APPLICATION MANUAL DSP-015 11/76

DETECTING, ENHANCING OR ELIMINATING SIGNALS

IN NOISE

Scientific-Atlanta
Spectral Dynamics Division
United States: P.O. Box 23575 • San Diego, CA 92123-0575 • (619) 268-7100 • TWX 910-335-2022
Europe: Horton Manor • Stanwell Road • Horton, Slough, SL3 9PA, England • Telephone: Colnbrook 3211 • Telex: 849406

Detecting, Enhancing or Eliminating Signals in Noise

SIGNAL-TO-NOISE ENHANCEMENT

Reasons for improving signal clarity

One of the goals of many signal processing tasks is to improve a desired signal's clarity with respect to background noise. The reasons for wanting to improve signal-to-noise ratios vary widely. For example, they might include establishing the existence of a signal; determining the harmonic structure of a signal when only its fundamental component is known; improving speech intelligibility; determining the precise frequency of a signal when by nature it is only barely visible in the spectrum; or removing extremely high level signals of a precisely known frequency, but of an unwanted nature, such as an unbalance component, hum, line frequency component, etc.

Signal enhancement techniques can be applied in the time and frequency domains. To compare the effectiveness of various signal-to-noise enhancement techniques, a common test signal will be used for each approach. The equipment setup for this example is below.

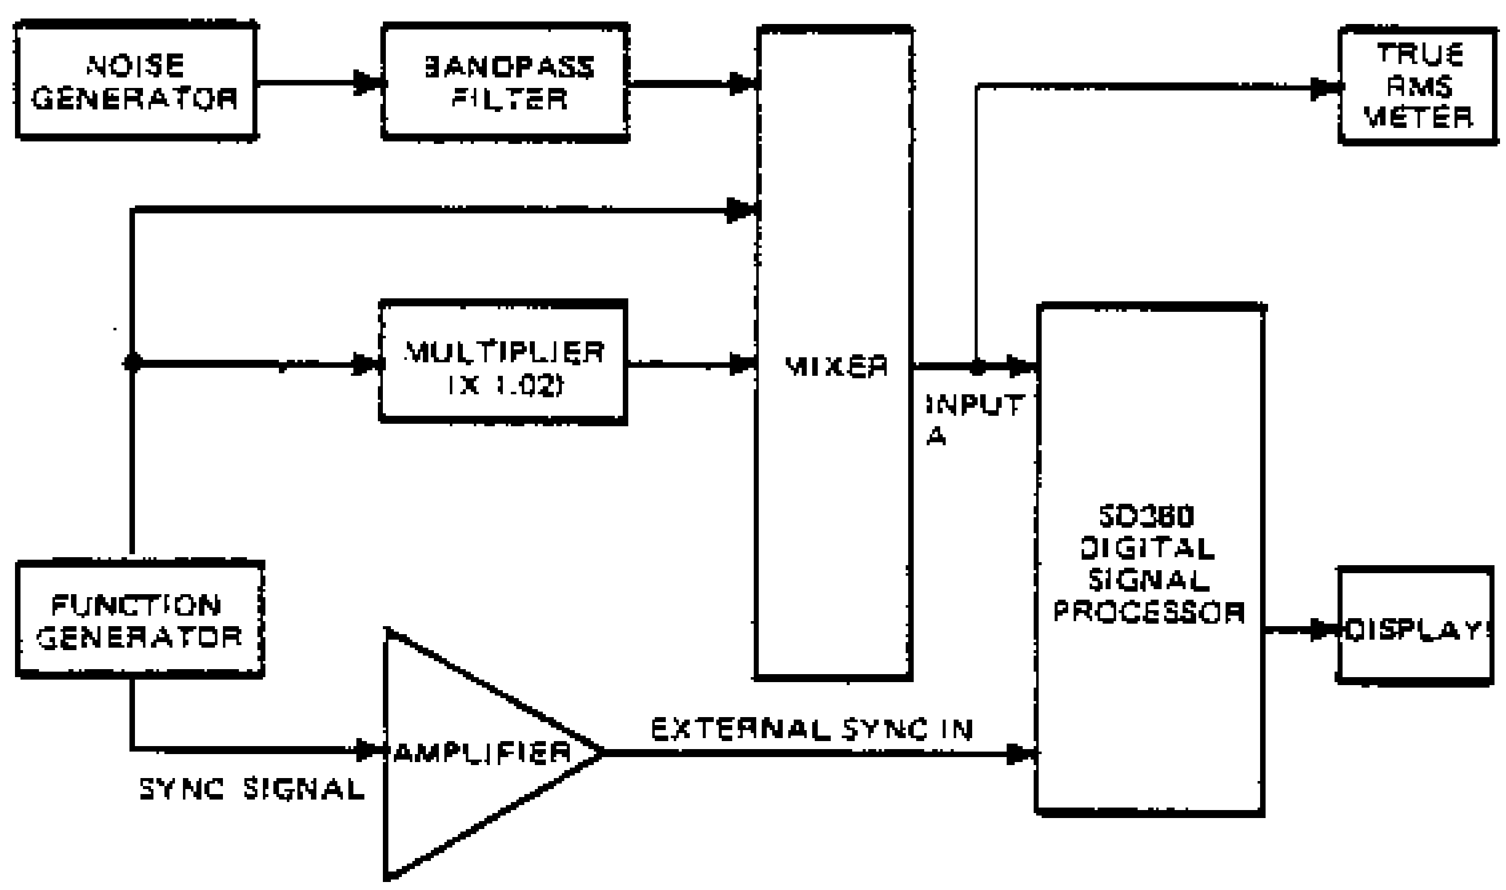

The test signal is several periodic waveforms mixed with a shaped noise component. The time waveform of the signal is below. Regardless of the type of signal processing which is to be performed, it is good practice to first observe the raw data signal in the time domain. For certain types of signals, which include transient characteristics or occasional extremely high level values, observation of the time waveform will often reveal characteristics which will not be at all obvious after some type of signal processing or averaging is performed.

RAW DATA WAVEFORM

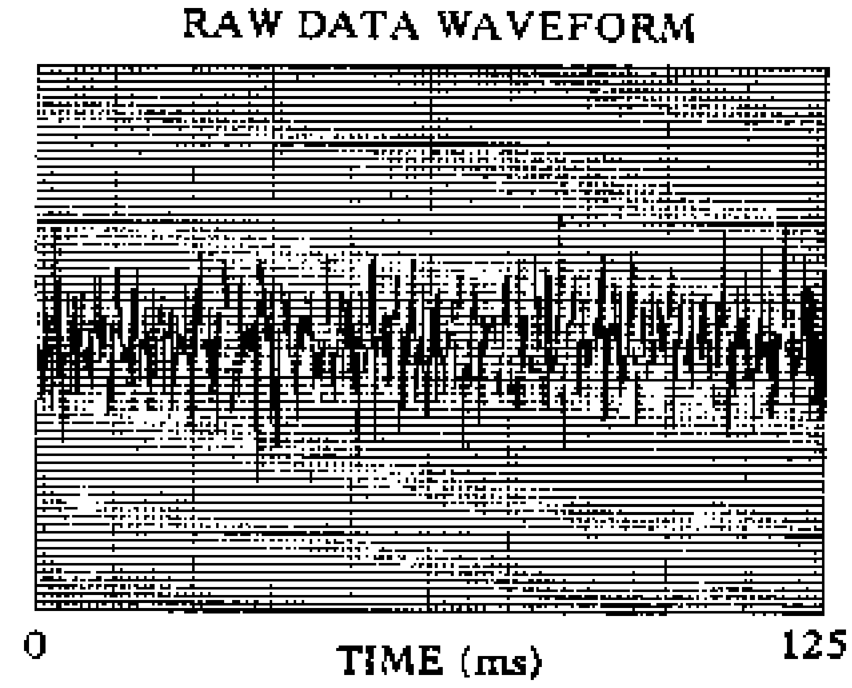

SPECTRUM ANALYSIS

Averaging versus analysis bandwidth reduction

For many applications, the process which is most meaningful in terms of relating signals back to physical occurrences is that of performing a spectrum analysis. In addition to giving information about the frequencies of occurrence in the signal being measured, a simple spectrum analysis gives some signal-to-noise enhancement, which is related to the ratio of the signal bandwidth to the analysis bandwidth. If an analyzer is used which uses a sampling process, such as a real time analyzer (RTA) or fast Fourier transform (FFT) processor, the signal to be analyzed is normally passed through a low-pass or aliasing filter. In this case, the signal bandwidth corresponds to the bandwidth of the aliasing filter. The analysis bandwidth is determined by the number of lines (number of data points/sampling index) in the analyzer or processor. In general, the signal-to-noise enhancement produced under these conditions is

$$S/N_{enh} = \sqrt{\frac{BW_{signal}}{BW_{processor}}}$$

In the composite signal being used here, the levels of the various signals included are 0.088V rms for the 287 Hz periodic signal, 0.084V rms for the square wave of 299 Hz, and 0.74V rms for the shaped noise signal. The result of performing a single-length transform on the test signal is seen below. This is the equivalent of providing a "real time" snapshot of the spectrum present in the waveform over a frequency range of 0 to 4 kHz. Note that in this spectrum display, a broad, overall shape can be seen, but with the possible exception of the spectrum component at 290 Hz, no evidence of periodicities can clearly be determined.

If spectrum averaging is now performed, the existence of the periodic tones in the noise background is more easily verified. This is seen in the figure below which is the average of 16 ensembles of the test signal. Several of the tones are now visible.

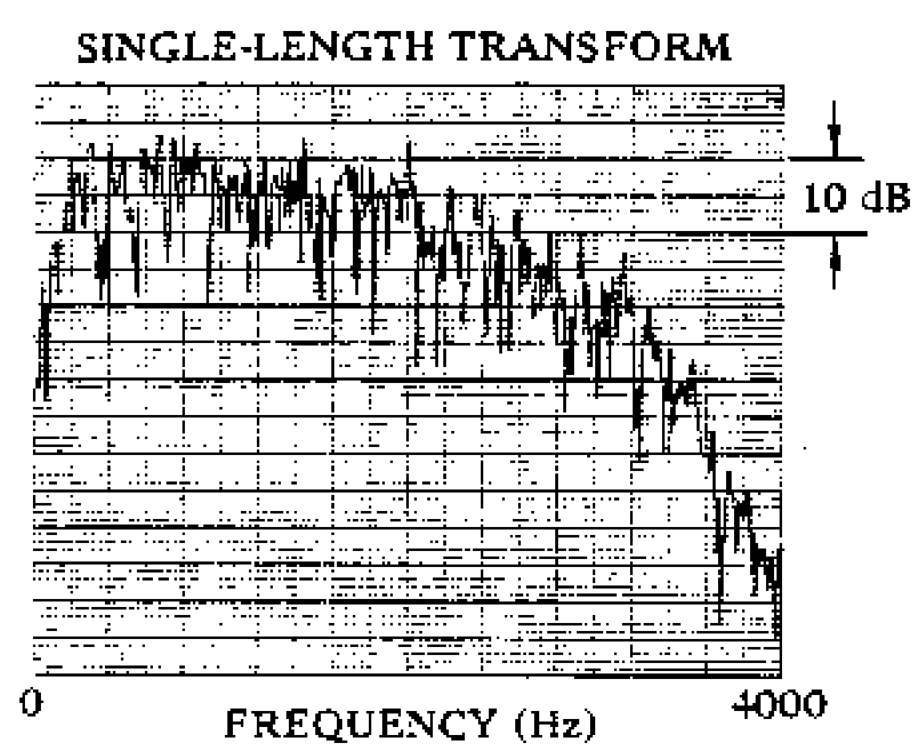

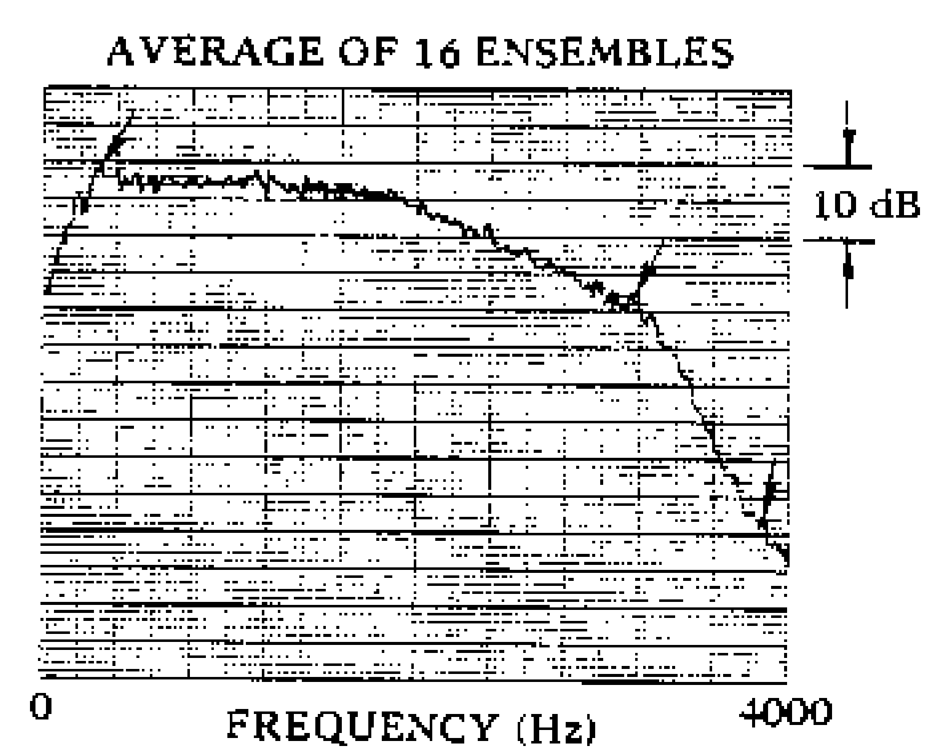

If additional spectrum averaging is performed, what improvements can be expected? The next figures show the same signal spectrum after 256 ensemble averages. Note that the increased averaging time has smoothed out statistical variations of the noise characteristic, and in this sense, has caused the periodicities to stand out somewhat more clearly than the previous spectrum of 16 averages. However, from the previous definition of signal-to-noise enhancement in the frequency domain, it is clear that enhancement is directly related to the ratio of bandwidths and is not directly related to averaging-time considerations. It is evident that additional signal-to-noise enhancement in the frequency domain is only possible by reducing analysis bandwidth at the expense of increasing analysis time.

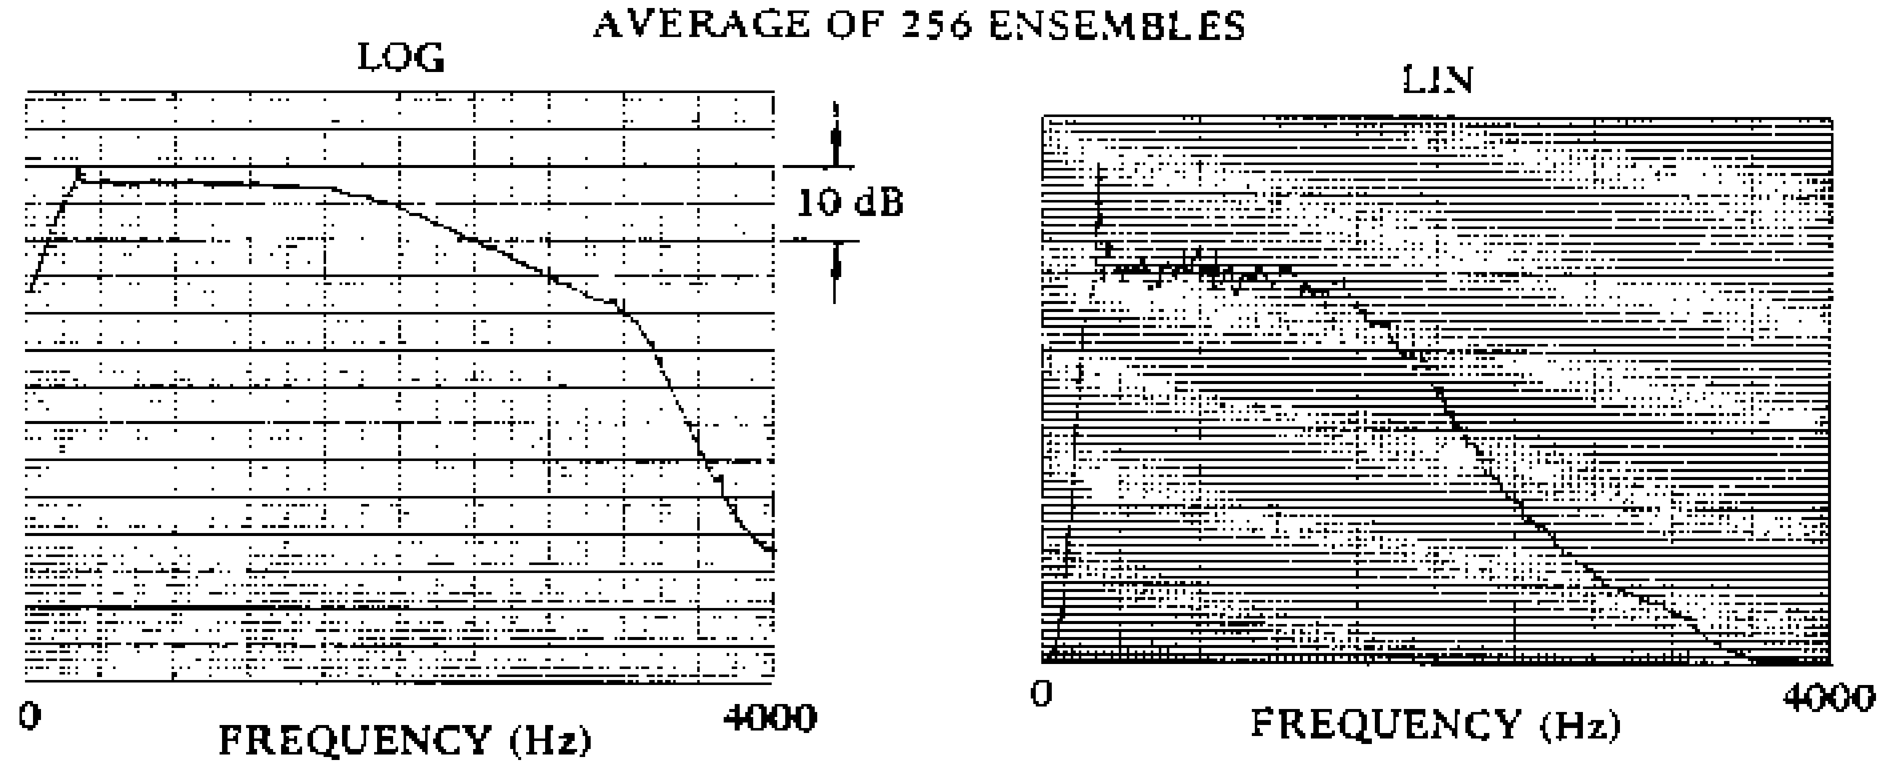

Effect of transform length

Increased enhancement is also achieved by increasing the number of spectrum lines developed within any given frequency interval. In an FFT processor, this is equivalent to increasing the length of the transform. The previous figures showed spectrum analyses produced with 512-line resolution. Processing the same signal over the same frequency range with twice the resolution (1024 lines) gives the figure below.

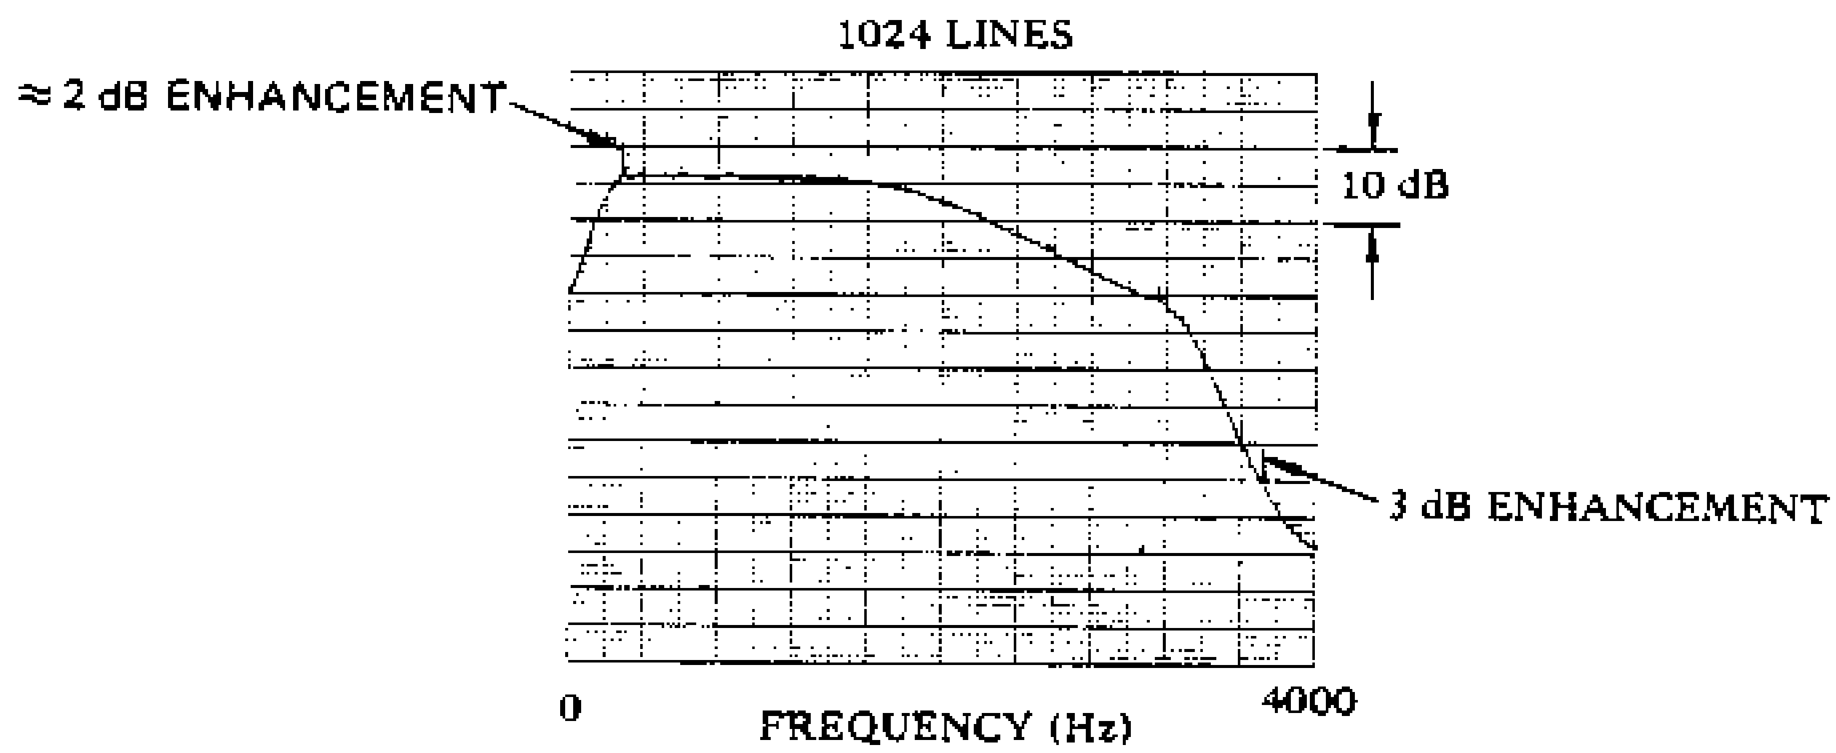

A maximum signal-to-noise increase of 3 dB can be achieved, which is dependent on the relationship between the exact frequencies of any periodicities present and the precise center of the analysis windows created by the FFT process. In this figure, it is obvious that there is the full 3 dB enhancement at the harmonic shown at 3730 Hz. Whereas, in the region of the fundamental, there is approximately 2 dB enhancement.

FREQUENCY TRANSLATION (ZOOM)

Analysis filter bandwidth reduction

Since signal-to-noise enhancement in the frequency domain is seen to vary in direct proportion to analysis filter bandwidth, an obvious way to improve signal enhancement is to provide the narrowest possible analysis window over a given selected frequency range. This can be done by either increasing the number of analysis lines, usually at the expense of increased memory size and increased analysis time, or by using a frequency *zoom* technique, where the entire resolution of a given analyzer or processor (such as, 400, 500, 800 or 1,000 lines) is concentrated in a selected frequency interval. As the effective filtering or analysis window is made narrower, the total processing time required becomes longer. One advantage of this approach is that virtually any signal-to-noise improvement is possible if a sufficiently narrow window can be created. A potential disadvantage is that only one frequency region at a time can thus be examined and the narrower the *zoom* window becomes, the more discrete step intervals must be examined to cover a specified frequency range.

In the standard 500-line spectrum display of the example there is an indication of a primary periodicity in the region of 290 Hz with no evidence of any other harmonics or periodicities. However, having some prior knowledge of the signals being investigated, it is obvious that there must be third harmonic information out in the region of 900 Hz. In the figure below, a frequency *zoom* was used (100 Hz window centered at 900 Hz). The existence of two tones is now seen clearly at 863 and 897 Hz. Whereas, these harmonics were initially 10 dB below the background noise level after processing, the signal-to-noise ratio has now been improved to nominally +3 dB after *zoom*. The increased signal-to-noise ratio is seen even more clearly on the linear amplitude display shown below.

RANGE TRANSLATION

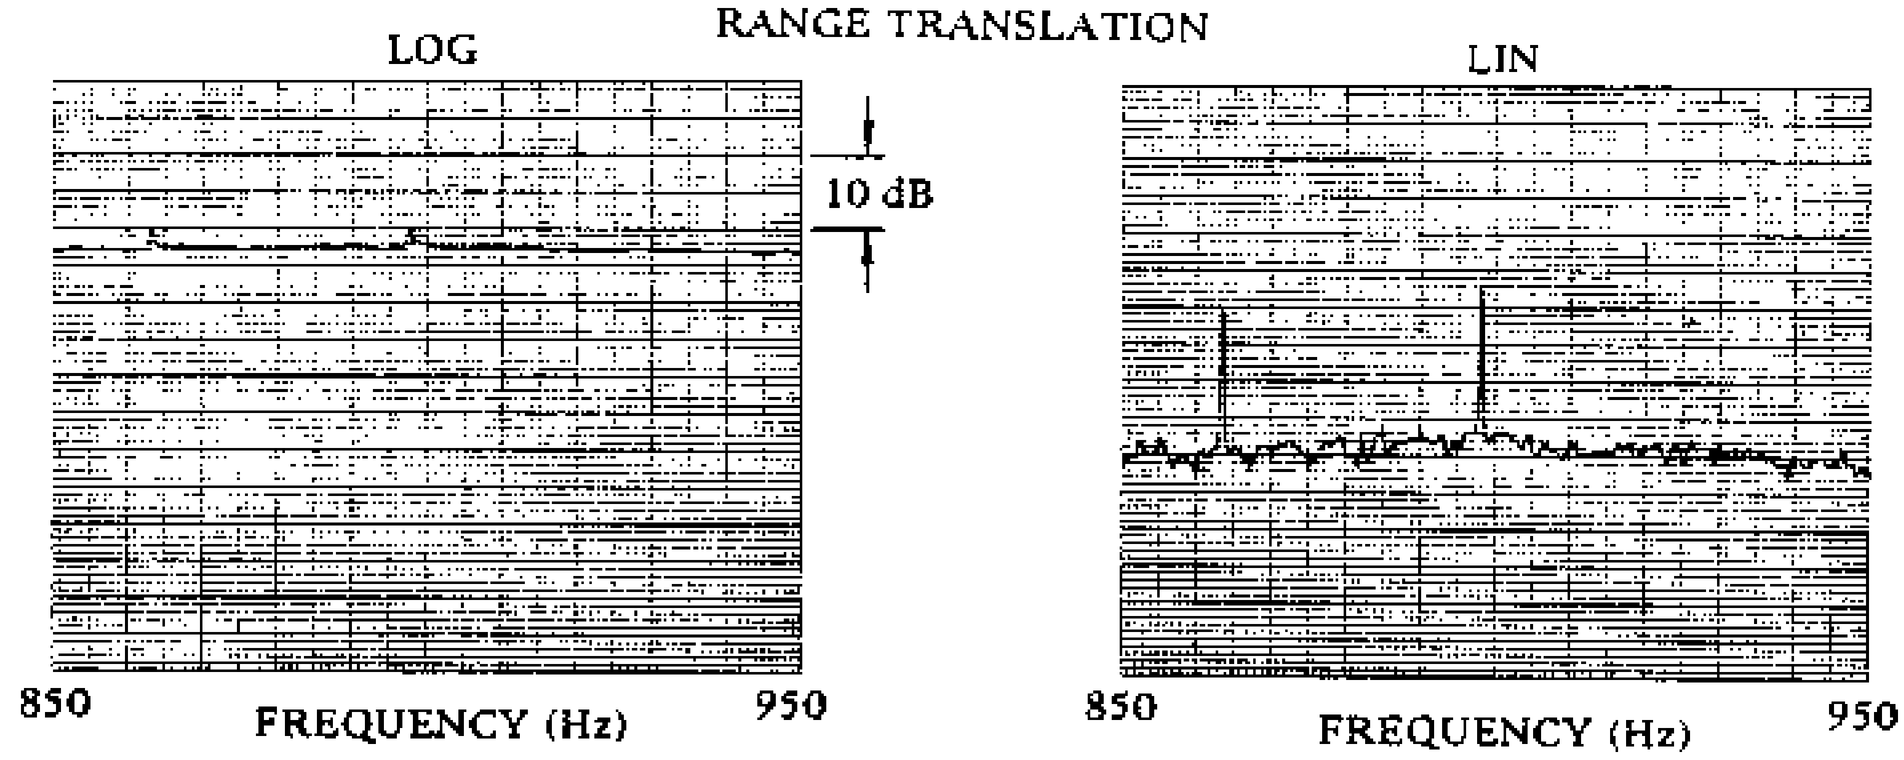

Although a 100 Hz translated range is displayed, in this case the total processing range during *zoom* was 140 Hz wide. Since the original frequency range processed was from 0 to 4000 Hz, this represents a bandwidth reduction, by using *zoom*, of approximately 28:1. Theoretically, the maximum signal-to-noise enhancement which should be achievable with this type bandwidth reduction is 14 dB. The apparent signal-to-noise improvement achieved in the given example was nominally 13 dB, or close to the theoretical optimum.

CORRELATION ANALYSIS

Effect of time difference on periodicities

Just as a spectrum analysis can give a direct determination of the amplitude and frequency content in a given signal, a correlation analysis can give a direct indication of how periodicities in a signal behave with respect to time difference. Formerly, correlation techniques were used quite extensively as a means for detecting signals in noise and "cleaning up" otherwise noisy signals, such as encountered in radar returns. There are actually many types of correlation function analysis, usually a variation of cross-correlation.

The simplest type of correlation is autocorrelation analysis in which a data signal is correlated against itself. The figure below shows the autocorrelation function of the test signal being analyzed.

AUTOCORRELATION FUNCTION

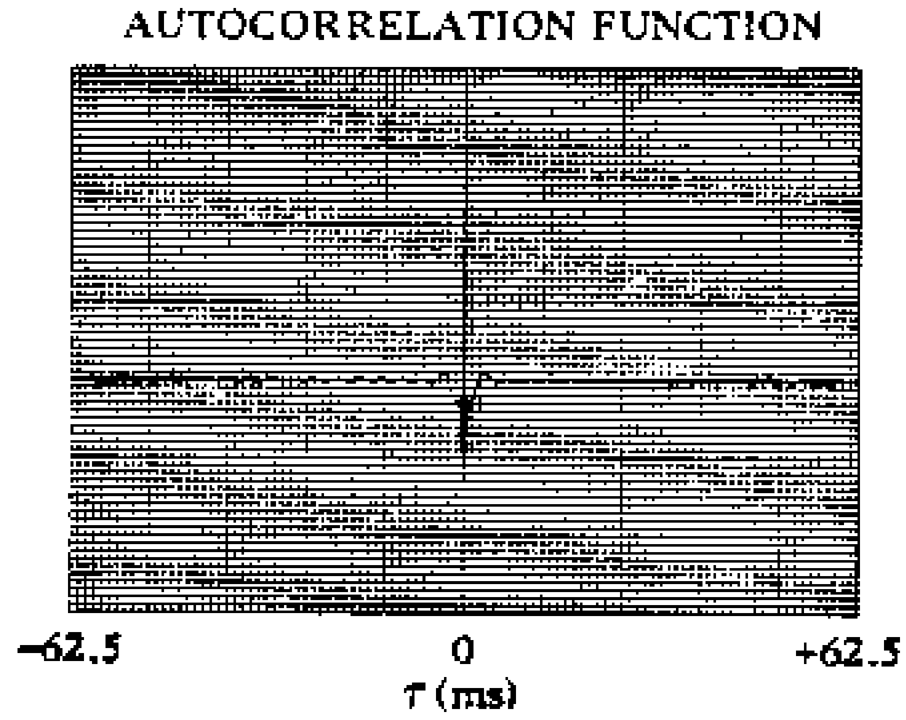

The display range shown here is from −62.5 to +62.5 milliseconds with zero time delay shown at the center of the graph. Note the symmetry of the correlation function, a characteristic of autocorrelation analysis. The primary feature of the display in this case is the obvious excellent correlation at zero time delay, but with no apparent periodicities seen. The steepness of the decay at the center of the display, falling through zero in less than 1 millisecond, indicates a fairly broad input spectrum. The absence of any continuous functions which appear periodic as time delay is increased in either direction indicates that within some nominal dynamic range, no periodicities appear to exist. However, since it is known that periodic signals do, in fact, exist in this data, it is obvious that a different type of correlation processing will have to be performed if effective signal-to-noise enhancement is to be achieved.

CROSS-CORRELATION WITH A PERIODIC REFERENCE

Amplitude-sensitive recovery of data signals

One of the theorems of cross-correlation analysis is that the only expected result from an averaged cross-correlation analysis is something which was common to both signals being correlated. If one of the signals being used is a periodic function free from noise, then the resulting cross-correlogram must either be a periodic signal free from noise or a function which averages to zero. If the reference signal being correlated against a complex data signal is a sine wave, the resulting correlation function will be a sinusoidal type wave of some arbitrary phase at the origin, if the frequency of the reference sine wave is contained within the complex data signal.

If another type of periodic waveform, such as a square wave, is used as the reference input and a complex periodic waveform of the same frequency is present in the data signal under analysis, it is necessary to have a good understanding of correlation functions to be able to properly interpret the resulting periodic correlogram. The data signal of the example consisted of three components; a square wave at 287 Hz, another at 299 Hz and a shaped broadband noise component. The figure below shows the result of correlating the data signal with a reference sine wave of 287 Hz. Note that a clean periodicity is recovered and that precisely 35 cycles of the waveform are seen. In this display, the average time difference between peaks is 3.9 ms, indicating a periodicity of 287 Hz. The same data signal was then correlated against a sine wave of 299 Hz. The resulting correlogram is seen below.

REFERENCE OF 287 Hz

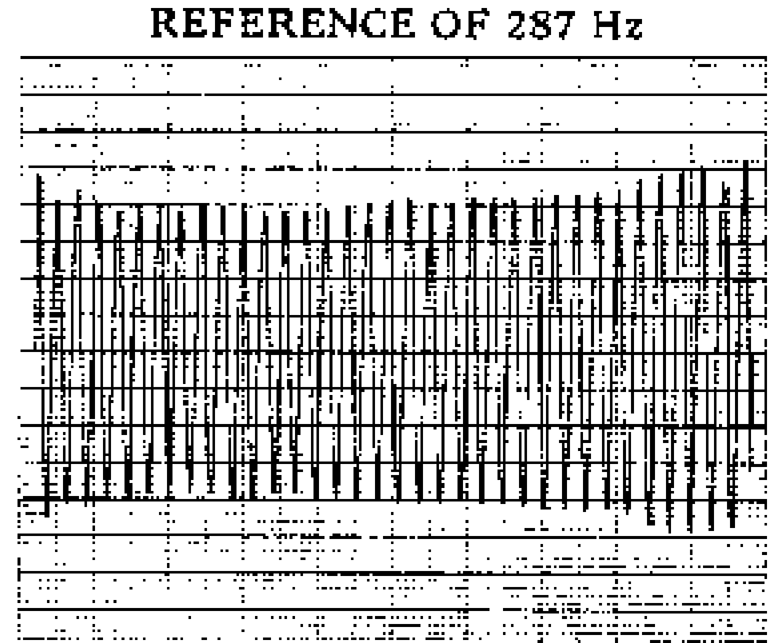

REFERENCE OF 299 Hz

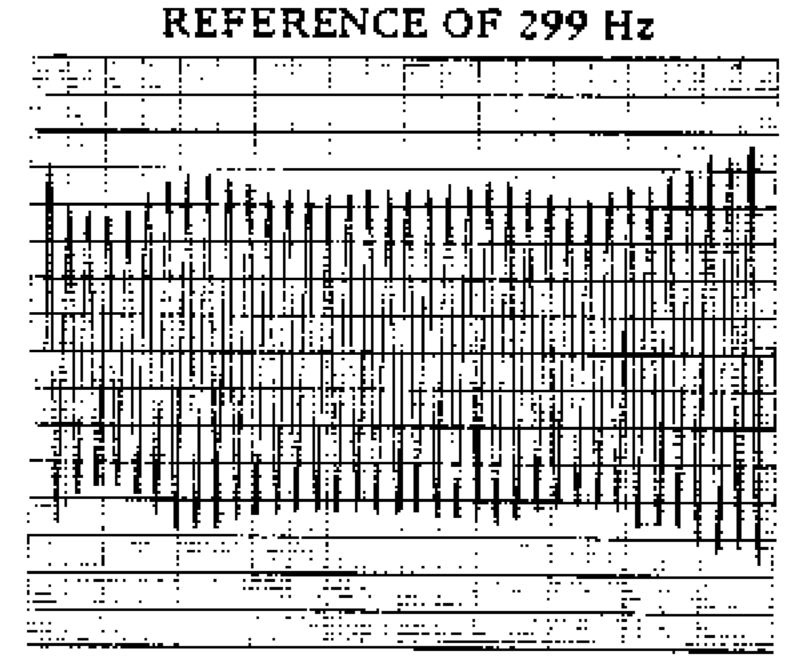

Note that in this case exactly 36 periods of the waveform are recovered, with a time difference between peaks of 3.34 ms. Other sinusoidal frequencies corresponding to harmonics of the fundamental rate can also be used in this application, resulting in an amplitude-sensitive recovery of any data signals occurring at precisely the harmonic frequency. Note that if a square wave were used as the reference input, the recovered periodicities would not be sinusoidal, but rather represent the correlation function of a square wave since there was, in fact, a square wave at the two fundamental frequencies present in the data signal. The resulting correlogram would then appear triangular.

Periodicities as unwanted data components

There may also be cases when the true data signal of interest is actually the broadband random component and the periodicities represent unwanted data components. In this case, the periodicities can, in general, be removed by correlating against a broadband random reference signal. The figure below shows just this situation where the reference signal was the broadband noise prior to bandpass filtering. This was used as Input B to the processor with the composite data signal (representing a system output) applied to Input A. Note the complete absence of periodicities in the correlogram. Note, also, that the resulting correlogram appears as a time-reversed impulse response for the bandpass filter used. This is to be expected since the two data signals were reversed from their normal input sequence. Although in this case the application might be the removal of periodic components, the analysis also points out the ease with which system impulse response functions can be measured using broadband excitation even if periodic data is also present.

System impulse response functions

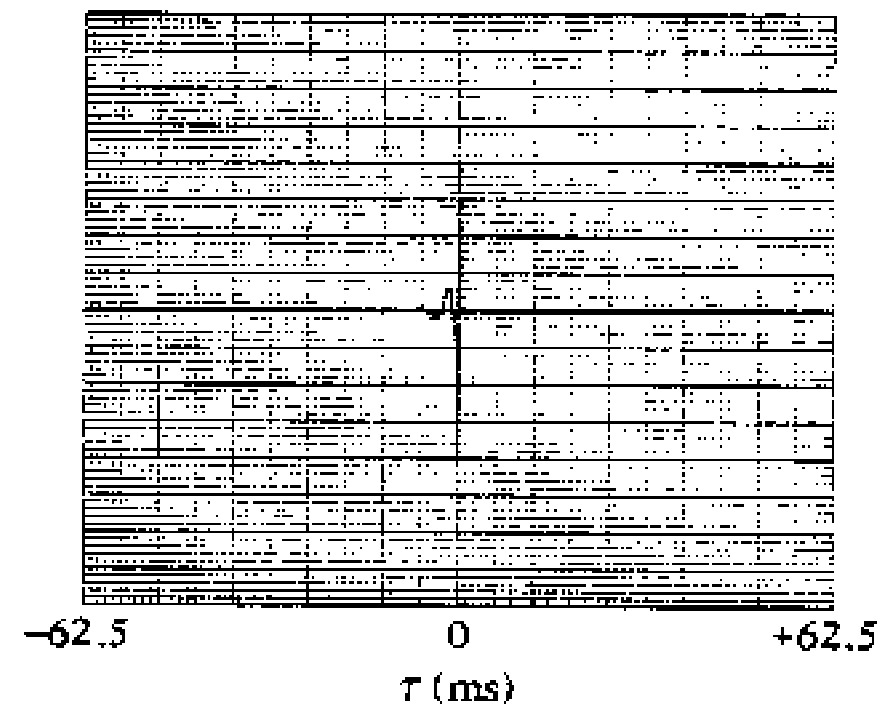

CROSS-SPECTRUM PROCESSING

Determining existence and amplitude of periodics

Cross-spectrum analysis by classic definition involves the processing of two data signals to determine their mutual properties in the frequency domain. Since the cross-spectrum is calculated by forming the product between two Fourier transforms, by its very nature it is an extremely sensitive measurement. Because of this, it can be used in several ways as a signal-enhancement tool.

Where it is of primary interest to determine the existence, relative amplitude and, possibly, exact amplitude of periodic signals surrounded by noise, the cross-spectrum can be formed between a data signal and itself. This is equivalent to calculating the power spectrum, but will often give an increased sensitivity to periodic tones when compared with a standard voltage spectrum analysis. The squaring operation of the power spectrum tends to "multiply up" periodic signals which might be slightly higher than the background noise (after spectrum processing) and "multiply down" normally low level noise components. Usually, up to 3 dB signal-to-noise improvement can be expected by using this multiplication or squaring technique. However, where specific tones of interest are involved, the cross-spectrum can be formed between a composite signal and a sinusoidal reference input in the same fashion that cross-correlation was previously performed. The figure below shows the cross-spectrum between the total test signal and a pure 287 Hz tone.

Specific tones of interest

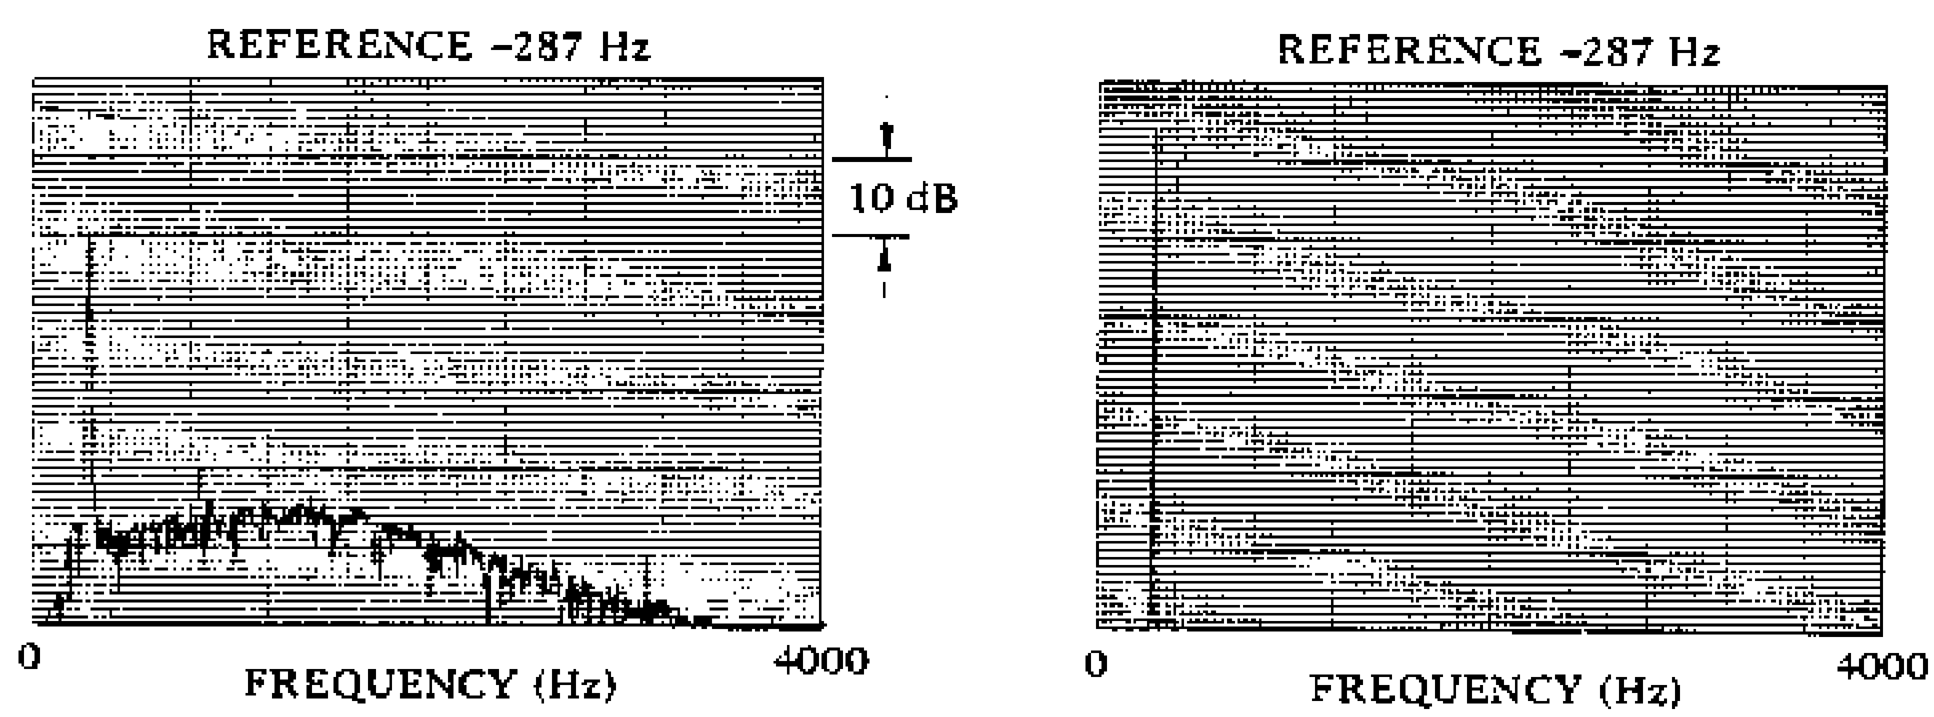

While in this case the reference signal was derived from the same 287 Hz source which is contained in the overall data, it is sometimes possible to establish a reference sinusoidal signal from a completely independent source based partially on the stability of the signal that is to be enhanced. The cross-spectrum result is shown here in both linear and log amplitude. Note in this figure and in all subsequent spectrum figures how much more clearly signals of interest are seen using linear spectrum amplitudes.

A similar test was then performed using as a reference input a fixed amplitude sinusoid at 299 Hz, again, derived from the same source as the periodic signal at that frequency obtained in the overall data signal. Note that in each case the effects of the other periodic component only 1.5 cells (1.5Δf) away from the reference sinusoidal input has been completely eliminated. This shows that the cross-spectrum as a signal-to-noise enhancement tool is extremely selective. Note also that the level of the signal recovered at 287 Hz is approximately 6% higher than that recovered at 299 Hz which, again, is consistent with the individual levels of the signal components.

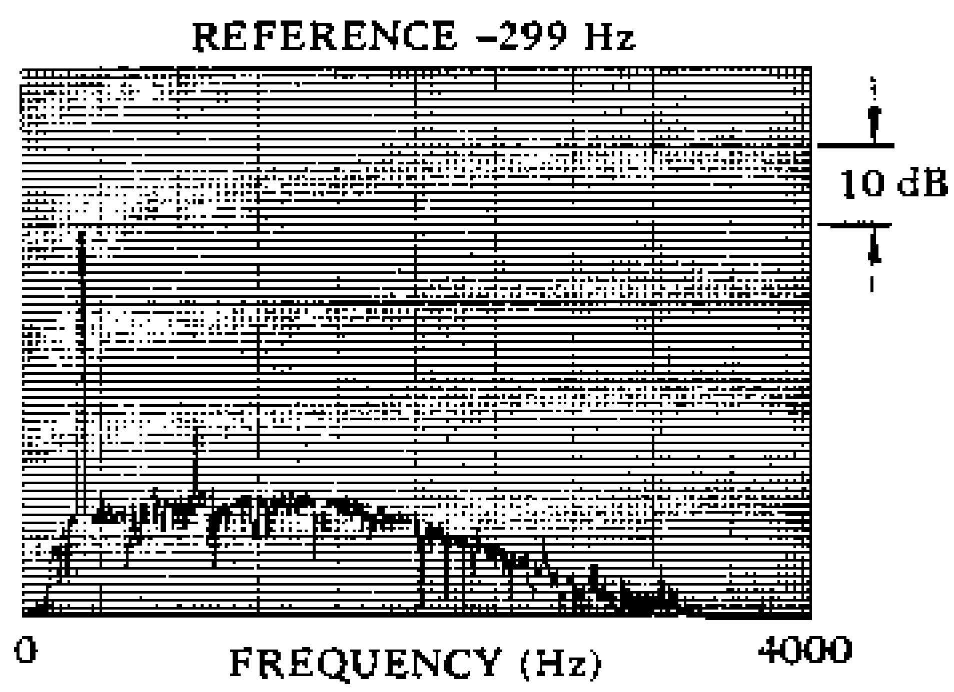

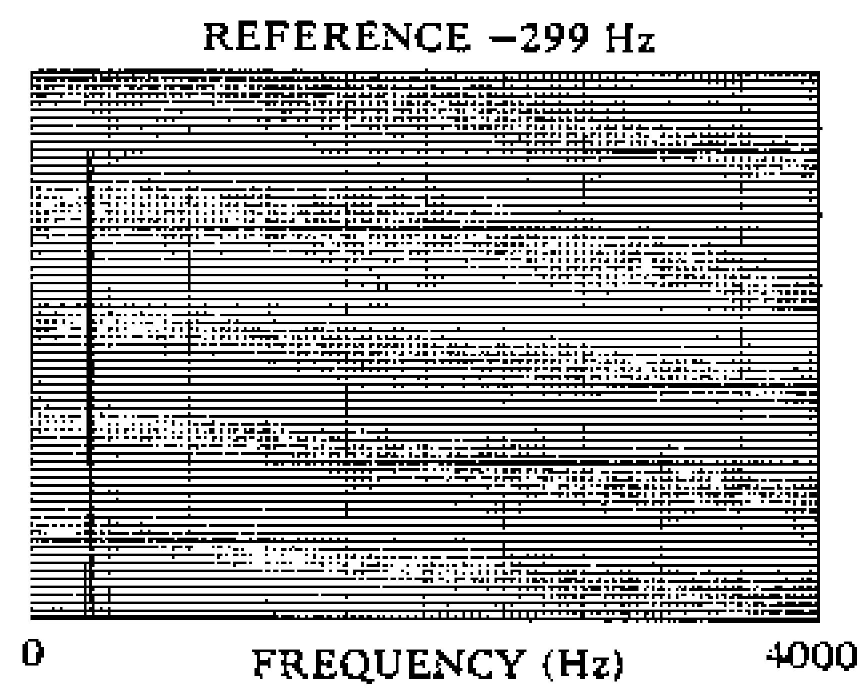

Unwanted periodics

Using random noise as the reference signal input produces the figures below. As in the case of cross-correlation with the reference random signal, all evidence of periodicities has disappeared. When using cross-spectrum processing in this manner as a signal-to-noise enhancement device, it must be kept in mind that the level of the resulting cross-spectrum is greatly affected by the level of each signal input being used in the cross-spectrum process.

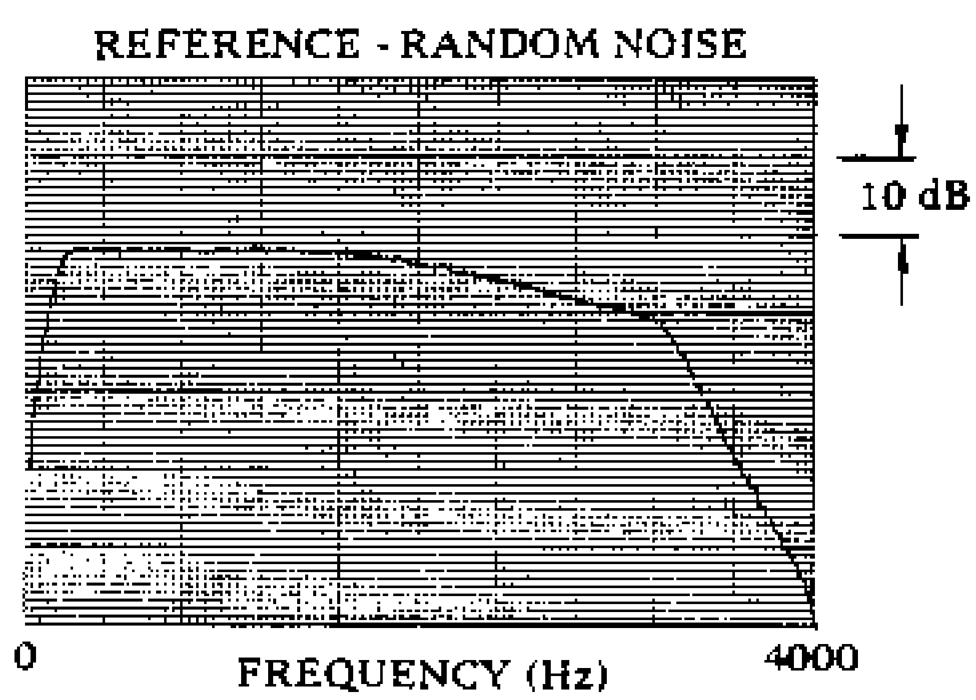

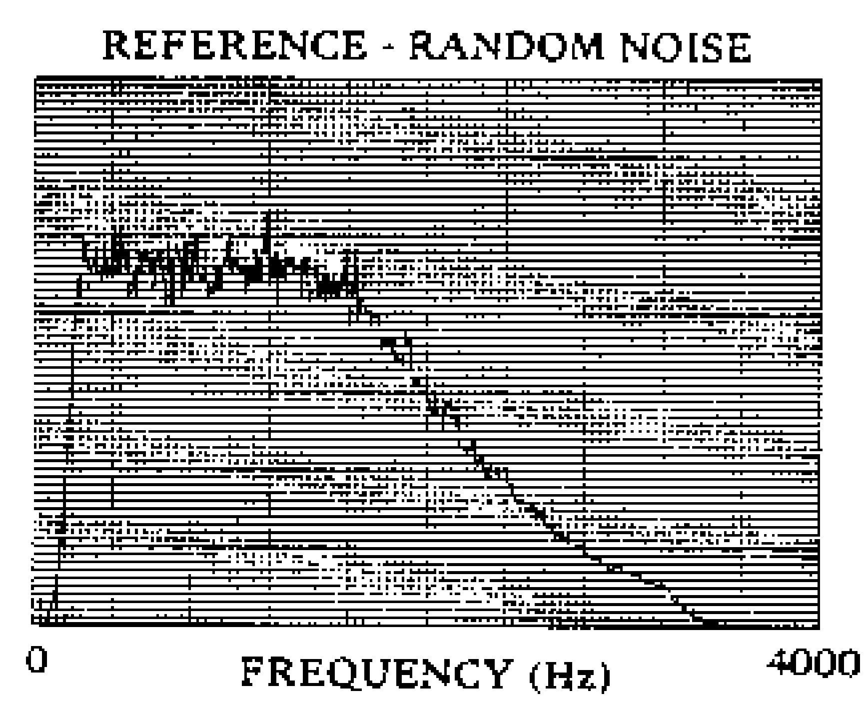

ENHANCEMENT THROUGH SIGNAL DELAY

Establishing the existence of periodicities with no prior knowledge and no reference signal

In one of the previous displays showing the correlation function of a signal which was predominantly random noise, it was noted that the correlogram consisted almost exclusively of the large spike at zero time delay. This is so because random noise correlates perfectly with itself at that one instant in time when there is no delay between the signal and its delayed replica; however, the correlation factor drops off very rapidly to zero for even a small amount of delay time. Since periodic signals have correlation factors which repeat over and over again, we have an important clue as to another potential signal-to-noise enhancement approach.

If a signal which consists of both periodic and random components were to be delayed with respect to itself and then processed for its correlation function, the preprocessing delay would tend to uncorrelate the random components of the signal, but only serve to introduce a phase shift to any periodic components. The amount of phase shift would depend on the sampling rate employed and the actual data frequency, but be of little consequence, since neither the nature of the periodicity, nor the time for one period to occur, would be changed. The use of this approach on the data signal of the example is seen in the figure which follows.

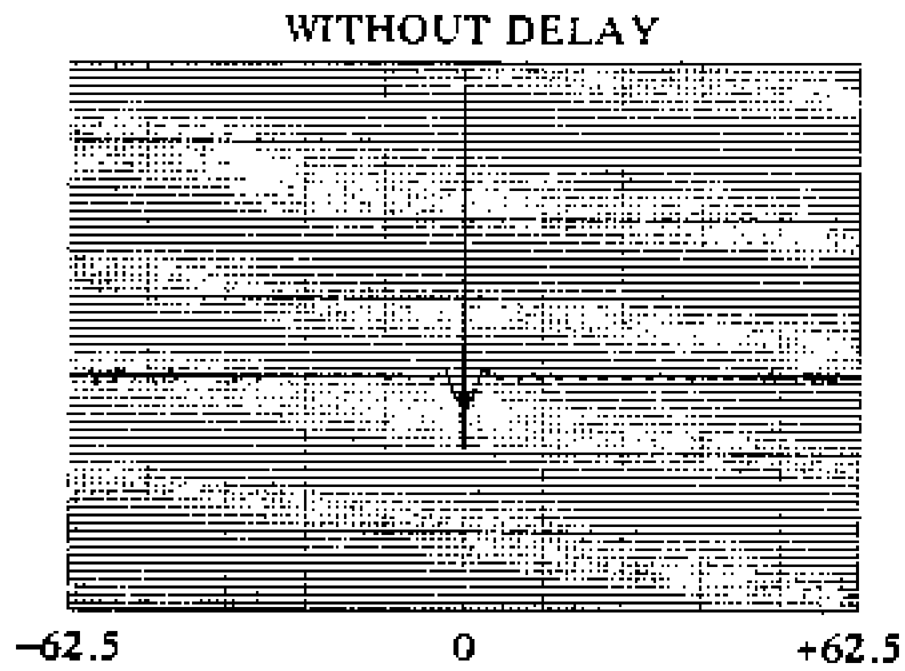

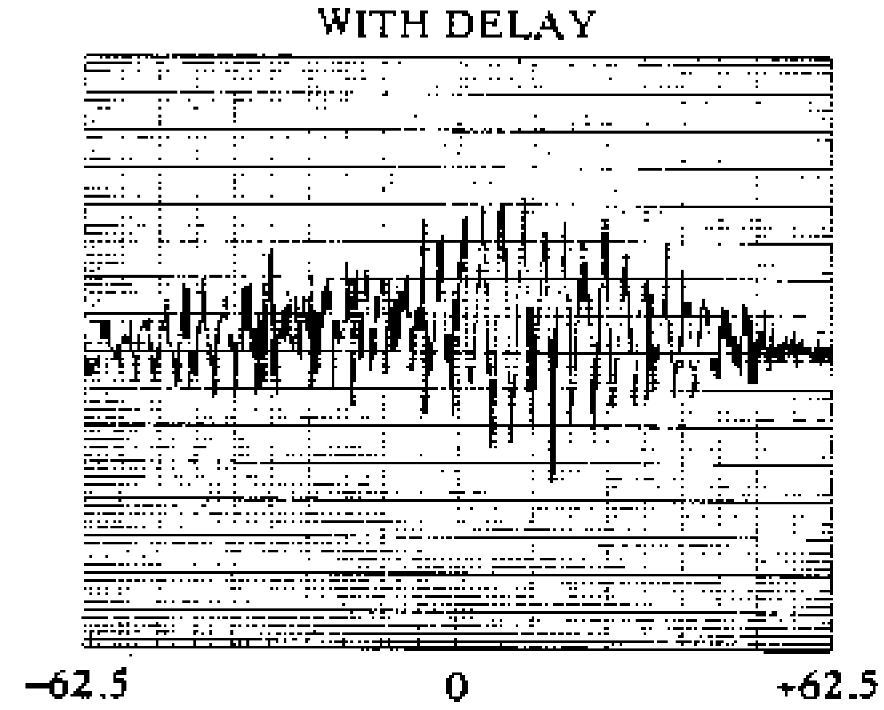

The first figure is a repetition of one previously seen and shows simply the correlation of the composite signal with itself with the characteristic large correlation factor at zero time delay. No evidence of periodicities is seen.

The same signal was then processed with respect to itself, but 62.5 ms of delay (1/2 memory period) was inserted in one path of the signal prior to computation. The effect this has on the resulting process is immediately obvious. The large data spike at zero time delay has completely disappeared and, now, the nature of the periodicity is seen clearly emerging.

Naturally, one of the primary purposes of this type processing is to establish the very existence of periodicities in the otherwise random signal and, clearly, this is done quite adequately. One of the key points to keep in mind for this type of processing is that it requires absolutely no prior knowledge about the data signals being processed. No reference signal is employed and the same single-channel input signal is merely processed through both input channels.

CROSS-SPECTRUM PROCESSING WITH DELAY

Enhancing periodicities by artificial delay

The same concept of introducing an artificial delay into one path of a signal process can also be used in the frequency domain using the cross-spectrum. Whereas, during correlation processing, the artificial delay in one signal path tended to uncorrelate random signals while enhancing periodic data, in the frequency domain the delay factor will again tend to enhance periodicities while giving cross products which average to very low values. The figures below show the results of the cross-spectrum process with and without signal delay.

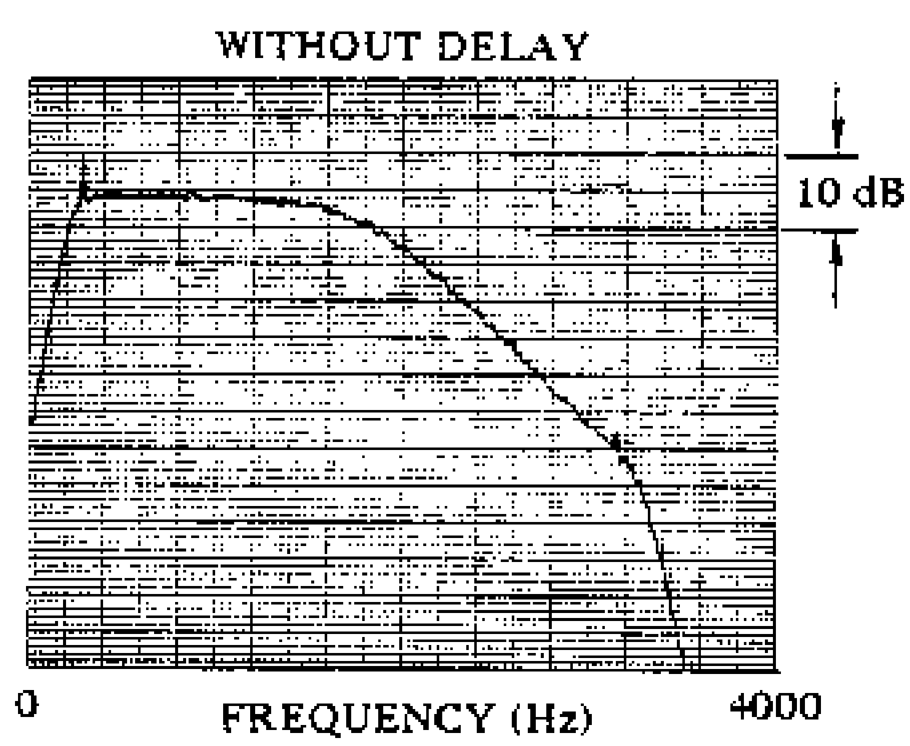

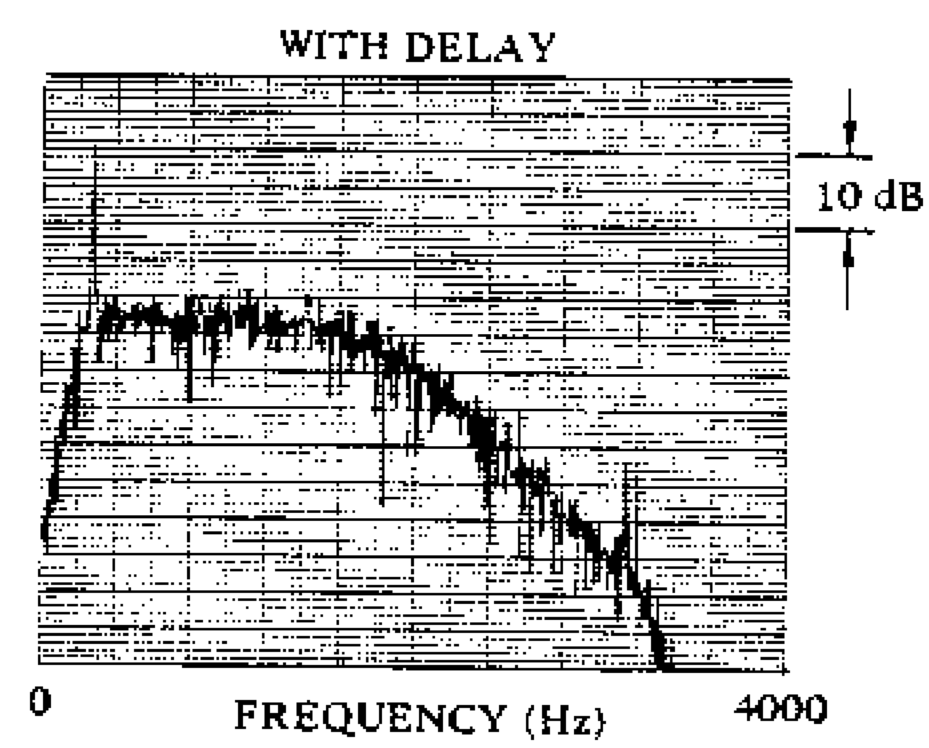

In the original cross-spectrum display, only a single fundamental component with one harmonic at 3130 Hz was visible above the noise. With delay, these two tones are now visible clearly above the noise and, once again, it is obvious that the linear amplitude display is optimum for signal detection applications. In the spectrum case, the total signal-to-noise improvement achieved simply by using the delay technique with no other reference inputs, was 16 dB. Signal-to-noise improvement in this mode depends upon the nature of the noise and periodicities, but, in general, can vary from an improvement of 3 dB to an improvement of nominally 20 dB.

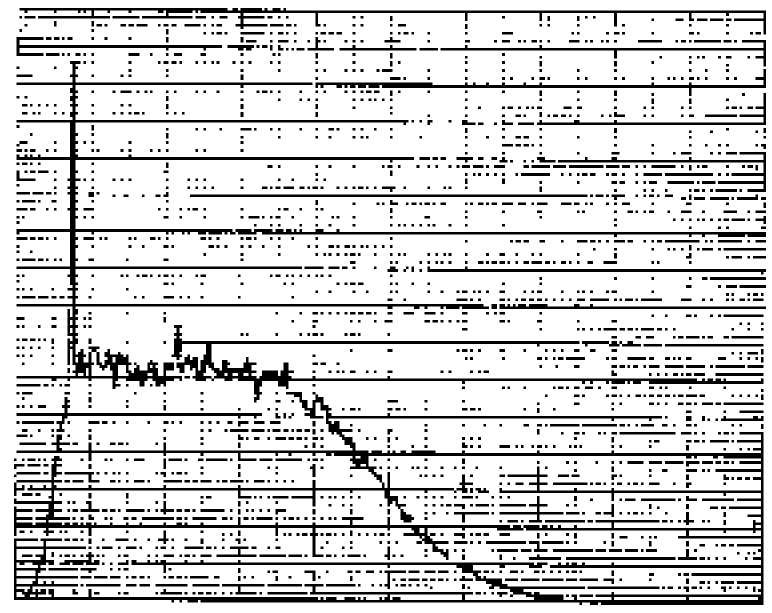

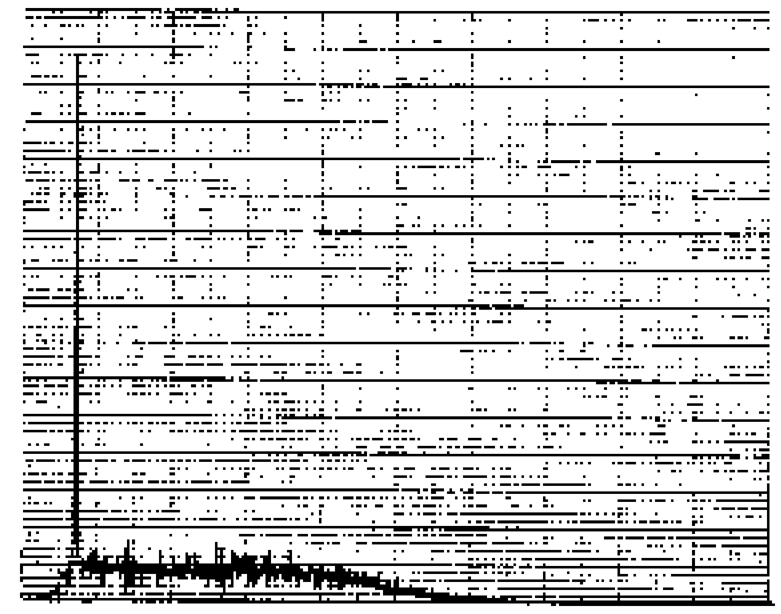

TIME-DOMAIN AVERAGING

Complete recovery of a periodic waveform and all its harmonics

There are many applications in which a repetitive signal is either generated (radar transmission) or simply exists by the nature of the mechanism being analyzed (rotating machinery applications). If it is desired to recover a periodic function and its harmonics from a composite signal which may include noise plus other periodic, but non-related, components, there are several good frequency-domain techniques. However, if a reference signal is also available which occurs in synchronism with the desired periodicity, an additional optimum technique can be used. This technique, synchronous averaging, is also sometimes called signal averaging or time-domain averaging. In practice, time-domain averaging is nothing more than synchronous sampling and, then, data summation.

The loading of data into the input memory is started in synchronism with a reference periodic pulse. When the memory is filled, the data is transferred to an accumulating memory. The input memory with the occurrence of the next sync pulse begins to load once again. After the input memory is filled the second time, its contents are summed with the first memory load in the accumulating memory. In this manner, any components which are synchronous in time with the reference sync pulse will tend to add up at whatever memory location they occupy. Other non-synchronous signals which are either random or periodic will average down and approach zero as a limit. By processing with this technique in the time domain, an entire periodic waveform can be recovered with the lowest frequency of the periodicity equal to the reference sync frequency and with all the harmonics of that signal also retained.

Although signal averaging is a form of digital sampling, it is identical mathematically to forming the cross-correlation product between the arbitrary data signal and a sequence of unit impulses occurring at the sync rate. It represents virtually the only known technique for recovering completely a periodic waveform with all its harmonics intact. The figure below shows the periodicity of the 287 Hz square wave recovered from the composite signal seen previously. For this type of processing, the signal-to-noise enhancement which can be achieved is directly proportional to the square root of the amount of averaging time. Thus, for signal averaging, the following relationship holds.

$$S/N_{enh} = 20 \log_{10}(n)^{1/2} \text{ dB}$$
$$= 10 \log_{10} n$$

where n is the number of independent ensembles averaged.

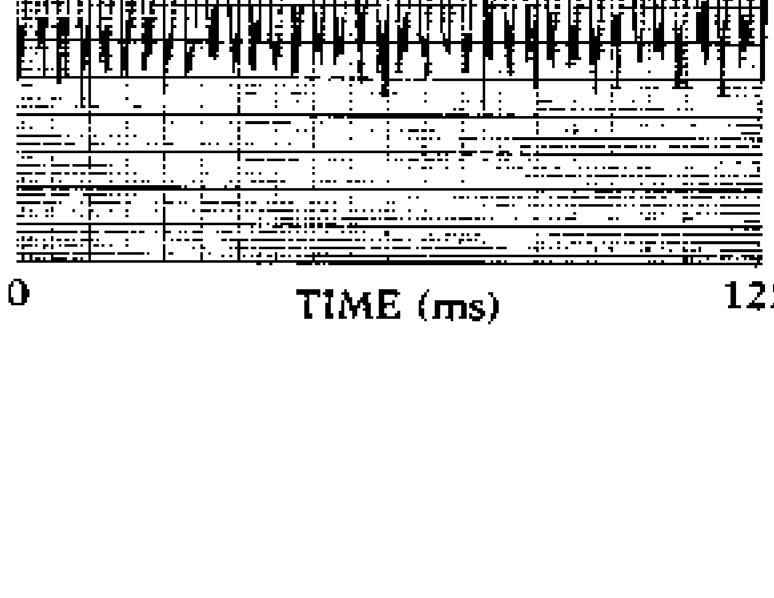

Isolating specific noise sources

An ideal application for time-domain or synchronous averaging is to isolate specific vibration or noise characteristics associated with the operation of rotating machinery. This technique can be very effective; for example where several large machines are operating in the same area and it is necessary to isolate the machine which is the source of a dynamic characteristic. The instrumentation below was used to observe the time waveform and extract from the overall sound pressure level waveform the component of this waveform associated first with N_1 operation and then with N_2 operation. When performing a signal averaging process, an external sync signal must be supplied as a reference to formulate the correct average. Signal averaging can then be done using the external sync input and internal sampling. If it is necessary to track changing speed characteristics, external synchronization with external sampling can be employed.

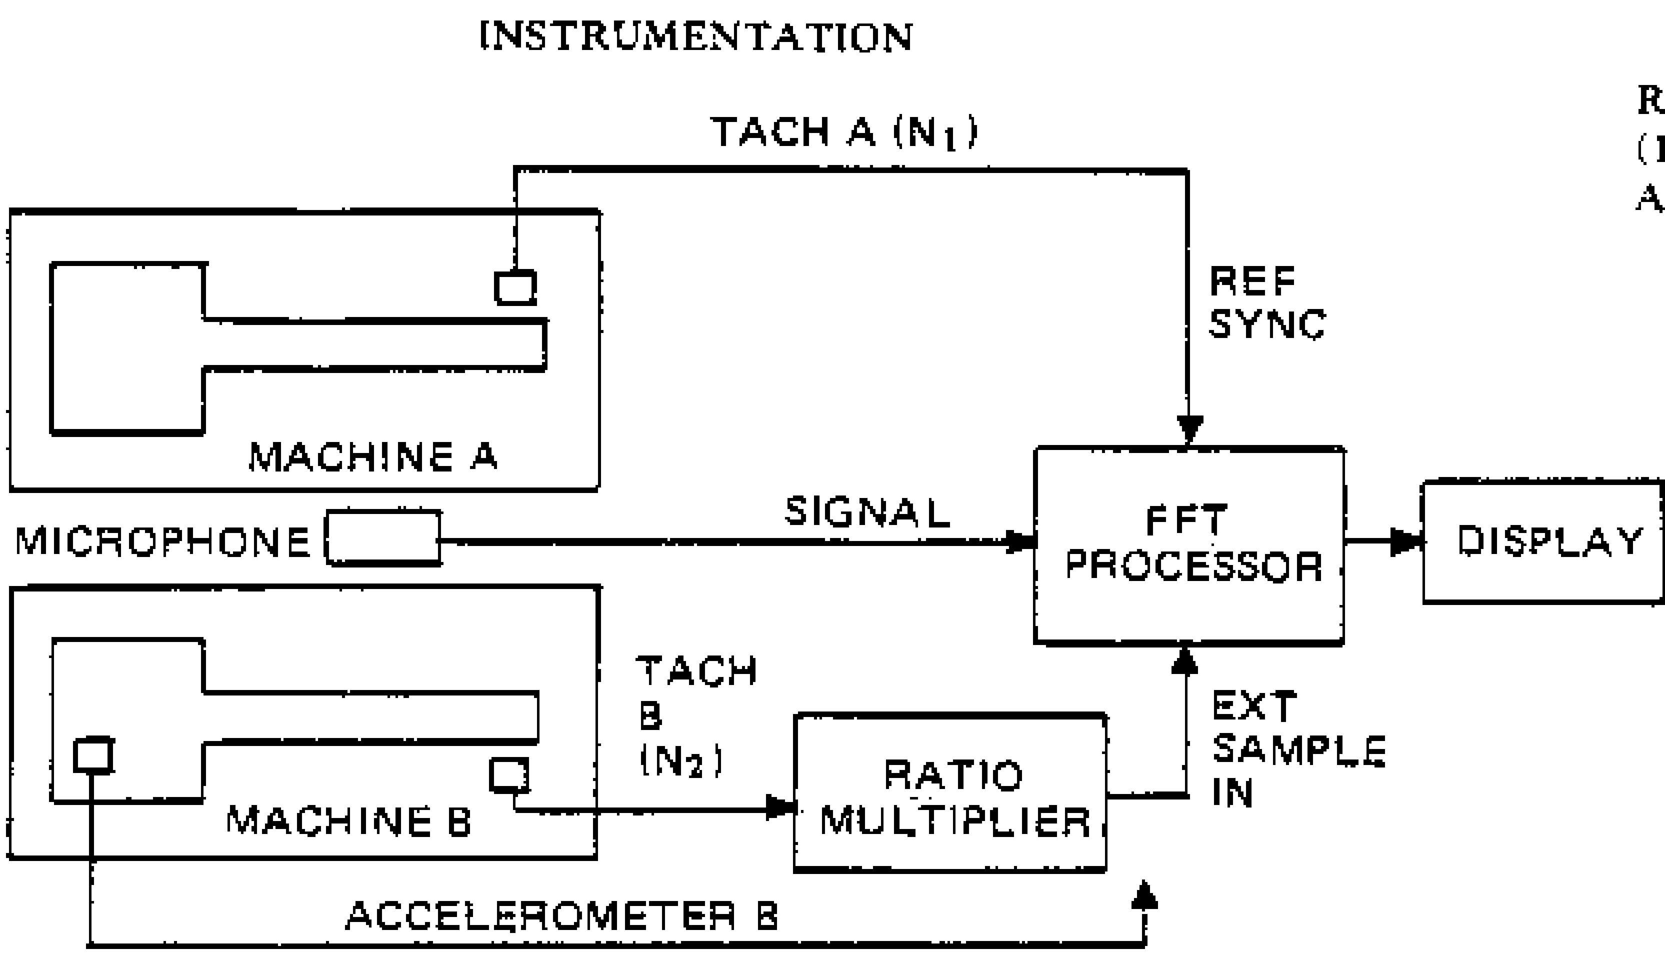

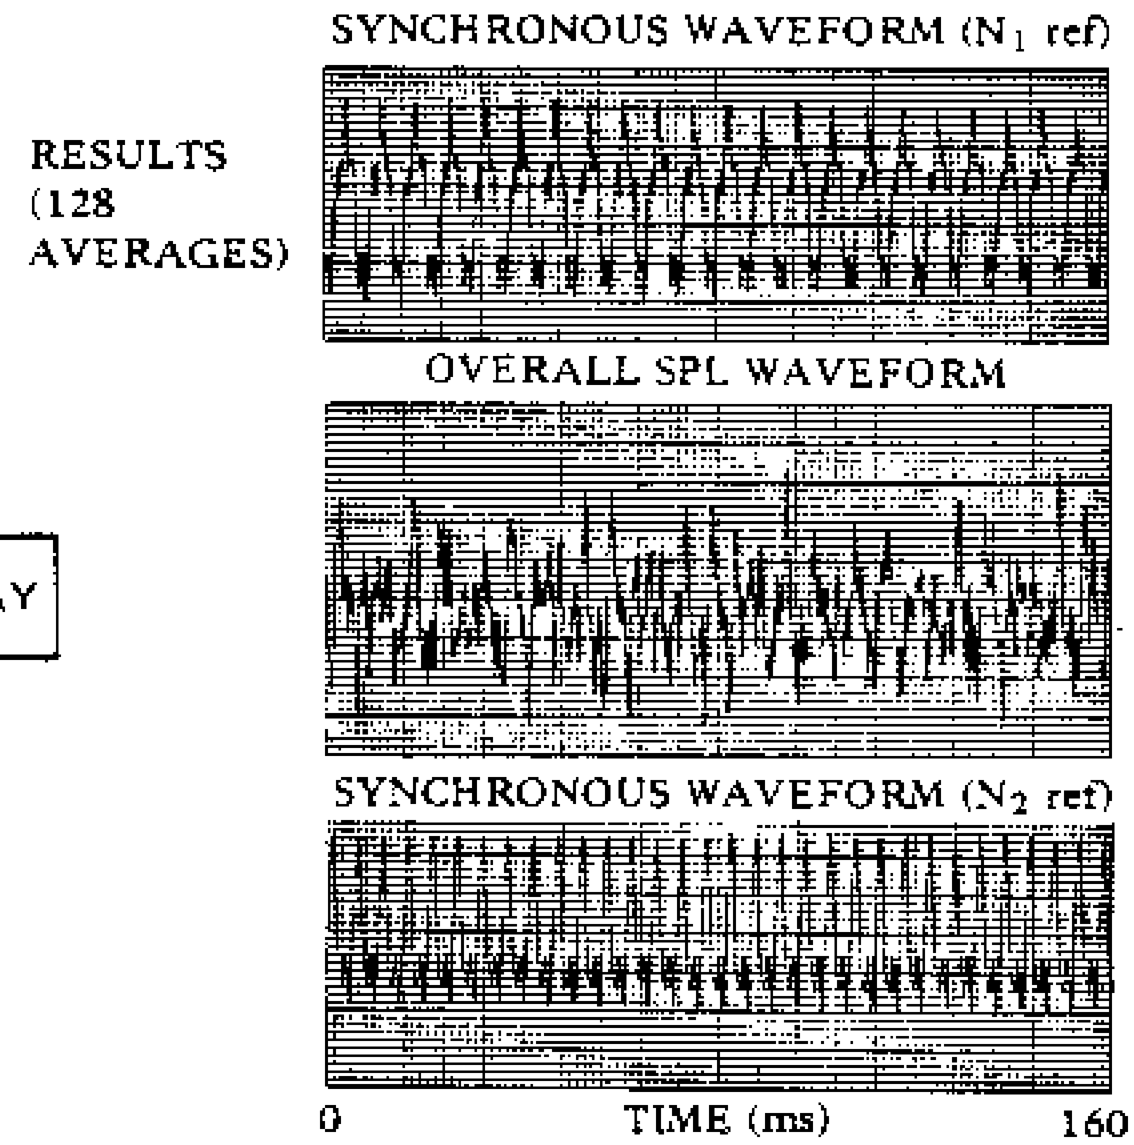

SYNCHRONOUS SPECTRA

Eliminating unrelated harmonics

After time-domain averaging, it is often extremely useful to Fourier transform the time-averaged result. Spectrum averaging is not necessary since synchronous time-domain averaging has already been performed. Both the overall standard spectrum and synchronous spectrum associated with the machine operation seen previously are shown below. The fundamental shaft speed of the two machines nearest the microphone was 5640 rpm (94 Hz) and 7440 rpm (124 Hz). The contribution of these two terms with their harmonics and broadband noise components to the composite spectrum is seen in the figure below. The two synchronous spectra have effectively eliminated the effects of these harmonic components generated by the non-referenced machine and, in addition, have reduced broadband noise contributions by an average of greater than 40 dB, revealing the presence of many harmonic components previously masked. If an external sync or tach pulse is available, one of the best possible signal-to-noise enhancement procedures is to synchronously time average a raw data waveform and then Fourier transform the result to secure the synchronous spectrum.

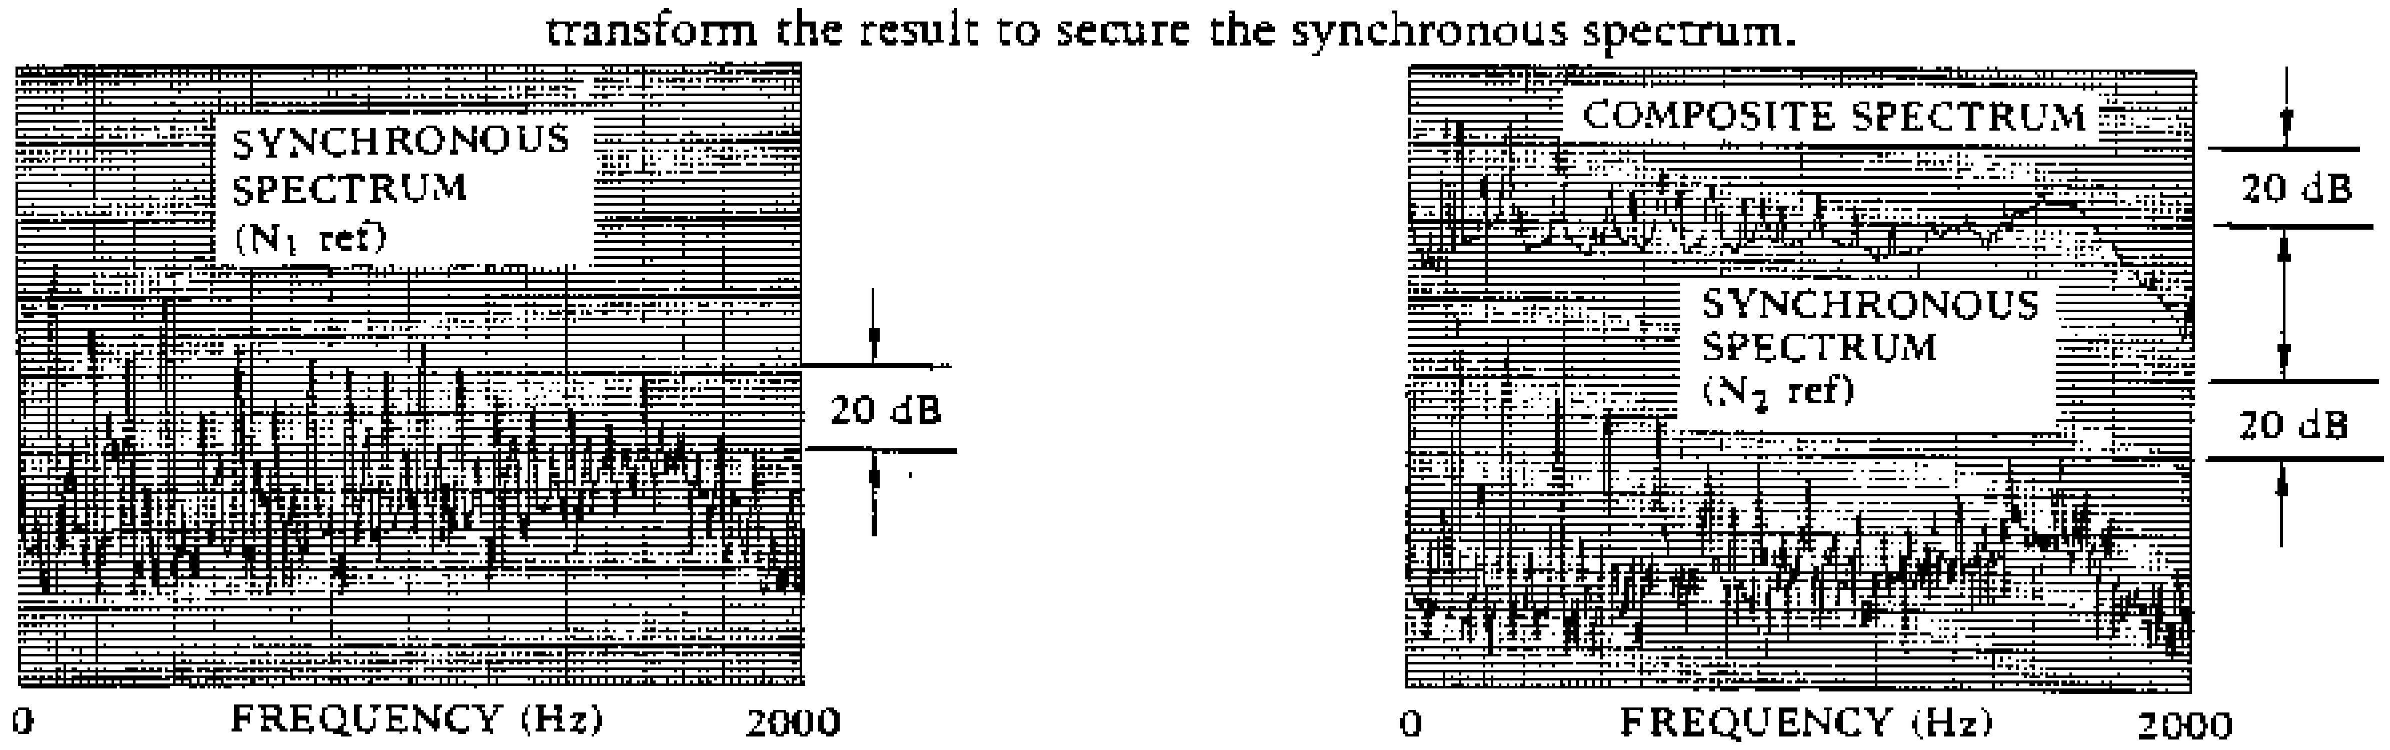

STANDARD AND SYNCHRONOUS SPECTRA REFERENCED TO N_1 AND N_2

PRIME SPECTRA

Enhancing a specific harmonic

The previous section showed that using a fundamental tach pulse the computation of synchronous spectra can provide a powerful signal-processing tool. Interesting characteristics are obtained if some manipulation of a 1/rev tach pulse is performed prior to synchronous spectrum computation. The use of the synchronous spectrum with a fundamental or once-per-rev tach reference produces a spectrum which, in effect, has the characteristics of a ***comb filter*** in which the teeth of the comb are spaced at multiples of their fundamental tach rate. If the one-per-rev tach signal is multiplied by a constant prior to synchronous spectrum averaging, each of the teeth of the resulting comb filter shifts in accordance with the new sync reference. If the multiplication factor is such that it corresponds with a harmonic of the signal being analyzed, that specific harmonic will tend to be greatly enhanced.

Effects of multiplication by prime, and not prime, constants

In a synchronous displacement spectrum taken using a 1/rev reference tach signal, and a noncontacting eddy current probe mounted in a gas turbine, it is obvious that the predominant components seen reflect a large unbalance situation with little higher order harmonic content. Multiplying the tach signal by 7 and then using it as the reference sync input produces a synchronous spectrum in which the 7th harmonic has been significantly enhanced. The fundamental signal is still present, but the harmonics between the fundamental and 7th have been significantly reduced on a relative basis (also representing a gain of 256 from the original scaling). If, on the other hand, the tach signal is multiplied by a constant which is not a prime number, submultiples of the multiplier will also contribute to the resulting synchronous spectrum. A tach multiplier of 8 was used to generate a reference sync signal. Not only was the 8th harmonic enhanced, but also the second, fourth, etc. As a final example of synchronous spectra, the tach signal was multiplied by 13 and clearly the 13th harmonic of running speed is the primary component which has thus been enhanced. Although the relatively high level of the fundamental signal still shows up in this resulting spectrum, it is obvious that other harmonics between the fundamental and 13 have, in effect, been discarded as the 13th harmonic has been enhanced, a classical example of a prime spectrum.

SYNCHRONOUS SPECTRA

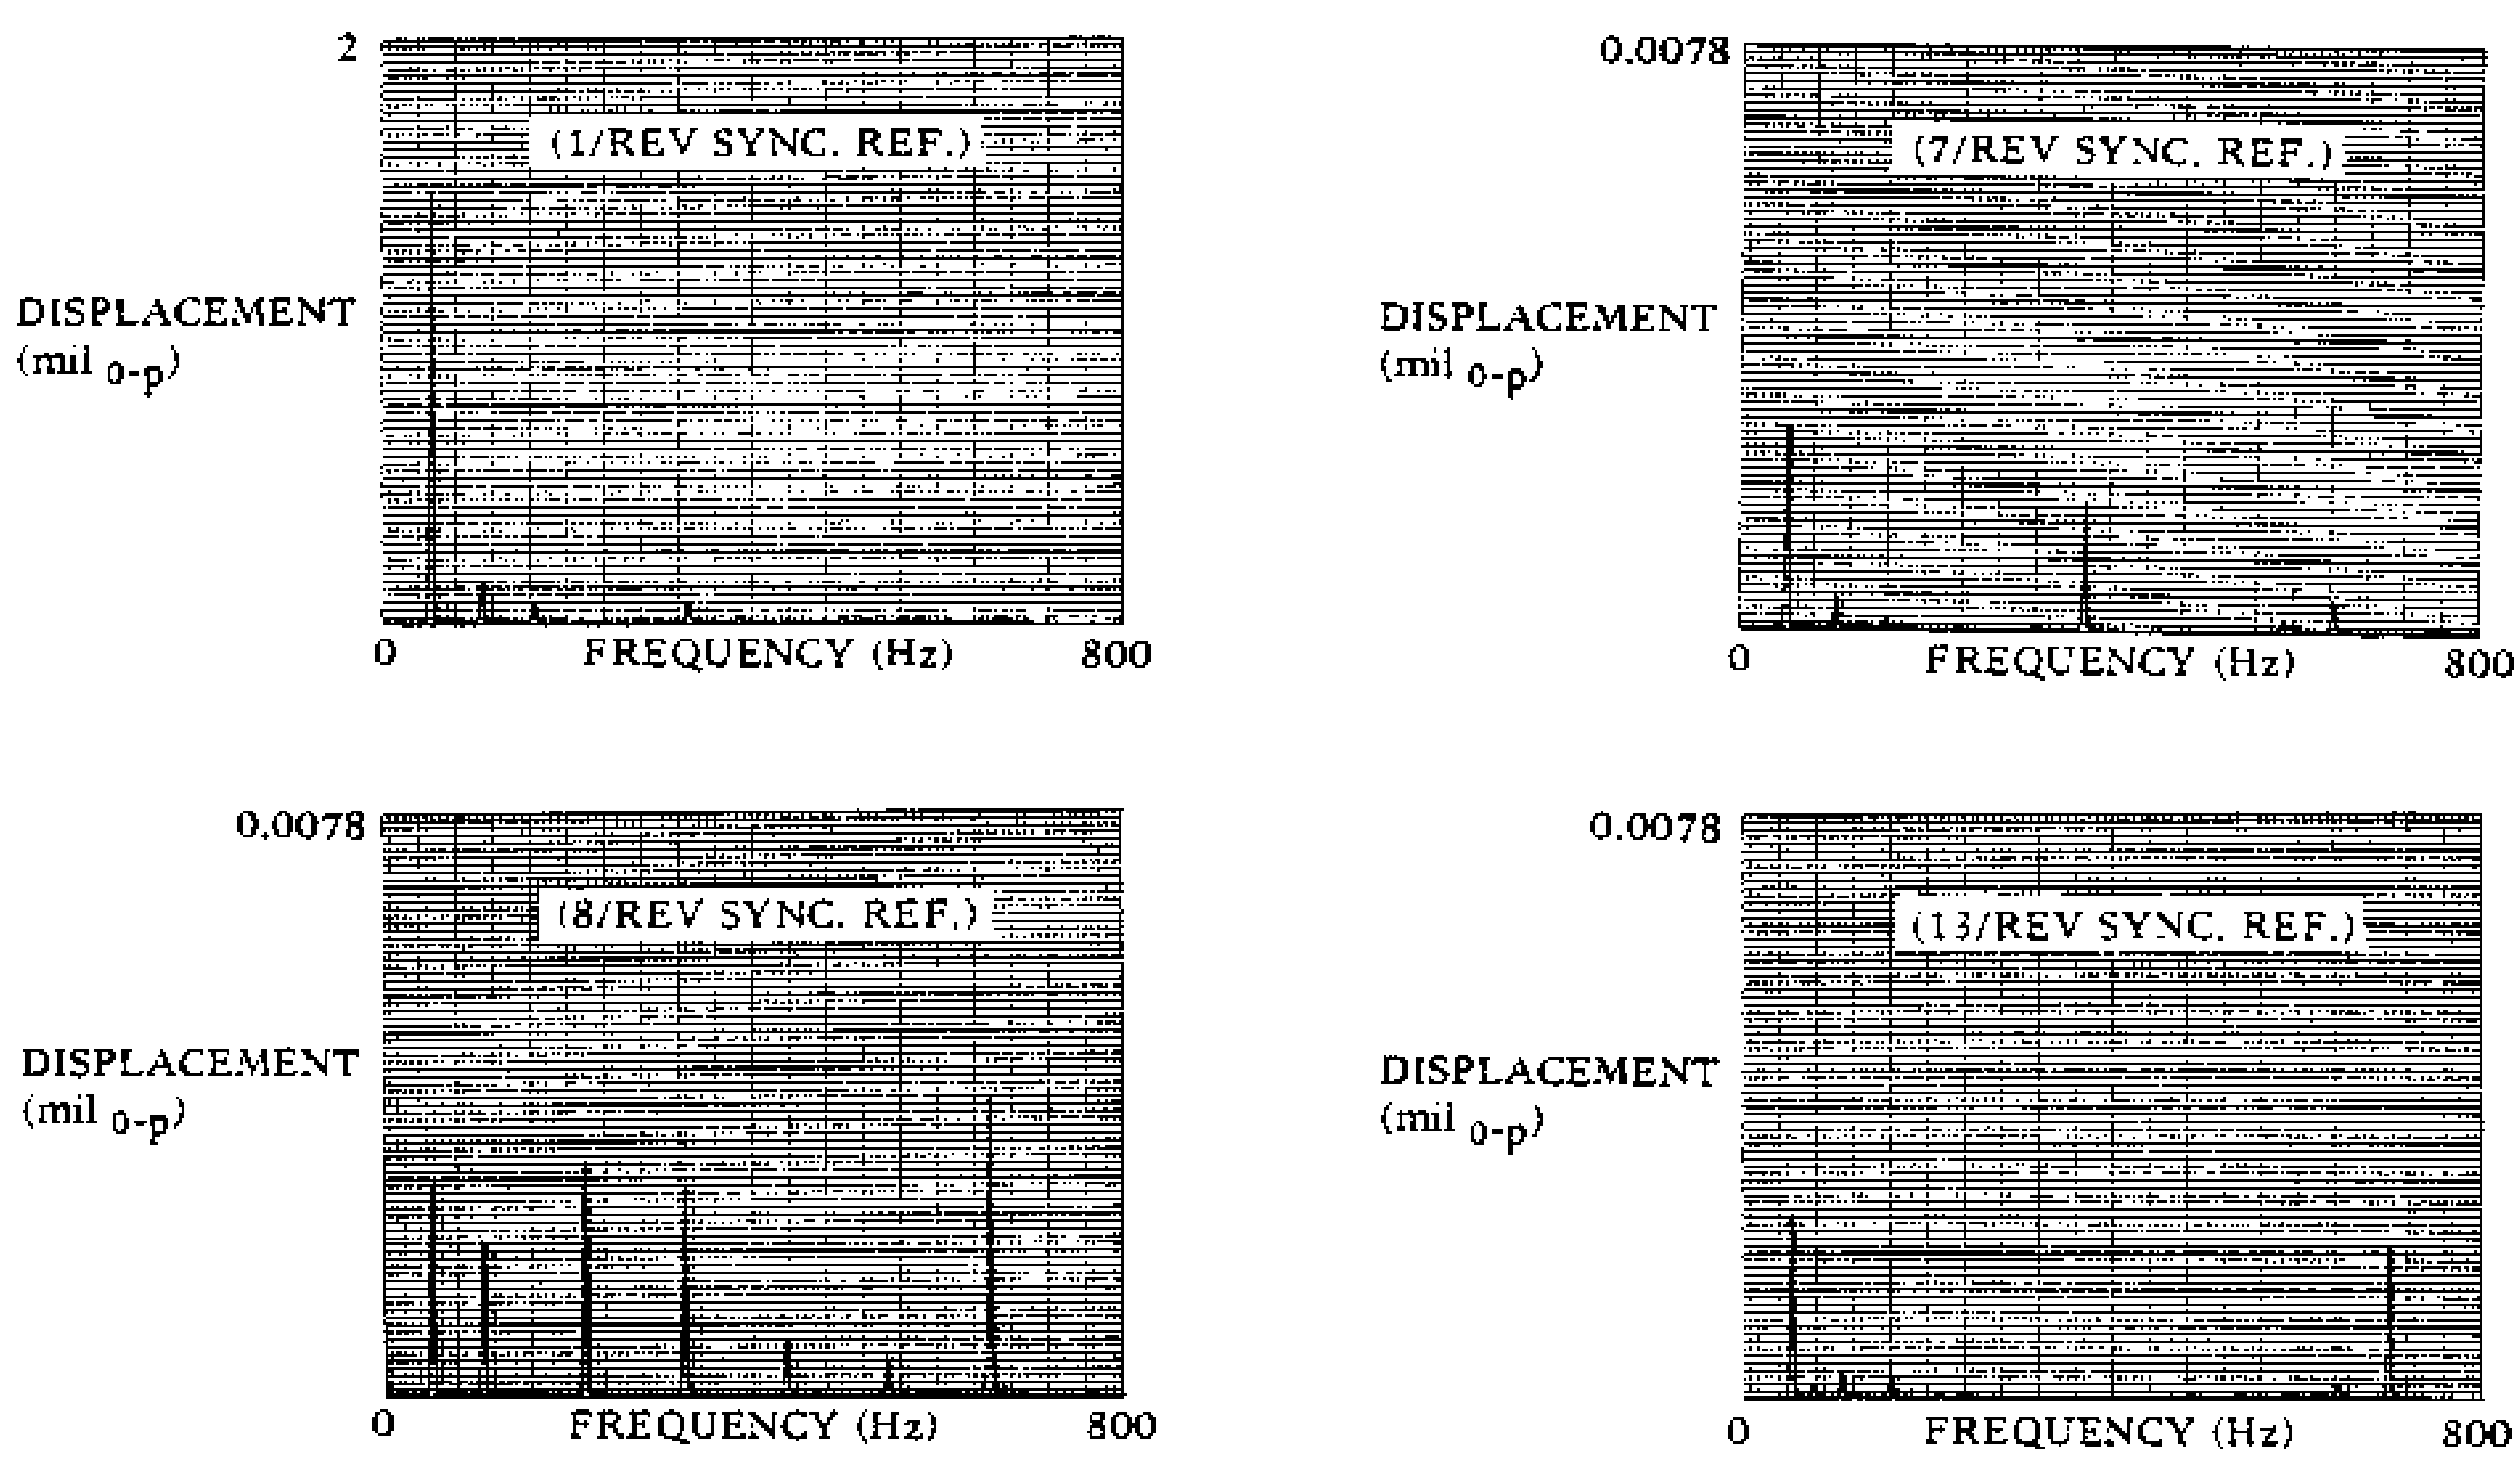

DATA FILTERING

Isolating signal characteristics

In many cases, an important characteristic of a signal may be masked by some unwanted repetitive or broadband random signal. If the characteristics of the portions of the signal which are wanted are distinct enough from the unwanted, a simple filtering process can be used to isolate the desired signal characteristic. This process could be either high pass, low pass or bandpass filtering such that the filtered signal can then be observed or processed in either the time or frequency domain. An example of this is seen below, in which a condenser microphone was used to monitor the noise produced by an operating, motor-driven fan where there was also broadband wind noise present. Applying a low pass filter, a simple act, has isolated the predominant time domain characteristic, a noise signal dominated by the fundamental and second harmonic of the machine operating speed. Naturally, adjustable filtering could be used in this case to select a specific portion of the signal's spectrum and observe its characteristic in either the time or frequency domain.

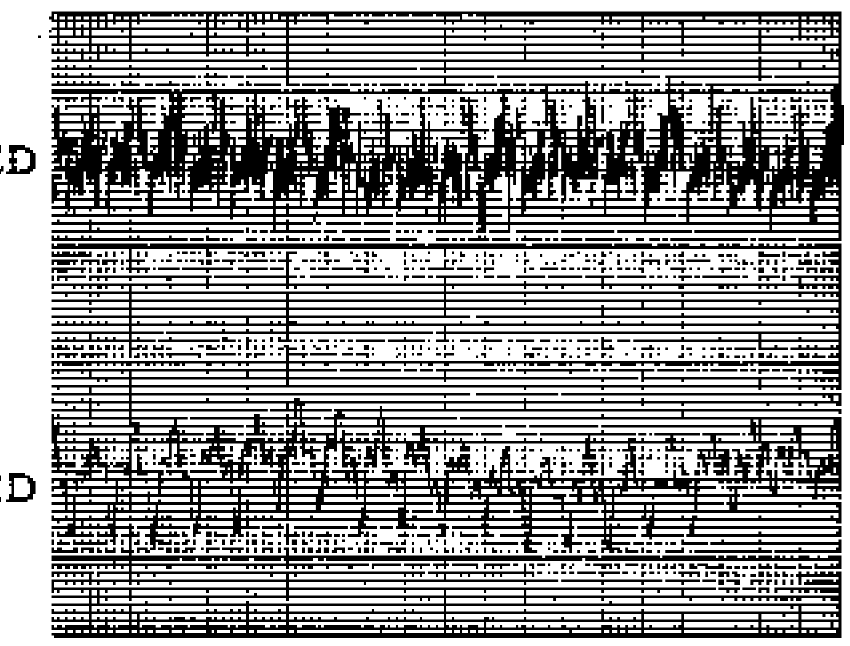

UNFILTERED AND LOW PASS FILTERED SOUND SIGNAL

DIGITAL FILTERING

Optimizing signal filtering

In the preceding section, it was seen that using a low pass or bandpass filter can often enhance a signal to the extent that the information which can be obtained from the signal is greatly increased and the signal can be much more clearly visualized. There are many cases where a signal is available in sampled value form and where a certain type of signal cleanup would be extremely valuable in interpreting time-domain signal representation. In many of these cases, it would be desirable to optimize a signal filtering or signal processing task to the actual data signal. In these situations, a digital filtering technique offers an optimum processing condition in terms of defining the filter shape, stability and accuracy of the filtering process, and convenience both in changing filter characteristic and in observing the resulting filtered waveform. Conceptually, digital filtering is equivalent to observing an output waveform which is the convolution product of an input signal to be filtered and the impulse response of the filter to be implemented, represented by

$$o(t) = \int_0^T f(\tau)\ h(t-\tau)\, d\tau$$

where $o(t)$ is the resulting filtered or convolved waveform, $f(t)$ is the signal to be filtered and $h(t)$ is the impulse response of the filter characteristic.

In digital-processing terms, a linear digital filter can be defined using the principle of super-position as follows;

$$y(nT) = \sum_{m=o}^{n} h(mT)\ X(nT - mT)$$

where $x(nT)$ is the input signal; $h(mT)$ defines the filter weights; $y(nT)$ represents the output sequence; and T is the sampling interval.

Eliminating unwanted tonals

An example of a signal requiring some extensive cleanup in order to assist interpretation of a physical phenomena is given in the figure below. Here an impact waveform was recorded, but was heavily corrupted and virtually obliterated by the presence of an unwanted tonal. A filter shape in the frequency domain was then specified to eliminate the tonal, but optimize the characteristic of the impact spectrum. The equivalent specification of this filter in the time domain is seen below. The resulting filtered waveform which has significantly reduced the effects of the periodicity within the time frame of the pulse itself is given below. Here, not only the initial impact, but also the characteristics of some of the rebounds previously masked by the periodicity are now clearly seen. In terms of signal cleanup, especially where data signals have been previously sampled or available in digital form, the concept of digital filtering is extremely useful.

SAMPLED AND STORED CORRUPTED WAVEFORM

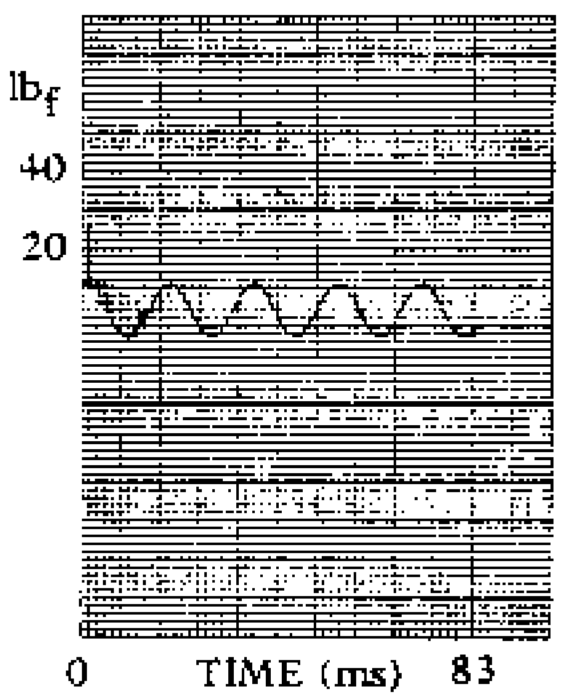

DIGITAL FILTER CHARACTERISTIC EXPRESSED IN TIME DOMAIN

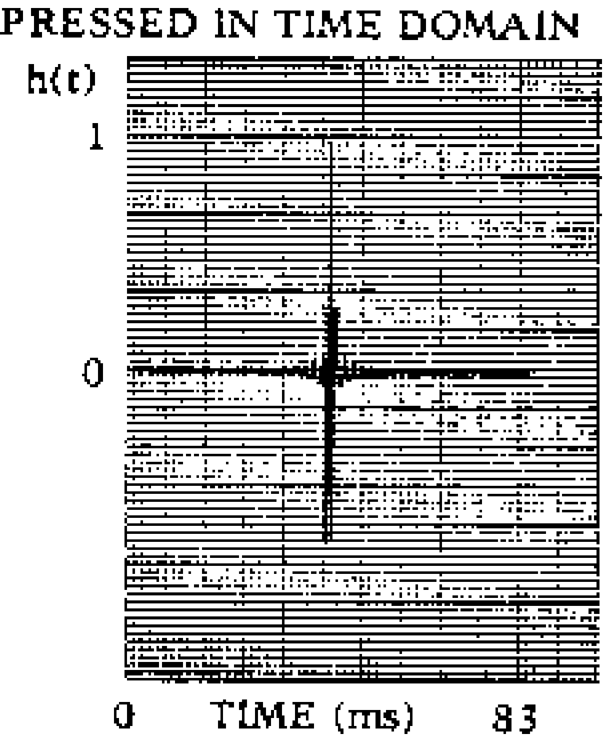

FILTERED OUTPUT FORCE PULSE

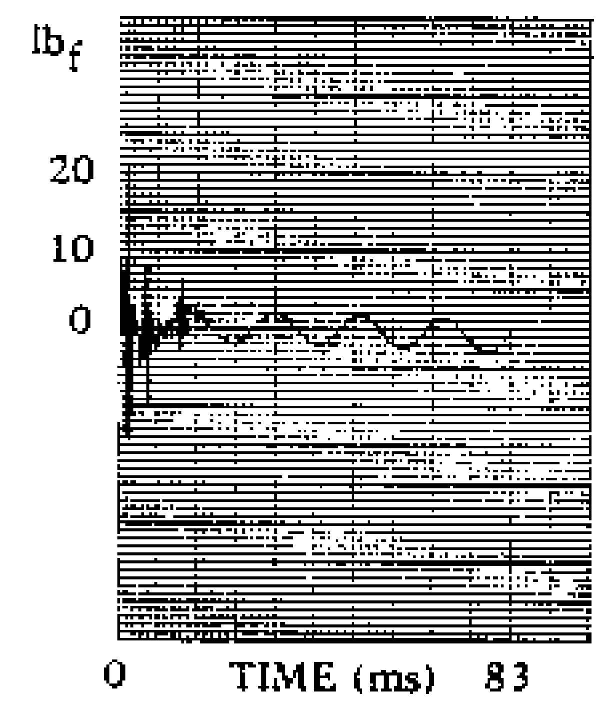

WEIGHTING CONSIDERATIONS

External sampling and rectangular weighting for optimizing signal-to-noise enhancement

The effective filter shape which is achieved when doing a spectrum analysis using a sampling type of system is either partially or exclusively determined by the type of windowing function employed in the analysis. Selection of a window alters the selectivity characteristic of the Fourier transform process and, as a result, can affect signal-to-noise enhancement considerations. Many weighting functions exhibit nulls in their side-lobe characteristics at appreciable amplitude levels, producing an effective filter shape which is perfectly rectangular at the frequency cells (Δf) of the nulls. The most obvious example of this type characteristic is the rectangular weighting window. In most cases, the rectangular window cannot be employed because of the unwanted truncation and side-lobe effects which it produces. However, if an external sampling input, synchronous to a periodic signal input, is available, rectangular weighting can be used to give an extremely selective analysis characteristic. An example of this is seen in the figures below.

KAISER-BESSEL WEIGHTING

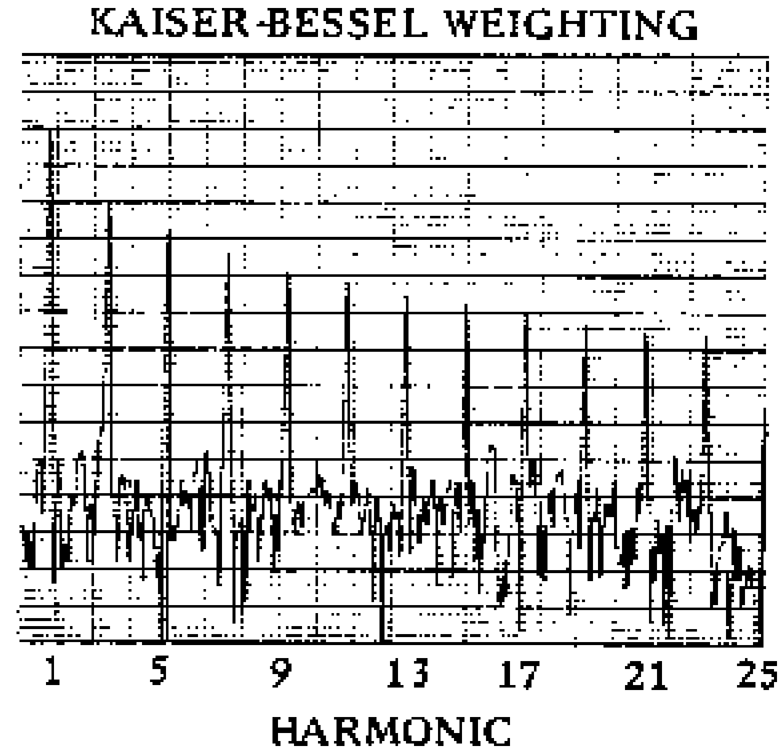

RECTANGULAR WEIGHTING

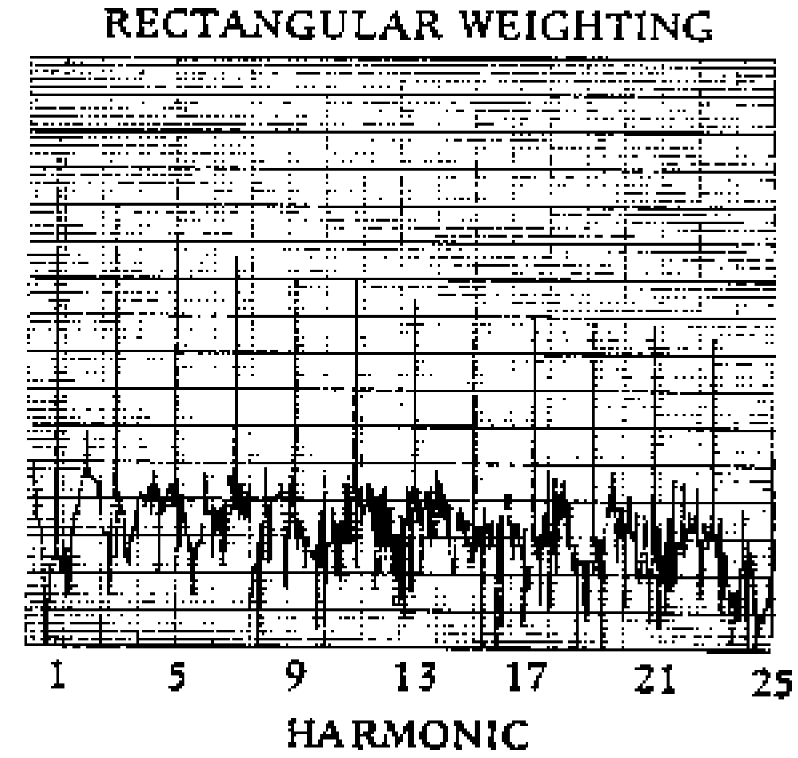

The first figure is that of a synchronous spectrum analysis performed up to the 25th harmonic, using external sampling inputs to an FFT processor with Kaiser-Bessel weighting. Excellent signal-to-noise enhancement is achieved; however, although the Kaiser-Bessel window has nearly an optimum filter shape, its actual filter characteristic is clearly visible.

In the next figure, the same analysis is performed using rectangular weighting. Note the extremely sharp filter shape, and also a slight apparent improvement in signal-to-noise ratio over the previous condition. In situations which permit external sampling, especially where synchronous averaging will not be used, this can be an important factor to consider in optimizing signal-to-noise enhancement.

ENHANCEMENT BY THREE-DIMENSIONAL DISPLAYS

Detection of buried periodics which are changing in frequency with time

One of the most difficult signal-to-noise enhancement tasks is the detection of periodic signals nearly buried by a random background when the periodic signals are moving about in frequency as time passes. A pure-tone signal moving about in this random background will often be missed if it is merely observed on a real-time oscilloscope display. For these applications, especially when no synchronizing or reference signal is available, averaging techniques are often inappropriate inasmuch as they would tend to smear the very signal that is under observation into the background noise. Sometimes, a display in terms of a *waterfall* (moving spectrum history) is most appropriate. The real-time spectrum of a signal, which consists predominantly of random noise, is also a very random display. However, visual enhancement can be achieved by displaying spectra of this nature in a continuous three-dimensional presentation. The figure below shows a typical presentation of this type in which a tone is swept back and forth through a predominantly random background. In this case, no spectrum averaging was performed as it is important to know precisely the total excursion in frequency of the signal being observed. It would be extremely difficult, if not impossible, to extrapolate the total signal excursion characteristic in this instance if using only a standard oscilloscope display.

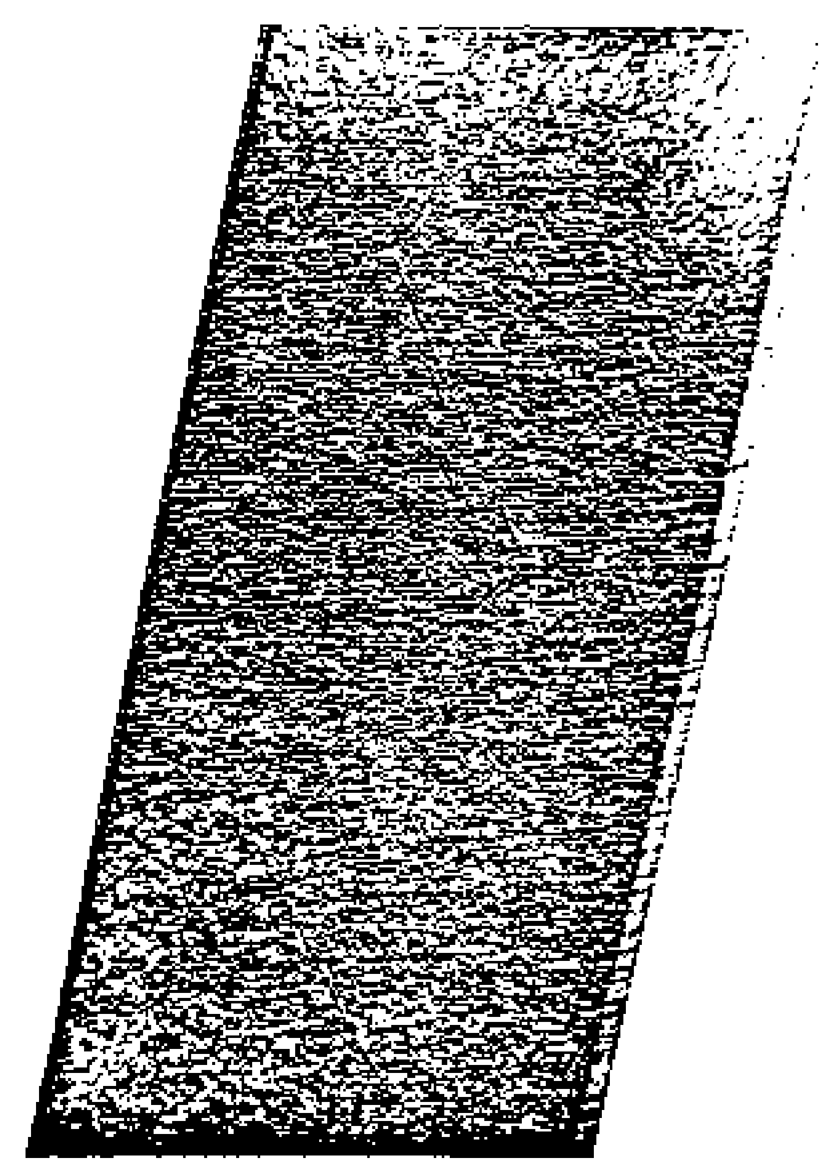

OVERLAP PROCESSING WITH VISUAL ENHANCEMENT

Obtaining faster visual presentations

In addition to observing a three-dimensional presentation of amplitude, frequency and time, enhancement is also available through the use of overlap processing. In this context, overlap processing is defined as the ability to produce a spectrum display in a period of time faster than the time samples were originally accumulated under the condition of continuous accumulation and updating. For example, if 4 transforms per memory period are used, the processor will perform 4 transforms in a 1 memory period time frame rather than just one transform. Initially a block of data is read from input memory to processor memory and a transform is performed. One-fourth memory period later another complete block of data is read from input memory to processor memory and a transform is performed. However, the second block of data contains three-fourths of the data collected in the first block plus one-fourth of new data. This is overlap processing. By its very nature, overlap processing is not, in general, designed to increase statistical confidence of a resulting analysis, but if it is done correctly, can give some spectrum averaging still within the time frame of 1 memory-load period. However, one of the primary advantages and applications of overlap processing is to obtain faster visual presentations of exactly what is happening and, as in the example, additional visual signal enhancement.

Determining precise signal excursion

The figure below shows a spectrum scan from 0 to 2000 Hz in which a 500-line analysis was performed with no overlap or redundant processing. Each spectrum display represents one non-redundant spectrum scan in which a totally new memory sample was employed each time. Contained within this background noise spectrum is a sinusoidal signal at an rms level of −17.5 dB with respect to the rms level of the noise which was moving across the spectrum at a rate of 42 Hz per second. No evidence of this moving tone can be seen in the first display.

The next display shows the identical signal condition, but with an overlap factor of 2.5:1. The same analysis range of 2 kHz and the same effective bandwidth filter was used in both cases, but the number of spectrum displays per unit time has been increased by a factor of 2.5. An obvious periodicity is now apparent in the spectrum.

NO OVERLAP

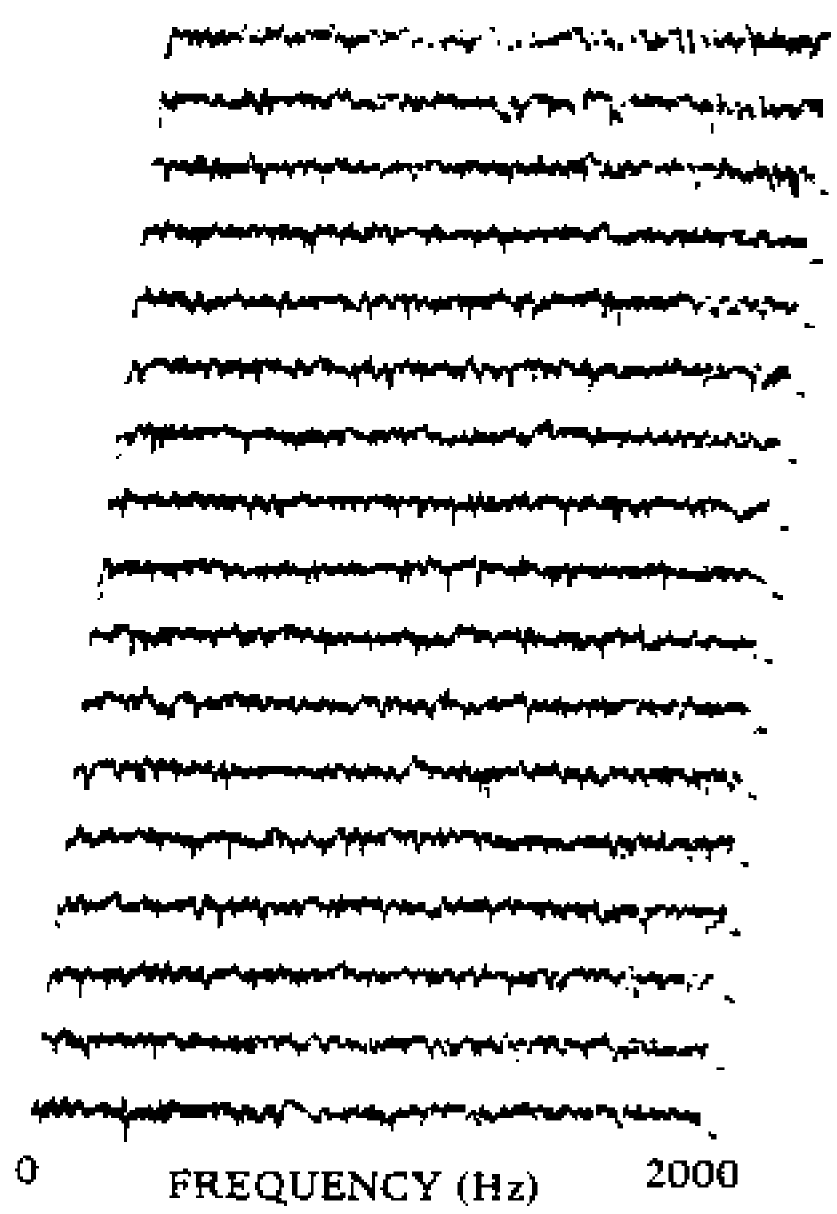

OVERLAP OF 2.5

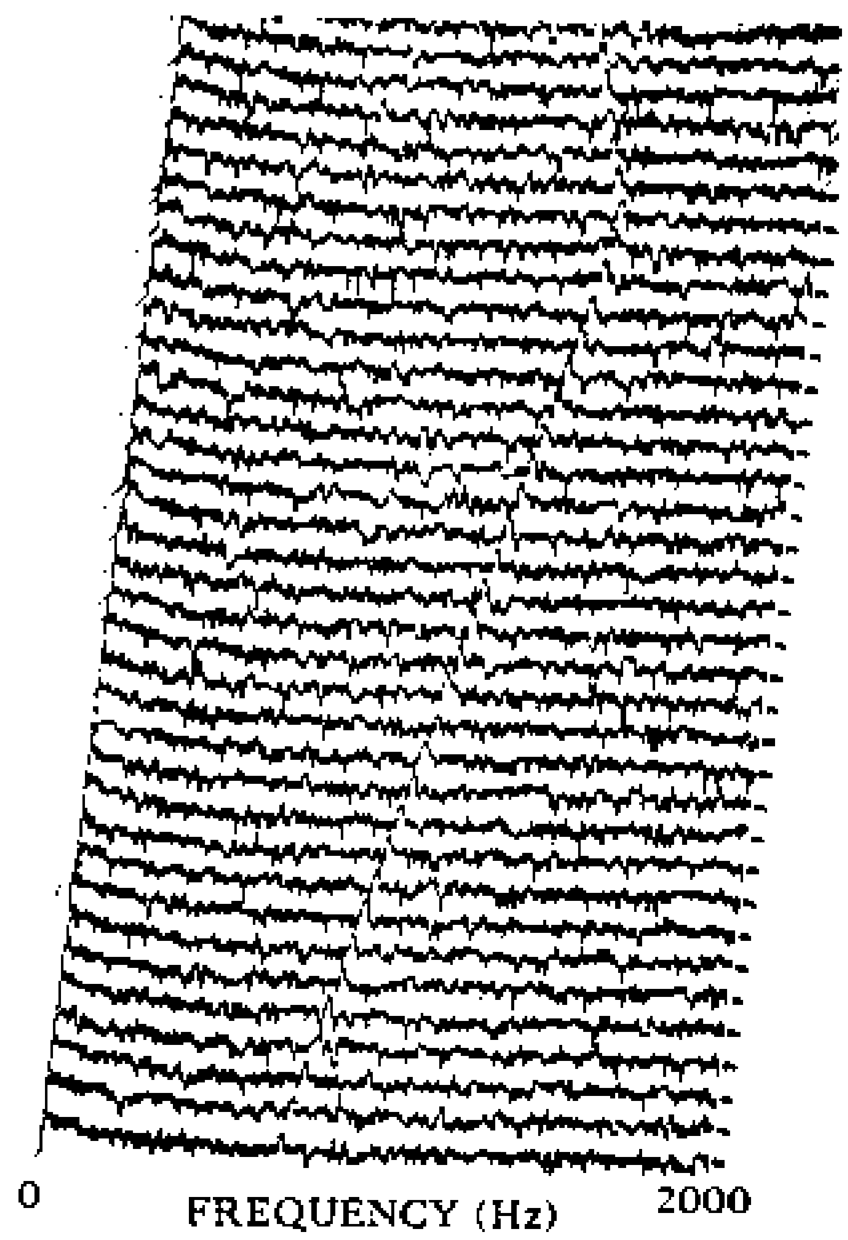

OVERLAP OF 5

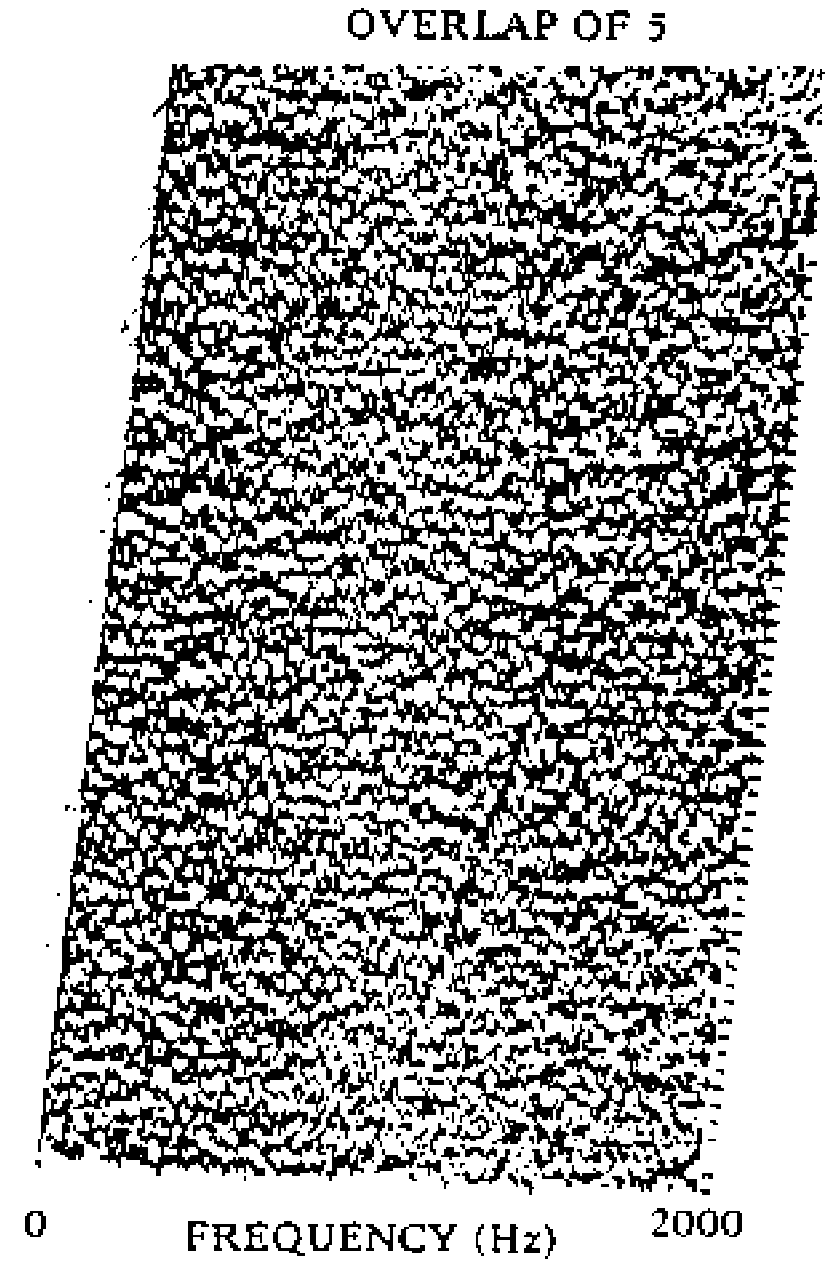

Finally, an overlap factor of 5 was employed. This is, once again, shown by the increase of number of displays per unit time. Again the signal stands out quite clearly and even though no averaging has been employed, sufficient signal-to-noise enhancement is achieved to permit an accurate determination of precise signal excursion. Previous investigations have shown that the overlap factor cannot be increased indefinitely. Typically, for this type of display, optimum overlap-processing factors range between 2 and 8, with no appreciable increase in apparent visual enhancement using more overlaps.

CONCLUSIONS

1. Signal enhancement can be achieved in the time and frequency domains.

2. For straight spectrum processing, signal-to-noise enhancement improves as the analysis filter bandwidth decreases.

3. For correlation processing, signal-to-noise enhancement relates to the square root of total averaging time.

4. If no reference or synchronizing signal is available, single-channel data signal enhancement can be achieved using two-channel signal processing techniques with an adjustable artificial delay between data signals.

5. If it is desired to recover periodic signals from noise, and a reference or synchronizing signal is available, there is maximum signal-to-noise enhancement in either the frequency or time domain by using synchronous averaging.

6. Overlap processing can effectively be used to achieve visual enhancement and, in some conditions, actual signal-to-noise improvement through averaging.

7. Three-dimensional displays, where available, can significantly improve signal-to-noise conditions from a signal-detection standpoint.