

COMPRESSING LARGE-SCALE WAVE PROPAGATION MODELS VIA PHASE-PRECONDITIONED RATIONAL KRYLOV SUBSPACES*

VLADIMIR DRUSKIN[†], ROB F. REMIS[‡], MIKHAIL ZASLAVSKY[†], AND
JÖRN T. ZIMMERLING[‡]

Abstract. Rational Krylov subspace (RKS) techniques are well-established and powerful tools for projection-based model reduction of time-invariant dynamic systems. For hyperbolic wavefield problems, such techniques perform well in configurations where only a few modes contribute to the field. RKS methods, however, are fundamentally limited by the Nyquist–Shannon sampling rate, making them unsuitable for the approximation of wavefields in configuration characterized by large travel times and propagation distances, since wavefield responses in such configurations are highly oscillatory in the frequency-domain. To overcome this limitation, we propose to precondition the RKSs by factoring out the rapidly varying frequency-domain field oscillations. The remaining amplitude-functions are generally slowly varying functions of source position and spatial coordinate and allow for a significant compression of the approximation subspace. Our one-dimensional analysis together with numerical experiments for large-scale two-dimensional acoustic models shows superior approximation properties of preconditioned RKS compared with the standard RKS model-order reduction. The preconditioned RKS results in a reduction of the frequency sampling well below the Nyquist–Shannon rate, a weak dependence of the RKS size on the number of inputs and outputs for multiple-input/multiple-output problems, and, most importantly, in a significant coarsening of the finite-difference grid used to generate the RKS. A prototype implementation indicates that the preconditioned RKS algorithm is competitive in the modern high performance computing environment.

Key words. model reduction, wave equation, Nyquist rate, geometrical optics, seismic exploration

AMS subject classifications. 65M60, 65M80, 93B11, 37M05, 78A05

DOI. 10.1137/17M1156848

1. Introduction. Numerical modeling of wave propagation is fundamental to many applications in design optimization and wavefield imaging. In the oil and gas industry, for instance, the solution of the Maxwell equations is required to invert electromagnetic measurements, while in seismic imaging the solution to the elastodynamic wave equation is needed to ultimately image the subsurface of the Earth.

Finite-difference discretization of the governing wave equations leads to large-scale linear systems, whose solution is computationally intense. Imaging and optimization often use multiple frequencies, sources, and receivers, which leads to systems that need to be evaluated for multiple right-hand sides, time steps, or frequencies, depending on whether the problem is solved in the time- or frequency-domain. Therefore, these so-called multiple-input/multiple-output (MIMO) systems have a high demand on memory and computational power, causing long runtimes. To be more specific, let us consider a surface seismic imaging problem in a k dimensional space ($1 \leq k \leq 3$),

*Received by the editors November 14, 2017; accepted for publication (in revised form) July 23, 2018; published electronically October 4, 2018.

<http://www.siam.org/journals/mms/16-4/M115684.html>

Funding: This research was supported by the Dutch Technology Foundation STW (project 14222), which is part of the Netherlands Organisation for Scientific Research (NWO), and which is partly funded by the Ministry of Economic Affairs.

[†]Schlumberger-Doll Research, Cambridge, MA 19104-2688 (druskin@slb.com, mzaslavsky@slb.com).

[‡]Circuits and Systems Group, Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, Mekelweg 4, 2628 CD Delft, The Netherlands (r.f.remis@tudelft.nl, jzimmerl@umich.edu).

with maximal propagation distance of N wavelengths. This would require the solution of a discretized system with $O(N^k)$ state variables, $O(N^{k-1})$ sources and receivers, and $O(N)$ frequencies or time steps [22]. Model-order reduction aims to reduce the complexity and computational burden of large-scale problems and here we target all three of these factors.

Recently, promising results were obtained in the time-domain via multiscale model reduction [8, 10]. The time-domain multiscale algorithms can be efficiently parallelized via domain-decomposition, but time stepping still needs to be carried out sequentially, while frequency-domain problems can be solved in parallel for different frequencies. Here we consider interpolatory projection-based model reduction in the frequency-domain; e.g., see [1]. The essence of this approach is the projection of the underlying system onto a rational Krylov subspace (RKS; originally introduced by Ruhe for eigenvalue computations [24]), which produces good-quality, low-order approximations if the spectrum of the system is well separated from the frequency interval of interest, as in the case of diffusion PDEs, e.g., [2, 13, 20, 16]. In the context of wavefield modeling, such a separation of the clustered eigenvalues is introduced by losses present in the media or by the use of absorbing boundary conditions for the truncation of unbounded domains. Projection-based reduced-order models (ROMs) for wavefield problems may therefore exhibit fast convergence, especially for resonant configurations with few isolated resonant eigenmodes [5, 11, 12]. Some modifications of the RKS projection method can also be competitive for problems with smooth initial conditions leading to effective suppression of highly oscillatory eigenmodes [15].

Usually, the computational cost of projection-based reduced-order modeling is dominated by the generation of a suitable projection basis, e.g., see [4, 23]. In the interpolatory projection-based ROM the Helmholtz equation has to be solved at different frequencies (shifts) and the span of these solutions forms the RKS basis. The solution obtained from the Galerkin projection onto this subspace interpolates at the shifts, which are therefore also known as interpolation points. Moreover, for coinciding sources and receivers the transfer-function and its first derivative are interpolated at these points. A general drawback of an RKS approach is that the number of interpolation points can become large when wavefield solutions with large travel times or propagation distances are of interest. Such wavefields are highly oscillatory in the frequency-domain and the Nyquist–Shannon sampling theorem states that this oscillatory field should be sampled with at least one point per oscillation (two points per wavelength (ppw)). Consequently, the number of interpolation points required to accurately represent the wavefield increases with the propagation distance. Moreover, discretization grids in Helmholtz solvers must also resolve wavefield oscillations. This requirement has an even more dramatic effect on the computational cost due to poor scalability of the available solvers. In favorable situations the best sampling rates approaching the Nyquist limit can be achieved with high-order spectral methods and their outgrowths. However, their cost per unknown can be significantly higher compared with less accurate low-order methods due to loss of sparsity. In this paper, we show that the sampling demand can be significantly lowered by adding phase information to the model-order reduction technique leading to phase-preconditioned RKSs (PPRKS). Preconditioning of Krylov subspaces for model reduction is a tough and still open problem in general. However, to achieve it for particular applications one can try to incorporate the underlying physics and asymptotic analysis to arrive at PPRKS. Our approach is related to other known approaches in the field of oscillatory wave problem computation, such as preconditioners for Helmholtz solvers [14, 17], Filon quadrature [19], and a recent approach to data compression using phase-tracking [21].

In particular, we construct RKSs using polar decompositions of frequency-dependent basis functions. These decompositions consist of a product of *smooth* amplitude-functions and a known frequency-dependent *oscillatory* phase term. The phase term is determined from high-frequency asymptotic expansions such as the Wentzel–Kramers–Brillouin (WKB) approximation [3]. The amplitude-functions are computed by splitting the RKS into incoming and outgoing waves (by applying one-way wave operators) and factoring out the corresponding phase terms. Analogous to Filon quadrature, we handle the highly oscillatory phase-functions analytically and the smooth amplitude-function numerically. By developing a block version of PPRKS for MIMO problems, we are also able to factor out the main dependence of the RKS on the input (source) location. This feature, and the reduction of the number of interpolation points mentioned above, leads to a significant compression of the approximation space.

Finally, the resulting phase-preconditioned ROMs can also extrapolate to frequencies outside the interval of interpolation points, since the basis functions are frequency-dependent and the amplitude-functions are smooth for smoothly varying wave speed profiles. This enables us to coarsen the second-order finite-difference grid used for the RKS generation.

In conclusion, with PPRKS we can effectively reduce all of the above mentioned factors contributing to the complexity of the MIMO wavefield problem. The overall goal is to approximate the transfer-functions from multiple sources to multiple receivers with a small ROM that honors the physics of the underlying wave equation. The approach uses a coarse grid and low-frequency interpolation points to build an RKS and to obtain smooth amplitude-functions. Using high-frequency asymptotic expansions, this RKS is extrapolated to high frequencies and evaluated on a fine grid. The projection of a fine grid wave operator onto the extrapolated RKS gauges the ROM to the fine scale we intend to model. In this way, fine-scale wave scattering and large-scale wave propagation can be combined, which allows us to obtain a ROM valid for all time scales. RKS algorithms for wavefield problems are at a disadvantage compared with polynomial and extended Krylov subspace algorithms when it comes to computational memory consumption as the basis needs to be saved for RKS methods. The compression of the approximation space to a small number of amplitudes and phases, however, leads to a reduction in the computational memory demand of the proposed method.

In section 2, we start with a short discussion on the wave equation and formulate the wavefield problem of interest for a single-input/single-output (SISO) configuration. Subsequently, we introduce a standard RKS in section 3 and construct field approximations in the frequency-domain. We show that this RKS approach is structure-preserving and that the transfer-function of ROMs based on this RKS is a Hermite interpolant of the transfer-function for a coinciding source-receiver pair. In section 4, we take the RKS approach of section 3 as a starting point and introduce the PPRKS for one-dimensional SISO configurations. We show that phase-preconditioning is structure-preserving and retains the interpolation properties of standard RKS Galerkin projection. The main result of this section is that for a piecewise constant wave speed profile, the new method yields the exact solution with the number of interpolation points equal to the number of homogeneous layers, i.e., this number plays the same role as the problem dimensionality in a conventional RKS approach. Section 5 discusses the algorithm for higher spatial dimensions in a MIMO setting using a block version of PPRKS. Finite-difference implementation via a two-grid algorithm is discussed in section 6. In section 7 we illustrate the performance

of the proposed RKS techniques through a number of two-dimensional numerical experiments. Section 8 discusses the implementation of the proposed method on parallel computation architectures, and the conclusions can be found in section 9. Throughout this manuscript, quantities in the time-domain are denoted in roman font, while quantities in the Laplace domain are in italics.

2. Problem formulation. In this paper we address the problem of solving the Green's function for wave equations within a spectral interval of interest. We start the discussion by considering the scalar, isotropic, continuous wave equation on $\mathbb{R}^k \times [0, \infty[$

$$(2.1) \quad \Delta u - \frac{1}{\nu^2} u_{tt} = -\frac{1}{\nu^2} \delta(t) \delta(x - x_S), \quad u|_{t=0} = 0, \quad u_t|_{t=0} = 0.$$

In this equation, Δ denotes the k -dimensional Laplace operator and the position vector is $x \in \mathbb{R}^k$ ($1 \leq k \leq 3$). Furthermore, $\nu(x) > 0$ is a wave speed distribution in $L^\infty[\mathbb{R}^k]$, and $u(x, t)$ is the wavefield with a compact support for all finite times.

After Laplace transformation, (2.1) becomes

$$(2.2) \quad \Delta u - \frac{s^2}{\nu^2} u = -\frac{1}{\nu^2} \delta(x - x_S),$$

where s is the complex Laplace parameter with $\Re(s) \geq 0$. The Laplace domain wavefield u satisfies the limiting absorption principle, i.e., u vanishes at infinity for $\Re(s) > 0$ and converges to the solution of Helmholtz's equation that satisfies the outgoing radiation condition as the Laplace parameter s approaches the imaginary axis via the right half of the complex s -plane.

Let Ω be a bounded subdomain of \mathbb{R}^k such that $x_S \in \Omega$. We now equivalently reduce the original problem on the unbounded domain to a problem on Ω by considering (2.2) in the weak formulation and testing this equation with a testing function p . This gives

$$(2.3) \quad \int_{\Omega} \bar{p} \left(\Delta - \frac{s^2}{\nu^2} \right) u \, dx = -\frac{1}{\nu(x_S)^2} \bar{p}(x_S),$$

where the overbar denotes complex conjugation. After integration by parts, we obtain

$$(2.4) \quad - \int_{\Omega} (\nabla \bar{p}) \cdot (\nabla u) \, dx - \int_{\Omega} \bar{p} \frac{s^2}{\nu^2} u \, dx + \int_{\partial\Omega} \bar{p} \frac{\partial u}{\partial n} \, dx = -\frac{1}{\nu(x_S)^2} \bar{p}(x_S)$$

with $\frac{\partial u}{\partial n}$ the derivative of u in the direction of the outward-pointing normal on $\partial\Omega$. Finally, introducing the Dirichlet-to-Neumann (DtN) map $D(s)$ on $\partial\Omega$ such that $\frac{\partial u}{\partial n} = D(s)u$, the above equation can be written as

$$(2.5) \quad - \int_{\Omega} (\nabla \bar{p}) \cdot (\nabla u) \, dx - \int_{\Omega} \bar{p} \frac{s^2}{\nu^2} u \, dx + \int_{\partial\Omega} \bar{p} D(s) u \, dx = -\frac{1}{\nu(x_S)^2} \bar{p}(x_S).$$

Without the boundary integral (third term on the left-hand side of the above equation) this equation is linear in s^2 ; the DtN map, however, is a nonlinear function of frequency s [9].

Notation. To better draw similarities between continuous and discrete formulations, we will treat the complex-valued functions u and p as vectors from \mathbb{R}^∞ in our linear algebraic derivations and introduce the inner product

$$(2.6) \quad p^H u = \int_{\Omega} \bar{p} u \, dx.$$

We note that u and p for $k > 1$ have singularities at x_S that may make this inner product divergent. To avoid this, we assume by default that instead of $\delta(x - x_S)$ we have some regular approximation of the delta function. After discretization, u and p become finite-dimensional vectors from \mathbb{R}^N and the issue of diverging integrals due to singularities disappears. In this notation, superscript H denotes the Hermitian transpose for vectors and an inner product with complex conjugation for functions. Operators are as capital italic letters like A , and for linear combinations such as

$$q_m = \alpha_1 g^{[1]} + \alpha_2 g^{[2]} + \cdots + \alpha_m g^{[m]},$$

with coefficients α_i and expansion functions $g^{[i]}$, we write $q_m = G_m \mathbf{z}$ with $\mathbf{z} = [\alpha_1, \dots, \alpha_m]^T$, and the expansion functions are stored as columns in the function array G_m , i.e., $G_m \in \mathbb{R}^{\infty \times m}$ (sometimes called a quasi matrix [27]). Finally, finite-dimensional matrices are given in a capital sans serif font (like A).

Using the notation outlined above, we now introduce the wave operator $Q(s)$ to rewrite (2.5) as

$$(2.7) \quad p^H Q(s) u = -\frac{1}{\nu(x_S)^2} \bar{p}(x_S).$$

First, we note that real and imaginary parts of $Q(s)$ are self-adjoint. In the time-domain, the wavefield is obviously real-valued and consequently operator $Q(s)$ and the field $u(s)$ satisfy the Schwarz reflection principle

$$(2.8) \quad Q(\bar{s}) = \overline{Q(s)} \text{ and } u(\bar{s}) = \overline{u(s)},$$

from which it immediately follows that the spectrum of $Q(s)$ is symmetric under complex conjugation.

Global energy conservation for problem (2.1) leads to passivity of $Q(s)$, which can be defined via its nonlinear numerical range $\mathcal{W}\{Q(s)\}$ (also known as nonlinear field of values); e.g., see [18], [13]. Specifically, for a nonlinear operator-valued function $A(s)$, the nonlinear numerical range $\mathcal{W}\{A(s)\}$ (e.g., see [18], [13]) is defined as

$$(2.9) \quad \mathcal{W}\{A(s)\} = \{s \in \mathbb{C} : x^H A(s) x = 0 \text{ for some nontrivial } x \in \mathbb{C}^k\}.$$

Passivity of dynamic system (2.7) is equivalent to the condition

$$(2.10) \quad \Re \mathcal{W}\{Q(s)\} \leq 0.$$

The nonlinear problem can formally be solved in the Laplace domain as $u = Q^{-1}(s)b$ and quadrature rules for evaluating the Bromwich integral in the inverse Laplace transform can be used to obtain a time-domain solution. To compute the inverse Laplace transform we use equidistant quadrature points on the imaginary s -axis in the spectral region where the wavelet is active. In the area where the Laplace transform of the applied wavelet is small we use a few quadrature points with increasing steps on the imaginary axis.

In the case of imaging or remote sensing applications one is typically not interested in the wavefield in the entirety of Ω but at certain receiver locations only. With r denoting the receiver function, the transfer-function from source to receiver can be defined as

$$(2.11) \quad f(r, b, s) = r^H Q^{-1}(s) b.$$

This is the SISO transfer-function, and since $Q(s)$ is symmetric in the transpose bilinear form we find that $f(r, b, s) = f(b, r, s)$ hold for real receiver functions. Since an imaging application often uses multiple sources and receivers, we define the MIMO transfer-function later. Fundamentally, we try to compress and approximate this transfer-function in this work. In the following sections we discuss a reduced-order modeling technique that approximates the transfer-function and preserves the above-mentioned symmetry properties, the Schwarz reflection principle, and passivity.

3. Structure preserving rational Krylov subspace reduction. As a first step toward an efficient rational Krylov methodology for multifrequency wavefield problems, we construct field approximations or ROMs based on an interpolatory RKS containing single-frequency solutions (snapshots) of the problem as trial and testing space. Specifically, our approach is to define an RKS of order m as

$$(3.1) \quad \mathcal{K}^m(\kappa) = \text{span} \{u(s_1), u(s_2), \dots, u(s_m)\}$$

with m distinct shifts $\kappa = [s_1, \dots, s_m]$ and to use its real form, the RKS

$$(3.2) \quad \mathcal{K}_R^{2m}(\kappa) = \text{span} \{\Re \mathcal{K}^m(\kappa), \Im \mathcal{K}^m(\kappa)\},$$

as a test and trial space. The real and imaginary parts of the snapshots $u(s_i)$ spanning $\mathcal{K}_R^{2m}(\kappa)$ are always linearly independent, since the eigenfunction expansion of the Dirac distribution appearing on the right-hand side of (2.2) has an infinite number of terms. Furthermore, from the symmetry given in (2.8) it follows that $\mathcal{K}^m(\kappa) \subset \mathcal{K}_R^{2m}(\kappa)$ and $\mathcal{K}^m(\bar{\kappa}) \subset \mathcal{K}_R^{2m}(\kappa)$ and a projection onto the subspace $\mathcal{K}_R^{2m}(\kappa)$ will therefore preserve the Schwarz reflection principle leading to real-valued, time-domain wavefield approximations. In the following subsections we will construct the reduced-order wavefield approximations, discuss their structure, and show the interpolation properties of these approximations.

3.1. Reduced-order solution. We start by approximating the weak solution of (2.7) by an element from the space $\mathcal{K}_R^{2m}(\kappa)$. To this end, let the functions $v^{[1]}, v^{[2]}, \dots, v^{[2m]}$ form a real basis $V_m \in \mathbb{R}^{\infty \times 2m}$ of $\mathcal{K}_R^{2m}(\kappa)$. The reduced-order solution is now expanded as $u_m = V_m z$ with expansion coefficients α_i collected in vector $z = [\alpha_1, \dots, \alpha_{2m}]^T$. These coefficients can be obtained from a standard Galerkin procedure defined through the weak form of (2.5) leading to

$$(3.3) \quad z = [V_m^H Q(s) V_m]^{-1} V_m^H b \quad \text{or} \quad z = R_m^{-1}(s) V_m^H b$$

with $b = -\delta(x - x_S)/\nu(x_S)^2$ and where $R_m(s)$ is the $2m \times 2m$ reduced-order operator given by $R_m(s) = V_m^H Q(s) V_m$. The ROM is structure preserving as show in the following proposition.

PROPOSITION 3.1. *The reduced-order operator $R_m(s)$ preserves the structure of the full-order operator $Q(s)$, that is, $R_m(s)$ is symmetric and satisfies the Schwarz reflection principle $R_m(s) = R_m(\bar{s})$ and its numerical range is contained in the numerical range of $Q(s)$, that is, $\mathcal{W}\{R_m(s)\} \subseteq \mathcal{W}\{Q(s)\}$.*

Proof. The symmetry of $R_m(s)$ follows from the symmetry of $Q(s)$ since $R_m(s) = V_m^H Q(s) V_m = V_m^T Q(s)^T V_m = R_m(s)^T$, since V_m is a real. The same argument shows that the Schwarz reflection principle holds as $R_m(\bar{s}) = V_m^H Q(\bar{s}) V_m = V_m^H \bar{Q}(s) V_m = \bar{R}_m(s)$. Moreover, from the definition of the numerical range of the reduced order we find that

$$(3.4) \quad x_m^H R_m(s) x_m = (V_m x_m)^H Q(s) (V_m x_m).$$

Now, every s that satisfies $\mathbf{x}_m^H \mathbf{R}_m(s) \mathbf{x}_m = 0$ also satisfies $(V_m \mathbf{x}_m)^H Q(s) (V_m \mathbf{x}_m) = 0$, such that every point in the numerical range of $\mathbf{R}_m(s)$ is also included in the numerical range of $Q(s)$. \square

Thus, preservation of passivity (2.10) is guaranteed by Proposition 3.1, and so is the preservation of causality and stability. Therefore, existence and uniqueness of the reduced-order solution in (3.3) is also guaranteed by Proposition 3.1, since it guarantees that $\mathbf{R}_m(s)$ is invertible for all Laplace parameters s with $\Re(s) \geq 0$. Finally, we mention that the time-domain counterpart of $u_m(s)$ can be obtained by evaluating the inverse Laplace transform. To compute the inverse Laplace transform, equidistant quadrature points on the imaginary s -axis are used in the spectral region where the wavelet is active. In the area where the Laplace transform of the wavelet is small we use a few quadrature points with increasing steps on the imaginary axis.

We end this section by introducing an alternative way of representing the reduced-order solution, which will be useful in the development of PPRKS methods. In expansion form the reduced-order solution can be written as

$$(3.5) \quad u_m = \sum_{i=1}^m \begin{bmatrix} d_i \\ \delta_i \end{bmatrix}^T \begin{bmatrix} u(s_i) \\ u(\bar{s}_i) \end{bmatrix},$$

where the expansion coefficients d_i and δ_i follow from the Galerkin condition. Due to the above mentioned linear independence of the real and imaginary parts of the snapshots, this representation is algebraically equivalent to $u_m = V_m \mathbf{z}$, i.e., there exists a transform from the $2m$ coefficients α_i to the coefficients d_i and δ_i of (3.5).

3.2. Interpolation properties. The standard theory of Galerkin interpolatory-projection model reduction of passive, self-adjoint, dynamic systems yields the following interpolation properties (e.g., see [1]).

PROPOSITION 3.2. *The projected RKS solution $u_m(s)$ interpolates at the shifts, i.e.,*

$$(3.6) \quad u_m(s) = u(s) \quad \forall s \in \kappa \cup \bar{\kappa},$$

and the SISO reduced-order transfer-function $f_m(s)$ is a Hermite interpolant of the SISO transfer-function $f(s)$ at the shifts, that is,

$$(3.7) \quad f_m(s) = f(s) \quad \text{and} \quad \frac{d}{ds} f_m(s) = \frac{d}{ds} f(s) \quad \text{with } s \in \kappa \cup \bar{\kappa}.$$

Proof. Since $\mathcal{K}^m(\kappa) \subset \mathcal{K}_R^{2m}(\kappa)$ and $\mathcal{K}^m(\bar{\kappa}) \subset \mathcal{K}_R^{2m}(\kappa)$, property (3.6) follows directly from the uniqueness of the Galerkin condition for passive problems. To prove (3.7), we first introduce the field error and residual as

$$e_m(s) = u(s) - u_m(s) \quad \text{and} \quad r_m(s) = b - Q(s)u_m(s),$$

respectively. From the Galerkin condition we obtain the relation

$$(3.8) \quad u_m^H(\bar{s}) r_m(s) = 0,$$

since $u_m(\bar{s}) \in \mathcal{K}_R^{2m}$. The error of the transfer-function can now be written as

$$(3.9) \quad f(s) - f_m(s) = b^H e_m(s) = u^H(\bar{s}) Q(s) e_m(s),$$

where we have used Schwarz's reflection principle. Since $Q(s)e_m(s) = r_m(s)$ and the Galerkin condition of (3.8) holds, we can write

$$(3.10) \quad f(s) - f_m(s) = u^H(\bar{s})r_m(s) = e_m^H(\bar{s})r_m(s),$$

which has double zeros at $s = \kappa \cup \bar{\kappa}$, since the error and residual vanishes for these frequencies due to relation (3.6). \square

The outlined approach is most efficient if only a few singular Hankel values of the system contribute to the solution, as is the case for resonating structures with a few excited and observable modes [4]. Then the frequency-domain response is well-described by a low-degree rational function and a rational Krylov technique will therefore quickly capture the desired wavefield response. For waves characterized by large travel times, however, this may no longer be the case, since such responses are highly oscillatory in the frequency-domain and sampling should at least take place at half the Nyquist–Shannon sampling rate. As an illustration, consider a source-receiver pair with an arrival at T^{arr} such that the source wavelet convolved with $\delta(t - T^{\text{arr}})$ is measured. In the Laplace domain this translates to multiplication by $\exp(-sT^{\text{arr}})$, which means that according to the Nyquist sampling theorem the maximum frequency-domain sampling distance is $\Delta s = \pi/T^{\text{arr}}$ on the imaginary axis. Clearly, the number of required frequency-domain samples increases as the travel time increases, leading to prohibitory large RKSs. In the next section we will incorporate travel time information to obtain basis functions that are less oscillatory to lower this sampling demand.

4. Field parametrization for SISO problems. To enhance the convergence of an RKS approach for travel time dominated structures, we need to incorporate travel time information into the Krylov subspace, and thus into our basis functions. To this end, we assume that variations of the medium take place on a scale much larger than the wavelength at the considered frequencies, since this allows us to use a geometrical optics ansatz. Every basis vector belonging to the RKS is now split into an incoming and an outgoing wave and for each of these waves we factor out a strongly oscillating phase term $\exp(\pm sT_{\text{eik}})$, where $T_{\text{eik}} = T_{\text{eik}}(x)$ is the eikonal time that solves the eikonal equation $|\nabla T_{\text{eik}}(x)|^2 = \frac{1}{\nu(x)^2}$. Splitting of the fields is realized using one-way wave equations. First we introduce this splitting for one-dimensional systems in section 4.1 and then generalize to higher dimensions in section 4.2.

4.1. One-dimensional field parametrization. We decompose the field into an incoming and an outgoing component by writing

$$(4.1) \quad u(s_j) = \exp(-s_j T_{\text{eik}}) c_{\text{out}}(s_j) + \exp(s_j T_{\text{eik}}) c_{\text{in}}(s_j).$$

For each component an oscillating phase term has been factored out and the amplitudes are determined from the single-frequency snapshot solutions $u(s_j)$ via one-way wave equations as

$$(4.2a) \quad c_{\text{out}}(s_j) = \frac{\nu}{2s_j} \exp(s_j T_{\text{eik}}) \left(\frac{s_j}{\nu} u(s_j) - \frac{\partial}{\partial |x - x_S|} u(s_j) \right)$$

and

$$(4.2b) \quad c_{\text{in}}(s_j) = \frac{\nu}{2s_j} \exp(-s_j T_{\text{eik}}) \left(\frac{s_j}{\nu} u(s_j) + \frac{\partial}{\partial |x - x_S|} u(s_j) \right).$$

In (4.2a), the incoming wave component of $u(s_j)$ is filtered out, leaving an outgoing component for which outgoing oscillations can be factored out. In (4.2b) the situation is reversed and the outgoing component of $u(s_j)$ is filtered out. Finally, we note that using the above one-way wave equations for decomposition is equivalent to enforcing the condition

$$(4.3) \quad \exp(s_j T_{\text{eik}}) \frac{\partial}{\partial |x - x_S|} c_{\text{out}}(s_j) + \exp(-s_j T_{\text{eik}}) \frac{\partial}{\partial |x - x_S|} c_{\text{in}}(s_j) = 0,$$

and the amplitudes c_{out} and c_{in} are spatially much smoother than the wavefield u , since the highly oscillatory phase term has been factored out.

Now to obtain a field approximation at frequency s , instead of projecting our operator onto single-frequency solutions $u(s_j)$, we project it onto the phase-corrected basis functions $\exp(-s T_{\text{eik}}) c_{\text{out}}(s_j)$ and $\exp(s T_{\text{eik}}) c_{\text{in}}(s_j)$. This is the central idea of our approach, which preserves the interpolation properties of the RKS. In particular, by introducing the phase-preconditioned subspace as

$$(4.4) \quad \mathcal{K}_{\text{EIK}}^{2m}(\kappa, s) = \text{span}\{\exp(-s T_{\text{eik}}) c_{\text{out}}(s_1), \dots, \exp(-s T_{\text{eik}}) c_{\text{out}}(s_m), \\ \exp(s T_{\text{eik}}) c_{\text{in}}(s_1), \dots, \exp(s T_{\text{eik}}) c_{\text{in}}(s_m)\},$$

and its symmetry-preserving real form

$$(4.5) \quad \mathcal{K}_{\text{EIK};\text{R}}^{4m}(\kappa, s) = \text{span}\{\Re \mathcal{K}_{\text{EIK}}^{2m}(\kappa, s), \Im \mathcal{K}_{\text{EIK}}^{2m}(\kappa, s)\}$$

we can construct ROMs in the usual way, but now in terms of frequency-*dependent* basis functions. More precisely, let $M \leq 4m$ be the dimension of $\mathcal{K}_{\text{EIK};\text{R}}^{4m}(\kappa, s)$ and let vectors $v^{[1]}(s), v^{[2]}(s), \dots, v^{[M]}(s) \in \mathbb{R}^\infty$ form an orthonormal basis of $\mathcal{K}_{\text{EIK};\text{R}}^{4m}$; then the field approximation drawn from this subspace can be written as

$$(4.6) \quad u_m(s) = \sum_{i=1}^M \alpha_i(s) v^{[i]}(s)$$

and the coefficients $\alpha_i(s) \in \mathbb{C}$ can again be determined from the Galerkin condition. Note that m denotes the number of snapshots used to construct $\mathcal{K}_{\text{EIK};\text{R}}^{4m}$, while $M \leq 4m$ denotes the dimension of this subspace. The factor of 4 in the upper bound on M is due to splitting into incoming and outgoing fields, which can lead to a twice-as-large approximation subspace compared with unpreconditioned RKS with the same shifts. However, as we shall see in subsection 5.2, this increase can be circumvented in our implementation; the overall dimension of the preconditioned RKS is usually comparable to the dimension of a standard unpreconditioned RKS for the same accuracy, while using a smaller number of snapshots.

With $V_{m;\text{EIK}}(s) \in \mathbb{R}^{\infty \times M}$ the real, orthonormal basis matrix of $\mathcal{K}_{\text{EIK};\text{R}}^{4m}(\kappa, s)$, the ROM that follows from the Galerkin condition can be written as a self-adjoint, time-invariant dynamic system

$$(4.7) \quad V_{m;\text{EIK}}(s) \mathbf{R}_{m;\text{EIK}}(s) V_{m;\text{EIK}}^H(s) u_{m;\text{EIK}}(s) = b_m$$

with

$$b_m = V_{m;\text{EIK}}(s) V_{m;\text{EIK}}^H(s) b \quad \text{and} \quad \mathbf{R}_{m;\text{EIK}}(s) = V_{m;\text{EIK}}^H(s) Q(s) V_{m;\text{EIK}}(s).$$

The ROM of (4.7) is the phase-corrected counterpart of the ROM of (3.3). Furthermore, since $\bar{c}_{\text{out/in}}(s_i) = c_{\text{out/in}}(\bar{s}_i)$ holds because the Schwarz reflection principle is

satisfied, we can also express the ROM of (4.6) in terms of the amplitude-functions $c_{\text{in}}(s)$ and $c_{\text{out}}(s)$ as (cf. (3.5))

$$(4.8) \quad u_m(s) = \sum_{i=1}^m \begin{bmatrix} a_i(s) \\ \alpha_i(s) \end{bmatrix}^T \begin{bmatrix} \exp(-sT_{\text{eik}}) c_{\text{out}}(s_i) \\ \exp(-\bar{s}T_{\text{eik}}) c_{\text{out}}(\bar{s}_i) \end{bmatrix} + \sum_{i=1}^m \begin{bmatrix} d_i(s) \\ \delta_i(s) \end{bmatrix}^T \begin{bmatrix} \exp(sT_{\text{eik}}) c_{\text{in}}(s_i) \\ \exp(\bar{s}T_{\text{eik}}) c_{\text{in}}(\bar{s}_i) \end{bmatrix}$$

with expansion coefficients a_i , α_i , d_i , and $\delta_i \in \mathbb{C}$ and where we have assumed that $M = 4m$. This formulation clearly shows that we use frequency-independent amplitudes, preconditioned by frequency-dependent phase-functions, employing conjugation to preserve the symmetry of the wave equation.

Finally, for field evaluations on the imaginary axis ($s \in i\mathbb{R}$) the above expansion can be written more compactly as

$$(4.9) \quad u_m(s) = \sum_{i=1}^{2m} \begin{bmatrix} a_i(s) \\ \alpha_i(s) \end{bmatrix}^T \begin{bmatrix} \exp(-sT_{\text{eik}}) c(s_i) \\ \exp(sT_{\text{eik}}) c(-s_i) \end{bmatrix} \quad \text{with } s \in i\mathbb{R}$$

and where $c(s_i) = c_{\text{out}}(s_i)$ for $i = 1, 2, \dots, m$ and $c(s_i) = \bar{c}_{\text{in}}(s_{i-m})$ for $i = m + 1, \dots, 2m$.

The following results show that the PPRKS retains the structure-preserving interpolatory-projection properties of standard RKS.

LEMMA 4.1. *The system of (4.7) is structure preserving, i.e., $W\{\mathbf{R}_{m;\text{EIK}}(s)\} \subseteq W\{Q(s)\}$ on the range (column space) of $V_{m;\text{EIK}}(s)$.*

Proof. Let a nontrivial \mathbf{x}_m be in the range of $V_{m;\text{EIK}}(s)$, that is $\mathbf{x}_m = V_{m;\text{EIK}}(s)\mathbf{y}_m$. Then

$$\mathbf{y}_m^H \mathbf{R}_{m;\text{EIK}}(s) \mathbf{y}_m = \mathbf{y}_m^H V_{m;\text{EIK}}^H(s) Q(s) V_{m;\text{EIK}}(s) \mathbf{y}_m = \mathbf{x}_m^H Q(s) \mathbf{x}_m. \quad \square$$

Thus, phase-preconditioned ROMs can restrict the numerical range as \mathbf{x}_m is in the range of $V_{m;\text{EIK}}(s)$; however, the spectrum is always contained in the projected operator.

PROPOSITION 4.2. *The SISO reduced-order transfer-function retains the interpolation properties of the unpreconditioned RKS with the same shifts stated in Proposition 3.2.*

Proof. By construction, $\mathcal{K}_{\text{EIK},\text{R}}^{4m}(\kappa, s) \supset \mathcal{K}_R^{2m}(\kappa)$ when $s \in \kappa \cup \bar{\kappa}$. According to Lemma 4.1, the ROM is passive given that the Galerkin problem has a unique solution. Therefore, the proof of Proposition 3.2 applies. \square

One of the motivations to use this method is the expected fast convergence, when the parametrization of (4.1) is valid. In that case only a few phase-corrected, smooth amplitude-functions $c_{\text{out/in}}$ are required to approximate the wavefield. Furthermore, in the RKS method discussed in the previous section, the number of required shifts or frequencies is dependent on the largest arrival time; however, in the PPRKS discussed above, the arrival times are factored out and the number of shifts is dependent on the complexity of the wave speed model $\nu(x)$ rather than the largest arrival time. We make this explicit in the following proposition.

PROPOSITION 4.3. *Let a one-dimensional problem have ℓ homogeneous layers. Then there exist $m \leq \ell + 1$ noncoinciding interpolation points such that $u_{m;\text{EIK}}(s) = u$.*

Proof. We start by noting that if the regions to the left and to the right of the source are considered as separate layers, then the solution to the one-dimensional wave equation consists of a superposition of left- and right-going waves of constant

amplitude in each of the $\ell + 1$ layers. For one-dimensional problems, the decomposition direction coincides with the travel direction of the wave; thus, $c_{\text{out}}(\kappa)$ and $c_{\text{in}}(\kappa)$ are piecewise constant. A piecewise constant function with $\ell + 1$ layers can be exactly represented by at most $\ell + 1$ linear-independent piecewise constant functions with the same jump locations. Let us prove from the opposite, i.e., assume that there are no $m \leq \ell + 1$ noncoinciding shifts κ_i such that $c_{\text{out}}(\kappa_i)$ form a basis for all possible $c_{\text{out}}(\kappa)$. Then the number of shifts m yielding linear-independent solutions should be less than $\ell + 1$. But by assumption there must be at least a single $c_{\text{out}}(\kappa)$ not from the subspace. Then one can add this solution to the subspace, i.e., the true number of linearly independent solutions is $m + 1$, which contradicts the assumption that this number is m . Analogously, we can prove the same statement for $c_{\text{in}}(\kappa_i)$.

In conclusion, we proved that there exist $m \leq \ell + 1$ noncoinciding shifts κ_i such that $c_{\text{out}}(\kappa_i)$ and $c_{\text{in}}(\kappa_i)$ form respective bases for all possible $c_{\text{out}}(\kappa)$ and $c_{\text{in}}(\kappa)$, i.e., the exact solution u will be in the projection subspace. Finally, due to Lemma 4.1 the exact solution u will be the unique solution of the Galerkin problem. \square

The proposition can be extended to almost all arbitrary $\ell + 1$ interpolation points, as the interpolation points that lead to $\ell + 1$ linear-dependent functions have measure zero. Thus, phase-preconditioning allows us to obtain the exact solution with the number of interpolation points equal to the number of homogeneous layers, i.e., this number plays the same role as the problem dimensionality in a conventional RKS approach.

Remark 1. Interpretation as Filon quadrature. Filon quadrature deals with the evaluation of oscillatory integrands. Using a standard quadrature to evaluate

$$(4.10) \quad y(s) = \int \exp(st) f(t) dt$$

with a smooth function $f(t)$ that varies slowly compared to $\exp(st)$ leads to

$$(4.11) \quad y(s) = \Delta t \sum_n a_n \exp(sn\Delta t) f(n\Delta t)$$

with quadrature weights a_n which requires $|s\Delta t| < \pi$. In Filon quadrature one makes the weights a_n dependent on s to arrive at quadrature formulas whose convergence depends merely on the properties of $f(t)$. Similarly, in phase-preconditioned rational Krylov we try to approximate an oscillatory frequency-domain transfer-function. By making our projection basis frequency-dependent we enhance the convergence of the ROM.

4.2. Generalization to higher dimensions. In higher spatial dimensions we again split the field in incoming and outgoing wave components and use a function $g(z)$ to factor out a strongly varying phase. Specifically, for two- and three-dimensional problems we write

$$(4.12) \quad u(s_j) = g(s_j T_{\text{eik}}) c_{\text{out}}(s_j) + g(-s_j T_{\text{eik}}) c_{\text{in}}(s_j)$$

and project the problem onto the real and imaginary parts of the phase-preconditioned subspace

$$(4.13) \quad \mathcal{K}_{\text{EIK}}^{2m}(\kappa, s) = \text{span}\{g(s T_{\text{eik}}) c_{\text{out}}(s_1), \dots, g(s T_{\text{eik}}) c_{\text{out}}(s_m), \\ g(-s T_{\text{eik}}) c_{\text{in}}(s_1), \dots, g(-s T_{\text{eik}}) c_{\text{in}}(s_m)\},$$

where $g(z) = \exp(-z)/z$ in three dimensions, while $g(z) = K_0(z)$ in two dimensions, with K_0 the modified Bessel function of the second kind and order zero. The singular behavior of the field at the source location is factored out, leading to a weaker dependence of the amplitude-functions $c_{\text{out/in}}$ on the source location.

In higher spatial dimensions, the field amplitudes are again obtained via one-way wave equations, but this time along the eikonal rays leading to decomposition directions $\pm \nabla T_{\text{eik}}$. For two-dimensional applications with $\mathcal{K}_0(z)$ as an incoming and $\mathcal{K}_0(-z)$ as an outgoing phase-function, we obtain the amplitude-functions

$$(4.14) \quad c_{\text{out}}(s_j) = \frac{s_j T}{\text{sign}(\mathfrak{Im}(s_j))i\pi} \left[\mathcal{K}_1(-s_j T) u(s_j) - \mathcal{K}_0(-s_j T) \frac{v^2}{s_j} \nabla T \cdot \nabla u(s_j) \right]$$

and

$$(4.15) \quad c_{\text{in}}(s_j) = \frac{s_j T}{\text{sign}(\mathfrak{Im}(s_j))i\pi} \left[\mathcal{K}_1(s_j T) u(s_j) + \mathcal{K}_0(s_j T) \frac{v^2}{s_j} \nabla T \cdot \nabla u(s_j) \right].$$

For three-dimensional applications with $\exp(sT)/sT$ as incoming and $\exp(-sT)/sT$ as outgoing phase-functions the corresponding amplitude-functions are given by

$$(4.16) \quad c_{\text{out}}(s_j) = \frac{v}{2} \exp(sT) \left[\frac{s}{v} \left(1 - \frac{1}{sT}\right) u(s_j) - \nabla T \cdot \nabla u(s_j) \right]$$

and

$$(4.17) \quad c_{\text{in}}(s_j) = \frac{v}{2} \exp(-sT) \left[\frac{s}{v} \left(1 + \frac{1}{sT}\right) u(s_j) + \nabla T \cdot \nabla u(s_j) \right].$$

Analogous to the one-dimensional reduced-order solution of (4.8), we can write the reduced-order solution in higher spatial dimensions as

$$(4.18) \quad u_m = \sum_{i=1}^m \begin{bmatrix} a_i(s) \\ \alpha_i(s) \end{bmatrix}^T \begin{bmatrix} g(sT_{\text{eik}})c_{\text{out}}(s_i) \\ g(\bar{s}T_{\text{eik}})c_{\text{out}}(\bar{s}_i) \end{bmatrix} + \sum_{i=1}^m \begin{bmatrix} d_i(s) \\ \delta_i(s) \end{bmatrix}^T \begin{bmatrix} g(-sT_{\text{eik}})c_{\text{in}}(s_i) \\ g(-\bar{s}T_{\text{eik}})c_{\text{in}}(\bar{s}_i) \end{bmatrix},$$

where the coefficients follow from the Galerkin condition.

Lemma 4.1 and Proposition 4.2 can be straightforwardly extended to the multidimensional case. However, Proposition 4.3 is not directly extendable to higher dimensions. As opposed to the one-dimensional case, a decomposition direction does not necessarily coincide with the travel direction of the wave and the field parametrization may be poor in such cases. This problem can be resolved, however, by considering MIMO wavefield systems with multiple sources and receivers, since in this case we have a decomposition direction for each source and the span of these directions may properly capture the propagation direction of the waves. In the next section, we therefore focus on wavefield systems with multiple sources and multiple receivers. The problem may be additionally complicated by multivalued solutions of the eikonal equation. In most situations it is sufficient to use the rays corresponding to the minimal travel time; however, as we shall see for the case of internal resonant structures in section 7.4, it can be beneficial to split the subspace along multiple rays.

5. Phase-preconditioning for MIMO systems.

5.1. Formulation of a block method. The time-domain equations governing MIMO systems are given by

$$(5.1) \quad \Delta u^{[l]} - \frac{1}{\nu^2} u_{tt}^{[l]} = -\frac{1}{\nu^2} \delta(t) \delta(x - x_S^l), \quad u|_{t=0} = 0, \quad u_t|_{t=0} = 0,$$

on $\mathbb{R}^k \times [0, \infty[$, where the superscript l is the source index with $l = 1, \dots, N_{\text{src}}$. The weak formulation of the corresponding s -domain equations is (cf. (2.7))

$$(5.2) \quad p^H Q(s) u^{[l]} = -\bar{p}(x_S^l) \frac{1}{\nu(x_S^l)^2} \quad \text{for } l = 1, 2, \dots, N_{\text{src}}.$$

Assuming possibly coinciding source-receiver pairs, we can define the source/receiver array as $B_s = [b^{[1]}, \dots, b^{[N_{\text{src}}]}]$, with individual source contributions $b^{[l]} = -\delta(x - x_S^l)/\nu^2(x_S^l)$ as columns. Equivalently, we define the array containing the fields as $U_s(s) = [u^{[1]}(s), \dots, u^{[N_{\text{src}}]}(s)]$ to write the MIMO equation of (5.2) as

$$(5.3) \quad P^H Q(s) U_s = P^H B_s.$$

Finally, we define the MIMO transfer-function of size $N_{\text{src}} \times N_{\text{src}}$ as

$$(5.4) \quad F(s) = B_s^H U_s(s),$$

which is symmetric due to reciprocity of the wavefields. To introduce the reduced-order transfer-function, we define a block-RKS

$$(5.5) \quad \mathcal{K}_B^{mN_{\text{src}}}(\kappa) = \text{span} \{U_s(s_1), U_s(s_2), \dots, U_s(s_m)\}$$

and its real counterpart containing $\mathcal{K}_B^{mN_{\text{src}}}(\kappa)$ and $\mathcal{K}_B^{mN_{\text{src}}}(\bar{\kappa})$ given by

$$(5.6) \quad \mathcal{K}_{B;R}^{2mN_{\text{src}}}(\kappa) = \text{span} \left\{ \Re \mathcal{K}_B^{mN_{\text{src}}}(\kappa), \Im \mathcal{K}_B^{mN_{\text{src}}}(\kappa) \right\}.$$

The ROM for the fields can now be constructed completely analogous to the SISO case. Specifically, with V_m a basis array that spans $\mathcal{K}_{B;R}^{2mN_{\text{src}}}(\kappa)$, we have

$$(5.7) \quad U_{s;m}(s) = V_m R_m(s)^{-1} V_m^H B_s \quad \text{with} \quad R_m(s) = V_m^H Q(s) V_m.$$

The reduced-order transfer-function now follows as

$$(5.8) \quad F_m(s) = B_s^H U_{s;m}(s),$$

and it is straightforward to show that the MIMO reduced-order transfer-function $F_m(s)$ is a Hermite interpolant of the MIMO transfer-function $F(s)$. The proof of this statement is completely analogous to the proof of Proposition 3.2.

To formulate the phase-corrected extensions of the block-RKS method, we note that the block-RKS field approximation $u_m^{[l]}(s)$ due to a source l can be written as

$$(5.9) \quad u_m^{[l]}(s) = \sum_{r=1}^{N_{\text{src}}} \sum_{i=1}^m \begin{bmatrix} a_i^{[l]} \\ \alpha_i^{[l]} \end{bmatrix}^T \begin{bmatrix} u^{[r]}(s_i) \\ u^{[r]}(\bar{s}_i) \end{bmatrix}$$

with $s_j \in \kappa$. In other words, the field approximation $u_m^{[l]}(s)$ due to source l is a linear combination of single-frequency solutions from *all* sources. A straightforward generalization of phase-preconditioning to MIMO systems is to use a field approximation

$u_m^{[l]}(s)$ that is a linear combination of phase-corrected incoming and outgoing fields from *all* sources. We write the field approximation as

$$(5.10) \quad u_m^{[l]}(s) = \sum_{r=1}^{N_{\text{src}}} \left(\sum_{j=1}^m \begin{bmatrix} a_{rj}^{[l]} \\ \alpha_{rj}^{[l]} \end{bmatrix}^T \begin{bmatrix} g(sT_{\text{eik}}^{[r]})c_{\text{out}}^{[r]}(s_j) \\ g(\bar{s}T_{\text{eik}}^{[r]})c_{\text{out}}^{[r]}(\bar{s}_j) \end{bmatrix} + \sum_{j=1}^m \begin{bmatrix} d_{rj}^{[l]} \\ \delta_{rj}^{[l]} \end{bmatrix}^T \begin{bmatrix} g(-sT_{\text{eik}}^{[r]})c_{\text{in}}^{[r]}(s_j) \\ g(-\bar{s}T_{\text{eik}}^{[r]})c_{\text{in}}^{[r]}(\bar{s}_j) \end{bmatrix} \right),$$

where $T_{\text{eik}}^{[r]}$ is the eikonal solution corresponding to the r th source. The coefficients $a_{rj}^{[l]}$, $\alpha_{rj}^{[l]}$, and $d_{rj}^{[l]}$, $\delta_{rj}^{[l]}$ are found via the block-Galerkin condition. For $N_{\text{src}} > 1$, this approach accounts for multidirectional scattering by representing the field as a linear combination of phase-corrected functions with multiple directions $\nabla T_{\text{eik}}^{[r]}$.

The idea that we followed to justify the use of block-Krylov methods is that the field caused by one source contains information about the field caused by a source with a different location (and frequency). In the context of phase-preconditioning, we can apply this idea a second time to obtain a block-preconditioned algorithm. This means that instead of using the phase-correction function $g(T_{\text{eik}}^{[r]})$ to only correct $c_{\text{out}}^{[r]}(s_j)$ for each source location r separately, we cross combine all phase-functions $g(sT_{\text{eik}}^{[r_2]})$ with all amplitudes $c_{\text{out}}^{[r_1]}(s_j)$ as indicated by the summations over r_1 and r_2 . This leads to a field approximation $u_m^{[l]}(s)$ due to the l th source given by

$$(5.11) \quad u_m^{[l]}(s) = \sum_{r_2=1}^{N_{\text{src}}} \sum_{r_1=1}^{N_{\text{src}}} \left(\sum_{j=1}^m \begin{bmatrix} a_{r_1 r_2 j}^{[l]} \\ \alpha_{r_1 r_2 j}^{[l]} \end{bmatrix}^T \begin{bmatrix} g(sT_{\text{eik}}^{[r_2]})c_{\text{out}}^{[r_1]}(s_j) \\ g(\bar{s}T_{\text{eik}}^{[r_2]})c_{\text{out}}^{[r_1]}(\bar{s}_j) \end{bmatrix} + \sum_{j=1}^m \begin{bmatrix} d_{r_1 r_2 j}^{[l]} \\ \delta_{r_1 r_2 j}^{[l]} \end{bmatrix}^T \begin{bmatrix} g(-sT_{\text{eik}}^{[r_2]})c_{\text{in}}^{[r_1]}(s_j) \\ g(-\bar{s}T_{\text{eik}}^{[r_2]})c_{\text{in}}^{[r_1]}(\bar{s}_j) \end{bmatrix} \right).$$

The expansion coefficients are found from the block-Galerkin condition. Basis vectors in this expression can be viewed as a tensor-product of the amplitudes $c_{\text{out}}^{[r_1]}$ and the phase terms $g(-\bar{s}T_{\text{eik}}^{[r_2]})$, while the Hadamard product is used spatially.

Similarly to the multidimensional SISO case, Lemma 4.1 and Proposition 4.2 can be straightforwardly extended for the multidimensional MIMO case. Our experiments presented in section 7 also indicate that the number of interpolation points needed for multidimensional MIMO configurations is dependent on the complexity of the wave speed model and not the largest travel time (as proven for the one-dimensional case in Proposition 4.3).

Remark 2. Application to resonant cavities. In the case that resonant cavities are present in the wave speed model, it is essential to model the resonance modes of the cavity as well as wave propagation from sources and receivers toward the cavity. This can be done by taking multivalued eikonal travel times into account. In the result section we give an example of a layered configuration that contains a resonant cavity. To extend our approach to this configuration we factored out two phase-functions, one that approximates the cavity modes and one that approximates the wave propagation from the cavity to the source and receivers.

5.2. SVD truncation of the expansion amplitudes. In the case of many sources and receivers it is possible to compress the field amplitudes $c_{\text{out}}^{[r_1]}(s_j)$ and $c_{\text{in}}^{[r_1]}(s_j)$ a posteriori using a thin singular value decomposition (SVD). Using this SVD we compress the ROM and remove redundancy in the expansion of (5.11). In

the phase-preconditioned approach redundancy of the basis occurs in two ways. First, smooth amplitude-functions corresponding to different frequencies can be close to linear dependent due to the frequency-dependent basis vectors in (5.11). Second, the amplitude-functions of multiple sources can be similar, which leads to redundancy once we use cross combinations of amplitudes and phase-functions as in (5.11). The original block-Krylov basis does not have these redundancies.

To realize SVD truncation, the amplitudes $c_{\text{in}}^{[r_1]}(s_j)$ and $c_{\text{out}}^{[r_1]}(s_j)$ are first normalized in pairs to have unit Euclidean norm. To be more specific, $c_{\text{in}}^{[r_1]}(s_j)$ and $c_{\text{out}}^{[r_1]}(s_j)$ are normalized by

$$\sqrt{\|c_{\text{in}}^{[r_1]}(s_j)\|^2 + \|c_{\text{out}}^{[r_1]}(s_j)\|^2}$$

such that the sum of the squared singular values is $2mN_{\text{src}}$. In this way the ratio between the incoming and outgoing amplitudes is preserved. The SVD of the $2mN_{\text{src}}$ incoming and outgoing amplitudes is then computed separately and truncated after $M_{\text{SVD}}^{\text{out}}$ and $M_{\text{SVD}}^{\text{in}}$ left singular vectors to obtain the compressed amplitudes $c_{\text{in};\text{SVD}}^j$ and $c_{\text{out};\text{SVD}}^j$. The original amplitudes $c_{\text{in}}^{[r_1]}(s_j)$ and $c_{\text{out}}^{[r_1]}(s_j)$ are associated with a specific frequency s_j and source r_1 , whereas the compressed amplitudes $c_{\text{in};\text{SVD}}^j$ and $c_{\text{out};\text{SVD}}^j$ are associated with a singular value. The amplitudes are therefore no longer associated with a source or frequency and the corresponding subscripts are dropped and replaced by the singular value index j . The resulting reduced-order solution expressed in terms of these compressed amplitudes is given by

$$(5.12) \quad u_m^{[l]}(s) = \sum_{r=1}^{N_{\text{src}}} \left(\sum_{j=1}^{M_{\text{SVD}}^{\text{out}}} \begin{bmatrix} a_{rj}^{[l]} \\ \alpha_{rj}^{[l]} \end{bmatrix}^T \begin{bmatrix} g(sT_{\text{eik}}^{[r]})c_{\text{out};\text{SVD}}^j \\ g(\bar{s}T_{\text{eik}}^{[r]})\bar{c}_{\text{out};\text{SVD}}^j \end{bmatrix} + \sum_{j=1}^{M_{\text{SVD}}^{\text{in}}} \begin{bmatrix} d_{rj}^{[l]} \\ \delta_{rj}^{[l]} \end{bmatrix}^T \begin{bmatrix} g(-sT_{\text{eik}}^{[r]})c_{\text{in};\text{SVD}}^j \\ g(-\bar{s}T_{\text{eik}}^{[r]})\bar{c}_{\text{in};\text{SVD}}^j \end{bmatrix} \right),$$

where $M_{\text{SVD}}^{\text{out/in}} \ll mN_{\text{src}}$.

If we contract the outgoing amplitudes and the conjugate of the incoming amplitudes into one amplitude basis c_{SVD}^j , compute the SVD after pairwise normalization, and evaluate on the imaginary line ($s \in i\mathbb{R}$) we can expand the field as

$$(5.13) \quad u_m^{[l]}(s) = \sum_{r=1}^{N_{\text{src}}} \sum_{j=1}^{M_{\text{SVD}}} \begin{bmatrix} a_{rj}^{[l]} \\ \alpha_{rj}^{[l]} \end{bmatrix}^T \begin{bmatrix} g(sT_{\text{eik}}^{[r]})c_{\text{SVD}}^j \\ g(-sT_{\text{eik}}^{[r]})\bar{c}_{\text{SVD}}^j \end{bmatrix} \quad \text{with } s \in i\mathbb{R},$$

where $M_{\text{SVD}} \ll mN_{\text{src}}$.

Here the c_{SVD}^j are first M_{SVD} left singular vectors of

$$\begin{bmatrix} c_{\text{in}}^{[r_1]} & c_{\text{out}}^{[r_1]} \end{bmatrix} / \sqrt{\|c_{\text{in}}^{[r_1]}\|^2 + \|c_{\text{out}}^{[r_1]}\|^2}.$$

In our numerical experiments we show that the singular values of the contracted amplitudes c_{SVD}^j decay much faster than the ones of the block-Krylov basis. Later in Figure 3(a) we plot the decay of singular values of a matrix with pairwise normalized vectors c_{SVD}^j and normalized vectors $u^{[r_1]}(s_j)$ as columns. Each SVD trace is normalized to the largest singular value to emphasize the decay. The singular values associated with c_{SVD}^j show a strong decay with a plateau at the level of the finite-difference error, whereas the singular values associated with the wave field solutions $u^{[r_1]}(s_j)$ barely show any decay before reaching the Nyquist sampling rate.

The compressibility of the amplitudes confirms that the chosen parametrization of the wavefield is valid. Phase-preconditioning enhances the convergence of RKS

not by increasing the subspace but by preconditioning a small basis of problem specific amplitudes in dependence of the evaluation frequency using phase-functions. In our experiments we show that the number of contributing amplitudes is only weakly dependent on the number of sources.

The compression of the amplitudes offers several advantages. First, it significantly reduces the cost of evaluating the ROM, since it reduces the amount of inner products that need to be computed to obtain the reduced-order operator. Second, the cost associated with communicating and storing the ROM is reduced as well. The compressed amplitude basis is only very weakly dependent on the source locations used to construct it. Therefore, we can reduce the number of sources (right-hand sides) for which the basis vectors need to be computed, since their response can be approximated from their eikonal travel time and the basis vectors computed from other sources. It is in line with our effort to reduce the computation at every stage of our algorithm.

Remark 3. Relation to existing approaches. Phase-tracking was recently introduced to decompose a seismic response into smooth amplitudes and phase-functions via optimization [21]. This method also uses the idea that once a physically meaningful parameterization has been found, it can be used to extrapolate to higher or lower frequencies. This method is entirely data driven, whereas our method is model driven and we can therefore use the eikonal time as phase. However, using optimization to obtain even smoother amplitude-functions may even be possible in our method.

In the fast Helmholtz solver proposed in [17], a plane wave ansatz is substituted into the Helmholtz equation and the resulting advection-diffusion-reaction equation is solved using multigrid methods. While in this method the Helmholtz equation is basically preconditioned with a single asymptotic function, in our proposed approach we effectively use a different phase-function for every source and receiver pair. In addition, we don't solve the resulting equations directly but project them onto a subspace.

6. Discrete formulation. In this section we consider the discrete implementation of the introduced reduced-order modeling technique. Discretization and selection of the grid accuracy are addressed first, followed by a discussion on how we handle numerical dispersion.

6.1. Finite difference discretization. Our basic approach is to solve (2.5) (restated here as (6.1)) using phase-corrected single-frequency solutions as expansion functions and to obtain a reduced-order solution for a complete spectral interval of interest via the Galerkin condition.

We consider a rectangular domain $\Omega \in \mathbb{R}^k$ with constant $\nu(x)$ on $\mathbb{R}^k \setminus \Omega$ and discretize the equation

$$(6.1) \quad \int_{\Omega} \nabla \bar{p} \cdot \nabla u d\Omega - \int_{\Omega} \bar{p} \frac{s^2}{\nu^2} u d\Omega + \int_{\partial\Omega} \bar{p} D(s) u d\partial\Omega = - \frac{1}{\nu(x_S)^2} \bar{p}(x_S),$$

using finite differences to obtain a linear shifted system with a matrix nonlinearly depending on s . Discretization of the first two terms in the above equation using a second-order accurate finite-difference scheme with constant step sizes is straightforward. To discretize the third term, we approximate the DtN map $D(s)$ using nearly optimal discrete perfectly matched layers (PML) according to [9]. The optimal Zolotarev rational approximants used for the PML construction make the size of the finite-difference problem (necessary for accurate approximation of $D(s)$) in $\mathbb{R}^k \setminus \Omega$

negligible compared to the grid in Ω 's interior. In two dimensions, for instance, the resulting equations that need to be solved in the PML can be solved efficiently with a block-cyclic solver [25] or with a band solver after sorting the PML system to a bandwidth of $2k + 1$ (in two dimensions). This discretization leads to the matrix equation

$$(6.2) \quad \mathbf{p}^H \mathbf{Q}(s) \mathbf{u} = -\mathbf{p}^H \mathbf{b},$$

with \mathbf{b} the discrete approximation of the scaled delta function. The matrix $\mathbf{Q}(s)$ of order N inherits all properties of the continuous operator $Q(s)$ and thus follows the Schwarz reflection principle, is symmetric in a bilinear form ($\mathbf{Q}(\bar{s})$ is the adjoint of $\mathbf{Q}(s)$ in the Hermitian inner product), and has a nonlinear numerical range in the left half of the complex s -plane. The single-frequency solutions $\mathbf{u}(s_j)$ needed to build the RKS can be obtained using iterative solvers or Gaussian elimination. The eikonal equation $|\nabla T|^2 = \frac{1}{\nu^2}$ is solved on the same grid using a fast marching method [26]. The associated computational cost is negligible with respect to the cost of solving the Helmholtz equation.

We use uniform finite differences in order to keep our approach simple, although single-sided finite differences may lead to better parameterizations near caustics and jumps in the medium parameters.

6.2. Realization on two grids. For smooth media, the amplitudes $c_{\text{out}}^{[r_1]}$ and $c_{\text{in}}^{[r_1]}$ are smooth functions of the spatial coordinates, since the highly oscillatory part of the frequency-domain wavefield is factored out together with the source singularity. Therefore, linear combinations of these amplitudes can form reasonable approximations of amplitude distributions at higher frequencies. Consequently, the phase-corrected ROMs can extrapolate to frequencies outside the convex hull of the interpolation points. Building a ROM that extrapolates to higher frequencies has the advantage that the amplitudes can be determined on grids that are much coarser than grids required by a direct method at these high frequencies. This significantly reduces the computational cost of solving the Helmholtz equation to obtain the amplitude-functions $c_{\text{out/in}}$, since the main cost of the algorithm is associated with solving shifted systems.

To be specific, let $\mathbf{Q}_{\text{fine}}(s)$ and $\mathbf{Q}_{\text{coarse}}(s)$ denote the matrix operators obtained by discretizing the wave operator $Q(s)$ on a fine and a coarse grid, respectively. Writing the transfer-function and field approximations obtained with these fine and coarse grid operators as $F_m(s)$ and $\mathbf{U}_m(s)$ and $F_{c;m}(s)$ and $\mathbf{U}_{c;m}(s)$, respectively, we have

$$(6.3) \quad F_m(s) - F_{c;m}(s) = [\mathbf{U}_m(s) - \mathbf{P}_{\text{interp}} \mathbf{U}_{c;m}(s)]^H \mathbf{Q}_{\text{fine}}(s) [\mathbf{U}_m(s) - \mathbf{P}_{\text{interp}} \mathbf{U}_{c;m}(s)],$$

which is essentially a block version of (3.10). \mathbf{U}_m and $\mathbf{U}_{c;m}$ don't live on the same grid and the interpolation matrix $\mathbf{P}_{\text{interp}} \in \mathbb{R}^{N_{\text{fine}} \times N_{\text{coarse}}}$ makes the above equation consistent. In this case, however, $\mathbf{U}_m(s) - \mathbf{P}_{\text{interp}} \mathbf{U}_{c;m}(s)$ signifies the difference between the fine and coarse grid field solutions and the interpolation property as presented in Proposition 3.2 obviously does not hold here. In other words, using a coarse grid for construction and a fine grid for projection leads to a loss of the interpolation property of the ROM. On the other hand, we do increase the accuracy of the coarse transfer-function as the errors introduced by interpolation and the coarse grid solution get squared at the interpolation points $s \in \kappa \cup \bar{\kappa}$.

The main drawback of using coarser grids is that the numerical dispersion error increases and the analytic phase term $\exp(\pm s T_{\text{eik}})$ does not match the phase term of

TABLE 1
Summary of the phase-preconditioned RKS algorithm.

1.	Solve the eikonal equation to obtain T_{eik}
2.	Solve the coarse grid wave problems $Q_{\text{coarse}}(\kappa_i)u(\kappa_i) = b$
3.	Decompose the fields $u(\kappa_i)$ into the amplitudes c_{in} and c_{out}
4.	Compress the amplitudes using an SVD
5.	Interpolate the amplitudes to a fine grid
6.	For each evaluation frequency project the operator Q_{fine} onto the phase-corrected subspace
7.	Compute the resulting reduced-order transfer-function
8.	Compute the inverse Laplace transform to obtain the time-domain transfer-function

$u(s)$ for large imaginary shifts. Fortunately, we can correct for this phase mismatch. To be precise, in the analytic case the phase term $\exp(-sT_{\text{eik}})$ is used to cancel the high-frequency dominant term $-s^2/\nu^2$ in the wave equation. To guarantee that this cancellation takes place in the discrete case and to match the discrete and analytic phases, we introduce the discrete, second-order, finite-difference gradient matrix D_{x_i} (see, e.g., [7]) mapping from a primary to a dual grid in all spatial directions $i = 1, \dots, k$ and adjust the wave speed model from ν to ν' , where ν' follows from the requirement

$$(6.4) \quad \exp\left(2sT_{\text{eik;p}}^{[l]}\right) \sum_{i=1}^k \left[D_{x_i}^T \exp\left(-sT_{\text{eik;d}_{x_i}}^{[l]}\right) \right]^2 = \frac{s^2}{\nu'^{[l]2}},$$

where $T_{\text{eik;p}}$ is the eikonal time on the primary grid, and $T_{\text{eik;d}_{x_i}}$ is the eikonal time on the primary grid in all spatial directions except for x_i . This equation is thus an equation on the primary grid. This is the discrete counterpart of the (continuous) relation $\exp(2sT_{\text{eik}}) [\nabla \exp(-sT_{\text{eik}})]^2 = \frac{s^2}{\nu^2}$. This equation ensures that the high-frequency dominant term $-s^2/\nu^2$ vanishes and the numerical dispersion error is minimized.

Finally, we note that this dispersion correction is accurate only in the dominant direction of $\nabla T_{\text{eik}}^{[l]}$ and works only in the reduced-order modeling framework. In our block approach multiple directions are taken into account by incorporating multiple source locations. Therefore, errors occurring in the directions orthogonal to $\nabla T_{\text{eik}}^{[l]}$ are corrected in the block projection framework by projection onto sources with different dominant directions.

We summarized our algorithm in Table 1.

Remark 4. Grid selection. The main goal of our method is to lower the asymptotic complexity of the basis computation and shift work to the highly parallelizable stage of computing inner products. In our current implementation we discretize all Helmholtz systems on the same coarse grid; another step toward reducing computational efforts is to adapt the coarse grid discretization for each solve to the frequency of the RKS shift.

7. Two-dimensional experiments. In this section we illustrate the performance of the developed solution methods using three different two-dimensional numerical experiments. In our first set of experiments, we show the performance of the proposed preconditioning technique for wavefields in a smooth layered configuration. We simulate the same structure with and without grid coarsening to show the effects of both concepts. As a second example, we consider a nonsmooth medium with jumps in the wave speed profile to illustrate that the effectiveness of preconditioning decreases as the high-frequency geometrical optics argument is no longer valid. However, the

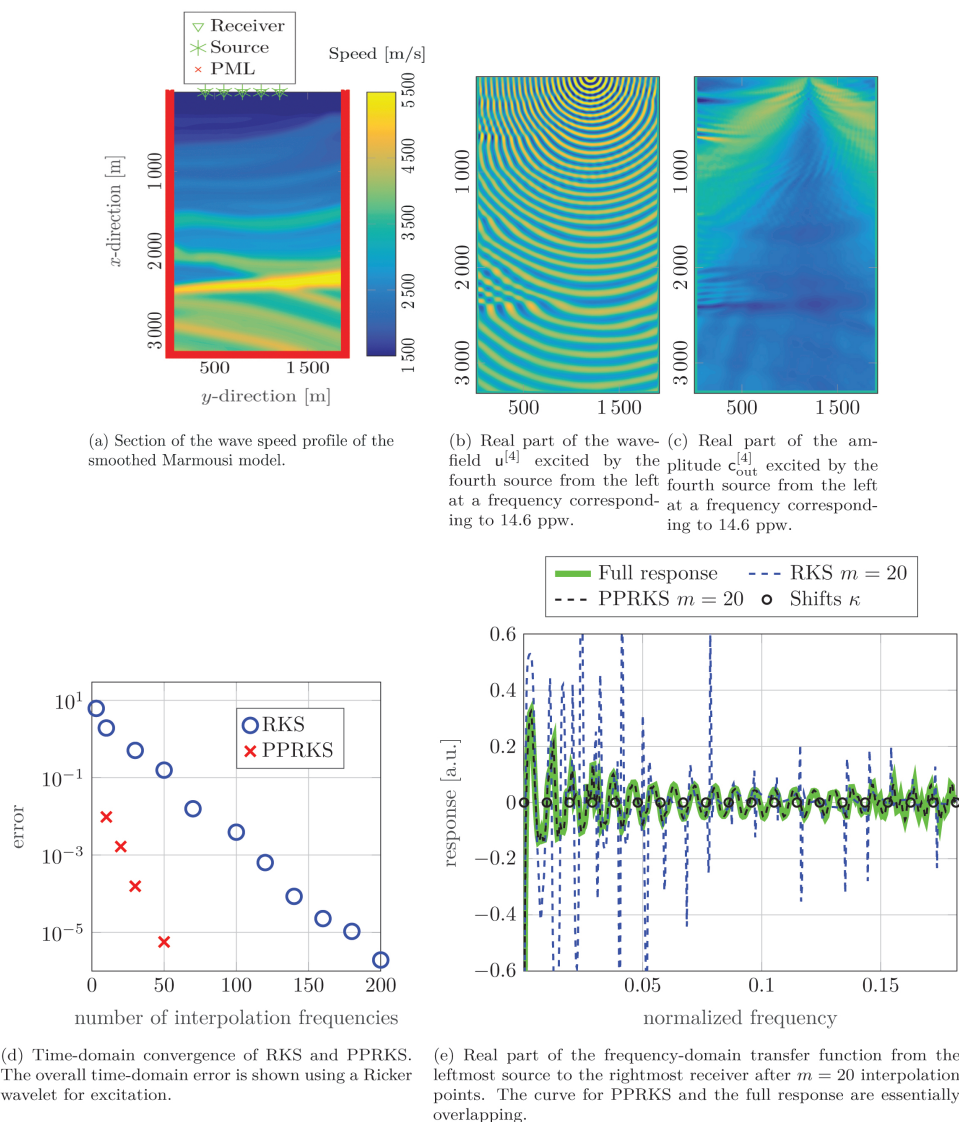


FIG. 1. The smoothed Marmousi layers test configuration. Shown are the configuration (a), the contour lines of the eikonal times (b), and a solution vector prior to (c) and after parametrization (d).

method still exhibits excellent approximation properties even for nonsmooth media. Finally, in the third experiment, a configuration with a resonant cavity present in a smooth geology is considered.

All experiments are performed using uniformly distributed shifts on the imaginary axis and we show that significant reductions can be achieved using this simple choice of shifts. The theory of optimal shift selection for RKSs has been developed for many applications; however, it is not clear how to extend these approaches to a priori optimal shift selection for wave equations and phase-preconditioning, due to the complex spectrum.

7.1. A geophysical structure with a smooth wave speed profile. To illustrate the effect of phase-preconditioning, we consider the smoothed geophysical structure illustrated in Figure 1(a). This model is obtained by smoothing a layered section of the acoustic Marmousi model [6] with a Hanning window of width $h_{\text{Han}} = 200$ m leading to a discretized model of order $N = 4 \cdot 10^5$. Five coinciding source-receiver pairs are placed at the top boundary, where a perfectly reflecting boundary condition is imposed to model a water-air interface. A Ricker wavelet with a maximum in its spectrum at $\omega^{\text{peak}} = 8$ Hz (13 ppw at 1% cut-off frequency) is used as a source signature and a near optimal eight-layer PML [9] is applied on the remaining outer boundaries to simulate outward wave propagation toward infinity. Finally, a fast marching method [26] is adopted to obtain the eikonal solution for this configuration. The true solution $u^{[4]}$ corresponding to the fourth source from the left at the frequency corresponding to 14.6 ppw is visualized in Figure 1(b). This solution shows “diving wave behavior” and a caustic can be seen at a depth of about 750 m in the left half of the configuration. The real part of the outgoing amplitude $c_{\text{out}}^{[4]}$ is depicted in Figure 1(c) and is clearly spatially much smoother than the original wavefield. Reflections of the wavefield can easily be identified in this amplitude plot.

The overall time-domain errors of the block-RKS and preconditioned block-RKS ROMs without grid coarsening are shown in Figure 1(d), where we used a 500-point Fourier method to obtain a comparison solution. The overall time-domain error is defined by the ratio of the root mean square (r.m.s.) error of all traces and the r.m.s. of the signals over *all* traces. Preconditioning the RKS method significantly decreases the number of interpolation points needed to reach a certain error level. To obtain an error of 1 percent in the time-domain, for example, the RKS algorithm needs about 80 interpolation points, while only 10 interpolation points are required in the phase-preconditioned algorithm. This fast convergence is due to the construction of the WKB-like field approximations at high frequencies in PPRKS, which already provide an accurate approximation of the Green’s function at high frequencies and in smooth structures as considered in this example.

The real part of the transfer-function of the leftmost source to the rightmost receiver is shown in Figure 1(e) for structure-preserving and preconditioned block-RKS ROMs of order $m = 20$. The phase-preconditioned model coincides with the comparison solution on the complete frequency interval of interest. The main oscillations present in this Green’s function response are due to the direct arrival of the wave and its first reflection from the salt layer located at a depth of about 2500 m. Typically, the PPRKS method provides a smooth approximation to the field response showing only small errors in the amplitudes or at highly oscillatory reflections. The structure-preserving RKS method, on the other hand, overshoots after every interpolation point causing spiking behavior as can be clearly seen in Figure 1(e).

7.2. Grid coarsening and SVD. The amplitude-functions c_{in} and c_{out} are spatially much smoother than the wavefield and therefore we expect that a coarser spatial grid can be employed. To investigate the effects of grid coarsening, we consider the same wave speed profile as in the previous example and place 12 coinciding source-receiver pairs at the top water-air interface instead of 5. For excitation, we use a modulated Gaussian pulse with a center frequency ω^{peak} and its support essentially given by $[0, 2\omega^{\text{peak}}]$. The pulse is shifted in time such that it starts at $t = 0$. Spatial discretization is now chosen such that we have about 5.5 points per smallest wavelength, where the wavelength corresponds to the center frequency of the pulse

and 2.7 ppw at the cut-off frequency of the pulse. With this choice, the step sizes of the grid are 4 times larger than in the previous example leading to a system that is 16 times smaller with $N = 2.5 \cdot 10^4$ unknowns. Using such a coarse grid to model wavefields without phase-preconditioning is obviously insufficient, but here we expect that the smoothness of the amplitude-functions c_{in} and c_{out} allows us to use a much coarser grid. During the evaluation of the ROM we project an operator corresponding to a fine grid onto the phase-corrected RKS. For this example, we choose a fine operator using half the step size compared to the previous operator in order to show that the projection gauges the ROM to the operator used during projection.

For MIMO systems with grid coarsening, we define the error as the error averaged over all source-receiver combinations. We denote the elements of the finite difference matrix transfer-function $F_F(s)$ by $f_F^{[ij]}(s) = b^{[i];H} u^{[j]}(s)$, while the elements of the ROM transfer-function $F_m(s)$ are given by $f_m^{[ij]}(s) = b^{[i];H} u_m^{[j]}(s)$. Having introduced these elements, the average MIMO error as a function of frequency is defined as

$$(7.1) \quad \text{err}_{\text{ROM}}^{\text{average}}(m, s) = \frac{\sqrt{\omega_{\max}}}{N_{\text{src}}^2} \sum_{j=1}^{N_{\text{src}}} \sum_{i=1}^{N_{\text{src}}} \frac{|f_F^{[ij]}(s) - f_m^{[ij]}(s)|}{\left(\int_{\omega=0}^{\omega_{\max}} |f_F^{[ij]}(i\omega)|^2 d\omega \right)^{1/2}}.$$

It is assumed that the comparison solution $f_F^{[ij]}(t)$ is computed with a spatial discretization of sufficient accuracy. This averaged frequency domain error definition gives a higher error, yet delivers more insight, than computing the overall error. The overall error is dominated by the mono static elements $f_F^{[ii]}$, whose direct arrival contains most energy and is well approximated. Furthermore, the above error definition allows us to study the error as a function of frequency.

The phase-corrected RKS is build using $m = 40$ equidistant shifts on the imaginary axis from $\omega = 2.4 \cdot 10^{-3} \omega^{\text{peak}}$ (2383 points per smallest wavelength) to $\omega = 1.1 \omega^{\text{peak}}$ (5 ppw). In other words, the RKS interpolation frequencies uniformly cover the lower half of the support of the spectrum of the source wavelet. With $m = 40$ interpolation points and $N_{\text{src}} = 12$ source-receiver pairs, we have 480 amplitude-functions c_{in} and an additional 480 amplitude-functions c_{out} . Computing the SVD of the 960 amplitude-functions $[c_{\text{out}} \ c_{\text{in}}]$, we observe that for this example, essentially only the first 100 singular functions contribute to the ROM for the contracted amplitudes. We therefore use a truncated SVD that uses the first 100 SVD basis functions to represent the amplitudes.

The resulting time-domain trace from the leftmost source to the rightmost receiver is shown in Figure 2, compared to the trace obtained via a 500-point Fourier method using an operator with step sizes eight times smaller than the step sizes used in the coarse operator. Both responses clearly coincide on the considered time window and the first arrival of the pulse, the complex interaction between the pulse and the upper layered medium, and the reflection of the pulse at the high contrast salt layer around $t = 3000$ can be observed. The multiple reflection from source to salt layer, water-air interface, salt layer, and back to the receiver can be seen around $t = 6000$.

In Figure 3(a) we plot the decay of the singular values of a matrix with pairwise normalized vectors. The decay of the singular values of the single-frequency solutions that make up the RKS is shown in black. The singular values associated with the amplitude matrix are shown in blue. Prior to the SVD the vectors $u^{[r_1]}$ were normalized, such that the square of the singular values sum up to mN_s . The vectors $c_{\text{out}}^{[r_1]}(s_j)$ and $\bar{c}_{\text{in}}^{[r_1]}(s_j)$ were normalized in pairs together to reflect the ratio of the incoming and

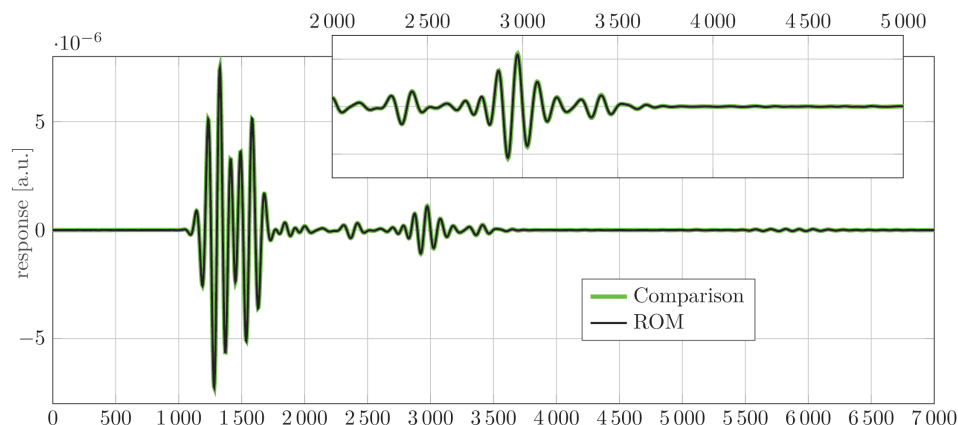


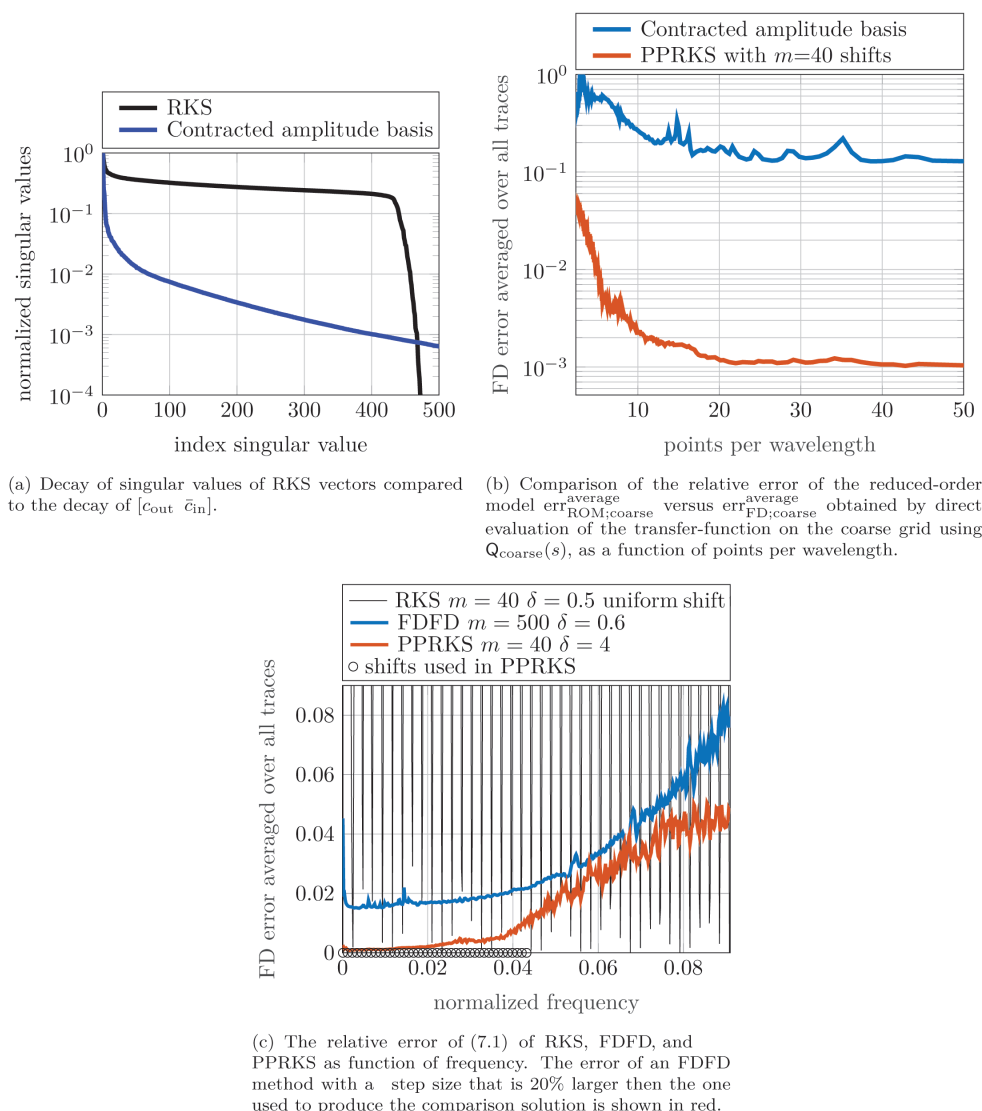
FIG. 2. Time-domain trace from the leftmost source to the rightmost receiver after $m = 40$ interpolation points.

outgoing waves at each frequency. Their SVD is computed together such that the sum of the squares of all singular values adds up to mN_s as well. Finally, to show the decay in singular values we normalize the largest singular value to one for each of the shown curves. This figure clearly shows that incoming and outgoing amplitudes are significantly compressible, whereas the RKS vectors are not. The singular values associated to RKS drop by less than a factor of 2 between the index 50 and the index 400, indicating that the RKS can hardly be compressed. To show that after the compression the basis is essentially independent of the number of sources, we computed the SVD of the amplitudes $[c_{\text{out}}^{[r_1]}(s_j) \quad \bar{c}_{\text{in}}^{[r_1]}(s_j)]$ for $N_{\text{src}} = 12, 24, 48$, and 96. The number of (normalized) singular values larger than 0.01 is shown in Table 2. It shows that the number of contributing singular vectors is basically independent of the number sources, and so are the left-hand side singular vectors.

In Figure 3(c), the averaged error $\text{err}_{\text{ROM}}^{\text{average}}$ over all 12^2 traces is shown along with the interpolation points used in the construction of the ROM. The same Fourier method that was used to compute the comparison solution in Figure 2 is used here to compute the errors in the transfer-function. Furthermore, the figure shows the error of a finite-difference frequency-domain (FDFD) method with a normalized step size of 0.6, which is 20% larger than the step size used to compute the comparison solution and used for the operator that was projected onto the PPRKS. For higher frequencies such an operator has increasing dispersion, such that the solutions between the comparison FDFD method with normalized step size of 0.5 and the one of 0.6 don't match anymore. For low frequencies there is a small discrepancy due to the inability of both grids to approximate a delta source.

When introducing grid coarsening, the ROM no longer interpolates the transfer-function, but the error remains small and below 1% on the frequency interval covered by the interpolation points. In addition, we observe that the phase-preconditioned ROMs can extrapolate to higher frequencies to a certain extent, since the basis in PPRKS is frequency-dependent. The error only gradually grows outside the interpolation interval, which covers the lower half of the spectrum of the pulse, and at 2.4 points per smallest wavelength we end up with an error of about 5%.

Finally, in Figure 3(b) the averaged error in the transfer-function is shown as a function of the number of points per wavelength. Again, the PPRKS with 40 interpolation points is compared to the 500-point Fourier method, but this time the

FIG. 3. *Smooth Marmousi test configuration with grid coarsening.*

latter method uses the same coarse-grid operator that is used during construction of the PPRKS (instead of the operator that uses step sizes 0.6 as in the previous figure). Clearly, the RKS approach that uses a Galerkin condition to select optimal linear combinations with respect to a fine operator outperforms the direct Fourier method that uses the same operator to construct the field approximations. The RKS approach is gauged to the operator by using the Galerkin condition.

7.3. A geophysical structure with a nonsmooth wave speed profile. The justification of the phase-preconditioned algorithm is based on a geometrical optics argument. This asymptotic argument is applicable for smooth media with spatial variations that take place on a scale much smaller than the wavelength. On the other

TABLE 2

Number of (normalized) singular values larger than 0.01 in dependence of the number of sources.

N_{src}	12	24	48	96
Number of (normalized) singular values > 0.01 for $[c_{\text{out}}, \bar{c}_{\text{in}}]$	69	72	73	73
Number of (normalized) singular values > 0.01 for u	457	833	1369	1741
$m \cdot N_{\text{src}}$	480	960	1920	3840

hand, RKS reduced-order modeling is a valid approach independent of the medium considered and Proposition 4.3 shows that one-dimensional problems with piecewise constant wave speeds need not be a problem for this approach to work. Therefore, let us turn to an unsmoothed variant of the layered geophysical structure from the Marmousi model considered earlier as depicted in Figure 4(a).

For this structure, we essentially follow the same procedure as before. Specifically, we again position 12 source-receiver pairs at the top air-water interface and use the same coarse grid operator as in the previous example to construct a phase-corrected RKS ROM of order $m = 40$ with interpolation points on the imaginary axis covering the lower half of the spectrum of the pulse such that we have 5 points per smallest wavelength for the highest interpolation frequency. The center frequency of the pulse is again chosen at 5.5 ppw. The only difference in model construction compared with the previous example is that here we use a truncated SVD that takes 150 SVD basis functions into account, instead of the 100 basis functions in the previous example. Here, more basis functions are required, since the amplitude-functions are less smooth due to the nonsmooth wave speed profile of the present Marmousi model. Finally, the comparison solution is computed using a direct 500-point Fourier method using a spatially discretized operator with step sizes that are four times smaller than the step sizes used in the coarse operator. The coarse operator has a normalized step size of 4 and the operator used to compute the comparison FDFD response has a normalized step size of 1.

The resulting error is shown in Figure 4(b) along with the corresponding error curve for an FDFD method which used the coarse operator that constructed the PPRKS. In addition, the error of an ordinary RKS method is shown, which uses the fine operator for construction and evaluation. It uses $m = 40$ shifts uniformly distributed on the whole spectral interval. An RKS method on the fine grid interpolates the FDFD response on the shifts, which leads to a strongly oscillatory error curve. This error curve clearly shows the advantage of phase-preconditioning with a dual grid approach—a lower error is achieved while solving considerably smaller shifted systems, and projecting on the same operator. The performance of the algorithm for smooth profiles is better than for nonsmooth wave speed profiles, especially for extrapolated frequencies. We also observe that the error decreases for lower frequencies, since lower frequencies have larger wavelengths and variations in the wave speed profile take place on a scale smaller than these wavelengths of operation. Furthermore, compared to an FDFD method that uses a 20% coarser grid than the comparison solution, the PPRKS achieves lower errors across the whole spectral interval while the systems that need to be solved are much smaller. Especially in the area where the phase-preconditioned method has shifts it reproduces the comparison solution remarkably well. A similar error comparison is shown in Figure 4(c), where the error is plotted against points per wavelength.

To illustrate the effects of an increased error in the time-domain, we show the time trace for the most distant source-receiver pair in Figure 4(d) (for the same Gaussian pulse as used before in Figure 2) along with a comparison solution obtained with

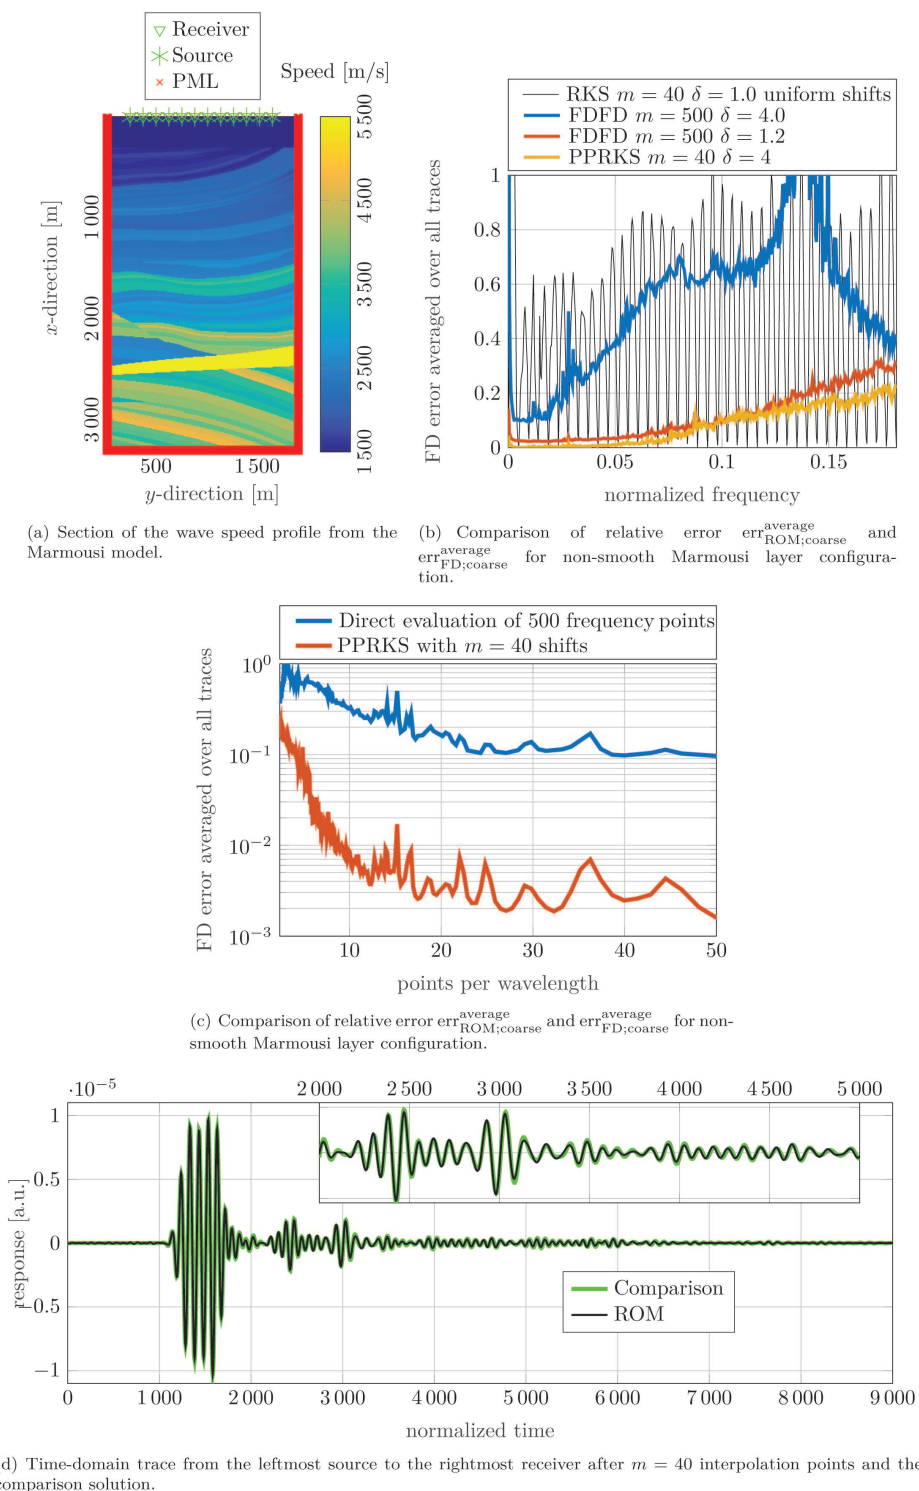


FIG. 4. Nonsmooth Marmousi test configuration with grid coarsening.

the 500-point Fourier method. We observe that the arrival times are approximated well; only the amplitudes are slightly off. Throughout our numerical work, we have found that this result is typical for nonsmooth problems. Furthermore, compared with the same trace computed for smooth media as shown in Figure 2, it is clear that a larger part of the pulse is scattered back to the receiver, as arrivals are visible on the complete time interval of observation for the nonsmooth velocity profile considered here.

7.4. A resonant cavity embedded in a smooth geology. In this section we investigate the performance of our algorithm in a configuration with a resonant cavity. Figure 5(a) shows the wave speed profile, which is inspired by borehole exploration. Coinciding source-receiver pairs are placed at the surface and inside a borehole of slow acoustical wave speed.

The grid coarsening procedure and wavelet selection are equivalent to the previous examples. A coarse grid operator with the same accuracy as selected in the previous example is used to construct a phase-corrected RKS ROM of order $m = 40$ with interpolation points on the imaginary axis covering the lower half of the spectrum of the pulse such that 5 points per smallest wavelength are used for the highest interpolation frequency. The center frequency of the pulse is again chosen at 5.5 ppw.

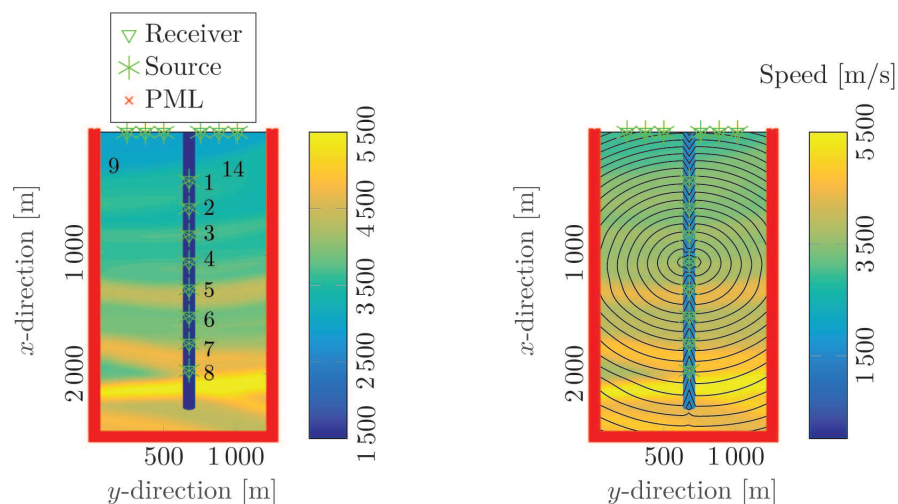
To approximate the cavity-resonances with few interpolation points, we extend the approach discussed in this paper and factor out oscillations of resonance modes along the borehole. To do so, we take the fact that the eikonal time T_{eik} is multivalued into account. More specifically, each solution $u(\kappa_i)$ is split using two different phase terms, a cavity-mode phase term and a propagation phase term. The eikonal phase term shown in Figure 5(b) shows a caustic inside the borehole, which has a low wave speed compared to its surrounding. In this experiment we also factor out the cavity-mode phase term $g(sT_{\text{eik;CM}})$, where $T_{\text{eik;CM}}$ follows the borehole as shown in Figure 5(c). The eikonal time of the cavity-mode $T_{\text{eik;CM}}$ is not the second arrival, but it is chosen to correctly factor out resonances present in the borehole. At every interpolation point we split the field into four amplitudes as

$$(7.2) \quad u^{[l]}(s_j) = g(s_j T_{\text{eik}}^{[l]}) c_{\text{out;eik}}^{[l]}(s_j) + g(-s_j T_{\text{eik}}^{[l]}) c_{\text{in;eik}}^{[l]}(s_j),$$

$$(7.3) \quad u^{[l]}(s_j) = g(s_j T_{\text{eik;CM}}^{[l]}) c_{\text{out;CM}}^{[l]}(s_j) + g(-s_j T_{\text{eik;CM}}^{[l]}) c_{\text{in;CM}}^{[l]}(s_j).$$

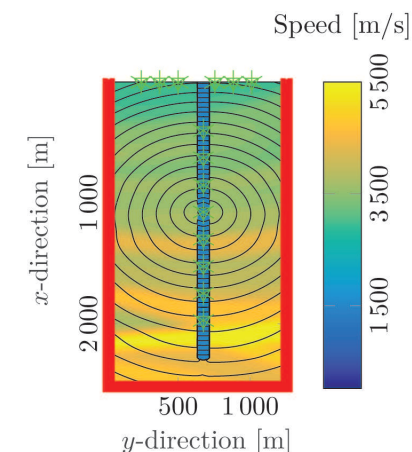
With 14 sources and 40 interpolation points we end up with 560 amplitudes for each of the four amplitudes $c_{\text{in/out;eik/CM}}$, which we compress to 30 each using an SVD. These compressed amplitudes are then used to construct the phase-preconditioned RKS on which the fine operator is projected.

For these types of configurations, the time window of interest tends to be very long due to the resonant nature of the configuration. Finite-difference time-domain therefore requires very long runtimes, whereas the proposed algorithm just needs to evaluate the ROM on more frequencies to avoid aliasing. The time-domain trace of the topmost source-receiver pair within the borehole is shown in Figure 6(a), where the emitted pulse bounces back and forth within the cavity. The ROM captures this resonant behavior, showing that the resonant modes are well approximated. In Figure 6(b) a trace is shown from a source within the borehole to a surface receiver. In this trace it can be seen that the interaction of the pulse with the smooth geology is modeled correctly next to the repetitive trace shape caused by the resonant cavity. An ordinary RKS method with no grid coarsening would perform well on this problem, since it is mainly dominated by the resonant cavity; however, this would require



(a) Simulated configuration with labeled sources. Eight source-receiver pairs are placed inside an elongated resonant cavity and six pairs are placed at the surface.

(b) Contour lines of the eikonal time T_{eik} used to approximate the propagative part of the solution.

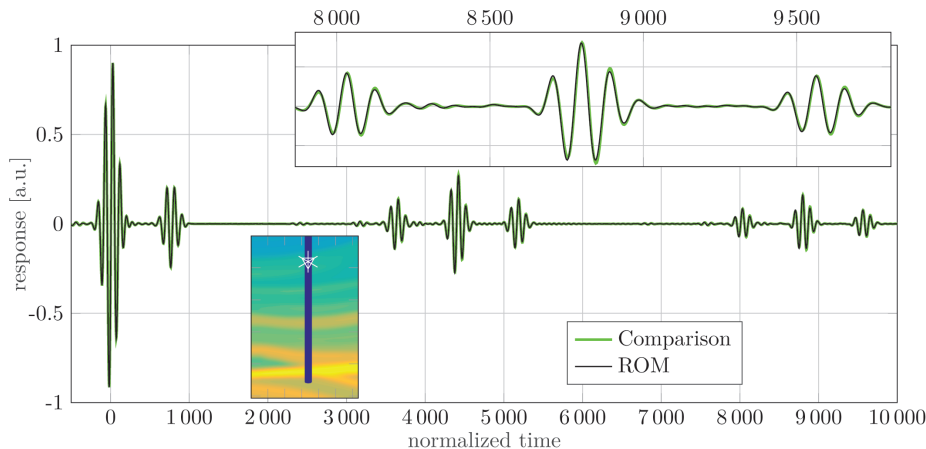


(c) Contour lines of the eikonal time $T_{\text{eik;CM}}$ used to approximate the cavity-modes of the solution.


FIG. 5. The source-receiver setup (a), and the eikonal times (b) and (c) used to decompose the fields into smooth amplitude-functions are depicted.

solutions of the wave equation on a much finer grid than the proposed approach. Furthermore, contrary to the proposed approach, the RKS approximation deteriorates as the configuration size and thus the propagation distance from the cavity to the receiver increase.

In this experiment we show that the developed algorithm shows potential for reduced-order modeling of resonant cavities within slowly varying media. The combination of an RKS method together with phase-preconditioning can approximate both



(a) Time-domain trace of the coinciding source receiver pair number 1 after $m = 40$ interpolation points, together with the comparison solution.



(b) Time-domain trace from source number 7 inside the borehole to the rightmost surface receiver number 14 after $m = 40$ interpolation points.

FIG. 6. Time-domain responses of the resonant cavity test-case.

resonant eigenmodes as well as propagative modes. We point out that this is just a first approach in order to include resonant structures into ROMs that are travel time dominated.

8. Discussion on parallel implementation. The numerical experiments of the previous section (using a serial MATLAB prototype code) showed significant compression of large-scale wave propagation due to phase-preconditioning. To see how observed dimensionality reduction can be translated to computational cost reduction using modern high-performance platforms, we consider the simplest parallel implementation, known in computer science literature as an “embarrassingly parallel workflow.”¹

Like the majority of the projection-based model reduction methods, the PPRKS can be split into basis construction and ROM evaluation stages, as summarized in

¹Term used for parallelization not requiring horizontal communication between nodes.

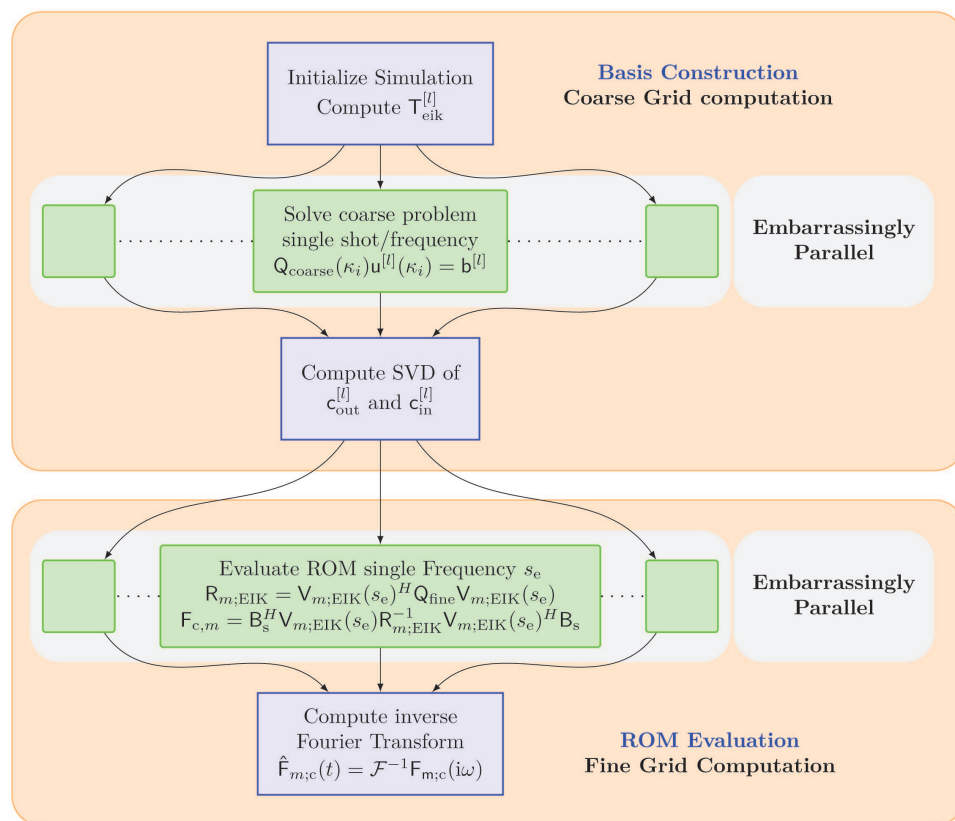


FIG. 7. Overview of the proposed algorithm. External embarrassing parallelism is symbolized by parallel blocks. Internal parallelism within a block is also possible.

TABLE 3
Cost estimates for PPRKS and RKS.

Step	PPRKS		Legend
	Computations per worker	# workers	
Eikonal	$O(N_f \log N_f)$	N_{src}	N_{src} # sources
Basis comp.	$O(N_c \psi(N_c))$	$N_{\text{src}} m_{\text{PPRKS}}$	$N_{c/f}$ # coarse-/fine-grid notes
SVD	$O(N_c [2N_{\text{src}} m_{\text{PPRKS}}]^2)$	1	m # ROM interpolation points
Eval	$O(N_f N_{\text{src}}^2 M_{\text{SVD}}^2)$	N_{eval}	N_{eval} # evaluation frequencies
	RKS		M_{SVD} size SVD compressed vectors
Basis comp.	$O(N_f \psi(N_f))$	$N_{\text{src}} m_{\text{RKS}}$	
SVD	$O(N_c [2N_{\text{src}} m_{\text{RKS}}]^2)$	1	m # ROM interpolation points
Eval	$O(N_f N_{\text{src}}^2 m_{\text{RKS}}^2)$	1	$\psi(N_{c/f})$ scaling function of Helmholtz solver

Figure 7. This figure is complemented by a summary in Tables 1 and 3, where we compare computational cost estimates for PPRKS with standard RKS neglecting $O(N_f)$ terms and considering only parallelism on the external level.

For both standard RKS and PPRKS the main cost of the first stage consists of the computation of the block-RKS and the rank-revealing subspace truncation via SVD. Phase-preconditioning adds the negligible cost of solving the eikonal equation and decomposing the waves into incoming/outgoing amplitudes via (4.14) and (4.15).

In Table 3 we assume that the block-RKS is computed by assigning solutions of Helmholtz problems for different frequencies and right-hand sides to separate workers, so that the PPRKS and RKS require $N_{\text{src}}m_{\text{PPRKS}}$ and $N_{\text{src}}m_{\text{RKS}}$ nodes, respectively. The PPRKS obviously reduces the number of the Helmholtz solves; however, in the parallel implementation the most important cost reduction lies in a single solve. In our case, this cost is critical due to the high complexity and poor internal parallel scalability of available Helmholtz solvers. In the table, $\psi(N)$ reflects this (usually faster than logarithmic) growth of the computational complexity of the Helmholtz solver.

Thus, the observed reduction in grid nodes from N_f to N_c , which lies between one and two orders of magnitude, can result in even stronger reductions of computation time. Subspace truncation is another poorly scalable bottleneck of the basis generation stage (e.g., see [23]) and the compound effect of the reduction of m_{PPRKS} and N_c compared to m_{RKS} and N_f is more than two orders.

The main cost of the second stage is the evaluation of the ROM frequency response at quadrature points in the frequency domain. In particular, the computation of the orthogonal basis and the Galerkin projection are the main bottlenecks with costs that grow linearly with respect to the fine grid dimension. The dimension of the PPRKS approximation space is the product of the size of the compressed amplitude space and the number of sources, which is usually of the same order as the dimension needed for standard RKS.² Nonetheless, storage of the space is reduced by a factor of $N_{\text{src}}N_f/N_c$ and the computation of the coarse grid amplitudes is obviously cheaper. However, the phase-preconditioned subspace is frequency-dependent, unlike the standard RKS. Therefore, the Galerkin projection should be computed for every frequency for the entire operator $Q(s)$. This is not a significant disadvantage thanks to the possibility of an embarrassingly parallel implementation; for every evaluation frequency a separate worker can be assigned. Moreover, the compressed tensor-product representation (5.12) allows efficient lower-level parallelization for the evaluation phase of PPRKS, i.e., columnwise, elementwise, and via domain-decomposition of the inner products. Solving the Galerkin system as well as carrying out inverse Fourier transforms to the time-domain do not depend on the grid size and their costs can be neglected.

We choose to benchmark the prototype implementation of algorithm in two parts. The basis construction is benchmarked on a CPU and the ROM evaluation on a GPU, as our algorithm is intended for the modern high-performance computing environment. Efficient Helmholtz solvers or solvers for large, sparse matrix systems are generally developed for CPUs. GPUs, however, are designed for fast, parallel computation of large inner products and therefore are excellent for the evaluation stage of the proposed model order reduction technique.

For the smooth geophysical structure example given in this paper with $N_f = 4 \cdot 10^5$, $N_{\text{src}} = 12$ we compare the most important computation times in Table 4. The basis computation is performed on a CPU³ and the ROM evaluation on a GPU.⁴ In the proposed algorithm we solve the wave equation on a coarse grid only, leading to a much lower cost in basis construction than standard RKS where fine-scale systems need to be solved. This is especially important considering that for large three-dimensional applications it can become infeasible to solve the wave equation on a

²Recall that this is due to the tensor-product structure of the PPRKS approximation space given by (5.12).

³Solved using UMFPACK v 5.4.0 on a 4-Core Intel i5-4670 CPU@3.40 GHz with parallel BLAS level-3 routines.

⁴Double precision python implementation on an Nvidia GTX 1080 Ti.

TABLE 4
Cost of the basis computation and evaluation of the ROM.

Basis computation comparison	Computation	Time	
Block solve fine grid	$\mathbf{Q}_{\text{fine}}(s_i)^{-1}\mathbf{B}$	10.3 s	
Single solve fine grid	$\mathbf{Q}_{\text{fine}}(s_i)^{-1}\mathbf{b}$	4.1 s	
Block solve coarse grid	$\mathbf{Q}_{\text{coarse}}(s_i)^{-1}\mathbf{B}$	0.6 s	
Single solve coarse grid	$\mathbf{Q}_{\text{coarse}}(s_i)^{-1}\mathbf{b}$	0.2 s	
Evaluation step	Computation	Time	Scaling
Computing phase-functions	$\exp(i\omega T_{\text{eik}})$	0.00546 s	$N_{\text{src}}N_{\text{f}}$
Hadamard products	$\exp(i\omega T_{\text{eik}})c_{\text{SVD}}$	0.01496 s	$M_{\text{SVD}}N_{\text{src}}N_{\text{f}}$
Galerkin inner product	$\mathbf{V}_{m;\text{EIK}}(s_e)^H \cdot \mathbf{Q}_{\text{fine}}\mathbf{V}_{m;\text{EIK}}(s_e)$	1.752 s	$N_{\text{f}}M_{\text{SVD}}^2N_{\text{src}}^2$

fine grid as the scaling function $\psi(\cdot)$ is much worse for three-dimensional systems than for two-dimensional systems. To show the cost of the evaluation of the ROM we benchmarked the evaluation of a single frequency s_e on a GPU. The used model parameters are $N_{\text{f}} = 4 \cdot 10^5$, $M_{\text{SVD}} = 100$, $N_{\text{src}} = 12$ and the results are given in Table 4. The computationally most involving part is the computation of the Galerkin inner product of left-hand vectors $\mathbf{V}_{m;\text{EIK}}(s_e)^H$ with the vectors $\mathbf{Q}_{\text{fine}}\mathbf{V}_{m;\text{EIK}}(s_e)$. Even for this relatively small example the computational cost of the solving a coarse system and projecting the ROM is smaller than evaluating the equation on a fine grid. We infer that especially for large-scale models these computation times become negligible with respect to basis construction, which scales worse. The phase-preconditioning approach drastically reduces the vertical communication of the algorithm as only coarse-grid amplitudes and phases need to be transferred to all workers instead of fine-grid RKS vectors. The ROM is essentially compressed and the storage is drastically reduced.

In summary, the computational cost is shifted from the poorly scalable basis construction to the highly scalable evaluation stage where inner products can be computed in an embarrassingly parallel fashion on multiple GPUs. We should also mention significant storage reduction due to phase-preconditioning as the amplitude basis is smaller than the standard RKS basis for the same accuracy and is stored on the coarse grid only, which significantly reduces vertical communication.

9. Conclusions. In this paper we have introduced the PPRKS for model order reduction and compression of wave propagation in unbounded domains targeting problems with large propagation distances. Preconditioning is achieved by splitting the RKS into incoming and outgoing waves and factoring out strongly oscillating phase terms using the WKB approximation. The remaining slowly varying amplitude terms are SVD-compressed and then used in the construction of the preconditioned projection space via combinations of the singular vectors of the compressed space and the WKB phase terms computed for different inputs (sources). Finally, the ROM is evaluated via structure-preserving model reduction.

Phase-preconditioning has multiple objectives, namely, reduction of the number of required interpolation points, right-hand sides, and spatial discretization. The number of interpolation points needed for a nonpreconditioned RKS method is fundamentally limited by the Nyquist frequency. However, phase-preconditioning weakens this dependence of the interpolation points on the Nyquist limit. We quantified this effect for one-dimensional SISO problems with piecewise constant coefficients, where the PPRKS solution is exact with the number of the RKS shifts equal to the number of the homogeneous layers, i.e., this number plays the same role as the problem dimensionality in a conventional RKS approach. Thus, in one dimension the number

of interpolation points needed is independent of the Nyquist rate. We do not have a rigorous estimate for the general case of multidimensional MIMO problems. However, numerical experiments show that the positive effects of preconditioning can increase due to simultaneous reduction of interpolation points and right-hand sides. In addition, factorization significantly relaxes requirements on the discretization grids for subspace computations, which is critical for large-scale problems due to the poor scalability of available Helmholtz solvers.

Furthermore, factorization reduces the computation cost and increases the model-reduction compression factor. More specifically, for a given approximation accuracy the SVD compressed amplitude space is much smaller than the RKS basis. Numerical experiments for sections of two-dimensional acoustic benchmark Marmousi problem show that the best cost reduction in subspace generation and compression is achieved for smooth wave speed profiles; however, our approach is still competitive for the discontinuous models and can even be adapted to include resonant substructures.

Finally, we point out that due to the tensor-product-like structure of the MIMO preconditioned projection space, it is larger than the space of compressed amplitudes and can even be comparable to the conventional block-RKSs required for the same accuracy.

However, unlike the subspace generation and compression, the projection is generally highly scalable and can be easily implemented in parallel on GPUs, leaving its computation time insignificant.

In this paper, we presented a prototype implementation of PPRKS for two-dimensional problems using serial computation; however, our eventual target is high-performance computing of large-scale two-dimensional seismic problems. In future work, we will also focus on optimal placement of the interpolation points. Specifically, we plan to investigate the approximation quality of the ROMs when we move the interpolation points away from the imaginary axis and into the complex plane. This can potentially improve both the approximation properties of the preconditioned RKS for the case of bounded time intervals and the performance of Helmholtz iterative solvers used for RKS construction. As a natural extension of PPRKS, we will also focus on the modeling of wave propagation in dispersive media using PPRKS, since this will not add additional costs to the evaluation stage. Finally, we note that WKB-like asymptotic solutions are available for many discrete and continuous dynamical systems, which opens up a number of possibilities to extend phase-preconditioning to such problems and related matrix-function computations, in particular if the cost of the solution of the shifted systems is dominant in the RKS algorithm.

REFERENCES

- [1] A. ANTOUNAS, C. BEATTIE, AND S. GUGERCIN, *Interpolatory model reduction of large-scale dynamical systems*, in Efficient Modeling and Control of Large-Scale Systems, J. Mohammadpour and K. Grigoriadis, eds., Springer, New York, 2010, pp. 3–58, <https://doi.org/10.1007/978-1-4419-5757-3>.
- [2] B. BECKERMANN AND L. REICHEL, *Error estimates and evaluation of matrix functions via the Faber transform*, SIAM J. Numer. Anal., 47 (2009), pp. 3849–3883.
- [3] C. M. BENDER AND S. A. ORSZAG, *Advanced Mathematical Methods for Scientists and Engineers*, Springer, New York, 1999.
- [4] P. BENNER, S. GUGERCIN, AND K. WILLCOX, *A survey of projection-based model reduction methods for parametric dynamical systems*, SIAM Rev., 57 (2015), pp. 483–531, <https://doi.org/10.1137/130932715>.
- [5] D. S. BINDEL, Z. BAI, AND J. W. DEMMEL, *Model reduction for RF MEMS simulation*, in Applied Parallel Computing. State of the Art in Scientific Computing: 7th International Workshop, PARA 2004, Lyngby, Denmark, 2004, J. Dongarra, K. Madsen, and

- J. Waśniewski, eds., *Lecture Notes in Comput. Sci.* 3732, Springer, New York, 2006, pp. 286–295, https://doi.org/10.1007/11558958_34.
- [6] A. BOURGEOIS, M. BOURGET, P. LAILLY, M. POULET, P. RICARTE, AND R. VERSTEEG, *Marmousi, model and data*, in *The Marmousi Experience*, European Association of Exploration Geophysicists, 1991, pp. 5–16.
 - [7] W. C. CHEW, *Electromagnetic theory on a lattice*, *J. Appl. Phys.*, 75 (1994), pp. 4843–4850, <https://doi.org/10.1063/1.355770>.
 - [8] E. T. CHUNG, Y. EFENDIEV, AND W. T. LEUNG, *Generalized multiscale finite element methods for wave propagation in heterogeneous media*, *Multiscale Model. Simul.*, 12 (2014), pp. 1691–1721, <https://doi.org/10.1137/130926675>.
 - [9] V. DRUSKIN, S. GÜTTEL, AND L. KNIZHNERMAN, *Near-optimal perfectly matched layers for indefinite Helmholtz problems*, *SIAM Rev.*, 58 (2016), pp. 90–116, <https://doi.org/10.1137/140966927>.
 - [10] V. DRUSKIN, A. V. MAMONOV, AND M. ZASLAVSKY, *Multiscale S-fraction reduced-order models for massive wavefield simulations*, *Multiscale Model. Simul.*, 15 (2017), pp. 445–475, <https://doi.org/10.1137/16M1072103>.
 - [11] V. DRUSKIN AND R. REMIS, *A Krylov stability-corrected coordinate-stretching method to simulate wave propagation in unbounded domains*, *SIAM J. Sci. Comput.*, 35 (2013), pp. B376–B400, <https://doi.org/10.1137/12087356X>.
 - [12] V. DRUSKIN, R. REMIS, AND M. ZASLAVSKY, *An extended Krylov subspace model-order reduction technique to simulate wave propagation in unbounded domains*, *J. Comput. Phys.*, 272 (2014), pp. 608–618, <https://doi.org/10.1016/j.jcp.2014.04.051>.
 - [13] V. DRUSKIN AND M. ZASLAVSKY, *On convergence of Krylov subspace approximations of time-invariant self-adjoint dynamical systems*, *Linear Algebra Appl.*, 436 (2012), pp. 3883–3903, <https://doi.org/10.1016/j.laa.2011.02.039>.
 - [14] B. ENGQUIST AND L. YING, *Sweeping preconditioner for the Helmholtz equation: Moving perfectly matched layers*, *Multiscale Model. Simul.*, 9 (2011), pp. 686–710.
 - [15] T. GÖCKLER AND V. GRIMM, *Convergence analysis of an extended Krylov subspace method for the approximation of operator functions in exponential integrators*, *SIAM J. Numer. Anal.*, 51 (2013), pp. 2189–2213.
 - [16] S. GÜTTEL, *Rational Krylov Methods for Operator Functions*, Ph.D. thesis, Institut für Numerische Mathematik und Optimierung, Technische Universität Bergakademie Freiberg, 2010.
 - [17] E. HABER AND S. MACLACHLAN, *A fast method for the solution of the Helmholtz equation*, *J. Comput. Phys.*, 230 (2011), pp. 4403–4418, <https://doi.org/10.1016/j.jcp.2011.01.015>.
 - [18] I. GOHBERG, P. LANCASTER, AND L. RODMAN, *Matrix Polynomials*, *Classics in Appl. Math.* 58, SIAM, Philadelphia, 2009.
 - [19] A. ISERLES, S. NØRSETT, AND S. OLVER, *Highly oscillatory quadrature: The story so far*, in *Numerical Mathematics and Advanced Applications*, Springer, New York, 2006, pp. 97–118.
 - [20] L. KNIZHNERMAN, V. DRUSKIN, AND M. ZASLAVSKY, *On optimal convergence rate of the rational Krylov subspace reduction for electromagnetic problems in unbounded domains*, *SIAM J. Numer. Anal.*, 47 (2009), pp. 953–971, <https://doi.org/10.1137/080715159>.
 - [21] Y. E. LI AND L. DEMANET, *Phase and amplitude tracking for seismic event separation*, *Geophysics*, 80 (2015), pp. WD59–WD72.
 - [22] W. A. MULDER AND R.-E. PLESSIX, *Time- versus frequency-domain modelling of seismic wave propagation*, in *Proceedings of the 64th Annual Conference and Exhibition*, 2002, pp. 27–30, http://aniso.citg.tudelft.nl/Private/Papers/Conferences/EAGE02_E015.pdf.
 - [23] M. O’CONNELL, M. E. KILMER, E. DE STURLER, AND S. GUGERCIN, *Computing reduced order models via inner-outer Krylov recycling in diffuse optical tomography*, *SIAM J. Sci. Comput.*, 39 (2017), pp. B272–B297, <https://doi.org/10.1137/16M1062880>.
 - [24] A. RUHE, *Rational Krylov sequence methods for eigenvalue computation*, *Linear Algebra Appl.*, 58 (1984), pp. 391–405, [https://doi.org/10.1016/0024-3795\(84\)90221-0](https://doi.org/10.1016/0024-3795(84)90221-0).
 - [25] A. A. SAMARSKIJ AND E. S. NIKOLAEV, *Numerical Methods for Grid Equations, Volume I, Direct Methods*, Birkhauser Basel, Basel, Switzerland, 1989.
 - [26] J. A. SETHIAN, *A fast marching level set method for monotonically advancing fronts*, *Proc. Nat. Acad. Sci. U.S.A.*, 93 (1996), pp. 1591–159.
 - [27] A. TOWNSEND AND L. N. TREFETHEN, *Continuous analogues of matrix factorizations*, *Proc. A*, 471 (2014), <https://doi.org/10.1098/rspa.2014.0585>.