Available online at www.sciencedirect.com
ScienceDirect
Journal of Computational Physics 226 (2007) 1196–1210
JOURNAL OF COMPUTATIONAL PHYSICS
www.elsevier.com/locate/jcp

An algebraic multigrid based shifted-Laplacian preconditioner for the Helmholtz equation

Tuomas Airaksinen [a,*], Erkki Heikkola [b], Anssi Pennanen [a], Jari Toivanen [c]

[a] *Department of Mathematical Information Technology, University of Jyväskylä, P.O. Box 35 (Agora), FI-40014 University of Jyväskylä, Finland*
[b] *Numerola Oy, P.O. Box 126, FI-40101 Jyväskylä, Finland*
[c] *Institute for Computational and Mathematical Engineering, Building 500, Stanford University, Stanford, CA 94305, USA*

Received 26 January 2007; received in revised form 11 May 2007; accepted 14 May 2007
Available online 26 May 2007

Abstract

A preconditioner defined by an algebraic multigrid cycle for a damped Helmholtz operator is proposed for the Helmholtz equation. This approach is well suited for acoustic scattering problems in complicated computational domains and with varying material properties. The spectral properties of the preconditioned systems and the convergence of the GMRES method are studied with linear, quadratic, and cubic finite element discretizations. Numerical experiments are performed with two-dimensional problems describing acoustic scattering in a cross-section of a car cabin and in a layered medium. Asymptotically the number of iterations grows linearly with respect to the frequency while for lower frequencies the growth is milder. The proposed preconditioner is particularly effective for low-frequency and mid-frequency problems.

Keywords: Algebraic multigrid method; Finite element method; GMRES; Helmholtz equation; Preconditioner

1. Introduction

Acoustic scattering problems have applications in many disciplines. These problems can be typically modeled using wave equation and often it is sufficient to consider only time-harmonic solutions which are described by the Helmholtz equation (reduced wave equation). For numerical simulation the equations can be discretized using the finite difference method or the finite element method, for example. The solution of resulting systems of linear equations can be a computationally challenging problem.

During the past few decades, numerical methods for acoustics have been under active research. Finite element method has emerged as a generic tool for discretizing the Helmholtz equation in complex geometries. A recent review [1] offers a glance at research efforts in this field. The efficiency of these methods still often limits the feasible size of scattering problems in mid-frequency and high-frequency regime. Particularly the phase shift (pollution)

* Corresponding author. Tel.: +358 14 260 2743; fax: +358 14 260 2771.
E-mail address: tuomas.airaksinen@jyu.fi (T. Airaksinen).

0021-9991/$ - see front matter
doi:10.1016/j.jcp.2007.05.013

error in discretizations necessitates finer meshes for high-frequency problems [2]. The finite element method have been used successfully for interior problems like scattering in a car cabin [3] as well as for exterior problems. Since the paper [4] the research on the construction of absorbing boundary conditions and absorbing layers at the truncation boundary of the exterior domain has been active; see [1] and references therein.

The resulting systems of linear equations from the discretization of the Helmholtz equation are non-Hermitian and indefinite. Furthermore, for mid-frequency and high-frequency problems, the systems can be extremely large. These reasons make them a challenge for the current solvers. Often it is feasible to use direct methods for solving these systems for two-dimensional problems, but three-dimensional problems lead to systems which cannot be solved by these methods with affordable computing effort. Hence, it is necessity to use iterative methods such as the GMRES method [5] or the BICGSTAB method [6]. However, these methods require a good preconditioner for the discretized Helmholtz equations in order to have reasonably fast convergence.

Various preconditioners and iterative solution techniques have been proposed for the discrete Helmholtz equation. Several domain decomposition methods have been proposed; see [7–11], for example. Multigrid methods have been considered in [12–14]. With multigrid methods, it is difficult to define a stable and sufficiently accurate coarse grid problems and smoothers for them. For problems in homogenous medium, domain imbedding/fictitious domain methods in [15–17] have been fairly effective. An incomplete factorization preconditioner has been considered in [18], for example, and in [19] a tensor product preconditioner is used. An alternative iterative approach for solving the Helmholtz equation has been proposed in [20] and further studied in [21]. The basic idea is to find a time-periodic solution to the wave equations using a controllability method, which leads to preconditioned conjugate gradient iterations for initial data.

In this paper, we consider shifted-Laplacian preconditioners which are obtained from the Helmholtz operator by adding damping. A recent review [22] gives a glance at this class of preconditioners utilizing the properties of the differential equation rather than the algebraic form of the discretized problem. The Laplace operator was proposed as a preconditioner for the Helmholtz equation in [23]. A shifted-Laplacian preconditioner obtained by changing the sign of the zeroth-order term in the Helmholtz operator was described in [24]. As a generalization the Laplacian with a complex shift was studied in [25]. Following this approach a multigrid preconditioner based on a damped Helmholtz operator was considered in [14]. There, the scattering problems were posed in a rectangular domain, they were discretized using low-order finite differences, and a geometric multigrid method was used. In [26], this preconditioner was used with a finite element discretization and in [27] with an absorbing perfectly matched layer. This paper extends this approach for general shaped domains using linear, quadratic, and cubic finite element discretizations. Particularly quadratic and cubic finite elements help to reduce the number of unknowns in order to reach prescribed accuracy, as they have much smaller interpolation and phase shift errors than linear basis functions [2]. Our preconditioner is based on an algebraic multigrid method which can be constructed fully algebraically when the matrix for the zeroth-order terms is also available.

There is a wide range of applications for acoustic scattering in the industry and sciences. In many applications the aim is to reduce the noise level. As an example of such a problem we consider the noise in a car cabin; see also [3]. Geophysical surveys employ acoustic/elastic backscattering from different layers to reconstruct a model for the subsurface. These problems lead to very large-scale scattering problems. We consider a three layer wedge model [19,14] in our numerical experiments. Acoustic scattering simulations have also many applications in medicine, sonar, and sound preproduction, for example.

This paper is organized as follows. In Section 2 we describe the Helmholtz model problem and its discretization. The iterative solution and preconditioning based on shifted-Laplacian preconditioners are discussed in Section 3. The algebraic multigrid method employed in the preconditioning is described in Section 4. Then numerical results are presented in Section 5 and finally, conclusions are given in Section 6.

2. Scattering problem and finite element discretization

Under suitable assumptions on medium, acoustic scattering can be described by the wave equation

$$c^2 \nabla \cdot \frac{1}{\rho} \nabla p - \frac{1}{\rho} \frac{\partial^2 p}{\partial t^2} = 0, \tag{1}$$

where $p(\mathbf{x}, t)$ is pressure field, $\rho(\mathbf{x})$ is the density of the material, $c(\mathbf{x})$ is the speed of sound and t is time. For a time-harmonic pressure $p(\mathbf{x}, t) = u(\mathbf{x})e^{-i\omega t}$, where ω is angular velocity and $i = \sqrt{-1}$, (1) leads to the Helmholtz equation

$$-\nabla \cdot \frac{1}{\rho}\nabla u - \frac{k^2}{\rho}u = 0, \tag{2}$$

where $k(\mathbf{x}) = \omega/c(\mathbf{x})$ is the wave number. In inhomogeneous medium the wave number k varies depending on location as the sound speed c varies.

We consider three types of boundary conditions. In order to describe them, we decompose the boundary $\Gamma = \partial\Omega$ into three non-overlapping parts Γ_d, Γ_s, and Γ_a such that $\Gamma = \Gamma_d \cup \Gamma_s \cup \Gamma_a$. Some of these boundary sets can be empty. The first type of boundary is sound-soft which is described by the Dirichlet boundary condition

$$u = g(\mathbf{x}) \quad \text{on } \Gamma_d, \tag{3}$$

where g describes the sound source, for example, an incident field. The second type is the impedance boundary condition

$$\frac{\partial u}{\partial \mathbf{n}} = i\gamma k u \quad \text{on } \Gamma_a, \tag{4}$$

where $\gamma(\mathbf{x})$ is an absorbency coefficient in the range $[0, 1]$ describing the amount of absorption on the boundary Γ_a. The specific case $\gamma = 0$, leading to a Neumann boundary condition, corresponds to a sound-hard boundary without any absorption.

Exterior problems are truncated into a bounded domain Ω with Γ_s as the truncation boundary. The boundary condition on Γ_s should let outgoing waves propagate out of the domain without any reflection, as the Sommerfeld radiation condition describes. Such a perfect absorbing boundary condition is a non-local operator which is computationally difficult. Instead, it is usual to approximate it by a local operator [28,29,4]. Here we use an absorbing boundary condition

$$\frac{\partial u}{\partial \mathbf{n}} = iku \quad \text{on } \Gamma_s. \tag{5}$$

The methods studied here can be also used with higher-order absorbing boundary conditions.

For the weak formulation of the Helmholtz equation, we define the test function space

$$V = \{v \in H^1(\Omega) : v = 0 \text{ on } \Gamma_d\} \tag{6}$$

and the solution space

$$V_g = \{v \in H^1(\Omega) : v = g(\mathbf{x}) \text{ on } \Gamma_d\}. \tag{7}$$

Now the weak form of (2) reads: Find $u \in V_g$ such that

$$\int_\Omega \frac{1}{\rho}(\nabla u \cdot \nabla v - k^2 uv)\,dx - \int_{\Gamma_a} \frac{1}{\rho}(i\gamma kuv)\,ds - \int_{\Gamma_s} \frac{1}{\rho}(ikuv)\,ds = 0 \tag{8}$$

for all $v \in V$.

For a finite element discretization, we define a triangulation given by a set of non-overlapping triangles K_h, such that $\Omega_h = \bigcup_{\tau \in K_h} \tau$. Here h denotes the diameter of the largest triangle and Ω_h is an approximation of Ω. An example of a coarse triangulation (also called mesh) for a cross-section of a car cabin is shown in Fig. 1. For the finite elements of order m a discrete test function space is

$$V_h = \{v_h \in H^1(\Omega_h) : v_h|_\tau \in P^m, \forall \tau \in K_h, v = 0 \text{ on } \Gamma_{d,h}\}, \tag{9}$$

where P^m denotes polynomials of order m. A discrete solution space $V_{g,h}$ is obtained similarly by approximating g on $\Gamma_{d,h}$ instead of zero. In this paper, we employ linear, quadratic, and cubic finite elements, that is, $m = 1$, 2, or 3.

In the following, we discuss briefly the errors in the finite element solutions and the influence of order of finite elements. The best approximation u_h^{opt} on V_h for u is given by

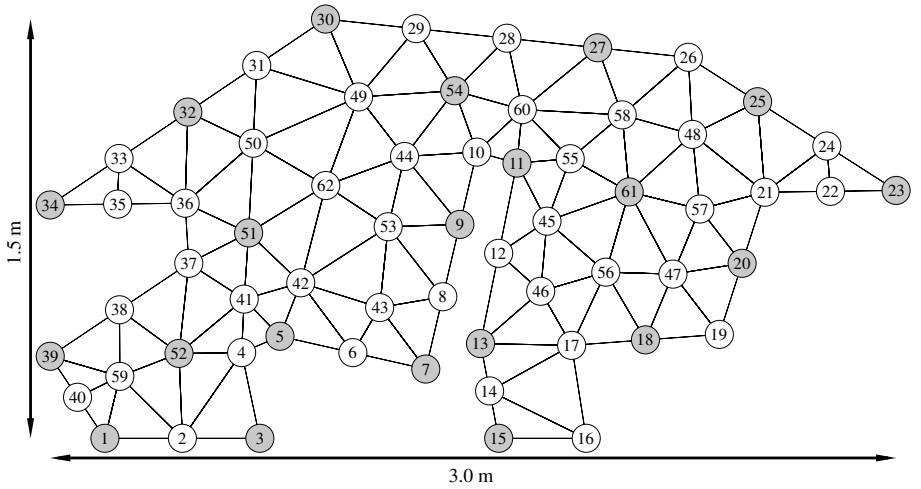

Fig. 1. A mesh for a cross-section of a car cabin. Gray nodes are selected to next coarse level.

$$u_h^{\text{opt}} = \underset{v_h \in V_h}{\operatorname{argmin}} \|v_h - u\|. \tag{10}$$

Here and in the following, we use the L^2-norm. According to error estimates [30,31,2] for the finite element solution u_h and the best approximation u_h^{opt}, we have

$$e_h = \|u_h - u\| \leqslant Ck(kh)^{2m} \quad \text{and} \quad \|u_h^{\text{opt}} - u\| \leqslant C(kh)^{2m}, \tag{11}$$

where C is a constant. The additional k in the right-hand side for the finite element error e_h is caused by the pollution error. In order to have the error e_h below some given $E < C$, the mesh step size h must satisfy

$$h \leqslant (C/E)^{2m} k^{-1} k^{-1/(2m)}. \tag{12}$$

The last term $k^{-1/(2m)}$ is due to the pollution error. Based on this inequality higher orders of the finite elements allow us to attain the same level of accuracy with larger mesh steps. Furthermore, it shows that the pollution error has reduced influence on the maximum mesh step size for larger m, that is, higher-order elements help to decrease the pollution error.

We use Lagrangian basis functions for the spaces V_h and $V_{g,h}$. Let the vector $\mathbf{u}$ contain the nodal values of u. Then using the discrete spaces instead of V and V_g in (8) and integrating over the discrete counterparts of the domain and boundaries, we obtain a system of linear equations

$$\mathbf{A}\mathbf{u} = \mathbf{f}, \tag{13}$$

where $\mathbf{A}$ is a sparse matrix and $\mathbf{f}$ is a non-zero vector due to the inhomogenous Dirichlet boundary condition. The approximation properties of such finite element discretizations for the Helmholtz equation have been studied in [2]. For an algebraic definition of the preconditioner described in Section 3, we define a mass matrix like $\mathbf{M}$ which includes the term k^2/ρ, that is, it corresponds to the integral

$$\int_{\Omega_h} \frac{k^2}{\rho} u_h v_h \, \mathrm{d}x, \tag{14}$$

where $u_h \in V_{g,h}$ and $v_h \in V_h$. Furthermore, we define the matrix $\mathbf{K} = \mathbf{A} + \mathbf{M}$ which contains the rest of the terms in the weak form.

3. Iterative solution and shifted-Laplacian preconditioner

The matrix $\mathbf{A}$ in (13) is indefinite and symmetric, but not Hermitian. Hence, the generalized minimal residual (GMRES) method [5] and the BICGSTAB method [6] are suitable iterative methods for the solu-

tion of the system (13). For these and other applicable iterative methods, see [32], for example. At each iteration, the GMRES method minimizes the norm of the residual vector on a Krylov subspace associated to the iteration. This is a desirable property leading to a monotonic reduction of residual norm over iterations, but a disadvantage is that a basis for the Krylov subspace needs to be formed and stored. Due to this the computational cost of the GMRES methods grows quadratically with iterations and the memory requirement grows linearly. With the BICGSTAB method the computational cost grows linearly and the memory requirement is independent of number of iterations, but convergence can be erratic and slower than with the GMRES method. In the numerical experiments we use the full GMRES method without restarts.

For medium- and large-scale scattering problems, the system (13) is badly conditioned, which leads to a very slow convergence of Krylov subspace methods when applied directly to the system (13). In order to improve the conditioning and the speed of convergence, we use a right preconditioner denoted by $\mathbf{B}$. This leads to a preconditioned system

$$\mathbf{A}\mathbf{B}^{-1}\tilde{\mathbf{u}} = \mathbf{f}. \tag{15}$$

Once $\tilde{\mathbf{u}}$ is solved from this system, the solution $\mathbf{u}$ is obtained as $\mathbf{u} = \mathbf{B}^{-1}\tilde{\mathbf{u}}$. Our aim is to find such a preconditioner $\mathbf{B}$ that the matrix $\mathbf{A}\mathbf{B}^{-1}$ is well conditioned and that vectors can be multiplied by $\mathbf{B}^{-1}$, that is, solve systems with $\mathbf{B}$, with a small computational effort. These properties would lead to a fast convergence of the iterative method and to a small overall computational cost.

In 2004, Erlangga et al. suggested in [25] to construct a preconditioner $\mathbf{B}_{\mathrm{SL}}$ by discretizing a shifted-Laplace operator

$$\mathscr{B}_{\mathrm{SL}} = -\nabla\cdot\frac{1}{\rho}\nabla - (\beta_1 + \beta_2\mathrm{i})\frac{k^2}{\rho}, \tag{16}$$

where we have added the density $\rho(\mathbf{x})$ in the operator. Using the notations defined in Section 2, the preconditioner can be defined algebraically as

$$\mathbf{B}_{\mathrm{SL}} = \mathbf{K} - (\beta_1 + \beta_2\mathrm{i})\mathbf{M} \tag{17}$$

which includes boundary term in $\mathbf{K}$.

By choosing $\beta_1 = 1$ and β_2 to be positive, $\mathscr{B}_{\mathrm{SL}}$ is the original Helmholtz operator with some additional damping. Such damping leads to good conditioning of $\mathbf{A}\mathbf{B}_{\mathrm{SL}}^{-1}$ and it is easier to solve systems with $\mathbf{B}_{\mathrm{SL}}$ than with $\mathbf{A}$ [25]. In [14], Erlangga et al. approximated the inverse of the shifted-Laplacian preconditioner $\mathbf{B}_{\mathrm{SL}}$ using one cycle of a geometric multigrid method; see [33], for example. We denote such multigrid based preconditioners by $\mathbf{B}_{\mathrm{MG}}$. This leads to a good conditioning of $\mathbf{A}\mathbf{B}_{\mathrm{MG}}^{-1}$ for low-frequency problems, while the number of BICGSTAB iterations appeared to grow linearly with frequency for high-frequency problems. They also showed that this preconditioner is well suited for problems with a highly varying speed of sound. In this paper, we replace the geometric multigrid method with a more generic algebraic multigrid method described in Section 4.

For the GMRES method, convergence estimates can be derived based on the spectrum of a matrix and its non-normality [5,32]. Similarly to [25,14], we study numerically the spectrum of the preconditioned matrices. For small problems, it is possible to compute the spectrum, while for larger problems we can only approximate it. The GMRES method forms the basis for a Krylov subspace using the Arnoldi iteration. After m iterations it has generated an $m \times m$ upper Hessenberg matrix which is usually denoted by $\mathbf{H}_m$. The eigenvalues of $\mathbf{H}_m$ approximate the eigenvalues of the system matrix.

4. Algebraic multigrid method

The preconditioner $\mathbf{B}_{\mathrm{MG}}$ is based on an algebraic multigrid (AMG) method which approximates the multiplication by the inverse of $\mathbf{B}_{\mathrm{SL}}$. We use an AMG method introduced by Kickinger in [34] with modifications proposed in [35]. This method uses a graph based on the system matrix to construct coarse spaces. Furthermore, it eliminates the degrees of freedom associated to the Dirichlet boundaries after forming the matrices for the

coarse spaces. Under these choices, the AMG method can be constructed in such a way that the coarse problems coincide with the ones obtained using a geometric multigrid method on a hierarchical linear finite element mesh.

The AMG initialization procedure is described by Algorithm 1. For linear finite elements, the initial graph G_0 is the graph defined by the sparse matrix $\mathbf{B}_{\mathrm{SL}}$. Alternatively it can be seen as the graph defined by the triangulation. For quadratic and cubic elements, the graph is defined by a refined triangulation in which quadratic elements are divided into four triangles, and cubic elements are divided into nine triangles. The reason not to use directly the graph defined by $\mathbf{B}_{\mathrm{SL}}$ for higher-order elements is that the coarsening procedure would coarsen the graph too much, leading to a slower convergence, that is, to not so well conditioned $\mathbf{A}\mathbf{B}_{\mathrm{SL}}^{-1}$. The nodes (vertices) in the graph G_0 associated with the Dirichlet boundaries are marked. On the coarser graphs also the nodes which were marked as Dirichlet nodes on the finer graphs are marked.

The nodes onto a coarser graph G_{k+1} are chosen from the nodes of G_k as follows. Find the node in G_k which has the smallest degree, that is, the smallest number of edges associated to it. If there are several such nodes, choose the first one according to the used node numbering. This node is included onto the graph G_{k+1}. Eliminate this node and all its neighbors from the graph G_k. There is one exception to this: if the node has a Dirichlet marked neighbor then the neighbor is not eliminated from G_k. This increases the stability of the procedure by making sure that there are sufficiently many Dirichlet nodes selected to the coarse levels. Repeat this procedure until there are no nodes left in G_k. Fig. 1 shows an example of this coarsening strategy when all boundaries are of Dirichlet type.

After choosing the nodes on G_{k+1}, they are numbered following their order on G_k. Then the restriction matrix $\mathbf{R}_k$ is defined by

$$(\mathbf{R}_k)_{ij} = \begin{cases} 1 & \text{for a fine node } j \text{ which is a coarse node } i, \\ \frac{1}{k} & \text{for a fine node } j \text{ which is a neighbor of coarse,} \\ & \text{node } i \text{ and has } k \text{ neighboring coarse nodes,} \\ 0 & \text{otherwise,} \end{cases}$$

where fine and coarse refers to the graphs G_k and G_{k+1}, respectively. The edges of the coarse graph G_{k+1} are formed using the restriction matrix $\mathbf{R}_k$. Each coarse graph node corresponds to a row in the restriction matrix and there is an edge between two nodes if and only if the corresponding rows of the restriction matrix have a non-zero element in the same column.

The prolongation matrix is the transpose of the restriction matrix, and a coarser grid system matrix is constructed by Galerkin method, that is, the fine grid matrix is multiplied by R_k from the left side and by R_k^{T} from the right side. The AMG cycle described by Algorithm 2 is a usual multigrid cycle given here in the general μ-cycle form. The choices $\mu = 1$ and $\mu = 2$ correspond to V-cycle and W-cycle, respectively. When the algorithm is used in preconditioning it is called with the approximate solution $\mathbf{x}_0$ being zero.

Algorithm 1. AMG initialization

Input: Matrix $\mathbf{B}_0$, initial graph G_0, the maximum size of the coarsest system n_c
1. $k = 0$
2. **Do while** the size of $\mathbf{B}_k$ is greater than n_c
3. Select the set of coarse nodes from the graph G_k
4. Form the restriction matrix $\mathbf{R}_k$
5. Create the graph G_{k+1}
6. Calculate the next system matrix $\mathbf{B}_{k+1} = \mathbf{R}_k\mathbf{B}_k\mathbf{R}_k^T$
7. Eliminate the rows and columns of $\mathbf{B}_k$ marked in the graph G_k
8. Eliminate the columns of $\mathbf{R}_k$ marked in the graph G_k
9. Eliminate the rows of $\mathbf{R}_k$ marked in the graph G_{k+1}
10. k=k+1
11. **End do**
12. Eliminate the rows and columns of $\mathbf{B}_k$ marked in the graph G_k
13. Factorize $\mathbf{B}_k$

Algorithm 2. Recursive algorithm for the AMG cycle

Input: Matrix $\mathbf{B}_l$, approximate solution $\mathbf{x}_l$, right-hand side vector $\mathbf{f}_l$
Output: Improved approximate solution $\mathbf{x}_l$
1. **If** on the coarsest level, that is, $l = k$
2. Solve $\mathbf{x}_l$ from $\mathbf{B}_l\mathbf{x}_l = \mathbf{f}_l$
3. **Else**
4. Presmooth $\mathbf{x}_l = \mathbf{x}_l + \mathbf{S}_l(\mathbf{x}_l,\mathbf{f}_l)$
5. Restrict the residual $\mathbf{f}_{l+1} = \mathbf{R}_l(\mathbf{f}_l - \mathbf{B}_l\mathbf{x}_l)$
6. Set $\mathbf{x}_{l+1} = 0$ and call μ times the cycle for the next level $l + 1$
7. Prolong the correction $\mathbf{x}_l = \mathbf{x}_l + \mathbf{R}_l^{\mathrm{T}}\mathbf{x}_{l+1}$
8. Post-smooth $\mathbf{x}_l = \mathbf{x}_l + \mathbf{S}_l(\mathbf{x}_l,\mathbf{f}_l)$
9. **End if**

5. Numerical results

5.1. Model problems with homogenous medium

We use two different model geometries with homogenous medium: the unit square and a cross-section of a car cabin. Fig. 2 shows typical solutions for these geometries. The same problem in the unit square was considered also in [14]. It has a point source at the middle and the absorbing boundary condition (5) is posed on all boundaries. The car cabin problem with a non-convex geometry resembles more real-world applications. The height of the car cabin is 1.5 m and its width is 3 m. The noise source is modeled using the Dirichlet boundary condition (3) with $g = 1$ on the wall behind pedals and on other boundaries the impedance boundary condition (4) with $\gamma = 0.2$ is used. The meshes for the car cabin problem were generated using Netgen [36] by refining a coarse mesh depicted in Fig. 1.

5.1.1. Eigenvalue study

We study the eigenvalues for both problems by computing both the full spectrum and Arnoldi approximations discussed in the end of Section 3. First we consider the Helmholtz problem with $k = 20$ in the unit square domain discretized on a 31×31 structured mesh. Figs. 3 and 4 demonstrate the influence of β_2, that is, the amount of damping to the spectrum of $\mathbf{AB}_{\mathrm{SL}}^{-1}$. We use the values $\beta_2 = 0.5$ and $\beta_2 = 1.0$. In these and all following results we have $\beta_1 = 1$. Fig. 3 shows that with the Neumann boundary condition the real parts of the eigenvalues are between zero and one with both β_2s, while the density of eigenvalues near zero is higher with $\beta_2 = 1.0$. The differences are more pronounced with the absorbing boundary condition. With a smaller β_2, the matrix $\mathbf{AB}_{\mathrm{SL}}^{-1}$ is closer to identity and this is seen as tighter clustering around

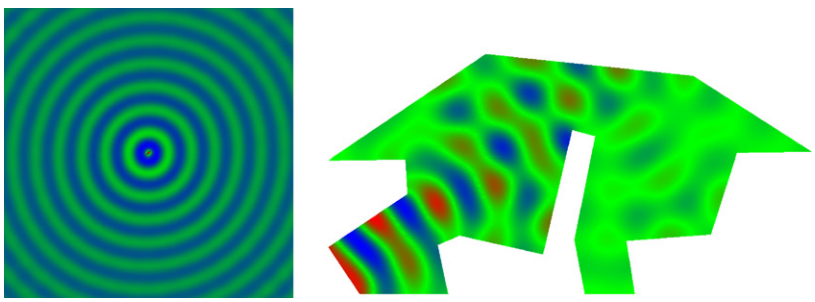

Fig. 2. In the left plot a solution for the unit square problem. In the right plot the solution for the car cabin problem with the wave number $k = 18.3$ which corresponds to the frequency $f \approx 1$ kHz.

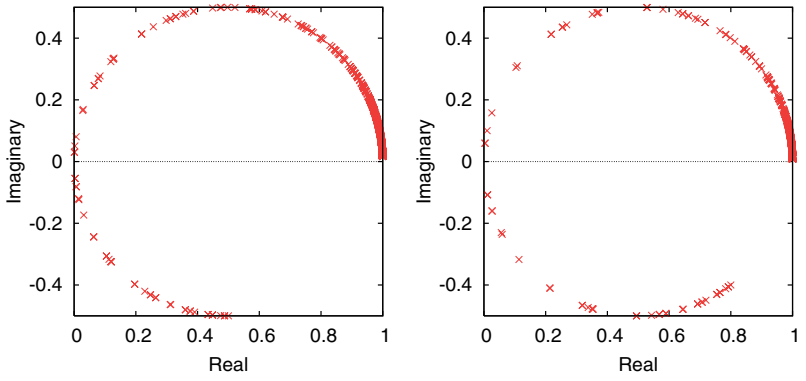

Fig. 3. The eigenvalues of $\mathbf{A}\mathbf{B}_{\mathrm{SL}}^{-1}$ with $\beta_2 = 1.0$ in the left plot and $\beta_2 = 0.5$ in the right plot for the unit square with the Neumann boundary condition discretized using linear finite elements.

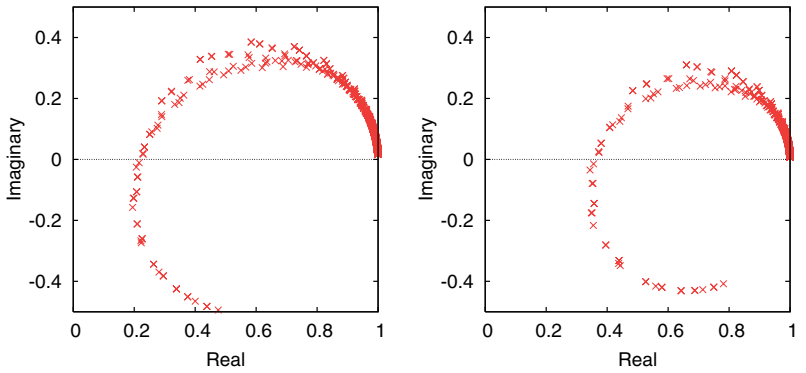

Fig. 4. The eigenvalues of $\mathbf{A}\mathbf{B}_{\mathrm{SL}}^{-1}$ with $\beta_2 = 1.0$ in the left plot and $\beta_2 = 0.5$ in the right plot for the unit square with the absorbing boundary condition discretized using linear finite elements.

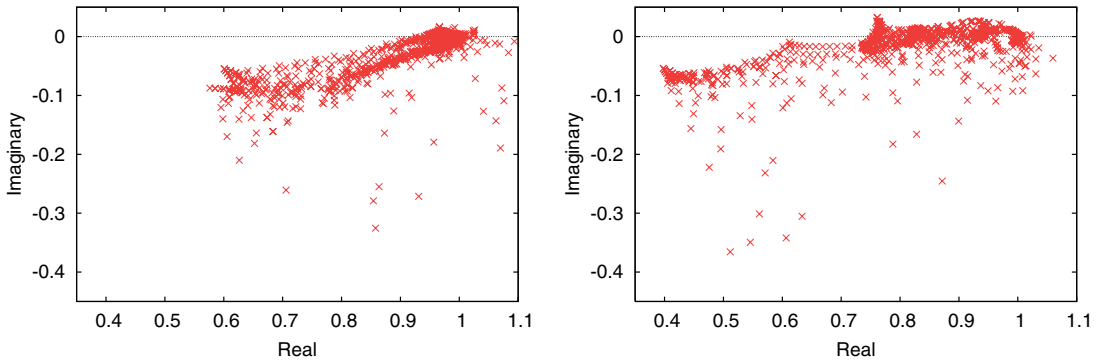

Fig. 5. The eigenvalues of $\mathbf{B}_{\mathrm{SL}}\mathbf{B}_{\mathrm{MG}}^{-1}$ for the unit square problem with the absorbing boundary condition. The left and right plots are computed using linear and quadratic elements, respectively.

one in Fig. 4. The spectrum of $\mathbf{A}\mathbf{B}_{\mathrm{SL}}^{-1}$ is very similar with quadratic and cubic elements to one with linear elements shown in these figures.

Next we study the AMG approximations of the inverse of the discrete shifted-Laplacians. We consider the quality of these approximations and the influence of order of finite elements. For this, we use one W-cycle

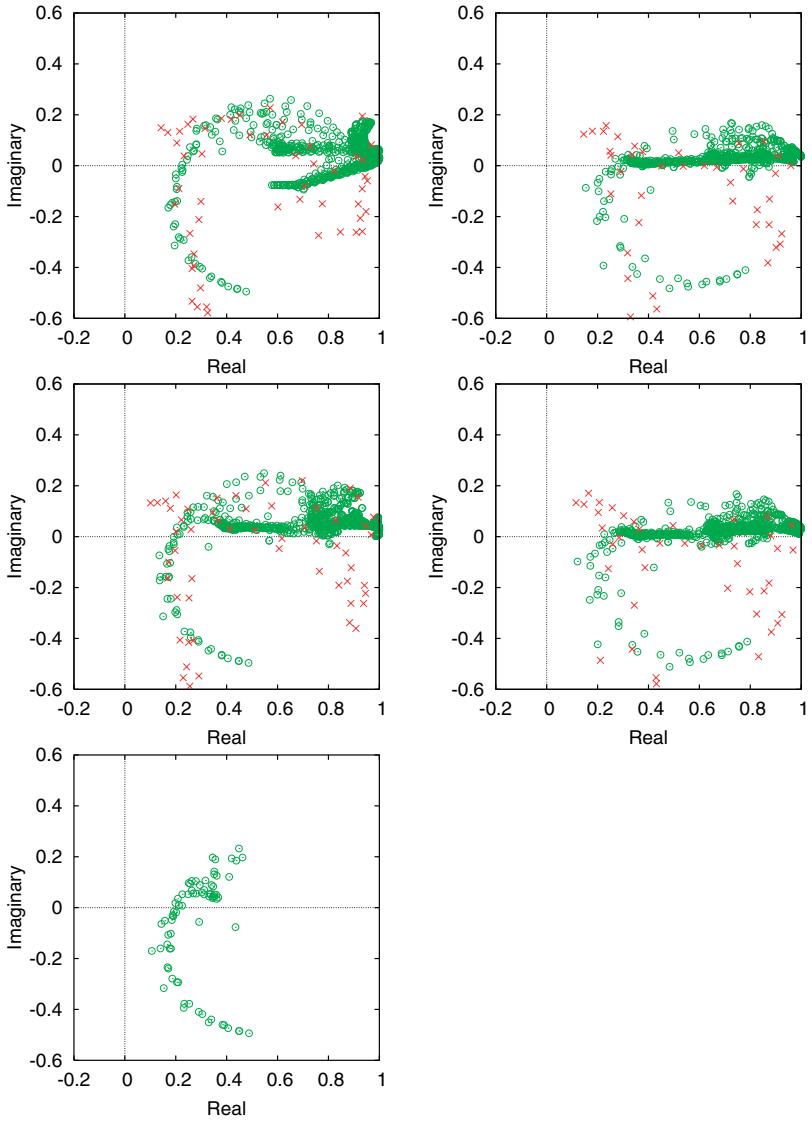

based on one presmoothing and postsmoothing iteration performed by the underrelaxed Jacobi with the relaxation parameter chosen according to Table 2. Fig. 5 plots the spectrums of $\mathbf{B}_{\mathrm{SL}}\mathbf{B}_{\mathrm{MG}}^{-1}$ for linear and quadratic finite elements with $\beta_2 = 1.0$. With linear elements these eigenvalues are more clustered around one than with quadratic elements which indicates that the AMG approximation of the inverse is better with linear elements.

Fig. 6 depicts the eigenvalues of $\mathbf{A}\mathbf{B}_{\mathrm{MG}}^{-1}$ with the absorbing boundary conditions for linear, quadratic and cubic finite elements when the parameter $\beta_2 = 0.5$ and 1.0. The eigenvalues are fairly similar for different elements and with $\beta_2 = 0.5$ some of them are closer to the real axis than with $\beta_2 = 1.0$. For the car cabin problem with $k = 15$, similar plots of eigenvalues are given in Fig. 7. The figures show that the spectrums are fairly similar for the unit square and car cabin problems. This suggest that the quality of the AMG preconditioner is not particularly sensitive to the geometry.

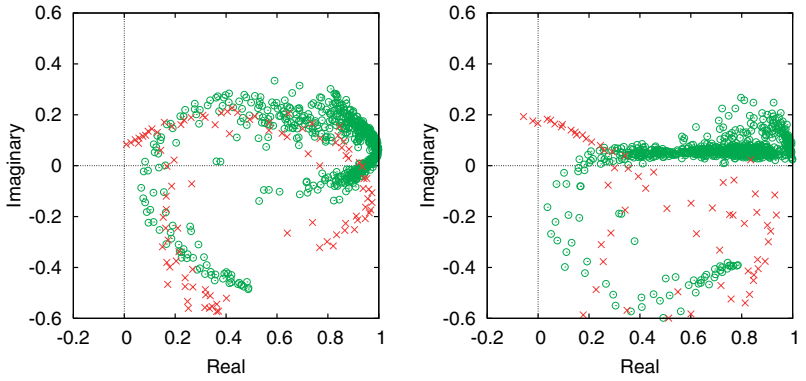

Fig. 7. For the car cabin problem discretized with linear elements the eigenvalues of $\mathbf{AB}_{\mathrm{MG}}^{-1}$ marked with ∘ and their Arnoldi approximations marked with ×. The left and right plots are based on $\beta_2 = 1.0$ and $\beta_2 = 0.5$, respectively.

5.1.2. Performance analysis

We mainly use the car cabin problem to study the performance of the iterative solver while we also report some results for the unit square problem. Fig. 2 shows usual time-harmonic scattering patterns for these problems. Table 1 describes the different meshes used with the car cabin problem.

The preconditioner $\mathbf{B}_{\mathrm{MG}}^{-1}$ is defined by one algebraic multigrid cycle described in Section 3. Our aim is to choose the parameters defining the preconditioner in such a way that the overall performance is optimal. We use the value $\beta_1 = 1$ while for β_2 we examine the values 0.5 and 1.0. There are several choices related to the AMG method. Our smoother is the underrelaxed Jacobi iteration with the relaxation parameter ω chosen according to Table 2. These relaxation parameters minimize the overall solution time in numerical tests. We use the W-cycle in the AMG method as it leads to shorter solution times than the V-cycle and F-cycle in experiments. The V-cycle led to five times and the F-cycle to slightly less than two times greater overall CPU consumption in our tests.

In Table 3, the number of iterations are reported for the unit square problem. We have chosen the wave numbers k and the mesh step sizes h to be the same as in [14]. With higher-order elements we could have used larger h without reducing accuracy. Their discretization was performed using a finite difference method and they employed a tuned geometric multigrid method instead of the AMG method used here. Furthermore, they used the BICGSTAB method with a slightly more strict stopping criterion. The number of iterations required here are higher. We need up to 1.5 times more iterations with $k = 40$ while the difference grows with k. Nevertheless the results in here and in [14] suggest that for higher frequencies the number of iterations roughly doubles when the frequency is doubled. The AMG preconditioner leads to particularly good results with cubic finite elements.

The convergence results for the car cabin problem are presented in Tables 4 and 5. In these tables, the wave number doubles from a row to the next and the mesh step size h is halved from a column to the next. Thus, the columns correspond to refined meshes in which the mesh step h is 2^{-n_r} times smaller than in the coarsest mesh. This leads to constant khs on diagonals. Along them, we can again observe that for higher frequencies the number of iterations roughly doubles when the wave number is doubled. The lower triangles of the tables cor-

Table 1
The number of elements and nodes in the car cabin meshes

	No. of refinements, n_r						
	0	1	2	3	4	5	6
Elements	82	328	1312	5248	20,992	83,962	335,872
Nodes, linear	62	205	737	2785	10,817	42,625	169,217
Nodes, quadratic	205	737	2785	10,817	42,625	169,217	674,305
Nodes, cubic	430	1597	6145	24,097	95,425	379,777	1,515,265

Table 2
The optimal choice of the Jacobi relaxation parameter ω for different finite elements and different values of β_2 found by extensive numerical experiments

β_2	ω		
	Linear	Quadratic	Cubic
0.5	0.4	0.4	0.4
1.0	0.8	0.7	0.7

Table 3
The number of iterations for the unit square problem for different element types when the iterations are terminated once the norm of the residual is reduced by the factor 10^{-6}

Element type	β_2	$k = 40$	$k = 50$	$k = 80$	$k = 100$	$k = 150$
Linear	1.0	43	51	76	93	137
	0.5	37	47	82	111	210
Quadratic	1.0	44	53	79	97	140
	0.5	32	41	70	95	177
Cubic	1.0	45	56	83	103	149
	0.5	31	37	59	77	126

Table 4
For the car cabin problem discretized with linear elements, the number of GMRES iterations required to reduce the norm of the residual by the factor 10^{-6} as a function of the number of refinements n_r and the wave number k

k	$\beta_2 = 1.0$							$\beta_2 = 0.5$						
	$n_r = 0$	$n_r = 1$	$n_r = 2$	$n_r = 3$	$n_r = 4$	$n_r = 5$	$n_r = 6$	$n_r = 0$	$n_r = 1$	$n_r = 2$	$n_r = 3$	$n_r = 4$	$n_r = 5$	$n_r = 6$
2	**11**	11	11	11	11	10	10	**8**	12	13	13	13	12	12
4	19	**19**	18	18	17	16	15	13	**14**	15	15	15	14	14
8	32	35	**34**	34	34	32	30	20	25	**22**	21	21	21	20
16	40	84	65	**63**	65	64	59	32	58	52	**43**	35	34	35
32	8	125	184	130	**127**	133	127	7	144	172	124	**91**	75	73
64	4	9	366	409	264	**262**	270	4	13	433	408	279	**194**	148

Table 5
For the car cabin problem discretized with quadratic (left table) and cubic elements (right table), the number of GMRES iterations required to reduce the norm of the residual by the factor 10^{-6} as a function of the number of refinements n_r and the wave number k

k	Quadratic elements							Cubic elements						
	$n_r = 0$	$n_r = 1$	$n_r = 2$	$n_r = 3$	$n_r = 4$	$n_r = 5$	$n_r = 6$	$n_r = 0$	$n_r = 1$	$n_r = 2$	$n_r = 3$	$n_r = 4$	$n_r = 5$	$n_r = 6$
2	**12**	13	13	14	13	12	12	**16**	15	15	15	15	15	15
4	15	**15**	15	15	14	14	13	20	**17**	17	16	16	16	16
8	27	23	**21**	21	21	20	19	33	25	**23**	22	22	21	21
16	61	59	45	**35**	34	35	33	77	56	46	**37**	36	38	34
32	172	190	141	95	**75**	71	66	193	195	123	87	**70**	70	69
64	17	496	464	315	199	**148**	149	250	> 500	459	234	161	**149**	137

In both tables, $\beta_2 = 0.5$.

respond to discretizations which do not have sufficiently high number of nodes per wavelength to capture the oscillatory behavior of solutions. This shows up as unusually high number of iterations. Based on these results, the value $\beta_2 = 0.5$ leads to much faster convergence on higher wave numbers.

With a large number of iterations, the time spent in forming the basis vectors in the GMRES method takes a larger part of the CPU time. This effect is seen in Table 6 which reports the time spent in different parts of the

Table 6
The time spent in the solver with the car cabin problem discretized using quadratic elements on the $n_r = 5$ mesh

Wave number k	GMRES iterations	CPU time in seconds	% of time spent in AMG
2	10	8	93.9
8	29	23	87.4
32	126	129	64.6

The AMG uses W-cycle and $\beta_2 = 1.0$.

Fig. 8. The plot on the left shows the mesh for the frequency $f = 5$ Hz and the definition of the problem. On the middle and right the solutions for $f = 30$ Hz and $f = 50$ Hz, respectively, are shown.

Table 7
For the wedge problem the number of GMRES iterations required to reduce the norm of the residual by the factor 10^{-6} as a function of the frequency

Element type	$f = 5$ Hz	$f = 30$ Hz	$f = 50$ Hz
Linear, $\beta_2 = 1.0$	24	97	148
Linear, $\beta_2 = 0.5$	17	83	124
Cubic, $\beta_2 = 1.0$	24	105	171
Cubic, $\beta_2 = 0.5$	19	62	97

solver for three problems discretized using quadratic elements; linear and cubic elements lead to similar results. The computations were performed on a PC with an 1.2 GHz Intel Core Duo U2500 processor.

5.2. Wedge problem with inhomogenous medium

The wedge problem is defined by three layers with different speed of sound c in the rectangle 600×1000 m^2, as shown in Fig. 8. This model problem was considered in [14,19]. Meshes for different frequencies were constructed with Comsol Multiphysics mesh generator in such a way that the mesh step size h was approximately one tenth of one wavelength, that is, $h \approx \lambda/10$. The wedge model has a point source at the middle of the top boundary. The absorbing boundary condition (5) is posed on all boundaries. A coarse mesh for the frequency $f = 5$ Hz and the solutions for the frequencies $f = 30$ Hz and $f = 50$ Hz are shown in Fig. 8.

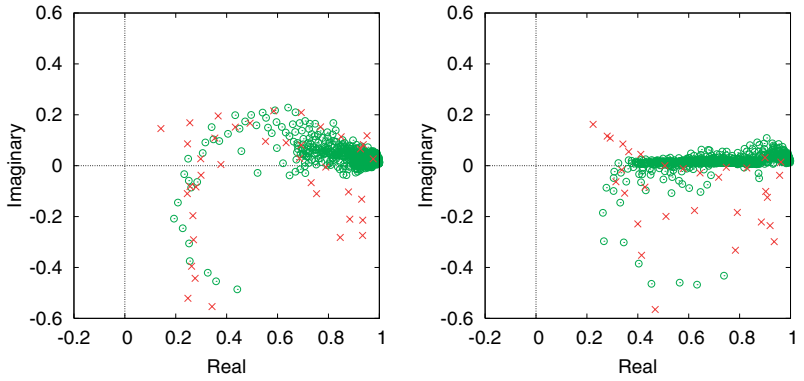

Fig. 9. The eigenvalues of $\mathbf{AB}_{\mathrm{MG}}^{-1}$ marked with ∘ and their Arnoldi approximations marked with × for the wedge problem with the frequency $f = 5$ Hz discretized using linear elements. The left and right plots are based on $\beta_2 = 1.0$ and $\beta_2 = 0.5$, respectively.

The performance results for different frequencies are presented in Table 7. According to these iteration counts, the convergence with $\beta_2 = 0.5$ is about 15% faster than with $\beta_2 = 1.0$ with linear elements and over 40% faster with cubic elements. We need again more iterations when compared to the results in [14], but our finite element discretizations have three advantages over the finite differences used in there: we can accurately model the interface between layers, we can use coarser meshes with higher-order elements, and we can use coarser mesh where the speed of sound is higher. Fig. 9 plots the eigenvalues of the preconditioned system for the low frequency $f = 5$ Hz with $\beta_2 = 1.0$ and $\beta_2 = 0.5$.

6. Conclusions

We have studied a preconditioner based on an algebraic multigrid (AMG) approximation of the inverse of a shifted-Laplacian for the Helmholtz equation. This is a generalization of the preconditioner proposed by Erlangga et al. in [14]. They used finite difference discretizations on rectangular domains and a geometrical multigrid. With our finite element discretizations and the AMG method we can solve problems in complicated domains and use higher-order finite elements. A big advantage of the AMG method is that the solver does not need hierarchical meshes nor operators discretized on different meshes. When the matrix for the zeroth-order term in a discretized Helmholtz equation or the mass matrix for a constant wave number problem is also available, the preconditioner can be constructed fully algebraically. Thus, in this case the preconditioned iteration can be seen as a "black box solver".

The numerical results demonstrated the capability to solve efficiently problems in complicated domains and varying wave numbers using the proposed preconditioner. Furthermore, the preconditioner was shown to be effective with linear, quadratic, and cubic finite elements. The proposed approach is especially well suited for low-frequency and mid-frequency problems while for high-frequency problems the number of iterations roughly doubles when the frequency is doubled. The same behavior was also observed in [14].

Acknowledgments

We thank Dr. Janne Martikainen for making his C++ FEM class library available to us and Prof. Tuomo Rossi for fruitful discussions. We also thank the referees for their comments which helped to improve the results and presentation. The research was supported by the Academy of Finland Grant #207089 and the US Office of Naval Research Grant N00014-06-1-0067.

References

[1] L.L. Thompson, A review of finite-element methods for time-harmonic acoustics, J. Acoust. Soc. Am. 119 (3) (2006) 1315–1330.
[2] F. Ihlenburg, Finite element analysis of acoustic scattering, Applied Mathematical Sciences, vol. 132, Springer-Verlag, New York, 1998.

[3] D.J. Nefske, J. Wolf, L. Howell, Structural-acoustic finite element analysis of the automobile passenger compartment: a review of current practice, J. Sound Vib. 80 (2) (1982) 247–266.
[4] B. Engquist, A. Majda, Absorbing boundary conditions for numerical simulation of waves, Math. Comput. 31 (1977) 629–651.
[5] Y. Saad, M.H. Schultz, GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems, SIAM J. Sci. Statist. Comput. 7 (3) (1986) 856–869.
[6] H.A. van der Vorst, Bi-CGSTAB: a fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems, SIAM J. Sci. Statist. Comput. 13 (2) (1992) 631–644.
[7] C. Farhat, A. Macedo, M. Lesoinne, F.-X. Roux, F. Magoulés, A. de La Bourdonnaie, Two-level domain decomposition methods with Lagrange multipliers for the fast iterative solution of acoustic scattering problems, Comput. Methods Appl. Mech. Eng. 184 (2–4) (2000) 213–239.
[8] C. Farhat, P. Avery, R. Tezaur, J. Li, FETI-DPH: a dual-primal domain decomposition method for acoustic scattering, J. Comput. Acoust. 13 (3) (2005) 499–524.
[9] M.J. Gander, F. Magoulès, F. Nataf, Optimized Schwarz methods without overlap for the Helmholtz equation, SIAM J. Sci. Comput. 24 (1) (2002) 38–60.
[10] K. Ito, J. Toivanen, A fast iterative solver for scattering by elastic objects in layered media, Appl. Numer. Math. 57 (5–7) (2007) 811–820.
[11] K. Otto, E. Larsson, Iterative solution of the Helmholtz equation by a second-order method, SIAM J. Matrix Anal. Appl. 21 (1) (1999) 209–229.
[12] A. Brandt, I. Livshits, Wave-ray multigrid method for standing wave equations, Electron. Trans. Numer. Anal. 6 (1997) 162–181.
[13] H.C. Elman, O.G. Ernst, D.P. O'Leary, A multigrid method enhanced by Krylov subspace iteration for discrete Helmholtz equations, SIAM J. Sci. Comput. 23 (4) (2001) 1291–1315.
[14] Y.A. Erlangga, C.W. Oosterlee, C. Vuik, A novel multigrid based preconditioner for heterogeneous Helmholtz problems, SIAM J. Sci. Comput. 27 (4) (2006) 1471–1492.
[15] O.G. Ernst, A finite-element capacitance matrix method for exterior Helmholtz problems, Numer. Math. 75 (2) (1996) 175–204.
[16] E. Heikkola, Y.A. Kuznetsov, P. Neittaanmäki, J. Toivanen, Fictitious domain methods for the numerical solution of two-dimensional scattering problems, J. Comput. Phys. 145 (1) (1998) 89–109.
[17] E. Heikkola, T. Rossi, J. Toivanen, A parallel fictitious domain method for the three-dimensional Helmholtz equation, SIAM J. Sci. Comput. 24 (5) (2003) 1567–1588.
[18] M.M. Monga Made, Incomplete factorization-based preconditionings for solving the Helmholtz equation, Int. J. Numer. Methods Eng. 50 (2001) 1077–1101.
[19] R.E. Plessix, W.A. Mulder, Separation-of-variables as a preconditioner for an iterative Helmholtz solver, Appl. Numer. Math. 44 (3) (2003) 385–400.
[20] M.O. Bristeau, R. Glowinski, J. Périaux, Controllability methods for the computation of time-periodic solutions; application to scattering, J. Comput. Phys. 147 (2) (1998) 265–292.
[21] E. Heikkola, S. Mönkölä, A. Pennanen, T. Rossi, Controllability method for acoustic scattering with spectral elements, J. Comput. Appl. Math. 204 (2) (2007) 344–355.
[22] E. Turkel, Numerical methods and nature, J. Sci. Comput. 28 (2–3) (2006) 549–570.
[23] A. Bayliss, C.I. Goldstein, E. Turkel, An iterative method for the Helmholtz equation, J. Comput. Phys. 49 (3) (1983) 443–457.
[24] A.L. Laird, M.B. Giles, Preconditioned iterative solution of the 2D Helmholtz equation, Tech. Rep. 02/12, Oxford Computer Laboratory, Oxford, UK, 2002.
[25] Y.A. Erlangga, C. Vuik, C.W. Oosterlee, On a class of preconditioners for solving the Helmholtz equation, Appl. Numer. Math. 50 (3–4) (2004) 409–425.
[26] E. Turkel, Y. Erlangga, Preconditioning a finite element solver of the exterior Helmholtz equation, in: P. Wesseling, E. Oñate, J. Périaux (Eds.), ECCOMAS CFD, ECCOMAS, 2006.
[27] Y. Erlangga, A preconditioner for the Helmholtz equation with perfectly matched layer, in: P. Wesseling, E. Oñate, J. Périaux (Eds.), ECCOMAS CFD, ECCOMAS, 2006.
[28] A. Bamberger, P. Joly, J.E. Roberts, Second-order absorbing boundary conditions for the wave equation: a solution for the corner problem, SIAM J. Numer. Anal. 27 (2) (1990) 323–352.
[29] A. Bayliss, M. Gunzburger, E. Turkel, Boundary conditions for the numerical solution of elliptic equations in exterior regions, SIAM J. Appl. Math. 42 (2) (1982) 430–451.
[30] I.M. Babuška, S.A. Sauter, Is the pollution effect of the FEM avoidable for the Helmholtz equation considering high wave numbers? SIAM J. Numer. Anal. 34 (6) (1997) 2392–2423.
[31] A. Bayliss, C.I. Goldstein, E. Turkel, On accuracy conditions for the numerical computation of waves, J. Comput. Phys. 59 (3) (1985) 396–404.
[32] Y. Saad, Iterative Methods for Sparse Linear Systems, second ed., Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2003.
[33] W.L. Briggs, V.E. Henson, S.F. McCormick, A Multigrid Tutorial, second ed., Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2000.
[34] F. Kickinger, Algebraic multi-grid for discrete elliptic second-order problems, in: Multigrid Methods V (Stuttgart, 1996), Springer, Berlin, 1998, pp. 157–172.

[35] J. Martikainen, A. Pennanen, T. Rossi, Application of an algebraic multigrid method to incompressible flow problems, Tech. Rep. B2/2006, Department of Mathematical Information Technology, University of Jyväskylä, Jyväskylä, Finland, 2006.
[36] J. Schöberl, NETGEN – an advancing front 2D/3D-mesh generator based on abstract rules, Comput. Visual. Sci. 1 (1997) 41–52.