

No. 669

December 2023

**Augmented Lagrangian acceleration of
global-in-time Pressure Schur complement
solvers for incompressible Oseen equations**

C. Lohmann, S. Turek

ISSN: 2190-1767

Augmented Lagrangian acceleration of global-in-time Pressure Schur complement solvers for incompressible Oseen equations

Christoph Lohmann, Stefan Turek

December 7, 2023

Abstract

This work is focused on an accelerated global-in-time solution strategy for the Oseen equations, which highly exploits the augmented Lagrangian methodology to improve the convergence behavior of the Schur complement iteration. The main idea of the solution strategy is to block the individual linear systems of equations at each time step into a single all-at-once saddle point problem. By elimination of all velocity unknowns, the resulting implicitly defined equation can then be solved using a global-in-time pressure Schur complement (PSC) iteration. To accelerate the convergence behavior of this iterative scheme, the augmented Lagrangian approach is exploited by modifying the momentum equation for all time steps in a strongly consistent manner. While the introduced discrete grad-div stabilization does not modify the solution of the discretized Oseen equations, the quality of customized PSC preconditioners drastically improves and, hence, guarantees a rapid convergence. This strategy comes at the cost that the involved auxiliary problem for the velocity field becomes ill conditioned so that standard iterative solution strategies are no longer efficient. Therefore, a highly specialized multigrid solver based on modified intergrid transfer operators and an additive block preconditioner is extended to solution of the all-at-once problem. The potential of the proposed overall solution strategy is discussed in several numerical studies as they occur in commonly used linearization techniques for the incompressible Navier-Stokes equations.

Keywords. Oseen equations; parallel-in-time; pressure Schur complement; augmented Lagrangian; waveform relaxation

1 Introduction

To numerically simulate unsteady flow phenomena, a commonly used approach is to first apply some time integration technique and then solve space-only problems at discrete time instances, which are coupled forward in time matching the nature of the physical process. This inherently sequential workflow naively prevents the full potential of today's supercomputers, which

consist of an enormous amount of computing units tuned to execute many computational tasks in parallel. To circumvent this barrier and increase the scaling capabilities especially when spatial parallelization techniques saturate, so called parallel-in-time solution strategies come into play, which parallelize the solution process in time so that the overall time-to-solution can be further reduced.

One famous representative of this class is Parareal first introduced by Lions, Maday, and Turinici [LMT01], where the sequential process of time integration is relaxed in an iterative manner by applying the solution strategy on subintervals in parallel and then synchronizing the auxiliary solutions in an inexpensive corrector step. This algorithm was first applied to the solution of the incompressible Navier-Stokes equations in [TP04; TP06] using a first order finite volume scheme for space discretization, whereas finite element and spectral approximations were considered by Fischer, Hecht, and Maday [FHM05]. An extension of the approach to non-isothermal flows was investigated in [CK09; MJY19], while Croce, Ruprecht, and Krause [CRK14] and Steiner et al. [Ste+15] focused on the performance of practical implementations and discussed the influence of the Reynolds number on the convergence behavior. But also other parallel-in-time methods were successfully applied to the Navier-Stokes equations. For example, the Paraexp algorithm was used in [KBG18] for rather low viscosity parameters, while Falgout et al. [Fal+14] and most recently Christopher and coworkers [Chr19; Chr+20] studied the influence of MGRIT for the Reynolds-averaged Navier-Stokes equations and in case of an adaptive space-time mesh refinement strategy, respectively.

The algorithm discussed in this work for solution of the Oseen equations is based on [DSW22; LT23] and accelerates the convergence behavior of the all-at-once pressure Schur complement (PSC) solution strategy presented therein using the augmented Lagrangian (AL) methodology. The latter approach was first introduced in [Hes69; Pow69] and modifies the momentum equation by means of the residual of the continuity equation. In this way, the solution remains the same, but very accurate and efficiently applicable PSC preconditioners can be constructed. These preconditioners are most accurate, if already existing global-in-time PSC preconditioners for the original system are incorporated in the solution process. In this work, global-in-time extensions of commonly used preconditioners like the *pressure convection diffusion (PCD)* [DSW22; LT23] and the *least squares commutator (LSC) preconditioner* [LT23] are considered for these purposes. The resulting solution strategy then converges rapidly to the solution of the Oseen equations even without the need of a coarse grid correction as proposed by the authors in [LT23].

Unfortunately, the AL methodology also comes with a disadvantage: The velocity problem involved in each Schur complement iteration becomes ill conditioned as more stabilization is introduced due to the fact that a singular matrix is added. Therefore, commonly used solution strategies fail for this problem and highly specialized versions are mandatory for an efficient solution algorithm. For example, Benzi and Olshanskii [BO06] and Schöberl [Sch99] introduced a geometric multigrid approach with an adapted relaxation and intergrid transfer for triangular meshes. In this work, an extension to quadrilateral meshes [SSW23] based on the Q_2 - P_1 finite element (FE) discretization is combined with the multigrid waveform relaxation method developed by [LO87] for efficiently solving the aforementioned global-in-time velocity problems.

The contents of the following sections are as follows: The unsteady incompressible Navier-Stokes equations are introduced in Section 2. The problem is then discretized in space and time to construct a global-in-time problem for the solution at all time instances (cf. Section 3). Section 4 summarizes the pressure Schur complement solver as presented in [LT23]. This approach is extended by the augmented Lagrangian methodology in Section 5 and numerically investigated in several linear and nonlinear test cases in Section 7. The work concludes in Section 8 with a discussion of future challenges, especially for convection-dominated fluid flows.

2 Continuous setting

The flow of a viscous and incompressible fluid through a domain $\Omega \subset \mathbb{R}^d$, $d = 2, 3$, can be described by the incompressible Navier-Stokes equations

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{v} - \nabla \cdot (2\mu(\dot{\gamma}(\mathbf{v}))\mathbf{D}(\mathbf{v})) + \text{grad}(p) = \mathbf{g}, \quad (1a)$$

$$\text{div}(\mathbf{v}) = 0, \quad (1b)$$

where $\mathbf{v} \in \mathbb{R}^d$ and $p \in \mathbb{R}$ denote the unknown velocity field and pressure variable, respectively. Here, the body force density acting on the fluid is given by $\mathbf{g} \in \mathbb{R}^d$, while the dynamic viscosity $\mu > 0$ may depend nonlinearly on the effective shear rate $\dot{\gamma} = \sqrt{2}\|\mathbf{D}(\mathbf{v})\|_{\text{F}}$ of the strain rate tensor $\mathbf{D}(\mathbf{v}) = \frac{1}{2}(\nabla \mathbf{v} + \nabla \mathbf{v}^{\top})$. In the special case of a Newtonian fluid, the viscosity parameter μ is constant so that the viscous part of the momentum equation (1a) simplifies to $-\mu\Delta \mathbf{v}$ due to the validity of the continuity equation (1b).

In the literature, there exist several different numerical solution techniques for problem (1). One common approach is to first apply a linearization technique, such as the Picard iteration or Newton's method, and then solve the resulting Oseen equations in each nonlinear iteration. For the former approach, the linearized system of partial differential equations reads

$$\frac{\partial \mathbf{v}}{\partial t} + (\bar{\mathbf{v}} \cdot \nabla) \mathbf{v} - \nabla \cdot (2\mu(\dot{\gamma}(\bar{\mathbf{v}))\mathbf{D}(\mathbf{v})) + \text{grad}(p) = \mathbf{g}, \quad (2a)$$

$$\text{div}(\mathbf{v}) = f \quad (2b)$$

for some known initial guess $\bar{\mathbf{v}}$ and a vanishing right hand side $f = 0$ of the continuity equation. In case of Newton's method, the solution update solves (2) for the right hand sides \mathbf{g} and f corresponding to the residuals of the momentum and continuity equations, respectively, while the additional term

$$\mathcal{R}(\bar{\mathbf{v}}; \mathbf{v}) = (\mathbf{v} \cdot \nabla) \bar{\mathbf{v}} - \nabla \cdot \left(4\mu'(\dot{\gamma}(\bar{\mathbf{v}))\dot{\gamma}(\bar{\mathbf{v}})^{-1}(\mathbf{D}(\bar{\mathbf{v}}) : \mathbf{D}(\mathbf{v}))\mathbf{D}(\bar{\mathbf{v}}) \right)$$

is involved in the definition of the momentum equation (2a). This linearization technique can significantly improve the nonlinear convergence behavior compared to the Picard iteration if the initial guess is sufficiently good. In the numerical examples, we therefore focus on test problems for the linear Oseen equations as occurring in the Picard iteration as well as on the nonlinear convergence behavior of Newton's method.

3 Global-in-time discrete Oseen problem

In this section, we discretize the Oseen equations in space and time and construct a global linear system of equations to compute the numerical approximation in all time instances simultaneously. For this purpose, the problem is discretized in space using the isoparametric Q_2 - P_1 FE pair, where each component of the velocity field is approximated by a continuous and piecewise biquadratic function while the pressure space is given by all piecewise linear functions on a quadrilateral mesh [BF91]. As we will see, this discretization technique allows the construction of efficient pressure Schur complement solvers using the augmented Lagrangian methodology and provides several advantages compared to an Taylor-Hood FE pair as considered in [LT23].

After discretization in space, the linear system of ordinary differential equations reads

$$\mathbf{M}_u \dot{\mathbf{u}}(t) + \mathbf{A}_u(t)\mathbf{u}(t) + \mathbf{B}\mathbf{p}(t) = \mathbf{g}(t), \quad \mathbf{B}^\top \mathbf{u}(t) = \mathbf{f}(t) \quad (3)$$

where $\mathbf{u}(t) \in \mathbb{R}^{N_u}$, $N_u \in \mathbb{N}$, and $\mathbf{p}(t) \in \mathbb{R}^{N_p}$, $N_p \in \mathbb{N}$, are the time-dependent vectors of degrees of freedom associated with the velocity field and pressure variable, while $\mathbf{g}(t) \in \mathbb{R}^{N_u}$ and $\mathbf{f}(t) \in \mathbb{R}^{N_p}$ correspond to the right hand sides of the momentum and continuity equation. The velocity mass matrix $\mathbf{M}_u \in \mathbb{R}^{N_u \times N_u}$ is computed using the Simpsons rule and is therefore diagonal. Furthermore, the matrices $\mathbf{B} \in \mathbb{R}^{N_u \times N_p}$ and $\mathbf{B}^\top \in \mathbb{R}^{N_p \times N_u}$ are the discrete counterparts of the gradient and divergence operator, respectively, while all other contributions to the momentum equation are summarized in $\mathbf{A}_u(t) \in \mathbb{R}^{N_u \times N_u}$.

To numerically approximate the solution to (3), we eventually apply the Crank-Nicolson scheme for time integration, where the time step size $\delta t > 0$ is assumed to be constant for the sake of simplicity. Then the approximations $\mathbf{u}^{(n+1)} \approx \mathbf{u}((n+1)\delta t)$ and $\mathbf{p}^{(n+1)} \approx p((n+\frac{1}{2})\delta t)$ solve the linear system of equations (c.f. [Tur99])

$$\mathbf{A}_i^{(n+1)} \mathbf{u}^{(n+1)} + \delta t \mathbf{B} \mathbf{p}^{(n+1)} = \frac{\delta t}{2} (\mathbf{g}^{(n+1)} + \mathbf{g}^{(n)}) - \mathbf{A}_e^{(n)} \mathbf{u}^{(n)}, \quad (4a)$$

$$\mathbf{B}^\top \mathbf{u}^{(n+1)} = \mathbf{f}^{(n+1)} \quad (4b)$$

at each time step $n = 0, \dots, K$, where $K \in \mathbb{N}$ denotes the total number of time steps and the system matrices $\mathbf{A}_i^{(n+1)}$ and $\mathbf{A}_e^{(n)}$ are defined by

$$\mathbf{A}_i^{(n+1)} = \mathbf{M}_u + \frac{\delta t}{2} \mathbf{A}_u^{(n+1)}, \quad \mathbf{A}_e^{(n)} = -\mathbf{M}_u + \frac{\delta t}{2} \mathbf{A}_u^{(n)}.$$

Compactly written, the velocity field $\mathbf{u}^{(n+1)}$ and the scaled pressure unknown $\tilde{\mathbf{p}}^{(n+1)} = \delta t \mathbf{p}^{(n+1)}$ solve the saddle point problem

$$\begin{pmatrix} \mathbf{A}_i^{(n+1)} & \mathbf{B} \\ \mathbf{B}^\top & \end{pmatrix} \begin{pmatrix} \mathbf{u}^{(n+1)} \\ \tilde{\mathbf{p}}^{(n+1)} \end{pmatrix} = \begin{pmatrix} \tilde{\mathbf{g}}^{(n+1)} - \mathbf{A}_e^{(n)} \mathbf{u}^{(n)} \\ \mathbf{f}^{(n+1)} \end{pmatrix}, \quad n = 0, \dots, K \quad (5)$$

using the right hand side $\tilde{\mathbf{g}}^{(n+1)} = \frac{\delta t}{2} (\mathbf{g}^{(n+1)} + \mathbf{g}^{(n)})$. The solution of this system is usually advanced in time step-by-step due to the fact that the solution $(\mathbf{u}^{(n+1)}, \tilde{\mathbf{p}}^{(n+1)})$ depends on the previous velocity field $\mathbf{u}^{(n)}$. To circumvent this barrier and provide the possibility to

parallelize the solution procedure in time, we treat all K time steps simultaneously and block the individual subproblems into a single all-at-once system of equations

$$\left(\begin{array}{cccc|ccc} \mathbf{A}_i^{(1)} & & & & \mathbf{B} & & \\ \mathbf{A}_e^{(1)} & \mathbf{A}_i^{(2)} & & & \mathbf{B} & & \\ & \ddots & \ddots & & & \ddots & \\ & & \mathbf{A}_e^{(K-1)} & \mathbf{A}_i^{(K)} & & & \mathbf{B} \\ \hline \mathbf{B}^\top & & & & & & \\ & \mathbf{B}^\top & & & & & \\ & & \ddots & & & & \\ & & & \mathbf{B}^\top & & & \end{array} \right) \begin{pmatrix} \mathbf{u}^{(1)} \\ \mathbf{u}^{(2)} \\ \vdots \\ \mathbf{u}^{(K)} \\ \tilde{\mathbf{p}}^{(1)} \\ \tilde{\mathbf{p}}^{(2)} \\ \vdots \\ \tilde{\mathbf{p}}^{(K)} \end{pmatrix} = \begin{pmatrix} \tilde{\mathbf{g}}^{(1)} - \mathbf{A}_e^{(0)} \mathbf{u}^{(0)} \\ \tilde{\mathbf{g}}^{(2)} \\ \vdots \\ \tilde{\mathbf{g}}^{(K)} \\ \mathbf{f}^{(1)} \\ \mathbf{f}^{(2)} \\ \vdots \\ \mathbf{f}^{(K)} \end{pmatrix}. \quad (6)$$

Obviously, problem (6) is equivalent to (5) and, hence, leads to the same velocity and pressure solution at all time steps. However, efficient and very robust solution techniques have to be applied to reduce the overall computational time compared to a sequential solution procedure. In the following section, a PSC approach is described, whose convergence behavior can significantly be improved by exploiting the augmented Lagrangian methodology.

4 Global-in-time pressure Schur complement iteration

In this section, we briefly summarize the PSC solution technique for the global-in-time saddle point problem (6) as it is described in [DSW22; LT23]. For this purpose, we interpret the linear system of equations as a blocked saddle point problem

$$\begin{pmatrix} \mathbf{A}_K & \mathbf{B}_K \\ \mathbf{B}_K^\top & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \tilde{\mathbf{p}} \end{pmatrix} = \begin{pmatrix} \tilde{\mathbf{g}} \\ \mathbf{f} \end{pmatrix}, \quad (7)$$

where all involved quantities are defined as expected. Multiplying the first equation of (7) by \mathbf{A}_K^{-1} and inserting the resulting expression for the velocity field $\mathbf{u} = \mathbf{A}_K^{-1}(\tilde{\mathbf{g}} - \mathbf{B}_K \tilde{\mathbf{p}})$ into the discrete counterpart of the continuity equation, we obtain the PSC equation

$$\mathbf{B}_K^\top \mathbf{A}_K^{-1} \mathbf{B}_K \tilde{\mathbf{p}} = \mathbf{B}_K^\top \mathbf{A}_K^{-1} \tilde{\mathbf{g}} - \mathbf{f} \quad (8)$$

for the unknown global-in-time pressure variable $\tilde{\mathbf{p}}$ only. Unfortunately, the involved PSC matrix $\mathbf{P}_K = \mathbf{B}_K^\top \mathbf{A}_K^{-1} \mathbf{B}_K$ is generally a dense matrix and, hence, practically impossible to compute. Therefore, iterative solution techniques are commonly exploited to solve this problem without explicitly knowing the entries of the system matrix \mathbf{P}_K . For example, the preconditioned Richardson iteration reads

$$\tilde{\mathbf{p}} \mapsto \tilde{\mathbf{p}} + \mathbf{q}, \quad \mathbf{q} = \mathbf{C}_K^{-1} (\mathbf{B}_K^\top \mathbf{A}_K^{-1} (\tilde{\mathbf{g}} - \mathbf{B}_K \tilde{\mathbf{p}}) - \mathbf{f}) \quad (9)$$

for some initial guess $\tilde{\mathbf{p}}$ while \mathbf{C}_K^{-1} is a suitable approximation to the inverse of \mathbf{P}_K . In each iteration of the solution procedure, a global linear system of equations has to be solved for computation of the auxiliary velocity field $\tilde{\mathbf{u}} = \mathbf{A}_K^{-1}(\tilde{\mathbf{g}} - \mathbf{B}_K \tilde{\mathbf{p}})$. This quantity approximates

the exact velocity solution of (7) if the norm of the global-in-time PSC residual $\mathbf{r} = \mathbf{B}_K^\top \tilde{\mathbf{u}} - \mathbf{f}$ is sufficiently small. While the efficient solution of the velocity problem is (partially) discussed in Section 6, we now define three possible preconditioners \mathbf{C}_K^{-1} , where the first two have in common that they are exact for infinitely small time step sizes δt while all three can be applied very efficiently to the space-time problem under investigation.

PCD preconditioner The first candidate of a Schur complement preconditioner is a global-in-time counterpart of the *pressure convection diffusion (PCD) preconditioner* [KLW02] and was first introduced by Danieli, Southworth, and Wathen [DSW22]. Compactly written, the preconditioner has the form [LT23]

$$\mathbf{C}_K^{-1} = (\mathbf{I}_K \otimes \mathbf{M}_p^{-1}) \mathbf{A}_{K,p} (\mathbf{I}_K \otimes \hat{\mathbf{D}}_p^{-1}), \quad (10)$$

where $\mathbf{A}_{K,p}$ is a suitable approximation of \mathbf{A}_K defined in the pressure FE space, while $\mathbf{M}_p \in \mathbb{R}^{N_p \times N_p}$ and $\hat{\mathbf{D}}_p \in \mathbb{R}^{N_p \times N_p}$ denote the pressure mass matrix and pressure Poisson matrix, respectively. According to the use of a discontinuous pressure approximation, the latter matrix is readily not defined. However, a mixed formulation of the Poisson problem can be used to justify the choice of $\hat{\mathbf{D}}_p = \mathbf{B}^\top \mathbf{M}_u^{-1} \mathbf{B}$. Note that this matrix can be explicitly determined due to the fact that \mathbf{M}_u is computed using the Simpsons rule and is therefore diagonal. For further details on the derivation and application of the PCD preconditioner, we refer to [DSW22; LT23].

LSC preconditioner An all-at-once extension of the *least-squares commutator (LSC) preconditioner* [QV17] is defined by [LT23]

$$\mathbf{C}_K^{-1} = (\mathbf{I}_K \otimes (\hat{\mathbf{D}}_p^{-1} \mathbf{B}^\top \mathbf{M}_u^{-1})) \mathbf{A}_K (\mathbf{I}_K \otimes (\mathbf{M}_u^{-1} \mathbf{B} \hat{\mathbf{D}}_p^{-1})). \quad (11)$$

This preconditioner explicitly exploits the definition of the system matrix \mathbf{A}_K and, hence, can be applied to more general Oseen equations where different components of the velocity field possibly interact with each other. This fact is especially of interest when the deformation tensor is used in the definition of the momentum equation or Newton's method is applied to linearize the incompressible Navier-Stokes equations. In case of the PCD preconditioner, a practically well-established workaround is to simply neglect these terms.

Uzawa preconditioner For a constant viscosity parameter μ and sufficiently large time increments, it is easy to show that the inverse of the PSC matrix \mathbf{P}_K is spectrally equivalent to

$$\mathbf{C}_K^{-1} = \frac{\delta t}{2} \mu \begin{pmatrix} \mathbf{M}_p^{-1} & & & & \\ \mathbf{M}_p^{-1} & \mathbf{M}_p^{-1} & & & \\ & \ddots & \ddots & & \\ & & \mathbf{M}_p^{-1} & \mathbf{M}_p^{-1} & \\ & & & & \mathbf{M}_p^{-1} \end{pmatrix} = \frac{\delta t}{2} \mu \begin{pmatrix} 1 & & & & \\ 1 & 1 & & & \\ & \ddots & \ddots & & \\ & & & 1 & 1 \end{pmatrix} \otimes \mathbf{M}_p^{-1}.$$

However, for small up to moderate time step sizes, this approximation is quite bad and, hence, will be considered only briefly in the below numerical examples.

At this point, we emphasize that all preconditioners can be applied very efficiently on modern hardware architectures due to the fact that all involved global linear systems of equations can be decomposed into K independent subproblems. The spatial system matrices to be inverted therein are the same over the whole time interval, so each global problem is equivalent to a single space-only problem with K right hand sides. This allows the use of very efficient solution strategies based on vector operations with reduced data communication.

5 Augmented Lagrangian methodology

In the previous section, we summarized the idea of a global-in-time pressure Schur complement iteration, which solves the discretized Oseen equations for all time steps simultaneously. The involved preconditioners for the pressure-only problem are designed as straightforward extensions of their sequential counterparts and provide a high degree of parallelism. However, we will observe in Section 7 that the convergence behavior of this basic Schur complement iteration is quite slow and even deteriorates as the spatial resolution increases or more and more time steps are blocked. To improve the solution methodology, the authors proposed in [LT23] the incorporation of a coarse grid correction in the framework of a hierarchical multigrid approach. While this technique relaxes the dependency of the rate of convergence on the above mentioned discretization parameters, the efficient solution of the resulting coarse grid problem is still an unsolved problem.

In this work, we extend the PSC iteration by the augmented Lagrangian methodology. The approach was originally developed independently by Hestenes [Hes69] and Powell [Pow69] and aims to improve the convergence behavior without the need of a coarse grid correction by modifying the discrete momentum equation in a strongly consistent manner. This guarantees that the solution of the discrete problem remains the same, while (theoretically) arbitrarily accurate PSC preconditioners can be constructed. The success of this technique is due to the *Sherman-Morrison-Woodbury identity* [Hag89]

$$(A + \gamma UCV)^{-1} = A^{-1} - A^{-1}U(\gamma^{-1}C^{-1} + VA^{-1}U)^{-1}VA^{-1} \quad (12)$$

holding true for any invertible matrices A , C , and $C^{-1} + VA^{-1}U$ while γ is some positive parameter. Roughly speaking, equation (12) describes the influence of a (possibly singular) perturbation on the inverse of a non-singular matrix in an additive fashion. If A is chosen as a spatial system matrix A_i (we drop the superscript $(n+1)$ for the time being) and the identity is multiplied from left and right by B^\top and B , respectively, the first expression on the right hand side of (12) coincides with the PSC matrix of (5) while the left hand side is the PSC matrix using a modified system matrix $A_{i,\gamma} = A_i + \gamma UCV$. Therefore, the PSC matrix $P_{i,\gamma} = B^\top A_{i,\gamma}^{-1} B$ is the sum of $P_i = B^\top A_i^{-1} B$ and some additive term. The crucial question is now how to define the matrices U , V , and C so that the inverse of $P_{i,\gamma}$ can be approximated more accurately than the inverse of P_i while at the same time the action of the approximation can be applied very efficiently. One possibility is the choice of $U = B$, $V = B^\top$, and $C = \delta t M_p^{-1}$, which does not even extend the sparsity pattern of the system matrix A_i due to the special choice of a discontinuous pressure FE space. For this definition,

simple algebraic manipulations and another application of the Sherman-Morrison-Woodbury identity (12) lead to the ingenious identity (cf. [Wec19, Lemma 5.2])

$$\mathbf{P}_{i,\gamma}^{-1} = \mathbf{P}_i^{-1} + \gamma\delta t \mathbf{M}_p^{-1}.$$

Thus, the inverse of $\mathbf{P}_{i,\gamma}$ converges to $\gamma\delta t \mathbf{M}_p^{-1}$ as γ increases. This additive term can be applied exactly and does not need to be approximated in any way due to the fact that \mathbf{M}_p is block-diagonal. This guarantees that very accurate approximations $\mathbf{C}_{i,\gamma}^{-1} = \mathbf{C}_i^{-1} + \gamma\delta t \mathbf{M}_p^{-1}$ of $\mathbf{P}_{i,\gamma}^{-1}$ can be constructed by adapting the value of γ even if the approximation \mathbf{C}_i^{-1} of \mathbf{P}_i^{-1} is quite inexact. However, adding $\gamma\delta t \mathbf{B} \mathbf{M}_p^{-1} \mathbf{B}^\top$ to the system matrix \mathbf{A}_i generally modifies the solution of the saddle point problem, which is especially not acceptable for large values of γ . This downside can be avoided by simultaneously adapting the right hand side of the momentum equation in a strongly consistent way. More precisely, we replace (5) by

$$\begin{pmatrix} \mathbf{A}_{i,\gamma}^{(n+1)} & \mathbf{B} \\ \mathbf{B}^\top & \end{pmatrix} \begin{pmatrix} \mathbf{u}^{(n+1)} \\ \tilde{\mathbf{p}}^{(n+1)} \end{pmatrix} = \begin{pmatrix} \tilde{\mathbf{g}}_\gamma^{(n+1)} - \mathbf{A}_e^{(n)} \mathbf{u}^{(n)} \\ \mathbf{f}^{(n+1)} \end{pmatrix} \quad (13)$$

using the velocity system matrix $\mathbf{A}_{i,\gamma}^{(n+1)} = \mathbf{A}_i^{(n+1)} + \gamma\delta t \mathbf{B} \mathbf{M}_p^{-1} \mathbf{B}^\top$ and the modified right hand side $\tilde{\mathbf{g}}_\gamma^{(n)} = \tilde{\mathbf{g}}^{(n)} + \gamma\delta t \mathbf{B} \mathbf{M}_p^{-1} \mathbf{f}^{(n)}$. Then the first equation of (13) reads

$$\mathbf{A}_i^{(n+1)} \mathbf{u}^{(n+1)} + \mathbf{B} \tilde{\mathbf{p}}^{(n+1)} = \tilde{\mathbf{g}}^{(n+1)} - \mathbf{A}_e^{(n)} \mathbf{u}^{(n)} + \underbrace{\gamma\delta t \mathbf{B} \mathbf{M}_p^{-1} (\mathbf{f}^{(n+1)} - \mathbf{B}^\top \mathbf{u}^{(n+1)})}_{=0},$$

where the last expression vanishes due to the fact that $\mathbf{B}^\top \mathbf{u}^{(n+1)} = \mathbf{f}^{(n+1)}$ is satisfied by the second equation. Therefore, the saddle point problem (13) is equivalent to (5) while the inverse of its PSC matrix can be approximated very efficiently for large values of γ .

The augmented Lagrangian methodology described above can readily be extended to the global-in-time saddle point problem (6). In this case, only the block-diagonal of \mathbf{A}_K is modified and results in the all-at-once linear system of equations

$$\begin{pmatrix} \mathbf{A}_K + \gamma\delta t \mathbf{B}_K \mathbf{M}_{K,p}^{-1} \mathbf{B}_K^\top & \mathbf{B}_K \\ \mathbf{B}_K^\top & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \tilde{\mathbf{p}} \end{pmatrix} = \begin{pmatrix} \tilde{\mathbf{g}} + \gamma\delta t \mathbf{B}_K \mathbf{M}_{K,p}^{-1} \mathbf{f} \\ \mathbf{f} \end{pmatrix}$$

$$: \Leftrightarrow \left(\begin{array}{cccc|ccc} \mathbf{A}_{i,\gamma}^{(1)} & & & & \mathbf{B} & & \\ \mathbf{A}_e^{(1)} & \mathbf{A}_{i,\gamma}^{(2)} & & & & \mathbf{B} & \\ & \ddots & & & & & \ddots \\ & & \mathbf{A}_e^{(K-1)} & \mathbf{A}_{i,\gamma}^{(K)} & & & \mathbf{B} \\ \hline \mathbf{B}^\top & & & & & & \\ & \mathbf{B}^\top & & & & & \\ & & \ddots & & & & \\ & & & \mathbf{B}^\top & & & \end{array} \right) \begin{pmatrix} \mathbf{u}^{(1)} \\ \mathbf{u}^{(2)} \\ \vdots \\ \mathbf{u}^{(K)} \\ \tilde{\mathbf{p}}^{(1)} \\ \tilde{\mathbf{p}}^{(2)} \\ \vdots \\ \tilde{\mathbf{p}}^{(K)} \end{pmatrix} = \begin{pmatrix} \tilde{\mathbf{g}}_\gamma^{(1)} - \mathbf{A}_e^{(0)} \mathbf{u}^{(0)} \\ \tilde{\mathbf{g}}_\gamma^{(2)} \\ \vdots \\ \tilde{\mathbf{g}}_\gamma^{(K)} \\ \mathbf{f}^{(1)} \\ \mathbf{f}^{(2)} \\ \vdots \\ \mathbf{f}^{(K)} \end{pmatrix},$$

where $\mathbf{M}_{K,p} = \mathbf{I}_K \otimes \mathbf{M}_p$ denotes the global-in-time pressure mass matrix. Using the same argumentation as above, we conclude that the inverse of the all-at-once PSC matrix can be approximated by

$$\mathbf{P}_{K,\gamma}^{-1} = \mathbf{P}_K^{-1} + \gamma\delta t \mathbf{M}_{K,p}^{-1} \approx \mathbf{C}_K^{-1} + \gamma\delta t \mathbf{M}_{K,p}^{-1} =: \mathbf{C}_{K,\gamma}^{-1},$$

Space-only Smoother The relaxation method presented in [SSW23] is based on an additive block preconditioner, where local subproblems for each spatial node of the triangulation have to be solved. More precisely, the subproblems refer to the parts of the linear system of equations $A_{i,\gamma}u = h$ whose degrees of freedom are associated with points in the interior of nodal patches. For their definition, we introduce the subset $\varpi_j \subseteq \Omega$ as the composition of all elements K^e containing mesh node \mathbf{x}_j , i.e.,

$$\varpi_j = \bigcup_{e, \mathbf{x}_j \in K^e} K^e, \quad j \in \{1, \dots, N_n\},$$

where $N_n \in \mathbb{N}$ is the total number of mesh nodes. Then the index set \mathcal{N}_j is defined as

$$\mathcal{N}_j = \{k \in \{1, \dots, N_u\} : \tilde{\mathbf{x}}_k \in \hat{\varpi}_j\}, \quad j \in \{1, \dots, N_n\},$$

where $\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_{N_u}$ are the locations associated with the biquadratic basis functions $\varphi_1, \dots, \varphi_{N_u}$ of the velocity FE space, i.e., $\varphi_j(\tilde{\mathbf{x}}_k) = \delta_{jk}$, and $\hat{\varpi}_j$ is the interior of ϖ_j , i.e., $\hat{\varpi}_j = \varpi_j \setminus \partial\varpi_j$. Now the additive block preconditioner $D_i^{-1} \approx A_{i,\gamma}^{-1}$ solves local subproblems for the degrees of freedom in \mathcal{N}_j and sums up the local solutions for all spatial nodes of the mesh, i.e.,

$$D_i^{-1} = \sum_{j=1}^{N_n} \mathcal{I}_j (\mathcal{I}_j^\top A_{i,\gamma} \mathcal{I}_j)^{-1} \mathcal{I}_j^\top. \quad (16)$$

Here, the injection operator $\mathcal{I}_j \in \mathbb{R}^{N_u \times |\mathcal{N}_j|}$ maps a vector of unknowns associated with the local patch ϖ_j to a vector of global degrees of freedoms and introduces zero entries in components that are not included in \mathcal{N}_j . Note that all degrees of freedoms associated with cells or edges of the mesh are treated several times and the updates are *not* averaged in the gathering progress.

Space-only intergrid transfer For definition of the specialized intergrid transfer operators, let P be the prolongation operator naturally induced by the Q_2 FE space. This quantity should not be mixed up with the PSC matrices $P_{i,\gamma}$ and P_i introduced in previous sections. Then the robust intergrid transfer operators are defined by

$$\tilde{P} = \left(I - \sum_{j=1}^{\bar{N}_n} \mathcal{I}_j (\mathcal{I}_j^\top A_{i,\gamma} \mathcal{I}_j)^{-1} (\mathcal{I}_j^\top (\gamma \delta t B M_p^{-1} B^\top) \mathcal{I}_j) \mathcal{I}_j^\top \right) P, \quad \tilde{R} = \tilde{P}^\top \quad (17)$$

where $\mathbf{x}_1, \dots, \mathbf{x}_{\bar{N}_n}$ are assumed to be the common mesh nodes of the coarse and fine triangulation. This definition is consistent in the way that $P = \tilde{P}$ holds true for $\gamma = 0$. Furthermore, the local subproblems occurring in the definition of \tilde{P} are also involved in the definition of D_i^{-1} and, hence, may only need to be computed once.

We now extend the relaxation technique as well as the intergrid transfer operators to the global-in-time problem (14), so they can be exploited in the context of multigrid waveform relaxation.

Global-in-time smoother The additive block preconditioner (16) is extended to the all-at-once problem by solving subproblems on local patches for all time steps simultaneously without synchronizing the solutions in between. This guarantees that all subproblems can be solved independently of each other and preserves the methodology of the time-simultaneous (multigrid) approach providing a high degree of parallelization. In detail, the *additive block preconditioner* reads

$$\mathbf{D}_K^{-1} = \sum_{j=1}^{N_n} (\mathbf{I}_K \otimes \mathcal{I}_j) ((\mathbf{I}_K \otimes \mathcal{I}_j) \mathbf{A}_K (\mathbf{I}_K \otimes \mathcal{I}_j^\top))^{-1} (\mathbf{I}_K \otimes \mathcal{I}_j^\top).$$

In the numerical examples presented below, this preconditioner is embedded into the GMRES method without restart to increase the robustness of the solution strategy.

Global-in-time intergrid transfer The robust space-time prolongation and restriction operators are given by

$$\tilde{\mathbf{P}}_K = \begin{pmatrix} \tilde{\mathbf{P}}^{(1)} & & & \\ & \tilde{\mathbf{P}}^{(2)} & & \\ & & \ddots & \\ & & & \tilde{\mathbf{P}}^{(K)} \end{pmatrix}, \quad \tilde{\mathbf{R}}_K = \tilde{\mathbf{P}}_K^\top$$

exploiting the sequential versions of the robust intergrid transfer operators as defined in (17). In contrast to the relaxation scheme, these operators only use the block diagonal entries of \mathbf{A}_K , but must be assembled for each time step individually. This makes the customized grid transfer more expensive than the one naturally induced by the FE space.

7 Numerical examples

In this section, we solve several numerical test problems using the augmented Lagrangian technique to accelerate the convergence behavior of the pressure Schur complement iteration as described in Section 4. This iterative solver is embedded into a GMRES method, which is restarted after four inner iterations. To quantify the efficiency of the solver at hand, the total number of *inner* iterations is counted. This value (approximately) measures how often the implicitly defined system matrix and the involved Schur complement preconditioner are applied till a certain tolerance in accuracy is reached for a vanishing initial guess.

The numerical examples are computed on the ‘flow around a cylinder’ domain as illustrated in Fig. 1, where all external forces are set to zero. Furthermore, homogeneous Neumann (‘do nothing’) boundary conditions are prescribed on $\Gamma_{\text{out}} = \partial\Omega \setminus (\Gamma_{\text{in}} \cup \Gamma_{\text{wall}})$ while no slip boundary conditions are enforced on Γ_{wall} . Therefore, the dynamics of the fluid are solely governed by the inflow boundary data

$$\mathbf{v}_{\text{in}} = U(t) \frac{4y(0.41 - y)}{0.41^2} \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad U(t) = U_0 |\sin(\frac{\pi}{8}t)|, \quad U_0 = 0.3 \quad \text{on } \Gamma_{\text{in}}$$

and the initial data $\mathbf{v}_0 = \mathbf{0}$ of the velocity field \mathbf{v} (cf. [LT23]).

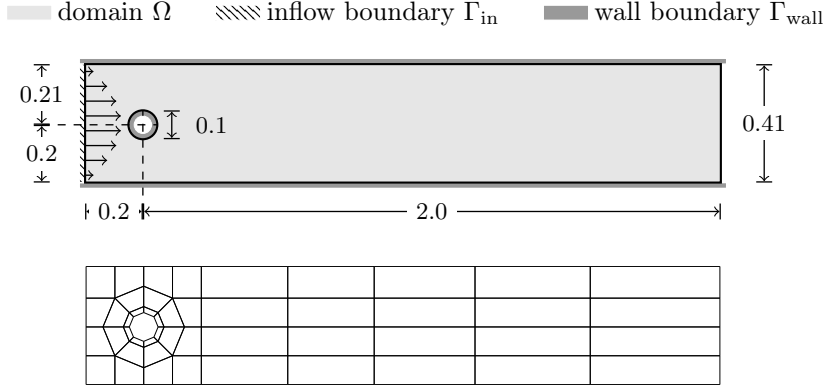


Figure 1: Computational domain including partition of boundary and triangulation on coarsest mesh level 0.

As mentioned above, the problem is discretized in space using the Q_2 - P_1 FE pair, while the time integrator is given by the Crank-Nicolson scheme using a fixed time step size of $\delta t = \frac{1}{25} \cdot 2^{-\text{lvl}}$, where $\text{lvl} \in \{0, \dots, 5\}$ denotes the number of uniform refinements of the triangulation at hand (cf. Fig. 1). Note that the convergence behavior of the iterative solution strategy using the PCD or LSC preconditioner can be significantly improved by choosing smaller time increments because these preconditioners then become more accurate. However, this property of the solver will not be investigated in this work and we refer to [LT23] for numerical results in this regard.

7.1 Stokes problem

The first test case is given by the Stokes problem, where no convective contribution is involved in the definition of the momentum equation and the viscosity parameter is fixed to $\mu = 10^{-2}$. In this case, we first discuss the influence of the AL parameter on the norm of the PSC residual (for a fixed solution) to define a reasonable stopping criterion for the iterative solution procedure. Figure 2 illustrates the convergence behavior of the GMRES method for the PCD preconditioner and different values of γ , where the preconditioned PSC residual is evaluated using another choice of the AL parameter, denoted by $\tilde{\gamma}$. While a rapid convergence can be observed for all considered variants of the PSC residual if γ is chosen sufficiently large, the relative norm of the preconditioned PSC residual depends only marginally on $\tilde{\gamma}$. Thus, this quantity seems to be an adequate measure for the quality of the solution and reducing it by a certain value is a reasonable stopping criterion. Fortunately, the preconditioned PSC residual is already determined in GMRES implementations frequently used in the literature so that no additional computations are required. In all linear test cases considered in this work, the GMRES method therefore terminates if

$$\left\| \mathbf{C}_K^{-1} (\mathbf{B}_K^\top \mathbf{A}_K^{-1} (\tilde{\mathbf{g}} - \mathbf{B}_K \tilde{\mathbf{p}}) - \mathbf{f}) \right\|_2 < \text{tol}_{\text{rel}} \left\| \mathbf{C}_K^{-1} (\mathbf{B}_K^\top \mathbf{A}_K^{-1} \tilde{\mathbf{g}} - \mathbf{f}) \right\|_2 \quad (18)$$

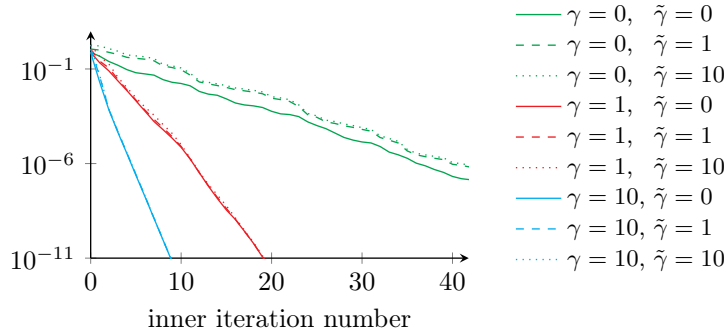


Figure 2: Stokes problem: History of norm of preconditioned PSC residual evaluated using different AL parameters $\tilde{\gamma}$ for PCD preconditioner on mesh level 2 for 800 blocked time steps.

Table 1: Stokes problem: Total number of iterations on mesh level 2 for 800 blocked time steps and different PSC preconditioners.

preconditioner \ γ	10^{-1}	10^0	10^1	10^2	10^3	10^4
PCD precon.	31	20	9	6	4	3
LSC precon.	63	36	18	7	5	3
Uzawa precon.	21 323	2198	314	48	12	6

is satisfied for a tolerance of $\text{tol}_{\text{rel}} = 10^{-11}$.

In Table 1, the total number of (inner) iterations required to reach this stopping criterion is summarized for different PSC preconditioners. As expected, this value decreases for all solution strategies under investigation as γ increases. However, the GMRES method exploiting the Uzawa preconditioner is not competitive with the other algorithms for moderate choices of the AL parameter γ because this preconditioner approximates the inverse of the PSC matrix \mathbf{P}_K less accurately. Although only (block-diagonal) pressure mass matrices have to be inverted for its application, each GMRES iteration is still costly due to the fact that an auxiliary velocity field must be computed whenever the PSC residual is determined. This makes the Uzawa preconditioner less interesting in practical applications and explains why it is not considered further in the following investigations.

Next, we vary the total number of blocked time steps K and discuss its influence on the convergence behavior for the Stokes problem (cf. Fig. 3). While the GMRES method using the PCD preconditioner requires more and more iterations as the length of the time interval increases, the rate of convergence for the LSC preconditioner seems to be bounded above independently of this value even for the unstabilized Schur complement equation. However, introduction of AL stabilization significantly reduces the total number of iterations and especially relaxes the dependency on K for the PCD preconditioner. The same behavior can be observed for the solution technique using the LSC preconditioner as the spatial resolution

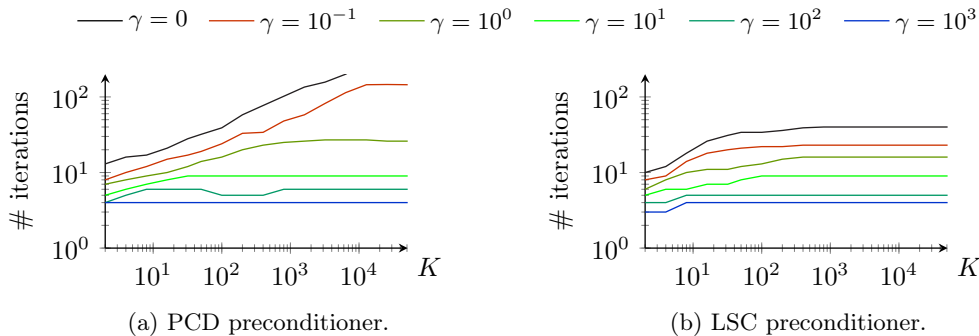


Figure 3: Stokes problem: Total number of iterations on mesh level 0 for different numbers of blocked time steps.

Table 2: Stokes problem: Total number of iterations on different mesh levels for 800 blocked time steps.

lvl \ γ	PCD preconditioner						LSC preconditioner					
	0	10^{-1}	10^0	10^1	10^2	10^3	0	10^{-1}	10^0	10^1	10^2	10^3
0	102	48	25	9	6	4	40	23	16	9	5	4
1	76	35	23	9	6	4	66	39	26	13	6	4
2	65	31	20	9	6	4	102	63	36	18	7	5
3	48	22	17	9	6	4	155	96	49	25	9	5
4	37	20	14	9	6	5	219	139	67	36	12	6
5	30	17	12	9	6	5	280	190	100	46	17	7

increases, as shown in Table 2. While the convergence behavior deteriorates on finer meshes, the influence becomes less significant for larger values of γ . In contrast to this, the algorithm exploiting the PCD preconditioner even improves as the spatial resolution increases. However, a rapid convergence can be observed for both preconditioners only for a great amount of AL stabilization.

7.2 Oseen problem for Newtonian fluid

After investigating the Stokes equations, we now add convectivity to the definition of the momentum equation and study the resulting Oseen problem for two different viscosity parameters, namely $\mu = 10^{-2}$ and $\mu = 10^{-3}$. In these cases, the velocity field involved in the convective contribution is set to the solution of the incompressible Navier-Stokes equations (1), so the linear problem at hand coincides with the final Picard iteration to solve the underlying nonlinear problem. Although no vortex shedding is to be expected behind the cylinder for both viscosity parameters as the Reynolds number does not exceed $\text{Re} = 2$ or $\text{Re} = 20$ (cf. [Sch+96]), the convective contribution is more dominant for $\mu = 10^{-3}$ and

significantly affects the convergence behavior of the GMRES method (cf. Table 3). Without any AL stabilization, the total number of inner iterations using the PCD preconditioner increases as more and more time steps are blocked. This dependency becomes more dominant as the viscosity parameter goes to zero and possibly even leads to a non-converging solution procedure. For the LSC preconditioner, this effect is reversed and the solver tends to improve for smaller viscosity parameters. However, more and more iterations are required as the spatial resolution increases, which perfectly fits the observations for the Stokes equations.

These dependencies become less significant as AL stabilization is introduced in the solution procedure and the total number of inner GMRES iterations goes to zero as γ increases for all configurations under investigation. In case of the largest value of the AL parameter $\gamma = 10^3$, only 3 – 6 iterations are required to solve the linear system of equations. In this case, the use of the PCD preconditioner seems to be preferable due to the fact that its application is two times less expensive than the application of the LSC preconditioner (cf. [LT23]).

7.3 Oseen problem for Carreau-Yasuda viscosity model

In the final linear test problem, the Newtonian fluid rheology exploiting a constant viscosity parameter μ is replaced by a non-Newtonian one as described by the Carreau-Yasuda model [YAC81]

$$\mu(\dot{\gamma}) = \mu_\infty + (\mu_0 - \mu_\infty)(1 + (\lambda\dot{\gamma})^a)^{(n-1)/a}. \quad (19)$$

This nonlinear definition of the viscosity parameter simplifies to the well-known Carreau model [Car72] for $a = 2$ and simulates shear thinning effects for $n < 1$. In the following numerical studies, the other material parameters are chosen as $\mu_\infty = 0$, $\lambda = 1$, and $n = 0.31$, while the maximum viscosity parameter μ_0 is set to 10^{-2} and 10^2 leading to an effective viscosity parameter $\mu(\dot{\gamma}) \in [0.0004, 0.01]$ and $\mu(\dot{\gamma}) \in [5.8126, 100]$, respectively, as it might occur in the flow behavior of thermoplastics. Therefore, the ratio between the maximum and minimum viscosity parameters is approximately the same and only the type of nonlinearity changes by modifying μ_0 . Another consequence of the varying viscosity parameter is the fact that the viscous part of the momentum equation does not simplify to $\mu\Delta\mathbf{v}$. Therefore, different components of the velocity field are coupled with each other even in case of the Picard iteration and only the LSC preconditioner is readily applicable to this fluid rheology.

In Table 4, the total number of inner GMRES iterations is summarized for this preconditioner and different amounts of AL stabilization. As already observed in the previous test problems, the total number of iterations required to reach the stopping criterion grows on finer meshes, while the length of the time horizon has only a minor influence on the rate of convergence. As μ_0 increases, solution of the Oseen problem becomes more expensive on mesh level 4. For both values of μ_0 , however, the total number of iterations can be drastically reduced by introducing AL stabilization and goes down to 4 – 6 for the considered values of γ . Note that the range of AL parameters is shifted by four orders of magnitude for $\mu_0 = 10^2$ compared to $\mu_0 = 10^{-2}$ to achieve a similar effect in the convergence behavior. Therefore, the amount of required AL stabilization to significantly reduce the number of GMRES iterations is mainly coupled to the viscous term of the momentum equation for this test case.

Table 3: Oseen problem for Newtonian fluid: Total number of iterations.

(a) PCD preconditioner.													
K	lvl \ γ	$\mu = 10^{-2}$						$\mu = 10^{-3}$					
		0	10^{-1}	10^0	10^1	10^2	10^3	0	10^{-1}	10^0	10^1	10^2	10^3
100	0	41	24	18	10	5	4	88	19	12	7	5	3
100	2	29	17	12	10	6	4	20	9	8	7	5	4
100	4	25	13	9	8	7	5	17	8	7	6	5	4
400	0	106	56	27	10	6	4	–	31	14	7	5	3
400	2	48	24	18	9	6	4	–	19	12	7	5	3
400	4	35	17	12	9	6	5	26	9	8	7	5	4
1600	0	706	102	33	10	6	4	–	63	16	7	5	3
1600	2	104	54	26	10	6	4	–	29	15	7	5	4
1600	4	53	23	18	9	6	5	–	19	13	7	5	4

(b) LSC preconditioner.													
K	lvl \ γ	$\mu = 10^{-2}$						$\mu = 10^{-3}$					
		0	10^{-1}	10^0	10^1	10^2	10^3	0	10^{-1}	10^0	10^1	10^2	10^3
100	0	31	21	12	8	5	4	24	9	6	5	4	3
100	2	71	47	24	15	8	5	30	16	9	7	5	3
100	4	99	75	43	19	12	6	43	22	14	8	6	4
400	0	34	22	15	9	5	4	33	10	7	5	4	3
400	2	87	57	30	17	7	5	41	20	11	8	5	4
400	4	168	114	60	33	12	6	67	36	16	11	6	4
1600	0	35	22	15	9	5	4	43	10	7	5	4	3
1600	2	100	64	40	18	8	5	51	23	16	8	5	4
1600	4	254	160	75	38	12	6	100	47	23	13	6	4

Table 4: Oseen problem for Carreau-Yasuda viscosity model: Total number of iterations for LSC preconditioner.

K	lvl \ γ	$\mu_0 = 10^{-2}$						$\mu_0 = 10^2$					
		0	10^{-1}	10^0	10^1	10^2	10^3	0	10^3	10^4	10^5	10^6	10^7
100	0	32	16	10	7	5	4	30	16	12	7	5	4
100	2	65	28	17	13	7	5	49	44	36	14	7	4
100	4	94	58	36	18	12	6	154	148	110	37	10	5
400	0	42	19	13	8	5	4	30	17	12	7	5	4
400	2	78	33	22	14	7	5	60	45	36	14	7	4
400	4	106	75	40	28	12	6	221	166	126	34	10	6
1600	0	44	19	13	8	5	4	30	17	12	7	5	4
1600	2	86	43	32	15	7	5	60	46	36	14	7	4
1600	4	159	89	52	29	12	6	236	185	127	33	10	6

7.4 Nonlinear Carreau-Yasuda viscosity model

So far, only linear test problems were studied as they occur in the final iteration of a Picard iteration for solution of the incompressible Navier-Stokes equations. However, this nonlinear solver converges only slowly to the solution, especially for the considered space-time problems (cf. [LT23]), so more efficient linearization techniques are required to outperform sequential solution strategies. In case of Newton's method, additional terms must be considered in the definition of the momentum equation resulting in a stronger coupling between different components of the velocity field. However, the above mentioned linear Schur complement iteration as well as the augmented Lagrangian methodology are still applicable in this case and result in an efficient solution strategy as we will see in this section.

To analyse the nonlinear convergence for (an inexact) Newton's method, we again solve the incompressible Navier-Stokes equations for the Carreau-Yasuda viscosity model and the maximum viscosity coefficient $\mu_0 = 10^2$. Firstly, all involved linear subproblems are solved exactly to illustrate the performance of Newton's method and measure the minimum number of nonlinear iterations required for this test problem. In Fig. 4, the nonlinear convergence behavior is illustrated for different values of the AL parameter by plotting the norm of the (unpreconditioned) PSC residual. Due to the fact that the AL stabilization does not modify the solution of the Oseen equations, the nonlinear convergence behavior is exactly the same for all AL parameters under investigation. However, the norm of the residual is shifted for different values of γ and may lead to an inaccurate solution caused by ill-conditioned problems if γ is chosen too large. Therefore, the AL parameter cannot practically be chosen too large without compromising the accuracy of the solution.

The exact Oseen solver is now replaced by the above mentioned GMRES method based on the Schur complement equation and the augmented Lagrangian approach, while the (inexact) Newton method is stopped if the relative norm of the PSC residual is smaller than 10^{-10} . In contrast to the above investigations for linear test problems, different relative

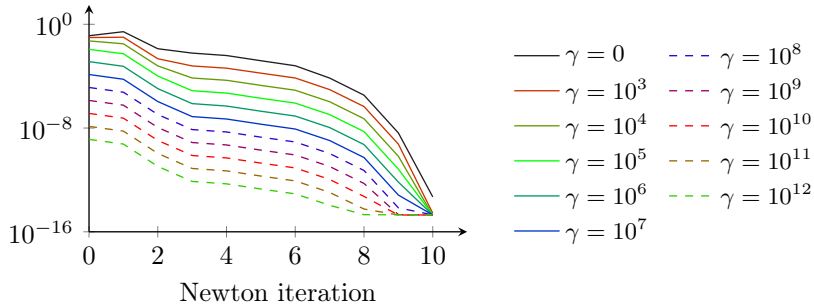


Figure 4: Carreau-Yasuda problem for $\mu_0 = 10^2$: History of norm of PSC residual for Newton’s method on mesh level 2 using 800 blocked time steps.

Table 5: Carreau-Yasuda problem for $\mu_0 = 10^2$: Total number of nonlinear (left column) and linear (right column) iterations on mesh level 2 for 800 blocked time steps and different linear stopping criteria.

$\text{tol}_{\text{rel}} \setminus \gamma$	0		10^3		10^4		10^5		10^6		10^7	
	nonl.	lin.	nonl.	lin.	nonl.	lin.	nonl.	lin.	nonl.	lin.	nonl.	lin.
10^{-1}	14	69	14	54	12	38	12	17	10	10	10	10
10^{-2}	10	96	10	55	11	47	10	17	10	10	10	10
10^{-3}	11	156	10	77	10	56	10	22	10	13	10	10
10^{-4}	10	186	10	102	10	68	10	29	10	17	10	13
10^{-5}	10	239	10	120	10	87	10	36	10	23	10	16
10^{-6}	10	287	10	148	10	101	10	42	10	24	10	21
10^{-7}	10	341	10	179	10	120	10	53	10	33	10	23

tolerances tol_{rel} are used in the stopping criterion (18) to determine the most efficient overall solution strategy.

In Table 5, the number of nonlinear Newton steps as well as the total number of inner GMRES iterations are summarized for different values of γ and tol_{rel} . For an overresolved computation of the solution to the Oseen problem, the number of nonlinear iterations coincides, as expected, with the total number of (exact) Newton steps for this problem. However, these configurations without AL stabilization require an enormous number of linear GMRES iterations and, hence, are practically very inefficient. By relaxing the tolerance for the linear subproblems, more nonlinear steps may be required, but the total number of GMRES iterations reduces. Therefore, the overall solution strategy becomes more efficient. This procedure can be further accelerated by introducing AL stabilization. In this case, all linear subproblems are solved with sufficient accuracy even with a single inner GMRES iteration for $\gamma \geq 10^6$ and $\text{tol}_{\text{rel}} \leq 10^{-2}$. Under the assumption that the cost for each GMRES iteration is independent of the choice of the AL parameter, this is the most efficient solver considered in this work.

Table 6: Velocity problem for Carreau-Yasuda viscosity model and $\mu_0 = 10^2$: Total number of specialized multigrid iterations on mesh level 2 for 200 blocked time steps.

$\nu_1 = \nu_2 \setminus \gamma$	0	10^{-2}	10^{-1}	10^0	10^1	10^2	10^3	10^4	10^5	10^6	10^7
1	50	50	51	51	54	279	–	–	–	–	–
2	19	19	19	19	17	55	78	171	210	174	126
4	9	9	9	9	10	23	24	42	53	49	52
8	6	6	6	6	6	12	13	18	21	22	22
16	5	5	5	5	5	8	8	8	10	10	10
32	4	4	4	4	4	6	7	6	6	6	6

7.5 Velocity solver for Carreau-Yasuda viscosity model

As mentioned above, the linear solution strategy proposed in this work can only perform well if the auxiliary velocity field involved in the definition of the PSC residual can be computed efficiently. In Section 6, a highly specialized multigrid solver is presented, which solves the velocity problems using customized techniques for relaxation and intergrid transfer. This solver is now applied to the Carreau-Yasuda test problem as occurring in the final step of Newton’s method. This means that the exact solution vanishes due to homogeneous boundary data and a zero vector for the right hand side. For validation purposes, the initial guess is given by the exact velocity field of the incompressible Navier-Stokes equations and the solution is accepted if 11 digits in the Euclidean norm of the residual are gained within a maximum number of 500 iterations.

In Table 6, the number of multigrid iterations is summarized for different numbers of pre- and post-smoothing steps ν_1 and ν_2 , respectively. Here, the solution is considered on mesh level 2 for $K = 200$ blocked time steps. While the specialized solver diverges for a large amount of AL stabilization and few smoothing steps, the convergence behavior significantly improves for more relaxation steps. Especially for $\nu_1, \nu_2 \geq 16$, the numerical effort for solving the space-time problem depends only marginally on the AL parameter γ and solution of the stabilized velocity problem is at most two times as expensive as in case of $\gamma = 0$.

This observation stays valid on finer meshes and for longer time intervals, as summarized in Table 7. While no more than 10 multigrid iterations are required to solve the problem at hand if the highly specialized relaxation and intergrid transfer are used, the rate of convergence for a commonly used multigrid waveform relaxation algorithm using the canonical prolongation and restriction operators as well as the block Jacobi relaxation (cf. [Dün+21]) drastically deteriorates for $\gamma \geq 10^0$. Thus, the latter approach fails in the context of the Oseen equations with AL stabilization, where $\gamma \geq 10^6$ seems to be a reasonable choice for efficiently reducing the overall numerical complexity (cf. Table 5).

8 Outlook

The numerical investigations presented in Section 7 focused on test cases for small up to moderate Reynolds numbers. In these cases, the dynamics of the fluid are mainly prescribed

Table 7: Velocity problem for Carreau-Yasuda viscosity model and $\mu_0 = 10^2$: Total number of iterations for different multigrid solvers using 16 pre- and post-smoothing steps.

K	lvl \ γ	standard multigrid						specialized multigrid					
		0	10^{-1}	10^0	10^1	10^2	10^3	0	10^3	10^4	10^5	10^6	10^7
100	2	7	7	7	9	57	—	5	8	8	9	9	9
100	4	8	8	8	9	71	—	5	10	8	10	9	9
400	2	7	7	9	26	401	—	5	8	9	10	10	10
400	4	8	8	8	42	—	—	5	9	9	10	10	10
1600	2	7	8	13	139	—	—	6	8	9	10	10	10
1600	4	9	10	32	—	—	—	6	8	9	10	10	10

Table 8: Oseen problem for Bench3 configuration: Total number of iterations.

K	lvl \ γ	PCD preconditioner						LSC preconditioner					
		0	10^{-1}	10^0	10^1	10^2	10^3	0	10^{-1}	10^0	10^1	10^2	10^3
100	0	—	95	33	11	6	4	112	21	9	5	4	3
100	2	67	15	10	8	5	4	34	15	9	6	5	3
100	4	19	8	7	6	5	4	43	22	14	8	6	4
400	0	—	181	40	11	6	4	369	27	11	6	4	3
400	2	—	75	32	11	6	4	307	20	10	7	5	4
400	4	—	15	10	8	5	4	64	34	16	11	6	4
1600	0	—	—	53	11	6	4	—	32	12	6	4	3
1600	2	—	143	36	11	6	4	450	23	12	8	5	4
1600	4	—	76	32	11	6	4	132	38	21	12	6	4

by the viscous part of the momentum equation and the initial condition only slightly influences the final solution. For this kind of problems, many parallel-in-time algorithms perform well and significant speedups compared to sequential solution strategies can be observed. This behavior changes as the convective term becomes more dominant, for instance, when focusing on the third test case of the well-known flow around a cylinder benchmark for $\mu = 10^{-3}$ and $U_0 = 1.5$ [Sch+96]. In this case, the Reynolds number attains the maximum value of $Re = 100$ at $t = 4$ and a von Kármán vortex street occurs behind the cylinder.

Although the unstabilized Schur complement iteration converges hardly at all for this test problem, again the augmented Lagrangian stabilization significantly improves the convergence behavior and only 3 – 4 inner GMRES iterations are required to reach the stopping criterion for $\gamma = 10^3$, no matter how many time steps are blocked or how fine the mesh is resolved (cf. Table 8).

However, the time-simultaneous multigrid algorithm for the velocity subproblems is not able to preserve the performance as observed for the Carreau-Yasuda viscosity model. The

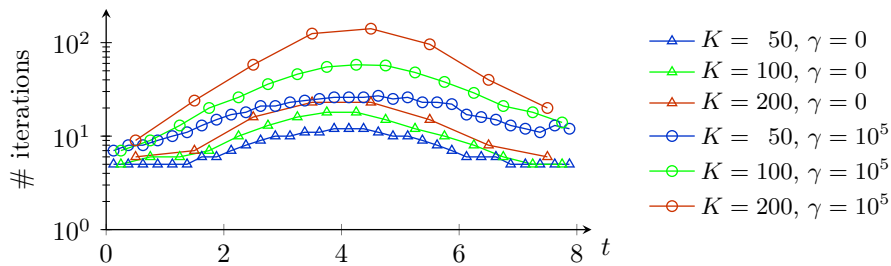


Figure 5: Velocity problem for Bench3 configuration: Total number of specialized multigrid iterations using 16 pre- and post-smoothing steps for solving velocity problem on mesh level 3 and time intervals $[K\delta t(i-1), K\delta ti]$, $i = 1, \dots, 1600K^{-1}$.

line plot in Fig. 5 shows the total number of multigrid iterations for the global-in-time velocity problem as occurring in different temporal subintervals. Especially in the middle of the time interval, where the magnitude of the velocity field attains its maximum and, hence, the convectivity is most dominant, the rate of convergence deteriorates significantly as more time steps are blocked. This behavior is even more pronounced if AL stabilization is included in the linear system of equations. Therefore, the so far considered time-simultaneous multigrid solver is not well suited for convection-dominated problems and calls for other global-in-time solution strategies. One candidate possibly performing better might be a space-time velocity solver as proposed by Gander and Neumüller [GN16], which exploits time coarsening for discontinuous Galerkin discretizations in time. Unfortunately, this approach has not yet been studied in combination with the augmented Lagrangian approach and, hence, is part of future investigations.

9 Conclusions

The augmented Lagrangian methodology is a well known approach to improve the convergence behavior of Schur complement solvers for the Oseen equations. It has already proved itself in several numerical test cases for stationary or sequential solution strategies. To the best knowledge of the authors, the present work provides the first attempts to extend this approach to the global-in-time solution of the Oseen equations. In combination with already existing pressure Schur complement (PSC) preconditioners for the unstabilized system and embedded into a GMRES method, the resulting solver guarantees a rapid convergence no matter how fine the spatial domain is resolved or how many time steps are considered simultaneously. The numerical results presented in this work illustrate the performance in several numerical test cases and suggest that a speedup can indeed be achieved compared to sequential solution strategies.

Unfortunately, the AL methodology comes at the cost that the velocity problems involved in each computation of the PSC residual become arbitrarily ill conditioned and, hence, require customized solution techniques. For this purpose, the specialized multigrid algorithm proposed in [SSW23] was successfully extended to solution of the global-in-time velocity

problem. If enough smoothing steps are performed, the rate of convergence seems to be independent of the amount of AL stabilization and the number of blocked time steps, at least if the viscosity parameter is sufficiently large. However, for high Reynolds numbers, the time-simultaneous solution procedure deteriorates for longer time intervals. Therefore, future work will focus on more sophisticated solution strategies for the velocity problem in the regime of dominant convective contributions. Furthermore, nonlinear solvers for the incompressible Navier-Stokes equations have to be investigated. As considered above, one promising candidate is (a global-in-time) Newton’s method, which requires some globalization technique to guarantee convergence for arbitrary initial guesses and large time intervals.

Acknowledgments

The authors gratefully acknowledge the computing time provided on the Linux HPC cluster at TU Dortmund University (LiDO3) and the Fritz HPC cluster at FAU Erlangen-Nürnberg, partially funded in the course of the Large-Scale Equipment Initiative by the German Research Foundation (DFG) as project 271512359 and 440719683, respectively. This work was supported by the Federal Ministry of Education and Research (BMBF) through the project “StroemungsRaum” 16ME0706K, which is part of the initiative “Neue Methoden und Technologien für das Exascale-Höchstleistungsrechnen” (SCALEXA).

References

- [BF91] F. Brezzi and M. Fortin. “Incompressible Materials and Flow Problems”. In: *Mixed and Hybrid Finite Element Methods*. Ed. by F. Brezzi and M. Fortin. New York, NY: Springer New York, 1991, pp. 200–273.
- [BO06] M. Benzi and M. A. Olshanskii. “An Augmented Lagrangian-Based Approach to the Oseen Problem”. In: *SIAM Journal on Scientific Computing* 28.6 (2006), pp. 2095–2113.
- [Car72] P. J. Carreau. “Rheological Equations from Molecular Network Theories”. In: *Transactions of the Society of Rheology* 16.1 (1972), pp. 99–127.
- [Chr+20] J. Christopher, R. D. Falgout, J. B. Schroder, S. M. Guzik, and X. Gao. “A space-time parallel algorithm with adaptive mesh refinement for computational fluid dynamics”. In: *Computing and Visualization in Science* 23 (1 2020).
- [Chr19] J. Christopher. *Space-time adaptivity with multigrid reduction in time for the compressible Navier-Stokes equations*. Jan. 2019.
- [CK09] E. Celledoni and T. Kvamsdal. “Parallelization in time for thermo-viscoplastic problems in extrusion of aluminium”. In: *International Journal for Numerical Methods in Engineering* 79.5 (2009), pp. 576–598.

- [CRK14] R. Croce, D. Ruprecht, and R. Krause. “Parallel-in-Space-and-Time Simulation of the Three-Dimensional, Unsteady Navier-Stokes Equations for Incompressible Flow”. In: *Modeling, Simulation and Optimization of Complex Processes - HPSC 2012*. Ed. by H. G. Bock, X. P. Hoang, R. Rannacher, and J. P. Schlöder. Cham: Springer International Publishing, 2014, pp. 13–23.
- [DSW22] F. Danieli, B. S. Southworth, and A. J. Wathen. “Space-Time Block Preconditioning for Incompressible Flow”. In: *SIAM Journal on Scientific Computing* 44.1 (2022), A337–A363.
- [Dün+21] J. Dünnebacke, S. Turek, C. Lohmann, A. Sokolov, and P. Zajac. “Increased space-parallelism via time-simultaneous Newton-multigrid methods for nonstationary nonlinear PDE problems”. In: *The International Journal of High Performance Computing Applications* (2021).
- [Fal+14] R. D. Falgout, A. Katz, T. V. Kolev, J. B. Schroder, A. M. Wissink, and U. M. Yang. *Parallel Time Integration with Multigrid Reduction for a Compressible Fluid Dynamics Application*. Lawrence Livermore National Laboratory, 2014.
- [FHM05] P. F. Fischer, F. Hecht, and Y. Maday. “A Parareal in Time Semi-implicit Approximation of the Navier-Stokes Equations”. In: *Domain Decomposition Methods in Science and Engineering*. Ed. by T. J. Barth, M. Griebel, D. E. Keyes, R. M. Nieminen, D. Roose, T. Schlick, R. Kornhuber, R. Hoppe, J. Périaux, O. Pironneau, O. Widlund, and J. Xu. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 433–440.
- [GN16] M. J. Gander and M. Neumüller. “Analysis of a New Space-Time Parallel Multigrid Algorithm for Parabolic Problems”. In: *SIAM Journal on Scientific Computing* 38.4 (2016), A2173–A2208.
- [Hag89] W. W. Hager. “Updating the Inverse of a Matrix”. In: *SIAM Review* 31.2 (1989), pp. 221–239.
- [Hes69] M. R. Hestenes. “Multiplier and gradient methods”. In: *Journal of Optimization Theory and Applications* 4 (5 1969), pp. 303–320.
- [KBG18] G. L. Kooij, M. A. Botchev, and B. J. Geurts. “An Exponential Time Integrator for the Incompressible Navier–Stokes Equation”. In: *SIAM Journal on Scientific Computing* 40.3 (2018), B684–B705.
- [KLW02] D. Kay, D. Loghin, and A. Wathen. “A Preconditioner for the Steady-State Navier–Stokes Equations”. In: *SIAM Journal on Scientific Computing* 24.1 (2002), pp. 237–256.
- [LMT01] J.-L. Lions, Y. Maday, and G. Turinici. “Résolution d’EDP par un schéma en temps «pararéel »”. In: *Comptes Rendus de l’Académie des Sciences - Series I - Mathematics* 332.7 (2001), pp. 661–668.
- [LO87] C. Lubich and A. Ostermann. “Multi-grid dynamic iteration for parabolic equations”. In: *BIT Numerical Mathematics* 27 (2 1987), pp. 216–234.
- [LT23] C. Lohmann and S. Turek. “On the Design of Global-in-Time Newton-Multigrid-Pressure Schur Complement Solvers for Incompressible Flow Problems”. In: *Journal of Mathematical Fluid Mechanics* 25 (3 2023), p. 64.

- [MJY19] Z. Miao, Y.-L. Jiang, and Y.-B. Yang. “Convergence analysis of a parareal-in-time algorithm for the incompressible non-isothermal flows”. In: *International Journal of Computer Mathematics* 96.7 (2019), pp. 1398–1415.
- [Pow69] M. J. D. Powell. “A method for nonlinear constraints in minimization problems”. In: *Optimization* (1969), pp. 283–298.
- [QV17] A. Quarteroni and A. Valli. *Numerical Models for Differential Problems*. Berlin, Heidelberg: Springer Berlin Heidelberg, Jan. 2017.
- [Sch+96] M. Schäfer, S. Turek, F. Durst, E. Krause, and R. Rannacher. “Benchmark computations of laminar flow around a cylinder”. In: *Flow simulation with high-performance computers II*. Springer, 1996, pp. 547–566.
- [Sch99] J. Schöberl. “Multigrid methods for a parameter dependent problem in primal variables”. In: *Numerische Mathematik* 84 (1 1999), pp. 97–119.
- [SSW23] Y.-h. Shih, G. Stadler, and F. Wechsung. “Robust Multigrid Techniques for Augmented Lagrangian Preconditioning of Incompressible Stokes Equations with Extreme Viscosity Variations”. In: *SIAM Journal on Scientific Computing* 45.3 (2023), S27–S53.
- [Ste+15] J. Steiner, D. Ruprecht, R. Speck, and R. Krause. “Convergence of Parareal for the Navier-Stokes Equations Depending on the Reynolds Number”. In: *Numerical Mathematics and Advanced Applications - ENUMATH 2013*. Ed. by A. Abdulle, S. Deparis, D. Kressner, F. Nobile, and M. Picasso. Cham: Springer International Publishing, 2015, pp. 195–202.
- [TP04] J. M. F. Trindade and J. C. F. Pereira. “Parallel-in-time simulation of the unsteady Navier-Stokes equations for incompressible flow”. In: *International Journal for Numerical Methods in Fluids* 45.10 (2004), pp. 1123–1136.
- [TP06] J. M. F. Trindade and J. C. F. Pereira. “Parallel-in-Time Simulation of Two-Dimensional, Unsteady, Incompressible Laminar Flows”. In: *Numerical Heat Transfer, Part B: Fundamentals* 50.1 (2006), pp. 25–40.
- [Tur99] S. Turek. *Efficient Solvers for Incompressible Flow Problems: An Algorithmic and Computational Approach*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999.
- [Wec19] F. Wechsung. “Shape optimisation and robust solvers for incompressible flow”. PhD thesis. University of Oxford, 2019.
- [YAC81] K. Yasuda, R. C. Armstrong, and R. E. Cohen. “Shear flow properties of concentrated solutions of linear and star branched polystyrenes”. In: *Rheologica Acta* 20 (2 1981), pp. 163–178.