Imaging the earth with passive seismic arrays

GARY L. PAVLIS, Indiana University, Bloomington, U.S.

Earthquake and exploration seismology share a common foundation in elastic wave propagation. Most of what we know about earthquakes is known from the elastic waves they emit, while seismic methods in exploration are all about using scattered elastic waves to image earth structure. Seismologists have been using elastic waves to probe the earth since the very earliest instruments. After all, it was early seismic measurements that led to the spherical shell model of earth structure (crust, mantle, and core) taught to grade school students.

Until the emergence of the IRIS-PASSCAL facilities just over a decade ago, however, there was a drastic difference in the approach our different subsets of the field used to infer earth structure from seismic data. For almost 50 years exploration geophysicists have been developing the technology of direct wave imaging, which for the readers of *TLE* should require no explanation. Earthquake seismology research, in contrast, has been heavily focused on parametric measurement extracted from seismic signals. Examples include traveltimes and amplitudes of one or more seismic "phases," attenuation parameters, or (more recently) S-wave splitting. All parametric data I know of require the solution of an inverse problem to make inferences about earth structure. As a result, for the past 30 years academic research in seismology has viewed geophysical inverse theory as the cornerstone of most data processing. As the papers in this special issue demonstrate, the new technology of the portable broadband station and the availability of significant numbers of these instruments through the IRIS-PASSCAL facilities (and other international seismological initiatives) are rapidly changing this. We are seeing a shift toward a style of data processing more like that used in reflection seismology. The primary goal of this paper is to demonstrate these relationships and present them in the language of the exploration geophysics community.

There are two fundamental reasons why earthquake seismology has traditionally focused on parametric measurements: data quality and data density. Before IRIS-PASSCAL most data we dealt with came from analog, short-period instruments that were spaced at intervals of tens to thousands of kilometers. Early digital data provided an incremental improvement. We strapped relatively low dynamic range digitizers onto global seismic network stations or on the back end of analog telemetry systems feeding data to a centralized, regional, seismic network recording system. IRIS-PASSCAL brought a fundamental change in data quality through the acquisition of standardized, low-power field digitizers that could be deployed in any geometry the research team chose. The early 1990s witnessed a major revolution in seismic instrumentation that I will refer to as the portable "broadband" sensor. These relatively low-power devices have a flat frequency response across the traditional long-period band (20-50 s period) into the lower end of the exploration band (30-40 Hz), yet are smaller than a breadbox. Small, low-power sensors and digitizers made it feasible to deploy portable "broadband" seismic stations at a cost orders of magnitude lower than previously possible. (I should point out that the term "broadband" has become deeply entrenched in the jargon of earthquake seismology and implies this type of instrumentation. Most exploration recording systems are equally "broadband" to an electrical engineer.)

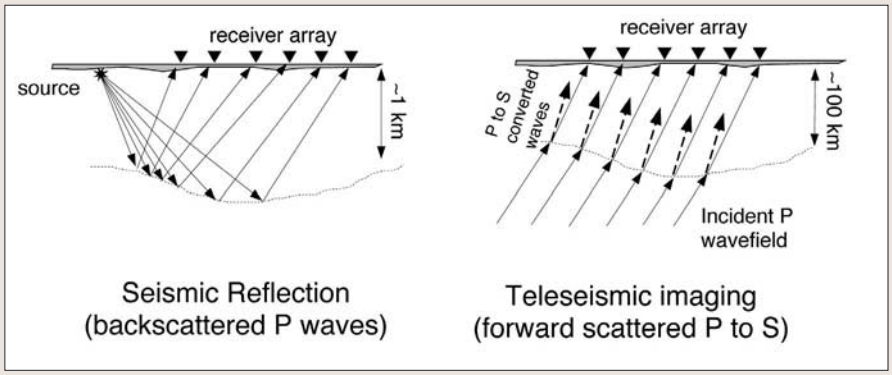

Figure 1. *Comparison of reflection and passive array geometries. The reflection method uses artificial sources to illuminate the subsurface from above while teleseismic body wave imaging has focused on forward scattered P-to-S conversions. Note the difference in scale that is related directly to the wavelength of the wavefield being imaged. The irregular, gray feature at the top of each figure compares the relative scale of two different features that have a comparable effect on data. In the reflection seismology case this symbolizes the weathered layer. At passive array imaging scales, a similar scale feature relative to the wavelength is an entire sedimentary basin.*

Broadband seismic instrumentation provides high-quality data critical to any direct imaging method. The reason is that the other limitation of older narrow-band data noted above, data density, scales directly with this capability. This is illustrated geometrically in Figure 1. The key issue is the difference in the primary recording band. Most seismic reflection data focus on the 10-100 Hz band while the major band for passive array imaging is around 0.03-2 Hz. The bandwidths are comparable, but the scales differ by two orders of magnitude. A typical seismic crew/ship has inline and crossline receiver spacings that range between 25 and 100 m. Direct scaling of this geometry upward for wavelengths a hundred times larger requires station spacing of 3-10 km to have comparable recording density. This has, in fact, become routine practice for passive arrays deployed in linear profiles. With areal arrays the situation is in a transition. The coverage required to record unaliased measurements of the wavefield (based on the above scaling model) exceeds the total pool of instruments currently available from IRIS-PASSCAL for anything but very shallow targets. USArray (see article by Meltzer in this volume) promises to change this with uniform coverage over the entire country. Although the USArray geometry is not consistent with the scaling described above (planned station spacing is approximately 70 km) direct imaging methods will still be applicable to these data at lower resolution/wavelength scales. For example, we may be forced to use only data below 0.1 Hz. Fortunately, there are good reasons to believe the earth is not as complex at this scale and the scaling relation described above may be overly pessimistic.

A key difference is the weathered layer. I do not think it is widely recognized that the most heterogeneous part of the earth is right under our feet. A boulder embedded in soil commonly

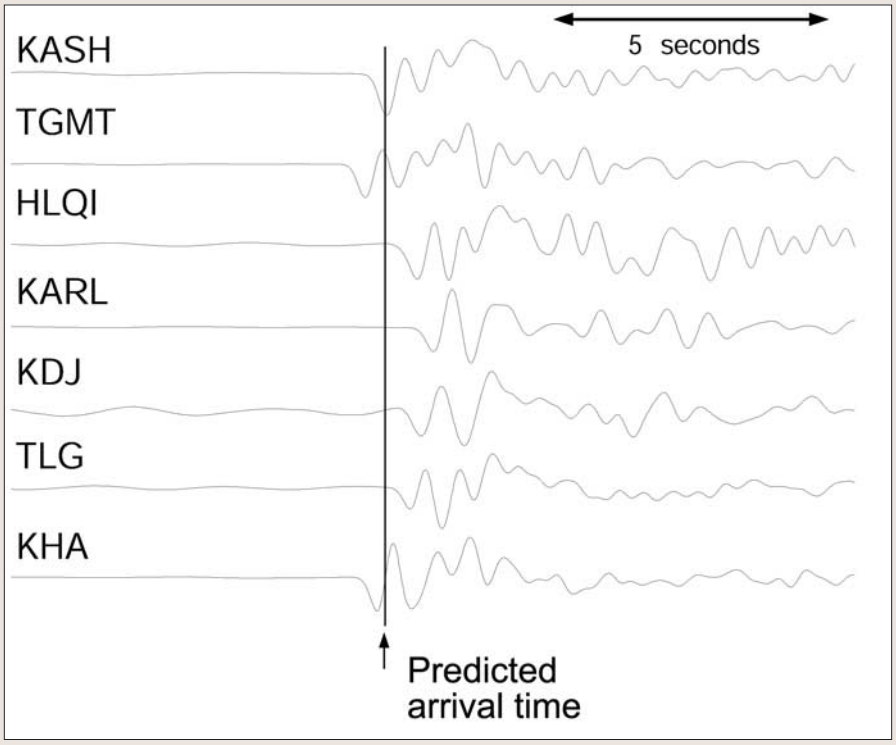

Figure 2. *Example teleseismic P waveforms. Shown here are vertical component seismograms of an intermediate depth (211 km) event from Sumatra recorded by an array of portable stations in the Tien Shan region of central Asia. The seismograms have been arranged in a roughly north-south profile from Kashgar (KASH) in China to KHA in southern Kazakhstan, which is a distance of more than 500 km. Note the similarity of waveforms and that they vary systematically from south to north. The traces have been aligned with the computed arrival times using the IASP91 earth model. The vertical lines show an estimate of the actual arrival times. Note that the arrivals match the theoretical times to within about 1 s. The variations from the predictions, called traveltime residuals, are the input data for P-wave tomography. Note that the moveout correction to align these signals is about five times the total time span of this figure. This illustrates that P residuals are small compared to the overall propagation delay. The similarity of the waveforms demonstrates why correlation methods are commonly used to measure relative timing.*

has a seismic velocity more than an order of magnitude larger than its surroundings, and the bedrock-soil interface is commonly a very complex boundary produced by chemical and biological agents that attack it. Anyone who has worked with both marine and land reflection data knows what a profound difference this makes in reflection data quality. I assert that broadband seismic data recorded on modern passive arrays are more like marine data than land exploration data. It is common to observe waveforms that hardly change at all for distances of hundreds of kilometers (Figures 2 and 3).

The fundamental reason for this is that for frequencies we commonly observe from teleseismic earthquakes (< 2 Hz) the weathered layer can nearly always be neglected. As Figure 1 illustrates at these frequencies entire sedimentary basins have a scale relative to a wavelength that is comparable to the weathered layer in land exploration data. Your signal is our noise. To us, sedimentary basins produce statics we have to work around just as the weathered layer produces statics that have to be handled in land data.

A final introductory remark. I chose to limit the scope of this paper to approaches that work with teleseismic body waves, and especially what I will call PdS waves. That is, I will focus primarily on P-to-S converted phases that can be viewed as forward scattered waves from an incident wavefield that is well approximated as a near-vertically incident plane wave (Figure 1). The reason for this focus is that to date this is the only component of the wavefields recorded by passive arrays that has shown promise of being amenable to direct imaging methods. I trust the reasons for this will be clear after you have read this paper, but you should recognize that it neglects a whole series of approaches that are widely applied to modern broadband data. These include a whole series of techniques centered on surface wave modes, S-wave splitting measurements from SKS phases, attenuation tomography, and a large class of methods based on waveform modeling. These other approaches have yielded, and will no doubt continue to yield, important scientific results, but I will not discuss them further because they utilize drastically different approaches than PdS methods. PdS phases are amenable to imaging approaches that have a very close tie to seismic reflection, which is a language familiar to most readers of this journal. Consequently, I will describe how these data are processed with direct parallels to algorithms familiar to all exploration geophysicists.

Preprocessing. My experience in reflection processing is that about 90% of human time spent on processing the data is spent on preprocessing steps like organizing the data, sorting out the geometry, editing dead channels, etc. Broadband seismic data are no different in this regard. It is just the problems that are different. Low-level battles with bits and bytes and quirky information systems are a universal problem and deserve no further comment. What is special about passive array data compared to reflection data is geometry. Some important differences are:

1) In reflection acquisition both the source and receiver arrays are always moving. With passive arrays the receiver geometry stays fixed for most, if not all of the experiment.
2) In a passive array we have no control over where the sources are, how big they are, or of the output wavelet. We take what we can get and try to deal with it after the fact. As a result we are required to link into a global system that provides information about earthquakes worldwide. We depend on agencies like the USGS, IMS, or ISC to provide reliable source information.
3) The difference in scale makes information systems used in reflection processing packages inappropriate for earthquake analysis. Every reflection package I know of uses local Cartesian coordinates relative to a geographic origin to define geometry. At the scales we work at, neglecting the curvature of the earth often results in a poor approximation, making it difficult to work directly with these systems.

Tomography = velocity analysis. Since pioneering work in the mid 1970s (Aki et al., 1977), body wave tomography has become a standard technique for imaging deep earth structure. The basis for the technique is illustrated in Figure 2. Notice how similar these waveforms are even though the group of stations shown span nearly 500 km. On the other hand, although I've applied the equivalent of a normal moveout correction (the raw moveout is calculated from an earth

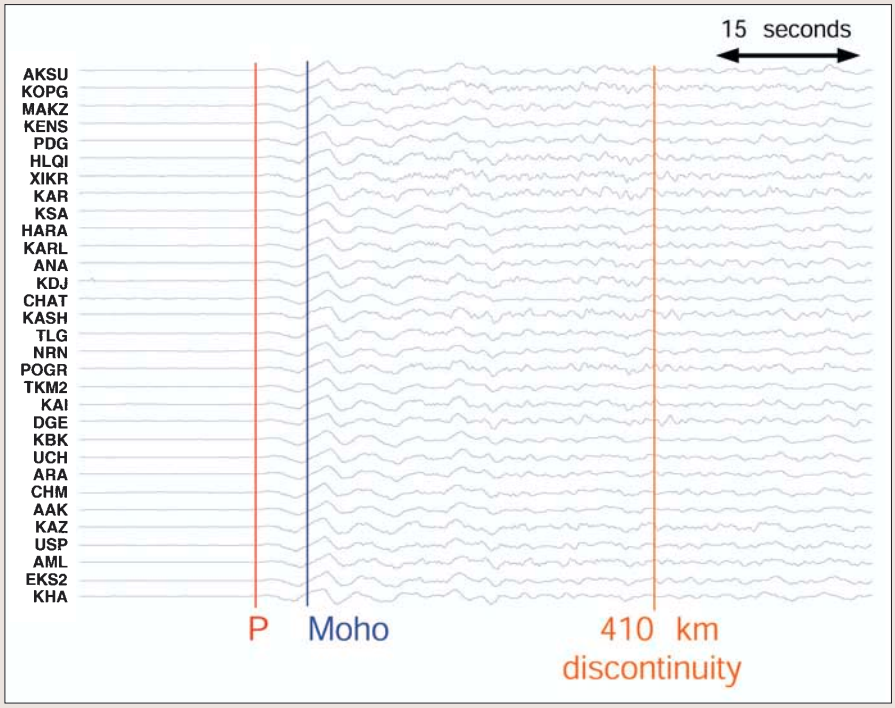

Figure 3. *Challenging aspects of deconvolution of passive array data. This figure shows vertical component seismograms from a shallow, magnitude 7 earthquake in Taiwan recorded by a broadband array in the Tien Shan of central Asia. The traces have been arranged by increasing distance from top to bottom and aligned on the predicted arrival time based on the IASP91 earth model. This figure shows how remarkably coherent P-waves on large aperture arrays like this can be (this array has a diameter of more than 500 km). These appear more coherent than those in Figure 2 largely because the dominant frequency for this large event is much lower than the deeper, smaller event shown there. It also demonstrates a fundamental problem. Large events have a long radiation time that makes the effective source wavelet very long and more like a vibrator signal than an impulse. This figure is a graphic demonstration of why deconvolution is so fundamental to success of converted wave imaging with passive arrays. I show the approximate time lag for a P-to-S conversion from crust-mantle boundary (blue line) and the major upper mantle discontinuity at 410 km depth (orange line). We design deconvolution operators to effectively compress the average wavelet seen here to yield an impulse response with a pulse width on the order of 1-3 s. Errors in constructing such an operator can easily produce false signals that are only sidebands of the deconvolution operator impulse response.*

model and a source location estimate from the USGS), the traces clearly do not line up. The timing corrections required to line up these traces, which can also be viewed as a form of static correction, are the input data of P-wave tomography. The techniques involved in using these data to infer earth structure have evolved significantly since the original work by Aki et al. and define a voluminous literature. Fundamentally, however, they all involve solution of a large, sparse nonlinear system of equations that relate the observed traveltime deviations to variations in seismic velocities under the array. If P phases are used, the result is a model of compressional-wave velocities while measurements of S can be used to build a shear-wave velocity model. Thus, for body-wave imaging, P- and S-wave tomography are analogous to velocity analysis in reflection imaging. They provide an analytical tool to produce a smoothed model of seismic velocity variations in space. I must emphasize that for many studies this is the end result of a passive array experiment. That is, the velocity model itself can and often does provide a valuable subsurface image that can be used to make inferences about geologic processes in the lower crust and upper mantle.

Although tomography plays the same role to us as velocity analysis, the methods used to make the measurements are closer to those used for measuring refraction statics. The earliest measurements used hand-picked times made by an analyst. Since work by VanderCarr and Crosson (1990), however, it has become common practice to use pairwise cross-correlations methods in the time or frequency domain to provide a more precise and objective way to align the waveforms.

More recently Lorie Bear and I (1999) developed a wavelet technique based on a set of special functions called multiwavelets (Lilly and Park, 1995) that is a hybrid of traditional time and frequency domain methods. This technique has the advantage of not only estimating arrival times; it simultaneously measures the best-fit particle motion ellipse from three-component data while at the same time estimating objective uncertainties in all these parameters. Using measured particle motions rather than depending on model-based predictions may often be important as a number of recent studies suggest P-wave particle motions often depart significantly from theoretical predictions (Bear et al., 1999; Schulte-Pelkam et al., 2003).

Deconvolution and wavefield interpolation. The most critical step in passive array processing to move beyond traveltime tomography is tied up in these two related steps. I discuss them together because they are elements of a series of three inseparable problems we face:

1) We have no control over the source. At best we know the approximate location of the source and we sometimes have a good estimate of the moment tensor (i.e., size and focal mechanism).
2) Source sizes vary by orders of magnitude. Remember, a magnitude 6 earthquake produces an amplitude 10 times that of a magnitude 5 at the same location.
3) We measure the vector wavefield at a relatively sparse collection of points constrained to lie at the earth's surface. We share this limitation with the standard seismic reflection method, but we have the added problem that the source is thousands of kilometers away and we cannot possibly model propagation accurately along the entire path. Furthermore, the density of coverage of most arrays deployed to date is low relative to the wavelength.

This may seem like an impossible problem and it probably would be intractable were it not for three fundamental things about the earth that make us lucky.

1) Evidence is accumulating from global tomography studies that the lower mantle is relatively homogeneous compared to the crust and upper mantle. Successful direct imaging methods have all focused on teleseismic body waves that pass through the lower mantle before propagating at relatively steep angles through the upper mantle and crust beneath the observing array.
2) Diffraction of transmitted waves, which is the curse of reflection seismology because it limits data resolution, helps tremendously because of the way it links to item 1. That is, if the lower mantle is indeed relatively

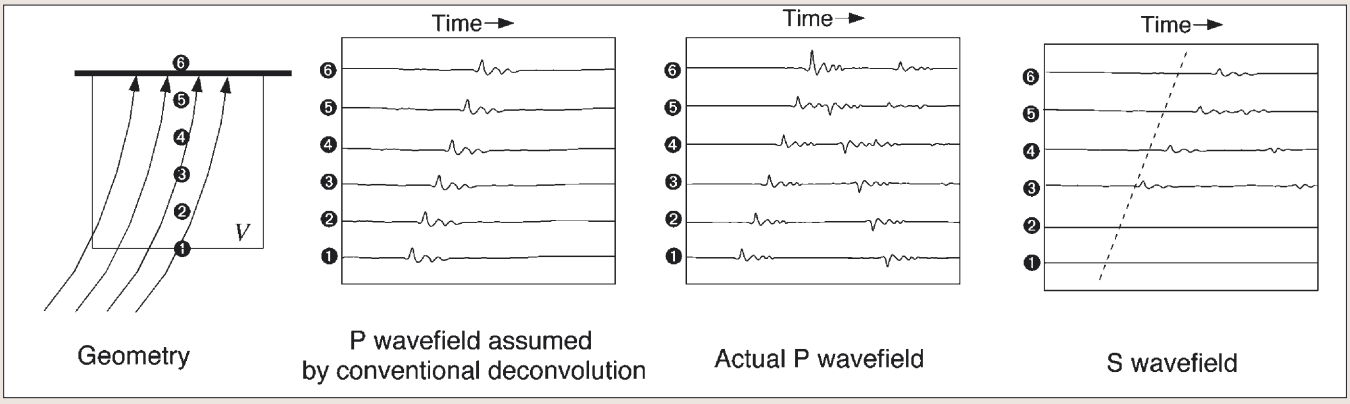

Figure 4. *Illustration of the fundamental assumption of all existing deconvolution schemes for body wave imaging. This figure can be viewed as a thought experiment. Suppose we had the equivalent of a string of sensors in a very deep well that penetrated well into the mantle. These imagined locations are shown as the numbered symbols. Data recorded by a such a collection of receivers looks like an upside down vertical seismic profile. Existing deconvolution schemes implicitly assume the earth through the study volume (where our imagined sensors are located) is approximately transparent yielding seismograms like that shown (second panel). The real situation is undoubtedly more like that shown on the right two panels. In this simple example we have a single horizon generating a P-to-S conversion near receiver 3. Deconvolution methods assume the S wavefield can be reconstructed from the signal measured in the longitudinal direction. The fact that the approach works at all suggests that the transparency approximation is not drastically far from reality.*

homogeneous, diffraction effects will homogenize the wavefield as it propagates through the lower mantle yielding a relatively constant, although perhaps ringy incident wavefield at the base of the transition zone.

3) At wavelength scales on the order of 5 km and greater, earth structure is dominated by near horizontal layering. At these scales pressure and temperature variations are overpowering such that the first-order structure is 1D. If the earth were strongly scattering in 3D, the deconvolution methods described below could not possibly work.

Unfortunately, mathematics is the only concise way to define this problem and those who wish to follow our reasoning are directed to the appendix. Although the approximations described in the appendix may seem crazy, at least as I've presented them, the undeniable fact is that the method works most of the time (Figure 4). Hundreds of papers have been published in the past decade that use "receiver functions" in one form or another to make inferences about earth structure to the point that this approach is now considered a standard analysis technique. The approach has consistently yielded reasonable images of the crust-mantle boundary and major upper mantle discontinuities. The hidden fact that colleagues reveal to me in the hallway, however, is that much of the reason this approximation seems to work is judicious data editing and a brutal use of stacking. Figure 3 shows an example of an event that is challenging to deconvolve cleanly. These vertical components are amazingly coherent across this entire array. This is remarkable considering this array was located in the Tien Shan of central Asia, a tectonically complex region, and the stations shown form an array with an aperture of approximately 500 km in diameter. This illustrates how important the deconvolution problem is, and some of the challenges. Much of the ringy, coherent signal seen here is a source effect generated by a large, shallow earthquake that occurred in Taiwan. The long, coherent signal is primarily a result of the long radiation time for a source this size and of source-side reverberations. Figure 3 shows the approximate lag for a P-to-S conversion from the base of the crust and the global-scale discontinuity at 410 km depth. Obviously deconvolution of this signal is critical because the duration of the effective source-time function is larger than the lag of a primary feature we would want to image (the Moho). (This is somewhat analogous to vibrator data except, in our case, the "sweep function" is not known.) On the other hand, on most data sets this type of event has a very good signal-to-noise ratio. Data like these do not, however, always yield the best results with current methods. Normally the best results come from larger, deep earthquakes that tend to have simple, effective source time functions that are relatively easy to deconvolve. I inflict this on the reader because it is my strong opinion that deconvolution is the single biggest barrier to improving this technology. The wavefield inversion/migration described below also has a lot of room to grow, but my experience in reflection processing is that if you have lousy data to start with, applying a procedure like migration usually only makes it worse.

Deconvolution is an essential first step to imaging. Otherwise it is like trying to migrate vibrator data without first correlating the data. Errors in deconvolution creep into results as a form of signal-generated noise that is not well understood. This is an area of overlap with industrial expertise that could yield important advances in the technology in the future.

A residual issue I want to address before closing this section is the issue of wavefield interpolation. Most "receiver function" studies have utilized what can be called a single-station technique for deconvolution. That is, deconvolution operators are commonly constructed from three-component records for each station and each event independently. Figure 3 illustrates why this may not be the best approach with an array. Even though this array is huge, the signals we observe across the array are remarkably consistent. This suggests immediately that multichannel deconvolution methods may be advised. Li and Nabalek (1999) were the first to apply multichannel methods to large aperture arrays as a way to improve deconvolution processing. More recently, Scott Neal and I (1999, 2001) generalized this concept with what we called pseudostation stacking.

Pseudostation stacking uses a weighted stack (we used a circular, Gaussian weight function) around a target image point to interpolate the wavefield onto a regular grid with a fixed smoothing function. Space-variable deconvolution operators can then be constructed across the surface of the earth that smoothly vary throughout the region covered by the array. Besides providing a way to interpolate the deconvolved wavefield onto a regular grid, we showed in our 2001 paper that this could be viewed as an efficient way to

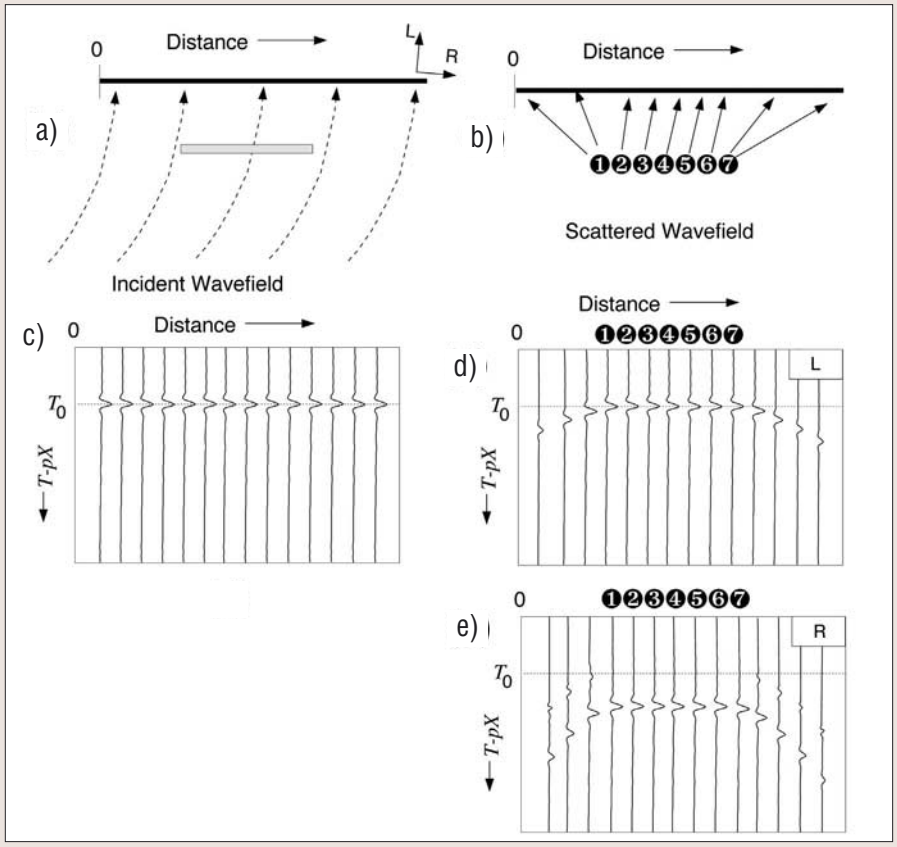

Figure 5. *The exploding converter model. (a) illustrates that the incident wavefield is a plane wave with a phase velocity of 1/p where p is the ray parameter. (b) is a ray diagram illustrating how this incident wavefield would be scattered by an inhomogeneity consisting of a thin slab of material—gray box in (a)—with a finite extent. The total wavefield is the superposition of (a) and (b). (c) to (e) illustrate theoretical components of the total wavefield that would be recorded by a hypothetical array of seismograms equally spaced along the top surface of the region illustrated in (a) and (b). All are illustrated as a reduced traveltime plot to simplify the record section displays. The standard pX correction aligns the wavefield to correct for the moveout of the incident plane wave. (c) illustrates an idealized L component record for the incident wavefield. The R component would be zero everywhere and is not shown. (d) and (e) illustrate the scattered wavefield for the L and R components, respectively. The exploding converter concept would replace the real geometry in (a) with a series of sources illustrated in discrete form in part (b) as seven circles with numbers. When treated as exploding converters, the scattered wavefield is equivalent to that generated by ripple firing these sources from left to right with a timing equivalent to a velocity of 1/p. The scattered P-wave seen on L is identical to the incident wavefield except for the tails caused by the edges, which is the reason P-waves are less useful for body wave imaging. The converted S-waves seen on the R component, in contrast, arrive with a lag relative to the incident wavefield arrival time (dashed line). P-waves bleed into the R component in this example because the scattered compressional waves from the edges of this object have a projection in the R direction. Similarly, scattered S-waves bleed into L as shown in (d) for the same reason. This illustrates the fundamental vector nature of this problem that makes it intrinsically more complicated than the scalar approximation commonly used in seismic reflection.*

construct a space-variable deconvolution operator that was appropriately regularized in both space and time.

The exploding converter model. The exploding reflector model is a well known construct for explaining imaging concepts in seismic reflection. There is a comparable construct in passive array imaging that I will refer to as the exploding converter model (Figure 5). Some key concepts to get from Figure 5 are:

1) In reflection seismology the illuminating wavefield comes from above (Figure 1) while here the illumination is from below. As a result there is a difference in timing. In the exploding reflector model the reflectors are conceived as secondary sources that fire at the arrival time of the incident wavefield. The idea here is the same. The incident P wavefield excites secondary S sources at the time the incident P wavefield sweeps by that point.
2) Any horizontal layer will shed a converted S with the same ray parameter (expressed as a linear moveout correction) as the incident P-wave, but with a lag proportional to its depth.
3) S-waves arrive with a polarization difference from the incident wavefield, which is predicted from elastic wave theory, and with a time lag caused by slower propagation of S that can be easily worked out from the exploding converter concept. Forward scattered P-waves, in contrast, may or may not be separable. Forward scattered P-waves shed from flat horizons arrive simultaneously with the incident wavefield and with identical polarizations. Only P-waves shed from the edges for inhomogeneities or dipping features will produce P-wave signals that have polarization or timing that differ from the incident wavefield.
4) Dipping layers (not shown in Figure 5) will shed scattered P and S phases with a linear moveout relative to the incident plane wave. Large dips are required to produce significant polarization or time lags for P-waves. In contrast, for small dips the polarization of an S phase will be close to that for a horizontal layer, but with a variable lag that scales directly with the dip. That is, dip is directly scaled with time lag.

The exploding converter model helps explain why P-to-S converted phases are an appropriate focus for passive array imaging. Because this is forward scattering instead of back scattering we have virtually no leverage to separate scattered P-waves from the incident P wavefield. S-waves, in contrast, have a simple scaling relationship between lag and depth, and are amenable to separation by polarization alone.

Wavefield separation. In reflection processing, one of the first steps is to apply top mutes to remove the direct and refracted arrivals. Most people don't think of it this way, but a major reason for muting data is to separate the scattered wavefield from the incident wavefield. The direct and refracted waves do not match the imaging condition that is most commonly applied for either simple CDP stacking or any form of prestack migration. Mutes remove them from the data in a very simple way.

In passive array processing this process is more complex. As Figure 3 shows, the incident and scattered wavefields commonly overlap completely. Deconvolution is essential. Without it the dilemma shown in Figure 3 is an inescapable prison. With S phases we have both polarization and timing to work with. In fact, Poppeliers and Pavlis (2003) showed that a simple top mute applied to deconvolved receiver functions provides an effective separation operator within the limits of the impulse response of the deconvolution operator.

With P phases this separation cannot be done at all for primaries except at extreme scattering angles. Bostock et al. (2001) pioneered the use of the free-surface reflection as a secondary source in passive array imaging. The incident P wavefield sheds reflected P and S phases from the free surface that can, in principle, be used as a secondary source. This yields a secondary wavefield indistinguishable from a surface P and S plane wave source. A current gap in the technology is that no existing deconvolution procedure can correctly account for this effect in estimating the scattered wavefield. Doing this problem right could significantly enhance the capabilities of the technique as it can, in principle, provide simultaneous illumination from above and below. This is loosely equivalent to recent work in exploration to use multiples as an integrated part of imaging instead of working hard to eliminate them.

Imaging. The densities of experimental arrays, in combination with the array apertures, have made true direct imaging methods begin to be feasible in only the last five years. The technology has developed rapidly during this time, but it is a long way from mature. One of the most important potential keys to making USArray successful, in fact, is likely to lie in building the bridges to successfully mix the extensive expertise in imaging found in the exploration community with scientists in the academic community.

With the background of the exploding converter model described above, it should be clear that the basic imaging condition for P-to-S conversions, at least, does not differ dramatically from that for backscattered P-waves in conventional reflection. In fact, the imaging condition for converted S-waves from a single event is nearly identical to that of conventional poststack migration. The reason is the well known concept that one can view CDP stacking as a focusing operator that converts the field data to the equivalent of the impulse response of the medium to a downward propagating, normal incidence, plane wave. The only difference in the passive array P-to-S conversion case is that the incident wavefield is propagating upward instead of downward, and the angle of incidence is not vertical. As Figure 5 illustrates, however, once the data are corrected for the incident wave moveout (i.e., lining up the first arrival P-waves), the basic imaging condition amounts to a form of back projection along the incident S-wave propagation path. The lag relative to the P-wave time scales with depth roughly as the difference between the S and P propagation time from that point to the surface (the exact relation is slightly more complex because the P and S phases for a given ray parameter travel along different paths described by classic ray theory).

The earliest forms of direct wave imaging applied to this problem were simple back projection schemes based on ray theory and the imaging condition described above. This technique, which Dueker and Sheehan (1997) called common conversion point (CCP) stacking, is somewhat analogous to old ray-based migration schemes. Signals from multiple sources are simply back projected along a computed S-wave raypath, binned, and summed to produce an image of the subsurface. There are, in my view, misleading statements in the literature that make an analogy between CCP stacking and CMP stacking of reflection data. CCP stacking is better thought of as a crude form of prestack migration based on pure ray theory that completely neglects diffraction effects.

In the past several years the concepts, at least, have been laid down for migration/inversion processing of passive array data. Developments to date can be grouped into two general classes: (1) extensions of reflection migration methods to this problem and (2) inverse scattering methods. The former are relatively straightforward to describe given the exploding converter model and the discussion above. Ryber and Weber (2000) extended CCP stacking to develop an approach that is similar to poststack migration of reflection data. That is, they show that CCP stacked data can be migrated with a standard Kirchhoff migration method using a migration velocity of twice the actual S-wave velocity of the medium. There are a number of abstracts of recent papers that describe similar concepts.

Inverse scattering methods, in theory, provide a more precise formalization for passive array imaging. The foundations of the theory were developed in the exploration geophysics community more than a decade ago, but recognition of their applicability to this problem was very recent. The critical distinction between inverse scattering in this problem and in conventional seismic processing is that conventional seismic formulations mostly focus on the acoustic approximation while the full elastic wave equation is essential here. The reason is simply that conventional seismic processing is designed to enhance and image primary P-wave reflections while S-waves are normally viewed as noise.

In our case, it is the conversion points we are trying to image, so a full elastic wave formulation is essential. Anyone who has attempted to read the inverse scattering literature knows that the mathematics quickly gets long and involved. I will not inflict this on the reader, but refer you instead to the suggested readings below. Instead, let me state some of the fundamentals and try to summarize the current state-of-the-art.

1) The existing theory is founded on the Born approximation. The most fundamental limitation this imposes is that it limits the imaging condition to primary conversions just as most migration methods are designed to focus primary reflections.
2) As a single scattering theory the equations boil down to three elements: propagate the wavefield from the source to the image point, scatter, and propagate the scattered wavefield from the image point to the receiver. All propagators used to date are based on ray theory. There are more fundamental issues about how to define the scattering term. Technically, the elastic scattering equations that fall out of the Born approximation involve radiation pattern terms that are the same

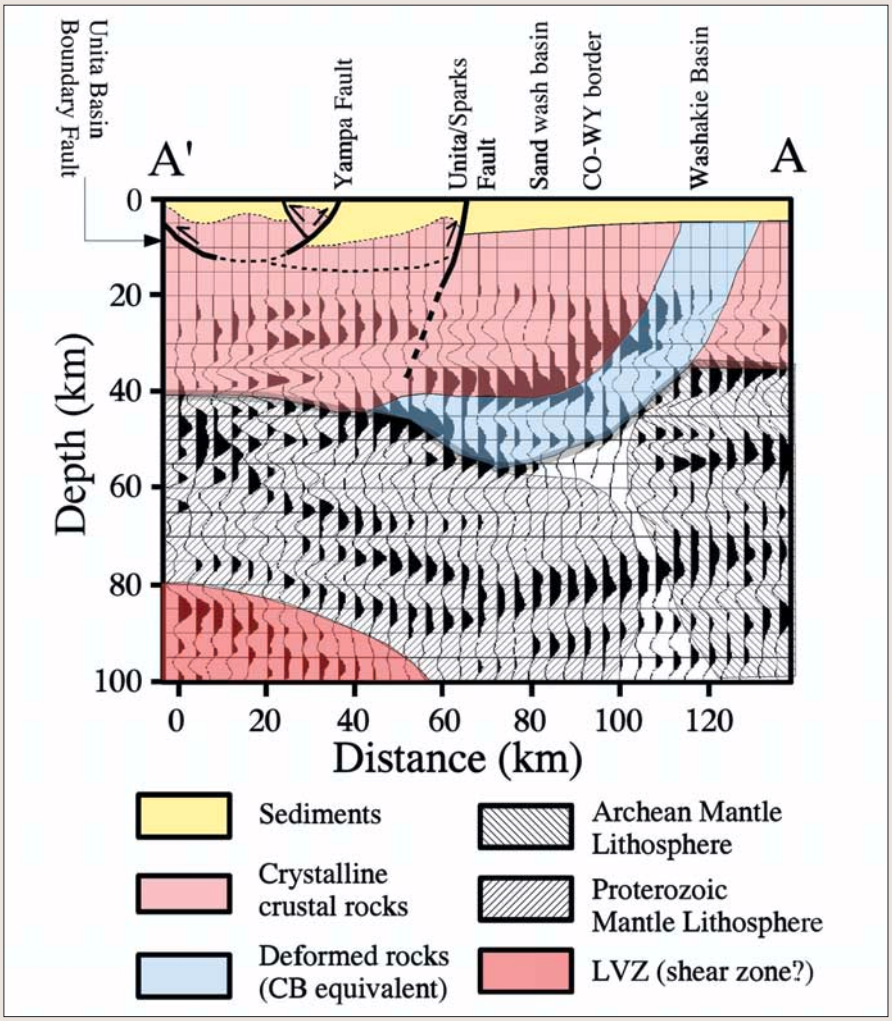

Figure 6. *Interpreted, converted-wave image of lithospheric scale structure. This is a converted-wave image produced by processing data from what has been called the Lodore array because of its location near the Canyon of Lodore in northwest Colorado. The array was deployed to image the suture between the Achean Wyoming province to the north (right) and the Proterozoic lithosphere Colorado Plateau to the south. This image was produced from 20 events recorded by this array over a 12 month period. Each event was deconvolved and interpolated onto a regular grid along a north-south profile using pseudostation stacking/deconvolution. A prestack migration method was then applied to each event and the results were stacked using a weighting scheme based on signal-to-noise ratio and azimuth normalization (needed to balance the stack from differing illumination directions). An interpretation, overlain on the data, demonstrates the potential imaging capabilities of this technology for understanding the link between crust and upper mantle deformations.*

as those in the classic elastic wave representation theorem. Bostock et al. (2001) include these terms in a formalized inversion for elastic wave perturbations within the entire medium. Rondenay et al. (2001) applied this approach to image the subducting slab under Oregon with spectacular results. This asks far more of the data than a simple approach like CCP stacking, however, and whether this approach will be as universally successful remains to be established.

3) The "inverse" in inverse scattering formulations generally involves an analytical inversion formula commonly called the inverse, generalized, Radon transform. It is generalized because the original Radon transform equations are for line integrals through a 2D area. The generalization to the 2D and 3D scattering equations involve a complex association of scattered wave "isochrons" with line integrals in the forward Radon transform. Ultimately, this relates to an important principle of an imaging condition. Each sample of a trace that is to be projected into the image has to satisfy a causality principle that it match the imaging condition. That is, each sample of every seismic trace can be associated with a curved, isocron surface that defines the locus of all points that satisfy the P-to-S conversion imaging condition. Bostock et al. and Poppeliers and Pavlis utilize this approach to develop an inverse scattering formulation for imaging passive array data. Bostock et al.'s approach is limited to 2D, but utilizes the full elastic equations to formally invert for P velocity, S velocity, and density fluctuations. Our paper was limited to estimating P-to-S scattering potential, but it is the only existing fully 3D formulation of this problem.

I close this overview of current passive seismic imaging techniques with an example of the current state-of-the-art in direct-wave imaging. The result in Figure 6 is from a 33-element array deployed in northwest Colorado and southeastern Wyoming on the eastern edge of the Uinta Mountains. The array was deployed to image the suture between the Colorado Plateau and Wyoming province commonly called the Cheyenne Belt. These data were processed with the prestack, plane wave migration procedure of Poppeliers and Pavlis (2002) and stacked using a source equalization technique described by Poppeliers and Pavlis (2003). The interpretation suggests a dipping, crustal scale feature that we interpret as the Cheyenne Belt suture. There are also distinct differences in the features observed in the upper mantle on the left and right sides of this image. This is consistent with the fact that the lithosphere on opposite sides of this boundary was formed at drastically different geologic times. (The lithosphere of the Wyoming province is Archean in age while lithosphere of the Colorado plateau was formed in the Proterozoic.)

Suggested reading. A more theoretical review paper that covers roughly the same ground as this one with a more complete series of references is "Direct imaging of the coda of teleseismic P waves" by Pavlis (in *Data Analysis and Imaging with Global and Local Arrays*, AGU monograph series, 2003). The full theory for inverse scattering in elastic media can be found in "Linearized inverse scattering problems in acoustics and elasticity" by Beylkin and Burridge (*Wave Motion*, 1990). Other papers cited in the text are "Determination of the three-dimensional seismic structure of the lithosphere" by Aki et al. (*Journal of Geophysical Research*, 1977); "Determination of teleseismic relative phase arrival times using multichannel cross-correlation and least squares" by VanDecar and Crosson (*Bulletin of the Seismological Society of America*, 1990); "Multichannel estimation of time residuals from seismic data using multiwavelets" by Bear and Pavlis (*Bulletin of the Seismological Society of America*, 1999); "Multiwavelet analysis of three-component seismic arrays: Application to measure effective aniso-

tropy at Pinon Flats, California" by Bear et al. (*Bulletin of the Seismological Society of America*, 1999); "Large teleseismic P wavefront deflections observed with broadband arrays" by Schulte-Pelkam et al. (submitted to the *Bulletin of the Seismological Society of America*); "Imaging P-to-S conversions with multichannel receiver functions" by Neal and Pavlis (*Geophysical Research Letters*, 1999); "Imaging P-to-S conversions with broadband seismic arrays using multichannel, time-domain deconvolution" by Neal and Pavlis (*Geophysical Research Letters*, 2001); "Deconvolution of teleseismic body waves for enhancing structure beneath a broadband seismic array" by Li and Nabelek (*Bulletin of the Seismological Society of America*, 1999); "Three-dimensional plane wave migration of teleseismic P-to-S converted phases: Part 1, theory" by Poppeliers and Pavlis (*Journal of Geophysical Research*, 2002); "Mantle discontinuity structure from midpoint stacks of converted P-to-S waves across the Yellowstone hotspot" by Duecker and Sheehan (*Journal of Geophysical Research*, 1997); "Receiver function array: a reflection seismic approach" by Ryber and Weber (*Geophysical Journal International*, 2000); "Multiparameter two-dimensional inversion of scattered teleseismic body waves 3. Application to the Cascadia 1993 data set" by Rondnay et al. (*Journal of Geophysical Research*, 2001); "Three-dimensional plane wave migration of P-to-S converted phases: Part 2, stacking multiple events" by Poppeliers and Pavlis (*Journal of Geophysical Research*, in press); "Multiparameter two-dimensional inversion of scattered teleseismic body waves 1. Theory for oblique incidence" by Bostock et al. (*Journal of Geophysical Research*, 2001).

Acknowledgments: Our work on converted wave imaging has been supported by the National Science Foundation under grants EAR-9904043 and EAR-0106062. Our work was greatly enhanced through the University Grant's program of Landmark Graphics Corporation. Finally, I thank the Cecil and Ida Green Foundation for Earth Sciences at the Scripps Institute of Oceanography for supporting me as a Green Scholar while I wrote this paper during a sabbatical there.

Corresponding author: pavlis@epicenter.ucsd.edu

Appendix. The data we need to work with are N three-component seismograms, u_i, recorded by an array that were created by a distant earthquake (Figure 3 shows an example of the vertical components of such an event). The fundamental theorem of seismology says that

$$u_i(t) = G_{ij,k} * M_{jk}(t) \qquad (1)$$

where $M_{jk}(t)$ is the moment tensor, $G_{ij,k}$ is Green's function that describes wave propagation from the source to our receiver, and * denotes convolution. We may have an estimate of M_{jk} but its details are generally not important. Because we are thousands of kilometers from the source, fluctuations across the array, induced by the radiation pattern, are rarely important because the range of azimuths span only a few degrees. The source time function (the t variable) is rarely known, but the deconvolution procedure we apply ultimately absorbs this term. This is a classic inverse problem. We measure u and we may know a bit about M_{jk} but we know little about $G_{ij,k}$. $G_{ij,k}$ is a function of source position, receiver position, and the elastic wave parameters of a large fraction of the earth. It is the latter that we would like to know, but its relationship to the data is complex.

To make the problem tractable we assume: (1) the illuminating wavefield is pure compressional-wave energy. (The scale makes this easier than one might think. The separation between P and S for teleseismic earthquakes is minutes while the portion we aim to image is within the first minute following the P wave time.); (2) the signal we record at the surface rotated to the longitudinal direction (L), and parallel to the raypath at the point of emergence, is a direct measure of the effective source wavelet that has illuminated the subsurface; and (3) the signal we record in directions perpendicular to the L component (R=radial parallel to the raypath and T=transverse or perpendicular to the plane of the raypath) are dominated by P-to-S conversions induced by the illuminating wavefield. With these assumptions the problem can be cast in a relatively simple mathematical framework. We begin by using the transformation $u' = Ru$, where R rotates the data in the measurement coordinates (east, north, vertical) to ray coordinates (L, R, and T). Let u_3' denote the longitudinal (L) component. We assume that we can write

$$u_3' * r_i = u_i' \qquad (2)$$

where i=1,2 are components in directions perpendicular to u_3'. That is, we assume that we can express the data observed on the transverse directions, u_1' and u_2' as a convolution of the longitudinal component data, u_3', with "receiver functions," r_1 or r_2. This receiver function is assumed to be an estimate of the impulse response of the medium for P-to-S conversions created by an effective source wavelet equal to the recorded signal in the longitudinal direction.

Figure 4 demonstrates graphically that this is equivalent to assuming the earth beneath the array is completely transparent to P-waves and the only significant interaction that takes place are P-to-S conversions. It neglects, for example, even first-order effects like the total reflection off the free surface. I know of no published paper on this subject that does not make this assumption as a foundation for deconvolving the data to produce estimates of the functions r_1 and r_2. The methods used for computing these deconvolution operators are well known in reflection seismology, but they are not what most people would think of as deconvolution in reflection processing. The methods we use are variants of wavelet shaping filters used, for example, to correct marine data based on a measured source wavelet. The primary difference is that our measured source wavelet is not independently measured as it might be, for example, on a marine vessel. Instead, we make the approximations described above and utilize an algorithm that is virtually identical to wavelet shaping filters but applied to this problem. TLE