Chapter 1

Introduction

> If Anisotropy is such a Big Deal, why has NMO worked as the basis for time-to-depth conversions for 50 years? – F. Muir, 1985

A medium is anisotropic if its properties depend on direction. To a philosopher, anisotropy is a fundamental property of a medium that says something deep about how it is put together. This basic symmetry property of the medium can then cause observable directional effects as a secondary phenomenon.

In this thesis I will consider only seismic anisotropy, meaning that sound waves propagate differently depending on direction. Like the applied physicist who does not need a grand unified theory of gravity in order to measure the mass of a beam, I will not be overly concerned as to exactly *how* anisotropy happens. It could be due to regular patterns of different kinds of atoms, irregular layers of sandstone and shale, or something else entirely. If a medium looks anisotropic to the waves I send through it, I will say that that medium is anisotropic.

Is the complex and impure Earth really an "anisotropic medium", or is it just a mess?

1.1 Why anisotropy

Most geophysicists have always ignored anisotropy, and they have gotten away with doing so for a long time. If anisotropy does commonly occur in the Earth, why have geophysicists been able to ignore it so successfully?

1.1.1 Does anisotropy exist in the Earth?

The answer to this question is subtle and significant. Anisotropy has been observed, even extremely strong anisotropy.

Historical observations of anisotropy

"Common knowledge" as far back as the 1930's was that P-wave velocities determined from refraction surveys were typically 10-20% higher than those determined from reflection surveys in the same area. For example, Levin (1978) reports that an unpublished study by McCollum and Snell in 1932 found horizontal P-wave velocities 40% faster than vertical ones in the Lorraine shales in Canada.

Such observations of anisotropy in exploration Geophysics became more respectable in the 1950's following theoretical work by researchers such as Postma (1955), who pointed out that even a completely isotropic layered Earth could *appear* anisotropic if the layering were on a finer scale than the wavelengths of seismic waves. (Rock outcrops show this should often be the case.) Around this time reports of anisotropy also started appearing more often in the exploration literature. For one notable example, Jolly (1956) reported finding horizontal SH waves traveling twice as fast as vertical ones in near-surface sediments.

Meanwhile, evidence for "intrinsic" anisotropy came from laboratory studies of rock samples. Nur (1969) found that anisotropy could be created in rocks by subjecting them to the sorts of pressures found in the Earth. Bachman (1979) found transverse isotropy in cores from the deep-sea drilling project. Jones and Wang (1981) found strong transverse isotropy in cores from the Williston basin in North Dakota.

Why has NMO worked?

This still leaves the question of how anisotropy has been successfully ignored. The answer lies in the sorts of reflection surveys that were typically done. Before the early 1980's only P-wave surveys with relatively narrow offsets were regularly recorded. Observations show that at small offsets P-wave NMO[1] velocities are usually quite accurate, even in the presence of strong anisotropy. (I.e. the reflector depths calculated using the recorded arrival times and best-fitting NMO velocities will be close to the actual depths found by drilling. For a field data example showing the relative (in)accuracies of P and SH-wave NMO and how they varied with lithology, see Winterstein (1986).) If this narrow-offset P-wave insensitivity to anisotropy is generally true, then the traditional narrow-offset P-wave recording methods and the standard isotropic imaging operations of P-wave NMO and Stack will have usually worked, despite any anisotropy present (except for the wide-angle near surface, which tends to get ignored anyway). Theoretical results show that such narrow-offset P-wave insensitivity to anisotropy would result if the earth predominantly consisted of thin isotropic layers all with about the same Vp/Vs velocity ratios (Helbig, 1984).

1.1.2 Is anisotropy worth trying to measure?

If P waves usually don't exhibit noticeable anisotropic effects, then how about shear waves?

Shear waves

Until the 1980's, when exploration Geophysicists recorded shear-wave data at all they would typically record only an "SH" (Yy) section, and if they were ambitious perhaps an "SV" (Xx) section as well. Such sections were usually of markedly poorer quality than P-wave data recorded in the same location, giving S-wave exploration in general a poor reputation. Starting in the mid-1980's, however, more attention began to be directed towards the theoretically predicted benefits of recording multi-component data, which renewed interest in trying to record usable shear-wave data.

A particularly enlightening observation was presented by Alford (1986), who showed an early multi-component survey that strikingly illustrated one reason why "SH" and "SV" sections are often of such variable and low quality. His dataset suffered from significant

[1] Note for non-Geophysical readers: NMO stands for "Normal MoveOut". The Encyclopedic Dictionary of Exploration Geophysics (Sheriff, 1984) defines the term as: "The additional time required for energy to travel from a source to a flat, reflecting bed and back to a geophone at some distance from the source point compared with the time to return to a geophone at the source point." Notice the implicit assumption in this definition of an isotropic layer, a flat reflector at the bottom, and sources and receivers on a flat surface at the top. If the variation of traveltime with distance fits this model the moveout is "normal", and the velocity of the isotropic layer is the "NMO velocity".

azimuthal anisotropy. As a result, both split shear waves were present on each of the standard "SH" and "SV" sections, and the two incoherent signals summed together to make a confusing hash out of the reflectors. Using the cross-terms (SH source into SV receiver (Yx) and SV source into SH receiver (Xy)) previously almost never recorded he was able to perform a coordinate-system rotation to minimize the confusion between the two split shear waves. The result was two new orthogonal shear-wave sections of dramatically higher quality. Willis, Rethford and Bielanski (1986) showed that such scrambled shear-wave effects were the rule rather than the exception at a variety of field locations in Texas, New Mexico, and California, strong evidence that significant azimuthal anisotropy is a commonplace phenomenon.

NMO velocities calculated from shear-wave data have also proven much less reliable for time-to-depth conversion. Winterstein (1986) found that for a given geological layer SH-wave stacking-velocity discrepancies tended to be much more variable with surface location than the corresponding P-wave data.

The silver lining

However, he found that the shear-wave stacking-velocity errors were *strongly correlated with lithology*: for shales he found that the SH-wave stacking velocity was significantly greater than the true vertical velocity, while for sandstones the two were in agreement. Similarly, other researchers have successfully used Alford's shear-section rotation method to estimate the extent and orientation of fractures in Austin Chalk (Mueller, 1990). The variability of shear waves is a liability, but it is also an opportunity.

1.2 Understanding anisotropy

To make use of these new possibilities, Geophysicists will need to develop an intuition for anisotropic propagation to extend and replace the basic intuition for isotropic propagation we have already developed.

1.2.1 Previous work

Of course, anisotropy has been studied by many researchers over more than a century. As early as 1837 McCullagh (1837) described the basic relationships between slowness surfaces and wavefront shapes. More recently, the properties of elastic waves in general anisotropic media have been described by authors such as Synge (1957), Helbig (1958), Federov (1968), Musgrave (1970) and Auld (1973). While these fundamental references do specify the mathematics of anisotropy completely, the great bulk of their analyses are confined to symmetry planes.

This is because the analytical methods they use are (for the most part) only tractable on the symmetry planes. Unfortunately, while symmetry planes are an understandable and mathematically tractable case, they are also a misleadingly special case. Some important anisotropic phenomena such as shear-wave singularities are inherently three-dimensional, and must be studied as such (Crampin and Yedlin, 1981). While the precise shapes of such complex three-dimensional features can be described mathematically, the analytical solutions are usually too messy to allow any easy comprehension.

1.2.2 Goal of the thesis

My goal, then, is to attack the subject in a new way. Whenever possible, I will rely on geometrical arguments in preference to mathematics. More importantly, I will verify my results using model examples computed *numerically* from basic principles.

The numerical models themselves will be calculated using two fundamentally different techniques. For the first method, I calculate theoretical impulse-response surfaces as the polar reciprocals of slowness surfaces. The slowness surfaces are calculated by solving the Christoffel equation (a classic eigenvalue-eigenvector problem) repeatedly for a large number of plane-wave propagation directions. The second method is just brute-force finite differences, which has the advantage of calculating the complete waveform (but is correspondingly more expensive to use).

In the course of the thesis this same basic philosophy will be applied again and again to several different problems in theoretical anisotropy. In the next section I briefly mention the major topics covered.

1.3 Summary of the thesis

I begin in Chapter 2 by examining elliptical anisotropy. Several authors have previously discussed elliptical anisotropy, such as Levin (1978), Helbig (1983), and Blair and Korringa (1987); I show how the properties of elliptical anisotropy these authors report are the natural consequences of a simple linear relation between isotropy and elliptical anisotropy.

With a slight generalization, the linear transform idea can be extended to encompass multiple layers with arbitrarily oriented elliptical anisotropy in each layer. The linear transformation is always kinematically exact, and is even dynamically exact in the case of pure SH waves propagating in a symmetry plane (Schoenberg and Costa, submitted).

Continuing Chapter 2, I present numerous examples bracketing the possible ranges of behavior of transverse isotropy. In the process I present formulas that can be used to predict the behavior of the impulse-response surface given the elastic constants. These results confirm the previous analytical work of Lyakhovitskiy (1984) and Helbig and Schoenberg (1987).

For a more applied transversely isotropic model example, I attempt to simulate a laboratory core-sample measurement by Vernik and Nur (1990). Vernik had noted some anomalies when performing the laboratory measurement; this computer model suggests why the anomalies occurred, and how the design of the experiment might be improved.

I conclude Chapter 2 by examining how the divergence and curl operators' wavetype-separation properties can be generalized to the case of arbitrary two-dimensional anisotropy. These anisotropic operators allow a vector wavefield to be collapsed onto two disjoint scalar pure-mode fields. I have primarily used this method to isolate one wavetype from a finite-difference model output; for example the qP wave can be extinguished before it hits the edges of the model to allow a more detailed study of the qS wave for the same size model grid. This method could also prove useful for two-dimensional vector-wavefield inversions, although I have not yet used it for that purpose.

Chapter 3 begins by attempting to extend the two-dimensional wavetype-separation algorithm to three dimensions. The three-dimensional algorithm works for qP waves, but runs into fundamental trouble for other wavetypes. The difficulties encountered provide a convenient platform for studying some of the peculiarities of three-dimensional anisotropic wave propagation. Foremost among these are shear-wave singularities. This interesting phenomena has been studied by several authors, notably Musgrave (1981) and Crampin and Yedlin (1981). I use geometrical arguments to show why singularities must always exist and to predict some of their properties. This investigation method in turn suggests a new classification nomenclature for singularities based on their topological properties.

To conclude Chapter 3 I construct several model examples for studying the significance of singularities. Topologically singularities prove to be slightly more complex than had been previously indicated in the literature; there are two kinds of "lid" and both may occur together. This is an esoteric point of unknown practical relevance to geophysicists, but some of the other three-dimensional finite-difference model results might be geophysically significant. In contradiction to the arguments of previous authors (for example (Garmany, 1989)), some of the effects of singularities prove to be of significant amplitude. In particular, one seemingly innocuous three-dimensional anisotropic effect called a “ring pinch” can unexpectedly provide a nearly ideal vector for coupling a vertical source into an SH receiver (a Zy section).

1.3.1 A sample application

The finite-difference modeling codes used in Chapters 2 and 3 prove useful for the study of wave-propagation effects in anisotropic three-dimensional media. Their primary selling point is that they are extremely accurate (Etgen and Dellinger, 1989) and model the complete vector wavefield. Unfortunately they are also computationally expensive (although they are quite cheap if the accuracy they deliver is taken into account). Often such completeness and accuracy is neither required nor wanted. For algorithms such as travel-time inversion we want to know the arrival time of the wavelet, not its precise amplitude or shape.

There are lots of methods already in use tailored for just such situations. One current favorite is ray tracing. While ray methods are efficient and can handle arbitrary three-dimensional inhomogeneities, they are quite tricky to code and debug and can become inaccurate if rays pass through or near caustics, shadow zones, or singularities (Gajewski and Pšenčík, 1987). Reflectivity methods are also popular, but these methods are basically limited to planar layered media and rapidly become expensive if the number of layers is large (Booth and Crampin, 1983).

The numerous examples in Chapters 2 and 3 of impulse-response surfaces overlayed on finite-difference wavefields suggest another alternative. The computationally inexpensive impulse-response-surface program appears to predict the position of the finite-difference wavefronts with great accuracy (see for example Figure 2.7 on page 32); why not use that? Unfortunately, that program only handles homogeneous media.

In Chapter 4, however, I present an extremely inexpensive approximate method based on the method of finite-difference traveltimes introduced by Van Trier and Symes (1991). At its core this method extends the homogeneous techniques used to calculate impulse-response surfaces in Chapters 2 and 3 to include heterogeneities.