HEWLETT-PACKARD JOURNAL

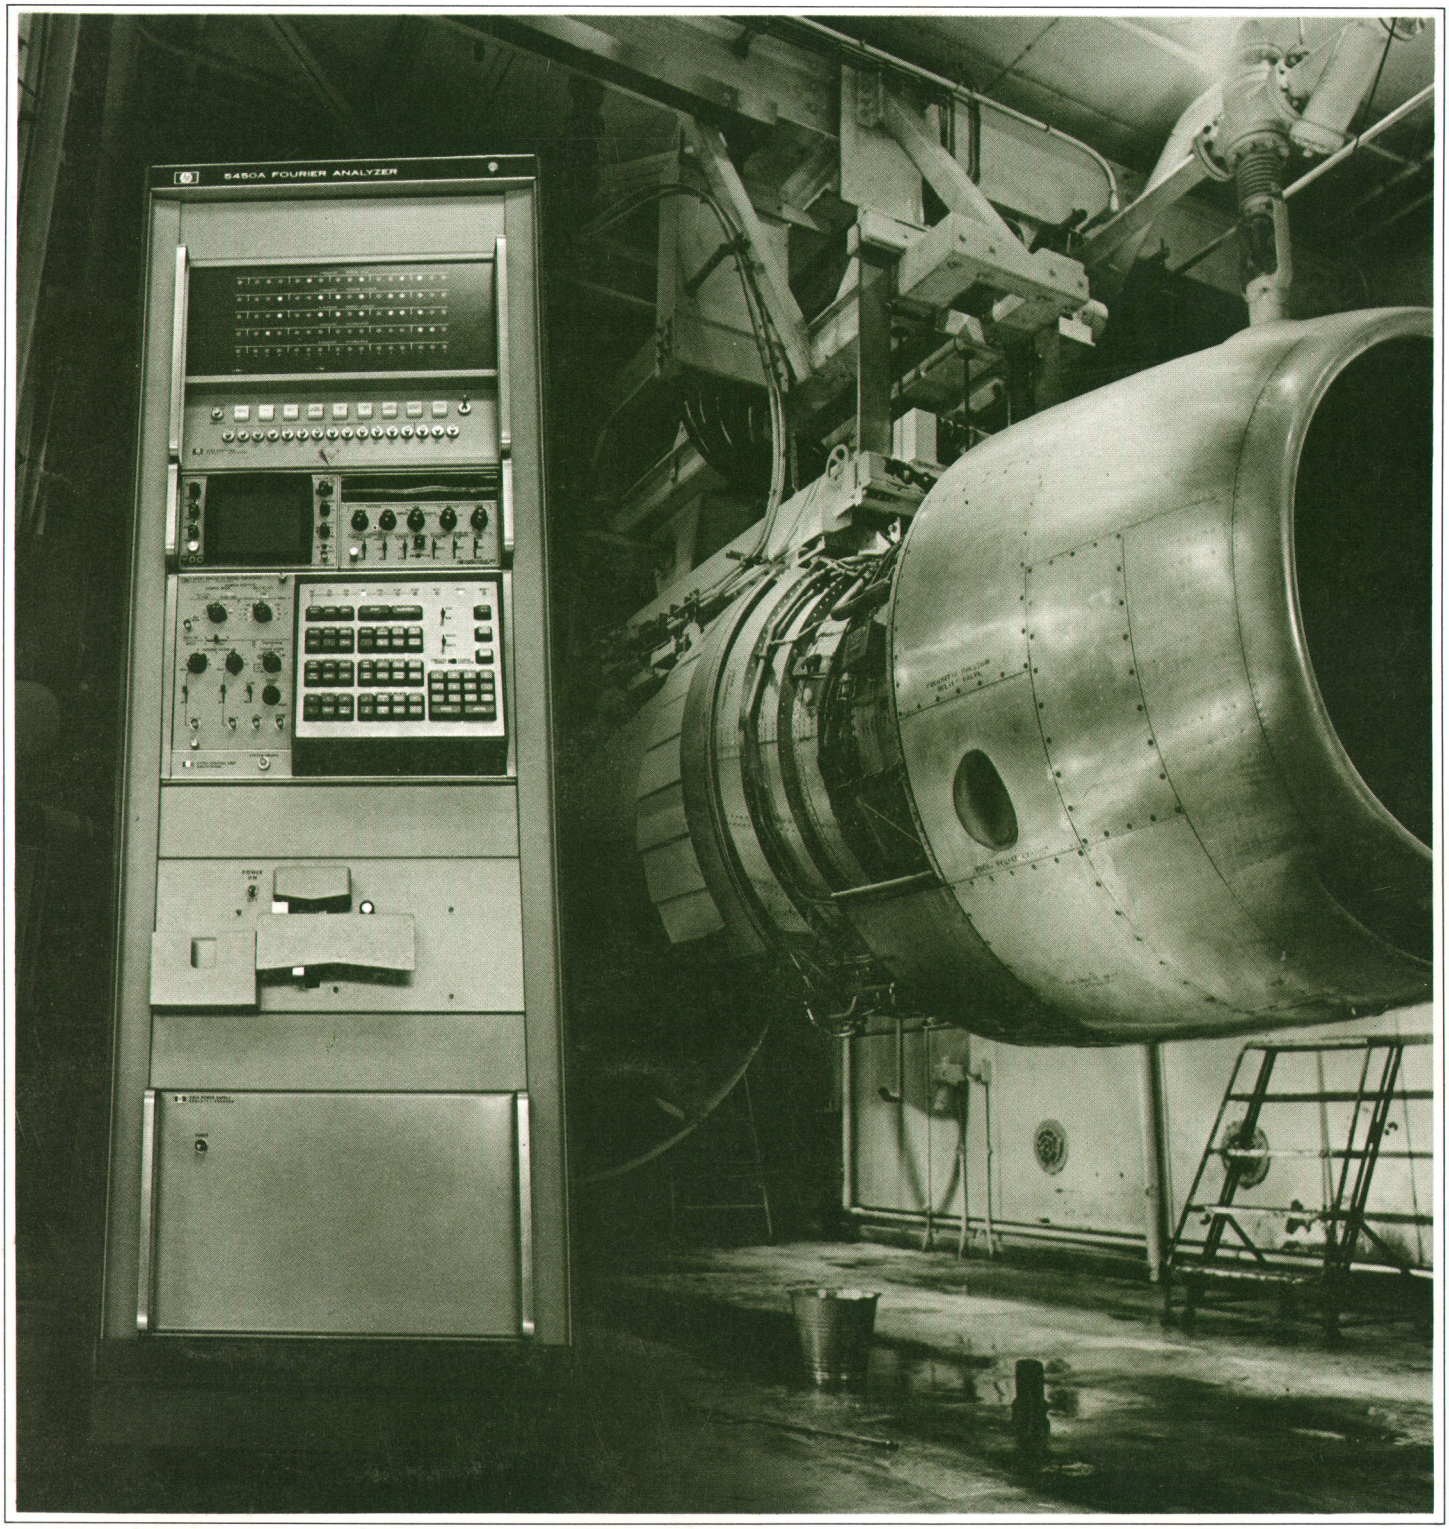

JUNE 1970

Digital Fourier Analysis

Some of the theoretical and practical aspects of measurements involving Fourier analysis by digital instrumentation.

By Peter R. Roth

WHEN THE CHARACTERISTICS OF A SIGNAL OR SYSTEM ARE MEASURED, the measurements most often made are the spectrum of the signal and the transfer function of the system. For example, if the transfer function of the landing gear and wing structure of an aircraft is known, and if the spectrum of the vibrations from typical runways can be determined, then the roughness of a landing can be evaluated. Or if the spectrum of the vibrations caused by typical roads can be determined, an automobile suspension system may be designed and tested to maximize ride comfort.

It is the questions of how to measure spectra and transfer functions, especially when signals more complex than simple sine waves are involved, that we will examine in this article.

The techniques to be described are based upon computation of the Fourier integral

$$S_x(f) = \int_{-\infty}^{+\infty} x(t) exp\{-i2\pi ft\} dt. \qquad (1)$$

While in principle the methods that will be examined are not new and have been partially implemented using analog instruments, their full development has waited on the availability of digital processors with sufficient speed and flexibility.*

How does computation of the Fourier integral help us make meaningful measurements? Consider the Fourier transform written in its sine-cosine form:

$$S_x(f) = \int_{-\infty}^{+\infty} x(t)\{\cos 2\pi ft - i \sin 2\pi ft\} dt. \qquad (2)$$

This equation states that the transform averages a time function input x(t) with a set of sines and cosines to determine the content of x(t) at some frequency f. Thus the transform resolves the time function into a set of components at various frequencies much as a set of analog filters would. However, it not only yields the amplitude at each frequency, but also resolves the in-phase (real, cosine) component and the quadrature (imaginary, sine) component, thereby giving magnitude and phase information which is difficult to obtain in any other way.

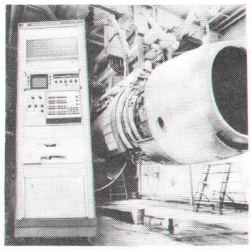

Cover: *Model 5450A Fourier Analyzer makes a variety of measurements fundamental to the analysis of waveforms and time series. Among its many uses is analysis of sound and vibration, symbolized here by a jet engine mounted in a test chamber. We are grateful to United Air Lines for allowing their engine to pose for us.*

In this Issue:

* The new HP Model 5450A Fourier Analyzer is one of these digital instruments. See article, page 10.

PRINTED IN U.S.A.

The Fourier transform is also valuable when it is applied to measurements on systems. The result of the operation of a linear system on any input signal in the time domain may be determined from the convolution of the system impulse response h(t) with the input signal x(t) to give the output y(t):

$$y(t) = \int_{-\infty}^{+\infty} h(\tau)\, x\,(t - \tau) d\tau. \qquad (3)$$

Visualizing the result of this operation is all but impossible for anything other than a simple case. But if the Fourier transform is applied to this convolution integral, a simple, easily understood relationship results. The output spectrum S_y is the product of the input spectrum S_x and the transfer function H:

$$S_y(f) = S_x(f) \cdot H(f). \qquad (4)$$

The simplest implementation of a measurement technique based on this relation is the use of a sine-wave input for x(t). Since the sine wave contains but one frequency component it provides a simple way of measuring the transfer function using voltmeters and phasemeters. However, not all systems may be measured using sine waves, either because there is no way of inserting such a signal into the system, or because the sine wave is not a realistic signal form.

A more general measurement method is to measure the input and output time series, in whatever form they may be, and to calculate H using S_x, S_y, and the Fourier transform. This method has several advantages, as I will show. But first, because the most powerful computational techniques available today are digital, it's necessary to say something about the nature of the Fourier transform when it is implemented on a digital processor.

Digital Fourier Transforms

Digital techniques make us realize very clearly that all measurements are discrete (i.e., have finite resolution) and of finite duration. All digital memories are obviously discrete and finite in size. Therefore, the equation for the Fourier transform must be changed to a finite sum for digital processing. This means, first of all, that the time function to be transformed must be sampled at discrete intervals, say Δt. It also means that only a finite number, say N, of such samples may be taken and stored. The record length T is then

$$T = N\Delta t. \qquad (5)$$

The effect of finite Δt is well known; it limits the maximum frequency that may be sampled without 'aliasing' error to

$$f_{max} = \frac{1}{2\Delta t}. \qquad (6)$$

Any components above this Nyquist frequency or its multiples are folded back onto frequencies below f_{max}. In practical measurement situations this aliasing presents little or no difficulty, since f_{max} can be chosen to include all significant components of the input signal, or a filter may be used before the sampler to eliminate any strong components above f_{max}.

The effect of finite record length T is also important. When a Fourier integral is taken over a finite record length T the result is a Fourier series, and the spectrum has discrete lines and finite resolution. A Discrete Finite Transform (DFT), which must be used whenever a Fourier transform is computed digitally, is more like a Fourier series than a transform, since it assumes that the input is periodic in the interval T and has a spectral resolution of

$$\Delta f = \frac{1}{T}. \qquad (7)$$

The DFT is written as

$$S_x(m\Delta f) = \frac{1}{N}\sum_{n=0}^{N-1} x(n\Delta t)\, exp\{-i\frac{2\pi}{N}mn\}. \qquad (8)$$

It yields in the frequency domain $\frac{N}{2}$ real (cosine) components and $\frac{N}{2}$ imaginary (sine) components from a sampled time record of N points. I will refer to this result as the *linear spectrum* to keep it sorted out from certain other spectrum forms.

While this raw form of spectrum has certain uses, it it of limited value because of its dependence on the time position of the input record. A waveform of constant shape will always have the same energy at any one frequency, but how this energy is distributed between the sine and cosine terms depends on the phase shift or time position of the waveform. Fig. 1 gives an example of this. Fig. 1a is the real part of the linear spectrum of a square pulse. It has the expected sin x/x form. However, the real part of the linear spectrum of the same pulse delayed a small amount, Fig. 1b, does not have this form. The linear phase shift given the spectrum by delaying the waveform has changed the distribution of the spectrum between its real and imaginary parts. On the other hand,

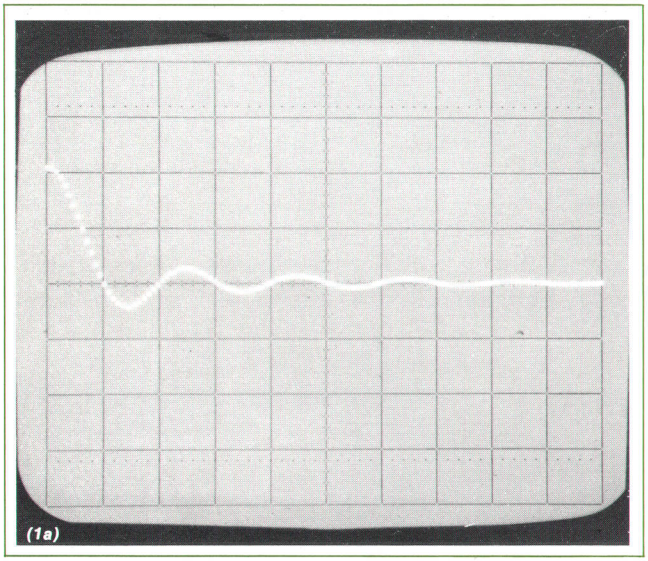

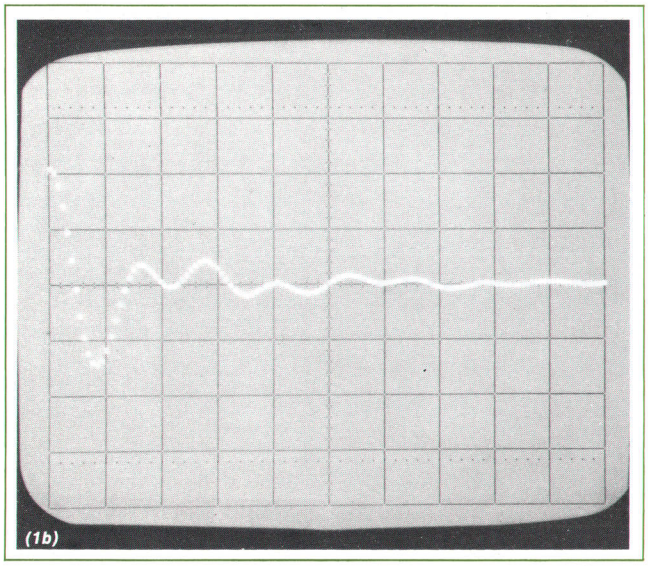

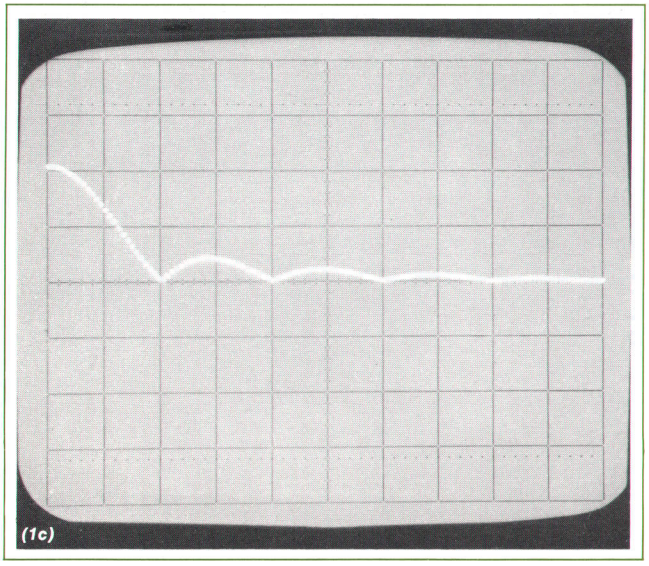

Fig. 1. *The real parts of the Fourier transforms of two rectangular pulses which are identical except for a time shift are shown in (a) and (b). (c) is the magnitude of the Fourier transforms of both pulses. Because it is independent of time position, the magnitude-and-phase form of the transform is more useful than the real-and-imaginary form. However, the square of the magnitude, or the power spectrum, is even more widely used.*

if we examine the *magnitude* of either the undelayed pulse spectrum or the delayed pulse spectrum, we see that it is constant (Fig. 1c) and that the energy in any line is the same no matter what the time position of the input waveform is.

It is clear, then, that to obtain a constant linear spectrum independent of time position it is at least necessary to convert the real and imaginary components of the spectrum into magnitude and phase. While the linear magnitude spectrum is a valid and perfectly acceptable way to achieve a useful spectrum it is cumbersome from a computational standpoint. A closely related function, the 'auto' spectrum or 'power' spectrum, gives the same basic information, is faster to compute, and can be applied to measurements which the linear magnitude spectrum cannot.

Power or Auto Spectrum

The auto spectrum, $G_{xx}(f)$, is formed by multiplying the value of the linear spectrum, $S_x(f)$, by its own complex conjugate.

$$G_{xx}(f) = S_x(f) \cdot S_x^*(f) = [A(f) + iB(f)]\,[A(f) - iB(f)] \tag{9}$$

$$G_{xx}(f) = A^2(f) + B^2(f) \tag{10}$$

Each spectral line of $G_{xx}(f)$ is proportional to the voltage squared at frequency f, or more exactly to the variance of the input waveform at frequency f. The auto spectrum is useful because it is the magnitude squared of the linear spectrum. For this reason, and because it has no imaginary part, it is independent of the time position of the input waveform. It is the square-law auto spectrum that is usually implied when the term 'spectral analysis' is used.

Analyzing Random Signals

The auto spectrum, because of its independence of time and phase, is a useful tool for analyzing signals that are deterministic, that is, for signals that do not change in spectral form from sample record to sample record, or only change in a predictable way. However, the auto spectrum is an even more useful tool for the analysis of signals that are stationary and random, that is, signals whose spectra will vary from sample record to sample record but will have a measurable mean or average value. Many processes generate signals whose spectra cannot be predicted for any single sample record, but whose spectra are stable on the average. Examples of such processes are 1/f noise in an amplifier, or the sea state noise in a sonar system. On the other hand, the process being measured may be a combination of deterministic and random spectra. For example, consider the fine-line components of the noise due to the rotating members of a turbojet hidden by the random noise of the combustion, or the tonal components of an acoustic signal hidden in the random noise of the ocean.

What is more important about a random spectrum is that for a single sample record of length T (i.e., of spectral resolution $\Delta f = 1/T$) the spectral lines are just as random as the time series that generated the spectrum no matter how long T is. In basic engineering and mathematical texts on the Fourier transform, the transition from periodic functions whose spectra are described by Fourier series to totally aperiodic functions whose spectra are described by the Fourier transform is made by making the record length T go to infinity in the limit. That this procedure does not work for the ultimate in aperiodic functions, random signals, can be intuitively demonstrated in two ways.

First consider a wave analyzer with a bandwidth Δf and a meter with very small damping. The response time of this analyzer is about $1/\Delta f$ or T seconds. If a random signal is applied to this wave analyzer an independent reading can be made about every T seconds. Now, if the bandwidth of the wave analyzer is cut to $\Delta f/2$, the response time of the filter and hence the time between independent readings becomes 2T. While the meter will move half as fast in this case, the randomness of the reading as expressed by the variance of the independent readings will be unchanged, since independent readings are twice as far apart. Thus no matter how long a record (i.e., how narrow a bandwidth) is used, no improvement in statistical certainty can be made. The only way to improve the reading is to put an integrating circuit on the meter that is much slower than the response time due to the reciprocal of the bandwidth $1/\Delta f$. Then the final reading will be the result of averaging many independent readings.

To show this effect for a DFT consider a spectrum computed from N equally spaced time samples over a sample record of length T, yielding $\frac{N}{2}$ real and $\frac{N}{2}$ imaginary frequency components. From N time points, exactly N values are obtained in the spectrum, and since no new information about the signal is added by the DFT, each spectral line will have no more statistical certainty than a sample point in the time function from which the spectrum was computed. In fact, for a spectrum of Gaussian noise of any spectrum shape, the variance of a spectral line for one sample record is equal to the expected value for the measurement. Such a measurement is so uncertain that it is no measurement at all.

However, if a number of independent samples of the spectrum are averaged, the variance of the resulting estimate of the spectrum will be reduced in a fashion analogous to integrating readings from the wave analyzer meter. Such a case is demonstrated in Fig. 2. Fig. 2a is a spectrum computed from a single sample record of a signal consisting of a sine wave plus random noise. Because the variance of one sample is equal to the expected value for each line it is impossible to tell which of the spikes is the spectrum of the sine wave and which is due to the variability of the estimate. Fig. 2b shows a spectrum computed from an average of 100 samples. Here the variability of the estimate is reduced to the point where it is perfectly clear where the single tone lies. It is also clear what the spectral shape of the Gaussian noise is. In fact, a statistical certainty for the estimate of the random spectrum is easily computed from the relationship that one standard deviation σ is

$$\sigma = \frac{1}{\sqrt{K}}, \tag{11}$$

where K is the number of sample spectra averaged.

For the case of 100 spectral averages 3σ is 1.1 dB. Thus one would expect that only one estimate in a thousand would fall farther away than 1.1 dB from the measured value in Fig. 2b. To achieve this degree of statistical stability using an analog wave analyzer with a 1 Hz bandwidth would require a 100-second integration.

Two Input Waveforms

So far we have considered measurements on one time series only. However, we often have to take measurements from two signals simultaneously so the relationship

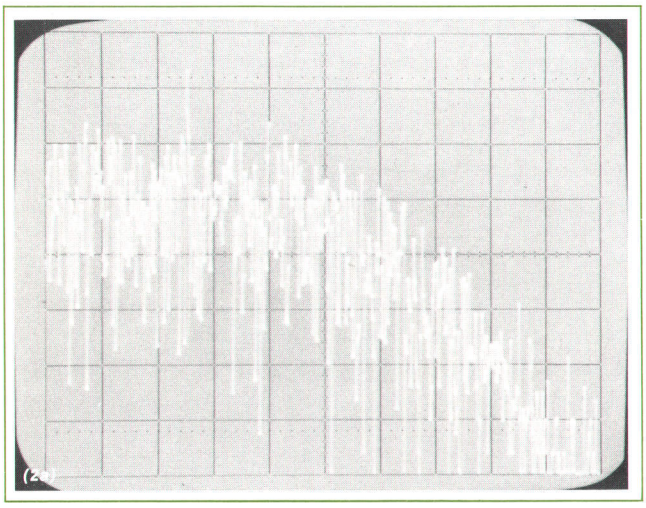

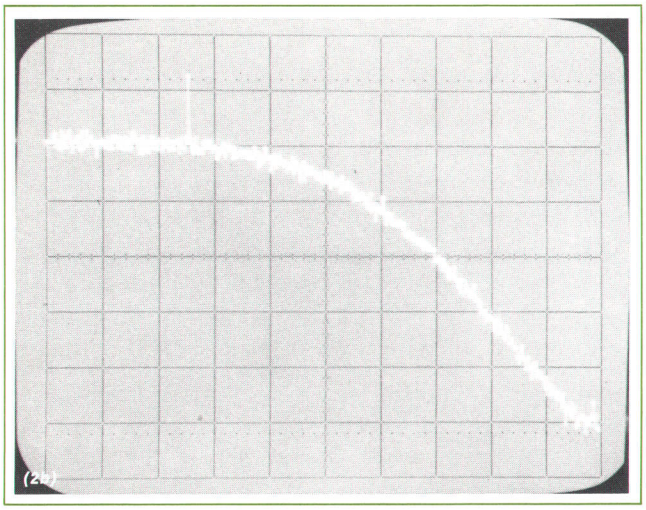

Fig. 2. *The power spectrum (a) computed from a single sample record of a random signal is as random as the signal itself. But when 100 such spectra are averaged, the result (b) shows not only the spectral shape of the random signal, but also that there was a sinusoid hidden in the signal.*

between two points in some process may be determined. For example, in the situation shown in Fig. 3 the relationship between input x(t) and output z(t) might be of interest. There are two distinct quantities that can be measured in such a situation. The first is the degree to which the output depends on the input. That is, is z(t) caused by x(t) or is z(t) due in part to some unrelated signal such as n(t)? Second, if z(t) is caused at least partly by x(t), what is the form of this relationship?

It is important to be aware that, although neither causality nor relationship can exist without the other, each contains different information about the process. It is also important to note that no single measurement of the correlation between two signals, either in the time domain

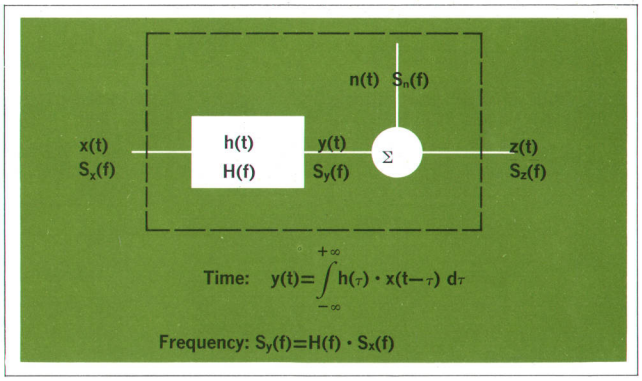

Fig. 3. *Is the output z(t) caused entirely by x(t) or is there also unrelated noise n(t)? What is the form of the relationship between x(t) and y(t)? The transfer function H(f) provides an answer to the second question. A quantity called the coherence function answers the first.*

or the frequency domain, is capable of separating these two quantities.

Cross Spectrum or Cross Power Spectrum

The cross spectrum, also known as the cross power spectrum, illustrates these points. The cross spectrum $G_{yx}(f)$ between two signals y(t) and x(t) in a process or system is formed by multiplying the linear spectrum of y(t) by the complex conjugate of the linear spectrum of x(t) measured at the same time.

$$G_{yx} = S_y S_x^* = (A_y + iB_y)(A_x - iB_x) \quad (12)$$

$$G_{yx} = (A_y A_x + B_y B_x) + i(B_y A_x - B_x A_y) \quad (13)$$

These relationships show that the cross spectrum is not a positive real quantity like the auto spectrum, but in general is both complex and bipolar. A physical interpretation of this function is quite straightforward. If there are components at a given frequency in both x(t) and y(t), the cross spectrum will have a magnitude equal to the product of the magnitudes of the components and a phase equal to the phase difference between the components.

While this interpretation is exactly true when x(t) and y(t) are uncontaminated by noise, an additional dimension must be added when unrelated signals are added to the process. Any single sample of the cross spectrum G_{zx} between the output z(t) of the linear system of Fig. 3 and the input x(t) will show the combined effects of x(t) and n(t) merged into z(t). However, if n(t) is unrelated to x(t) (i.e., random and uncorrelated), its contribution to the magnitude of G_{zx} will not have a constant phase from sample record to sample record as will that of x(t). If many sample records are averaged, the random

phase of the contribution of n(t) will ultimately cause it to have a negligible contribution to the cross spectrum.

How many independent samples of the cross spectrum it takes to achieve a result of a given accuracy cannot be determined without some further information beyond the cross spectrum itself. Also required is information about the relative contributions of the various signals to the measurement. This makes an important point: *a simple cross spectrum measurement does not differentiate between causality and relationship*. Without more information than is contained in the simple cross spectrum it cannot be determined if a high value in a cross spectrum is due to a strong gain of the measured system at that frequency, or to a large input x(t), or to a strong contaminating signal n(t). In the time domain, the crosscorrelation function also suffers from this same inability to discriminate between causality and relationship in a measurement.

Transfer Functions

While the cross spectrum does not give a definite measurement it leads to two measurements which not only separate relationship and causality but also give quantitative results. The first of these functions measures the relationship between x(t) and z(t). It is a familiar function, the transfer function H(f) of the system (Fig. 3).

The transfer function is the ratio of the output linear spectrum for zero noise to the input linear spectrum.

$$H(f) = \frac{S_y(f)}{S_x(f)}. \tag{14}$$

Multiplying the numerator and denominator of this ratio by S_x^* shows that the transfer function can also be expressed as the ratio of the cross spectrum to the input auto spectrum.

$$H = \frac{S_y S_x^*}{S_x S_x^*} = \frac{G_{yx}}{G_{xx}} \tag{15}$$

There are two important points with regard to transfer functions measured in this way. The first is that this technique measures phase as well as magnitude since the cross spectrum contains phase information. Second, this measurement procedure is not limited to any particular input, such as sinusoids. In fact, the input signal may be random noise, or whatever signals are normally processed by the system being measured. For example, a telephone transmission system might be tested while in use with the normal traffic providing the test signal.

Coherence Functions

The major error in transfer function measurements develops when the output z(t) is not totally caused by the input x(t) but is contaminated by internal system noise n(t). Consider the input-output cross spectrum when there is uncorrelated noise with spectrum S_n added to the output.

$$G_{zx} = (S_y + S_n)S_x^* = G_{yx} + G_{nx} \tag{16}$$

If the noise n(t) is truly uncorrelated with x(t), and if enough averages of $\overline{G_{zx}}$ are taken, the contribution of $\overline{G_{nx}}$ to $\overline{G_{zx}}$ will approach zero, and $\overline{G_{zx}}$ will approach $\overline{G_{yx}}$. How rapidly the average of $\overline{G_{zx}}$ will approach $\overline{G_{yx}}$ depends upon how much noise there is in the output spectrum, that is, to what degree z(t) is caused by x(t).

To measure this coherence between x(t) and z(t) it is necessary to compute a new quantity, the *coherence function,* defined as[1]

$$\overline{\gamma^2} = \frac{\overline{G_{zx}}\,\overline{G_{zx}^*}}{\overline{G_{zz}}\,\overline{G_{xx}}} = \frac{|\overline{G_{zx}}|^2}{\overline{G_{zz}}\,\overline{G_{xx}}} \tag{17}$$

The horizontal bars denote ensemble averages.

After a number of records are averaged the numerator of the coherence function will reduce to $\overline{G_{yy}}\,\overline{G_{xx}}$. The denominator of the coherence function will be the auto spectrum of the normal output plus the noise, times the input auto spectrum. The output-plus-noise auto spectrum is

$$G_{zz} = (S_y + S_n)(S_y + S_n)^* = G_{yy} + G_{yn} + G_{ny} + G_{nn}. \tag{18}$$

After averaging, the cross terms in equation 18 disappear because they are uncorrelated with S_y, leaving

$$\overline{G_{zz}} = \overline{G_{yy}} + \overline{G_{nn}} \tag{19}$$

for the output auto spectrum. The coherence function then has an averaged value of

$$\overline{\gamma^2} = \frac{\overline{G_{xx}}\,\overline{G_{yy}}}{(\overline{G_{yy}} + \overline{G_{nn}})\overline{G_{xx}}} = \frac{\overline{G_{yy}}}{\overline{G_{yy}} + \overline{G_{nn}}} \tag{20}$$

Equation 20 shows that the coherence function $\overline{\gamma^2}$ has a value between 0 and 1, depending on the degree to which the output of the system is causally related to the input. This number not only defines the degree of causality, a useful quantity in itself, but it also defines the number of averages of the cross spectrum and input auto

spectrum that are required to define the transfer function to a given degree of accuracy.

An Example

Fig. 4 is an example of the separation of causality and relationship in a measurement. The system under test had a second-order highly damped transfer function. The input signal was Gaussian noise band-limited to the Nyquist folding frequency (10 kHz in this case).

$\overline{\gamma^2}$ for this measurement was about 0.8 out to the point where the transmission attenuation was about 20 dB. Beyond this frequency the data had too small a value to compute $\overline{\gamma^2}$ with any accuracy and it fell off to zero. The midband value of 0.8 for $\overline{\gamma^2}$ indicates that there was uncorrelated noise added to the system at some point other than the input. This could have been due either to real noise or to nonlinearities.

The transfer function, on the other hand, is a smooth well-defined function whose 3 dB and 90° phase points are at the same frequency. This indicates a good measurement of the relationship between input and output in spite of a fairly high uncorrelated noise environment.

Fig. 5 points up even more clearly the difference between a simple cross spectrum measurement and a transfer function measurement. Here the magnitude of the cross spectrum $\overline{G_{zx}}$ and the transfer function $\overline{H}$ are pictured on the same dB scale. Twenty-five sample records were averaged to determine system response, using a white noise input. One standard deviation on the input spectrum for this measurement is 20%, and since the cross spectrum does not employ information about the input its statistical certainty is poor. However, calculating the transfer function using the input power spectrum measured simultaneously with the cross spectrum reduces the statistical variation and gives a result with a few tenths of a dB of variation rather than 3 or 4 dB. In spite of the fact that a flat noise source is used, measurement of the transfer characteristics using a cross relationship alone is both inefficient and inaccurate.

Figs. 4 and 5 also illuminate a number of advantages of calculating the transfer function from the input spectrum and the cross spectrum. Clearly a good measurement can be made in spite of system noise. Also a measurement can be made using realistic test signals such as band-limited random noise. The phase measurement is unaffected by harmonic distortion and can be accurately made over wide dynamic ranges between input and output. The measurement can be made even more rapidly when there is no contaminating noise present. Thus, digital techniques of Fourier analysis offer powerful methods

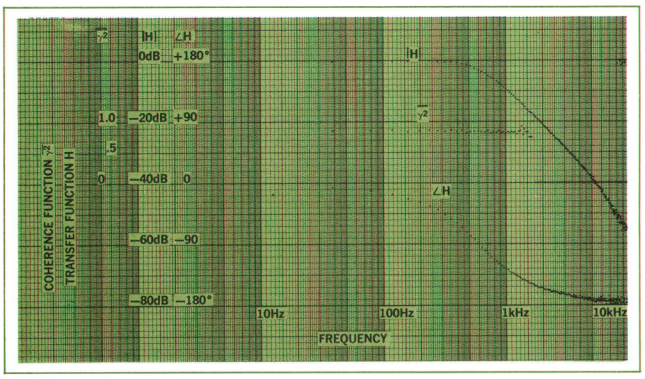

Fig. 4. *Transfer function of a second-order highly damped system measured by digital analyzer. Coherence function $\overline{\gamma^2} = 0.8$ indicates the presence of uncorrelated noise in the system (1.0 would indicate no noise), but transfer function is smooth and well defined, indicating a good measurement in spite of the noise.*

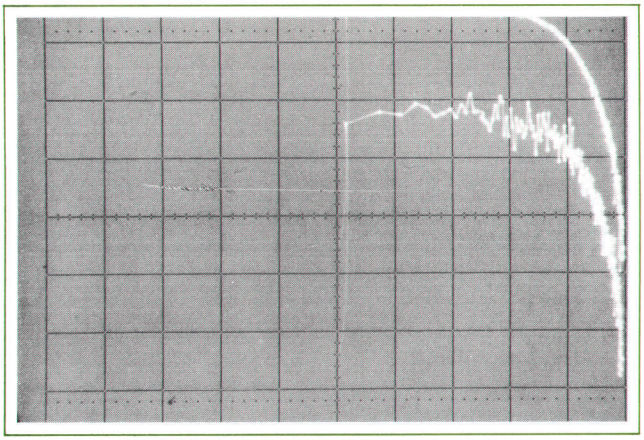

Fig. 5. *Transfer function of second-order highly damped system and cross power spectrum of input and output measured by digital analyzer. Spectra of twenty-five sample records were averaged. Not only do the magnitudes of $\overline{H}$ and $\overline{G_{zx}}$ differ, but also the statistical uncertainty in $\overline{G_{zx}}$ is much greater. This is because the computation for H takes into account the input power spectrum G_{xx}, whereas the computation for G_{zx} does not.*

for transfer-function determination that are unavailable with analog instruments.

Correlation Functions

So far I have described measurements that produce functions of frequency as their results. There are also functions of time which can be used in some of the same ways as spectra to clarify the nature of linear processes. These are correlation functions. The crosscorrelation function for two functions x(t) and y(t) is

$$R_{yx}(\tau) = \frac{1}{T}\int_0^T x(t)y(t-\tau)dt. \qquad (21)$$

The autocorrelation function R_{xx} is the same function with $y(t) = x(t)$. Naturally, when implemented on a digital processor the integral is replaced by a sum.

The computation proceeds as follows. First the average value of the sample-by-sample product of the two functions is computed over some interval T. Then the functions are displaced relative to each other and the process is repeated for the new value of the displacement τ. This is repeated for all values of τ and the results plotted as a function of τ.

The result of all this is a function R_{yx} which peaks when the functions y and x displaced by τ match each other well. The best use of the crosscorrelation function is to determine the delay between y(t) and x(t). The autocorrelation function, on the other hand, is used to determine periodicities in a single function, since it will peak every time the displacement is equal to the period.

It is interesting to consider several alternatives to a direct calculation of correlation functions. First of all, it can be shown that the auto spectrum and the autocorrelation function are Fourier transforms of each other. The same holds true for crosscorrelation and cross spectrum.[1]

$$G_{xx} = F[R_{xx}] \text{ and } R_{xx} = F^{-1}[G_{xx}] \quad (22)$$

$$G_{yx} = F[R_{yx}] \text{ and } R_{yx} = F^{-1}[G_{yx}] \quad (23)$$

Thus it is possible to calculate a correlation function by transforming a waveform to find the appropriate spectrum, complex conjugate multiplying the spectrum by itself or another spectrum, and then taking the inverse transform. While this may appear to be the long way around, it actually requires fewer multiplications to find a correlation function than calculating the average displaced products directly. Certain precautions must be observed because the discrete Fourier transform always assumes the sampled function is periodic with period T. However, it is possible to calculate an *exact* correlation function of $\pm N/2$ displacements (points) if 2N time points are available.

The lower trace in Fig. 6 shows the results of a crosscorrelation and impulse-response measurement on a damped second-order system with a white random noise input.[2] The measurement is the average of 25 sample records, but it still shows considerable statistical variation. It is difficult to determine if the ripple in the waveform is due to external noise, normal statistical variation, or the characteristics of the system being measured.

The upper waveform in Fig. 6 is the inverse Fourier transform of the transfer function computed from $\overline{G_{zx}}$ and $\overline{G_{xx}}$. This result shows much less statistical variation and is a more efficient way to compute the system impulse response, although it still does not give information about the effect of uncorrelated noise. For this we still need the coherence function.

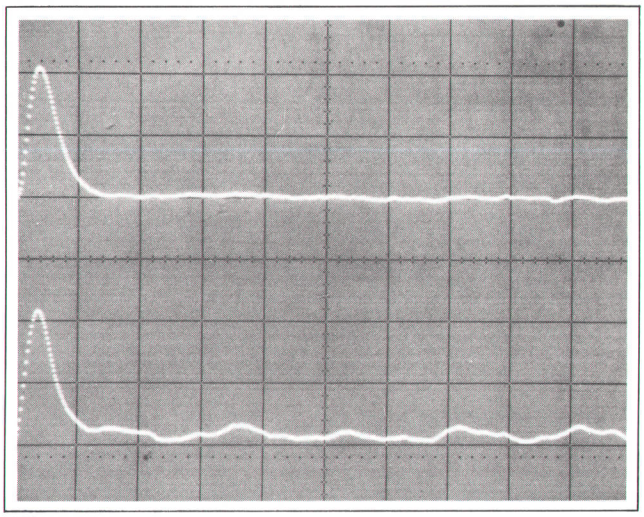

Fig. 6. *Crosscorrelation between white noise input and the output of a fourth-order linear system has the shape of the system impulse response. Lower trace is the crosscorrelation function computed directly. Upper trace was computed by inverse transforming the system transfer function, which was calculated by dividing the input-to-output cross power spectrum by the input power spectrum. Both traces are the average of 25 measurements. Upper trace is smoother as a result of taking into account the actual input power spectrum.*

References

[1]. J. S. Bendat and A. G. Piersol, 'Measurement and Analysis of Random Data,' John Wiley and Sons, 1966.

[2]. R. L. Rex and G. T. Roberts, 'Correlation, Signal Averaging, and Probability Analysis,' **Hewlett-Packard Journal,** November 1969.

Peter R. Roth

Peter Roth is project leader for the 5450A Fourier Analyzer. Before joining HP in January of 1965, Peter was an engineering officer in the U.S. Coast Guard for three years. Assigned to the Coast Guard's electronics laboratory in Alexandria, Virginia, he worked on high-frequency communications and on Loran A and Loran C equipment design. At HP, he has been responsible for development of the 5210A Frequency Meter, and has worked on the 5216A and 5221A, HP's first IC counters.

Peter received his BS and MS degrees in electrical engineering from Stanford University in 1959 and 1961. He is a member of IEEE and Tau Beta Pi.

A Calibrated Computer-Based Fourier Analyzer

This pushbutton-controlled digital measuring instrument performs complex analytical operations on input signals or time series. As a bonus, the user gets a general-purpose digital computer.

By Agoston Z. Kiss

ONE HEARTBEAT IN EVERY ¾ SECOND — 80 heartbeats per minute: these are the time-domain and frequency-domain descriptions of the same phenomenon. Neither contains more information than the other but to different people or to the same people in different circumstances one description may have more meaning or clarity than the other.

This duality between the time domain and the frequency domain is the basis of many important theorems and useful methods in signal and time-series analysis. Autocorrelation and crosscorrelation, power spectral density, cross power spectra, impulse response and transfer function, coherence, probability distribution and characteristic functions, convolution and filtering — these are examples of such methods. Since the principal theoretical bridge between the time domain and the frequency domain is the Fourier transform theorem, the methods of signal analysis that are based on the time-frequency duality are often called Fourier analysis.

Digital Fourier Analysis

Since 1965, the year of the Cooley-Tukey algorithm[1], Fourier analysis has been done more and more by digital techniques. The Cooley-Tukey algorithm, also called the fast Fourier transform, reduces the lengthy and cumbersome calculation of the Fourier coefficients by digital computer to a manageable, relatively rapid procedure. Computations that used to take hours can now be done in seconds. As a result, Fourier analysis is now becoming fashionable in many fields where it has not been used before because it took too long.

A version of the Cooley-Tukey algorithm is implemented in the new HP Model 5450A Fourier Analyzer, a calibrated, pushbutton-controlled instrument that can perform almost any Fourier-transform-based or related signal analysis (see Fig. 1). At the push of a button, the analyzer becomes a power spectrum analyzer, or a correlator, or an averager, or a digital filter, or any of a number of other instruments. *No knowledge of computer programming is required to operate it.* However, it can be converted into a general-purpose digital computer simply by moving a front-panel switch.

Model 5450A Fourier Analyzer combines a small general-purpose computer and some peripheral hardware into a flexible, user-oriented general-purpose instrument. An HP 2115A or 2116B computer with 8K memory is interfaced with a keyboard (Fig. 2), a dual-channel analog-to-digital converter (Fig. 3), a special display unit (Fig. 4), a teleprinter, and a punched-tape photoreader. An additional 8K memory can be installed to increase both the internal range and the number of peripherals. The analyzer has two basic modes of operation, i.e., as a Fourier analyzer or as a general-purpose computer. In the analyzer mode, it is either under keyboard control or under the remote control of another general-purpose computer.

The basic operations the Model 5450A can perform in the analyzer mode can be categorized as:

- data input/output
- transform related operations
- arithmetic operations
- data manipulations
- writing and editing of analysis routines.

Specific mathematical functions under keyboard control are:

- forward and inverse Fourier transform
- power spectrum
- cross power spectrum

- auto and crosscorrelation
- convolution
- histogram
- Hanning and other weighting functions
- real and complex multiplication and standard arithmetic operations
- integration and differentiation
- ensemble averaging

These can be executed separately or combined into complex routines.

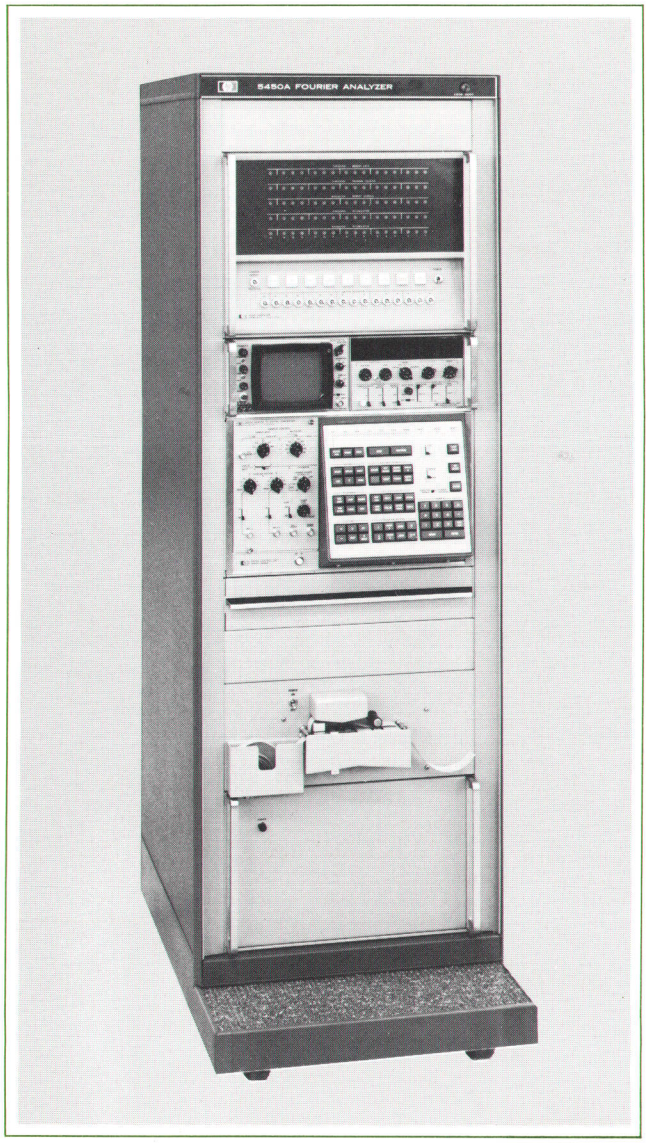

Fig. 1. *Model 5450A Fourier Analyzer is a flexible, push-button-controlled, modular, digital instrument, useful for analyzing waveforms and time series in a wide variety of systems and processes. It uses a standard HP computer for memory and computation, but requires no knowledge of computer programming. When it isn't doing Fourier analysis, the computer can be used separately.*

Fig. 2. *All Fourier analyzer operations are keyboard controlled. The principal operations—Fourier transforms, convolution, correlation, complex multiplication, coordinate transformations, and so on—can be called for by single keystrokes, or strung together using the programming and editing features to form routines to be run automatically later on. Typical routines can change the analyzer into a spectrum analyzer, a correlator, an averager, and many other instruments.*

Data Input/Output

There are 3K words available for data storage (8K words in the 16K version of the analyzer). This storage space can be filled up with data records; the shortest record is 64 words long and the longest is 1024 words long (4096 in the 16K version). Record lengths are push-button selectable in powers of two between these limits. The number of records which can be stored is the size of the data storage divided by the record length. Consequently, the 8K version can store 3 records of 1024 points each, or 6 records of 512 points each, and so on up to 48 records of 64 points each. These records are addressable as data block 0, 1, 2, . . . in every keyboard command.

Data Input

Analog data records can be read in via the analog-to-digital converter, which has two input channels with separate input attenuators. It can be switched to single-channel mode when only channel A is operational, or to dual-channel mode when channels A and B are both operational. Channels A and B are sampled simultaneously, then sequentially converted into digital values — channel

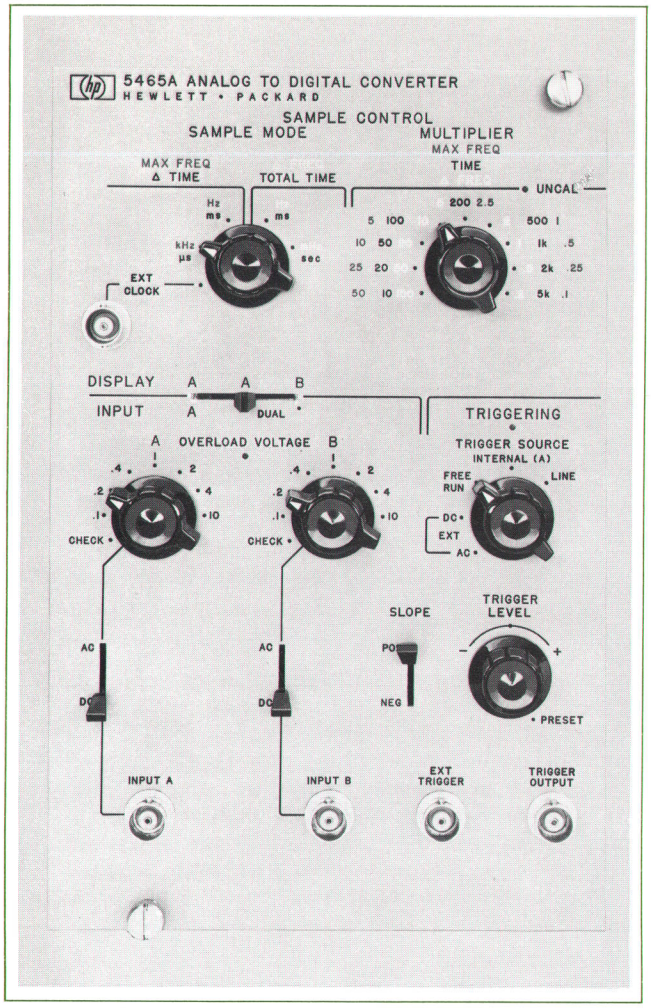

Fig. 3. *Analog-to-digital converter is the principal input device for analog signals. It can be operated as a single-channel unit or a dual-channel unit. The maximum sample rate for single-channel operation is 20μs per data point; the minimum rate is one sample in every five seconds. Data can also be read into the analyzer via peripheral devices, such as a tape reader or a teleprinter.*

A first — and stored in separate data blocks. The sample rate can be varied from 20 μs per data point (50 μs for dual-channel input) down to one sample in every five seconds.

There are some obvious but important relations between sampling time Δt, record length T, number of samples in a record (or data block size) N, frequency resolution Δf and upper frequency limit f_{max}:

$$T = N\Delta t \tag{1}$$

$$f_{max} = \frac{1}{2\Delta t} \tag{2}$$

$$\Delta f = \frac{1}{T} \tag{3}$$

Equation 1 says simply that a data record of length T seconds has been sampled N times with Δt seconds between samples.

Equation 2 is sometimes called the Shannon or Nyquist criterion of sampling, which states that to avoid loss of information, the highest frequency in a signal must be sampled at least twice per cycle.

Equation 3 really says that better frequency resolution requires a longer record. The analog equivalent of this statement is the observation that narrower-band filters take a longer time to reach steady state conditions.

A-D Converter

The analog-to-digital converter is a 10-bit ramp-type device with a 100 MHz clock. Because of its high differential linearity (3% as opposed to 25–50% for a typical successive-approximation-type A-D converter), the 60 dB dynamic range of the 10-bit converter will be appreciably improved, in some cases to as much as 90 dB, when any

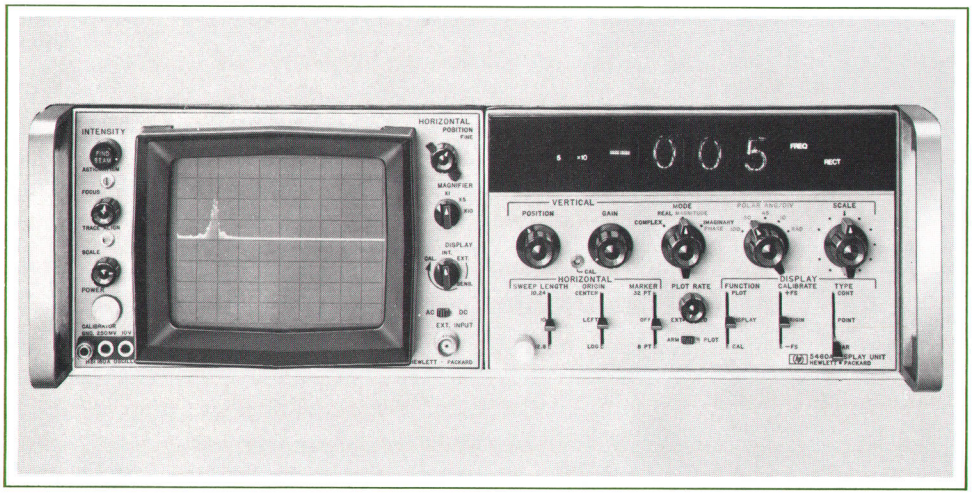

Fig. 4. *Built-in display unit is the principal analyzer output device when the recipient of the data is human. The digital display and annunciation indicate the vertical scale factor and the type of display. Data in the analyzer are always absolutely calibrated.*

averaging is done. The Fourier transform is a weighted average, of course. We have consistently observed dynamic ranges of 80 dB or more in computed transforms. How differential linearity and averaging affect dynamic range is quite a complex subject, and we hope to publish a paper on it soon.

Once the keyboard command is given for analog input, the actual record will be started by an internal or external sync signal with positive or negative slope, as selected by the user. After the last sample of the record has been stored, the analyzer calibrates the data, taking the input-attenuator setting into account, and establishes a scale factor for the record, which will follow it through all calculations. This absolutely calibrated input/output is one of the most important basic features of the analyzer.

Data can also be introduced into the analyzer through the numeric keys of the keyboard, through the teletype, through the photoreader (if they are on a punched paper tape), through the double binary I/O channels from another computer, and from digital magnetic tape (16K version only). The common feature of all the data input modes is that they can be initiated by keyboard control and that they establish calibrated data records in the analyzer.

Data Output

The most often used data output device is the display unit. Any stored data record can be displayed on the CRT by keyboard command. Also, when the analyzer is idle, it automatically reverts to a display mode, generally displaying the data record which was the subject of some I/O or analytical operation just before the idle period.

The display unit has many convenient features. It can display a time record or a frequency spectrum. When a spectrum is being displayed, its real part or imaginary part — or its amplitude or phase — can be displayed as a function of frequency, or the imaginary part can be displayed as function of the real part (Nyquist plot). Fig. 5 illustrates the possibilities. In every mode of data display, the calibration factor is also displayed as a power of 10, facilitating the readout of absolute values. Besides showing calibration, display lights also show whether the record displayed is in the time or frequency domain, whether the amplitudes are linear or logarithmic, and whether they are calculated in rectangular (real and imaginary) or polar (amplitude and phase) coordinates.

Other features of the display unit are: digital up or down scaling in ten steps, linear or logarithmic horizontal scale, markers on every 8 or 32 points, point display, continuous curve display or bars drawn from display points to the zero level horizontal axis. It also has a calibration mode, and a plotter mode in which it can drive an X-Y recorder to plot exactly what is being displayed on the CRT.

Data records can also be printed out on the teletype, punched out on paper tape either on the punch unit of the teletype or on an optional fast punch, transferred on the double binary I/O channels to another computer, plotted on a digital plotter, or stored on digital magnetic tape (the last two features on the 16K version only). Common features of all data output modes are that they can be initiated by keyboard command and that the data are always calibrated.

A final remark about the calibrated input-output feature. The analyzer, being a binary device, carries the calibration in radix two. In every output operation where the recipient is non-human (binary I/O, paper tape, digital magnetic tape), the calibration remains in radix two to retain maximum accuracy. However, in every human-related output operation (display, data printout, plotting) the calibration is changed to radix 10 for maximum user convenience.

Transform-Related Operations

The most important transform-related operations are, of course, the forward and inverse Fourier transforms. The definitions of these operations are:

$$S_x(m\Delta f) = \frac{1}{N}\sum_{n=0}^{N-1} x(n\Delta t)\, exp\left\{-i\frac{2\pi}{N}\, nm\right\} \qquad (4)$$

and

$$x(n\Delta t) = \sum_{m=0}^{N-1} S_x(m\Delta f)\, exp\left\{i\,\frac{2\pi}{N}\, nm\right\} \qquad (5)$$

where N is the number of samples (points) in the time record x(t) or frequency record $S_x(f)$.

Although the time function x(t) is always real, the spectrum, $S_x(f)$, is generally complex. In a complex spectrum, every spectral value (except dc) has to be described by two quantities, either amplitude and phase, or real (cosine or in-phase) and imaginary (sine or quadrature) components. The former is the polar-coordinate representation and the latter is the rectangular-coordinate representation. In the analyzer, all calculations are carried out in rectangular coordinates, but the results can be converted into polar coordinates by keyboard command.

Since the Fourier transform does not create new information, the Fourier spectrum of a time record with N independent data points will also contain exactly N inde-

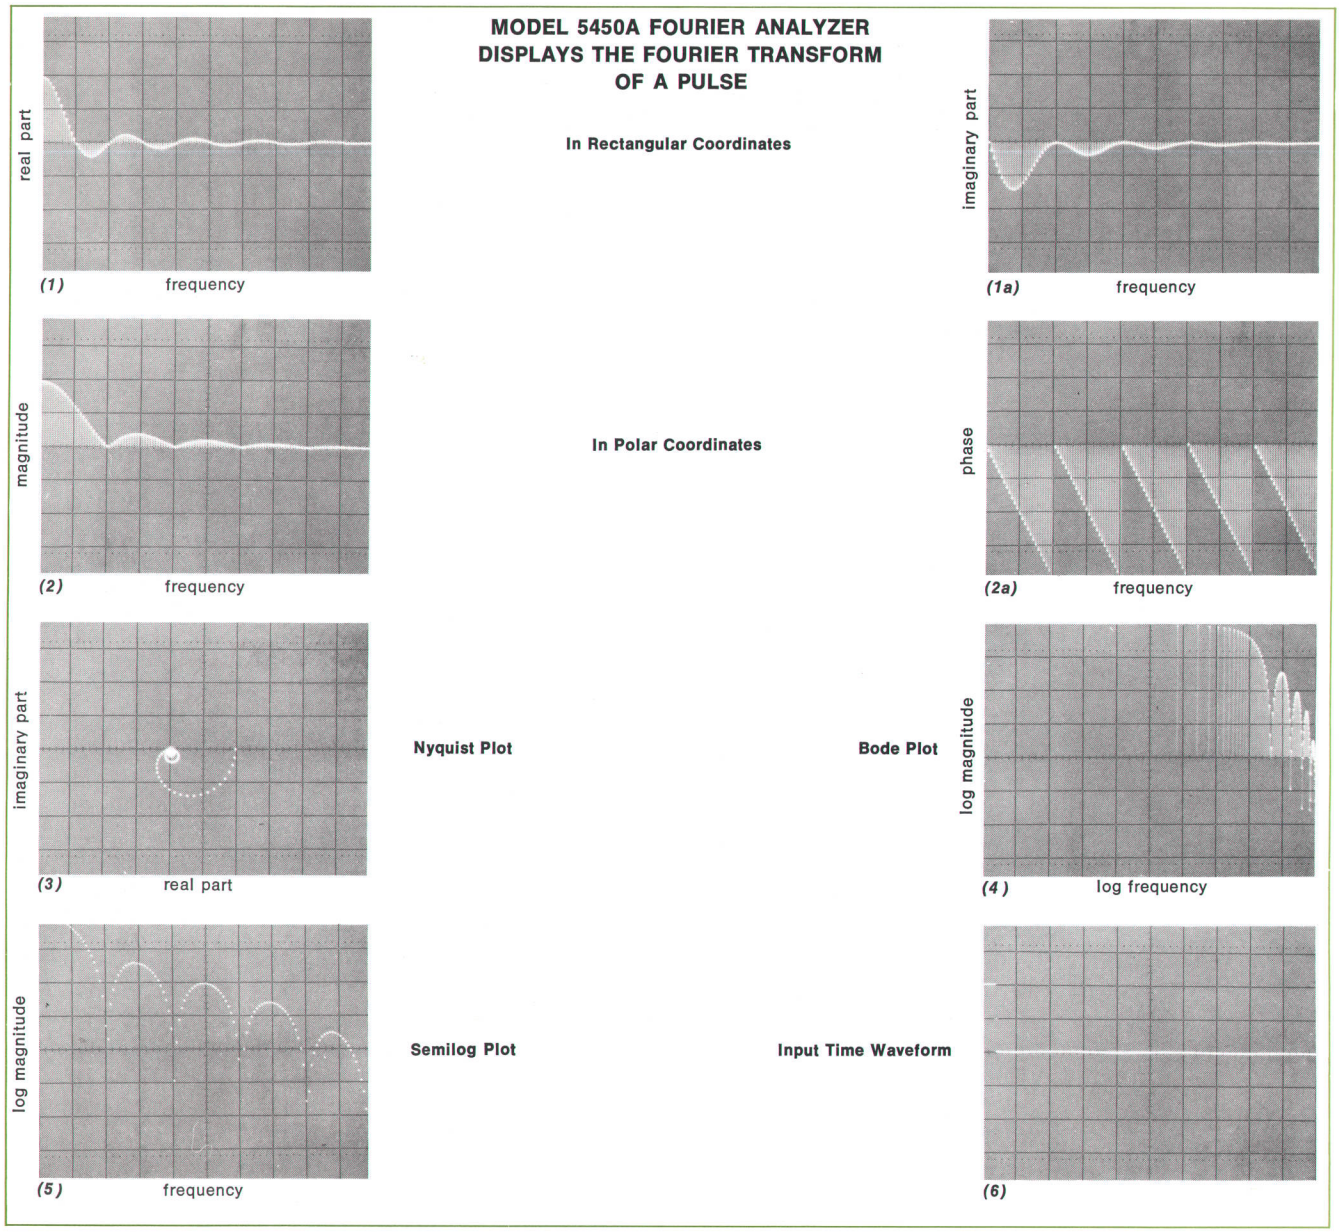

Fig. 5. *The analyzer has a display mode to suit every need, and it changes from one to another at the touch of a button or the flick of a switch. In every case the readouts on the display unit and the A-D converter indicate scale factors and type of display.*

pendent data points. But since every frequency point has to be described by two independent data values — except dc, which has no phase, and the highest frequency, which by definition has zero phase — the Fourier spectrum of a time record with N points will contain N/2 frequency-value pairs (for counting purposes dc and the highest frequency are counted as one frequency-value pair).

Equation 4 actually defines a spectrum for negative as well as positive frequencies. However, the analyzer is restricted to the analysis of physically realizable, and therefore real, time functions only. The spectra of real time functions are Hermitian (i.e., even real part and odd

A Fourier Analyzer Makes Fundamental Measurements

The measurements a Fourier analyzer makes are useful to behavioral scientists, psychophysicists, biomedical researchers, process control system designers, analytical chemists, and oceanographers, and to people working in vibration analysis, structural mechanics, acoustics, geophysics, control system design and analysis, component testing, system identification, sonar, and many other fields. The reason a Fourier analyzer is so widely useful is that, like a voltmeter, it makes fundamental measurements. For the same reason, no finite list of applications can convey a true picture of its capabilities. Here are just a few examples of applications.

Analytical chemists can use it to measure nuclear magnetic resonance (NMR) spectra, and as an averager to improve the sensitivity of their spectrum measurements.

Structural designers, e.g. of airframes, can use it to determine the transfer function and vibration modes of a structure, the spectra of vibrations induced at various points by various inputs, and the degree of coherence between vibrations at different points.

Behavioral scientists can use it to determine the transfer function of a driver, and the degree of coherence between his responses and various input stimuli.

Brain researchers can use it to measure the spectra of brain waves, and the degree of coherence between waves at different points in the brain.

Designers of process control systems and other systems—power plants, servomechanisms, etc.—can use it to determine transfer functions, impulse responses, coherence between signals, power spectra, and cross power spectra.

In application after application, the measurements are the same—transfer function, coherence function, power spectrum, cross power spectrum, and combinations of these fundamental measurements. End uses of the data differ, of course. To the designer of a structure or a control system, it's accurate information that he couldn't have obtained without the analyzer, and he uses it to optimize his design, avoid overdesign, and optimize performance adjustments. The physician analyzing an electromyogram (EMG) is looking for evidence of muscle disease. What these and other users and potential users of Fourier analyzers have in common is that they are working with time series—voltages, vibrations, sound waveforms, or perhaps just a series of data points obtained at regular intervals and punched on paper tape. On such inputs the Fourier analyzer makes measurements and computes functions that would be difficult to do by any other means. It does these things with the convenience of keyboard control, rapidly, and with great flexibility.

imaginary part), so the negative-frequency parts of the spectra of real time functions contain no additional information. Partly to increase the effective transform speed of the analyzer and partly to avoid the confusion that the mentioning of the existence of negative frequencies generally creates, we used a version of the fast Fourier algorithm that applies only to time signals that are real.

Like other versions of the fast Fourier algorithm, ours is an 'in-place' algorithm. Intermediate and final results of computations are stored in the same data block as the original data.

Correlation and Convolution

Auto and crosscorrelation are well known and widely used methods in signal analysis. They are used to improve signal-to-noise ratio, to find hidden periodicities, and so on. Convolution, on the other hand, is in most cases a mean trick nature plays on us. When we use any measuring equipment to measure an event, the result is never the phenomenon we want to observe but its convolution with the impulse response of the equipment used. Sometimes, however, even convolution can be useful. For example, smoothing a record by taking a K-point running average can be performed in the analyzer by convolving the record in question with another record containing K unit impulses.

Correlation and convolution are generally defined on the time domain, although they do depend on the frequency content of the functions in question. Since both correlation and convolution involve an enormous number of multiplications and additions — N^2 of them to be exact — to perform either of them within a reasonable time requires special hardware. But according to the convolution theorem, convolution (correlation) in one domain is multiplication (conjugate complex multiplication) in the other domain. Therefore convolution (correlation) can be reduced to two Fourier transforms and one point-by-point multiplication (conjugate complex multiplication) of the two records involved. If x(t) and y(t) are two time functions and their respective spectra are $S_x(f)$ and $S_y(f)$, the analyzer performs the following calculations: for crosscorrelation,

$$x(t)\bigstar y(t) = F^{-1}\,[S_x(f) \cdot S_y{}^*(f)] \qquad (6)$$

and for convolution,

$$x(t)*y(t) = F^{-1}\,[S_x(f) \cdot S_y(f)] \qquad (7)$$

Here the superscript * stands for complex conjugate, the $*$ between two time functions for convolution, the $\bigstar$ for crosscorrelation, and F^{-1} for inverse Fourier transform. These operations can be performed step by step on the

analyzer, but for convenience the keyboard has a correlation key and a convolution key.

Hanning

Physically realizable devices can act only on signals which are limited in duration and in bandwidth. If infinitely long signals or signals with infinite bandwidth are passed through any physical device, they will be time and frequency-band limited by the device itself.

The simplest kind of time-limiting is the application of a square time window. If we have a function x(t) and we take a T-second long record of it, say from $t = 0$ to $t = T$, then we have really multiplied x(t) by a square pulse T seconds long with unity amplitude (see Fig. 6).

What happens to the spectrum of this function? The convolution theorem says that multiplication in one domain is convolution in the other domain. Since we multiplied x(t) by the window function $\sqcap$ (t/T), the spectrum of the time-limited function x(t) • $\sqcap$ (t/T) will be the convolution of the spectrum of the original function and the spectrum of the time window. Let us say that the spectrum of x(t) is $S_x(f)$. The spectrum of $\sqcap$ (t/T) is (Fig. 7) $\text{sinc}\pi Tf = \frac{\sin\pi Tf}{\pi Tf}$, and the spectrum of x(t) • $\sqcap$ (t/T) is $S_x(f) * \text{sinc}\pi Tf$. The maximum value of the sinc function is unity at $f = 0$, it has zero crossings at $f = 1/T$, $2/T, \ldots$, and the amplitude of the sidelobes decreases at 6 dB per octave. If $S_x(f)$ has spectral lines exactly at $f = 0, 1/T, 2/T, \ldots$, that is, if x(t) was periodic in the time window $\sqcap$ (t/T), then convolving $S_x(f)$ with $\text{sinc}\pi Tf$ will simply result in $S_x(f)$. But if x(t) was not periodic in the time window, then the spectral lines of $S_x(f)$ and the zero crossings of the window spectrum will not coincide and the convolution process will smear each spectral line of $S_x(f)$ all over the spectrum. Even if $S_x(f)$ contains one spectral line only, the result will be a series of spectral lines spaced 1/T apart and having an amplitude decay of 6 dB per octave. This phenomenon is often referred to as the leakage effect.

Leakage can be avoided only by making sure that the function x(t) is periodic in the time window. Obviously, this condition can seldom be met. Therefore, in order to reduce the effect of leakage, different window shaping ideas have been proposed. The idea of the window shaping is to make x(t) somehow 'quasi-periodic' in the time window with the least possible loss of information. Among these window-shaping methods the Hanning window has proved most popular. It is a $\frac{1}{2}\left(1 \pm \cos\frac{2\pi t}{T}\right)$ window, where both the window and its derivative approach zero at the two ends of the record. Its effect on the spectrum is that the main lobe of each line is widened by an additional 1/T, but the sidelobes decay by an additional 12 dB per octave.

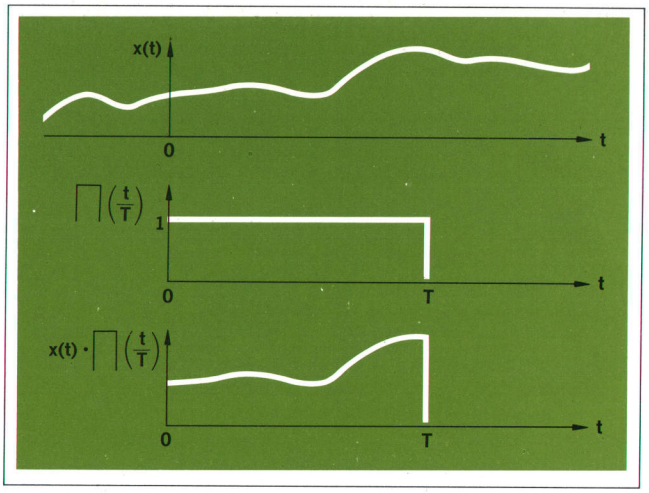

Fig. 6. *When a T-second record is taken of an analog input, the effect is to multiply the input by a square window function. If the input isn't periodic with period T, the spectral lines of the input will not be lines but will have the sin x/x shape shown in Fig. 7. To reduce this effect, Model 5450A Fourier Analyzer has built-in Hanning window-shaping functions.*

Two other window-shaping methods are the Chebyshev window and the Parzen window. The Chebyshev window achieves a faster sidelobe decay than the Hanning window but is much more cumbersome to implement. The triangular Parzen window is fairly easy to implement but not as effective as the Hanning window.

In Model 5450A two different Hanning windows can be applied by pushbutton command. The interval-centered Hanning window, HI, is used to reduce leakage as described above. The origin-centered Hanning window, HO, can be used to form a 3-point running average of records with ¼, ½, ¼ weighting.

Integration and Differentiation

There is a keyboard command to integrate any data record between any two chosen data points or to differentiate any chosen data record. The defining equations for integral and differential are:

$$\int D_k = D_{k-1} + D_k \quad , k = 0, 1, \ldots, N-1 \qquad (8)$$

$$\frac{d}{dx} D_k = D_k - D_{k-1} \quad , k = 0, 1, \ldots N-1 \qquad (9)$$

By definition, $D_{0-1} = 0$.

The integral routine is especially useful for calculating integral power spectra, cumulative probability distribu-

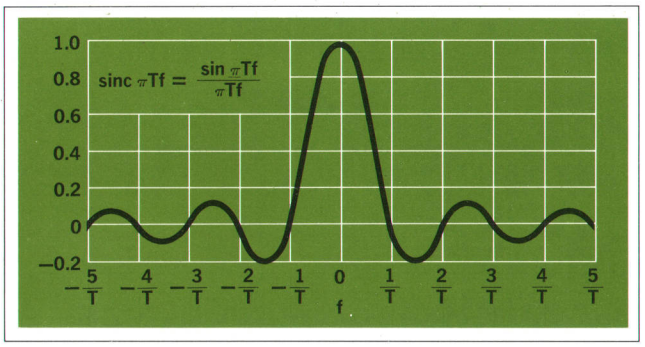

Fig. 7. *Spectral lines of a sampled function will have this sin x/x shape if the function isn't periodic in the record length T. Hanning weighting doubles the width of the main peak, but causes later peaks to fall off at 18dB per octave instead of 6dB per octave.*

tion functions or third-octave, half-octave, or full-octave filters. The differential routine can be used to calculate higher moments of probability density functions by differentiating their Fourier transforms (i.e. their characteristic functions).

Arithmetic Operations

There are keyboard commands for the addition, subtraction, multiplication, and division of data records. These operations are performed on a point by point basis. Addition is especially useful for ensemble averaging of data records either in the time domain or in the frequency domain, thereby improving the statistics of the measurement. There are separate commands for complex multiplication and conjugate complex multiplication of two selected data records. Both multiplications result in real multiplication if the records are in the time domain.

The division of a data record by another selected data record is performed as real or complex division in the time or frequency domain, respectively. In addition to these data block operations, any selected data block can be multiplied or divided by any positive or negative constant whose magnitude is less than 32767.

Data Manipulations

Since there can always be more than one data record stored in the analyzer, 'Store,' 'Load' and Interchange' keyboard commands were established to effect data transfers among them.

As I have mentioned, all transform-related and arithmetic operations are performed in rectangular coordinates. However, spectral results are often desired in polar coordinates (amplitude and phase). There are keyboard commands to change the coordinate system of any chosen data record from rectangular to polar or from polar to

PROGRAMMING THE FOURIER ANALYZER

A Power Spectrum Averaging Program

	Key	Meaning
	LABEL	Identifies starting point of sequence.
*	ANALOG IN	Take in sample of analog data.
*	F	Take Fourier transform of sample.
*	MULT*	Complex conjugate multiply Fourier transform, yielding power spectrum.
	+	Add power spectrum to sum of previous power spectra.
	STORE	Store new power spectrum sum.
	COUNT	Repeat above process (from the label point) the number of times desired.
	÷	Divide final power spectrum sum by number of spectra taken, yielding average.
	END	End of sequence.

Fig. 8. *An often-used routine is the power spectrum averaging program. After the steps are entered into the analyzer's memory, the program can be listed on the teleprinter. Errors can be corrected by adding, deleting, or modifying steps. Another keystroke makes the program execute. Model 5450A will compute one 1024-point spectral estimate (the three steps marked*) in 2.4 seconds or less.*

rectangular. Further keyboard commands can change linear amplitudes to logarithmic or logarithmic amplitudes to linear. The execution of these commands ('Rectangular,' 'Polar,' 'Logarithmic Amplitude,' 'Exponential Amplitude') are based on a power-series technique in which the coefficients are calculated by Chebyshev expansion of the function desired.

Writing and Editing Routines

The power of Model 5450A Fourier Analyzer is not only in the easy access it offers to the most important basic signal analytical operations, but also, and perhaps even more so, in its capability of building automatic routines using these operations. Programming the analyzer to carry out a sequence of computations actually transforms it into a different measuring instrument — a spectrum analyzer, for example, or a signal averager, or a correlator.

Keyboard commands can be assembled into routines up to 100 steps long (200 steps in the 16K version). The routines can incorporate labels, jump instructions, subroutines, and loops, thereby providing an extremely flexible and easily learned high-level instruction set for almost any type of signal analysis. The assembled rou-

tines reside in the analyzer. They can be listed on the teletype, punched out on paper tape, re-edited using 'Delete,' 'Replace,' and 'Insert' edit commands, and can be run under keyboard control.

Here are some of the most often used routines.

Power Spectral Analysis. Ensemble averaging to improve the signal-to-noise ratios of power spectral estimates can be simply performed by:

1. reading in a time record
2. taking its Fourier transform
3. conjugate complex multiplying the spectrum by itself, thereby creating a power spectral estimate
4. summing the power spectral estimate into a second record
5. repeating operations 1–4 any desired number of times
6. after summing a given number of power spectral estimates, dividing the result by the number of estimates.

Fig. 8 illustrates the program, which computes

$$G_{xx}(f) = \overline{S_x(f) \cdot S_x{}^*(f)} \qquad (10)$$

Cross Power Spectra. The cross power spectrum contains the frequencies common to the individual spectra of two signals. It is the Fourier-transform of the crosscorrelation function. To create the ensemble average of cross power spectral estimates, one can follow the instructions for power spectral averaging, except in step 1 take two simultaneous records, and in step 3 conjugate complex multiply one spectrum by the other. The function computed is

$$G_{xy}(f) = \overline{S_x(f) \cdot S_y{}^*(f)} \qquad (11)$$

In Equations 10 and 11 $G_{xx}(f)$ stands for power spectrum, $G_{xy}(f)$ for cross power spectrum, $S_x(f)$ and $S_y(f)$ are the Fourier spectra of functions x(t) and y(t) respectively, the superscript * stands for complex conjugate, and the upper bar for ensemble averaging.

Digital Filtering. Let us consider a filter as a black box with one input and one output:

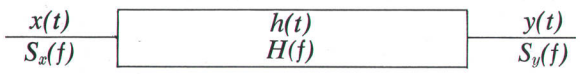

The black box can be characterized by its impulse response, h(t) or its transfer function, H(f). They are Fourier-transform pairs. The input function is x(t), and the output is y(t). $S_x(f)$ and $S_y(f)$ are their respective Fourier spectra.

The filter equation simply states that the output spectrum is the product of the input spectrum and the transfer function of the filter:

$$S_y(f) = S_x(f) \cdot H(f) \qquad (12)$$

Filtering can be easily performed in the Model 5450A by storing the filter transfer function in one of the data records and block-multiplying the spectrum of the input signal by it. Taking the inverse transform of the product results in the output function, y(t).

Inverse Filtering or Deconvolution. Equation 12 can be rewritten in the time domain using the convolution theorem (multiplication in one domain equals convolution in the other domain):

$$y(t) = x(t) * h(t) \qquad (13)$$

that is, the output of the black box is the convolution of its impulse response and the input function. Now if this black box happens to be some measuring equipment, it is x(t) that we are interested in, not y(t). The inverse operation of convolution is pretty difficult to produce, but Equation 12 can be rewritten as:

$$S_x(f) = \frac{S_y(f)}{H(f)} \qquad (14)$$

Since $S_y(f)$ and H(f) are known, the division can be performed point by point. Taking the inverse Fourier transform of the quotient results in x(t).

Transfer Function and Coherence. A method based on Equation 12 can be worked out to measure the transfer functions of unknown black boxes and to find causal relationships between inputs and outputs. This extremely important and interesting subject is discussed by Peter Roth elsewhere in this issue.

Measurement of Statistical Behavior. Random data can be characterized by their statistics: probability density functions, distribution functions, and the moments of the probability distribution. The analyzer can collect amplitude histograms or, with an optional input box, time-interval histograms. The histograms are really frequency curves; the independent variable is amplitude or (time interval) and the dependent variable is the frequency of occurrence. Histograms can be easily normalized to give probability density functions and integrated to calculate distribution functions.

The Fourier transforms of distributions are called characteristic functions; they are used mainly in theoretical work in statistics. The differentials of the characteristic functions can be used to calculate the moments and central moments of the distributions.[3]

SPECIFICATIONS

HP Model 5450A
Fourier Analyzer

ANALOG INPUT

The Analog-to-Digital Converter accepts one or two inputs. In two-channel operation both inputs are sampled simultaneously. Resolution of the ADC is 10 bits.

INPUT IMPEDANCE: 1 MΩ ±1% shunted by 45 pF max.
SENSITIVITY: 30 μV rms (sine wave).
CONVERSION GAIN (CHANNEL A):
ACCURACY (as a function of frequency):
±.2% ±1 × 10^{-4}%/Hz.
TEMPERATURE STABILITY: 0.005%/°C.
LINEARITY: Integral, ±0.05%; Differential, ±3%.
GAIN AND PHASE CHANNEL A TO B:
CONVERSION GAIN A/B: ±0.2% ±4 × 10^{-4}%/Hz.
TEMPERATURE STABILITY: 0.01%/°C.
PHASE AND DELAY A TO B: ±0.2°, ±5 μs.
SAMPLE RATE CONTROL:
MAXIMUM FREQUENCY/TIME BETWEEN SAMPLES MODE: Maximum frequency is selectable from 0.1 Hz to 25 kHz (0.1 Hz to 10 kHz in two-channel operation).
FREQUENCY RESOLUTION/TOTAL TIME MODE: Frequency resolution is selectable from 0.2 mHz to 100 Hz.

DISPLAY UNIT

Data may be displayed on the 8 x 10 cm oscilloscope or output to a plotter or remote oscilloscope in the following forms:

Y AXIS	X AXIS
Real Part Amplitude	Time
Real Part Amplitude	Frequency (Linear or Log)
Imaginary Part Amplitude	Frequency (Linear or Log)
Magnitude (Linear or Log)	Frequency (Linear or Log)
Phase	Frequency (Linear or Log)
Imaginary Part Amplitude	Real Part Amplitude

ANALOG DISPLAY ACCURACY: ±1%.
TYPES OF DISPLAY: Points, bars, or continuous (interpolation).
AMPLITUDE SCALE: Data in memory is automatically scaled to give a maximum on-screen calibrated display. The scale factor is given in volts/division, volts²/division, or in dB offset.
LINEAR DISPLAY RANGE: ±4 divisions with scale factor ranging from 1 × 10^{-150} to 5 × 10^{+150} in steps of 1, 2, 5 and 10.
LOG DISPLAY RANGE: 4 decades with a scale factor ranging from 0 to −998 dB.
TIME AND FREQUENCY SCALE:
LINEAR SWEEP LENGTH: 10, 10.24, or 12.8 divisions.
LOG HORIZONTAL: 0.5 decade/division.
ANALOG PLOTTER OUTPUT:
AMPLITUDE: 0.5 V per oscilloscope display division.
LINEARITY: 0.1% of full scale.

BLOCK SIZES FOR TYPICAL MEASUREMENTS

The following table indicates some of the measurements made by the 5450A as well as the *maximum* block size available for these measurements.

MEASUREMENT	BLOCK SIZE N (Points/Ensemble) 8K MEMORY	16K MEMORY
Power Spectral Density — Ensemble Average	1024	4096
Voltage Spectrum — Ensemble Average	1024	4096
Cross Power Spectral Density — Ensemble Average	1024	2048
Transfer Function	512	2048
Coherence Function	512	1024
Autocorrelation of N/2 Lags	1024	4096
Crosscorrelation of N/2 Lags	1024	4096
Crosscorrelation of N/2 Lags — Ensemble Average	1024	2048
Autocorrelation of N/2 Lags — Ensemble Average	1024	2048
Power Spectral Density of One Shot Transient	1024	4096
Voltage Spectrum of One Shot Transient	1024	4096

COMPUTATIONAL SPEED

The speeds shown are based on using the 2116B Digital Processor.
FOURIER TRANSFORM:
Block Size 64: 52 ms
Block Size 1024: 1.4 s
POWER SPECTRUM ENSEMBLE AVERAGE:
Block Size 64: 110 ms/Spectral Estimate
Block Size 1024: 2.1 s/Spectral Estimate
CROSS POWER SPECTRUM ENSEMBLE AVERAGE:
Block Size 64: 210 ms/Spectral Estimate.
Block Size 1024: 4.4 s/Spectral Estimate.

DIGITAL ACCURACY AND RESOLUTION

All calculations using floating point arithmetic on a block basis. Data overflow does not occur. Amplitude resolution is 1 part in 16,000 worst case.
DATA MEMORY SIZE: 3072 words (8192 for a 16,384 word memory).
DATA BLOCK SIZE: Any power of 2 from 64 to 1024 (to 4096 with a 16,384 word memory).
DATA WORD SIZE: 16 bit real and 16 bit imaginary or 16 bit magnitude and 16 bit phase.
COMPUTATIONAL RANGE: ±150 decades.
TRANSFORM ACCURACY: 0.1% worst case error during the forward or inverse calculation.

SPECTRAL RESOLUTION

The element of spectral resolution is the frequency channel width, the maximum frequency divided by ½ the data block size.
MAXIMUM FREQUENCY: 25 kHz single channel; 10 kHz dual channel.
FREQUENCY CHANNEL WIDTH: <3.2% down to <0.2% of the maximum frequency (down to <0.05% for 16,384 word processor).
SPECTRAL RESOLUTION OF TWO EQUAL AMPLITUDE SINE WAVES: If separated by 3 frequency channel widths, there will be a null of at least 3 dB between them; if separated by 7 frequency channel widths the relative magnitudes will be correct to within 0.1%. The power spectrum for two equal amplitude sine waves separated by 5 frequency channels will have the correct relative magnitude to within 0.1%.
DYNAMIC RANGE: 4 decades over ±150 decades.
ENVIRONMENTAL CONDITIONS: 0°C to 55°C using 2116B Digital Processor (10°C to 40°C using 2115A Digital Processor).

PRICE: Systems start at approximately $50,000, depending upon choice of computer and other required options.

MANUFACTURING DIVISION: HP Santa Clara Division
5301 Stevens Creek Boulevard
Santa Clara, California 95050

Possibilities Unlimited

The analytical operations available in Model 5450A Fourier Analyzer can be combined in many, many ways, and only the best known were mentioned here. But the analyzer will cater to the most esoteric tastes, including, for example, *cepstrum, saphe cracking,* and *liftering.*[4]

Acknowledgments

The original idea of Model 5450A and much of the initial groundwork is due to Ron Potter. The engineering development was under the capable leadership of Peter Roth as project leader, who also developed the analog-to-digital converter. Hans-Jurg Nadig and Evan Deardorff developed the keyboard control and display units, respectively. Steve Cline wrote most of the software. Bill Katz and Dick Cavallaro gave valuable assistance in the development work and documentation. The industrial designer was Roger Lee, and the mechanical designers were Wally Mundt and Chuck Lowe. Special thanks are due to Fred London and his group for their marketing support, to Pete Schorer for the excellent manuals, and to Jim Shea and his service group — especially Al Linder — for their service support. Production is in the capable hands of Don Lien's group and Walt Noble, production engineer.

The fast-Fourier algorithm used in Model 5450A is based mainly on those of Cooley and Tukey[1] and Gentleman and Sande[5]. However, we also benefited from the ideas of I. F. Good, R. C. Simpleton, R. B. Blackman, G. C. Danielson, C. Lanczos, R. Shively, and others.

References

[1]. V. W. Cooley and J. W. Tukey, 'An algorithm for the machine calculation of complex Fourier series,' Math. Comp. Vol. 19, pp. 297-301, 1965.

[2]. R. L. Rex and G. T. Roberts, 'Correlation, Signal Averaging, and Probability Analysis,' **Hewlett-Packard Journal,** November 1969.

[3]. R. M. Bracewell, 'The Fourier Transform and Its Applications,' McGraw-Hill Book Company, 1965, Chapter 8.

[4]. B. P. Bokert, M. J. R. Healy, and J. W. Tukey, 'The Quefrency Alanysis of Time Series for Echoes: Cepstrum, Pseudo-Autocovariance, Cross-Cepstrum, and Saphe Cracking,' Proceedings of the Symposium on Time Series Analysis: Brown University, June 11–14, 1962 (M. Rosenblatt, ed.), John Wiley and Sons, Inc., 1963.

[5]. W. M. Gentleman and G. Sande, 'Fast Fourier Transform for Fun and Profit,' 1966 Fall Joint Computer Conference, AFIPS, Proceedings, Vol. 29, pp. 563-578.

Agoston Z. Kiss

Ago Kiss began his varied professional career in his native Hungary doing research in nuclear particles and electromagnetic interactions. After five years he moved to England, where he spent another five years, first as a research physicist and then as a consultant in control systems. In 1962 he came to the United States and spent three years developing bandwidth reduction and coding techniques before joining the staff of Hewlett-Packard Laboratories in 1965. At HP, Ago has worked on character recognition and gas chromatograph control, and since 1967 has been engineering group manager in charge of the development of Fourier Analyzers. Ago wes responsible for defining the 5450A Fourier Analyzer and developing its algorithms.

Ago studied electrical engineering at the Polytechnical University of Budapest, and physics and math at the University of Sciences in Budapest. He has also done postgraduate work at the University of Pittsburgh. He holds several British and U.S. patents related to control circuits. For relaxation, Ago chooses the out-of-doors. Skiing and sailing are his favorite leisure-time activities.

TECHNICAL INFORMATION FROM THE LABORATORIES OF THE HEWLETT-PACKARD COMPANY PUBLISHED AT 1501 PAGE MILL ROAD, PALO ALTO, CALIFORNIA 94304
Editor: R. H. Snyder Editorial Board: R. P. Dolan, L. D. Shergalis Art Director: Arvid A. Danielson Assistant: Maridel Jordan