

ON SOLVING MAXWELLIAN EIGENVALUE PROBLEMS FOR ACCELERATING CAVITIES

Peter Arbenz, Institute of Scientific Computing, Swiss Federal Institute of Technology (ETH)
8092 Zurich, Switzerland

Stefan Adam, Paul Scherrer Institute, 5232 Villigen, Switzerland

Abstract

We investigate algorithms for solving large sparse symmetric matrix eigenvalue problems resulting from finite element discretizations of steady state electromagnetic fields in accelerating cavities. The methods have been applied to the new design of the accelerating cavity for the PSI 590 MeV ring cyclotron.

The solutions of this kind of eigenvalue problems can be polluted by so-called spurious modes if the divergence-free condition is not treated properly. In this paper we deal with a method that suppresses spurious modes by adding a penalty term to the basic quadratic form. This is the method we had the best experience with [1, 2].

The large sparse eigenvalue problems have been solved with the implicitly restarted Lanczos algorithm. Numerical results obtained on a HP Exemplar X-Class System are reported.

1 INTRODUCTION

For the production of high intensity proton, neutron and meson beams used in a wide range of research activities the cyclotron facility of the Paul Scherrer Institute accelerates proton beams to an energy of 590 MeV with an intensity of 1.5 mA in routine operation. The four accelerating cavities in the separate sector cyclotron provide a 50 MHz peak accelerating voltage of 750 kV each. Aiming at a reduction of the power losses in the cavity walls a project for a new design of the accelerating cavities has been started, cf. Fig. 1.

During design the RF engineer is interested in the field shapes and the frequencies of the fundamental and of some few higher modes of the standing electromagnetic waves in a cavity. This corresponds to the wish for access to a Maxwellian eigenvalue solver for some of the lowest eigenvalues and corresponding eigenvectors.

At the Institute for Scientific Computing at ETH Zurich the key questions related to the creation of such an eigensolver have been studied. In order to be able to handle complex geometries with locally fine details, the formulation of the problem was based on finite elements and emphasis has been given to the efficient solution of large problems. It proved to be important to understand the occurrence and to find means to avoid the so-called spurious modes (or ghost modes) that can occur in finite element formulations of Maxwell's equations. This study compared different finite element formulation schemes each used with different

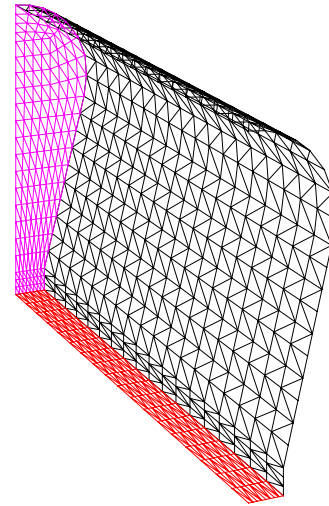


Figure 1: Surface triangulation of an eighth of the new PSI cavity design.

orders of approximation as well as several algorithms for the solution of the corresponding large sparse matrix eigenvalue problems [1, 2].

Different algorithms have been investigated to compute a few of the eigenvalues of the resulting large sparse matrix eigenvalue problems [1, 2]. Here we restrict ourselves to the implicitly restarted Lanczos algorithm that is, together with the Jacobi-Davidson algorithm, the best suited algorithm for this kind of problems. The linear systems of equations that occur in the shift-and-invert mode of the Lanczos algorithm have been solved by the conjugate gradient method in conjunction with a hierarchical basis preconditioner.

2 THE MAXWELLIAN EIGENVALUE PROBLEM

Without changing the basic structure of the problem one can assume that the metallic surfaces are perfectly conducting and that the inside of the cavity Ω is all in vacuum. The electromagnetic field in the cavity is described by the Maxwell equations [13]. After separation of time and space variables and after elimination of the magnetic field inten-

sity the differential equations

$$\mathbf{curl} \mathbf{curl} \mathbf{e}(\mathbf{x}) = \lambda \mathbf{e}(\mathbf{x}), \quad \mathbf{x} \in \Omega, \quad \lambda := \omega^2/c^2, \quad (1a)$$

$$\operatorname{div} \mathbf{e}(\mathbf{x}) = 0, \quad \mathbf{x} \in \Omega, \quad (1b)$$

$$\mathbf{n} \times \mathbf{e} = 0, \quad \mathbf{x} \in \Gamma := \partial\Omega. \quad (1c)$$

are obtained for the electric field intensity \mathbf{e} . Equations (1) have solutions only for certain values ω called *eigenfrequencies* of the electromagnetic oscillations of the resonator. These solutions are called *eigenmodes*.

The eigenvalue problem (1) can be solved analytically for a few domains, in particular for the case where Ω is a rectangular box [1]. We conducted our numerical experiments by means of this simple model which allows to compare the numerical results to an absolute reference. Moreover, for separate sector cyclotrons like the 590 MeV ring cyclotron installed at the Paul Scherrer Institute (PSI) in Villigen, Switzerland, the gross shape of a cavity is actually well approximated by a rectangular box.

The variational form of the eigenvalue problem (1) is given by [1]

$$\begin{aligned} \text{Find } (\lambda, \mathbf{u}) \in \mathbb{R} \times W_0 \text{ such that } \mathbf{u} \neq \mathbf{0} \text{ and} \\ (\mathbf{curl} \mathbf{u}, \mathbf{curl} \mathbf{v}) = \lambda(\mathbf{u}, \mathbf{v}), \quad \forall \mathbf{v} \in W_0. \end{aligned} \quad (2)$$

Let $L^2(\Omega)$ be the Hilbert space of square-integrable functions over the 3-dimensional domain Ω with inner product $(\mathbf{u}, \mathbf{v}) = \int_{\Omega} \mathbf{u}(\mathbf{x}) \cdot \mathbf{v}(\mathbf{x}) \, d\mathbf{x}$ and norm $\|\mathbf{u}\|_{0,\Omega} := (\mathbf{u}, \mathbf{u})^{1/2}$. In (2), $W_0 := \{\mathbf{v} \in W \mid \operatorname{div} \mathbf{v} = 0\}$ with

$$\begin{aligned} W := \{\mathbf{v} \in L^2(\Omega)^3 \mid \mathbf{curl} \mathbf{v} \in L^2(\Omega)^3, \\ \operatorname{div} \mathbf{v} \in L^2(\Omega), \mathbf{n} \times \mathbf{v} = 0 \text{ on } \partial\Omega\}. \end{aligned} \quad (3)$$

The difficulty with (2) stems from the condition $\operatorname{div} \mathbf{v} = 0$ as it is hard to find divergence-free finite elements. Therefore, ways have been looked for to get around this *divergence-free condition*. In this process care has to be taken in order not to introduce so-called *spurious modes*, i.e. eigenmodes that have no physical meaning.

In [1, 2] we considered two approaches free of spurious modes, a penalty method and a mixed method. The former was implemented with Lagrange finite elements the latter with Nédélec elements [9, 11]. In this note we restrict ourselves to the penalty method. Results concerning the Nédélec elements will be reported upon later. From our earlier experiments we expect them to be less efficient. In those experiments the computing time needed to get a prescribed accuracy for eigenvalues and vectors was about an order of magnitude smaller with the penalty method than with the mixed method. In the penalty method approach (2) is replaced by [14, 16]

$$\begin{aligned} \text{For } s > 0, \text{ find } (\lambda, \mathbf{u}) \in \mathbb{R} \times W, \mathbf{u} \neq \mathbf{0}, \text{ such that} \\ (\mathbf{curl} \mathbf{u}, \mathbf{curl} \mathbf{v}) + s(\operatorname{div} \mathbf{u}, \operatorname{div} \mathbf{v}) = \lambda(\mathbf{u}, \mathbf{v}), \quad (4) \\ \text{holds for all } \mathbf{v} \in W, \end{aligned}$$

In (4), s is a positive, usually small parameter. The eigenmodes $\mathbf{u}(\mathbf{x})$ corresponding to eigenvalues $\lambda < \mu_1 s$ are

eigenmodes of (2). μ_1 is the smallest eigenvalue of the negative Laplace operator $-\Delta$ on Ω . Therefore, s has to be chosen such that $\mu_1 s$ is above the largest eigenvalue we want to compute.

3 FINITE ELEMENT FORMULATION

We triangulate the open, bounded domain $\Omega \subset \mathbb{R}^3$ with polyhedral boundary Γ by tetrahedrons Ω_e . We approximate W by the well known Lagrange (or node-based) finite elements spaces [9],

$$W_h^{(k)} := \{\mathbf{v} \in W \mid \mathbf{v}|_{\Omega_e} \in (P_k(\Omega_e))^3\}, \quad (5)$$

where $P_k(\Omega_e)$ is the space of polynomials in Ω_e of degree $\leq k$. In our numerical experiments [2] we found that trial spaces made up of piecewise quadratic ($k=2$) polynomials are superior to trial spaces made up of piecewise linear ($k=1$) with respect to computing cost relative to accuracy. We will therefore restrict ourselves to $W_h^{(2)}$ equipped with a hierarchical basis [6, 23],

$$W_h^{(2)} = W_h^{(1)} \oplus V_h^{(2)}. \quad (6)$$

The piecewise linear polynomials in $W_h^{(1)}$ are determined by their values in the four vertices of the tetrahedron. The piecewise quadratic polynomials in $V_h^{(2)}$ are determined by the function values in the midpoints of the six edges.

Using $W_h^{(2)}$ in the Rayleigh-Ritz method for discretizing (4) gives the finite dimensional problem

$$\begin{aligned} \text{For } s > 0, \text{ find } (\lambda_h, \mathbf{e}_h) \in \mathbb{R} \times W_h^{(2)}, \mathbf{e}_h \neq \mathbf{0}, \text{ s.t.} \\ (\mathbf{curl} \mathbf{e}_h, \mathbf{curl} \Psi_h) + s(\operatorname{div} \mathbf{e}_h, \operatorname{div} \Psi_h) = \lambda(\mathbf{e}, \Psi) \quad (7) \\ \text{holds for all } \Psi_h \in W_h^{(2)}. \end{aligned}$$

Let $\{\Phi_i\}_{i=1}^n$ be a basis of $W_h^{(2)}$ and $\mathbf{e}_h = \sum_{i=1}^n \Phi_i \xi_i$. Then (7) becomes the matrix eigenvalue problem

$$A\mathbf{x} = \lambda M\mathbf{x}, \quad \mathbf{x} = (\xi_1, \dots, \xi_n), \quad (8)$$

where the elements of A and M are given by

$$\begin{aligned} a_{i,j} &= (\mathbf{curl} \Phi_i, \mathbf{curl} \Phi_j) + s(\operatorname{div} \Phi_i, \operatorname{div} \Phi_j), \\ m_{i,j} &= (\Phi_i, \Phi_j). \end{aligned}$$

For positive s , both A and M are symmetric positive definite n -by- n matrices. If the basis functions are chosen properly, then A and M are sparse.

Let us consider a rectangular box that is divided into $m_1 \times m_2 \times m_3$ boxes each of which is in turn subdivided into six tetrahedrons. Then $\dim W_h^{(2)} \approx 24m_1m_2m_3$. The matrices M and A have up to 65 and 134 nonzero elements per row, respectively. Notice that $\dim W_h^{(1)} \approx 3m_1m_2m_3$. We will use $W_h^{(1)}$ in our ‘coarse grid’ correction in the hierarchical preconditioner of section 5.

4 THE IMPLICITLY RESTARTED LANCZOS ALGORITHM (IRL)

For computing a few, say p , eigenvalues of a sparse matrix eigenvalue problem

$$A\mathbf{x} = \lambda M\mathbf{x}, \quad A = A^T \geq 0, \quad M = M^T > 0, \quad (9)$$

closest to a number τ it is advisable to make a shift-and-invert spectral transformation with a *shift* σ close to τ and solve [12, 17]

$$C\mathbf{x} := (A - \sigma M)^{-1}M\mathbf{x} = \mu\mathbf{x}, \quad \mu = \frac{1}{\lambda - \sigma}. \quad (10)$$

instead of (9). $C = (A - \sigma M)^{-1}M$ is M -symmetric, i.e., it is symmetric with respect to the inner product $\mathbf{x}^T M \mathbf{y}$. The spectral transformation leaves the eigenvectors unchanged. The eigenvalues of (9) close to the shift become the largest absolute of (10). In addition they are relatively well-separated which improves the speed of convergence of Krylov-type subspace methods [20]. The cost of the improved convergence rate is the necessity to solve a linear system of equations involving $A - \sigma M$.

In [1, 2] we compared four algorithms for computing a few of the smallest eigenvalues of (10): (i) subspace iteration [20], (ii) the block Lanczos algorithm [12], (iii) the implicitly restarted Lanczos algorithm [22, 8], and (iv) the Jacobi-Davidson algorithm [21, 10]. Subspace iteration was not competitive because of its low convergence rate. The block Lanczos algorithm was performing best for problems of limited size. However, the memory space available was not sufficient for solving our largest problems. The eigensolvers that were able to handle all problem sizes were the implicitly restarted Lanczos algorithm and the Jacobi-Davidson algorithm. The former was in our experiments slightly faster than Jacobi-Davidson by about 10-20%. We therefore only consider IRL in this study. We used the Fortran subroutines from the publicly available ARPACK [15].

To overcome the limitation of the Lanczos algorithm with respect to memory consumption, Sorensen proposed an elegant way to restart the iteration process [22, 8]. (These ideas apply to the Arnoldi algorithm for non-symmetric eigenvalue problems as well.)

The Lanczos iteration process (here with the shift-and-invert approach)

$$\begin{aligned} \mathbf{q}_{j+1}\beta_{j+1} &= \mathbf{r}_{j+1} = C\mathbf{q}_j - \alpha_j\mathbf{q}_j - \beta_j\mathbf{q}_{j-1}, \\ \mathbf{q}_i^T M \mathbf{q}_j &= \delta_{ij}, \quad \alpha_j = \mathbf{q}_j^T M (A - \sigma M)^{-1} M \mathbf{q}_j, \end{aligned} \quad (11)$$

is executed until $j = p+k$, where k is some positive integer. Often $k = p$ is chosen. Then, the Lanczos relation

$$CV_{p+k} = V_{p+k}T_{p+k} + \mathbf{r}_{p+k+1}\mathbf{e}_{p+k}^T \quad (12)$$

holds where $V_{p+k} = [\mathbf{q}_1, \dots, \mathbf{q}_{p+k}]$ and T_{p+k} is a symmetric tridiagonal matrix. Now, k sweeps of the QR algorithm [20] with shifts μ_1, \dots, μ_k are applied to T_{p+k} to

obtain $\hat{T} = \hat{Q}^T T_{p+k} \hat{Q}$ where \hat{Q} contains the cumulated rotations of the QR sweeps. Multiplication of (12) by \hat{Q} from the right and considering only the first p columns yields a new Lanczos relation

$$C\hat{V}_p = \hat{V}_p \hat{T}_p + \hat{\mathbf{r}}_{p+1} \mathbf{e}_p^T, \quad (13)$$

with which the computations are continued. In (13), \hat{V}_p are the first p columns of $V_{p+k}\hat{Q}$, \hat{T}_p is the $p \times p$ principal submatrix of $\hat{T} = \hat{Q}^T T_{p+k} \hat{Q}$ and $\hat{\mathbf{r}}_{p+1}$ is the p -th column of $\mathbf{r}_{p+k+1}\mathbf{e}_{p+k}^T \hat{Q}$. In ARPACK, the shifts μ_1, \dots, μ_k are chosen as the k eigenvalues of \hat{T}_{p+k} furthest away from the desired target value τ . Then, the eigenvalues of \hat{T}_p are the p eigenvalues of T_{p+k} closest to the target value.

Besides the storage for the matrices A and M , the memory requirements of the IRL algorithm are essentially the space needed to store \mathbf{q}_j and $M\mathbf{q}_j$, $j = 1, \dots, p+k$, i.e., $2n(p+k)$ floating point numbers.

5 TWO-LEVEL HIERARCHICAL BASIS PRECONDITIONERS

With ARPACK it is possible to use any algorithm to solve the indefinite system of equations $(A - \sigma M)\mathbf{x} = \mathbf{y} = M\mathbf{q}_j$ in (11). We chose the iterative solver SYMMLQ of Paige and Saunders [19], a variant of the conjugate gradients method designed to handle symmetric indefinite systems of equations that we have to expect if the shift σ is inside the spectrum of A relative to M . Preconditioners for SYMMLQ have to be positive definite.

The accuracy to which the linear system in (11) is solved has to be at least as high as the desired accuracy in the eigenvalue calculation in order that the coefficients of the Lanczos three term recurrence are sufficiently accurate [15]. In our experiments, when we computed the transformed eigenvalues to an accuracy of about ε , we set the convergence tolerance for the system solver to $\varepsilon/100$.

In [2] we experimented with ILU and SSOR preconditioners. In general, they reduced the number of iteration steps. But taking the execution times into account we obtained the best results with diagonal preconditioning. In [4] we study hierarchical basis preconditioners as discussed by Bank [6] and Brenner [7] for box-shaped cavities. Hierarchical basis preconditioners are most natural if finite elements of higher orders are employed. Let us assume that the basis elements $\{\Phi_i\}_{i=1}^n$ of $W_h^{(2)}$ are arranged according to the direct sum decomposition (6). Then the equation $(A - \sigma M)\mathbf{x} = \mathbf{y}$ can be written in the form

$$\begin{pmatrix} A_{11} - \sigma M_{11} & A_{12} - \sigma M_{12} \\ A_{21} - \sigma M_{21} & A_{22} - \sigma M_{22} \end{pmatrix} \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{pmatrix} = \begin{pmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{pmatrix} \quad (14)$$

where the order of the $(1, 1)$ block is $n_1 = \dim W_h^{(1)}$. Let

$$\begin{aligned} D &:= \begin{pmatrix} A_{11} - \sigma M_{11} & 0 \\ 0 & A_{22} - \sigma M_{22} \end{pmatrix}, \\ L &:= \begin{pmatrix} 0 & 0 \\ A_{21} - \sigma M_{21} & 0 \end{pmatrix}. \end{aligned} \quad (15)$$

n	n_1	t_C	diagonal		BJ-pj1		BJ-pj2		BJ-pgs1		BJ-pgs2	
			[msec]	t [sec]	\bar{n}_{it}	t [sec]	\bar{n}_{it}	t [sec]	\bar{n}_{it}	t [sec]	\bar{n}_{it}	t [sec]
17571	2019	14.1	286	187	232	89	272	65	317	64	453	62
29211	3387	27.8	515	189	348	74	447	61	530	60	740	59
39899	4625	45.5	788	194	506	71	612	59	731	58	1051	57
56715	6609	71.0	1188	202	755	71	915	58	1075	56	1527	55
71155	8287	90.0	1627	212	983	72	1112	57	1338	56	1862	55
93147	10887	128.0	2264	228	1374	75	1502	57	1744	55	2486	54

n	n_1	t_C	BGS-pj1		BGS-pj2		BGS-pgs1		BGS-pgs2	
			[msec]	t [sec]	\bar{n}_{it}	t [sec]	\bar{n}_{it}	t [sec]	\bar{n}_{it}	t [sec]
17571	2019	14.1	355	70	279	42	274	36	330	33
29211	3387	27.8	496	56	420	35	462	34	558	31
39899	4625	45.5	710	55	570	34	640	33	787	30
56715	6609	71.0	1084	55	871	34	927	32	1121	29
71155	8287	90.0	1441	57	1100	35	1127	31	1384	29
93147	10887	128.0	1978	59	1497	36	1487	31	1778	28

Table 1: Execution times in seconds and average iteration numbers of outer iteration for various preconditioners. BS (Block Jacobi) and BGS (Block Gauss-Seidel) denote the outer iteration method. pjm (Point Jacobi) and pgsm (Point Gauss-Seidel) denote the m -step interior iteration method.

Preconditioning $A - \sigma M$ with D or with $(D + L)D^{-1}(D + L^T)$, respectively, amounts to applying a block Jacobi or symmetric block Gauss-Seidel iteration to (14) [6]. We however do not properly apply these block preconditioners. We do solve systems involving $D_{11} = A_{11} - \sigma M_{11}$ directly with a multifrontal method [5] as the order of D_{11} is only $n_1 = \dim W_h^{(1)} \approx n/8$. The much larger systems involving $D_{22} = A_{22} - \sigma M_{22}$ are solved by m steps of a stationary iteration called a smoother, in our case either damped Jacobi or symmetric Gauss-Seidel (i.e. SSOR with $\omega = 1$). The damping factor in the Jacobi iteration was chosen to be $2/3$. Notice that with the Conrad-Wallach trick [18] m steps of symmetric Gauss-Seidel cost only as much as $m+1$ steps of the ordinary Gauss-Seidel iteration.

6 NUMERICAL EXPERIMENTS

We computed the 10 smallest eigenvalues of the eigenvalue problem (7) with a domain closely approximating the shape of the new cavity for the 590 MeV ring cyclotron at the Paul Scherrer Institute, cf. Fig. 1. We computed the eigenvalues to a relative accuracy $\varepsilon = 10^{-8}$. In (10) we chose the spectral shift such that $A - \sigma M$ is positive definite. The penalty parameter in (4) was set to $s = 1.5$. The computational results have been obtained on one processor of the PA-8000 processor-based HP Exemplar X-class system (180 MHz cycle, 720 MFlop/s peak) at ETH Zurich. We used the multiprocessor computer for its large memory. Results on a parallel implementation of the code are reported in [3].

In Tab. 1 we present timings for six problem sizes ranging from $n = \dim W_h^{(2)} = 17571$ up to $n = 93147$. n_1 denotes the dimension of the ‘coarse space’ $W_h^{(1)}$. t is the overall time to solve a problem; the time for the con-

struction of the matrices is not included. In each iteration of the (restarted) Lanczos algorithm a system of equations has to be solved. All the eigenvalue problems listed in Tab. 1 required 44 steps in the IRL algorithm, i.e. 44 calls to SYMMLQ. SYMMLQ in turn needed in the average \bar{n}_{it} iterations until convergence which means that $\bar{n}_{it} + 3$ linear systems of equations have been solved with the preconditioner [19]. The column in Tab. 1 labelled by t_C shows the times for solving a system with the matrix D_{11} in (15). Therefore, the time consumed by all coarse grid corrections is about $44(\bar{n}_{it} + 3)t_C$. This is 25-32% with BJ-pj1 and 9-13% with BGS-pgs1. The percentage increases with the problem size.

Diagonal preconditioning is satisfactory only for problem sizes smaller than those given in Tab. 1, cf. [4]. As the number of iteration steps increases with the problem size the performance of diagonal preconditioning deteriorates with increasing n .

As predicted by the theory, the iteration counts of the two-level methods are hardly affected by the problem size. In all problem sizes we observe that increasing the number m of smoothing steps lowers \bar{n}_{it} however not to the extent that the additional computing time is outweighed. In our examples, block Jacobi smoothed by one step of damped point Jacobi (BJ-pj1) performed best in all problem sizes. This may be surprising as this combination gave the highest average iteration count of all two-level methods. However, as we start the inner iteration with the zero vector, with Jacobi, the first smoothing step is just a multiplication of a diagonal matrix with a vector. Further steps of the Jacobi smoother as well as all steps of the Gauss-Seidel smoother require the access of two sparse triangular matrices, an operation which has a low flop count per memory access and thus performs at a low Mflop/s rate. There-

fore, the timings obtained with 2-step Jacobi and in particular with the Gauss-Seidel smoothers are higher than with BJ-pj1 although the average iteration count is significantly lower. If we should succeed in increasing the performance of the sparse matrix-vector multiplication, block Gauss-Seidel smoothed with one step of point Gauss-Seidel can be expected to perform best. Notice that because of the symmetry of $A - \sigma M$ the computation of the residuals in SYMMLQ has a doubled flop count per memory reference.

Our results show that hierarchical preconditioners are well suited for the systems of equations that have to be solved in IRL. The number of iteration steps is independent of the problem size. The experiences made with point Jacobi indicates that a good smoother does not need to access much of $A - \sigma M$ to be effective. An overlapping Schwarz smoother with small subdomains may therefore help to further improve the convergence rate. The time for the coarse grid correction could be reduced by moving from the two-level to a multilevel algorithm. With larger problem sizes the coarse grid correction may dominate the solution time and memory consumption of the eigensolver. We therefore will in future work investigate more economic ways to solve the coarse grid problem.

7 ACKNOWLEDGEMENT

We would like to thank Roman Geus of the Institute of Scientific Computing at ETH Zurich for performing the numerical experiments.

8 REFERENCES

- [1] St. Adam, P. Arbenz, and R. Geus, *Eigenvalue solvers for electromagnetic fields in cavities*, Technical Report #275, ETH Zürich, Institute of Scientific Computing, Computer Science Department, October 1997, (Available at URL <http://www.inf.ethz.ch/publications/>)
- [2] P. Arbenz and R. Geus, *A comparison of solvers for large eigenvalue problems originating from Maxwell's equations*, Numer. Linear Algebra Appl. **5** (1998).
- [3] ———, *Parallel solvers for large eigenvalue problems originating from Maxwell's equations*, Euro-Par '98 Parallel Processing (D. Pritchard and J. Reeve, eds.), Springer-Verlag, 1998, (Lecture Notes in Computer Science, 1470).
- [4] ———, *Two-level hierarchical basis preconditioners for computing eigenfrequencies of cavity resonators with the finite element method*, Recent Advances in Numerical Methods and Applications. Proceedings of the 4th International Conference on Numerical Methods and Applications (O. Iliev, M. Kaschiev, S. Margenov, Bl. Sendov, and P. Vassilevski, eds.), World Scientific, Singapore, 1999.
- [5] C. Ashcraft and R. Grimes, *The influence of relaxed supernode partitions on the multifrontal method*, ACM Trans. Math. Softw. **15** (1989), 291–309.
- [6] R. E. Bank, *Hierarchical bases and the finite element method*, Acta Numerica **5** (1996), 1–43.
- [7] S. C. Brenner, *Preconditioning complicated finite elements by simple finite elements*, SIAM J. Sci. Comput. **17** (1996), 1269–1274.
- [8] D. Calvetti, L. Reichel, and D. C. Sorensen, *An implicitly restarted Lanczos method for large symmetric eigenvalue problems*, Electronic Trans. Numer. Anal. **2** (1994), 1–21.
- [9] Ph. G. Ciarlet, *The finite element method for elliptic problems*, North-Holland, Amsterdam, 1978, (Studies in Mathematics and its Applications, 4).
- [10] D. R. Fokkema, G. L. G. Sleijpen, and H. A. van der Vorst, *Jacobi-Davidson style QR and QZ algorithms for the partial reduction of matrix pencils*, SIAM J. Sci. Comput. **20** (1999), 94–125.
- [11] V. Girault and P.-A. Raviart, *Finite element methods for the Navier-Stokes equations*, Springer-Verlag, Berlin, 1986, (Springer Series in Computational Mathematics, 5).
- [12] R. Grimes, J. G. Lewis, and H. Simon, *A shifted block Lanczos algorithm for solving sparse symmetric generalized eigenproblems*, SIAM J. Matrix Anal. Appl. **15** (1994), 228–272.
- [13] J. D. Jackson, *Classical electrodynamics*, 2nd ed., Wiley, New York, 1975.
- [14] F. Kikuchi, *Mixed and penalty formulations for finite element analysis of an eigenvalue problem in electromagnetism*, Comput. Methods Appl. Mech. Eng. **64** (1987), 509–521.
- [15] R. B. Lehoucq, D. C. Sorensen, and C. Yang, *ARPACK users' guide: Solution of large-scale eigenvalue problems by implicitly restarted Arnoldi methods*, SIAM, Philadelphia, PA, 1998, (The software and this manual are available at URL <http://www.caam.rice.edu/>).
- [16] R. Leis, *Zur Theorie elektromagnetischer Schwingungen in anisotropen Medien*, Math. Z. **106** (1968), 213–224.
- [17] K. Meerbergen and D. Roose, *Matrix transformations for computing rightmost eigenvalues of large sparse non-symmetric eigenvalue problems*, IMA J. Numer. Anal. **16** (1996), 297–346.
- [18] J. Ortega, *Introduction to parallel and vector solution of linear systems*, Plenum Press, New York, 1998.
- [19] C. C. Paige and M. A. Saunders, *Solution of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal. **12** (1975), 617–629.
- [20] B. N. Parlett, *The symmetric eigenvalue problem*, Prentice Hall, Englewood Cliffs, NJ, 1980.
- [21] G. L. G. Sleijpen and H. A. van der Vorst, *A Jacobi-Davidson iteration method for linear eigenvalue problems*, SIAM J. Matrix Anal. Appl. **17** (1995), 401–425.
- [22] D. Sorensen, *Implicit application of polynomial filters in a k-step Arnoldi method*, SIAM J. Matrix Anal. Appl. **13** (1992), 357–385.
- [23] H. Yserentant, *Hierarchical bases*, ICIAM 91: Proceedings of the Second International Conference on Industrial and Applied Mathematics (R. E. O'Malley, ed.), SIAM, Philadelphia, PA, 1992, pp. 256–276.