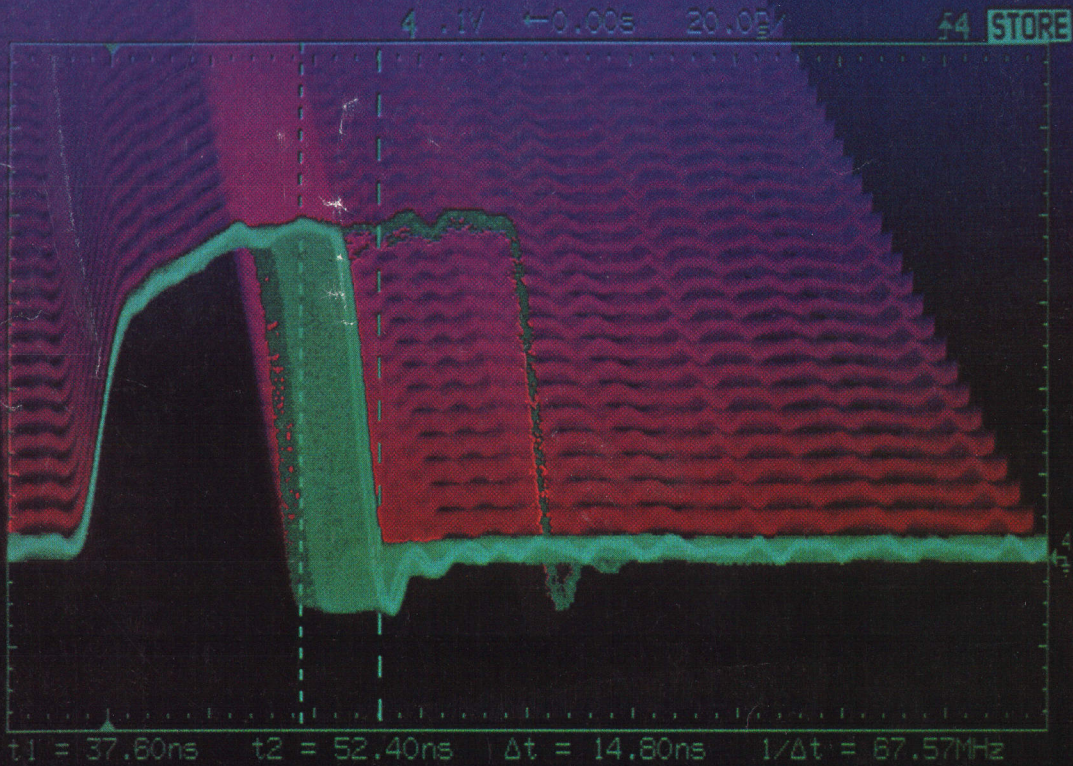
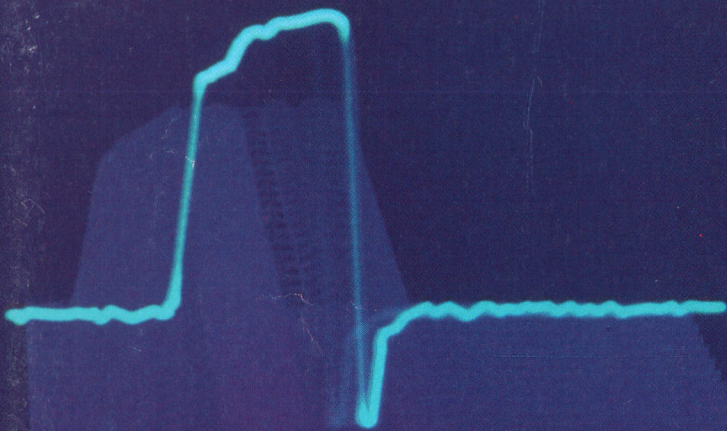


HEWLETT-PACKARD JOURNAL

February 1992



Articles

-
- 6 **Low-Cost, 100-MHz Digitizing Oscilloscopes**, by Robert A. Witte
-
- 11 **A High-Throughput Acquisition Architecture for a 100-MHz Digitizing Oscilloscope**, by Matthew S. Holcomb and Daniel P. Timm
-
- 18 **Sample Rate and Display Rate in Digitizing Oscilloscopes**
-
- 21 **A Fast, Built-In Test System for Oscilloscope Manufacturing**, by Stuart O. Hall and Jay A. Alexander
-
- 22 **Verification Strategy**
-
- 27 **Stimulus/Response Defect Diagnosis in Production**
-
- 29 **Measuring Frequency Response and Effective Bits Using Digital Signal Processing Techniques**, by Martin B. Grove
-
- 34 **Calculating Effective Bits from Signal-To-Noise Ratio**
-
- 36 **Mechanical Design of the HP 54600 Series Oscilloscopes**, by Robin P. Yergenson and Timothy A. Figge
-
- 41 **EMC Design of the HP 54600 Series Oscilloscopes**, by Kenneth D. Wyatt
-
- 45 **Digital Oscilloscope Persistence**, by James A. Kahkoska
-
- 48 **A High-Resolution, Multichannel Digital-to-Analog Converter for Digital Oscilloscopes**, by Grosvenor H. Garnett
-
- 54 **Using the High-Resolution, Multichannel DAC in the HP 54601A Oscilloscope**
-
- 57 **Comparing Analog and Digital Oscilloscopes for Troubleshooting**, by Jerald B. Murphy
-

62 **An Introduction to Neural Nets**, by *John McShane*

66 **Design Challenges for Distributed LAN Analysis**, by *William W. Crandall*

75 **Poor Network Partitioning**

Departments

- 4** **In this Issue**
- 5** **Cover**
- 5** **What's Ahead**
- 60** **Authors**

The Hewlett-Packard Journal is published bimonthly by the Hewlett-Packard Company to recognize technical contributions made by Hewlett-Packard (HP) personnel. While the information found in this publication is believed to be accurate, the Hewlett-Packard Company disclaims all warranties of merchantability and fitness for a particular purpose and all obligations and liabilities for damages, including but not limited to indirect, special, or consequential damages, attorney's and expert's fees, and court costs, arising out of or in connection with this publication.


Subscriptions: The Hewlett-Packard Journal is distributed free of charge to HP research, design and manufacturing engineering personnel, as well as to qualified non-HP individuals, libraries, and educational institutions. Please address subscription or change of address requests on printed letterhead (or include a business card) to the HP address on the back cover that is closest to you. When submitting a change of address, please include your zip or postal code and a copy of your old label. Free subscriptions may not be available in all countries.

Submissions: Although articles in the Hewlett-Packard Journal are primarily authored by HP employees, articles from non-HP authors dealing with HP-related research or solutions to technical problems made possible by using HP equipment are also considered for publication. Please contact the Editor before submitting such articles. Also, the Hewlett-Packard Journal encourages technical discussions of the topics presented in recent articles and may publish letters expected to be of interest to readers. Letters should be brief, and are subject to editing by HP.

Copyright © 1992 Hewlett-Packard Company. All rights reserved. Permission to copy without fee all or part of this publication is hereby granted provided that 1) the copies are not made, used, displayed, or distributed for commercial advantage; 2) the Hewlett-Packard Company copyright notice and the title of the publication and date appear on the copies; and 3) a notice stating that the copying is by permission of the Hewlett-Packard Company.

Please address inquiries, submissions, and requests to: Editor, Hewlett-Packard Journal, 3200 Hillview Avenue, Palo Alto, CA 94304 U.S.A.

In this Issue



The analog oscilloscope, all but obsolete for laboratory analysis applications by the digital or digitizing oscilloscope, refuses to die. It remains the first choice of engineers and technicians for troubleshooting because of its low cost, its easy-to-use controls, and its real-time display. Taking this as a challenge, engineers at HP's Colorado Springs Division set out to design a digitizing oscilloscope that troubleshooters would not only find equal to their analog oscilloscopes, but would actually prefer. The HP 54600 Series digitizing oscilloscopes have all the features normally associated with the full-featured 100-MHz analog oscilloscopes most often used for troubleshooting. They have the same bandwidth—100 MHz—and are comparable in cost and ease of use. While they're

clearly digitizing oscilloscopes—displayed waveforms are made up of dots rather than continuous lines—the HP 54600 Series oscilloscopes are as quick to respond to circuit adjustments as analog oscilloscopes in most applications and are actually better for some tasks. What makes them preferable to analog oscilloscopes, not just comparable, is the array of storage and measurement capabilities that only a digitizing oscilloscope can offer. Since waveform data is sampled and stored in a memory, it's possible to see data both before and after a trigger event, manipulate the data mathematically, and display waveforms indefinitely without fading. Beginning with an introductory article on page 6 and ending with a head-to-head comparison with analog oscilloscopes for troubleshooting (page 57), nine articles in this issue deal with the design of the HP 54600 Series oscilloscopes. They describe how the cost issue was addressed by a high level of circuit integration, the use of surface mount technology for loading printed circuit boards, a cost-effective mechanical package, and careful attention to the manufacturing process, including the cost of test time and test equipment. Ease of use was addressed in part by providing dedicated knobs for the main control functions instead of a menu-driven softkey user interface, although menus and softkeys were retained for control of the digitizing oscilloscope functions. The display rate capability was increased to a million points per second, fifty to one hundred times that of other digitizing oscilloscopes, by means of a new architecture and two dedicated integrated circuits. Waveform smoothness was improved by quadrupling the number of points displayed per trace. You'll find the details of the architecture and custom ICs in the article on page 11, the mechanical design on page 36, and the test strategy and test system on page 21. The production test strategy of verification rather than characterization greatly reduces the number of parameters that need to be measured, and new FFT-based measurement algorithms (page 29) further improve efficiency. The production test system is partly built-in and uses only two signal sources and one external instrument, a digital multimeter. On page 41 you can read about the steps taken to ensure that the HP 54600 Series oscilloscopes meet international and military standards for electromagnetic compatibility—important for a troubleshooting instrument. A new way to use the storage and infinite persistence abilities of digitizing oscilloscopes is described in the article on page 45. Called autostore, it displays the latest trace at full intensity and earlier traces at half intensity so the user can see the effects of adjustments more easily. The analog-to-digital converter used in the HP 54600 Series and other HP digitizing oscilloscopes is a 16-channel, 16-bit, indirect type (page 48). In addition to converting waveform samples to digital data, it's used for calibrating the vertical gain.

Neural nets are computing architectures that represent an attempt to build processors modeled on the human brain. While current technology is far from being able to emulate even simple brain functions, neural nets are effective for solving some problems that are difficult to solve with more conventional processor architectures. Unlike conventional methods, which provide algorithms for computing an output given an input, neural net methods specify a procedure by which the net can learn how to produce an output for a given input. In the article on page 62, John McShane, formerly with HP Laboratories in Bristol, England, gives us a brief but interesting introduction to neural nets and their applications.

A computer failure is frustrating for the user of that computer, but a computer network problem can result in dozens of irate users. Distributed network monitoring, with monitors spread throughout a network continuously tracking use, detecting and logging errors and other events, and raising alarms when problems occur, allows a network manager to react much more quickly to network problems than the technique of dispatching a cable tester or protocol analyzer to look for a problem. However, the design of a distributed monitoring system presents challenges. As the article on page 66 explains, several megabytes of data can flow across each segment of a typical Ethernet local area network (LAN) every minute. It's impossible to keep track of it all, so the design of a monitoring system must address the questions of what data to capture, how to process it, how to transmit it to the central station, and how to format and display it for the network manager's use. The article tells how these questions were answered in the design of the HP Lan-Probe monitors and the HP ProbeView software, a distributed Ethernet LAN monitoring system. A case study of a hypothetical company network clarifies the concepts and demonstrates how the various ProbeView tools are used to track down an elusive fault.

R.P. Dolan
Editor

Cover

An artist's rendition of an analog oscilloscope display (background) and an HP 54600A oscilloscope display (foreground) of the output of a circuit designed to synchronize an asynchronous event such as a keypress to a microprocessor clock. The asynchronous event should appear as a high output level for one cycle of the system clock. However, this circuit is sensitive to the timing of the event and occasionally produces an output that is two clock cycles wide. This wide pulse is shown clearly by the HP 54600A but missed by the analog oscilloscope. The jitter is caused by system noise near the circuit's input threshold level. The HP 54600A displays the true peak-to-peak limits of the jitter while the analog oscilloscope displays only the part of the jitter that occurs more frequently.

What's Ahead

The April issue will have several articles on the design and implementation of HP's family of modular, plug-in instruments conforming to the VXIbus instrumentation standard. This standard makes it possible to create a system containing instruments from different manufacturers simply by plugging cards into a mainframe. Also featured in the April issue will be the HP 8990A peak power analyzer, a new microwave instrument that makes accurate, calibration-free measurements on pulsed microwave signals, the HP 8751A 500-MHz network analyzer, a fast, accurate analyzer for evaluating filters and resonators in production and the laboratory, and the HP 82324A measurement coprocessor, a plug-in card for HP Vectra personal computers that runs HP BASIC programs in the DOS environment.

Low-Cost, 100-MHz Digitizing Oscilloscopes

The HP 54600 Series oscilloscopes combine the convenience, familiarity, and display responsiveness of analog oscilloscopes with the features, accuracy, and measurement power of a digital architecture.

by Robert A. Witte

The oscilloscope has been around for a long time and is one of the most basic and versatile electronic test and measurement instruments. Fundamentally, an oscilloscope displays voltage as a function of time, which gives it the ability to view a wide variety of signals. While an oscilloscope is primarily a viewing instrument, many oscilloscopes are now capable of performing automatic quantitative measurements on a waveform.

The analog oscilloscope has been the most common type of oscilloscope because of its low cost and good display quality. Digitizing oscilloscopes have been growing in popularity in the high-performance arena, but their relatively high prices have limited their acceptance in general-purpose applications. As digitizing technology (sampler circuits, analog-to-digital converters, and digital memories) has steadily improved, the cost of digitizing oscilloscopes has decreased. With the introduction of the HP 54600 Series, the cost of a 100-MHz digitizing oscilloscope is now comparable to a full-featured 100-MHz analog oscilloscope.

The HP 54600A two-channel 100-MHz oscilloscope (Fig. 1) and the HP 54601A four-channel 100-MHz oscilloscope represent a major improvement in digitizing oscilloscope technology and product design. The two oscilloscopes are identical in capability except for the number of channels. Both oscilloscopes have two full-range inputs (2 mV/div to 5V/div). In addition, the HP 54601A has two limited-attenuation inputs (100 mV/div and 500 mV/div) optimized for use with logic signals, while the HP 54600A has an external trigger input. The bandwidth of all channels is 100 MHz. A maximum sample rate of 20 megasamples per second provides a 2-MHz bandwidth for capturing single-shot events. The 8-bit analog-to-digital converter has a vertical resolution of 0.4%.

These oscilloscopes also have features normally associated with full-featured analog oscilloscopes. These include delayed sweep, bandwidth limit, ac/dc coupling, adjustable vernier on both horizontal and vertical axes, trigger conditioning (ac/dc coupling, low-frequency reject, high-frequency reject, and noise reject) and a full-screen XY (channel versus channel) mode.




Fig. 1. The HP 54600A is a two-channel 100-MHz general-purpose digitizing oscilloscope.

Analog Barriers

The HP 54600 design team was given the challenge of developing an oscilloscope that would bring the advantages of a true digitizing oscilloscope architecture to the analog oscilloscope user. (Some oscilloscopes have added a digitizer to an analog oscilloscope system, requiring the user to choose between analog and digital operation.) The team identified three main barriers that discourage analog oscilloscope users from converting to digitizing oscilloscopes:

- Cost
- Ease of use
- Display confidence.

Cost. The cost goals for the HP 54600 Series products required that these digitizing oscilloscopes be comparable in price to a full-featured analog oscilloscope of the same bandwidth and channel count. Throughout the project, the manufacturing cost was carefully evaluated. A high level of circuit integration, the use of surface mount technology, a cost-effective mechanical package, and careful attention to manufacturability produced a design that met this goal. While surface mount technology is not inherently less expensive—the cost of the parts and the machine placement of these parts is about the same as the older through-hole technology—the use of surface mount technology allows the oscilloscope circuitry to reside on one printed circuit board (not counting the keyboard, power supply, and display modules). For the acquisition processor, the design team took advantage of HP's CMOS-34 1-micrometer process to produce a high-performance but economical design.

Ease of Use. Previous HP digitizing oscilloscopes have used a totally softkey-driven approach to the user interface. This type of user interface is similar to many other HP instruments and does a good job of giving the oscilloscope user consistent access to a large number of features. The disadvantage of such a front-panel design is that the user is required to navigate through several keystrokes and menu changes to set up the oscilloscope. This type of user interface has been well-accepted by many users, particularly those who use oscilloscopes for analysis and characterization tasks. However, analog oscilloscope users sometimes had trouble adjusting to the menu-driven interface. Especially in troubleshooting applications, the analog oscilloscope user wants easy and direct access to the main controls of the oscilloscope.

The HP 54600A and 54601A use dedicated knobs for their primary control functions. These functions include the three main controls (channel 1 volts/division, channel 2 volts/division, and horizontal time/division) as well as the vertical position controls, horizontal delay (position), trigger level, and trigger holdoff. In addition, a general-purpose entry knob provides control of the measurement cursors and other advanced functions. Dedicated front-panel keys are assigned to the **Autoscale** and storage keys (**Run**, **Stop**, **Autostore**, and **Erase**).

Softkey menus are used for the next level of features, which includes ac/dc coupling, bandwidth limit, trigger source, trigger mode, trigger slope, and main/delayed sweep. Softkey menus also control the advanced set of features associated with digitizing oscilloscopes: automatic measurements, cursors, trace storage, setup storage, and printer control.

Overall, the user interface is a good compromise, providing easy access to the main controls while still supporting the advanced features that digitizing oscilloscope users have come to expect.

Display Confidence. Most digitizing oscilloscope displays are not as responsive to signal changes as their analog counterparts. Typically, the waveform is processed in software using a microprocessor that can cause dead times between acquisitions, limiting the number of waveforms displayed per second. Other display problems such as aliasing leave the analog oscilloscope user wondering




Fig. 2. The waveform processing technology used in the HP 54600 family of oscilloscopes accurately displays complex waveforms such as amplitude modulated signals.

whether the displayed waveform is telling the real story. This is especially important in troubleshooting applications, in which the user may not know what the signal looks like. (Compare this with applications in which the signal is known, such as the characterization of a pulse's rise time.)

In the HP 54600 Series oscilloscopes, display confidence is improved through the use of two integrated circuits. These two ICs provide a display update rate that is unmatched in a digitizing oscilloscope in this price range. With sufficient trigger rate, the display is updated as fast as one million data points per second.

Why is display update rate important? One reason is that many applications involve the adjustment of the circuit being tested. A responsive waveform display lets the user make the adjustment very quickly. Another reason involves complex waveforms. Waveforms such as amplitude modulated signals and video signals are multivalued functions that change very rapidly. A fast display update rate produces a much better picture of such a waveform. To the analog oscilloscope user, the "correct" oscilloscope display is defined as the display that an analog oscilloscope would produce. Fig. 2 shows the HP 54600A oscilloscope display of an amplitude modulated signal. Aliasing has always been a problem in sampled data systems, and oscilloscopes are no exception. When aliasing occurs, an input signal can appear on the oscilloscope display as a signal of a different frequency. Since this is clearly an undesirable situation, the HP 54600 oscilloscopes use HP proprietary sampling and display algorithms to reduce the effects of aliasing.

Digital Storage

Removal of these barriers was required to attract the analog oscilloscope user, but what could motivate such a oscilloscope user to *prefer* a digitizing oscilloscope? What are the advantages of a digitizing oscilloscope when compared to an analog oscilloscope? In a word, *storage*. Storage lets the oscilloscope user capture a waveform, measure it, print or plot it, and transfer it to a computer.

Since the waveform in a digitizing oscilloscope is acquired in digital form and resides in digital memory, measurement possibilities exist that are not available on an analog oscilloscope. The most obvious and simplest use of storage is to acquire a waveform and freeze it onscreen. Then the user can view the waveform or perform cursor or automatic measurements at a leisurely pace. In cases where the signal being viewed is transient in nature or is difficult to probe, the user can obtain the waveform, stop the oscilloscope, and then worry about the measurement. Storage is particularly useful for measuring single-shot events (Fig. 3), but it is also handy when making repetitive measurements. Analog oscilloscopes with special storage tubes can store waveforms, too, but the waveform tends to bloom or fade after a short period of time. A digitally stored waveform can remain on the display indefinitely without losing intensity or becoming distorted.

Alternatively, the oscilloscope user might want to transfer the waveform to a printer or plotter to provide a permanent hard-copy record. In this way, measurement data can be recorded in a laboratory notebook or used as production line or service documentation. If the oscilloscope is connected to a computer, the data can be transferred to the computer for further analysis or for incorporation into a document.

The applications mentioned so far are clearly storage operations. If we expand the notion of storage we can identify a few other storage advantages of a digitizing oscilloscope. The digitizing system of these oscilloscopes is normally running all the time. When a trigger occurs, there are already samples stored in memory and these samples represent the portion of the signal before the trigger event. This pretrigger viewing capability (also known as negative time viewing) lets the oscilloscope user look back in time before the trigger occurred. This has obvious utility for troubleshooting a system. Often the symptoms of the failure trigger the oscilloscope after the cause of the problem has disappeared.

Another important digital storage advantage is fade-free viewing. A waveform on an analog oscilloscope has a




Fig. 3. A digitizing oscilloscope can capture a single-shot transient.

tendency to fade out when the sweep speed is increased, while digital oscilloscopes do not suffer from this problem (see article, page 57). Again going beyond the capabilities of an analog oscilloscope (and most digital oscilloscopes), a special storage mode, called *autostore*, provides current and historical waveform information simultaneously (see article, page 45). The historical information (shown in half-bright intensity) shows the worst-case excursions of the waveform while the latest waveform is shown in full intensity. Thus, an oscilloscope user can monitor the waveform as it exists at that moment and still not miss any temporary aberrations that may have occurred in the past. Again, since the waveform storage is totally digital, there is none of the blooming or fading associated with analog storage oscilloscopes.

Peak Detect

Most digitizing oscilloscopes reduce the effective sample rate at slow sweep speeds by simply throwing away sample points. This causes a problem because narrow pulses or glitches that are easily viewable on fast time-base settings can disappear as the sweep speed is reduced. A special acquisition mode called peak detect (also known as glitch detect) maintains the maximum sample rate at all sweep speeds. In peak detect mode, any glitch that is 50 ns wide or wider is captured and displayed, regardless of the sweep speed.

Automatic Measurements

The days of having to count graticule lines and compute measured values are gone. Since the waveform is stored in digital form, the microprocessor in the oscilloscope can accurately and consistently perform automatic measurements on the data. Maximum, minimum, peak-to-peak, rms, and average voltages can be calculated. Automatic timing measurements include frequency, period, pulse width, duty cycle, rise time, and fall time. Cursors can be automatically placed at the measurement points so the user knows where the measurement is being made on the waveform (Fig. 4). As mentioned previously, the waveform can also be transferred to an external computer if more complex analysis is required.

Digital Interfaces

Three digital interfaces are available as options. The HP-IB (IEEE 488, IEC 625) interface and the RS-232 interface connect to an external computer or hard-copy device (printer or plotter). The parallel (Centronics) interface provides only hard-copy capability.

The parallel interface provides compatibility with printers (available from a variety of vendors) that support HP PCL (HP Printer Control Language) or Epson protocols. The parallel interface is easy to connect. There are no complicated configuration problems with either the oscilloscope or the printer.

RS-232 is the industry standard interface for the personal computer. Oscilloscope users who want instant compatibility with their PC will prefer this interface. All functions of the oscilloscope can be programmed via this interface using easy-to-understand commands. The RS-232 interface can also be used to transfer displayed waveforms to HP




Fig. 4. The oscilloscope automatically measures the frequency of a sine wave and places the cursors automatically to show the measurement location on the waveform.

PCL and Epson-compatible printers and HP-GL (HP Graphics Language) plotters.

The HP-IB interface is the standard instrumentation interface for large automated test systems. Its data transfer rate is better than that of the RS-232 interface and is the preferred interface when measurement time is critical. The HP-IB readily supports multiple instruments connected to a single computer interface while RS-232 usually requires a separate computer interface for each instrument. While RS-232 is standard on PCs, the HP-IB must usually be added as an option, although it is standard on some high-performance computer workstations. Besides offering full programmability, the HP-IB interface supports hard-copy dumps to HP printers and HP-GL plotters.

The digital interfaces are implemented as option modules that attach to the rear of the oscilloscope. Since the interface to the option module includes a portion of the microprocessor bus, ROM and RAM can reside on the option module. Thus, firmware associated with the particular interface can reside on the module. Additional features can be added to the oscilloscope at a later date by adding a module. For example, the HP 54655A/56A test automation module gives the oscilloscope 100 additional stored setups with mask testing for automated test applications. The HP 54657A/8A measurement/storage module adds more channel math capability, up to 100 stored traces, more automatic measurements, a real-time clock, and mask testing.

HP ScopeLink Software

The full programmability of the RS-232 and HP-IB interfaces equips the oscilloscope user with enough power to implement complex test systems, provided that the user is willing to write the necessary software. For customers who don't want to bother with writing custom software just to transfer data from the oscilloscope, an optional software package, called HP ScopeLink, is available. This PC-compatible software package provides a user-friendly computer link from the oscilloscope to the PC, allowing

the oscilloscope user to transfer oscilloscope display images, waveform data, and instrument setup information from the oscilloscope to the computer. Graphical images can be translated into industry standard formats (TIFF and PCX) to be incorporated into documents produced using word processing or desktop publishing software. Graphical images can be stored on the computer disk, viewed on the computer screen and dumped to the computer's printer. The numerical data associated with the waveform (time and voltage pairs) can also be transferred to the computer and stored in a format directly readable by spreadsheet and data analysis programs. These programs can be used for further analysis of the oscilloscope waveform data.

Electrical Design

The electronic hardware in the HP 54600 oscilloscopes includes a high degree of custom integrated circuit technology coupled with the use of conventional surface mount components. A simplified block diagram of the two-channel oscilloscope is shown in Fig. 5.

Analog Section. The signal at the channel input is scaled in amplitude by the attenuator and preamp circuits. Various combinations of gain (in the preamp) and loss (in the preamp and the attenuator) change the amplitude of the signal depending on the volts/division control on the front panel. All these gain and loss settings are under microprocessor control.

The preamp also splits off a replica of the signal for use by the trigger system. The trigger system consists of the trigger multiplexer (MUX), the signal-conditioning circuit, and the high-speed trigger comparator. The trigger multiplexer selects the trigger source (either of the channels, the external trigger, or an ac line sync signal). The signal conditioning circuit provides ac coupling, low-frequency reject, and high-frequency reject. All of these circuits are under microprocessor control.

Returning to the channel section of the block diagram, the output of the preamp goes to the track-and-hold circuit. When a sample needs to be acquired, the track-and-hold circuit freezes the instantaneous voltage of the signal and holds it long enough for the analog-to-digital converter (ADC) to convert the analog voltage to digital form. The postamp increases the signal level out of the track-and-hold circuit so that the full range of the ADC is used.

Acquisition Processor. Both channels of ADC data are fed to the acquisition processor, which takes the sampled data, correlates it in time with the trigger, and places it in waveform memory. Since the oscilloscope uses random-repetitive sampling, the placing of data points into the waveform memory can be computationally complex. The waveform samples are not necessarily acquired in sequential order, so the acquisition processor must determine their proper placement by measuring when the sample was taken relative to the trigger signal. The acquisition processor uses HP CMOS technology to implement dedicated logic that processes sample points at a very high rate.

Waveform Translator. The waveform translator IC takes the time-ordered samples from the waveform memory and writes the data to the display. Basically, this means that the voltage and time values associated with a data point in waveform memory must be translated into vertical and horizontal pixel locations on the display. Once these locations are derived, the waveform translator sets the appropriate memory location in the video RAM. The raster-scan display is refreshed from video RAM under the control of the waveform translator IC.

The two custom ICs (acquisition processor and waveform translator) are dedicated to the task of processing the waveform data for the display. Such dedicated hardware produces display throughput and responsiveness that are equivalent to an analog oscilloscope. A side benefit of this architecture is that the main processor is freed to scan the keyboard and respond quickly to control changes. Together with careful software design, this gives the oscilloscope user immediate feedback when making oscilloscope adjustments.

Microprocessor. A conventional microprocessor system using a 68000 CPU, ROM, and RAM is used to control the oscilloscope hardware. Virtually every circuit in the oscilloscope is under microprocessor software control. This includes simple on-off controls like bandwidth limit and ac/dc coupling as well as control of circuit voltages via a 16-channel digital-to-analog converter (DAC). The microprocessor responds to the user's actions by reading the front-panel keyboard, adjusting the hardware controls, and writing status and error messages on the display. The

microprocessor software also performs a wide range of advanced functions such as measurement cursors, automatic voltage and time measurements, saving and recalling setups and traces, hard-copy control, and averaging of waveforms.

Product Design

With outer dimensions of 7 inches (18 cm) high by 14 inches (36 cm) wide by 12 inches (30 cm) deep, the HP 54600 oscilloscopes fit easily on almost any workbench. At 14 pounds (6.4 kilograms) in weight, they are conveniently portable. An optional cover protects the front panel from accidental damage, and an optional pouch can hold the oscilloscope probes and operating manual.

Cost of Ownership

The HP 54600 oscilloscopes have a standard three-year warranty (five-year warranty optional). They are designed for high reliability and easy calibration to produce a low cost of ownership. Most of the usual adjustments have been eliminated through the use of software calibration. (The 4-channel oscilloscope has only eight manual adjustments inside it.) The microprocessor uses the 16-channel DAC to control various points throughout the instrument. During the software calibration procedure (which can be performed in minutes without opening the oscilloscope), measurement errors such as offset, gain, and channel-to-channel skew are detected by the microprocessor and removed under software control. The only required piece

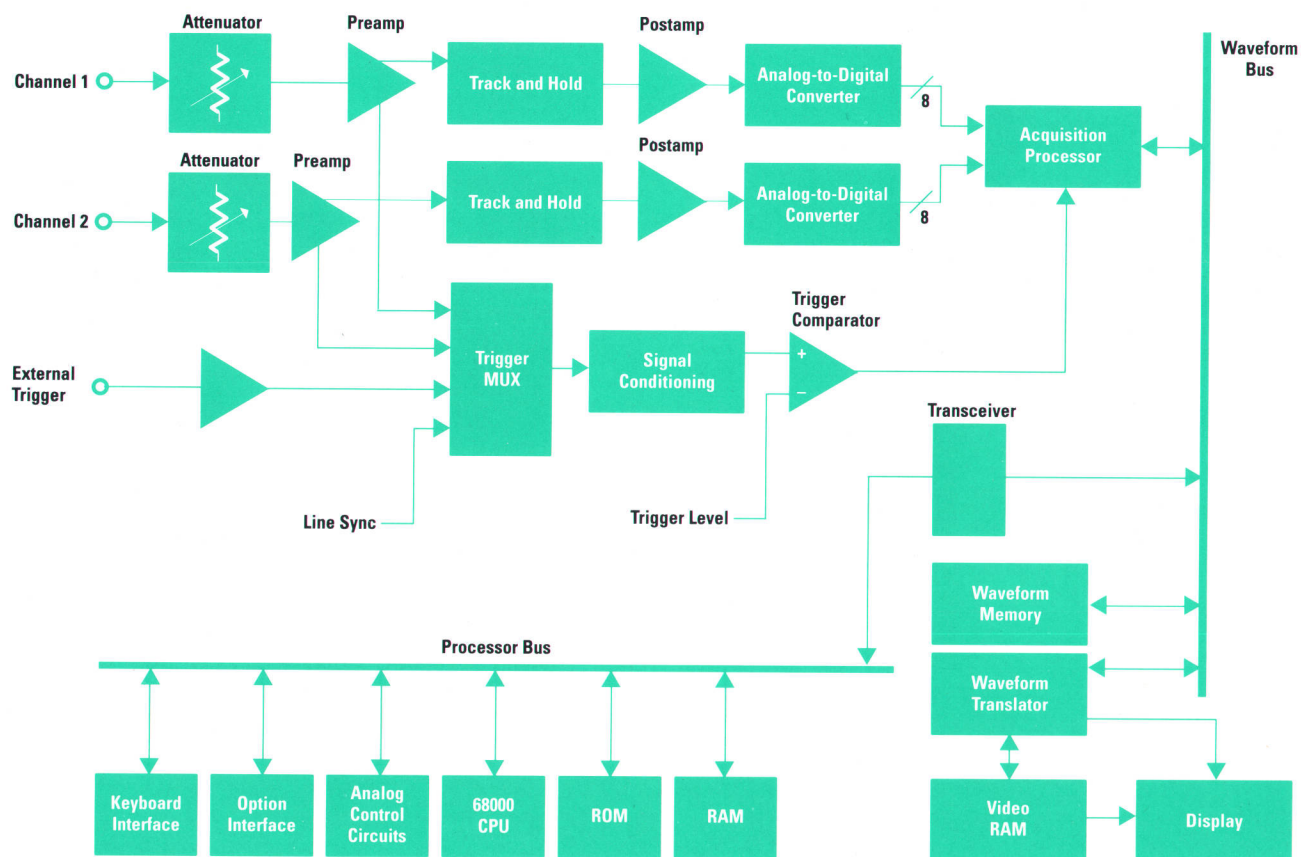


Fig. 5. Simplified block diagram of the HP 54600A two-channel oscilloscope.

of external test equipment is a pulse or function generator.

Acknowledgments

Recognition must go to the HP 54600 R&D team. Bob Kimura, Ken Gee, and Steve Roach each designed portions of the analog acquisition circuitry. Dan Timm, Matt

Holcomb, Yvonne Utzig, and Warren Tustin contributed to various portions of the digital hardware. James Kahkoska and Mark Schnaible designed the software and Rob Yergenson and Tim Figge produced the mechanical design. Don Henry did the industrial design. Dennis Weller led the original investigation of this product concept.

A High-Throughput Acquisition Architecture for a 100-MHz Digitizing Oscilloscope

Two custom integrated circuits offload functions from the system microprocessor to increase waveform throughput and give the HP 54600 digitizing oscilloscopes the "look and feel" of an analog oscilloscope.

by **Matthew S. Holcomb and Daniel P. Timm**

The key objective in the development of the HP 54600 Series digitizing oscilloscopes was to design a low-cost digitizing oscilloscope that has the "look and feel" of an analog oscilloscope. In other words, the new oscilloscopes were to have familiar controls and traces on the screen that look almost as good as analog oscilloscope traces, while maintaining the advantages of a digitizing architecture.

The acquisition design team investigated why the present digitizing oscilloscopes aren't widely used by analog oscilloscope users in troubleshooting environments. The result was not surprising: the display on digitizing oscilloscopes is notably slower. Dots are drawn on their raster displays no faster than about ten or twenty thousand points per second. In contrast, analog oscilloscopes update their screens at a rate of hundreds of thousands of *waveforms* per second.

On many digitizing oscilloscopes, waveform record lengths are only 500 points or so. These relatively short records sometimes lead to sparse, disconnected traces on the screen. Edges don't fill in, aliasing occurs, and familiar waveforms become unfamiliar. Users become uncertain of what is going on between samples.

Needless to say, one of the key design goals for the HP 54600 family was a significant increase in displayed signal quality. To achieve this, the designers selected absolute waveform throughput, from the probe tip to the CRT, as the fundamental performance goal. An update rate target of 1,000,000 points per second was chosen. To optimize

waveform throughput, the single-shot record size was increased to 2000 points. These records are interleaved into a final record size of 4000 points per channel, allowing eight times as many points per trace on the screen.

The theory was simple: if you just go fast enough, the display will start to approach the quality of an analog display. Multivalued signals (such as TV signals, eye diagrams, and modulated carriers) will look familiar. As more and more dots per second are drawn on the screen, intensity information will again be available. And all of the advantages of a digitizing architecture will be provided, such as the ability to view signals before the trigger, hard copy, storage, crystal-controlled time-base accuracy, and highly accurate automatic measurements. In some cases, the display is even better. Infrequent events, or very slow sweep speeds, are no longer dim, eliminating the need for viewing hoods. Single-shot events can be stored cleanly and clearly, without the blooming and smearing of an analog storage oscilloscope. In fact, at many of the common sweep speed settings, the HP 54600 processes more waveforms per second than even the fastest analog oscilloscope.

Previous Architecture

Fig. 1 shows a simplified topology for a traditional digitizing oscilloscope. Fig. 2 shows a task flow diagram for such an instrument. Notice that the system microprocessor controls everything. It unloads points from the acquisition memories, reads the time interpolation logic,

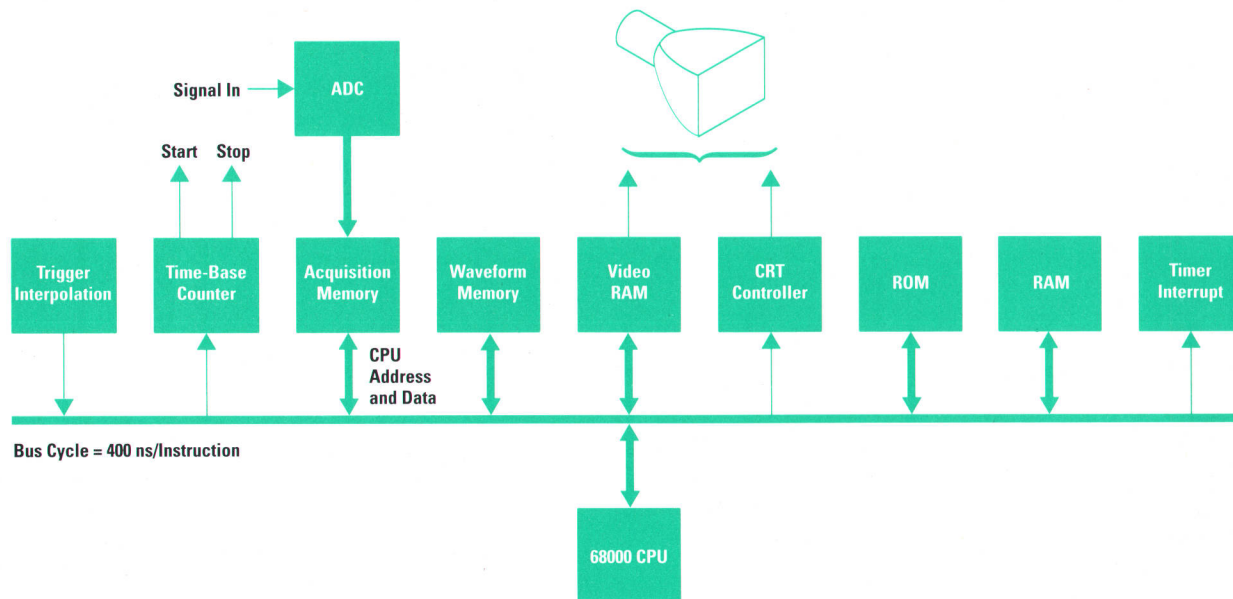


Fig. 1. Previous raster digital oscilloscope architecture, simplified.

plots and erases points on the screen, services the front panel and the HP-IB connector, moves markers, manages memory, performs measurements, and performs a variety of system management tasks. Because the system microprocessor is in the center of everything, everything slows down, especially when the system needs to be the fastest (during interactive adjustments). This serial operation gives rise to an irritatingly slow response. The keyboard feels sluggish, traces on the screen lag noticeably behind interactive adjustments, and even a small number of measurements slows the waveform throughput significantly.

HP 54600 Architecture

The differences between the traditional digitizing oscilloscope architecture and the architecture developed for the HP 54600 family are shown in Figs. 3 and 4. In contrast to the central microprocessor topology discussed above, the new topology separates the acquisition and display functions—that is, the oscilloscope functions—from the host processor, and places these functions under the control of two custom integrated circuits. The *acquisition processor* IC manages all of the data collection and placement mathematics. The *waveform translator* IC is responsible for all of the waveform imaging functions. This leaves the main processor free to provide very fast interaction between the user and the instrument.

Acquisition Processor

The acquisition processor is a custom integrated circuit, designed in HP's 1.0-micrometer CMOS-34 process. It contains roughly 200,000 FETs in an 84-pin PLCC package. This IC is essentially a specialized 20-MHz digital signal processor. It pipelines the samples from the analog-to-digital converter into an internal acquisition RAM, and later writes these points as time-ordered pairs into the external waveform RAM. It also contains the system trigger and time-base management circuitry.

Fig. 5 shows a logical block diagram for the acquisition processor. An external 80-MHz clock is brought onto the chip and divided down to 20 MHz. This signal is used as the system clock for the majority of the digital logic on the IC. The 20-MHz clock is also converted to quasi-ECL levels and sent off-chip as the sample clock for the track-and-hold circuits and the analog-to-digital converters.

The external analog-to-digital converters are always run at 20 megasamples per second (MSa/s), regardless of the sweep speed setting of the oscilloscope. The digitized samples are brought onto the chip continuously over an 8-bit TTL bus. When the acquisition processor is set up to look at two channels, it alternately enables the analog-to-digital converter output from each channel (still at a 20-MSa/s rate), yielding a sampling rate of 10 MSa/s per channel.

Internally, the vertical sampled data first passes through a dual-rank peak detection circuit. When the HP 54600 is set to a very slow sweep speed, there are many more samples available than can be stored into the final waveform record (2000 points). For example, since the analog-to-digital converters are always digitizing at a 20-MSa/s rate, the acquisition processor may only be storing, say, one sample out of every hundred, yielding an effective sample rate of $(20 \text{ MSa/s})/100 = 200 \text{ kSa/s}$.

During this condition (known as oversampling), the oscilloscope user may turn on peak detect mode, enabling the acquisition processor's peak detection logic. Instead of simply ignoring the oversampled points, the peak detection logic keeps track of the minimum and maximum analog-to-digital converter values and passes these pairs downstream. This allows the user to watch for narrow pulses (glitches) even at the slowest sweep speeds. With one channel enabled, the minimum detectable glitch width is 50 ns.

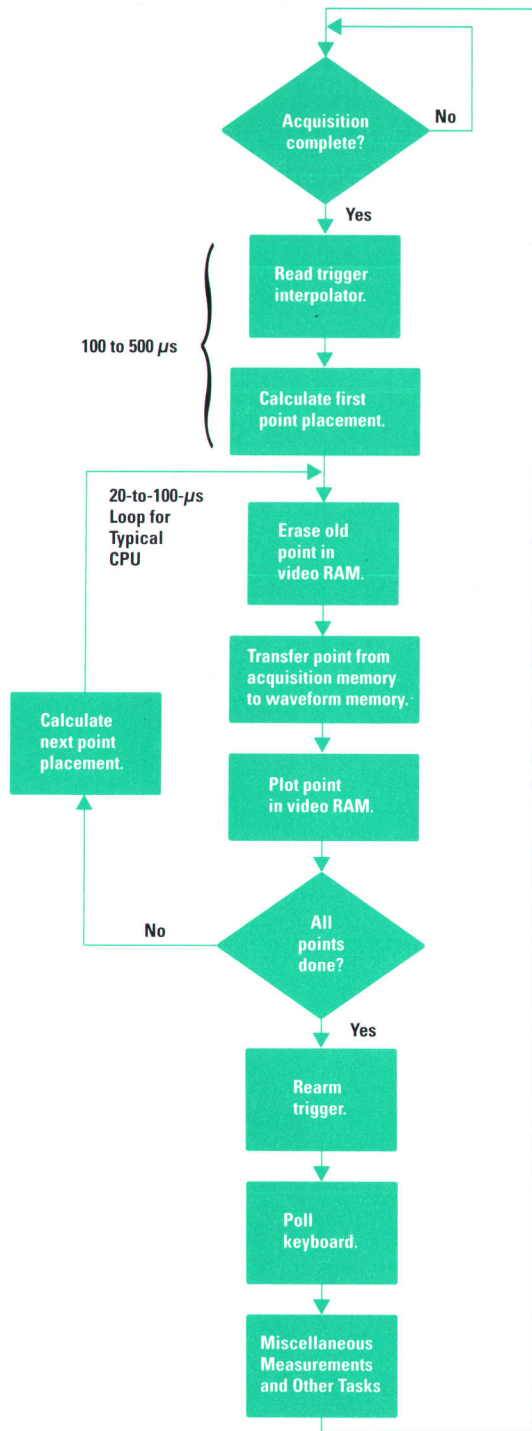


Fig. 2. Simplified task flows in the previous raster digital oscilloscope architecture.

The samples to be stored are put into an internal $2K \times 8$ -bit acquisition RAM. This memory is simply a circular scratchpad RAM. When an acquisition is started, the writing logic starts at address 0 and keeps on writing at the effective sample rate until it has collected all of the samples up to the right edge of the screen. In any two-channel mode, the data is aligned such that channel one's samples are stored at even addresses and channel two's samples are stored at odd addresses. When peak detection is enabled, the minimums are held in even addresses and the maximums are stored at odd addresses.

Meanwhile, the trigger block tracks the trigger input signal. An internal 28-bit counter generates the oscilloscope's trigger holdoff (holdoff by time) function. Considerable design effort went into the development of these front-end trigger flip-flops to minimize their mean time between synchronizer failures. Running continuously at 100 MHz, the synchronizer metastability failure rate is predicted to be on the order of once every 10,000 years.

Delay counters are used to implement the trigger delay function (pretrigger and post-trigger delay). Essentially, this allows the user to move the trigger point to the right (showing time before the trigger event) or to the left (showing time much later than the trigger event). The acquisition processor allows at least one screen of negative time (pretrigger delay), and at least 256 screens of positive delay.

As soon as the main acquisition logic is ready, it arms the main system trigger flip-flop. When the next trigger event is detected, a time interval counter external to the chip measures the time between the trigger event and the second rising 80-MHz edge. Using this value and the phase of the 20-MHz sample clock with respect to the 80-MHz clock, an internal six-stage microcoded processor calculates the effective screen address of where the first point stored into the acquisition RAM should be placed on the screen.

After this address has been calculated, the postprocessing logic starts unloading data as parallel time/voltage pairs simultaneously from the RAM and the address calculation block. Access to the acquisition RAM is internally arbitrated: the writes have unconditional priority. As soon as the address calculation block signals that it knows the screen address of the next available point in acquisition RAM, the postprocessing logic unloads that point at the next available clock edge. This dual-access behavior allows the sampled trace to be displayed before the entire waveform has been acquired.

In the postprocessing section, either or both channels can be digitally inverted (implementing the channel invert coupling option). Additionally, any channel can be shifted left or right up to three bits. This allows digital magnification of the analog-to-digital converter samples (used by the oscilloscope in 5-mV/div and 2-mV/div settings) and split screen displays (used in the delayed sweep mode).

Finally, the collected data is written asynchronously into the external waveform memory at about a 1-MHz rate. The vertical samples are written to the waveform data bus, while the horizontal (time) information is written to the address bus. Both buses can be configured to control exactly where and how the data is written into memory.

The host 68000 can program the acquisition processor by writing directly into its internal registers, using a traditional chip enable, read/write, register select protocol. Similarly, the 68000 can determine the acquisition status by reading a series of status registers. After the acquisition processor has been set up, the host can then issue various commands to the IC by writing to a write-only section of address space. There are commands to start an acquisition, stop, calibrate the time interpolator, and scramble (dither) all sample clocks. Additionally, the host

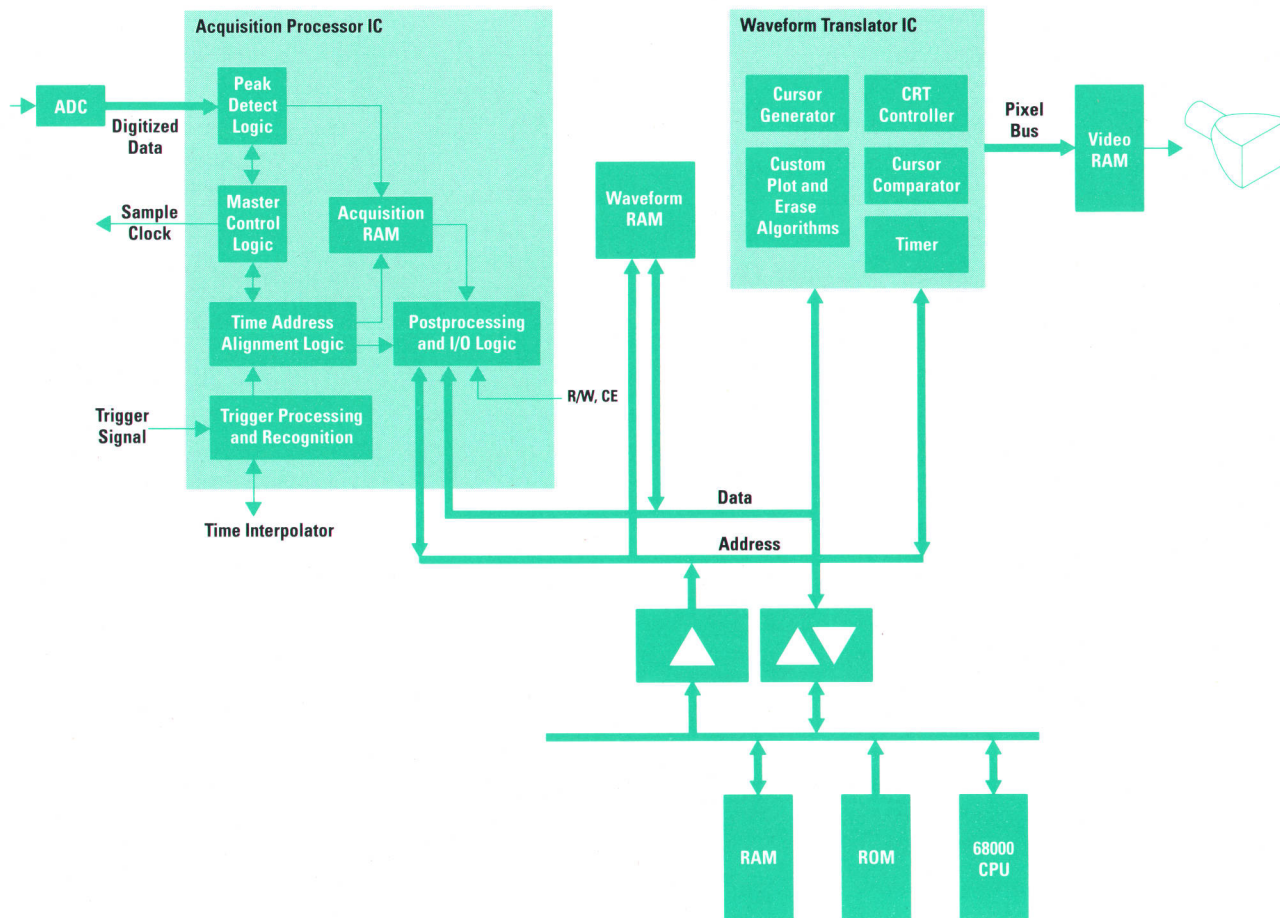


Fig. 3. New raster digital oscilloscope architecture.

can enable interrupts, allowing the acquisition IC to signal the host when an event has occurred, such as the completion of an acquisition.

The acquisition processor can also be used to acquire Y-versus-X data. In this XY mode, channels one and two are collected as (X,Y) pairs and passed directly from the input to the output section of the IC, bypassing the intermediate storage in the internal RAM. These points are simply written into incrementing locations in the waveform RAM. Thus the waveform RAM address no longer bears any time information. Since XY mode is inherently untriggered, the acquisition processor samples the trigger clock input to provide a simple Z-axis blanking (XYZ mode) capability. Only samples that occur when the trigger input (Z signal) is at a logic high level are passed from input to output.

The acquisition processor also controls a variety of the peripheral acquisition functions. Most significant, it controls two powerful clock dithering algorithms, which greatly reduce the potential for aliasing at low and medium sweep speed settings. Additionally, it contains the necessary logic to keep the external time interpolation circuitry calibrated (the HP 54600 time interpolator is continuously calibrated about once every 20 ms). When the oscilloscope isn't getting any triggers, the acquisition processor can also generate its own triggers and run in autotrigger mode. It contains a dedicated trigger counter and special peak detection modes for performing very

fast autoscale and autolevel functions. It also contains a good deal of self-test and testability logic and circuitry.

Physically, the bulk of the IC consists of two standard library data paths, controlled by two programmable logic arrays (PLAs). One data path controls the vertical signal processing, while the other controls the horizontal signal processing. About a third of the the chip area is dedicated to the 2K×8-bit acquisition RAM (see Fig. 6). The trigger section consists of about 100 gates of custom 100-MHz synchronizing trigger flip-flops and clock generation logic.

The design of the chip was remarkably automated. ASCII-text style schematics were created, showing connectivity between the low-level data path cells. The logic control was similarly described in a high-level PLA description language known as PLADO. From these two sources, a behavioral model was generated, which was used extensively to verify that the chip actually did what the oscilloscope needed.

From this very same source, an awk script was written to generate artwork automatically. These tools allowed logical changes to be made to the design while the physical artwork was kept up to date automatically and transparently. Switch-level models were then extracted from the artwork, simulated hierarchically, and compared with the behavioral description. All of these tools allowed a very rapid design-simulate-modify design cycle.

Waveform Translator

The waveform translator, implemented in a 1.5-micrometer CMOS process, is a gate array with a main clock speed of 40 MHz. The gate array has 6500 gates of custom logic packaged in an 84-pin PLCC package. The main function of the waveform translator is to take the time/voltage pairs from waveform memory and turn on the corresponding pixels on the display. The gate array also includes other features that reduce the number of parts in the microprocessor system and the number of microprocessor tasks on a system level. This results in a very responsive display system, while keeping the total system cost low.

The HP 54600 family uses a seven-inch-diagonal monochrome raster-scan display organized as 304 vertical by

512 horizontal pixels. The area used for the oscilloscope graticule and waveform display measures 256 vertical pixels by 500 horizontal pixels, providing eight-bit vertical resolution and nine-bit horizontal resolution. The waveform translator acts as the display controller and provides the vertical sync, horizontal sync, and video signals (all TTL levels) to the display.

The waveform translator has control over a $256K \times 4$ -bit video DRAM for image capture. Each nibble of the video RAM is divided into two pixels of full-bright and two pixels of half-bright information representing two video image planes. Live waveforms are plotted into the full-bright plane at a rate of approximately 1.75 million points per second (7.5 waveforms/video frame). The half-bright plane is used for things such as graticules, text, and

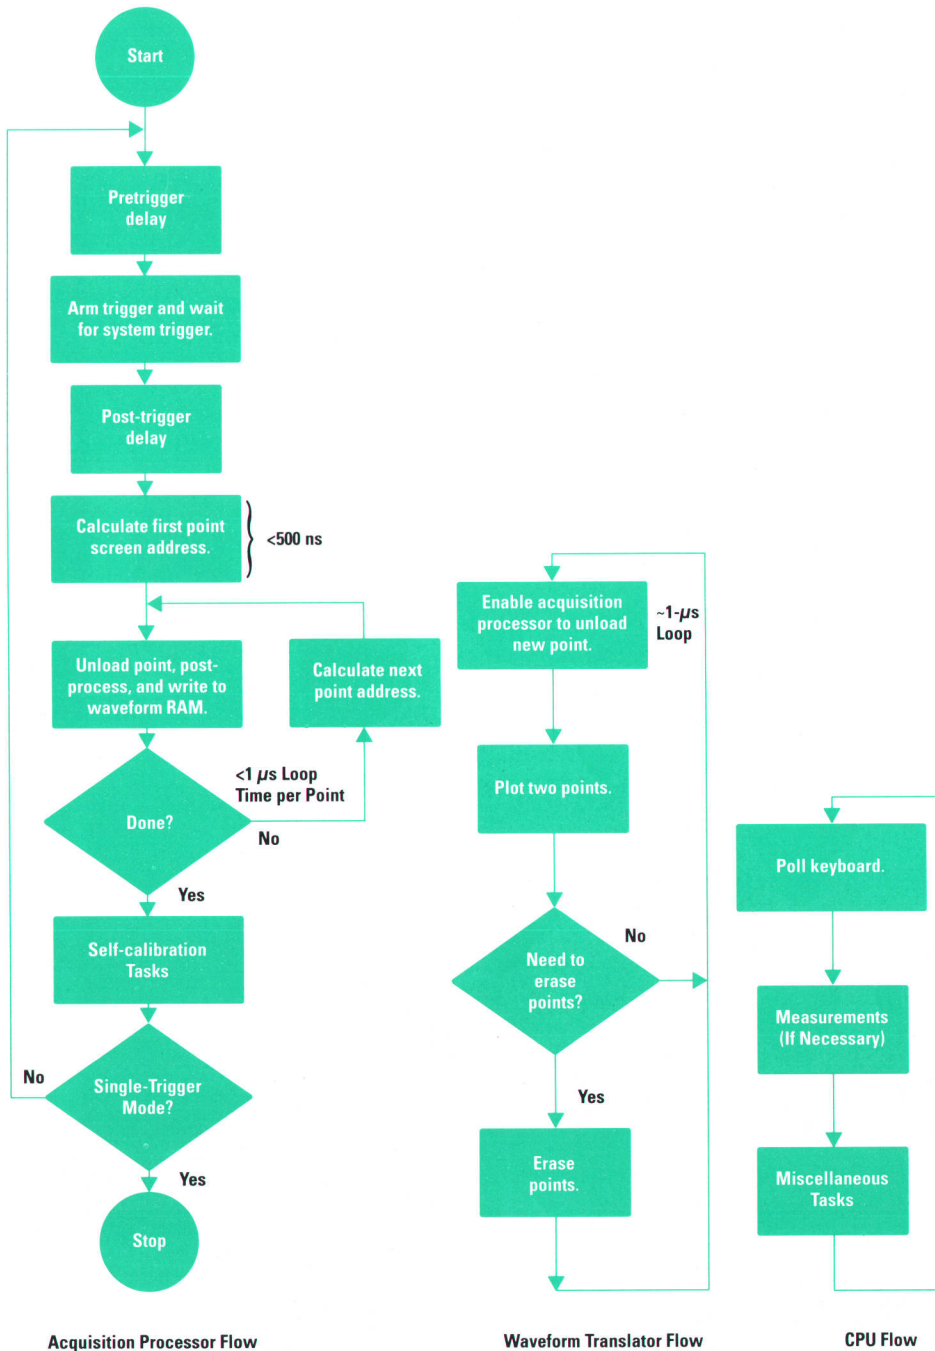


Fig. 4. Simplified task flows in the new raster digital oscilloscope architecture.

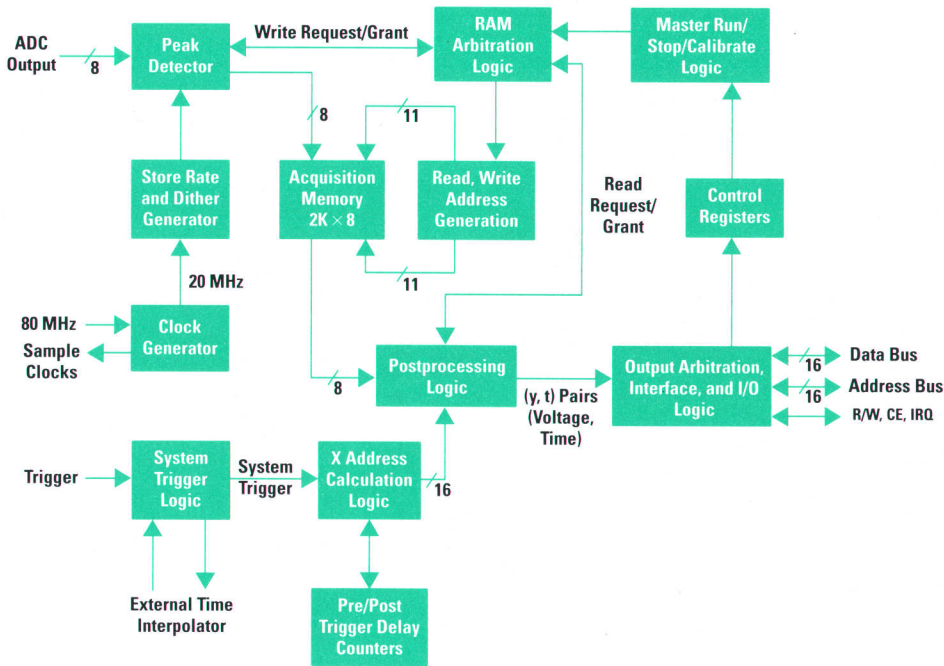


Fig. 5. Logical block diagram of the acquisition processor.

stored waveforms. All accesses to and from the video RAM are controlled by the waveform translator.

The waveform translator handles the waveform bus arbitration (Fig. 7), generating the handshake signals required to unload the acquisition processor, to allow microprocessor accesses, and to read waveform points for plotting. For each waveform bus cycle, there is one

unload request to the acquisition processor for new data. This is followed by two reads. In each read, two points are read from the waveform RAM, translated, and then plotted into the video RAM (see Fig. 7b). The waveform RAM is organized into eight 4K-byte waveform records. Each record is a time-ordered list of voltage data values,

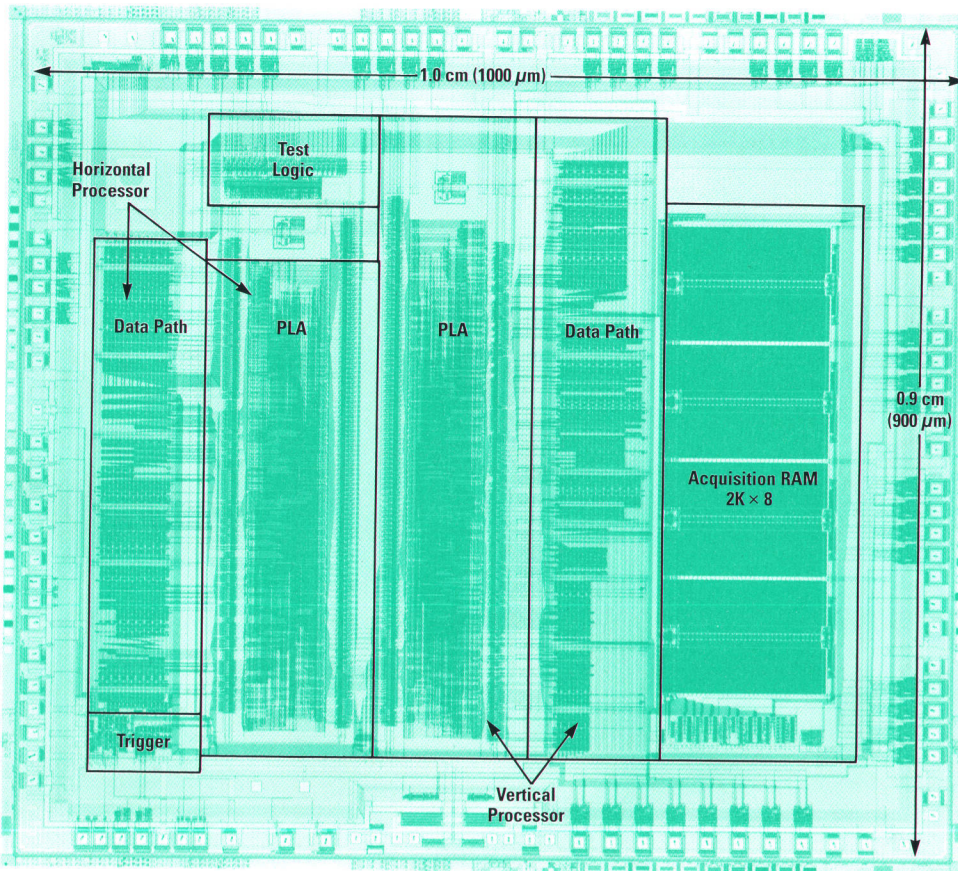


Fig. 6. Acquisition processor chip layout.

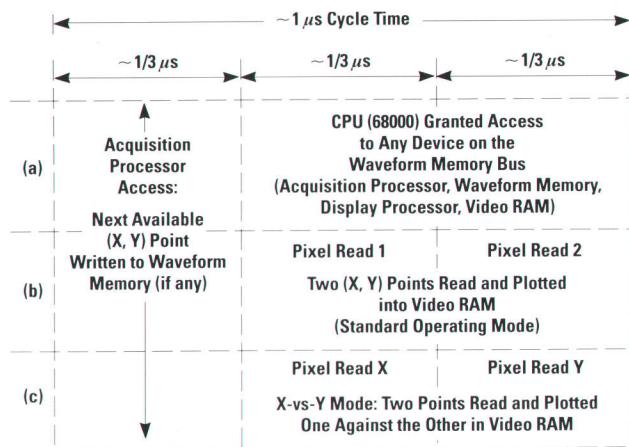


Fig. 7. Waveform memory bus arbitration.

with the smallest address representing the left edge of the graticule.

In its normal operating mode, the oscilloscope plots voltage versus time. In this mode, the address of the waveform RAM represents the horizontal coordinate, while the data at that address represents the vertical coordinate of the pixel memory bit to set in the video memory. The waveform translator plots 4000-point records into a 500 (horizontal) by 256 (vertical) pixel array. This overmapping results in eight voltage points per horizontal pixel column, which simultaneously improves the sample density on the waveform and allows all 4000 points to be displayed at once.

Y-versus-X (channel versus channel) waveforms are plotted in a somewhat different manner. Two data points (one from each channel) are required to form the horizontal and vertical coordinates. These two points are supplied by the two consecutive reads in the waveform bus cycle (Fig. 7c). Once the horizontal and vertical coordinates are read, the setting of a pixel in the video RAM can occur. Both values are checked to ensure their validity. The X data has eight bits of information, which is mapped into a nine-bit-resolution array, leaving the least-significant horizontal bit undefined. The waveform translator then adds digital noise for this bit in a pseudo-random fashion. This fills every other pixel column that would have formerly been empty, resulting in a brighter, connected, more appealing image.

A fundamentally new display mode called *autostore* displays the live data in full brightness while maintaining previous waveforms in half brightness. Similar to infinite persistence mode on previous-generation products, autostore lets the user view the accumulated samples without obscuring the most recent waveform. In autostore mode, the waveform translator plots the live waveform into both the full-bright and half-bright planes. While the full-bright image gets erased every frame, the half-bright image does not. It accumulates, providing a history of all the captured waveforms in a discernibly different intensity. This contrasts with previous digital oscilloscopes, which display both the most recent and all previous waveforms in the same intensity, making it difficult if not impossible to identify the live waveform. This simultaneous plotting

is possible because both the full-bright and half-bright bits reside in the same nibble of the video RAM.

The waveform translator performs several functions that the system microprocessor has performed in previous architectures. In addition to plotting waveforms, it erases images, generates multiple cursors, and provides the microprocessor with several special access modes to the video RAM. The erase cycle occurs periodically during each video frame, allowing the erasure of all pixel locations once per frame. Full-bright and half-bright images can be individually enabled for erasure. Live waveform erasure is always enabled. This relieves the microprocessor of the (rather significant) overhead of tracking all points on screen. Since the waveform RAM retains the last waveform indefinitely, and since the waveform translator can reconstruct the full-bright waveform image several times every video frame, a flicker-free image appears onscreen. Half-bright image erasure is required during a number of setting changes. For the microprocessor to erase the graticule area directly would take hundreds of milliseconds. With the waveform translator performing the erase, the microprocessor simply needs to enable the half-bright erasing and wait one video frame (16.6 ms), speeding up these operations considerably.

Cursor generation in previous digital oscilloscopes is a difficult and time-consuming task. The microprocessor had to compute the location of the cursor line and plot the dashed line across the display, individually turning the appropriate pixels on and off. In the HP 54600 family, the waveform translator generates up to eleven line cursors, six vertical and five horizontal, in hardware. These cursors are used as a trigger level indicator, as delayed sweep indicators, and as measurement cursors. The lines can be solid, short-dashed, or long-dashed so that the different cursors can be identified. Additionally, the six vertical markers can be half screen height, in either the upper or lower half of the graticule. The cursor generation is made possible by having the CRT controller on-chip. The CRT controller counters track the raster beam position, allowing the cursor comparators to generate a serial video data stream that is 0Red off-chip with the full-bright data coming from the video RAM.

Several specialized modes for microprocessor access to the video RAM have been incorporated to speed software algorithms. For example, there are special modes for character generation and image drawing. There is even a mode that allows the microprocessor to address the nibble-wide video RAM as a byte.

To reduce system cost, board space, and power requirements, as much functionality was placed on-chip as the package I/O pins allowed. This includes the CRT controller counter system, a 16-bit programmable timer, and the microprocessor handshake and strobe generation. The CRT counter system has eight counters, four for the horizontal sync system clocked by the dot clock and four for the vertical sync system clocked by the horizontal sync pulse. The programmable timer runs at approximately 1 MHz and is used by the microprocessor to time various software activities.

(continued on page 19)

Sample Rate and Display Rate in Digitizing Oscilloscopes

There is considerable confusion about how sample rate relates to typical oscilloscope measurements. Most oscilloscope users understand that for truly single-shot events, there is no substitute for a high sample rate. When an event is single-shot or repeats so slowly that it is impractical to wait for subsequent occurrences, the oscilloscope only gets one chance (one trigger) to acquire the waveform. The sampling theorem states that to be able to reconstruct a baseband signal it must be sampled at greater than twice the signal bandwidth.¹

$$f_s > 2BW.$$

Practical limitations require that an oscilloscope sample at an even higher rate, typically four times the bandwidth, to maintain good pulse response.² Therefore, single-shot measurements require a very fast, expensive analog-to-digital converter (ADC).

If a signal is a repetitive waveform, the oscilloscope has multiple chances to acquire and digitize the waveform. The entire waveform content does not have to be acquired on a single trigger, but can be collected from multiple triggers. Thus, a slower, more economical ADC can be used. The type of repetitive sampling used in the HP 54600 family oscilloscopes is called *random repetitive sampling*. Fig. 1 shows how random repetitive sampling works. The waveform is constantly sampled at 20 MSA/s. At fast time/division settings, this sample rate will produce several samples on each occurrence of a waveform, but will not sample fast enough to satisfy the sampling theorem for 100-MHz bandwidth. As the waveform repeats, more samples are acquired at different points on the waveform and the displayed waveform fills in. With a high enough trigger rate and display update rate, this happens so fast that the oscilloscope user perceives it as instantaneous. The sample rate must not be correlated with the frequency of the waveform or the samples will occur consistently at the same locations on the waveform, preventing the displayed waveform from filling in. To prevent this, the sample rate is dithered slightly. Hence the term random repetitive sampling. The spacing of the collected samples is much closer than the original sample rate, resulting in a new *effective sample rate* that satisfies the requirements of the sampling theorem.

Random repetitive sampling is the clear choice for an economical digitizing oscilloscope that competes in the 100-MHz analog oscilloscope market. Analog oscilloscopes are good for viewing repetitive waveforms but are not very effective at measuring single-shot events. Therefore, the repetitive-only nature of random repetitive sampling allows the oscilloscope to measure the same signals that an analog oscilloscope can. In addition, since the digitizing oscilloscope has a sample rate of 20 MSA/s, low-frequency single-shot measurements can be made, but not to the full bandwidth of the oscilloscope. HP specifies the single-shot bandwidth of the HP 54600 family at one tenth of the sample rate, which provides ten points per period of the highest frequency.

Display Rate

If sample rate is a confusing concept to digitizing oscilloscope users, the display update rate is even more confusing. Sometimes the oscilloscope user should be concerned about the sample rate and sometimes display update rate is the real issue. Consider the waveforms shown in Fig. 2. The first waveform, a single-shot pulse, occurs only once and must be measured with sufficient sample rate to capture the bandwidth of interest. The next waveform in Fig. 2 is a repetitive waveform that repeats exactly each time, never varying. As previously discussed,




Fig. 1. Random repetitive sampling.

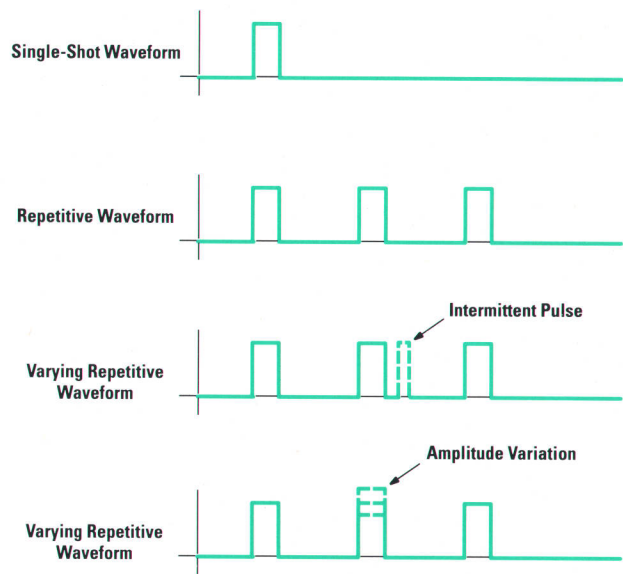


Fig. 2. Waveform types.

this type of waveform can be acquired using random repetitive sampling. Oscilloscope users expect the display to update instantaneously when they connect a signal, so the display update rate needs to be fast mostly in human terms (i.e., less than 100 ms or so). Fast display update rate is not required to track any changes in the waveform because, by definition, the waveform never changes.

A more interesting case is when the waveform is basically repetitive, but exhibits a change once in a while. Two examples of such *varying repetitive waveforms* are shown in Fig. 2. Here the waveform basically repeats over time, but does change occasionally. In one waveform the occasional change is an intermittent pulse (or glitch) while in the other waveform it is a variation in pulse amplitude. An occasional change in the waveform is often of interest to the oscilloscope user, so it is desirable that these temporary changes be viewable. On an analog oscilloscope, these temporary excursions might be easily viewable, invisible, or somewhere in between depending on how often they occur. (Infrequently occurring events do not light the phosphor on a cathode ray tube often enough to be viewed.) On a digitizing oscilloscope, the display update rate must be fast enough to capture and display the occasional change in the waveform.

Consider the situation shown in Fig. 3. The waveform is a varying repetitive signal that has an intermittent pulse occurring in it (on average) every t_{rep} seconds. (An intermittent pulse is used as an example, but the same principles apply to other signal variations.) Samples are taken for a period of time t_a , the acquisition time (the same as the time duration displayed on the oscilloscope) and then a dead time t_d occurs as the oscilloscope processes those samples. Processing of the samples includes reading the data from memory, comparing the samples in time to when the trigger occurred, and writing the data to the display memory. During this dead time, variations in the waveform may be missed.

Define the oscilloscope duty cycle k_d as the ratio of the sampling time to the total time, giving a measure of what portion of time the oscilloscope is actually sampling the waveform:

$$k_d = \frac{t_a}{t_d + t_a}$$

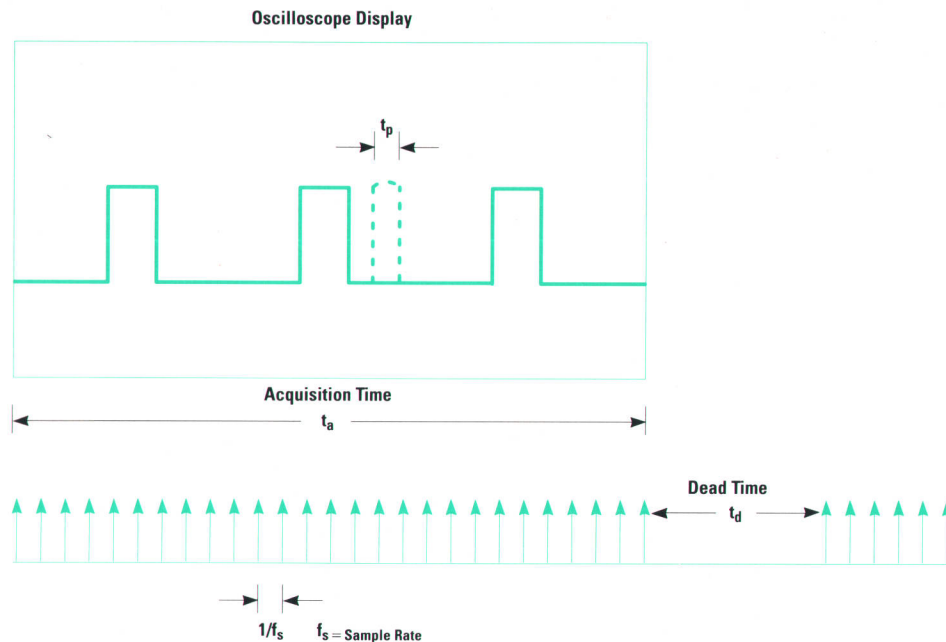


Fig. 3. Acquisition time, sample rate, and dead time.

The product of the oscilloscope duty cycle and the sample rate is equal to the display rate:

$$f_d = k_d f_s$$

The display rate f_d is the number of samples processed and displayed per second.

To maximize the probability of sampling an intermittent event, the oscilloscope should have a high display rate. High sample rate alone is not sufficient. The oscilloscope duty cycle can be thought of as a measure of how efficiently the oscilloscope circuitry turns samples into displayed points.

It is important to note the assumptions about the measurement situation described above. The oscilloscope was triggered on a repetitive waveform that had an infrequent glitch in it. This scenario fits oscilloscope applications in which the user

doesn't necessarily know that a glitch is occurring—the oscilloscope is being used to find that out. In such a case, display rate is critical. However, some digitizing oscilloscopes such as the HP 54500 family provide a means of triggering on such a glitch. If the user suspects that a glitch is occurring and triggers on it, display update rate is much less an issue. In fact, the glitch generally can be captured and measured more accurately by triggering on it.

References

1. A.V. Oppenheim and A.S. Willsky, *Signals and Systems*, Prentice-Hall, Inc., 1983.
2. *Voltage and Time Resolution in Digitizing Oscilloscopes*, Application Note 348, Publication number 5954-2652, Hewlett Packard Company, November 1986.

Robert A. Witte
R&D Project Manager
Colorado Springs Division

Results

In the final design, the displayed waveforms look surprisingly lively. Qualitatively, the traces on the screen respond instantaneously to front-panel changes, as well as to changes in the input signal. Fast edges tend to fill in quickly, especially when the autostore mode is enabled. Multivalued and modulated signals are easily recognizable, and rarely aliased. Unfortunately, however, the discrete samples are often identifiable, giving an apparent granularity, or graininess, to the traces. And even though the dot density on the screen now shows rudimentary intensity information, a few users will still prefer an analog display for these signals.

In many situations, the display looks significantly better than an analog display—and certainly better than conventional digitizing oscilloscopes. Low-duty-cycle signals and slow sweep speeds are no longer dimmed by the fading phosphor. The waveform image is maintained on the screen indefinitely (and in internal memory for analysis) until it is overwritten by a later acquisition. Additionally, even at fast sweep speeds, the 60-Hz refresh rate of the raster display adds a secondary visual enhancement. Any

displayed point will light up on the screen for at least 16.6 ms (one frame). Since over eight 2000-point waveforms are displayed every frame, this effectively extends the visual record size from 4000 points to over 16,000 points. This adds brightness to dim traces and depth to multivalued waveforms.

Quantitatively, the screen update rate for the HP 54600 oscilloscopes is shown in Fig. 8. Fig. 8a shows the number of samples displayed on the screen every second and the number of triggers processed per second. The roll-off toward the left edge of the graph is a result of slowing down the effective sample rate for very slow sweep speeds. At the very fastest sweep speeds, the update rate becomes limited by the acquisition rearm time (similar to an analog oscilloscope's retrace and rearm time).

Fig. 8b shows the relative waveform throughput of the HP 54600 versus typical analog and other digitizing oscilloscopes. For this graph, waveform throughput is defined as full records for digitizing oscilloscopes (2000 points for the HP 54600) and as sweeps per second for

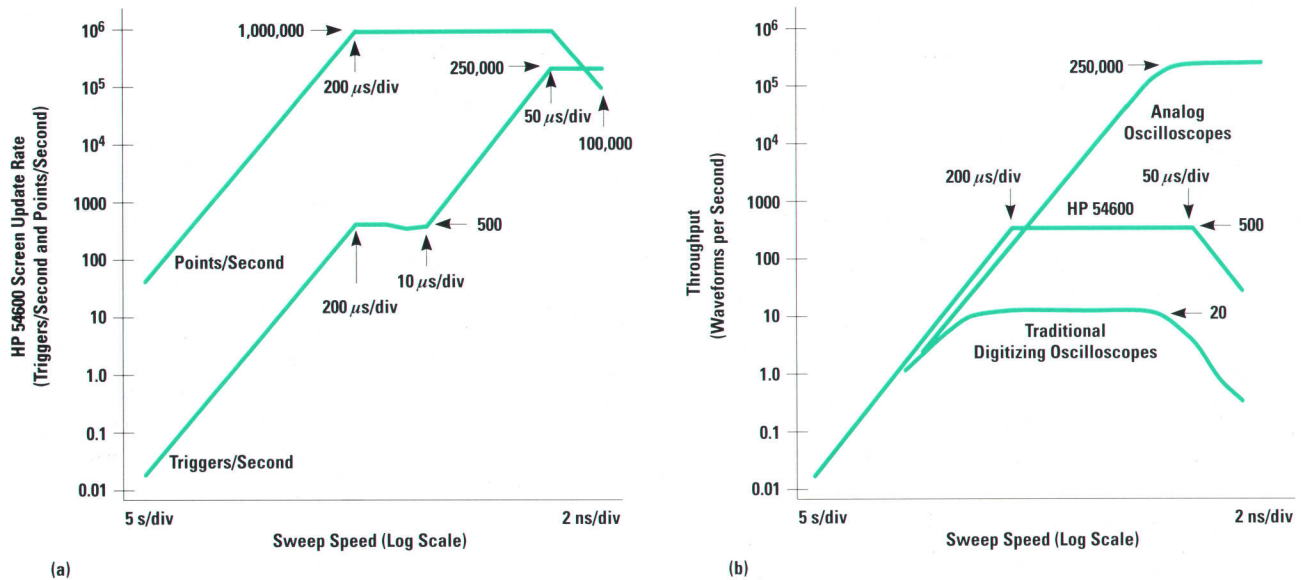


Fig. 8. (a) HP 54600 screen update rate. (b) HP 54600 throughput compared to analog oscilloscopes and traditional digitizing oscilloscopes.

analog oscilloscopes. For all of the intermediate and slow sweep speeds, the HP 54600 has the highest throughput of any other oscilloscope. (Analog oscilloscopes are limited by a rather significant 25% retrace and rearm time). As the sweep speed increases, the analog oscilloscopes begin displaying waveforms significantly faster than the digital oscilloscopes. However, at these sweep speeds, each trace becomes increasingly dim. Whereas it might take the HP 54600 about 500 triggers for the waveform to be displayed (at one point per pixel column), it probably takes the analog oscilloscope just as many triggers to brighten the trace on the screen. Notice

in Fig. 8a that at these fast sweep speeds, analog oscilloscopes process triggers at a very similar rate to the HP 54600.

Summary

The architecture developed for the HP 54600 family sets a new level of performance for display update rates in digitizing oscilloscopes. By offloading the acquisition and plotting algorithms from the system processor, the overall responsiveness—both to the front panel and to the HP-IB interface—has been improved.

A Fast, Built-In Test System for Oscilloscope Manufacturing

Following a verification strategy instead of a screening or characterization strategy, a special module was designed to replace the computer input/output option module of the HP 54600 Series oscilloscopes. The resulting test system has reduced both equipment costs and test times to one tenth those of previous test systems.

by **Stuart O. Hall and Jay A. Alexander**

The HP 54600A and 54601A general-purpose oscilloscopes have the lowest manufacturing cost of any oscilloscopes in HP's history. A large contributor to this is a test strategy that minimizes the cost of calibration and verification.

Traditionally, digital storage oscilloscopes have been tested with a system that evolved from a prototype characterization system or an environmental qualification system. The R&D team described a set of parameters that they wanted to see tested in changing environmental conditions, and the test engineer developed a system that tested those parameters as accurately as possible. This process resulted in a long list of parameters and a set of test algorithms and equipment optimized for accuracy. This test implementation is required when attempting to understand every aspect of a new product design.

The test strategy for the HP 54600 family of oscilloscopes was developed to minimize the test process cycle time, using the least expensive set of equipment possible (see "Verification Strategy," page 22). Developing this strategy required a complete understanding of the hardware of the product, a commitment on the part of the design team to design margins into the product, and a willingness to recognize the difference between testing to characterize a design and testing to verify a product. Here we will describe this test strategy, from prototype characterization to production testing, and discuss the benefits of the resulting system.

Prototype Characterization

Like most new products, the HP 54600 Series oscilloscopes required extensive performance characterization at various points in the design cycle. The characterization test system developed for this was used both for characterizing prototype units and for qualifying the product to HP environmental specifications. Data from this system was used to influence the design team to increase margins on some specifications. In other cases the system verified that a circuit had more margin than the designer expected.

A goal was to characterize prototype instruments early in the design process. This had several positive influences

on the success of the project. Because the characterization system was developed as an HP-IB-controlled automated test system, it was necessary to install the HP-IB parser software at an earlier stage than usual. This caused the parser code to be developed as an integral part of the instrument, rather than as an additional feature. The characterization process uncovered hardware and software problems that might otherwise have been missed until a point in the project when it would have been much more expensive to correct them. The system was also used to test new prototype test algorithms and to corroborate the techniques used in the production test process described later in this article. It was also recognized that the characterization test system would be required as an auditing system in production for ongoing verification of key assumptions regarding parameter relationships.

The standard tests performed with this system were developed using the traditional process described above. The system uses the following equipment:

- HP 8131A programmable pulse generator
- HP 8656B synthesized signal generator
- HP 3458A programmable system multimeter
- HP 437B programmable power meter
- HP 4262A LCR meter
- HP 3235A programmable multiplexer
- HP 9000 Model 360 computer.

Thirty-one different types of tests are performed, generating more than 200 test parameters for the oscilloscope. (The term "parameter", as used in this article, means the result of a given test under one specific set of test conditions. Some tests have many different sets of conditions. For example, a trigger test that measures performance for two different input signal frequencies, on eight different vertical ranges, on two channels, will generate $2 \times 2 \times 8$ or 32 separate result parameters.) By the completion of the production prototype phase, the characterization system had been used to verify that each of the more than 200 test parameters had specification margins of at least four standard deviations (4σ).

Verification Strategy

Traditional oscilloscope tests typically require one to two hours of test time and capital equipment costing many thousands of dollars. Even if the tests are largely automated and thus free of excessive labor, there is still a cost penalty associated with the cycle time for each unit. Test times impact capital outlays—slower test times mean that more stations must be purchased to support even moderate production volumes. Multiple stations in turn drive up floor space and maintenance costs. All of these issues ultimately translate into higher costs for customers. For a high-volume product family such as the HP 54600 Series oscilloscopes, process modeling showed that between six and ten stations would have to be purchased to meet capacity requirements if a traditional test strategy were employed. Concern for delivering the best value to the customer and a desire for a simpler process led to a new strategy focused on reducing unit test time. Our goal for this strategy was 10 minutes per unit. This figure would include the time for all adjustments, calibrations, self-tests, and final tests, and represented a tenfold improvement over the predecessor product, the HP 54501A.

We began with the belief that the typical 200 to 300 test parameters are not all independent. If a small set of independent parameters could be identified, they could then be tested and used to estimate the performances of the other, dependent parameters. We labeled this idea the minimum set method. Others have developed similar techniques called predictive subset testing¹ and QR factorization.² Our approach was to combine knowledge about which parameters one might expect to be related (gained by in-depth scrutiny of the product hardware design) with empirical correlation studies performed during prototype characterization.

Production Test

The production test system was developed according to three fundamental concepts. First, the system would be controlled by the device under test, or DUT (the oscilloscope currently being tested) to alleviate the cost and throughput disadvantages of an HP-IB controller system. Second, the system would use only two signal sources: a high-frequency source and a precision low-frequency source. Third, the test process would be designed to verify that the product is statistically normal. It is important to recognize the distinction between this last concept and the process of ensuring that all specifications are met. Of course, for this method to be sufficient to verify production instruments, all normal instruments must meet specification limits. This is the fundamental reason to include margins in the design of a product.

The sequence of test operations carried out in production is as follows:

1. Functional self-test: comprehensive test for basic functionality
2. Software calibration: self adjustment of dc parameters
3. Hardware adjustments: flatness and pulse response adjustments
4. Parametric final test: final product verification before assembly.

Production Calibration Module

The HP 54600 Series' optional computer interfaces gave the test system designers a method of developing custom hardware and software for use in production test. This custom hardware and software is contained in the production calibration module. Fig. 1 shows how the production calibration module fits into the overall production test system architecture.

Key to the success of this method were the performance margins designed into the product. These margins allowed us to streamline the parameter set without compromising quality, and to move from a screening philosophy to a verification philosophy.

In addition to reducing the number of test parameters, the team focused on determining the most efficient ways to measure the parameters actually tested. This effort led to the use of new measurement algorithms, such as the FFT techniques described in the article on page 29. Combining the new algorithms with the reduced parameter set, we realized that a significant reduction in the variety of test stimuli required was possible. Simplifying the stimuli allowed us to build it all in, further lowering capital costs. Finally, executing all tests in firmware eliminates time-consuming I/O operations and makes the system almost entirely self-contained.

Details of the HP 54600 test strategy development and its results are described in the accompanying article. While this strategy is not universally applicable, it does represent a good example of the kind of thinking that can contribute to manufacturing competitiveness, and we hope to leverage the knowledge gained on these products to other Hewlett-Packard programs.

References

1. J.B. Brockman and S.W. Director, "Predictive Subset Testing: Optimizing IC Parametric Performance Testing for Quality, Cost, and Yield," *IEEE Transactions on Semiconductor Manufacturing*, Vol. 2, no. 3, August 1989, pp. 104-113.
2. T.M. Souders and G.N. Stenbakken, "Cutting the High Cost of Testing," *IEEE Spectrum*, March 1991, pp. 48-51.

Hardware. A large portion of the 68000 microprocessor bus is accessible at a 50-pin connector on the back panel of the oscilloscope. The customer can connect one of three interface modules to this connector to define the product's I/O personality. The three interface modules are RS-232, HP-IB, and Centronics.

The HP 54600 production calibration module was developed using the HP 54650A HP-IB interface module as a platform. The production calibration module consists of an HP-IB module with an additional 32K bytes of nonvolatile RAM (64K bytes total). A fixed-amplitude, variable-frequency pulse generator was added for high-frequency testing and pulse response adjustments (Fig. 2). The

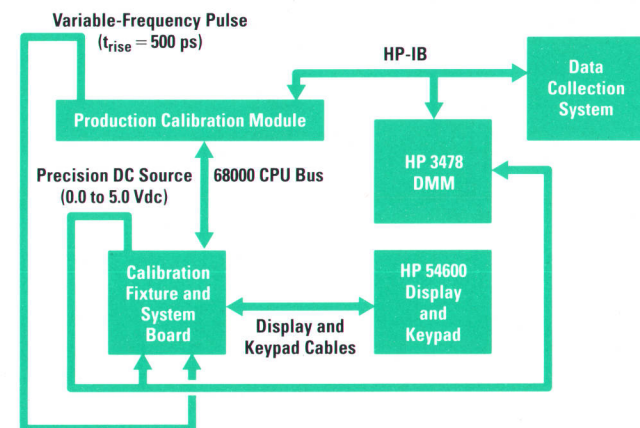


Fig. 1. HP 54600 Series oscilloscope production test system, showing the production calibration module.

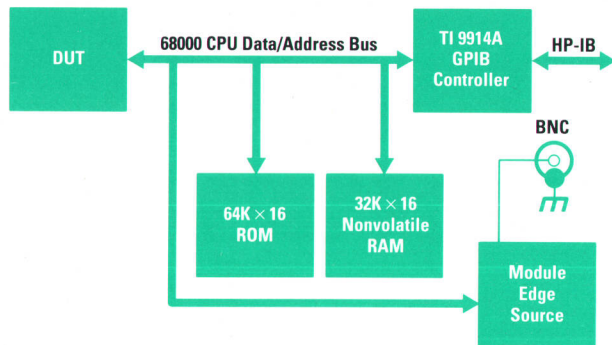


Fig. 2. Block diagram of the production calibration module.

500-picosecond rise time of this pulse generator is sufficiently fast to adjust and measure the high-frequency response of the 100-MHz oscilloscopes. The effect of using a 500-ps source to measure the rise time of a 100-MHz (3.5-ns rise time) system can be estimated as follows:

$$\begin{aligned}
 t_{\text{err}} &= \sqrt{t_{\text{system}}^2 + t_{\text{source}}^2} - t_{\text{system}} \\
 &= \sqrt{3.5^2 + 0.5^2} - 3.5 \\
 &= 40 \text{ ps} .
 \end{aligned}$$

This represents an error of approximately 1 percent, which is acceptable given the performance margins designed into these oscilloscopes.

The pulse generator is driven by the CPU system clock. This allows it to be used for testing the oscilloscope's time-base system, which is driven by a different oscillator. The CPU system clock drives an array of counters with three different countdown values: 1 clock cycle, 16 clock cycles, and 256 clock cycles. This yields output frequencies of 10 MHz, 625 kHz, and 39 kHz. The output can also be put into a nontoggling or dc state. The outputs of the counters are decoded and sent through a series of Motorola ECLIPS gates to decrease the rise time from approximately 30 ns to 500 ps. This method of developing a fast rising edge results in a large amount of period jitter in the signal, but the jitter is sufficiently low not to affect the testing of the product. The pulse tops of the resulting signal are flat within less than one percent.

Since the production calibration module is based on the HP-IB interface module, it contains a TI 9914A GPIB controller IC. This chip is normally used to allow an external controller to program the oscilloscope, but in the case of the production calibration module it is used to allow the oscilloscope to function as a bus master during the test process. This capability provides some significant advantages which are described later.

Software. The main benefit of using the production calibration module for test control is that the DUT's microprocessor executes all of the control and measurement actions. Executing the tests from firmware resident in the production calibration module eliminates many slow HP-IB data transfers to a controller. The test system software resides in the 128K x 8 ROM on the production calibration module, and must have access to system code

on the main board. To allow compatibility with all revisions of system software, the interface between the module software and the system software consists of a series of look-up tables. Calls to system code from module code are accessed via ROM "hooks" located in system ROM. Calls to module code from system code are accessed via ROM hooks located in module ROM. These function pointer look-up tables do not change for minor revisions of module and system software. To accommodate custom module software like the production calibration module, many of the more important action routines within the system software are accessed through RAM hooks. This table of function pointers can be modified by the module software to implement many custom features (Fig. 3).

The test system software takes advantage of this architecture and virtually takes over the operation of the DUT. When the instrument is powered up with a production calibration module attached to the back panel, a function called `initmodule` is called. This function resides in the module software and replaces the front-panel key-press parsing function `parsekey` with a custom key-press parsing function `parseProdKeyHook`. This function is a state machine that displays the test system user interface, from which all the tests can be accessed.

Built-in Self-Test

Normal functionality of an instrument can be verified in large part by built-in self tests requiring no external stimulus. This test step executes nine built-in self-tests designed to verify the workings of the instrument comprehensively. The objective of these tests is to identify defective units as quickly and as early in the test process as possible. The tests are designed to operate properly regardless of whether the software calibration has been performed. The self-tests are:

- Functional test
- Analog-to-digital converter (ADC) test
- Video IC test
- Acquisition IC test
- Digital-to-analog converter (DAC) test
- Waveform RAM test

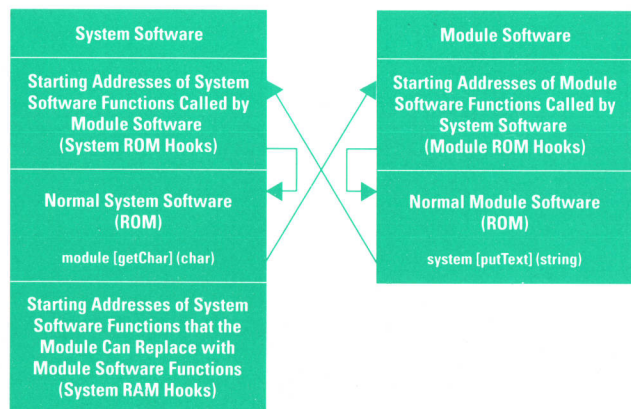


Fig. 3. RAM and ROM look-up tables contain "hooks" that are used for communication between the module software and the software of the system under test.

- System RAM test
- ROM checksum test
- Nonvolatile RAM test.

These tests execute in less than 30 seconds.

Software Calibration

The HP 54600 family of oscilloscopes is equipped with a multichannel 16-bit DAC, one channel of which is output to a BNC connector on the rear panel. This source is used for gain, offset, and trigger calibration and for low-frequency parametric testing of the DUT. To verify that this dc output is accurate enough for calibration and test, an HP 3478A DMM is used to measure two points on the source's output range. The digital multimeter connects to the DUT via the HP-IB, and the DUT controls the entire operation using the TI 9914A chip located in the production calibration module. An external controller is not used. To meet the test time goals for this product family, it was necessary to improve the software calibration time over previous products. The high-throughput acquisition system of the product makes it possible to measure and compensate for errors very quickly. The extensive characterization performed in the prototype phase also contributed to our understanding of how rigorous the calibrations had to be. The techniques used for calibration in the HP 54600 family are discussed in "Using the High-Resolution, Multichannel DAC" on page 54. The algorithms achieve an average calibration time of approximately 3.0 minutes.

Hardware Adjustments

Because of the aggressive throughput goals for this system, traditional means of executing hardware adjustments were not sufficient. To minimize the variability of an adjustment, an automated measurement is necessary. However, feeding this measurement back to the calibrator instantaneously is imperative if the adjustment is to be made quickly. Since the test system is built-in and operates using the CPU of the DUT, adjusted parameters can be measured and displayed in real time. To eliminate subjectivity, the test system uses a graphical feedback adjustment scheme (Fig. 4). A grid and a parameter indicator are displayed on the screen and updated using the HP 54600 family's high-performance waveform viewing architecture. The system CPU is involved in automated measurements performed during the adjustments, but the update rate is still superior to an external controller-based system. It is important to note the distance between the two adjustment limit indicators. Adjustment operators have demonstrated a tendency to overadjust when given the opportunity. In other words, if a procedure is worded "Adjust until indicator is between the markers," then the operator will tend to try to adjust the indicator to the median point between the markers. If the adjustment limit markers are correctly set this is unnecessary. The finite resolution of a raster display can be advantageous in dealing with this phenomenon. In this case, there are only two pixels between the markers, so there are only two locations to adjust to, neither appearing superior to the other.

The software calibrations described so far adjust all of the dc parameters of the oscilloscope. Two adjustments

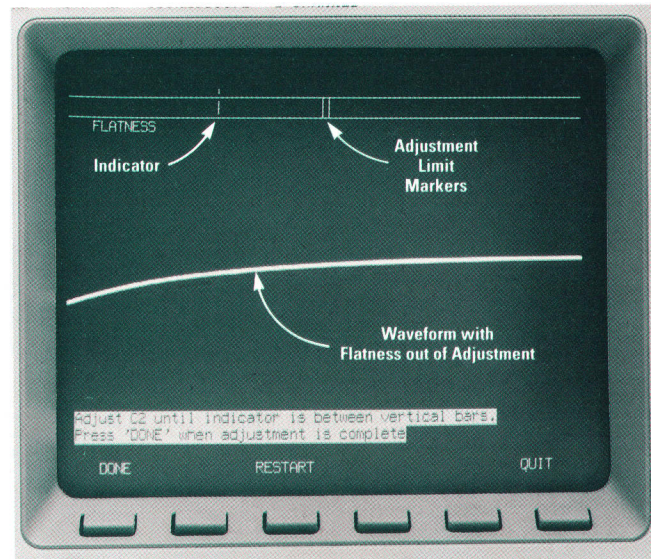


Fig. 4. Flatness adjustment display.

per channel are necessary to optimize the higher-frequency parameters of each unit. The low-frequency gain of the front-end attenuator is tuned by adjusting the flatness of a square wave. This adjustment is complete when the measured nonflatness is between 0.5% and -0.5%. The high-frequency response of the DUT is optimized by adjusting the overshoot with a 500-ps rising edge input. The range for this adjustment is 4.85% to 5.65%. Both of these adjustments use the production calibration module pulse source for input stimulus. The graphical adjustment method described above results in a high degree of repeatability for both adjustments. Table I shows the statistical results from 193 HP 54600A 2-channel oscilloscope boards for these adjustments.

Table I
Adjustment Results

Adjustment	Desired Range	Mean	Sigma
Flatness Channel 1	-0.5% to +0.5%	-0.01%	0.16%
Flatness Channel 2	-0.5% to +0.5%	-0.03%	0.17%
Overshoot Channel 1	4.85% to 5.65%	5.39%	0.16%
Overshoot Channel 2	4.85% to 5.65%	5.36%	0.19%

Measurements taken on the production line indicate that the adjustments typically take 1.5 minutes on the HP 54600A and 3.0 minutes on the HP 54601A.

Parametric Final Test

As its name implies, final test is the most critical step since it checks final instrument performance. The set of tests performed on the HP 54600 family was driven largely by the two-signal-source design goal. Tests that traditionally were tested with a low-frequency sine wave were changed to use the programmable dc source described above. An example of this is trigger hysteresis, which is usually tested by using a 1-kHz sine wave and expanding the DUT time base so that the signal on the screen appears as a flat line. The "dc" value of this flat




Fig. 5. Frequency response curve showing the area where hysteresis can be used as a measure of trigger sensitivity.

line is then computed for both positive and negative trigger slopes to obtain the hysteresis. In our case, the dc source is slewed through a trigger level for both slopes, and the difference in trigger values equals the hysteresis. On the other end of the frequency spectrum, tests that normally use a high-frequency sine wave are replaced with tests using the fast rising edge of the production calibration module. Delta-T accuracy and time-base linearity are examples of tests that were easily converted to use the repetitive pulse train from the production calibration module. Trigger bandwidth is another example. Trigger bandwidth is typically computed by stepping a sine input to higher and higher frequencies until the DUT fails to trigger. For the HP 54600 products, this parameter is computed in the following manner. First, the top and base of the square wave are measured using automated measurements in the acquisition path of the DUT. Using these voltages, the 10% and 90% voltage levels are determined. The trigger level is then programmed to the 10% voltage and the zero-crossing point is measured in the acquisition path. This procedure is repeated for the 90% voltage. Changing the trigger level from the 10% point to the 90% point causes a time slew in the acquisition path that is equal to the 10% to 90% rise time of the trigger path. Assuming that the response of the trigger path is second-order or simpler, the bandwidth of the trigger path can be calculated using:

$$BW(\text{trigger path}) = 0.35/t_{\text{rise}}(\text{trigger path}).$$

Perhaps the best example of using new test algorithms for the stimulus sources available is found in the bandwidth test. Bandwidth has traditionally been the most critical specification for an oscilloscope. It is usually measured with a high-frequency sine wave and a power meter by measuring the frequency at which the oscilloscope's gain is 3 dB lower than the gain at dc. Alternatively, an oscilloscope can be deemed as passing the bandwidth test if its response is less than 3 dB down at the specified frequency. In the case of the HP 54600 products, bandwidth is computed by acquiring the step response to the production calibration module edge source, differentiating the step response to approximate the impulse response, and taking the FFT (fast Fourier transform) of the impulse response to arrive at the

bandwidth. (See the article on page 29 for details of this technique.) Recalling that all the test hardware and software is self-contained, this means that the oscilloscope can test its own bandwidth, an interesting result.

Statistical Correlation

In addition to using new methods to test standard parameters, the HP 54600 test scheme uses statistical correlation to replace the measurement of one parameter with another. An example of this technique can be found in the trigger tests. The trigger sensitivity is a measure of the DUT's ability to trigger on a small, high-frequency signal. This test is expensive in that it requires a leveled high-frequency signal generator and a time-consuming search for the signal while adjusting the trigger level. Fortunately, the value of the trigger sensitivity is proportional to the amount of hysteresis in the trigger comparator. Until the frequency response of the trigger path begins to affect the measurement, the trigger sensitivity is exactly equal to the trigger hysteresis (Fig. 5). As long as the trigger sensitivity is specified at a frequency below f_{tb} , then the hysteresis can be measured instead of the sensitivity. Data acquired with the characterization system illustrates the correlation between these two parameters (Fig. 6). For both 25-MHz and 100-MHz test frequencies, the two parameters are linearly related over the normal operating region.

Another general example of correlation can be found in the number of vertical ranges used for the parametric tests. The single biggest reason for the traditionally long test times of oscilloscopes is that vertical range combinations are exhaustively tested. To see how this works, refer to Fig. 7. Let the first block represent an oscilloscope attenuator, the second block a coarse gain control, and the third block another level of gain control. (This architecture is often used in oscilloscopes to achieve a wide range of vertical sensitivities.) If nothing is known or assumed about the interrelationships of the blocks, a

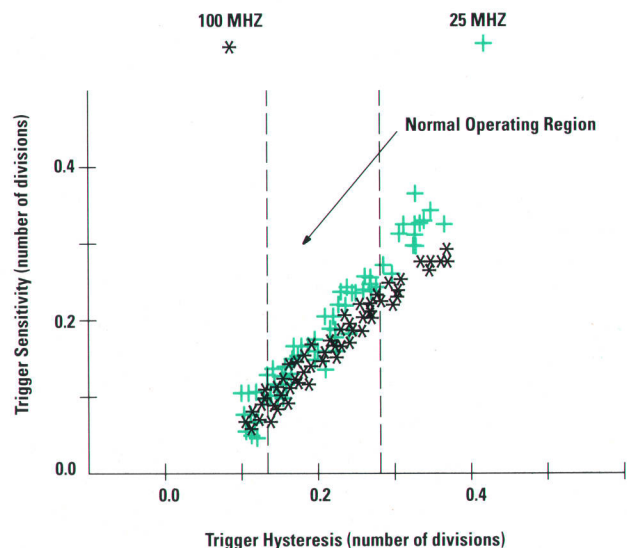


Fig. 6. Plot showing the linear relationship between hysteresis and trigger sensitivity.




Fig. 7. There are $3 \times 4 \times 2 = 24$ possible settings for this chain of three blocks.

given test will require $3 \times 4 \times 2$ or 24 combinations. In general the formula for N blocks is:

$$\text{Exhaustive combinations} = \prod_{i=1}^N M_i$$

where M_i is the number of settings of the i th block. On the other hand, if the blocks can be modeled as linear systems, then many of the above combinations will provide redundant information, and the the number of unique combinations can be reduced to:

$$\text{Unique combinations} = \sum_{i=1}^N M_i - (N - 1),$$

while still varying only one block's setting per combination. For the example above, using this technique results in 7 unique combinations instead of 24. Sometimes the number of combinations can be reduced even further by making use of a priori knowledge regarding relative performances across combinations. For example, by analyzing the underlying physics of a device's behavior, one can often determine which setting will produce the worst relative results. For tests with one-sided limits, it is then possible to test only the known worst-case setting, and infer the performances of the other settings. For this method to work, the true mechanism behind the device's behavior must be known and repeatable. Another way of stating this is that correlations must be confirmed with causative knowledge. In the HP 54600 products, both of the above techniques were used to arrive at a very small set of vertical ranges for most tests, resulting in as little redundant testing as possible.

Using the test minimization strategy detailed above, the final design for the production test system results in 14 tests being performed. For the 4-channel HP 54601A this translates to 53 different parameters, compared to the over 200 parameters tested by the characterization system. Table II lists the parametric tests performed on each product and the signal source used.

Table II
Parametric Tests

Test	Source
TV trigger verify	Pulse
Bandwidth	Pulse
Rise time	Pulse
Delta-T accuracy	Pulse
Time-base linearity	Pulse
Trigger bandwidth	Pulse
Hysteresis	DC
Trigger range	DC
Offset	DC
Offset intercept	None
Peak-to-peak noise	None

Test times for this step average 1.5 and 3.0 minutes for the HP 54600A and 54601A respectively.

Data Collection

For parametric data tracking, an HP 9000 Series 300 computer is used to collect data from multiple test systems (Fig. 8). Since the production calibration module uses a TI 9914A GPIB controller, the production calibration module is the bus controller in this architecture. The data collection controller uses HP 98624A HP-IB cards. The test system transfers the data to an HP-IB card as if it were a printer, and the data collection system parses the ASCII data into floating-point numbers and stores them in a database. This implementation is particularly convenient in that the system can be duplicated in a low-volume version with only hard-copy data tracking. Also, troubleshooting technicians use this feature to print test results. The specification limits residing in the test system software can be transferred to the data collection system using a function that operates in a similar manner. This makes it unnecessary to recode the limits in the

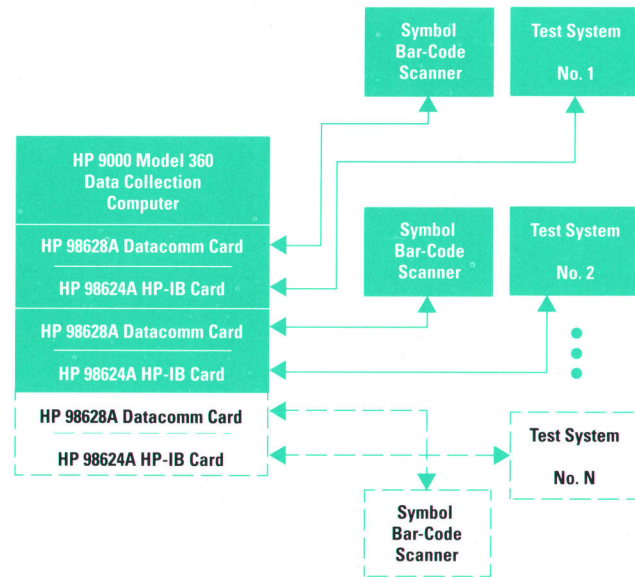


Fig. 8. Data collection scheme for the HP 54600 production test system.

data collection software. The serial number of the DUT is scanned from a bar-code label on the printed circuit board. The bar-code scanner interfaces with the data collection system via an HP 98628A Datacomm card. The serial number is transferred from the data collection system to the DUT, where it is stored in nonvolatile RAM. The data collection system is capable of generating test statistics in the form of histograms, \bar{x} -R control charts,

and c_{pk} analyses, as well as basic raw data dumps. Test cycle time data is also available through this system for capacity planning purposes. Since the data collection system is offline from the main test process, reports can be generated at any time while the process is running. This provision allows either production engineers or production operators to assess the state of any of the test parameters at will.

Stimulus/Response Defect Diagnosis in Production

Since production processes are less than perfect there will always be defective parts that must be repaired or discarded. Automated board testers help reduce the percentage of defects but technicians are needed to troubleshoot any that remain. Simple signal tracing will generally suffice for diagnosis of analog circuitry, but diagnosis of digital circuitry can be difficult.

For troubleshooting defective digital circuitry, there are two viable implementation extremes. Either the defective assembly diagnoses itself, or external means (equipment and labor) are applied to diagnose the defect. Self-diagnosis presupposes that the assembly is at least partially functional and that sufficient up-front engineering effort has been invested to produce useful diagnostics. The use of external stimulus/response test equipment also requires up-front investment and does not require functional circuitry, but does require labor for interconnections and equipment. Somewhere between the extremes is a viable, cost-effective solution to the problem of production line repair.

If the product is in the early stages of development there will be more flexibility in choosing which method to employ. A bottom-up procedure for verifying the functionality of each circuit block will aid in revealing troublesome portions of the design. Based on this procedure some changes can be incorporated to make troubleshooting easier. Even the most mundane issues such as part spacing for probing are often overlooked, to the detriment of the repair process. A review of the design by an experienced technician can reveal these kinds of oversights.

For high-volume production of complex products, a logic analyzer and a microprocessor emulator can provide the needed stimulus/response capability for fast fault isolation in a cost-effective manner (Fig. 1). The emulator executes code designed to exercise the address/data bus and control lines while the logic analyzer looks for the proper edges and patterns. Both timing and state analysis can be employed. A computer uses the bus programming features of the emulator and logic analyzer to automate much of the diagnosis.

For example, the emulator can execute code that writes to an address guaranteed to produce an output pulse from an address decoder. The logic analyzer waits for an edge in the proper direction and an error condition is signaled if the expected line does not assert. In like fashion, the neighboring pin of the decoder can be set to trigger before the write occurs. If a trigger is received then there is reason to believe that a short is preventing proper operation.

With this scheme even surface mount technology boards can be probed. High-quality probe clips are commercially available for surface mount devices including SOIC and PLCC packages. The emulator probe head can be clipped on top of a PLCC-packaged microprocessor on the target board by means of a suitable adapter, provided that the proper pull-up resistors have been designed into the product. The logic analyzer probe leads are configured such that the logic analyzer probe pods do not require different connections to the probes to perform the diagnostics. In other words, the logic analyzer lead connections to the probe clips for a 20-pin SOIC are made in such a way that the assembly can be clipped onto any 20-pin SOIC part on the product, with the computer programming the logic analyzer differently as necessary to reflect the different pinouts for the integrated circuits. In this

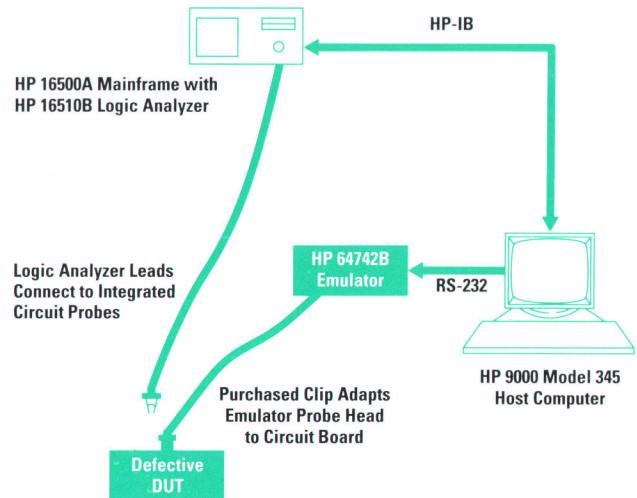


Fig. 1. System used in HP 54600 Series oscilloscope production to troubleshoot defective digital boards.

manner lengthy setup times for the tests can be eliminated since the probe configuration is static and all that is required is clipping the probes to the device under test (DUT).

The emulator can also be used independently of the logic analyzer. The attending technician can invoke a routine that writes a logic high onto one bit of the data bus while holding all other pins low. A "walking ones" stimulus can also be employed to find bus lines stuck low or high.

Digital-to-analog converters (DACs) can be programmed to ramp over their entire range. All of these effects can be observed by monitoring the proper lines with a logic probe or oscilloscope. Since the emulator has full control over the microprocessor, product-resident self-tests are not necessary. The product can even be booted using its own program code running on the emulator to determine if the microprocessor is at fault.

Depending on the level of preliminary design consideration given to troubleshooting and the amount of stimulus/response program code written, it is possible to give program outputs such as "No edge found U42 pin 5" or "Address bit 3 stuck." The value of these messages is obvious.

Chris J. Magnuson
Manufacturing Development Engineer
Colorado Springs Division





Fig. 9. HP 54600 production test system (top) compared with the test system for an earlier oscilloscope family (bottom).

Results

Results on this project have exceeded our goals in many ways. The average total test time for both HP 54600 products is less than 9 minutes, meeting the $10\times$ improvement goal set at the onset of test development. In terms of capital equipment, a single station for HP 54600 production test has a single-shift capacity of 35 units per day and costs one tenth as much as test stations for previous products. Fig. 9 illustrates how dramatic the equipment reduction is. The test strategy described here results in considerable savings in production cost, which translates to a difference of several hundred dollars in selling price, significantly improving the value delivered to customers. Several other potential benefits of the built-in test strategy have also become apparent. Because of the low cost of duplicating the test system, each production technician is supplied a test system for use in troubleshooting. The technicians can easily verify any corrections made to a DUT by executing the test that originally failed (see sidebar for digital troubleshooting techniques). Another benefit is the relative ease of transfer to other manufacturing sites because of the low equipment requirements. This latter benefit is especially important for high-volume, low-cost products such as the HP 54600 family. A final benefit is the experience gained regarding the ultimate limits of built-in self-test. For test and measurement instruments, this product serves as a good example of what can be done in the test arena when a coordinated strategy is employed. Although this strategy cannot be used for all products, we hope to leverage the ideas to new designs wherever possible.

Acknowledgments

The authors would like to thank Marty Grove of the Colorado Springs Division for his design of the production calibration module (especially the 500-ps edge source). This device enabled us to extend our built-in test concepts farther than many believed possible. We would also like to acknowledge Lee Edelmaier of the Colorado Springs Division for his work on statistical correlations and margin analysis tools. Also, thanks to Mark Schnaible and James Kakhoshka for their vital help in the development of the HP 54600 test system software.

Measuring Frequency Response and Effective Bits Using Digital Signal Processing Techniques

Frequency response and effective bits are informative measurements of digital oscilloscope performance, and can be calculated by efficient algorithms using the fast Fourier transform.

by **Martin B. Grove**

Two important measures of digital oscilloscope performance are frequency response and effective bits. The magnitude of the frequency response shows the voltage gain of the oscilloscope versus frequency. Effective bits is a measure of the ratio of signal power to noise power. Any source of noise or distortion in the oscilloscope will reduce the number of effective bits. Thus, effective bits is a wideband measurement that gives visibility to several performance limitations.

The digitizing architecture of the HP 54600A and 54601A oscilloscopes allows efficient implementations of these measurements using digital signal processing techniques. Both algorithms use the fast Fourier transform (FFT) to convert a digital time-domain waveform to a discrete frequency-domain representation.

Frequency Response Algorithm

A flow diagram of the algorithm for estimating frequency response is shown in Fig. 1a. A square wave generator is used to apply a fast step input to the oscilloscope. The resulting digital waveform represents the step response of the oscilloscope. The step response of the channel 1 input of an HP 54600A at 100 mV/div sensitivity is shown in Fig. 1b. The time derivative of the step response is the impulse response. A discrete version of the true derivative is performed on the digital waveform using the first backward difference operator. Fig. 1c shows the impulse response of the HP 54600A. The jagged lines result from noise in the step response amplified by the discrete derivative operation. This noise is attenuated by setting the HP 54600A in averaging mode.

The frequency response is the Fourier transform of the impulse response. A discrete version of the Fourier transform is implemented using a fast Fourier transform (FFT) algorithm.¹ The estimated frequency response magnitude of the HP 54600A is shown in Fig. 1d. The response is plotted out to 200 MHz. The bandwidth can be read from this plot and is equal to 140 MHz. Noise in the frequency response results from step noise in the impulse response.

Radix-2 FFT algorithms require the number of points to be an even power of two. To achieve this, the impulse

response is padded with zeros at the end of the waveform. This does not affect the accuracy of the frequency response estimate.

Sources of Error

Several factors, however, do limit the accuracy. An understanding of the sources of error can help minimize these limitations.

The most obvious source of error is the step input. The overall frequency response is a product of the contributions from the step input and from the oscilloscope. Thus, the contribution from the step input should be flat beyond the bandwidth of the oscilloscope. A first-order approximation for bandwidth is 0.35 divided by the 10%-to-90% rise time.² To make an accurate estimate of the frequency response out to the bandwidth of the oscilloscope (BW) with less than 5% error at the bandwidth frequency, the rise time t_r of the input should be:

$$t_r < 0.35/(3BW).$$

The step input used for the HP 54600A has a rise time of 500 ps, giving a reasonable frequency response estimate out to 233 MHz.

The frequency resolution of the FFT is also a limitation. The FFT algorithm produces samples of the continuous Fourier transform of the digital impulse response. The frequency resolution Δf of the FFT is a function of the effective sample rate f_s and the number of points N from the oscilloscope:

$$\Delta f = f_s/N.$$

In a random repetitive oscilloscope, the effective sample rate, not the sample rate of the analog-to-digital converter (ADC), determines the frequency limitations of the FFT. The HP 54600A can place points 100 ps apart and thus has a maximum effective sample rate of 10 GHz:

$$f_s = 1/\Delta t,$$

where Δt is the sample period of the oscilloscope and is equal to the time base range divided by the number of points. A 2000-point waveform from the HP 54600A acquired at the maximum sample rate gives a frequency

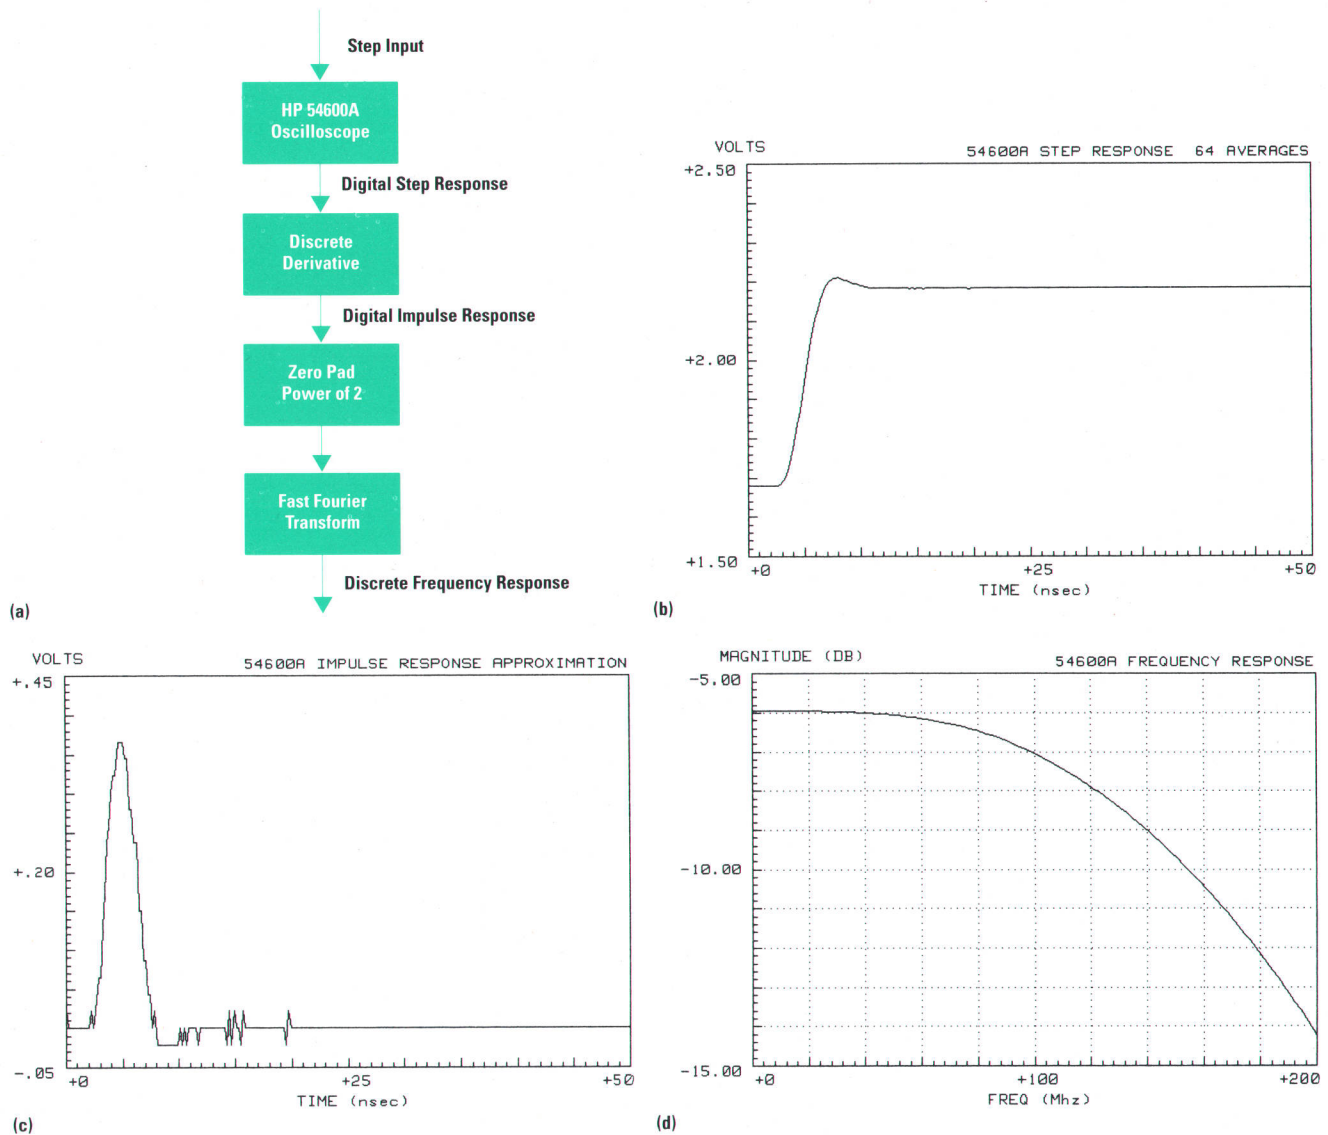


Fig. 1. Discrete frequency response estimation algorithm. (a) Flow diagram. (b) Digital step response of an HP 54600A oscilloscope, channel 1, 100 mV/div. (c) Approximate impulse response using the discrete derivative. (d) Estimated frequency response using a 1000-point step response and 2-GHz sample frequency.

resolution of 5 MHz. To obtain finer resolution requires either increasing the number of points or reducing the effective sample rate by slowing the sweep speed. In Fig. 1d, the effective sample frequency is 2 GHz and the number of points is 1000, giving a frequency resolution of 2 MHz.

Zero padding the impulse response can be used to increase the number of points. However, this does not add new information, and the resolving capability of the FFT is not improved. For example, if zero padding extends the sequence length by a factor of two, then every other point of the FFT of the zero padded sequence has the same value as the FFT of the unpadded sequence.¹ Thus zero padding has the same effect as interpolation: it fills in points between frequency samples giving a better visual image of the frequency response curve.

Aliasing can cause error in the frequency response measurement. Aliasing occurs for frequencies equal to or

greater than half the sample frequency (Nyquist frequency). The power in the signal at frequencies above the Nyquist frequency folds back and adds to the power at lower frequencies, distorting the estimated frequency response. Frequencies higher than the Nyquist frequency will fold back to a particular frequency f_1 as follows:¹

$$mf_s - f_1 \text{ and } mf_s + f_1 \text{ alias to } f_1 \quad m = 1, 2, \dots$$

For example, in Fig. 1d, the effective sample rate is 2 GHz, so the power at 200 MHz includes the power at 1.8 GHz and at 2.2 GHz. However, this power is small because of the two-pole gain roll-off of the HP 54600A at frequencies above the bandwidth.

Aliasing errors can be minimized by attenuating the power in the signal at frequencies greater than the Nyquist frequency or by increasing the sample rate. Single-shot digital oscilloscopes usually incorporate a low-pass filter to cut off power at frequencies greater than the Nyquist frequency. In the HP 54600A random

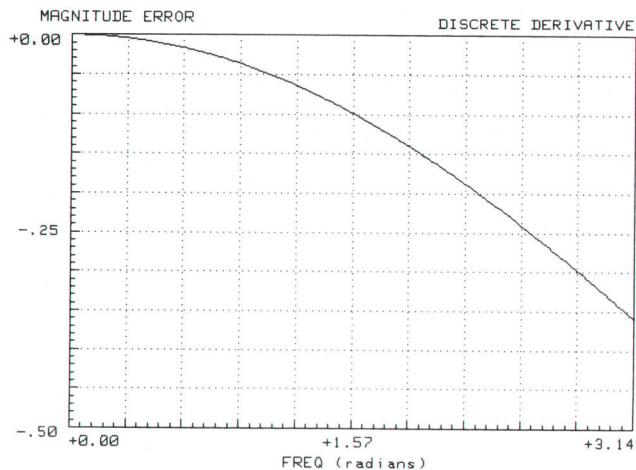


Fig. 2. Frequency response error resulting from the discrete derivative.

repetitive mode, the sweep speed is set fast enough to guarantee that the effective sample rate is much greater than the bandwidth of the oscilloscope.

The discrete derivative is an approximation of the true derivative and causes errors in the estimated frequency response. This error can be quantified. Let $x(n)$ represent a sample of the impulse response. The discrete derivative is simply:

$$[x(n+1) - x(n)]/\Delta t.$$

The error increases with larger sample interval Δt , or equivalently, with smaller sample frequency. The error in the frequency domain can be determined by letting $x(n)$ equal the complex exponential $\exp(j\omega n)$. Then the discrete derivative is:

$$[\exp(j\omega(n+1)) - \exp(j\omega n)]/\Delta t$$

where $\omega = 2\pi f\Delta t$ radians. The true derivative is:

$$\frac{d}{dn} \exp(j\omega n)/\Delta t = j\omega \exp(j\omega n)/\Delta t.$$

The magnitude of the ratio of the discrete derivative to the true derivative is:

$$\sin(\omega/2)/(\omega/2) = \text{sinc}(\omega/2) = \text{sinc}(\pi f/f_s).$$

The error in the discrete derivative versus ω is plotted in Fig. 2. The Nyquist frequency is equivalent to $\omega = \pi$ radians. The error is negative and its absolute value increases with increasing frequency. This error can be reduced by increasing the sample frequency f_s . However, there is a trade-off. Increasing the sample frequency degrades the frequency resolution of the FFT.

The step input to the ADC will have a noise component from the signal generator and from the signal conditioning circuitry of the oscilloscope. This noise is generally small enough that it causes only the least-significant bit (LSB) of the ADC to change randomly. The random fluctuation can be modeled as a white noise component added to the step response. This noise is amplified and high-pass filtered by the discrete derivative operation

before passing on to the FFT. An upper bound for the noise is:³

$$(2/M) \sqrt{(\pi N/2) \sin(\pi f/f_s)},$$

where M is the number of ADC codes. The noise increases as the number of points N increases. Reducing the number of points to reduce this noise must be traded off with more error in the discrete derivative and poorer frequency resolution from the FFT.

Fig. 3 shows the HP 54600A frequency response estimate for 200-point and 1000-point waveforms. The 1000-point waveform has 7 dB more noise at the Nyquist frequency. The upper noise bound is also plotted. The noise approaches the upper bound as predicted. The noise in the lower frequencies is caused by aliasing. Noise at frequencies greater than the Nyquist frequency fold back to the lower frequencies.

One method for reducing the noise without sacrificing the number of points is to average several time-domain waveforms. The HP 54600A has a very useful averaging mode that can average 8, 64, or 256 waveforms. The averaging occurs in real time and is updated after each


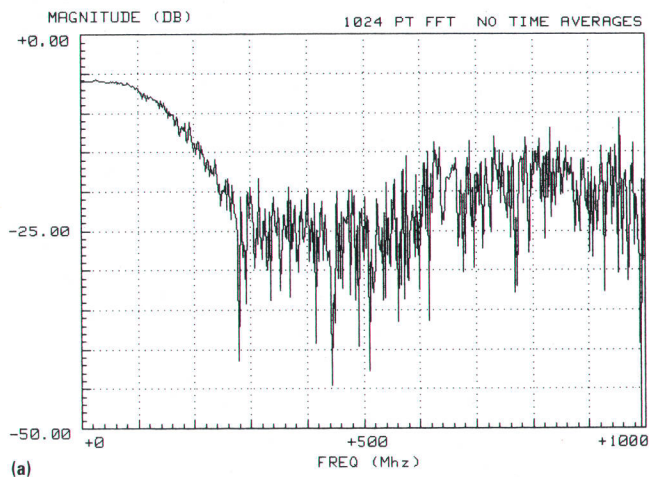



Fig. 3. Noise in the frequency response of an HP 54600A. (a) 200 points. (b) 1000 points.



(a)



(b)

Fig. 4. Noise in the frequency response of an HP 54600A. (a) No averages. (b) 64 averages.

acquisition. Fig. 4 shows the frequency response for no averages and for 64 averages. Averaging in the time domain reduces the noise in the frequency domain by 20 dB.

Overall Performance

Fig. 5 shows the estimated frequency response for a second-order system. The second-order system was chosen to represent the best fit for channel 1 of the HP 54600A at 100-mV/div sensitivity. This best-fit model was derived from the rise time and overshoot of the step response.² The model was then used to create a 1000-point second-order step response waveform. White noise with variance equal to the noise variance in the HP 54600A was added to the step response. This step was then passed through the frequency response estimation algorithm and plotted in Fig. 5. The true second-order frequency response was also plotted. Noise dominates the response at 500 MHz or one-fourth of the sample frequency. At the 140-MHz room-temperature bandwidth of the HP 54600A, the estimation is repeatable to within ± 0.1 dB of the true response.

Effective Bits Algorithm

Effective bits in a digital oscilloscope is a measure of the signal-to-noise ratio (SNR). The power spectrum can be used to calculate the SNR of a signal. A plot of the power spectrum shows the power in a signal as a function of frequency. The block diagram for a power spectrum estimation algorithm is shown in Fig. 6a. A spectrally pure sine wave is input to the oscilloscope. The resulting digitized sine wave from the oscilloscope is windowed with a Hanning window. The end of the waveform is then padded with zeros. The zero pad points should extend the waveform to at least four times its original length and the number of points to an even power of two. The extra zero padding improves power estimation accuracy as will be explained later. The FFT is then performed to convert from the time domain to the frequency domain. The magnitude of the FFT output is squared to give a scaled estimate of the power spectral density.

The area under the power spectral density is proportional to the total power in the waveform. Fig. 6c shows the power spectrum for a 20-MHz sine wave digitized by the HP 54600A. The digital sine wave is shown in Fig. 6b. The area under the lobe at 20 MHz is proportional to the signal power. The area under the noise floor is proportional to the noise power. The ratio of these two areas is an estimate of the SNR.

Effective bits can be calculated from the SNR (see page 34):

$$\text{Eff bits} = (\text{SNR in dB} - 1.8)/6.02$$

where SNR in dB = $10\log(\text{SNR})$.

The 6.02 factor is equal to $20\log(2)$, which represents one bit of an ADC. The effective bits calculated for the HP 54600A using 8 averages and a 20-MHz input is 7.5 bits.

Sources of Error

Factors that limit the accuracy of the effective bits estimate include the spectral purity of the sine wave input, the resolution of the FFT, aliasing, and spectral leakage. The input should have an SNR at least 10 dB




Fig. 5. True frequency response of a second-order system compared with the estimated response.


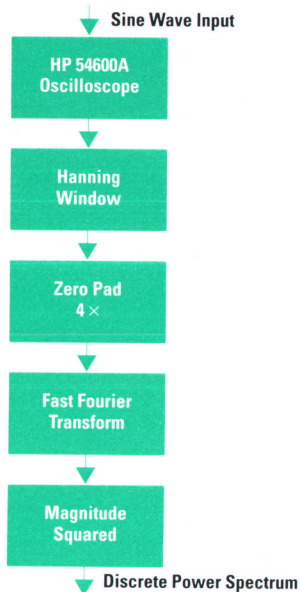


Fig. 6. Discrete power spectrum estimation algorithm. (a) Flow diagram. (b) 20-MHz, 1000-point sine wave from an HP 54600A. (c) Estimated power spectrum using a Hanning window.

greater than the SNR of the oscilloscope. Using the effective bits equation above for the HP 54600A with 8 bits, this would require:

$$\text{input SNR in dB} > 10 + 8(6.02) + 1.8 = 60 \text{ dB.}$$

The ability to estimate signal and noise power from the discrete power spectrum depends directly on the resolution of the FFT. The frequency resolution of the FFT was given earlier as:

$$\Delta f = f_s/N.$$

Increasing the number of waveform points gives finer frequency resolution, but also increases the execution time. For a radix-2 FFT, the number of math operations is proportional to:¹

$$N \log_2(N).$$

Reducing the sample frequency also provides greater resolution at the expense of increased aliasing error.

Spectral leakage refers to the spectral line lobes in the frequency spectrum. The frequency width of the lobes is inversely proportional to the length of the time-domain waveform. If the Fourier transform could be applied to an infinite-duration waveform, the width of the lobes would be reduced to single spectral lines. The shape of the lobes depends on the window applied to the time-domain waveform. A window is applied by multiplying each point in the window by each point in the time-domain waveform. Multiplication in the time domain becomes convolution in the frequency domain.

Fig. 7 shows the frequency spectrum for two types of windows, each with 100 points. The windows have been zero padded to give better visual displays. The rectangular window has a magnitude of 1 in the 100-point interval and a magnitude of zero outside this interval. This corresponds to truncating the time-domain waveform. The main lobe width in the frequency spectrum is:¹

$$\text{main lobe width} = 2f_s/N$$

The first sidelobe is 13 dB below the main lobe. Each successive sidelobe falls off very gradually. Since the power of individual noise components from the HP 54600A is 60 dB below the signal component (see Fig. 6c), the rectangular window cannot be used for power spectral estimation. The noise components cannot be resolved from the signal component.

The other window shown in Fig. 7 is the Hanning window. The discrete time-domain Hanning window $w(n)$ is given by:¹

$$w(n) = 0.5[1 - \cos(2\pi n/(N-1))].$$

The width of the main lobe in the frequency spectrum is:

$$\text{main lobe width} = 4f_s/N.$$

The first sidelobe is 31 dB below the main lobe. Each successive sidelobe falls off very sharply. The sixth sidelobe from the main lobe is attenuated more than 60 dB. The area or power under the main lobe represents 99.9998% of the total spectral power at a particular frequency. Thus the area of the main lobe is proportional

Calculating Effective Bits from Signal-To-Noise Ratio

Analog-to-digital converters (ADCs) have limited voltage resolution. Sampling the analog waveform produces quantization error noise. The signal-to-noise ratio (SNR) from the ADC is proportional to the number of quantization levels. For an ADC with b bits, there are 2^b quantization levels. Effective bits is a measure of the SNR. This measure can be applied to any signal. The effective bits of the signal refers to the number of ADC bits required to produce an equivalent SNR.

Random signal theory can be used to derive a formula for effective bits as a function of SNR. An ideal ADC has a maximum roundoff error of $\pm \frac{1}{2}$ quantization level. Each quantization level has a magnitude of:

$$Q = A2^{-b},$$

where A represents the range of the ADC. Assuming that any roundoff error from the ADC between $\pm Q/2$ is equally likely, the probability density function for the error is uniform, as shown in Fig. 1. The magnitude of the probability density function is such that the area under the probability density function is equal to one. This ensures that there is a 100% probability of getting a roundoff error between $+Q/2$ and $-Q/2$.

The variance of a signal is proportional to its power. The SNR can be written:

$$\text{SNR in dB} = 10\log[\text{var}(\text{signal})/\text{var}(\text{error})].$$

Let the input $x(t)$ to the ADC be a sine wave that covers the full ADC range:

$$x(t) = (A/2)\sin(\omega t)$$

where ω is frequency in radians per second. The variance of $x(t)$ is:

$$\begin{aligned} \text{var}(x(t)) &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \left[\frac{A}{2}\sin(\omega t)\right]^2 d(\omega t) \\ &= \frac{A^2}{8}. \end{aligned}$$

to the power. Overlap of the main lobes in the noise spectrum causes error in the estimated power. This error is reduced by increasing the frequency resolution.

In Fig. 6c, the sample frequency, which is adjusted by the sweep speed of the oscilloscope, is set so the 20-MHz signal lobe is at the very left of the spectrum. This cuts off the left half of the spectral leakage centered at 20 MHz, minimizing the noise power covered up by this leakage. To estimate the area under the signal lobes accurately, zero padding is used. By experimenting, the number of zero pad points required was found to be at least three times the number of points in the original waveform. These zero pad points fill in the shape of the lobe but do not alter it. The SNR is computed by summing the values under the main lobe of the power spectrum and dividing by the sum of the values under the rest of the spectrum. The noise summation starts at the point where the sidelobes of the main signal are below the noise floor.

For applications that require estimating spectral power rather than a ratio of powers, it is necessary to solve for the proportionality constant that relates the mean-square value of a signal to the power in the discrete power spectrum. If a Hanning window is applied to the digital

The variance of the quantization error $e(t)$ is:

$$\begin{aligned} \text{var}(e(t)) &= \frac{1}{A2^{-b}} \int_{-\frac{A2^{-b}}{2}}^{\frac{A2^{-b}}{2}} e^2 de \\ &= \frac{A^2 2^{-2b}}{12}. \end{aligned}$$

Substituting these variances into the SNR equation gives:

$$\begin{aligned} \text{SNR in dB} &= 10\log[(12/8)2^{2b}] \\ &= 1.761 + 6.021b \text{ dB}. \end{aligned}$$

Solving for effective bits b gives:

$$b = (\text{SNR in dB} - 1.761)/6.021.$$

Thus the effective bits is a scaled version of the SNR as expected.




Fig. 1. Probability density function assumed for quantization noise of an ADC with b bits and a range of A .

signal, then the proportionality constant is the inverse of the Hanning window power. The Hanning window power can be found using Parseval's theorem:¹

$$\begin{aligned} \sum_{k=0}^{N+N_z-1} |W(k)|^2 &= (N + N_z) \sum_{n=0}^{N-1} |w(n)|^2 \\ &= (3/8) (N - 1) (N + N_z) \end{aligned}$$

where N is the number of window points, N_z is the number of zero-pad points and $W(k)$ is the discrete Fourier transform of $w(n)$.

Overall Accuracy

Fig. 8 is a table of actual bits versus calculated effective bits using the power spectrum estimation algorithm. A sine wave was created mathematically and then quantized. The resulting waveform was passed to the algorithm. The SNR and effective bits were calculated from the estimated power spectrum. The accuracy increases for an increasing number of points since the frequency resolution of the spectrum estimate improves. Accuracy falls off for ADC resolution greater than 10 bits because spectral leakage from the main signal interferes with the noise spectrum.





Fig. 7. Frequency spectrum for (a) a 100-point rectangular window and (b) a 100-point Hanning window.

Conclusion

Both the frequency response and effective bits algorithms are efficient in terms of lines of code and execution time. Both measurements provide a great deal of information about the performance characteristics of a digital

Actual Bits	Estimated Effective Bits		
	200 Points	1000 points	2000 Points
4	4.15	4.08	4.07
6	6.31	6.05	6.03
8	8.19	7.96	8.01
10	10.11	9.97	9.98
12	11.14	11.62	11.62

Fig. 8. Actual number of bits versus estimated number using power spectrum estimation method.

oscilloscope. With an understanding of the limitations and sources of error, repeatable, accurate results can be obtained. The algorithms rely on the FFT to convert digital time data to a discrete frequency-domain representation. Care should be taken in applying the FFT because it is subject to finite-precision arithmetic errors.¹ The ability of the FFT to represent the continuous spectrum improves with increasing effective sample rate, ADC resolution, and number of points. The HP 54600A and 54601A oscilloscopes are very well-suited for these kinds of measurement techniques. The fast display update rate, 8-bit ADC resolution, 4000-point waveforms, average mode, and 10-GHz maximum effective sample rate allow fast, accurate results. The frequency response magnitude estimates have been shown to be accurate within ± 0.1 dB for frequencies less than or equal to the bandwidth of the oscilloscope. The effective bits estimates have been shown to be accurate within ± 0.1 effective bit for signal-to-noise ratios less than or equal to 62 dB (10 effective bits).

References

1. A.V. Oppenheim and R.W. Schaffer, *Digital Signal Processing*, Prentice-Hall, Inc., 1975.
2. B.C. Kuo, *Automatic Control Systems*, Prentice-Hall, Inc., 1982.
3. M. Souders and D.R. Flatch, "Accurate Frequency Response Determinations from Discrete Step Response Data," *IEEE Transactions on Instrumentation and Measurement*, Vol. IM-36, no. 2, 1987.

Mechanical Design of the HP 54600 Series Oscilloscopes

Simplicity of manufacture and a minimum of parts were the approaches taken to achieve high quality and reliability. Robotic assembly wasn't a consideration, so rotating motions were often chosen to mate components in final assembly.

by Robin P. Yergenson and Timothy A. Figge

The mechanical design team for the HP 54600 Series oscilloscopes shared the overall design objective of achieving the highest quality and reliability of any oscilloscope available. Quality and reliability must be designed in from the start. The mechanical design team felt that the most effective way to ensure that the instrument attained these goals was to make it simple to manufacture with as few parts as possible to go wrong in the field.

With this in mind, fasteners have been eliminated where possible. Snap fits and sliding and rocking component lock mechanisms are employed. As a result, the HP 54600A and 54601A require a total of only 22 threaded

fasteners for final assembly. Fig. 1 shows the components that go into the final assembly.

It was decided early on that robotic assembly would not be cost-effective in any portion of the assembly operation. This allowed us to free our thinking a bit and forget the straight-line-motion assembly techniques often used to evaluate manufacturability. As a result, over half of the assembly operations consist of a simple rotational attachment technique. Typically, one mating end has fingers and slots that interlock so that when one component is rotated into place, a slight interference at the joint results in a retention load on the two components. The other end is aligned by component details and held together by

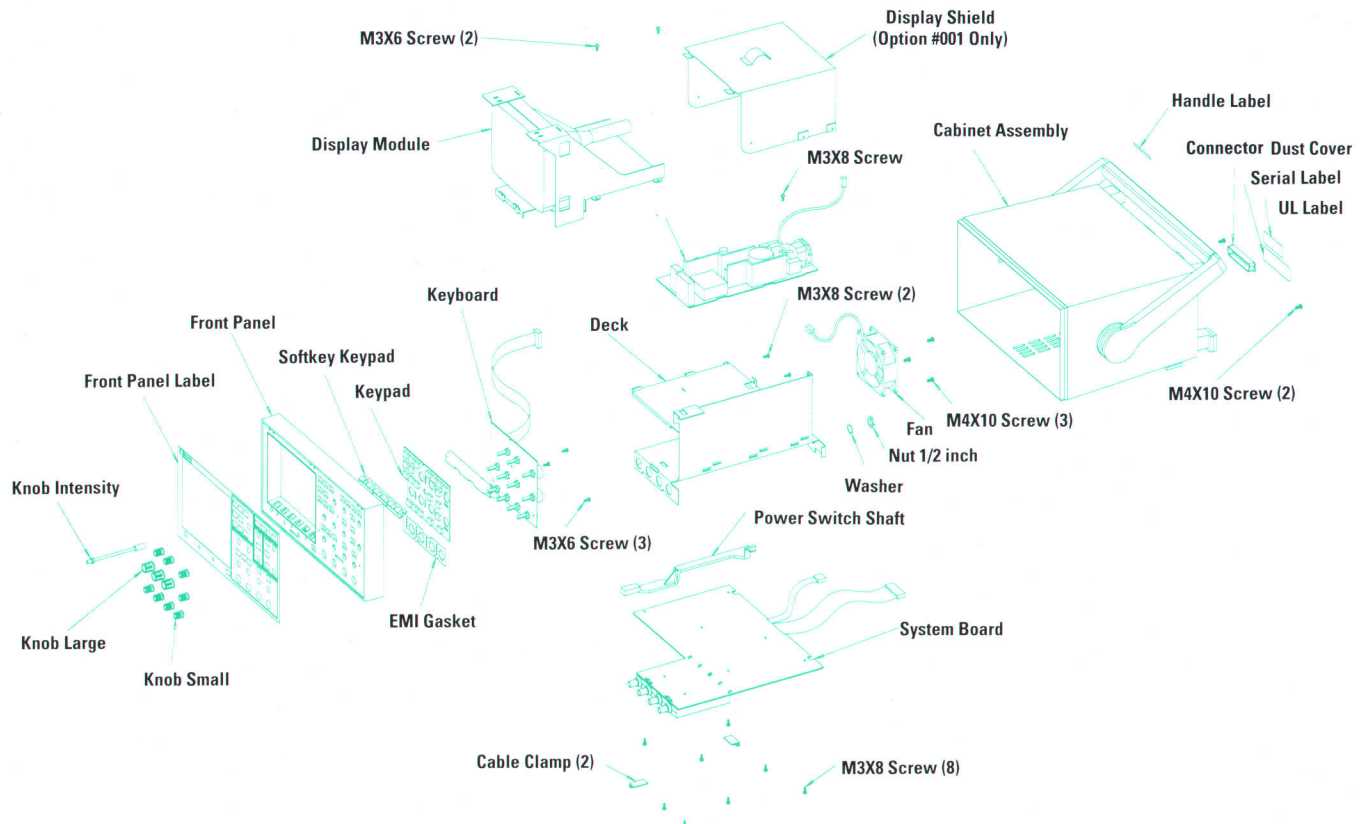


Fig. 1. The components of the final assembly process of the HP 54600 Series oscilloscopes. Only 22 threaded fasteners are used.




Fig. 2. Examples of HP 54600 assembly techniques.

a screw or snap feature. Fig. 2 shows examples of this concept. This method of assembly can reduce threaded fasteners by a factor of two. The rotary motions of the parts, which would be difficult and time-consuming for a robot, are simple operations for manual assemblers.

The use of keyhole standoffs also helps reduce fasteners. For example, the power supply mounts with a 9-mm sliding motion to engage the keyhole standoffs and one screw to eliminate sliding. Also, one side of the power supply printed circuit board mounts to a card guide formed in the side of the sheet-metal deck with the same motion required by the keyhole standoffs.

The knobs on the front panel of the HP 54600 use an internal interference protrusion with a reverse draft angle to ensure that they stay fixed on the encoder shafts. This eliminates the installation time and complexity of set screws. The feature works amazingly well, resulting in knobs that push on quite easily, but require a substantial increase in force to remove.

The electronics of the HP 54600 family consists of four assemblies: power supply, display, keyboard, and system board. Each of these components arrives at the assembly line fully tested and ready for installation. To keep

interconnection problems to a minimum, only two ribbon cables are used inside the HP 54600. These features result in a total assembly time of less than 12 minutes, one-third that of previous designs.

Even with an estimated mean time between failures of 50,000 hours it still seemed appropriate to include serviceability as an objective. Removing two screws (see Fig. 3) from the cabinet allows the remaining chassis assembly to be slid out of the cabinet to permit removal of the front panel, system board, display module, power supply, or fan. Although caution must be used when the cabinet is removed, the instrument remains fully operational to allow easy access for troubleshooting. With the cabinet removed, access holes in the front panel allow calibration for flatness and overshoot (see Fig. 4).

To satisfy the instrument's cooling requirements the enclosure is pressurized with an 80-mm dc brushless fan. The air flows across the power supply and display modules and then wraps around the sides to flow over the system board and out the ventilation holes, which are along the bottom of the cabinet. The narrow ducting area along the bottom of the cabinet results in a venturi effect which provides a relatively high airflow velocity across




Fig. 3. Two screws allow removal of the chassis assembly from the cabinet. For ruggedness, the rear feet are integrated into the cabinet and have four load-bearing walls.




Fig. 4. With the cabinet removed, four holes in the front panel allow calibration for flatness and overshoot.




Fig. 5. Airflow in the HP 54600 Series.

the system board (see Fig. 5). The 20 cubic feet per minute of air flowing through the box results in an average air temperature rise of only about 8°C above ambient. To satisfy environmental conditions requiring additional cooling, a thermal sensor inside the instrument ramps the fan voltage from 8.5V to 14V depending on incoming air temperature. This way the instrument is extremely quiet at standard room temperatures without compromising component life when used in harsh conditions.

Portability

General-purpose oscilloscope users want to be able to move an oscilloscope easily, so portability ranked high on the list of mechanical design objectives. The size evolved out of an earlier decision to put the whole oscilloscope on one board and to minimize raw board costs by having four boards per panel. This gave us a tentative board size of about 8.5 by 11 inches. The project's first manager suggested putting the oscilloscope board opposite the power supply and display modules with the sheet metal deck separating them. This configuration helps isolate sensitive electronics from the magnetic fields associated with the power supply and display modules and to a large degree defines the footprint of the instrument. A 3/4-width cabinet (about 12.5 inches wide) meets HP industrial design standards and works well with the 11-inch board dimension. The height is based on clearances required for the display module. Although the final configuration diverges from the deep aspect ratio of some other oscilloscopes, the resulting display area is nearly twice that of these other products.

The relatively lightweight plastic resins that make up the cabinet along with the single-board configuration result in a weight of about 14 lb. Light weight, a generously radiused bail handle for maximum carrying comfort, and cord wrap feet for the power cord amply satisfy the design objective for portability.

Ruggedness

A further design objective for a portable product is that it be rugged. By making the instrument light so that it satisfies the design objectives for portability we created a

product that could be thrown higher too, and that increases the need for ruggedness. HP's standard Class B-1 environmental tests require that lighter instruments survive higher levels of shock. The HP 54600 family is subjected to a 50% energy increase in shock tests compared to a 20-pound instrument of the same class.

On the other hand, there are some obvious advantages to having a small package size. Compare a short instrument and a tall instrument, both positioned on their rear feet and both with a similar base. The short instrument offers greater stability against tipping because of its lower center of mass, and when tipped over experiences lower impact forces than the taller instrument. The U.S. MIL-T-28800D standard requires just such impact tests on this type of instrument.

To enhance the product's resistance to impact, a highly impact-resistant ABS/polycarbonate resin blend was chosen for the enclosure. While this material exhibits only moderate stiffness, its resistance to impact is two to five times that of most of the stiffer modified resins. Ribs were added where a higher level of stiffness is required.

The rear feet were quite vulnerable to impact. By integrating them into the cabinet (see Fig. 3), providing four load-bearing walls on each foot, and adding generous fillets to reduce stress concentrations, the vulnerable foot problem was eliminated. The handle is molded of an even more impact-resistant polycarbonate resin and was tested to support the cabinet with a static load of over 400 pounds.

Although cosmetic blemishes resulting from rough field handling do not affect the performance of the instrument, we wanted to minimize them. By ruling out any external paint operations and by molding in custom colors, the final textured surface was made resistant to scratches.

To meet our objectives for ruggedness on the inside of the cabinet, we chose a tempered steel alloy for the deck which allows it to withstand higher levels of stress while preventing permanent deformation. Rubber bumpers serving as shock isolators are strategically located to help protect some of the more sensitive internal components from impacts. The resulting product not only passes HP's own internal ruggedness tests but also meets U.S. MIL-T-28800D shock and vibration requirements for Type III, Class 3, Style D equipment.

Low Cost

It was apparent that the package for the HP 54600 Series would have to bear its share of the cost reductions. The strategy to achieve this goal was to integrate as many parts as possible, and to hard-tool* all mechanical parts.

One example is the handle design. Most bail-type handles consist of a number of separate parts requiring several steps to assemble. The HP 54600 handle is one molded piece. Its size and design make it comfortable to use, and it provides six commonly used detent positions. The

*Hard-tooled parts are those for which tools of high-quality steel have been made in anticipation of long production runs.




Fig. 6. The HP 54600 front panel layout resembles that of a typical 100-MHz analog general-purpose oscilloscope.

handle mates to the side wall of the cabinet, which has handle removal and position defining features designed in (again only one part). The cabinet also has an integral fan grill and has most of the rear-panel text molded in to avoid as many add-on labels as possible.

Much of the internal structure has also been integrated to a high degree. All main subassemblies in the HP 54600 family, excluding the handle, are held in place by a one-piece deck assembly (see Fig. 1). The deck provides a rigid platform for the instrument and all the mounting structure and EMI isolation for the system board. Every feature except one fold in the deck is done on hard-tool sheet-metal dies.

To provide shielding between input channels, die-cast attenuator covers are used. These covers combine the required shielding with BNC connectors cast in place, all mounted to the deck with the previously mentioned hook and rotate motion and one screw.

No machined parts were designed for the HP 54600. The hard-tooled parts require no secondary machining other than sprue removal on the large plastic parts. Castings are designed to be flash-free, and sheet-metal dies with burr-free tolerancing are used. As a result, the cost goals set early in the project for the package were achieved.

User Interface

Much time and consideration were put into the user interface for the HP 54600. The layout of the controls and nomenclature (see Fig. 6) was painstakingly scrutinized by the entire project team. Focus group studies were done to discover the types of knobs and buttons customers prefer most on a 100-MHz oscilloscope. Of course, the final choices had to be tempered by our aggressive cost goals. The result is a keyboard/keypad combination that provides very intuitive control at an extremely low cost. The keys hinge about their upper edge to eliminate the annoying wobble of previous elastomeric keys. These keys also have a tactile "pop" that gives the user a clear signal that the key has been depressed, without the use or expense of an audio beeper.

Two types of mechanical rotary encoders are used behind the front-panel knobs: detented and smooth. The detented version has a custom spring and lubricant to provide the

user with good positive feedback without requiring too much torque. The smooth feel of the other encoder type is provided by heavy lubricants inside the encoder. This gives the user the feel of a quality damped control.

The front-panel knobs were designed with the customer's mode of operation in mind. The large knobs on the detented encoders are intended to be grasped by the user and turned slowly in definite steps, as a user turning vertical or time base settings on an analog oscilloscope would do. Raised flutes on the knob facilitate this mode of operation. The small knob used on the smooth encoder has no flutes so the knob and encoder can be spun in the fingers of the user.

Familiarity was another goal for the front panel. Customers are comfortable with analog or hybrid oscilloscope interfaces. To many customers, the HP 54600 looks like an oscilloscope they learned to use in school or on their first job. Customers expect to find the intensity control, ground reference, probe compensation, and power switch on the front of a 100-MHz oscilloscope. The front panel of the HP 54600 attempts to meet these expectations instead of trying to change them.

EMC Design

EMC designs of the past have depended upon gasketing techniques at the seams that are typically expensive and labor intensive. The injection-molded enclosure used in our design allowed us to integrate the features required to provide reliable mating contact into the parts themselves and helped reduce part count, labor, and cost.

The HP 54600 cabinet design is a "deep can" configuration that minimizes the number of enclosure seams and at the same time simplifies the type of relative motion required for proper engagement of the cabinet and the front panel. To provide reliable contact force every inch or so along the perimeter of the seam, the knife-edge of the cabinet slides in between the stiffening ribs and the outer wall of the front panel (see Fig. 7). Interfering bumps along the wall of the front panel force the knife-edge of the cabinet to "snake" back and forth between the interfering bumps and the stiffening ribs, resulting in reliable contact force between the two mating parts. This design also minimizes the insertion force required to mate




Fig. 7. Details of the contact method used to mate the front panel to the cabinet. The amount of deformation in the cabinet knife-edge is exaggerated for clarity. Some deformation also occurs in the wall of the front panel. Also see Fig. 1 on page 41.

the two parts. When the parts are fully engaged, an 8-mm overlap between them offers additional capacitive coupling to help minimize the potential difference between the two parts at higher frequencies.

Once adequate mating contact was achieved between the two nonconducting materials, a method was needed for making the enclosure electrically conductive. The design required good surface conductivity, which ruled out conductive-filled plastic. We were trying to avoid external painting operations so that process rejects, high cost, and ugly scratches resulting from customer use could be avoided. This ruled out double-sided plating (although we knew that we could fall back on it later if our shielding requirements became too high). Conductive paints didn't offer the kind of wear characteristics or the consistent coating thickness that we desired. The two remaining options that were still under consideration were selective plating and vacuum metallization. Selective plating was considered to be higher in cost and was believed to have certain process problems that could result in delivery problems. Our early data suggested that the vacuum metallization process under consideration would provide a low-cost alternative while providing the good surface conductivity, adequate wear characteristics, and consistent coating thickness that were desired. The supplier who would be molding our enclosure parts had an aluminum vacuum metallization process in-house, so we chose the vacuum metallization process.

This particular coating process results in aluminum thickness measurements of 0.5 to 5 micrometers (20 to

200 microinches). For the frequencies that we are trying to attenuate, the enclosure is in the near field of the source and the field is primarily magnetic. Normally, the primary means of attenuating these high-frequency magnetic fields is through absorption loss, but that requires a much greater thickness than the vacuum metallization process provides. Reflection loss then becomes the remaining means of attenuation. With the majority of the oscilloscope electronics on one board and some careful EMC design on the part of the electrical designers, the shielding thickness proved to be sufficient.

Other aspects of the HP 54600 EMC design are covered in the article on page 41.

Acknowledgments

The success of the HP 54600 mechanical design is to a large degree a result of the teamwork and the healthy "can do" attitude that existed not only in R&D but also throughout a number of other departments. In particular, the early brainstorming sessions were greatly enhanced by the creative efforts of Rich Huber, manufacturing engineer. Jerry Case, tooling engineer, added years of experience on injection-molding process and tooling issues and provided assistance on part design as well. Finally, the Colorado Springs Division model shop provided high-quality prototypes that proved to be an invaluable means of communicating and verifying design concepts.

EMC Design of the HP 54600 Series Oscilloscopes

By a combination of electronic circuit design and mechanical shielding techniques, the design meets German FTZ standards and, with optional shielding, most U.S. military standards for electromagnetic compatibility.

by **Kenneth D. Wyatt**

The EMC design of the HP 54600 Series digitizing oscilloscopes consisted of a combination of circuit board suppression and mechanical design techniques. Since the entire product (including enclosure) was a completely new design, we had an opportunity to design in RFI suppression techniques from the very start of the development. This article describes the design and test methods employed to ensure that the products met international and military EMC standards.

RFI standards for HP products typically include the German FTZ 1046 Class B limit (similar to the U.S.A. FCC Class B limit).^{*} Additionally, we decided to attempt to meet the U.S. MIL-STD-461C RFI requirements as specified in the environmental standard MIL-T-28800D. Meeting these military standards imposed much more stringent design goals than normal. Since one of the primary drivers of the HP 54600 design was cost, early EMC/RFI integration was necessary. By considering EMC early in the system design, we hoped to minimize the cost of achieving the multiple goals of both international and military standards. In the end we were forced to offer two of the more costly suppression methods for some of the military requirements as options. We felt that the additional cost involved should not be imposed on the majority of customers, who would not require the additional shielding.

Enclosure Design

We knew that older enclosure designs included too many seams for easy RFI suppression. Therefore, we decided to start with fresh ideas. Several concepts were suggested and we finally settled on a two-part molded design, coated on the interior with a 2.5- μ m-thick layer of vapor-deposited aluminum.^{**} One of the parts is the main cabinet and the other is the front panel. This two-part enclosure concept reduces the number of seams to the minimum required for access to the electronics. The aluminum is applied to all sides of the seam area so that good contact is made between the cabinet and the front panel.

^{*}Starting in 1992, the European EN55011 standard becomes effective. The HP 54600 Series oscilloscopes meet the Class A limit of this new standard.

^{**}More information about this design can be found in "Mechanical Design of the HP 54600 Series Oscilloscopes" on page 36.




Fig. 1. The main seam that joins the front panel to the cabinet is formed by a knife-edge on the cabinet that fits between a series of raised bumps and stiffening ribs on the front panel. Also see Fig. 7 on page 40.

The seam design is a continuous, overlapping "knife-edge", which fits in a slightly zigzag manner between a series of raised bumps and stiffening ribs. These stiffening ribs are spaced about 3 to 6 cm apart all the way around the seam, providing a distributed pressure to ensure a good electrical connection. Molded bumps between the contact fingers act as additional contact points. Fig. 1 shows the main seam details. This design provides uniform contact while allowing easy disassembly without the need for special tools. The maximum spacing of about 3.5 cm between contact points yields a theoretical shielding effectiveness of 20 dB at 500 MHz. We performed several tests on both the seam itself and the enclosure as a whole to confirm the shielding performance.

Enclosure Testing

Several evaluation techniques were used to assess the RFI performance of the enclosure before building a finished electronic prototype. A number of proof-of-concept tests were performed on a glued-together prototype enclosure. These tests included near-field seam leakage




Fig. 2. This photo shows an option module cover being measured using the shielding effectiveness fixture and a network analyzer. This yields a continuous plot of shielding effectiveness versus frequency for thin shielding materials. The technique allows a comparison of different shielding techniques.

tests using an HP 11940A close-field probe and far-field tests with a harmonic comb generator placed within the enclosure and measured as if it were a finished product.¹

The first enclosure testing consisted of a measurement of the basic shielding properties of various conductive plating choices. For this test, a pair of HP 11940A close-field probes were used tip to tip, one being the transmitting source and the other the receiver (Fig. 2). These probes were connected to an HP 8753B network analyzer and the transfer characteristic of the measurement system (without shielding material) was normalized to zero dB. The shielding samples were then placed between the probe tips and the near-field magnetic shielding effectiveness plot was obtained directly from the screen. This method allowed a quick comparison of several different shielding materials. Since the close-field probes were held by posts on the fixture, it was possible to measure various points of the actual molded parts to verify shielding performance and consistency.

Once molded prototypes were available, the comb generator was mounted within the enclosure. The comb generator simulated a very noisy product by emitting strong harmonics every 5 MHz from 30 to 1000 MHz. With the generator in place, seam emissions were measured with the 11940A close-field probe. A composite emissions characteristic was recorded using a spectrum analyzer and the signal leaks were noted. This emission recording was repeated for each mold trial to track the general progress of the design and to allow improvements to be incorporated into the mold before the electronic prototype was available.

Finally, far-field testing was performed, again using the comb generator as a simulated product. By measuring the difference in signal strength between the generator in the enclosure and the generator without the enclosure, a rough idea of the system shielding effectiveness (including all apertures and seams) was obtained. The video display module, keyboard, and steel deck were mounted

within the enclosure to fill the larger apertures and more closely resemble the finished product. The prototype enclosure was tested at a 3-meter distance in an anechoic chamber. First, the harmonic levels of the generator were measured. Then, the generator was placed inside the enclosure and the harmonics were remeasured. For each measurement, the product was rotated in azimuth and the highest harmonic levels were recorded. The difference in readings then indicated the worst-case shielding effectiveness. Fig. 3 shows a typical plot of shielding effectiveness versus frequency. This technique exposed weak areas of the total system design before installing the electronics and allowed the mechanical and electronic design efforts to proceed in parallel.

Together, these three evaluation techniques allowed us to determine weak areas of the enclosure design before completion of a prototype oscilloscope. The net result was a good enclosure design early enough in the product cycle so that the costs for RFI reduction were minimized.

Option Module Design

The HP 54600 oscilloscope includes option modules which plug onto the rear of the instrument and provide custom I/O and computer control. These modules are made of molded plastic and coated with the same vacuum metallization process as the main enclosure. The challenge from an EMC point of view was how to connect the module shield and main enclosure shield together well enough to prevent RFI when both coatings were on the inside of their respective moldings.

The option module is connected to the main enclosure by a 50-pin connector and three metal-coated hooks which fit into three openings in the main cabinet. The connector (similar in style to an HP-IB connector) includes a grounded shell extension which overlaps the raised shell of the mating connector. We had originally hoped that these overlapping ground connections and the three hooks would adequately join the module and main cabinet shields together. Unfortunately, the connector half in the option module was soldered directly to the circuit




Fig. 3. Measured shielding effectiveness of the HP 54600 oscilloscope enclosure including seams and apertures. The harmonic comb generator was used to obtain the plot.

board inside and thus was not grounded well enough to the module shield, so the initial result was very strong emissions once the module was attached to the oscilloscope.

When dealing with enclosure design, it is useful to consider each enclosure (in this case, the module and oscilloscope) as a separately shielded EMC environment. Any untreated penetration of either environment (by a wire, or by a connector in this case) will allow internally generated noise currents to escape, thus reducing the effectiveness of the environment's shield. Since the connector of the module was not directly connected to the module shield, internal noise currents were allowed outside the shield at the connector resulting in an increase in emission level. To fix this, an EMI finger stock strip was staked directly to the module enclosure half and against the metallic coating as shown in Fig. 4. The flexible beryllium copper fingers press firmly to the ground shell of the module connector and reduce the emission level by 3 to 8 dB.

Circuit Suppression Design

One of the real keys to our radiated emissions success was the source suppression techniques used on the circuit board. Three methods were used: decoupling of the crystal oscillators, proper bypassing of each IC, and RC filtering of clock signals at the primary buffer gate.

It is well-known that crystal oscillators or other high-frequency periodic signals (such as clock lines) are among the primary sources of radiated emissions on circuit boards.² Three crystal oscillators are used in the circuit design: 10.73, 40, and 80 MHz. These high frequencies required that close attention be paid to circuit trace layout and board design. All oscillator circuits are well-bypassed and decoupled from the rest of the circuit board with an L filter consisting of a small 77-ohm ferrite bead and a 0.1- μ F multilayer ceramic chip capacitor.




Fig. 4. A finger stock strip is used to reduce the radiated emissions from the option module. The fingers contact the module connector ground shell.

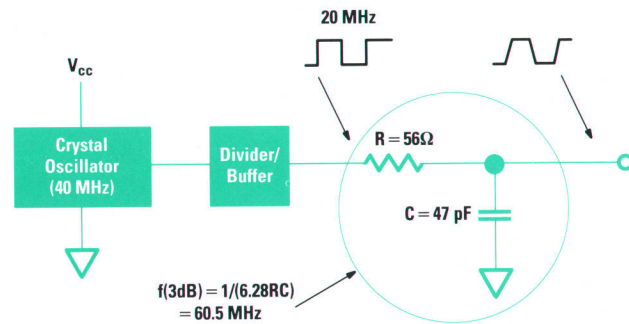


Fig. 5. An RC filter is installed following the microprocessor divider/buffer circuit. The filter slows down the rise time of the clock signal and greatly reduces the clock harmonics.

All ICs are liberally bypassed from the supply pin to ground with 0.01- μ F or 0.1- μ F chip capacitors. Each package has one or more as appropriate. The larger ASICs have both bypass capacitors and filter networks to control noise generation.

The last primary RFI reduction technique used is a simple RC filter network on the microprocessor clock buffer/divider (Fig. 5). Clock rise times are a major cause of emissions and are often 2 ns or less, depending upon the logic family used. Very often, these fast edges are unnecessary.

If the Fourier series of a trapezoidal waveform is analyzed, it will be noted that the envelope containing the noise current harmonics starts rolling off at a 20-dB/decade rate and then decreases at a 40-dB/decade rate with a break frequency that depends only on the rise (or fall) time of the square wave (Fig. 6). By slowing the edges of the clock signal, the harmonic content can be substantially reduced. Radiated clock harmonics have been reduced by 6 to 12 dB merely by adding this simple RC network.




Fig. 6. The envelope of the Fourier series of a trapezoidal waveform has a breakpoint that depends purely upon the rise time (τ_r). Slowing down the rise time shifts this breakpoint down in frequency and reduces the amplitude of the high-frequency harmonics.

Related to the edge speed problem is the use of fast devices. In one case, a transistor with an f_T of 5 GHz was used to amplify an 80-MHz trigger signal. The rise time for this device was 2 ns and we were measuring a strong harmonic at 800 MHz. By changing to a different device with an f_T of 1 GHz and a rise time of 3.5 ns, the harmonic amplitude was reduced by 6 dB.

Printed Circuit Board Design

The circuit board is a six-layer, single-sided, surface mount design. Full ground planes are used to provide both a ground plane and additional isolation between trace layers. All secondary power is isolated with L filters consisting of a multiturn ferrite core and multiple chip capacitors at the board power connector. This prevents circuit board noise currents from traveling out the power supply wiring.

Another suppression technique is component grouping by function and speed. The power input and filtering are located at the rear of the board, well away from other sensitive or noisy connectors. The crystal oscillators are located close to the power filter section and away from sensitive analog inputs. The analog circuitry is located on one side of the board, while the high-speed digital circuitry is located on the other. Plenty of filtering and decoupling is used on all supply traces to prevent noise currents from flowing from noisy areas to sensitive areas of the board. Much of the analog circuitry is contained within a cast shield to prevent interference from external fields.

The circuit board is also attached to the steel deck, which serves three purposes. Besides serving as a structural portion of the system, the deck provides about 10 dB of low-frequency (60 Hz) shielding and also serves as an image plane.³ An image plane acts to confine external fields generated on the circuit board to the area between the board and the plane. The large metal surface forces any electric (E) and magnetic (H) field lines impinging on its surface to rotate orthogonally. As the image plane is moved close to the board, this forced rotation tends to cancel both the E and H fields that originate on the board. This technique (which is free) further reduces the radiated emission sources within the oscilloscope and helps lower system RFI.

Power Supply Filtering

Design constraints on the power supply filter circuit made for another major challenge. The oscilloscope must be able to use any power source worldwide, so it has to operate from 90 to 270 volts and from 44 to 440 Hz, automatically, without user-settable switches. We also had to meet safety leakage current limits (3.5 mA, maximum) at any voltage/frequency combination, and we had to meet both FTZ (Vfg) 1046 Class B (based upon the German VDE 0871 standard) and MIL-STD-461C, CE-03 (U.S. military) conducted emission limits. Finally, because of cost constraints, we could not use shielded power entry modules (PEMs). The entire filter circuit had to be designed with discrete components mounted on the power supply board. To say that this was a challenge would be an understatement.

One major difference between FTZ 1046 and MIL-STD-461 conducted emissions tests is that for FTZ 1046, noise voltages are measured, while for MIL-STD-461, noise currents are measured. Thus, the filter design for FTZ 1046 needs to look like a low impedance (shunt capacitor to limit the voltage), while the design for MIL-STD-461 needs to look like a high impedance (series inductance to limit the current). Thus, the filter topologies chosen for the two standards tend to be diametrically opposed.

Our filter vendor ultimately proposed a filter that resembles a typical FTZ 1046 topology with series inductors to reduce the MIL-STD-461 noise currents. The circuit is shown in Fig. 7.

System Design

Since the circuit board assembly contains connectors at each end, the usual nuts that attach the oscilloscope probe BNC connectors to the front panel were eliminated to avoid stresses on the enclosure. The BNC connectors are part of a cast shield assembly that attaches directly to the main circuit board and then to the steel deck with a screw. These connectors are simply pushed through matching holes in the molded plastic front panel as the circuit board is installed. This was another EMC design challenge in that we were originally counting on a fairly long path to ensure an adequate ground connection (or cable shield termination) between the BNC ground and the shielded enclosure of the oscilloscope. The result was a higher level of probe cable radiation than we could tolerate.

The solution to this is a thin, nickel-coated, phosphor-bronze shim placed between the cast BNC shield assembly and the front panel. This shim includes four connection points (90 degrees apart) which press around the perimeter of each BNC connector. Other tabs press between the deck and front panel. The shim provides a good connection between the circuit board and the enclosure, thereby reducing probe cable radiation.

The other system problem stems from the use of a raster-scanned CRT display. Large magnetic fields (60-Hz vertical and 25-kHz horizontal) are used to drive the electron beam, and these fields radiate from the oscilloscope display. The 60-Hz field can interfere with other nearby CRT displays, and vice versa. Also, other low-frequency fields (from power transformers, typically) can interfere with the oscilloscope display, causing the display to swim or wiggle. Finally, the 25-kHz field radiates strongly out the front of the CRT and can be picked up




Fig. 7. This is the final circuit used for the switching power supply filter. It will allow the HP 54600 Oscilloscope to meet both FTZ 1046 and MIL-STD-461 conducted emission limits.

by the oscilloscope probe tip as far away as one to two feet. If the user measures sensitive circuits directly in front of the CRT, they will tend to pick up this horizontal scanning signal. Passage of the radiated emission and susceptibility portions of MIL-STD-461C requires both of these fields to be considerably reduced. The optional shields previously mentioned were designed to help with both of these problems. While the product meets commercial EMC standards worldwide, some customers may need the additional shielding provided by these options.

Option 001 is a mumetal shield that snaps over the CRT deflection coils. This both reduces the external fields exiting the oscilloscope and, more important, reduces the susceptibility of the display to low-frequency external fields (such as power transformers). Unfortunately, even with the shield in place, we are unable to comply fully with the RE01 (magnetic radiated) or RS02 (magnetic susceptibility) tests of MIL-STD-461C because the fields from these tests can enter through the front of the CRT display and upset the displayed waveform by moving the electron beam. This is a disadvantage of all raster-scanned CRT displays.

Option 002 is a conductive filter shield that is attached in front of the display screen. This shield reduces the E fields at the horizontal scan frequency of 25 kHz exiting the front of the CRT display. This filter screen allows the oscilloscope to meet the RE02 emission requirements of MIL-STD-461C. Additionally, it solves the problem of 25-kHz CRT emissions entering the oscilloscope probe tip.

Acknowledgments

I would like to acknowledge the entire HP HP 54600 design team for their assistance in preparing this article and in persevering throughout the EMC design effort. I would also like to thank Dean Chaney for all the EMC testing.

References

1. A more detailed description of the harmonic comb generator measurement techniques along with construction details can be found in K. Wyatt and D. Chaney, "RFI Measurements Using a Harmonic Comb Generator," *RF Design Magazine*, January 1991, pp. 53-58.
2. H. Ott, *Noise Reduction Techniques in Electronic Systems*, 2nd edition, Prentice-Hall, 1988.
3. R. German, H. Ott, and C. Paul, "Effect of an Image Plane on Printed Circuit Board Radiation," *IEEE International Symposium on EMC*, Washington, D.C., August 21-23, 1990, pp. 284-291.

Digital Oscilloscope Persistence

Autostore, a storage technique for monochrome digital storage oscilloscopes, displays historical traces at half intensity and the most recent, or live, trace at full intensity. The technique allows new ways of viewing signals.

by James A. Kahkoska

The advent of digital oscilloscopes brought new techniques of simulating analog persistence and storage using raster CRT technology. Digital raster persistence in digital storage oscilloscopes simulates the screen persistence found on analog storage oscilloscopes. To date, various forms of digital raster persistence have existed. The latest form for monochrome digital storage oscilloscopes allows the user to view both the current signal and all worst-case signal excursions. This storage technique can be applied to a wide variety of real-world problems, ranging from its use as a simple "electronic grease pencil" to complex CPU system analysis.

To simulate analog persistence, digital variable persistence is used. Each pixel that is excited as the signal is displayed is time-tagged. The time tag allows the pixel erasure time to be determined based on the user-selected persistence time.

The most powerful use of digital persistence is infinite persistence. Pixels are never erased, so all signal excursions are accumulated indefinitely without the fading or blooming commonly associated with analog storage oscilloscopes. Infinite persistence is invaluable for catching infrequent events or making worst-case measurements. One drawback of infinite persistence is that the live trace information is lost in the historical signal excursions, making it difficult to see the current signal.

Autostore

To overcome this problem, Hewlett-Packard has introduced a new form of infinite persistence, called autostore. In autostore mode, half intensity is used to display all historical signal excursions while the current signal is shown at full intensity. Simple one-key operation




Fig. 1. Comparison of an autostored good unit signal at half intensity with a bad unit signal at full intensity.

allows the user to move into and out of autostore mode quickly.

A common troubleshooting technique used by technicians is comparison of a good unit with a bad unit. Either operator memory or a grease pencil is used to remember the signal characteristics of the good unit. Autostore allows fast, accurate signal comparison with the minimum number of keystrokes. The user simply enables autostore briefly to create a half-intensity image of the good signal, then probes the bad unit for comparison. Small signal details that might otherwise go unnoticed are easily identified. To remove the electronic grease pencil, the user simply presses **Erase**. With three keystrokes, the user can record, compare, then erase the image. Fig. 1 shows the comparison of an autostored good unit signal with the bad unit signal at full intensity. The slight variation in pulse width might go unnoticed if the two units were compared visually.

The main strength of autostore is the ability to view all signal excursions while distinguishing the current live trace. This allows the user to adjust the unit under test while capturing the entire range of the adjustment. For making adjustments to minimize overshoot, jitter, or noise, autostore provides visual feedback without which it would be difficult to see the effects of adjustments when the signal changes become small. For example, for a fine adjustment to minimize overshoot, adjustments are made so that the current full-intensity trace is always bounded by the historical trace information. Autostore provides details down to pixel resolution to verify that the overshoot is truly minimized (see Fig. 2).

Autostore is also useful for building a set of characteristic curves. Coupled with the digital oscilloscope's ability to capture single-shot events and view negative time, autostore allows a family of single-shot events to be captured on the screen, the most recent at full intensity. This technique can be applied when characterizing single-shot events over a range of specified conditions. For example, an important criterion for power supply testing is the power-up characteristics of the supply over a wide range of loads. Fig. 3 shows the power-up characteristics of a power supply over the specified range of current capacity, with the most recent power-up cycle at full intensity.




Fig. 2. Minimizing overshoot by adjusting a circuit so that the current live trace at full intensity is always within the autostored historical traces at half intensity.

Worst-case signal variations, such as jitter or noise, are easy to analyze using conventional infinite persistence, but the current state of the system becomes lost as the traces accumulate. This makes it difficult to analyze the effect of any change to the system, such as adjustments, component changes, or temperature, relative to the worst case. Autostore solves this by displaying all but the current trace at half intensity. Fig. 4 shows a jittery signal with the range of jitter at half intensity and the current signal at full intensity.

Digital System Troubleshooting

For digital system troubleshooting, the ability of autostore to capture a full spectrum of valid and invalid transitions is invaluable. For example, the normal read cycle of a CPU system results in exactly one device at a time being enabled to drive the data bus. If the CPU occasionally reads bad data, two likely causes are either a bus conflict or a floating bus.

Bus conflict results when two different devices attempt to drive a data line at the same time. This conflict may be caused by either a bad bus device or bad decoding of the chip select signals so that two bus drivers are enabled at once. The result is an invalid voltage level between logic




Fig. 3. Power-up characteristics of a power supply over the specified load current range, with the most recent cycle at full intensity.




Fig. 4. A jittery signal is shown with the range of jitter at half intensity and the most recent signal at full intensity.

levels. This conflict may only occur when different output levels are present on each bus driver, resulting in a "soft error" that may be difficult to diagnose because of its context dependence. The upper trace in Fig. 5 represents the read enable signal and the lower trace represents the data line. Normal bus operation results in the data line's being valid—either a one or zero as indicated by the markers—on the positive edge of the read enable signal. Exercising the system, perhaps with a memory test, shows that one of the read cycles generated an invalid level because of a bus conflict. With autostore, not only can the anomaly be captured with half-intensity infinite persistence, but also the frequency of occurrence can be seen as full-intensity variations. Normal infinite persistence makes it impossible to see the frequency of occurrence once the anomaly occurs because there is no intensity variation between the infinite persistence traces and the current trace. Seeing the current trace at full

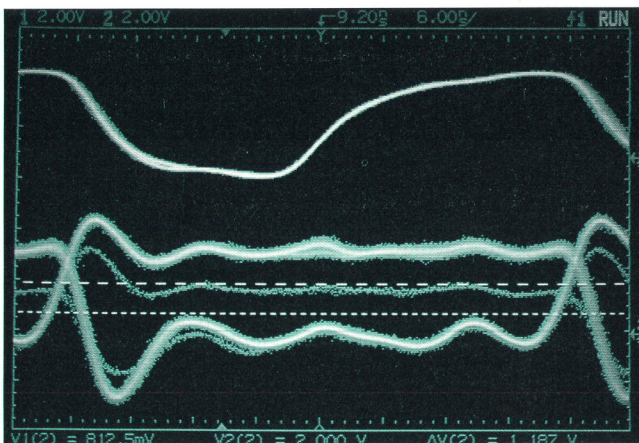


Fig. 5. The upper trace is the read enable signal in a CPU system and the lower traces represent data on the data bus. One of the read cycles resulted in an invalid level on the data bus.




Fig. 6. VCR power supply current as a function of supply voltage during power-up.

intensity allows the user to attempt to localize the problem, perhaps by using a software test to stimulate different CPU addresses and data values.

XY Mode

Autostore can be used in any horizontal mode, such as delayed sweep or XY. In XY mode, one application of autostore is to measure the peak power used by a device at power-up or as the load is changed. A current or differential probe is used on the Y axis, allowing a voltage-versus-current graph to be drawn. A given power level is represented by a diagonal on the graph, since power is voltage multiplied by current. Autostore allows the complete power-up characteristic and the live power consumption to be graphed. For example, if a VCR power supply were being monitored, power-up and other VCR modes could be selected while autostore captures the peak voltage, current, and power excursions. Seeing the live trace at full intensity allows the live voltage, current, and power to be distinguished from the historical data. Fig. 6 shows the power-up cycle of a digital circuit. The power-off point is at the lower left. The current drawn, represented by the Y axis, begins to climb as the capacitors on the board charge up. Once they are charged, the current begins to fall as the voltage stabilizes. The peak power consumption during power-up is captured in the upper right corner. Once the system has stabilized, at the lower right in Fig. 6, the user can go on to observe various system features, such as rewinding a tape, to examine the voltage, current, and power effects.

Summary

Autostore, with its ability to view both historical signal changes and the current signal, has numerous applications in adjustment, troubleshooting, and design. Simple one-button operation eliminates the complexity of analog or digital persistence controls, allowing all levels of operators a new way of viewing signals.

A High-Resolution, Multichannel Digital-to-Analog Converter for Digital Oscilloscopes

This 16-bit, 16-channel DAC is used for microprocessor adjustment of fourteen dc signals that control the analog section of the main oscilloscope board in the HP 54601A digitizing oscilloscope. It also provides a high-accuracy dc reference for calibrating the vertical gain.

by Grosvenor H. Garnett

At the time of manufacture, oscilloscopes usually require several internal factory calibration adjustments, such as vertical gain, trigger hysteresis, and others. Many of these calibration adjustments involve applying a precise dc voltage to the circuit being calibrated. In the past, these dc voltages were developed across potentiometers and the calibration was accomplished by having production calibration personnel adjust these potentiometers while observing instrument parameters or output.

The operation of an oscilloscope often requires operator input to control such functions as trigger level and offset. In the past this often involved adjustment of front-panel potentiometers that applied accurate dc voltages to the internal circuits of the oscilloscope.

Digital oscilloscopes have made it desirable to bring calibration and operating controls under microprocessor control. The digital-to-analog converter (DAC) makes this possible, since it provides a precision dc output voltage

from a digital input signal. This article describes a high-resolution, multichannel DAC developed for HP digital oscilloscopes.

General Description

The multichannel DAC is an indirect type that takes in digital serial data, converts this data to a pulse signal of variable rate, width, and level, and then applies this signal to low-pass filters to produce the analog voltage output. The DAC is implemented in three circuit blocks—a digital section, an output voltage stabilizer section, and an output filter section for each channel—as shown in Fig. 1.

The specifications and characteristics of this DAC are:

- Number of channels: 16.
- Maximum resolution: 16 bits. Each channel will operate as an n-bit DAC if that channel is loaded with n bits of data, where $n \leq 16$.
- Linearity: 12 bits.

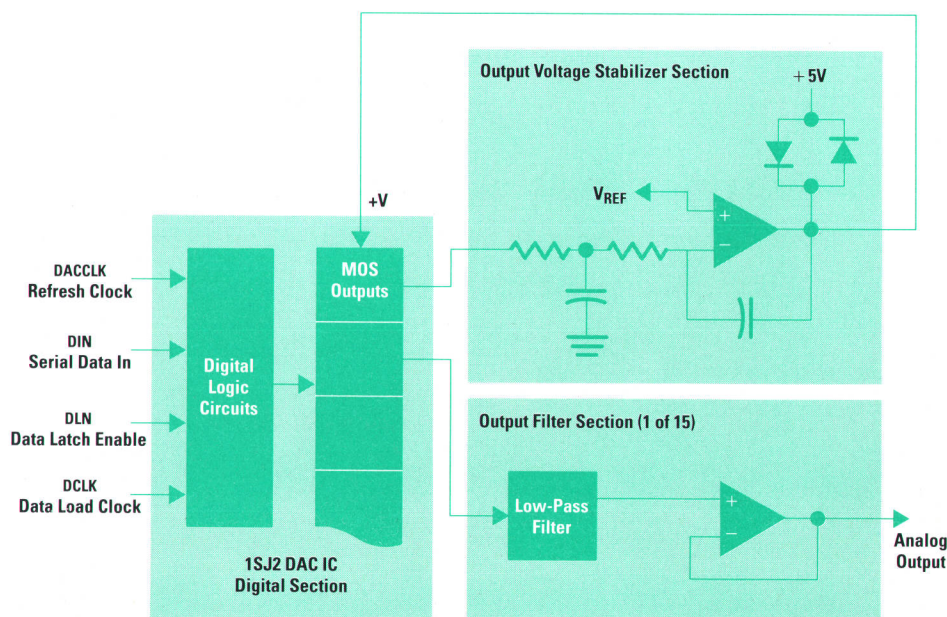


Fig. 1. Block diagram of the high-resolution, multichannel, indirect type digital-to-analog converter.

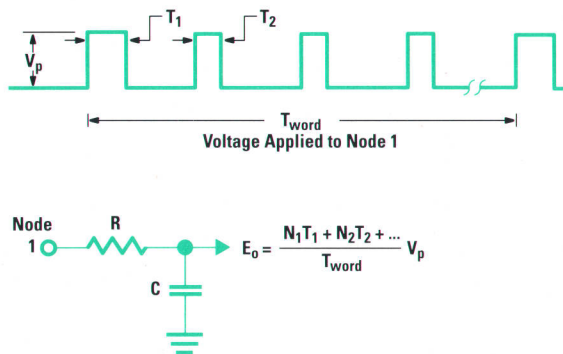


Fig. 2. Filter section input waveform and output voltage.

- Maximum refresh clock frequency input: 20 MHz.
- Maximum data load clock frequency input: 20 MHz.
- Processor interface: addressed serial data bus.
- Self-refreshed (requires processor service only when new data is loaded).
- Input digital signal compatibility: standard TTL.
- Operating speed: settles to within $\pm 1/2$ LSB in approximately 50 milliseconds for 16 data bits.
- Digital section: custom HP 1SJ2 CMOS DAC IC in a 0.3-inch-wide 24-pin plastic dual in-line package.

The 16-channel digital section is implemented in a custom 24-pin CMOS integrated circuit designated 1SJ2. This IC

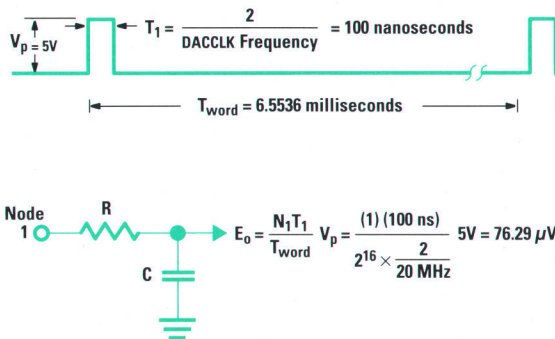


Fig. 3. Filter section input waveform and output voltage for data bit D0 only high, 16-bit data, and a DACCLK frequency of 20 MHz. The input consists of one 5V, 100-ns pulse per T_{word} period.

provides all digital signal processing and outputs a precisely controlled variable pulse stream to the output voltage stabilizer and the output filters.

The output voltage stabilizer circuit stabilizes the filter output voltages against variations of power supplies, clock duty cycle, and digital output gate timing. It does this by regulating the +5Vdc supply voltage to the 1SJ2, which controls the level of the pulses out of the 1SJ2. This stabilizer circuit, which is composed of external

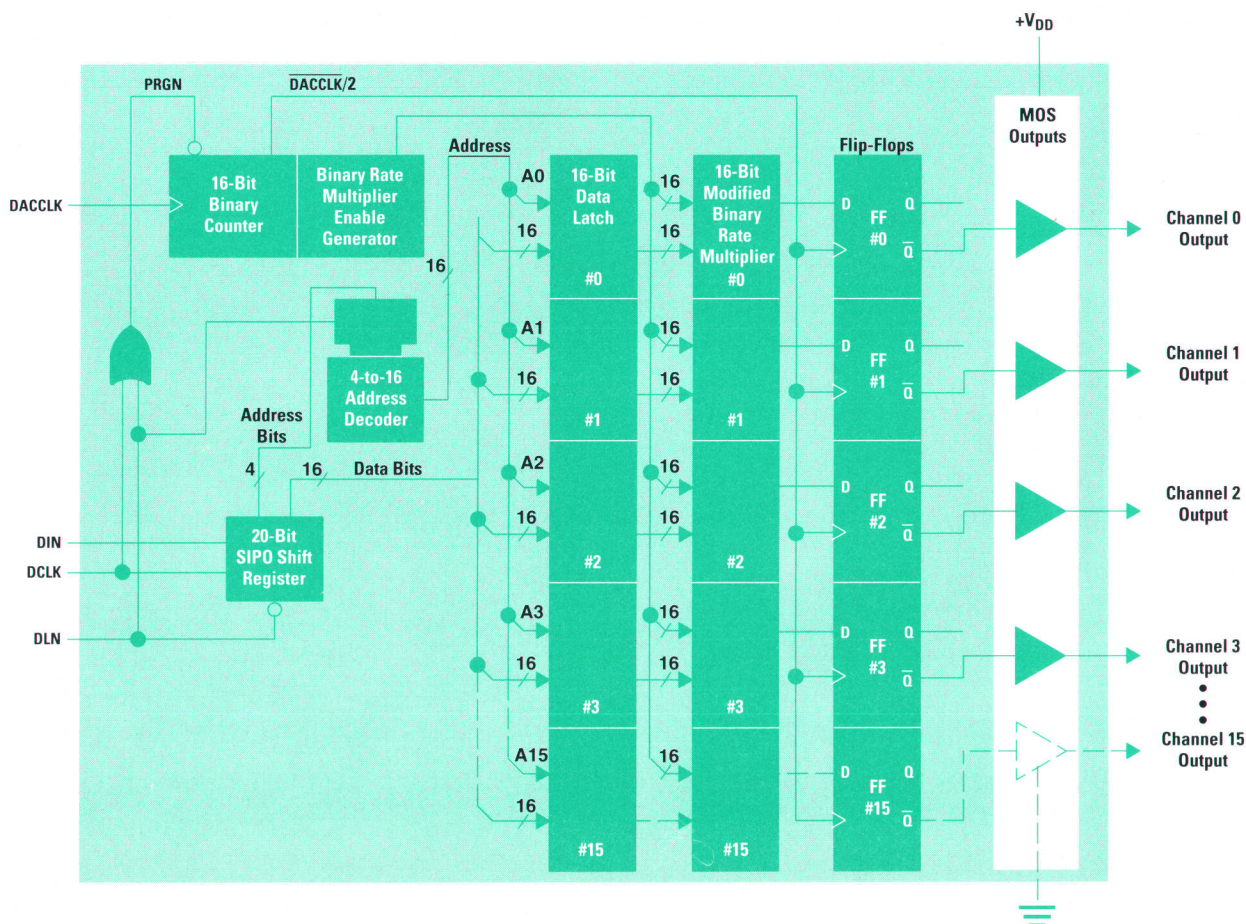


Fig. 4. Block diagram of the 1SJ2 integrated circuit, which is the digital section of the multichannel DAC.

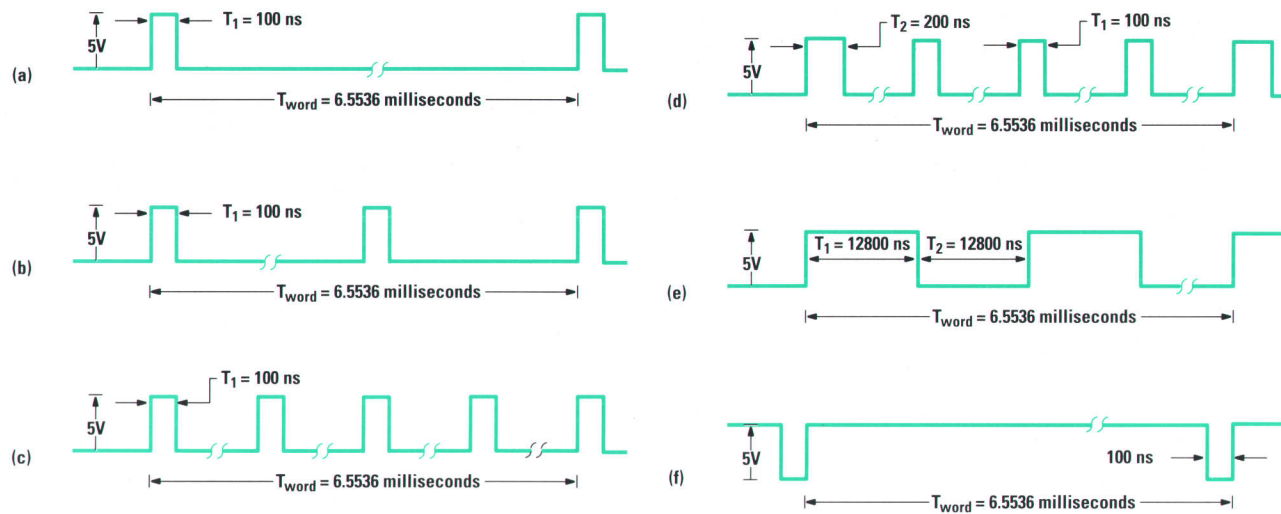


Fig. 5. 1SJ2 DAC IC output for 16-bit data and a 20-MHz DACCLK frequency. (a) D0 only high gives one 5V, 100-ns pulse per T_{word} period. (b) D1 only high gives two 5V, 100-ns pulses per T_{word} period. (c) D8 only high gives 256 5V, 100-ns pulses per T_{word} period. (d) D8 and D0 only high gives 255 5V, 100-ns pulses and one 5V, 200-ns pulse per T_{word} period. (e) D15 only high gives 256 5V, 12800-ns pulses per T_{word} period (50% duty cycle). (f) D0 through D15 all high gives one 5V, 6.5535-ms pulse per T_{word} period.

off-chip discrete components, also provides a means of calibrating the DAC.

One of the 16 DAC channels is used to provide an input to the output voltage stabilizer circuit. There is an output filter section for each of the other 15 channels.

The output filter section receives the variable rate/width/level pulse string output from the 1SJ2 and integrates it to provide the analog output voltage. The filters are external off-chip discrete components. This type of implementation provides the opportunity to reduce capacitively or inductively coupled interference by locating the output filter at the receiving end of the DAC output and using the filter for both output smoothing and interference reduction. Depending on the needs of the channel, this section may only contain a passive low-pass filter. It can also contain active elements to provide gain and offset and to lower the filter output impedance.

Operation

To understand how this type of DAC works, refer to Figs. 1 and 2. The 1SJ2 DAC IC and the output voltage stabilizer in Fig. 1 generate a variable rate/width/level pulse string in response to the input data. This continuously refreshed pulse string is applied to the output filter, which integrates the pulse string to produce the analog voltage output.

Fig. 2 shows the output filter voltage for a general pulse string, where N_1 is the number of pulses of width T_1 , N_2 is the number of pulses of width T_2 , and so on. Fig. 3 shows the output filter voltage for input data bit D0 (LSB) only high in a 16-bit system with a 20-MHz refresh clock (DACCLK). The output filter average voltage output is independent of the DACCLK frequency and the values of R and C of the output filter.

With only a few restrictions, any one of a number of different types of low-pass filters can be used for the output filter.

Digital Section Integrated Circuit

Fig. 4 shows a block diagram of the 1SJ2 IC. All digital inputs are TTL-compatible. The processor interface is an addressed serial data bus. The input signals are:

- DACCLK. Refresh clock.
- DCLK. Data load clock.
- DIN. Serial data in. The data bits are serially shifted into a 20-bit series in, parallel out (SIPO) shift register. The first bit shifted in is the data LSB (least-significant bit). The last bit shifted in is the output address MSB (most significant bit). For 16-bit data, the sequence is:
A3 A2 A1 A0 D15 D14 D13 D12 D11 D10 D9 D8 D7 D6 D5 D4 D3 D2 D1 D0
- DLN. Data latch enable. Latches data bits into the appropriate transparent data latch inside the IC as defined by the address bits. Also, on its rising edge, sets up the SIPO data and address register inside the IC to be cleared by the first rising edge of DCLK.

The variable pulse string out of the 1SJ2 IC is pulse width modulated by the eight most-significant input data bits and pulse rate modulated by the remaining input data bits.

The maximum number of output pulses from the 1SJ2 IC is 256 per DAC word. The period of the DAC word, T_{word} , is the time required for a counter in this IC to count 16 data bits. Since the DACCLK clock is divided by two in this IC:

$$T_{\text{word}} = 2^{16} \frac{2}{\text{DACCLK frequency}}$$

The period of a 16-bit DAC word with a 20-MHz DACCLK is:

$$T_{16} = 2^{16} \frac{2}{20 \text{ MHz}} = 6.553600 \text{ milliseconds}$$

Fig. 5 shows the pulse outputs from the 1SJ2 DAC IC with different data inputs and a 20-MHz DACCLK. With all

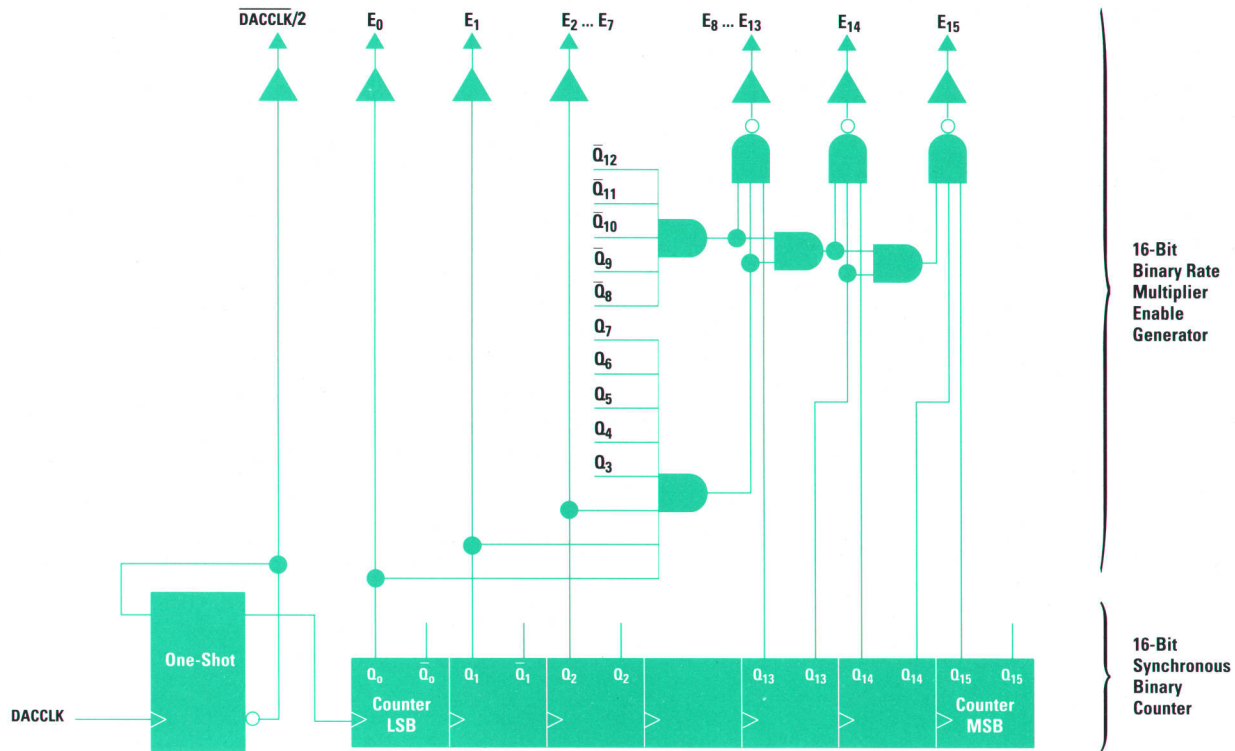


Fig. 6. 16-bit synchronous binary counter and binary rate multiplier enable generator.

data bits low, the 1SJ2 will not output any pulses (not shown in Fig. 5). When the LSB, data bit D0, only is high, the 1SJ2 will output one 5V, 100-nanosecond pulse per T_{word} period. When data bit D1 only is high, the 1SJ2 will output two 5V, 100-nanosecond pulses per T_{word} period.

For each data increase of one LSB, the 1SJ2 will output one additional 100-nanosecond pulse per T_{word} period until data bit D8 only is high. With data bit D8 only high, the 1SJ2 will output 256 100-nanosecond pulses (the maximum number of pulses) per T_{word} period. When data bits D8 and D0 only are both high, the 1SJ2 will output 255 100-nanosecond pulses and one 200-nanosecond pulse per T_{word} period.

As the input data increases, one of the 1SJ2 output pulses will become 100 nanoseconds wider for each input data LSB increase. When MSB data bit D15 only is high, the 1SJ2 will output 256 5V, 12800-nanosecond pulses per T_{word} period, which is the 50 percent duty cycle pulse output. When all input data bits D0 through D15 are high, the 1SJ2 will output one 6.5535-millisecond pulse per T_{word} period.

As already explained, up to 16 data bits followed by four channel output address bits are serially shifted into the 20-bit SIPO shift register by the data load clock DCLK (see Fig. 4). The first bit shifted in is the data LSB, and the last bit shifted in is the channel output address MSB. Since the register is automatically cleared each time before data and address bits are loaded, fewer than 16 data bits can be used.

The data latch enable signal DLN is cycled one time to latch the data bits from the SIPO shift register into the appropriate 16-bit transparent data latch as directed by

the address bits. The output of each data latch provides the data bits to one of the 16-bit modified binary rate multipliers. The binary rate multiplier enable generator provides the enable signals to the modified binary rate multipliers.

Fig. 6 shows a block diagram of the 16-bit synchronous binary counter and binary rate multiplier enable generator. The refresh clock DACCLK is divided by two and converted to a 50% duty cycle in this block. The output count of the 16-bit synchronous counter drives the 16-bit binary rate multiplier enable generator, which generates the binary rate multiplier enable pulses. All 16 of these pulses are applied to each of the modified binary rate multipliers.

The modified binary rate multiplier is shown in Fig. 7. The lower eight data bits along with the upper eight enable pulses are applied to a standard binary rate multiplier to provide pulse rate modulation. The upper eight data bits along with the lower eight enable pulses are applied to two four-bit magnitude comparators to provide pulse width modulation.

Each output of a modified binary rate multiplier is clocked through a D-Q flip-flop by the negative phase of the DACCLK/2 pulse (Fig. 4). The D-Q flip-flops drive the MOS outputs, which provide the final 1SJ2 DAC IC pulse modulated output.

Digital signal processing occurs on the positive phase of DACCLK/2, while output pulse gating occurs with the negative phase of DACCLK/2. This timing scheme makes the digital signal processing timing noncritical and provides precisely gated output pulses from the 1SJ2.

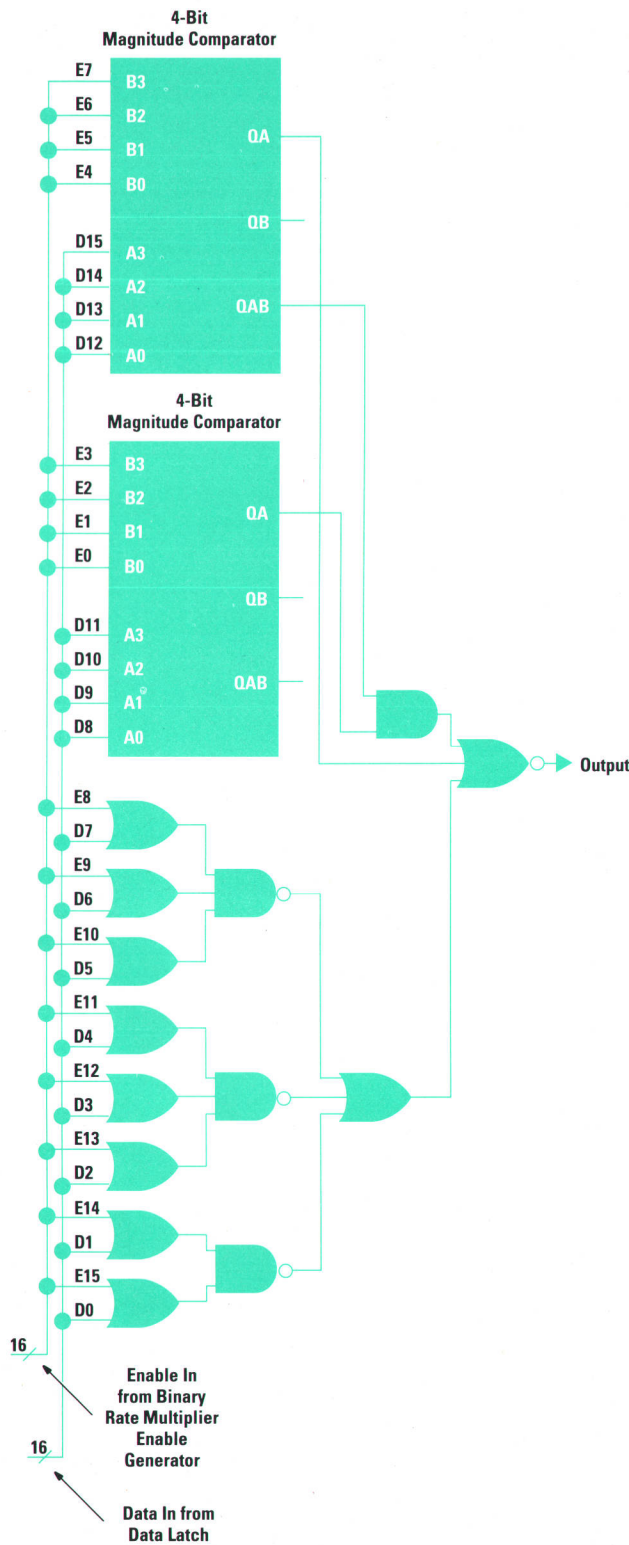


Fig. 7. Modified binary rate multiplier.

Output Voltage Stabilizer

The output voltage stabilizer circuit plays the major role in setting and maintaining the high accuracy of the DAC. This circuit uses a high-gain operational amplifier to compare the integrated pulse output voltage of channel 7

(the control channel) of the 1SJ2 to a very accurate and stable voltage reference. The output of this operational amplifier drives a transistor, which completes the negative feedback loop that regulates the positive supply voltage to the 1SJ2. Any small voltage changes at the output of channel 7 of the 1SJ2 are greatly amplified and then applied back to the 1SJ2. This negative feedback to the 1SJ2 positive power supply controls the pulse level or amplitude out of the 1SJ2, thereby effectively canceling the voltage variations in all channels of the 1SJ2.

The design of the output voltage stabilizer circuit can vary depending on the application, but the gain of the operational amplifier should ideally be 150V/mV or greater.

Fig. 8 shows an output voltage stabilizer circuit designed to use the Linear Technology LT1021CCN8-5 precision voltage reference. The accuracy and stability of the DAC depend on the accuracy and stability of this voltage reference. With this circuit, the DAC can be calibrated by loading the control channel (channel 7) with all 16 data bits high immediately after each instrument power-up.

Output Filters

The output filters receive the variable rate/width/level pulse string output from the 1SJ2 DAC IC, integrate this pulse string, and produce the analog voltage output.

Although a number of different types of low-pass filters can be used, the one necessary requirement is that each of these filters present a load of no less than 100,000 ohms to the 1SJ2 DAC IC channel output. If the filter input impedance at 39.06 kHz is lower than 100,000 ohms, DAC linearity will be degraded. A typical inverting filter that might be used is shown in Fig. 9. A typical noninverting filter is shown in Fig. 10.

The filter output voltage change for an input data change of one LSB is given in Fig. 11 for several data inputs. In a typical application, the output voltage bit size is usually specified and the number of data bits is calculated. The output voltage bit size is given by:

$$V_{\text{LSB}} = \frac{\text{Output Voltage Range}}{2^n} = \frac{(\text{Max. Volts Out}) - (\text{Min. Volts Out})}{2^n}$$

where n is the number of data bits. The number of data bits can be calculated from:

$$n(\text{bits}) = \frac{\log \frac{\text{Output Voltage Range}}{V_{\text{LSB}}}}{\log 2}$$

Output Filter Attenuation Characteristics

Since the output filters receive a pulse string as an input, and since these filters have finite attenuation, the analog output voltage from these filters will contain some ripple frequencies of the input pulse string.

(continued on page 55)

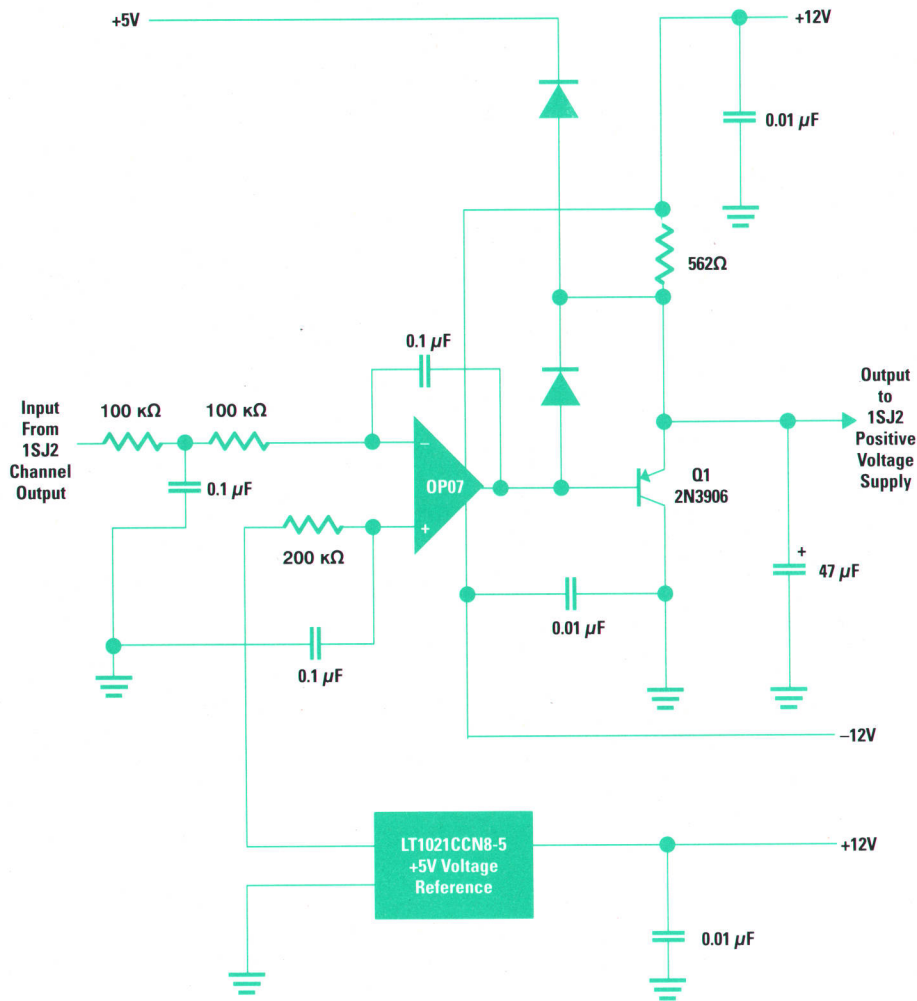
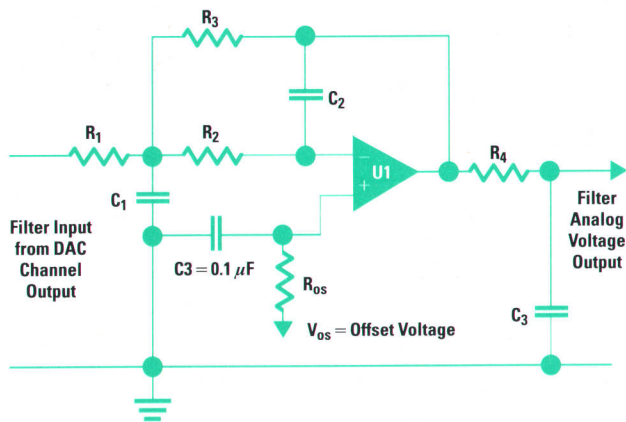


Fig. 8. DAC output voltage stabilizer.

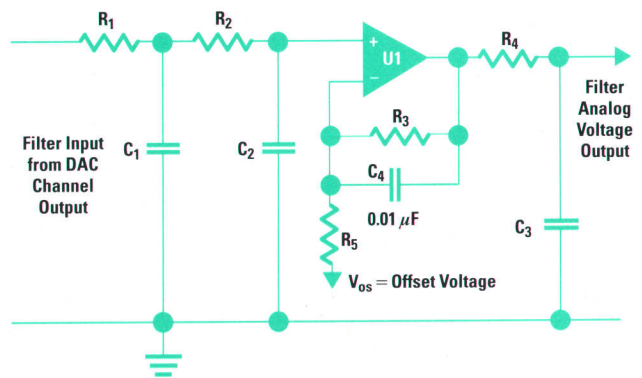


$$\text{Gain} = \frac{R_3}{R_1}$$

$$Q = \frac{(C_1/C_2)^{1/2}}{(R_3/R_2)^{1/2} + (R_2/R_3)^{1/2} + (R_3R_2/R_1)^{1/2}}$$

$$\omega_0^2 = \frac{1}{R_2 R_3 C_1 C_2}$$

Fig. 9. Typical inverting DAC filter.



$$\text{Gain} = 1 + \frac{R_3}{R_5}$$

$$Q = \frac{R_1^{1/2} R_2^{1/2} C_1^{1/2} C_2^{1/2}}{R_1 C_1 + R_2 C_2 + R_1 C_2}$$

$$\omega_0^2 = \frac{1}{R_1 R_2 C_1 C_2}$$

Fig. 10. Typical noninverting DAC filter.

Using the High-Resolution, Multichannel DAC in the HP 54601A Oscilloscope

The HP 54601A digitizing oscilloscope requires the microprocessor adjustment of 14 dc signals that control the analog section of the main oscilloscope board (see Fig. 1). In addition, the HP 54601A needs a high-accuracy adjustable dc reference so that the vertical accuracy can be calibrated within 1.5%. The high-resolution multichannel DAC described in the accompanying article meets both requirements. This single IC, with its associated filtering circuitry, allowed the HP 54601A designers to reduce board cost, self-calibration time, and cost of ownership while ensuring that the oscilloscope can meet accuracy specifications over a wide temperature range. The 14 dc voltages the multichannel DAC generates control the ADC references (2), postamp offsets (2), input offsets (4), programmable preamp gains (2), trigger level (1), trigger hysteresis (1), time-interpolator threshold (1), and time-interpolator slope (1).

Offset voltages are generated by the DAC from an equation of the form $y = mx + b$ (see Fig. 2a). Since the offset circuit has excellent linearity, any voltage between $\pm 2V$ can be generated precisely as long as the slope (DAC counts/volt) and intercept (DAC value for 0V) are known. The firmware generates these values during a self-calibration by outputting 0V and 3V (the offset circuit has some headroom) from the dc calibration output into a channel's input. Once the offset circuit's characteristics are known, it can be used as the stimulus for the gain portion of the vertical calibration instead of the dc calibration output. The reason for using the offset input as the stimulus is that it is bipolar. Therefore, the gain can be calibrated over a larger portion of the screen. Not using the dc calibration output makes it possible to place the front-end attenuator in the divide-by-20 position to reduce the effects of noise coupled into the input via the cable.

The gain for various volts/div settings is controlled through the programmable preamp—a linear transistor array. The preamp has an attenuation stage ($\div 1$, $\div 5$) and a gain stage ($\times 1$, $\times 2$, $\times 4$), which are controlled by digital inputs, as well as a continuously variable gain adjustment controlled by an analog input. This analog vernier input is used to calibrate the vertical gain for various 1-2-5 volts/div settings. By changing the offset input between ± 3 times the volts/div setting, the calibration routine can adjust the vernier control voltage until the two voltages show up onscreen exactly 6 divisions apart. The same procedure is followed to calibrate the full-down-vernier position of the preamp—that is, the preamp gain that results in the next higher 1-2-5 volts/div setting but with the discrete attenuation and gain stages in the same state. Given these two calibration factors (DAC values), the firmware can generate vernier control voltages for any volts/div setting between the nominal 1-2-5 gains with the formula:

$$VCV = VCV_{fdv} + (\text{Gain Ratio}) (VCV_{nv} - VCV_{fdv})$$

where VCV = vernier control voltage (DAC value)
 VCV_{fdv} = vernier control voltage at full-down-vernier
 VCV_{nv} = vernier control voltage at no vernier

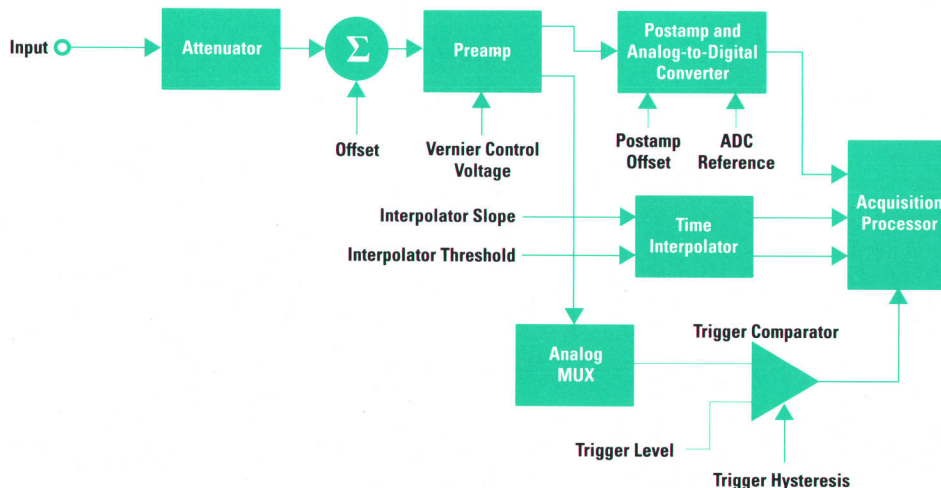


Fig. 1. DAC-generated control voltages in the HP 54601A digitizing oscilloscope.

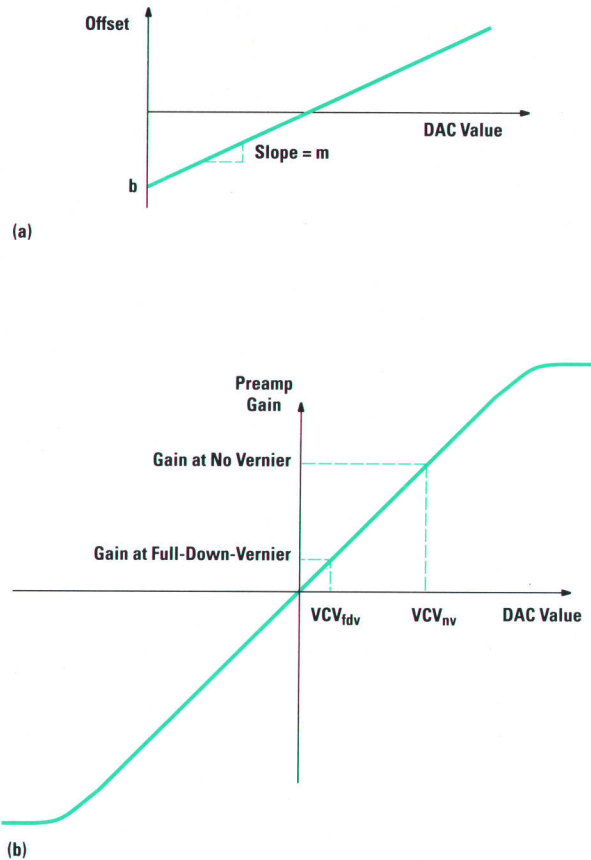


Fig. 2. (a) Offset voltages are generated by the DAC from an equation of the form $y = mx + b$. (b) Preamp gain as a function of DAC value.

and

$$\text{Gain Ratio} = \frac{\frac{\text{no-vernier volts/div}}{\text{desired volts/div}} - \frac{\text{no-vernier volts/div}}{\text{full-down-vernier volts/div}}}{1 - \frac{\text{no-vernier volts/div}}{\text{full-down-vernier volts/div}}}$$

These formulas assume the gain is linear between the no-vernier and full-down-vernier gains of the preamp. In reality, there is some nonlinearity in this curve (Fig. 2b). The gain accuracy specification is degraded in this region ($\approx 3.0\%$).

Calibrating the entire vertical section (including the ADC references, postamp offsets, trigger levels for each channel, and trigger hysteresis) takes less than 3.5 minutes in our manufacturing process. No external source is necessary and minimal operator intervention is needed. The user can calibrate the HP 54601A by first connecting a BNC cable between the dc calibration output BNC connector on the back of the oscilloscope and the channel 3 input. Then, when the oscilloscope prompts, the user moves the cable to the channel 1 input, channel 4 input, and finally to the channel 2 input.

Another calibration using the multichannel DAC is performed every 40 milliseconds during normal oscilloscope operation. This is the time-interpolator calibration. The time interpolator measures the time between the trigger event and the sample clock with 100-ps resolution by stretching the interval and counting cycles of a high-frequency clock. In this calibration, a known pulse width is input to the time interpolator by the acquisition processor. The firmware adjusts the time interpolator's slope and threshold until the pulse is stretched and measured at exactly the correct width. Calibrating this circuit (instead of just characterizing it) eliminates the need for a look-up table or equation that tells the acquisition processor where to place the data, thereby increasing throughput. This calibration takes approximately 150 μ s.

The HP 54600A uses the same DAC IC and algorithms but uses one DAC channel fewer. The channel 3 offset channel is unused, and the channel 4 offset is used as the trigger level for the external trigger source.

Mark P. Schnaible
Development Engineer
Colorado Springs Division

Fig. 11 gives the filter output ripple voltage frequencies (critical frequencies) for the case when the LSB D0 only is high (and every 256 counts above the LSB only high) and when the MSB D15 only is high. These ripple voltage frequencies will vary with the frequency of the DAC refresh clock. These output frequencies are the fundamental frequencies of a rectangular pulse string. The ripple frequency that occurs when the LSB D0 only is high is given by:

$$f_{\text{LSB}} = \frac{\text{DACCLK Frequency}}{2 \times 2^n}$$

The ripple frequency that occurs when all data bits are high is given by:

$$f_{\text{All}} = \frac{\text{DACCLK Frequency}}{2 \times 2^n} = f_{\text{LSB}}$$

The ripple frequency that occurs when the MSB D15 only is high is the 50 percent duty cycle pulse output. This ripple frequency is given by:

$$f_{\text{MSB}} = \frac{\text{DACCLK Frequency}}{2 \times 2^8} = f_{50\%}$$

To determine the output filter attenuation needed at the critical frequencies, one needs to know the critical

frequencies (already given), the peak-to-peak input voltage to the filter at each critical frequency, and the maximum peak-to-peak ripple voltage allowed at the filter output. If the input voltage to the filter were sinusoidal, then the calculation for filter attenuation would be straightforward. Since this is not the case, some additional calculation is required to determine filter attenuation at the critical frequencies.

Fig. 12 represents an input voltage $f(t)$ to the filter. The time value of voltage $f(t)$ can be expressed as:

$$f(t) = \begin{cases} 0, & -T + T_0 < t < 0 \\ 5, & 0 < t < T_0 \\ 0, & T_0 < t < T \end{cases}$$

For a 16-data-bit system, when the MSB D15 only is high, $f(t)$ will be such that

$$t_0 = T/2.$$

When the LSB D0 only is high, $f(t)$ will be such that

$$t_0 = 1 \text{ and } T = 2^{16}.$$

For each of these two cases the fundamental frequency component of this rectangular pulse string must be attenuated sufficiently to give the maximum peak-to-peak filter output ripple voltage allowed at each of these fundamental frequencies (critical frequencies). The peak-to-peak filter input voltages at these two fundamental frequencies can be determined with the use of Fourier analysis. Using the exponential form of the Fourier series,

$$f(t) = \sum_{k=-\infty}^{\infty} a_k e^{jk\omega t}$$

where

$$a_k = \frac{1}{T} \int_0^{T_0} A e^{-jk(2\pi/T)t} dt.$$

Hence

$$f(t) = \frac{5}{\pi} \sum_{k=-\infty}^{\infty} \frac{1}{k} \sin(k\pi T_0/T) e^{jk(2\pi/T)(t - T_0/2)}$$

The peak-to-peak amplitude of the k th harmonic is:

$$4|a_k| = \frac{(4)(5)}{k\pi} |\sin[(k\pi(T_0/T))]|$$

volts peak-to-peak. When the MSB D15 only is high, $A = 5V$, and $T_0 = T/2$, then the peak-to-peak fundamental frequency ($k = 1$) voltage amplitude is:

$$\begin{aligned} 4|a_k| &= \frac{(4)(5)}{\pi} |\sin[(1)(\pi)(1/2)]| \\ &= 6.366197 \text{ volts peak-to-peak.} \end{aligned}$$

Number of Data Bits	DACCLK Frequency	1SJ2 Pulse Amplitude Output	1SJ2 Ripple Frequency Output		One-LSB Filter Output Voltage for Unity-Gain Filter	Filter Attenuation for $\pm 1/2$ LSB Filter Output Ripple Voltage Amplitude (peak-to-peak)	
			Data LSB Only High	Data MSB Only High		For All Data Bits High	For MSB Only High
	Refresh Clock	Volts peak-to-peak	f_{LSB}	f_{MSB} (kHz)	Output Voltage Bit-Size	dB at f_{LSB}	dB at f_{MSB} 39.06 kHz
16	20 MHz	5V	152.6 Hz	39.06	76.29 μ V	-12.04	-98.43
15	20 MHz	5V	305.2 Hz	39.06	152.59 μ V	-12.04	-92.41
14	20 MHz	5V	610.4 Hz	39.06	305.18 μ V	-12.04	-86.39
13	20 MHz	5V	1.221 kHz	39.06	610.35 μ V	-12.04	-80.37
12	20 MHz	5V	2.441 kHz	39.06	1.2210 mV	-12.04	-74.35
11	20 MHz	5V	4.883 kHz	39.06	2.4414 mV	-12.04	-68.32
10	20 MHz	5V	9.766 kHz	39.06	4.8828 mV	-12.04	-62.30
9	20 MHz	5V	19.31 kHz	39.06	9.7656 mV	-12.04	-56.28
8	20 MHz	5V	39.06 kHz	39.06	19.531 mV	-12.04	-50.26

Fig. 11. DAC output filtering requirements.

This peak-to-peak voltage is used to calculate the filter attenuation required to keep the filter output ripple voltage that occurs at f_{MSB} to some desired level. The required filter attenuation to give a maximum peak-to-peak ripple voltage of $\pm 1/2$ LSB = V_{LSB} at f_{MSB} with a 20-MHz refresh clock is shown in Fig. 11. In general,

$$\text{Filter Attenuation at } f_{MSB} = 20 \log \frac{V_{LSB}}{6.366197 V_{pp}}$$




Fig. 12. A generic filter input voltage.

When the LSB D0 only is high, $A = 5V$, $T_0 = 1$, and $T = 2^{16}$, then the peak-to-peak fundamental frequency ($k = 1$) voltage amplitude is:

$$4|a_k| = \frac{(4)(5)}{\pi} \left| \sin \left[(1)(\pi) \left(\frac{1}{2^{16}} \right) \right] \right|$$

$$= 305.1757 \text{ microvolts peak-to-peak.}$$

This peak-to-peak voltage is used to calculate the filter attenuation required to keep the filter output ripple voltage that occurs when LSB D0 is only high (and also every 256 counts above D0 only high and when all data bits are high) to some desired level. The required filter attenuation to give a maximum peak-to-peak ripple voltage of $\pm 1/2$ LSB = V_{LSB} when all data bits are high is given in Fig. 11. In general,

$$\text{Filter Attenuation at } f_{All} = 20 \log \frac{V_{LSB}}{4|a_k|}$$

where $4|a_k|$ is calculated for all data bits high.

Comparing Analog and Digital Oscilloscopes for Troubleshooting

The analog oscilloscope has remained the troubleshooter's instrument of choice even though the digital oscilloscope has replaced it for laboratory analysis. However, the analog oscilloscope has limitations, especially in digital troubleshooting.

by **Jerald B. Murphy**

The oscilloscope is the troubleshooting tool of choice for engineers and technicians because it gives a picture of a circuit's operation. This picture, a display of instantaneous voltage as a function of time, reveals details about the circuit's operation that can't be obtained from other measuring instruments. The analog oscilloscope has remained the troubleshooter's instrument of choice, even though the digital oscilloscope has replaced it in laboratory analysis applications. Early in the definition phase of the HP 54600 Series oscilloscope project, we set out to find out the real reasons behind the remaining use of what we believe to be an obsolete technology. A series of market research tools were applied to this problem and the results of this research gave us a clear picture of what was needed to solve these very real needs.

The research tools applied to this problem were focus groups, mail surveys, telephone interviews, and personal visits. This research covered oscilloscope users in the U.S.A., Japan, and the Far East. The findings in all of these geographic areas were the same. There are very small differences in the application of oscilloscopes in these geographically diverse markets.

Analog Oscilloscopes

Analog oscilloscopes remain the troubleshooter's tool of choice because of three key characteristics attributed to this technology. The first of these is the highly interactive display, which results from the use of the direct-view vector CRT. The vector display shows changes in the signal in real time. Adjustments are simplified because there is no delay between a change and the observation of the effect on the waveform. Another characteristic of the highly interactive display is that complex waveforms, such as TV video and amplitude modulated RF, are displayed as expected.

The second attribute of the analog oscilloscope that keeps it so popular is its direct-access controls. For example, to change the size of the display, simply turn the volts/division control.

The third key attribute of the analog oscilloscope is its low price. Analog oscilloscope technology is mature and price competition between suppliers has kept prices low.

Analog oscilloscopes are not without limitations. The heart of the analog oscilloscope is the vector CRT. This is the component that gives the analog oscilloscope its highly interactive display. It is also the weak link in the analog oscilloscope's ability to keep up with the shift to digital system troubleshooting and analysis. The light produced by the phosphor has a very short lifetime after the electron beam has passed.¹ In the case of the most commonly used phosphor, P31, the light output decays to a level that is just discernible in normal room light in 38 microseconds. This means that the electron beam must reexcite the phosphor at a rate that is fast enough to keep the light output at a level that is visible. If the refresh rate drops below 26 kHz there will be a dramatic drop in light output. This may not be a problem until something like a high-resolution time interval measurement of a control or handshake line is encountered. In this case a fast sweep speed is required, and the signal repetition rate is less than the rate needed to keep the phosphor excited. The faster sweep speeds result in a lower light output from the phosphor because the electron beam does not spend enough time at any given location to cause the phosphor to produce its maximum light output. Thus, any decrease in refresh rate will have a greater effect on light output. Another factor that can reduce the light output of the CRT is the use of delayed sweep. As the magnification increases, the display repetition rate drops.² Even in the case where the main sweep display is bright and crisp, the phosphor's light output may drop below the visible level as delayed sweep is applied to increase the timing resolution of the display.

Another phenomenon that limits the usefulness of the CRT is flicker. When the refresh rate drops below a certain minimum, the display will appear to flash on and off (flicker).¹ Although the flicker frequency is a function of the specific phosphor used in the CRT, the apparent rate seen by the eye is independent of the actual rate and is approximately 15 to 20 Hz. Flicker becomes more annoying as the display gets brighter. This often occurs when a fast signal must be viewed at a slow sweep speed, as is the case in many digital circuit applications.¹




Fig. 1. Analog oscilloscope operating area.

Taking these two characteristics of the analog CRT into account, a map of useful operating area can be developed. Fig. 1 shows a plot of sweep speed as a function of input trigger repetition rate. Let's examine this map to see how the limits were established. Holding the input signal repetition rate constant at a point where a bright crisp display will be obtained and decreasing the sweep speed, we will reach a point where the display starts to flicker. Continuing to decrease the sweep speed will cause the flicker to increase to a point where the display is no longer useful. It will eventually decay to a moving bright dot. The point at which the flicker becomes so bad as to cause the display to be unusable is the left-hand limit of the operating area. Holding the sweep speed constant at a point where the light output and flicker problems are minimized and varying the signal repetition rate, we see that as the input signal repetition rate is decreased the display becomes dimmer. Once a point is reached where the display is no longer viewable in room light, the lower limit of the operating area is reached. This limit increases as the sweep speed is increased.

Another limitation of analog oscilloscopes is their inability to look ahead of the trigger point. Laboratory-quality analog oscilloscopes have delay lines in their vertical systems to ensure that the trigger point is displayed, but this does not solve the problem when the trigger event is at the end of the event to be displayed. This is a very common case in digital system analysis.

Another limitation of analog based oscilloscopes is that their measurements are based on analog ramps with accuracy specified as a percent of full scale. Often greater accuracy is required to determine the cause of problems in digital systems.

Digital Oscilloscopes

The digital oscilloscope has addressed these problems. Digital oscilloscopes make their measurements with digital precision, and they have the ability to view events that occur before the trigger. In most cases they have the ability to place the trigger event at the end of the display window, giving the user a full screen of negative-time or pretrigger information.

With their digital storage displays, digital oscilloscopes have solved the flicker and low light output problems. The present generation of digital oscilloscopes offers a bright display but it is rather static. The displays of these

oscilloscopes are more like a slide projector show than a video show. That is, they are very high-fidelity, but changes in the waveform are difficult to observe. Higher sample rates offer a partial solution to this problem, but the limiting factor is the processing time required to create the image. The present state of the art in digital oscilloscopes results in an image throughput rate of less than 100 waveforms per second. For this reason these instruments have been widely accepted in research and development labs, but they have not replaced the analog oscilloscope in troubleshooting and adjustment applications.

Effect on HP 54600 Design

The market research gave us a clear picture of the needs of oscilloscope users who were continuing to use analog oscilloscopes even when the inherent limitations of these oscilloscopes became problems. These users placed such high value on the interactive display and direct access controls of analog oscilloscopes that they were willing to live with their other limitations. The design goals of the HP 54600 project were to produce a digital oscilloscope that would have the look and feel of analog with the power of digital and at a price that was comparable to high-quality analog oscilloscopes. The HP 54600 Series oscilloscopes are designed for users who know that they need the advantages of a digital oscilloscope but are unwilling to sacrifice the real-time performance of analog oscilloscopes.

This design goal resulted in a three-processor architecture that relieves the CPU of the acquisition and display functions and places those tasks under the control of two custom integrated circuits. This new architecture provides a display throughput of more than one million points per second, producing a digital oscilloscope with display and control responsiveness equal to that of analog oscilloscopes.

Fig. 2 is a plot of sweep speed as a function of trigger repetition rate for the HP 54600. The digital display system never flickers, so the oscilloscope is fully useful at sweep speeds as slow as 5 seconds per division. Since the display is refreshed by a custom processor instead of the sweep ramp, the brightness does not decay as the sweep speed is increased or as delayed sweep is used. The only part of the map that is not covered is the high-speed single shot area. This coverage is limited by




Fig. 2. HP 54600 operating area.




Fig. 3. Analog oscilloscope delayed sweep display.

the 20-megasample-per second analog-to-digital converter system.

The logarithmic vertical axis of Fig. 2 masks the single-shot limitation of the HP 54600 analog-to-digital converter system. Single-shot captures are useful for signals with bandwidths of 2 MHz or less.

HP 54600 displays of slowly changing signals are presented free of flicker. Precision timing of control and handshake signals can be performed at sweep speeds that required a viewing hood or camera with analog oscilloscopes. Figs. 3 and 4 are analog oscilloscope and HP 54600 displays of the same delayed sweep application.




Fig. 4. HP 54600 delayed sweep display.

References


1. *Principles of Cathode-Ray Tubes, Phosphors, and High-Speed Oscillography*, Application Note 115, Publication number 5952-2044, Hewlett-Packard Co., 1970.
2. *Eliminating Time-Base Errors from Oscilloscope Measurements*, Application Note 262, Publication number 5953-3812, Hewlett-Packard Co., 1978.

Authors

February 1992

6 100-MHz Digitizing Oscilloscopes

Robert A. Witte




Bob Witte was R&D project manager for the HP 54600 Series oscilloscopes at HP's Colorado Springs Division. He received his BSEE degree from Purdue University in 1978 and his MSEE degree from Colorado State University in 1981. In 1978 he

joined HP's Loveland Instrument Division and contributed to the design of the HP 3577A network analyzer. When the Lake Stevens Instrument Division was formed he moved to Washington state and worked on the HP 3588A spectrum analyzer. After moving back to Colorado, this time to Colorado Springs, he worked on the HP 54501A digitizing oscilloscope. A member of the IEEE, he has written two books on electronic instruments and measurements and has authored many magazine articles. Bob comes from Fort Wayne, Indiana. He's married, has two children, and his hobbies are amateur radio (KB0CY), hiking, camping, and skiing.

11 Acquisition Architecture


Matthew S. Holcomb



Matt Holcomb is a hardware development engineer with HP's Colorado Springs Division. He joined the company in 1986, and has contributed to the design of the HP 54600 Series oscilloscopes, particularly the acquisition processor IC. He received his

BS degree in electrical engineering from Colorado State University in 1985 and his MS degree in electrical engineering from the California Institute of Technology in 1986.

Daniel P. Timm




A native of St. Paul, Minnesota, Dan Timm attended the University of Minnesota, graduating in 1982 with a BSEE degree. Later that year, he joined HP's Colorado Springs Division. He has done firmware design for the HP 54100A/D, 54110D,

54120T, and 54112D oscilloscopes and was responsible for the design of the digital processor and graphics subsystems for the HP 54600 Series oscilloscopes. He's currently pursuing graduate studies at the University of Colorado at Boulder. Dan's list of interests includes photography, mountain biking, home brewing, volleyball, hockey, hiking, camping, climbing fourteeners (mountains over 14,000 feet high), and telemark, cross-country, and downhill skiing.

21 Built-in Test System


Stuart O. Hall



Now an R&D software engineer at HP's Colorado Springs Division, Stuart Hall was formerly a test development engineer in manufacturing engineering, and was responsible for developing the test strategy for the HP 54600 oscilloscopes and the

hardware and software to implement it. Previously, he had developed test systems for the HP 54502A and 54504A oscilloscopes. A member of the IEEE, he received his BSEE degree from Texas A&M University in 1987 and joined HP in 1988. Stu is a native of Bryan, Texas. He's married to another Colorado Springs Division engineer, and he enjoys USVBA volleyball, skiing, and sports cars.


Jay A. Alexander



Currently pursuing an MS degree at the University of Colorado at Boulder, Jay Alexander was test development manager at HP's Colorado Springs Division, responsible for test engineering for all new products, including the HP 54600 Series oscilloscopes. He joined HP in 1986 after receiving his BSEE degree from Northwestern University. He has served as a manufacturing engineer for the HP 54201A/D oscilloscopes and as a test engineer for the HP 16530/31A, 16515/16A, and 16540/41A logic analyzers. He has also done robotic workcell development for IBM Corporation. A member of the IEEE, he has held several local-section offices and is a registered professional engineer in Colorado. Jay was born in Wausau, Wisconsin. He is married, does volunteer work with the disabled, and enjoys snowboarding, boardsailing, and hiking.

29 Frequency Response & Effective Bits

Martin B. Grove



Marty Grove is a manufacturing engineer at HP's Colorado Springs Division. He did the new-product manufacturing engineering for the HP 54600 Series oscilloscopes. A native of White Salmon, Washington, he received his BSEE degree from the University of Portland in 1986 and first came to HP as a student in 1985 at the Vancouver Division. At Colorado Springs, he has also served as manufacturing support engineer for HP 54100 Series digital oscilloscopes, and in 1991 he received his MSEE degree from the University of Colorado, specializing in digital signal processing. Marty is married and has two children. He's a glider pilot and enjoys hiking and tennis.


36 Oscilloscope Mechanical Design

Robin P. Yergenson



Now an R&D development engineer with HP's Boise Printer Division, Rob Yergenson joined HP's Colorado Springs Division in 1984 after receiving his BS degree in mechanical engineering from Virginia Polytechnic Institute and State University. At Colorado Springs, he served as a development engineer for the HP 54111D oscilloscope, the HP 10400 Series probes, and the HP 54601A oscilloscope. He was responsible for molded part design for the HP 54601A, and his earlier probe work resulted in a U.S. patent on an IC grabber design. Rob was born in Emmett, Idaho and served in the U.S. Navy for four years. He's married, has three children, and enjoys his family, woodworking, and studying the aerodynamics of bird flight.


Timothy A. Figge



Tim Figge is an R&D hardware designer with HP's Colorado Springs Division. He received his BS degree in mechanical engineering from the Illinois Institute of Technology in 1985 and joined HP that same year. From 1985 to 1988 he was responsible for the printed circuit assembly area, and developed an on-line assembly documentation system. As a mechanical designer on the HP 54600 oscilloscope project, he was responsible for the elastomeric keypads, the keyboard, RFI and sheet-metal design, the modules, the display, the power supply, and mechanical testing. He is now working for his MSEE degree at the University of Colorado and is doing both mechanical and electrical hardware design. A special interest is micro-electromechanical machines. Tim was born in Minooka, Illinois. He is married and has a son. He's an avid mountain biker and runner, and also enjoys backpacking, skiing, and Roller Blading.

41 Oscilloscope EMC Design

Kenneth D. Wyatt




Ken Wyatt received a BS degree in biology from the University of California at Irvine in 1973 and a BSEE degree from California State University at Long Beach in 1977. Following school, he worked for ten years in the aerospace industry as an R&D engineer, specializing in fields—from microwave design to switching power converters. During this time, he developed and sold a popular RF circuit analysis program. With HP since 1987, he has been leading the EMC design and research program for the Colorado Springs Division, where he is currently the product regulations manager. A prolific author, he has

written several magazine articles and technical papers on various aspects of RF and EMC design. Ken is married, has two children, and lives in Woodland Park, Colorado. His outside interests include ham radio (WA6TTY), computing, and mineral collecting.

45 Digital Oscilloscope Persistence


James A. Kahkoska



James Kahkoska was responsible for the system and module software of the HP 54600 Series oscilloscopes. He joined HP's Colorado Springs Division in 1979. After ten years in software quality assurance, he moved to R&D in 1989. His work on the HP 54600 Series has led to two patent applications. Born in Weisbaden, Germany, James is married, has twin girls, and enjoys running, cycling, and skiing. He holds a BS degree in computer science from the University of Colorado at Colorado Springs.

48 Multichannel ADC


Grosvenor H. Garnett



Grove Garnett developed the multichannel DAC for HP digitizing oscilloscopes and is named as an inventor on two patents resulting from that work. Born in San Antonio, Texas, he received his BSEE degree from The University of Texas at Austin in 1961 and his MSEE degree from Southern Methodist University in 1967. In 1969, after nine years in telecommunications R&D with Collins Radio, he joined HP's Colorado Springs Division. Besides the multichannel DAC, he did the hardware design for the HP 16532A oscilloscope and has contributed to the design of the HP 54501/02/03A oscilloscopes, HP 1300 Series displays, and HP 1700 Series oscilloscopes. Grove serves on the Department of Electrical and Computer Engineering Visiting Committee of The University of Texas. His public service credits include 4½ years in the U.S. Navy and periods with the local sheriff's posse and volunteer fire department. His interests include reading, observing nature, and ways of motivating and inspiring young people to learn. He's married and has three children.

57 Oscilloscopes for Troubleshooting

Jerald B. Murphy




Jerry Murphy received his BSEE degree from the University of Texas at Arlington in 1964 with an emphasis on microwave circuits and devices. After five years in the aerospace industry as an avionics engineer, he joined HP as a field engineer in

1969. Since 1972 he has been a member of the marketing staff of the Colorado Springs Division and has worked in nearly every functional area of marketing. Most recently, as lead marketing engineer for the HP 54600 Series oscilloscopes, he was responsible for the market research that led to the definition of these products and for the marketing plan. He has authored several papers on oscilloscope applications and time-domain reflectometry. Jerry's hobbies are photography, hiking and camping in the Colorado Mountains, and free-flight model airplanes—he is the current U.S. national champion in the F1A glider class. He also serves as an assistant scoutmaster for a local Boy Scout troop. A native of Tyler, Texas, he has two grown children.

62 Neural Networks


John McShane



John McShane received his BA degree in psychology from Trinity College Dublin in 1974 and his PhD degree in psychology from Cambridge University in 1977. From 1977 to 1980 he was a university lecturer at St. Andrews University and from 1981 to 1988 he held a similar post at London University. From 1989 to 1991 he was a member of the technical staff of HP Laboratories in Bristol, England, where he worked on an expert configuration support system and participated in a neural networks technology watch program. He is now a senior lecturer in cognitive science at Hatfield Polytechnic University. John has authored two books and is currently coauthoring a third. He is the author or coauthor of numerous papers in scientific journals and has edited books in the area of cognitive science.

66 Distributed LAN Analysis

William W. Crandall



A software engineer at HP's Intelligent Networks Operation, Bill Crandall helped design and implement new features in HP ProbeView and HP LanProbe, including node traffic, autopolling, and protocol analysis. Previously, with the HP Information Networks Division, he developed an Asian-language version of an IBM 3270 terminal emulator and an IBM RJE emulator for HP 9000 computers. A native of San Francisco, Bill received his AB degree in public and international affairs from Princeton University in 1987 and joined HP the same year. He is a member of the ACM and is interested in software quality and productivity in addition to networking. He works as a volunteer on open space issues and political campaigns, and enjoys reading, hiking, skiing, and studying Japanese.

An Introduction to Neural Nets

Unlike conventional algorithms, neural net algorithms can learn the mapping between input and output. Neural nets represent information in a distributed, rather than local, way, and can have different topologies depending on the application. This paper explains these features, lists major application areas, and briefly discusses hardware and software for development.

by **John McShane**

Neural nets are a family of architectures used for processing information. Neural nets are often said to process information in the same way that the brain does. This is an overstatement. The architecture of a neural net is loosely inspired by the architecture of the brain, but there are enormous differences between a brain and a neural net. The current technology of neural nets is very far from being able to simulate even simple brain functions.

From a technological viewpoint, neural nets are of interest because they offer a computational approach that may prove to be a very effective way of solving certain problems that are difficult to solve by conventional means. Conventional problem-solving methods require an explicit representation of the mapping between input and output. Expert systems, for example, work by first making the problem-solving procedure explicit as a set of rules, then implementing these rules in a program. The rules constitute an explicitly formulated procedure for mapping from input to output. This approach works well so long as the rules that implement a procedure can be discovered. However, it is often very difficult to discover a set of rules that define a problem-solving procedure. Neural nets offer an alternative approach to an explicitly formulated input-output mapping.

Neural nets do not require that the mapping from input to output be explicitly represented in a program. For this reason they have been called "nonalgorithmic computing." This is misleading, not to say paradoxical. There are, of course, algorithms, but they are different from conventional algorithms. Conventional algorithms specify an explicit procedure for mapping between input and output. Neural net algorithms specify a learning procedure for adjusting the mapping between input and output. Effectively, the job of specifying an explicit input-output mapping in a conventional system is replaced by the job of specifying a mechanism that is capable of learning how to map between input and output. To illustrate this it is necessary to consider the structure of a neural net.

Neural Net Structure

A neural net consists of a set of processing nodes that are connected together. The nodes can be arranged in one single layer, with the same set of nodes acting to receive input and produce output. Alternatively, they can be arranged in two or more layers, with different sets of




Fig. 1. A three-layer neural net.

nodes receiving input and producing output. Fig. 1 shows a three-layer net.

Understanding how the connections between nodes work is the key to understanding neural nets. Each node receives input, either from an external source or from other nodes in the net, or sometimes from both. The input that a node receives from any other node is a function (usually a simple product) of the output of that node and the weight of the connection between the nodes. All inputs are combined (usually by being summed). Each node has an activation value. The input to a node and its current state of activation are passed through a function f to produce an output value for the node. In Fig. 2a the activation value is shown as a threshold that must be exceeded before the node produces any output. Two typical output functions are shown in Figure 2b.

The arrangement of nodes in a net and the connections between the nodes define the architecture of a neural net. There are many different types of connectivity. In some nets every node is connected to every other node. This is particularly common in single-layer nets. In nets with more than a single layer, the connections are usually




Fig. 2. (a) The structure of a typical neural net node. Inputs are provided by $a_1 \dots a_n$. Each input has a corresponding weight w_{ij} . The threshold is θ_j . (b) Two typical output functions f .

between layers, with each node receiving input from some of the nodes in the previous layer.

Neural nets learn how to create a mapping from input to output by adjusting the weights on the connections between nodes until the output of the net meets some criterion.

The major features of neural nets are:

- Neural nets can learn the mapping between input and output, rather than needing to have it specified in advance.
- Each node in a net receives inputs from several other nodes, which are combined together to determine the output of the node.
- There are different net topologies (or architectures), which typically have different applications.

Distributed Representation

Given some entities to be represented, the most straightforward scheme is a local representation that uses one computing element per entity. Conventional systems largely use local representation. Neural nets do not, except at the lowest level at which each node responds to some particular input feature. Above this level, each entity is represented by a pattern of activity distributed over many computing elements (nodes), and each computing element is involved in representing many different entities.

A representation of an entity is built up over time in a neural net by presenting exemplars of the entity to the net and adjusting the strength of the connection between nodes that respond to the exemplar (i.e., nodes that respond to a feature contained in the exemplar). Fig. 3 is an example of a two-layer neural net that receives input from a source external to the net and provides output from the net. Each node in the input layer responds to the presence of a particular feature of the input. Typically this will be some very low-level feature (such as a line at a particular orientation in a net that is designed to recognize letters). Thus, the input of any entity will be distributed over several nodes. If the feature to which a node responds is present in the input, the node is activated; otherwise it remains inactive. Each input node is connected to various output nodes and activates the

output nodes as a function of its own activation and the weight on the connection between it and the output node. Each output node sums all the inputs that it receives and produces an output if the sum of the inputs exceeds a threshold value.

Learning in a neural net occurs by adjusting the weights between nodes according to principles to be discussed below. When a net has learned to represent an entity, presentation of the entity results in a particular pattern of activation over a set of nodes. This type of representation has advantages and disadvantages. One of its greatest advantages lies in the fact that no single input feature is critical to the pattern of activation. This means that the whole pattern of representation can be recalled given only partial input about the entity. This is of particular value in pattern recognition tasks.

Learning

Learning in a neural net occurs by adjusting the weights between nodes. In practice, the weights are assigned some initial value, which may be random or may be determined by some knowledge of what the approximate weights are likely to be for a particular representation. The net is then given some training input and an output is produced. There are two types of learning: supervised and unsupervised.

In supervised learning, there is an external criterion for what constitutes a correct output and the net is provided with feedback about how close its output is to the criterion. The difference between the desired output and that produced is computed. The weights are then adjusted according to some algorithm, which is designed to reduce the error, typically by a small amount only, between the desired output and the output produced. The net then receives new input and the process is repeated. Over a large number of training trials, the net learns to produce stable output from the input that it receives. A number of theorems about convergence and stability have been proven for a variety of learning algorithms.

In unsupervised learning, there is no external criterion of performance. Instead, the learning algorithm is designed




Fig. 3. A two-layer neural net. Input is provided from outside the net to $a_1 \dots a_n$. Output from the net is provided by $b_1 \dots b_p$. The weight on the connection between a_i and b_j is w_{ij} .

to search for patterns in the input data. Many unsupervised learning algorithms are essentially parallel implementations of clustering algorithms. Unsupervised learning is useful for exploring patterns that may exist in complex data, but this method of learning is not usually the basis of applications.

In summary:

- In supervised learning, neural nets learn by comparing the output produced with some desired output and then adjusting the weights on connections to reduce the difference between the output produced and that desired.
- The method of reducing the difference is specified by the learning algorithm.

Net Architectures

The learning algorithm and the net architecture are the two major features of a neural net. The net architecture is the arrangement and interconnections of the nodes. The nodes in a neural net are arranged in layers. There can be a single layer that serves as both input and output, there can be two layers, one for input and one for output, or there can be three or more layers, in which case there are layers that do not have direct connections to either input or output. These layers are said to be hidden. Their main function is to create intermediate representations of the input.

Single-Layer Nets. The best known single-layer net is the Hopfield net. In 1982, John Hopfield¹ published a seminal paper in which he showed how to analyze the global dynamics of a fully connected, single-layer net, using techniques borrowed from spin-glass theory. This introduced a new set of mathematical techniques to the area of neural nets, and had a revolutionary impact on a field that had attracted relatively little research interest since Minsky and Papert² published their critique of perceptrons (see below) in 1969.

A Hopfield net initially stores a number of patterns using a simple learning rule that increases the weight of the connection between units that are simultaneously active when an input pattern is presented. Hopfield showed that the dynamics of such a system could be modeled by Lyapunov energy functions³ and that as the system evolves the energy dissipates so that the net converges to a local minimum. This represents a stable point of storage for the pattern. Once the patterns have been stored the weights are frozen. Hopfield showed that these patterns are stable and can be retrieved given fragments of the pattern as input. The net can thus act as a content-addressable memory.

Hopfield nets have also been used to solve optimization problems.⁴ In these cases, it is first necessary to create an energy function that describes a specific combinatorial optimization problem, and then construct a net from the energy function. Hopfield and Tank⁵ have obtained almost optimal solutions to the traveling salesman problem using this technique.

Two-Layer Nets. Two-layer nets have a direct connection between the input and the output layers. Historically, this architecture is of considerable importance because it was used by Frank Rosenblatt⁶ in the 1950s and 1960s to

investigate the ability of neural nets to learn to recognize patterns. Rosenblatt called his nets perceptrons. Rosenblatt's learning algorithm for a two-layer net* adjusts the weights between the input and output layers in proportion to the error between the desired and computed outputs. Rosenblatt attempted to extend the same learning procedure to three-layer nets, but failed to find a satisfactory way of adjusting the weights in the hidden layer of units. This failure, together with the fact that Minsky and Papert showed in 1969 that there are fundamental limitations on the kinds of patterns that a two-layer net can learn to recognize, was largely responsible for a loss of interest in neural nets as computational systems between the late 1960s and the early 1980s.

Multilayer Nets. In the mid-1980s the problem of adjusting the weights in a hidden layer of units, which had defeated Rosenblatt, was solved, apparently independently, in at least three different locations. The best-known solution is that of Rumelhart, Hinton, and Williams⁷, which is called back-propagation. The back-propagation learning technique paved the way for serious exploration of the power of nets with hidden layers. These are usually called multilayer perceptrons. A large number of applications have been developed around a multilayer perceptron architecture and some variant of back-propagation learning. There is also considerable theoretical interest in these nets.

With a multilayer net the number of hidden units and the connectivity of the net become issues of considerable importance. These are probably the areas in which the greatest theoretical advance is possible. At present there are few general principles to guide the construction of a multilayer net. Many applications simply guess at an appropriate number of hidden units, or compare a few different specifications for performance. Similarly, the degree of connectivity is often determined by ad hoc principles. These issues are of considerable importance in considering large-scale applications because the learning time required is a nonlinear function of the degree of connectivity and the number of hidden units.

In summary:

- Single-layer nets can be used to store patterns of information and can function as content-addressable memories. These nets have been extensively analyzed.
- Two-layer nets have also been extensively analyzed. Their learning capabilities and their limitations are well-understood.
- Multilayer nets have aroused great interest. They have given rise to a range of successful applications. However, their properties are not well-understood. At present, there is considerable theoretical research into these nets.

Neural Net Applications

The main application area for single-layer nets to date is in pattern recognition tasks. Because of their content-addressable memory, single-layer nets could also potentially revolutionize database applications. However, in practice

*The perceptron is also sometimes confusingly called the single-layer perceptron. The single layer in this case is a single layer of weights between the input and output layers of nodes. The growing convention is to use the term "layer" to refer to layers of nodes rather than layers of weights.

the performance of a single-layer net degrades when the number of patterns exceeds 15% of the number of nodes, which makes the system unsuitable for large-scale storage. Attempts continue to improve the storage capacity of single-layer nets.

Two-layer nets have the longest established line of applications, largely because of the work of Widrow and Hoff⁸ in the early 1960s. They developed a system called the Adaline, which is similar to Rosenblatt's perceptron, and which has been used for adaptive signal processing and control systems. For example, the Adaline is widely used as an adaptive equalizer in high-speed digital modems.

Multilayer nets have given rise to a wide variety of applications. These include:

- Identification of radar patterns
- Target recognition in sonar images
- Detection of explosive substances at airports
- Recognition of Kanji handwriting by PCs
- Parts inspection in manufacturing
- Fault diagnosis of engines
- Classification and interpretation of mass spectroscopy data
- Optical character recognition
- Insurance risk assessment
- Mortgage risk assessment.

Despite the diversity of application domains, all of these problems are examples of pattern recognition. This class of applications may be where neural nets will prove most useful. In many pattern-recognition problems it is impossible to specify precise criteria for membership of the class to be detected (because there are none). Thus, a probabilistic approach is required. However, a further difficulty is that the features used in identifying a pattern often have context-sensitive probabilities. Obtaining a precise representation of this information, as is required in conventional approaches to pattern recognition, is extremely difficult. The ability of a neural net to learn an input-output mapping may prove a more practical and more robust approach. A good example of this approach is MSnet, which was built by Bo Curry⁹ at the Stanford Science Center, and which classifies mass spectra more reliably than other methods.

Multilayer nets are also being extensively explored for various aspects of speech processing. The debates in this area, in particular, are vigorous between the proponents and critics of neural net approaches.^{10,11} The nature of the debate usually revolves around what it might mean to have learned some aspect of a language. Many neural net demonstrations have been criticized as restricted to local problems of language while missing the principles that give language its power. These criticisms have some force when applied to the demonstrations to date, but many of the proponents of neural nets believe that neural nets will be built that can learn to process language significantly better than other systems.

The major problem with using neural nets for applications is setting up the architecture of the net in the first place. In most existing applications the architecture of a neural net has been selected in advance of using the net. At

present the choice of an architecture is to some extent a matter either of inspired guesswork or of trial and error. There is relatively little guidance in the neural net literature on how to match an architecture—how many layers and what degree of connectivity—to a problem. However, recent developments indicate that it may be possible to allow a net to modify its own architecture during the course of learning.

Hardware and Software

Neural nets employ the principle of parallel computation by a layer of nodes (and sometimes of several layers with bidirectional connections simultaneously activating each other). To date, most of this parallelism has been simulated on serial computers. However, it is now possible to obtain specialized hardware for neural net applications. There are special-purpose accelerators to speed up run time on serial computers, implementations of particular networks on VLSI chips, and neurocomputers containing arrays of parallel processors.

There is an increasing amount of neural net software available for rapid prototyping of neural net applications. A typical software environment will contain a selection of architectures and learning rules. The application developer need only select the architecture desired and provide the input for the net and the exemplars on which the net is to be trained.

References

1. J.J. Hopfield, "Neural Networks and Physical Systems with Emergent Collective Computational Abilities," *Proceedings of the National Academy of Sciences*, Vol. 79, 1982, pp. 2554-2558.
2. M.A. Minsky, and S. Papert, *Perceptrons*, MIT Press, 1969.
3. M.W. Hirsch and S. Smale, *Differential Equations, Dynamical Systems, and Linear Algebra*, Academic Press, 1974.
4. J.B. Shackelford, "Neural Data Structures: Programming with Neurons," *Hewlett-Packard Journal*, Vol. 40, no. 3, June 1989, pp. 69-78.
5. J.J. Hopfield and D.W. Tank, "Neural Computation and Constraint Satisfaction and the Travelling Salesman," *Biological Cybernetics*, Vol. 55, 1985, pp. 141-152.
6. F. Rosenblatt, "The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain," *Psychological Review*, Vol. 65, 1958, pp. 368-408.
7. D.E. Rumelhart, G.E. Hinton, and R.J. Williams, "Learning Internal Representations by Error Propagation," in D.E. Rumelhart and J.L. McClelland, *Parallel Distributed Processing: Volume 1*, MIT Press, 1986.
8. G. Widrow and M.E. Hoff, "Adaptive Switching Circuits," *IRE WESCON Convention Record*, 1960, pp. 96-104.
9. B. Curry and D.E. Rumelhart, *MSnet: A Neural Network that Classifies Mass Spectra*, HP Laboratories Technical Report, 1990.
10. D.E. Rumelhart and J. McClelland, "On Learning the Past Tense of English Verbs," in J.L. McClelland and D.E. Rumelhart, *Parallel Distributed Processing: Volume 2*, MIT Press, 1986.
11. S. Pinker and A. Prince, "On Language and Connectionism: Analysis of a Parallel Distributed Processing Model of Language Acquisition," *Cognition*, Vol. 28, 1988, pp. 73-193.

Bibliography

Many of the seminal papers, including several of those mentioned above, can be found in:
J.A. Anderson and E. Rosenfeld, eds., *Neurocomputing*, MIT Press, 1988.
Two useful introductions to the field are:
J.D. Cowan and D.H. Sharp, "Neural Nets and Artificial Intelligence," *Daedalus*, Winter 1988, pp. 85-121.
R.P. Lippmann, "An Introduction to Computing with Neural Nets," *IEEE ASSP Magazine*, April 1987, pp. 4-22.

Design Challenges for Distributed LAN Analysis

The design of a distributed local area network management system is primarily a problem of data reduction, data transmission, and data presentation. HP ProbeView software and LanProbe monitors continuously monitor the health of an Ethernet or IEEE 802.3 network to allow the diagnosis of complicated problems without dispatched equipment.

by William W. Crandall

Intermittent faults that cripple Ethernet local area networks (LANs) often cannot be detected by traditional network problem-solving techniques. By the time a traditional tool like a cable tester or protocol analyzer is rolled out to monitor the network, the problem has vanished. Users suffer from unreliable LANs while network managers suffer from the wrath of angry users for problems the managers are unable to diagnose.

One solution is distributed LAN analysis. A distributed LAN analysis tool taps into a network segment and continuously monitors it, tracking use, logging important events, detecting errors, and raising alarms over major problems. To minimize this tool's interaction with the network, it is independent of other nodes on the network. Using a network management console, network managers can retrieve information gathered by the distributed LAN analysis monitors that have been placed throughout the

network. Because the monitors provide a history of network activity, the sources of transient faults can be discovered even after the network has returned to good health. The data gathered by the monitors can be analyzed to plan for the future of the network.

Designing a good distributed LAN monitoring system is difficult. With several megabytes of data flowing across a typical Ethernet segment every minute, keeping track of all of the data that appears on the network would be impossible. Thus the distributed monitor must be able to filter and store only important information and send it quickly and efficiently to the network management console. The console must be able to combine the reports of many distributed monitors into an expansive view of the network's health. Only then will the network manager be able to solve specific problems and plan for the network's future growth.

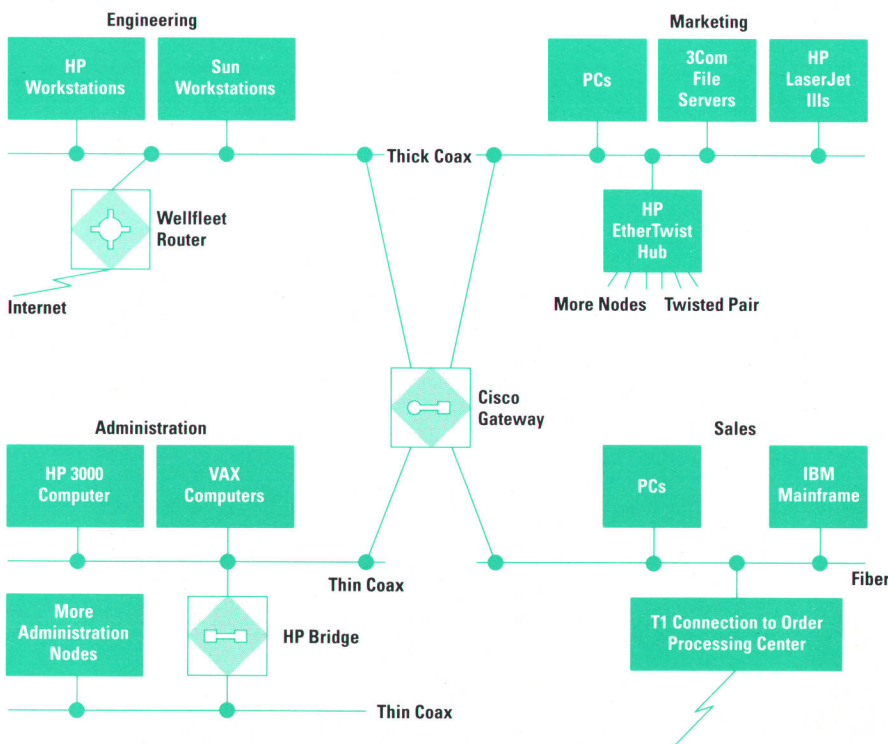


Fig. 1. Acme, Inc.'s hypothetical network.

To demonstrate the utility of distributed LAN analysis, we will describe how a typical customer of Hewlett-Packard's 4990S LanProbe distributed LAN analysis system might use the system to find intermittent network faults. We will then discuss the challenges that were faced in designing the LanProbe system.

Case Study: Acme, Inc.

We begin our study of distributed LAN analysis by looking at a typical large Ethernet-based network run by a hypothetical firm: Acme, Inc. Acme's networking problems are a conglomeration of the problems seen on the networks of HP's medium-size and large customers. As shown in Fig. 1, Acme's network is made up of a variety of cabling technologies (thick and thin coaxial cable, twisted pair, and fiber optic cable) that are tied together by different interconnect devices (repeaters, bridges, routers, and gateways). The network is spread around the world in engineering, marketing, manufacturing, and sales offices.

The Acme network is managed with tools like protocol analyzers, cable test equipment, and local diagnostic and monitoring software. Most of these tools are reactive, that is, they are not used until someone calls in with a complaint. Only then, for example, is the protocol analyzer wheeled out of the closet, connected to a network segment in the next building, and the problem found and fixed. Uncovering the source of many problems requires physically breaking the network so that the strategy of "divide and conquer" can be used to pinpoint the exact location of the physical fault or errant node.

This type of network management is costly. Users suffer the expense of frequent and unpredictable network downtime. The company pays the price of having the network management group dedicated solely to debugging the network rather than planning for its growth and improving its efficiency.

Acme decided to place an HP 4991A LanProbe on every segment in the network. The LanProbe is a passive monitor that listens to all of the traffic flowing past it on the network. It collects and analyzes this data and sends it to a central management console running the HP ProbeView software. HP 4990A ProbeView is a PC-based, Microsoft® Windows 3.0 application that provides a number of tools for examining the data that ProbeView gets from LanProbe. These tools include network and segment maps, statistics, packet trace, event log, cable test, and alert manager.

Each of these tools provides useful information. The network map is a logical picture of the layout of the Ethernet segments, repeaters, bridges, routers, and gateways that make up the network. Each segment on the network map can be examined in detail by opening the segment map. This window shows the name and type of each node attached to that segment.

The statistics tool is made up of four charts that provide different views of traffic on the network over time. The QuickView (Fig. 2), trends (Fig. 3), and packet size distribution charts show what has been happening on the segment as a whole while the node traffic chart (Fig. 4) breaks down the traffic on a node-by-node basis.

The network manager can capture and analyze raw data flowing across the network with the packet trace tool. Important events and changes on a segment are recorded in the event log tool. The network manager can test the integrity of a coaxial cable with the cable test tool; errors like cable breaks and shorts are quickly found. Finally, the alert manager lists high-priority alert messages sent to ProbeView by LanProbes scattered throughout the network.

Each of these tools is most useful when used interactively and in combination with the other tools, as we will see below.

Finally, an HP 4992A NodeLocator was added to each coaxial segment. Working together, the NodeLocator and LanProbe accurately determine the locations of nodes on coaxial cables. This feature allows ProbeView's segment map to show nodes exactly where they are physically attached to the segment. This information lets the network manager quickly find the locations of bad nodes and physical network faults like cable breaks and shorts.

Finding an Intermittent Problem

While installing LanProbe and ProbeView, the Acme network manager spent time working with the QuickView, part of the statistics tool (see Fig. 2.). QuickView is a "network dashboard" which shows the current, average, and peak values for packets sent on the network, bytes on the network, broadcasts, errors, and so on. The network manager can set thresholds for each category. For example, the manager can tell LanProbe to alert ProbeView if the number of errors on the network is more than 20 in a one-second interval. An icon will ring and flash on the PC's screen, calling the network manager's attention to the event. With threshold alerts, problems come to the manager's attention immediately, as soon as they happen; the manager does not have to go out and pin them down. This is crucial if the network problems are transient.

Soon after the Acme network manager set the thresholds, an alert appeared. Opening the alert manager, the network manager found that the error threshold on the Yellow Segment had been exceeded. The manager connected to the offending segment and looked at the trends window (another part of the statistics tool) to see network activity over the past few minutes. As Fig. 3 shows, the error rate had already dropped back to its normal level. Switching to the daily trends window, the manager saw a disturbing trend: the error level jumped sharply in even intervals throughout the day. Anticipating that the problem would soon reappear, the manager configured node traffic (yet another component of the statistics tool) to gather data for the next several hours and then went out for a long lunch.

The node traffic data identified the source of the errors. As shown in Fig. 4, node traffic provides network statistics broken down on a node-by-node basis: packets sent and received, bytes, broadcasts, multicasts, and errors sent, and utilization. The manager saw that node FFFFFFFF-FFFFFFF was generating almost all of the errors. Looking in the ProbeView log, the manager read the log entry "Bad source address FFFFFFFF-FFFFFFF seen on segment"




Fig. 2. QuickView shows the current, average, and peak values for packets on the network, bytes on the network, broadcasts, errors, and so on. The network manager can set alarm thresholds for each category.

and realized that this Ethernet address is valid only as a broadcast destination address. Suspecting that a node might have been misconfigured with the bad source address, the manager set up the trace tool to capture packets sent from address FFFFFFFF-FFFFFFF, then started the trace and went to a network planning meeting.

By the end of the meeting, trace had found many packets whose source address was the invalid address. As presented in Fig. 5, all of the bits in these packets were set to one, the CRC error-checking value was wrong, and the length of the packets was below the minimum length for an Ethernet packet (a runt). With this new information, the Acme manager now suspected a cable problem or a piece of faulty hardware connected to the net. Since the problems were occurring on a coaxial segment, the manager invoked the cable test tool. A one-shot test uncovered no problems, so the manager set the cable test to run continuously, hoping to catch a transient fault.

LanProbe's cable test is nondisruptive, unlike other TDR cable tests. The LanProbe can accurately find the location and nature of the coaxial cable problem and reports this information in the ProbeView log. The log will include a message like "Weak short 153 feet away from LanProbe" or "Cable open near LanProbe. (Connected? Missing terminator?)." When used with the NodeLocator, which accurately maps the physical locations of nodes on coaxial segments onto ProbeView's segment map, ProbeView can show exactly where the fault lies, such as "between Mike's workstation and the LaserJet III print server" (see Fig. 6). Traditional cable test techniques leave it up to the user to segment the cable manually to try to figure out where the problem is. It is a time-consuming, tedious procedure that can force network administrators to crawl into the rat's nest of wires stuffed under the floors and up the columns of modern office buildings.

LanProbe and ProbeView do away with the mess and provide quick answers.

Unfortunately, while the trends screen continued to show surges in error rates, the cable test continued to report no errors. To see if the problems experienced by the Yellow Segment were appearing on other parts of the network, the Acme manager connected to LanProbes on other segments that were attached to the Yellow Segment by a bridge. The manager used the same tools (trends, daily trends, node traffic, trace, and cable test) but did not see the same errors. Suspecting that the bridge might be corrupting packets that it was passing between segments, the manager replaced the bridge and the problem went away. The bridge manufacturer identified a faulty board as the source of the problems.

What the Scenario Shows

This scenario shows many things about distributed LAN analysis tools. The Acme network manager was able to diagnose a complicated problem easily without leaving the office. Solving the problem did not require dispatched equipment; in fact, without distributed equipment in place to monitor the health of the network continuously, the problem may not have been discovered until it had escalated into a serious, disruptive problem. The value of an early warning system is clear.

Furthermore, since ProbeView uses a graphical user interface, many tools can be placed side-by-side on the central console. Thus, different views of the network can be compared simultaneously, permitting faster problem solving as correlations become readily apparent.

Design Issues

The Acme scenario demonstrates how powerful distributed LAN analysis systems can be. It also should be clear

that they require careful design. On a typical Ethernet, several megabytes of data can flow across the wire every minute. Unlike a dispatched protocol analyzer, the user's console (ProbeView) is detached from a segment monitor (LanProbe) so the design is constrained by the speed of the link between the console and the monitor. It might seem that the best way to connect those two devices would be over the Ethernet itself. But, because a network manager will most want to connect to the segment

monitor when the network is down, the console must be able to reach the monitor via an out-of-band serial line or modem connection during those times. Thus, while the network manager may choose to connect the console to the monitor over a fast 10-Mbit/s Ethernet connection or over a slow 1200-baud modem connection, depending on the network's topology and state, the product design must always assume the worst: a 1200-baud link.




Fig. 3. (a) The trends window shows network activity over the past few minutes. Here the error rate has already dropped back to normal. (b) The daily trends window in this case shows that the error rate has been jumping sharply at even intervals.

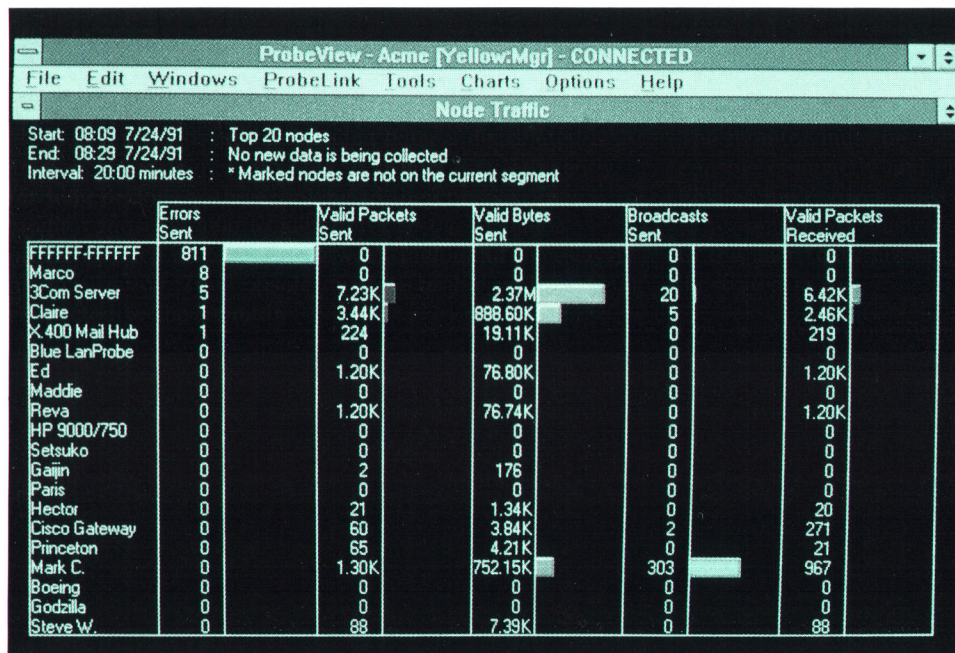


Fig. 4. The node traffic tool provides network statistics on a node-by-node basis. Here the data shows that node FFFFFFFF-FFFFFFF is generating most of the errors.

Furthermore, since the console and monitor are not always connected and exchanging data, a distributed LAN analysis monitor must be able to collect data efficiently for long periods without talking to the console. The monitor must be able to alert the console of major problems such as physical network faults and excessively high utilization. The console must be usable without being connected to the monitor.

Thus, the design of a distributed network management system is primarily a problem of data reduction, transmission, and presentation. The network monitor must be able to digest large volumes of data from the network quickly

and store only that which is important (data reduction). The monitor must be able to upload stored data to the console compactly and efficiently (data transmission). Finally, the console must be able to uncompress the data to present an expansive and useful view of the network to the user (data presentation). These design issues are illustrated in Fig. 7.

Cost is also a major design factor because the product is distributed and permanent. A network manager can afford a protocol analyzer like the HP 4972A LAN analyzer or HP 4981A network advisor because only a few are needed for even the largest of networks. A protocol

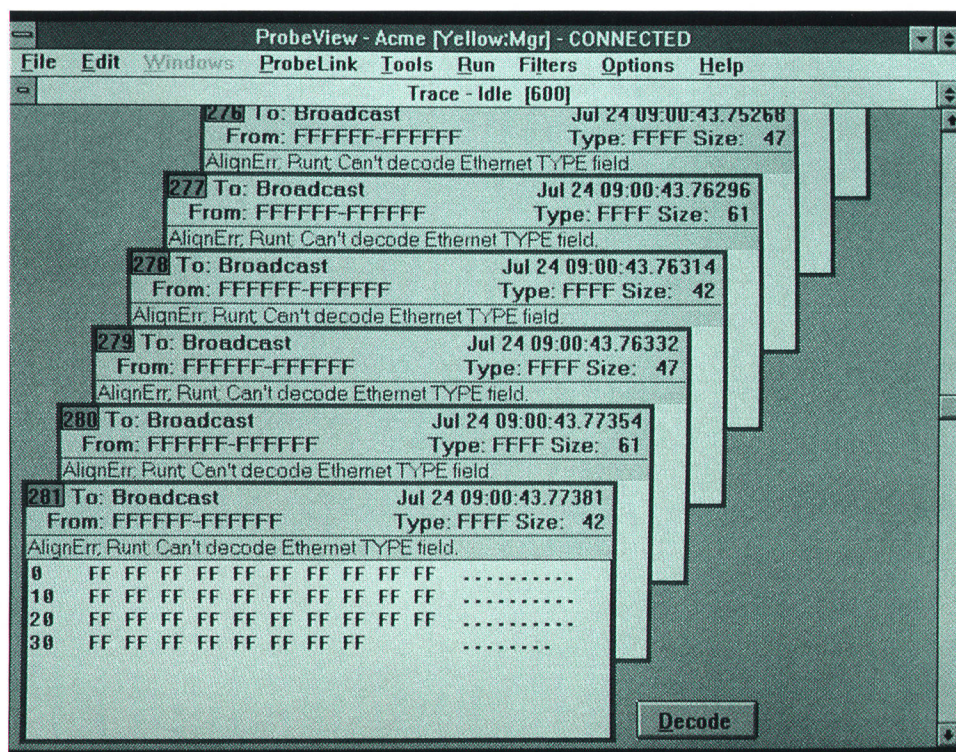


Fig. 5. The trace tool can be set up to capture packets sent from a specific node, in this case FFFFFFFF-FFFFFFF.

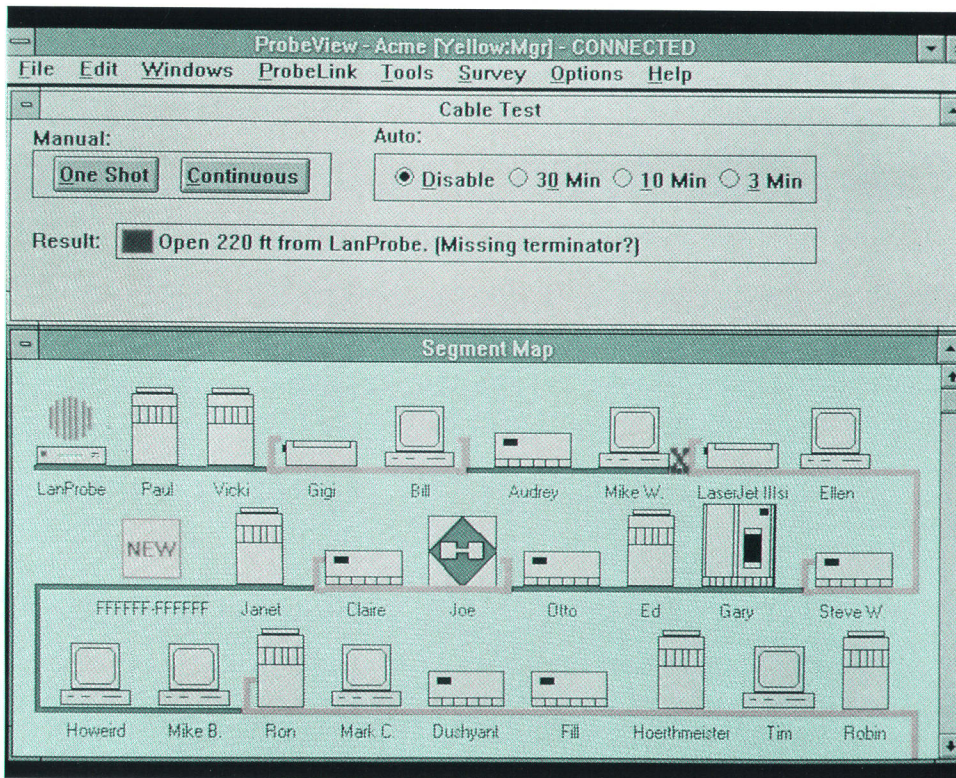


Fig. 6. The cable test tool can be used on a one-shot basis or set to run continuously to capture intermittent faults. The test is nondisruptive and shows the location and nature of a cable problem.

analyzer can be rolled to the scene of a problem, while distributed monitors are most useful when they are in place ahead of time. To keep the price down to a point where customers will be able to afford a monitor on most or all of the tens or hundreds of Ethernet segments that one might find in a medium to large network, the design must be constrained by the need to keep the cost of implementation and manufacture low.

Data Reduction in the LanProbe

Consider the problem of node table management on the LanProbe. The LanProbe runs a multitasking, dual-processor, real-time operating system. As packets are pulled off the Ethernet, they are passed to various tasks, which generally correspond to the visible tasks on ProbeView (see Fig. 8).

One of these tasks is the node table manager. This task checks the source address of each packet to see if that node's address already exists in its node table. If not, it adds it. Since Ethernet addresses are six bytes long, there are potentially about 2.8×10^{14} Ethernet addresses, many more than one could fit into a node table. Thus, the number of nodes that the LanProbe can know about must be a subset of that total. Even if the node table is very large, it still can fill up as stray or invalid source addresses are sent across the network and as network equipment is added and moved. Therefore the monitor needs a scheme that periodically thins out the node table to keep it from overflowing.

LanProbe's solution is provided by node aging. The network manager sets a time limit (from 3 minutes to 30 days). If LanProbe does not see a packet from a node within that time period, the node is marked for deletion from the node table. By setting a long time limit, the manager makes sure that outdated nodes are pruned from

the network over time. By setting a short time limit, the manager can see if important nodes like a file server have stopped transmitting.

Every 30 minutes, LanProbe also marks for deletion any "stray nodes" in the table. Stray nodes are nodes from which LanProbe has seen less than four packets. This generally means that the source address was part of a garbage or stray packet. When the table is nearly full, the marked nodes will be deleted by a background garbage collection process. If the network is extremely big, the node table will lock when full and no new nodes will be added. Balancing HP's experience with large networks with LanProbe's memory constraints, the size of LanProbe's node table was set at about 2,300 nodes.

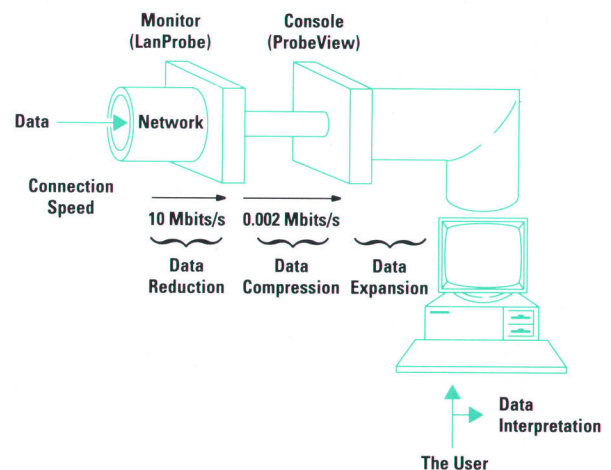


Fig. 7. Design issues for distributed LAN analysis: data reduction, compression, and expansion.

LanProbe's node table management is a good example of data reduction in a distributed network monitor. It shows that the monitor needs to know what is important and how to store it neatly. The monitor must have good strategies for dealing with information overload. Furthermore, it must be able to identify what is most and least important so that it can know what can be thrown away safely.

But perfect data reduction is impossible. For example, instead of locking the node table when it fills, it might be better on some networks to toss out some old entries to make room for new ones. The choice is subjective. As LanProbe has been installed at more customer sites and on more varied network topologies, HP has changed some of its data reduction algorithms to fit a wider class of networks.

In an early implementation of the LanProbe firmware, when the node table filled, if the nodal aging scheme did not prune enough nodes from the table, the whole table would be flushed and rebuilt based on data seen later on the network. Since Ethernet segments on well-planned networks generally contain traffic from less than 2,000 nodes, it was difficult for HP engineers to test this "fill and flush" node table management algorithm on live networks. But one early customer had a very large network that was poorly partitioned. Thus LanProbe saw packets from more than 4,000 different source addresses. When the node table filled, the table was flushed, only to be rapidly refilled and flushed again. The LanProbe wasted time building up and tearing down the node table and ProbeView was consumed trying to keep up with the volume of changes on the LanProbe.

Based on this customer's experience, new algorithms were explored. One alternative was to keep track of the

time at which each node had last sent a packet on the network. When the node table was nearly full, the LanProbe would delete nodes in the reverse order of their last transmission: the nodes that had not sent data for the longest period of time would be deleted first to make room for newer nodes. This algorithm is similar to the least recently used (LRU) algorithm used by many operating system page table managers.

But there were drawbacks to this approach. Keeping track of the time at which each node last transmitted took about 8K bytes of valuable LanProbe memory. Storing and sorting this data and purging nodes from the table was time-consuming. Furthermore, as in the "fill and flush" approach, if data from the nodes that were purged did reappear on the network, log entries indicating that a "new" node had been seen on the network would be generated. ProbeView would recognize these "new" nodes as ones that the LanProbe had seen before and would ignore them. But this meant that ProbeView and LanProbe would still be bogged down exchanging information that did not help the network manager.

Another alternative was to increase the size of the node table. While this change does not improve the "fill and flush" algorithm, it reduces the chance of the node table's filling in the first place. Unfortunately, cost constraints on the LanProbe kept this change from being implemented.

Therefore, the alternative of locking the node table when it fills was implemented. The great disadvantage to this "fill and lock" approach is that new nodes transmitting on the network will not be tracked by the LanProbe after the node table has filled. But it avoids the memory and performance problems of the other options and keeps the LanProbe from being overwhelmed by processing "new" nodes and then communicating that information to ProbeView. HP's experience is that if data from more than 2,000 nodes is seen on a segment, the network should be repartitioned, using intelligent bridges and routers to replace dumb bridges and repeaters. Given this opinion, this node table management algorithm is a good one.

Data Transmission

As stated earlier, LanProbe and ProbeView can communicate on either a fast network connection or a slow serial or modem link. Having the out-of-band serial or modem connection is vital. Without it, the network manager cannot connect to the distributed monitors when the network is down. However, having a slow-speed alternative means that the distributed LAN analysis system must be designed to work well regardless of the speed of the connection. This requirement holds true especially when large amounts of data are being uploaded from the LanProbe to ProbeView for immediate display to the user.

One tool that needs to display its information quickly is packet trace. With the packet trace tool, the network manager can tell ProbeView to have LanProbe capture all or a filtered set of the packets being sent across the network. For example, if a user calls to complain that a diskless workstation is not booting from a server, the network manager can trace all of the packets sent to or

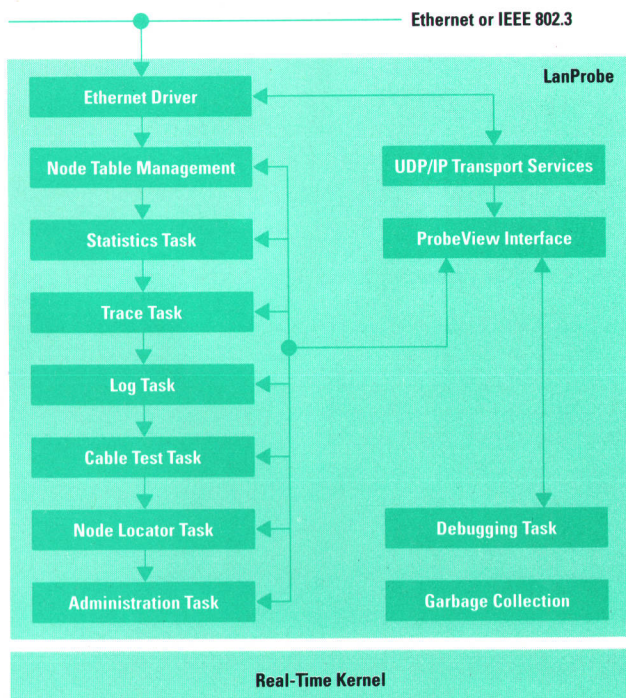


Fig. 8. HP LanProbe internals and tasks.

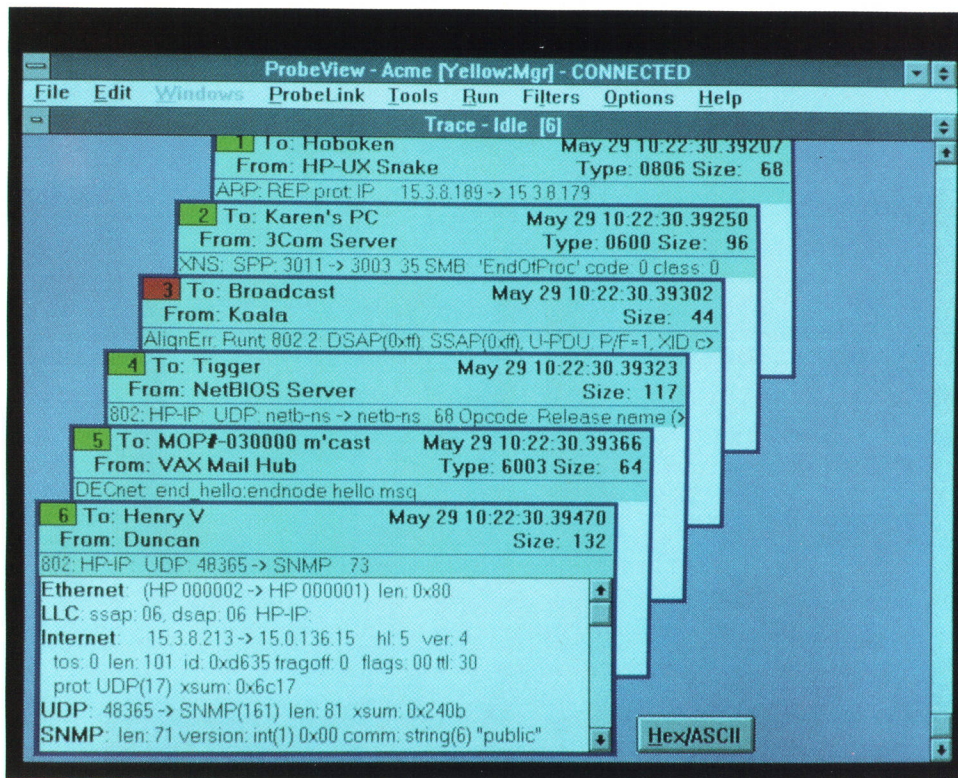


Fig. 9. The LanProbes do not sequentially upload entire captured packets to ProbeView. They first send header information, which ProbeView uses to fill in the header and the body of the notecard in the trace display.

from that workstation, decode the packets, and diagnose the problem.

LanProbe can capture up to 600 packets or 250K bytes of data. Depending on the volume of traffic on the network and the types of packets the user has chosen to capture, the LanProbe trace buffer can fill quickly (in less than a second) or slowly (over hours or days). If the buffer fills slowly, packets can be uploaded to ProbeView as they are captured. The speed of the connection between LanProbe and ProbeView is unimportant.

However, if the buffer fills quickly, the connection speed is crucial. 250K bytes of data can be sent over a typical TCP/IP Ethernet file transfer connection in less than a minute. Over a 1200-baud connection, it could take 34 minutes or longer to upload the data to ProbeView, especially if other information is being exchanged between LanProbe and ProbeView. 34 minutes is too long to wait when the network manager wants to see the packets immediately.

Thus the LanProbe system designers had to find a faster way to present the trace data to the user. The solution: as packets are captured, they are not uploaded sequentially in their entirety to ProbeView. Instead:

- The LanProbe first uploads the time the packet was seen, the total length of the packet, whether it was a valid packet or an error packet, and the first 138 bytes of the packet (which contain the Ethernet source and destination addresses and the packet type or length, depending on whether the packet is an Ethernet or IEEE 802.3 packet, and probably all of the protocol layer headers). ProbeView uses this information to fill in the header and the body of the "notecard" in the trace display (see Fig. 9).

- When the LanProbe finishes sending the header information for each captured packet, it starts sending the data in the packets. It divides each packet further into 138-byte chunks and, should there be more data in the packet to send, uploads the second 138-byte chunk of the first packet, then of the second packet, the third packet, and so on. It then goes back to the first packet and, if there is more data from the packet to upload, sends the third 138-byte chunk, and so on.
- If new packets are captured while data chunks are being uploaded, header information for the new packets are sent to ProbeView before any more data chunks are sent.

This technique works because it gives the network manager the most important information first. In trying to diagnose the diskless workstation's failure to boot, the network manager wants to know right away if the workstation is sending correctly addressed boot requests and if it is getting any responses. This information is included in the header data sent to ProbeView.

Only when those questions are answered will the network manager want to start looking through the contents of the packet. Since the most revealing data (the headers for the transport and higher layers in the network protocol stack) are at the front of the packet, most network managers will want to look at the data in the first part of each packet before looking at the data in the subsequent part of any packet. Thus uploading the data in chunks makes sense.

Note the difference in the rate at which the network manager sees the header information using a straight upload versus the LanProbe algorithm. If ten 1500-byte

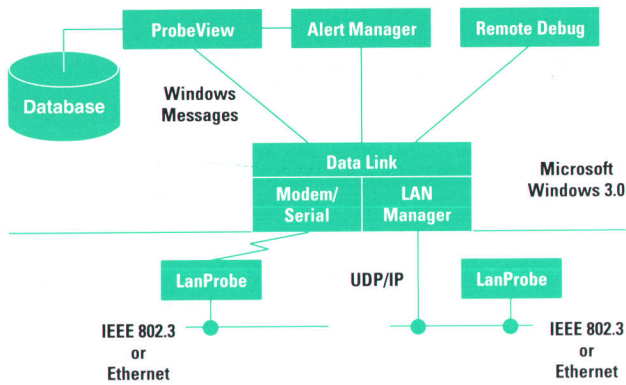


Fig. 10. HP ProbeView is a collection of Microsoft Windows 3.0 applications.

packets and then one 64-byte packet were captured and uploaded sequentially in full over a 1200-baud link, the network manager would have to wait for more than two minutes to see anything about the eleventh packet. With the LanProbe algorithm, the manager waits only 12 seconds. Furthermore, if the manager scrolls directly to a packet whose header has not been updated, ProbeView will request that packet's header immediately from the LanProbe.

But there are times that the network manager would like to see the full contents of a packet immediately, rather than waiting for the upload of the entire trace buffer to complete. Accordingly, the system is designed so that if the user is in Hex/ASCII mode and scrolls through a particular notecard in the trace window, and if all of the data that the user wants to see in that packet has not been fully received yet, ProbeView will request the remaining data from the LanProbe. LanProbe will respond to the urgent request by uploading the data immediately.

The trace upload algorithm was implemented because the out-of-band serial or modem connection between LanProbe and ProbeView is too slow for a simpler "upload everything in the order it was received" approach to work. The design of the algorithm was dictated by the needs of network managers: information is sent in order of importance to the manager. The algorithm is flexible enough to recognize when the manager wants to see data that has not been sent yet and to send that data right away. Providing this sort of intelligence and efficiency is key to making data transmissions between distributed LAN monitors and the network management console work well.

Data Presentation

Presentation of the data gathered by the LanProbes is handled by ProbeView. ProbeView is a collection of Microsoft Windows® 3.0 applications (see Fig. 10). When the network manager starts ProbeView, the ProbeView data link is also started; this application provides reliable network transport connections over UDP/IP and reliable serial link connections. The data link runs in the background and is invisible to the user. The alert manager, which displays warning messages that are generated when

LanProbes detect serious errors, is started when the manager selects the alert manager tool in ProbeView or when ProbeView gets an alert from a LanProbe. The applications that make up ProbeView use a message-based interface to share information.

ProbeView has a modular design. Multiple applications on the PC share the common data link transport. The transport works over different media (Ethernet, serial, and modem connections) but is transparent to the higher-level applications. The transport supports a primary connection between ProbeView and one LanProbe, but it can also support several inbound alert connections from other LanProbes. Thus the manager can get alerts sent by LanProbes throughout the network even while looking at a single segment. The separation of the alert manager from ProbeView makes it easier and faster for the user to get and manage alerts. ProbeView's modular design means that it is easier to enhance, test, and support, which leads to a more reliable and useful product.

The network manager works with the ProbeView window shown on the screen. While graphical user interfaces are easy to use, most do not have ways of neatly coupling related applications that run in different windows. ProbeView solves these problems with its "tools" model. A tool is one or more windows that provide a particular feature (for example, the network map and segment map windows together form the map tool). When the manager selects a tool, ProbeView's menu changes to reflect the menu choices available for that tool. Some menu selections such as exit ProbeView, connect to LanProbe, help, and ProbeView options are common to all tools. The tools model is a convenient way of tying together different actions and displays; it saves memory on the PC and improves performance relative to having all of the tools run as stand-alone applications. It also means that several tools can be aligned on the screen so that information from many tools can be seen at once, permitting faster problem solving as correlations become immediately apparent.

All of the tools share a common database. This database is built from information supplied by the network manager and by the LanProbes to which ProbeView connects. Each time ProbeView connects to a LanProbe, the two systems identify changes in common parts of the database and synchronize those areas. The ability of ProbeView and LanProbe to synchronize quickly and accurately is an important part of good distributed network management. Without synchronization, the ProbeView tools might present incomplete or outdated information to the manager.

The ProbeView tools provide different views of common information in the database. For example, each node on a network is physically attached to only one segment. But if data sent by that node is passed to other segments and is heard by several LanProbes, that node will appear on the segment maps of several segments. (The LanProbes have no way of telling whether the node is connected to the local segment or if data from that node is being forwarded by a bridge, repeater, etc.) Thus ProbeView depends on the network manager to tell it where a node physically resides.

Microsoft is a U.S. registered trademark of Microsoft Corp.

Therefore, the network manager is allowed to place a node on a segment with the map tool. Placing a node means it is physically attached to that one segment. ProbeView automatically removes that node from the other segments on which it used to be visible. By pulling this placement information out of the database, the node traffic window of the statistics tool can display which nodes have been placed on the current segment. This information is useful for identifying if heavy users of a segment are on the current segment. If heavy users are not located on the current segment, there may be a need to reconfigure the network topology. This is all part of the data presentation done by ProbeView; from a common store of data, the ProbeView tools present multiple views and interpretations.

Distributed LAN Analysis as a Planning Tool

Traditional network management tools are not good planning tools. Most Ethernet management tools are fault management tools and help with fault detection, isolation, and recovery and perhaps some performance monitoring. General network management needs to be more comprehensive to include not only fault management but planning, configuration, security, performance, and more.

Distributed LAN analysis can fill one of these needs: planning. Because distributed network monitors provide historical data, it is easy to collect data over time and analyze past network performance to plan for the future. This planning can include capacity growth, changes in network topology, and reliability improvements such as the identification and replacement of error-prone nodes, segments, and interconnects. But as always, the distributed LAN analysis system must know which data is important and how to store it compactly. One distributed network analysis product collects data in a central database. A customer reported that this database grew by four gigabytes per month! Even if enough disk space were available to store that volume of data, there would probably be too much of it to analyze well. This experience emphasizes the importance of good data reduction at the source.

ProbeView is designed to be a planning tool, as the following example illustrates. In laying out a large network, network managers use many strategies to improve network performance and reliability. For example, related nodes are clustered together—engineering nodes are located on the engineering segment, separated from the rest of the network by a router. By partitioning the network, users get faster access to local resources like servers and printers, network backbones are used less, and errors and crashes are localized. For example, if the marketing segment goes down, users on the other segments will not be affected.

Well-designed networks work fine at first, but over time, entropy sets in. For example, a new support group might move into space vacated by engineering. The marketing

group might begin to run graphics-intensive demos on local workstations booted from engineering and support servers. Administrative functions might be consolidated so the volume of traffic and number of nodes on the administration segment increases sharply, along with wide area network (WAN) traffic. The benefits of partitioning the network are lost; local performance drops, the backbone becomes clogged, and network or node downtime on one segment will take down the whole company, not just one department.

Distributed LAN analysis can come to the rescue in such situations. The network manager already uses the distributed LAN analysis system to connect to various segments by hand to diagnose existing problems and get a current view of the network. But to plan for the future, the manager needs to build up a good store of information over time and analyze it offline.

The network manager can use ProbeView's AutoPolling feature to automate this data collection. AutoPolling can be configured to connect to each segment on the network in turn each day. After uploading new log entries and synchronizing the distributed database, ProbeView gathers and exports segment map, log, trends, daily trends, and node traffic data. This new data is appended to the end of the existing export files. These export files can be read into a database manager or into a spreadsheet like Microsoft Excel for further analysis. HP provides some Excel macros to solve some problems. For example, simple analysis can find configuration errors like duplicate IP addresses. More sophisticated, network-specific analysis can detect problems like:

- A large volume of client/server traffic running across a backbone segment (solution: partition the network to isolate clients and servers onto the same segments).
- Overburdened segments and devices like the administration segment and the WAN gateway (solution: repartition segments, add new devices, or increase their capacity).
- Overly high levels of broadcast and multicast traffic, which can significantly slow network performance (solution: change configurations or software on nodes sending too many broadcasts or multicasts).

ProbeView provides network managers with the information they need to analyze and plan for the future of their networks. It does so in a way that cannot be matched by a dispatched protocol analyzer or performance tool that provides a short-term view of a small part of the network.

Poor Network Partitioning

On a poorly partitioned network, when a node sends a packet, the packet will be seen on nearly every segment in the network, rather than being limited only to the segments necessary to get the packet from its source to its destination. This packet filtering and routing capability can be provided by segment interconnect devices such as intelligent bridges and routers. These interconnects watch the traffic on the network, figure out where nodes are located, and then filter traffic accordingly. Thus if a bridge hears node A and node B transmitting on port 1, and if it then hears a packet sent from A to B, the bridge will not send a copy of that packet on to port 2. This reduces the level of unnecessary traffic on the Ethernet segment attached to port 2.

When a network grows quickly without extensive planning or when cost is very important, low-cost repeaters and dumb bridges may be included in the network. While easy and cheap to add, these devices do not have the same sophisticated packet filtering capabilities as do their more expensive cousins. They pass on all of the data that they hear on one segment to all other segments to which they are attached.

Conclusion

Distributed LAN analysis can solve network problems. By continually monitoring an Ethernet LAN and by compiling a record of network activity, the HP LanProbe distributed LAN analysis system allows network managers to solve intermittent faults and plan network growth. This system does so with distributed analysis monitors that watch network traffic, store important information and events, and compactly transmit this data to a network management console that presents a multifaceted, expansive view of network activity. LanProbe and ProbeView allow

network managers to easily identify and fix network problems as well as analyze the network as an aid to planning its future.

LanProbe and ProbeView work because they have been designed to handle the problems of distributed LAN analysis: data reduction, transmission, and presentation. LanProbe continuously analyzes the volumes of data flowing across the network, capturing only the most vital parts of it. It has well-defined strategies for avoiding and handling information overflow. LanProbe compactly transmits its data to ProbeView using adaptive algorithms designed to work well over both fast and slow connections. ProbeView, taking advantage of its graphical user interface and common central database, gives the network manager many tools with which to solve network ailments and to plan for growth.

Acknowledgments

I would like to acknowledge and congratulate the whole team from Eon Systems and HP's Intelligent Networks Operation for their creativity and hard work in creating LanProbe and ProbeView. The design challenge for distributed LAN analysis is great but the success of the LanProbe system is testimony to how well this team met that challenge. In writing this article, I received help from many people. I would like to thank all of them, especially Mike Watters for his guidance, Joe Adler for his editorial criticism, and Gigi Chu for her historical perspective. Paul Taira took the excellent photographs that illustrate this piece.

Hewlett-Packard Company, P.O. Box 51827
Palo Alto, CA 94303-0724

ADDRESS CORRECTION REQUESTED

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

HEWLETT-PACKARD
JOURNAL

February 1992 Volume 43 • Number 1

Hewlett-Packard Company

Hewlett-Packard Company, P.O. Box 51827
Palo Alto, California, 94303-0724 U.S.A.

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

Hewlett-Packard (Canada) Ltd.
6877 Goreway Drive, Mississauga, Ontario L4V 1M8 Canada

01029627 LOAD
MR. GEORGE PONTIS
SUITE 409
24 COVE LN
REDWOOD CITY CA 94065-1528

CHANGE OF ADDRESS:

To subscribe, change your address, or delete your name from our mailing list, send your request to Hewlett-Packard Journal, P.O. Box 51827, Palo Alto, CA 94303-0724 U.S.A. Include your old address label, if any. Allow 60 days.