Chapter 3

Anisotropy in 3 dimensions

> When we try to pick out anything by itself, we find it hitched to everything else in the universe. – John Muir

In the previous chapter I showed how waves in anisotropic two-dimensional media can display such phenomena as large amplitude variations, cusps, and anomalous polarizations. Although such things are certainly possible, we also know that geophysics got along pretty well assuming P waves were isotropic or at worst elliptically anisotropic for a long time. So we can hope that the sorts of media we are likely to encounter in the real Earth are usually not too far removed from isotropy. From the examples in the previous chapter, we might then feel justified in expecting that we can usually count on having waves that at least approximately fit the standard isotropic P, SV, and SH designations.

In this chapter we show that while the two-dimensional story is complete as far as it goes, it is not a good preview of what happens in three dimensions.

3.1 The problem with only looking at symmetry planes

We begin with a simple example of how two-dimensional intuition can be misleading. Refer to Figure 3.1, which shows three symmetry-plane slices through a simple orthorhombic anisotropic slowness surface. Each slice shows what seem to be pure P, SV, and SH modes;

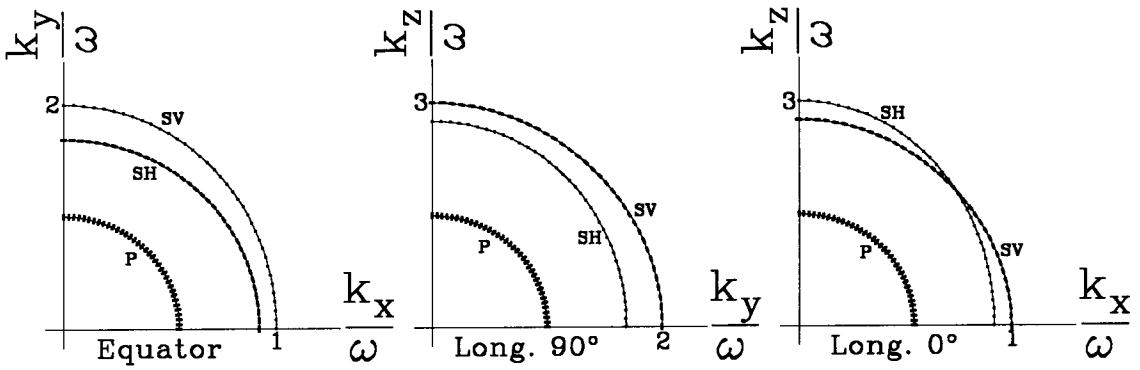

FIG. 3.1. Symmetry-plane slices through the three-dimensional slowness surfaces of a simple orthorhombic medium. The fat bars and dots show particle-motion direction. (Dots represent motion in and out of the plane.) Following the numbers from 1 through 3 seems to show the shear surfaces have a mobius-like topology. To see what's really happening, look at Figure 3.2.

furthermore, on each slice the P mode and one of the S modes are circular, while the other S mode is elliptical. Does this mean this medium supports pure P, SV, and SH modes?

No, as we can see by attempting to follow one shear surface around through all three plots. Start at "1" on the left plot where the "SV surface" intersects the (k_x/ω) axis. Trace from there up to "2". Now jump to "2" on the center plot (remembering that in three dimensions both "2"s mark the same point). Continuing on in the same way, we get to "3" and jump to the right plot. Here there is a surprise: we are now on the "SH surface", and the "SH surface" appears not to reconnect back to "1" where we started. What is going on?

Figure 3.2 shows a more accurate three-dimensional view. (Musgrave (1981) shows similar pictures for several physical media.) At "3", the outer surface is "SV" in the k_y–k_z plane, but "SH" in the k_x–k_z plane. This is possible because the designations "SH" and "SV" are relative concepts that depend on the orientation of the slice under consideration. Is the labeled "crossing" really a point where a "qSH" and "qSV" surface intersect? No, as we can see by closely examining the slices cut at 10° and 45° angles to the k_x–k_z plane. Clearly trying to shoehorn shear modes like these into "qSV and qSH" designations is not going to work in three dimensions. (Crampin (1981) also remarks on many of these points.)

Perhaps this example is atypically perverse. Do the standard isotropic shear modes harbor any similar misbehaviors? In Figure 3.3, note that by defining the shear surfaces as SH and SV we have introduced a nonphysical particle-motion direction discontinuity

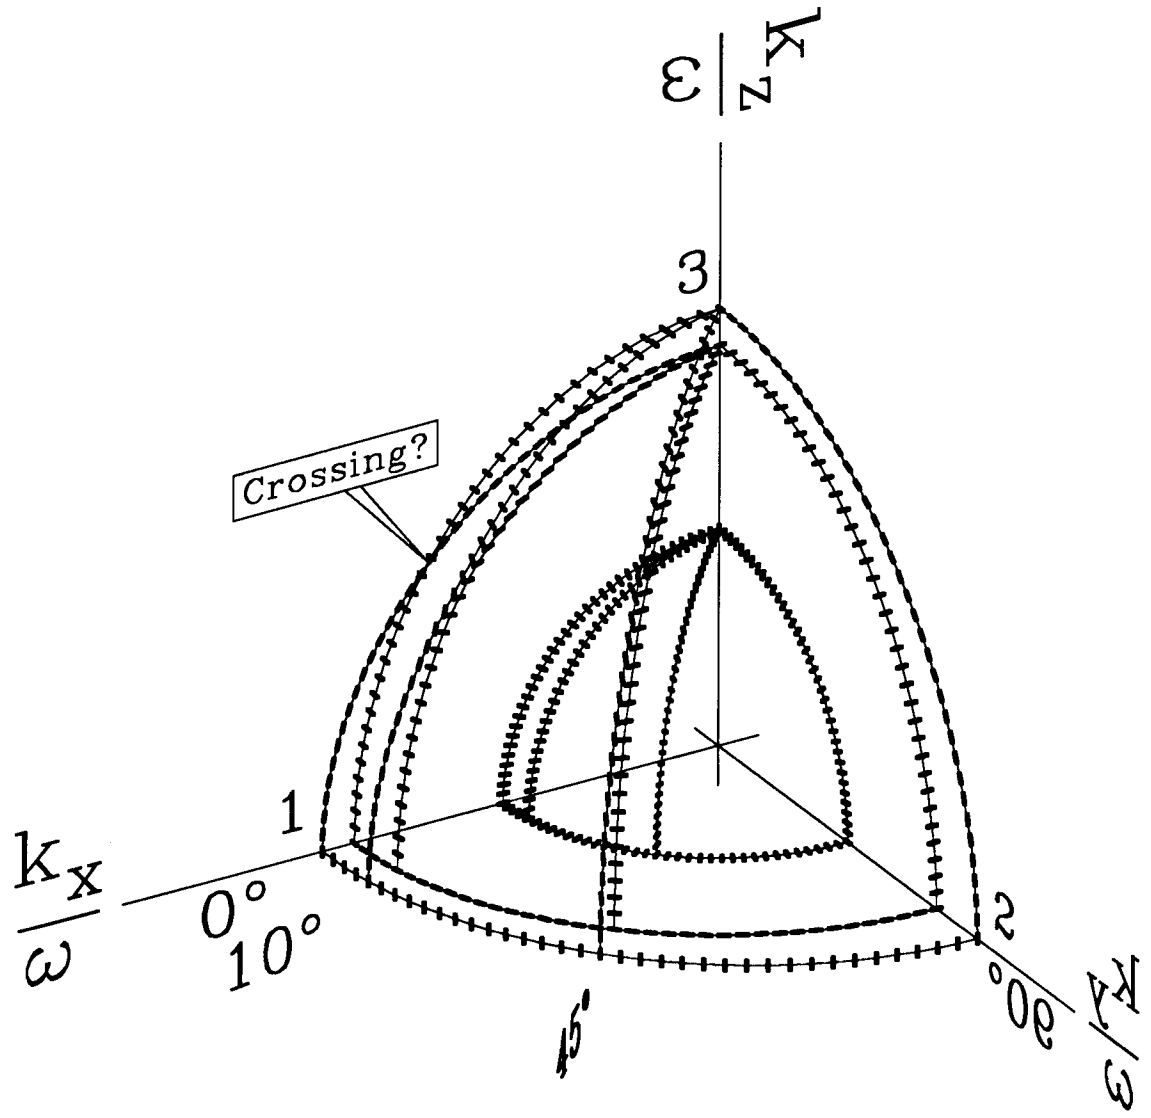

FIG. 3.2. Slices through the three-dimensional slowness surfaces of a simple orthorhombic medium. The fat bars show particle-motion direction. The "crossing" is not what it appeared to be in Figure 3.1. For another view of the same surface, look ahead to Figure 3.10.

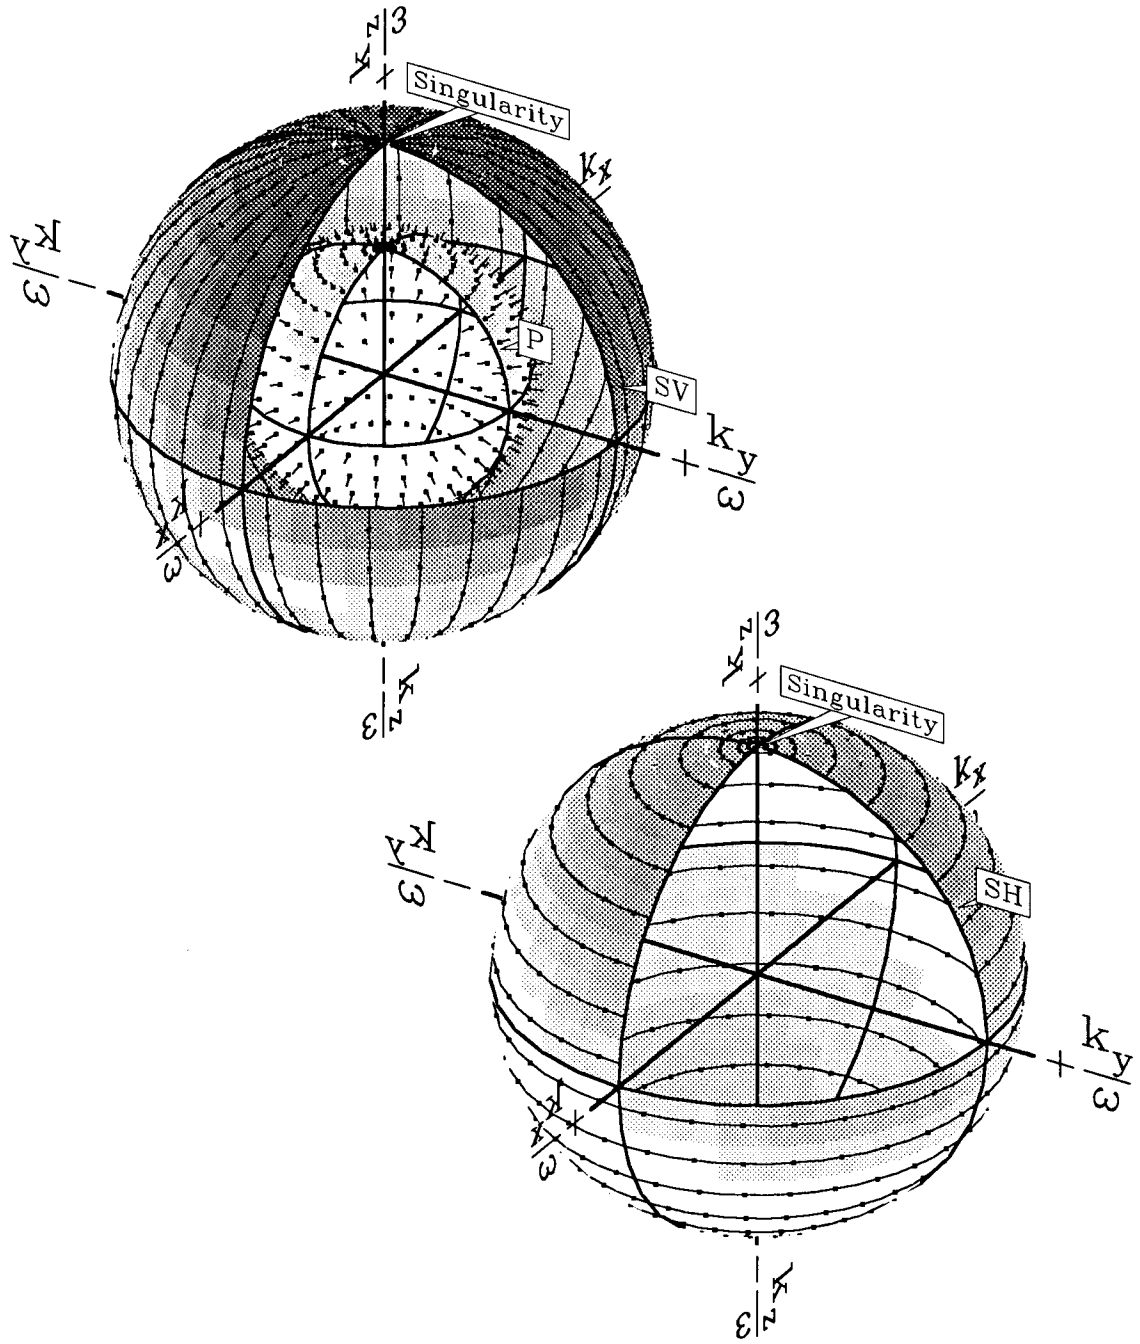

FIG. 3.3. Three-dimensional slowness surfaces for the standard {P,SV,SH} isotropic wave modes. Slowness surfaces are plots of phase slowness versus plane-wave propagation direction; they can also be considered as dispersion relation plots. We show particle-motion direction by the "sticks" attached to the surfaces. All three surfaces could not be plotted together because the two shear surfaces are everywhere coincident.

for vertically propagating waves. This is a consequence of the basic "in-plane" and "cross-plane" definitions of SV and SH; the same vertically traveling wave with North-South particle motion is called "SH" on an East-West section but "SV" on a North-South one.

3.2 Wavetype separation in three dimensions

Is it possible there is some other way of separating the degenerate isotropic shear mode into two well-defined orthonormal shear wavetypes, one that avoids the artificial singularity the standard "SV, SH" definition creates for vertically traveling waves? The answer is no, as we will discover in the next two sections in the course of attempting to extend the two-dimensional wavetype-separation algorithm from section 2.5 to three dimensions.

3.2.1 Extending 2D to 3D

The two-dimensional algorithm in section 2.5.1 seems to extend to three dimensions in a natural way. As in two dimensions, we use the solutions of the Christoffel equation (already presented as equation (2.4) on page 17) to find the propagation velocity and particle-motion direction for each plane-wave component in the Fourier domain. The problem is still a well behaved symmetric eigenvalue-eigenvector problem; we are still always guaranteed three orthogonal wavetypes for each $\mathbf{k}$. The only difference should be that there are three solutions now instead of two. Following this recipe, we should quickly be able to sketch out a three-dimensional wavefield-separation algorithm modeled after the two-dimensional version previously presented on page 53.

Remember back on page 17 I warned that things would get interesting later? This time around I will be a little more careful about how I describe the wavefield-separation algorithm; the whole point of this section is to show where it goes wrong in three dimensions.

First, construct the operator (I'll go over each step in more detail later):

(1.) Decide which mode this operator will pass.

(2.) For all (k_x, k_y, k_z):

{

(A.) Substitute (k_x, k_y, k_z) into equation (2.4) (on page 17) and find the eigenvector solutions $\mathbf{v}_1$, $\mathbf{v}_2$, and $\mathbf{v}_3$.

(B.) Decide which solution corresponds to the desired mode chosen in step (1). Call this solution $\mathbf{v}$.

(C.) Both $\mathbf{v}$ and $-\mathbf{v}$ are equally valid normalized eigenvectors. Choose the one that is consistent with particle-motion directions already determined at adjacent (k_x, k_y, k_z) points.

(D.) Store the result for this (k_x, k_y, k_z) in $\mathbf{V}(k_x, k_y, k_z)$.

}

The vector field $\mathbf{V}$ is then used to calculate the scalar pure-mode component M of an input elastic wavefield $\mathbf{U}$ by the following:

(3.) Fourier transform $\mathbf{U}(x, y, z; t)$ over x, y, and z, obtaining $\hat{\mathbf{U}}(k_x, k_y, k_z; t)$.

(4.) For all (k_x, k_y, k_z):

{

$$\hat{M}(k_x, k_y, k_z; t) = \mathbf{V}(k_x, k_y, k_z) \cdot \hat{\mathbf{U}}(k_x, k_y, k_z; t)$$

}

(5.) Inverse Fourier transform $\hat{M}(k_x, k_y, k_z; t)$ over k_x, k_y, and k_z, obtaining $M(x, y, z; t)$.

We will collectively refer to this algorithm as a "wavetype-separation operator": this operator accepts as input a vector field $\mathbf{U}$ and outputs a pure-mode scalar component M.

3.2.2 Possible pitfalls

Where can things go wrong? The algorithm attempts to treat each plane-wave component (k_x, k_y, k_z) separately. This is misleading. The values of $\mathbf{V}$ chosen at each vector wavenumber (k_x, k_y, k_z) do not exist in isolation, but must together describe a single global wavetype. Furthermore, the global wavetype described must agree with the way waves actually propagate in the given medium.

Mathematically, given that the matrix of elastic constants $\underset{\sim}{\mathbf{C}}$ really does correspond to the medium under consideration, these requirements boil down to two kinds of Fourier-domain continuity: First, the vector field $\mathbf{V}$ must be a continuous function of plane-wave direction (k_x, k_y, k_z); a true Fourier-domain discontinuity in particle-motion direction

would imply there existed an infinite plane wave in the space domain. This would violate causality, among other things.[1] Second, the phase velocity (ω/k) must also be a continuous function of plane-wave direction (k_x, k_y, k_z), since a discontinuous phase velocity would represent a non-physical discontinuous wavefront.

Our naive algorithm can fail at several steps: At step (2A) there is no guarantee of three *unique* solutions. Steps (1) and (2B) presuppose we know what "the global pure modes" look like, and that they are well-defined and continuous. At step (2C) there is no guarantee that it is possible to choose a sign that is continuous with all the previous adjacent choices.

In section 2.5.2 I showed several successful applications of the two-dimensional version of this wavetype-separation algorithm. No special care was taken to avoid the possible difficulties outlined above; they simply didn't occur. Is there any reason to expect the same algorithm that worked in two dimensions should fail in three?

3.2.3 Why two dimensions worked

In two dimensions there are two in-plane wave modes. (There is also a single uncoupled out of the plane mode, if "2 dimensions" is construed to include particle motion out of the plane. Here I will use "2 dimensions" in the strict mathematical sense, in which case there is no such thing as "out of the plane".)

Except for certain anomalous cases (see section 2.3.3 page 30, or Helbig and Schoenberg (1987)), the two in-plane modes are clearly analogous to the P and SV modes in isotropic media. If this is the case, there are no problems; the qP wave is faster than the qSV wave for all propagation directions. Mathematically, since the two modes never have the same velocity (ω/k), they must have distinct eigenvalues ($\rho\omega^2$). Distinct eigenvalues have unique orthogonal eigenvectors, so there is no difficulty finding the particle-motion directions in step (2A). The phase-velocity inequality also trivially solves the problem of identifying which mode is which in step (2B); the faster one is qP, the slower one is qS. Finally, in step (2C), as long as the "qP" mode does not have pure S particle motion for any propagation direction, defining the "outward" particle-motion direction on the qP mode to be positive will not lead to an inconsistency. The right hand rule then gives a consistent sign convention for the qSV mode.

[1] Later on in Figure 3.33 we will see how the limited sort of Fourier-domain discontinuity that occurs at a singularity can cause a *finite* plane wave that is actually physical.

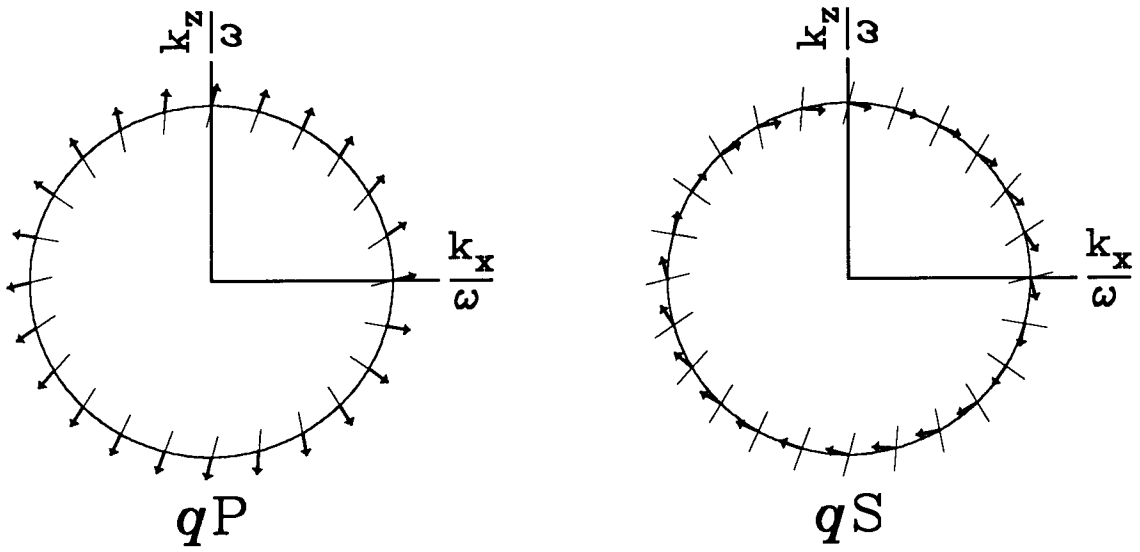

FIG. 3.4. Polar plots of particle-motion direction versus plane wave propagation direction for "normal" two-dimensional media. The shafts of the arrows indicate particle-motion direction. Each shaft has an arrowhead on one end, thus labeling that direction as the "positive" one. This choice is arbitrary, but must be done so that the choice at each arrow is consistent with the choice at its neighbors. Left: A qP mode. Right: A qS mode.

This is demonstrated graphically in Figure 3.4. On the left is a qP mode with the "outward" direction everywhere chosen to be "positive". On the right is a qS mode with the sign chosen by a 90° clockwise rotation from the qP one.

What if the modes are anomalous, neither qP nor qS? In section 2.3.2, we saw that the fast and slow solutions for transversely isotropic media in general never touched, except for special values of C_{13} and C_{55} when they could just touch at isolated points. So the arguments from the previous paragraphs about steps 1, 2A, and 2B still go through (although we might have to perturb C_{13} or C_{55} to resolve the point degeneracies if we happen to have elastic constants corresponding to one of the special cases).

What about the choice of positive particle-motion direction in step 2C? Figure 3.5 demonstrates why this also always works. On the left is a typical "anomalous" mode. The "outward = positive" method used in Figure 3.4 does not work for this medium. However, if we simply work around the loop keeping consistent at each step of the way, the loop still closes properly. On the right is a mode where this method fails. If we attempt to proceed as before, we find the vector has performed a "half flip" and the loop will not close. However, the example on the right is impossible: The solutions to equation (2.4)

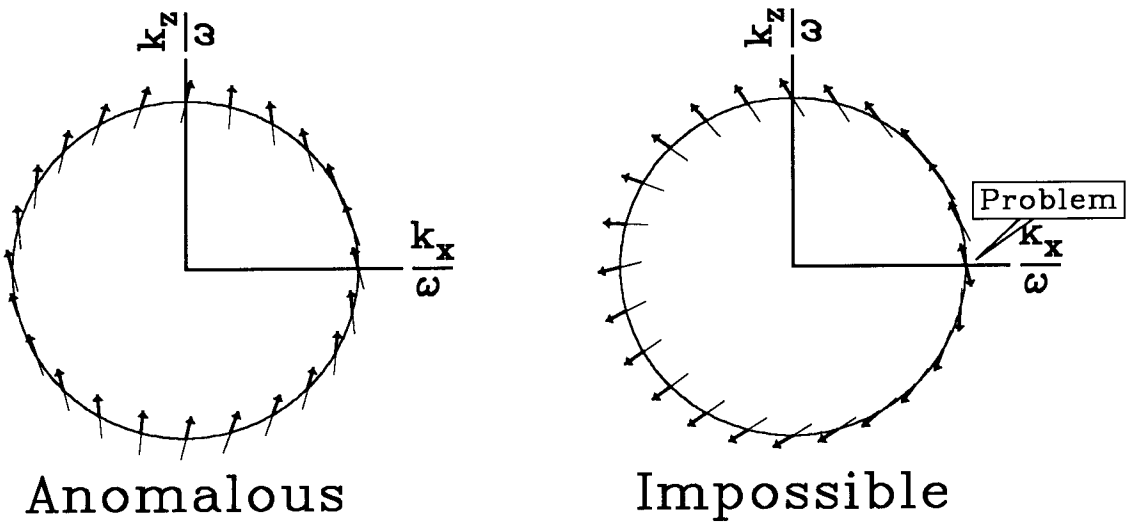

FIG. 3.5. Polar plots of particle-motion direction versus plane-wave propagation direction for exotic two-dimensional media. Left: An "anomalous" mode, neither qP nor qS. Right: An "impossible" mode. No consistent choice of sign for the particle-motion direction is possible.

are unchanged when $\mathbf{k}$ is replaced by $-\mathbf{k}$, and so slowness surfaces for physically realizable media must have point symmetry about the origin. For any physically possible medium, only an even number of "half flips" can be performed in one loop about the origin, and so the loop always closes consistently.[2] This "impossible" slowness surface lacks this symmetry.

3.2.4 qP works in 3D

In two dimensions we found we could always produce two global pure modes; in three dimensions we intuitively expect there should be three. The good news is that three-dimensional qP modes are just as well behaved as two-dimensional ones.

Figure 3.3 shows the slowness surfaces for the familiar isotropic P, SV, and SH wave modes. Note that as in the two-dimensional isotropic case, the P mode is faster than the shear modes for all propagation directions. For media not too far removed from isotropy, there still ought to be a distinct "qP" mode that never intersects the others, and so as in two dimensions there should be no trouble with the qP mode at steps (2A) or (2B).

What about the choice of sign on the particle-motion direction? First I must precisely

[2]Later in this chapter in Figure 3.11 we encounter another kind of loop where half-flips are the usual case.

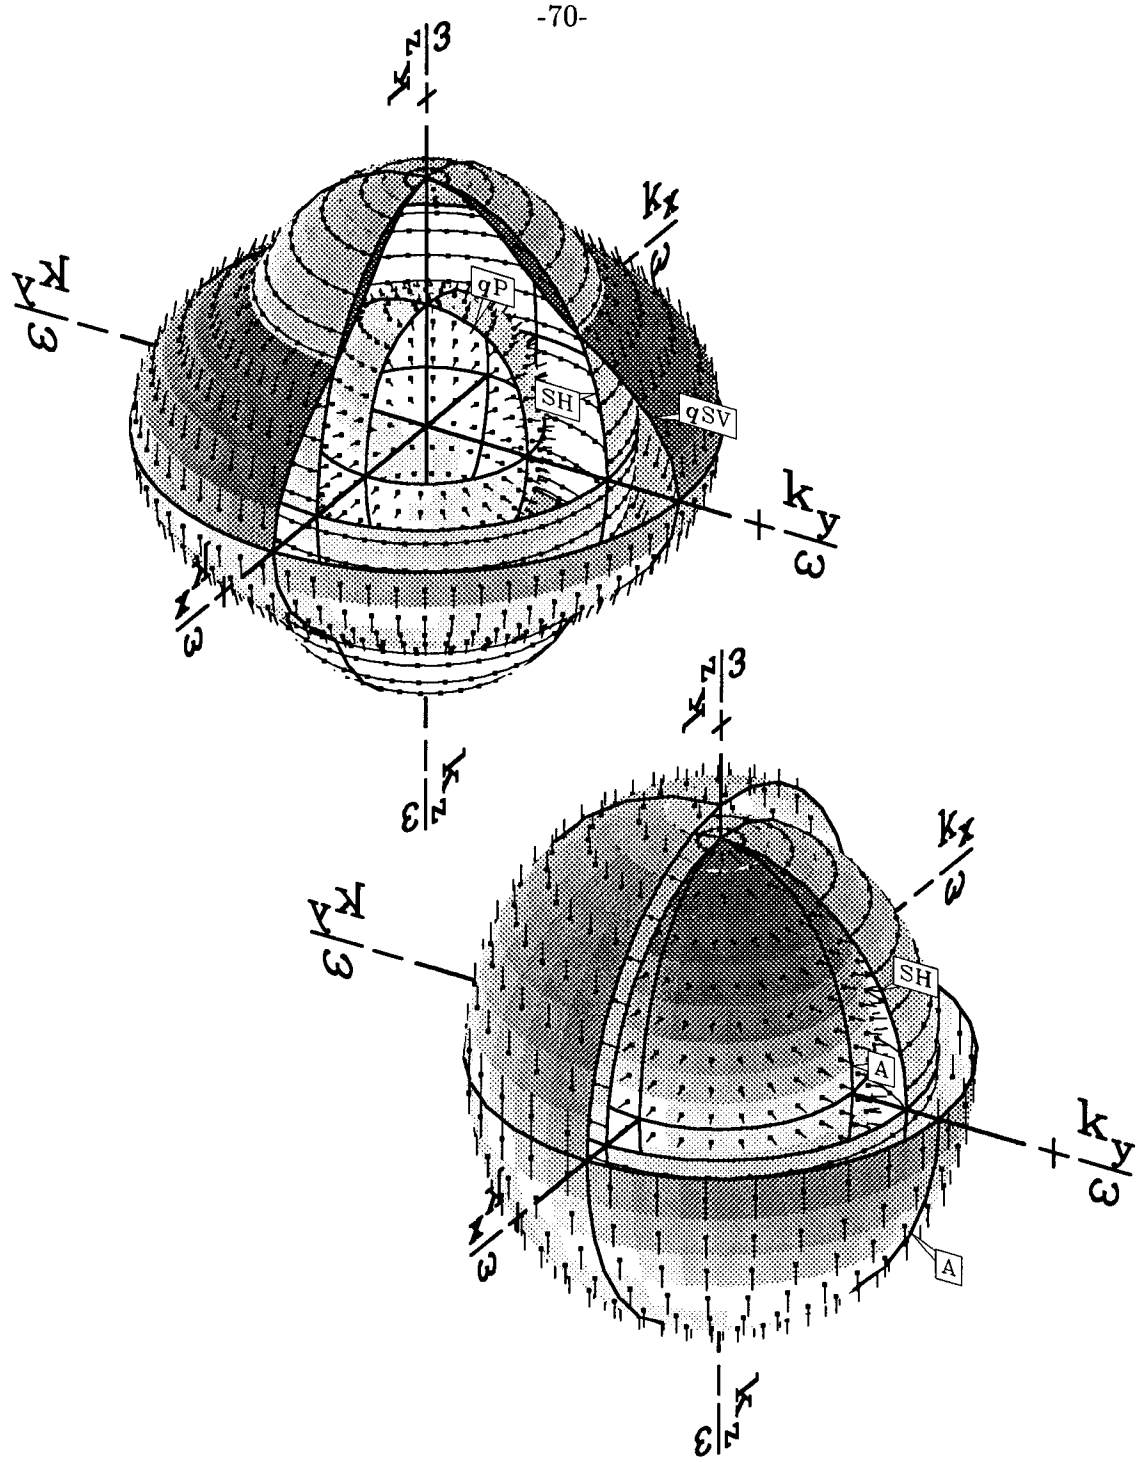

FIG. 3.6. Slowness surfaces for two different transversely isotropic (TI) media. On the left is "Greenhorn Shale" (Jones and Wang, 1981); this medium displays the standard TI "{qP,qSV,SH}" modes. On the right is an "anomalous" transversely isotropic media. There is still an "SH" mode, but the other two modes (both labeled "A" for Anomalous) are qP-like for some propagation directions and qSV-like for others.

define "qP mode". A global wave mode is a qP mode if the particle-motion direction $\mathbf{V}(k_x, k_y, k_z)$ is *not perpendicular* to the plane-wave propagation direction $\mathbf{k} = (k_x, k_y, k_z)$ *for any* $(k_x, k_y, k_z) \neq \mathbf{0}$. Similarly, a global wave mode is a qS mode if the particle-motion direction $\mathbf{V}(k_x, k_y, k_z)$ is *not parallel* to the plane-wave propagation direction $\mathbf{k} = (k_x, k_y, k_z)$ *for any* $(k_x, k_y, k_z) \neq \mathbf{0}$. Modes that are neither qP nor qS are called "anomalous". See Figure 3.6 for examples of "qP", "qS", and "anomalous" modes. (Note my definitions define the various mode types by what they are *not*: qP modes are nowhere pure S; qS modes are nowhere pure P; anomalous modes are neither qP nor qS modes.)

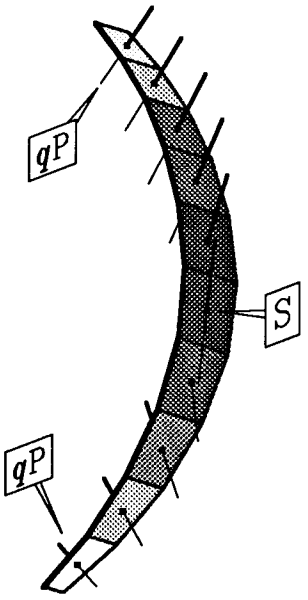

FIG. 3.7. A slice out of a 3-D slowness surface for an anomalously polarized mode. Each particle-motion direction vector (the "sticks" projecting through the surface) has one end drawn thick to indicate that direction has been chosen as "positive". At the top and bottom the mode is qP-like, but the "positive" direction points away from the origin at the top and towards the origin at the bottom.

For trouble to occur at step (2C) of our algorithm there would have to be some closed path on the slowness surface along which the signed particle-motion direction vector performed a 180° rotation. However, the defining property of a qP mode is that it does not have pure S-direction particle motion anywhere. Given a qP mode, then, such a flip in direction cannot occur: somewhere halfway through the P-direction component of the particle motion would have to pass through zero as it changed sign. See Figure 3.7 for a graphical demonstration. (Remember that the particle-motion direction vector $\mathbf{v}$ always has unit magnitude.) Thus there should also not be any difficulty at step (2C) for qP modes.

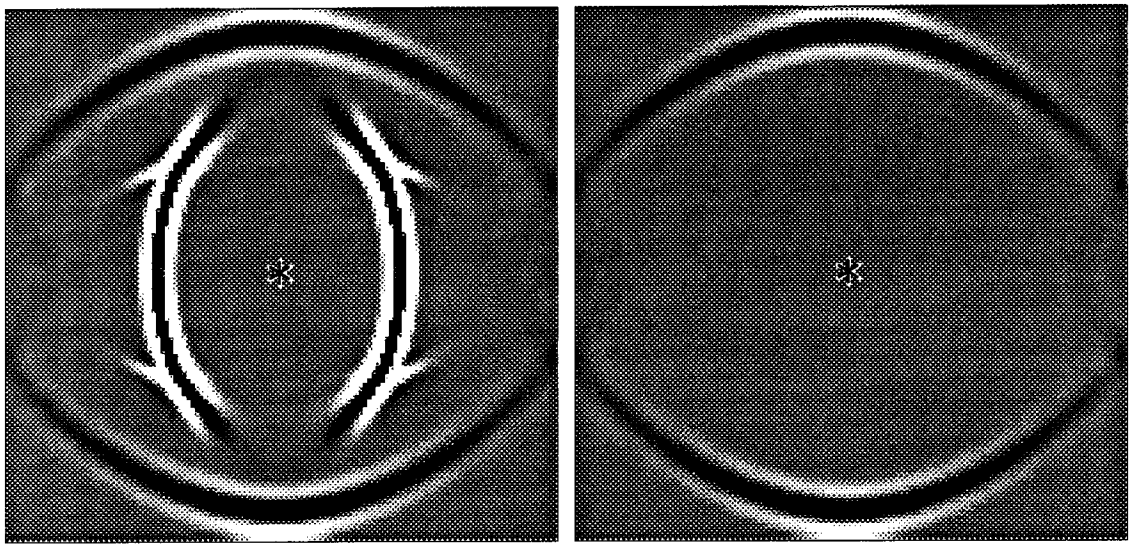

FIG. 3.8. A three-dimensional qP mode-separation example. The medium is orthorhombic, "Cracked Greenhorn Shale" (see Table C.2). The x-z symmetry plane containing the source is shown. (As we might expect from the discussion at the beginning of this chapter, the symmetry plane slices don't show much evidence of the cracks.) The source is a z point force; the z component of displacement is shown. The model is periodic. Wavetype separation works for 3D qP modes. (There is no problem calculating a scalar mode for this example. It was just convenient to use the same vector algorithm described on page 78 for all the examples in this chapter.)

We conclude that three-dimensional qP modes present no new problems over the two-dimensional case. (Of course there is no guarantee that a qP mode exists; Figure 3.32 shows an example of a medium for which all three wave modes are anomalous, neither qP nor qS.)

3D qP mode-separation example

Figure 3.8 shows a three-dimensional qP mode-separation example. The algorithm works as advertised, even though the medium is orthorhombic.

3.3 Shear Singularities

In the previous sections I have shown why mode separation in two dimensions should always work, and that mode separation in three dimensions should always work for qP

modes. Is there any reason to expect mode separation in three dimensions to fail for qS modes?

3.3.1 Defining "the" isotropic shear modes

In isotropic media, it is customary to consider the degenerate shear wave as a single vector pure mode. We typically resolve this globally degenerate mode into "SV" and "SH" scalar modes, but this choice is arbitrary. Isotropy, with its infinite symmetry, lies at the intersection of all the various anisotropic symmetry systems. Perturb the elastic constants away from isotropy and the perfect global symmetry is broken; the shear waves become nondegenerate.

The standard "{SV,SH}" labels result if we perturb in the direction of transverse isotropy with a vertical symmetry axis. There are many other possibilities, resulting in other equally valid definitions of "the isotropic shear modes". If the SV and SH modes as shown in Figure 3.3 are arbitrary, why are they so widely used? SV and SH are popular because they are easy to understand in terms of vertical two-dimensional slices, which is how we usually look at the Earth. Figure 3.9 shows an example of the bizarre sorts of modes that usually result if the elastic constants are perturbed randomly.

Previously we saw that the standard SV and SH modes shown in Figure 3.3 have an artificial discontinuity for vertically traveling waves. Such a point on the slowness surface where the particle-motion direction becomes a discontinuous function of phase direction is called a "singularity". They are usually associated with S or qS waves, and so are often called "shear singularities". The arbitrary modes shown in Figure 3.9 likewise display shear singularities in the particle-motion direction field for waves traveling in several apparently random directions. Singularities are annoying, because they cause discontinuities in the vector field $\mathbf{V}$ at the core of our wavetype-separation algorithm.

3.3.2 Furry ball theorem

Finally we are back to our original question from page 65: Are singularities avoidable? Pure S-type motion requires that the particle-motion direction be perpendicular to the direction of plane-wave propagation. In Figure 3.3, the particle-motion direction is represented by "hairs" attached to the slowness surface. For a pure shear wave, each hair must lie flat against the ball's surface. Imagine what happens if you try to comb the hairs on a furry ball flat and regular everywhere. No matter how you proceed, there must always be

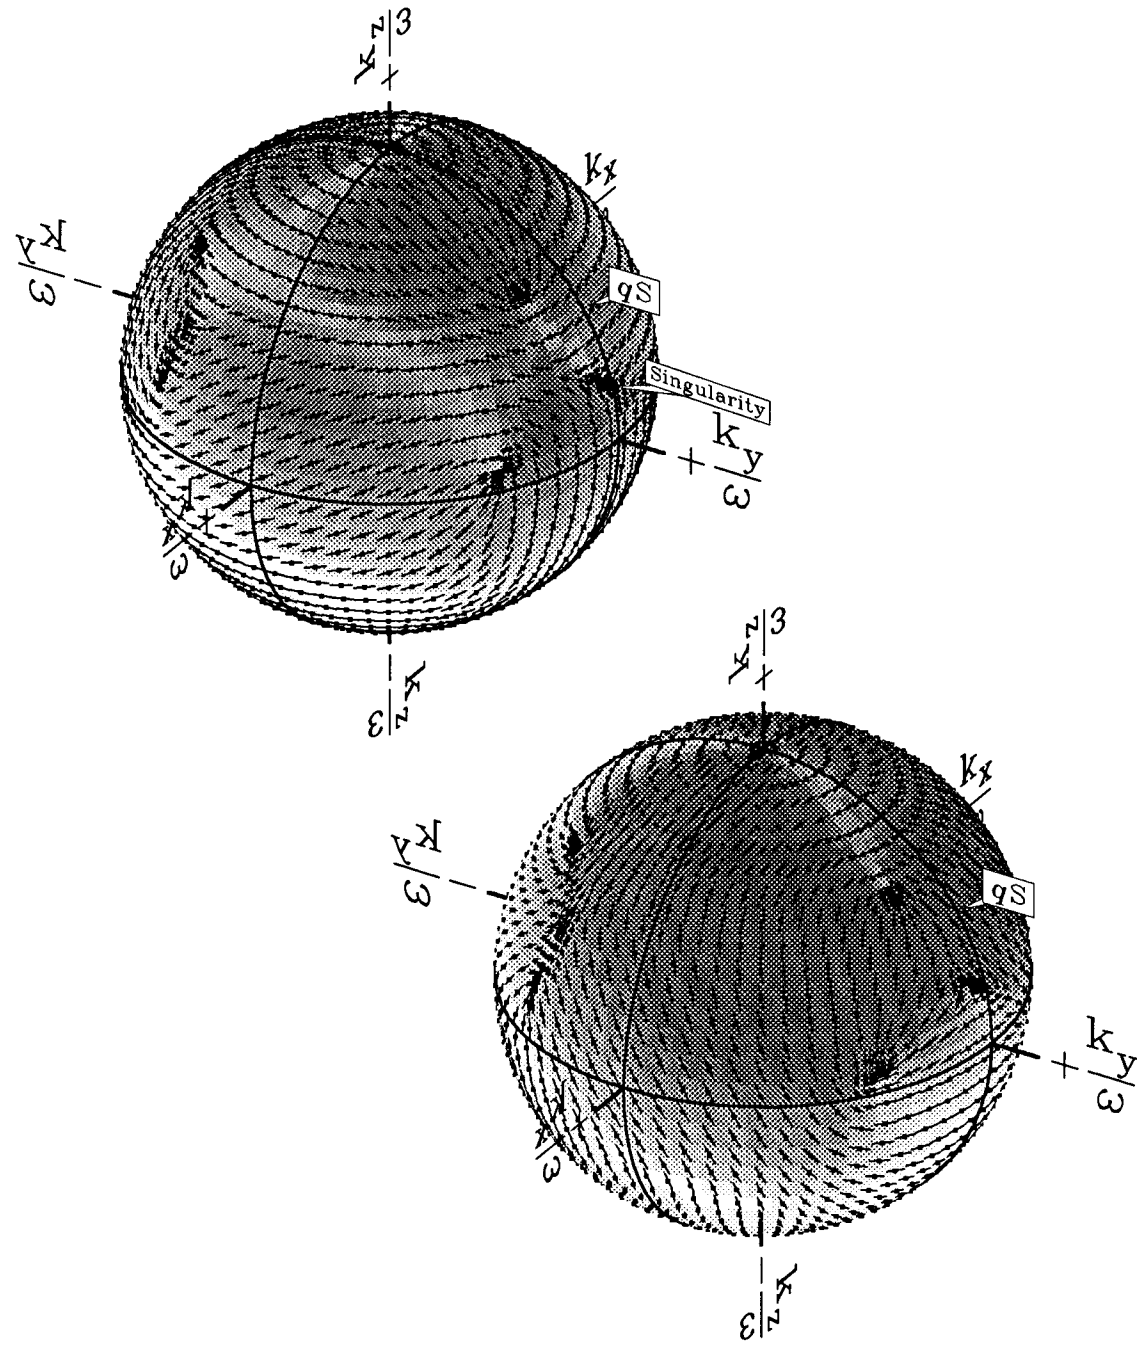

FIG. 3.9. Shear slowness surfaces for arbitrary isotropic shear modes. Mathematically these modes are just as valid a decomposition of the degenerate isotropic shear modes as the standard method shown in Figure 3.3. One singularity is labeled, although several are visible. (The top plot is the slow (qS2) split shear wave; Figure 3.15 shows another view.)

a "whorl" or "cowlick" somewhere, where adjacent hairs point in different directions.[3] We now have the answer: these discontinuous points are shear singularities, and they cannot be eradicated. There is no way to define two globally continuous isotropic shear modes.

Shear-wave singularities must always occur for pure S modes, but what about qS modes? Simply apply the previous argument for pure S modes to the pure S component of the particle-motion direction at each point (i.e., the component of the particle motion perpendicular to the plane-wave propagation direction). The only way to escape the singularity is by letting the pure S component of the particle-motion direction go to zero where we would have had trouble (i.e. the "qS mode" has a pure P-mode direction there: the "hair" at the center of the whorl sticks straight up). Since the defining property of a qS mode is that it does not have pure P-direction particle motion anywhere, this cannot happen.

3.3.3 What do shear singularities represent?

We know that physical wavefronts do not support discontinuous particle-motion directions. How can we reconcile this observation with the ubiquity of singularities?

For any phase direction there are three *orthogonal* modes, so it is impossible for one wavetype to have a discontinuous particle-motion direction in isolation; at least one of the other wavetypes must share in the discontinuous particle-motion directions at the singular point (although rotated by 90°). Precisely at the singular point the particle-motion directions of these two modes must become indeterminate. This can only happen if the two modes are degenerate there.

A singularity, then, must correspond to a point where two otherwise continuous and distinct modes touch and temporarily lose their unique identities (Crampin and Yedlin, 1981). We can see this illustrated in Figure 3.6; the singularity for vertically traveling S waves (along the k_z axis) always occurs on two intersecting surfaces. On the axis particle motions in the x and y directions are symmetrically equivalent. Figure 3.10 shows the nature of singular points even more clearly. In this orthorhombic example the qS modes are clearly two separate shells that pinch together at a finite number of singular points.

A shear singularity does not cause a discontinuous particle motion in reality because for some range of angles around the singular phase direction two orthogonal wavetypes are

[3]This is a special case of a well known result from differentiable manifold theory: no even-dimensional sphere has any continuous nonvanishing field of tangent vectors (Boothby, 1975).

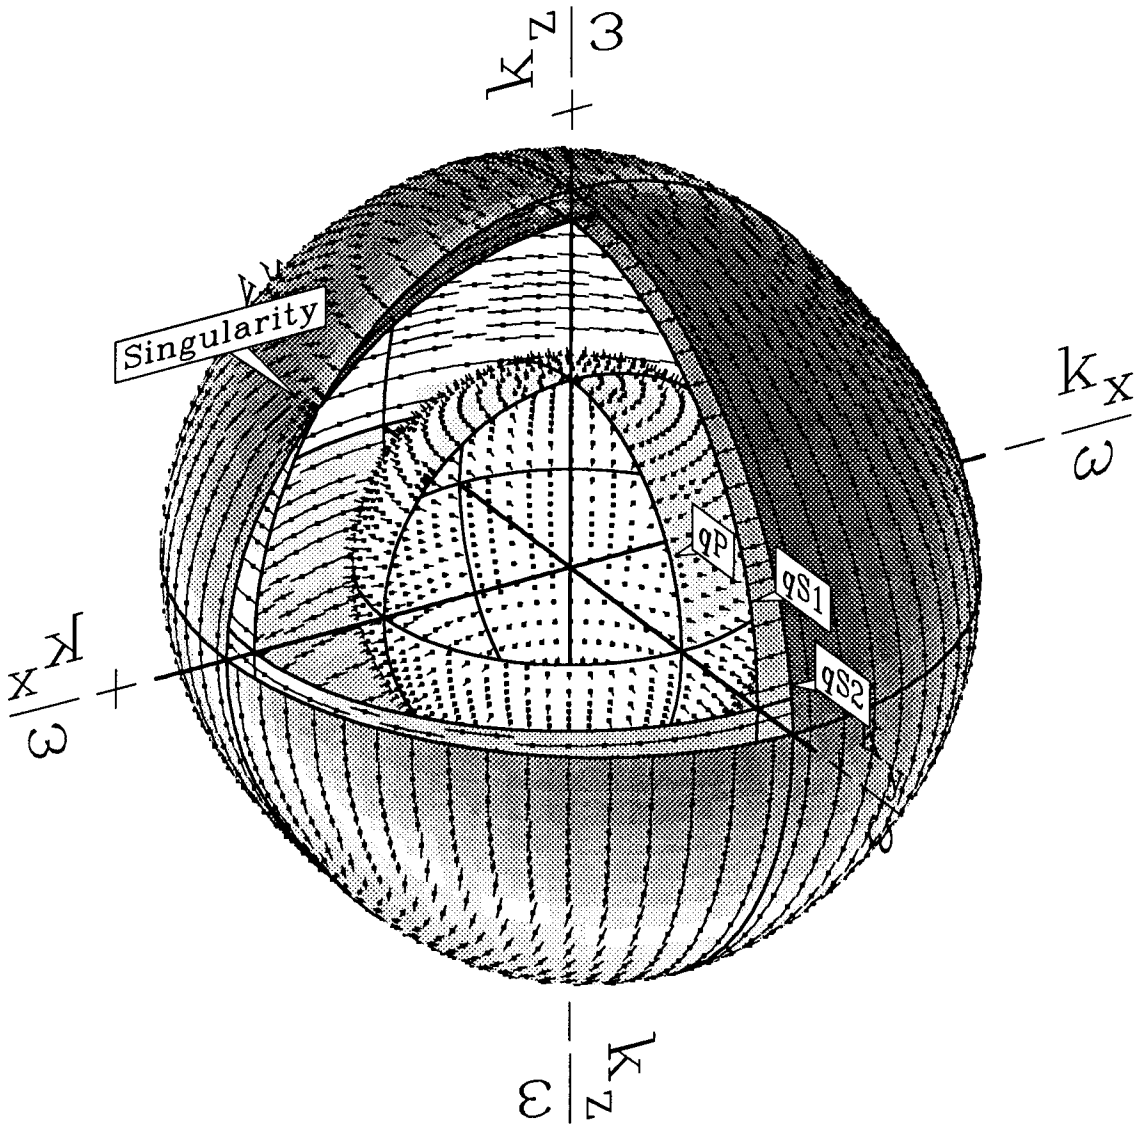

FIG. 3.10. The three-dimensional slowness surfaces of the same orthorhombic medium shown in Figure 3.2. One octant has been removed to show the true topology: two nested qS surfaces that just touch at the singularities. (There are four singularities in all; the other three are symmetrically equivalent to the labeled one.) Neither of the surfaces can be adequately described as "qSV" or "qSH".

strongly coupled together. While individually their particle-motion vectors rapidly twist around, for any finite frequency their sum remains smooth. (Singularities can produce other visible effects, however, as will be demonstrated in section 3.5.1.)

3.3.4 qS modes are inseparable

I have now shown that any qS mode has singularities on it, and that singularities are places where two modes touch. If there are two qS modes that are both everywhere slower than a single qP mode, then the two qS modes must be touching each other. Does this automatically mean mode separation of qS waves is impossible? How does trouble occur?

At step (2A), we must expect occasional trouble with nonunique qS solutions in three dimensions, because shear-wave singularities represent degenerate directions. A typical way around such problems in ray tracing is to perturb the elastic constants so that the singularity moves off our current ray direction (Jech and Pšenčík, 1989). This procedure shouldn't create significant errors since a tiny perturbation in the elastic constants should only cause a tiny change in the global wavefield. (We are still left with the problem of figuring out which of the barely split modes is the one we want.)

Steps (1) and (2B) of our algorithm presuppose we know what the global qS modes look like. (I haven't even yet proven that there is such a thing as a global qS mode for general anisotropy!) The isotropic modes are usually separated into {P,SV,SH}, and the transversely isotropic modes into {qP,qSV,SH}, but we have seen these designations are arbitrary at best. Is there any sensible way to separate the connecting qS surfaces into two qS pure modes for completely general anisotropic media?

The example in Figure 3.10 suggests one possibility. We wish to cut the two-sheeted qS surface in this figure into two pure modes. We cannot do this arbitrarily. Step (2B) of our algorithm requires that the three plane wavetypes at each (k_x, k_y, k_z), $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$, must be identified as a permutation of the three global wave modes (usually {qP,qS1,qS2}) evaluated for that phase direction. The three resulting pure modes must be unique single-valued *continuous* functions of phase velocity versus plane-wave propagation direction. (In fact I would say this must be the definition of a global pure mode.) As should be clear from Figure 3.10, the only possible place to cut is at the shear singularities.

This definition makes the modes very easy to identify computationally; the modes are simply sorted on their phase velocity (ω/k) into an inner (fastest) qP, a middle qS1, and an outer (slowest) qS2 mode (Crampin, 1981). (Of course while this method is the most

general, it still may make sense to use notations such as {qP,qSV,SH} when appropriate — although these symmetric notations are misleading, as will be shown in section 3.5.1.)

Unfortunately, as can be seen in Figure 3.11, sorting the modes by phase velocity does nothing to ameliorate the particle-motion discontinuity at the shear singularities. As I will demonstrate in the next section, this local discontinuity in the Fourier domain produces a large planar artifact in the space-domain version of the corresponding wavetype-separation operator, although it can be suppressed in certain special cases.

We have seen there are troublesome problems at step (2B). What about step (2C)? Figure 3.11 shows most of the particle-motion direction vectors as two-sided, since mathematically there is nothing in equation (2.4) to distinguish any eigenvector $\mathbf{v}$ from its opposite $-\mathbf{v}$. I have attempted to pick one sign as "positive" for each particle-motion vector along a path (shown as a dotted line) around the singularity; these vectors are shown as one-sided rays. Unfortunately if we pick the vectors to preserve continuity as we move along the path, we find the signs don't match when we complete the loop. There is an inescapable branch cut in step (2C).[4]

This branch cut makes a mess of any attempt to construct a pure-mode scalar field. If we are satisfied merely to separate the wavetypes, however, we can avoid the sign-choice problem by using a slightly modified algorithm:

(3.′) Fourier transform $\mathbf{U}(x,y,z;t)$ over x, y, and z, obtaining $\hat{\mathbf{U}}(k_x,k_y,k_z;t)$.

(4.′) For all (k_x,k_y,k_z):

{

$$\hat{\mathbf{M}}(k_x,k_y,k_z;t) = \mathbf{V}(k_x,k_y,k_z)\Big(\mathbf{V}(k_x,k_y,k_z)\cdot\hat{\mathbf{U}}(k_x,k_y,k_z;t)\Big)$$

}

(5.′) Inverse Fourier transform $\hat{\mathbf{M}}(k_x,k_y,k_z;t)$ over k_x, k_y, and k_z, obtaining $\mathbf{M}(x,y,z;t)$.

The branch cut does not adversely affect this operator because the sign of $\mathbf{v}$ enters the equation twice, and so cancels out.

[4]This is an example of a phenomenon of much current interest in quantum physics, anholonomy (Berry, 1990). The basic idea is that after some parameters of a system have been perturbed around a closed loop in a continuous way, some other seemingly unrelated parameters are found to have changed sign or otherwise shifted in phase.

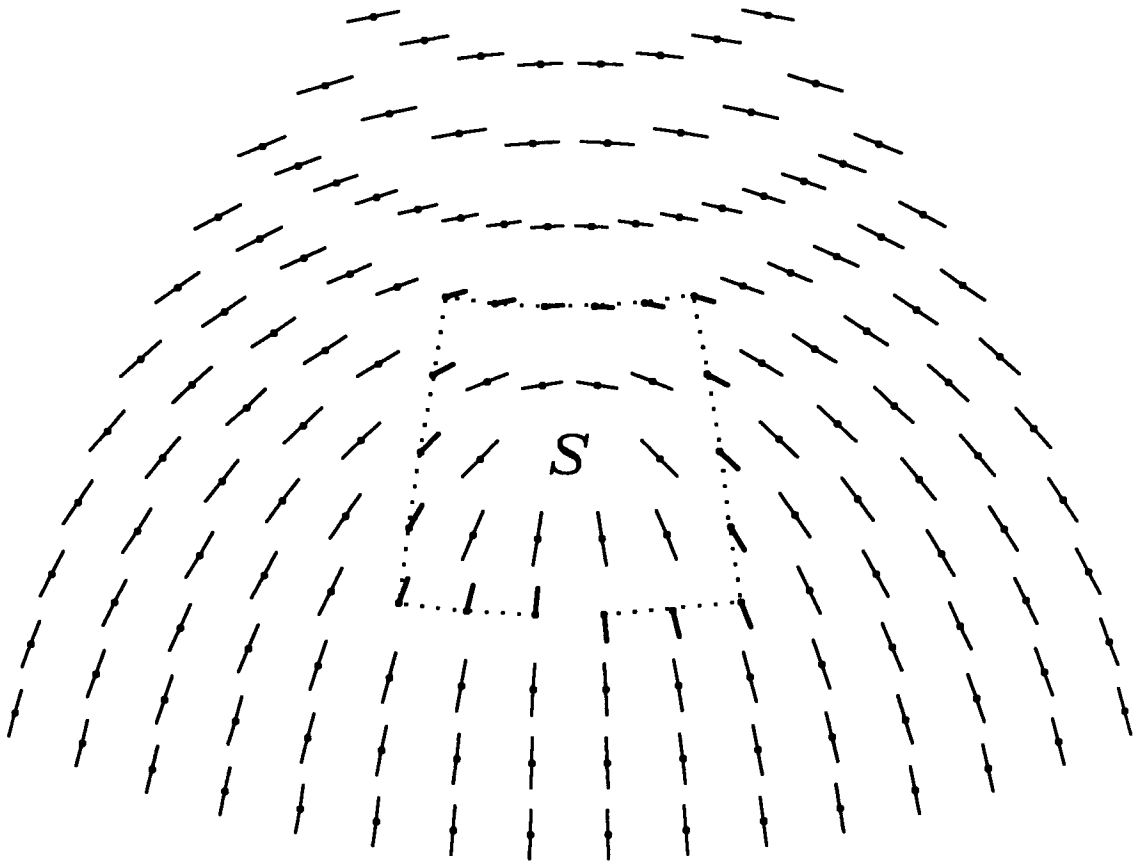

FIG. 3.11. A close-up view of the outer surface in Figure 3.10, centered on the shear singularity (here indicated by an "S"). Note that the particle-motion directions become discontinuous at the singularity. There is a more subtle disturbance at the singularity as well. Most of the particle-motion direction vectors are drawn two-sided. For the vectors along the dotted path looping around the singularity, however, we have attempted to pick a preferred "positive" direction and so these are drawn one-sided. Once we make a choice of sign at the beginning, we are constrained at each step to choose the direction consistent with the previous choice. Unfortunately, the continuity we have carefully maintained going around the loop is lost at the end; the choices at the beginning and end of the loop do not agree. Contrast this situation with the one in two dimensions shown in Figure 3.5.

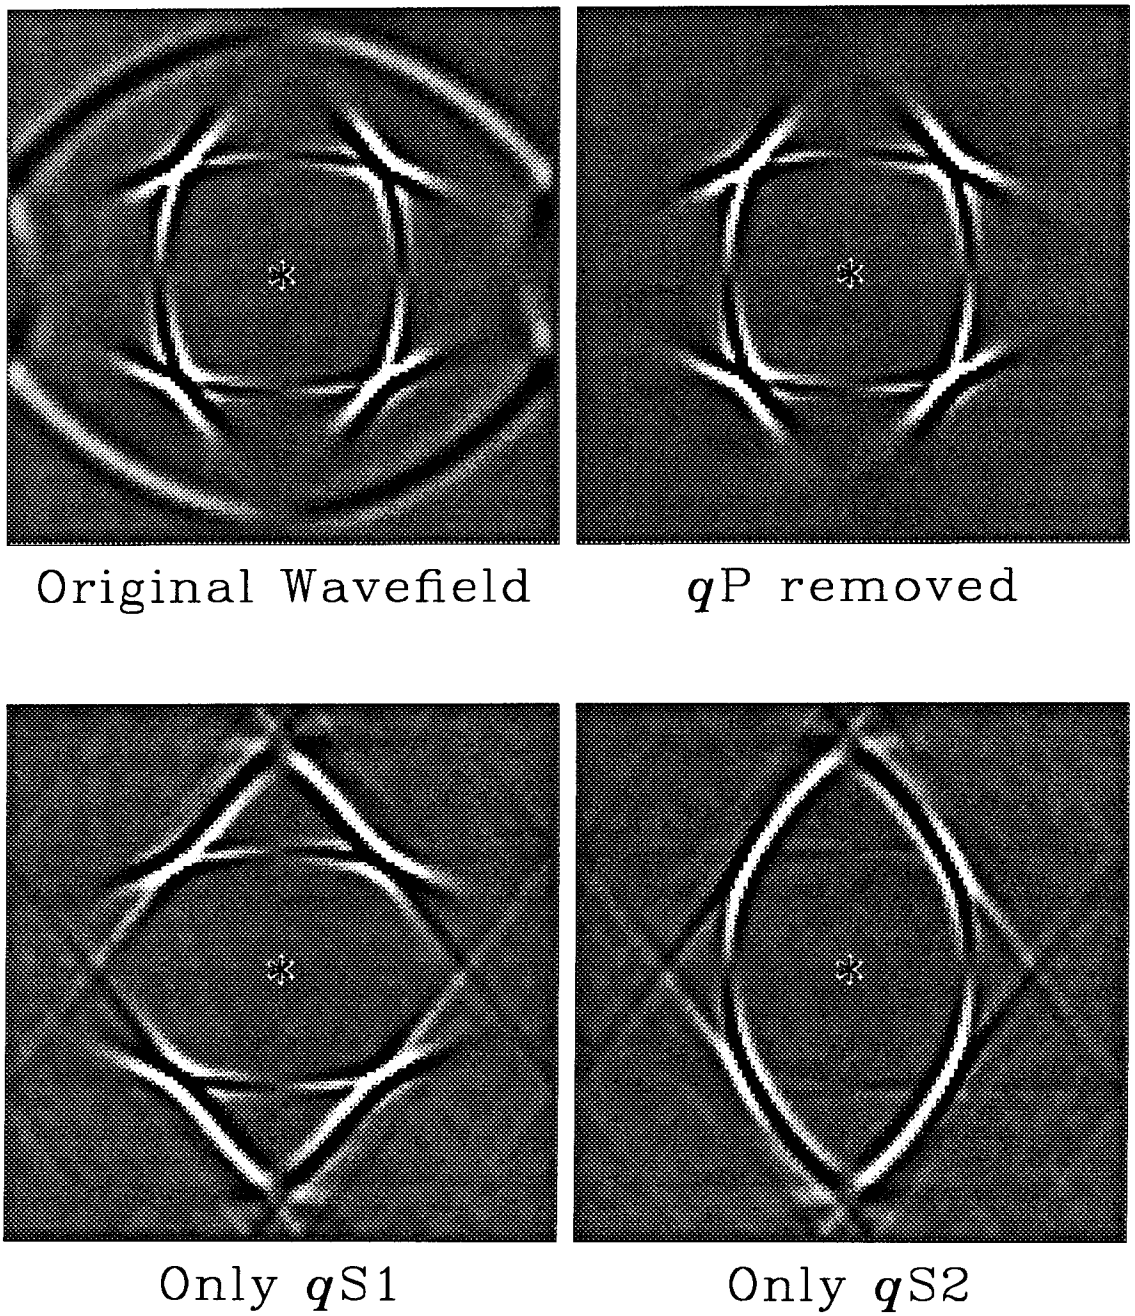

FIG. 3.12. Three-dimensional qS mode-separation examples. The plot and parameters are like those in Figure 3.8, except this time the x component of displacement is shown and a gpow of .8 has been applied to make weak events more easily visible. Upper left: The original wavefield. Upper right: The qP mode has been successfully nulled. Lower left: Trying to pass only the qS1 wave causes artifacts. Lower right: The qS2 wave likewise doesn't work by itself. The sum of the two lower plots gives the upper right plot.

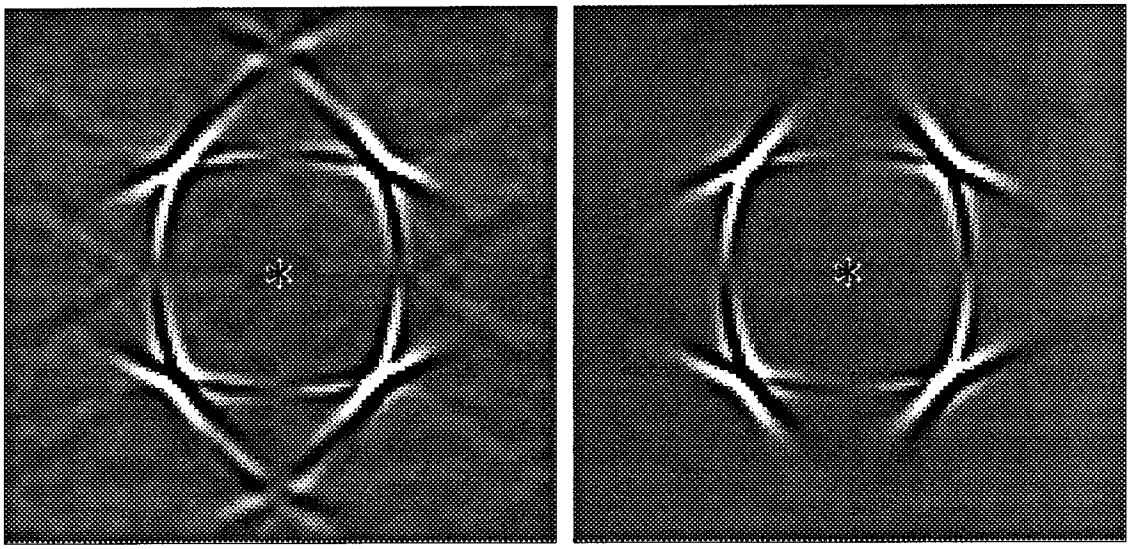

FIG. 3.13. Two attempts at defining and separating a three-dimensional qSV mode. The plot and parameters are like those in Figure 3.12. Left: For each (k_x, k_y, k_z) the qS plane-wave mode with particle motion more nearly "SV" was chosen. This pseudo-mode does not have a continuous phase-velocity function. Right: The orthorhombic elastic constants were perturbed to become transversely isotropic, and the wavetype-separation operator appropriate for the qSV mode for the perturbed TI medium was applied to the original orthorhombically anisotropic wavefield. As a result the qP mode is not completely nulled.

This new wavetype-separation operator does not attempt to collapse the vector wavefield onto a scalar pure mode M. Instead it nulls the two unwanted wave types while passing the desired pure-mode component through unchanged. Since $\mathbf{M}$ is a vector wavefield like $\mathbf{U}$, some important manipulations become easier than they would be with a scalar mode like M. For example, we can construct the difference $\mathbf{U} - \mathbf{M}$; this *nulls* a single wavetype instead of passing a single wavetype.

Unfortunately, the fundamental problem of particle-motion discontinuities at shear singularities remains. There is no solution to this problem; it is a consequence of the basic inconsistency between the "furry ball" theorem and the definition of a qS "pure mode".

I conclude that in general anisotropic media it is not possible to cleanly separate the "qS1 and qS2 modes". The particle-motion discontinuity at the singularities will always cause the two modes to leak energy into each other. This mode-mode coupling is especially bothersome for elastic ray tracing in three dimensions. It is possible to handle the coupling, but only by propagating the two not-quite-distinct shear modes together

(Chapman and Shearer, 1989). While it may often prove useful to differentiate two qS waves for some set of propagation angles, globally they are inextricably linked and cannot be cleanly separated.

3.3.5 Examples

Figure 3.12 shows how this works in practice. The upper left plot shows the trial input wavefield. In the upper right plot the qP wave has been successfully removed. In the lower two plots I have attempted to pass only the qS1 or qS2 pure modes. The Fourier-domain particle-motion discontinuities at the shear singularities create strong planar artifacts that can be seen to wrap around several times. The artifacts have opposite signs on the qS1 and qS2 plots; when the two qS modes are summed the artifacts cancel and the "qP removed" plot results.

For these models a z source was used. Since the only particle-motion discontinuity for SV modes is for vertically traveling shear waves, which are unexcited by a z source, we might expect an SV mode-separation operator to be more tractable. There are two possible approaches to use; Figure 3.13 shows both of them.

In the left example in Figure 3.13 for each (k_x, k_y, k_z) plane-wave direction I have chosen the mode that most nearly fits the designation "SV". Sure enough, this operator picks out something resembling an SV mode, but the resulting wavefield is cluttered by the artifacts caused by both the particle-motion direction *and* phase-velocity discontinuities.

In the right example in Figure 3.13 I have perturbed the elastic constants of the wavetype-separation operator to become transversely isotropic. (This was easy to do, since the original orthorhombic elastic constants were generated by numerically "cracking" (Schoenberg and Muir, 1989) transversely isotropic Greenhorn Shale (Jones and Wang, 1981) in the first place.) For the perturbed medium there is a true qSV mode, and so the separation works (given the z source used in this example). The qP mode does slightly leak through because the elastic constants don't quite fit the medium, however.

3.4 Classifying singularities

Clearly singularities are one of the more important features of slowness surfaces; how do they change when the underlying elastic constants are perturbed?

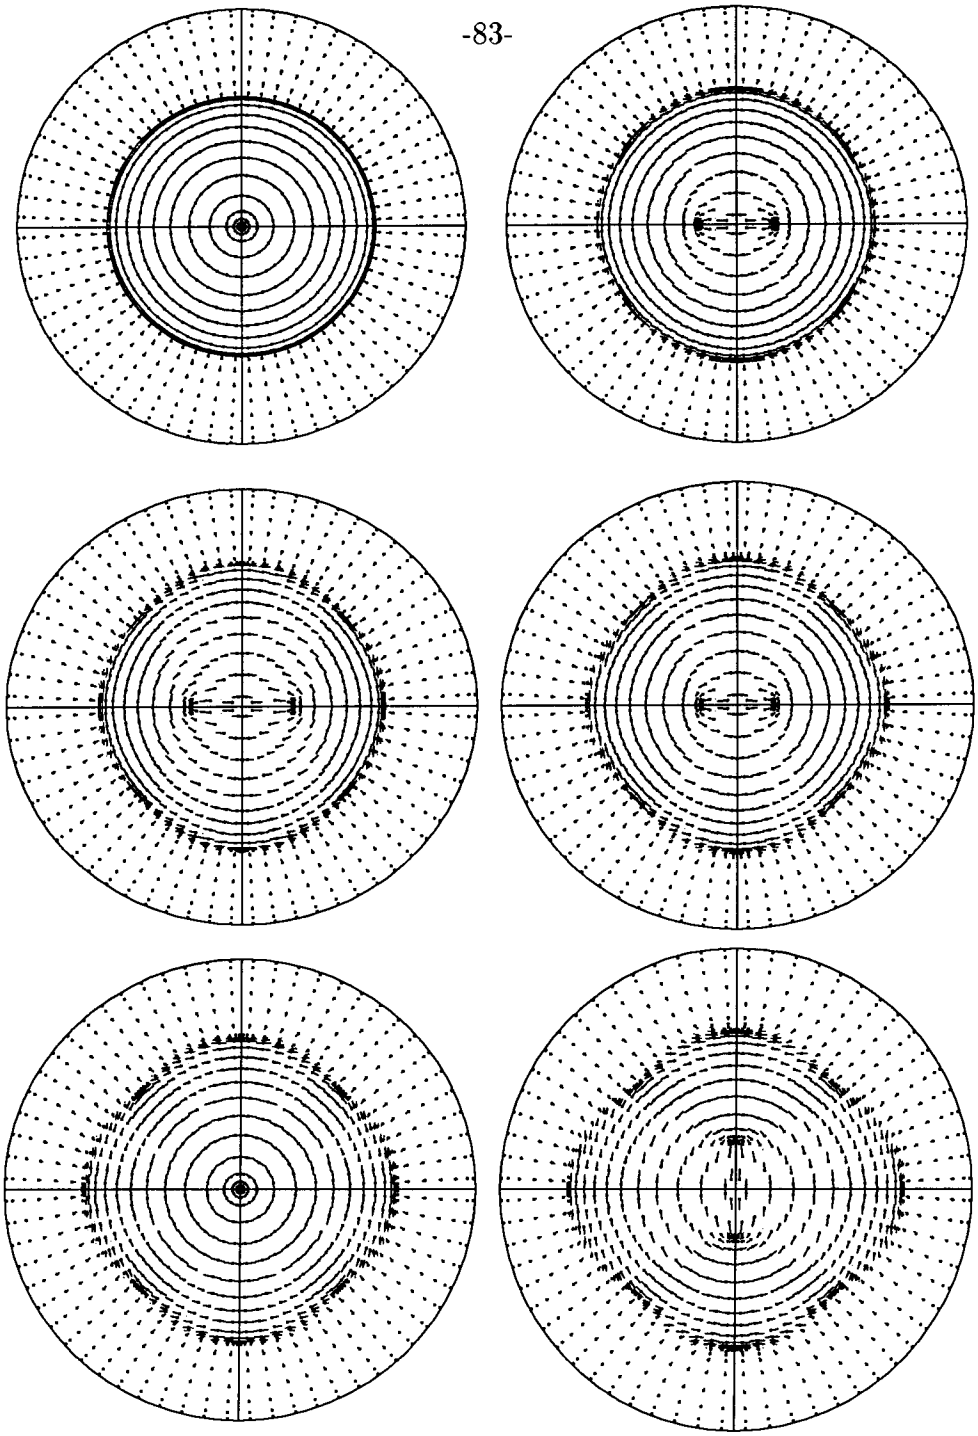

FIG. 3.14. Successive perturbations away from transverse isotropy. (The plots are ordered from left to right and then top to bottom.) The first (top left) plot corresponds to the TI medium shown in the top plot in Figure 3.6. Only the outer qS2 wave is shown here. The view is from the "North Pole": the $+k_z$ axis points directly out of the page, the $+k_x$ axis points to the bottom of the page and the $+k_y$ axis points to the right. For the second (top right) and third (middle left) plot successively more cracks are added in the x-z plane. Starting with the fourth (middle right) plot the first set of cracks are held constant and more and more cracks are added in the y-z plane instead. For the fifth plot (lower left) the amount of x-z and y-z cracking are the same.

Figure 3.14 shows a succession of qS2 slowness surfaces as the elastic constants are successively perturbed away from transverse isotropy. The first plot is transversely isotropic. Cracks are then added in the x-z plane to perturb the elastic constants to become orthorhombic (Nichols, Muir, and Schoenberg, 1989). The TI kiss singularity at the center of the plot splits into two point singularities on the k_y axis, a typical configuration for orthorhombic media (Crampin and Kirkwood, 1981). Starting with the fourth plot in the sequence the x-z cracks are held constant and another crack set is added in the y-z plane. The singularities move back towards the center and converge back into a single kiss singularity when the amount of cracking in the two perpendicular directions is equal. As the cracking in the y-z plane becomes dominant in the last plot, the kiss singularity bifurcates again, this time along the k_x axis.

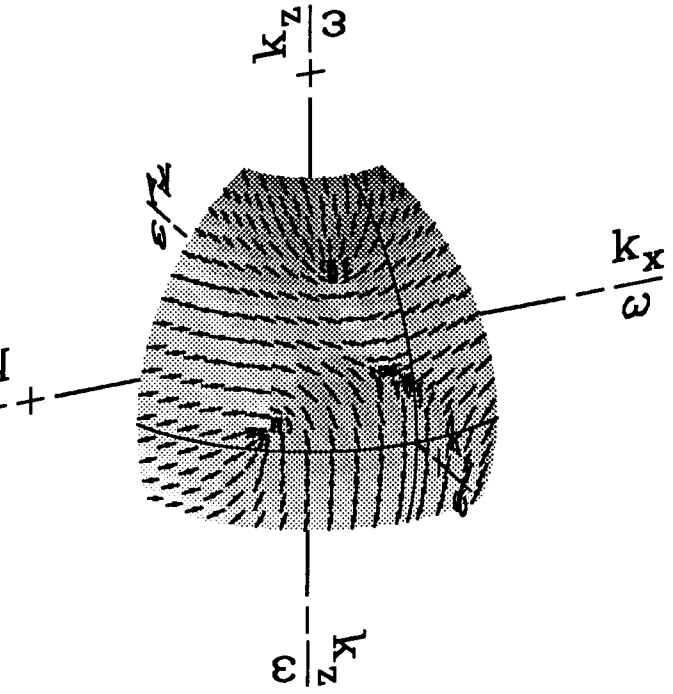

FIG. 3.15. Another view of the upper plot in Figure 3.9. A Saskatchewan shaped portion of the slowness surface containing the three point singularities visible near the $+k_y$ axis has been cut out and rotated to the front. (Ideally, Saskatchewan is bounded by two lines of constant latitude and two lines of constant longitude.)

This behavior suggests to me an alternative notation for singularities that has some advantages over the "{point, kiss, intersection}" terminology now in common use. (This notation was first introduced by Crampin and Yedlin (1981), and more recently has been discussed by Winterstein (1990).) I suggest that singularities could be classified by the number of half-loops the particle-motion direction vector completes in one loop around the

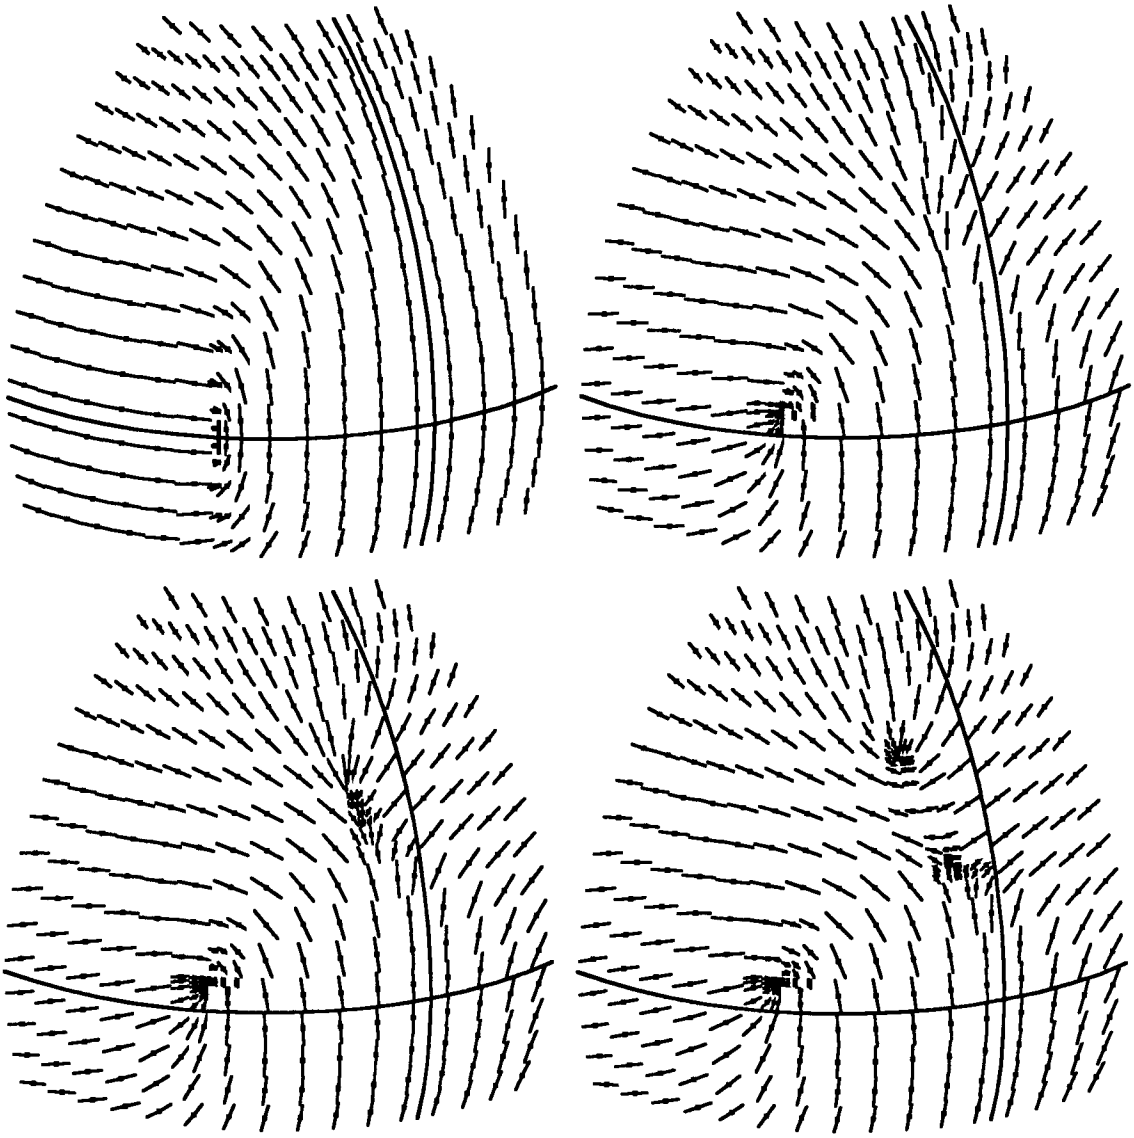

FIG. 3.16. This sequence of four plots (left to right and then top to bottom) shows how the particle-motion directions change as an orthorhombically anisotropic medium is perturbed to become generally anisotropic. Two singularities of opposite sign appear together out of nowhere and then separate. Figure 3.15 forms the fifth plot in the sequence and shows the orientation of all of the plots. The elastic constants are given in Note 3 on page 169.

singularity, reminiscent of integration around residuals in complex analysis. In particular it is possible for the particle-motion vector to rotate in the same sense as the loop is traversed (+) or in the opposite sense (−). This seems to be a fundamental property of singularities that is robust against perturbations in the elastic constants; the sum about a group of singularities also appears to be conserved when two of them merge or a double one splits.

Figure 3.11 shows a canonical example of an order +1 singularity; the particle-motion direction vector performs a half flip in the same direction as the traverse around the singularity. In Figure 3.14 the order +2 singularity on the k_z axis at the center of the plot splits into two order +1 singularities. Although harder to see, the infinity of order 0 intersection singularities (the dark circle halfway out from the center in the first plot in Figure 3.14) splits into 4 order +1 singularities (on the k_x and k_y axes) and 4 order −1 singularities (in between).

The labeled "point" singularity in Figure 3.9 shows a clearer example of an order −1 singularity (the two adjacent ones are each of order +1).[5] Figure 3.16 shows how this singularity spontaneously appears as part of a singularity-antisingularity pair as originally orthorhombic elastic constants are perturbed more and more. The original order +1 singularity of the orthorhombic medium moves slightly but cannot split, since the particle-motion direction vector must complete an integral number of half-flips.

3.5 Some canonical modeling examples

In this section I use elastic finite-difference-modeled wavefields to illustrate some of the more interesting possible properties of wave propagation in three-dimensional anisotropic media.

3.5.1 Perturbing TI

In section 3.4 we saw that relatively small changes to the elastic constants could cause fairly dramatic topological changes in the singularities. Intuitively we expect that small

[5] A closely related classification scheme for "umbilic points" on surfaces is described by Berry and Hannay (1977). My order −1 singularity they would call a "Star" singularity, which has "index −1/2". My order +1 singularity they would call a "Lemon" singularity, which has "index +1/2". They also describe one more possible singularity type, the "Monstar", which also has index +1/2. I have not been able to generate an example of a "monstar" shear singularity.

perturbations in the material properties of a medium should only be able to cause small changes in the associated wavefield. Is it possible for topological changes in the singularities to cause noticeable effects in the resulting wavefield, even when the change in the elastic constants is small?

Figure 3.17 shows two slices through a three-dimensional wavefield snapshot. (The qP mode has been nulled so I can squeeze as much resolution out of the 128^3 finite-difference model as possible.) One slice is oriented lies in the x-z plane, the other is oriented at 30° to the x-z plane. For each orientation two plots are shown; one shows displacement in the SV direction, the other in the SH direction. (Note I am using "SV" and "SH" here as directions, not as labels for modes.) The source is a pure z point force.

There are two things to note in Figure 3.17. First, the two sets of plots appear to be identical. This is because the medium is transversely isotropic, and so the two orientations are in fact symmetrically equivalent. Second, there appears to be no SH motion at all. Since the medium is transversely isotropic there is a pure SH mode; this mode is orthogonal to the z source and so is not excited.

We now take Figure 3.17 and perturb the elastic constants by adding cracks in the x-z plane. (The cracking is not severe, only enough to slow the qSV wave propagating perpendicular to the plane of the cracks by about 5%.) Given the qSV and SH modes of the original transversely isotropic medium shown in Figure 3.17 and the relatively small perturbation to the elastic constants, we would expect to see something like qSV and qSH modes in the cracked medium. Figure 3.18 shows the actual results. As expected, the SV plots look much the same as before. The SH plot in the x-z symmetry plane is again zero. The SH plot at 30° appears quite strange, though. There is a "qSH" wavefront dimly perceptible (maybe not, depending on the reproduction quality of your copy!), but far stronger is a perplexing nearly planar event that connects the expected "qSV" and "qSH" wavefronts, but corresponds to nothing in transverse isotropy. What is going on?

The problem is that our intuition has been learned from studying highly symmetric two-dimensional examples (or, even worse, isotropy). In Chapter 2 we encountered several examples where events would disappear just at certain exceptional values of the elastic constants. (For example, compare cases 1 and 10 in Figure 2.4 on page 24 with the corresponding finite-difference results in Figures 2.7 and 2.8 on pages 32 and 33.) One could not understand transverse isotropy by studying only such special cases. Unfortunately two-dimensional anisotropy itself, with many elastic constants held equal, is just such a

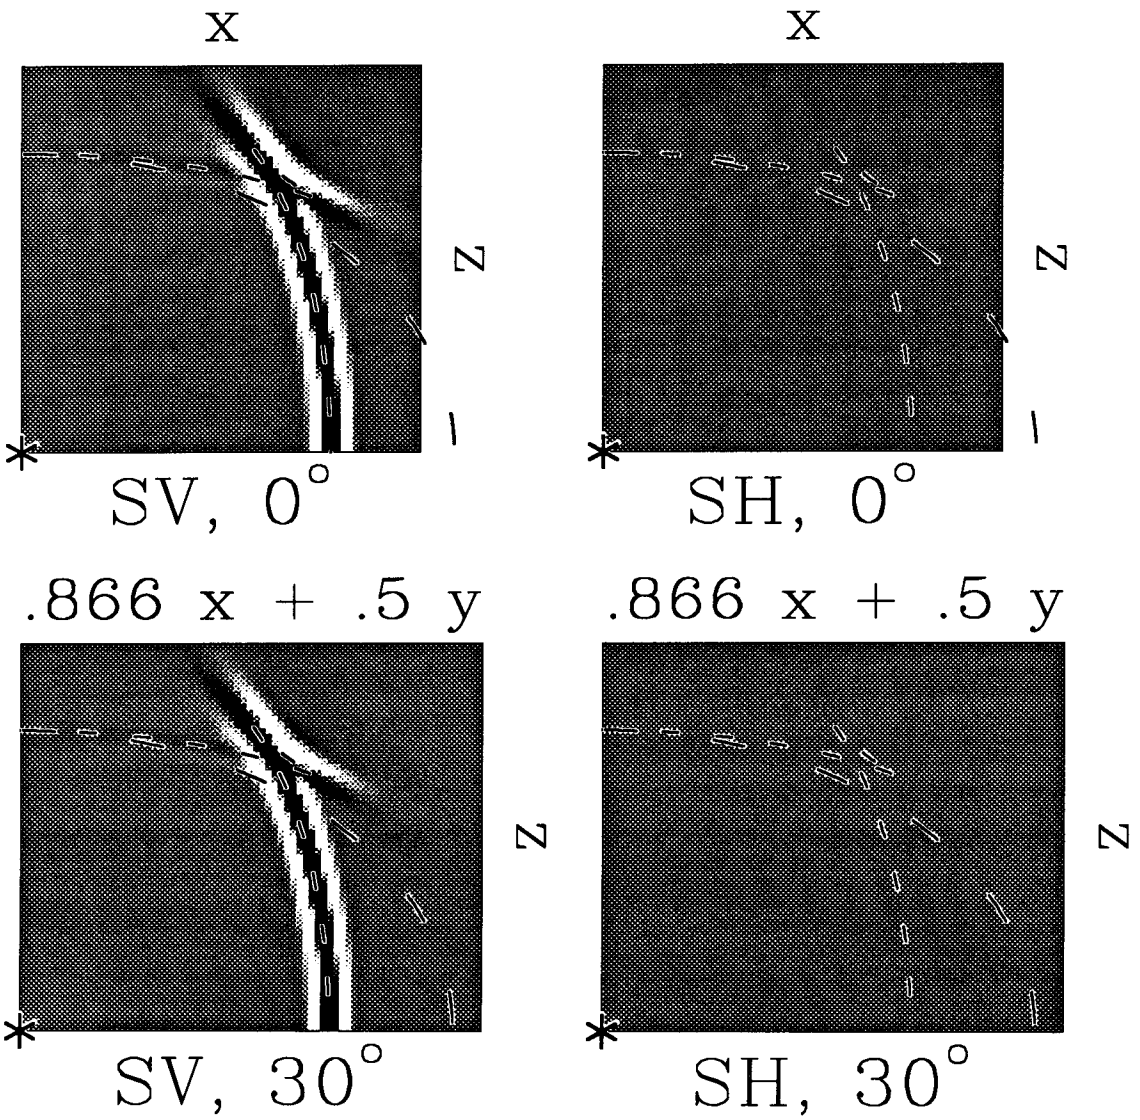

FIG. 3.17. Four slices through a three-dimensional Greenhorn Shale (Jones and Wang, 1981) snapshot, corresponding to the medium shown in the top left plot in Figure 3.16. The qP wave has been nulled to allow the maximum resolution on the shear waves. (The original model had 128^3 gridpoints; the slices are through one octant.) The source is a z point force at the lower left corner. The top row shows slices through the x-z plane, the bottom row shows slices rotated 30° about the z axis out of the x-z plane. (This slice is wider because the snapshot cube is being cut at an oblique angle.) The left column shows the component of displacement in the SV direction (i.e. in the plane of the slice and perpendicular to a vector pointing away from the source). The right column shows the component of displacement in the SH direction (i.e. perpendicular to the plane of the slice). Theoretical impulse-response curves for the qSV (finer dashes) and SH (coarser dashes) modes for this medium are overlayed on the plots. Since the medium is transversely isotropic and the source is aligned with the symmetry axis, the 0° and 30° slices are equivalent.

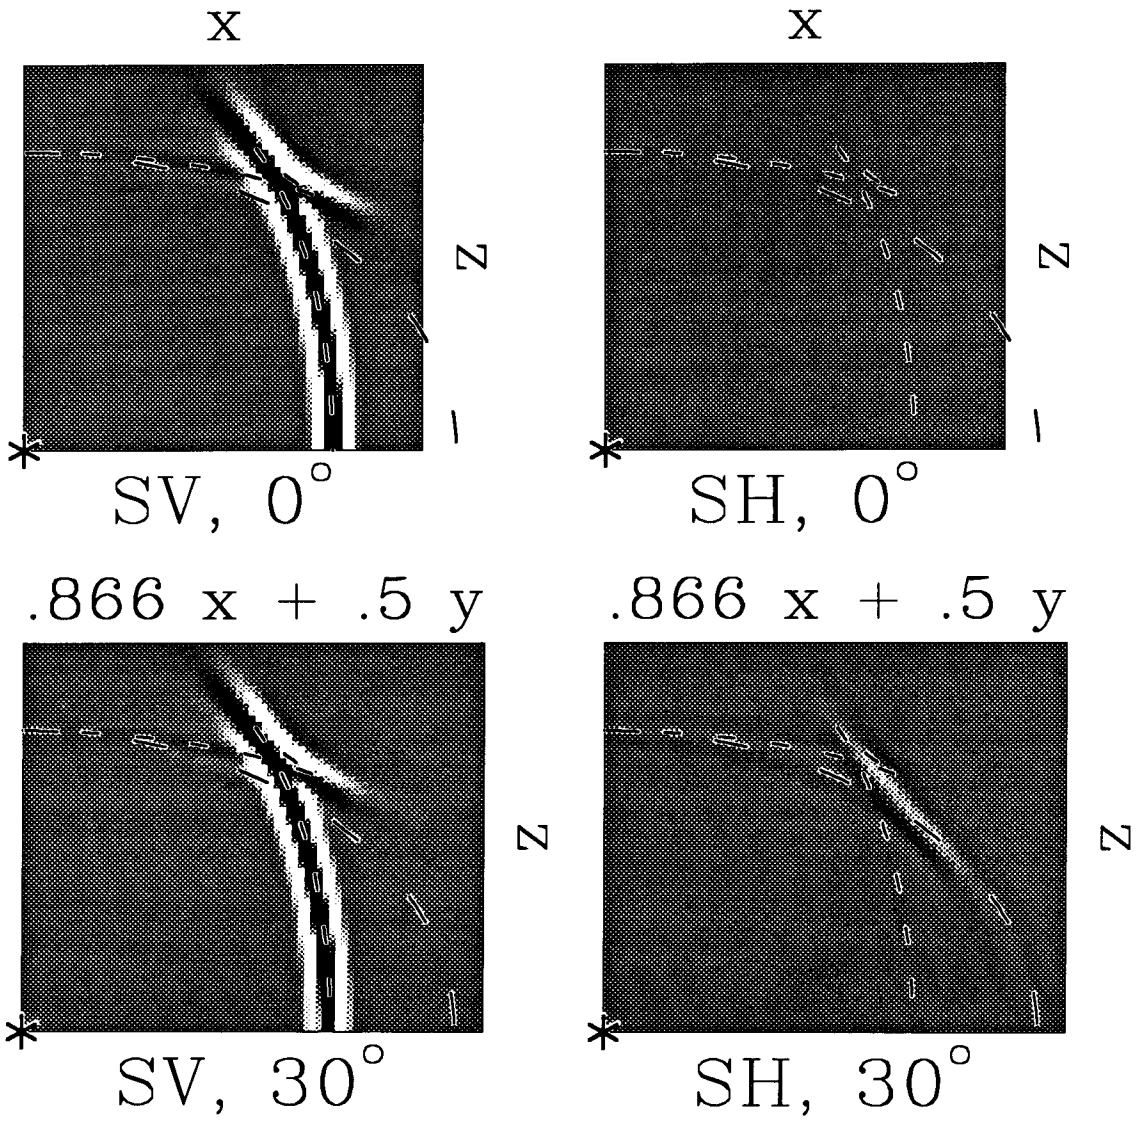

FIG. 3.18. Four slices through a three-dimensional Cracked Greenhorn Shale snapshot, corresponding to the medium shown in the middle left plot in Figure 3.16. (The elastic constants are listed in Table C.2.) The parameters of the four plots are the same as for the corresponding plots in Figure 3.17; the only difference is the underlying elastic constants have been perturbed to break the axisymmetry. The same Greenhorn Shale theoretical impulse-response curves used in the previous figure are overlayed again here. This time the elastic constants are somewhat different, so there's no guarantee they'll fit as well as they did in the previous figure. For the 0° (x-z) slice the overlayed curves fit exactly anyway, and for the 30° SV-component plot they are only slightly off. For the 30° SH-component plot, though, something unexpected happens. Instead of some sorts of "qSV" or "qSH" waves as we might expect, mostly we see a short nearly planar wavefront. (Its amplitude is about 15% of that of the primary "qSV" wave.)

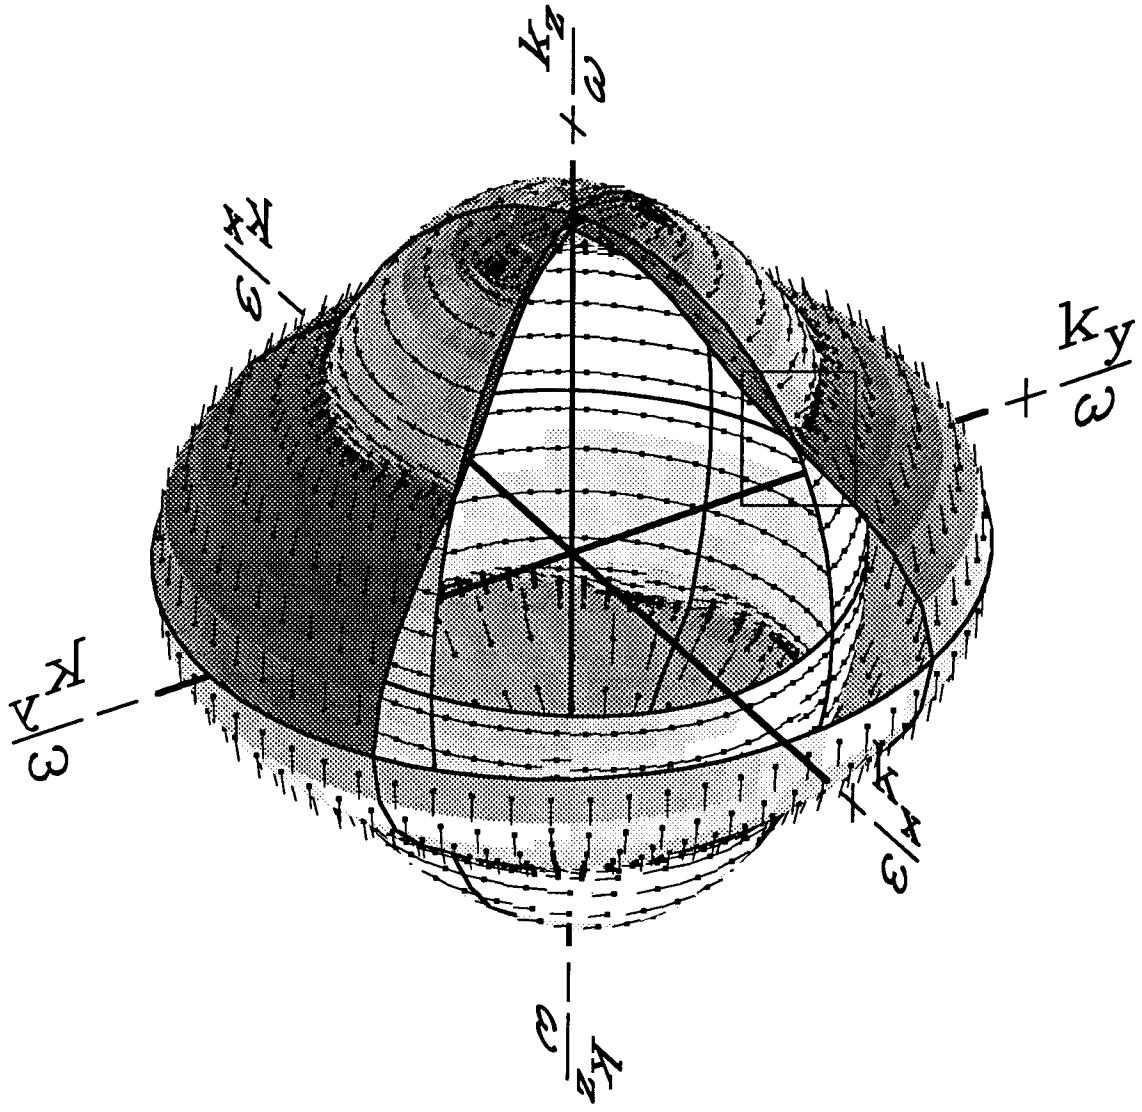

FIG. 3.19. The three-dimensional shear slowness surfaces for "Cracked Greenhorn Shale", which is just Greenhorn Shale perturbed by adding fractures in the x-z plane. Compare this figure with the unperturbed version, the top plot in Figure 3.6 on page 70. This time I have rotated the medium 60° so that the edges of the cut-out octant do not line up with the symmetry axes of the medium. At first glance, the topology appears the same as before. Figure 3.20 shows a high-resolution view of the region inside the rectangle, which shows the true topology. (The middle left plot in Figure 3.16 shows yet another view of this same three-dimensional surface.)

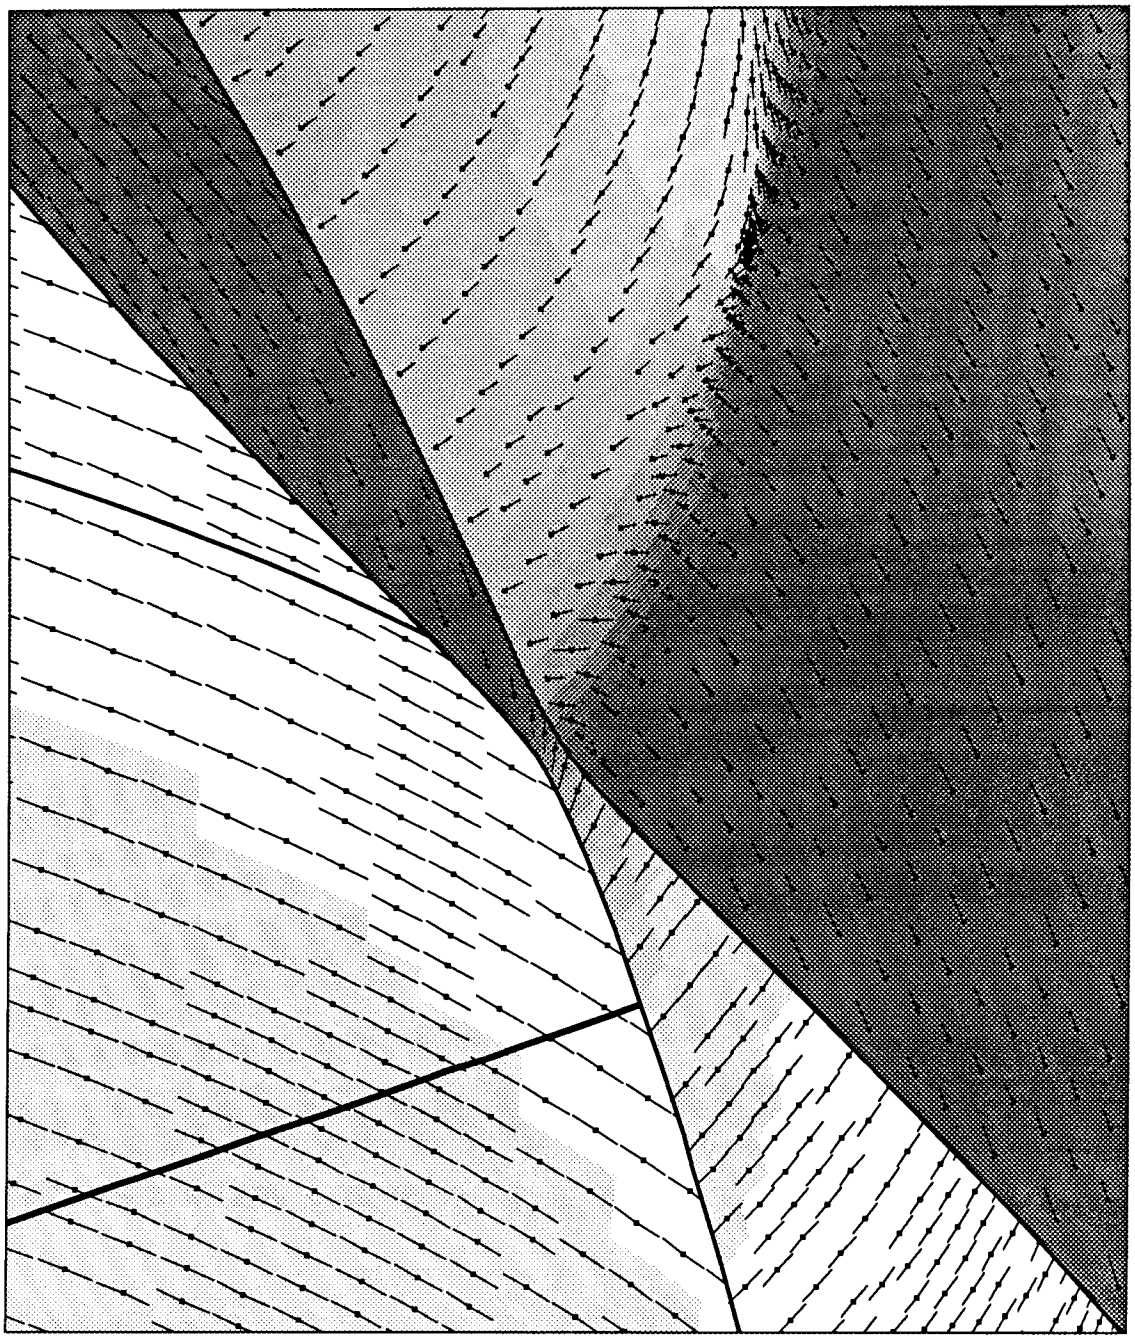

FIG. 3.20. A close-up view of Figure 3.19 showing the true topology. (Remember, I use the darkness of the surfaces to show the associated particle-motion direction, SV motion being darker and SH motion lighter.) Like the orthorhombic medium in Figure 3.10, this orthorhombic medium displays inner qS2 and outer qS1 surfaces that touch at a small number of discrete points (singularities). One such point, a singularity of order -1, is visible as the black dot a little to the right and below top center.

special case.[6]

"Connections"

So where does the mysterious event come from?

Figure 3.19 shows the three-dimensional slowness surfaces for the cracked medium. It appears very much like the unperturbed TI medium shown in the top plot in Figure 3.6, but there is a crucial difference, as can be seen clearly in Figure 3.20. The formerly distinct qSV and SH modes have crossed over and recombined to form new qS1 and qS2 modes. For the most part, these new shear surfaces can be thought of as virtually unchanged patched-together fragments of the old qSV and SH surfaces, and hence I like to call them "chimeras"[7].

This "exchanging identity without quite touching" phenomenon is quite common in physics. Spheroidal normal modes of the Earth are another good geophysical example. A plot of frequency versus angular order for these modes seems to show distinct sets of crisscrossing parallel lines. Upon close examination, however, the parallel lines prove to be an optical illusion formed by a series of noncrossing curves each shaped like a set of staircase steps. (See for example Figure 17 in Gilbert and Dziewonski (1975).)

This sort of behavior can be modeled using extremely simple mathematics; take for example the equation $y^2 = x^2$. We would probably write the solutions to this equation as

$$y = \begin{cases} x \\ -x \end{cases} . \tag{3.1}$$

However, if we used instead the perturbed equation $y^2 = x^2 + \epsilon$, we would write the solutions as

$$y = \begin{cases} \sqrt{x^2 + \epsilon} \\ -\sqrt{x^2 + \epsilon} \end{cases} . \tag{3.2}$$

[6]Thus we find that *all* two-dimensional anisotropy is inherently dishonorable! (Recall the epigraph on page 9.)

[7]I'm using a biological analogy here; the operative definition of "chimera" from Webster's (Woolf, 1975) is "an individual, organ, or part consisting of tissues of diverse genetic constitution and occurring esp. in plants at a graft union". (Although one of the other listed definitions, "an imaginary monster compounded of incongruous parts", might be more appropriate here!)

By letting $\epsilon \rightarrow 0$ we find an alternate solution set for the original equation,

$$y = \begin{cases} |x| \\ -|x| \end{cases} \quad . \tag{3.3}$$

For our cracked Greenhorn Shale example, "ϵ" is the perturbation of the Christoffel matrix $\underset{\sim}{\mathbf{C}}$ away from transverse isotropy. This "ϵ" effectively vanishes in the orthorhombic symmetry planes.

These topological details of solution sets are usually a mere curiosity, but for our perturbed transversely isotropic example they turn out to be significant. It is the novel parts of the qS1 and qS2 surfaces in Figure 3.20 where they approach but don't quite touch that create the new event seen in Figure 3.18. (Kawasaki (in preparation) similarly finds that for an azimuthally anisotropic earth such "chimeric" effects need to be taken into account when inverting earthquake surface-wave data.) To see how this relatively small change in the slowness surfaces can be significant, we will now examine the relevant three-dimensional *impulse-response surfaces*.

We begin with Figure 3.21, which shows the three-dimensional shear impulse-response surfaces for unperturbed Greenhorn Shale. The elliptical SH mode in three dimensions becomes a prolate ellipsoid, while the cusped qSV mode becomes a sort of flanged barrel. Figures 3.24 and 3.22 show the corresponding qS1 and qS2 modes for Cracked Greenhorn Shale. The thick lines correspond to the thick lines in Figure 3.19. Note that although they lie in a plane in the slowness domain, they definitely do not in the group domain. (I would have shown both surfaces together, but found the resulting tangle of intersecting surfaces too complex to visualize easily. To give some idea of the relative positions of the qS1 and qS2 surfaces, I have shown the thick lines for both surfaces in both sets of plots.)

The impulse-response surfaces shown in Figures 3.24 and 3.22 combine gross elements of the surfaces visible in Figure 3.21, but some new features are also visible. The surface shown in Figure 3.24 is adorned with thin disks hovering just above the remainder of the top surface, and attached only at a relatively small region at the center. The disks fit into the holes visible in Figure 3.22. The outer rims of the disks (and the inner edges of the holes) correspond to the point singularities visible in Figure 3.19. (The outer surface of the disks are usually called "lids" in the literature (Crampin (1981), but there seems to be some confusion about what the term "lid" precisely means. I will avoid using the term until after I have given better examples of the topology of the singularities, notably

in Figures 3.30, 3.33, and 3.34.)

Another apparent difference is the larger cusps in Figure 3.24. The topology of the two intersecting shear surfaces is complex; Figure 3.26 shows where the extended cusp comes from more clearly. The concavity visible on the qS2 surface in Figure 3.20 corresponds to a cusp on the associated impulse response, just like the one on the qSV surface of Greenhorn Shale did in Chapter 2. This cusp comes from a small region of the slowness surface, and so should not be too significant energetically. However, this cusp is exactly what we were seeing back in Figure 3.18; Figure 3.27 shows this more clearly. How is this possible?

Figure 3.26 very much appears like the unperturbed qSV and SH surfaces of Greenhorn Shale, plus a new event I will call a "connection". The connection comes from the ring along which the original qSV and SH surfaces intersected, but where the new qS1 and qS2 surfaces merely appulse in a "ring pinch" (except at the few point singularities) (Crampin and Yedlin, 1981). Points in the phase-slowness domain correspond to plane-wave components in the impulse-response domain, so the connection lies where there was a common plane-wave component of the original qSV and SH modes (i.e. it is tangent to both).

A connection can have a significant effect by providing a channel for coupling the ersatz "qSV" and "qSH" modes. In the Cracked Greenhorn Shale example used here, the connection roots into the qSV mode in the middle of the high-amplitude cusp, and acts to provide a channel along which some of this energy can drain away into the "qSH" mode. Since at the same time the particle-motion direction also slowly rotates to include more and more SH motion, it manages to show up strongly on the SH snapshot slice. This sort of anisotropic behavior provides an ideal vehicle for coupling a vertical P source into a horizontal SH receiver, even for media that otherwise are virtually transversely isotropic. Note that connections should only be significant well off of the symmetry planes. To me this suggests that deliberately trying to record seismic data only along suspected symmetry planes is a questionable practice.

Ray tracers should note that while the SH component in the 0° symmetry plane in Figure 3.18 is exactly zero for all wavetypes, as it must be by symmetry, the SV component of the connection in the symmetry plane is nonzero (although quite weak). This event is in fact a non-ray wave. Unfortunately the connection on the 0° SV plot in Figure 3.18 is not visible on the hardcopy, although it is dimly visible on a good screen.

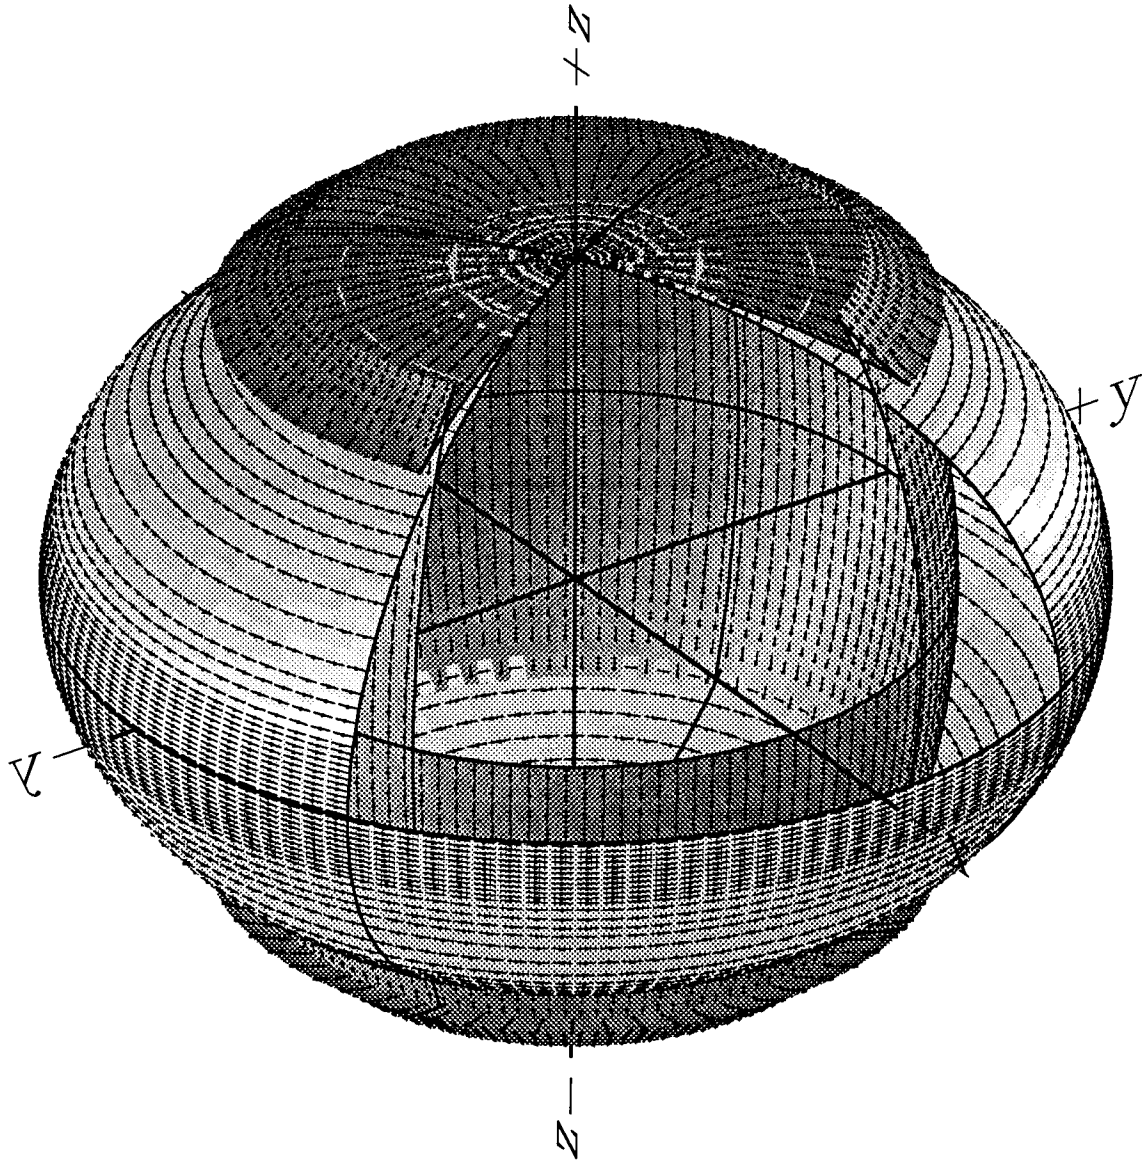

FIG. 3.21. The three-dimensional shear impulse-response surfaces for Greenhorn Shale (Jones and Wang, 1981). The figure shows a light-colored prolate-ellipsoid pure SH mode intersecting a flanged barrel-shaped qSV mode.

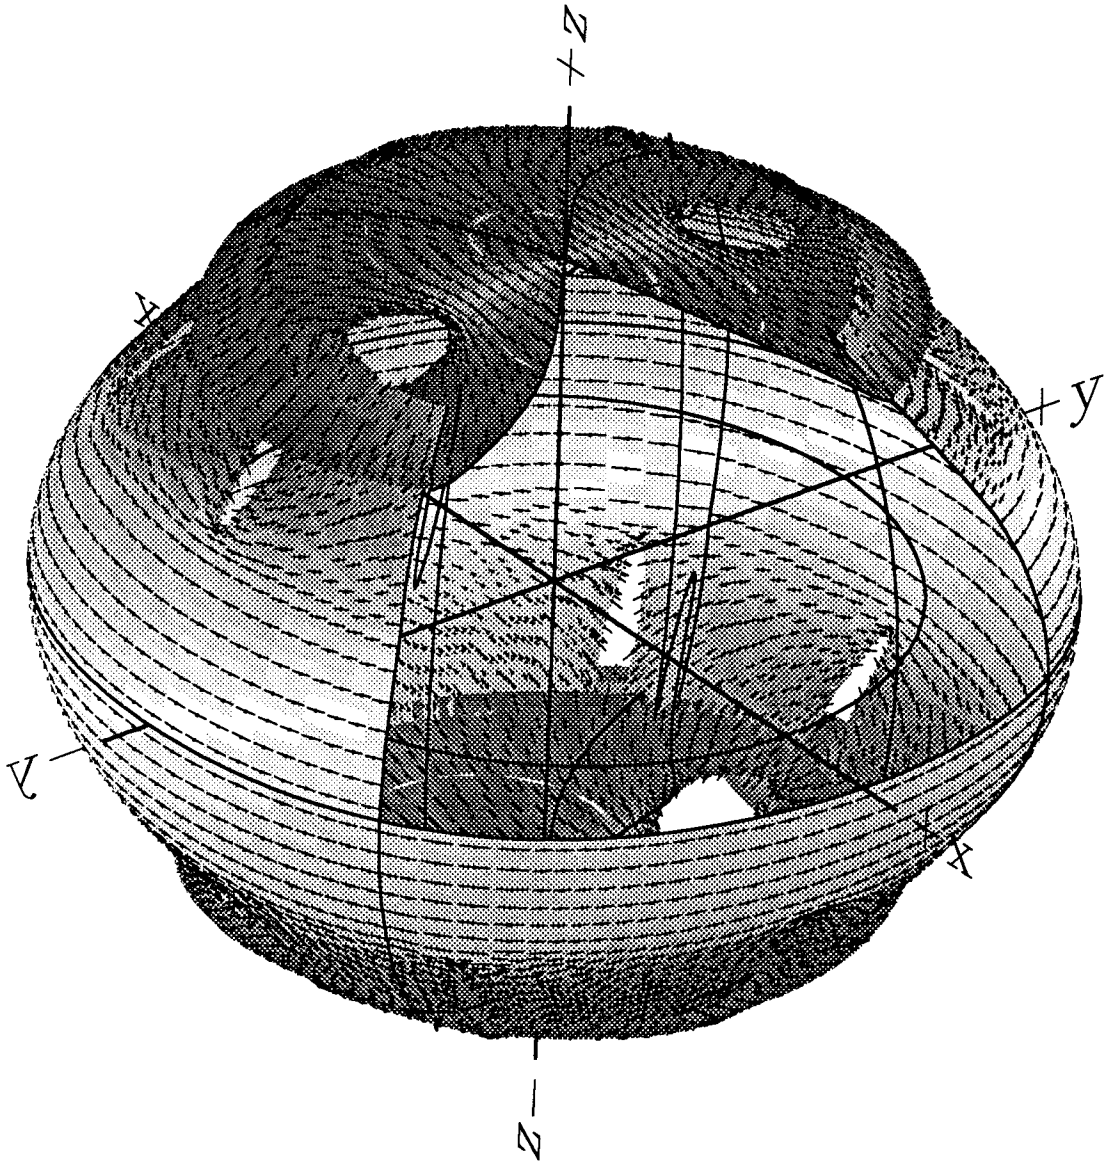

FIG. 3.22. The three-dimensional qS1 impulse-response surface for Cracked Greenhorn Shale corresponding to the slowness surfaces shown in Figure 3.19. This figure is the left element in a stereo pair; Figure 3.23 is the corresponding right element. The isolated thick lines show the position of the omitted qS2 surface.

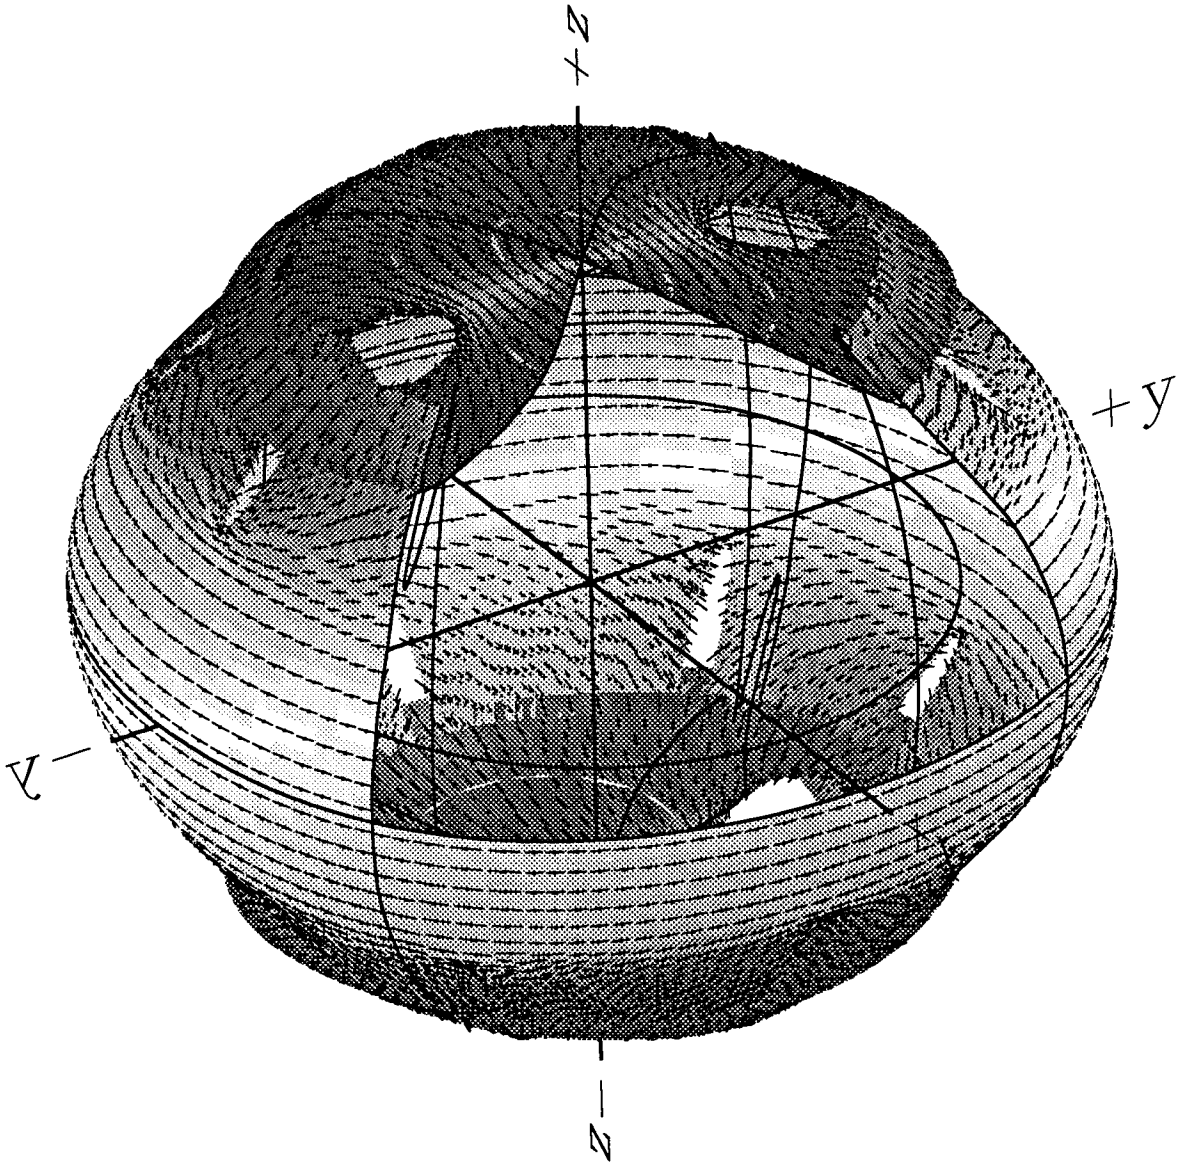

FIG. 3.23. The three-dimensional qS1 impulse-response surface for Cracked Greenhorn Shale corresponding to the slowness surfaces shown in Figure 3.19. This figure is the right element in a stereo pair; Figure 3.22 is the corresponding left element. The isolated thick lines show the position of the omitted qS2 surface.

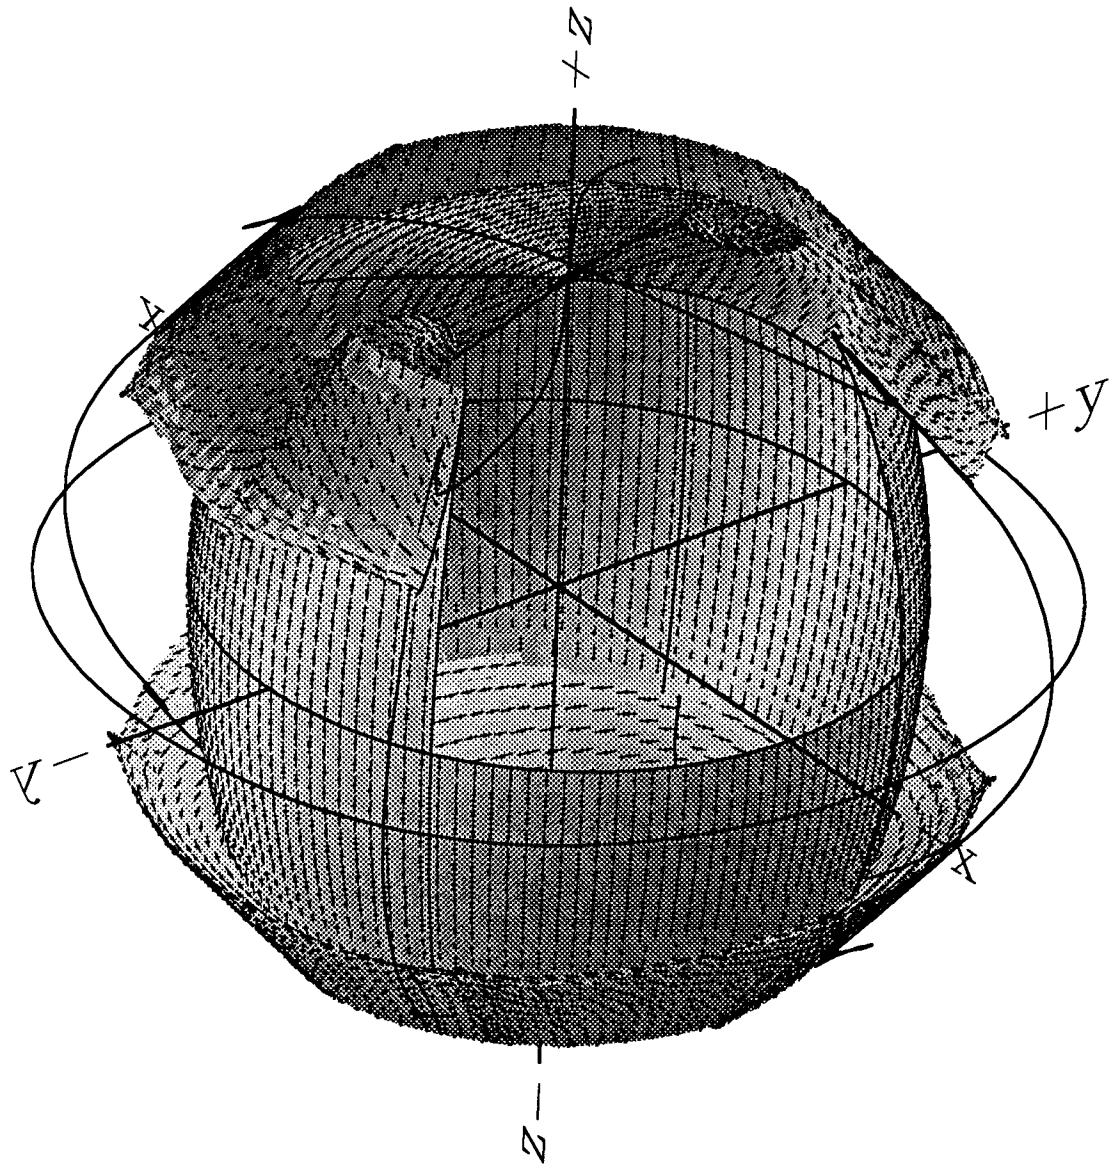

FIG. 3.24. The three-dimensional qS2 impulse-response surface for Cracked Greenhorn Shale corresponding to the slowness surfaces shown in Figure 3.19. This figure is the left element in a stereo pair; Figure 3.25 is the corresponding right element. The isolated thick lines show the position of the omitted qS1 surface.

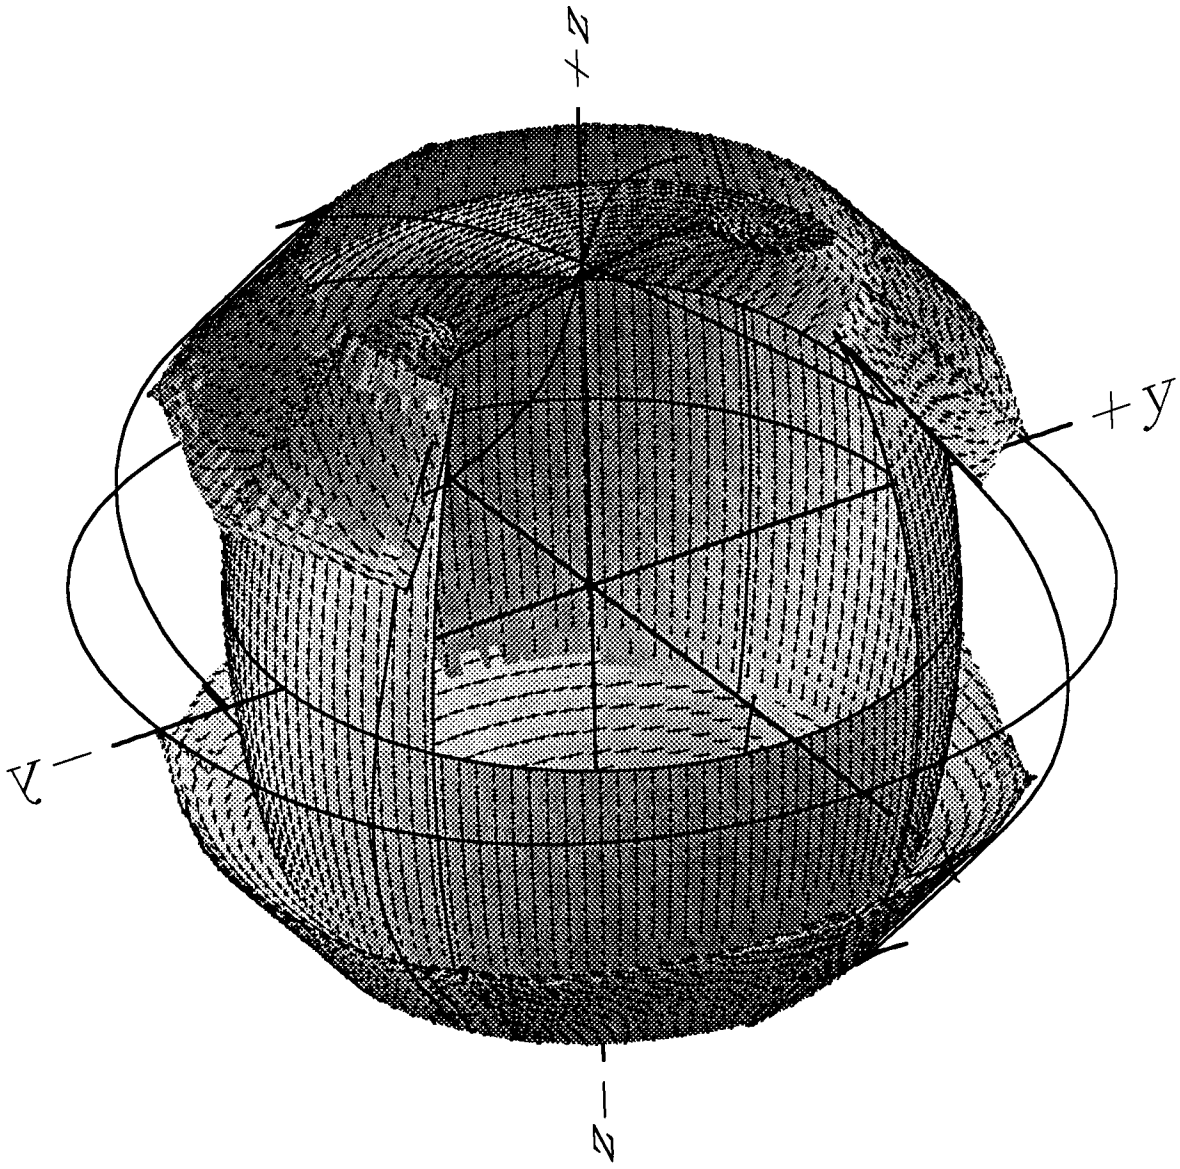

FIG. 3.25. The three-dimensional $qS2$ impulse-response surface for Cracked Greenhorn Shale corresponding to the slowness surfaces shown in Figure 3.19. This figure is the right element in a stereo pair; Figure 3.24 is the corresponding left element. The isolated thick lines show the position of the omitted $qS1$ surface.

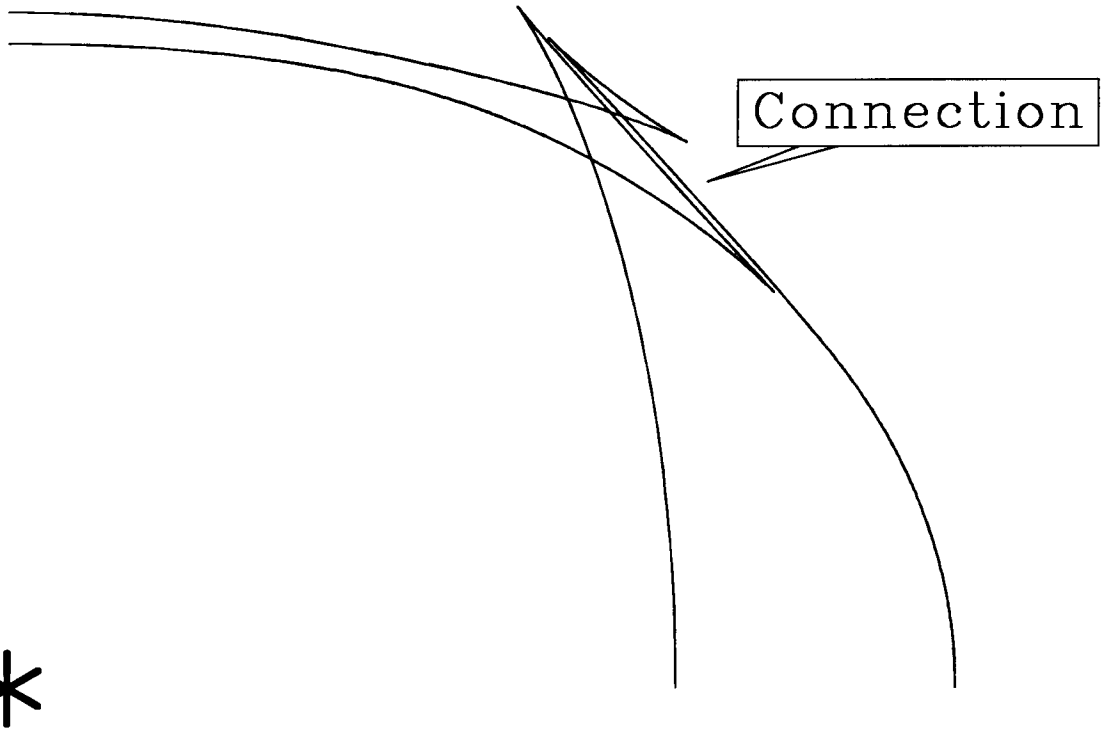

FIG. 3.26. A slice through the 3D impulse response surface for Cracked Greenhorn Shale, oriented to correspond to the 30° slices in Figure 3.18. Topologically there are distinct qS1 and qS2 surfaces, but the effect is very much like that of qSV and qSH surfaces with a "connection" spanning the tangents to those surfaces. This figure also illustrates a couple of other interesting counterexamples. Note that the qS2 wave arrives *first* for a small range of angles near the upper edge of the old qSV cusp, even though by definition the qS2 mode has *slower* phase velocity for all phase directions. The connection itself consists of two closely spaced linear events; the two parts are portions of the qS1 and qS2 modes. Particle-motion directions for a given phase direction are guaranteed to be perpendicular for distinct modes. This is usually almost true in the group domain as well. For this example, however, at the top of the connection, where it roots into the former qSV mode, both modes have nearly pure SV polarization. At the bottom of the connection, where it roots into the former SH mode, both modes have nearly pure SH polarization. (In between, the particle motions perform the transition by rotating in opposite directions.)

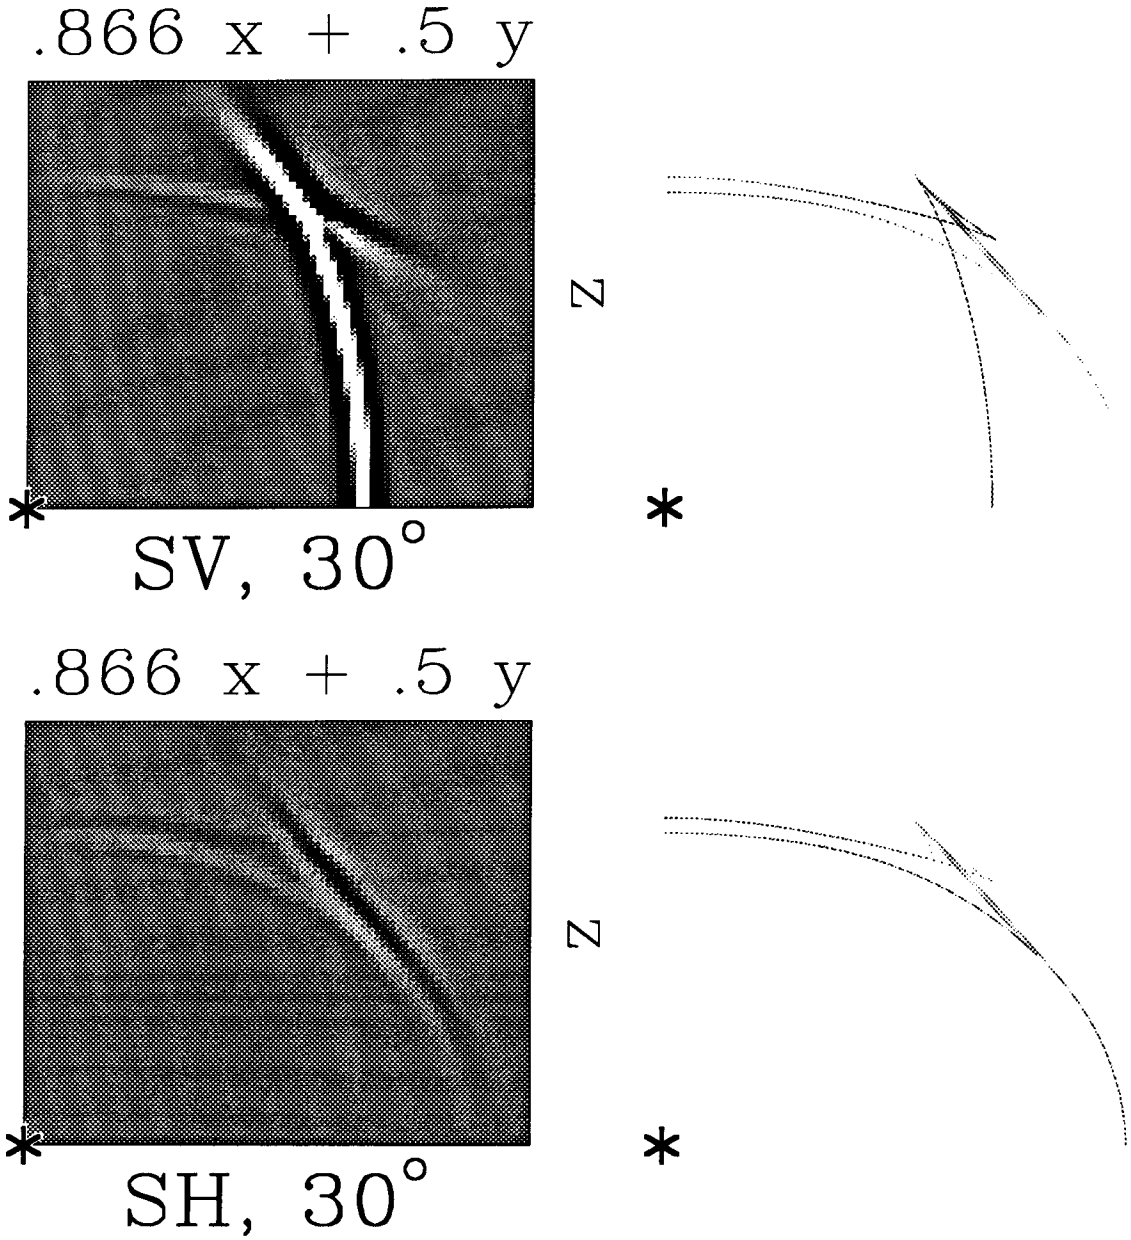

FIG. 3.27. A reprise of Figure 3.18, this time showing the correct impulse-response surfaces for Cracked Greenhorn Shale (on the right). A gpow of .7 on the snapshots (left) has been applied to make weaker events more visible. An attempt has been made to show the proportion of the theoretical modes' particle-motion direction aligned with the pure SV and SH directions by shading the curves appropriately (although the shading has somewhat run afoul of the halftone dithering pattern).

3.5.2 Anatomy of a singularity

We have seen that singularities are important topologically, and that in certain cases the topology of the wave modes case create significant events such as "connections". But can the singularities themselves cause any significant propagation effects? I begin by examining the topology of point singularities in the impulse-response domain.

A canonical orthorhombic example

Figure 3.28 shows the qS slowness surfaces for a simple orthorhombic medium. This medium has four widely spaced singularities of order $+1$, all lying in the k_y-k_z plane. For the example here I have rotated the medium about the k_x axis to line up one of the singularities with the $+k_z$ axis. Figure 3.29 shows the corresponding impulse-response surfaces. The geometries of the slowness and impulse-response surfaces seen here seem to be much the same: two surfaces intersecting at a point forming a shallow double cone. In one case the center of the cone is on the k_z axis; in the other the center is off to one side.

In either domain one could imagine cutting the inner and outer surfaces apart at the intersection point and labeling those as the two qS pure modes. Unfortunately the resulting classification would not be consistent between the two domains, because the intersection point in the group domain does not correspond to the singularity. What on the impulse-response surfaces does correspond to the singularity?

As is shown in Appendix B, the group direction is perpendicular to the tangent to the slowness surface at each point. As can be seen in Figure 3.28, at the singularity the tangent becomes discontinuous. In fact the point singularity in the slowness domain maps to a cone of directions in the impulse reponse. (This phenomenon is well known in crystals, where it can be directly observed as the phenomenon of internal conical refraction (de Klerk and Musgrave, 1955). For our example in Figure 3.28, one of the possible tangents at the singularity is for a constant-velocity surface (circular in two-dimensional cross section). Since for that case the group and phase directions are the same, the $+z$ axis in the impulse-response domain in our example must mark one point along the cone of singular directions. The singularity in Figure 3.29 is the annular gap containing the $+z$ axis.

In fact there is more to the singularity than that. As shown by Burridge (1967), the "dimple" on the impulse-response surface is made convex by the additional presence of a planar event stretching across the concavity. This event, called a "plane lid" by Burridge,

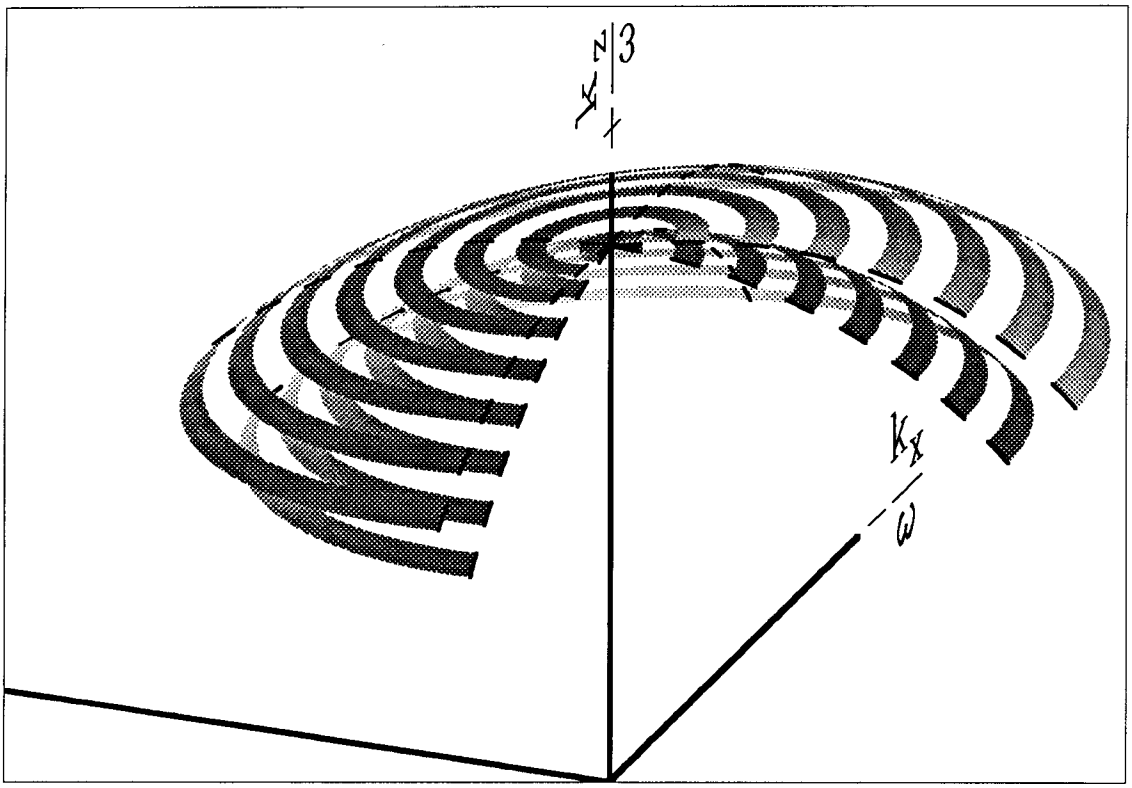

FIG. 3.28. The qS slowness surfaces for a particularly symmetric orthorhombic medium. This medium has a global pure P mode and pure SH and SV modes in each symmetry plane. In each symmetry plane one of the shear modes is elliptical and the other is circular. The medium has been rotated about the x axis to position one of the singularities on the $+k_z$ axis; thus, only the k_y-k_z axial plane is still one of the three planes of symmetry. In order to show the topology of the singularity more clearly, only alternate latitude bands are shown, and those only from latitudes 45° to 90°. One 90° sector has been omitted to allow a view of the singularity.

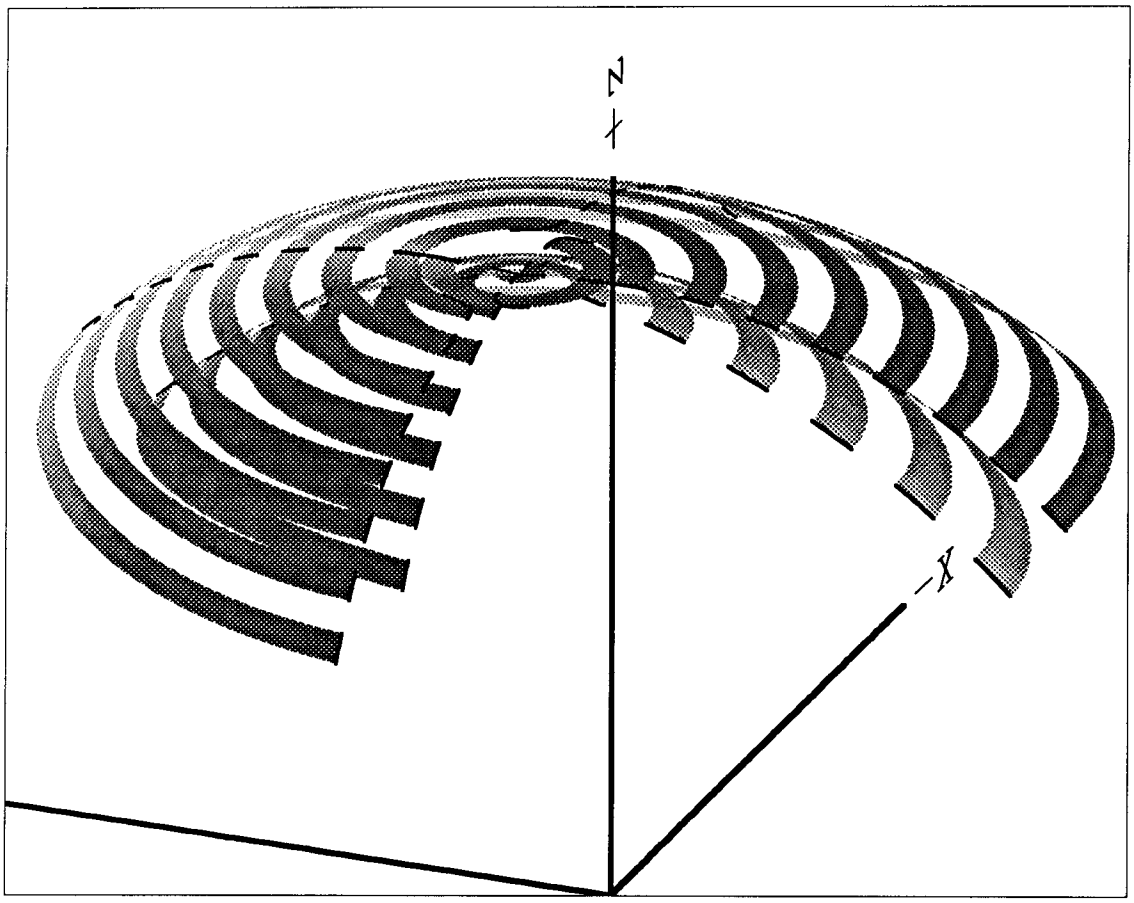

FIG. 3.29. The qS impulse-response surfaces corresponding to the slowness surfaces in Figure 3.28. The point singularity in the previous figure corresponds to the annular gap containing the $+z$ axis. (The line leaving the figure to the left is the $-y$ axis.) Note that the dark axial lines marking the k_x-k_z plane abruptly bend as they approach the annulus marking the location of the singularity. This is because the singularity in the slowness domain was a fraction of a degree out of the k_x-k_z plane. The axial lines do not cross the ring of the singularity in the group domain but detour around it instead.

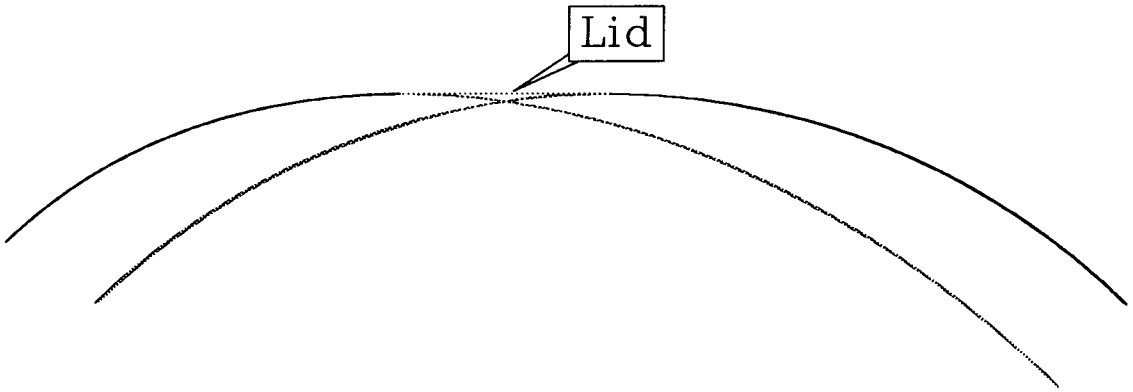

FIG. 3.30. A cross section through the y-z plane of the surface shown in Figure 3.29. The qS1 wave mode is shown by a black line, the qS2 wave mode is shown by a dark stippled line, and the plane lid by a thin dotted line.

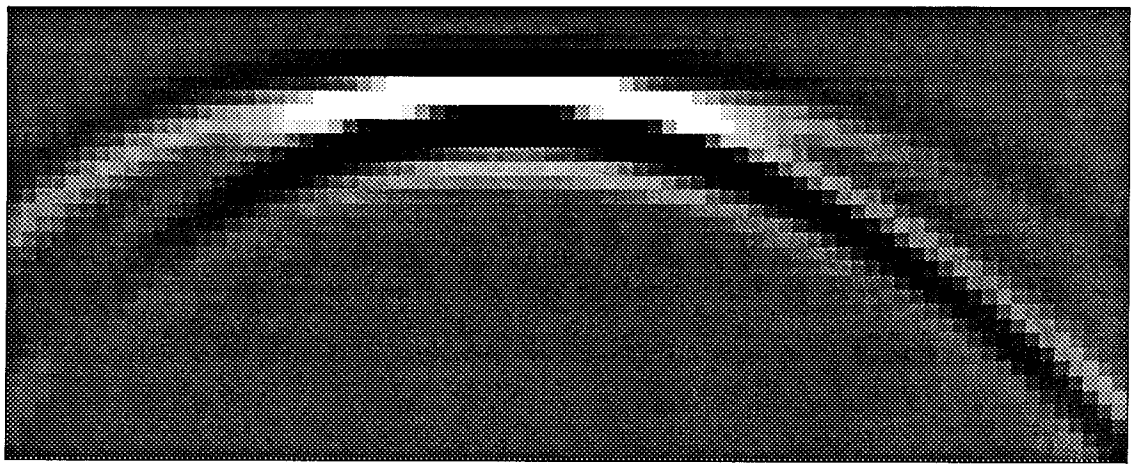

FIG. 3.31. The finite-difference model result corresponding to Figure 3.30. The modeled wavefield was generated by a y point source, but the x component of particle motion is shown. This emphasizes wave components with intermediate polarization directions, such as occur around the singularity. The slice is displaced slightly out of the plane containing the source (in the symmetry plane the x component is zero).

is the plane-wave component of the impulse response corresponding to the singularity. The cone of singular directions previously described marks the edges of the plane lid, in this case shaped like a circular disk. Figure 3.30 shows the position of the plane lid on a cross section through the y-z symmetry plane of the impulse-response surface. (The plane lid is omitted in Figure 3.29.) Figure 3.31 shows a corresponding cross section through a finite-difference model snapshot.

A triply connected example

Burridge (1967) found that the singularities of cubic nickel had "plane lids", and derived their amplitudes. Later Crampin (1981) pointed out that the "lids" of orthopyroxene are shallow cones, not planar. Are they both right?

Figure 3.32 shows all three impulse-response surfaces for a complex orthorhombic medium. This medium is interesting in several ways. All three wave modes are anomalously polarized. All three wave modes are connected by singularities; there is no separable qP mode. In the next few figures I will examine the properties of the singularity in the y-z symmetry plane on the *fastest* wavetype; it is particularly amenable to study since it is part of the first-arriving wavefront.

Figure 3.33 shows an annotated slice of the impulse-response surface in Figure 3.32 along the y-z symmetry plane. Note this example shows two distinct kinds of "lids": Burridge and Crampin were both right, but were talking about different things. I will follow their lead by calling one kind a "plane lid" versus just plain "lid" for the other.

I define a *plane lid* as the portion of the impulse-response surface corresponding to the point singularity in the slowness domain. It is a planar surface that stretches across what otherwise would be a concavity in the impulse-response surface, and is bounded by the cone of singular directions associated with the singularity. Its presence keeps the surface convex. I define a *lid* as the outermost part of the impulse-response surface within the region bounded by the cone of singular directions associated with the point singularity not counting the "plane lid" generated by the singularity itself. This part of the impulse-response surface will be concave, and although faster in the group domain will be part of the slower of the two wavetypes in the phase-slowness domain. Some highly symmetric media, such as Burridge's cubic nickel example and the example in Figure 3.28 do not exhibit a lid because it has contracted to a point in the slowness domain and merged with the plane lid.

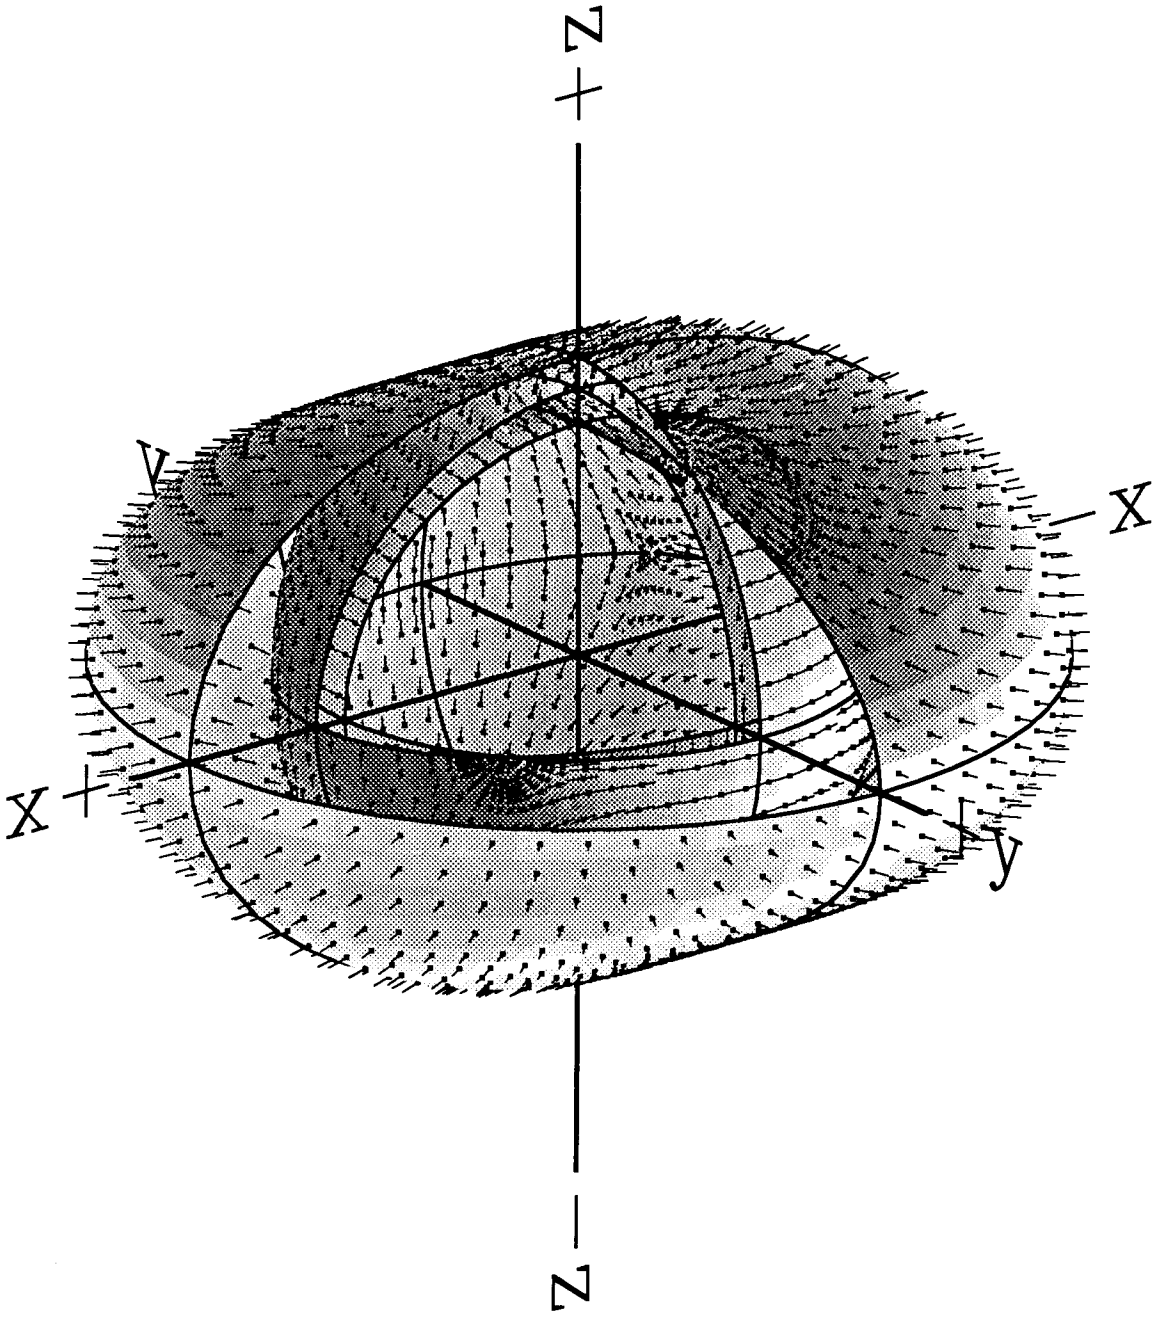

FIG. 3.32. Impulse-response surfaces for a "triply connected" anomalously polarized orthorhombic medium. All three surfaces have singularities, even the first-arriving one. Since the surfaces are anomalously polarized, the standard anisotropic notation "{qP, qS1, qS2}" does not apply. A general notation proposed by Muir (1989) extends the original meaning of "P" and "S" to sort the three wavetypes noncommittally by phase velocity as "{*P, S, T*}", for *primero*, *secundo*, and *tertio*.

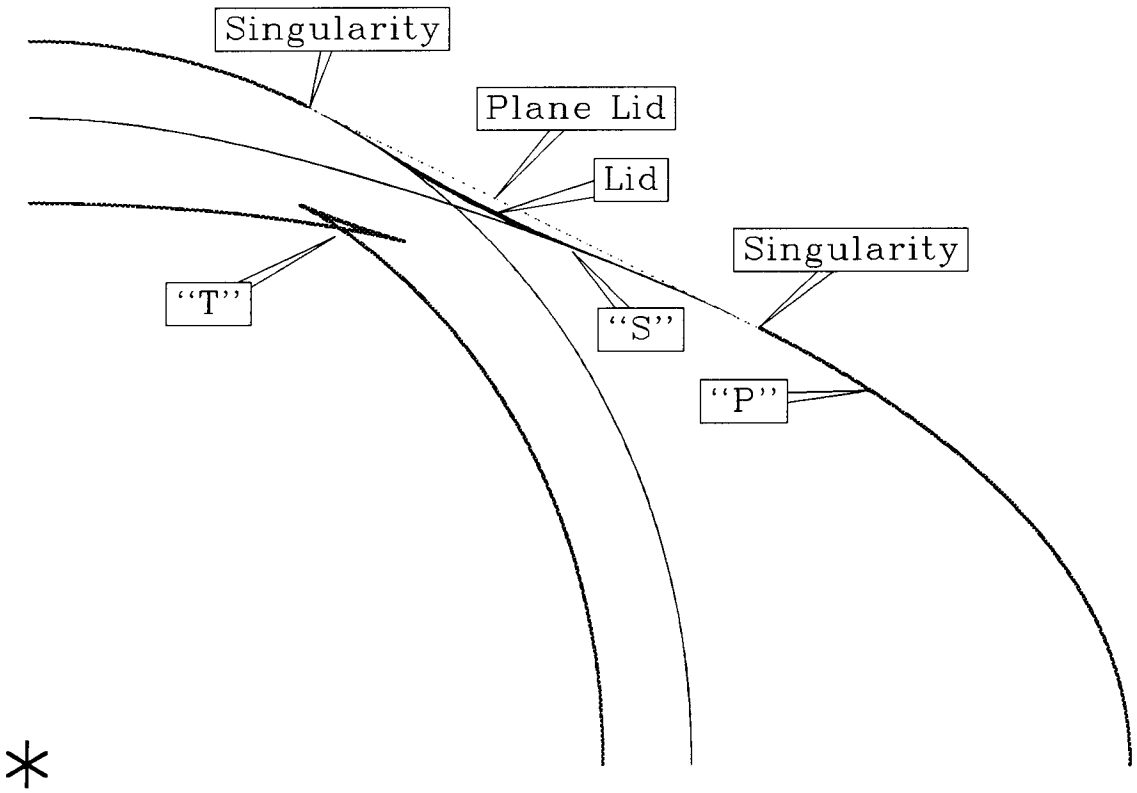

FIG. 3.33. An annotated slice through the y-z plane of the surfaces in Figure 3.32. The three surfaces are labeled following Muir's $\{P, S, T\}$ convention. The P and T surfaces are drawn thick and stippled, while the intermediate S surface is drawn thin and black. The lid is part of the S wavefront. (The plane lid is omitted in Figure 3.32, and is shown here with a thin dotted line.) Note the two points marked "Singularity" correspond to the same point singularity in the slowness domain.

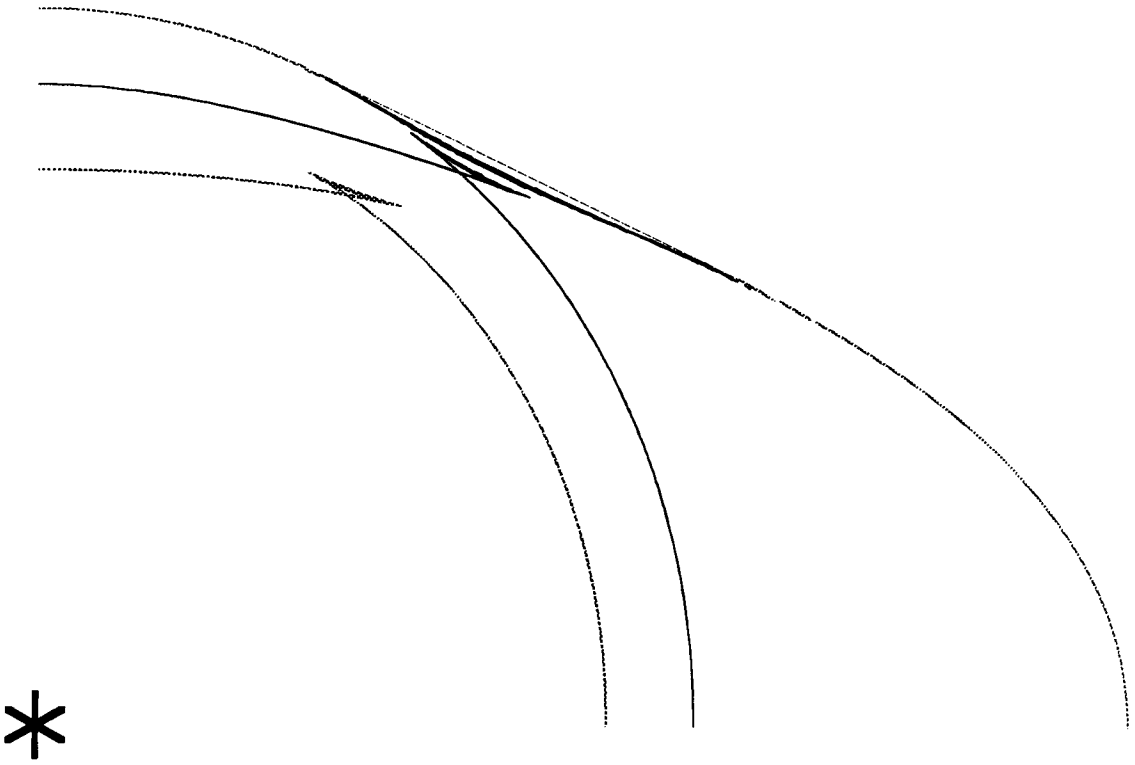

FIG. 3.34. As Figure 3.33, but rotated 5° out of the y-z symmetry plane to better show the true topology. The intermediate-velocity S surface appears as two disjoint pieces in this slice, although in three dimensions the parts form one continuous surface that intersects itself in the y-z plane.

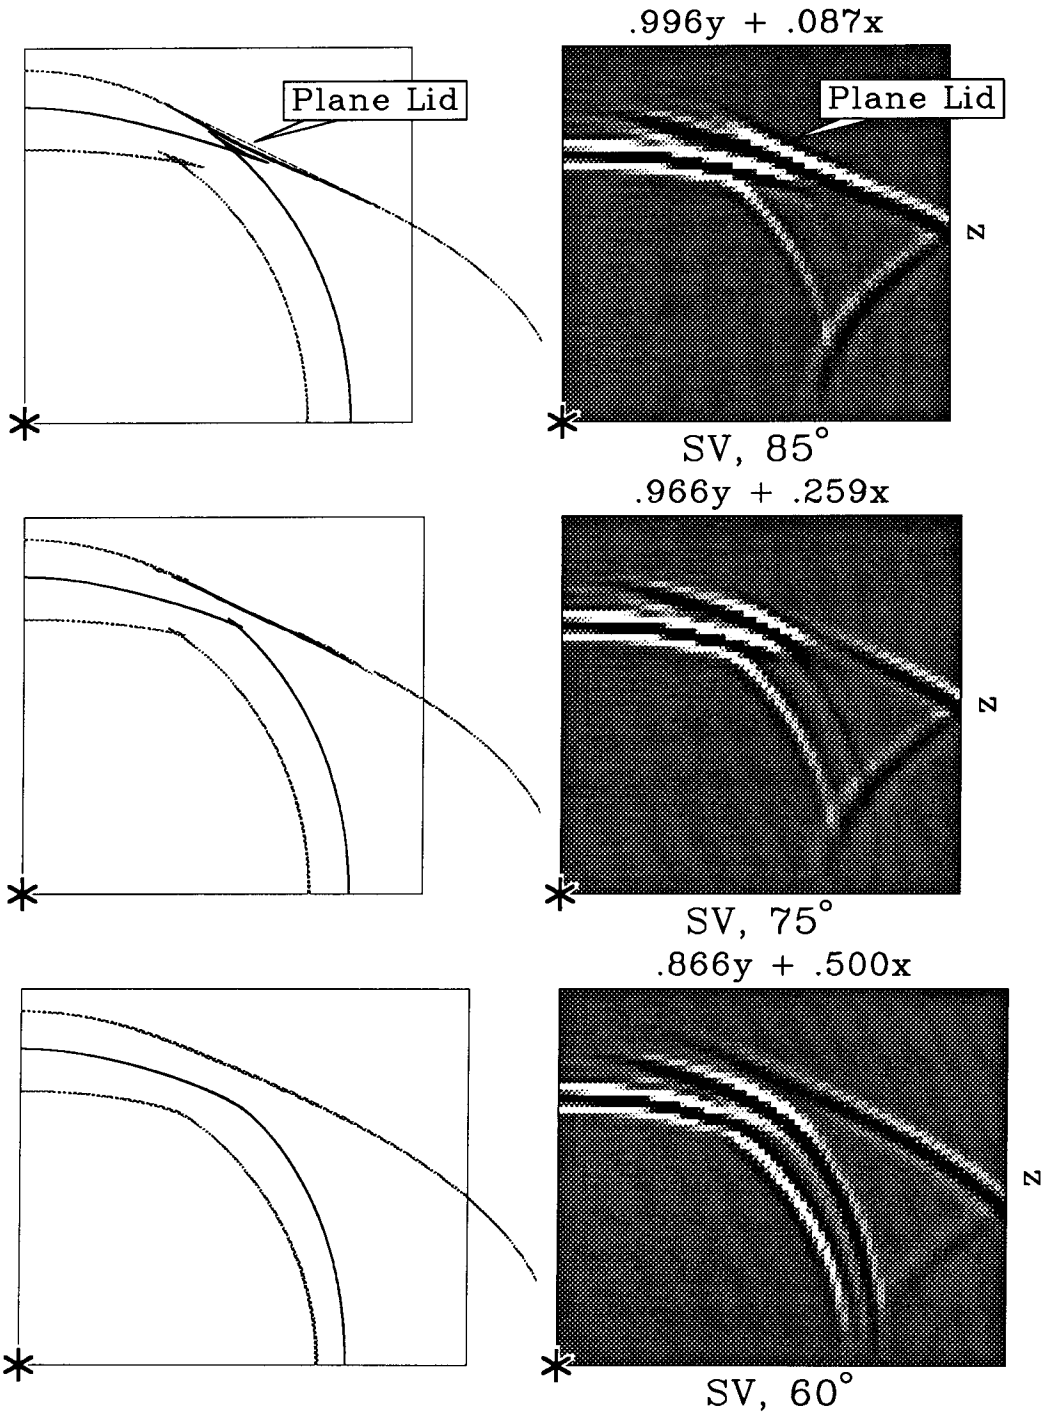

FIG. 3.35. Various non-symmetry-plane slices through the impulse-response surface in Figure 3.32 with corresponding finite-difference results. The source for the finite-difference model is a y point force. The SV component of the particle-motion direction is shown. The plane lid also shows exists in the longitude 75° slice, but is too close to the lid to be visible in either representation.

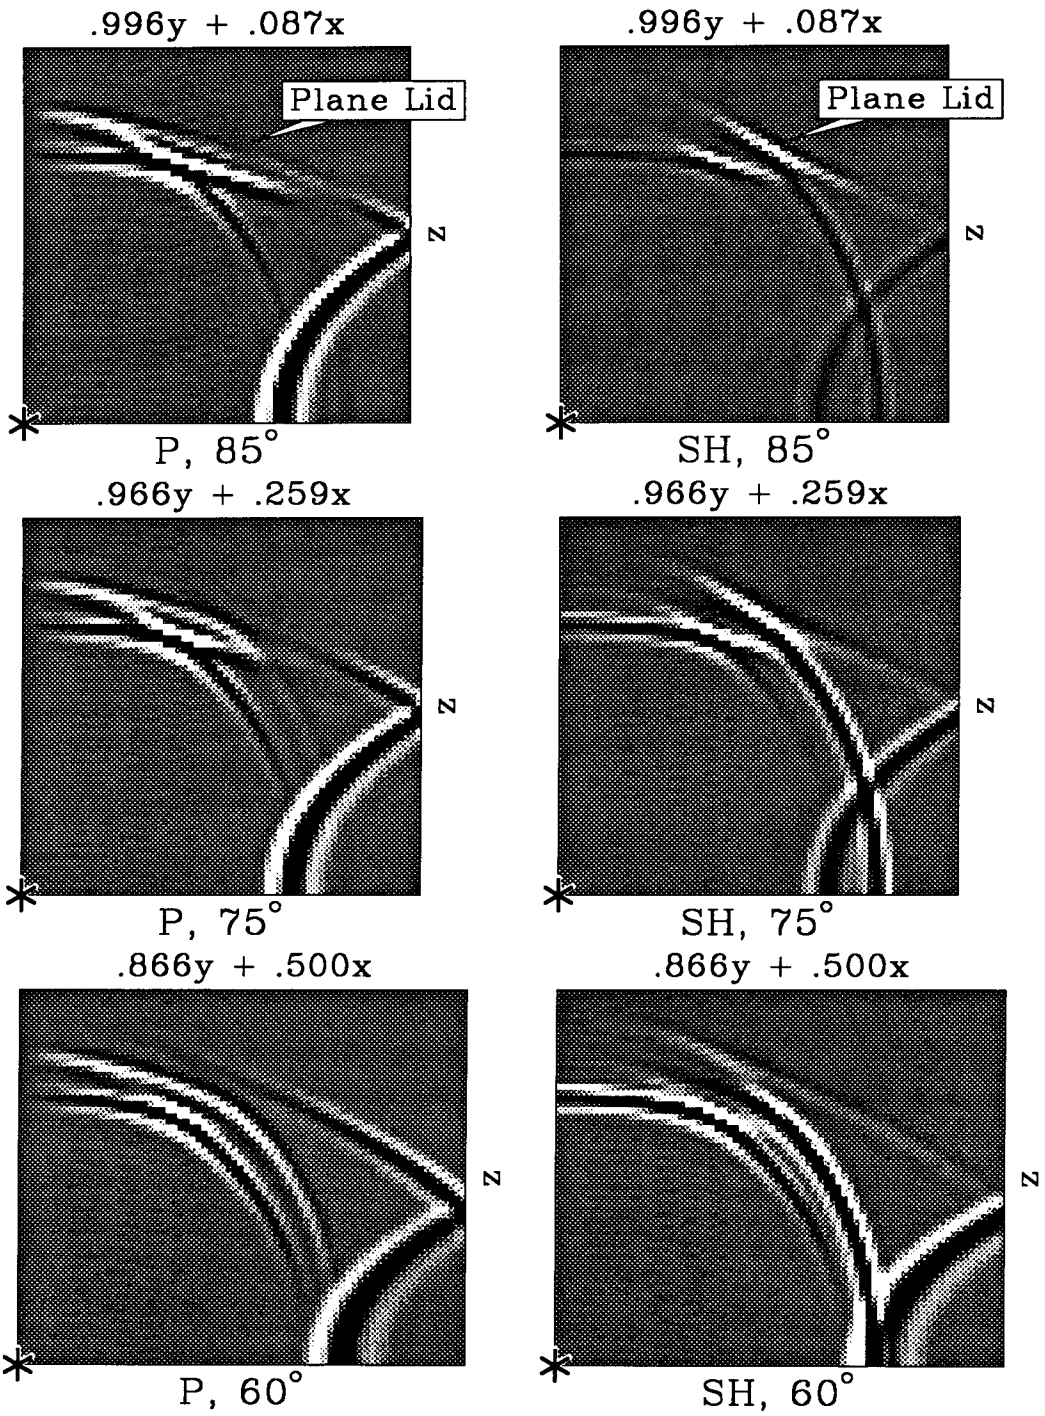

FIG. 3.36. Finite-difference results corresponding to three non-symmetry-plane slices through the impulse-response surface in Figure 3.32. The slices are the same ones used in Figure 3.35, but for this plot the P and SH components of the particle-motion direction are shown.

What about the three-dimensional topology of the singularity shown in Figure 3.33? In the example in the previous section (Figure 3.29), we saw that there was a point where the impulse-response surfaces crossed through themselves forming a shallow double cone. This was an accident of the high symmetry of that example; Figure 3.34 shows the true topology of the surfaces in Figure 3.32 revealed by cutting on a non-symmetry plane. The intermediate S surface in this example intersects itself in a line, not a point.

Finite-difference model results

The lid and plane lid are distinguishable in Figure 3.34, but can they also be distinguished in a finite-difference model result? Figure 3.35 shows several non-symmetry-plane slices of the impulse-response surfaces in Figure 3.32 along with the corresponding finite-difference snapshot slices. The top right plot shows a dark (−) event preceding the main lid (which appears white (+) here) and filling in the concavity in it. This should be the plane lid, since it occurs in the right place and seems to exhibit the necessary properties, but the result needs to be backed up by further analytical calculations to be sure. Figure 3.36 shows the other two particle-motion components of the finite-difference model results at the same clip levels as in Figure 3.35. The plane lid exists in these plots as well, but is too weak to show up on these plots.

Can singularities cause visible effects?

I still have not answered the original question of this section: can the singularities themselves cause significant propagation effects?

In Figures 3.35 and 3.36 the direct effect of the singularity, the plane lid, was a subtle feature of the wavefront. On the other hand, the singularity showed up very strongly in Figure 3.31. In that example the y source and x receiver are at right angles; the swirl of particle-motion directions around the singularity provides a good coupling mechanism and makes the singularity stand out.

These examples indicate that direct effects such as plane lids are not very significant, but that the local disturbance to the particle-motion direction field caused by a singularity can be. This latter effect will show up again in the next section.

3.5.3 What should 3D anisotropy look like?

So far in this chapter I have shown only model snapshots. In this section I will show some seismograms to go with two of the models examined in the previous section. This is meant to give some idea how the effects previously noted might appear on surface data.

What should a connection look like?

Figure 3.37 shows model sections recorded over a homogeneous half-space of "Cracked Greenhorn Shale" (described in section 3.5.1) at two different azimuths, along the cracks (azimuth 0°) and at an acute angle to the cracks (azimuth 30°). The model and plot parameters of this plot are the same as those for Figure 2.21 on page 48, although the example in Chapter 2 is not directly comparable. That model medium had a different value of C_{13} and so does not triplicate as strongly as the example here.

The diagnostic orthorhombic effect, the "connection", shows up strongly on the "SH, 30°" plot. The other three plots, however, reveal very little evidence that the model medium is not transversely isotropic Greenhorn Shale. The connection event also nearly disappears for sections oriented along azimuths 90° and 45°. This suggests that to observe orthorhombic effects in the field it is necessary to record both along suspected symmetry directions as inferred from the regional crack directions and at other arbitrary angles unlikely to correspond to any special symmetry. It is also beneficial to record the "null" sections such as Z source into SH receiver. These "off diagonal" sections are most likely to show energy coupling from one mode into another via the particle-motion direction disturbance associated with singularities.

What should a singularity look like?

Figures 3.38 and 3.39 show model sections recorded over a homogeneous half-space of the medium from Figure 3.29 (described in section 3.5.2). This time instead of showing sections recorded along lines passing over the source at differing azimuths, a series of sections recorded along parallel lines are shown. Instead of showing off-diagonal sections, for all plots both the source and receiver are oriented along the y axis.

Figure 3.38 shows lines parallel to the x axis. (Since the source and receiver are oriented along the y axis, these could be called "SH" lines.) Four different lines are shown. They were selected to cut through different parts of the cone of singular directions. In terms of

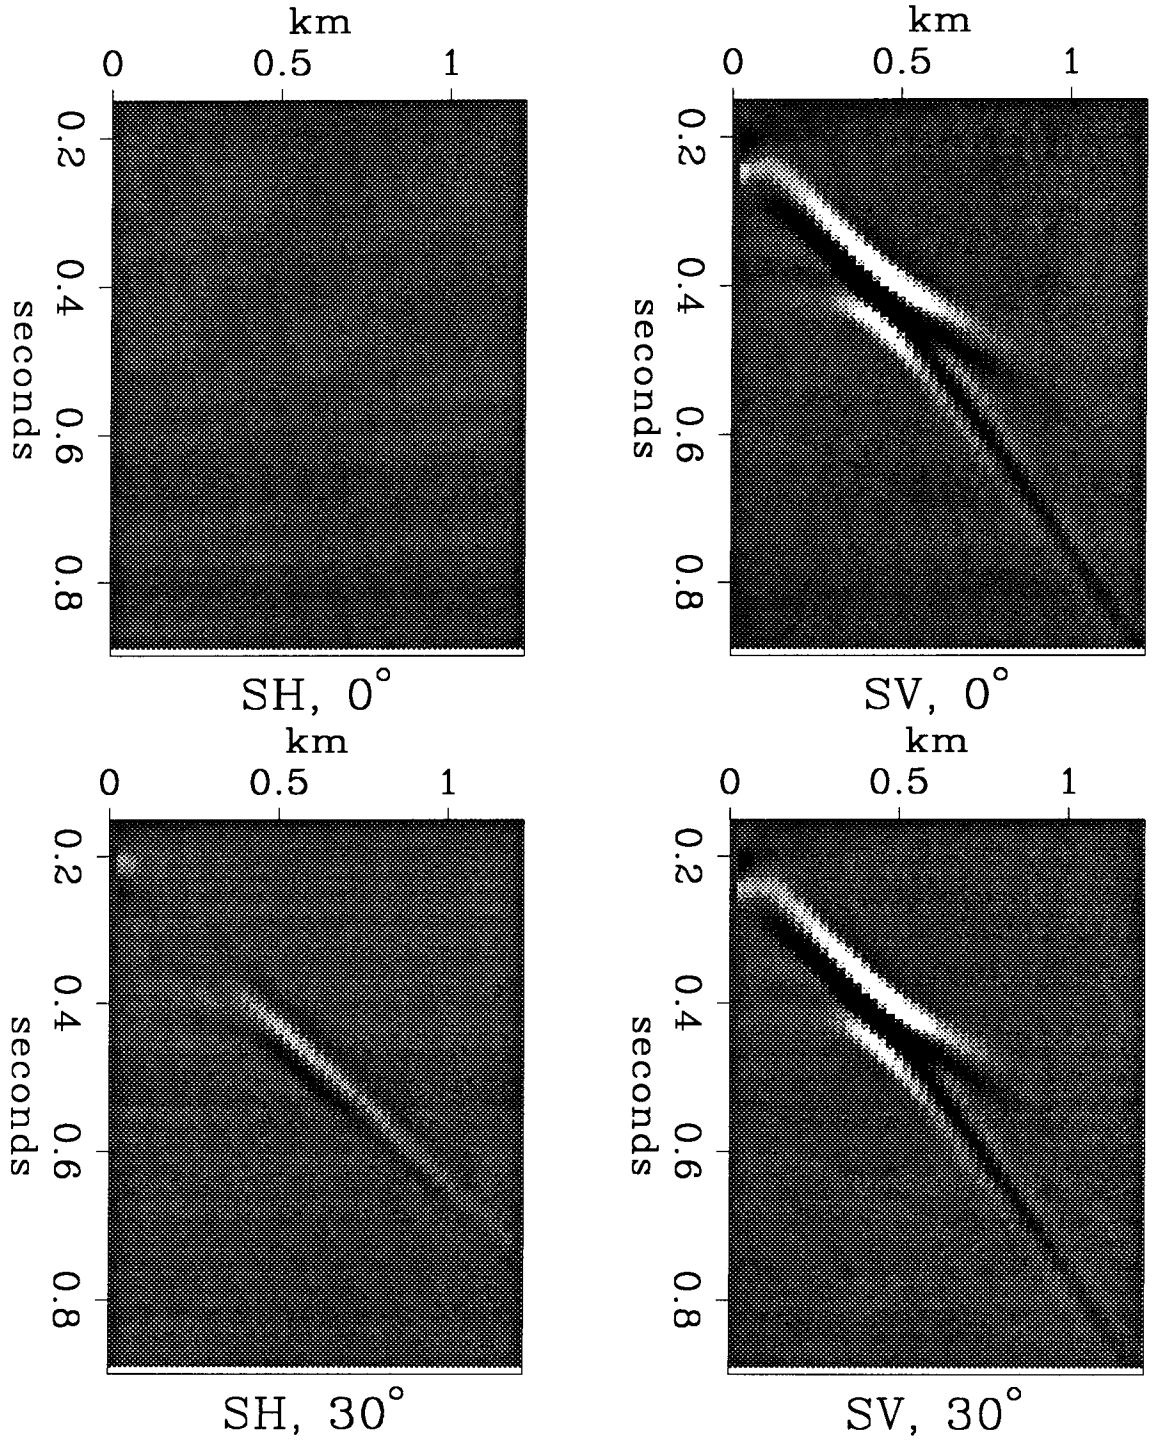

FIG. 3.37. Model sections for "cracked Greenhorn Shale" recorded along two different lines of azimuth. The cracks run along azimuth 0° (the x-z plane). The source is a vertical point force buried at a depth of 500 meters. The qP wave has been removed (leaving a slight artifact near the top left of each plot). Only the "SH, 30°" plot shows any distinctive orthorhombic effects.

Figure 3.29 on page 104, the $y = .25$ plot completely misses the singularity to the right. The $y = 0.$ plot was recorded along the x axis and passes tangentially through one edge of the cone of singular directions at $x = 0$, directly over the source. The $y = .125$ plot passes through the center of the cone of singular directions, and the $y = -.25$ plot passes tangentially through the left edge of the cone. The singularity mostly shows up in the strange flatness of the top of the hyperbola, especially on the $y = -.125$ plot.

Figure 3.39 shows lines parallel to the y axis. (Since the source and receiver are oriented along the y axis, these could be called "SV" lines.) In terms of Figure 3.29, the $x = 0.$ plot passes over the source and through the center of the cone of singular directions. The $x = .125$ plot passes tangentially through the near edge of the cone of singular directions. The $x = .313$ and $x = .656$ plots completely miss the singularity. Here the singularity shows up in the strange breaks in the events as the qS1 and qS2 modes reassert their separate identities after we move past the singularity, especially on the $x = .313$ plot.

Unfortunately from these examples it appears the direct effects of a singularity likely to show up on standard SV and SH sections are not very clear cut. Stacking such mangled events along hyperbolas is likely to produce something that looks like merely poor-quality or noisy data, even as the corresponding qP events stack consistently and well.

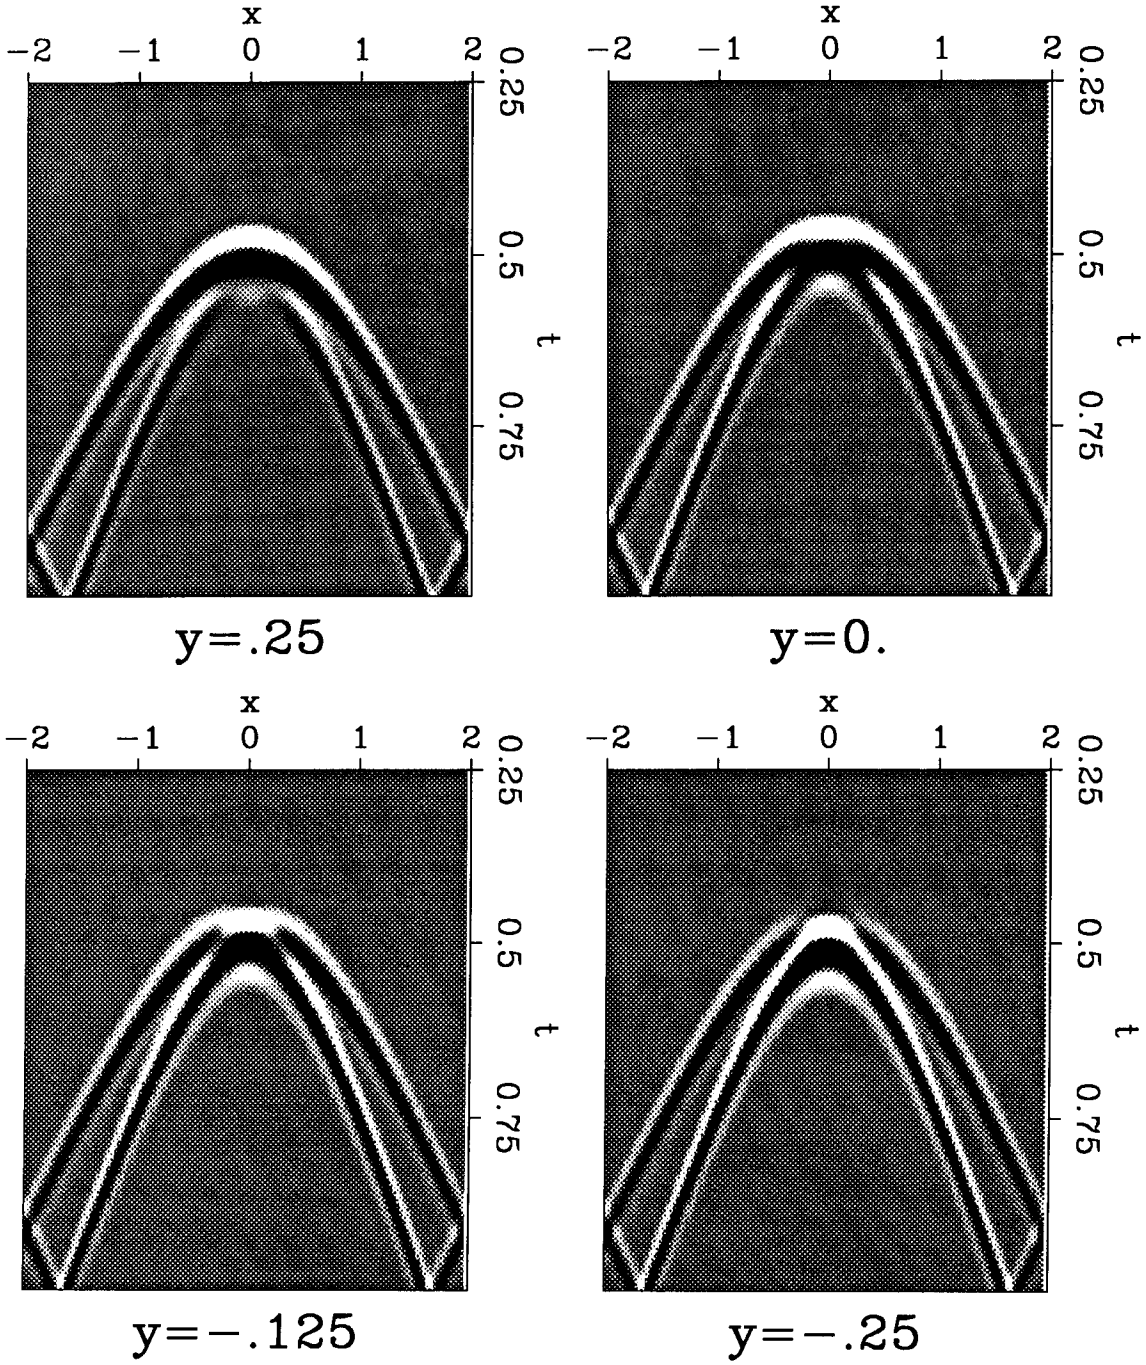

FIG. 3.38. Modeled x sections of the medium from Figure 3.30. The source is a Y point force buried at a depth of 1 unit; the y component of particle-motion direction is shown.

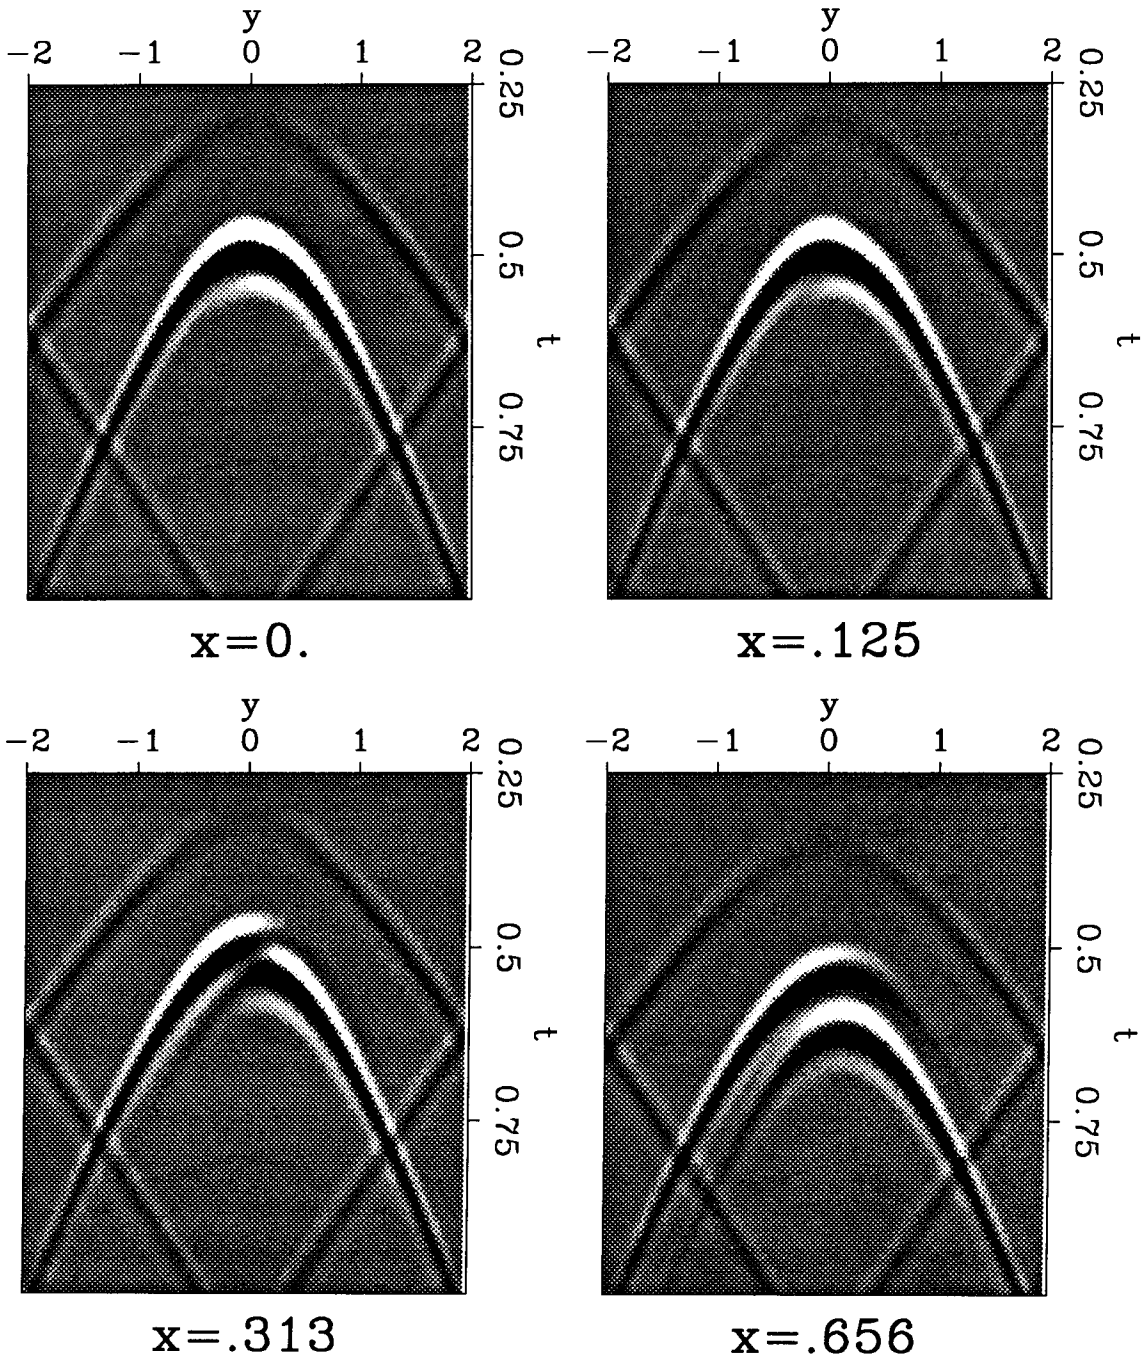

FIG. 3.39. Modeled y sections of the medium from Figure 3.30. The source is a Y point force buried at a depth of 1 unit; the y component of particle-motion direction is shown.