work, is based on variational principles, which require that the functions that constitute the solution of a problem should maximize or minimize certain functionals. A benchmark study for viscoplastic flows also using variational principles is the work by Beris et al. [7] where the sedimentation of a sphere in a viscoplastic fluid is studied. In addition, variational inequalities have been used in many studies by Muravleva and Muravleva [38–40]. Furthermore, Dubash and Frigaard [20] used Prager's variational principles in order to estimate the conditions under which the rise of a bubble in an HB fluid stops. The most famous variational approach for the solution of viscoplastic problems is the Augmented Lagrangian (AL) method, which is based on Variational Inequalities, either of rate–of–strain minimization or stress maximization. This method has been introduced by Hestenes [34] and Powell [48] for the solution of optimization problems. Glowinski and co-workers [21,28,29] advanced and applied it to the linear Stokes problem and to non–linear problems, such as the flow of Bingham-type fluids.

Briefly, according to the AL method, in order to solve the viscoplastic flow problem, a functional is defined, which is minimized by the correct solution. The singularity that arises at the yield surface is remedied by the introduction of a Lagrange multiplier, related to the stress tensor, which aids in the relaxation of computing the velocity gradient. Additionally, a “quadratic” term is incorporated in the functional to accelerate the convergence of the iterative scheme. The two aforementioned terms are multiplied by certain coefficients which are updated in each iteration according to a conjugate gradient method. Dimakopoulos et al. [15] simulated the steady bubble rise in Herschel–Bulkley fluids following either the PR or the AL approach. The mass and momentum conservation equations along with the constitutive equations were solved in order to study the buoyancy-driven rise of a bubble in a viscoplastic medium, mainly in order to examine the deformation and shape of the bubble and also the conditions for its entrapment. From a physical perspective, an important conclusion was that shear–thinning does not affect the entrapment conditions and, additionally, viscoplasticity cannot lead to the formation of cusps in the south pole of the bubble giving it an inverted tear–drop shape. This left

elastic effects to be their only possible cause. As for the computational techniques, the regularized problem was solved using direct Newton–Raphson iterations while for the AL formulation a three-stage Picard scheme was employed with adjustment of parameters. Overall, the accuracy of the AL method was found to be superior in capturing the yield surfaces either away from the bubble or around its equatorial plane and its back side. An inherent disadvantage of the Papanastasiou regularization is that it predicts a very slow flow even beyond the yield surfaces which contradicts the physical implications of the model, although a critical Bingham number can be defined relatively accurately. On the other hand, the main drawback when using the AL method is the very high computational cost, especially when considering 2D problems. In fact, it was found that the AL was about 1.000–10.000 times slower than the Papanastasiou approximation method! This observation serves as the motivation for the present work. More specifically a modified AL method is proposed which remedies the slow convergence associated with the original AL method and simultaneously improves the accuracy in comparison to the PR method.

1.2. Other works involving the AL method and the current effort to accelerate it

Apart from the benchmark problem of the steady bubble rise, the AL method has been successfully employed in various studies involving simulations of viscoplastic fluids. Wachs [65] studied the steady Bingham flow through an eccentric annulus, motivated by engineering applications like mud or cement slurry flow. Another problem of great importance encountered in geophysics and engineering applications is the dam break flow, which has been studied numerically by Liu et al. [36]. A two–dimensional dambreak has been simulated using a regularization technique and/or an Augmented Lagrangian scheme. It was found that the two methods provided similar results. However, the AL algorithm was overall slower and was mostly used when examining the flow close to failure or when the motion was tending to stop. Another application in the geophysics area is the study of lava flows. In the work of Robertson and Kerr [49] lava is assumed to behave like a Bingham viscoplastic material flowing inside a rectangular channel and was determined via a multigrid based AL scheme. In this case, the Papanastasiou approximation could not predict true plug regions arising in the flow and the AL method is consequently preferred. Using this method, the authors could capture plug regions even for small Bingham numbers, thus pointing out that the AL has a closer connection to the physics of this fluid flow than the regularization technique.

In the recent computational rheology literature, there has been an effort for the development of accelerated algorithms for the solution of viscoplastic flows, while maintaining the high-accuracy of the results. A numerical method for faster resolution of viscoplastic fluid problems was proposed by Saramito [52], based on the Newton method. As already stated, solving the original non-regularized viscoplastic law via the Newton method produces a singular Jacobian matrix, which needs to be regularized, compromising the super-linear convergence of the method. Saramito [52] achieved the solution of the singular Jacobian problem by using a preconditioner and avoiding the need for regularization. The methodology he used is similar to that described in the work of Alart and Curnier [1]. In terms of computational cost, this damped Newton method was found to be better than the AL approach by reducing the necessary computational time, while producing equally satisfactory results in predicting the yield surface. Additionally, a quite fast and efficient modification of the AL scheme is presented in the work of Treskatis et al. [62]. They studied the steady state, creeping flow of a non-regularized viscoplastic material in a cavity. The method used a class of proximal gradient algorithms, called FISTA (Fast Iterative Shrinkage–Thresholding) algorithm, which achieves minimization of the cost function at a fast rate, while keeping the computational cost relatively low. Using this algorithm, they were able to efficiently calculate the yielded/unyielded regions. Last but not least, in a very recent publication by Bleyer [9], an interior point method is employed showing promising convergence properties even in 3-Dimensional problems, which are traditionally very costly computationally. These works highlight the fact that the need for new accelerated methods is a topic of intense research and that there are still opportunities for new ideas to emerge.

The present work explores the application of a new modification of the AL method with accelerated converge properties combined with improved accuracy of the results to several benchmark problems. This method is based on a monolithic Newton solver and is described in §2. Subsequently, in §3, it is tested with the free boundary problem of the steady rise of a bubble in a viscoplastic medium. In §4, §5 and §6, the steady entrance flow in a rectangular channel, the lid-driven cavity flow and the fully developed flow in a square duct are solved, respectively, and in §7 the transient filament stretching is examined. In all cases, our results with the PAL method are compared to those from the PR and/or the original AL either computed by us or reported in the literature. Conclusions are drawn in §8.

2. The penalized augmented Lagrangian method

The new numerical scheme is based on coupling the Augmented Lagrangian method with the penalty procedure in order to update the Lagrangian multipliers. Hence, it will be called PAL (Penalized Augmented Lagrangian) method. Our previous implementation of the AL algorithm, presented in Dimakopoulos et al. [15], involved a multistage procedure. The AL algorithm begins with the application of the Cauchy–Schwarz inequality on two symmetric, second order tensors. Thus, a point wise inequality is obtained, valid in both the solid and liquid regions. The yield stress term is replaced by its upper bound, since this term is not a differentiable function of its argument $\underline{\tilde{\dot{\gamma}}}$, when $\tilde{\dot{\gamma}} \to 0$, because of the square root in the definition of the second invariant. The variational inequality is derived by

(i) Taking the inner product of the equation of motion with the deviation of a trial velocity from the exact velocity and integrating over the entire domain, Ω,
(ii) Introducing the point-wise, variational inequality,
(iii) Performing integration by parts,
(iv) Using the incompressibility constraint and finally,
(v) Introducing the various boundary conditions on the stress or the velocity field.

In this way, the constitutive relation is combined with the equation of motion resulting in a large system of equations, which is solved iteratively. Because the variational inequality is defined over the entire flow domain, its solution simultaneously provides the location of the yield surface. This derivation resembles the standard steps taken to derive the weak form of the governing equations in the standard Galerkin/Finite Element method. Let $\underline{u}$ be the correct velocity field and $\underline{v}$ a trial velocity field satisfying only the velocity boundary conditions. Then the variational inequality for this problem reduces to

$$c(\tilde{u}, \tilde{v} - \tilde{u}) + a(\tilde{u}, \tilde{v} - \tilde{u}) + j(\tilde{v}) - j(\tilde{u}) \geq L(\tilde{v} - \tilde{u}) \tag{5}$$

In the above, c stands for the functional of the inertia terms, a is the functional for viscous dissipation, j is the yield stress dissipation functional and L is the functional of the gravitational forces in the bulk and capillary forces on gas/liquid interfaces. These are given by the following analytical expressions:

$$\begin{aligned}
c(\tilde{u}, \tilde{v} - \tilde{u}) &= \tilde{\rho} \int_\Omega (\tilde{u} \cdot \nabla \tilde{u}) \cdot (\tilde{v} - \tilde{u}) d\Omega \\
a(\tilde{u}, \tilde{v} - \tilde{u}) &= \tfrac{1}{2} \int_\Omega \tilde{\eta}\,(|\underline{\tilde{\dot{\gamma}}}(\tilde{u})|)\, \underline{\tilde{\dot{\gamma}}}(\underline{u}) : \underline{\tilde{\dot{\gamma}}}(\tilde{v} - \tilde{u}) d\Omega \\
j(\tilde{v}) &= \tilde{\tau}_y \int_\Omega (|\underline{\tilde{\dot{\gamma}}}(\tilde{v})|)\, d\Omega \\
L(\tilde{v} - \tilde{u}) &= -\int_\Omega \tilde{g}(\tilde{v} - \tilde{u}) d\Omega - \int_{\partial\Omega_b} 2H \underline{n} \cdot (\tilde{v} - \tilde{u}) ds
\end{aligned} \tag{6}$$

where $\tilde{\eta}(|\underline{\tilde{\dot{\gamma}}}(\underline{u})|)$ is the plastic viscosity and $\partial\Omega_b$ the free surface boundary. The dependence of the plastic viscosity on $\underline{\tilde{\dot{\gamma}}}$ makes the former non-differentiable as well, in Herschel-Bulkley fluids. The inertia and yield stress functionals are nonlinear in the unknown velocity. The inequality is satisfied, if the following functional is minimized:

$$J(\underline{\tilde{u}}) = c(\underline{\tilde{u}}, \underline{\tilde{u}}) + \frac{1}{n+1}a(\underline{\tilde{u}}, \underline{\tilde{u}}) + j(\underline{\tilde{u}}) - L(\underline{\tilde{u}}) \tag{7}$$

The singularity introduced by the second invariant of the rate-of-strain tensor $\underline{\tilde{\dot{\gamma}}}$, is approximated by employing a projection tensor $\underline{\tilde{q}}$. The equivalence to the previous functional is imposed by introducing a Lagrangian tensor $\underline{\tilde{\lambda}}$. This leads to a new functional containing $\underline{\tilde{\lambda}}$:

$$G(\underline{\tilde{u}}, \underline{\tilde{q}}, \underline{\tilde{\lambda}}) = J(\underline{\tilde{u}}) + (\underline{\tilde{\lambda}}, \underline{\tilde{\dot{\gamma}}} - \underline{\tilde{q}}) \,, \; (\underline{\tilde{\lambda}}, \underline{\tilde{\dot{\gamma}}} - \underline{\tilde{q}}) = \int_{\Omega} \underline{\tilde{\lambda}} : (\underline{\tilde{\dot{\gamma}}} - \underline{\tilde{q}}) d\Omega \tag{8}$$

The convergence of this iterative scheme is extremely slow, however. A significant acceleration can be achieved by adding a quadratic term where an adjustable, positive augmentation parameter is introduced, $\tilde{r} \geq 0$, hence the term "Augmented Lagrangian" method:

$$G_r(\underline{\tilde{u}}, \underline{\tilde{q}}, \underline{\tilde{\lambda}}) = G(\underline{\tilde{u}}, \underline{\tilde{q}}, \underline{\tilde{\lambda}}) + \frac{1}{2}\tilde{r}\int_{\Omega} |\underline{\tilde{\dot{\gamma}}} - \underline{\tilde{q}}|^2 d\Omega \tag{9}$$

The solution procedure by Dimakopoulos et al. using ALG2 [15] consisted of 3 distinct stages. The new idea in the present work is to relax the constraint of stage 3 (i.e. the updating of the Lagrange multiplier in the solution procedure) in the aforementioned paper, using a penalty method, thus enabling us to obtain results quickly and to the required accuracy. Moreover, in this way, the multistage approach of [15] is avoided and a monolithic scheme is applied based on Newton–Raphson iterations. In short, our application consists of writing the continuity and momentum balances, along with the constitutive model, and modifying them by introducing penalty terms, so that the diagonal elements of the Jacobian do not become equal to zero and thus leading to a destructive singularity. The idea of penalizing equations to avoid numerical difficulties is quite common. For example, Bercovier and Engelman [6] have used a penalty method for the continuity equation in a Bingham fluid problem, but this approach has been used even for Newtonian fluids. In fact, the addition of the quadratic term to $G(\underline{\tilde{u}}, \underline{\tilde{q}}, \underline{\tilde{\lambda}})$ to obtain $G_r(\underline{\tilde{u}}, \underline{\tilde{q}}, \underline{\tilde{\lambda}})$ and the AL method has been called penalization in [21, 28]. However, because of the functional $(\underline{\tilde{\lambda}}, \underline{\tilde{\dot{\gamma}}} - \underline{\tilde{q}})$, the minimization can be achieved without making the parameter $\tilde{r}$ tend to zero, which, using ordinary penalization methods, would have the effect of causing a deterioration in the conditioning of the system to be solved. The same authors have called the ALG procedure "regularization" of the original problem. In recent years however, this term has been used very extensively when the constitutive model is directly modified, resulting in the Papanastasiou, biviscosity and Bercovier and Engelman models. Following this trend and to keep things clear, we consider our new method a modification of the AL procedure and not a regularization method.

The new method will be presented with the equations in dimensional form, because different nondimensionalization is required for the specific problems to be examined. Clearly, the new method can be easily implemented on any other problem. The continuity equation and the momentum conservation equation, written in terms of the Cauchy stress tensor, are:

$$\tilde{\nabla}\cdot\underline{\tilde{u}} = 0 \tag{10}$$

$$\tilde{\rho}\frac{D\underline{\tilde{v}}}{D\tilde{t}} = \nabla\cdot\underline{\tilde{\sigma}} + \tilde{\rho}\,\underline{\tilde{g}} \tag{11}$$

The total stress tensor is split into the pressure and the extra stress tensor:

$$\underline{\tilde{\sigma}} = -\tilde{p}\,\underline{I} + \underline{\tilde{\tau}} \tag{12}$$

The extra stress tensor is further split into two "parts". One part is the Lagrangian tensor $\underline{\tilde{\lambda}}$, and the other part is a term involving a difference between the "real" rate-of-strain tensor $\underline{\tilde{\dot{\gamma}}}$ and the projection of the rate-of-strain tensor $\underline{\tilde{q}}$:

$$\underline{\tilde{\tau}} = \tilde{r}(\underline{\tilde{\dot{\gamma}}} - \underline{\tilde{q}}) + \underline{\tilde{\lambda}} \tag{13}$$

Upon convergence, the Lagrangian tensor should be equal to the actual extra stress tensor, satisfying both the continuity and the momentum equations, since the "difference" term disappears. There are two algorithm-related parameters in PAL. The first one is a so-called augmentation parameter $\tilde{r}$ which also appears in the original AL method and it multiplies the "difference" term in Eq. (13). The elements of the projected rate-of-strain tensor are computed by satisfying the constitutive model:

$$\underline{\tilde{q}} = \max\left(0, \frac{|\underline{\tilde{\tau}}| - \tilde{\tau}_y}{\tilde{k}}\right)^{1/n} \frac{\underline{\tilde{\tau}}}{|\underline{\tilde{\tau}}|} \tag{14}$$

Note that the same expression of the extra stress tensor, Eq. (13), is used in both the momentum and the constitutive model equations. The only thing remaining now is to force the $(\underline{\tilde{\dot{\gamma}}} - \underline{\tilde{q}})$ term, "difference" or "projection" term, to vanish and force the Lagrange multipliers tensor to become the actual extra stress tensor. Since the constitutive model equations were used to compute the projected rate-of-strain tensor, we need an equation for the Lagrange multipliers tensor. To this end, we introduce the second algorithm-related parameter, the so-called penalty parameter, denoted by $\tilde{r}_p$, which determines the accuracy of the entire scheme via the "strictness" of the penalty method. This is attained by a penalty-type equation:

$$\underline{\tilde{\lambda}} = \tilde{r}_p(\underline{\tilde{\dot{\gamma}}} - \underline{\tilde{q}}) \tag{15}$$

For sufficiently large values of $\tilde{r}_p$, since $\underline{\tilde{\lambda}}$ is finite, the difference $(\underline{\tilde{\dot{\gamma}}} - \underline{\tilde{q}}) \to 0$ and, therefore the Lagrangian tensor converges to the true extra stress tensor, since it should simultaneously satisfy both the momentum and constitutive model equations. One last thing to note is that, as we will see shortly, this so-called penalty parameter generates similar results to a regularization parameter, if a direct comparison to the Papanastasiou method is made. Therefore, it is also interesting to explore if the new method is superior to the Papanastasiou regularization.

The grid is generated by solving a system of quasi-elliptic partial differential equations, Eqs. (16) and (17), in order for the mesh to follow any deformation of the domain occupied by the liquid [16]. These equations determine the connection between the node coordinates in the physical domain (X, Y) to those defined in a simpler (usually fixed) computational one (ξ, ζ). Indicatively, they are given below for Cartesian coordinates and in dimensionless form:

$$\underline{\nabla}\cdot\left\{\left(\varepsilon_1\sqrt{\frac{X_\zeta^2 + Y_\zeta^2}{X_\xi^2 + Y_\xi^2}} + (1-\varepsilon_1)\right)\underline{\nabla}\zeta\right\} = 0 \tag{16}$$

$$\underline{\nabla}\cdot\underline{\nabla}\xi = 0 \tag{17}$$

In the above, ε_1 is an adjustable parameter to optimally adjust the coordinate orthogonality to mesh informity in each problem.

The entire equation set is solved using the mixed finite element method for the discretization of the velocity, node positions in the mesh, pressure and stress fields. The velocity and position vectors are approximated with 6-node Lagrangian basis functions, while the stress (or the Lagrange multipliers) and pressure fields by 3-node basis functions. The procedure has been described in Dimakopoulos and Tsamopoulos [15] and successfully used in many studies, e.g. Pettas et al. [46], Chatzidai et al. [11], Papaioannou et al. [42], Fraggedakis et al. [27], and Dimakopoulos et al. [19] where the material undergoes large deformations and free surfaces are present. Recently, Fraggedakis et al. [26], extended this method to three-dimensional flows.

Using the PAL scheme, the residual vector $\underline{Res}$ consists of the weak form residuals of the governing equations. More details on the form of the residuals can be found in previous works , e.g. [41, 63]

$$\underline{Res} = \begin{bmatrix} \underline{Res}^M \\ \underline{Res}^G \\ \underline{Res}^C \\ \underline{Res}^Q \\ \underline{Res}^\Lambda \end{bmatrix} \tag{18}$$

The primary unknown of the momentum residuals $\underline{Res}^M$ is the velocity field $\underline{v}$, of the grid residuals, $\underline{Res}^G$ is the node position vector $\underline{x}$, of the continuity residuals $\underline{Res}^C$ is the pressure $\tilde{p}$, of the constitutive model residuals$\underline{Res}^Q$ is the projection rate of deformation tensor $\underline{\tilde{q}}$, and of the Lagrangian residuals $\underline{Res}^\Lambda$ are the entries of the Lagrangian stress tensor $\underline{\tilde{\lambda}}$. Therefore, the Jacobian matrix will consist of blocks and will have the following form:

$$\begin{array}{c} \begin{matrix} \tilde{\underline{v}} & \tilde{x} & \tilde{p} & \tilde{\underline{q}} & \tilde{\underline{\lambda}} \end{matrix} \\ \begin{bmatrix} \frac{\partial \underline{Res}^M}{\partial \tilde{\underline{v}}} & \frac{\partial \underline{Res}^M}{\partial \tilde{x}} & \frac{\partial \underline{Res}^M}{\partial \tilde{p}} & \frac{\partial \underline{Res}^M}{\partial \tilde{\underline{q}}} & \frac{\partial \underline{Res}^M}{\partial \tilde{\underline{\lambda}}} \\ \underline{0} & \frac{\partial \underline{Res}^G}{\partial \tilde{x}} & \underline{0} & \underline{0} & \underline{0} \\ \frac{\partial \underline{Res}^C}{\partial \tilde{\underline{v}}} & \frac{\partial \underline{Res}^C}{\partial \tilde{x}} & \underline{0} & \underline{0} & \underline{0} \\ \underline{0} & \frac{\partial \underline{Res}^Q}{\partial \tilde{x}} & \underline{0} & \frac{\partial \underline{Res}^Q}{\partial \tilde{\underline{q}}} & \frac{\partial \underline{Res}^Q}{\partial \tilde{\underline{\lambda}}} \\ \underline{0} & \frac{\partial \underline{Res}^\Lambda}{\partial \tilde{x}} & \underline{0} & \frac{\partial \underline{Res}^\Lambda}{\partial \tilde{\underline{q}}} & \frac{\partial \underline{Res}^\Lambda}{\partial \tilde{\underline{\lambda}}} \end{bmatrix} \begin{matrix} \tilde{\underline{v}} \\ \tilde{x} \\ \tilde{p} \\ \tilde{\underline{q}} \\ \tilde{\underline{\lambda}} \end{matrix} \end{array} \tag{19}$$

It can be clearly seen that with the additional Eq. (15), the only zero entry in the diagonal of the Jacobian is the one generated by the continuity equation. The choice of the PAL parameters, the augmentation parameter, $\tilde{r}$ and the penalty parameter, $\tilde{r}_p$, is crucial and is made considering the following. Ideally, the "projection" term vanishes under two conditions: (i) $\tilde{r} \to 0$ or (ii) $\tilde{r}_p \to \infty$ (or both happening simultaneously). Although these extreme values seem tempting to use, they do not allow the method to converge, because the original discontinuous problem is recovered. For the computations in this work, the assumption that was made with ALG2 [15] is followed, where $r = 1$ (dimensionless for that matter), since it plays a more neutral role in the accuracy of the scheme and its actual value does not affect the solution as explained earlier. In fact, we have performed some tests with smaller values of r, e.g. $r = 0.1$, and we found that this parameter has a two-fold impact. First of all, the smaller r is, the "faster" the contribution of $\underline{\dot{\gamma}} - \underline{q}$ vanishes in the first iterations. A Newton iteration procedure is more likely to diverge when r attains small values, since the Jacobian matrix becomes more sensitive to disturbances of the two tensors $\underline{\dot{\gamma}}\&\underline{q}$. In addition, we found out that for smaller r, we could not significantly raise the value of the penalty parameter r_p as desired for better accuracy. Ultimately, it was seen that more accurate solutions were obtained by increasing the penalty parameter with a fixed r parameter rather than decreasing r for a fixed r_p. Thus, we used $r = 1$ throughout the rest of this study and varied only the penalty parameter.

A comparison of our new method to the ones proposed by Saramito [52] and Treskatis et al. [62] recently in order to accelerate the original AL method without resorting to regularization of the constitutive model, can be found in Table 1. In what follows, the proposed method is put to the test by solving four steady state problems and one transient problem and comparing the results to other studies and numerical methods using either the PR or the AL method.

A brief description is given here of the test cases examined in this paper. Each case has a different geometry and/or driving force generating the flow. Schematically they are given in Fig. 1:(A) the buoyancy-driven steady rise of a gas bubble of constant volume $\widetilde{V}_b$ and equivalent radius $\widetilde{R}_b$, in viscoplastic quiescent fluid, (B) the steady entrance flow of a viscoplastic material with plug velocity profile ($\widetilde{U}_o$) across the entrance of a Cartesian channel having a half - width equal to $\widetilde{L}$, (C) the steady flow of a viscoplastic fluid in a square cavity with side lengths equal to $\widetilde{L}$, due to horizontal motion of its lid with constant velocity $\widetilde{U}_o$, (D) the steady fully-developed flow of a viscoplastic fluid in a duct with square cross-section, driven by a constant pressure gradient, $\widetilde{f} \equiv \underline{e}_z \cdot \nabla \widetilde{P}$, in the $\tilde{z}$-direction, and (E) the transient stretching of viscoplastic filament with initial uniform radius equal to $\widetilde{R}$ and height equal to $\widetilde{L}_o$, and a constant stretching velocity $\widetilde{U}_o$. The governing equations Eqs. (10)–((17)) of each problem is non-dimensionalized by a group of scales summarized in Table 2, which leads to the formation of a set of dimensionless numbers, shown in Table 3. All variables have been defined already except for the viscocapillary time-scale, which is defined as $\tilde{t}_{vc} = (\tilde{k}\widetilde{R}/\tilde{\gamma})^{1/n}$.

The common characteristic in each case is that irrespective of the scaling that is followed, the Bingham number can be defined and thus the dimensionless constitutive equation for the projection of the rate of strain rate, Eq. (14), becomes:

$$\max(0\,,\ |\underline{\tau}| - Bn)^{1/n} \frac{\underline{\tau}}{|\underline{\tau}|} = \underline{q} \tag{20}$$

The momentum and continuity equations are scaled accordingly following the non-dimensionalization of each case into account. Each section that follows presents a brief description of the problem along with boundary conditions.

3. Bubble rise

The first problem examined is the steady rise of a bubble in a Viscoplastic fluid. This problem has been solved both for a Bingham fluid by Tsamopoulos et al. [63] using the PR method and by Dimakopoulos et al. [15] using both the PR and the AL methods for Bingham and HB fluids. We compare our results with those obtained by Dimakopoulos et al. [15] for a HB fluid for relatively large exponents of PR. Thus these results are ideal for a thorough comparison with the present results. The dimensionless bubble velocity and the drag coefficient are plotted for $(Ar,\ Bo,\ n) = (1,\ 20,\ 0.7)$ and for different Bingham numbers. The results were obtained using a First-Order Continuation (FOC) scheme. In order to compare on equal basis with the computation times of Dimakopoulos et al. [15], the mesh initially consisted of 50 and 100 elements in the azimuthal and radial direction respectively, but had multiple refinement stages. The elements are doubled once up to $r = 4$ in both directions, then once again up to $r = 2$ and finally once more up to $r = 1.25$. This is exactly the mesh used in [15]. The same convergence criterion is also adopted in the current work by terminating the Newton-Raphson iterations, if the residual norm drops below 10^{-9}. Two critical quantities of interest are the rise velocity of the bubble, which determines whether the bubble can be considered motionless and the drag coefficient, defined by the following expression:

$$C_d = \frac{2\widetilde{F}}{\tilde{\rho}\widetilde{U}^2 \pi \widetilde{R}_b^2} = \frac{2F}{\pi Ar U^2} = \frac{8}{3Ar U^2} \tag{21}$$

This last expression results from substituting the force exerted on the bubble by buoyancy; see [15]. Following the aforementioned work, the augmentation parameter is kept constant and equal to $r = 1$, while the value of the penalty parameter is varied. Results with three different values of r_p are presented in Fig. (2).

As Bn increases, the fluid attains a more "plastic" like nature. Therefore, material yielding is obstructed more, until a critical Bingham number is reached where $U_{rise} \to 0$ and the bubble is entrapped inside the fluid. This also increases the drag coefficient on the bubble via Eq. (21). Both the drag coefficient and the rise velocity predictions depend on the choice of r_p in a manner similar to their dependence on the Papanastasiou exponent. Theoretically, as $r_p \to \infty$, the solution obtained will be more accurate, i.e., closer to the predictions of the original AL method. It turns out that for this particular problem, the

Table 1
Comparison of the method proposed by Saramito [52] and Treskatis et al. [62] to our method.

	Main solution idea	Finite elements space	Solver	Tested cases
Saramito (2016) [52]	Use a preconditioner for the singular Jacobian matrix	Standard Taylor Hood elements for $\underline{v}$ and p & linear discontinuous approximations of stresses	Newton Raphson & Iterative solver for the linearized problem	Flow in a square duct
Treskatis et al. (2016) [62]	Accelerated dual proximal gradient method	P_1-iso-P_2/P_1 element	Stage – wise solution procedure using conjugate gradients.	Lid-driven cavity
Present work	Augmented Lagrangian method along with penalization of the augmented equations	Standard Taylor Hood elements for all variables	Sparse direct solver or iterative solver for the NR solution of the linearized problem (e.g., PARDISO [45], Watson [31] & MUMPS [2])	Bubble rise (A) Channel entrance (B) Lid-driven cavity (C) Square duct (D) Filament stretch (E)

solution obtained with the PAL method gives similar results with the Papanastasiou regularization for a penalty parameter one order of magnitude lower than the corresponding Papanastasiou exponent. Examining Fig. 2, we observe that the rise velocity decreases linearly in the semi-log scale for low Bingham numbers, however as the critical Bn is approached it decreases abruptly. This is caused by the quicker approach of the yield surface to the bubble surface. Similarly, the PAL method predicts a very steep increase of the drag coefficient, which, as the penalty parameter increases, asymptotically approaches infinity at the critical Bn. However, both the PR and the PAL methods predict that after the critical number there is a second change of the slope of the curve where both the velocity and the drag coefficient vary at a slower rate and seem to approach an asymptotic value, which is not physically acceptable. In fact, $U_{rise} \to m$ and $C_d \to M$ where m and M are a very small and a very large number, respectively, and they depend on the choice of the penalty parameter in PAL or the exponent in PR. The performance of the algorithm is evaluated through the required Newton–Raphson iterations for convergence, following the aforementioned continuation procedure. For the first order continuation on Bingham number, a step of $\Delta Bn = 10^{-3}$ was used in order to have a "good" initial guess. The number of iterations is reported for $0 \leq Bn \leq 0.16$, i.e., up to slightly larger than the critical Bn, accounting for a total of 160 continuation steps.

As the penalty criterion becomes "stricter", the required iterations increase, but not very much. For $r_p = 10^4$ each steady state solution for a different Bn required about 10 iterations, for $r_p = 10^5$ the iterations increased to about 20 and for $r_p = 2.5 \cdot 10^5$ they increased further to about 25-30. Moreover, as the penalty parameter increases, the required iterations for convergence fluctuate more intensely. Another practical result that is shown in Fig. 3 is the total elapsed time for convergence of the NR scheme for representative Bingham numbers. It can be seen that the required time per NR iteration is pretty consistent and close to a half minute. The differences in the total time arise due to the difference in the number of required iterations.

It is also interesting to present how the magnitude of the norm of the residual vector, which provides the convergence criterion, evolves with each subsequent iteration. Fig. 4 shows this dependence for a typical value of $Bn = 0.05$, which is in the region where the iterations for convergence present the maximum spikes. The L_2 norm of the residual is defined as

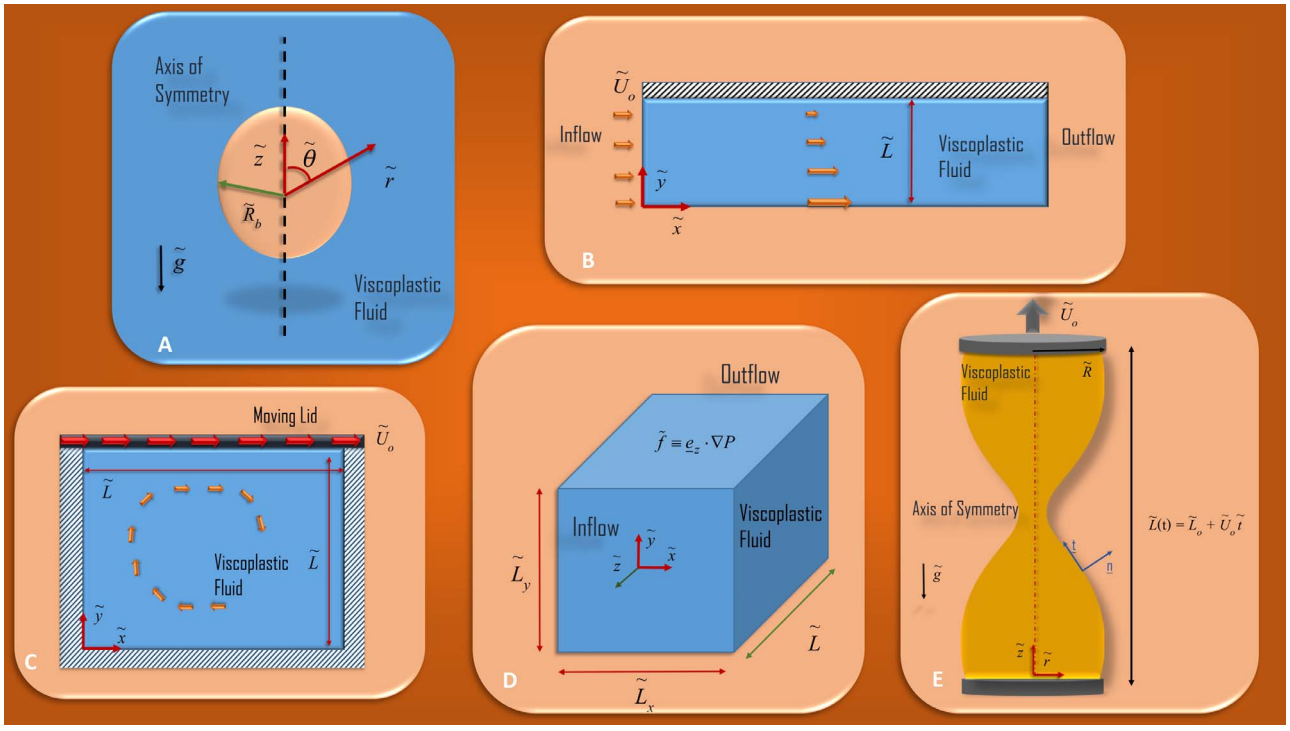

Fig. 1. Problems solved in the paper in order (A–E).

Table 2
The characteristic scales arising in each of the five test cases.

Flow problem\Mechanism& variable	Length	Velocity	Stress	Mechanism generating the flow	References
Bubble rise (A in Fig. 1)	$\widetilde{R}_b$	$(\widetilde{\rho}\widetilde{g}\widetilde{R}_b^{n+1}/\widetilde{k})^{1/n}$	$\widetilde{\rho}\widetilde{g}\widetilde{R}_b$	Buoyancy, $\widetilde{g}$	[15,63]
Channel entrance (B in Fig. 1)	$\widetilde{L}$	$\widetilde{U}_o$	$\widetilde{k}(\widetilde{U}_o/L)^n$	Plug velocity field at entrance, $\widetilde{U}_o$	[47]
Lid-driven cavity (C in Fig. 1)	$\widetilde{L}$	$\widetilde{U}_o$	$\widetilde{k}(\widetilde{U}_o/L)^n$	Moving lid with constant velocity, $\widetilde{U}_o$	[58-60]
Square duct (D in Fig. 1)	$\widetilde{L}$	$\widetilde{L}(\widetilde{f}\widetilde{L}/\widetilde{k})^{1/n}$	$\widetilde{f}\widetilde{L}$	Pressure gradient $\widetilde{f} \equiv \varrho_z \cdot \nabla\widetilde{P}$	[52]
Filament stretch (E in Fig. 1)	$\widetilde{R}$	$\widetilde{R}/(\widetilde{k}\widetilde{R}/\widetilde{\gamma})^{1/n}$	$\widetilde{k}/\widetilde{t}_{ve}^n$	Moving upper plate with constant velocity, $\widetilde{U}_o$	[3-4]

Table 3
The dimensionless numbers that arise in each case.

Flow Problem	*Bn* (Bingham)	*Re* (Reynolds)	*Ar* (Archimedes)	*Bo* (Bond)
Bubble rise (A in Fig. 1)	$\widetilde{\tau}_y/\widetilde{\rho}\widetilde{g}\widetilde{R}_b$		$\left(\widetilde{\rho}^2\widetilde{g}^{(2-n)}\widetilde{R}_b^{(n+2)}/k^2\right)^{1/n}$	$\widetilde{\rho}\widetilde{g}\widetilde{R}_b^2/\widetilde{\gamma}$
Channel entrance (B in Fig. 1)	$\widetilde{\tau}_y/\widetilde{k}(\widetilde{U}_o/L)^n$	$\widetilde{\rho}\widetilde{U}_o^{2-n}\widetilde{L}^n/\widetilde{k}$		
Cavity (C in Fig. 1)	$\widetilde{\tau}_y/\widetilde{k}(\widetilde{U}_o/L)^n$	$\widetilde{\rho}\widetilde{U}_o^{2-n}\widetilde{L}^n/\widetilde{k}$		
Square duct (D in Fig. 1)	$\widetilde{\tau}_y/\widetilde{f}\widetilde{L}$	$\widetilde{\rho}\widetilde{U}_o^{2-n}\widetilde{L}^n/\widetilde{k}$		
Filament stretch (E in Fig. 1)	$\widetilde{\tau}_y\widetilde{R}/\widetilde{\gamma}$	$(\widetilde{\rho}\widetilde{R}^3/\widetilde{\gamma})(\widetilde{\gamma}/\widetilde{k}\widetilde{R})^{2/n}$		$\widetilde{\rho}\widetilde{g}\widetilde{R}^2/\widetilde{\gamma}$

$$||\underline{Res}||_{L_2} = \sqrt{\underline{Res}\cdot\underline{Res}} < tolerance \tag{22}$$

Even in the first iteration, the residual norm is relatively low, indicating that the FOC scheme produces a good initial guess for the next continuation step. In a semi-logarithmic plot, the residual norm decreases linearly with the number of elapsed iterations. Close to convergence with residual in the order of $O(10^{-8})$, it decreases at a slower rate for a few iterations until it ultimately drops below the desired threshold for convergence. More practically, it is interesting to evaluate the real speed-up of the algorithm. In real time, the results for $r_p = 10^4$ were obtained in about 10 h, those of $r_p = 10^5$ in about 20 h and those of $r_p = 2.5{\cdot}10^5$ roughly in a day. The computations were performed on a single computational node. For the record, Dimakopoulos et al. [15] needed about 1 day with $N = 10^5$, 7 days for $N = 10^6$ and 32 days (!) using the Augmented Lagrangian Method, because it required several thousands of iterations per step. Therefore the results with the new method have the same precision and are obtained in a day, as those obtained in a week with the PR, while the AL achieves superior accuracy, but takes much more time.

As it was previously stated, an important aspect in any study of viscoplastic fluid flow is capturing the yield surface accurately. In [15], it was shown that the Papanastasiou approximation overestimated the yield surface, and a more accurate solution was reached as the Papanastasiou exponent increased. Physically this is anticipated, because increasing N, makes the unyielded material more solid-like, which in turn decelerates the bubble motion. In Fig. 5, the yield surface around the bubble along with the bubble shape are compared to the predictions of the original AL method and the strictest Papanastasiou case.

In the left panel of Fig. 5, the similarity of the yield surfaces and the bubble shapes is quite remarkable. The yield surface forms the usual envelope surrounding the bubble and a larger yielded region at the equatorial plane of the bubble, for reasons explained in [15]. The size and shape of the bubble are also nearly identical. The prediction of the Papanastasiou regularization in the right panel is somewhat worse than that by the PAL method, as it predicts a larger yield surface. The appearance of an unyielded region on the bubble surface around the equatorial plane is observed in all cases. With regard to the two velocity components, their profiles obtained by the PAL method are compared with those of Dimakopoulos et al. [15] for the Papanastasiou regularization.

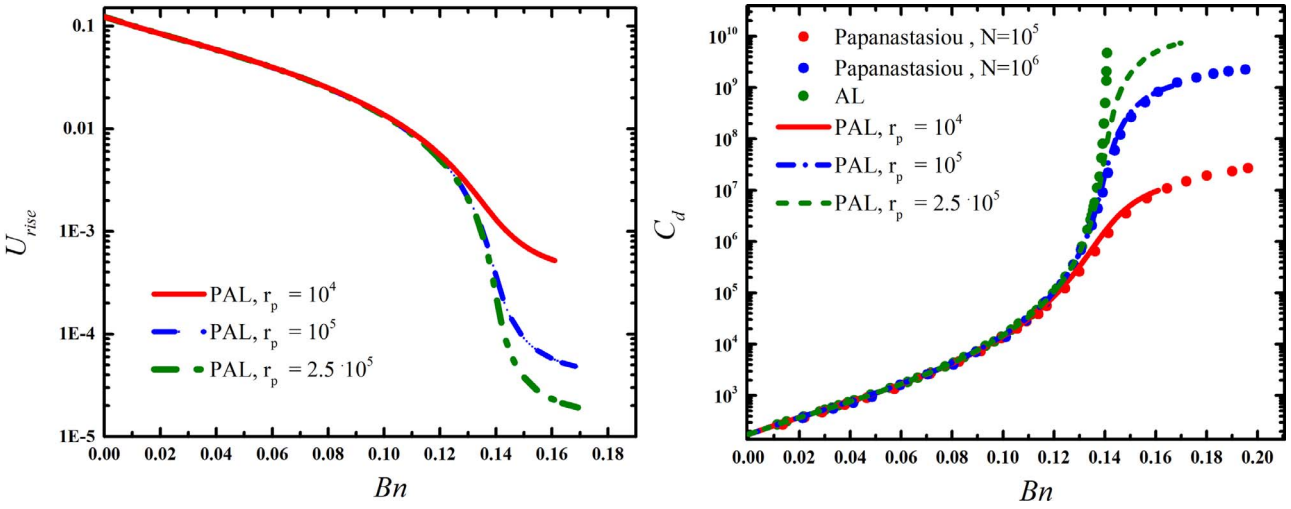

Fig. 2. Dimensionless rise velocity and drag coefficient for a rising bubble vs. Bn. The other parameters are $(Ar,\ Bo,\ n) = (1,\ 20,\ 0.7)$.The current results are compared to those of Dimakopoulos et al. [15].

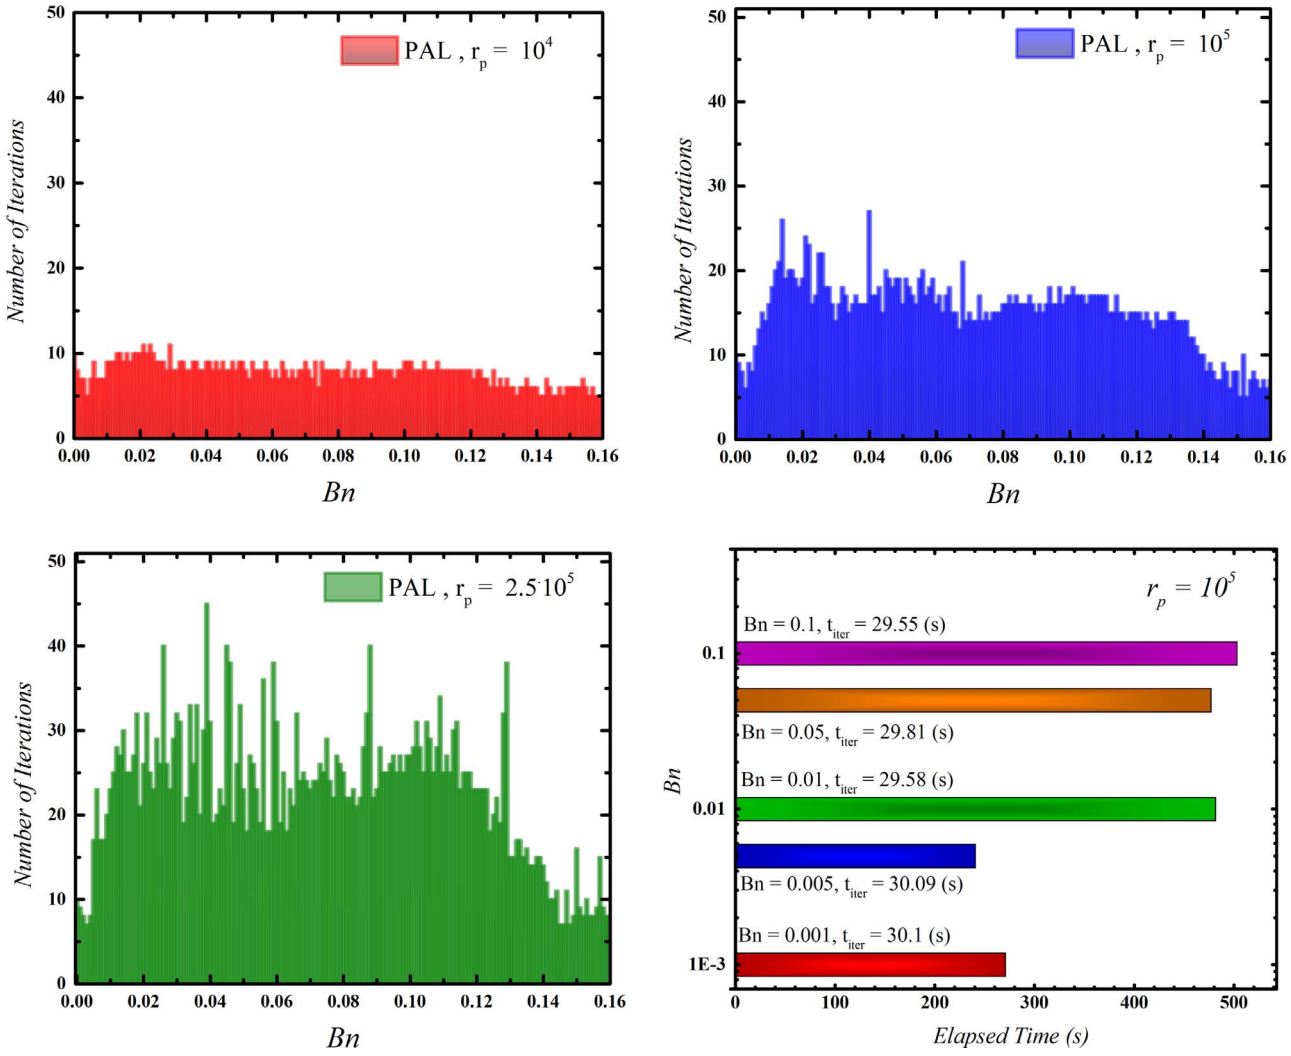

Fig. 3. Newton–Raphson iterations and elapsed time for the PAL method results of Fig. 2.

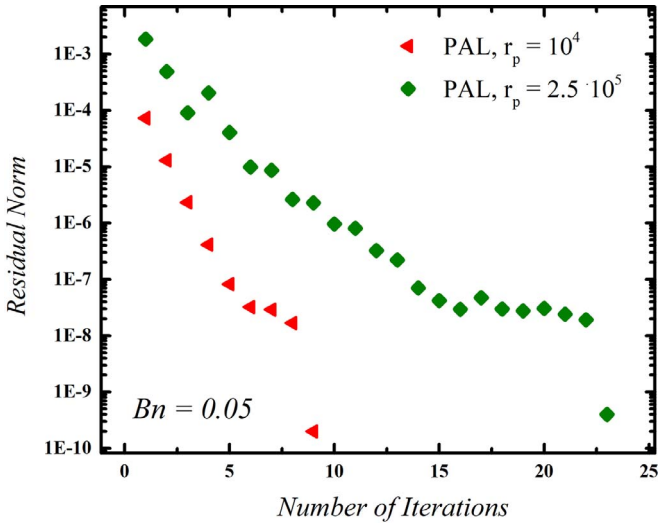

Fig. 4. Decrease of the residual norm at each NR iteration for $(Ar,\ Bo,\ n,\ Bn) = (1,\ 20,\ 0.7,\ 0.05)$, i.e. in one step of the computations in Fig. 3.

The direct comparison of the PAL method velocity predictions with $r_p = 2.5 \cdot 10^5$ to the PR ones with $N = 10^6$ close to the bubble entrapment conditions, i.e. when the velocities decrease considerably and the deviations between the two methods are expected to be maximized, reveals that they are qualitatively similar but not quantitatively. As observed in Fig. 6, the PAL method predicts a smaller value for both the azimuthal and the radial velocity. The azimuthal velocity at the equatorial plane near the bubble and the radial velocity near the axis of symmetry predicted by the PAL method are 35% lower than the corresponding ones predicted by PR. Since the results depicted here portray the conditions near criticality, these results show that indeed the PAL method produces more accurate results.

4. Channel entrance

This is a particularly important problem due to the fact that confined flows either in straight or corrugated channel allow the testing of various phenomena, such as in Tsouka et al. [64]. The boundary conditions impose plug flow at the inlet, no slip and no penetration at the upper plate and symmetry at the mid-plane. Finally, at the exit boundary, the Open Boundary Condition (OBC) is employed as discussed in Papanastasiou et al. [44] and Dimakopoulos et al. [18]. Inertia is neglected, $Re = 0$, because our main interest is in determining the yield surfaces. The total length of the long channel is $x = 10$ and is divided into three regions extending, respectively, up to $x_1 = 2$, $x_2 = 6$ and$x_3 \equiv x = 10$. The first region, where the flow develops, is the most important one and is discretized using 130 elements in the x direction. Each of the other two regions is discretized using 50 elements in each

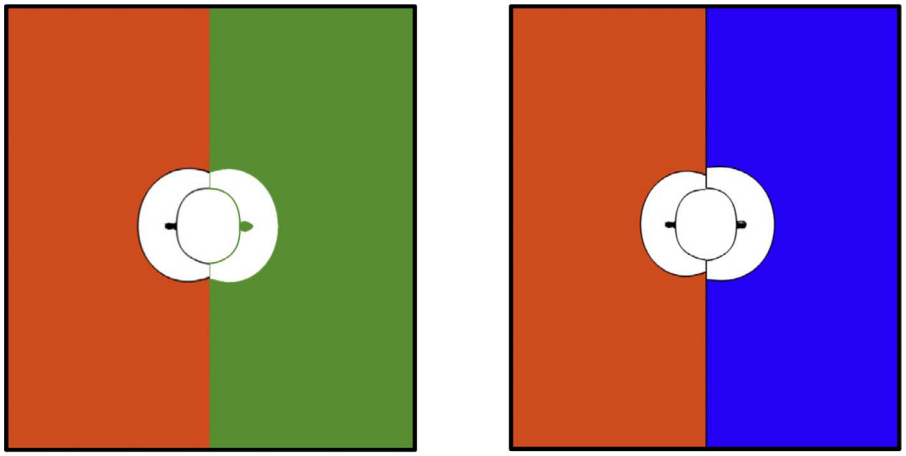

Fig. 5. Comparison of yield surfaces and bubble shapes predicted by the PAL method to those predicted by Dimakopoulos et al. [15] near the critical *Bn*. The left panel shows the comparison of the results with the original AL method (left half) to those from the present work with $r_p = 2.5{\cdot}10^5$(right half). The right panel is the comparison of AL (left half) to Papanastasiou results (right half) with $N = 10^6$ (from [15]). The parameters are $(Bn, Ar, Bo, n) = (0.135, 1, 20, 0.7)$.

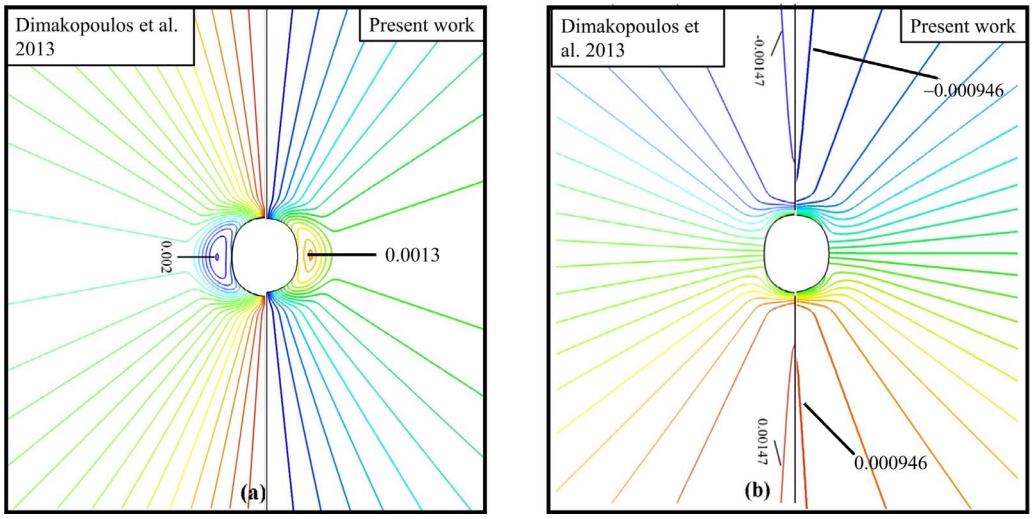

Fig. 6. Comparison of the azimuthal (a) and radial (b) velocity profiles predicted by the PAL method and the Papanastasiou regularization with $N = 10^6$. The parameters are $(Bn, Ar, Bo, n) = (0.135, 1, 20, 0.7)$.

region in the *x*-direction. Throughout, the *y*-direction is tessellated using 100 elements. The flow has been recently studied by Philippou et al. [47].

Fig. 7 depicts the yield surfaces for various Bingham numbers and methods of solution. For the PAL method, the augmentation parameter was set equal to $r = 1$, as in the bubble case, and the penalty parameter is given in the figure caption. For the Papanastasiou approximation, the dimensionless exponent is set equal to $N = 500$, similar to the values

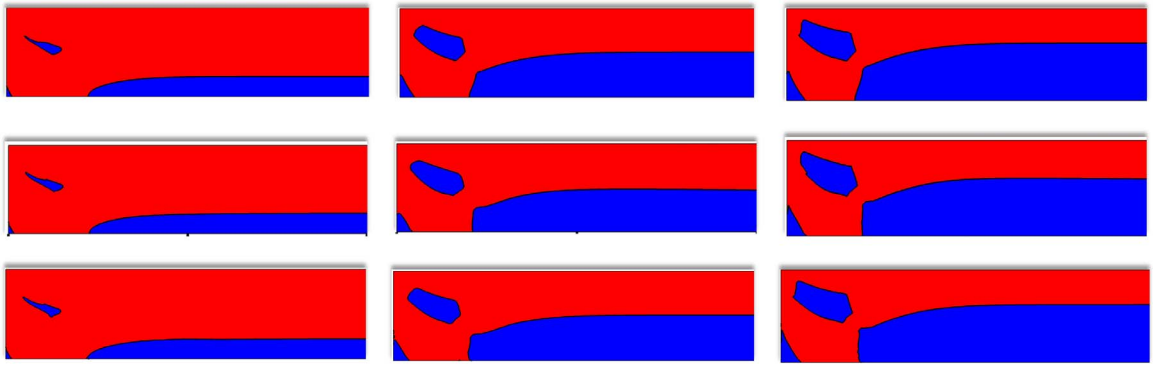

Fig. 7. Yielded (red) and unyielded (blue) regions of entrance flow in a rectangular channel for three values of *Bn*. The results are compared with the predictions of the Papanastasiou model. The first row of panels includes results of the Papanastasiou approximation, the second-row results of the PAL method with $r_p = 10^3$and the third row results of the PAL method with $r_p = 10^4$. Each column corresponds to $Bn = (1, 5, 10)$respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

used in the work of Karapetsas and Tsamopoulos [35] for the squeeze flow of a viscoplastic fluid. In each panel, there are three unyielded regions. There is the "core" unyielded region in the fully developed part of the flow, an unyielded "island" closer to the entrance and between the mid-plane and the wall and a small region at the entrance and at the symmetry plane, which is hardly visible for $Bn = 1$. As expected, an increase in Bingham number causes all three unyielded zones to increase in size. The solution using either method produces similar results for the same Bingham number. The shapes, locations, and sizes of the yield surfaces are very much alike. One noticeable difference is observed for larger Bingham numbers. Here the "core" unyielded region intersects the symmetry plane almost vertically in the case of the PAL method, but not so with the PR method.

In order to obtain the above results a zero order continuation scheme on Bingham number was implemented, where the initial guess for the next step is simply the solution of the previous step. The Bingham number is varied in the range $0 \le Bn \le 10$. For the Papanastasiou method, a value of $N = 500$ was used which is generally considered small. Here in order for the method to converge, the continuation step could not exceed$\Delta Bn = 10^{-2}$, and it had to be even lower in order to increase this Papanastasiou exponent. Alternatively, in order for the PR to converge with a step of $\Delta Bn = 10^{-1}$ we had to use $N \sim 100$ or lower. Another approach would have been to initialize our computations for a specific value of Bn and low N exponent and perform a continuation on the latter, but this is beyond the scope of this work. The NR procedure converged in a fixed number of 7 iterations per step with PR, while PAL needed more iterations per continuation step, as seen in Fig. 8. However, the PAL algorithm had better convergence properties overall, as the continuation step was an order of magnitude larger, $\Delta Bn = 10^{-1}$. Therefore, the PAL method needed a total of 100 continuation steps, while the Papanastasiou regularization method required 1000 steps. The corresponding number of iterations can be seen in the Fig. 8.

The PAL method, using $r_p = 10^3$ required on average 15 iterations per Bingham step and for $r_p = 10^4$ about 20 iterations (Fig. 8). This number of iterations follows the same pattern as in the bubble problem. Ultimately, the real time required for the entire set of continuations in Bingham number was $\sim$ 7 h using the PAL method and about $\sim$25 h using the Papanastasiou approximation method, due to the much lower continuation step. It is worth noting that when the Papanastasiou exponent is increased to $N = 10^3$, the scheme diverged after several steps. Fig. 9 shows the required time for a specific Bingham to be reached, during the continuation procedure, for the cases of PR with $N = 500$ and PAL with $r_p = 10^3$. In addition, Fig. 9 shows the dependence of the norm of the residuals on the number of iterations. For the Papanastasiou method, the convergence is smooth and decreases roughly an order of magnitude with each iteration, while for the PAL method the residual norm varies linearly with an abrupt decrease near convergence.

The main reason for the larger computation time with the PR method is the well-known problem of lack of convergence of the Newton–Raphson scheme, when the exponential factor in PR increases. The only meaningful basis of comparison of the two methods is to get results which are as close as possible to the "real" solution for the minimum time needed for computations for a given range of Bingham numbers. Under this perspective again, the PAL method was found to be superior.

An additional aspect of the solution is the stiffness of the linear problem that occurs by employing the NR method. One way to assess this property is by the condition number of the Jacobian matrix at each NR cycle. The corresponding feature of the Watson [31] library has been used in order to calculate the condition number at each iteration. The threshold for NR iterations is 100 per continuation step as they are much more than the number of iterations in which the PAL method typically converges. The continuation scheme is initiated with the Newtonian case.

The results of Fig. 10 correspond to small (10(a)) and large (10(b)) Bingham numbers, respectively. The problem becomes stiffer as the plastic nature of the material increases, mainly due to the presence of a larger portion of unyielded areas where the numerical difficulties arise.

The results for the small Bingham numbers are summarized in Table 4. The plot for small Bingham numbers is shown for up to $\sim$40 iterations for clarity, however the divergent cases surpassed the 100 iterations threshold. As expected, the condition number of the Jacobian is roughly two orders of magnitude higher in the case of the PAL method when compared to all cases of the PR solution. For the two values of the penalty parameter, the condition numbers that are computed are identical and the corresponding points in Fig. 10 overlap. The higher values of the condition number in the PAL case, when compared to PR, indicate that the matrix tends to become singular, a characteristic of the original AL problem. As for the PR method, 3 cases are shown of which only one converges. When using a value of $N = 500$, the convergence depends on the continuation step and the divergent case halts at $Bn = 0.1$. By keeping the same continuation step, when $N = 1000$ the continuation scheme diverged at $Bn = 0.03$. One possible explanation of the inferior convergence behavior of the PR method, lies in the fact that the exponential term of the PR method renders the problem very sensitive to small perturbations. In other words, both when the solution is updated following the ZOC (Zero-order-Continuation) scheme and on each NR iteration, the presence of the

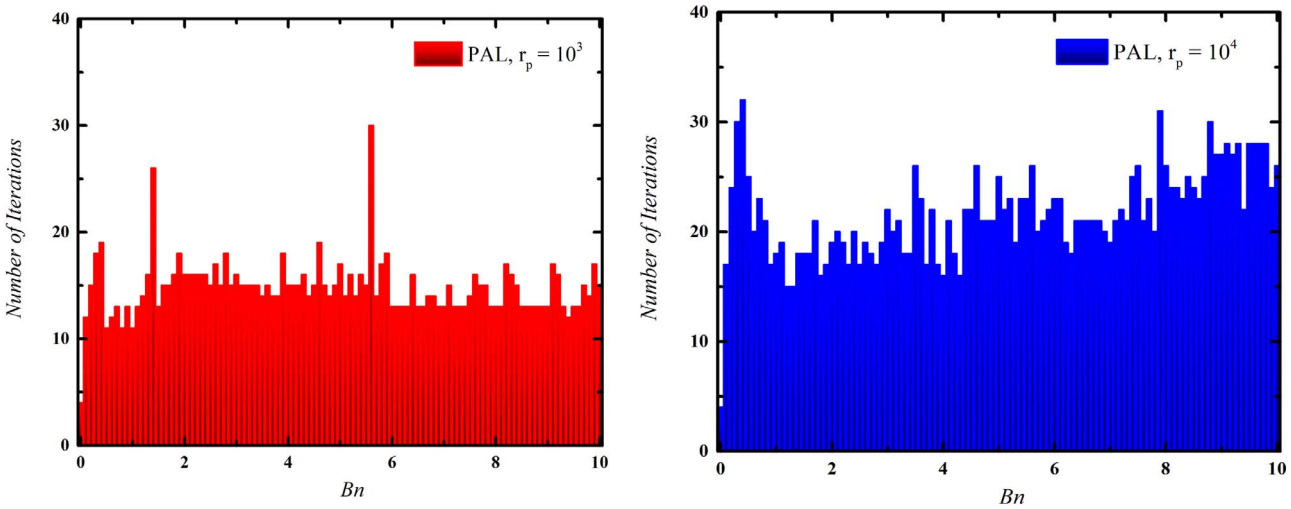

Fig. 8. Number of Newton–Raphson iterations for the channel problem using the PAL method for $r_p = 10^3$ (left), and $r_p = 10^4$(right).

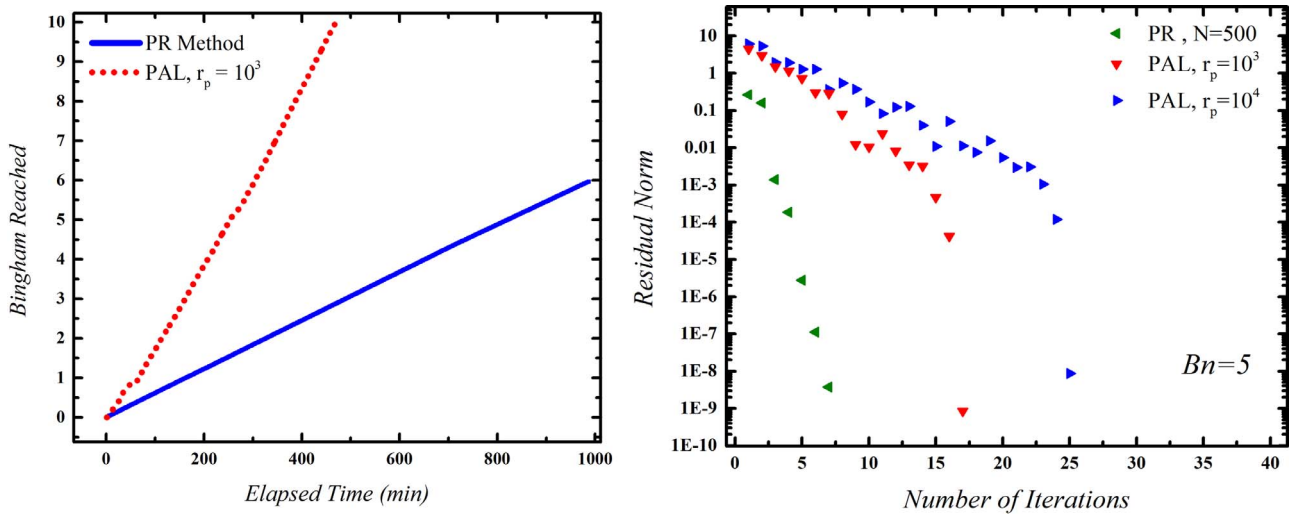

Fig. 9. Required time for reaching a specific Bingham number using ZOC and reduction of the residual norm as a function of NR iterations for $Bn = 5$.

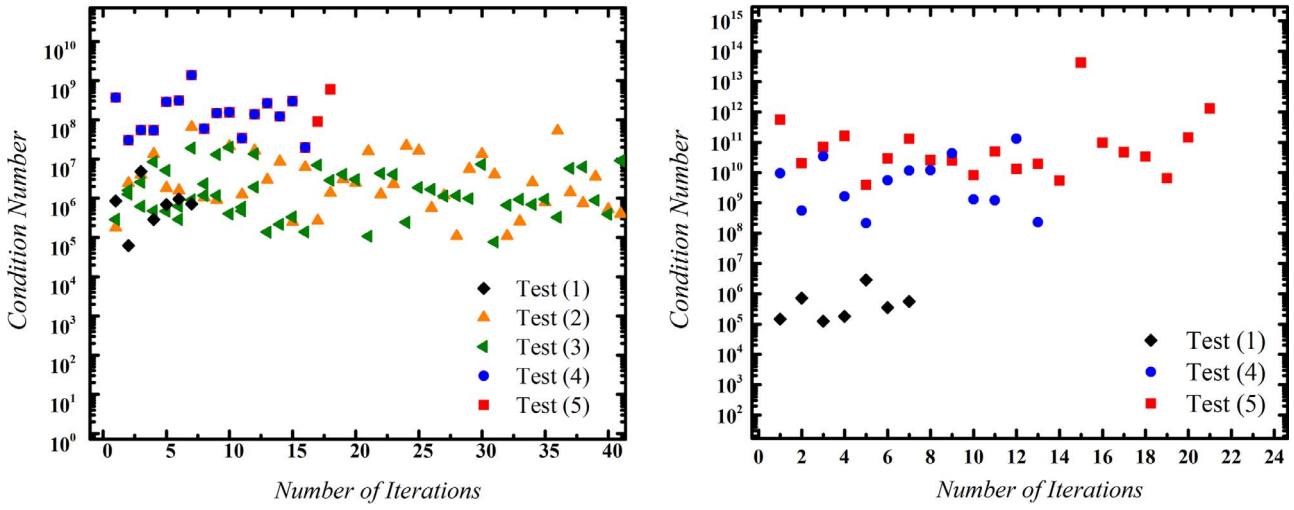

Fig. 10. Plot of the condition number of the Jacobian matrix for small (left) and large (right) Bingham numbers. See text for description.

exponential in the viscosity function with a large value of N changes the correction vector abruptly and thus the solution vector of the NR is "moving away" from the correct direction.

For the high Bingham number tests, the non-convergent cases have been omitted and thus, only the tests (1), (4) and (5) are given in Table 5 and Fig. 10(b). The deviation between PAL and PR have been intensified in this scenario. Typically, the condition number of the Jacobian using the PAL formulation are 5-6 orders of magnitude higher than the PR method. The penalty parameter value in the PAL method in this case gives rise to deviations on the condition number for Tests 1 and 4 which were absent before. In addition, the results qualitatively agree with those reported for the bubble problem for the drag coefficient. In other words, for lower Bingham numbers, the PAL and PR methods usually do not differ very noticeably, but the efficiency of PAL is revealed as plasticity increases.

Table 4
Tested cases of low the Bingham numbers presented in Fig 10(a).

Case	Method	Regularization Factor	Bn	ΔBn	Convergence
Test (1)	PR	$N = 500$	0.1	0.01	Yes
Test (2)	PR	$N = 500$	0.1	0.1	No
Test (3)	PR	$N = 1000$	0.03	0.01	No
Test (4)	PAL	$r_p = 1000$	0.1	0.1	Yes
Test (5)	PAL	$r_p = 10000$	0.1	0.1	Yes

5. Lid-driven cavity

The lid-driven cavity problem is probably the most often used benchmark problem for testing novel numerical methods in terms of accuracy and overall performance. In this section, a comparison between our results and those of Treskatis et al. [62] is presented in order to contrast some characteristics of our method and the FISTA algorithm presented by these authors. The boundary conditions imposed are no-slip and no-penetration in each side of the cavity. Motion is induced by a constant translation of the lid. The problem is solved using the PAL algorithm with $r_p = 10^3$ and $r_p = 10^4$ in order to examine the effect of

Table 5
Tested cases of high Bingham number presented in Fig 10(b).

Case	Method	Regularization factor	Bn	ΔBn	Convergence
Test (1)	PR	$N = 500$	10	0.01	Yes
Test (4)	PAL	$r_p = 1000$	10	0.1	Yes
Test (5)	PAL	$r_p = 10000$	10	0.1	Yes

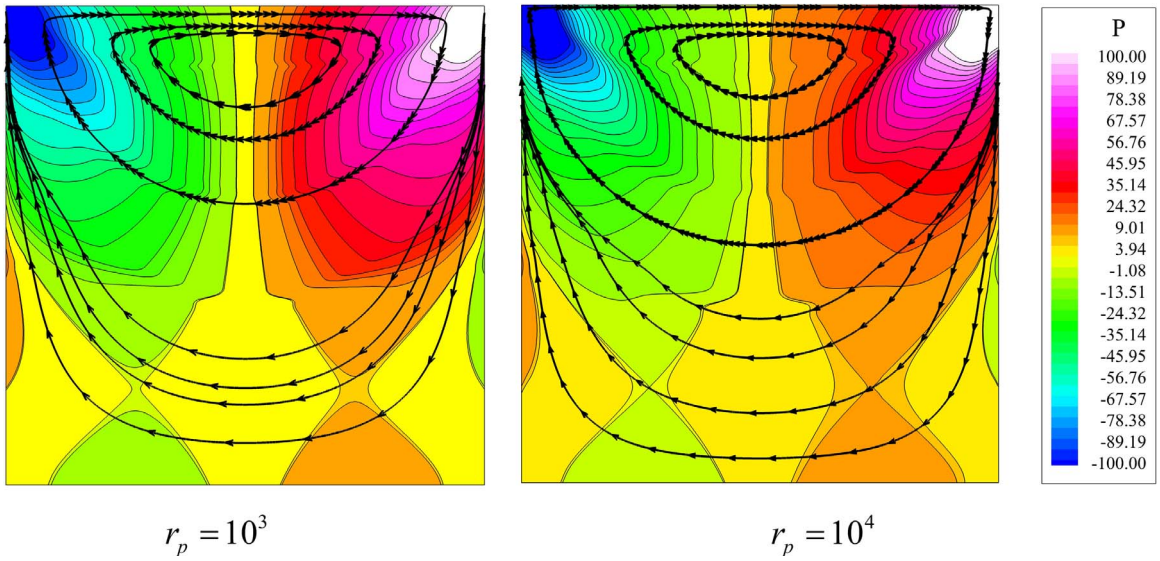

Fig. 11. Pressure contours and streamlines predicted by the PAL method for two different values of the penalty parameter, with dimensionless parameters (Bn, Re, n) = (20, 0, 1).

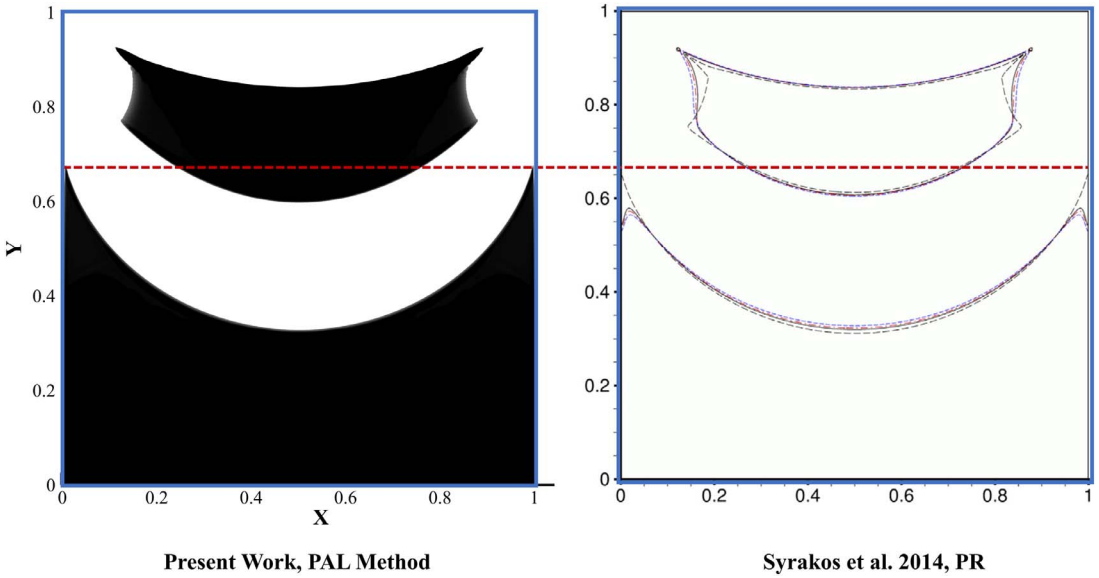

Fig. 12. Yielded (white) and Unyielded (black) regions for (Bn, Re, n) = (10, 0, 1)using the PAL method with $r_p = 10^4$ and compared to the predictions of Syrakos et al. [58] using the PR method with N = (100, 200, 400) for the blue–dash , red–dash and black-solid lines respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the penalty parameter to the flow variables and, especially, the yield surfaces that are formed.

The pressure field inside the cavity along with the streamlines is given in Fig. 11. The pressure field is very similar to that in Syrakos et al. [60], thus validating that the present results are qualitatively correct. It should be noted that our contours are generally sharper, which is a characteristic of the presence of discontinuity in the constitutive model. The penalty parameter affects the pressure contours and the streamlines slightly, but without a noticeable difference.

Of more importance are the yield surfaces that are formed. First of all, it is interesting to examine the predictions of our algorithm in comparison to those by Syrakos et al. [58], where the PR method was employed using various values for the dimensionless exponent. Before doing so, it should be mentioned that the discontinuous yield stress model has been found to have some intrinsic numerical "noise" near the yield surface in this flow, Treskatis et al. [62]. Therefore, the contour plots, i.e. the yielded/unyielded regions, are determined by boundaries defined via $|\underline{\tau}| = (1 \pm \varepsilon)Bn$, where a very small parameter $\varepsilon \simeq 10^{-3}$ is used to include several contours that produce the yield surfaces shown. However, this is not observed with the PR approximate model. In addition, following ref. [62], a uniform grid of 128 finite elements in each direction has been used, although these authors also used adaptive refinement of the mesh. In Fig (12), the set of dimensionless parameters is (Bn, Re, n) = (10, 0, 1) and the penalty parameter is $r_p = 10^4$.

At first glance, the plots of Fig. 12 are qualitatively similar. Generally, there are two major unyielded regions. One in the bottom half of the domain, away from the moving lid, and one resembling an island in the middle of the upper half of the cavity. The upper surface of the former has a parabolic shape. Particularly, near the side walls of the cavity and the unyielded island, the values of the PR exponent affect the solution. It can be seen from Syrakos et al. [58] that the unyielded surface near the cavity side-walls has a maximum slightly away from the walls. The black dashed line in the same figure (right in Fig. 12), which is the numerically extrapolated solution for very large PR exponents, does indeed predict a perfect parabolic shape. In fact, this is the shape obtained directly with the PAL method. Not only this, but also

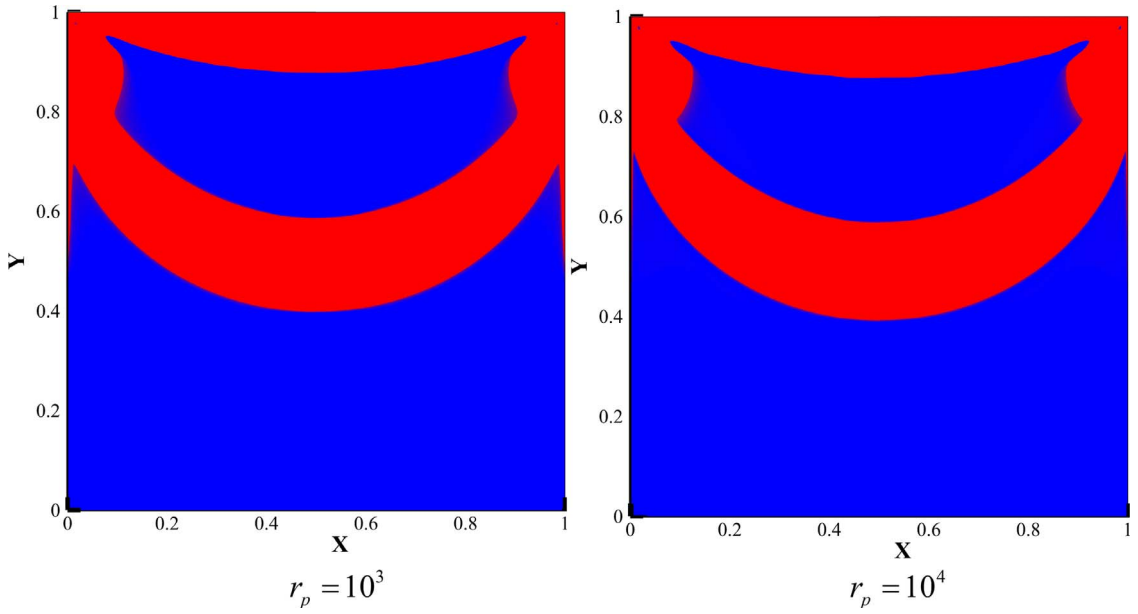

Fig. 13. Yielded and unyielded regions for (Bn, Re, n) = (20, 0, 1).

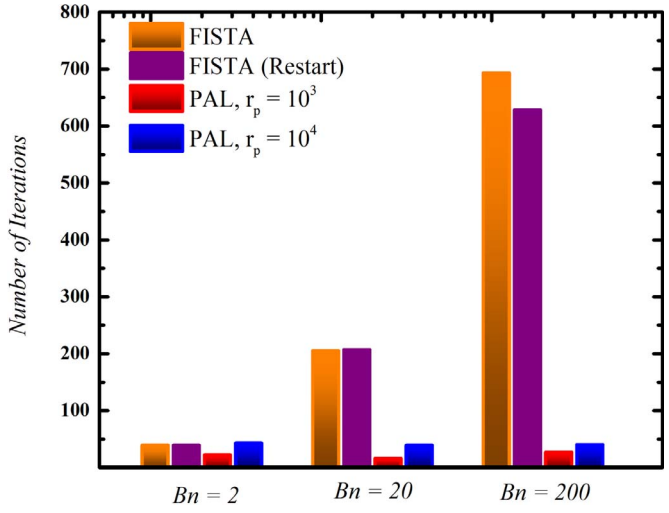

Fig. 14. Iterations for convergence of the lid-driven cavity problem in various cases. The orange and purple bars were obtained by Treskatis et al. [62]. The red and blue bars are the iterations using the PAL method. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the point of intersection of the parabolic yield surface with the side walls is identical to the corresponding point obtained with the extrapolated solution of [58]. Furthermore, for the given exponents of the PR solution of [58], the "island" – like unyielded region is overly smooth, in contrast to the sharper edges of the extrapolated solution. Once again, the solution using the PAL method predicts a sharper yield surface, similarly to the extrapolated solution. Thus, the PAL method is shown to be a very accurate method for obtaining those unique characteristics of viscoplastic fluids.

Next, it is useful to examine the effect of the penalty parameter to the predicted yield surfaces in Fig. (13). The results are given for (Bn, Re, n) = (20, 0, 1) in order to be the same with those in Treskatis et al. [62]. As mentioned earlier, the largest deviations are located near the edges of the unyielded regions, although the penalty parameter is increased considerably. One noticeable difference arises from the parabolic profile of the lower unyielded region near the wall. It can be seen that, for the lower penalty parameter, the yield surface does not approach and eventually intersect the wall monotonically. A sharp maximum arises, which approaches further the wall for the larger penalty parameter. It is anticipated that this trend will continue until the maximum in the yield surface coincides with the intersection with the wall, which is the case when the AL method is used, like in Glowinski and Wachs [30]. Those sharp "tips" are located at a distance of 0.013 (when $r_p = 10^3$) and 0.0031 (when $r_p = 10^4$) from the wall. Hence, this artificial "boundary layer" has thickness about 300 times smaller than the domain's length and, thus, this deviation from the predictions of the original AL method is negligible. In addition, the "island" is smoother for the smaller penalty parameter, as expected from the correspondence mentioned earlier between the penalty parameter and the PR exponent. The profiles are very similar to those of [62], for $Bn = 20$.

Comparing the computational cost between our new method and that by Treskatis et al. [62], we observe the following from Fig. 14. The new PAL algorithm converges much faster, especially for larger Bingham numbers, needing for $Bn = 20$ just 20 iterations, whereas the FISTA scheme requires about 200 iterations. Moreover, the number of iterations with PAL remains fairly constant, whereas it increases exponentially with FISTA, as the Bingham number increases.

An issue that may raise concerns about the PAL solution method is its performance with iterative solvers which in many cases are preferable, such as in 3D simulations. To this end, the problem is revisited by using an iterative method of solving the linear system. A GMRES solver is employed with an incomplete factorization preconditioner. The ILU (τ_{if},γ_{if}) / GMRES algorithm is preferred, because it is applicable in saddle-point problems (e.g. incompressible flows). The index *if* stands for incomplete factorization. The ILU(τ_{if},γ_{if})/GMRES algorithm is preferred due to its speed, if properly selected parameters are used. Out of the two parameters associated with it, the first one is the τ_{if} drop-off parameter. All entries(i, j) and (j, i) are dropped, if they are τ_{if} times smaller than the diagonal entry (i, i), for $i > j$. The second one is the γ_{if} parameter, which is the upper threshold for the zero entries in each row and column. More details are given in the solver manual [31] and in Gupta [32], Saad [50] and Saad and Schultz [51].

In general, the γ_{if} parameter should take large values in order that the ILU(τ_{if},γ_{if}) is efficiently applied in the system of equations formed by (18) and (19). Its typical values range from 3 to 5. In all cases examined in the linear-log Fig. 15, $\gamma_{if} = 4.5$. It can be seen that the reduction of the τ_{if} parameter increases almost linearly the average time for the assembly of the precoditioner matrix. However, the total time eventually

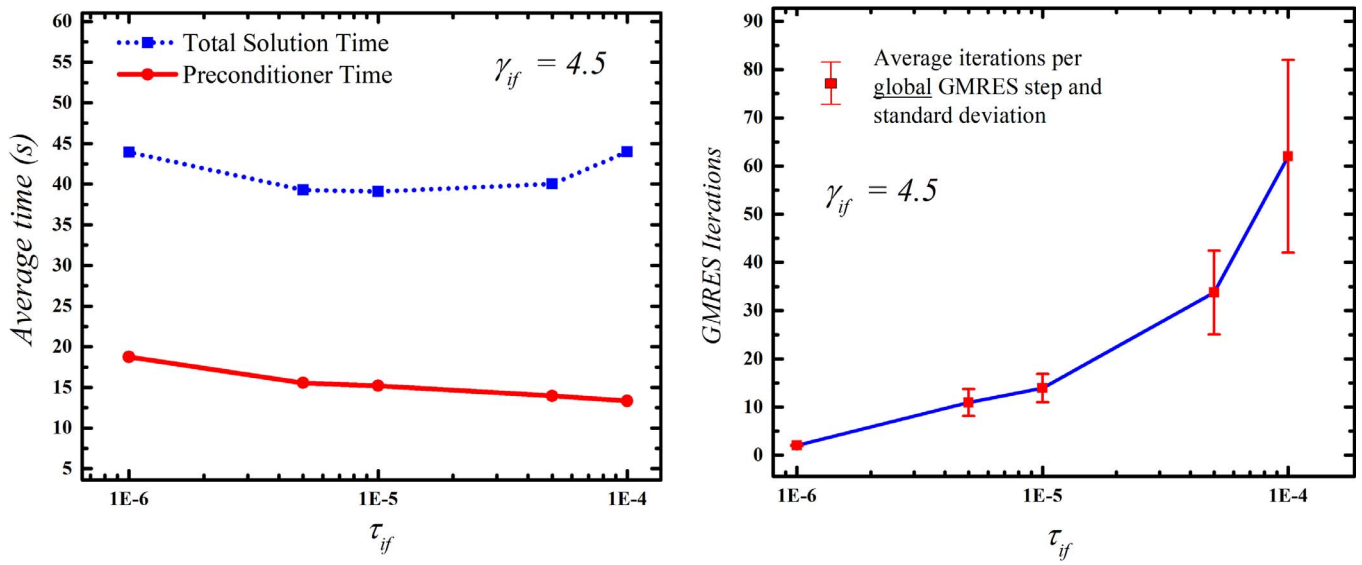

Fig. 15. Average time(left) and iterations(right) for each GMRES step for various values of τ_{if}. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

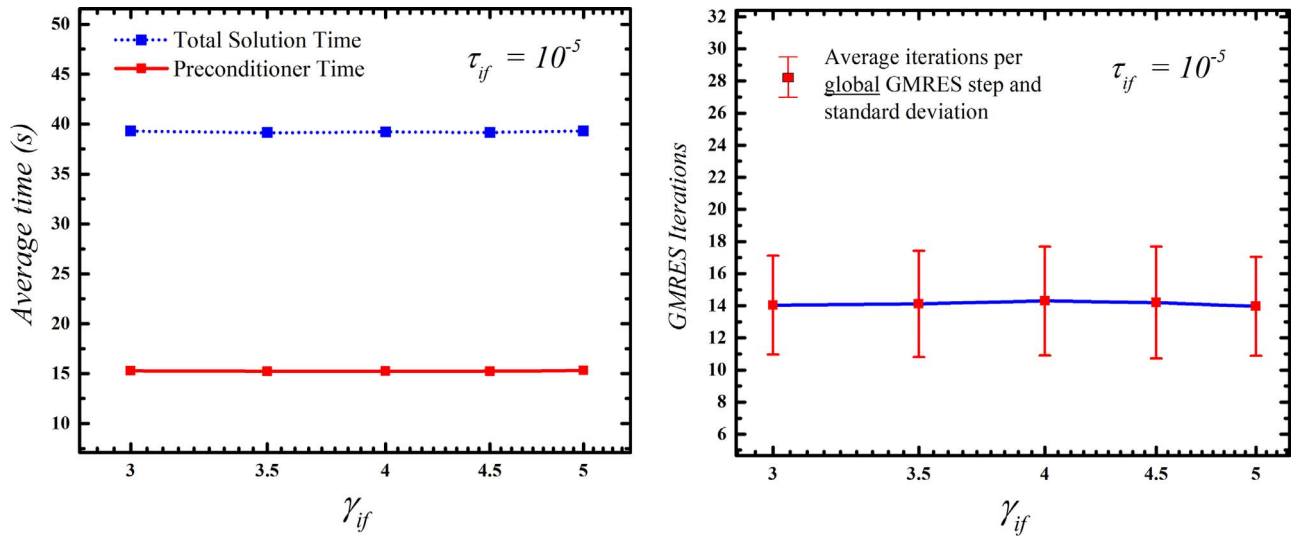

Fig. 16. Average time and iterations for each GMRES step for various values of γ_{if}.

increases, since more iterations are needed for the solution of the linear problem. It can be seen that there is an optimal value of the parameter, approximately $\tau_{if} \simeq 10^{-5}$ in this case, for which the total completion time for a GMRES step is minimized. Larger simulations, such as 3D cases can exploit this finding, since the τ_{if} can be tuned for speedup. Another characteristic is that by comparing the results for required times in Fig. 15a it is seen that the generation of the preconditioner is approximately $\sim$ 40% of the total solution time. It should be noted that the time of a global GMRES step is related to the fact that the required number of iterations was practically the same regardless the NR iteration and the Bingham number, especially for lower values of τ_{if}.

Similar tests were undertaken for examining the effect of the γ_{if} parameter. From the results depicted in Fig. 16, it is clear that γ_{if} practically has no impact to the iterations needed for convergence of the linear problem, the corresponding total time and the time required for the generation of the preconditioner. Finally, it should be mentioned that WATSON has a built-in utility that adjusts dynamically both these parameters depending on a preliminary analysis of the input matrix. This was not used in this work, but in general it is suggested that when implementing the PAL scheme with this method for the solution of the linear problem, it is a good idea to do preliminary runs and decide what is optimal speed-wise.

6. Square duct

As mentioned before, Saramito [52] proposed a solution method for Viscoplastic flows using a preconditioner to eliminate the discontinuity of the Jacobian matrix. The method was showcased by solving the square duct problem, which is briefly presented in this section. As mentioned before, a uniform uni-directional pressure gradient is applied, driving the flow in a square duct. Therefore, the velocity vector is $\tilde{\underline{v}} = (0, 0, \tilde{v}_z(\tilde{x}, \tilde{y}))$. The results are validated by comparison to both those of Saramito [52] and Taylor and Wilson [61]. A basic difference between the present work and Saramito's [52] is that here the geometric symmetries are not exploited, therefore the problem is solved in the entire square domain. The governing equations remain the same as in the previous problems. As for the boundary conditions, no-slip is imposed on all sides of the square.

The results for a Herschel-Bulkely fluid are shown in Fig. 17 and are plotted for the upper-right “quadrant” for greater resolution. Of course,

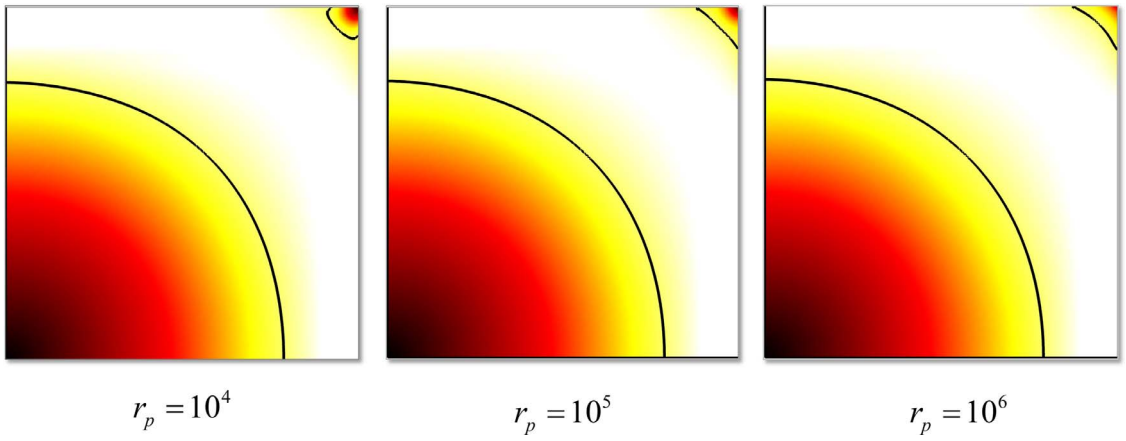

Fig. 17. Second invariant of stresses (color band) and the corresponding yield surfaces (black lines) for (Bn, Re, n) = (0.2, 0, 0.5)and varying penalty parameter values. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

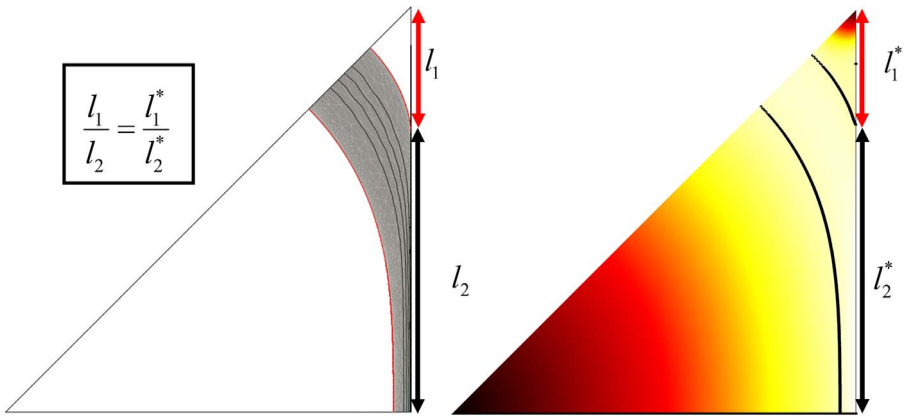

Fig. 18. Yield surfaces as predicted by Saramito (left) and the current method (right) for (Bn, Re, n) = (0.25, 0, 0.5). The corresponding Bingham number for Saramito is Bn = 1 (see text).

the solution is identical in all four corners of the square. The present results resemble those of Taylor and Wilson [61] with regard to the shape of the yield surfaces, their locations, and the critical Bingham number, for which the dead regions in the corners merge with the core unyielded region. In fact, the latter is found to be $Bn \simeq 0.27$ and the maximum velocity is $v_{z,max} = 1.69 \cdot 10^{-8}$ in this case. This critical Bingham value has been shown not to be affected by shear – thinning [52]. The computational grid used consists of 100 equispaced elements in each direction.

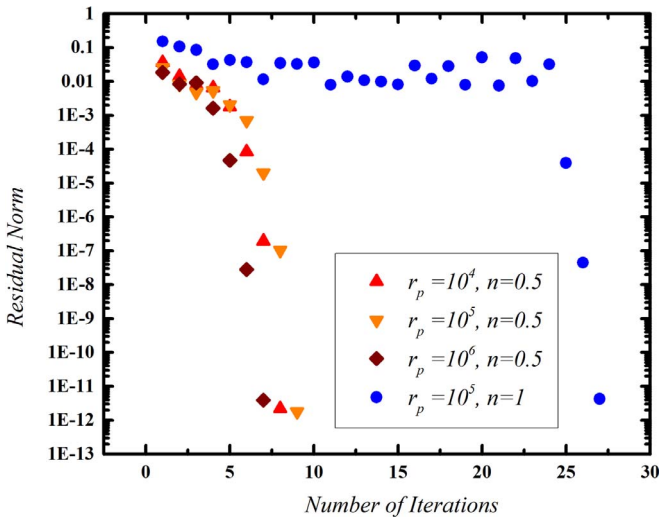

Fig. 19. Decrease of residual norm with iterations for PAL method using various values of the penalty parameter for the parameters of Fig. 17.

Once again, the importance of the penalty parameter is highlighted in Fig. 17, since it considerably affects the shape of the yield surfaces near the corner. For the lowest value of the penalty parameter, the shape of the yield surface is convex with respect to the corner, which is not the correct shape of the yield surface according to the findings of Saramito [52]. He predicted unyielded regions in the same locations (for a different value of Bingham number due to a different selection of characteristic length scales; see below), but the boundary of his unyielded region in the corner is concave with respect to the corner. This is found with our code by increasing the penalty parameter. For $r_p = 10^5$ the curvature of the line changes to concave and for the highest value used, which is $r_p = 10^6$, the boundary aligns with the inner yield line, forming a circular section as expected.

Comparing our results for the yield surfaces predicted by the inexact Newton's method of Saramito, it can be seen in Fig. 18 that the results in both cases are in excellent agreement. Saramito uses Bn = 1 which is 4 times larger than the Bingham number used in our formulation. If the square has a side length $\tilde{L}$, we use this value as characteristic length, whereas Saramito defines his Bingham number using the hydraulic length, i.e. $0.25\tilde{L}$ as a characteristic length. Therefore, a rescaling should take place since the yield stress and the pressure drop are identical, in order to have comparable results.

As far as the performance of the algorithm is concerned, the residual norm decreases with the elapsed iterations according to Fig. 19. Again we have found that the PAL method is fast and converges in about 7 to 9

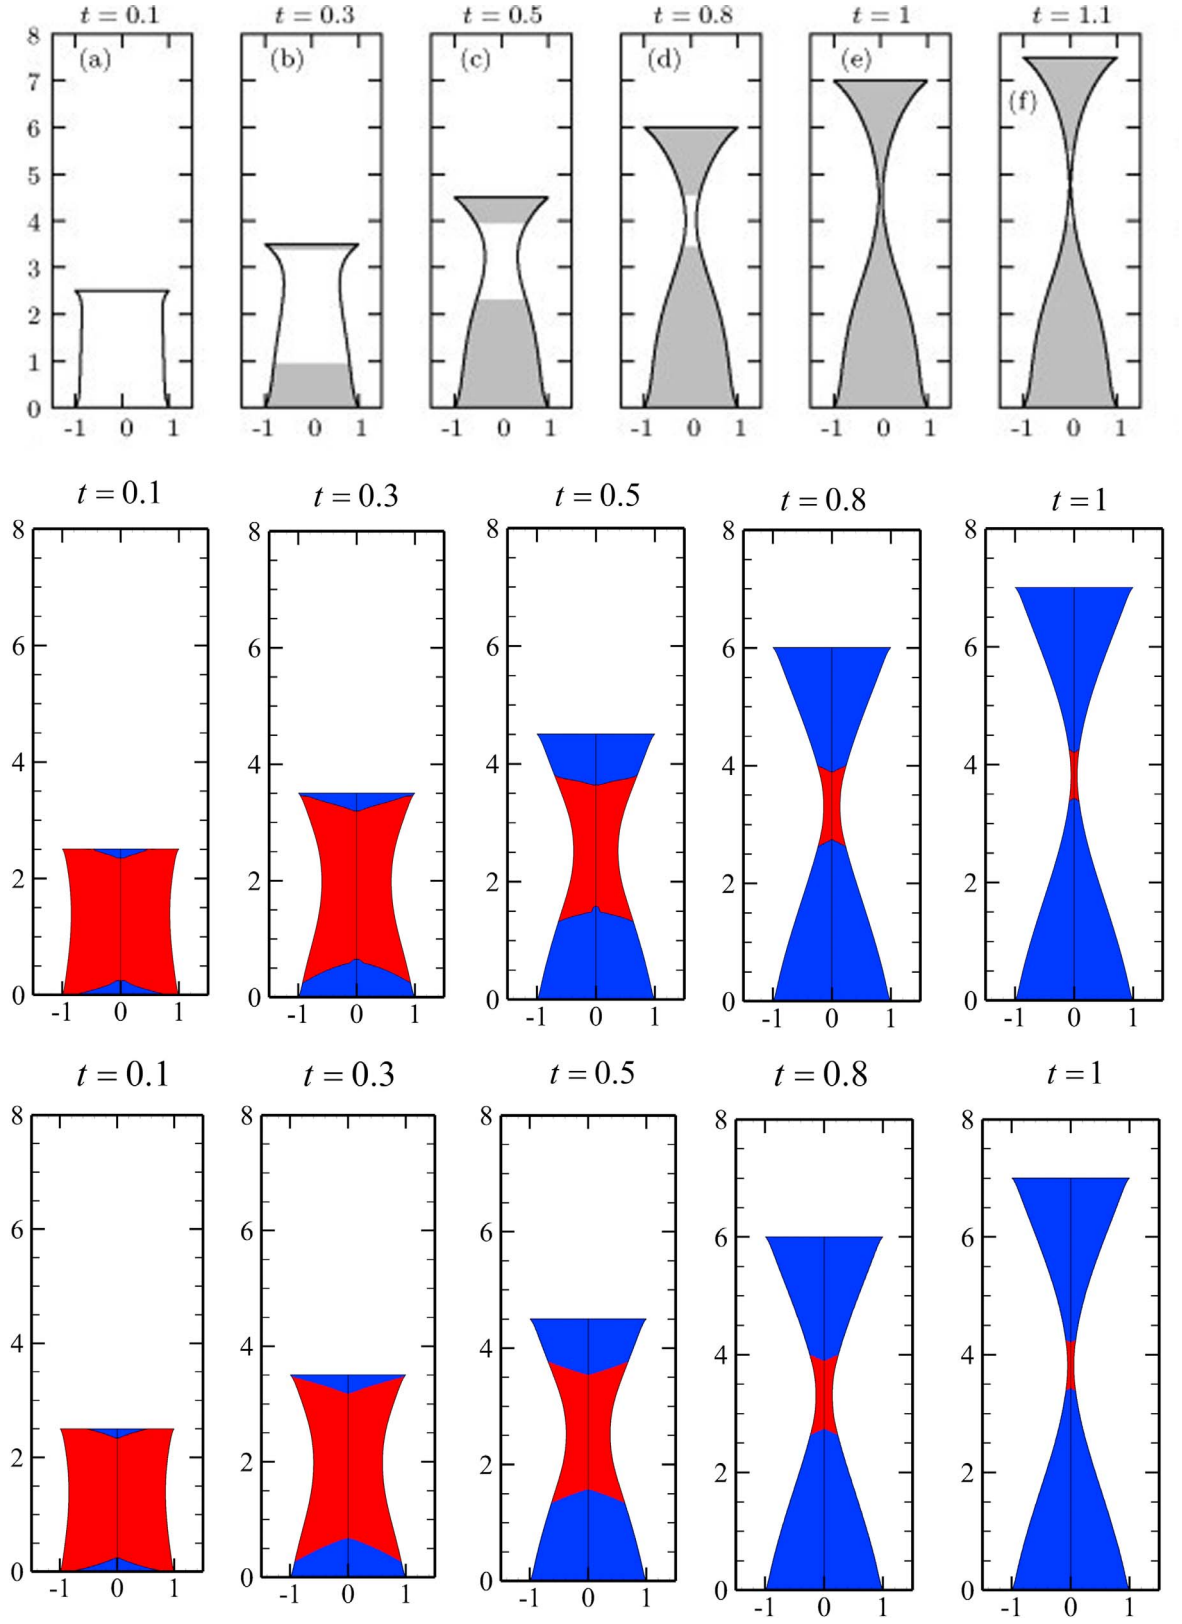

Fig. 20. Time evolution of the filament shape and yielded/unyielded regions. The panels in the upper row show the predictions of the slender body approximation by Balmforth et al [3]. The panels in the middle row show the predictions of the PAL method with $r_p = 10^4$. The panels in the lower row show the predictions of the Papanastasiou approximation with $N = 10^3$. The dimensionless numbers used are $(U_o, Bn, Bo, Re) = (5,\ 5,\ 1,\ 0)$ and $L_o = 2$.

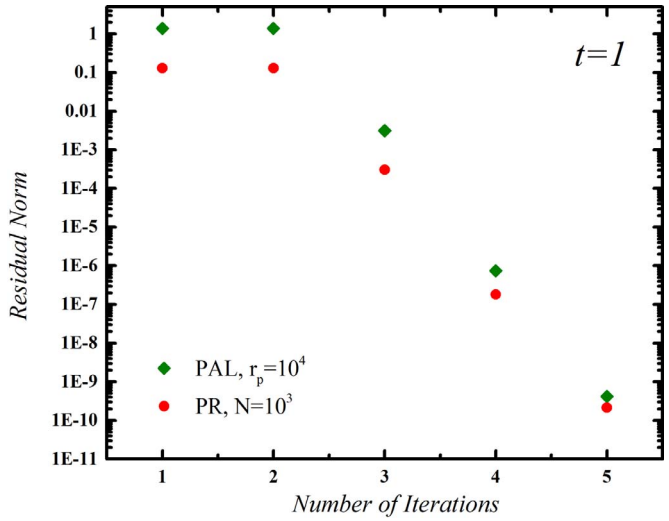

Fig. 21. Reduction of the residual norm as a function of NR iterations for $t = 1$ of Fig. 20.

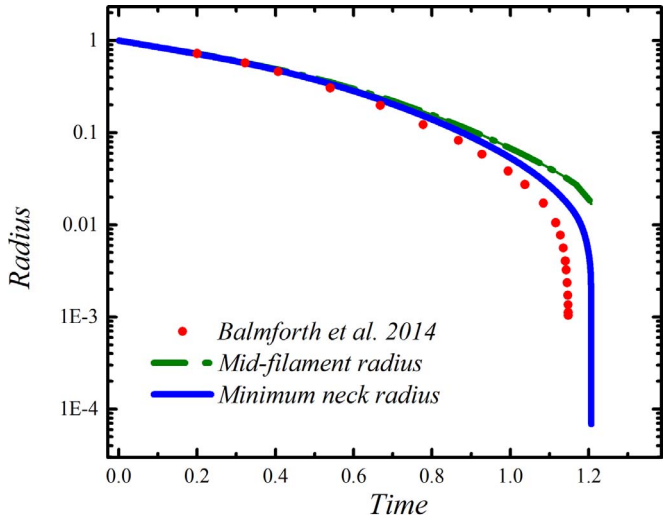

Fig. 22. Evolution of the mid-filament and minimum radius with time and comparison to the minimum radius computed by Balmforth et al. [3].

iterations for the set of (Bn, Re, n) = (0.2, 0, 0.5). However, as it was also observed in [52], the convergence for a Bingham plastic ($n = 1$) is slower. In fact in the present case, the solution for (Bn, Re, n, r_p) = (0.2, 0, 1, 10^5) needs roughly 3 times the number of iterations needed for the shear – thinning fluid. Saramito [52] shows the required iterations for converge for the inexact Newton's method, for a typical value of $Bn = 0.5$ (according to his non-dimensionalization) and various computational grids. For the grid consisting of 80 elements in each direction, it was seen that the required iterations were roughly 17 when $n = 0.3$ and about 28 when $n = 1$. Therefore, the PAL algorithm converges at a rate comparable (when $n = 1$) or even higher than (when $n < 1$) the method of [52]. However, a really direct comparison cannot be made either for the exact number of iterations or for the required time, partially, because Saramito [52] also performs refinement to increase the accuracy near the yield surfaces and solves the problem taking into account all its symmetries.

7. Filament stretching

This is the only transient problem in the present work with a moving liquid/air boundary. The non-dimensionalization follows the work of Balmforth et al. [3,4] who studied this problem previously, using the slender filament approximation and neglecting variations in the radial direction. Therefore, the present analysis is the first computational work to our knowledge, where viscoplastic filament stretching is studied fully in two dimensions. Here a viscoplastic fluid is confined between two coaxial disks and it is "pinned" in their edges. The origin of the coordinate system is at the center of the lower disk. Stretching begins at $t = 0$ by pulling the upper disk. As time passes, the filament forms a "neck" the width of which constantly decreases until the filament breaks. This constant stretching causes the stresses in the material to dynamically vary continuously changing the yield surfaces. A similar problem for a Newtonian or viscoelastic filament including deformable bubbles has been studied by Foteinopoulou et al. [22,23] and Papaioannou et al. [41]. The boundary conditions are the usual no-slip and no-penetration conditions on the disks, symmetry is exploited as seen in Fig. (1) and the force balances along the material/air interface. The formulation is completed by specifying the initial conditions. To this end, the material is stationary and the shape of the free surface is cylindrical.

In this problem, the mesh construction is very crucial for the successful convergence of the algorithm. This necessity mainly arises near conditions of filament pinch-off, where a neck has been clearly formed. In the necking region, the deformation of the mesh is very high. Therefore, care should be taken in order to keep a reasonably good distribution and shape of the elements in the highly deformed region, so that there is not much computational error involved and remeshing steps are avoided. For the current computations, the mesh consists of 230 elements in the z-direction and 30 elements in the r-direction, properly distributed initially. For the PAL method, the parameters used are (r, r_p) = (1, 10^4). Our results for the full axisymmetric model are compared to those of the lubrication model [3]. More specifically, for a given set of parameters, it is interesting to explore the evolution of the shape of the filament and the solid and liquid regions. This is done in Fig. 20 for the set of (U_o, Bn, Bo, Re) = (5, 5, 1, 0), which is identical to the one used in [3].

The transient nature of the problem adds a significant computational burden, since a consistent numerical scheme should be used for the time integration. To this end, a predictor–corrector algorithm is constructed in order to ensure the stability of the entire scheme. Depending on the discrepancy between the prediction and the following correction, the time-step is adjusted accordingly, since near "pinch-off" different time scales vary differently and a very small time-step is needed for convergence. The integration begins with an extremely small time-step $\Delta t_{\min} = 10^{-9}$, which increases up to $\Delta t_{\max} = 10^{-4}$, whereas further increasing it results in convergence difficulties of the NR scheme. When the necking region becomes very thin, the time-step is automatically decreased to $\Delta t_{\max} \approx 10^{-6}$. In general, the required small time-steps result in approximately ~15,000 time-steps per run. Of course, this depends on the time-scales and the pinch-off time. Typically, the implicit time-integration step, which requires a Newton-Raphson cycle converges in 5 iterations for the PAL method, when the max time-step is achieved, with several time-steps requiring up to 8-10 iterations. Therefore, the real time for the complete solution was typically 2 days. As for the Papanastasiou method, the solution of the same problem was attempted both using $N = 10^3$ and $N = 10^4$, which are relatively large values of the exponent. The solution was possible only when $N = 10^3$, with approximately the same number of time steps required and 5 iterations per Newton–Raphson cycle, while when $N = 10^4$ was used, the scheme diverged after $t \simeq 0.27$. The dependence of the residual norm on elapsed iterations is shown in Fig. 21 for a specific time-step.

Comparing our results to those of Balmforth et al. [3], it can be seen that the PAL method produces qualitatively similar results both for the instantaneous shape of the filament and the yielded/unyielded regions. For early times, there is an initial global yielding of the filament, except for two regions near the axis of symmetry and where the filament is "pinned" to the disks. This is not observed in [3]. Generally, an anticipated difference between the predictions of the present work and

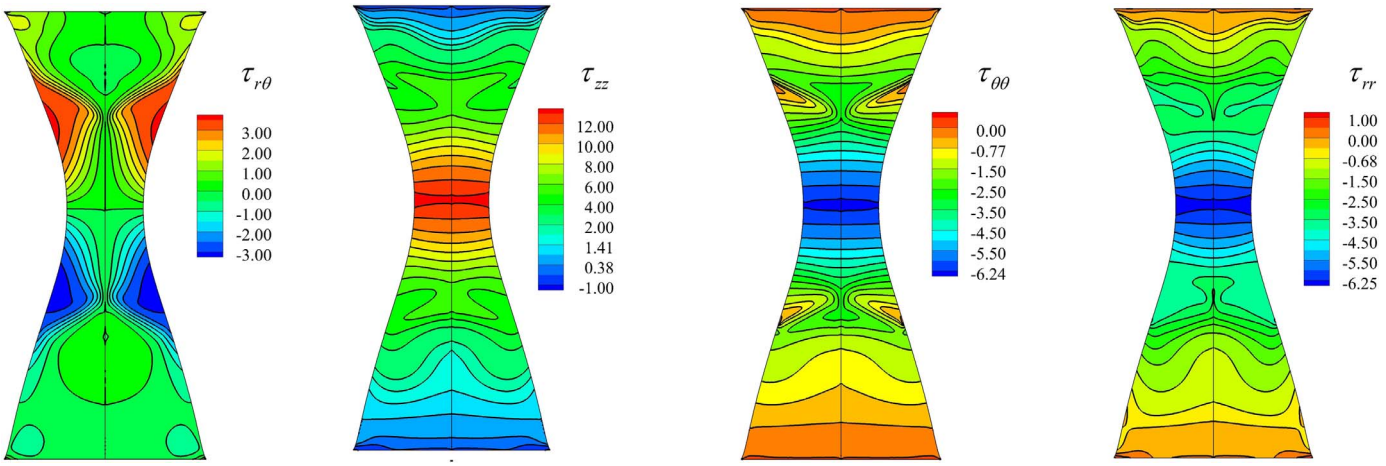

Fig. 23. Stress distribution for $t = 0.5$ where the neck has started to form.

those of the slender approximation is the variation of the yield surface (or any other variable) along the radial direction. Those variations cannot be predicted in the 1-dimensional approach, however, they seem to be relatively small and the phenomenon is described similarly between the two approaches. For larger times, the yielded region shrinks around the neck. This does not happen symmetrically, because gravity pulls the material towards the bottom of the filament. Regarding the gravitational effects, another difference between our results and those of [3] is that, in the latter, gravity seems to cause larger sagging in the filament. Finally near pinch-off, the shape of the filament obtains this sharp conical arrangement near the neck, which is characteristic of viscoplastic fluids as shown for Carbopol and kaolin suspensions in [3]. The results for the Papanastasiou approximation are very close to those by the PAL method.

Another important quantity in this problem is the change of the minimum radius at the neck of the filament, which occurs at different axial locations and is compared to the radius at its middle, Fig. 22. Although not very clear, the effect of gravity can be appreciated by the fact that the mid-filament radius is different from the minimum radius, though not as visibly as in the case of the 1-D approximation. The time required for the minimum filament to reach a very small value namely about $r_{\min} \sim 10^{-3}$ is predicted as $t \simeq 1.2$ using the present method and $t \simeq 1.15$ in the slender body approximation case. These values are quite close.

The distribution of the stress components at $t = 0.5$ is given in Fig. 23. As expected, the magnitudes of the normal stresses are much larger than the shear stresses. The shear stresses have a larger magnitude near the free surface and are antisymmetric with respect to the neck region, while they are nearly zero at the neck. The primary normal τ_{zz} stress is dominant, justifying, to some extent, the slender filament approximation. It has its maximum value at the necking region where the stretching is mainly felt and decreases towards the stagnation points. Furthermore, the importance of the other normal stresses is also clear and they follow the same pattern as τ_{zz}, i.e., they attain their maximum absolute value around the necking region and tend to decrease near the disks.

8. Conclusions

Flows of Viscoplastic fluids are tremendously important in many applications and, thus, accelerated algorithms that provide accurate solutions in a reasonable time frame are required. In the present work, the algorithm presented by Dimakopoulos et al. [15] is modified, by penalizing the equation, which updates the Lagrangian multipliers and the validity of the results obtained is tested in five problems. The dependence of the results on the penalty parameter is a central subject of discussion throughout this work. We call this method the Penalized AL (PAL) method and we show that it can predict with high accuracy the yield surfaces that are produced. More specifically, its higher accuracy and overall faster convergence are more profound when the problem is solved for higher Bingham numbers, where larger unyielded regions arise, with respect to the PR method. The PAL scheme needs only a few iterations for the Newton-Raphson scheme to converge and the number of required iterations does not significantly change when the penalty parameter is increased. Regarding the NR algorithm employed, the initial guess that is used in each iteration has a large impact on the final convergence. Therefore, we have utilized the techniques of zero or first order continuation in order to obtain a “good” initial approximation of the solution when the Bingham number varied with a fixed step. Then, the PAL method has been found to be more robust, because larger continuation steps could be used than with the PR method, leading to an overall speedup of the simulation. The PAL method comes with several limitations in its current form. Although it can achieve more accurate solutions compared to PR without suffering from non-convergence problems, the solution achieved is inherently less accurate than the one with the AL method as implemented in ALG2. Of course, one could perform continuation in the penalty parameter, bringing its predictions closer to those of AL.

In the future, we would like to apply this methodology in 3-Dimensional problems and observe how PAL compares to the AL and PR methods. The scalability of the method and its performance with iterative linear solvers is of a great importance. Finally, this method can be improved by finding a consistent way to update the PAL parameters, i.e. (r, r_p) in each NR iteration so that they can attain values as close as possible to the optimal ones.

Acknowledgments

This research was supported by Grant #E656 from the Research Committee of the University of Patras via “K. Karatheodori” program. G.M. has been supported by LIMMAT foundation under the project MuSiComPS. We would also like to thank the reviewers for their constructive comments

References

[1] P. Alart, A. Curnier, A mixed formulation for frictional contact problems prone to Newton like solution methods, Comput. Methods Appl. Mech. Eng. 92 (3) (1991) 353–375.
[2] P.R. Amestoy, I.S. Duff, J. Koster, J.Y. L'Excellent, A fully asynchronous multifrontal solver using distributed dynamic scheduling, SIAM J. Matrix Anal. A 23 (1) (2001) 15–41.
[3] N.J. Balmforth, N. Dubash, A.C. Slim, Extensional dynamics of viscoplastic filaments: II. Drips and bridges, J. Non-Newton. Fluid Mech. 165 (19–20) (2010)

1147–1160.
[4] N.J. Balmforth, N. Dubash, A.C. Slim, Extensional dynamics of viscoplastic filaments: I. Long-wave approximation and the Rayleigh instability, J. Non-Newton. Fluid 165 (19–20) (2010) 1139–1146.
[5] N.J. Balmforth, I.A Frigaard, G. Ovarlez, Yielding to stress: recent developments in viscoplastic fluid mechanics, Annu. Rev. Fluid Mech. 46 (2014) 121–146.
[6] M. Bercovier, M. Engelman, A finite-element method for incompressible non-Newtonian flows, J. Comput. Phys. 36 (1980) 313–326.
[7] A.N. Beris, J.A. Tsamopoulos, R.C. Armstrong, R.A. Brown, Creeping motion of a sphere through a Bingham plastic, J. Fluid Mech. 158 (1985) 219–244.
[8] E.C. Bingham, Fluidity and Plasticity 2 McGraw-Hill, 1922.
[9] J. Bleyer, Advances in the simulation of viscoplastic fluid flows using interior point methods, Comput. Methods Appl. Mech. Eng. 330 (2018) 368–394.
[10] G.R. Burgos, A.N. Alexandrou, V. Entov, On the determination of yield surfaces in Herschel–Bulkley fluids, J. Rheol. 43 (1999) 463.
[11] N. Chatzidai, A. Giannousakis, Y. Dimakopoulos, J. Tsamopoulos, On the elliptic mesh generation in domains containing multiple inclusions and undergoing large deformations, J. Comput. Phys. 228 (2009) 1980–2011.
[12] I. Cheddadi, P. Saramito, F. Graner, Steady Couette flows of elastoviscoplastic fluids are nonunique, J. Rheol. 56 (1) (2012) 213.
[13] Y. Damianou, G. Georgiou, Viscoplastic Poiseuille flow in a rectangular duct with wall slip, J. Non-Newton. Fluid Mech. 214 (2014) 88–105.
[14] Y. Damianou, M. Philippou, G. Kaoullas, G. Georgiou, Cessation of viscoplastic Poiseuille flow with wall slip, J. Non Newton. Fluid Mech. 203 (2013) 24–37.
[15] Y. Dimakopoulos, M. Pavlidis, J. Tsamopoulos, Steady bubble rise in Herschel–Bulkley fluids and comparison of predictions via the Augmented Lagrangian Method with those via the Papanastasiou model, J. Non-Newton. Fluid Mech. 200 (2013) 34–51.
[16] Y. Dimakopoulos, J. Tsamopoulos, A quasi-elliptic transformation for moving boundary problems with large anisotropic deformations, J. Comput. Phys. 192 (2) (2003) 494–522.
[17] Y. Dimakopoulos, J. Tsamopoulos, Transient displacement of a viscoplastic material by air in straight and suddenly constricted tubes, J. Non-Newton. Fluid Mech. 112 (2003) 43–75.
[18] Y. Dimakopoulos, G. Karapetsas, N. Malamataris, E. Mitsoulis, The free (open) boundary condition at inflow boundaries, J. Non Newton. Fluid Mech. 187–188 (2012) 16–31.
[19] Y. Dimakopoulos, J. Tsamopoulos, Gas-assisted injection molding with fluids partially occupying straight or complex tubes, Polym. Eng. Sci. 46 (1) (2006) 47–68.
[20] N. Dubash, I. Frigaard, Conditions for static bubbles in viscoplastic fluids, Phys. Fluids 16 (2004) 4319.
[21] M. Fortin, R. Glowinski, Augmented Lagrangian methods: application to the numerical solution of boundary value problems, Studies in Mathematics and Its Applications, Elsevier Science Publishers B.V., 1983.
[22] K. Foteinopoulou, V.G. Mavrantzas, Y. Dimakopoulos, J. Tsamopoulos, Numerical simulation of multiple bubbles growing in a Newtonian liquid filament undergoing stretching, Phys. Fluids 18 (4) (2006) 042106.
[23] K. Foteinopoulou, V.G. Mavrantzas, J. Tsamopoulos, Numerical simulation of bubble growth in Newtonian and viscoelastic filaments undergoing stretching, J. Non Newton. Fluid Mech. 122 (1–3) (2004) 177–200.
[24] D. Fraggedakis, Y. Dimakopoulos, J. Tsamopoulos, Yielding the yield-stress analysis: a study focused on the effects of elasticity on the settling of a single spherical particle in simple yield-stress fluids, Soft Matter 12 (24) (2016) 5378–5401.
[25] D. Fraggedakis, Y. Dimakopoulos, J. Tsamopoulos, Yielding the yield stress analysis: a thorough comparison of recently proposed elasto-visco-plastic (EVP) fluid models, J. Non Newton. Fluid Mech. 238 (2016) 170–188.
[26] D. Fraggedakis, J. Papaioannou, Y. Dimakopoulos, J. Tsamopoulos, Discretization of three-dimensional free surface flows and moving boundary problems via elliptic grid methods based on variational principles, J. Comput. Phys. 344 (2017) 127–150.
[27] D. Fraggedakis, M. Pavlidis, Y. Dimakopoulos, J. Tsamopoulos, On the velocity discontinuity at a critical volume of a bubble rising in a viscoelastic fluid, J. Fluid Mech 789 (2016) 310–346.
[28] R. Glowinski, P. Le Tallec, Augmented Lagrangian and Operator Splitting Methods in Non-Linear Mechanics, Society for Industrial and Applied Mathematics, Philadelphia, 1989.
[29] R. Glowinski, Lectures on Numerical Methods for Non-Linear Variational Problems, Tata Institute of Fundamental Research, Springer – Verlag, 1980.
[30] R. Glowinski, A. Wachs, On the numerical simulation of viscoplastic fluid flow, Handbook of Numerical Analysis, Numerical Methods for Non-Newtonian Fluids 16 (2011), pp. 483–717.
[31] A. Gupta, WSMP: Watson Sparse Matrix Package (Part-I & III), IBM T. J. Watson Research Center, Yorktown Heights, NY, 2000 Technical Report RC 21886November http://www.research.ibm.com/projects/wsmp.
[32] A. Gupta, Enhancing performance and robustness of ILU preconditioners by blocking and selective transportation, SIAM J. Sci. Comput. 39 (1) (2017) A303–A332.
[33] W.H. Herschel, R. Bulkley, Konsistenzmessungen von Gummi-Benzollösungen, Kolloid Zeitschrift 39 (1926) 291–300.
[34] M. Hestenes, Multiplier and gradient methods, J. Optim. Theory App. 4 (5) (1969) 303–320.
[35] G. Karapetsas, J. Tsamopoulos, Transient squeeze flow of viscoplastic materials, J. Non Newton. Fluid Mech. 133 (2006) 35–56.
[36] Y. Liu, N.J. Balmforth, S. Hormozi, D.R. Hewitt, Two-dimensional Viscoplastic dambreaks, J. Non Newton. Fluid Mech. 238 (2016) 65–79.
[37] E. Mitsoulis, J. Tsamopoulos, Numerical simulations of complex yield-stress fluid flows, Rheol. Acta 56 (3) (2017) 231–258.
[38] E. Muravleva, L. Muravleva, Unsteady flows of a viscoplastic medium in channels, Mech. Sol. 44 (5) (2009) 792–812.
[39] L. Muravleva, E. Muravleva, G. Georgiou, E. Mitsoulis, Numerical simulations of cessation flows of a Bingham plastic with the augmented Lagrangian method, J. Non Newton. Fluid 165 (2009) 544–550.
[40] L. Muravleva, E. Muravleva, G. Georgiou, E. Mitsoulis, Unsteady circular Couette flow of a Bingham plastic with the Augmented Lagrangian Method, Rheol. Acta 49 (2010) 1197–1206.
[41] J. Papaioannou, A. Giannousakis, Y. Dimakopoulos, J. Tsamopoulos, Bubble deformation and growth inside viscoelastic filaments undergoing very large extensions, Ind. Eng. Chem. Res. 53 (18) (2014) 7548–7569.
[42] J. Papaioannou, G. Karapetsas, Y. Dimakopoulos, J. Tsamopoulos, Injection of a viscoplastic material inside a tube or between two parallel disks: conditions for wall detachment of the advancing front, J. Rheol. 53 (1155) (2009).
[43] T.C. Papanastasiou, Flow of materials with yield, J. Rheol. 31 (1987) 385.
[44] Papanastasiou, T.C., Malamataris, N., Ellwood, K., "A new outflow boundary condition" , Int. J. Numer. Methods Fluids 14(5), 587–608.
[45] C.G. Petra, O. Schenk, M. Lubin, K. Gärtner, An augmented incomplete factorization approach for computing the Schur complement in stochastic optimization, SIAM J. Sci. Comput. 36 (2) (2014) 139–162.
[46] D. Pettas, G. Karapetsas, Y. Dimakopoulos, J. Tsamopoulos, On the degree of wetting of a slit by a liquid film flowing along an inclined plane, J. Fluid Mech 820 (2017) 5–41.
[47] M. Philippou, Z. Kountouriotis, G.C. Georgiou, Viscoplastic flow development in tubes and channels with wall slip, J. Non Newtonian Fluid Mech 234 (2016) 69–81.
[48] M.J.D. Powell, Algorithms for non-linear constraints that use Lagrangian functions, Math. Program. 14 (1978) 224–248.
[49] J.C. Robertson, R.C. Kerr, Isothermal dynamics of channeled Viscoplastic lava flows and new methods for estimating lava rheology, J. Geophys. Res 117 (2012).
[50] Y. Saad, ILUT: a dual threshold incomplete LU factorization, Numer. Linear Algbr. 1 (1994) 387–402.
[51] Y. Saad, M.H. Schultz, GMRES: a generalized minimal residual algorithm for solving non-symmetric linear systems, SIAM J. Sci. Stat. Comput. 7 (1986) 856–869.
[52] P. Saramito, A damped Newton algorithm for computing viscoplastic fluid flows, J. Non Newton. Fluid Mech. 238 (2016) 6–15.
[53] P. Saramito, A new elastoviscoplastic model based on the Herschel-Bulkley viscoplastic model, J. Non Newton. Fluid Mech. 158 (1–3) (2009) 154–161.
[54] P. Saramito, A new constitutive equation for elastoviscoplastic fluid flows, J. Non Newton. Fluid Mech. 145 (1) (2007) 1–14.
[55] P. Saramito, A. Wachs, Progress in numerical simulation of yield stress fluid flows, Rheol. Acta 56 (3) (2017) 211–230.
[56] D. Smyrnaios, J. Tsamopoulos, Squeeze flow of Bingham plastics, J. Non Newton. Fluid Mech. 100 (1–3) (2001) 165–190.
[57] A. Syrakos, Y. Dimakopoulos, G.C. Georgiou, J. Tsamopoulos, Viscoplastic flow in an extrusion damper, J. Non Newton. Fluid Mech. 232 (2016) 102–124.
[58] A. Syrakos, G.C. Georgiou, A.N. Alexandrou, Performance of the finite volume method in solving regularised Bingham flows: inertia effects in the lid-driven cavity flow, J. Non Newton. Fluid Mech. 208 (2014) 88–107.
[59] A. Syrakos, G.C. Georgiou, A.N. Alexandrou, Cessation of the lid-driven cavity flow of Newtonian and Bingham fluids, Rheol. Acta 55 (2016) 51–66.
[60] A. Syrakos, G.C. Georgiou, A.N. Alexandrou, Solution of the square lid-driven cavity flow of a Bingham plastic using the finite volume method, J. Non.Newton. Fluid Mech. 195 (2013) 19–31.
[61] J.A. Taylor, S.D.R. Wilson, Conduit flow of an incompressible, yield-stress fluid, J. Rheol. 41 (1997) 93.
[62] T. Treskatis, M. Moyers-González, C.J. Price, An accelerated dual proximal gradient method for applications in viscoplasticity, J. Non Newton. Fluid 238 (2016) 115–130.
[63] J. Tsamopoulos, Y. Dimakopoulos, N. Chatzidai, G. Karapetsas, M. Pavlidis, Steady bubble rise and deformation in Newtonian and viscoplastic fluids and conditions for bubble entrapment, J. Fluid Mech. 601 (2008) 123–164.
[64] Tsouka S., Dimakopoulos Y., Mavrantzas V., Tsamopoulos J., "Stress-gradient induced migration of polymers in corrugated channels", J. Rheol.. 58 (4), 911–947.
[65] A. Wachs, Numerical simulation of steady Bingham flow through an eccentric annular cross-section by distributed Lagrange multiplier/fictitious domain and augmented Lagrangian methods, J. Non Newton. Fluid Mech. 142 (2007) 183–198.