Chapter 2

ROBUST SIGNAL AND NOISE SEPARATION

2.1 Introduction

The dominant events in the data from the OGS experiment are noise: either uncorrelated noise from trace to trace, or strong coherent noise from sources off the vertical plane containing the seismic array. The goal of the processing sequence described in this chapter is to enhance events in the range of moveouts expected for the wavefield emitted by a source at depth.

Examination of the data suggests a classification of noise patterns into three groups depending on the length of their spatial correlation. The different groups include first, noise that occurs on isolated traces; second, narrowband noise on a few adjacent traces; and, third, noise present over a large part of the array.

To suppress each of these three types of noise, I follow three different approaches. I attenuate the spatially uncorrelated noise with weighted median filters, which unlike linear filters are known to remove high-amplitude noise bursts without smearing. The weights of the median filter are chosen to limit the effect of the filter on the spatially correlated events. Next, to avoid suppressing drill-bit signal in further steps of the processing, I apply velocity filtering to separate the data into two parts. Events outside a range of moveouts expected for the drill-bit signal contain only noise. From these events, I estimate first the narrowband noise by linear prediction along the time axis. Then I estimate the noise with large spatial correlation by linear prediction along the offset axis. Finally, I use the two noise estimates from the velocity-filtered data as regressors in a least-squares problem, in order to attenuate the noise in the original data.

To illustrate each step of the processing, I display all intermediate results for data in 6 control windows, chosen for their different noise patterns. The final results of the processing show that the noise has been well suppressed and reveal events with hyperbolic coherency that were hard to see in the raw data. To emphasize that these hyperbolas are not an artifact of the processing, I apply the processing sequence to field data that have no correlation along offset; then, reassuringly, the output is also spatially uncorrelated.

To check the effect of the processing on a large volume of data, I compare amplitude and dip spectra before and after noise attenuation. After noise attenuation, the amplitude spectra are smoother. In the dip spectra, computed as a function of offset and dip, significant energy remains only in the range of dips where drill-bit signal is expected.

2.2 Examples of data

A typical window of seismic data recorded during drilling is displayed in Figure 2.1. The most conspicuous events are the spatially uncorrelated (random) noise and the spatially correlated noise at the far offsets. The spatially correlated noise at the far offsets is present on only a few traces, and has distinctive linear moveout and narrowband frequency content. Events with hyperbolic moveout, as expected from the drill-bit signal, are not readily apparent. The only processes that have been applied to these data are power equalization (each trace has the same power in a window of 25 seconds), highpass filter with cutoff frequency[1] at 8 Hz, and an adaptive notch filter that attenuates strong monochromatic noise[2] generated by the mud pumps (Kostov, 1988).

Figure 2.3 displays six other windows of data chosen to illustrate significant changes that have occurred in the noise patterns over a period of 4 h. During this period of time, the depth of the drill bit increased by about 50 m (from 780 m to 840 m). Seismic data were recorded for a total time of 10 min during this 4 h interval.

To understand better the signal and noise separation method and to be able to establish limits of validity depending on the signal-to-noise ratios, I apply the same processing sequence to synthetic data (Figure 2.2). The window of synthetic data is 1 s long, and has

[1]The goal of the 8 Hz cutoff is to attenuate ground-roll with apparent velocity as low as 0.3 km/s; noise with wavenumbers below 25 km^{-1} is not attenuated by the group arrays (Figure B.2).

[2]The variation of the frequency component removed by the notch filter is seen around 20 Hz in the right-hand side panels of Figure 2.17.

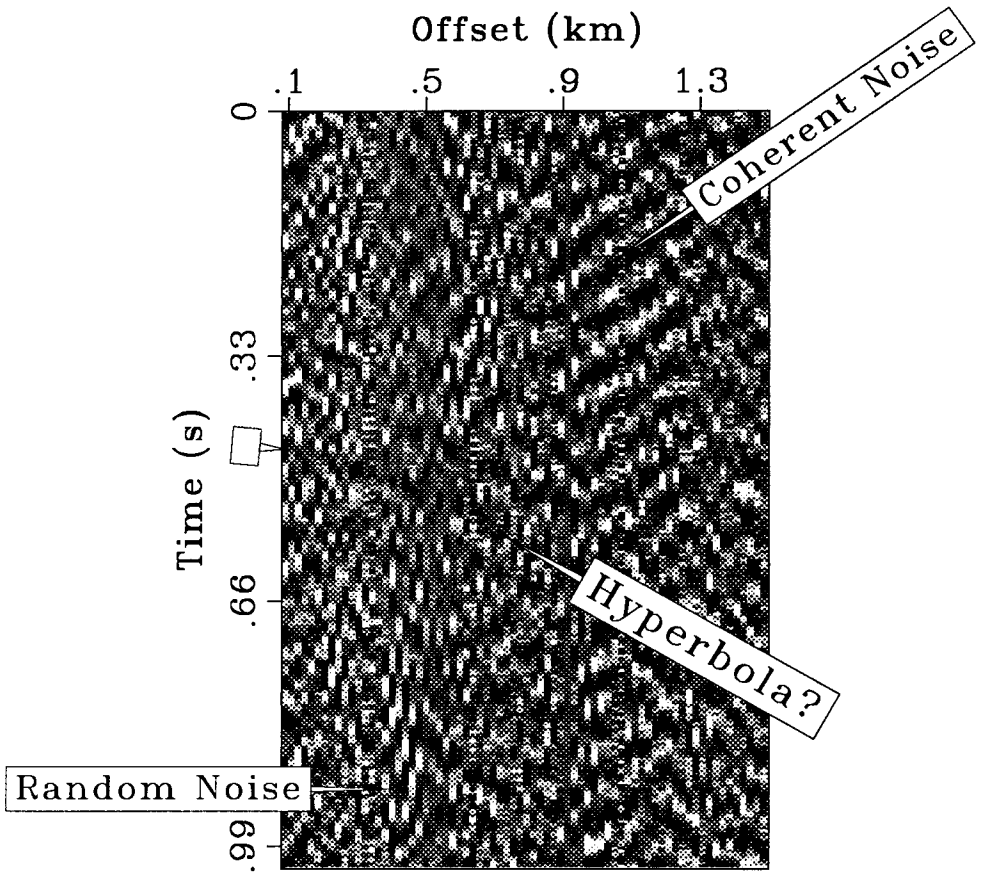

FIG. 2.1. A typical window of data illustrates random noise (uncorrelated from trace to trace), strong coherent noise at the far offsets, and weak events with possibly hyperbolic moveout. The coherent noise is propagating toward the well, its dominant frequency is 16 Hz, and its apparent velocity is about 4 km/s. The presumed hyperbolic event also has a high apparent velocity, greater than 3 km/s.

the same sampling parameters as the windows of data from the OGS experiment shown in Figures 2.1 and 2.3. The synthetic data are the sum of signal and coherent noise. I model the signal as the sum of two hyperbolas with equal amplitudes and convolved with the same wavelet (minimum phase Ricker wavelet, fundamental frequency of 40 Hz); one of the hyperbolas has a moveout[3] as expected for a point source at depth of 0.8 km in a medium of RMS velocity of 2 km/s, while the second hyperbola corresponds to a point source at depth of 1 km and an RMS velocity of 3 km/s. The noise is the sum of strong narrowband noise (16 Hz) on 7 adjacent traces in the middle of the gather, and, at the far

[3]See also the discussion on "Moveout correction," page 16.

offsets, of broadband noise with hyperbolic moveout corresponding to a source located off the vertical plane containing the array. The broadband noise data are a window of marine seismic data (Yilmaz and Cumro, 1983).

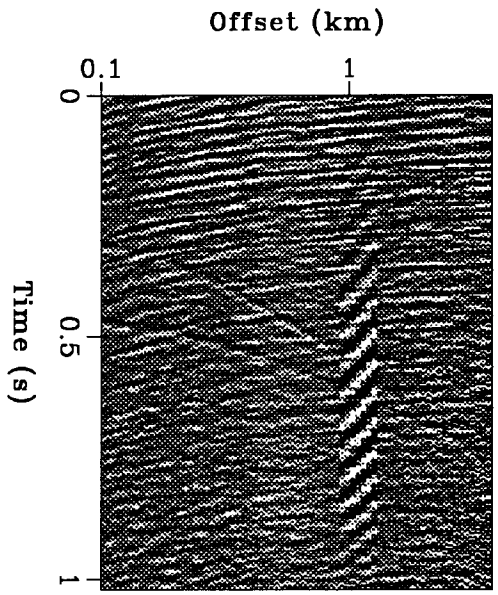

FIG. 2.2. The synthetic data are the superposition of signal (the two hyperbolas with tops at zero offset), and narrowband noise in a window of seven adjacent traces, as well as broadband noise with large spatial correlation.

2.3 Attenuation of spatially uncorrelated noise by median smoothing

2.3.1 Overview

To attenuate the spatially uncorrelated noise, I apply weighted-median smoothing along offset. I prefer median filters to linear filters in the early stages of the processing, because median filters reject noise spikes rather than smear them as linear filters do. Although median filters are non-linear, their spectral properties can be controlled to some extent through the choice of weighting functions. I smooth the data in the frequency-offset domain, where the spatially uncorrelated noise appears as spikes. Other datasets may contain noise bursts in the time domain because of factors such as recording errors. For these bursts, median smoothing should be applied first in the time-offset domain.

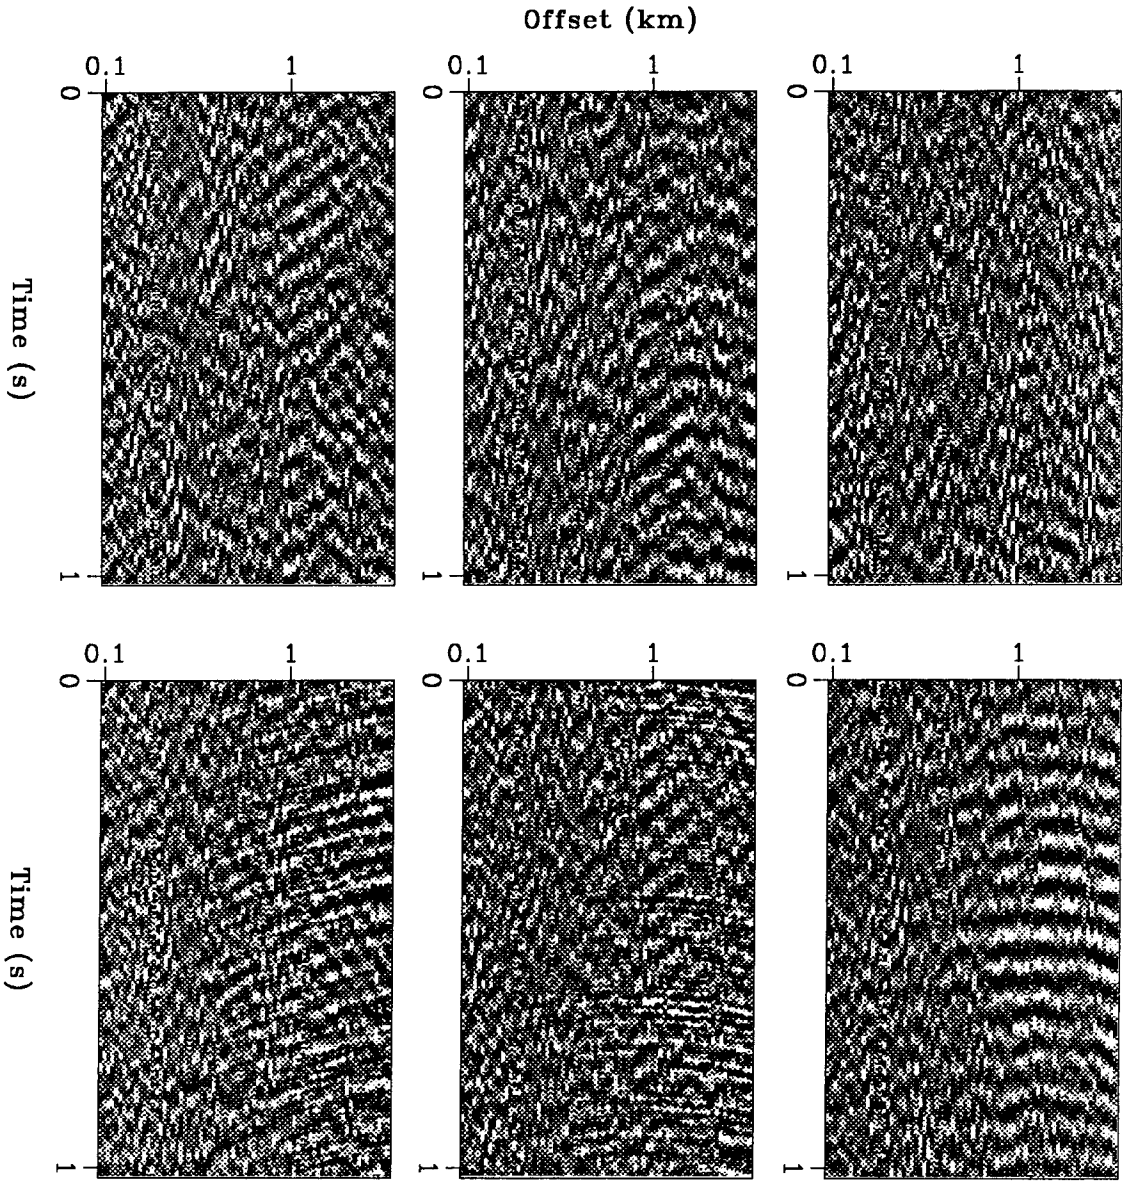

FIG. 2.3. These windows of data illustrate the varying noise patterns, especially strong at the far offsets. The only processes applied to these data are power equalization along offset, highpass filtering with a cutoff frequency of 8 Hz, and adaptive notch filtering. These windows of data are part of a 10 min data sequence, recorded within a period of 4 h. During that time the depth of the drill-bit increased by about 50 m.

2.3.2 Weighted-median filters: review of definition and properties

Definition of the median

The median of N samples $\{s_i\}_{1 \le i \le N}$ is defined by Claerbout and Muir (1973) as a solution to an L_1 minimization problem,

$$\min \sum_{i=1}^{N} |s_i - \text{median}|. \tag{2.1}$$

For sequences of real-valued samples, the median is any value between the middle two values in the sequence $\{s_i\}$: this property follows from a straightforward computation of the gradient in Equation 2.1. More specifically, the computation of the median depends on the parity of the number of samples N : for an odd number of samples, the median is unique and equal to the sample of rank $N/2+1$, while for an even number of samples the median can be any number between the samples of ranks $N/2$ and $N/2+1$. In numerical computations a convention for a "unique" median may be required; a common choice[4] is the average of the middle two values of the samples.

The median is less sensitive to outliers than is the mean; for instance, the median values of $\{1,2,3\}$ and $\{1,2,1000\}$ are identical, while their mean values are very different. Insensitivity to, or also, robustness to outliers is one reason for choosing the median as an estimator of the central value of a distribution.

Median filters

The running mean of length n is a linear lowpass filter that averages n consecutive samples from an input sequence. The output of the running median filter of length n is defined similarly as the median of n consecutive samples from an input sequence. Although the running median is a non-linear filter, it will also have a "lowpass effect" when the distribution of amplitudes of the input sequence is symmetric, because for such distributions mean and median values coincide.

To illustrate the lowpass properties of the median filter, I apply a running median filter to a sinc function (Figure 2.4), as was done originally by Claerbout and Muir (1973).

[4]For a more systematic definition of a unique median, that is valid also for data in higher-dimensional spaces, see Dellinger (1984) and Zhang et al. (1988).

The sidelobes of the sinc function have constant widths, equal to half the width of the main lobe. Thus the sinc function appears locally as a sinusoidal signal of varying frequency. A running median preserves the low-frequency main lobe while attenuating the high-frequency sidelobes.

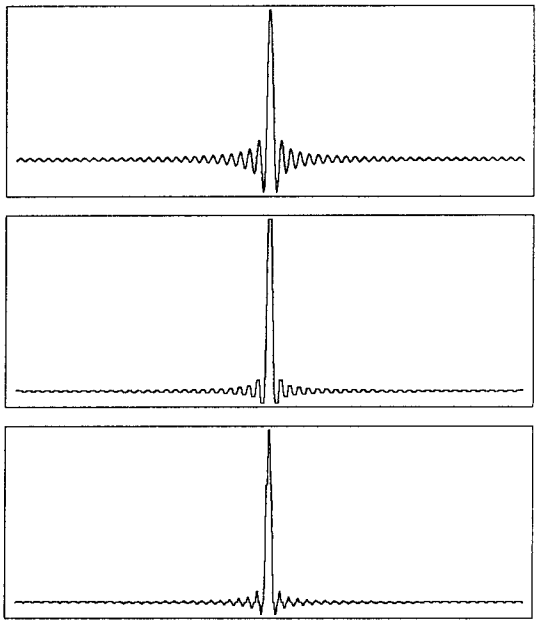

FIG. 2.4. Median filtering of a discrete sinc function:
Top: Discrete sinc function.
Middle: Output of a running median with constant weights.
Bottom: Output of a running median with triangular weights.
The running median filters (middle and bottom panels) attenuate the sidelobes of the sinc function, while preserving its main lobe. A median filter with a triangular weighting function (bottom) introduces fewer "edges" in the waveform than the uniformly weighted median (middle).

Weighted median of real numbers

The concept of a weighted median (Claerbout and Muir, 1973) extends the previous definition by weighting the terms in the objective function of Equation 2.1,

$$\min \sum_{i=1}^{N} w_i |s_i \; - \; \text{median}|, \tag{2.2}$$

with a sequence of positive weights $\{w_i\}$.

Claerbout (1985a) relates the computation of the weighted median of real numbers to the L1-norm solution of a system of linear equations, and gives a computer program based on Hoare's sorting algorithm.

Figure 2.4 illustrates also that a median filter with a triangular weighting function

outputs a smoother signal than a uniformly weighted median filter.

Weighted median of complex numbers

The definition of the weighted median (Equation 2.2) is applicable not only to real numbers, but more generally to samples in any vector space with a norm. In particular, I will need to compute in later sections the median of complex numbers. However, the cost of computing a weighted median by minimizing the objective function in Equation 2.2 using general optimization methods, would be prohibitive in applications to large volumes of data. Therefore, I approximate the weighted median of complex numbers by the complex number whose real and imaginary parts are the weighted medians of the real and imaginary parts of the input sequence.

These two definitions for a complex median will agree for symmetric distributions, where median and mean are identical. I expect that the two definitions will yield similar results when the data are modeled as a mixture of a symmetric distribution and a few outliers. Indeed, outliers in a complex distribution will have in general a real or an imaginary part that is an outlier in the corresponding sequence of real or imaginary parts of the data.

2.3.3 Median smoothing of drill-bit data

Moveout correction

In a medium of uniform velocity v, the traveltime delays from a source — located at depth z and at zero offset — to receivers at offsets x along the surface are given by:

$$t(x) = \frac{\sqrt{z^2 + x^2}}{v}. \tag{2.3}$$

For a horizontally layered medium, the above equation still approximates the traveltime delays, provided that the parameter v is interpreted as the RMS velocity of the medium (Dix, 1955).

I apply a moveout correction, determined using the hyperbolic approximation for the traveltimes (Equation 2.3), the known depth of the borehole, and the RMS velocities from surface seismic data (Appendix B), to align approximately the direct arrivals from the

drill bit. Thus, all data recorded for depths of the borehole between 0.8 km and 1 km, have been moveout-corrected with the same parameters (Equation 2.3):

$$z = 0.8 \text{ km} \quad \text{and} \quad v = 1.8 \text{ km/s}.$$

A more precise, depth-dependent moveout correction is unnecessary before smoothing, and might even introduce a bias in subsequent measures of changes in moveout as a function of depth.

Median smoothing

The strong, spatially uncorrelated noise gives a "speckled" appearance to the data in Figures 2.1 and 2.3. Fourier transform over the time axis separates better the spatially uncorrelated noise from the spatially correlated noise shown in Figure 2.6, where strong amplitude events with frequencies above 20 Hz are mainly spatially uncorrelated noise.

The effect of increasing filter length is shown in Figure 2.5, where the synthetic data of Figure 2.2 have been smoothed, after moveout correction that flattens the hyperbola with lower velocity. The weights of the two median filters are $(1, 2, 1)$ and $(1, 2, 3, 2, 1)$. The 3-point-long median filter passes the data without significant changes, while the 5-point-long filter has strongly attenuated dips that are high with respect to the moveout correction. The amplitudes along the hyperbolas have not been modified significantly by the smoothing.

I smoothe the data shown in Figure 2.6 by applying a a weighted median filter along offset in the frequency-offset domain, after moveout correction. The weighting function is a 5-point-long triangle, and the result of the smoothing is displayed in Figure 2.7. The improved continuity along offset is clearly apparent in each of the three domains: time-offset, frequency-offset, and frequency-wavenumber. The same robust smoothing improves also the spatial continuity for each of the 6 control windows, shown before smoothing in Figure 2.3 and after smoothing in Figure 2.8.

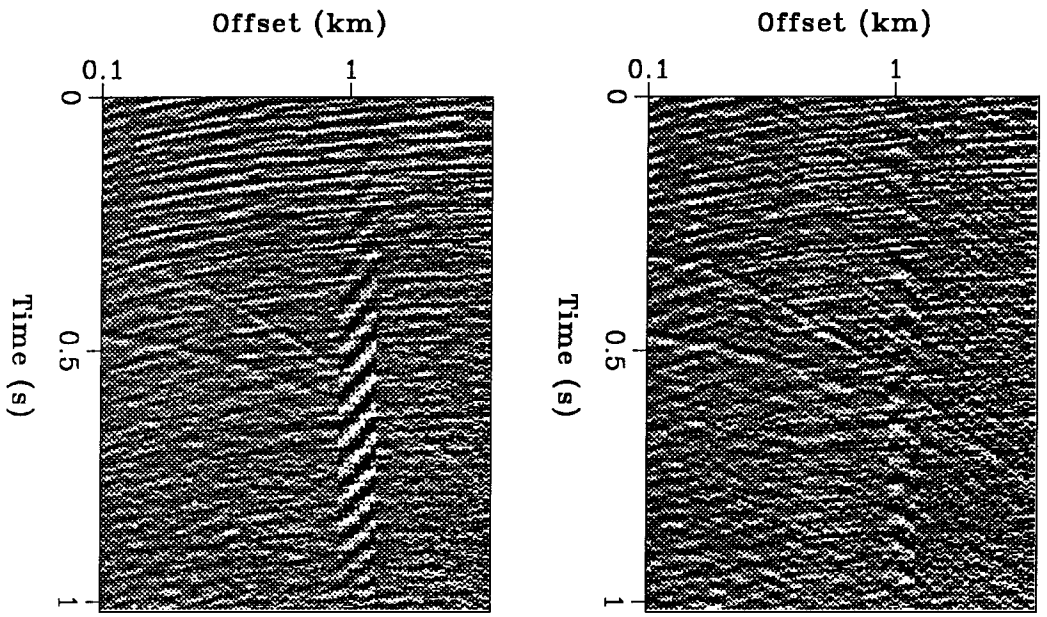

FIG. 2.5. Synthetic data after application of median filters with triangular weighting functions.
Left: 3-point-long triangular weighting;
Right: 5-point-long triangular weighting.
The 3-point-long median filter passes the data without significant changes, while the 5-point-long filter has strongly attenuated dips, that are high with respect to the moveout correction. The amplitudes along the hyperbolas have not been significantly modified by the smoothing.

Other approaches to robust smoothing

The weighted median smoothing is a procedure that is simple to implement and that helped to attenuate the strong, spatially uncorrelated noise in the OGS data. The drawback of median filters is that they tend to distort the waveforms — an undesirable effect that might outweight their advantages, especially on data with high signal-to-noise ratios.

Alternative approaches to the weighted median smoothing would be the two-step process of detection and editing of noisy traces (Mavko, 1988; Anderson and McMechan, 1989), or smoothing with an alpha-trimmed filter (Bednar, 1983).

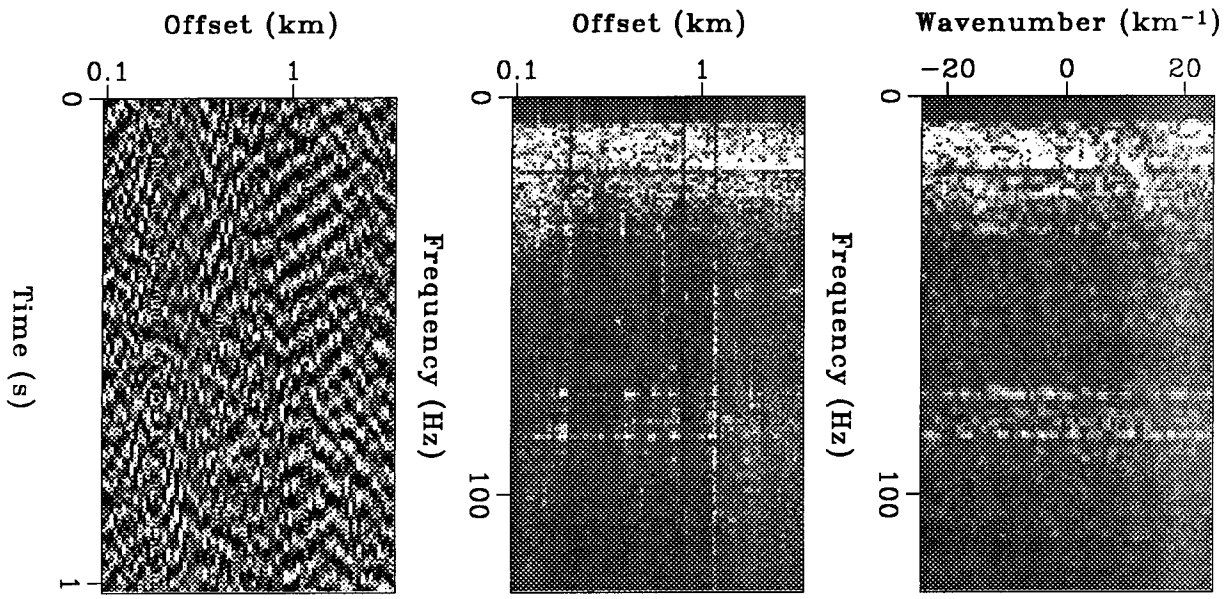

FIG. 2.6. Window of data displayed in three domains before median smoothing. Large positive amplitudes are plotted in white.
Left: Time-offset domain.
Middle: Frequency-offset domain.
Right: Frequency-wavenumber domain.

2.4 Attenuation of coherent noise

2.4.1 Overview

The goal in this section is to attenuate time-coherent or space-coherent events whose moveout differs significantly from the moveout expected from the drill-bit signal. To avoid suppressing drill-bit signal along with the noise, I start with a crude separation of signal and noise: I apply a velocity filter that removes events — signal or noise — whose moveout is close to the expected moveout of the drill-bit signal. The parameters of the velocity filter are such that no drill-bit signal remains in the passband of the velocity filter.

After velocity filtering, I observe two types of coherent noise with different correlation lengths along offset: noise that is narrowband in time and present on only a few adjacent traces, or noise that affects a large number of traces. In both cases, Iestimate the noise by linear prediction from the velocity-filtered data — along the time axis for the first type of noise, and along the offset axis for the second. Finally, to subtract the estimated coherent

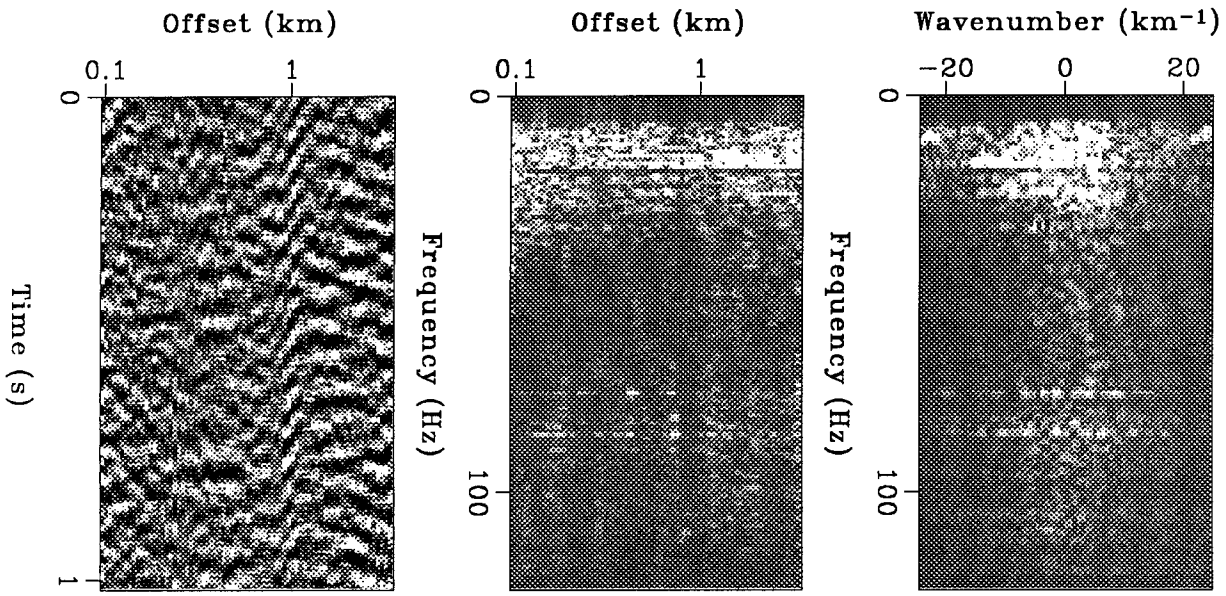

FIG. 2.7. Window of median-smoothed data displayed in three different domains. Large positive amplitudes are plotted in white.
Left: Time-offset domain.
Middle: Frequency-offset domain.
Right: Frequency-wavenumber domain.

noise from the original (non-velocity-filtered) data I solve a small regression problem that compensates for amplitude differences between the estimated noise and the noise in the original data.

2.4.2 Velocity filtering

I apply a velocity filter that rejects drill-bit signal and noise, while passing only noise. After a hyperbolic traveltime correction that aligns the direct arrivals from the drill-bit source (as discussed on page 17), I reject events with dips between -0.2 s/km and 0.1 s/km. The stopband is chosen so that drill-bit energy reflected at depths greater than the depth of the well, or scattered at impedance contrasts along the borehole (Samec and Kostov, 1988) would also be rejected by the dip filter. Reflected drill-bit energy will have higher apparent velocity than the direct arrivals from the drill bit, and the corresponding dips will be negative after moveout correction. Conversely, scattering at shallower depths than the depth of the drill bit will generate events with lower apparent velocity and positive

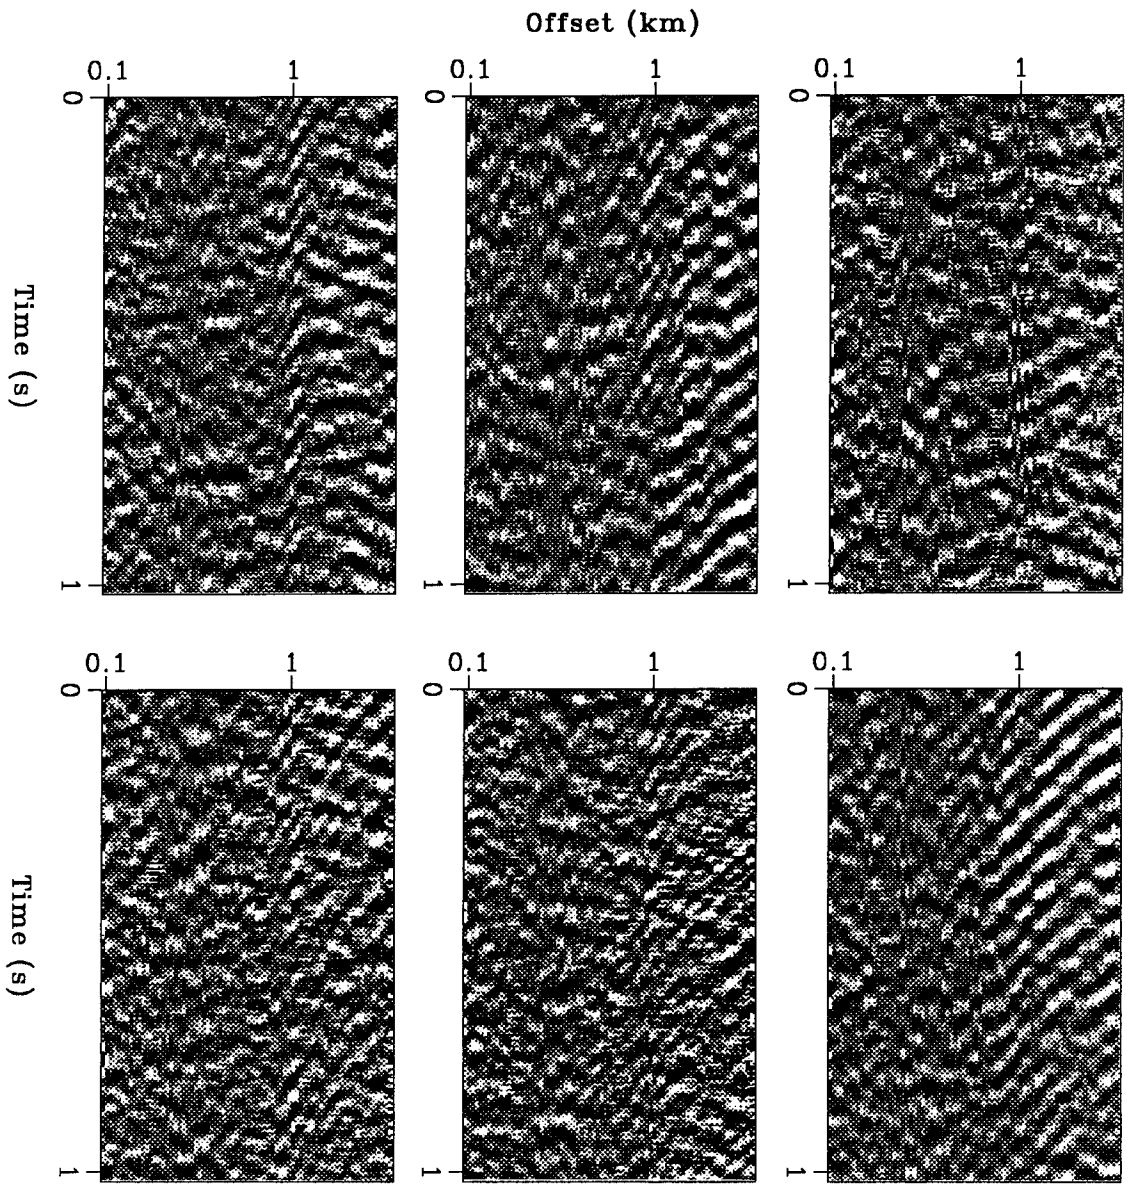

FIG. 2.8. Results of median smoothing for the 6 control windows shown in Figure 2.3. The data are displayed in the time-offset domain, after moveout correction that should align direct arrivals from the drill bit along the offset axis. The strongest events are still coherent noise at far offsets. Such noise has high apparent velocity, and is generated most likely by sources off the plane of the array.

dips.

The implementation of the dip filter consists of three steps: transformation to the frequency-dip domain, setting to zero the energy in a range of dips, and inverse transforming to the time-offset domain. The transforms between the time-offset and frequency-dip domains are least-squares inverses of each other, as described in Appendix C. Figure 2.9 displays the data after the velocity filtering has removed events within the range of move-outs expected for the drill-bit signal. The events in the passband of the velocity filter (that is, the preliminary estimate of signal) are shown in Figure 2.10.

2.4.3 Attenuation of coherent noise by linear prediction

Linear prediction along the time axis

The strong narrowband noise, present in several control windows at far offsets (Figure 2.9) affects a few adjacent traces and introduces sharp amplitude contrasts along offset. Because such noise is continuous along offset, it is not suppressed well by median filtering. On the other hand, linear prediction along offset would smear the noise along offset. Therefore I chose to suppress such noise by gapped linear prediction along the time axis (Peacock and Treitel, 1969; Claerbout, 1989). In the gapped prediction-error model, the sample at time $(t + GAP)$ is predicted from L past samples at times $\{t, t-1, \ldots, t-L+1\}$:

$$d_{t+GAP} \;=\; a_0 d_t \;+\; a_1 d_{t-1} \;+ \ldots + a_{L-1} d_{t-L+1} \;+\; n_t, \tag{2.4}$$

where n_t are the residuals or also prediction-errors, and $\{a_0, a_1, \ldots, a_{L-1}\}$ are the coefficients of a prediction-error filter of length L, determined so that the power of the prediction-error is minimum.

The normal equations for the gapped prediction model (Equation 2.4) establish the following relationship for the samples of the autocorrelation of the data $\rho_d(t)$:

$$\rho_d(GAP+\tau) \;=\; a_0 \rho_d(\tau) + \ldots \;+\; a_{L-1} \rho_d(\tau - L + 1), \tag{2.5}$$

where τ ranges from 0 to $L-1$.

Narrowband signal components have an autocorrelation that oscillates up to large lags,

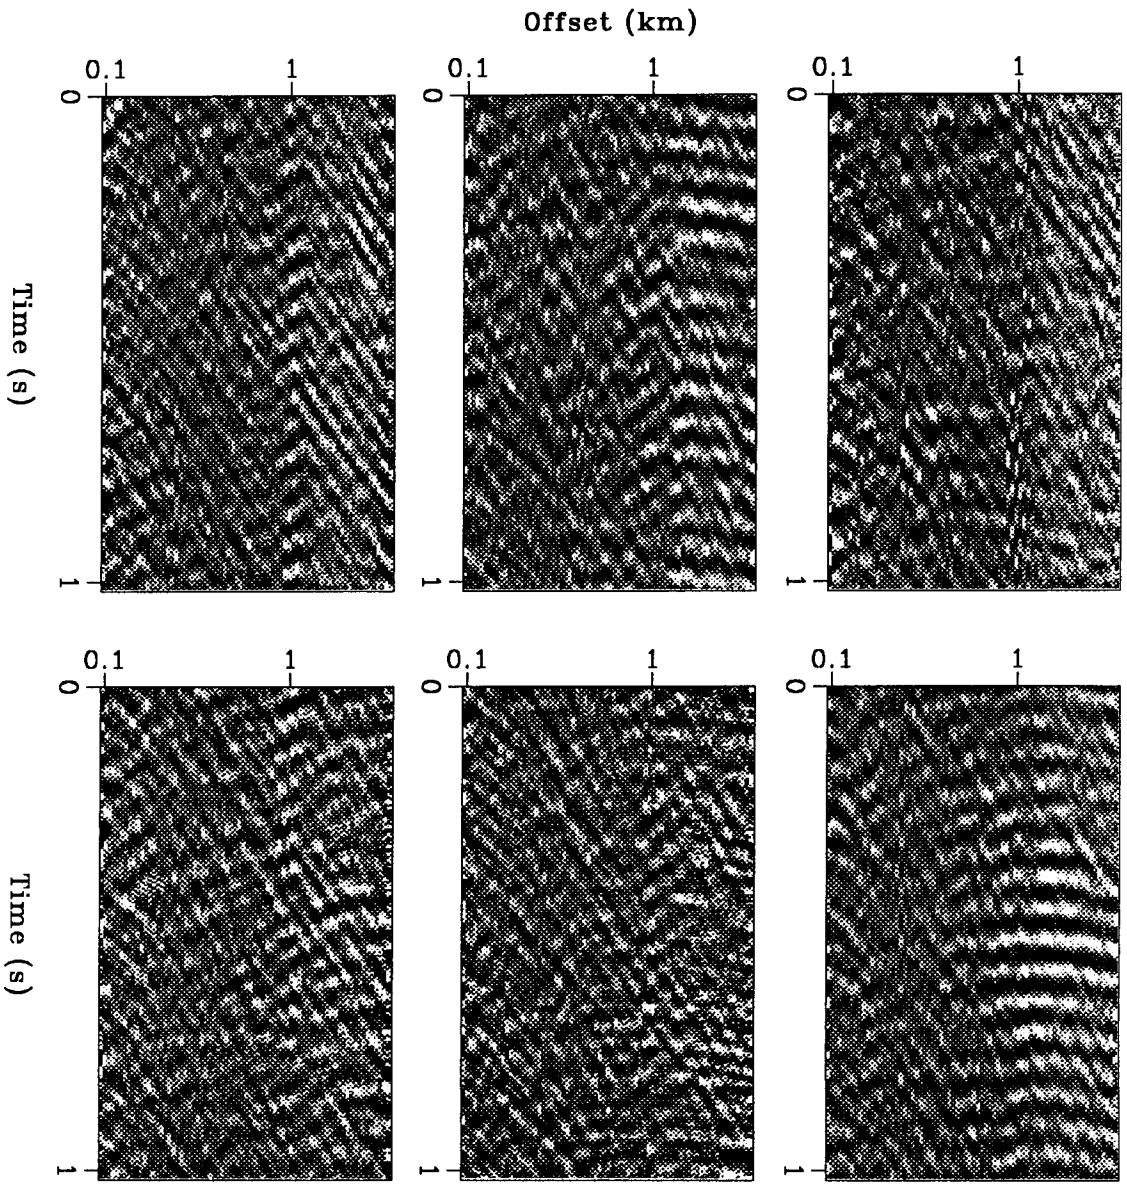

FIG. 2.9. Residuals of the velocity filtering, free from coherent events with moveout close to the expected moveout from the drill-bit signal. The moveout correction applied to the data before median smoothing (Figure 2.8) has been undone; therefore the moveout in these data is comparable to the moveout in the original data (Figure 2.3).

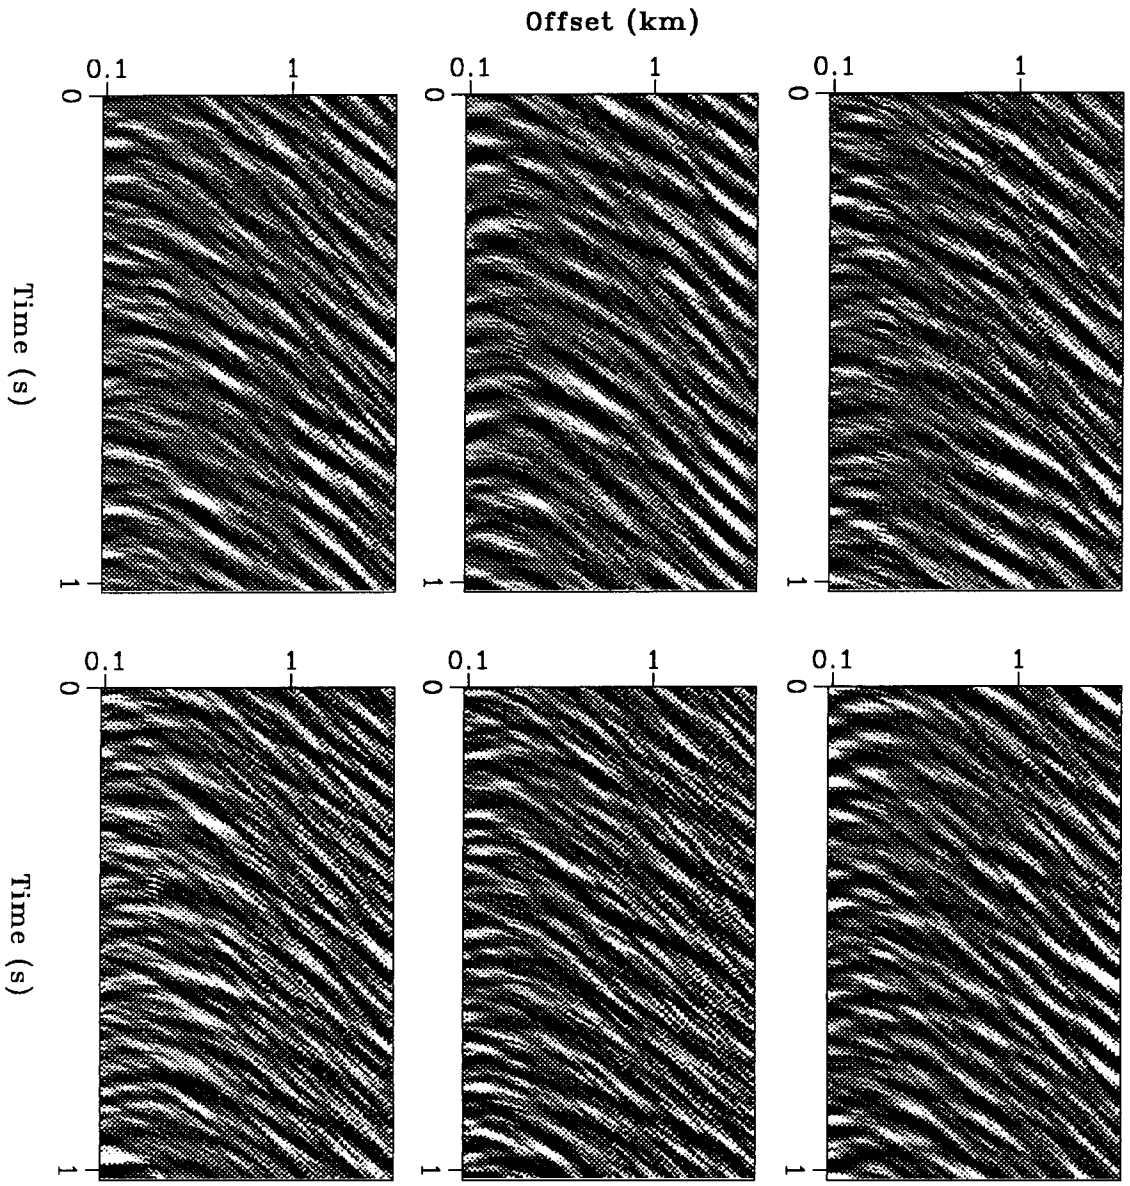

FIG. 2.10. Coherent events with moveout close to the expected moveout from the drill-bit signal have been separated from the rest of the data by velocity filtering. The moveout correction applied to the data before median smoothing (Figure 2.8) has been undone; therefore the moveout in these data is comparable to the moveout in the original data (Figure 2.3).

and are predictable for large values of GAP. On the other hand, when the autocorrelation of the data becomes zero for lags greater than GAP, the prediction-error filter becomes zero also. Thus, the gapped-prediction filter has a "detection" feature; it acts as an adaptive notch filter on some traces, while it leaves others unchanged.

In the current processing sequence, I computed independent prediction-error filters for each trace in windows of data 8 s long; the length of the prediction-error filter was chosen as 0.05 s, and the gap was 0.08 s. The choice of the gap was such that only narrowband noise was predictable. The estimated noise is displayed in Figure 2.11; it appears indeed narrowband, and affects groups of a few traces along offset.

Linear prediction along the offset axis

The steeply dipping events with linear moveout, readily apparent in the velocity-filtered data (Figure 2.9), have not been estimated by the previous single-channel gapped prediction (Figure 2.11). To estimate and suppress such spatially coherent events, I will follow a method introduced by Canales (1984) and known as f-x smoothing.

In this method, the data are modeled as the superposition of plane waves and spatially uncorrelated noise:

$$d(x,t) \;=\; \sum_i f_i(s_i x \;-\; t) \;+\; n(x,t),$$

where x denotes offset, t denotes time, and, $f_i(t)$ and s_i are the waveform and the apparent slowness across the array of the i^{th} plane wave. After a Fourier transform along the time axis, each plane wave becomes a complex exponential, and the data are represented as the sum of complex exponentials and noise,

$$d(x,\omega) \;=\; \sum_i A_i e^{j\,\omega\, s_i\, x} \;+\; n(x,\omega). \tag{2.6}$$

A particular property holds for the above model when the noise term is zero and the data are regularly sampled along offset; then there exists a linear filter such that any sample is exactly predictable from past samples of the sequence. In other words, knowledge of the data, or of the frequencies of the sinusoids in Equation 2.6, or of the coefficients of the linear filter, are equivalent (Burg, 1975; Marple, 1987). The f-x smoothing method estimates the predictable part of the data (sum of sinusoids) by applying Burg's algorithm

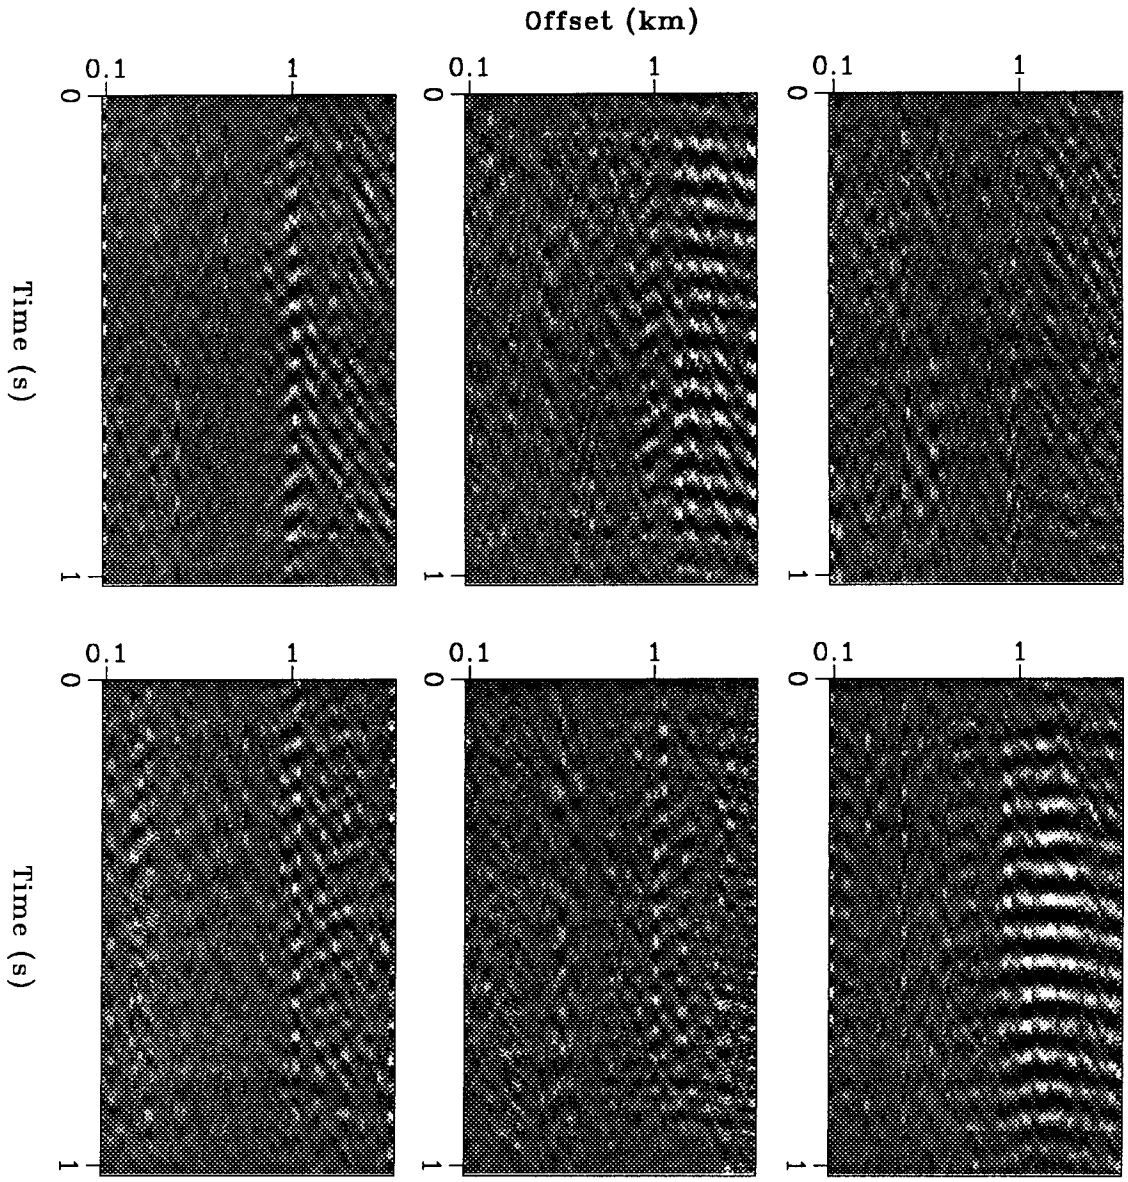

FIG. 2.11. The noise displayed in this figure is strong, narrowband, and present on groups of adjacent traces. This noise is the part of the data in Figure 2.9, that is predictable by gapped linear prediction along each trace.

for unit-lag prediction.

I apply the f-x smoothing method to the residuals of the single-channel gapped prediction (Figure 2.11), and also estimate the predictable part of the signal with Burg's algorithm. I choose Burg's algorithm over other unit-lag prediction algorithms because this algorithm makes no assumptions about the values of missing samples and is familiar in surface seismic exploration (Claerbout, 1985a). Figure 2.12 displays the results of linear prediction along the offset axis; steeply dipping coherent noise is determined by this prediction model in several control windows.

Attenuation of coherent noise

Subtraction of the estimated coherent noise (Figures 2.11 and 2.12) from the original data (Figure 2.3) shows that the amplitudes of the noise have been under-estimated. This under-estimation of amplitudes is a well known property of least-squares models; an additional factor in the current processing method could be also the use of velocity filtering to suppress the noise.

To compensate for the under-estimation of amplitudes and better subtract the coherent noise from the data, I solve a small regression problem to determine the scalar gain factors a and b that minimize the power of the residuals in the following model:

$$\text{data}(t,x) \;-\; a \times \text{nt}(t,x) \;-\; b \times \text{nx}(t,x) \;\approx\; 0, \tag{2.7}$$

where t denotes time, x denotes offset, data refers to the original data (Figure 2.3), nt refers to the estimated narrowband noise (Figure 2.11), and nx refers to the spatially correlated noise (Figure 2.12). Because the coherent noise components nt and nx are estimated from data which contain no drill-bit signal (data in the stop-band of the velocity filter), the drill-bit signal in the data will be temporally uncorrelated with these noise components, and therefore it will not be suppressed by the regression.

The results of the signal and noise separation are displayed in Figure 2.13, where several hyperbolic events are now clearly apparent, while the noise patterns (Figures 2.11 and 2.12) have been suppressed. The final results (Figure 2.13) show that the spreading of random and coherent noise has been reduced with respect to the preliminary estimate of drill-bit signal obtained by velocity filtering (Figure 2.10); as a consequence the visual

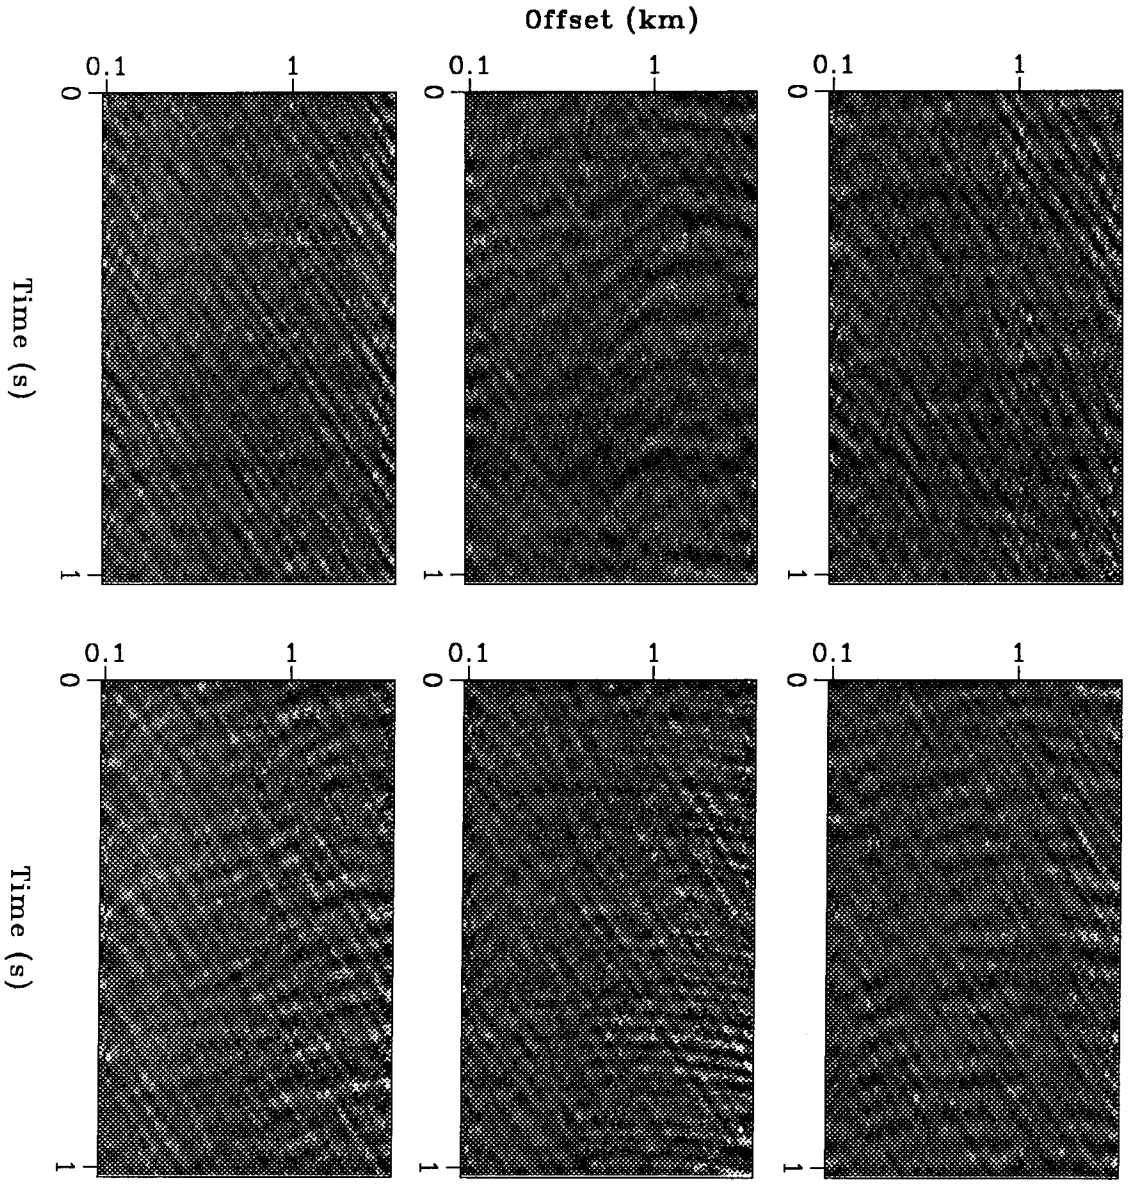

FIG. 2.12. The noise displayed in this figure consists mainly of steeply dipping events. This noise is the part of the residuals from gapped linear interpolation (data in Figure 2.9 minus data in Figure 2.11) that is predictable by unit-lag linear prediction along along offset.

detection of coherent hyperbolic events is improved.

2.4.4 Signal-to-noise ratios

Synthetic data

To understand better why the signal and noise separation method improves the results of velocity filtering, and when such an effect could be expected, I apply the processing sequence described above to the synthetic data shown first in Figure 2.2.

On the synthetic data, I define the signal-to-noise ratio as the ratio between the power of the signal (the two hyperbolas) and the power of the noise[5] (sum of broadband and narrowband components). In both cases, the power is computed in a 1 s long window of data. The results of processing data with a signal-to-noise ratio of -20 db are shown in Figure 2.14. This computation confirms that the signal and noise separation algorithm improves the results of the velocity filtering, mainly by reducing the smearing of random and coherent noise.

The results of the processing for signal-to-noise ratios of -10 db, -20 db and -30 db are shown in Figure 2.15; at -30 db the signal hyperbolas are barely visible. The improvement over the results of preliminary velocity filtering (not shown here) are most important around -20 db; for high signal-to-noise ratios the preliminary velocity filtering is sufficient, while for low signal-to-noise ratio neither method enhances the hyperbolas well.

Signal-to-noise ratios from drill-bit data

On field data, the processing does not separate the data into drill-bit and noise; therefore the signal-to-noise ratios are better described as ratios between the power of the input and the power of the output of the processing. The processing sequence has decreased the power in each of the 6 control windows (Figure 2.13) by about 10 db with respect to the raw data (Figure 2.3). Both smoothing and coherent noise filtering decrease the power by about 5 db each.

[5] $\mathrm{SNR(db)} = 10 \log_{10}(P_{\mathrm{Signal}}/P_{\mathrm{Noise}})$.

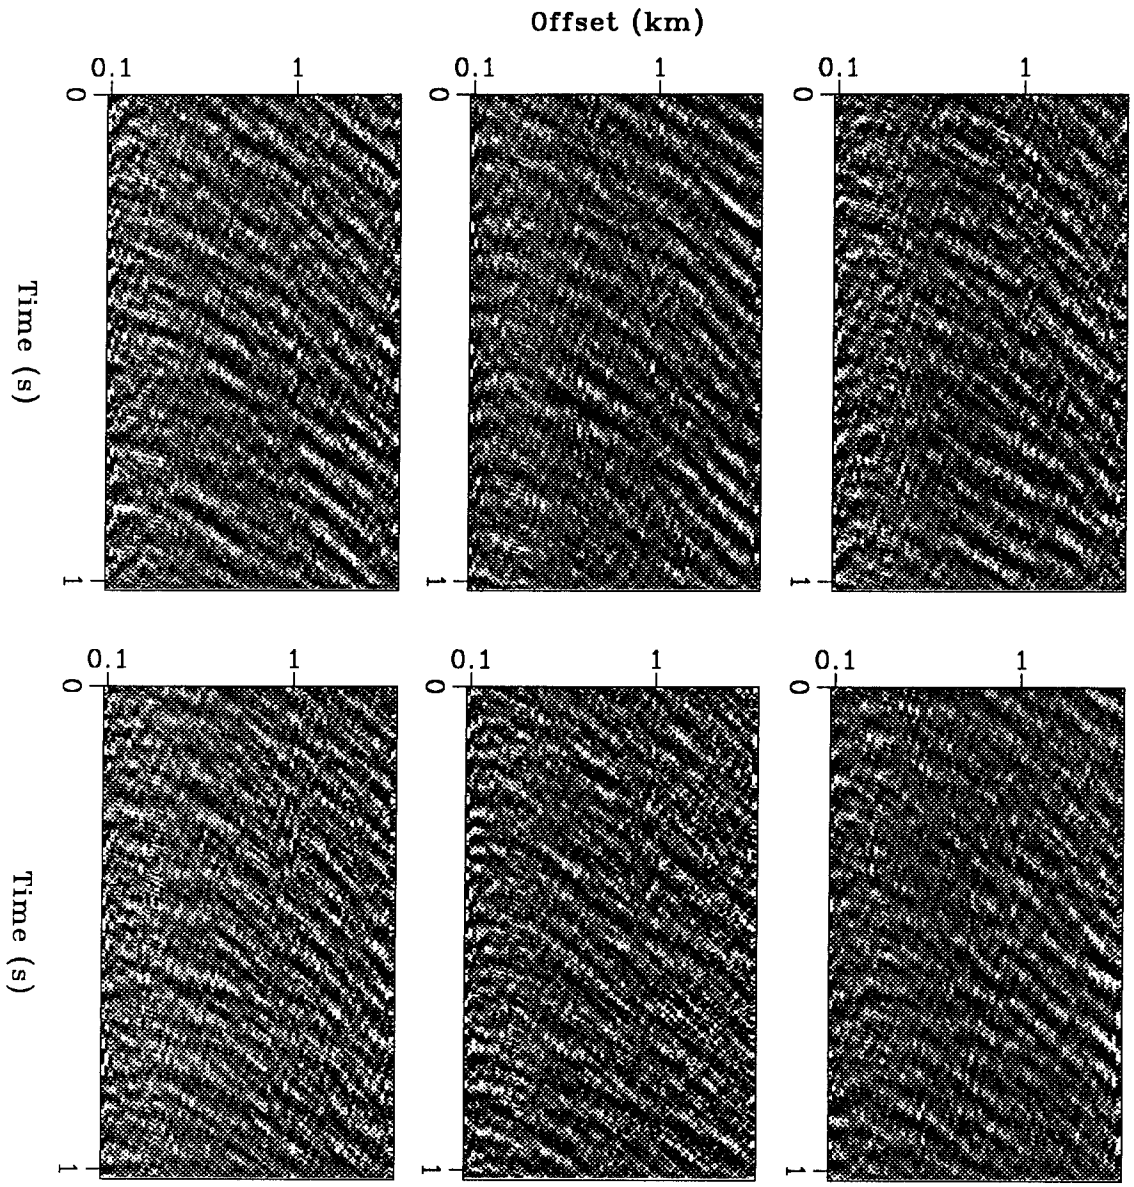

FIG. 2.13. Results of the processing; several hyperbolic events are now apparent above the background noise level.
These figures are obtained by subtraction of a linear combination of the estimated noise (Figures 2.11 and 2.12) from the median smoothed data (Figure 2.8).

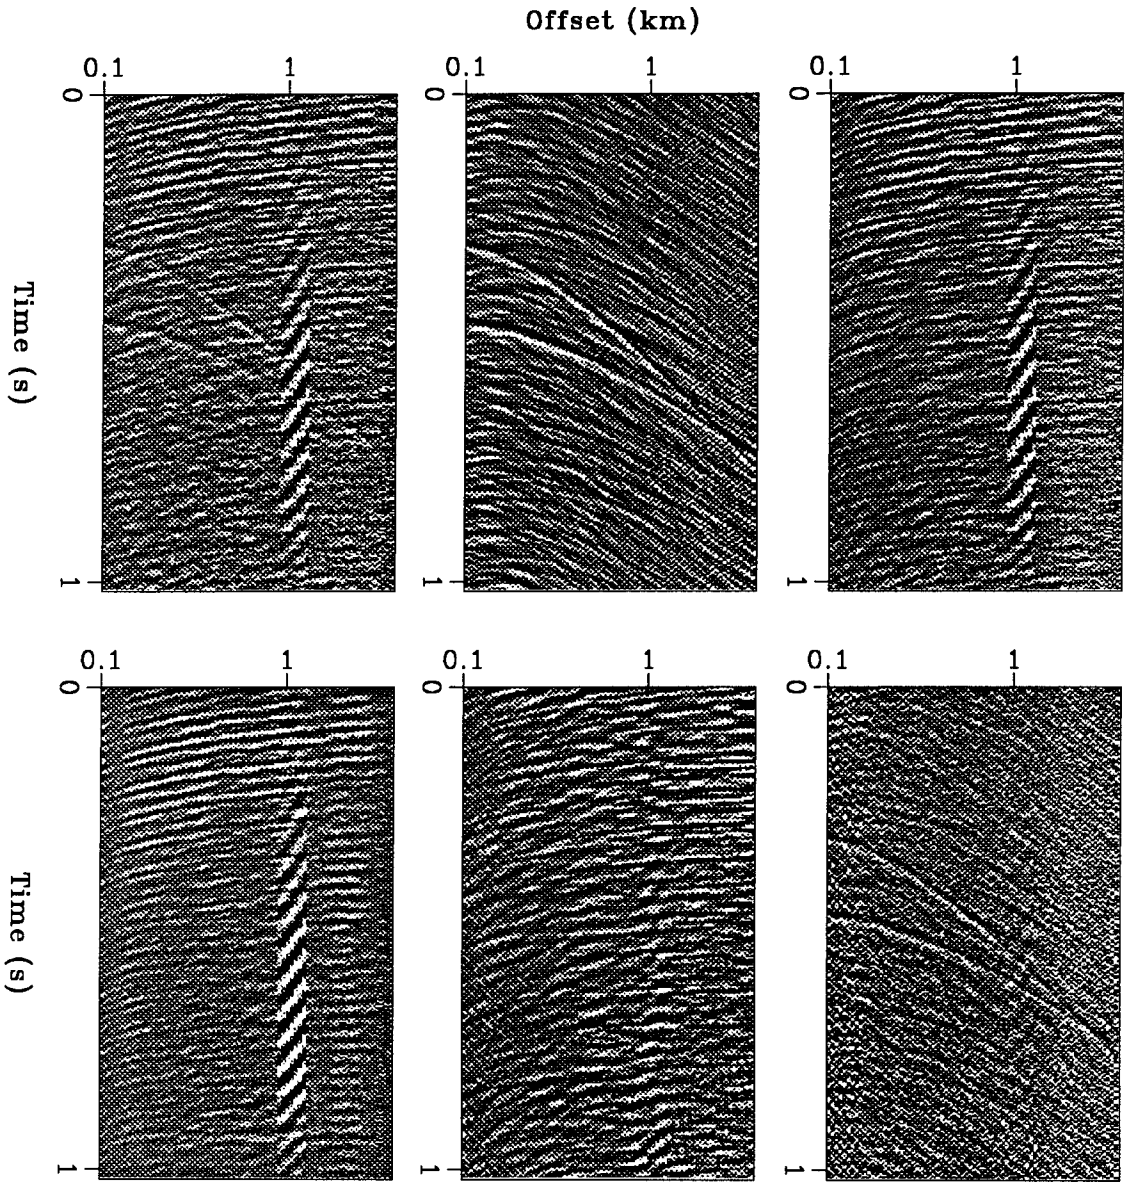

FIG. 2.14. Results of signal and noise separation applied to synthetic data:
Top row: (from left to right) Input data, preliminary estimate of signal (passband of the velocity filter), preliminary estimate of noise (stopband of velocity filter).
Bottom row: (from left to right) Estimate of coherent noise by prediction along the time axis, estimate of the coherent noise by prediction along the offset axis, final estimate of the signal and random noise.

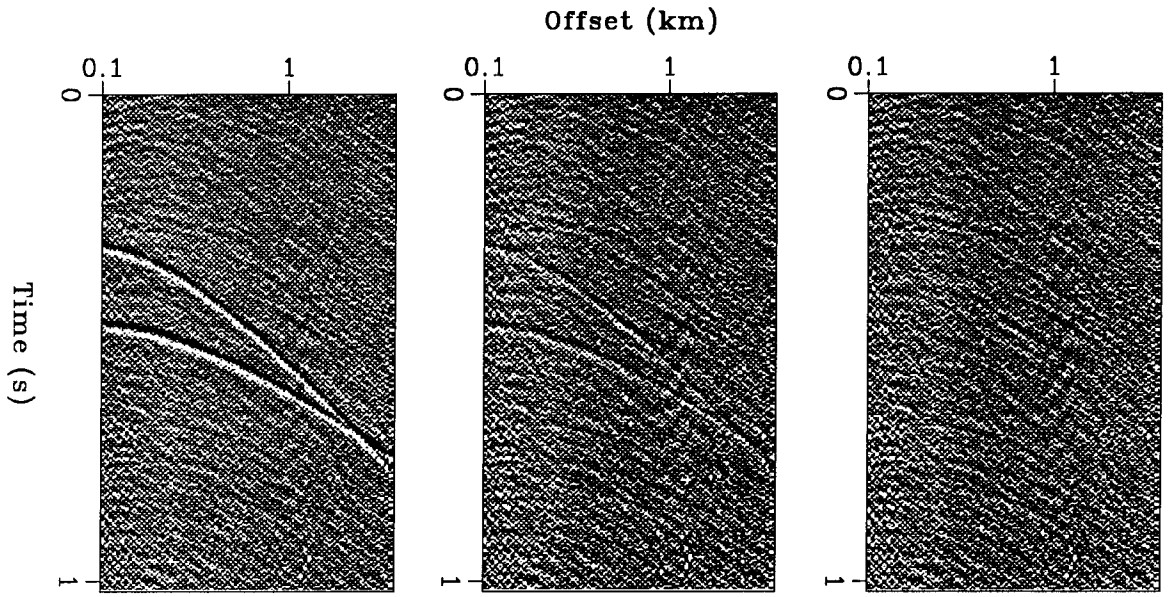

FIG. 2.15. Results of signal and noise separation for different signal-to-noise ratios (from left to right) of -10 db, -20 db and -30 db.

2.5 Tests for the attenuation of coherent noise

In this section, I discuss three tests of the signal and noise separation method presented in this chapter. The first test establishes a limit of validity of the processing: no hyperbolas become apparent when the input data consist of spatially uncorrelated noise. The second test compares amplitude spectra before and after the processing. The comparison illustrates some of the steps in the processing and provides a reference for the interpretation of average velocity spectra in the next chapter. The third test computes dip spectra as a function of offset, in order to check for residual coherent noise in the data.

Processing of spatially uncorrelated data

In the first test, I apply the processing sequence for the separation of signal and noise to spatially uncorrelated data; these data are obtained from the first three control panels (top row in Figure 2.8) by random permutation of the traces along offset. Figure 2.16 shows the results of the processing; no hyperbolas become apparent when the input data are random.

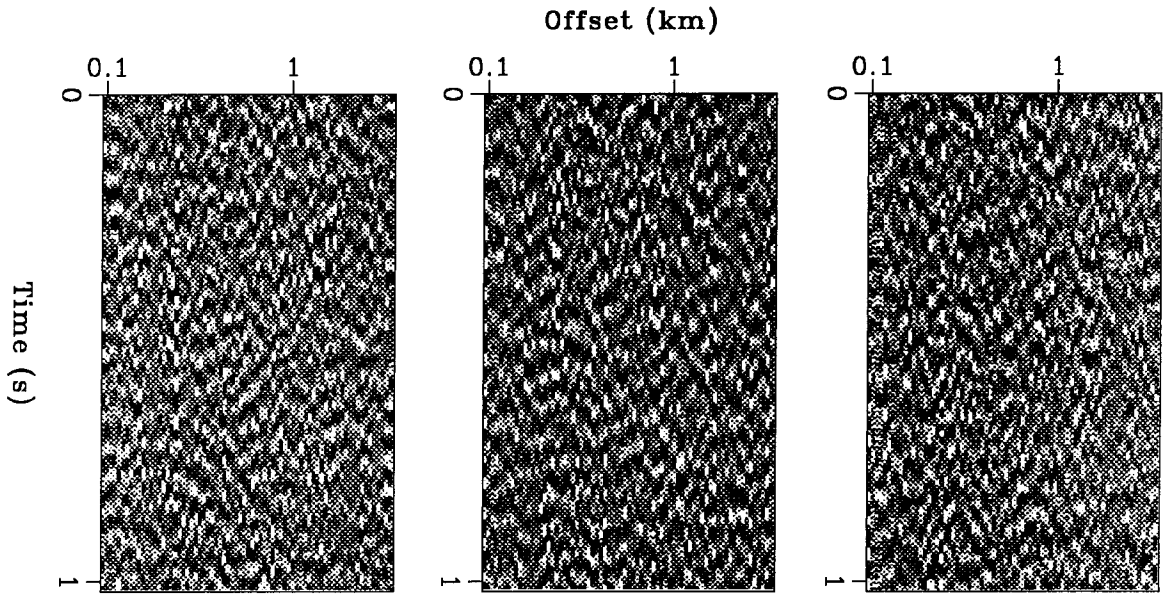

FIG. 2.16. No spatially coherent signal is apparent in the above panels, which represent the signal estimated from spatially uncorrelated data. The input data are obtained by random permutation of the traces in three of the control windows (top row of Figure 2.8).

Average amplitude spectra

The second test compares average amplitude spectra before processing for noise attenuation. I consider the data as a cube with coordinate axes offset, frequency (frequency component in windows each 8 s long) and time; each of the spectra in Figure 2.17 is obtained by averaging the amplitudes of the data along one of the coordinate axes.

The leftmost column in Figure 2.17 shows the amplitude spectra averaged over frequency, as a function of time and offset. After processing the amplitudes appear uniform along time and offset; such uniform amplitudes are required before averaging velocity spectra over time.

The middle column in Figure 2.17 shows the amplitude spectra averaged over time, as a function of frequency and offset. Processing has increased the smoothness of the spectrum, eliminated a few traces with anomalous amplitude spectra, and increased somewhat the higher frequencies with respect to the low frequencies.

The rightmost column in Figure 2.17 shows the amplitude spectra averaged over offset, as a function of time and frequency. The effect of the adaptive notch filter (near 20 Hz) is

clearly apparent. The variations of several other narrowband components can be tracked over time; such narrowband components are due to noise from the engines on the drilling platform.

Local dip spectra along offset

In this third test, I compare dip spectra computed as a function of offset before and after noise attenuation. Figure 2.18 displays dip spectra computed from 10 min of data in windows each 4 s long and 10 traces wide; the zero dip reference is the expected moveout of the drill-bit signal, and the range of dips centered at zero dip is 1.4 s/km. Before processing the energy appears uniformly distributed across the range of dips, whereas after processing the energy remains as expected near zero dip. The spectra show little variation as a function of the recording time, and thus indicate a good suppression of noise over the whole volume of data.

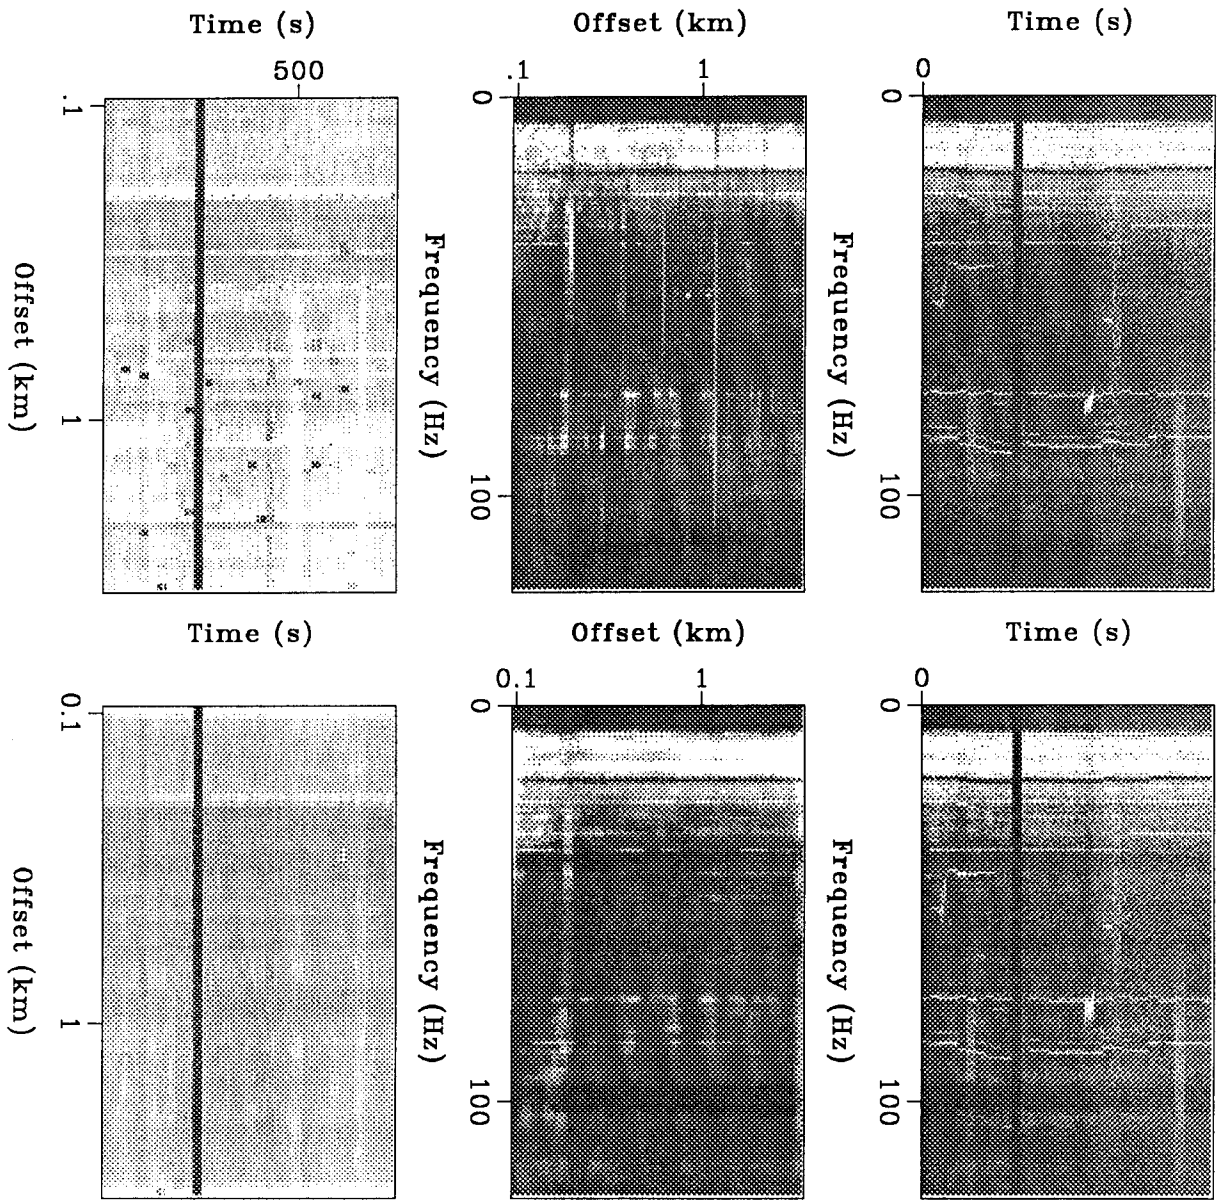

FIG. 2.17. The amplitude in the data is computed from 8 s long windows, as function of frequency, offset, and time (or equivalently depth). The upper three panels are computed before the processing is applied for signal and noise separation; the bottom three panels are computed after the processing.
From left to right the panels display the averaged amplitudes:

Left: Spectra as function of offset and time, averaged over frequency.

Middle: Spectra as function of frequency and offset, averaged over time.

Right: Spectra as function of frequency and time, averaged over offset.

Dark vertical lines indicate missing (zero) records. Large positive amplitudes are plotted in white.

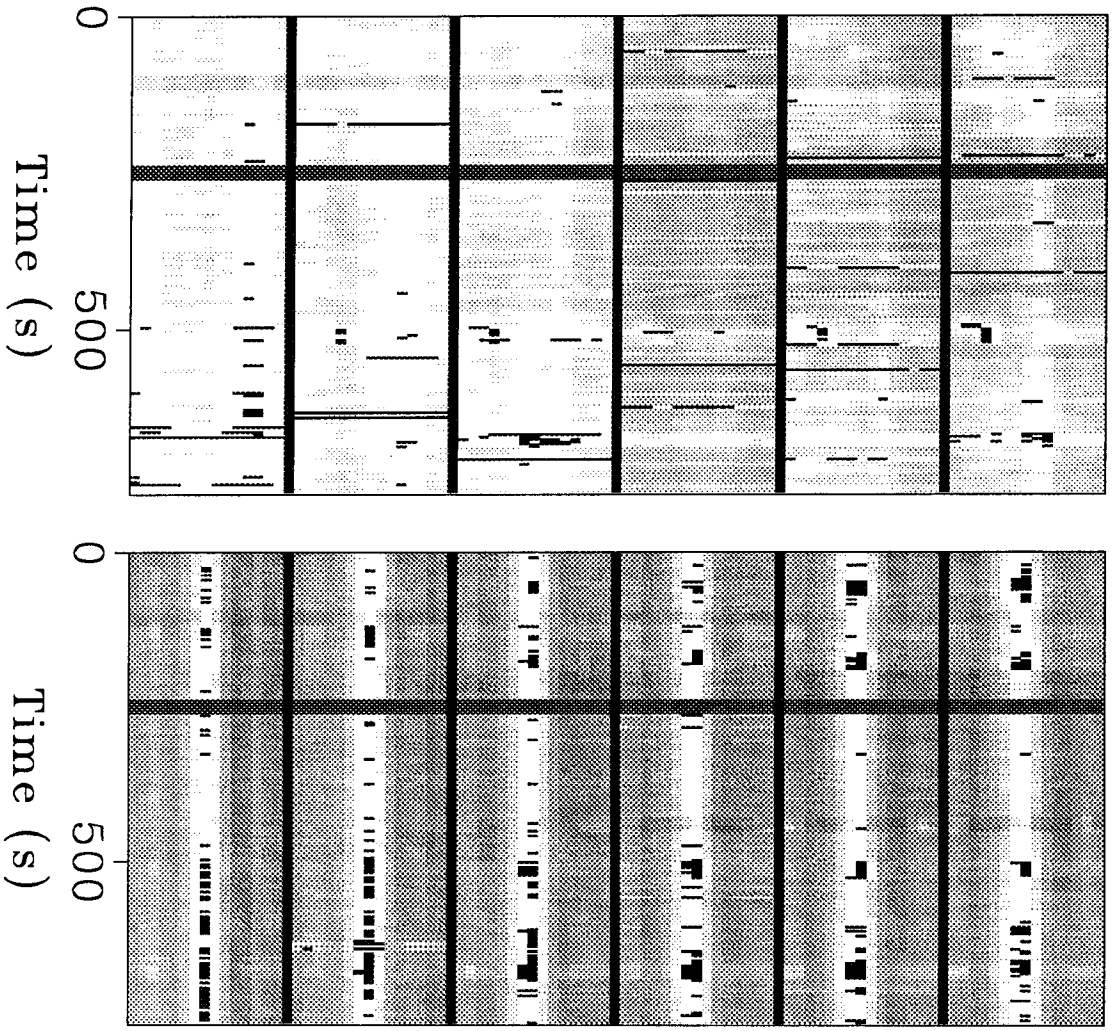

FIG. 2.18. Each of the panels displays a local dip spectrum computed in windows of 10 traces each and for dips in the interval from -0.7 to 0.7 s/km, centered around the expected moveout for the drill-bit signal (positive dips correspond to lower apparent velocity).

Top: Dip spectra from median-smoothed data.

Bottom: Dip spectra computed after coherent noise attenuation has been done.

Dark vertical lines represent a boundary between spectra; dark horizontal lines extending across all spectra correspond to missing (zero) data; black squares (within white regions) indicate amplitudes above the 98^{th} percentile, that have been clipped.