Electronic Transactions on Numerical Analysis.
Volume 22, pp. 114-121, 2006.

ISSN 1068-9613.

ETNA
Kent State University
etna@mcs.kent.edu

PRECONDITIONERS FOR SADDLE POINT LINEAR SYSTEMS WITH HIGHLY SINGULAR (1,1) BLOCKS*

CHEN GREIF† AND DOMINIK SCHÖTZAU‡

Abstract. We introduce a new preconditioning technique for the iterative solution of saddle point linear systems with (1,1) blocks that have a high nullity. The preconditioners are block diagonal and are based on augmentation, using symmetric positive definite weight matrices. If the nullity is equal to the number of constraints, the preconditioned matrices have precisely two distinct eigenvalues, giving rise to immediate convergence of preconditioned MINRES. Numerical examples illustrate our analytical findings.

Key words. saddle point linear systems, high nullity, augmentation, block diagonal preconditioners, Krylov subspace iterative solvers

AMS subject classifications. 65F10

1. Introduction.

Consider the following saddle point linear system

$$\mathcal{K}x \equiv \begin{pmatrix} F & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} c \\ d \end{pmatrix} \equiv b, \tag{1.1}$$

with $F \in \mathbb{R}^{n\times n}$ and $B \in \mathbb{R}^{m\times n}$, where $m \leq n$. The matrix F is assumed to be symmetric and have a high nullity. We further assume that $\mathcal{K}$ is nonsingular, from which it follows that

$$\operatorname{rank}(B) = m \quad \text{and} \quad \operatorname{null}(F) \cap \operatorname{null}(B) = \{0\}. \tag{1.2}$$

From (1.2) it also follows that F has rank at least $n - m$, and hence its nullity can be at most m. Saddle point linear systems of the form (1.1) appear in many applications; see [1] for a comprehensive survey. Frequently they are large and sparse, and iterative solvers must be applied. In recent years, a lot of research has focused on seeking effective preconditioners. For example, 2×2 block diagonal preconditioners have been successfully used in the simulation of incompressible flow problems; see [2] and references therein. Those preconditioners typically have a (1,1) block that approximates the (1,1) block of the original saddle point matrix, and a (2,2) block that approximates the Schur complement.

However, when F is singular, it cannot be inverted and the Schur complement does not exist. In this case, one possible way of dealing with the system is by augmentation, for example by replacing F with $F + B^T W^{-1} B$, where $W \in \mathbb{R}^{m\times m}$ is a symmetric positive definite weight matrix; see [3] and references therein.

In this paper we consider the case of a (1,1) block with a high nullity, and introduce a Schur complement-free preconditioner based on augmentation that leads to an effective iterative solution procedure. We show that if the nullity of F is m, then preconditioned MINRES [8] converges within two iterations. The approach presented in this paper is motivated in part by the recent work [4], where a block diagonal preconditioner is proposed for solving the time-harmonic Maxwell equations in mixed form.

The remainder of this paper is structured as follows. In Section 2 we present the preconditioners and analyze their spectral properties. In Section 3 numerical examples that validate our analytical findings are given. We conclude with brief remarks in Section 4.

*Received September 20, 2005. Accepted for publication November 14, 2005. Recommended by M. Benzi. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada.

†Department of Computer Science, University of British Columbia, Vancouver, BC V6T 1Z4, Canada (greif@cs.ubc.ca).

‡Mathematics Department, University of British Columbia, Vancouver, BC V6T 1Z2, Canada (schoetzau@math.ubc.ca).

2. The proposed preconditioners. We start with a general form of our preconditioners, and then discuss a specific choice that is particularly suitable for matrices with a semidefinite (1,1) block with high nullity. We end the section with a brief discussion of computational costs.

2.1. The preconditioner $\mathcal{M}_{U,W}$. Consider the following block diagonal matrix as a preconditioner:

$$\mathcal{M}_{U,W} = \begin{pmatrix} F + B^T U^{-1} B & 0 \\ 0 & W \end{pmatrix},$$

where $U, W \in \mathbb{R}^{m \times m}$ are symmetric positive definite.

PROPOSITION 2.1. *Suppose $\mathcal{M}_{U,W}$ is symmetric positive definite. Let $\{z_i\}_{i=1}^{n-m}$ be a basis of the null space of B. Then the vectors $(z_i, 0)$ are $n - m$ linearly independent eigenvectors of $\mathcal{M}_{U,W}^{-1}\mathcal{K}$ with eigenvalue* 1.

Proof. The eigenvalue problem for $\mathcal{M}_{U,W}^{-1}\mathcal{K}$ is

$$\begin{pmatrix} F & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} v \\ q \end{pmatrix} = \mu \begin{pmatrix} F + B^T U^{-1} B & 0 \\ 0 & W \end{pmatrix} \begin{pmatrix} v \\ q \end{pmatrix}.$$

From the nonsingularity of $\mathcal{K}$ it follows that $\mu \neq 0$. Substituting $q = \frac{1}{\mu} W^{-1} B v$, we obtain for the first block row

$$\mu F v + B^T W^{-1} B v = \mu^2 (F + B^T U^{-1} B) v. \tag{2.1}$$

Suppose that $v = z_i \neq 0$ is a null vector of B. Then (2.1) simplifies into

$$(\mu^2 - \mu) F z_i = 0,$$

and since by (1.2) a nonzero null vector of B cannot be a null vector of F, it follows that $F z_i \neq 0$ and hence we must have $\mu = 1$. Since $B z_i = 0$, it follows that $q = 0$ and therefore $\mu = 1$ is an eigenvalue of $\mathcal{M}_{U,W}^{-1}\mathcal{K}$ of algebraic multiplicity (at least) $n - m$, whose associated eigenvectors are $(z_i, 0)$, $i = 1, \ldots, n - m$. $\square$

2.2. The preconditioner $\mathcal{M}_W$. From Proposition 2.1 it follows that regardless of U and W, we have at least $n - m$ eigenvalues equal to 1. Stronger clustering can be obtained for specific choices of those two weight matrices. For the case of a (1,1) block with high nullity it is possible to obtain a preconditioner with improved spectral properties by making the choice $U = W$. Let us define

$$\mathcal{M}_W = \begin{pmatrix} F + B^T W^{-1} B & 0 \\ 0 & W \end{pmatrix}. \tag{2.2}$$

If in addition F is semidefinite, it follows from (1.2) that the augmented (1,1) block is positive definite, making it possible to use the preconditioned conjugate gradient.

The next theorem provides details on the spectrum of the preconditioned matrix $\mathcal{M}_W^{-1}\mathcal{K}$.

THEOREM 2.2. *Suppose that F is symmetric positive semidefinite with nullity r. Then $\mu = 1$ is an eigenvalue of $\mathcal{M}_W^{-1}\mathcal{K}$ of algebraic multiplicity n and $\mu = -1$ is an eigenvalue of multiplicity r. The remaining $m - r$ eigenvalues of $\mathcal{M}_W^{-1}\mathcal{K}$ are all strictly between -1 and 0 and satisfy the relation*

$$\mu = -\frac{\lambda}{\lambda + 1}, \tag{2.3}$$

where λ are the $m-r$ positive generalized eigenvalues of

$$\lambda F v = B^T W^{-1} B v. \tag{2.4}$$

A set of linearly independent eigenvectors for $\mu = \pm 1$ can be found as follows. Let $\{z_i\}_{i=1}^{n-m}$ be a basis of the null space of B, $\{x_i\}_{i=1}^{r}$ a basis of the null space of F, and $\{y_i\}_{i=1}^{m-r}$ a set of linearly independent vectors that complete $\mathrm{null}(F) \cup \mathrm{null}(B)$ to a basis of $\mathbb{R}^n$. Then the $n-m$ vectors $(z_i, 0)$, the r vectors $(x_i, W^{-1}Bx_i)$ and the $m-r$ vectors $(y_i, W^{-1}By_i)$ are linearly independent eigenvectors associated with $\mu = 1$, and the r vectors $(x_i, -W^{-1}Bx_i)$ are eigenvectors associated with $\mu = -1$.

Proof. Let μ be an eigenvalue of $\mathcal{M}_W^{-1}\mathcal{K}$ with eigenvector (v, q). Then

$$\begin{pmatrix} F & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} v \\ q \end{pmatrix} = \mu \begin{pmatrix} F + B^T W^{-1} B & 0 \\ 0 & W \end{pmatrix} \begin{pmatrix} v \\ q \end{pmatrix}.$$

Since $\mathcal{K}$ is nonsingular, we have $\mu \neq 0$. Substituting $q = \frac{1}{\mu} W^{-1} B v$ we obtain

$$(\mu^2 - \mu) F v + (\mu^2 - 1) B^T W^{-1} B v = 0. \tag{2.5}$$

If $\mu = 1$, then (2.5) is satisfied for any arbitrary nonzero vector $v \in \mathbb{R}^n$, and hence $(v, W^{-1}Bv)$ is an eigenvector of $\mathcal{M}_W^{-1}\mathcal{K}$.

If $x \in \mathrm{null}(F)$ then from (2.5) we obtain

$$(\mu^2 - 1) B^T W^{-1} B x = 0,$$

from which it follows that $(x, W^{-1}Bx)$ and $(x, -W^{-1}Bx)$ are eigenvectors associated with $\mu = 1$ and $\mu = -1$ respectively.

Next, suppose $|\mu| \neq 1$. We divide (2.5) by $\mu^2 - 1$, which yields (2.3), with v defined in (2.4). Since F and $B^T W^{-1} B$ are positive semidefinite, the remaining generalized eigenvalues λ must be positive and hence μ must be strictly between -1 and 0, as stated in the theorem.

A specific set of linearly independent eigenvectors for $\mu = \pm 1$ can be readily found. The vectors $\{z_i\}_{i=1}^{n-m}$ and $\{x_i\}_{i=1}^{r}$ defined above are linearly independent by (1.2) and span a subspace of $\mathbb{R}^n$ of dimension $n - m + r$. Let $\{y_i\}_{i=1}^{m-r}$ complete this set to a basis of $\mathbb{R}^n$ as stated above. It follows that $(x_i, W^{-1}Bx_i)$, $(z_i, 0)$, and $(y_i, W^{-1}By_i)$ are eigenvectors of $\mathcal{M}_W^{-1}\mathcal{K}$ associated with $\mu = 1$. The r vectors $(x_i, -W^{-1}Bx_i)$ are eigenvectors associated with $\mu = -1$. $\square$

A convenient choice for the weight matrix is $W^{-1} = \gamma I$, where $\gamma > 0$ is a parameter that takes into account the scales of the matrices F and B [3]. In this case, notice that from Theorem 2.2 it follows that the $m - r$ eigenvalues μ of $\mathcal{M}_W^{-1}\mathcal{K}$ that are not equal to ± 1 are given by

$$\mu = -\frac{\lambda\gamma}{\lambda\gamma + 1},$$

where λ are the generalized eigenvalues defined by

$$\lambda F v = B^T B v.$$

Thus, as γ increases these eigenvalues tend to -1, and further clustering is obtained. We note, however, that choosing γ too large may result in ill-conditioning of $\mathcal{M}_W$.

ETNA
Kent State University
etna@mcs.kent.edu

By (1.2) the nullity of F must be m at most. From Theorem 2.2 we conclude that the higher it is, the more strongly the eigenvalues are clustered. In fact, for nullity m we have the following result.

COROLLARY 2.3. *Suppose that F is positive semidefinite with nullity m. Then the preconditioned matrix $\mathcal{M}_W^{-1}\mathcal{K}$ has precisely two eigenvalues: $\mu = 1$, of multiplicity n, and $\mu = -1$, of multiplicity m.*

Corollary 2.3 implies that a preconditioned minimal residual Krylov subspace solver is expected to converge within two iterations, in the absence of roundoff errors. Each preconditioned MINRES iteration with $\mathcal{M}_W$ includes a matrix-vector product with $\mathcal{K}$, a solve for W, and a solve for $F_W = F + B^T W^{-1} B$. If F_W is formed explicitly, then solving for it includes n solves for W, one for each column of B. However, typically F_W is not formed explicitly since W^{-1} is dense. In this case one can apply the (preconditioned) conjugate gradient method, and then the number of solves for W is equal to the number of iterations. Hence, since the number of MINRES iterations is guaranteed to be small by the analysis of this section, the key for an effective numerical solution procedure overall is the ability to efficiently solve for F_W.

3. Numerical examples. In this section we illustrate the performance of our preconditioning approach on two applications in which the (1,1) blocks of the associated matrices are highly singular.

3.1. Maxwell equations in mixed form. We first consider a finite element discretization of the time-harmonic Maxwell equations with small wave numbers [9, Section 2]. The following two-dimensional model problem is considered: find u and p that satisfy

$$\begin{aligned} \nabla \times \nabla \times u - k^2 u + \nabla p &= f \qquad \text{in } \Omega, \\ \nabla \cdot u &= 0 \qquad \text{in } \Omega, \\ u \cdot t &= 0 \qquad \text{on } \partial\Omega, \\ p &= 0 \qquad \text{on } \partial\Omega. \end{aligned}$$

Here u is the electric field and p is a Lagrange multiplier, $\Omega \subset \mathbb{R}^2$ is a simply connected polygonal domain, and t denotes the tangential unit vector on $\partial\Omega$. The datum f is a given generic source. We assume that the wave number k^2 is small and is not a Maxwell eigenvalue.

We employ a standard finite element discretization on uniformly refined triangular meshes of size h. The lowest order two-dimensional Nédélec elements of the first kind [6, 7] are used for the approximation of the electric field, along with standard nodal elements for the multiplier. This yields a saddle point linear system of the form

$$\begin{pmatrix} A - k^2 M & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} g \\ 0 \end{pmatrix},$$

where now $u \in \mathbb{R}^n$ and $p \in \mathbb{R}^m$ are finite arrays representing the finite element approximations, and $g \in \mathbb{R}^n$ is the load vector associated with the datum f. The matrix $A \in \mathbb{R}^{n\times n}$ is symmetric positive semidefinite with rank $n - m$, and corresponds to the discrete curl-curl operator; $B \in \mathbb{R}^{m\times n}$ is a discrete divergence operator. Due to the zero Dirichlet boundary conditions, B has full row rank: $\mathrm{rank}(B) = m$. Indeed, the discretization that we use is inf-sup stable [6, p. 179]. The matrix $F = A - k^2 M$ is positive semidefinite for $k = 0$ and indefinite for $k > 0$. When the mesh size is sufficiently small, the saddle point matrix $\mathcal{K}$ is nonsingular [6, Chapter 7].

For the purpose of illustrating the merits of our approach, we will deliberately avoid exploiting specific discrete differential operator-related properties, and focus instead on purely algebraic considerations. To that end, we pick a scaled identity matrix, $W^{-1} = \gamma I$. Based on scaling considerations, we set $\gamma = \frac{\|A\|_1}{\|B\|_1^2}$.

We consider five meshes; the number of elements and dimension of the resulting systems are given in Table 3.1.

TABLE 3.1
Example 3.1: number of elements and sizes of the linear systems for five meshes.

Mesh	Nel	$n+m$
G1	64	113
G2	256	481
G3	1024	1985
G4	4096	8065
G5	16384	32513

Experiments were done with several right hand side functions, and the iteration counts were practically identical in all experiments. In the tables below we report the results that were obtained by setting $f = 1$. (Note that in this case the datum is divergence free.)

Table 3.2 validates the analysis of Section 2. It shows the iteration counts for preconditioned MINRES, applying exact inner solves, for various values of k and meshes G1–G5. The (outer) iteration was stopped using a threshold of 10^{-6} for the relative residual. We observe that for $k = 0$ convergence is always reached within a single iteration, which is better than two iterations guaranteed by Theorem 2.2. (This behavior might be related to special properties of the underlying differential operators, that allow for decoupling the problem using the discrete Helmholtz decomposition [4].) As k grows larger, and/or as the mesh is refined, Theorem 2.2 does not apply anymore and convergence is slower. However, Proposition 2.1 holds and the solver is still remarkably robust, at least for small values of k. Preconditioning the same problem with high wave numbers introduces additional computational challenges and is not considered here.

TABLE 3.2
Example 3.1: iteration counts for various values of k and meshes G1–G5 using exact inner solves. The iteration was stopped using a threshold of 10^{-6} for the relative residual.

Mesh	$k=0$	$k=0.25$	$k=0.5$	$k=0.75$	$k=1$
G1	1	1	1	1	1
G2	1	2	2	3	3
G3	1	2	2	3	3
G4	1	2	2	3	3
G5	1	2	2	3	3

Figure 3.1 depicts the negative eigenvalues of the preconditioned matrix $\mathcal{M}_W^{-1}\mathcal{K}$ for the mesh G2 and $k = 0.5$. They are extremely close to -1. This shows the potential of the preconditioner even for cases of an indefinite (1,1) block, in which case Theorem 2.2 does not hold.

In practice the preconditioner solves need to be done iteratively. Efficient multigrid solvers that exploit the properties of the differential operators are available and can be used (see [6, Chapter 13] and references therein). Here we simply consider the conjugate gradient iteration, preconditioned using the incomplete Cholesky decomposition, IC(0). It should be noted that the use of a non-stationary iteration (like PCG) in the inner solves means that a non-constant, nonlinear preconditioning operator is introduced for the outer solver. In such

ETNA
Kent State University
etna@mcs.kent.edu

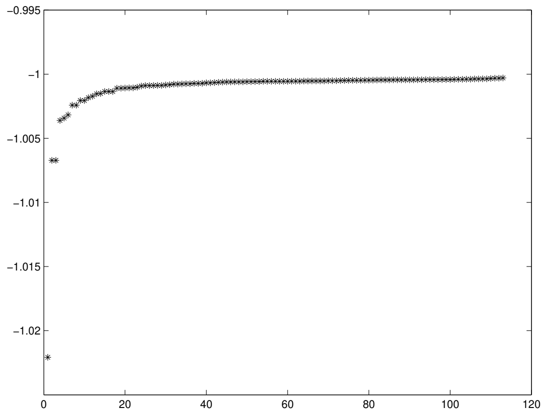

FIG. 3.1. *Example 3.1: the negative eigenvalues of the preconditioned matrix* $\mathcal{M}_W^{-1}\mathcal{K}$ *for* $k = 0.5$ *and grid G2. All the* positive *eigenvalues of* $\mathcal{M}_W^{-1}\mathcal{K}$ *are identically equal to 1.*

settings flexible Krylov methods for the outer iteration are commonly used. However, we have used MINRES and experimented with a fixed loose inner tolerance, and our conclusion is that this inner solve strategy works well.

Table 3.3 shows the performance of MINRES, preconditioned with our preconditioner, with a fixed inner tolerance of 10^{-2} and an outer tolerance of 10^{-6}. Naturally, there is an increase of iterations as the inner tolerance is loosened, as is evident when Tables 3.2 and 3.3 are compared to each other. Nevertheless, the speed of convergence of the inner solves, resulting from loosening the stopping criterion, more than compensates for the increase in the number of outer iterations, and results in significant savings.

TABLE 3.3
Example 3.1: iteration counts for various values of k *and meshes G1–G5 using inexact inner solves. The stopping criterion for the inner iterations was a threshold of* 10^{-2} *for the relative residual. For the outer iterations, the threshold was* 10^{-6}.

Mesh	$k = 0$	$k = 0.25$	$k = 0.5$	$k = 0.75$	$k = 1$
G1	4	4	4	6	6
G2	6	6	6	6	6
G3	6	6	6	6	7
G4	6	6	6	6	7
G5	6	6	6	7	7

3.2. An inverse problem. As a second numerical example we consider a nonlinear minimization problem taken from [5], which arises in geophysics, electromagnetics, and other areas of applications. Suppose the vector b represents observations of a field v at some discrete locations and r is the underlying model to be recovered. Suppose further that Q projects the field v onto the measurement locations. The constrained problem formulation in [5] is based on minimizing $\|Qv - b\|_2$, subject to a forward problem (typically a discretized second order PDE) $C(r)v = f$ that needs to be solved exactly. Upon regularization, the following minimization problem is obtained:

$$\text{minimize} \;\; \phi(v,r) = \frac{1}{2}\|Qv - b\|_2^2 + \frac{\beta}{2}\|S(r - r_0)\|_2^2$$
$$\text{subject to} \;\; C(r)v = f$$

ETNA
Kent State University
etna@mcs.kent.edu

where r_0 is a reference model, $S^T S$ is a smoothing operator (typically a diffusion operator), and β is a regularization parameter.

The constraints are incorporated using a Lagrange multiplier approach, and a Gauss-Newton iteration is applied. At each step, an indefinite linear system of the following block form has to be solved:

$$\begin{pmatrix} Q^T Q & 0 & C^T \\ 0 & \beta S^T S & G^T \\ C & G & 0 \end{pmatrix} \begin{pmatrix} \delta v \\ \delta r \\ \delta\lambda \end{pmatrix} = \begin{pmatrix} L_v \\ L_m \\ L_\lambda \end{pmatrix}.$$

Here $\delta\lambda$ are the increments of the Lagrange multipliers and G is the Jacobian of C with respect to r. The matrix Q can be extremely sparse, in particular in situations of undersampling.

A three-dimensional problem on the unit cube is considered, discretized by standard finite volumes. The regularization parameter β is equal to 10^{-2}. The operator S is the discretized gradient and C is a discrete diffusion operator with diffusivity depending on r. Finally, f is a vector obtained from sampling a smooth analytical function. A full description of models of this type is given in [5].

We consider the performance of preconditioned MINRES on three uniformly refined meshes M1–M3. Since there is no obvious scaling strategy, we set $W = I$. The dimensions of the associated linear systems, the nullities of the $(1, 1)$ blocks, and the iteration counts are given in Table 3.4. As is evident, our solver performs extremely well. Numerical experiments for other values of the regularization parameter β have shown similar iteration counts.

TABLE 3.4
Example 3.2: sizes of the linear systems and iteration counts for meshes M1–M3 using exact inner solves. The iteration was stopped using a threshold of 10^{-6} for the relative residual.

Mesh	n	m	Nullity	Iterations
M1	189	125	72	4
M2	1241	729	530	3
M3	9009	4913	4538	2

Finally, in Figure 3.2 we show the distribution of the eigenvalues of the preconditioned matrix for mesh M2. As expected, we observe strong clustering of the eigenvalues at ± 1.

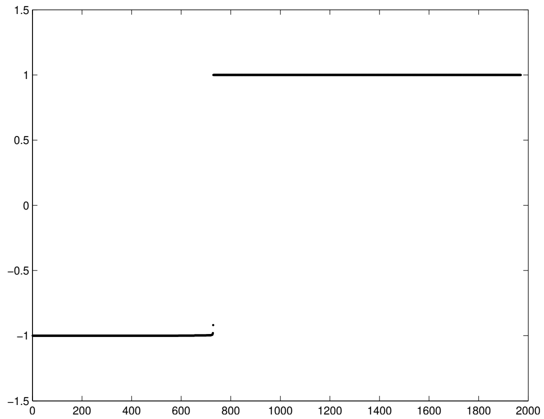

FIG. 3.2. *Example 3.2: Eigenvalues of the preconditioned matrix for mesh M2.*

4. Conclusions. We have presented a new Schur complement-free preconditioning approach based on augmenting the (1,1) block and using the weight matrix applied for augmentation as the matrix in the (2,2) block. As we have shown, this approach is very effective, and specifically, in cases where the (1,1) block has high nullity, convergence is guaranteed to be almost immediate. We have shown the high potential of our approach for the time-harmonic Maxwell equations in mixed form and for an inverse problem.

Acknowledgments. We thank two anonymous referees for their valuable comments and suggestions. We also thank Michele Benzi for helpful remarks related to inner-outer iterations, and Eldad Haber for providing the code used in our second numerical example.

REFERENCES

[1] M. BENZI, G.H. GOLUB, AND J. LIESEN, *Numerical solution of saddle point problems*, Acta Numer., 14 (2005), pp. 1–137.
[2] H.C. ELMAN, D.J. SILVESTER, AND A.J. WATHEN, *Finite Elements and Fast Iterative Solvers*, Oxford University Press, 2005.
[3] G.H. GOLUB AND C. GREIF, *On solving block structured indefinite linear systems*, SIAM J. Sci. Comput., 24 (2003), pp. 2076–2092.
[4] C. GREIF AND D. SCHÖTZAU, *Preconditioners for the discretized time-harmonic Maxwell equations in mixed form*, Technical Report TR-2006-01, Department of Computer Science, The University of British Columbia, 2006.
[5] E. HABER, U.M. ASCHER, AND D. OLDENBURG, *On optimization techniques for solving nonlinear inverse problems*, Inverse Problems, 16 (2000), pp. 1263–1280.
[6] P. MONK, *Finite element methods for Maxwell's equations*, Oxford University Press, New York, 2003.
[7] J.C. NÉDÉLEC, *Mixed finite elements in* $\mathbb{R}^3$, Numer. Math., 35 (1980), pp. 315–341.
[8] C.C. PAIGE AND M.A. SAUNDERS, *Solution of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629.
[9] I. PERUGIA, D. SCHÖTZAU, AND P. MONK, *Stabilized interior penalty methods for the time-harmonic Maxwell equations*, Comput. Methods Appl. Mech. Engrg., 191 (2002), pp. 4675–4697.