Median spectra

Marta Woodward and *Joe Dellinger*

INTRODUCTION

This paper addresses the problem of finding a single representative spectrum for a group of traces. While motivated primarily as an academic exercise, its observations are pertinent wherever a single spectrum might be used to generate a minimum phase source wavelet or prediction error filter for a larger data set—e.g., in the design both of source wavelets for synthetic modeling and of prediction error filters for surface consistent deconvolution. Six possible solutions to the problem are implemented below on the three dynamite-generated shot profiles illustrated in Figure 1. The methods fall into two main groups according to spectral estimation technique. The four methods composing the first group produce spectral estimates for each trace through zero padding and fast Fourier transformation; they assume the traces are zero outside the recorded interval. The two methods composing the second group produce spectral estimates using Burg's maximum entropy spectral estimation algorithm; they make as few assumptions as possible about the data outside the recorded interval. Within each group the methods are distinguished by the ways in which they combine spectral estimates from individual traces.

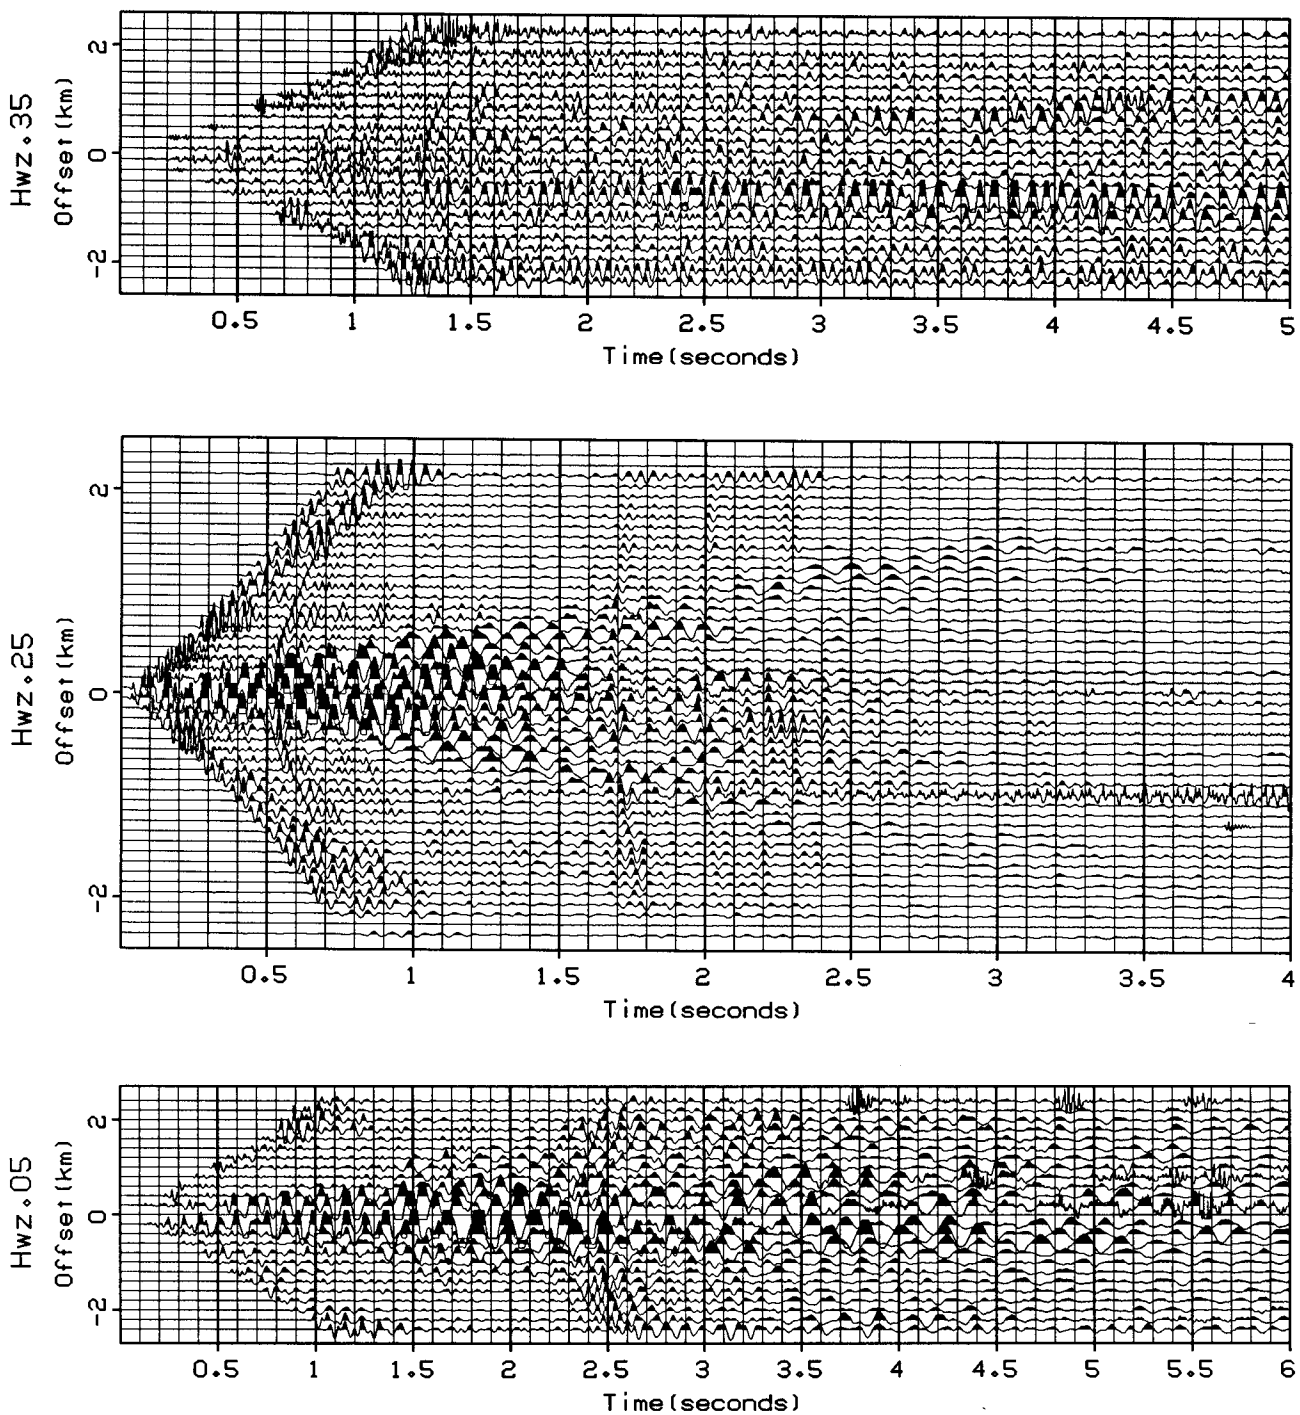

FIG. 1. Three dynamite-generated, .002 msec sampled shot profiles donated by Western Geophysical. The traces have been gained through multiplication by t^2.

FOURIER TRANSFORM/WINDOW METHODS

Mean versus median

The first two methods in this group combine raw spectral estimates from individual traces into a single representative estimate by selecting either the mean or median power at each frequency. Because power is always positive, these operations guarantee formation of a power spectrum. The results of implementing these methods on the data sets of Figure 1 are shown in Figures 2, 3 and 4; the final spectra have all been normalized such that their integrated power equals one. Comparison of the illustrated mean and median spectra yields two conclusions: first, the mean method clearly demonstrates a greater contribution from low frequency ground roll and high frequency background noise; second, the mean appears smoothed relative to the median—varying less from sample point to sample point. Both these results are typical of the classic differences between means and medians.

The first contrast is most striking for Hwz.25, that data set which exhibits the strongest contribution of low frequency ground roll on its inner traces. (See Figure 5 for illustrations of the median spectra corresponding to the ground roll-dominated four inner traces of Hwz.05 and Hwz.25.) Because ground roll appears as overwhelming, anomalous energy on the frequency spectra of the affected traces, it significantly increases the mean spectral power at frequencies characteristic of ground roll relative to those characteristic of reflected energy; because ground roll affects only a minority of the traces, its impact on the median spectral power is much smaller. Similar arguments explain the relatively larger high frequency content of the mean as compared to the median spectra. Indeed, the small high frequency spikes appearing on the mean spectrum for Hwz.25 may each be attributed to a single anomalous spike appearing on one of several otherwise normal trace spectra. The robustness of the median—its insensitivity to anomalous frequency content affecting only a small proportion of the traces—makes it superior to the mean as an estimator of a group spectrum.

The roughness of the median relative to the mean results from the fact that a mean weights and sums data while a median selects a single value. Jumping discontinuously as it picks one data point from one spectrum and the next from another, the median preserves the true trace to trace variability in frequency content. Figure 6 demonstrates just how little inclination the median has to follow any particular trace spectrum: the top plot shows which trace's spectrum contributed the median value at each frequency for data set Hwz.25; the bottom plot shows how often each trace's spectrum was chosen in the 0 to 125 Hz range. It is hard to find more than a few cases in the entire figure where the

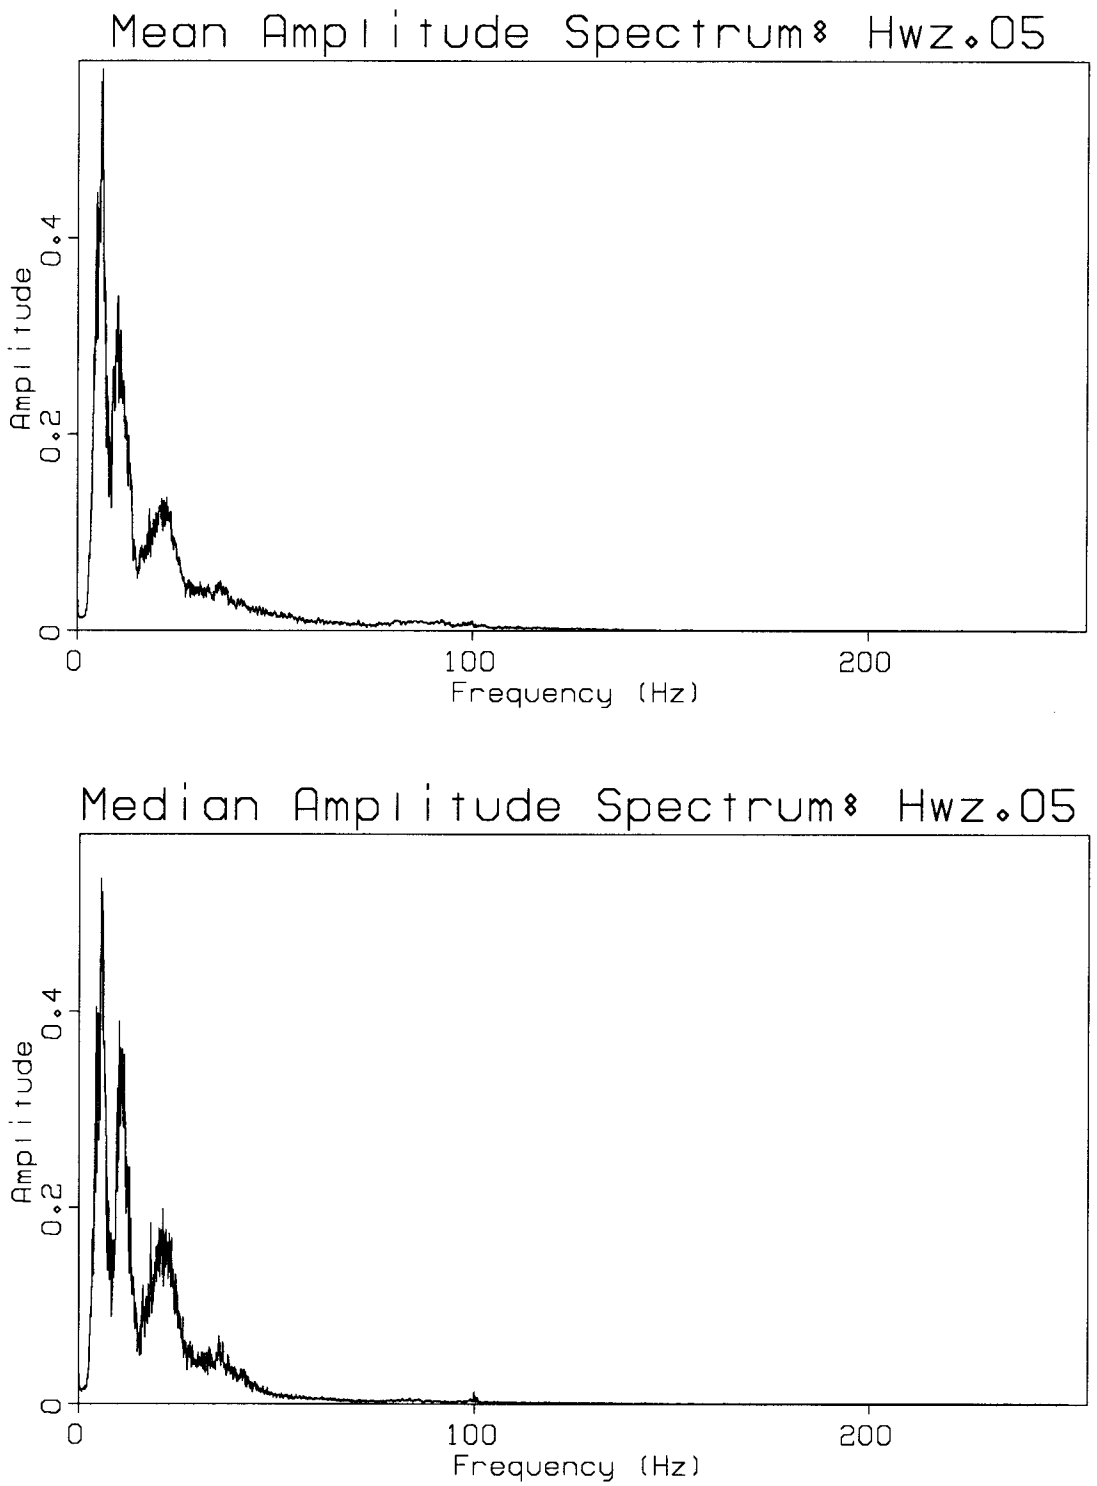

FIG. 2. Mean and median amplitude spectra for data set Hwz.05, calculated from unnormalized trace spectra. The plotted spectra have been normalized such that their integrated powers equal one.

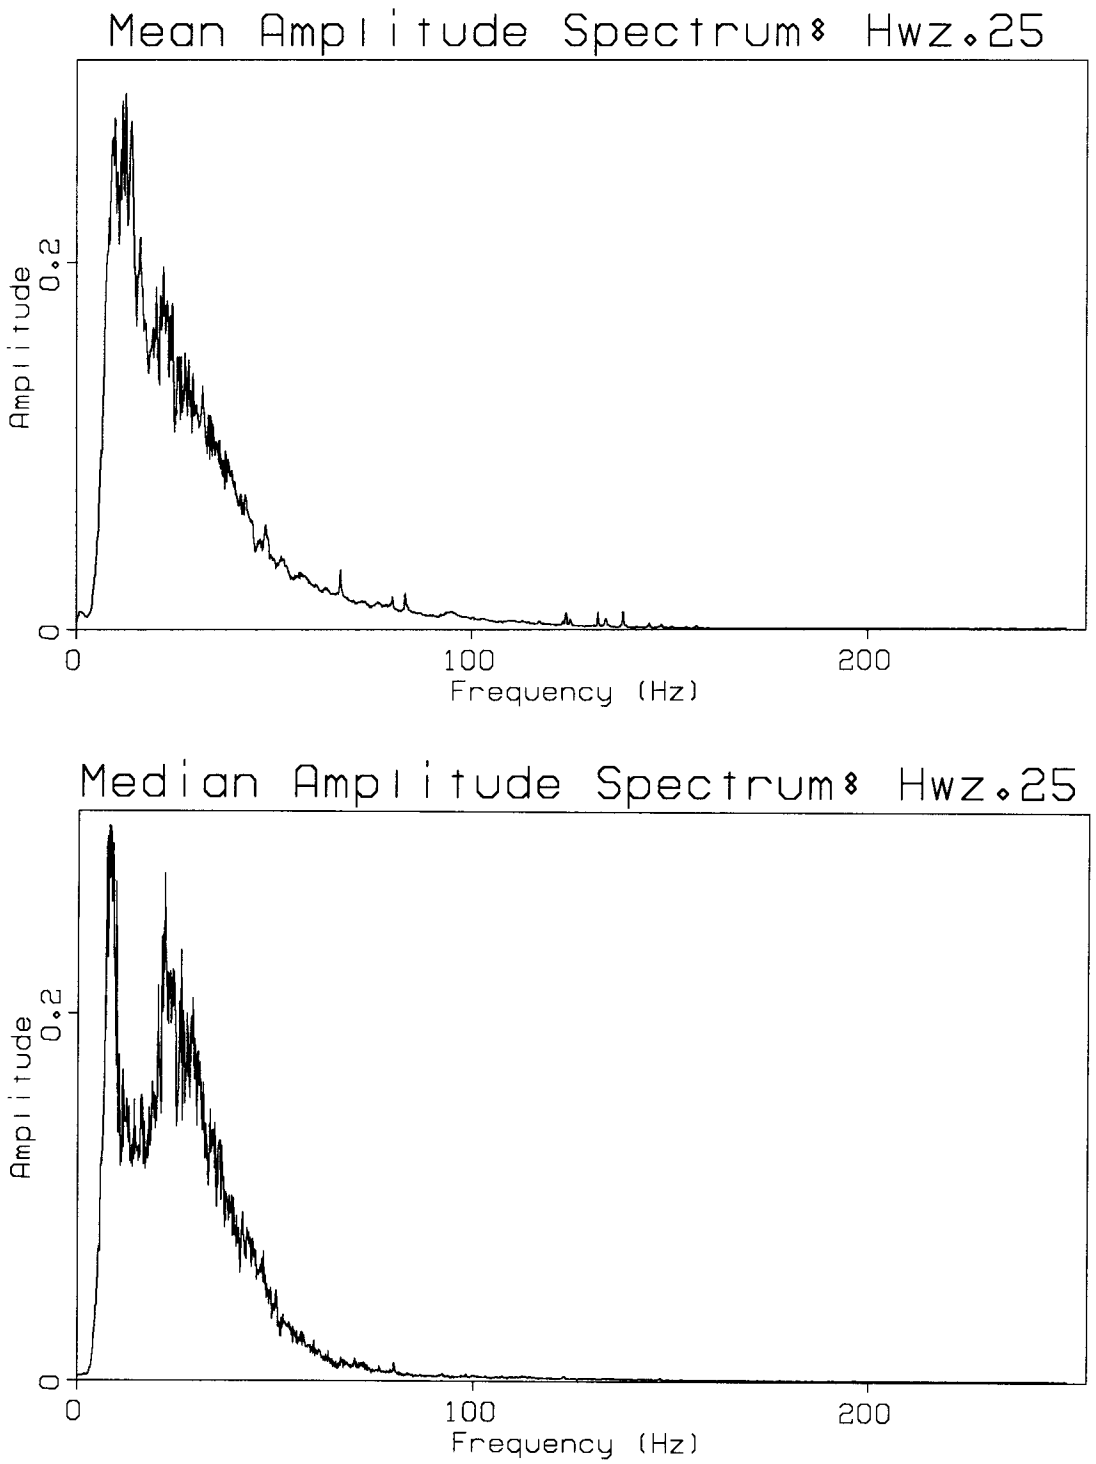

FIG. 3. Mean and median amplitude spectra for data set Hwz.25, calculated from unnormalized trace spectra. The plotted spectra have been normalized such that their integrated powers equal one.

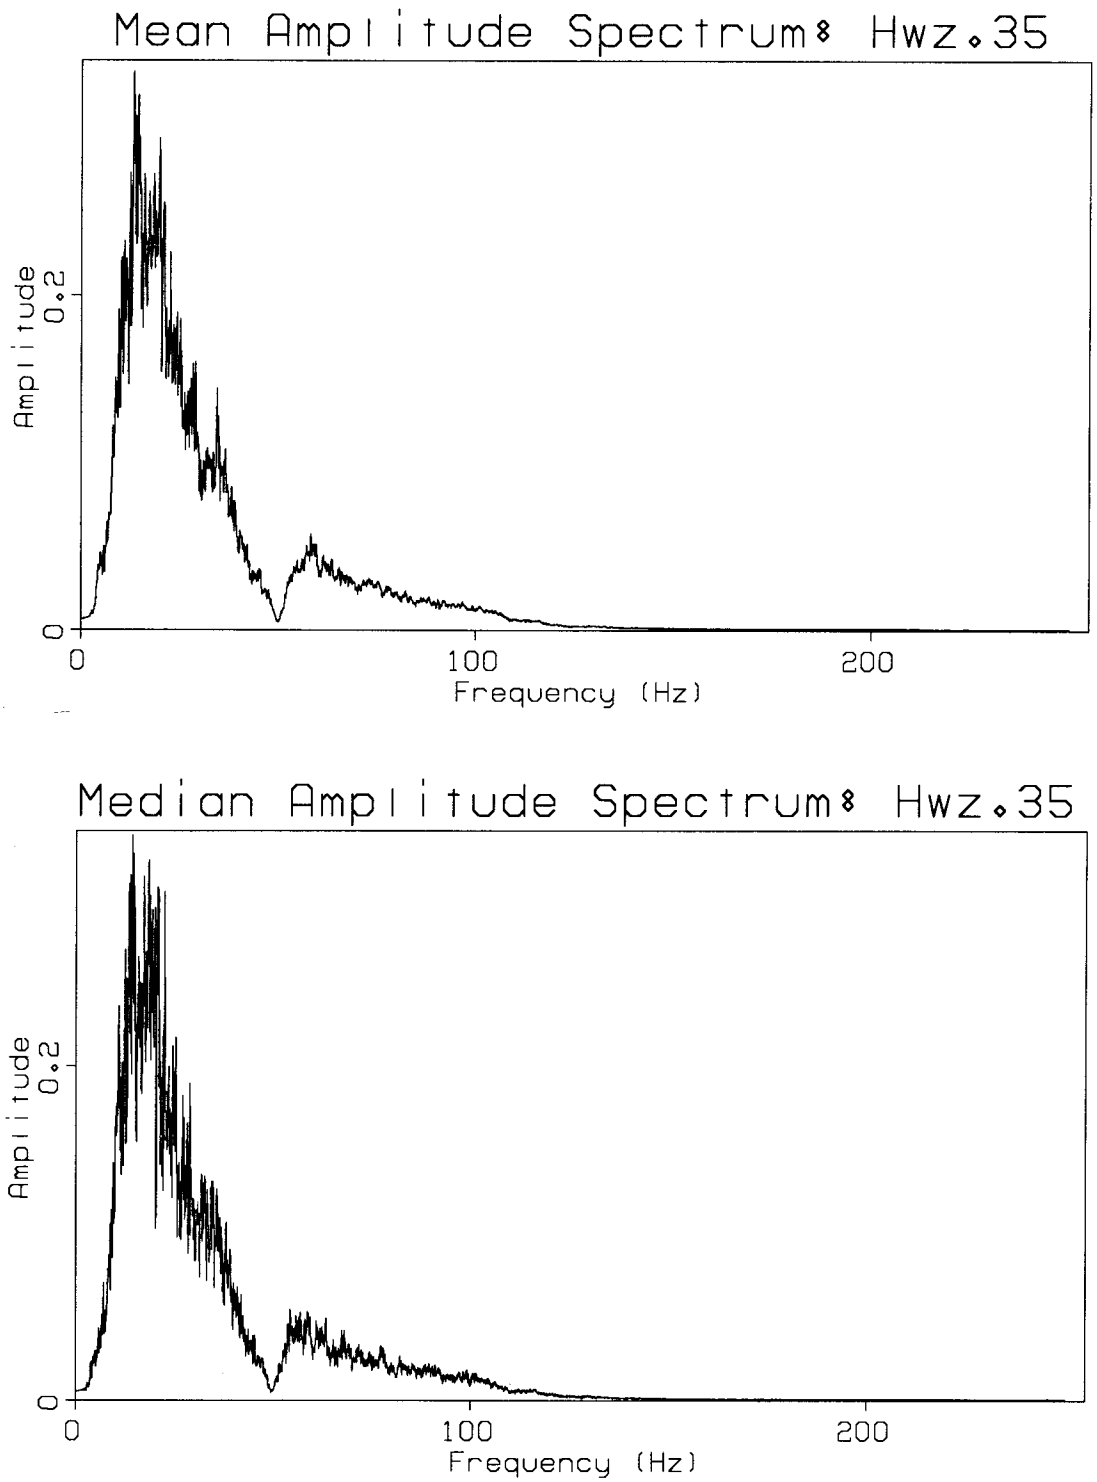

FIG. 4. Mean and median amplitude spectra for data set Hwz.35, calculated from unnormalized trace spectra. The plotted spectra have been normalized such that their integrated powers equal one.

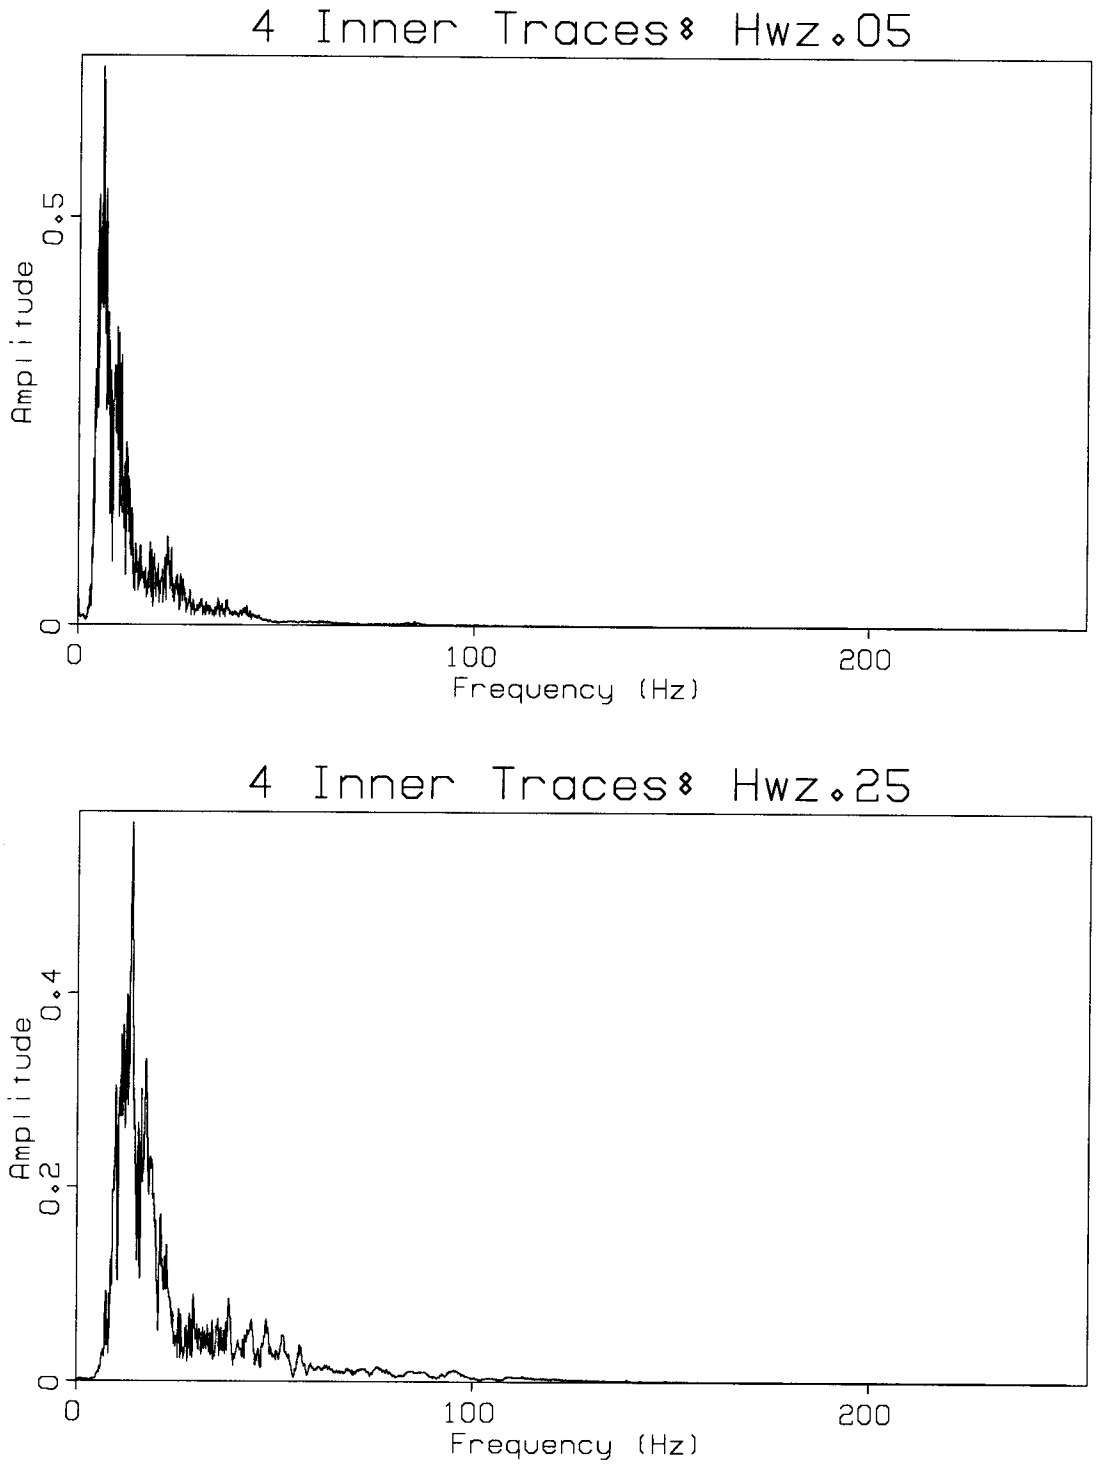

FIG. 5. Median amplitude spectra derived from the unnormalized, ground roll-dominated four inner traces of data sets Hwz.05 and Hwz.25, respectively. The plotted spectra have been normalized such that their integrated powers equal one.

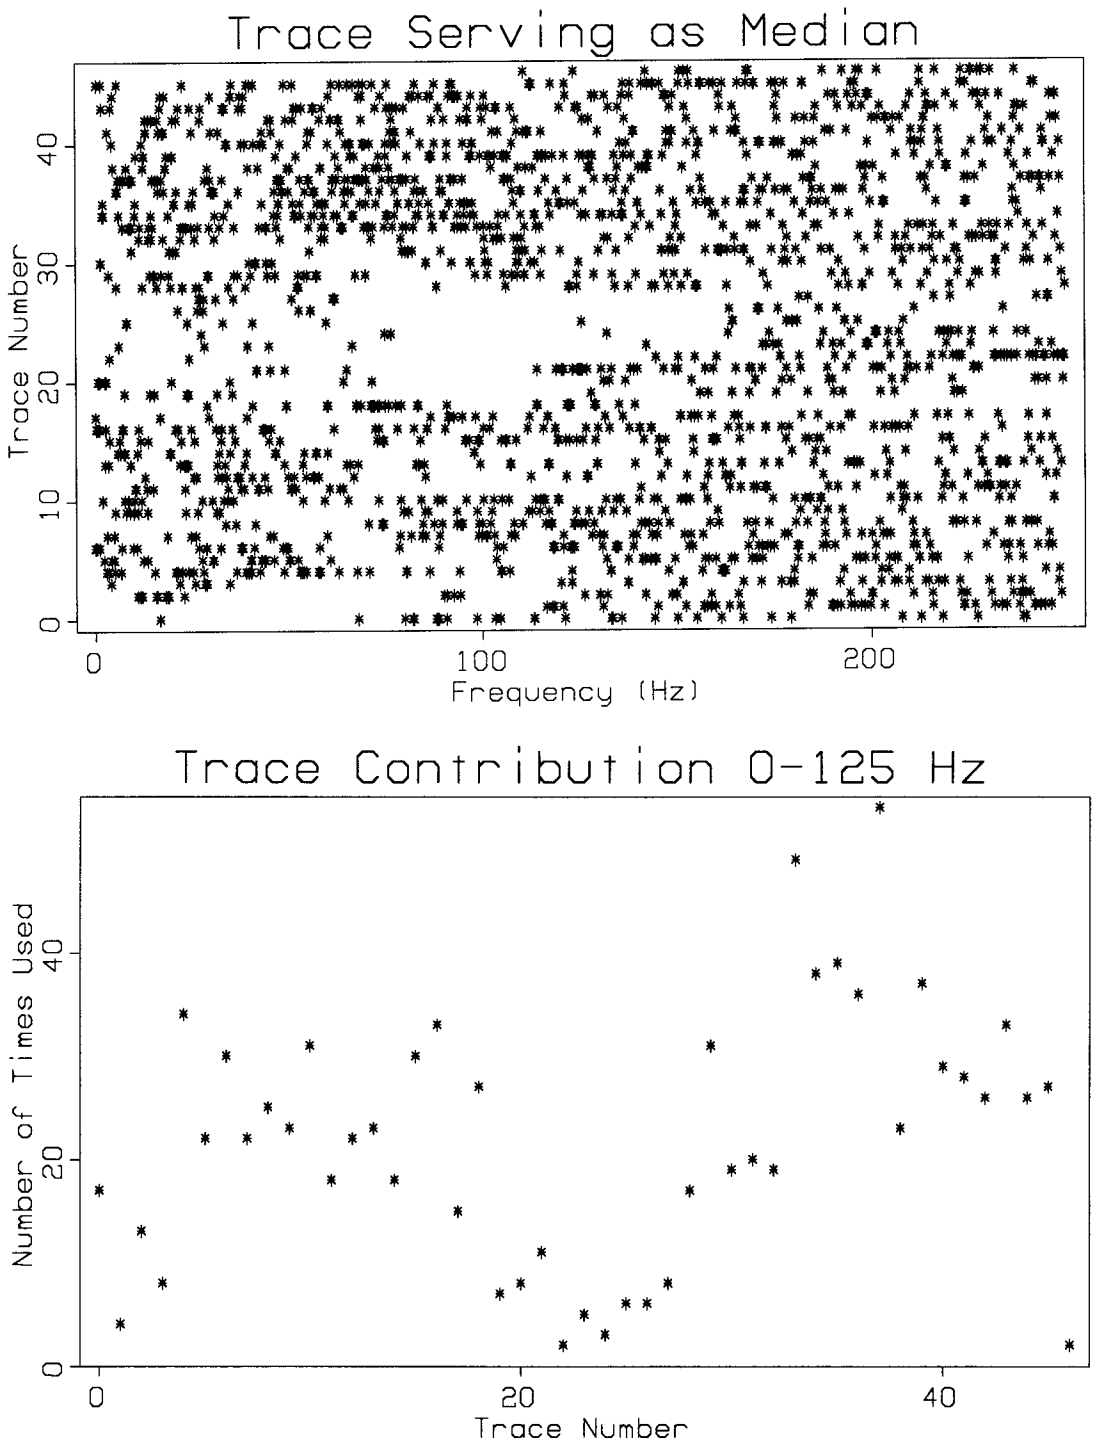

FIG. 6. Contribution of different trace spectra from data set Hwz.25 to the median spectrum for that data set. The top plot shows which trace's spectrum served as the median value at each frequency; the bottom plot shows how many times each trace's spectrum value was picked as the median for the range of frequencies from 0 to half the Nyquist.

same trace served as the median twice in a row. The figure also shows the median operator discriminating against ground roll and garbage traces (such as traces 2 and 48). Given that all of the median and mean spectral estimates derived above would require smoothing for transformation into suitably short deconvolution operators, the roughness of the median is of very little significance.

Normalized versus unnormalized input trace spectra

The second two methods in this group correspond to implementation of the median and mean methods on trace spectra that are individually normalized prior to combination. Comparison of the resulting median and mean spectra shown in Figures 7, 8 and 9 reveals fewer pronounced differences than in the previous case at low frequencies. Because normalization of individual trace spectra with respect to integrated spectral power is equivalent to trace balancing, the occasional spectral peak reflecting large amounts of anomalous ground roll energy has been reduced in magnitude—causing the mean to more closely resemble the median. Accompanying this advantage are several disadvantages equally associated with trace balancing: first, the extra step requires extra computation; second, the method depresses valid frequency content along with high magnitude noise components on affected traces; third, the technique introduces the risk of amplifying bad (low signal to noise ratio) traces. This last disadvantage appears most noticeably in the mean spectrum for Hwz.25 (Figure 8). Due to the presence of several near dead traces in the data set, the high frequency noise content has been amplified instead of reduced.

BURG ALGORITHM METHODS

Burg median amplitude spectra

The first method in this group is the Burg algorithm equivalent to taking the median of normalized spectra described above. Because the Burg algorithm generates prediction error filters in the time domain, the spectral estimates exhibited in the tops of Figures 10 and 11 were calculated by taking one over the amplitude spectrum of the appropriate, normalized inverse filter. Because these filters have only been calculated out to 128 points—a more than reasonable length for deconvolution operators—their 128 point inverse spectra are much smoother than those derived in the previous section. Nevertheless, it is clear that their overall shapes correspond closely to those of the median spectra of Figures 2, 3,7 and 8.

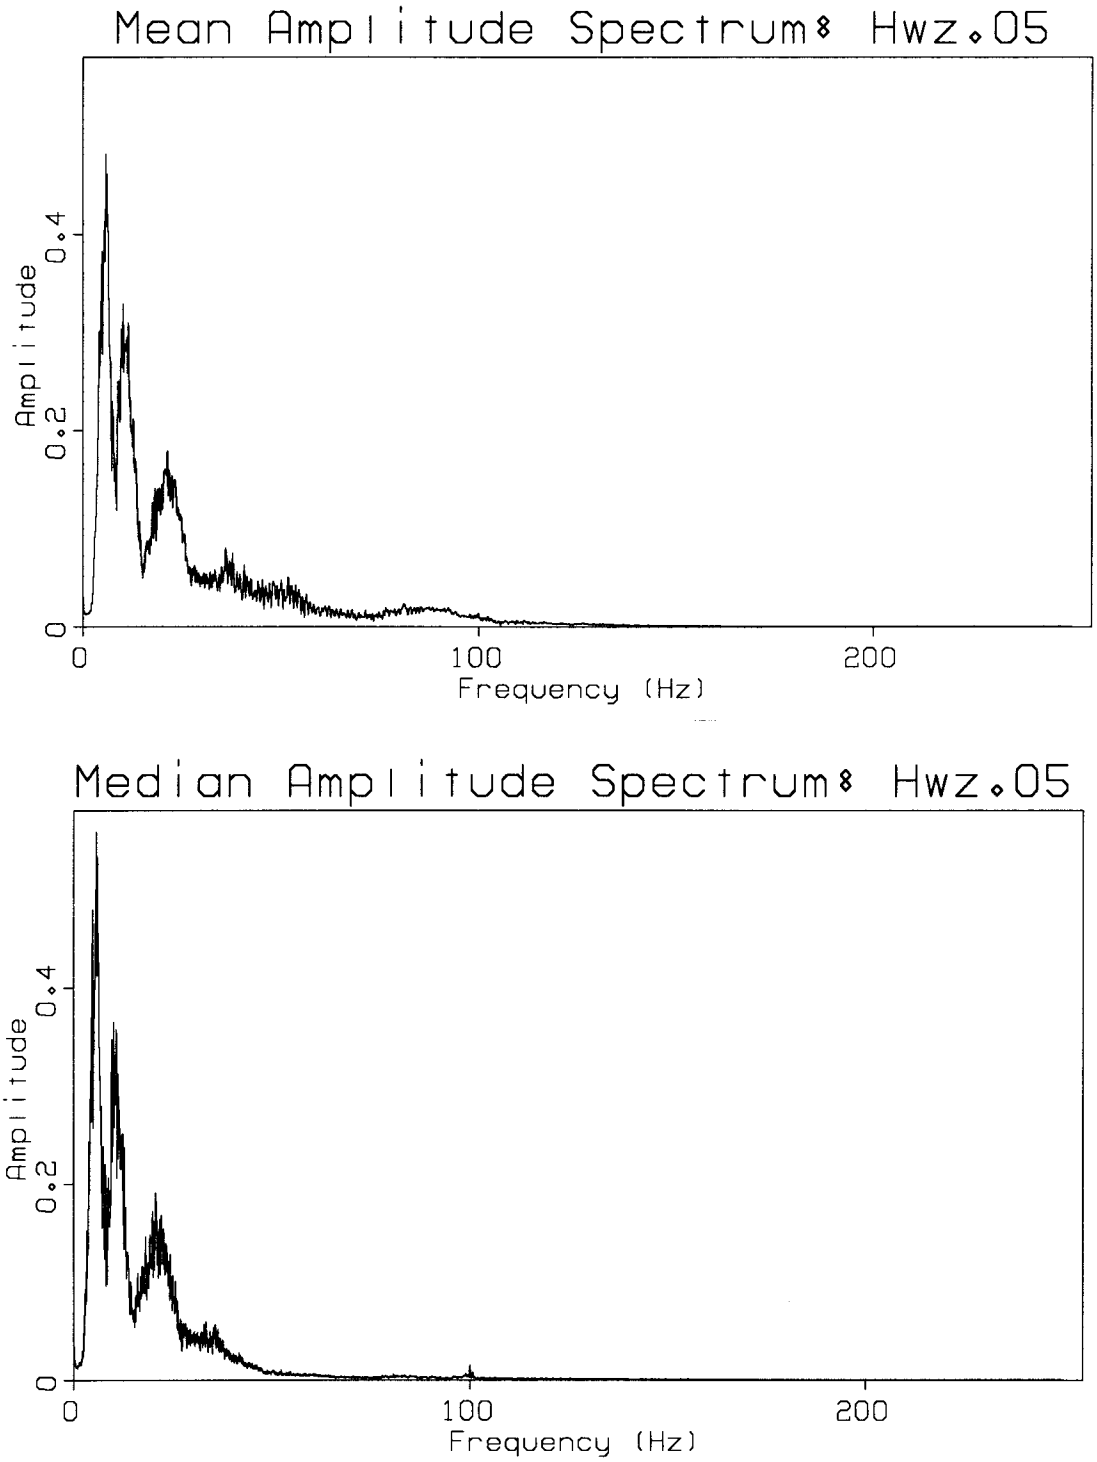

FIG. 7. Mean and median amplitude spectra for data set Hwz.05, calculated from normalized trace spectra. The plotted spectra have been normalized such that their integrated powers equal one.

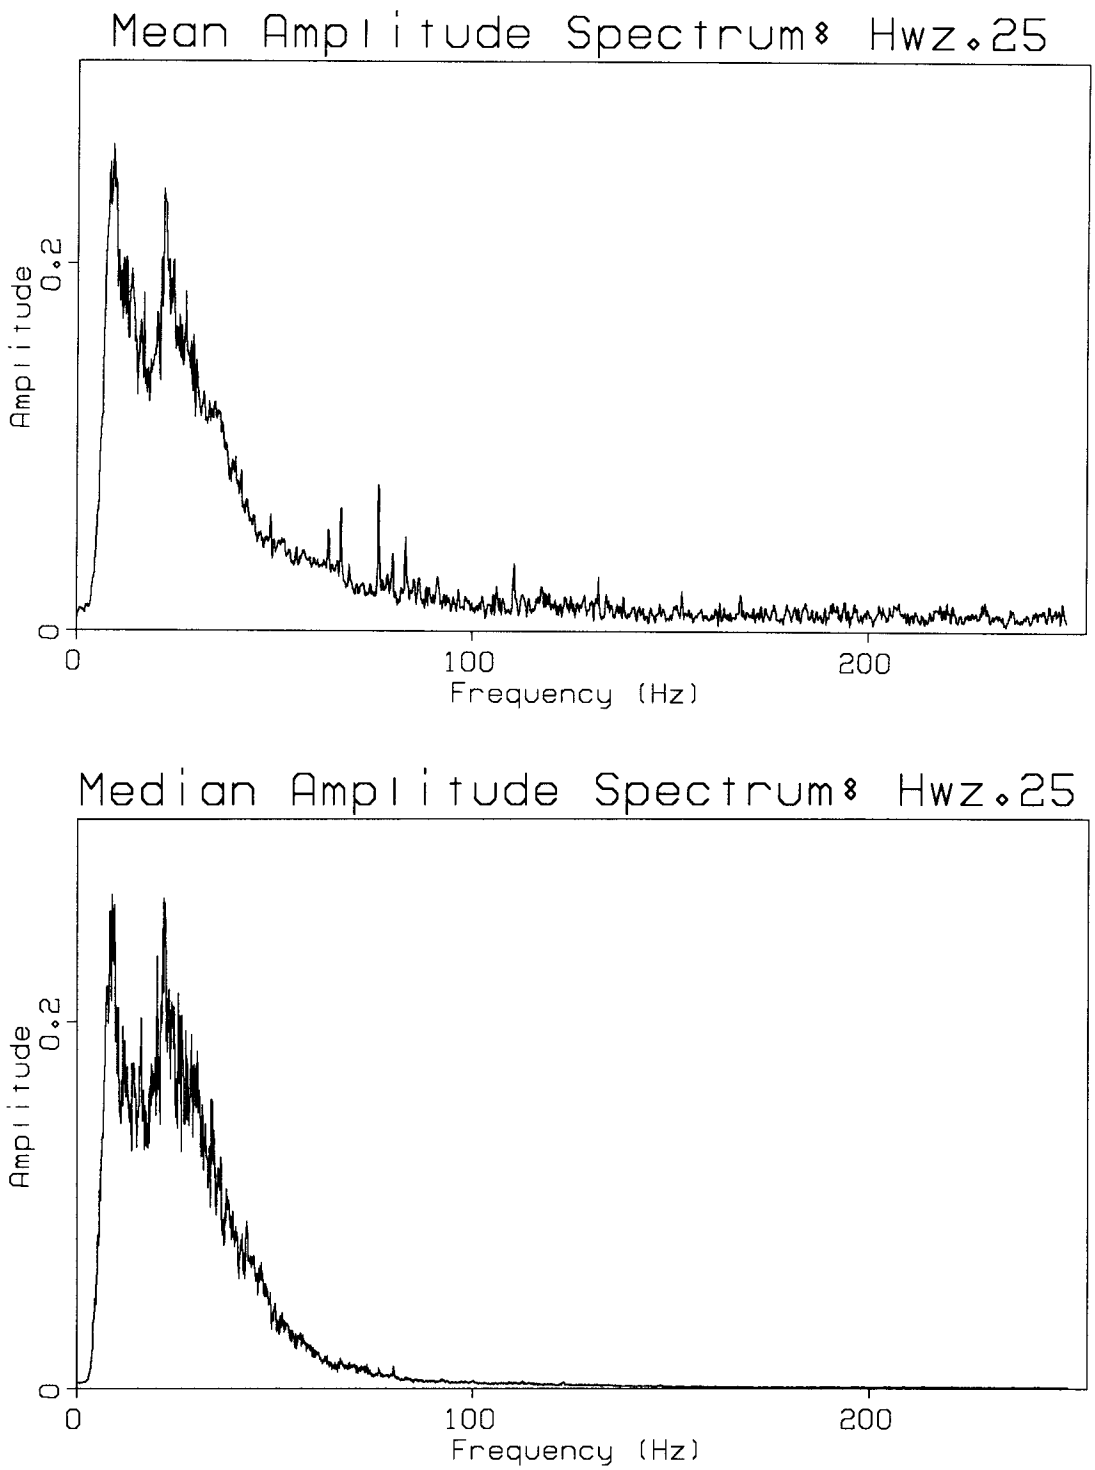

FIG. 8. Mean and median amplitude spectra for data set Hwz.25, calculated from normalized trace spectra. The plotted spectra have been normalized such that their integrated powers equal one.

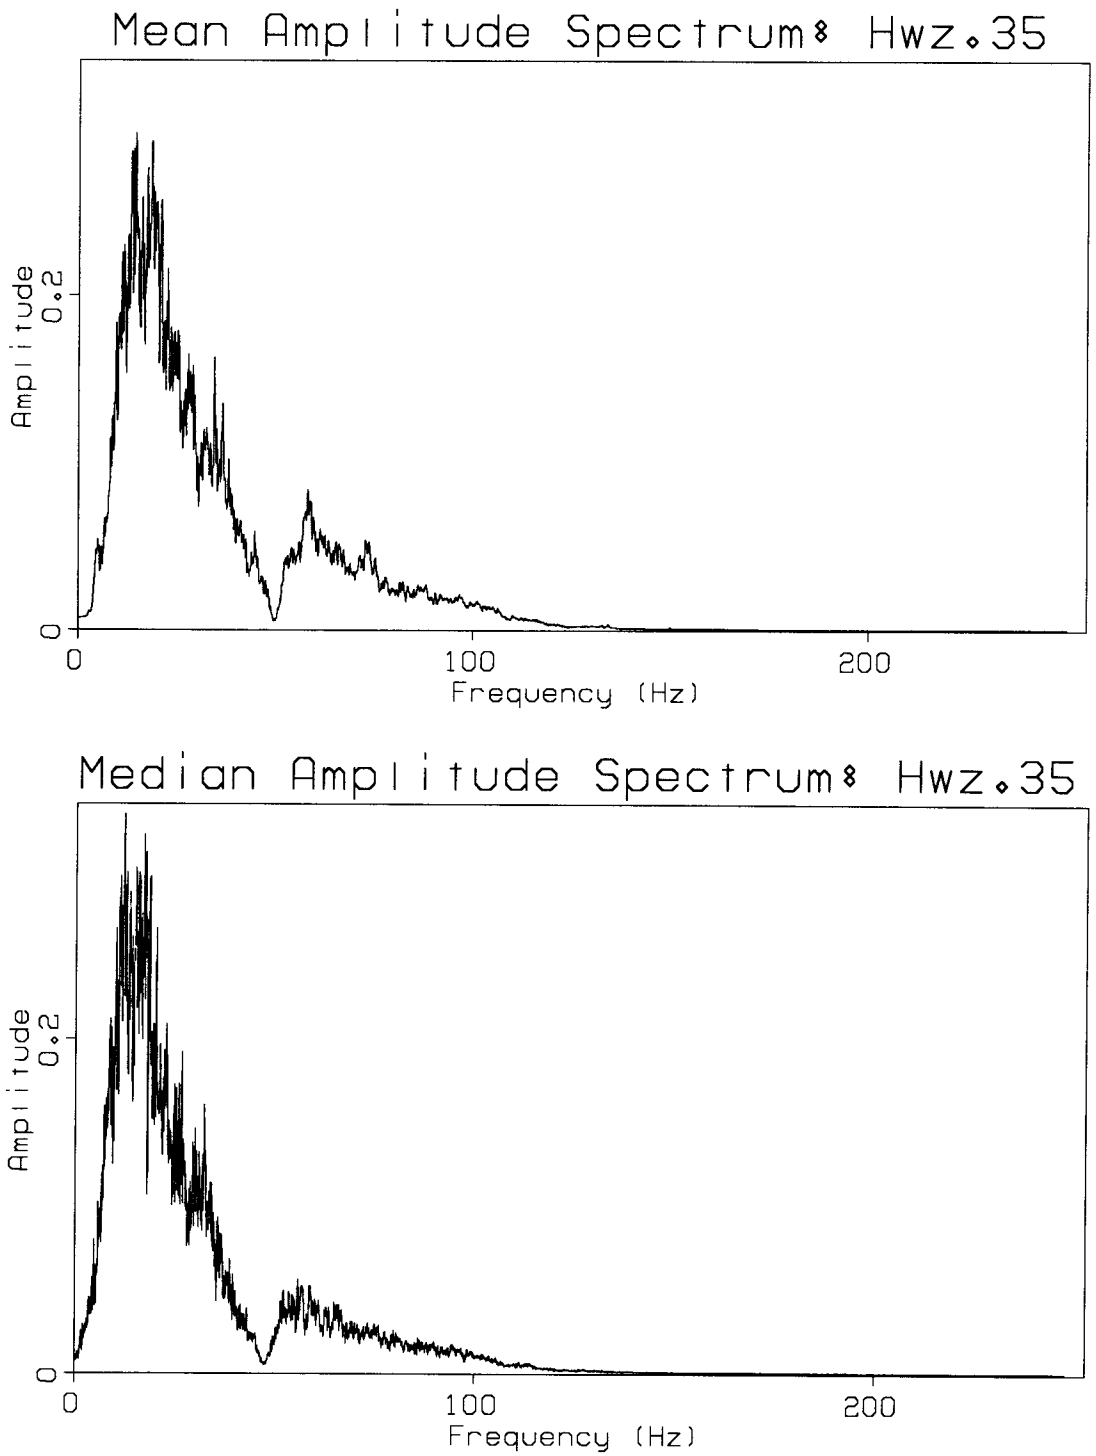

FIG. 9. Mean and median amplitude spectra for data set Hwz.35, calculated from normalized trace spectra. The plotted spectra have been normalized such that their integrated powers equal one.

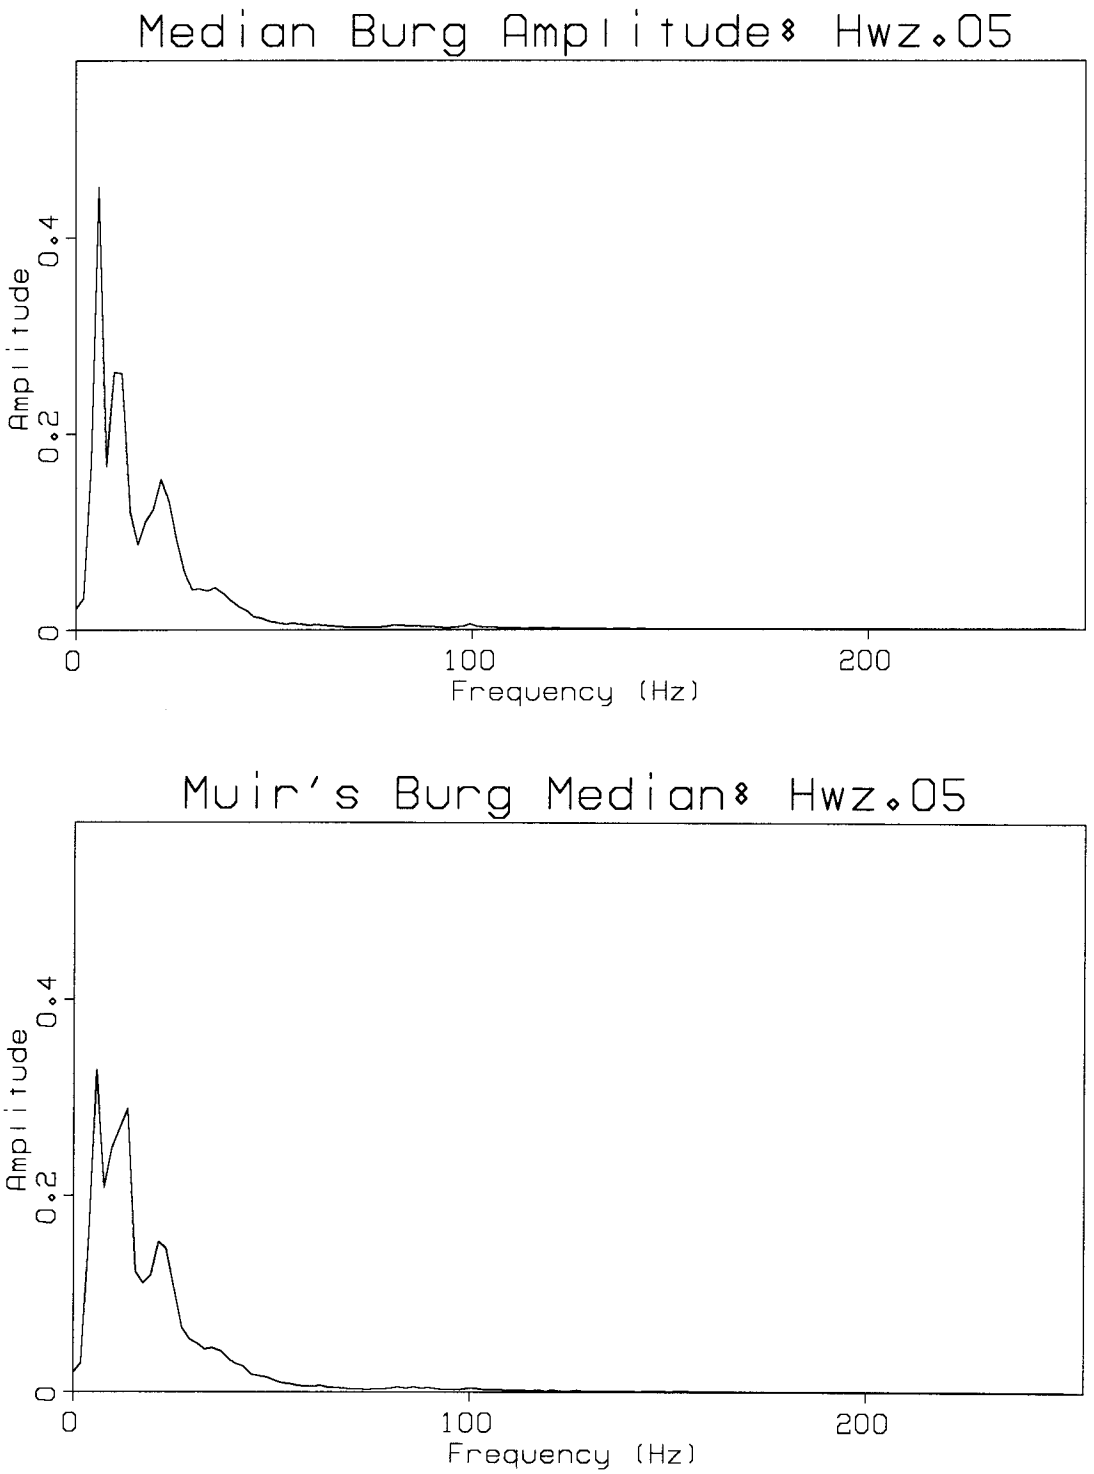

FIG. 10. Two different median amplitude spectra calculated using Burg's method for data set Hwz.05. The median of the Burg amplitude spectra calculated trace by trace is on top. The lower spectrum was calculated using Muir's interleaved median/Burg algorithm method. The plotted spectra have been normalized such that their integrated powers equal one.

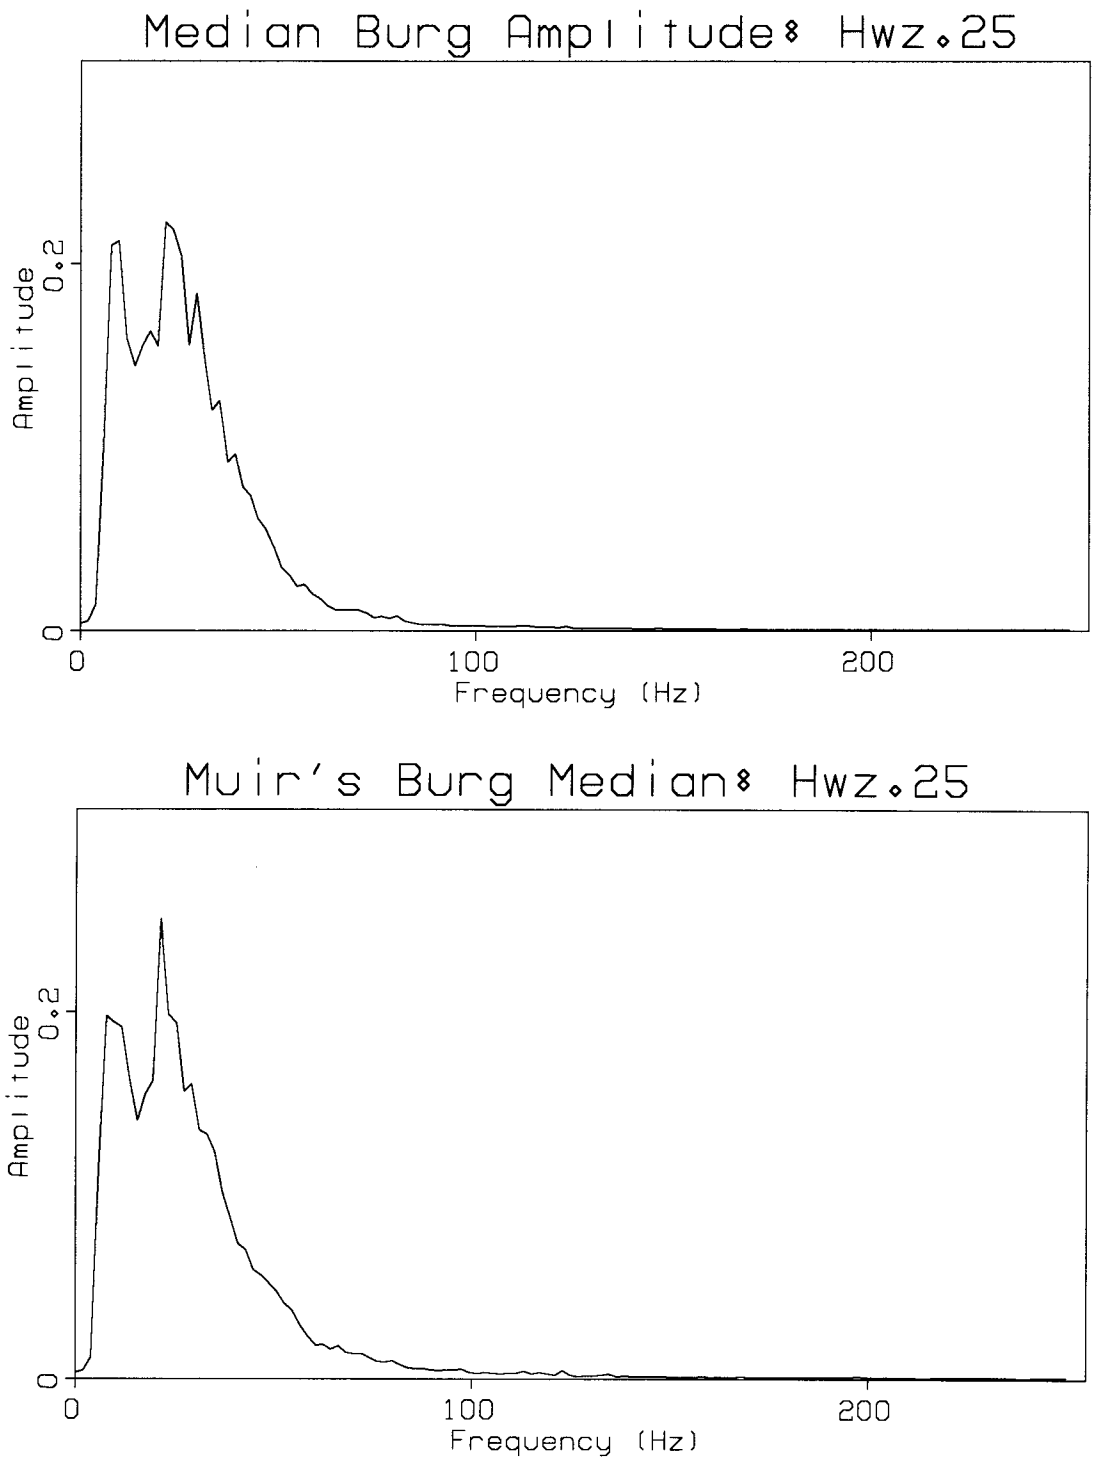

FIG. 11. Two different median amplitude spectra calculated using Burg's method for data set Hwz.25. The median of the Burg amplitude spectra calculated trace by trace is on top. The lower spectrum was calculated using Muir's interleaved median/Burg algorithm method. The plotted spectra have been normalized such that their integrated powers equal one.

Muir's Burg median

A second, more elegant method for combining Burg algorithm generated trace spectra has been suggested by Francis Muir. Recognizing that the motivation for all these methods is the determination of a common spectrum shared by a group of related traces, he has pointed out that instead of constructing separate prediction error filters for each trace independently, an ideal algorithm should construct a single prediction error filter for all traces simultaneously.

The recursive nature of Burg's algorithm makes this construction possible. Burg's method produces minimum phase, minimum-squared-error prediction error filters of increasing length by invoking the Levinson recursion. Each step of the Levinson recursion combines a prediction error filter of length n with a parameter called a reflection coefficient (or C for short) and returns a prediction error filter of length n+1. This filter is then used in Burg's algorithm to construct the C required for determination of the next longer prediction error filter. For the filter produced by the Levinson recursion to be minimum phase the C's must all lie between plus and minus one; the Burg algorithm guarantees that every C has this property.

Francis Muir's method proceeds by sequentially applying Burg's algorithm to every trace in a data set at each recursion, thus generating a vector of C's representing the entire group. It then selects a median C from this vector and plugs it into the Levinson recursion, producing a single prediction error filter for common use on all the traces in the subsequent step. Because every C in the domain of the median falls between plus and minus one, this median C also falls between plus and minus one, and the new prediction error filter is again guaranteed to be minimum phase. The procedure is repeated until the filter achieves some predetermined required length.

The results from application of this method to data sets Hwz.05 and Hwz.25 are illustrated in the lower halves of Figures 10 and 11. The process produces a median estimate of the spectrum very similar to the other median methods discussed above. As with the other Burg method, comparison with the Fourier transform/window methods is difficult given the short lengths of the prediction error filters.

CONCLUSIONS

The first section of this paper produced practical examples showing the superiority of robust median methods over mean methods in generating single spectral estimators from large numbers of related trace spectra. Distinction between the median methods on similar grounds is more difficult. For this purpose prediction error filters of several

lengths generated both by the straightforward median approach of the first section and by Francis Muir's interleaved median/Burg algorithm method of the second section were used to deconvolve the original data sets. Unfortunately, no significant differences were immediately apparent between the resulting "deconvolved" shot profiles.

As far as computational efficiency, note that in calculating the C's Muir's method requires the same amount of computation as the first Burg method, where the full array of C's are generated for each trace. For large data sets Muir's method requires more disk I/O than the first method, since the former can work on one trace and then discard it, while the latter uses every trace at each step of the iteration. Whether the extra run time required is worth the better model remains an object of research. Since the fast Fourier transform methods were done in an array processor and the Burg methods were not, it is not clear which of these two classes is more computationally efficient.

ACKNOWLEDGMENTS

We would like to thank Francis Muir for many fruitful discussions on the subject, and for the use of his idea concerning interleaving a median step with the Burg algorithm. We would also like to thank Jon Clærbout for suggesting that we look at median spectral estimators.

REFERENCES

Burg, J.P., 1975, Maximum entropy spectral analysis: SEP-6.

Clærbout, J.F., 1976, Fundamentals of geophysical data processing with applications to petroleum prospecting: McGraw Hill, New York.