

Technical Report

TR-2010-011

Analysis of Augmented Lagrangian-Based Preconditioners for the Steady Incompressible Navier-Stokes Equations

by

Michele Benzi, Zhen Wang

MATHEMATICS AND COMPUTER SCIENCE

EMORY UNIVERSITY

Augmented Lagrangian Preconditioners for the Incompressible Navier-Stokes Equations

Michele Benzi^{1,*}, Maxim A. Olshanskii² and Zhen Wang³

¹*Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA*

²*Department of Mechanics and Mathematics, Moscow State M. V. Lomonosov University, Moscow 119899, Russia*

³*Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA*

SUMMARY

We study different variants of the augmented Lagrangian-based block triangular preconditioner introduced by the first two authors in [SIAM J. Sci. Comput., 28 (2006), pp. 2095–2113]. The preconditioners are used to accelerate the convergence of the Generalized Minimal Residual (GMRES) method applied to various finite element and MAC discretizations of the Oseen problem in two and three space dimensions. Both steady and unsteady problems are considered. Numerical experiments show the effectiveness of the proposed preconditioners for a wide range of problem parameters. Implementation on parallel architectures is also considered. The augmented Lagrangian-based approach is further generalized to deal with linear systems from stabilized finite element discretizations. Copyright © 2000 John Wiley & Sons, Ltd.

KEY WORDS: preconditioning; saddle point problems; Oseen problem; augmented Lagrangian method; Krylov subspace methods; parallel computing

1. Introduction

We consider the solution of the incompressible Navier–Stokes equations governing the flow of viscous Newtonian fluids. For an open bounded domain $\Omega \subset \mathbb{R}^d$ ($d = 2, 3$) with boundary $\partial\Omega$, time interval $[0, T]$, and data \mathbf{f} , \mathbf{g} and \mathbf{u}_0 , the goal is to find a velocity field $\mathbf{u} = \mathbf{u}(\mathbf{x}, t)$ and

*Correspondence to: M. Benzi, Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA. Email: benzi@mathcs.emory.edu

Contract/grant sponsor: U. S. National Science Foundation; contract/grant number: DMS-0511336

Contract/grant sponsor: Russian Foundation Basic Research; contract/grant number: 08-01-00415, 09-01-00115

pressure field $p = p(\mathbf{x}, t)$ such that

$$\frac{\partial \mathbf{u}}{\partial t} - \nu \Delta \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u} + \nabla p = \mathbf{f} \quad \text{on } \Omega \times (0, T] \quad (1)$$

$$\operatorname{div} \mathbf{u} = 0 \quad \text{on } \Omega \times [0, T] \quad (2)$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \partial\Omega \times [0, T] \quad (3)$$

$$\mathbf{u}(\mathbf{x}, 0) = \mathbf{u}_0(\mathbf{x}) \quad \text{on } \Omega \quad (4)$$

where ν is the kinematic viscosity, Δ is the Laplacian, ∇ is the gradient and div the divergence. Implicit time discretization and linearization of the Navier–Stokes system (1)–(4) by Picard fixed-point iteration result in a sequence of (generalized) Oseen problems of the form

$$\sigma \mathbf{u} - \nu \Delta \mathbf{u} + (\mathbf{v} \cdot \nabla) \mathbf{u} + \nabla p = \mathbf{f} \quad \text{in } \Omega \quad (5)$$

$$\operatorname{div} \mathbf{u} = 0 \quad \text{in } \Omega \quad (6)$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \partial\Omega \quad (7)$$

where \mathbf{v} is a known velocity field from a previous iteration or time step (the ‘wind’) and σ is proportional to the reciprocal of the time step ($\sigma = 0$ for a steady problem). When $\mathbf{v} = \mathbf{0}$ we have a (generalized) Stokes problem.

Spatial discretization of (5)–(7) using finite differences or finite elements results in large, sparse saddle point systems of the form

$$\begin{pmatrix} A & B^T \\ B & -C \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} f \\ g \end{pmatrix}, \quad (8)$$

where u and p represent the discrete velocity and pressure, respectively, A is the discretization of the diffusion, convection, and time-dependent terms, B^T is the discrete gradient, B the (negative) discrete divergence, C is a stabilization matrix, and f and g contain forcing and boundary terms. If the discretization satisfies the LBB (‘inf-sup’) stability condition, no pressure stabilization is required and we can take $C = 0$. If the LBB condition is not satisfied, the stabilization matrix $C \neq 0$ is symmetric and positive semidefinite and the actual choice of C depends on the particular finite element pair being used; see, e.g., [1].

The efficient solution of systems of the form (8) necessitates rapidly convergent iterative methods. The two main approaches available are preconditioned Krylov subspace methods [2] and multigrid methods [3, 4]. The two approaches can be combined by using one or more multigrid cycles as preconditioners for Krylov methods. In this paper, we focus on block preconditioners for GMRES [5]; for related work on multigrid solvers, see [6].

In recent years, a considerable amount of work has been done in developing efficient preconditioners for Krylov subspace methods applied to incompressible flow problems; see the comprehensive treatments in [1] and [7]. More recent contributions include the papers [8] and [9] on preconditioning based on the augmented Lagrangian (AL) approach [10], and [11] where the least squares commutator preconditioner (see [1, Section 8.2.2]) is generalized to stabilized finite element discretizations of the Oseen problem. Other relevant work includes the development of ILU-type preconditioners for saddle point problems [12] and SIMPLE-type block preconditioners [13]. In these papers, these preconditioners are analyzed and compared with other preconditioners; see also [14], where the AL-based preconditioner of [8] was found to compare favorably with other approaches.

In this paper we further develop the AL-based block triangular preconditioner in various directions. It was shown in [8] that the AL-based approach results in preconditioners that are independent of the mesh size h and fairly insensitive to the viscosity ν , resulting in a remarkably robust and nearly optimal solver for the steady Oseen problem up to Reynolds numbers of about 10,000. It was further shown in [9] that the preconditioner performs quite well also for challenging linear systems arising from the linear stability analysis of linearized (Newton) solutions to the incompressible Navier-Stokes equations. As shown in these papers, the crucial ingredient for the AL-based preconditioner was an efficient multigrid cycle used as an approximate solver for the velocity subproblem associated with the (1,1) block of the preconditioner. For stable discretizations of the steady Oseen problem, excellent results were obtained in [8] with a geometric multigrid scheme based on the method in [15].

Here we introduce a variant of the AL-based block triangular preconditioner, which we refer to as the *modified AL preconditioner*, which can be more readily implemented for general discretizations and geometries using off-the-shelf algebraic solvers for scalar elliptic problems. In particular, state-of-the-art parallel algebraic multilevel solvers can be used to solve the subsystems arising in the application of the preconditioner. Moreover, we extend the AL preconditioner to saddle point systems arising from stabilized finite element methods ($C \neq 0$ in (8)). This requires a different approach to augment the linear system before constructing and applying the preconditioner. In addition, we consider three-dimensional (3-D) Oseen problems discretized by the Marker-and-Cell (MAC) [16] finite difference method. The staggered grid in MAC makes it a stable discretization method, so the original and modified AL preconditioners can be applied directly. Our numerical experiments show that in 3-D one can expect similar results to those obtained in 2-D using stable finite elements.

This paper is organized as follows. In Section 2, the modified AL preconditioner is described and applied to solve linear systems arising from 2-D stable finite element and 3-D stable MAC discretizations. We consider both steady and unsteady problems. AL-based preconditioners for 2-D stabilized (Q1-Q1) finite elements are developed in Section 3. Numerical experiments (including a few on a parallel computer) are used throughout to illustrate the behavior of the different methods. Some brief conclusive remarks are given in Section 4.

While some theoretical analysis is given to motivate the proposed preconditioners and to guide in the choice of parameters, the paper is mostly about practical computational issues and numerical results.

2. Modified AL-based preconditioner for stable finite elements

In this Section we present the modified AL-based block triangular preconditioner for the Oseen problem discretized by stable finite element pairs, such as Q2-Q1 or Q2-P1; see, e.g., [1]. We begin by recalling the original AL-based preconditioner described in [8]. We discuss first the 2-D case; the extension to the 3-D case is given in section 2.8.

2.1. Problem formulation

Here we consider solving the steady-state Oseen equations ($\sigma = 0$ in (5)). An LBB-stable finite element discretization gives rise to the following system:

$$\begin{pmatrix} A & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} f \\ g \end{pmatrix}, \quad \text{or} \quad \mathcal{A}x = b. \quad (9)$$

The equivalent augmented Lagrangian (AL) formulation [10] is given by

$$\begin{pmatrix} A + \gamma B^T W^{-1} B & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} \hat{f} \\ g \end{pmatrix}, \quad \text{or} \quad \hat{\mathcal{A}}x = \hat{b}, \quad (10)$$

where $\hat{f} := f + \gamma B^T W^{-1} g$, W is symmetric positive definite, and $\gamma > 0$. A good choice of W is the pressure mass matrix, M_p ; in practice, we use the main diagonal of M_p instead, in order to maintain sparsity in $A + \gamma B^T W^{-1} B$. The choice of γ is important and will be discussed below.

The use of the AL formulation (10) instead of the original one (9) can be justified in various ways; see for instance the discussion in [8, 9]. Here we justify this choice by the observation that preconditioning (10) allows us to circumvent the delicate issue of finding good approximations for the pressure Schur complement $BA^{-1}B^T$ or its inverse, which is crucial when constructing preconditioners for the non-augmented system (9).

2.2. The original AL-based preconditioner

Letting $A_\gamma := A + \gamma B^T W^{-1} B$, an ideal preconditioner for problem (10) is given by the block triangular matrix

$$\mathcal{P} = \begin{pmatrix} A_\gamma & B^T \\ 0 & -\frac{1}{\gamma}W \end{pmatrix}. \quad (11)$$

Because of the identity

$$\mathcal{P}^{-1} = \begin{pmatrix} A_\gamma^{-1} & \gamma A_\gamma^{-1} B^T W^{-1} \\ 0 & -\gamma W^{-1} \end{pmatrix} = \begin{pmatrix} A_\gamma^{-1} & 0 \\ 0 & I_m \end{pmatrix} \begin{pmatrix} I_n & B^T \\ 0 & -I_m \end{pmatrix} \begin{pmatrix} I_n & 0 \\ 0 & \gamma W^{-1} \end{pmatrix},$$

applying \mathcal{P}^{-1} to a vector requires one solve with W and one with A_γ . The right-preconditioned matrix is

$$\hat{\mathcal{A}}\mathcal{P}^{-1} = \begin{pmatrix} I_n & 0 \\ BA_\gamma^{-1} & \gamma BA_\gamma^{-1} B^T W^{-1} \end{pmatrix}, \quad (12)$$

showing that $\lambda = 1$ is an eigenvalue of multiplicity at least n . Additionally, there are m eigenvalues λ that satisfy the generalized eigenvalue problem

$$BA_\gamma^{-1} B^T p = \lambda \left(\frac{1}{\gamma} W \right) p.$$

Lemma 4.1 in [8] states that if all the relevant matrices are invertible,

$$(BA_\gamma^{-1} B^T)^{-1} = (B(A + \gamma B^T W^{-1} B)^{-1} B^T)^{-1} = (BA^{-1} B^T)^{-1} + \gamma W^{-1}. \quad (13)$$

Setting $W = M_p$ (the pressure mass matrix), we thus have

$$\begin{aligned} \lambda^{-1} p &= \frac{1}{\gamma} (BA_\gamma^{-1} B^T)^{-1} M_p p \\ &= \frac{1}{\gamma} \left((BA^{-1} B^T)^{-1} + \gamma M_p^{-1} \right) M_p p \\ &= \frac{1}{\gamma} (BA^{-1} B^T)^{-1} M_p p + p \\ &= \frac{1}{\gamma \mu} p + p, \end{aligned}$$

where μ satisfies the generalized eigenproblem

$$(BA^{-1}B^T)q = \mu M_p q.$$

Hence, we obtain

$$\lambda = \frac{\gamma\mu}{1 + \gamma\mu}. \quad (14)$$

Next, we derive bounds for the real and imaginary parts of λ . Writing $\lambda = a_\lambda + i b_\lambda$ and $\mu = a_\mu + i b_\mu$, where $i = \sqrt{-1}$, we have

$$a_\lambda + i b_\lambda = \frac{\gamma(a_\mu + i b_\mu)}{1 + \gamma(a_\mu + i b_\mu)} = \frac{\gamma a_\mu(1 + \gamma a_\mu) + \gamma^2 b_\mu^2 + i \gamma b_\mu}{(1 + \gamma a_\mu)^2 + \gamma^2 b_\mu^2},$$

which implies

$$a_\lambda = \frac{\gamma a_\mu(1 + \gamma a_\mu) + \gamma^2 b_\mu^2}{(1 + \gamma a_\mu)^2 + \gamma^2 b_\mu^2} \quad \text{and} \quad b_\lambda = \frac{\gamma b_\mu}{(1 + \gamma a_\mu)^2 + \gamma^2 b_\mu^2}.$$

Following the argument in [17], it can be shown that the μ 's are contained in a rectangle which lies in the positive half plane $\text{Re}(z) > 0$ and which does not depend on h . Easy manipulations give

$$0 < \min_{\mu} \frac{\gamma a_\mu}{1 + \gamma a_\mu} \leq a_\lambda \leq 1, \quad (15)$$

$$|b_\lambda| \leq \max_{\mu} \min \left\{ \gamma |b_\mu|, \frac{1}{\gamma |b_\mu|} \right\} \leq 1. \quad (16)$$

Note that the real part of λ is bounded away from zero uniformly in h , for all fixed $\gamma > 0$. Likewise, the imaginary part of λ is bounded uniformly in h , for all fixed $\gamma > 0$, and is bounded by 1 for $\gamma = 1$. Furthermore, for $\gamma \rightarrow \infty$ all eigenvalues tend to 1; see [8, Theorem 4.2].

In Table I we show the maximum and minimum of the real part and the maximum imaginary part of μ and λ corresponding to a uniform Q2-Q1 discretization of the lid driven cavity problem described in [1, page 316]. The value $\gamma = 1$ is used in the (ideal) AL preconditioner. Notice that 0 is also an eigenvalue (corresponding to the hydrostatics pressure mode), which does not affect the convergence of preconditioned GMRES and can be excluded from further consideration (see [1, page 83]). One can clearly see the independence of λ with respect to h , and the near-independence with respect to ν . The real and imaginary parts of μ show a weak dependence on h , which according to the theory must disappear in the limit of $h \rightarrow 0$; note that the minimum of the real part of μ is already h -independent even for these rather coarse grids. On the other hand, there is a strong dependence of μ on the viscosity ν , as already observed in [17].

Because of the expensive solve associated with the velocity submatrix A_γ , the 'ideal' preconditioner \mathcal{P} in (11) is not practical. It is necessary to replace the exact solves with inexact ones, leading to preconditioners of the form

$$\mathcal{P} = \begin{pmatrix} \hat{A}_\gamma & B^T \\ 0 & \hat{S} \end{pmatrix}, \quad (17)$$

where $\hat{A}_\gamma \approx A + \gamma B^T W^{-1} B$ and \hat{S} are implicitly defined via the action of their inverses on vectors. In [8], \hat{A}_γ^{-1} was implemented as a W-cycle of a non-standard geometric multigrid

Table I. Eigenvalue bounds with the ideal AL preconditioner.

Grid	$\max a_\mu$	$\min a_\mu$	$\max b_\mu $	$\max a_\lambda$	$\min a_\lambda$	$\max b_\lambda $
ν	0.1					
16×16	15.677	1.259	2.274	0.9411	0.5573	0.0127
32×32	19.355	1.277	4.323	0.9519	0.5608	0.0121
64×64	21.147	1.278	4.973	0.9553	0.5610	0.0185
ν	0.01					
16×16	132.77	9.16	38.22	0.9925	0.9016	0.0275
32×32	159.89	11.57	64.60	0.9938	0.9204	0.0292
64×64	192.85	12.65	88.61	0.9948	0.9267	0.0255
ν	0.001					
16×16	1279.6	2.3	148.9	0.9992	0.6961	0.0586
32×32	1477.7	2.2	301.3	0.9993	0.6914	0.0529
64×64	1584.5	2.3	452.2	0.9994	0.6968	0.0270

method based on [15]; for \widehat{S}^{-1} , a few Richardson iterations preconditioned with the diagonal \widehat{M}_p of M_p were used to approximately solve linear systems associated with M_p . (In practice, another acceptable choice is to simply use $\widehat{S}^{-1} = -\gamma \widehat{M}_p^{-1}$.) Theory and numerical experiments in [8] show that the preconditioner (17) is nearly optimal, meaning that iterative methods with this AL preconditioner is independent of the grid and almost insensitive to viscosity.

In the next subsection we consider some variants of this approach that do not require sophisticated geometric multigrid techniques in the implementation of \widehat{A}_γ^{-1} .

2.3. The modified AL preconditioner

As already observed, the main cost associated with the application of the block triangular preconditioner \mathcal{P} is associated with the (approximate) solution of linear systems with matrix $A_\gamma = A + \gamma B^T W^{-1} B$. Hence, how to simplify the action of A_γ^{-1} becomes the critical issue.

Using the fact that A is block diagonal: $A = \text{diag}(A_1, A_2)$ and writing $B = (B_1, B_2)$ we have

$$\begin{aligned}
A_\gamma &= A + \gamma B^T W^{-1} B \\
&= \begin{pmatrix} A_1 & 0 \\ 0 & A_2 \end{pmatrix} + \gamma \begin{pmatrix} B_1^T \\ B_2^T \end{pmatrix} W^{-1} \begin{pmatrix} B_1 & B_2 \end{pmatrix} \\
&= \begin{pmatrix} A_1 + \gamma B_1^T W^{-1} B_1 & \gamma B_1^T W^{-1} B_2 \\ \gamma B_2^T W^{-1} B_1 & A_2 + \gamma B_2^T W^{-1} B_2 \end{pmatrix} \\
&=: \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}.
\end{aligned}$$

Consider the following block triangular approximation to A_γ :

$$\tilde{A}_\gamma = \begin{pmatrix} A_1 + \gamma B_1^T W^{-1} B_1 & \gamma B_1^T W^{-1} B_2 \\ 0 & A_2 + \gamma B_2^T W^{-1} B_2 \end{pmatrix} = \begin{pmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{pmatrix},$$

in which the (2,1) block A_{21} is dropped. We refer to the block triangular matrix

$$\tilde{\mathcal{P}} = \begin{pmatrix} \tilde{A}_\gamma & B^T \\ 0 & \hat{S} \end{pmatrix} = \begin{pmatrix} A_{11} & A_{12} & B_1^T \\ 0 & A_{22} & B_2^T \\ 0 & 0 & \hat{S} \end{pmatrix} \quad (18)$$

as to the *modified AL preconditioner*. Because of the block triangular structure of $\tilde{\mathcal{P}}$, most of the work in computing the action of $\tilde{\mathcal{P}}^{-1}$ on a vector lies in the solution of the two linear systems with coefficient matrices A_{11} and A_{22} . Observing that $A_{ii} = A_i + \gamma B_i^T W^{-1} B_i$ ($i = 1, 2$), in which the A_i 's represent discrete scalar convection-diffusion operators and the B_i 's are discretizations of the partial derivatives with respect to x and y , we immediately see that the A_{ii} 's can be interpreted as discrete scalar anisotropic convection-diffusion operators with anisotropy ratio $\approx 1 + \frac{\gamma}{\nu}$. Thus, applying A_{11}^{-1} and A_{22}^{-1} requires solving two scalar anisotropic convection-diffusion problems (three in 3-D). These subsystems can be solved exactly or inexactly. For 2-D problems, sparse direct solvers can be used to solve the subsystems efficiently. For 3-D problems, however, direct solvers become prohibitively expensive for sufficiently fine meshes, and some inner iterative method should be used instead. We will return on this issue later on.

The fact that the modified AL preconditioner is obtained from the ideal one by dropping a block of the form $\gamma B_2^T W^{-1} B_1$ suggests that the parameter γ should not be taken too large. This is also desirable in view the fact that when ν is small, a small value of γ ensures that the anisotropy ratio $1 + \frac{\gamma}{\nu}$ in the diagonal blocks A_{ii} does not grow too large; this can be important if inner iterations are used to solve the two linear systems associated with A_{11} and A_{22} , since strong anisotropies may cause difficulties for iterative solvers. This can be verified by numerical experiments; see Section 2.4.

2.4. Numerical experiments: steady problems

In this section we compare the original ('ideal') and modified AL preconditioners on a standard 2-D test problem, in order to assess the effect of dropping the A_{21} block from A_γ in the preconditioner and to get some insight on the choice of γ . Based on numerical evidence, we state an empirical rule to choose γ for the modified preconditioner and then apply this rule to other types of problems with more complicated geometries, showing its effectiveness. Finally, we present results of some tests on a generalized (unsteady) Oseen problem, showing that the modified AL preconditioners work well on this kind of problems. We use IFISS [18, 19] to generate the Oseen problems in this section and the next on 2-D stabilized finite elements. These computations are performed on a Sun Microsystems SunFire V40z with 4 Dual Core AMD Opteron Processors and 32 GB of memory.

The basic Krylov solver used in all our experiments is full GMRES [2], starting from a zero initial guess. The iteration is stopped when the relative residual norm is reduced below 10^{-6} .

First, the original and modified AL preconditioners are compared by using the lid driven cavity test problem in on different grids and for different values of the viscosity ν . The linearized Oseen system is the one in the first step of the Picard iteration immediately following the initial one (which reduces to a Stokes problem) and is discretized by uniform Q2-Q1 finite elements. We take W equal to the diagonal of the pressure mass matrix. For the original AL preconditioner, we use $\gamma = 1$ for all the experiments; little is gained in tuning γ . For the modified preconditioner, in contrast, the choice of γ has a considerable influence on the

Table II. Number of GMRES iterations with the ideal and modified AL preconditioners for steady 2D Oseen problems (Q2-Q1 FEM). The optimal γ is in parentheses.

Viscosity	0.1		0.01		0.001	
Grid	Original	Modified	Original	Modified	Original	Modified
16×16	9	14 (0.5)	7	18 (0.08)	8	32 (0.04)
32×32	9	16 (0.4)	7	21 (0.06)	8	46 (0.03)
64×64	10	18 (0.3)	6	23 (0.04)	8	53 (0.02)
128×128	10	19 (0.3)	7	25 (0.03)	7	65 (0.02)

performance and cannot be made in a cavalier way. In our first set of experiments, we find the value of γ that minimizes the number of GMRES iterations by performing experiments with different values of γ . For $\nu = 0.1$ we take values of γ in the interval $[0.1, 1]$ with 0.1 increments. For $\nu = 0.01$ and $\nu = 0.001$ we take $\gamma \in [0.01, 0.1]$ with increments of 0.01; experiments show that the optimal γ lies in these intervals.

When applying the preconditioners we need to solve linear systems with A_γ (for the ideal preconditioner) and with A_{11} and A_{22} (for the modified preconditioner). Here we performed exact solves by first applying a column approximate AMD permutation [20] (using the MATLAB function `colamd`) to A_γ (respectively, to A_{11} and A_{22}) followed by sparse LU factorization in MATLAB.

Results are shown in Table II. From these results we can see that the performance of the ideal AL preconditioner is independent on both the mesh size h and the viscosity ν . The eigenvalues of the preconditioned matrices corresponding to $\nu = 0.01$ and $\nu = 0.001$ on the 32×32 grid are plotted on the top two figures in Fig. 1. The non-zero eigenvalues are bounded away from 0, and cluster around 1 (note the different scales used for the horizontal and vertical axes). The only 0 eigenvalue comes from the hydrostatic pressure mode, which makes the saddle point system singular. Note that for ν smaller, the imaginary parts of the eigenvalues are slightly larger, as already shown in Table I. The strong clustering of the eigenvalues is in agreement with the fast convergence of the preconditioned GMRES iteration.

The convergence rate for the modified AL preconditioner depends on the mesh size and viscosity, and some deterioration is observed when either decreases. The deterioration with respect to h is mild when the viscosity is not too small, but it becomes more noticeable as ν becomes smaller. There is a clear deterioration with respect to decreasing viscosity for all h . Nonetheless, even the ‘worst’ result (65 iterations for the 128×128 mesh with $\nu = 0.001$) is still quite acceptable, and competitive with the results reported in [1, page 357] using some of the best available preconditioners, with a similar cost per iteration.

The eigenvalues of the preconditioned matrices using the modified AL preconditioner tend to scatter as ν gets small, as can be seen from the bottom two figures in Fig. 1.

The advantage of the modified AL preconditioner over the ideal one is that the cost of forming and applying the preconditioner is greatly reduced, both in terms of computing time and in terms of storage. This reduction in costs more than compensates the increase in the number of iterations. Indeed, the LU factorization of the highly coupled block matrix A_γ requires far more time and storage than the factorization of its diagonal blocks A_{11} and A_{22} . For the 128×128 grid, factoring A_γ dominates the computation, accounting for about 75% of




Figure 1. Plots of eigenvalues of matrix with the AL original preconditioner (top) and the modified one (bottom) on 32×32 grid, as $\nu = 0.1$ (left) and 0.001 (right) (Q2-Q1).

the total solution time. In contrast, the time taken by the iterative phase dominates the overall cost when using the modified AL preconditioner; storage needs are substantially reduced, since the LU factors of A_{11} and A_{22} contain much less fill-in overall than the LU factors of A_γ . Hence, each iteration is much less expensive than with the ideal AL preconditioner.

Some timings for the 128×128 grid are reported in Table III. Here, “Iter” denotes the number of iterations required to reduce the relative residual norm below 10^{-6} , and “Time” includes two parts: the setup time for constructing and factoring the two preconditioners, and the computing time for solving the linear system by preconditioned GMRES. For the modified AL preconditioner, we exploit the fact that the diagonal blocks A_{11} or A_{22} are independent and can be factored in parallel. We point out that the Matlab codes used for these experiments are not highly optimized, so the timings are shown here just to give an indication of the relative performance of the ideal and modified preconditioners.

2.5. Choosing the parameter γ

We now turn to the important issue of how to choose γ when using the modified AL preconditioner. The ‘best’ values of γ (in terms of iteration counts) in the previous set of

Table III. Timings for the original and modified AL preconditioners on 128×128 grid (Q2-Q1).

Viscosity	Original		Modified	
	Iter	Time	Iter	Time
0.1	10	180.6/78.8	19	15.1/52.3
0.01	7	191.5/61.2	25	15.6/68.8
0.001	7	189.2/59.9	65	16.7/174.8

Table IV. Iteration counts for different γ , 2-D driven cavity (Q2-Q1 FEM).

Viscosity	0.1		0.01		0.001	
	$\gamma = 0.3$	Exp.	$\sqrt{2}$ rule	Exp.	$\sqrt{2}$ rule	Exp.
16×16	16	14 (0.5)	18 (0.08)	18 (0.08)	32 (0.04)	32 (0.04)
32×32	16	16 (0.3)	21 (0.0566)	21 (0.06)	47 (0.028)	46 (0.03)
64×64	18	18 (0.3)	23 (0.04)	23 (0.04)	53 (0.02)	53 (0.02)
128×128	19	19 (0.3)	25 (0.0283)	25 (0.03)	60 (0.0141)	60 (0.015)

Table V. Additional test results for the choice of γ (Q2-Q1 FEM).

Lid driven cavity (stretched grid)					Backward facing step (uniform grid)				
Grid	$\sqrt{2}$ rule		Exp.		Grid	$\sqrt{2}$ rule		Exp.	
	γ	Iter	γ	Iter		γ	Iter	γ	Iter
16×16	0.04	29	0.04	29	16×48	0.1	25	0.1	25
32×32	0.0283	37	0.03	37	32×96	0.0707	25	0.07	25
64×64	0.02	47	0.02	47	64×192	0.05	28	0.06	27
128×128	0.0141	56	0.015	56	128×384	0.0354	32	0.05	29

experiments are displayed in the 3rd, 5th and 7th columns of Table IV under “Exp.”, since they were determined experimentally. We make the following observations. First, when $\nu = 0.1$ the modified AL preconditioner is not sensitive to γ , and taking $\gamma = 0.3$ turns out to be a good choice. Second, for other values of the viscosity, the experimentally found optimal γ for the coarsest grid is divided by a factor of about $\sqrt{2}$ when the mesh size h is divided by 2. Hence, for fixed ν , letting γ_0 be the γ for the coarsest grid, we use $\gamma_0/\sqrt{2}$, $\gamma_0/2$ and $\gamma_0/2\sqrt{2}$ for each refined grid. We call this the “ $\sqrt{2}$ rule.” The 2nd column of Table IV displays the iteration count for the case $\nu = 0.1$ using $\gamma = 0.3$, whereas the 4th and 6th columns display the number of GMRES iterations for $\nu = 0.01$ and $\nu = 0.001$, respectively, using the value of γ given in parentheses. One can see that these γ ’s give results that are essentially identical to those obtained with the optimal choice of γ . Note that for the 128×128 grid and $\nu = 0.001$, $\gamma = 0.015$ is used in the last column to compare with $\gamma = 0.0141$.

Next, we test the “ $\sqrt{2}$ rule” on two other problems from [1], namely, the lid driven cavity discretized on a stretched grid and the flow over a backward facing step on a uniform grid.

Table VI. GMRES iterations with the original and modified AL preconditioners for unsteady driven cavity problem (Q2-Q1).

Viscosity	0.1		0.01		0.001	
Grid	Original	Modified	Original	Modified	Original	Modified
16×16	11	11	10	12	8	13
32×32	13	15	13	16	13	17
64×64	17	18	17	20	17	22
128×128	22	23	22	25	22	27

Here we only consider the cases $\nu = 0.001$ for the driven cavity and $\nu = 0.005$ for the backward facing step, since the performance of the modified AL preconditioner is less sensitive to the choice of γ for larger values of ν . For the backward facing step we do not consider smaller values of ν since, as pointed out in [1, page 316], the flow becomes unstable already for fairly moderate Reynolds numbers and computing a steady solution for smaller values of ν would not be meaningful. The results obtained with the choice of γ based on the “ $\sqrt{2}$ rule” are compared with the best ones found by numerical experiment in Table V. It is clear that the “ $\sqrt{2}$ rule” works remarkably well also in these cases.

2.6. Numerical experiments: unsteady problems

Here we present some results for unsteady (generalized) Oseen problems. In this type of problem the (1,1) block A of the saddle point matrix contains an additional term of the form σM_u where σ is the reciprocal of the time step Δt and M_u is the velocity mass matrix. Linear systems of this type tend to be easier to solve than the ones arising in the steady case, since the presence of the additional positive definite term σM_u makes the (1,1) block more diagonally dominant, especially when Δt is sufficiently small. On the other hand, many such systems (one for each time step) have to be solved during a simulation, so fast solvers are absolutely indispensable.

We have performed numerical experiments for both the driven cavity and the backward facing step problems. In our experiments we let $\sigma = h^{-1}$ where h is the mesh size. It turns out that the simple choice $\gamma = 1$ works well in all cases. The results are shown in Table VI. Clearly, the modified AL preconditioner is comparable to the original one in terms of iteration numbers (and therefore much faster in terms of computing times). This is not surprising, since the presence of the additional term σM_u makes the diagonal sub-blocks A_{11} and A_{22} much larger in norm than the off-diagonal ones, and consequently dropping one of the off-diagonal blocks has only a marginal effect on the quality of the preconditioner, while drastically reducing the associated costs. Notice that both preconditioners are essentially insensitive to the viscosity ν . There is, however, a dependence on h . To improve the convergence behavior as $h \rightarrow 0$, a different choice for the (2,2) block of the preconditioner is needed (currently, it is just a diagonal matrix). Similar to the well-known Cahouet-Chabard [21] preconditioner for the unsteady Stokes problem, we modify the (2,2) block of the preconditioner (implicitly defined by its inverse) by using

$$\widehat{S}^{-1} = -\nu \widehat{M}_p^{-1} - \sigma (B \widehat{M}_u^{-1} B^T)^{-1}$$

Table VII. GMRES iterations for unsteady driven cavity problem with improved \hat{S}^{-1} .

Viscosity	0.1		0.01		0.001	
Grid	Original	Modified	Original	Modified	Original	Modified
16×16	8	15	9	18	10	19
32×32	9	16	11	20	12	23
64×64	8	15	9	20	11	23
128×128	7	14	7	18	7	20

Table VIII. GMRES iterations for unsteady flow over a step problem with improved \hat{S}^{-1} .

Viscosity	0.1		0.01		0.001	
Grid	Original	Modified	Original	Modified	Original	Modified
16×48	11	18	15	25	16	27
32×96	9	18	15	30	17	35
64×192	9	18	17	34	20	44
128×394	9	17	18	36	22	53

instead of $\hat{S}^{-1} = -\gamma \widehat{M}_p^{-1}$. Here \widehat{M}_u denotes the diagonal of the velocity mass matrix; moreover, the action of the (pseudo-)inverse of the scaled discrete Laplacian $B \widehat{M}_u^{-1} B^T$ is implemented inexactly via a single iteration (“K-cycle”) of Yvan Notay’s algebraic multigrid method [22, 23].

Iteration counts are shown in Table VII for the lid driven cavity and Table VIII for the backward facing step problem. As before, the ‘Original’ preconditioner requires exact solves with A_γ , the modified one only with the diagonal sub-blocks A_{11} and A_{22} . Both preconditioners are now quite robust and stable with respect to h on sufficiently fine meshes. In the case of the backward facing step, we observe some dependence on ν . Also, for $\nu = 0.001$ there is a noticeable dependence on h with the modified preconditioner. This dependence persists even when the action of the improved inverse Schur complement \hat{S}^{-1} is computed ‘exactly’.

2.7. Parallel results

In this subsection we show the results of a few numerical experiments on a compute cluster. The test problem is the unsteady 2-D Oseen problem for the lid driven cavity discretized with Q2-Q1 finite elements, using $\sigma = h^{-1}$. The value of the viscosity is $\nu = 0.01$. The preconditioner is the modified AL preconditioner, where the linear systems associated with the diagonal sub-blocks A_{ii} are solved by a parallel iterative solver, using 10^{-5} as the inner relative residual tolerance. For the (2,2) block, we use $\hat{S}^{-1} = -\gamma \widehat{M}_p^{-1}$. We experimented with an inner flexible GMRES (FGMRES, see [24]) iteration preconditioned with the smoothed aggregation AMG preconditioner ML from the Trilinos package [25]. The reason why FGMRES is needed is because the number of cycles with the AMG preconditioner is not fixed and therefore the preconditioner is a variable one. Of course, this means that the outer iteration used to solve the whole saddle-point system must also be a flexible one.

The experiments are performed on a cluster consisting of 32 nodes and 128 processor cores. Each node has two dual core AMD 2214 2.2 GHz Opteron CPUs, 4 GB RAM and an 80

Table IX. Timings / Iteration numbers for the parallel implementation with Trilinos ML AMG solver

Problem size	Number of processors			
	2	4	8	16
2466	0.048/18	0.031/18	0.13/18	0.44/18
9538	0.34/19	0.16/20	0.19/20	0.23/20
37506	1.97/26	1.17/26	0.77/27	0.72/27
148738	11.9/33	6.64/33	3.72/33	2.57/27

GB drive. The program is compiled and run with Open MPI. Timings and (outer) FGMRES iteration counts using $p = 2, 4, 8$ and 16 processors are shown in Table IX for four different problem sizes.

The results show fairly good scalability when the problem is sufficiently large. Clearly, using more than 16 processors would not be warranted, even for the larger problem size. The observed variability in the iteration count for different numbers of processors (especially going from 8 to 16 processors) can be explained by the fact that the ML preconditioner changes with p , therefore the inner solves produce different approximate solutions for different p .

2.8. Extension to 3-D case

In this subsection we evaluate the performance of the ideal and modified AL preconditioners for Stokes and Oseen problems in 3-D. Besides the usual (convective) form of the Oseen equations, we also consider linearizations of the Navier-Stokes equations in rotation form; see, e.g., [26] and references therein for details.

For a stable discretization, the saddle point system is again of the form

$$\begin{pmatrix} A & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} f \\ g \end{pmatrix}, \quad \text{or} \quad \mathcal{A}x = b, \quad (19)$$

where for 3-D problems $A = \text{diag}(A_1, A_2, A_3)$ and $B = (B_1, B_2, B_3)$. Therefore the coefficient matrix for the equivalent augmented Lagrangian formulation is

$$\begin{aligned} A_\gamma &= A + \gamma B^T W^{-1} B \\ &= \begin{pmatrix} A_1 & 0 & 0 \\ 0 & A_2 & 0 \\ 0 & 0 & A_3 \end{pmatrix} + \gamma \begin{pmatrix} B_1^T \\ B_2^T \\ B_3^T \end{pmatrix} W^{-1} (B_1 \ B_2 \ B_3) \\ &= \begin{pmatrix} A_1 + \gamma B_1^T W^{-1} B_1 & \gamma B_1^T W^{-1} B_2 & \gamma B_1^T W^{-1} B_3 \\ \gamma B_2^T W^{-1} B_1 & A_2 + \gamma B_2^T W^{-1} B_2 & \gamma B_2^T W^{-1} B_3 \\ \gamma B_3^T W^{-1} B_1 & \gamma B_3^T W^{-1} B_2 & A_3 + \gamma B_3^T W^{-1} B_3 \end{pmatrix} \\ &=: \begin{pmatrix} A_{11} & A_{12} & A_{13} \\ A_{21} & A_{22} & A_{23} \\ A_{31} & A_{32} & A_{33} \end{pmatrix}. \end{aligned}$$

The ideal AL preconditioner is again given by (11). In the modified variant we replace A_γ

Table X. Iteration numbers of GMRES with the original and modified AL preconditioners for steady 3D Stokes problems (MAC).

	Original	Modified
Grid	Iter	Iter
$8 \times 8 \times 8$	9	12
$16 \times 16 \times 16$	9	12
$24 \times 24 \times 24$	-	13

with the block triangular approximation

$$\tilde{A}_\gamma = \begin{pmatrix} A_1 + \gamma B_1^T W^{-1} B_1 & \gamma B_1^T W^{-1} B_2 & \gamma B_1^T W^{-1} B_3 \\ 0 & A_2 + \gamma B_2^T W^{-1} B_2 & \gamma B_2^T W^{-1} B_3 \\ 0 & 0 & A_3 + \gamma B_3^T W^{-1} B_3 \end{pmatrix} = \begin{pmatrix} A_{11} & A_{12} & A_{13} \\ 0 & A_{22} & A_{23} \\ 0 & 0 & A_{33} \end{pmatrix}.$$

Note that in the 3-D case we drop three blocks: the (2,1), (3,1) and (3,2) blocks of A_γ , so the performance could be affected more than in the 2-D case. As in the 2-D case, each diagonal block A_{ii} represents a discrete scalar convection-diffusion operator. For the Stokes problem and for the rotation form of the Navier–Stokes equations, no convective term is present and each sub-block A_{ii} is symmetric and positive definite (SPD). On the other hand, in the rotation form each off-diagonal block A_{ij} (with $i \neq j$) contains additional coupling terms not present in the standard form of the Oseen problem. Finally, for unsteady problems an additional reaction term (also SPD) is present in each diagonal sub-block.

2.9. Numerical experiments: 3-D examples

We use a Marker-and-Cell (MAC) scheme, see [16], to discretize the Stokes and Oseen problems on the unit cube $\Omega = [0, 1] \times [0, 1] \times [0, 1]$. This scheme is known to be div-stable, hence no pressure stabilization is needed and the (2,2) block C in the saddle point problem (8) is zero. Homogeneous Dirichlet boundary conditions are imposed on the velocity components.

In the first experiment, we compare the original and the modified AL preconditioners on a steady Stokes problems. We use $\gamma = 1$ for both preconditioners for simplicity, and because it gives good results. Iteration counts for three grids of increasing size are shown in Table X. A “-” means that the problem has exceeded the memory constraint. Using a complete sparse Cholesky factorization on the sub-matrix A_γ makes the ideal AL preconditioner unfeasible already for rather coarse grids in 3-D. On the other hand, the modified AL preconditioner is able to handle larger problems, since only the diagonal blocks A_{ii} of A_γ need to be factored.

Our results indicate that for this problem, the rate of convergence with the modified AL preconditioner is independent of the mesh size.

Next, we proceed to the more challenging case of 3-D Oseen problems, in both convection and rotation forms. The divergence-free wind function (\mathbf{v} in (5)) is taken to be

$$\mathbf{v} = \begin{pmatrix} (2y-1)x(1-x) \\ (2x-1)y(1-y) \\ -2z(1-2x)(2y-1) \end{pmatrix}.$$

Table XI. GMRES iterations with the original and modified AL preconditioners for steady 3-D Oseen problem in convection form (MAC).

Viscosity	0.1		0.01		0.001	
Grid	Original	Modified	Original	Modified	Original	Modified
$8 \times 8 \times 8$	6	11 (1)	5	17 (0.1)	5	59 (0.01)
$16 \times 16 \times 16$	6	11 (1)	5	16 (0.1)	5	63 (0.01)
$24 \times 24 \times 24$	-	13 (0.1)	-	16 (0.1)	-	65 (0.01)
$32 \times 32 \times 32$	-	13 (0.1)	-	16 (0.1)	-	65 (0.01)

Table XII. GMRES iterations with the original and modified AL preconditioners for steady 3-D Oseen problem in rotation form (MAC).

Viscosity	0.1		0.01		0.001	
Grid	Original	Modified	Original	Modified	Original	Modified
$8 \times 8 \times 8$	6	12 (1)	5	20(0.1)	5	152 (0.01)
$16 \times 16 \times 16$	6	12 (1)	5	16 (0.1)	5	127 (0.01)
$24 \times 24 \times 24$	-	15 (0.1)	-	18 (0.1)	-	116 (0.01)
$32 \times 32 \times 32$	-	15 (0.1)	-	18 (0.1)	-	114 (0.01)

Table XIII. GMRES iterations with the original and modified AL preconditioners for unsteady 3-D Oseen problem in convection form (MAC).

Viscosity	0.1		0.01		0.001	
Grid	Original	Modified	Original	Modified	Original	Modified
$8 \times 8 \times 8$	7	9	7	19	6	26
$16 \times 16 \times 16$	8	10	7	17	7	26
$24 \times 24 \times 24$	-	10	-	15	-	22
$32 \times 32 \times 32$	-	12	-	15	-	19

The results are shown in Tables XI and XII. The original preconditioner is unfeasible on $24 \times 24 \times 24$ or larger grids due to exceeded memory limits. We use $\gamma = 1$ for the original preconditioner; for the modified one, γ is given in parentheses after the number of GMRES iterations.

While the ideal AL preconditioner is independent of the mesh size and viscosity, the modified preconditioner appears to be independent of h but shows a degradation of convergence rate when ν becomes very small. Also notice that γ must be taken smaller when $\nu = 0.001$. Also for small ν the modified preconditioner is less effective for the rotation form than for the convection form. This is due to the presence, in the rotation form, of relatively stronger coupling terms in the off-diagonal blocks A_{ij} ($j > i$) that are neglected when forming the preconditioner.

Next, we consider unsteady Oseen problems with $\sigma = h^{-1}$. Because of the dominance of the block diagonal part of the (1,1) block of the saddle point system, we can expect good performance of the modified AL preconditioner. As in 2-D, we set $\gamma = 1$ for both preconditioners. The (2,2) block is given by $\hat{S}^{-1} = -\gamma \hat{M}_p^{-1}$. Iteration counts are shown in

Table XIV. GMRES iterations with the original and modified AL preconditioners for unsteady 3-D Oseen problem in rotation form (MAC).

Viscosity	0.1		0.01		0.001	
Grid	Original	Modified	Original	Modified	Original	Modified
$8 \times 8 \times 8$	7	10	6	20	6	25
$16 \times 16 \times 16$	8	10	7	18	6	25
$24 \times 24 \times 24$	-	11	-	16	-	22
$32 \times 32 \times 32$	-	11	-	13	-	19

Tables XIII (convection form) and XIV (rotation from).

The performance of the modified AL preconditioner is quite satisfactory though some deterioration is observed with respect to ν . On the other hand, for $\nu = 0.1$ the iteration number is essentially independent of the mesh size, while for smaller ν the rate of convergence actually improves when the grid becomes finer. This is likely due to the fact that only sufficiently fine meshes are able to resolve the flow.

3. AL-type preconditioners for stabilized finite elements

In this section, we consider the Oseen problem discretized by stabilized finite elements, e.g., Q1-Q1 or Q1-P0 elements. We generalize the original AL preconditioner to stabilized elements, aiming to achieve robustness with respect to h and ν .

3.1. Problem formulation

Discretizing (1)–(4) with LBB unstable elements requires additional pressure stabilization terms in the finite element problem formulation, see e.g. [28]. In this case, the (2,2) block of the saddle point matrix is no longer zero; it is replaced by $-C$ with some symmetric positive semidefinite matrix C , see the remark below.

Remark 1. For the case of continuous pressure elements, the matrix C corresponds to the bilinear form $\sum_{K \in \mathcal{T}} \delta \int_K \nabla p \nabla q \, d\mathbf{x}$, where the summation runs over all triangles K in the triangulation \mathcal{T} of Ω . For P0 pressure elements the corresponding bilinear form $\sum_{\Gamma \in \mathcal{T}} \delta |\Gamma|^{-1} \int_{\Gamma} [p][q] \, ds$ penalizes the pressure jumps $[p]$ over all internal edges Γ of the triangulation \mathcal{T} ; δ is a stabilization parameter. In either case the spectral properties of the matrix C are somewhat similar to those of a scaled Laplacian discretization. Note also that C is symmetric.

We use the following choice of the stabilization parameter (see [29]): $\delta = \frac{\beta h^2}{\nu + h \|w\|}$, where h is the mesh size and w is the wind function, and $\beta = 1$. Assuming (for ease of notation) that

$g = 0$, the resulting linear system reads

$$\begin{pmatrix} A & B^T \\ B & -C \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} f \\ 0 \end{pmatrix}. \quad (20)$$

3.2. Augmented linear systems and an AL preconditioner

Due to the presence of a nonzero (2,2) block, the augmentation of system (20) must be done differently than in the case of stable finite elements. As before, let $\gamma > 0$ and let W be a symmetric positive definite matrix. Then from $Bu - Cp = 0$ it follows that

$$\gamma B^T W^{-1} Bu - \gamma B^T W^{-1} Cp = 0.$$

Adding the above equation to $Au + B^T p = f$ gives

$$(A + \gamma B^T W^{-1} B)u + (B^T - \gamma B^T W^{-1} C)p = f.$$

Therefore, the (first) augmented linear system is

$$\begin{pmatrix} A_\gamma & B_\gamma^T \\ B & -C \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} f \\ 0 \end{pmatrix}, \quad \text{or} \quad \hat{\mathcal{A}}x = b, \quad (21)$$

where $A_\gamma = A + \gamma B^T W^{-1} B$ and $B_\gamma^T = B^T - \gamma B^T W^{-1} C$.

Notice that in (21), the (1,2) block B_γ^T is not equal to the transpose of the (2,1) block B . To get a more ‘symmetric’ augmented linear system, we can obtain from $Bu - Cp = 0$ the equation

$$-\gamma CW^{-1} Bu + \gamma CW^{-1} Cp = 0.$$

Then, combining this equation with $Bu - Cp = 0$, we have

$$(B - \gamma CW^{-1} B)u - (C - \gamma CW^{-1} C)p = 0.$$

Letting $C_\gamma = C - \gamma CW^{-1} C$, we obtain the second augmented system

$$\begin{pmatrix} A_\gamma & B_\gamma^T \\ B_\gamma & -C_\gamma \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} f \\ 0 \end{pmatrix}. \quad (22)$$

Numerical experiments suggest that applying the AL preconditioner to (21) produces almost the same results as to (22). Therefore, all results of numerical experiments will be shown for (21) in the sequel. However, our analysis will apply to both (21) and to the ‘symmetrized’ system (22).

Similar to the stable FE case we build the block-triangular preconditioner in the form

$$\mathcal{P} = \begin{pmatrix} A_\gamma & 0 \\ B & -\hat{S} \end{pmatrix}. \quad (23)$$

Here in order to simplify the action of \mathcal{P}^{-1} we use a block lower triangular AL preconditioner so as to have B in the (2,1) block rather than the more cumbersome B_γ^T in the (1,2) block. For the ‘symmetrized’ system (22), B should be replaced by B_γ ; both lower and upper triangular preconditioners are essentially equivalent in this case.

For LBB-stable elements, in Section 2.2 we set $\hat{S} = \gamma^{-1} W = \gamma^{-1} \widehat{M}_p$. As we will see, for the case of $C \neq 0$ the choice of \hat{S} and W is more delicate. From (12) one notices that \hat{S} intends

to approximate the pressure Schur complement of the augmented system, i.e. the matrices $S_\gamma := B_\gamma A_\gamma^{-1} B_\gamma^T + C_\gamma$ for (22) and $\tilde{S}_\gamma := B A_\gamma^{-1} B^T + C$ for (21). Recall the notation for the pressure Schur complement matrix of the non-augmented problem: $S = B A^{-1} B^T + C$. The following result, which extends the representation in (13) to the case of $C \neq 0$, will help us to set W , build the preconditioner \hat{S} , and analyze the spectrum of the preconditioned system.

Lemma 1. *Assuming all the relevant matrices are invertible, it holds*

$$S_\gamma^{-1} = S^{-1} + \gamma(W - \gamma C)^{-1}, \quad (24)$$

$$\tilde{S}_\gamma^{-1} = S^{-1}(I - \gamma C W^{-1}) + \gamma W^{-1}. \quad (25)$$

Proof. The matrix $X := -S_\gamma^{-1}$ is the (2,2) block of the inverse of the coefficient matrix in (22). Denoting by Y the (1,2) block of this inverse matrix we get the following system of matrix equations:

$$(A + \gamma B^T W^{-1} B)Y + B^T(I - \gamma W^{-1} C)X = 0, \quad (26)$$

$$(I - \gamma C W^{-1})BY - (I - \gamma C W^{-1})CX = I. \quad (27)$$

From (27) we get $BY = (I - \gamma C W^{-1})^{-1} + CX$. Substituting this into (26) and applying A^{-1} lead to

$$Y = -\gamma A^{-1} B^T W^{-1} (I - \gamma C W^{-1})^{-1} - A^{-1} B^T X.$$

Now substituting Y to (27) gives, after simple manipulations,

$$-(BA^{-1} B^T + C)X(I - \gamma C W^{-1}) = I + \gamma BA^{-1} B^T W^{-1}.$$

By straightforward computations one verifies that the last equation is solved by matrix $X = -(BA^{-1} B^T + C)^{-1} - \gamma(W - \gamma C)^{-1}$. Thus (24) is proved. The result in (25) follows from the obvious identity $(I - \gamma C W^{-1})\tilde{S}_\gamma = S_\gamma$. \square

The expressions (24) and (25) suggest that the auxiliary matrix W should be such that $W - \gamma C$ is positive definite. Below we consider the following two choices of W satisfying this constraint:

$$W_1 := M_p + \gamma C,$$

$$W_2 := M_p \quad \text{with} \quad 0 < \gamma \leq (2\|M_p^{-1}C\|)^{-1}.$$

Similar to the LBB-stable case, in practical computations M_p is replaced by its diagonal approximation \widehat{M}_p . Let us briefly comment on both choices of W .

Remark 2. Setting $W = W_1$ will lead to a simple choice of preconditioner \hat{S} such that the preconditioned system enjoys the same eigenvalue bounds as in the LBB-stable case, i.e. (15)–(16). At the same time, $W = W_1$ involves the Laplacian-type matrix C , cf. Remark 1. Hence the inverse W^{-1} may become an (almost) full matrix, resulting in that $A_\gamma = A + B^T W^{-1} B$ is an (almost) full matrix and consequently making the solution of linear systems with A_γ much more difficult. This happens for example with Q_1 - Q_1 elements. For Q_1 - P_0 elements, however, the matrix C has a special block-diagonal structure, which leads to a relatively cheap solve with A_γ . In numerical experiments we will use W_1 only with Q_1 - P_0 elements.

The choice $W = W_2$ preserves the sparsity of A_γ . However, the restriction on γ yields the decrease of γ when ν is small and h tends to zero, since for small ν it holds $\|M_p^{-1}C\| = O(h^{-1})$, see the discussion on matrix C in Remark 1. Thus less augmentation is introduced and the performance of the solver becomes more sensitive to the variation in ν and h .

In the next Section we present the eigenvalue analysis and show the corresponding choices of \hat{S} for both cases $W = W_1$ and $W = W_2$. We shall also discuss a third (intermediate) alternative of setting the augmentation and preconditioning, which is not covered by our analysis, but shows stable and almost ν and h -independent convergence behavior while keeping the matrix A_γ sparse.

3.3. Eigenvalue analysis and the choice of \hat{S}

We consider the following generalized eigenvalue problem:

$$\begin{pmatrix} A_\gamma & B_\gamma^T \\ B & -C \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \lambda \begin{pmatrix} A_\gamma & 0 \\ B & -\hat{S} \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix}. \quad (28)$$

For the symmetrized system, matrix B in the (2,1) block is replaced by B_γ and matrix C in the (2,2) block by C_γ . As in the case of stable finite elements, we consider the eigenvalue problem

$$Sq = \mu M_p q, \quad (29)$$

where $S = BA^{-1}B^T$, and obtain bounds on λ in terms of μ .

For $W = W_1$ one immediately gets from (24) and (25)

$$\begin{aligned} S_\gamma^{-1} &= S^{-1} + \gamma M_p^{-1}, \\ \tilde{S}_\gamma^{-1} &= S^{-1} M_p (M_p + \gamma C)^{-1} + \gamma (M_p + \gamma C)^{-1}. \end{aligned}$$

Therefore, setting

$$\hat{S} := \gamma^{-1} M_p + C \quad \text{for (21)} \quad \text{or} \quad \hat{S} := \gamma^{-1} M_p \quad \text{for (22)}, \quad (30)$$

we obtain with the same arguments as for the case of $C = 0$ in Section 2.2 that all non-unit eigenvalues of (28) satisfy $S^{-1}p + \gamma M_p^{-1}p = \lambda^{-1} \gamma M_p^{-1}p$, where $p \neq 0$, and thus

$$\lambda = \frac{\gamma\mu}{1 + \gamma\mu}.$$

This representation is identical to the one in (14). Therefore, we obtain the following theorem.

Theorem 1. *Assume $W = W_1$ and \hat{S} is defined as in (30). The preconditioned matrix $\mathcal{P}^{-1}\hat{\mathcal{A}}$ has the eigenvalue 1 of multiplicity at least n . All other (nonunit) eigenvalues satisfy the following bounds:*

$$0 < \min_{\mu} \frac{\gamma a_\mu}{1 + \gamma a_\mu} \leq a_\lambda \leq 1, \quad |b_\lambda| \leq \max_{\mu} \min \left\{ \gamma |b_\mu|, \frac{1}{\gamma |b_\mu|} \right\} \leq 1,$$

where $\lambda = a_\lambda + i b_\lambda$ and $\mu = a_\mu + i b_\mu$.

We noted already that the choice $W = W_1$ is not always practical. The following theorem shows eigenvalue bounds for the case $W = W_2$.

Theorem 2. Assume $W = W_2$, $0 < \gamma \leq (2\|M_p^{-1}C\|)^{-1}$ and $\hat{S} = \gamma^{-1}M_p$. The preconditioned matrix $\mathcal{P}^{-1}\hat{A}$ has the eigenvalue 1 of multiplicity at least n . All other (nonunit) eigenvalues satisfy the following bounds: for the non-symmetric augmentation (21),

$$0 < \min_{\mu} \frac{\gamma a_{\mu}}{1 + \gamma a_{\mu}} \leq a_{\lambda} \leq 1, \quad |b_{\lambda}| \leq \max_{\mu} \min \{\gamma |b_{\mu}|, 1\} \leq 1, \quad (31)$$

and for the symmetrized augmentation (22),

$$0 < \min_{\mu} \frac{(\gamma a_{\mu})^2}{2 + (\gamma a_{\mu})^2} \leq a_{\lambda} \leq 1, \quad |b_{\lambda}| \leq 1, \quad (32)$$

where $\lambda = a_{\lambda} + i b_{\lambda}$ and $\mu = a_{\mu} + i b_{\mu}$.

Proof. From (28), we immediately get that $\lambda = 1$ is eigenvalue of multiplicity (at least) n and any vector $[u; 0]$ with $u \neq 0$ is a corresponding eigenvector. The remaining eigenvalues λ satisfy

$$\tilde{S}_{\gamma} p = \lambda \hat{S} p \quad \text{or} \quad S_{\gamma} p = \lambda \hat{S} p.$$

First we consider the case of non-symmetrized augmentation. For $W = M_p$, $\hat{S} = \gamma^{-1}M_p$, using representation (24), we obtain

$$S p = \frac{\lambda}{1 - \lambda} \left(\frac{1}{\gamma} M_p - C \right) p. \quad (33)$$

For brevity, we let $\eta = \frac{\lambda}{1 - \lambda}$ and $Q = \frac{1}{\gamma} M_p - C$. It follows from Bendixson's Theorem [27] that

$$\min_p \frac{p^T D p}{p^T Q p} \leq a_{\eta} \leq \max_p \frac{p^T D p}{p^T Q p}, \quad |b_{\eta}| \leq \max_p \frac{|p^T R p|}{p^T Q p},$$

where $\eta = a_{\eta} + i b_{\eta}$, $D = B \frac{A^{-1} + A^{-T}}{2} B^T + C$ is the symmetric part of S , and $R = B \frac{A^{-1} - A^{-T}}{2} B^T$ is its skew-symmetric part. Using (33) we shall obtain bounds for λ in terms of μ from (29). Since γ satisfies $0 < \gamma \leq (2\|M_p^{-1}C\|)^{-1}$, it holds

$$\frac{1}{2\gamma} M_p \leq \frac{1}{\gamma} M_p - C \leq \frac{1}{\gamma} M_p. \quad (34)$$

Therefore, we have

$$\begin{aligned} \gamma \min_p \frac{p^T D p}{p^T M_p p} &\leq \min_p \frac{p^T D p}{p^T Q p} \leq a_{\eta} \leq \max_p \frac{p^T D p}{p^T Q p} \leq 2\gamma \max_p \frac{p^T D p}{p^T M_p p}, \\ |b_{\eta}| &\leq \max_p \frac{|p^T R p|}{p^T Q p} \leq 2\gamma \max_p \frac{|p^T R p|}{p^T M_p p}. \end{aligned}$$

Applying Bendixson's Theorem also to the generalized eigenvalue problem (29) gives

$$\min_p \frac{p^T D p}{p^T M_p p} \leq a_{\mu} \leq \max_p \frac{p^T D p}{p^T M_p p}, \quad |b_{\mu}| \leq \max_p \frac{|p^T R p|}{p^T M_p p}.$$

There are extreme eigenvalues for which the Bendixson bounds are attained. **The last statement is incorrect!** Hence

$$a_{\eta} \geq \gamma \min_{\mu} a_{\mu} > 0, \quad |b_{\eta}| \leq 2\gamma \max_{\mu} |b_{\mu}|. \quad (35)$$

Solving $\eta = \frac{\lambda}{1-\lambda}$ for a_λ and b_λ , we have

$$a_\lambda = \frac{a_\eta(1+a_\eta) + b_\eta^2}{(1+a_\eta)^2 + b_\eta^2}, \quad b_\lambda = \frac{b_\eta}{(1+a_\eta)^2 + b_\eta^2}.$$

From this and (35) the result in (31) follows.

Now consider the case of the symmetrized augmentation. For $W = M_p$, $\hat{S} = \gamma^{-1}M_p$, using representation (25), we obtain for $\lambda \neq 1$

$$S^{-1}p + \gamma(M_p - \gamma C)^{-1}p = \lambda^{-1}\gamma M_p^{-1}p.$$

Denoting by η the generalized eigenvalue of

$$S^{-1}p + \gamma M_p^{-1}p = \eta^{-1}\gamma M_p^{-1}p, \quad (36)$$

we get using Bendixson's Theorem, (34) and $\eta^{-1} = \gamma^{-1}\mu^{-1} + 1$, the following bounds:

$$\begin{aligned} 1 \leq \min_{\eta} \Re(\eta^{-1}) &\leq \Re(\lambda^{-1}) \leq 2 \max_{\eta} \Re(\eta^{-1}) \leq 2((\gamma \min_{\mu} a_{\mu})^{-1} + 1), \\ |\Im(\lambda^{-1})| &\leq \max_{\eta} |\Im(\eta^{-1})| \leq (\gamma \min_{\mu} a_{\mu})^{-1}. \end{aligned} \quad (37)$$

Finally, from estimates (37) we deduce (32). \square

The bounds for λ in theorems 1 and 2 are written in terms of the eigenvalues μ from (29). Following the same argument as in [17], we can prove that the largest and smallest real and imaginary parts of μ are independent of h , but depend on ν . The resulting eigenvalue bounds for λ are very similar to those for the LBB-stable case ($C = 0$) from Section 2 and [8]. This suggests that the choice $\gamma = O(1)$ leads to a method which is essentially insensitive to variations of parameters ν and h . However, for the practical choice of $W = \widehat{M}_p$ we have the restriction on γ . Numerical experiments show that with the setting of W and \hat{S} from Theorem 2 the restriction is indeed important and prohibits the choice $\gamma = O(1)$ for all values ν and h of interest. The situation looks better from the numerical viewpoint if one sets (for non-symmetrized case (21))

$$W = \widehat{M}_p, \quad \hat{S} := \gamma^{-1}M_p + C. \quad (38)$$

This combination of the augmentation and preconditioning, which is intermediate between those in Theorems 1 and 2, is not covered by the eigenvalue analysis above. However, in practice it leads to performance which is nearly robust with respect to ν and h .

3.4. Numerical experiments

In this subsection, we study the numerical behavior of the AL preconditioners. Two choices of γ are tested. The first one is to set $\gamma = 1$ as in the stable case, while the second follows the restriction of Theorem 2, i.e. $\gamma = (2\|M_p^{-1}C\|)^{-1}$.

The first set of experiments is to use the AL preconditioner to solve the lid driven cavity problem discretized by Q1-P0 finite elements. With this choice of elements one can set $W = \widehat{M}_p + \gamma C$, $\hat{S} = \gamma^{-1}\widehat{M} + C$ for the AL preconditioner, cf. Remark 2. In this case no restriction on γ applies (see Theorem 1), so we set $\gamma = 1$. The results are shown in Table XV

Table XV. GMRES iterations with the AL preconditioner for steady 2-D lid driven cavity problem (Q1-P0), $W = \widehat{M}_p + \gamma C$, $\widehat{S} = \gamma^{-1}\widehat{M} + C$, $\gamma = 1$.

Viscosity	0.1	0.01	0.001	0.1	0.01	0.001
Grid	uniform grid			stretched grid		
16×16	7	6	5	7	6	5
32×32	8	6	6	8	6	6
64×64	8	7	7	9	7	7
128×128	9	6	8	9	7	7

Table XVI. GMRES iterations for steady 2-D driven cavity problem (Q1-Q1), $W = \widehat{M}_p$, $\widehat{S} = \gamma^{-1}\widehat{M}$.

Viscosity	0.1		0.01		0.001	
Grid	$\gamma = 1$	$\gamma = \frac{1}{2\ \widehat{M}_p^{-1}C\ }$	$\gamma = 1$	$\gamma = \frac{1}{2\ \widehat{M}_p^{-1}C\ }$	$\gamma = 1$	$\gamma = \frac{1}{2\ \widehat{M}_p^{-1}C\ }$
16×16	14	16 (0.1467)	16	19 (0.0880)	16	22 (0.0821)
32×32	18	16 (0.1061)	25	22 (0.0473)	23	28 (0.0415)
64×64	23	16 (0.0857)	36	25 (0.0269)	36	36 (0.0211)
128×128	25	17 (0.0755)	49	30 (0.0167)	58	52 (0.0109)

Table XVII. The smallest real parts of the eigenvalues from (28), $W = \widehat{M}_p$, $\widehat{S} = \gamma^{-1}\widehat{M}$.

Grid	The smallest real part
16×16	0.163
32×32	0.0929
64×64	0.0481

for the lid driven cavity problem on both uniform and stretched grids. The h -independence and ν -independence of the AL preconditioner are obvious from the data.

For the case of Q1-Q1 elements the choice of $W = \widehat{M}_p + \gamma C$ is not practical, cf. Remark 2. Thus, in all further experiments we set $W = \widehat{M}_p$ and show results with only Q1-Q1 elements. For Q1-P0 elements, the results were found to be quite similar. First we try the simple choice of preconditioner $\widehat{S} := \gamma^{-1}\widehat{M}_p$. From Table XVI, we can see that the AL preconditioner performs quite well with this stabilization, but no h -independence is achieved, except for the case $\nu = 0.1$. Moreover, $\gamma = 1$ cannot be recommended, as has been predicted by the eigenvalue analysis in the previous section. The values for $\gamma = (2\|\widehat{M}_p^{-1}C\|)^{-1}$ are shown in parentheses, and they decrease as the grid is refined. This explains why we cannot get h -independence. As γ becomes smaller, the smallest real parts of the eigenvalues of the preconditioned coefficient matrix also gets smaller. We can observe this from Table XVII for $\nu = 0.001$.

We also experiment with replacing $\widehat{S} = \gamma^{-1}\widehat{M}$ with $\widehat{S} = \gamma^{-1}\widehat{M} + C$ in the preconditioner and show the results in Tables XVIII for the lid driven cavity and XIX for the backward facing step. The iteration counts with the exact solve of A_γ are presented in the “Original” columns, while the “Modified” ones shows using the block triangular approximation to solve A_γ , as

Table XVIII. GMRES with the ideal AL preconditioner for steady 2D Oseen problem (Q1-Q1),
 $W = \widehat{M}_p$, $\widehat{S} = \gamma^{-1}\widehat{M} + C$, $\gamma = 1$.

Viscosity	0.1		0.01		0.001	
Grid	Original	Modified	Original	Modified	Original	Modified
16×16	10	11	13	19	15	37
32×32	10	12	14	18	18	41
64×64	11	12	16	16	20	41
128×128	11	12	18	15	22	40

Table XIX. GMRES iterations with the ideal AL preconditioner for backward facing step problem (Q1-Q1), $W = \widehat{M}_p$, $\widehat{S} = \gamma^{-1}\widehat{M} + C$, $\gamma = 1$.

Viscosity	0.1		0.01		0.005	
Grid	Original	Modified	Original	Modified	Original	Modified
16×48	10	15	13	23	14	25
32×96	11	15	15	23	16	28
64×192	11	16	17	25	18	32
128×384	12	17	19	28	20	39

we did in Section 2. We choose $\gamma = 1$ for the original preconditioner; for the modified one, first we determine the optimal γ experimentally on the coarsest grid, then we use the $\sqrt{2}$ rule for the finer grids. The iteration counts show the excellent behavior of both preconditioners, with (near) independence on the mesh size and a mild dependence on ν . In the case of small viscosity in the backward facing step problem, a modest deterioration in the performance of the modified AL preconditioner is observed as the grid is refined, but even in that case the convergence is quite fast.

4. Conclusions

In this paper, we have introduced and studied modifications of the augmented Lagrangian-based preconditioner for Oseen-type problems. In addition, we have extended the AL approach to the case of stabilized finite element pairs, and heuristics for the choice of the augmentation parameter γ have been investigated. Theoretical analysis and extensive numerical experimentation on a variety of test problems indicate that overall, the modified AL preconditioners are very robust and perform quite well on both steady and unsteady problems over a wide range of values of the viscosity ν .

One advantage of the modified AL preconditioners proposed in this paper is that they can be readily implemented using standard off-the-shelf algebraic multilevel solvers developed for elliptic PDEs, in particular parallel AMG-type solvers. Our experiments show that this can be an effective approach for the parallel solution of finite element discretizations of the incompressible Navier-Stokes equations.

REFERENCES

1. Elman HC, Silvester DJ, Wathen AJ. *Finite Elements and Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*. Oxford Series in Numerical Mathematics and Scientific Computation, Oxford University Press: Oxford, 2005.
2. Saad Y. *Iterative Methods for Sparse Linear Systems* (2nd edn). SIAM: Philadelphia, PA, 2003.
3. Wesseling P. *Principles of Computational Fluid Dynamics*. Springer Series in Computational Mathematics 29, Springer: New York, 2001.
4. Trottenberg U, Oosterlee C, Schuller A. *Multigrid*, Academic Press: San Diego, 2001.
5. Saad Y, Schultz MH. *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, *SIAM Journal on Scientific and Statistical Computing* 1986; **7**:856–869.
6. Hamilton SP, Benzi M, Haber E. New multigrid smoothers for the Oseen problem. Submitted to *Numerical Linear Algebra with Applications*, May 2009.
7. Benzi M, Golub GH, and Liesen J. Numerical solution of saddle point problems. *Acta Numerica* 2005; **14**:1–137.
8. Benzi M and Olshanskii MA. An augmented Lagrangian-based approach to the Oseen problem. *SIAM Journal on Scientific Computing* 2006; **28**:2095–2113.
9. Olshanskii MA, Benzi M. An Augmented Lagrangian approach to linearized problems in hydrodynamic stability. *SIAM Journal on Scientific Computing* 2008; **30**:1459–1473.
10. Fortin M, Glowinski R. *Augmented Lagrangian Methods: Applications to the Numerical Solution of Boundary-Value Problems*. North-Holland: Amsterdam-New York, 1983.
11. Elman H, Howle VE, Shadid J, Silvester D, Tuminaro R. Least squares preconditioners for stabilized discretizations of the Navier-Stokes equations. *SIAM Journal on Scientific Computing* 2007; **30**:290–311.
12. ur Rehman M, Vuik C, Segal G. Preconditioners for the steady incompressible Navier-Stokes problem. *International Journal of Applied Mathematics* 2008; **38**:223–232.
13. ur Rehman M, Vuik C, Segal G. SIMPLE-type preconditioners for the Oseen problem. *International Journal for Numerical Methods in Fluids* 2009; to appear.
14. ur Rehman M, Vuik C, Segal G. A comparison of preconditioners for incompressible Navier-Stokes solvers. *International Journal for Numerical Methods in Fluids* 2008; **57**:1731–1751.
15. Schöberl J. Multigrid methods for a parameter dependent problem in primal variables. *Numerische Mathematik* 1999; **84**:97–119.
16. Harlow FH, Welch JE. Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface. *Physics of Fluids* 1965; **8**:2182–2189.
17. Elman HC, Silvester DJ. Fast nonsymmetric iterations and preconditioning for the Navier-Stokes equations. *SIAM Journal on Scientific Computing* 1996; **17**:33–46.
18. Elman HC, Ramage A, Silvester DJ. IFISS: A Matlab toolbox for modelling incompressible flow. *ACM Transactions on Mathematical Software* 2007; **33**: Article 14.
19. Silvester DJ, Elman HC, Ramage A. IFISS: Incompressible Flow Iterative Solution Software, available online at <http://www.manchester.ac.uk/ifiss>. [June 2009].
20. Amestoy PR, Davis TA, Duff IS. An approximate minimum degree ordering algorithm. *SIAM Journal on Matrix Analysis and Applications* 1996; **17**:886–905.
21. Cahouet J, Chabard JP. Some fast 3D finite element solvers for the generalized Stokes problem. *International Journal for Numerical Methods in Fluids* 1998; **8**:869–895.
22. Notay Y. An aggregation-based algebraic multigrid method. Report GANMN 08-02, Université Libre de Bruxelles, Brussels, Belgium, 2008 (revised 2009).
23. Notay Y. Aggregation-based algebraic multigrid method, available online at <http://homepages.ulb.ac.be/~ynotay/AGMG/index.html>. [March 2008].
24. Saad Y. A flexible inner-outer preconditioned GMRES algorithm, *SIAM Journal on Scientific Computing* 1993; **14**:461–469.
25. Heroux MA, Bartlett RA, Howle VE, Hoekstra RJ, Hu JJ, Kolda TG, Lehoucq RB, Long KR, Pawlowski RP, Phipps ET, Salinger AG, Thornquist HK, Tuminaro RS, Willenbring JS, Williams A, Stanley KS. An overview of the Trilinos project, *ACM Transactions on Mathematical Software* 2005; **31**:397–423.
26. Benzi M, Liu J. An efficient solver for the Navier–Stokes equations in rotation form. *SIAM Journal on Scientific Computing* 2007; **29**:1959–1981.
27. Stoer J, Bulirsch R. *Introduction to Numerical Analysis* (2nd edn). Springer-Verlag: New York, 1993.
28. Hughes TJR, Franca LP, Balestra M. A new finite element formulation for computational fluid dynamics. V: Circumventing the Babuska-Brezzi condition: a stable Petrov-Galerkin formulation of the Stokes problem accommodating equal-order interpolations. *Computer Methods in Applied Mechanics and Engineering* 1986; **59**:85–99.
29. Codina R. A stabilized finite element method for generalized stationary incompressible flows. *Computer Methods in Applied Mechanics and Engineering* 2001; **190**:2681–2706.