JOURNAL OF THE CANADIAN SOCIETY OF EXPLORATION GEOPHYSICISTS
VOL. 21, NO. 1 (DEC. 1985), P. 54-63

MEDIAN FILTERING: REVIEW AND A NEW F/K ANALOGUE DESIGN

ROBERT R. STEWART[1]

ABSTRACT

Median filters operate by selecting the middle value of an ascending-ordered sequence of numbers. These numbers are taken from a moving window on the data to be processed. The operation can be used in seismic data processing to reject glitches on data as well as enhance discontinuities in the data. This paper reviews the definitions and some of the properties of median filters. The median filter is very useful in Vertical Seismic Profile (VSP) data processing and automatic editing of surface seismic data. It is also used in seismic processing to enhance linear events. A new median filter design is discussed here which uses a box of numbers (2-D array) and weights to achieve some of the properties of an f/k filter while retaining median filter characteristics. Early tests indicate that this new filter rejects spikes, smears discontinuities less than standard f/k filters, and passes a range of dipping events.

INTRODUCTION

Seismic data processing has been concerned for many years with extracting or enhancing particular events on recorded seismic traces. One of the most powerful and commonly used techniques is stacking, or the summing together of adjacent traces to produce another trace which, hopefully, has undesirable effects excluded or minimized.

The output "stacked" trace or mean trace is, however, only one of a number of possible estimates of the statistical properties of a set of traces (Naess and Bruland, 1985). The reader may recall the three fundamental estimates or descriptions of a sequence of numbers — the mean, mode and median. Let us first summarize these estimates.

The mean $\bar{a}$ is the familiar average of a sequence of numbers a_i

$$\bar{a} = \frac{1}{N} \sum_{i=1}^{N} a_i$$

where N is the number of points in the sequence.

The mode is the most common number (most repeated value) in the sequence. The mode may or may not be present. There also could be several numbers that are repeated an equal number of times in the sequence. Thus the mode, if it exists, may not be unique. The median is the middle point of a sequence ordered in ascending value. If there is an even number of points in the sequence, then the median is generally taken to be the average of the two middle numbers.

Often these statistical descriptors are roughly similar. However, for some sequences these three values can be radically different.

Consider the sequence

$(10^6, -.5, -.5, 8.0, 8.5, 9.0, 7.5, 7.0, 9.1, 9.2, 7.1, 7.3, 8.2)$

Neither the mean (76929.2) nor the mode (−.5) is very representative of what might be the perceived nature of this sequence — a value somewhere near the median value of 8.0. If the negative values and the large value are considered to be errors, then the median would be a more robust estimate. That is, it is less sensitive to these errors. Thus, there is a certain motivation to investigate filtering operations based on the median sample as opposed to other statistics.

Median filters (or running median selections) do not appear to have been used a great deal in seismic processing, but have been known for some time (Tukey, 1971). In fact, the least absolute value criterion (which is shown later in this paper to select the median) was apparently introduced by Boscovitch in 1757, somewhat before the first reports on least-square methods (Denoel and Solvay, 1985). More recently, Claerbout and Muir (1973) discussed the properties of medians and absolute error minimization (some of these properties are summarized in the next section). They also proposed the use of medians in a number of geophysical problems, including first-break analysis and magnetic mapping. Nonetheless, as late as 1979, it was stated (Huang *et al.*, 1979) that while ". . . it is well-known that median filtering is useful for reducing random noise (especially when the noise amplitude probability density has large tails) and periodic patterns, theoretical

[1]Veritas Software Ltd., 615 - 3rd Avenue S.W., Calgary, Alberta T2P 0G6

Numerous discussions with Dan Hampson and Brian Russell on digital filtering have been insightful and useful. Kim Andersen assisted in testing the computer code and in data processing. Ben Dickson of BP Canada Inc. provided helpful suggestions on the manuscript.

results on its behavior are non-existent in the open literature." The median filter is currently enjoying greater theoretical attention (Lee and Kassam, 1985) and practical application. Nonetheless, there are yet relatively few papers published on median filter behaviour. Some of these papers are briefly discussed below.

Taylor (1981) wrote a thorough article on the uses of the L_1 norm in seismic processing. Again, the L_1 norm (or absolute value) minimization of a sequence of numbers selects the median sample. In particular, he considered the problems of statistical errors, deconvolution and sparse spike processing. One of the interesting aspects of L_1 deconvolution that he discussed is that it results in a sparse spike representation of the reflectivity.

Evans (1982) considered the approximate frequency domain behaviour of median filters as well as their spike removal (deglitching) characteristics. He showed that the median filters under consideration had roughly the same frequency pass behaviour as mean filters of similar length. Frequency domain analysis is somewhat qualitative, however, as the median filter is nonlinear. He found that the median filter performed more resiliently than mean filters in the face of spiky noise, and was very useful for filtering coded square-wave data (WWVB signals).

In another article concerned with median filter applications, Bednar (1983) compared the acoustic log editing performance of Markov processes, running means and median filters. He suggested that, at least for certain cases, the median filter was the preferred choice. He also considered median filtering in the cepstral domain as a way to enhance a deconvolution process. He concluded that there is some potential in this approach.

Hardage (1983) gave considerable emphasis to the use of median filters in Vertical Seismic Profiling (VSP) data processing. He showed how they could be used to considerable advantage to enhance events of interest. One method of coherent noise rejection and thus signal enhancement proceeds by first aligning the event considered to be noise. This event is enhanced via a median filter. The enhanced event is then subtracted from the aligned raw data (the pre-enhancement data). These raw aligned data with the coherent noise subtracted are then shifted back to their original position. They now have the coherent noise removed. In fact, median filtering of flat or linearly aligned events has become a standard practice in VSP processing (Hardage, 1983; DiSiena and Stewart, 1985).

Fitch *et al.* (1984) described a decomposition technique that would allow a more simplified analysis of median filters. They reduce a sequence of numbers (or traces) to binary signals (rectangular waves) via given thresholds. This is accomplished by constructing a set of rectangular waves (zero values if the raw trace amplitude is below a given threshold, and a value of one if the amplitude is above the threshold) from the input trace. These binary signals are repeatedly median filtered to arrive at a "root trace" — that is, a trace that is invariant under further median filtering. These root traces can be recombined to give the response of median filtering the original signal. They suggested that this approach may lead to faster algorithms.

M.J. Woodward of the Stanford Exploration Project has discussed median "stacked" sections and median-based velocity analyses (*pers. comm.*, 1985). R.L. Kirlin and others associated with the University of Wyoming Volcanic Reflection Research Project have also considered median filters in some detail (Kirlin *et al.*, 1985). They used dipping linear constant weights in 2-D arrays to enhance dipping events on the seismic data. Hardage (1985) discussed using "dip-steered" median selection to enhance dipping events. In this type of filter, a median is extracted for various selected dips. The median sample associated with the smallest variance for the range of dips is selected as the appropriate filtered point.

A process that displays properties of both the median and the mean is called the alpha-trimmed mean (Watt and Bednar, 1983). In this type of selection, the data points are ordered but a range of values around the median are averaged. Equivalently, this process discards outlying points but smooths (averages) the interior points. This type of filtering shows considerable promise in seismic data processing (Farmer and Haldorsen, 1985).

The median filter also has been used extensively outside geophysics in image processing (Huang *et al.*, 1979; Nodes and Gallagher, 1982). The edge enhancement (step-function pass) properties of the median filters are of fundamental importance in this application. After this brief overview, the next section discusses some of the properties of median filters in more detail.

Median Properties

As noted previously, the median is the middle point of an ascending-value sequence. For example,

$(5, 1000, 6, -1, -1, 7, 9)$

when ordered, becomes

$(-1, -1, 5, 6, 7, 9, 1000)$

The median is 6. Several aspects of this selection are notable. First, the median is an actual value of the odd-numbered sequence (as opposed to, say, the average, which is a combination of the input numbers and has a value of 146.42). Second, the inclusion of a very aberrant number, 1000, disturbs the median selection very little. If 1000 were dropped from the sequence, the median would change only slightly. These particular properties may or may not be desirable, depending on the application.

If a sequence is symmetric about some point (for example, an event has a symmetric noise distribution), then the median and mean selections will be the same. For example, the sequence

$(0, 5, 6, 7, 8, 9, 14)$

has a median and mean equal to 7 (there is no mode).

Another interesting property of the median is that it is the value of a sequence that is a minimal "distance" from all other points (Claerbout and Muir, 1973; Claerbout, 1976). That is, the median value minimizes the absolute value of the sum of differences between it and other points of the sequence (the l_1 norm).

In mathematical terms, this can be expressed as follows:
Given $x_1, x_2, \ldots\ldots\ldots, x_N$
Find $\hat{x}$ which minimizes

$$J = \sum_{L=1}^{N} |\hat{x} - x_i| = \sum_{i=1}^{N} (\hat{x} - x_i)\ \mathrm{sgn}\ (\hat{x} - x_i)$$

$$\text{Now} \quad \frac{dJ}{d\hat{x}} = \sum_{i=1}^{N} \mathrm{sgn}\ (\hat{x} - x_i)$$

$$\text{And setting} \quad \frac{dJ}{d\hat{x}} = 0 \ \text{then} \ \sum_{i=1}^{N} \mathrm{sgn}\ (\hat{x} - x_i) = 0$$

But this is true when there are equal numbers of positive and negative signs of $(\hat{x} - x\ )$ or an equal number of points on either side of $\hat{x}$. Thus the required $\hat{x}$ is the median. Note that if there is an even number of points, another point is conceptually added between the two middle samples. This new point is taken to be the median.

Recall that the mean is that value which minimizes the sum of the squared differences between it and the other points of the sequence (the l_2 norm).

The median selection is also invariant to scaling. That is, if all the numbers in the sequence are scaled or exponentiated (assuming positive numbers), then a given number will occupy the same ordered position before and after scaling. For example, if a positive sequence of numbers has a median x_m, and then the sequence values are all squared, the new median will be x^2_m. Note that the mean does not have this exponentiation property

(*e.g.* $\left[\frac{1}{N} \sum_{i=1}^{N} x_i\right]^2 \neq \left[\frac{1}{N} \sum_{i=1}^{N} x^2_i\right]$)

Median Filtering

The median filter M is a data-dependent operator which extracts a value from a sequence of numbers.

$M\ (5, 1000, 6, -1, -1, 7, 9) = 6$

It is a nonlinear operator, as it does not conform to the definition of linearity:

$$F\ (a\vec{x} + \vec{b}) = aF\ (\vec{x}) + F\ (\vec{b})\ (1) \qquad (1)$$

For example, take

$\vec{x} = (1,2,3) \quad b = (2,4,3) \quad a = 2$

M = Median filter

$M(\vec{x}) = 2 \quad M(\vec{b}) = 3$

$aM(\vec{x}) + M(\vec{b}) = 7$

$a\vec{x} + \vec{b} = (2,4,6) + (2,4,3)$

$= (4,8,9)$

$M\ (a\vec{x} + \vec{b}) = 8$

$8 \# 7$

Thus the median filter is nonlinear. Recall that the l_2 norm selection (mean) is linear. If F were a mean filter in the above example, then equation (1) would hold (with $7 = 7$).

Most filters move (run) along or across a sequence of numbers. In seismic processing these numbers represent seismograms. The running median filter can be applied along a given trace, or across adjacent traces. In general, though, a window of values will be used, a median extracted and placed on the output trace at the middle of the window, the window shifted one data point along the sequence, a median extracted, window shift, etc. This filtering process is shown for a simple case using 5-point filter in Figure 1.

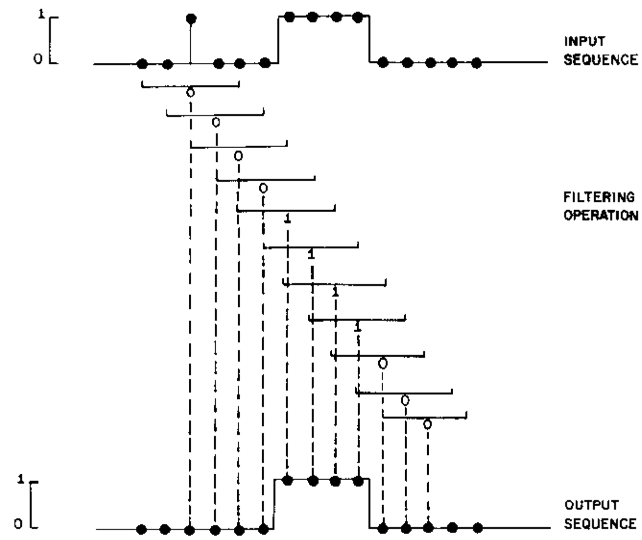

Fig. 1. Schematic diagram of the median filtering operation. A 5-point filter is used on the input sequence. Note the despiking and step-passing effects.

Several observations are evident. The median entirely rejected the single spike, and totally passed the boxcar values. Generally, any "spike" or perturbation on the trace that is less than N/2 points long, where N is the filter length, will be rejected. Consistent sets of values longer than this will be passed. Because of this rejection property, the median filter has a somewhat similar response to a high-cut filter. As the filter is nonlinear, the traditional tools of linear analysis (sinusoidal superposition) cannot be exactly used (Fitch *et al.*, 1984). However, consider a boxcar M points long and a sample interval of Δt: then the frequency content of the boxcar function is largely less than $1/(M\Delta t)$ or the first spectral zero-crossing. If a median filter of length N is to reject the box, then $M < N/2$. Thus, very approximately, when the box is of length such that an appreciable amount of its spectral power is above $2/(N\Delta t)$, it will be rejected by the N-length median filter. On other more complex waveforms, the reject characteristics of the filter are different. For example, a single sine wave cycle of period $N\Delta t$ has a dominant frequency of about $1/(N\Delta t)$. This sine cycle would be rejected by a median filter of length N.

Note that the edge or step on the boxcar in Figure 1 is not smoothed. The jump from zero sample to the unity sample is retained exactly. Both of the above properties are useful. especially in VSP processing.

While the median filter may be used to filter a trace in time, it is also often used as a spatial process. That is, it is often possible to align events (such as through NMO correction), then apply the filter in the spatial (z or x) direction. The filter will then attenuate dipping events. The dips that are rejected are a function of the spatial window width of the filter and the temporal length of the dipping events. For the simple case of a dipping boxcar event with temporal length T(ms) and a dip, m(ms/trace), then a filter of length (number of traces) L will be required to reject the dipping event where

$$L > 2T/m \tag{2}$$

Examples of Median Filtering

Shot-Gather Example

In this example a single, gained, trace-equalized, NMO-corrected shot gather is median-filtered. The input data (96-channel, dynamite source) are shown in Figure 2. Low-velocity events are evident on the near-offset traces. The gather also shows evidence of multiples, especially at about 1030 ms. For this example, the median filter is applied at each time (0 to 2000 ms) and selects the amplitude values from a specified number of adjacent traces. A median is extracted from this window at a specific time and assigned to the output trace at the midwindow position at that time. The window shifts over one trace and extracts the median sample again. This process is repeated until all traces at that time have been filtered. The process then moves to the next time sample and runs across the adjacent traces. This procedure is repeated for all times. Figure 3 shows the results of a 9-trace median filter applied to the shot gather displayed in Figure 2. Generally, the median-filtered data appear much less noisy. The low-velocity events in the interior traces have also been attenuated. There is, however, some very high frequency "chatter", or sample-to-sample discontinuity, evident on the section. This is caused by the median process, which has independent estimates from one time on the trace to the next. A high-cut filter (with high cut above the signal band) is generally applied to remove this sample-to-sample discontinuity (see Fig. 4). Figure 5 shows the raw data of Figure 2 median-filtered again, but this time with a 17-trace window. The results of filtering with the two different filter widths are similar, but the larger filter width has reduced the dipping events (multiples) more. This would be expected in the light of equation (2). The multiple event at 1030 ms has been attenuated at the larger offsets but the near offset traces still have multiple energy on them. At least visually, the filter has reduced noise as well as attenuated much of the dipping energy.

VSP Examples

As mentioned previously, the median filter is currently being used in a production mode for the processing of VSP data. Generally, any event that can be seen can be aligned and then enhanced by the median. If this enhanced event is desirable, it can be interpreted or further processed at this point. If the event is undesirable, it can be subtracted from the pre-filtering (input) data. In part, because the median filter honours the data (*i.e.* selects actual values), this subtraction process is ext-

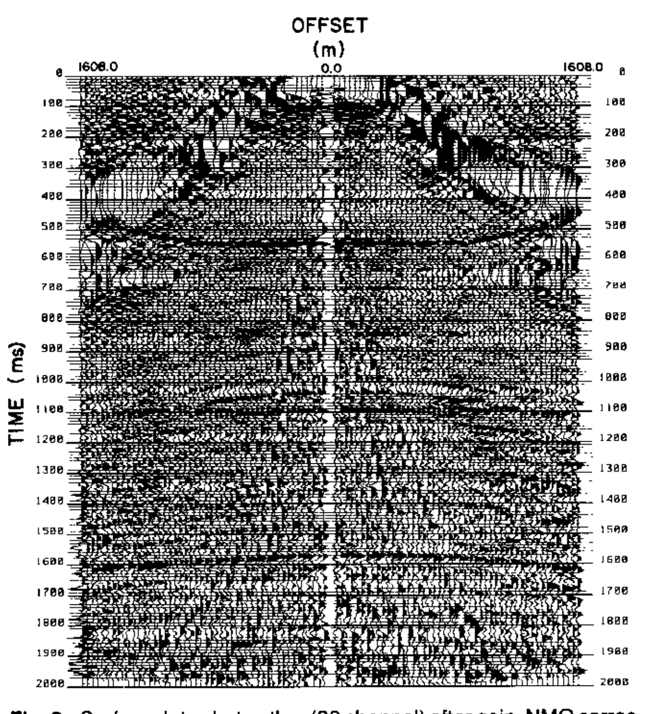

Fig. 2. Surface data shot gather (96 channel) after gain, NMO correction and trace equalization.

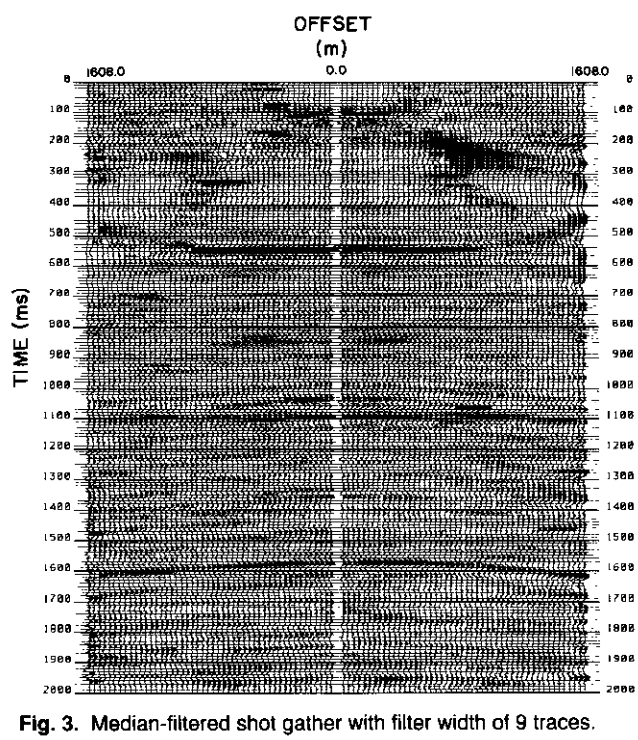

Fig. 3. Median-filtered shot gather with filter width of 9 traces.

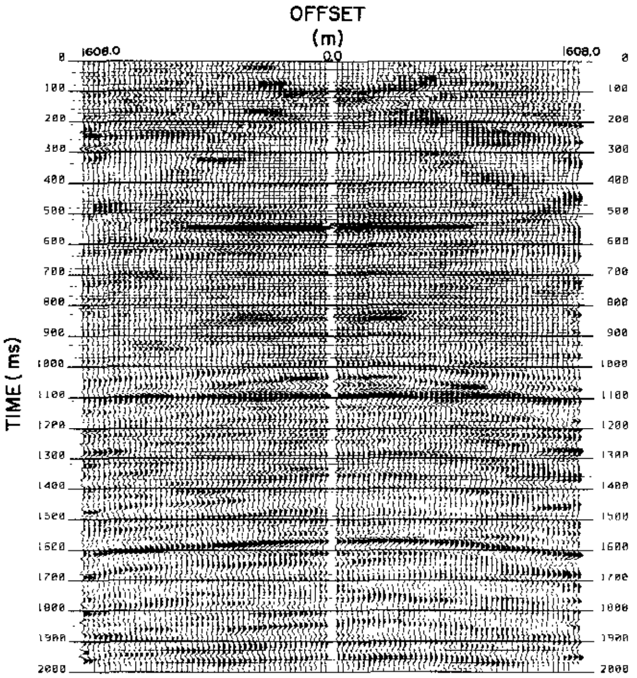

Fig. 4. Bandpass and median-filtered shot gather. The data of Figure 3 were bandpass-filtered with a bandpass (0 -10 -60 -75 Hz) filter.

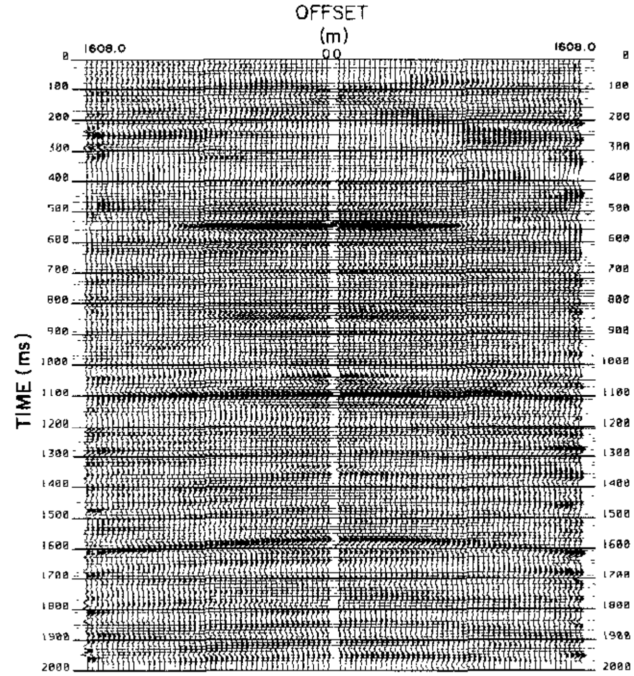

Fig. 5. Median- and bandpass-filtered shot gather. The data shown in Figure 2 were filtered with a 17-trace median filter, then bandpass (0 -10 -60 -75 Hz) filtered.

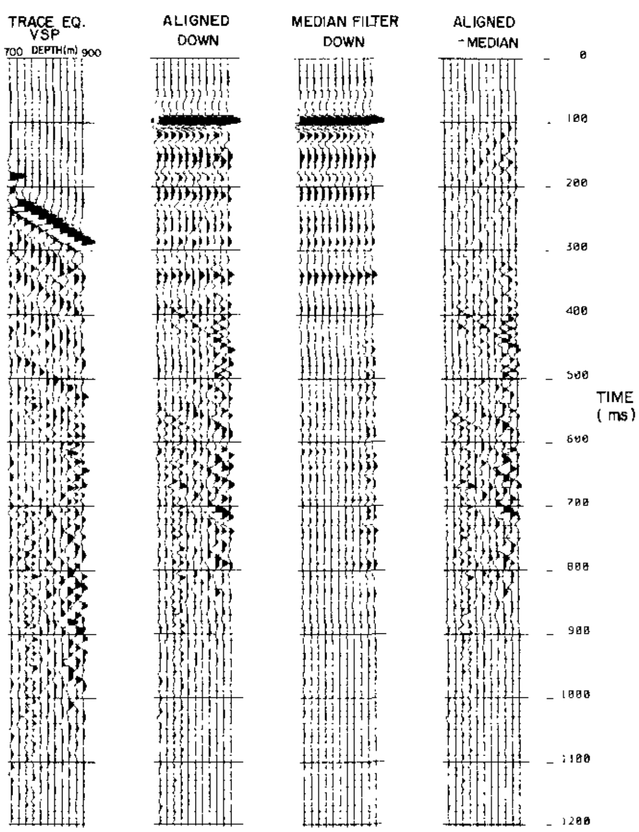

Fig. 6. Group of VSP traces from a western Canada VSP. The first set of traces is the trace-equalized raw data. The next group has the traces aligned on their first breaks. The third group shows the median-filtered aligned traces. The fourth group is the result of subtracting the median-filtered output (3rd group) from its input (2nd group). Note the excellent rejection of the aligned (downgoing) waves.

remely effective. Several VSP examples are briefly discussed to illustrate the effectiveness of the median filter. Figure 6 shows a processing example for a limited number of VSP traces from a VSP conducted in western Canada. The median filter is applied to the aligned input data to enhance the first arrivals (the downgoing waves). This enhanced downgoing event is then subtracted from the aligned input data. The subtraction result contains the upgoing waves as well as other events. These upgoing waves can be further processed. Note how well the large downgoing wave has been removed from the aligned input data.

Figure 7 shows the upgoing waves of another VSP from western Canada shot over a larger depth. There is a large low-velocity event, starting at about 450 ms on the leftmost trace (possibly a source-generated shear wave) in the near surface. There are also a number of very noisy traces. Figure 8 shows the results of an 11-trace median filter applied to the preceding data. Note the glitch rejection as well as the low-velocity event rejection. Another important feature of this median-filtering process is that the first arriving energy is not smeared out into the pre-first-break region. This is an important feature in VSP processing, as preservation of the exact first-break and upgoing event times is critical in the later stages of VSP interpretation.

2-D Median Filter

The median filter as previously described provides excellent despiking and dip-rejection capabilities. However, it is sometimes desirable to retain a range of dips in data processing. Velocity or f/k filters have been used for some time in production data processing to pass or attenuate specific dipping events (Embree *et al.*, 1963; March and Bailey, 1983). It would be very useful to develop a hybrid filter that rejects spikes but passes step functions and a range of dips. It is to this end that the median selection is expanded to a box of data points as opposed to a line of points (Huang *et al.*, 1979; Kirlin *et al.*, 1985). If all points were considered equally in the box, then the median selection could be from a dipping event. It is possible, furthermore, to give different importances (weights) to different numbers in the window (Claerbout and Muir, 1973). Thus if it were felt that the centre number were most important it could be given a weight of 10, say, with the other numbers given weights of 1. A weight of 10 would simply mean that this centre number would be repeated 10 times in the ordering process. This would change the sample selected for the median.

If the f/k filter is viewed in the time domain, then it is possible to use the time domain f/k weights to apply to the data at the appropriate points in the box window. The f/k filter used in this paper employs user-specified event dips (ms/trace) which are to be passed and rejected.

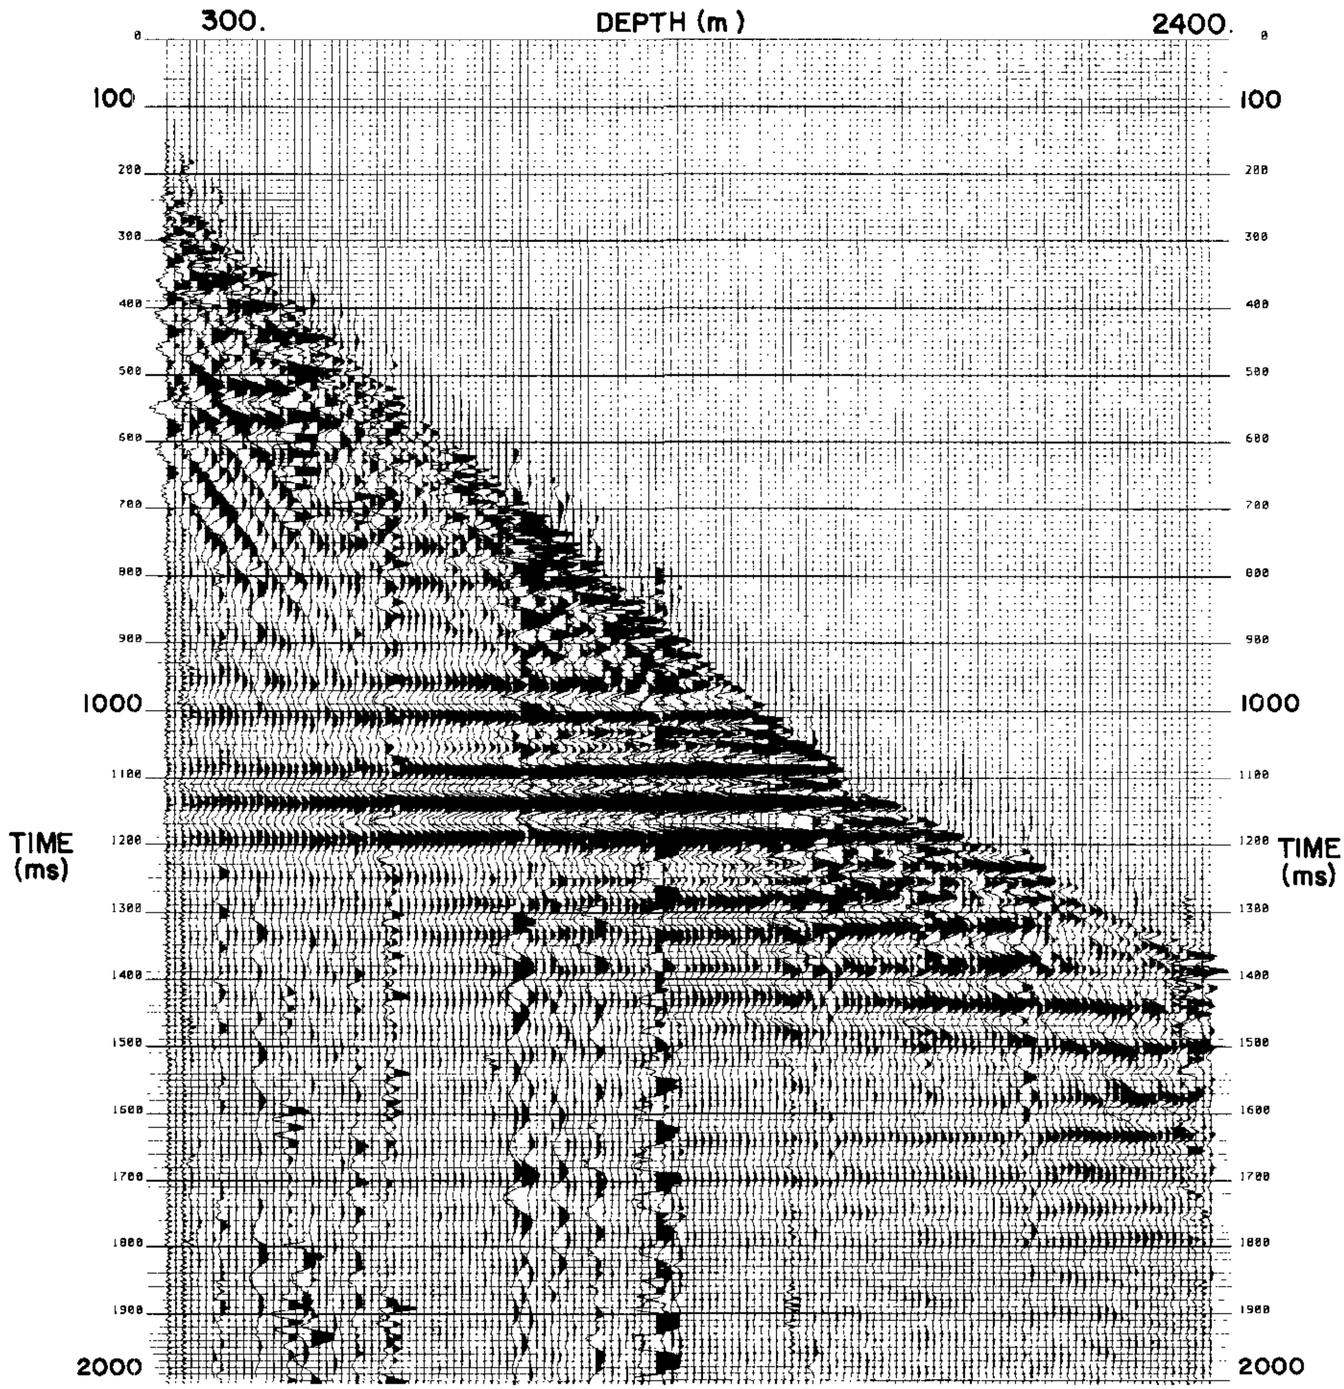

Fig. 7. Aligned upgoing VSP waves in a western Canada VSP.

These dips are converted to pie-slice regions in the f/k domain. This pie-slice in the f/k domain is then transformed back to the x/t (space-time) domain, where it can be applied to the seismic traces.

Figure 9 shows the time response (truncated spatially and temporally) of an f/k filter designed to pass dips of less than 1 ms/trace and reject dips of more than 2 ms/trace. In standard time-domain f/k filtering these weights would be multiplied by the associated data values, and the average of all these products would be the filtered output value. How can these weights be applied in a median filter? For a given box of points (say 11 adjacent traces) over 19 ms, each point is assigned its corresponding weight as given in Figure 9. The points of the box are then sorted into ascending order, always keeping their respective weights with them. When the weights of the ordered sequence total to greater than one-half the cumulative weight of the whole box, this point is selected as the f/k weighted median. The box centres on every point of the input assemblage of traces and selects the weighted median point as the filtered point to be output at the centre of the box. It is some-

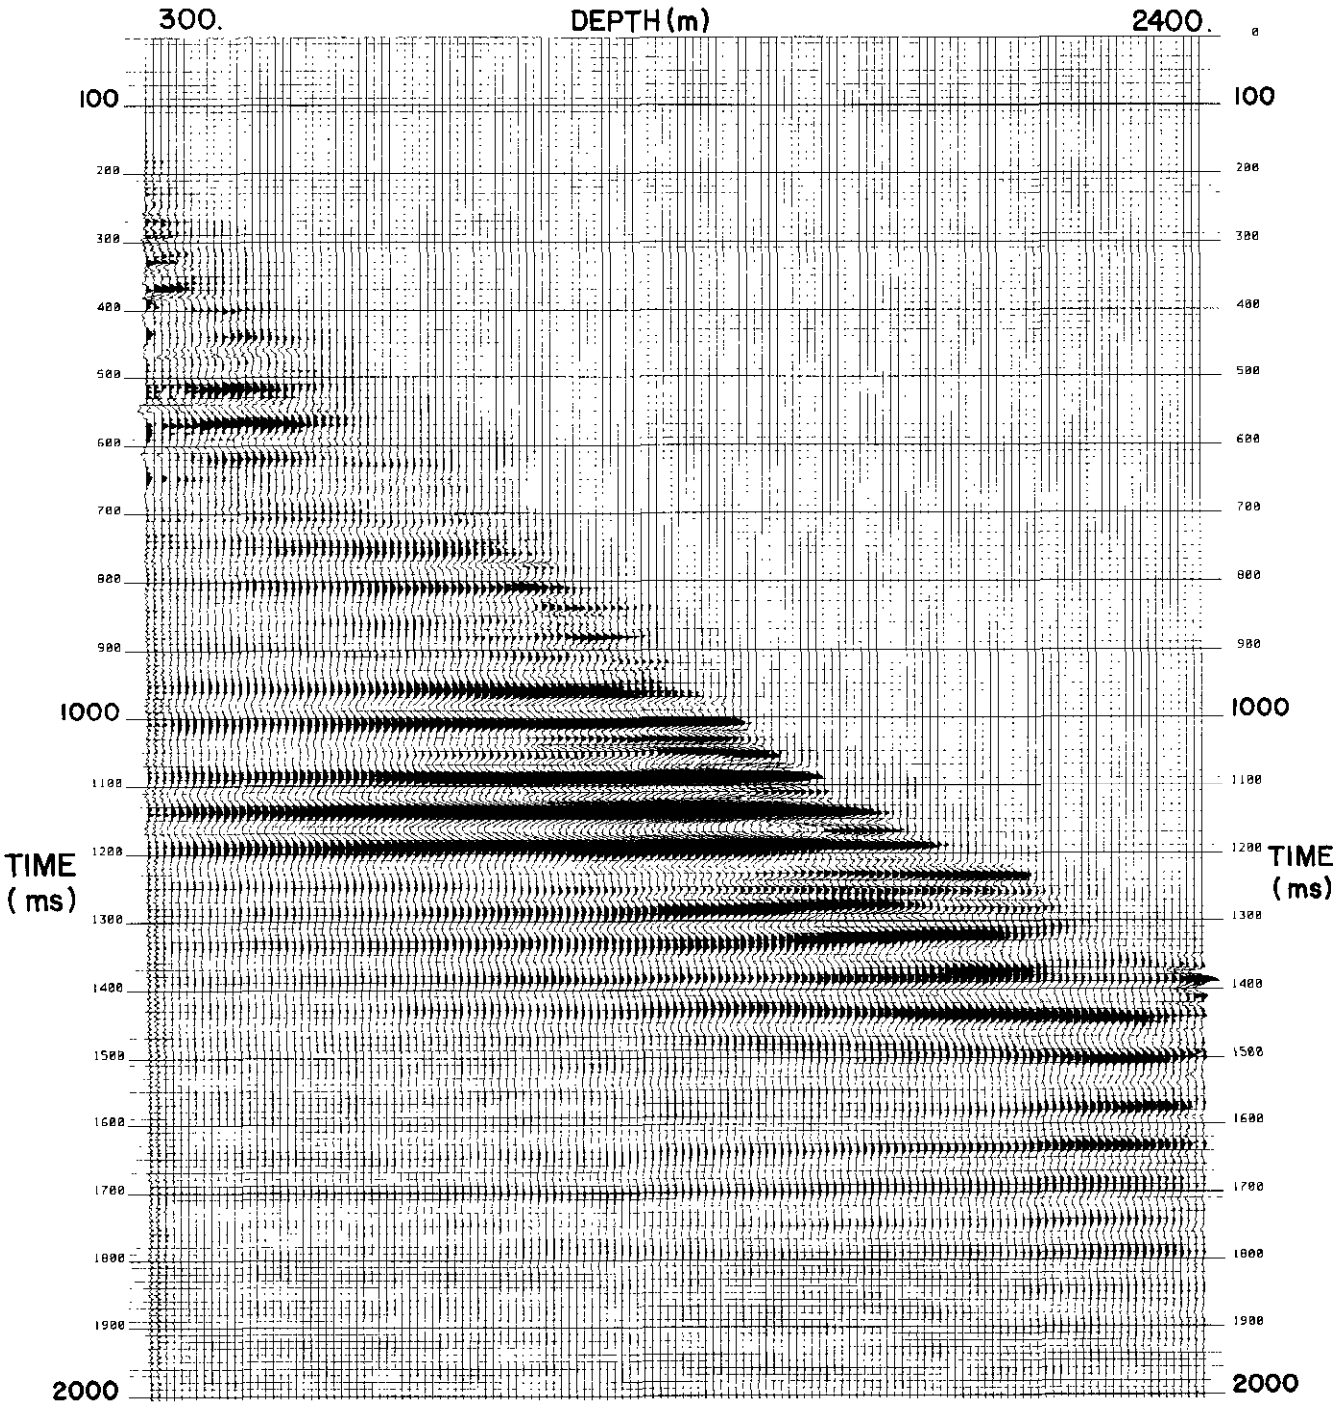

Fig. 8. Median filtered upgoing waves from Figure 7 (11-trace filter).

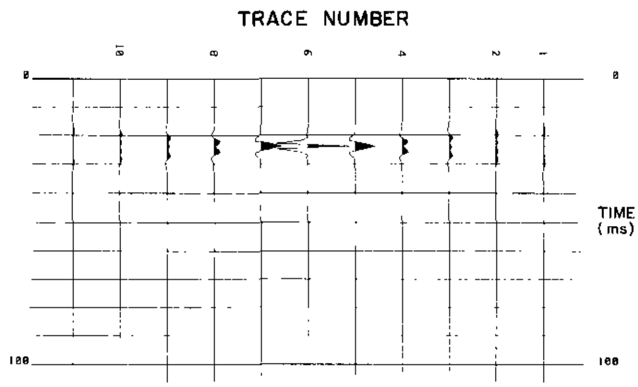

Fig. 9. Truncated time-domain response of an f/k filter designed to pass dips of 1 ms/trace and reject dips greater than 2 ms/trace.

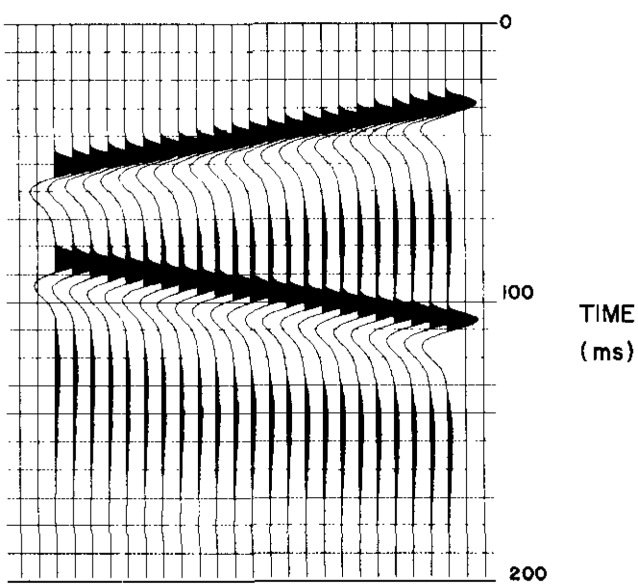

Fig. 10. Synthetic traces with linear dips of 1 ms/trace and −1 ms/trace.

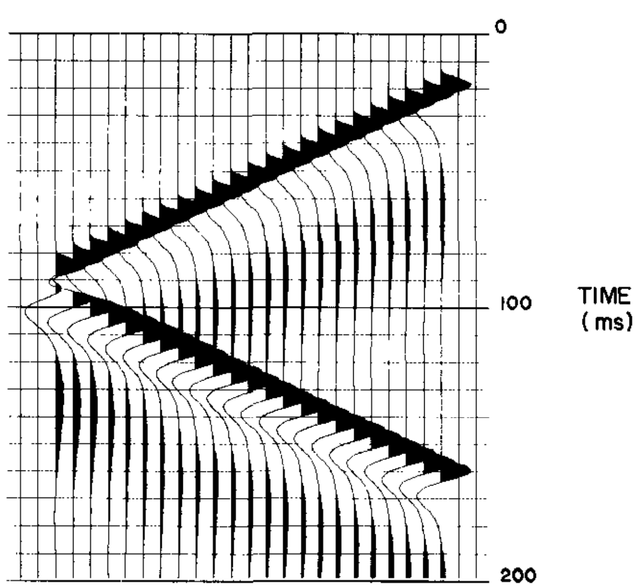

Fig. 11. Synthetic traces with linear dips of 3 ms/trace and −3 ms/trace.

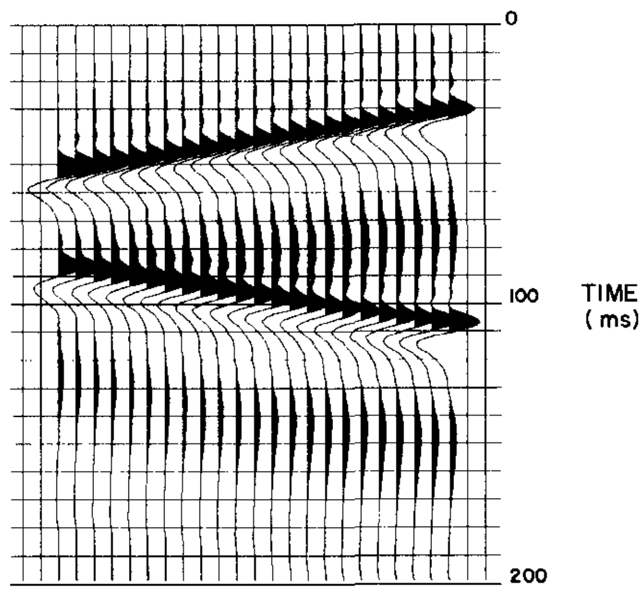

Fig. 12. Linear dipping data (± 1 ms/trace) filtered with f/k median filter that passes dips less than 1 ms/trace and rejects dips greater than 2 ms/trace.

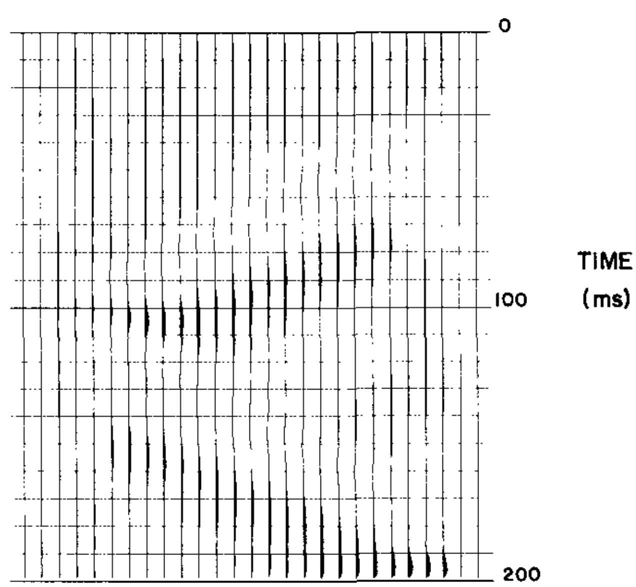

Fig. 13. Fk median filter as in Figure 12 used on data of Figure 11 (synthetic traces with dips of ± 3 ms/trace).

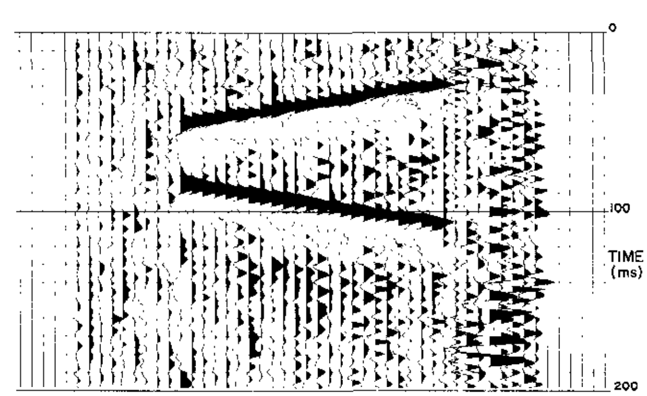

Fig. 14. Linear dips (± 1 ms/trace) of Figure 10 with additive band-limited, field-recorded noise.

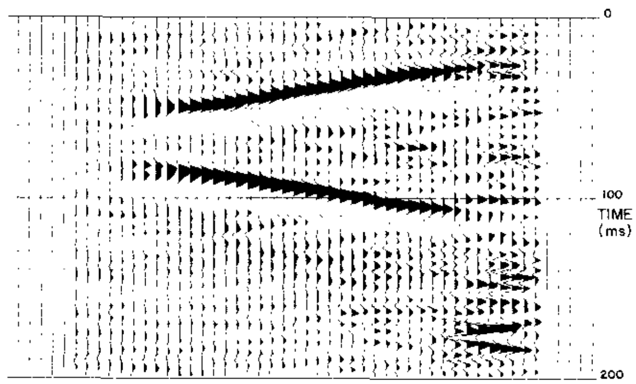

Fig. 15. Results of filtering noisy data of Figure 14 (± 1 ms/trace dipping events) with standard f/k filter.

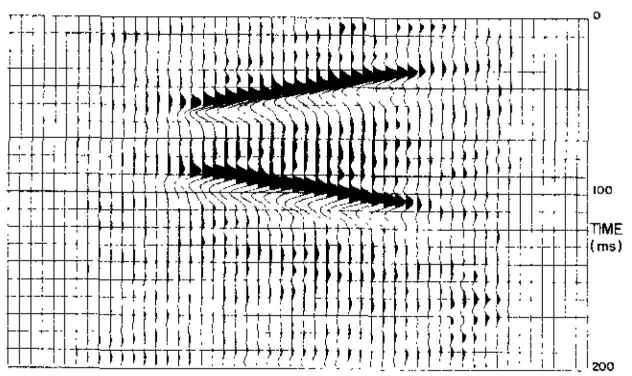

Fig. 16. Results of filtering noisy data of Figure 14 (± 1 ms/trace dipping events) with the f/k median filter.

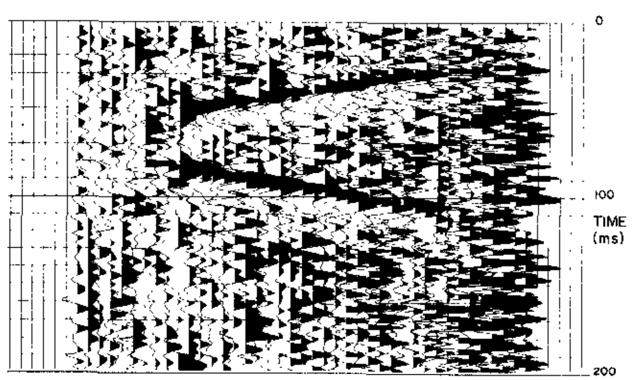

Fig. 17. Linear dips of Figure 10 (± 1 ms/trace dipping events) with three times the noise level of Figure 14.

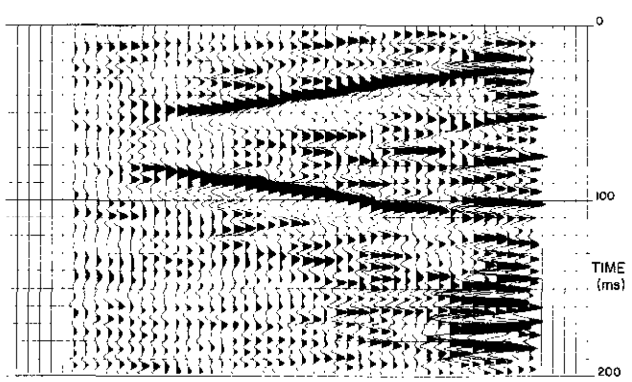

Fig. 18. Results of filtering noisy data of Figure 17 with standard f/k filter.

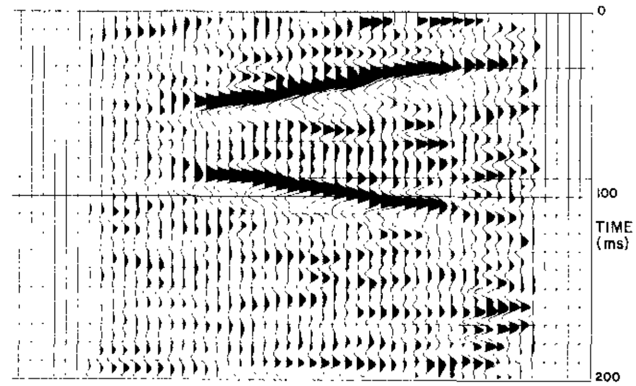

Fig. 19. Results of filtering noisy data of Figure 17 with f/k median filter.

what remarkable that this f/k weighted selection appears to preserve the dip pass and rejection properties of the f/k filter but also has the spike rejection property of the median filter. This point is discussed in the next series of figures.

To test the dip pass and dip rejection properties of the f/k median filter, two synthetic dipping events are filtered. The first set of synthetic traces has two events dipping at 1 ms/trace and − 1 ms/trace (Fig. 10). The event itself is constructed by convolving a spike with 10 Hz to 60 Hz bandpass 3rd order Butterworth filter. The second set has events dipping at 3 ms/trace and −3 ms/trace (Fig. 11). An f/k filter is designed to pass dips up to 1 ms/trace and reject dips greater than 2 ms/trace. Only a portion of the time domain weights are used in the median filter process to improve computer run time (an 11-trace by 19-ms array). The results of the f/k weighted median filter applied to these two cases are shown in Figures 12 and 13. Note that the f/k median filter has largely passed the ± 1 ms/trace dipping traces while the events dipping with ± 3 ms/trace have been attenuated. Thus, the dip pass and reject properties of the f/k filter have been preserved. Note that the amplitudes of the plots in these and the following cases have been adjusted to make their signal amplitudes approximately equal. More work remains to be done on how sensitive the median filter result is to the size of the box window. Preliminary tests suggest that increasing the window size (with corresponding weights) provides better results.

More important tests for the median filter lie in its performance on noisy data. The two following examples compare the standard f/k filter and the f/k median filter on noisy data sets. Figure 14 displays the two dipping events of Figure 10, but now with noise added. The noise was taken from actual VSP field records before the first-breaking arrivals. Application of a standard f/k filter (1 ms/trace - pass, 2 ms/trace - reject) gives the output shown in Figure 15. In general, the f/k filter has reduced the noise. However, large noise events (as in the bottom right corner of the plot) are smeared laterally. The ends of the event lines are also smeared laterally. The result of f/k median filtering (Fig. 16) is significantly different. The large noise events have been largely rejected everywhere. Also, the ends of the event lines suffer less lateral smearing.

Figure 17 shows the same dipping events as Figure 10 but now with three times the noise added in Figure 14. The result of standard f/k filtering is shown in Figure 18. Once again large amplitude noise is smeared laterally (as are the discontinuous events), resulting in a "wormy" appearance of the data. The f/k median filter result (Fig. 19) attenuates the high-amplitude noise significantly better than the standard f/k filter. Once again the f/k median filter has left the event ends less smeared laterally. Note also that, if there were a few very large glitches on the traces, the median algorithm would reject them entirely while the standard f/k filter would spread the energy in the glitch over numerous traces. For these cases, the f/k median filter appears to perform better than the standard f/k filter.

Conclusions

Median filters are used routinely to advantage in VSP data processing, and to a lesser extent in producing alternative CDP sections and velocity analyses. The useful properties of the median filter include its glitch rejection (automatic editing) and its ability to pass step functions (not smear edges). The median filter behaves similarly, in some ways, to a low-pass filter. The linear median filter also attenuates dips that are not linearly aligned in the direction of its application. This is a very useful characteristic for enhancing desired energy and attenuating other events. Sometimes, however, it is desirable to pass a range of dipping energy. The f/k filter is well-known for its utility in dip enhancement or rejection. By using the time-domain responses of the f/k filter as weights for a median selection, a hybrid f/k median filter has been designed here. Early tests with this filter indicate that it combines the properties of the median filter (glitch rejection, edge preservation) as well as those of the f/k filter (specific dip attenuation). Preliminary results with this filter suggest that it will be useful in processing noisy data in which it is desirable to preserve dipping events. It may outperform standard f/k filtering in a noisy environment. Further tests need to be conducted to describe quantitatively the characteristics of this f/k median filter, and how much better it performs in a noisy environment than the standard f/k filter.

References

Bednar J.B., 1983, Applications of median filtering to deconvolution, pulse estimation and statistical editing of seismic data: Geophysics, v. 48, p. 1598-1610.

Claerbout, J.F., 1976, Fundamentals of geophysical data processing: New York, McGraw-Hill.

——— and Muir, F., 1973, Robust modelling with erratic data: Geophysics, v. 38, p. 826.

Denoel, E. and Solvay, J.P., 1985, Linear prediction of speech with a least absolute error criterion: IEEE Trans. Acoust., Speech, Signal Processing, v. ASSP-33, no. 6.

DiSiena, J.P. and Stewart, R.R., 1985, Personal communication, an introduction to median filtering: S.E.G. Continuing Education, V.S.P. Course Notes, Section 3, p. 44-54.

Embree, P., Burg, J.P. and Backus, M.M., 1963, Wide-band velocity filtering — the pie-slice process: Geophysics, v. 28, no.6.

Evans, J.R., 1982, Running median filters and a general despiker: Bull. Seis. Soc. America, v. 72, no. 1, p. 331-338.

Farmer, P. and Haldorsen, J., 1985, Suppression of high energy noise using an alternate stacking procedure: Paper S 20.1, Expanded Abstracts of 55th Annual International SEG Meeting, Oct. 6-10, Washington, D.C.

Fitch, E.P., Coyle, E.J. and Gallagher, N.C., 1984, Median filtering by threshold decomposition: IEEE Trans. Acoust., Speech, Signal Processing, v. ASSP-32, no. 6, p. 1183-1188.

———, Coyle, E.J. and Gallagher, N.C., 1985, Root properties and convergence rates of median filters: IEEE Trans. Acoust., Speech, Signal Processing, v. ASSP-33, no. 1, p. 230-239.

Gallagher, N.C. and Wise, G.L., 1981, A theoretical analysis of the properties of median filters: IEEE Trans. Acoust., Speech, Signal Processing, v. ASSP -29, no. 6, December.

Hardage, B.A., 1983, Vertical Seismic Profiling, Part A: Principles: Amsterdam, Geophysical Press, p. 182.

———, 1985, Personal communication: S.E.G. Continuing Education, V.S.P. Course Notes, Section 3 (by J.P. DiSiena and R.R. Stewart).

Huang, T.S., Yang, G.J. and Tang, G.Y., 1979, Fast, two-dimensional median filtering algorithm: IEEE Trans. Acoust., Speech, Signal Processing, v. ASSP-20, no. 1, February.

Kirlin, R.L., Cudzilo, B. and Wilson, S., 1985, Two-dimensional orthogonal median filters and applications: Proc. International Conf. Acoustics, Speech, Signal Processing, Tampa, Florida, March, p. 1325 -1328.

Lee, Y.H. and Kassam S.A., 1985, Generalized median filtering and related non-linear filtering techniques: IEEE Trans. Acoust., Speech, Signal Processing, v. ASSP-33, no. 3, p. 672-683.

March, D.W. and Bailey, A.D., 1983, A review of the two-dimensional transform and its use in seismic processing: First Break, v. 1, no. 1.

Naess, O.E. and Bruland L., 1985, Stacking methods other than simple summation. *In:* Fitch, A.A. (Ed.), Developments in Geophysical Exploration Methods - 6, London, Elsevier Appl. Sci. Publishers.

Nodes, T.A. and Gallagher, N.C., 1982, Median filters: some modifications and their properties: IEEE Trans., Acoust., Speech, Signal Processing, v. ASSP-30, no. 5, p. 739-746.

Taylor, H.L., 1981, The l_1 norm in seismic data processing. *In:* Fitch, A.A. (Ed.), Developments in Geophysical Exploration Methods - 2, London, Elsevier Appl. Sci. Publishers.

Tukey, J.W., 1971, Exploratory Data Analysis (preliminary ed.). Reading, MA: Addision-Wesley.

Watt, T. and Bednar, J.B., 1983, Role of the alpha-trimmed mean in combining and analysing seismic common-depth-point gathers: 53rd Annual SEG Meeting, Las Vegas, Nevada, Sept. 11-15.