Iterative Solution of Double Saddle-Point Problems

Michele Benzi
Scuola Normale Superiore, Pisa

Iterative Solution of Large-Scale Saddle-Point Problems
Cortona, 9–20 May, 2022

Outline

- Double saddle point problems
- Existence and uniqueness conditions
- Two examples
- Block preconditioners
- Numerical experiments
- Conclusions

Acknowledgements

- Joint work with Fatemeh Beik
- Special thanks to Alison Ramage and Miroslav Tůma

F. A. P. Beik and M. Benzi, *Iterative methods for double saddle point systems*, SIAM J. Matrix Anal. Appl., 39 (2018), pp. 902–921.

F. A. P. Beik and M. Benzi, *Block preconditioners for saddle point systems arising from liquid crystal directors modeling,* Calcolo, 55 (2018), art. 29 .

Double saddle point problems

Consider the linear system

$$\mathcal{A}u \equiv \begin{bmatrix} A & B^T & C^T \\ B & 0 & 0 \\ C & 0 & -D \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix} \equiv b, \tag{1}$$

where $A \in \mathbb{R}^{n \times n}$ is symmetric positive definite (SPD), $B \in \mathbb{R}^{m \times n}$, $C \in \mathbb{R}^{p \times n}$, and $D \in \mathbb{R}^{p \times p}$ is symmetric positive semidefinite (SPS) and possibly zero. Throughout the paper we assume that $n \geq m + p$.

Note that $\mathcal{A}$ is symmetric indefinite.

Linear systems of the form (1) arise in numerous applications. The term double saddle point systems is sometimes used to descibe (1).

A. Ramage and E. C. Gartland, Jr., *A preconditioned null space method for liquid crystal director modeling*, SIAM J. Sci. Comput., 35 (2013).

Double saddle point problems (cont.)

Applications leading to double saddle point problems include:

- Mixed-Hybrid FEM for second order elliptic PDEs;
- Finite element discretization of liquid crystal director modeling;
- Numerical modeling of non-Newtonian fluids;
- Linear elasticity;
- Coupled systems of PDEs (for example, Stokes-Darcy);
- Interior point methods in constrained optimization.

Note that the block matrices

$$\mathcal{B} = \begin{bmatrix} A & C^T & B^T \\ C & -D & 0 \\ B & 0 & 0 \end{bmatrix} \quad \text{and} \quad \mathcal{C} = \begin{bmatrix} -D & C & 0 \\ C^T & A & B^T \\ 0 & B & 0 \end{bmatrix};$$

can be brought into the same form as matrix $\mathcal{A}$ by means of symmetric permutations.

Double saddle point problems (cont.)

Matrix $\mathcal{A}$ can be regarded as a 2×2 block matrix in two different ways, according to which of the following partitioning strategies is used:

$$\mathcal{A} = \left[\begin{array}{c|cc} A & B^T & C^T \\ \hline B & 0 & 0 \\ C & 0 & -D \end{array}\right] \quad \text{or} \quad \mathcal{A} = \left[\begin{array}{cc|c} A & B^T & C^T \\ B & 0 & 0 \\ \hline C & 0 & -D \end{array}\right].$$

The first partitioning highlights the fact that problem (1) can in principle be treated as a "standard" saddle point problem, possibly stabilized (or regularized) when $D \neq 0$.

On the other hand, the second partitioning highlights the double saddle point structure of $\mathcal{A}$, since the (1,1) block is itself the coefficient matrix of a saddle point problem.

Our main goal is to develop efficient solvers that exploit the 3×3 block structure of $\mathcal{A}$.

Double saddle point problems (cont.)

Some additional references:

D. Boffi, F. Brezzi, and M. Fortin, *Mixed Finite Element Methods and Applications*, Springer, 2013.

A. Aposporidis, E. Haber, M. Olshanskii, and A. Veneziani, *A mixed formulation of the Bingham fluid flow problem: Analysis and numerical solution*, Computer Meth. Appl. Mech. Engrg., 200 (2011).

P. Chidyagway, S. Ladenheim, and D. B. Szyld, *Constraint preconditioning for the coupled Stokes–Darcy system*, SIAM J. Sci. Comput., 38 (2016).

B. Morini, V. Simoncini, and M. Tani, *Spectral estimates for unreduced symmetric KKT systems arising from Interior Point methods*, Numer. Linear Algebra Appl., 23 (2016).

J. Maryška, M. Rozložník, and M. Tůma, *Schur complement systems in the mixed-hybrid finite element approximation of the potential fluid flow problem*, SIAM J. Sci. Comput., 22 (2000).

Conditions for a unique solution

Proposition 1. Assume that $A \succ 0$ and $D \succ 0$. Then $\mathcal{A}$ is nonsingular if and only if B^T has full column rank.

Proposition 2. Let $A \succeq 0$ and $D \succeq 0$, with $D \neq 0$. Assume that at least one of A and D is positive definite and that B^T has full column rank. The following statements hold:

Case 1. Suppose that $A \succ 0$ and $D \not\succ 0$.

- If $\ker(C^T) \cap \ker(D) = \{0\}$ and $\mathrm{range}(B^T) \cap \mathrm{range}(C^T) = \{0\}$, then $\mathcal{A}$ is nonsingular.
- If $\mathcal{A}$ is nonsingular then $\ker(C^T) \cap \ker(D) = \{0\}$.

Case 2. Suppose that $A \not\succ 0$ and $D \succ 0$.

- If $\ker(A) \cap \ker(B) \cap \ker(C) = \{0\}$ and $\mathrm{range}(B^T) \cap \mathrm{range}(C^T) = \{0\}$, then $\mathcal{A}$ is nonsingular.
- If $\mathcal{A}$ is nonsingular then $\ker(A) \cap \ker(B) \cap \ker(C) = \{0\}$.

Conditions for a unique solution (cont.)

Proposition 3. Let $A \succ 0$ and assume that B^T and C^T have full column rank. Consider the linear system (1) with $D = 0$. Then

$$\text{range}(B^T) \cap \text{range}(C^T) = \{0\}$$

is a necessary and sufficient condition for the coefficient matrix $\mathcal{A}$ to be nonsingular.

Remark. It should be stressed that in the case $D = 0$, both B^T and C^T must have full column rank for $\mathcal{A}$ to be invertible. In contrast, in the case that $D \succeq 0$ and $D \neq 0$, only the matrix B^T is required to have full column rank while the matrix C^T can be rank deficient.

Example 1: The potential fluid flow problem

Our first example is a low-order Raviart–Thomas mixed-hybrid finite element approximation of Darcy's law and continuity equation describing the 3D potential fluid flow problem in porous media:

$$k(\mathbf{x})\mathbf{u} = -\nabla p, \quad \nabla \cdot \mathbf{u} = q \qquad \text{in } \Omega \subset \mathbb{R}^3,$$

where $\mathbf{u}$ is the fluid velocity, p is the piezometric potential (fluid pressure), $k(\mathbf{x})$ is the symmetric and uniformly positive definite second-rank tensor of the hydraulic resistance of the medium with $[k(\mathbf{x})]_{ij} \in L^\infty(\Omega)$ for $i, j = 1, 2, 3$, and q represents the density of potential sources in the medium.

The boundary conditions are given by

$$p = p_D \quad \text{on} \quad \partial\Omega_D, \qquad \mathbf{u} \cdot \mathbf{n} = \mathbf{u}_N \quad \text{on} \quad \partial\Omega_N,$$

where $\partial\Omega = \overline{\partial\Omega_D} \cup \overline{\partial\Omega_N}$ with $\partial\Omega_D \neq \emptyset$, $\partial\Omega_D \cap \partial\Omega_N = \emptyset$, and $\mathbf{n}$ is the outward normal vector defined (a.e.) on $\partial\Omega$.

Example 1: The potential fluid flow problem (cont.)

The discrete system is of the form

$$\begin{bmatrix} A & B^T & C^T \\ B & 0 & 0 \\ C & 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}.$$

The solution vectors x and y correspond to velocity and pressure degrees of freedom (respectively), while z is a vector of Lagrange multipliers. For this problem we have that $A \succ 0$ and B^T, C^T have full column rank, and that $\mathcal{A}$ is nonsingular.

Moreover, the A matrix is block diagonal, with 5×5 blocks.

As a consequence, both matrices $BA^{-1}B^T$ and $CA^{-1}C^T$ retain a great deal of sparsity and can be formed explicitly.

See J. Maryška, M. Rozložník, and M. Tůma, *Schur complement systems in the mixed-hybrid finite element approximation of the potential fluid flow problem*, SIAM J. Sci. Comput., 22 (2000).

Example 1: The potential fluid flow problem (cont.)

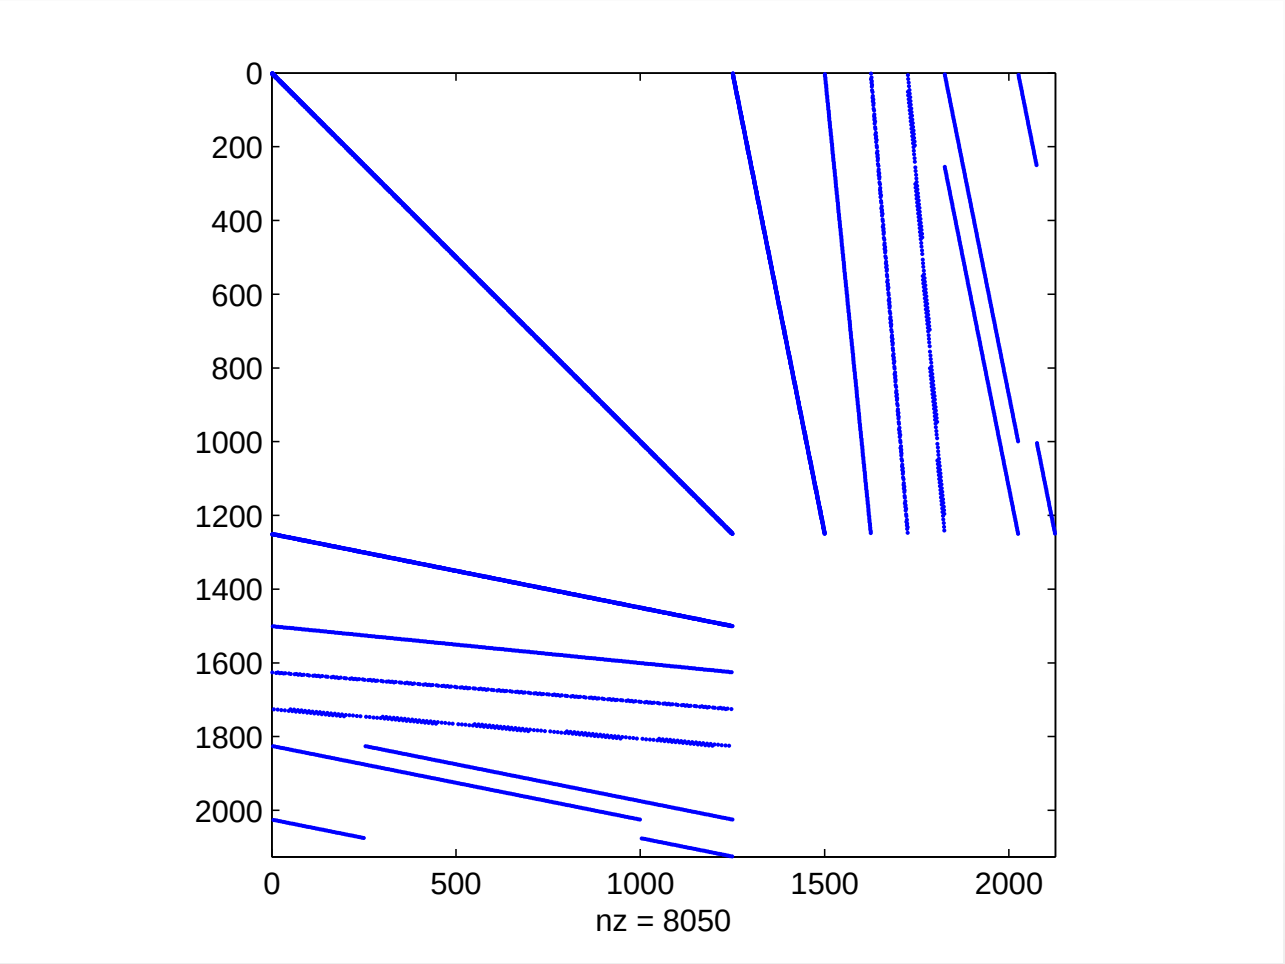

Sparsity pattern for potential fluid flow problem on the unit cube.

Example 2: The liquid crystal director model ($D \neq 0$)

Mathematical models for understanding the orientational properties of liquid crystals lead to the minimization of the free energy functional

$$\mathcal{F}[u, v, w, U] = \frac{1}{2} \int_0^1 \left[(u_z^2 + v_z^2 + w_z^2) - \alpha^2 (\beta + w^2) U_z^2 \right] dz \qquad (2)$$

where u, v, w and U are functions of the space variable $z \in [0, 1]$ that are required to satisfy suitable end-point conditions, $u_z = \frac{du}{dz}$ (etc.), and α, β are prescribed (positive) parameters. The problem is discretized by means of a uniform piecewise-linear finite element method.

With a mesh-size $h = \frac{1}{k+1}$ (with $k + 1$ the number of cells), using nodal quadrature and the prescribed boundary conditions leads to replacing the functional $\mathcal{F}$ with a function f of $4k$ variables:

$$\mathcal{F}[u, v, w, U] \approx f(u_1, \ldots, u_k, v_1, \ldots, v_k, w_1, \ldots, w_k, U_1, \ldots, U_k).$$

Example 2: The liquid crystal director model (cont.)

The free energy functional (2) must be minimized subject to the unit vector constraint. For the discrete (finite-dimensional) problem, this constraint is enforced by imposing that the solution components u_j, v_j and w_j are such that

$$u_j^2 + v_j^2 + w_j^2 = 1, \quad j = 1, \ldots, k.$$

Introducing Lagrange multipliers $\lambda_1, \ldots, \lambda_k$, the problem reduces to finding the critical points of the Lagrangian function

$$L = f + \frac{1}{2} \sum_{j=1}^{k} \lambda_j (u_j^2 + v_j^2 + w_j^2 - 1).$$

Setting the partial derivatives of L to zero results in the nonlinear system of $5k$ equations in as many unknowns $\nabla L(\mathbf{x}) = \mathbf{0}$, where the unknown vector $\mathbf{x} \in \mathbb{R}^{5k}$ contains the values (u_j, v_j, w_j) with $1 \leq j \leq k$, λ_j with $1 \leq j \leq k$, and U_j with $1 \leq j \leq k$ (in this order).

This nonlinear system is solved by Newton's method.

Example 2: The liquid crystal director model (cont.)

Newton's method leads to a linear system of the form

$$\nabla^2 L(\mathbf{x}^{(\ell)})\, \delta\mathbf{x}^{(\ell)} = -\nabla L(\mathbf{x}^{(\ell)}) \tag{3}$$

at each step ℓ, where $\nabla^2 L(\mathbf{x}^{(\ell)})$ denotes the Hessian of L evaluated at the current approximate solution $\mathbf{x}^{(\ell)}$.

As shown by Ramage and Gartland (SISC 2013), the Hessian has the following block 3×3 structure:

$$\nabla^2 L = \begin{bmatrix} A & B^T & C^T \\ B & 0 & 0 \\ C & 0 & -D \end{bmatrix} = \mathcal{A},$$

where A is $n \times n$, B is $m \times n$, C is $p \times n$ and D is $p \times p$ with $n = 3k$ and $m = p = k$. Hence, it is necessary to solve a system of the form (1) at each Newton step.

All the nonzero blocks in $\mathcal{A}$ are sparse and structured.

Example 2: The liquid crystal director model (cont.)

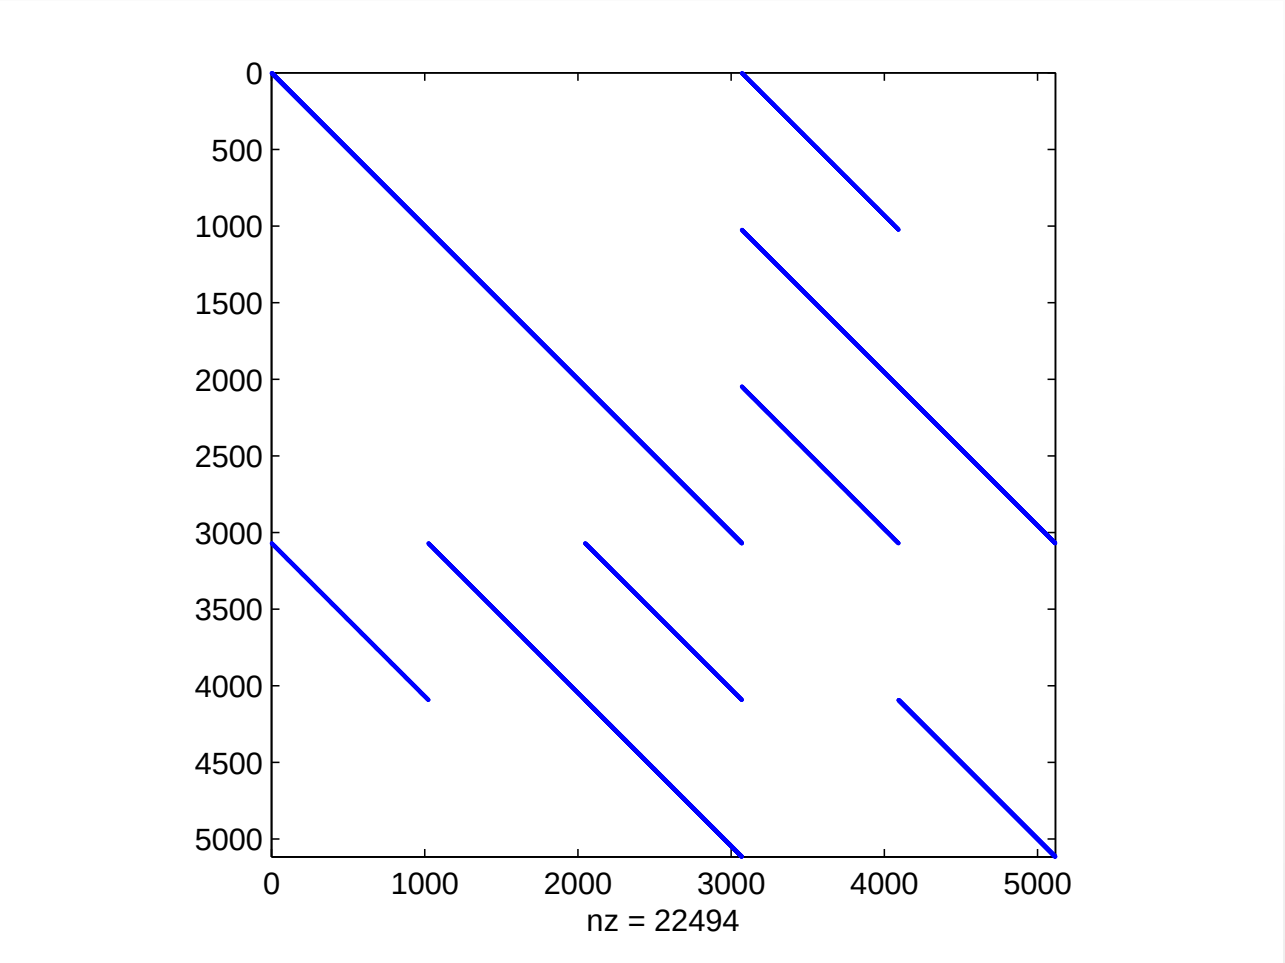

Sparsity pattern for discretized liquid crystal director model.

Block preconditioners for the case $D = 0$

Block diagonal and block triangular preconditioners have been extensively studied in the saddle point literature; see, e.g., Elman, Silvester, and Wathen (2005, 2014); B., Golub, and Liesen (2005).

These 'classical' preconditioners suggest different extensions to the 3×3 block case, depending on the partitioning used for $\mathcal{A}$. It is unclear, *a priori*, which of these several possible preconditioners is better for a given problem.

We first investigated the following block diagonal and block triangular preconditioners for solving systems of the form (1) with $D = 0$:

$$\mathcal{P}_{D} = \begin{bmatrix} A & 0 & 0 \\ 0 & BA^{-1}B^T & 0 \\ 0 & 0 & CA^{-1}C^T \end{bmatrix}, \; \mathcal{P}_{T} = \begin{bmatrix} A & B^T & C^T \\ 0 & -BA^{-1}B^T & 0 \\ 0 & 0 & -CA^{-1}C^T \end{bmatrix}.$$

In the sequel, we let $S_B := BA^{-1}B^T$ and $S_C = CA^{-1}C^T$.

Block preconditioners for the case $D = 0$ (cont.)

Theorem 1. Suppose that $A \succ 0$, B^T and C^T have full column rank, and that $D = 0$ in (1). Then

$$\sigma(\mathcal{P}_D^{-1}\mathcal{A}) \subset \left(-1, \frac{1-\sqrt{1+4\gamma_*}}{2}\right) \cup \{1\} \cup \left(\frac{1+\sqrt{1+4\gamma_*}}{2}, 2\right),$$

with

$$\gamma_* = \min \frac{x^T(B^T S_B^{-1} B + C^T S_C^{-1} C)x}{x^T A x} > 0,$$

where the minimum is taken over all $x \in \mathbb{R}^n$, $x \notin \ker(B) \cap \ker(C)$, such that $(x; y; z)$ is an eigenvector of $\mathcal{P}_D^{-1}\mathcal{A}$.

The set $\{1\} \cup \left(\frac{1+\sqrt{1+4\gamma_*}}{2}, 2\right)$ contains n eigenvalues and the negative interval $\left(-1, \frac{1-\sqrt{1+4\gamma_*}}{2}\right)$ contains $m + p$ eigenvalues. Furthermore, if $\lambda \neq 1$ is an eigenvalue of $\mathcal{P}_D^{-1}\mathcal{A}$, then $1 - \lambda$ is also an eigenvalue (though not necessarily with the same multiplicity).

Block preconditioners for the case $D = 0$ (cont.)

The previous result shows that the positive eigenvalues of $\mathcal{P}_D^{-1}\mathcal{A}$ are nicely bounded between 1 and 2.

The negative eigenvalues, on the other hand, lie between -1 and

$$\frac{1 - \sqrt{1 + 4\gamma_*}}{2},$$

and thus the upper bound approaches 0 (from the left) if

$$\gamma_* = \min \frac{x^T(B^T S_B^{-1} B + C^T S_C^{-1} C)x}{x^T A x} \to 0+,$$

which may well happen in practice.

For instance, for the potential fluid flow problem we found that the smallest negative eigenvalue (and thus γ_*) approaches 0 as $h \to 0$, and the performance of preconditioned MINRES deteriorates as the mesh is refined.

Block preconditioners for the case $D = 0$ (cont.)

For the block triangular preconditioner $\mathcal{P}_{_T}$ we have the following result.

Theorem 2. Suppose that $A \succ 0$, B^T and C^T have full column rank, and that $D = 0$. Then $\sigma(\mathcal{P}_{_T}^{-1}\mathcal{A}) \subset (0, 2)$, with $\lambda = 1$ being an eigenvalue of multiplicity at least n. Moreover, the spectrum of $\mathcal{P}_{_T}^{-1}\mathcal{A}$ is symmetric with respect to $\lambda = 1$; i.e., if $\lambda_1 \neq 1$ and $\lambda_2 \neq 1$ are two eigenvalues of $\mathcal{P}_{_T}^{-1}\mathcal{A}$, then $\lambda_1 + \lambda_2 = 2$.

Note: The proof of this result shows that the preconditioned matrix has eigenvalues of the form

$$\lambda = \frac{x^T(B^T S_B^{-1} B + C^T S_C^{-1} C)x}{x^T A x},$$

which in some cases may approach zero as the mesh is refined (as in the case of the potential fluid flow problem).

Block preconditioners for the case $D = 0$ (cont.)

Other “natural" (ideal) preconditioners, corresponding to the first of the two possible block partitionings of $\mathcal{A}$, include

$$\mathcal{P}_{GD} = \left[\begin{array}{c|cc} A & 0 & 0 \\ \hline 0 & BA^{-1}B^T & BA^{-1}C^T \\ 0 & CA^{-1}B^T & CA^{-1}C^T \end{array}\right]$$

and

$$\mathcal{P}_{GT,1} = \left[\begin{array}{c|cc} A & 0 & 0 \\ \hline B & -BA^{-1}B^T & -BA^{-1}C^T \\ C & -CA^{-1}B^T & -CA^{-1}C^T \end{array}\right].$$

Note: In the potential fluid flow problem, all the blocks are sparse and can be formed explicitly if needed.

Block preconditioners for the case $D = 0$ (cont.)

Using instead the second possible block partitioning of $\mathcal{A}$ we obtain the following block triangular preconditioner:

$$\mathcal{P}_{GT,2} = \left[\begin{array}{cc|c} A & B^T & 0 \\ B & 0 & 0 \\ \hline C & 0 & -S \end{array}\right],$$

where

$$S = C(A^{-1} + A^{-1}B^T S_B^{-1} B A^{-1})C^T.$$

Note: The standard Murphy–Golub–Wathen results apply to these three ideal preconditioners.

Block preconditioners for the case $D \neq 0$

We focus on the following two block triangular preconditioners:

$$\tilde{\mathcal{P}}_{_T} = \begin{bmatrix} A & B^T & C^T \\ 0 & -BA^{-1}B^T & 0 \\ 0 & 0 & -(D + CA^{-1}C^T) \end{bmatrix},$$

and

$$\hat{\mathcal{P}}_{_T} = \begin{bmatrix} A & B^T & C^T \\ 0 & -BA^{-1}B^T & -BA^{-1}C^T \\ 0 & 0 & -(D + CA^{-1}C^T) \end{bmatrix}.$$

Note: These preconditioners wil be nonsingular if $A \succ 0$, B^T has full column rank, $D \succeq 0$ and $\ker(D) \cap \ker(C^T) = \{0\}$. We recall that these conditions also guarantee the invertibility of $\mathcal{A}$.

Block preconditioners for the case $D \neq 0$ (cont.)

Theorem 3. Assume that $A \succ 0$, B^T has full column rank, $D \succeq 0$ and $\ker(D) \cap \ker(C^T) = \{0\}$. Then

$$\sigma(\tilde{\mathcal{P}}_T^{-1}\mathcal{A}) \subset (0, 1-\sqrt{\xi}] \cup \{1\} \cup [1+\sqrt{\xi}, 2) \subset (0,2),$$

where

$$\xi = \frac{\bar{\lambda}_{\min}(CA^{-1}C^T)}{\lambda_{\max}(D) + \bar{\lambda}_{\min}(CA^{-1}C^T)}$$

(here $\bar{\lambda}_{\min}$ is used to denote the smallest nonzero eigenvalue).
Moreover, $\lambda = 1$ is an eigenvalue of multiplicity at least n.

Note: This result implies that if $\lambda_{\max}(D) \approx 0$, then the preconditioned matrix may have tiny eigenvalues.

Block preconditioners for the case $D \neq 0$ (cont.)

A much stronger result holds for $\hat{\mathcal{P}}_T$ if we require $D \succ 0$:

Theorem 4. Assume that $A \succ 0$, B^T has full column rank, and $D \succ 0$. Then $\sigma(\hat{\mathcal{P}}_T^{-1}\mathcal{A}) \subseteq \{1\} \cup [1+\xi, 1+\eta] \subset [1,2)$, where ξ is as in Theorem 3 and

$$\eta = \frac{\lambda_{\max}(CA^{-1}C^T)}{\lambda_{\min}(D) + \lambda_{\max}(CA^{-1}C^T)}.$$

Moreover, if C^T has full column rank then $\lambda = 1$ is an eigenvalue of multiplicity at least n.

Note: Of course, as usual, "exact" application of all these preconditioners may be very expensive or even impossible, and inexact versions are usually necessary and/or faster (see below).

Numerical experiments

We consider two model problems of the type (1), one with $D = 0$ and the other with $D \neq 0$:

- Mixed-hybrid FEM discretization of 3D potential fluid flow problems
- Liquid crystals director modeling

All of the reported numerical results were performed on a 64-bit 2.45 GHz core i7 processor and 8.00GB RAM using MATLAB version 8.3.0532.

In all of the experiments we have used right-hand sides corresponding to random solution vectors, performing ten runs and then averaging the CPU-times.

The iteration counts reported in the tables (under “Iter") are also averages (rounded to the nearest integer).

Numerical experiments (cont.)

All of the methods require repeated solution (whether “exact" or inexact) of SPD linear systems as subtasks. These are either solved by sparse Cholesky factorization with symmetric approximate minimum degree (SYMAMD) reordering or by the preconditioned conjugate gradient (PCG) method. When using PCG, unless otherwise specified, the preconditioner used is a drop tolerance-based incomplete Cholesky factorization with drop tolerance $\tau = 10^{-2}$.

In all of the numerical tests below, the initial guess is taken to be the zero vector. The outher iterations are stopped once

$$\|b - \mathcal{A}u^{(k)}\|_2 < 10^{-10}\|b\|_2.$$

For the inner PCG iterations (whenever applicable), the stopping tolerances used are specified below.

Implementation details

For the potential fluid flow problem, the best results were obtained with inexact versions of the preconditioners $\mathcal{P}_{GD}$, $\mathcal{P}_{GT,1}$ and $\mathcal{P}_{GT,2}$.

The first two of these preconditioners require solving linear systems with the $(1,1)$ submatrix A (easy!) and with the Schur complement matrix

$$\mathcal{S} = \begin{bmatrix} BA^{-1}B^T & BA^{-1}C^T \\ CA^{-1}B^T & CA^{-1}C^T \end{bmatrix} = \begin{bmatrix} B \\ C \end{bmatrix} A^{-1} \begin{bmatrix} B^T & C^T \end{bmatrix}.$$

This matrix is sparse and inexpensive to form. Also, $S_B = BA^{-1}B^T$ turns out to be a scalar multiple of identity, see Maryška et al. (SISC 2000).

Sparse direct solution of linear systems involving $\mathcal{S}$, however, is costly. We solve these linear systems with CG preconditioned by incomplete Cholesky factorization.

The PCG convergence tolerance is 10^{-15} for "exact" solution, 10^{-4} for inexact.

Implementation details (cont.)

Each application of $\mathcal{P}_{GT,2}^{-1}$ requires solving a saddle point problem

$$\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} r_1 \\ r_2 \end{bmatrix}, \tag{4}$$

followed by solution of a linear systems with the coefficient matrix

$$\bar{S} = C(A^{-1} + A^{-1}B^T S_B^{-1} B A^{-1})C^T.$$

The solution of (4) is very cheap and can be obtained in two steps as follows:

- **Step I.** Solve $S_B w_2 = BA^{-1}r_1 - r_2$, to find w_2.
- **Step II.** Set $w_1 = A^{-1}(r_1 - B^T w_2)$.

The Schur complement matrix $\bar{S}$ is also sparse for this problem. We use PCG as before to solve linear systems involving $\bar{S}$.

Test results for 3D potential fluid flow problem

size	Iter	CPU-time
2125	3	0.0125
17000	3	0.0947
57375	3	0.4829
136000	3	1.6226
265625	3	3.9002
459000	3	8.8899

Table: Results for MINRES with block diagonal preconditioner $\mathcal{P}_{GD}$.

No surprise here...

Test results for 3D potential fluid flow problem (cont.)

	Preconditioner			
	$\mathcal{P}_{GT,1}$		$\mathcal{P}_{GT,2}$	
size	Iter	CPU-time	Iter	CPU-time
2125	2	0.0191	2	0.0180
17000	2	0.1284	2	0.1180
57375	2	0.5247	2	0.4516
136000	2	1.5425	2	1.2936
265625	2	3.6811	2	3.1080
459000	2	7.9861	2	6.8368

Table: Results for GMRES with ideal block triangular preconditioners.

... or here. But what about the inexact case?

Test results for 3D potential fluid flow problem (cont.)

	Preconditioner					
	$\mathcal{P}_{GT,1}$			$\mathcal{P}_{GT,2}$		
size	Iter	Iter$_{pcg}$	CPU-time	Iter	Iter$_{pcg}$	CPU-time
2125	5	25	0.0085	5	25	0.0073
17000	6	47	0.0575	6	53	0.0534
57375	6	66	0.2361	6	72	0.2265
136000	6	87	0.7480	6	95	0.6563
265625	6	108	1.8190	6	112	1.5220
459000	6	134	4.2658	5	117	3.0442

Table: Results for FGMRES with inexact block triangular preconditioners.

Note: Stopping tolerance for inner PCG iterations: 10^{-4}.

Implementation details (cont.)

For the liquid crystal director model we use the preconditioners $\tilde{\mathcal{P}}_T$ and $\hat{\mathcal{P}}_T$. Both preconditioners require the solution of linear systems with SPD coefficient matrices A, $BA^{-1}B^T$, and $D + CA^{-1}C^T$ at each (outer) GMRES or FGMRES iteration.

The matrices $BA^{-1}B^T$ and $D + CA^{-1}C^T$ are full and we never form them explicitly.

Systems with these matrices are solved by PCG, with the preconditioners described next.

Implementation details (cont.)

- For $D + CA^{-1}C^T$ we use the Cholesky factors of $D + CD_A^{-1}C^T$, where $D_A = \text{diag}(A)$;
- For $BA^{-1}B^T$ we use the approximate inverse

$$BAB^T \approx (BA^{-1}B^T)^{-1},$$

which only requires sparse matrix-vector products.

Note: For this problem, B has (nearly) orthogonal rows. We observe that if B were an orthogonal matrix, then $(BA^{-1}B^T)^{-1} = BAB^T$.

Numerical experiments indicate that the spectrum of

$$(BAB^T)(BA^{-1}B^T)$$

is contained in the interval $[1, 1.3944]$, independent of mesh size.

Implementation details (cont.)

Preconditioner:	$\tilde{\mathcal{P}}_T$			$\hat{\mathcal{P}}_T$		
Coefficient matrix:	A	S_B	$\tilde{S}_C$	A	S_B	$\tilde{S}_C$
GMRES	$\star$	1e–12	1e–12	$\star$	1e–12	1e–12
FGMRES	1e–02	1e–02	1e–01	1e–03	1e–03	1e–01

Table: Inner tolerance in PCG method used inside the preconditioned methods, liquid crystal problem. A '$\star$' means that a sparse direct solver was used.

Test results for liquid crystals model

	Method			
	GMRES		FGMRES	
size	Iter	CPU	Iter	CPU
5115	10	0.53144	9	0.08461
10235	9	0.98055	9	0.14981
20475	9	1.9727	8	0.27168
40955	9	3.6642	8	0.53854
81915	9	8.4782	8	1.2848
163835	9	17.947	8	3.1012
327675	9	42.743	8	7.4957

Table: Results for ideal/inexact preconditioner $\tilde{\mathcal{P}}_T$

Note that the inexactly preconditioned FGMRES iteration may converge faster than the exactly preconditioned GMRES one!

Test results for liquid crystals model (cont.)

	Method			
	GMRES		FGMRES	
size	Iter	CPU	Iter	CPU
5115	6	0.32469	6	0.02981
10235	6	0.62408	6	0.05179
20475	6	1.2614	6	0.08967
40955	6	2.4815	6	0.17196
81915	6	5.4721	6	0.33430
163835	6	11.879	6	0.69814
327675	6	28.196	6	1.5901

Table: Results for ideal/inexact preconditioner $\hat{\mathcal{P}}_T$

Concluding remarks

- The 3×3 block structure in double saddle point problems leads to a number of possible block preconditoners
- Important to choose best partitioning for a given problem
- Some eigenvalue estimates derived; more theory needed
- Good approximations to multiple Schur complements may be required
- Inexact variants lead to efficient solvers
- Scalable (or nearly so) solvers for model problems of potential fluid flow and crystal directors (alternative to null space-based methods); these problems display additional structure
- See the paper for additional results and experiments for Uzawa-type methods