

Seismic interferometry—turning noise into signal

ANDREW CURTIS, University of Edinburgh, UK

PETER GERSTOFT, University of California at San Diego, USA

HARUO SATO, Tokoku University, Japan

ROEL SNIEDER, Colorado School of Mines, USA

KEES WAPENAAR, Delft University of Technology, The Netherlands

Turning noise into useful data—every geophysicist's dream? And now it seems possible. The field of seismic interferometry has at its foundation a shift in the way we think about the parts of the signal that are currently filtered out of most analyses—complicated seismic codas (the multiply scattered parts of seismic waveforms) and background noise (whatever is recorded when no identifiable active source is emitting, and which is superimposed on all recorded data). Those parts of seismograms consist of waves that reflect and refract around exactly the same subsurface heterogeneities as waves excited by active sources. The key to the rapid emergence of this field of research is our new understanding of how to unravel that subsurface information from these relatively complex-looking waveforms. And the answer turned out to be rather simple. This article explains the operation of seismic interferometry and provides a few examples of its application.

A simple thought experiment. Consider an example of a horizontally stratified (one-dimensional) acoustic medium, and for the moment let us imagine that it has only a single internal interface. Now, say horizontally planar pressure waves are emitted by two impulsive sources, one after the other, and that one source is above the interface and one below. Vibrations from the resulting propagating waves are recorded at two receivers which can be placed anywhere between the two sources (Figure 1, left).

The recordings are shown in the center of the figure. At each receiver a direct and a reflected wave is recorded for source 1, whereas only one transmitted wave is recorded for source 2.

Seismic interferometry of these data involves only two simple steps: The two recorded signals from each source are crosscorrelated and the resulting crosscorrelograms are summed (stacked). The result, shown on the right of Figure 1, is surprising; for positive times it is the seismogram that would have been recorded at either receiver if the other had been a source, and at negative times it is the time reverse of this seismogram. In other words, by this simple, two-step operation we have constructed the seismic trace from a virtual source—a source that did not exist in our initial experiment, and a source that is imagined to be at the location of one of our receivers.

To generalize, this simple example placed no constraint on where the receivers were placed, provided they were between the sources. By moving either or both of them (or by using many distributed receivers from the start), it is therefore possible to construct the trace from an infinite number of virtual source and receiver pairs placed at any locations, by recording the signal from only two actual sources. What is more, provided one of the active sources is above the interface and receivers and the other is below, the location of the active sources is also arbitrary, and in order to carry out the process above we do not even need to know where these sources are.

Seismic interferometry steps. The fundamental steps of the

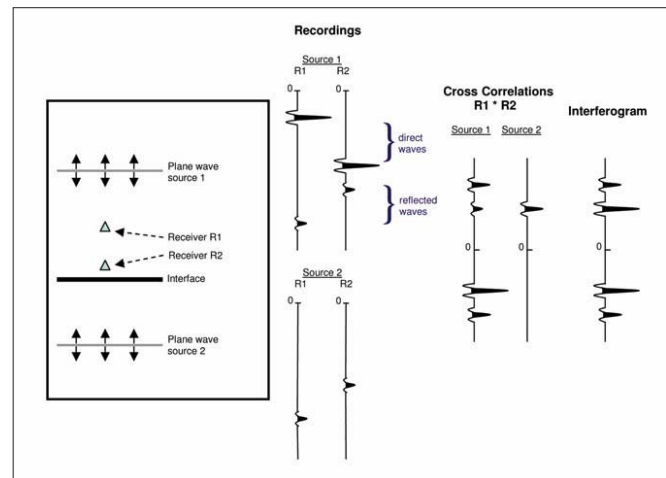


Figure 1. Interferometric construction of a virtual source. (left) One-dimensional acoustic medium consisting of single interface between two half-spaces, with two plane-wave sources and two receivers. (center) Traces recorded at each receiver for each source. (right) crosscorrelations between pair of traces for source 1 and for source 2, and the sum of these crosscorrelations. At positive times, the final summed trace turns out to be the trace that would be recorded at one receiver if the other had been a source. (Note that the virtual source wavelet in the three traces on the right should in fact be the autocorrelation of the recorded source wavelet shown in the center; we have omitted this change in source wavelet in the figure for simplicity.)

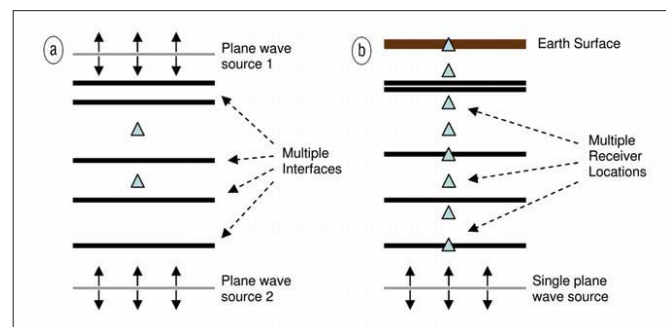


Figure 2. Alternative, more Earth-like models for which the process in Figure 1 works equally well. (left) Multiple layers with no free surface—still two sources required. (right) Multiple layers with a free surface—only one source required. The right plot also shows that any receiver locations can be used for the virtual source and receiver reconstruction.

operation are simple: crosscorrelation (we can understand this operation as detecting the traveltime difference of the recorded waves between the pair of receivers), then stacking (i.e., integration over all actual sources; a few details required to get the dynamics correct have been omitted for clarity). Yet, the technique is powerful and so far we have barely scratched the surface.

The result above holds for any horizontally stratified medium, still using only two actual sources (Figure 2a). The important criterion for the distribution of actual sources is that they completely surround the medium of interest (a portion of a one-dimensional medium is “surrounded” by two

points, at the top and at the bottom). However, if any part of the boundary is a surface of total reflection (like the free surface of the Earth), it turns out that no source is required on that boundary. Hence, in 1D Earth-like models, only a single actual source is required to construct seismic traces between *any* source-receiver pair, including sources or receivers placed on the free surface (Figure 2b). Now, consider a case in which a complex, multilayered medium is situated below the region of the model of interest (between the sources in Figure 2a); this is probably realistic for the Earth. In that case, if source 2 is moved below this complex part of the medium, its contribution to the received signals becomes virtually zero due to transmission losses. In that case, the lower source in Figure 2a can be neglected and again only a single active source is necessary to construct the inter-receiver seismic traces.

The above example also shows us how to make sense of noise and codas (the long, multiply reflected tails of data observed on seismic traces). It turns out that impulsive sources on the boundary can be replaced by uncorrelated noise sources that emit continually and simultaneously. Any pair of extensive noise records from any two receivers can then be crosscorrelated and, remarkably, the result will be approximately the same as above: The crosscorrelation will approximate the impulse response (the measured wavefield at one location if an impulsive source is placed at the location of the other) on the right of Figure 1, convolved with a source time function that is the autocorrelation of the noise. In fact, the first ever seismic interferometry theory, derived by Claerbout in 1968, was for the emergence of the reflection response in cases like Figure 2b where both receivers are placed on the surface and the actual source below (e.g., microseismic activity) emitted random noise.

Generalizing to 3D. What is even more remarkable about the current theory is that with a little modification it is applicable to waves propagating through any lossless (nonattenuating) one-, two-, or three-dimensional medium, and some empirical applications in attenuating media have also been successful (see below). Hence it can also be applied to three-dimensional elastic Earth models. Other than some numerical adjustments, the main modification is that the sources must surround the medium entirely; the stacking or integration step is then performed over the crosscorrelations from that entire set of sources.

Notice something important about the 1D results: all of the information we need to calculate wave propagation from any source location within the medium is contained within the waves propagating from a single source on the lower boundary in Figure 2b (or from a single source in Figure 2a if the medium is complex outside of the part depicted, as discussed earlier). To record data from multiple sources in a 1D medium is to store redundant data. As the theory generalizes to multiple dimensions, so also does this result: storing data from sources on the boundary of a medium is sufficient to construct data from any other source placed within the medium or on the free surface.

Requirements of 3D seismic interferometry. It is worth noting the main assumptions behind seismic interferometry theory which currently impose limitations on its domain of applicability. First, for an exact application using noise sources on the bounding surface, the medium must be lossless (nonattenuating). Second, if random noise sources (e.g., sources of background noise in the Earth) are to be used in two or three dimensions, then the distribution of that noise must be “even” in senses to be made clear later. Third, as

we stated above, if active sources are used, then to obtain dynamically correct responses (i.e., with the correct amplitudes) the sources must completely surround the portion of the medium of interest (other than along completely reflecting boundaries). On this last point, notice that the only place that we do not require sources to completely surround the medium is precisely where physical and cost constraints force us to put them—on the Earth’s surface!

Definition of seismic interferometry. The term *interferometry* generally refers to the study of interference phenomena between pairs of signals in order to obtain information from the differences between them. Seismic interferometry simply refers to the study of interference of seismic-related signals. The principal mathematical operation used to study this interference is crosscorrelation of pairs of signals, but one could equivalently consider convolution as the principal operation because crosscorrelation is simply convolution with the reverse of one of the two signals. The signals themselves may come from background-propagating waves or reverberations in the Earth, from earthquakes, from active artificial seismic sources, from laboratory sources, or from waveforms modeled on a computer—examples of using all of these data types will be given below.

The above definition is fairly general and covers a multitude of different “types” of interferometry. For instance, one example presented below uses the crosscorrelation of only two signals recorded at a single receiver from pairs of repeated artificial seismic sources to obtain information about the difference in average seismic velocity of the crust before and after an earthquake occurred. Another example integrates a suite of crosscorrelations of synthetic signals from sources that span the entire boundary of a medium of interest in order to efficiently model waveforms between any sources and receivers contained in the interior of that medium. In each example, various different subsequent processing steps are applied to the data in order to obtain different types of information, but the fundamental initial operation is the same: crosscorrelation.

The majority of different theories and applications using seismic interferometry can be divided into two distinct classes of techniques: those that are predominantly used to obtain information about the medium through which waves have propagated, and those that reconstruct information about the propagated waves themselves. As we will see, each has tangible relevance to the industrial seismic community.

The rise of seismic interferometry into mainstream geophysics has been fueled by a rapid sequence of advances. The majority of the most significant advances have been reported outside of the exploration literature. We will review these important pieces of work (many of which are listed in the appendix) and then show some recent, exciting examples of how this new theory can be applied. The examples include the evaluation of building responses to seismic waves, tomography for crustal properties, changes in crustal properties over time, ground-roll removal from seismic data, and waveform modeling. Finally, we discuss some key challenges for this field in the future.

Historical development of interferometry. The development of what we now think of as interferometry has occurred through several distinct breakthroughs, each of which changed our understanding of how interferometry works. We now review these as they form the fundamental literature behind the supplement to the July-August issue of *GEOPHYSICS*.

The extraction of the impulse response of a system from

noise is known in physics as the *fluctuation dissipation theorem* (e.g., Callen and Welton, 1951). A system near equilibrium behaves in the same way to an external force as it does to fluctuations within the system (noise); it relaxes to the equilibrium state. However, the birth of interferometry in seismic applications can be clearly identified in a paper by Jon Claerbout (1968). He showed that if a 1D medium is bounded on top by a free surface (like the surface of the Earth) and is bounded below by a half-space (homogeneous, infinitely extensive Earth), then the plane-wave reflection response of a horizontally layered medium (what we would record at the surface of the Earth given a source also at the surface) can be obtained from the plane-wave transmission response of the same medium (what we would record if instead the source was in the half-space below). Indeed, he showed that the reflection response is obtained directly by autocorrelation of the transmission response—crosscorrelation of the transmission response with itself. When the source in the half-space below is a noise source, the source wavelet of the reconstructed reflection response is the autocorrelation of the noise.

Claerbout's conjecture. This in itself was of great theoretical interest, but Claerbout also made a phenomenal conjecture: that the crosscorrelation of noise traces recorded at two different receiver locations in three-dimensional, heterogeneous media gives the response that would be observed at one of the locations if there was a source at the other. In other words, simply by listening to noise at two receivers, we can construct the signal that would have been observed if we had used a source at one of the receiver locations. This method to construct artificial seismic sources was to be demonstrated and proven years later (e.g., Rickett and Claerbout demonstrated its application to helioseismological data in 1999).

Weaver and Lobkis (2001) demonstrated this conjecture for ultrasound waves by calculating that the long-term average of random noise in an aluminum block yields the time-domain impulse response between the two points. In 2002, they also provided one of the first proofs of Claerbout's conjecture. However, in order to prove it, they assumed that the noise wavefield was diffuse (i.e., waves arrive from all angles with equal strength). Diffusivity might be created approximately in nature by multiple scattering in a finite body with an irregular bounding surface, multiple scattering between randomly distributed scatterers within the body, or due to a random distribution of uncorrelated sources distributed throughout the medium. Nevertheless, the assumption of a diffuse wavefield imposes significant restriction on the domain of application of seismic interferometry.

In 2004, one of us (KW) proved the generalization of Claerbout's conjecture for 3D (acoustic and elastodynamic) media without assumptions about randomness of the medium, noise sources, or diffusivity of the wavefield. The derivation is based on reciprocity theory, and applies to any inhomogeneous, lossless, anisotropic medium. It uses independent responses of many sources (transient or noise) recorded at each pair of receivers to construct the inter-receiver impulse response, and for uncorrelated noise sources the expressions reduce to a single crosscorrelation of observations at the two receivers. This validated the interferometric conjecture of Claerbout.

Time reversal. In a pair of 2003 articles, Derode et al. showed how the principle of interferometry is related to time-reversed wavefields; in so doing they provided an intuitive and elegant derivation of impulse response reconstruction

based entirely on physical, and not mathematical, arguments. They showed how a time-reversed mirror (a mirror that reflects any signal but with the time axis flipped) can be used to time-reverse a wavefield emitted from a single source, such that it converges on the original source location (imagine playing a movie backward such that ripples in a pond contract back to the point where a stone had been dropped into the water). If that wavefield is recorded at a second receiver location, the time-reversed impulse response between the source and receiver is recorded. The wavefield response to the original source of energy (the stone) is then reconstructed so the waves (ripples) begin to expand again; during this phase of the experiment the time-forward impulse response is recorded. It turns out that the operation of time reversal at the mirror is precisely the crosscorrelation operation that is used in seismic interferometry if instead sources were placed on the boundary rather than inside of the medium. Note that in Figure 1 (right), both time-forward and time-reversed impulse responses were constructed.

For direct source-receiver arrivals and for singly reflected waves, Snieder (2004, 2006) provided an explanation of the mechanism of impulse response reconstruction by crosscorrelation. Snieder argued that observable arrivals usually occur when the signal phase is stationary (approximately, does not change) with respect to perturbations of raypaths. Making this approximation allowed him to analyze the location of actual sources that contribute to the interferometric signal. He demonstrated that the dominant contribution to an interferometrically constructed, inter-receiver impulse response comes from scatterers (or noise sources) in two cones around extensions of the line that connects the two receivers. This implies that sources in other locations may be (approximately) redundant in the reconstruction.

Schuster (2001) and Schuster et al. (2004) showed how crosscorrelations of seismic responses from man-made or natural sources at the surface or in the subsurface can be used to form an image of the subsurface. The crosscorrelated data (interferograms) are downward extrapolated and imaged in much the same way as in standard prestack migration. The method is more robust than methods that first reconstruct the reflection response, and they make no assumptions with respect to wavefield diffusivity or sufficient source coverage. On the downside, most multiple reflections are incorrectly handled, and unwanted virtual multiples are generated—the latter appear as ghosts in the images.

Campillo and Paul (2003) produced the first account of geophysical, interferometric impulse response reconstruction using noise. They correlated seismic coda waves (which were assumed to be approximately diffuse due to multiple scattering) to estimate inter-receiver impulse responses. The reconstruction was only effective for the surface wave and direct arrival part of the impulse response (probably due to the near-surface distribution of common noise sources). Nevertheless, this was the main, practical breakthrough paper that encouraged geophysicists to consider using interferometric techniques in their fields.

Bakulin and Calvert (2004) produced the first practical application of seismic interferometry in an exploration setting. They showed for the first time that it is possible to create a virtual source at a subsurface receiver location (in a well) in practice. They measure the response of surface sources at downhole receivers. Using the virtual source methodology (crosscorrelating traces), they moved the sources (artificially) to downhole locations. This method has the potential to eliminate the imprint of the complex

overburden on the reflection response. The main restriction of the method is that it relies on the assumption that the downgoing primary wave at the buried receiver can be isolated from the total response which includes multiples.

The application of interferometrically constructed impulse responses for tomographic velocity inversion was shown by Shapiro et al. (2005) and by Sabra et al. (2005b). In applying the virtual source method, they crosscorrelated noise observations between many pairs of stations, thereby reconstructing the surface-wave components of impulse responses between these stations. Through tomographic velocity inversion, these surface waves were subsequently used to produce phase velocity maps indicating properties of the subsurface of southern California.

Computational wavefield modeling. Van Manen et al. (2005, 2006) showed how constructing impulse responses interferometrically on a computer provides an efficient and flexible, numerical waveform modeling scheme for heterogeneous, lossless media. By modeling the response of sources on a contour (or surface) around an arbitrary 2D (or 3D) medium, the response of any source recorded at any receiver entirely within this contour (or surface) can be obtained by crosscorrelation. This is particularly useful for nonlinear inversion or for modeling multiples, which require impulse responses between many combinations of source and receiver points within the medium. The authors also demonstrated that in a medium of any dimensionality it was necessary to store only waveforms generated by boundary sources in order to infer signals from wave propagation from any other source within the medium, and that to do otherwise was to store redundant data.

This resume of the history of interferometry shows that in a relatively small number of papers and mainly within the past six years, seismic interferometry has moved from being a field that was not popular and was based only on conjecture, to an active and extremely dynamic field in which huge theoretical and practical advances are being made almost monthly. Although we reported the advances as though they occurred in series, in fact the ideas developed more or less independently in different fields of science (ultrasonics, seismology, exploration geophysics, and oceanography). It was only in around 2003 that researchers in the different disciplines became aware of each other's work and began to see the links between them. The special supplement of *GEOPHYSICS* with its many additional advances is therefore timely, and below we present still further exciting examples that are published elsewhere.

Recent examples. We begin this section with two studies that focus on estimating aspects of Earth structure and properties, followed by three that focus on extracting and using novel types of waveforms using interferometry.

Example 1: Temporal change in velocity revealed from man-made explosions. In some applications, the goal of seismic interferometry is not to extract the impulse response of the medium from incoherent waves, but to determine the temporal change in the medium from these waves. Snieder et al. (2002) used the phrase *coda wave interferometry* for the interferometric measurement of temporal changes in the medium from strongly scattered waves. In order to do this one must have a repeatable source, but the source location and its signature need not be known.

Earthquake doublets are pairs of earthquakes with almost identical source mechanisms. The crosscorrelation of earthquake doublet records is a particular type of coda interferometry that is sensitive to differences in the medium that

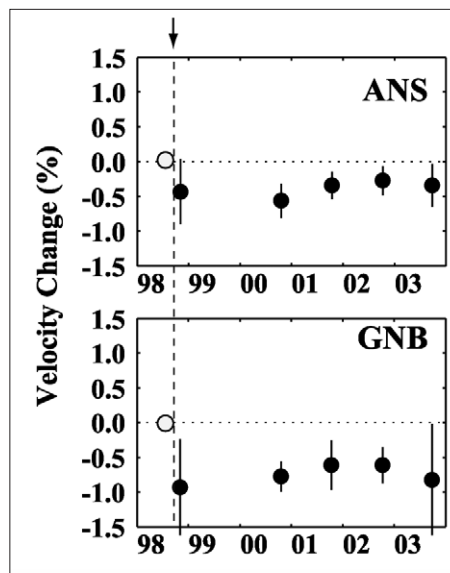


Figure 3. Temporal change in velocity for 3–6 Hz at two sites near the M6.1 earthquake focal area of Iwate volcano, northeastern Japan, revealed from the repeated artificial explosion experiments. Solid circles represent the velocity change in % relative to a shot on 10 August 1998 (open circle), where vertical lines indicate one standard deviation. (Modified from Figure 8 of Nishimura, 2005.)

took place in between the occurrence of the pair of earthquakes.

Applying such an analysis to coda wave records of earthquake doublets before and after the occurrence of the 1979 Coyote earthquake in California, Poupinet et al. (1984) first revealed a velocity drop from the phase difference plot against lapse time. Taking this a step further, here we introduce recent observation of velocity change associated with an earthquake, and a recovery process during an interval of five years after the earthquake, as revealed by interferometry of records of repeated artificial explosions.

At Iwate volcano in northeastern Honshu, Japan, in early 1998, geophysicists found an inflation of the mountain topography and an increase of seismicity associated with a dike intrusion. An earthquake took place on the southwestern flank of the volcano on 3 September 1998.

Six similar, artificial explosions were detonated at the south end of the area beneath which the earthquake ruptured from August 1998 until 2003. Applying interferometry to seismic records from the artificial explosions recorded at eight seismograph stations around the focal area, Nishimura et al. (2000, 2005) found that the average seismic velocity of the crust in the frequency range of 3–6 Hz decreased by about 1% around the earthquake focal region. This velocity drop can be explained by the dilatation caused by the earthquake, if stress sensitivity of the percentage velocity change is of the order of 0.1/MPa. From the set of successive artificial explosion experiments, a gradual recovery of the seismic velocity toward its original value was observed over the next four years. Figure 3 shows temporal change in velocity at two sites near the focal area.

It is interesting this experimental study demonstrates the extreme sensitivity of interferometry to changes in the medium. While interferometry detected this 1% change in velocity, it was unidentifiable from traveltime analysis of first arrivals.

Example 2: Teleseismic surface-wave tomography across California. In crustal seismology, most seismic noise propagates as surface waves. By crosscorrelating long noise records from pairs of seismic stations it is possible to obtain the impulse response for surface waves traveling between the two stations. Seismic noise data from 148 broadband seismic stations in Southern California (Figure 4a) were used to extract the surface-wave arrival times between all station pairs in the network in the frequency band 0.05–0.2 Hz

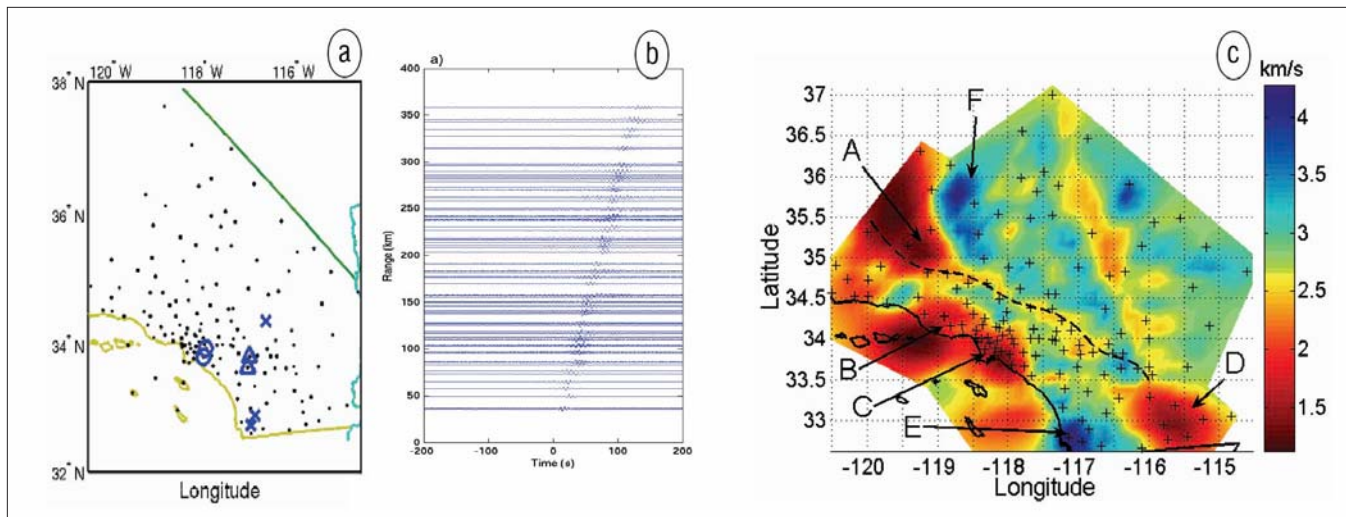


Figure 4. (a) Map of the 150 online stations in the Southern California Seismic Network. (b) Shot record generated from crosscorrelation of one month of the noise. (c) Surface-wave group velocity map for Southern California. (A = San Joaquin valley; B = Ventura; C = Los Angeles; D = Salton Sea Trough).

(Shapiro et al., 2005; Sabra et al., 2005a,b; Gerstoft et al., 2006). The station-pairs can be ordered into a so-called “passive shot” record indicating an outward-traveling wave (Figure 4b). In this area the seismic noise is strongly directionally biased, originating from the Pacific Ocean as ocean microseisms. For this reason a one-sided impulse response is obtained from interferometry of the noise (rather than an impulse response plus its time reverse as in Figure 1).

The seismic data were then used in a simple, but densely sampled tomographic procedure to estimate the surface-wave velocity structure for a region in Southern California (Figure 4c). The result compares favorably with previous estimates obtained using more conventional and elaborate inversion procedures. This study demonstrates that coherent ambient noise without an identifiable source can be used to create virtual sources at each of pairs of stations, and thereafter for seismic imaging purposes.

Example 3: Estimating building responses. Seismic interferometry is useful for extracting the response of structures, such as buildings, from the motion generated by an incoherent excitation. The left panel of Figure 5 shows the Robert A. Millikan Library in Pasadena, California. The right panel shows the horizontal component of the acceleration measured in the basement and the 10 floors after the Yorba Linda earthquake of 3 September 2002. At around $t = 10$ s, an impulsive S-wave strikes the building (blue box); at later times the building is excited by an extended surface-wave train (red box).

The waveforms in Figure 5 depend on (1) the mechanical properties of the building, (2) the excitation by the earthquake, and (3) the coupling of the building to the subsurface. Snieder and Safak (2006) separate the mechanical properties of the building from the excitation and the ground coupling by deconvolving the motion at all levels with respect to the motion at a given target level. (Here deconvolution is used rather than correlation because this operation eliminates the characteristics of the incoherent excitation more effectively than does correlation.) The result of deconvolving the waveforms of Figure 5 with respect to the waves recorded at the top floor is shown in Figure 6. In contrast to the original incoherent waves in Figure 5, the deconvolved waves in Figure 6 are coherent; they consist of the superposition of one upgoing wave and one downgoing wave. These waves can be used to measure the shear velocity and attenuation in the building.

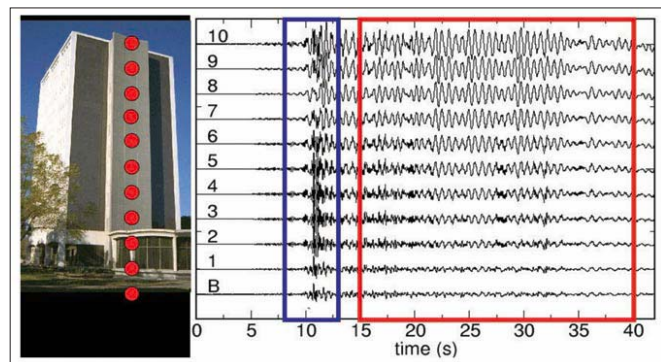


Figure 5. (left) The Robert A. Millikan Library in Pasadena, California, and the location of the accelerometers (red dots). The horizontal component of the acceleration after an earthquake is shown on the right. The red and blue boxes indicate the time intervals used for deconvolution.

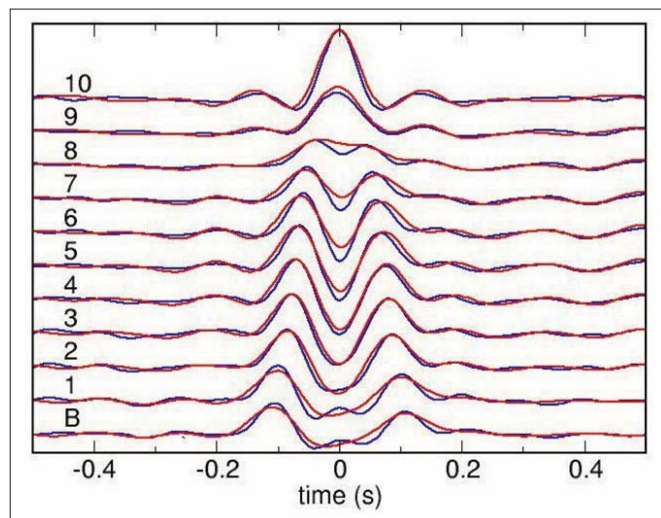


Figure 6. The waveforms obtained by deconvolving the waves at every level with the waves recorded at the top floor. Blue (red) lines indicate the waveforms obtained from the waves shown in the blue (red) boxes of Figure 5.

In fact, Snieder and Safak (2006) carried out the deconvolution process for the waves in the red box and the blue box in Figure 5 separately, and the corresponding deconvolved waves are shown in Figure 6 by red and blue lines,

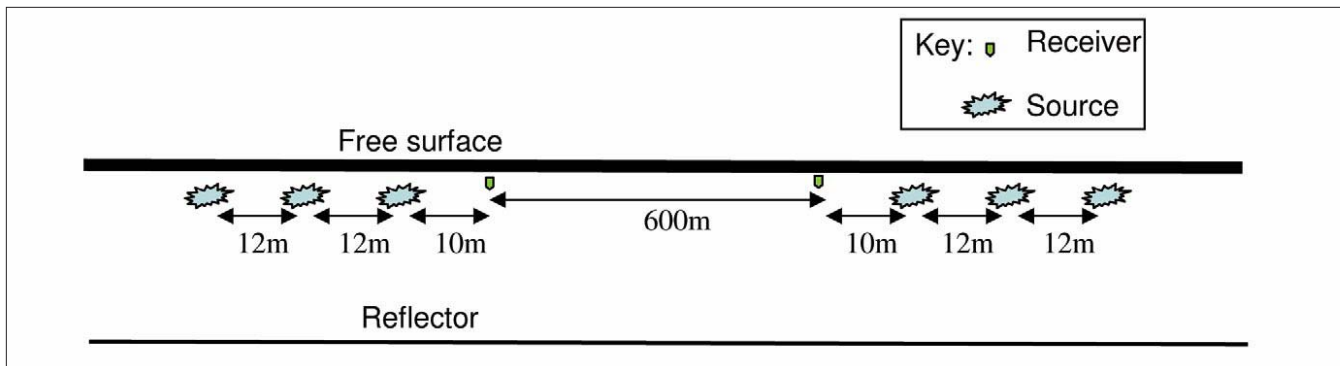


Figure 7. Source and receiver pattern for the ground-roll experiments. The earth model consisted of a strongly attenuating medium including a single horizontal layer at depth of 20 m.

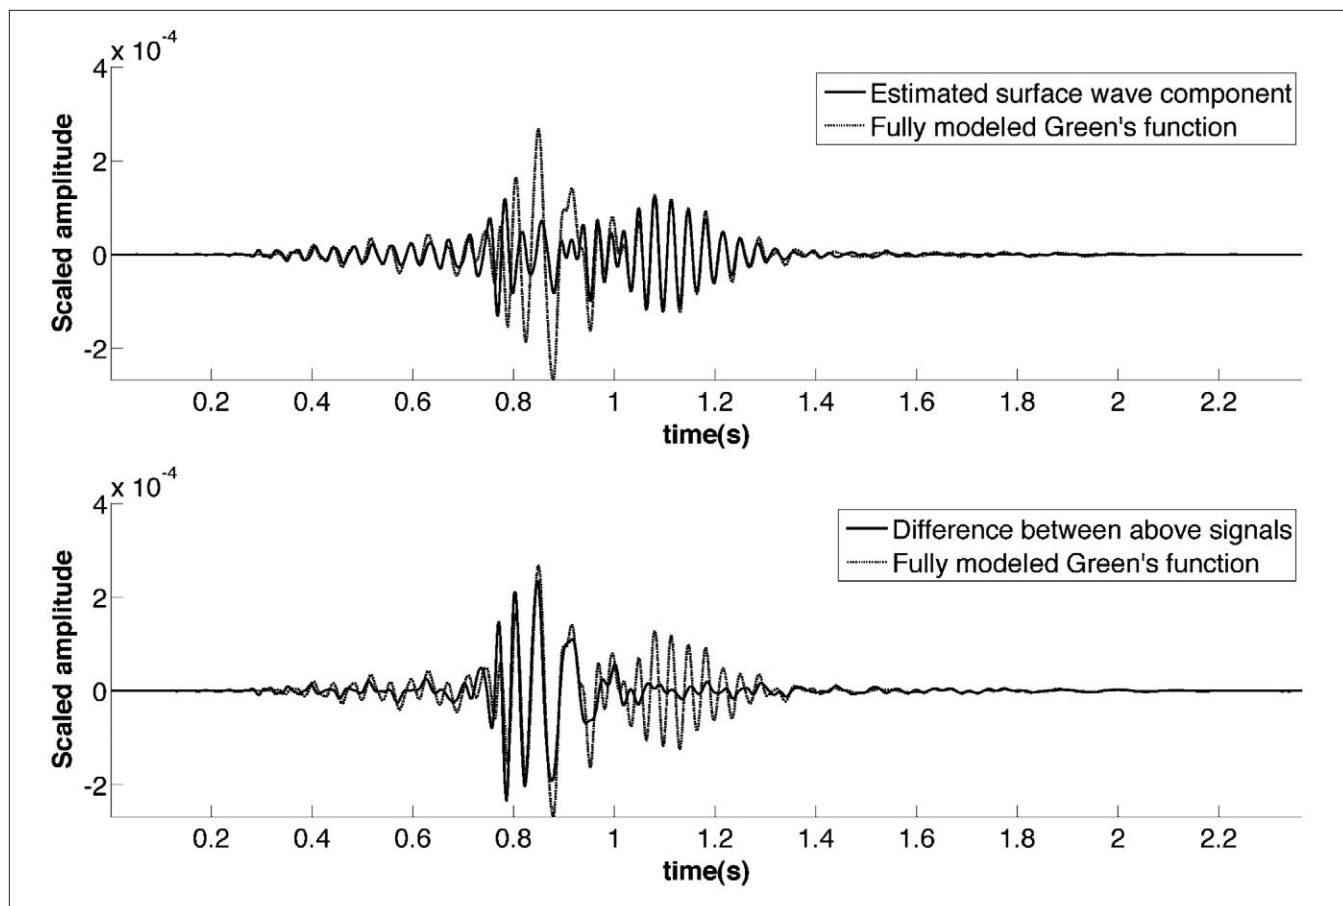


Figure 8. Traces from experiment in Figure 7. (top) Dotted line is the true inter-receiver impulse response (Green's function), the solid line is the interferometrically constructed trace using only the sources in Figure 7. (bottom) Dotted line is same as top plot, the solid line is the difference between the curves in the top plot, i.e., the full impulse response with the interferometric trace subtracted.

respectively. The original waveforms in the red and blue boxes of Figure 5 are quite different, yet the deconvolved waves are highly similar. These results show that this type of seismic interferometry does not depend strongly on the nature of the employed input signals.

Example 4: Ground-roll estimation and removal. We noted earlier that most practical studies using interferometry of background noise in the Earth successfully construct mainly surface and direct waves only, rather than reflected waves. This fact can be turned to our advantage.

Halliday et al. (2006) show that if the reason why this occurs is that typical noise sources in the Earth occur in the crust (i.e., near the surface), then it might be possible to reproduce this effect using artificial surface seismic sources. This

would be achieved by crosscorrelating and integrating traces recorded at any pair of receivers from a set of surface seismic sources that completely surround the receivers on the surface only. The result would be seismic traces between a virtual source and receiver pair at the surface that contain dominantly the surface wave train, or ground roll.

This might seem like an odd thing to do—surface waves tell us nothing about the deep subsurface where reservoirs are situated. However, for exactly this reason ground-roll removal is a common and critical processing step for seismic data. After applying the method above, the interferometrically constructed ground-roll traces can be subtracted from real traces recorded when actual sources are placed at each of the virtual source locations. Thus, we should obtain

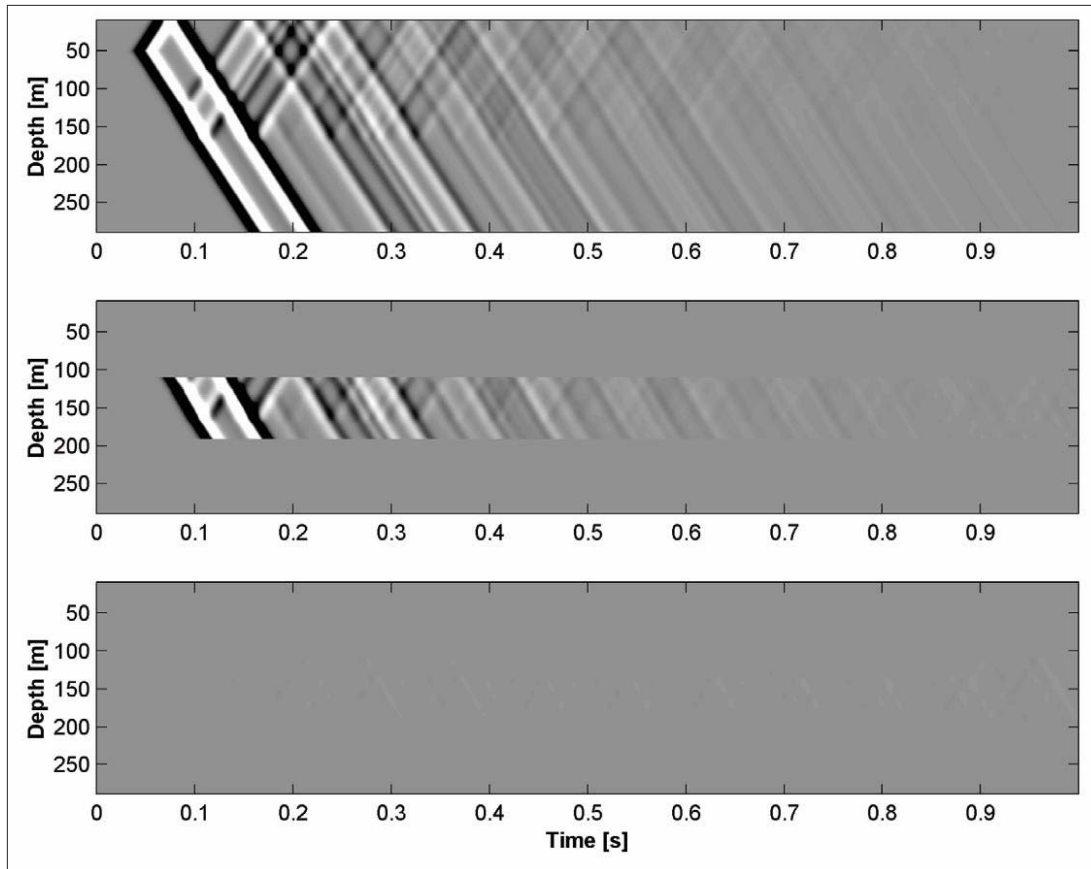


Figure 9. Shot records for a vertical array in a 1D medium. (top) Full simulation over entire model. (middle) Local simulation around strong perturbation. (bottom) Difference between the two within the local region.

seismic traces with ground roll removed.

Halliday et al. tested this on 2D synthetic acoustic and elastic examples of ground roll constructed in horizontally layered media, and media with complex near-surface topography, with and without realistic attenuation. Figures 7 and 8 show an acoustic example with a realistically attenuative medium that has only a single subsurface interface. Note that an equivalent theorem for interferometry in attenuative media has not been published to date; hence the interest in this example. Three surface sources were used to “surround” the receiver pair. The records from each source at the two receivers were crosscorrelated, and the results for the three sources were summed. The resulting interferogram, scaled by a positive constant factor, is shown at the top of Figure 8.

The arrivals in Figure 8 consist of a large body wave arrival that is the primary reflection (just before 1 s) and ground roll (after 1 s). The interferogram constructs the ground-roll component almost exactly, but not the body waves. Hence, the subtraction of the two top traces, shown in the lower plot, leaves the body waves intact but removes the ground roll.

There are several points to note about this example. First, although the structure is simple, it was not at all clear a priori that this method would work in an attenuative medium. Nevertheless, the results are impressive, as are those when using a more complex subsurface. The equivalent results for a nonattenuative medium are almost perfect.

Usually, so-called f - k filters are used to remove surface waves. However, these become increasingly ineffective as the recorded ground roll becomes more scattered from lateral reflectors or diffractors, because such scattering has the effect of spreading the ground roll over all wavevectors. On the other hand, given an appropriate source distribution, interferometry is equally effective for direct and scattered waves. This suggests that the new method of ground-roll removal could be equally effective in areas with complex near surface structures

creating strong horizontal scattering—those in which f - k filtering fails.

Example 5: Exact waveform modeling after arbitrary changes in earth models. The example in Figure 1 shows that waveforms from sources distributed around the boundary of a medium contain sufficient information to obtain impulse responses from any source point in the interior of the medium. Van Manen et al. (2005, 2006a) show how this method can be generalized to both 2D and 3D, acoustic and elastic media, producing a flexible and efficient technique for waveform simulation.

Geophysical objectives often involve generating earth models for which signals from computer-simulated seismic wave propagation fit measured seismic waveform data. We simulate waveforms for a variety of earth models and check the data fit for each. Because each simulation typically requires powerful computing resources, the computational task of performing all of the simulations quickly becomes intractable unless intelligent tricks are employed and unrealistic limitations imposed on the modeling.

However, often we are interested to test data-fit for many models that differ only within some localized region, within a modeled reservoir for example. A class of sophisticated modeling techniques exists that allows approximate waveforms for the whole model to be calculated by simulating wave propagation only in a small subvolume, just large enough to contain the region that changed (these are called *wavefield injection methods*). While this reduces the computational task dramatically, the problem with such methods is that they employ something similar to what is commonly known as the Born approximation: they assume that if the model changes within the localized region, there will be no waves that interact with (reflect off or pass through) this altered part of the model more than once. Such multiple interactions always exist in the Earth and can significantly alter wave propaga-

tion, so estimated waveforms are always incorrect.

In a new development, van Manen et al. (2006b) have shown how to model exact waveforms after any local area of a model has been perturbed arbitrarily. This is significant because it means that the only times when it is necessary to simulate waveforms through a complete model are (a) once at the beginning through any reference model, and (b) for any model that differs from the reference model over the majority of its volume. If any model differs from the reference model only in a local area, then a local simulation around only that area is sufficient to obtain *exact* waveforms for the entire model.

Importantly, it turns out that the impulse responses between internal sources and receivers illustrated in Figure 1 are exactly those required to use this new method; constructing these impulse responses using the interferometric modeling methods of van Manen et al. makes this new method efficient and flexible.

Figure 9 shows an example in which waves in a 1D medium, similar to that in Figure 2b, were simulated after a large model perturbation within a region in the middle of the model's depth range. In the top panel the complete wavefield across the whole, perturbed model was simulated to obtain the true solution. In the middle panel, only a local simulation was performed around the perturbed area using the new method.

Provided the simulation in this local area matches that in the full simulation, the new method allows us to obtain the wavefield at any other point in the medium. All previous injection techniques would produce errors in the local simulation due to multiple interactions (e.g., between the perturbed area and the Earth's surface). However, the lower panel shows the difference between the top two panels within the local region, showing that the simulation using the new method is perfect.

The future. While the examples above and in the GEOPHYSICS supplement illustrate many theoretical developments and several applications that have already been found, some significant questions and challenges remain, and several currently uninvestigated areas of application show great promise for exploration and monitoring.

The challenges mainly concern extending the theory to account for real-world media and noise. First, we obtain only correct amplitudes in the traces derived from interferometry if the distribution of impulsive or noise sources is ideal in the sense that they illuminate the region of interest from all sides equally, and if the sources are uncorrelated. Natural sources do not satisfy these geometrical requirements and are often correlated, while active sources can never surround any portion of the Earth entirely (although this may not be necessary, see the discussion around Figure 2). Consequently a major challenge is to devise methods to obtain (approximately) dynamically correct waveforms from imperfect noise sources.

Second, all the theory described and used above assumes that the medium of propagation is lossless. No such medium exists in the natural Earth, so there is a need to generalize the theory to account for attenuative media.

A third theoretical question is the following: most of the theory above was derived using crosscorrelations. However, there are situations in which deconvolution seems a more natural operation to apply—indeed it was applied in the Millikan Library example. Given that the only general proof of Claerbout's conjecture that exists was based around correlational representation theorems, is there a version of the theory based around deconvolution instead? If so, under what conditions should it be applied?

There are three main areas of application that would seem to be promising for the exploration industry. First, in principle the theory above can be applied directly to obtain impulse responses for multicomponent (MC) source and receiver surveys across a region using only MC receivers. The MC sources are created virtually by interferometry. Campillo and Paul (2003) showed that for crustal surface waves the 3×3 impulse response tensor can be determined from interferometry. Since one of the main barriers to nine-component surveys is the cost of MC sources, this would seem to be a way forward to reduce this cost. Clearly experiments are needed to test this hypothesis.

Second, the examples above show that interferograms constructed by crosscorrelating coda waves recorded at different times are extremely sensitive to small changes in the medium that occurred between recordings. This would seem to have obvious application to time-lapse seismic monitoring where changes in reservoir fluids may have very small impact on the acoustic impedance.

Third, we have mainly presented seismic applications, but presently the theory applies equally to (lossless) electromagnetic wave propagation. This possibility has barely been developed within the exploration industry but would seem promising given the recent upsurge in interest of time-domain EM methods for fluid discrimination, and its use for near-surface geophysics.

Principal message. The principal message in this article is that seismic interferometry is worth watching closely. This field shows great promise, is developing rapidly compared to other areas of seismology, and is still in its infancy. Given this context and the current level of interest, it is a safe bet that further theoretical breakthroughs and the first commercially attractive applications are around the corner. [TJE](#)

Acknowledgments: We thank the office of Basic Energy Sciences of the Department of Energy for sponsoring a workshop on "Advanced Noninvasive Monitoring Techniques" in Houston in November 2005, at which this publication was proposed.

Corresponding author: Andrew.Curtis@ed.ac.uk

Appendix. Because this is a review article which cites several papers originally published in journals not specifically dedicated to geophysics, the list of references is more extensive than for the typical TLE article and are in the form typically used by an archival journal rather than TLE's standard "Suggested reading" style.

Bakulin, A., and R. Calvert, 2004, Virtual source: new method for imaging and 4D below complex overburden: 74th Annual International Meeting, Society of Exploration Geophysicists, *Expanded Abstracts*, 2477–2480.

Callen, H. B. and T. A. Welton, 1951, Irreversibility and generalized noise, *Physical Review*, 83, 1231–1235.

Campillo, M., and A. Paul, 2003, Long-range correlations in the diffuse seismic coda: *Science*, 299, 547–549.

Claerbout, J. F., 1968, Synthesis of a layered medium from its acoustic transmission response: *GEOPHYSICS*, 33, 264–269.

Derode, A., E. Larose, M. Campillo, and M. Fink, 2003a, How to estimate the Green's function of a heterogeneous medium between two passive sensors? Application to acoustic waves: *Applied Physics Letters*, 83, no. 15, 3054–3056.

Derode, A., E. Larose, M. Tanter, J. de Rosny, A. Tourin, M. Campillo, and M. Fink, 2003b, Recovering the Green's function from field-field correlations in an open scattering medium (L): *Journal of the Acoustical Society of America*, 113, no. 6, 2973–2976.

Gerstoft, P., K. G. Sabra, P. Roux, W. A. Kuperman, and M. C. Fehler, 2006, Green's functions and surface wave tomography from microseisms in Southern California, in press *GEOPHYSICS*.

Halliday, D., A. Curtis, D.-J. van Manen, J. O. A. Robertsson, 2006, On Interferometric surface wave isolation, *Geophysical Research Letters*, Submitted.

Nishimura, T., N. Uchida, H. Sato, M. Ohtake, S. Tanaka, and H. Hamaguchi, 2000, Temporal changes of the crustal structure associated with the M6.1 earthquake on September 3, 1998, and the volcanic activity of Mount Iwate, Japan, *Geophysical Research Letters*, 27, 269–272.

Nishimura, T., S. Tanaka, T. Yamawaki, H. Yamamoto, T. Sano, M. Sato, H. Nakahara, N. Uchida, S. Hori, and H. Sato, 2005, Temporal changes in seismic velocity of the crust around Iwate volcano, Japan, as inferred from analyses of repeated active seismic experiment data from 1998 to 2003, *Earth Planets Space*, 57, 491–505.

Poupinet, G., W. L. Ellsworth, and J. Frechet, 1984, Monitoring velocity variations in the crust using earthquake doublets: An application to the Calaveras fault, California, *Journal of Geophysical Research*, 89, 5719–5731.

Rickett, J., and J. Claerbout, 1999, Acoustic daylight imaging via spectral factorization: Helioseismology and reservoir monitoring: *THE LEADING EDGE*, 18, 957–960.

Sabra, K. G., P. Gerstoft, P. Roux, W. A. Kuperman, and M. C. Fehler, 2005b, Surface wave tomography from microseisms in Southern California: *Geophysical Research Letters*, 32, L14311–L14311-4.

Sabra, K. G., P. Gerstoft, P. Roux, W. A. Kuperman and M. C. Fehler, 2005a, Extracting time-domain Green's function estimates from ambient seismic noise, *Geophysical Research Letters*, 32, doi:10.1029/2004GL021862.

Schuster, G. T., 2001, Theory of daylight/interferometric imaging: tutorial: 63rd Meeting, European Association of Geoscientists and Engineers, *Extended Abstracts*, Session: A32.

Schuster, G. T., J. Yu, J. Sheng, and J. Rickett, 2004, Interferometric/daylight seismic imaging: *Geophysical Journal International*, 157, 838–852.

Shapiro, N. M., M. Campillo, L. Stehly and M. H. Ritzwoller, 2005, High-resolution surface wave tomography from ambient seismic noise, *Science*, 307, 1615–1617.

Snieder, R., 2004, Extracting the Green's function from the correlation of coda waves: A derivation based on stationary phase: *Physical Review E*, 69, 046610-1–046610-8.

Snieder, R., A. Gret, H. Douma, and J. Scales, 2002, Coda wave interferometry for estimating nonlinear behavior in seismic velocity, *Science*, 295, 2253–2255.

Snieder R. and E. Safak, Extracting the building response using seismic interferometry; theory and application to the Millikan Library in Pasadena, California, *Bulletin of the Seismological Society of America*, 96, 586–598, 2006.

Snieder, R., K. Wapenaar, and K. Larner, Spurious multiples in seismic interferometric imaging of primaries, *GEOPHYSICS*, in press, 2006 (Special issue).

van Manen, D.-J., J. O. A. Robertsson, and A. Curtis, 2005, Modeling of wave propagation in inhomogeneous media: *Physical Review Letters*, 94, 164301-1–164301-4.

van Manen, D.-J., A. Curtis, and J. O. A. Robertsson, 2006a, SI183–SI196.

van Manen, D.-J., J. O. A. Robertsson, and A. Curtis, 2006b, On long range interactions of waves in heterogeneous media: *Science*, Submitted.

Wapenaar, K., 2004, Retrieving the elastodynamic Green's function of an arbitrary inhomogeneous medium by crosscorrelation: *Physical Review Letters*, 93, 254301-1–254301-4.

Weaver, R. L., and O. I. Lobkis, 2001, Ultrasonics without a source: Thermal fluctuation correlations at MHz frequencies: *Physical Review Letters*, 87, 134301-1–134301-4.