DELFT UNIVERSITY OF TECHNOLOGY

REPORT 17-03

AN EFFICIENT TWO-LEVEL PRECONDITIONER FOR MULTI-FREQUENCY WAVE PROPAGATION PROBLEMS

M. BAUMANN AND M.B. VAN GIJZEN

ISSN 1389-6520

Reports of the Delft Institute of Applied Mathematics

Delft 2017

Copyright © 2017 by Delft Institute of Applied Mathematics, Delft, The Netherlands.

No part of the Journal may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission from Delft Institute of Applied Mathematics, Delft University of Technology, The Netherlands.



Available online at www.sciencedirect.com



Journal of Applied Numerical Mathematics (2018) 1-18

Journal of Applied Numerical Mathematics

An Efficient Two-Level Preconditioner for Multi-Frequency Wave Propagation Problems

Manuel Baumann^a, Martin B. van Gijzen^a

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

Abstract

We consider wave propagation problems that are modeled in the frequency-domain, and that need to be solved simultaneously for multiple frequencies within a fixed range. For this, a single shift-and-invert preconditioner at a so-called *seed* frequency is applied. The choice of the seed is crucial for the performance of preconditioned multi-shift GMRES and is closely related to the parameter choice for the Complex Shifted Laplace preconditioner. Based on a classical GMRES convergence bound, we present an optimal seed parameter that purely depends on the original frequency range. The new insight is exploited in a two-level preconditioning strategy: A shifted Neumann preconditioner with minimized spectral radius is additionally applied to multi-shift GMRES. Moreover, we present a reformulation of the multi-shift problem to a matrix equation solved with, for instance, global GMRES. Here, our analysis allows for rotation of the spectrum of the linear operator. Numerical experiments for the time-harmonic visco-elastic wave equation demonstrate the performance of the new preconditioners.

Keywords: multi-shift GMRES, shift-and-invert preconditioner, preconditioner design, visco-elastic wave equation, frequency-domain formulation

1. Introduction

We consider the efficient iterative solution of a sequence of $n_s > 1$ shifted systems of the form,

$$(\mathcal{K} - s_k \mathcal{M})\mathbf{x}_k = \mathbf{b}, \quad \text{for } k = 1, ..., n_s, \tag{1}$$

where the matrices \mathcal{K} and \mathcal{M} depend on the specific problem discretization, and $\{s_k\}_{k=1}^{n_s}$ is a sequence of (possibly complex) shifts. Problems of the form (1) arise, for instance, in oscillatory hydraulic tomography [1] and lattice quantum chromodynamic [2]. Moreover, the extension of acoustic Helmholtz problems [3, 4, 5, 6] to a multi-frequency setting results in the framework (1). The focus of the present work, however, lies on situations where the discretization matrices \mathcal{K} , \mathcal{M} stem from a discretization of the time-harmonic elastic wave equation [7]. Depending on the specific choice of boundary conditions the structure of the matrices varies, and the shifts s_k are either equal to the (angular) wave frequencies [9]. For both situations, we will consider viscous damping by substituting $s_k \mapsto (1 - \epsilon i) s_k$, where $\epsilon > 0$ is the damping parameter and $i \equiv \sqrt{-1}$, cf. [4, 8, 9, 10].

Email address: m.m.baumann@tudelft.nl (Manuel Baumann)

URL: http://www.manuelbaumann.de (Manuel Baumann)

Throughout this document we put emphasis on the case when the set of shifts $\{s_1, ..., s_{n_s}\}$ in (1) is distinct, and multiple frequencies are considered, i.e. $n_s > 1$. Without loss of generality, we assume the frequencies to be ordered, and, in particular, $s_1 = \min_k \{s_k\} =: s_{\min}$ and $s_{n_s} = \max_k \{s_k\} =: s_{\max}$ are the extreme frequencies. A large area of application where the fast solution of (1) at multiple frequencies is required is the so-called Full-Waveform Inversion modeled in frequency-domain; cf. [11, 12, 13, 14].

If the matrices in (1) are large and sparse, Krylov subspace methods are the common choice for the iterative numerical solution of (1). When re-formulating problem (1) to problems with shifted identity, Krylov methods can be particularly efficient, and variants of almost all popular Krylov methods have been derived for this type of problems (such as GMRES(k) [2], FOM(k) [15], BiCGstab(ℓ) [16] and IDR(s) [8, 17] among others). It is, however, difficult to apply a preconditioner and, at the same time, preserve the shifted structure: Most recently, polynomial preconditioners [18], flexible preconditioners [1], nested methods [8], and multi-preconditioned methods [19] have successfully been developed. Alternative approaches to solve sequences of linear systems such as (1) are the reformulation as a matrix equation [20], and the usage of information of previous solves called *recycling* [21, 22].

In most cases, a single preconditioner of the form,

$$\mathcal{P}(\tau) := (\mathcal{K} - \tau \mathcal{M}), \text{ with seed shift } \tau \in \mathbb{C},$$
 (2)

is applied where the choice of $\tau \in \mathbb{C}$ for a given set $\{s_1, ..., s_{n_s}\}$ is crucial for the convergence behavior of the overall algorithm, as has been pointed out in [1, 19, 20]. The present paper addresses the following:

- 1. We present an *optimal* choice for the seed parameter τ in (2) when a single shift-and-invert preconditioner is applied to (1). Our proposed choice is based on spectral analysis and the minimization of a classical GMRES convergence bound that also holds in the multi-shift framework.
- 2. Once a preconditioner of the form (2) is applied, the spectra are known to be bounded by circles which gives rise to the efficient application of a shifted Neumann preconditioner [18] as a second-level preconditioner. Our choice for τ minimizes its spectral radius.
- 3. The spectral analysis of the multi-shift framework is exploited for an equivalent matrix equation formulation of (1) studied in [20]. A simple post-rotation of the block spectrum yields a second-level preconditioner for the matrix equation and significantly speeds up convergence of global GMRES [23].

We point out that the analysis of an efficient seed parameter τ is fundamentally different from the single-frequency case studied in [4] since there is no trivial solution that needs to be excluded from the optimization. Moreover, note that we can not include directly frequency-dependent Sommerfeld boundary conditions in the preconditioner (2), as recommended in [24]. When (2) is applied inexactly using a multi-grid algorithm [25] or deflation [26], the choice of τ is usually combined with damping such that the multi-grid solver works well, cf. [6, 27] for an analysis in the Helmholtz case. We do not consider this aspect in this paper but note that the MSSS preconditioner developed in [20] allows to apply the inverse of (2) fast, even for large frequencies. We conclude with numerical examples obtained from a finite element discretization of the time-harmonic visco-elastic wave equation at multiple wave frequencies.

2. The time-harmonic elastic wave equation at multiple frequencies

The aim of this work is the efficient iterative solution of the elastic wave equation in a frequency-domain formulation. The displacement vector $\mathbf{u}(t, x)$ at time *t* and with spatial component *x* satisfies the elastic wave equation,

$$\rho \ddot{\mathbf{u}} = \nabla \cdot \sigma(\mathbf{u}) + \mathbf{s}, \quad \text{in } (0, T] \times \Omega, \quad \Omega \subset \mathbb{R}^d, \quad d \in \{2, 3\}, \tag{3}$$

with inhomogeneous material density $\rho = \rho(x)$, stress tensor σ , and source term s, cf. [7]. We consider the following set of boundary conditions,

$$\rho \dot{\mathbf{u}} = \sigma(\mathbf{u})\mathbf{n} \text{ for } x \in \partial \Omega_a \quad \text{and} \quad \sigma(\mathbf{u})\mathbf{n} = \mathbf{0} \text{ for } x \in \partial \Omega_r, \tag{4}$$

where the condition on $\partial \Omega_r$ models reflection of waves, and the Sommerfeld radiation condition on $\partial \Omega_a$ is one way to model absorption. For the time-harmonic ansatz $\mathbf{u}(t, x) = \hat{\mathbf{u}}(x)e^{-i\omega t}$ substituted into (3)-(4) we obtain:

$$-\omega^2 \rho \hat{\mathbf{u}} - \nabla \cdot \sigma(\hat{\mathbf{u}}) = \hat{\mathbf{s}}, \quad \text{in } \Omega, \tag{5a}$$

$$i\omega\rho B(c_p, c_s) \,\hat{\mathbf{u}} + \sigma(\hat{\mathbf{u}})\mathbf{n} = \mathbf{0}, \quad \text{on } \partial\Omega_a,$$
(5b)

$$\sigma(\hat{\mathbf{u}})\mathbf{n} = \mathbf{0}, \quad \text{on } \partial\Omega_r. \tag{5c}$$

Note that, in the frequency-domain formulation (5a)-(5c), the Sommerfeld condition yields a term that is proportional to the frequency. For the definition of $B(c_p, c_s)$ in (5b) we refer to [10, 20]. In [20] a detailed derivation of a finite element discretization of (5a)-(5c) is presented that yields linear systems of the form,

$$(K + i\omega_k C - \omega_k^2 M)\hat{\mathbf{u}}_k = \hat{\mathbf{s}}, \quad k = 1, ..., n_s,$$
(6)

with *K* being a stiffness matrix, *M* being a mass matrix and *C* includes Sommerfeld boundary conditions. The angular frequencies ω_k appear quadratic in (6). Therefore, we apply a linearization [28] that results in block-systems of doubled dimensions,

$$\begin{pmatrix} \begin{bmatrix} iC & K \\ I & 0 \end{bmatrix} - \omega_k \begin{bmatrix} M & 0 \\ 0 & I \end{bmatrix} \end{pmatrix} \begin{bmatrix} \omega_k \hat{\mathbf{u}}_k \\ \hat{\mathbf{u}}_k \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{s}} \\ \mathbf{0} \end{bmatrix}, \quad k = 1, ..., n_s.$$
(7)

Let $\epsilon > 0$. We formally add *viscous damping* to (6) by introducing the set of complex frequencies $\hat{\omega}_k := (1 - \epsilon i)\omega_k$. The damped problem then reads, $(K + i\hat{\omega}_k C - \hat{\omega}_k^2 M)\hat{\mathbf{u}}_k = \mathbf{b}$, for $k = 1, ..., n_s$, where the ansatz $\mathbf{u}(t, x) = \hat{\mathbf{u}}(x)e^{-i\hat{\omega}t} = \hat{\mathbf{u}}(x)e^{-i\omega t}e^{-\epsilon\omega t}$ now includes a damping term. When damping is added to the problem, spectral properties change and, in particular, the bounding circles that we describe in Section 4 do no longer touch the origin. The systems (7) are of the form (1), see Problem 2.1.

We also consider different types of absorbing boundary conditions on $\partial \Omega_a$: The case of purely non-mixed boundary conditions in (4) trivially yields $C \equiv 0$ in (6) and we, again, obtain a problem of the form (1). Absorption can also be modeled by introducing a *sponge* layer [9] or using perfectly matched layers (PML) [29]. The frequency-independent PML derived in [30] yields a term $C(\omega_0)$. In general, however, including PML boundary conditions yields a nonlinear term $C(\omega)$ that is not considered here. We summarize the above derivations by the following two problem statements.

Problem 2.1. Consider the discretized time-harmonic elastic wave equation (5a)-(5c) with Sommerfeld radiation boundary conditions [10, 20] on $\partial \Omega_a \neq \emptyset$, and multiple angular frequencies $\omega_k \equiv 2\pi f_k$,

$$(K + i\omega_k C - \omega_k^2 M) \hat{\mathbf{u}}_k = \hat{\mathbf{s}}, \quad k = 1, ..., n_s, \quad \{K, C, M\} \in \mathbb{C}^{N \times N},$$

where the matrices K and C are symmetric positive semi-definite, and M is symmetric positive definite. The linearization (7) yields a shifted problem of the form (1) with, in particular, block matrices,

$$\mathcal{K} := \begin{bmatrix} iC & K \\ I & 0 \end{bmatrix} \in \mathbb{C}^{2N \times 2N}, \quad \mathcal{M} := \begin{bmatrix} M & 0 \\ 0 & I \end{bmatrix} \in \mathbb{C}^{2N \times 2N}, \quad shifts \ s_k := \omega_k, \tag{8}$$

and right-hand side vector $\mathbf{b} := [\mathbf{\hat{s}}, \mathbf{0}]^{\mathsf{T}}$.

Remark 2.2. In reformulation (7), dimensions are doubled compared to the original problem size in (6). For the preconditioner (2), however, the following decomposition holds,

$$\mathcal{P}(\tau)^{-1} = (\mathcal{K} - \tau \mathcal{M})^{-1} = \left(\begin{bmatrix} iC & K \\ I & 0 \end{bmatrix} - \tau \begin{bmatrix} M & 0 \\ 0 & I \end{bmatrix} \right)^{-1} = \begin{bmatrix} I & \tau I \\ 0 & I \end{bmatrix} \begin{bmatrix} I & 0 \\ 0 & (K + i\tau C - \tau^2 M)^{-1} \end{bmatrix} \begin{bmatrix} 0 & I \\ I & -iC + \tau M \end{bmatrix},$$
(9)

where the (inexact) inversion is required only at the original problem size.

Problem 2.3. Consider the discretized time-harmonic elastic wave equation (5a)-(5c) with frequencyindependent PML [30] or 'sponge layer' [9] boundary conditions replacing (5b), and angular frequencies ω_k ,

$$(K - \omega_k^2 M) \mathbf{\hat{u}}_k = \mathbf{\hat{s}}, \quad k = 1, ..., n_s$$

where K is symmetric positive semi-definite, and M is symmetric positive definite. The problem is, trivially, of the form (1) for $\mathcal{K} := K$, $\mathcal{M} := M$ and shifts $s_k := \omega_k^2$ equal to the squared angular frequencies.

Remark 2.4. For $\epsilon > 0$, the damped problem corresponding to Problem 2.1 and Problem 2.3, respectively, is given by the substitution,

$$(\mathcal{K} - \hat{s}_k \mathcal{M})\hat{\mathbf{x}}_k = \mathbf{b}, \quad where \ \hat{s}_k := (1 - \epsilon i)s_k \ for \ k = 1, ..., n_s.$$
(10)

In our following notation, we indicate quantities related to the damped problem with a hat.

3. The shift-and-invert preconditioner for multi-shift GMRES

In this section we briefly review the multi-shift GMRES method introduced in [2]. Throughout this paper we always consider the case when multi-shift GMRES is right-preconditioned by a preconditioner of the form (2). When applying a (scaled) shift-and-invert preconditioner as a right preconditioner to systems (1) the resulting preconditioned systems are shifted linear systems. Moreover, the shift parameter (sometimes called *seed frequency*) gives some freedom. Recall that we consider a sequence of problems of the form (1),

$$(\mathcal{K} - s_k \mathcal{M})\mathbf{x}_k = \mathbf{b}, \quad \text{for } k = 1, ..., n_s,$$

where the matrices \mathcal{K} , \mathcal{M} are defined in Problem 2.1 or Problem 2.3, respectively. For $\tau \in \mathbb{C} \setminus \{0\}$, we define the shiftand-invert preconditioner $\mathcal{P}(\tau) = (\mathcal{K} - \tau \mathcal{M})$ as in (2). Right preconditioning of (1) with the scaled preconditioner $\mathcal{P}_k := 1/(1 - \eta_k)(\mathcal{K} - \tau \mathcal{M}) = 1/(1 - \eta_k)\mathcal{P}(\tau)$ yields,

$$(\mathcal{K} - s_k \mathcal{M}) \mathcal{P}_{\nu}^{-1} \mathbf{y}_k = \mathbf{b} \quad \Leftrightarrow \quad (\mathcal{K} (\mathcal{K} - \tau \mathcal{M})^{-1} - \eta_k I) \mathbf{y}_k = \mathbf{b}, \tag{11}$$

where $\eta_k := s_k/(s_k - \tau)$. Note that the latter is a (preconditioned) shifted linear system with (possibly complex) shifts η_k and system matrix $\mathcal{KP}(\tau)^{-1}$. Note further that the back-substitution $\mathbf{x}_k = \mathcal{P}_k^{-1} \mathbf{y}_k = (1 - \eta_k)(\mathcal{K} - \tau \mathcal{M})^{-1} \mathbf{y}_k$ can be computed efficiently for $k = 1, ..., n_s$. A similar equivalence as (11) that yields shifted systems with base matrix $\mathcal{M}(\mathcal{K} - \tau \mathcal{M})^{-1}$ is used in [1].

Remark 3.1. It is well-known that Krylov subspaces are shift-invariant, i.e. for $\mathcal{A} := \mathcal{K}(\mathcal{K} - \tau \mathcal{M})^{-1}$ it holds,

$$\mathcal{K}_m(\mathcal{A}, \mathbf{b}) \equiv span\{\mathbf{b}, \mathcal{A}\mathbf{b}, ..., \mathcal{A}^{m-1}\mathbf{b}\} = \mathcal{K}_m(\mathcal{A} - \eta I, \mathbf{b}) \quad \forall \eta \in \mathbb{C}, \forall m \in \mathbb{N}.$$

As a consequence, the shifted Arnoldi relation holds,

$$(\mathcal{A} - \eta I)V_m = V_{m+1}(\underline{H}_m - \eta \underline{I}),$$

where the columns of V_m are an orthonormal basis of $\mathcal{K}_m(\mathcal{A}, \mathbf{b})$ that are computed by the Arnoldi method only once for all shifted systems, cf. Algorithm 1.

In [4] the authors analyze spectral properties of the shifted Laplace preconditioner in the single-frequency case, $n_s = 1$, and exploit their analysis within preconditioned GMRES [31]. One of the results of [4] is the fact that the preconditioned spectrum (11) lies within circles of radius *R* and center *c*. Both are, in the single frequency case, a function of s_1 and τ (denoted by z_1 and z_2 in [4], respectively). Moreover, the authors of [4] show that in the absence of viscous damping the circles touch the origin, i.e. R = |c|. We state the following convergence bound for the GMRES residual norm that is in the absence of damping (compare Remark 2.4) of little practical use.

Theorem 3.2 (Classical convergence bound for GMRES, [15]). Let the eigenvalues of the preconditioned matrix be enclosed by a circle with radius *R* and center *c*. Then the GMRES-residual norm after *j* iterations $\|\mathbf{r}^{(j)}\|$ satisfies,

$$\frac{\left\|\mathbf{r}^{(j)}\right\|}{\left\|\mathbf{r}^{(0)}\right\|} \le c_2(X) \left(\frac{R}{|c|}\right)^j, \quad where \ \mathbf{r}^{(0)} = \mathbf{b} \ if \ \mathbf{x}_0 \equiv \mathbf{0}$$

and where X is the matrix of eigenvectors, and $c_2(X)$ denotes its condition number in the 2-norm.

Algorithm 1 Multi-shift GMRES with right preconditioning for (11), cf. [2]

1: Set $\mathbf{r}^{(0)} = \mathbf{b}, \ \beta = \|\mathbf{r}^{(0)}\|, \ \mathbf{v}_1 = \mathbf{r}^{(0)}/\beta$ ▶ Initialization with zero initial guess 2: **for** j = 1 to *m* **do** Apply $\mathbf{w} = (\mathcal{K} - \tau \mathcal{M})^{-1} \mathbf{v}_i$ ▶ Apply preconditioner (2) (cf. Section 4) 3: Compute $\mathbf{w} = \mathcal{K}\mathbf{w}$ 4: ▶ Arnoldi method 5: for *i* = 1 to *j* do $h_{i,j} = \mathbf{w}^{\mathsf{H}}\mathbf{v}_i$ 6: $\mathbf{w} = \mathbf{w} - h_{i,j}\mathbf{v}_i$ 7: end for 8: Set $h_{j+1,j} = ||\mathbf{w}||$ and $\mathbf{v}_{j+1} = \mathbf{w}/h_{j+1,j}$ Set $\underline{\mathbf{H}}_m = [h_{i,j}]_{i=1,\dots,m+1}^{j=1,\dots,m}$ and $V_m = [\mathbf{v}_1, \dots, \mathbf{v}_m]$ 9: 10: ▷ Orthogonal basis of $\mathcal{K}_m(\mathcal{A}, \mathbf{b}), \mathcal{A} := \mathcal{K}(\mathcal{K} - \tau \mathcal{M})^{-1}$ 11: end for 12: **for** k = 1 to n_s **do** Solve $\mathbf{y}_k = \operatorname{argmin}_{\mathbf{v}} \| \beta \mathbf{e}_1 - (\underline{\mathbf{H}}_m - \eta_k \underline{\mathbf{I}}) \mathbf{y} \|$ ▶ Solve shifted Hessenberg systems 13: Compute $\mathbf{x}_k = (1 - \eta_k)(\mathcal{K} - \tau \mathcal{M})^{-1} V_m \mathbf{y}_k$ ▶ Back-substitution 14: 15: end for

The convergence bound for GMRES described in Theorem 3.2 can be extended to the preconditioned multi-shift variant presented in Algorithm 1 in a straight-forward way. That is because in multi-shift GMRES optimality for the shifted residuals holds for all individual systems.

Corollary 3.3 (Convergence bound for multi-shift GMRES, [2, 15]). An extension of the bound described in Theorem 3.2 to the (preconditioned) multi-shift GMRES-residual norms computed by Algorithm 1 is given by,

$$\frac{\left\|\mathbf{r}_{k}^{(j)}\right\|}{\left\|\mathbf{r}^{(0)}\right\|} \le c_{2}(X) \left(\frac{R_{k}}{|c_{k}|}\right)^{j}, \quad k = 1, ..., n_{s}, \quad j \le m,$$
(12)

where the spectrum of the k-th shifted system after preconditioning is assumed to be enclosed by a circle of radius R_k and center point c_k , respectively.

The following section gives detailed explanations on the suitable choice of the bounding circles (c_k, R_k) in terms of the seed frequency τ . In particular, we will derive explicit formulas for these quantities, and make use of the fact that $|c_k| > R_k$ when viscous damping is added such that the bound in Corollary 3.3 can be exploited.

4. Spectral analysis and optimal seed shift parameter τ^*

...

We describe the main result of this paper: The efficiency of the preconditioner in (11) highly depends on the choice of the seed parameter $\tau \in \mathbb{C}$. The following lemma provides insight how to choose this parameter such that the bound of Corollary 3.3 is minimized. The result yields an explicit formula for τ in terms of the considered frequency range $[s_{\min}, s_{\max}]$, and the damping parameter $\epsilon \ge 0$ as introduced in Remark 2.4.

Lemma 4.1 (Optimal seed frequency for preconditioned shifted GMRES). Let $\{s_k\}_{k=1}^{n_s} \subseteq [s_{min}, s_{max}] \subset \mathbb{R}$. Consider the sequence of problems $(\mathcal{K} - s_k \mathcal{M})\mathbf{x}_k = \mathbf{b}$ with a right preconditioner $\mathcal{P}_k := 1/(1 - \eta_k)(\mathcal{K} - \tau \mathcal{M})$,

$$(\mathcal{K} - s_k \mathcal{M}) \mathcal{P}_k^{-1} \mathbf{y}_k = \mathbf{b}, \quad \mathbf{x}_k = \mathcal{P}_k^{-1} \mathbf{y}_k, \quad for \ k = 1, ..., n_s.$$
(13)

For $\eta_k = s_k/(s_k - \tau)$, problem (13) is equivalent to,

$$(\mathcal{K}(\mathcal{K}-\tau\mathcal{M})^{-1}-\eta_k I)\mathbf{y}_k=\mathbf{b}, \quad k=1,...,n_s,$$
(14)

where $\mathcal{P}(\tau) = (\mathcal{K} - \tau \mathcal{M})$ is the shift-and-invert preconditioner at seed frequency τ , and the matrices \mathcal{K} and \mathcal{M} and shifts $\{s_k\}_{k=1}^{n_s}$ are defined in Problem 2.1 and Problem 2.3, respectively. The following statements give guidance on choosing $\tau \in \mathbb{C}$ in an optimal sense:

- (i) For λ from the spectrum of \mathcal{KM}^{-1} , $\lambda \in \Lambda[\mathcal{KM}^{-1}]$, it holds $\mathfrak{I}(\lambda) \geq 0$.
- (ii) Let $\tau = \Re(\tau) + i\Im(\tau)$. The preconditioned spectra of (14) are enclosed by circles of radii R_k and center points c_k ,

$$c_{k} = \left(\frac{1}{2} - \frac{s_{k}(s_{k} - \Re(\tau))}{(s_{k} - \Re(\tau))^{2} + \Im(\tau)^{2}}, \frac{\Re(\tau)}{2\Im(\tau)} - \frac{s_{k}\Im(\tau)}{(s_{k} - \Re(\tau))^{2} + \Im(\tau)^{2}}\right), \quad R_{k} = \frac{1}{2}\sqrt{1 + \left(\frac{\Re(\tau)}{\Im(\tau)}\right)^{2}} =: R(\tau)$$

The preconditioned spectra of (14) with viscous damping $\hat{s}_k := (1 - \epsilon i)s_k$ are enclosed by circles of radii $\hat{R}_k = R(\tau)$ and center points \hat{c}_k ,

$$\hat{c}_k = \left(\frac{1}{2} - \frac{(1+\epsilon^2)s_k^2 + (\epsilon\mathfrak{I}(\tau) - \mathfrak{R}(\tau))s_k}{(s_k - \mathfrak{R}(\tau))^2 + (\epsilon s_k + \mathfrak{I}(\tau))^2}, \frac{\mathfrak{R}(\tau)}{2\mathfrak{I}(\tau)} - \frac{(\mathfrak{I}(\tau) + \epsilon\mathfrak{R}(\tau))s_k}{(s_k - \mathfrak{R}(\tau))^2 + (\epsilon s_k + \mathfrak{I}(\tau))^2}\right).$$

(iii) The set of points $\{\hat{c}_k\}_{k=1}^{n_s} \subset \mathbb{C}$ described in statement (ii) lie on a circle with center \underline{c} and radius \underline{R} given by,

$$\underline{c} = \left(0, \frac{\epsilon |\tau|^2}{2\Im(\tau)(\Im(\tau) + \epsilon \Re(\tau))}\right), \quad \underline{R} = \sqrt{\frac{|\tau|^2(\epsilon^2 + 1)}{4(\Im(\tau) + \epsilon \Re(\tau))^2}}$$

In the undamped case, $\epsilon = 0$, this center is equal to the origin and, therefore, $\underline{R} = R(\tau)$.

(iv) Consider the preconditioner $\mathcal{P}(\tau^*) = (\mathcal{K} - \tau^* \mathcal{M})$ in (14). An optimal seed frequency τ^* for Algorithm 1 that minimizes the GMRES-bound in Corollary 3.3 is given by,

$$\tau^*(\epsilon) = \underset{\tau \in \mathbb{C}}{\operatorname{argmin}} \max_{k=1,\dots,n_s} \left(\frac{R(\tau)}{|\hat{c}_k|} \right) = \frac{2s_{\min}s_{\max}}{s_{\min} + s_{\max}} - i \frac{\sqrt{\left[\epsilon^2 (s_{\min} + s_{\max})^2 + (s_{\max} - s_{\min})^2\right] s_{\min}s_{\max}}}{s_{\min} + s_{\max}}, \quad (15)$$

where $s_{min} := \min_k \{s_k\}$ and $s_{max} := \max_k \{s_k\}$.

Proof. We prove this lemma 'step-by-step'.

(i) For \mathcal{K}, \mathcal{M} as in Problem 2.3 the statement trivially holds since $\lambda \in \mathbb{R}$. With the assumption on K, C and M in Problem 2.1 we know that all eigenvalues μ_k of the quadratic eigenvalue problem,

$$(K + \mu_k C + \mu_k^2 M)v_k = 0, \quad k = 1, ..., 2N,$$

are stable, i.e. $\Re(\mu_k) \leq 0$, cf. [Table 1.1][32]. Eigenvalues of \mathcal{KM}^{-1} are then given by $\lambda_k := -i\mu_k$, and the corresponding eigenvectors are $\mathbf{v}_k := \mathcal{M}[-i\mu_k v_k, v_k]^{\mathsf{T}}$:

$$\begin{bmatrix} iC & K \\ I & 0 \end{bmatrix} \begin{bmatrix} M^{-1} & 0 \\ 0 & I \end{bmatrix} \mathbf{v}_k = \lambda_k \mathbf{v}_k \iff \begin{bmatrix} iC & K \\ I & 0 \end{bmatrix} \begin{bmatrix} -i\mu_k v_k \\ v_k \end{bmatrix} = -i\mu_k \begin{bmatrix} M & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} -i\mu_k v_k \\ v_k \end{bmatrix} \iff \begin{bmatrix} \mu_k C v_k + K v_k \\ -i\mu_k v_k \end{bmatrix} = \begin{bmatrix} -\mu_k^2 M v_k \\ -i\mu_k v_k \end{bmatrix}.$$

The definition of λ_k together with $\Re(\mu_k) \le 0$ imply that $\Im(\lambda_k) \ge 0$ for Problem 2.1.

(ii) Consider the system matrix of (14),

$$(\mathcal{K}(\mathcal{K}-\tau\mathcal{M})^{-1}-\eta_k I)=\mathcal{K}\mathcal{M}^{-1}(\mathcal{K}\mathcal{M}^{-1}-\tau I)^{-1}-\eta_k I, \quad \text{with } \eta_k=\frac{s_k}{s_k-\tau}.$$

The latter is a Möbius transformation with complex shift, hence the spectrum satisfies the mapping,

$$\boldsymbol{\Lambda}[\mathcal{K}\mathcal{M}^{-1}] \ni \lambda \mapsto \frac{\lambda}{\lambda - \tau} - \frac{s_k}{s_k - \tau}.$$

Since $\mathfrak{I}(\lambda) \ge 0$, it is well-known [4] that for $s_k \equiv 0$ the Möbius transformation maps the spectrum within a circle of radius $R = \left|\frac{\tau}{\tau-\overline{\tau}}\right|$ and center $c_0 = \frac{-\overline{\tau}}{\tau-\overline{\tau}}$. In the shifted case, it holds:

$$\hat{R}_{k} \equiv R(\tau) = \left| \frac{\tau}{\tau - \bar{\tau}} \right| = \frac{1}{2} \sqrt{1 + \left(\frac{\Re(\tau)}{\Im(\tau)}\right)^{2}},$$

$$\hat{c}_{k} = \frac{-\bar{\tau}}{\tau - \bar{\tau}} - \frac{\hat{s}_{k}}{\hat{s}_{k} - \tau} = \frac{-\Re(\tau) + i\Im(\tau)}{2i\Im(\tau)} - \frac{s_{k} - i\epsilon s_{k}}{(s_{k} - \Re(\tau)) - i(\epsilon s_{k} + \Im(\tau))}$$

$$= \left(\frac{1}{2} - \frac{(1 + \epsilon^{2})s_{k}^{2} + (\epsilon\Im(\tau) - \Re(\tau))s_{k}}{(s_{k} - \Re(\tau))^{2} + (\epsilon s_{k} + \Im(\tau))^{2}} \right) + i \left(\frac{\Re(\tau)}{2\Im(\tau)} - \frac{(\Im(\tau) + \epsilon\Re(\tau))s_{k}}{(s_{k} - \Re(\tau))^{2} + (\epsilon s_{k} + \Im(\tau))^{2}} \right),$$
(16)

where we again write $\tau = \Re(\tau) + i\Im(\tau)$. The case $\epsilon = 0$ yields the corresponding result for c_k in the absence of viscous damping. Note that the radii (16) are independent of k.

(iii) We prove this fact by first constructing a center point \underline{c} . Therefore, we use two points \hat{c}_k that are opposite to each other with real part zero,

$$\Re(\hat{c}_k) = 0 \quad \Leftrightarrow \quad s_k^{1/2} = \pm |\tau| (\epsilon^2 + 1)^{-1},$$

where we note that negative frequencies are not considered. Substituting $s_k^{1/2}$ into the imaginary part of (17) and computing the middle point yields,

$$\underline{\mathbf{c}} = \left(0, \frac{\epsilon |\tau|^2}{2\mathfrak{I}(\tau)(\mathfrak{I}(\tau) + \epsilon \mathfrak{R}(\tau))}\right).$$
(18)

We use Maple to show that every point \hat{c}_k has a constant distance from <u>c</u>. This distance is the squared radius <u>R</u>,

$$\underline{\mathbf{R}}^2 = \left\| \hat{c}_k - \underline{\mathbf{c}} \right\|_2^2 = \frac{|\tau|^2 (\epsilon^2 + 1)}{4(\Im(\tau) + \epsilon \Re(\tau))^2}, \quad \text{independent of } s_k.$$
(19)

(iv) In part (iii) of this proof, we have shown that the center points \hat{c}_k (17) of the preconditioned spectra lie on a circle with center \underline{c} (18) and radius \underline{R} (19). Therefore, an alternative parametrization of the distance to the origin $|\hat{c}_k|$ is given by,

$$|\hat{c}_k|^2 = \underline{\mathbf{R}}^2 + \Im(\underline{\mathbf{c}})^2 + 2\Im(\underline{\mathbf{c}})\underline{\mathbf{R}}\sin(\varphi_k) = \underline{\mathbf{R}}^2 - \Im(\underline{\mathbf{c}})^2 + 2\Im(\underline{\mathbf{c}})\Im(\hat{c}_k).$$

where the imaginary part of \hat{c}_k is given explicitly by (17), and φ_k is the corresponding phase angle, cf. Figure 1 (right). The expression for the GMRES bound in Corollary 3.3 can, hence, be simplified,

$$\tau^{*}(\epsilon) = \underset{\tau \in \mathbb{C}}{\operatorname{argmin}} \max_{k=1,\dots,n_{s}} \left(\frac{R(\tau)}{|\hat{c}_{k}|} \right) = \underset{\tau \in \mathbb{C}}{\operatorname{argmin}} \max_{k=1,\dots,n_{s}} \left(\frac{R(\tau)^{2}}{|\hat{c}_{k}|^{2}} \right) = \underset{\tau \in \mathbb{C}}{\operatorname{argmin}} \max_{k=1,\dots,n_{s}} \left(\frac{R^{2}}{\underline{R}^{2} + \mathfrak{I}(\underline{c})^{2} + 2\mathfrak{I}(\underline{c})\underline{R}\sin(\varphi_{k})} \right)$$

$$\stackrel{(*)}{=} \underset{\tau \in \mathbb{C}}{\operatorname{argmin}} \max_{k \in \{1,n_{s}\}} \left(\frac{R^{2}}{\underline{R}^{2} - \mathfrak{I}(\underline{c})^{2} + 2\mathfrak{I}(\underline{c})\mathfrak{I}(\hat{c}_{k})} \right) \stackrel{(**)}{=} \left\{ \tau \in \mathbb{C} \left| \mathfrak{I}(\hat{c}_{1}) = \mathfrak{I}(\hat{c}_{n_{s}}) \right\},$$

$$(20)$$

where in step (*) we use that $\sin(\varphi_k)$ obtains its minimum at the boundary (the case $\sin(\varphi_k) = -1$ has been excluded in part (iii)), and in step (**) we use that the minimum of the maximum of two functions occurs when the two functions are equal. Setting the imaginary parts equal (20) yields:

We next express $\tau^* = s^* e^{i\varphi^*}$ in polar coordinates, with length given by $s^* = \sqrt{s_{\min}s_{\max}(1+\epsilon^2)}$. Expressing the objective function $\mathcal{J} := R^2 [\underline{\mathbf{R}}^2 - \mathfrak{I}(\underline{\mathbf{c}})^2 + 2\mathfrak{I}(\underline{\mathbf{c}})\mathfrak{I}(\hat{c}_k)]^{-1}$ in terms of (s, φ) , we can use Maple¹ to solve $\frac{\partial \mathcal{I}(s^*,\varphi)}{\partial \varphi} = 0$ for the unique minimum,

$$\varphi^* = \varphi^*(s_{\min}, s_{\max}, \epsilon) = \arctan\left(-\sqrt{\frac{\epsilon^2(s_{\min} + s_{\max})^2 + (s_{\max} - s_{\min})^2}{4s_{\min}s_{\max}}}\right).$$
(21)

The conversion to Cartesian coordinates completes the proof.

¹We added corresponding Maple [v 18.02] scripts to our public repository [33].

Remark 4.2. Note that for $\epsilon = 0$, the ratio (15) equals to 1, cf. [4]. In this limit case, the optimal seed frequency yields,

$$\tau^*(0) = \frac{2s_{min}s_{max}}{s_{min} + s_{max}} - i\frac{|s_{max} - s_{min}|\sqrt{s_{min}s_{max}}}{s_{min} + s_{max}} = \sqrt{s_{min}s_{max}}\exp\left(i\arctan\left(-\frac{|s_{max} - s_{min}|}{2\sqrt{s_{min}s_{max}}}\right)\right),\tag{22}$$

which is a function of the geometric mean of the extreme frequencies, and their distance.

Remark 4.3. For $s_{min} = s_{max} \equiv s$, we get $\tau^*(\epsilon) = s\sqrt{1+\epsilon^2}e^{i\arctan(-\epsilon)} = (1-\epsilon i)s = \hat{s}$, and $\tau^*(0) = s$.

We illustrate the results of Lemma 4.1 in Figure 1: The left figure demonstrates the optimality of τ^* as stated in (15). Moreover, we plot the angle along which we have optimized in (21) as a dashed line. In the right figure, we show the corresponding preconditioned spectrum with bounding circles for a surrogate problem. This distribution of the circles corresponds to the case when τ minimizes the bound (15). Since the radii of all preconditioned spectra have the same magnitude, we see that the two extreme frequencies are expected to converge slowest because the respective distances to the origin is smallest which yields a worst case for the bound in Corollary 3.3.



Figure 1: Left: Convergence bound and optimal seed frequency τ^* according to (15). Along the dashed line the function is differentiable. Right: Preconditioned spectra and surrounding circles for positive damping $\epsilon = 0.7$ added to Problem 2.1. Here, we use $n_s = 20$ frequencies equally spaced within the interval [1,9]Hz. The imaginary parts that belong to the extreme frequencies are equal as imposed in (20), but not equal to zero.

As a result of Lemma 4.1, we see that the GMRES-bound in Corollary 3.3 and the location of the optimal seed frequency (15) are explicit functions of the damping parameter ϵ and the extreme frequencies $[s_{\min}, s_{\max}]$. This is both illustrated in Figure 2. The optimization of the seed parameter is obtained based on the damped problem. Because of continuity in Figure 2 (left), we note that τ^* smoothly depends on the damping factor ϵ which motivates the choice for the seed parameter in Remark 4.2 in the limit case $\epsilon \rightarrow 0$.

Lemma 4.1 does not give information about the actual value of the multi-shift GMRES bound (12) other than $R_k/|c_k| \equiv 1$ when $\epsilon = 0$. In the following corollary we show that when τ^* is chosen according to (15), the bound $R_k/|c_k|$ evaluated at τ^* is a function of the damping parameter $\epsilon > 0$ and the ratio $\rho := s_{\text{max}}/s_{\text{min}}$ only.

Corollary 4.4. Let $\epsilon > 0$, and $\tau^* = \tau^*(\epsilon, s_{min}, s_{max})$ as in (15) for a frequency interval $[s_{min}, s_{max}]$. Then there exists a function $f(\epsilon, s_{max}/s_{min})$ such that,

$$\frac{R_k(\tau^*)}{|c_k(\epsilon,\tau^*)|} = f(\epsilon,\rho), \quad \text{where } \rho := s_{max}/s_{min}, \tag{23}$$



Figure 2: Left: Imaginary part of τ^* scaled by s_{max} for different intervals $[s_{min}, s_{max}]$ and varying damping parameter ϵ . Right: GMRES-bound (12) at τ^* as a function of the ratio s_{max}/s_{min} and ϵ . Note that in (15) the real part of τ^* is independent of ϵ and is, therefore, not plotted here.

i.e. the bound in (12) depends only on the damping parameter ϵ and the ratio of the interval boundaries ρ . The quantities R_k and c_k are,

$$R_k(\tau^*) = R(\tau^*) = \frac{1}{2} \sqrt{1 + \left(\frac{\Re(\tau^*)}{\Im(\tau^*)}\right)^2}, \quad c_k(\epsilon, \tau^*) \stackrel{\epsilon > 0}{=} \hat{c}_k(\tau^*) = \begin{pmatrix} \frac{1}{2} - \frac{(1+\epsilon^2)s_k^2 + (\epsilon\Im(\tau^*) - \Re(\tau^*))s_k}{(s_k - \Re(\tau^*))^2 + (\epsilon s_k + \Im(\tau^*))^2} \\ \frac{\Re(\tau^*)}{(s_k - \Re(\tau^*))^2 + (\epsilon s_k + \Im(\tau))^2} \end{pmatrix} \in \mathbb{C},$$

according to Lemma 4.1(ii).

Proof. We show that

$$\frac{R(\tau^*(\epsilon, s_{\min}, s_{\max}))}{|c_k(\epsilon, \tau^*(\epsilon, s_{\min}, s_{\max}))|} = \frac{R(\tau^*(\epsilon, d \cdot s_{\min}, d \cdot s_{\max}))}{|c_k(\epsilon, \tau^*(\epsilon, d \cdot s_{\min}, d \cdot s_{\max}))|}, \quad k \in \{1, n_s\}$$

for any scalar d. First, note that

$$\tau^*(\epsilon, s, t) = \frac{2st}{s+t} - i \frac{\sqrt{[\epsilon^2(s+t)^2 + (t-s)^2] st}}{s+t},$$

is a homogeneous function of degree 1 with respect to the second and third argument, i.e., $\tau^*(\epsilon, ds_{\min}, ds_{\max}) = d \cdot \tau^*(\epsilon, s_{\min}, s_{\max})$. Therefore, the real and imaginary part of τ^* scale with *d* in the same way. This implies,

$$R(d \tau^*) = \frac{1}{2} \sqrt{1 + \left(\frac{d \mathfrak{R}(\tau^*)}{d \mathfrak{I}(\tau^*)}\right)^2} = R(\tau^*),$$

and, moreover,

$$\begin{aligned} \Re(c_1(\epsilon, d\tau^*)) &= \frac{1}{2} - \frac{d^2(1+\epsilon^2)s_1^2 + d^2(\epsilon\mathfrak{I}(\tau^*) - \mathfrak{R}(\tau^*))s_1}{d^2(s_1 - \mathfrak{R}(\tau^*))^2 + d^2(\epsilon s_1 + \mathfrak{I}(\tau^*))^2} = \Re(c_1(\epsilon, \tau^*)), \\ \mathfrak{I}(c_1(\epsilon, d\tau^*)) &= \frac{d\mathfrak{R}(\tau^*)}{2d\mathfrak{I}(\tau^*)} - \frac{d^2(\mathfrak{I}(\tau^*) + \epsilon\mathfrak{R}(\tau^*))s_1}{d^2(s_1 - \mathfrak{R}(\tau^*))^2 + d^2(\epsilon s_1 + \mathfrak{I}(\tau^*))^2} = \mathfrak{I}(c_1(\epsilon, \tau^*)), \end{aligned}$$

and, in the same way, $c_{n_s}(\epsilon, d\tau^*) = c_{n_s}(\epsilon, \tau^*)$, where we associate s_{\min} with s_1 (k = 1), and s_{\max} with s_{n_s} ($k = n_s$), cf. (20) in Lemma 4.1(iv). Thus, s_1 and s_{n_s} scale with d as the interval boundaries do.

9



Figure 3: Number of iterations *j* such that the relative residual norm is bounded by {1e-4, 1e-6, 1e-8}. The residual \mathbf{r}_k is associated with the (angular) wave frequency $s_k \in [s_{\min}, s_{\max}]$ with fixed ratio $\rho = s_{\max}/s_{\min} = 2$, and $c_2(X)$ set to 1 in (12).

The bound (23) is plotted as a function of the damping parameter ϵ and the ration $\rho = s_{\text{max}}/s_{\text{min}}$ in Figure 2 (right). If ρ is kept constant, the GMRES bound implies an a priori known maximum iteration number for a fixed relative residual tolerance, cf. Figure 3. The numerical experiment 6.4 exploits the insight of Corollary 4.4 when splitting a frequency interval [s_{min} , s_{max}] into subintervals in a balanced way.

5. Areas of application within a two-level preconditioning framework

We present two examples in which the insight of the previous section is exploited for the design of efficient two-level preconditioners: In Section 5.1 we make use of the spectral bounds of Lemma 4.1 in order to choose the parameter of the shifted Neumann polynomial preconditioner [18] such that its spectral radius is minimal. In Section 5.2 we present a reformulation of the multi-frequency problem (1) as a matrix equation, cf. [20, 34]. Here, the convergence behavior depends on the union of the spectra of all considered frequencies and, hence, a suitable rotation of the bounding circles yields an efficient second-level preconditioner for global GMRES [23].

5.1. Shifted Neumann preconditioning techniques

For a set of frequencies $\{s_k\}_{k=1}^{n_s}$ and viscous damping parameter $\epsilon > 0$, the preconditioned shifted problems (14) at an optimal seed frequency τ^* given explicitly by (15) read,

$$(\mathcal{A} - \eta_k I)\mathbf{y}_k = \mathbf{b}, \quad \text{with } \mathcal{A} := \mathcal{K}(\mathcal{K} - \tau^* \mathcal{M})^{-1} \text{ and } \eta_k := \frac{\hat{s}_k}{\hat{s}_k - \tau^*} = \frac{(1 - \epsilon i)s_k}{(1 - \epsilon i)s_k - \tau^*}.$$

For $s_k = 0$, the spectrum of the matrix \mathcal{A} is bounded by a circle of radius R and center c_0 as stated in part (ii) of Lemma 4.1,

$$R = \frac{1}{2} \sqrt{1 + \left(\frac{\Re(\tau^*)}{\Im(\tau^*)}\right)^2} \quad \text{and} \quad c_0 = \frac{\overline{\tau}^*}{\overline{\tau}^* - \tau^*} = \left(\frac{1}{2}, \frac{\Re(\tau^*)}{2\Im(\tau^*)}\right)^2$$

We consider a Neumann preconditioner [15, Chapter 12.3] of degree n as an approximation to the inverse of \mathcal{A} ,

$$\mathcal{A}^{-1} \approx \sum_{i=0}^{n} \left(I - \xi \mathcal{A} \right)^{i} =: p_{n}(\mathcal{A}), \quad \text{with parameter } \xi = \frac{1}{c_{0}} = \frac{\overline{\tau}^{*} - \tau^{*}}{\overline{\tau}^{*}}.$$
(24)

and its monic representation $p_n(\mathcal{A}) = \sum_{i=0}^n \gamma_i \mathcal{A}^i$. Shift-invariance, cf. Remark 3.1, can be preserved by the Neumann preconditioner if the following holds,

$$(\mathcal{A} - \eta_k I) p_{n,k}(\mathcal{A}) = \mathcal{A} p_n(\mathcal{A}) - \tilde{\eta}_k I,$$
(25)

where $p_{n,k}(\mathcal{A}) = \sum_{i=0}^{n} \gamma_{i,k} \mathcal{A}^{i}$ is a polynomial preconditioner of degree *n* for the *k*-th shifted matrix $(\mathcal{A} - \eta_k I)$. Substitution yields,

$$\sum_{i=0}^{n} \gamma_{i,k} \mathcal{A}^{i+1} - \sum_{i=0}^{n} \eta_k \gamma_{i,k} \mathcal{A}^i - \sum_{i=0}^{n} \gamma_i \mathcal{A}^{i+1} + \tilde{\eta}_k I = 0.$$
(26)

The latter (26) is a difference equation and can be solved backwards [18]:

$$\begin{split} \gamma_{n,k} &= \gamma_n, \\ \gamma_{i-1,k} &= \gamma_{i-1} + \eta_k \gamma_{i,k}, \quad \text{for } i = n, ..., 1, \\ \tilde{\eta}_k &= \eta_k \gamma_{0,k}. \end{split}$$

As a result, we solve the shifted systems on the right-hand side in (25) with Algorithm 1 using the Neumann preconditioner (24) of degree n. An alternative polynomial preconditioner for shifted systems is derived in [35].

Remark 5.1. Note that $p_n(\mathcal{A}) = c_0 \mathcal{A}^{-1}$, as $n \to \infty$, converges to the scaled inverse of \mathcal{A} . From (25) we conclude that $p_{n,k}(\mathcal{A}) = (c_0 - \tilde{\eta}_k)(A - \eta_k I)^{-1}, n \to \infty$, and, hence, $p_{n,k}$ is a polynomial preconditioner of degree n for the k-th shifted problem $(A - \eta_k I)$.

5.2. Matrix equation formulation with a spectral scaling strategy

An alternative approach to efficiently solve multi-frequency wave problems is to rewrite the discretized problem as a matrix equation $A(\mathbf{X}) = B$, where the block unknown \mathbf{X} is the stacked numerical solution at different frequencies [20] and A is a linear operator of the form, $A(\mathbf{X}) = \sum_{j=1}^{J} A_j \mathbf{X} B_j^{\mathsf{T}} = C$, where J = 2 in (27). Consider, for instance, Problem 2.3 and the reformulation,

$$A(\mathbf{X}) := K\mathbf{X} - M\mathbf{X}\Omega^2 = B, \quad \text{where } \Omega := \text{diag}(\omega_1, ..., \omega_{n_*}), \ B := \mathbf{\hat{s}}\mathbb{1}^\mathsf{T}, \tag{27}$$

and where the unknown is $\mathbf{X} := [\hat{\mathbf{u}}_1, ..., \hat{\mathbf{u}}_{n_s}]$. The spectrum of the operator A is equal to the spectrum of the block diagonal matrix of the corresponding vectorized problem,

$$\begin{bmatrix} (K - \omega_1^2 M) & & \\ & \ddots & \\ & & (K - \omega_{n_s}^2 M) \end{bmatrix} \begin{pmatrix} \mathbf{\hat{u}}_1 \\ \vdots \\ \mathbf{\hat{u}}_{n_s} \end{bmatrix} = \begin{pmatrix} \mathbf{\hat{s}} \\ \vdots \\ \mathbf{\hat{s}} \end{pmatrix},$$

and, hence, the spectrum of A equals the union of the spectra of the shifted systems in Problem 2.3. Consider the iterative solution of (27) with a global Krylov method such as global GMRES [23]. The analysis of Section 4 can be used in order to improve the convergence of global GMRES for the matrix equation (27). Therefore, we define the preconditioners,

$$P_1^{-1}(\mathbf{Y}) := (K - \tau^2 M)^{-1} \mathbf{Y} \Gamma, \quad \text{where } \Gamma := \text{diag} \left(1 - \eta_1, ..., 1 - \eta_{n_s} \right), \tag{27a}$$

$$P_2^{-1}(\mathbf{Y}) := \mathbf{Y}R, \qquad \text{where } R := \text{diag}\left(1, e^{-i(\phi_2 - \phi_1)}, ..., e^{-i(\phi_{n_s} - \phi_{n_s - 1})}\right), \qquad (27b)$$

12

where ϕ_k is the angular component of \hat{c}_k in Lemma 4.1(ii), and $\eta_k := \omega_k^2/(\omega_k^2 - \tau)$ as in (11). Because of the correspondence of the spectrum of the matrix operator with the shifted systems, the above preconditioners have the following interpretation that are illustrated in Figure 4: The application of P_1^{-1} as a right preconditioner in global GMRES is equivalent to (13) and the left-hand side in (11) and yields a spectrum that is equal to the union of the circles described in Lemma 4.1. For the fast convergence this spectrum is not favorable. Therefore, we apply P_2^{-1} as a second-level preconditioner that yields a rotation of the spectrum to the right half plane.



Figure 4: Left: Block spectrum of the linear operator (27) with shift-and-invert preconditioner (27a) that is equivalent to the multi-shift approach. Right: Block spectrum after additional rotation with (27b). We use a surrogate problem with $n_s = 5$ frequencies equally spaced within $f_k \in [1, 9]$ Hz.

Note that the preconditioners (27a) and (27b) commute, i.e. $(A \circ P_1^{-1}) \circ P_2^{-1} = (\mathcal{A} \circ P_2^{-1}) \circ P_1^{-1}$. For global GMRES, this allows to apply P_1 inexactly and, trivially, apply P_2 exactly.

Remark 5.2. In the next section we also consider the case when Problem 2.1 is reformulated as the matrix equation $A(\mathbf{X}) := K\mathbf{X} + iC\mathbf{X}\Omega - K\mathbf{X}\Omega^2 = B$, cf. Experiment 6.3.

6. Numerical experiments

We present numerical examples for a finite element discretization² of the time-harmonic elastic wave equation (5a)-(5c) in 2D. The problem setting shown in Figure 5 (left) is an inhomogeneous *wedge* problem inspired by the acoustic analogue proposed in [36] which has been used for the demonstration of spectral analysis of Helmholtz problems in [4]. The computational domain is $\Omega = [0, 600] \times [0, 1000]$ meter. Whenever Problem 2.1 is solved, we place a point source at (300, 0) meter and prescribe absorbing Sommerfeld conditions on the upper boundary, cf. Figure 5 (middle). Generally speaking, Problem 2.3 is easier to solve numerically. In our numerical examples we consider Problem 2.3 with a point source at the center of Ω and consider reflecting walls on the entire boundary, cf. Figure 5 (right).

Experiment 6.1 (Proof-of-concept). We numerically demonstrate the key findings of Lemma 4.1. In particular, we show that $\tau^*(\epsilon, \omega_{\min}, \omega_{max})$ is independent of the number of frequencies within $[\omega_{\min}, \omega_{max}]$. Moreover, we show mesh-independency of τ^* and demonstrate the direct connection between the bounding circles described in Lemma 4.1 (ii) and the convergence behavior of multi-shift GMRES applied to Problem 2.1.

The experiments reported in Table 1 are performed at fixed frequency range and for a fixed damping parameter. The shift-and-invert preconditioner is applied to multi-shift GMRES at the optimal seed frequency corresponding to (15). We consider n_s equidistantly-spaced frequencies. The results show the expected result that more frequencies

²For the finite element discretization we use the Python package nutils (http://nutils.org).

13



Figure 5: Numerical set-up: Material density ρ in $[kg/m^2]$ (left), and numerical solution of Problem 2.1 at $\epsilon = 0$ (middle) and Problem 2.3 with damping $\epsilon = 0.05$ (right) at different frequencies. See [20] for a detailed description of the test problem.

within the same interval can be solved at no extra iterations, and at low extra computational costs. When repeating some of the experiments on a finer mesh we conclude mesh-independence of the optimal seed parameter, cf. Table 2.

Table 1: Multi-shift GMRES using the optimal seed parameter τ^* according to Lemma 4.1, and fixed damping $\epsilon = 0.05$. Discretization size of $h_x = h_z = 5m$ implies N = 48,642 dofs.

| $\omega_{\rm min}/2\pi$ [Hz] | $\omega_{\rm max}/2\pi$ [Hz] | n_s | # iterations | CPU time [s] |
|------------------------------|------------------------------|-------|--------------|--------------|
| | | 5 | 106 | 45.6 |
| 1 | 5 | 10 | 106 | 48.7 |
| | | 20 | 106 | 47.3 |
| | | 5 | 251 | 205.1 |
| 1 | 10 | 10 | 252 | 223.7 |
| | | 20 | 252 | 243.5 |

Table 2: Setting as in Table 1 and discretization size half compared to Table 1.

| $\omega_{\rm min}/2\pi$ [Hz] | $\omega_{\rm max}/2\pi$ [Hz] | n_s | # iterations | CPU time [s] |
|------------------------------|------------------------------|-------|--------------|--------------|
| 1 | 5 | 10 | 103 | 189.4 |
| 1 | 10 | 10 | 246 | 770.10 |

We next demonstrate the close relation between the spectral bounds derived in Lemma 4.1 and the convergence behavior of multi-shift GMRES preconditioned with a shift-and-invert preconditioner at optimal seed frequency τ^* . We, therefore, consider the same multi-frequency setting at two different seed frequencies. Figure 6 shows the respective convergence curves next to the bounding circles described in Lemma 4.1. When comparing the two choices for τ , we note that the circles corresponding to the optimal τ are further away from the origin which yields a smaller bound in Corollary 3.3. Moreover, the *outlier* in Figure 6a motivates the min-max criterion chosen for the optimization in (15).

Experiment 6.2 (Shifted Neumann preconditioner). *In this numerical experiment we study the effect of the shifted Neumann preconditioner* (24) *on the convergence behavior of multi-shift GMRES within the two-level preconditioning technique described in Section 5.1.*

A major drawback of GMRES (and its multi-shift variant) is the increasing computational work and memory requirement when the number of iterations grows. This can be overcome by restarting Algorithm 1, cf. [2]. If

the matrix-vector operation is relatively cheap, polynomial preconditioners [37] are an important alternative. If a Neumann polynomial (24) of degree *n* is applied, the number of matrix-vector products per iteration is n + 1. The experiment in Table 3 shows that GMRES iteration numbers and computation times can be reduced by approximately a factor of 4 compared to the case without shifted Neumann preconditioner (n = 0).



Figure 6: Relation between convergence of multi-shift GMRES and spectral bounds of Lemma 4.1. Here, we chose $\epsilon = 0.7$ which yields a value for the multi-shift GMRES bound of 0.812 at $\tau = (0.3 - 0.7i)\omega_{max}$ which is significantly larger than 0.659 obtained at the optimum.

Table 3: Multi-shift GMRES without restarting using optimal seed parameter τ^* according to Lemma 4.1, $\epsilon = 0.05$, and a Neumann polynomial preconditioner (24) of degree *n*. We consider $n_s = 10$ frequencies in Problem 2.1 equally spaced within the interval $\omega_k/2\pi = [1, 10]$ Hz. The problem size is N = 48,642 dofs.

| <i>n</i> = | 10 | 5 | 4 | 3 | 2 | 1 | 0 |
|--------------|-------|-------|-------|-------|-------|-------|--------|
| # iterations | 45 | 64 | 94 | 80 | 121 | 150 | 252 |
| CPU time [s] | 52.10 | 46.88 | 64.69 | 47.39 | 73.21 | 89.74 | 213.18 |

In Figure 7 we compare convergence for different values of the seed frequency, and Neumann preconditioners of

degree n = 0 (only shift-and-invert) and n = 5. We observe that the seed parameter τ^* that minimizes the bound in Corollary 3.3 yields an iteration number close to the optimum. Moreover, we note in Figure 7a that in the case of large damping the bound in Corollary 3.3 is more descriptive and the choice of τ has a larger influence on the convergence of multi-shift GMRES. For small damping ($\epsilon = 0.05$) on the other hand, the parameter choice in the Neumann preconditioner at degree n = 5 gains importance and iteration numbers can be reduced up to a factor of 2, cf. Figure 7b.



Figure 7: Optimality of the shifted Neumann preconditioner. Left: Without preconditioner. Right: Neumann polynomial preconditioner at n = 5. The optimal seed parameter τ^* is marked with a cross, and compared against values of τ with varying imaginary part. All values are scaled by ω_{max} . In the present experiment, the frequency range is fixed at $f_k \in [1, 5]$ Hz with $n_s = 5$ and 12, 322 dofs.

Experiment 6.3 (Matrix equation (27) with spectral rotation). *In this experiment we use global GMRES [23] to solve the matrix equation* (27) *preconditioned by* (27a). *The experiment demonstrates the benefit of spectral rotation* (27b) *as a second-level preconditioner for the matrix equation approach described in Section 5.2.*

We solve Problem 2.3 reformulated as a matrix equation (27) using global GMRES [23]. As explained in Section 5.2 the block preconditioner P_1 (27a) yields the spectral situation of Lemma 4.1 in a matrix equation framework. When C = 0 the eigenvalues of the preconditioned linear operator lie on the bounding circles described in Lemma 4.1(ii), see Figure 4. In Table 4 we evaluate the effect of the rotation P_2 (27b) for two different frequency ranges and for different total number of frequencies $n_s = \{5, 15\}$. The comparison shows clearly the benefit of rotating the spectrum. This becomes more evident when the number of frequencies is increased from 5 to 15. Since all circles are rotated on top of each other, the clustering of the spectrum is the same in both cases and we observe an almost equal iteration number for global GMRES. Moreover, we consider the case where Sommerfeld boundary conditions are present and the matrix equation in Remark 5.2 is solved. The effects with respect to P_2 are similar.

Experiment 6.4 (An interval splitting strategy). We consider an interval $I = [\omega_{\min}, \omega_{\max}]$ of equidistantly spaced frequencies, and assume $n_p > 1$ available parallel processors. This experiment investigates a strategy for splitting I into n_p subintervals based on the choice for τ^* according to (15) for each subinterval such that a balanced load in agreement with Corollary 4.4 is achieved.

For fixed damping parameter $\epsilon > 0$, the optimal seed frequency in (15) is a function of the frequency range of the original problem only, i.e. $\tau^* = \tau(\epsilon, \omega_{\min}, \omega_{\max})$. Consider first the case where $n_p = 2$ CPUs are present and the interval of frequencies can be split into two parts, $I = [\omega_{\min}, \omega_{\max}] = [\omega_{\min}, \omega_{\min}] \cup [\omega_{\min}, \omega_{\max}]$ with seed parameter chosen optimally according to (15) for both subinterval. In Figure 8a this splitting point ω_{\min} is varied, and the largest iteration number (marked by crosses) and the larger bound (12) at the respective optimum for the two subintervals is reported. We conclude that the best splitting point is when the boundary ratios are equal, i.e. $\omega_{\min}/\omega_{\min} = \omega_{\max}/\omega_{\min}$

Table 4: Solution of (27) with global GMRES: The preconditioners \mathcal{P}_1 and \mathcal{P}_2 are as defined in (27a) and (27b), respectively. Damping is introduced via $\omega_k^2 \mapsto (1 - i\epsilon)\omega_k^2$, with $\epsilon = 0.1$ in this table. The considered problem has 48,642 dofs, and global GMRES is restarted after 200 iterations.

| | frequency range | n_s | $A(P_1^{-1}(\mathbf{X})) = B$ | $A(P_1^{-1}(P_2^{-1}(\mathbf{X}))) = B$ |
|--------------|---------------------|-------|-------------------------------|---|
| | $f_k \in [1, 3]$ Hz | 5 | 220.2 (301 iter.) | 93.8 (164 iter.) |
| <i>C</i> = 0 | $f_k \in [1, 3]$ Hz | 15 | 2296.9 (702 iter.) | 203.2 (171 iter.) |
| | $f_k \in [6, 9]$ Hz | 5 | 1356.5 (983 iter.) | 22.2 (66 iter.) |
| | $f_k \in [6, 9]$ Hz | 15 | no convergence | 53.8 (65 iter.) |
| <i>C</i> ≠ 0 | $f_k \in [1,3]$ Hz | 5 | 94.9 (203 iter.) | 24.8 (72 iter.) |
| | $f_k \in [1, 3]$ Hz | 15 | 502.1 (300 iter.) | 66.0 (76 iter.) |
| | $f_k \in [6, 9]$ Hz | 5 | 499.4 (566 iter.) | 18.5 (54 iter.) |
| | $f_k \in [6,9]$ Hz | 15 | 1827.2 (627 iter.) | 42.8 (53 iter.) |

which is obtained at the geometric mean at $\omega_{mid} = 3$ Hz in Figure 8a. In the subsequent experiments in Figure 8b and 8c, we report the upper interval boundary of the first subinterval and apply the previously derived splitting strategy *inductively* to the $n_p - 1$ remaining subintervals. In conclusion, a splitting equidistantly on a logarithmic scale yields best results, and in the present test case the number of iterations can be reduced by this strategy from 43 (at $n_p = 1$) to 14 (at $n_p = 4$).



Figure 8: Effect of interval splitting when $I = 2\pi [1,9]$ Hz and $\epsilon = 0.5$ for $n_p = \{2,3,4\}$. We investigate the splitting point for the first out of n_p subintervals. From left to right, the splitting strategy is applied 'recursively' to the $n_p - 1$ remaining subintervals.

Experiment 6.5 (The undamped ($\epsilon \equiv 0$) case). In this experiment we study the quality of $\tau^*(0)$ as in (22) for the case when no viscous damping is present, cf. Remark 4.2. This choice is compared to choices found in different literatures.

The optimality of τ^* in (15) is derived for positive damping parameter $\epsilon > 0$ because only then the circles that bound the preconditioned spectra do not touch the origin, i.e. $|c_k| > R_k$ in (12). The graph in Figure 2 (left), however, shows a smooth dependence of $\tau^*(\epsilon, s_{\min}, s_{\max})$ on ϵ and, in particular, yield an optimal value in the case of $\epsilon = 0$ stated in Remark 4.2. In Table 5 we compare the seed parameter $\tau^*(0)$ with two choices found in the literature: $\tau = (1 - 0.5i)\omega_{\max}$ in [4] and the hand optimized value $\tau = (0.7 - 0.3i)\omega_{\max}$ used in [20, 34]. Moreover, we use an alternative optimization criteria for minimizing the bound in Corollary 3.3 as a third comparison value. The results in Table 5 show that optimality in terms of GMRES iteration numbers is lost but, on the other hand, comparable results to the established choices in literature are obtained.

Table 5: Multi-shift GMRES without restarting using different seed parameters τ and no damping ($\epsilon = 0$). We consider $n_s = 10$ frequencies equally spaced within the interval $\omega_k/2\pi = [5, 10]$ Hz. The problem size is 48,642 dofs.

| seed shift | $\tau^* = (0.66 - 0.26i)\omega_{\max}$ | $\tau = (0.7 - 0.3i)\omega_{\rm max}$ | $\tau = (1 - 0.5i)\omega_{\rm max}$ | $\tau = \min_{\tau} \operatorname{mean}_k(R/ c_k)$ |
|--------------|--|---------------------------------------|-------------------------------------|---|
| # iterations | 226 | 201 | 295 | 300 |
| CPU time [s] | 139.5 | 109.3 | 247.6 | 257.6 |

7. Conclusions

We have derived an optimal seed parameter τ^* for the shift-and-invert preconditioner $\mathcal{P}(\tau^*) = (\mathcal{K} - \tau^* \mathcal{M})$ applied to a sequence of shifted systems. For a given set of frequencies $s_k \in [s_{\min}, s_{\max}]$ the optimal seed is an explicit function of the extreme frequencies, and the viscous damping parameter ϵ , cf. (15) in Lemma 4.1. The optimality of the parameter is derived with respect to a well-known GMRES convergence bound that has been extended to the multi-shift setting, and in the presence of viscous damping, i.e. $\epsilon > 0$. Our numerical experiments, however, prove the usefulness even for the case without damping (Experiment 6.5). Comparisons with shift-and-invert preconditioners with parameter different from τ^* show a slower convergence behavior of multi-shift GMRES and, therefore, numerically prove optimality of τ^* (see Experiment 6.1 and Experiment 6.2).

The spectral analysis that has been carried out for the derivation of τ^* gives valuable insight that we exploit within two applications: In Section 5.1, a shifted Neumann preconditioner is derived that has minimum spectral radius. The numerical examples in Experiment 6.2 show that an increase of the degree of the Neumann polynomial leads to a significant reduction of GMRES iteration numbers and, hence, of memory requirements. Moreover, numerical tests have shown that the Neumann preconditioner based on τ^* yields fast convergence especially in the case of a small damping parameter (cf. Experiment 6.2). In Section 5.2, we apply global GMRES to a matrix equation reformulation of the shifted problem. Then, the bounding circles of the shifted spectra can be rotated which yields a more favorable spectrum for the matrix equation approach, cf. Experiment 6.3.

We have also considered the situation when more than one CPU is present, and a sequence of shifted problems within a fixed interval $[s_{\min}, s_{\max}]$ can be split into subintervals that are solved simultaneously on each available CPU. In Experiment 6.4 we give strong numerical evidence that an optimal interval splitting strategy is to split equidistantly on a logarithmic scale. With respect to future work, we would like to point out that multi-shift GMRES with a shift-and-invert preconditioner (as in Algorithm 1) fits the more general framework of rational Krylov methods [38]. The fact that we apply a single shift-and-invert preconditioner corresponds to the situation where a rational Krylov space with denominator degree equals to one is chosen. The recent RKFIT algorithm [39] provides a strategy for pole selection in rational Krylov methods and can, thus, be used for comparison. Moreover, the presented interval splitting strategy yields (optimal) seed parameters for each subinterval and can be exploited in the framework of multi-preconditioned GMRES for shifted systems [19].

Acknowledgments

We want to thank René-Édouard Plessix for helpful discussions on computational geophysics, and Christian Schröder for pointing the authors to [32, Table 1.1]. Shell Global Solutions International B.V. is gratefully acknowledged for financial support of the first author.

References

- 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.
- [2] A. Frommer, U. Glässner, Restarted GMRES for Shifted Linear Systems, SIAM J. Sci. Comput. 19 (1998) 15-26.
- [3] Y. Erlangga, C. W. Oosterlee, C. Vuik, A novel multigrid based preconditioner for heterogeneous Helmholtz problems, SIAM J. Sci. Comput. 27 (2006) 1471–1492.
- [4] 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.
- [5] A. Sheikh, D. Lahaye, L. G. Ramos, R. Nabben, C. Vuik, Accelerating the shifted Laplace preconditioner for the Helmholtz equation by multilevel deflation, J. Comput. Phys. 322 (2016) 473–490.
- [6] P.-H. Cocquet, M. J. Gander, How Large a Shift is Needed in the Shifted Helmholtz Preconditioner for its Effective Inversion by Multigrid?, SIAM J. Sci. Comput. 39 (2017) A438–A478.
- [7] A. De Hoop, Handbook of Radiation and Scattering of Waves, Academic Press, London, United Kingdom, 1995.
- [8] M. Baumann, M. B. van Gijzen, Nested Krylov methods for shifted linear systems, SIAM J. Sci. Comput. 37 (2015) S90–S112.
- [9] G. Rizzuti, W. Mulder, Multigrid-based 'shifted-Laplacian' preconditioning for the time-harmonic elastic wave equation, J. Comput. Phys. 317 (2016) 47–65.
- [10] T. Airaksinen, A. Pennanen, J. Toivanen, A damping preconditioner for time-harmonic wave equations in fluid and elastic material, J. Comput. Phys. 228 (2009) 1466–1479.
- [11] J. Virieux, S. Operto, An overview of full-waveform inversion in exploration geophysics, Geophysics 73 (2009) VE135-VE144.

- [12] W. A. Mulder, R.-E. Plessix, How to choose a subset of frequencies in frequency-domain finite-difference migration, Geophys. J. Int. 158 (2004) 801–812.
- [13] R. Pratt, Seismic waveform inversion in the frequency domain, Part 1: Theory and verification in a physical scale mode, Geophysics 64 (1999) 888–901.
- [14] R.-E. Plessix, C. A. Pérez Solano, Modified surface boundary conditions for elastic waveform inversion of low-frequency wide-angle active land seismic data, Geophys. J. Int. 201 (2015) 1324–1334.
- [15] Y. Saad, Iterative Methods for Sparse Linear Systems: Second Edition, Society for Industrial and Applied Mathematics, 2003.
- [16] A. Frommer, BiCGStab(ℓ) for families of shifted linear systems, Computing 70 (2003) 87–109.
- [17] M. van Gijzen, G. Sleijpen, J.-P. Zemke, Flexible and multi-shift Induced Dimension Reduction algorithm for solving large sparse linear systems, Numer. Linear Algebra Appl. 1 (2014) 1–25.
- [18] M. I. Ahmad, D. B. Szyld, M. B. van Gijzen, Preconditioned multishift BiCG for \mathcal{H}_2 -optimal model reduction, Technical Report 12-06-15, Department of Mathematics, Temple University, 2013.
- [19] T. Bakhos, P. K. Kitanidis, S. Ladenheim, A. K. Saibaba, D. B. Szyld, Multipreconditioned GMRES for Shifted Systems, Technical Report 16-03-31, Department of Mathematics, Temple University, 2016.
- [20] M. Baumann, R. Astudillo, Y. Qiu, E. Ang, M. B. van Gijzen, R.-E. Plessix, An MSSS-Preconditioned Matrix Equation Approach for the Time-Harmonic Elastic Wave Equation at Multiple Frequencies, Technical Report 16-04, Delft University of Technology, 2016.
- [21] K. M. Soodhalter, Block Krylov Subspace Recycling for Shifted Systems with Unrelated Right-Hand Sides, SIAM J. Sci. Comput. 38 (2016) A302–A324.
- [22] K. M. Soodhalter, D. B. Szyld, F. Xue, Krylov subspace recycling for sequences of shifted linear systems, Appl. Numer. Math. 81C (2014) 105–118.
- [23] K. Jbilou, A. Messaoudi, H. Sadok, Global FOM and GMRES algorithms for matrix equations, Appl. Numer. Math. 31 (1999) 49-63.
- [24] T. A. Manteuffel, S. V. Parter, Preconditioning and boundary conditions, SIAM J. Numer. Anal. 27 (1990) 656-694.
- [25] 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 (2006) E57–E63.
- [26] D. Lahaye, C. Vuik, How to Choose the Shift in the Shifted Laplace Preconditioner for the Helmholtz Equation Combined with Deflation, in: D. Lahaye, J. Tang, C. Vuik (Eds.), Modern Solvers for Helmholtz Problems, Springer International Publishing, 2017, pp. 85–112.
- [27] M. J. Gander, I. G. Graham, E. A. Spence, Applying GMRES to the Helmholtz equation with shifted Laplacian preconditioning: What is the largest shift for which wavenumber-independent convergence is guaranteed?, Num. Math. 131 (2015) 567–614,.
- [28] 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.
- [29] E. Bécache, S. Fauqueux, P. Joly, Stability of Perfectly Matched Layers, Group Velocities and Anisotropic Waves, Technical Report RR-4304, INRIA, 2006.
- [30] V. Druskin, R. Remis, Coordinate-stretching method to simulate wave propagation in unbounded domains, SIAM J. Sci. Comput. 35 (2013) B376–B400.
- [31] Y. Saad, M. Schultz, GMRES: a generalized minimal residual algorithm for solving nonsymetric linear systems, SIAM J. Sci. Stat. Comput. 7 (1986) 856–869.
- [32] F. Tisseur, K. Meerbergen, The quadratic eigenvalue problem, SIAM review 43 (2001) 235-286.
- [33] M. Baumann, An efficient two-level preconditioner for multi-frequency wave propagation problems: Numerical examples and MAPLE derivations, https://github.com/Manuel/MBaumann/opt_tau, April 2017. doi:10.5281/zenodo.508357.
- [34] M. Baumann, M. van Gijzen, Efficient iterative methods for multi-frequency wave propagation problems: A comparison study, in: Proceedings of the International Conference on Computational Science (ICCS), 2017 (to appear).
- [35] 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.
- [36] R.-E. Plessix, W. A. Mulder, Seperation-of-variables as a preconditioner for an iterative Helmholtz solver, Appl. Numer. Math. 44 (2004) 385–400.
- [37] M. B. van Gijzen, A polynomial preconditioner for the GMRES algorithm, J. Comput. Appl. Math. 59 (1995) 91-107.
- [38] A. Ruhe, Rational Krylov sequence methods for eigenvalue computation, Linear Algebra Appl. 58 (1984) 391-405.
- [39] M. Berljafa, S. Güttel, The RKFIT algorithm for nonlinear rational approximation, MIMS EPrint 2015.38, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, 2015.

Supplementary material

An interactive visualization using Bokeh (http://bokeh.pydata.org/) demonstrates the findings of Lemma 4.1. The visualization is purely browser-based and can be obtained from:

http://www.manuelbaumann.de/opt_tau

and is authored by: Manuel Baumann.

The numerical experiments in Section 6 are available from the author's github repository [33].