



An extended Krylov subspace model-order reduction technique to simulate wave propagation in unbounded domains

Vladimir Druskin^a, Rob Remis^{b,*}, Mikhail Zaslavsky^a

^a Schlumberger Doll Research, 1 Hampshire St., Cambridge, MA 02139, USA

^b Circuits and Systems Group, Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, Mekelweg 4, 2628 CD Delft, The Netherlands

ARTICLE INFO

Article history:

Received 26 February 2014

Received in revised form 10 April 2014

Accepted 28 April 2014

Available online 6 May 2014

Keywords:

Reduced-order modeling
Extended Krylov subspaces
Coordinate stretching
Scattering poles

ABSTRACT

In this paper we present a novel extended Krylov subspace reduced-order modeling technique to efficiently simulate time- and frequency-domain wavefields in open complex structures. To simulate the extension to infinity, we use an optimal complex-scaling method which is equivalent to an optimized perfectly matched layer in which the frequency is fixed. Wavefields propagating in strongly inhomogeneous open domains can now be modeled as a non-entire function of the complex-scaled wave operator. Since this function contains a square root singularity, we apply an extended Krylov subspace technique to construct fast converging reduced-order models. Specifically, we use a modified version of the extended Krylov subspace algorithm as proposed by Jagels and Reichel [14], since this algorithm allows us to balance the computational costs associated with computing powers of the wave operator and its inverse. Numerical experiments from electromagnetics and acoustics illustrate the performance of the method.

© 2014 Elsevier Inc. All rights reserved.

1. Introduction

Frequency- and time-domain wave problems in unbounded domains emerge in a wide range of applications: from seismic processing in geophysics to metamaterial development, from bioelectromagnetics to radar and antenna design. Furthermore, developing fast and robust forward modeling methods is not only a significant topic by itself, it is also of great importance in inversion and optimization algorithms. In all these different disciplines, conventional forward solution techniques are usually the method of choice. For time-domain problems, it is typically the Finite-Difference Time-Domain method (FDTD method, see, for example, [19]) that is used, while discretized frequency-domain problems are usually solved using linear system solvers. When solutions for multiple frequencies are required, conventional workflow consists of solving frequency-domain problems for each frequency separately. In many applications, such methods quite often leave much room for improvement. Indeed, discretization grids for 3D problems may consist of up to a billion nodes and, even with state-of-the-art preconditioners, obtaining the solution of discretized frequency-domain problems results in rather computationally intensive tasks.

Projection-based model reduction is a well-established tool and allows us to obtain time- and frequency-domain solutions by projecting the large scale dynamic system on a small Krylov or rational Krylov subspace. In an FDTD setting, for

* Corresponding author.

E-mail addresses: Druskin1@slb.com (V. Druskin), R.F.Remis@TUDelft.NL (R. Remis), mzaslavsky@slb.com (M. Zaslavsky).

example, optimal time steps can be determined, though rather implicitly, via polynomial Krylov subspace (PKS) projection methods. These Krylov methods have been shown to provide significant speed up for lossy diffusion-dominated problems, e.g., for the diffusive Maxwell equations (see [21]). For lossless problems, however, e.g. wave problems in lossless media on bounded domains, such an optimality hardly provides any acceleration compared with time-stepping at the Courant limit [4].

The principal difference between wave problems in bounded and unbounded domains is that the latter one is essentially lossy. Indeed, loosely speaking, infinity acts like an absorber and for any bounded subdomain there is no wave permanently traveling inside it without energy loss. Therefore, an attempt to construct a Krylov subspace projection method to speed up the solution of such problems is not hopeless. In fact, two of the authors showed that for wavefield problems on unbounded domains, standard polynomial Krylov methods may indeed outperform FDTD on large time intervals provided we respect the delicate spectral properties of the corresponding continuous problem (see [7]). To be more precise, in any practical implementation the domain of computation is necessarily bounded and boundary conditions have to be imposed for domain truncation. If selfadjoint boundary conditions (Dirichlet boundary conditions, for example) are used then the solution to a selfadjoint problem with a continuous spectrum is approximated by a solution to a selfadjoint problem with a discrete spectrum. In the latter case, the solution consists of an infinite summation of periodic eigenfunctions that do not decay. Enlarging the domain of computation does not help, since the eigenfunctions remain periodic and do not show any attenuation. For general sources, polynomial Krylov methods used on discretized bounded domains with selfadjoint boundary conditions cannot outperform FDTD running at the Courant limit. However, for non-selfadjoint boundary conditions that simulate the extension to infinity, polynomial Krylov methods may outperform FDTD on large time intervals. This can be explained by considering the scattering poles expansion of a wavefield u satisfying the wave equation on an unbounded domain [15]. In particular, Lax and Phillips showed that for any fixed point in \mathbb{R}^n , n odd, the solution $u(t)$ can be expanded for large t in a sum of time-exponential modes. These modes decay exponentially in time and correspond to the so-called scattering resonances. The resonances are the poles of the resolvent for the wave equation and are located on the second Riemann sheet of the square root function. In practice, only a finite number of the corresponding modes essentially contribute to the solution for large t . Consequently, polynomial Krylov methods may outperform FDTD for large times if we can capture these resonant poles via a Krylov subspace method.

The scattering poles can be found by truncating the domain of interest via complex coordinate scaling also known as Aguilar–Balslev–Combes–Simon theory [12]. This coordinate stretching technique was introduced in the 1970s and is since then used extensively in atomic and molecular physics to calculate energies, widths, and cross-sections of open quantum systems. The technique coincides with Berenger’s perfectly matched layer (PML, [2]) using a fixed frequency for coordinate stretching. The problem with complex scaling is, however, that resonances and anti-resonances of the system are found simultaneously resulting in instability of the traditionally computed solution via a sin or cos function of the operator. Two of the authors resolved this issue in [7] by introducing a so-called stability-corrected exponential function that depends on the square root of the wave operator. Any type of Krylov subspace method can now be applied and we split them into two groups depending on the application. The first group targets to applications where the offline cost of constructing the subspace is more important compared to the online cost of subsequent projection. For example, as a stand-alone forward solver, the latter cost is typically negligible. *Polynomial* Krylov subspace methods can be quite competitive here in terms of computation time, especially when the frequency range doesn’t include small frequencies. However, when applied within an inversion/optimization framework, not only the forward problem solution is required but also the derivatives with respect to the parameters that are being optimized or inverted. Depending on the number of these parameters, the computation of the derivatives can make the projection cost dominating. For these applications, multiple solutions of the problem are required, and the same reduced order model can be reused for this purpose. In this case, it is beneficial to construct a subspace of smaller size for a given accuracy even at a larger offline cost. We put the methods of this kind into the second group. Taking this into account, we note that for non-entire operator functions (such as the stability-corrected operator function considered here), *polynomial* Krylov subspace methods are known to be not the best choice in terms of convergence speed. Indeed, *rational* Krylov subspace (RKS) methods are known to be much better suited for non-entire functions (see, e.g., [5]). Moreover, it has recently been shown that RKSs are beneficial for the solution of regular enough wave problems on bounded domains [9].

In this paper, we follow the approach based on the stability-corrected exponential function and construct optimal rational field approximants via a so-called Extended Krylov Subspace Method (EKSM). This Krylov method generates field approximations using both positive and negative matrix powers and produces rational field approximants with restricted poles at the origin and at infinity. The original EKSM of [5] allowed for short recursions and a structured projection matrix only for a fixed number of positive or negative matrix powers. More flexible EKSM modifications, allowing for sequences of nested extended subspaces with increasing positive and negative matrix powers, were suggested in [18,13]. Here we use a modified version of the algorithm proposed in [13], since it allows for an arbitrary ratio between positive and negative matrix powers. We illustrate the performance of our proposed solution technique in Section 5, where we present a number of numerical experiments related to the solution of the wave equation and solutions of Helmholtz’s equation for a wide range of frequencies. In our experiments for time-domain problems, we show that our solution technique significantly outperforms polynomial Krylov subspace projection in terms of subspace dimension, while for Helmholtz’s equation we show that with the PKS method we obtain solutions for a complete frequency range at the cost of an unpreconditioned linear solver applied to the problem with the smallest frequency. Our results clearly show that EKS significantly outperforms a polynomial Krylov approach when solutions for small frequencies are required.

2. Problem formulation and polynomial Krylov subspace projection

Let us consider the hyperbolic initial value problem

$$Au + u_{tt} = 0 \quad \text{with } u|_{t=0} = 0 \text{ and } u_t|_{t=0} = b \quad (1)$$

for $t \geq 0$. Its frequency-domain counterpart is

$$Au + s^2u = b. \quad (2)$$

In these equations, A is a selfadjoint nonnegative partial differential equation operator on an unbounded domain having an absolutely continuous spectrum. The solution of Eq. (1) can be written in terms of operator functions as

$$u(t) = \eta(t)A^{-1/2} \sin(A^{1/2}t)b, \quad (3)$$

where $\eta(t)$ is the Heaviside unit step function and $A^{1/2}$ is a principal square root of operator A (see [17]). In [7,3] we discretized the above partial differential equation on a uniform second-order grid and used Zolotarev's optimal rational approximation of the square root to implement an optimal complex scaling method for domain truncation. (The optimal scaling approach of [7] deals with propagative modes only, while in [3] a more advanced implementation is presented in which evanescent and propagative modes are taken into account simultaneously.) Writing the resulting system as

$$\tilde{A}_N u_N + u_{N;tt} = 0 \quad \text{with } u_N|_{t=0} = 0 \text{ and } u_{N;t}|_{t=0} = b_N, \quad (4)$$

where the tilde indicates that complex-stretching has been applied and the subscript N signifies that we are dealing with a discretized system, it is tempting to take the solution

$$u_N(t) = \eta(t)\tilde{A}_N^{-1/2} \sin(\tilde{A}_N^{1/2}t)b_N \quad (5)$$

as an approximation to the exact solution as given by Eq. (3). The approximation of Eq. (5) is unstable, however, since the eigenvalues of \tilde{A}_N approximate both the resonances and the anti-resonances of the system. Therefore $\sin(\tilde{A}_N^{1/2}t)$ increases exponentially as t increases and u_N is unstable.

Similarly, important frequency-domain field properties are lost if we take

$$u_N(s) = (\tilde{A}_N + s^2I)^{-1}b_N \quad (6)$$

as an approximation to $u(s)$. In particular, the symmetry property $u(\bar{s}) = \bar{u}(s)$ is broken and the limit $\lim_{s \rightarrow 0} u_N(s)$ is not purely real as it should be, of course.

Fortunately, we can correct for these failures by exploiting the symmetry of the exact solution with respect to $(-\tilde{A}_N)^{1/2}$ and $(-\tilde{A}_N)^{1/2}$, where the overbar denotes complex conjugation. Specifically, setting

$$B_N = (-\tilde{A}_N)^{1/2} \quad (7)$$

and taking the principle value for the square root, a stabilized approximation to the continuous open problem of Eq. (1) is given in [7] by

$$\tilde{u}_N(t) = -\eta(t) \operatorname{Re}[B_N^{-1} \exp(-B_N t)]b_N. \quad (8)$$

This approximation is unconditionally stable because of the principle value convention. We note that if we replace \tilde{A}_N by A in Eq. (8), then the expression for the stability-corrected field becomes identical to the operator function expression for u in Eq. (3). In addition, by applying a Laplace transform to Eq. (8), we obtain the frequency-domain result

$$\tilde{u}_N(s) = -\frac{1}{2} [B_N^{-1}(B_N + sI_N)^{-1} + \bar{B}_N^{-1}(\bar{B}_N + sI_N)^{-1}]b_N, \quad (9)$$

where I_N is the identity matrix of order N . This stability-corrected field satisfies the symmetry property $\tilde{u}_N(\bar{s}) = \bar{\tilde{u}_N}(s)$ and becomes purely real as $s \rightarrow 0$. We note that for frequency-domain wave fields operating at an angular frequency ω we take $s = i\omega$.

Having the stability-corrected field approximations available, we can now construct polynomial reduced-order models in the usual way. In particular, the models are drawn from the Krylov subspace

$$\mathbb{K}_m = \operatorname{span}\{b_N, \tilde{A}_N b_N, \tilde{A}_N^2 b_N, \dots, \tilde{A}_N^{m-1} b_N\} \quad (10)$$

and are obtained by following a Petrov–Galerkin procedure. We remark that a basis for \mathbb{K}_m can be generated very efficiently via a Lanczos-type algorithm using a short three-term recurrence relation, since matrix \tilde{A}_N is essentially complex-symmetric (see Section 4). Furthermore, the error of the reduced-order models based on the Krylov space of Eq. (10) is of the same order as the error of the PML discretization in \tilde{A}_N on the frequency interval of the signal. Finally, we note that at each Lanczos iteration we need matrix \tilde{A}_N only to form matrix–vector products and a single Lanczos iteration requires approximately the same amount of work as a single FDTD iteration. Further details can be found in [7], where it is also shown that the reduced-order models generated in this way eventually do indeed outperform FDTD on large time intervals.

3. Extended Krylov subspace reduced-order models

Given that the stability-corrected field approximations of Eqs. (8) and (9) have a square root singularity, we observe that the reduced-order models taken from the Krylov subspace of Eq. (10) may not provide us with the fastest convergence. The reason is that models taken from this space are polynomial field approximations, but rational functions are usually much better at approximating functions near singularities (see [20]). This then naturally leads us to consider rational Krylov subspaces for stability-corrected field approximations. However, the main drawback of utilizing rational Krylov spaces is that large shifted systems of equations (Helmholtz systems) need to be solved to generate bases for these spaces. Without any effective preconditioners, this may simply be unfeasible.

To handle the square root singularity, we therefore resort to the extended Krylov subspaces

$$\mathbb{K}_{m_1, m_2} = \text{span}\{\tilde{A}_N^{-m_1+1}b_N, \dots, \tilde{A}_N^{-1}b_N, b_N, \tilde{A}_N b_N, \dots, \tilde{A}_N^{m_2-1}b_N\}. \tag{11}$$

An extended Krylov subspace can be seen as a special case of a rational Krylov subspace with one expansion point at zero and one at infinity. Obviously, the action of \tilde{A}_N^{-1} on a vector is required to generate a basis for such a space. This essentially amounts to solving a Poisson-type equation for which efficient solution techniques are available [10]. Computing matrix–vector products with the inverse of the system matrix is generally still more expensive than computing matrix–vector products with the system matrix itself. Therefore, from a computational point of view, we prefer to deal with extended Krylov subspaces \mathbb{K}_{m_1, m_2} with $m_1 < m_2$. In addition, the basis vectors should be constructed via short term recurrence relations, since storage of all basis vectors is generally not practical for large-scale real-world problems.

In the original EKS method as introduced for Hermitian matrices in [5], the number of inverse powers m_1 is fixed and nested subspaces

$$\mathbb{K}_{m_1, 1} \subset \mathbb{K}_{m_1, 2} \subset \dots \subset \mathbb{K}_{m_1, m_2}$$

are recursively generated via short term recurrence relations. Furthermore, it is shown that for the matrix square root and $m_2 = m_1$, the number of iterations required for convergence of the EKS reduced-order models is of $O(\sqrt[4]{\|A_N\| \|A_N^{-1}\|})$. We stress that this estimate holds for Hermitian matrices A_N and diagonal ($m_1 = m_2$) subspaces only. Generating a basis for such spaces obviously requires an equal number of iterations with the system matrix and its inverse. For poorly conditioned operators A_N , the size of EKS can be rather large: in our numerical examples below it reaches an order that runs into the thousands. As was shown in [6], for Hermitian problems an RKS approach with optimally chosen shifts allows us to reduce the size of the subspace significantly thanks to the logarithmic dependence of the error bound on the condition number of A_N . For the non-Hermitian case, an adaptive approach (see [8]) can be employed, however, it requires multiple solves of frequency-domain problems for nonzero frequencies which is a rather computationally intensive task by itself.

In the EKS method proposed in [18], again a diagonal extended Krylov subspace is considered, but this time the sequence of nested subspaces

$$\mathbb{K}_{1, 1} \subset \mathbb{K}_{2, 2} \subset \dots \subset \mathbb{K}_{m_1, m_1}$$

is generated via short recurrence relations. The method converges at a rate as predicted by theory, but its computational costs are very large.

In [13,14], Jagels and Reichel developed a more generalized approach in which the sequence of subspaces

$$\mathbb{K}_{1, i+1} \subset \mathbb{K}_{2, 2i+1} \subset \dots \subset \mathbb{K}_{k, ki+1}$$

is generated again via short term recurrences. Here, i is an integer that allows us to optimize the computational costs. In particular, since matrix–vector products with the system matrix are usually more easily computed than matrix–vector products with its inverse, it may be advantageous to apply the extended Krylov approach of Jagels and Reichel, since the number of iterations with the system matrix is essentially i times larger than the number of iterations k that are carried out with the inverse of the system matrix. Different values of i do produce different extended Krylov subspace sequences, however, and the reduced-order models based on these spaces may have different approximating properties. Furthermore, the Krylov projection matrix is a pentadiagonal matrix of order $k(i + 1)$ with a block-tridiagonal structure and consequently the EKS reduced-order models can be evaluated very efficiently.

In this paper, we use a variant of the extended Krylov scheme as proposed by Jagels and Reichel to construct EKS reduced-order models for wavefield problems in open domains. We exploit the complex symmetric structure of the system matrix to generate a basis of the extended Krylov space $\mathbb{K}_{k, ki+1}$. The basis vectors are complex orthogonal by construction and can be computed via short term recurrence relations. This scheme is presented in Section 4 and provides us with an extended Krylov decomposition of the system matrix \tilde{A}_N . The computed decomposition then allows us to construct the EKS reduced-order models for the stability-corrected field approximation of Eqs. (8) and (9).

4. Extended Krylov subspace algorithm

As mentioned above, we are interested in constructing EKS reduced-order models for wavefield quantities in unbounded domains that satisfy the wave equation with variable coefficients. To be specific, let us consider wavefields propagating in two- or three-dimensional unbounded domains. In the interior, we discretize the wave equation on an equidistant second-order grid and use optimal Zolotarev complex-scaling to simulate the extension to infinity. We do not provide any details about this discretization process here, since it is completely described in [7,3]. We do mention, however, that our approach is not limited to second-order schemes in the interior. Other higher-order discretization schemes can be used as well. In fact, since the optimal scaling-method has spectral accuracy, it would be preferable to use a discretization scheme that has spectral accuracy in the interior as well. High-order spectral methods can be applied for example (see [11]) or optimal grids for interior domains as described in [1]. Since our focus is on constructing EKS reduced-order models, we use a straightforward second-order accurate scheme for simplicity only.

After the discretization procedure, we arrive at the semi-discretized system of Eq. (4). Due to complex-scaling, the system matrix \tilde{A}_N belongs to $\mathbb{C}^{N \times N}$ and is not selfadjoint (Hermitian). It does, however, satisfy the symmetry relation

$$\tilde{A}_N^T M = M \tilde{A}_N, \quad (12)$$

where M is a diagonal matrix with products of the spatial stepsizes and medium parameters on its diagonal (for details, see [7]). Note that matrix M is not positive definite and does not even belong to $\mathbb{R}^{N \times N}$, since the optimal stepsizes used for complex-scaling are complex.

Introducing now the bilinear form

$$\langle x, y \rangle = y^T M x \quad (13)$$

we have $\langle \tilde{A}_N x, y \rangle = \langle x, \tilde{A}_N y \rangle$ for any $x, y \in \mathbb{C}^N$ and two (nonzero) vectors x and y belonging to \mathbb{C}^N are said to be complex orthogonal if $\langle x, y \rangle = 0$. Finally, we stress that Eq. (13) does not define an inner product, since M is not positive definite.

Given the symmetry of matrix \tilde{A}_N , we can now generate a complex orthogonal basis for $\mathbb{K}_{k,ki+1}$ by modifying the orthogonalization algorithm for $\mathbb{K}_{k,ki+1}$ and real symmetric matrices of Jagels and Reichel [13,14]. In particular, if we replace the inner products for vectors from \mathbb{R}^N by the bilinear form of Eq. (13), we arrive at an algorithm that generates a complex orthogonal basis of the Krylov space $\mathbb{K}_{k,ki+1}(\tilde{A}_N, b_N)$, where each basis vector has a Euclidean norm equal to one. With $\|\cdot\|$ denoting the Euclidean norm, the algorithm is as follows:

Algorithm Complex Orthogonalization Process for $\mathbb{K}_{k,ki+1}(\tilde{A}_N, b_N)$.

Input: k, i , and source vector b_N

Output: complex orthogonal basis $\{v_j\}_{j=-k}^{ki+1}$ of $\mathbb{K}_{k,ki+1}(\tilde{A}_N, b_N)$

$v_{-1} = 0$;

$\beta_0 = \|b_N\|$; $v_0 = b_N/\beta_0$; $\delta_0 = \langle v_0, v_0 \rangle$;

for $p = 0, 1, \dots, k-1$ **do**

$u := \tilde{A}_N v_{-p}$;

$\alpha_{-p,ip} := \langle u, v_{ip} \rangle$; $u := u - \alpha_{-p,ip}/\delta_{ip} v_{ip}$;

$\alpha_{-p,-p} := \langle u, v_{-p} \rangle$; $u := u - \alpha_{-p,-p}/\delta_{-p} v_{-p}$;

$\beta_{ip+1} := \|u\|$; $v_{ip+1} := u/\beta_{ip+1}$; $\delta_{ip+1} = \langle v_{ip+1}, v_{ip+1} \rangle$;

$u := \tilde{A}_N v_{ip+1}$;

$\alpha_{ip+1,ip} := \langle u, v_{ip} \rangle$; $u := u - \alpha_{ip+1,ip}/\delta_{ip} v_{ip}$;

$\alpha_{ip+1,-p} := \langle u, v_{-p} \rangle$; $u := u - \alpha_{ip+1,-p}/\delta_{-p} v_{-p}$;

$\alpha_{ip+1,ip+1} := \langle u, v_{ip+1} \rangle$; $u := u - \alpha_{ip+1,ip+1}/\delta_{ip+1} v_{ip+1}$;

$\beta_{ip+2} = \|u\|$; $v_{ip+2} = u/\beta_{ip+2}$; $\delta_{ip+2} = \langle v_{ip+2}, v_{ip+2} \rangle$;

for $j = 3, \dots, i$ **do**

$u := \tilde{A}_N v_{ip+j-1}$;

$\alpha_{ip+j-1,ip+j-2} := \langle u, v_{ip+j-2} \rangle$; $u := u - \alpha_{ip+j-1,ip+j-2}/\delta_{ip+j-2} v_{ip+j-2}$;

$\alpha_{ip+j-1,ip+j-1} := \langle u, v_{ip+j-1} \rangle$; $u := u - \alpha_{ip+j-1,ip+j-1}/\delta_{ip+j-1} v_{ip+j-1}$;

$\beta_{ip+j} := \|u\|$; $v_{ip+j} = u/\beta_{ip+j}$; $\delta_{ip+j} = \langle v_{ip+j}, v_{ip+j} \rangle$;

end

$w := \tilde{A}_N^{-1} v_{i(p+1)}$;

$\delta_{i(p+1),-p} := \langle w, v_{-p} \rangle$; $w := w - \delta_{i(p+1),-p}/\delta_{-p} v_{-p}$;

for $j = 1, \dots, i$ **do**

$\delta_{i(p+1),ip+j} := \langle w, v_{ip+j} \rangle$; $w := w - \delta_{i(p+1),ip+j}/\delta_{ip+j} v_{ip+j}$;

end

$\beta_{-(p+1)} = \|w\|$; $v_{-(p+1)} = w/\beta_{-(p+1)}$; $\delta_{-p+1} = \langle v_{-p+1}, v_{-p+1} \rangle$;

end

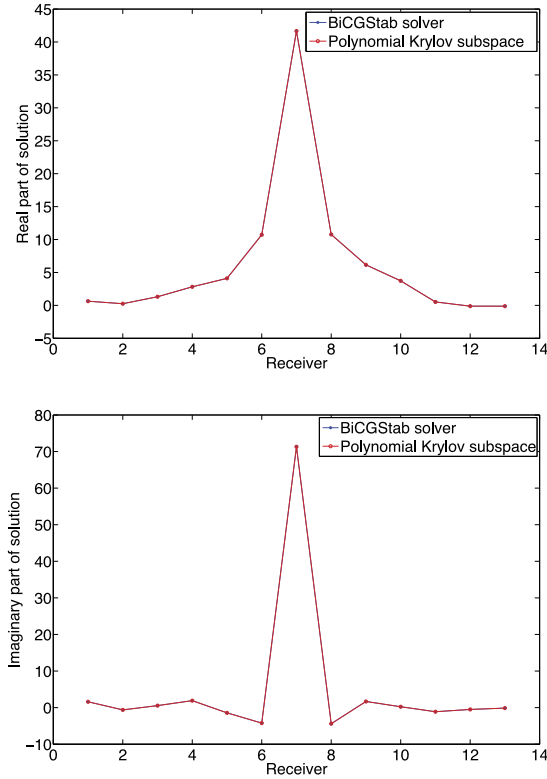


Fig. 1. SEG/EAGE Salt model. Real (top) and imaginary (bottom) parts of solution for frequency 7.5 Hz computed using the BiCGStab and PKS solvers with 1600 matrix–vector multiplications and 1800 iterations, respectively. Both solutions are almost indistinguishable.

After a successful completion of this algorithm, we store the vectors of the generated basis in the N -by- $k(i + 1)$ matrix

$$V_{k(i+1)} = [v_0, v_1, \dots, v_i, v_{-1}, \dots, v_{-k+1}, \dots, v_{ik}] \tag{14}$$

and all iterations can be summarized into the equation

$$\tilde{A}_N V_{k(i+1)} = V_{k(i+1)} H_{k(i+1)} + z_{k(i+1)} e_{k(i+1)}^T, \tag{15}$$

where

$$z_{k(i+1)} = h_{k(i+1)+1, k(i+1)} v_{-k} + h_{k(i+1)+2, k(i+1)} v_{ik+1}$$

and h_{ij} is the (i, j) entry of matrix $H_{k(i+1)}$. Furthermore, matrix $H_{k(i+1)}$ is a pentadiagonal matrix (see [13]) that satisfies

$$D_{k(i+1)} H_{k(i+1)} = V_{k(i+1)}^T \tilde{A}_N V_{k(i+1)}, \tag{16}$$

where $D_{k(i+1)}$ is diagonal matrix with entries $\delta_0, \delta_1, \dots, \delta_i, \delta_{-1}, \dots, \delta_{-k+1}, \dots, \delta_{ik}$. The entries of matrix $H_{k(i+1)}$ can be easily obtained in explicit form from the ones given in [13].

Eq. (15) is the extended Krylov decomposition of the system matrix \tilde{A}_N and can be used to construct the EKS reduced-order models both in the time- and frequency-domain. Specifically, setting $d = k(i + 1)$ and introducing

$$B_d = (-D_d^{1/2} H_d D_d^{-1/2})^{1/2}, \tag{17}$$

where again the principle value for the square root is taken, the time-domain EKS reduced-order model is given by

$$\tilde{u}_d(t) = -\eta(t) \delta_0 \operatorname{Re}[V_d B_d^{-1} \exp(-B_d t)] e_1, \tag{18}$$

where e_1 is the first column of the identity matrix of order d . Its frequency-domain counterpart is given by

$$\tilde{u}_d(s) = -\frac{1}{2} \delta_0 [V_d B_d^{-1} (B_d + sI_d)^{-1} + \bar{V}_d \bar{B}_d^{-1} (\bar{B}_d + sI_d)^{-1}] e_1. \tag{19}$$

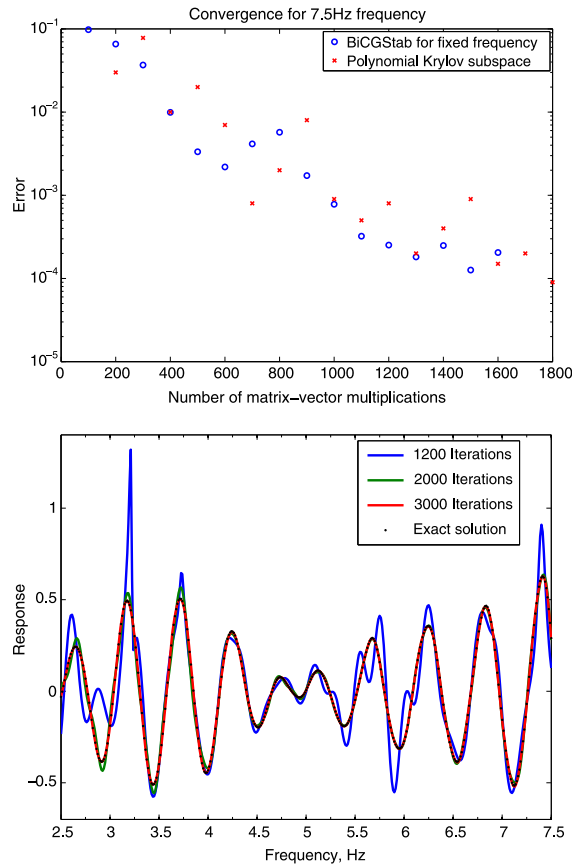


Fig. 2. SEG/EAGE Salt model. Convergence rates for the problem with a fixed frequency of 7.5 Hz (top) and convergence behavior of PKS for a frequency range running from 2.5 Hz to 7.5 Hz (bottom). After 2000 iterations the solution has converged for high frequencies only, while after 3000 iterations the solution completely coincides with the exact solution on the entire frequency range.

5. Numerical experiments

To illustrate the performance of the EKS method in the time- and frequency-domain, two sets of experiments are presented. In the first set an acoustic wavefield problem in the frequency-domain is considered, while in the second set we consider an electromagnetic photonic waveguide problem in the time-domain.

For the acoustic problem, we have considered a 3D SEG/Salt model and benchmarked our solver against an independently developed unpreconditioned BiCGStab algorithm that is a competitive conventional iterative Helmholtz solver (see [16] for details). Here we took the case $m_1 = 1$ corresponding to PKS. Fig. 1 shows the solutions at different receiver locations for a fixed frequency of 7.5 Hz. Both the BiCGStab and PKS model reduction solutions coincide with each other very well. When comparing the number of iterations, we note that the PKS approach requires one matrix-vector multiplication per iteration, while BiCGStab needs two. Since this part constitutes the most computationally intensive part, it makes sense to compare the performance of the two methods in terms of matrix-vector multiplications rather than in terms of iterations. To this end, we have plotted the convergence rates of BiCGStab and model order reduction against the number of matrix-vector multiplications in Fig. 2 (top) again for a frequency of 7.5 Hz. We observe that for this single frequency problem both rates are rather close. In Fig. 2 (bottom), we plotted the convergence of Krylov approximations of a fixed order for different frequencies within a frequency range running from 2.5 Hz to 7.5 Hz. We observe that the solution obtained after 3000 iterations is virtually indistinguishable from the exact solution on the entire frequency range. On the other hand, after 2000 iterations the solution has converged for high frequencies only, while for small frequencies there is still a noticeable error. We remark that the BiCGStab-based approach converges similarly: convergence is faster for higher frequencies and slower for lower ones. At this point, we stress that the principal difference between the PKS method and BiCGStab is that the former approach obtains solutions for the whole given frequency range at the convergence cost of BiCGStab for the lowest frequency (from that range).

Subsequently, we consider what adding negative powers gives us in terms of performance. In Fig. 3 (top), we have plotted convergence curves (with respect to increasing k) for approximants obtained using EKS $\mathbb{K}_{k,ki+1}$ for different values of fixed i ($i = \infty$ corresponds to PKS) and for a frequency range that runs from 2.5 Hz to 7.5 Hz. From this figure it is clear that adding negative powers ($i < \infty$) visibly improves EKS convergence compared with the PKS method. In terms of computation

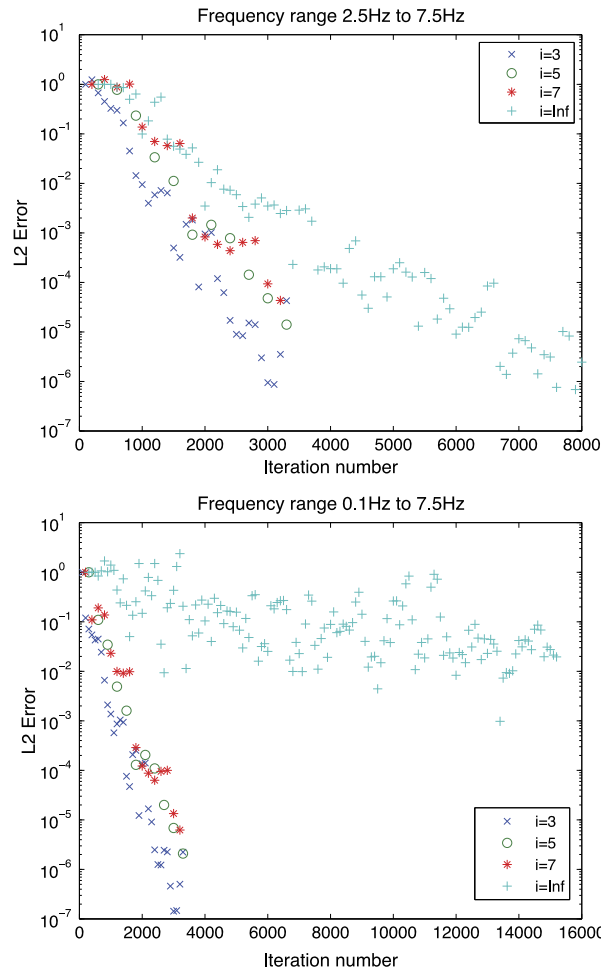


Fig. 3. SEG/EAGE Salt model. Convergence of EKS for different values of i and for frequency ranges from 2.5 Hz to 7.5 Hz (top) and from 0.1 Hz to 7.5 Hz (bottom).

time, it took PKS 157 500 seconds (4500 iterations) to reach 10^{-4} accuracy, while EKS with $i = 3$, $i = 5$, and $i = 7$ required 154 000 seconds (2300 iterations), 189 000 seconds (3000 iterations), and 156 000 seconds (300 iterations), respectively, to reach the same level of accuracy. From these experiments, it is clear that PKS converges faster for higher frequencies, while EKS improves convergence for smaller frequencies and, consequently, makes the convergence more uniform for the whole frequency range. This statement is confirmed in Fig. 3 (bottom), where we have plotted the convergence curves for the EKS method with the same subspace parameters as before, but this time for a frequency range that runs from 0.1 Hz to 7.5 Hz. We observe that PKS significantly slows down in this case, while EKS performs almost as efficient as for the narrower frequency range. For example, to reach a 10^{-4} accuracy, it takes about the same time for EKS to converge, while PKS hardly reached 10^{-2} accuracy in 525 000 seconds of computational time (15 000 iterations).

In our second set of experiments, we consider E-polarized time-domain electromagnetic wave field propagation in the two-dimensional configuration shown in Fig. 4. This configuration was also considered in [7], where the PKS method was used to compute the field distributions. We revisit this problem here so that we can compare the performances of EKS and PKS. As an independent check, we also compare our results with FDTD computations. Since FDTD is based on the first-order Maxwell system, we use a first-order formulation of the EKS method in our time-domain experiments as well.

The configuration of Fig. 4 consists of a set of dielectric rods with vacuum as a background medium. The rods have a relative permittivity of 11.56 and the distance a between the rods is approximately $0.58 \mu\text{m}$. The radius of each rod is $0.18a$. By removing rods inside this crystal we create a bend as is illustrated in Fig. 4. A source is placed at the corner of the bend and we record the electric field strength halfway one of the bends. The source pulse is a modulated Gaussian with its spectrum essentially contained in the frequency interval $9.8 \cdot 10^{14} \leq \omega \leq 1.44 \cdot 10^{15}$ rad/s, which coincides with the ideal frequency bandgap of this structure.

In Fig. 5, we present the electric field response at the receiver location for a diagonal extended Krylov subspace ($m_1 = m_2$) after 100 iterations with the system matrix and 100 iterations with its inverse (black line). The field response obtained after 8197 iterations of the FDTD method run at the Courant limit is also shown in Fig. 5. Clearly, the EKS reduced-order

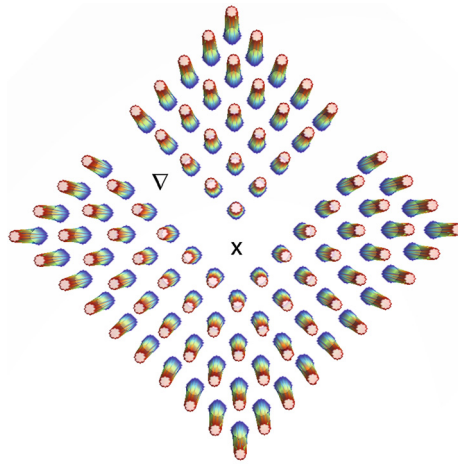


Fig. 4. Crystal consisting of dielectric rods. The cross and the triangle indicate the location of the source and receiver, respectively. The rods have a relative permittivity of 11.56. The distance a between the rods is $0.58 \mu\text{m}$ and the radius of each rod is $0.18a$.




Fig. 5. Electric field strength at the receiver location. Black line: diagonal EKS reduced-order model after 100 iterations with the system matrix and 100 iterations with its inverse. Grey line: FDTD simulation obtained after 8197 iterations at the Courant limit.

model has not converged yet on the time interval of interest. We therefore double the number of iterations and carry out 200 iterations with the system matrix and 200 with its inverse. The resulting EKS reduced-order model for the electric field strength is plotted in Fig. 6 along with the corresponding FDTD result. Even though there is some improvement, the reduced-order model still has not converged on our time interval of observation. Finally, carrying out 400 iterations with the system matrix and 400 with its inverse we obtain the reduced-order model as shown in Fig. 7 (black line). The model essentially overlaps with the FDTD result in this case and the number of iterations required by the EKS method is clearly much smaller than the number of FDTD iterations (800 for EKS, 8197 for FDTD). However, the amount of work that is involved at each iteration is not the same of course, since EKS requires matrix inversion, while FDTD does not. We therefore compare the computation times of both methods as well and find that FDTD requires 1764 seconds to solve this problem, while the EKS method requires only 888 seconds. In terms of offline costs, for this example we have achieved a speed up of approximately a factor of two compared with FDTD. If we compare this data with the results presented in [7], we observe that EKS is actually slower than PKS in this case, since according to [7], PKS requires only 606 seconds to solve this problem. However, the order of the PKS reduced-order model is 3000 (see [7]), which is much larger than the order of the EKS reduced-order model (which is 800). We conclude that for this problem both EKS and PKS are faster than FDTD. When reduction of the offline costs is important then PKS should be used, since it has the shortest computation time, but if the online costs are dominant then the EKS model should be used, since it provides us with the smallest reduced-order model.




Fig. 6. Electric field strength at the receiver location. Black line: diagonal EKS reduced-order model after 200 iterations with the system matrix and 200 iterations with its inverse. Grey line: FDTD simulation obtained after 8197 iterations at the Courant limit.




Fig. 7. Electric field strength at the receiver location. Black line: diagonal EKS reduced-order model after 400 iterations with the system matrix and 400 iterations with its inverse. Grey line: FDTD simulation obtained after 8197 iterations at the Courant limit.

6. Conclusions

Based on model reduction techniques, we have developed a powerful tool for solving large-scale multi-frequency wave problems. The PKS method allows us to obtain wave field solutions on a whole frequency range at the cost of an unpreconditioned linear solver applied to the problem with the smallest frequency. Moreover, we have shown that EKSs significantly outperform the polynomial approach when solutions for small frequencies are required and we have demonstrated that the EKS approach allows for substantial dimensionality reduction in the time-domain. The time-domain models produced by EKS have a much smaller order than the models produced by PKS. On the other hand, subspace creation via EKS takes more time than subspace creation via PKS. Based on these findings, we conclude that PKS should be used if reduction of the offline computational costs is desired, while EKS is the method of choice when the online costs need to be minimized.

Finally we note that our approach is not limited to second-order schemes in the interior. In fact, since the optimal discrete PML has spectral accuracy (see [7,3]), it would be preferable to use a discretization scheme that has spectral accuracy in the interior as well. High-order spectral methods can be applied, for example, (see [11]) or optimal grids for interior domains as described in [1].

Acknowledgements

We'd like to thank Dr. Guangdong Pan for providing us with results for his BiCGStab finite-difference solver.

References

- [1] S. Asvadurov, V. Druskin, L. Knizhnerman, Application of the difference Gaussian rules to the solution of hyperbolic problems, *J. Comput. Phys.* 158 (2000) 116–135.
- [2] J.P. Berenger, A perfectly matched layer for the absorption of electromagnetic waves, *J. Comput. Phys.* 114 (1994) 185–200.
- [3] V. Druskin, S. Guettel, L. Knizhnerman, Near-optimal perfectly matched layers for indefinite Helmholtz problems, University of Manchester, 2013, MIMS preprint 2013.53.
- [4] V. Druskin, L. Knizhnerman, Two polynomial methods of calculating functions of symmetric matrices, *USSR Comput. Math. Math. Phys.* 29 (1989) 112–121.
- [5] V. Druskin, L. Knizhnerman, Extended Krylov subspaces: approximation of the matrix square root and related functions, *SIAM J. Matrix Anal. Appl.* 19 (3) (1998) 755–771.
- [6] V. Druskin, L. Knizhnerman, M. Zaslavsky, Solution of large scale evolutionary problems using rational Krylov subspaces with optimized shifts, *SIAM J. Sci. Comput.* 31 (5) (2009) 3760–3780.
- [7] V.L. Druskin, R.F. Remis, A Krylov stability-corrected coordinate stretching method to simulate wave propagation in unbounded domains, *SIAM J. Sci. Comput.* 35 (2) (2013) B376–B400.
- [8] V.L. Druskin, V. Simoncini, Adaptive rational Krylov subspaces for large-scale dynamical systems, *Syst. Control Lett.* 60 (2011) 546–560.
- [9] T. Goeckler, V. Grimm, Convergence analysis of an extended Krylov subspace method for the approximation of operator functions in exponential integrators, *SIAM J. Numer. Anal.* 51 (4) (2013) 2189–2213.
- [10] G.H. Golub, C.F. Van Loan, *Matrix Computations*, 4th edition, The Johns Hopkins University Press, Baltimore, 2013.
- [11] J.S. Hesthaven, S. Gottlieb, D. Gottlieb, *Spectral Methods for Time-Dependent Problems*, Cambridge University Press, Cambridge, 2007.
- [12] P.D. Hislop, I.M. Sigal, *Introduction to Spectral Theory*, Springer, New York, 1996.
- [13] C. Jagels, L. Reichel, The extended Krylov subspace method and orthogonal Laurent polynomials, *Linear Algebra Appl.* 431 (2009) 441–458.
- [14] C. Jagels, L. Reichel, Recursion relations for the extended Krylov subspace method, *Linear Algebra Appl.* 434 (2011) 1716–1732.
- [15] P. Lax, R. Phillips, *Scattering Theory*, second edition, Academic Press, London, 1989.
- [16] G. Pan, A. Abubakar, T. Habashy, An effective perfectly matched layer design for acoustic fourth-order frequency-domain finite-difference scheme, *Geophys. J. Int.* 188 (2012) 211–222.
- [17] M. Reed, B. Simon, *Methods of Modern Mathematical Physics: Functional Analysis*, vol. 1, Academic Press, San Diego, 1980.
- [18] V. Simoncini, A new iterative method for solving large-scale Lyapunov matrix equations, *SIAM J. Sci. Comput.* 29 (3) (2007) 1268–1288.
- [19] A. Taflov, S.C. Hagness, *Computational Electrodynamics. The Finite-Difference Time-Domain Method*, third edition, Artech House, Boston, 2005.
- [20] L.N. Trefethen, *Approximation Theory and Approximation Practice*, SIAM, Philadelphia, 2013.
- [21] M. Zaslavsky, V. Druskin, L. Knizhnerman, Solution of 3D time-domain electromagnetic problems using optimal subspace projection, *Geophysics* 76 (2011) F339–F351.