Applied Numerical Mathematics 60 (2010) 1157–1170

Contents lists available at ScienceDirect

Applied Numerical Mathematics

www.elsevier.com/locate/apnum

Application of the shifted-Laplace preconditioner for iterative solution of a higher order finite element discretisation of the vector wave equation: First experiences

P.B. Hooghiemstra [a,*,1], D.R. van der Heul [b,2], C. Vuik [b]

[a] Netherlands Institute for Space Research (SRON), Institute for Marine and Atmospheric research Utrecht (IMAU), Princetonplein 5, 3584CC Utrecht, The Netherlands
[b] Delft University of Technology, Faculty of Electrical Engineering, Mathematics and Computer Science, Department of Applied Mathematics, Mekelweg 4, 2628 CD, Delft, The Netherlands

ARTICLE INFO

Article history:
Received 30 June 2008
Received in revised form 15 October 2009
Accepted 7 July 2010
Available online 16 July 2010

Keywords:
Finite elements
Shifted-Laplace preconditioner
Maxwell equations
Electromagnetic scattering

ABSTRACT

The analysis of electromagnetic scattering by electrically large one-side open–ended cavities remains a challenge. Finite element discretisation of the vector wave equation to solve for the electric field inside the cavity leads to ill-conditioned indefinite linear systems of large dimension, a result of the requirement of a fine nearly uniform discrete sampling in the computational domain. Direct methods based on frontal solution techniques to solve the resulting linear system have been used because efficient iterative methods were not available. With the arrival of the shifted-Laplace preconditioner of Erlangga, iterative solution of the indefinite system becomes tractable. This paper discusses the modifications required for application of the shifted-Laplace preconditioner to cavity scattering and some preliminary results of this approach. The shifted-Laplace preconditioner is shown to be very effective for improving the convergence rate of the iterative solution algorithm. However, to be able to handle problems with larger number of degrees of freedom, it is necessary to include a multigrid algorithm to solve the preconditioner system, as this will allow the use of short recurrence Krylov subspace methods, as opposed to the currently employed long-recurrence method, for which the storage requirements of the Krylov basis become unpractically large.

1. Introduction

The radar signature of a military platform remains one of the key parameters that determines its deployment envelope. For jet powered fighter aircraft at near nose on observation angles the contribution to the radar signature of the jet engine is generally dominant. The jet engine air intake duct, together with the compressor fan form a large and deep one-side open cavity that scatters the incident electromagnetic waves over a large cone of observation angles. Therefore, the modelling and analysis of the electromagnetic scattering properties of large open–ended cavities remains a challenging task of continuing interest to the military aerospace engineering community.

When the characteristic dimensions of a scattering body are many times larger than the wavelength of the incident electromagnetic waves, such bodies are referred to as *electrically large*, as is the case for a fullsize jet engine air intake duct.

* Corresponding author.
E-mail address: p.b.hooghiemstra@uu.nl (P.B. Hooghiemstra).

[1] The work was performed as an internship at the National Aerospace Laboratory NLR, Amsterdam, The Netherlands to become MSc.

[2] Present address: Delft University of Technology, Faculty of Electrical Engineering, Mathematics and Computer Science, Department of Applied Mathematics, Mekelweg 4 2628 CD, Delft, The Netherlands.

0168-9274/$30.00 © 2010 IMACS. Published by Elsevier B.V. All rights reserved.
doi:10.1016/j.apnum.2010.07.004

For convex scattering bodies, the interaction between the electric current distribution induced on different regions of the scattering body, weakens and eventually nearly vanishes with increasing electrical size. This phenomenon is the basis for the use of so-called *high-frequency asymptotic analysis method*s that, to a certain extend, discard this interaction, and will give an accurate description of the first order scattering phenomena for electrically large convex scattering bodies. However, the concave shape of an electrically large cavity retains a strong interaction between the electric current distribution induced on different parts of the cavity interior, which means that beforementioned asymptotic methods cannot be used to perform an accurate analysis. Although a number of dedicated analysis techniques are available, these are not universally applicable because of the underlying assumptions and/or limitations on the geometry of the cavity (consult the review paper of Anastassiu [1] for more details). Only so-called *full-wave* methods, that are based on direct discretisation of the Maxwell equations can by used to analyse large, noncylindrical curved cavities. In the current application the response to a monochromatic excitation is analysed, using a formulation in the frequency domain. Frequency domain full-wave methods fall in two distinct categories:

- Methods based on a finite element discretisation of the Combined Field Integral Equation (CFIE), essentially an integral form of the Maxwell equations.
- Methods based on a finite element discretisation of the vector wave equation, essentially a differential form of the Maxwell equations.

With the introduction of Multi Level Fast Multipole Algorithm (MLFMA) based acceleration techniques (see e.g. Ref. [2]), the former approach has established a position as a standard workhorse for electromagnetic analysis. However, because of the strong interaction inside the concave domain, application of MLFMA for the analysis of a jet engine air intake is not completely straightforward. Additionally, it is often desirable to be able to analyse the contribution of the isolated engine air intake and combine it with a separate (possibly asymptotic) analysis of the other components of the aircraft that comprise it's outer mould shape. In this case only the contribution of the *inside* of the air intake is required. Although integral equation formulations for the contribution of the inside of the jet engine air intake do exist, this formulation is not present in a standard MLFMA solver. In the second approach based on discretisation of the vector wave equation inside the jet engine air intake cavity the contribution of the exterior is excluded automatically in the formulation.

In modern fighters extensive use is made of Radar Absorbing Coatings (RAC) and Radar Absorbing Materials (RAM) to absorb and dissipate the incident electromagnetic energy and thereby reduce the aircraft radar signature. Because of the multiple reflections that occur inside the cavity, this application of RAC can be very effective. Whereas within the integral equation formulation the RAC have to be approximated by imposing an impedance boundary condition on the surface, the differential formulation also allows discretisation of the volume occupied by the RAC, imposing the appropriate material interface conditions through properly chosen basis functions for the electromagnetic fields.

1.1. A jet engine air intake scattering model based on a finite element discretisation of the vector wave equation

For these reasons the National Aerospace Laboratory NLR has started development of a jet engine air intake scattering model based on a finite element discretisation of the vector wave equation. The initial approach was based on the work of Jin et al. [10,9,11,14]. They combined a higher order finite element discretisation with an efficient (direct) frontal solver that takes full advantage of the special geometry of the air intake, where the typical diameter-to-depth ratio is 6–10. The latter fact combined with an optimised ordering of the elements limits the memory requirements of the solution algorithm to just more than is required to store a complex valued matrix of dimension equal to the number of degrees of freedom within a cross section of the cavity. An additional reason to choose a direct method is that the system typically has to be solved for a couple of hundred right-hand sides, corresponding to the many different directions of excitation and two different field polarisations. However, it was found that the computational work involved in the direct solver prohibits it's practical application for problems with very large number of degrees of freedom ($> 10^9$). Initially iterative solution methods had been considered but our experiences with Krylov subspace methods accelerated with 'classic' preconditioners, e.g. ILU, showed that either the preconditioner was not very effective or required excessive memory storage. With the introduction of the *shifted-Laplace* preconditioner by Erlangga [6] for the discretised Helmholtz equation, the question arose if this approach is equally effective for iterative solution of the discretised vector wave equation, as both equations are obviously very similar. In this paper our first experiences with this approach are discussed. Our aim was to make an initial evaluation of the algorithm, investing a limited effort. Now the value of the algorithm for this application has been established it is currently further developed by including algebraic multigrid to solve the preconditioner system and adapting the continuous formulation to improve the convergence rate of the solution algorithm.

1.2. Outline

Both the continuous formulation and the finite element discretisation of the cavity scattering problem are briefly discussed. Next the solution procedure of the linear system using the original direct solver is reviewed. This is followed by a description of how the shifted-Laplace preconditioner is integrated in a block-preconditioner, to efficiently handle the fully populated submatrix that results from discretisation of the imposed global radiation boundary condition on the cavity

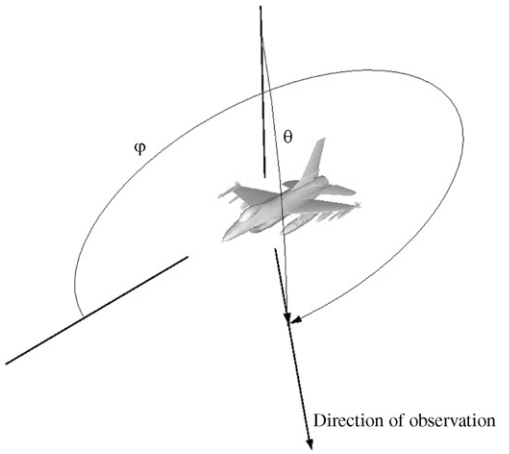

Fig. 1. Definition of the observation angles θ and ϕ.

aperture. Some numerical experiments have been performed to explore the sensitivity to the parameters that define the preconditioners. Finally, a number of recommendations are made to further improve the algorithm.

2. Continuous formulation

Because the excitation by the electromagnetic waves is assumed to be time-harmonic and monochromatic, the problem can be formulated in the frequency domain, by considering the *phasor* $\mathbf{E}$ or complex valued amplitude of the electric field $\mathcal{E}$ with rotational frequency ω, defined in the following way:

$$\mathcal{E}(\mathbf{x},t) = \Re\big(\mathbf{E}(\mathbf{x})e^{i\omega t}\big). \tag{1}$$

The quantity $\mathbf{E}$ will from hereon be referred to as the electric field, because the actual electric field $\mathcal{E}$ will not be introduced in the problem formulation. The contribution to the radar cross section of the deep one-side open–ended cavity formed by the jet engine air intake and the compressor is quantified by the *radar cross section* σ, that is dependent on the two observation angles (ϕ,θ) (defined in Fig. 1) and the polarisation of the incident and scattered time-harmonic electric field:

$$\sigma_{\hat{\mathbf{E}}_{\text{in}}\hat{\mathbf{E}}_{\text{sc}}}(\phi,\theta) = \lim_{r\to\infty} 4\pi r^2 \frac{|\mathbf{E}_{\text{sc}}(\mathbf{x}_{\text{obs}})|}{|\mathbf{E}_{\text{in}}|}, \quad r = |\mathbf{x}_{\text{obs}}|, \tag{2}$$

where $\mathbf{x}_{\text{obs}}$ is the point of observation and $\mathbf{E}_{\text{in}}$ and $\mathbf{E}_{\text{sc}}$ are the incident electric field on the cavity and the electric field scattered by the cavity, respectively. In turn the scattered electric field observed at a location sufficiently far removed from the cavity can be accurately approximated by:

$$\mathbf{E}_{\text{sc}}(\mathbf{x}) = \frac{ik_0 e^{ik_0|\mathbf{x}|}}{2\pi|\mathbf{x}|} \iint_{S_{\text{ap}}} \mathbf{M}(\mathbf{x}') e^{-k_0(\mathbf{x}\cdot\mathbf{x}')/|\mathbf{x}|}\, dS', \tag{3}$$

where $k_0 = \omega\sqrt{\mu_0\epsilon_0}$, with μ_0, ϵ_0 the permeability and permittivity in vacuum, respectively and the magnetic current $\mathbf{M}(\mathbf{x})$ defined as:

$$\mathbf{M}(\mathbf{x}) \equiv -\hat{\mathbf{n}} \times \mathbf{E}(\mathbf{x}), \tag{4}$$

where $\hat{\mathbf{n}}$ is the unit normal vector directed outward on the aperture surface S_{ap}.

Inside the volume V_{cv} defined by the exterior surface $S_{\text{cv}} + S_{\text{ap}}$ (Fig. 2) the electric field phasor $\mathbf{E}(\mathbf{x})$ obeys the homogeneous vector wave equation

$$\nabla \times \nabla \times \mathbf{E}(\mathbf{x}) - k_0^2 \epsilon_r \mathbf{E}(\mathbf{x}) = 0, \tag{5}$$

where ϵ_r is the relative permittivity of the material (vacuum or RAM). To define a well-posed problem, it suffices to prescribe either the tangential electric field or the tangential magnetic field on the boundary surface of the domain. The former is prescribed on the cavity surface S_{cv} and the latter on the aperture surface S_{ap} by imposing a homogeneous Dirichlet boundary condition and a global radiation boundary condition, respectively:

$$\mathbf{E}_t(\mathbf{x}) = 0, \quad \mathbf{x} \in S_{\text{cv}}, \tag{6}$$

$$\mathbf{H}_t(\mathbf{x}) = i\frac{Z_0}{k_0}\left[k_0^2 \iint_{S_{\text{ap}}} \mathbf{M}(\mathbf{x}') g(\mathbf{x},\mathbf{x}')\, dS' + \nabla \iint_{S_{\text{ap}}} \nabla' \cdot \mathbf{M}(\mathbf{x}') g(\mathbf{x},\mathbf{x}')\, dS'\right], \quad \mathbf{x} \in S_{\text{ap}}, \quad g(\mathbf{x},\mathbf{x}') = \frac{e^{ik_0|\mathbf{x}-\mathbf{x}'|}}{2\pi|\mathbf{x}-\mathbf{x}'|}. \tag{7}$$

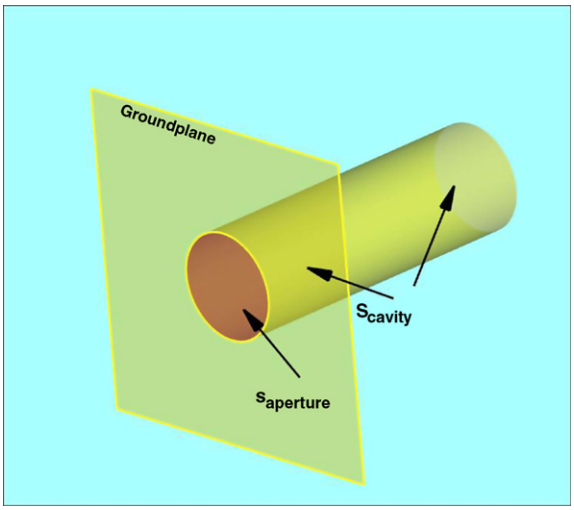

Fig. 2. Schematic layout of the cavity scattering problem.

In formulating the global radiation boundary condition it is assumed the cavity aperture is located in a perfectly conducting groundplane of infinite extend (Fig. 2), which eliminates the electric current distribution on the aperture surface. Using (5), (6) and (7), the following functional can be defined:

$$F = \frac{1}{2}\iiint_{V_c}\left[\big(\nabla\times\mathbf{E}(\mathbf{x})\big)\big(\nabla\times\mathbf{E}(\mathbf{x})\big) - k_0^2\epsilon_r\mathbf{E}(\mathbf{x})\cdot\mathbf{E}(\mathbf{x})\right]dV - k_0^2\iint_S \mathbf{M}(\mathbf{x})\cdot\left[\iint_{S_a}\mathbf{M}(\mathbf{x})g\big(\mathbf{r},\mathbf{r}'\big)\,dS'\right]dS$$
$$+\iint_{S_{\text{aperture}}}\nabla\cdot\mathbf{M}(\boldsymbol{x})\left[\iint_{S_{\text{aperture}}} g\big(\boldsymbol{x},\boldsymbol{x}'\big)\nabla'\cdot\mathbf{M}\big(\boldsymbol{x}'\big)\,dS'\right]dS - 2ik_0Z_0\iint_{S_{\text{aperture}}}\mathbf{M}(\boldsymbol{x})\cdot\mathbf{H}^i(\boldsymbol{x})\,dS. \tag{8}$$

The electric field inside the cavity $\mathbf{E}(\mathbf{r})$ can now be found by locating the stationary point of the functional (8).

3. Finite element discretisation

The weak formulation of the vector wave equation is discretised using higher order edge-based test functions on tetrahedral elements, following the approach of Graglia et al. [7], who combined standard Nedelec basis functions with polynomial interpolation functions. Tetrahedral elements are chosen because of the straightforward definition of curl conforming basis functions and the possibility of accurately boundary conforming automatic grid generation. The higher order formulation leads to a significant reduction of the number of degrees of freedom in the problem, while the edge-based formulation will guarantee the correct field behavior across material interfaces with discontinuous permittivity (continuous tangential and discontinuous normal component of the field) and suppresses the occurrence of spurious solutions generally encountered when using node-based test functions. Because the computational domain in this application does not have highly curved boundaries, rectilinear elements are used as opposed to the curvilinear elements of Graglia et al. [7]. This allows a more efficient evaluation of the element matrices. The zeroth-order curl conforming test functions are given by:

$$\boldsymbol{\Omega}_{\gamma\beta}(\mathbf{r}) = \xi_n\nabla\xi_m - \xi_m\nabla\xi_n, \tag{9}$$

where $\{\gamma,\beta,m,n\}$ are the six even permutations of $\{1,2,3,4\}$ such that $\gamma<\beta$. $\boldsymbol{\Omega}_{\gamma\beta}(\mathbf{r})$ is associated to the edge that connects face γ to face β of the tetrahedron. The curl conforming higher order basis functions of order p can be expressed as:

$$\boldsymbol{\Omega}^{\gamma\beta}_{ijkl} = N^{\gamma\beta}_{ijkl}\frac{(p+2)^2\xi_\gamma\xi_\beta\hat{\alpha}_{ijkl}(\boldsymbol{\xi})}{i_\gamma i_\beta}\boldsymbol{\Omega}_{\gamma\beta}(\mathbf{r}), \tag{10}$$

where $i_{\gamma,\beta}$ is taken to be i,j,k,l for $\gamma,\beta=1,2,3,4$ respectively. The Silvester–Lagrange interpolation polynomial $\hat{\alpha}_{ijkl}(\boldsymbol{\xi})$ is defined as:

$$\hat{\alpha}_{ijkl}(\boldsymbol{\xi}) = \hat{R}_i(p+2,\xi_1)\hat{R}_j(p+2,\xi_2)\hat{R}_k(p+2,\xi_3)\hat{R}_l(p+2,\xi_4), \tag{11}$$

where the shifted Sylvester polynomials $\hat{R}_i$ of order p are defined as

$$\hat{R}_i(p,\xi) = \begin{cases} \frac{1}{(i-1)!}\prod_{k=1}^{i-1}(p\xi - k), & 2 \leqslant i \leqslant p+1, \\ 1, & i = 1, \\ \lim_{i\to 0}\frac{\hat{R}_i(p,\xi)}{i} \equiv \frac{1}{p\xi}. \end{cases} \tag{12}$$

The normalization factor $N_{ijkl}^{\gamma\beta}$ is defined as:

$$N_{ijkl}^{\gamma\beta} = \frac{p+2}{p+2-i_\gamma - i_\beta}|\boldsymbol{l}_{\gamma\beta}|, \tag{13}$$

where

$$\boldsymbol{l}_{\gamma\beta} = \mathcal{J}(\nabla\xi_j \times \nabla\xi_k), \quad \mathcal{J} = \boldsymbol{l}^i \cdot \boldsymbol{l}^j \times \boldsymbol{l}^k, \quad \boldsymbol{l}^i = \frac{\partial \boldsymbol{r}}{\partial \xi_i}. \tag{14}$$

Following the approach of Graglia et al. [7] test functions of arbitrary order can be defined. As a fair balance between complexity and efficiency, second order test functions (cubic polynomials) have been chosen. Because of the uniform discretisation of the computational domain, even the use of higher order basis functions still leads to problems with very large number of degrees of freedom: $> 10^7$.

The electric field $\boldsymbol{E}(\boldsymbol{r})$ and the magnetic current are expanded in a series of these test functions:

$$\begin{aligned} \mathbf{E}(\mathbf{x}) &= \sum_{i=1}^{\text{Ndof int}} E_i \boldsymbol{\Omega}_{\mathbf{i}}(\mathbf{x}) + \sum_{i=\text{Ndof int}+1}^{\text{Ndof int}+\text{Ndof ap}} E_i \boldsymbol{\Omega}_{\mathbf{i}}(\mathbf{x}), \\ \mathbf{M}(\mathbf{x}) &= \sum_{i=1}^{\text{Ndof ap}} E_i \boldsymbol{\Lambda}_{\mathbf{i}}(\mathbf{x}) | \mathbf{x} \in S_{\text{ap}}, \\ \boldsymbol{\Lambda}_{\mathbf{i}}(\mathbf{x}) &= -\hat{\mathbf{n}} \times \boldsymbol{\Omega}_{\mathbf{i}+\text{Ndof int}}(\mathbf{x}), \quad i = 1, \ldots, \text{Ndof ap}, \end{aligned} \tag{15}$$

where *Ndof int* and *Ndof ap* are the number of internal degrees of freedom and the number of degrees of freedom located on the aperture, respectively. Substitution of the expansion (15) in the functional (8) and application of the Rayleigh–Ritz procedure leads to a system of linear equations for the unknown complex coefficients E_i in the expansion. The dimension of the linear system will be *Ndof*, the total number of degrees of freedom.

4. Properties of the linear system

The linear system that results from discretisation of the weak form of the vector wave equation for our frequency band of interest, combined with the beforementioned global radiation boundary condition on the aperture has the following properties, that should be considered when choosing a solution method:

- The system is complex valued through the contribution of the radiation boundary condition. Furthermore, when a radar absorbing liner is fitted in the air intake, the complex value of the permittivity of the liner material will result in additional complex valued coefficients.
- The system is nearly symmetric (non-Hermitian). Although the Galerkin FE discretisation of the vector wave equation results in a symmetric system, this does not hold for the discretisation of the radiation boundary condition. More specifically, the 'inner' and 'outer' integrals of the operator are handled differently for the (nearly) singular case of two triangular elements that (nearly) coincide. This is common practice in the discretisation of boundary integral equations. Although $\max(|A_{ij} - A_{ji}|/|A_{ij}|) \approx 10^{-4}$–$10^{-5}$ this is sufficient to stall convergence of linear solvers that assume symmetry.
- For the frequency band of interest the high wave number k_0 will make the matrix highly indefinite.
- Within the system matrix three unique submatrices can be identified, where for future reference the explicit dependence of the (sub)matrices on the wavenumber k_0 is introduced:

$$\mathbf{A}(k_0) = \begin{pmatrix} \mathbf{A}_{11}(k_0) & \mathbf{A}_{12}(k_0) \\ (\mathbf{A}_{12}(k_0))^T & \mathbf{A}_{22}(k_0) \end{pmatrix}, \tag{16}$$

where
· $\mathbf{A}_{11}(k_0) \in \mathbb{C}^{\text{Ndof int}\times\text{Ndof int}}$, sparsely populated,
· $\mathbf{A}_{12}(k_0) \in \mathbb{C}^{\text{Ndof int}\times\text{Ndof ap}}$, sparsely populated,
· $\mathbf{A}_{22}(k_0) \in \mathbb{C}^{\text{Ndof ap}\times\text{Ndof ap}}$, fully populated.
The discretisation of the global radiation boundary condition on the aperture results in the matrix being partially sparsely and partially fully populated, as indicated in Fig. 3.

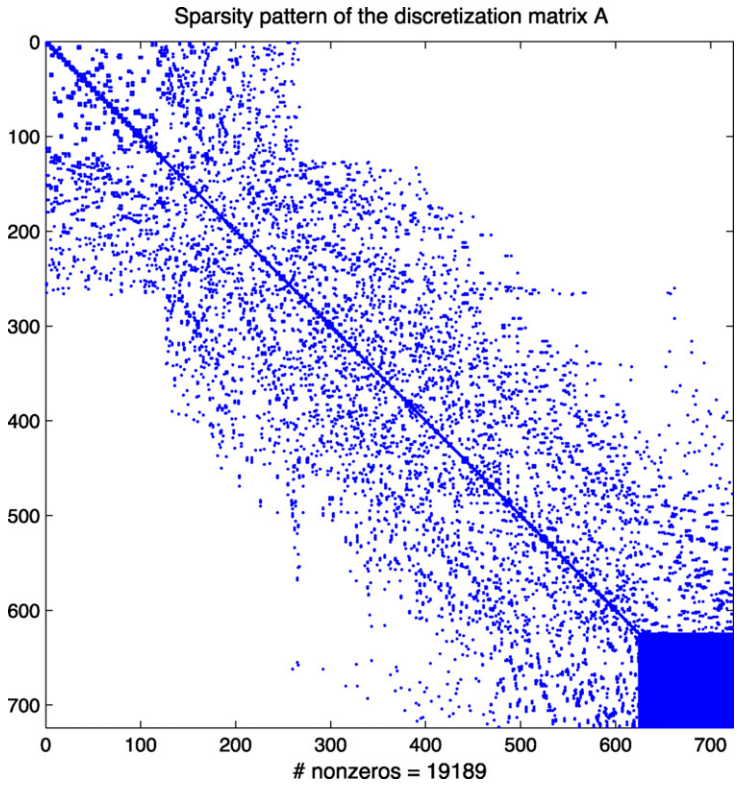

Fig. 3. Typical sparsity pattern for the system of equations derived from discretisation of the vector wave equation with a global radiation boundary condition imposed on the aperture. The length-diameter ratio of the cavity $L/d \approx 0.4$, with 723 degrees of freedom.

5. Direct solution of linear system

Jin et al. [10,9,11,14] use a direct solution method to solve the system of linear equations that result from the finite element discretisation of the vector wave equation. There are multiple reasons for choosing a direct solution method for this application:

- The geometry of the problem allows formulation of a relatively efficient frontal solution method with limited memory requirements.
- The system has to be solved repeatedly, for many different right-hand sides, corresponding to different angles of excitation.
- 'Standard' preconditioners as e.g. ILU() or AI, are either not effective in improving the convergence rate of Krylov subspace methods for this application or require nearly as much work as is involved in factorisation of the matrix for direct solution of the system.
- 'Standard' multigrid algorithms are not easily set up to handle high frequency wave equations.

Because of the specific geometry of the computational domain, characterized by one dominant principal dimension, the degrees of freedom can be ordered (starting from the bottom of the cavity to the aperture) such that a relatively small bandwith is achieved when the system matrix would be assembled. Clearly, the theoretically possible minimal bandwith is equal to the number of degrees of freedom in a typical cross section of the cavity. Instead of first assembling the system matrix, a so-called frontal solution algorithm can be designed that simultaneously assembles the finite element matrices into the *frontal matrix* and eliminates only those degrees of freedom for which *all* element matrices they belong to are fully assembled. This simultaneous approach limits the storage requirements of the algorithm to roughly only $(Ndof\,ap)^2$, as opposed to $Ndof \cdot Ndof\,ap$ when the system would be assembled. Using a subdivided prismatic grid, an ordering of the elements can be achieved that is indeed close to the optimal value, see Van der Heul et al. [17]. Jin et al. [10] used the classical method of Irons [8] while in the current algorithm use was originally made of a state-of-the-art frontal solver MA42, developed by Duff and Scott [3].

However, in using the direct solution algorithm it became clear that computational work becomes unacceptably high for problems with a very large number of degrees of freedom (10^6–10^7). Therefore, an alternative solution algorithm was required.

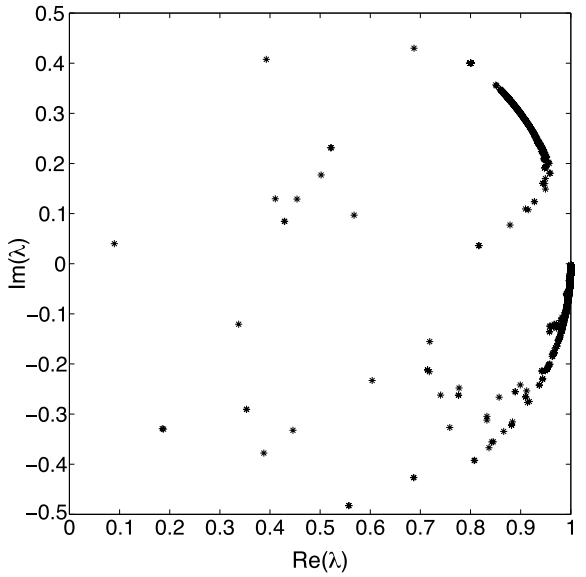

Fig. 4. Spectrum of the preconditioned system for the discretisation of the vector wave equation on a small cavity geometry, using the shifted-Laplace preconditioner of Erlangga and Erlangga's optimal shift parameters for the discretised Helmholtz equation. The eigenvalues are contained within a circle, and removed from the origin.

6. Iterative solution of the linear system

As mentioned in the introductory section our aim is to evaluate the use of the shifted-Laplace preconditioner of Erlangga, that has established a very good reputation for improving the convergence rate for iterative solution of the discretised Helmholtz equation, to solve the system resulting from discretisation of the vector wave equation. The vector wave equation (5) can be explicitly expressed as a vector form of the Helmholtz equation by using the vector identity:

$$\nabla \times \nabla \times \mathbf{A} = \nabla(\nabla \cdot \mathbf{A}) - \nabla^2 \mathbf{A}, \tag{17}$$

and Gauß' electric field law to give:

$$\nabla \times \left(\frac{1}{\mu_r} \nabla \times \mathbf{E} \right) - k_0^2 \varepsilon_r \mathbf{E} = -\frac{1}{\mu_r} \nabla^2 \mathbf{E} - k_0^2 \varepsilon_r \mathbf{E}. \tag{18}$$

The basic idea of the shifted-Laplace preconditioner of Erlangga is to use as a preconditioner the original operator $\mathbf{A}()$, but evaluated for a *shifted* value of the wavenumber: $\boldsymbol{\beta} k_0 = \tilde{k}_0$; $\beta, \tilde{k}_0 \in \mathbb{C}$, where the complex valued shift $\beta = \beta_1 + i\beta_2$ is chosen in such a way the preconditioner system $\mathbf{A}(\tilde{k}_0)$ now is positive definite and can be (approximately) solved efficiently, preferably using multigrid. To reduce the effort involved, we will start by using a Krylov subspace method for solution of the preconditioner system $\mathbf{A}(\tilde{k}_0)$ and postpone the realisation of multigrid for the preconditioner solve. This means that the choice for the iterative method to solve the original system $\mathbf{A}(k_0)$ is limited to those methods that allow a *variable* preconditioner, e.g. flexible GMRES or GCR [15,4,18], which is chosen here to solve both the original and the preconditioner systems. A similar approach was taken by Turkel and Erlangga for the 2D Helmholtz equation [16]. Fig. 4 shows the spectrum of the linear system resulting from discretisation of the small cavity test problem (2nd order discretisation using 5976 degrees of freedom), preconditioned with Erlangga's shifted-Laplace preconditioner, where the shift parameters (β_1, β_2) are chosen equal to the optimal values for the Helmholtz equation, i.e. such that $(\beta_1 + i\beta_2)^2 = 1 - 0.5i$.

7. Handling the global radiation boundary condition

One difference between the system matrix in Erlangga et al. [6] and the present matrix is the presence of the fully populated submatrix $\mathbf{A}_{22}(k_0)$, as Erlangga et al. [6] used a local radiation boundary condition, the discretisation of which leads to a completely sparsely populated system. In order to use a preconditioner like the shifted-Laplace preconditioner [6], we have to get rid of the fully populated block. Once this is achieved, we can use the shifted-Laplace preconditioner denoted by $\mathbf{A}_{11}(\tilde{k}_0)$ for the block $\mathbf{A}_{11}(k_0)$. Hence, to solve the system

$$\begin{pmatrix} \mathbf{A}_{11}(k_0) & \mathbf{A}_{12}(k_0) \\ \mathbf{A}_{21}(k_0) & \mathbf{A}_{22}(k_0) \end{pmatrix} \mathbf{x} = \mathbf{b} \tag{19}$$

we proceed as follows:

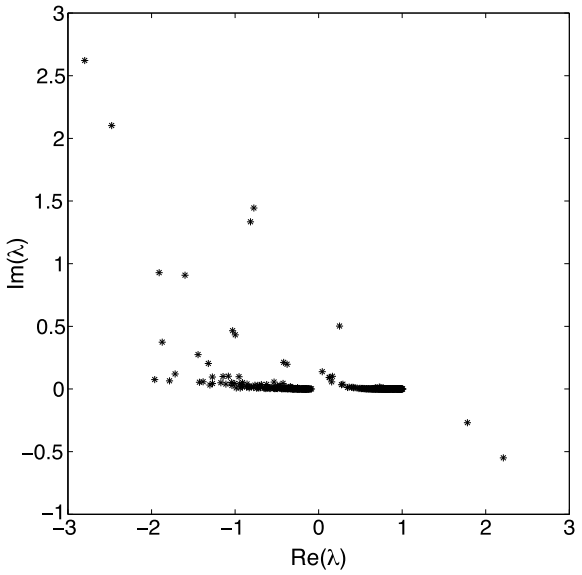

Fig. 5. Spectrum of the preconditioned system for the discretisation of the vector wave equation in a small cavity geometry, using the preconditioner given by (20).

(1) Apply the Krylov method GCR [18] with preconditioner $\mathbf{M}$ (see Algorithm 1):

$$\mathbf{M} = \begin{pmatrix} \mathbf{A}_{11}(k_0) & \mathbf{A}_{12}(k_0) \\ \mathbf{0} & \mathbf{A}_{22}(k_0) \end{pmatrix}. \tag{20}$$

In the GCR method the system $\mathbf{M}\mathbf{s}^k = \mathbf{r}^{k-1}$ has to be solved.

(a) Solve $\mathbf{A}_{22}(k_0)\mathbf{s}_2^k = \mathbf{r}_2^k$ by using a precomputed and stored *LU* decomposition of $\mathbf{A}_{22}(k_0)$.

(b) Solve $\mathbf{s}_1^k$ from

$$\mathbf{A}_{11}(k_0)\mathbf{s}_1^k = \mathbf{r}_1^k - \mathbf{A}_{12}(k_0)\mathbf{s}_2^k = \xi_1 \tag{21}$$

using GCR again with the shifted-Laplace preconditioner $\mathbf{A}_{11}(\tilde{k}_0)$. Now the system $\mathbf{A}_{11}(\tilde{k}_0)\mathbf{z}^j = \mathbf{y}^j$ has to be solved within GCR and this is done using the GCR method for the third time.

The pseudo code of the algorithm is included in Appendix A as Algorithm 2.

Solving the system (19) with this approach leads to a nested iterative method consisting of three GCR loops. In the outer loop (k), the system $\mathbf{A}\mathbf{x} = \mathbf{b}$ is solved using preconditioner $\mathbf{M}$ (see Algorithm 1) hence, iterates $\mathbf{x}^k$ are computed. The middle loop (j) computes the search direction $\mathbf{s}^k$ for $\mathbf{x}^k$ by solving $\mathbf{M}\mathbf{s}^k = \mathbf{r}^{k-1}$ iteratively, where $\mathbf{r}^{k-1}$ is the residual given by $\mathbf{r}^{k-1} = \mathbf{b} - \mathbf{A}\mathbf{x}^{k-1}$. In the j-loop, the two steps (a) and (b) are performed and the system $\mathbf{A}_{11}(k_0)\mathbf{s}_1^k = \xi_1$ is solved using preconditioned GCR with the shifted-Laplace preconditioner $\mathbf{A}_{11}(\tilde{k}_0)$. Finally, the system $\mathbf{A}_{11}(\tilde{k}_0)\mathbf{z}^j = \mathbf{y}^j$ is solved iteratively with unpreconditioned GCR: the (i) loop.

Algorithm 1. GCR algorithm for $\mathbf{A}\mathbf{x} = \mathbf{b}$ with preconditioner $\mathbf{M}$.

Input: $\mathbf{A}$, $\mathbf{b}$ and $\mathbf{M}$
Output: $\mathbf{x}$
  Calculate initial residual $\mathbf{r}^0 = \mathbf{b} - \mathbf{A}\mathbf{x}^0$
  for $k = 1, 2$ **do**
    Solve $\mathbf{M}\mathbf{s}^k = \mathbf{r}^{k-1}$
    $\mathbf{v}^k = \mathbf{A}\mathbf{s}^k$
    Orthogonalize $\mathbf{s}^k$ ($\mathbf{v}^k$) w.r.t. $\mathbf{s}^1, \ldots, \mathbf{s}^{k-1}$ ($\mathbf{v}^1, \ldots, \mathbf{v}^{k-1}$)
    $\mathbf{x}^k = \mathbf{x}^{k-1} + (\mathbf{v}^k)^\top \mathbf{r}^{k-1}\mathbf{s}^k$
    $\mathbf{r}^k = \mathbf{r}^{k-1} - (\mathbf{v}^k)^\top \mathbf{r}^{k-1}\mathbf{v}^k$
  end for

From inspection of the spectrum (Fig. 5) it becomes clear, that omitting the single lower block matrix $(\mathbf{A}_{12}(k_0))^\top$ from the preconditioner has a significant impact on the distribution of the eigenvalues. However, we still use the shifted-Laplace preconditioner for fast solution of (21).

This method is tested for a series of problems with a rectangular cavity of dimensions $1.5\lambda \times 1.5\lambda \times L$, where $L \in \{\lambda, 2\lambda, 4\lambda, 8\lambda\}$. The results of these experiments are given in Table 1. The number of iterations performed in each loop is

Table 1
Performance of GCR using the nested preconditioner **M** (Algorithm 2 in Appendix A). *Ndof* is the number of degrees of freedom in the problem, $\bar{It}_i$ is the average number of inner iterations. The CPU time is normalized with respect to the first experiment (top row). The value of the shift parameters $(\beta_1, \beta_2) = (1.0, 2.5)$.

Dimensions	*Ndof*	It_k, It_j, $\bar{It}_i$	CPU
$1.5\lambda \times 1.5\lambda \times \lambda$	19614	584, 10, 34	1
$1.5\lambda \times 1.5\lambda \times 2\lambda$	39840	832, 10, 35	3.1
$1.5\lambda \times 1.5\lambda \times 4\lambda$	80283	1898, 10, 36	13.9
$1.5\lambda \times 1.5\lambda \times 8\lambda$	160599	4000, 10, 44	142.7

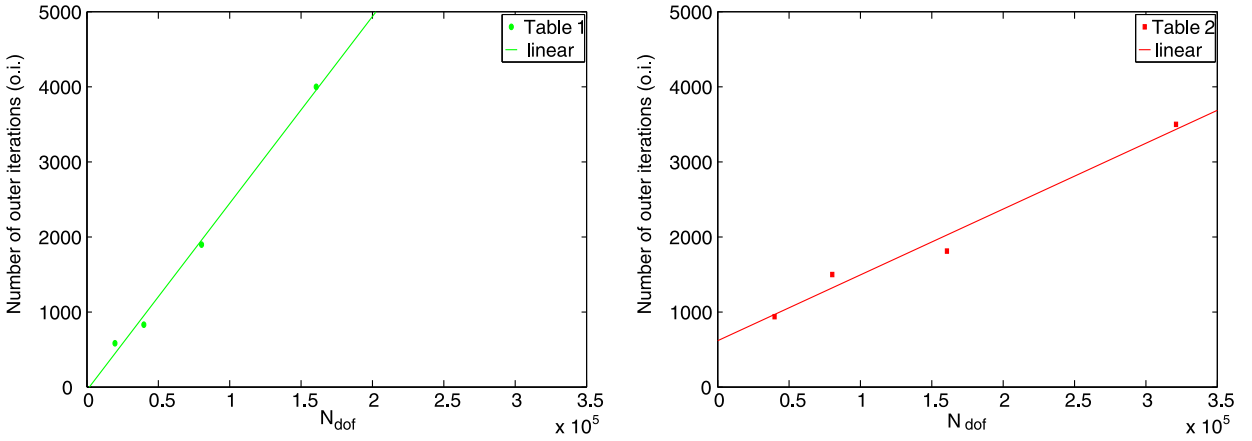

(a) data from Table 1 (b) data from Table 2

Fig. 6. The number of outer loop (k) iterations as a function of the number of degrees of freedom (*Ndof*) for a series of testproblems given in Table 1 (left panel) and Table 2 (right panel).

given in the third column. For the inner loop (i) the average number of iterations per middle loop (j) iteration is given to reach the preset tolerance $\varepsilon_i = 10^{-1}$. The number of middle loop (j) iterations is constant here to test the dependence of the number of outer loop (k) iterations. It is observed that the average number of inner loop (i) iterations increases slightly with the size of the matrix (*Ndof*). However, the number of outer loop (k) iterations increases linearly with *Ndof* leading to very high CPU times as can be seen in the fourth column and Fig. 6 (left panel). Here the CPU times are normalized to the time needed to perform the top-row experiment.

Since the iterative method with three nested loops gives unfeasible high CPU times for deep cavities ($L \geqslant 8\lambda$), a new preconditioner is proposed to accelerate the rate of convergence for solving $\mathbf{Ax} = \mathbf{b}$. This preconditioner is constructed by coupling the former upper diagonal block matrix **M** with the shifted-Laplace preconditioner $\mathbf{A}_{11}(\tilde{k}_0)$:

$$\mathbf{M}_{new} = \begin{pmatrix} \mathbf{A}_{11}(\tilde{k}_0) & \mathbf{A}_{12}(k_0) \\ 0 & \mathbf{A}_{22}(k_0) \end{pmatrix}. \tag{22}$$

Using this new preconditioner reduces the amount of work involved in solving $\mathbf{Ax} = \mathbf{b}$, because the middle loop (j) is not executed anymore. The solution of $\mathbf{Ax} = \mathbf{b}$ is now iteratively computed using preconditioned GCR with only two nested loops: In the outer loop (k) iteration we solve $\mathbf{Ax} = \mathbf{b}$. In the inner loop (i) iteration, we solve the system $\mathbf{M}_{new}\mathbf{s}^k = \mathbf{r}^{k-1}$. Again this system is solved in two steps. First the precomputed *LU* decomposition of the fully populated block $\mathbf{A}_{22}(k_0)$ is used to compute $\mathbf{s}_2^k$. The remaining system $\mathbf{A}_{11}(\tilde{k}_0)\mathbf{s}_1^k = \mathbf{r}_1^k - \mathbf{A}_{12}(k_0)\mathbf{s}_2^k$ is solved with unpreconditioned GCR. The pseudo code of the algorithm is included in Appendix A as Algorithm 3.

To compare the performance of the new preconditioner with the first one, almost the same experiments are performed although the cavity size is increased a little since the new preconditioner yields much better convergence as can be seen in Table 2. The first two columns show the dimensions of the cavities used and the corresponding matrix size. The number of outer loop (k) iterations is given and the average number of inner loop (i) iterations to reach a tolerance level ε_i. Since the preconditioned system $\mathbf{M}_{new}\mathbf{s}^k = \mathbf{r}^{k-1}$ does not need to be solved exactly (because the solution is the search direction for iterates $\mathbf{x}^k$) we use $\varepsilon_i = 10^{-1}$ in first instance. However, for larger cavities with increasing depth ($L \geqslant 8\lambda$) it turns out that the number of outer loop (k) iterations increases fast as we keep the inner loop (i) tolerance to this value of 10^{-1}. Therefore, $\varepsilon_i = 10^{-2}$ for $L = 8\lambda$ and $L = 16\lambda$. With these settings of the inner loop (i) tolerance it is observed that the number of outer loop (k) iterations increases linearly with *Ndof*, but the slope of the curve is far less steep compared to the former method (Fig. 6). The fourth column of the table summarizes the CPU times used by the method, again normalized with respect to the first experiment in Table 1. It is observed that for relatively small cavities we gain a factor of 10 in CPU time. However, for cavities with increasing depth, this factor increases to 25 or more.

Table 2
Performance of GCR using preconditioner $\mathbf{M}_{new}$ (Algorithm 3 in Appendix A). *Ndof* is the number of degrees of freedom in the problem, $\bar{It}_i$ is the average number of inner iterations. The CPU time is normalized with respect to the first experiment in Table 1 (top row). The value of the shift parameters $(\beta_1, \beta_2) = (1.0, 2.5)$.

Dimensions	*Ndof*	It_k, $\bar{It}_i$	CPU
$1.5\lambda \times 1.5\lambda \times 2\lambda$	39 840	938, 28	0.33
$1.5\lambda \times 1.5\lambda \times 4\lambda$	80 283	1500, 28	0.81
$1.5\lambda \times 1.5\lambda \times 8\lambda$	160 599	1812, 68	5.24
$1.5\lambda \times 1.5\lambda \times 16\lambda$	321 066	3500, 75	23.61

Table 3
The average number of inner iterations $\bar{It}_i$, the number of outer iterations It_k and the CPU time (min) for various values of the shift pair (β_1, β_2).

β_1, β_2	$\bar{It}_i$	It_k	CPU	β_1, β_2	$\bar{It}_i$	It_k	CPU
(1.0, 2.5)	57	412	51	(1.0, 0.5)	293	201	168
(1.0, 2.25)	63	416	60	(1.0, 0.25)	496	175	226
(1.0, 2.0)	70	394	61	(1.0, 0.1)	$\geqslant 500$	217	304
(1.0, 1.75)	80	380	59.5	(−1.0, 2.5)	57	459	50
(1.0, 1.5)	91	343	63	(−0.5, 2.5)	53	459	53
(1.0, 1.25)	107	288	66	(0.0, 2.5)	52	553	61
(1.0, 1.0)	139	260	83.5	(0.75, 2.5)	55	432	51
(1.0, 0.75)	187	224	107	(0.5, 2.5)	53	424	49

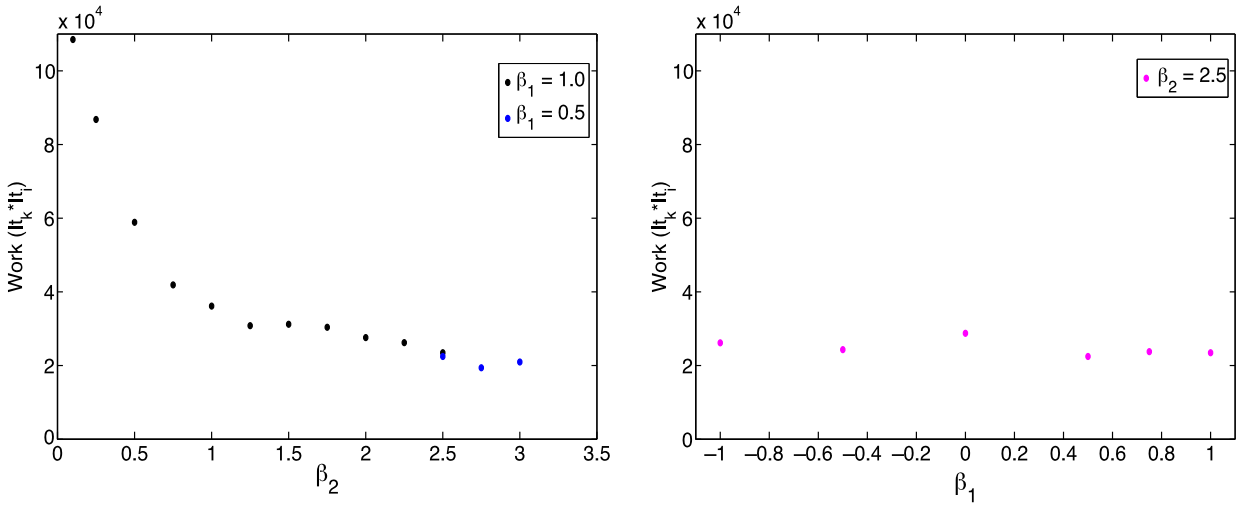

(a) Vary β_2 for β_1 fixed (b) Vary β_1 for β_2 fixed

Fig. 7. The amount of work (approximated by Work $= It_k \cdot \bar{It}_i$) as a function of the shift parameters (β_1, β_2).

8. Tuning parameters

In the previous section the preconditioned GCR method has been introduced. The performance of this method depends strongly on a set of parameters. In this section the sensitivity of the method to these parameters is investigated. First, the shift pair (β_1, β_2) is considered, followed by the inner tolerance parameter ε_i also affecting the average number of inner iterations $\bar{It}_i$.

8.1. Tuning shift pair

According to Erlangga et al. [6], the optimal value for the shift pair $(\beta_1, \beta_2) = (1.0, 0.5)$ resulting in a robust preconditioner for the (scalar) Helmholtz equation. Using these values, the innerloop is only executed a few times using a multigrid approach. At the same time the preconditioner is close enough to the original matrix to reduce the number of outer iterations significantly. To investigate if this is also the case for the preconditioner $\mathbf{M}_{new}$ considered here, the effect on both the number of inner and outer iterations for a cavity of $(1.5\lambda \times 1.5\lambda \times 1.0\lambda)$ and $Ndof = 22\,407$ is investigated for a range of alternative values of (β_1, β_2).

At first the complex shift β_2 is tuned while β_1 is fixed. Then we use the most promising value for β_2 to tune the real shift β_1. The results of this experiment are summarized in Table 3 and visualized in Fig. 7.

According to Table 3 the average number of inner iterations $\overline{It}_i$ depends heavily on the complex shift β_2. While $\overline{It}_i$ increases rapidly when $\beta_2 \to 0$, the corresponding number of outer iterations It_k decreases slowly resulting in an increase in CPU time (column 4). This increment in CPU time is easily explained by considering the amount of work for the method approximated by the formula

$$\text{Work} \approx \overline{It}_i \cdot It_k. \tag{23}$$

Using this estimate for the work involved in the method, it is clear that the couple $(\beta_1, \beta_2) = (1.0, 2.5)$ (work = 23 484) results in less work (and hence lower CPU) than for example the couple $(\beta_1, \beta_2) = (1.0, 0.5)$ (work = 49 224), see Fig. 7.

Considering Table 3 once more it is observed that a much larger value for β_2 is found here in comparison with the optimal value found by Erlangga et al. [6,5,19]. By applying a multigrid method to solve the preconditioned system, only a few iterations are necessary to solve this system accurately with $\beta_2 = 0.5$. Our preconditioned system is much harder to solve as can be seen by the value of $\overline{It}_i$. From Table 3 it is clear that we have to reduce the amount of inner iterations (even) more to decrease the amount of work and eventually CPU. Also it is clear from the table that the tuning of β_1 does not have a great influence on the performance of the method.

8.2. Tuning inner loop tolerance

In the last section we tuned the shift parameters (β_1, β_2) in order to reduce the CPU time. The CPU time does not depend on these two parameters solely, since the tolerance in the inner loop is also very important. In this section we will investigate the dependence of the inner loop tolerance ε_i on the number of outer iterations and hence on the dependence of the CPU time.

The expectation is that solving the preconditioner system $\mathbf{A}(\tilde{k}_0)\mathbf{z}_1 = \mathbf{v}_1$ very accurately is necessary nor efficient to reduce the number of outer iterations. This is justified by the fact that the matrix $\mathbf{A}(\tilde{k}_0)_{11}$ in Eq. (22) is only an approximation of the submatrix $\mathbf{A}(k_0)_{11}$ of the original matrix $\mathbf{A}$. To check whether this is really the case, the following experiment is performed. Using a model problem of dimensions $1.5\lambda \times 1.5\lambda \times 4.0\lambda$ with 80 283 degrees of freedom, the inner loop tolerance ε_i is decreased from 10^{-1} to 10^{-8} and the effect on the number of inner and outer iterations is investigated. The results are summarized in Table 4 and Fig. 8.

From Table 4 and Fig. 8 it is clear that the number of outer loop (k) iterations does not decrease significantly if we solve the preconditioned system with high accuracy whereas the average number of inner loop (i) iterations increases dramatically. However, a significant difference in the number of outer loop (k) iterations is observed for $\varepsilon_i = 10^{-1}$ and $\varepsilon_i = 10^{-2}$ (see Fig. 8). Up to 600 outer loop (k) iterations the convergence rate is comparable. If more outer loop (k) iterations are needed, it turns out solving the system $\mathbf{A}(\tilde{k}_0)\mathbf{z}_1 = \mathbf{v}_1$ should be done more accurate for a better convergence. Note here also that the average number of inner loop (i) iterations and hence the total amount of work involved increases rapidly when going from $\varepsilon_i = 10^{-1}$ to $\varepsilon_i = 10^{-2}$. Therefore, as a rule of thumb the system can be solved with a tolerance of $\varepsilon_i = 10^{-1}$ for shallow cavities with a depth up to 4λ. For cavities with $L \geqslant 8\lambda$ it is advised to use an inner loop tolerance $\varepsilon_i = 10^{-2}$.

9. Limitations of the current approach

As mentioned before, the use of GCR to solve the preconditioner system, dictates the original system be solved with a Krylov method that allows the preconditioner to be nonconstant. In practice this means a long-recurrence method has to be used, requiring storage of the Krylov base vectors for all iterations. We found this to be unpractical for problems with a very large number of degrees of freedom. Although many restarting and truncating strategies have been tried, none of these had satisfactory convergence properties, most likely a result of the indefiniteness of the original system. Therefore, it is necessary to use multigrid to solve the preconditioner system. Currently, an algebraic multigrid algorithm is being incorporated in the algorithm to perform this task.

10. Recommendations for future research

A slightly more complicated problem that is strongly related to cavity scattering is the scattering by dielectric bodies. Inside the dielectric body the vector wave equation is discretised, like it is done inside the cavity. However, on the surface of a dielectric body both an electric and a magnetic current are induced, and both have to be taken into account in the radiation boundary condition. Generally, like in the cavity scattering problem this boundary condition has a formulation based on the Electric Field Integral Equation (EFIE). For the dielectric body scattering problem an alternative formulation has been derived by Liu and Jin [12], where the radiation boundary condition is based on the Combined Field Integral Equation (CFIE). For the latter problem it can be shown that the eigenvalue spectrum of the CFIE formulation is significantly more favourable than the original EFIE based formulation. This is not that surprising as similar behavior is commonly encountered for integral equation solvers for perfectly electrically conducting scatterers based on either of the two formulations. On top of that the new formulation allows a preconditioner system to be formulated based on a local instead of a global radiation boundary condition (see Liu and Jin [13]). The reduced accuracy of the local boundary condition in the preconditioner

Table 4
The influence of a high accuracy in the inner loop on the average number of inner ($\bar{It}_i$) and outer iterations (It_k). For these experiments we have used an upper limit of 1000 outer iterations. For $\varepsilon_i = 10^{-1}$ the relative residual norm is not converged to the outer tolerance $\varepsilon_o = 10^{-4}$. The model problem has dimensions $1.5\lambda \times 1.5\lambda \times 4.0\lambda$ and 80 283 degrees of freedom.

ε_i	CPU	$\bar{It}_i$	It_k
10^{-1}	22	28.18	> 1000
10^{-2}	41	59.81	≈ 1000
10^{-3}	69	97.18	896
10^{-4}	104	137.13	891
10^{-8}	247	298.73	880

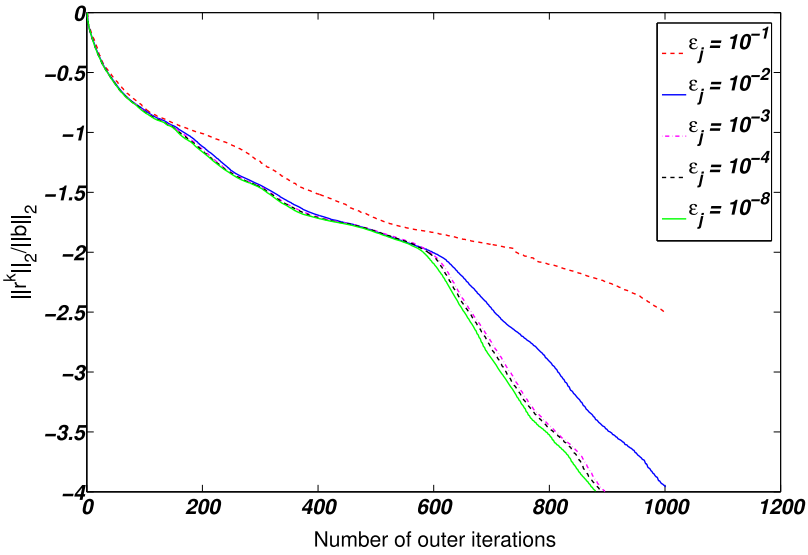

Fig. 8. The number of outer iterations for a decreasing inner loop tolerance ε_j. For this experiment a rectangular cavity of dimensions $1.5\lambda \times 1.5\lambda \times 4.0\lambda$ and 80 283 degrees of freedom.

system does not affect the solution of the original system, but will make the preconditioner solve more efficient, because the block triangular formulation of the preconditioner is no longer required.

Furthermore, to be applicable for problems with large number of degrees of freedom the algorithm has to be modified to be able to handle multiple right-hand sides simultaneously.

11. Conclusions

The shifted-Laplace preconditioner of Erlangga et al. [6,5,19] provides an effective means to improve the convergence rate of iterative solution of linear systems derived from the discretised vector wave equation. A start has been made with exploring the sensitivity of the computational work to the parameters that define the preconditioner. However, to be able to solve problems with a large number of degrees of freedom it is found to be essential to use multigrid to solve the preconditioner system. Not only will this reduce the computational work involved in the preconditioner solve, but it will also allow the use of short recurrence Krylov methods to solve the original system, as the preconditioner then has become constant. Currently, an algebraic multigrid solver is being incorporated in the algorithm. Two additional propositions are made to improve the efficiency of the method: an alternative formulation of the radiation boundary condition on the aperture that will improve the condition number of the original system and the use of a Krylov method that is able to simultaneously solve for multiple right-hand sides.

Appendix A. Algorithm pseudo code

In this paper we discuss two algorithms to solve a large linear system $\mathbf{Ax} = \mathbf{b}$. The main difference between Algorithm 2 and Algorithm 3 is that in the latter we have only 2 loops. This is due to the fact that we use only 1 preconditioner instead of 2 in the former. In this way a reduction of the work compared to the 3-loop GCR is obtained. The loop in Algorithm 2 is not executed in Algorithm 3 anymore.

Algorithm 2. Nested GCR algorithm for $\mathbf{Ax} = \mathbf{b}$ with preconditioner $\mathbf{M}$ from Eq. (20) and the shifted-Laplace preconditioner $\mathbf{A}_{11}(\tilde{k}_0)$.

Input: $\mathbf{A}$, $\mathbf{b}$ and preconditioners $\mathbf{M}$ and $\mathbf{A}_{11}(\tilde{k}_0)$
Output: $\mathbf{x}$

Start GCR-1 for solving $\mathbf{Ax} = \mathbf{b}$
Calculate initial residual $\mathbf{r}^0 = \mathbf{b} - \mathbf{Ax}^0$
for $k = 1, 2$ **do**
  Solve $\mathbf{A}_{22}(k_0)\mathbf{s}_2^k = \mathbf{r}_2^{k-1}$ using stored *LU* composition
  Start GCR-2 for solving $\mathbf{A}_{11}(k_0)\mathbf{s}_1^k = \mathbf{r}_1^{k-1} - \mathbf{A}_{12}(k_0)\mathbf{s}_2^{k-1} = \xi_1$
  Calculate initial residual $\tilde{\mathbf{r}}^0 = \xi_1 - \mathbf{A}_{11}(k_0)\tilde{\mathbf{s}}^0$
  for $j = 1, 2$ **do**
    Start GCR-3 for solving $\mathbf{A}_{11}(\tilde{k}_0)\tilde{\mathbf{s}}^j = \tilde{\mathbf{r}}^{j-1}$
    Calculate initial residual $\bar{\mathbf{r}}^0 = \tilde{\mathbf{r}}^{j-1} - \mathbf{A}_{11}(\tilde{k}_0)\bar{\mathbf{s}}^0$
    for $i = 1, 2$ **do**
      $\bar{\mathbf{s}}^i = \bar{\mathbf{r}}^{i-1}$
      $\bar{\mathbf{v}}^i = \mathbf{A}_{11}(\tilde{k}_0)\bar{\mathbf{s}}^i$
      Orthogonalize $\bar{\mathbf{s}}^i$ $(\bar{\mathbf{v}}^i)$ w.r.t. $\bar{\mathbf{s}}^1, \ldots, \bar{\mathbf{s}}^{i-1}$ $(\bar{\mathbf{v}}^1, \ldots, \bar{\mathbf{v}}^{i-1})$
      $\mathbf{w}^i = \mathbf{w}^{i-1} + (\bar{\mathbf{v}}^i)^\top \bar{\mathbf{r}}^{i-1}\,\bar{\mathbf{s}}^i$
      $\bar{\mathbf{r}}^i = \bar{\mathbf{r}}^{i-1} - (\bar{\mathbf{v}}^i)^\top \bar{\mathbf{r}}^{i-1}\,\bar{\mathbf{v}}^i$
    end for
    $\tilde{\mathbf{s}}^j = \mathbf{w}^i$
    $\tilde{\mathbf{v}}^j = \mathbf{A}_{11}(k_0)\tilde{\mathbf{s}}^j$
    Orthogonalize $\tilde{\mathbf{s}}^j$ $(\tilde{\mathbf{v}}^j)$ w.r.t. $\tilde{\mathbf{s}}^1, \ldots, \tilde{\mathbf{s}}^{j-1}$ $(\tilde{\mathbf{v}}^1, \ldots, \tilde{\mathbf{v}}^{j-1})$
    $\mathbf{y}^j = \mathbf{y}^{j-1} + (\tilde{\mathbf{v}}^j)^\top \tilde{\mathbf{r}}^{j-1}\,\tilde{\mathbf{s}}^j$
    $\tilde{\mathbf{r}}^j = \tilde{\mathbf{r}}^{j-1} - (\tilde{\mathbf{v}}^j)^\top \tilde{\mathbf{r}}^{j-1}\,\tilde{\mathbf{v}}^j$
  end for
  $\mathbf{s}_1^k = \mathbf{y}^j$
  $\mathbf{s}^k = (\mathbf{s}_1^k, \mathbf{s}_2^k)^\top$
  $\mathbf{v}^k = \mathbf{As}^k$
  Orthogonalize $\mathbf{s}^k$ $(\mathbf{v}^k)$ w.r.t. $\mathbf{s}^1, \ldots, \mathbf{s}^{k-1}$ $(\mathbf{v}^1, \ldots, \mathbf{v}^{k-1})$
  $\mathbf{x}^k = \mathbf{x}^{k-1} + (\mathbf{v}^k)^\top \mathbf{r}^{k-1}\,\mathbf{s}^k$
  $\mathbf{r}^k = \mathbf{r}^{k-1} - (\mathbf{v}^k)^\top \mathbf{r}^{k-1}\,\mathbf{v}^k$
end for

Algorithm 3. Nested GCR algorithm for $\mathbf{Ax} = \mathbf{b}$ with preconditioner $\mathbf{M}_{new}$ from Eq. (22).

Input: $\mathbf{A}$, $\mathbf{b}$ and preconditioner $\mathbf{M}_{new}$
Output: $\mathbf{x}$

Start GCR-1 for solving $\mathbf{Ax} = \mathbf{b}$
Calculate initial residual $\mathbf{r}^0 = \mathbf{b} - \mathbf{Ax}^0$
for $k = 1, 2$ **do**
  Solve $\mathbf{A}_{22}(k_0)\mathbf{s}_2^k = \mathbf{r}_2^{k-1}$ using stored *LU* composition
  Start GCR-2 for solving $\mathbf{A}_{11}(\tilde{k}_0)\mathbf{s}_1^k = \mathbf{r}_1^{k-1} - \mathbf{A}_{12}(k_0)\mathbf{s}_2^{k-1} = \xi_1$
  Calculate initial residual $\bar{\mathbf{r}}^0 = \xi_1 - \mathbf{A}_{11}(\tilde{k}_0)\bar{\mathbf{s}}^0$
  for $i = 1, 2$ **do**
    $\bar{\mathbf{s}}^i = \bar{\mathbf{r}}^{i-1}$
    $\bar{\mathbf{v}}^i = \mathbf{A}_{11}(\tilde{k}_0)\bar{\mathbf{s}}^i$
    Orthogonalize $\bar{\mathbf{s}}^i$ $(\bar{\mathbf{v}}^i)$ w.r.t. $\bar{\mathbf{s}}^1, \ldots, \bar{\mathbf{s}}^{i-1}$ $(\bar{\mathbf{v}}^1, \ldots, \bar{\mathbf{v}}^{i-1})$
    $\mathbf{w}^i = \mathbf{w}^{i-1} + (\bar{\mathbf{v}}^i)^\top \bar{\mathbf{r}}^{i-1}\,\bar{\mathbf{s}}^i$
    $\bar{\mathbf{r}}^i = \bar{\mathbf{r}}^{i-1} - (\bar{\mathbf{v}}^i)^\top \bar{\mathbf{r}}^{i-1}\,\bar{\mathbf{v}}^i$
  end for
  $\mathbf{s}_1^k = \mathbf{w}^i$
  $\mathbf{s}^k = (\mathbf{s}_1^k, \mathbf{s}_2^k)^\top$
  $\mathbf{v}^k = \mathbf{As}^k$
  Orthogonalize $\mathbf{s}^k$ $(\mathbf{v}^k)$ w.r.t. $\mathbf{s}^1, \ldots, \mathbf{s}^{k-1}$ $(\mathbf{v}^1, \ldots, \mathbf{v}^{k-1})$
  $\mathbf{x}^k = \mathbf{x}^{k-1} + (\mathbf{v}^k)^\top \mathbf{r}^{k-1}\,\mathbf{s}^k$
  $\mathbf{r}^k = \mathbf{r}^{k-1} - (\mathbf{v}^k)^\top \mathbf{r}^{k-1}\,\mathbf{v}^k$
end for

References

[1] H. Anastassiu, A review of electromagnetic scattering analysis for inlets, cavities and open ducts, IEEE Trans. Antennas and Propagation 45 (6) (2003) 27–40.
[2] R. Coifman, V. Rohklin, S. Wandzura, The fast multipole method for the wave equation: A pedestrian description, IEEE Trans. Antennas and Propagation 35 (3) (1993) 7–12.
[3] I. Duff, J. Scott, The design of a new frontal code for solving sparse, unsymmetric systems, ACM Trans. Math. Software 22 (1) (1996) 30–45.
[4] S. Eisenstat, H. Elman, M. Schultz, Variational iterative methods for nonsymmetric systems of linear equations, SIAM J. Numer. Anal. 20 (1983) 345–357.
[5] Y. Erlangga, C. Oosterlee, C. Vuik, A novel multigrid based preconditioner for heterogeneous Helmholtz problems, SIAM J. Sci. Comput. 27 (2006) 1471–1492.
[6] Y. Erlangga, C. Vuik, C. Oosterlee, On a class of preconditioners for solving the Helmholtz equation, Appl. Numer. Math. 50 (2004) 409–425.
[7] R. Graglia, D. Wilton, A. Peterson, Higher order interpolatory vector bases for computational electromagnetics, IEEE Trans. Antennas and Propagation 45 (3) (1997) 329–342.
[8] R. Irons, A frontal method solution program for finite element analysis, Internat. J. Numer. Methods Engrg. 2 (1970) 5–32.

[9] J. Jin, Electromagnetic scattering from large, deep and arbitrarily shaped open cavities, Electromagnetics 18 (1998) 3–34.
[10] J. Jin, J. Liu, Z. Lou, C. Liang, A fully high-order finite-element simulation of scattering by deep cavities, IEEE Trans. Antennas and Propagation 51 (9) (2003) 2420–2429.
[11] J. Liu, J.-M. Jin, A special higher order finite-element method for scattering by deep cavities, IEEE Trans. Antennas and Propagation 48 (5) (2000) 694–703.
[12] J. Liu, J.-M. Jin, A novel hybridization of higher order finite element and boundary integral methods for electromagnetic scattering and radiation problems, IEEE Trans. Antennas and Propagation 49 (12) (2001) 1794–1806.
[13] J. Liu, J.-M. Jin, A highly effective preconditioner for solving the finite element-boundary integral matrix equation of 3-D scattering, IEEE Trans. Antennas and Propagation 50 (9) (2002) 1212–1221.
[14] J. Liu, J.-M. Jin, Scattering analysis of a large body with deep cavities, IEEE Trans. Antennas and Propagation 51 (6) (2003) 1157–1167.
[15] Y. Saad, A flexible inner–outer preconditioned GMRES, SIAM J. Sci. Stat. Comput. 14 (1993) 461–469.
[16] E. Turkel, Y. Erlangga, Preconditioning a finite element solver of the exterior Helmholtz equation, in: Proceedings of ECCOMAS CFD 2006, 2006.
[17] D. van der Heul, H. van der Ven, J.-W. van der Burg, Full wave analysis of the influence of the jet engine air intake on the radar signature of modern fighter aircraft, in: Proceedings of ECCOMAS CFD 2006, 2006.
[18] H. van der Vorst, C. Vuik, GMRESR: a family of nested GMRES methods, Numer. Linear Algebra Appl. 1 (1994) 369–386.
[19] 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 (2007) 1942–1958.