Instantaneous spectral analysis: Detection of low-frequency shadows associated with hydrocarbons

JOHN P. CASTAGNA and SHENGJIE SUN, University of Oklahoma, Norman, U.S.
ROBERT W. SIEGFRIED, Gas Technology Institute, Des Plaines, Illinois, U.S.

Instantaneous spectral analysis (ISA) is a continuous time-frequency analysis technique that provides a frequency spectrum for each time sample of a seismic trace. ISA achieves both excellent time and frequency localization utilizing wavelet transforms to avoid windowing problems that complicate conventional Fourier analysis. Applications of the method include enhanced resolution, improved visualization of stratigraphic features, thickness estimation for thin beds, noise suppression, improved spectral balancing, and direct hydrocarbon indication. We have seen four distinct ways in which ISA can help in the detection of hydrocarbons: (1) anomalously high attenuation in thick or very unconsolidated gas reservoirs, (2) low-frequency shadows in reservoirs where the thickness is not sufficient to result in significant attenuation, (3) preferential illumination at the "tuning" frequency which can be different for gas or brine-saturated rocks, and (4) frequency-dependent AVO. In this paper, we describe the ISA technique, compare it to other spectral decomposition methods, and show some examples of the use of ISA to detect low-frequency shadows beneath gas reservoirs.

The ISA method involves the following steps: (1) decompose the seismogram into constituent wavelets using wavelet transform methods such as Mallat's matching pursuit decomposition, (2) sum the Fourier spectra of the individual wavelets in the time-frequency domain to produce "frequency gathers," and (3) sort the frequency gathers to produce common (constant) frequency cubes, sections, time slices, and horizon slices. The results can be viewed using animation techniques available in commercial interpretation and visualization packages.

Figure 1 shows a synthetic seismic trace and the corresponding ISA time-frequency analysis. The time-frequency plot shows amplitude spectra for each time sample. We refer to this kind of plot as a "frequency gather." The first arrival on the synthetic seismogram results from an isolated reflector. The frequency spectrum is the spectrum of the wavelet. Note that the duration of the spectrum is identical to the duration of the arrival in the time domain as opposed to Fourier-based methods in which the time duration is equal to the window length. The second event is a composite of two events of differing center frequency arriving precisely at the same time. The frequency spectrum indicates a low-frequency arrival spread over time and a higher-frequency arrival that is more localized in time. The third event is caused by two interfering arrivals of the same frequency. Although the presence of two arrivals is not immediately apparent on the seismogram, the time-frequency decomposition clearly shows two distinct arrivals. The fourth event is a composite of four waveforms arriving at two distinct times evident on the time-frequency analysis. The final event consists of three arrivals of the same frequency that are very closely spaced in time. The three distinct arrivals cannot be resolved at low frequencies, but the separation is clearly evident on the time-frequency analysis at high frequency. It is apparent that ISA provides a useful representation of the information contained in a seismic trace.

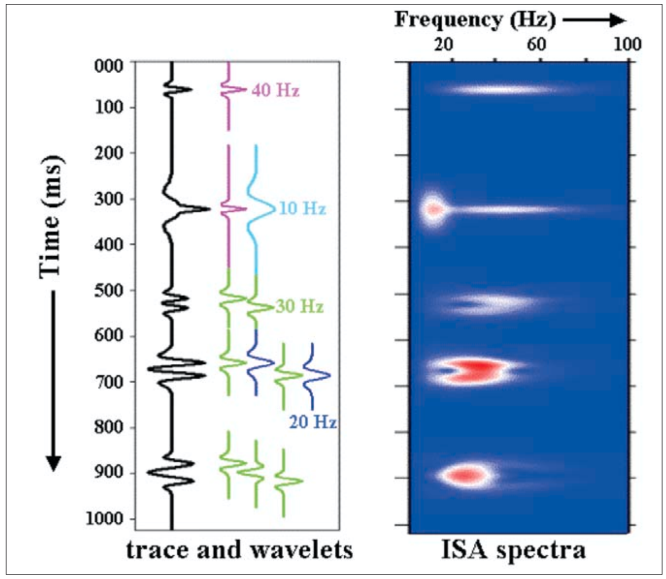

Figure 1. Synthetic waveform with transient arrivals (black seismogram), constituent wavelets (color coded by center frequency), and time-frequency analysis (red represents high amplitude). At any given time, the spectrum of the seismogram is the weighted superposition of the wavelet spectra. Since the exact wavelets summing to form the seismogram are known, exact spectra versus time may be computed.

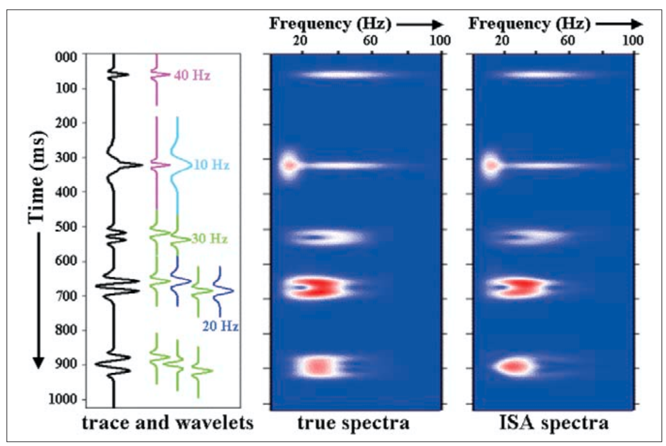

Figure 2. Comparison of true spectra and ISA for a synthetic waveform.

Spectral decomposition methods. Various techniques have been utilized in time-frequency analysis. Traditionally, the Fast Fourier transform (FFT) and discrete Fourier transform (DFT) have been applied (DFT having the advantages of greater speed and not mandating transform lengths that are a power of 2 as required by FFT). Both techniques have limited vertical resolution because the seismogram must be windowed. The spectral energy is distributed in time over the length of the window, thereby limiting resolution. If the time window is too short, the spectrum is convolved with the transfer function of the window, and frequency localization is lost (i.e., the frequency spectrum is smeared). This can be mitigated to some extent by tapering the window, but it is obviously preferable to avoid windowing altogether. Another disadvantage of a short window is that side lobes of arrivals appear as distinct events in the time-frequency analysis. If the time window is lengthened to improve frequency resolution, mul-

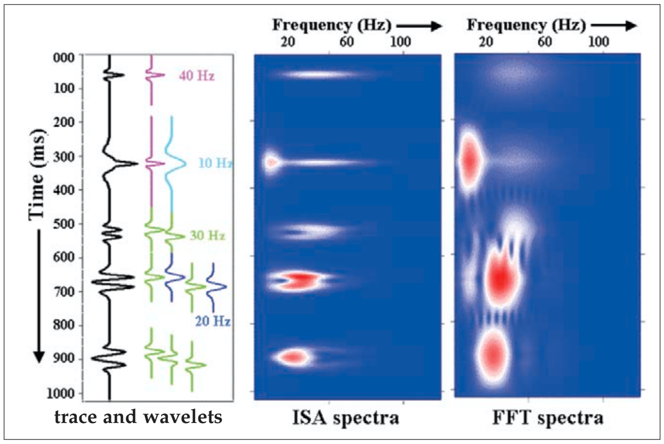

Figure 3. Comparison of ISA and time-frequency decomposition obtained with FFT using a 200-ms window. Notice the poor vertical resolution and spectral distortion (ribs and notches) caused by FFT windowing.

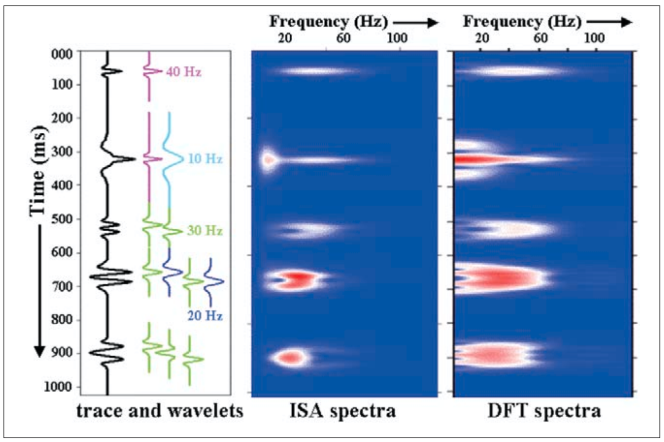

Figure 4. Comparison of ISA and time-frequency analysis obtained with DFT using a short time window. Although DFT with a short window has excellent vertical resolution, the frequency spectrum has been smoothed by convolution with the spectrum of the window and false events are associated with side lobes of transient arrivals.

tiple events in the window will introduce notches that dominate the spectrum. Long windows thus make it very difficult to ascertain the spectral properties of individual events. This is one reason why Q-attenuation over short intervals is difficult to measure quantitatively using Fourier-based techniques.

Recently, the maximum entropy method (MEM) has been used for spectral decomposition. This technique can achieve excellent frequency resolution but can be unreliable if the signal violates the assumptions of the method or if the window is too short. The major disadvantage, in our experience, is that it seems to be unstable especially for less-than-expert users. Note in the example presented in this article that, even for the first arrival, there are two peak frequencies.

Wavelet transforms decompose a seismogram into constituent wavelets. The superposition principle tells us that the frequency spectrum of a seismogram is the sum of the frequency spectra of the wavelets that sum to produce that seismogram. At any given time, the frequency spectrum is the superposition of weighted wavelet spectra in the vicinity of that time sample. In a somewhat overlooked paper published in GEOPHYSICS in 1995, Chakraborty and Okaya showed how wavelet transforms can be used in time-frequency analysis. They utilized matching pursuit decomposition to produce high-resolution time-frequency analyses. However, our experience with this method suggests that it introduces artifacts into the time-frequency analysis manifested as high-amplitude bursts at a given time over a wide frequency band or at a given

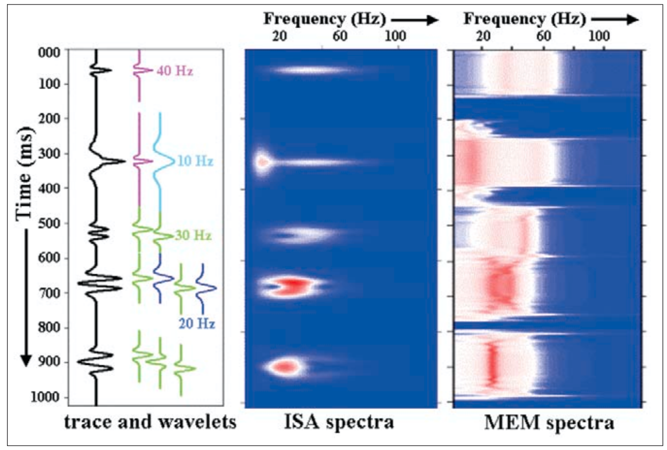

Figure 5. Comparison of ISA and time-frequency analysis obtained with MEM using a long time window. Although MEM at times has superb frequency resolution, vertical resolution is limited by the window length necessary for reliable spectral determination.

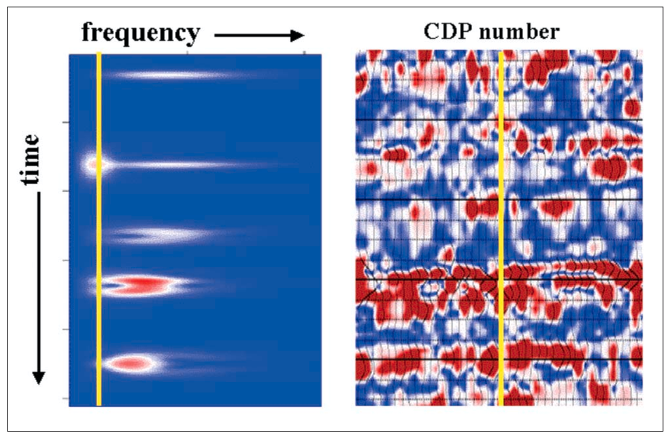

Figure 6. A frequency gather (left) and a common frequency section (right) obtained by sorting many frequency gathers according to frequency. The common frequency gathers can be thought of as instantaneous amplitude at a given frequency.

frequency over a long time interval, thereby superimposing a cross-hatch pattern on the time-frequency analysis. This is a consequence of computational shortcuts designed specifically for rapid computation made possible by selecting a particular wavelet dictionary. We prefer to select the wavelet dictionary to better capture the features of the seismogram while selecting parameters judiciously and avoiding as many cross-correlation operations as possible to achieve reasonable computation time. Thus, wavelet-transform based instantaneous spectral analysis (ISA) can be done accurately with acceptable speed while simultaneously achieving excellent time and frequency resolution.

Comparison of methods. When comparing different spectral decomposition techniques, it is useful to have "the right answer" as a guide. When dealing with synthetic data, constructed as a superposition of wavelets, the "true" time-frequency analysis is readily calculated as the sum of the spectra of the known wavelets. As shown in Figure 2, the ISA technique does not yield the true spectrum precisely because wavelet decomposition is not unique in very much the same way that seismic inversion is not unique. This fact can be used to prove that time-frequency analysis itself is not unique. Many different time-frequency decompositions can result from the same seismogram and, conversely, can be inverse transformed to produce the same seismogram. This then leads to the question: How can a particular frequency-decomposi-

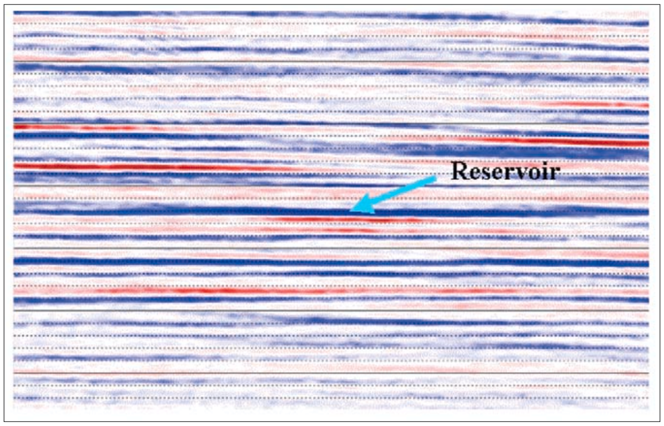

Figure 7. Broad-band migrated stacked section for offshore Tertiary clastic section. Troughs are blue, and peaks are red. The reservoir (arrow) is a classic bright spot (low-impedance gas sands with a characteristic leading trough). No shadowing beneath the reservoir is apparent. Timing lines represent 20 ms.

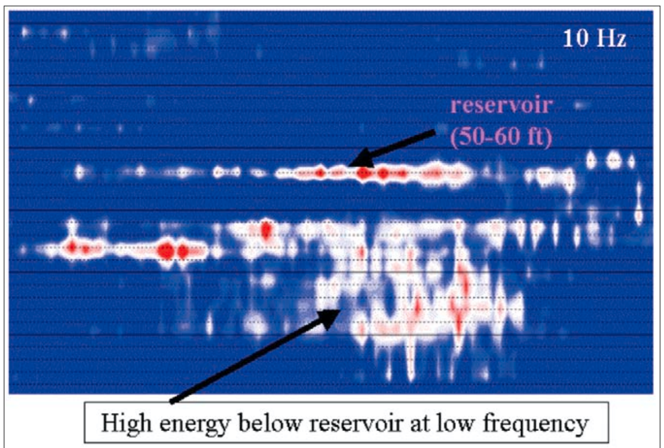

Figure 8. 10-Hz common frequency section corresponding to the broad-band section in Figure 7. Significant low-frequency energy occurs beneath the reservoir but is absent elsewhere. Timing lines are 20 ms.

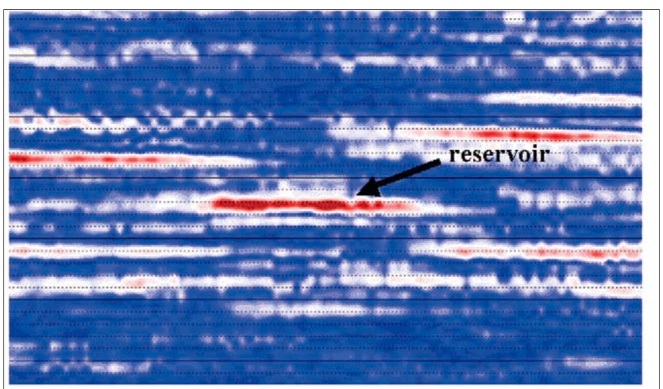

Figure 9. 30-Hz common frequency section corresponding to the broad-band section in Figure 7. The low-frequency shadow in Figure 8 has disappeared. Events immediately below the reservoir appear somewhat attenuated. Timing lines are 20 ms.

tion be judged as better than another if there is no unique answer? In our opinion, the important question is not which decomposition is right or wrong, but which captures the essential features important in interpretation. We believe it is important that the following criteria are met:

1) The sum of the time-frequency analysis over frequency should approximate the instantaneous amplitude of the seismic trace.

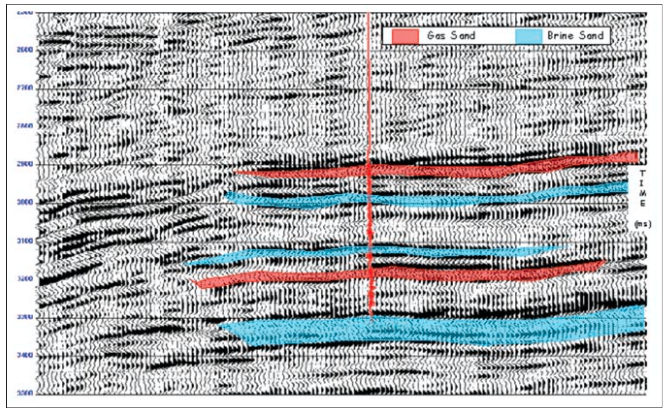

Figure 10. Broad-band seismic section from NW Shelf of Australia. Gas sands are pink and brine sands are blue.

2) The sum of the time-frequency analysis over time should approximate the spectrum of the seismic trace.
3) Distinct seismic events should appear as distinct events on the time-frequency analysis. In other words, the vertical resolution of the time frequency analysis should be comparable to the seismogram. The time duration of an event on the time-frequency analysis should not differ from the time duration on the seismogram.
4) Side lobes of events on the seismogram should not appear as separate events on the time-frequency analysis.
5) The amplitude spectrum of an isolated event should be undistorted. The spectrum should not be convolved with the spectrum of the window function.
6) There should be no spectral notches related to the time separation of resolvable events.

A sure sign of improper spectral decomposition is a large dc (zero frequency) component for given events on data that have had low-cut filters applied and should have little low-frequency energy. This is usually caused by spectral smear due to windowing.

By definition, the ISA technique is designed to mathematically conform to criteria (1) and (2). As the ISA method involves no windowing of the seismogram it also does well at meeting criteria (3) to (6). As can be seen in Figure 2, ISA does a good job of capturing the essential features of the true time-frequency spectra. One caveat: the more appropriate the selection of the wavelet dictionary, the better the time-frequency representation. An inappropriate wavelet dictionary will cause the ISA method to fail in the sense that criteria (3) and (4) will not be achieved.

Figure 3 compares ISA and FFT spectra for the same synthetic trace discussed previously. The FFT was calculated with a 200-ms window. Notice that the FFT spectral energy is spread out over the window length. When multiple arrivals occur in the same window, severe spectral notches appear as well as "ribs" between events. The location of the notches and periodicity of the ribs depend entirely on the location of events in time and tell little about the spectral characteristics of individual reflectors. Certainly, closely spaced arrivals cannot be resolved in time on the FFT time-frequency analysis. Their spacing in time can only be inferred from location of spectral notches—this being a strongly model dependent inference. The temporal resolution of the FFT can be improved by reducing the window length, the cost being loss of frequency resolution and smearing of the spectrum (Figure 4). In this case, DFT was used rather than FFT because the window length was not a power of 2. At first glance, DFT appears to have resolution comparable to ISA. Part of this improvement is illusory, however, as false events (that can be associated with side lobes of

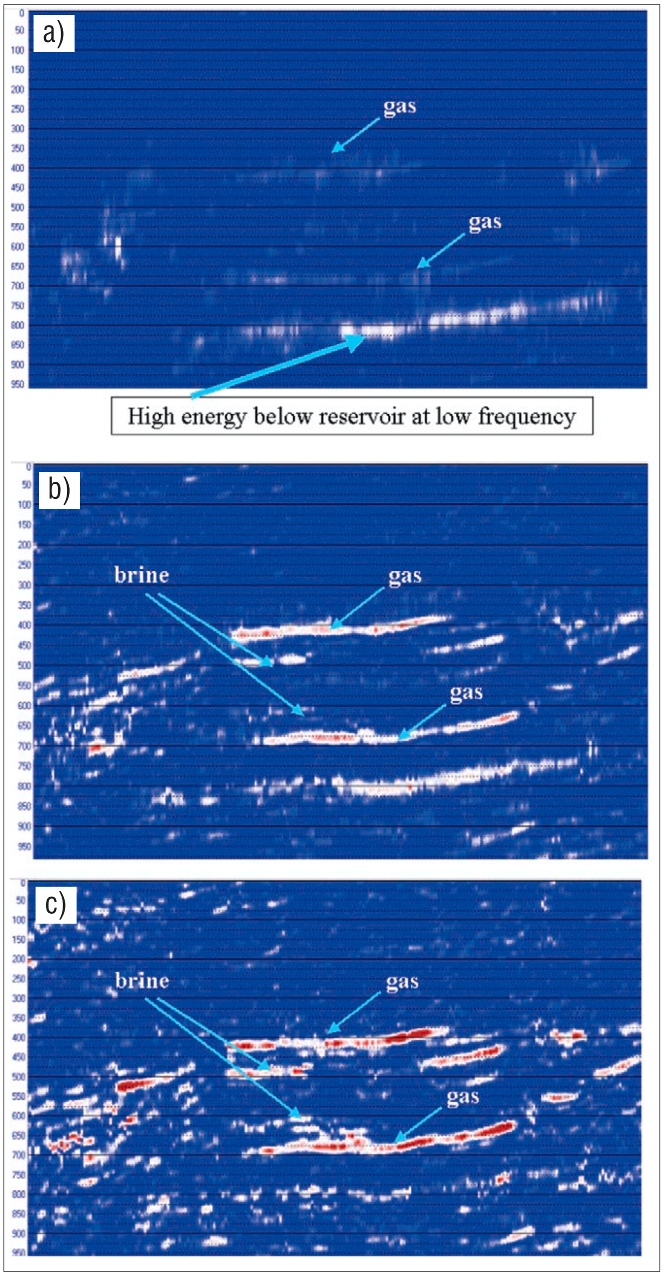

Figure 11. *(a) 10-Hz common frequency section corresponding to the broad-band seismic section in Figure 10. The low-frequency shadow beneath the lower gas sand is the strongest event at 10 Hz. (b) 20-Hz common frequency section corresponding to the broad-band seismic section in Figure 10. The low-frequency shadow beneath the lower gas sand is now weaker than the overlying gas sands. (c) 30-Hz common frequency section corresponding to the broad-band seismic section in Figure 10. The low-frequency shadow beneath the lower gas sand is now gone and the gas sands are the strongest events on the section.*

arrivals) are evident at low frequencies. Furthermore, it can be seen that frequency resolution has been lost and the spectra have been spread out over a broader frequency band than is actually present in the data.

Superb frequency resolution can be achieved with MEM (Figure 5); however, the method becomes unstable when short windows are used. Thus, in practice, the temporal resolution that can be achieved is relatively poor.

Low-frequency shadows. Since the inception of bright spot technology in the 1960s, low-frequency shadows beneath amplitude anomalies have been used as a substantiating hydrocarbon indicator. These shadows are often attributed by explorationists to abnormally high attenuation in gas-filled reservoirs. However, it is often difficult to explain observed shadows under thin reservoirs where there is insufficient travel path through absorbing gas reservoir to justify the observed shift of spectral energy from high to low frequencies. At the 1996 SEG/EAGE Summer Research Workshop, Dan Ebrom summarized at least 10 mechanisms that can explain these low-frequency shadows. Other than intrinsic attenuation, we think that one or more of the following mechanisms for introducing low-frequency shadows may be at work at any given time: stacking in of locally converted shear-waves and peg-leg multiples; NMO stretch of far-offset information; improper moveout correction and consequent loss of

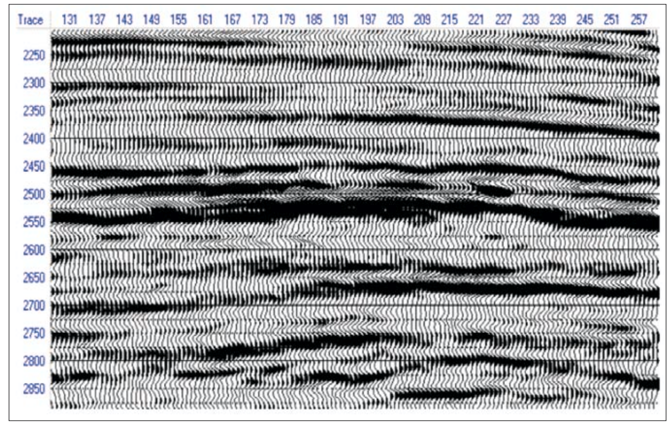

Figure 12. *Broad-band seismic section for offshore Gulf of Mexico bright spot. The top of reservoir is the trough at 2500 ms at trace 190. Low-frequency shadowing is not particularly apparent on the broad-band data.*

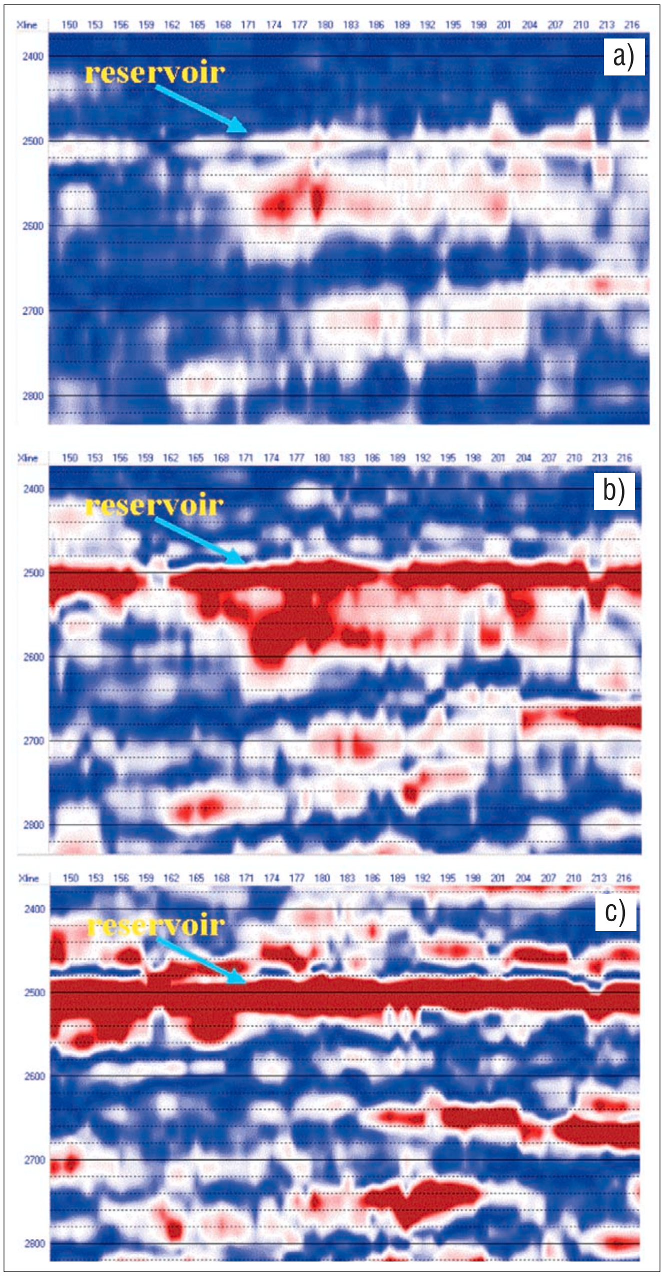

Figure 13. *(a) 8-Hz common frequency section corresponding to the broad-band seismic section in Figure 12. The low-frequency shadow just beneath the reservoir is the strongest event on the section. The reservoir extends from xline 165 to 200. (b) 12-Hz common frequency section corresponding to the broad-band seismic section in Figure 12. The low-frequency shadow and the reservoir have comparable amplitude, but the extent of the shadow better defines the reservoir dimensions (from xline 165 to 200). (c) 20-Hz common frequency section corresponding to the broad-band seismic section in Figure 12. The low-frequency shadow is completely attenuated.*

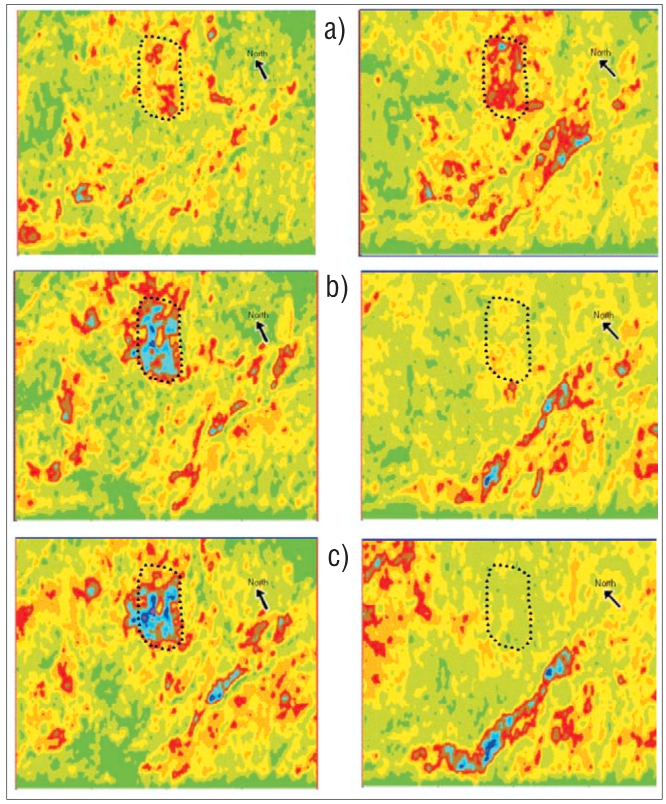

Figure 14. *(a) 6-Hz common frequency horizon slice (left) on the top of a reservoir (dotted black line) and for a 50-ms window immediately below the reservoir base (right). At 6 Hz, two strong events are apparent—one associated with the reservoir dimensions (a shadow) and another indeterminate elongate feature toward the lower left. (b) 14-Hz common frequency horizon slice on the top of a reservoir (left) and for a 50-ms window immediately below the reservoir base. At 14 Hz, the reservoir is now bright, the low-frequency shadow is gone, and the indeterminate elongate feature toward the lower left remains. (c) 21-Hz common frequency horizon slice on the top of a reservoir (left) and for a 50-ms window immediately below the reservoir base. At 21 Hz, the reservoir remains bright, the low-frequency shadow is gone and the indeterminate elongate feature is starting to look like a channel.*

high frequencies upon stacking; and time-varying deconvolution where the gas bright spot is in the design window.

In the examples which follow, the frequency gathers are sorted into common frequency cubes, sections, and horizon slices (Figure 6). Each common frequency display is thus the spectral amplitude for that frequency versus time.

The first example is a bright spot from the Gulf of Mexico (Figure 7). The reservoir has a characteristic leading-trough (blue) on the broadband seismic data and is slightly brighter than nearby events. Figure 8 shows the corresponding ISA section at 10 Hz. The reservoir is anomalously bright at this frequency, but what is most intriguing is the zone of abnormally strong low-frequency energy beneath the reservoir. At 30 Hz (Figure 9) the reservoir is clearly defined, though less anomalous in amplitude, and the energy under the reservoir apparent at 10 Hz is gone.

The next example (Figure 10) from the NW Shelf of Australia exhibits two distinct gas reservoirs. At 10 Hz (Figure 11a) the brightest event on the section is beneath the deeper gas pay. We believe this to be a low-frequency shadow. At 20 Hz, the gas reservoirs are brighter than the shadow which still persists. The shadow has completely disappeared at 30 Hz (Figure 11c).

Returning to the Gulf of Mexico, a weak amplitude anomaly can be seen at the crest of the structure in Figure 12. Figure 13 shows the relative variations in event amplitudes as frequency increases. At 8 Hz the strongest event is a low-frequency shadow under the reservoir which extends to the reservoir limits (Figure 13a). At 12 Hz, the top reservoir-sand reflection is now strong and the shadow persists (Figure 13b). At 20 Hz the shadow is gone (Figure 13c).

Viewing frequency-dependent effects in map view is also very revealing. Figure 14 shows frequency-dependent horizon slices at the top of the reservoir (left) and for a 50-ms time window immediately beneath the reservoir. The reservoir dimensions are outlined by the dashed contour. At 6 Hz (Figure 14a) the reservoir amplitude is not particularly bright. The deeper window shows a strong shadow under the reservoir but also other strong energy to the lower left of the reservoir. At 14 Hz (Figure 14b) the reservoir is a clear bright spot, the shadow is gone, and the high-frequency energy to the lower left persists (indicating that this energy has another origin). At 21 Hz, the energy to the lower left has developed a crisp channel-like character showing that it is a stratigraphically older geologic feature unrelated to the reservoir.

Discussion and conclusions. In this article we have shown that the ISA method has a much better combination of temporal and frequency resolution than conventional spectral decomposition methods. This enables the use of ISA as a direct hydrocarbon indicator. Low-frequency shadows are much more apparent on spectrally decomposed data than on broadband seismic sections.

For every example shown, the shadow was stronger than the reservoir reflection at lower frequencies, suggesting that shadows are not necessarily a simple attenuation phenomenon because low-frequency energy must have been added or amplified by some physical or numerical process. Attenuation alone should simply attenuate higher frequencies, not boost lower frequencies. Furthermore, attenuation should be less localized in time than what we are observing. We believe that these shadows are caused by one or more of the mechanisms described by Ebrom.

Conventional amplitude analysis is made at the arbitrary and largely accidental dominant frequency of the seismic data resulting from a complex interaction of acquisition parameters, earth filtering, and data processing. When one sees how amplitudes change with frequency, the inadequacy of "broadband" amplitude analysis becomes immediately apparent. Explorationists who construct conventional horizon amplitude maps need to ask themselves if they should continue to generate maps at the accidental dominant frequency of the data when their prospect may in fact be more anomalous at some other frequency. In our opinion, such broad-band amplitude anomaly mapping will become obsolete in the not too distant future.

It should also be apparent that the old idea of "tuning thickness" is also obsolete. Because we can investigate the data at any frequency, there is no single tuning thickness. Explorationists need to think in terms of the "tuning frequency" for a given reservoir, not the tuning thickness for a given seismic data set. This will be discussed in a future article.

Suggested reading. As a general introduction to spectral decomposition we recommend various papers by Greg Partyka or Kurt Marfurt in GEOPHYSICS. For an introduction to wavelet-transform based spectral decomposition, Chakraborty and Okaya's 1995 GEOPHYSICS paper "Frequency-time decomposition of seismic data using wavelet-based methods" is a must read. "The low-frequency gas shadow on seismic section" by Ebrom is available in the proceedings for the 1996 SEG/EAGE Summer Workshop on Wave Propagation in Rocks. TLE

Acknowledgments: Thanks to the Gas Technology Institute and Fusion Geophysical for financial backing in the development and evaluation of ISA. Seismic data were provided by ChevronTexaco. Thanks are also owed to Sven Treitel who doesn't even know that he gave me the clue I needed to do this in a conversation relating the Fourier transform to correlation with sine waves. Special thanks are due to my late father John F. Castagna. I didn't believe him when he told me years ago that geophysicists were doing things all wrong and that we needed to look at things a single frequency at a time. Thanks, Dad, for being so arrogant in suggesting the ridiculous and for teaching me to do the same.

Corresponding author: castagna@ou.edu