Modified Augmented Lagrangian Methods for Incompressible Flow Problems

London Mathematical Society Durham Symposium
Computational Linear Algebra for Partial Differential Equations
Durham, UK
16 July 2008

Michele Benzi
Department of Mathematics and CS
Emory University
Atlanta, GA, USA

Acknowledgements

Thanks to:

- Maxim Olshanskii (Moscow State University)
- Alessandro Veneziani (Emory University)
- NSF

Outline

- Problem formulation
- The augmented Lagrangian (AL) approach
- Block triangular preconditioning
- Numerical results, I
- Modified AL preconditioner
- Numerical results, II
- Conclusions and future work

Some papers available at

http://www.mathcs.emory.edu/∼benzi

Problem Formulation

We are concerned with the solution of large linear systems of the form

$$\begin{bmatrix} A & B^T \\ B & O \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathbf{0} \end{bmatrix} \tag{1}$$

where

- A is $n \times n$, B is $m \times n$, $m < n$;
- A represents reactive, diffusive and advective terms;
- B represents the discrete divergence, B^T the gradient;

We further assume that system (1) is nonsingular.

Systems of this kind arise from discretizations of various linearizations of the incompressible Navier–Stokes equations (Oseen, Newton).

The Augmented Lagrangian Formulation

In the Oseen problem, $A \approx \sigma - \nu\Delta + \mathbf{w} \cdot \nabla$ (where $\sigma = 0$ for steady problems).

Typically, many such linear systems have to be solved in the course of a simulation. Especially in 3D, iterative methods are virtually mandatory. Efficient preconditioning techniques are essential for fast convergence.

Notice that the original system is equivalent to

$$\begin{bmatrix} A + \gamma B^T W^{-1} B & B^T \\ B & O \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathbf{0} \end{bmatrix}. \tag{2}$$

Here $\gamma > 0$ is a parameter and W an arbitrary invertible matrix. This is called an Augmented Lagrangian formulation (Fortin & Glowinski, 1982).

Block Triangular Preconditioner for AL System

We are interested in preconditioners for the AL system (2).

Let $\mathcal{A}$ be the coefficient matrix of the AL system, and let

$$\mathcal{P} = \begin{bmatrix} A + \gamma B^T W^{-1} B & B^T \\ O & -\frac{1}{\gamma} W \end{bmatrix}.$$

For LBB-stable discretizations, we have the following result:

Theorem (B. & Olshanskii, 2006): Assume $W = M_p$ (the pressure mass matrix). Then for all $\gamma > 0$ the preconditioned matrix $\mathcal{P}^{-1}\mathcal{A}$ has the eigenvalue 1 of multiplicity n; the remaining m eigenvalues λ_i lie inside a rectangle $\mathcal{R}$ in the right half-plane which does not depend on the mesh size h. Furthermore, γ can be chosen so that $\mathcal{R}$ does not depend on ν, and all the eigenvalues tend to 1 for $\gamma \to \infty$.

Block Triangular Preconditioner for AL System

In the case of Galerkin FE methods (with no stabilization), a result by Elman and Silvester implies that it is sufficient to set $\gamma = O(\nu^{-1})$ to ensure that all non-unit eigenvalues of $\mathcal{P}^{-1}\mathcal{A}$ are contained in a box $[a,\, b] \times [-c,\, c]$, $a > 0$, with a, b, c independent of both ν and h.

However, a very large value of γ makes the solution of the (1,1) block more difficult, and it is better to use moderate or small values of γ in the computations.

Indeed, for very large γ the highly singular term $\gamma B^T W^{-1} B$ will 'swamp' the other terms in the (1,1) block.

Block Triangular Preconditioner for AL System

Motivated by the previous result, we look for a practically feasible block triangular preconditioner for the AL system, of the form

$$\mathcal{P} = \begin{bmatrix} \widehat{A}_\gamma & B^T \\ O & \widehat{S} \end{bmatrix}$$

where $\widehat{A}_\gamma \approx A + \gamma B^T W^{-1} B$ and $\widehat{S} \approx -\frac{1}{\gamma} W$.

Clearly, the critical component is $\widehat{A}_\gamma$. We implement the action of $\widehat{A}_\gamma^{-1}$ by a single iteration (W-cycle) of a suitable multigrid method.

In the FEM context it is natural to take $W = M_p$ (pressure mass matrix) or, in practice, a diagonal approximation $\bar{M}_p$ of it, and to set $\widehat{S}^{-1} := -\nu M_p^{-1} - \gamma \bar{M}_p^{-1}$.

Block Triangular Preconditioner for AL System

The matrix $A + \gamma B^T \bar{M}_p^{-1} B$ can be regarded as a FEM discretization of a (non-standard) differential operator of the form "linear elasticity + convection".

For large values of ν, this is essentially the operator arising from the linear elasticity equations. Note that γ can be tuned to provide the appropriate scaling.

We have combined a multigrid method for linear elasticity problems due to Schöberl (Numer. Math., 1999) with an appropriate block smoother with overlapping blocks.

See paper by M. B. and M. Olshanskii (SISC, 2006) for details.

Numerical Experiments, I

We experimented with the AL-based preconditioner on a few steady problems on the unit square $\Omega = [0,1] \times [0,1]$:

- A constant wind problem (Oseen)
- A recirculating flow (vortex) problem (Oseen)
- A lid-driven cavity problem with different winds (Oseen)
- Two indefinite problems ($\sigma = -\alpha < 0$) from linear stability analysis (Newton)

We implemented two uniform FEM discretizations (isoP2-P0 and isoP2-P1) and an unstructured one (P2-P1). These discretizations satisfy the inf-sup condition: no pressure stabilization is needed.

SUPG stabilization is used for the velocity.

Numerical Experiments, I

The Krylov subspace methods used are Bi-CGStab and GMRES, except for the indefinite problem where we used FGMRES.

This is because for the indefinite problem we used a non-stationary multigrid method (similar to the one proposed by Elman, Ernst and O'Leary for the Helmholtz equation) to approximately solve linear systems involving the (1,1) block. Hence, the preconditioner varies from one iteration to the next, necessitating a flexible outer iteration.

The cost of each iteration is linear in the number of unknowns.

Numerical Experiments, I

Results for AL approach, isoP2-P0 FEM.

mesh size h	viscosity ν				
	1	0.1	0.01	10^{-3}	10^{-4}
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
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

Number of preconditioned Bi-CGStab iterations ($\widehat{A}_\gamma^{-1}$ is one W(1,1)-cycle, $\gamma = 1$).

Numerical Experiments, I

Results for AL approach, isoP2-P1 FEM.

mesh size h	viscosity ν				
	1.	0.1	0.01	10^{-3}	10^{-4}
	parameter γ				
	1.	1.	1.	0.1	0.02
constant wind					
1/16	6	6	7	8	24
1/32	7	6	10	8	22
1/64	7	6	8	7	19
1/128	7	6	9	9	18
rotating vortex					
1/16	6	6	7	13	25
1/32	5	6	9	11	32
1/64	4	5	10	11	37
1/128	4	4	10	12	34

Number of preconditioned Bi-CGStab iterations ($\widehat{A}_\gamma^{-1}$ is one W(1,1)-cycle).

Numerical Experiments, I

# elem	$h \downarrow \backslash Re \rightarrow$	1	10	100	1000	10000
656	1/8	13/14	13/15	13/21	13/24	13/28
2596	1/16	13/14	13/15	13/20	13/23	13/26
10480	1/32	13/14	13/15	13/20	13/22	12/24
41852	1/64	13/14	13/15	13/20	13/22	12/22

Lid-driven cavity problem, unstructured mesh.

Test case with "rotating-vortex" wind: number of outer GMRES iterations for different values of the space discretization and of the Reynolds number. The first number is for $\gamma = 1$. The second one is for $\gamma = 0.1$.

Numerical Experiments, I

FGMRES its/timings for 2D indefinite problems, Newton linearization about U, inexact solves via a single multigrid V(1,1)-cycle, isoP2-P0 FEM.

Note: Here $\gamma = \alpha = 1$; the problems become more indefinite as $Re \to \infty$.

h	Reynolds number $Re = \nu^{-1}$			
	1	10	100	1000
U=Poiseuille flow				
1/256	13 (57s)	13 (57s)	16 (71s)	31 (140s)
1/512	13 (268s)	13 (269s)	16 (339s)	26 (545s)
U=rotating vortex				
1/256	13 (56s)	12 (53s)	18 (79s)	45 (203s)
1/512	13 (264s)	12 (242s)	18 (370s)	46 (976s)

Note the near-perfect scaling with respect to CPU time.

Numerical Experiments, I

Same as before, but now using isoP2-P1 FEM.

Note: Here $\alpha = 1$; the problems become more indefinite as $Re \to \infty$. We use $\gamma = 1$ in all cases except for $Re = 1000$, where we use $\gamma = 0.1$.

h	Reynolds number $Re = \nu^{-1}$			
	1	10	100	1000
U=Poiseuille flow				
1/256	13 (59s)	13 (59s)	16 (92s)	31 (148s)
1/512	13 (271s)	13 (254s)	16 (444s)	26 (554s)
U=rotating vortex				
1/256	13 (58s)	12 (59s)	18 (102s)	45 (221s)
1/512	13 (273s)	12 (253s)	18 (458s)	46 (995s)

Again, the scaling with respect to CPU time is excellent.

Modified AL Preconditioner

The previous test cases indicate that the AL-based preconditioner is effective and robust.

Moreover, experimental comparisons show that this approach is superior to block preconditioners based on approximations of the pressure Schur complement, especially for small ν.

The main issue is the approximate solution of linear systems involving the (1,1) block, i.e., the block matrix

$$A_\gamma = \begin{bmatrix} A_1 + \gamma B_1^T W^{-1} B_1 & \gamma B_1^T W^{-1} B_2 \\ \gamma B_2^T W^{-1} B_1 & A_2 + \gamma B_2^T W^{-1} B_2 \end{bmatrix}. \qquad (3)$$

The challenge is to make this approach viable for general geometries and discretizations.

Modified AL Preconditioner

Here we consider a simple modification of the AL-based preconditioner that results in subproblems of simpler form. For simplicity, we consider the 2D case only.

Dropping the (2,1) sub-block in A_γ results in a block triangular matrix of the form

$$\tilde{A}_\gamma = \begin{bmatrix} A_1 + \gamma B_1^T W^{-1} B_1 & \gamma B_1^T W^{-1} B_2 \\ O & A_2 + \gamma B_2^T W^{-1} B_2 \end{bmatrix}. \qquad (4)$$

Applying the preconditioner with the approximation given by (4) only requires approximately solving linear systems with coefficient matrices of the form $A_i + \gamma B_i^T W^{-1} B_i$; these matrices are discrete analogues of scalar anisotropic convection-diffusion (or convection-diffusion-reaction) operators, with anisotropy ratio $\approx 1 + \gamma/\nu$.

Modified AL Preconditioner

For the generalized Stokes case ($Re = 0$), the continuous counterpart of the previous block triangular matrix is the operator matrix

$$\begin{bmatrix} \sigma - \nu\Delta - \gamma\partial_{xx}^2 & -\gamma\partial_{xy}^2 \\ O & \sigma - \nu\Delta - \gamma\partial_{yy}^2 \end{bmatrix}. \tag{5}$$

Note that the diagonal 'blocks' are

$$\sigma - \nu\Delta - \gamma\partial_{xx}^2 = \sigma - \nu\left(\left(1 + \frac{\gamma}{\nu}\right)\partial_{xx}^2 + \partial_{yy}^2\right)$$

and

$$\sigma - \nu\Delta - \gamma\partial_{yy}^2 = \sigma - \nu\left(\partial_{xx}^2 + \left(1 + \frac{\gamma}{\nu}\right)\partial_{yy}^2\right).$$

The dropped term is $\gamma B_2^T W^{-1} B_1 \approx -\gamma\partial_{yx}^2$.

All this suggests taking γ small.

Numerical Experiments, II

$h \downarrow \backslash Re \rightarrow$	10	20	100	200	1000
1/32	16/20	16/20	17/21	17/24	17/26
1/64	19/23	20/24	21/25	21/26	21/30
1/128	24/28	25/28	26/29	26/30	26/36

Unsteady Oseen problem, $\sigma = 1/\Delta t = O(h^{-1})$, MAC spatial discretization, original vs. modified AL preconditioner, $\gamma = 0.1$. The first number is the number of GMRES iterations with the original AL approach. The second one refers to the modified AL approach.

In terms of timings, the modified AL method is about 50% faster on average. It also requires less storage.

Numerical Experiments, II

$h \downarrow \backslash Re \rightarrow$	10	20	100	200	1000
1/32	14/22	16/27	19/36	29/46	32/111
1/64	14/22	16/26	18/35	19/45	31/104
1/128	14/22	15/25	18/35	19/45	30/102

Same as previous slide, but now steady case ($\sigma = 0$). Original vs. modified AL scheme, GMRES iterations, $\gamma = 0.1$ except for $Re = 1000$, where $\gamma = 0.05$.

In terms of timings, the modified AL method is always faster (between 20% and 40%) for the finest grid, except for $Re = 1000$.

Note the degradation in the performance of the modified AL preconditioner for large Re.

Conclusions and Future Work

- AL-based block preconditioner is effective and robust
- Optimal behavior with respect to mesh size observed in all our experiments; very mild dependency on ν
- Modified AL preconditioner is much cheaper/easier to implement
- Tests indicate good performance for unsteady Oseen
- Performance not as good for very small ν for steady Oseen
- Spectral analysis for modified AL is being worked out
- Next: 3D unstructured implementation

References

1. M. Benzi and M. A. Olshanskii, *An augmented Lagrangian-based approach to the Oseen problem*, SIAM J. Sci. Comput., 28(6), pp. 2095–2113, 2006.

2. M. Benzi, G. H. Golub and J. Liesen, *Numerical solution of saddle point problems*, Acta Numerica, 14 (2005), pp. 1–137.

3. M. A. Olshanskii and M. Benzi, *An augmented Lagrangian approach to linearized problems in hydrodynamic stability*, SIAM J. Sci. Comput., 30(3), pp. 1959–1973, 2008.

4. H. C. Elman, D. Silvester, and A. Wathen, *Finite Elements and Fast Iterative Solvers, with Applications in Incompressible Fluid Dynamics*, Oxford Science Publications, 2005.