

Errata

Document Title: Signal Averaging With the 3582A Spectrum Analyzer (AN 245-1)

Part Number: 5952-8767

Revision Date: January 1979

HP References in this Application Note

This application note may contain references to HP or Hewlett-Packard. Please note that Hewlett-Packard's former test and measurement, semiconductor products and chemical analysis businesses are now part of Agilent Technologies. We have made no changes to this application note copy. The HP XXXX referred to in this document is now the Agilent XXXX. For example, model number HP8648A is now model number Agilent 8648A.

About this Application Note

We've added this application note to the Agilent website in an effort to help you support your product. This manual provides the best information we could find. It may be incomplete or contain dated information, and the scan quality may not be ideal. If we find a better copy in the future, we will add it to the Agilent website.


Support for Your Product

Agilent no longer sells or supports this product. You will find any other available product information on the Agilent website:

www.agilent.com

Search for the model number of this product, and the resulting product page will guide you to any available information. Our service centers may be able to perform calibration if no repair parts are needed, but no other support from Agilent is available.

SIGNAL AVERAGING WITH THE HP 3582A SPECTRUM ANALYZER



FOREWORD

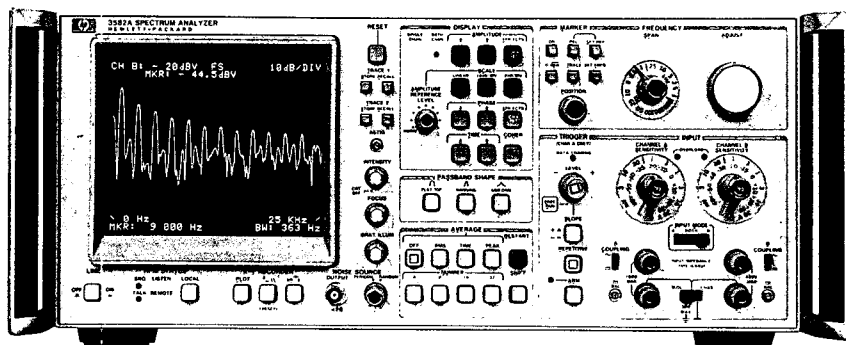
Low frequency spectrum analyzers are fast turning into digital instruments. A primary reason for this trend is the frequency analysis speed achieved through the startlingly efficient Fast Fourier Transform computing algorithm. With the additional help of large-scale integrated circuits, compact low-frequency analyzers are emerging which are often 10 to 100 times faster than traditional swept-tuned-filter spectrum analyzers. The HP 3582A is a low-frequency analyzer which continues in a long historical line of HP wave and spectrum analyzers. However, it is totally digital except for input filtering and cathode ray display.

The purpose of this application note is to develop an understanding, by theory and example, of two kinds of signal averaging commonly used in digital signal analysis of the kind employed by the 3582A. These averaging

routines are called "power spectrum averaging" and "time averaging." The corresponding swept analyzer techniques are video filtering (for smoothing measurements of random signals) and narrow-band analysis (when used to enhance signal-to-noise ratios). Two other related digital techniques are discussed which do not have direct equivalents in swept analyzers. These are signal-to-noise enhancement of recurrent transients and a peak-hold feature.

This application note is largely tutorial because of our belief that, while many readers will not have had experience in the area of digital signal processing, all will want a basic understanding of certain of its ideas in order to get maximum benefit from an instrument like the 3582A. This note and others in the series are intended to help provide this understanding.

THE HEWLETT-PACKARD MODEL 3582A SPECTRUM ANALYZER



The HP 3582A is a spectrum analyzer covering the frequency range of DC to 25 kHz. Although it is a FFT-based, digital instrument, a special design effort has made it as straightforward to use as a conventional swept analyzer. With dual measurement channels it is possible to measure transfer function gain and phase, as well as the coherence function. A built-in random or pseudo-random

noise source, whose spectrum tracks the analysis range, is a useful measurement stimulus. Band Selectable Analysis enables narrowband, high resolution analysis to be applied to any portion of the frequency range. The instrument comes equipped with a flexible HP-IB interface for control and two-way data transfers.

TABLE OF CONTENTS

SECTION		PAGE
1	The reasons for signal averaging Deterministic and random signals Classes of random signals The kinds of signals which benefit from signal averaging are defined and discussed.	1
2	The spectrum of a random process Some properties of a random process Analyzing a random process with the FFT Certain properties of random spectra influence the way measurements are made. In addition, the FFT of a random process has some intrinsic limitations which are overcome by averaging.	2
3	Power spectrum averaging Power spectrum averaging compared with video filtering This is a light treatment of the origin of power spectrum averaging, and a discussion of the analogy between this and video filtering.	3
4	Time averaging When to use time averaging How it works There is another kind of averaging which can be used whenever a synchronizing signal is available.	4
5	Summary of averaging properties Power spectrum averages Time averages This is a tabular summary of the principle features of both power spectrum and time averaging.	5
6	Characteristics of 3582A averaging routines Power spectrum averaging; Exponential averaging Time averaging Peak "averaging" The way signal averaging is carried out in the 3582A is explained, along with a couple of related techniques also available in the analyzer.	6
7	Examples Two power (RMS) averaging examples Two time averaging examples Peak averaging example Here are some documented experiments designed to illustrate each of the 3582A averaging techniques.	7
	Appendix 1: Probability distributions of power spectrum estimates Gaussian random processes The averaged power spectrum as a random process	12
	Appendix 2: 90% confidence limits for power averages	
	Bibliography	

SECTION 1:

The reasons for signal averaging

Deterministic and random signals. A deterministic waveform, or signal, is one which can be described by an explicit mathematical function of time. Deterministic signals are easy to visualize: sinusoids and pulse trains are good examples.

Signals derived from real, physical processes are not often deterministic. More likely they are random, or a mixture of deterministic and random. One good reason for this state of affairs is succinctly stated by Shannon's information theory: deterministic signals are not information-bearing, since they are predictable. Some common examples of real-world signals which are more or less random are speech, music, digital data, seismic data, and mechanical vibrations.

Classes of random signals. The technical term used to describe a signal which is a random function of time is "random process." Two useful classes of random processes are:


- a. Pure random process: a signal with no deterministic portion. An example is the sound emitted by a compressed gas escaping from a nozzle.
- b. Mixed random process: a composite signal, the sum of a pure random process and a deterministic signal. For instance, the output of a noisy amplifier with a sinewave input is a mixed random process.

An example of each of these is shown in Figure 1.


Measuring the frequency spectra of random processes involves some difficulties not encountered in measuring deterministic signals. The source of the difficulties is discussed in the next section. **The purpose of signal averaging techniques is to improve the measurement and analysis of random processes.**

Figure 1.

Typical random processes



a. Pure random process



b. Mixed random process

Each of the examples shown is a 50-millisecond segment of a process which, theoretically, continues for all time.

SECTION 2:

The spectrum of a random process

Although the exact time waveform of a random process cannot be described by an explicit mathematical function of time, there are certain ways to describe a random process which are explicit and quantitative. Among these ways are the probability density function, the auto-correlation function, and the power spectrum. These parameters provide valuable information about the process, and instruments are available to measure these (and other) parameters of any particular random process. In this note, we are concerned with the measurement of the parameter called the power spectrum. (The 3582A actually displays the square root of the power spectrum, called the amplitude spectrum.)

Some properties of the spectrum of a random process:

- The spectrum is a continuous function of frequency.** This follows from the fact that a random process is not periodic.
- The phase of the linear spectrum of a random process is a random function of frequency.** For this reason, a phaseless spectrum, called the power spectrum, is generally used to describe random processes. This is the same as the magnitude-squared spectrum, obtained by multiplying the linear spectrum by its complex conjugate. Sometimes the square root of the power spectrum is used, as in the 3582A.
- Each segment of the infinite-duration time waveform, being different from any other segment, makes a unique contribution to the spectrum.** Thus, the spectrum resulting from any finite-time measurement is only an approximation to the true spectrum of the random process.

In this last property is found the primary difficulty in measuring the spectrum of a random process. Because we are limited to making finite-time measurements, **the spectrum calculated from one finite segment of the random process only approximates the true spectrum.** It follows that the spectrum derived from one time segment will differ from that of the next, and so on. How inaccurate is such a single measurement? How can we increase the accuracy of measuring the power spectrum? These questions will be taken up next.


Analyzing a random process with the FFT.

The Discrete Fourier Transform, in the form of the Fast Fourier Transform (FFT) implementation, translates a finite segment of discrete time data into a discrete frequency spectrum. For instance, the 3582A operates (in single-channel mode) by sampling the signal on its input terminals to produce 1024 binary numbers representing a segment of the input time function. These numbers are transformed by the FFT into 512 complex values in the frequency domain. Because of possible aliasing, only 256 of these are used. The amplitude display consists of 256 connected points representing the numerical contents of 256 data storage locations called "bins." The numbers in the bins are calculated from the FFT output and can represent either spectrum magnitude or log magnitude.

What if the input time signal is a random process? A common example of this is a Gaussian process. In this case, the magnitude spectrum turns out to be 256 independent random variables with Rayleigh distribution (see appendix). It can be shown that the standard deviation of each bin's contents is about the same as the true value being measured. This obviously poor measurement is inherent in the FFT, regardless of the number of points in the transform. Figure 2 illustrates this with a white noise source as the Gaussian process.

Figure 2.

The spectrum of a 10-millisecond segment of a Gaussian random process



SECTION 8:

Power spectrum averaging

If we watch a particular bin while the 3582A analyzes successive segments of a random process, we will begin to suspect that there is a way to improve the measurement accuracy. While there is considerable variation among the values of amplitude, there appears to be a central tendency (or mean value) for the data. More exactly, if we take K measured values and calculate their average (mean), and then repeat this several times, we will have a collection of numbers (the individual averages) whose variance is considerably less than that of the individual data points. Or, putting it another way, the average of K independent measurements is a better statistical "estimator" of the true magnitude of a bin than any single measurement. The accuracy of the average improves (in the sense that the variance decreases) as K gets larger. This kind of average is called a power average, since it is derived from a series of power spectra.

Since employing large values of K can be very time-consuming, especially in narrow-band analysis, it is necessary to know what accuracy to expect for a given value of K . In dealing with random data, a useful form in which to present such accuracy specifications is the "confidence interval." This is a chart or table listing a numerical interval to be attached to a measured value. With a given confidence, the true value can be stated to be within the interval. For instance, a 90% confidence table states that, in 90% of the measurements, the true value will lie within the interval given. Such a table and an example of its use is given in appendix 2 to aid 3582A users in choosing values of K appropriate to their accuracy needs.

Power spectrum averaging compared with video filtering. Many readers will have had experience with conventional swept spectrum analyzers. In these instruments, post-detection low-pass filtering (video filtering) is provided to smooth the measurement of random process spectra. As it happens, this technique is directly related to power averaging in digital analysis, and it is interesting to examine the relationship briefly.

Conventional analyzers: Assume the input to the analyzer is a random process whose power spectrum is nearly constant across the analyzer bandwidth (this assumption is necessary so that the estimate will be "unbiased"; that is, tend in the limit to the true value of the spectrum). Then the spectrum estimate will have this statistical accuracy:

normalized std. dev. of the estimate of the amplitude spectrum =

$$\frac{1}{\sqrt{B_a T_a}}$$

where

B_a = analysis bandwidth

T_a = effective averaging time, equal to two time constants for single pole filters

Digital analyzers: Using the same assumption as above, the spectrum estimate of a single bin is:

normalized std. dev. of the estimate of the amplitude spectrum =

$$\frac{1}{\sqrt{B_d K T_d}}$$

where

B_d = bandwidth of one bin $\geq 1/T_d$

K = number of records averaged

T_d = length of time record

Comparison: By comparing these results on the basis of equal analysis bandwidths ($B_a = B_d$), it is plain that equal averaging times ($T_a = K T_d$) produce statistically equivalent measurements of the amplitude spectrum at each frequency. **However, the N-point FFT makes N such estimates, covering a total analysis range of $N B_d$, in the same time that the conventional analyzer makes one estimate.**

SECTION 4:

Time averaging

When power averaging is applied to a mixed random process, the deterministic portion of the signal is unaffected, since its variance is zero to start with. **Power averaging will smooth only the estimate of the random portion of the spectrum.** It will not, for instance, uncover the deterministic spectrum if it is "buried in the noise."

When to use time averaging. If there exists an independent signal, free from noise, which is synchronous with the periodic part of a mixed random process, then we can use another kind of averaging, called "time averaging" in the 3582A, which will enhance signal-to-noise ratios. Of course, the need for the synchronizing signal is rather restrictive, although there are numerous situations in which one is available. For example, in biological stimulus-response measurements, the stimulus signal itself will serve as the synchronizer while analyzing the noise-contaminated response.

How it works. The principle of time averaging is straightforward. The operation may be explained from the point of view of either time or frequency domains, although perhaps the time domain view is intuitively clearer: select K equal length intervals of a mixed random process. The intervals must be chosen so that the first point of each occurs at the same position in the cycle of the periodic component. If the corresponding points of all the intervals are added together, and then divided by K to produce an average, we can deduce the following about this averaged waveform:

- 1) The amplitude of the normalized periodic component is $(1/K)K = 1$ times the amplitude of the periodic component in one interval. This is because the synchronized components add directly.

- 2) The random components, being uncorrelated, add on an RMS basis; their normalized sum is $(1/K)\sqrt{K} = 1/\sqrt{K}$ times the amplitude of the random component in one interval.

Thus, in the average, the ratio of the periodic component to the random component is

$$\frac{1}{1/\sqrt{K}} = \sqrt{K}$$


times higher. This means a S/N improvement of $20 \log \sqrt{K} = 10 \log(K)$ dB. The maximum K is 256 in the 3582A, so the S/N improvement for time averaging approaches 24 dB.

Time averaging does *not* reduce the normalized standard deviation in measuring the random portion of the spectrum. In the averaged waveform, the ratio of the standard deviation to the mean value (of the random portion) is the same as that of a single measurement. Therefore, if one wants both improved measurement of the random portion and enhanced S/N ratio, both spectrum and time averaging should be used in succession. This is possible either for stored time data, or for signals whose statistics don't change with time. Of course, a synchronizing signal must be used for the time average.

Time averaging can also be performed on some mixed random processes whose deterministic components are not strictly periodic. This is most easily seen in the case of transient analysis in which the transient is stimulated by a non-periodic signal. It is necessary that the transient diminish to insignificance between applications of the stimulus so that the averaging process will be performed on identical samples of the deterministic signal. See Figure 3 and the example in Section 7.

Figure 3.

Time averaging on a repetitive but not periodic transient



Each interval T begins at the same point in the waveform. An actual example of this is given in Section 7.

SECTION 5:

Summary of averaging properties

The important properties of the two forms of signal averaging in the 3582A are listed here in summary:

Power spectrum averaging

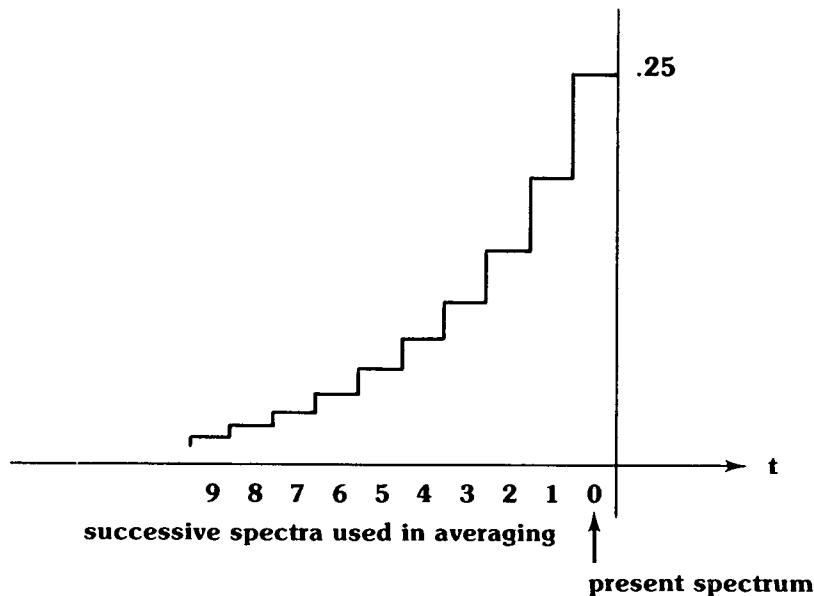
- a) Power averaging is applicable to either pure or mixed random processes.
- b) Power averaging reduces the inherent variance resulting when the FFT is used to determine the spectrum of a random process. Thus, a power averaged spectrum is a statistically more accurate estimate of the true spectrum.
- c) When applying power averaging to a mixed random process, the deterministic portion of the signal is unaffected. Thus, the S/N ratio (the ratio of the deterministic to the random parts of the signal) is *not* changed.
- d) No synchronizing signal is needed.
- e) Phase is not available in the power average routine. (The phase information given by the 3582A is calculated with another routine; see Section 6.)

Time averaging

- a) Time averaging is useful only with mixed random processes (the signal must have a deterministic component).
- b) Time averaging increases the ratio of the deterministic to random portions of the signal (i.e., it improves the S/N ratio).
- c) The normalized standard deviation of the random portion of the spectrum is unchanged.
- d) A synchronizing signal, in fixed relation to the deterministic portion of the signal, must be used with time averaging.
- e) In the 3582A, the results of time averaging on a signal may be displayed in both time and frequency domains.
- f) A time averaged spectrum is complex; both amplitude and phase spectra are derived from the same linear averaging algorithm.

Figure 4.

Relative weighting of spectra in exponential average routine



SECTION 6:

Characteristics of 3582A averaging routines

Power spectrum averaging. The 3582A computes a power average as follows: the FFT produces both a real and an imaginary spectrum component at each analysis frequency. These components are squared and added. For each successive transform, the same operation is performed, and a cumulative sum is maintained for each data bin. Thus, for a K-sized average (K is a power of 2, chosen between 4 and 256), each bin contains the sum of 2K squared values at the end of processing.

This sum is then divided by K, which converts it to the power spectral value for that bin. This process, performed independently for each of the 256 bins, generates the power spectrum average. Before display, the square root of each sum is extracted to produce the amplitude spectrum. The main reason for this step is convenience of units; for instance, volts is usually more appropriate than (volts)². Because of the square root operation, the control button which calls up this averaging routine is labeled "RMS" rather than "POWER."

With the power averaging routine, there is also available a phase display, but it is not a power-averaged quantity, since the power spectrum is phaseless. Rather, for each transform, the phase angle of each bin is computed conventionally as

$$\text{phase} = \tan^{-1} \left(\frac{\text{imaginary component}}{\text{real component}} \right)$$

and the K resulting numbers are simply averaged for the display.

While power averaging is proceeding, the user can watch the interim averages on the display. (However, the averaging process will proceed more quickly if the display is turned off, since this eliminates the need for intermediate root-taking and display formatting.) Also, if further averaging seems desirable after the K spectra are averaged, pressing a higher-numbered button will continue the process to the new value of K.

Exponential averages. Another variation of power averaging is the "moving" or exponential average. As the name suggests, this form of average gives more weight to the most recent measurements. In the 3582A, moving averages are calculated by weighing the latest spectrum by 1/4 and adding it to the previous average, weighted by 3/4. The result is that the Mth spectrum before the present one is given the weight $1/4(3/4)^M$. The standard deviation of an estimate from this averaging routine is about 8 dB less than a single FFT spectrum estimate. The weighting function is shown in Figure 4.

Moving averages are especially useful in cases where the random process is not stationary; that is, when the mean and/or variance of the random process changes with time.

Time averaging. The 3582A computes this type of average as follows: K successive, *synchronized* time records are added, and the sum is divided by K (K is chosen as a power of 2 from 4 to 256). The result is the time domain average. Both the time average and its transform are available for display, so that the result of the S/N enhancement may be observed in either time or frequency domains.

Two methods are available for applying the necessary synchronizing (triggering) signal used with time averaging:

- a) Internal triggering. The trigger signal is applied to Channel A, using the polarity and level controls to establish reliable triggering. (The internal trigger circuitry is only connected to Channel A.) Then the signal to be averaged is connected to Channel B.
- b) External triggering. The TTL-compatible trigger signal is applied to the rear panel connector, and the adjacent switch is set to "ext. trig." Then any combination of channels may be used for the signal(s) to be averaged.

Time averaging is a relatively fast procedure, requiring only slightly more time than that necessary to acquire the K time records. This is because only one FFT operation is performed, rather than K, as in the case of power averaging. (In the 3582A an FFT takes about 350 milliseconds.)

There is no exponential (moving) average routine for the time average procedure. Attempting this causes an error message to be displayed on the CRT.

Peak "averaging." This procedure is not truly averaging, but rather a peak holding process: K transforms are made (K is 4, 8, . . . , 256), and each set of data is compared, bin for bin, with the previous set. The larger number is then retained, with the result that, at the end of the procedure, each bin contains the largest data value encountered during the processing of the K signal segments.

Applications include noise monitoring, measurement of frequency drift, and the like. If continuous monitoring is wanted (that is, no limit on the size of K), pressing the "EXP" key will cause peak averaging to continue indefinitely.

SECTION 7:

Examples

In this section we have included the results of several actual measurements using the 3582A. These were chosen in order to illustrate the material discussed so far, and enough information is included to enable the reader to try similar experiments if he chooses. Each example includes a measurement block diagram and photos of the 3582A display screen, as well as relevant control settings and discussion.

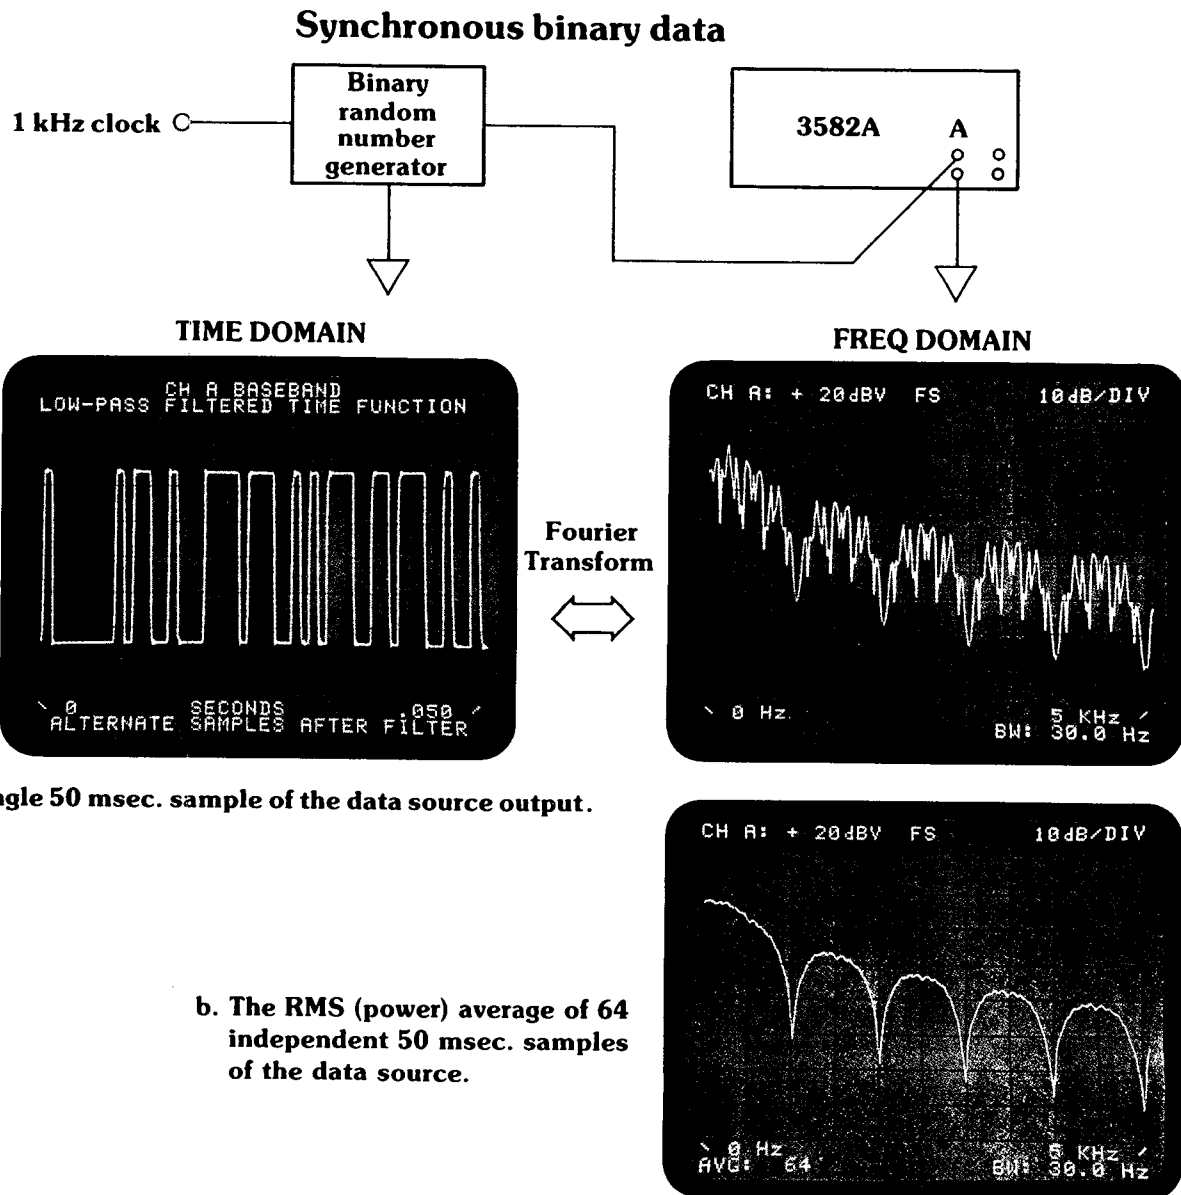
Examples of RMS (power) averages. In Section 3 we stated that averaging a number of FFT spectra of a random process gives an estimate of the true spectrum which is more accurate than any single transform. Figure 5 illustrates this point. The signal being analyzed is a random binary data stream, whose transitions are clocked at a 1-kHz rate. A 50-bit segment of the signal is

shown in (a), along with the Fourier transform of the segment. Although the spectrum is an accurate frequency domain representation of that *particular* sample of the entire signal, one sample alone gives a poor and misleading indication of the characteristics of the whole process. A much better estimate of the spectrum of the process is shown by the power average of 64 spectra, in (b).

The control settings can generally be inferred from the self-documenting display. Other data of interest are:

- The Hanning passband was used for good resolution.
- For the single sample; the REPETITIVE button was out (off) and a single record was captured by pressing ARM.

Figure 5.



Another power averaging example.

There are some random processes whose statistics vary considerably over short intervals of time, but which are more stable (statistically "stationary") in the long run. Human speech is a good example. In the experiment of Figure 6, the object was to determine quantitative spectral differences between adult male and female voices. A common-speech script was chosen, which each speaker read for about two minutes, enough time to process 256 spectra. The individual spectra varied widely; some corresponded to time segments between words and had little energy. After 40 or 50 spectra were averaged, however, the long-term trends became evident. Several 256-spectrum averages from the same speaker showed differences of less than 3 dB.

In performing the experiment, the first step was to acquire and process the time records from the first speaker. After completing the RMS averaging, the display amplitude reference level was shifted 30 dB higher

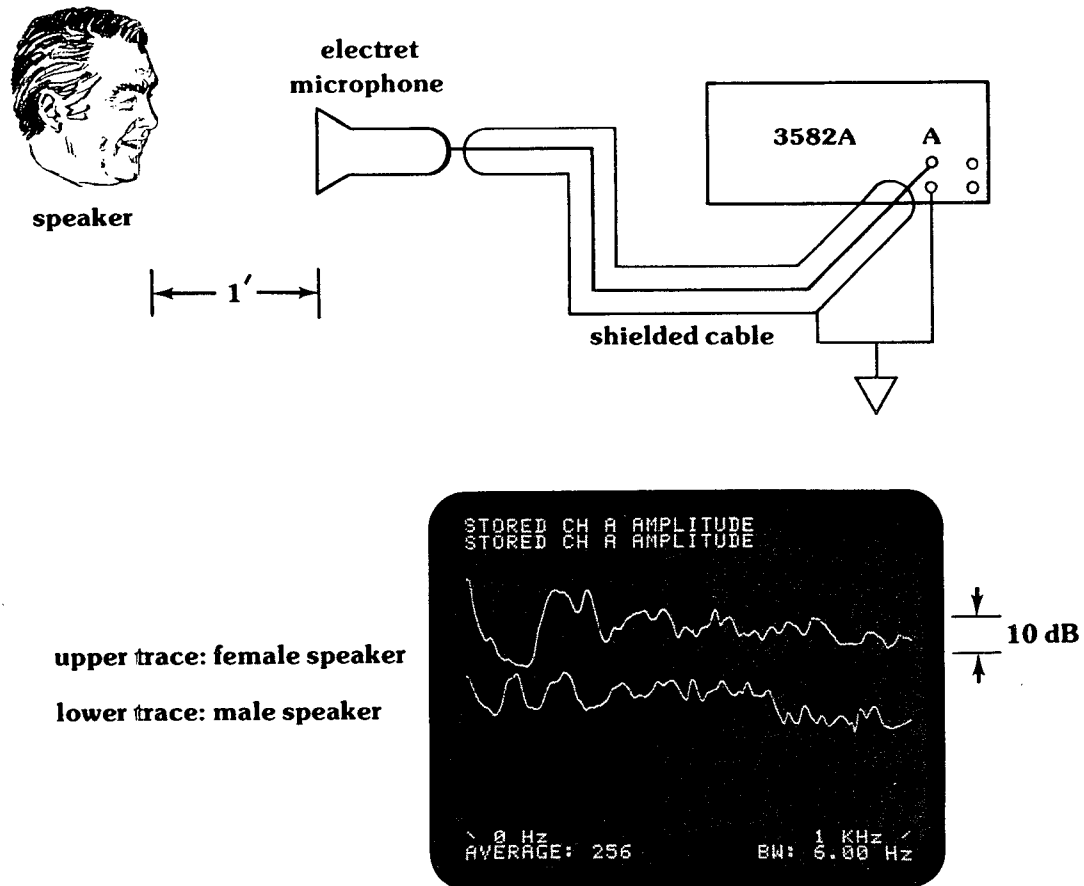
(to separate this data from the next, when displayed simultaneously) and then stored in TRACE 1. The second speaker's voice was then averaged in the same way. When this was completed, the two traces were displayed together for comparison. Hanning passband shape was used throughout the experiment.

Examples of time averaging. As we discussed in Sections 1 and 2, a principal feature of time averaging is the use of a synchronizing signal to insure that each time record used in the average contains the deterministic waveform in the same relative position.

Figure 7 is the block diagram and measurement results showing how a signal can be extracted from noise by this technique. Since Channel A in the 3582A is the only one from which an internal trigger may be derived, it is used in this experiment to trigger the data acquisition. The analysis is carried out in Channel B, to which the noisy signal is connected.

Figure 6.

Human voice spectrum



traces were separated 30 dB for clarity

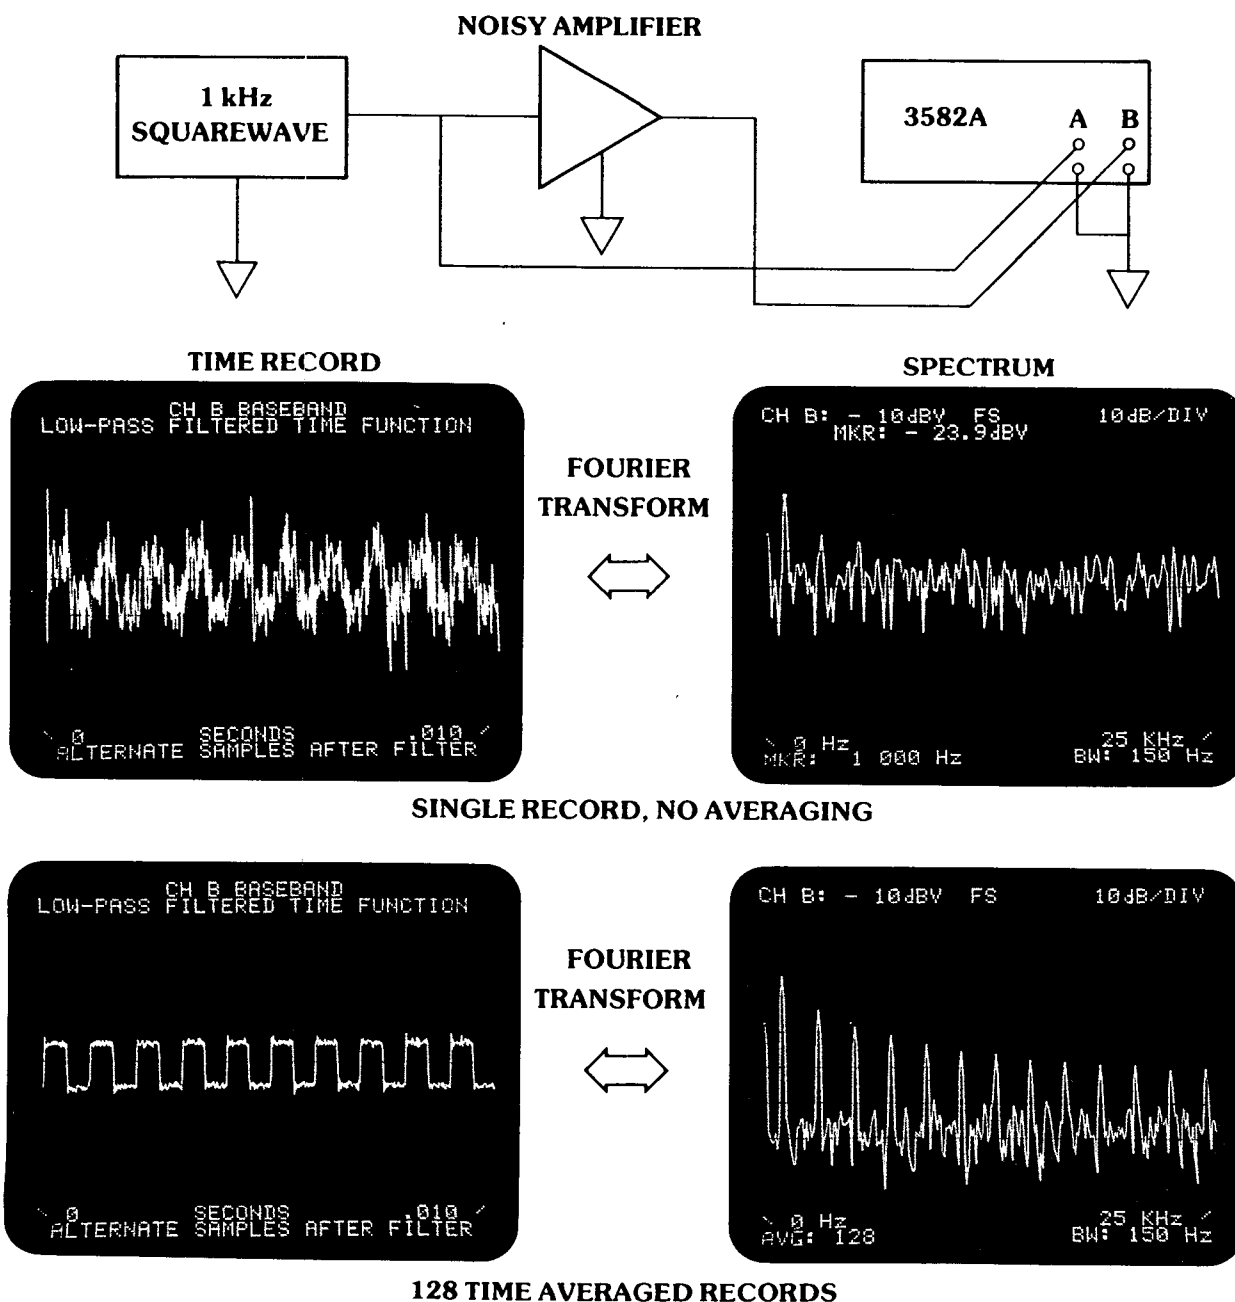
The procedure followed was:

- a) Set up the controls as follows:
 Display: Amplitude A, 10 dB/div, ref. level normal
 Passband shape: Hanning
 Average: off
 Marker: off
 Span: 0-25 kHz
 Trigger: + slope, repetitive
 Input: A, AC coupling A & B
- b) Connect the pure squarewave to A. Holding in the Time A button, adjust Channel A Sensitivity for about half scale display. Then adjust Trigger level for reliable sync.
- c) Switch display and input mode to B. Connect

- d) the noisy signal to Channel B. Adjust B Sensitivity as in (b).
- e) At this point, the single record photographs were made by momentarily turning off the Repetitive control (button out).
- f) In the Average block of controls, push Time, 8/128, and Shift keys. This starts the averaging process.

Some interesting observations may be made from these results. First, the spectrum photos show a noise reduction of roughly 20 dB, which agrees well with the theoretical value of $10 \log 128$. Second, the noise is not "smoothed" by time averaging: its relative standard deviation appears about the same in both the single and the averaged spectra.

Figure 7. Squarewave plus noise



Self-Synchronizing

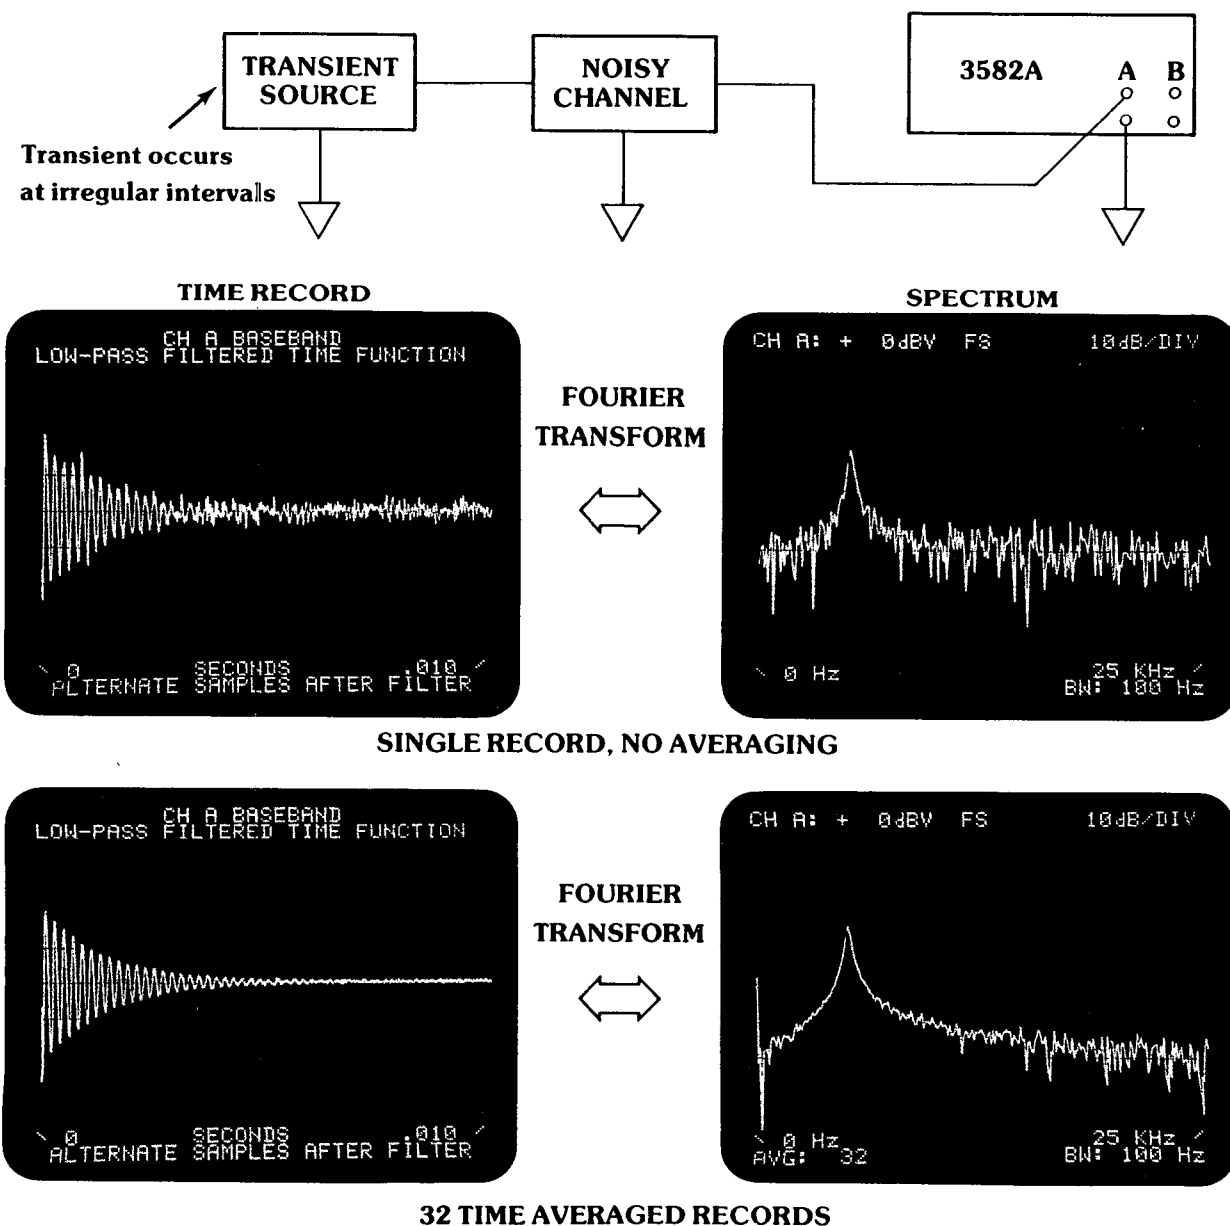
Sometimes the required synchronizing signal can be derived from the signal to be averaged. This is true when there is a portion of the deterministic waveform which is large compared with the peak noise. Such a situation is shown in Figure 8 which also demonstrates that time averaging can be used with a non-periodic signal. In this case, a tuned circuit is impulsed at irregular intervals, resulting in an identical transient each time. The negative leading edge of the transient was sufficiently higher than the noise to serve as the trigger. Control settings were similar to the previous example, except that Channel A was used for both trigger and analysis, and 32 averages were taken. Also, the uniform passband was used, as is normal with transient analysis. (The uses of the three available passbands are explained in the operating manual.)

A test to see whether a signal can self-trigger in time averaging is to observe the time waveform after some averaging has occurred; if the deterministic waveform seems to diminish or change form as averaging proceeds, there is too much jitter, and a separate trigger must be used.

Example of peak averaging. As we mentioned already, peak averaging is not truly an averaging process; rather, it is a means of comparing successive spectra and saving the largest amplitudes encountered at each analysis frequency.

A popular use of this feature is monitoring the frequency drift of some device which nominally operates at a constant frequency. For instance, a motor's speed may vary due to load, temperature, etc. Using a tachometer as a speed-to-frequency transducer, the peak average routine of the 3582A will reveal the maximum excursions of the speed over the test time.

Figure 8. Transient signal in noisy channel



The example used here is similar. An electronic signal generator varies in frequency in a slow, periodic way. When this action is deliberate, it is called a "sweep generator." In Figure 9 such a source is shown connected to the 3582A for analysis of its peak-to-peak frequency excursion. Since the source was a programmable frequency synthesizer whose excursion could be accurately set, the experiment was really intended to check the 3582A. The result, shown in the spectrum photo, is accurate. Using the marker dot would improve the accuracy to 2 Hz resolution.

Some information on the setup is:

- a) The Hanning passband was used, since its narrow peak is more useful for frequency measurements than the broader peak of the flat top passband.
- b) The 3582A was operated to process an unbounded number of spectra; this is done by pressing both 32/EXP and SHIFT keys in the averaging section. When the test was stopped, more than 5000 spectra had been examined.

One word of caution for this kind of measurement: the signal being analyzed should not change frequency too fast, or the analysis will be smeared and inaccurate. A good rule is that the frequency change should not exceed 2% of the analysis span during the time record. For the 3582A, this rule can be formulated:

$$\text{frequency rate-of-change} < \frac{(\text{analysis span})^2}{12500} \text{ Hz/sec}$$

The experiment of Figure 9 met this criterion, since

$$\text{frequency rate} = \frac{2 \times (22150 - 21950)}{33.3} = 12 \text{ Hz/sec}$$


is less than

$$\frac{(500)^2}{12500} = 20 \text{ Hz/sec}$$

The rule is not as restrictive as it sounds. Higher sweep rates cause the appearance of distinct sidebands in the analysis. The frequency excursion can then be calculated from frequency modulation theory. This is beyond the scope of the application note, however.

Figure 9.

Slowly swept source




Sweep generator settings:

$$f_{min} = 21950 \text{ Hz}$$

$$f_{max} = 22150 \text{ Hz}$$

$$\text{sweep rate} = 33.3 \text{ sec/sweep}$$

Peak hold "averaged" spectrum



5,162 separate spectra were measured and compared to derive this peak value envelope.

APPENDIX 1:

Probability distributions of power spectrum estimates

Gaussian random processes. If the input time signal to an FFT analyzer is a Gaussian random process, the output frequency variables are also Gaussian. This is because the Discrete Fourier Transform is a **linear** operation on the input. It can be shown that, at each frequency, there are two independent Gaussian random variables, which are the real and imaginary spectral components at that frequency. The two components have zero means and identical variances. From this fact, we conclude that the spectral power is equally divided between the real and imaginary spectral components.

The averaged power spectrum as a random process. The power at any frequency is the sum of the squares of the real and imaginary spectral components. This sum itself is a random process, being a function of two independent random processes. However, the prob-

ability distribution is no longer Gaussian, but "chi-squared." The sum of K independent, zero mean, unity variance, *squared* Gaussian random variables is the chi-squared variable of order K . The square root of the second-order chi-squared variable is called the Rayleigh variable. Hence the magnitude spectrum of the D.F.T. is Rayleigh distributed. In the power spectrum averaging procedure described in Section 3, the average of K spectra is computed from the sum of $2K$ squared components. From the above discussion, it is apparent that the probability distribution of the sum is chi-squared, of order $2K$. Since the chi-squared variable is a standard, tabulated quantity, it is possible to calculate whatever statistical parameters one needs to know. The confidence table below was calculated on the basis of the 0.05 and 0.95 tails of the appropriate chi-squared distribution.

APPENDIX 2:

90% confidence limits for power averages

To use the table, first decide on the allowable statistical tolerance in dB. Then find the number of averages whose 90% limits are within the tolerance bounds. For instance, if we can tolerate a ± 2 dB accuracy band, then the 16-average routine is what we need. Remember that the limits given are statistical, not absolute. That is, they state that, *on the average*, the true value of amplitude will lie within the stated bounds in 9 out of 10 measurements.

	K = number of averages						
	4	8	16	32	64	128	256
Upper limit dB	+4.7	+3.0	+2.0	+1.4	+1.0	+0.7	+0.5
Lower limit dB	-2.9	-2.2	-1.6	-1.2	-0.8	-0.6	-0.4

As an example of use, suppose that a random noise source has been measured, and that 32 spectra have been power averaged. Using the marker readout, the 1000 Hz bin shows a signal level of -55 dBV. The table can be interpreted in this case to indicate that the true signal amplitude has a 90% probability of being in the range of -53.6 dBV to -56.2 dBV.

Bibliography

1. "Random Data: Analysis and Measurement Procedures," J. S. Bendat and A. G. Piersol, John Wiley 1971. This is a fundamental text covering both analog and digital analysis of random processes. Chapters 3 and 4 provide the background for understanding the nature of random processes and measurements made on them. Chapter 6 details the errors encountered in several kinds of measurements on random processes.
2. "Digital Time Series Analysis," R. K. Otnes and L. Enochson, John Wiley 1972. This deals specifically with computer analysis of random processes.



For more information, call your local HP Sales Office or East (301) 948-6370 • Midwest (312) 255-9800 • South (404) 955-1500 • West (213) 970-7710. Or write: Hewlett-Packard, 1501 Page Mill Road, Palo Alto, California 94304. In Europe: P.O. Box 85, CH-1217 Meyrin 2, Geneva, Switzerland. In Japan: YHP, 1-59-1, Yoyogi, Shibuya-Ku, Tokyo, 151.

Printed in U.S.A.

5952-8767