

HEWLETT-PACKARD JOURNAL



APRIL 1968

What Is Signal Averaging?

Repetitive waveforms buried in noise can often be pulled out by a signal averager, an instrument that takes advantage of the redundant information provided by repetition.

By Charles R. Trimble

AS MAN CONTINUES TO EXPLORE THE WORLD AROUND HIM, he frequently finds it necessary to develop new methods and more sophisticated tools to help him uncover the hidden 'why's' of nature. Noise, for example, is a perpetual problem. As he investigates phenomena characterized by low-level signal amplitudes, the scientist finds himself hindered—and sometimes stymied—by uninteresting, unwanted random disturbances. The search for new and better ways to deal with this noise seems to have no end.

Cover: *The dots represent sample values of human brain waves, displayed on the CRT of the new HP Model 5480A Signal Averager. The apparently random white dots represent the brain wave observed after a single stimulus (a light flash). The black dots represent the average of 64 responses. See 'Where Averaging Helps,' page 9.*

In this Issue: *What is Signal Averaging?; page 2. Calibrated Real-Time Signal Averaging; page 8. Where Averaging Helps; page 9. Off-Line Analysis of Averaged Data; page 14.*

One common method of attenuating noise is to limit the bandwidth of the monitoring instruments by conventional filtering. If the instrument bandwidth is wider than the desired signal's bandwidth, the extra bandwidth only admits extra noise and can be discarded without losing any of the signal. Useful as filtering is, it is of little value when the signal and noise occupy the same part of the frequency spectrum.

If a waveform is repetitive, signal-to-noise ratio can sometimes be improved by making use of the redundant information inherent in repetition. Multiple-exposure photography has been used for years to enhance the signal-to-noise ratios of waveforms displayed on oscilloscopes. More recently, variable-persistence oscilloscopes have been able to do the same job without the camera.

Even more convenient and more precise enhancement of noisy repetitive signals can be obtained through the use of special-purpose digital instruments. Generally classified as signal averagers, these instruments have opened many new areas to scientific investigation. Operating in real time, 'on line,' they allow the researcher to see experimental results while the experiment is in progress, even though the signal may be so obscured by noise that the raw data seem to contain little or no useful information. Averagers demand only that some means be available for synchronizing them to the repetitive signal that is to be pulled out of the noise.

In the description of signal averaging which follows, it will be assumed that the signal is a repetitive voltage waveform. It's important to remember, though, that the physical phenomenon being observed can be anything that can be expressed as a voltage waveform. Signal averaging can be applied to frequency spectra, to probability distributions, or to biomedical, chemical, nuclear, spectroscopic, or mechanical phenomena — to anything that can be translated by transducers or other means to repetitive electrical waveforms.

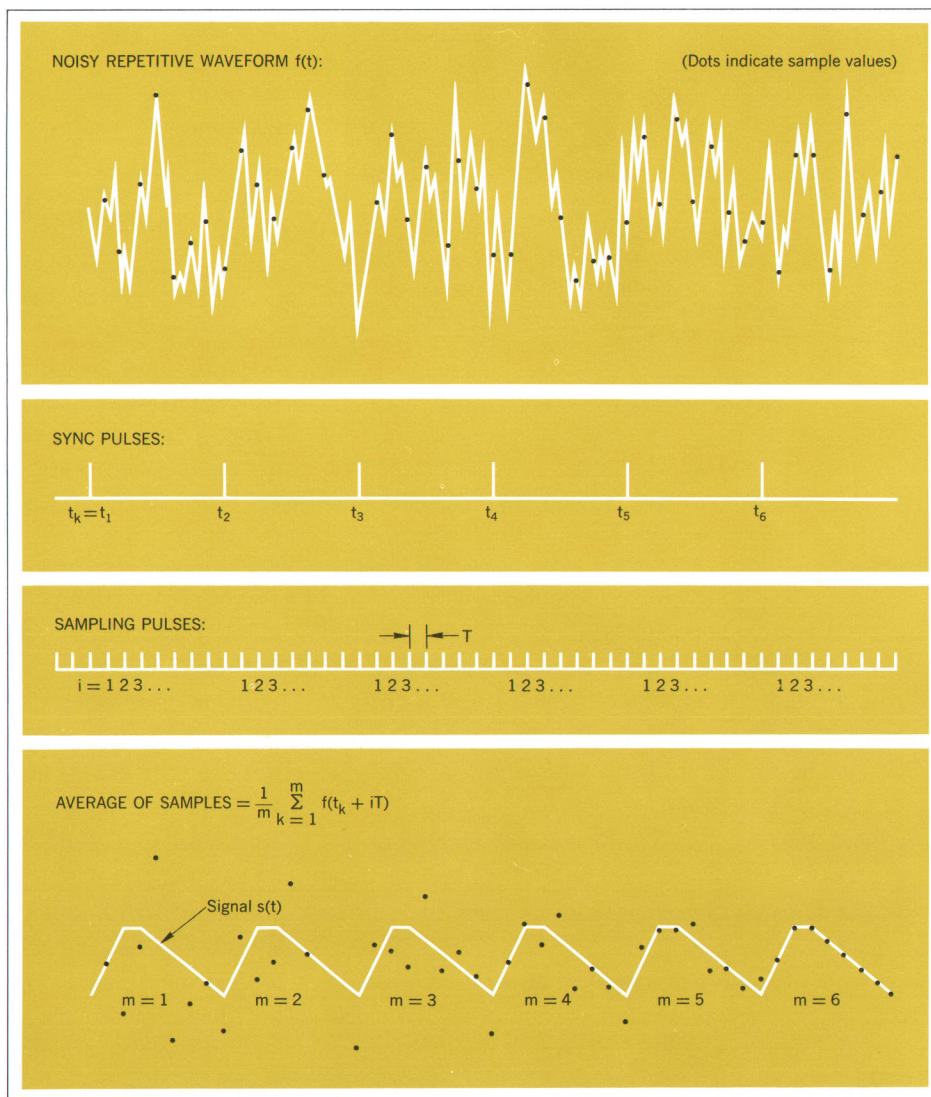
Averaging in the Time Domain

Signal averagers sample input signals at fixed time intervals (see Fig. 1), convert the samples to digital form, and store the sample values at separate locations in a memory. The sampling theorem tells us that *no information is lost by this discrete representation*, provided that

the sampling rate is at least twice as fast as the highest frequency present in the input signal. Notice that when the signal is band-limited it is not necessary to take the average value of the signal during the period between samples; it is only necessary to sample at a sufficiently high rate.

The sampling process is continued for a preset number of repetitions of the desired signal. During the first repetition, sample values are stored in memory, with each memory location corresponding to a definite sample time. Then, during subsequent repetitions, the new sample values are added algebraically to the values accumulated at the corresponding memory locations. After any given number of repetitions, the sum stored in each memory location is equal to the number of repetitions times the average of the samples taken at that point on the desired waveform.

Fig. 1. Signal averaging can recover repetitive waveforms from noise even when the raw data (top) seem to contain little or no useful information. A sync signal tells the signal averager where to find the start of each repetition. The averager then samples the input $f(t)$ every T seconds and stores the sample values at separate locations in a memory. It then computes the average sample value for each m and i (bottom). The noise portion of the average gradually dies out, since the noise makes both positive and negative contributions to each sample point on successive repetitions. For random noise, the signal-to-noise voltage ratio improves as \sqrt{m} , where m is the number of repetitions.






Fig. 2. Top: Noisy triangular wave viewed on conventional oscilloscope. Vertical Scale: 0.5 V/cm. Bottom: Same wave after averaging. Vertical Scale: 0.05 V/cm. Number of repetitions averaged: 4096. Improvement in signal-to-noise ratio: 36 dB.

To tell the averager where the beginning of each signal repetition is, a synchronizing signal must be available. The waveform of interest needn't be periodic, but it must repeat exactly following each sync pulse.

This simple summation process tends to enhance the signal with respect to the noise. The signal portion of the input is a constant for any sample point, so its contribution to the stored sum is multiplied by the number of repetitions. On the other hand, the noise—which is random and not time-locked to the signal—makes both positive and negative contributions at any sample point during successive repetitions. Therefore the noise portion of the stored sum grows more slowly than the signal portion.

More formally, the averaging or summation process can be described as follows. Let the input be $f(t)$, composed of a repetitive signal portion $s(t)$ and a noise portion $n(t)$. Say the k^{th} repetition of $s(t)$ begins at time

t_k (and let $t_1 = 0$). Finally, let samples be taken every T seconds. We have then

$$f(t) = s(t) + n(t).$$

This signal is sampled, and the sample values are

$$\begin{aligned} f(t_k + iT) &= s(t_k + iT) + n(t_k + iT) \\ &= s(iT) + n(t_k + iT) \end{aligned}$$

For a given i and k , $n(t_k + iT)$ is a random variable. It's reasonable to assume that in a real situation, where the noise is thermal noise, shot noise, $1/f$ noise, or the like, all the $n(t_k + iT)$ have a mean value of zero and the same rms value, say σ . And for different k 's, the noise samples are usually statistically independent.

Now consider the i^{th} sample point. A measure of the noise masking the signal is the signal-to-noise voltage ratio, S/N . On any particular repetition,

$$S/N = \frac{s(iT)}{\sigma}.$$

After m repetitions, the value stored at the i^{th} memory location is

$$\begin{aligned} \sum_{k=1}^m f(t_k + iT) &= \sum_{k=1}^m s(iT) + \sum_{k=1}^m n(t_k + iT) \\ &= m s(iT) + \sum_{k=1}^m n(t_k + iT). \end{aligned}$$

Since the noise is random and the m samples are independent, the mean square value of the sum of the m noise samples is $m\sigma^2$, and the rms value is $\sqrt{m}\sigma$. Therefore the signal-to-noise ratio after summation is

$$(S/N)_m = \frac{ms(iT)}{\sqrt{m}\sigma} = \sqrt{m}(S/N).$$

Thus summing m repetitions improves the signal-to-noise ratio by a factor of \sqrt{m} . If m is 2^{19} , the improvement in S/N is 57 dB. Fig. 2 shows how an averager can find a signal in noise when there seems to be no signal there.

Often the noise masking the signal isn't random at all. For example, 60-Hz hum can effectively obscure a signal. In such cases, the original informal argument still applies. At any sample point, a signal which isn't time-locked to the desired signal will eventually contribute both positive and negative values to the sum, and will therefore contribute less than the desired signal. However, it's difficult to determine exactly how much the undesired signal will be attenuated. To get an answer to this question, we can look at signal averaging as a filtering process and attack the problem in the frequency domain.

Averaging in the Frequency Domain

In the digital implementation of signal averaging, sampling is required. However, if the conditions of the sampling theorem are met, no new information is gained by considering sampling as part of the averaging problem. For simplicity, therefore, sampling and averaging will be treated separately in the discussion that follows.

Summing m repetitions of the signal $s(t)$ is equivalent, mathematically, to convolving the input $f(t) = s(t) + n(t)$ with a train of m unit impulses, that is, with the sync-pulse train. A short calculation will demonstrate this. For simplicity, assume the signal $s(t)$ is periodic, and let its period be τ . Convolution of the input $f(t)$ with m unit impulses spaced τ seconds apart gives

$$\begin{aligned} a(t) &= \int_{-\infty}^{\infty} f(t - \zeta) \sum_{k=1}^m \delta(\zeta - k\tau) d\zeta \\ &= \sum_{k=1}^m f(t - k\tau), \end{aligned}$$

i.e., the sum of m periods of $s(t)$.

Since the averager effectively convolves the input with m sync pulses, the effective impulse response of the averager is a train of m impulses, i.e.,

$$h(t) = \sum_{k=1}^m \delta(t - k\tau).$$

Translating this to the frequency domain, we get the transfer function of the averager, $H(j\omega)$.

$$|H(j\omega)| = \left| \frac{\sin\left(\frac{\omega m \tau}{2}\right)}{\sin\left(\frac{\omega \tau}{2}\right)} \right|$$

Notice that $|H(j\omega)| = m$ whenever $\omega\tau$ is an integral multiple of 2π .

It's a Comb Filter

The averager's transfer function $|H(j\omega)|$ is illustrated in Fig. 3 for various values of m . In each case the function was calculated by a computer and displayed on a CRT. The vertical scales are all normalized, that is, the transfer function has been divided by m .

As Fig. 3 shows, the averager is, in effect, a comb filter with all 'teeth' having the same height and 3-dB points. The bandwidth of each tooth gets progressively narrower as the number of repetitions increases. It's important to notice that, since the sync pulses are time-locked to the signal $s(t)$, every frequency component of $s(t)$ coincides with the center of one of the teeth of the comb filter.

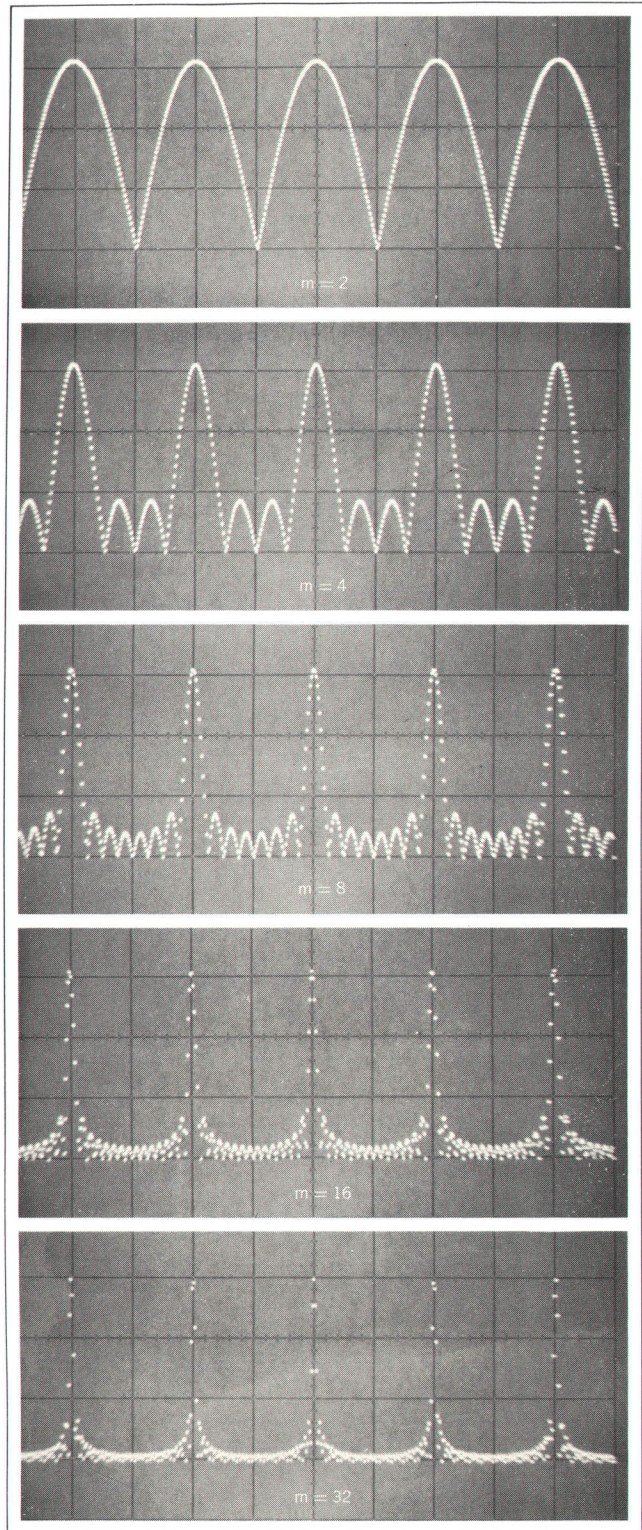


Fig. 3. Effective transfer function of signal averager, calculated by digital computer. Averager acts like comb filter for repetitive waveforms. Since averager is time-locked to repetitive waveform, frequency components of waveform coincide exactly with centers of comb-filter teeth. Width of teeth is inversely proportional to number of repetitions averaged.

Therefore, as the number of repetitions increases and the filter becomes more and more selective, the filter increasingly favors the signal with respect to the noise.

For large m , the 3-dB width of one tooth of the comb filter is approximately

$$\Delta f = \frac{0.886}{m\tau} \text{ Hz.}$$

Fig. 4 shows how an averager can effectively separate a 200.0-Hz rectangular wave from an interfering 200.1-Hz sine wave. Here m was 2^{14} , $1/\tau$ was 50 Hz, and $\Delta f = 0.0027$ Hz.

It is the averager's comb-filter behavior that makes this type of instrument so eminently suitable when the desired signal and the noise are in the same frequency range. The noise power remaining in the averaged waveform is the sum of the powers contained in each comb tooth. Since the tooth width is inversely proportional to m (the number of repetitions) the total noise bandwidth is also inversely proportional to m . Here, then, is another

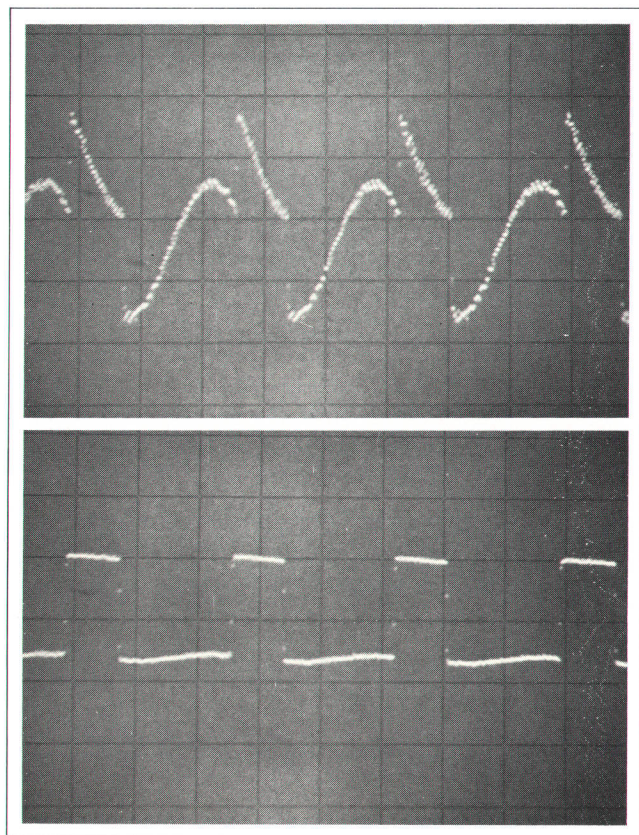


Fig. 4. Top: 200.0-Hz rectangular wave plus 200.1-Hz sine wave, before averaging. Bottom: 200.0-Hz rectangular wave recovered by signal averager after 2^{14} repetitions. In this case, averager had the effect of a comb filter whose teeth had 3-dB bandwidths of 0.0027 Hz.

demonstration that the rms noise contracts by a factor of $1/\sqrt{m}$ with respect to the signal.

Instrument noise limits the maximum output signal-to-noise ratio that an averager can produce to about 57 dB. Subject to this limitation, an averager can produce improvements in S/N of as much as 60 dB (for $m = 2^{20}$). At 100 repetitions per second, this takes 174 minutes.

Some Cautions About Sampling

In the time domain, sampling is essentially the process of multiplying a continuous waveform $f(t)$ by a train of impulses spaced at fixed equal intervals, say T . The Fourier transform of a sampled signal is the convolution of $F(j\omega)$ —the transform of $f(t)$ —with a train of impulses spaced at intervals of $1/T$ (see Fig. 5). If $1/T$ is at least twice the highest-frequency component in $f(t)$, the sampling process duplicates the spectrum $F(j\omega)$ around each harmonic of the sampling frequency $1/T$, and no information about $f(t)$ is lost.

In the discussion so far, it has been assumed that the sampling rate of the averager is at least twice the highest frequency present in the input. If the signal we are looking for meets this requirement, the averager will not distort the signal even if the noise does not meet the requirement. However, noise above half the sampling frequency will be folded back onto the low-frequency noise (the duplicate spectra in Fig. 5 will overlap). This means that the input signal-to-noise ratio will be smaller, and so the averaging time for a given output S/N will be longer. Therefore, conventional high-frequency noise filtering ahead of the averager will increase the efficiency of the system.

How to Get the Most Out of an Averager

Here are three points to watch when using a signal averager.

- Conventional bandpass filtering to eliminate noise outside the signal passband will improve the input S/N and reduce the total averaging time. Filtering can be helpful even when there is no danger of the high-frequency noise spectrum's being folded back onto the low-frequency spectrum during the sampling process, as described in the preceding section. If we can reduce the input noise by 10 dB by restricting the bandwidth, the averaging time for a given output S/N decreases by a factor of 10.
- One of the consequences of the ever narrowing comb tooth is that the synchronizing or stimulating signal must be solidly time-locked to the waveform of interest. Jitter of the synchronizing signal will appear as

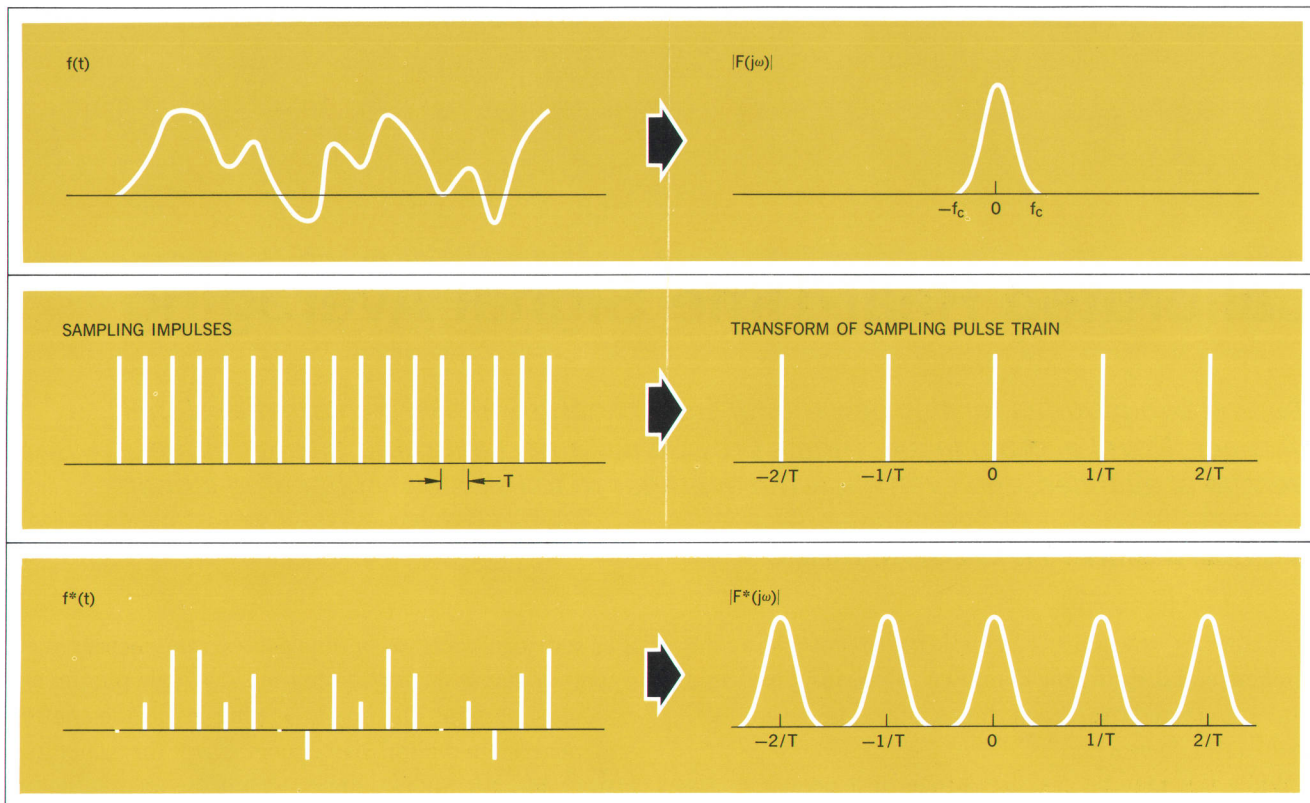



Fig. 5. Band-limited signal (top) sampled by infinite impulse train (middle) has frequency spectrum consisting of original spectrum duplicated at multiples of sampling frequency $1/T$. No information about signal is lost as long as duplicate spectra do not overlap, i.e., as long as $1/T > 2f_c$, where f_c is highest frequency present in signal $f(t)$.

smearing of the averaged waveform. In many cases, this will become more pronounced as the averaging proceeds. The same applies to the waveform to be averaged. Different parts of the waveform should not jitter with respect to each other. If they do, the average will actually deteriorate as the number of sweeps is increased. Hence there are cases where excessive averaging will deteriorate the signal rather than improve it. In some cases, like evoked biological responses, truly stationary signals are not available and the average is of questionable value if we average too long. (On the other hand, evoked responses are often impossible to measure without signal averaging. See 'Where Averaging Helps,' page 9.)

- If the signal to be averaged is really quasi-repetitive, that is, if it is the response to an operator-controlled stimulus, then the stimulus frequency should be chosen to minimize 60 Hz, its harmonics, or any other troublesome frequency. Often this can be accomplished by having the stimulus frequency change in a pseudo-random fashion.

Variations on the Basic Process

Throughout this article, the terms 'averaging' and 'summation' have been used interchangeably. While this is mathematically valid, there is a distinction from the user's point of view. The average is the sum divided by the number of repetitions. It is difficult to implement a running average, because it requires a fast division by a different number during each repetition. Therefore, most averagers only sum, which means they display a signal which grows with time until the preset number of repetitions has occurred. The new HP signal averager described in the following article does perform a running normalization, and therefore gives a stable, calibrated display of the average signal at all times.

Another variation on the basic process, also used in the new HP averager, is a method of averaging which, after a preset number of repetitions, keeps the width of the comb-filter's teeth constant. (Recall that in the basic process this width is inversely proportional to the number of repetitions.) This variation allows the averager to follow slowly-changing signals. 

Calibrated Real-Time Signal Averaging

The first two plug-ins for this new digital signal analyzer make it a versatile signal averager. Novel averaging algorithms provide a stable, calibrated display of the average at all times, and even allow the averager to follow slowly changing signals.

By J. Evan Deardorff and Charles R. Trimble

AVERAGING N NUMBERS is a simple process of adding the numbers and dividing the sum by N. Repeating the same process over and over, each time with one more number, is still a simple process.

While this may be easy for a man, it is not very easy for a machine. Machines that can perform a series of fast divisions — each time by a different integer — are complex and expensive.

Traditionally, signal averagers have left the division to the operator. The machine samples and sums, as the preceding article explains. After the preset number of signal repetitions have been summed, the operator adjusts the controls to get an on-screen display, then reads his data from the CRT and the control settings. During the averaging process, measurements are impossible because the displayed signal grows with each repetition, often going off the screen entirely.

Using a new algorithm which we call 'stable averaging,' we have designed a signal averager (see Fig. 1) which relieves the operator of the need to adjust controls after an experiment has started. Division is internal and automatic. Thus the instrument gives a stable, calibrated display of the average signal at all times during the averaging process. The only change that takes place on the CRT during averaging is the gradual attenuation of noise.

Another new algorithm provides signal-to-noise-ratio improvements for slowly changing signals. Exponentially weighted averaging gradually de-emphasizes old information with respect to new information to produce a continuous running average of the input waveform.

Actually, the new instrument is more than an averager, it is a plug-in digital signal analyzer, an instrument that

uses statistical tools for on-line analysis and measurement of input data. It takes two plug-ins — a logic plug-in and an analog plug-in. The first logic plug-in has the controls and programming for signal averaging, for generating histograms, for multi-channel scaling (sequential counting), and for conventional summing. The first analog plug-in is designed to optimize the averaging function, but it can also be used for the other functions.

Although it processes digitally, the new instrument is an analog-in/analog-out machine. This, along with the fact that the information of interest doesn't change during




Fig. 1. New HP Model 5480A Signal Analyzer is an instrument that applies statistical principles to the on-line, real-time analysis of input data. The first two plug-ins make it a signal averager. Features include a display which doesn't flicker or grow off-scale during averaging, 100 kHz maximum sampling rate, memory of 1024 24-bit words, and ability to enhance slowly varying waveforms.

averaging, means that the experimenter is always closely coupled visually with the noisy experimental situation. He can see what's happening all the time.

Flexibility for additional signal processing or storage is available, too. An I/O coupler makes the averager compatible with a computer and with many kinds of peripheral equipment, such as tape readers and teleprinters. The article on page 14 describes the many intriguing possibilities of such a system.

Horizontal Sweeps are Calibrated

The signal averager can be thought of as an oscilloscope for noisy waveforms. Horizontal sweeps are calibrated, and range from 1 ms/cm to 50 s/cm.

Memory capacity is 1024 24-bit words. Of these, 1000 words are used for data storage, so the displayed waveform can be represented by up to 100 points per horizontal centimeter. This much resolution is often unnecessary; hence the memory can be divided into halves or quarters. When this is done, a full 10-cm-wide display

with the same sweep speed is presented, but the number of points per centimeter is reduced to 50 or 25.

The rate at which the averager samples the input signal is determined by the sweep rate and the number of points per centimeter. The maximum sampling rate is 100 kHz.

Display is Flicker Free

An annoying aspect of studying low-repetition-rate signals is having to look at a flickering CRT display. If the repetition rate is very low, say 1 Hz, the display may be only a dot moving across the screen.

This problem comes from displaying the summed waveform synchronously with the sampling of the input, a procedure that is completely unnecessary when the signal is being processed digitally. In the new signal averager, flickering is eliminated by treating the input of raw data separately from the display of processed information. Input and processing are handled on an interrupt basis. Thus, except for brief interruptions that are invisible to the operator, the averager is always cycling


Where Averaging Helps

Signal averaging has the extraordinary ability to extract repetitive signals from noise of approximately the same frequency content. It is a powerful aid in a variety of disciplines.

In high-resolution spectroscopy (e.g., microwave, NMR, and mass spectroscopy) averaging can help overcome stability problems. In NMR studies, a signal averager like the new HP Model 5480A can improve resolution by an order of magnitude. An averager used in conjunction with a frequency synthesizer can improve the resolution of the system by almost another order of magnitude.

In the biological sciences signal averaging finds numerous applications. One such application is a recently developed technique for detecting heart defects. Two electrocardiograms (ECG) are taken, one while the patient is at rest and another while he is walking on a treadmill or doing some other kind of exercise. Normally, an ECG waveform is quite clean. However, when the patient is exercising, muscle voltages obscure much of the useful information, and signal averaging becomes necessary. The HP 5480A Signal Averager should be especially helpful in this application because it can enhance slowly varying signals as well as stationary ones.

Brain-wave responses to stimuli such as sound, light, or touch would be virtually impossible to extract from electroencephalograms without signal averaging. The two traces shown here illustrate how the HP Model 5480A can enhance evoked responses. In this experiment the stimulating instrument, a flashing light, was triggered by the averager's sync output. The response was processed by the averager after each stimulus. Then, before the next sync pulse, the aver-



Human brain-wave response to visual stimulus (flashing light). Top trace: Signal-averager display after one repetition seems to contain no useful information. Bottom trace: Averager display after 64 repetitions. Signal-to-noise ratio improvement was 18 dB.

ager paused for a 'post-analysis delay' to allow the subject (who happened to be author Trimble) to calm down. There was also a pre-analysis delay between each sync pulse and the beginning of processing to allow for the delay inherent in the brain's response.

In vibration analysis, averaging in conjunction with a pseudo-random noise generator (e.g., the HP Model 3722A) can help diagnose mechanical faults.

Other candidates for averager assistance are seismology, fluorescent-decay studies, and numerous electronic laboratory studies. Whenever there is a repetitive signal and a synchronizing signal, an averager can improve signal-to-noise ratio by as much as 60 dB.




Fig. 2. The 'stable averaging' algorithm used in the new signal averager takes a few more repetitions than summation to achieve the same noise attenuation. The big advantage of stable averaging is that the display remains calibrated at all times instead of growing off-scale. The new averager also has a summation mode.

through memory, displaying the latest information about the signal.

The Noise Can be Monitored

Sometimes it is helpful to be able to see how the input deviates from the average. Therefore, the new averager can be set to display the difference between the input and the stored average. The difference signal is also brought out to a rear-panel connector.

Difference monitoring is equivalent to turning the comb-filter* associated with the averaging upside down, so that it rejects a portion of the signal. This technique might be used to get a statistical analysis of the noise or, if the signal of interest is riding on a large periodic signal (e.g., 60-Hz hum), difference monitoring can be used to see what the interfering signal looks like.

The ability to monitor the noise is a direct consequence of the stable averaging technique used in the new averager. It could not be done if only the sum were stored.

Stable Averaging Gives Stable Display

As the article on page 2 explains, the averager samples the input signal $f(t)$ every T seconds, and accumulates in its memory the sample values taken on successive repetitions of the input signal. If the sync pulse marking the beginning of the k^{th} repetition of the signal occurs at time t_k , then the average stored in the i^{th} memory location after m repetitions should be

* See article, page 2.

$$M_m^i = \frac{1}{m} \sum_{k=1}^m f(t_k + iT) = M_{m-1}^i + \frac{f(t_m + iT) - M_{m-1}^i}{m} \quad (1)$$

This turns out to be a difficult algorithm to implement because of the large amount of hardware needed for a fast division by m . Even if it were implemented, round-off errors could build up to be a significant problem.

In stable averaging, we approximate equation (1) by

$$M_m^i = M_{m-1}^i + \frac{f(t_m + iT) - M_{m-1}^i}{2^N} \quad (2)$$

where $2^{N-1} < m < 2^N + 1$. That is, after the first repetition we divide by 1, after the second by 2, after the third by 4 (instead of 3), after the fourth by 4, after the fifth by 8 (instead of 5), and so on. This is easy to implement because dividing a binary number by 2^N is just shifting its binary point N places to the left. Roundoff errors are also eliminated by this algorithm.

Equations (1) and (2) both give the same value for the signal portion of the average; the second term in each expression is zero when there is no noise. The only difference is in the averaging of the noise.

The important question is, how efficient is equation (2) in getting rid of noise? Figure 2 gives the theoretical improvements in S/N provided by the two methods. In the ideal case, that is, using equation (1), doubling the number of repetitions increases S/N by 3 dB. As you might expect, the efficiency of equation (2)—stable averaging—is less; for a given number of repetitions we don't get quite so much noise attenuation. However, the loss is surprisingly small—it grows asymptotically to 0.77 dB as m becomes large, as shown in Fig. 3. Visually, the difference is insignificant; for proof, look at the CRT displays of Fig. 4, which compare stable averaging with

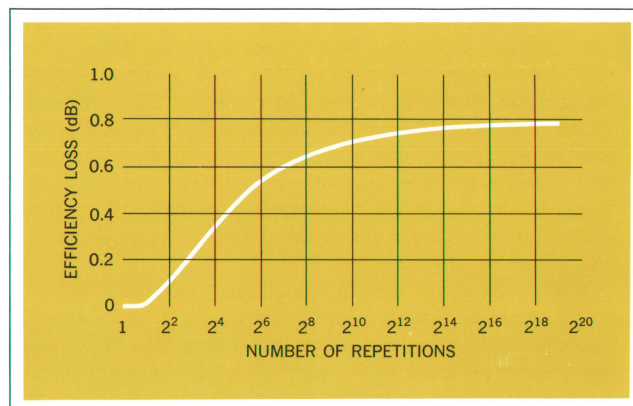


Fig. 3. For same number of repetitions, noise attenuation given by stable averaging is slightly less than that given by summation. Difference, or efficiency loss, grows asymptotically to 0.77 dB.

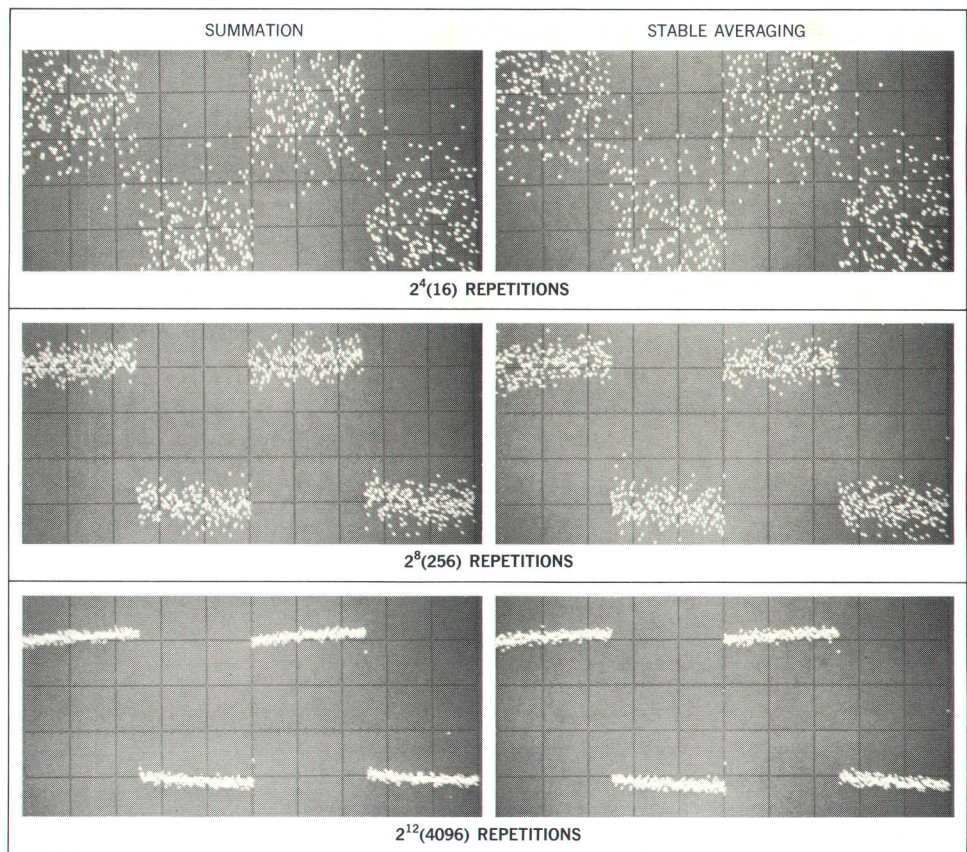


Fig. 4. It's difficult to see any difference between the results of stable averaging and the results of summation.

conventional summation in pulling a square wave out of random noise.

Actually, the price paid for a slight loss in averaging efficiency is only one of time. Stable averaging can give any S/N improvement within the limitations of memory size* — it just takes a little longer than summation.

A block diagram of the system we use for implementing stable averaging is shown in Fig. 5. Notice that the difference between the input and the old stored average is taken before the data is digitized. This means that the analog-to-digital converter is looking at the noise, which has an average value of zero. If there are noise spikes large enough to exceed the range of the A-to-D converter, the resulting roundoff, or clipping, errors will be symmetrical about zero and will not lead to amplitude distortion of the averaged signal.

Decaying Memory Follows Changing Signals

It is difficult to monitor slowly varying noisy signals using a strict averaging or summation technique. It is difficult because the averager's transfer function (specifically, the width of the comb-filter teeth) changes with each signal repetition. About the best that can be done

using conventional techniques is to take 'snapshot views' — that is, average a number of repetitions, look, reset, and average again.

What is really needed for slowly changing signals is an averaging algorithm that doesn't change with each repetition. Such an algorithm is

$$M_m^i = M_{m-1}^i + \frac{f(t_m + iT) - M_{m-1}^i}{X}$$

$$= \frac{1}{X} \sum_{k=1}^m \left(\frac{X-1}{X} \right)^{m-k} f(t_k + iT)$$

where X is a fixed integer, the same for all m .

As X becomes large, the factor $\left(\frac{X-1}{X} \right)^{m-k}$ approaches $e^{-(m-k)/X}$. Hence this algorithm produces an exponentially weighted average; at the m^{th} repetition, information obtained on the k^{th} repetition is weighted less heavily than the latest information by a factor of approximately $e^{-(m-k)/X}$. We don't derive it here, but the S/N enhancement for a signal mixed with random noise, using this weighted average, approaches a factor of $\sqrt{2X}$ as m becomes large.

In the new signal analyzer $X = 2^N$, where N can be chosen by turning a front-panel switch. N , called the

* Approximately 57 dB improvement in S/N can be obtained. However, the output S/N is limited to a maximum of about 57 dB by machine noise.

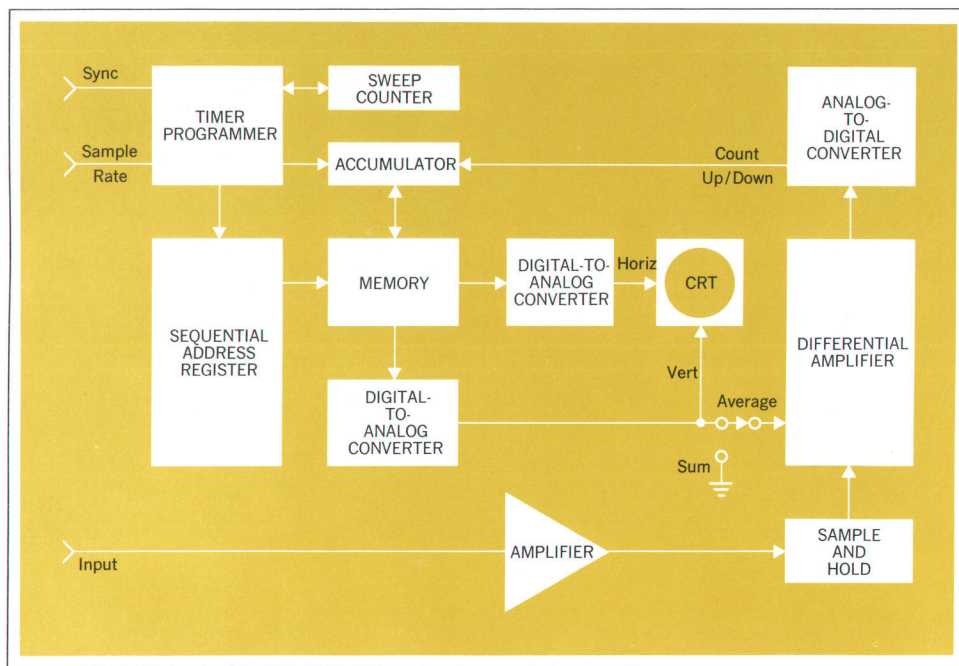



Fig. 5. Method of implementing stable averaging and summation. In stable averaging, input to A-D converter is difference between input and stored average; thus roundoff errors average to zero.



J. Evan Deardorff

Evan Deardorff graduated from Swarthmore College in 1963 with a BS degree in electrical engineering. He then joined a private research foundation, and from November 1963 to January 1965 he operated a cosmic-ray monitor in the antarctic. To recuperate from that experience, he spent the next several months traveling around the world.

During the 1965-66 academic year, Evan earned his MSEE degree at Stanford University and worked part-time in HP's nuclear instrumentation group. Since leaving Stanford, he has stayed with HP, helping to develop the 5480A Signal Analyzer.



Charles R. Trimble

Shortly after joining the HP Frequency and Time Division in June of 1964 Charlie Trimble was asked to look into signal averaging techniques. After proving the feasibility of many of the ideas embodied in the HP Signal Analyzer system, he was appointed group leader and given the responsibility for developing the HP 5480A and associated digital signal-processing equipment.

Charlie graduated with honors from the California Institute of Technology in June of 1963 and received his MSEE there a year later. He has several patents pending in the digital signal-processing field.

'sweep number,' determines the tradeoff between S/N improvement and the speed at which the averaged signal will follow the input. An increase of one in N will enhance S/N by 3 dB more, but the average will take twice as long to adapt to changes in the input. The time constant for adapting to changes is approximately 2^N divided by the signal repetition rate. For example, if the repetition rate is 50 Hz and $N = 10$, the time constant of change is about 20 seconds. If the signal changes abruptly, say from a square wave to a triangular wave, the average will take about three time constants, or 1 minute in this case, to catch up with the input again.

When exponential averaging is being used, the analyzer averages the first 2^N repetitions in the stable-averaging mode. This means the displayed average is calibrated from the beginning, instead of building up exponentially to its final value. Maximum efficiency in enhancing S/N is also obtained.

The exponentially weighted average or 'decaying memory' mode can be very helpful in setting up experiments. Often an improvement of only 3 to 6 dB in S/N is enough to show the experimenter whether any experimental parameters need adjusting, and with this small S/N improvement, the average will follow the adjustments quite rapidly. Thus the experimenter can zero in quickly on the parameters he wants.

Organization and Controls

The mechanical configuration of a special-purpose machine determines in large part its usefulness for tasks

other than those for which it was specifically designed. Ability to perform a given task is important, but more important is the ability to do it conveniently. Ideally, front-panel controls should be there when necessary and disappear when unnecessary. One practical way to tailor a machine to a particular class of problems while retaining flexibility is to use a plug-in approach. This philosophy has been followed in designing the new signal analyzer.

The analog plug-in (right-hand plug-in in Fig. 1) provides the interface with the specific class of experiments. The first analog plug-in to be designed is a dual-channel averager. It holds input amplifiers, attenuators, multiplexers, sample-and-hold circuitry, and special-purpose display controls. The functions of this plug-in's controls should be apparent in the light of what has already been said about how the machine works.

The logic plug-in holds program logic and special-purpose knob controls. In the first logic plug-in, programming is provided for summation, averaging, histogram generation, and multi-channel scaling. Controls on this plug-in select the sweep rate (which also determines the sampling rate), the number of repetitions to be averaged, whether the average is to be stable (PRESET) or weighted (NORMAL), and which 10 bits of the 24-bit

memory words will be displayed (VERTICAL EXPANDER). For histogram generation, a PRESET TOTALIZER fixes the number of points to be accumulated. This plug-in also contains the sync circuitry, along with a pre-analysis delay circuit which delays the beginning of processing with respect to the sync pulse, and a post-analysis delay circuit which determines the timing of an output sync pulse and the instrument 'dead time.'

The main frame of the analyzer contains the memory, general purpose registers, the CRT display, and the power supply. These are general elements used in most signal processing and display operations.

Acknowledgments

The consistent support of Alfred Low, first as a technician and later as a product designer, was crucial to the development of the signal analyzer. Charles N. Taubman's theoretical and practical contributions to the electrical design were invaluable. Richard V. Cavallaro, David A. Bottom, and in recent weeks James D. Nivison and John W. Nelson, have played an important role in tying together the innumerable details that make or break a project. Finally, appreciation must be expressed to our hard working, cheerful, and able wiring girls, Barbara M. Ahrens, Erika F. Leger, and Jean Cypert.

SPECIFICATIONS

HP Model 5480A Signal Analyzer

(with HP 5485A Dual Channel Averager plug-in
and HP 5486A Process Control plug-in)

AVERAGING (3 methods): Up to 60 dB signal-to-noise ratio improvement.

STABLE AVERAGING: Continuous calibrated on-line display. Signal amplitude remains the same as noise is attenuated.

WEIGHTED AVERAGING: Permits signal enhancement of slowly varying waveforms by exponential weighting of previous information with respect to new information. SWEEP NUMBER setting determines speed at which the average signal follows input.

SUMMATION AVERAGING: Algebraic summation process. Signal will grow from stable base line. If placed in AUTO mode, display will be automatically calibrated at the end of the preset number of sweeps.

SWEEP NUMBER: Manually selected. Dial is arranged in binary sequence (2^N) from single sweep (0 dial position) to 2^{19} (524,288) sweeps.

SWEEP TIME (horizontal sweep): Internally generated sweep time is calibrated in s/cm. Adjustable in 15 steps, in a 1, 2, 5 sequence, from 1 millisecond/cm to 50 s/cm. External sweep can be either sawtooth or triangular wave.

PRE-ANALYSIS DELAY: Variable in 15 steps from zero to 0.5 second.

POST-ANALYSIS DELAY: Continuously variable from 0.01 to 10 seconds.

SYNCHRONIZATION (Three modes):

INTERNAL: Pulse available at back panel. Can be used to trigger stimulus, and is controllable by POST-ANALYSIS DELAY.

EXTERNAL: Requires 100 millivolt rms signal with rise time greater than 10 milliseconds.

LINE: Synchronized to power line frequency.

NOISE MONITORING: CRT display of difference between raw input signal and memory stored average signal. Also available at back panel connector for variance analysis.

HISTOGRAMS: Probability density generation with respect to time interval and frequency.

TIME INTERVAL: Time between synchronization pulses. Horizontal calibration by time base.

FREQUENCY: Start and stop determined by time base. Horizontal calibration by time base.

TOTALIZING: Total count can be preset from 100 to 10,000,000 in magnitudes of 100 (10^2 , 10^3 , 10^4 , ... 10^7).

MULTICHANNEL SCALING: Counting rate up to 10 MHz. Horizontal calibration by time scale.

INPUT CHARACTERISTICS: Two channels with polarity switch for each channel. Channels can be used individually or their inputs can be summed.

COUPLING — ac or dc.

INPUT IMPEDANCE: Exceeds 1 M Ω shunted by 25 pF.

BANDWIDTH: From dc (2 Hz ac coupled) to more than 50 kHz.

SAMPLING RATE: Up to 100 kHz.

INPUT SENSITIVITY: Adjustable from 50 millivolts/cm to 20 volts/cm in 12 steps (1, 2, 5 sequence) with $\pm 3\%$ accuracy.

ANALOG-TO-DIGITAL CONVERTER: Ramp type with variable resolution 1 ms/cm sweep time has 5 bit resolution. 2 ms/cm sweep time has 7 bit resolution. 5 ms/cm or slower sweep time has 9 bit resolution.

MEMORY: 1024 word x 24 bit magnetic core memory, 1000 words (addresses) are used for data storage. Can be divided in half or quarter with independent selection for each channel. CRT display is always full scale regardless of memory sectioning.

DISPLAY: 8 x 10 cm rectangular display CRT with internal graticule. 500 kHz bandwidth. 10 bit horizontal and vertical resolution of digital display. 1, 2, or 4 trace display depending on input channels used and memory sectioning. Independent vertical position and gain adjust for each channel. Vertical expander permits selection of suitable 10 bit vertical display.

SYSTEMS CONTROLS

RESET: Clears displayed memory sections. Requires pushing two buttons simultaneously to avoid accidental erasure of memory.

START: Initiates data accumulation.

STOP: Stops accumulation at end of cycle.

CONTINUE: Resume data accumulation.

DISPLAY: Continuous display of processed data. Goes automatically in this mode at the completion of preset sweeps.

OUTPUT: Cycles through memory at rate determined by SWEEP TIME control.

BACK PANEL CONNECTION: Complete access to analog and stored digital information. Also provides for remote control. Convenient interface with other equipment.

GENERAL

POWER: 115/230 V, 50-400 Hz, 175 W.

WEIGHT: 76 lbs (34.5 kg) net.

PRICE: HP 5480A including HP 5485A and HP 5486A plug-ins \$9,500.00.

HP MODEL 5495A Input/Output Coupler (See article, page 14)

MANUFACTURING DIVISION: HP FREQUENCY AND TIME DIVISION
1501 Page Mill Road
Palo Alto, California 94304

Off-Line Analysis of Averaged Data

This new input/output coupler makes the new HP signal averager compatible with a computer and peripheral equipment.

By Francis J. Yockey

IMPORTANT AS ON-LINE SIGNAL PROCESSING IS FOR MONITORING the progress of an experiment, final results are often obtained only after further, off-line, processing of the collected data. The new signal analyzer described in the preceding article is an on-line, analog-in/analog-out machine. To facilitate off-line processing of the averaged data stored in the analyzer's memory after an experiment, an input/output coupler has been designed.

The coupler is a special-purpose digital instrument. It has three principal functions.

- It converts the 24-bit binary words in the analyzer's memory to decimal equivalents (7 digits plus sign) for display on built-in Nixie® tubes or for output to a printer or teleprinter.
- It provides for permanent storage of averaged waveforms and for reading-in waveforms or parts of waveforms. In this mode of operation, the coupler acts as a buffer between the analyzer's memory and a teleprinter, a tape punch, and a tape reader. It transmits characters in standard ASCII code.
- It provides for processing the waveforms in the analyzer's memory and for moving waveforms from one

© Registered TM, Burroughs Corp.

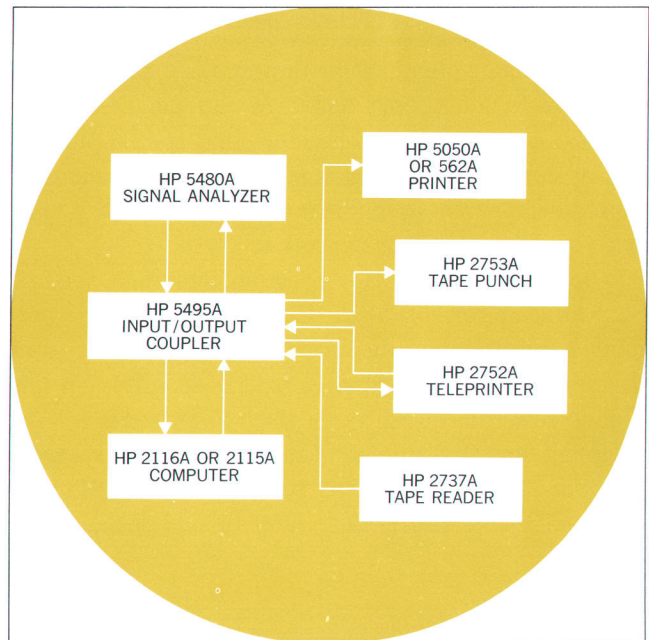


Fig. 1. HP Model 5495A Input/Output Coupler makes HP 5480A Signal Analyzer compatible with computer, printer, teleprinter, tape reader, and tape punch. Any or all peripherals can be added to system simply by plugging card/cable assemblies into back of coupler.

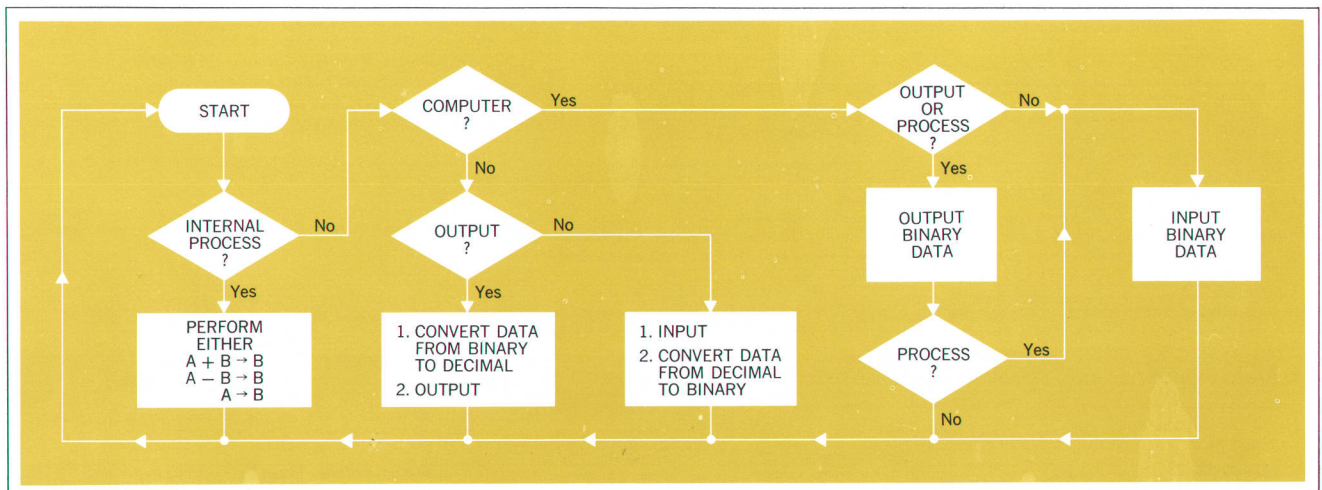



Fig. 2. I/O coupler provides digital display, input and output data transfers, and simple processing of data stored in memory of HP 5480A Signal Analyzer.

Fig. 3. Typical signal-analyzer/computer system contains HP 5480A Signal Analyzer for on-line signal processing, and computer and peripherals for off-line processing and storage. New Model 5495A I/O Coupler is instrument with digital display tubes. Small box on top of coupler is specially built program selector which provides pushbutton selection of computer programs listed in Table I.



part of memory to another. The coupler can add or subtract two waveforms (to 24-bit precision) and return the result to memory. For more complex processing, the coupler interfaces the analyzer with a general-purpose digital computer.

A minimal system includes only the analyzer and the coupler. This provides the digital display and the ability to do simple processing. The computer and other peripherals—any or all of the devices shown in Fig. 1—can be added at any time, simply by plugging card/cable assemblies into the back of the coupler.

Input/Output Functions

In outputting a data point to the display tubes or to one of the peripherals, the I/O coupler converts the 24-bit binary words in the analyzer's memory to signed seven-digit decimals. The decimal output is calibrated in centimeters.

A marker control allows the operator to choose one memory location to be displayed or to be the first one transmitted. The chosen point is intensified on the analyzer's CRT. When the digital display is being used, the display tubes show either the address of the chosen memory location or the value stored there. As the marker is moved, the digital display follows.

When the printer, the tape punch, or the teletypewriter is the output device, the address of the chosen memory location is transmitted first, followed by the value of the chosen point and the values of all succeeding points until outputting is manually stopped or the end of the waveform is reached. The teletypewriter output is formatted in

blocks of 5, 10, and 50 words so that data points can be identified easily. The punched tape is in the same format as the teletypewriter data; hence it can be listed on any standard ASCII-coded machine.

When data is being read into the analyzer's memory, the coupler converts the signed seven-digit decimals from the input device to 24-bit binary words. Data is stored in consecutive memory locations starting with the address specified by the first piece of data transmitted by the input device. When data is being entered from the teletypewriter, the operator can change individual points of a waveform, or type in an entire waveform.

A Typical Computer System

Complex computer processing of the averaged data in the analyzer's memory is accomplished by transmitting

Table I

Fourier transform	}	Display real part, imaginary part, magnitude, or phase of result.
Inverse Fourier transform		
Complex multiply		
Complex divide		
Square root		
Integrate		
Differentiate		
Low-pass filter	}	Choice of 8 time constants.
Low-pass restore		
High-pass filter		
High-pass restore		
Rotating storage (8 quarter-memory waveforms or 4 half-memory waveforms)		
Dc, average, peak, and rms computation with alpha- numeric display.		
Removal of dc component from waveform.		

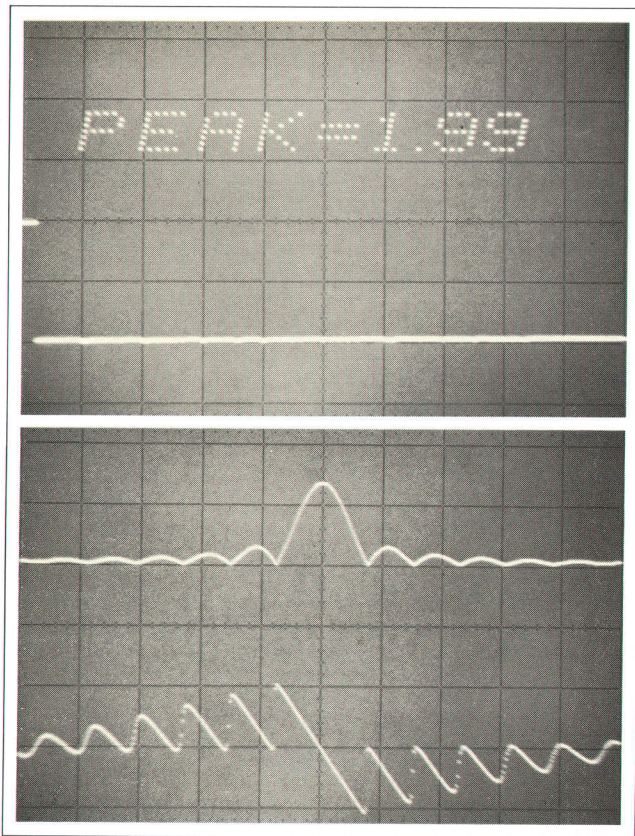


Fig. 4. Top: Rectangular pulse and its peak value. Pulse was read into analyzer from paper tape, then transmitted via I/O coupler to computer, where peak value was calculated. Alphanumeric peak information was then transmitted back to analyzer's memory for display on CRT. Bottom: Magnitude and phase of Fourier transform of same pulse. System of Fig. 3 produced both displays.

an entire memory segment to the computer, via the I/O coupler. After being processed by the programs stored in the computer, information is transferred back to the analyzer — via the coupler — for display on the CRT.

The coupler has been designed to provide data lines which are compatible with a wide variety of general-purpose computers. Each 24-bit binary word is divided into three eight-bit segments, which are then transmitted in sequence, eight bits at a time. Thus the input from the computer can be made to look like the input from a high-speed tape reader, and the output to the computer can be treated like the output to a high-speed tape punch.

Fig. 3 shows a typical analyzer/computer system, including the I/O coupler. The computer is an HP 2116A with an 8192-word memory. To make this system more user-oriented, a special control box provides push-button program selection. Once the master programs have been loaded into the computer and the computer set to its

starting address, the operator can forget about the computer controls. The computer becomes, in effect, a 'black box'.

The program-selector box contains only a set of switch closures to ground and a few diodes for programming. Because each user may have different needs, there is no standard configuration. Each user can easily make up his own program selector.

Table I lists the operations that can be called for by pushing buttons in the system of Fig. 3. They include such things as Fourier transforms, integration and filtering. Additional functions — e.g., auto- and cross-correlation, inverse filtering, special filter functions, finding transfer characteristics, and so on — can be generated by combining two or more of the functions listed. The speed of execution of these functions ranges from about four seconds to take the Fourier transform of a 500-point waveform, to well under a second for most of the other functions. The Fourier-transform time can be reduced by a factor of about five by fitting the computer with its hardware multiply and divide option.

Fig. 4 shows the results of taking the Fourier transform and the peak value of a rectangular pulse. The pulse was read into the analyzer from a paper tape, then transmitted to the computer by the I/O coupler. The computed Fourier transform and alphanumeric peak value were then transmitted back to the analyzer for display on the CRT.

Acknowledgments

I would like to thank Lawrence A. Lim for his valuable support in designing the mechanical parts of the Model 5495A Input/Output Coupler.



Francis J. Yockey

Frank Yockey received B.S. and M.S. degrees in electrical engineering from the University of Michigan in 1964 and 1965. He joined HP in 1965.

The first instrument Frank worked on at HP was the 5105A Frequency Synthesizer. His next responsibility was his current one, designing the 5495A I/O Coupler and the associated interface cards and programs.

Frank is a member of IEEE, Tau Beta Pi, Eta Kappa Nu, and Phi Kappa Phi.