Shifted-Laplacian Preconditioners for Heterogeneous Helmholtz Problems

C.W. Oosterlee, C. Vuik, W.A. Mulder, and R.-E. Plessix

Abstract We present an iterative solution method for the discrete high wavenumber Helmholtz equation. The basic idea of the solution method, already presented in [18], is to develop a preconditioner which is based on a Helmholtz operator with a complex-valued shift, for a Krylov subspace iterative method. The preconditioner, which can be seen as a strongly damped wave equation in Fourier space, can be approximately inverted by a multigrid method.

1 Introduction

The efficient numerical solution of the Helmholtz equation with spatially-dependent high wavenumbers is a difficult task. The prescription of reasonable boundary conditions, on a domain which is truncated for computational reasons, further complicates the numerical treatment in a real-life setting. A thorough overview of the issues regarding the numerical solution of the Helmholtz equation was presented in [56].

In a series of papers, [16, 17, 18, 47, 48, 20, 44, 57], supported by Royal Dutch Shell and Philips via the Dutch Ministry of Economic Affairs, project BTS01044, we have systematically developed a robust and efficient numerical solution technique for the heterogeneous high wavenumber Helmholtz equation. We have subsequently looked into the method's performance for spatially-dependent wavenum-

C.W. Oosterlee
CWI, Center for Mathematics and Computer Science, Amsterdam, and Delft University of Technology, Delft, the Netherlands, e-mail: c.w.oosterlee@cwi.nl

C.Vuik
Delft University of Technology, Delft, e-mail: c.vuik@tudelft.nl

W.A. Mulder
Shell International Exploration and Production, Rijswijk, NL e-mail: Wim.Mulder@shell.com

R-E Plessix
Shell International Exploration and Production, NL, e-mail: ReneEdouard.Plessix@shell.com

B. Koren and C. Vuik (eds.), *Advanced Computational Methods in Science and Engineering*, Lecture Notes in Computational Science and Engineering 71,
DOI 10.1007/978-3-642-03344-5_2,

bers, for high wavenumber problems, for the so-called absorbing boundary layer boundary conditions, for fourth-order discretizations, for 2D and 3D applications, academic and industrial applications. Former PhD student Yogi Erlangga, one of the driving forces behind this series of papers, has published an overview article on the topic, see [19], and a theoretical analysis has been presented in [20].

In [18], a preconditioned Bi-CGSTAB method has been introduced, in which the preconditioner is based on a second Helmholtz equation with an imaginary shift. This preconditioner is the basis of our work. It appears as a member of the family of shifted Laplacian operators, introduced in [34]. An interesting aspect is that its inverse can be efficiently approximated by means of a multigrid iteration, which is somewhat surprising as the original Helmholtz equation cannot be solved efficiently with off-the-shelf multigrid solvers. The particular preconditioner presented can be viewed as a generalization of the work by Bayliss, Goldstein, and Turkel [4] from the 1980s, where the Laplacian was proposed as a preconditioner for Helmholtz problems.

The idea of preconditioning the indefinite operator with a shifted Laplacian has now been considered also at other places in the research community. In [55] for example, the idea is highlighted. Further, a Finnish group at Jyväskylä has adopted the approach for their Helmholtz applications in [25, 26, 28, 1]. The preconditioner is also discussed and considered, for other Helmholtz applications in [12, 45, 7], and in different research areas that need to deal with indefinite problems, like electromagnetics or optics, in [6, 22, 59, 36, 42].

As an example of the solver, we discuss in this paper a version of the preconditioner for a fourth-order 2D finite-difference discretization of the Helmholtz operator. In the multigrid preconditioner we replace the point-wise Jacobi smoother, from [18], by a variant of the incomplete lower-upper factorization smoother, ILU(0). Furthermore, we show the performance of a prolongation scheme that originates from algebraic multigrid (AMG) [49]. We show that these enhancements to the iterative solver, proposed in [57], can reduce both the number of iterations and the total CPU time needed for convergence. Moreover, we aim to reduce the size of the imaginary shift parameter in the shifted Laplacian preconditioner, compared to the choice in [18], so that an even faster solution method is obtained. A fourth-order Helmholtz discretization enables us to use fewer grid points per wavelength compared to a second-order discretization.

Finally, the overall solution method with these algorithmic improvements is not limited to structured Cartesian grids, as it can be set up fully algebraically (a similar goal has been pursued in [1]). Although our method extends to solving problems on unstructured grids, we focus here on heterogeneous Helmholtz problems on Cartesian grids. We focus on the two-dimensional case; however, all of the method's ingredients can be easily generalized to three dimensions. Previously obtained 3D Helmholtz results, with the shifted Laplacian multigrid preconditioner, with a point-wise smoother, can be found in [48] (academic test problems) and [44] (industrial test problems).

This article is set up as follows. In Section 2, we discuss the Helmholtz equation, its field of application, and the discrete finite-difference formulations of second-

and fourth-order. The iterative solution method, including the preconditioner and its components, is presented in Section 3. Numerical results are presented in Section 4.

2 The Helmholtz equation and a seismic application

Accurately imaging the Earth is one of the major challenges in the hydrocarbon industry. Subsurface formations are mapped by measuring the time required for a seismic pulse to return to the surface after reflection from interfaces between formations with different physical properties. Variations in these reflection times, as recorded on the Earth's surface, usually indicate structural features in the strata below. Depths to reflecting interfaces can be determined from the times, using velocity information that can be obtained from the reflected signals themselves. At the same time the amplitudes of the signals provide valuable information.

In geophysics, numerical solutions for the wave equation are used in seismic imaging to map in depth the information recorded in time in the seismic data [13]. In the oil and gas industry, until approximately 1995, three-dimensional applications mainly relied on solutions of a high-frequency approximation of the wave equation, due to computer constraints. When the medium is very complex (containing heterogeneities that result in strong contrast) or in the case of crossing rays (multivaluedness), however, these so-called ray-based high-frequency migration techniques reach their limits.

With the increase of computer power and the need to image increasingly complex geological terrains, the so-called par-axial approximation of the wave equation [13] became popular. The accuracy of this approximation is however limited to certain angles of incidence which does not allow one to image steep reflectors accurately [40]. Nowadays, the three-dimensional (full) wave equation can be solved on a large cluster of computers with a finite-difference scheme and a time marching approach. This leads to the so-called reverse-time imaging algorithms. While still expensive, these imaging algorithms are now routinely applied in industry.

Typically an imaging algorithm requires the computation of thousands of wave equation solutions for given velocity and density fields. The complexity of a time-domain solution for a given source position is $O(n^4)$ with n the number of grid points in one direction (We can assume that the number of time steps is proportional to n). When the full time response is required, the time-domain solution has an optimal complexity if only one source is considered. However, it was noticed [39] that the matrix of the linear system associated with the discretization of the *frequency-domain wave equation*, i.e., the time-harmonic wave equation obtained by Fourier transformation, is independent of the source term. This implies that, when we can decompose this matrix, for instance by an LU decomposition, the frequency-domain approach can be more efficient than the time-domain approach when many solutions have to be computed with different source terms. This is the case in two dimensions [39]. Unfortunately, in three dimensions, the fill-in with an LU decomposition is very unfavorable, which makes the frequency-domain approach less attractive.

However, attempts to use massively parallel linear solvers are pursued in the context of full-wave form inversion when only a limited number of frequency responses are processed [43].

When the imaging algorithms require to localize the information in depth with a high resolution, as with the so-called migration algorithm, a full band-limited frequency response is needed, according to the Nyquist theorem. In this case, in three dimensions, it is more efficient to solve the time-domain wave equation than the Helmholtz equation. However, when only the solution at a limited number of frequencies is required, the choice of either the time or frequency domain remains open. This is the case for the full waveform inversion algorithm [46]. A competitive frequency-domain solver would need to be more efficient than a time-domain solver followed by a Fourier transform [44]. Since a direct solver is a-priori too expensive because of the filling of the LU decomposition, an iterative solver should be considered. For a given frequency and a given source, the optimal complexity of the iterative frequency-domain solver would be $O(n^3)$, which is better than the $O(n^4)$ complexity of the time-domain solver. Unfortunately, an optimal frequency-solver at seismic frequencies does not yet exist. The aim of the research presented here is to develop a robust and efficient solver of the Helmholtz equation at high wavenumbers. The Helmholtz equation corresponds to the frequency-domain acoustic wave-equation with constant density.

From the exploration-seismology point of view, the Earth is a heterogeneous semi-infinite medium. The wavenumber can be large, which implies that the discretized Helmholtz operator gives rise to both positive and negative eigenvalues and, therefore, the discretization matrix, A_h, is indefinite. For 2D problems, however, the computation can be performed efficiently by using, for example, direct methods combined with nested-dissection reordering [21]. Only one LU decomposition is needed to calculate the solutions at multiple source locations. The result can be used for the computation of all of the wavefields, for all shots and, also, for the back-propagated receiver wavefields [41]. However, for 3D problems, the matrix sizes and bandwidths rapidly become too large and one has to fall back on iterative methods. In that case, one no longer has the advantages in the frequency domain related to the LU decomposition.

For the Helmholtz equation, unfortunately, many iterative methods suffer from slow convergence, especially if high frequencies need to be resolved, due to the indefiniteness. The development of fast iterative methods for high-frequency Helmholtz problems remains a subject of active research. One approach to iteratively solving this equation is presented below. We focus on the 2D case, but provide a solution method which can easily be generalized to 3D.

2.1 Mathematical problem definition

We start with the description of the 2D Helmholtz problem of interest,

$$-\nabla^2 u(x,\omega) - k(x)^2(1-\alpha i)u(x,\omega) = g(x,\omega),\ x \in \Omega. \tag{1}$$

Unknown $u(x,\omega)$ represents the pressure field in the frequency domain, ∇^2 is the Laplacian operator, $k(x) = \omega/c(x)$ is the wavenumber, with $c(x)$, the acoustic-wave velocity, which varies with position, and $\omega = 2\pi f$ denotes angular frequency, a scalar measure of rotation rate (f is the frequency in Hertz). Wavenumber k depends on x because of a spatially dependent speed of sound, $c(x)$. The source term is denoted by g. The medium is called barely attenuative if $0 \leq \alpha \ll 1$, with α indicating the fraction of damping in the medium (and $i = \sqrt{-1}$, the imaginary unit). In geophysical applications, which are of our main interest, this damping can be up to 5% ($\alpha = 0.05$). While Equation (1) arises through the Fourier transform of a wave equation with a very simple model of damping, $(-\nabla^2 + (1-\alpha i)\partial_t^2)u = g$, it is closely related to the Fourier transform of the strongly damped wave equation, $(-\nabla^2 + \tau\partial_t\nabla^2 + \partial_t^2)u = g$ that yields, after scaling, $(-\nabla^2 - \frac{k^2}{1+\tau i k})u$. For small values of $\alpha = \tau k$, Equation (1) is an accurate approximation of the Fourier-domain strongly damped wave equation.

The semi-infinite physical domain needs to be truncated for a numerical treatment. A popular approach in geophysics in order to obtain a satisfactory near-boundary solution, without artificial reflections, is to use the absorbing boundary layer (ABL) approach; see, for example, [29] or [37]. This unphysical boundary layer is used to gradually damp the outgoing waves by adding dissipation in the equation outside the domain of interest. An efficient numerical solution technique should be robust with respect to this kind of feature. The absorption layers (denoted by Ω^e) are attached to the physical domain, Ω, (see Figure 1). In Ω^e, a damped Helmholtz equation (1) should be satisfied [51], with

$$\alpha = 0.25\frac{\| x - x_d \|^2}{\| x_e - x_d \|^2},\ x \in \Omega^e, \tag{2}$$

where point x_d is a point at the boundary, Γ, and x_e a point at Γ^e (see Figure 1). The boundary conditions at the boundary Γ^e are in the form of first- or second-order absorbing boundary conditions. We use approximate radiation (or non-reflecting) boundary conditions at the artificial boundary. The well-known second-order radiation boundary condition [15], to avoid unphysical reflections at boundaries, reads

$$\mathscr{A}_\Gamma u := \frac{\partial u}{\partial \nu} - iku - \frac{i}{2k}\frac{\partial^2 u}{\partial \tau^2} = 0 \quad \text{on } \Gamma^e, \tag{3}$$

with ν the outward normal direction to the boundary and τ pointing in the tangential direction. At the corner-points the suggestions in [3] to avoid corner reflections have been adopted.

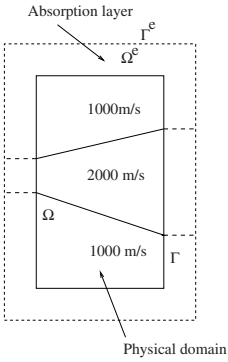

Fig. 1 A 2D domain with ABL in the case of a regular heterogeneous wedge medium.

2.2 Discretization

The equations are discretized here either by a second- or a fourth-order finite-difference scheme, resulting in the linear system:

$$A_h \phi_h = b_h \,, \tag{4}$$

where ϕ_h and b_h represent the discrete frequency-domain pressure field and the source, respectively.

In a heterogeneous medium, the smallest velocity is usually selected based on the representative wavelength, λ_f. The number of wavelengths in a domain of size L equals L/λ_f. A dimensionless wavenumber, k, on a non-dimensional $[0,1]^2$ domain is defined by $k = 2\pi f L/c$, and a corresponding mesh size by $h = \lambda_f/(n_w L)$, with n_w the number of points per wavelength.

The usual 5-point stencil related to a second-order accurate discretization reads:

$$A_h^{2o} \triangleq \frac{1}{h^2} \begin{bmatrix} & -1 & \\ -1 & 4 - k^2 h^2 (1 - \alpha i) & -1 \\ & -1 & \end{bmatrix}_h . \tag{5}$$

With domain size $L = 1$, an accuracy requirement, for second-order discretizations, is that $kh \leq \pi/5 (\approx 0.63)$ for $n_w = 10$ points per wavelength, and $kh \leq 0.53$ with $n_w = 12$ points per wavelength. The number of grid points used assumes a linear connection between k and h. In order to avoid a reduction of accuracy due to *pollution of the solution*, however, $k^2 h^3$ should be chosen constant, as stated in [5, 27]. For an iterative solution method, the requirement that kh should be constant is more severe and, so, this is the constraint that we consider here.

For the absorbing boundary conditions at Γ_e, we also apply central differences. Another discretization that we consider in this work is the $O(h^4)$ accurate discretization based on the Padé approximation. It is called the HO discretization in [50], with stencil

$$A_h^{HO} \triangleq \frac{1}{h^2}\begin{bmatrix} -\frac{1}{6} & -\frac{2}{3}-\frac{(kh)^2(1-\alpha i)}{12} & -\frac{1}{6} \\ -\frac{2}{3}-\frac{(kh)^2(1-\alpha i)}{12} & \frac{10}{3}-\frac{2(kh)^2(1-\alpha i)}{3} & -\frac{2}{3}-\frac{(kh)^2(1-\alpha i)}{12} \\ -\frac{1}{6} & -\frac{2}{3}-\frac{(kh)^2(1-\alpha i)}{12} & -\frac{1}{6} \end{bmatrix}. \tag{6}$$

An important reason for choosing a higher-order discretization method is that the number of grid points per wavelength can be reduced compared to a second-order discretization. Here, for example, we will show numerical experiments in which $kh = 0.8$ is set. This results in smaller matrices for the same level of accuracy and, thus, may lead to an algorithm that is more efficient overall, if the matrices associated with the higher-order discretization can be solved efficiently.

These matrices remain positive definite as long as k^2 is smaller than the first eigenvalue of the discrete Laplacian. The wavenumber in geophysical applications can, however, be large, which implies that the discretized Helmholtz equation gives rise to both positive and negative eigenvalues and, therefore, the discretization matrix, A_h, is indefinite. The size of the system of linear equations (4) gets very large for high frequencies. So, A_h in (4) is a large but sparse matrix, with complex-valued entries, because of the absorbing boundary conditions and the attenuative medium. It is symmetric but non-Hermitian.

2.2.1 Validation of the discretization

In order to validate the choice of boundary condition, ABL, and discretization, we first compute the solution for a constant wavenumber, and $\alpha = 0$, problem in a homogeneous medium with the source function, representative for a seismic pulse, chosen as

$$g_h = \frac{1}{h^2}\delta(x_1 - \frac{1}{2}, x_2 - \frac{1}{32}).$$

Here, $\delta(\cdot,\cdot)$ represents the Dirac delta function, which is 1 when its argument is $(0,0)$, and 0 elsewhere. The scaling by h^2 guarantees that the solutions on fine and coarse grids are of the same amplitude, giving a discrete approximation of a δ-function distribution. For this problem at constant wavenumber, the analytic solution is known, as the Green's function is available. The 2D solution reads:

$$u(r) = \frac{i}{4}H_0^{(1)}(k|r|),\ r = \sqrt{(x-1/2)^2 + (y-1/32)^2}, \tag{7}$$

where $H_0^{(1)}$ is the Hankel function of the first kind of order 0. So, we can compare a numerical solution with this analytic solution.

Two formulations of the boundary discretization are compared here. In the first, we prescribe the second-order absorbing boundary conditions directly at the phys-

ical boundaries, whereas, in the second formulation, the boundary discretization is based on an extra absorbing boundary layer, placed along all physical domain boundaries. An ABL of $n/4$ points is added to each side.

In this first numerical experiment, we fix the wavenumber, $k = 40$, and use a model domain, $(0,1)^2$, covered by a fine grid consisting of 256^2 points ($kh = 0.156$). Figure 2 presents the two corresponding solutions with the second-order discretization. An unphysical damping of the solution without the ABL can be observed near the domain boundaries.

Turning to the fourth-order discretization we drastically reduce the number of grid points. Figure 3 compares, for wavenumber $k = 40$, by means of a *solution profile* at $x = 0.125$, the second- and fourth-order discretizations. Coarse grids consisting of $32^2, 48^2$ and 64^2 points are chosen. As shown, the profile lines generated by the fourth-order discretization converge very nicely towards the physical solution (7), whereas the profile with the second-order discretization on the 32^2-grid is too inaccurate to show, and the solutions on the 48^2- and 64^2-grids are shifted in phase. Thus, the fourth-order discretization with the ABL leads, in this case, to an accurate numerical solution, already on relatively coarse grids, although the accuracy requirement (for second-order discretizations) $kh = 0.625$ is not satisfied there.

3 Iterative solution method

Before we discuss the solution method of choice, we outline the general convergence problems when using the multigrid method directly for the discrete Helmholtz equation of interest. This serves as an illustration of the difficulties one can encounter when solving this equation in a robust and efficient way.

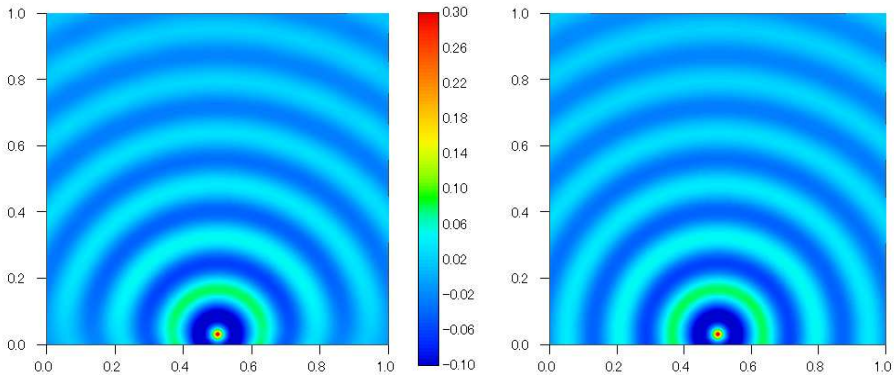

Fig. 2 Numerical solutions for $k = 40$ and $h = 1/256$, without (left-side) and with (right-side) ABL.

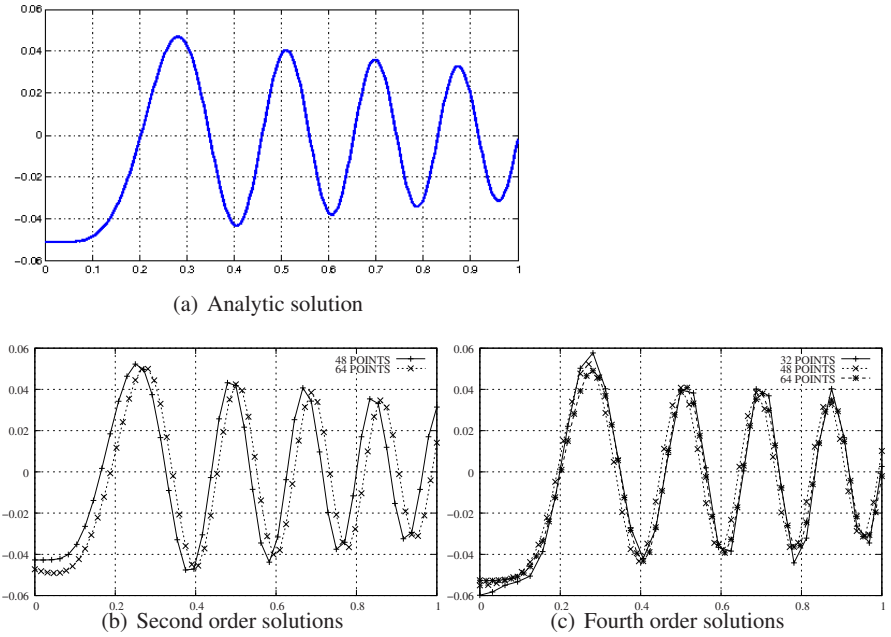

(a) Analytic solution

(b) Second order solutions

(c) Fourth order solutions

Fig. 3 Comparison of the vertical line solutions on coarse grids, at $x = 0.125$, $k = 40$ with the analytic solution, left side: second-order discretization, right side: fourth-order discretization. ABL is included.

3.1 Multigrid for the Helmholtz equation

Many authors, e.g. [23, 10, 14, 32], have contributed to the development of appropriate multigrid methods for the Helmholtz equation, but an efficient direct multigrid treatment of heterogeneous problems with high wavenumbers arising in engineering settings has not yet been proposed in literature. The multigrid method [8, 24] is known to be a highly efficient iterative method for discrete Poisson-type equations, even with fourth-order accurate discretizations [11, 54]. The Helmholtz equation, however, does not belong to the class of PDEs for which off-the-shelf multigrid methods perform efficiently. Convergence degradation and, consequently, loss of $O(N)$ complexity are caused by difficulties encountered in the smoothing and coarse-grid correction components.

Textbook multigrid methods are typically set up so that a smoothing method reduces high frequency components of an error, between the numerical approximation and the exact discrete solution, and a coarse-grid correction component handles the low frequency error components. Whereas such methods are easily defined for elliptic Poisson-like equations, this is not the case for the Helmholtz equation without any damping in (1), i.e., $\alpha = 0$. Depending on the particular value of the wavenumber, this equation gives rise to both smoothing and coarse grid correction difficul-

ties. For analyzing multigrid algorithms quantitatively, Fourier smoothing, two-, and three-grid analysis [8, 54, 61] are the tools of choice. They have given the following insights. The matrix, resulting from a discretization of the Helmholtz equation, has eigenvalues in only the right half-plane as long as k^2 is less than the smallest eigenvalue of the Laplacian. For larger values, this matrix does not have positive eigenvalues only. Point-wise Jacobi iteration with under-relaxation does not *converge* in that case, but since its smoothing properties are satisfactory, the multigrid convergence will deteriorate only gradually for increasing k^2. By the time it approaches a certain eigenvalue (the 6th), the standard multigrid method diverges. The Jacobi relaxation now diverges also for some smooth eigenfrequencies. Consequently, the multigrid method will still converge as long as the coarsest level used is fine enough to represent these smooth eigenfrequencies sufficiently. So, the coarsest level chosen limits the convergence. When the wavenumber gets larger more variables need to be represented on the coarsest level for standard multigrid convergence. Eventually, this does not result in an $O(N)$ iterative method.

In addition to this feature, the Helmholtz equation also suffers from its coarse-grid correction components. Eigenvalues close to the origin may undergo a sign change after discretization on a coarser grid. If a sign change occurs, the coarse-grid solution does not give a convergence acceleration but gives a severe convergence degradation (or even divergence) instead. In [14] this phenomenon is analyzed and a remedy for the coarse-grid correction related to these problematic eigenvalues is proposed. The efficient treatment in [14] is that the multigrid method is combined with Krylov subspace iteration methods on both fine and coarse grids.

Standard multigrid will also fail for k^2-values very close to eigenvalues. In that case subspace correction techniques should be employed [9].

3.2 Shifted Laplacian preconditioned Krylov subspace method

Iterative solution methods for complex-valued indefinite systems based on Krylov subspace methods [53] are typically generalizations of the conjugate-gradient (CG) method. The Bi-conjugate gradient stabilized (Bi-CGSTAB) algorithm [58] is one of the better known Krylov subspace algorithms for non-Hermitian problems, which has been used for Helmholtz problems, for example, in [18, 1]. One of the advantages of Bi-CGSTAB, compared to full GMRES [52], is its limited memory requirements. Bi-CGSTAB requires only 7 vectors to be stored. Bi-CGSTAB is based on the idea of computing two mutually bi-orthogonal bases for the Krylov subspaces based on matrix, A_h, and its conjugate transpose, A_h^H and is easy to implement.

Without a preconditioner, however, the Krylov subspace methods converge very slowly, or not at all, for the problems of interest [17]. By preconditioning with a matrix, M_h^{-1}, we solve an equivalent linear system,

$$A_h M_h^{-1} \tilde{\phi}_h = b_h, \qquad \tilde{\phi}_h = M_h \phi_h. \tag{8}$$

The challenge, then, is to find a matrix, M_h, such that $A_h M_h^{-1}$ has a spectrum that is favorable for iterative solution with Krylov subspace methods, and whose inverse, M_h^{-1}, can be efficiently approximated.

In [18], a shifted-Laplacian operator was proposed as a preconditioner for the Helmholtz equation, with M_h defined as a discretization of

$$\mathscr{M} = -\nabla^2 - k^2(x)(\beta_1 - \beta_2 i). \tag{9}$$

Equation (9) looks like equation (1), but is much easier to solver with multigrid methods. It will serve as the preconditioner here. Boundary conditions were set identically to those for the original Helmholtz equation. The influence of parameters β_1 and β_2 was evaluated in [18], and the optimal values for the solver proposed there were $(\beta_1, \beta_2) = (1, 0.5)$. Here, we will also consider $\beta_2 = 0.4$, as in [57]. Smaller values of β_2 do not lead to a converging algorithm with the components to be introduced below. The matrix after discretization of (9), M_h, is obtained from either the 5-point, $O(h^2)$, or the 9-point, $O(h^4)$, finite-difference discretization.

3.3 Fourier analysis

The discrete Helmholtz matrix, A_h, as well as the preconditioner, M_h, allow us, assuming a constant wavenumber and Dirichlet boundary conditions, to apply Fourier analysis on the basis of discrete sine-eigenfunctions,

$$v_h^{p,q} = \sin(p\pi x)\sin(q\pi y), \tag{10}$$

to gain insight into the spectrum of $A_h M_h^{-1}$. With these discrete sine functions, $A_h M_h^{-1}$ is diagonalized, and the eigenvalues can easily be determined. As long as k^2 is not equal to any of the eigenvalues of the discrete Laplace operator, $A_h M_h^{-1}$ is nonsingular. Otherwise, the matrix is singular and its nullspace is spanned by the corresponding eigenfunctions (10). The assumption of homogeneous Dirichlet boundary conditions greatly simplifies the analysis (for radiation boundary conditions the Helmholtz operator is non-normal so that eigenvalue analysis alone would not be sufficient for analyzing preconditioned Krylov subspace methods).

We perform Fourier analysis here to visualize the effect of the choice of the parameters, (β_1, β_2), as well as the choice of discretization on the clustering of the eigenvalues of the preconditioned system. This analysis gives only a first indication of what we can expect from the solver. For both A_h and M_h, we use either the second-order discretizations or the fourth-order, HO stencils. Initially, we do not include damping in A_h in the analysis (we take $\alpha = 0$ in (1)).

First, we visualize the effect of the choice of (β_1, β_2) in the preconditioner on the clustering of the eigenvalues of the preconditioned system. For both A_h and M_h we choose the 5-point stencil. Figure 4 presents the spectra of $A_h M_h^{-1}$ for $(\beta_1, \beta_2) = (0, 0)$ (Laplacian preconditioner [4]), $(\beta_1, \beta_2) = (-1, 0)$ (Laird precondi-

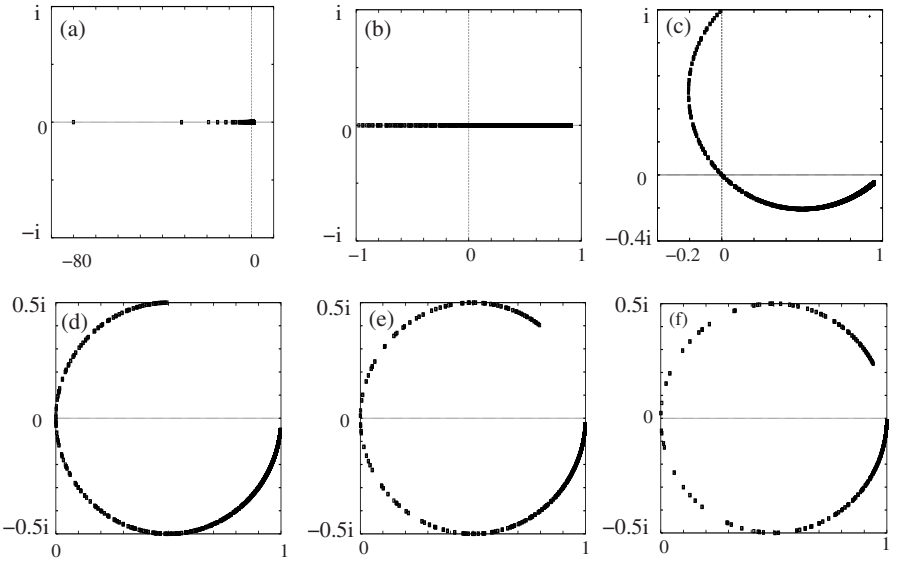

Fig. 4 Spectral pictures of $A_h M_h^{-1}$ with $\alpha = 0$ and different values of (β_1, β_2) in (8), see also [18]. (a) $(\beta_1, \beta_2) = (0,0)$, (b) $(-1,0)$, (c) $(0,1)$, (d) $(1,1)$, (e) $(1,0.5)$, and (f) $(1,0.3)$.

tioner [34]), $(\beta_1, \beta_2) = (0,1)$ (preconditioner from [16]), $(\beta_1, \beta_2) = (1,1)$ (basic parameter choice in [18]), $(\beta_1, \beta_2) = (1,0.5)$, and $(\beta_1, \beta_2) = (1,0.3)$ (more advanced parameters). The results are for $k = 40$ ($k^2 = 1600$) and $h = 1/64$.

From the spectra presented, the lower pictures of Figure 4 are favorable as their real parts vary between 0 and 1. The Laplacian preconditioner in Figure 4(a) exhibits large isolated eigenvalues; for the Laird preconditioner the eigenvalues in Figure 4(b) are distributed between -1 and 1 on the real axis. The preconditioners with complex Helmholtz terms give rise to a curved spectrum, see also [20]. Whereas the real part of the spectrum in Figure 4(c) still includes a part of the negative real axis, this is not the case for the (β_1, β_2)-preconditioners with $\beta_1 = 1$. The difference between Figures 4(d), 4(e), and 4(f) is that, with a smaller value of β_2, fewer outliers close to the origin are observed. This is favorable for the convergence of the preconditioned Krylov method. The approximate inversion of the preconditioner itself by multigrid, however, will be harder for smaller values of β_2. In Figure 5 the spectra for $k = 100 (k^2 = 10^4)$ are presented on a grid with $h = 1/160$ for $\beta_1 = 1$ and β_2 varying between 1 and 0.3. The spectra are very similar to those in Figure 4. More eigenvalues lie, however, in the vicinity of the origin due to the higher wavenumber and the correspondingly finer grid. Figure 6 then presents the distribution of eigenvalues for the case that 5% damping ($\alpha = 0.05$) is set in $\mathscr{A}$ from (1). Parameters in the preconditioner are $(\beta_1, \beta_2) = (1,0.5)$. Again the 5-point stencil as in (5) is used for discretization. Figure 6(a) presents the spectrum for $k = 40, h = 1/64$, and Figure 6(b) presents the spectrum for $k = 100, h = 1/160$. An interesting observation is that now the eigenvalues move away from the origin into the right half-plane.

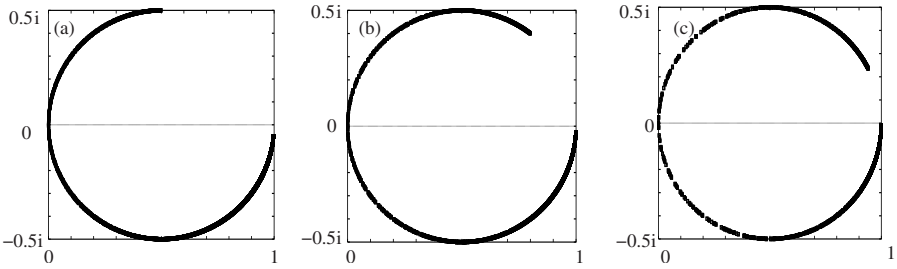

Fig. 5 Spectral pictures of $A_h M_h^{-1}$ for $k = 100$, $h = 1/160$, and $\alpha = 0$, see also [18]; (a) $(\beta_1, \beta_2) = (1,1)$, (b) $(\beta_1, \beta_2) = (1, 0.5)$, and (c) $(\beta_1, \beta_2) = (1, 0.3)$.

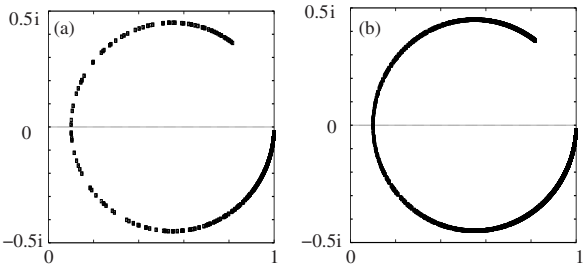

Fig. 6 Spectral pictures of AM^{-1} with 5 % damping in A and $(\beta_1, \beta_2) = (1, 0.5)$, see also [18]; (a) $k = 40, h = 1/64$ and (b) $k = 100, h = 1/160$.

This is beneficial for iterative solution methods. From the spectra in Figure 6 it is expected that the Bi-CGSTAB (and GMRES) convergence in the case of damping in the original equation will be considerably faster than in the undamped case. As the circles have moved away from the origin it is possible to apply the classical theory of the GMRES convergence [52, 53], for example.

Returning to the undamped case, $\alpha = 0$, we concentrate on the choice of discretization, and fix $k = 100$ $(k^2 = 10^4), h = 1/160$. Figure 7 presents the curved spectrum of $A_h M_h^{-1}$ for $(\beta_1, \beta_2) = (1, 0.4)$ in M_h, where both operators, A_h and M_h, are discretized by the fourth-order stencil. A very similar eigenvalue distribution is obtained as for the second-order discretization above.

The eigenvalues closest to the origin are the most problematic ones for the convergence of the Krylov subspace method. Figure 8a shows a zoom of the spectrum near the origin, comparing (for $\beta_2 = 0.5$) the location of the eigenvalues near the origin with the second- and the fourth-order discretizations. In Figure 8b, we also compare the location of the eigenvalues near the origin for $\beta_2 = 0.4$ and $\beta_2 = 0.5$, with the fourth-order discretization, keeping $k = 100$, $h = 1/160$. With the fourth-order discretization, the eigenvalues stay further from the origin as compared to the second-order discretization. This should have a positive effect on the convergence of the Krylov subspace method. The same is true when comparing the eigenvalues with $\beta = 0.4$ and $\beta_2 = 0.5$ where, as expected, the clustering with $\beta_2 = 0.4$ is more favorable for iterative solution.

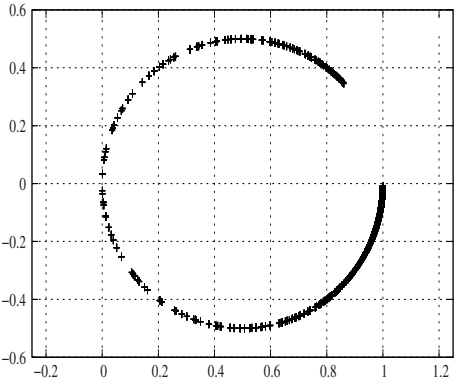

Fig. 7 Spectral picture of $M_h^{-1}A_h$ with $\alpha = 0$, $(\beta_1, \beta_2) = (1, 0.4)$, $k = 100, h = 1/160$. Both operators are discretized by the fourth-order, HO discretization.

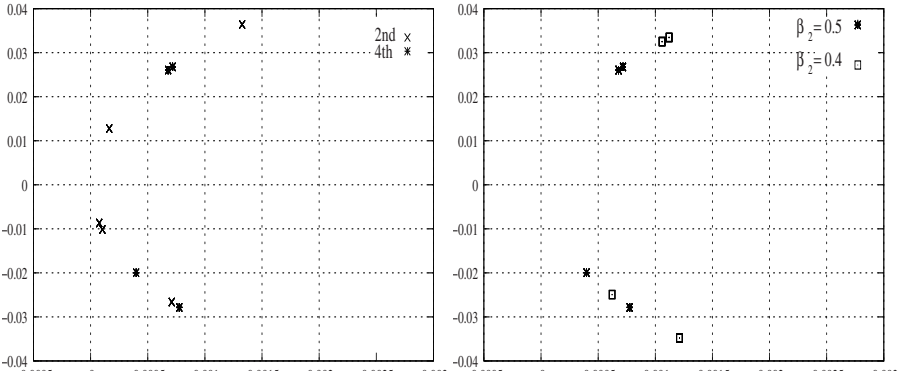

(a) 2nd versus 4th order, $\beta_2 = 0.5$ (b) 4th order, $\beta_2 = 0.4$ vs. $\beta_2 = 0.5$

Fig. 8 Comparison of zoomed spectral pictures of $M_h^{-1}A_h$ with $\alpha = 0$, $k = 100, h = 1/160$. (a): Comparison of eigenvalues near the origin discretized with second-order and with fourth-order discretizations; (b) Comparison of eigenvalues for $\beta_2 = 0.4$ and $\beta_2 = 0.5$, fourth-order discretizations.

When discretized with second-order finite differences, M_h, with $(\beta_1, \beta_2) = (1, 0.5)$, can be relatively easily handled by a multigrid method, which is confirmed by Local Fourier Analysis, a quantitative multigrid analysis tool [54]. This is due to the imaginary term, $\beta_2 i$ in the shifted Laplacian. LFA also indicates that M_h based on the fourth-order discretization can be dealt with in multigrid as efficiently as the second-order discretization. With the multigrid components from [57], including an ILU(0) smoother within the preconditioner, we aim to decrease the value of β_2, to $\beta_2 = 0.4$, and obtain an efficient preconditioned Krylov subspace method.

3.4 Multigrid preconditioner

One multigrid cycle, based on standard grid coarsening and point-wise smoothing, can be used as an approximation to M_h^{-1} with $(\beta_1, \beta_2) = (1, 0.5)$. In [18], an F(1,1)-cycle [54], with one pre- and one post-smoothing iteration, with a point-wise Jacobi smoother with under-relaxation parameter $\omega = 0.5$ was chosen for the high wavenumber problems. The other multigrid components were:

- Restriction operators, I_h^H, based on 2D full weighting, whose stencil [54] reads:

$$I_h^H \triangleq \frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}_h^H, \tag{11}$$

  with h denoting the fine-, and H denoting the coarse-mesh size.
- Prolongation operators, I_H^h, were 2D matrix-dependent interpolations, based on de Zeeuw's interpolation weights [63]. The interpolation weights are especially tailored to the symmetric complex Helmholtz equation, i.e., the asymmetric components in [63] have been removed. As for symmetric problems with jumping coefficients, the prolongation operator by de Zeeuw [63] is very similar to the operator-dependent prolongation in [2]. For satisfactory convergence it is, however, important to consider the *moduli* of the complex-valued operator elements in the definition of the interpolation weights.
- Coarse-grid matrices were based on Galerkin coarse-grid discretizations, defined as $M_H = I_h^H M_h I_H^h$.

In [18], it was shown that the full-weighting restriction combined with the matrix-dependent prolongation resulted in a stable convergence for a variety of problems with irregular heterogeneities and strong contrasts. The inclusion of an ABL in the discretization does not lead to any multigrid convergence difficulties, as the multigrid components chosen are especially designed for problems with varying coefficients.

With a more powerful smoother, however, a robust multigrid method can be developed for approximately inverting matrices M_h that originate from a fourth-order discretization. As the smoother in the multigrid preconditioner, the point-wise Jacobi smoother was replaced by an ILU smoother in [57]. ILU smoothing is well-known in the multigrid literature [30, 31, 62, 60, 64]. We choose here the ILU(0) variant, meaning that we do not allow any additional fill-in in the lower- and upper-triangular factors outside of the nonzero pattern of matrix M_h. An ILU(0) smoother is known to be more powerful than a point-wise Jacobi smoother for a number of test problems [60]. Strictly speaking, ILU methods do not only have a smoothing effect on the errors. A lexicographical version may also reduce low-frequency errors, especially when the entries of the remainder matrix, R_h, in

$$M_h = \hat{L}_h \hat{U}_h - R_h,$$

are relatively small. Parallelization of an ILU smoother is, however, less trivial, but possible, compared to a point-wise Jacobi smoother.

3.5 AMG type interpolation

An efficient multigrid scheme relies on the effective complementarity of the chosen relaxation and interpolation procedures in reducing the error components in an approximate solution. The coarse-grid correction operator is designed to reduce errors that the chosen smoother is slow to attenuate. Such errors should lie in the range of interpolation, so that the coarse-grid correction may be effective. We consider a fixed choice of coarse grid, i.e., Cartesian (doubling the mesh size in each direction) as in geometric multigrid and in [18], but employ an interpolation operator that is chosen based on algebraic multigrid (AMG) principles, evaluated in [57]. The interpolation developed is largely based on the real-valued AMG interpolation from [49], and discussed for complex-valued equations in [33, 38].

Consider, then, an error, e_h, that is not quickly reduced by relaxation. For many standard problems and smoothers, these errors coincide with those vectors that yield small residuals. For the purpose of interpolation, AMG assumes that the error, e_h, is much larger than its residual when measured point-wise, $(A_h e_h)_j \ll (e_h)_j$, for each fine-grid index j. Based on this property, we have

$$(A_h e_h)_j \approx 0 \Rightarrow a_{jj}(e_h)_j \approx -\sum_{k \neq j} a_{jk}(e_h)_k, \tag{12}$$

meaning that the value of the error at a fine-grid node, j, can be accurately approximated by the values from its neighboring nodes. If all neighboring nodes are also coarse-grid nodes, then (12) is easily turned into an interpolation formula.

With the fixed coarsening considered here, fine-grid node j will have both fine-grid and coarse-grid nodes as neighbors. Designing an interpolation procedure can, then, be thought of as modifying the balance in (12) in such a way as to remove connections to other fine-grid neighbors of j while preserving the overall balance. This is typically done by applying a partition to the neighboring nodes of j that identifies some nodes as important, or strong, connections and other nodes as unimportant, or weak connections. That is, we write the set, $\{k \neq j\} = C_j \cup F_j^s \cup F_j^w$, where C_j is the set of strongly connected coarse-grid neighbors of j, and the disjoint sets, F_j^s and F_j^w, denote the strong fine-grid and weak connections, respectively.

The matrix arising from the Helmholtz equation is complex and, typically, the sum of the moduli of the off-diagonal elements is larger than that of the diagonal element in each row. In this case, a different criterion should be considered as a measure of the strong connections. Here, we give two common criteria for defining the set, S_j, of strong connections for node j, defining

$$S_j = \left\{ k : |a_{jk}| \geq \theta \max_{l \neq j} |a_{jl}| \right\},$$

or

$$S_j = \left\{ k : -Re(a_{jk}) \geq \theta \max_{l \neq j} -Re(a_{jl}) \right\}.$$

Parameter θ allows some adjustment of the number of connections chosen as strong (relative to the strongest connection); for many problems, $\theta = 0.25$ is considered to be a standard choice. Numerical experiments with the discrete complex-valued shifted Laplacian have revealed that sometimes divergence is observed for high wavenumber problems if we use the measure based on the norm. The measure based on the real part of the matrix elements gave a satisfactory multigrid performance over a large range of wavenumbers and, thus, is used in the numerical results that follow.

It is expected that the weak connections of fine-grid node j can be discarded from the balance in (12). To remove these terms (in particular, the appearance of $(e_h)_k$ for $k \notin S_j$) without upsetting the balance, these terms are "lumped to the diagonal". In effect, this means that we make the approximation that $(e_h)_k \approx (e_h)_j$ for $k \in F_j^w$; while this approximation may not be very accurate, it is not harmful to make such a choice, since the connections involved are not important. Treating the strongly connected fine-grid neighbors of j is much more important, as these are connections that (by definition) cannot be easily dropped. In classical AMG methods, one assumes that these connections are well-represented on the coarse grid, by their values at neighboring points. Then, an approximation may be made by considering the weighted average of the values at common coarse-grid neighbors of node j and its fine-grid neighbor, node k, resulting in the expression,

$$(e_h)_k \approx \frac{\sum_{l \in C_j} a_{kl} (e_h)_l}{\sum_{l \in C_j} a_{kl}}.$$

If there is no point in C_j such that $a_{kl} \neq 0$ (or if $\sum_{l \in C_j} a_{kl} = 0$), then node k is neglected in the interpolation formula. Making these substitutions in (12) and choosing for equality, we then have

$$a_{jj}(e_h)_j = -\sum_{k \in C_j} a_{jk}(e_h)_k - \sum_{k \in F_i^s} a_{jk} \frac{\sum_{l \in C_j} a_{kl}(e_h)_l}{\sum_{l \in C_j} a_{kl}} - \sum_{k \in F_i^w} a_{jk}(e_h)_j,$$

or $(e_h)_j = \sum_{k \in C_j} w_{jk}(e_h)_k$, for

$$w_{jk} = -\frac{a_{jk} + \sum_{m \in F_i^s} \frac{a_{jm} a_{mk}}{\sum_{l \in C_j} a_{ml}}}{a_{jj} + \sum_{m \in F_i^w} a_{jm}}.$$

With these weights, we can form the coarse-to-fine transfer matrix, W, from which we can express the overall prolongation matrix, I_H^h, as

$$I_H^h = \begin{bmatrix} W \\ I \end{bmatrix}.$$

We stress that while we only investigate the use of this interpolation for structured grids in this work, the use of these multigrid components enable the solution of unstructured-grid Helmholtz problems.

4 Numerical experiments

In this section, we perform several numerical 2D experiments of increasing complexity. We start with the constant wavenumber problem, which serves as a benchmark for the algorithmic choices, after which we evaluate the method's performance for a Helmholtz problem with a wedge heterogeneity, and a model of the Sigsbee field.

The experiments have been performed with a C++ Helmholtz code on an Intel Core2 Duo 1.66GHz CPU, with 1.0GB RAM memory.

4.1 Homogeneous problem

The first numerical experiments are based on the *homogeneous* Helmholtz problem on the square domain, $(0,1)^2$, to gain insight into the overall performance of the solvers. The pulse source, g, is located near the surface, at $(\frac{1}{2}, \frac{1}{32})$ and is represented by the scaled delta function.

We set $\alpha = 0$ in equation (1) in all numerical experiments, unless stated otherwise.

We will evaluate the following two multigrid preconditioners:

1 A multigrid V(1,1)-cycle with de Zeeuw's prolongation operator, FW restriction and Jacobi smoothing with relaxation parameter $\omega = 0.5$. This is the solver from [18].
2 A multigrid V(0,1)-cycle with AMG's prolongation operator, FW restriction and ILU(0) post-smoothing. This is the solver from [57].

The quality of the multigrid preconditioner was assessed in [18] and [57], respectively. These preconditioners are combined with the Bi-CGSTAB Krylov subspace solver. The value of β_1 in the shifted Helmholtz preconditioner equals 1, β_2 is set to either 0.4 or 0.5.

4.1.1 Second- and fourth-order discretizations

We compare the convergence of the Krylov subspace solvers for the second- and fourth-order discretizations of both the original operator and the preconditioner. The ABL is not included in this experiment. We test each setting with a random initial guess and a point source as a right-hand side. The iteration is terminated as soon as

preconditioner	β_2	h: 1/64	1/128	1/256	1/512
ω-Jacobi	0.4	39 (0.26)	75 (1.7)	139 (12.6)	266 (99)
Zeeuw-V(1,1)	0.5	36 (0.22)	69 (1.5)	125 (11.3)	236 (88)
ILU(0)	0.4	26 (0.17)	43 (0.94)	88 (7.9)	218 (83)
AMG-V(0,1)	0.5	27 (0.19)	42 (0.92)	83 (7.4)	162 (61)

Table 1 Bi-CGSTAB performance for the homogeneous model (*second-order discretization*) in terms of number of iterations and CPU time in seconds (in brackets).

multigrid preconditioner	β_2	h: 1/64	1/128	1/256	1/512
ω-Jacobi	0.4	35 (0.30)	70 (2.0)	122 (14.1)	215 (103)
Zeeuw-V(1,1)	0.5	32 (0.25)	62 (1.8)	110 (13.0)	201 (96)
ILU(0)	0.4	16 (0.13)	26 (0.78)	45 (5.7)	84 (45)
AMG-V(0,1)	0.5	19 (0.16)	30 (0.91)	52 (6.6)	95 (51)

Table 2 Bi-CGSTAB performance for the homogeneous model (*fourth-order discretization*), $kh = 0.625$, in terms of number of iterations and CPU time in seconds (in brackets).

multigrid preconditioner	β_2	h: 1/64	1/128	1/256	1/512
ω-Jacobi	0.4	48 (0.42)	111 (3.2)	267 (31.1)	> 500
Zeeuw-V(1,1)	0.5	46 (0.44)	107 (3.1)	242 (28.1)	> 500
ILU(0)	0.4	21 (0.17)	35 (1.06)	65 (8.3)	151 (81.1)
AMG-V(0,1)	0.5	23 (0.25)	41 (1.25)	73 (9.3)	132 (71.0)

Table 3 Bi-CGSTAB performance for the homogeneous model (*fourth-order discretization*), $kh = 0.8$, in terms of number of iterations and CPU time in seconds (in brackets).

the relative residual is reduced to a prescribed tolerance of 10^{-6},

$$\frac{||r_i||}{||r_0||} \leq 10^{-6}. \tag{13}$$

In the Tables 1 and 2 we present, for $kh = 0.625$, the Bi-CGSTAB performance on four meshes, with the two multigrid preconditioners, for the second- and fourth-order discretizations, respectively. Furthermore, in Table 3 we repeat the experiment with the fourth-order discretization, now setting $kh = 0.8$.

For all solvers, we observe a more-or-less linear increase in the number of iterations for increasing wavenumbers. The performance of the V(0,1) multigrid preconditioner with the ILU(0) smoothing appears to be the most robust among these choices. Its convergence for both values of β_2, and for both discretizations, is very satisfactory. An interesting observation, comparing the results in the Tables 1 and 2, is that the performance of the AMG interpolation (as in Table 2) is, especially on the finest grid, significantly better for the fourth-order discretization. Setting $kh = 0.8$ in Table 3 gives us an increase in the number of iterations, compared to the perfor-

number of points		discretization	
n	β_2	2nd order	4th order
32	0.4	109 (0.42)	56 (0.31)
	0.5	120 (0.48)	67 (0.31)
64	0.4	34 (0.45)	32 (0.33)
	0.5	37 (0.48)	39 (0.55)
128	0.4	29 (1.4)	28 (2.0)
	0.5	30 (1.5)	29 (2.1)
256	0.4	27 (5.8)	25 (7.6)
	0.5	29 (6.3)	27 (8.2)

Table 4 Bi-CGSTAB performance for second- and fourth-order discretizations for the homogeneous model in terms of number of iterations and CPU time in seconds (in brackets).

mance in Table 2. Note, however, that we deal with higher wavenumbers in Table 3, when we fix kh to 0.8 and fix h.

The CPU times per iteration reported for the fourth-order discretization are always somewhat higher than for the second-order problem, as we deal with 9-point discretization stencils on all grids.

4.1.2 Fixed wavenumber, increasing mesh sizes

We reconsider the 2D homogeneous model from Subsection 2.2.1, discretized on the unit square with the ABL. Wavenumber $k = 60$ is set in this experiment, and $kh = 0.8$. Note that this choice of the linear relation between k and h does not lead to accurate solutions with the second-order discretization on the coarse meshes.

The number of grid points increases in order to confirm the asymptotic grid independent convergence of the preconditioned Bi-CGSTAB solver for a fixed continuum problem. We terminate the iterations as soon as the relative residual is less than 10^{-6}.

The iterative solver in this experiment is based on the Bi-CGSTAB method with a V(0,1) multigrid preconditioner, with $\beta_1 = 1$, in which the ILU(0) smoother and the AMG prolongation are incorporated. Table 4 presents the number of Bi-CGSTAB iterations, plus the CPU time to reach the termination criterion, for two values of β_2, $\beta_2 = 0.4$ and $\beta_2 = 0.5$, with the second- and fourth-order discretizations. We observe the asymptotic h-independent convergence rate for the iterative solver on the finer grids; with fixed wavenumber and h decreasing, approximately the same number of iterations is needed to satisfy the termination criterion, for both discretizations. However, on the coarse grids, where we have $kh > 1$, we see a drastic increase in the number of iterations needed to converge, especially for the second-order discretization.

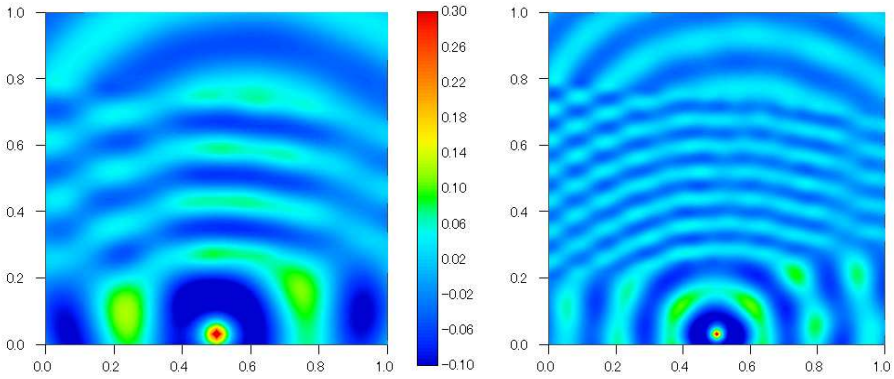

(a) $k = 40$ (b) $k = 80$

Fig. 9 Numerical solutions for the wedge problem with $k = 40$ and $k = 80$ with ABL.

4.2 The wedge problem

In this section, we present numerical results for the wedge problem. The domain, as in Figure 1, is a box, $(0,1)^2$, in which a wedge-shaped heterogeneity is placed, and the location of the source is $(1/2, 1/32)$. The wave number inside the wedge region is k, and outside the wedge it is set to $k/2$. We employ the fourth-order discretization with the ABL here, with $n/4$ points on both sides in the ABL.

Figure 9 presents the solutions of the wedge problem for $k = 40$ and $k = 80$.

We also examine the convergence of the preconditioned Bi-CGSTAB method, with the shifted Laplacian V(0,1)-multigrid preconditioner with the ILU(0) smoother and the AMG prolongation. We again set $kh = 0.8$ here. We also present results with some damping included in the original Helmholtz equation. Parameter α in (1) varies between 0 and 0.05. The number of iterations and CPU time (in seconds) are presented in Table 5. We notice a significant improvement of the method's convergence, already when 1% damping is included in the original problem. With 5% damping, we even observe a more-or-less constant number of iterations, for varying k. Compared to the performance of the solution method with the damped Jacobi smoother (not shown here), the results in Table 5 are significantly improved, both in terms of the number of iterations and in terms of the CPU time.

4.3 The Sigsbee problem

The Sigsbee2A synthetic data set models the geologic setting found on the Sigsbee escarpment in the deep-water Gulf of Mexico. There is a substantial uniform layer of water at the top of the model. Here, we use a version of the original Sigsbee model to test our iterative Helmholtz solver, see Figure 10. The size of the domain

	damping	number of points (kh constant)				
β_2	α	32	64	128	256	512
0.4	0.0%	18 (0.14)	27 (0.45)	48 (3.39)	88 (26.8)	175 (209)
	1.0%	17 (0.14)	24 (0.41)	38 (2.67)	54 (16.4)	86 (103)
	2.5%	14 (0.08)	21 (0.36)	29 (2.03)	38 (11.6)	53 (63.6)
	5.0%	12 (0.06)	16 (0.28)	20 (1.41)	24 (7.31)	28 (33.6)
0.5	0.0%	20 (0.14)	33 (0.58)	55 (3.88)	101 (30.7)	189 (227)
	1.0%	20 (0.16)	31 (0.53)	44 (3.09)	64 (19.4)	94 (113)
	2.5%	17 (0.09)	23 (0.39)	34 (2.39)	43 (13.1)	52 (62.5)
	5.0%	14 (0.08)	20 (0.34)	24 (1.69)	28 (8.52)	30(26.0)

Table 5 Bi-CGSTAB performance for the fourth-order discretization of the wedge model, with α % damping, in terms of number of iterations and CPU time in seconds (in brackets).

is $15000^2 m$ and a source is placed at $(7500, 117)$, near the top wall. The frequency chosen for this computation is $5Hz$.

The grid size consists of 512^2 points with an ABL of 128 points on each edge. The largest value of kh is 0.6135. Figure 11 presents the solution of this Helmholtz problem, where the fourth-order discretization is used. With the linear solver based on preconditioned Bi-CGSTAB with the V(0,1) multigrid preconditioner for the shifted Laplacian, using the ILU(0) smoother and AMG-based interpolation as essential components, we solve this problem in 61 iterations and 74.2 CPU seconds with $\beta_2 = 0.4$, and in 68 iterations and 85.5 CPU seconds for $\beta_2 = 0.5$. This convergence is highly satisfactory. As a comparison, the solver with $\beta_2 = 0.5$, the multigrid V(1,1)-cycle preconditioner, point-wise Jacobi smoothing and de Zeeuw's interpolation needed 216 iterations and 237 seconds CPU time.

Fig. 10 Domain for the scaled Sigsbee problem and the distribution of velocity, $c(x)$, and wavenumber, k.

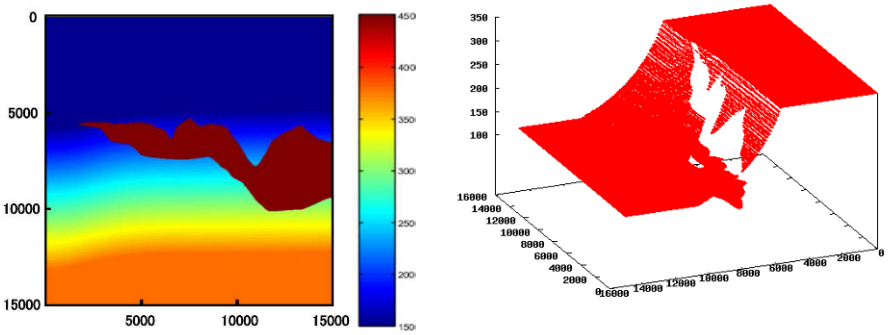

(a) Speed of sound, $c(x)$ for Sigsbee scaled domain

(b) Distribution of k

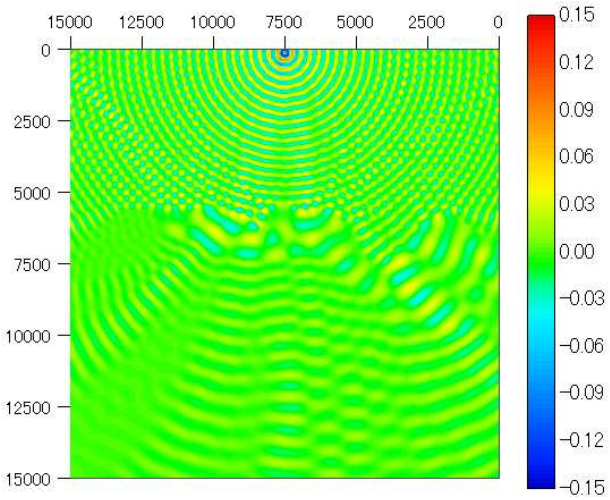

Fig. 11 Solution of the scaled Sigsbee problem with ABL, frequency $5Hz$.

5 Conclusion

In this paper, we have discussed the ingredients of a robust and efficient iterative solver for heterogeneous high wavenumber Helmholtz problems. A preconditioned Bi-CGSTAB solver has been developed in which the preconditioner is based on a shifted Laplacian with a complex-valued shift. We have shown that it is possible to work with fourth-order finite differences, both in the discrete original problem, as well as in the preconditioner. An absorbing boundary layer improves the quality of the solution significantly and does not pose difficulties to the solution method proposed. We have focused on fourth-order discretizations mostly obeying a linear relation between the wavenumber and the mesh size, $kh = 0.8$, here.

The fourth-order accurate shifted Laplacian preconditioner can be approximated by one V(0,1)-cycle of multigrid. In the multigrid preconditioner, we have compared a point-wise Jacobi smoother with an ILU(0) smoother and we included an AMG-based prolongation scheme. This enables us to choose a small imaginary shift parameter ($\beta_2 = 0.4$) in the preconditioner, which improves the solver's convergence (especially for high wavenumbers on fine meshes). The generalization to 3D is straightforward; the parallelization of this solver requires, however, some considerations as an ILU smoother is not easily parallelized.

Acknowledgment The authors would like to thank Shell International Exploration and Production and Philips for their financial support to this BTS project. Furthermore, they would like to thank Y.A. Erlangga, Ch. Dwi Riyanti, A. Kononov, M.B.

van Gijzen for their input in the project. In particular, we would also like to thank S.P. MacLachlan and N. Umetani for their contribution to this article, as the present article relies on paper [57] of which they are the co-authors. N. Umetani performed some additional numerical experiments, especially for this paper.

References

1. T. Airaksinen, E. Heikkola, A. Pennanen, J. Toivanen, An algebraic multigrid based shifted-Laplacian preconditioner for the Helmholtz equation. *J. Comp. Physics,* 226: 1196–1210, 2007.
2. R. E. Alcouffe, A. Brandt, J. E. Dendy Jr., J. W. Painter, *The multi-grid method for the diffusion equation with strongly discontinuous coefficients*, SIAM J. Sci. Comput., 2 (1981), pp. 430–454.
3. A. Bamberger, P. Joly, J.E. Roberts, Second-order absorbing boundary conditions for the wave equations: A solution for the corner problem. *SIAM J. Numer. Anal.,* 27: 323–352, 1990.
4. A. Bayliss, C. I. Goldstein, E. Turkel, *An iterative method for Helmholtz equation*, J. Comput. Phys., 49 (1983), pp. 443–457.
5. A. Bayliss, C.I. Goldstein E. Turkel, On accuracy conditions for the numerical computation of waves. *J. Comput. Phys.,* 59: 396–404, 1985.
6. M. Benzi, D. Bertaccini, Block Preconditioning of Real-Valued Iterative Algorithms for Complex Linear Systems, October 2006, revised September 2007, 21 pages. To appear in IMA Journal of Numerical Analysis.
7. M. Bollhöfer, M. Grote, O. Schenk, Algebraic multilevel preconditioning for Helmholtz equation, *In:* Proc. Europ. Conf. on Comput. Fluid Dynamics (ECCOMAS CFD 2006), Egmond aan Zee, The Netherlands, Sept. 5-8, 2006.
8. A. Brandt, Multi-level adaptive solutions to boundary-value problems. *Math. Comput.* 31: 333–390, 1977.
9. A. Brandt, S. Ta'asan, *Multigrid method for nearly singular and slightly indefinite problems*, in Proceedings EMG'85 Cologne, Multigrid Methods II, W. Hackbusch, U. Trottenberg, eds., Springer, Berlin, 1986, pp. 99–121.
10. A. Brandt, I. Livshits, Wave-ray multigrid methods for standing wave equations. *Elect. Trans. Numer. Anal.* 6: 162–181, 1997.
11. W.L. Briggs, V.E. Henson, S.F. McCormick, A multigrid tutorial: SIAM, Philadelphia, USA, 2000.
12. I. Duff, S. Gratton, X. Pinel, X. Vasseur, Multigrid based preconditioners for the numerical solution of two-dimensional heterogeneous problems in geophysics. *Intern. J. Computer Math.,* 84(8):1167–1181, 2007.
13. J. Claerbout, *Imaging the Earth's interior*, Blackwell Science Inc., 1985.
14. H. R. 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: 1291–1315, 2001.
15. B. Engquist, A. Majda, Absorbing boundary conditions for the numerical simulation of waves, *Math. Comput.*, 31: 629–651, 1977.
16. Y.A. Erlangga, C. Vuik, C.W. Oosterlee, *On a class of preconditioners for the Helmholtz equation*, Appl. Numer. Math., 50 (2004), pp. 409–425.
17. Y.A. Erlangga, C. Vuik, C.W. Oosterlee, Comparison of multigrid and incomplete LU shifted-Laplace preconditioners for the inhomogeneous Helmholtz equation. *Applied Num. Math.* 56: 648-666, 2006.
18. Y.A. Erlangga, C.W. Oosterlee, C. Vuik, A novel multigrid based preconditioner for heterogeneous Helmholtz problems. *SIAM J. Sci. Comput.* 27: 1471-1492, 2006.

19. Y.A. ERLANGGA, Advances in Iterative Methods and Preconditioners for the Helmholtz Equation *Archives Comput. Methods in Engin.,* 15: 37-66, 2008.
20. M. B. VAN GIJZEN, Y. A. ERLANGGA, C. VUIK, Spectral Analysis of the Discrete Helmholtz Operator Preconditioned with a Shifted Laplacian *SIAM J. Sci. Comput.*, 29: 1942-1958, 2007.
21. A. GEORGE, J.W. LIU, *Computer solution of large sparse positive definite systems,* Prentice-Hall, New Jersey, 1981.
22. S.GHEORGHE, *On Multigrid Methods for Solving Electromagnetic Scattering Problems.* PhD Thesis, Univ. Kiel, Germany, 2006.
23. J. GOZANI, A. NACHSHON, E. TURKEL, Conjugate gradient coupled with multigrid for an indefinite problem, in *Advances in Comput. Methods for PDEs V*, 425–427, 1984.
24. W. HACKBUSCH, *Multi-grid methods and applications*. Springer, Berlin, 1985.
25. E. HEIKKOLA, S. MÖNKÖLÄ, A. PENNANEN, T. ROSSI, Controllability method for acoustic scattering with spectral elements, *J. Comput. Appl. Math.,* 204(2): 344-355, 2007.
26. E. HEIKKOLA, S. MÖNKÖLÄ, A. PENNANEN, T. ROSSI, Controllability method for the Helmholtz equation with higher-order discretizations. *J. Comp. Phys.,* 225(2): 1553-1576, 2007.
27. F. IHLENBURG, I. BABUSKA, Finite element solution to the Helmholtz equation with high wave numbers. *Comput. Math. Appl.,* 30: 9-37, 1995.
28. K. ITO, J. TOIVANEN, A Fast Iterative Solver for Scattering by Elastic Objects in Layered Media. *Appl. Numerical Math.,* 57: 811-820, 2007.
29. C.-H. JO, C. SHIN, J.H. SUH, An optimal 9-point, finite-difference, frequency space, 2-D scalar wave extrapolator, *Geophysics* 61(2): 529–537, 1996.
30. R. KETTLER, Analysis and comparison of relaxation schemes in robust multigrid and preconditioned conjugate gradient methods. *In:*W. Hackbusch, U. Trottenberg (eds.), Multigrid methods, *Lecture Notes in Mathematics* 960: 502–534, Springer, Berlin, 1982.
31. M. KHALIL, *Analysis of linear multigrid methods for elliptic differential equations with discontinuous and anisotropic coefficients*. Ph.D. Thesis, Delft University of Technology, Delft, Netherlands, 1989.
32. S. KIM, S. KIM, Multigrid simulation for high-frequency solutions of the Helmholtz problem in heterogeneous media, *SIAM J. Sci. Comput*. 24: 684–701, 2002.
33. D. LAHAYE, H. DE GERSEM, S. VANDEWALLE, K. HAMEYER, *Algebraic multigrid for complex symmetric systems*, IEEE Trans. Magn., 36 (2000), pp. 1535–1538.
34. A. L. LAIRD, M. B. GILES, *Preconditioned iterative solution of the 2D Helmholtz equation.* Report NA 02-12, Comp. Lab., Oxford Univ., 2002.
35. B. LEE, T. A. MANTEUFFEL, S. F. MCCORMICK, J. RUGE, *First-order system least-squares for the Helmholtz equation*, SIAM J. Sci. Comput., 21 (2000), pp. 1927–1949.
36. S.S. LI, X.W. PING, R.S. CHEN, A Kind of Preconditioners Based on Shifted Operators to Solve Three-Dimensional TVFEM Equations *In:* Int. Symp. on Microwave, Antenna, Propagation and EMC Technologies for Wireless Communications, 842-844, 2007
37. Q. LIAO, G.A. MCMECHAN, Multifrequency viscoacoustic modeling and inversion. *Geophysics* 61(5): 1371–1378, 1996.
38. S.P. MACLACHLAN, C.W. OOSTERLEE, Algebraic multigrid solvers for complex-valued matrices, *SIAM J. Sci. Comput.*, 30:1548-1571, 2008.
39. K.J. MARFURT, Accuracy of finite-difference and finite-element modeling of the scalar and elastic wave-equations, *Geophysics,* 49: 533-549, 1984.
40. W.A. MULDER, R.-E. PLESSIX, A comparison between one-way and two-way wave-equation migration, *Geophysics*, 69: 1491–1504, 2004.
41. W.A. MULDER, R.-E. PLESSIX, How to choose a subset of frequencies in frequency-domain finite-difference migration, *Geophys. J. Int.* 158: 801–812, 2004.
42. O.V. NECHAEV, E.P. SHURINA, M.A. BOTCHEV, Multilevel iterative solvers for the edge finite element solution of the 3D Maxwell equation. *Comp. Mathem. Applications.* doi:10.1016/j.camwa.2007.11.003, to appear 2008.

43. S. OPERTO, J. VIRIEUX, P. AMESTOY, J.Y. L'EXCELLENT, L. GIRAUD, B. BEN HADJ ALI, 3D finite-difference frequency-domain modeling of visco-acoustic wave propagation using a massively parallel direct solver: A feasibility study, *Geophysics,* SM195-SM211, 2007.
44. R.-E. PLESSIX, A Helmholtz iterative solver for 3D seismic-imaging problems, *Geophysics,* 72: 185–194, 2007.
45. P. POULLET, A. BOAG, Incremental unknowns preconditioning for solving the Helmholtz equation. *Num. Methods Partial Diff. Equations,* 23(6): 1396-1410, 2007.
46. R.G. PRATT, Seismic waveform inversion in frequency domain. Part I: theory and verification in a physical scale domain, *Geophysics*, 64: 888–901, 1999.
47. C.D. RIYANTI, Y.A. ERLANGGA, R.-E. PLESSIX, W.A. MULDER, C.W. OOSTERLEE, C. VUIK, A new iterative solver for the time-harmonic wave equation, *Geophysics* 71: 57-63, 2006.
48. C.D. RIYANTI, A. KONONOV, Y.A. ERLANGGA, C. VUIK, C.W. OOSTERLEE, R-E PLESSIX W.A. MULDER, A parallel multigrid-based preconditioner for the 3D heterogeneous high-frequency Helmholtz equation. *J. Comp. Physics* 224: 431-448, 2007.
49. J.W. RUGE, K. STÜBEN, Algebraic Multigrid (AMG). *In:* S.F. McCormick (ed.), Multigrid Methods, Frontiers in Appl. Math., SIAM Philadelphia, 5: 73–130, 1987.
50. I. SINGER, E. TURKEL, High Order Finite Difference Methods for the Helmholtz Equation, *Comp. Meth. Appl. Mech. Eng.* 163:343-358, 1998.
51. I. SINGER, E. TURKEL, A Perfectly Matched Layer for the Helmholtz Equation in a Semi-infinite Strip. *J. Comput. Phys.* 201 (2004) 439-465.
52. Y. Saad, M.H. Schultz, GMRES: A generalized minimal residual algorithm for solving non-symmetric linear system, *SIAM J. Sci. Comput.,* 7: 856–869, 1986.
53. Y. SAAD, *Iterative Methods for Sparse Linear Systems*. SIAM, Philadelphia, 2003
54. U. TROTTENBERG, C.W. OOSTERLEE, A. SCHÜLLER, *Multigrid*, Academic Press, London, 2001.
55. E. TURKEL, Numerical Methods and Nature. *J. of Scientific Computing,* 28 2/3: 549-570, 2006.
56. E. TURKEL, *Numerical difficulties solving time harmonic equations*, in Multiscale Computational Methods in Chemistry and Physics, A. Brandt, et. al., eds., IOS Press, Tokyo, 2001, pp. 319–337.
57. N.UMETANI S.P.MACLACHLAN, C.W. OOSTERLEE, *A multigrid-based shifted-Laplacian preconditioner for a fourth-order Helmholtz discretization.* Working paper, Delft Univ. Techn. Delft, the Netherlands, 2008. Submitted for publication.
58. 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. Comput.*, 13: 631-645, 1992.
59. X. WEI, H. P. URBACH, A. J. WACHTERS, Finite-element model for three-dimensional optical scattering problems. *J. Opt. Soc. Am. A,* 24: 866-881, 2007.
60. P. WESSELING, *An introduction to multigrid methods*. John Wiley, Chichester, 1992.
61. R. WIENANDS, C. W. OOSTERLEE, *On three-grid Fourier analysis for multigrid*, SIAM J. Sci. Comput., 23 (2001), pp. 651–671.
62. G. WITTUM, On the robustness of ILU smoothing. *SIAM J. Sci. Comput.* 10: 699–717, 1989.
63. P. M. DE ZEEUW, Matrix-dependent prolongations and restrictions in a blackbox multigrid solver, *J. Comput. Appl. Math.*, 33: 1–27, 1990.
64. P.M. DE ZEEUW, Incomplete line LU as smoother and as preconditioner. *In:* W. Hackbusch, G. Wittum (eds.), Incomplete decompositions (ILU) - algorithms, theory, and applications. Vieweg, Braunschweig, 215–224, 1993. *Int. J. Num. Methods in Fluids* 20: 59–74, 1995.