


NOVEMBER 1982

HEWLETT-PACKARD JOURNAL



5180A WAVEFORM RECORDER
HEWLETT-PACKARD



Contents:

- 3** **Waveform Recording with a High-Dynamic-Performance Instrument**, by James L. Sorden and Mark S. Allen *A large memory and flexible triggering capabilities complement its tested, fully specified performance.*
- 9** **Designing a Ten-Bit, Twenty-Megasample-per-Second Analog-to-Digital Converter**, by Arthur S. Muto, Bruce E. Peetz, and Robert C. Rehner, Jr. *Custom hybrid and integrated circuits accurately sample and digitize a signal in forty nanoseconds.*
- 21** **Measuring Waveform Recorder Performance**, by Bruce E. Peetz, Arthur S. Muto, and J. Martin Neil *Realistic dynamic tests are the key to user confidence in the quality of the recorded waveshape.*
- 29** **Time Base Requirements for a Waveform Recorder**, by Steven C. Bird and Jack A. Folchi *Time base instability causes the time between samples to vary. Amplitude errors are the result.*
- 34** **Display and Mass Storage for Waveform Recording**, by Christina M. Szeto and Michael C. Detro *This waveform recorder companion provides functions needed in many applications.*

In this Issue:



Suppose you're a manufacturer of high-quality flash bulbs. It's important for you to know just how the light output of your bulbs varies with time. You know the light doesn't appear and disappear instantly, because physical things never work that way. More likely, the light output starts from zero, builds up rapidly to a peak, and then decays back to zero at some rate. But how do you find the exact shape of that curve?

If you're the owner of an HP 5180A Waveform Recorder, you can aim the flash at a photoelectric detector that converts light to voltage, connect the voltage output of the detector to one of the 5180A's inputs, and fire the bulb. The waveform recorder will sample the detector voltage at a rate of 20 million times per second, convert each sample to a binary number using a circuit called an analog-to-digital converter, and store the numbers—up to 16,384 of them—in a digital memory. If enough samples are taken, the waveform can be reconstructed from them. Since the samples are stored in a memory, the waveform can be reconstructed over and over again. It can be displayed on a screen with the time scale expanded so you can see the shape of the light output curve, or you can transmit the numbers to a computer you've programmed to analyze this data.

Examples of transient signals abound in industry, and waveform recorders to capture them can be obtained from several manufacturers. The major contribution of the 5180A, HP's first waveform recorder, is that its performance is not only very good, but is fully tested and specified for realistic, dynamic input signals. Most users buy waveform recorders based on static specifications, mainly resolution and sampling rate, and rely on the sampling theory of Swedish-born physicist/engineer Harry Nyquist for assurance that the data taken and stored by the waveform recorder accurately represents the input signal. As the article on page 21 demonstrates, this isn't necessarily so. The 5180A's designers had to break new ground to develop test methods and specifications that would provide the assurance the user needs and the Nyquist theory can't give for real waveform recorders under realistic operating conditions. There's now an Institute of Electrical and Electronics Engineers committee working on industry-wide standards for testing analog-to-digital converters, but at present no standards exist, and to its designers' knowledge, the 5180A is the only recorder that is tested and specified so completely.

The 5180A makes other contributions, too, including its large memory, its facilities for precise user control of what it records, and several other features new to waveform recorders. Pages 3 to 34 of this issue cover its design and testing. A companion product, Model 5181A Display/Tape Storage Module, is described beginning on page 34. The 5180A has no built-in CRT display. The 5181A provides one, along with a cartridge tape unit for recording 5180A data permanently or bringing field-recorded signals to the 5180A for digitizing.

Our cover photo this month shows the 5180A Waveform Recorder and the digitizer hybrid circuit, a key component of the 5180A's analog-to-digital converter.

-R. P. Dolan

Waveform Recording with a High-Dynamic-Performance Instrument

This new waveform recorder digitizes and stores single-shot or repetitive signals. Its ten-bit, 20-MHz analog-to-digital converter delivers exceptional performance that is fully specified and characterized under realistic operating conditions

by James L. Sorden and Mark S. Allen

WAVEFORM RECORDERS are general-purpose, time-domain measuring instruments that take samples of analog waveforms and store the sample values as digital data suitable for analysis by a computer. Waveform recorders offer performance far beyond that of other measurement techniques for analog signal capture and analysis. Waveform recorders are generally suited to measurements of both transient or single-shot waveforms and continuous or repetitive waveforms. Applications are found in virtually every area of science and technology. The following is a partial list.

- Electrical and electronic engineering (power supply testing, computer disc testing, servomechanism design, radar and communications system testing, waveform analysis)
- Materials science (pressure vessel testing, acoustic emission)

- Electromagnetic compatibility studies
- Energy research
- Physics and chemistry (laser spectroscopy, flash photolysis, reaction kinetics)
- Ultrasonics
- Explosives testing
- Ballistics
- Biomedical research
- Medical ultrasound

The new HP 5180A Waveform Recorder, Fig. 1, is a high-accuracy instrument that digitizes input waveforms at a rate of 20 million samples per second. It is the first waveform recorder whose dynamic performance is tested and fully specified so the user knows just how accurately the measurement results represent an input waveform, even a rapidly changing one with frequency components to 10 MHz. The 5180A offers 60-dB dynamic range, pretrigger




Fig. 1. Model 5180A Waveform Recorder digitizes input waveforms with high accuracy at a rate of 20 million samples per second. It has many features that make it easy to use and incorporate into automatic test systems.

waveform capture, and the ability to transfer the data to a computer. Its dynamic range and fast data transfer make the 5180A ideal for automatic test systems.

Elements of a Waveform Recorder

The block diagram of the 5180A Waveform Recorder, Fig. 2, is typical of waveform recorders in general. Elements common to all such instruments are:

- Input signal conditioning, often including multiple sensitivity ranges and high impedance.
- High-speed analog-to-digital converter (ADC). Many different and often novel techniques have been applied here. Scan-converter CRTs, analog shift registers using charge-coupled devices (CCDs), flash conversion ADCs without sample-and-hold, and other techniques are used in some of today's instruments.
- High-speed memory system capable of following the ADC output, or a slower ADC and memory system in scan-converter or CCD applications.
- Internal and external triggering circuitry to define the beginning and end of a measurement.
- Time base and oscillator system for control of the ADC sample interval.
- Digital system controller. All recent waveform recorder systems use microprocessor control.
- Digital I/O, usually HP-IB,* often supplemented with other methods such as RS-232-C/CCITT V.24 or full parallel direct memory access (DMA).
- Analog output capable of driving an external directed-beam or raster display and/or an internal CRT display.

In the 5180A, a 20-MHz-sample-rate, 10-bit ADC is followed by 16,384 10-bit words of memory (Fig. 2). The ADC samples the input signal at 20 million samples per second and the memory stores the resulting 10-bit digital codes at user-selectable rates.

Since the waveform data is captured in a digital memory, all digital waveform recorders offer some form of pretrigger, a feature found also in logic analyzers. In this mode, the memory records in a circular configuration.** The trigger is then used to stop the memory after an adjustable delay. If the delay is set to be the same as one sweep width, the result

*Hewlett-Packard Interface Bus, HP's implementation of IEEE Standard 488 (1978).

**Samples are stored sequentially in memory starting with the first word. When the last word of memory is filled, the next sample value is stored in the first word and recording continues as before.

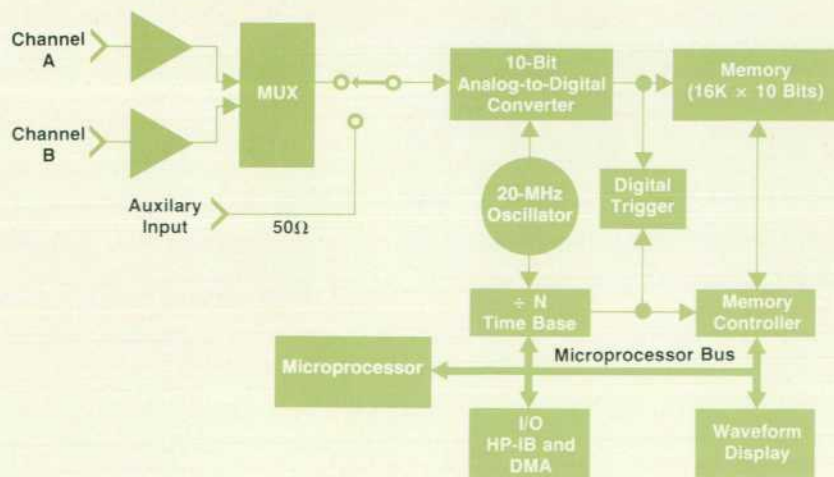


Fig. 2. The 5180A Waveform Recorder has a 20-MHz, 10-bit analog-to-digital converter and 16K words of memory. Channels A and B are high-impedance channels and the auxiliary input is a precision 50Ω channel.

is the same as a storage oscilloscope—the waveform is visible only after the trigger point. However, with adjustable pretrigger, the recorder can be set to show any amount of time ahead of the trigger so the cause of the trigger can be seen. With the 5180A, the pretrigger is adjustable from 0 to 100% of the record length.

Post-trigger delay can also be added to view a segment of the waveform well after the trigger. This capability is useful for characterizing various parts of a long waveform measured from a common trigger point. A television video signal is a good example of a signal that is easily measured using the delay features of the 5180A. The vertical frame sync pulse can be used as the trigger signal, and the 5180A trigger delay can be set to capture any of the 525 scan lines of data.

An obvious difference between a waveform recorder and a conventional storage oscilloscope is that the recorder stores data in a digital memory, so a waveform can be manipulated following its capture. It can be displayed on a standard X-Y display by replaying the stored data through a low-speed digital-to-analog converter (DAC). This gives a display that neither flickers nor fades with time. With a proper digital interface the digital data may be sent to a computer for analysis and plotted or returned for display.

Inside the 5180A

As shown in Fig. 2, the 5180A includes the elements typical of other waveform recorders, along with some new ones. The two 1-MΩ high-impedance channels can be used together in the Chop A,B mode or either can be used as a single input. For the finest performance there is an extra precision 50Ω channel, **AUX**.

The **AUX** channel is relay-switched in place of the high-impedance channels for applications requiring the highest bandwidth, lowest drift, and lowest noise. This input is optimized for use with 50Ω systems or customized input signal conditioning such as special anti-aliasing filters. The design of such filters is made considerably easier because the system bandwidth is far beyond anti-aliasing requirements. The **AUX** channel is overvoltage and power protected by an internal clamp circuit and a front-panel fuse.

Analog-to-Digital Converter

The ADC in the 5180A is designed to give an accurate




Fig. 3. Adjustable hysteresis makes it possible to set the 5180A to trigger reliably on small signals or on large noisy signals. (In this example, the signal must cross the lower threshold and then the upper threshold to trigger the recorder.)

representation of the input waveform (see article, page 9). Input bandwidths of the amplifiers and sample-and-hold circuit are much higher than the maximum sample rate to reduce the possibility that these circuits will influence the measured waveform. A complete set of specifications and tests to verify them provides the confidence that the 5180A will do what it claims (see article, page 21).

Digital Trigger

Historically, the second most difficult assignment for any measurement instrument of this class (after digitizing accurately) has been to trigger properly. The 5180A trigger circuit uses the digital data from the ADC. This makes the trigger level easily settable and repeatable. It is also possible to set the sensitivity of the trigger circuit. The 5180A has a hysteresis control which varies the width of the band that the signal must go through before a trigger can occur. This makes it possible to trigger reliably on both small signals (by setting the hysteresis small) and large signals with large amounts of noise. Fig. 3 illustrates this feature. In the top trace, the hysteresis is small and the machine triggers on noise. In the bottom trace the hysteresis has been adjusted to account for the noise.

The selectable hysteresis feature means that the 5180A actually has two settable trigger levels, one for the trigger and one for the hysteresis. By rearranging the circuit slightly, the 5180A can offer bittrigger, the ability to trigger on a signal if the absolute value exceeds a certain range. This feature allows triggering on transients where the exact amplitude and direction are not known. For example, it is possible to detect spikes on either the positive or the negative cycle of the 60-Hz line (Fig. 4).

High-Speed Memory

The 5180A has one of the largest memories available in a waveform recorder. While a conventional storage oscilloscope exhibits about 1% time resolution for any stored signal and many other waveform recorders offer 1-to-2K memories for 0.5 to 0.1% time resolution, the 16K samples of the 5180A provide 0.006% resolution. This improvement in time resolution results in a much higher probability that

the transient to be investigated will be fully captured. At the fastest sample rate of 50 ns per point, a 0.8-ms record is achieved.

The memory is implemented using fast NMOS static RAM (random-access memory) chips of the commercially available 2147 family. To ensure a reliable 20-MHz range, two ranks of 10 bits each are written into sequentially. The memory may be partitioned into 1 to 32 blocks. This allows the user to measure a long transient or divide the memory so that consecutive short events may be captured. An autoadvance feature is available to advance the memory location following each measurement so that up to 32 events may be recorded without operator intervention. Each event is captured with a corresponding clock time so that it may be reviewed later along with its time of occurrence.

Microprocessor

The 5180A's microprocessor section is based on the 8-bit 6800. There are four major subsections: processor, program storage, variable storage, and input/output (I/O). The microprocessor has the ability to read status from and control various attributes of the measurement sections (time base, memory control, input amplifier, trigger).

The processor design is based on a design used in the HP 5370A Universal Time-Interval Counter and the HP 5359A Time Synthesizer. The program storage consists of 36K bytes of ROM (read-only memory). 1408 bytes of NMOS RAM are used to store temporary variables such as numbers used in computations, and there are 512 bytes of battery-backed-up CMOS RAM for storage of the current front-panel settings and up to four additional setups.

The I/O subsection contains two interfaces for computers—HP-IB and parallel DMA—and two interfaces for bench use—XYZ display and front-panel keyboard. The HP-IB is implemented using a large-scale integrated (LSI) circuit chip. The microprocessor sends and receives data under software control. The parallel DMA interface was designed in conjunction with the memory control block to allow the processor to be shut off and data to be transmitted from the memory control, over the processor bus, and out the parallel DMA interface at a maximum of 1M words per second. Since speed is important, the design uses low-power and standard Schottky TTL ICs. The XYZ display has its own refresh RAM, address counters, and X and Y DACs and amplifiers. The microprocessor places data in the XYZ display RAM and the circuit continuously refreshes the display by sending RAM data to the DACs. A significant advantage of this is that the picture remains on the screen even during the next measurement, instead of blanking out.

(continued on page 8)




Fig. 4. Bittrigger mode captures transients of unknown slope.

Waveform Recording with the 5180A

Applications for waveform recorders range from subsonic seismic surveying and biomedical research to microwave frequencies and beyond, including radar and laser experiments. The following examples of electronic applications demonstrate the capabilities of the HP 5180A Waveform Recorder described in the accompanying article.

Power Supply Testing

A common transient in an instrument is the turn-on characteristic. For example, a reliable, well regulated and electrically quiet power supply is a must for a well designed instrument. The turn-on and load change responses of the power supply are extremely important. Detailed knowledge of the behavior of the supply becomes even more desirable when the supply has automatic over-voltage and overcurrent shut-down. Input surge currents can cause the supply to shut down as the instrument is powered on. A waveform recorder can be triggered on the shut-down signal and look backwards to find the cause of the shut-down and to show how the power supply voltage decays.

Microprocessor Applications

Another application involves the interaction of digital microprocessor circuits and an asynchronous interrupt. Fig. 1 illustrates the power-down sequence in the 5180A. A powerfail signal gives advance warning that the power is going down so that the microprocessor can calculate a checksum for the CMOS RAM that stores the front-panel settings. When the interrupt occurs, the microprocessor calculates the checksum and then write-protects the RAM. In a good design, the checksum routine is short enough that it can be completed before the power drops below the level required by the microprocessor.

In Fig. 1 the measurement was triggered on the powerfail signal. The 5180A was set to pretrigger about -5%. In other words, it was set to load about 5% of its memory with samples of the waveform before the trigger point. The measurement was made in the dual-channel (Chop A,B) mode with channel A connected to

the CMOS RAM disable line and channel B connected to the +5V power supply. The CMOS RAM disable signal is the result of the completion of the checksum routine.

Disc Memory Testing

In the area of computer disc testing the 5180A is useful for electromechanical analysis of the head positioning servo system and for analysis of the read/write electronics. The top trace of Fig. 2 shows a 16K record of data read from a disc. In this case, the disc is divided into sectors of 256 bytes each. At the end of each sector is a check word. If there has been an error in reading the sector, the check word will not agree with the one computed from the data read. This condition is used to trigger the 5180A in a -100% trigger measurement. The record then contains 16K points that represent the electronic signal from the read head during the reading of the sector in question. The data can be sent to a computer and a pattern recognition routine can be run to look for an error.

The bottom trace in Fig. 2 shows the error, a pulse that did not completely cross the logic threshold. There are many causes for errors in reading: imperfections in the magnetic material itself, noise in the amplifiers, poor head positioning, pulse crowding, and in flexible discs, head bounce. By isolating the problem with the waveform recorder, the disc can be exercised repeatedly for many hours and the errors saved for study. The characteristic shape of the error can give a clue to the cause of the problem.

Use with Spectrum Analyzer

A powerful measurement capability can be realized by combining a spectrum analyzer and the 5180A. The versatility of the spectrum analyzer in receiving and down-converting incident microwave signals matches the 5180A's ability to digitize and store an input. The setup is shown in Fig. 3. Whether the 5180A is connected to the spectrum analyzer's IF or video output, time-domain information from the RF input can be stored in the 5180A and processed as desired. Since the 5180A is oriented towards




Fig. 1. A 5180A measurement of the power-down sequence of another 5180A.

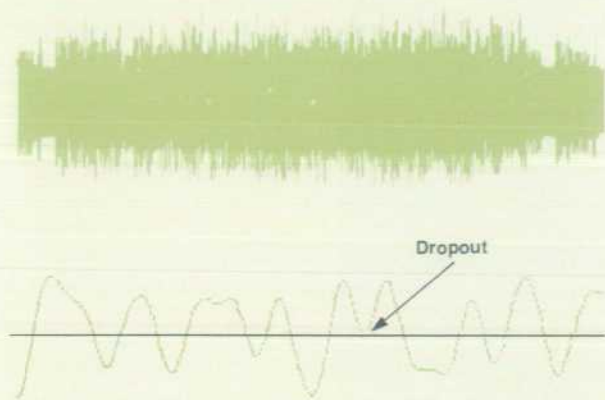


Fig. 2. Top trace shows data read from a disc. Expanded bottom trace shows an error, a pulse that did not cross the threshold.

single-shot data capture, transient or short-duration signals can be reliably monitored with this setup. Because the 5180A has pretrigger measurement capability, an RF burst can be completely captured, not just the portion following the trigger. Applications for the 5180A/spectrum-analyzer combination are found in radar system testing, in noise measurements, and in the processing of communications signals.

True-rms measurements can also be made with the setup of Fig. 3. In the zero-span mode, the spectrum analyzer becomes a variable-bandwidth, superheterodyne tunable receiver. The video output in zero-span mode is the time-domain envelope of the RF input. The IF output from the spectrum analyzer is the down-converted version of the input signal. For the HP 8566A (and many other HP spectrum analyzers), the IF center frequency is 21.4 MHz, and the IF bandwidth is controlled by the selected resolution bandwidth.

Although it is often written that the sample rate required to recover the information from a given input is $f_s > 2f_{max}$, where f_s is the sample rate and f_{max} is the highest frequency component in

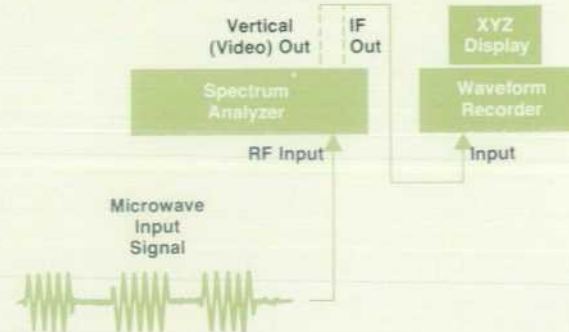


Fig. 3. Interconnections for using the 5180A Waveform Recorder with a spectrum analyzer to down-convert signals.

the input, this inequality is not completely general. To recover completely any input signal from the sampled input data, the minimum required sample rate is

$$f_{s(min)} > \frac{2f_{max}}{\text{Integer part of } [f_{max}/(f_{max}-f_{min})]}$$

This expression degenerates into the less general form when $f_{min}=0$. In words, this general expression states that it may be possible to recover completely a band-limited signal with a sample frequency smaller than any frequency in the input signal.

Choosing the widest possible bandwidth (3 MHz), and making the assumption that the skirts of the spectrum analyzer IF system frequency response are infinitely sharp, the minimum sample rate for complete recovery of the IF is

$$\frac{2(22.9)}{\text{Integer } [22.9/(22.9-19.9)]} = \frac{45.8}{7} = 6.5429 \text{ MHz}$$

The technique of using a waveform recorder as a sampling down-converter will only work if the recorder's bandwidth greatly exceeds $2f_{max}$. In this case, the 5180A's bandwidth is adequate.

Fig. 4 shows examples of communications signals captured by this method. The signals are television broadcast signals. Figs. 4a, 4b and 4c show vertical interval test signals, and Fig. 4d shows two lines of the picture.




Fig. 4. Television broadcast signals captured using a waveform recorder and a spectrum analyzer. (a), (b), (c) Vertical interval test signals. (d) Two lines of the picture.




Fig. 5. The 5180A's line display mode connects sample dots and provides a continuous display. (a) Square wave, dot mode. (b) Square wave, line mode. (c) Sine wave, dot mode. (d) Sine wave, line mode.

The front-panel displays and keyboard are driven from several LSI ICs. One IC drives the small digits and scans the 64 keys, interrupting the microprocessor each time a key is pressed. Two other LSI ICs drive the large digits and light the annunciators.

Software structure in the 5180A consists of a real-time operating system with foreground (interrupt) and background (measurement) divisions. I/O operates in the foreground under interrupts. Measurements and tasks that require the high-speed memory occur in the background. Communication takes place through locations in both the temporary and front-panel RAMs. Internal libraries supply HP-IB I/O routines, stack-oriented 32-bit and 48-bit floating-point math routines, integer math routines, and 4K bytes of HP-GL (Hewlett-Packard Graphics Language) routines to send an annotated graph of waveform data to a plotter.

XY Display

Waveform recorders sample waveforms, but the user is most interested in viewing a waveform as if it were continuous. When the input frequency is low with respect to the sample rate, the dots appear close together and the human eye connects them automatically. However, when the input frequency is high (20 samples per cycle) the dots become too far apart, and the eye has trouble. Various means can be used to connect the dots into a continuous display. The 5180A low-pass-filters the output of the display driver to obtain a smooth curve. The filter is optimized to provide linear phase response so that square waves will be displayed properly. It also has a sharp cutoff to reproduce sine waves faithfully. Fig. 5 shows sine and square waves before and after the line mode has been enabled.

Other Operating Features

The 5180A is completely programmable. All front-panel

controls are under microprocessor control. Light-emitting diode (LED) displays and annunciators reflect the current state of the instrument. These displays facilitate remote programming of the front-panel controls by echoing commands as they are sent. All fundamental settings have displays so that it is easy to monitor their current states.

A computer may obtain front-panel settings from the 5180A by sending one of many teach/learn commands. The settings are returned in the same ASCII* form used to program the 5180A. This form of output is more easily read by human operators, who can verify the settings. High-speed memory data is available in either ASCII for ease of use and readability or binary code for speed. A 16-bit parallel interface inputs or outputs waveform data at up to one million 10-bit words per second. This is extremely useful when a rapid series of measurements is needed. For example, 1K records may be read into an HP 9826A Computer with less than 4 ms between measurements.

Many uses of the 5180A require some form of hard-copy output. There is not always a need, however, for a computer to read and manipulate the data. The 5180A can output a completely annotated graph to any HP 9872B/C/S/T, 7225A, 7245A/B, 7470A, or 7580A/85A Plotter. The output uses the HP-GL plotter language and is sent over the HP-IB. Fig. 1 on page 6 is an example of this kind of output.

Like an oscilloscope, the 5180A is a complicated instrument with many choices of settings. To make it easier to use, the **AUTOSET** key causes the 5180A to search for and automatically acquire a repetitive signal. To do this, the 5180A makes measurements at two different sample rates and searches for a signal in the input channels. It adjusts the gain of the channel so the signal nearly fills the vertical dimension of the CRT and then adjusts the time base so that

* ASCII = American Standard Code for Information Interchange

from two to five cycles of the input signal are visible. The trigger level is set at the 50% point of the waveform, and the hysteresis is set to 25% of the peak-to-peak amplitude. Periodic waveforms from 40 Hz to 10 MHz can be acquired.

Once the desired settings have been achieved, they may be saved in the nonvolatile CMOS memory. The current settings are always saved when the power is shut off. Four more settings can also be saved and recalled. This provides benchtop convenience when multiple standard measurements are made, and reduces the number of bytes that need to be sent when the 5180A is used on the HP-IB.

James L. Sorden



Jim Sorden joined HP in 1964. Currently an R&D section manager at the Santa Clara Division, he was project manager during the early development of the 5180A. He has been project manager for several R&D assignments, including the 5345A Frequency Counter. Having served in the U.S. Army from 1955 to 1958, Jim graduated from the University of Wisconsin in 1962 with a BSEE, then did graduate studies at Wisconsin and Stanford. Jim has authored technical articles in a variety of publications on the subject of frequency counters. His work has resulted in one U.S. patent

and other patent applications relating to frequency counter and waveform recorder instrumentation. He is presently the U.S. technical advisor to the IEC for electronic measurement instrumentation. Away from HP, Jim is involved in stock and real estate management and in family activities, including skiing, sailing, home remodeling and coaching Little League baseball and youth soccer. A native of Wisconsin, Jim is married, has two children, and lives in Saratoga, California.

Acknowledgments

A large and intensely dedicated group of engineers worked long, hard and creatively to make the 5180A a reality. We wish to thank the authors of other articles in this issue and the dedicated efforts of the people they acknowledge. In addition, there are many others who dedicated their engineering talents to the project. In R&D, special recognition must go to Phil Deaver, Bill English, Ron Felsenstein, Kenzo Ishiguro, Kreg Martin, Ralph Smith and Ron Young. The product introduction was made especially successful by the efforts of the following engineers in other departments: Mark Brouder, Tom Carrico, Kim Gray, Gary Ikari, Doug Nicols, Art Weigel and Eric Youngberg. Finally, we must acknowledge Ian Band, Jim McNeish and Jack Lieberman for their vision and wisdom, their confidence that we could do the job and that the job was worth doing.

Mark S. Allen



Mark Allen is a project manager at HP's Santa Clara Division. A native of Los Angeles, he received the BS degree in 1973 and the Master of Engineering degree in 1974 from Harvey Mudd College. With HP since 1974, he's worked on the 5370A Time Interval Counter and the 5180A Waveform Recorder. His work has resulted in a patent on an automatic arming control scheme for time interval measurements, and he's co-author of a paper on IEEE 488 and microprocessors. He is married, has a daughter, and is remodeling his home in Saratoga, California. He also likes

woodworking, home computers, swimming, skiing, ice skating, and sailing.

Designing a Ten-Bit, Twenty-Megasample-per-Second Analog-to-Digital Converter System

by Arthur S. Muto, Bruce E. Peetz, and Robert C. Rehner, Jr.

THE ANALOG-TO-DIGITAL CONVERTER (ADC) in HP's new 5180A Waveform Recorder is a custom-designed, high-performance converter capable of 20 million conversions per second with 10-bit (2-mV) amplitude resolution, equivalent to 60 dB of dynamic range. Performance typically degrades by only one bit from dc to the Nyquist or "folding" frequency (input frequency equal to one-half the conversion rate). Credit for this reli-

able performance is due in large part to two hybrid substrate assemblies that contain 12 custom integrated circuits.

High-speed ADCs that can encode or digitize continuously are usually "flash" or all-parallel designs. These converters are extremely fast because an N-bit converter uses $2^N - 1$ comparators to compare the analog input voltage to $2^N - 1$ precision reference voltages. For a high-resolution converter the amount of circuitry required can become pro-

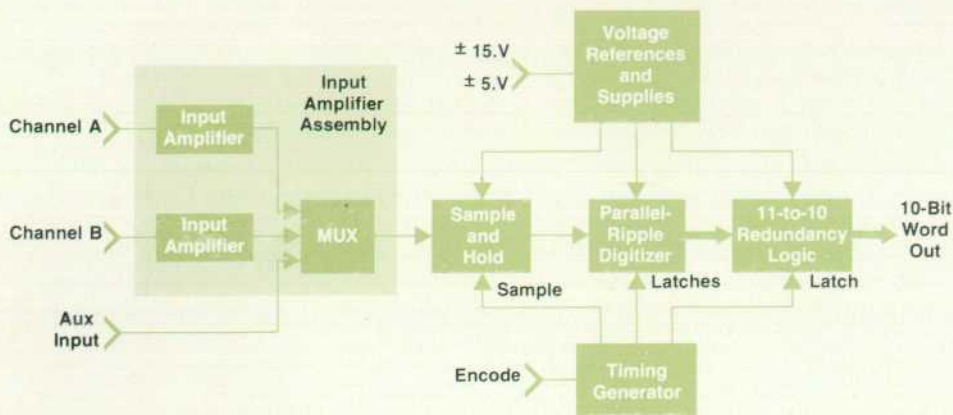


Fig. 1. 5180A Waveform Recorder parallel-ripple analog-to-digital converter (ADC) block diagram.

hibitive. A 10-bit flash converter, for example, requires 1023 comparators in addition to the necessary logic circuitry to decode the comparator outputs to a 10-bit binary code. The large number of high-speed comparators leads to excessively high input capacitance and to prodigious power supply requirements.

A technique that trades conversion speed for circuit simplification is the parallel-series or parallel-ripple approach. In a two-pass converter using the parallel-ripple technique, a coarse flash encode of the input voltage is performed to get the most-significant bits. Then the analog voltage equivalent of the coarse encoder's digital output is generated and this is subtracted from the input voltage. The difference voltage, or residue, is then digitized by a fine flash encoder to get the least-significant bits.

The parallel-ripple design was selected for the 5180A ADC instead of the flash converter because of the tremendous reduction in the number of comparators needed. A 10-bit, two-pass, parallel-ripple converter that includes redundancy (whose purpose will become evident when the digitizer is described in detail later) can be built with less than 100 comparators—an order-of-magnitude reduction. However, there are significant tradeoffs. Not only is the parallel-ripple converter slower, but it also requires several difficult-to-build components: a high-speed digital-to-analog converter (DAC) that transforms the first-pass word into an analog level, a high-speed operational amplifier that performs the subtraction, and a high-performance sample-and-hold circuit that guarantees that the voltage input to the parallel-ripple digitizer does not change between the first-pass and second-pass conversions.

The complete parallel-ripple ADC in the 5180A is contained in a module consisting of five printed circuit board assemblies that perform the following functions: sample-

and-hold, parallel-ripple digitizer, precision reference and power supply, precision timing generator, and redundancy logic circuitry. The sample-and-hold and digitizer assemblies include hybrid circuits that contain five of six custom integrated circuits specifically designed for the 5180A. The sixth custom design, found in the input amplifier hybrid, is described in the box on page 18.

Fig. 1 is a block diagram of the ADC module.

Sample-and-Hold Design

In HP's parallel-ripple ADC, a sample-and-hold circuit (S/H) maintains the digitizer's input voltage constant while both encoder passes digitize the input. A simplified S/H block diagram appears in Fig. 2. The S/H is directed to acquire a point on the S/H input waveform by a positive pulse on the sample line. This pulse forces the sampling switch to close, causing the hold capacitor to charge to the voltage at the input to the switch and then follow or track it. This voltage is a buffered attenuated version of the S/H input. When the pulse returns low, the switch opens and the voltage on the capacitor is held because of the high impedance at this node. The postamplifier buffers the hold capacitor voltage from the digitizer and amplifies it by a factor of two to compensate for the resistive attenuator at the preamplifier output. The function of the attenuator will be explained later.

Four custom ICs were designed for the S/H using HP's proprietary 5-GHz f_T process, highlights of which are presented in the box on page 15. Two of the ICs are used in the preamplifier and postamplifier. The remaining two ICs—the sample gate and gate driver—were designed for the sampling switch circuitry shown in Fig. 3. The sampling switch is a monolithic Schottky diode bridge that is contained on the sample gate chip, a photomicrograph of

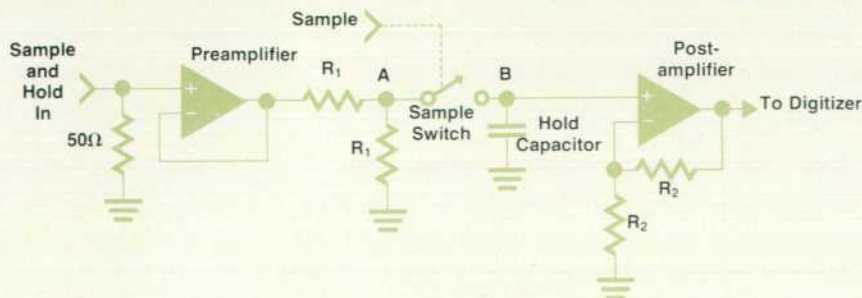
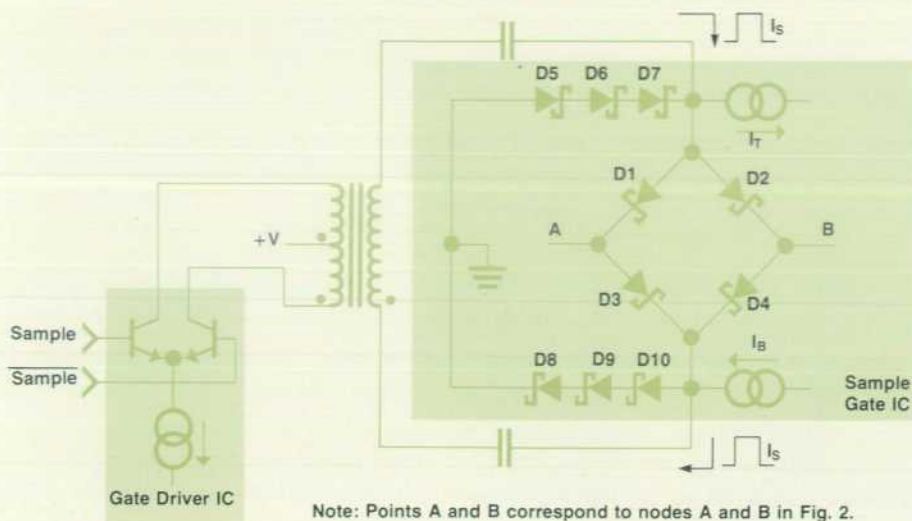


Fig. 2. Simplified sample-and-hold circuit block diagram.



Note: Points A and B correspond to nodes A and B in Fig. 2.

Fig. 3. Sampling switch circuit contains two custom integrated circuits.

which is shown in Fig. 4a. Since the bridge is monolithic, diodes D1 through D4 are better matched and have smaller parasitic elements than bridges made from discrete diodes. As a result, when the switch is closed by passing a current I_S through the bridge, the bridge remains fully balanced. For example, if the voltage at point A is initially much more negative than the hold capacitor voltage, then the hold capacitor will be linearly discharged by the full I_S current through D4. Since the bridge is well balanced, D1 also conducts I_S while D2 and D3 are nonconducting. The resistor attenuator between the preamplifier output and the bridge prevents this I_S current, which can be as large as 50 mA, from disturbing the linear operation of the preamplifier. High slew rates are characteristic of this linear portion of the acquisition of the input sample because of the large magnitude of the I_S current. When the hold capacitor voltage is within about 200 mV of the voltage at point A, then D2 and D3 begin to conduct and a slower RC-limited discharging takes place. This slower portion of the acquisition, controlled by the hold capacitor and the resistance it sees—one-half R_1 plus the resistance of the diode—continues until all four diodes conduct equally. At this point acquisition of the input sample is complete and the voltage on the hold capacitor continues to track the input waveform until the sampling switch is opened by reducing the current I_S to zero as quickly as possible to prevent contamination of the hold capacitor voltage by undesirable diode current mismatches.

The sample gate IC also contains circuitry needed during the 40-ns hold time, when the voltage on the hold capacitor must be kept constant. When I_S is equal to zero, the matched npn current sink and pnp current source drive 1-mA currents through the series diodes, D5 through D10. This reverse-biases the bridge diodes and lowers the impedance at the top and bottom of the bridge, attenuating the feed-through or coupling from point A to the hold capacitor across the 1-pF junction capacitances of the reverse-biased bridge diodes.

The current I_S that turns the bridge on is supplied by the gate driver IC, a photomicrograph of which is shown in Fig. 4b. The gate driver is simply two cascaded emitter-coupled

amplifiers designed to deliver large current pulses having subnanosecond rise and fall times. Fast transitions are needed to define the sampling time unambiguously. When the S/H is directed to sample through the differential sample and sample lines, the open-collector outputs of the gate driver IC deliver a 10-ns-wide current pulse to the primary of the center-tapped pulse transformer. The secondary current of the trifilar-wound transformer is nearly double the primary current (core losses prevent exact doubling). The secondary current is capacitively coupled to the sampling bridge and can be thought of as a floating 50-mA current source connected across the top and bottom of the bridge during the 10-ns sampling time. The large current amplitude results in fast slewing of the hold capacitor and 100% sampling efficiency. This level of sampling efficiency, which indicates that the accuracy of each sample is independent of the value of the previous sample, is an absolute requirement for one-shot applications like the capture of transient phenomena.

The third special IC designed for the S/H is a high-speed operational amplifier (op amp), whose photomicrograph appears in Fig. 5. This amplifier was designed to satisfy both the preamplifier and postamplifier functions, and the same basic op amp is also used in the digitizer and the input amplifier. Fig. 6 is a simplified schematic of the op amp. Q1 and Q4 are emitter-follower inputs to the gain stage formed by emitter-coupled transistors Q2 and Q3. To minimize the phase delay through the amplifier, important for high-frequency phase margin, only a single stage of gain is used. The open-loop gain of the amplifier is adjusted with a metal mask option which selects $R=10, 50, \text{ or } 75$ ohms. The lower the desired closed-loop gain, the larger the emitter resistors must be to increase phase margin by lowering the open-loop gain. Q5 and Q6 are cascode transistors that increase the amplifier's bandwidth by minimizing the Miller capacitance of the emitter-coupled pair. The three outboard pnp transistors, Q7 through Q9, are connected as an active current mirror and form a high-impedance active load to maximize the amplifier's open-loop gain, typically 60 dB. Transistor Q10 and Zener diode D1 form an emitter-follower buffer and level-shift stage. The output is also an

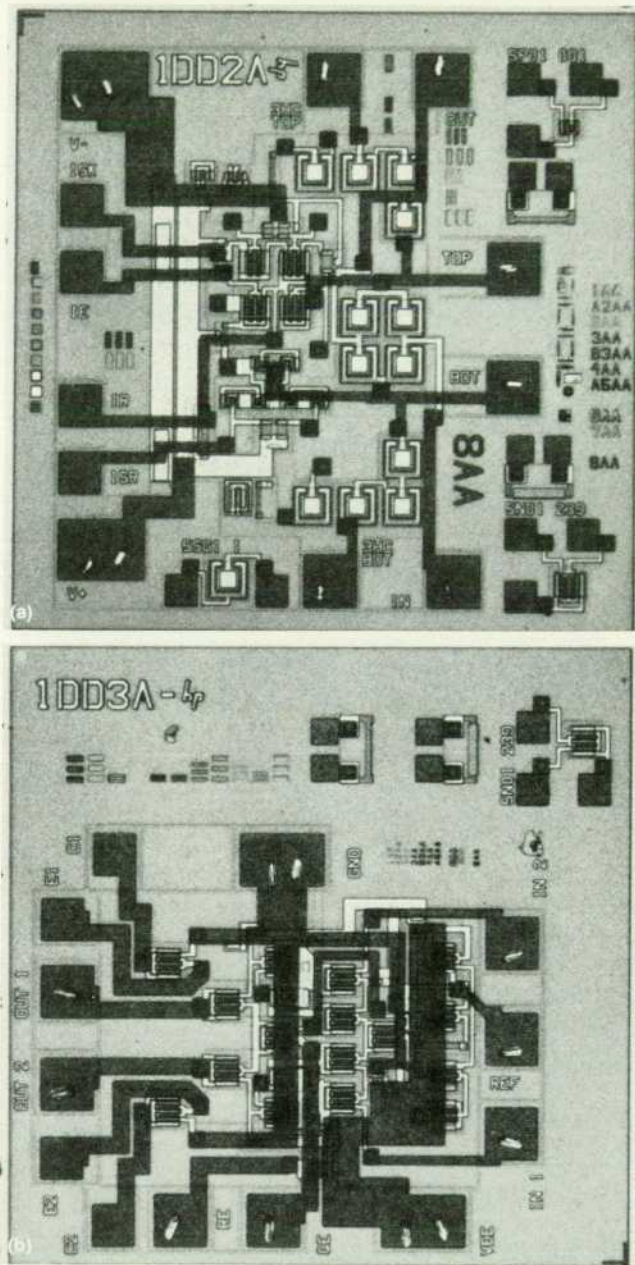


Fig. 4. (a) Sample gate IC contains a monolithic Schottky diode bridge sampling switch. The chip size is 1.3×1.3 mm. (b) Gate driver IC generates a high-current pulse that has subnanosecond transition times. The chip size is 1.28×1.28 mm.

emitter-follower stage with an emitter current of 40 mA, which is needed for driving 50-ohm loads to ± 1 -volt levels with low distortion. Q11 is actually three parallel transistors with emitter resistors to ensure that they share current equally. Diodes D2 and D3 are Schottky diode clamps that minimize the recovery time for input overload situations. Variable capacitor C1 can be adjusted to give critical damping of the amplifier's step response, desirable for high slew rates and rapid settling applications as described later. The chip dissipates nearly 1 watt of power.

For the preamplifier, the op amp is used with 75-ohm




Fig. 5. Three different versions of this high-speed operational amplifier are used in the ADC system. It serves as a preamplifier and a postamplifier in the sample-and-hold circuit and as an error amplifier in the digitizer. The chip size is 0.93×1.21 mm.

emitter resistors in a closed-loop, positive gain-of-one configuration. The drift, gain flatness, and linearity of this amplifier are far superior to the open-loop buffers often found in other S/H designs. The preamplifier typically has a full-power bandwidth of 150 megahertz when driving the 50-ohm load represented by the resistive attenuator to 2-volt peak-to-peak levels. The harmonic distortion components generated with a full-scale input are greater than 66 dB below the fundamental amplitude, equivalent to one-half least-significant bit, at any frequency within the 10-megahertz Nyquist bandwidth of this ADC system.

For the postamplifier, the op amp is used with 50-ohm emitter resistors in a closed-loop, positive gain-of-two configuration. The postamplifier is capable of slewing at 1.0 volt per nanosecond, compatible with the high slew rates associated with the hold capacitor voltage. Within a few nanoseconds after the sampling switch is opened, the postamplifier has settled enough for the digitizer to perform the first-pass conversion. The bipolar op amp includes pnp current mirrors to provide input bias current compensation that reduces the voltage droop of the hold capacitor. The postamplifier's output stage, capable of ± 1 -volt levels, is connected to the digitizer through a short coaxial cable that is terminated in 50 ohms at the digitizer to minimize voltage errors caused by line reflections.

The value of the hold capacitor represents a compromise. A small value lowers the gate driver current requirement and lowers the RC time constant of the sampling switch circuitry for wider bandwidth. A large hold capacitor reduces the feedthrough across the reverse-biased bridge diodes and reduces the droop of the hold capacitor voltage caused by the finite input bias current of the bipolar postamplifier.

Hybrid packaging completes the S/H design. As the picture of the hybrid assembly in Fig. 7 shows, the four integrated circuits are interconnected on a thin-film ceramic substrate to reduce parasitic reactances. These unwanted

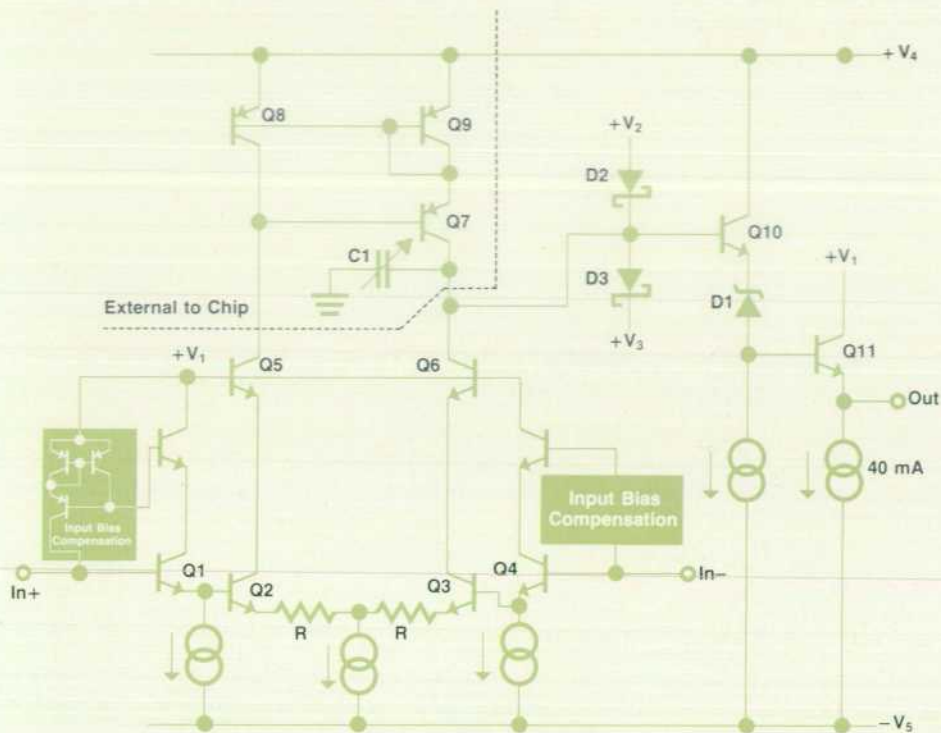


Fig. 6. Simplified operational amplifier schematic diagram.

capacitances and inductances can limit fast, accurate acquisition and contribute to contamination of the hold capacitor voltage. This packaging concept, previously used in the HP 9825A Desktop Computer, is described in more detail later in this article.

The result of the IC and hybrid design efforts is a high-

performance sample-and-hold circuit. The bandwidth, including the input protection, sampling switch, and both amplifiers, is typically 80 megahertz. This factor-of-eight margin over the 10-megahertz Nyquist frequency minimizes analog gain and phase aberrations and allows an input anti-aliasing filter, if needed, to be designed independently of the less precisely controlled S/H transfer function poles. The high sampling efficiency of the S/H is demonstrated by its ability to acquire the positive and negative extremes of the full-scale voltage range ($\pm 1V$) in consecutive sample cycles with full 10-bit resolution (0.1%).

Digitizer Design

The digitizer portion of the analog-to-digital converter system receives the sampled analog signal from the sample-and-hold circuit and encodes it to obtain a digital representation. It does this in a multistep operation as shown in Fig. 8. A coarse estimate of the incoming signal is made by the first-pass encoder, which is accurate to about five significant bits. The digital estimate is then reconverted to analog form by the digital-to-analog converter IC and subtracted from the original signal by an error amplifier. The error amplifier output represents the error inherent in quantizing the signal to only 5 bits. This error is digitized by the second-pass encoder, this time to 6-bit accuracy. The first-pass estimate and the second-pass result are then combined to give the full 10-bit output of the digitizer subsystem.

The performance of the digitizer is made possible by the custom integrated circuits used in the digitizer hybrid. The key circuit is the four-bit quantizer used in the first and second passes. The first pass uses two quantizers for five-bit operation while the second pass uses four quantizers to obtain six-bit resolution. The quantizer is capable of digitiz-




Fig. 7. Sample-and-hold hybrid mounted in its heat sink.

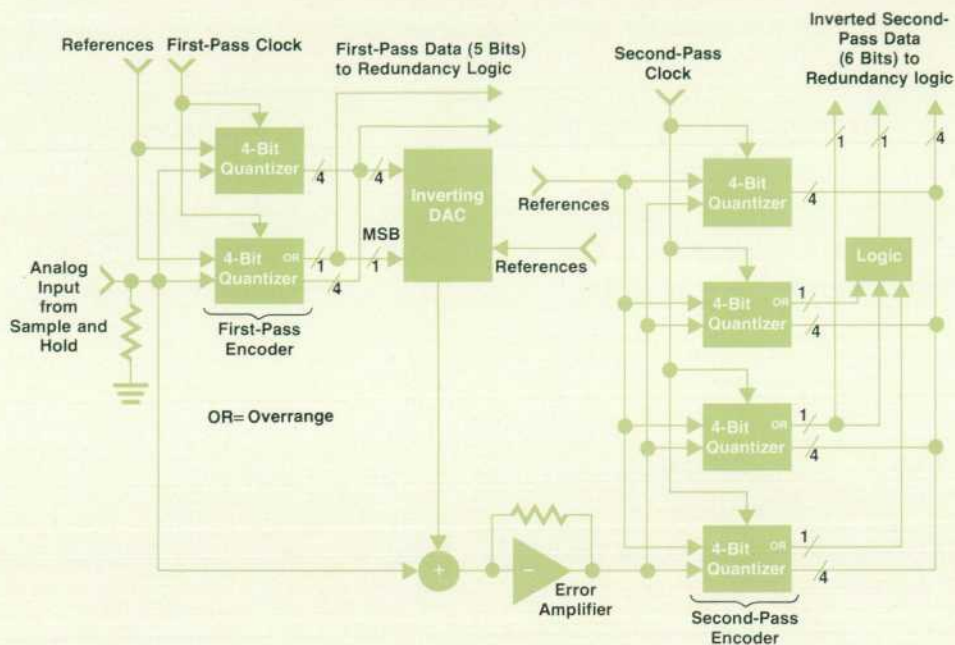


Fig. 8. Digitizer encodes the sample-and-hold circuit output in a two-pass operation.

ing an analog signal into a four-bit word every four nanoseconds.

As shown in Fig. 9, the quantizer input signal is received by a resistor ladder. Since the signal is applied to both ends of the ladder, it appears on each input of the 16 amplifier/buffers. A reference voltage is applied to each end of the other (reference) resistive ladder. This ladder establishes a unique reference voltage at each reference input of the 16 amplifier/buffers. As a result, the bank of amplifier/buffers generates a thermometer code that describes the value of the input voltage compared to the reference voltages. Buffers whose inputs are greater than their reference voltages give

high outputs while buffers whose inputs are less than their reference voltages give low outputs.

The amplifier/buffer is designed to be a high-bandwidth, high-slew-rate circuit. The input is buffered by emitter followers (Fig. 10) for high isolation between individual differential switches and the input signal. The buffering also produces a low input capacitance for the chip, approximately 7.5 pF. The remainder of the design is an emitter-coupled, low-gain amplifier. Overall 3-dB bandwidth is about 1 GHz.

Following the first stage, the latch stage contains circuitry for decoding and latching functions. The initial




Fig. 9. The four-bit quantizers used in the digitizer are custom integrated circuits that can convert an analog signal into a four-bit word every four nanoseconds.

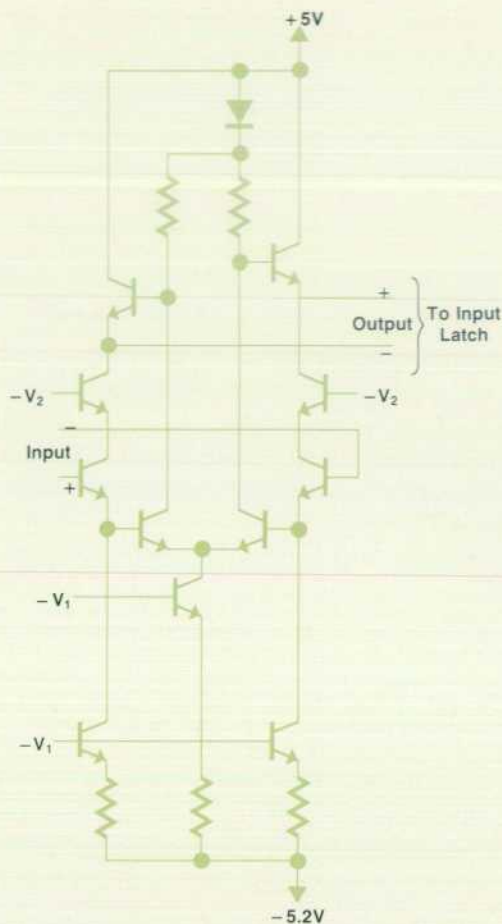


Fig. 10. Quantizer input amplifier/buffer schematic diagram.

thermometer code is converted by an input cross-coupled stage (Fig. 11) to an 8-bit cyclical code with an overrange bit. The cross-coupled stage is an AND gate with one of the inputs inverted. For the three valid input states, this stage behaves like an EXCLUSIVE-OR circuit, from which a cyclical code is derived. Since the cross-coupled input to the latch does a preliminary level of decoding when receiving the signal, only nine latch circuits are required to hold the information. The latch portion of this stage is a conventional ECL (emitter-coupled logic) latch design, with the latch command applied to a current-steering pair of transistors that selects either the input (learn mode) or regenerative (hold mode) transistor pair. The latches use positive feedback to regenerate logic levels from the analog signals at their inputs when the regenerative pair is selected by the latch command. The digital information then advances to the logic stage, which converts the information to four-bit binary code with an overrange bit.

The logic design employs standard ECL configurations to realize EXCLUSIVE-OR and AND functions. An additional circuit combines these functions into a single stage for overrange decoding. Following the logic stage, Zener diodes level-shift the signal to the ground-reference standard for use outside the chip. The data latch that follows is another conventional ECL design. The ECL output is distributed differentially by four emitter followers for

Custom IC Processes

Two custom HP integrated circuit processes played a significant part in developing the 5180A ADC module: A 5-GHz bipolar process, used for all of the special ICs in the ADC, and a 1-GHz bipolar LSI process, used in the input amplifier (see page 18).

The 5-GHz process¹ was developed to minimize τ_F and R_B , raising the f_T of transistors to over 5 GHz at $I_C = 10$ mA. The process features an ion-implanted base and pn-junction isolation.

The 1-GHz LSI process is designed for high density and medium frequency response. Very high yield is achieved on arrays with 300 to 400 gates. This process allows clock rates of 175 MHz on low-power circuits. It has been used extensively for both digital and linear circuits where standard npn and pnp cells as well as I²L (integrated injection logic) cells with 1000 gates are fabricated.

The table below summarizes the characteristics of the two processes.

Parameter	5-GHz Process	1-GHz LSI Process	Units
τ_F	31.8	120	ps
R_B	40.0	100	Ω
C_{JBEO}	0.57	0.57	pF
C_{JBEO}	0.56	0.52	pF
C_{CS}	0.20	0.50	pF
β_F	60.0	100	--
I_S	3.2×10^{-16}	6.2×10^{-16}	A
Metal	Moly Gold	Aluminum	--
Line Center-to-Center Space	5	12	μm
Emitter Length	40×2	40×1	μm

Reference

1. D.M. DiPietro, "A New 5-GHz Process for High-Speed Digital Circuits," IEEE Digest, 1975 International Solid-State Circuits Conference.

maximum versatility.

Special design features of the quantizer enhance its utility. One is input bias current compensation. A pnp current mirror is employed on both the signal and reference inputs. Each mirror detects the amount of base current flowing into all of the amplifier/buffer stages and adds that much current to the respective input. Another error-canceling feature is the use of dual resistive ladders. The resistive ladder on the reference side provides the linearly weighted reference voltages used to define the code bins. The side effect of using an equal-weight ladder for this function is a bowed error created by base current supplied to the input by the ladder. By using an identical ladder on the signal side, the error is made common-mode and is rejected by the amplifier/buffers.

The use of 16 amplifier/buffers makes parallel configuration of quantizers possible by a coding scheme that preserves the top transition in the form of an overrange bit. When the signal is above the range of the quantizer, this fifth bit is on and all the other bits are zero. This five-bit output simplifies decoding when more than one quantizer is used. Since the four low-order bits are zero when the signal is out of range of a quantizer, simple wired-OR decoding is used for these four bits. For the five-bit first pass, the

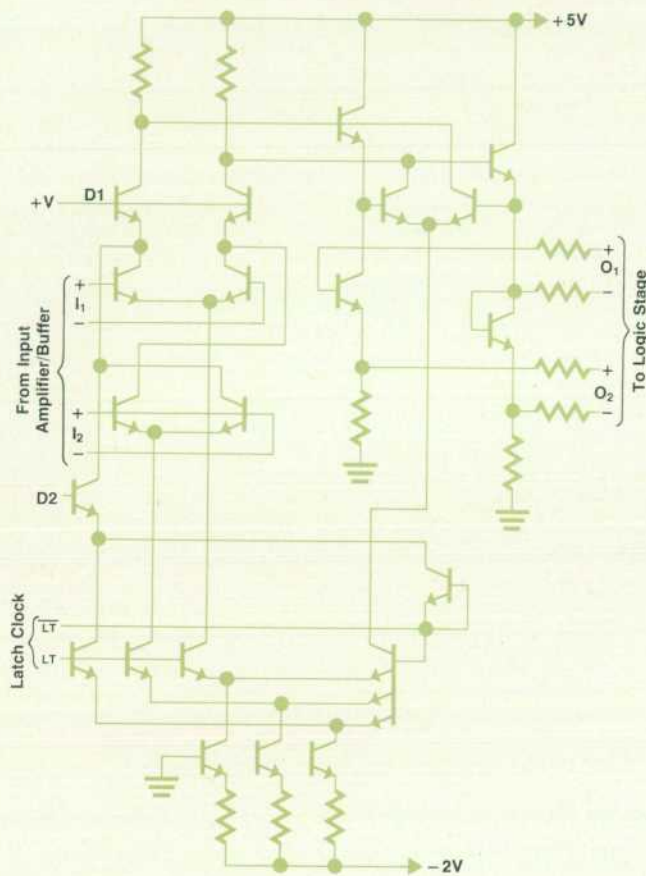


Fig. 11. Quantizer input latch schematic diagram.

overrange bit of the lower quantizer becomes the most-significant bit. For the six-bit second pass, the central overrange bit becomes the MSB, and the second MSB is simply decoded from the overrange bits for the three lower quantizers. The quantizer contains 600 transistors and dissipates a maximum of 1.8 watts on a 2.98-mm-by-3.05-mm chip (Fig. 12).

Another key integrated circuit in the digitizer is the 10-bit digital-to-analog converter. The DAC converts five bits from the first pass to an analog signal with 10-bit accuracy. It uses differentially buffered transistor switch cells to route a fixed current into an R-2R ladder (Fig. 13). The digital portion of a cell receives a bit of information, converting it to a low-level differential signal. This differential signal drives the analog current-routing portion of the cell in a high-speed, low-noise fashion. The R-2R ladder combines the currents from all the cells in a binary fashion to form the analog voltage. The cell switches and R-2R ladder are all integrated to reduce parasitic reactances, allowing settling to 10-bit accuracy in less than 15 ns. Trimming for dc errors is done by external adjustment of the fixed currents. The DAC dissipates one-half watt on a 2.72-mm-by-1.65-mm chip (Fig. 14).

The signal from the DAC is subtracted from the sample-and-hold output by the error amplifier. The error amplifier




Fig. 12. Quantizer chip contains 600 transistors and measures 2.98 x 3.05 mm.

is an operational amplifier identical in design to that used in the sample-and-hold. For use in the digitizer, a version with 10-ohm emitter resistors is used; this is appropriate to the large closed-loop gain required of this amplifier. Setting the gain is an integrated network of silicon resistors fabricated on the same wafer as the DAC to provide good matching over temperature and time. The signal is amplified by these components to cover one-half the second-pass range.

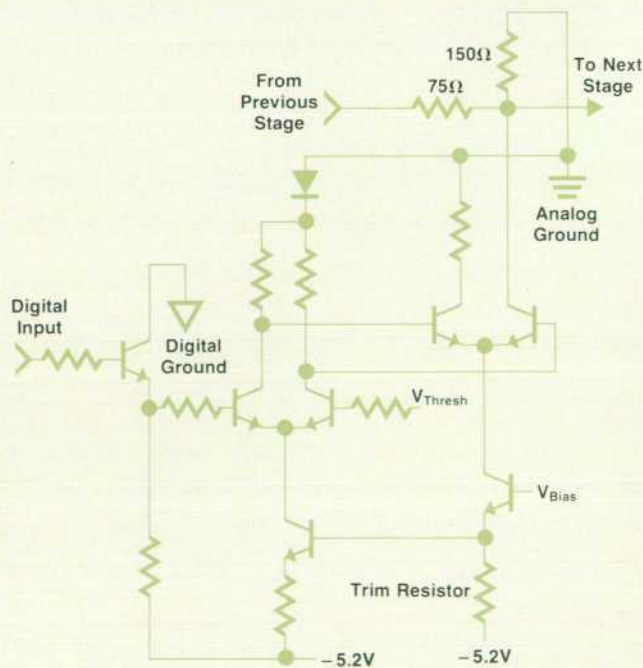


Fig. 13. The digital-to-analog converter (DAC) used in the parallel-ripple ADC consists of ten of these cells. Only five inputs are used in the 5180A.




Fig. 14. DAC chip measures 2.72 x 1.65 mm.

Error Analysis

To achieve 10-bit, 20-megasample/second operation, careful allocation of timing and error margins for each subsystem was done. The time budget for the digitizer consists of 40 ns of hold time. The first 5 ns is allocated for the sample-and-hold to settle to the six-bit level and for concurrent acquisition of the signal by the first pass. The next 35 ns is allotted for the DAC and error amplifier to receive the signal and settle. Of this 35 ns, the final 4 ns is concurrent with acquisition of the signal by the second pass.

The error budget for the digitizer revolves around the use of an extra bit of resolution in the second pass. It is this redundancy bit that allows the error signal to be scaled for one-half the total second-pass range (the middle half). The




Fig. 15. The error signal or residue from the first pass is scaled and offset to fall in the middle half of the second-pass range. Thus the first-pass LSB overlaps the second pass MSB. If the first pass makes an error, the overlap of the second pass detects and corrects the error.

other half (the two outer quarters) of the range is used for error correction.

To make the scaled error signal from the first pass fall in the middle half of the second-pass range, the error signal is offset by one-half of a first-pass LSB. This means that the first-pass LSB overlaps the second-pass MSB. The overlapping makes the final output code somewhat complicated to extract. As shown in Fig. 15, when the code falls in the overlap region, it is equivalent to a point 32 LSB into the main region with the first-pass code in an adjacent bin. Mathematically, this means that the result should be:

$$\text{code} = 32F + S - 16,$$

where F is the first-pass code and S is the second-pass code. In this digitizer, the mathematics are further complicated by an inversion done by the error amplifier; the resultant second-pass code is the complement of the desired second-pass code. This is handled by altering the formula:

$$\text{code} = 32F - S + 47.$$

This code computation is done on another board, the redundancy board (see Fig. 1).

A pair of examples may clarify the error correction scheme.

Example 1: Correct code = 0110011010

	1st-pass error	No error
1st-pass code	01101	01100
2nd-pass code	001010	101010
Computation	01101 + 001010 ----- 0110101010	01100 + 101010 ----- 0110101010
Result	- 1 ----- 0110011010	- 1 ----- 0110011010

Example 2: Correct code = 0110000111

	1st-pass error	No error
1st-pass code	01011	01100
2nd-pass code	110111	010111
Computation	01011 + 110111 ----- 0110010111	01100 + 010111 ----- 0110010111
Result	- 1 ----- 0110000111	- 1 ----- 0110000111

The errors that the redundancy range corrects are the incomplete settling of the sample-and-hold signal at the time the first pass samples it and the nonideal digitizing of the signal by the first pass. The assumption is made that the value ultimately digitized by the second pass is correct, so that differences between the first- and second-pass values are used to correct first-pass errors.

The error attributable to the first pass allows one-quarter-bit errors in any given threshold, which is equivalent to 8 LSB in the second pass. This is within normal manufacturing tolerance of the quantizers. The time budget of 5 ns allows over twice the time normally needed for

A 40-MHz Input Amplifier

Any analog-to-digital converter system has limited usefulness without appropriate signal conditioning circuitry at the input to perform the functions of scaling and offsetting the input signals. The input amplifier of the 5180A performs these functions under HP-IB control. There are two 1-megohm input channels, which can be used individually or chopped. For users who wish to supply their own signal conditioning, an auxiliary channel provides a direct 50-ohm connection to the ADC module's sample-and-hold circuit.

The heart of the input amplifier is a thin-film hybrid (Fig. 1). On this hybrid are two preamplifier ICs, a level-shifting chip, and an operational amplifier chip, all designed and built by Hewlett-Packard. The preamp ICs are built using a 2-GHz f_T bipolar process.¹ The op amp IC is described in the accompanying article, and the level-shifting IC is the first circuit in production using HP's MAGIC chip.*

*MAGIC = multiple access generalized interconnect circuit.




Fig. 1. 5180A Waveform Recorder input amplifier hybrid circuit includes four HP integrated circuits and 27 thin-film resistors.

acquisition by a quantizer.

The error budget calls for full 10-bit accuracy of the DAC and error amplifier combination. This depends on the time allotted for the pair to settle. The time budget of 35 ns makes the amplifier the limiting speed component in the system. Since the sample-and-hold signal reaches the error amplifier before the DAC signal, the amplifier is briefly driven out of the active range. Schottky clamps in the error amplifier keep it out of saturation. Recovery is fast enough for the closed-loop bandwidth of 40 MHz to meet settling requirements in the allotted time. The second pass uses quantizers at their resolution limit: peak bin size deviations up to 1/3 LSB may occur. As with the first pass, the time requirement of 4 ns is easily met by the quantizers.

To summarize, the digitizer's performance is the result of four different custom IC designs used in nine places on a

Built using a 1-GHz bipolar LSI process, the MAGIC chip contains an array of transistors and resistors. The designer defines the function of the circuit by specifying a custom second metal mask. This approach allows fast development time at low cost.

Signal voltages presented to one of the two high-impedance inputs pass through a discrete attenuator and JFET impedance converter assembly which drives the hybrid (see Fig. 2). On the hybrid, signals proceed to the preamp IC where the voltages are attenuated and the signal information is converted to differential currents. These differential currents are summed with other differential currents containing the offset information. These combined currents are converted to a differential voltage that is presented to the MAGIC chip, which level-shifts the voltage and buffers the preamp stage from the output stage. The op amp output stage converts the differential voltage to a single-ended voltage. The hybrid output drives the 50-ohm input to the sample-and-hold circuit.

The hybrid also contains 27 thin-film resistors. Because the gain of the circuit is determined by the ratio of several of these resistors, trimming some of the resistor values is a must. An HP 9825A Desktop-Computer-based automatic trimming system² heats each resistor individually by running a high current through it and monitors the change in resistance value resulting from the oxidation of resistor material. Resistors are trimmed to 0.25% tolerance with low values of drift with time.

The 5180A input amplifier has a typical bandwidth of 60 MHz. Coupled with the ADC system this results in an amplifier bandwidth greater than 40 MHz for a flat frequency response past the Nyquist frequency. Each channel can be ac coupled. Each channel input can be internally grounded or connected to a precision 100-mV reference for precise operational verification of the amplifier/ADC system.

Reference

1. J.K. Millard, "An Oscilloscope Vertical-Channel Amplifier that Combines Monolithic, Thick-Film Hybrid, and Discrete Technologies," *Hewlett-Packard Journal*, December 1975.
2. C. Worley, "Precision Current Trimming of Thin-Film Resistors," *Hewlett-Packard Hybrid Packaging Conference*, June 22, 1981.

-Pat Deane

thin-film hybrid microcircuit (Fig. 16). The output is provided on ECL lines for high-speed transmission of information with a minimum of noise.

Design Considerations

The nature and complexity of the sample-and-hold and digitizer circuitry dictated hybrid circuits. Pulse rise times less than 1 ns and a sample-conversion process (sample-and-hold, first pass, DAC settling, second pass) that must take place in 50 ns fit well with hybrid geometries and their microwave transmission line structures. Line widths and lengths are controlled so the timing is repeatable and well defined.

The high density attainable with hybrid techniques has two other important features. By placing all components close together, parasitic elements can be reduced, ground

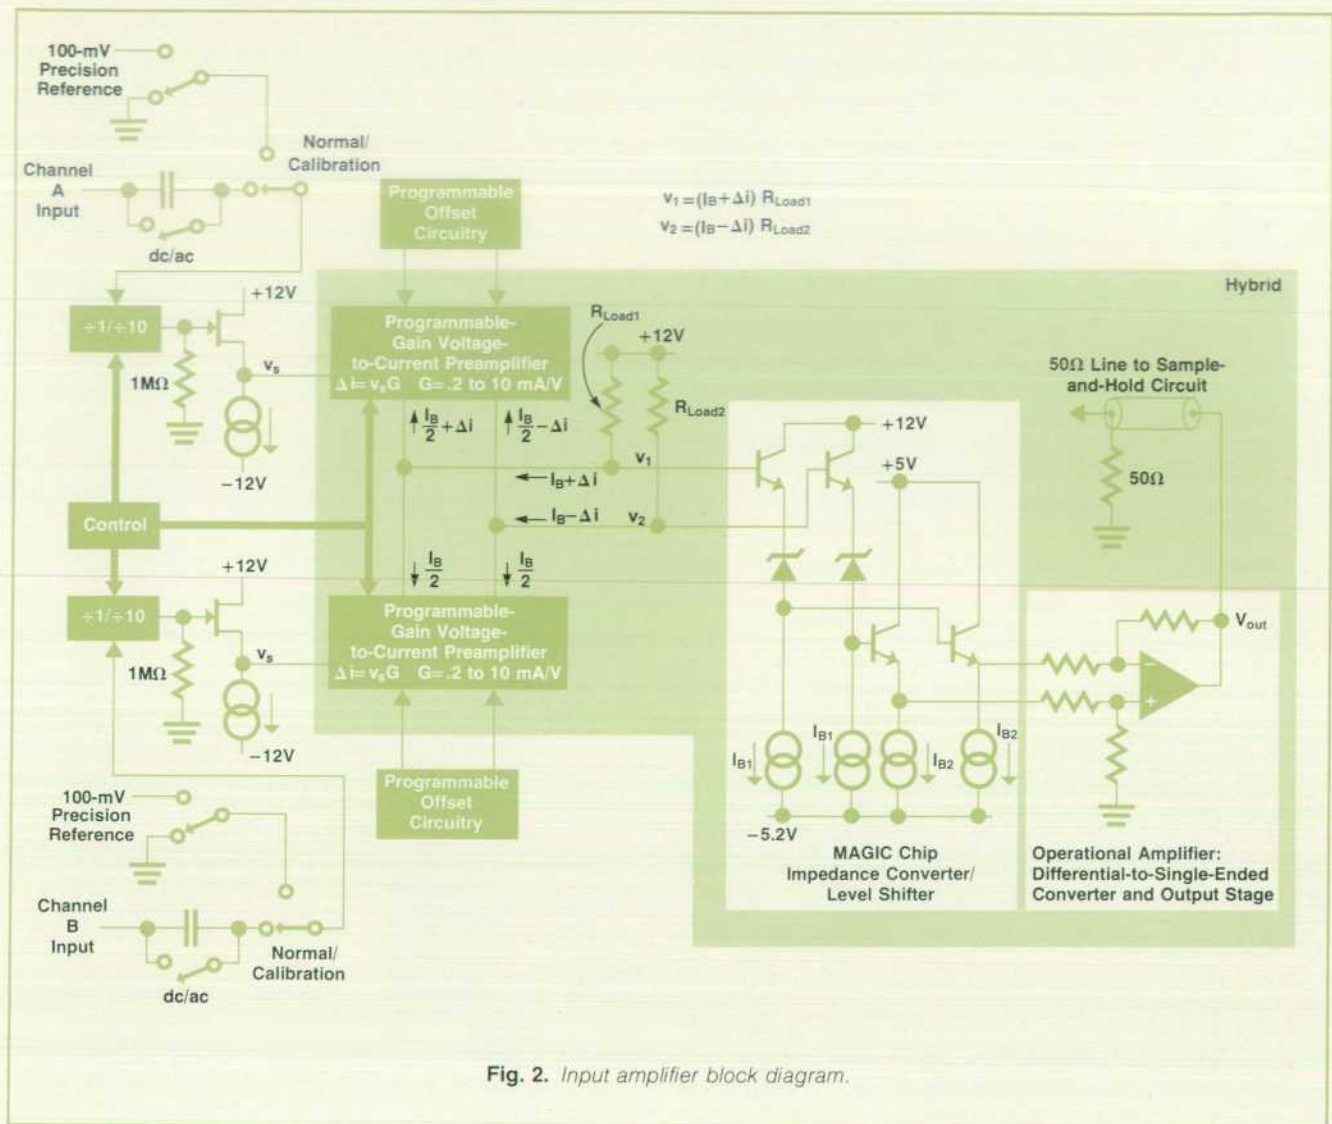


Fig. 2. Input amplifier block diagram.

loops can be minimized and references can be bypassed to the "same" node. A quantization resolution of 2 mV places very stringent requirements on the cleanness of the analog ground references and power supply. Digital switching currents must be isolated and bypassed so references do not change. A thermal environment that allows the DAC and feedback resistor to track each other is also important, and close proximity with a minimum of packaging interfaces is necessary.

Twelve watts of power must be removed from the digitizer hybrid with a minimum temperature rise at the IC junctions. The hybrid is mounted in a finned aluminum heat sink. Heat sink grease is applied in dots just under the ICs. This technique minimizes the effects of substrate camber and results in the lowest thermal impedance. Infrared microscope measurements on this structure have confirmed a thermal impedance of 15°C/watt. This indicates that the hottest transistor junction of a quantizer IC is only 20°C above the ambient measured at the back surface of the heat sink. When the instrument is at maximum rated temperature, 55°C, the transistor junction is at 105°C, which is excellent for reliability considerations.

An automatic gold thermosonic bonder is the assembly technique for fabricating the 275 wire interconnects in the digitizer circuit. This bonder was chosen more for repeatability, quality, and correct bond placement than cost of circuit construction.

The digitizer and sample-and-hold hybrids are exactly the same size and mounted in similar ways on their respective boards. A 1.65-mm-diameter tube of silicon rubber covered by a Kapton™ film with etched conductors 0.076 mm wide and 0.102-mm spaces is sandwiched between mirrored pads on the board and substrate. This gives mechanical alignment, current-handling capability for the high dc bias currents, and a good RF match.

The sample-and-hold and digitizer printed circuit boards are part of the ADC module, which is a stand-alone five-board set that requires only dc power and a single-phase clock to perform the analog-to-digital conversion. The other three boards in the module are timing, redundancy, and voltage reference boards.

The timing board generates six different one- or two-phase clock signals. Three of these, the sample-and-hold, first-pass, and second-pass clocks, require a very precise




Fig. 16. Digitizer hybrid.

time relationship. The time spacings are obtained by measured coaxial delay lines and high-speed logic ICs made in a 5-GHz (f_T) IC process. These gates are used to generate the required fast rise times and to minimize the jitter that would cause instrument errors.

The redundancy board takes the eleven data bits and two extra carry bits from the digitizer and performs the mathematical operation already described to arrive at the 10-bit answer.

The voltage reference board is a series of op amp voltage sources followed by emitter followers for current output capability. This board excludes noise generated by the rest of the instrument system and generates precision references required by the sample-and-hold and digitizer hybrids.

Acknowledgments

The authors wish to acknowledge the support of Dexter Hartke, under whose lab management the project was begun, and Steve Schultz, who developed the first prototype digitizer, including the initial design of several of the integrated circuits. Credit goes to Frank Chu and Ron Young for helping to refine the ADC so that it can be easily produced. This project could not have been successful without the

cooperation of Santa Clara Division's Integrated Circuit Operation and Hybrid Department for developing the processes and devices necessary. Special thanks must go to Dave DiPietro for his expert IC consultation and to Kathi Luiz for her patience in assembling the endless number of "impossible" experimental hybrids throughout the development.

Arthur S. Muto



Art Muto joined HP in 1970 after receiving his BS and MS degrees in electrical engineering and computer science from the University of California at Berkeley. He developed the 500-MHz main gate for the 5345A Counter and was project leader for the 5353A Plug-in. For the 5180A Waveform Recorder, he was project leader for development of the prototype ADC module. He's now project manager for high-performance universal counters. He's co-authored several magazine articles on counters and ADCs. Born in Sacramento, California, Art is married, has a daughter, and lives in Cupertino. He likes to tinker with high-fidelity components and take nature walks, and he enjoys live music, motion pictures, and gourmet food.

Robert C. Rehner, Jr.




Bob Rehner received BS and MS degrees in electrical engineering from the University of Michigan in 1965 and 1966. With HP since 1966, he's served as design engineer, production engineer, production engineering manager, project leader, and project manager, mostly for high-frequency hybrid microelectronics and handheld calculators. He designed the hybrid for the HP01 Watch/Calculator. He's a member of the IEEE and the International Society for Hybrid Microelectronics and has taught microwave measurements at a local community college. Bob was born in Anniston, Alabama and now lives in Cupertino, California. He has two sons, is active in the YMCA, and enjoys tennis, skiing, sailing, hiking, and running.

Bruce E. Peetz



Bruce Peetz helped develop the ADC for the 5180A Waveform Recorder and is now a project leader with HP's Santa Clara Division. A 1973 BSEE graduate of Massachusetts Institute of Technology, he joined HP in 1977 with experience in radar system design. He's a member of the IEEE and serves on the IEEE waveform analysis standards committee. Bruce was born in Santa Monica, California and now lives in San Jose. He's married, has a daughter, plays guitar and piano, and enjoys woodworking.



Measuring Waveform Recorder Performance

by Bruce E. Peetz, Arthur S. Muto, and J. Martin Neil

THE KEY TO CONFIDENCE in the quality of a waveform recorder is assurance that the analog-to-digital converter (ADC) encodes the signal without degrading it. Dynamic tests that cover the frequency range over which the converter is expected to operate can provide that assurance. The results of the dynamic tests give the user a model of resolution versus frequency for the recorder. More elaborate models of failure mechanisms can be obtained by varying the conditions of the tests.

All of the dynamic tests used for the 5180A Waveform Recorder use sine waves as stimulus. Sine waves were chosen primarily because they are the easiest to generate in practice at the frequencies of interest with adequate fidelity. While it may be possible to generate a square wave, for example, whose function is known to the 10-bit resolution of the 5180A, no square wave generators exist that can guarantee the same waveshape to 10-bit resolution at 10 MHz from unit to unit. Another motivation for choosing a sine wave stimulus is the simple mathematical model a sine function provides for analysis. This benefit greatly simplifies the algorithms used for data analysis.

Four dynamic tests for waveform recorder characterizations are presented here: beat frequency testing,¹ histogram analysis,² sine wave curve fitting,^{3,4} and discrete finite Fourier transform.⁵ The last three tests operate in the same way. A sine wave source is supplied to the waveform recorder and one or more records of data are taken. A computer is then used to analyze the data. The tests differ primarily in the analysis algorithms and consequently in the sort of errors brought to light. Critical to the success of these tests is the purity of the sine wave source. Synthesized sources are necessary to provide the short-term and long-term stability required by the dynamic range of the ADC. Passive filters (a six-pole elliptical filter is used for 5180A tests) are required to eliminate harmonic distortion from the source.

These tests provide the most stressful conditions for the ADC with the input signal amplitude at full scale. Generally

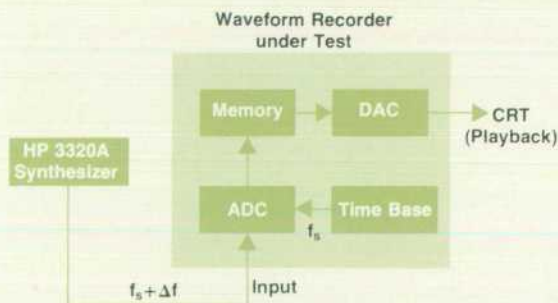


Fig. 1. Beat frequency test setup.



Fig. 2. When the input frequency is close to the sample rate f_s , the encoded result is aliased to the difference or beat frequency Δf .

speaking, nonlinear effects increase more quickly than the signal level increases because of the nonideal large-signal dc behavior of the ADC components and the higher slew rates large amplitudes imply.

Beat Frequency Testing

The beat frequency and envelope tests are qualitative tests that provide a quick, simple visual demonstration of ADC dynamic failures. An input frequency is selected that provides worst-case range changes and maximal input slew rates that the ADC is expected to see in use. The output is then viewed on a display in real time.

The name "beat frequency" describes the reasoning behind the test. The input sinusoid is chosen to be a multiple of the sample frequency plus a small incremental frequency (Fig. 1). Successive samples of the input waveform step slowly through the sine wave as a function of the small difference or beat frequency (Fig. 2). Ideally, the multiplicative properties of sampling would yield a sine wave of the beat frequency displayed on the waveform recorder's CRT. Errors can be seen as deviations from a smooth sine function. Missing codes, for example, appear as local discontinuities in the sine wave. The oversize codes that accompany missing codes are seen as widening in the individual codes appearing on the sine wave.

By choosing an arbitrarily low beat frequency, a slow accurate DAC may be used for viewing the test output. For best results, the upper limit on the beat frequency choice is set by the speed with which the beat frequency walks through the codes. It is desirable to have one or more successive samples at each code. This alleviates the settling




Fig. 3. When the input frequency is near one-half the sample rate, the envelope of the difference frequency results.




Fig. 4. A beat frequency display produced by the 5180A Waveform Recorder with a 10.0031-MHz input frequency and a 10-MHz sample rate. The smooth sine wave indicates freedom from dynamic errors.

constraint on the DAC and ensures that the display covers all possible code failures. For a 20-MHz sample rate and a 10-bit ADC, this implies a 3-kHz maximum beat frequency for a minimum of one sample per code bin.

Although the usual input frequency for a beat frequency test is near the sample rate, the analog bandwidth of the ADC may be measured by setting the carrier to a number of different multiples of the sample rate. The band limit is observed as a rolloff in amplitude as the carrier frequency is increased.

The envelope test differs from the beat frequency test in the choice of input frequency that the ADC encodes. Instead of a multiple of the sample frequency, an input frequency near one-half the sample rate is used. Now the ideal output is two out-of-phase sinewaves at the beat frequency (Fig. 3). This means that successive samples can be at the extreme ends of the ADC range, which is useful for examining slew problems that might not appear when successive samples are at adjacent codes. To avoid placing the same stress on the DAC used for display, a bank of D flip-flops removes every other sample before the data arrives at the DAC. Thus only one phase of the beat frequency remains.




Fig. 5. A beat frequency display for a commercially available 10-MHz, 8-bit ADC with a 10.0031-MHz input.




Fig. 6. Setup for histogram test.

Fig. 4 shows 5180A beat frequency test results for a 10.0031-MHz input sine wave sampled at 10 MHz. For comparison, Fig. 5 shows a 10.0031-MHz sine wave being sampled at 10 MHz by a commercially available 8-bit, 20-MHz ADC.

Histogram Testing

A sine-wave-based histogram test provides both a localized error description and some global description of the A-to-D converter. Using the histogram test, it is possible to obtain the differential nonlinearity of the ADC, to see whether any missing codes exist at the test frequency, and to get a measure of gain and offset at the test frequency. Of the sine-wave-based tests presented here, the histogram test yields the best information about individual code bin size at an arbitrary frequency.

A statistically significant number of samples of the input sinusoid are taken and stored as a record (Fig. 6). The frequency of code occurrence in the record is then plotted as a function of code. For an ideal ADC, the shape of the plot would be the probability density function (PDF) of a sine wave (Fig. 7) provided that the input and sample frequencies are relatively independent. The PDF of a sine wave is given by:

$$p(V) = \frac{1}{\pi \sqrt{A^2 - V^2}}$$

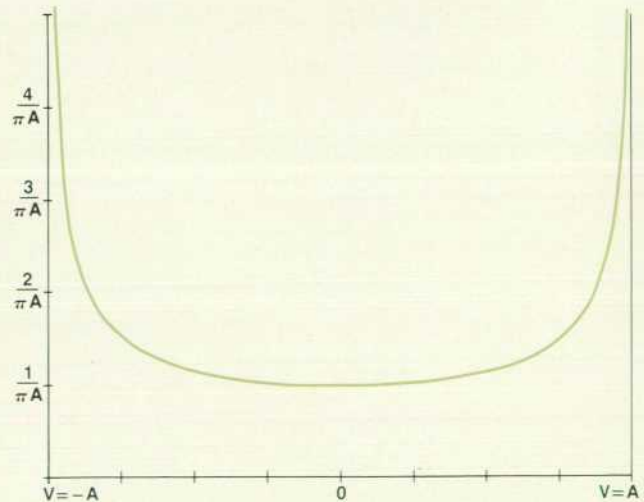


Fig. 7. Sine wave probability density function.

where A is the sine wave amplitude and V is the independent variable (voltage). For a real ADC, fewer than the expected number of occurrences for a given code bin indicates that the effective code bin width is smaller than ideal at the input frequency.* No occurrences indicates that the code bin width is zero for that input. A greater-than-expected number of occurrences implies a larger-than-ideal code bin width.

What is a statistically significant number of samples? We can determine significance from probability theory. For a given input PDF and record size, each bin of an ideal ADC has an expected number of occurrences and a standard deviation around that expectation. The confidence that the number of occurrences is close to the expectation is equal to the probability that the occurrences fall within the appropriate number of standard deviations. The ratio of the standard deviation to the expectation (and thus the error for a given confidence) decreases with more samples. To get the confidence for the entire range, the probabilities for all codes lying within the desired error are multiplied together.

For an ideal 10-bit ADC, 100,000 samples would give us a 12% confidence that the peak deviation from the input PDF is less than 0.3 least-significant bit (LSB) and a 99.9% confidence that the peak deviation is less than 0.5 LSB. The notion of confidence relies on the input's being a random process. We can model the sine wave input as random process only if the input and sample frequencies are relatively independent.

The specification of greatest interest that can be calculated using the histogram test is differential nonlinearity. Differential nonlinearity is a measure of how each code bin varies in size with respect to the ideal:

$$\text{Differential nonlinearity} = \frac{\text{actual } P(n^{\text{th}} \text{ code})}{\text{ideal } P(n^{\text{th}} \text{ code})} - 1$$

*Histogram testing can be thought of as a process of sampling and digitizing the input signal and sorting the digitized samples into bins. Each bin represents a single output code and collects samples whose values fall in a specific range. The number of occurrences or samples collected in each bin varies according to the input signal. If N is the number of ADC bits, there are 2^N bins. Ideally, if B is the full-scale range of the ADC in volts, each bin corresponds to a range of sample sizes covering $B/2^N$ volts. In a real ADC, the bins may not all have the same width.

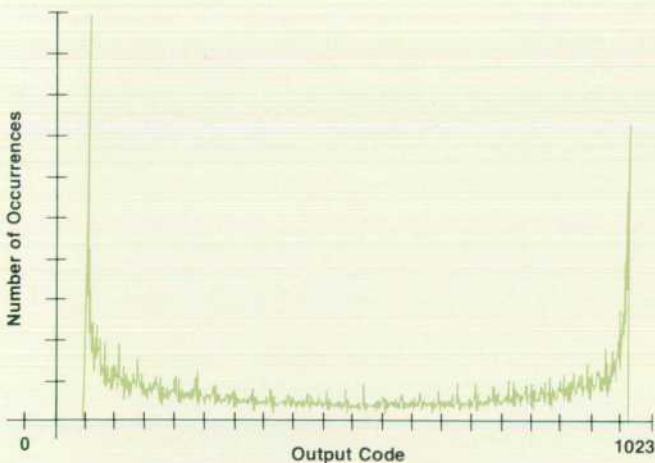


Fig. 8. A 100,000-sample histogram for a 5180A with a 9.85-MHz sine wave input. All discontinuities are less than one least-significant bit (LSB).

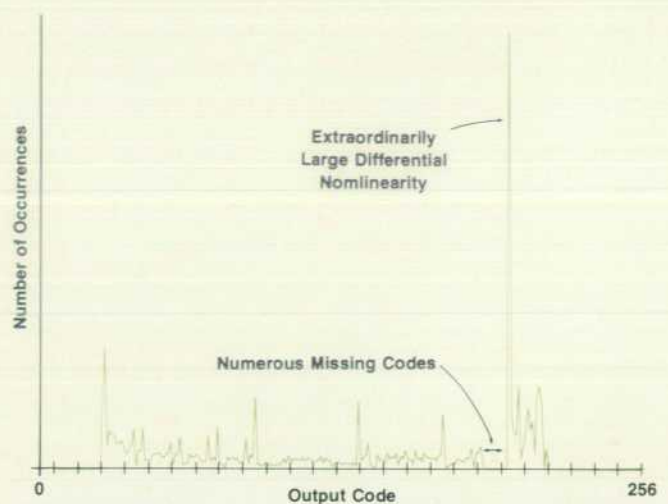


Fig. 9. A 100,000-sample histogram plot for a commercially available 20-MHz, 8-bit ADC with a 9.85-MHz input. Large differential nonlinearities and numerous missed codes are apparent.

where actual $P(n^{\text{th}} \text{ code})$ is the measured probability of occurrence for code bin n , and ideal $P(n^{\text{th}} \text{ code})$ is the ideal probability of occurrence for code bin n . The code bin number n goes from 1 to 2^N , where N is the number of ADC bits. Using the probability of occurrence eliminates dependence on the number of samples taken. To calculate the probability for each code in the actual data record, the number of occurrences for each code is divided by the number of samples in the record. The ideal probability of occurrence is what an ideal ADC would generate with a sine wave input. For each code bin, this is the integral of the probability density function of a sine wave over the bin:

$$P(n) = \frac{1}{\pi} \left(\sin^{-1} \left(\frac{B(n-2^{N-1})}{A2^N} \right) - \sin^{-1} \left(\frac{B(n-1-2^{N-1})}{A2^N} \right) \right)$$

where n is the code bin number, B is the full-scale range of the ADC, and N is the number of ADC bits. To avoid large differences in code probability caused by the sinusoid cusp, a sine wave amplitude A is chosen that slightly overdrives the ADC.

A judicious choice of frequency for the input sinusoid in this test is necessary for realistic test results. An input frequency that is a submultiple of the sample frequency violates the relative independence criterion and will result in sampling of the same few codes each input cycle. Using an input frequency that has a large common divisor with the sample frequency generates similar problems since the codes repeat after each cycle of the divisor frequency. Ideally the period of the greatest common divisor should be as long as the record length.

A 5180A histogram is shown in Fig. 8 for an input sine wave at 9.85 MHz. For comparison, Fig. 9 shows data from a commercially available, 8-bit 20-MHz ADC for an input sine wave at 9.85 MHz, while Fig. 10 shows data from an 8-bit, 100-MHz ADC taken at 9.85 MHz.

Curve Fitting

The curve-fit test is a global description of the ADC. This




Fig. 10. A 100,000-sample histogram plot for a 100-MHz, 8-bit ADC with a 9.85-MHz input sampled at 20 MHz. Extremely large differential nonlinearities and numerous missed codes are apparent.

means that the errors measured by the test are averaged to give a general measurement of the ADC transfer function. The result of this test is a figure of merit called the number of effective bits for the ADC. The effective bit number is a general measure of how much an ADC's nonlinearity has impaired its usefulness at a given frequency.

The number of effective bits is obtained by analyzing a record of data taken from a sine wave source (Fig. 11). The analysis consists of generating a sine wave in software that is a best fit to the data record. Any difference between the data record and the best-fit sine wave is assumed to be error (Fig. 12). The standard deviation of the error thus calculated is compared to the error an ideal ADC of the same number of bits might generate. If the error exceeds the ideal the number of effective bits exhibited by the ADC is less than the number of bits it digitizes. Errors that cause degradation in this test are nonlinear effects such as harmonic distortion, noise, and aperture uncertainty. Gain, offset, and phase errors do not affect the results since they are ignored by the curve-fit process.

The number of effective bits is computed using expressions for average error as follows:

$$\text{Effective bits} = N - \log_2 \left(\frac{\text{actual rms error}}{\text{ideal rms error}} \right)$$

where N is the number of ADC bits. The ideal rms error is not actually computed for the input waveform, but is assumed to be the quantization noise exhibited by an ideal ADC with a uniform-probability-density (UPD) input such as a perfect triangle wave. The ideal error is found from the expectation of squared error for a rectangular distribution. A rectangular distribution is used since that represents a UPD taken over an ideal code bin. The result thus obtained is:

$$\text{Ideal rms error} = \frac{Q}{\sqrt{12}}$$

where Q is the ideal code bin width. Although the input

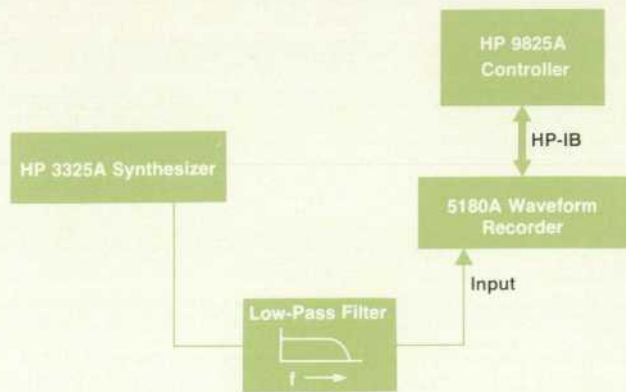


Fig. 11. Setup for the curve-fit test and the discrete finite Fourier transform (DFT) test.

sine wave is not a UPD function, the UPD assumption is still valid since it is locally applied over each code bin. The deviation from a UPD over each code bin is very small, so the errors in using sine waves to approximate UPD inputs are negligible.

The actual rms error is simply the square root of the sum of the squared errors of the data points from the fitted sine wave. The actual rms error is given by:

$$E = \sqrt{\sum_{k=1}^m [x_k - A \cos(\omega t_k + P) - C]^2} \quad (1)$$

where E is the actual rms error, x_k and t_k are the data points, m is the number of data points in the record, and the fitted sine wave parameters are amplitude A , frequency ω , phase P , and offset C .

Equation 1 is also used to find the best-fit sine wave by minimizing the error E . The error is minimized by adjusting the fit parameters: frequency, phase, gain, and offset. This is done by taking the partial derivative of E in equation 1 with respect to each of the four parameters. The error minimum occurs when all of the derivatives are equal to zero. This gives the four simultaneous equations:

$$\sum_{k=1}^m x_k \cos(\omega t_k + P) = A \sum_{k=1}^m \cos^2(\omega t_k + P) + C \sum_{k=1}^m \cos(\omega t_k + P) \quad (2)$$

$$\sum_{k=1}^m x_k = A \sum_{k=1}^m \cos(\omega t_k + P) + nC \quad (3)$$

$$\sum_{k=1}^m x_k t_k \sin(\omega t_k + P) = A \sum_{k=1}^m t_k \cos(\omega t_k + P) \sin(\omega t_k + P) + C \sum_{k=1}^m t_k \sin(\omega t_k + P) \quad (4)$$

$$\sum_{k=1}^m x_k \sin(\omega t_k + P) = A \sum_{k=1}^m \cos(\omega t_k + P) \sin(\omega t_k + P) + C \sum_{k=1}^m \sin(\omega t_k + P) \quad (5)$$

Equations 2 and 3 result from gain and offset adjustments. These are substituted into the other two equations, 4 and 5,




Fig. 12. The first 20 points of the curve-fit data record and the error residue from a fitted sine wave.

giving two nonlinear equations:

$$\frac{\sum_{k=1}^m (x_k - \bar{x}) t_k \sin(\omega t_k + P)}{\sum_{k=1}^m (x_k - \bar{x}) \cos(\omega t_k + P)} = \frac{\sum_{k=1}^m [\cos(\omega t_k + P) - \bar{a}] t_k \sin(\omega t_k + P)}{\sum_{k=1}^m [\cos(\omega t_k + P) - \bar{a}] \cos(\omega t_k + P)} \quad (6)$$

$$\frac{\sum_{k=1}^m (x_k - \bar{x}) \sin(\omega t_k + P)}{\sum_{k=1}^m (x_k - \bar{x}) \cos(\omega t_k + P)} = \frac{\sum_{k=1}^m [\cos(\omega t_k + P) - \bar{a}] \sin(\omega t_k + P)}{\sum_{k=1}^m [\cos(\omega t_k + P) - \bar{a}] \cos(\omega t_k + P)} \quad (7)$$

$$\text{where } \bar{a} = \frac{1}{m} \sum_{k=1}^m \cos(\omega t_k + P)$$

These are solved iteratively to give values for the parameters. The difference between the right and left sides of equation 6 is defined as error parameter R and the difference between the right and left sides of equation 7 is defined as error parameter S. An approximation algorithm using a first-order Taylor series expansion drives R and S to zero. This approximation algorithm requires an initial guess for frequency and phase close to the solution to ensure convergence to the best-fit sine wave. For frequency, the frequency of the generator output in Fig. 11 is used as a guess. For phase, a guess is based on an examination of the data record by a software routine.

Although the result of this process is a single figure of merit, some enlightenment can be gained about the error components in the ADC by varying the test conditions. White noise produces the same degradation regardless of input frequency or amplitude. That is, the error term in equation 1 is independent of test conditions for this sort of error. Another way of identifying noise in this test is by the randomness in the error residue, or the difference between the best-fit sine wave and the data taken.

Aperture uncertainty is identifiable because it generates an error that is a function of input slew rate. When this is the dominant error causing a low number of effective bits, the

number of effective bits will vary linearly with both input frequency and amplitude. If the input waveform is sampled only at points of constant slew rate, such as zero crossings, then the aperture uncertainty corresponds to the amount that the effective bits decline as a function of slew rate.

Harmonic distortion is usually a nonlinear function of amplitude and frequency. Its distinguishing characteristic is the presence of the harmonics (or aliased harmonics if the fundamental is close to the Nyquist frequency) in the error residue. The amplitudes of the harmonics can be extracted by fitting the error residue with best-fit sine waves of the important harmonic frequencies. The impact of noise and aperture uncertainty in the presence of large distortion errors can be assessed by effective bit values and error residues with the fitted harmonics removed.

The greatest pitfall in the curve-fit test is using an input frequency that is a submultiple of the sample frequency. Since the same codes are sampled at exactly the same voltages each cycle, the locally uniform probability distribution assumption is violated. In the worst case, a submultiple of one-half, the quantization error would not be measurable at all. From a practical standpoint, this also defeats the global description of the test by sampling only a handful of codes.

Another potential pitfall is lack of convergence of the curve-fit algorithm. There are a few occasions where this can become a problem, such as when the data is very poor or the computational resolution is inadequate.

Fig. 12 shows the error plot for a 5180A curve-fit test taken at a 9.85-MHz input frequency. The number of effective bits associated with this error is 8.51.

FFT Testing

The fast Fourier transform (FFT) is used to characterize an ADC in the frequency domain in much the same way that a spectrum analyzer is used to determine the linearity of an analog circuit. The data output for both techniques is a presentation of the magnitude of the Fourier spectrum for the circuit under test. Ideally the spectrum is a single line that represents the pure sine wave input and is devoid of distortion components generated by the circuit under test. There are, however, significant differences between the spectrum analyzer and ADC spectra because of the sampling operation of the ADC.

The Fourier transform of a signal $x(t)$ that is continuous for all time is defined as

$$X(f) = \int_{-\infty}^{\infty} x(t) e^{-i2\pi ft} dt$$

and includes the amplitude and phase of every frequency in $x(t)$. The Fourier transform cannot be used in this form for an ADC, however, because $x(t)$ is only digitized at a finite number of points, M , spaced Δt apart. Instead, the discrete finite transform (DFT) must be used. It is defined as

$$XD(f) = \sum_{m=0}^{M-1} x(m\Delta t) e^{-i2\pi f(m\Delta t)\Delta t}$$

There are significant differences between $X(f)$ and $XD(f)$.




Fig. 13. Time-domain and frequency-domain representations of rectangular and Hanning windows.

While $X(f)$ has infinite spectral resolution, $XD(f)$ has a discrete frequency resolution of $\Delta f = 1/M\Delta t$ because of the finite number of points in the data record. The finite record size also accounts for another difference between $X(f)$ and $XD(f)$ whenever a nonintegral number of cycles of $X(t)$ is contained in the record. Since the DFT assumes that the record repeats with a period of $M\Delta t$ (to satisfy the Fourier transform condition that $x(t)$ be continuous for all time) sharp discontinuities at the points where the start of one record joins the end of the preceding record cause the spectral components of $X(f)$ to be spread or smeared in $XD(f)$.

The smearing, called leakage, can be explained as follows. The finite record size of $x(t)$ can be considered the consequence of multiplying $x(t)$ by a rectangular function having unity amplitude during the time period $M\Delta t$ that the record is acquired and zero amplitude elsewhere. Since multiplication of two functions in one domain (time, in this case) is equivalent to convolution in the other, the spectrum of $XD(f)$ is derived by convolving $X(f)$ with $W(f)$, the Fourier transform of the rectangular function. $W(f)$ is the familiar $\text{sinc}(x)/x$ function (see Fig. 13 for $|W(f)|$), consisting of a main lobe surrounded by a series of sidelobes whose amplitudes decay at a 6-dB-per-octave rate. It is these sidelobes that are responsible for leakage. Even if the spectrum of $X(f)$ is a single line, the sidelobes of $W(f)$ during the convolution smear the energy in the single line into a series of spectral lines spaced $1/M\Delta t$ apart whenever the frequency of $x(t)$ is not an integral multiple of $1/M\Delta t$.

Leakage can be reduced by multiplying the data in the record by a windowing function that weights the points in the center of the record heavily while smoothly suppressing the points near the ends. Many different windowing functions exist that offer various tradeoffs of amplitude resolution versus frequency resolution. A function commonly used with sine waves is the Hanning window, defined by $|(1/2)(1 - \cos 2\pi t/M\Delta t)|$. Notice in Fig. 13 that both the window and its derivative approach zero at the two ends of the record and that the transform's main lobe is twice as wide as that of the rectangular function, while the amplitudes of the sidelobes decay by an additional 12 dB per octave. The reduced level of the sidelobes reduces leakage, but the

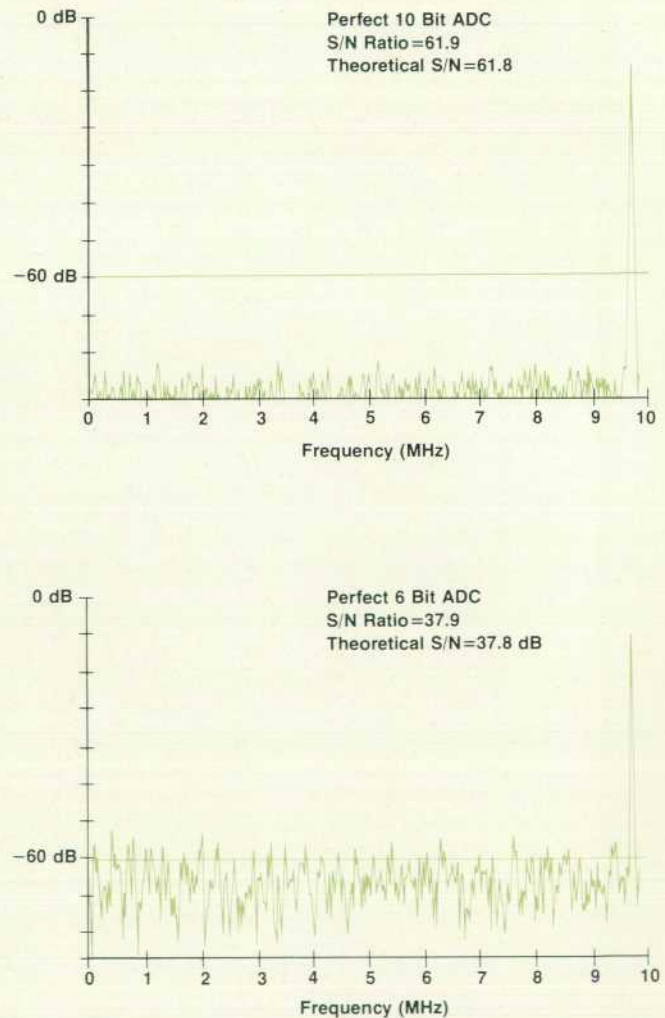


Fig. 14. FFT plots for 0.85-MHz data quantized by perfect 10-bit (top) and 6-bit (bottom) ADCs. The signal-to-noise ratio computed in each case agrees closely with the theoretical value of $6N + 1.8$ dB where N is the number of ADC bits.

wider main lobe limits the ability to resolve closely spaced frequencies. Furthermore, the shape of the main lobe can attenuate the spectral amplitudes of $X(f)$ by as much as 1.5 dB. However, for the DFT testing to be described here, the Hanning window was selected as a good compromise between frequency and amplitude resolution.

The third difference between the spectra of $X(f)$ and $XD(f)$ is the limited range of frequencies displayed for $XD(f)$. The sampling process causes the two-sided spectrum of $X(f)$, symmetrical about the origin, to be replicated at the sampling frequency f_s and at all of its harmonics. If $X(f)$ contains components that exceed $f_s/2$, then these components are folded back, or aliased, onto spectral lines below f_s , causing aliasing errors. The frequency $f_s/2$ is sometimes called the Nyquist frequency, referring to the Nyquist criterion, which requires the sampling rate to be twice the highest frequency present in the input signal to define the waveform uniquely. The result is that the spectrum of $XD(f)$ is displayed only from dc to $f_s/2$ and the maximum input frequency must be limited to less than $f_s/2$ to avoid aliasing.

Fig. 14 presents the magnitude of the spectra derived from the DFT for perfect 10-bit and 6-bit ADCs given a pure sinusoidal input. Useful information about the ADCs' performance can be derived from three features of the spectra:




Fig. 15. DFT plots for the 5180A with input frequencies of 9.85 MHz (top) and 0.95 MHz (bottom). The low harmonic distortion indicates very low integral nonlinearity.

the noise floor, the harmonic level, and the spurious level.

Two classes of noise sources determine the level of the noise floor. The first is called quantization noise. This is the error, bounded by $\pm 1/2$ LSB, that is inherent in the quantization of the input amplitude into discrete levels. As can be seen in Fig. 14, even ideal ADCs have noise floors determined by quantization noise. The higher the number of bits, the smaller the error bound and, therefore, the lower the noise floor.

All real-life ADCs have noise floors that are higher than that solely from quantization noise. The second class of noise source includes wideband noise generated within the ADC, along with other sources. In a parallel-ripple ADC, for example, such things as misadjustment between the first-pass and second-pass ranges (exceeding the redundancy range) or inadequate DAC settling can cause localized code errors or differential nonlinearities in the ADC's static transfer function. Furthermore, localized code errors can increase in amplitude and in the number of codes affected under dynamic input conditions. Aperture jitter is another major source of dynamic error; the magnitude of this localized code error is dependent upon the slew rate of the input at the time of sampling. Each of these localized code errors can be modeled as a sharp discontinuity in the time domain that when transformed into the frequency domain results in a broad spectrum that raises the height of the noise floor above that caused by quantization noise alone.

The second feature of the DFT-derived spectrum that indicates an ADC's level of dynamic performance is the harmonic content. Static and dynamic integral nonlinearities cause curvature in the ADC's transfer function. If the input frequency f_{in} is much lower than the Nyquist frequency ($f_s/2$), then the harmonic components will be in the expected locations: $2f_{in}$, $3f_{in}$, etc. If, on the other hand, the harmonics of f_{in} are greater than $f_s/2$, then these frequencies will be aliased onto components below $f_s/2$. Take, for instance, a 20-megasample-per-second (f_s) ADC with an input of 9.85 MHz. The second harmonic at 19.7 MHz is aliased to 0.3 MHz, the third harmonic at 29.55 MHz is aliased to 9.55 MHz, the fourth at 39.4 MHz is aliased to 0.6 MHz, and so on.

Care must be exercised in selecting the input frequency for the DFT test. An incorrectly chosen frequency can alias one of its harmonic components onto the fundamental and thereby understate the harmonic distortion (in the example above, an input of exactly 5 MHz would place the third harmonic at the fundamental frequency). The input frequency should be chosen so that the harmonics are far enough away to be easily resolvable from the fundamental, whose energy has been spread into several adjacent bins ($1/M\Delta t$ locations) by the Hanning window. This accounts for the 0.15-MHz offset from 10 MHz used in the example of Fig. 14.

The third feature of the DFT-based spectrum that is indicative of the ADC's level of dynamic performance is the spurious content. Spurious components are spectral components that are not harmonically related to the input. For example, a strong signal near the ADC may contaminate the ADC's analog ground somehow and thereby appear in the spectrum. The nearby signal will not only appear as itself, but because of nonlinearities within the ADC, can combine




Fig. 16. DFT plots for a 20-MHz, 8-bit ADC (top) and a 100-MHz, 8-bit ADC (bottom). Full-scale input sine waves at 9.85 MHz were sampled at a rate of 20 MHz. The high levels of harmonic distortion indicate severe integral nonlinearities.

with the input signal to form sum and difference terms resulting in intermodulation distortion.

The combined effects of noise floor, harmonic distortion and spurious errors are reflected in the ADC's rms signal-to-noise ratio, which can be derived from the DFT magnitude spectrum. The signal energy is determined by summing the energy in all the bins associated with the fundamental. The noise energy is the sum of the energy in all other bins. By taking the logarithm of the ratio of signal energy to noise energy and multiplying by 20, the signal-to-noise ratio for the ADC can be calculated. An ideal N-bit ADC having quantization noise only is theoretically known to have a signal-to-noise ratio equal to $(6N+1.8)$ dB, which sets an upper bound. A signal-to-noise ratio below this ideal limit is indicative of errors of all types that the ADC produces.

The FFT test setup is presented in Fig. 11. A full-scale sine wave of a properly chosen frequency is applied to the ADC under test. The low-pass filter ensures a spectrally pure input. A 1024-point record sampled at the maximum sampling rate is then taken and given to the computer, which calculates the DFT using an FFT algorithm. The spectral magnitude is plotted as a function of frequency.

Fig. 15 shows the graphical outputs for the 5180A for full-scale sine wave inputs at 0.95 MHz and 9.85 MHz. As might be expected, the distortion increases with increasing frequency. Harmonic and spurious components are typically better than -60 dBc below 1 MHz and -54 dBc at 9.85 MHz. The latter spectrum at 9.85 MHz is the frequency-domain representation for one of the most demanding tests of an ADC, called the envelope test, which was described earlier.

Fig. 16 presents, for comparison, the test results for commercially available digitizers: a 20-megasample-per-second, 8-bit ADC and a 100-megasample-per-second, 10-bit ADC with a full-scale, 9.85-MHz sine wave input, sampled at 20 megasamples per second. The numerous large harmonic components, both odd and even, are indica-

Error	Histogram	DFT	Sine Wave Curve-Fit	Beat Frequency Test
Differential Nonlinearity	Yes—shows up as spikes	Yes—shows up as elevated noise floor	Yes—part of rms error	Yes
Missing Codes	Yes—shows up as bins with 0 counts	Yes—shows up as elevated noise floor	Yes—part of rms error	Yes
Integral Nonlinearity	Yes (could be measured directly with a highly linear ramp waveform)	Yes—shows up as harmonics of fundamental aliased into baseband	Yes—part of rms error	Yes
Aperture Uncertainty	No—averaged out. Can be measured with "locked" histogram	Yes—shows up as elevated noise floor	Yes—part of rms error	No
Noise	No—averaged out. Can be measured with "locked" histogram	Yes—shows up as elevated noise floor	Yes—part of rms error	No
Bandwidth Errors	No	No	No	Yes—used to measure analog bandwidth
Gain Errors	Yes—shows up in peak-to-peak spread of distribution	No	No	No
Offset Errors	Yes—shows up in offset of distribution average	No	No	No

Fig. 17. Summary of the errors exposed by the dynamic tests.

tive of severe harmonic distortion errors resulting from integral nonlinearity in the transfer functions of both of these ADCs.

A rule of thumb has evolved that uses the DFT-based spectrum as a quick overview of an N-bit ADC's dynamic performance. If all harmonic and spurious components are at least 6N dB below the full-scale amplitude of the fundamental, then the ADC is performing satisfactorily, since each error component has a peak-to-peak amplitude smaller than an LSB. If, on the other hand, harmonic or spurious components are less than 6N dB down, or if the noise floor is elevated, then other tests can be performed that are better at isolating the particular integral and differential nonlinearity errors. In particular, the FFT test can be followed by the histogram test or the beat frequency test (or envelope test), as conditions warrant.

Conclusion

The four sine-wave-based ADC tests described provide information about the quality of any recorder. The tests may be used to isolate specific failures, even at high speed and fine resolution (Fig. 17). The tests are simple to run, requiring only a synthesized generator and an HP-IB computer.

References

1. D.J. Packard, "Beat Frequency Testing of Real Time A/D Converters," Workshop on High Speed A/D Conversion, Portland, Oregon, February 1980.
2. H.U. Koller, "New Criterion for Testing Analog-to-Digital Con-

verters for Statistical Evaluation," IEEE Transactions on Instrumentation and Measurement, Vol. IM-22, pp. 214-17, September 1973.

3. R. Potter, "Least-Squares Estimation of Sinusoidal Parameters from Measured Data," Hewlett-Packard internal memo, June 1974.

4. L. Ochs, "Measurement and Enhancement of Waveform Digitizer Performance," IEEE International Convention, Boston, May 1976.

5. R.N. Bracewell, *The Fourier Transform and Its Applications*, McGraw-Hill, New York, 1965.

J. Martin Neil

A native of Astoria, Oregon, Marty Neil received the BES degree in electrical engineering from Brigham Young University in 1970 and the MSEE degree from Stanford University in 1971. After two years as a member of the technical staff of Bell Telephone Laboratories, he joined HP as service engineer for microwave counters. He then served as product manager for waveform recorders and is now product marketing manager with HP's Santa Clara Division. He's written several magazine articles on counter applications and on waveform recording and analysis.

Marty is an avid basketball player, a Little League and youth soccer coach, and an active member of his church. He's married, has four children and lives in Cupertino, California.



Time Base Requirements for a Waveform Recorder

by Steven C. Bird and Jack A. Folchi

FOR AN ANALOG-TO-DIGITAL CONVERTER (ADC) to digitize an input waveform accurately, there must be a time base or reference oscillator to provide a continuous, stable, sampling or ENCODE clock. The 5180A Waveform Recorder contains a room-temperature, low-noise, fundamental-mode crystal oscillator operating at a frequency of 20 MHz. There is also provision for selecting an external, user-supplied ENCODE clock.

Time Base Stability Requirements

If the stability of the sampling clock is not good enough, then one or more of the least-significant bits of the digitized result will be in error. In specifying the stability of an oscillator, two areas need to be identified. They are long-term stability, called drift or aging, and short-term stability. Other names for short-term stability are jitter, cycle-to-cycle jitter, close-in phase noise, $\Delta t/t$, and $\Delta f/f$. All of these terms

mean the same thing. They specify the degree to which an oscillator will produce exactly the same frequency over a given short amount of time, usually less than one second. This article will concentrate mainly on short-term stability, since that has the greatest effect on the sampling process of a waveform recorder.

Since the cycle-to-cycle jitter of an oscillator alters the time between samples in an ADC, there is a direct relationship between phase noise of an oscillator and amplitude errors in the ADC conversion process. Ideally this phase noise should not produce more than one-half least-significant bit (LSB) of amplitude error under worst-case conditions. These conditions include maximum sample rate f_s , maximum record length or maximum amount of data storage M, and maximum frequency and amplitude of the input signal $v_{in}(t)$. According to Nyquist's sampling theorem, if a signal is being sampled at a rate f_s , then it is

possible to reconstruct the signal only if the bandwidth of the sampled signal is no greater than $f_s/2$. Thus for worst-case analysis the input signal $v_{in}(t)$ can be described as:

$$v_{in}(t) = A \sin(2\pi f_s t/2) \quad (1)$$

where A is equal to one-half the full-scale dynamic range of the ADC. The maximum amplitude error occurs when sampling the input signal at the zero crossings, where its slope is maximum. The slope of the input signal is given by:

$$dv_{in}(t)/dt = (2\pi f_s/2) A \cos(2\pi f_s t/2) \quad (2)$$

At $t=0$, the slope of $v_{in}(t)$ is $A2\pi f_s/2$ volts/second.

As the number of ADC bits increases, the change of $v_{in}(t)$ between bit representations decreases and we can approximate the maximum slope between samples of $v_{in}(t)$ as $A2\pi f_s/2$ with increasing accuracy. As an example, for seven bits, the difference between the actual slope and the approximation is only 0.001%. Thus the amplitude error in volts is:

$$\Delta v_{in}(t) = \Delta t(A2\pi f_s/2) \quad (3)$$

Since the peak-to-peak amplitude of the input signal is $2A$ and is equal to the full dynamic range of the N -bit ADC, the maximum allowable voltage error of $1/2$ LSB is:

$$\Delta v_{in}(t) \text{ maximum} = 1/2 \text{ LSB} = 2A/2^{N+1} \quad (4)$$

Equating the last two equations above and solving for Δt yields an expression for the maximum allowable time base error or jitter:

$$\Delta t \text{ maximum} = 1/(2^{N+1}\pi f_s/2) \quad (5)$$

Δt is in units of seconds and is a peak value. Since the actual oscillator jitter is a random statistical process with a Gaussian distribution, we need to convert Δt to an rms value to compare the theoretical phase noise requirements with the measured phase noise. If we equate Δt maximum to $2.5\sigma_x$, where σ_x is the standard deviation of the Gaussian random variable $x = \Delta t$, then the actual value of Δt can be expected to be less than our maximum Δt 98.76% of the time.

The actual Δt is a function of the time interval τ over which Δt is measured. The waveform recorder needs an oscillator for which $\sigma_x(\tau)$ is less than $0.4 \times (\Delta t \text{ maximum})$ for τ equal to the maximum duration of a waveform recorder measurement using the maximum ADC sampling rate. This duration is the product of the minimum sample period and the maximum amount of data storage. Since the maximum amount of memory is fixed, longer measurement durations are obtained only by using slower sample rates. Since equation 5 shows a linear decrease in $\sigma_x(\tau)$ as the sample rate is lowered and there is a linear increase in measurement duration, the requirement on $\sigma_x(\tau)$ relaxes linearly with unity slope for slower sample rates.

In Fig. 1a σ_x is plotted as a function of τ , which corresponds to measurement duration. The line with zero slope is called a white phase noise floor process. The unity slope line is called a flicker-of-frequency process. Note

that varying the sample rate affects the white phase noise floor process only, not the flicker-of-frequency process.

Fig. 1b shows a family of requirement curves obtained by holding the sample rate and memory length constant and varying the number of ADC bits. This affects both the white phase noise floor process and the flicker-of-frequency process. Fig. 1c shows a family of requirement curves obtained by holding the sample rate and number of bits in the ADC constant and varying the memory length from 8K samples to 64K samples. Note that varying the memory length affects only the flicker-of-frequency process.

Dividing $\sigma_x(\tau)$ by τ , thereby normalizing to the measurement duration, yields a parameter known as $\sigma_y(\tau)$, where y is equal to $\Delta t/t$. Transforming $\sigma_x(\tau)$ to $\sigma_y(\tau)$ is necessary to

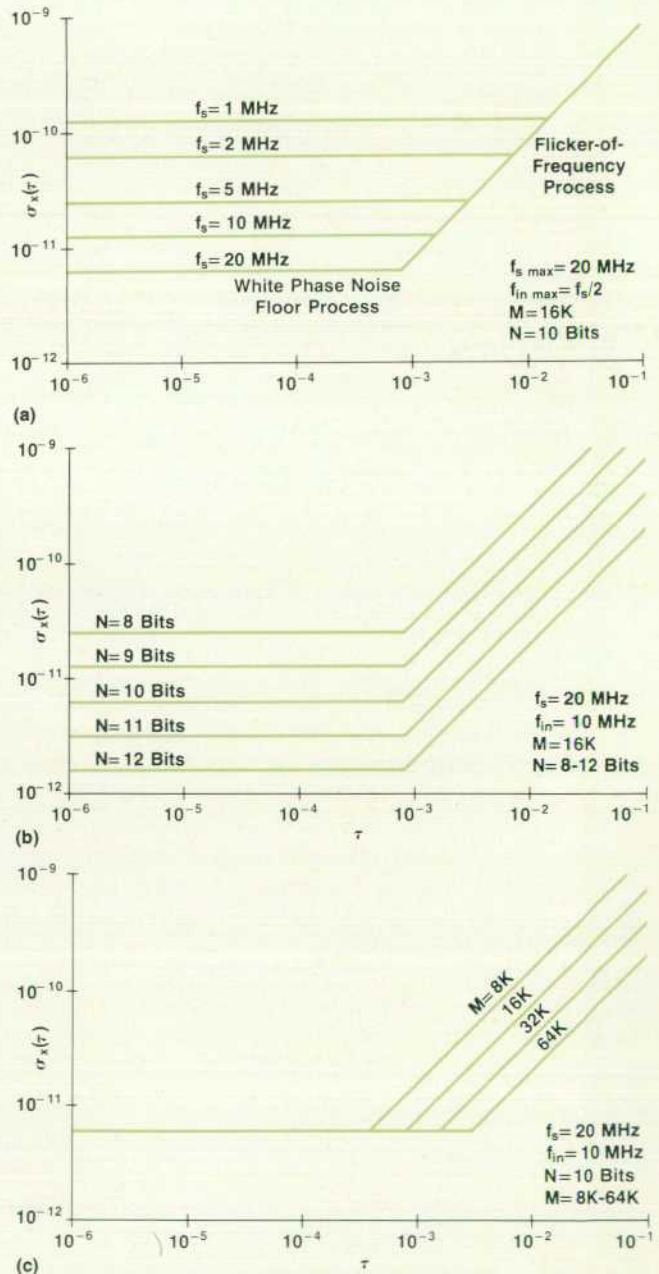


Fig. 1. Time base stability requirement as a function of (a) sample rate, (b) number of ADC bits, and (c) memory length.

put the jitter requirement curves into a form that can be compared to measured values. Figs. 2a, 2b, and 2c show the same time-domain requirements as Figs. 1a, 1b, and 1c expressed using $\sigma_y(\tau)$. The white phase noise floor process now has a slope of -1 and the flicker-of-frequency process has a slope of zero. A technique for measuring $\sigma_y(\tau)$ is described later in this article.

The most common representation of phase noise is in the frequency domain, where it is known as $\mathcal{L}(f)$. $\mathcal{L}(f)$ is described as the power ratio of a single phase-modulation sideband to the total signal power, referenced to the carrier on a spectral-density basis. $\mathcal{L}(f)$ is measured in units of dBc/Hz and is most commonly plotted as a function of

frequency offset from the carrier. Since $\Delta f/f$ is equivalent to $\Delta t/t$, the ADC time-domain jitter requirement curves, Fig. 1, can be transformed into frequency-domain $\mathcal{L}(f)$ requirement curves.¹

Figs. 3a, 3b, and 3c show a waveform recorder's oscillator requirement curves transformed into the frequency domain. It is instructive to note that when the sample rate is varied by a factor of two, there is a 6-dB shift in the noise floor process. Each time the record length changes by a factor of two, there is a 6-dB shift in the flicker-of-frequency process. There is a 6-dB shift in both processes for each additional bit in the ADC.

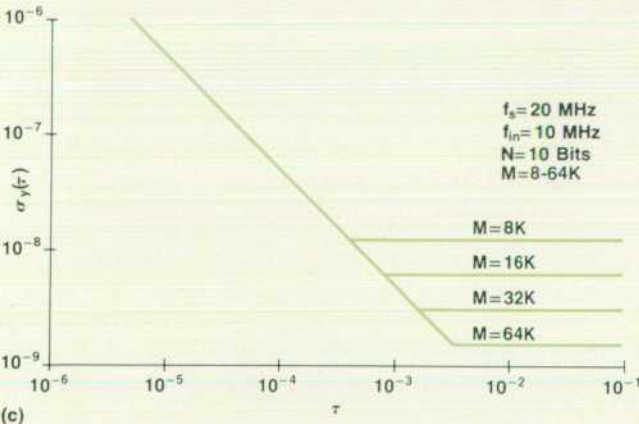




Fig. 2. $\sigma_y(\tau)$ representation of time base stability requirement as a function of (a) sample rate, (b) number of ADC bits, and (c) memory length.

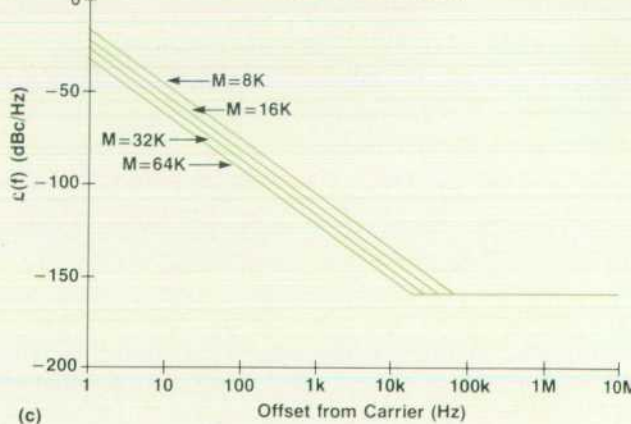

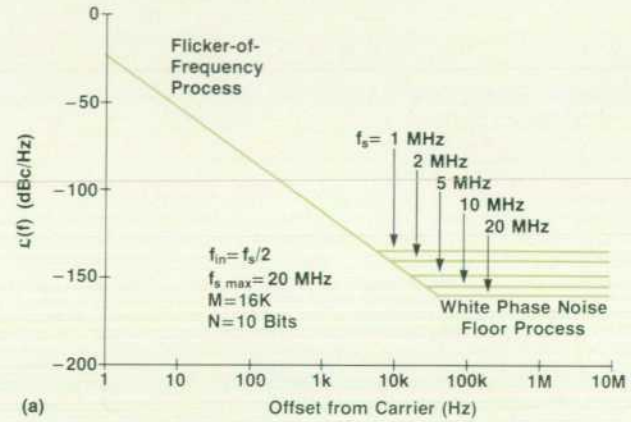


Fig. 3. $\mathcal{L}(f)$ representation of time base stability requirement as a function of (a) sample rate, (b) number of ADC bits, and (c) memory length.

Measuring Oscillator Phase Noise

Now that the phase noise requirements for an oscillator used in a waveform recorder are known, some measurement techniques will be described. There is no easy way to measure cycle-to-cycle jitter of an oscillator directly since it is a short-term phenomenon. First let us define cycle-to-cycle jitter, better known as instantaneous fractional frequency deviation. Any frequency source can be described as,

$$f(t) = [A_o + A_d] \cos[2\pi f_o t + \phi(t)] \quad (6)$$

where A_o = nominal amplitude (volts)
 A_d = amplitude fluctuations (volts)
 f_o = nominal frequency (Hz)
 $\phi(t)$ = phase fluctuations in radians.

The instantaneous frequency can be expressed as the time derivative of the instantaneous phase.

$$\omega_o = \frac{d}{dt}[2\pi f_o t + \phi(t)] = 2\pi f_o + \frac{d}{dt}[\phi(t)] \quad (7)$$

Normalizing the above to ω_o , the instantaneous fractional frequency deviation can be expressed as

$$y(t) = \frac{1}{2\pi f_o} \frac{d}{dt}[\phi(t)] \quad (8)$$

One way of measuring this is by making use of statistics. This entails measuring the average frequency f_{avg} for n measurements at a given time interval τ . The fractional frequency deviation for that time interval is then

$$y_i = (f_{avg} - f_i)/f_i.$$

$\sigma_y(\tau)$ is then given by:

$$\sigma_y(\tau) = \sqrt{\frac{1}{2n} \sum_{i=1}^n (y_{i+1} - y_i)^2} \quad (9)$$

Repeating this for several different averaging times yields a graph of $\sigma_y(\tau)$ versus τ . The HP 5390A Frequency Stability Analyzer performs this measurement. The system requires mixing the output of the oscillator under test with a user-supplied reference to produce a beat frequency of 10 Hz to 100 kHz. The output of the mixer contains sum and difference frequencies. The sum is filtered out while the difference is passed to a frequency counter. Since the measured phase noise is a result of the noise of both oscillators (total noise is equal to the square root of the sum of the squared noise contributions), the reference oscillator noise must be equal to or better than that of the oscillator under test. For example, if the two are equal in quality, the oscillator under test is 1.414 times quieter than the measurement indicates. The HP 5390A system can display either $\sigma_y(\tau)$ in the time domain or single-sideband phase noise $\mathcal{L}(f)$ in the frequency domain. Converting measured results between the two domains demonstrates that the same information is obtained by measurements taken in either domain. When converting, one must be careful to use the same bandwidth

called for by the 5390A software. One must also realize that the bandwidth of the circuit driven by the oscillator under test may be quite different from the measurement bandwidth. This can be compensated for by scaling the white phase noise floor requirement by a factor of

$$\sqrt{BW_2/BW_1},$$

where BW_2 is the bandwidth of the driven circuit and BW_1 is the measurement bandwidth.

Fig. 4a shows the time-domain requirement for an oscillator in a 20-MHz waveform recorder with a 10-bit ADC and a maximum record length of 16K samples. The circuit driven by the oscillator has a bandwidth of 1 GHz while the HP 5390A has a measurement bandwidth of 100 kHz. Therefore, the white phase noise floor process must be scaled by a factor of 100. The flicker-of-frequency process is bandwidth-independent. The dashed line shows the sensitivity limit of the HP 5390A. Fig. 4b shows that both of the curves in Fig. 4a transform to the same frequency-domain requirement. Since the HP 5390A normalizes its frequency-domain measurements to a 1-Hz bandwidth, one can directly compare this measurement to the oscillator performance requirements curve without scaling any process.

Fig. 4 illustrates that the HP 5390A system cannot measure the full range of interest for an oscillator used in a waveform recorder. A recently introduced product that can




Fig. 4. (a) Scaling the noise floor process to compensate for measurement bandwidth. (b) $\mathcal{L}(f)$ normalized to a 1-Hz bandwidth.




Fig. 5. 5180A Waveform Recorder time base oscillator performance.

measure the full range is the HP 3047A Automatic Spectrum Analyzer System. It measures only $\mathcal{L}(f)$ and provides a fast continuous frequency measurement system with a much lower system noise floor. The continuous frequency measurement allows much easier detection of bright lines in an oscillator's output. The 3047A double-balanced mixer requires a +15-dBm drive level on one input port. The other port can have a lower drive level, but the system noise floor degrades proportionally with drive levels below +15 dBm. In measuring the oscillator in the 5180A Waveform Recorder, external low-noise amplifiers are required to increase the -0.5 dBm oscillator output to the +15 dBm required by the 3047A mixer inputs.

A second 5180A oscillator was modified to allow it to be phase-locked. The measured results are shown in Fig. 5.

Once the requirements for $\mathcal{L}(f)$ are derived (Fig. 4b), and the oscillator's phase noise performance is measured (Fig. 5), one can easily compare the two. In general, meeting the corner frequency requirements for a family of curves such as Fig. 4b is not a problem, and this is the case for the 5180A. On a single-measurement basis, phase-locking the oscillator to a reference for better long-term stability will not improve the ADC's performance, since this only improves the oscillator's performance inside the corner frequency (i.e., close to the carrier). Eliminating everything closer to the carrier than the corner frequency leads to an overall simplification of the relationship between bit resolution and the white phase noise floor process; the relationship becomes 6 dB/bit. Knowing this, one can derive the theoretical bit resolution as a function of input frequency for a constant (measured) noise floor, as shown in Fig. 6. This can be thought of as bit compression caused by the noise floor process.

The ADC performance curve of Fig. 6 can be measured by making use of a sine wave curve-fit test (see article, page 21). This test is set up such that all of the ADC output codes (bit representations) are exercised. This requires the

amplitude of $v_{in}(t)$ to be full-scale and the input frequency to be nonharmonically related to the sample frequency f_s . The beat frequency and the memory record length must be such that all codes are captured at least once. Since the slope of $v_{in}(t)$ is a cosine wave and is sampled throughout the full cycle, the slope used in equations 2 and 3 should be an rms value. Using an rms factor of $\sqrt{2}$ shifts both processes of the time-domain requirements of Figs. 1 and 2 up by a factor of $\sqrt{2}$. This shifts both processes of $\mathcal{L}(f)$ in Fig. 3 up by 3 dB. Fig. 6 shows a one-half-bit increase in bit resolution that one can expect as a result of the sine wave curve fit. This assumes no noise from other sources of rms error such as differential nonlinearities, missing codes, integral nonlinearities, or aperture uncertainty. The test results are a combination of all these sources of uncertainty.

Improving Stability

What happens when one measures an oscillator and it doesn't meet specifications in some way? What can one do to improve each process? Knowing what circuitry affects each process enables the designer to improve each process. The white phase noise floor is caused solely by the noise floor of the active device. Thus, if the noise floor is to be improved, design with an active device that has a lower noise floor. It is important to note that any successive stages must have a noise floor at least as low as that of the oscillator stage to avoid degrading the noise floor performance of the oscillator. The flicker-of-frequency process, $1/f^3$, has two contributing processes. The first is a $1/f^2$ process caused by the crystal, which can be modeled as a second-order bandpass filter. By using a crystal with a higher Q or an active device with improved linearity one can improve the $1/f^3$ process and thus the corner frequency; this may be required when one increases memory length. The other contributing process is $1/f$ noise. This process is not well understood, but is generally significant around dc in the spectrum. Because of nonlinearities in the active device or




Fig. 6. Theoretical worst-case effective bit resolution of the 5180A for its measured noise floor of -148 dB (Fig. 5). Effective bits measured in the sine wave curve-fit test are $\frac{1}{2}$ bit better.

devices, a mixing action causes this $1/f$ noise to appear near the carrier. Using an active device with improved linearity will also improve corner frequency. Nonlinearities in successive stages can also cause mixing and degraded noise by mixing $1/f^3$ noise further with the signal.

Acknowledgments

We would like to acknowledge the contributions and efforts of many people who helped in the design and understanding of the time base, especially Bryce Jeppsen, who did some initial oscillator requirements investigations, and Mike Fischer, who helped with the HP 5390A system and in

converting time-domain requirements to the frequency domain.

Reference

1. M.C. Fischer, "Frequency Stability Measurement Procedures," Eighth Annual Precise Time Interval Applications and Planning Meeting, December 1976.



Steven C. Bird

Steve Bird designs oscillators, time bases, and HP-IB interfaces and manages multi-user computer resources for his R&D section at HP's Santa Clara Division. A member of the IEEE, he was born in Cut Bank, Montana, received his BSEE degree from California State Polytechnic Institute at Pomona in 1979, and is now working for his MSEE degree at Stanford University. He's married, has a son, and lives in San Jose, California.



Jack A. Folchi

Jack Folchi received his BSEE and MSEE degrees in 1973 and 1975 from the University of Santa Clara. With HP since 1975, he designed the oscillator and time base and wrote much of the firmware for the 5180A Waveform Recorder. He's now a project leader with HP's Santa Clara Division and is a member of the IEEE Computer Society. Jack was born in Portola, California and now lives in Santa Clara. He's married, has two daughters, and enjoys photography and skiing.

Display and Mass Storage for Waveform Recording

by Christina M. Szeto and Michael C. Detro

THE 5181A DISPLAY/TAPE STORAGE MODULE, Fig. 1, is a companion instrument to the 5180A Waveform Recorder. It provides a high-resolution CRT for waveform display and a cartridge tape unit for recording digitized waveforms onto magnetic tape data cartridges. The 5181A is designed for both bench and field use. With mass storage and a display, it is possible to make measurements on location and bring them back into a lab for analysis.

In the 5181A, two OEM instruments from Hewlett-

Packard, combined with careful product and software definitions, provide the capabilities needed. The first of these two instruments is the 1332A Display. The 1332A is a high-quality display with a spot size of 0.30 mm. This gives a pleasing, sharp picture of most waveforms.

The second of the two instruments is the 9915A Modular Computer.¹ It contains both the data cartridge and the operating system of the HP-85 Desktop Computer. By adding a dedicated program resident in ROM inside the 9915A, the computer is customized to control the flow of data




Fig. 1. The HP 5181A Display/Tape Storage Module is a companion instrument to the 5180A Waveform Recorder.

between the 9915A and the 5180A. This dedicates the 9915A to the 5181A task.

Features

One data cartridge stores 64K 5180A data words plus associated front-panel setups. Waveform data is stored in any of four files on the tape with each file providing storage space for 16K 5180A data points (Fig. 2). Each file is structured to mirror the 16K in the 5180A. For convenience, the 5181A can store waveform data of different lengths simultaneously on the same tape file.

The tape is overwrite-protected so that important data won't be accidentally destroyed. The 5181A responds to an attempt to record over existing data with an error message. If the user really does want to destroy the old data, the PURGE LOC function can be used.

5180A waveform data that has been recorded onto a 5181A cartridge can be read back to any 5180A using the PLAYBACK function. Besides writing the waveform data to the 5180A, the 5181A also returns the 5180A front-panel settings to the values used to record the waveform. This means that the original measurement may be as easily interpreted in the laboratory as it was when the measurement was made.

Because the user may want to analyze the data by computer, the 5181A data tape can also be read by an HP-85, HP 9835, or HP 9845 Desktop Computer. This allows post-processing of the data at a site remote from the 5180A/5181A.²

Inside the 5181A

The mechanical design combines the two OEM units, extends the package length to that of the 5180A, adds mechanical strength, and makes it possible to tie the 5181A and 5180A together. The additional length allows room to store the cables that interconnect the 5181A to the 5180A.

The design of the dedicated operating program was constrained to use the four 9915A softkeys plus the shift key. Using relatively simple interaction with the 5180A, the 5181A can enable the 5180A so that the user can set recording parameters using the 5180A's data entry knob and then initiate action with the 5181A. To help the first-time user, the 5180A has pull-out instruction cards that indicate the correct operation of both the 5180A and the 5181A.

The BASIC operating system of the 9915A offered several features to facilitate the 5181A design. For example, the language makes the softkeys easy to use. The 5181A takes advantage of this by setting up a routine that corresponds to each key (Fig. 3a). The main dedicated operating program establishes the linkages between the keys and the execution routines.

The 9915A treats the program ROM as a mass storage device like a tape or disc. To conserve memory, one of the data transfer routines (STORE) and the tape initialization routine are brought into the machine from the ROM only when needed (Fig. 3b). This makes it possible to transfer 16K-word records to and from the 5180A with only the standard 16K memory in the 9915A.

When a softkey is pressed, the program jumps to the specified execution subroutine, where two things happen. First, all keys are disabled except the CLEAR FUNCTION key. This makes sure that the function completes or aborts cleanly. Second, the execution routine either completes the function itself or loads the necessary data transfer routine from the ROM and executes it. Because of the powerful HP-IB* and mass storage features of 9915A BASIC, the entire 5181A program fits in less than 10K bytes.

The 5181A must communicate with the 5180A to send and receive data and to control which block of memory gets

*Hewlett-Packard Interface Bus, HP's implementation of IEEE Standard 488 (1978).




Fig. 2. 5181A file structure. Several records of different lengths can be stored in the same 16K file.

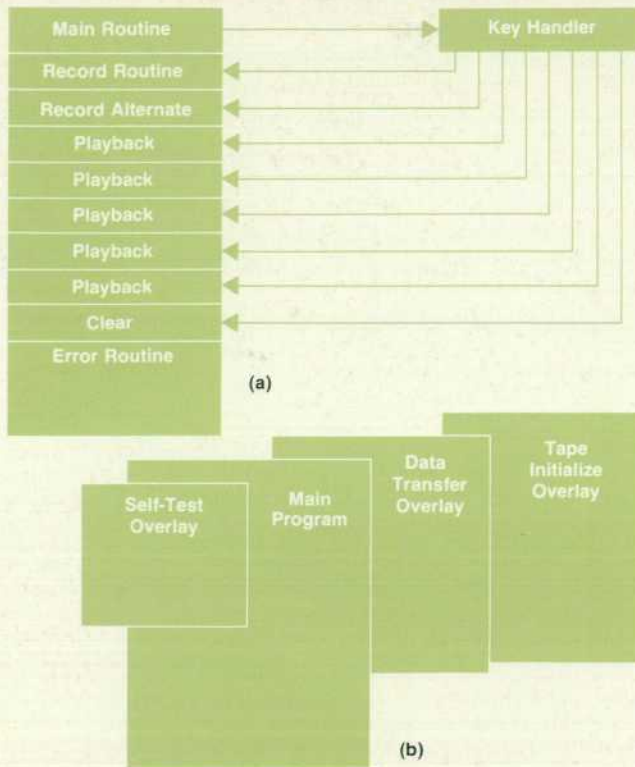


Fig. 3. (a) The 5181A executes a routine for each front-panel softkey. (b) To conserve memory, some other routines are brought into main memory from the program ROM only when needed.

transferred. To do this, it makes use of the 5180A's teach and learn commands and its binary data input and output. Since it is possible for the link connecting the two units to be disconnected, the 5181A checks the link as part of its power-up self-test. If it finds an error here or at any other time, it reports the problem to the user by calling an error routine that flashes an error code on the eight light-emitting diodes (LEDs) on the front panel. Interpretation of the error codes is on the pull-out instruction cards.

Conclusion

The 5181A design required 11 months from initial conception to release. Its speedy development was made possible by keeping the definition a simple, easily understandable analog of the 5180A, and by making use of OEM components. The powerful language and I/O capabilities of the 9915A Modular Computer made the programming job fast and lowered the risk of software errors.

References:

1. V.C. Jones, "Desktop Computer Redesigned for Instrument Automation," *Hewlett-Packard Journal*, Vol. 32, no. 7, July 1981.
2. HP Product Note 5181A-1.

Michael C. Detro



Mike Detro was born in Berkeley and graduated from the University of California there, receiving the BSME degree in 1978. He joined HP's Santa Clara Division in 1978 and contributed to the product design of the 5180A Waveform Recorder and the 5181A Display/Storage Module. He lives in Los Gatos, California. When he's not sailing, diving, rafting, camping, or engaging in other outdoor activities, he keeps busy with home improvement projects.

Christina M. Szeto



Chris Szeto received her BS degree in electrical engineering and computer science from the University of California at Berkeley in 1981. Joining HP's Santa Clara Division as a development engineer, she developed the firmware for the 5181A Display/Storage Module and firmware enhancements for the 5180A Waveform Recorder. She's a member of the Society of Women Engineers, a native of Hong Kong, and enjoys tennis, aerobic dancing, and traveling—her dream is to live in Europe for a year.

Hewlett-Packard Company, 3000 Hanover Street, Palo Alto, California 94304

Bulk Rate
U.S. Postage
Paid
Hewlett-Packard
Company

HEWLETT-PACKARD JOURNAL

NOVEMBER 1982 Volume 33 • Number 11

Technical Information from the Laboratories of
Hewlett-Packard Company

Hewlett-Packard Company, 3000 Hanover Street
Palo Alto, California 94304 U.S.A.

Hewlett-Packard Central Mailing Department
Van Heuven Goedhartlaan 121

1181 KK Amstelveen, The Netherlands

Yokogawa-Hewlett-Packard Ltd., Suginami-Ku Tokyo 168 Japan

Hewlett-Packard (Canada) Ltd.

6877 Goreway Drive, Mississauga, Ontario L4V 1M8 Canada

020002 070766 & BLAC & CA00
MR C A BLACKBURN
JOHN HOPKINS UNIV
APPLIED PHYSICS LAB
JOHNS HOPKINS RD
LAUREL MD 20707

CHANGE OF ADDRESS: To change your address or delete your name from our mailing list please send us your old address label. Send changes to Hewlett-Packard Journal, 3000 Hanover Street, Palo Alto, California 94304 U.S.A. Allow 60 days