Chapter 1

Introduction

Traditionally, geophysical and geological information have been used separately in seismic exploration: geophysicists process the raw seismic data and produce a seismic image and velocity model of the subsurface; geologists subsequently interpret the image—incorporating the velocity model and geological information—and identify target zones for oil exploration. If the velocity model and seismic image were unique, this separation of information would be appropriate. However, the limited amount of data available in a typical seismic reflection survey forces geophysicists to make assumptions in the processing stage. In areas with complex geology, these assumptions often introduce errors into the velocity model, resulting in errors in the seismic image. This, in turn, can lead to mistakes in the interpretation.

To avoid these mistakes, geological information must be used in the processing stage to help the geophysicist constrain the velocity model. In other words, the initial velocity model and its associated seismic image must no longer be considered as end products of seismic processing. Instead, they represent intermediary steps in an iterative process, in which a geological interpretation of the seismic image helps constrain and improve the velocity model. Likewise, well-log and other geological information must be included in the velocity estimation.

1.1 GOAL

In this thesis I describe a method for estimating seismic velocities that uses such an integration of processing and interpretation. More specifically, this method interprets events in depth-migrated constant-offset sections, and uses the result of this interpretation to update the velocity model. The method is useful in areas with complex geology, especially in regions where the geology is dominated by salt intrusions. These regions are important for oil exploration because the impermeable salt boundaries often serve as traps below which hydrocarbons accumulate.

For example, Figure 1.1 shows a salt structure imaged by time migration, a standard migration method that assumes laterally invariant velocities (although time migration is often successfully applied in regions where velocities are weakly varying in the lateral direction). An initial interpretation of the lower left image below the salt structure reveals sediments bending upward and truncating against the salt boundary, a potentially interesting environment for oil exploration. The preliminary interpretation seems to indicate that the migration has imaged the salt and sediments well: the steeply dipping salt boundary coincides more or less with the truncation of the sediments, a sign that the migration velocity is approximately correct.

However, the laterally invariant velocity model used for migrating the data is likely to be inaccurate, because salt velocity is normally much higher than sediment velocity. Therefore, the migrated image may very well be wrong, and the above interpretation based on that image may draw the wrong conclusions. A puzzling feature in the image, which may indicate possible imaging errors, is the weak, steeply dipping reflector imaged inside the salt body. Is this feature a migration artifact, or is it a genuine reflection event, and if so, what causes it?

The goal of the velocity-estimation method presented in this thesis is to answer these questions by finding a velocity model that accurately images the salt boundaries and the surrounding sediments. Before I describe the method in more detail, I discuss some of the problems with seismic imaging in salt regions and the reasons why conventional velocity-analysis techniques often fail in these areas.

1.2 PROBLEMS WITH IMAGING SALT STRUCTURES

Although many migration algorithms are now capable of accurately imaging seismic data, imaging under salt structures remains a difficult task. Complications in the imaging arise for several reasons. First, the high velocity of salt structures causes strong velocity variations, both vertically and laterally. Moveout in reflection events is therefore nonhyperbolic, and, as I discussed above, time migration does not produce satisfactory results. Instead,

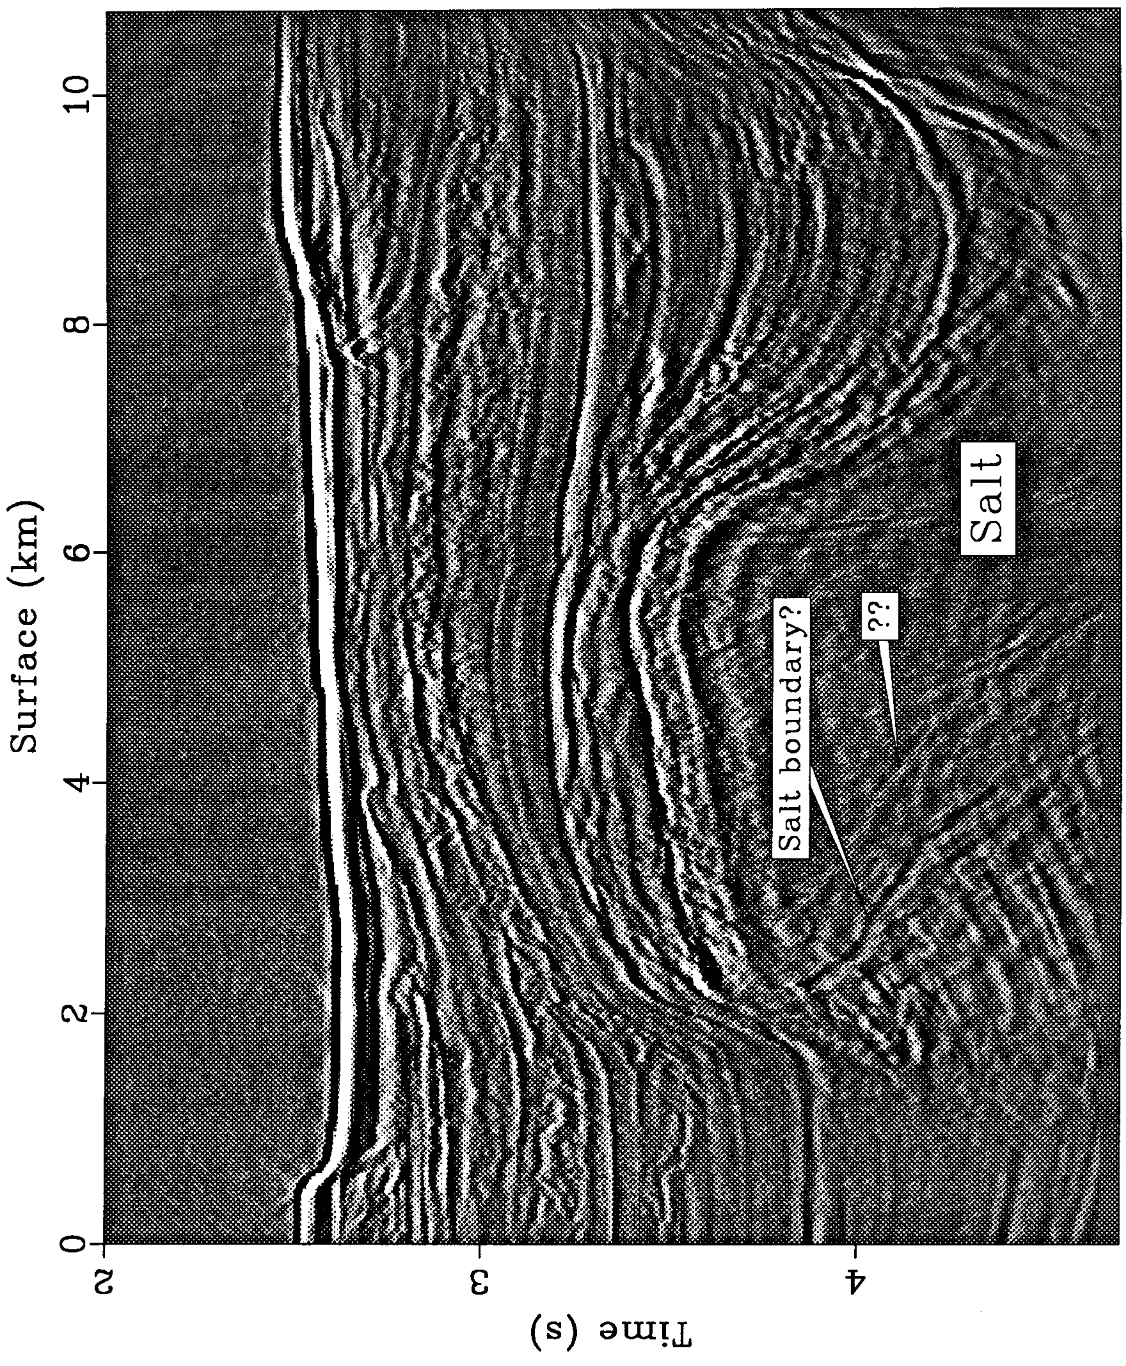

FIG. 1.1. Time-migrated image of a data set recorded above a salt structure in the Gulf of Mexico Data courtesy of Arco Oil and Gas Co. (See section 6.1 for a more detailed description of the data.)

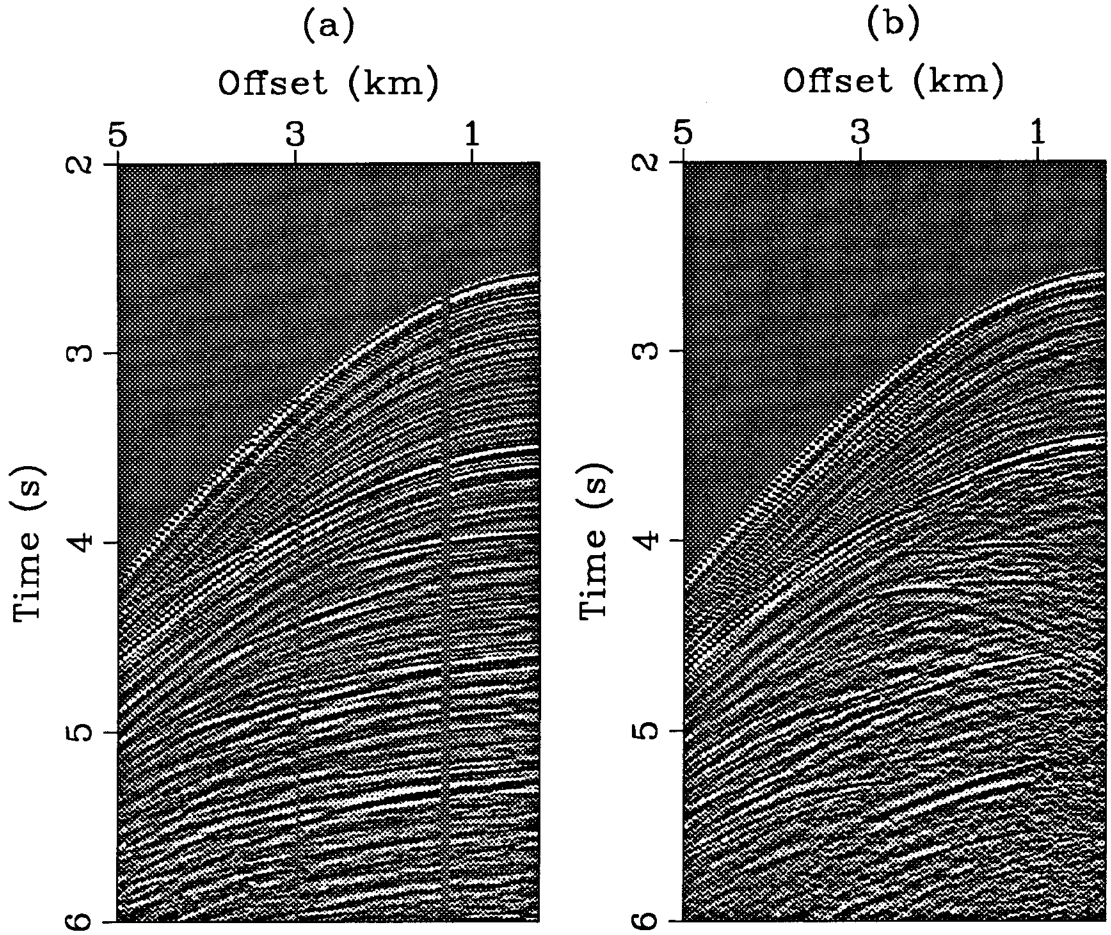

FIG. 1.2. Data gathers from the data set shown in Figure 1.1: (a) CMP gather at a midpoint of 4.8 km; (b) Shot gather at a shot location of 4.8 km. The CMP gather clearly shows the refracted wave off the top of the salt. This refracted wave is not present in the shot gather because large offsets in that gather were recorded left of the salt structure. The data are gained by AGC.

depth migration is needed, not only for focusing the data, but also for correctly locating salt boundaries (Larner, 1987). Second, because their density is lower than that of surrounding sediments, salt bodies tend to move upwards, deforming and faulting the sediments as they intrude. This salt movement can lead to complex structural settings and steeply dipping beds, all of which are hard to image. A related problem is poor data quality: the intrusive salt flow disrupts surrounding sediments, and these fractured sediments may not generate clear reflections. Third, because of the high reflection coefficient of salt-sediment boundaries, little seismic energy penetrates the salt structure, and reflections from sediments below the salt are weak.

Figure 1.2 illustrates some of these problems; the figure shows a common-midpoint (CMP) and a shot gather from the data set in Figure 1.1. The gathers are located directly above the salt structure. The salt reflection is clearly visible, starting at 3.6 s at near-offset. Below the salt, reflections are nonhyperbolic and fragmented; events overlap with conflicting dips. Chapter 2 gives a more detailed description of the different events.

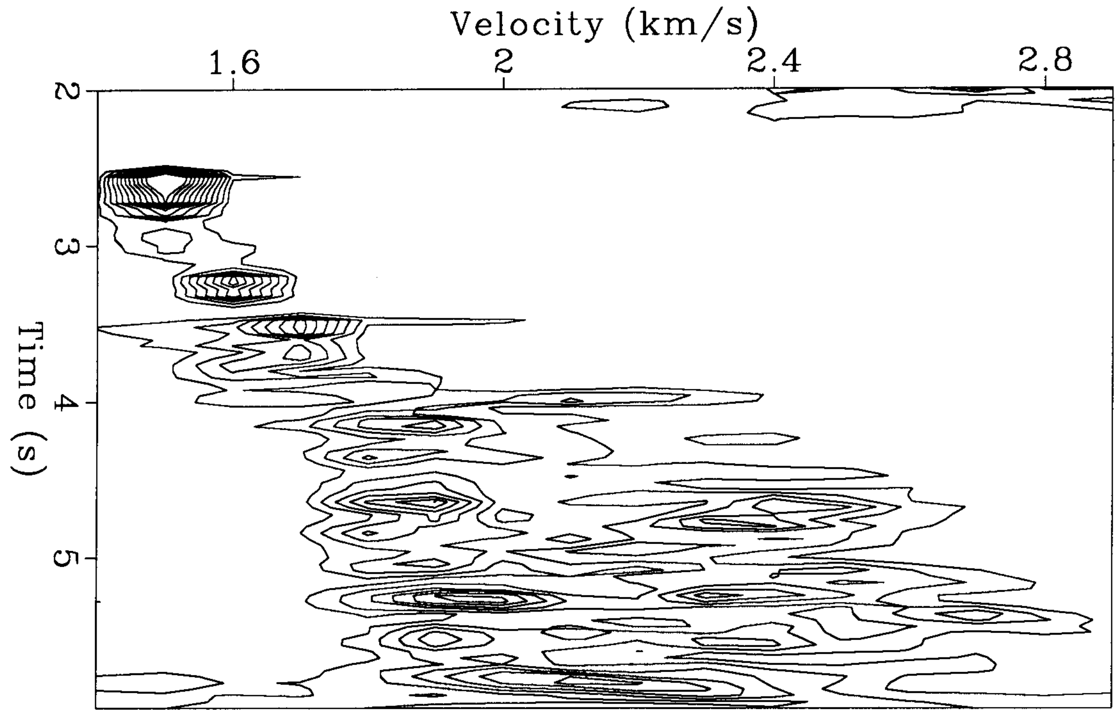

FIG. 1.3. NMO velocity analysis of the CMP gather in Figure 1.2a.

Associated with the imaging problem is the problem of estimating seismic velocities in salt regions. Figure 1.3 shows an example of conventional stacking or NMO velocity analysis (Taner and Koehler, 1969). The analysis does well in the sediment sequence above the salt, where the plot reveals a well-defined velocity trend, but it fails drastically below the salt structure. The reason for this failure is the inability of stacking-velocity analysis to handle lateral velocity variation and complex structure.

Toldi (1985) addresses the first problem, and extends NMO analysis to deal with lateral variations in velocity. More sophisticated methods of migration-velocity analysis (Faye and Jeannot, 1986; Al-Yahya, 1987; Fowler, 1988) solve the problems with structure, especially if they use depth migration (Etgen, 1990; Cox et al., 1988). However, most of these methods use a stacking measure in the migrated data to estimate velocities, and may therefore lack the resolution to detect strong lateral velocity variations. Local stacking methods that use unmigrated data (Sword, 1987; Biondi, 1990), on the other hand, face just the opposite problem: they resolve local variations in velocity well, but are more sensitive to data quality and have difficulties in dealing with some aspects of the unmigrated data caused by complex structure (such as crossing events and triplications).

Tomographic methods (Dines and Lytle, 1979; Bishop et al., 1985; Stork and Clayton, 1986; Nolet, 1987) are velocity-analysis methods that avoid stacking altogether. In seismic tomography, traveltimes recorded in a seismic experiment are picked and compared with traveltimes computed for an assumed velocity model; the differences are used to calculate an update to the model. The main drawback of tomography is that picking prestack data is cumbersome and difficult if data quality is poor. (Picking events is virtually impossible in the lower part of the gathers in Figure 1.2.) However, tomographic methods have high resolution in velocity, and can handle strong lateral velocity variations. Also, they have the advantage that they are event-driven, and therefore well-suited for combination with an interactive interpretation in structurally complex regions with few strong reflectors.

1.3 TOMOGRAPHY AFTER DEPTH MIGRATION

The velocity-estimation method presented in this thesis uses an interactive, event-driven approach, but analyzes constant-offset sections after they have been depth migrated. Because these constant-offset sections all image the same subsurface area, they should be identical after migration if the correct velocity was used (Gardner et al., 1974). This observation serves as the basis for the velocity-estimation method: the method is an iterative optimization scheme that tries to find a velocity model for which all reflectors are imaged at identical positions on different constant-offset sections. For this purpose, I pick reflectors in the migrated sections, and from these picks I estimate perturbations in reflector depth as a function of offset. The optimization then determines how the velocity model must change so that the perturbations are minimized. The minimization is done through a conjugate-gradient algorithm, where the calculation of the gradient (or backprojection) operator is based on a linearization of the relation between the positions of the migrated reflectors and the velocity model.

As I show in Chapter 2, the migrated constant-offset sections are better suited for interpretation than the unmigrated data gathers are. First, they closely resemble the geology, so geological constraints can guide the interpretation. Second, seismic energy is focused by the migration, so that the local signal-to-noise ratio in the data is improved. Migration backs out the effects on the data of wave propagation through the overlying velocity structure; the residual effects that are left after migration are clearer and easier to analyze. Thus, migration greatly simplifies the picking of events in the data, and thereby solves one of the main problems in tomography.

1.4 COMPARISON WITH TRAVELTIME TOMOGRAPHY

Although the velocity-estimation method is similar to traveltime tomography—both use picked events and project perturbations in these events back onto the velocity model—there is one fundamental difference: in traveltime tomography the events are picked in the unmigrated data, whereas here the events are determined after the data have been migrated. I have already discussed the merits of using migrated data in the interpretation. However, there is another, equally important, reason why migration is useful in the velocity analysis.

In tomography, traveltime observations are independent of the velocity model, whereas here the depths found after migration depend on the velocity model. It is this dependency, and the fact that migration is supposed to produce identical images at all offsets, that provide the ingredients for an efficient optimization scheme. In other words, unlike tomography, my optimization does not minimize the mismatch between modeled and observed events, but instead uses the unique relation between the migrated reflectors and the velocity model to check the validity of that velocity model.

By using migrated data, I address one of the major flaws in reflection tomography: the assumption that reflector positions are known and can be held fixed in the computation of the backprojection operator. This assumption underlies the application of Fermat's principle, which assumes fixed ray endpoints in the inversion. In transmission tomography, where ray endpoints are known, this approach works well. However, in reflection tomography—where reflector positions are unknown and often incorrectly guessed—this approach may lead to errors in the velocity estimation. Even when tomographic methods estimate reflector positions simultaneously with velocity in an iterative scheme, traveltimes are still modeled for inaccurate reflector positions at each iteration.

Tomography after migration does not depend on a linearization with respect to an *assumed* velocity and structural model. Rather, it obtains reflector positions after migration, and, in the calculation of the backprojection operator, takes into account the positions' dependency on the ***known*** migration-velocity model. The linearization of the inverse problem around the known migration-velocity model is well determined, so that an efficient optimization scheme results.

Some other tomographic methods also use migration to obtain reflector positions (Stork, 1988; Verprat et al., 1988). However, these methods are still based on the comparison of modeled traveltimes with the observed ones, and still compute the backprojection operator for fixed, incorrect reflector positions. Thus, they do not take reflector movement into account. (In fact, the approach of these methods is paradoxical: if one migrates the data with an initial velocity model, and then models reflection events for the migrated horizons with the same velocity model, the traveltimes of the events should be identical to the ones in the unmigrated data. Therefore, traveltime perturbations should be zero; the only reason they are not is that the horizons are normally determined from a stacked migrated image, which cannot reflect the exact offset behavior of the events.)

To summarize, migration is important not only as a tool that aids the interpretation, but also in providing a consistent framework for finding reflector positions and verifying seismic velocities.

1.5 OVERVIEW

In Chapter 2 I show the structural interpretation of the migrated data. Residual moveout in events identified in the interpretation is used as input to the velocity-estimation method.

Chapter 3 discusses a new method for calculating seismic traveltimes by finite differences, which is at the computational core of the velocity-estimation method. The traveltimes are used in the migration of both the raw data and the horizons, and in the velocity-estimation algorithm itself. The finite-difference calculations are more efficient than ray tracing, the traditional method for calculating traveltimes in tomography. They are also better suited for structural velocity models, in which rays may behave erratically because of strong contrasts in velocities.

Chapter 4 is devoted to residual event migration, the forward operator that describes how reflectors in the migrated data move as a function of velocity. The residual migration handles horizons in prestack data, in contrast to conventional map-migration methods that only migrate events in zero-offset or stacked sections. I discuss both nonlinear and linear residual-event-migration operators.

In Chapter 5 I describe the inverse process of estimating velocities from residual moveout in the horizons, in which I use the linear residual-event-migration operator from Chapter 4.

Finally, Chapter 6 shows a field data example, where I estimate seismic velocities for the salt data set of Figure 1.1.