Comput. Methods Appl. Mech. Engrg. 241–244 (2012) 155–171

Contents lists available at SciVerse ScienceDirect

Comput. Methods Appl. Mech. Engrg.

journal homepage: www.elsevier.com/locate/cma

Iterative solution applied to the Helmholtz equation: Complex deflation on unstructured grids

R. Aubry [a,*], S. Dey [b], R. Löhner [c]

[a] Sotera Defense Solutions, 1501 Farm Credit Drive, Suite 2300, McLean, VA 22102-5011, USA
[b] US Naval Research Laboratory, 4555 Overlook Ave., SW, Washington, DC 20375, USA
[c] CFD Center, Dept. of Computational and Data Science, College of Science, George Mason University, M.S. 6A2, Fairfax, VA 22030-4444, USA

ARTICLE INFO

Article history:
Received 23 January 2012
Received in revised form 9 May 2012
Accepted 5 June 2012
Available online 21 June 2012

Keywords:
Iterative solvers
Helmholtz equation
Deflation
Complex GMRES
Pollution effect

ABSTRACT

Extensions of deflation techniques developed for the Poisson and Navier equations (Aubry et al., 2008; Mut et al., 2010; Löhner et al., 2011; Aubry et al., 2011) [1–4] are presented for the Helmholtz equation. Numerous difficulties arise compared to the previous case. After discretization, the matrix is now indefinite without Sommerfeld boundary conditions, or complex with them. It is generally symmetric complex but not Hermitian, discarding optimal short recurrences from an iterative solver viewpoint (Saad, 2003) [5]. Furthermore, the kernel of the operator in an infinite space typically does not belong to the discrete space. The choice of the deflation space is discussed, as well as the relationship between dispersion error and solver convergence. Similarly to the symmetric definite positive (SPD) case, subdomain deflation accelerates convergence if the low frequency eigenmodes are well described. However, the analytic eigenvectors are well represented only if the dispersion error is low. CPU savings are therefore restricted to a low to mid frequency regime compared to the mesh size, which could be still relevant from an application viewpoint, given the ease of implementation.

1. Introduction

The Helmholtz equation is the archetype of the wave equation in the time domain. Applications of this equation are numerous and include acoustic scattering, geophysical seismic imaging, wireless communications. Due to this vast area of applications, a substantial effort has been invested in their numerical resolution. After discretization, these methods give rise to large, possibly sparse, matrices and their inversion may be time consuming. For large three dimensional problems, iterative methods in a broad sense (geometric or algebraic multigrid, iterative solvers, domain decomposition methods) represent the methods of choice due to memory (and possibly CPU) requirements. However their robustness is often criticized beyond the elliptic case. The main motivation of this work is given by the challenge of solving iteratively the coupled elastodynamic acoustic problem. The first building block for the elastic part consisted in extending the results of the scalar Poisson solver to the static elastic system. This was presented in [4]. The present work constitutes the first departure from the symmetric definite positive (SPD) case for scalar equation. In the literature, the Helmholtz equation has been mainly studied from two apparently different viewpoints, either from a discretization and accuracy viewpoint, or from an algebraic solver viewpoint. However, in both approaches plane waves, which are solutions of the Helmholtz equation in free space play a special part.

From an accuracy viewpoint, the oscillating nature of the Helmholtz equation gives rise to the famous pollution effect for high wave numbers [6–9]. Beside refining the mesh in the Finite Element h-refinement approach and increasing the polynomial order in the p-refinement paradigm, numerous methods have been designed to stabilize the Helmholtz equation. The Generalized Finite Element Method (GFEM) [10] modifies a bilinear stencil to have minimum pollution effect by minimizing the distance between the zeros of the discrete symbol and the one of the continuous symbol. The Partition of Unity Method (PUM) [11], uses analytical functions in the shape function definition. Extensions in three dimensions are presented in [12] in a Finite Element context and [13] in a Boundary Element context. The ultraweak method [14] relies on test functions that are solutions of the adjoint problem. The least square method [15] uses plane waves or Bessel functions in a discontinuous manner inside each element. The Galerkin Least Square (GLS) method [16] intends to stabilize the Helmholtz equation by adding consistently new weighted terms. The Residual Free Method (RFM) [17] relies on a bubble which verifies the analytical solution of the Helmholtz equation inside each element. A discontinuous Galerkin method [18] enriches the classical polynomial space with plane waves, and continuity is enforced weakly through Lagrange multipliers. The

* Corresponding author.
E-mail address: raubry@gmu.edu (R. Aubry).

0045-7825/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.cma.2012.06.007

Discrete Singular Convolution (DSC) [19] method applies singular convolution with a special kernel for high wave numbers. The residual based method [20] belongs to the variational multiscale methods but includes the residual on inter-element boundaries. As underlined, most of these methods rely on a continuous or discontinuous introduction of plane waves to improve the classical finite element discretization.

From an algebraic solver viewpoint, numerous techniques have also been attempted, including preconditioning, geometric and algebraic multigrid, and domain decomposition methods. A review is given by Erlangga in [21]. For the first class of methods, Magolu Monga Made et al. [22,23] proposes to use an imaginary perturbation of the original matrix as preconditioner. A large class of preconditioners begins with the work in Bayliss et al. [24], where the normal equation is solved with a symmetric successive over relaxation (SSOR) preconditioner relying only on the discrete Laplacian part of the Helmholtz operator. The preconditioner is now symmetric positive definite and a few sweeps of a multigrid solver may be used. Later, Laird and Giles [25] add a mass matrix part in the preconditioner while the mass matrix is associated with a negative sign in the original equation, and do not include boundary conditions for this matrix. Recently, Erlangga et al. [26,27] add a complex shift for the mass matrix, still allowing the possibility of a multigrid solve for the the fast inversion of the preconditioner, as shown in an unstructured context by Airaksinen et al. [28]. Bollhöfer et al. [29] present an algebraic multilevel preconditioner in heterogeneous media. Finally, Osei-Kuffuor and Saad [30] combine imaginary diagonal shifts with an algebraic recursive multilevel preconditioner. Recently, deflation has been applied to the Helmholtz equation discretized by an Integral Formulation [31]. The deflation space is composed of the eigenvectors of a coarse grid operator interpolated on the fine grid. It is clearly shown that deflation may improve drastically convergence, and as a by product, demonstrates the weak non normality of the Helmholtz equation [32]. However, the coarse mesh size presented is of the order of 40 percent of the size of the fine mesh, which is not affordable for large problems. Regarding multigrid techniques, a major breakthrough comes from Brandt and Livshits [33–35]. There are at least two reasons for the bad convergence of standard multigrids for the Helmholtz equation. First, standard smoothers diverge due to the non SPD behavior of the operator. Secondly, due to the oscillatory nature of the Helmholtz solution, standard restriction operators put a heavy constraint on the size of the coarse grid. To alleviate the latter major drawbacks, exponential restriction is performed, and only the smooth part of the solution is transfer to the ray grids, where it is solved efficiently. A similar approach is followed in Lee et al. [36] for a first order system least-squares formulation. Kim and Kim [37] use a Gauss Seidel (GS) or Conjugate Gradient for the Normal equation (CGNR) as a smoother. The large coarse grid problem is solved by a domain decomposition method. Elman et al. [38] use GMRES as a smoother in an outer flexible loop for robustness. Another approach is proposed in Vanek et al. [39], where aggregation is first performed to obtain a coarse level. Relying on the free space solution of the Helmholtz equation, a tentative prolongation is then build and smoothed through a polynomial matrix iteration, whose main aim is to minimize the energy of the columns of the prolongation. Oscillatory functions are interpolated with a constant value and with their gradients, as they do not belong exactly to the discrete space. As a final remark on multigrids relying on exponential interpolation, only two dimensional examples with very simple geometries have been shown to illustrate their numerical performances. Domain decomposition methods have also been applied to the Helmholtz equation by Farhat et al. for continuous [40] and discontinuous discretizations [41] relying on plane waves for the Lagrange multiplier space as well as for the primary variable.

It is therefore obvious that plane waves play a special part in the numerical resolution of the Helmholtz equation, as much from an accuracy as from an efficiency viewpoint. The shift produced by the wave number in the Laplacian spectrum implies that Laplacian eigenmodes associated with higher and higher eigenvalues become low energy modes impeding convergence. Furthermore, as noted in [36], though not from an algebraic viewpoint, the density of these modes increases with the wave number, as the Laplacian spectrum is much denser at its upper end. As plane waves do not belong to low order discretization, it may be foreseen that they will approximate well the low energy modes only for low to mid frequencies. The dispersion induced by the discretization will create a larger mismatch as the wave number increases.

In this paper, deflation applied to the Helmholtz equation is presented. Deflation has been shown to possess various computational advantages compared to other algebraic solvers for large unstructured meshes [1–4]. Whereas the multigrid approach gives a sound basis to tackle the problem, the geometric multigrid hierarchy is awkward to treat in an unstructured context with moving bodies, and the algebraic set up is slow. As noted in [38], the wave ray multigrid is "considerably more difficult to implement" than the multigrid proposed in the latter paper, even though it may be more efficient. It was hoped that deflation may achieve this efficiency in a three dimensional context with the ease of implementation of the deflation technique. However, this aim is only partially met. After this introduction, the deflation technique is reviewed in Section 2, and differences between the SPD and the non SPD case are highlighted. The complex deflated GMRES algorithm is recalled. It will be the method of choice for the next section as the Helmholtz equation gives rise to a symmetric complex but non Hermitian matrix. The Helmholtz equation is then presented in Section 3. Deflation applied to the discrete Helmholtz equation is considered. Finally, numerical results are provided in Section 4.

2. Complex deflation

In this section, the complex deflated GMRES used in this paper is presented. First, deflation applied to iterative solvers is succinctly reviewed. The Hermitian case is then recalled, followed by the non Hermitian case. Finally, the complex deflated GMRES is derived.

2.1. Deflation applied to iterative solvers

Deflation is an old and common technique in iterative solvers for eigenvalues [42,43]. In his seminal paper [44], Nicolaides accelerates an iterative solver for symmetric positive definite matrices, the widely utilized preconditioned conjugate gradient (PCG) [45], through a deflation technique (see [1] for other references). In a non symmetric context, Morgan [46] considers deflation to improve the GMRES restart. More recently, deflation has been extended to non symmetric solvers with success in [47,48]. The deflated preconditioned GMRES is at the crossroads of various iterative solvers for large matrices such as multigrid, either geometric or algebraic, domain decomposition, and of course Krylov subspace methods, as all these methods may be interpreted as projection methods [49,5,50]. Even though the core algorithm is constituted by a Krylov iterative solver, its main aim is to remove from the residual eigenvector components that are difficult to remove by standard iterative solvers. Convergence of GMRES for symmetric positive definite matrices can be shown to strongly rely on the con-

dition number κ of the matrix [5]. Even though GMRES convergence is much more complex for the general case [51,32], and may strongly depend on the conditioning of the eigenbasis, it has been repeatedly reported that removing eigenmodes associated with low eigenvalues may improve convergence [46,52,53]. From a practical viewpoint, modes associated with low eigenvalues should be identified, either analytically or algebraically, and should be well represented in the deflated subspace. Compared to a multigrid approach, the deflated subspace plays the role of the prolongation in a two level coarse grid correction. The use of zero energy modes in a multigrid context is discussed in [54], but was already noted in [55]. In [1], the deflated subspace was constructed by subdomain agglomeration on the mesh to try to represent constant modes in each subdomain, as constant modes belong to the kernel of the continuous operator. Subdomain agglomeration or plain aggregation has been used since the fifties (see [56] p. 68, [57] and references therein) in economic modeling and was also present in the original paper of Nicolaides [44].

2.2. Deflation for Hermitian matrices

For Hermitian matrices, there are many ways of deriving the deflated conjugate gradient [44,58,59,1]. Here, the approach of [59] is followed as it is more easily extended to the non Hermitian case. Let us suppose the algebraic system to be solved reads:

$$\mathbf{A}\mathbf{x} = \mathbf{b} \tag{2.1}$$

where $\mathbf{A}$ is a square matrix and $\mathbf{b}$ the right hand side. Given a deflation space $\mathbf{W}$ constituted by to the modes of the residual to be removed, let us define the projector $\mathbf{P}$ as:

$$\mathbf{P} = \mathbf{I} - \mathbf{A}\mathbf{W}(\mathbf{W}^T\mathbf{A}\mathbf{W})^{-1}\mathbf{W}^T \tag{2.2}$$

$\mathbf{P}$ is an $\mathbf{A}^{-1}$-orthogonal projector onto $\mathbf{W}^{\perp}$ along span{$\mathbf{A}$ $\mathbf{W}$} as:

$$(\mathbf{P}\mathbf{V}, \mathbf{A}\mathbf{W})_{\mathbf{A}^{-1}} = (\mathbf{P}\mathbf{V}, \mathbf{W}) = (\mathbf{V}, \mathbf{P}^T\mathbf{W}) = 0 \quad \forall \mathbf{V} \tag{2.3}$$

Its transpose is:

$$\mathbf{P}^T = \mathbf{I} - \mathbf{W}(\mathbf{W}^T\mathbf{A}\mathbf{W})^{-1}\mathbf{W}^T\mathbf{A} \tag{2.4}$$

and is an $\mathbf{A}$-orthogonal projector onto $\mathbf{W}^{\perp_{\mathbf{A}}}$ along span{$\mathbf{W}$} as:

$$(\mathbf{P}^T\mathbf{V}, \mathbf{W})_{\mathbf{A}} = (\mathbf{A}\mathbf{P}^T\mathbf{V}, \mathbf{W}) = (\mathbf{V}, \mathbf{P}\mathbf{A}\mathbf{W}) = 0 \quad \forall \mathbf{V} \tag{2.5}$$

It is easily verified that $\mathbf{P}$ and $\mathbf{P}^T$ are projectors, as:

$$\mathbf{P}^2 = \mathbf{P} \tag{2.6}$$

As $\mathbf{A}$ is symmetric, the following relation holds true:

$$\mathbf{A}\mathbf{P}^T = \mathbf{P}\mathbf{A} \tag{2.7}$$

The solution $\mathbf{x}$ is therefore obtained as [60]:

$$\mathbf{x} = \mathbf{P}^T\mathbf{x} + (\mathbf{I} - \mathbf{P}^T)\mathbf{x} = \mathbf{P}^T\mathbf{x}_1 + \mathbf{x}_2 \tag{2.8}$$

with:

$$\mathbf{P}^T\mathbf{x}_1 = \mathbf{P}^T\mathbf{x} \tag{2.9}$$

such that:

$$\mathbf{P}\mathbf{A}\mathbf{x}_1 = \mathbf{P}\mathbf{b} \tag{2.10}$$

and:

$$\mathbf{x}_2 = (\mathbf{I} - \mathbf{P}^T)\mathbf{x} = \mathbf{W}(\mathbf{W}^T\mathbf{A}\mathbf{W})^{-1}\mathbf{W}^T\mathbf{b} \tag{2.11}$$

In details, $\mathbf{x}_2$ is computed as:

$$\mathbf{W}^T\mathbf{A}\mathbf{W}\mu = \mathbf{W}^T\mathbf{b} \tag{2.12}$$

and

$$\mathbf{x}_2 = \mathbf{W}\mu \tag{2.13}$$

In order to obtain $\mathbf{x}$, $\mathbf{x}_1$ is multiplied by $\mathbf{P}^T$ and is added to $\mathbf{x}_2$ to form the solution.

2.3. Deflation for non Hermitian matrices

For the non Hermitian case, some modifications must be conducted. A new projector $\mathbf{Q}$ must be defined as:

$$\mathbf{Q} = \mathbf{I} - \mathbf{W}(\mathbf{X}\mathbf{A}\mathbf{W})^{-1}\mathbf{X}\mathbf{A} \tag{2.14}$$

where $\mathbf{X}$ is another subspace of appropriate dimensions, which would ideally be composed of the left eigenvectors, as commented in the next section. The projector $\mathbf{P}$ remains the same, apart from the introduction of $\mathbf{X}$:

$$\mathbf{P} = \mathbf{I} - \mathbf{A}\mathbf{W}(\mathbf{X}\mathbf{A}\mathbf{W})^{-1}\mathbf{X} \tag{2.15}$$

As $\mathbf{A}$ is not Hermitian, Eq. (2.7) does not hold anymore. However the following relation substitutes it:

$$\mathbf{A}\mathbf{Q} = \mathbf{P}\mathbf{A} \tag{2.16}$$

The solution $\mathbf{x}$ is obtained as [47]:

$$\mathbf{x} = \mathbf{Q}\mathbf{x} + (\mathbf{I} - \mathbf{Q})\mathbf{x} = \mathbf{Q}\mathbf{x}_1 + \mathbf{x}_2 \tag{2.17}$$

with:

$$\mathbf{Q}\mathbf{x}_1 = \mathbf{Q}\mathbf{x} \tag{2.18}$$

such that:

$$\mathbf{P}\mathbf{A}\mathbf{x}_1 = \mathbf{P}\mathbf{b} \tag{2.19}$$

and:

$$\mathbf{x}_2 = (\mathbf{I} - \mathbf{Q})\mathbf{x} = \mathbf{W}(\mathbf{X}\mathbf{A}\mathbf{W})^{-1}\mathbf{X}\mathbf{b} \tag{2.20}$$

In details, $\mathbf{x}_2$ is computed as:

$$\mathbf{X}\mathbf{A}\mathbf{W}\mu = \mathbf{X}\mathbf{b} \tag{2.21}$$

and

$$\mathbf{x}_2 = \mathbf{W}\mu \tag{2.22}$$

In order to obtain $\mathbf{x}$, $\mathbf{x}_1$ is multiplied by $\mathbf{Q}$ and is added to $\mathbf{x}_2$ to form the solution, similarly to the Hermitian case.

2.4. The deflated GMRES algorithm

As a choice of iterative solver, the classical GMRES algorithm has been chosen [61]. It consists in building a basis for the Krylov subspace by the Arnoldi process and solving the minimizer of the residual on this subspace. It uses Givens rotations for the rank one update of the Hessenberg matrix to convert it into an upper triangular matrix. The solution to the linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$ is sought as:

$$\mathbf{x} = \mathbf{x}_0 + \mathbf{V}_m\mathbf{y} \tag{2.23}$$

where $\mathbf{V}_m$ is a subspace of dimension m where the solution is sought, so that:

$$\begin{aligned} \mathbf{b} - \mathbf{A}\mathbf{x} &= \mathbf{b} - \mathbf{A}(\mathbf{x}_0 + \mathbf{V}_m\mathbf{y}) && (2.24)\\ &= \mathbf{r}_0 - \mathbf{A}\mathbf{V}_m y && (2.25)\\ &= \beta\mathbf{v}_1 - \mathbf{V}_{m+1}\mathbf{H}_m\mathbf{y} && (2.26)\\ &= \mathbf{V}_{m+1}(\beta\mathbf{e}_1 - \mathbf{H}_m\mathbf{y}) && (2.27) \end{aligned}$$

where $\mathbf{H}$ is an upper Hessenberg matrix, r_0 is the initial residual, and β the norm of the first residual. The GMRES approximation is defined as:

$$\mathbf{x} = \mathbf{x}_0 + \mathbf{V}_m \mathbf{y}_m \tag{2.28}$$

with the constraint that:

$$\mathbf{y}_m = \mathbf{argmin} \|\beta \mathbf{e}_1 - \mathbf{H}_m \mathbf{y}\|_2 \tag{2.29}$$

For the complex version, the restarted algorithm reads [5,62]:

Algorithm 1. Given $\mathbf{x}$:

while (not converged) **do**
  $\mathbf{x}_0 = \mathbf{x}$
  $\mathbf{r} = \mathbf{b} - \mathbf{A}\mathbf{x}$
  $\beta = \|\mathbf{r}\|_2,\ \mathbf{v}_1 = \mathbf{r}/\beta$
  for $i = 1$ to m **do**
    $\mathbf{w} = \mathbf{A}\mathbf{v}_i$
    for $k = 1$ to i **do**
      $h_{k,i} = \mathbf{v}_k^T \mathbf{w},\ \mathbf{w} = \mathbf{w} - h_{k,i}\mathbf{v}_k$
    end for
    $h_{i+1,i} = \|\mathbf{w}\|_2,\ \mathbf{v}_{i+1} = \mathbf{w}/h_{i+1,i}$
    if (converged) **then**
      exit
    end if
  end for
  compute the minimizer of $\|\beta \mathbf{e}_1 - \mathbf{H}y\|_2$
  $\mathbf{x} = \mathbf{x}_0 + \mathbf{V}\mathbf{y}$
  if (converged) **exit**
end while

Compared to the real case, the complex GMRES uses a complex scalar product, and rotations have to be modified, as commented in [5,62]. Following Vuik's approach, the GMRES algorithm applied to $\mathbf{PAx} = \mathbf{Pb}$ reads:

Algorithm 2. Given $\mathbf{x}$:

while (not converged) **do**
  $\mathbf{x}_0 = \mathbf{x}$
  $\mathbf{r} = \mathbf{Pb} - \mathbf{PAx}_0$
  $\beta = \|\mathbf{r}\|_2,\ \mathbf{v}_1 = \mathbf{r}/\beta,\ \tilde{\beta} = \beta \mathbf{e}_1$
  for $i = 1$ to m **do**
    $\mathbf{w} = \mathbf{A}(\mathbf{v}_i - \mathbf{W}\mu)$ where $\mathbf{XAW}\mu = \mathbf{XAv}_i$
    !Same as before
  end for
end while

However, the final unknown is given by:

$$\mathbf{x}_f = (\mathbf{I} - \mathbf{Q})\mathbf{x} + \mathbf{Q}\mathbf{x} \tag{2.30}$$
$$= \mathbf{W}(\mathbf{XAW})^{-1}\mathbf{Xb} + \mathbf{x} - \mathbf{W}(\mathbf{XAW})^{-1}\mathbf{XAx} \tag{2.31}$$
$$= \mathbf{W}(\mathbf{XAW})^{-1}\mathbf{Xb} + \mathbf{x}_0 + \mathbf{V}_m\mathbf{y} - \mathbf{W}(\mathbf{XAW})^{-1}\mathbf{XAx}_0 - \mathbf{W}(\mathbf{XAW})^{-1}\mathbf{XAV}_m\mathbf{y} \tag{2.32}$$
$$= \mathbf{x}_0 + \mathbf{W}(\mathbf{XAW})^{-1}\mathbf{Xr}_0 + \tilde{\mathbf{V}}_m\mathbf{y} \tag{2.33}$$
$$= \tilde{\mathbf{x}}_0 + \tilde{\mathbf{V}}_m\mathbf{y} \tag{2.34}$$

where:

$$\tilde{\mathbf{V}}_m = \mathbf{V}_m - \mathbf{W}(\mathbf{XAW})^{-1}\mathbf{XAV}_m \tag{2.35}$$

It is then possible to modify the algorithm to obtain a more compact version:

Algorithm 3. Given $\mathbf{x}$:

while (not converged) **do**
  $\mathbf{x}_{-1} = \mathbf{x}$
  $\mathbf{r} = \mathbf{Pb} - \mathbf{PAx}_{-1}$
  $= \mathbf{r}_{-1} - \mathbf{AW}\mu$
  $= \mathbf{b} - \mathbf{Ax}_0$ where $\mathbf{XAW}\mu = \mathbf{Xr}_{-1}$, $\mathbf{r}_{-1} = \mathbf{b} - \mathbf{Ax}_{-1}$ and $\mathbf{x}_0 = \mathbf{x}_{-1} + \mathbf{W}\mu$
  $\beta = \|\mathbf{r}\|_2,\ \mathbf{v}_1 = \mathbf{r}/\beta,\ \tilde{\beta} = \beta \mathbf{e}_1$
  for $i = 1$ to m **do**
    $\mathbf{w} = \mathbf{A}(\mathbf{v}_i - \mathbf{W}\mu)$ where $\mathbf{XAW}\mu = \mathbf{XAv}_i$
    !Same as before
  end for
  for $k = n_r - 1$ to 1 **do**
    $y_k = (\tilde{b}_k - \sum_{i=k+1}^{n_r} r_{k,i} y_i)/r_{k,k}$
  end for
  $\mathbf{z} = V\mathbf{y}$
  $\mathbf{x} = \mathbf{x}_0 + \mathbf{z} - \mathbf{W}(\mathbf{XAW})^{-1}\mathbf{XAz}$
  if (converged) **exit**
end while

3. Deflation applied to the Helmholtz equation

In this section, the deflated complex GMRES is applied to the Helmholtz equation. First, the equation is recalled, along with its boundary conditions. Then, some spectral properties of the discrete system are commented to better understand the behavior of the solver. The deflation space is presented thereafter. Finally, some possible discrete improvements are discussed compared to the basic version.

3.1. The Helmholtz equation

The Helmholtz equation is derived from the linearized equations for compressible flows [63]. For a bounded domain Ω with boundary Γ, it reads:

$$-\Delta u - k^2 u = 0 \quad \text{in } \Omega \tag{3.1}$$

$$\frac{\partial u}{\partial \mathbf{n}} + \beta u = g \quad \text{on } \Gamma_R \tag{3.2}$$

$$u = u_0 \quad \text{on } \Gamma_D \tag{3.3}$$

where k is the wave number:

$$k = \frac{\omega}{c} \tag{3.4}$$

and ω is the studied frequency, c is the fluid velocity, $\mathbf{n}$ is the surface normal pointing outwards and β is a complex number. Γ_D is the part of the boundary where the Dirichlet boundary condition applies and Γ_R is the part of the boundary where Neumann or Robin boundary conditions applies. For the former, $\beta = 0$ and for the latter $\beta = -ik$. Eq. (3.2) for Robin boundary conditions is an approximation of the Sommerfeld boundary condition:

$$\lim_{r\to\infty} r^{\frac{d-1}{2}} \left(\frac{\partial u}{\partial r} - iku \right) = 0 \tag{3.5}$$

where d is the space dimension. Supposing Eq. (3.1) is discretized with finite differences or low order finite elements with lumped mass matrix, it is clearly seen that the mass term adds a negative contribution to the diagonal, which may rise difficulties for iterative methods. The matrix is symmetric but indefinite without Sommerfeld boundary conditions. It is symmetric complex when the Sommerfeld boundary conditions are used. After discretization, the algebraic system reads:

$$(\mathbf{K} - k^2\mathbf{M} - ik\mathbf{M}_b)\mathbf{u} = \mathbf{f} \tag{3.6}$$

where $\mathbf{K}$ is the Laplacian matrix, $\mathbf{M}$ is the mass matrix, $\mathbf{M}_b$ is the boundary mass matrix where the Robin boundary condition applies, and $\mathbf{f}$ is given by the Neumann boundary condition and possible external volume forces.

3.2. Spectral properties of the discrete Helmholtz equation

From a discrete spectral viewpoint, the eigenvalues of the discrete second order one dimensional finite difference Laplacian with homogeneous Dirichlet boundary conditions on a unit length domain are [5,64]:

$$\lambda_i = \frac{2}{h^2}\left(1 - \cos\left(\frac{i\pi}{N+1}\right)\right) \tag{3.7}$$

For SPD matrices, the convergence of the iterative solver is related, in particular, to the condition number:

$$\kappa_2(\mathbf{A}) = \|\mathbf{A}\|_2\|\mathbf{A}^{-1}\|_2 = \frac{\lambda_{max}}{\lambda_{min}} \tag{3.8}$$

Therefore, if the next eigenvalue λ_{min+1} after λ_{min} is such that the ratio $\lambda_{max}/\lambda_{min+1}$ is much smaller than the condition number of $\mathbf{A}$, the removal of λ_{min} will lead to an increase speed in convergence. Equivalently, it means that convergence will speed-up if the ratio $\lambda_{min+1}/\lambda_{min}$ is large. Fig. 3.1 displays the ratio λ_{i+1}/λ_i for $N = 10^5$. It is clearly seen that the largest ratio occurs for the smallest eigenvalues, meaning that removal of a few eigenvalues will have a drastic effect. This explains the drastic improvement due to deflation where it is not necessary to treat all the spectrum, as a multigrid would do. Furthermore, this effect is stronger when the dimension increases as illustrated in Fig. 3.3. For a one dimensional operator, the eigenvalues are evenly distributed from the smallest to the largest, while as the dimension increases, the medium part of the spectrum increases. Therefore the gains of removing the smallest eigenvalues are still larger in three dimensions as they are more scattered at both ends of the spectrum.

For the one dimensional finite difference Helmholtz operator with homogeneous Dirichlet boundary conditions on a unit length domain, the eigenvalues are given by [38]:

$$\lambda_i = \frac{2}{h^2}\left(1 - \cos\left(\frac{i\pi}{N+1}\right)\right) - k^2 \tag{3.9}$$

By imposing the engineering constraint of having ten mesh points per wavelength, namely:

$$kh = \frac{2\pi}{10} \tag{3.10}$$

and using the relationship:

$$h = \frac{1}{N+1} \tag{3.11}$$

one obtains:

$$ih = 0.2 \tag{3.12}$$

which provides an estimate of the number of negative eigenvalues on the finest grid. As the mesh is refined, more and more negative eigenvalues appear in the system for the given constraint. Fig. 3.2 displays the function λ_{i+1}/λ_i for $N = 10^2$. It is clearly seen that this ratio is highest for the eigenvalues close to the wave number squared. However, the ratio of the largest to the smallest absolute eigenvalue is not a factor influencing convergence for non SPD matrices. This is illustrated in Tables 3.1 and 3.2 for a three dimensional cavity example presented below with Neumann and Robin boundary conditions respectively, where the iteration number has been reported against the ratio of the largest to the smallest eigenvalue in absolute value. The eigenvectors associated with the lowest eigenvalues have been displayed for both cases on Figs. 3.4 and 3.5 respectively. As expected, they become more and more oscillatory as the wave number increases

3.3. Plane waves

For the deflation to be effective, modes associated with low eigenvalues should be identified, either analytically or algebraically, and should be well represented in the deflated subspace. A natural extension for deflation techniques to the Helmholtz equation relies on introducing plane waves in the deflation space. Therefore $\mathbf{W}$, the deflation space, is given by:

$$\mathbf{W}_j^k = e^{i\mathbf{k}\cdot\mathbf{x}} \tag{3.13}$$

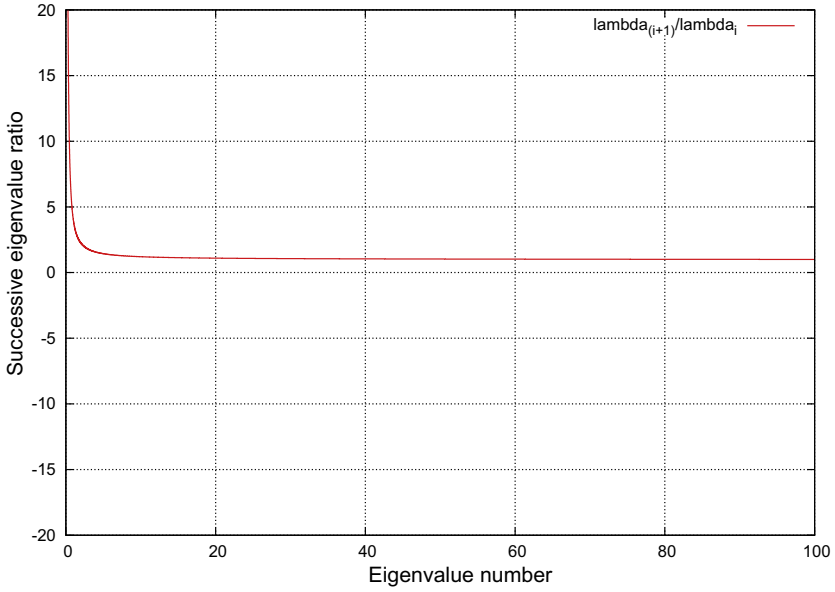

Fig. 3.1. Eigenvalue ratio for successive eigenvalues of the Laplacian operator. The highest ratio occurs at the low end of the spectrum. Therefore, the gains obtained by removing the lowest eigenvalues are high.

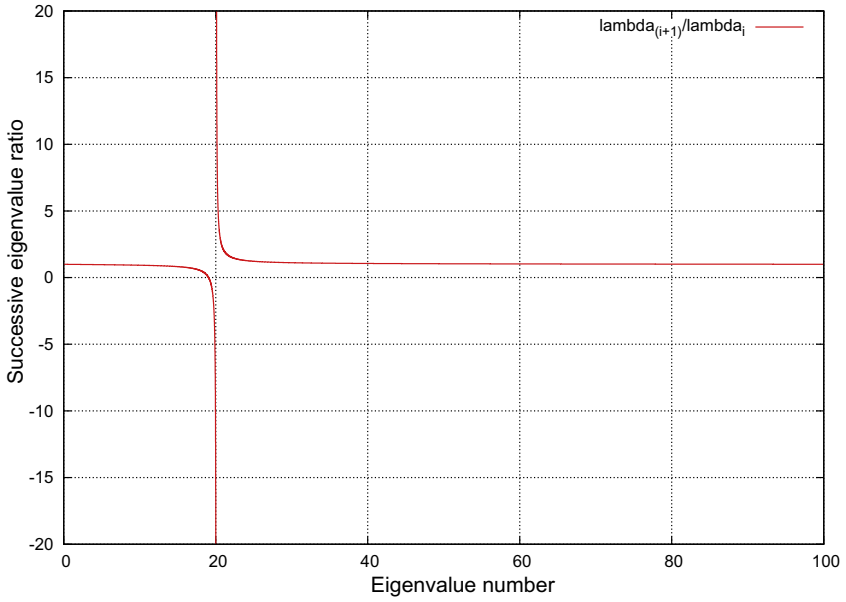

Fig. 3.2. Eigenvalue ratio for successive eigenvalues in absolute value of the Helmholtz operator. The highest ratio occurs for eigenvalues close to the wave number squared.

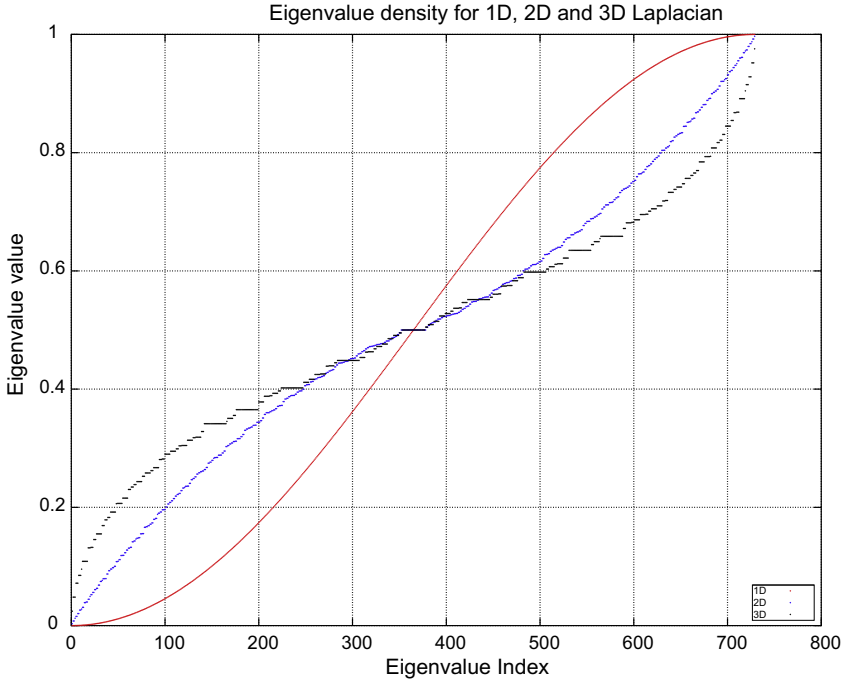

Fig. 3.3. Normalized eigenvalue distribution for the 1D, 2D and 3D Laplacian operator. The operator is scarcer at both ends when the dimension increases. The ratio of successive eigenvalues therefore increases at the low end of the spectrum with the dimension.

Table 3.1
GMRES iteration number and condition number with various wave numbers for the cavity example with Neumann boundary conditions. The condition number does not drive convergence.

	GMRES ite.	Condition number $\lvert\lambda_{max}\rvert/\lvert\lambda_{min}\rvert$ (10^3)
$k = 0$	308	26
$k = 1$	322	44
$k = 2$	339	42
$k = 4$	368	18
$k = 8$	521	44
$k = 16$	666	22

Table 3.2
GMRES iteration number and condition number with various wave numbers for the cavity example with Robin boundary conditions. Again, the condition number is not a relevant factor to convergence.

	GMRES ite.	Condition number $\lvert\lambda_{max}\rvert/\lvert\lambda_{min}\rvert$ (10^3)
$k = 0$	307	26
$k = 1$	315	32
$k = 2$	334	44
$k = 4$	353	8
$k = 8$	493	83
$k = 16$	668	4

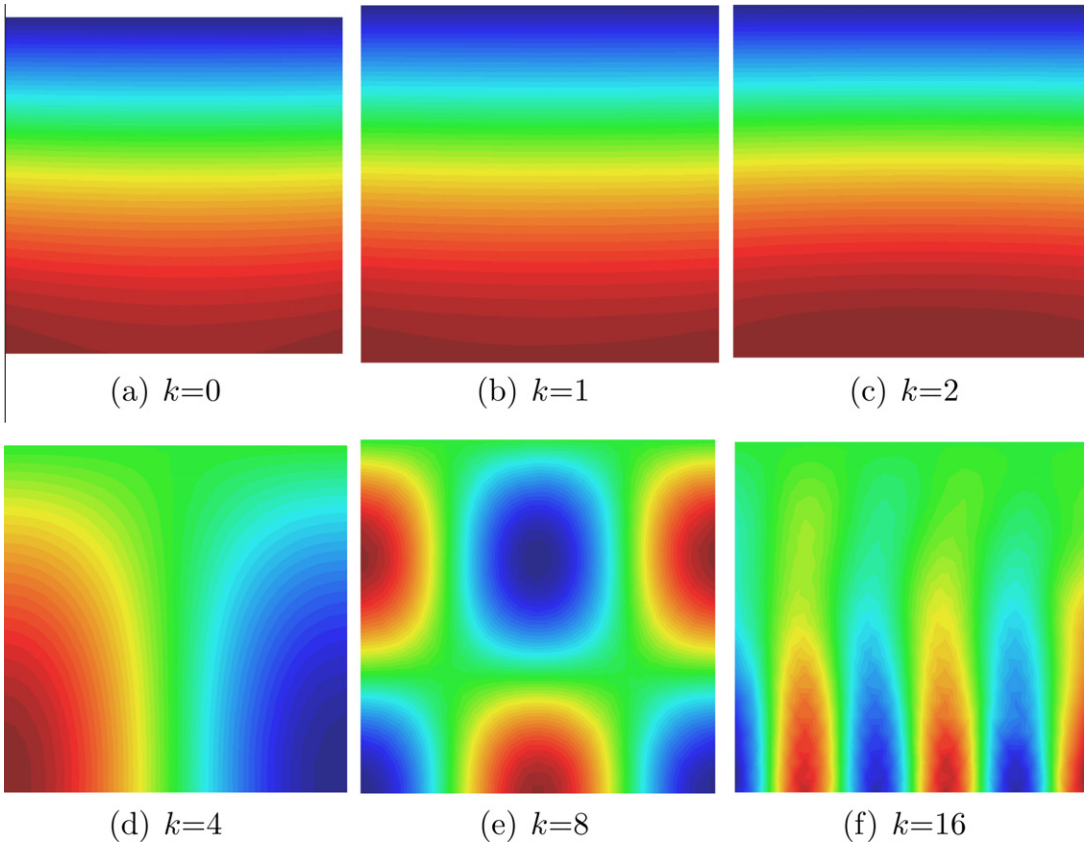

Fig. 3.4. Eigenmodes associated with lowest eigenvalue for the cavity example at various wave numbers with Neumann boundary conditions.

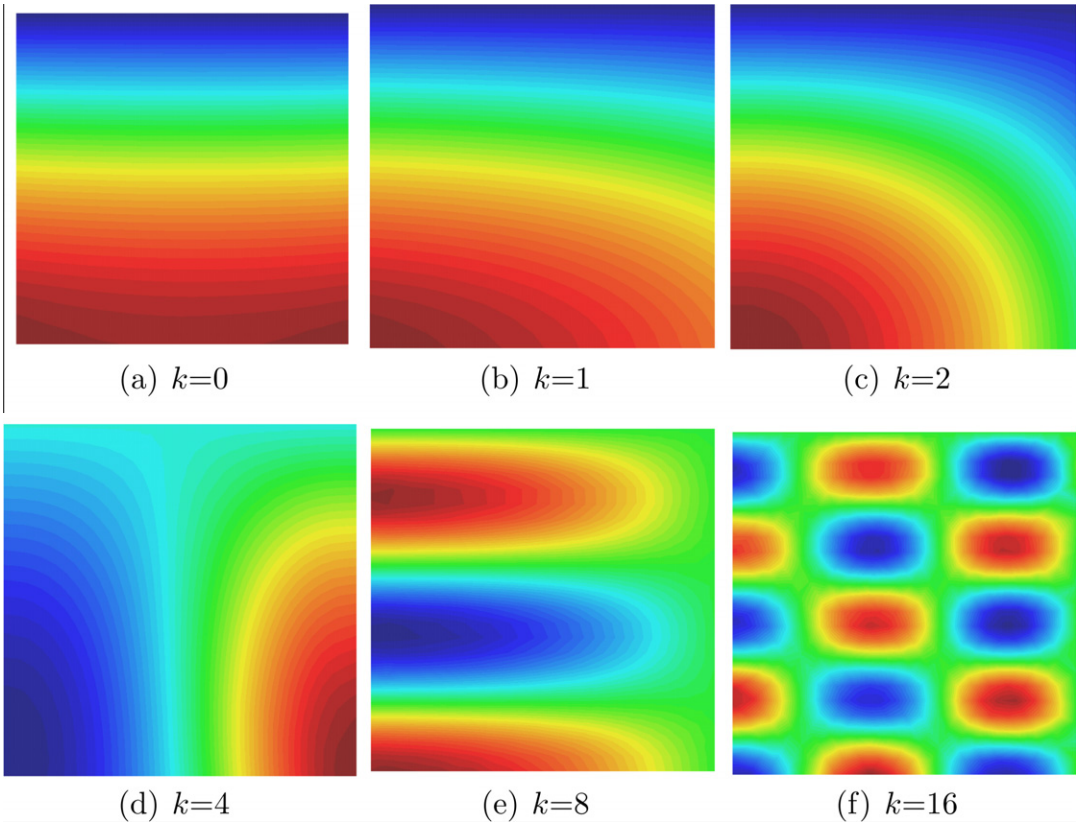

Fig. 3.5. Eigenmodes associated with lowest eigenvalue for the cavity example at various wave numbers with Robin boundary conditions. The condition number is not a relevant factor to convergence.

where a node **M** of coordinates **x** belongs to group j with direction **k**. The small system to be solved after an agglomeration of size `ngroup` has dimension `ngroup*ndir` where `ndir` is the number of directions chosen for the plane waves. Compared to the scalar Laplace deflation case, the small subsystem has been multiplied by `ndir`. This deflation space is obviously reminiscent from the work in [33,39].

If the definition of **W** is rather straightforward, the definition of **X** is not. In the three references relying on a multigrid approach mentioned in the introduction, namely [33,36,39], differences do appear for the definition of the restriction operator, played by **X** in a deflation context as seen in Eq. (2.14). In [39], the symmetric approach is applied to the Helmholtz equation. In [36], the restriction operator is defined as the adjoint of the interpolation operator, therefore suggesting the conjugate transpose for **X**. Relying on exponential interpolation allows to define a prolongation operator that trivially transfers the nearly singular eigenvectors from the coarse grid to the fine grid [50] but the lack of Hermitian character of the Helmholtz operator does not allow for an easy restriction. Finally, [33] certainly provides the answer for the restriction operator in what is called the separation procedure. Given the residual after a classical V cycle, problematic components have the form:

$$\mathbf{r} = \sum_k \mathbf{r}_k e^{i\mathbf{k}\cdot\mathbf{x}} \tag{3.14}$$

The iterative process will be much more efficient at reducing the geometric smooth part $\mathbf{r}_k$ than the highly geometric oscillatory component, or algebraic smooth part $\mathbf{r}_k e^{i\mathbf{k}\cdot\mathbf{x}}$. Obtaining the value $\mathbf{r}_k$ is equivalent to perform a convolution on the complete residual, which physically illustrates the non Hermitian behavior of the Helmholtz operator, as it would require building an orthogonal basis of the oscillatory components, which is not $\overline{\mathbf{W}}^T$. Therefore, the residual is multiplied by $e^{-i\mathbf{k}\cdot\mathbf{x}}$ so that its new value is:

$$\mathbf{r} = \mathbf{r}_k + \sum_{j\neq k} \mathbf{r}_j e^{i(\mathbf{j}+\mathbf{k})\cdot\mathbf{x}} \tag{3.15}$$

Now a full weighting is applied to simulate the action of a low pass filter in order for $\mathbf{r}_k$ to be extracted from **r**. It seems therefore that a conjugate transpose would be more appropriate than only the transpose, even though the convolution is only approximated with the conjugate transpose. As the Helmholtz operator is not Hermitian, there is no necessity to use $\mathbf{X} = \overline{\mathbf{W}}^T$ as the restriction operator, as the eigenvectors do not form an orthogonal basis with respect to the complex scalar product.

A strict application of the non Hermitian deflation would require **X** to be composed of the left eigenvectors, relying on the well known Hotelling deflation [42], in order to leave the part of the spectrum and the eigenvectors that do not belong to the deflation unchanged. As mentioned before, the deflation projector reads:

$$\mathbf{P} = \mathbf{I} - \mathbf{AW}(\mathbf{XAW})^{-1}\mathbf{X} \tag{3.16}$$

In a matrix by vector multiplication, the action of the operator **PA** will be applied to the Krylov basis. Therefore, considering the residual $\mathbf{r} = \mathbf{b} - \mathbf{Ax}$, we obtain:

$$\mathbf{PAr} = \mathbf{Ar} - \mathbf{AW}(\mathbf{XAW})^{-1}\mathbf{XAr} \tag{3.17}$$

Supposing **W** is composed of the right eigenvectors and **X** of the left eigenvectors, we have the relation $\mathbf{AW} = \mathbf{WD}$ and $\mathbf{XW} = \mathbf{I}$, where **D** is the diagonal matrix composed of the eigenvalues, and the left and right eigenvectors have been normalized. Writing a modal decomposition of the initial residual, we obtain:

$$\mathbf{PAr} = \mathbf{Ar} - \sum_i \mathbf{AW}(\mathbf{XWD})^{-1}\mathbf{XA}\alpha_i\mathbf{w}_i \tag{3.18}$$

Therefore:

$$\mathbf{PAr} = \mathbf{Ar} - \sum_i \mathbf{AW}(\mathbf{D})^{-1}\mathbf{X}\lambda_i\alpha_i\mathbf{w}_i \tag{3.19}$$

And:

$$\mathbf{PAr} = \mathbf{Ar} - \sum_i \mathbf{A}\alpha_i\mathbf{w}_i \tag{3.20}$$

Optimally, the small matrix inversion should only provide a scaling by the eigenvalues and, as said before, the application of **X** should provide the value of the components of the residual with respect to their modal decomposition.

Ultimately, from a practical viewpoint, there are no physical reasons not to treat the propagation vectors in **W** in both directions so that $\mathbf{W}^T\mathbf{AW}$ and $\overline{\mathbf{W}}^T\mathbf{AW}$ applied to opposite directions will generate the same coarse operator anyway. Therefore, in this paper, the choice $\mathbf{X} = \mathbf{W}$ has been followed for efficiency and storage purposes.

3.4. Discrete comparisons

As already pointed out, there are numerous difficulties in solving the Helmholtz equation. One of them lies in the fact that exponential oscillatory functions do not belong to the discrete space, contrarily to the symmetric positive definite case, where constant functions do belong to the discrete space. Another difficulty is related to the increasing number of problematic eigenvalues with the wave number.

Regarding the pollution effect and considering the one dimensional finite element stencil, the characteristic equation reads [16,63]:

$$(2\alpha+1)x^2 + (2(4\alpha-1))x + (2\alpha+1) = 0 \tag{3.21}$$

with:

$$\alpha = \frac{(kh)^2}{12} \tag{3.22}$$

A solution for an outgoing wave reads:

$$x_1 = \frac{(1-4\alpha) + \sqrt{12\alpha(\alpha-1)}}{2\alpha+1} \tag{3.23}$$

The dispersion for the discrete wave number reads:

$$k_h h = \arccos\frac{1-4\alpha}{1+2\alpha} \tag{3.24}$$

Or:

$$k_h = k - \frac{k^3h^2}{24} + O(k^5h^4) \tag{3.25}$$

The last equation shows that discrete waves are dispersive, but more importantly that discrete wave numbers are slowed down as k increases. This argument has been used to try to stabilize the standard FEM [16]. If k is given by the problem at hand, it may be interesting to use k_{num}, a numerical wave number with a slightly lower magnitude for interpolation, and coarse grid operator evaluation. An improvement in convergence would then show that the discrete operator has substantially moved from the continuous operator. Some tests were conducted in this direction leading to minor improvements. The main reason is due to the fact that in higher dimensions, the discrete eigenvectors depart from their continuous counterpart in a much more complex manner.

Finally, it is interesting to compare the coarse grid operator obtained by the deflation coarse grid operator, and the wave-ray multigrid. In the former, the coarse grid correction is given by the operator **XAW** which, for a set of n directions `ndir`, takes into account the coupling between directions. The size of the small matrix is therefore the size of the graph of the subdomains multiplied

by `ndir` [2]. Clearly, the overhead is noticeable for a large set of directions. In the latter, the operator derived for each direction is solved at each ray level, without taking the couplings into account. However, experiments with deflation revealed that removing the coupling in the direct solve was affecting negatively and in a substantial manner the convergence rate.

4. Numerical examples

In this section, numerical examples are studied to support the theoretical arguments presented before, and illustrate the proposed method. First, a simple example with analytical solution is thoroughly studied. Then realistic examples are presented. Contrarily to the SPD case, numerous parameters influence the result. Apart from the group number for the agglomeration procedure, the direction number per group has to be chosen. Due to the non Hermitian character, the restart parameter should also be determined in the GMRES procedure for realistic examples.

4.1. Square cavity

This example has been extracted from [24]. It allows to compare the computed solution to an analytical solution with and without the approximated Sommerfeld condition. The square is given by $[0:1]\times[0:1]$, and has been slightly extruded in the third direction to run in 3D. The analytical solution is:

$$u(x,y) = e^{i\left(\sqrt{k^2-\frac{\pi^2}{4}}\right)x}\cos\frac{\pi}{2}y \tag{4.1}$$

At the top of the cavity, u is set to 0. At the bottom, a Neumann boundary condition is set so that $\frac{\partial u}{\partial n}=0$. In the first configuration, $\frac{\partial u}{\partial n}$ is set to the analytical value on both sides. In the second configuration, the $\frac{\partial u}{\partial n}=iku$ is set on the right boundary. Fig. 4.1 illustrates the boundary conditions. A nice feature of this example relies on the fact that the analytical solution is in fact a sum of two waves:

$$u(x,y) = e^{i\left(\sqrt{k^2-\frac{\pi^2}{4}}\right)x}\frac{\left(e^{i\frac{\pi}{2}y}+e^{-i\frac{\pi}{2}y}\right)}{2} \tag{4.2}$$

or:

$$u(x,y) = \frac{1}{2}e^{i\left(\sqrt{k^2-\frac{\pi^2}{4}}x+\frac{\pi}{2}y\right)}+\frac{1}{2}e^{i\left(\sqrt{k^2-\frac{\pi^2}{4}}x-\frac{\pi}{2}y\right)} \tag{4.3}$$

Therefore, the two wave numbers are given by:

$$\mathbf{k}_1=\begin{pmatrix}\sqrt{k^2-\frac{\pi^2}{4}}\\ \frac{\pi}{2}\end{pmatrix},\quad \mathbf{k}_2=\begin{pmatrix}\sqrt{k^2-\frac{\pi^2}{4}}\\ -\frac{\pi}{2}\end{pmatrix} \tag{4.4}$$

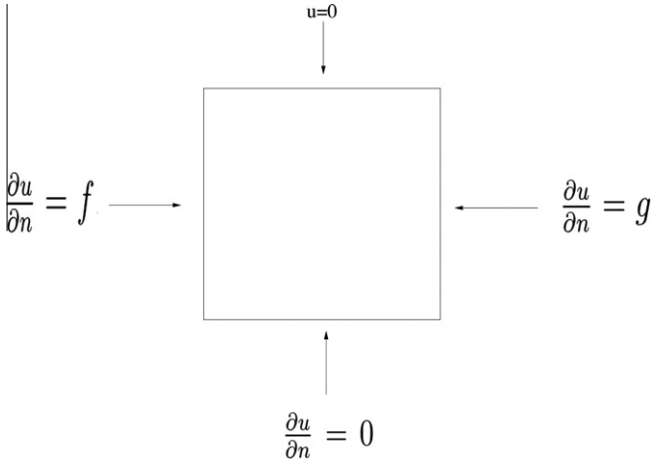

Fig. 4.1. Cavity example.

Supposing the exponential functions would belong to the discretization space, the solver would converge in one iteration providing that the deflated space is composed of two waves with exactly these directions. Unfortunately, the exponential does not belong to the discrete space. This is a major difference with the Poisson equation where the low energy modes are the constant one, which as commented before, do belong to the discrete space. Furthermore, a general solution will typically be composed of more waves, where each direction is generally unknown.

The solution is displayed on Fig. 4.2 for different wave numbers. It is clearly seen that as k increases, both wave directions align more and more along the x abscissa, as suggested by Eq. (4.4). Note that for $k=1$, evanescent waves are present in the solution while the solution is purely oscillatory for higher values of k. The mesh contains $23\cdot10^3$ points and $70\cdot10^3$ elements.

4.1.1. Convergence of Robin and Neumann bc with conjugate transpose

The iterations necessary to converge with a relative tolerance of 10^{-5} are given in Fig. 4.3 without deflation, with deflation composed by the two exact waves and one group, and by only one of these waves and one group. The case with one wave was added to show that the increase in the iteration count for the case of the two exact waves is not due to the worse and worse collinearity of the two waves. Compared to the symmetric definite positive case, a very annoying phenomenon appears, as subdomain deflation may actually worsen convergence. The insertion of the two waves gives rise to a small 2 by 2 system to be solved. It is only performed for illustration purposes as the dimension of the deflated space is obviously too small regarding efficiency. However, this should give the exact solution *in one iteration* if the separation process or convolution would be perfect and the exponential would describe exactly the discrete eigenvector. It is clear that as the wave number increases, the exponential moves away from a representation of a discrete eigenmode. The decrease in iteration counts for small to medium values is noticeable though. For this small system, the restart of GMRES was not used in order to simplify the analysis. Note the increase in the iteration count for this simple example. GMRES needs 10 percent of the point number to converge for $k=32$, which is clearly prohibitive for large problems. It clearly illustrates the difficulties that iterative methods meet with indefinite operators.

The same experience is conducted with Robin boundary conditions at the right boundary. Iterations counts are reported on Fig. 4.4. From an iteration viewpoint, as the left boundary is real valued, complex values will not appear in the inner iterations of GMRES until the number of iterations equals the graph depth to come from the left side to the right side. Larger iteration counts can therefore be forecast for Robin boundary conditions.

4.1.2. Convergence of Robin and Neumann bc with transpose only

Next, $\mathbf{X}=\mathbf{W}^T$, is replaced by $\mathbf{X}=\mathbf{W}^T$ to study the effects of the conjugacy and to validate the assumptions made on the separation process. The iterations necessary to converge with a tolerance of 10^{-5} are given in Fig. 4.5 without deflation, with deflation composed by the two exact waves and one group, and by only one of these waves and one group, as before. Except for $k=2$, the choice $\mathbf{X}=\mathbf{W}^T$ consistently outperforms the choice $\mathbf{X}=\mathbf{W}^T$. The case $k=2$ is easily explained by the fact that the wave number considered is low, and no Robin boundary condition is introduced. Therefore the character of the matrix is almost real SPD, for which the eigenvectors are orthogonal if $\mathbf{X}=\mathbf{W}^T$ is considered. Apart from that, the same characteristics as before are met, namely an increase in the iteration count with an increase in the wave number.

The same experience is conducted with Robin boundary conditions at the right boundary and is reported on Fig. 4.6. Now the

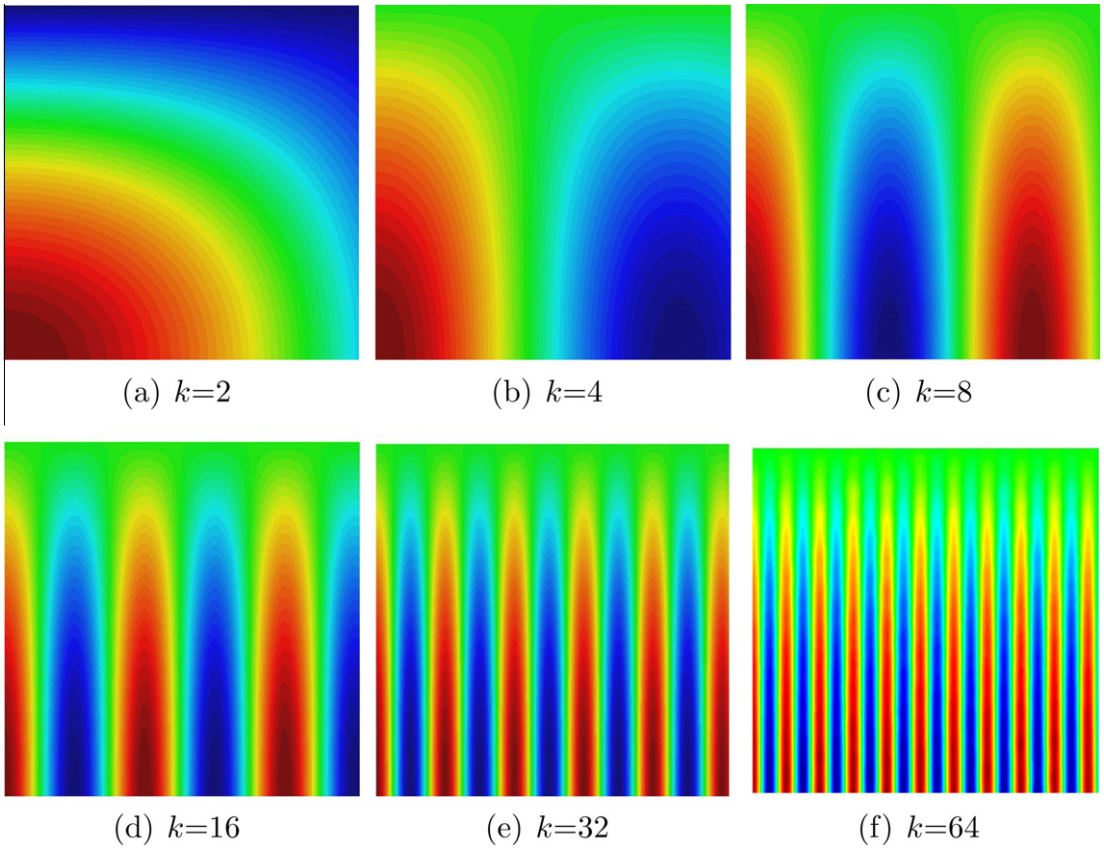

Fig. 4.2. Pressure disturbance for the cavity example at various wave numbers.

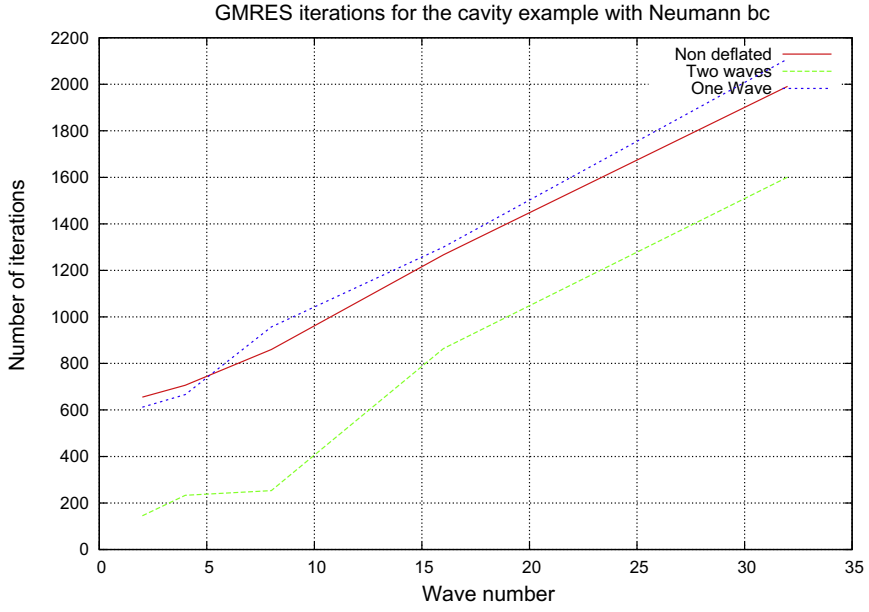

Fig. 4.3. Iteration count for the cavity example with conjugate transpose and Neumann boundary condition.

choice $\mathbf{X} = \overline{\mathbf{W}^T}$ consistently outperforms the choice $\mathbf{X} = \mathbf{W}^T$ for all wavenumber value. The introduction of the Robin boundary condition reinforces the symmetric complex character, for which the choice $\mathbf{X} = \mathbf{W}^T$ is not adapted. It appears to be clear that $\mathbf{X} = \overline{\mathbf{W}^T}$ is a reliable choice for oscillatory solutions, as discussed in Section 3.3.

4.1.3. Discrete eigenvector

After having investigated the choice of space for $\mathbf{X}$, the gap between the discrete and the continuous eigenvector is studied. Table 4.1 gives the nodal contribution of $\mathbf{AW}$ for a point fully surrounded by the same group normalized by $\mathbf{K}_{ii} + k^2\mathbf{M}_{ii}$ where $\mathbf{K}$ is the Laplacian matrix and $\mathbf{M}$ the mass matrix, in order to obtain a positive

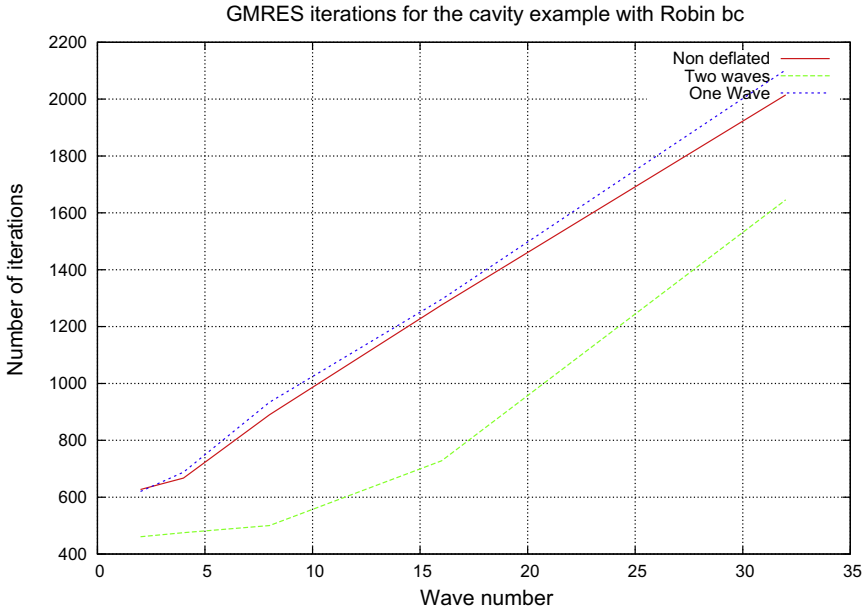

Fig. 4.4. Iteration count for the cavity example with conjugate transpose and Robin boundary condition.

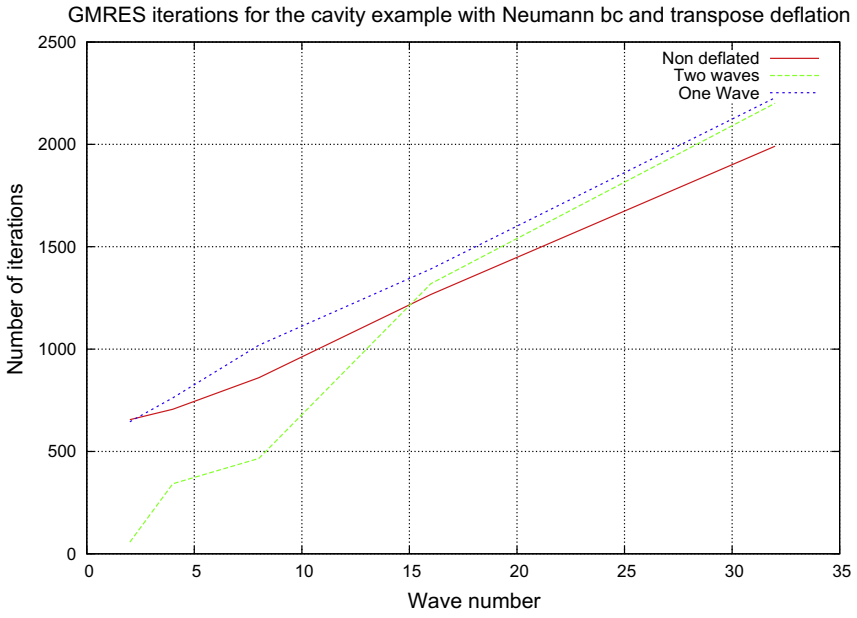

Fig. 4.5. Iteration count for the cavity example with transpose only and Neumann boundary condition.

non zero value. As **W** is a free space solution of the Helmholtz operator, **AW** should be machine zero for all the directions as it is for a constant eigenvector for the Laplacian. It is not, as its discrete counterpart does not belong to the discrete space considered here, but as expected, it is clearly seen that it increases as the wave number increases.

4.1.4. Influence in the group and direction number

Finally, the same example is studied with respect to the group and direction number. Regarding the direction, three stencils were chosen, namely 6, 26 and 56 directions. Results are reported in Table 4.2 for a constant interpolation used typically for SPD matrices, in Table 4.3 for a stencil with 6 directions, in Table 4.4 for a stencil of 26 directions and in Table 4.5 for a stencil with 56 directions with Neumann boundary conditions. Very similar results are obtained with Robin boundary conditions. A dash indicates an error in the factorization. A simple complex skyline LU direct solver is used in this study without pivoting. A more advanced solver certainly performs better but does not alter the main characteristics of the discussion. As observed, an error in the factorization appears when the direction number increases, or when the wave number decreases. It is fully consistent with the fact that the increase in the direction number gives rise to a more and more ill-conditioned system if the wave number does not increase. However, the stencil with 6 directions already provides a substantial acceleration in the GMRES convergence. Typically, the iteration number does not seem to be affected by the increase in the wave number, until it reaches a value for which it begins to increase. By increasing the

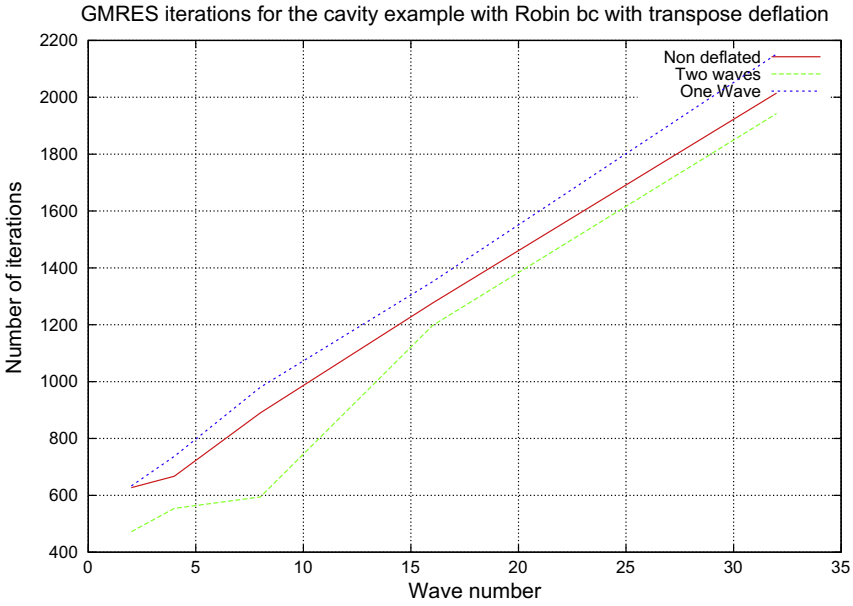

Fig. 4.6. Iteration count for the cavity example with transpose only and Robin boundary condition.

Table 4.1
Nodal contribution of **AW** for a point fully surrounded by the same group for the square example.

Wave number	2	4	8	16	32
1st wave Real	$-6\cdot10^{-5}$	$1\cdot10^{-4}$	$-1\cdot10^{-3}$	$2\cdot10^{-4}$	$-4\cdot10^{-3}$
1st wave Imag.	$2\cdot10^{-7}$	$-2\cdot10^{-4}$	$1\cdot10^{-4}$	$-4\cdot10^{-3}$	$1\cdot10^{-2}$
2nd wave Real	$5\cdot10^{-5}$	$-1\cdot10^{-4}$	$1\cdot10^{-3}$	$-4\cdot10^{-4}$	$4\cdot10^{-3}$
2nd wave Imag.	$-1\cdot10^{-7}$	$2\cdot10^{-4}$	$-5\cdot10^{-4}$	$4\cdot10^{-3}$	$-1\cdot10^{-2}$

Table 4.2
Iteration number for various wave numbers and group numbers with a constant interpolation and Neumann boundary conditions for the cavity example.

ngrou	0	10	100	500
$k=2$	655	423	192	100
$k=4$	706	506	223	109
$k=8$	860	783	361	177
$k=16$	1267	1159	629	313
$k=32$	1991	2017	1479	825

Table 4.3
Iteration number for various wave numbers and group numbers with a stencil of 6 directions and Neumann boundary conditions for the cavity example.

ngrou	0	10	100	500
$k=2$	655	171	71	31
$k=4$	706	174	71	31
$k=8$	860	239	76	31
$k=16$	1267	567	87	31
$k=32$	1991	1657	336	48

wave number, the system becomes less and less singular. For each of the three stencils and for each number of group tested, there is first an increase in the iteration number, certainly due to the almost singular system, and then a decrease. It also does not seem to pay off to increase the direction number compared to the group number, as the assembly cost is noticeably higher, the solver is less

Table 4.4
Iteration number for various wave numbers and group numbers with a stencil of 26 directions and Neumann boundary conditions for the cavity example.

ngrou	0	10	100
$k=2$	655	–	–
$k=4$	706	199	–
$k=8$	860	81	–
$k=16$	1267	98	28
$k=32$	1991	515	38

Table 4.5
Iteration number for various wave numbers and group numbers with a stencil of 56 directions and Neumann boundary conditions for the cavity example.

ngrou	0	10	100
$k=2$	655	–	–
$k=4$	706	–	–
$k=8$	860	–	–
$k=16$	1267	210	–
$k=32$	1991	225	–

robust for low wave number, and the iteration count is not substantially reduced. However, very good convergence is obtained overall, even for high wave numbers. Furthermore, an almost perfect monotonic behavior is observed between the increase of the group number and the decrease of the iteration count, including the case without deflation. This will not be the case for real three dimensional examples, besides the impediment to use full GMRES, and illustrates the fact that it is not the same to solve small and large problems with iterative solvers.

4.2. Cylinder scattering

This example is a classical benchmark for the Helmholtz equation with analytical solution. Two concentric cylinders define an inner volume. The inner cylinder diameter is 0.5 and the outer one is 2. This example was again run in three dimensions with a slight extrude of the two dimensional mesh. The mesh contains $280\cdot10^3$ points and $860\cdot10^3$ elements. The full GMRES is given

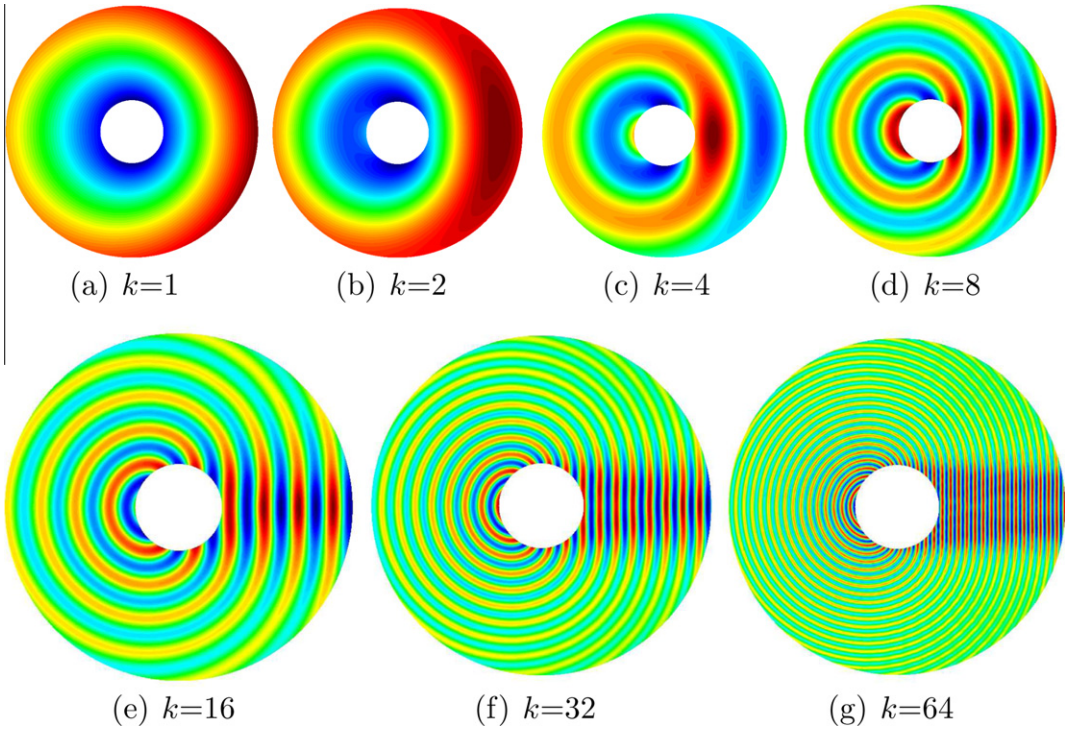

Fig. 4.7. Pressure disturbance for the cavity example at various wave numbers.

up for the restarted GMRES to run in a more applied context. The Krylov subspace used for this example was fixed to 100. Results obtained for different wave numbers are displayed on Fig. 4.7.

Iterations obtained for different wave numbers and group numbers are reported on Table 4.6 for a constant interpolation, on Table 4.7 for a stencil of 6 directions, and on Table 4.8 for a stencil of 26 directions. The same trends as in the previous example are observed. For a fixed k, increasing the group number or the direction number produces a drastic reduction of the iteration number. However, as k increases, there seems to be a value from which the fast convergence stops suddenly and the deflation produces a larger iteration number than GMRES without deflation. Note the non monotonic behavior of GMRES with respect to the wave number. Again, the system is not invertible for a large number of directions and a low wave number, as expected. Apart from this drawback, the increase in direction does come with a faster convergence. Finally, in order to converge fast at high wave number (k = 64), it is seen that the coarse grid size is of the order of 20 percent the size of the fine grid, almost the double as one coarsening level in three dimensions, and the beginning of slow down of

Table 4.6
Iteration number and CPU time for various wave numbers and group numbers with a constant interpolation for the cylinder example.

ngrou	0	$6 \cdot 10^2$	$6 \cdot 10^3$	$6 \cdot 10^4$
k = 1	3096 (1385)	118 (38)	120 (50)	20 (17)
k = 2	3516 (1007)	184 (62)	56 (25)	22 (13)
k = 4	5392 (2681)	496 (218)	120 (77)	35 (22)
k = 8	2264 (718)	2380 (1061)	444 (237)	71 (37)
k = 16	2433 (770)	10^4	1832 (1036)	247 (180)
k = 32	2437 (774)	10^4	10^4	1055 (955)
k = 64	2586 (897)	10^4	10^4	5982

Table 4.7
Iteration number and CPU time for various wave numbers and group numbers with a stencil of 6 directions for the cylinder example.

ngrou	0	10^2	10^3	10^4
k = 1	3096 (1385)	64 (14)	22 (7)	6 (12)
k = 2	3516 (1007)	74 (15)	25 (8)	7 (13)
k = 4	5392 (2681)	95 (23)	30 (9)	9 (14)
k = 8	2264 (718)	296(89)	37 (11)	12 (17)
k = 16	2433 (770)	10^4	81 (27)	14 (18)
k = 32	2437 (774)	10^4	3528 (1237)	20 (28)
k = 64	2586 (897)	10^4	10^4	182 (139)

Table 4.8
Iteration number and CPU time for various wave numbers and group numbers with a stencil of 26 directions for the cylinder example.

ngrou	0	10^2	10^3
k = 1	3096 (1385)	–	–
k = 2	3516 (1007)	–	–
k = 4	5392 (2681)	30 (116)	–
k = 8	2264 (718)	45 (117)	14 (114)
k = 16	2433 (770)	77 (162)	16 (116)
k = 32	2437 (774)	10^4	23 (122)
k = 64	2586 (897)	10^4	234 (433)

Table 4.9
Nodal contribution of **AW** for a point fully surrounded by the same group for the cylinder scattering.

ndir	1	2	3	4	5	6
k = 1	$8 \cdot 10^{-7}$	$8 \cdot 10^{-7}$	$3 \cdot 10^{-6}$	$3 \cdot 10^{-6}$	$-4 \cdot 10^{-7}$	$-4 \cdot 10^{-7}$
k = 2	$8 \cdot 10^{-7}$	$8 \cdot 10^{-7}$	$1 \cdot 10^{-3}$	$1 \cdot 10^{-3}$	$-2 \cdot 10^{-6}$	$-2 \cdot 10^{-6}$
k = 4	$-2 \cdot 10^{-5}$	$-2 \cdot 10^{-5}$	$4 \cdot 10^{-3}$	$4 \cdot 10^{-3}$	$-8 \cdot 10^{-6}$	$-8 \cdot 10^{-6}$
k = 8	$4 \cdot 10^{-5}$	$4 \cdot 10^{-5}$	$2 \cdot 10^{-2}$	$2 \cdot 10^{-2}$	$-4 \cdot 10^{-5}$	$-4 \cdot 10^{-5}$
k = 16	$4 \cdot 10^{-5}$	$4 \cdot 10^{-5}$	$4 \cdot 10^{-2}$	$4 \cdot 10^{-2}$	$-1 \cdot 10^{-4}$	$-1 \cdot 10^{-4}$
k = 32	$-1 \cdot 10^{-3}$	$-1 \cdot 10^{-3}$	$4 \cdot 10^{-2}$	$4 \cdot 10^{-2}$	$-4 \cdot 10^{-4}$	$-4 \cdot 10^{-4}$
k = 64	$4 \cdot 10^{-3}$	$4 \cdot 10^{-3}$	10^{-1}	10^{-1}	$3 \cdot 10^{-3}$	$3 \cdot 10^{-3}$

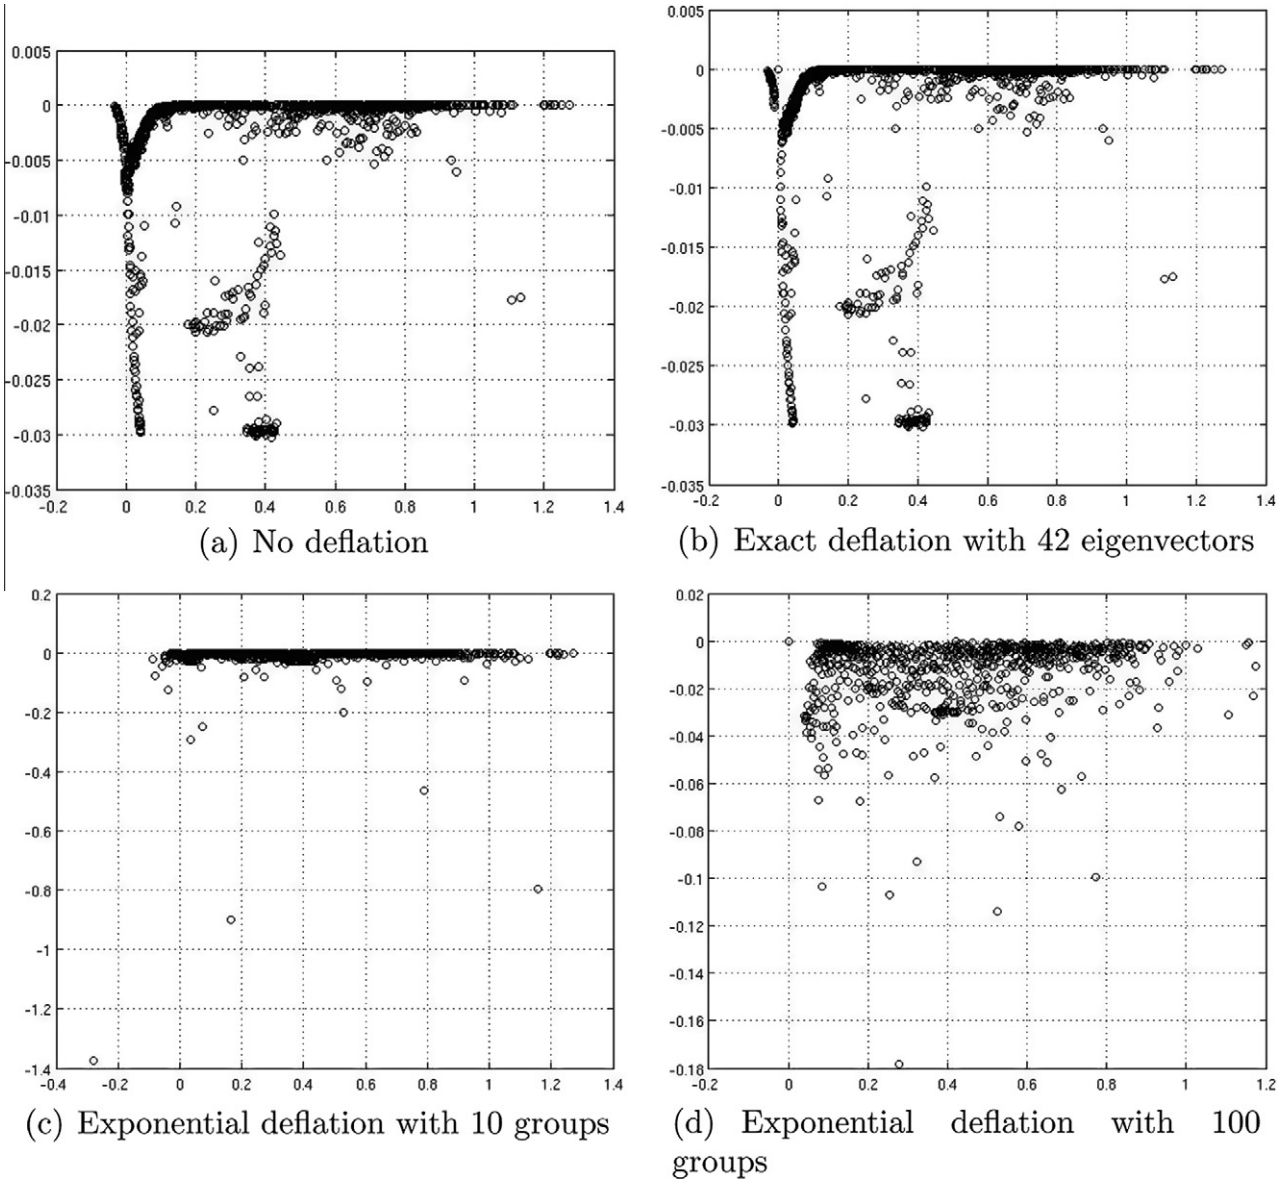

Fig. 4.8. Eigenvalue distribution with and without deflation, with exact eigenvectors and complex exponentials for the cylinder example.

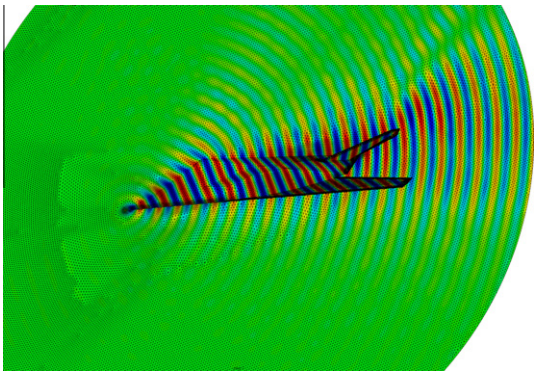

Fig. 4.9. Mesh and solution for an F117 flyer for $k = 16$.

convergence rises. Table 4.9 reports the nodal contribution of an interior node normalized as before by the diagonal of $K + k^2M$. It is again clearly seen how the analytical solution deviates from eigenvectors as k increases.

Table 4.10
Iteration number and CPU time for various wave numbers and group numbers with constant interpolation for the F117 flyer example.

ngrou	0	$6 \cdot 10^3$	$6 \cdot 10^4$
$k = 1$	1669 (9409)	196 (1067)	82 (568)
$k = 2$	1563 (7802)	750 (3230)	283 (2025)
$k = 4$	1544 (7891)	3657 (25666)	1294 (12946)
$k = 8$	1700 (9434)	10000	10000
$k = 16$	1998 (10635)	10000	10000

Compared to the SPD case, it is seen that subdomain deflation worsens convergence. After investigating with different parameters, it was found that even though convergence was worse with deflation, it was the restart procedure that could not accommodate deflation with a high wave number. The same example with a small mesh was conducted with MATLAB [65]. The eigenvalue distribution for $k = 8$ for a much smaller mesh representing 10 points per wave lengths is shown in Fig. 4.8 without deflation, with deflation with the exact eigenvectors associated to the smallest

Table 4.11
Iteration number and CPU time for various wave numbers and group numbers with a stencil of 6 directions for the F117 flyer example.

ngrou	0	10^3	10^4	10^5
$k = 1$	1669 (9409)	71 (638)	39 (1076)	20 (62120)
$k = 2$	1563 (7802)	87 (784)	44 (1031)	21 (63340)
$k = 4$	1544 (7891)	1840 (20568)	72 (1373)	24 (41852)
$k = 8$	1700 (9434)	5453 (50901)	1550 (16045)	55 (52706)
$k = 16$	1998 (10635)	7811 (55371)	10000	1324(75360)

Table 4.12
Iteration number and CPU time for various wave numbers and group numbers with a stencil of 26 directions for the F117 flyer example.

ngrou	0	10^3
$k = 1$	1669 (9409)	201 (4509)
$k = 2$	1563 (7802)	59 (1112)
$k = 4$	1544 (7891)	62 (1098)
$k = 8$	1700 (9434)	6892 (71306)
$k = 16$	1998 (10635)	10000

eigenvalue, and with the approximate exponentials for 10 groups and 100 groups. It is clearly shown that the exponential deflation spreads more the eigenvalues off of the real axis for 10 groups. For 100 groups, the same phenomenon is observed compared to the original spectrum, but all the negative eigenvalues have been removed, as well as the closest to zero. The eigenvalues are much more clustered, leading to fast convergence. For this small problem, full GMRES without deflation converges in 277 iterations, deflation with 10 groups in 262 iterations, and with 100 groups in 26 iterations. As the Krylov subspace diminishes, deflation is worse and worse for 10 groups, until becoming much worse than the non deflated GMRES. For a Krylov subspace of 10, the non deflated GMRES converges in 486 iterations while the deflated GMRES with 10 groups converges in 928 iterations, and the deflated GMRES with exact eigenvectors converges in 403 iterations for 42 eigenvectors. This illustrates that the cause of divergence of the deflation process is due to the approximate eigenvectors constituting the deflation subspace. For exact eigenvectors, deflation removes the eigenvalue from the spectrum, giving rise to the same or to a simpler polynomial approximation for the Krylov subspace.

4.3. F117 flyer

In this example, the geometry of the F117 flyer is studied. This is the first example of mildly complex three dimensional geometry. A planar wave hits the surface of the F117 and is scattered in the volume of study. Dirichlet boundary conditions are applied on the surface, and approximated Robin boundary conditions are applied on the outer surface. The mesh is composed of $29 \cdot 10^6$ elements and $5.1 \cdot 10^6$ points. The solution for $k = 16$ is displayed on Fig. 4.9. A Krylov subspace of 100 was used for this example. The iteration count for constant interpolation, as in the SPD case, is reported on Table 4.10. The iteration count for various wave numbers and group numbers is reported on Table 4.11 for a stencil of 6 directions and on Table 4.12 for 26 directions. The iteration history is reported in Fig. 4.10 for the non deflated GMRES and for deflation with 10^4 groups with 6 directions. It is clearly seen that exponential interpolation provides better results for the same number of degrees of freedom for the coarse operator compared to the constant interpolation for the same coarse matrix size. Regarding a stencil of 26 directions, no substantial improvement appears. As noted before, it is less robust for low wave number and does not give a substantial acceleration for high wave number. For the three interpolations, the same pattern appears, namely a fast converge for low wave numbers and a sudden slow down up to a complete divergence as the wave number keeps on increasing.

4.4. Submarine

This example depicts a submarine and was chosen as representing a realistic example. Again, a plane wave impinges on the surface of the submarine where a Dirichlet boundary condition has been applied. An approximated Robin boundary condition has been

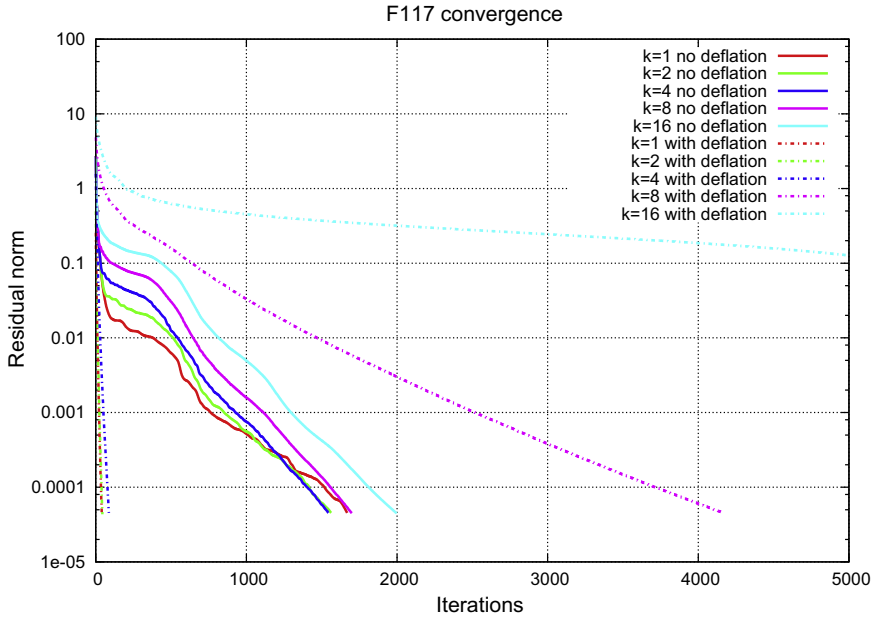

Fig. 4.10. Convergence history for the F117 flyer with and without deflation. The deflation reported here has been performed with 10^4 groups and 6 directions per group. The three steepest curves correspond to the deflated solutions for $k = \{1,2,4\}$. The five next curves correspond to the non deflated GMRES for $k = \{1,2,4,8,16\}$. Finally, the two less steep curves correspond to the deflated solution for $k = \{8,16\}$.

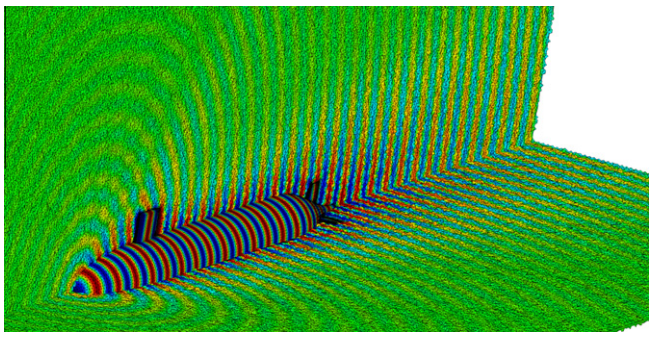

Fig. 4.11. Mesh and solution for the submarine example at $k = 8$.

Table 4.13
Iteration number for various wave numbers and group numbers with a stencil of 6 directions for the submarine example.

ngrou	0	10^4
$k = 1$	3842 (55848)	69 (2984)
$k = 2$	3063 (45478)	678(18742)
$k = 4$	3391 (46804)	10000
$k = 8$	4936 (114869)	10000

applied to the outer surface, simulating an infinite domain. The mesh contains $262 \cdot 10^6$ elements and $44 \cdot 10^6$ points, and the solution for $k = 8$ is depicted on Fig. 4.11. A Krylov subspace of 50 was used for this example. Results for 6 directions are displayed on Table 4.13, as the same behavior as the previous examples is observed. Again, a fast convergence for low wave numbers takes place, namely a CPU speed up of almost 20. Then a slow down in convergence appears for $k = 2$, and then divergence for higher k. Compared to previous examples, the increased size of this example stretches out the need for a larger coarse grid. However, the size of the coarse grid becomes too large to be solved efficiently by the direct solver, and clearly illustrates the limitation of the method for large examples and high wave numbers.

5. Conclusion

Complex subdomain deflation applied to the Helmholtz equation was presented on large and complex examples. The first major departure from the SPD case is characterized by the constraint that the coarse grid should be large enough to even guarantee convergence. As the wave number increases, the size of the coarse grid should increase accordingly to maintain fast convergence. The same convergence pattern was observed for all the examples presented here, namely a fast convergence for a small coarse grid and a low wave number until a sudden divergence if the coarse grid size is maintained constant while the wave number increases. The main drawback in the deflation procedure reveals itself in the coarse grid inversion, solved by a direct solver, becoming prohibitively expensive for high wave numbers on complex examples. Iterative inversion should be performed accurately, as underlined in [49], giving still worse performances than with a direct solver.

It was nevertheless shown that, when the coarse grid size was in the convergence range, exponential interpolation was consistently better than constant interpolation, as advocated in [33] in a multigrid context. An interesting procedure proposed in the former work, namely the separation procedure, may be performed easily on structured grids but has, to the authors' knowledge, no direct extensions on unstructured grids. In the deflation context, it would provide decoupling of the ray equations and smaller coarse matrices, but does not change the previous conclusions asymptotically.

Regarding the number of directions for the deflation space, it was shown that a large number of directions did not seem to be an effective choice. On one hand, a large number of directions produces a coarse matrix with very high condition number, if invertible, for low wave numbers. On the other hand, convergence did not seem to pay off the higher assembly cost above 6 directions.

The sensibility to the choice of deflation vectors, or prolongation, was also mentioned. Changing the value of the wave number in the prolongation could lead to improvements, although not substantials. However, it gives a hint on what can go wrong, namely the coarse grid solve or coarse grid correction. The main reason of the lack of convergence at high wave numbers is given in [66], where it is shown that the coarse grid is increasingly sensitive to approximate eigenvectors as the wave number increases.

The strong coupling between the discretization accuracy and the solver efficiency was highlighted. For indefinite problems, the coupling between the stability of the discrete operator and the convergence of the iterative solver has been highlighted for the Stokes problem in [67]. A very similar situation appears with Helmholtz.

There are different possible ways of improving convergence to higher wave numbers. One solution may consist in looking for more accurate eigenvectors of the discrete operator. Without relying on an explicit eigenvector computation, a smoothed approximation [39] could certainly improve convergence. Also, a larger coarse grid solved only approximately iteratively may provide enough acceleration to prove efficient on a CPU basis. Finally, the removal of the coarse grid operator for a more algebraic approach may also provide an answer. However, the deflation technique would only benefit from the first approach, as an approximated solve is not appropriate for deflation. Finally, the use of high order elements [68–70] should contribute to a better approximation of plane waves for higher wave number, and therefore a better convergence.

Ultimately, concerning low order approximations, one may wonder if it worth solving a discrete problem with a mesh too coarse to see fast convergence with deflation, while it is already known that accuracy will be low, as discrete eigenmodes deviate substantially from their analytical counterpart.

Acknowledgments

This work was partly supported by ONR and the DoD HPCMP PETT Program

References

[1] R. Aubry, F. Mut, R. Löhner, J. Cebral, Deflated preconditioned conjugate gradient solvers for the pressure-Poisson equation, J. Comput. Phys. 227 (24) (2008) 10196–10208.
[2] F. Mut, R. Aubry, R. Löhner, J. Cebral, Fast numerical solution of patient based blood flow in 3D arterial systems, Int. J. Numer. Methods Bio. Eng. 26 (1) (2010) 73–85.
[3] R. Löhner, F. Mut, J. Cebral, R. Aubry, G. Houzeaux, Deflated preconditioned conjugate gradient solvers for the pressure-Poisson equation: extensions and improvements, Int. J. Numer. Methods Eng. 87 (1) (2011) 2–14.
[4] R. Aubry, F. Mut, S. Dey, R. Löhner, Deflated preconditioned conjugate gradient solvers for linear elasticity, Int. J. Numer. Methods Eng. 88 (2011) 1112–1127.
[5] Y. Saad, Iterative Methods for Sparse Linear Systems, second ed., SIAM, 2003.
[6] F. Ihlenburg, I. Babuska, Finite element solution of the Helmholtz equation with high wave number part I: the h-version of the FEM, Comput. Math. Appl. 30 (9) (1995) 9–37.
[7] F. Ihlenburg, I. Babuska, Finite element solution of the Helmholtz equation with high wave number part II: the h-p version of the FEM, SIAM J. Numer. Anal. 34 (1) (1997) 315–358.
[8] I. Babuska, S. 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.
[9] A. Bayliss, C. Goldstein, E. Turkel, On accuracy conditions for the numerical computation of waves, J. Comput. Phys. 59 (1985) 396–404.

[10] I. Babuska, F. Ihlenburg, E. Paik, S. Sauter, A generalized finite element method for solving the Helmholtz equation in two dimensions with minimal pollution, Comput. Methods Appl. Mech. Eng. 128 (3) (1995) 325–359.
[11] J. Melenk, I. Babuska, The partition of unity finite element method: basics theory and applications, Comput. Methods Appl. Mech. Eng. 139 (1996) 289–314.
[12] E. Perrey-Debain, O. Lagrouche, P. Bettess, J. Trevelyan, Plane wave basis finite-elements for wave scattering in three dimensions, Commun. Numer. Methods Eng 19 (2003) 715–723.
[13] E. Perrey-Debain, O. Lagrouche, P. Bettess, J. Trevelyan, Plane-wave basis finite elements and boundary elements for three-dimensional wave scattering, Philos. Trans. R. Soc. Lond. 362 (1816) (2004) 561–577.
[14] O. Cessenat, B. Despres, Application of an ultra weak variational formulation of elliptic PDEs to the two-dimensional Helmholtz problem, SIAM J. Numer. Anal. 35 (1) (1998) 255–299.
[15] P. Monk, D. Wang, A least squares method for the Helmholtz equation, Comput. Methods Appl. Mech. Eng. 175 (1998) 121–136.
[16] I. Harari, T. Hughes, Finite element methods for the Helmholtz equation in an exterior domain: model problems, Comput. Methods Appl. Mech. Eng. 87 (1991) 59–96.
[17] L. Franca, C. Farhat, A. Macedo, M. Lesoinne, Residual-free bubbles for the Helmholtz equation, Int. J. Numer. Methods Eng. 40 (1997) 4003–4009.
[18] C. Farhat, I. Harari, U. Hetmaniuk, A discontinuous Galerkin method with Lagrange multipliers for the solution of Helmholtz problems in the mid-frequency regime, Comput. Methods Appl. Mech. Eng. 192 (2003) 1389–1419.
[19] G. Bao, G. Wei, S. Zhao, Numerical solution of the Helmholtz equation with high wave numbers, Int. J. Numer. Methods Eng. 59 (2004) 389–408.
[20] A.A. Oberai, P.M. Pinsky, A residual-based finite element method for the Helmholtz equation, Int. J. Numer. Methods Eng. 49 (3) (2000) 399–419.
[21] Y.A. Erlangga, Advances in iterative methods and preconditioners for the Helmholtz equation, Arch. Comput. Methods Eng. 15 (1) (2008) 37–66.
[22] M. Magolu Monga Made, R. Beauwens, G. Warze, Preconditioning of discrete Helmholtz operators perturbed by a diagonal complex matrix, Commun. Numer. Methods Eng. 16 (11) (2000) 801–817.
[23] M. Magolu Monga Made, Incomplete factorization-based preconditionings for solving the Helmholtz equation, Int. J. Numer. Methods Eng. 50 (2001) 1077–1101.
[24] A. Bayliss, C. Goldstein, E. Turkel, An iterative method for the Helmholtz equation, J. Comput. Phys. 49 (1983) 443–457.
[25] A. Laird, M. Giles, Preconditioning harmonic unsteady potential flow calculations, AIAA J. 44 (11) (2006) 2654–2662.
[26] Y. Erlangga, C. Oosterlee, C. Vuik, A novel multigrid based preconditioner for heterogeneous Helmholtz problems, SIAM J. Sci. Comput. 27 (2006) 1471–1492.
[27] Y. Erlangga, C. Vuik, C. Oosterlee, On a class of preconditioners for solving the Helmholtz equation, Appl. Numer. Math. 50 (2004) 409–425.
[28] T. Airaksinen, E. Heikkola, A. Pennanen, J. Toivanen, An algebraic multigrid based shifted-Laplacian preconditioner for the Helmholtz equation, J. Comput. Phys. 226 (2007) 1196–1210.
[29] M. Bollhöfer, M.J. Grote, O. Schenk, Algebraic multilevel preconditioner for the Helmholtz equation in heterogeneous media, SIAM J. Sci. Comput. 31 (5) (2009) 3781–3805.
[30] D. Osei-Kuffuor, Y. Saad, Preconditioning Helmholtz linear systems, Appl. Numer. Math. 60 (4) (2010) 420–431.
[31] J. Sifuentes, Preconditioning the integral formulation of the Helmholtz equation via deflation, Master's thesis, Rice University, Houston, Texas, 2006.
[32] L. Trefethen, M. Embree, Spectra and Pseudospectra: The Behavior of Non-normal Matrices and Operators, Princeton University Press, 2005.
[33] A. Brandt, I. Livshits, Wave-ray multigrid method for standing wave equations, Trans. Numer. Anal 6 (1997) 162–181.
[34] I. Livshits, A. Brandt, Accuracy properties of the wave-ray multigrid algorithm for Helmholtz equations, SIAM J. Sci. Comput. 28 (4) (2006) 1228–1251.
[35] I. Livshits, An algebraic multigrid wave-ray algorithm to solve eigenvalue problems for the Helmholtz operator, Numer. Linear Algebra Appl. 11 (2) (2004) 229–239.
[36] B. Lee, T.A. Manteuffel, S.F. McCormick, J. Ruge, First-order system least-squares for the Helmholtz equation, SIAM J. Sci. Comput. 21 (5) (2000) 1927–1949.
[37] S. Kim, S. Kim, Multigrid simulation for high-frequency solutions of the Helmholtz problem in heterogeneous media, SIAM J. Sci. Comput. 24 (2) (2002) 684–701.
[38] 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.
[39] P. Vanek, J. Mandel, M. Brezina, Two-level algebraic multigrid for the Helmholtz problem, in: Tenth International Conference on Domain Decomposition, Contemporary Mathematics, vol. 218, American Mathematical Society, 1998, pp. 349–356.
[40] C. Farhat, A. Macedo, M. Lesoinne, A two-level domain decomposition method for the iterative solution of high frequency exterior Helmholtz problem, Numer. Math. 283 (85) (2000) 4003–4009.
[41] C. Farhat, R. Tezaur, J. Toivanen, A domain decomposition method for discontinuous Galerkin discretizations of Helmholtz problems with plane waves and Lagrange multipliers, Int. J. Numer. Methods Eng. 78 (2009) 1513–1531.
[42] J. Wilkinson, The Algebraic Eigenvalue Problem, Oxford University Press, Inc., New York, NY, USA, 1988.
[43] B. Parlett, The Symmetric Eigenvalue Problem, Classics in Applied Mathematics, SIAM, Philadelphia, 1998.
[44] R. Nicolaides, Deflation of conjugate gradients with applications to boundary value problems, SIAM J. Numer. Anal. 24 (2) (1987) 355–365.
[45] M. Hestenes, E. Stiefel, Methods of conjugate gradients for solving linear systems, J. Res. Nat. Bur. Stand. 49 (1952) 409–436.
[46] R. Morgan, A restarted GMRES method augmented with eigenvectors, SIAM J. Matrix Anal. Appl. 16 (4) (1995) 1154–1171.
[47] Y. Erlangga, R. Nabben, Deflation and balancing preconditioners for Krylov subspace methods applied to nonsymmetric matrices, SIAM J. Matrix Anal. Appl. 30 (2) (2008) 684–699.
[48] Y. Erlangga, R. Nabben, Multilevel projection-based nested Krylov iteration for boundary value problems, SIAM J. Sci. Comput. 30 (3) (2008) 1572–1595.
[49] J. Tang, R. Nabben, C. Vuik, Y. Erlangga, Theoretical and numerical comparison of various projection methods derived from deflation, domain decomposition and multigrid methods, Report 07-04, Delft University of Technology, Delft Institute of Applied Mathematics, Delft, 2007.
[50] Y. Shapira, Matrix-Based Multigrid: Theory and Applications, Springer, 2008.
[51] A. Greenbaum, V. Pták, Z. Strakos, Any nonincreasing convergence curve is possible for GMRES, SIAM J. Matrix Anal. Appl. 17 (3) (1996) 465–469.
[52] A. van der Sluis, H. van der Vorst, The rate of convergence of conjugate gradients, Numer. Math. 48 (1986) 543–560.
[53] M. Eiermann, O.G. Ernst, O. Schneider, Analysis of acceleration strategies for restarted minimal residual methods, J. Comput. Appl. Math. 123 (1–2) (2000) 261–292.
[54] P. Vanek, M. Brezina, J. Mandel, Convergence of algebraic multigrid based on smoothed aggregation, Computing 56 (1996) 179–196.
[55] J. Ruge, K. Stüben, Algebraic multigrid, in: S. McCormick (Ed.), Multigrid Methods, SIAM, Philadelphia, PA, 1987, pp. 73–130.
[56] M. Brezina, Robust iterative methods for unstructured meshes, Ph.D. thesis, University of Colorado, 1997.
[57] V. Gravemeier, M. Gee, W. Wall, An algebraic variational multiscale-multigrid method based on plain aggregation for convection–diffusion problems, Comput. Methods Appl. Mech. Eng. 198 (2009) 3821–3835.
[58] Y. Saad, M. Yeung, J. Erhel, F. Guyomarc'h, A deflated version of the conjugate gradient algorithm, SIAM J. Sci. Comput. 21 (5) (2000) 1909–1926.
[59] J. Frank, C. Vuik, On the construction of deflation-based preconditioners, SIAM J. Sci. Comput. 23 (2001) 442–462.
[60] F. Vermolen, C. Vuik, A. Segal, Deflation in preconditioned conjugate gradient methods for finite element problems, in: M. Křižek, P. Neittaanmätki, R. Glowinski, S. Korotov (Eds.), Conjugate Gradient and Finite Element Methods, Springer, Berlin, 2004, pp. 103–129.
[61] Y. Saad, M.H. Schultz, GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems, SIAM J. Sci. Comput. 7 (1986) 856–869.
[62] H. van der Vorst, Iterative Krylov Methods for Large Linear Systems, Cambridge University Press, Cambridge, 2003.
[63] F. Ihlenburg, Finite Element Analysis of Acoustic Scattering, Applied Mathematical Sciences, vol. 132, Springer, New York, 1998.
[64] U. Trottenberg, C.W. Oosterlee, A. Schuller, Multigrid, Elsevier, 2000.
[65] MATLAB, version 7.10.0 (R2010a), The MathWorks Inc., Natick, Massachusetts, 2010.
[66] J. Fish, Y. Qu, Global-basis two-level method for indefinite systems. Part 1: convergence studies, Int. J. Numer. Methods Eng. 49 (3) (2000) 439–460.
[67] S. Norburn, D. Silvester, Stabilised vs stable mixed methods for incompressible flow, Comput. Methods Appl. Mech. Eng. 166 (1997) 131–141.
[68] S. Dey, D.K. Datta, A parallel hp-FEM infrastructure for three-dimensional structural acoustics, Int. J. Numer. Methods Eng. 68 (2006) 583–603.
[69] S. Dey, Evaluation of p-FEM approximations for mid-frequency elasto-acoustics, J. Comput. Acoust. 11 (2) (2003) 195–225.
[70] S. Dey, J.J. Shirron, L.S. Couchman, Mid-frequency structural acoustic and vibration analysis in arbitrary, curved three-dimensional domainsevaluation of p-FEM approximations for mid-frequency elasto-acoustics, Comput. Struct. 79 (6) (2001) 617–629.