

Robust preconditioning for finite element discretizations of the steady incompressible Navier-Stokes equations

Michele Benzi

Department of Mathematics and Computer Science
Emory University

Sparse Days at CERFACS
15-17 June, 2010

Joint work with Maxim Olshanskii (Moscow State University)
and Zhen Wang (Emory University)

- 1 Introduction
- 2 Augmented Lagrangian-based preconditioner for stable finite elements
- 3 Generalization of AL-based preconditioners for stabilized finite elements
- 4 Conclusions

Outline

- 1 Introduction
- 2 Augmented Lagrangian-based preconditioner for stable finite elements
- 3 Generalization of AL-based preconditioners for stabilized finite elements
- 4 Conclusions

Navier–Stokes equations

The steady incompressible Navier–Stokes equations are:

$$-\nu \Delta \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u} + \nabla p = \mathbf{f} \quad \text{on } \Omega \quad (1)$$

$$\operatorname{div} \mathbf{u} = 0 \quad \text{on } \Omega \quad (2)$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \partial\Omega \quad (3)$$

- $\Omega \subset \mathbb{R}^d$ ($d = 2, 3$): an open bounded domain with boundary $\partial\Omega$.
- \mathbf{f} : a force field , \mathbf{g} : boundary data.
- ν : the kinematic viscosity (proportional to Re^{-1}) .
- Δ : vector Laplacian, ∇ : gradient, div : divergence.
- Goal: Find a velocity field $\mathbf{u} = \mathbf{u}(\mathbf{x})$ and pressure field $p = p(\mathbf{x})$ that satisfy (1)–(3).

Oseen equations

Linearization of the Navier–Stokes system by Picard iteration results in a sequence of Oseen problems:

$$-\nu\Delta\mathbf{u} + (\mathbf{v} \cdot \nabla)\mathbf{u} + \nabla p = \mathbf{f} \quad \text{on } \Omega \quad (4)$$

$$\operatorname{div} \mathbf{u} = 0 \quad \text{on } \Omega \quad (5)$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \partial\Omega \quad (6)$$

- \mathbf{v} is the velocity field computed from a previous iteration.
- This fixed-point iteration is convergent under mild assumptions.
- Sometimes used to get a good initial guess for Newton's method.

A good reference: Elman, Silvester, and Wathen's book (2005) and accompanying software (IFISS).

Saddle point system

Discretization of the Oseen equations using finite differences, finite elements or other discretization methods results in a large and sparse saddle point system

$$\begin{pmatrix} A & B^T \\ B & -C \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} f \\ g \end{pmatrix}. \quad (7)$$

- u and p represent the discrete velocity and pressure, respectively.
- $A \in \mathbb{R}^{n \times n}$ is the discretization of the diffusion and convection terms.
- $B^T \in \mathbb{R}^{n \times m}$ is the discrete gradient, and B is the (negative) discrete divergence.
- $C \in \mathbb{R}^{m \times m}$ is a stabilization matrix:
 - For LBB (“inf-sup”) stable finite elements, $C = 0$;
 - Otherwise, $C \neq 0$ is usually symmetric positive semidefinite.
- f and g contain forcing and boundary terms.

Wanted: scalable solution methods that are robust with respect to type of discretization, mesh size, and viscosity.

Outline

- 1 Introduction
- 2 Augmented Lagrangian-based preconditioner for stable finite elements**
- 3 Generalization of AL-based preconditioners for stabilized finite elements
- 4 Conclusions

Problem Formulation

An LBB (“inf-sup”)-stable finite element discretization of the Oseen problem gives rise to the following system:

$$\begin{pmatrix} A & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} f \\ g \end{pmatrix}. \quad (8)$$

The equivalent augmented Lagrangian formulation (Fortin and Glowinski, 1982) is given by

$$\begin{pmatrix} A + \gamma B^T W^{-1} B & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} f + \gamma B^T W^{-1} g \\ g \end{pmatrix}, \quad \text{or} \quad \hat{A}x = \hat{b}, \quad (9)$$

where W is symmetric positive definite (e.g., the pressure mass matrix M_p or its diagonal \hat{M}_p), and $\gamma > 0$.

AL-based preconditioner

Letting $A_\gamma := A + \gamma B^T W^{-1} B$ and $\hat{f} := f + \gamma B^T W^{-1} g$,

$$\begin{pmatrix} A_\gamma & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} \hat{f} \\ g \end{pmatrix}. \quad (10)$$

B. and Olshanskii (SISC, 2006) proposed the following block preconditioner:

$$\mathcal{P} = \begin{pmatrix} A_\gamma & B^T \\ 0 & \hat{S} \end{pmatrix} \quad \text{with} \quad \hat{S}^{-1} = -\nu \hat{M}_p^{-1} - \gamma W^{-1}. \quad (11)$$

Applying \mathcal{P}^{-1} to a vector mainly requires one solve with A_γ and one with \hat{S} .

The rate of convergence of Krylov subspace methods with this preconditioner is optimal with exact solves:

- independent of the grid;
- independent of viscosity.

AL-based preconditioner

In practice, however, exact solves with $A_\gamma = A + \gamma B^T W^{-1} B$ are too expensive. Inexact solves must be used.

Can we do this without spoiling the convergence properties too much?

Three main options:

- 1 Solve systems with the complete matrix A_γ inexactly;
- 2 Replace A_γ with a simpler matrix and solve exactly with sparse LU;
- 3 Replace A_γ with a simpler matrix and solve inexactly with (A)MG.

All three options have been tested; only options 1 and 3 lead to scalable algorithms.

For option 1, a geometric MG algorithm by Schöberl is adapted to A_γ .

A few results follow.

Numerical results

Viscosity	1	0.1	0.01	0.001	0.0001
Mesh size	Constant wind				
1/16	7	5	5	6	6
1/32	7	5	6	7	8
1/64	5	5	6	5	7
1/128	5	5	5	5	6
Mesh size	Rotating vortex				
1/16	5	5	6	10	15
1/32	4	4	5	10	21
1/64	4	4	5	9	18
1/128	4	5	5	7	14

Table: Bi-CGStab iterations (isoP2-P0, \widehat{A}_γ^{-1} is one W(1,1)-cycle, $\gamma = 1$)

Solution cost is linear in the number of unknowns, and the algorithm is scalable. However, geometric MG is not well suited for general grids and geometries.

Modified AL-based preconditioner

Observe that for 2D problems we have:

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

Using the following block triangular approximation to A_γ :

$$\tilde{A}_\gamma = \begin{pmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{pmatrix},$$

we define

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

This is the modified AL-based preconditioner; see B., Olshanskii and Wang (IJNMF, 2010). The 3D case is analogous.

Modified AL-based preconditioner

The application of the preconditioner is greatly simplified, since only linear systems associated with the diagonal blocks

$$A_{11} = A_1 + \gamma B_1^T W^{-1} B_1 \quad \text{and} \quad A_{22} = A_2 + \gamma B_2^T W^{-1} B_2$$

need to be solved. These matrices can be interpreted as discretizations of anisotropic convection-diffusion operators.

Two main questions:

- How to efficiently solve linear systems with A_{ii} , and how accurately?
- How to pick γ ?

For 2D problems we tested both sparse LU (exact) and AMG (inexact). For 3D, only AMG (inexact) is tested, as there is not enough memory for LU.

We choose γ using Fourier analysis; see B. and Wang (SISC, submitted).

Numerical results

Viscosity	0.1		0.01		0.005		0.001	
Grid	FA	Opt	FA	Opt	FA	Opt	FA	Opt
16×16	9	9	12	12	26	15	42	23
32×32	10	9	11	11	20	14	37	29
64×64	9	9	11	11	13	13	33	27
128×128	9	9	10	10	13	12	25	24

Table: GMRES iterations with modified AL preconditioner. Lid driven cavity from IFISS, Q2-Q1 FEM, uniform grids.

- **Opt:** choice of γ that minimizes iteration counts (by trial-and-error);
- **FA:** choice of γ by Fourier analysis.

The iterations are essentially h -independent and only mildly dependent of ν .

Eigenvalues of preconditioned matrices

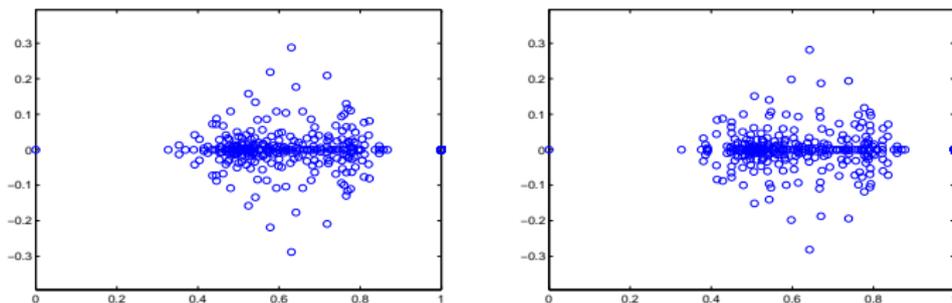


Figure: Plots of the eigenvalues of the preconditioned Oseen matrix (lid driven cavity, Q2-Q1, 32×32 uniform grid, $\nu = 0.01$). Left: with optimal γ . Right: with γ chosen by Fourier analysis.

The two values of γ are very close: 0.05 vs. 0.056.

The eigenvalue $\lambda = 1$ has multiplicity n (for all γ).

Results for stretched grids and comparison with PCD/LSC/mPCD

Viscosity	0.1		0.01		0.005		0.001	
	FA	Opt	FA	Opt	FA	Opt	FA	Opt
16×16	9	9	11	11	21	13	35	20
32×32	9	9	11	11	17	14	31	23
64×64	8	8	11	11	14	14	29	25
128×128	8	7	11	11	14	13	26	26

Table: GMRES with modified AL preconditioner (cavity, Q2-Q1, stretched grids).

Grid	Uniform			Stretched		
	PCD	LSC	mPCD	PCD	LSC	mPCD
16×16	81	56	80	79	50	81
32×32	104	79	105	105	78	201
64×64	119	90	110	117	117	135
128×128	104	86	99	117	174	144

Table: GMRES with PCD, LSC and modified PCD preconditioners, steady Oseen problem (cavity, Q2-Q1 FEM, uniform and stretched grids), viscosity $\nu = 0.001$.

Comparison of exact and inexact solves. Scalability.

In the following Table we present a comparison of modified AL preconditioning with exact and inexact solves.

For the exact solves we use the sparse LU factorization with column AMD reordering available in Matlab.

For the inexact solves we use one iteration of the AMG method with ILU smoother in IFISS (Boyle, Mihajlovic and Scott, 2010).

We perform tests for both Picard and Newton linearizations of the lid driven cavity problem discretized with Q2-Q1 elements (Newton is harder), using the same value of γ from Fourier analysis in both cases. The viscosity is $\nu = 0.005$.

The experiments are performed in Matlab on a Sun Microsystems SunFire with four Dual Core AMD Opteron Processors and 32 GB of memory.

Using inexact solves does not affect the convergence rates, and leads to a **scalable** algorithm.

Comparison of exact and inexact solves. Scalability.

Grid	Picard		Newton	
	Exact	Inexact	Exact	Inexact
32 × 32	20	21	37	37
Setup time	0.15	3.63	0.15	3.61
Iter time	0.18	0.48	0.38	0.83
Total time	0.33	4.11	0.54	4.44
64 × 64	13	14	35	35
Setup time	1.93	11.88	1.95	12.28
Iter time	0.62	1.01	1.75	2.48
Total time	2.55	12.89	3.70	14.76
128 × 128	13	14	39	39
Setup time	34.90	50.70	34.34	51.41
Iter time	4.44	4.08	10.94	10.23
Total time	39.34	54.78	45.28	61.64
256 × 256	13	14	43	42
Setup time	856.74	222.71	673.29	213.71
Iter time	40.22	18.44	85.84	52.30
Total time	896.96	241.15	759.12	266.01

Table: GMRES iterations and timing, exact and inexact modified AL preconditioner.

Inexact solves for 3D case

3D lid driven cavity, P2-P1 FEM, uniform grids (code by Leo Rebholz).

Outer iteration: FGMRES with inexact modified AL-based preconditioner.

Inner iteration: AMG-preconditioned GMRES ($tol = 10^{-2}$).

Viscosity	0.1	0.05
$8 \times 8 \times 8$	25 (7.4)	32 (7.4)
$12 \times 12 \times 12$	21 (24.5)	29 (30.6)
$16 \times 16 \times 16$	20 (64.1)	28 (75.3)
$20 \times 20 \times 20$	20 (135.1)	27 (152.2)
$24 \times 24 \times 24$	19 (233.3)	27 (272.1)

Table: FGMRES iterations and timings with modified AL preconditioner (cavity, 3D, P2-P1, uniform grid)

Notice again the good scalability with respect to time.

Important: The AMG set-up can be reused over several Picard or Newton steps!

Outline

- 1 Introduction
- 2 Augmented Lagrangian-based preconditioner for stable finite elements
- 3 Generalization of AL-based preconditioners for stabilized finite elements**
- 4 Conclusions

Problem description

Using LBB unstable finite elements discretization, e.g., Q1-Q1, requires additional pressure stabilization terms $-C$ in the (2,2) block:

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

where $C = \delta \widehat{C}$.

- \widehat{C} is symmetric positive semidefinite.
- We use

$$\delta = \frac{\beta h^2}{\nu + h \|\mathbf{v}\|},$$

where h is the mesh size, \mathbf{v} is the wind function, and β depends on the choice of finite elements; $\beta = 1$ for Q1-Q1.

Augmented linear system

$$\begin{pmatrix} A & B^T \\ B & -C \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} f \\ 0 \end{pmatrix}. \quad (12)$$

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 augmented linear system is

$$\begin{pmatrix} A_\gamma & B_\gamma^T \\ B & -C \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} f \\ 0 \end{pmatrix}, \quad \text{or} \quad \widehat{A}x = b, \quad (13)$$

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

An AL-based preconditioner and choices of parameters

An ideal AL-based preconditioner is

$$\mathcal{P} = \begin{pmatrix} A_\gamma & 0 \\ B & \widehat{S} \end{pmatrix} = \begin{pmatrix} A + \gamma B^T W^{-1} B & 0 \\ B & \widehat{S} \end{pmatrix}.$$

In the paper by B., Olshanskii and Wang (IJNMF, 2010) we proposed

- $W = \widehat{M}_p = \text{diag}(M_p)$ to maintain the sparsity in A_γ .
- $\widehat{S} = -\gamma^{-1} \widehat{M}_p - C$.

The parameter γ is found by Fourier analysis (FA):

- Assume periodic b.c.'s and discretize the diffusion and frozen convection terms in A by centered differences.
- Discretize the divergence term in B by one-sided differences.
- $W = \widehat{M}_p$ and C both scale as h^2 .
- Express \mathcal{P} and \widehat{A} in terms of “Fourier eigenvalues”, and minimize the average distance of the eigenvalues of the preconditioned matrix $\mathcal{P}^{-1} \widehat{A}$ from 1.

Iteration counts

All 2D experiments are again done using IFISS.

Viscosity	0.1		0.01		0.005		0.001	
Grid	FA	Opt	FA	Opt	FA	Opt	FA	Opt
16×16	7	7	7	7	7	7	7	7
32×32	6	6	7	6	7	7	7	7
64×64	6	6	6	6	6	6	7	7
128×128	5	5	6	5	6	6	7	6

Table: GMRES iterations (lid driven cavity, Q1-Q1, stretched grid)

Observations:

- The iteration counts for GMRES with γ chosen by Fourier analysis are almost the same as those with optimal one.
- The iteration counts with both sets of γ do not depend on grid or viscosity.
- Similar behavior for the backward facing step problem.

Modified AL-based preconditioner

Since

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

Using the following block triangular approximation to A_γ :

$$\tilde{A}_\gamma = \begin{pmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{pmatrix},$$

we define

$$\tilde{\mathcal{P}} = \begin{pmatrix} \tilde{A}_\gamma & 0 \\ B & \hat{S} \end{pmatrix} = \begin{pmatrix} A_{11} & A_{12} & 0 \\ 0 & A_{22} & 0 \\ B_1 & B_2 & \hat{S} \end{pmatrix}.$$

This is the modified AL-based preconditioner introduced in B., Olshanskii and Wang (IJNMF, 2010).

Eigenvalue analysis

A straightforward calculation gives

$$\tilde{\mathcal{P}}^{-1}\hat{\mathcal{A}} = \begin{pmatrix} I_{n/2} - D & 0 & E \\ * & I_{n/2} & * \\ F & 0 & I_m - G \end{pmatrix}.$$

The eigenvalues of $\tilde{\mathcal{P}}^{-1}\hat{\mathcal{A}}$ are 1 of multiplicity $n/2$ and those of the matrix

$$\begin{pmatrix} I_{n/2} - D & E \\ F & I_m - G \end{pmatrix} = I_{n/2+m} - \begin{pmatrix} D & -E \\ -F & G \end{pmatrix}$$

Generally, the multiplicity of the eigenvalue $\lambda = 1$ is only $n/2$.

However, letting $\hat{S}^{-1} = -\gamma W^{-1}$, the matrix on the right-hand side is rank deficient by $n/2$, so $\tilde{\mathcal{P}}^{-1}\hat{\mathcal{A}}$ has the eigenvalue $\lambda = 1$ of multiplicity at least n .

Iteration counts for lid driven cavity on stretched grids

Viscosity	0.1		0.01		0.005		0.001	
Grid	FA	Opt	FA	Opt	FA	Opt	FA	Opt
16×16	9	9	14	13	32	15	82	24
32×32	9	8	12	11	22	14	43	24
64×64	9	8	12	11	14	14	35	25
128×128	8	7	10	10	13	13	26	25

Table: GMRES iterations with modified AL preconditioner (cavity, Q1-Q1, stretched grid, $\widehat{S}^{-1} = -\gamma \widehat{M}_p^{-1}$)

Observations:

- The iteration counts with γ chosen by Fourier analysis are close to those with optimal γ (for h small enough if ν is small).
- The convergence rates with both γ 's are essentially grid-independent and mildly dependent on ν .

Iteration counts for backward facing step problem

Viscosity	0.1		0.01		0.005	
Grid	FA	Opt	FA	Opt	FA	Opt
16×48	16	12	67	16	97	18
32×96	13	11	40	17	57	20
64×192	12	11	20	16	35	20
128×384	11	11	17	17	27	23

Table: Results for modified AL preconditioner (step, Q1-Q1, uniform grid, $\widehat{S}^{-1} = -\gamma \widehat{M}_p^{-1}$)

Observations:

- Similar performance to lid driven cavity.
- But for $\nu = 0.005$ with optimal γ , the iteration count is not quite h -independent due to the choice of \widehat{S} .

Iteration counts with $\widehat{S} = -\gamma^{-1}\widehat{M}_p - C$

Now we have h -independent convergence:

Viscosity	0.1		0.01		0.005	
Grid	FA	Opt	FA	Opt	FA	Opt
16×48	13	13	24	17	42	20
32×96	13	12	24	17	42	20
64×192	12	12	19	16	28	20
128×384	11	11	15	14	20	20

Table: GMRES iterations with modified AL preconditioner ($\widehat{S} = -\gamma^{-1}\widehat{M}_p - C$)

The additional costs incurred by the inclusion of C into \widehat{S} are very small.

Newton step on stretched grids: comparison with PCD/LSC/mPCD

Viscosity	0.1		0.01		0.005		0.001	
Grid	FA	Opt	FA	Opt	FA	Opt	FA	Opt
16×16	13	13	24	21	28	25	107	47
32×32	13	13	23	22	32	30	96	62
64×64	14	14	24	24	33	33	87	81
128×128	14	14	25	24	40	34	89	77

Table: GMRES with modified AL preconditioner (lid driven cavity, Newton, Q1-Q1, stretched grids).

	Uniform			Stretched		
Grid	PCD	LSC	mPCD	PCD	LSC	mPCD
64×64	82	111	232	248	246	> 250
128×128	80	168	> 250	> 250	> 250	> 250

Table: GMRES with PCD, LSC and modified PCD preconditioners (lid driven cavity, Newton, Q1-Q1 FEM, stretched grids), viscosity $\nu = 0.001$.

Comparison of exact and inexact solves. Scalability.

Grid	Picard		Newton	
	Exact	Inexact	Exact	Inexact
Timings				
Setup time	0.10	1.44	0.10	1.40
Iter time	0.16	0.77	0.34	1.49
Total time	0.26	2.21	0.44	2.89
64×64	13	13	34	34
Setup time	0.96	7.80	0.89	8.64
Iter time	0.52	2.57	1.33	7.08
Total time	1.49	10.37	2.22	15.72
128×128	13	13	41	41
Setup time	19.52	32.54	19.14	37.53
Iter time	3.24	11.10	10.88	33.00
Total time	22.76	43.64	30.02	70.53
256×256	12	13	42	42
Setup time	158.62	121.10	158.97	128.82
Iter time	15.93	37.98	63.55	121.36
Total time	174.55	159.08	222.52	250.18

Outline

- 1 Introduction
- 2 Augmented Lagrangian-based preconditioner for stable finite elements
- 3 Generalization of AL-based preconditioners for stabilized finite elements
- 4 Conclusions**

Conclusions and future work

- The numerical solution of the incompressible Navier–Stokes equations remains an important source of problems with many challenges
- The classical Augmented Lagrangian (AL) approach has been around for over 30 years, but exact solves are not feasible for large problems
- Suitably modified and combined with an AMG-type inner solver, the AL approach yields a robust and scalable preconditioner for the Oseen problem
- Both stable and stabilized elements can be accommodated
- Stretched grids do not pose any serious difficulties
- Preconditioner works well also for Newton linearization
- The Fourier-based approach gives very good estimates of the parameter γ
- Compares favorably with state-of-the-art methods
- More work needed for 3D problems, parallel implementations