

Introduction

The world's remaining petroleum reservoirs will be discovered using information obtained by interpreting 3D seismic images. Because petroleum explorationists have purportedly located, and are currently producing, the majority of the world's 'easy oil' - i.e. nearer to surface, in simple geologic settings, and accessible to public companies - future reservoir discoveries will be located increasingly in frontier areas characterized by more complex geology. Accordingly, the search for new hydrocarbon reserves continuously motivates the development of new, and the refinement of existing, 3D seismic imaging techniques able to better image complex subsurface structure.

The first 3D acoustic seismic imaging algorithms implemented in production settings were Kirchhoff approaches (??), which model the wave-equation using paraxial approximations. Although these techniques have been (and still are) applied successfully to seismic data from locations around the world, Kirchhoff migration routinely proves inadequate for areas of complex geology (?). Its main problems are inherited from the underlying ray-theoretic approximation: an inability to model all multi-pathed wavefield phases, and a difficulty in handling boundaries of strong velocity contrast. Both of these issues commonly occur at, say, salt-sediment interfaces. Generating more accurate images that facilitate 3D seismic interpretation thus requires introducing higher fidelity migration algorithms.

Full wave-equation (FWE) methods are another class of seismic imaging techniques that circumvent many problems associated with Kirchhoff migration. There are two key differences between WE and Kirchhoff approaches. First, FWE methods employ operators derived from the full acoustic wave equation - not asymptotic approximations of it - and are thus finite-frequency approaches. Second, FWE methods extrapolate seismic wavefields throughout the migration domain, not just along an incomplete set of rays. These two factors combine to ensure that FWE methods better handle multi-pathed wavefield arrivals and strong velocity contrasts.

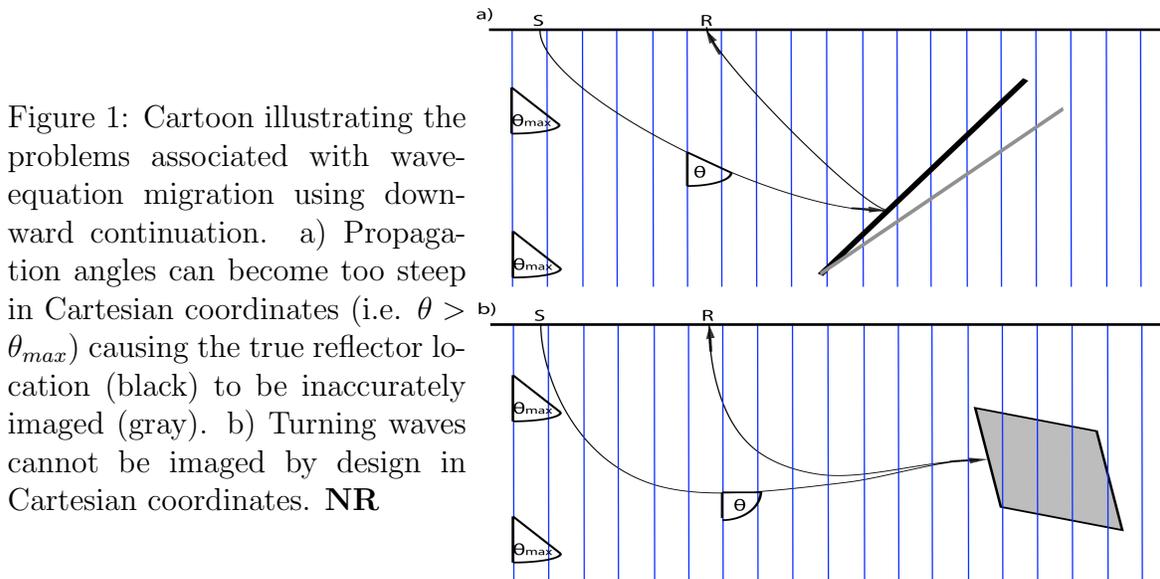
The most accurate FWE method is reverse-time migration (??), which propagates and images individual shot-profiles using operators derived from the full acoustic wave equation. Although this approach generates the highest-quality FWE migration results, the computational cost, I/O throughput, and memory requirements for performing reverse-time migration on typical industry-sized 3D data sets are still too onerous for all but the most powerful computer clusters. Thus, one must turn to approximate FWE methods that retain much of the accuracy of reverse-time migration, but are less computationally demanding and have lower I/O and memory requirements.

One-way wave equations are some of the more commonly used approximations that realize these objectives (??). Conventional migration methods based on one-way equations recursively extrapolate surface-recorded wavefields step-wise in depth. Images are generated at each level by evaluating a physical imaging condition. Relative to reverse-time migration, the computational costs of one-way extrapolation approaches are significantly lower [i.e. approximately 30-50 times (?)], which affords

efficient migration of industry-sized 3D data sets on relatively modest clusters.

Wave-equation migration using one-way extrapolation operators, though, has some significant implementation and conceptual limitations. Figure 1 introduces two of these problems. The first issue is that one-way wavefield extrapolation in laterally varying media becomes inaccurate at steep propagation angles (i.e. where $\theta > \theta_{max}$ in Figure 1a). The angular limit of accurate propagation, with respect to a vertical extrapolation axis, is usually given between $\theta_{max} \approx 45^\circ$ to $\theta_{max} \approx 85^\circ$, depending on the particular implementation. For cases exhibiting nearly horizontal geology, this assumption seldom greatly affects the seismic imaging result. However, inaccurately propagating energy originating from steeply dipping or discontinuous structure (e.g. faults, salt flanks) can generate erroneous subsurface images that may lead to incorrect geologic interpretations. Figure 1a depicts this problem by showing a second gray reflector imaged at the incorrect location.

A second conceptual issue is that downward continuation cannot propagate the upgoing paths of turning waves by design because the extrapolation direction is always oriented downward (see Figure 1b). Thus, conventional migration in Cartesian geometry precludes imaging complex structure with turning wavefield components, which can be detrimental to imaging and interpretation in areas of otherwise poor illumination [though two-pass migration approaches (?) somewhat obviate these concerns]. Figure 2 illustrates these two problems by exhibiting propagation differences between



two-way finite-difference modeling (panel 2a) and one-way Cartesian wavefield extrapolation (panel 2b). The two panels show the wavefield after propagation through the BP velocity model at the 1, 2, 3 and 4s time steps. The Cartesian wavefield is similar in many respects to the two-way modeling, suggesting that it is sufficient for correctly imaging most of the seismic wavefield. However, the two-way modeling in panel 2a contains many additional upward-propagating events not present in the Cartesian one-way panel. The Cartesian wavefields also become inaccurate at steep

propagation angles, in particular through the salt body to the right-hand-side. These Cartesian one-way propagation errors will generate incorrectly positioned (or absent) reflectors and lead to increased interpretation uncertainty. Hence, overcoming the problems of inaccurate high-angle and turning-wave propagation - while maintaining the computational advantages of one-way wave equations - is an important seismic imaging research goal, and motivates most of the work (i.e. generating the more accurate propagation physics in Figures 2c and d) reported in this thesis.

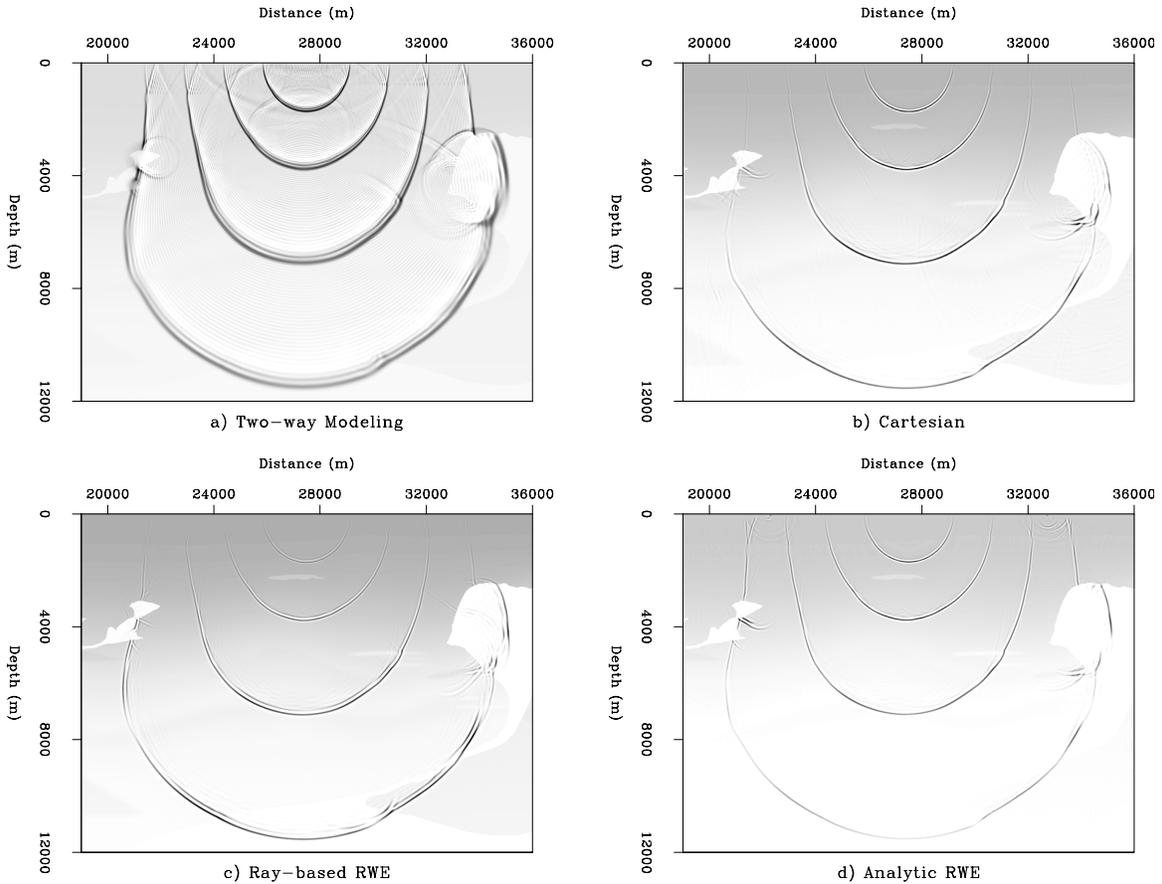


Figure 2: Comparisons between four different acoustic wave-propagation techniques through the BP velocity model. a) Two-way finite-difference modeling. b) Cartesian one-way wavefield extrapolation. c) Ray-coordinate-based one-way Riemannian wavefield extrapolation. d) Analytic coordinate one-way Riemannian wavefield extrapolation. **CR**

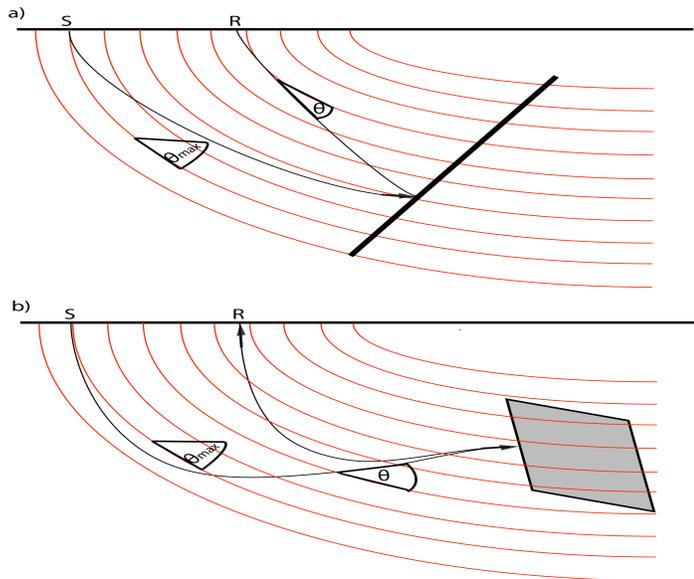
IMPROVING ONE-WAY EXTRAPOLATION

Much geophysical research in the past few decades has been devoted to mitigating problems associated with inaccurate one-way steep- and turning-angle propagation. Because most issues with one-way equations arise where assuming depth-oriented

extrapolation axes, many authors have directly (or indirectly) reevaluated whether or not to use Cartesian coordinates as the geometric basis for FWE migration.

Consider again the ray paths in the Cartesian reference frames in Figure 1. In both cases, the angle between the ray and the vertical extrapolation axis increases beyond the maximum extrapolation angle, leading to inaccurate downward continuation. Defining wavefield extrapolation on coordinate meshes more conformal to the ray paths (Figure 3) would eliminate this problem because the effective extrapolation angle always obeys $\theta < \theta_{max}$. Hence, one central concept involved with extrapolating wavefields on non-Cartesian geometry is to find coordinate systems that lower the relative angle between the propagation direction and extrapolation axis orientation.

Figure 3: Cartoon illustrating how making the migration geometry more conformal to the wave-propagation direction can lead to imaging improvements. a) Propagating wavefields on meshes more conformal to ray-path direction reduces the relative extrapolation angle. b) Improved coordinate system designs enable turning-wave propagation with one-way wave-equations. **NR**

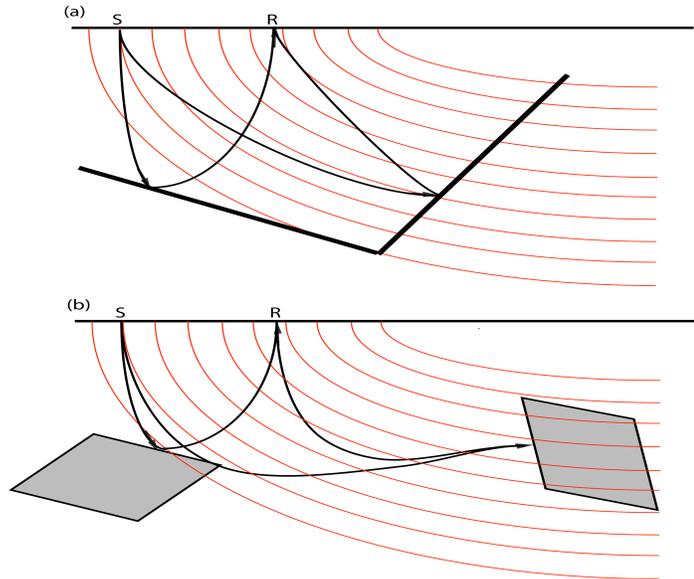


One important caveat is that the migration geometry in Figure 3 is optimized only for individual ray paths. One obvious question is, thus, how can this approach be extended to cases where multiple ray paths interrogate subsurface structures of opposing dip? Figure 4 illustrates this for the steep dip and the turning-wave imaging problems. Evidently, alternate coordinate systems are not, alone, a panacea to the problems of one-way wavefield extrapolation.

One effective way to resolve to this question involves decomposing the total data volume into subsets on the basis of wavefield dips or some other data localization scheme. These partial data volumes then can be extrapolated on separate coordinate systems designed to be more conformal to the wave-propagation direction of the individual data subset.

Figure 5 illustrates the four steps associated with this approach. The first step is deciding which data decomposition scheme and coordinate mesh are optimally matched for the particular imaging task. The next step is performing the data decomposition and setting up the different coordinate meshes. Third, individual migration images are computed on the different coordinate systems. These images are

Figure 4: Cartoon illustrating the problems associated with imaging substructure using wavefields with conflicting dips. a) Coordinate system where the right-hand structure can be imaged using one-way extrapolation, but the left-hand reflector cannot. b) Example where a coordinate enables imaging of part, but not all, subsurface structure. **NR**



stacked into the final image in the last step. Accordingly, imaging in non-Cartesian geometry requires coordinating two important concepts: wavefield decomposition and coordinate systems conformal to propagation directions.

Established migration strategies

A number of established migration methods follow these two basic concepts: directional depth migration (?), turning-wave migration (?), Gaussian beam migration (??), coherent states (?), beam waves (?), plane-wave migration in tilted coordinates (?), and Riemannian wavefield extrapolation (??). Although these approaches share many commonalities, they differ in a number of respects:

- Data domain decomposition - numerous ways of data decomposition exist, including beams, shot profiles, and synthesized plane-wave sections;
- Mesh generation technique - coordinate systems can be generated in a variety of ways, ranging from ray tracing to specifying analytic grids; and
- Extrapolation localization - propagation domains vary greatly in size, ranging from narrow beams to alternate full migration domains.

Two common examples of this approach are Gaussian beams and plane-wave migration in tilted coordinates. Gaussian beam migration is a hybridization of ray and wavefield methods. Generally, beam migrations use ray tracing to generate a skeleton mesh for a suite of take-off angles at each point. The corresponding wave-packets are propagated and imaged along domains defined by a narrow beamwaist around each traced ray. Final images are generated by superposing individual beam images. These powerful approaches have the advantage of coupling ray methods with

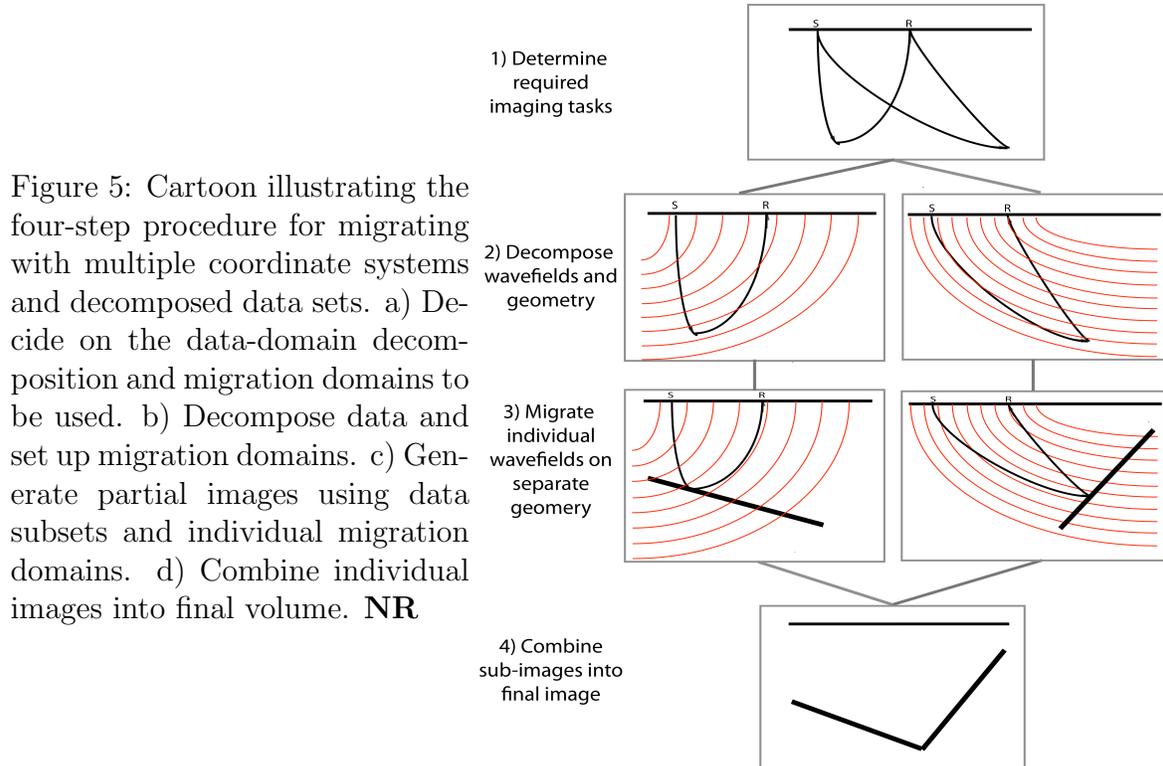


Figure 5: Cartoon illustrating the four-step procedure for migrating with multiple coordinate systems and decomposed data sets. a) Decide on the data-domain decomposition and migration domains to be used. b) Decompose data and set up migration domains. c) Generate partial images using data subsets and individual migration domains. d) Combine individual images into final volume. **NR**

wavefield techniques that inherently produce multi-pathing and other band-limited properties. Beams also can be propagated to steep and overturning angles. Some disadvantages include leaving model space shadow zones, not handling diffractions from sharp velocity model features due to localized beam domains, and introducing beam superposition artifacts such as beam boundary effects.

Another successful approach is plane-wave migration in tilted coordinates (?). This method exploits the fact that full 3D data volumes can be synthesized into plane-wave sections. The resulting number of plane-wave sections for migration is significantly fewer than the corresponding number of shot profiles, leading to improved computational efficiency. Because a plane wave is defined by a single take-off ray parameter, Cartesian meshes easily can be rotated to that orientation to be more conformal to the propagation direction. Individual plane-wave migration images are computed separately on the different rotated Cartesian meshes, and interpolated/stacked to form the final Cartesian image volume. Advantages of this approach include accurate large-angle propagation and a significant reduction in the number of required migrations. Some disadvantages are that the aperture for 3D plane-wave migration can be substantial and impose a significant memory burden, and that the image quality degrades as plane-wave sampling becomes increasingly sparse. An additional concern is that common acquisition geometries do not lend themselves well to this approach - in particular where the number of acquired sail lines exceeds the number of plane waves required in the cross-line direction to achieve a non-aliased image.

RIEMANNIAN WAVEFIELD EXTRAPOLATION

Riemannian wavefield extrapolation (RWE) is another method for propagating wavefields on generalized coordinate meshes (?). As before, the key idea is to globally transform the computational domain from Cartesian to a geometry where the extrapolation axis conforms to the bulk wavefield propagation direction. The main difference in RWE, though, is that the transformation is global and not confined to narrow propagation domains like beams. Changing the migration geometry requires properly formulating one-way wave equations in the transformed domain. ? and ? demonstrate how this can be done by writing the governing 3D Helmholtz equation on general Riemannian manifolds (?), and generating corresponding extrapolation wavenumbers using conventional one-way wave-equation approximations (?).

An instructive RWE example is generating a 2D point-source Green’s function using a coordinate system formed by a suite of rays traced through a velocity model. Figure 6 presents a RWE-generated Green’s function example. The upper (lower) panels represent the velocity (image) domains, while the left (right) panels show the Cartesian (Riemannian) domains. The first procedural step is to generate the smooth ray-coordinate mesh (Figure 6a) using Huygens’ wavefront tracing (?) on a smoothed BP synthetic velocity model (?). Figure 6b shows the mesh transformed into the ray-coordinate domain overlying the interpolated velocity model. The extrapolation axis in this reference frame is parameterized by travel time along a ray, while the orthogonal axis is shooting angle.

Generating a point-source Green’s function in Cartesian coordinates requires introducing an impulsive wavefield at the first extrapolation step. The corresponding wavefield state in the first step of the ray-coordinate system is a plane wave, which is equivalent to outwardly extrapolating equal energy at all shooting angles. The wavefield is then propagated through the velocity model to generate the ray-coordinate wavefield (Figure 6c) at four different travel times, which are then interpolated back to Cartesian (Figure 6d) using the known and invertible mapping relationship between two grids. The resulting Cartesian wavefield has energy at steep, vertical and overturning angles, which illustrates the potential for RWE to improve upon the conventional limits of wide-angle and turning-wave propagation by one-way extrapolation.

Figures 2 show the RWE wavefield extrapolation improvements, relative to that in a Cartesian coordinate system and the two-way finite difference modeling benchmark. Note the improved accuracy at large propagation angles of ray-based RWE extrapolation (panel 2c), especially within the right-hand salt body, relative to Cartesian extrapolation (panel 2b).

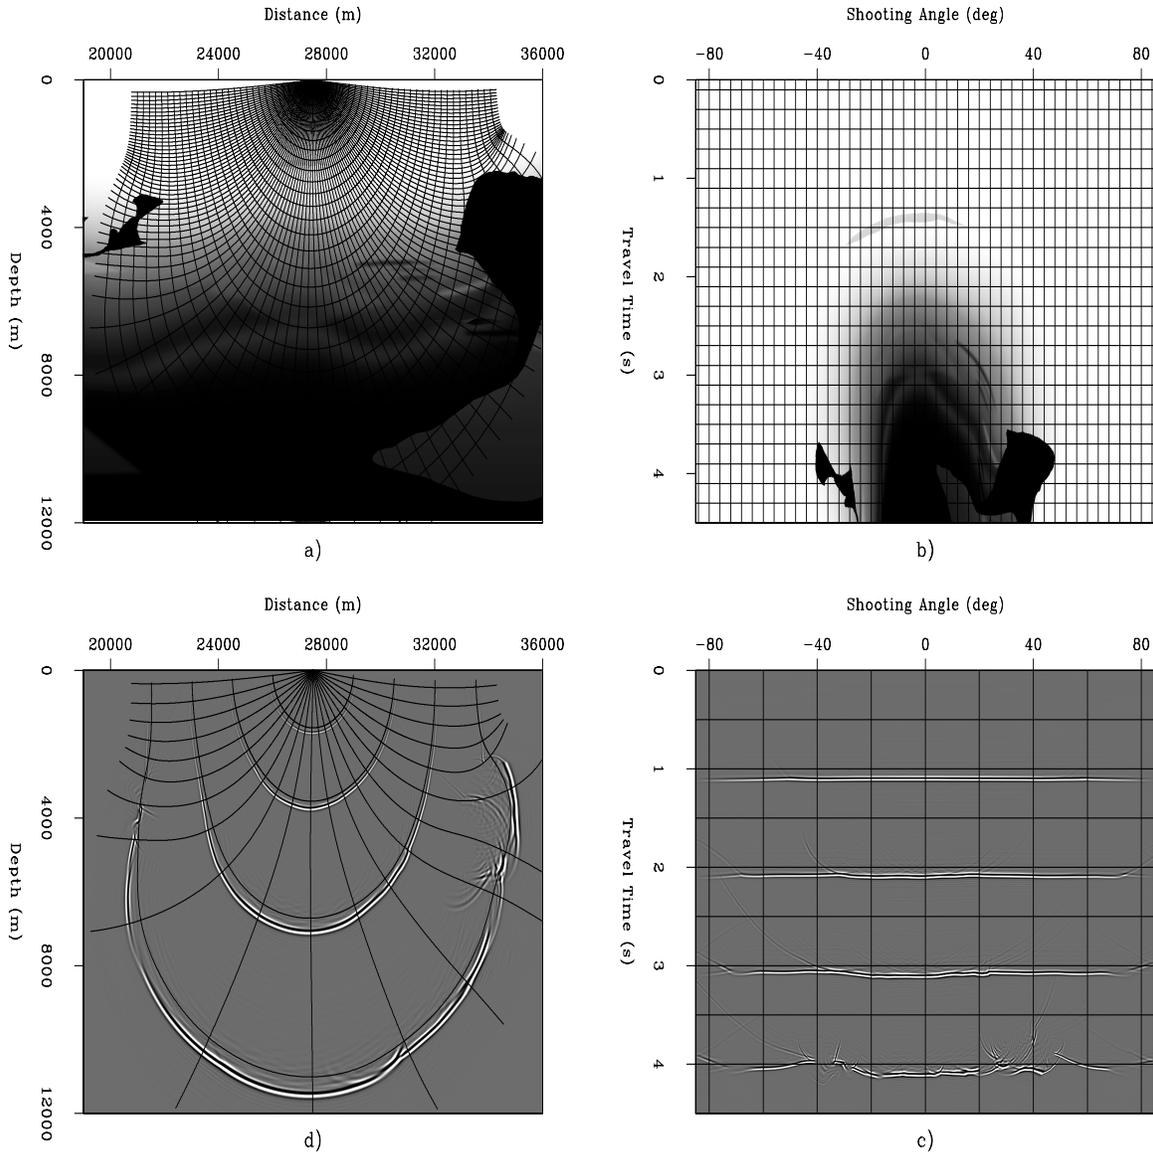


Figure 6: Illustration of the RWE approach using Green's functions calculated on ray-coordinate meshes through the BP velocity model. a) Velocity model with a smooth coordinate mesh overlain. b) Velocity model in a) interpolated into ray coordinates. c) Point-source Green's function in ray coordinates at four time steps. d) Wavefield in c) interpolated to Cartesian. Note the step, vertical and overturning waves, illustrating RWE's imaging potential. **ER**

Challenges with existing RWE implementations

Although the RWE naturally adapts to propagation in 2D ray-coordinate systems, the approach described by ? has numerous numerical implementation and conceptual challenges that need to be resolved before RWE can be applied successfully in 3D prestack migration scenarios. One major implementation issue is how to handle ray-coordinate triplications. Standard ray theory predicts infinite amplitudes in the limit where distances between neighboring rays goes to zero. Similarly, RWE generates unstable amplitudes at triplications because the formulation effectively normalizes amplitudes by a measure related to the inter-ray distance. A related issue is that significant ray-coordinate bunching or spreading can occur, even where meshes are triplication-free, which generate spurious grid reflection noise that degrades wavefield extrapolation quality.

A second implementation issue is that meshes generated through 2D ray-tracing are orthogonal grids, because the extrapolation direction is always orthogonal to the other (shooting-angle) axis. Similarly, meshes formed by 3D ray-tracing are limited to partial orthogonality because the extrapolation direction is orthogonal to the two other, not necessarily mutually orthogonal (shooting angle) axes. The assumption of (partial) orthogonality is unnecessarily restrictive as it precludes using smoother, non-orthogonal, triplication-free meshes.

Two additional conceptual challenges have made it difficult to apply RWE effectively and efficiently in more general prestack migration settings. (Herein, I will be assuming a shot-profile migration style unless otherwise specified.) First, the receiver wavefields used in shot-profile migration are usually broadband in plane-wave dip spectrum and cannot be easily represented by a single coordinate system. (That is, reflections from opposing dips propagate in opposing directions, as illustrated in Figure 4.) A second issue is that the coordinate systems optimal for point source and receiver wavefields seldom share a common geometry. For example, polar (spherical) coordinate systems are well-suited for propagating 2D (3D) point source wavefields, while elliptic (ellipsoidal) meshes are more appropriate for 2D (3D) receiver wavefields. This factor is detrimental to algorithmic efficiency where images are generated by correlating source and receiver wavefields: by existing on different grids they must both be interpolated to a common Cartesian reference frame prior to imaging. This leads to a significant number of interpolations, which renders the algorithm computationally unattractive, except in target-oriented imaging situations.

Establishing RWE-based migration as a viable seismic imaging technique will require resolving these implementation and conceptual challenges. Overcoming these problems represents the main contributions of this thesis.

THESIS CONTRIBUTIONS

The central goal of this thesis is to demonstrate that RWE-based migration is a viable 3D seismic imaging technique. A second goal is to prove that RWE-based approaches afford significant imaging improvements over conventional one-way extrapolation-based techniques at modest additional (and sometimes a reduced) computational cost.

The first major contribution is a new RWE formulation, more general than that of ?, that opens up new approaches for coordinate system design including non-orthogonal meshes. This extension leads to a more explicit connection of coordinate geometry in one-way wave equations, and helps define analytical extrapolation wavenumbers that improve the accuracy of RWE operator implementations. I also explore more wavefield-centric coordinate design approaches derived from ray-tracing algorithms. Propagation on these coordinate systems leads to extrapolated wavefields more accurate than those calculated in Cartesian coordinates, though of lower accuracy than analytic coordinate approaches. Both the analytic and ray-derived mesh generation techniques discussed herein avoid the problems associated with triplicating coordinate systems discussed in ?. Overall, I argue that analytic coordinates represent a more optimal trade-off between the competing constraints of extrapolation axes conforming to wavefield propagation directions, and the numerical accuracy and computational efficiency of the extrapolation operator implementation. This assertion is illustrated in Figure 2d, which shows the potential accuracy increases afforded by wavefield extrapolation in analytic coordinates relative to Cartesian and ray-traced coordinates in panels 2b and 2c, respectively.

The second contribution is extending the RWE approach to prestack migration, which was rendered conceptually challenging by the issues discussed in the above section. For 2D examples I show that elliptic coordinate systems have useful geometric properties, and are an appropriate geometry for propagating both the source and receiver wavefields. I demonstrate that elliptic-coordinate migration is an example where the trade-off between the competing constraints of large-angle accuracy, ease of numerical implementation, and computational cost overhead is excellent, if not optimal. Tests on the BP velocity synthetic data set demonstrate that elliptic-coordinate migration results offer significant imaging improvements over conventional Cartesian WE migration algorithms.

Next, I examine whether angle-domain common-image gather (ADCIG) theory remains valid in generalized 2D coordinate systems. I demonstrate that ADCIGs can be calculated directly using Fourier-based methods for a particular class of coordinate system that include elliptic meshes. I show that computing ADCIGs in elliptic coordinates offers imaging advantages over doing so in Cartesian grids. In particular, I argue that the spatially varying extrapolation axis leads to more accurate large-angle propagation while minimizing the insensitivity of the ADCIG calculation to steep structural dips commonly observed in conventional implementations.

Finally, I examine what combination of coordinate system geometry and wavefield

decomposition provides an optimal match for 3D prestack migration. I argue that an inline delayed-shot migration strategy is an effective strategy for situations where the sources are well-sampled inline, but have a limited number of sail lines or the sail-line sampling is too coarse. Using line sources, though, leads to impulse responses with more conical-like geometries that remain more linear in one direction, and more cylindrical in the other. Extending the analytic approach above to 3D, I detail a RWE migration strategy for performing inline delayed-shot migration using tilted elliptical-cylindrical (TEC) meshes that conform fairly well to the shape of a general linear-source impulse response. This approach retains the efficiency of plane-wave migration, while affording the migration of most steep-dip and turning-wave components to all azimuths. I present wide-azimuth migration results to validate the theory for a wide-azimuth synthetic 3D data set computed from a realistic Gulf of Mexico geologic model. The imaging results indicate that migration in TEC geometry offers imaging improvements over Cartesian meshes, especially for steeply dipping geologic structures such as salt flanks, at a reduced computational cost. The approach is applied to a 3D narrow-azimuth Gulf of Mexico data set to demonstrate the imaging advantages in a field data test.

THESIS OVERVIEW

Chapter 2: RWE: Non-orthogonal coordinate systems - I demonstrate how the RWE approach can be extended to include modeling one-way wave propagation on generalized coordinate meshes. The RWE implementation of ? assumes that coordinate systems are defined by either orthogonal or semi-orthogonal geometry. This restriction leads to situations where coordinate meshes suffer from problematic bunching and singularities. I develop a procedure for avoiding many of these problems by posing wavefield extrapolation on smooth, generally non-orthogonal, but singularity-free, coordinate meshes. The resulting extrapolation operators include additional terms that describe non-orthogonal propagation effects. These extra degrees of complexity, however, are offset by smoother coefficients that are more accurately implemented in one-way extrapolation operators. I validate my theory of non-orthogonal propagation with two analytic coordinate system examples, and present a method for eliminating any remaining singularities from coordinate systems. I demonstrate non-orthogonal RWE through numerical calculation of 2D and 3D Green's functions for cylindrical and near-spherical geometry. Results from 2D benchmark testing suggest that the computational overhead associated a mixed space- and Fourier-domain RWE implementation is roughly 35% greater than Cartesian-based extrapolation. However, I show that the computational overhead in analytic coordinate systems, even in 3D applications, is likely to be less than 6% greater than the corresponding cost in Cartesian coordinates. Results from this chapter have been published as ?.

Chapter 3: Shot-profile migration in elliptic coordinates - I extend the Riemannian wavefield extrapolation (RWE) formulation of Chapter 2 to 2D prestack migration using analytically defined elliptic-coordinate systems. I show that the cor-

responding 2D elliptic extrapolation wavenumber introduces only an isotropic slowness model stretch to the single-square-root operator. This enables the use of existing Cartesian finite-difference extrapolators for propagating wavefields on elliptic meshes. A post-stack migration example illustrates the ability of elliptic coordinate migration to image with turning waves. A 2D imaging test using a velocity benchmark data set demonstrates that the RWE prestack migration algorithm generates high-quality migration images that are more accurate than those generated by Cartesian operators of the equivalent accuracy. I note that even in situations where RWE geometries are employed, a high-order implementation of the one-way extrapolator operator is required for accurate propagation and imaging. Results from this chapter have been published as ?.

Chapter 4: Generalized coordinate ADCIGs - Chapter 4 extends the theory of 2D angle-domain common-image gathers (ADCIGs) to migrations performed in generalized coordinate systems. I develop an expression linking the definition of reflection opening angle to various geometric factors. I demonstrate that generalized coordinate ADCIGs can be calculated directly using Fourier-based offset-to-angle approaches for locally isotropic coordinate systems. Tilted Cartesian, polar and elliptic coordinate examples are provided to help illustrate theory. I validate the ADCIG theory by comparing analytically and numerically generated image volume results for a set of elliptically shaped reflectors. Experiments with the BP velocity synthetic data set demonstrate that elliptic-coordinate ADCIGs better-resolve steeply dipping structure relative to Cartesian ADCIGs. Results from this chapter have been accepted for publication as ?.

Chapter 5: Inline delayed-shot migration in TEC coordinates - Chapter 5 extends the 2D analytic RWE approach of Chapter 3 to 3D coordinate systems. I show how to perform inline delay-shot migration in tilted elliptical-cylindrical (TEC) coordinate systems. The elliptic geometry is oriented in the cross-line direction, which naturally allows the inline oriented linear source to propagate to steep angles and overturn as necessary. When inline coordinate tilt angles are well-matched to the inline plane-wave ray parameters, the TEC coordinate extension affords accurate propagation of most steep-dip and turning-wave components of inline-source phase-encoded wavefields to all azimuths. I show that wavefield extrapolation in TEC coordinates is no more complicated than propagation in elliptically anisotropic media. Impulse response tests using 80° finite-difference operators illustrate the implementation's large-angle accuracy and lack of numerical anisotropy. I apply this approach to a 3D wide-azimuth synthetic and a 3D narrow-azimuth Gulf of Mexico data set to demonstrate the imaging advantages made possible through 3D RWE implementations.