

An efficient coupled acoustic-elastic local solver applied to phase inversion

Bram Willemsen¹ and Alison Malcolm²

ABSTRACT

In characterizing reservoirs, we are often interested in retrieving detailed elastic parameters for only a very limited part of the subsurface. One way of doing so involves studying the seismic reflection response within the region of interest. To that end, we introduce a local solver that uses an acoustic solver to propagate the wavefield to a subdomain on which we use a local elastic solver. This avoids the use of an expensive full-domain elastic solver while still incorporating elastic physics in the region where it is most important. We then study whether this modeled phase is sufficiently accurate for recovering important subsurface reservoir properties in an inversion procedure.

INTRODUCTION

When seismic waves reflect from interfaces in the earth, their reflected amplitude depends on the incidence angle of the incoming wave. A physical model for describing the reflection coefficient in terms of the incidence angle is the [Zoeppritz \(1919\)](#) equations. The Zoeppritz equations and their various approximations ([Aki and Richards, 2002](#)) assume an incident plane wave. As a result, the derived plane-wave reflection coefficients (PRCs) are only accurate for plane waves. In most physical experiments, however, waves radiate outward from point sources in (approximately) spherical shells. [O'Brien \(1963\)](#) notes that plane-wave theory poorly describes the reflection coefficient of spherical waves close to the critical incidence angle. [Červený \(1959\)](#) shows that the spherical-wave reflection coefficients (SRCs) are related to PRCs as a weighted integral of PRCs over a range of incidence angles. The specifics of how the reflection amplitude varies with incidence angle is the subject of study in amplitude variation with angle (AVA) and

amplitude variation with offset (AVO). The P-wave velocity and density, and S-wave velocity of the material on both sides of the interface together determine the particular shape of the reflection coefficient curve when plotted against angle (AVA) or offset (AVO). Because oil and gas content in the pores of rocks influences these seismic material properties, studying AVA/AVO is an important tool for seismic interpreters.

Although it is less well-studied, seismic waves also experience phase variation with angle (PVA) when they are postcritical (i.e., arrive at angles larger than the critical angle). The PRCs of these postcritical reflections are complex (see, for instance, [Downton and Ursenbach \[2006\]](#) for a recent discussion) and therefore introduce a phase shift to the reflected wave. In contrast, PRCs for incidence angles smaller than the critical angle have zero phase shift. We use the term phase shift exclusively for shifts that change the shape of the wavelet. Time shifts are technically also phase shifts, but these are not the focus of PVA. Similar to AVA, the characteristics of this PVA curve depend on the contrast in elastic properties along the interface. This makes PVA a useful quantity for inversion either by itself or in addition to AVA. These angle-dependent effects are not limited to PP or SS reflections, but also occur in converted waves such as PS and SP reflections. For instance, [Gong and McMechan \(2016\)](#) use AVA and PVA for the combination of PP and PS reflections to estimate material properties at an interface. We focus on PP reflections only because they are present in land and marine surveys and typically have the best signal-to-noise ratio. Recall that the SRC for a specific incidence angle contains PRCs of all incidence angles through a weighted integration. This means that the SRC for a precritical incidence angle already includes a phase shift ([Červený and Hron, 1961](#)), unlike the PRC for the same precritical incidence angle. If the P-wave velocity contrast is insufficient to generate PVA in the PP reflection because there is no critical angle, then the SP reflection may be used as an alternative because a critical angle will almost always be present. Furthermore, the phase appears to remain relatively stable when propagating through a heterogeneous overburden, unlike amplitude, which is

Manuscript received by the Editor 8 November 2016; revised manuscript received 31 January 2017; published online 21 April 2017.

¹Massachusetts Institute of Technology, Earth Resources Lab, Cambridge, Massachusetts, USA. E-mail: lawillem@mit.edu.

²Memorial University of Newfoundland, Department of Earth Sciences, St. John's, Newfoundland, Canada. E-mail: amalcolm@mun.ca.

© 2017 Society of Exploration Geophysicists. All rights reserved.

affected by energy partitioning at each reflector (Zhu and McMechan, 2012).

AVA for precritical reflections has traditionally received a lot of attention in the literature; see Buland and Omre (2003) and Smith and Gidlow (1987) for example. The analysis by O'Brien (1963) suggests that this modeling, which is based on plane waves, will break for more realistic spherical waves in the neighborhood of the critical angle. There are only a few studies making use of the AVA/PVA response of large-angle reflections for determining material properties at a reflector of interest. One of the issues for such studies is the selection of an appropriate forward model. Arguably, the most accurate forward model is a wavefield simulation on an earth model containing the reflectors of interest. The amplitude and phase of the modeled reflections can then be compared directly with the observed reflections in the data. At the very least, such a simulation requires P-wave velocity, density, and S-wave velocity, to include the parameters to which reflection coefficients are sensitive. The simplest type of wavefield simulation for reflection modeling is, therefore, elastic, but these simulations are computationally expensive. Even though acoustic simulations are significantly cheaper, they cannot substitute for elastic solvers because phase shifts of PP reflections are influenced significantly by the S-wave velocity contrast, as we demonstrate in this paper.

Some studies have used alternative forward models for the inversion of large-angle reflections. For example, Zhu et al. (2012b) and Zhu and McMechan (2014) use the τ - p transform to decompose the recorded data into plane waves. A ray tracer is then used to index these plane waves for multiple reflection angles at reflector locations of interest. From these indexed plane waves, the AVA and PVA of the corresponding PRC are determined. These PRCs at the reflection points are then compared with PRCs from Zoeppritz modeling to invert for the material contrast. Because it used a ray-based method, this approach does not model finite-frequency effects of the wavefield as it propagates through the overburden. The method also requires the explicit picking of reflectors and their dip angles from migrated images. Instead of using a ray tracer to measure PVA inversion, Zhu et al. (2012a) use a finite-difference (FD) wavefield simulator to generate angle-domain common image gathers (ADCIGs) (Sava et al., 2001; Jin et al., 2014). Zhu et al. (2012a) show that the phase of the imaged reflector varies as a function of the reflection incidence angle in these ADCIGs. In theory, one would expect that these observations could subsequently be used to invert for material parameters, although to our knowledge this has not been done. As with any imaging method, such an approach is sensitive to defocusing as a result of inaccuracies in the velocity model used for migration. The imperfections in constructing ADCIGs and defocusing due to inaccurate velocity models are expected to lead to artifacts when used in an inversion, which provides the impetus to the development of full-waveform methods that directly fit the phase of the data themselves, rather than first computing the ADCIGs.

Both this proposed approach of directly fitting the phase of the waveform and the ADCIG approach of Zhu et al. (2012a) require an FD wavefield simulator. It is this FD simulator that results in the largest expense of these methods. There is a lot of fundamental research to improve their efficiency; see for instance the work of Wang et al. (2011) and Zepeda-Núñez and Demanet (2016). For the purpose of modeling reflections, there are additional ways to reduce the cost of the wavefield simulation. Hobro et al. (2014)

and Chapman et al. (2014) show that elastic PP reflections can be approximated using an acoustic FD solver, but the inability of this method to accurately model postcritical reflections and their associated phase shifts makes it less applicable to PVA inversion. An alternative way of bringing down the cost of the wavefield simulation of the reflection is by using a local solver centered around the reflector of interest.

Local solvers reduce computation time by limiting the size of the computational domain to a smaller region of interest. This usually requires a precomputation step in which several full-domain simulations are performed. After these precomputations, the local wavefields can be generated at much faster speed than full-domain simulations. These localized wavefields can then be used for inversions for a salt boundary (Willemssen et al., 2016), time lapse (Robertsson et al., 2000; Borisov et al., 2015; Malcolm and Willemssen, 2016), and reflector properties from PVA as we do in this paper.

Broadly speaking, local solvers can be subdivided into exact local solvers (van Manen et al., 2007; Willemssen et al., 2016) and inexact local solvers. By relaxing accuracy constraints, the inexact local solvers attain higher computational efficiency. The inexact (potentially anisotropic) elastic local solver formulation of Robertsson and Chapman (2000) allows the model to be perturbed within the local domain. Computational gains are achieved by implicitly discarding waves scattering from those perturbations that exit the local domain, then scatter in the inhomogeneous background model and subsequently reenter the local domain. Modeling primary PP reflections used in the AVA/PVA analyses will not result in significant accuracy loss, making it possible to change the elastic material properties around the reflector and still accurately model the reflection response we are interested in for inversion.

The elastic local solver of Robertsson and Chapman (2000) requires many elastic full-domain Green's function computations to connect the local domain and the acquisition. Recalling the observation of Zhu and McMechan (2012) that the phase of PP reflections is approximately unchanged by propagation through the overburden, we try to reduce the computational cost of this step by using acoustic Green's functions. Because elastic simulations are significantly more expensive than acoustic simulations (see, e.g., Hobro et al., 2014), this approach has the potential to reduce the precomputation cost of the local solver while still accurately modeling the phase of the reflected wave. The acoustic approximation is not always good, as Cance and Capdeville (2015) demonstrate. However, their figures 4b and 5b demonstrate that even with highly heterogeneous elastic models, acoustic propagation gives wavelets of similar shape although they have minor arrival time differences. We therefore expect that, for the purpose of propagating the local elastic reflections through the overburden, acoustic propagation is sufficient for preserving the shape of the wavelet we invert for. The local solver we introduce in this paper therefore mixes physics by using an acoustic solver to propagate the wavefield through the overburden to a local elastic solver around the reflector, in which the elastic behavior determines the phase shift of interest. Although we focus on PVA inversion in this paper, potentially combining this with AVA better constrains the inverse problem for the elastic material contrast as shown by Zhu and McMechan (2014).

There are a few other seismic studies in the literature where different solvers are combined. Monteiller et al. (2013) use an efficient spherically symmetric wave solver to propagate teleseismic body waves through the earth. Only in the relatively small region of in-

terest, they are coupled with a 3D heterogeneous solver. Capdeville et al. (2003) use a similar approach, where the earth is again separated into two parts. This approach again limits the expensive solver to a small region, using a simplified solver for the remainder. Finally, Yu et al. (2016) derive a coupled acoustic-elastic set of equations to aid elastic reverse time migration using ocean-bottom geophones.

This paper is organized as follows. In the “Methods” section, we explain how computationally efficient acoustic Green’s functions are used to couple the elastic local solver with the sources and receivers. In the “Results” section, we start with a relatively simple two-layer model to verify the local solver and explain its functionality. We then apply the local solver to a more realistic synthetic North Sea model. After discussing the shot gathers produced by the local solver, we apply it to PVA inversion.

METHODS

In this paper, we efficiently model primary PP reflections by surrounding the reflector(s) of interest with a local elastic solver within an acoustic background. We use the local elastic solver of Robertsson and Chapman (2000) and replace the full-domain elastic (EL) Green’s functions with computationally efficient constant density acoustic (CDA) Green’s functions to couple the local solver to the acquisition. This significantly reduces the cost of precomputing the required Green’s functions for the local solver. This replacement is motivated by the observation of Zhu and McMechan (2012) that the phase of a PP-reflection is primarily affected by scattering at the reflector and less so by propagation through the (heterogeneous) elastic overburden. We expect that this observation remains valid when restricting ourselves to heterogeneous overburdens with constant density and zero S-wave velocity (i.e., a CDA simulation) because it is a simpler case.

Before introducing the local solver, we first introduce the following definitions. We define the background wavefield as the wavefield in background model $m^0(\mathbf{x})$. The perturbed wavefield is the wavefield in the perturbed model $m(\mathbf{x})$. Finally, the scattered wavefield is the difference between the perturbed and the background wavefield, caused by model update $\Delta m(\mathbf{x})$ limited to the interior of the local solver domain. For example, the z -component of the velocity would be denoted $v_z^0(\mathbf{x}, t)$, $v_z(\mathbf{x}, t)$, and $v_z^{\text{sc}}(\mathbf{x}, t)$. The local solver computes the various components of the elastic wavefield on a staggered grid, as is seen in Figure 1. This figure is a useful guide in Appendix A in which we explain to the interested reader the implementation of the local solver as presented by Robertsson and Chapman (2000).

The local solver generates an approximation of the scattered field due to an elastic model perturbation within the local domain. The only component of the wavefield (implicitly) discarded by the local solver is the wavefield scattering on the perturbation, leaving the local domain, and then reentering before propagating to the receivers (i.e., long-range interactions). Such a discarded event is displayed in gray in Figure 2. This limitation is inconsequential for the purpose of modeling primary reflections.

CDA coupling

In this section, we discuss how we replace the full-domain elastic Green’s functions with computationally efficient CDA Green’s functions. The scalar CDA solver from the numerical code PySIT (Hewett and Demanet, 2013) is used to compute these Green’s func-

tions. Concretely, the CDA background Green’s functions are obtained by defining the background model $m_0(\mathbf{x})$ as exactly the background P-wave velocity on the full computational domain

$$m^0(\mathbf{x}) = \{V_{p,0}(\mathbf{x})\}. \quad (1)$$

We then solve the wave equation,

$$\frac{1}{V_{p,0}(\mathbf{x})^2} \frac{\partial^2 p(\mathbf{x}, t)}{\partial t^2} - \nabla^2 p(\mathbf{x}, t) = f(\mathbf{x}, t), \quad (2)$$

with the source defined by Dirac delta functions for Green’s function computations:

$$f(\mathbf{x}, t) = \delta(\mathbf{x} - \mathbf{x}_0) \delta(t - t_0). \quad (3)$$

The elastic local solver discussed in Appendix A is second-order accurate in time and fourth-order accurate in space. We match these orders of accuracy in PySIT when solving for the CDA background

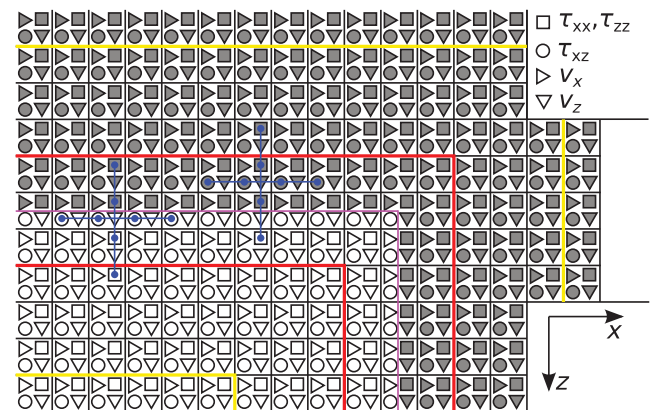


Figure 1. Staggered elastic grid around the top-right boundary of the elastic local solver. The various colored lines are used at different points in the paper, where their purposes are also described.




Figure 2. Schematic drawing of a local solver (red box) within a full-domain solver (black box). The local solver used in this study is a modification of the local solver introduced by Robertsson and Chapman (2000). It implicitly discards waves scattering on local model updates $\Delta m(x)$, which then exit and subsequently reenter the local domain, like the gray event. Multiple scattering within the local domain, like the blue event, is included naturally.

Green's functions on a regular FD grid. To expand our method to higher orders of accuracy, the stencil size will change, but the required modifications are straightforward.

The CDA propagation on the background model $m_0(\mathbf{x})$ physically corresponds to propagation through a constant density liquid (i.e., no S-wave velocity). But inside the local domain, away from the boundaries, the perturbed model $m(\mathbf{x})$ is elastic:

$$m(\mathbf{x}) = \{V_P(\mathbf{x}), \rho(\mathbf{x}), V_S(\mathbf{x})\}. \quad (4)$$

Figure 3 schematically shows the combination of the CDA full-domain solver with the EL local solver. The pink layers represent the perfectly matched layer (PML) boundary conditions, although free-surface conditions are also possible on the full domain. The PML around the elastic local solver is much thinner than the one surrounding the full computational domain. The red boundary corresponds with the region between the red lines in Figure 1 where the wavefield is “injected.” Within the blue boundary of Figure 3 is the elastic region, in which the model update, $\Delta m(\mathbf{x})$, contains the reflector whose PP reflection we want to model. Between the blue and red boundaries, the material gradually transitions from a fully elastic medium to the CDA state of the exterior over an exponential taper with size proportional to the dominant wavelength. The result of this taper is that the material properties at the injection boundary of the elastic local solver remain CDA, ensuring physical consistency with the injected CDA wavefield. The star and triangle symbols represent a source and receiver, respectively. The lines connecting these to the red boundary represent background CDA Green's functions. With the geometry defined, we are now ready to discuss how the CDA field is injected on the staggered EL grid of the local solver.

Injection approach using CDA Green's functions

The injection approach we introduce in this section is flexible. It will correctly compute the injection field for the tapered geometry above, which preserves CDA material properties at the red injection boundary of Figure 3. The taper increases the size of the local solver, reducing the speedup factor compared with a full-domain simulation. For this reason, we also investigate an alternative approach in which we remove the taper and directly inject the CDA wavefield on




Figure 3. Schematic of the local solver. The gray lines represent the required background CDA Green's functions among the sources, receivers, and boundaries of the local solver. The model update within the local solver is elastic, with the region between the red and the blue lines representing ρ and V_P taper. The blue box contains exactly the part of the elastic model that we want to model.

a local elastic model. This injection approach is approximate in nature because of the inconsistency in material parameters at the injection boundary. The background pressure wavefield is computed using CDA material parameters, whereas the model in the local solver is in this case no longer CDA; it has the same P-wave velocity model but also spatially varying $\rho(\mathbf{x})$ and $V_S(\mathbf{x})$ at the injection boundary. The following injection approach can accommodate both approaches using the same precomputed CDA Green's functions. Specifically, the steps of the approach are:

- In a precomputation step, store the CDA pressure Green's functions on the nine layers between the two yellow lines in Figure 1. This precomputation step is the only time we use a full-domain CDA solver.
- First we compute both velocity components by evaluating equations A-1 and A-2, using the density $\rho(\mathbf{x})$ of the elastic model. For the stresses τ_{xx} and τ_{zz} , we use the pressure obtained by convolving the precomputed Green's functions with the source wavelet. This CDA pressure field has zero shear stress τ_{xz} . Because equations A-1 and A-2 require the evaluation of derivatives, we can only compute them at interior layers and the velocity components will therefore be available at fewer than the nine layers where the CDA pressure is available.
- Using the velocity components computed in the previous step and the elastic material properties around the injection boundaries, we evaluate equations A-3–A-5. This will naturally result in $\tau_{xx} \neq \tau_{zz}$ and $\tau_{xz} \neq 0$ when $V_S \neq 0$. Just like for the velocity components, equations A-3–A-5 require evaluation of the derivative stencils. The shear stresses can be computed exactly at the three layers between the red boundaries in Figure 1.

With this scheme, we turn the nine layers of precomputed CDA background pressures into the velocity and stress tensor at all the required injection locations for the elastic local solver. Even though the CDA Green's function is stored along nine layers, it requires less storage than the full five-parameter elastic injection wavefields on the three layers between the red lines. In addition, the stored elastic wavefield may not be consistent if the local model changes along the injection boundary (e.g., trying a different S-wave velocity taper). Instead, the procedure we introduce is flexible in that the elastic injection fields can easily be recomputed if the elastic properties along the boundary of the local solver are modified. No CDA Green's functions need to be recomputed when injecting on a model with a different choice of $\rho(\mathbf{x})$ and $V_S(\mathbf{x})$.

Even a smooth variation in S-wave velocity and density influences the amplitude of a P-wave. When injecting a CDA wavefield on a material with varying S-wave velocity and density, we therefore want to approximately correct its amplitude by using compensation factors $l_\rho^{\text{inj}}(\mathbf{x})$ and $l_{V_S}^{\text{inj}}(\mathbf{x})$ for density and S-wave velocity, respectively. The compensation factors are derived by Willemssen (2017) and summarized in Appendix B. They only account for the instantaneous effect of the density and S-wave velocity at the injection point. They do not compensate for transmission losses that are observed if an elastic solver is used for the background Green's function computations on a fully elastic model. The S-wave velocity scaling, $l_{V_S}^{\text{inj}}(\mathbf{x})$, is automatically applied to the stress fields during the injection procedure listed above, as is explained by Willemssen

(2017). We then scale the entire elastic injection field with the density term $l_{\rho}^{\text{inj}}(\mathbf{x})$. Note that the procedure described above can be applied with or without tapering the density and S-velocity between the blue and red lines in Figure 3.

Exact boundary integral

The local solver computes a wavefield within the local domain, corresponding to the perturbed model $m(\mathbf{x})$. In general, for an inversion, the wavefield is also required at the receiver locations where it is compared with the observed data. We use Green's functions between the boundary of the local solver and the receivers, shown with gray lines in Figure 3, to propagate the scattered wavefield to the receiver locations. This is done by evaluating a boundary integral. We first address the scenario of Figure 3 where the model update, $\Delta m(\mathbf{x})$, at the boundary is zero due to tapering, preserving the CDA properties there. In this scenario, the boundary integral along the red box in Figure 3 assumes the standard CDA form (Fokkema and van den Berg, 2013):

$$p^{\text{sc}}(\mathbf{y}, t) = \int_S \left[G_0(\mathbf{x}, \mathbf{y}, t) * \frac{\partial p^{\text{sc}}(\mathbf{x}, t)}{\partial \mathbf{x}_n} - p^{\text{sc}}(\mathbf{x}, t) * \frac{G_0(\mathbf{x}, \mathbf{y}, t)}{\partial \mathbf{x}_n} \right] dS_x, \quad (5)$$

where $*$ is convolution; $G_0(\mathbf{x}, \mathbf{y}, t)$ is the background CDA Green's function between boundary position \mathbf{x} and evaluation position \mathbf{y} by substituting the source in equation 3 into equation 2 with $\mathbf{x}_0 = \mathbf{y}$. We denote receiver locations with index \mathbf{y} throughout. The boundary integral in equation 5 accurately propagates the scattered field as long as the elastic model update, $\Delta m(\mathbf{x})$, is restricted to the interior of the integration boundary. We evaluate the normal derivatives in equation 5 using a fourth-order discretization to match the discretization order of the FD stencils. The integral boundary goes through the shaded area in Figure 1 and all around the local domain where the local solver computes the scattered wavefield to be evaluated in equation 5.

Approximate boundary integral for variable ρ/V_s along the boundary of the local domain

As mentioned when introducing the injection routine, we are interested in relaxing the constraint that the perturbed model should taper to CDA propagation along the injection boundary. In this case, we still want to evaluate the pressure boundary integral of equation 5, even though the requirement of restricting $\Delta m(\mathbf{x})$ to the interior of the integral is no longer honored and the wavefield is now elastic. As we did in the previous section, we do this by applying density and S-wave velocity scaling factors to the CDA Green's functions used to propagate the wavefield to the receiver locations in equation 5 to get

$$p^{\text{sc}}(\mathbf{y}, t) = \int_S \left[l_{\rho}^{\text{rec}}(\mathbf{x}) l_{V_s}^{\text{rec}}(\mathbf{x}) \left[G_0(\mathbf{x}, \mathbf{y}, t) * \frac{\partial p^{\text{sc}}(\mathbf{x}, t)}{\partial \mathbf{x}_n} - p^{\text{sc}}(\mathbf{x}, t) * \frac{G_0(\mathbf{x}, \mathbf{y}, t)}{\partial \mathbf{x}_n} \right] \right] dS_x, \quad (6)$$

where $l_{\rho}^{\text{rec}}(\mathbf{x})$ and $l_{V_s}^{\text{rec}}(\mathbf{x})$ are approximate inverses of the injection compensation factors. The pressure at location \mathbf{x} is affected by the

density and S-wave velocity at location \mathbf{x} . The factors $l_{\rho}^{\text{rec}}(\mathbf{x})$ and $l_{V_s}^{\text{rec}}(\mathbf{x})$ approximately restore the amplitude of the elastic pressure wavefield back to the amplitude under CDA conditions.

The scattered pressure wavefield computed at the receiver locations through equation 6 is then added to the background wavefield to obtain the perturbed wavefield. The background wavefield is obtained by convolving the background Green's function between source and receiver grid nodes with the source wavelet. We must also correct the amplitudes of the scattered and perturbed wavefields for differences in density and S-wave velocity between the source and receiver grid nodes. For the scattered wavefield, these scalar compensation factors are folded into compensation factors $l_{\rho}^{\text{rec}}(\mathbf{x})$ and $l_{V_s}^{\text{rec}}(\mathbf{x})$. As a result, the boundary integral automatically takes care of this for the scattered field. For the background wavefield, we multiply Green's functions between source and receiver with the scalar $l^0(\mathbf{x}_s, \mathbf{x}_r)$. Finally, both wavefields are also compensated for S-wave velocity at the source grid node by $l_{V_s}^{\text{inj}}(\mathbf{x}_s)$. All these amplitude compensation factors, together with a schematic overview, are found in Appendix B.

Phase measuring

To measure the phase of a wavelet, we follow Fomel et al. (2010) who use the observation that a zero-phase wavelet has maximum correlation with its envelope. To measure the phase in this way, we use a grid-search procedure, with a grid over phase. We shift the wavelet in 1° increments, and we compute the correlation with the envelope of this shifted wavelet. The phase angle that gives the largest correlation is then the phase-shift of the input wavelet. This algorithm is summarized by Zhu et al. (2011).

Phase inversion

For PVA inversion, we minimize the following least-squares objective equation:

$$\chi(\mathbf{m}) = \frac{1}{2} \langle \mathbf{d} - \mathbf{u}(\mathbf{m}), \mathbf{d} - \mathbf{u}(\mathbf{m}) \rangle = \frac{1}{2} \langle \mathbf{r}(\mathbf{m}), \mathbf{r}(\mathbf{m}) \rangle \quad (7)$$

between observed phase vector \mathbf{d} and simulated phase vector \mathbf{u} . The simulated phase \mathbf{u} is a function of the elastic material parameters at the layers of interest, grouped together into the model parameters \mathbf{m} . The phase residual $\mathbf{r}(\mathbf{m})$ is corrected for phase wrap to get absolute phase differences. In each iteration i of the inversion, we form the normal equations

$$\mathbf{J}_i^T \mathbf{J}_i \mathbf{p}_i = -\mathbf{J}_i^T \mathbf{r}_i \quad (8)$$

for current model estimate \mathbf{m}_i to minimize equation 7. Here, \mathbf{p}_i is the search direction we want to obtain. The $n \times n$ Jacobian matrix \mathbf{J}_i is evaluated discretely using n forward models. In each forward model, we perturb one material parameter by a small amount (e.g., we find 25.0 works well both for velocities in m/s and density in kg/m³) and then obtain one of the n columns of \mathbf{J}_i . In the inversion example in this paper, we invert for two layers with three elastic material parameters each, so n is six. The search direction is then obtained by inverting the small $n \times n$ matrix $\mathbf{J}_i^T \mathbf{J}_i$:

$$\mathbf{p}_i = -(\mathbf{J}_i^T \mathbf{J}_i)^{-1} \mathbf{J}_i^T \mathbf{r}_i. \quad (9)$$

With this search direction, a line search is initiated to update model \mathbf{m}_i to \mathbf{m}_{i+1} :

$$\mathbf{m}_{i+1} = \mathbf{m}_i + \alpha_i \mathbf{p}_i. \quad (10)$$

Because a Newton method is used to compute the search direction \mathbf{p}_i in equation 9, we first attempt step length $\alpha_i = 1.0$. If this does not decrease the objective function equation 7, then the step length is reduced and equation 7 is evaluated again. This continues until an improved model \mathbf{m}_{i+1} is found.

RESULTS

Only when we taper to constant density and zero S-wave velocity along the boundary of the local solver do we satisfy the theoretical requirements that make the boundary integral and injection exact. We first investigate the local solver in this condition and verify its exactness on a simple two-layer model. We then apply the same configuration to a more realistic earth model with varying elastic material properties around the boundary of the local solver. We also empirically investigate the effect of partially removing the taper and injecting on a variable density and S-wave velocity medium. Finally, we show synthetic phase inversion results using the local solver.

Two-layer example

We now investigate the local solver on a simple elastic two-layer model. In the local solver, we taper the elastic true model to constant density and S-wave velocity at its boundary. We then verify that the phase of the reflection in this tapered simulation is similar to that of the true elastic model.

The P-wave velocity of the true elastic model is shown in Figure 4, with the elastic parameters given in the caption and source location at $x = 1.0$ km and $z = 0.1$ km. The receivers are placed along the entire horizontal extent of the model also at $z = 0.1$ km. The red and blue boxes in the P-wave velocity model correspond to those in Figure 3. A Ricker wavelet with peak frequency of 25.0 Hz is used as the source. We verify that with a grid spacing of 3.25 m, there is no appreciable numerical dispersion in our setup by confirming that the phase of the direct wave is independent of offset. With negligible dispersion, we can focus entirely on phase changes due to reflection. The entire model is surrounded by PML boundary conditions.

The two-layer elastic model with P-wave velocity shown in Figure 4 is considered the “true” model in this experiment. We now investigate whether we can correctly model the elastic PP reflection at the layer boundary while tapering the density and S-wave veloc-




Figure 4. P-wave velocity of the two-layer true elastic model, with the source at $x = 1.0$ km and $z = 0.1$ km. Material properties in the top layer are $V_p = 2000.0$ m/s, $\rho = 2000.0$ kg/m³, and $V_s = 880.0$ m/s. In the bottom layer, these material properties are with the parameters 4000.0 m/s, 2300.0 kg/m³, and 1540.0 m/s, respectively. The density and S-wave velocity models are tapered from fully elastic inside the blue box to CDA at the red box. Specifically, we taper to zero S-wave velocity (CDA) and to the density of the top layer. For this simple model, all scaling factors summarized in Figure B-1 are constants equal to 1.0, with the exception of $l_{V_s}^{\text{inj}}(\mathbf{x}_s)$, which is $2000^2 - 880^2 / 2000^2 = 0.8064$.

ity to CDA properties. Within the blue box, the full elastic model (i.e., the true V_p , ρ , and V_s) is used in the local solver. We implement an exponential taper in the region between the blue and red boxes to transition from the EL model in the blue box to the CDA model in the red box and the exterior. The size of the taper is the same on all sides and is chosen conservatively as 240 m, which is three times the peak wavelength in the top layer. For this simple model, all amplitude compensation factors are independent of position.

Figure 5 compares the results for the full-domain elastic solver in the true model (Figure 5a) with the P-wave velocity shown in Figure 4 with the full-domain elastic solver on the tapered model (Figure 5b) and the coupled CDA-EL local solver (Figure 5c). The full-domain tapered shot gather and the local shot gather are multiplied by the scalar $l_{V_s}^{\text{src}}(\mathbf{x}_s)$ to compensate for the S-wave velocity at the source grid node.

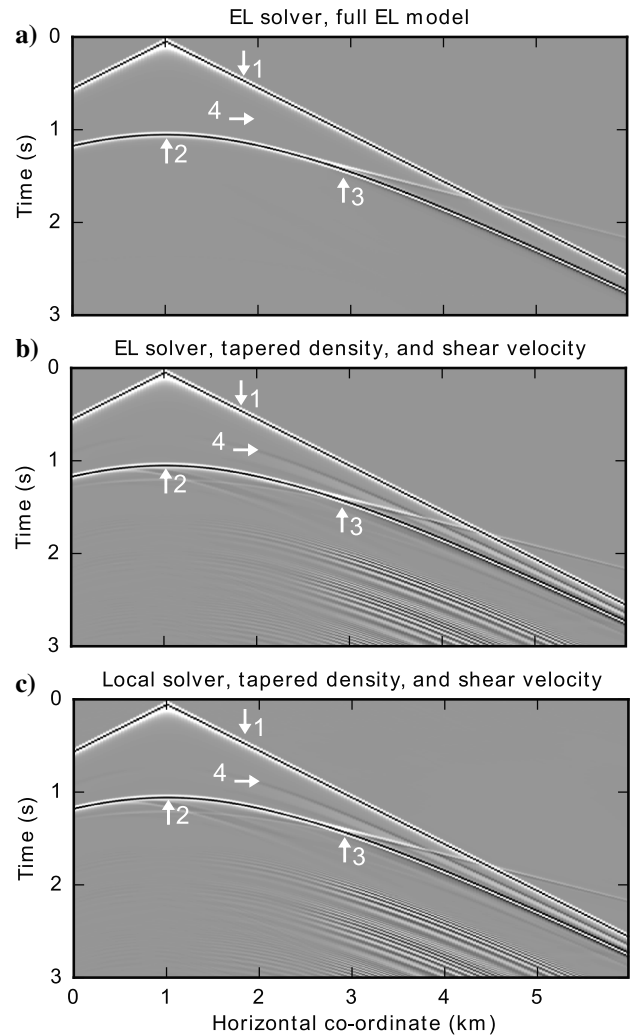


Figure 5. Gathers: The number 1 marks the direct wave, the number 2 marks the reflection, and the number 3 marks the location in which the head wave peels off and where we clearly enter the post-critical regime (to the right, larger offset). The number 4 marks the reflection of the incident pressure wavefield on the top taper boundary. This event is not present in the true simulation because there is no taper. The same amplitude clipping is applied to all three simulations so that the color scales correspond.

In these figures, the event marked with number 1 is the direct wave. Event 2 is the reflection from the layer boundary 1.0 km below the source position. This is the event we want to model accurately and efficiently using the CDA-EL coupled local solver. Event 3 at the horizontal coordinate $x = 3.0$ km shows the head wave separating from the reflection. This head wave is generated at the critical incidence angle, which is 30° in this model. At larger horizontal coordinates (i.e., larger offsets), the reflection is postcritical and a noticeable phase shift is observed. Using the full-domain elastic solver but replacing the true density and S-wave velocity models with their tapered equivalents introduces several extra events in Figure 5b. The density is tapered to the density of the top layer. As a result, the top taper only contains an S-wave velocity contrast. Below the source (i.e., $x = 1.0$ km), the reflection of the incident P-wave from the S-wave velocity taper is negligible. The reflection from the top taper only becomes significant at grazing incidence angles as indicated by event 4.

Just below the desired reflection, two other minor reflection events can be seen, both crossing the arrow for event 2. These events are reflections on the left and bottom of the taper. In contrast with the reflection from the top taper, the bottom taper reflection happens at normal incidence angles as well because it includes a density contrast. Because the material properties in the bottom layer are larger than those of the top layer but we taper to the same density and S-wave velocity on all sides, the taper contrast is larger on the bottom. In addition, the P-wave velocity in the bottom layer is twice as large as in the top layer. Even though the size of the taper is three peak wavelengths on top, it is therefore only half of that on the bottom. Despite all these reinforcing factors, the amplitude of the bottom taper reflection is relatively small compared with the reflection event 2 that we want to model.

In Figures 6 and 7, we investigate whether the taper reflections influence the reflection event whose phase we aim to model. We interpret the late arrivals in Figure 5b and 5c as reverberations caused by the introduced taper. They do not interfere with reflection event 2 because they arrive significantly later. Figure 6 shows a trace at $x = 3.67$ km. The black, blue, and red curves plot the traces of the full-domain true, full-domain tapered, and local domain simulations. Because there are no waves leaving the local solver and subsequently reentering, the elastic solver of Robertsson and Chapman (2000) does not discard any events. It correctly computes the wavefield resulting from the injected CDA background wavefield scattering on the elastic model updates. The boundary integral described in equation 6 is in theory also exact in this case because the amplitude compensation factors are 1.0 and the integral simplifies to equation 5. That the red and blue curves in Figure 6 almost perfectly overlap is therefore not a surprise. It only demonstrates that the local solver is correctly implemented, at least for this geometry. The next thing to notice is that the black curve of the true simulation is also matched quite accurately. The only appreciable deviation happens at $t \approx 1.6$ s. Looking at $x = 3.67$ km and $t \approx 1.6$ s in the shot gathers in Figure 5, this deviation appears to happen at the intersection between the head wave and event 4 (i.e., the top-taper reflection).

In this example, the arrival time of reflection event 2 is known exactly. We use a 0.15 s window and we measure the phase of the reflection in the true simulation, the full-domain tapered simulation, and the local simulation. Figure 7 shows that the phase of the true simulation is quite accurately matched by that of the local domain simulation. On the horizontal axis, we convert the horizontal coordinate to incidence angle.

The three colors correspond to those in Figure 6. The solid curves represent the measured phase of the windowed reflection. They start at a phase of $\approx 315^\circ$, or equivalently $\approx -45^\circ$ due to wrapping, which corresponds to the well-known phase shift in 2D simulation. From these three solid curves, the dashed phase shift curves are obtained by subtracting the zero incidence angle phase and removing the 360° phase wrapping. The tapered full-domain simulation and the local domain simulation differ by at most 1° , as expected. The phase of the true simulation is also matched quite accurately with a maximum error of 4° and average error of 2° . We used the same material properties in our two layer study as Zhu and McMechan (2012), and our results show good agreement with the simulations in Figure 1b of their paper. With this example, we have demonstrated that the coupled CDA-EL local solver can be used to accurately model phase shifts.




Figure 6. Traces at $x = 3672.5$ m for a shot at $x = 1.0$ km. Applying the described single scalar amplitude correction to the full-domain taper and local shot gathers. Notice the excellent agreement between the different models. Only at the intersection of events 3 and 4 (between 1.50 and 1.65 s) do we see some differences. The simulation on the true model does not have the taper influence.

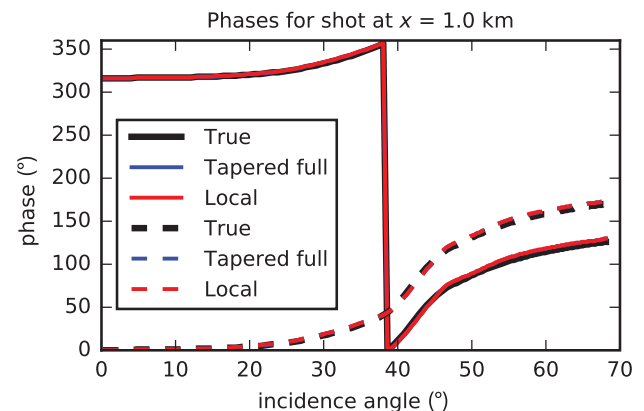


Figure 7. Phase shift with respect to incidence angle at the reflector. The dashed lines are obtained by subtracting the zero-offset reflection phase (which is approximately shifted by -45° because we are using a 2D model). The phase error between the estimated phase in the true and tapered/local results is 4° at most.

Synthetic North Sea model

The synthetic model we focus on in this section, shown in Figure 8, has features similar to what is observed in the North Sea. We investigate whether we can model postcritical reflections, from a chalk layer at a depth slightly below 3.0 km with velocities sometimes exceeding 5000 m/s (Japsen, 1998), with acceptable accuracy. This chalk layer is composed of two regions, each with its own uniform P-wave velocity, density, and S-wave velocity. One region is the bulk chalk, and the other region centered at approximately $x = 5.75$ km is a small (potentially hydrocarbon-bearing) anomaly. The top 80.0 m of the model are water with zero S-wave velocity. The density is derived from the P-wave velocity model using the relation from Gardner et al. (1974). The S-wave velocity is obtained from the P-wave velocity assuming a Poisson solid. The only deviations from these relations are present in the ocean (fluid), top 700 m sediment (low V_S , unconsolidated), and the potential hydrocarbon-bearing anomaly. For the latter, we reduce the P-wave velocity and S-wave velocity by 300.0 and 200.0 m/s and the density by 200.0 kg/m³ compared with the bulk chalk. The grid spacing used in this study is 10.0 m in both coordinate directions. We use a relatively low-frequency wavelet of 6.0 Hz to prevent any appreciable dispersion. In practice, an elastic solver with higher order accuracy would be used and the frequency of the seismic wavelet could then be increased. PML boundary conditions are placed on all sides of the model for simplicity. Sources and

receivers are placed at a depth of 10.0 m. We model six sources with locations 2.0, 2.5, 3.0, 3.5, 4.0, and 4.5 km. Receivers are continuously placed between $x = 5.0$ and 12.0 km. With this geometry for the sources and receivers, any chalk reflection recorded by the receivers will happen within the fully elastic blue box of Figure 8. Figure 9 shows the true shot gather simulated on the full elastic model of Figure 8 for a shot at $x = 2.5$ km. The blue lines window the top of the reservoir reflection whose phase we want to model using the local solver. The red box highlights the region we focus on in later figures for more detailed comparisons.

We first investigate whether the local solver can still accurately model the local reflection in this more complicated North Sea model, with an emphasis on phase modeling. For this, we investigate two different local solver configurations, called “local 1” and “local 2.” The configurations only differ in how they taper the elastic interior of the local solver, which contains the object (e.g., in this case chalk) whose reflection we want to model, to the CDA propagation of the Green’s functions in the exterior. Then, as a demonstration of the value of postcritical reflections, we perform a simple phase inversion using the local solver.

Table 1 summarizes the two taper configurations that we investigate. The first configuration local 1 is the same configuration as in the two-layer example, where we taper to constant density and zero S-wave velocity. In the local 2 configuration, we no longer taper to zero S-wave velocity and the taper is partially removed. We use the amplitude compensation factors from the theory section for both local solver configurations, but the values will differ for each configuration. In the two-layer case, almost all of them evaluated to 1.0 for local 1; the same is true for local 1 on the North Sea model of Figure 8. In the two-layer example, only $l_{V_S}^{\text{src}}(\mathbf{x}_s)$ was unequal to one because the S-wave velocity at the source grid node was nonzero. In the North Sea example, this final compensation factor is 1.0 because the source is in the sea with zero S-wave velocity. But now the density factors $l_{\rho}^{\text{inj}}(\mathbf{x})$ and $l_{\rho}^{\text{rec}}(\mathbf{x})$ are constants unequal to one for local 1

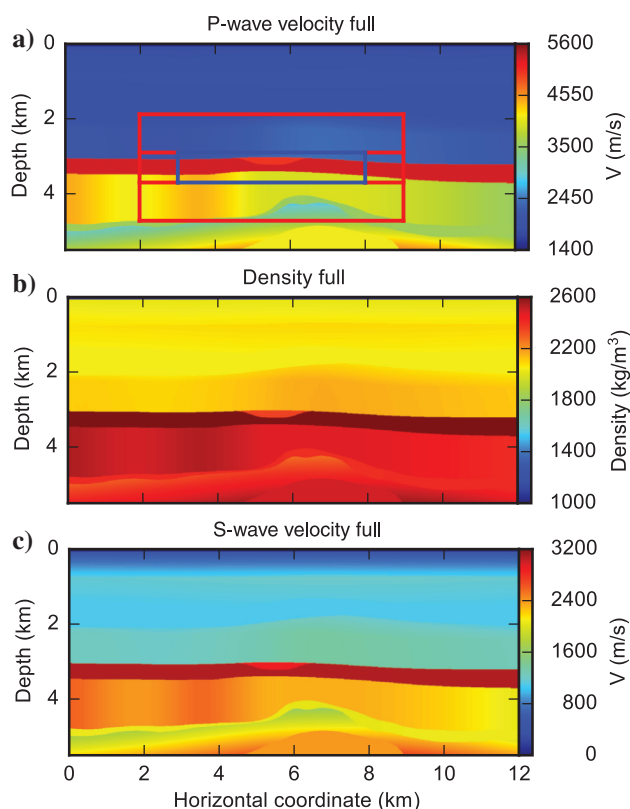


Figure 8. True V_P , ρ , and V_S for the North Sea model. The image of V_P is the background velocity model on which the CDA Green’s functions are computed. (a) The outer red box and inner blue box correspond with those in Figure 3. The lines connecting the outer red box and inner blue box are used in heuristic local solver configuration, in which a part of the taper is removed.




Figure 9. Shot gather for a source at $x = 2.5$ km with the red box indicating the zoomed region shown for the local solvers in Figure 10. This gather is generated on the true elastic model. The blue line represents the window later used to measure the phase of the reflection.

because the density at the source and receiver locations is not the same as the density taper value.

Local 1: Tapering the local model to CDA

In the local 1 configuration, we taper the fully elastic model inside the blue box of Figure 8 to CDA at the outer red box. This is also the boundary of the local solver in this experiment, with taper values given in Table 1. The region in the red box of Figure 9 is expanded in Figure 10a. Figure 10 shows shot gathers for different simulations for a source at $x = 2.5$ km. Figure 10b shows the full-domain elastic simulation on the model, which is fully elastic in the blue box and is tapered toward CDA on the outer red box. Figure 10c shows the coupled CDA-EL local solver on the same model within the outer red box. Figure 10d shows the full-domain elastic simulation on a CDA model using the true P-wave velocity from Figure 8a with constant density and zero S-wave velocity. We obtain exactly the same result as Figure 10d when we use a highly efficient dedicated CDA solver, instead of an elastic solver with only P-wave velocity. Arrow 1 shows the position of the top of reservoir reflection. Arrow 2 shows the location where the head wave and the reflection separate. Arrow 3 shows the horizontal coordinate along which traces are extracted for comparison. Figure 11 plots traces extracted from Figure 10 at $x = 8.0$ km and displays them in the same order. Notice how, as expected, the full-domain tapered result is the same as the local 1 result, just as in the two-layer scenario. The local 1 simulation quite accurately models the top of the reservoir reflection, although some small differences are visible, particularly in amplitude. Compared with the CDA simulation in green, it is clear that using the elastic properties in the region of interest improves the accuracy of the modeled reflection.

Local 2: Heuristic tapering approach

Configuration local 2 only uses a thin horizontal taper, marked by red lines, on either side of the blue box in Figure 8a. It tapers to the density and S-wave velocity values listed in Table 1. It is worth noting that because the earth in Figure 8 is a Poisson solid along the top and bottom boundaries of the blue box, the amplitude compensation factors $l_{V_S}^{\text{inj}}(\mathbf{x}) = 2/3$ and $l_{V_S}^{\text{rec}}(\mathbf{x}) = 3/2$ are constant. These compensation factors vary spatially along the red portions of the boundary due to the tapering of the S-wave velocity. The density amplitude compensation factors vary less than those for the S-wave velocity.

Table 1. Summarizations of the taper properties for the two local solver configurations. Entries for ρ and V_S indicate the taper end value at the boundary. The entry N/A means that the particular taper is not present in this configuration. Local 1 tapers from the blue box to the outer red box in Figure 8. Local 2 only retains the horizontal tapers. Graphically, this corresponds to the thin red box extending sideways from the blue box in Figure 8.

Solver configuration	Taper sides			
	Top	Bottom	Left	Right
Local 1	ρ : 2100.0 kg/m ³ V_S : 0.0 m/s	ρ : 2100.0 kg/m ³ V_S : 0.0 m/s	ρ : 2100.0 kg/m ³ V_S : 0.0 m/s	ρ : 2100.0 kg/m ³ V_S : 0.0 m/s
Local 2	N/A	N/A	ρ : 2200.0 kg/m ³ V_S : 1300.0 m/s	ρ : 2200.0 kg/m ³ V_S : 1300.0 m/s
—	—	—	—	—

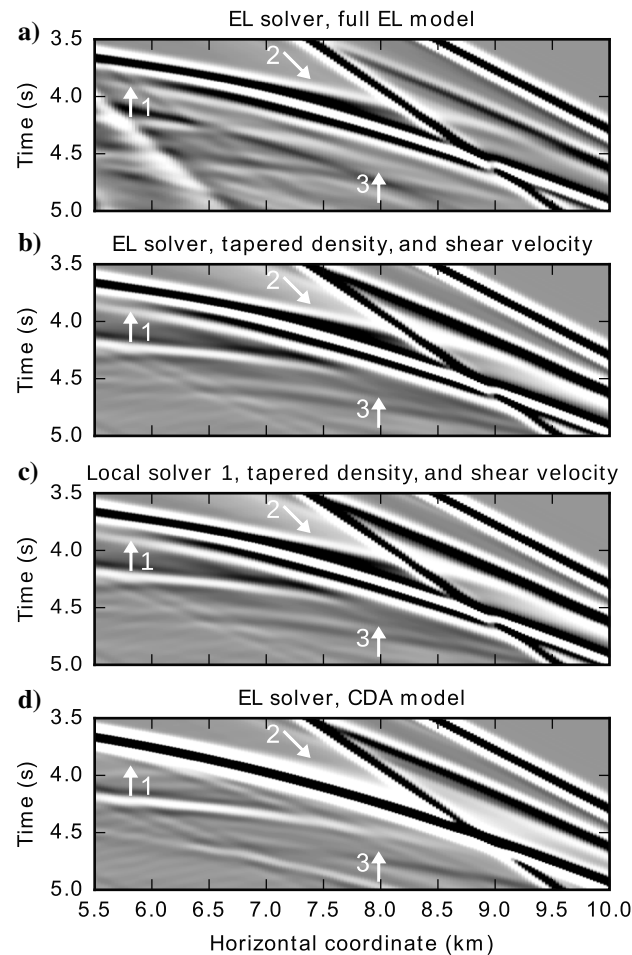


Figure 10. Sections of shot gathers for a source at $x = 2.5$ km in four different scenarios, within the region marked with the red box in Figure 9. The horizontal axis is the receiver coordinate. Three different locations of interest are marked. (1) The top of reservoir reflection. (2) The head wave separating from this reflection, marking the transition to the postcritical regime. (3) Location where the traces are plotted in Figure 11. The phase has obviously rotated with respect to the nearer offset reflections before the critical angle. The amplitude scale is the same in all figures.

For the local 1 simulation, we used the region between the blue box and the outer red box in Figure 8a to taper the elastic model to the CDA model. The size of the taper substantially adds to the size of the local domain, reducing the computational benefits. Also in local 1, the S-wave velocity is tapered from approximately 1300 m/s for the overlying sediment in the blue box to 0.0 m/s for the CDA properties at the outer red boundary. Although we expect that the effect of transmission on phase is small, it is not zero especially when the taper contrast is large. In this section, we do not taper to CDA zero S-wave velocity anymore, reducing the contrast across the taper. The background model on which the Green's functions are computed is still CDA. This means that the elastic model perturbations to the CDA background now intersect the boundary of the local solver. This causes inaccuracies both when injecting the

pressure wavefield from the CDA Green's function and when extrapolating the scattered pressure field to the physical receivers using the boundary integral equation 6. When injecting the wavefield in the local solver, the factors $l_{V_S}^{inj}(\mathbf{x})$ and $l_{\rho}^{inj}(\mathbf{x})$ are no longer constants. These scaling factors approximately compensate for smooth variations of their respective material properties. This scaling is not perfect, in particular when there are lateral variations in these constants along the injection boundary. In addition, there are some inherent inaccuracies when transforming the pressure field stored at nine layers around the boundary of the local solver into a velocity field and stress tensor with $\tau_{xx} \neq \tau_{zz}$ and $\tau_{xz} \neq 0$ when $V_S \neq 0$. By no longer tapering to CDA, the model perturbation to the CDA background model is no longer limited to the interior of the boundary integral of equation 6 as is required by theory. In addition to inaccuracies in the injected wavefield, there is therefore also an inaccuracy in propagating the scattered pressure wavefield of the chalk PP-reflection from the local solver to the receiver locations. By modeling the chalk PP-reflection in the configuration local 2, we investigate the impact for the model in Figure 8. Further theoretical developments are needed to quantify the errors in injection and extrapolation in general models.

Figure 12a shows the shot gather for this heuristic approach called local 2, plotted in the same window as those in Figure 10. The results using the local 2 configuration are similar to that of local 1 in Figure 10c, even though local 2 uses empirical amplitude compensation factors. To emphasize the effect of the amplitude compensation factors, we also investigate the scenario when they are not applied. To be precise, all amplitude compensation factors are set to 1.0 except for $l_{V_S}^{inj}(\mathbf{x})$, which is applied automatically in the injection procedure. The shot gather in Figure 12b therefore represents the scenario of not paying attention to the effect of density and S-wave velocity on the amplitude of the wavefield.

Figure 13 shows traces extracted at $x = 8.0$ km from the simulations in Figure 12. The vertical scale is different than in Figure 11, but the true trace in black is exactly the same. We see that the shape of the reflection computed by local 2 is similar to that of the true simulation at $x = 8.0$ km. The shape of the unscaled reflection is clearly not satisfactory, emphasizing the importance of the amplitude compensation factors in the heuristic local 2 configuration. To

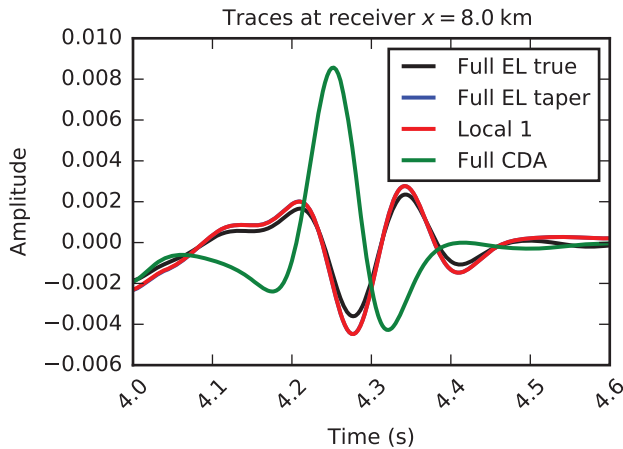


Figure 11. Traces corresponding to the four scenarios in Figure 10a–10d in the same order at $x = 8.0$ km marked by the number 3 in the first panel. The red line almost perfectly matches the blue line. It also matches the shape of the true elastic response in black quite well. The green line is almost reversed in polarity, which clearly shows that CDA physics do not accurately model the phase of elastic reflections.




Figure 12. Plots of the wavefield simulations for configurations local 2 and local 2 without amplitude compensation factors. Similar to Figure 10, the source is placed at $x = 2.5$ km. The amplitude range used for plotting is the same in all figures.

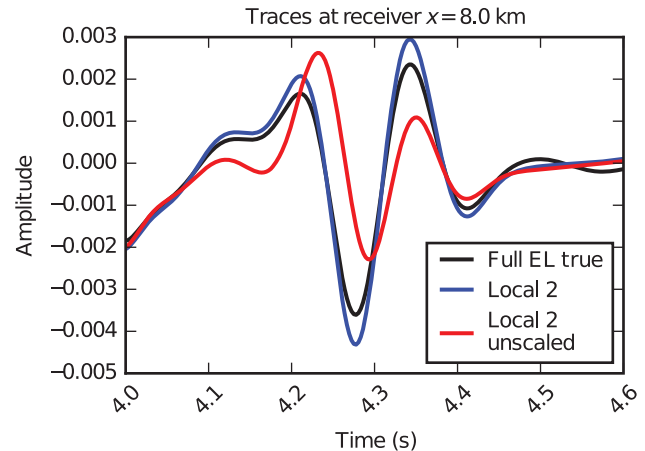


Figure 13. The black trace is the same true trace as in Figure 11. The other two traces correspond to the solver configurations in Figure 12. The red line shows that, without amplitude compensation factors, the phase of the modeled reflection is clearly less accurate.

quantify the match of local 1 and local 2 with the true reflection, we measure the phase of the reflection within a 0.2 s window, marked by the blue lines in Figure 9. The three extracted gathers are plotted in Figure 14.

For each trace in the windowed shot gathers of Figure 14, the phase is determined by using the algorithm of Fomel et al. (2010), summarized by Zhu et al. (2011), with results plotted in Figure 15a. Similar to Figure 7, the receivers nearest to the source at $x = 2.5$ km have measured phase close to 315° . This -45° shift for near offset is the result of 2D effects. At certain offsets, events intersect with the chalk reflection and as a result the solid red line oscillates rapidly. See the region between $x = 8.25$ and 9.75 km, for instance, in all three plots. At these locations, the combination of the direct wave and ocean bottom reflection intersects with the chalk reflection. The ocean bottom reflection is different in the true simulation in which it is affected by the elastic contrast at the ocean bottom and in the local simulations in which the ocean bottom reflection is CDA. The superposition of this event with the chalk reflection will have a different picked phase than that of the chalk reflection alone. This change will not be reproduced by the local solver because the ocean bottom reflection has different physics than in the true simulation.

Another example is the region between $x = 5.0$ and 5.5 km for the true simulation. Here, a surface wave intersects the chalk reflection, perturbing the picked phase. One approach to remove some of these oscillations is to enforce smoothness on the picked curve using regularization (Fomel et al., 2010). Instead, we decide to mute certain positions where events crossing the chalk reflection perturb its phase measurement. The mute windows for the source at $x = 2.5$ km are shown with gray shading in Figure 15a. A more desirable option would be to use processing to remove these events, but this is outside the scope of this paper.

To get the phase shift of the reflection as a function of position, we must add approximately 45° to correct for 2D effects. Including this shift, at large offsets (large incidence angles) the phase shift approaches 180° . At intermediate angles, the shift is more than 180° . This very large phase shift is caused by the presence of two critical angles. The first critical angle originates from the P-wave velocity in the chalk, which is larger than that of the overlying sediment. In addition, the S-wave velocity in the chalk is also larger than the P-wave velocity in the overburden. This causes the second critical angle, now for the PS converted wave. This angle also influences the phase of the PP reflection we focus on. Figure 15b shows the absolute (wrap corrected) phase error with respect to the true phase for the different simulations.

We repeat the experiment for the remaining shots at 2.0, 3.0, 3.5, 4.0, and 4.5 km. The average phase error is 8.93° for local 1 and 4.09° for local 2. From which we conclude that local 2 models the chalk reflection more accurately than local 1. In addition, local 2 is also faster than local 1 because the local solver is smaller. For these reasons, we use configuration local 2 in the inversion experiment.

Inversion

We now investigate whether the phase of the reflection event, generated by the local solver, can be used to invert for material properties at a reflector of interest. We do this by minimizing the least-squares objective function of equation 7. Figure 15b shows the phase residual for a source at $x = 2.5$ km where the model \mathbf{m} on which \mathbf{u} is simulated is the true model. Because neither of the local solver configurations model the reflections exactly, there is a phase

residual even on the true model. Phase residuals in the muted regions of each shot are not included in vector $\mathbf{r}(\mathbf{m})$. The nonmuted phase residuals for all six sources, equally spaced between 2.0 and 4.5 km, are concatenated to form the vector $\mathbf{r}(\mathbf{m})$.




Figure 14. The extracted gathers within the window indicated by blue lines in Figure 9. Plot (a) shows the true gather, and plots (b and c) show the local 1 and local 2 simulations, respectively. The amplitude scale is the same in all figures.

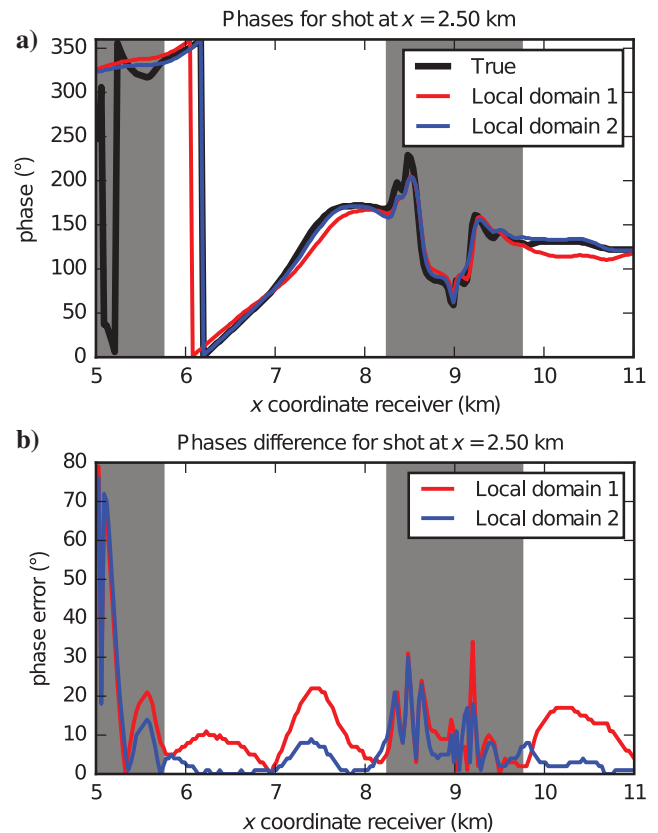


Figure 15. (a) Comparing the phases of the windowed reflections in Figure 14. Intersecting events cause abrupt deviations from a smooth phase transition. We, therefore, choose to mute the gray regions. (b) Plotting the difference with the “true” phase for all simulations in Figure 15a. The phase difference is corrected for phase wrapping. Average misfits in the nonmuted regions: local 1: 10.07° and local 2: 2.67° .

Table 2. The P-wave velocity, density, and S-wave velocity for the potentially hydrocarbon-bearing anomaly and the bulk chalk. The initial guess is improved significantly during this phase-only inversion.

Hydrocarbon anomaly	True	Initial		Inverted	
		Value	Error	Value	Error
V_P (m/s)	5000.0	4200.0	16.0%	4862.1	2.8%
ρ (kg/m ³)	2350.0	2000.0	14.9%	2314.4	1.5%
V_S (m/s)	2860.0	2200.0	23.1%	2832.3	1.0%
Bulk chalk	True	Initial		Inverted	
		Value	Error	Value	Error
V_P (m/s)	5300.0	4200.0	20.8%	5485.9	3.5%
ρ (kg/m ³)	2650.0	2000.0	24.5%	2721.6	2.7%
V_S (m/s)	3060.0	2200.0	28.1%	3056.4	0.1%




Figure 16. Normalized model evolution during inversion. Number “1” refers to the potentially hydrocarbon bearing anomaly, and number “2” refers to the bulk chalk.




Figure 17. Phase for source at $x = 2.5$ km for selected inversion iterations. The grayed-out region is muted, and the phase does not influence the inversion.

The question of how to parameterize the model is important. [Zhu and McMechan \(2012\)](#) show that, in general, the amplitude and phase of a reflection depend on four ratios of material properties at a reflector. Because the full elastic material contrast at a reflector is defined by six quantities (P- and S-wave velocities and density on both sides of the reflector), at least two of them need to be fixed. [Zhu and McMechan \(2012\)](#) fix the P-wave velocity and density in the upper layer and demonstrate that the method is sensitive to bias in these picks. In this study, we assume the overburden is known and invert only for the material properties of the bulk chalk and the potentially hydrocarbon bearing anomaly at $x = 5.75$ km. These regions are defined by the P-wave velocity, density, and S-wave velocity, so the model vector \mathbf{m} contains six elements.

Table 2 shows the initial, inverted, and true model parameters for the potentially hydrocarbon-bearing anomaly and the bulk chalk. The evolution of the model \mathbf{m}_i as a function of iteration number i is plotted in Figure 16. Each curve is normalized by the true value, so all curves should approach 1.0 in case of a perfect inversion. The subscript 1 refers to the material properties of the potentially hydrocarbon bearing anomaly in the chalk anticline at $x = 5.5$ km, and subscript 2 refers to the bulk chalk. The first five iterations show large and sometimes oscillatory updates. In particular, the density appears to diverge during initial iterations. During later iterations, the convergence behavior becomes smooth and incremental. The same precomputed Green’s functions are used for each iteration; no recomputation is needed when the local model is updated.

Figure 17 shows the evolution of the phase for a collection of selected iteration numbers. We see that the phase residual gets progressively smaller and smaller as the inversion progresses. The phase curve of the true model (black) and that of the true local model (red) are not the same. The CDA-EL local solver is not perfect, and as a result, we do not expect to get a perfect model recovery. As a result of this mismatch, we find models that fit our modeled data better than the true model. This is seen in Figure 18 where we plot the objective function during inversion and mark the objective function on the true model with the dashed black line. Continuing the inversion beyond 23 iterations will therefore never result in the true

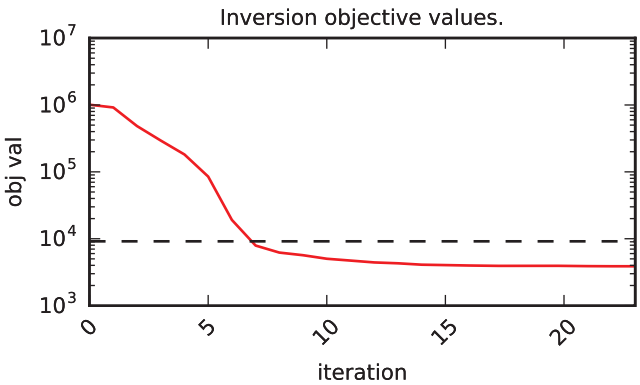


Figure 18. Log plot of the objective value of equation 7. The black dashed line is the objective value from the true model using local 2.

model because this has a higher objective value associated with it. Despite these inaccuracies, we see in Table 2 that all the model parameter estimates at iteration 23 have errors of less than 4%.

DISCUSSION

The local solver introduced in this paper is based on an elastic solver, which is fourth-order accurate in space. With higher order accuracy stencils or other dispersion optimized stencils, the number of nodes per wavelength can be reduced, but stencil evaluation may be more expensive. Depending on the particular implementation, this may reduce the time required for computing the full-domain Green's functions and performing the local solves. The increase in grid spacing and time step will reduce the storage cost for Green's functions, but an increased stencil size requires that more than nine layers of CDA Green's functions are stored around the boundary of the local solver. All these trade-offs need to be considered when determining which particular stencil achieves the desired level of dispersion at the lowest computational cost.

In its local 1 configuration, the local solver generates exactly the same results as a full-domain elastic solver on the same tapered elastic-CDA model. The amplitude compensation factors are all constants and the boundary integral in equation 6 is just a scalar multiple of the exact CDA integral in equation 5. The amplitude compensation factors only scale the resulting wavefield to approximately compensate for elastic properties at the source and receiver locations in the true model. In contrast to local 1, the wavefields generated by local 2 do not correspond to an equivalent full-domain elastic simulations on some tapered model. The injected wavefield is modified on a node by node basis through the amplitude compensation factors $l_{V_s}^{\text{inj}}(\mathbf{x})$ and $l_{\rho}^{\text{inj}}(\mathbf{x})$. When extrapolating the wavefield back to the receivers, we use similar spatially varying amplitude compensation factors within the boundary integral of (equation 6). The scaled injection and extrapolation are not exact. For the particular North Sea model of Figure 8, we observe a good result for local 2, but without a rigorous theoretical framework, it is hard to predict whether the local 2 configuration transfers well to other models. Willemssen (2017) develops some preliminary empirical guidelines in the absence of a theoretical justification: (1) We apply a taper to the local solver on those sides where the layer of interest crosses its boundary, as in local 2. (2) We pick the contour of the local domain in such a way that layers with strong elastic material contrasts do not cross the boundary. (3) A more generally applicable routine would be to smooth the entire P-wave velocity, S-wave velocity, and density models except for the reflector of interest.

It should be noted that the chalk reflection is in fact the combination of the top and bottom chalk reflection with the 6 Hz Ricker we use. From a practical point of view, this superposition does not matter to our approach because the true data and the simulated data contain this superposition. When minimizing the misfit in phase in the inversion, we observe in Figure 16 that the P-wave velocity and density are not as well-recovered as the S-wave velocity. This good recovery of the S-wave velocity may be helped by the chalk layer having a PS-critical angle. Zhu and McMechan (2014) observe something similar and show that a combined inversion using amplitude and phase information gives slightly better results. In Figure 13, we see that the modeled amplitude of the local 2 trace is close to that of the true trace. It should be noted that the overburden sediment in this paper includes several layers, but that the contrast at these layers is weak. Most variation in material parameters happens smoothly

within these sediments. Future studies focusing more on the amplitude should investigate whether the amplitude of the local reflection remains relatively accurately modeled in a more heterogeneous overburden, perhaps using methods introduced by Hobro et al. (2014).

In this paper, we have investigated coupling a CDA background model to a local elastic model. The reason for doing this is to reduce the cost of precomputing Green's functions by leveraging the high efficiency of CDA wave solvers. The price we pay for this is that the tapers make the local domain slightly larger than the region of interest (i.e., the blue box in Figure 8), which makes the local solves more expensive. In addition, there is a small reduction in accuracy. An intermediate approach would use variable density acoustic Green's functions, which increases computational cost but potentially also the accuracy because no density compensation factors are required anymore. If precomputing a large set of full-domain elastic Green's functions is not a problem, then the original local solver introduced by Robertsson and Chapman (2000) could be used unaltered. In this case, no taper is needed and there is no loss of accuracy. This approach may be less beneficial if no detailed elastic overburden model is available, but it may be the better path when there are many small-scale perturbations in the overburden precluding acoustic simulation (Cance and Capdeville, 2015). If many small-scale perturbations with respect to the wavelength are present in the true earth, then the effective anisotropy this induces cannot be accurately modeled using an acoustic solver (Cance and Capdeville, 2015). If finite frequency effects in the overburden are insignificant, it would be interesting to investigate whether the local elastic solver can be coupled to the acquisition using a ray-based approach. Not only would this drastically reduce the precomputation cost, but it would also reduce the storage cost significantly.

The computational cost of advancing the local solver one time step is cheaper than advancing the full-domain solver. The speedup factor is proportional to the difference in the size of the full domain and the local domain. An additional cost savings is available to the local solver because the reflection of interest would not enter the local domain until some time $t_0 > 0$. Even though we have not done so in this paper, we only need to model between this time t_0 and the time t_1 so that the wavefield leaves the local domain.

CONCLUSIONS

In this paper, we introduce a mixed local solver where CDA Green's functions are used to couple a local elastic solver to the rest of a larger model. We first investigate this method on a simple two-layer model in which this mixed local solver matches the full-domain elastic reflection response almost perfectly. After this, we try the local solver in two different configurations on a more complicated heterogeneous velocity model. The local 1 configuration directly corresponds to an equivalent tapered elastic full-domain simulation. The local 2 configuration has no full-domain equivalent and is heuristic in nature. For this configuration, we use spatially varying amplitude compensation factors. These take into account the effect of density and S-wave velocity on the pressure of the background Green's function. As a result, these amplitude compensation factors make the reflection response from the mixed local solver approximate the full-domain elastic response. We observe that the local 2 configuration, in which we use tapers only on the vertical boundaries of the local domain, is computationally efficient and that it also models the reflection of interest with high accuracy on the investigated model. We show that for a phase-only

inversion, we are able to quite accurately retrieve the model parameters at the reflector of interest. In combination with the computational efficiency of the local solver, this demonstrates the potential of this approach.

ACKNOWLEDGMENTS

We thank X. Zhu and F. Broggini for insightful discussions. We are also grateful to the reviewers, whose suggestions have improved this paper. This work is supported by Chevron and with grants from the Natural Sciences and Engineering Research Council of Canada Industrial Research Chair Program (IRCPJ 491051-14) and the Research and Development Corporation of Newfoundland and Labrador (5405.1085.104) and by the Hibernia Management and Development Corporation. We are grateful to F. ten Kroode for providing us with a North Sea velocity model, which we modified for the example in this paper. We would also like to thank MIT and the ERL consortium members for their support.

APPENDIX A

STANDARD ELASTIC LOCAL SOLVER

In this appendix, we introduce the elastic local solver of [Robertsson and Chapman \(2000\)](#). In their approach, the full-domain Green's functions and the local wavefields are computed using the same elastic staggered grid implementation. In this study, we use the algorithm introduced by [Levander \(1988\)](#) for elastic wavefield simulations. The equations of motion in 2D are

$$\rho \frac{\partial v_x(\mathbf{x}, t)}{\partial t} = \frac{\partial \tau_{xx}(\mathbf{x}, t)}{\partial x} + \frac{\partial \tau_{xz}(\mathbf{x}, t)}{\partial z} \quad (\text{A-1})$$

$$\rho \frac{\partial v_z(\mathbf{x}, t)}{\partial t} = \frac{\partial \tau_{zx}(\mathbf{x}, t)}{\partial x} + \frac{\partial \tau_{zz}(\mathbf{x}, t)}{\partial z}, \quad (\text{A-2})$$

and the constitutive equations for an isotropic elastic medium are

$$\frac{\partial \tau_{xx}(\mathbf{x}, t)}{\partial t} = (\lambda(\mathbf{x}) + 2\mu(\mathbf{x})) \frac{\partial v_x(\mathbf{x}, t)}{\partial x} + \lambda(\mathbf{x}) \frac{\partial v_z(\mathbf{x}, t)}{\partial z} \quad (\text{A-3})$$

$$\frac{\partial \tau_{zz}(\mathbf{x}, t)}{\partial t} = (\lambda(\mathbf{x}) + 2\mu(\mathbf{x})) \frac{\partial v_z(\mathbf{x}, t)}{\partial z} + \lambda(\mathbf{x}) \frac{\partial v_x(\mathbf{x}, t)}{\partial x} \quad (\text{A-4})$$

$$\frac{\partial \tau_{xz}(\mathbf{x}, t)}{\partial t} = \mu(\mathbf{x}) \left(\frac{\partial v_x(\mathbf{x}, t)}{\partial z} + \frac{\partial v_z(\mathbf{x}, t)}{\partial x} \right). \quad (\text{A-5})$$

A bold font is used to represent vector quantities; i.e., \mathbf{x} is (x, z) . The elastic solver updates the velocities and the stresses in a turnwise manner. First, the velocities in equations [A-1](#) and [A-2](#) are updated. These are then used to update the stresses in equations [A-3](#) to [A-5](#). Equations [A-1](#) to [A-5](#) are solved on a staggered grid with second-order accuracy in time and fourth-order accuracy in space.

Figure 1 shows a section of a staggered elastic grid zoomed in on part of the top-right boundary of the local solver. The four different symbols represent the locations where the elastic wavefield quantities from equations [A-1](#) to [A-5](#) are evaluated, as described in the

legend. On this staggered grid, these wavefield components are evaluated at different physical locations. The black grid lines denote grid cells, each including the four staggered grid symbols. The local solver computes the perturbed wavefield within the interior of the magenta line in Figure 1. This magenta line intersects grid cells on the boundary by design. The model perturbations $\Delta m(\mathbf{x})$ are restricted to inside the inner red line. At the shaded symbols outside the magenta line, the scattered wavefield quantities are computed instead of the perturbed wavefield quantities. These scattered wavefields become nonzero only when $\Delta m(\mathbf{x})$ is nonzero. The local solver of [Robertsson and Chapman \(2000\)](#) requires storage of the background wavefields at all locations between the two red lines, as discussed below.

The order of accuracy with which the spatial derivatives in equations [A-1](#) to [A-5](#) are discretized determines the width of the stencils. As a concrete example of how the governing equations are discretized on the staggered grid, we now give the update equation for the perturbed z -component of the velocity in equation [A-2](#). Graphically, this is represented by the leftmost blue stencil in Figure 1, which is centered around the unfilled (i.e., perturbed) v_z symbol. For this example, we ignore that this particular update stencil involves some shaded symbols and we assume all of the symbols are unfilled. In other words, we assume for now that the stencil only involves perturbed wavefield components. We update the velocity,

$$v_z(\mathbf{x}_{(i,k)}, t) = v_z(\mathbf{x}_{(i,k)}, t - \Delta t) + \Delta v_z(\mathbf{x}_{(i,k)}, t), \quad (\text{A-6})$$

using

$$\begin{aligned} \Delta v_z(\mathbf{x}_{(i,k)}, t) = & \frac{\Delta t}{\rho} \left[\frac{c_1}{\Delta x} \left(\tau_{zx}(x_{i+1}, z_k, t - \frac{1}{2}\Delta t) - \tau_{zx}(x_i, z_k, t - \frac{1}{2}\Delta t) \right) \right. \\ & + \frac{c_2}{\Delta x} \left(\tau_{zx}(x_{i+2}, z_k, t - \frac{1}{2}\Delta t) - \tau_{zx}(x_{i-1}, z_k, t - \frac{1}{2}\Delta t) \right) \\ & + \frac{c_1}{\Delta z} \left(\tau_{zz}(x_i, z_{k+1}, t - \frac{1}{2}\Delta t) - \tau_{zz}(x_i, z_k, t - \frac{1}{2}\Delta t) \right) \\ & \left. + \frac{c_2}{\Delta z} \left(\tau_{zz}(x_i, z_{k+2}, t - \frac{1}{2}\Delta t) - \tau_{zz}(x_i, z_{k-1}, t - \frac{1}{2}\Delta t) \right) \right]. \end{aligned} \quad (\text{A-7})$$

Notice how the staggered spatial derivatives are evaluated at the staggered time step (i.e., $t - 1/2\Delta t$). In equation [A-7](#), the fourth-order accurate staggered grid coefficients c_1 and c_2 are $9/8$ and $-1/24$, respectively. The subscript on the position vector $\mathbf{x}_{(i,k)}$ represents the corresponding location of the wavefield quantity in cell number (i, k) . On the right side, the vector \mathbf{x} is decomposed in its x and z coordinates in this 2D example. Extensions to 3D are trivial. An update equation completely analogous to equation [A-7](#) can be derived for the propagation of the scattered v_z^{sc} in the shaded region. This analogous update equation replaces all perturbed wavefield quantities with their scattered counterparts. Updated equations for elastic wavefield quantities other than v_z use analogous derivative stencils to those shown in Figure 1.

Sufficiently far away from the magenta boundary in Figure 1, the updated equations for the perturbed and scattered wavefields are trivial to compute. But within the two red lines, the derivative stencils cross the magenta line separating the locations where the perturbed and scattered wavefields are computed and stored. The updated equations for nodes between the red lines, therefore, require a modification. We now go back to the specific case of the leftmost stencil in

Figure 1. In our previous discussion of equation A-7, we made the assumption that the perturbed wavefield quantities were available at all required locations. With i, k centered at the central cell of the stencil, we observe that the updated stencil accesses the scattered τ_{zz}^{sc} at the shaded symbols in cells i, k and $i, k-1$ instead of the required perturbed τ_{zz} . By adding the background field τ_{zz}^0 to the scattered τ_{zz}^{sc} at these locations, we obtain by definition the perturbed τ_{zz} and consequently the update equation A-7 can be evaluated correctly. This background field τ_{zz}^0 is obtained by convolving the precomputed τ_{zz}^0 Green's function with the source wavelet for the particular source that we simulate locally. Applying this correction procedure to all nodes between the red boundaries automatically introduces the background wavefield into the local domain in the case where there are no model updates $\Delta m(\mathbf{x})$.

Similarly, the second stencil in Figure 1 has the opposite problem when updating v_z^{sc} at the shaded symbol. Part of its stencil crosses the magenta line and accesses the perturbed field τ_{xx} stored at the unfilled symbols. At these locations, we need to subtract the background wavefield from the perturbed wavefield so that we can correctly evaluate the stencil for updating the scattered wavefield at the shaded symbol. When the model is updated with $\Delta m(\mathbf{x})$ in the unshaded region this correctly computes the perturbed fields at the unfilled symbols and the scattered field at the filled symbols. Multiple scattering between the background model at the shaded nodes and the potentially perturbed model at the unfilled nodes is computed correctly. By definition, when we talk about the local domain in this paper, we mean the entire local solver grid. Therefore, the local domain is the combination of the shaded and unshaded nodes, surrounded by its PML boundary condition outside the outer yellow line.

APPENDIX B

AN OVERVIEW OF THE AMPLITUDE COMPENSATION FACTORS

In Table B-1, we list the heuristic amplitude compensation factors that are required when $\rho(\mathbf{x})$ and $V_S(\mathbf{x})$ are variable along the injection boundary, or between source and receiver locations. These amplitude compensation factors are derived and discussed in more detail by Willemssen (2017). All these factors compensate the amplitude for smooth variation in density and S-wave velocity which the precomputed CDA (i.e., constant density, zero shear) Green's functions cannot account for. The factors in Table B-1 do not compensate for transmission losses encountered when parts of the pressure wavefield reflect on abrupt density and S-wave velocity perturbations.

The first two terms in Table B-1 modify the amplitude of the precomputed CDA pressure wavefield, which is injected in the local elastic solver. Factor $l_{V_S}^{inj}(\mathbf{x})$ compensates the precomputed CDA wavefield for the, potentially spatially varying, S-wave velocity along the boundary of the local elastic solver. The effect of smooth density variations on the amplitude is introduced by $l_{\rho}^{inj}(\mathbf{x})$.

To propagate the scattered wavefield to the receiver locations in an elastic model, we use the weighted boundary integral equation 6. In this boundary integral, we use the compensation factors $l_{V_S}^{rec}(\mathbf{x})$ and $l_{\rho}^{rec}(\mathbf{x})$. The first fraction within $l_{V_S}^{rec}(\mathbf{x})$ is the inverse of $l_{V_S}^{inj}(\mathbf{x})$. It brings back the amplitude of the scattered wavefield to an S-wave velocity state corresponding to the CDA model. Similarly, factor $l_{\rho}^{rec}(\mathbf{x})$ approximately corrects the pressure from a variable density state to a constant density state. The weighted boundary integral equation 6, therefore, approximately uses a CDA representation of the pressure field along

Table B-1. Definitions of the amplitude compensation factors used in this paper. The terms are discussed in Appendix B. Detailed derivations are available in Willemssen (2017); \mathbf{x}_s and \mathbf{x}_r represent the source and receiver locations, respectively. The position \mathbf{x} is along the boundary of the local solver.

	Definition
$l_{V_S}^{inj}(\mathbf{x})$	$\frac{V_P^2(\mathbf{x}) - V_S^2(\mathbf{x})}{V_P^2(\mathbf{x})}$
$l_{\rho}^{inj}(\mathbf{x})$	$\sqrt{\frac{\rho(\mathbf{x})}{\rho(\mathbf{x}_s)}}$
$l_{V_S}^{rec}(\mathbf{x})$	$\frac{V_P^2(\mathbf{x})}{V_P^2(\mathbf{x}) - V_S^2(\mathbf{x})} \frac{l_{V_S}^{inj}(\mathbf{x}_r)}{l_{V_S}^{inj}(\mathbf{x}_s)}$
$l_{\rho}^{rec}(\mathbf{x})$	$\sqrt{\frac{\rho(\mathbf{x}_r)}{\rho(\mathbf{x})}}$
$l^0(\mathbf{x}_s, \mathbf{x}_r)$	$\sqrt{\frac{\rho(\mathbf{x}_r)}{\rho(\mathbf{x}_s)}} \frac{l_{V_S}^{inj}(\mathbf{x}_r)}{l_{V_S}^{inj}(\mathbf{x}_s)}$

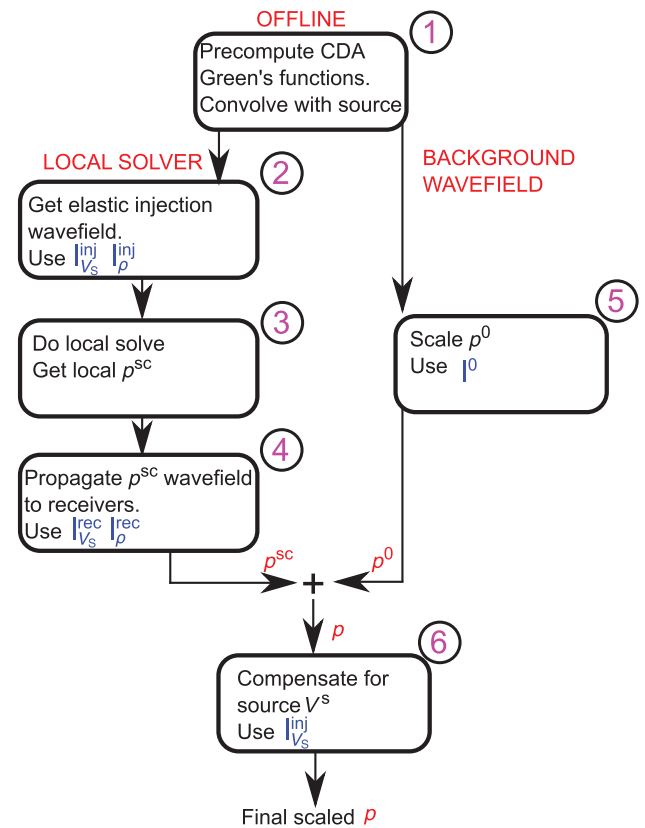


Figure B-1. Workflow of the proposed algorithm. Step 1 computes CDA Green's functions, which is the only time we use a full-domain solver. Step 2 uses injection scaling factors $l_{V_S}^{inj}(\mathbf{x})$ and $l_{\rho}^{inj}(\mathbf{x})$. Step 3 computes the wavefield on the local solver grid. Step 4 uses $l_{V_S}^{rec}(\mathbf{x})$ and $l_{\rho}^{rec}(\mathbf{x})$ for propagating the wavefield to the receivers. Step 5 scales the background wavefield with $l^0(\mathbf{x}_s, \mathbf{x}_r)$ to compensate for differences in V_S and ρ between the source and receiver location. Finally, step 6 scales the perturbed wavefield for the S-wave velocity at the source location with $l_{V_S}^{inj}(\mathbf{x}_s)$.

its boundary. We now need to compensate for differences in the V_s compensation factor $l_{V_s}^{\text{inj}}$ between source and receiver grid nodes. This amplitude variation due to interactions between P-wave velocity and S-wave velocity is not modeled by the CDA Green's functions. The second fraction in $l_{V_s}^{\text{rec}}(\mathbf{x})$ takes care of this.

The wavefield in the perturbed model at the receiver locations is the sum of the background wavefield and the scattered wavefield. The scattered wavefield is scaled using the amplitude compensation procedure above. The same scaling needs to be applied to the background wavefield for this addition to be consistent. This will compensate for the amplitude variation that would be observed in the elastic model if the wavefield would smoothly transition from the source density and S-wave velocity to that at the receiver. This compensation factor is $l^0(\mathbf{x}_s, \mathbf{x}_r)$, where the density and S-wave velocity values at the sources and receivers are extracted from the elastic model from which we want to approximate the pressure wavefield.

So far, we have compensated the amplitude of the perturbed (i.e., background + scattered) wavefields for contrasts in density and S-wave velocity between source and receiver grid nodes. In the elastic solver, we use (Levander, 1988), the amplitude of the perturbed wavefield scales directly with the value of $l_{V_s}^{\text{inj}}(\mathbf{x})$ at the source grid node. So the last amplitude compensation multiplies the perturbed gather by $l_{V_s}^{\text{inj}}(\mathbf{x}_s)$. Figure B-1 shows the total workflow and illustrates where each amplitude compensation factor is applied.

REFERENCES

- Aki, K., and P. G. Richards, 2002, Quantitative seismology, 2nd ed.: University Science Books.
- Borisov, D., S. C. Singh, and N. Fuji, 2015, An efficient method of 3-D elastic full waveform inversion using a finite-difference injection method for time-lapse imaging: *Geophysical Journal International*, **202**, 1908–1922, doi: [10.1093/gji/ggv268](https://doi.org/10.1093/gji/ggv268).
- Buland, A., and H. Omre, 2003, Joint AVO inversion, wavelet estimation and noise-level estimation using a spatially coupled hierarchical Bayesian model: *Geophysical Prospecting*, **51**, 531–550, doi: [10.1046/j.1365-2478.2003.00390.x](https://doi.org/10.1046/j.1365-2478.2003.00390.x).
- Cance, P., and Y. Capdeville, 2015, Validity of the acoustic approximation for elastic waves in heterogeneous media: *Geophysics*, **80**, no. 4, T161–T173, doi: [10.1190/geo2014-0397.1](https://doi.org/10.1190/geo2014-0397.1).
- Capdeville, Y., E. Chaljub, and J. Montagner, 2003, Coupling the spectral element method with a modal solution for elastic wave propagation in global earth models: *Geophysical Journal International*, **152**, 34–67, doi: [10.1046/j.1365-246X.2003.01808.x](https://doi.org/10.1046/j.1365-246X.2003.01808.x).
- Červený, V., 1959, On the reflection of spherical waves at a plane interface with refractive index near to one — Part I: *Studia Geophysica et Geodaetica*, **3**, 116–134, doi: [10.1007/BF02585557](https://doi.org/10.1007/BF02585557).
- Červený, V., and F. Hron, 1961, Reflection coefficients for spherical waves: *Studia Geophysica et Geodaetica*, **5**, 122–132, doi: [10.1007/BF02585356](https://doi.org/10.1007/BF02585356).
- Chapman, C., J. Hobro, and J. Robertsson, 2014, Correcting an acoustic wavefield for elastic effects: *Geophysical Journal International*, **197**, 1196–1214, doi: [10.1093/gji/ggu057](https://doi.org/10.1093/gji/ggu057).
- Downton, J. E., and C. Ursenbach, 2006, Linearized amplitude variation with offset (AVO) inversion with supercritical angles: *Geophysics*, **71**, no. 5, E49–E55, doi: [10.1190/1.2227617](https://doi.org/10.1190/1.2227617).
- Fokkema, J. T., and P. M. van den Berg, 2013, *Seismic applications of acoustic reciprocity*: Elsevier.
- Fomel, S., and M. van der Baan, 2010, Local similarity with the envelope as a seismic phase detector: 80th Annual International Meeting, SEG, Expanded Abstracts, 1555–1559.
- Gardner, G., L. Gardner, and A. Gregory, 1974, Formation velocity and density — The diagnostic basics for stratigraphic traps: *Geophysics*, **39**, 770–780, doi: [10.1190/1.1440465](https://doi.org/10.1190/1.1440465).
- Gong, T., and G. A. McMechan, 2016, Target-oriented linear least squares and nonlinear, trust-region Newton inversions of plane waves using AVA and PVA data for elastic model parameters: *Geophysics*, **81**, no. 5, R325–R338, doi: [10.1190/geo2015-0471.1](https://doi.org/10.1190/geo2015-0471.1).
- Hewett, R., and L. Demanet, and the, PySIT Team, 2013, PySIT: Python seismic imaging toolbox v0.5. (Release 0.6).
- Hobro, J. W., C. H. Chapman, and J. O. Robertsson, 2014, A method for correcting acoustic finite-difference amplitudes for elastic effects: *Geophysics*, **79**, no. 4, T243–T255, doi: [10.1190/geo2013-0335.1](https://doi.org/10.1190/geo2013-0335.1).
- Japsen, P., 1998, Regional velocity-depth anomalies, north sea chalk: A record of overpressure and neogene uplift and erosion: *AAPG bulletin*, **82**, 2031–2074.
- Jin, H., G. A. McMechan, and H. Guan, 2014, Comparison of methods for extracting ADCIGS from RTM: *Geophysics*, **79**, no. 3, S89–S103, doi: [10.1190/geo2013-0336.1](https://doi.org/10.1190/geo2013-0336.1).
- Levander, A. R., 1988, Fourth-order finite-difference P-SV seismograms: *Geophysics*, **53**, 1425–1436, doi: [10.1190/1.1442422](https://doi.org/10.1190/1.1442422).
- Malcolm, A., and B. Willemssen, 2016, Rapid 4D FWI using a local wave solver: *The Leading Edge*, **35**, 1053–1059, doi: [10.1190/le35121053.1](https://doi.org/10.1190/le35121053.1).
- Monteiller, V., S. Chevrot, D. Komatitsch, and N. Fuji, 2013, A hybrid method to compute short-period synthetic seismograms of teleseismic body waves in a 3-D regional model: *Geophysical Journal International*, **192**, 230–247, doi: [10.1093/gji/ggs006](https://doi.org/10.1093/gji/ggs006).
- O'Brien, P., 1963, A note on the reflection of seismic pulses with application to second event refraction shooting: *Geophysical Prospecting*, **11**, 59–72, doi: [10.1111/j.1365-2478.1963.tb02025.x](https://doi.org/10.1111/j.1365-2478.1963.tb02025.x).
- Robertsson, J., and C. Chapman, 2000, An efficient method for calculating finite-difference seismograms after model alterations: *Geophysics*, **65**, 907–918, doi: [10.1190/1.1444787](https://doi.org/10.1190/1.1444787).
- Robertsson, J. O., S. Ryan-Grigor, C. M. Sayers, and C. H. Chapman, 2000, A finite-difference injection approach to modeling seismic fluid flow monitoring: *Geophysics*, **65**, 896–906, doi: [10.1190/1.1444786](https://doi.org/10.1190/1.1444786).
- Sava, P., B. Biondi, and S. Fomel, 2001, Amplitude-preserved common image gathers by wave-equation migration: 71st Annual International Meeting, SEG, Expanded Abstracts, 296–299.
- Smith, G., and P. Gidlow, 1987, Weighted stacking for rock property estimation and detection of gas: *Geophysical Prospecting*, **35**, 993–1014, doi: [10.1111/j.1365-2478.1987.tb00856.x](https://doi.org/10.1111/j.1365-2478.1987.tb00856.x).
- van Manen, D.-J., J. O. A. Robertsson, and A. Curtis, 2007, Exact wave field simulation for finite-volume scattering problems: *The Journal of the Acoustical Society of America*, **122**, EL115–EL121, doi: [10.1121/1.2771371](https://doi.org/10.1121/1.2771371).
- Wang, S., M. V. de Hoop, and J. Xia, 2011, On 3D modeling of seismic wave propagation via a structured parallel multifrontal direct Helmholtz solver: *Geophysical Prospecting*, **59**, 857–873, doi: [10.1111/gpr.2011.59.issue-5](https://doi.org/10.1111/gpr.2011.59.issue-5).
- Willemssen, B., 2017, Problem with a localized nature in exploration seismology: Ph.D. thesis, Massachusetts Institute of Technology.
- Willemssen, B., A. Malcolm, and W. Lewis, 2016, A numerically exact local solver applied to salt boundary inversion in seismic full-waveform inversion: *Geophysical Journal International*, **204**, 1703–1720, doi: [10.1093/gji/ggv547](https://doi.org/10.1093/gji/ggv547).
- Yu, P., J. Geng, X. Li, and C. Wang, 2016, Acoustic-elastic coupled equation for ocean bottom seismic data elastic reverse time migration: *Geophysics*, **81**, no. 5, S333–S345, doi: [10.1190/geo2015-0535.1](https://doi.org/10.1190/geo2015-0535.1).
- Zepeda-Núñez, L., and L. Demanet, 2016, The method of polarized traces for the 2D Helmholtz equation: *Journal of Computational Physics*, **308**, 347–388, doi: [10.1016/j.jcp.2015.11.040](https://doi.org/10.1016/j.jcp.2015.11.040).
- Zhu, X., and G. McMechan, 2011, Comparison of methods for modeling phase variation with angle: 81st Annual International Meeting, SEG, Expanded Abstracts, 3053–3058.
- Zhu, X., and G. McMechan, 2012a, Amplitude and phase in angle-domain common-image gathers: 82nd Annual International Meeting, SEG, Expanded Abstracts, doi: [10.1190/segam2012-0161.1](https://doi.org/10.1190/segam2012-0161.1).
- Zhu, X., and G. McMechan, 2012b, Analysis of phase variation with ray parameter in the τ - p domain: 82nd Annual International Meeting, SEG, Expanded Abstracts, doi: [10.1190/segam2012-0161.1](https://doi.org/10.1190/segam2012-0161.1).
- Zhu, X., and G. A. McMechan, 2012, Elastic inversion of near- and postcritical reflections using phase variation with angle: *Geophysics*, **77**, no. 4, R149–R159, doi: [10.1190/geo2011-0230.1](https://doi.org/10.1190/geo2011-0230.1).
- Zhu, X., and G. A. McMechan, 2014, Amplitude and phase versus angle for elastic wide-angle reflections in the τ - p domain: *Geophysics*, **80**, no. 1, N1–N9, doi: [10.1190/geo2013-0191.1](https://doi.org/10.1190/geo2013-0191.1).
- Zoepprit, K., 1919, VII b. über reflexion und durchgang seismischer wellen durch un-stetigkeitsflächen: *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, **1919**, 66–84.