

Attenuation and Deconvolution

Gary F. Margrave, Devon Canada Corporation

SUMMARY

All seismic waves experience some degree of anelastic attenuation as they propagate because the earth is not a perfect, homogeneous, elastic solid. This causes the seismic wavelet to evolve as it propagates in a way characterized by progressive diminishment of high frequencies and progressive phase rotations. As a result, there is no single “wavelet” embedded in a seismic record, instead, there is a changing and evolving wavelet which has progressively lower bandwidth as traveltimes increase. Mathematically, data with this property are said to be nonstationary. In contrast, the major wavelet shaping step in seismic data processing remains stationary spiking deconvolution, an algorithm that has changed very little since its introduction some 70 years ago. This algorithm explicitly assumes that seismic data are stationary, or equivalently, that the wavelet does not evolve. Data processors often cope with this conflict between physics and algorithmic assumptions by designing the deconvolution operator over a limited time window containing the exploration target. While this can optimize the image at target, the essential nonstationarity of the data is not addressed and the result is data with characteristic wavelet distortions at times outside the design window. I present a systematic study of this issue using a sophisticated synthetic dataset created with finite-difference modelling over a 2D earth whose stratigraphy comes from well logs and whose attenuation is prescribed by constant- Q theory. The study reveals the characteristic wavelet distortions that are present in real seismic data when the processing includes only stationary methods.

INTRODUCTION

Deconvolution has been an essential step in nearly every seismic data processing flow since it was introduced by Enders Robinson (1954, 1967). It has been remarkably successful in improving seismic resolution and enabling high quality well ties. Even today, most deconvolution software closely follows Robinson’s original ideas and the term “spiking deconvolution” essentially means Robinson deconvolution. However, the purpose of this paper is to point to a major shortcoming of Robinson’s method and to argue that spiking deconvolution is largely in disagreement with the established theory of seismic attenuation (e.g. Kjartansson, 1979). Put another way, spiking deconvolution is built upon the assumption of stationarity but seismic attenuation is very nonstationary. The result is a conflict between data processing and physical theory that leads to a characteristic set of seismic image distortions. The assumption of stationarity was absolutely necessary when Robinson first proposed the method and has indeed remained so until recently due largely to limitations in computing power. That is no longer the case and it is time to move to nonstationary algorithms that come closer to matching the physics of attenuation.

The term “stationary” when applied to physical theory usually refers to a process that is invariant in either time or space. In the present context, it is time invariance that is implied and the stationary assumption in deconvolution is that the seismic wavelet is unchanging as it propagates. That is, the wavelet that encounters a shallow reflector is the same wavelet that encounters a deep reflector. There are physical effects other than attenuation that

falsify this assumption such as transmission losses. The wavelet that encounters the deeper reflector has transmitted across many more interfaces than the shallow reflector and at each one it incurs a transmission loss. These transmission losses are quantified by the transmission coefficients of each interface and these are normally frequency-independent scalars. Hence the wavelet at the deeper reflector is smaller than at the shallow reflector by a frequency-independent scale factor. However, this sort of nonstationarity is relatively easy to address by suitable time-variant gain (scaling) and is not what is of concern here. Attenuation causes a much more serious nonstationarity in which the wavelet changes simultaneously in both time and frequency. This is a much greater problem because it is computationally (and mathematically) more difficult to filter (i.e. change) a signal in a deterministic way in both time and frequency. Our concept of frequency arises largely from the Fourier transform which is the tool that moves a signal from the time domain to the frequency domain or the reverse. As a result, we tend to think of time and frequency as alternative, even opposite, domains such that we can work in one or the other but not both at the same time. Certainly, a bandpass filter defined in the frequency domain and applied there causes the signal to change in the time domain. However, an example of something that we cannot do with standard Fourier theory is to apply a filter whose pass band changes with time.

The nonstationarity induced by attenuation is progressive in the sense that the propagating wavelet continuously experiences high frequency attenuation and associated phase shifts as it propagates. The longer the propagation time the greater the attenuation and phase distortion. Figure 1a shows a theoretical prediction of a wavelet, which is initially minimum phase, propagating through a 1D attenuating medium. Careful examination shows both strength (amplitude) and shape (phase) changes as the propagation proceeds. In Figure 1b are the corresponding amplitude spectra and the progressive high frequency attenuation is clearly evident. Figure 1c shows the amplitude spectra of the attenuation process which is described by $e^{-\pi|f|t/Q}$. When spiking deconvolution is applied to a seismic trace with such an evolving wavelet, the best that we can hope for is the estimation and inversion of some sort of average wavelet. Often, the deconvolution operator design is restricted to a short time interval encompassing the zone of interest, and then the operator is applied to the entire trace. In this case, we might imagine designing an inverse for the middle wavelet in Figure 1a and then applying that inverse to all of the wavelets. The result will be optimal only for the middle wavelet and there will be systematic distortions for the others. For the earlier wavelets, the operator will be trying to remove too much attenuation while for the later wavelets it will not remove enough.

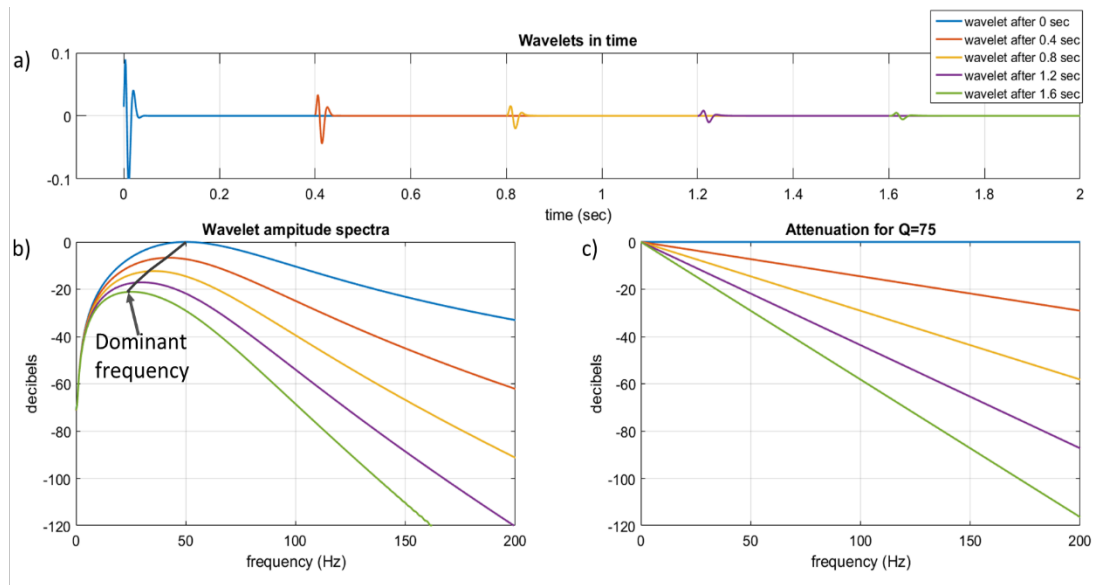


Figure 1: a) Five snapshots of an initially minimum-phase wavelet propagating through a 1D attenuating medium. b) The amplitude spectra of the wavelets. c) The amplitude spectra of the attenuation process, as described by constant Q theory, for $Q = 75$.

A motivating example

Figure 2a shows two very different appearing seismic traces which both have the property known as a white spectrum. Figure 2b shows the Fourier amplitude spectrum of both traces and illustrates what is meant by white spectrum. The term comes from analogy to white light which has equal power at all frequencies. In Figure 2b, the spectra have this property in an average sense if we were to smooth them slightly. The story behind these figures is that both traces are the result of stationary deconvolution applied to synthetic seismograms. The seismograms have the same reflectivity but one had a stationary wavelet and the other had a nonstationary wavelet. A primary design goal of deconvolution is to whiten the spectrum and clearly that has been done here but equally clearly only one of the traces in Figure 2a is an acceptable reflectivity estimate. Figure 3 shows the synthetic seismograms, stationary and nonstationary, as well as the reflectivity, and the initial wavelet. In the stationary case, the initial wavelet propagates unchanged and hence each reflector is illuminated by an identical wavelet. In the nonstationary case, the wavelet attenuates as it propagates and progressively later reflectors encounter a progressively diminished wavelet. What has happened after deconvolution, is that both traces have been spectrally whitened but the nonstationarity has not been removed. A stationary algorithm cannot remove nonstationarity.

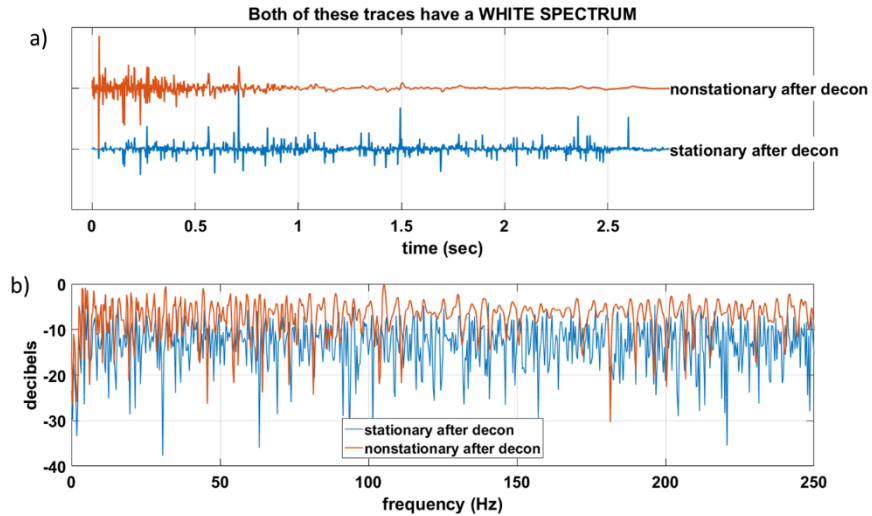


Figure 2: a) Two very different traces that both have a white spectrum. b) The Fourier amplitude spectra of the two traces.




Figure 3: The traces of Figure 2a were created by applying stationary deconvolution to the nonstationary trace and the stationary trace shown here. Both traces have the same reflectivity and same initial wavelet. The nonstationary trace also had a $Q = 75$ attenuation process applied.

The ordinary Fourier spectrum does not reveal this essential nonstationarity. However local spectra computed in small windows are more enlightening. Figure 4 shows the total spectra and the local spectra in three different windows of the two seismograms before deconvolution. In the stationary case, all four spectra have essentially the same shape the only difference being the detail induced by the reflectivity and the local spectra have lower overall power. In contrast, for the nonstationary case all four spectra have unique shapes and the increasing attenuation with increasing time is evident. Comparing the three windowed spectra, the earliest window has greater power and the most strength at higher frequencies. The later windows show the progressive diminishment of the higher

frequencies. Comparing the total spectrum to the local ones, we see that the total spectrum, rather than being an average, is most similar to the earliest window. This is because the greater strength of the early samples dominates. Figure 5 shows the same spectra but after deconvolution. In this case, the entire trace was used to design the deconvolution operator which is why the total spectrum has been whitened nicely (flattened). However, only in the stationary case are the local spectra also whitened. In the nonstationary case we can see that the local spectra are very far from the desired flattened state. The earliest window comes closest because its spectra is most like the total spectra but the other spectra are very “under whitened”. Under whitening means lower resolution in the time domain because the higher frequencies lack the power that they should have. Had we followed the common practice of using the middle window for design, then that window would be flat after deconvolution, the earliest window and the total spectrum would be “over whitened” and the last window under whitened.

To summarize what we have learned from this example:

- 1) A spectrum can be white and far from optimal even with no noise.
- 2) Local spectra can be very non white even when the total spectrum is white.
- 3) This non-ideal circumstance results from the interaction between attenuation and stationary deconvolution.
- 4) Nonstationary attenuation cannot be fully compensated for by stationary deconvolution.

This paper documents a detailed simulation of this conflict between attenuation and stationary deconvolution. First a 2D stratigraphic cross section is created by horizon-guided interpolation between well logs. Then an acoustic finite-difference modelling process was used to simulate 60 shots along a 4km line. Using a $Q = 50$ forward Q filter, both a lossless and a dataset with significant attenuation were created. Random noise with an $s2n$ (signal-to-noise ratio) of 3 was added to the dataset with attenuation. The resulting three datasets have the same geometry with receivers every 5m and shots distributed evenly along the 4km line resulting in a full fold of 30. These will be called the stationary dataset, and nonstationary dataset, and the nonstationary-noisy dataset. Then all three datasets were processed with deterministic gain, stationary deconvolution, normal-moveout removal, mute and common-midpoint (CMP) stack. The data are analyzed in time and frequency and the effects suggested above are documented. Post-stack processing is also examined.

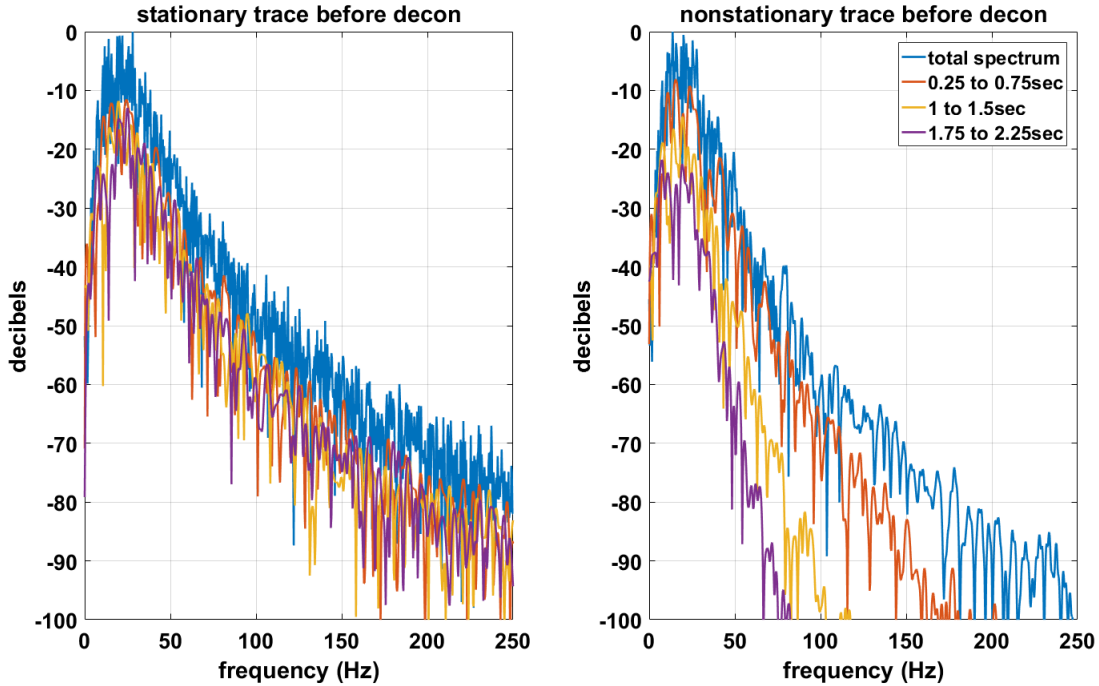


Figure 4: (left) For the stationary trace of Figure 3, the total Fourier spectrum and the local spectra in three different time windows. All four spectra have essentially similar shapes. (right) For the nonstationary trace of Figure 3, the total Fourier spectrum and the local spectra in the same three time windows. The spectra are all dramatically different.

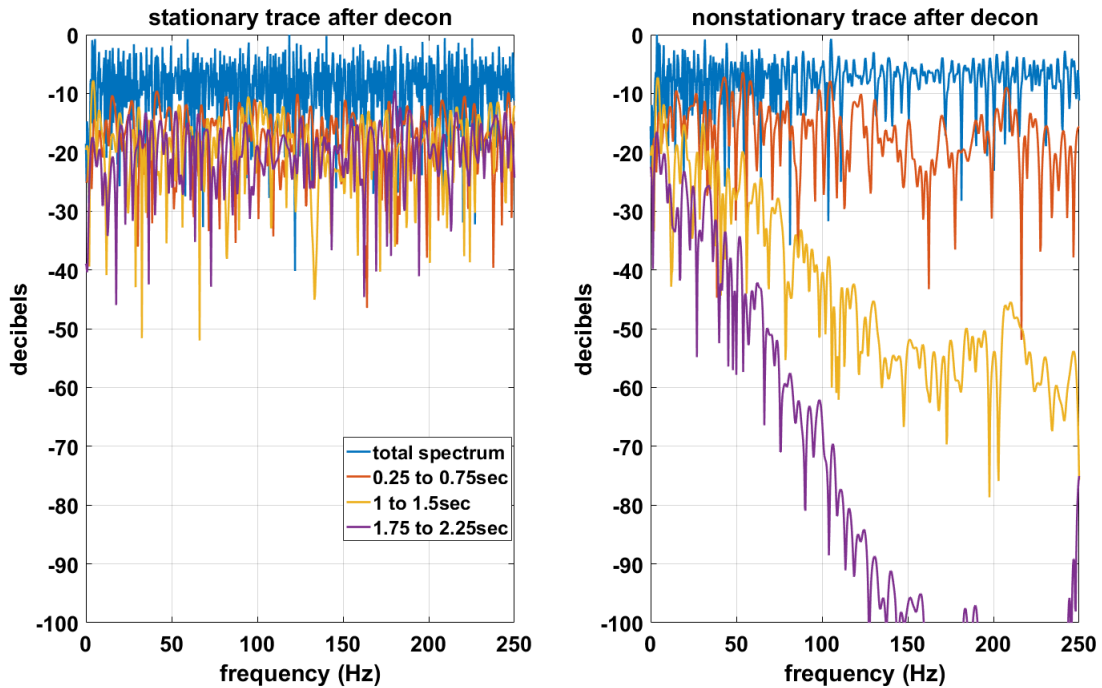


Figure 5: (left) For the stationary trace after deconvolution (Figure 2a), the total spectrum and the local spectra in three different windows are shown. All four spectra are white. (right) For the nonstationary trace after deconvolution, the total spectrum and three local spectra are shown. Only the total spectrum is white. The deconvolution operator design windows was the entire trace in both cases.

DATA CREATION, PROCESSING, AND ANALYSIS

Figure 6 shows the velocity model used in the simulation and a similar density model was also created and used but is not shown. Both velocity and density models are specified on a 2.5m square grid for 0→4000m in x (horizontal position) and 0→2000m in z (depth). As mentioned in the introduction, a total of 60 shots were simulated at positions spaced evenly across the model on the free surface. The source wavelet was a 90Hz minimum-phase Butterworth low-pass filter. All shots were recorded at identical geophone positions spaced every 5m across the top of the model. The acoustic finite difference code used for the simulation is second-order in time and fourth-order in space and can be found in the CREWES Matlab release. After simulation, a $Q = 50$ forward Q filter (including phase effects) was applied to each shot. Finally, random noise was added to the Q filtered shots with an $s2n$ (signal-to-noise ratio) of 3 measured in the deconvolution design window (see below). In this way, three datasets were created with identical geometry:

1. **Stationary:** Acoustic, noise-free data without attenuation.
2. **Nonstationary:** The stationary data with a forward $Q = 50$ filter applied.
3. **Nonstationary-noisy:** The nonstationary dataset with $s2n = 3$ random noise added.

Figure 7 shows sample shots from the stationary and nonstationary datasets. The nonstationary-noisy dataset is not shown at this stage because the noise is too low to discern.

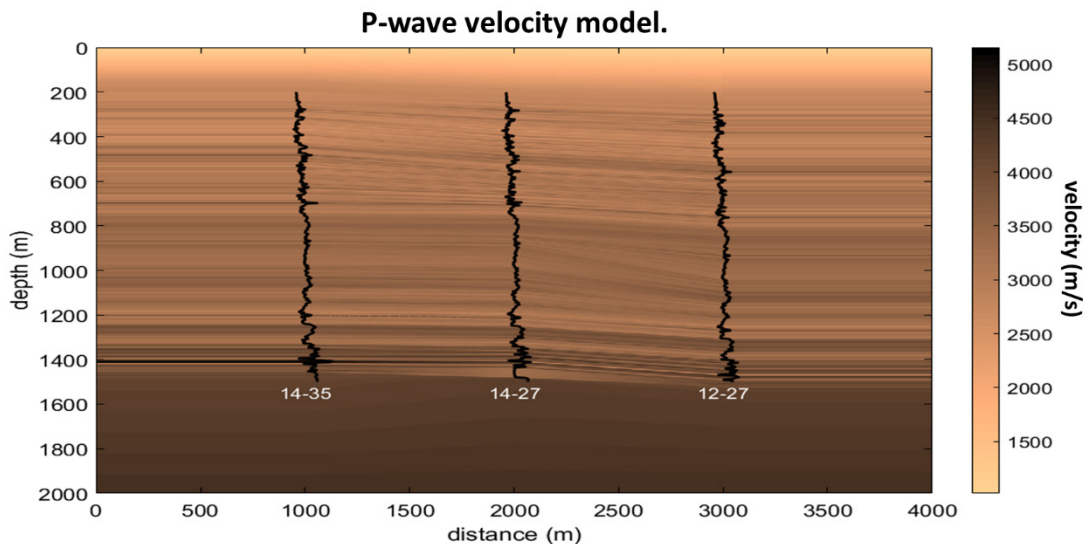


Figure 6: The velocity model used in the simulation. Shown in black are the p-wave velocity logs from the 3 wells used in the simulation. These wells are near one another on the Hussar field. In copper tones is the interpolated velocity model created from the three logs. Interpolation was guided by formation tops at each well resulting in the stratigraphic detail seen. An overburden and an underburden were attached by linear interpolation from the top or bottom of each interpolated log to fixed values at the top or bottom of the model. The result as shown is sampled on a 2.5m square grid. A similar density model was also created (not shown).




Figure 7: Sample shot records from the stationary and nonstationary datasets. The nonstationary-noisy dataset is not shown. The broad assortment of quasi-linear events radiating from the shot position are diving waves in the linear-gradient overburden. The fine spatial sampling means that these are easily removed by f-k filter.

The quasi-linear coherent noise emanating from the shot locations seen in Figure 7 are diving waves in the linear-gradient overburden. These were easily eliminated by f-k filtering. Figure 8 shows the resulting data after also deterministic gain and a mute to eliminate the direct wave. Also shown in this figure are the boundaries of spectral analysis windows. Figure 9 shows the same two shots after stationary spiking deconvolution where the operator was designed in the window indicated. This window corresponds to a target zone that is at the base of the logs used to define the stratigraphy. The stationary data has responded to the deconvolution in a uniform way at all times. Resolution (frequency content has been boosted and the amplitudes are relatively similar at all times. The nonstationary data, in contrast, shows a very nonuniform response with the amplitudes in the shallow section being most strongly boosted (this figure is displayed in “true relative amplitude” to make the amplitude distortion obvious).

Figure 10 compares the spectra of the example shot records before and after deconvolution (Figures 8 and 9). The stationary dataset behaves exactly as expected by stationary deconvolution theory. Before deconvolution the spectra are essentially identical in each window. This means that it really does not matter which window is used for operator design. After deconvolution, the spectra in all windows are whitened essentially identically. For the nonstationary dataset, things are very different. The spectra before deconvolution show the clear signature of attenuation (spectral decay) and after deconvolution only the spectrum in the third window, which is also the deconvolution design window, is whitened. The spectra in the two earlier windows receive the same boost at each frequency, but because they have less attenuation, the result is over whitening with the shallowest window being the most overdone. Also, it is important to realize that because attenuation also has phase effects, the over whitening is also accompanied by phase errors.

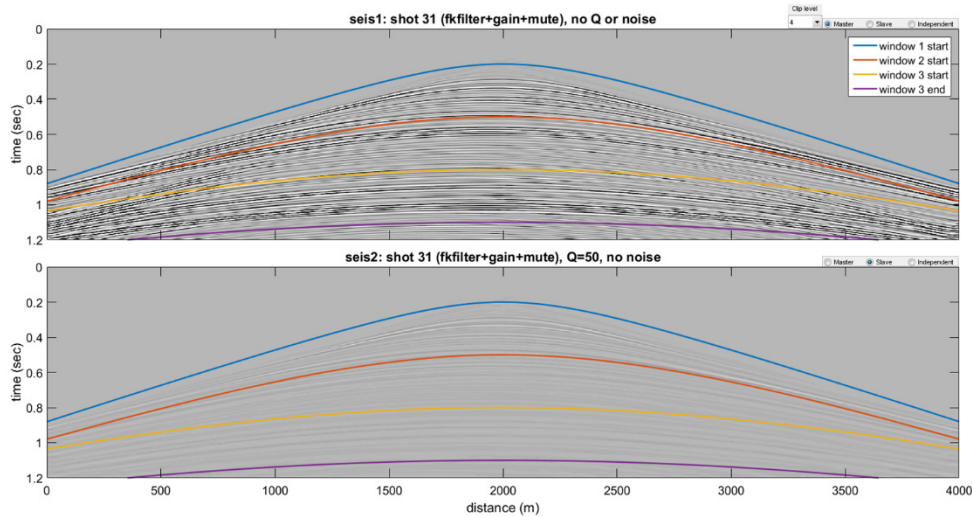


Figure 8: The shot records of Figure 7 are shown after f-k filter, deterministic gain, and mute. Also shown are the boundaries of three spectral analysis windows. These are the data input to deconvolution.

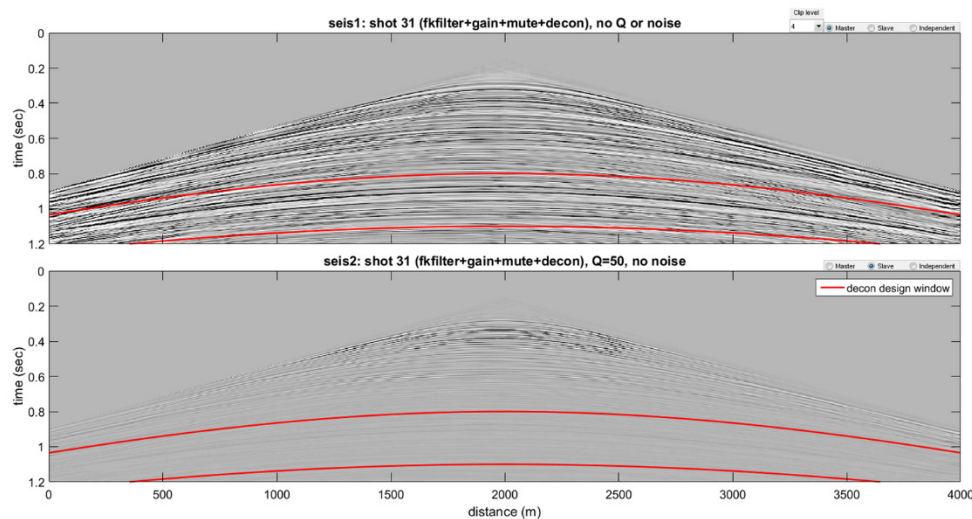


Figure 9: The shot records of Figure 8 are shown after stationary deconvolution. The deconvolution operator was 0.1 seconds long and was designed in the window shown. The design window is the same as the third spectral analysis window in Figure 8.

Figure 11 compares CMP stacks for the 3 datasets with and without deconvolution. Normal moveout removal was done using the known RMS velocities. The stacks were all properly normalized by dividing each sample by the true time-variant fold. The deconvolution was the same in all cases as just described. The stationary dataset is well balanced both before and after deconvolution with the main effect of decon being the uniform increase in resolution. The nonstationary datasets show much the same effect as the shot records with the shallow portions of the stacked traces being strongly over whitened resulting in amplitudes that are much too strong and phase distortions to be described. A major conclusion to draw from this figure is that the essential nonstationarity in the second two datasets has not been addressed by the data processing. Figure 12 shows an enlarged trace to aid in the comparison. The high-amplitudes at the early times of the

nonstationary datasets are similar to those routinely seen in almost all seismic data processing and are often called (in my experience) a “hot front-end”. This is a characteristic sign of uncompensated attenuation that is often cosmetically hidden by time-variant scaling. However, as we will see, the embedded wavelets are time-variant

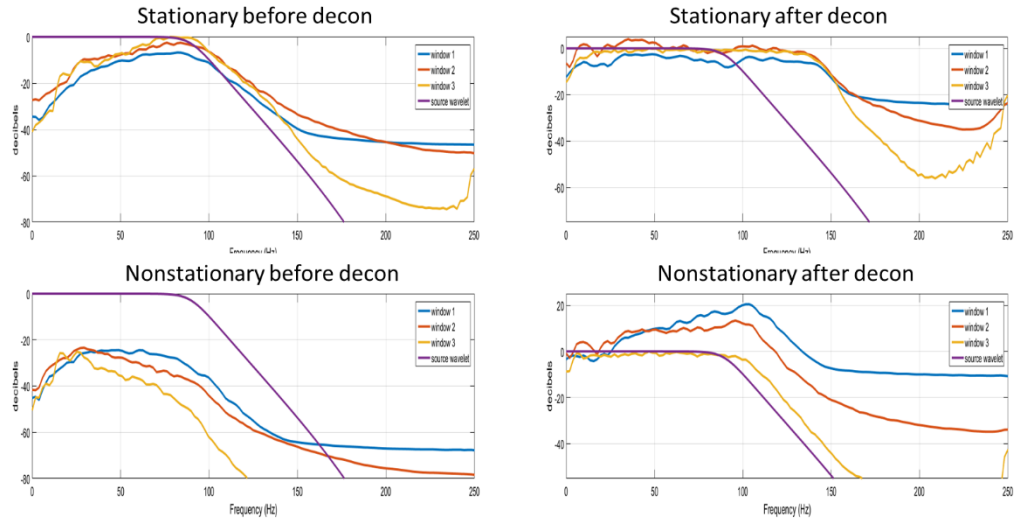


Figure 10: Spectra before and after deconvolution for the shot records shown in Figures 8 and 9. The spectra are averages over all traces computed in the windows defined in Figure 8. Also shown is the spectrum of the source wavelet which was a 90 Hz low-pass Butterworth filter.

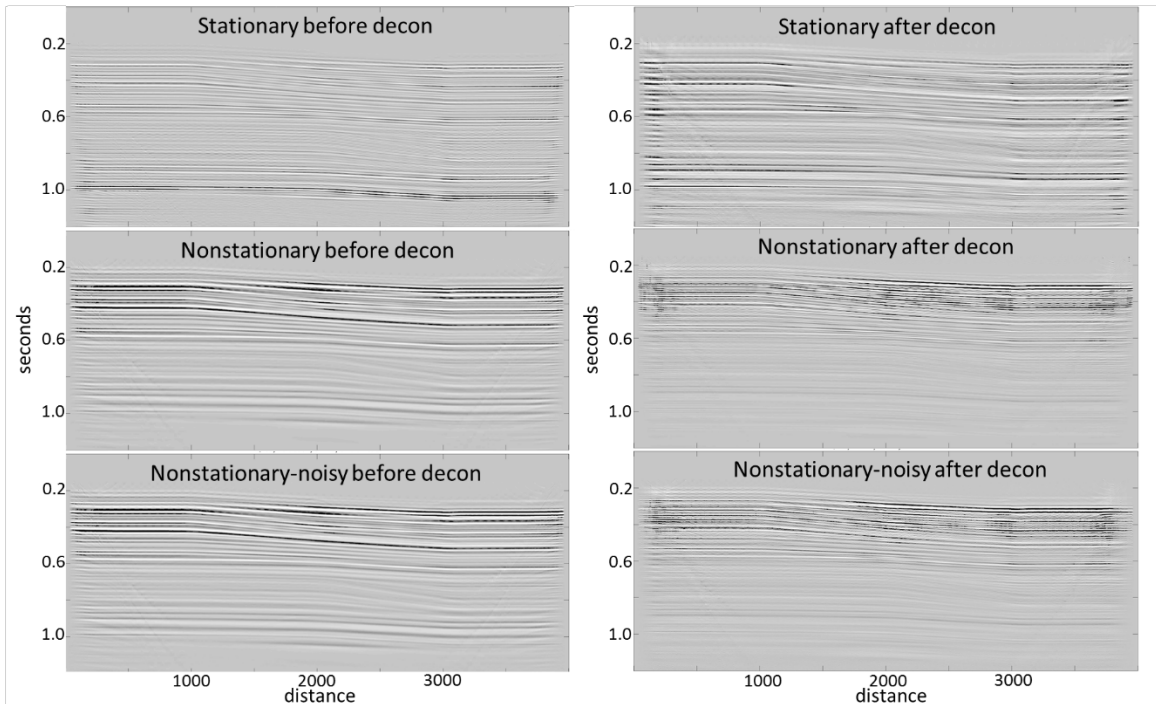


Figure 11: Comparison of the CMP stacks for the three datasets before and after stationary deconvolution. The essential signature of the nonstationary datasets is the amplitude imbalance which is not removed by deconvolution. The physics of attenuation ensures that this amplitude imbalance will also be accompanied by a phase errors.

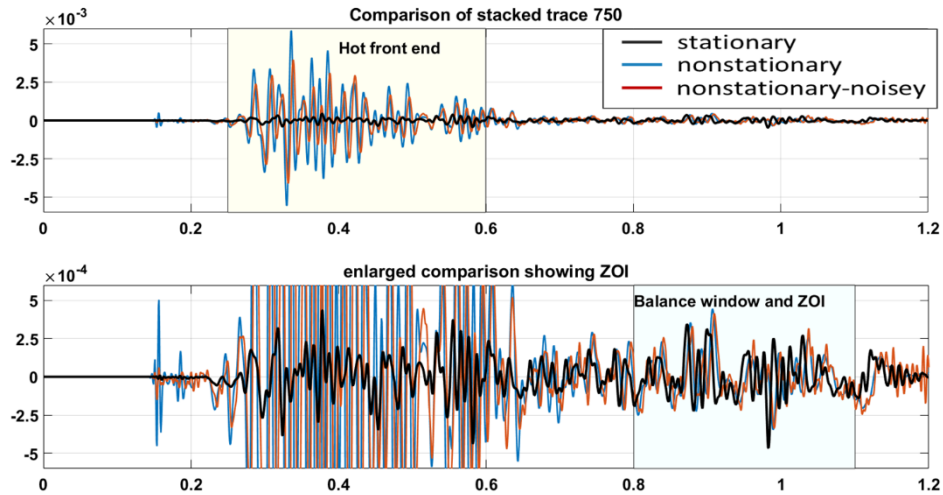


Figure 12: A comparison of a particular stacked trace after deconvolution from the three datasets. (top) A true amplitude display showing the so-called “hot front end” of the nonstationary results. (bottom) A clipped display allowing inspection of the ZOI (zone of interest) and the differing phases there.

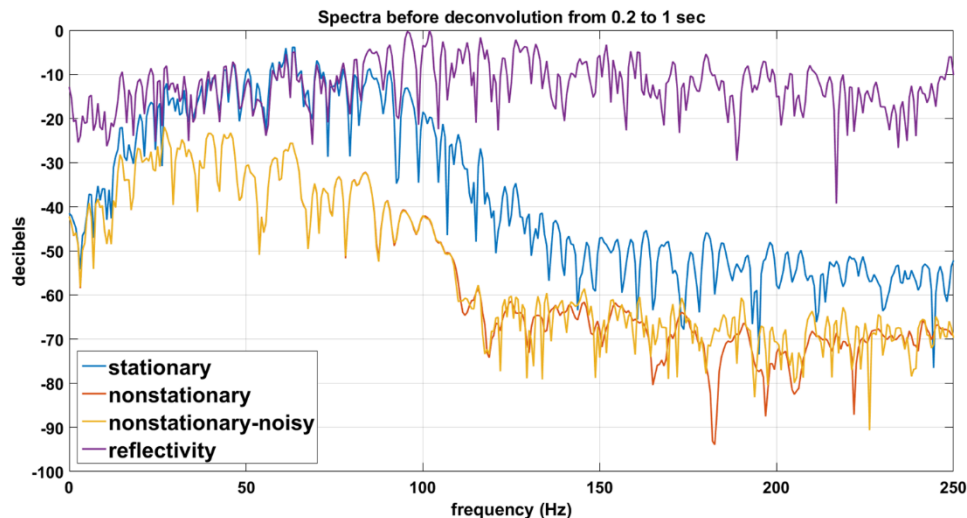


Figure 13: A comparison of the reflectivity spectrum with the spectra of the three undeconvolved stacks of Figure 11. The comparison is done at the center trace of each stack and uses the corresponding reflectivity trace computed from the velocity and density models and converted to time. The reflectivity spectrum is not white but shows a characteristic 20db roll-off between its maximum at 100Hz and 0Hz.

Figures 14-17 show the results of wavelet estimation performed on the six stacks of figure 11. Wavelet estimation was done at the center trace of each stack comparing the stacked trace to the known reflectivity in three different time zones. The wavelets were estimated by least-squares match filtering and were constrained to be causal. Since the reflectivity in this model comes from real well logs, it has spectral color (Figure 13) and so the estimated wavelets take this into account. The spectrum of the stationary wavelet before deconvolution can be visualized from Figure 13 by mentally dividing the stationary spectrum by the reflectivity spectrum. Since the spectra are roughly equal around 60Hz

we expect a dominant frequency there and a decrease in strength for frequencies above and below this value. Figures 14-15 show the wavelets in the undeconvolved stacks in both the time and frequency domains. The stationary case is again the easy and expected behaviour as intuited previously. The wavelets are the same in each time zone. The nonstationary cases show clear wavelet evolution as expected from attenuation theory. There is essentially no difference between the noisy and noise-free cases. Figure 16-17 show the wavelets after deconvolution in both domains. In the stationary case, the wavelet is still the same in all time zones but is compressed in time when compared to the wavelets before deconvolution. This is exactly what we expect from theory. It is also interesting to notice the wavelet's spectral shape before deconvolution is not the shape of a 90 Hz Butterworth low pass filter, which has a completely flat spectrum below 90Hz (Figure 10). This is because a source that radiates a wavelet $w(t)$ when placed in a volume, will radiate $w'(t)$ (the time-derivative wavelet) when placed on a free surface. In the frequency domain, the amplitude spectrum of $w'(t)$ is $f|\widehat{w}(f)|$ where $|\widehat{w}(f)|$ is the spectrum of $w(t)$. These remarks are for 3D and in 2D the free surface effect gives the so-called "square-root time derivative wavelet" which has the amplitude spectrum $\sqrt{f}|\widehat{w}(f)|$. This is what is seen here and is most easily apparent in the stationary case.

In the nonstationary case, the wavelets before deconvolution look exactly as we would expect. The shallow wavelet has the broadest spectrum and the spectra of deeper wavelets are related to the shallower ones by a factor like $e^{-\pi f \Delta t / Q}$ where Δt is the time interval between the wavelets. Notice that this factor is 1.0 at $f = 0$ and decreases exponentially as f increases. If we take $\Delta t \sim 0.3$ then an easy calculation shows that the separation between two wavelet spectra in Figure 15 is about $.164f$ decibels. So, at $f = 50\text{Hz}$ we expect about 8db separation and at $f = 100\text{Hz}$ about 16db and this is just about what is seen in Figure 15. There is very little difference between the noisy and noise-free cases.

After deconvolution, the nonstationary wavelets still show unequal spectra. The deepest wavelet, which is from the deconvolution design window, has a quasi-flat spectrum out to about 70Hz while the shallower wavelets are over whitened. The separation between the spectra of the different wavelets is similar to the separation before deconvolution (in simple theory it should be the same). In the time domain, the over whitened wavelets show far more ripples than the deepest wavelet. Are any of these wavelets "optimal"? The stationary wavelets can be taken as a measure of what is optimal and the deepest wavelet is arguably most similar to the stationary case. Perhaps the most important thing to realize is that there is no single embedded wavelet in the nonstationary case either before or after deconvolution. While the wavelet in the deconvolution design window has been handled reasonably well, outside this window there are systematic errors in both amplitude and phase. These errors will be more significant after inversion because there the solution at the target is dependent upon the trace everywhere above the target.

Finally, it is observable that none of the estimated wavelets are anything like zero phase. In fact they are all minimum phase. This follows because the source wavelet was minimum phase, the attenuation process is nonstationary minimum phase, and the deconvolution operator is minimum phase. Hence the resulting wavelets are all formed from minimum phase operators and are minimum phase. This is exactly what happens with a dynamite source in real data. Deconvolution compresses the wavelet towards time zero but it never

quite gets there. If the source wavelet was minimum phase then the embedded wavelet after deconvolution is guaranteed to be minimum phase and, with real data, it is guaranteed to be nonstationary. Since these wavelets are all minimum phase, a second pass of deconvolution, preferably after stack, is a very good option. With vibroseis data, if the common phase correction has been applied to convert the Klauder wavelet to its minimum-phase equivalent, then these remarks apply to vibroseis as well.

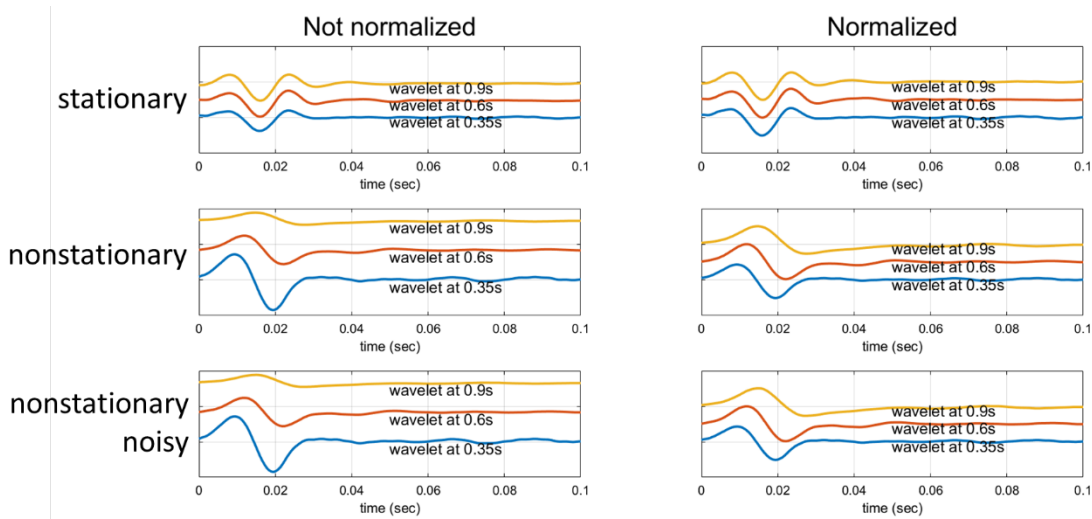


Figure 14: Estimated wavelets before deconvolution. These were estimated from the undeconvolved stacks of Figure 11 using Wiener match filtering to known reflectivity at the center trace in the stack. “Normalized” simply means that all of the wavelets have their maximum amplitude equalized to 1.0 so allow the wavelet shapes to be compared.

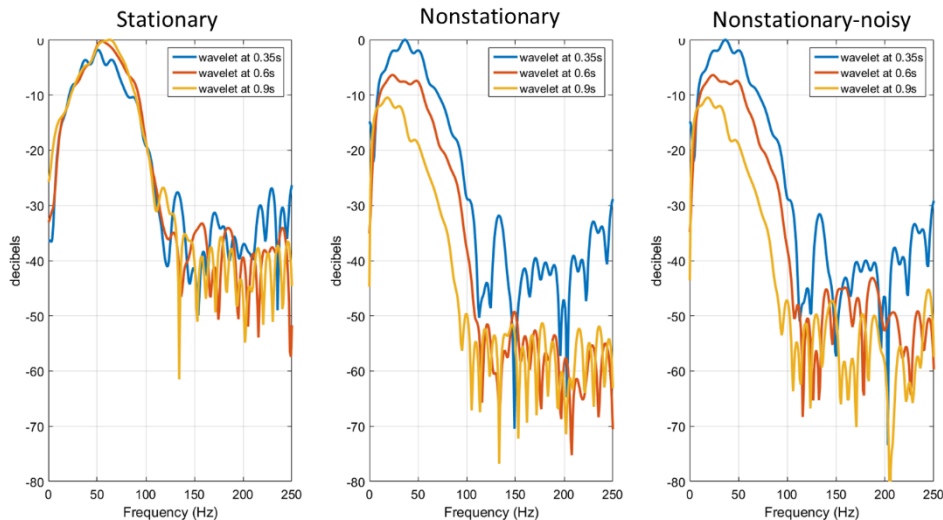


Figure 15: The spectra of the wavelets in Figure 14 are shown.

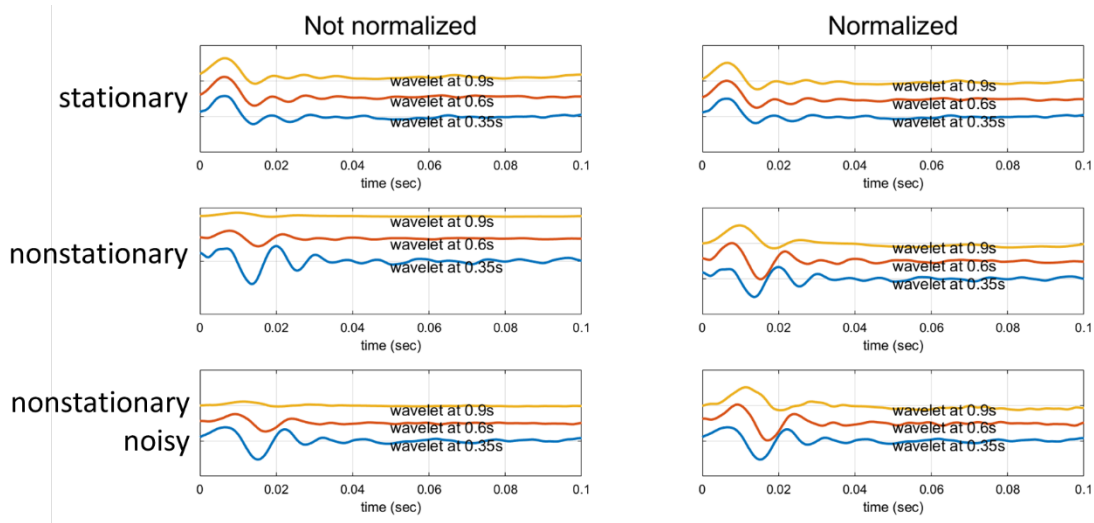


Figure 16: Estimated wavelets after deconvolution. These were estimated from the deconvolved stacks of Figure 11 using Wiener match filtering to known reflectivity at the center trace in the stack.

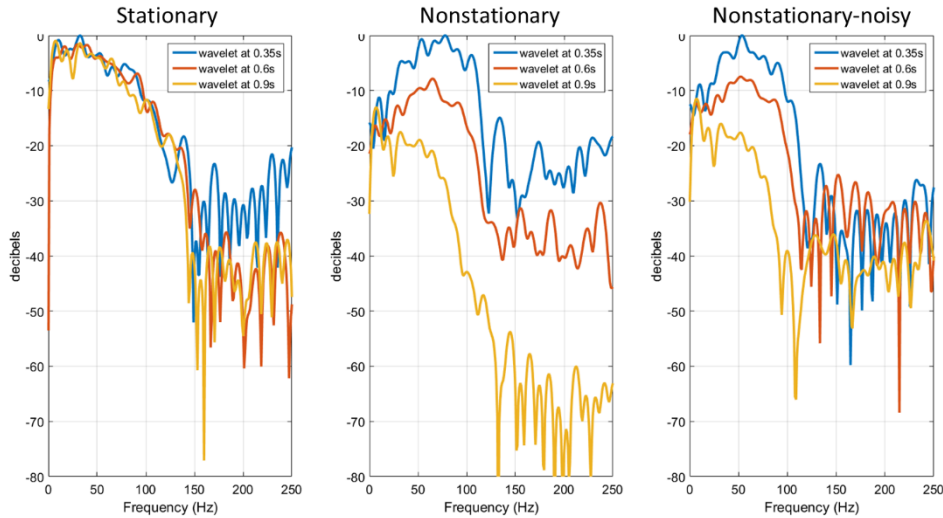


Figure 17: The amplitude spectra of the wavelets in Figure 16 are shown. Only the stationary case is optimal at all times. In the nonstationary case, the deepest wavelet (gold color) has a nearly optimal spectrum (similar to the stationary case) while the shallower

DISCUSSION AND CONCLUSIONS

Anelastic attenuation of seismic data is a nonstationary process that is always present and which causes the seismic wavelet to continuously evolve as it propagates. The wavelet evolution is characterized by progressive loss of amplitude that is an exponential decay as a function of both time and frequency. The process is known to be nonstationary and minimum phase so that the amplitude loss is accompanied by phase changes. The standard wavelet shaping tool in data processing is stationary spiking deconvolution and it has been demonstrated here that spiking deconvolution is fundamentally incapable of addressing the nonstationarity caused by attenuation. Post-decon, or post-stack, amplitude scaling is merely a cosmetic adjustment that hides the fundamental problem of a nonstationary wavelet. While predictable from theory, these conclusions were demonstrated here by a

very realistic synthetic dataset that used real well logs for stratigraphy and a spatially constant Q . Shot records (60 of them) were created using acoustic finite-difference modelling. Three distinct datasets were created (1) stationary (no attenuation), (2) nonstationary without noise, and (3) nonstationary with noise. The datasets were then processed through to CMP stack using stationary deconvolution as well as other standard tools. CMP stacks were generated with and without deconvolution to demonstrate the effects of that tool. Spectra were examined before and after deconvolution and before and after stack. Wavelets were estimated after stack by matching to the known reflectivity. It was conclusively demonstrated that the stationary data responds properly to stationary data yielding an optimal result but the nonstationary data results are far from optimal. The wavelets estimated post-stack without deconvolution show a clearly evolving wavelet as expected from attenuation theory. Wavelets estimated post-stack with deconvolution show an optimal wavelet only in the decon design window while elsewhere there are systematic and predictable amplitude and phase distortions. This distortions will not be resolvable with a standard well-tying that estimates only a single wavelet.

ACKNOWLEDGEMENTS

I thank the sponsors of CREWES for their support. I'm especially grateful to Devon Energy for their support and to my colleagues at Devon for their insight and suggestions.

REFERENCES

- Kjartansson, E., 1979, Constant Q -wave propagation and attenuation: *J. Geophys. Res.*, 84, 4737–4748.
Robinson, E. A., 1954, Predictive decomposition of time series with applications to seismic exploration: Ph.D. Thesis, MIT.
Robinson, E. A., 1967, Predictive decomposition of time series with applications to seismic exploration: *Geophysics*, 32, 428-484.