Convergence and complexity study of GMRES variants for solving multi-frequency elastic wave propagation problems

Baumann, Manuel; van Gijzen, Martin B.

DOI
10.1016/j.jocs.2018.03.004

Publication date
2018

Document Version
Final published version

Published in
Journal of Computational Science

Citation (APA)
Baumann, M., & van Gijzen, M. B. (2018). Convergence and complexity study of GMRES variants for solving multi-frequency elastic wave propagation problems. *Journal of Computational Science*, *26*, 285-293. https://doi.org/10.1016/j.jocs.2018.03.004

Journal of Computational Science 26 (2018) 285–293

Contents lists available at ScienceDirect

Journal of Computational Science

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

Convergence and complexity study of GMRES variants for solving multi-frequency elastic wave propagation problems

Manuel Baumann*, Martin B. van Gijzen

Delft Institute of Applied Mathematics, Delft University of Technology, The Netherlands

ARTICLE INFO

Article history:
Received 14 October 2017
Received in revised form 5 February 2018
Accepted 6 March 2018
Available online 8 March 2018

Keywords:
Time-harmonic elastic wave equation
Global GMRES
Multi-shift GMRES
Nested multi-shift Krylov methods
Additive coarse grid correction

ABSTRACT

In this paper we present a comparison study of three different frameworks of iterative Krylov methods that we have recently developed for the simultaneous numerical solution of frequency-domain wave propagation problems when multiple wave frequencies are present. The three approaches have in common that they require the application of a single shift-and-invert preconditioner at a suitable *seed* frequency. In particular for three-dimensional problems, we present the efficient application of the elastic shift-and-invert preconditioner by means of an additive coarse grid correction. The focus of the present work lies, however, on the performance of the respective iterative method. We conclude with numerical examples that provide guidance concerning the suitability of the three methods.

1. Introduction

After spatial discretization using, for instance, the finite element method with N degrees of freedom [6,Section 2], the time-harmonic wave equation has the form,

$$(K + i\omega_k C - \omega_k^2 M)x_k = b, \quad \omega_k := 2\pi f_k, \quad k = 1, \ldots, n_\omega, \tag{1}$$

with stiffness matrix K, mass matrix M, and C consisting of Sommerfeld boundary conditions [2] modelling absorption. Note that (1) yields a sequence of n_ω linear systems of equations. One way to solve the systems (1) simultaneously is to define the block matrix of unknowns, $\mathbf{X} := [x_1, \ldots, x_{n_\omega}] \in \mathbb{C}^{N \times n_\omega}$, and to note that (1) can be rewritten as a linear matrix equation,

$$\mathcal{A}(\mathbf{X}) := K\mathbf{X} + iC\mathbf{X}\Omega - M\mathbf{X}\Omega^2 = B, \tag{2}$$

with $\Omega := \mathrm{diag}(\omega_1, \ldots, \omega_{n_\omega})$ and $B := b\mathbb{1}^T$.

The matrix Eq. (2) can then be solved using a global Krylov method [18], for instance global GMRES [21] or global IDR(s) [4,6].

A second approach is to consider a linearization [31] of (1) of the form,

$$\left(\begin{bmatrix} iC & K \\ I & 0 \end{bmatrix} - \omega_k \begin{bmatrix} M & 0 \\ 0 & I \end{bmatrix} \right) \begin{bmatrix} \omega_k x_k \\ x_k \end{bmatrix} = \begin{bmatrix} b \\ 0 \end{bmatrix}, \quad k = 1, \ldots, n_\omega, \tag{3}$$

where the angular frequencies $\omega_1, \ldots, \omega_{n_\omega}$ appear as a (linear) shift. For the short-hand notation of (3), we define the two block matrices,

$$\mathcal{K} := \begin{bmatrix} iC & K \\ I & 0 \end{bmatrix} \in \mathbb{C}^{2N \times 2N} \quad \text{and} \quad \mathcal{M} := \begin{bmatrix} M & 0 \\ 0 & I \end{bmatrix} \in \mathbb{C}^{2N \times 2N}, \tag{4}$$

and write (3) as $(\mathcal{K} - \omega_k \mathcal{M})\mathbf{x}_k = \mathbf{b}$, for $\mathbf{x}_k := [\omega_k x_k, x_k]^T$ and $\mathbf{b} := [b, 0]^T$. Note that we will consider the case $C \equiv 0$ independently. The matrix Eq. (2) then reduces to two terms, and we can identify $\mathcal{K} = K$ as well as $\mathcal{M} = M$ and avoid doubling of dimensions in (3). In this paper, we review and compare the following recently developed algorithms:

- Global GMRES [21] suitable for the matrix equation approach (2) (cf. Algorithm 1 and [6]),
- Polynomial preconditioners [1,8] for multi-shift GMRES (cf. Algorithm 2),
- Nested multi-shift FOM-FGMRES as presented in [7] (cf. Algorithms 3 and 4).

The application of a single shift-and-invert preconditioner is crucial for the fast convergence of the above methods and yields the

* Corresponding author.
E-mail address: m.m.baumann@tudelft.nl (M. Baumann).
URL: http://www.manuelbaumann.de (M. Baumann).

https://doi.org/10.1016/j.jocs.2018.03.004
1877-7503/

main computational cost per iteration. Note that this list does not consider a comparison with the algorithms suggested by [5,29] and by [32]. Particularly, we do not consider a comparison with [5,29] because both frameworks require the exact application of multiple shift-and-invert preconditioners which is unpractical when n_ω is only in the order of ten. Moreover, we restrict ourselves to GMRES-variants of the respective algorithms, and refer to [4] for global IDR(s) and to [7] for the more memory-efficient combination of nested IDR-FQMRIDR(s). In [1] a shifted polynomial preconditioner is used within multi-shift BiCG. The derivations in Section 2 emphasize that the costs-per-iteration of each proposed algorithm are comparable, and are dominated by the shift-and-invert preconditioner. In Section 3, we evaluate the three approaches for a benchmark problem of the discretized time-harmonic elastic wave equation with respect to the convergence behavior under different circumstances. Because the shift-and-invert preconditioner is dominating the computations, we put special emphasis on the practical application of the preconditioner in three spatial dimensions in Section 4. The efficient application of the preconditioner in 3D is closely related to the work on the Complex Shifted Laplacian for the acoustic Helmholtz equation [15,26] using multigrid [14] or domain decomposition approaches [19]. In particular, we demonstrate that also the damped *elastic* preconditioner can be applied in linear computational complexity by means of an additive coarse grid correction.

2. Review of the developed algorithms

The review of the subsequent algorithms is based on our works [6–8].

2.1. The preconditioned matrix equation approach

The matrix Eq. (2) with right preconditioning reads,

$$\mathcal{A}(P(\tau)^{-1}\mathbf{Y}) = B, \quad \mathbf{X} = P(\tau)^{-1}\mathbf{Y}, \tag{5a}$$

$$\text{where } P(\tau)^{-1} := (K + i\tau C - \tau^2 M)^{-1}, \tag{5b}$$

and $\mathcal{A}(\cdot)$ as in (2). A similar reformulation has been suggested in [32] where the author exploits the flexibility of solving a Sylvester equation instead of shifted systems. We note that in (5a) and (5b), the preconditioner $P(\tau)$ can be applied inexactly using, for instance, an incomplete LU factorization. The (possibly complex) parameter τ in (5b) is called the *seed* frequency.

Algorithm 1. Right-preconditioned global GMRES [21] for the matrix Eq. (5a) and (5b).

Set $R_0 = B$, $V_1 = R_0/\|R_0\|_F$ ▷ Initialization ($\mathbf{X}_0 = \mathbf{0}$)
for $j = 1$ to m **do**
  Apply $W = \mathcal{A}(P(\tau)^{-1}V_j)$ ▷ Preconditioner (5b)
  for $i = 1$ to j **do** ▷ Block-Arnoldi method
    $h_{i,j} = \mathrm{tr}(W^H V_i)$
    $W = W - h_{i,j}V_i$
  end for
  Set $h_{j+1,j} = \|W\|_F$ and $V_{j+1} = W/h_{j+1,j}$
end for
Set $\underline{H}_m = [h_{i,j}]_{i=1,\dots,m}^{j=1,\dots,m+1}$, $\mathbf{V}_m = [V_1, \dots, V_m]$
Solve $\mathbf{y}_m = \mathrm{argmin}_{\mathbf{y}} \left\| \underline{H}_m \mathbf{y} - \|B\|_F \mathbf{e}_1 \right\|_2$ ▷ $\mathbf{e}_1 \in \mathbb{C}^{m+1}$
Compute $\mathbf{X}_m = P(\tau)^{-1}(\mathbf{V}_m {}^* \mathbf{y}_m)$ ▷ '$*$' denotes the so-called star product [21]

In Algorithm 1, we state the global GMRES method [18,21]. Note that in the block Arnoldi method the trace inner product is used, and norms are replaced by the Frobenius norm $\|\cdot\|_F$ for block matrices. After m iterations, an approximate solution to (2) in the block Krylov subspace $\mathcal{K}_m(AP(\tau)^{-1}, B)$ is obtained. An efficient preconditioner will yield convergence for $m \ll N$.

2.2. Preconditioners for shifted linear systems

The methods presented in this section are both two-level preconditioning approaches. As a first-level preconditioner, a shift-and-invert preconditioner of the form,

$$\begin{aligned} \mathcal{P}(\tau)^{-1} &:= (\mathcal{K} - \tau\mathcal{M})^{-1} \overset{(4)}{=} \left(\begin{bmatrix} iC & K \\ I & 0 \end{bmatrix} - \tau \begin{bmatrix} M & 0 \\ 0 & I \end{bmatrix} \right)^{-1} \\ &= \begin{bmatrix} I & \tau I \\ 0 & I \end{bmatrix} \begin{bmatrix} I & 0 \\ 0 & (K + i\tau C - \tau^2 M)^{-1} \end{bmatrix} \begin{bmatrix} 0 & I \\ I & -iC + \tau M \end{bmatrix}, \end{aligned} \tag{6}$$

is applied. Based on the decomposition (6) we note that $P(\tau)^{-1} = (K + i\tau C - \tau^2 M)^{-1}$ as defined in (5b) is the main computational work and, hence, the work-per-iteration is comparable to Algorithm 1. When applying a scaled shift-and-invert preconditioner to the block systems (3), the following equivalence holds,

$$(\mathcal{K} - \omega_k \mathcal{M})\mathcal{P}_k^{-1}\mathbf{y}_k = \mathbf{b} \quad \Leftrightarrow \quad (\mathcal{K}\mathcal{P}(\tau)^{-1} - \eta_k I)\mathbf{y}_k = \mathbf{b}, \tag{7}$$

if we introduce $\eta_k := \omega_k/(\omega_k - \tau)$, and the scaled preconditioners $\mathcal{P}_k^{-1} := (1 - \eta_k)\mathcal{P}(\tau)^{-1} = (1 - \eta_k)(\mathcal{K} - \tau\mathcal{M})^{-1}$. Note that the latter is a preconditioned shifted linear system with (complex) shifts η_k and a system matrix $\mathcal{C} := \mathcal{K}\mathcal{P}(\tau)^{-1} = \mathcal{K}(\mathcal{K} - \tau\mathcal{M})^{-1}$. Due to the equivalence in (7), the preconditioner (6) needs to be applied exactly. Moreover, right-preconditioning requires the back-substitution, $\mathbf{x}_k = \mathcal{P}_k^{-1}\mathbf{y}_k$.

2.2.1. A shifted Neumann polynomial preconditioner

After applying the shift-and-invert preconditioner (6) to (3), we still must solve,

$$(\mathcal{C} - \eta_k I)\mathbf{y}_k = \mathbf{b}, \quad \mathbf{x}_k = \mathcal{P}_k^{-1}\mathbf{y}_k, \quad k = 1, \dots, n_\omega, \tag{8}$$

where $\mathcal{C} = \mathcal{K}\mathcal{P}(\tau)^{-1}$, and with shifts $\eta_k = \omega_k/(\omega_k - \tau)$. Efficient algorithms for shifted linear systems such as (8) rely on the shift-invariance property, i.e. the identity, $\mathcal{K}_m(\mathcal{C}, \mathbf{b}) \equiv \mathcal{K}_m(\mathcal{C} - \eta I, \mathbf{b})$, for any shift $\eta \in \mathbb{C}$; cf. [17,30]. The (preconditioned) spectrum of the matrix $\mathcal{C}$ is known to be enclosed by a circle of radius R and center c [8,36]. Therefore, the Neumann preconditioner p_n [28,Chapter 12.3] of degree n,

$$\mathcal{C}^{-1} \approx \sum_{i=0}^{n} \left(I - \xi\mathcal{C}\right)^i =: p_n(\mathcal{C}), \quad \text{with} \quad \xi = \frac{1}{c} = -\frac{\tau - \bar{\tau}}{\bar{\tau}}, \tag{9}$$

has optimal spectral radius [8]. The polynomial preconditioner (9) can be represented in a different basis as, $p_n(\mathcal{C}) = \sum_{i=0}^{n} \alpha_i \mathcal{C}^i$. Shift-invariance can then be preserved if the following holds,

$$(\mathcal{C} - \eta_k I)p_{n,k}(\mathcal{C}) = \mathcal{C}p_n(\mathcal{C}) - \tilde{\eta}_k I, \tag{10}$$

where $p_{n,k}(\mathcal{C}) = \sum_{i=0}^{n} \alpha_{i,k}\mathcal{C}^i$ is a polynomial preconditioner for $(\mathcal{C} - \eta_k I)$. Substitution yields,

$$\sum_{i=0}^{n} \alpha_{i,k}\mathcal{C}^{i+1} - \sum_{i=0}^{n} \eta_k \alpha_{i,k}\mathcal{C}^i - \sum_{i=0}^{n} \alpha_i \mathcal{C}^{i+1} + \tilde{\eta}_k I = 0. \tag{11}$$

The latter (11) is a difference equation and can be solved in closed form [1]:

$$\alpha_{n,k} = \alpha_n, \tag{12a}$$

$$\alpha_{i-1,k} = \alpha_{i-1} + \eta_k \alpha_{i,k}, \quad \text{for} \quad i = n, \dots, 1, \tag{12b}$$

$$\tilde{\eta}_k = \eta_k \alpha_{0,k}. \tag{12c}$$

Algorithm 2. Multi-shift GMRES [17] with polynomial preconditioner (9) for systems (8), cf. [1,8].

Set $\mathbf{r}_0 = \mathbf{b}$, $\mathbf{v}_1 = \mathbf{r}_0/\|\mathbf{r}_0\|$ ▷ Initialization ($\mathbf{x}_0 = \mathbf{0}$)
for $j = 1$ to m **do**
  Apply $\mathbf{w} = \mathcal{C}p_n(\mathcal{C})\mathbf{v}_j$ ▷ Requires $n+1$ `MatVec`'s
  for $i = 1$ to j **do** ▷ Arnoldi method
    $h_{i,j} = \mathbf{w}^H\mathbf{v}_i$
    $\mathbf{w} = \mathbf{w} - h_{i,j}\mathbf{v}_i$
  end for
  Set $h_{j+1,j} = \|\mathbf{w}\|$ and $\mathbf{v}_{j+1} = \mathbf{w}/h_{j+1,j}$
end for
Set $\underline{H}_m = [h_{i,j}]_{i=1,\dots,m}^{j=1,\dots,m+1}$ and $V_m = [\mathbf{v}_1, \dots, \mathbf{v}_m]$
for $k = 1$ to n_ω **do**
  Solve $\mathbb{C}^m \ni \mathbf{z}_k = \operatorname{argmin}_{\mathbf{z}} \left\| (\underline{H}_m - \tilde{\eta}_k \underline{I}_m)\mathbf{z} - \|\mathbf{r}_0\|\mathbf{e}_1 \right\|$ ▷ New shifts $\tilde{\eta}_k$ acc. to (12c)
  Resubstitute $\mathbf{y}_k = p_{n,k}(\mathcal{C})V_m\mathbf{z}_k$ ▷ Coefficients of $p_{n,k}$ acc. to (12a) and (12b)
end for

2.2.2. Inner–outer Krylov methods

In our approach [7], we modify (8) by the substitutions, $\bar{\mathcal{K}} := \mathcal{K} - \omega_1\mathcal{M}$, $\bar{\mathcal{C}} := \bar{\mathcal{K}}\mathcal{P}(\tau)^{-1}$, and solve the equivalent systems,

$$(\bar{\mathcal{C}} - \tilde{\eta}_k I)\mathbf{y}_k = \mathbf{b}, \quad \tilde{\eta}_k := \frac{\omega_k - \omega_1}{\omega_k - \tau}, \quad k = 1, \dots, n_\omega, \tag{13}$$

with the advantage that for $k = 1$ we solve the *base* system $\bar{\mathcal{C}}\mathbf{y}_1 = \mathbf{b}$ (unshifted case). A nested multi-shift Krylov algorithm consists in general of m_i inner iterations and m_o outer iterations. The nested FOM-FGMRES algorithm [7] is a combination of inner multi-shift FOM (Algorithm 3) with outer flexible multi-shift GMRES (Algorithm 4). In [7], we derive that if the inner method yields collinear residuals in the sense,

$$\mathbf{r}_j^{(k)} = \gamma_j^{(k)}\mathbf{r}_j, \quad \gamma_j^{(k)} \in \mathbb{C} \quad \text{for} \quad k = 1, \dots, n_\omega, \tag{14}$$

Algorithm 3. Inner multi-shift FOM for (13), cf. [30].

Set $\mathbf{r}_0 = \mathbf{b}$, $\mathbf{v}_1 = \mathbf{r}_0/\|\mathbf{r}_0\|$ ▷ Initialization ($\mathbf{x}_0 = \mathbf{0}$)
for $j = 1$ to m_i **do**
  Apply $\mathbf{w} = \bar{\mathcal{K}}(\mathcal{K} - \tau\mathcal{M})^{-1}\mathbf{v}_j$ ▷ Application $\bar{\mathcal{C}}$ (13)
  for $i = 1$ to j **do** ▷ Arnoldi method
    $h_{i,j} = \mathbf{w}^H\mathbf{v}_i$
    $\mathbf{w} = \mathbf{w} - h_{i,j}\mathbf{v}_i$
  end for
  Set $h_{j+1,j} = \|\mathbf{w}\|$ and $\mathbf{v}_{j+1} = \mathbf{w}/h_{j+1,j}$
end for
Set $H_{m_i} = [h_{i,j}]_{i=1,\dots,m_i}^{j=1,\dots,m_i}$ and $V_{m_i} = [\mathbf{v}_1, \dots, \mathbf{v}_{m_i}]$
for $k = 1$ to n_ω **do**
  Solve $\mathbb{C}^{m_i} \ni \mathbf{y}_k = (H_{m_i} - \tilde{\eta}_k I_{m_i})^{-1}(\|\mathbf{r}_0\|\mathbf{e}_1)$
  Compute $\gamma_k = \mathbf{y}_k(m_i)/\mathbf{y}_1(m_i)$ ▷ Collinearity [7]
  Compute $\mathbf{x}_k = V_{m_i}\mathbf{y}_k$
end for

Algorithm 4. Outer multi-shift FGMRES for (13), cf. [7].

Set $\mathbf{r}_0 = \mathbf{b}$, $\mathbf{v}_1 = \mathbf{r}_0/\|\mathbf{r}_0\|$ ▷ Initialization ($\mathbf{x}_0 = \mathbf{0}$)
for $j = 1$ to m_o **do**
  $[\mathbf{z}_j^{(k)}, \{\gamma_j^{(k)}\}_{k=1}^{n_\omega}] = \texttt{msFOM}(\bar{\mathcal{C}}, \{\tilde{\eta}_k\}_{k=1}^{n_\omega}, \mathbf{v}_j, m_i)$ ▷ The inner method, i.e. call of Algorithm 3
  Apply $\mathbf{w} = \bar{\mathcal{K}}(\mathcal{K} - \tau\mathcal{M})^{-1}\mathbf{z}_j^{(k=1)}$
  for $i = 1$ to j **do** ▷ Arnoldi method
    $h_{i,j} = \mathbf{w}^H\mathbf{v}_i$
    $\mathbf{w} = \mathbf{w} - h_{i,j}\mathbf{v}_i$
  end for
  Set $h_{j+1,j} = \|\mathbf{w}\|$ and $\mathbf{v}_{j+1} = \mathbf{w}/h_{j+1,j}$
end for
Set $\underline{H}_{m_o} = [h_{i,j}]_{i=1,\dots,m_o}^{j=1,\dots,m_o+1}$ and $Z_{m_o}^{(k)} = [\mathbf{z}_1^{(k)}, \dots, \mathbf{z}_{m_o}^{(k)}]$
for $k = 1$ to n_ω **do**
  Set $\Gamma_k := \operatorname{diag}(\gamma_1^{(k)}, \dots, \gamma_{m_o}^{(k)})$ ▷ Collinearity factors
  Set $\underline{H}_{m_o}^{(k)} := \left(\underline{H}_{m_o} - \underline{I}_{m_o}\right)\Gamma_k + \underline{I}_{m_o}$ ▷ See (15)
  Solve $\mathbb{C}^{m_o} \ni \bar{z}_k = \operatorname{argmin}_{\mathbf{z}} \|\underline{H}_{m_o}^{(k)}\mathbf{z} - \|\mathbf{r}_0\|\mathbf{e}_1\|$
  Compute $\mathbf{y}_k = Z_{m_o}^{(k)}\bar{z}_k$
end for

for $\mathbf{r}_j$ being the residual of the base system after m_i inner iterations, we can preserve shift-invariance in the outer method. The consecutive collinearity factors of the inner method then appear on a diagonal matrix Γ_k of a modified shifted Hessenberg matrix in the outer loop (see line 14 in Algorithm 4 and [7], respectively). More precisely, after m_o outer iterations, the solution to,

$$\bar{z}_k = \operatorname{argmin}_{\mathbf{z} \in \mathbb{C}^{m_o}} \left\| \left((\underline{H}_{m_o} - \underline{I}_{m_o})\Gamma_k + \underline{I}_{m_o} \right)\mathbf{z} - \|\mathbf{r}_0\|\mathbf{e}_1 \right\|_2, \tag{15}$$

with $\mathbf{y}_k = Z_{m_o}^{(k)}\bar{z}_k$, yields approximate solutions to (13) in the search spaces $Z_{m_o}^{(k)} \in \mathbb{C}^{2N \times m_o}$ that minimize the 2-norm of the residual of the kth shifted system, cf. [7]. In (15), the Hessenberg matrix $\underline{H}_{m_o}$ corresponds to the base system, and $\Gamma_k := \operatorname{diag}(\gamma_1^{(k)}, \dots, \gamma_{m_o}^{(k)})$ is constructed from the collinearity factors in (14). Note that multi-shift FOM (cf. Algorithm 3) yields collinear residuals by default [7,30].

3. Comparison study and convergence behavior

We focus our numerical experiments on linear systems (1) that stem from a finite element discretization [2,6] of the time-harmonic elastic wave equation [12]:

$$-\omega_k^2\rho\mathbf{u}_k - \nabla\cdot\sigma(\mathbf{u}_k) = \mathbf{s}, \quad \mathbf{x} \in \Omega \subset \mathbb{R}^{d=\{2,3\}}, \tag{16a}$$

$$i\omega_k\rho \quad B(c_p, c_s)\mathbf{u}_k + \sigma(\mathbf{u}_k)\hat{\mathbf{n}} = \mathbf{0}, \quad \mathbf{x} \in \partial\Omega_a, \tag{16b}$$

$$\sigma(\mathbf{u}_k)\hat{\mathbf{n}} = \mathbf{0}, \quad \mathbf{x} \in \partial\Omega_r. \tag{16c}$$

The Stress tensor in (16a) fulfills Hooke's law, $\sigma(\mathbf{u}_k) = \lambda(\mathbf{x})(\nabla\cdot\mathbf{u}_k \;\; I_d) + \mu(\mathbf{x})(\nabla\mathbf{u}_k + (\nabla\mathbf{u}_k)^T)$, and we consider Sommerfeld radiation boundary conditions on $\partial\Omega_a$ that model absorption, and a free-surface boundary condition on $\partial\Omega_r$ (reflection). A finite element discretization[1] with basis functions that are B-splines [11,Chapter 2] of degree $p \in \mathbb{N}_{>0}$ yields linear systems,

$$(K + i\omega_k C - \omega_k^2 M)\mathbf{u}_k = \mathbf{s}, \quad k = 1, \dots, n_\omega, \tag{17}$$

where $\mathbf{u}_k$ contains FEM coefficients of the kth displacement vector, and $\mathbf{s}$ models a time-harmonic source term. In the case of purely reflecting boundary conditions, i.e. $\partial\Omega_a = \emptyset$, we obtain $C = 0$; cf. [6]. The inhomogeneous set of parameters $\{\rho, c_p, c_s\}$ is described in Fig. 1a. In Fig. 1b, we prescribe material-air boundary conditions at the upper boundary only, and a point source at $(L_x/2, 0)^T$.

When comparing convergence behavior of the matrix equation approach (2) with the shifted system reformulation (3), we make use of the identity,

$$\left\|\mathbf{R}_m\right\|_F = \sqrt{\sum_{k=1}^{n_\omega} \left\|\mathbf{r}_m^{(k)}\right\|_2^2}, \quad \mathbf{R}_m := [\mathbf{r}_m^{(1)}, \dots, \mathbf{r}_m^{(n_\omega)}] \in \mathbb{C}^{N \times n_\omega},$$

where $\{\mathbf{r}_m^{(k)}\}_{k=1}^{n_\omega}$ are the columns of $\mathbf{R}_m$ and *not* the residuals of the shifted systems. Since this way the block residual in Frobenius norm naturally is larger than an individual residual norm in 2-norm, we use the maximum 2-norm of the residuals of (3) as a *fair* stopping criteria. All numerical examples presented in Section 3 have been implemented in `Python-3`, and executed on a computer with 4 CPUs `Intel I5` with 32 GB of RAM.

3.1. Convergence study in the presence of viscous damping

As a first numerical experiment we consider the case when viscous damping is added to (17) via the substitution $\omega_k \mapsto (1 - \epsilon i)\omega_k$ for $\epsilon > 0$. As we explain in Section 2.2.1, the spectral radius of the

[1] For the finite element discretization of (16a)–(16c) we use the Python package `nutils` [37].

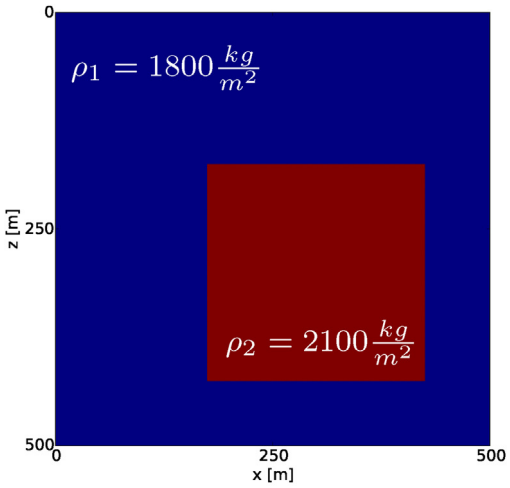

(a) Density distribution.

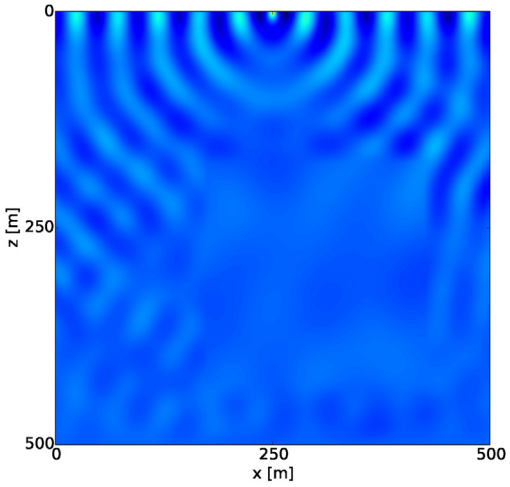

(b) $\Re(\mathbf{u}_z)$ at $f = 16\text{Hz}$, $C \neq 0$.

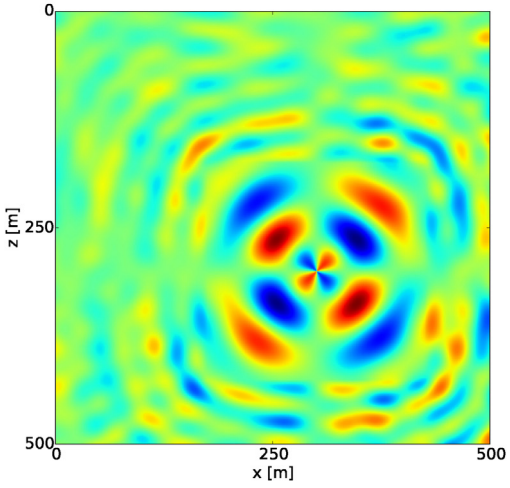

(c) $\Re(\mathbf{u}_x)$ at $f = 20\text{Hz}$, $C \equiv 0$.

Fig. 1. Set-up of the 2D numerical experiments: density distribution (top), and real part of the z-component of the displacement at $f = 16$ Hz (middle) and $f = 20$ Hz (bottom). The speed of pressure and shear waves are $c_p = \{2000, 3000\}\frac{m}{s}$ and $c_s = \{800, 1600\}\frac{m}{s}$, respectively, and the Lamé parameters $\{\lambda, \mu\}$ in Hooke's law are calculated accordingly. The case $C \equiv 0$ implies purely reflecting boundary conditions, and we apply viscous damping of 5%.

polynomial preconditioner (9) can be minimized as a result of the optimal seed frequency $\tau^*(\epsilon)$ derived in [8],

$$\tau^*(\epsilon) = \frac{2\omega_1\omega_{n_\omega}}{\omega_1 + \omega_{n_\omega}} - i\frac{\sqrt{[\epsilon^2(\omega_1 + \omega_{n_\omega})^2 + (\omega_{n_\omega} - \omega_1)^2]\omega_1\omega_{n_\omega}}}{\omega_1 + \omega_{n_\omega}}. \quad (18)$$

Table 1 demonstrates that an increase of the polynomial degree n reduces the number of iterations of Algorithm 2, cf. [16,39]. The best CPU time is obtained for $n = 3$ in (9).

Table 2 compares the performance of the three algorithms when viscous damping is present, i.e. $\epsilon > 0$; cf. Fig. 1c. Clearly, the shifted systems approaches outperform the matrix equation approach.

3.2. Suitability for wide frequency ranges

We next consider the undamped problem ($\epsilon = 0$) with Sommerfeld boundary conditions (see Fig. 1b) which is numerically more challenging. Here, we use $n = 0$ in Algorithm 2 because the spectral radius of the Neumann preconditioner is $R/|c| \equiv 1$, cf. [8,36]. The experiments in Tables 3 and 4 show that the matrix equation approach requires a large number of iterations, especially when the number of frequencies is increased. This is due to the fact that the union of the preconditioned spectra needs to be well approximated by the global GMRES method, cf. [8].

The equivalent, vectorized reformulation of the matrix Eq. (2),

$$\begin{bmatrix} (K + i\omega_1 C - \omega_1^2 M) & & \\ & \ddots & \\ & & (K + i\omega_{n_\omega} C - \omega_{n_\omega}^2 M) \end{bmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_{n_\omega} \end{pmatrix} = \begin{pmatrix} b \\ \vdots \\ b \end{pmatrix},$$

shows that the preconditioner (5b) acts on the block diagonals which demonstrates that the block Krylov subspace in Algorithm 1 needs to approximate the union of the spectra whereas in the shifted systems approach only one space is built due to shift-invariance. This drawback is partly overcome by applying appropriate rotations to the spectrum as we show in detail in [8,Section 5.2].

3.3. Inexact solves for the shift-and-invert preconditioner

In Section 2, we describe that the matrix equation approach (5a) and (5b) does not rely on shift-invariance and, thus, yields a less restrictive framework. In [32] the author makes use of this fact by applying recycling methods to a Sylvester equation. In Table 5 we exploit the use of an inexact LU factorization for the preconditioner (5b) in Algorithm 1. Therefore, we extend the test case in Fig. 1a to 3D by an expansion in y-direction, see Fig. 2a. Note that applying the shift-and-invert preconditioner inexactly is in general more time-efficient but the spectral analysis in [8] does not hold anymore which makes the rotation strategy described in the previous section more difficult. The measured CPU times indicate this trade-off between decomposition time and overall number of iterations. In practice, more advanced inexact preconditioners such as multigrid [26,27] or hierarchical matrix decompositions [3,6] are used for seismic applications.

4. Additive coarse grid correction for the 3D elastic preconditioner

We turn our attention to the shifted systems approach, i.e. Algorithm 2 or Algorithms 3 and 4, and consider numerical problems that stem from elasticity problems in three spatial dimensions, cf. Fig. 2. It has been noted in Section 2.2 that the shift-and-invert preconditioner (6), other than in the matrix equation approach in Section 3.3, needs to be applied to full accuracy due to the equivalence used in (7). In [6], we have proposed a 3D block-SSOR

Table 1
Performance of Algorithm 2 for the case $C=0$ and viscous damping parameter $\epsilon=0.05$. We consider a fixed frequency range of $n_\omega=5$ equally spaced frequencies in $f_k \in [8, 16]$ Hz, and $2\times200\times200$ dofs. The seed parameter τ is chosen according to (18), cf. [8].

$n=$	10	5	4	3	2	1	0
# Iterations	12	20	25	29	39	57	106
CPU time [s]	24.20	20.77	20.66	19.84	20.27	22.51	36.87

Table 2
Comparison of the three algorithms for the setup described in Table 1. The degree of the polynomial preconditioner is fixed at $n=3$. We report CPU time in seconds and in parenthesis the number of iterations until tol=1e-8 is reached.

Problem size	Frequency range	n_ω	Gl-GMRES	Poly-msGMRES	FOM-FGMRES
$2\times200\times200$	$\omega_k \in 2\pi[12, 16]$ Hz	5	29.3 (48)	12.65 (12)	12.63 (7·8)
$2\times200\times200$	$\omega_k \in 2\pi[10, 16]$ Hz	5	46.6 (75)	15.31 (19)	16.04 (12·8)
$2\times200\times200$	$\omega_k \in 2\pi[8, 16]$ Hz	5	79.9 (112)	19.80 (29)	19.90 (17·8)
$2\times200\times200$	$\omega_k \in 2\pi[12, 16]$ Hz	15	64.8 (47)	15.71 (12)	13.41 (7·8)
$2\times200\times200$	$\omega_k \in 2\pi[10, 16]$ Hz	15	115.9 (73)	18.37 (19)	16.86 (12·8)
$2\times200\times200$	$\omega_k \in 2\pi[8, 16]$ Hz	15	198.9 (109)	22.49 (29)	20.71 (17·8)

Table 3
Comparison for the undamped case, $\epsilon=0$, and increased frequency range at a fixed seed parameter $\tau=(0.7-0.3i)\omega_{\max}$, with $\omega_{\max}=2\pi\cdot8$ Hz in this table.

Problem size	Frequency range	n_ω	Gl-GMRES	Poly-msGMRES	FOM-FGMRES
$2\times100\times100$	$\omega_k \in 2\pi[7, 8]$ Hz	5	14.2 (111)	9.98 (96)	5.40 (20·8)
$2\times100\times100$	$\omega_k \in 2\pi[4, 8]$ Hz	5	16.3 (124)	10.81 (96)	5.55 (20·8)
$2\times100\times100$	$\omega_k \in 2\pi[1, 8]$ Hz	5	29.5 (193)	12.40 (106)	8.40 (20·11)
$2\times100\times100$	$\omega_k \in 2\pi[7, 8]$ Hz	15	42.6 (116)	11.42 (96)	5.86 (20·8)
$2\times100\times100$	$\omega_k \in 2\pi[4, 8]$ Hz	15	50.5 (127)	11.69 (96)	6.02 (20·8)
$2\times100\times100$	$\omega_k \in 2\pi[1, 8]$ Hz	15	148.9 (324)	13.68 (106)	8.97 (20·11)

Table 4
Setting as in Table 3 using quadratic B-splines ($p=2$).

Problem size	Frequency range	n_ω	Gl-GMRES	Poly-msGMRES	FOM-FGMRES
$2\times100\times100$	$\omega_k \in 2\pi[7, 8]$ Hz	15	86.9 (117)	18.74 (97)	13.86 (20·8)
$2\times100\times100$	$\omega_k \in 2\pi[4, 8]$ Hz	15	98.6 (130)	19.59 (97)	13.96 (20·8)
$2\times100\times100$	$\omega_k \in 2\pi[1, 8]$ Hz	15	267.4 (332)	28.87 (107)	18.95 (20·11)

Table 5
Inexact solves for the shift-and-invert preconditioner in Algorithm 1. We consider $n_\omega=10$ equally spaced frequencies with seed parameter $\tau=(0.7-0.3i)\omega_{\max}$. We use $\|\mathbf{R}_m\|_F<1e-8$ as stopping criteria.

Problem size	Frequency range	Preconditioner	Setup time	CPU time	# iter.
$3\times35\times35\times35$	$\omega_k \in 2\pi[1, 3]$ Hz	exact inverse	4533.9	5396.2	53
$3\times35\times35\times35$	$\omega_k \in 2\pi[1, 3]$ Hz	iLU(10.0)	332.9	2852.3	482
$3\times35\times35\times35$	$\omega_k \in 2\pi[1, 3]$ Hz	iLU(20.0)	559.2	2179.0	367
$3\times35\times35\times35$	$\omega_k \in 2\pi[1, 3]$ Hz	iLU(30.0)	1061.4	2129.8	197

preconditioner that efficiently solves a sequence of 2D problems. The numerical results in [6], however, yield unsatisfying convergence behavior even if the block-SSOR preconditioner is applied at the original frequency. We note that in Algorithm 2–4, the shift-and-invert preconditioner needs to be applied at a seed frequency τ that is chosen based on the spectral analysis performed in [8,Section 3], i.e. according to (18), and that contains damping.

Here, we present an extension of the block-SSOR preconditioner by an additive coarse grid correction (CGC). The application of an additive CGC is a standard tool used for Schwarz methods [13,25], i.e. in the situation where the numerical solution is computed separately on sub-domains and where the CGC yields an additional connectivity between the domains. In [13], an analysis of the CGC is presented based on the ratio h/H between the fine grid associated with grid size h and the coarse grid associated with grid size H. In [28,Section 13.4.2], coarse grid correction is interpreted as a two-grid multigrid method with no smoothing. The observation of [23,34,33] that a suitable choice for the deflation matrices Z and Z^T yields an additive CGC leads immediately to the applicability of additive CGC as a preconditioner. The comparison with multiplicative CGC in [24] concludes similar results for both variants.

We recall the splitting for the damped elastic operator $P(\tau^*)$ (5b) introduced in [6]

$$K+i\tau^*C-(\tau^*)^2M =: \underline{L}+\hat{S}+\bar{U}, \tag{19}$$

$$\text{with } \hat{S}=\text{blkdiag}(\hat{S}_1,\ldots,\hat{S}_{n_z}),$$

and with τ^* chosen according to (18). With appropriate splitting, the block entries $\{\hat{S}_i\}_{i=1}^{n_z}$ yield a sequence of 2D problems that can be solved efficiently using, for instance, level-2 MSSS techniques as in [6]. Based on the splitting (19), a block-SSOR preconditioner P_{SSOR} can be defined. If an additive coarse grid correction P_{CGC} is applied, the (damped) preconditioner reads,

$$P_h(\tau^*)^{-1} := P_{\text{SSOR}}^{-1}+P_{\text{CGC}} \tag{20a}$$

$$=\left[(\underline{L}\hat{S}^{-1}+I)\hat{S}(\hat{S}^{-1}\bar{U}+I)\right]^{-1}+PP_H(\tau^*)^{-1}R, \tag{20b}$$

with prolongation and restriction matrices P and R that realize a mapping between the fine and the coarse grid, respectively. If the

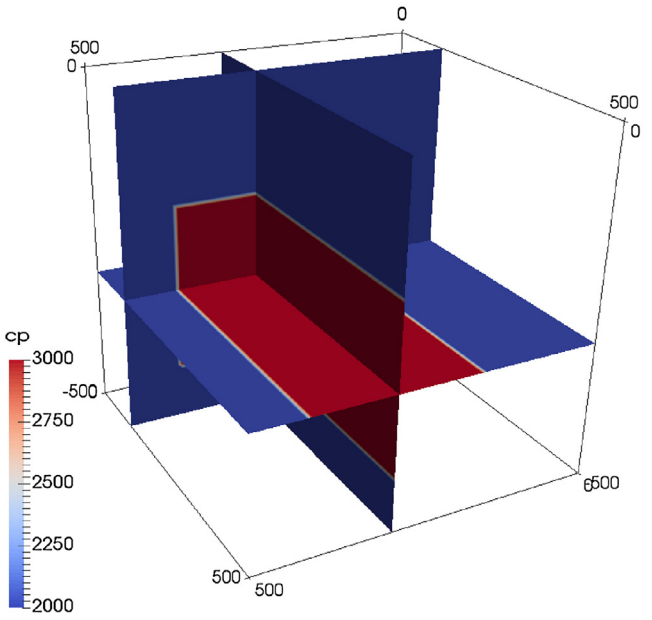

(a) Speed of P-waves in m/s.

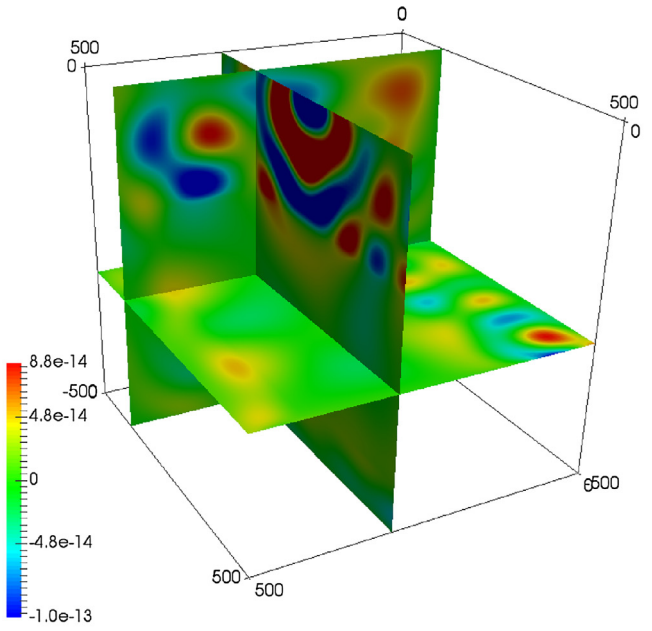

(b) $\Re(\mathbf{u}_x)$ at $f = 8\text{Hz}$, $C \neq 0$.

Fig. 2. Three-dimensional test case. Top: the underlying parameters $\{\rho, c_p, c_s\}$ are chosen according to Fig. 1a, and have been duplicated in y-direction. Bottom: x-component of the displacement vector at $f = 8$ Hz.

discretization size H for the coarse grid is much larger than h for the fine grid, the inverse in (20b) can be computed efficiently; for instance using a direct method. The coarse operator P_H here stems from a *direct* (finite element) discretization of the damped elastic wave equation on a coarse grid and, thus, replaces the so-called Galerkin operator $(RP_hP)^{-1}$ used in the multigrid literature [10,35]. A block Jacobi preconditioner has been combined with additive coarse grid correction in [38].

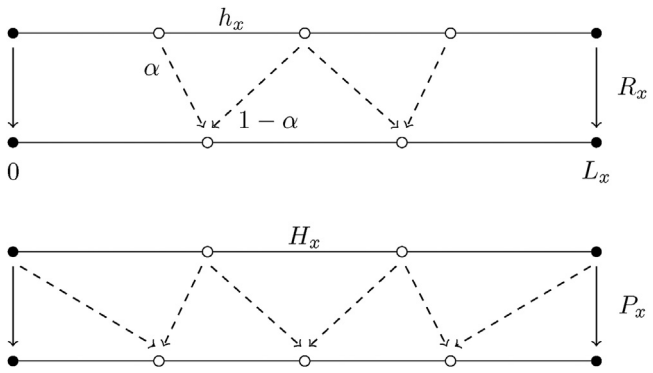

Fig. 3. Restriction and prolongation operators for a one-dimensional grid discretizing the interval $[0, L_x]$, exemplified for a fine grid of $n_x = 5$ and a coarse grid of $n_x^c = 4$ grid points. Dashed arrows indicate linear interpolation, and straight arrows visualize injection on the boundary. Note that here H_x is not a multiple of h_x.

Prolongation and restriction operators

The prolongation matrix P and restriction matrix R in (20b) realize a mapping between the fine and coarse grid. Here, we assume a three-dimensional grid ($d = 3$) and first-order B-splines for the FEM discretization (i.e. $p = 1$). Let the fine grid have $N = 3n_xn_yn_z$ grid points and the coarse grid $N^c = 3n_x^cn_y^cn_z^c$, with $N^c \ll N$ and, in general a grid spacing of $H \neq 2h$. In one spatial direction, the prolongation and restriction operators are rectangular matrices of the dimensions,

$$P_x \in \mathbb{R}^{n_x \times n_x^c} \quad \text{and} \quad R_x \in \mathbb{R}^{n_x^c \times n_x}.$$

In Fig. 3, we illustrate the definition of both operators in our implementation: since H_x is not necessarily a multiple of h_x, we suggest for interior grid points to use linear interpolation from nearest neighbors, and so-called *injection* at the boundary. Since the solution of the damped elastic wave equation is smooth, we expect linear interpolation to be sufficient. For the lexicographic numbering used in [37], the tensor product,

$$P = I_3 \otimes P_x \otimes P_y \otimes P_z \in \mathbb{R}^{N^c \times N},$$
$$R = I_3 \otimes R_x \otimes R_y \otimes R_z \in \mathbb{R}^{N \times N^c},$$

yields prolongation and restriction operators in 3D. This *grid-based* definition of restriction and prolongation is valid when first-order finite elements are considered. We refer to [20] for a more general framework of prolongation and restriction operators mapping between function spaces of higher-order B-splines.

4.1. An SSOR-MSSS preconditioner with additive coarse grid correction for damped 3D problems

The efficiency of an additive CGC preconditioner is demonstrated by means of numerical experiments performed for the damped three-dimensional elastic problem described in Fig. 2. In order to have an efficient algorithm for the shifted linear systems approach (using Algorithm 2 or Algorithms 3 and 4), we require the preconditioner (6) to have optimal computational complexity, that is $\mathcal{O}(n^3)$ if $n := n_x = n_y = n_z$. In order to achieve the required accuracy, we apply preconditioner (20a) to PGMRES. In Fig. 4, we demonstrate that the additive CGC (20a) leads to a constant number of PGMRES iterations if the (damped) seed frequency τ^* is chosen according to (18). The convergence results in Fig. 4, moreover, demonstrate grid-independence of the CGC since we double the maximum wave frequency in each experiment which in 3D implies an increase of the (fine) computational grid by a factor of 8.

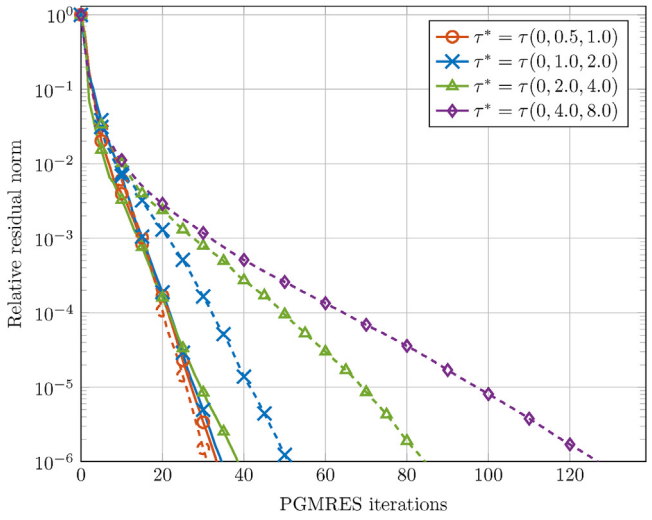

Fig. 4. Convergence of preconditioned GMRES applied to the damped 3D problem $(K + i\tau^* C - (\tau^*)^2 M)$ with preconditioner (20a). The grid size is increased proportional to the largest frequency $f_{\max}$, cf. Table 6. The 'seed' is chosen as $\tau^* = \tau(\epsilon, 2\pi f_{\min}, 2\pi f_{\max})$ according to (18). We compare convergence with (straight lines) and without (dashed lines) additive coarse grid correction in (20a).

Table 6
Iteration numbers and CPU times for PGMRES when the preconditioner (20a) is applied to the damped 3D problem. We report the effect of an additive CGC at the absence of viscous damping, i.e. $\epsilon = 0$.

$h_x = h_y = h_z$ Freq. range	40 m [0.5, 1] Hz	20 m [1, 2] Hz	10 m [2, 4] Hz	5 m [4, 8] Hz
n_{dofs} = $\mathcal{P}_H(\tau*)^{-1}$ at	5, 184 $H = 100$m	46, 875 $H = 50$m	375, 000 $H = 25$m	3, 000, 000 $H = 12.5$m
P_{SSOR}^{-1}	31 (0.43)	51 (6.2)	85 (87.0)	127 (1145.4)
$P_{\text{SSOR}}^{-1} + P_{\text{CGC}}$	34 (0.50)	35 (4.7)	39 (47.3)	–

Table 6 contains the CPU times corresponding to Fig. 4. The additive CGC has been implemented upon our Fortran 90 implementation used in [6], and a pre-computed *LU* decomposition using SuperLU [22] is used for the inversion at the coarse grid level. The experiments in Table 6 and Fig. 4 show a constant number of PGMRES iterations only when an additive CGC is applied. Since we double the number of grid points in every spatial direction, we observe a factor of about 8 in CPU time per iteration. A small increase of this optimal complexity is due to the fact that the *LU* factorization of the coarse grid operator needs to be applied. For the largest test case, the storage requirement of the SuperLU implementation exceeds the RAM memory. This can be overcome with a recursive strategy, i.e. a 3-grid cycle (or multigrid [26]). We repeat the experiment in the presence of viscous damping, $\epsilon > 0$, for the CGC. Again, the constant iteration numbers of PGMRES in Table 7 imply optimal computational complexity for applying the preconditioner in 3D.

4.2. Computational complexity study of the overall shifted algorithm

We next study the overall computational complexity of the shifted systems (3) approach for two-dimensional ($d = 2$) and three-dimensional ($d = 3$) elasticity problems at multiple frequencies. In order to develop an efficient multi-shift Krylov method, the computational costs per iteration are required to be of the order $\mathcal{O}(n^d)$ which is the order of a matrix–vector product. In Section 4.1, it has been shown that the shift-and-invert preconditioner (6) can be applied at optimal complexity in the 3D case. In order to analyze the performance of the overall multi-shift algorithms, we first

Table 7
Performance of PGMRES with SSOR and additive CGC preconditioner (20a) applied to the damped 3D problem with viscous damping $0 \leq \epsilon \ll 1$.

$h_x = h_y = h_z$ Freq. range	40 m [0.5, 1] Hz	20 m [1, 2] Hz	10 m [2, 4] Hz
$\epsilon = 0.5$	31 (0.43)	30 (4.1)	33 (42.5)
$\epsilon = 0.1$	34 (0.47)	35 (4.8)	38 (46.8)
$\epsilon = 0.05$	34 (0.46)	35 (4.7)	38 (46.1)
$\epsilon = 0.0$	34 (0.48)	35 (4.8)	39 (47.3)

Table 8
Performance of MSSS-preconditioned GMRES as developed in [6] applied to the 2D damped elastic wave equation. The semiseparable order [6,Definition 4] is bounded by $r^* = 5$ in all experiments.

$h_x = h_z$ Freq. range	10 m [2, 4] Hz	5 m [4, 8] Hz	2.5 m [8, 16] Hz	1.25 m [16, 32] Hz
$\epsilon = 0.5$	4 (0.02)	4 (0.09)	4 (0.37)	4 (1.5)
$\epsilon = 0.1$	5 (0.03)	5 (0.11)	4 (0.36)	6 (2.1)
$\epsilon = 0.05$	5 (0.03)	5 (0.11)	5 (0.45)	6 (2.1)
$\epsilon = 0.0$	5 (0.03)	5 (0.11)	5 (0.45)	6 (2.1)

Table 9
Preconditioned multi-shift GMRES (Algorithm 2 at $n = 0$) applied to $n_\omega = 10$ equidistantly spaced frequencies. The shift-and-invert preconditioner is applied as studied in Table 8.

$h_x = h_z$ Freq. range	10 m [2, 4] Hz	5 m [4, 8] Hz	2.5 m [8, 16] Hz	1.25 m [16, 32] Hz
$\epsilon = 0.5$	10 (0.62)	11 (2.8)	12 (12.8)	12 (49.4)
$\epsilon = 0.1$	20 (0.90)	28 (5.8)	33 (31.8)	37 (160.7)
$\epsilon = 0.05$	24 (1.02)	40 (8.0)	53 (47.8)	64 (264.9)
$\epsilon = 0.0$	31 (1.22)	77 (14.0)	163 (134.6)	358 (1391.2)

focus on the two-dimensional problem described in Fig. 1. For the shift-and-invert preconditioner, we use an approximate block-*LU* decomposition based on MSSS matrix computations described in detail in [6]. Table 8 shows that if the level-2 MSSS-preconditioner is applied to a damped problem at seed frequency $\tau^*(\epsilon = 0)$ (18), a constant off-diagonal rank (semiseparable order) of $r^* = 5$ is sufficient such that PGMRES convergences within few iterations. In particular, we measure a factor of 4 when the problem size is doubled in both spatial directions which is optimal.

We next apply preconditioned multi-shift GMRES (Algorithm 2 at $n = 0$) to multi-frequency problems in 2D up to a frequency of 32 Hz, cf. Table 9. As in the previous experiment, the preconditioner is applied at the (optimal) seed frequency τ^* (18). As a consequence of [8,Corollary 4.4], we compare convergence at frequency intervals such that the ratio $\rho := \omega_{\max}/\omega_{\min}$ is constant. If the viscous damping parameter ϵ is large, the convergence bound studied in [8] is a function of ρ only, and as a consequence we observe an almost constant number of iterations for the different experiments in Table 9. In particular, an upper bound for the multi-shift GMRES iteration number exists when ρ is fixed, $\epsilon > 0$, and a tolerance for the relative residuals is prescribed, cf. [8]. In CPU time, we observe an increase of a factor of 4 due to the increase in mesh size. In the absence of viscous damping, $\epsilon = 0$, the iteration number is proportional to the largest wave frequency [26,27] and we observe a factor of approximately 8 in CPU time in Table 9.

(Multi-shift) GMRES is a long-recurrence Krylov method, and the costs per iteration grow due to Arnoldi's orthogonalization method. Moreover, the storage requirement is proportional to the iteration number. A more memory-efficient multi-shift Krylov method is, therefore, the nested approach [7] presented in Algorithms 3 and 4. In Table 10, we show that in the case of $\epsilon = 0$ the same complexity $\mathcal{O}(n^3)$ can be obtained using short recurrences due to the inner-outer scheme.

Table 10
Comparison of multi-shift GMRES (Algorithm 2) with nested multi-shift FOM-FGMRES (Algorithms 3 and 4) for the test case without viscous damping, $\epsilon = 0$, in Table 9. The iteration numbers for the nested algorithm are reported as $m_o \times m_i$, where the inner method (Algorithm 3) is stopped if either the residual norm drops below `0.1`, or the maximum iteration number m_i is reached.

$h_x = h_z$ Freq. range	10 m [2, 4] Hz	5 m [4, 8] Hz	2.5 m [8, 16] Hz	1.25 m [16, 32] Hz
`msGMRES`	31 (`1.22`)	77 (`14.0`)	163 (`134.6`)	358 (`1391.2`)
`FOM-FGMRES`	3×15 (`2.1`)	6×15 (`19.0`)	7×30 (`179.7`)	16×30 (`1835.1`)

Table 11
Multi-shift GMRES with a block-SSOR and additive CGC preconditioner (20a) applied to the 3D problem shown in Fig. 2 at $n_\omega = 10$ equally spaced frequencies.

$h_x = h_y = h_z$ Freq. range	40 m [0.5, 1] Hz	20 m [1, 2] Hz	10 m [2, 4] Hz
$\epsilon = 0.5$	6 (`9.2`)	8 (`103.7`)	9 (`1201.2`)
$\epsilon = 0.1$	7 (`10.5`)	11 (`140.6`)	16 (`2020.8`)
$\epsilon = 0.05$	8 (`11.2`)	12 (`148.0`)	20 (`2376.3`)
$\epsilon = 0.0$	8 (`11.2`)	14 (`163.4`)	26 (`2889.4`)

Table 12
Comparison regarding memory requirements and costs-per-iteration when (1) has fixed problem size N, and n_ω distinct frequencies are present. Note that in all three algorithms, a single `MatVec` also requires a solve for the shift-and-invert preconditioner.

Algorithm	Leading memory requirement	# `MatVec`
`Gl-GMRES(m)`	$N \cdot n_\omega \cdot m$ for $\mathbf{V}_m$ (in Algorithm 1, line 10)	$n_\omega \cdot m$
`Poly-msGMRES(m,n)`	$2N \cdot m$ for V_m (in Algorithm 2, line 10)	$(n+1) \cdot m$
`FOM(`m_i`)-F GMRES(`m_o`)`	$2N \cdot n_\omega \cdot m_o$ for $Z_{m_o}^{(k)}$ (in Algorithm 4, line 11)	$m_i \cdot m_o$

Consider the extension of the previous test case to three-dimensions in Fig. 2. We have demonstrated in Section 4.1 that the shift-and-invert preconditioner can be applied at optimal computational complexity if an additive coarse grid correction is applied.

In Table 11, we repeat the numerical experiment for multi-shift GMRES for a three-dimensional test problem, and observe a similar behavior: in the presence of viscous damping the iteration number is (almost) constant and we obtain an algorithm in $\mathcal{O}(n^3)$. If $\epsilon = 0$, the iteration number doubles and we obtain $\mathcal{O}(n^4)$ complexity. In particular, this complexity is favorable to what we have observed in [6] when a block-SSOR preconditioner has been applied *directly* to the wave frequency of the original single-shift problem.

5. Conclusions

We have compared three GMRES-based algorithms for the simultaneous iterative solution of frequency-domain wave propagation problems at multiple frequencies that have the discretized form (1). The three approaches share that they require the application of a single shift-and-invert preconditioner at a so-called *seed* frequency. From our numerical experiments we draw the following conclusions:

- In the presence of viscous damping (experiments in Section 3.1) the optimal seed parameter derived in [8] implies a polynomial preconditioner which, depending on the degree n of the Neumann polynomial, leads to a significant reduction of the number of multi-shift GMRES iterations (Algorithm 2). Without viscous damping, however, the spectral radius of the polynomial preconditioner equals one and no improvement has been observed.
- The matrix equation approach (Algorithm 1) builds up a block Krylov space that needs to approximate the union of all preconditioned spectra. This leads to a much larger number of overall iterations, and a worse performance compared to the shifted systems approach especially when a wide range of frequencies is considered, and the ratio $\omega_{n_\omega}^2/\omega_1^2$ is much larger than $\omega_{n_\omega}/\omega_1$, cf. [8,Corollary 4.4]. Because of the less restrictive framework, however, the shift-and-invert preconditioner (5b) can be applied *inexactly* which leads to improvements especially for 3D problems. This has been demonstrated using an incomplete LU factorization in Section 3.3. Note that also the SSOR-CGC preconditioner (20a) described in Section 4 can be applied in the matrix equation framework. Moreover, the benefits of efficient block matrix–vector products when multiple sources are considered is demonstrated in [6].
- For a wide frequency range (experiments in Section 3.2) we observe that the nested Algorithms 3 and 4 outperforms the considered alternatives with respect to measured CPU time. This is due to shorter loops in the respective Arnoldi iterations. From the summary in Table 12 we note that the memory requirements for the search spaces $Z_{m_o}^{(k)}$ in the flexible outer Krylov method can be limited by the nested framework when m_o is small compared to m_i.
- In the numerical experiments in Sections 4.1 and 4.2, the computational complexity of the multi-shift approach with a shift-and-invert preconditioner in two and three spatial dimensions has been studied. Therefore, multiple frequencies ω_k in the intervals $\mathcal{I}_i = 2\pi[2^i, 2^{i+1}]$ Hz, $i = 0, 1, 2, \ldots$, have been considered. The largest frequency induces a computational grid with, hence, $n \sim 2^i/h$ grid points in each spatial direction. The parameter of the shift-and-invert preconditioner $\tau_i^*(\epsilon)$ is chosen based on the spectral analysis in [8]. Due to the damping in $\tau_i^*(\epsilon)$ according to (18), the 2D and 3D preconditioner can be applied in optimal complexitiy, i.e. in $\mathcal{O}(n^d)$. The convergence study of preconditioned multi-shift GMRES and preconditioned nested multi-shift FOM-FGMRES has shown computational complexity $\mathcal{O}(n^d)$ if the viscous damping parameter $\epsilon > 0$, and $\mathcal{O}(n^{d+1})$ if $\epsilon = 0$, where $d \in \{2, 3\}$ is the problem dimension. Both are independent of the number n_ω of frequencies within $\mathcal{I}_i$ due to shift-invariance in both Krylov methods.

Code availability

The source code of the implementations used to compute the presented numerical results can be obtained from: https://github.com/ManuelMBaumann/freqdom_compare and is authored by: Manuel Baumann. `doi:10.5281/zenodo.495915`. Please contact Manuel Baumann for licensing information.

Acknowledgments

We want to thank René-Édouard Plessix for helpful scientific discussions, and the organizers of the International Conference on Computational Science 2017 for providing the opportunity to publish this extended version of our conference proceedings [9]. Moreover, we wish to thank all referees for carefully reading our manuscript and for providing valuable comments on an earlier version of this manuscript. Shell Global Solutions International B.V. is gratefully acknowledged for financial support of the first author.

References

[1] M.I. Ahmad, D.B. Szyld, M.B. van Gijzen, Preconditioned multishift BiCG for $\mathcal{H}_2$-optimal model reduction, SIAM J. Matrix Anal. Appl. 38 (2017) 401–424.
[2] T. Airaksinen, A. Pennanen, J. Toivanen, A damping preconditioner for time-harmonic wave equations in fluid and elastic material, J. Comput. Phys. 228 (5) (2009) 1466–1479.

[3] P. Amestoy, C. Ashcraft, O. Boiteau, A. Buttari, J.-Y. L'Excellent, C. Weisbecker, Improving multifrontal methods by means of block low-rank representations, SIAM J. Sci. Comput. 37 (2015) A1451–A1474.
[4] R. Astudillo, M.B. van Gijzen, Induced dimension reduction method for solving linear matrix equations, Proc. Comput. Sci. 80 (2016) 222–232.
[5] T. Bakhos, P.K. Kitanidis, S. Ladenheim, A.K. Saibaba, D.B. Szyld, Multipreconditioned GMRES for shifted systems, SIAM J. Sci. Comput. 39 (2017) S222–S247.
[6] M. Baumann, R. Astudillo, Y. Qiu, E.Y.M. Ang, M.B. van Gijzen, R.-É. Plessix, An MSSS-preconditioned matrix equation approach for the time-harmonic elastic wave equation at multiple frequencies, Comput. Geosci. 22 (1) (2017) 43–61.
[7] M. Baumann, M.B. van Gijzen, Nested Krylov methods for shifted linear systems, SIAM J. Sci. Comput. 37 (5) (2015) S90–S112.
[8] M. Baumann, M.B. van Gijzen, An Efficient Two-Level Preconditioner for Multi-Frequency Wave Propagation Problems. Technical Report 17-03, Delft University of Technology, 2017.
[9] M. Baumann, M.B. van Gijzen, Efficient iterative methods for multi-frequency wave propagation problems: a comparison study, Proc. Comput. Sci. 108 (2017) 645–654.
[10] W.L. Briggs, V. Emden Henson, S.F. McCormick, A Multigrid Tutorial, 2nd ed., Society for Industrial and Applied Mathematics, 2000.
[11] J.A. Cottrell, T.J.R. Hughes, Y. Bazilevs, Isogeometric Analysis. Towards Integration of CAD and FEA, John Wiley & Son, Ltd, 2009.
[12] A.T. De Hoop, Handbook of Radiation and Scattering of Waves, Academic Press, London, United Kingdom, 1995.
[13] O. Dubois, M.J. Gander, S. Loisel, A. St-Cyr, D.B. Szyld, The optimized Schwarz method with a coarse grid correction, SIAM J. Sci. Comput. 34 (1) (2012) A421–A458.
[14] Y.A. Erlangga, C.W. Oosterlee, C. Vuik, A novel multigrid based preconditioner for heterogeneous Helmholtz problems, SIAM J. Sci. Comput. 27 (2006) 1471–1492.
[15] Y.A. Erlangga, C. Vuik, C.W. Oosterlee, Comparison of multigrid and incomplete LU shifted-Laplace preconditioners for the inhomogeneous Helmholtz equation, Appl. Numer. Math. 56 (2006) 648–666.
[16] B. Fischer, R.W. Freund, On adaptive weighted polynomial preconditioning for Hermitian positive definite matrices, SIAM J. Sci. Comput. 15 (2) (1994) 408–426.
[17] A. Frommer, U. Glässner, Restarted GMRES for shifted linear systems, SIAM J. Sci. Comput. 19 (1) (1998) 15–26.
[18] A. Frommer, K. Lund, D.B. Szyld, Block Krylov subspace methods for functions of matrices, Electron. Trans. Numer. Anal. 47 (2017) 100–126.
[19] M.J. Gander, H.A. Zhang, Class of Iterative Solvers for the Helmholtz Equation: Factorizations, Sweeping Preconditioners, Source Transfer, Single Layer Potentials, Polarized Traces, and Optimized Schwarz Methods. Technical Report, 2017 arXiv:1610.02270v2.
[20] C. Hofreither, B. Jüttler, G. Kiss, W. Zulehner, Multigrid methods for isogeometric analysis with THB-splines, Comput. Methods Appl. Mech. Eng. 308 (2016) 96–112.
[21] K. Jbilou, A. Messaoudi, H. Sadok, Global FOM and GMRES algorithms for matrix equations, Appl. Numer. Math. 31 (1999) 49–63.
[22] X.S. Li, An overview of SuperLU: algorithms, implementation, and user interface, ACM Trans. Math. Softw. 31 (2005) 302–325.
[23] R. Nabben, C. Vuik, A comparison of deflation and coarse grid correction applied to porous media flow, SIAM J. Numer. Anal. 42 (2004) 1631–1647.
[24] Y. Notay, A. Napov, Further comparison of additive and multiplicative coarse grid correction, Appl. Numer. Math. 65 (2013) 53–62.
[25] L.F. Pavarino, E. Zampieri, Overlapping Schwarz and spectral element methods for linear elasticity and elastic waves, J. Sci. Comput. 27 (1-3) (2006) 51–73.
[26] C.D. Riyanti, Y.A. Erlangga, R.-E. Plessix, W.A. Mulder, C. Vuik, C. Osterlee, A new iterative solver for the time-harmonic wave equation, Geophysics 71 (5) (2006) E57–E63.
[27] G. Rizzuti, W.A. Mulder, Multigrid-based 'shifted-Laplacian' preconditioning for the time-harmonic elastic wave equation, J. Comput. Phys. 317 (2016) 47–65.
[28] Y. Saad, Iterative Methods for Sparse Linear Systems: Second Edition, Society for Industrial and Applied Mathematics, 2003.
[29] A. Saibaba, T. Bakhos, P. Kitanidis, A flexible Krylov solver for shifted systems with application to oscillatory hydraulic tomography, SIAM J. Sci. Comput. 35 (2013) 3001–3023.
[30] V. Simoncini, Restarted full orthogonalization method for shifted linear systems, BIT Numer. Math. 43 (2003) 459–466.
[31] V. Simoncini, F. Perotti, On the numerical solution of $(\lambda^2 A + \lambda B + C)x = b$ and application to structural dynamics, SIAM J. Sci. Comput. 23 (2002) 1875–1897.
[32] K.M. Soodhalter, Block Krylov subspace recycling for shifted systems with unrelated right-hand sides, SIAM J. Sci. Comput. 38 (1) (2016) A302–A324.
[33] J.M. Tang, S.P. MacLachlan, R. Nabben, C. Vuik, A comparison of two-level preconditioners based on multigrid and deflation, SIAM J. Matrix Anal. Appl. 31 (4) (2010) 1715–1739.
[34] J.M. Tang, R. Nabben, C. Vuik, Y.A. Erlangga, Comparison of two-level preconditioners derived from deflation, domain decomposition and multigrid methods, J. Sci. Comput. 39 (3) (2009) 340–370.
[35] U. Trottenberg, C.W. Oosterlee, A. Schüller, Multigrid, Academic Press, 2001.
[36] 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 (5) (2007) 1942–1958.
[37] G. van Zwieten, C. Verhoosel, J. van Zwieten, T. van Opstal, W. Hoitinga, Nutils v2.0. Zenodo, 2016.
[38] C. Vuik, J. Frank, A parallel block preconditioner accelerated by coarse grid correction, in: M. Bubak, H. Afsarmanesh, B. Hertzberger, R. Williams (Eds.), High Performance Computing and Networking: 8th International Conference, HPCN Europe 2000, Springer Berlin Heidelberg, Amsterdam, 2000, pp. 99–108.
[39] G. Wu, Y.-C. Wang, X.-Q. Jin, A preconditioned and shifted GMRES algorithm for the Pagerank problem with multiple damping factors, SIAM J. Sci. Comput. 34 (2012) 2558–2575.

Manuel Baumann received his Ph.D. in applied mathematics from Delft University of Technology in 2018. Prior to that, he studied(18) mathematics at Technical University of Berlin (B.Sc.) and computer science at Royal Institute of Technology in Stockholm and at Delft University of Technology (M.Sc., double degree program). His research interests include numerical linear algebra, computational geophysics and model-order reduction for optimal control. He has published several journal papers in these fields including *SIAM Journal on Scientific Computing* and *Springer Computational Geosciences*. More information: www.manuelbaumann.de.

Martin van Gijzen received his Ph.D. degree in applied mathematics from the Delft University of Technology in the Netherlands, in 1994. He subsequently worked as a postdoc at Utrecht University, as a project leader at Dutch Organisation for Applied Scientific Research, and as a senior scientist at CERFACS, France. He presently holds a position as associate professor in numerical analysis at the Delft University of Technology. Van Gijzen is well-known for his contributions in numerical linear algebra, in particular on Krylov methods and preconditioners. He has co-authored more than 70 refereed journal and proceedings papers.