HEWLETT-PACKARD JOURNAL

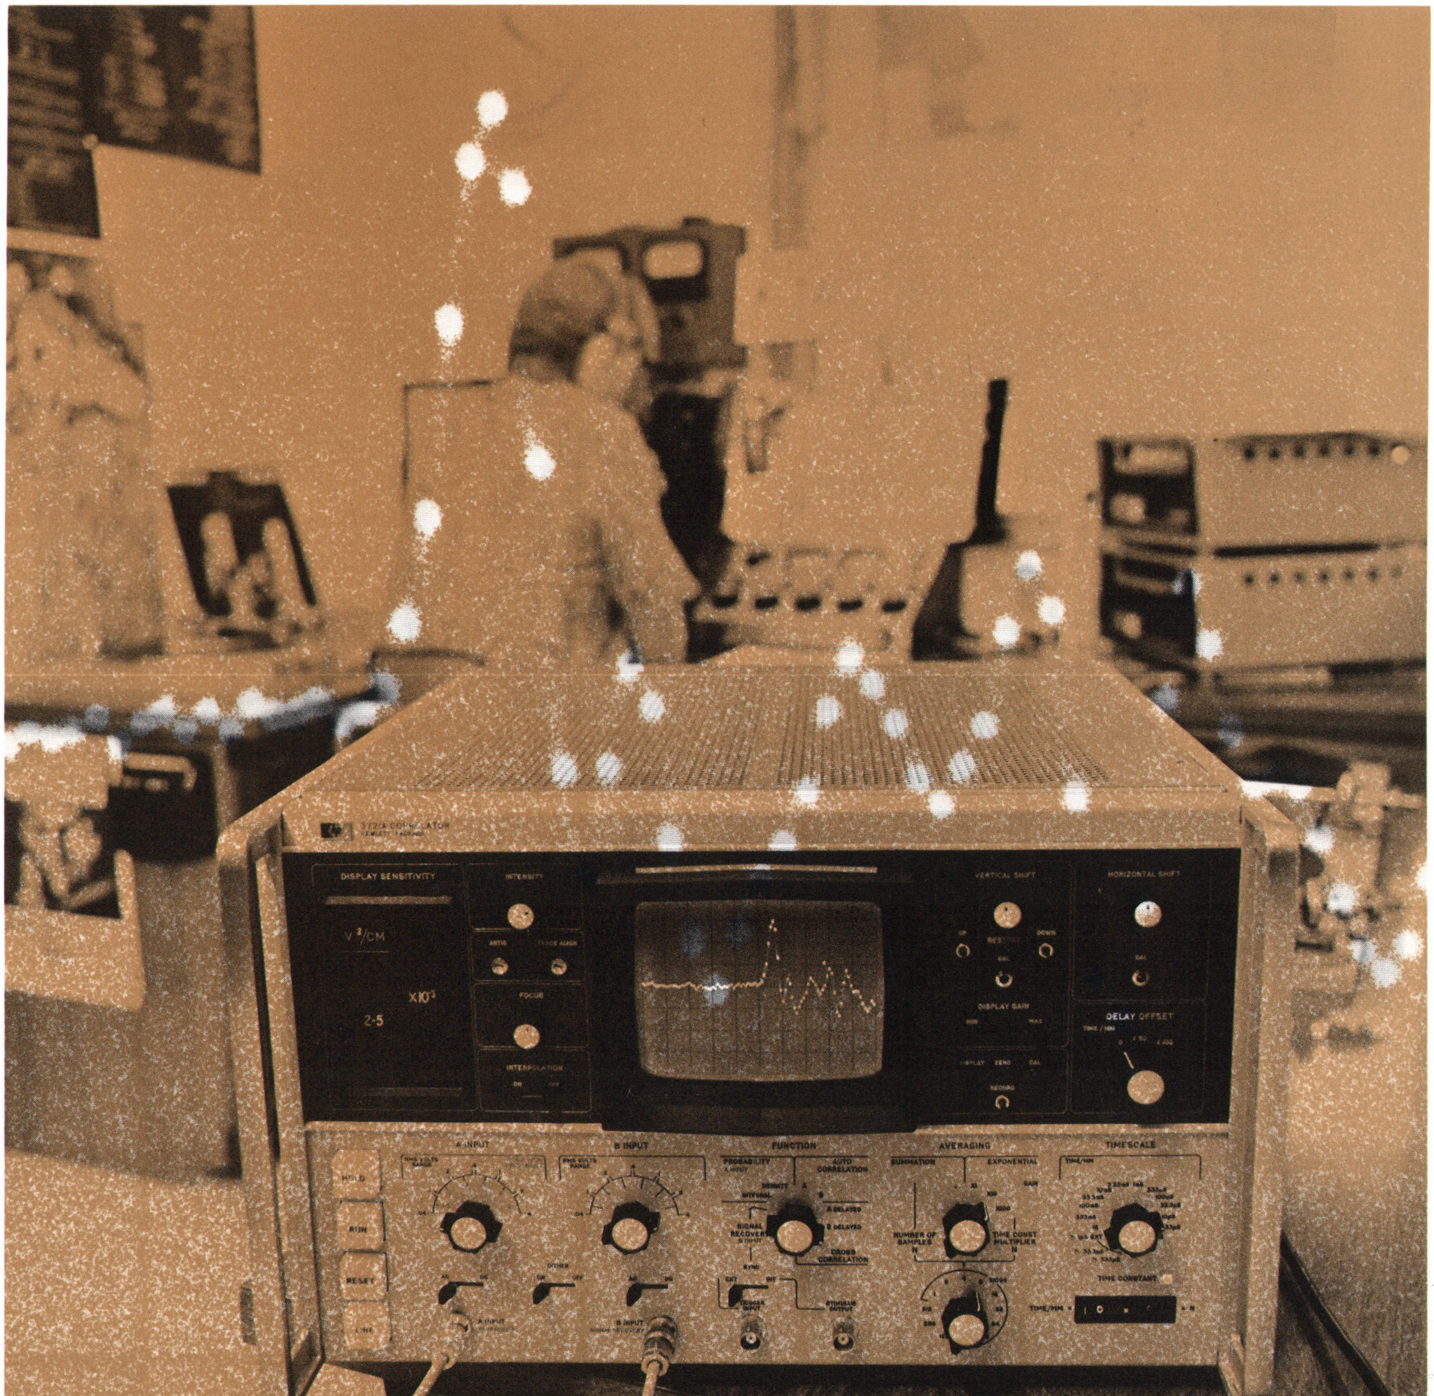

NOVEMBER 1969

Correlation, Signal Averaging, and Probability Analysis

Correlation is a measure of the similarity between two waveforms. It is useful in nearly every kind of research and engineering — electrical, mechanical, acoustical, medical, nuclear, and others. Two other statistical methods of waveform analysis — signal averaging and probability analysis — are also widely useful.

By Richard L. Rex and Gordon T. Roberts

SOME THINGS, LIKE TWO PEAS IN A POD, ARE SIMILAR; others, like chalk and cheese, are not. Throughout science, however, we find instances where the situation is not as clear as this, and where it is desirable to establish a *measure* of the similarity between two quantities. Correlation is such a measure.

As it applies to waveforms, correlation is a method of time-domain analysis that is particularly useful for detecting periodic signals buried in noise, for establishing coherence between random signals, and for establishing the sources of signals and their transmission times. Its applications range from engineering to radar and radio astronomy to medical, nuclear, and acoustical research, and include such practical things as detecting leaks in pipelines and measuring the speed of a hot sheet of steel in a rolling mill.

Mathematically, correlation is well covered in the existing literature,[1] and the use of correlation for research purposes has been established for more than a decade. Until recently, however, correlation in practice has been a complex and time-consuming operation involving, in most cases, two separate processes — data recording and computer analysis. For this reason, correlation techniques could hardly be considered for routine use. Today, correlation in *real time* is entirely practicable, and there seems little doubt that the techniques will soon take their place in all fields of engineering and scientific research.

In this article we present the basic principles of correlation theory. Also included are brief discussions of the concepts of signal averaging and probability density functions. The article on page 9 describes a new HP instrument which applies these ideas in real time, and the article on page 17 deals with applications of this instrument.

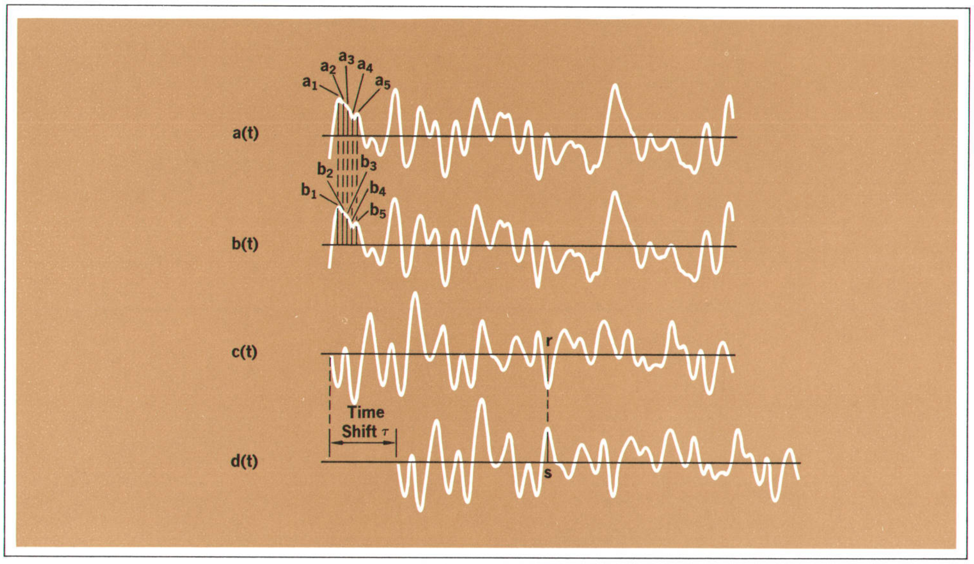

Fig. 1. *Correlation is a measure of the similarity between two waveforms. It is computed by multiplying the waveforms ordinate by ordinate and finding the average product. Here waveforms a(t) and b(t) are identical, so the correlation between them is large. Waveforms c(t) and d(t) are identical in shape, but there is a time shift between them, so the correlation between them is smaller than that between a(t) and b(t). Hence correlation is a function of the time shift between two waveforms.*

PRINTED IN U.S.A.

The Basic Propositions

How can we correlate — that is, test for similarity — two waveforms such as a(t) and b(t) of Fig. 1? Of the easily mechanizable processes, the most effective is multiplication of the waveforms, ordinate by ordinate, and addition of the products over the duration of the waveforms. To assess the similarity between a(t) and b(t) we multiply ordinate a_1 by ordinate b_1, a_2 by b_2, a_3 by b_3, and so on, finally adding these products to obtain a single number which is a measure of the similarity. In this example a(t) and b(t) are identical, with the result that every ordinate—positive or negative—yields a positive product. The final sum is therefore large. If, however, the waveforms were dissimilar, some products would be positive and some would be negative. There would be a tendency for the products to cancel, so the final sum would be smaller.

Now consider waveform c(t) of Fig. 1 and the same waveform shifted in time, d(t). If the time shift (usually denoted by the symbol τ) were zero then we would have the same conditions as before, that is, the waveforms would be in phase and the final sum of the products would be large. If the time shift τ is made large, the waveforms appear dissimilar (for example, ordinates r and s have no apparent connection) and the final sum is small.

Going one step farther, we can find the *average product* for each time shift by dividing each final sum by the number of products contributing to it. If now we plot the average product as a function of time shift, the resulting curve will show a positive maximum at $\tau = 0$, and will diminish to zero as τ increases. The peak at $\tau = 0$ is equal to the mean square value of the waveform. This curve is called the *autocorrelation function* of the waveform. The autocorrelation function $R(\tau)$ of a waveform is a graph of the similarity between the waveform and a time-shifted version of itself, as a function of the time shift. An autocorrelation function has:

- symmetry about $\tau = 0$, i.e., $R(\tau) = R(-\tau)$
- a positive maximum at $\tau = 0$ equal to the mean square value $(\overline{x^2})$ of the signal from which it is derived, i.e., $R(0) = \overline{x^2}$, and $R(0) \geq R(\tau)$ for all τ.

Also note the special case of the periodic waveform — the autocorrelation function of *any* periodic waveform is periodic, and has the same period as the waveform itself. An example of this is the pseudo-random binary sequence, Fig. 2, the autocorrelation function of which is a series of triangular functions.[2,3]

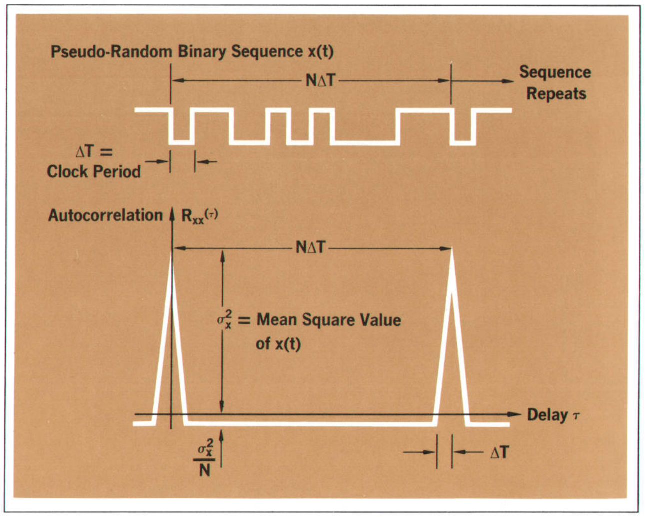

Fig. 2. *The autocorrelation function is a measure of the similarity between a waveform and a shifted version of itself. Here is a pseudo-random binary sequence and its autocorrelation function. The autocorrelation function of a periodic waveform has the same period as the waveform.*

The random noise-like signal shown in Fig. 3 is quite different from the periodic waveform. When compared with a time shifted version of itself, only a small time shift is required to destroy the similarity, and the similarity never recurs. The autocorrelation function is therefore a sharp impulse which decays from the central maximum to low values at large time shifts. Intuitively, the width of the 'impulse' can be seen to depend on the mean zero-crossing rate of the noise waveform, that is, on the bandwidth of the noise. The higher the zero-crossing rate, the smaller the time shift required to destroy similarity.

Two samples of noise of the same bandwidth might have quite different waveforms, but their autocorrelation functions could be identical. The autocorrelation function of any signal, random or periodic, depends not on the actual waveform, but on its frequency content.

Cover: *Model 3721A Correlator displays the crosscorrelation between wideband noise coming from a loudspeaker and the output of a microphone several feet from the speaker, both in the studio of FM station KTAO, Los Gatos, California. Each peak corresponds to a different speaker-to-microphone sound path, the first peak to the direct path and the others to various bounce paths. Acoustic absorption coefficients of the studio walls and furnishings could be found by analyzing the peak heights.*

In this Issue:

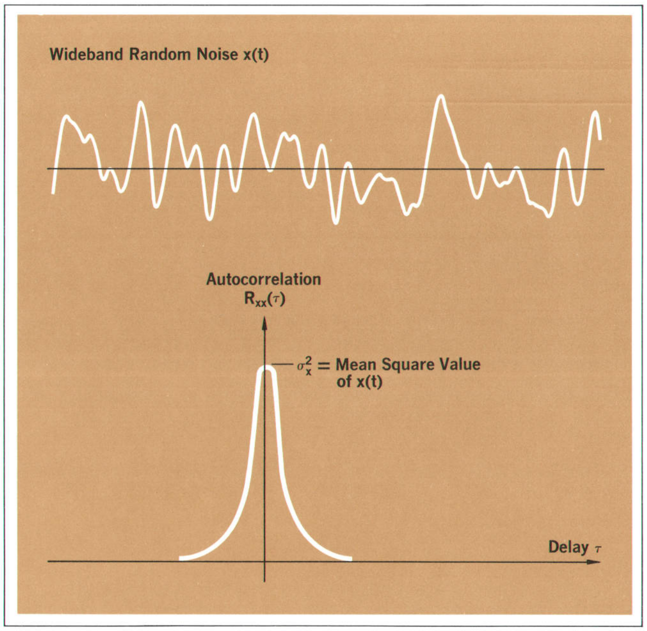

Fig. 3. *The autocorrelation function of a wideband non-periodic waveform is non-periodic and narrowly peaked. The wider the bandwidth, the narrower the peak.*

The Mathematics of Autocorrelation

The autocorrelation function of a waveform x(t) is defined as:

$$R_{xx}(\tau) = \lim_{T\to\infty} \frac{1}{T} \int_0^T x(t)\, x(t-\tau)\, dt \qquad (1)$$

that is, the waveform x(t) is multiplied by a delayed version of itself, $x(t-\tau)$, and the product is averaged over T seconds. Another way we can write this is $R_{xx}(\tau) = \overline{x(t)\,x(t-\tau)}$. The continuous averaging process implied by this expression for $R_{xx}(\tau)$ can be accomplished by analog methods, but in a digital system, it is more convenient to approximate this average by sampling the signal every Δt seconds, and then summing a finite number, N, of the sample products.

$$R_{xx}(\tau) = \frac{1}{N} \sum_{k=1}^{k=N} x(k\Delta t)\, x(k\Delta t-\tau) \qquad (2)$$

This would be computed for several values of τ. The range of τ over which $R_{xx}(\tau)$ is of interest depends on the bandwidth of the signal x(t). For example, the autocorrelation function of a 1 MHz signal could be computed for values of τ ranging from zero to, say, 10 μs with 100 ns resolution in τ. Likewise, a 100 Hz signal would perhaps be analyzed with τ going from zero to 100 ms with 1 ms resolution.

It might be implied from equation 2 that for a good approximation the interval between pairs of samples, Δt, should be of the same order as the chosen resolution or increment in τ. This is not true. The sampling interval Δt can be very large in relation to the resolution in τ. The point to note here is that we are computing signal *statistics*. In other words, we are looking for a measure of average behavior — we don't need to reconstruct the actual waveshape. Hence, the requirements of Shannon's sampling theorem (sampling rate greater than twice the highest signal frequency) need not necessarily be met. Provided the signal statistics do not change with time (i.e. provided the signal statistics are stationary), it does not matter how infrequently the pairs of samples are taken. They need not even be taken at regular intervals — sampling at random intervals is quite acceptable (the HP Model 3406A Random Sampling Voltmeter works on this principle). The important factor is the absolute number, N, of samples taken, and not the rate at which they are taken. The statistical error decreases as N increases.

Relationship of Autocorrelation Function and Power Density Spectrum

We saw in Fig. 3 that wideband signals are associated with narrow autocorrelation functions, and vice versa. There is in fact a specific relationship between a signal's power density spectrum and its autocorrelation function. They are a Fourier transform pair.

$$R_{xx}(\tau) = \int_0^\infty G_{xx}(f) \cos 2\pi f\tau\, df \qquad (3)$$

$$G_{xx}(f) = 2\int_{-\infty}^\infty R_{xx}(\tau) \cos 2\pi f\tau\, d\tau \qquad (4)$$

where $G_{xx}(f)$ is the *measurable* power density spectrum existing for positive frequencies only, that is, $G_{xx}(f)$ is what we could measure with a wave analyzer having a true square-law meter.

Crosscorrelation

If autocorrelation is concerned with the similarity between a waveform and a time shifted version of itself, then it is reasonable to suppose that the same technique could be used to measure the similarity between two *non-identical waveforms,* x(t) and y(t).

The *crosscorrelation* function is defined as

$$R_{xy}(\tau) = \overline{x(t-\tau)\, y(t)}.$$

Like the autocorrelation function, the crosscorrelation function can be approximated by the sampling method:

$$R_{xy}(\tau) = \frac{1}{N} \sum_{k=1}^{k=N} x(k\Delta t-\tau)\, y(k\Delta t)$$

Again the choice of Δt is not critical. Δt can be large compared with the chosen resolution in τ, and need not necessarily be constant throughout the process. Fig. 4 illustrates the process of crosscorrelation. Relationships similar to equations 3 and 4 exist for crosscorrelation functions and functions called cross power spectral densities.[1]

Uses of Correlation

Before discussing any more theory, we will pause now and talk about some of the things correlation is good for. We will be fairly general here; more specific applications will be covered in the article beginning on page 17.

Detection of Signals Hidden in Noise

The first practical example we consider represents the basic problem of all communications and echo ranging systems. A signal of known waveform is transmitted into a medium and is received again, unchanged in form but buried in noise. What is the best way of detecting the signal? The receiver output consists of two parts: the desired signal, and the unwanted noise. If we crosscorrelate the transmitted signal with the receiver output then the result will also have two components; one part is the autocorrelation function of the desired signal (which is common to both of the waveforms being crosscorrelated), and the other part results from the crosscorrelation of the desired signal with unwanted noise. Now in general there is *no correlation between signal and noise,* so the second part will tend to zero, leaving only the signal — in the form of its autocorrelation function. Crosscorrelation has thus rejected the noise in the received signal, with the result that the signal-to-noise ratio is dramatically increased. A simulation of typical transmitted and received signals is recorded in Fig. 5. Fig. 5(a) shows a swept-frequency 'transmitted' signal. To simulate the received signal this swept signal was delayed, then added to wide-band noise as shown in Fig. 5(b). The signal-to-noise ratio was about −10 dB. The result of crosscorrelating the transmitted and received signals is shown at 5(c). The signal appears clearly — in the form of its autocorrelation function — while the noise has been completely rejected. Note that the autocorrelation function is displaced from the time-shift origin at the left side of the correlogram; this is because of the delay between transmission and reception.

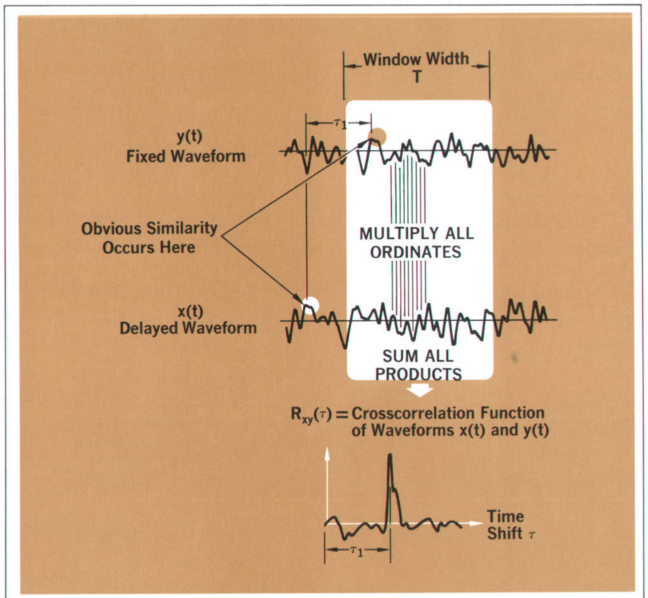

Fig. 4. *Crosscorrelation function shows the similarity between two non-identical waveforms as a function of the time shift between them. Here the peak in the crosscorrelation function of x(t) and y(t) shows that at a time shift τ, there is a marked similarity between the waveforms. The similarity is clearly visible in this example, but crosscorrelation is a very sensitive means of signal analysis which can reveal similarities undetectable by other methods.*

The crosscorrelation method of signal detection is not confined to swept-frequency waveforms. For example, the method (under the name of phase-sensitive, or coherent detection) has been used for many years in communications systems to recover sinusoidal signals from noise.

Crosscorrelation requires a reference signal which in most cases will be similar to the signal to be recovered. Hence the method is clearly unusable for the detection of unknown signals. However, for detecting unknown *periodic* signals, *autocorrelation* is uniquely successful. Autocorrelation reveals periodic components in a noisy signal without the need for a reference signal. Why don't we always use autocorrelation? Simply for the reason that an autocorrelation function contains no phase, or relative timing, information. Autocorrelation could not show, for example, the delay between the transmitted and received signals in Fig. 5. What it can show, however, is unsuspected periodicities. A striking application of autocorrelation is the detection of periodic signals from outer space — for example, emissions from pulsars.

Signal Averaging or Signal Recovery

Before leaving the subject of periodic signal detection, we shall consider a way in which crosscorrelation can recover actual *waveshape*. This can be done by crosscorrelating the noisy signal not with a replica of the hidden periodic waveform, but with a constant-amplitude pulse synchronized with the repetitions of the waveform. In a digital system this would be accomplished by taking a series of samples of the noisy signal, then multiplying each sample by the pulse amplitude. Corresponding products (that is, samples) from each series are then averaged. The periodic waveform is reinforced at each repetition, but any noise present in the signal — since it contributes random positive and negative amounts to the samples — averages to zero. This technique of wave-shape recovery — known as *signal averaging* or *signal recovery* — finds wide application in spectroscopy, biological sciences and vibration analysis.[4]

Finding Relationships between Random Signals

Suppose we have two random signals, neither of which in itself contains meaningful information, yet we know that the two signals are related and that there is information contained in the relationship. Examples would be a transmitted signal and a received signal, or the apparently random waveforms appearing at the input and output of an element in a complex system. Crosscorrelation provides a method of extracting information about the relationships between such signals.

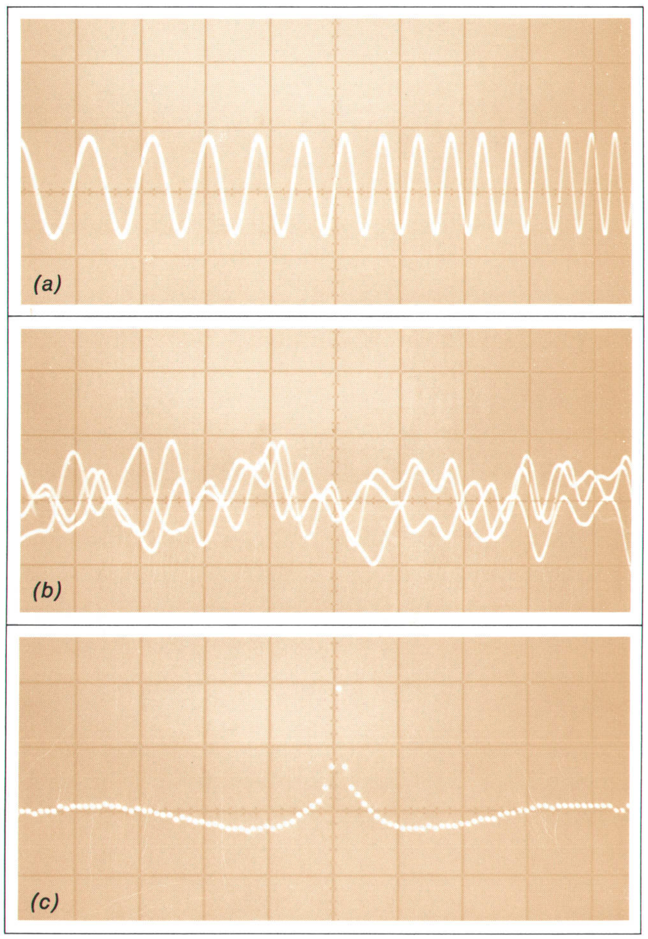

Fig. 5. *Simulation of the ability of crosscorrelation to detect a known signal buried in noise.*
(a) 'Transmitted' signal, a swept-frequency sine wave.
(b) 'Received' signal, the swept sine wave plus noise. S/N = −10 dB.
(c) Result of crosscorrelating the transmitted and received signals. Distance from left edge to peak represents transmission delay.

The simplest relationship is pure delay, which would be revealed as a displacement of the correlation function's peak from the time-shift origin (see Fig. 5). In real situations, however, more things happen than just delays. Distortion, smoothing, and other frequency-dependent changes also occur. The *shape* of the crosscorrelation function can reveal the nature of these changes. An important example is the use of crosscorrelation in system identification.

System Identification without Disturbing the System

By the system identification problem, we mean that any system, from a power station to a simple RLC network, is treated as an unknown black box, and we wish to identify the black box, that is, to find sufficient information about it to predict its response to any input. Note that we do not wish to know the precise nature of each individual component of the system. Normally we are not interested in such detail. If we can tell in advance the system's output in, say, megawatts in response to an arbitrary input, perhaps kilograms of uranium, then we are satisfied.

A complete description of a system may be given in either the frequency or the time domain. The frequency-domain type of description will already be familiar to many readers. If we are given the amplitude-vs-frequency and phase-vs-frequency characteristics of the system, then we can obtain a complete description of the system's output by multiplying the input amplitude spectrum by the system's amplitude characteristic and adding the input phase spectrum to the system's phase characteristic.

The equivalent description in the time domain is the impulse response: a unit impulse applied to a system produces an output signal called the impulse response of the system. If we know a linear system's impulse response, we can calculate its response to any input by convolution: the impulse response convolved with the input gives the output.[5]

How can we determine the impulse response of the system? The obvious method is to inject an impulse and observe the result; this method, however, has the disadvantage that the response can be very difficult to detect (perhaps buried in background noise) unless a large-amplitude pulse is applied — with, of course, a disturbing effect on the system's normal functioning. We are looking for a general method, workable with all systems from passive networks right through to process control installations which can *never* be taken off-line for evaluation purposes.

How Crosscorrelation Helps

An approximation to the *impulse response* of a linear system can be determined by applying a suitable *noise* signal to the system input, then *crosscorrelating* the noise signal with the system output signal. The equipment setup is illustrated on page 17 and the mathematical theory is given in the appendix, page 8.

The noise test signal can be applied at a *very low level,* resulting in almost no system disturbance and very small perturbations at the output. Crosscorrelation, which is essentially a process of accumulation, builds up the result over a long period of time. Hence, although the perturbations may be very small, a measurable result can be obtained provided that the averaging time is sufficiently long. Background noise in the system will be uncorrelated with the random test signal, and will therefore be effectively reduced by the correlation process.

A suitable noise signal for this technique is one whose bandwidth is very much greater than the bandwidth of the system, so the system impulse response is a relatively slowly changing function of its argument compared to the autocorrelation function of the noise. In other words, to the system the autocorrelation function of the noise should look like an impulse, or something very close to one. 'White' noise is one possibility. Another is binary pseudo-random noise, which has an autocorrelation function that is very close to an impulse. Pseudo-random noise has the advantage that the averaging time T for the correlation system only needs to be as long as one period of the pseudo-random waveform, i.e., as long as one complete pseudo-random sequence. Unlike random noise, pseudo-random noise introduces no statistical variance into the results, as long as the averaging time T is exactly one sequence length, or an integral number of sequence lengths.[3]

Probability Density Functions

We have discussed two ways of describing a signal — the autocorrelation function and its equivalent in the frequency domain, the power spectrum. Neither of these, however, gives any indication of the waveshape or amplitude-vs-time behavior of the signal (except in the case of a repetitive signal for which a synchronized pulse train is available).

A means of characterizing a random signal's amplitude behavior is to determine the *proportion* of time spent by the signal at all possible amplitudes during a finite period of time. In practical terms, this means totalizing the time spent by the signal in a selection of narrow (δx) amplitude windows, and then dividing the total for each window by the measurement time (T). The curve obtained by plotting the window totals against amplitude is known as the probability density function (pdf) of the signal (see Fig. 6).[1] The area under the pdf between any two amplitudes x_1 and x_2 is equal to the proportion of time that the signal spends between x_1 and x_2. This area is also equal to the probability that the signal's amplitude at any arbitrary time will be between x_1 and x_2. A pdf is always normalized so the total area under it is exactly one.

The most commonly encountered pdf for naturally occurring signals is the Gaussian or normal distribution, Fig. 7. The amplitude (horizontal) scale of the pdf is calibrated in terms of σ, a symbol used in statistics to denote *standard deviation,* a measure of the spread of a set of values about their mean. In general σ is equal to the rms amplitude of the ac component of the signal.

The probability density function can yield important information about nonlinearities in a system. If, for example, a Gaussian-type signal were applied to an amplifier having an insufficient dynamic range, the distorted output would have a pdf with 'ears' indicating that clipping had occurred.

Cumulative Probability Function

The integral of the pdf of a signal is the cumulative probability distribution function (cdf). The cdf represents the probability that a signal will be at or below a certain amplitude. The final value of any cdf is one, for the reason that a signal must spend *all* the time at or below its maximum level. Fig. 8 shows the cdf of a Gaussian signal.

The cdf is sometimes a more convenient function than the pdf for a clear description of a signal's amplitude properties. A square wave, for example, has a step cdf which can present evidence of distortion more clearly than the corresponding pdf, which is just a pair of impulses.

References

1. See, for example, J. S. Bendat and A. G. Piersol, 'Measurement and Analysis of Random Data,' John Wiley and Sons, 1966.
2. G. A. Korn, 'Random Process Simulation and Measurements,' McGraw-Hill Book Company, 1966.

3. G. C. Anderson, B. W. Finnie, and G. T. Roberts, 'Pseudo-Random and Random Test Signals,' **Hewlett-Packard Journal,** September 1967.
4. C. R. Trimble, 'What Is Signal Averaging?,' **Hewlett-Packard Journal,** April 1968.
5. See, for example, M. Schwartz, 'Information Transmission, Modulation, and Noise,' McGraw-Hill Book Company, 1959. Or see any textbook on linear system theory.

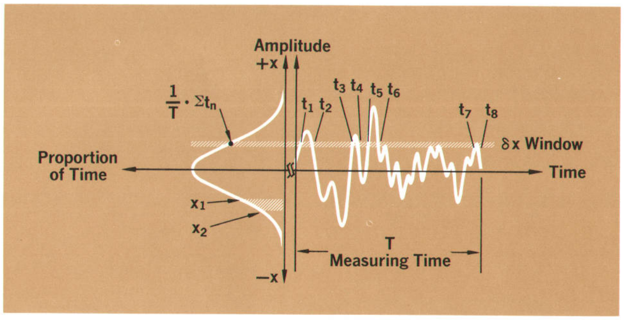

Fig. 6. *The probability density function tells what proportion of time is spent by a signal at various amplitudes. Each point on the density function represents the proportion of the total measurement time T that the signal is in a particular window δx wide, or the probability that the signal will lie in that window. The area under the curve between any two amplitudes x_1 and x_2 is equal to the proportion of time that the signal's amplitude is between x_1 and x_2.*

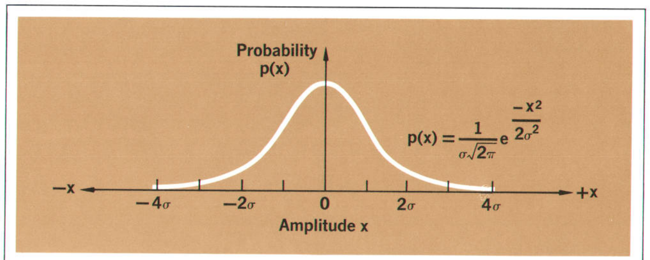

Fig. 7. *Gaussian probability density function is characteristic of many natural disturbances.*

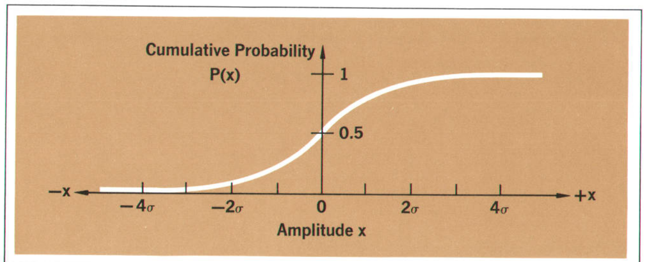

Fig. 8. *Cumulative probability distribution function gives the proportion of time or probability that the signal lies below any given amplitude. It is the integral of the probability density function. $P(-\infty)$ is always zero; $P(\infty)$ is always one. This is the Gaussian distribution function.*

Appendix

Determination of Impulse Response by Crosscorrelation

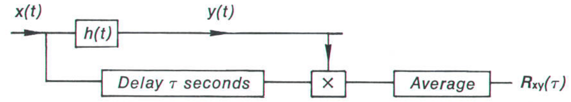

$$R_{xy}(\tau) = \lim_{T\to\infty} \frac{1}{T}\int_0^T x(t-\tau)\, y(t)\, dt$$

$$= \lim_{T\to\infty} \frac{1}{T}\int_0^T x(t-\tau) \int_{-\infty}^{+\infty} h(u)\, x(t-u)\, du\, dt.$$

Interchange order of integration

$$R_{xy}(\tau) = \int_{-\infty}^{+\infty} h(u) \left\{ \lim_{T\to\infty} \frac{1}{T}\int_0^T x(t-\tau)\, x(t-u)\, dt \right\} du.$$

The second integral is seen to be the autocorrelation of x(t) with argument $(\tau-u)$

$$R_{xy}(\tau) = \int_{-\infty}^{+\infty} h(u)\, R_{xx}(\tau-u)\, du.$$

Now, if we choose the bandwidth of the noise to be much greater than the system passband, then h(u) will be a relatively slowly changing function in comparison with $R_{xx}(\tau-u)$. The term h(u) will be almost constant over the small range of values of u around $u = \tau$ for which $R_{xx}(\tau-u)$ has significant values. The integral then becomes:

$$R_{xy}(\tau) = h(\tau) \int_{-\infty}^{+\infty} R_{xx}(\tau-u)\, du.$$

Comparing this with equation 4, page 4, we see that the integral term gives us the power spectral density for f = 0, hence

$$R_{xy}(\tau) = \tfrac{1}{2}\, h(\tau)\, G_{xx}(0)$$
$$= \text{constant} \times \text{impulse response.}$$

$G_{xx}(f)$ is the physically realizable, one sided power density spectrum. $G_{xx}(0)$ is the value of this function in volts squared per Hertz, at very low frequencies.

Richard L. Rex

Until his recent departure from HP, Dick Rex was the product manager concerned with marketing the 3721A Correlator, the 3722A Noise Generator, and related products. He joined HP in 1966.

Dick attended Derby Technical College, then joined the RAF, where he taught radar and communications for two years. After leaving the RAF, he held several industrial positions, mainly in technical marketing. Just before coming to HP, Dick spent three years with a commercial agency specializing in product support functions.

Gordon T. Roberts

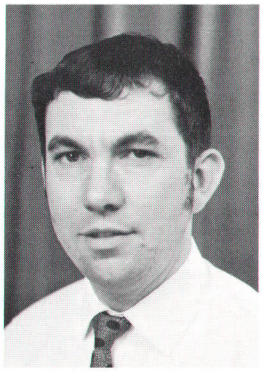

Gordon Roberts is engineering manager of HP Limited, South Queensferry, Scotland. Before he joined HP in 1965, Gordon was at Edinburgh University, where he lectured in control theory and headed a research team investigating the uses of noise signals in systems evaluation.

Gordon graduated from the University of North Wales in 1954 with a B.Sc. degree in electrical engineering, and three years later he received his Ph.D. degree from Manchester University. His doctoral research was concerned with noise in nonlinear systems.

A Calibrated Real-Time Correlator/Averager/Probability Analyzer

This digital signal analyzer computes and displays 100-point autocorrelation functions, crosscorrelation functions, waveshapes of signals buried in noise, probability density functions, and probability distributions.

By George C. Anderson and Michael A. Perry

Ever since the theory of correlation was developed and the potential advantages of the technique brought to light, people have been looking for practical ways to apply it. Because correlation requires prodigious computation, a common way of getting correlation functions has been to record data and process them later, off-line, in a digital computer. The problem with this method is that it takes too much time. If the data are inadequate or if procedures or programs need modification, it takes a long time to find out. In the meantime, a lot of equipment may be tied up.

On-line correlation has the advantage of providing answers where they are needed, that is, where the measurements are being made. But while it isn't unknown, on-line correlation isn't very common, either, the reason being that instruments that can correlate in real time haven't been very accurate, versatile, or easy to use.

The HP Model 3721A Correlator, Fig. 1, is designed to be a truly practical way of getting correlation functions in real time. This new instrument processes analog signals — on line and in real time — and presents the results on a calibrated CRT display. It is more than just a correlator; it has two analog inputs and it can compute

- the autocorrelation function of one input signal
- the crosscorrelation between two signals
- the probability density function of an input signal, or its integral, the probability distribution
- the waveshape of a repetitive signal buried in noise (by signal recovery or signal averaging).*

* The word 'averaging' has two meanings in this article. **Signal averaging** is a technique of signal analysis; it is useful for recovering repetitive signals from noise (see reference 1). **Averaging** is a mathematical process; it is used in all of the modes of operation of the new correlator: autocorrelation, crosscorrelation, probability display, and signal averaging. To avoid confusion we will often refer to signal averaging as 'signal recovery.'

In general, using the new correlator is similar to using an oscilloscope, and in some ways, it is easier to use than an oscilloscope. It has a wide selection of measurement parameters — sampling rates, averaging times, and so on — and the vertical scale factor of the display, which is affected by several of these selectable parameters, is automatically computed and displayed on the front panel. Like an oscilloscope, and unlike off-line computers, the correlator can follow slowly varying signals, can give a 'quick-look' analysis to show the need for and the results of adjustments, and can easily be carried from place to place (it weighs 45 pounds).

Digital Techniques Used

Model 3721A is primarily a digital instrument, although it does use some analog techniques. Long-term stable averaging, which is necessary for the very-low-frequency capability, can only be done digitally; capacitive averaging, which is used in some correlators, simply can't provide long enough time constants with capacitors of reasonable size. A simple, accurate, and stable digital multiplier which is fast enough to allow sampling rates up to 1 MHz is another benefit of digital operation.

Model 3721A displays the computed function on its CRT using 100 points. Typical correlation displays are shown in Fig. 2.

100 points are sufficient for many measurements; however, if more resolution is required, there is an option (for the correlation mode) which provides additional delay in batches of 100 points, up to a maximum of 900 points. Using this pre-computational delay, the time scale can be compressed and the correlation function examined one section at a time.

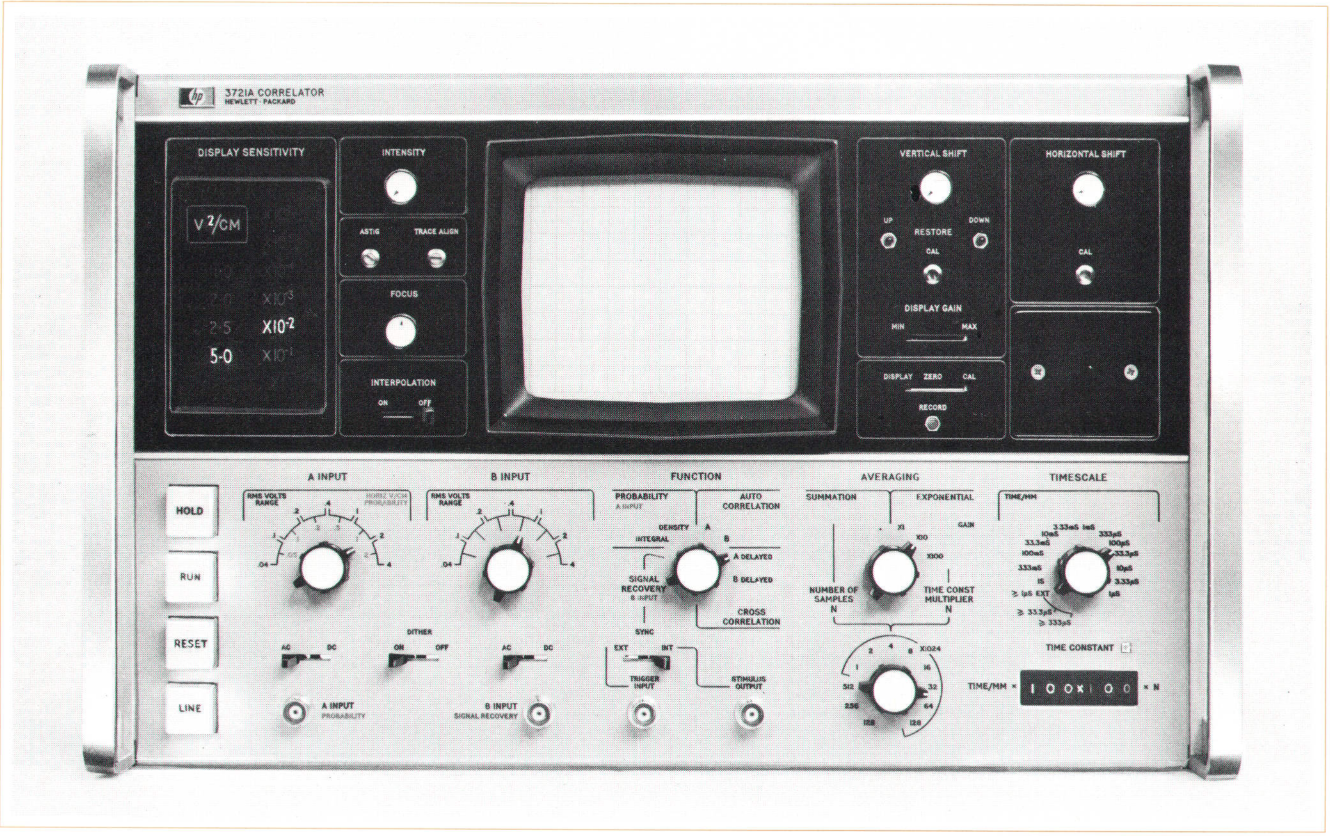

Fig. 1. *Despite its versatility, Model 3721A Correlator is as easy to use as an oscilloscope. An illuminated panel shows the display sensitivity in V^2/cm for correlation and in V/cm for signal averaging (signal recovery). The correlator has a wide selection of measurement parameters and modes of operation, including a 'quick-look' mode which minimizes delays in setting up experiments.*

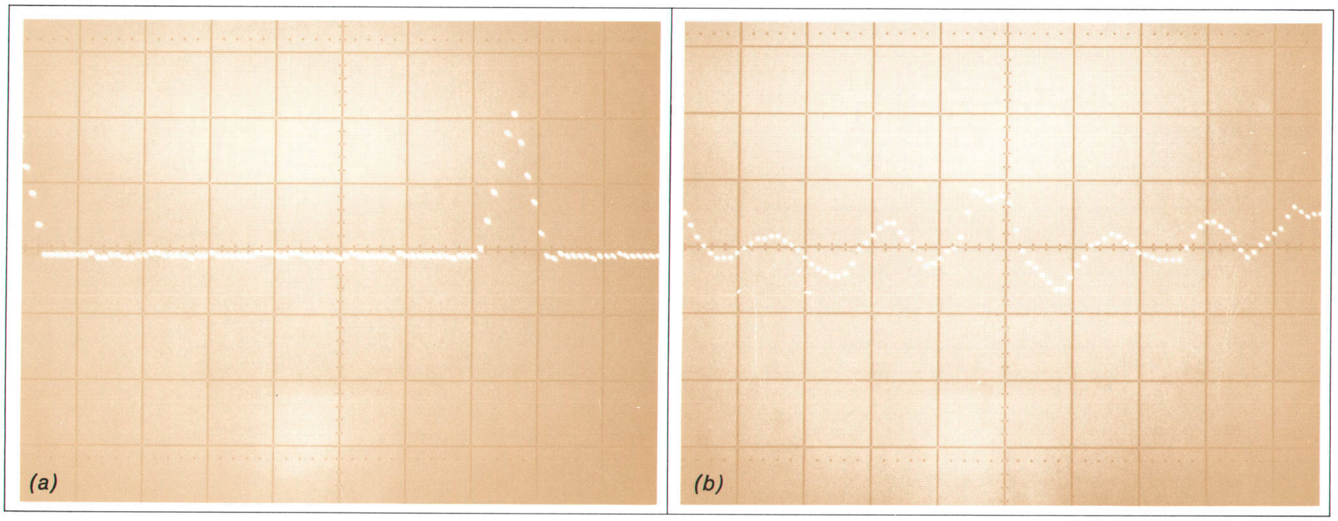

Fig. 2. *Examples of 100-point correlation functions computed and displayed in real time by Model 3721A Correlator. (See also page 20 for autocorrelation functions of speech sounds.)*
(a) Autocorrelation function of 15-bit pseudo-random binary sequence, clock period 5 ms. Time scale 1 ms/mm.
(b) Crosscorrelation between noise in a machine shop and noise from a nearby ventilating fan. Double hump in center shows there are two principal transmission paths from fan to machine shop, differing in propagation delay by 4 ms.

The horizontal axis of the display is scaled ten points to each centimeter and is calibrated in time per millimeter. The time per millimeter is also the period between successive samples of the analog input. The 'sweep rate' can be switched from 1 μs/mm to 10 s/mm. It is controlled by a crystal clock and is accurate within 0.1%. Lower sweep rates may be provided by inputs from an external source.

The vertical axis is accurately calibrated,* and an illuminated display beside the CRT avoids the 'numbers trouble' which can easily occur in an instrument which uses both analog and digital techniques. The vertical scale factor of the display depends upon the product of four numbers. These are the settings of the two analog input amplifiers, a gain control for the averaging algorithm, and a control which expands the trace in the vertical axis. The first two variables follow a 1, 2, 4, 10 sequence, the third follows a 1, 10, 100 sequence, and the fourth follows a binary progression. The scale factor is also affected by the function the instrument is performing. Autocorrelation requires the square of one analog channel gain setting, whereas crosscorrelation requires the product of two. Signal recovery and probability display have entirely different requirements. All of this might leave the user with some awkward mental arithmetic, were it not for the illuminated display, which shows the scale factor.

* The accuracy of a statistical analyzer is difficult to define, since the accuracy of the displayed result depends upon the statistics and bandwidth of the input signals and on the averaging time constant used. Systematic errors in the new correlator — such things as display nonlinearity and variations in quantizer gain (see appendix) — are typically less than 1 or 2 per cent at low frequencies.

Being a digital instrument, the correlator is easily interfaced with a computer; there is a plug-in option for this purpose. There are also rear-panel outputs for an X-Y recorder.

How It Correlates

The equation for the crosscorrelation function of two waveforms a(t) and b(t) which are both functions of time is

$$R_{ba}(\tau) = \overline{a(t)b(t-\tau)}$$

where the bar denotes taking an average. There are three important operations: delaying b(t) by an amount τ, multiplying the delayed b(t) by the current value of a(t), and taking the average value of these products over some time interval. The correlation function is a plot of $R_{ba}(\tau)$ versus the delay, τ. The new correlator computes 100 values of the correlation function for 100 equally spaced values of τ and does them simultaneously, as follows.

Input waveforms a(t) and b(t) are sampled and the current sample of b(t) is loaded into location 0 of the delay store (see Fig. 3). Locations 1 through 99 contain previous values of b(t), all stacked in order. To compute its estimate of the correlation function, the correlator calculates

$$R_{ba}(m\Delta\tau) = \overline{a(t)b(t-m\Delta\tau)}$$

for 100 values of m, that is, for m = 0, 1, 2, . . . , 99. For this operation, switch S1 in Fig. 3 is positioned so the contents of the delay store rotate until they arrive back in their original positions. During this rotate operation, each delayed value of b(t) passes the multiplier, where it is multiplied with the stored value of a(t) to

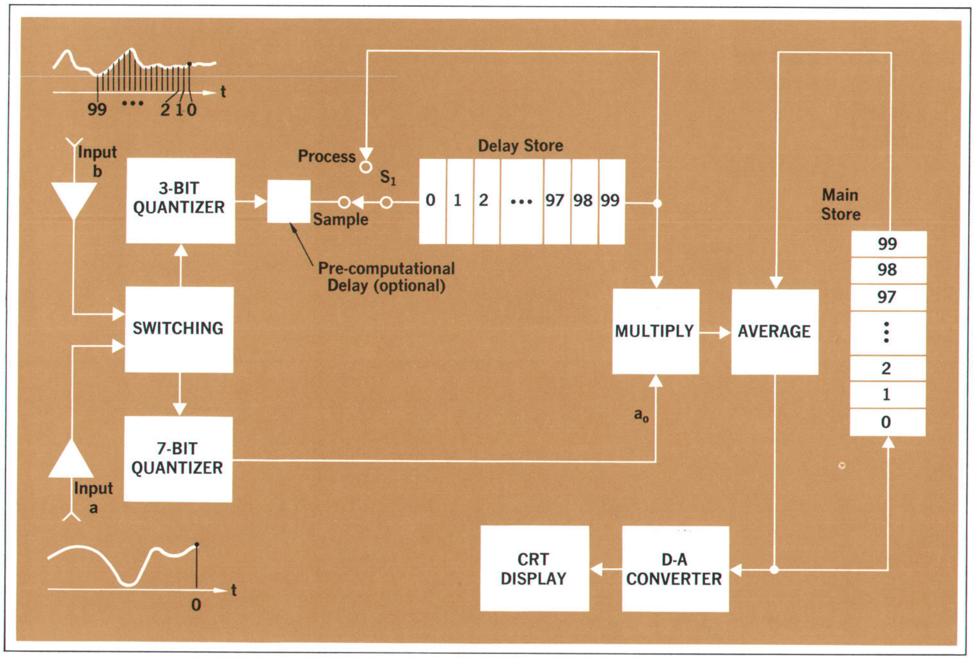

Fig. 3. *Correlator samples inputs, converts samples to digital form, delays one input with respect to the other, then multiplies the two and updates the 100 average products in the main store. Three-bit quantizer allows fast operation and doesn't impair accuracy (see appendix).*

produce a series of 100 products $a(t)b(t-m\Delta\tau)$. The products are fed into an averager where they interact with the contents of the serial main store, which are also undergoing a rotate operation. The contents of the main store are the 100 previously calculated values of $R_{ba}(m\Delta\tau)$. In the averager, the new product a(t)b(t) updates $R_{ba}(0)$, the product $a(t)b(t-\Delta\tau)$ updates $R_{ba}(\Delta\tau)$, and so on up to $R_{ba}(99\Delta\tau)$. The main store, therefore, always contains the most recent estimate of the correlation function. A complete sequence of rotation, multiplication, and averaging is known as a *process cycle*. When it is complete, the instrument reverts to a *data acquisition cycle* and new samples of a(t) and b(t) are taken. The new value of b(t) is put into location 0 of the delay store, and all others are shifted down one place. $b(t-99\Delta\tau)$ 'rolls off the end' and is discarded, since it is of no further value. When pre-computational delay is needed, it is added by increasing the length of the delay store.

The time taken up by a process cycle is 135.6 μs, the cycle time of the main store. How then, you may ask, can the instrument operate with a sampling rate as high as 1 MHz, yet enter a process cycle after each sample has been taken? The answer lies in the assumption that the input signals are stationary; this is a way of saying that the statistics of the signals are constant for all time. When this is true, we can wait as long as we wish between successive data samples without affecting the accuracy of the measurement, as long as we take enough samples eventually.

At high data sampling rates ($\Delta\tau < 333$ μs) the correlator operates in an ensemble sampling mode, or batch mode (see Fig. 4). In the batch mode, instead of entering one new sample of b(t) into the delay store at the instant a(t) is sampled, the sampling of a(t) is preceded by taking 99 fresh samples of b(t) and entering those into the delay store first.* We now have stored information showing the relationship between a(t) and 100 consecutive samples of b(t) as before, but now the speed with which data can be acquired is independent of the time taken for a process cycle. There is a tradeoff involved in batch sampling. For statistical accuracy, a large number of samples of a(t) must be taken, and this makes the measurement less efficient than the normal mode which was described first. However, batch sampling is only used at high sampling rates where the extra *real time* involved is too small to be significant in many cases.

* For $\Delta\tau = 1$, 3.33, or 10 μs, 99 samples of b(t) are taken, followed by simultaneous sampling of a(t) and b(t). For $\Delta\tau = 33.3$ or 100 μs, only 9 samples of b(t) precede the simultaneous sampling of a(t) and b(t). For $\Delta\tau \geq 333$ μs, a(t) and b(t) are always sampled simultaneously, i.e., the instrument operates in the normal mode.

Signal Recovery

It is the batch mode of operation which allows the Model 3721A Correlator to perform signal averaging or signal recovery. In this case, there is no signal a(t). A constant value is inserted into the a(t) storage register and, on receipt of a sync pulse, 100 sampled values of b(t), the signal being averaged, are entered into the delay store. A process cycle rotates the delay store and updates the information in the main store. Since the delay store has a much faster access time than the main store, this method of operation allows fast sampling at rates independent of the main-store cycle time.

Two Averaging Modes

Two kinds of averaging are available for correlation, probability measurements, or signal recovery. A front-panel switch selects one or the other.

One averaging mode is straighforward summation. The function computed after N samples of the inputs have been taken at intervals of $\Delta\tau$, starting at $t = 0$, is

$$R_{ba}(m\Delta\tau) = \frac{1}{N}\sum_{k=m}^{N+m-1} a(k\Delta\tau)b(k\Delta\tau - m\Delta\tau),$$

where $m = 0, 1, 2, \ldots, 99$. For signal averaging, or signal recovery, $a(k\Delta\tau)$ is held constant, while for correlation it varies with the input a(t).

The second averaging mode is a running average which forgets very old information. If A_{N-1} is the current value of the running average and a new input I_N arrives, the new average would be

$$A_N = A_{N-1} + \frac{I_N - A_{N-1}}{N},$$

where N is the number of samples that has been taken. Because it is difficult to divide by an arbitrary integer N at a reasonable cost in a high-speed system, the new correlator divides by the nearest power of 2, so it can divide simply by shifting. The algorithm becomes

$$A_N = A_{N-1} + \frac{I_N - A_{N-1}}{2^n}$$

No error is introduced into the averages by dividing by 2^n instead of by N; the averages just take a little longer to approach within a given percentage of their final values (assuming these final values aren't changing).[2]

The running-average algorithm acts like an RC averaging circuit; the output responds exponentially to a step input with a time constant approximately equal to $2^n\Delta\tau$ seconds. A powerful feature of this exponential averaging algorithm is the effect of the varying value of n during the measurement. Initially, n is zero. Then, as the measurement proceeds, N, the number of samples, becomes equal to the next higher value of 2^n, and n is incremented. This has the effect of averaging with a variable capacitor C whose value is initially small, and becomes larger as the measurement proceeds. High-frequency noise is averaged out immediately, leaving the lower frequencies for later and giving a rapid preliminary estimate of the correlation function or average. Once n has reach the terminal value determined by the front-panel time-constant switch, the system continues to operate with that time constant. A wide range of control over the value of n provides the user with a useful and flexible averaging capability.

Details of Operation

Figs. 3 and 5 together show the correlator's block diagram. Analog signals are scaled by two preamplifiers and fed into a changeover switching network, which routes signals into the reference or delay channels according to the function chosen by a front-panel switch. Next, the signals are digitized or quantized. The reference signal is converted into a seven-bit, two's complement code by a ramp-and-counter-type analog-to-digital converter whose conversion time has a maximum value of 42 μs. The signal which is to be delayed with respect to the reference signal is fed into a fast (<1 μs) three-bit converter which has a special encoding characteristic (see Fig. 6). Rigorous computer analysis showed that the coarseness of the three-bit encoding law would not impair the accuracy of the correlator.* The choice of a three-bit

* See the appendix, page 15.

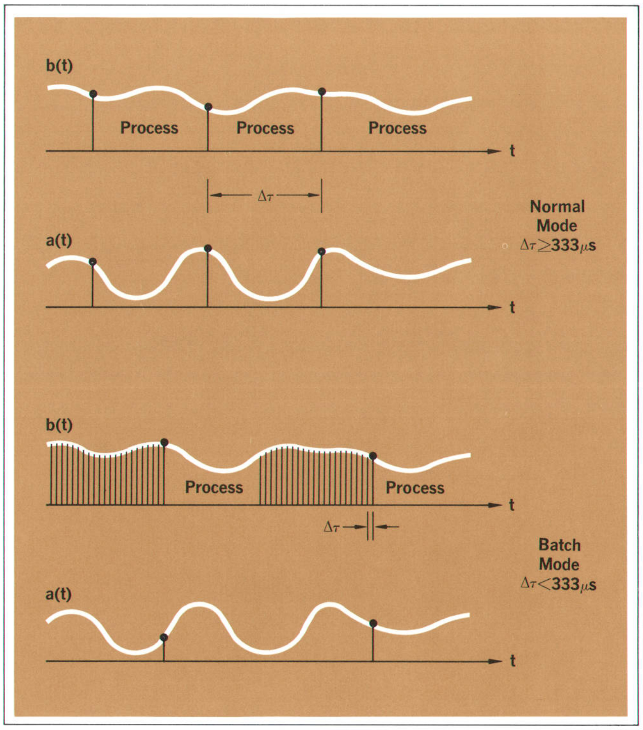

Fig. 4. *For delay increments $\Delta\tau$(time/mm) of 333 μs or more, the correlator operates in the normal mode, going through a process cycle — updating the 100 stored average products — after each sample of the inputs is taken. For shorter delay increments, the correlator uses the batch mode, taking 100 samples of input b for each sample of input a. Although much slower, the batch mode is used only at very short delay increments and isn't a practical limitation.*

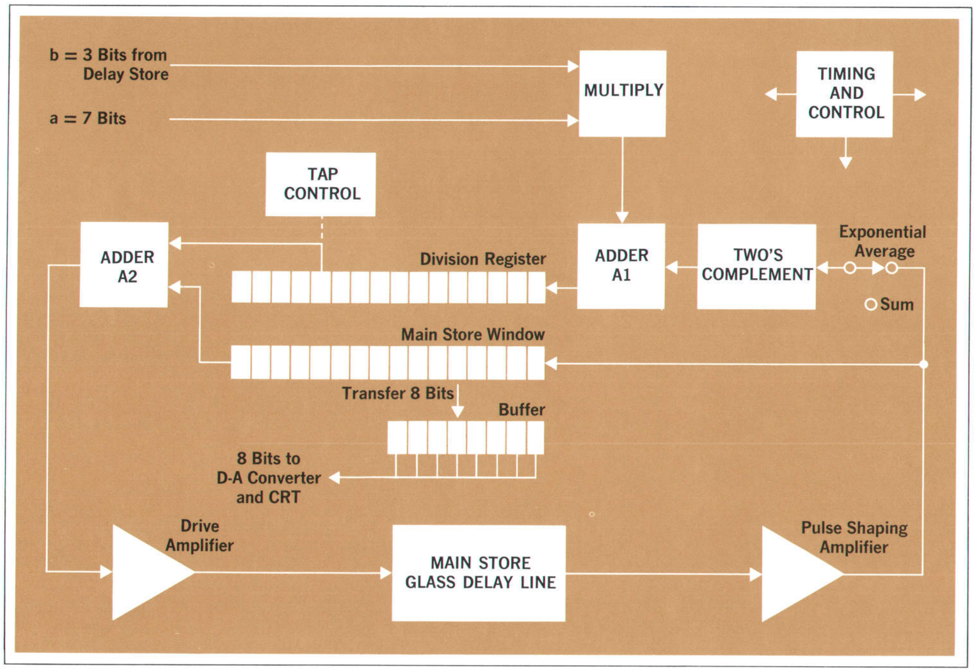

Fig. 5. *The arithmetic unit computes averages digitally, so it doesn't have the stability problems of capacitive averaging. Multiplications and divisions, all by powers of two, require only time shifts.*

code for the delay channel made it possible to provide a fast encoder at relatively low cost, and to keep the capacity of the delay store to a minimal 300 bits.

In the three-bit encoder the signal is coarsely quantized into values of 0, ±1, ±2, or ±4. The virtue of this simple law is that to multiply the quantized output by another binary word involves only shift operations; this saves hardware in the multiplier.

The current estimate of the correlation function (or average) is stored serially in 100 24-bit words in the main store, which is a glass ultrasonic delay line having a capacity of 102 24-bit words. The two extra words are used for bookkeeping and control. Information in the main store recirculates at a bit rate of approximately 18 MHz, one complete cycle taking about 135.6 μs. Operation of the glass delay line is similar to shouting down a long tunnel, catching the sound at the other end, and feeding it back to the beginning.

All arithmetic processing is carried out on the output of the main store using two's complement arithmetic. Subtraction is done by a 'two's complement and add' algorithm. Multiplications and divisions are all by powers of two, so they are implemented as time delays and time advances, respectively.

Display

Output from the memory to the display is via the shift register called the 'main store window'. For each of the 100 computed points, eight bits from this register are transferred to a buffer register, converted to analog form, and used to position the CRT beam vertically. The horizontal position of each dot is determined by a counter. Which eight of the 24 possible bits are displayed depends on the display gain setting, which has three values.

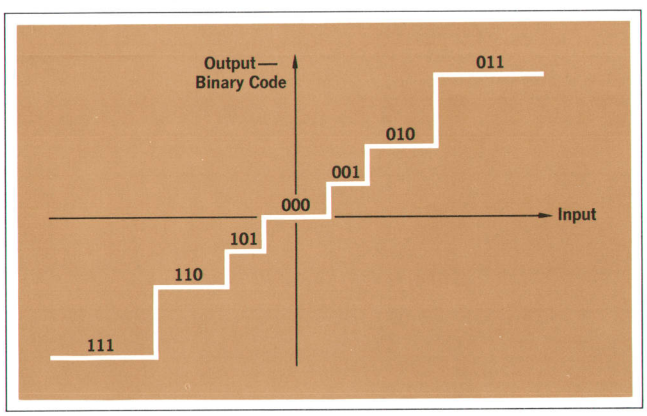

Fig. 6. *The three-bit quantizer's outputs are powers of two. Multiplication by them is simply a gating function.*

Probability Distributions

The measurement of probability density and probability distribution functions depends for its operation upon the constant cycle time of the main store. When the first word in the chain emerges from the main store, the input signal is sampled at the input to the seven-bit ramp-and-counter A-D converter. This time we are not concerned with converting the sampled voltage into binary code, but into a time delay. The sampling instant initiates a ramp which runs down linearly from the value of the sampled input signal to a reference voltage. When coincidence occurs, information is gated into adder A1 to increment the word which is currently in the arithmetic unit. In the probability density mode, no further words are incremented; but in the probability distribution mode, which is the integral of the density, all succeeding words are also incremented. The averaging algorithms operate as before. An example of a probability display is shown in Fig. 7.

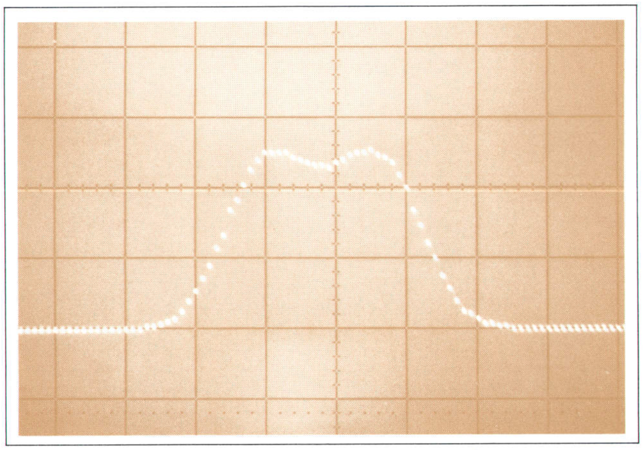

Fig. 7. *Probability density function of sine wave plus Gaussian noise. Model 3721A Correlator can also compute the cumulative distribution function, which is the integral of the probability density function.*

Acknowledgments

The feasibility studies for the Model 3721A Correlator were initiated by Brian Finnie. Associated with us in the electronic design were Glyn Harris, Alister McParland, David Morrison, Rajni Patel, and John Pickering. Peter Doodson did the industrial design. Product design was by Duncan Reid and David Heath. Jerry Whitman assisted with the feasibility studies.

References

1. C. R. Trimble, 'What Is Signal Averaging?,' **Hewlett-Packard Journal,** April 1968.
2. J. E. Deardorff and C. R. Trimble, 'Calibrated Real-Time Signal Averaging,' **Hewlett-Packard Journal,** April, 1968.

Appendix: Equivalent Gain of a Quantizer

A quantizer or analog-to-digital converter is a nonlinear device which has a staircase characteristic like the one shown in Fig. 6, page 14. Quantization has the effect of adding distortion in the form of 'corners' to an input signal x(t). This distortion, called quantization noise, can be shown to be uncorrelated with the input signal.* If the quantizer is assumed to be sufficiently wideband not to distort the signal in other ways, its output will have a component proportional to the input signal, Kx(t), plus uncorrelated quantization noise, n(t). The constant K is called the equivalent gain of the quantizer.

It has been known for some time that accurate correlation measurements can be made using very coarse quantization of data, i.e., very few bits. Correlation involves an averaging process, so a correlator is only affected by the *averaged* response of the quantizer to an input signal. This averaged response is the equivalent gain K. For accurate correlation measurements, K should be constant over the specified range of signal amplitudes. The curve below is the equivalent gain of the three-bit quantizer used in the Model 3721A Correlator, for a Gaussian input signal. The sharp drop in gain at low signal levels occurs because the signal barely climbs over the first step of the quantizer's staircase characteristic. This defect is overcome by adding uncorrelated Gaussian noise to the input signal, thereby keeping the equivalent gain nearly constant down to zero input. The falloff in gain at high signal levels can be attributed to saturation.

The fact that accurate correlation measurements can be made using only a three-bit quantizer is significant. In the case of the Model 3721A Correlator, it means reduced complexity and cost, and higher speed.

* D. G. Watts, 'A Study of Amplitude Quantisation with Application to Correlation Determination,' Ph.D. Thesis, University of London, January 1962.

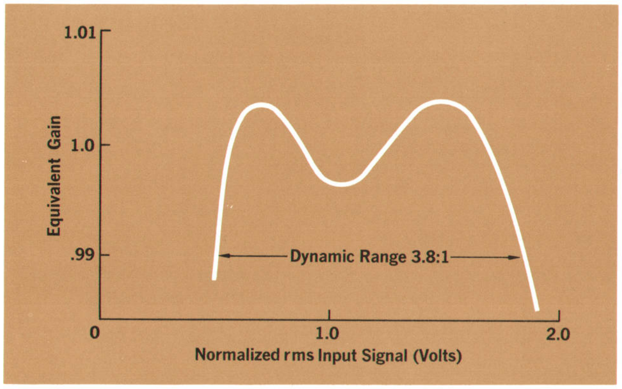

Michael A. Perry

Mike Perry studied electronic engineering at the University College of North Wales, obtaining the B.Sc. (Honours) degree in 1963 and the Ph.D. degree in 1966. His doctoral research was in electrometer measurement techniques, and a paper he co-authored on this research won an IEE Electronics Division premium for 1965. After joining HP in 1966, Mike worked on the 3722A Noise Generator and the 3721A Correlator. He is now involved in new product investigations.

George C. Anderson

George Anderson was project leader on the 3722A Noise Generator and on the 3721A Correlator. He has been with HP since 1966.

George graduated from Heriot-Watt University, Edinburgh, in 1954. After completing a two-year graduate course in electrical engineering, he did varied industrial work for a number of years. During the three years prior to his joining HP, he was with the Royal Observatory, Edinburgh, where he developed data recording systems for the seismology unit.

SPECIFICATIONS
HP Model 3721A
Correlator

INPUT CHARACTERISTICS

Two separate input channels, A and B, with identical amplifiers.

BANDWIDTH: The B input signal is sampled at a maximum rate of 1 MHz. According to the sampling theorem, this means that input signals must be *band-limited* to 500 kHz or less. A useful rule of thumb for random signals is to allow at least four samples per cycle of the upper 3-dB cutoff frequency; thus the maximum 3-dB cutoff frequency would be 250 kHz. Pure sine-wave inputs would probably require more samples per cycle to give a recognizable picture on the CRT; however, for computer processing of the data four samples per cycle would be entirely adequate. Model 3721A's lower cutoff frequency is selectable, dc or 1 Hz.

INPUT RANGE: 6 ranges, 0.1 V rms to 4 V rms, in 1, 2, 4, 10 sequence.

ANALOG-TO-DIGITAL CONVERSION: Fine quantizer; 7 bits. Coarse quantizer (feeds delayed channel): 3 bits. Coarse quantizer linearized by internally-generated wideband noise (dither).

OVERLOAD: Maximum permissible voltage at input: dc coupled 120 V peak, ac coupled 400 V = dc + peak ac.

INPUT IMPEDANCE: Nominally 30 pF to ground, shunted by 1 MΩ.

CORRELATION MODE

Simultaneous computation and display of 100 values of auto or cross-correlation function. Display sensitivity indicated directly in V^2/cm on illuminated panel. Non-destructive readout; computed function can be displayed for an unlimited period without deterioration. (Non-permanent storage; data cleared on switchoff.)

TIME SCALE: (Time/mm = delay increment $\triangle\tau$) 1 μs to 1 second (total delay span 100 μs to 100 seconds) in 1, 3.33, 10 sequence with internal clock. Other delay increments with external clock; minimum increment 1 μs (1 MHz), no upper limit.

DELAY OFFSET: Option series 01 provides delay offset (precomputation delay) facility. Without offset, first point on display represents zero delay; with offset, delay represented by first point is selectable from 100 $\triangle\tau$ to 900 $\triangle\tau$ in multiples of 100 $\triangle\tau$.

DISPLAY SENSITIVITY: 5 x $10^{-6}V^2/cm$ to 5 V^2/cm.
Calibration automatically displayed by illuminated panel.

VERTICAL RESOLUTION: Depends on display sensitivity. Minimum resolution is 25 levels/cm. Interpolation facility connects points on display.

AVERAGING: Two modes are provided: Summation (true averaging) and Exponential.

1. SUMMATION MODE
   Computation automatically stopped after N process cycles, at which time each point on the display represents the average of N products. N is selectable from 128 to 128 x 1024 (2^7 to 2^{17} in binary steps). Display calibration automatically normalized for all values of N. Summation time indicated by illuminated panel.
2. EXPONENTIAL MODE
   Digital equivalent of RC averaging, with time constant selectable from 36 ms to over 10^7 seconds. Approximate time constant indicated by illuminated panel. Display correctly calibrated at all times during the averaging process.

SIGNAL RECOVERY MODE (Channel B only)

Detects coherence in repeated events, when each event is marked by a synchronizing pulse. After each sync pulse, a series of 100 samples of channel B input is taken, and corresponding samples from each series are averaged. The 100 averaged samples are displayed simultaneously. Display sensitivity is indicated directly in V/cm on illuminated panel.

SYNCHRONIZATION: An averaging sweep is initiated either by a trigger pulse from an external source (EXT) or, in internally triggered mode (INT), by a pulse derived from the internal clock. In the INT mode, the start of each sweep is marked by an output pulse (stimulus) used to synchronize some external event.

TRIGGER INPUT: Averaging sweep initiated by negative-going step.

STIMULUS OUTPUT: Negative-going pulse at start of averaging sweep.

TIME SCALE: (Time/mm = interval between samples) 1 μs to 1 second (total display width 100 μs to 100 seconds) in 1, 3.33, 10 sequence with internal clock. Other intervals (hence other display widths) with external clock; minimum interval 1 μs (1 MHz), no upper limit.

DISPLAY SENSITIVITY: 50 μV/cm to 1 V/cm. Calibration automatically displayed by illuminated panel.

VERTICAL RESOLUTION: Depends on display sensitivity. Minimum resolution is 25 levels/cm. Interpolation facility connects points on display.

SIGNAL ENHANCEMENT: Improvement in signal-to-noise ratio equals square root of number of averaging sweeps.

NUMBER OF SWEEPS = N in summation mode; N x gain factor of 1, 10 or 100 in exponential mode.

PROBABILITY MODE (Channel A only)

Displays either (1) amplitude probability density function (pdf) or (2) integral of the pdf of channel A input. Signal amplitude represented by horizontal displacement on display, with zero volts at center; vertical displacement represents amplitude probability.

DISPLAY SENSITIVITY: Horizontal sensitivity 0.05 V/cm to 2 V/cm in 5, 10, 20 sequence.

HORIZONTAL RESOLUTION: 100 discrete levels in 10 cm wide display = 10 levels/cm.

VERTICAL RESOLUTION: 256 discrete levels in 8 cm high display = 32 levels/cm.

VERTICAL SCALING: Depends on averaging method used (summation or exponential).

SAMPLING RATE: 1 Hz to 3 kHz in 1, 3, 10 sequence with internal clock. Other sampling rates with external clock; maximum frequency 3 kHz, no lower frequency limit.

INTERFACING

X-Y RECORDER: Separate analog outputs corresponding to horizontal and vertical coordinates of the CRT display.

OSCILLOSCOPE: Separate analog outputs corresponding to the horizontal and vertical coordinates of the CRT display.

NOISE GENERATOR MODEL 3722A: Control of the Correlator from the Model 3722A Noise Generator. The gate signal from the 3722A is used to set the Correlator into RUN state; on termination of the gate signal, Correlator will go into HOLD state.

DIGITAL COMPUTER: Option 020 provides interface hardware (buffer card) for reading out displayed data to digital computer.

CLOCK

INTERNAL CLOCK: All timing signals derived from crystal-controlled oscillator: stability 40 ppm over specified ambient temperature range.

EXTERNAL CLOCK: Maximum frequency 1 MHz.

PROCESS CLOCK: 135 μs wide negative-going pulse. Normally +12 V, falls to 0 V at start of each process cycle and returns to +12 V after 135 μs.

REMOTE CONTROL & INDICATION

CONTROL: Remote control inputs for RUN, HOLD and RESET functions are connected to DATA INTERFACE socket on rear panel.

INDICATION: Remote indication of correlator RUN, HOLD or RESET states is available at the DATA INTERFACE socket on rear panel.

GENERAL

AMBIENT TEMPERATURE RANGE: 0° to +50°C.

POWER: 115 or 230 V ±10%, 50 to 1000 Hz, 150 W.

DIMENSIONS: 16¾ in. wide, 10¾ in. high, 18¾ in. deep overall (426 x 273 x 476 mm).

WEIGHT: 45 lb. (20.5 kg) net.

OPTIONS

DELAY OFFSET OPTION SERIES 01
- Option 011 Correlator with 100 Δt offset facility.
- Option 013 Correlator with 300 Δt offset facility.
- Option 015 Correlator with 500 Δt offset facility.
- Option 017 Correlator with 700 Δt offset facility.

The above options are extendable (factory conversion only) to 900 Δt offset, in multiples of 200 Δt.
- Option 019 Correlator with 900 Δt offset facility.

DATA INTERFACE: Option series 02
- Option 020 Correlator with interface for data output to computer.

PRICE: Model 3721A, $8350.00

MANUFACTURING DIVISION: HEWLETT-PACKARD LTD.
South Queensferry
West Lothian, Scotland

Correlation In Action

Selected Applications of Model 3721A Correlator

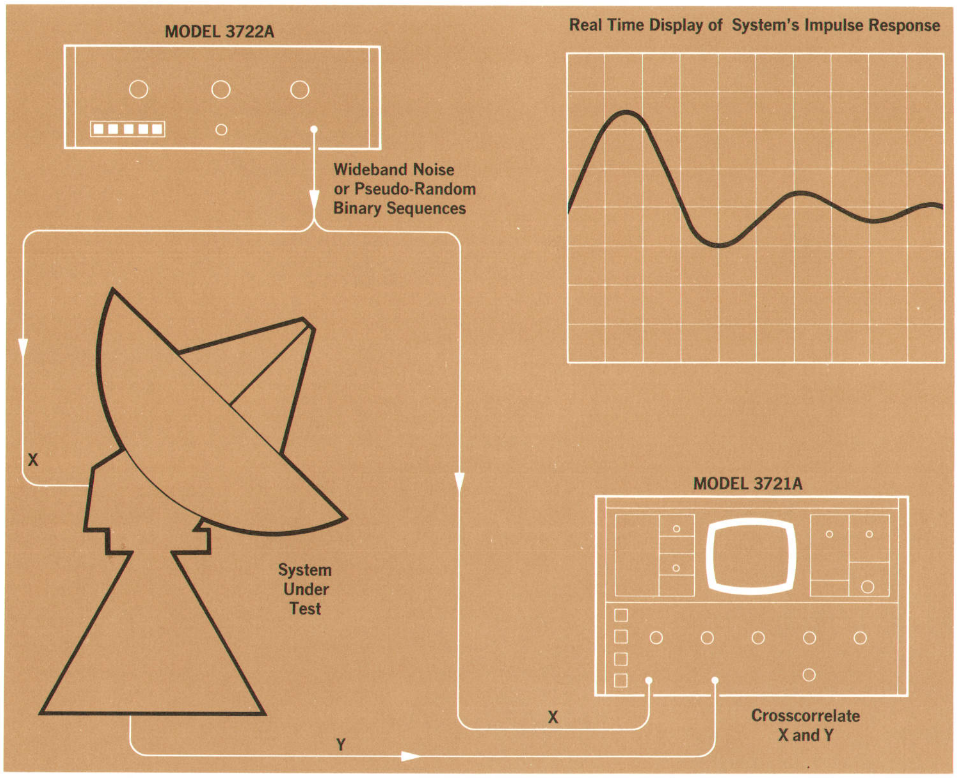

System Identification

Throughout engineering and science there is a requirement for system identification, or the determination of the laws which relate the outputs of a system to its inputs. The system in question can be anything from a simple electronic circuit with one input and one output to a large chemical process plant with many inputs and outputs. Previously used techniques of system identification — e.g., step, impulse, or sine-wave testing—all have limitations. Using a random or pseudo-random noise input and correlation techniques, the basic limitations can be overcome. The advantages of using correlation identification techniques are:

- The plant or system need not be closed down for testing.
- Test signals can be kept small, and they need not intefere with normal operation.
- Results can be obtained in the presence of random noise and parameter drifts.
- If a random or pseudo-random binary sequence is used as the input signal, it can be reproduced easily by a transducer (e.g., a solenoid valve cannot reproduce a wideband noise signal, but it can produce an on-off binary sequence).

A suitable noise source is the HP 3722A Noise Generator, which provides random or pseudo-random binary or Gaussian signals.

The technique is illustrated in the diagram. The noise signal is applied to the system input and the input and output are crosscorrelated. The result is a plot of the system's impulse response. Optionally, the Model 3721A Correlator can be interfaced with a computer for further analysis of the data; for example, computing the Fourier transform of the impulse response gives the frequency and phase responses of the system.

System identification by correlation can be used to determine the transfer functions of 'black boxes' so their performance in control systems can be predicted, or to determine process plant characteristics on line, so control can be optimized. In the latter case, the state of the art has not advanced very rapidly because no suitable on-line correlator was available. Model 3721A should change this. It is possible that the new correlator could be used as part of an adaptive control loop where, along with a digital computer, it would continuously determine the system's transfer functions and control the system to optimize its performance.

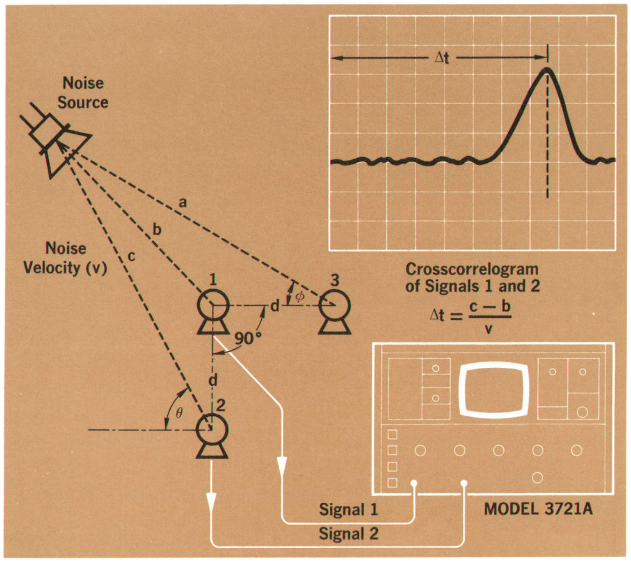

Noise Source Direction Finding

The location of a noise source can be determined by crosscorrelating the outputs of two closely-spaced detectors. The diagram depicts the technique applied to a two-dimensional problem, where the noise source lies in the same horizontal plane as the detectors. Initially the detectors are placed in positions 1 and 2 to determine the value of angle θ. Time delay Δt between the signal arriving at positions 1 and 2 will be displayed as a peak on the cross correlogram of the two detector output signals. For small distance d,

$$\sin\theta = \frac{c-b}{d} = \frac{vt}{d}.$$

Hence θ can be calculated assuming that v is known. By relocating detector 2 at position 3 and once again cross-correlating the two detector output signals, the value of ϕ can be calculated from

$$\cos\phi = \frac{d}{vt}.$$

From the values of θ and ϕ thus obtained, the location of the noise source can be determined. For a three-dimensional problem where the noise source does not lie in the same horizontal plane as detector positions 1, 2 and 3, at least one further measurement must be made with a detector relocated in the plane normal to that of positions 1, 2 and 3.

This technique is applicable to underground and underwater direction finding as well as to other types.

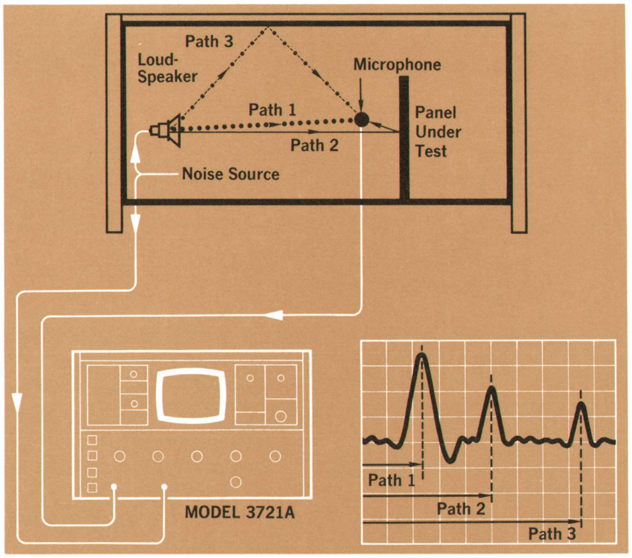

In-Situ Measurements of Acoustic Absorption Coefficients

When designing special acoustic environments such as broadcasting studios or concert halls, the acoustics engineer often finds it necessary to determine the sound absorbing powers or acoustic absorption coefficients of the materials (e.g., walls, ceilings, furniture, etc.) which make up the buildings. The conventional measurement technique requires that the material or item under test be placed in a special chamber, but this is sometimes an expensive or even impossible task in the case, say, of a wall! Further, results obtained by this method are often inconsistent. Correlation can be used to provide rapid, consistent measurements for the values of absorption coefficients of materials in their normal environment. In the *in-situ* test shown, sound from the loudspeaker can reach the microphone via many paths of different lengths, and the sound takes different times to traverse these paths. If the input to the loudspeaker is cross-correlated with the microphone output, the first correlation peak on the 3721A display indicates the shortest sound path (path 1), the second peak indicates the next shortest path (path 2), and so on. For measurement of the acoustic absorption coefficient of the panel under test, only the relative amplitudes of the sounds from paths 1 and 2 are of interest. The amplitudes can be read directly from the first two peaks of the correlogram and from this information the acoustic absorption coefficient of the panel can be determined rapidly.

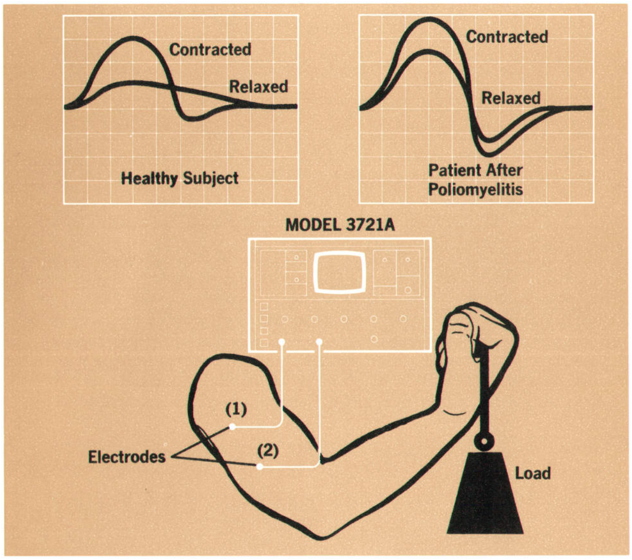

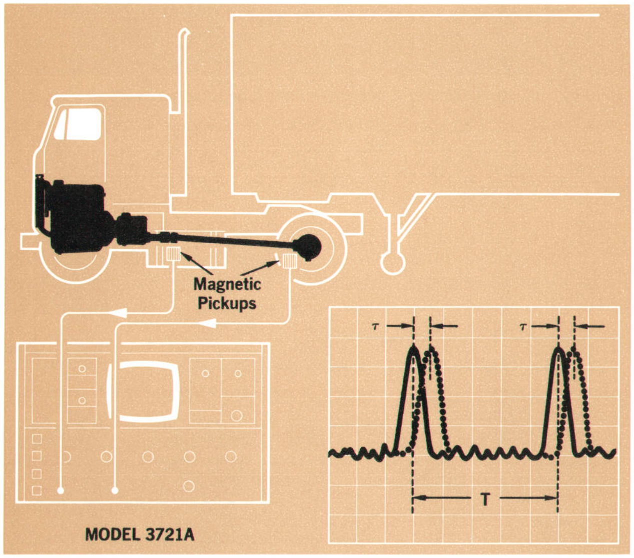

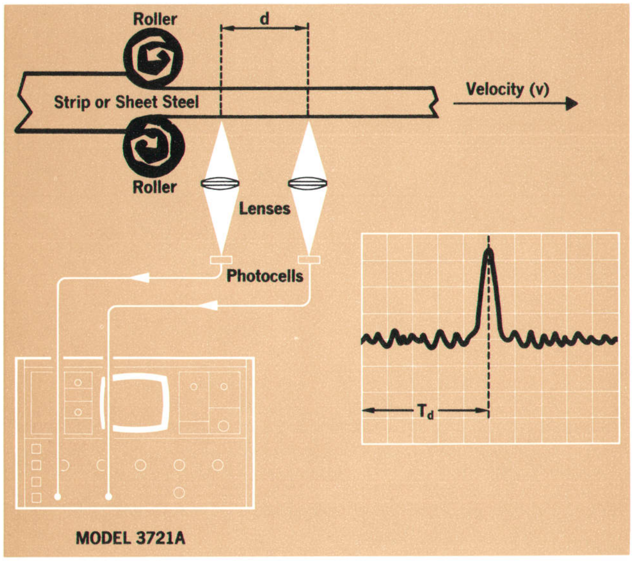

Crosscorrelation of Electromyograms

The electromyogram (EMG) is a record of muscular electrical activity, which can assist in the diagnosis of nervous and muscular diseases. When detected by a needle electrode, the EMG consists of a series of pulses representing the electrical activity of the muscle cells in the vicinity of the electrode. A useful measure of fatigue or disease in a muscle is the extent to which the cells fire independently of each other. The crosscorrelogram of the two EMG's from a muscle will show whether the cells are tending to fire together or randomly. The diagram contrasts the crosscorrelogram obtained from a healthy subject with that acquired from a patient suffering from the aftereffects of poliomyelitis. With the healthy subject, the crosscorrelogram of signals 1 and 2 shows independent firing of the cells when the muscle is relaxed, with correlation increasing as the muscle contracts on load. Where the patient is suffering from the aftereffects of poliomyelitis, this correlation is prominent even in the relaxed state of the muscle, and it increases rapidly on load as the muscle becomes tired. The crosscorrelogram displayed by the Model 3721A is, therefore, a useful guide to muscle condition.

Measurement of Torsion in Rotating Shafts

Correlation techniques can be used to measure torsion in a power transmission shaft under operating conditions, by accurately detecting a change in phase angle between the output signals of transducers placed at either end of the shaft. In the example illustrated, the twist in a long transmission shaft is being measured. Magnetic pickups are placed adjacent to the universal joints at the ends of the shaft so that when the shaft rotates each pickup produces a train of regular pulses with period proportional to the shaft speed. Crosscorrelation of the pickup output signals yields a correlogram whose peaks are spaced T (seconds) apart, T being the pulse repetition period. If a load (torque) change occurs, the phase angle between the output signals of the pickups will change resulting in an overall movement τ of the displayed peaks. This change τ will give a measure of the shaft twist and hence of the torque it is transmitting. This technique can be applied to any situation in which large powers are being transmitted along rotating shafts.

Contactless Velocity Measurement

Measurement of the velocity of steel strip or sheet from a rolling mill is a difficult problem when the metal is cold, but when the metal is white hot, the difficulty is increased greatly. Contactless measurement of the velocity is possible, however, using a crosscorrelation technique. The technique is illustrated in the diagram. When metal is rolled, its surface is not perfectly smooth and any irregularities will affect the output of a photocell which is focused on the surface. After a finite time, each irregularity will pass the focusing point of a second similar photocell placed downstream from the first one. Crosscorrelation of the two photocell output signals using the 3721A Correlator will indicate the time delay T_d directly. If the photocell separation d is known, the velocity v of the strip or sheet leaving the rollers can be determined simply from:

$$v = \frac{d}{T_d}.$$

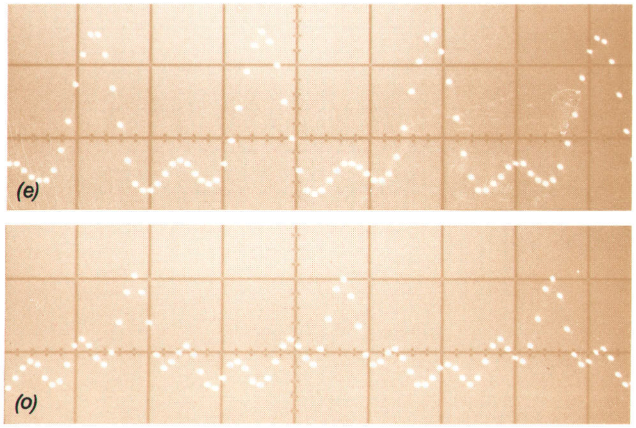

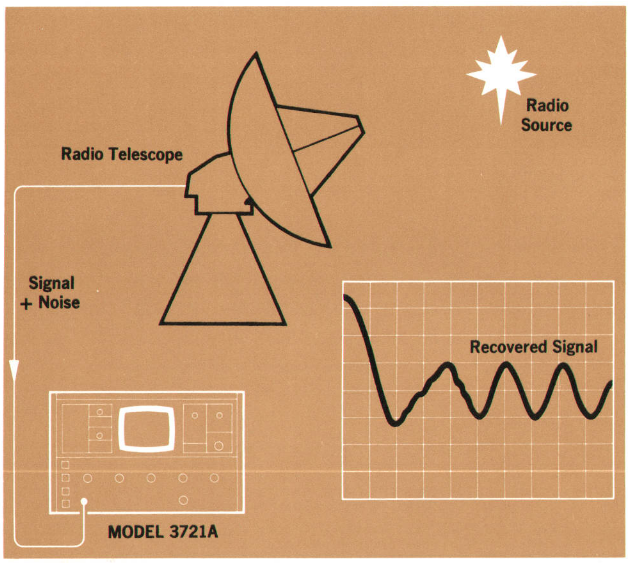

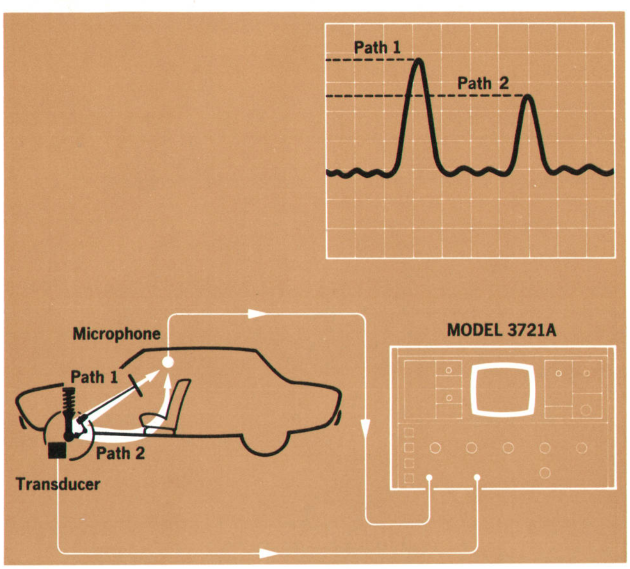

Speech Research

Among many topics of current interest in the audio field which demand more precise basic knowledge of human speech are the reduction of speech bandwidth for telecommunication purposes, the education of the deaf, and the provision of man-computer interfaces which speak human languages rather than machine languages. Although spectral analysis is a well-established tool for characterizing audio signals, correlation techniques, which present the same information in the time domain, are often more useful where speech is concerned. Using the 3721A Correlator, the autocorrelation of speech signals can be performed very simply. The two photographs show the autocorrelation functions of two common speech sounds: 'e' as in 'tree,' and 'o' as in 'mole.' These functions were obtained within seconds from direct speech into a microphone connected to the Model 3721A.

Radio Astronomy

The detection of very weak and very distant radio sources in space is carried out by high-gain radio telescopes. These telescopes are extremely sensitive to interference from random electrical disturbances, which tend to mask the small signals of interest. However, in some cases — e.g., pulsars — the signals of interest are periodic, and by using the Model 3721A and autocorrelation techniques, the periodic component of the received signal can be extracted from the noise. The method is based on the principle that, after an appreciable delay, the autocorrelation function of the noise component will have died away to zero, but the periodic component will have a periodic autocorrelation function which will persist at all delays. A simple detection system is illustrated in the diagram. There are no theoretical limits to the autocorrelation method. Any periodic signal can be detected in the presence of noise. Besides radio astronomy, this technique can be used in many other branches of research—medical, acoustical, vibration and communications.

Determination of Noise Transmission Paths

Unwanted noise and vibrations often travel to the ear of the hearer via several different paths. If these paths can be determined, steps can be taken to minimize or eliminate their effects. Correlation techniques can be used to determine noise transmission paths. For example, in the automobile industry it may be desirable to determine the noise transmission paths from the front axle of a car to the driver. This can be done by placing a transducer on the axle and a microphone inside the car in the vicinity of the driver's head, and crosscorrelating the outputs of the transducer and the microphone. In the example illustrated, the noise from the front axle is transmitted primarily through the steering column (path 1) and via the floor and driver's seat (path 2). The correlogram displays two peaks with time delays representative of the individual transmission path delays. Examination of the amplitude of the peaks will give a measure of the relative amounts of noise transmitted by each path.

HEWLETT-PACKARD JOURNAL **NOVEMBER 1969** *Volume 21 • Number 3*

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