

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/380535031>

Iterative PDE-Constrained Optimization for Seismic Full-Waveform Inversion

Article in Computational Mathematics and Mathematical Physics · May 2024

DOI: 10.1134/S0965542524700192

CITATIONS

0

READS

77

4 authors, including:



Mikhail Malovichko

Aramco Innovations

40 PUBLICATIONS 170 CITATIONS

[SEE PROFILE](#)



Nikolay I. Khokhlov

Moscow Institute of Physics and Technology

116 PUBLICATIONS 829 CITATIONS

[SEE PROFILE](#)



Petrov Igor

Moscow Institute of Physics and Technology

283 PUBLICATIONS 1,808 CITATIONS

[SEE PROFILE](#)

Iterative PDE-Constrained Optimization for Seismic Full-Waveform Inversion

M. S. Malovichko^a, A. Orazbayev^a, N. I. Khokhlov^{a,*}, and I. B. Petrov^a

^a*Moscow Institute of Physics and Technology, Moscow, Russia*

* *e-mail: khokhlov.ni@mipt.ru*

Received November 2, 2023; revised November 6, 2023; accepted February 6, 2024

Abstract—This paper presents a novel numerical method for the Newton seismic full-waveform inversion (FWI). The method is based on the full-space approach, where the state, adjoint state, and control variables are optimized simultaneously. Each Newton step is formulated as a PDE-constrained optimization problem, which is cast in the form of the Karush–Kuhn–Tucker (KKT) system of linear algebraic equations. The KKT system is solved inexactly with a preconditioned Krylov solver. We introduced two preconditioners: the one based on the block-triangular factorization and its variant with an inexact block solver. The method was benchmarked against the standard truncated Newton FWI scheme on a part of the Marmousi velocity model. The algorithm demonstrated a considerable runtime reduction compared to the standard FWI. Moreover, the presented approach has a great potential for further acceleration. The central result of this paper is that it establishes the feasibility of Newton-type optimization of the KKT system in application to the seismic FWI.

Keywords: seismic full-waveform inversion, PDE-constrained optimization, Karush–Kuhn–Tucker, Newton FWI

DOI: 10.1134/S0965542524700192

1. INTRODUCTION

The seismic full-waveform inversion (FWI) is a coefficient inverse problem aiming to estimate the subsurface distribution of material properties based on recorded seismic data. The FWI has emerged almost four decades ago [1–3], and has evolved into a powerful method for reconstruction of the subsurface properties, see reviews [4, 5].

The FWI is formulated as a Tikhonov-like minimization procedure,

$$\underset{\gamma}{\text{minimize}} \|d - \mathcal{F}(\gamma)\|^2 + \mathcal{R}(\gamma), \quad (1)$$

where γ is a function of material parameters (say, the squared P-wave slowness), d is a vector of observed data, \mathcal{F} is a nonlinear forward-problem operator, and \mathcal{R} is the stabilizing term. We are interested in the frequency-domain acoustic FWI and, thus, operator \mathcal{F} implies the solution of Helmholtz's equation. Many variants of basic formulation (1) exist: \mathcal{R} can take various forms or even be omitted, the residual can be weighted, and so on.

Optimization (1) is usually performed by either quasi-Newton or Newton methods. The quasi-Newton methods dominate industrial applications due to their relative numerical efficiency [6–10] (among many others). These methods account for nonlinearity partially, usually by approximating the Hessian with a diagonal matrix. Typical representatives of this class are the nonlinear conjugate gradient (NLCG) method, where a diagonal matrix scales the gradients [11, 12], and L-BFGS [13]. Though computationally tractable, the quasi-Newtonian methods may suffer from convergence stagnation.

The Newton methods achieve faster and more reliable convergence, which catch up with the nonlinearity through quadratic approximation of the target function, see [14–16]. The often-used Gauss–Newton variant computes the Hessian matrix approximately by linearizing the forward-problem operator. The main disadvantage of the Newton methods applied to (1) is an enormous computational burden stemming from the need to solve a large dense system of linear equations at each Newton iteration. The system is solved iteratively, usually with the conjugate gradients (CG). The matrix is ill-conditioned, and therefore each Newton (external, nonlinear) iteration may require hundreds of internal (linear) iterations. In addi-

tion, no efficient preconditioner is known to improve this matrix's conditioning. For these reasons, numerical implementations of the Newton FWI are far less common than the quasi-Newton ones.

It is natural to ask if we can obtain the fast Newton-like convergence with a less resource-demanding procedure. The ultimate goal of this paper is to demonstrate that the answer is yes.

Optimization (1) is known as *the reduced-space approach* in the optimization community. Its alternative is *the full-space approach*, also known as *all-at-once*. In the reduced-space approach, only material parameters γ are optimized, whereas the connection between the wavefield in the medium and γ is prescribed by the forward-problem operator \mathcal{F} . In the full-space approach, one simultaneously optimizes the control, state, and adjoint-state variables. Upon formulating the necessary conditions and discretizing the PDE, one obtains a nonlinear system of algebraic equations, known as the Karush–Kuhn–Tucker (KKT) system, which is sparse. Its sparsity is a clear advantage over formulation (1). The main challenge here is that the problem becomes many times bigger than the reduced-space approach because the state and adjoint-state variables for all sources are included in optimization.

The full-space approach has been introduced to geophysics in the seminal paper [17], followed by [18, 19]. Those authors cast the geophysical electromagnetic inversion as a full-space optimization problem, derived the KKT system and applied the Newton algorithm to it. Despite spectacular results, the full-space approach had been considered impractical for years due to the enormous size of the linear system, arising on each Newton step. Regarding FWI, the KKT optimality conditions were derived in [15, 16], but were viewed as a theoretical tool rather than a basis for numerical implementation.

Significant progress has been made in this direction during the last decade. Papers [20–22] report the application of the *penalty method* to the FWI, marking an important step toward the full-space approach. The penalty method has some theoretical drawbacks connected to the choice of the penalty parameter [23] (but see recent developments [24]). The *augmented Lagrangian* (AL) method is a viable alternative to the penalty method. A variant of it, known as the *alternating-direction method of multipliers* (ADMM) [25], was applied to the FWI with great success, see [26, 27]. In essence, the ADMM tackles the full-space optimization by making alternating steps in the state, adjoint-state, and control state spaces. Still, it is generally accepted that the main weakness of the ADMM compared to the Newton iterations is its slow convergence [25].

This paper studies the full-space optimization based on applying a Newton solver to the original KKT matrix. The crucial component of our approach is a suitable preconditioner based on a block-triangular factorization of the permuted KKT matrix combined with approximations of individual blocks. Similar ideas were exploited optimization with Maxwell's equations [18] and the Stokes equation [28]; thus, this paper extends this framework to Helmholtz's equation. We benchmarked our method against the standard reduced-space Gauss–Newton FWI with the internal CG solver. In the numerical experiments conducted the Marmousi model, we observed the $6\times$ speedup. The method has the potential for further acceleration. In particular, it will benefit from any method designed to approximate the reduced-space Hessian, such as L-BFGS. The central result of this paper is that it establishes the feasibility of Newton-type optimization of the KKT system in application to seismic FWI.

The paper is organized as follows. In Section 2, we formulate the forward problem. The reduced-space Gauss–Newton inversion is considered in Section 3. Sections 4 and 5 are dedicated to the full-space approach and its iterative solution. Numerical experiments are presented in Section 6. Concluding remarks are given in Section 7.

2. THE FORWARD PROBLEM

In this section, we briefly review the forward problem. Let $\Omega \subset \mathbb{R}^2$ be a bounded domain with boundary $\partial\Omega$. Consider the 2D Helmholtz's equation in the following form:

$$\begin{aligned} -Z \frac{\partial}{\partial x} \left(\frac{1}{X} \frac{\partial U}{\partial x} \right) - X \frac{\partial}{\partial z} \left(\frac{1}{Z} \frac{\partial U}{\partial z} \right) - \omega^2 \gamma X Z U &= X Z F, \quad \text{in } \Omega, \\ U &= 0, \quad \text{on } \partial\Omega. \end{aligned} \quad (2)$$

Here, $F(x, z)$ are the right-hand side, $U(x, z, \omega)$ are the acoustic field (monochromatic pressure), ω is the circular frequency, $\gamma(x, z)$ is the squared slowness, $\gamma = c^{-2}$, where $c(x, z)$ is the speed of sound. Complex-valued functions of one variable $X(x)$ and $Z(z)$ implement the perfectly-matched layers (PML). They deviate from 1.0 only in the damping region on the sides of the computational domain to suppress the side reflections (see Fig. 1). On the PML implementation, see [29] and references therein.

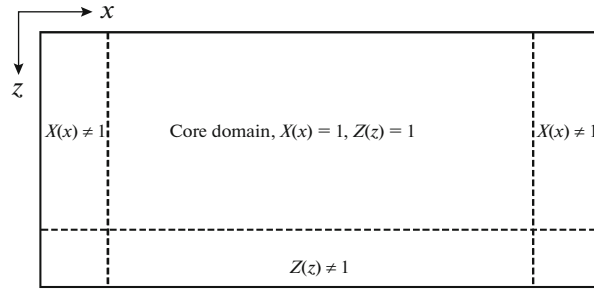


Fig. 1. Schematic representation of the computational domain with a PML region. In the PML region $X(x) \neq 1$, $Z(z) \neq 1$.

Let us rewrite (2) with a more compact notation as follows:

$$\begin{aligned} \mathcal{D}U &= F, \quad \text{in } \Omega, \\ U &= 0, \quad \text{on } \partial\Omega, \end{aligned} \quad (3)$$

where

$$\mathcal{D}(\cdot) := -Z \frac{\partial}{\partial x} \frac{1}{X} \frac{\partial}{\partial x} (\cdot) - X \frac{\partial}{\partial z} \frac{1}{Z} \frac{\partial}{\partial z} (\cdot) - \omega^2 \gamma XZ(\cdot). \quad (4)$$

Note that we replaced XZF with F in (3) because we assume that F is supported only in the core domain.

In a seismic survey, the acoustic field is excited by many monochromatic sources independent of each other. Thus, we have a series of problems,

$$\begin{aligned} \mathcal{D}U_k &= F_k, \quad \text{in } \Omega, \\ U_k &= 0, \quad \text{on } \partial\Omega, \end{aligned} \quad (5)$$

where index k designates the source index. Let us combine solutions U_k and right-hand sides F_k into vector variables to make the formulas even more compact. Starting from now on, U is a direct product of all U_k , F is a direct product of F_k :

$$U := \prod_{k=0}^{K-1} U_k, \quad F := \prod_{k=0}^{K-1} F_k, \quad k = 0, 1, \dots, K-1. \quad (6)$$

With these notations, a set of K problems (5) can be compactly written as

$$\begin{aligned} \mathcal{D}U &= F, \quad \text{in } \Omega, \\ U &= 0, \quad \text{on } \partial\Omega. \end{aligned} \quad (7)$$

Now we introduce the observation operator \mathcal{Q} which takes the acoustic field and returns a vector of data d . We define \mathcal{Q} through the convolution with the Dirac δ -function (although any linear functional can be used instead) as follows,

$$d = \mathcal{Q}U := \left\{ \int_{\Omega} \delta(x_{kj} - x) U_k dx \right\}, \quad k = 0, 1, \dots, K-1, \quad j = 0, 1, \dots, N_k - 1, \quad (8)$$

where x_{kj} is the position of the j th receiver of the k th source, N_k is the number of receivers for the k th source, $d \in \mathbb{C}^N$, $N = \sum_{k=0}^{K-1} N_k$ is the total number of receivers for all sources.

In what follows, we will use the discrete form of the forward problem. The finite-difference (FD) discretization is assumed because it is used in our numerical experiments. However, we might as well have performed discretization by the finite-element or spectral-element methods.

Let u_k , f_k , and s be finite-difference approximations of U_k , F_k , and γ , respectively. We introduce the following notation:

$$u = \prod_{k=0}^{K-1} u_k, \quad f = \prod_{k=0}^{K-1} f_k, \quad k = 0, \dots, K-1. \quad (9)$$

A set of discrete forward problems reads

$$Au = f, \quad (10)$$

where A is a block-diagonal matrix. Each diagonal block A_j of matrix A is a sparse complex non-Hermitian matrix. If the PMLs are not used, then A_j turns into the real symmetric indefinite matrix,

$$-\Delta - \omega^2 M, \quad (11)$$

where Δ is a discrete Laplacian, M is a diagonal mass matrix, containing values of s on its diagonal $M = Is$ (it can be non-diagonal in some FD schemes [30]).

3. THE REDUCED-SPACE GAUSS–NEWTON INVERSION

In this section, we review the reduced-space Gauss–Newton inversion (RSGN). Since we will eventually end up with the FD approximation, and thus discrete variables will be grid functions, it is easier to formulate the inverse problem in the discrete form (the *discretize-then-optimize* approach). In the finite-element framework, a more convenient way would be to formulate the optimality conditions in function spaces and then move on to discretization (the *optimize-then-discretize* approach) (see [31]).

Let S be a nonlinear operator that maps a given s to u ,

$$u = S(s). \quad (12)$$

Let Q be a finite-dimensional operator that maps a given discrete acoustic field u to data d ,

$$d = Qu. \quad (13)$$

Discrete coefficient s is iterated as $s_{n+1} = s_n + \delta s$, where δs , known as the model update. On each iteration, s is a solution to the following minimization problem,

$$\underset{\delta s}{\text{minimize}} \frac{1}{2} \|W(r - QG\delta s)\|_2^2 + \frac{\varepsilon}{2} \|s_n + \delta s\|_2^2. \quad (14)$$

Here, W is a real-valued diagonal matrix of weights designed primarily to balance signal attenuation with offset from the source. Vector r is the data residual,

$$r = d - QS(s_n), \quad (15)$$

G is a matrix of the Frechét operator of S at point $s = s_n$, i.e.

$$G = \left. \frac{\partial S}{\partial s} \right|_{s=s_n}. \quad (16)$$

Solution to (14) satisfies

$$(G^*Q^*W^TWQG + L)\delta s = W^TG^*Q^*r - Ls_n, \quad (17)$$

where $L = \varepsilon I$. The following formula defines matrix G ,

$$G = A^{-1}P, \quad (18)$$

where P is a diagonal matrix containing forward solution u_n at $s = s_n$,

$$P := \omega^2 Iu_n. \quad (19)$$

To verify (18), (19), one can subtract the two discrete Helmholtz's equations

$$\begin{aligned} (-\Delta - \omega^2 I(s_n + \delta s))(u_n + \delta u) &= f, \\ (-\Delta - \omega^2 Is_n)u_n &= f, \end{aligned} \quad (20)$$

although the derivation can be conducted directly in the Hilbert spaces. In geophysics, the matrix

$$J := QA^{-1}P, \quad (21)$$

is known as the Jacobian, the matrix

$$H := J^*W^TWJ + L \quad (22)$$

is the (regularized) Hessian, and vector

$$g := W^T J^* r - Ls_n \quad (23)$$

is called the gradient. Using this definitions, we rewrite (17) as

$$H \delta s = g. \quad (24)$$

For any $\varepsilon > 0$ matrix H is positive-definite, and thus (24) has a unique solution. It should be noted that, even if $\varepsilon = 0$ (singular H), then a Krylov solver will deliver a meaningful estimate of δs , that is, a projection to a corresponding Krylov subspace.

The system matrix in (24) is dense and large. In real-life 3D problems, the system (24) can be solved only iteratively, with the CG being a natural choice. We will abbreviate this algorithm as RSGN-CG. The CG multiplies the system matrix $J^* J$ by a vector once per iteration. Because of (21) it implies two linear solves per CG iteration—the forward simulation with A and the adjoint simulation with A^* . We can rewrite (24) as

$$\left(\sum_{j=0}^{K-1} J_j^* W_j^T W_j J_j + L \right) \delta s = \sum_{j=0}^{K-1} W_j^T J_j^* r_j - Ls_n. \quad (25)$$

The forward and adjoint simulations for different sources can be performed in parallel but these results must be combined at each CG iteration.

However, the most critical problem is a very slow convergence of the CG. The system matrix (24) has a very high condition number for all meaningful values of ε . Thus, tens or hundreds of forward simulations (for all sources) may be needed to generate a good update δs . In addition, preconditioning of a regularized normal system of linear equations is notoriously difficult.

4. THE FULL-SPACE GAUSS–NEWTON INVERSION

In this section, we describe the full-space Gauss–Newton (FSGN) inversion. We start by noting that the discrete optimization problem (14) is equivalent to the following constrained optimization problem,

$$\underset{\delta s}{\text{minimize}} \Psi := \frac{1}{2} \|W(r - Q\delta u)\|_2^2 + \frac{\varepsilon}{2} \|s_n + \delta s\|_2^2, \quad \text{s.t.} \quad A\delta u = P\delta s. \quad (26)$$

Here n is the Newton iteration number, δu is the update to u_n , i.e., $u_{n+1} \approx u_n + \delta u$. The target functional Ψ is quadratic and bounded away from zero for any $\varepsilon > 0$. The constraints are linear. For such problems, known as linear-quadratic, the existence of unique state and control variables are readily established, for example, ([32], Theorem 1.43). Moreover, since we start from the discrete formulation, the solvability is proved simply by establishing the equivalence between the matrix form and the normal system (24), see below.

Let us form the Lagrangian,

$$\mathcal{L} := \Psi + \lambda^*(A\delta u - P\delta s), \quad (27)$$

where λ is the adjoint-state variable (a complex-valued grid function). By applying the optimality conditions [32], we obtain the KKT system,

$$\begin{aligned} -Q^* W^T W(r - Q\delta u) + A^* \lambda &= 0, \\ L(s_n + \delta s) - P^* \lambda &= 0, \\ A\delta u - P\delta s &= 0. \end{aligned} \quad (28)$$

This system can be rewritten in the matrix form,

$$\underbrace{\begin{bmatrix} F & O & A^* \\ O & L & -P^* \\ A & -P & O \end{bmatrix}}_{\mathcal{M}} \underbrace{\begin{bmatrix} \delta u \\ \delta s \\ \lambda \end{bmatrix}}_{\xi} = \underbrace{\begin{bmatrix} Q^* W^T W r \\ -Ls_n \\ 0 \end{bmatrix}}_b, \quad (29)$$

where $F = Q^* W^T W Q$.

Let us prove that (29) is equivalent to (24) for it establishes solvability of the KKT system. Using the third equation of (29), we compute δu ,

$$\delta u = A^{-1}P\delta s. \quad (30)$$

Using this result, we compute λ from the first equation,

$$\lambda = A^{-*}(Q^*r - Q^*QA^{-1}P\delta s), \quad (31)$$

where notation A^{-*} means $(A^*)^{-1}$. Now we substitute (30) and (31) into the second equation of (29) and get

$$(P^*A^{-*}Q^*QA^{-1}P + L)\delta s = P^*A^{-*}Q^*r - Ls_n. \quad (32)$$

From (21) we see that (32) is equivalent to (24).

From the preceding discussion, it is clear that the model update δs computed from either the KKT system (29) or the normal system (24) is identical, provided that the linear systems are solved exactly. The difference between the two approaches lies in the structure of the system matrix. The system matrix in (24) is complex-valued, dense, Hermitian positive-definite, and relatively small. For small- to medium-size 3D problems with rigorously selected FD grids and frequencies, satisfactory results have been obtained with the use of the sparse LU decomposition of the system A [8, 33]. In general, however, a parallel on-the-fly application of A^{-1} and A^{-*} is the only option. The system matrix in (29) is sparse complex-valued Hermitian indefinite. It is many times larger than the normal matrix, and it precludes from applying direct solvers in any form. Its iterative solution is considered in the next section.

5. ITERATIVE SOLUTION OF THE KKT SYSTEM

In this section, we design an iterative solver for the KKT system (29). The system matrix \mathcal{M} is can be explicitly partitioned across sources as follows,

$$\begin{bmatrix} F_0 & & & O & A_0^* & & \\ & \ddots & & & & \ddots & \\ & & F_{K-1} & & & & A_{K-1}^* \\ O & & & L & -P_0^* & \dots & -P_{K-1}^* \\ A_0 & & & -P_0 & & & \\ & \ddots & & \vdots & & & O \\ & & A_{K-1} & -P_{K-1} & & & \end{bmatrix} \begin{bmatrix} \delta u_0 \\ \vdots \\ \delta u_{K-1} \\ \delta s \\ \lambda_0 \\ \vdots \\ \lambda_{K-1} \end{bmatrix} = \begin{bmatrix} Q_0^*W_0^TW_0r_0 \\ \vdots \\ Q_{K-1}^*W_{K-1}^TW_{K-1}r_{K-1} \\ -Ls_n \\ 0 \\ \vdots \\ 0 \end{bmatrix}. \quad (33)$$

Block F is a positive semidefinite matrix, which is very sparse. If receivers are located at the grid nodes, then F is a diagonal matrix where the number of non-zero entries in each F_j block equals the number of data points for the j th source. Block A is a block-diagonal matrix of full rank whose spectrum lies in the left and right half-planes. Block L is positive-definite here but may become positive semidefinite if L represents the Laplacian. In our case, $L = \epsilon I$ and so it is diagonal. Block P is a rectangular matrix consisting of sparse blocks. All submatrices of P are diagonal; see definition (19). The pattern of the matrix \mathcal{M} for a test problem consisting of 25 sources is presented in Fig. 2.

Our preconditioner is based on a block-triangular decomposition, which is one of the most efficient strategies for preconditioning saddle-point problems. In the context of the KKT matrix, this approach was considered in [18, 28].

The matrix \mathcal{M} has singular (1, 1) and (3, 3) blocks, and so the direct block-LU decomposition is not possible. Thus, we perform block-LU factorization of the permuted matrix as follows,

$$\begin{bmatrix} A & O & -P \\ F & A^* & O \\ O & -P^* & L \end{bmatrix} = \begin{bmatrix} I & O & O \\ FA^{-1} & I & O \\ O & -P^*A^{-*} & I \end{bmatrix} \begin{bmatrix} A & O & -P \\ O & A^* & FA^{-1}P \\ O & O & H \end{bmatrix}, \quad (34)$$

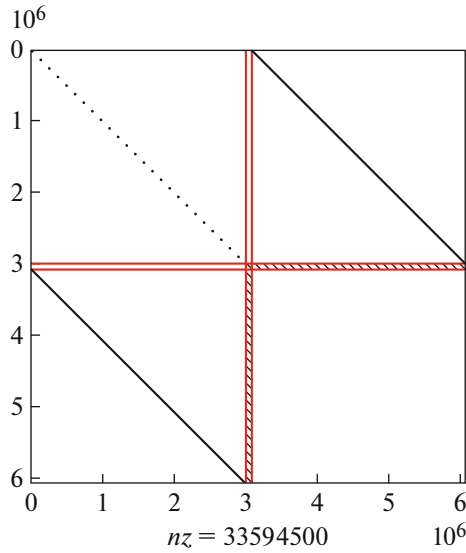


Fig. 2. Partitioning (red lines) and non-zero elements (black dots) of the KKT matrix for a test problem. The problem contains 25 sources, each system matrix is 120000×120000 . The size of the L block is 72000×72000 .

where $H = L + P^* A^{-*} F A^{-1} P$. The U -factor is an excellent preconditioner since applying its inverse to the initial (permuted) matrix results in a matrix whose spectrum consists of a single eigenvalue equal to 1. This approach would be equivalent to the reduced-space method and has no advantages because it requires a linear solve with the reduced Hessian H .

Let us approximate (34) by assuming $F \approx 0$, which implies $H \approx L$. Plugging $F = 0$, $H = L$ into (34), we get

$$\begin{bmatrix} I & O & O \\ O & I & O \\ O & -P^* A^{-*} & I \end{bmatrix} \begin{bmatrix} A & O & -P \\ O & A^* & O \\ O & O & L \end{bmatrix} = \begin{bmatrix} A & O & -P \\ O & A^* & O \\ O & -P^* & L \end{bmatrix}. \quad (35)$$

By reversing the permutation, we define the preconditioner to the original system matrix \mathcal{M} as follows,

$$\mathcal{P} := \begin{bmatrix} O & O & A^* \\ O & L & -P^* \\ A & -P & O \end{bmatrix}. \quad (36)$$

The preconditioned matrix, $\mathcal{P}^{-1} \mathcal{M}$ is non-symmetric; therefore, a non-symmetric Krylov solver is required to solve it. We selected the GMRes and abbreviated the resulting method as FSGN-GMRes.

The preconditioner (36) has several appealing properties. First, it requires three forward solves with matrices A , A^* , and L , i.e. almost the same work as needed for a single iteration of the CG with a reduced Hessian. Second, the A and A^* both split into K independent linear systems. Thus, the solutions with blocks (1, 3) and (3, 1) can be performed in parallel. Third, the preconditioner performs approximation $H \approx L$, but other approximations to H are possible. In particular, one can use the L-BFGS approximation in place of block L (see [28]).

Now, we define another preconditioner, which is obtained from \mathcal{P} by approximating A by another matrix \tilde{A} as follows:

$$\tilde{\mathcal{P}} := \begin{bmatrix} O & O & \tilde{A}^* \\ O & L & -P^* \\ \tilde{A} & -P & O \end{bmatrix}. \quad (37)$$

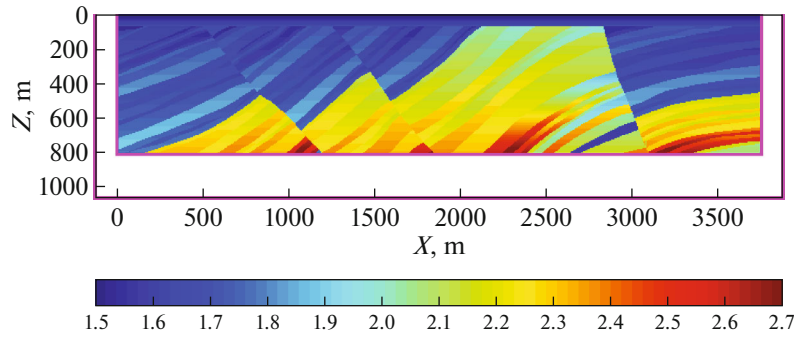


Fig. 3. The distribution of the P-wave velocity, c , used to simulate synthetic data. The numbers on the color scale are the velocity in km/s. The maroon lines on the sides indicate the boundaries of the PML region.

In this paper, we utilized the incomplete LU (ILU) factorization to approximate A . This choice is motivated by the fact that the ILU is widely available and requires no tuning apart from selecting the fill-in level.

We will demonstrate below that the FSGN-GMRes with the $\tilde{\mathcal{P}}$ preconditioner is superior to the RSGN-CG. Still, it is important to recognize that more efficient strategies to approximate A are likely to exist. In particular, the shifted-Laplacian preconditioner leveraged with the multigrid [34] deserves a dedicated study.

6. NUMERICAL EXPERIMENTS

This section is dedicated to numerical experiments. The optimization part of FWI was programmed in Python, whereas the forward and adjoint simulation problem is coded in C++. The forward problem was solved directly with a sparse LU factorization. All calculations were serial.

We conducted numerical experiments using a part of the Marmousi velocity model (Fig. 3). This P-wave velocity distribution was used in the data simulation, representing the geologic medium we want to reconstruct.

The purpose of the first experiment was to compare, inside a single GN step, the performance of the CG applied to the normal system matrix with the preconditioned GMRes applied to the KKT matrix. With the GMRes, we checked two preconditioners, \mathcal{P} and $\tilde{\mathcal{P}}$. The velocity distribution shown in Fig. 4a was the initial model for the inversion. The computed slowness updates are compared in Figs. 4b–4d.

The preconditioner $\tilde{\mathcal{P}}$ was created by making use of ILU with the fill-in parameter 21. It is not possible to compare the CG and GMRes directly, because the CG computes $E_{cg} := \|H\delta s_n - g\|_2$ at each iteration, whereas GMRes computes $E_{gmres} := \|\mathcal{M}\xi_n - b\|_2$. To make the comparison, we separately computed values of E_{cg} from the GMRes iterates ξ_n . The convergence of the iterative methods is presented in Fig. 5. The CG and GMRes achieved approximately the same tolerance on iterations 19 and 30, respectively (Fig. 5). As expected, all methods generated almost identical updates. The convergence rate of the CG is the same as that of the GMRes with the \mathcal{P} , although the GMRes happened to reach the desired tolerance by 3 iterations faster. The convergence of the GMRes with $\tilde{\mathcal{P}}$ is slower than that with \mathcal{P} , but the cost is much lesser. The performance of the algorithms is compared in Table 1. We see the dramatic acceleration of matrix-vector multiplication in the GMRes compared to the CG. This is because the linear solves with factorized matrices \tilde{A} and \tilde{A}^* inside $\tilde{\mathcal{P}}$ are much cheaper than the linear solves with A and A^* during the CG matrix-vector multiplication. Another observation is that we needed a high value of fill-in (we used 21) to have an acceptable approximation $A \approx \tilde{A}$.

In the second numerical experiment, we demonstrate that our approach serves as an efficient engine for the multi-frequency FWI. The data were simulated at frequencies from 5 to 40 Hz with a step of 2.5 Hz. At each frequency in the sequence, from lower to higher, a single GN step was performed, and the updated model was passed to the next frequency. The regularization parameter ε was changing linearly from 10 to 1E5. The number of internal GMRes iterations was set to 30. As indicated in Fig. 6, the data misfit was

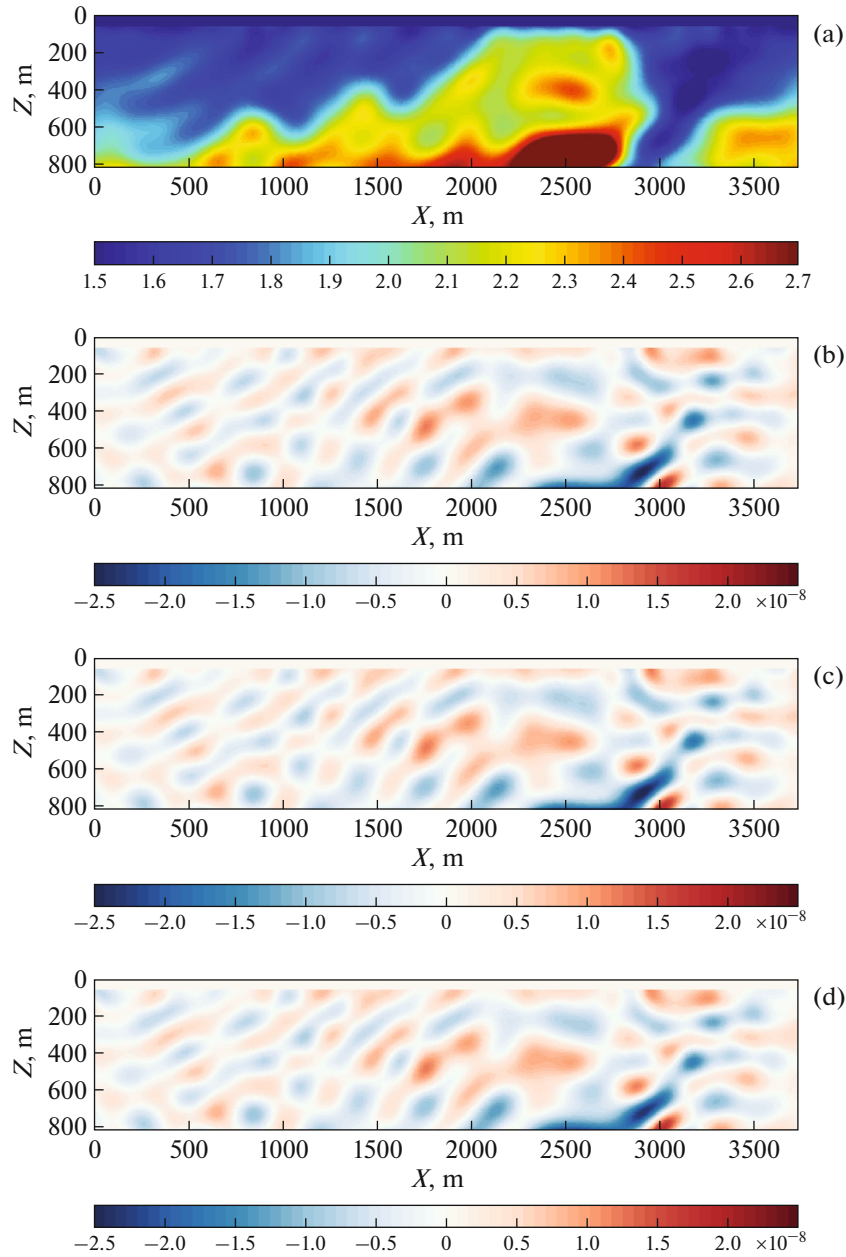


Fig. 4. Comparison of the two slowness updates computed at the end of a single GN iteration. Panel (a) shows the initial velocity distribution. Panel (b) depicts the slowness update obtained by applying the CG solver to the normal system of linear equations. Panel (c) shows the slowness update by the GMRes solver with \mathcal{P} applied to the KKT system. Panel (d) same as (c) but with preconditioner $\tilde{\mathcal{P}}$. Note that in (a) color represents the P-wave velocity in km/s, whereas in (b–d) color represents update to the squared slowness in s^2/m^2 .

improved after the inversion. The improvement becomes more pronounced at higher frequencies because the data magnitudes increase with the frequency and the initial model describes high-frequency data very poorly.

The initial velocity model, several intermediate models, and the final model are presented in Fig. 7. In this numerical experiment, we set the term $-Ls_n$ in (29) to zero and, thus, the algorithm was penalizing the norm of the update δs , not the updated model $s_n + \delta s$. This was the reason for the high-velocity build-up at $X = 2500$. We stress again that identical FWI results would have been obtained if we had applied the

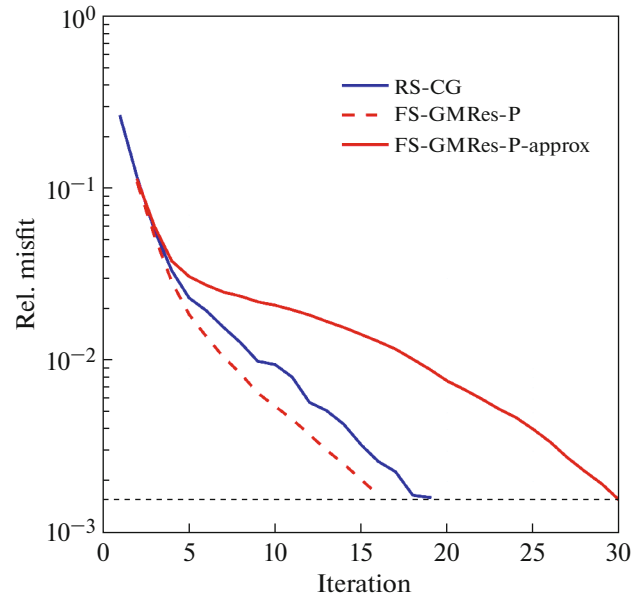


Fig. 5. Evolution of data misfit for the CG applied to the normal system and the preconditioned GMRes applied to the KKT matrix. The Y -axis represents the quantity $\|H\delta s_n - g\|_2 / \|g\|_2$.

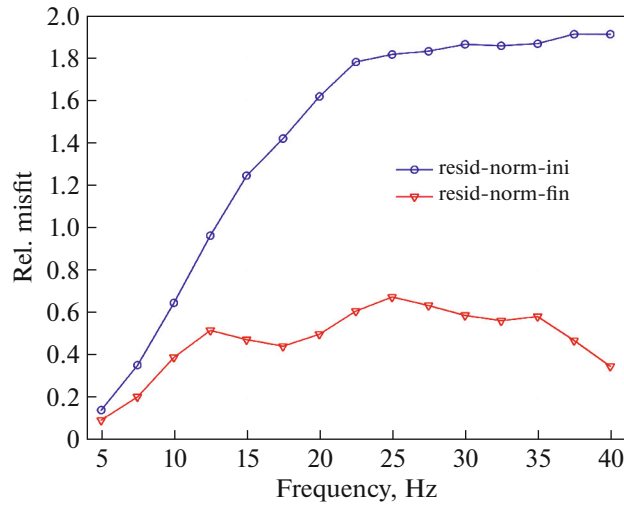


Fig. 6. The relative residual misfit on Newton iterations: “resid-norm-ini” is $\|d_{\text{initial}} - d\|_2 / \|d\|_2$, “resid-norm-fin” is $\|d_{\text{final}} - d\|_2 / \|d\|_2$.

Table 1. Linear solvers comparison. The column “CG” refers to the CG applied to the normal system of linear equations, “GMRes-full” is the GMRes applied to the KKT matrix preconditioned with \mathcal{P} , “GMRes-approx” is the GMRes preconditioned with $\tilde{\mathcal{P}}$

	CG	GMRes-full	GMRes-approx
Iterations	19	16	30
Time per iteration, s	264	267	18
ILU initialization, s	0	0	121
Total time, s	4800	5391	834

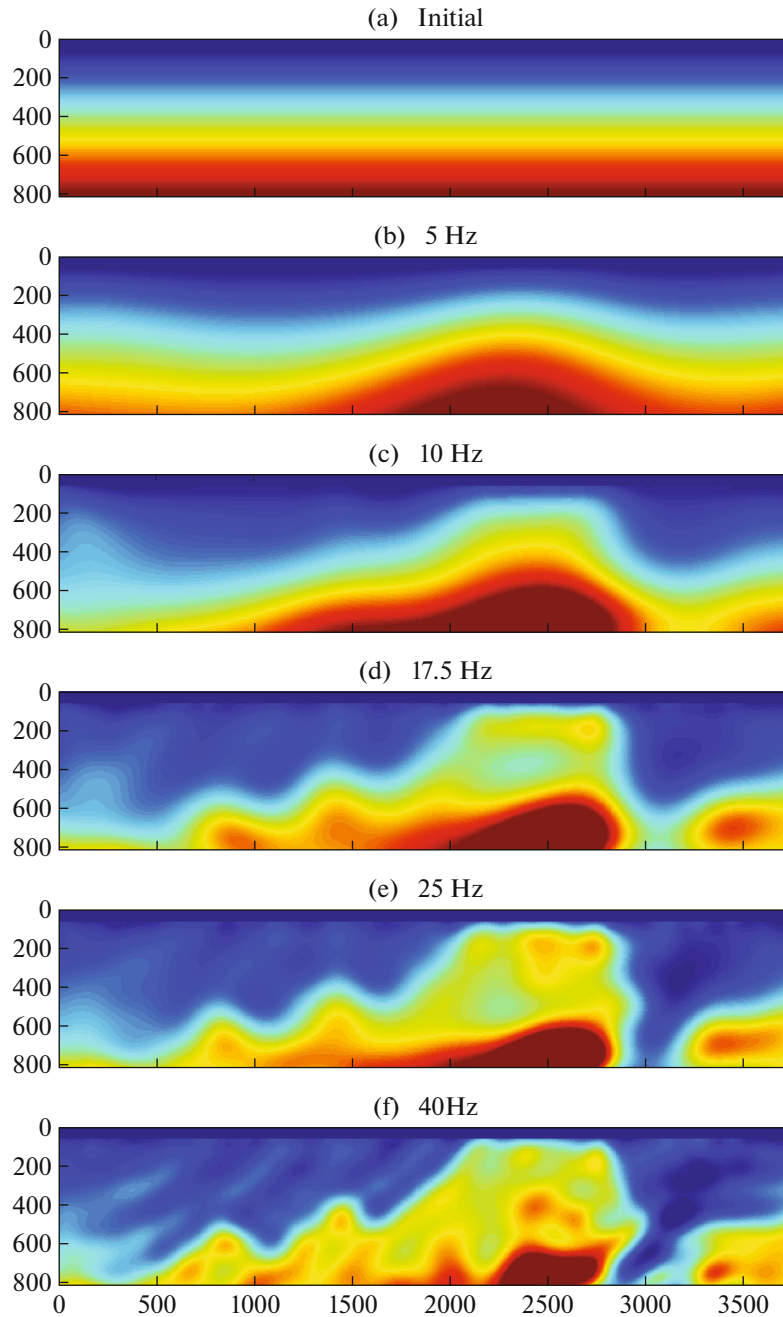


Fig. 7. The initial velocity distribution (a) and updated distributions (b)–(f) at different frequencies. The color scales here and in Fig. 3 are identical. The X - and Z -axis are marked in meters. Similar, if not identical, updates would have been obtained if we had performed the standard RSGN-CG optimization with the same accuracy of data fit.

standard RSGN-CG optimization (provided, of course, that the values of error E_{cg} match in both approaches). The difference between the two approaches is in their performance.

7. CONCLUSIONS

This paper proposes a new numerical method for the seismic FWI, based on the full-space approach. Prior work has demonstrated that the full-space approach can be used as a framework for efficient waveform inversion algorithms. Those works have focused either on the penalty method with the primal-dual

descend optimization or the AL approach with the ADMM optimization procedure. These approaches have demonstrated impressive results, but may suffer from the slow convergence compared to the Newton iterations [25].

This paper studies the iterative full-space FWI by the Newton method. A preconditioned Krylov solver is applied to the KKT system on each Newton iteration. Our method computes the material parameter updates identical to those calculated by inverting the reduced Hessian, provided that the linear systems are solved exactly. The critical element of our approach is a special preconditioner. The preconditioned iterative algorithm requires two forward simulations (exact or approximate) at each iteration plus an inexpensive solve with the constraint matrix L . We considered two variants of the preconditioner: preconditioner \mathcal{P} computes linear solves with A and A^* exactly, whereas preconditioner $\tilde{\mathcal{P}}$ computes the inexact solves. A single application of \mathcal{P} should therefore have almost the same cost compared to a single iteration of the CG applied to the normal system. Our numerical experiments indicated that these two methods indeed have very close performance. On the other hand, a single application of $\tilde{\mathcal{P}}$ is much cheaper. In our numerical experiments, this approach reduced the execution time considerably.

For a real-scale 3D FWI, several problems need to be addressed. First, the forward and adjoint simulations need to be based on the theory of linear elasticity, so it remains to be shown that our approach works for elastic FWI. Second, direct linear solvers become unfeasible in the 3D applications, even for the acoustic equations, not to mention the linear elasticity. Thus, a reliable preconditioned iterative solver should be applied for the forward and adjoint simulations. Substantial progress in this direction has been made; see [29] and references therein. In principle, any preconditioner designed for the forward simulation can be used inside $\tilde{\mathcal{P}}$. Third, the convergence can be improved by a more sophisticated approximation of the reduced Hessian, such as the L-BFGS or the background Hessian. Virtually any method designed to accelerate the reduced-space Newton method can be used to compute the H -block solve within our approach. Finally, using the BiCGStab instead of the GMRes will eliminate the memory needed to store GMRes iterates between restarts. Future work, therefore, should include a study of these options.

FUNDING

This research was partially supported by the Russian Science Foundation, project no. 21-11-00139. The authors acknowledge computational resources granted by Complex for Simulation and Data Processing for Mega-Science Facilities at NRC “Kurchatov Institute,” <http://ckp.nrcki.ru/>.

CONFLICT OF INTEREST

The authors of this work declare that they have no conflicts of interest.

REFERENCES

1. P. Lailly, “The seismic inverse problem as a sequence of before stack migrations,” in *Conference on Inverse Scattering, Theory, and Application* (SIAM, Philadelphia, 1983), pp. 206–220.
2. A. Tarantola, “Inversion of seismic reflection data in the acoustic approximation,” *Geophysics* **49** (8), 1259–1266 (1983).
3. P. R. Mora, “Nonlinear two-dimensional elastic inversion of multi-offset seismic data,” *Geophysics* **52** (9), 1211–1228 (1987).
4. J. Virieux and S. Operto, “An overview of full-waveform inversion in exploration geophysics,” *Geophysics* **74**, WCC1–WCC26 (2009).
5. J. Tromp, “Seismic wavefield imaging of Earth’s interior across scales,” *Nature Rev. Earth Environ.* **1**, 40–53 (2020).
6. R. Plessix, “Three-dimensional frequency-domain full-waveform inversion with an iterative solver,” *Geophysics* **74** (6), WCC149–WCC157 (2009).
7. M. Warner, A. Ratcliffe, T. Nangoo, J. Morgan, A. Umpleby, N. Shah, V. Vinje, I. Štekl, L. Guasch, C. Win, G. Conroy, and A. Bertrand, “Anisotropic 3d full-waveform inversion,” *Geophysics* **78** (2), R59–R80 (2013).
8. S. Operto, A. Miniussi, R. Brossier, L. Combe, L. Métivier, V. Monteiller, A. Ribodetti, and J. Virieux, “Efficient 3-D frequency-domain mono-parameter full-waveform inversion of ocean-bottom cable data: Application to Valhall in the visco-acoustic vertical transverse isotropic approximation,” *Geophys. J. Int.* **202** (2), 1362–1391 (2015).
9. S. Operto and A. Miniussi, “On the role of density and attenuation in three-dimensional multiparameter visco-acoustic VTI frequency-domain FWI: An OBC case study from the North Sea,” *Geophys. J. Int.* **213** (3), 2037–2059 (2018).

10. A. Fichtner, B. L. N. Kennett, H. Igel, and H.-P. Bunge, “Full seismic waveform tomography for upper-mantle structure in the Australasian region using adjoint methods,” *Geophys. J. Int.* **179** (3), 1703–1725 (2009).
11. C. Shin, S. Jang, and D.-J. Min, “Improved amplitude preservation for prestack depth migration by inverse scattering theory,” *Geophys. Prospect.* **49** (5), 592–606 (2001).
12. W. A. Mulder and R.-E. Plessix, “Exploring some issues in acoustic full waveform inversion,” *Geophys. Prospect.* **56** (6), 827–841 (2008).
13. R. Brossier, S. Operto, and J. Virieux, “Seismic imaging of complex onshore structures by 2D elastic frequency-domain full-waveform inversion,” *Geophysics* **74** (6), WCC105–WCC118 (2009).
14. R. G. Pratt, C. Shin, and G. J. Hick, “Gauss–Newton and full Newton methods in frequency–space seismic waveform inversion,” *Geophys. J. Int.* **133** (2), 341–362 (1998).
15. I. Epanomeritakis, V. Akçelik, O. Ghattas, and J. Bielak, “A Newton-CG method for large-scale three-dimensional elastic full-waveform seismic inversion,” *Inverse Probl.* **24** (3), 034015 (2008).
16. L. Métivier, R. Brossier, S. Operto, and J. Virieux, “Full waveform inversion and the truncated newton method,” *SIAM Rev.* **59** (1), 153–195 (2017).
17. E. Haber, U. M. Ascher, and D. Oldenburg, “On optimization techniques for solving nonlinear inverse problems,” *Inverse Probl.* **16** (4), 1263 (2000).
18. E. Haber and U. M. Ascher, “Preconditioned all-at-once methods for large, sparse parameter estimation problems,” *Inverse Probl.* **17**, 1847–1864 (2001).
19. E. Haber, U. M. Ascher, and D. W. Oldenburg, “Inversion of 3D electromagnetic data in frequency and time domain using an inexact all-at-once approach,” *Geophysics* **69** (5), 1216–1228 (2004).
20. A. Abubakar, W. Hu, T. M. Habashy, and P. M. van den Berg, “Application of the finite-difference contrast-source inversion algorithm to seismic full-waveform data,” *Geophysics* **74** (6), WCC47–WCC58 (2009).
21. T. van Leeuwen and F. J. Herrmann, “Mitigating local minima in full-waveform inversion by expanding the search space,” *Geophys. J. Int.* **195** (1), 661–667 (2013).
22. T. van Leeuwen and F. J. Herrmann, “A penalty method for PDE-constrained optimization in inverse problems,” *Inverse Probl.* **32**, 015007 (2016).
23. J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. (Springer, New York, NY, 2006).
24. G. Rizzuti, M. Louboutin, R. Wang, and F. J. Herrmann, “A dual formulation of wavefield reconstruction inversion for large-scale seismic inversion,” *Geophysics* **86** (6), R879–R893 (2021).
25. S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Found. Trends Mach. Learn.* **3** (1), 1–122 (2011).
26. H. S. Aghamiry, A. Gholami, and S. Operto, “Improving full-waveform inversion by wavefield reconstruction with the alternating direction method of multipliers,” *Geophysics* **84** (1), R125–R148 (2019).
27. K. Aghazade, A. Gholami, H. S. Aghamiry, and S. Operto, “Anderson-accelerated augmented Lagrangian for extended waveform inversion,” *Geophysics* **87** (1), R79–R91 (2022).
28. G. Biros and O. Ghattas, “Parallel Lagrange–Newton–Krylov–Schur methods for PDE-constrained optimization: I. The Krylov–Schur solver,” *SIAM J. Sci. Comput.* **27** (2), 687–713 (2005).
29. N. Yavich, N. Khokhlov, M. Malovichko, and M. S. Zhdanov, “Contraction operator transformation for the complex heterogeneous Helmholtz equation,” *Comput. Math. Appl.* **86**, 63–72 (2021).
30. B. Hustedt, S. Operto, and J. Virieux, “Mixed-grid and staggered-grid finite-difference methods for frequency-domain acoustic wave modelling,” *Geophys. J. Int.* **157** (3), 1269–1296 (2004).
31. M. S. Malovichko, A. V. Tarasov, N. B. Yavich, and K. V. Titov, “Application of optimal control to inversion of self-potential data: Theory and synthetic examples,” *IEEE Trans. Geosci. Remote Sens.* **60**, 1–11 (2021).
32. M. Hinze, R. Pinnau, M. Ulbrich, and S. Ulbrich, *Optimization with PDE Constraints* (Springer, Dordrecht, 2009).
33. V. Kostin, S. Solov'yev, A. Bakulin, and M. Dmitriev, “Direct frequency-domain 3D acoustic solver with intermediate data compression benchmarked against time-domain modeling for full-waveform inversion applications,” *Geophysics* **84** (4), T207–T219 (2019).
34. Y. A. Erlangga, C. Vuik, and C. W. Oosterlee, “On a class of preconditioners for the Helmholtz equation,” *Appl. Numer. Math.* **50**, 409–425 (2004).

Publisher’s Note. Pleiades Publishing remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.