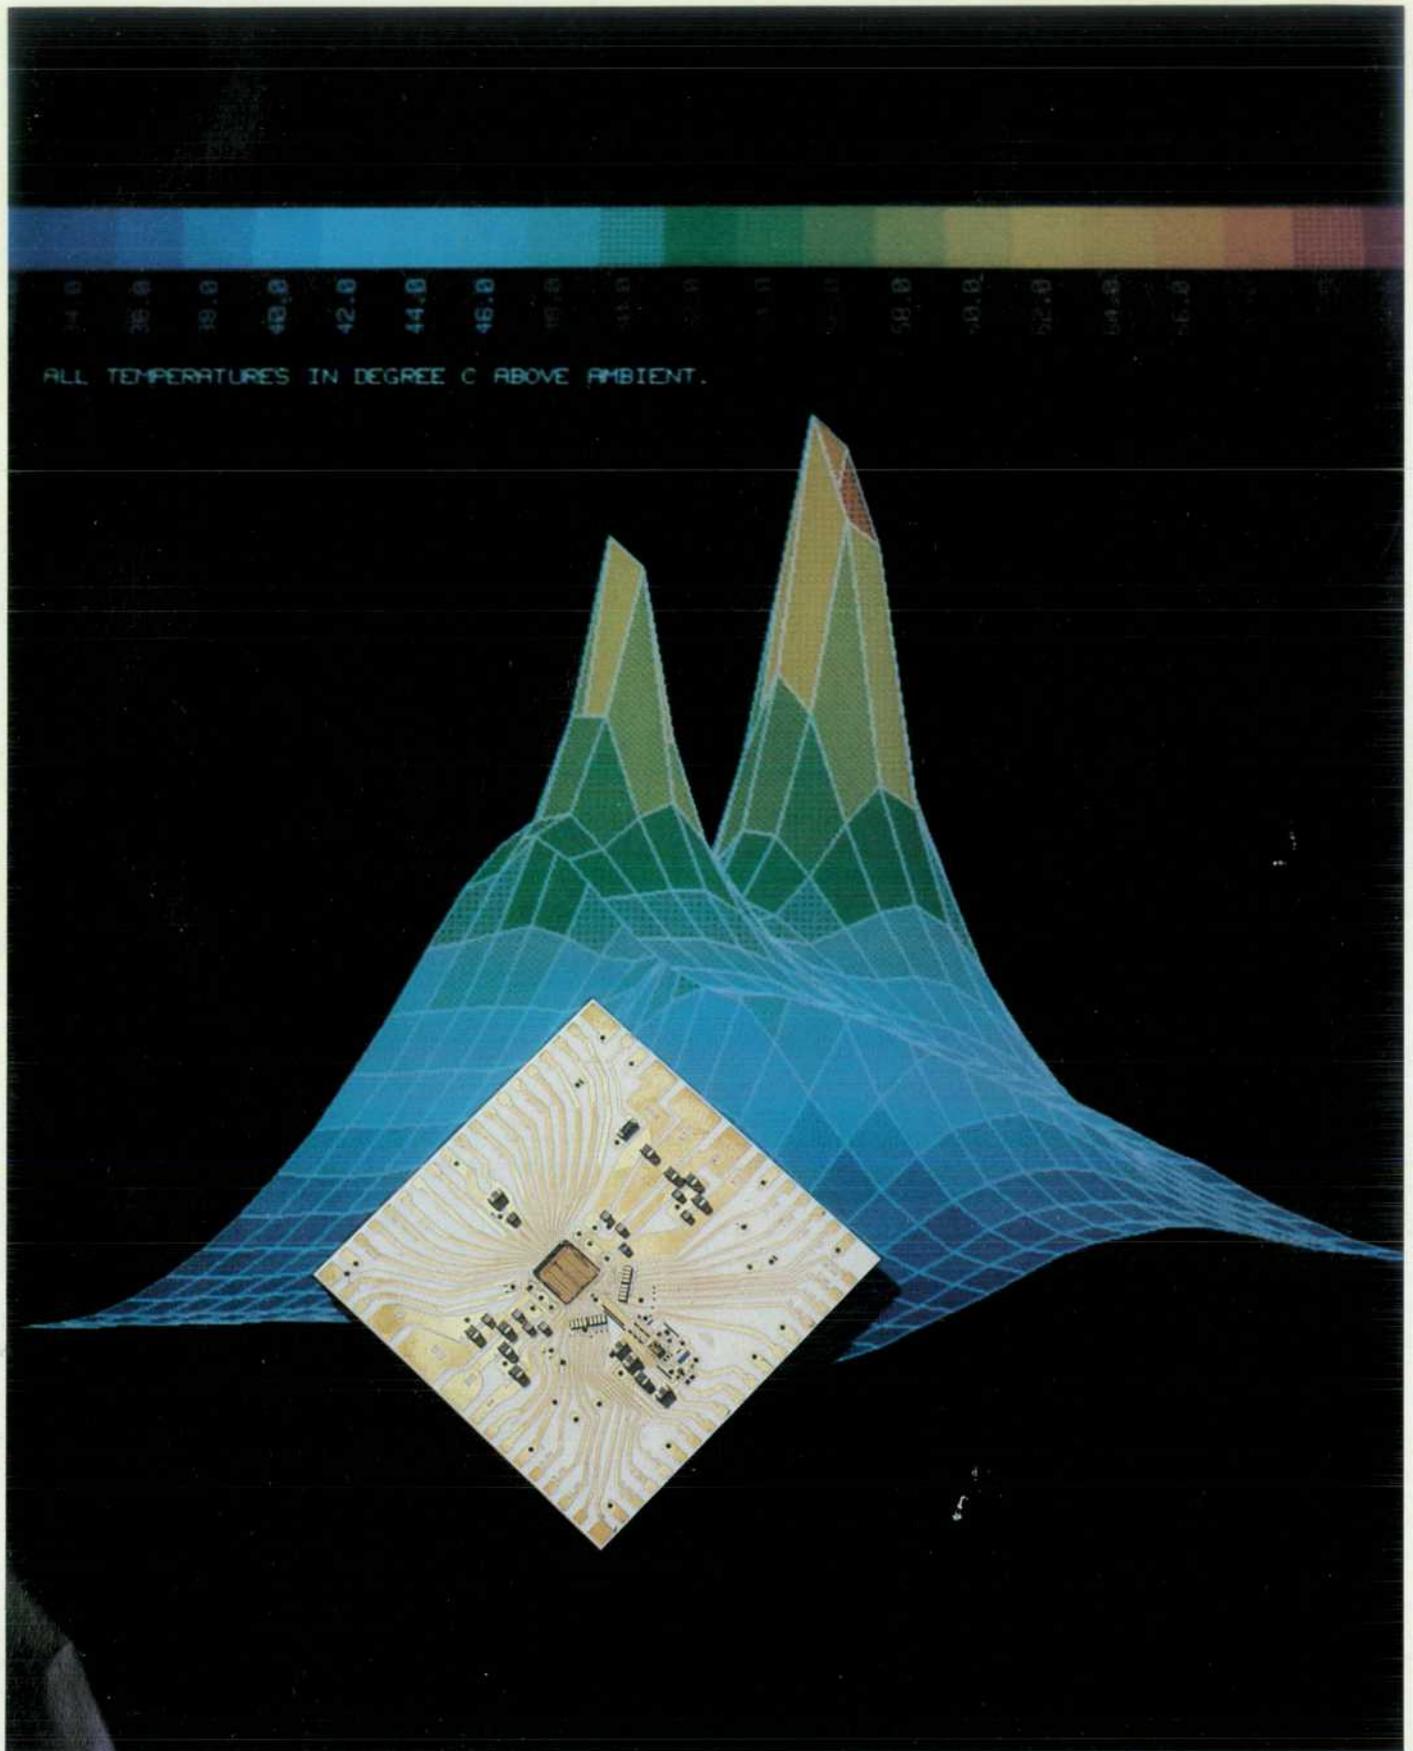


HEWLETT-PACKARD JOURNAL

FEBRUARY 1988



Articles

6 Precision Digital Oscilloscopes and Waveform Recorders, *by James L. Sorden*

12 Waveform Recorder Software Design

15 Signal Conditioning and Analog-to-Digital Conversion for a 4-MHz, 12-Bit Waveform Recorder, *by Albert Gee and Ronald W. Young*

23 Adaptive Sample Rate: A First-Generation Automatic Time Base, *by Richard W. Page and Nancy W. Nelson*

26 Waveform Reconstruction Techniques for Precision Digitizing Oscilloscopes, *by Richard W. Page and Allen S. Foster*

32 Digital Design of a High-Speed Waveform Recorder, *by Rayman W. Pon, Steven C. Bird, and Patrick D. Deane*

35 Printed Circuit Board Transmission Lines

39 Waveform Recorder Design for Dynamic Performance, *by Brian J. Frohring, Bruce E. Peetz, Mark A. Unkrich, and Steven C. Bird*

48 Fixed-Frequency Sine Wave Curve Fit

49 Packaging a High-Performance 250-Megasample-per-Second Analog-to-Digital Converter, *by Patrick D. Deane, Simcoe Walmsley, Jr., and Farid Dibachi*

53 Precision Digitizing Oscilloscope Waveform Analysis, Display, and Input/Output, *by Douglas C. Nichols*

63 Handling of Significant Digits

65 Developing a Printed Circuit Board Design System, *by Elaine C. Regelson*

68 Automating the Printed Circuit Board Design Process, *by Gary Jackoway*

70 Finding Paths in a Gridded Data Structure

71 Managing HP PCDS with the Design System Manager, *by Paul S. Reese and Mark E. Mayotte*

73 Use of Filesets in HP PCDS

75 Version Strings

77 A Multidevice Spooler for Technical Applications, *by Deborah A. Lienhart*

80 Integrating Applications in a Design Management System, *by Mark E. Mayotte*

82 HP PCDS Library Module

84 Software Quality Assurance on the HP Printed Circuit Design System Project, *by David E. Martin*

Research Report

87 Silicon-on-Insulator MOS Devices for Integrated Circuit Applications, *by Jean-Pierre Colinge*

Departments

4 In this Issue

5 What's Ahead

94 Authors

Editor, Richard P. Dolan • Associate Editor, Business Manager, Kenneth A. Shaw • Art Director, Photographer, Arvid A. Danielson • Support Supervisor, Susan E. Wright
Administrative Services, Typography, Anne S. LoPresti • European Production Supervisor, Michael Zandwijken

In this Issue



That most basic of electronic instruments, the oscilloscope, has gone digital in a big way. The *1988 Hewlett-Packard Test and Measurement Catalog* doesn't list a single traditional analog oscilloscope, but there are ten different kinds of digitizing oscilloscopes. In a digitizing scope, input signals are sampled, the samples are measured by an analog-to-digital converter (ADC), and the resulting numbers are stored in a waveform memory. This stored data can be used to create a display of the input signal that looks like a traditional oscilloscope display, or it can be manipulated digitally for various purposes. There are different digitizing scopes because there are design trade-offs. In general, if you sample faster you get more bandwidth but you can't measure as precisely. You get the most bandwidth if you deal only with repetitive input signals. Some digitizing scopes are designed for maximum bandwidth for repetitive inputs, some for maximum digitizing speed, and some for maximum fidelity and dynamic range.

In this last category are the HP 5180T/U, HP 5183T/U, and HP 5185T Precision Digitizing Oscilloscopes, described on pages 4 to 64. These scopes are waveform recorders paired with an analysis and display module. They rate the adjective "precision" because high-fidelity, fully specified dynamic performance is a distinguishing characteristic of the waveform recorders they're based on. Critical to that performance is the ADC, and two ADC designs are presented in this issue. One has a 250-megasample-per-second digitizing rate and 8-bit precision (page 39), and the other is a 4-megasample-per-second, 12-bit unit (page 15). Other waveform recorder design details can be found in the articles on pages 4, 32, and 49, and the article on page 53 explains the functions of the analysis and display module. Two unusual capabilities of these scopes are adaptive sample rate (page 23) and waveform reconstruction (page 26). The former saves waveform memory by automatically reducing the digitizing rate if the frequency content of the input signal is below a certain threshold. The latter applies sophisticated interpolation techniques when creating a display from the stored waveform data. While most digitizing scopes need ten samples to make a cycle of sine wave look right, these can do it with only two and a half samples.

What gets mounted on a printed circuit board and how it's put there has changed a lot over the years. Complexity and density have increased greatly. Manufacturing and testing are fully or partially automated. A board designer in this environment is seriously handicapped without an effective computer-aided design (CAD) system. Hewlett-Packard's offering in this area consists of two software applications for HP 9000 engineering workstations. The HP Engineering Design System (HP EDS) captures the designer's schematic diagram and simulates the operation of the logic to verify the design functionally. The HP Printed Circuit Design System (HP PCDS) provides tools for board layout and generates instructions for automatic manufacturing and testing equipment. The two applications can exchange data in either direction and they share the same library of standard components.

The story of HP PCDS is told by its designers on pages 65 to 86, beginning with an overview of the product and a description of the Design Module, which interacts with the board designer to define the board, place components on it, and generate manufacturing outputs. Algorithms for autoplacement of components and autorouting of the traces between components are discussed in the paper on page 68. Autorouting, considered the most critical element in a printed circuit board CAD system, is done by the HP PCDS Autorouter Module. The third of HP PCDS' three modules, the 8500-component, user-modifiable Library Module, is described on page 82. Quality assurance issues faced by the design team are covered in the article on page 84. Two special needs of HP PCDS and similar CAD systems were addressed by developing related software. First, because the productivity gains that a CAD system delivers can be greatly diminished by the extra management effort required to keep track of data and coordinate large project teams,

the Design System Manager was developed (see page 71). HP PCDS runs in the Design System Manager environment along with similar CAD applications. The Design System Manager provides application integration (page 80), network support, file security, and control of multiple versions of files. Second, because a CAD system has to send output to a wider-than-usual variety of devices, a special spooler was developed. The spooler is the subject of the paper on page 77.

Eleven years ago, HP produced its first silicon-on-sapphire (SOS) integrated circuit, a 16-bit microprocessor chip. SOS offered a combination of low power consumption, high speed, high circuit density, and static operation that was unmatched by any other IC technology. The hope was that the system-level advantages of SOS would offset its relatively high cost. Today, SOS is considered too costly for most commercial uses, although it is still used in some applications; for example, it's used in satellites because of its radiation hardness. On the other hand, the basic idea of building silicon integrated circuits on insulating substrates is alive and well. New silicon-on-insulator (SOI) technologies are being developed using calcium fluoride or silicon dioxide instead of sapphire (which is a form of aluminum oxide). Silicon dioxide is especially attractive because it's a common integrated circuit element and doesn't require special processing equipment. In the paper on page 87 we get an introduction to SOI technology and its advantages and learn about the SOI research being done at HP Laboratories.

-R. P. Dolan

Cover

The ADC hybrid microcircuit of the HP 5185A Waveform Recorder is pictured in front of a display from a thermal modeling program that was used to predict the heat transfer characteristics of the ADC hybrid and other microcircuits.

What's Ahead

Scheduled for the April issue are design papers on millimeter-wave components and the HP 8770A Arbitrary Waveform Synthesizer. There will also be three papers from the 1987 HP Software Engineering Productivity Conference describing a virtual user software testing tool, a system for controlling and measuring the load on the kernel of the HP-UX operating system, and process measures to improve R&D scheduling accuracy. Another paper discusses how to adapt the various logon mechanisms of AT&T's UNIX[®] operating system to the manufacturing environment.

Precision Digital Oscilloscopes and Waveform Recorders

This precision instrument family consists of five digitizing oscilloscopes based on three waveform recorders and an analysis, display, and I/O section.

by James L. Sorden

FOR MANY APPLICATIONS, the functions of an oscilloscope, voltmeter, counter, power meter, and spectrum analyzer can be provided by a single measurement instrument: the precision digitizing oscilloscope. In addition to measurements on repetitive signals (the generally assumed condition with conventional measurement instrumentation), the precision digitizing oscilloscope can make measurements on transient or single-shot signals, signals often impractical to analyze with conventional measurement tools.

To be able to function as all of the instruments mentioned, the precision digitizing oscilloscope must meet two conditions. First, the fundamental measurement device (transducer or analog-to-digital converter) must have a high degree of accuracy, fidelity, and dynamic range. Second, the analysis and input/output functions must be quick, comprehensive, and correct.

The HP 5180T/U, HP5183T/U and HP 5185T Precision

Digitizing Oscilloscopes combine analysis functions with excellent measurement fidelity to characterize single-shot or repetitive signals, either simple waveshapes or complex modulated analog signals. These instruments have two major subsections: (a) the waveform recorder with either two or four channels of data acquisition, and (b) the analysis, display, and I/O section. The HP 5180T/U (Fig. 1) uses the 20-MHz HP 5180A Waveform Recorder,¹ which has a 10-bit ADC (analog-to-digital converter) and a 16K-word memory. The HP 5183T/U (Fig. 2) uses the HP 5183A Waveform Recorder, which has dual independent 4-MHz, 12-bit ADCs and 64K words of memory (256K optional). The T versions are two-channel instruments and the U versions are four-channel instruments that use two of the two-channel waveform recorders. The HP 5185T (Fig. 3) uses the HP 5185A Waveform Recorder, which has dual independent 250-MHz, 8-bit ADCs.

The measurement strengths of these instruments,

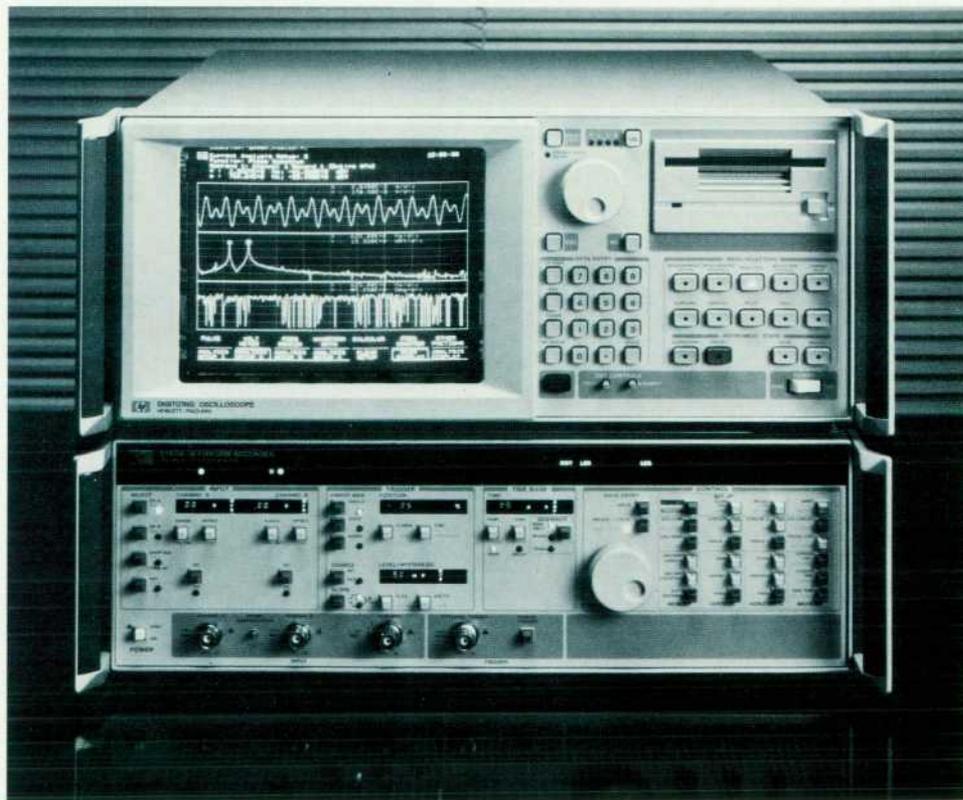


Fig. 1. The HP 5180T Precision Digitizing Oscilloscope consists of an HP 5180A Waveform Recorder and an analysis, display, and I/O section. The HP 5180A provides 20-MHz, 10-bit analog-to-digital conversion.

coupled with their built-in processing and analysis capabilities, make them well-suited for measurement systems in ATE, production, and R&D environments. In the laboratory, the digitizing oscilloscope is a stand-alone, interactive instrument. For many ATE and production systems, on the other hand, the waveform recorder can be used without an analysis, display, and I/O section to minimize rack space and eliminate front-panel clutter and confusion.

This article gives an overview of the design of the HP 5183A Waveform Recorder. The analysis, display, and I/O section that transforms the waveform recorder into digitizing oscilloscopes is described in the article on page 53. The waveform reconstruction techniques that maximize the effective bandwidths of the HP oscilloscopes are discussed in the paper on page 26. The design of the HP 5185A Waveform Recorder is the subject of the papers on pages 32, 39, and 49.

HP 5183A Waveform Recorder

The HP 5183A is a lower-cost waveform recorder. Its data acquisition, analog-to-digital conversion, and analysis techniques make many contributions to the state of the art, including adaptive sample rate, dropout trigger, and other concepts. These are detailed in the box on page 12 and the papers on pages 15 and 23.

The block diagram of the HP 5183A Waveform Recorder is typical for waveform recorders. Elements common to most recorders are:

- Input signal conditioning, including multiple sensitivity ranges and high impedance
- High-speed transducer (analog-to-digital converter)
- High-speed memory system capable of following the ADC output and later replaying data into a digital I/O

system

- Internal and external triggering to define measurement time endpoints
- Time base and oscillator system to control the ADC
- Digital system controller
- Digital HP-IB.

Shown in Fig. 4 is the signal flow functional block diagram for the HP 5183A recorder.

Input Amplifier System. The input amplifiers consist of a pair of 1-M Ω fully conditioned and programmable differential input amplifiers with switchable anti-aliasing filters. This high-impedance signal buffering makes it possible to monitor the user's system without disturbing its normal operation. Since this signal buffering system is built into the instrument, it is fully specified (this is not the case if external buffering is required). The switchable low-pass anti-aliasing filters prevent undesired aliasing of high-frequency signals and noise into the bandwidth of the recorded signal.

Analog-to-Digital Converter System. Each of the HP 5183A's two channels has its own independent 12-bit ADC with dc-to-4-MHz sample-and-hold signal conditioning. The maximum sample rate is 4 MHz. The two-pass converter uses an 8-bit flash ADC with a resultant 4-bit overlap (see article, page 15). From a system designer's point of view, a digitizer is most useful if it can be clocked at arbitrary rates and/or with nonuniform encode sequences. The HP 5183A ADC can be clocked at any rate from less than 1 Hz to 4 MHz.

Precision Digital Trigger

The second most difficult technical challenge in any waveform recorder design (after providing a high-fidelity

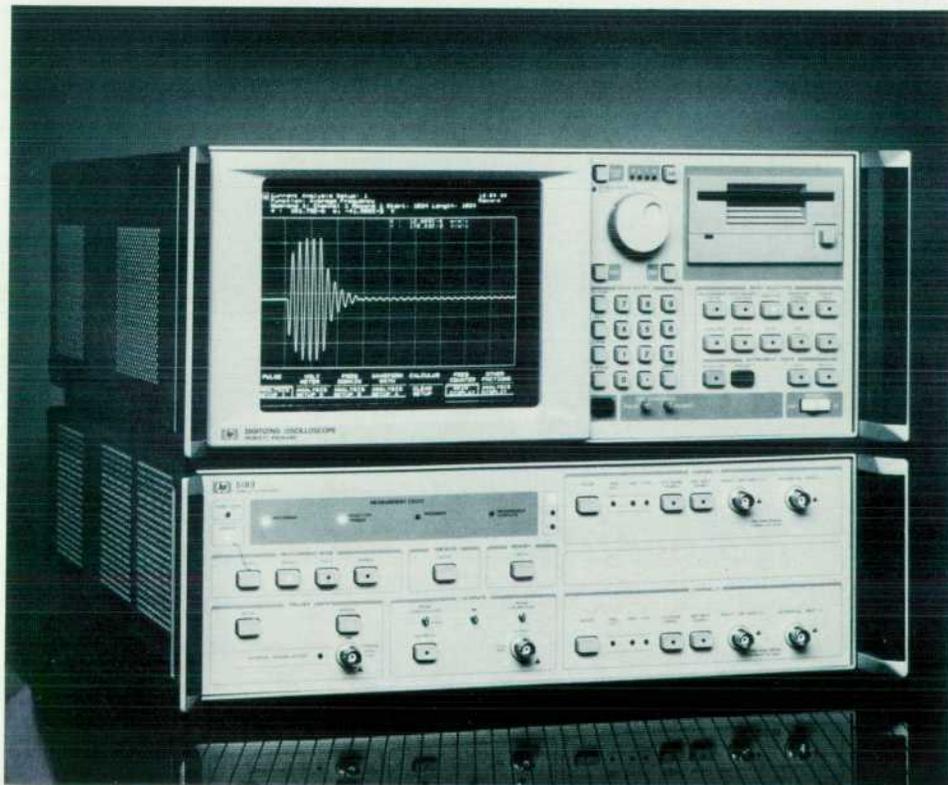


Fig. 2. The HP 5183T Precision Digitizing Oscilloscope consists of an HP 5183A Waveform Recorder and an analysis, display, and I/O section. The HP 5183A is a low-cost recorder that provides 4-MHz, 12-bit analog-to-digital conversion.

ADC) is to provide an accurate and reliable trigger system. All of the products in this family use digital triggering. Since the ADC is fully specified and tested, so is the trigger circuitry. The exact signal digitized by the ADC is seen by the precision digital trigger. There can be no drift or other inaccuracies relative to the digitized data. The trigger level control and the trigger hysteresis control are fully adjustable. This ensures optimum triggering for each application.

Dropout Trigger. A triggering technique that allows triggering in the absence of a signal has been implemented in the HP 5183 and HP 5185 systems. This feature allows triggering on and recording signal sag or total signal loss on signals that have a normally constant peak amplitude (Fig. 5). When the preset trigger amplitude conditions are not met for an interval set by the dropout delay timer, the trigger command occurs. The dropout trigger circuit consists of conventional trigger detection circuitry plus a special timer. The timer is reset with the trigger detection circuitry. If the timer circuit does not reset in the time specified by the user, then the trigger has not occurred during that time, and a dropout trigger signal is sent to the memory, activating the normal memory stop circuitry. Some examples of dropout trigger applications include:

- Ignition systems. Determine the cause of cylinder misfirings and record the time and waveshape of the signal pattern around the misfire.
- Communication systems. Determine the cause and characteristics of carrier loss or distortion.
- Power distribution system monitoring. Two classic anomalies on a power distribution system are brownouts or line sag, an application for dropout trigger (Fig. 5), and high-voltage spikes, an application for bi-trigger mode (Fig. 6).

- Oscillators. When an oscillator squeaks, a measurement can be made to determine the time and shape of the anomaly.
- Magnetic media testing. Digital tape or flexible disc systems can be precisely analyzed to see the exact magnetic position of data lost.

High-Frequency Trigger. A class of transients that conventional trigger circuits do not handle well is high-frequency noise transients that do not exceed the nominal peak-to-peak voltage level of the signal. An example of this class of noise transients is high-frequency transients on a power line system, often generated by switching spikes from inductive loads. With the adaptive sample rate system, HP 5183 Option 301, the included high-frequency digital filter can be used to detect these transients. Relatively low-energy, high-frequency transients can be detected in the presence of a much larger low-frequency signal. The detected signal is used to fire the trigger system and initiate recording of the transient (Fig. 7).

Time Base System

The time base system has the primary function of providing flexibility in the ADC sampling rate, permitting longer time records at slower sampling rates (assuming one operates within the Nyquist sampling criterion). The HP 5183 system can switch between its main rate and a second, or delayed, rate. This feature is most often used to conserve memory or increase measurement time when the operator has prior knowledge of the signal bandwidth. Either time base can set to any period from 250 ns to 4 seconds in 250-ns increments. In addition, two new functions, burst time base and adaptive sample rate, are available in the time base.

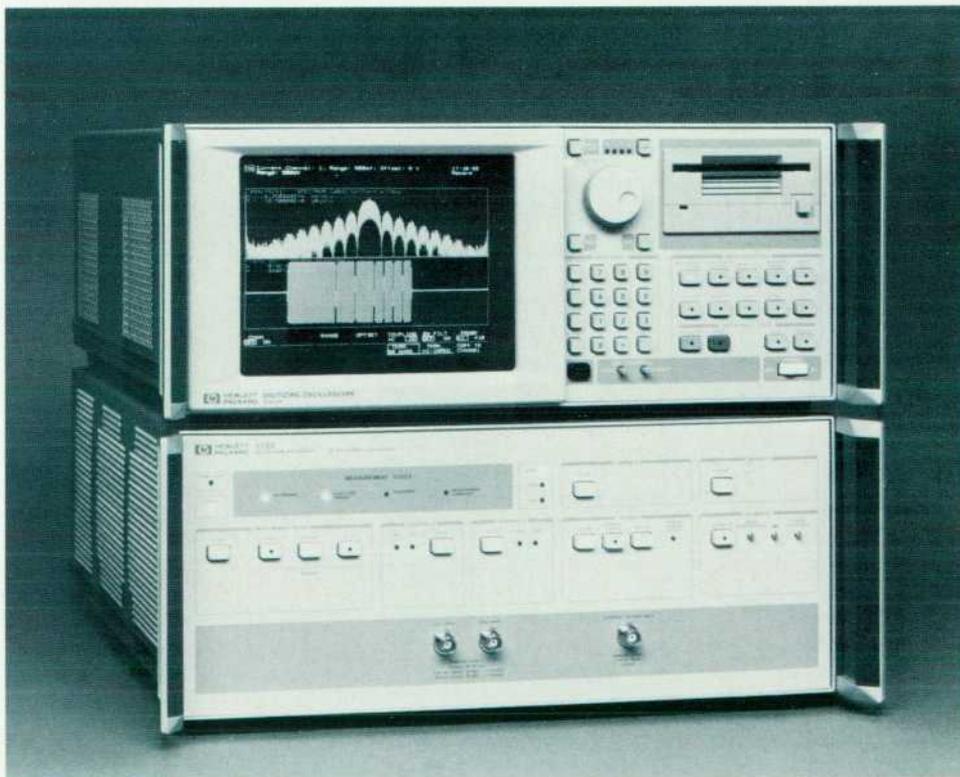


Fig. 3. The HP 5185T Precision Digitizing Oscilloscope consists of an HP 5185A Waveform Recorder and an analysis, display, and I/O section. The HP 5185A provides 250-MHz, 8-bit analog-to-digital conversion.

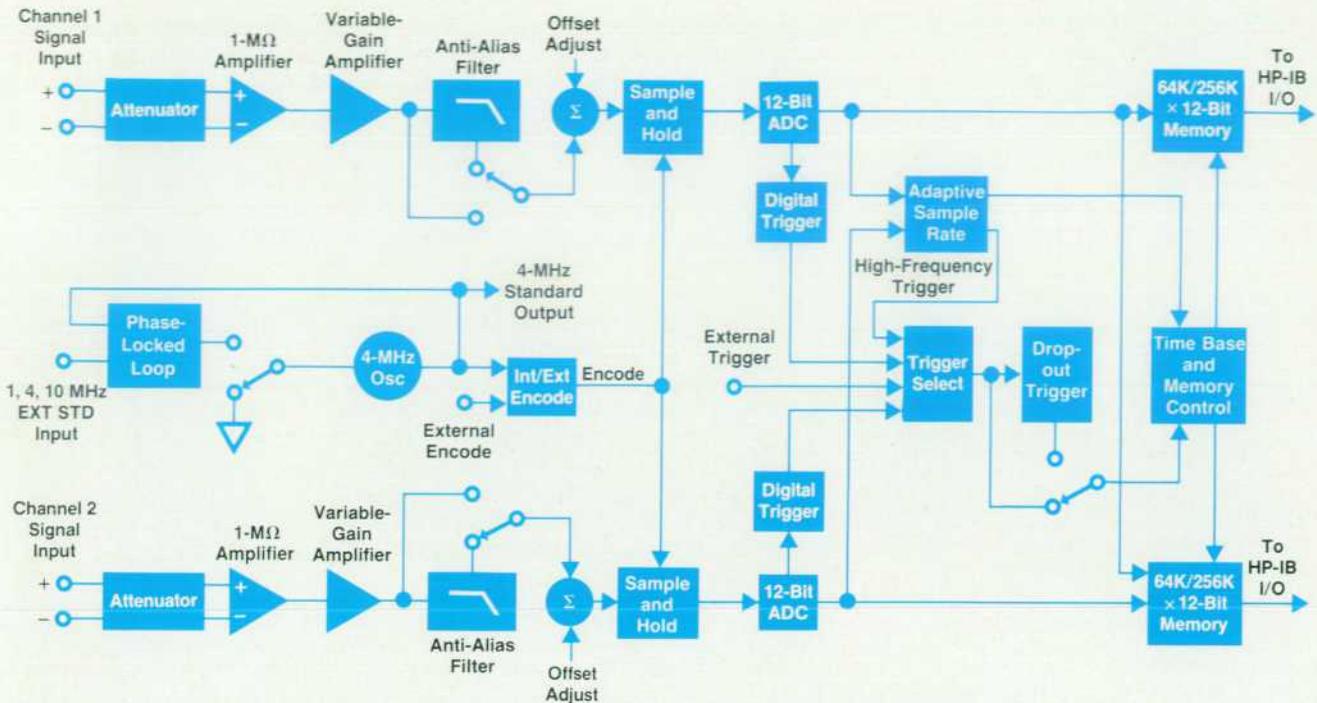


Fig. 4. HP 5183A Waveform Recorder simplified signal flow block diagram.

Burst Time Base. The burst time base mode allows the operator to record a selectable number of data samples automatically after a trigger event. The number of samples can be any integer from one to the total memory length. After a rearm delay of two sample points, the HP 5183A is ready to record the next burst immediately. This is an extremely efficient technique for recording hundreds or thousands of data bursts in a single record. No data is recorded between bursts (Fig. 8).

The circuit consists of a counter in the time base system. It is set for a specified number of sample points, then is rearmed and waits for the trigger event. A typical application is to analyze peak pulse amplitudes of a series of waveforms. Examples include radar return pulse analysis, Otto cycle engine instantaneous power peak analysis, and digital magnetic media error testing.

Adaptive Sample Rate. Adaptive sample rate (Option 301) automatically slows down the time base when higher sampling rates are not required, preserving all timing information and signal shape. Details of the signal processing are covered in the article on page 23. Adaptive sample rate effectively increases memory up to 60 times or up to 30 million samples.

Reference Oscillator. The 4-MHz crystal oscillator is an extremely critical time base subsystem. It must always maintain sufficient purity to ensure the fidelity of the ADC. For example, with a 1-MHz input signal and 12 bits of ADC resolution, $\frac{1}{2}$ bit of noise distortion will be induced by 39 picoseconds of jitter (Fig. 9). For this reason, the oscillator consists of an extremely quiet voltage-controlled oscillator in a sophisticated phase-locked loop with a 4-MHz crystal reference. The VCO can be locked to a 1-MHz, 4-MHz, or 10-MHz external high-stability oscillator or to the Option

010 internal oven oscillator to provide the stability needed for mathematically comparing waveform data taken at significantly different times or different temperatures.²

Memory System

The typical waveform recorder memory system consists of static read/write memory (SRAM) IC chips controlled by a counter and an address register. Although this is a satisfactory scheme, SRAM is more expensive, takes more power, and is less dense (fewer bits per IC) than dynamic read/write memory (DRAM). In a computer architecture,

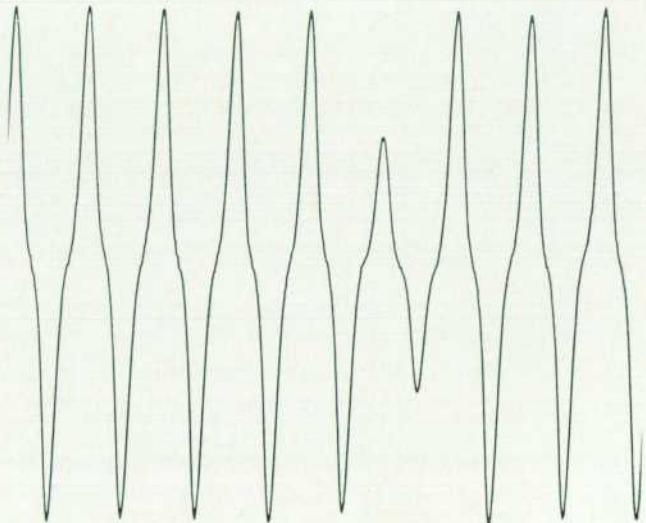


Fig. 5. Dropout trigger is a new feature that can capture brownouts or unexpected load increases.

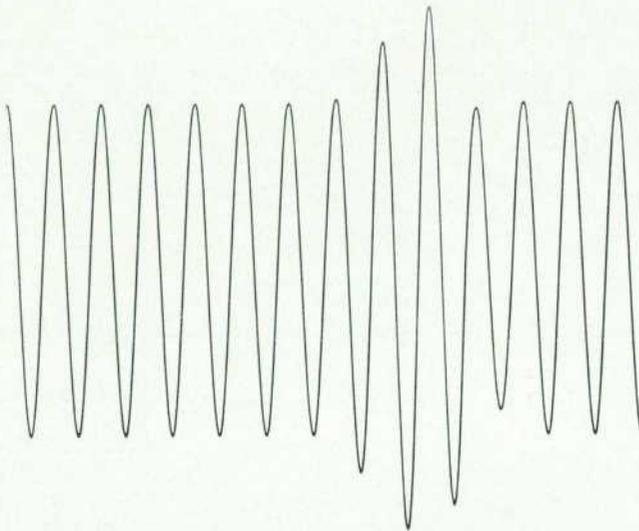


Fig. 6. Bi-trigger can be used to capture an unexpected load decrease or high-voltage spikes.

where the necessary cycle stealing control circuits for memory refresh (mandatory for preserving memory) are relatively straightforward, DRAM is much less expensive.

In waveform recorders, access to the memory must be assured whenever a sample point is taken. The HP 5183A data rate can vary from nanoseconds to seconds per sample point. The memory must start and stop as demanded by the user (a function of memory size and trigger time). In addition, the adaptive sample rate system may ask the memory system to decimate 63 out of every 64 sample points taken. Although a complex digital control problem, we chose a DRAM system for its advantage of greater memory size at far lower cost.

The HP 5183A memory system provides input synchronization, first in, first out (FIFO) buffering, shift register buffering (for delay of ADC data), and output synchronization with a minimum amount of hardware (Fig. 10, page 14). Among the design contributions in the memory system is a unified hardware implementation for a combination FIFO buffer and a programmable shift register having synchronous input and output capabilities. The design is not subject to errors caused by metastable states of the input and output flip-flops.

Clock Generator. The memory generates an internal clock of approximately 4.6 MHz. This frequency is used so that the instantaneous data rate of the memory is slightly higher than the 4-MHz maximum sample rate. This 0.6-MHz margin allows for refresh overhead (cycle stealing). The front end of the memory runs on the phase 2 clock and the rest runs on phase 1.

Data Synchronizer. Data from the ADCs (two 12-bit ADC data channels) is clocked alternately into 24-bit registers. The data is clocked in by the data clock coming from the ADC, and clocked out by the phase 2 clock. When a word is being read into one register, the previous word is being clocked out of the other register. This guarantees that the data will not be changing as it is read out, regardless of the relative timing of the data clock and the incoherent

phase 2 clock. An input flip-flop controls which of the registers accepts ADC input data and an output flip-flop controls which register outputs the data. The two flip-flops are connected through a synchronizing flip-flop which provides the transition between the two different clocks in the system. An exclusive-OR gate generates the output signal to indicate whether data is available from the data synchronizer during the current 4.6-MHz clock period. The adaptive sample rate speed and data valid signals are processed through similar synchronizers.

Master FIFO. The master FIFO stores the adaptive sample rate speed bit and control data while the DRAM memory is refreshing and then releases it as the DRAM is able to accept it. The buffer consists of static RAM. It is written to on the phase 2 clock and read out on the phase 1 clock. Four address bits are required. Multiplexers switch the data and address between read and write. The addresses are generated by two four-bit counters. The write counter increments whenever new data comes in and the read counter increments whenever data is read out by the DRAM. Hence, the read counter is constantly chasing the write counter. During refresh periods, the write counter continues to advance and store, while the read counter is held. When no new data is coming in, the write counter stands still while the read counter catches up. A comparator looks at the two addresses and generates an empty signal if they are equal. This instructs the DRAM not to try to take any more data. The FIFO has a capacity of 15 words. A maximum of 9 words are used during DRAM refresh delay.

The data control signal coming out of the FIFO is routed to a qualifier to determine whether the ADC data should be stored. This allows the adaptive sample rate system to discard redundant data. This happens downstream from the FIFO so as not to upset the counters. Hence, data will be written to the DRAMs unless the adaptive sample rate system says no, the FIFO is empty, or the DRAMs are in a refresh cycle. The adaptive sample rate speed bit is sent

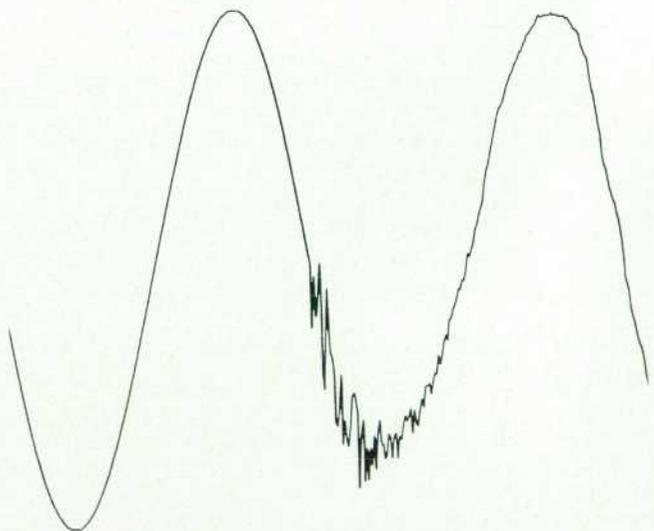


Fig. 7. High-frequency trigger can capture high-frequency transients that do not exceed the nominal peak-to-peak voltage of the signal.

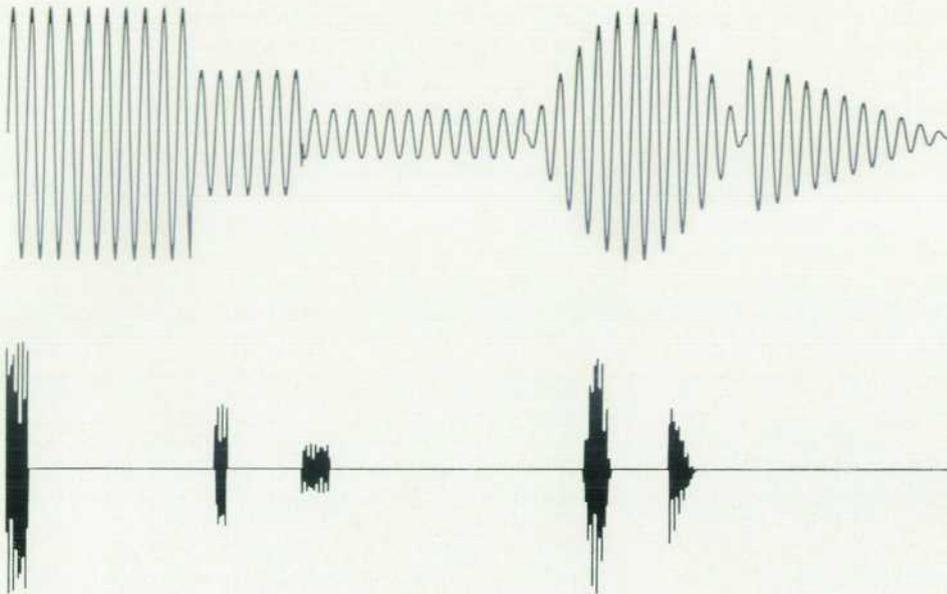


Fig. 8. Burst time base mode is used to record a selectable number of samples after a trigger and then stop, rearm, and begin recording again after another trigger. Memory is conserved because no samples are recorded during uninteresting periods between bursts. Recorded signal (top) for input signal (bottom).

directly to the DRAM array as a data bit.

ADC data passes through the slave FIFO. This FIFO is similar to the master except that it has ten-bit address counters and a 24×1024 -bit static RAM array. It is identified as a slave circuit because the count enable signals from the master control its address counters. The read counter is preset to either 2 or 548 before making a measurement. This offset is maintained indefinitely, resulting in a 2-word or 548-word delay pipeline in the data path. This circuit can be thought of as a programmable shift register of length 2 or 548. The delay of 2 compensates for trigger system detection delays. The 548-word delay is invoked during adaptive sample rate operation to compensate for the delaying effects of the adaptive sample rate low-pass and smoothing filters.

DRAM Array. After the incoming data is synchronized, buffered, and pipelined, it is routed to the memory, which consists of 64K or 256K 25-bit words of DRAM. The 25 bits are needed for the two 12-bit ADC data channels plus the adaptive sample rate speed bit. Memory refresh is done on one row every 10 microseconds.

DRAM Address Circuitry. The addresses for the DRAMs are generated by an 18-bit data address counter and an 8-bit refresh counter. A multiplexer selects between the refresh address and the data address. In a repeating cycle, the refresh address is strobed in. The row address is then latched and a series of column addresses is sent. This is the page mode of DRAM operation.

Read Registers. The data played back from the DRAMs is read out one byte at a time to the data bus. The DRAM control circuit generates a sequencing signal to clock new data from the DRAM into the read registers after the previous byte has been accepted.

DRAM Control Circuit. The DRAM control circuit is a 48-state repeating sequence circuit driven by the 4.6-MHz clock. Two periods are devoted to refresh, and two periods are devoted to latching the row address. The other 44 periods are available for writing or reading data. An eight-bit counter is hardwired to divide by 48. A decoder circuit

switches the address multiplexer at the appropriate time and enables the DRAM RAS and CAS appropriately.

Other Design Features

The HP-IB system is a conventional state machine design. The high-current $\pm 5V$ power supplies are secondary switching regulator designs.

The mechanical package design contains a conventional card cage but the cage is constructed as a monocoque welded aluminum frame. This design yields extremely stiff and strong but lightweight packaging. The thermal and acoustic management considerations dictated the use of three low-noise cooling fans for maximum reliability and low noise. Front-panel keys and annunciators are slaved to the display, I/O, and analysis module in the HP 5183T/U or to a remote computer when the HP 5183A recorder is used.

Acknowledgments

Many, many people have contributed to success of the HP 5180T/U and HP 5183T/U. I would especially like to thank Mark Allen, Phil Deaver, Mike Detro, John Fenwick, Jack Folchi, Rick Karlquist, Nancy Nelson, Doug Nichols, and Celia Vigil, who all served as R&D project managers. Jim Ammon, Jason Cotton, Bill Daley, Patty Damron, Nina Delu, Al Gee, Andy Gong, Ron Felsenstein, Al Foster, Dick

(continued on page 14)

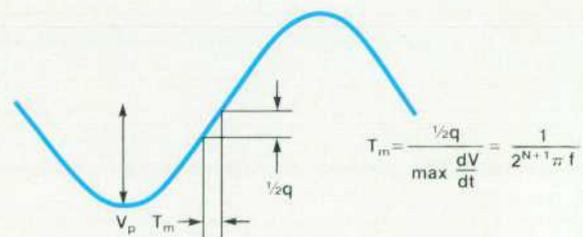


Fig. 9. With a 1-MHz input signal and 12-bit resolution, $1/2$ bit of noise distortion will be induced by 39 ps of jitter.

Waveform Recorder Software Design

The HP 5183A Waveform Recorder is tailored to the measurement requirements of a computer-controlled general-purpose waveform recorder and integrated with measurement application specific software. The HP 5183A consists of the precision waveform recorder hardware described in the accompanying article and measurement control software optimized for high measurement throughput.

While the HP 5183A Waveform Recorder and the HP 5183T Precision Digitizing Oscilloscope (an HP 5183A with a display module) are directed toward different applications, migration paths from one application to the other do exist. An example of moving from an HP 5183T to an HP 5183A is the development of a measurement procedure by a design engineer through interactive use and testing on the lab bench. When the procedure is implemented on an automated production line, the final algorithm can be implemented with a recorder. The HP 5183A offers faster performance and lower cost than the HP 5183T in the dedicated test system. Alternatively, measurements taken on remote recorders can be sent to a central HP 5183T to be processed using its powerful analysis capabilities. The HP 5183A would function as a lower-cost precision transducer, and the HP 5183T would be the central data reduction processor.

Software Architecture

The HP 5183A Waveform Recorder software is a library of procedures that allow the user to program the setup of the waveform recorder and control measurement cycles in the recorder. This software is typically used as procedures loaded with the user's software. Users can develop their own custom measurement systems and use the services provided by the HP 5183A software. Additional demonstration procedures allow the HP 5183A and its computer to function as a digitizing oscilloscope.

The primary contribution of the software is to provide an interface between easy-to-program high-level software and the difficult-to-program control register set that makes up the hardware. The software hides all the detailed information required to program the instruments; the user is presented with a standard procedure-oriented interface. This procedure library is provided in the BASIC language for HP 9000 Series 200/300 Technical Computers.

Architecture. The HP 5183A software is implemented as a set of standard procedures that are directly executed when called. Support of interrupts and multiprocess concurrency on the host computer are independent of the use of these libraries. These features can be implemented by the user in custom software.

Interfaces. Parameters passed to library procedures describe the operations to be performed by the library. The parameters are numeric or ordinal types that correspond to the physical parameters of the system where possible. More complex record type data structures supported in some languages are not used as parameter types. Advantages of this interface are that the same design can be ported to other languages or systems with minimum revision. It also allows users to program the instrument with the same numeric parameters they use to analyze data, rather than the ASCII string command embedding required by most programmable instruments.

To maximize usability, the interfaces are designed to be similar to other HP instrument controller software systems. In particular, the interfaces are similar to those used in the earlier-generation

HP 51800 Waveform Recorder Software and will be compatible with future waveform recorders and products. Since the software insulates the user from any machine specific interactions with the recorder, it will be source code compatible with other waveform recorder products.

Feature Set Extension. The software supports all hardware capabilities of the HP 5183A. Analysis or enhanced display capabilities can be added easily using third-party software.

Software Data Structures. The HP 5183A software can be modeled as two fundamental data structures. One data structure contains the current recorder state in the form of control codes written to the recorder when it is programmed. This low-level or bit-level description of the instrument setup is a software image of the recorder control registers. This type of data structure contributes to the high measurement throughput.

The second fundamental data structure prevents complex interactions between parameters from introducing any calling order dependence on the subprograms. This structure remembers the parameters that are interactive, so the final instrument state can be made independent of the library subprogram calling order.

Recorder State. The recorder state completely describes the state of the waveform recorder. It contains additional information required by the software, including instrument number, HP-IB select code, and bus address for multiple-instrument systems. This data structure is treated as a read-write variable. It is copied or rewritten if recorder setup information is internally saved or recalled to a file or to disc.

A library procedure is provided so the user can easily determine the current state of any HP 5183A in a system. This procedure formats the state information and displays it on the controller display.

For each recorder in a multiple-recorder system, the software instantiates a copy of the two data structures. All references to the particular instrument being programmed for setup or measurement are made by setup number. In this way, the number of recorders in the system is dynamic, and each recorder is separately controllable.

Software Operations

In controlling the waveform recorder, the software performs certain operations on its data structures. These operations include: setup, translate, measurement, initialize, calibrate, and verify.

Setup Input Parameters. These operations take user-entered values for recorder setup, enter them into the appropriate data structures, and update the hardware to reflect the new setup state. These operations are implemented as procedures called by the user software; value parameters describe the new setup state. This layer of software can be modeled as a shell that hides the complex hardware register programming task from the user.

The setup layer is implemented as BASIC subprograms; parameters to the subprograms define the hardware capabilities controlled by the software. Some checking of the input parameters is performed. Invalid parameters are logged as errors by the called subprogram. Error conditions cause the subprogram to exit without modifying the setup state. Consistency checks, which generate warnings when a parameter is limited or adjusted to the nearest allowed hardware value, are deferred to a separate subprogram. This approach, as opposed to having each subprogram issue warnings, avoids having the same warning issued

more than once and avoids issuing a warning that is not warranted by the final instrument state. Once the user has made all the changes to a setup, subprogram `B83check_setup` can be called to do a consistency check.

An example of a setup procedure implemented in BASIC is:

```
SUB B83set_input (INTEGER Setup, Channel, REAL Range, Offset,
                OPTIONAL INTEGER Coupling, Config, Filter)
```

Setup	1 to 32	index to data tables for this instrument
Channel	1 to 2	
Range	0.1 to 50	volts select in 1-2-5 sequence
Offset	$\pm 2 \times \text{range}$	entered in volts
Coupling	0,1	dc-coupled/ac-coupled
Config	1,2,3	single-ended/differential/disconnected
Filter	0,1	filter out/in

Measurement. These procedures initiate and control a measurement cycle in the waveform recorder, using the instrument setup previously entered. At the completion of these procedures, the sample data from the waveform recorder has been read into the computer memory and is ready for processing by the user software. An example of a measurement cycle using the BASIC software is as follows:

```
SUB B83start_meas (Setup)
  REPEAT
    Temp = B83meas_done (Setup)
  UNTIL Temp
  B83transfer_rec (Setup, Channel, Buffer, Header(*))
```

Setup	index to instrument being controlled
Channel	channel to be transferred
Buffer	I/O path name to predefined data buffers
Header(*)	pointer to an array to hold the waveform header information required to scale the data

Software is supplied in source code format for BASIC. The BASIC system consists of about 3000 noncommented source statements.

Performance Results

In many applications the ability to acquire many waveforms rapidly translates directly into reduced production test overhead. The measurement throughput numbers shown in the following table are typical characteristics and can vary depending on parameters such as sample rate, time spent waiting for a trigger, etc.

	Measurements per Second					
	HP Vectra*	HP 9000 Model 236U	HP 9000 Model 310	HP 9000 Model 320	HP 9000 Model 330	HP 9000 Model 350
1024-Point Blocks with Data Transfer to Controller	8	16	13	20	22	27
Auto Advance: 1024-Point Blocks Stored in Internal Memory	29	40	31	55	60	70

*Requires the HP 82300A BASIC Language Processor Card.

Other Systems

A goal of the HP 5183A project was to provide a path for porting the software into other systems. The use of a structured

software design and the implementation in a high-level language were key strategies in this process. Also, the use of HP language extensions to BASIC was minimized.

MS-DOS Systems. Many instrument control and automated test systems are being developed in the MS™-DOS operating system for use on personal computers, such as the HP Vectra PC or the IBM PC. Using the HP 82300A HP BASIC coprocessor card, HP BASIC 5.0 runs directly in the PC. The HP 5183A software can then be run under BASIC 5.0. All the functionality of the HP 5183A software can be exercised. This includes instrument addressing and control over the HP-IB, waveform display to the CRT, and the creation of waveform data files on disc. Users can combine the waveform recorder control software and the superior HP-IB I/O processing and fast development time of HP BASIC with the industry standard MS-DOS data processing and spreadsheet packages to develop their own optimum measurement systems and to leverage existing investments in HP-IB and MS-DOS systems.

HP-UX/UNIX Systems. HP BASIC 5.0 allows direct file transfers from BASIC to HP-UX or UNIX® operating system environments. Using these systems, data files from the waveform recorder taken with the HP 5183A software library can be passed directly into HP-UX. The full range of UNIX data manipulation and communications capabilities can then be used to process the data. For example, under the BASIC operating system, data from the waveform recorder can be saved in a disc file. From the HP-UX operating system, the file can then be processed by any of the standard UNIX utilities, such as `sort`, `awk`, or `grep`.

MS-DOS is a trademark of Microsoft Corp.

UNIX is a registered trademark of AT&T in the U.S.A. and other countries.

John Ketchum
Project Leader
Santa Clara Division

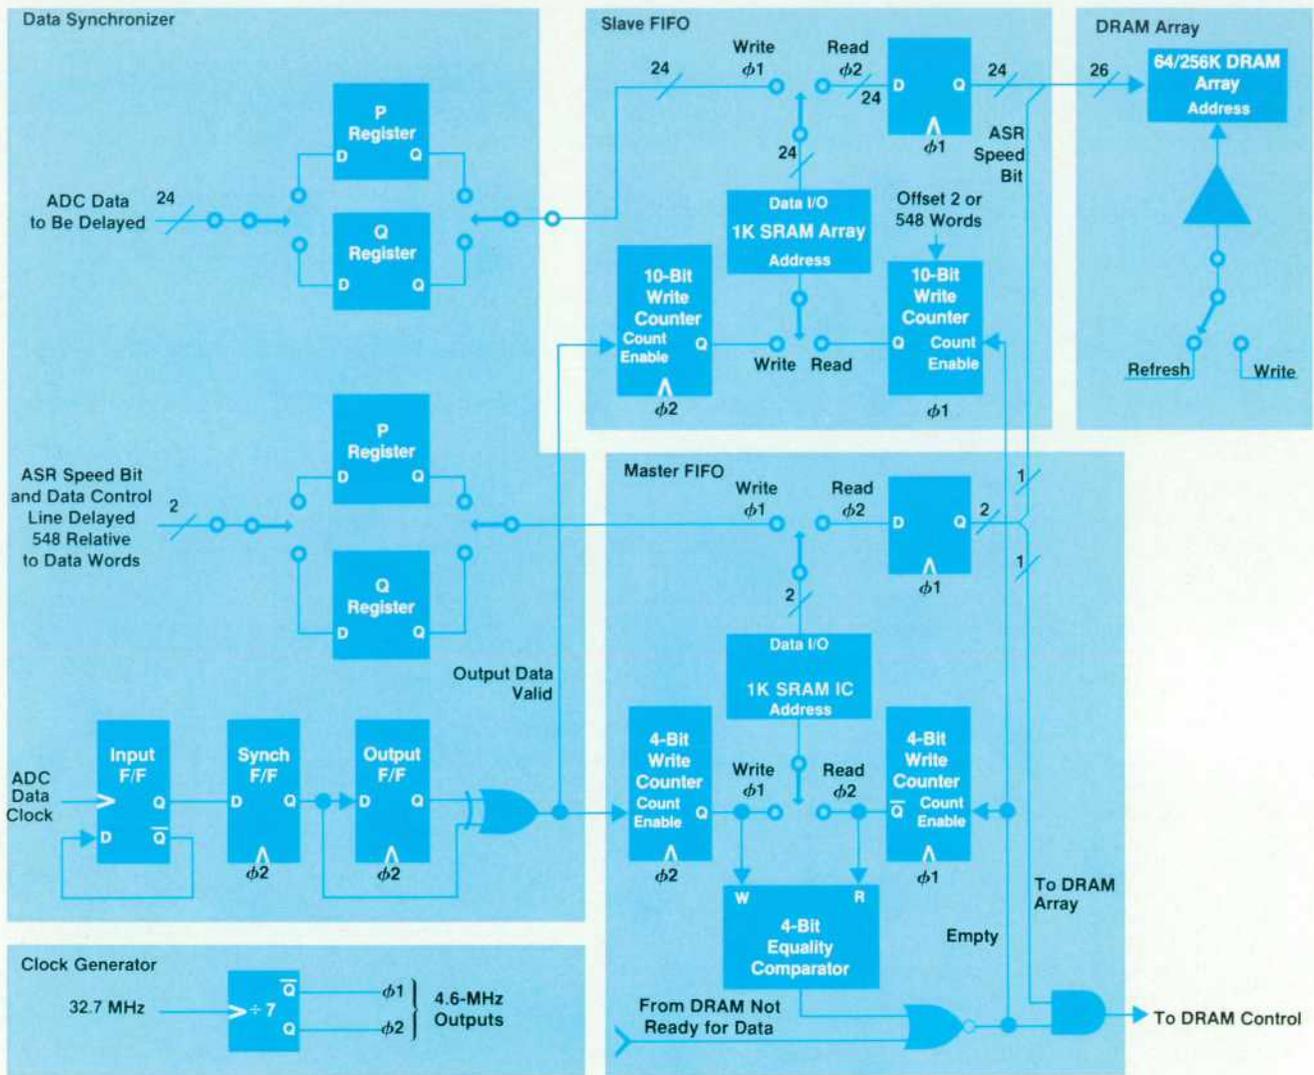


Fig. 10. HP 5183A memory system.

(continued from page 11)

Fowles, Ron Keeley, Alice Kwei, Rich Page, Don Schremp, Phil Scott, Ralph Smith, Dana Stoffers, Chris Szeto, Phil Vitale, Steve Will, Ron Young, and Russ Zandbergen are other dedicated R&D engineers who labored long and productively on these projects. Many others contributed years of work in the QA, production, and marketing departments. Special thanks to Al Scalise, key production engineer, who sorted out so many of the early production problems, and to Dan Hunsinger, Santa Clara Division R&D manager, for his continued support.

References

1. J.L. Sorden and M.S. Allen, "Waveform Recording with a High-Dynamic-Performance Instrument," *Hewlett-Packard Journal*, Vol. 33, no. 11, November 1982, pp. 3-9.
2. S.C. Bird and J. A. Folchi, "Time Base Requirements for a Waveform Recorder," *ibid.*, pp. 29-34.

Signal Conditioning and Analog-to-Digital Conversion for a 4-MHz, 12-Bit Waveform Recorder

by Albert Gee and Ronald W. Young

THE ANALOG-TO-DIGITAL converter (ADC) in the HP 5183A Waveform Recorder samples at a rate of four million samples per second and provides 12 bits of amplitude resolution. The 4-MHz sample rate exceeds the minimum 2-MHz rate (Nyquist rate) needed for adequate characterization of input signals with 1-MHz bandwidth. With 12 bits of amplitude data, the ideal signal-to-quantization-noise ratio is 74 dB and the resolution is $\pm 0.05\%$ for full-scale input signals.

Key features of the analog circuitry include the design of a modern discrete operational amplifier, a discrete sample-and-hold circuit, series-parallel ADC topology with pipelined ADC timing, a low-noise oscillator, and pseudo-random noise to improve the ADC's linearity. The performance of the input amplifier and ADC systems has been thoroughly evaluated and characterized with both static and dynamic input signals. The static characterization involves measuring the transfer function (with a tracking loop measurement) to determine the linearity. The dynamic performance is analyzed using the DFT (discrete Fourier transform), curve fit, and histogram tests.¹

Digital signal processing techniques can be used on the raw digital data to increase the signal-to-noise performance;

these techniques include averaging multiple measurements and oversampling (with respect to input signal bandwidth) followed by digital filtering. These issues are discussed in detail later in this paper.

Input Amplifier

Two identical attenuators, one for each input, handle the input configuration and coarse ranging while presenting 1-M Ω input impedance to the signal sources (see Fig. 1). In each channel, input clamp circuitry is followed by two high-impedance buffers that convert the differential signal to a low-impedance system. This signal is then converted to a single-ended signal. Two switchable gain stages in cascade implement the 1-2-5 variable-gain sequence, producing a $\pm 1\text{V}$ -full-scale signal which goes to a switchable 1-MHz anti-aliasing filter and then to the output buffer, which incorporates a $\pm 200\%$ offset control. The control interface circuitry allows input configuration and offset calibration from the HP-IB through the internal register bus.

The input configuration has four possible states: ac-coupled, dc-coupled, ground, and reference. In the ground and reference states, the input BNC is decoupled from the attenuator and left floating. The ground state connects the

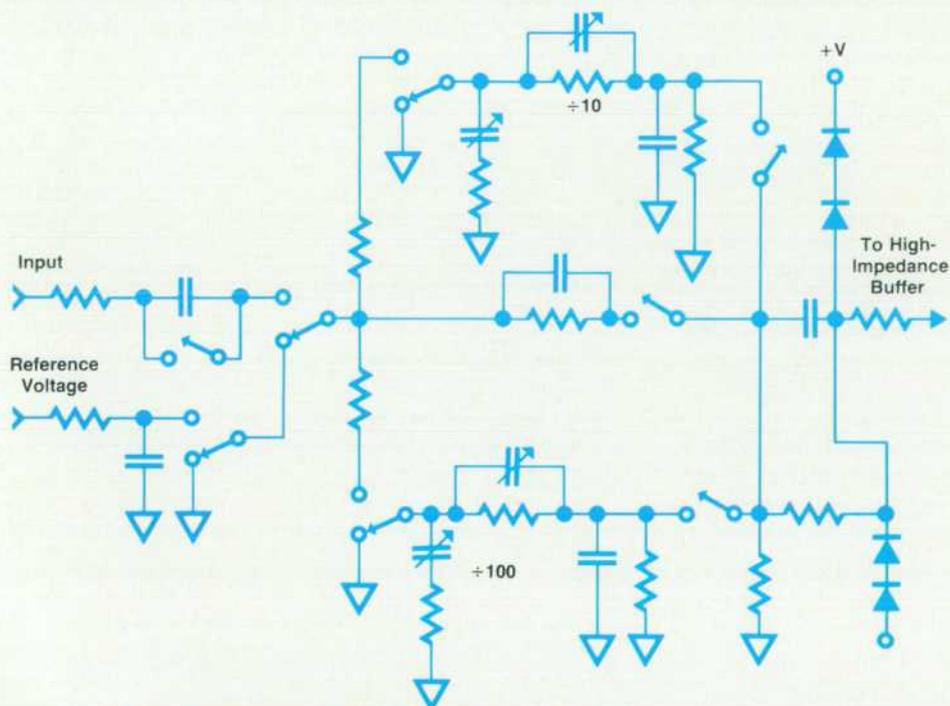


Fig. 1. HP 5183A Waveform Recorder input attenuator.

local ground to the input of the attenuator. The reference state connects a precision reference voltage to the input of the attenuator. The voltage reference is selectable from ± 100 mV to ± 10 V in a 1-2-5 sequence, plus 0V.

The attenuator provides coarse ranging as a selection of one of three decade attenuators: $\div 1$, $\div 10$, and $\div 100$. The selected range is terminated with a 1-M Ω resistor. The $\div 10$ and $\div 100$ attenuators have pole-zero compensation for parasitic capacitance. They also have an adjustment to match the input capacitance with the $\div 1$ attenuator. The nominal input impedance is 45 pF in parallel with 1 M Ω .

The clamp circuitry protects the FET in the high-impedance buffer from overvoltage. A large series resistor in front of the clamps limits the current in the clamp diodes for low-frequency overvoltages. For high-frequency overvoltages the resistor is bypassed with a capacitor to speed up the response of the clamp diodes. Two pairs of series diodes form the positive and negative clamps. The series diodes maintain the low-leakage characteristic of the circuit after repeated high-frequency clamping. The clamps limit the FET gate voltage to approximately ± 6.5 V.

The high-impedance buffer is a discrete FET-input op amp in unity-gain configuration. The discrete op amp is used throughout the input amplifier design. A detailed description is given later.

Up to this point in the system, there are two identical sets of circuitry, one for each differential input. The two signal paths are matched and have identical circuit layouts. Another discrete op amp, with similar topology, performs the differential-to-single-ended conversion. An active current source is used to bias the gain stage instead of a resistor. This increases the common mode rejection of the amplifier.

Each switchable gain stage is a discrete op amp in noninverting gain configuration. The first stage switches between $\times 5$ and $\times 1$. The second stage switches between $\times 2$ and $\times 1$. This allows the implementation of a 1-2-5-10 gain sequence. The partitioning of the gain is optimized for overall noise performance.

The anti-aliasing filter is a 1-MHz low-pass filter designed for overall frequency response of -3 dB at 1 MHz and below -60 dB at 3 MHz. It is a 10-pole matched-impedance elliptical filter, Gaussian to -6 dB. The filter can be bypassed by selecting the alternative signal path, which matches the insertion loss of the filter.

The output stage is another discrete op amp nominally in a noninverting gain-of-two configuration. The gain and offset adjustments occur at this stage. The offset control voltage is introduced at the summing node. The nominal output is ± 1 V with ± 2 V offset.

Discrete Operational Amplifier

Much of the performance of the input amplifier is dependent upon the performance of the discrete FET-input op amp (Fig. 2). The discrete op amp is designed to rely more on circuit linearity than feedback to achieve its distortion performance and bandwidth. The open-loop gain is approximately 50 dB, compared to over 100 dB for a typical commercial op amp. The FET input stage provides for low input bias current and the output stage can drive 50 Ω to 2V peak-to-peak at 1 MHz with low distortion.

The frequency compensation is symmetric, with one

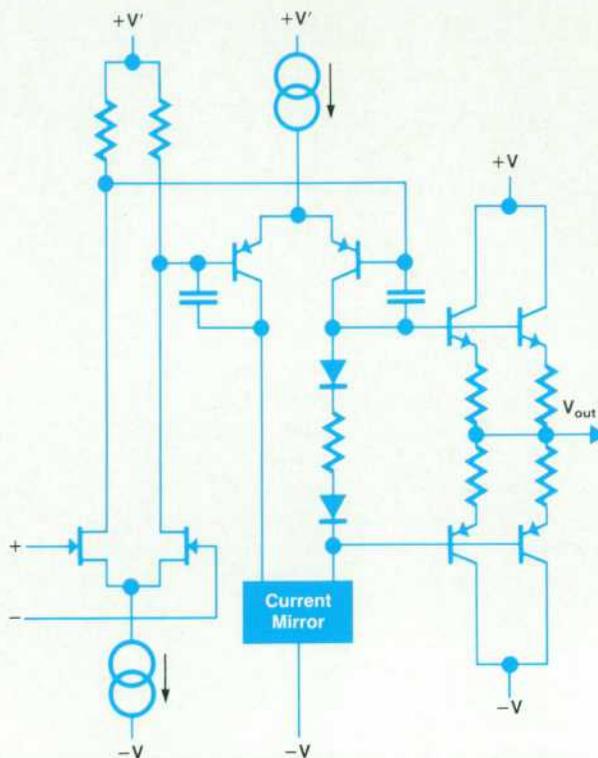


Fig. 2. Discrete FET-input operational amplifier used in the input amplifier.

capacitor from the base to the collector of each pnp transistor in the emitter-coupled pair in the gain stage. The capacitor connected to the low-impedance side of the active load introduces a nearly coincident pole-zero pair at low frequencies. The capacitor connected to the high-impedance side of the active load introduces the dominant pole for the open-loop frequency response.

The discrete op amps are optimized for 1-MHz distortion performance. The input bias current, which is determined by the FET, is less than 100 pA at 25°C. The input offset voltage is typically less than 10 mV and has a temperature coefficient of 50 μ V/°C. In a unity-gain configuration, the bandwidth is 40 MHz. The corresponding harmonic distortion is less than -75 dBc for a 10-dBm sinusoid input at 0.95 MHz. The slew rate is 400V/ μ s. The output noise in unity-gain configuration is 8 nV/ $\sqrt{\text{Hz}}$ at 500 kHz. The op amp dissipates 800 mW in its quiescent state with no load.

Analog-to-Digital Converter

Fig. 3 shows the simplified block diagram of the series-parallel ADC. The components of this subranging ADC include two sample-and-hold circuits, an analog multiplexer, an 8-bit A-to-D converter, a 12-bit D-to-A converter, and an error amplifier. The intermediate results of two 8-bit A-to-D conversions are combined to form the final high-resolution 12-bit result.

Sample-and-hold circuit #1 acquires and holds a sample of the input signal. This sample is routed via path A to the A-to-D converter, where a first-pass conversion produces a digital approximation to the analog signal. The precision D-to-A converter converts the digital approximation to an

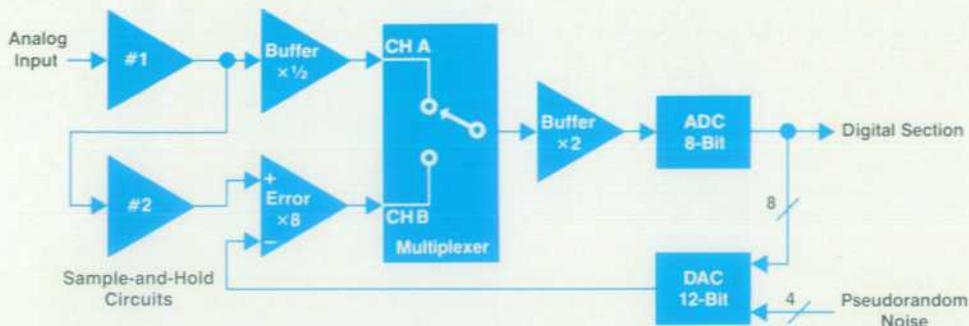


Fig. 3. Block diagram of the 4-MHz, 12-bit analog-to-digital converter in the HP 5183A.

analog signal, which is subtracted from the sample-and-hold output signal. This analog error signal is amplified and routed via path B to the A-to-D converter, where a second-pass conversion produces the digital error. The second-pass digital error is scaled to compensate for the error amplifier gain and is added to the first-pass approximation to generate the high-resolution, 12-bit representation of the analog input signal (Fig. 4).

An error analysis of this system shows that only the D-to-A converter and the sample-and-hold circuit need to have 12-bit accuracy. The rest of the system components need only have an accuracy consistent with the 8-bit A-to-D converter. A benefit of this converter, unlike successive approximation ADCs, is that errors in one pass can be corrected by subsequent passes if there are overlapping bits between the passes. This ADC has a speed advantage because the D-to-A converter needs to settle only once versus twelve times for a successive approximation ADC.

Pipelined Architecture. Even with its speed advantage, the series-parallel topology is not sufficient to achieve the desired 4-MHz sampling rate with the chosen components. The most time-consuming portion of the conversion cycle is waiting for the D-to-A converter and error signal to settle. The 4-MHz sampling rate is attained by pipelining the system. The main idea of pipelining is that a repetitive sequential process can be converted to a higher-frequency process by inserting an appropriate delay element so that operations can be performed concurrently. In this case, pipelining is implemented by using a second sample-and-hold circuit as the delay element, as shown in Fig. 3. The second sample-and-hold circuit makes both the current analog sample and the next analog sample available concurrently. While the error signal for the current sample is settling, a first-pass conversion of the next sample is done.

Fig. 5 gives a detailed example of unpipelined versus pipelined timing of the ADC. In the unpipelined case, the sequence of events is: acquire the signal, perform the first-pass conversion, wait for the error to settle, and then perform the second-pass conversion. In the pipelined case, there are two parallel paths. For the first path, the sequence of events is: acquire the current signal and then perform the first-pass conversion of the current sample. For the second path, the sequence is: acquire the previous signal and then perform the second-pass conversion of the previous sample. For the example given, the sampling rate increases from 2.5 MHz for the unpipelined case to 4 MHz for the pipelined case.

An additional benefit of pipelining is that the stepwise outputs of the second sample-and-hold circuit and the D-to-

A converter can be timed to change simultaneously. Thus the error amplifier is not overdriven and exhibits only a very short switching transient.

Pseudorandom Noise. Since the D-to-A converter determines the overall system accuracy, a 12-bit converter is used. The eight most-significant bits are used for the first-pass digital approximation. The four least-significant bits are used to add pseudorandom noise to dither the analog error signal as shown in Fig. 3.

The purpose of pseudorandom noise is to randomize the quantization error associated with each ADC code. For instance, referring to Fig. 6, assume that the error function of the ADC transfer function is represented by the first graph and that the pseudorandom noise probability density is represented by the second graph. If the pseudorandom noise is independent of the input signal, then the resulting averaged error function is given by the convolution of the two graphs. The net effect, as shown in the third graph, is to improve the average linearity of the ADC. Since the digitally added noise is known exactly, the same noise can be subtracted out of the digital result. Therefore, pseudorandom noise has all the benefits of analog dither plus the important advantage of not contributing any noise power to the digitized signal. Thus large amounts of pseudorandom noise can be used, and in this case, up to 16 LSBs (= 4 quantizer bits) of noise are used. Note that with the gain ratio (ratio of first-pass gain to second-pass gain) mismatched, an attenuated component of the pseudorandom noise is present at the digital output. Thus the pseudorandom noise can also be used to calibrate the gain ratio by grounding the analog input signal and nulling the second-pass output for minimum noise power.

Pseudorandom noise is especially useful for repetitive

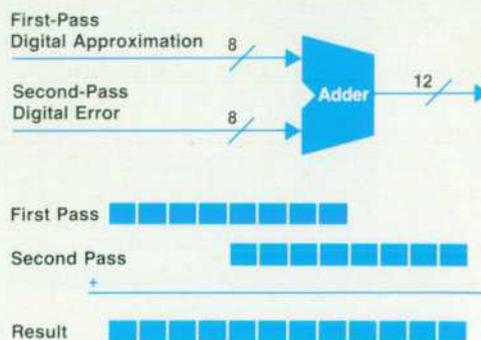


Fig. 4. The second-pass digital error is scaled and added to the first-pass approximation to form the 12-bit result.

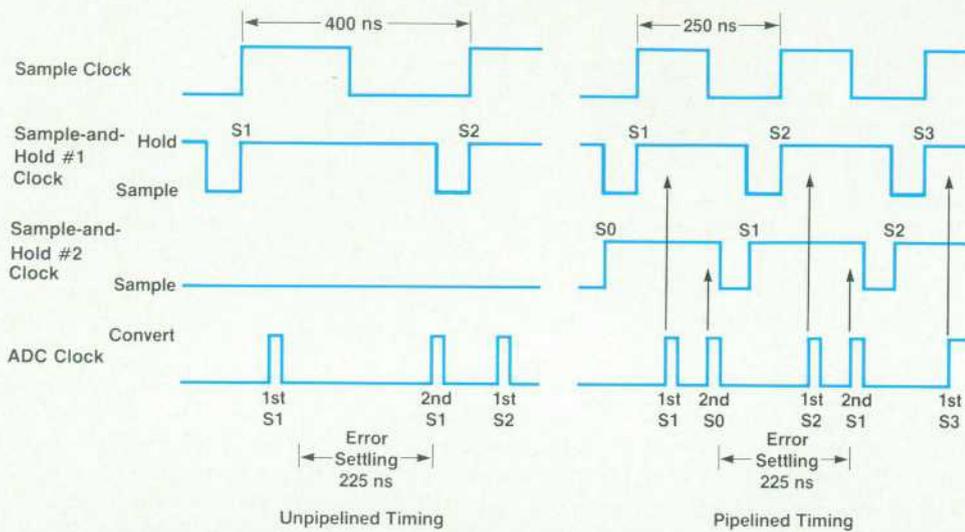


Fig. 5. Pipelining is used to achieve the 4-MHz sampling rate in the 12-bit ADC. The arrows show which sample-and-hold circuit is involved in each conversion in the pipelined case.

signals that can be averaged, signals that are phase coherent with the sample clock, or low-level signals where the quantization noise is less likely to be uniformly distributed. An example of the last case is shown in Fig. 7, which compares the spectrum of a digitized low-level sine wave with and without pseudorandom noise. The randomizing effect of the digital dither can be seen by the absence of the aliased harmonics in the pseudorandom noise case.

ADC Performance

In the curve fit test, a sine wave is digitized by the ADC, and from the resulting digital data, the best-fit sine wave that minimizes the mean squared error is calculated. Normally, the frequency used is one that will ensure testing a

majority of the ADC codes and is similar to the intended application bandwidth. In the dynamic testing of this particular ADC, full-scale sine waves with a nominal frequency of 1 MHz were used. Fig. 8 shows the resulting error in both the time and frequency domains of a curve fit test using a test frequency of 0.985 MHz; the number of effective bits is 10.65. The apparent performance of a high-resolution ADC can easily be limited by the noise of the input signal source. An HP 8662A low-noise synthesizer (with the output low-pass filtered) was needed to measure the number of effective bits accurately.

The DFT test is effective in separating the noise power into distortion (harmonics), environmental sensitivities (discrete lines), and white noise (flat noise floor) components. Analyzing the data in the frequency domain is particularly useful since nonlinearities present in the ADC produce harmonics of the input frequency that are readily detected. The harmonics are not restricted to multiples of the input frequency, since aliasing folds the higher-order harmonics ($f_x > f_s/2$) back down to between 0 and $f_s/2$ as given by the following equation:

$$f_{\text{alias}} = |f_x - f_s \times \text{round}(f_x/f_s)|$$

In Fig. 9 the spectrum of a digitized 0.985-MHz sine wave is computed. The highest spurious signal, second-harmonic distortion, is below -70 dBc.

Integral linearity measures the deviation of the transfer function data from the best fit (minimum mean squared error) line through the data. Integral nonlinearity highlights global nonlinearities, which can produce high distortion. Static differential nonlinearities greater than ± 1 LSB indicate potentially missing codes and nonmonotonicity. Fig. 10 shows both the integral and static differential nonlinearity of the ADC to be within $\pm 1/2$ LSB.

For dynamic signals, differential nonlinearity and missing codes can be best measured by using the histogram test with a sine wave input. The input test frequency should be chosen to ensure that all the ADC code levels are tested. The ADC raw histogram and linearized histogram results in Fig. 11 show the ac differential nonlinearity to be well within ± 1 LSB.

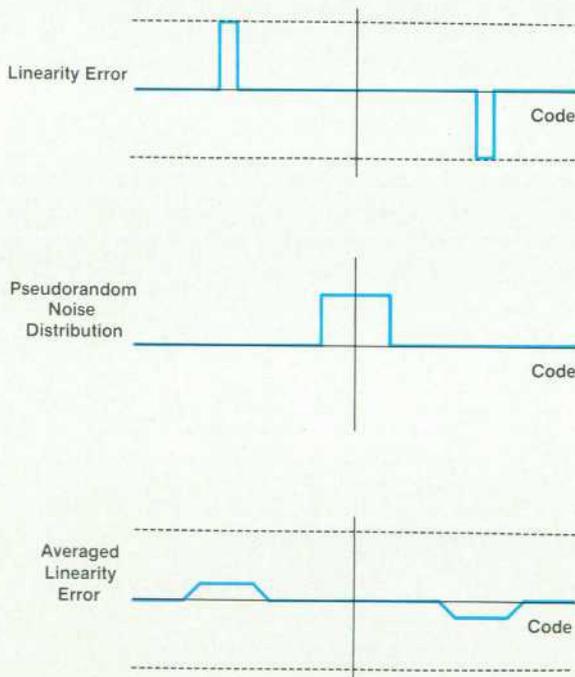


Fig. 6. Pseudorandom noise is used to dither the analog error signal and improve the average linearity of the ADC.

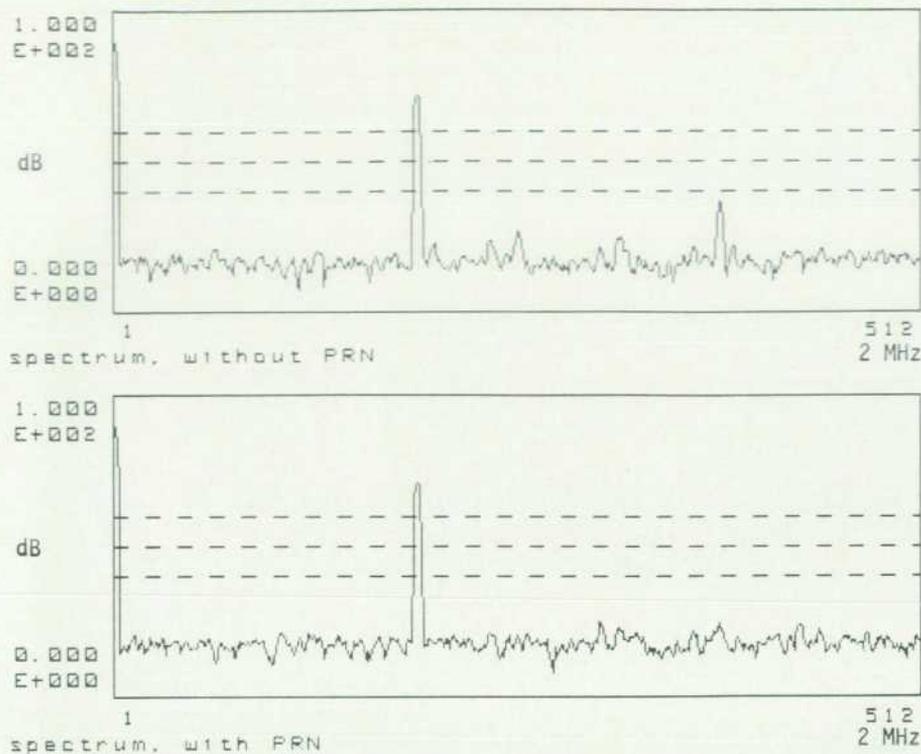


Fig. 7. Spectra of a digitized low-level sine wave with and without pseudorandom noise.

In an aperture jitter test, the input signal is phase-locked to the ADC sample clock, and the ADC is set to sample only on the zero crossings (regions of highest slew rate) of a high-frequency, high-amplitude sine wave. Experimental results for this converter indicate an rms aperture jitter of 28 ps or a degradation of 0.3 effective bit for a 1-MHz test input.

Maximizing Performance

In many instances, the signal-to-noise ratio performance of the HP 5183A can be improved by postprocessing the raw digital data with averaging and filtering algorithms. These signal processing techniques depend on random noise that is uncorrelated with the input signal. With the injection of pseudorandom noise into the ADC, as men-

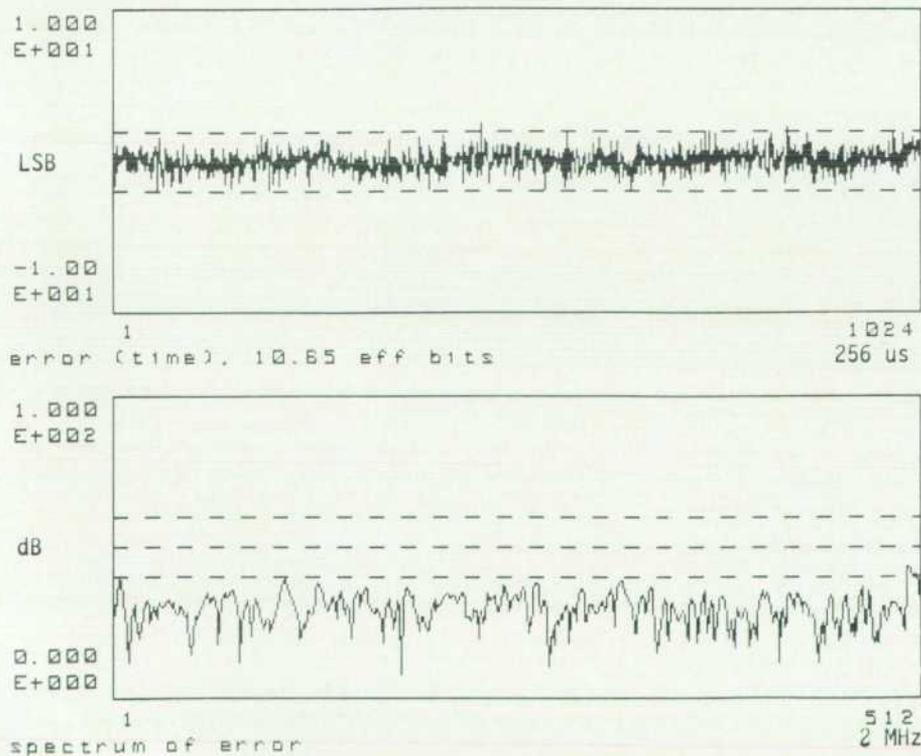


Fig. 8. (Top) Error in curve fit test for a test sine wave at 0.985 MHz. (Bottom) Frequency spectrum of the error.

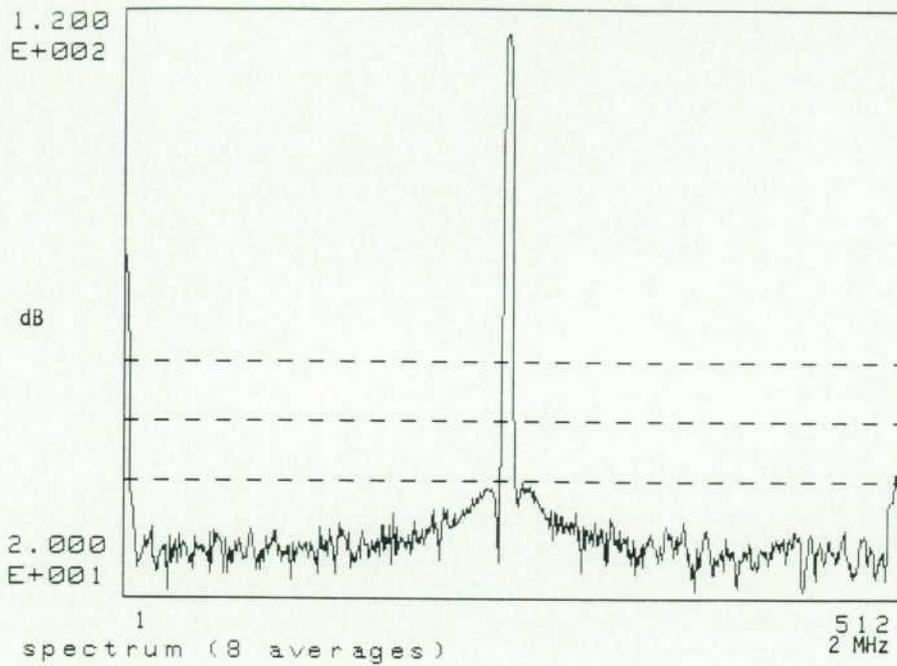


Fig. 9. Spectrum of a digitized 0.985-MHz sine wave.

tioned earlier, the quantization noise is randomized with respect to the input signal so that the noise can be considered white.

Averaging. If the input signal is repetitive and if the measurements can be phase-coherently triggered with the input signal, then averaging multiple measurements will decrease the noise power of a single measurement by a factor equal to the number of measurements averaged (n). Thus the signal-to-noise ratio of n averaged waveforms is:

$$\begin{aligned} \text{SNR}(n) &= 10 \log(\text{signal power}/(\text{noise power}/n)) \\ &= 10 \log(n) + \text{SNR}(1). \end{aligned}$$

For example, with sixteen averages, the SNR can be expected to increase by about 12 dB over the SNR of a single measurement. With averaging, the noise that is reduced can be the noise in the input signal and/or the noise generated internally in the digitization process. An example of the benefits of averaging is shown in Fig. 12. The upper

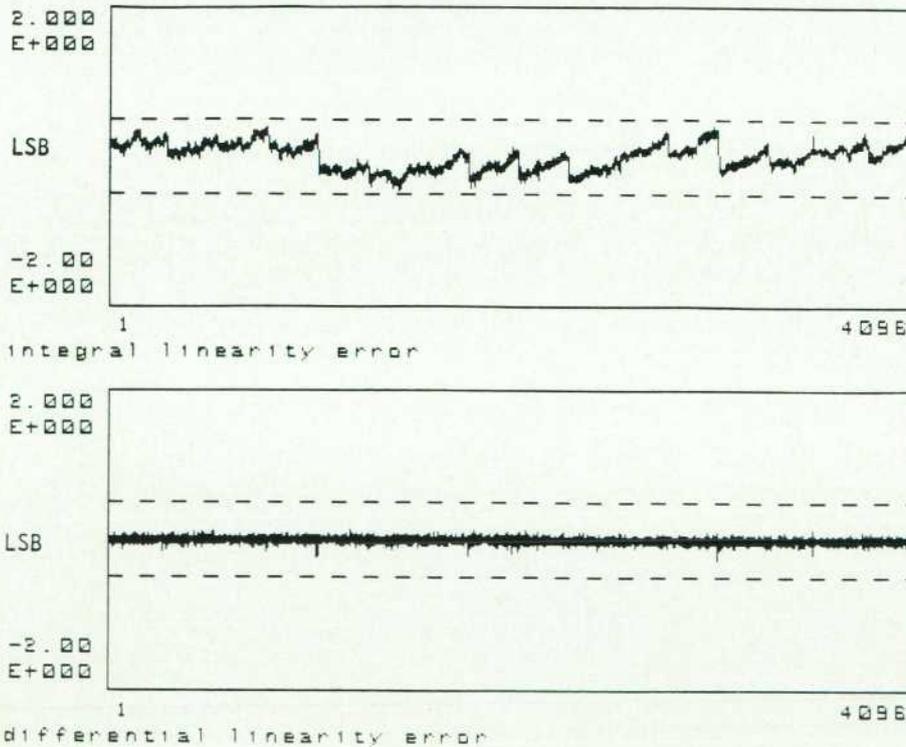


Fig. 10. Integral and static differential nonlinearity of the HP 5183A ADC.

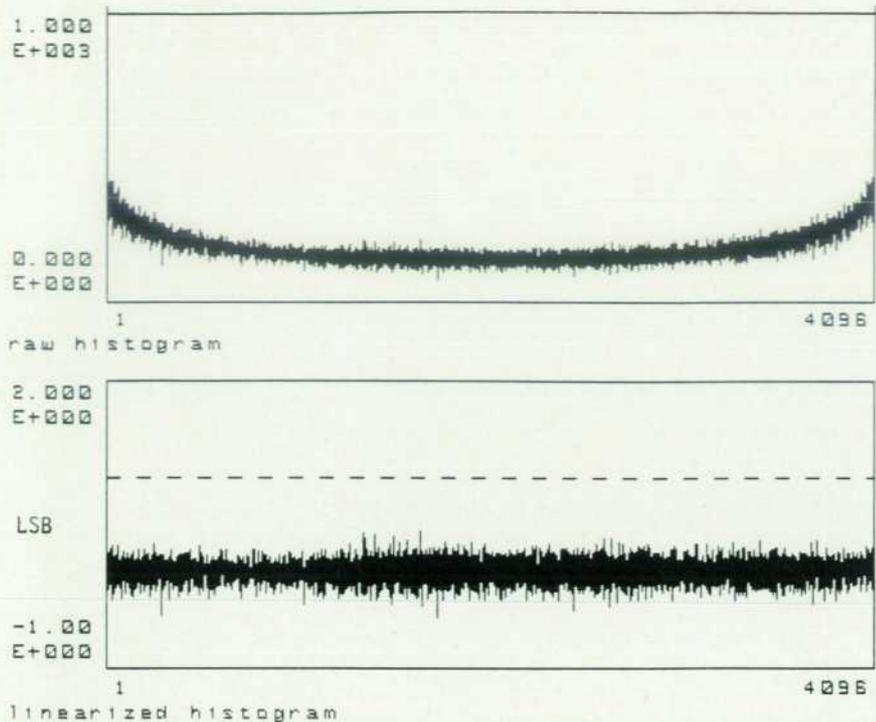


Fig. 11. Raw and linearized histograms show the ac differential nonlinearity of the ADC to be within ± 1 LSB.

plot shows the sum of the input signal and noise based on one measurement (no averaging). The lower plot shows the result of averaging 64 measurements. The original input signal characteristic is now easily discerned. It is important to note that averaging techniques to enhance SNR are ultimately limited by the linearity of the digitizer.

Oversampling and Postprocessing. The dynamic perfor-

mance of the HP 5183A is specified as 10 effective bits at 1 MHz. This specification is limited by the analog harmonic distortion performance at high frequencies. For lower input frequencies, the performance increases as the harmonic distortion decreases. Performance at lower frequencies is limited primarily by quantization noise. Thus, if the input frequency is lower than 1 MHz, the number of effective

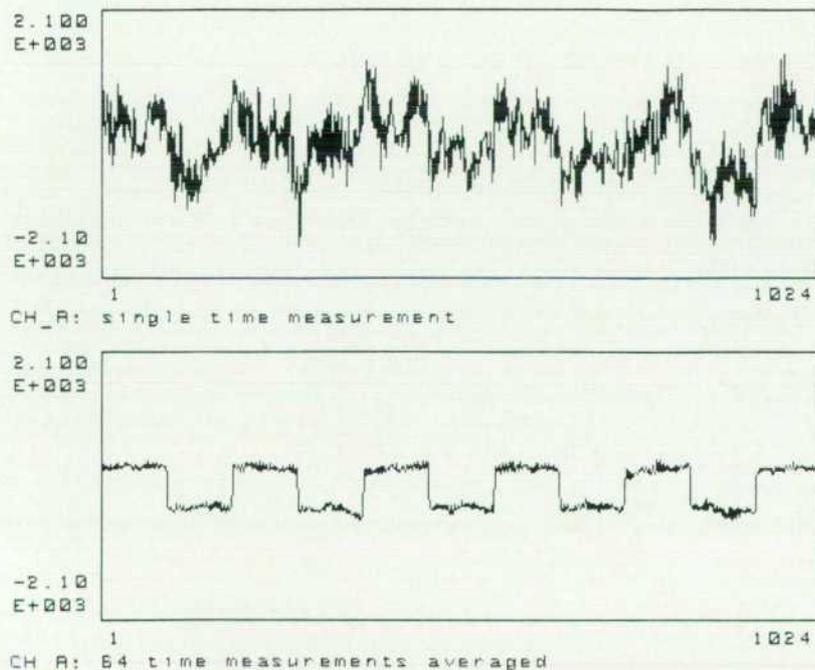


Fig. 12. An example of the benefits of averaging.

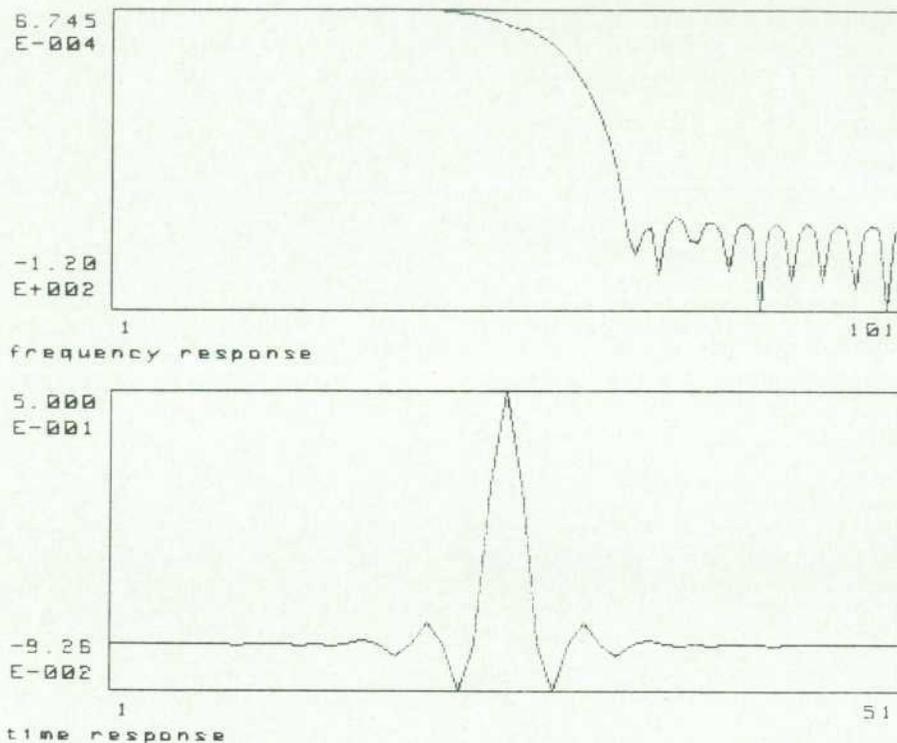


Fig. 13. Response of the digital filter used to increase effective bits for low-frequency signals by removing high-frequency quantization noise.

bits will be greater than 10 and lower harmonic distortion will be realized. However, an even better signal-to-noise ratio can be achieved by removing the high-frequency quantization noise as well. This can be accomplished by oversampling the signal and externally postprocessing the quantized data with a digital low-pass filter to reduce the noise bandwidth. If the input signal is oversampled ($f_s > n f_{in}$, $n > 2$) and if the quantization noise is white, then the noise power in the frequency range $f_s/2 > f > f_{in}$ can be digitally removed. The noise power left is $f_{in}/(f_s/2)$ of the original noise power, which corresponds to an increase in the number of effective bits as follows:

$$\text{Original noise power } N_o = \frac{q^2}{12} = \frac{A^2}{12(2^{2b_o})}$$

where q is the quantizer resolution and A is the full-scale amplitude. Rearranging terms,

$$b_o = \text{original effective bits} = \frac{1}{2} \log_2 \frac{A^2}{12N_o}$$

The filtered noise power is

$$N_f = N_o \left(\frac{BW_f}{BW_o} \right),$$

where BW_o and BW_f are the original and filtered bandwidths. Therefore, the number of effective bits after filtering is

$$b_f = \frac{1}{2} \log_2 \frac{A^2}{12N_f} = b_o + \frac{1}{2} \log_2 \frac{BW_o}{BW_f}$$

For every factor-of-two reduction in bandwidth (oversampling by a factor of 2) the number of effective bits can increase by $1/2$ bit (if quantization noise is the limitation and not ADC distortion). For example, with an input signal of 125 kHz maximum and a 12-bit ADC sampling at a rate of 4 MHz ($BW_o = 2$ MHz), this signal processing technique increases the ideal 12-bit ADC to a theoretical performance of 14 effective bits. Note that the noise floor in the frequency domain does not decrease, but rather the bandwidth is reduced. The costs of this technique are a reduction in maximum record length and input bandwidth, and additional computation time.

The signal is oversampled by setting the sample rate to $n f_{in}$ and setting the record length to n times the desired length. The digital filtering can be done in either the time domain or the frequency domain. We chose a linear-phase FIR (finite impulse response) filter with 50 coefficients (Fig. 13). The filter is designed to divide the input bandwidth by half the sample rate. To select bandwidths other than $1/2$, the filter is used iteratively in conjunction with data decimation (by a factor of two) to reduce the sample rate. Since the data is filtered before each decimation, there is negligible aliasing. With a 100-kHz signal (using a low-pass-filtered HP 8662A), the number of effective bits increases to 12.7; this compares with 11.0 effective bits without any digital filtering.

Reference

1. *Dynamic Performance Testing of A-to-D Converters*, Hewlett-Packard Product Note 5180A-2, 1982.

Adaptive Sample Rate: A First-Generation Automatic Time Base

by Richard W. Page and Nancy W. Nelson

WHEN CAPTURING TRANSIENT signals, it is desirable to use a high sample rate to preserve input signal details. This limits the maximum recording time of the measurement. In the HP 5183A Waveform Recorder with Option 301 Adaptive Sample Rate (ASR), the input signal is sampled at the selected fast rate only when there is significant detail present, and at a much slower rate when there is no significant detail. Thus ASR maximizes the recording time without compromising signal integrity. ASR continuously estimates the input signal's bandwidth and switches to a lower sample rate whenever possible.

The ASR circuit computes the ratio of the input signal's instantaneous spectral energy below the set threshold frequency to the total input signal energy. If this ratio is unity, all signal energy lies below the threshold frequency and the input signal can be down-sampled without any loss of information. The ASR circuit fixes the threshold frequency to be one hundredth of the Nyquist frequency and allows down-sampling by sixty-four.

Fig. 1 shows a typical ASR output and the input signal reconstructed from it, demonstrating that the ASR algorithm preserves the input signal.

ASR Implementation

Fig. 2 illustrates the data flow and computations within one channel of the ASR circuit. Two channels of ADC data are input at the user-selected time base rate. This data is

filtered by the spectral estimation filters which eliminate energy above the threshold frequency. Output from the estimation filters is used along with the input signal to compute instantaneous energy ratio. This ratio is computed by subtracting logarithms of the two signals. The resulting difference in logarithms is smoothed and compared to a set energy threshold. If the resulting smoothed ratio estimate is below the threshold limit, a slow sample rate is indicated for the current sample. Buffer control circuitry accumulates sample rate decisions over groups of 64 input samples and generates the composite fast/slow sample rate signal. This signal is used to control the signal that tells the waveform recorder's data memory whether to store data or discard data. It is also stored with the data waveform memory. Because this fast/slow signal is stored, it is possible, for a sufficiently low threshold limit, to reconstruct the input signal completely, including all timing information.

Although the ASR circuit independently examines and computes the required sample rate for each channel, both channels are stored at the same rate. Selecting the appropriate trigger source (either channel 1, channel 2, or composite trigger) causes the sample rate decision to be derived respectively from channel 1, channel 2, or the inclusive-OR of the decisions from channel 1 and channel 2. The ASR circuit always monitors the input signal and computes a sample rate decision. For down-sampling to begin, at least 256 consecutive ratio tests must indicate that down-sam-

Current Analysis Setup: 1
Function: Inverse ASR
Operand 1: Channel 1 Record 1 Start: 15787 ASR Len: 331

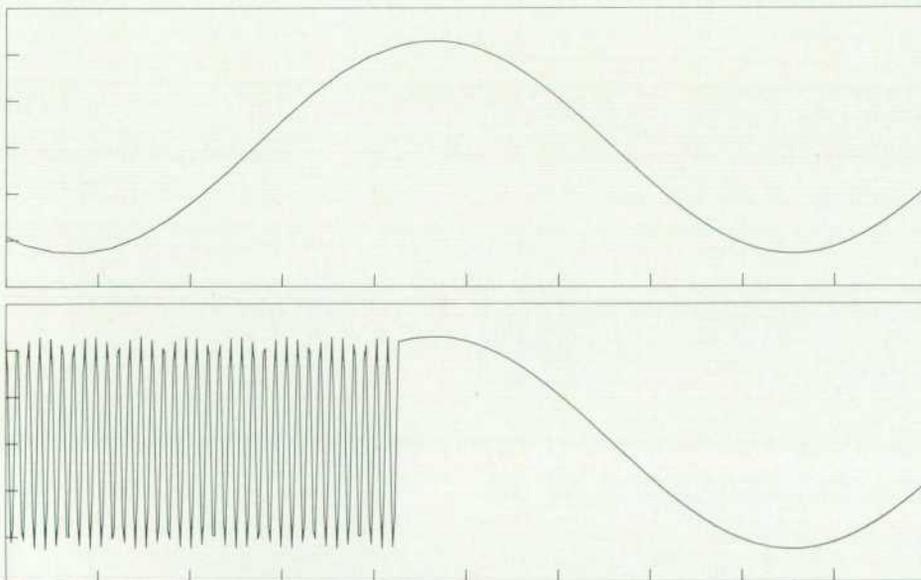


Fig. 1. HP 5183T display of the output of the adaptive sample rate circuit (bottom) and the reconstructed input waveform (top). The bottom waveform shows the actual memory contents. Before the trigger point, the input waveform is down-sampled; only one out of every 64 possible sample values is stored, making the display appear compressed. The top waveform shows that the original signal can be reconstructed from the stored data.

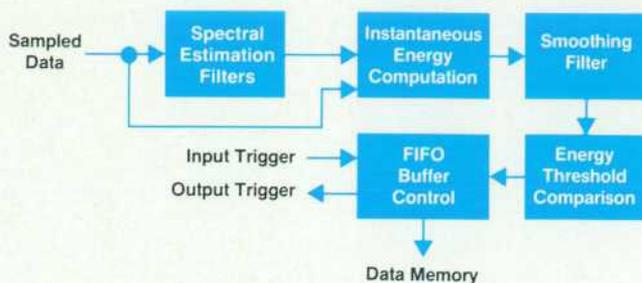


Fig. 2. Adaptive sample rate system single-channel data flow.

pling is possible. In addition, the ASR FIFO buffer must contain four consecutive blocks of 64 data samples that can be down-sampled. This guarantees that at least four samples will be stored in memory at the slow rate once down-sampling has been initiated. Down-sampling will continue as long as each consecutive block of 64 samples indicates that down-sampling is possible (Fig. 3). The ASR FIFO buffer delays the input signal to allow the ASR circuitry to change the sampling decision from slow to fast for the preceding 256 input samples. This ensures that data preceding the slow-to-fast sample rate decision is output at the fast sample rate, providing maximum detail around the sample rate transition.

The instrument trigger signal is sent to the ASR circuitry and used to ensure fast sampling around the trigger sample. This guarantees that the trigger sample will be stored in memory. The trigger signal is inclusive-ORed with the instantaneous sample rate decision derived from the energy ratio comparison. The ASR FIFO buffer responds to a trigger by forcing a slow-to-fast sample rate decision.

When ASR is disabled, the buffer control circuitry is held reset and the sample rate decision is still output to the trigger circuit. Setting the trigger condition to high-frequency trigger allows the trigger circuit to be activated by transitions on the sample rate decision line. This feature allows the instrument to trigger on transitions from slow to fast sample rate as computed by the ASR. This is useful because wideband noise is present in many transient capture applications before and after the occurrence of the transient event.

The user-selectable noise threshold prevents the ASR circuit from computing a fast sample rate decision for insignificant noise by forcing all signals below the noise

threshold to be ignored. The control is user-adjustable over 25% of the input signal range. Thus, the noise threshold control affects signals that are small and near zero volts, measured at the ADC. Clearly, if the input signal has a large dc bias, the noise threshold control will not be effective. A dc component can be eliminated by using ac input coupling feature or offset control. A large dc bias will also adversely affect the ASR circuit's ability to make the correct sample rate decision because this decision is based on an energy ratio.

The ASR circuitry computes partial short-time Fourier transforms using a custom CMOS circuit. This circuit is a 3- μm -process, 2000-gate gate array with 98% of the gates used. The gate array's pipelined arithmetic circuits are clocked at a 16-MHz rate. Four additions and two multiplications are computed in 250 ns. One custom CMOS circuit is used per input channel. The output of the custom CMOS circuit is a measure of the input signal energy below the threshold frequency. It is used to compute the energy ratio.

Instantaneous energy comparison results are computed by subtracting the logarithms of the squared magnitude of the input signal and the squared magnitude of the CMOS gate array output (short-time Fourier transform). These instantaneous results are statistically noisy. To remove the variance from the energy comparison, the instantaneous energy comparison results are passed through a time division multiplexed digital smoothing filter. This digital filter is a single-pole low-pass filter.

Signal Processing Theory

The ASR circuit computes spectral energy estimates by means of a partial calculation of the short-time Fourier transform (STFT), which is defined as:

$$X(k,n) = \sum_{t=-\infty}^{\infty} w(t-n)x(t)W_N(k,t) \quad (1)$$

where $w(t)$ is the N -sample window beginning at time n , $x(t)$ is the input sequence, and

$$W_N(k,t) = \exp(-j2\pi kt/N)$$

Here the sampling frequency is normalized to 1 Hz and $\Delta t = 1/f_s = 1$.

The STFT estimates spectral components by computing

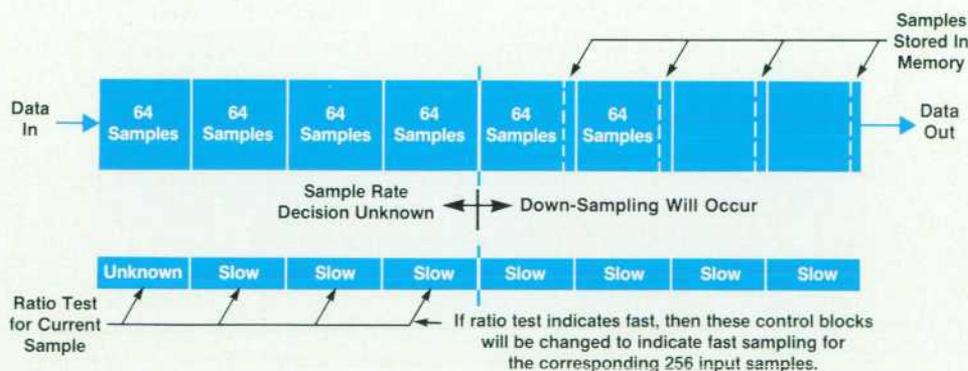


Fig. 3. Adaptive sample rate FIFO buffer operation.

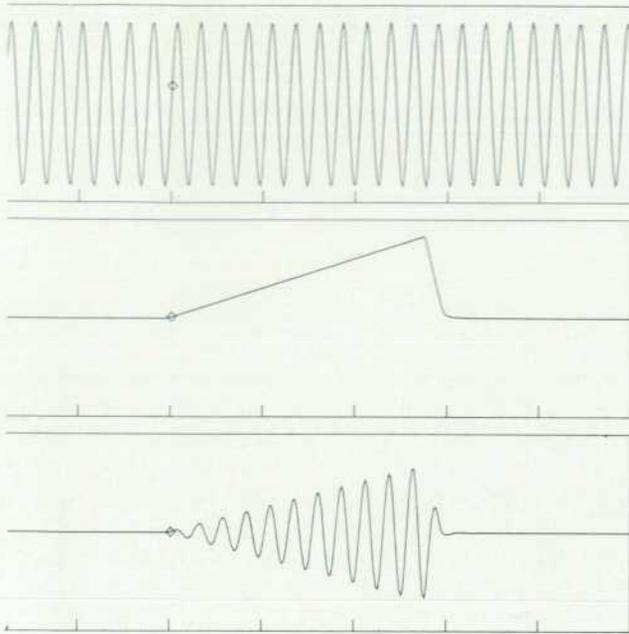


Fig. 4. Time function, sliding window, and windowed time function for short-term Fourier transform computations.

the discrete Fourier transform of the input signal multiplied by a sliding window (Fig. 4). The window $w(t)$ determines a filter shape that indicates how the spectral estimates are smeared by the finite observation time and how much spectral leakage will occur as a result of the window shape.

The STFT for bin k can be viewed as the output of a linear system whose input is the frequency-shifted (modulated) signal $x(t)W_N(k,t)$. The response is determined entirely by the window function. The STFT is equivalent to a bank of identical low-pass filters operating on the N sequences $x(t)W_N(k,t)$, $k = 0, 1, \dots, N-1$ (see Fig. 5). The output power from each low-pass filter is equal to the input power contained in the band centered about $\omega = 2\pi k/N$ as viewed through the spectral weighting function induced by the sliding window.

The filter bank analogy provides a convenient framework for examining spectral estimators. The analysis filter's main lobe width determines the ability to resolve components. Sidelobe behavior determines leakage (scallop loss or picket fence effect). The ability to track variations in the input spectrum depends on the estimation window and the filter bank output sample rate. Longer time windows result in improved spectral resolution but reduced ability to track nonstationary behavior.

The total signal energy for the ratio test beginning at time n is given by Parseval's relation:

$$E(n) = \sum_{k=0}^{N-1} |X(k,n)|^2 = \sum_{t=-\infty}^{\infty} |x(t)w(t-n)|^2 \quad (2)$$

The energy threshold T_L is:

$$T_L = \frac{1}{E(n)} \sum_{k=0}^{f_L(n)} |X(k,n)|^2 \quad (3)$$

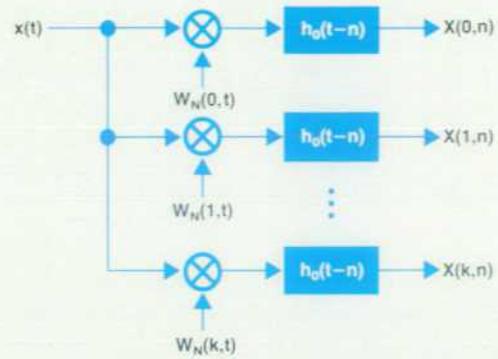


Fig. 5. The short-term Fourier transform is equivalent to a bank of identical low-pass filters operating on N weighted input sequences.

The threshold frequency $f_L(n)$ can be computed from equation 3 by solving for $f_L(n)$ with T_L fixed. Alternatively, T_L can be computed for a fixed value of $f_L(n)$. Fixing $f_L(n)$ and computing T_L requires less hardware. The current adaptive sample rate implementation assumes a fixed value of $f_L(n)$ and computes T_L . A comparison is made between the computed value of T_L and a fixed value. If T_L is greater than or equal to this fixed value, the sample rate is reduced. The fixed value is chosen such that 90% of the input energy will be included in the band from dc to $f_L(n)$.

Acknowledgments

The custom short-time Fourier transform CMOS circuit was designed by Patty Damron. Jim Sorden provided the initial concept of an automatic time base and supplied the resources and management support for this development.

Bibliography

1. M.R. Portnoff, "Time-Frequency Representation of Digital Signals and Systems Based on Short-Time Fourier Analysis," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. ASSP-28, February 1980, pp. 55-68.
2. J.B. Allen and L.R. Rabiner, "A Unified Approach to Short-Time Fourier Analysis and Synthesis," *Proceedings of the IEEE*, Vol. 65, no. 11, November 1977, pp. 1558-1564.
3. R.W. Shafer and L.R. Rabiner, "Design and Simulation of a Speech Analysis-Synthesis System Based on Short-Time Fourier Analysis," *IEEE Transactions on Audio and Electroacoustics*, Vol. AU-21, June 1973, pp. 165-173.
4. A. Paupolis, *Signal Analysis*, McGraw-Hill, 1977.
5. A.N. Kolmogorov and S.V. Fomin, *Introductory Real Analysis*, Dover Publications, 1970.
6. B. Widrow, *EE373 Adaptive System Course Notes*, Stanford University, 1977.
7. J.B. Allen, "Short-Term Spectral Analysis, Synthesis, and Modification by Discrete Fourier Transform," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. ASSP-25, June 1977, pp. 235-238.
8. J.O. Smith and J.B. Allen, "Variable Bandwidth Adaptive Delta Modulation," *Bell System Technical Journal*, Vol. 60, no. 3, May-June 1981, pp. 719-738.
9. J.S. Lim, et al, "Signal Reconstruction from Short-Time Fourier Magnitude," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. ASSP-31, August 1983, pp. 986-998.

Waveform Reconstruction Techniques for Precision Digitizing Oscilloscopes

by Richard W. Page and Allen S. Foster

WAVEFORM RECONSTRUCTION algorithms allow digitizing oscilloscopes to produce unambiguous displays, even when sampling near the Nyquist limit. In theory, a signal can be reconstructed from its samples if the sampling rate is greater than twice the highest frequency present in the signal. In other words, slightly more than two samples per cycle are theoretically enough to reconstruct a sine wave. In practice, one usually needs to sample at two or more times this Nyquist rate to get a reasonably faithful reproduction of a waveform from its samples. That is, the signal bandwidth usually cannot exceed about 50% of the digitizer bandwidth.

With special waveform reconstruction techniques, however, over 80% of the digitizer bandwidth can be viewed. The result is a lower-cost digitizing circuit and better use of high-speed acquisition memory. Fig. 1 shows an HP 5180T display of a sinusoid captured using 10 points per cycle without reconstruction. Fig. 2 shows a sinusoid captured using 2.5 points per cycle with reconstruction. The display of the higher-frequency sinusoid (2.5 points per cycle) is equal in fidelity to Fig. 1. In this case, waveform reconstruction increases the usable bandwidth by a factor of 4.

Comparing the reconstructed waveform to the raw data shown in Fig. 3, it is apparent that waveform reconstruction enhances postcapture processing. In particular, reconstruction allows examination of the fine structure in the data. Display measurement accuracy is increased for zero cross-

ing, peak value, and time interval determination measurements. Waveform reconstruction provides optimum viewable bandwidth for a given digitizing rate.

Without reconstruction, the display is typically an envelope of the minimum and maximum sample values. Fig. 4 illustrates what can happen. The signal is sampled at four points per cycle. When it changes phase, its amplitude appears to drop by 3 dB. This phenomenon arises because linear interpolation (connect the points) has been used to reconstruct the signal from its sample values. As the signal is oversampled by a larger margin, linear interpolation reconstruction becomes increasingly more reasonable and accurate. At five times the Nyquist rate, 10 points per cycle for sinusoids, linear interpolation is accurate within 5%. At this level of accuracy, reconstruction errors are acceptable. This has led to the "minimum ten sample points per cycle" rule in digitizing oscilloscopes for reasonable displays. However, the sampling theorem tells us that all of the signal information is preserved when we sample at slightly more than two points per cycle.

Representation of Signals

A signal can be modeled as a function of time, that is, a rule that assigns a value to the signal for all values of time, or by a collection of sample values taken at uniform intervals. The Nyquist-Shannon sampling theorem establishes a link between these representations. Nyquist showed that a sinusoid can be represented by its samples, provided

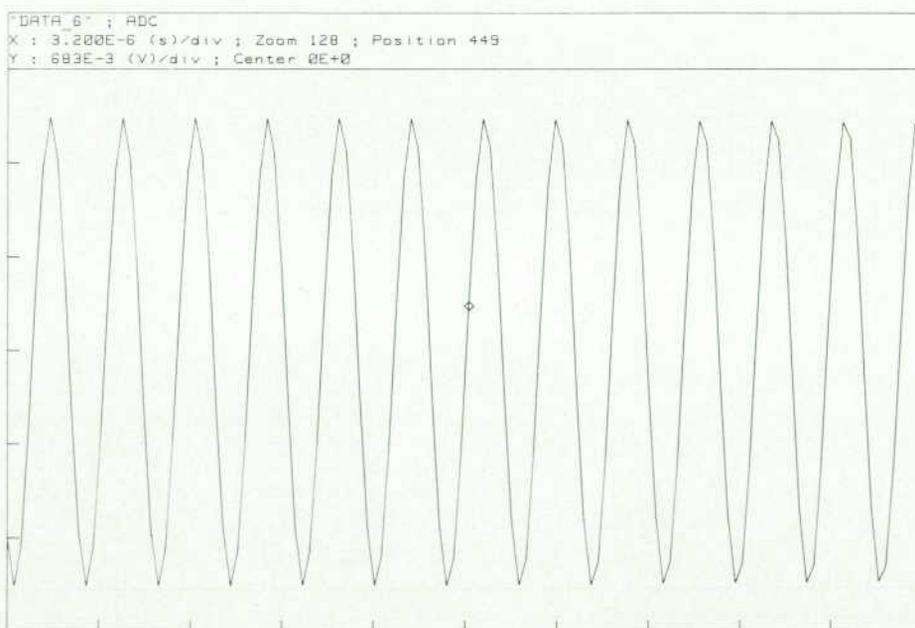


Fig. 1. HP 5180T display of a sinusoid captured using 10 points per cycle without reconstruction.

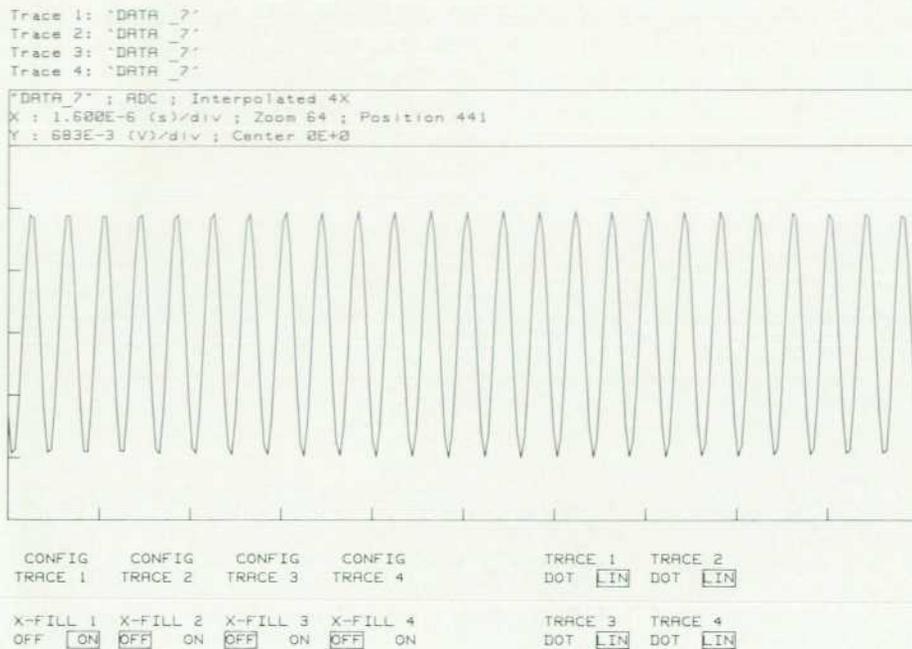


Fig. 2. HP 5180T display of a sinusoid captured using 2.5 points per cycle with reconstruction.

more than two samples per period are taken. Shannon showed that any band-limited signal can be represented by its samples, provided the sample frequency is greater than twice the frequency of the highest Fourier component present in the signal. Let $x(t)$ be a band-limited signal with Fourier transform $X(\omega)$. Since $x(t)$ is band-limited, $X(\omega) = 0$ for $\omega > \omega_{\max}$. The sampling theorem states that $x(t)$ is related to the sample sequence $x(nT)$ for all sample intervals $T < \pi/\omega_{\max}$ by the formula:

$$x(t) = \sum_{n=-\infty}^{\infty} x(nT) \operatorname{sinc} \left(\frac{1}{T}(t-nT) \right) \quad (1)$$

where

$$\operatorname{sinc} x = \frac{\sin \pi x}{\pi x}$$

According to the sampling theorem, a signal can be sparsely sampled, yet exact reconstruction is possible by applying equation 1. Without reconstruction, such a sample sequence is very difficult to interpret. For example, consider again the signal shown in Fig. 4:

$$x(t) = \cos(\omega t + \phi)$$

sampled four times per cycle (two times the Nyquist rate).

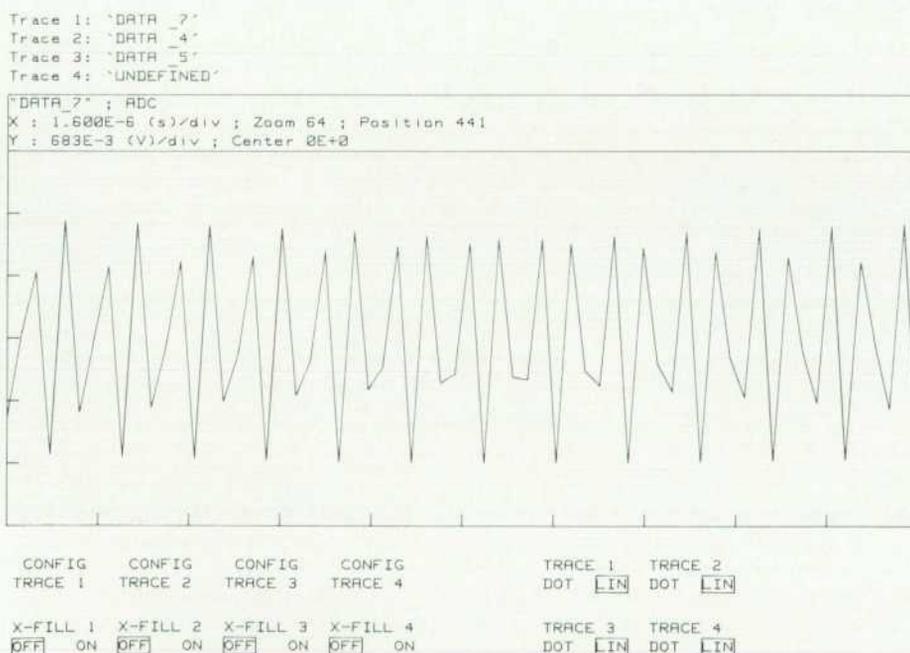


Fig. 3. Raw data for Fig. 2 displayed without reconstruction.

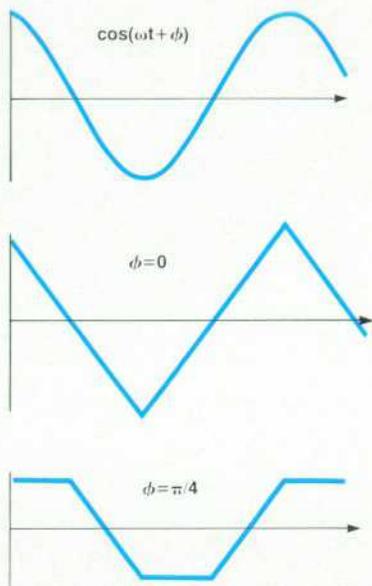


Fig. 4. When the top waveform is sampled at four points per cycle and displayed using linear interpolation (middle), its amplitude appears to drop by 3 dB when it changes phase (bottom).

When $\phi = 0$, the sample sequence is $\{1, 0, -1, 0, \dots\}$. Connecting the sample points with line segments yields a triangular waveform. When $\phi = \pi/4$, the sample sequence is $\{\sqrt{2}/2, \sqrt{2}/2, -\sqrt{2}/2, -\sqrt{2}/2, \dots\}$ which interpolates to a trapezoidal waveform whose amplitude differs from $x(t)$ by 3 dB. Thus the visual interpretation of an unreconstructed waveform sampled at twice the Nyquist rate is largely dependent on the phase relationship between the sampling process and the signal.

Interpolation as a Filtering Process

The spectrum of a sampled waveform is the sum of shifted replicas of the waveform's spectrum. Replication of the spectrum occurs every $2\pi/T$ radians per second, introducing symmetry about the points $\omega = n\pi/T$. The spectral images will not overlap if a signal is sampled according to the sampling theorem. The resulting spectrum in the interval 0 to π/T will be identical to the original spectrum. Exact reconstruction is possible in this case by passing the sample sequence through an ideal low-pass filter, as implied by the Shannon-Nyquist reconstruction formula.

As the sample frequency is increased above the Nyquist limit, spectrum images are spread farther apart. At 10 points per cycle, there is no significant energy between $\pi/5T$ and $9\pi/5T$, which eases the requirements on the reconstruction filter.

Interpolation can be used to alter the sampled signal spectrum given a signal sampled near the Nyquist limit. The resulting spectrum will be the same as if the original sample rate had been higher. The digital interpolating filter must compute, say, $R-1$ evenly spaced values between sample points and must not alter the original sample values.

It is convenient to model the process of interpolation as a polyphase filter. The impulse response for each subfilter and the interpolation error at each interpolated point can

be independently calculated. Each subfilter is a linear, time-invariant system, even though the general input-output relationship for a 1-to- R interpolator is a periodically time-varying filter with period R .¹

Each subfilter performs a computation of the form:

$$y_\lambda(n) = \sum_{k=-\infty}^{\infty} x(k)h_\lambda(n-k), \quad \lambda = 0, 1, \dots, R-1 \quad (2)$$

According to the interpolation condition requiring $y_o(n) = x(n)$, we have:

$$h_o(n) = \begin{cases} 1 & n = 0 \\ 0 & \text{otherwise} \end{cases}$$

The problem has been reduced to designing $R-1$ separate, linear, time-invariant systems. Interpolating filters can be broken down into two large classes:

- Filters having $h_\lambda(n)$ of finite duration
- Filters having $h_\lambda(n)$ of infinite duration.

These are called finite impulse response (FIR) and infinite impulse response (IIR) filters, respectively. The FIR structure offers the advantage of precisely linear phase independent of the magnitude response. The majority of interpolating filter design techniques generate FIR designs.

A finite impulse response interpolator is described by a convolution equation identical to equation 2 except that the limits of summation are finite. In the HP 5180T, HP 5183T, and 5185T Precision Digitizing Oscilloscopes, direct implementation of the convolution sum yields the tapped delay line structure shown in Fig. 5. At each sample time, $y(n)$ is the weighted sum of L input samples, where L is the number of weights or taps. The filter tap weights can be obtained by several methods including polynomial (Lagrange) or minimum norm methods.

Since the outputs of the R subfilters are orthogonal, the FIR polyphase interpolator can be modeled as a single filter. The input sequence must be augmented with $R-1$ zero values (zero padding) to maintain orthogonality of the subfilter outputs. The model clearly indicates the filter delay and the number of output samples generated before the filter impulse response dies out. Often, the data sequence to be interpolated is available and filter coefficients can be viewed as a weighting sequence rather than an impulse response. While the polyphase model is computationally more efficient, the single-filter model with zero padding illustrates the amount of data consumed by the interpolator.

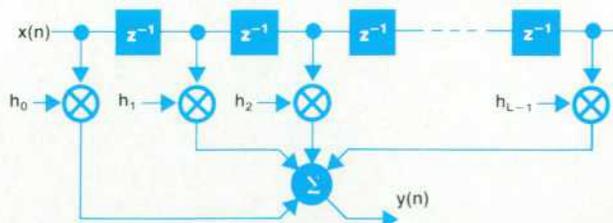


Fig. 5. Waveform reconstruction in the HP 5180T, HP 5183T, and HP 5185T oscilloscopes is implemented by a tapped delay line.

The longer the weighting sequence, the more data is consumed. This problem is alleviated by extrapolation of the data record. However, interpolation errors increase as extrapolated data is added to the computation.

Given a digitizer with dynamic performance specified up to the Nyquist limit, reconstruction algorithms can be used to increase the usable bandwidth. In the following examples, the HP 5180A and the HP 5183A digitizers were used.

Fig. 6a shows a sinusoid at 88% of the Nyquist frequency, or 2.3 samples per cycle. Without interpolation, the resulting display gives the appearance of an amplitude modulated sinusoid. Interpolation by four shows the correct envelope. Fig. 6b shows the same waveform expanded.

Notice that it is difficult to recognize that the input signal is sinusoidal without interpolation. Interpolation restores the signal and produces a displayed waveform typical of an analog oscilloscope display.

Fig. 7 shows the video sync portion of a single scan line sampled at four times the color burst frequency. Reconstruction allows the amplitude and duration of the sync signal to be measured accurately. Colors and their intensities are easy to measure with reconstruction. Similarly, phase change in the color burst signal can be seen with the phase transition expanded (Figs. 8 and 9).

It is often desirable to sample at less than the maximum digitizing rate. For example, storing a complete sector of flexible disc data in a digital oscilloscope memory may

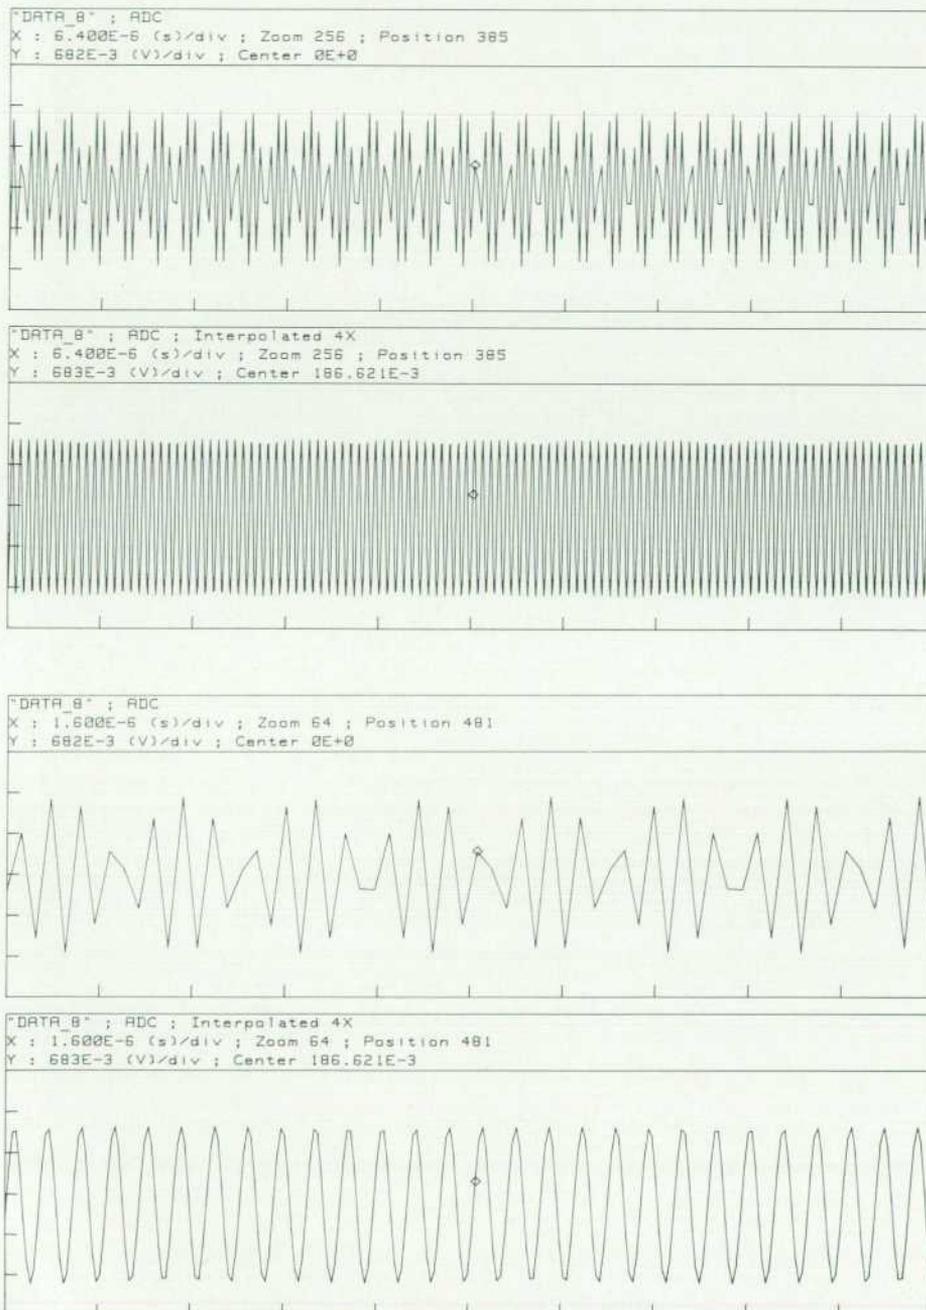


Fig. 6. (a) A sinusoid sampled at 88% of the Nyquist rate without interpolation and with interpolation by four. (b) Same waveform expanded.

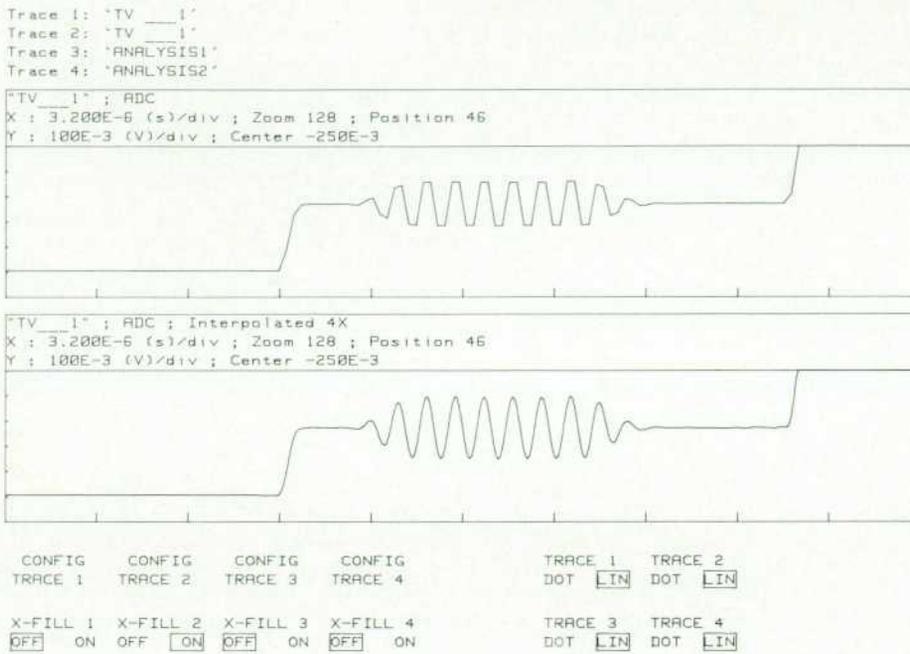


Fig. 7. Video sync portion of a single television scan line sampled at four times the color burst frequency without and with reconstruction.

necessitate sampling very close to the Nyquist limit. Fig. 10 shows a typical disc dropout captured from a flexible disc read head signal. The reconstructed waveform clearly indicates the exact voltage levels and zero crossing intervals.

Conclusion

The maximum sample rate, dynamic performance, and reconstruction strategy of a digitizing oscilloscope determine its bandwidth. The dynamic performance specification determines the sampling accuracy near the Nyquist rate. Waveform reconstruction allows signals sampled near the Nyquist limit to be displayed with accuracy approach-

ing the sampler-imposed limits. The reconstruction algorithm accuracy can be specified as a design parameter. Speed of operation, reconstruction filter impulse response length, and reconstruction accuracy are interrelated. It is easy to design a reconstruction filter that interpolates by four (10 points per cycle at 80% of the Nyquist rate) and is accurate to 0.5% over a band extending from 0 to 80% of the Nyquist rate. Reconstruction is also useful when memory constraints dictate sampling slower than the minimum digitizer rate. A good reconstruction algorithm allows the user to take maximum advantage of the digitizing oscilloscope's memory by allowing sampling near the Nyquist limit.

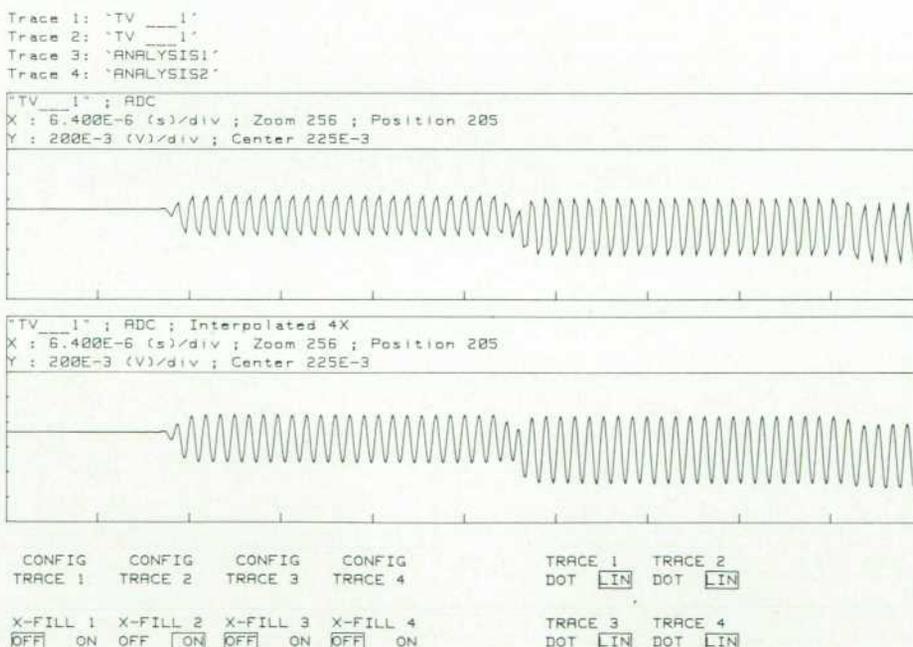


Fig. 8. Phase change in the color burst signal without and with reconstruction.

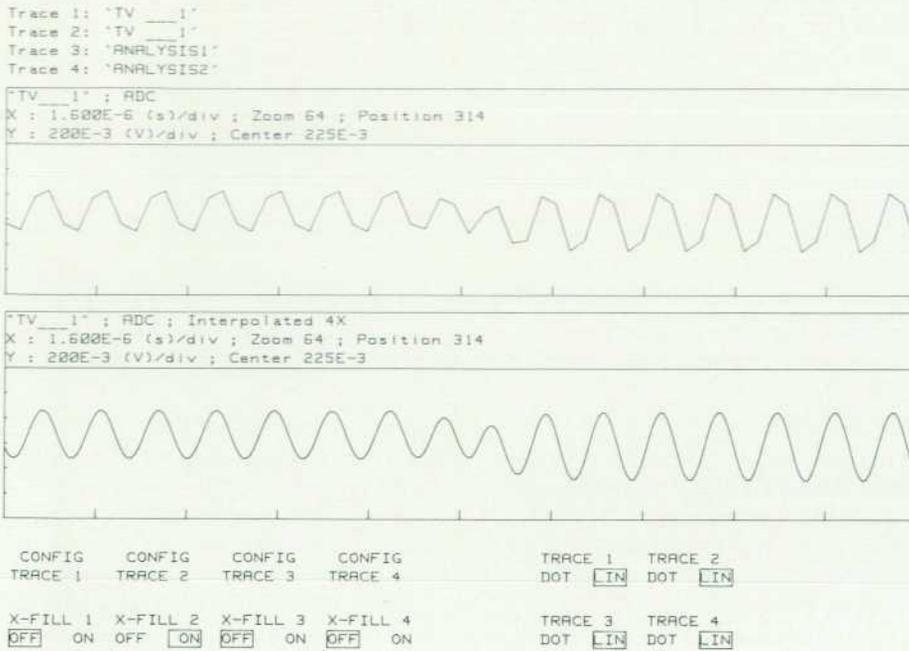


Fig. 9. Detail view of color burst signal.

Reference

1. R.E. Crochiere and L.R. Rabiner, *Multirate Digital Signal Processing*, Prentice-Hall, 1983.

Bibliography

1. G. Oetken, et al, "New Results in the Design of Digital Interpolators," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. ASSP-23, no. 3, June 1975, pp. 301-308.
2. T.A. Ramstad, "Digital Methods for Conversion between Arbitrary Sampling Frequencies," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. ASSP-32, no. 3, June 1984, pp. 577-591.

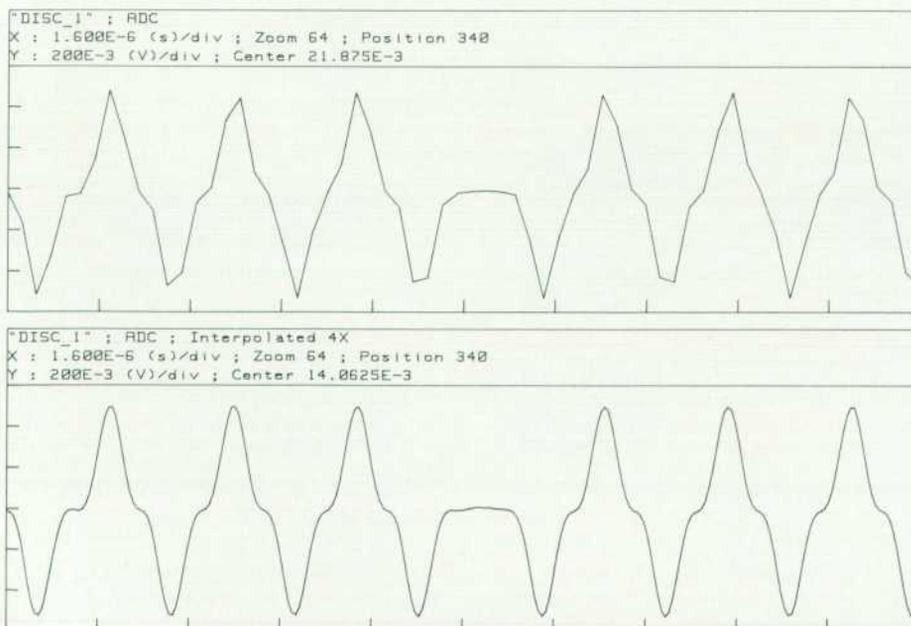


Fig. 10. A typical disc dropout without and with reconstruction.

channel offers ac or dc coupling operation and calibrated offset control. Each channel also has a selectable low-pass filter that reduces the noise floor and provides for anti-aliasing when digitizing at the highest sample rate.

Sampling Control. Several features allow the user more flexible control of the sampling process. The internal oscillator can be phase-locked to an external reference frequency of 1, 2, 5, or 10 MHz applied to the **EXT REF** input. The sampling process can be set to 250 megasamples per second or to $125/N$ megasamples per second, where $N = 1, 2, 3, \dots, 61441$, using the internal time base. Using an external gating signal, the sampling process can be turned on and off in the gated time base mode, or the user can seize control of the sampling process completely by supplying a signal to the **EXT ENCODE** input.

Signal Capture. One of the most useful features of a waveform recorder is the ability to capture a signal both before and after a trigger event. Waveform recorders allow the user to select the time relationship between when a trigger event occurs and the end of the digitized data record. Several triggering modes are available. Positive, negative, positive-negative, bi-trigger, and dropout trigger are analogous to the triggering on the HP 5183A Waveform Recorder (see article, page 6). The trigger level and hysteresis are selectable. A high-impedance external trigger channel with selectable trigger level is also available. The **TRIG OUT** signal available on the rear panel can be used to inform other instruments in a test system that a trigger event has occurred.

Digital System

The purpose of the digital circuitry in the HP 5185A (see Fig. 1) is to store, for each channel, at a rate determined by the user, up to 64K samples from the 250-megasample-per-second data stream created by the quantizer IC on the ADC hybrid until instructed to stop storage at a precisely known time correlated with a trigger event.

Given a steady stream of digitized data from the ADC, the basic measurement is defined by the process shown in Fig. 2. After a minimum number of data points are stored to memory to satisfy the desired pretrigger requirements, the trigger circuitry is enabled. Data continues to be stored while the trigger circuitry examines the analog input to see whether the trigger conditions have been met. Once the trigger event is detected, the post-trigger delay counters ensure that data continues to be stored to memory until the data record is full. At that point the measurement is complete and the record can be retrieved from the memory.

Some applications do not require the signal to be sampled at 250 MHz. A slower sample rate allows signal capture over a longer time period. An obvious approach would have been to operate the ADC hybrid clock at a lower rate. However, adding more circuitry between the SAW (surface-acoustic-wave) oscillator and the ADC hybrid clock input would significantly increase the clock jitter. Instead, the ADC output digital data stream can be undersampled to lower the effective sample rate without adding noise to the measurement. This undersampling of the data occurs at latches in the memory section by means of reduced-frequency memory clocks.

Portions of the digital system required special care in

design and implementation. Storing data at arbitrary rates up to 250 MHz in a useful manner requires careful consideration of each step along the path from the ADC to the memory. Since the speed with which the decisions must be made on a 250-megasample-per-second data stream eliminates the possibility of any microprocessor control, dedicated high-speed digital hardware, together with a special analog-digital high-speed trigger system, maintains precise real-time control during a measurement. Finally, starting and stopping the 250-megasample-per-second measurement in the two channels warrants special attention to ensure synchronous data storage.

Timing Issues

Consider the general data transfer loop shown in Fig. 3. The transmitter serves as the reference point for the timing analysis. Two separate paths, the data path and the clock path, are traversed before signals converge on a receiver. The timing analysis begins by determining the relationship between the transmitter data output and the clock path input. The intermediate element timing analysis includes everything between the transmitter and the receiver, including IC propagation delays and trace delays. The receiver imposes a set of timing requirements on arriving signals, such as setup and hold times at the receiver inputs.

This general model provides a framework for expressing a fundamental requirement of all synchronous digital designs:

The sum of timing uncertainties in a data transfer loop must be less than the data valid time at the transmitter, derated by the receiver setup and hold requirements, to guarantee error-free data transfer under all conditions.

The timing uncertainties in the data transfer loop include variances in transmitter clock-to-data propagation delay and variances in the propagation delay difference between the clock and data paths of all intermediate circuitry. The uncertainties encompass chip-to-chip variations from different IC lots and differential drift over temperature. Error-free data transfers can be achieved if the fundamental requirement is met, as long as the analysis of timing uncertainties does not stop with just the timing information given on the IC vendor's data sheets. Physical constraints usually dictate the need for some additional timing margin. For instance, when the data path is N bits wide, if the physical layout forces compromises such that the trace routing for all N lines is not identical, then the differences in trace

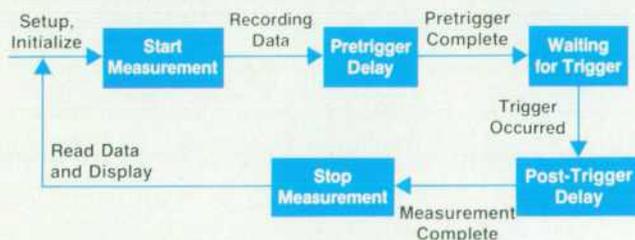


Fig. 2. Basic HP 5185A measurement sequence.

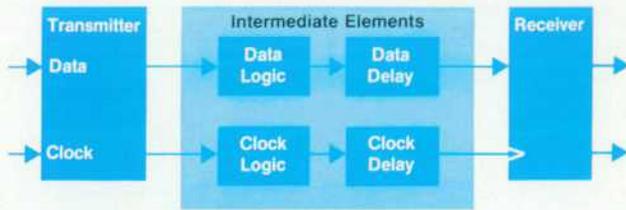


Fig. 3. General data transfer loop.

delays must be accounted for as an additional uncertainty. Similarly, unequal capacitive loading on lines will alter signal propagation delay in ways that are sometimes difficult to model, adding to timing uncertainty.

For low-speed designs, the timing uncertainties associated with presently available IC technologies are much less than the data valid time, so the requirement is easily satisfied. In high-speed designs, however, the significance of this requirement is much more pronounced. The difficulty of data transfer through a 250-MHz system with minimal processing can be illustrated with the following ECL parameters:¹

	Min	Max	
Transmitter clock-to-data delay	0.8	1.8	1.0
Data logic delay	0.4	1.4	1.0
Clock logic delay	0.4	1.4	1.0
Receiver setup time	0.7		
Receiver hold time	0.8		
Receiver setup-hold window			1.5
Required data valid time at transmitter			4.5 ns

The requirement that data be valid for 4.5 ns for even the simplest circuits in a system with data changing every 4 ns demonstrates the futility of attempting to design a system for 250-MHz operation using standard off-the-shelf ICs. Consequently, a custom high-speed deceleration IC was developed to receive the 250-MHz data stream generated by the quantizer IC on the ADC hybrid. Built in a proprietary HP 5-GHz bipolar process,² this IC reduces propagation delay uncertainty and integrates intermediate logic. Designed into this IC is circuitry to decelerate both the clock and data by a factor of two, turning a 250-MHz stream of data at the IC input into two 125-MHz streams of data at the output. The data valid times at the IC outputs approach 8 ns, easing the timing margins for the rest of the digital system.

DC-to-250-MHz Operation

The design of the HP 5185A digital system is complicated by the fact that it operates not only at a 250-megasample-per-second data storage rate but may also operate at lower rates, set either by the internal time base (which allows data storage at rates that are integer submultiples of 125 MHz) or by the user-supplied external encode signal. In fact, the external encode input makes possible full-spectrum sampling from dc to 250 MHz, allowing the user to tailor the sampling process to specific needs. Full-spectrum sampling implies not only sampling at any specific frequency between dc and 250 MHz, but also sampling in the

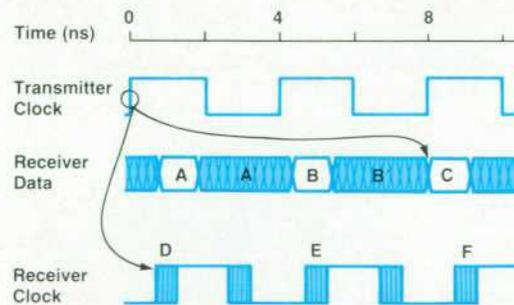
more general case with a signal whose spectrum lies within that range. This allows nonuniform sampling signals, such as linear or logarithmic sweeps or chirps. Satisfying the 4-ns fundamental timing requirement ensures reliable data transfers only at 250 MHz. Full-spectrum operation up to this rate imposes a second timing requirement: the digital system must be insensitive to the system clock period. Timing must be controlled so that clocks and data valid windows always have a fixed time relationship at the receiver input.

The timing diagram in Fig. 4 represents a circuit that is sensitive to the clock period. The transmitter clock at time 0 propagates through the clock path and arrives at the receiver input as the clock edge D. Similarly, assume that data correlated with the time 0 transmitter clock arrives at the receiver as data C. Since the clock edges always fall within the data valid windows, 250-MHz data transfers will be successful, in the sense that the clocks never attempt to capture uncertain data. However, as the clock frequency is reduced below 250 MHz (Fig. 5), the D-clock-to-C-data timing relationship stays fixed, while the clock edge E moves right and data A moves left, reflecting the increasing clock period. Data transfer errors begin to occur as clock E causes the receiver to capture the uncertain data at B'.

If the frequency is reduced further, data B, originally captured by clock E, will be captured by clock D. Although this results in the capture of valid data, the time relationship between the clock and the data has now changed. Since the data is eventually stored into a memory and the clock path also drives the address control for the memory system, the result will be that valid data samples will indeed be stored, but at memory locations that are off by one compared to the 250-MHz case.

Clock period dependent timing problems can be avoided by properly delaying the clock with respect to the data. In the example, either the data trace delay can be reduced or the clock trace delay can be increased, placing the clock edge D within the B data valid window. The functionality of the revised circuit is now insensitive to clock frequency changes.

The clock still has not quite been optimally positioned. Assume that the circuit has just been powered up. If the time 0 transmitter clock is the first clock in the system,



Shaded areas indicate uncertain data.

Fig. 4. Clock and data signals of a transfer loop. The transmitter clock at time 0 causes clock edge D and data C at the receiver.

Printed Circuit Board Transmission Lines

Unlike the HP 5185A, most digital systems today require nothing more than simple connectivity from IC to IC. However, in a 250-MHz system, every printed circuit board trace is as important to the system timing as the ICs performing the logic functions. Operation with wideband logic within these short clock cycles requires that the designer treat the interconnects as transmission lines, paying particular attention to propagation delay, reflections, and crosstalk.

Signals can become distorted by the presence of either reflections or crosstalk along the printed circuit traces. Reflections occur when a fast edge travels along a line and hits a discontinuity in the characteristic impedance of the line or in the termination of the line. Crosstalk occurs when two or more lines are close enough together that a portion of the signal on one line will capacitively couple onto the other lines. Both reflections and crosstalk reduce system noise margin, which can be as low as 130 mV for standard ECL. Since the logic swing of ECL is about 850 mV, a reflection of 15% that happens at the wrong time can wipe out this noise margin. The result can be latching an incorrect logic level or double clocking a flip-flop.

To minimize reflections and crosstalk and to have predictable trace delays, every high-speed interconnect in the HP 5185A is designed as a microstrip transmission line. Multilayer boards with traces on the outer layers and solid ground planes on the adjacent inner layers are used. The characteristic impedance Z_0 of a lossless microstrip transmission line is $\sqrt{L_0/C_0}$, and propagation delay per unit length, T_d , is $\sqrt{L_0C_0}$ where L_0 and C_0 are respectively, the inductance and capacitance per unit of trace length. The tolerances of both Z_0 and T_d depend on physical parameters such as trace width, trace height, distance of the trace from the ground plane, and dielectric constant ϵ_r .^{1,2}

Most high-speed systems are designed with Z_0 equal to 50, 68, or 75 ohms. Even though a 50 Ω system uses more power, it was selected over the others because:

- It is better suited to available board geometries (i.e., the board dimensional tolerances and dielectric uniformity have less of an effect on Z_0 and T_d of the transmission line).
- Higher capacitive coupling to ground means that for a given trace-to-trace spacing, the capacitive coupling from adjacent lines induces less crosstalk.
- Rise and fall times have less degradation for a given capacitive load on the line, C_p , because the time constant Z_0C_p is smaller. When a transmission line is terminated in its characteristic impedance, no reflections occur and the signal is undistorted.

Special termination resistor arrays in SIP packages were developed for the HP 5185A. These not only have a fraction of the parasitic inductance usually associated with termination resistors, but also include a built-in bypass capacitor between the -2V termination supply and ground. Using these components, the reflection coefficient is limited to 6% given the worst-case tolerances, for which Z_0 varies from 44 to 56 ohms.

Ideally, a printed circuit board transmission line termination resistor is located at the very end of the trace, with all gate inputs directly attached to the trace without significant branch length. Failure to do this results in stubs that can look either inductive or capacitive depending upon both the length of stub and the frequency involved. Since the HP 5185A is specified to work from dc to 250 MHz, special care was taken to detect and reduce the number of stubs and other undesired discontinuities by sweeping all high-speed traces over the entire frequency range of interest.

System timing can be compromised by having unpredictable propagation delays in the interconnect lines. All critical line delays were measured to uncover any timing skews resulting from physical layout compromises. This also uncovered effects, such as parasitic loading, that would alter the trace delays from their theoretical values. The controlled nature of these transmission lines allowed them to be used as delay elements in the clock distribution and transfer loop circuitry.

In one situation, the physical layout dictated the need for four small one-quarter-inch stubs on the 250-MHz 50 Ω clock line to the four surface-mounted custom deceleration ICs. The capacitive loading from the stubs and the loading from the inputs caused a nonuniform characteristic impedance along the line. By carefully modeling the capacitance of each stub and load and optimizing the spacing between loads, the capacitive parasitic can be treated as a distributed capacitance C_d that modifies the capacitance per unit length C_0 . The equation for Z_0 then becomes $\sqrt{L_0/(C_0 + C_d)}$. A narrower trace than usual was chosen to raise L_0 and lower C_0 . When combined with the distributed C_d , the effect was to bring the line back to a 50 Ω characteristic impedance, maintaining signal integrity.

References

1. *MECL System Design Handbook*, Fourth Edition, Motorola Inc., 1983.
2. I.J. Bahl and R. Garg, "Simple and Accurate Formulas for a Microstrip with Finite Strip Thickness," *Proceedings of the IEEE*, Vol. 65, no. 11, November 1977, pp. 1611-1612.

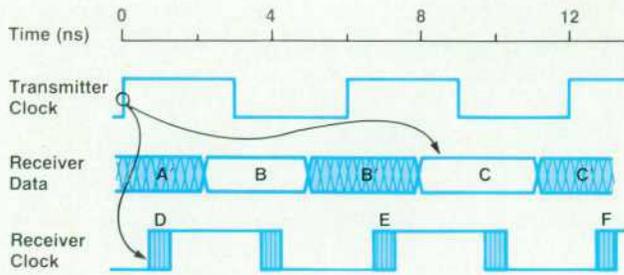
then data C is the first valid data. Any data points before point C represent the power-up characteristics of the circuitry and not true measurement information. If clock D, the first clock seen by the receiver, is allowed to act on data B, invalid data will be propagated throughout the system at power-up. Other initialization circuitry or software intelligence is then required to tell the system to ignore such points.

Proper positioning of the clock with respect to the data requires that the clock be delayed so that clock D captures data C. This type of pipelined system, which relies on the Nth clock edge acting upon the Nth data point, with both clock and data rippling through the system from stage to stage, ensures insensitivity to clock frequency changes and

also prevents the inadvertent storage of invalid data.

Pipeline Partitioning

The use of the custom deceleration ICs allows the remainder of the digital system to run on a 125-MHz clock. This still leaves only an 8-ns cycle time for data storage into memory, trigger detection, time base control, and overall measurement control. Even with high-speed ECL logic, the cumulative timing uncertainty associated with the amount of logic needed to achieve the required functionality and control quickly exceeds the 8-ns data valid time. The solution is to partition the pipeline into smaller segments to reduce the amount of intermediate logic present in each loop. This is done by judicious placement of registers



Shaded areas indicate uncertain data.

Fig. 5. As the frequency of the clock in Fig. 4 is reduced, data transfer errors occur.

throughout the pipeline in such a way that each of the smaller segments still satisfies the fundamental timing requirement.

On occasion, the intermediate logic may contain identical elements in both the data logic and clock logic blocks in Fig. 3. If these two logic elements originate from the same IC package (yielding propagation delay tracking) and if the IC's propagation delay from a low to a high state, T_{plh} , matches the propagation delay from a high to a low state, T_{phl} , then it is possible to eliminate consideration of both logic elements in the cumulative timing analysis of the loop. Element timing then tracks, and more functionality can be designed into this segment of the pipeline. In practice, gate propagation delays will not exactly track on an IC and T_{plh} will not exactly track T_{phl} . However, the error, if any, represents only a fraction of the cumulative timing saved by this technique.

A useful property of ECL logic is the ability to tie outputs directly together in a wired-OR connection.³ Provided that the outputs are close together on the board to minimize reflections, this feature makes it possible to add more functionality without extra logic gates. This technique, used extensively in the trigger detection circuitry, eliminates the timing uncertainty associated with physical gates, and thereby requires fewer partition stages.

Data Deceleration

Data deceleration is used to store the 250-MHz data stream from the ADC into a 64K-word memory. The deceleration scheme splits the 250-MHz data stream into 16 data streams, each at a reduced rate, 250/16 MHz, allowing the use of medium-speed static RAM.

Deceleration occurs in three stages. The first and most difficult stage, using the custom deceleration IC, takes the 250-MHz data stream from the ADC and splits it into two 125-MHz data streams. The custom deceleration IC data path is three bits wide. Four such ICs are required to handle the 12-bit-wide ADC output, which consists of eight data bits, two trigger bits, and underflow and overflow bits.

A one-bit representation of the first deceleration stage is shown in Fig. 6. A divider circuit uses the 250-MHz clock from the quantizer IC on the ADC hybrid to generate a 125-MHz clock. The rising edge of this clock latches one bit of the digitized word from the ADC into a flip-flop, FFa. Four nanoseconds later, the falling edge of the 125-MHz

clock latches the data from FFa into another flip-flop, FFb. At the same time, this falling edge latches a bit from the next digitized word into a third flip-flop, FFc. Since this process occurs in parallel with the other 11 bits of the ADC output, the result is two phases of data aligned in time, phase 1 and phase 2, each at 125 MHz and each 12 bits wide. A buffered version of the 125-MHz clock is also output, allowing the close timing relationship between the clock and data to be used by the succeeding data transfer loop.

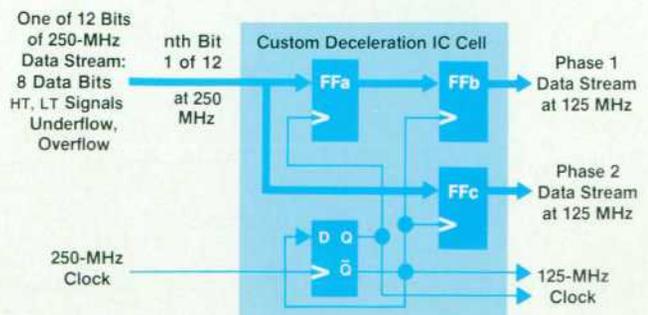
The second deceleration stage, operating only on the 8-bit digitized data for subsequent memory storage, partitions each of the decelerated phases by another factor of two. Off-the-shelf high-speed ECL ICs are used for this process, with special care in the design and implementation to accommodate the 8-ns cycle time.

This deceleration process can be continued for two more divide-by-two stages to provide the net division by 16 needed. However, each deceleration stage requires twice as many registers as the preceding stage, so the divide-by-eight and divide-by-16 stages would need 24 (= 8 + 16) additional 8-bit registers. Instead, another custom HP IC provides a direct factor-of-four deceleration and a built-in ECL-to-TTL translation to reduce the IC count dramatically. The combined deceleration multiplies the data cycle time to 64 ns (4 ns × 16), enough time to satisfy medium-speed static RAM requirements.

Analog-Digital Trigger

A versatile yet reliable trigger circuit is a key factor in any waveform recorder design. The HP 5180A/T and HP 5183A/T described elsewhere in this issue feature a digital triggering technique. Each sample of the digitized input signal is compared with programmed hysteresis and trigger levels using digital N-bit comparators. By comparison, typical analog trigger schemes are prone to errors caused by gain or offset drift differences between the trigger path and the measurement path. The asynchronous nature of an analog trigger also makes time correlation between the trigger event and the captured data samples uncertain.

Implementing a purely digital trigger for the HP 5185A with off-the-shelf technology is unrealistic. Propagation delays of the 8-bit digital comparators and trigger decoders required to do the job far exceed the 4-ns cycle time. A digital trigger using the two phases of 125-megasample-per-second data from the deceleration ICs would require four 8-bit comparators. These would increase the fanout of each



line on the high-speed data bus and would either compromise the bus routing and loading or require an additional buffer IC in an already tight situation.

The 250-MHz trigger implementation in the HP 5185A (Fig. 7) is both an analog and a digital trigger. The quantizer IC contains two additional comparators connected to the analog input line. The comparator reference levels are set according to the programmed trigger level and hysteresis. The accuracy of these dc levels, the intimate connection of the trigger comparators to the analog input line, and the excellent tracking of the trigger comparators with the comparators used in the analog-to-digital conversion process ensure that little error will exist. Latching of these threshold comparator outputs, called HT (high trigger) and LT (low trigger), occurs at the same time the analog input signal is sampled by the quantizer.

The HT and LT signals are split, along with the data, into two 125-megasample-per-second phases using the custom deceleration ICs. The decelerated trigger signals then serve as inputs to dedicated digital state machines which implement the various trigger modes. Each state machine contains a static control signal which selects the state machine corresponding to the trigger mode chosen. A high-speed control signal enables the selected state machine only after the pretrigger delay has been satisfied.

Each state machine contains two main sections, the hysteresis decoder/flip-flop and the trigger decoder/flip-flop. The hysteresis decoder scans the phase 1 and phase 2 signals for a hysteresis event. In the positive trigger mode, for example, a hysteresis event happens the first time the input waveform is sampled below the lower threshold setting. Once a hysteresis event is detected, the hysteresis flip-flop is latched and the trigger decoder is enabled. The trigger decoder then scans the phase 1 and phase 2 signals for a

trigger. The trigger decoder can also bypass the hysteresis flip-flop. If the phase 1 and phase 2 signals meet the trigger and hysteresis conditions simultaneously, then the trigger decoder causes a trigger to occur immediately.

Once a trigger is detected, the trigger flip-flop output is fed back to the trigger decoder to freeze the flip-flop in its current state. Information present at the instant of the trigger event is needed to decode which of the two phases of data caused the trigger.

Two-Channel Synchronous Data Storage

Within a single measurement channel of the HP 5185A, the clock, data, and trigger information are pipelined together throughout the entire high-speed digital system to preserve critical signal timing relationships, thereby reducing cumulative timing uncertainty. When two such pipelined measurement channels are combined to form a synchronous two-channel system, the unavoidable buildup of timing uncertainty between channels leads to the need for some special clocking techniques to ensure synchronous starting and stopping of a two-channel measurement.

Synchronous startup of the 250-MHz clock for both channels ensures synchronous startup of the measurement, but involves more than simply turning on an oscillator or gating a free-running clock. The inherent problem with these methods is the possibility of generating partial or incomplete clock pulses upon startup. If these "runt" pulses are interpreted differently by the two channels, the channels will lose synchronization.

The de-teeter circuit (Fig. 8) uses custom high-speed flip-flops to deliver identical full-width pulses to both channels. Any metastable condition appearing at the first flip-flop output as a result of the trailing edge of the asynchronous RECORD control line has an extremely short duration.

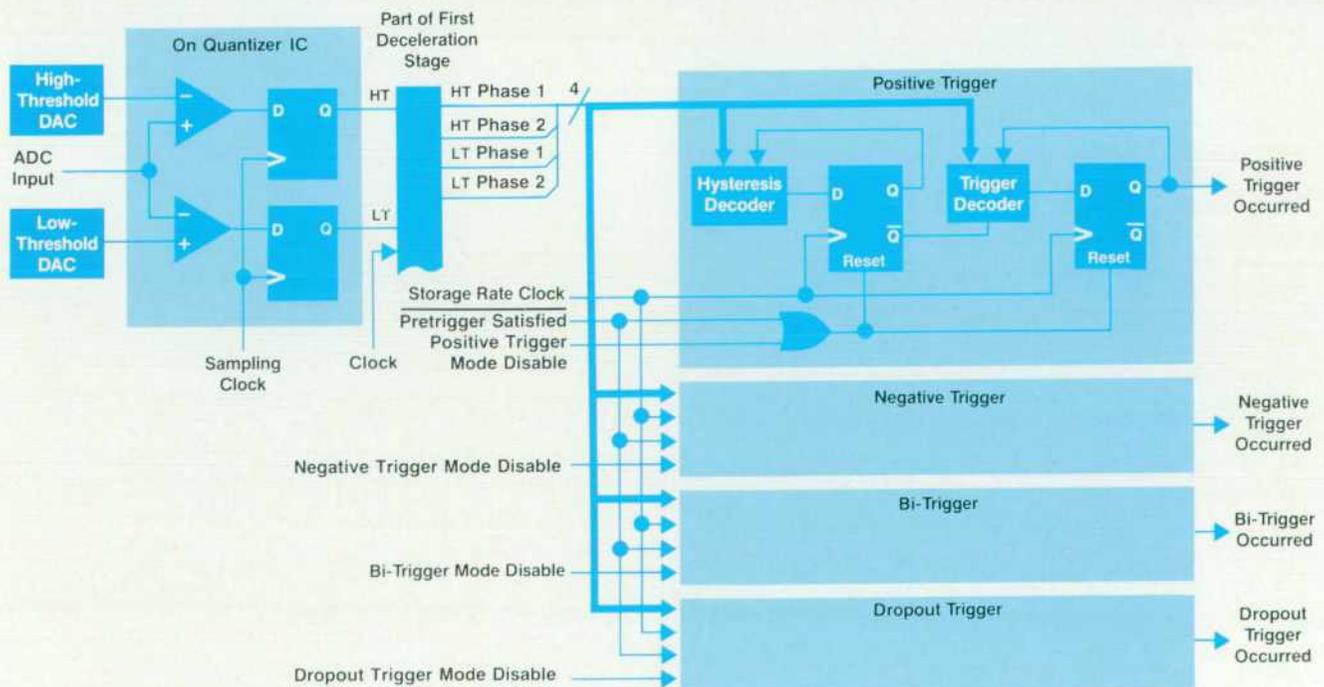


Fig. 7. Analog-digital trigger circuit.

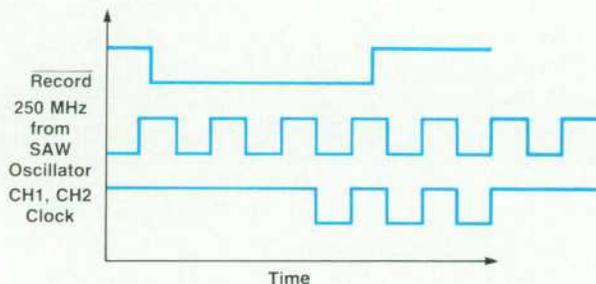
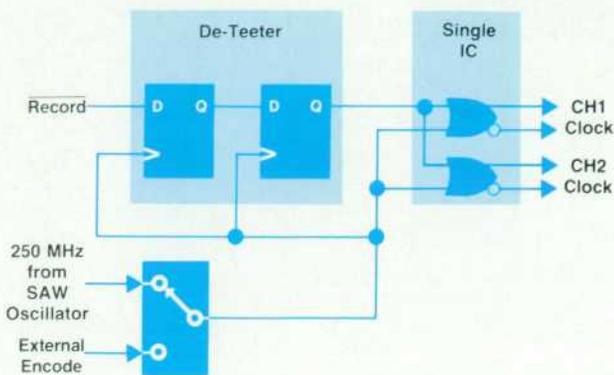


Fig. 8. Synchronous startup of the clocks in the two HP 5185A channels is guaranteed by a de-teeter circuit at the output of the 250-MHz SAW oscillator.

Specifications for the flip-flop guarantee that the propagation delay of the first flip-flop, the probable metastable duration, and the setup time of the second flip-flop sum to less than 1.8 ns, well within the 4-ns clock period. This ensures that the trailing edge of the second flip-flop output will be synchronous with the leading edge of the next clock.

The short propagation delay of the second flip-flop also ensures that the first clock pulse will have full width. Identical clock timing is provided to both channels by means of the OR gates, identically connected, which originate from the same IC package.

In light of the number of logic elements appearing within the clock paths of each channel, attempting to stop the dual-channel measurement synchronously within the 8-ns digital system clock period is a futile effort. Instead, a clock decelerated to one-sixteenth the sample rate forms the basis for the valid measurement stop intervals. Synchronization of the decelerated clock with the 125-MHz digital system clock within each channel maintains coherence within the channel.

Acknowledgments

The authors are by no means solely responsible for the design of the HP 5185A Waveform Recorder digital system. Mark Kotfila designed early versions of the memory and HP-IB circuitry. Custom integrated circuitry for data deceleration was designed by Brian Hill, Randy Tsang, and Gary Jacobsen. Gary Jacobsen was also responsible for valuable analysis and consultation on the rest of the digital system. Mike Detro, Bill Daley, Ron Keeley, and Jim Ammon all contributed to the thermal and mechanical design. Test engineering was performed by Molly Myers, Lisa Craig, and John Schmitz. Lisa Craig and Janet Anvick were invaluable in getting the HP 5185A into production. Keith Ferguson managed the project during the early design phases.

References

1. *MECL Device Data*, Motorola Inc., 1985.
2. I.V. Pecenco and A.S. Wang, "A High-Performance, Bipolar Integrated Circuit Process," *Hewlett-Packard Journal*, Vol. 33, no. 9, September 1982.
3. *MECL System Design Handbook*, Fourth Edition, Motorola Inc., 1983.

Waveform Recorder Design for Dynamic Performance

Quantizer on-chip delays and delay mismatches, low-order distortion in the input amplifier, and low levels of phase noise in the sampling clock can have significant effects on performance.

by Brian J. Frohring, Bruce E. Peetz, Mark A. Unkrich, and Steven C. Bird

IN THE PAST, fast, wideband waveform recorders have not kept pace with their lower-speed cousins in measurement fidelity. Therefore, the primary design objective of the HP 5185A Waveform Recorder was superior dynamic performance, requiring careful design of all subsystems affecting measurement performance. This article describes the key considerations involved in the design of the major analog components. We also discuss the most useful parameter for verifying dynamic performance, effective bits.

The essential elements of the analog-to-digital measurement system for one channel are shown in Fig. 1. The HP 5185A provides two identical digitizing channels driven by the common 250-MHz SAW (surface-acoustic-wave) reference oscillator. Each channel includes several custom integrated circuits designed to provide the required measurement performance and feature set.

The 8-bit quantizer IC is the foundation of the measurement system. Minimizing on-chip delays and delay mismatches was a major consideration in designing this device for optimal dynamic performance. Matching requirements were satisfied through the choice of converter architecture, cell design optimization, and chip layout. One of the performance characteristics resulting from this design effort is the absence of missing codes for full-scale input signals through 100 MHz. The relationship of on-chip matching to performance and the mechanisms that create harmonic distortion will be covered later in this article. Harmonic distortion resulting from the nonlinear nature of the quantizer input capacitance is also analyzed.

The attenuator, preamp hybrid, and op amp IC condition the analog input signal, providing continuously variable gain across all input ranges for optimum signal-to-noise performance. A selectable low-pass filter is useful to prevent aliasing of the input signal and high-frequency noise while digitizing at the maximum rate. The general design requirements for these analog signal conditioning components are discussed and an analysis of low-order distortion in the differential input stage of the preamp IC is presented.

Phase noise in the sampling clock can degrade dynamic performance. We therefore describe the relation between various types and sources of phase noise and the effects of this noise on amplitude errors when sampling high-slew-rate signals. We then discuss considerations related to the choice of oscillator topology, circuit design, and compo-

nent selection to achieve the required performance in the SAW reference oscillator, that is, less than 2 ps rms phase jitter.

Finally, we review how dynamic performance is measured and outline the closed-form sine wave curve fitting algorithm used in production testing of the HP 5185A.

Analog-to-Digital Conversion

The analog-to-digital conversion function is the most critical in the signal path of any waveform recorder. Realizing the dynamic performance objectives for the HP 5185A Waveform Recorder required careful minimization of the sources of error.

Many analog-to-digital converter architectures require the use of a separate sample-and-hold circuit to sample the input waveform for subsequent conversion to a digital value. In the HP 5185A, both the analog sampling and the analog-to-digital conversion functions are provided by a single monolithic integrated circuit, the quantizer. This circuit performs a new conversion every 4 ns.

ADC Hybrid. A custom bipolar operational amplifier IC functions as a fixed-gain amplifier to drive the 25-pF capacitive load presented by the quantizer input. It converts the differential output of the preamp hybrid to a single-ended signal at the quantizer input, realizing a gain of 6.7 with a bandwidth greater than 220 MHz. The custom operational amplifier and quantizer are mounted in a thin-film hybrid package which is described in the article on page 49.

Chip Architecture. The block diagram of the quantizer is shown in Fig. 2. This is the architecture of a classic flash analog-to-digital converter, having one comparator for each quantization threshold to be encoded. The buffer/latches, 256 for the 8-bit conversion, sample the analog input signal and output a "thermometer" code representation of the

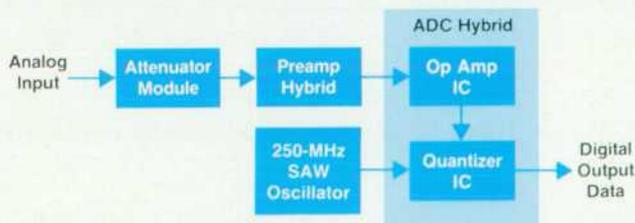


Fig. 1. Analog-to-digital measurement system (one channel).

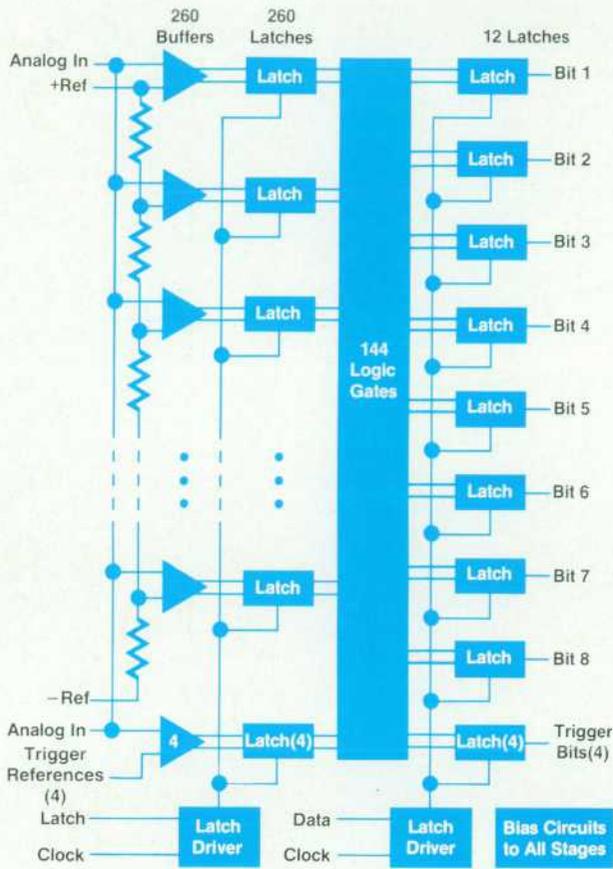


Fig. 2. 8-bit quantizer IC block diagram.

value of the input signal at the sampling instant. The decoding logic converts this rather inefficient (256-bit) digital code to an 8-bit binary code. Although circuit efficiencies could be achieved by shifting some of the encoding function ahead of the sampling comparator stage,¹ this more conventional approach equalizes the delay paths for all encoded states. Analog decoding before the first stage latches usually requires nonidentical cells with the resultant risk of adding code dependent delays to the input signal before sampling. To achieve a nearly square chip layout, the IC is divided into two 7-bit converter subsections.

The four extra comparators (the bottom ones in Fig. 2) represent a slight departure from the conventional flash architecture. These are identical to the sampling comparators and provide input signal threshold detection matched in time to the sampling process accomplished by the converter's main comparator bank. Two of these provide analog-digital trigger detection with hysteresis, and the other two are used for analog underflow and overflow detection.

Design Considerations

The key design considerations for maximizing the dynamic performance of the quantizer were:

- Layout of the input comparators and associated interconnect for optimum delay and load matching
- Selection of an error-tolerant encoding scheme providing

matched loading on the input comparators

- Comparator design to minimize the sampling latch regeneration time constant, thus minimizing the probability of internal ambiguous logic levels which may result in output bit errors.

Variable Comparator Delay. When a flash analog-to-digital converter is integrated, the propagation delay of two identical gates on the same circuit may differ. If the propagation delay variance among comparators is large, errors can occur when the analog input is sampled. The signal corruption becomes significant when the delay difference between two comparators becomes so great that, for high-slew-rate signals, two consecutive comparators switch in the wrong order, resulting in an invalid state of the thermometer code at the output of the sampling comparators. For a full-scale 125-MHz sinusoidal input converted to 8 bits, this mismatch amounts to about 10 ps.

To reduce the variability of the comparator delays, the buffer and latch transistors in the signal path are biased at currents below that of the f_T peak. At low collector currents, f_T is dependent on transit time, thermal voltage, and collector current (Fig. 3), while at higher collector currents, the Kirk effect² limits f_T . Biasing the transistors at the f_T peak would make f_T , and hence the delay, dependent on both transit time and Kirk effect. Although the reduced bias approach makes the average propagation delay greater, it reduces the comparator delay variance by eliminating the device matching requirements for the Kirk effect.

Encoding. To reduce further the effect of comparator delay mismatch, a compact encoding scheme was chosen that produces the most benign results for likely invalid states. Based on simulations of several techniques for binary and nonbinary encoding, the optimal approach produces an 8-bit binary code at the quantizer output.

As mentioned previously, the 8-bit converter is divided into two 7-bit subsections. The first-level encoding for each subsection converts the 128-bit thermometer code generated by the bank of comparators into a 32-bit intermediate code.

The first-level encoding scheme significantly reduces metal interconnect by allowing the comparator cells to be

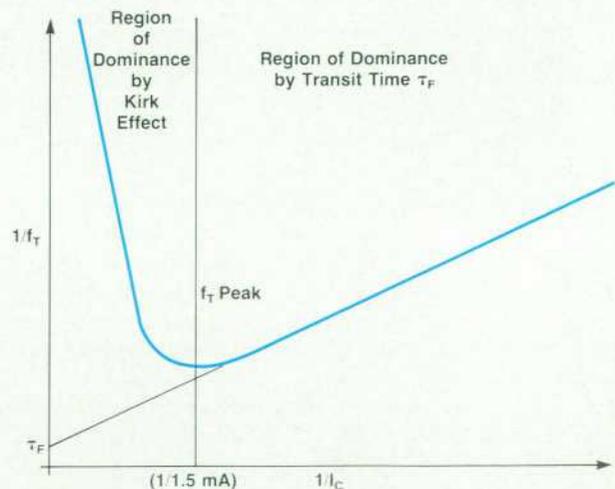


Fig. 3. Transistor f_T as a function of collector current.

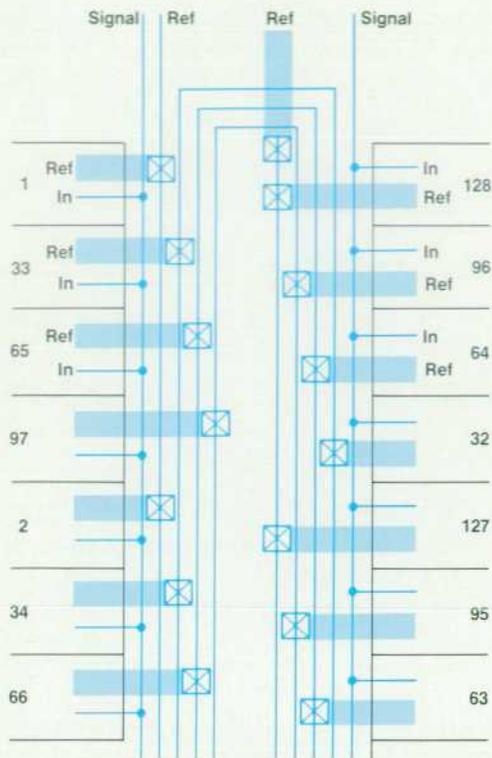


Fig. 4. Comparators are interleaved by four in the quantizer IC input network.

adjacent to the first-level logic cells, thus saving die area. In addition, the scheme maintains equal loading on the comparators to avoid mismatched propagation delay. The encoding for a 7-bit subsection is given by the following Boolean logic relation:

$$B[n] = (C[n]C[n + 32]) + (C[n + 64]C[n + 96])$$

where $B[n]$ is the n th bit of the partially encoded word, $C[x]$ is the output of the x th comparator, and n ranges from 1 through 32. This scheme allows comparators to be interleaved by four on the chip layout as shown in Fig. 4, yielding a benefit of longer metal runs on the resistor ladder to improve matching of the one-ohm-per-tap values. For the final conversion, the encoding algorithm that proved most tolerant to invalid states was implemented.

Regeneration Time Constant. For any binary encoding technique, two or more output bits may change for a single input threshold transition. If the input signal level is close to the first-stage sampling latch balance point at such a threshold transition, the comparator output may not have regenerated to a valid logic level when the output latch stage is latched. If the encoding logic for any two of the affected output bits interprets this level differently, an erroneous output code can result. Given a uniformly distributed input, the likelihood P of an error can be derived:³

$$P = \frac{V_a}{Aq} e^{-T/\tau} \quad (1)$$

where V_a is the ambiguous voltage range of the comparator output logic level, A is the analog gain of the comparator, q is the voltage span between two adjacent input comparator thresholds, T is the time available to latch, which is constrained by the time between the input comparator and output latch clocks, and τ is the comparator positive feedback (regeneration) time constant.

Since a binary scheme was chosen to minimize the variable comparator delay problem, the comparator (Fig. 5) is designed to minimize P in equation 1 above. In a typical latch, $R5$ and $R6$ are effectively zero-valued and only $R3$ and $R4$ can be varied. In such a configuration, larger values of $R3$ and $R4$ increase the positive feedback loop gain but have the undesirable effect of increasing latch recovery time, the time required for the comparator to go from a latched state to tracking the input with no memory of its previous state. Splitting the latch load resistors is key to achieving an optimal trade-off among latch gain, regeneration time constant τ , and latch recovery time. The sum of resistors $R3$ and $R5$ (identical to $R4$ and $R6$, respectively) is made large to increase the positive feedback loop gain, thus decreasing τ . An independent choice of $R3$ optimizes latch recovery time. For a fixed sampling period, reducing the latch recovery time increases the time available to latch, T .

Systematic Comparator Delay Mismatch. On a chip large enough to house 256 comparators, the mismatch in the propagation of the input and latch signals can generate significant amounts of error. These signals are routed on separate lumped RC lines. Because of the RC characteristic and the short length of the line, no attempt was made to terminate the line in its characteristic impedance. Model-

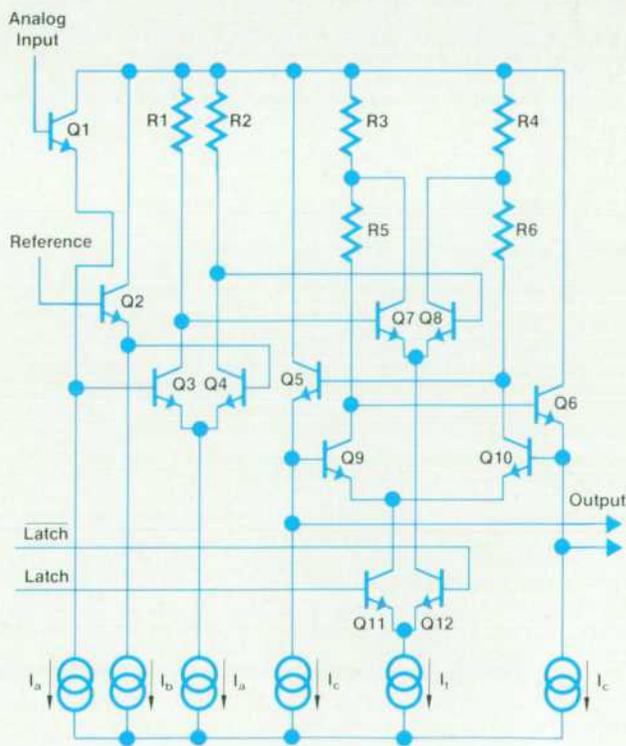


Fig. 5. Quantizer IC comparator cell schematic.

ing the lumped open-circuit line as a distributed system, the velocity of propagation for a line that is short compared to the wavelength of the applied frequency is:⁴

$$\text{velocity} = \frac{-s}{y\gamma^2} \quad (2)$$

where s is complex frequency, y is line position ($y = 0$ at the end of the line), and γ is the line characteristic. The line characteristic is dominated by distributed metal resistance, distributed metal capacitance, and lumped transistor capacitance. Inductance is negligible in the frequency range of interest on IC metal traces for the purposes of propagation delay. Distributing the transistor capacitance, the line characteristic is given by:

$$\gamma = \sqrt{RCs} \quad (3)$$

Substituting equation 3 into equation 2 gives the velocity:

$$\text{velocity} = -\frac{1}{RCy} \quad (4)$$

By integrating the inverse of equation 4 from some point y_1 on the line to the end of the line, the propagation delay is found:

$$\text{delay} = \frac{RCy_1^2}{2}$$

Optimal matching of RC between the input line and the latch line is achieved by running both lines on the same layer of metal and matching the transistors connected to the line. Nonideal matching would be seen as eight parabolic lumps over the input range of the converter. The number eight arises from the interleaving of comparators by four (Fig. 4) in each of the two 7-bit subsections. Beat

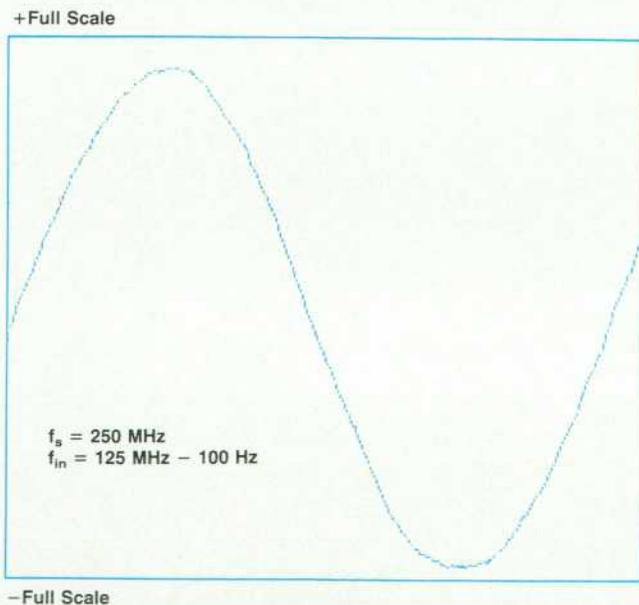


Fig. 6. Beat frequency plot for a signal at 125 MHz - 100 Hz sampled at 250 MHz.

frequency testing⁵ can be used to measure this effect.* The plot in Fig. 6 shows such a test under the most severe conditions—sampling at the maximum rate with an input frequency near half the sampling rate. The data has been decimated by plotting every other point to realize an effective sampling rate of 125 MHz. As expected, the nonlinear delay shows up as parabolic distortion, which can just be perceived riding on the rising edge of the sine wave. This amounts to mismatch on the order of 10 ps over the entire line. Analytically, this kind of nonlinearity can be treated as frequency modulation by Bessel functions,⁶ giving energy in the 21st, 23rd, and 25th harmonics for a full-scale input.

Distortion Sources. As can be seen in Fig. 7, both second and third harmonic distortion increase at high input frequencies. Second-order distortion HD_2 is a result of the nonlinear base-collector capacitance C_{jc} of Q1 in the input latch (Fig. 5) and the finite source resistance driving C_{jc} . This source resistance results from the op amp output impedance and series resistance in the op amp/quantizer interconnect network. Analysis using Volterra series gives the result:

$$HD_2 = \frac{V_0\omega C_{j1}R}{2\sqrt{1 + (2\omega C_{j0}R)^2}}$$

where V_0 is the magnitude of the fundamental, ω is the input frequency, C_{j1} is the linear voltage coefficient, that is, the coefficient of the V^1 term in a power series expansion of $C_{jc}(V)$, C_{j0} is the constant term in the power series expansion of $C_{jc}(V)$, and R is the source resistance. Since the source is inductive as well as resistive, R could be replaced with a more complex expression that takes the inductance of the source into account.

Third-order distortion arises from multiple sources. In addition to the nonlinear RC mechanism just discussed, the latching of high-slew-rate signals by finite-rise-time latch signals can generate third-order distortion. This latter effect can be understood by referring to the simple latch

*Beat frequency testing is a technique of undersampling a high-frequency input sine wave at a rate slightly offset from the input frequency. The input frequency is mixed to a lower frequency where it can be viewed in detail.

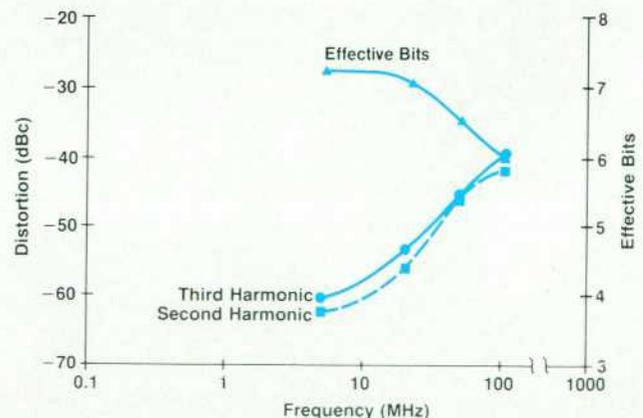


Fig. 7. Distortion and effective bits versus frequency for the analog-to-digital converter of the HP 5185A Waveform Recorder.

of Fig. 8. If the latch command is in the track mode, all of the current passes through the lefthand differential pair and the output follows the input. As the latch command is slowly changed, some current starts to flow in the righthand pair. The input-output transfer function starts exhibiting hysteresis because of the positive feedback of the righthand pair. As more current is switched, the amount of hysteresis increases until the circuit is fully latched. If the latch command is switched at full speed, the circuit still passes through the period of increasing hysteresis, although quickly.

A quasistatic analysis of the two upper pairs in Fig. 8 results in the hysteresis boundary:

$$V_i = 2V_t \tanh^{-1}[(I_1/I_1) - (I_2/I_1) \tanh(V_i/2V_t)] \quad (5)$$

where V_i is the hysteresis boundary as viewed from the input, V_t is the thermal voltage, I_1 is the total comparator stage current, I_1 is the sum of the emitter currents in the righthand pair, and I_2 is the sum of the emitter currents in the lefthand pair, and V_i is the logic level, $I_1 R_L$. Assuming that the currents I_1 and I_2 switch uniformly with time, equation 5 is plotted as a function of time in Fig. 9. Also plotted in Fig. 9 are lines representing two signals, one quickly varying, the other slowly. As long as a signal stays within the area labeled "latched region" after initially entering this area, the signal will be latched correctly. But if a signal enters the latched region and then leaves it before being latched, it will be latched incorrectly.

To see how errors arise, first consider a slow-moving signal. Since it enters the hysteresis boundary from below and later is latched low, we can think of the sample instant as occurring near the onset of the hysteresis, before point A in Fig. 9. Now consider the fast-moving signal. It also enters the hysteresis boundary from below but passes above the hysteresis boundary before being latched, resulting in a high value. In this case, the sampling instant is effectively after point B in Fig. 9, later than for the slow-moving signal. Since this speed dependent delay occurs equally for both rising and falling inputs, it results in odd-harmonic distortion.

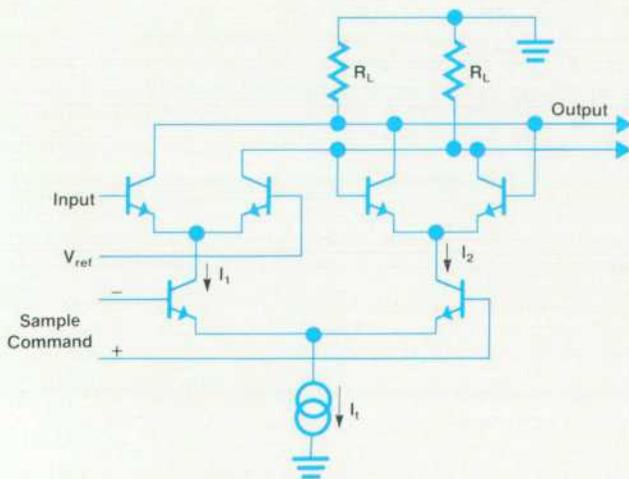


Fig. 8. Simple latch.

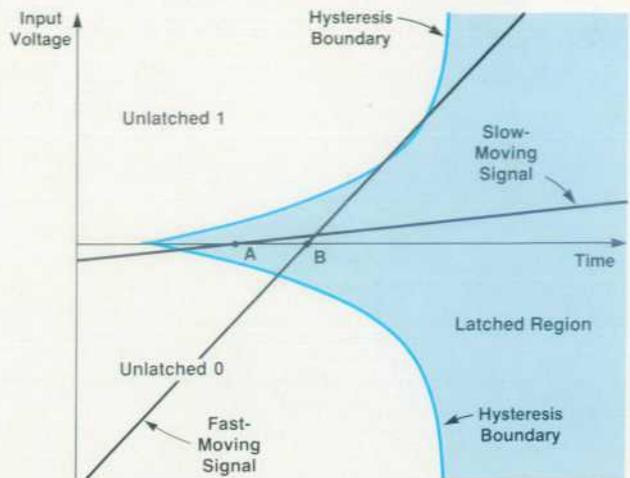


Fig. 9. Hysteresis boundary versus time for the latch of Fig. 8.

Assuming that signals are constant-slew ramps over the input range near the comparator threshold, and that tail currents switch uniformly with time, the projected error boundary based on equation 5 is :

$$V_{err} = V_i(t_1) - (t_1 - t_0) \frac{dV_i}{dt}$$

where V_i is given by equation 5, t_1 is the time that describes the states of currents I_1 and I_2 when a constant-slew signal is tangent to the hysteresis boundary, and t_0 is the time that describes the states of currents I_1 and I_2 at the onset of hysteresis.

This hysteresis-based source of distortion can be reduced by using lower values of R_1 through R_4 (Fig. 5), which reduces the slew rate seen by the sampling latch circuitry. The split load scheme of the latch portion is particularly helpful in keeping R_3 and R_4 low.

Distortion in the Signal Conditioning Path

Having taken great care to realize unsurpassed dynamic performance in the 8-bit quantizer, we could not allow the analog signal conditioning path to compromise this performance. Components in the analog chain are therefore designed with aggressive goals for low distortion, low noise, and wide bandwidth. The input stage of the preamp provides a useful vehicle for discussion of several design issues related to distortion.

A high signal-to-noise ratio is desirable to minimize the rms noise during a measurement. Increasing the signal levels to maximize the signal-to-noise ratio must be balanced with the errors introduced by the resulting increased levels of distortion. Second-order harmonic distortion is proportional to the input signal level and third-order distortion is proportional to the square of the input signal level.

Preamp Hybrid. The preamp (Fig. 10) is based on a pair of custom bipolar integrated circuits produced in HP's 5-GHz- f_T high-frequency bipolar process. The input cascode transconductance of the amplifier is selectable for three values in a 1-2-4 sequence by effectively switching the

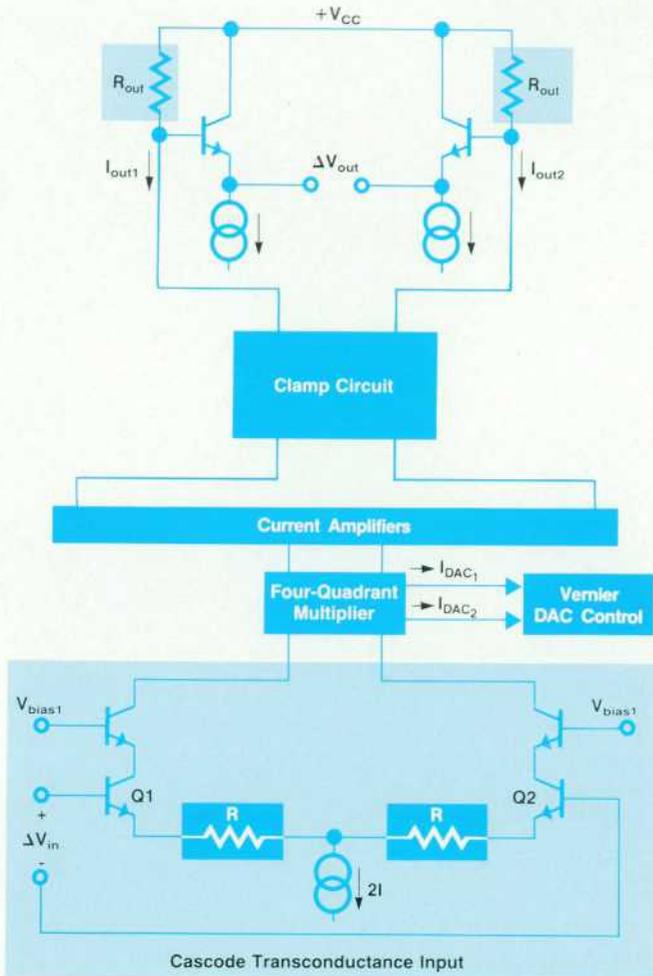


Fig. 10. Simplified input preamplifier schematic diagram.

value of the emitter degeneration resistance R . With one of the inputs incrementally grounded, the single-ended-to-differential conversion is performed. Examining distortion from this stage, the simplified transfer function for the emitter degenerated differential pair of transistors $Q1$ and $Q2$ can be derived from:

$$\Delta v = 2V_t \tanh^{-1} \left(\frac{\Delta i}{2I} \right) + R\Delta i$$

where Δv is the differential base voltage $V_{bQ1} - V_{bQ2}$, V_t is the thermal voltage, Δi is the differential collector current $I_{cQ1} - I_{cQ2}$, R is the emitter degeneration resistance, and I is the emitter and collector bias current ignoring β effects.

Using a Taylor series expansion and collecting the terms,

$$\Delta v = \left(\frac{V_t}{I} + R \right) \Delta i + \frac{V_t}{12I^3} \Delta i^3 + \frac{V_t}{80I^5} \Delta i^5 + \dots \quad (7)$$

Inverting this equation gives the transfer function:

$$\Delta i = \frac{I}{V_t + RI} \Delta v - \frac{V_t I}{12(V_t + RI)^4} \Delta v^3 + \dots \quad (8)$$

The output differential current is of the form

$$\Delta i = a_1 \Delta v + a_2 \Delta v^2 + a_3 \Delta v^3 + \dots \quad (9)$$

where

$$a_1 = \frac{I}{V_t + RI}$$

$$a_2 = 0$$

$$a_3 = \frac{-V_t I}{12(V_t + RI)^4}$$

Assuming a sine wave input, $\Delta v = A \sin(\omega t)$, no offset, and low distortion levels such that second and third-order harmonic distortion are dominated by the second and third-order nonlinearities, respectively, the distortion terms are:

$$HD_2 = \left| \frac{a_2}{2a_1} \right| A = 0 \quad (10)$$

$$HD_3 = \left| \frac{a_3}{4a_1} \right| A^2 = \frac{V_t}{48(V_t + RI)^3} A^2 \quad (11)$$

For the maximally stressed gain range and operating conditions of $I = 8 \text{ mA}$, $R = 17 \Omega$, $V_t = 26 \text{ mV}$ at 25°C , and maximum input level $A = 50 \text{ mV}$, the dominant distortion term is $HD_3 = -70 \text{ dB}$. In practice, HD_2 is generated by internal mismatches and offsets in the input signal resulting in distortion terms up to the level of HD_3 , but these have been ignored in the analysis.

SAW Oscillator

An important source of error in any high-speed analog-to-digital converter is the phase noise, or cycle-to-cycle jitter, of the oscillator that provides the master sampling clock. Time jitter of the clock alters the time between successive samples in a random fashion. This time jitter can be directly related to amplitude errors at the output of the ADC. If the jitter of the oscillator is too high, it contributes to sampling errors (amplitude noise) for points sampled on the high-slew-rate portions of input signals. As a practical design goal, the jitter of the oscillator should be low enough to produce no more than $1/2q$ of amplitude error under the worst-case conditions, where q is the average voltage span between adjacent comparator thresholds.

The worst-case phase noise requirements for the master reference oscillator of the HP 5185A are strict. Assuming that the Nyquist sampling theorem is not being violated, the worst-case conditions are:

- Sampling at sine wave zero crossings
- Maximum sample rate of 250 megasamples per second
- Input frequency of 125 MHz
- Full-scale signal at the input
- Maximum memory record length of 64K words.

Under these conditions the jitter requirement corresponding to amplitude noise of no more than $1/2 \text{ LSB}$ is 2 ps rms , which is computed from the following equation:

$$\sigma_x(\tau) < \frac{1}{2^{N+1} \pi f_{in} (2.5)}$$

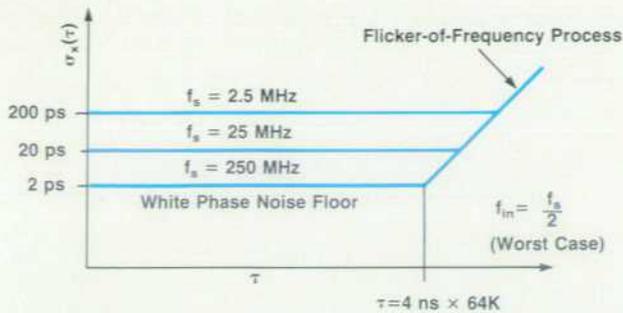


Fig. 11. Time-domain phase noise requirement for the HP 5185A.

where f_{in} is the input frequency, N is the number of bits, and 2.5 is taken to be a peak-to-rms conversion for a Gaussian distribution.⁷ This jitter requirement must be held for the duration of a worst-case measurement, which is four nanoseconds per sample times 64K samples per measurement, or approximately one quarter of a millisecond. The white phase noise floor, shown in Fig. 11, defines the jitter performance requirement under worst-case conditions. This plot shows rms jitter $\sigma_x(\tau)$ versus time τ , where τ is the total measurement duration. Note that the worst-case requirement is relaxed for lower input frequencies and sample rates.

The time-domain plot of Fig. 11 can be transformed into a frequency-domain plot of the phase noise requirements.⁸ Fig. 12 is such a plot showing the SSB phase noise spectral power density (in dBc/Hz) as a function of offset from the carrier (in Hz). Fig. 12 shows that the 2-ps noise floor requirement derived in the time domain transforms into a noise floor of 148 dB/Hz below the carrier in the frequency domain.

The requirements can be directly compared with the measured phase noise performance of the HP 5185A oscillator, also shown in Fig. 12. The overall system performance is dependent upon the integral of the phase noise margin over the full bandwidth of the system. The cumulative margin is large compared to the cumulative loss of the corner violation shown, so the corner violation has little effect on the overall system performance. (The narrowband spikes in the 100-kHz-to-40-MHz range are artifacts of the measurement system.)

SAW Oscillator Phase-Locked Loop. Fig. 13 shows a block diagram of the surface-acoustic-wave (SAW) phase-locked loop (PLL). The heart of this circuit is the voltage-controlled oscillator (VCO). The SAW VCO provides the sampling clock to the ADC in each channel.

The major components of the VCO circuit use custom technology. The crystal is a SAW device developed for this product at Hewlett-Packard's Santa Rosa Technology Center.⁹ The SAW device was selected because of its excellent phase noise characteristics. It is a fundamental-mode device, which can be modeled as an LC series-resonant circuit, allowing a simpler design than most RF oscillators.

The active device chosen for the SAW VCO is a differential amplifier made from custom emitter-coupled logic fabricated at the Santa Clara Technology Center. This differential amplifier was selected for its low phase noise floor and

because it can deliver fast edge speeds ($t_{20-80} \approx 350$ ps). Maintaining fast edge speeds throughout the sampling clock chain minimizes the phase noise contribution of each stage in the chain. To preserve the quality of the VCO output, the encode multiplexer, synchronizer, and output drivers are also implemented in custom emitter-coupled logic.

The phase-locked loop operates in the same fashion as most conventional phase-locked loops. Phase-locking allows the VCO to take on the superior aging and temperature stability of the reference (internal or external) while retaining its own short-term stability. Phase-locking to an external reference also allows the user to stabilize the phase relationship between the sample clock and the input waveform, provided that the input waveform is coherent with the reference.

Design Considerations. The noise process that dominates the SAW phase-locked loop jitter performance is the white phase noise process. This process is a result of the noise floor of the differential amplifier used in the SAW VCO. Buffer stages between the VCO and the ADC sampling clock input also degrade the noise floor of the sample clock. This degradation can be calculated using a root-sum-squares formula. For example, adding a second stage with a noise floor identical to the VCO noise floor will degrade the noise floor of the sample clock by 3 dB. A third stage will degrade the noise floor by an additional 1.77 dB, and so on. Therefore, minimizing the number of buffer stages is important.

The differential amplifier, like all amplifiers, has $1/f$ noise around dc and has some inherent nonlinearities. In an oscillator, these nonlinearities mix the $1/f$ noise onto the carrier. The SAW resonator, modeled as a second-order bandpass filter, produces a $1/f^2$ process. Mixing $1/f$ with $1/f^2$ produces the $1/f^3$ phase noise process shown in Fig. 12. Consideration must be given to both the amplifier linearity and the Q of the crystal since together they determine the corner frequency between the white phase noise floor process and the $1/f^3$ process. This corner frequency becomes more important for longer measurement durations.

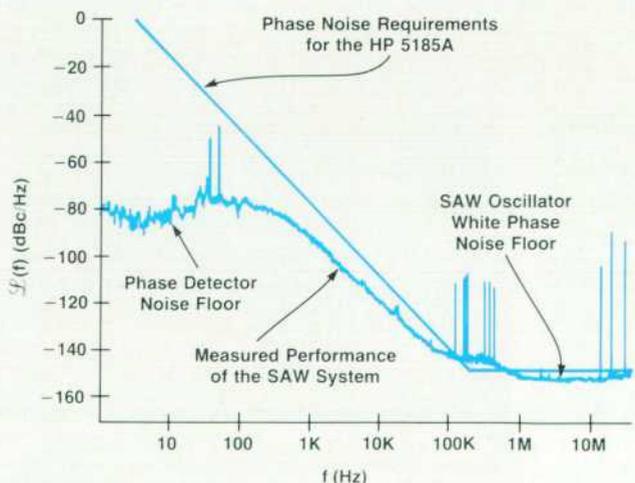


Fig. 12. Frequency-domain phase noise requirement and actual performance.

There is another noise floor closer to the carrier—the noise floor of the phase detector. This noise floor and the $1/f^3$ process meet at an offset from the carrier of approximately 100 Hz. This offset provides the optimal bandwidth for the low-pass filter that follows the phase detector. If the bandwidth is too wide, the $1/f^3$ performance of the SAW resonator will be degraded. If it is too narrow, the phase detector performance will be degraded.

The pulling range of a VCO is the amount by which the VCO frequency can be shifted from the nominal frequency by varying the input control voltage. The pulling range must be large enough to compensate for the aging rate of the SAW resonator, the temperature dependence of the SAW device's parameters, and center frequency tolerances in the SAW manufacturing process. However, when the pulling range is increased, it is at the expense of the overall Q of the circuit. This increases the phase noise in the $1/f^3$ region and pushes the corner frequency outward. The pulling range of the SAW oscillator is maximized through the use of back-to-back, low-capacitance varactor diodes in series with the SAW resonator in the oscillator feedback loop.

Making Effective Bits Measurements

Effective bits is a figure of merit used to quantify the contributions of all forms of distortion and noise added by the measurement process: quantization error (differential nonlinearities, missing codes, and aperture uncertainty), integral nonlinearity (harmonic distortion), and system noise (additive noise and spurious nonharmonic spectral components). The method of measuring effective bits described in previous publications^{5,10} and now generally (but not universally) used for specifying the performance of digitizing instruments uses sine wave curve fitting. The test technique involves making a measurement of a single-frequency sinusoidal input, estimating the actual input by fitting a sine wave to the data using a least-mean-squared-

error algorithm, then calculating the residual rms error in the measured data relative to the fit sinusoid. Effective bits is then defined as follows:

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

where N is the number of bits of resolution of the digitizer being tested, E_{actual} is the residual rms error in the measured data after subtracting the fit sine wave, and E_{ideal} is the expected value of the rms error an ideal N-bit quantizer adds to a uniformly distributed input signal, $q/\sqrt{12}$, where q is the voltage span between adjacent comparator thresholds.

Since the effective bits parameter is sensitive to all forms of distortion and noise, it is important that care be taken to ensure that the test system used to evaluate the performance of an instrument under test makes no significant contribution to the detected rms error.

The basic test setup required to measure the effective bits performance of a given instrument is shown in Fig. 14. The bandpass filter is present to reduce harmonic distortion components from the generator output to a negligible level and to ensure that no signal-generator-produced or coupled low-frequency noise is present at the input to the instrument under test. The major sources of potential errors in the measurement of effective bits are signal generator phase noise and amplitude errors (inaccuracy of the signal generator output, losses in the bandpass filter, and cable losses).

Signal generator phase noise represents a deviation from the ideal sine wave assumed to be present at the input. The nearly uniform sampling of the instrument under test translates signal generator phase noise into amplitude errors, which are most pronounced near zero crossings. These errors are readily detected in effective bits measurements.

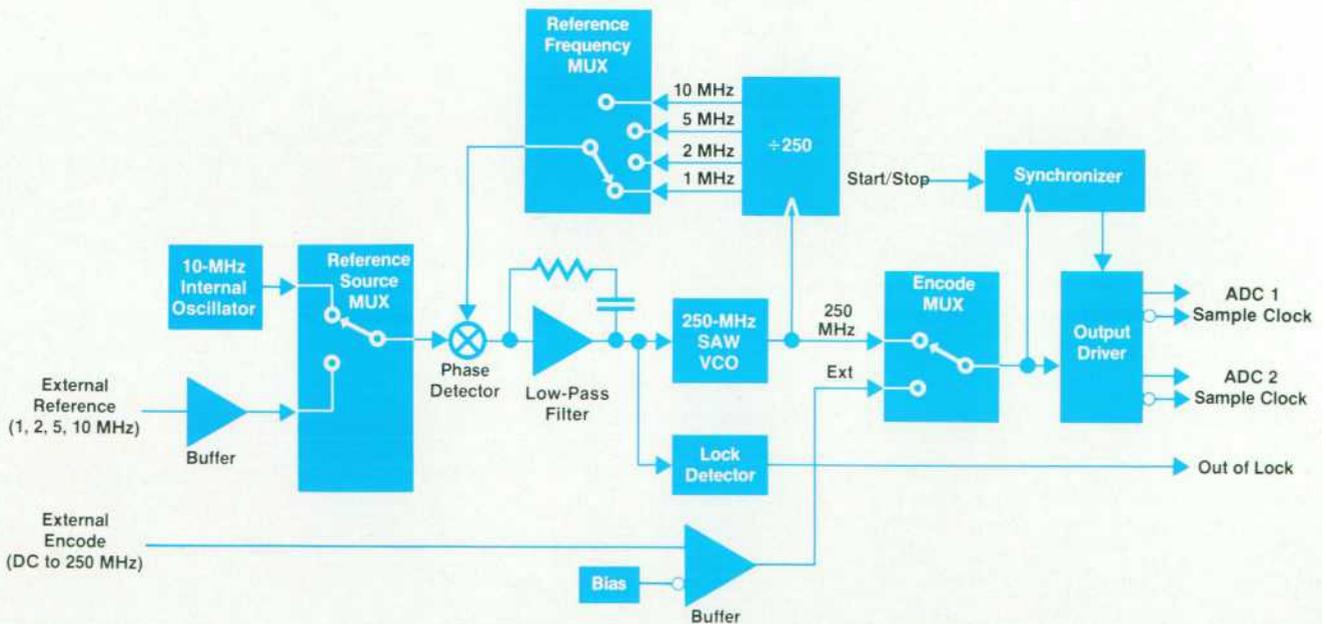


Fig. 13. 250-MHz SAW oscillator phase-locked loop system.

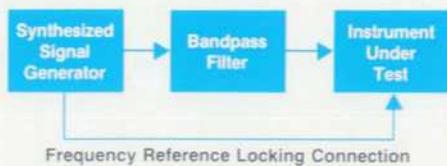


Fig. 14. Effective bits test setup.

An HP 8662A Synthesized Signal Generator is used in all HP 5185A test systems at the factory because of its excellent phase noise characteristics.

Since the effective bits parameter measures distortion as well as noise, testing will be sensitive to variations in the input amplitude when measurement distortion is not negligible. This generally occurs at higher input frequencies. The absolute accuracy of the synthesized signal generator, specified at ± 1 dB, plus attenuation of the fundamental by the bandpass filter, require amplitude compensation of the generator output to provide the proper amplitude at the input of the instrument under test. In production testing, these losses are corrected by identifying the amplitude settings that result in the proper signal amplitude at the instrument under test. This is accomplished through an automated calibration procedure using an HP 438A Power Meter with HP 8482A Power Sensors. Amplitude errors of a few dB are reduced to well under ± 0.5 dB through this calibration process.

Acknowledgments

The authors are greatly indebted to the following people who contributed to the performance of the HP 5185A Waveform Recorder. The input attenuator module was designed by Jim Johnson. Jeremy Sommer took time away from his switching power supply design to improve the dynamic performance of the integrated blocks. The opera-

tional amplifier in the ADC hybrid was designed by James Kang, who also contributed to the design of the quantizer. Most of the quantizer circuitry was designed by Brian Hamilton. John Schmitz refined the critical oscillator-to-quantizer timing. The assistance and technology of the Santa Clara Technology Center are greatly appreciated. The fixed-frequency sine wave curve fit method was derived by Ron Potter.

References

1. J. Corcoran, K. Knudsen, P. Clark, and D. Hiller, "A 400 MHz 6b ADC," *ISSCC Digest of Technical Papers*, IEEE Cat. No. 84CH1971-1, February 1984, pp. 294-295.
2. C.T. Kirk, Jr., "A Theory of Transistor Cutoff Frequency (f_T) Fallout at High Current Densities," *IRE Transactions on Electron Devices*, Vol. ED-9, 1962, p. 164.
3. B. Zojer, R. Peschacher, and W. Luschnig, "A 6 bit/200 MHz Full Nyquist A/D Converter," *IEEE Journal of Solid-State Circuits*, Vol. SC-20, no. 3, June 1985, pp. 780-786.
4. B. Peetz, B. D. Hamilton, and J. Kang, "An 8-bit 250 Megasample per Second Analog-to-Digital Converter: Operation Without a Sample and Hold," *IEEE Journal of Solid-State Circuits*, Vol. SC-21, no. 6, December 1986.
5. *Dynamic Performance Testing of A-to-D Converters*, Hewlett-Packard Product Note 5180A-2, 1982.
6. M. Schwartz, *Information Transmission, Modulation, and Noise*, McGraw-Hill, 1959, Chapter 3.
7. S.C. Bird and J.A. Folchi, "Time Base Requirements for a Waveform Recorder," *Hewlett-Packard Journal*, Vol. 33, no. 11, November 1982, pp. 29-34.
8. M.C. Fischer, "Frequency Stability Measurement Procedures," *Eighth Annual Precise Time Interval Applications and Planning Meeting*, December 1976.
9. P.S. Cross and S.S. Elliot, "Surface-Acoustic-Wave Resonators," *Hewlett-Packard Journal*, Vol. 32, no. 12, December 1981.
10. B.E. Peetz, A.S. Muto, and J.M. Neil, "Measuring Waveform Recorder Performance," *Hewlett-Packard Journal*, Vol. 33, no. 11, November 1982.

Fixed-Frequency Sine Wave Curve Fit

Estimation of the rms error contributed by the quantization process is the purpose of sine wave curve fitting. A sinusoid fit to the input data using a least-mean-squares algorithm is assumed to represent the actual input signal presented to the digitizer being tested. As discussed in the accompanying article, the effective bits parameter is based on the rms noise power present in the error signal—the difference between the digitized data and the fit sinusoid—over the period of the measurement record.

Any fitting algorithm fits a sine wave of the following form:

$$V_{in}(n) = V_o + V_a \cos(2\pi f_{in} n T_{sample} + \phi)$$

A generalized fitting algorithm requires fitting the parameters V_o , V_a , f_{in} , and ϕ in the above equation. Such fitting algorithms are iterative and therefore time-consuming and prone to undesirable convergence behavior because of the nonlinear relation between parameters.

Connecting the **EXT REF** input of the instrument under test to the signal generator reference output phase-locks the sampling process to the input signal. The frequency of the input sinusoid is thus known to sufficient accuracy to enable use of a simplified sinusoidal curve fitting algorithm.

To derive the simplified algorithm, fit a sinusoid of the following form to the input data:

$$x_n = A \cos(\omega t_n) + B \sin(\omega t_n) + C$$

where ω is $2\pi f_{in}$, f_{in} being the known input frequency, and the t_n are the sample times.

Given a data record y_n of M samples of the input sinusoid measured at times t_n (in practice, uniformly separated by the sampling interval, T_{sample}), the total residual error ϵ of the measured data relative to the fit sine wave is given by:

$$\epsilon = \sum_{k=1}^M (y_k - x_k)^2 = \sum_{k=1}^M [y_k - A \cos(\omega t_k) - B \sin(\omega t_k) - C]^2$$

Setting the partial derivatives with respect to the parameters being fit to zero gives:

$$0 = \frac{\partial \epsilon}{\partial A} = -2 \sum_{k=1}^M [y_k - A \cos(\omega t_k) - B \sin(\omega t_k) - C] \cos(\omega t_k)$$

$$0 = \frac{\partial \epsilon}{\partial B} = -2 \sum_{k=1}^M [y_k - A \cos(\omega t_k) - B \sin(\omega t_k) - C] \sin(\omega t_k)$$

$$0 = \frac{\partial \epsilon}{\partial C} = -2 \sum_{k=1}^M [y_k - A \cos(\omega t_k) - B \sin(\omega t_k) - C]$$

Defining $\alpha_k = \cos(\omega t_k)$ and $\beta_k = \sin(\omega t_k)$ gives

$$\sum_{k=1}^M y_k \alpha_k = A \sum_{k=1}^M \alpha_k^2 + B \sum_{k=1}^M \alpha_k \beta_k + C \sum_{k=1}^M \alpha_k$$

$$\sum_{k=1}^M y_k \beta_k = A \sum_{k=1}^M \alpha_k \beta_k + B \sum_{k=1}^M \beta_k^2 + C \sum_{k=1}^M \beta_k$$

$$\sum_{k=1}^M y_k = A \sum_{k=1}^M \alpha_k + B \sum_{k=1}^M \beta_k + CM$$

The fit parameters are thus given by the solution to the linear equation

$$\mathbf{Y} = \mathbf{U}\mathbf{X}$$

which is

$$\mathbf{X} = \mathbf{U}^{-1}\mathbf{Y}$$

where:

$$\mathbf{X} = \begin{bmatrix} A \\ B \\ C \end{bmatrix}$$

$$\mathbf{Y} = \begin{bmatrix} \sum_{k=1}^M y_k \alpha_k \\ \sum_{k=1}^M y_k \beta_k \\ \sum_{k=1}^M y_k \end{bmatrix}$$

$$\mathbf{U} = \begin{bmatrix} \sum_{k=1}^M \alpha_k^2 & \sum_{k=1}^M \alpha_k \beta_k & \sum_{k=1}^M \alpha_k \\ \sum_{k=1}^M \alpha_k \beta_k & \sum_{k=1}^M \beta_k^2 & \sum_{k=1}^M \beta_k \\ \sum_{k=1}^M \alpha_k & \sum_{k=1}^M \beta_k & M \end{bmatrix}$$

Adding one more sum,

$$\epsilon_0 = \sum_{k=1}^M y_k^2$$

The total residual error is given by:

$$\begin{aligned} \epsilon = \epsilon_0 - 2AY_1 - 2BY_2 - 2CY_3 \\ + A^2U_{11} + 2ABU_{12} + 2ACU_{13} \\ + B^2U_{22} + 2BCU_{23} + MC^2 \end{aligned}$$

and the rms error is then:

$$E_{actual} = \sqrt{\frac{\epsilon}{M}}$$

Using this fixed-frequency curve fit has reduced test computation time by a significant factor over the previously used variable-frequency curve fit.¹⁰ Fixed-frequency sine wave curve fitting is a closed-form, noniterative solution ensuring convergence. The fixed-frequency curve fit thus provides a faster, more reliable measure of effective bits performance.

Packaging a High-Performance 250-Megasample-per-Second Analog-to-Digital Converter

by Patrick D. Deane, Simcoe Walmsley, Jr., and Farid Dibachi

CUSTOM INTEGRATED CIRCUITS in the analog signal path make the superior measurement performance of the HP 5185A Waveform Recorder possible. As integrated circuits deliver higher levels of performance, particularly at high frequencies, greater care must be taken to ensure that this performance is not squandered by component packaging. The fundamental objective of the packaging scheme used in the HP 5185A was to achieve maximum performance from the custom chip set, taking into account not only electrical performance but also reliability, producibility, and serviceability.

Previous high-performance analog-to-digital converter (ADC) measurement systems have used either thick-film or thin-film hybrid technology. Several modifications of standard thin-film hybrid technology were developed to achieve the performance goals of the HP 5185A.

Packaging Requirements and System Objectives

The objective of the HP 5185A is to provide an 8-bit, 250-megasample-per-second analog-to-digital conversion system with superior distortion, noise, and pulse response characteristics.

Basic initial requirements imposed on the ADC packaging scheme were:

- Analog signal paths of the circuit must have frequency response significantly better than the system bandwidth of 125 MHz.
- Digital and clock paths of the circuit must accommodate 250-MHz signals with transition times as low as 500 ps and pulse widths of 1.3 ns.
- Integrated circuit junction temperatures must be less than 125°C at the maximum instrument ambient temperature of 55°C.
- Circuits must be removable from the large-area printed circuit board that supplies I/O paths, control, and power.

A block diagram of the circuitry on the ADC hybrid is shown in Fig. 1. As integrated circuits were designed to meet system objectives, each imposed its own requirements on the circuit packaging. The quantizer IC, the heart of the HP 5185A, was designed for maximum dynamic performance. Implemented as a flash converter, it consumes 14 watts. Its high current requirements (for example, one supply draws 2A) made it necessary to feed power to all four corners of the IC. The part is very large for a 5-GHz- f_T chip, 0.220 inch square. The subnanosecond edges on input clocks and data outputs call for careful consideration of transmission line termination and crosstalk issues.

An operational amplifier IC with a 2-GHz gain-bandwidth product was specifically designed to drive the

quantizer IC's large input capacitance of 25 pF. To achieve the required performance, the op amp requires external pnp transistors and small-area precision resistors. Its output is physically close to the quantizer input to minimize high-frequency distortion (see article, page 39).

Requirements for external pnp transistors and tight-tolerance resistors naturally led to a choice of either thick-film or thin-film hybrid technology. Aside from their cost advantages, thick-film circuits combine good resistors with the possibility of conductive vias to a backside ground plane and trace crossovers. Thick-film conductors also have lower series resistance than thin-film conductors. Thin-film technology offers resistors with better tolerance and drift specifications. Thin-film resistors also require less substrate area, which lowers parasitics and allows tighter packing of the circuits. The ideal would be a combination process offering the resistors and small features of thin-film technology along with the trace crossovers, conductive vias to substrate backside, and high-conductivity traces of thick film.

Substrate Technology

The substrate process developed for the ADC hybrid of the HP 5185A is a modification of thin-film technology that combines the advantages of thin-film resistors with conductive vias to a backside metallization pattern. This backside metal pattern can be used both as a large ground plane and as a cross-under for top-layer crossing traces.

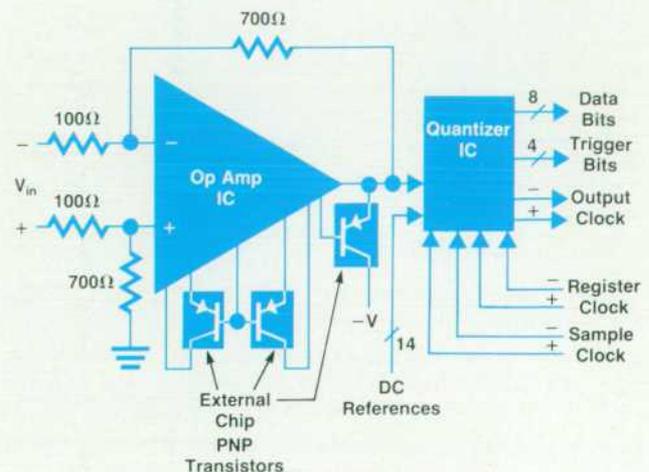


Fig. 1. ADC hybrid block diagram. The electrical and thermal requirements of the hybrid packaging scheme revolve around the quantizer IC.

The pattern plating process also allows unusually thick gold for a thin-film process.

Substrate processing begins with an unmetallized ceramic substrate. Top and bottom surfaces of the substrate are treated with materials to prevent laser slag from accumulating. A laser is then used to cut holes in the substrate where the vias will be located.

Creation of vias in thin-film technology has much stricter requirements than in thick-film technology. Sputtering thin metal films onto a substrate requires a smooth substrate surface. For sputtered metal to stick to the laser-drilled holes, the inside surfaces of the holes must be smooth and free of cracks. To achieve the desired surface, a cleaning and refining step is performed after laser drilling.

Thin films of Ta₂N resistor material, a TiW adhesion layer, and gold are in turn sputtered onto the substrate, taking special care to ensure that the films are sputtered inside the holes.

Negative photoresist is applied with a roller coater and then exposed through a mask with the conductor pattern on it. After the unexposed resist is washed away, the remaining resist is used as a plating mask. Metal areas not covered with developed photoresist are plated up from an initial gold thickness of a few hundred nanometers to 5 to 10 μm thickness. After plating, the resist is removed and a quick gold etch removes the gold from areas that are not conductors. This plating process produces a thin-film trace with very high conductivity and excellent line definition (Fig. 2).

Thin-film resistors are then formed using the resistor mask. After oven stabilization, these resistors are trimmed to their final values using a pulsed current technique.



Fig. 2. 420 \times magnified photograph of spiral inductor traces on the preamp hybrid. The traces are 50 μm wide on a 100- μm pitch and are 10 μm thick. Pattern plating gives the advantages of thin-film technology with a factor of 20 improvement in trace conductivity.

Thermal Performance

In addition to state-of-the-art substrate technology, a different approach had to be taken in the thermal design of the package. Standard chip-on-substrate hybrid technology cannot accommodate the 14W power dissipation of the quantizer.

The reliability of an electrical component is an inverse function of operating temperature. Since the quantizer IC dissipates 14W of power and is 0.220 inch square, it dominates the requirements for the thermal performance of the package. To maintain IC reliability, the design goal was to keep temperatures on the ICs below 125°C at a 55°C ambient. For a 14W IC, this requires that the junction-to-air thermal impedance be less than 5°C/W. By comparison, standard ECL flatpacks or ceramic chip carriers have thermal impedances of 50 to 100°C/W, depending on air flow. Standard thick-film or thin-film hybrid packaging techniques are also inadequate, since they typically involve epoxy die attach to a ceramic substrate in intimate contact with a large-area heat sink. The quantizer die temperature, when mounted to a 2-inch-square alumina substrate, can be approximated using a simple one-dimensional model (Table I), assuming that either nonconductive epoxy (Case #1) or conductive epoxy (Case #2) is used for quantizer die attach, and assuming that the rest of the circuitry on the hybrid dissipates 2W.

Table I

ADC Thermal Analysis Standard Hybrid Packaging

	Thermal Z °(C/watt)	Power (watts)	Temperature Rise Case #1	Case #2
Nonconductive epoxy	4.8	14	67°C	
Conductive epoxy	1.0	14		14°C
Alumina	0.9	14	13°C	13°C
Alumina/heat sink	4.2	14	59°C	59°C
Heat sink/air	2.0	16	32°C	32°C
Junction temperature rise above ambient:			171°C	118°C
Requirement:			<70°C	

The table shows that the elimination of nonconductive die attach material is a necessary but not sufficient condition to meet the requirement that the IC operates at a temperature no greater than 70°C above ambient. Given an optimum heat sink design and the use of conductive epoxy, the only remaining thermal impedances that can be improved upon are those of the alumina substrate and the interface between the substrate and the heat sink.

The thermal performance of both the substrate and the interface to the heat sink is a direct result of the mediocre heat conductivity of alumina (Table II). Substitution of a metal for the alumina in the thermal path has the obvious advantage of an order-of-magnitude better thermal conductivity, which lowers the heat rise across the bulk material. More important, the improved heat spreading angle that results leads to a much larger effective area at the interface with the heat sink. A similar scheme of direct attachment of the silicon die to a low thermal impedance via a metal spreader is used in several commercially available VLSI pin-grid array packages.

Table II

Thermal Conductivities of Selected Materials¹

Material	Thermal Conductivity (watts/cm-°C)
Alumina	0.276
Aluminum	2.17
Beryllia	1.969
Copper	3.937
Diamond	6.299
Glass	0.008
Silicon (doped)	0.984
Tungsten	1.969

The structure ultimately chosen for the HP 5185A ADC hybrid (Fig. 3) achieves optimum thermal performance while preserving the advantages of thin-film substrate technology. A thin-film alumina substrate, having a laser drilled hole slightly larger than the quantizer die, is processed to form the hybrid resistor and conductor patterns. This substrate is attached to a metal plate using nonconductive epoxy. The quantizer IC is attached through the hole in the substrate directly to the metal plate using a conductive epoxy. The metal plate acts as a heat spreader in intimate thermal contact with a very good heat sink to complete the thermal path.

Many alternatives were evaluated during the packaging investigation. Simple one-dimensional thermal analysis had its place in the design process; it eliminated techniques having no chance of success. However, reality takes the form of a three-dimensional hybrid with several heat sources and is much more difficult to analyze. Real heat transfer problems require a solution to differential equations for which closed-form solutions do not exist. By simplifying the boundary conditions and by using Green's function, it was possible to develop a program to solve the three-dimensional heat transfer problem for a multilayer structure using a convergent infinite series.^{2,3} Fig. 4 shows

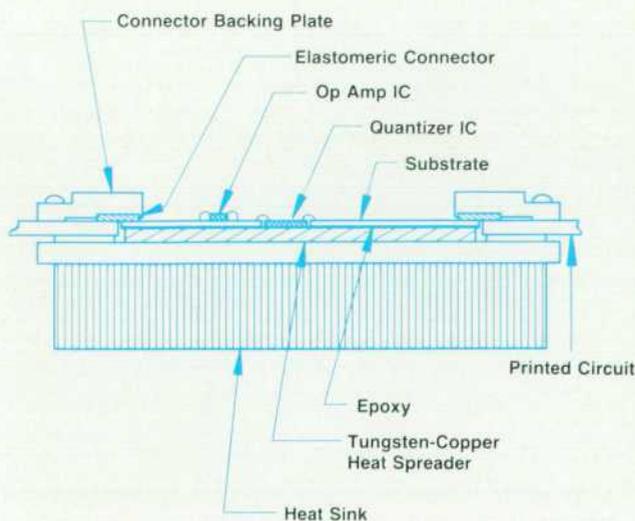


Fig. 3. The ADC hybrid assembly scheme provides excellent thermal conductivity and electrical performance and allows easy removal from the printed circuit board for service.

the type of result obtainable from this program (another example is shown in the cover photograph). Running under HP-UX, the program allowed the hybrid designers to test assumptions about heat sink design, chip placement, and chip attach techniques while monitoring component temperatures. Thermal data taken later from actual hybrids confirmed that the model was correct to within 5°C.

Attachment of a large-area IC to metal has another set of problems.⁴ The thermal coefficient of expansion of silicon is 3 ppm/°C. Typically, metals have coefficients of expansion of about 20 ppm/°C. If a hard solder is used to attach the IC to the metal, stresses built up as the temperature changes lead to cracking of the silicon die. If soft solder or conductive epoxy is used to attach the die to the metal, thermal cycling leads to work hardening of the die attach material. Eventually, cracks form in the die attach material, resulting in increased chip temperature because the thermal impedance of the joint has increased. For circuits that are sensitive to temperature, this will lead to performance drift over time. Eventually, the IC will either detach itself completely from the metal or fail because of elevated junction temperatures.

To meet the requirements of metal-like thermal conductivity and limited differential thermal expansion, a sintered tungsten-copper heat spreader was chosen. These metals are being used as heat spreaders in state-of-the-art electronic packages because of their combination of thermal conductivity and thermal coefficient of expansion. The coefficient of 9.5 ppm/°C is a good match to alumina (8 ppm/°C) and a much better match to silicon than other metals.

Interconnect

The ADC hybrid requires a large number of connections to the printed circuit board supporting it. The interconnect system must meet high-frequency and high-current electrical performance requirements. The realities of production and field repair of the instrument also require the ability to remove and replace hybrids without the risk of destroying either the hybrid or the printed circuit board.

Analog signals being measured by the HP 5185A have nominal bandwidths greater than 125 MHz. Clock and data signals to and from the ADC hybrid are 100K ECL-like signals with subnanosecond rise times. In a system with only a few I/O pins, special RF connectors might be a good

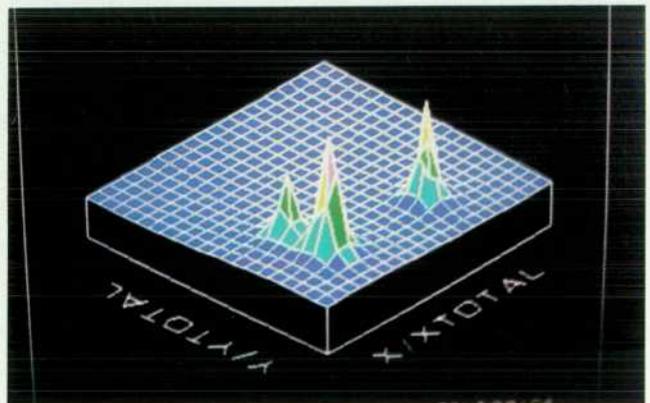


Fig. 4. A special thermal modeling program gives realistic thermal predictions.

approach. However, the ADC hybrid has two RF signal inputs, one RF signal output, four high-speed ECL clock inputs, 12 ECL data line outputs, two ECL clock outputs, and over 30 dc power and reference lines.

Since the only chance for success in a broadband measurement system is to preserve a uniform transmission environment on lines with any significant length, a planar connector was developed to preserve the characteristics of controlled-impedance lines. It uses a metal pattern of 0.005-in lines on 0.010-in centers that connects the 50 Ω transmission lines on the hybrid to the printed circuit board. This metal pattern is plated on a dielectric and backed by an elastomer. A metal backing plate presses the elastomer tightly against the contacts (see Fig. 3).

Past designs of connectors relying on compressed elastomeric pressure contacts demonstrated two problems. Insufficient strength in the backing plate can lead to contact failure caused by deformation of the backing member. The HP 5185A hybrid connector system uses a very strong metal backing plate with a span of less than an inch between the screws holding it in place. Another class of connector problems is caused by contaminants deposited on the printed circuit board copper during board fabrication. If a standard thin gold flash plating (about 3 μ in) is used, contaminants can work their way to the surface in time, causing a rise in connector resistance. This problem has been avoided by plating the contacts on the printed circuit board with over 40 μ in of gold.

The final connector design easily accommodates the 0.100-in pitch of the hybrid I/O lines. Early prototypes using the same concepts handled lines on a 0.025-in pitch, but required precise alignment during assembly. Other published connector designs⁵ can accommodate a 0.025-in pitch, but require unusually precise alignment between the circuit image and the holes on the printed circuit board. Since 50 Ω microstrip transmission lines on a 0.025-in alumina substrate are 0.025 in wide, and since a 3-to-1 space-to-trace ratio gives good crosstalk performance, 0.100 in is a satisfactory line pitch. The connector design allows

the ADC hybrid to be replaced by removing eight screws.

Conclusion

The resulting packaging system provides for good transmission line matching, good thermal matching, high pin count, and ease of removal, while maintaining less than a 5°C/W thermal impedance. Pictured in Fig. 5, the ADC hybrid has 46 conductive vias for low-impedance connections to ground and for power routing on the hybrid, and 71 I/O lines on 0.100-in centers. The HP 5185A preamp hybrid (Fig. 6) uses the same interconnect system. Although no conductive vias are required on this circuit, the same pattern plating process used on the ADC substrate allows the use of 50- μ m traces for the spiral inductors in the upper right corner, without sacrificing low series resistance.

Acknowledgments

Mike Detro and Bill Daley contributed to the design of the hybrid interconnect and heat sinking systems. Hagop Stephanian, the hybrid production engineer, participated in the development and testing of the hybrids. Lisa Craig's work as the hybrid test engineer exceeded the call of duty. Kathi Luiz routinely performed miracles in hybrid prototype assembly.

References

1. A. Fong and R. Knapp, *Thermal Design*, Hewlett-Packard Company, 1977.
2. W.M. Kays and M. Crawford, *Convective Heat and Mass Transfer*, Second Edition, McGraw-Hill, 1980.
3. Gordon N. Ellison, *Thermal Computations for Electronic Equipment*, Van Nostrand Reinhold, 1984.
4. D.R. Olsen and H.M. Berg, "Properties of Die Bond Alloys Relating to Thermal Fatigue," *IEEE Transactions on Components, Hybrids, and Manufacturing Technology*, Vol. CHMT-2, no. 2, June 1979, pp. 257-263.
5. D.L. Timmins, "An Elastomeric Interconnect System for Fine Pitch Leadless Chip Carriers," *Proceedings of the 34th Electronic Components Conference*, 1984.

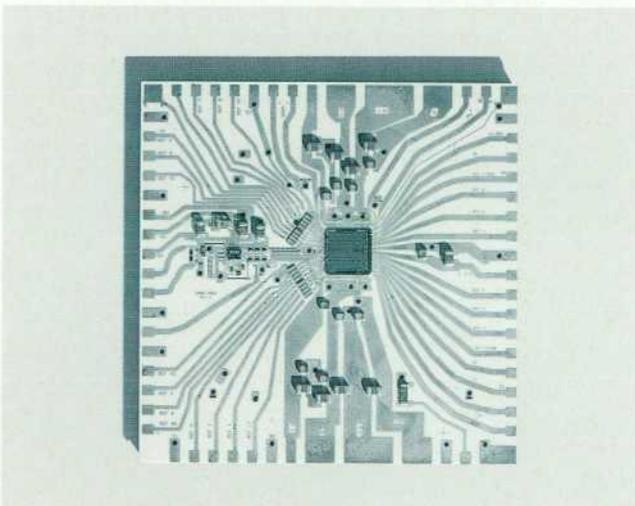


Fig. 5. ADC hybrid. Machining the substrate provides 46 conductive vias for ground and power distribution and allows the quantizer IC to be attached directly to a heat spreader.

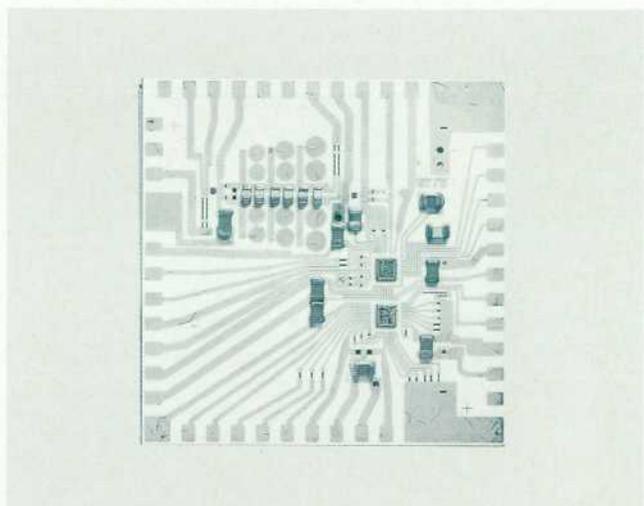


Fig. 6. Pre-amplifier hybrid. Spiral inductors etched onto the substrate (upper right) make a six-pole transitional filter realizable in a portion of the 1.5-in-square hybrid.

Precision Digitizing Oscilloscope Waveform Analysis, Display, and Input/Output

by Douglas C. Nichols

THE HP 5180T/U, HP 5183T/U, and HP 5185T Precision Digitizing Oscilloscopes each consist of a waveform recorder and an analysis, display, and I/O module. The waveform recorder is the HP 5180A, the HP 5183A, or the HP 5185A. For all of the precision digitizing oscilloscopes, the analysis, display, and I/O module is the HP 51089A.

The display is a 9-inch-diagonal 2048 × 2048-point addressable vector CRT display that generates high-resolution text and waveform images (Fig. 1). The precision, wide dynamic range, and high fidelity of this instrument family demanded a display of this quality. This display lets the user view fine detail that cannot be seen on a conventional limited-resolution raster-scan CRT. A fast refresh memory system supports the vector display, as shown in the simplified diagram of the hardware architecture of the analysis, display, and I/O module, Fig. 2.

System control is managed by a real-time multitasking operating system that simultaneously handles multiple processes such as keyboard entry, data analysis, and waveform display. Real-time refers to the machine's ability to respond to external events including waveform triggers, measurement complete signals, time clock interrupts, HP-IB commands, and keystrokes. These events cause interrupt-driven software routines to be executed internally. External user interfaces are high-level and user-friendly. Softkey menus and CRT status displays are used for local entries. High-level ASCII keyword strings are used for remote entry.

A 6803 microprocessor manages all front-panel inputs

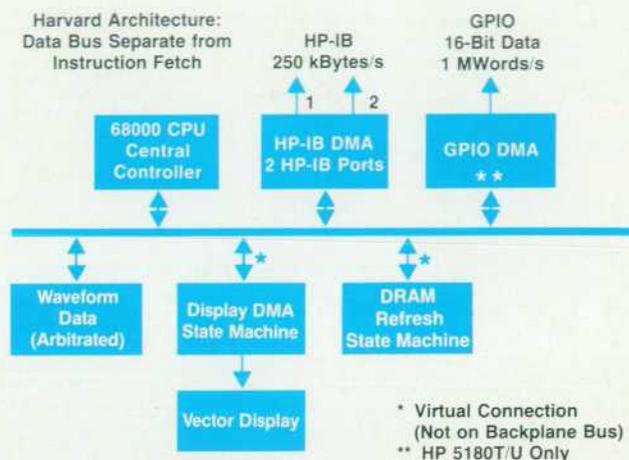


Fig. 2. Hardware architecture of the HP 51089A Analysis, Display, and I/O Module.

including hardkey and data entry knob control firmware and scanning control for the touchscreen display. Finger presses on the CRT display are detected when the finger interrupts a matrix of infrared LED sources and phototransistor detectors. The bezel is molded from an IR-transparent polycarbonate. The solid construction protects the photoelectronics from the effects of dust, fingerprints, and harsh ambient light conditions. For better legibility in bright sunlight, the CRT face is covered by a fine conductive mesh sunscreen. The sunscreen also provides greater resistance

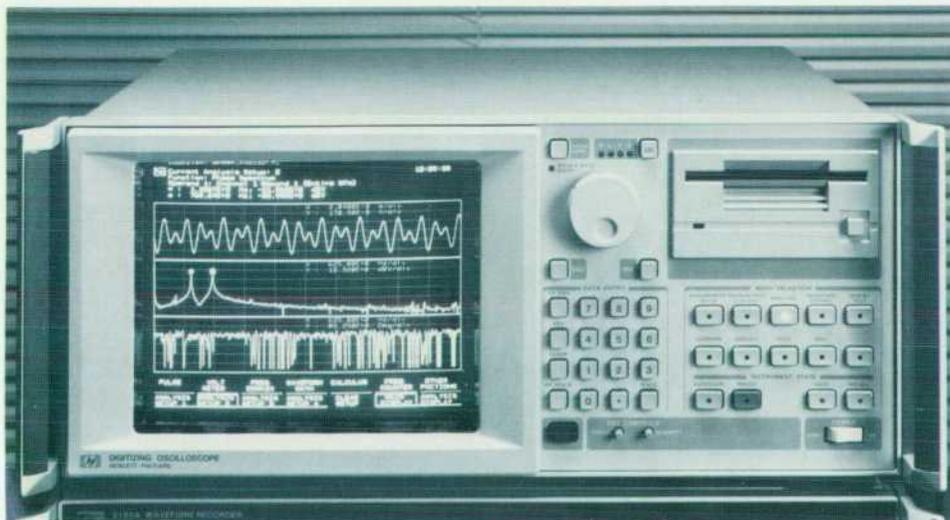


Fig. 1. The HP 51089A Analysis, Display, and I/O Module is used with the HP 5180A, HP 5183A, or HP 5185A Waveform Recorders for precision digitizing oscilloscope applications. The display is a 9-inch-diagonal, 2048 × 2048-point vector display capable of showing finer detail than a raster-scan CRT.

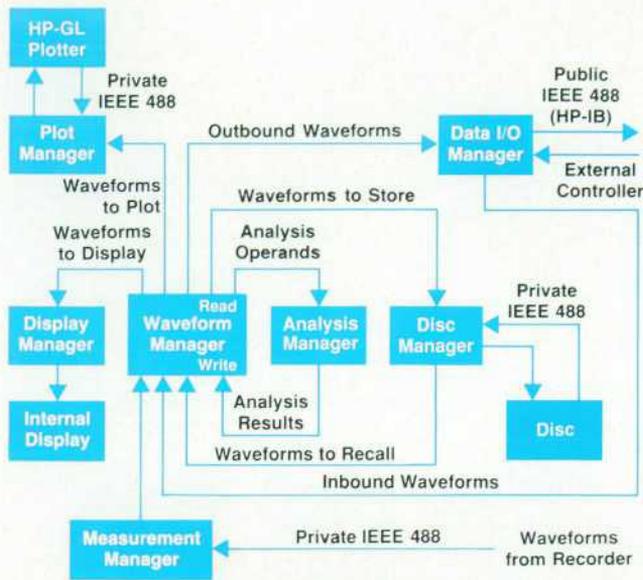


Fig. 3. The HP 51089A takes control of the measurement process. Waveform traffic is controlled by the waveform manager.

to electrostatic discharge.

When the analysis, display, and I/O module is used with the HP 5180A, HP 5183A, or HP 5185A Waveform Recorder, the HP 51089A's microcomputer takes control of the overall measurement process (Fig. 3). Following capture of a waveform, the result is transferred to the HP 51089A via a private HP-IB (IEEE 488/IEC 625) bus. The HP 51089A hardware architecture features a shared bus design that allows waveform data transfer from the waveform recorder to occur simultaneously with program execution. A waveform management subprogram directs the machine se-

quencing operations as ordered by the user. Because the firmware executes in a true multitasking environment, it is practical to do several things simultaneously, such as acquiring measurement data in the waveform recorder, transferring data over the HP-IB, plotting results, and performing analysis operations.

Systems Capabilities

All instrument states are programmable over the HP-IB. Programming commands are consistent throughout the product family. A program that runs on an HP 5183U will run without change on an HP 5180T, assuming it uses features common to both waveform recorders. Simple commands provide significant results. The BASIC program:

```
OUTPUT@Digitizer;"RMS?"
ENTER@Digitizer;Rms_value
```

will execute the rms analysis function on the current user-defined operand and return the result to the external controller. The result will be stored in the variable Rms_value. Teach and learn modes are supported using complementary commands. The input range is set using the teach string RANGE 2 V and is queried with RANGE?. Entire instrument setups can be read with a single SET? query. Waveform input and output can be accomplished using either ASCII or binary transfer mode. The HP 51089A maintains complete state information for all waveforms (i.e., sample rate, trigger position, y-axis range and offset, etc.). This information can be separately read over the bus for complete waveform knowledge. Using the binary waveform transfer mode, transfer rates of ~250 kbytes/second are achievable.

Because transfer of waveforms from an external controller is supported, it is possible to use the disc, display, plotter, and analysis even when no waveform recorder is connected. For example, waveform data, synthesized or

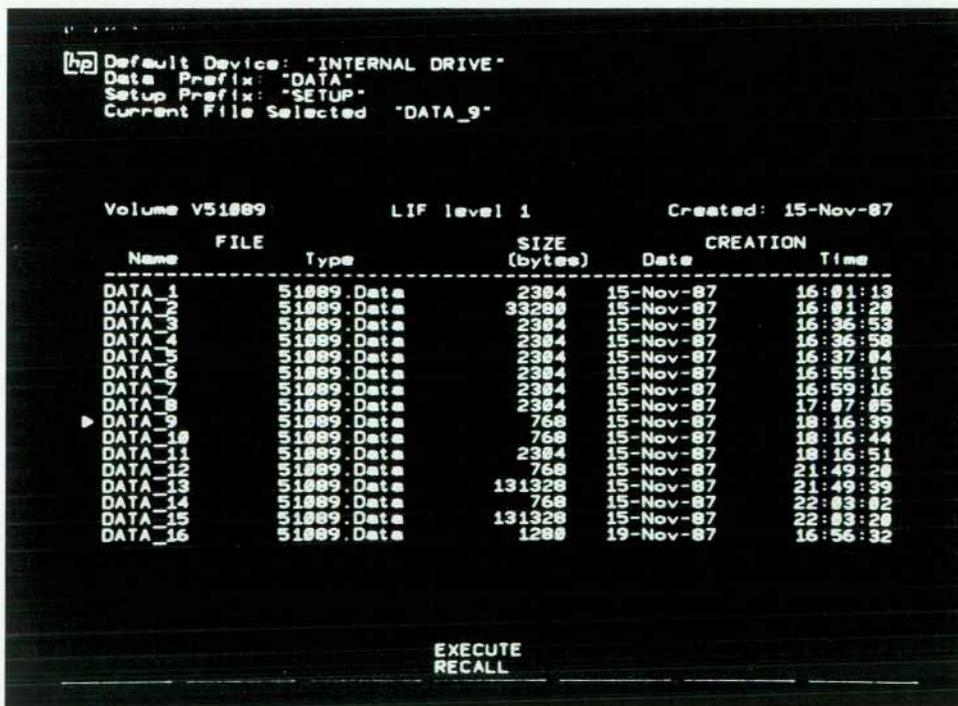


Fig. 4. The user interacts with the optional flexible disc via directories.

captured on some other device, can be transferred, and a power spectrum measurement can be performed, with the calibrated result (in either dBm or dBV) read back to the external controller. The programming interface is both forgiving on input and precise on output. Examples of forgiving input include: choice of abbreviated or complete program commands (e.g., use TRIGGER LEVEL 1 V or TRIG LEVE 1), case insensitivity, interchangeability of integer/real numbers, and acceptability of all reasonable terminations (CR/LF/EOI).

Built-In Disc Drive

The HP 51089A features an optional 3.5-in flexible disc drive that provides storage and recall of waveforms, scalar analysis results, and instrument setups. User interaction with the disc is via directories, such as the one shown in Fig. 4.

The HP 51089A autonames the stored results using an ascending numerical suffix, prepended with a character sequence chosen by the user. Using the HP 51089A's real-time clock and perpetual calendar, all stored results are time-stamped with one-second resolution. This provides a powerful record-keeping utility, especially when many results need to be stored.

Display Features

We have designed the display subsystem as a peripheral, which the user can configure to suit the application. For example, a user may elect not to view the raw measurement data in an automated testing application because this data is always directly transferred to an external computer. For a different application, a user may be interested only in a frequency-domain view of the input. By defining the raw measurement data as the source operand for a spectrum analysis operation and assigning just the analysis result to the display, a user converts the instrument to a multichannel spectrum analyzer for single-shot or repetitive signals.

The versatility of the display can be enhanced since it is decoupled from the measurement process. The waveform

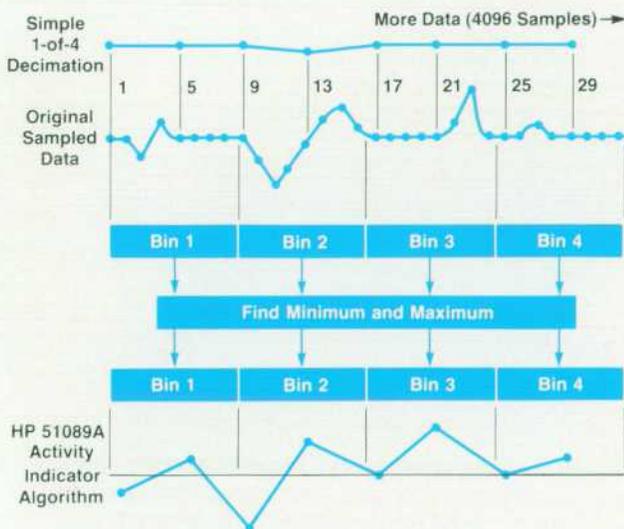


Fig. 5. The compression algorithm in the HP 51089A ensures that signal activity is never missed.

reconstruction algorithms, described in the paper on page 26, further enhance visual detail.

Large waveforms are processed through a compression algorithm before being displayed. The need for compression arises whenever the number of waveform sample points exceeds the actual displayable resolution. The simple 1-of-N sample decimation scheme often used is a poor choice, because the displayed waveform will appear to have been sampled at $1/R$ of the actual rate, where R is the ratio of waveform samples to displayed resolution. This phenomenon, technically termed display aliasing, is undesirable because it presents waveform images that are not a reasonable representation of the captured signal. Our compression method sorts through waveform data bins of width $2R$ and selects the maximum and minimum sample values from each. These extreme values are presented on the display. The effect is to present a display very much like an analog oscilloscope or a strip-chart recorder display. That is, even if the time base is too slow to show detail, the user can see the maximum and minimum waveform excursions. Unlike simple decimation, which can fail to show significant signal activity, this compression method assures that signal activity will appear on the screen.

Fig. 5 illustrates the distinction between the simple decimation and compression methods. For this example, a 4096-sample waveform was chosen. With horizontal display resolution of 1024, only one fourth of the samples can appear on the display. The upper portion of Fig. 5 shows the result created by simple decimation—a user could be misled by this result into thinking that the input remained unchanged during this time period. The lower portion of Fig. 5 shows the display result created by the minimum-maximum compression algorithm. Regardless of the compression ratio, this algorithm never fails to indicate the presence of signal activity. Seeing an area of interest on the screen, a user can then zoom in for a detailed examination of the waveform activity (Fig. 6).

Direct Hardcopy

The HP 51089A can produce fully annotated output on any of HP's family of digital plotters. A representation of the display can be plotted (see Fig. 7 for a typical example), or an individual waveform can be selected for output in its entirety. For plots of large waveforms, no more than 1024 samples per page are output; our rationale was that trying to squeeze more data on a page wouldn't be reasonable. If an autopaging plotter is being used, the HP 51089A recognizes the capability and automatically ejects the pages as consecutive segments of the large waveform plot are completed.

Analysis Capabilities

The analysis, display, and I/O module provides a wide range of functions categorized as pulse, voltmeter, waveform math, frequency domain, calculus, frequency counter, and iterative functions. Many specific functions are available within each of these categories, some of which are detailed later in this article. The user has access to four independent analysis setups. A setup consists of:

- The function to execute
- Input waveform(s) to be analyzed

- Any possible special modifiers such as a preprocessing window for frequency domain functions.

Fig. 7 illustrates an analysis setup. The top trace is the analysis result and the lower traces are the waveforms being analyzed. The user has selected the waveform add function, with operand 1 the captured output of a function generator in swept sinusoid (linear FM) mode, and operand 2 the modulation signal. Delimitation capability is used here to isolate different segments of the operands being added. The vertical lines on the lower traces show the portions being analyzed, while the status information near the top indicates the segment of each operand used in the analysis.

Analysis function operation is designed for ease of use. An operation so natural that it often goes unnoticed is the automatic execution of the analysis whenever conditions change to require reexecution. The most obvious condition causing automatic reexecution is updating of the source

operand. For example, if a new data acquisition occurs and the acquired waveform is the operand for a SPECTRUM in dBV function, the analysis result will automatically update using the newly acquired source waveform. Other conditions that invoke automatic analysis execution are:

- Change of operand (e.g., changing operand 1 from channel 1 record 1 to a waveform named Pulses).
- Change of position or length of delimit window. As the user moves the delimit window across the source waveform, the analysis result is continually updated.
- Change of a modifier that affects the selected analysis function (e.g., changing from HANNING to FLATTOP window will invoke automatic execution if the function is SPECTRUM in dBm, but won't cause execution if the function is AVERAGE FREQUENCY, since the window selection doesn't affect this result).

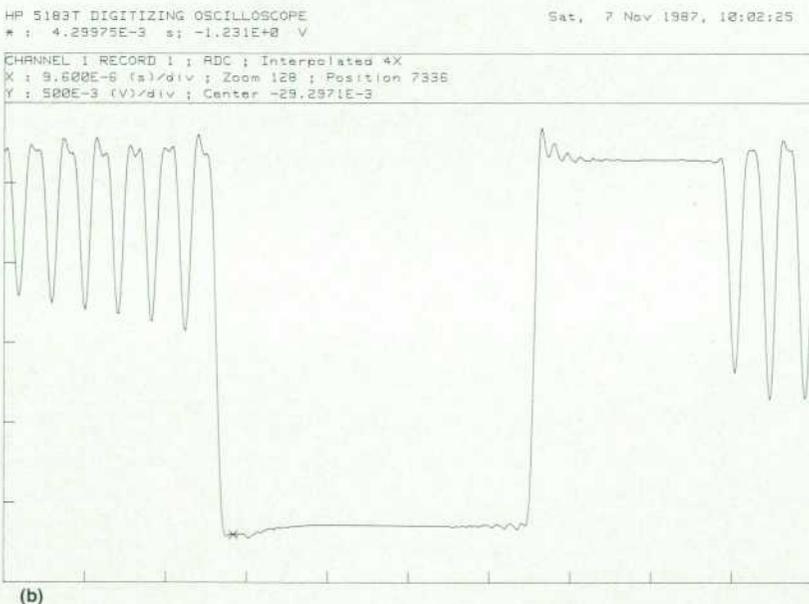
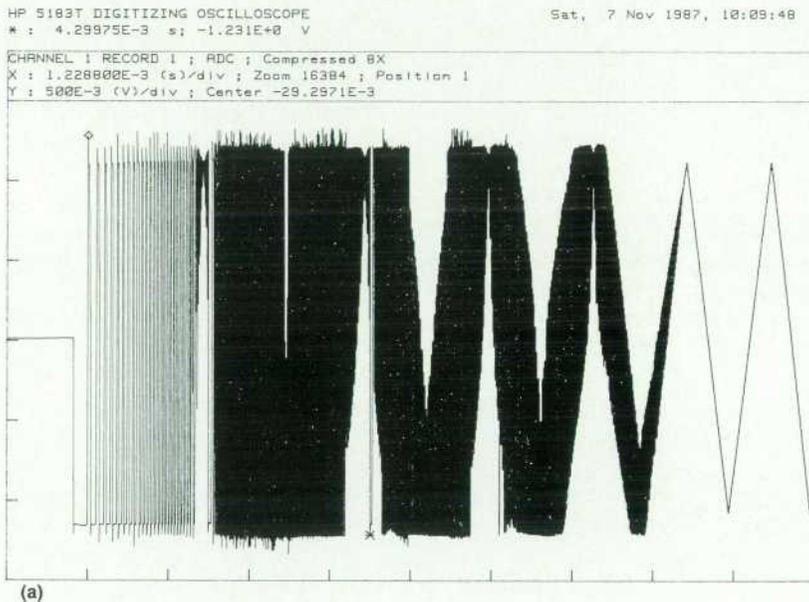
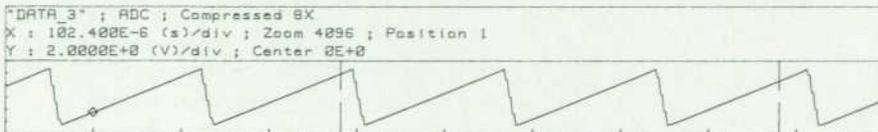
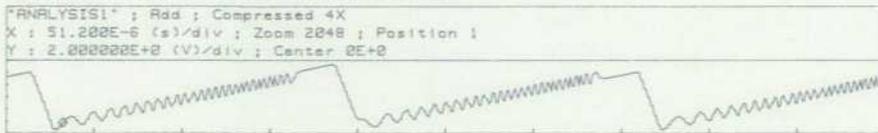


Fig. 6. (a) 16K-sample display shows a macro view of a function generator turn-on transient. (b) Display zoomed in on cursor, showing fine detail.

Current Analysis Setup: 1
 Function: Add
 Operand 1: "DATA_2" Start: 277 Length: 2048
 Operand 2: "DATA_3" Start: 1561 Length: 2048



PULSE	VOLT METER	FREQ DOMAIN	WAVEFORM MATH	CALCULUS	FREQ COUNTER	OTHER FNCTIONS
ANALYSIS SETUP 1	ANALYSIS SETUP 2	ANALYSIS SETUP 3	ANALYSIS SETUP 4	CLEAR SETUP	MAIN DISPLAY	ANALYSIS DISPLAY

Fig. 7. Analysis setup display. The top trace is the analysis result and the lower traces are the waveforms being analyzed.

Analysis User Interface

Fundamental rules for any machine to be easy to use include:

- Never surprise the user with internally invoked changes of instrument state.
- Requested results should always appear automatically on the display.

Strict adherence to one of the above objectives may lead to failure to meet the other. For example, a typical display presentation is two channels of waveform data. A user who then selects an analysis function probably expects to see the analysis result on the display, but this cannot happen without breaking the first rule. We have chosen to resolve this dilemma by providing two display modes, as shown

in Fig. 8. The main display mode can be configured by the user with complete generality; any custom configuration can be achieved. When the user selects an analysis function for execution, the firmware in the HP 51089A first checks to see if the analysis result (the waveform name ANALYSIS1 is reserved for the output of analysis setup 1) is part of the user-customized display setup. If it is, no display mode change is needed, since the analysis result is going to appear on the display. If it isn't, the display is automatically switched to analysis display mode to force the analysis result to be seen. The analysis display mode always puts the analysis result in the top trace and the associated operand(s) in lower trace(s). Recovery of the user-customized setup is accomplished by pressing the **MAIN** display

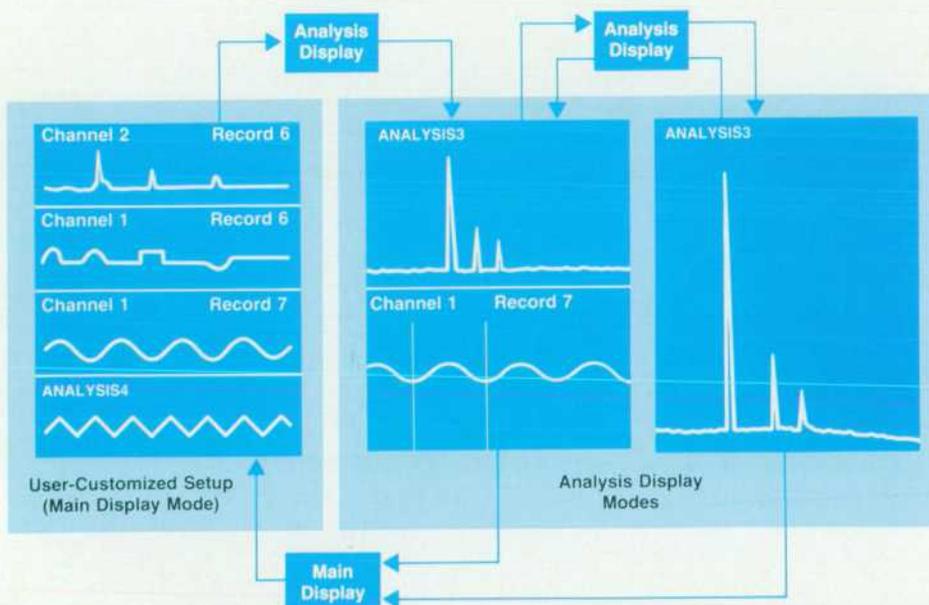


Fig. 8. Two display modes, main and analysis, ensure that the user sees the expected data.

key.

The creation of main and analysis display modes allows both objectives to be met. The interaction of these modes is explained by the following table:

User Expectation	HP 51089A Response
Results appear on display when analysis function is invoked.	Results always appear, either automatically or by display mode change to analysis.
Customized display configuration won't be destroyed when analysis function is invoked.	Expectation automatically met if result is on display. If change to analysis display occurs, user recovers custom display setup by pressing the MAIN display key.
User wants a visual representation of the analysis setup.	Pressing the ANALYSIS display touchkey will assign the analysis result to the top trace and the analysis operand(s) to lower traces. Any operand delimitation will appear visually as dashed vertical lines on the operand waveform(s).
User wants to view only the analysis result.	Pressing the ANALYSIS display key repeatedly causes it to toggle between a display showing just the analysis result and a display showing the analysis result and the analysis operand(s).

The general-purpose trace assignment capability provides the ability to customize the display for only the re-

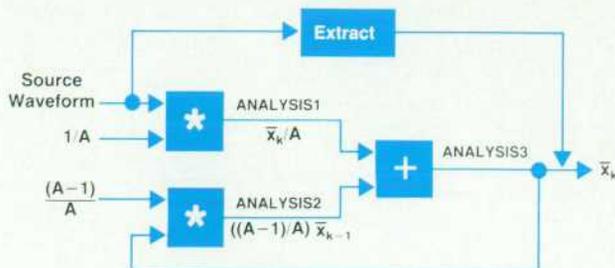


Fig. 9. A decaying weighted average computation is an example of a user-customized analysis.

sults of interest. For example, to use the instrument as a single-shot spectrum analyzer and frequency counter, the user can configure trace 1 to display the spectrum in dBm and trace 2 to display the average frequency. Whenever the time-domain source updates, these results will be updated on the display.

Multiple Sequence Processing

The user can perform analysis on analysis results, so it is possible to connect the four analysis setups to create functional results not directly available from a single function. As an example of how multiple sequence processing can be used to customize a desired measurement, consider trying to generate a decaying weighted average from a succession of waveform updates. One way to implement this is using the expression:

$$\bar{x}_k = \bar{x}_{k-1}(A-1)/A + x_k/A, \quad \bar{x}_1 = x_1$$

where \bar{x} indicates an average result waveform, x a source waveform, and A a weighting factor. Large values of A cause older updates to have greater influence over the current average. Values of A approaching 1 cause the current

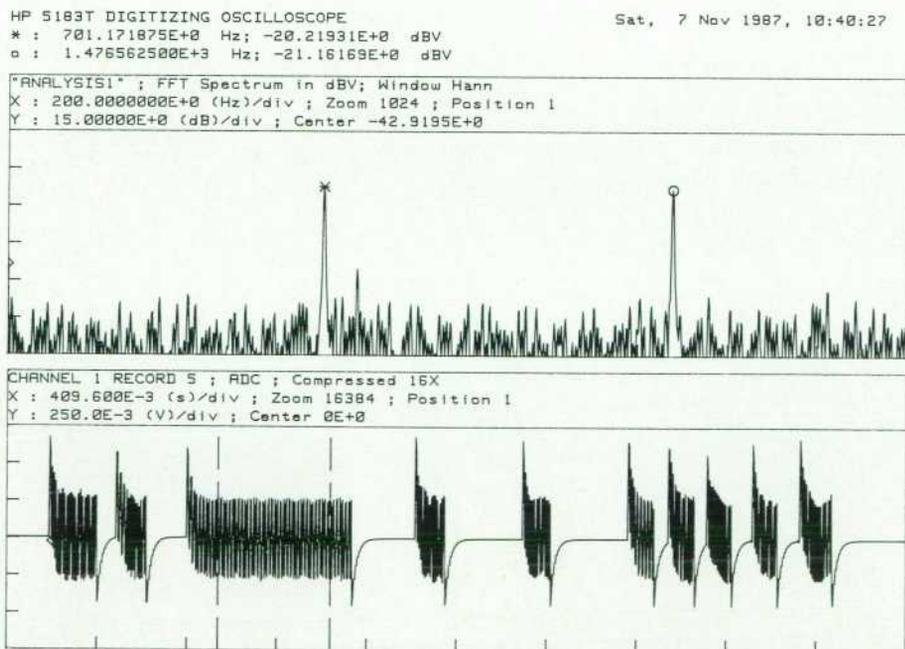


Fig. 10. (bottom) Time-domain display of a touchtone generator output. (top) Spectrum in dBV shows that the third key pressed was a 3.

average result to be dominated by recent updates. Fig. 9 illustrates the process. To begin the averaging process the first update is EXTRACTED (copied) to the average result. Upon each subsequent update the product $(1/A)x_k$ will be automatically computed. Pressing the ANALYSIS SETUP 2 key causes the product $\bar{x}_{k-1}(A-1)/A$ to be computed. Pressing ANALYSIS SETUP 3, which has been defined to sum the products just computed (by assigning as its operands the outputs of analysis setups 1 and 2), will produce the new average.

Keeping Track of Units

The analysis software has general units handling capability that automatically processes units information from one execution to the next. One advantage this provides is easy verification that the dimensionality of the sequence of analyses is correct. For example, suppose you want to subtract the third integral of channel 1 record 1 from the third integral of channel 2 record 1. If you've done it right, the Y-axis units of the sources and the result will appear as Vs^3 . If by mistake you subtract the second integral of channel 2 record 1, the resultant units will be UNKNOWN, an indication that the Y-axis units of the operands weren't dimensionally correct for subtraction.

The analysis software will always perform an indicated operation, regardless of the dimensional correctness of the operand units. The distinction between dimensional correctness and incorrectness is that the unit indicator UNKNOWN is used whenever the dimensions are incorrect for the indicated operation. Dimensional correctness is partly dependent upon the operation to be performed. Subtracting a second integral from a third integral doesn't make dimensional sense, but multiplying them does, and results in Y-axis units of V^2s^5 .

The machine also processes units on the X axis, a more complicated task than Y-axis units processing because the units must make sense together and the X-axis increments

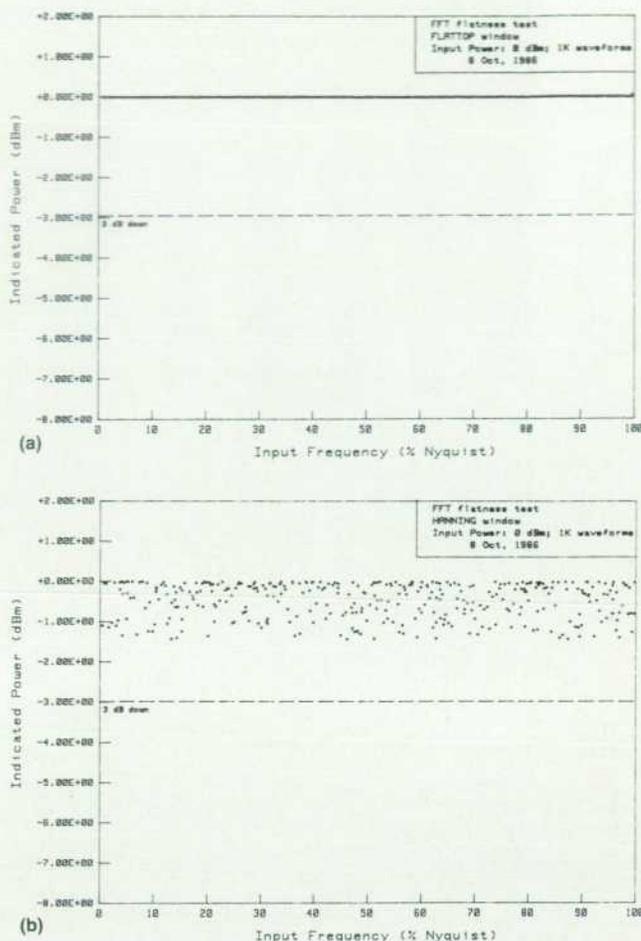


Fig. 11. (a) Accuracy of the flattop window is demonstrated by a series of spectra of 0-dBm sine waves at random frequencies. (b) The Hanning window is less accurate but has superior frequency resolving properties.

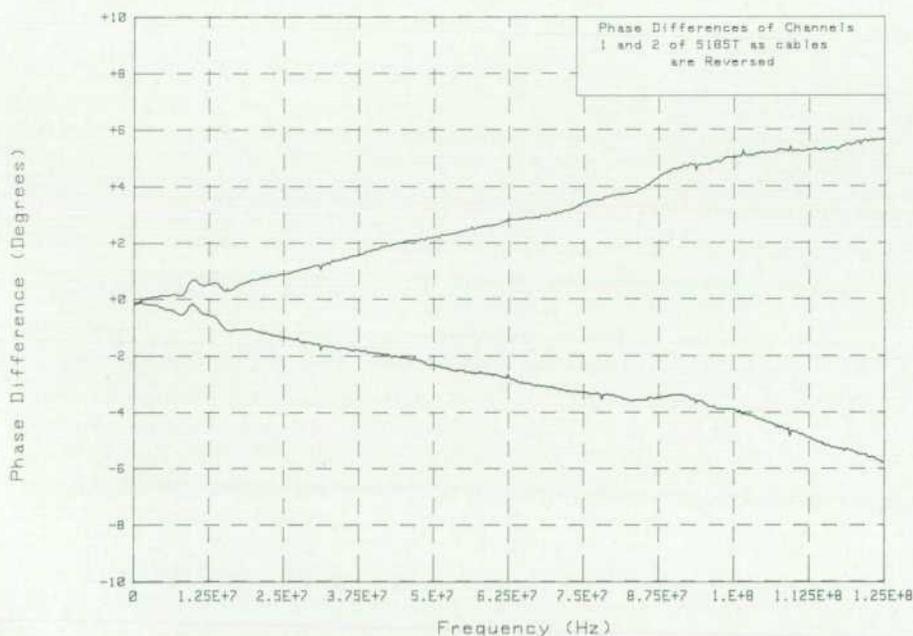


Fig. 12. A measurement with no device under test can be used to calibrate timing differences between HP 5180T, HP 5183T, or HP 5185T inputs.

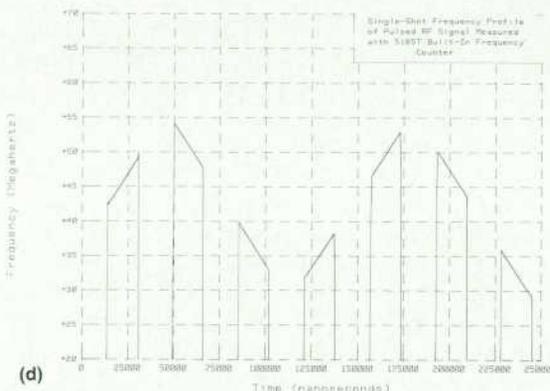
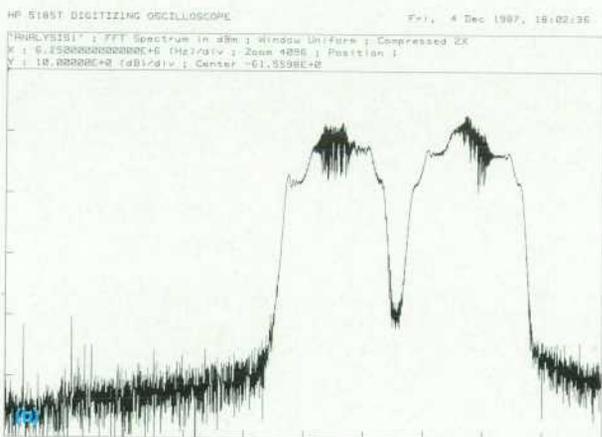
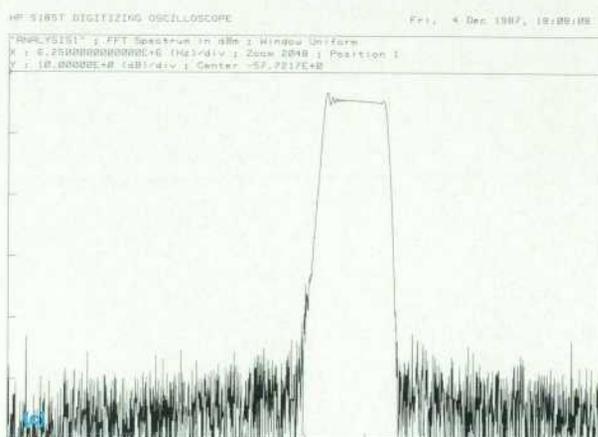
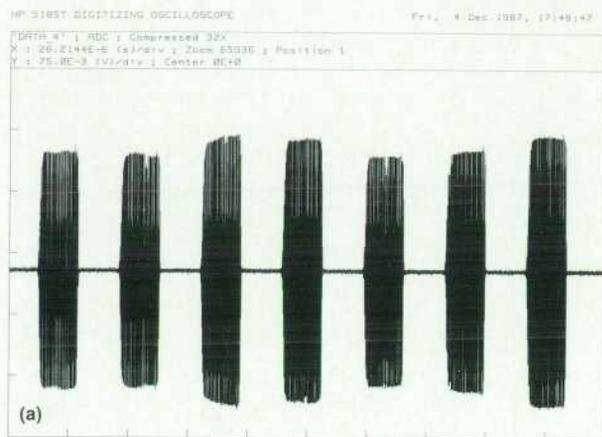


Fig. 13. Pulsed RF signal. (a) Time record. (b) Frequency spectrum. (c) Frequency spectrum of an individual burst. (d) Frequency profile measured using the built-in frequency counter.

(e.g., the sample rates for ADC waveforms) must be the same for both operands. For example, multiplying a waveform captured at 250 MHz with a waveform captured at 1 Hz is not sensible, so the resultant X-axis units become UNKNOWN.

Frequency-Domain Functions

Transformation of the source data (time domain) to the frequency domain is accomplished using either the dBm, dBV, or phase spectrum functions. For example, the lower waveform in Fig. 10 is a time-domain look at the output from a touchtone generator as a sequence of keys was pressed. The sequence was the digits one through nine, ending with zero. Timing information such as tone burst duration or time between bursts is readily determined from the time waveform. However, if we want to know which keys were pressed, the time-domain view is of little help. Even after zooming in on each burst the determination is nearly impossible. By selecting the spectrum in dBV function and using the general delimitation capability of analysis, the frequency content of each key press can be isolated and examined. The top waveform of Fig. 10 shows that the third tone burst is essentially two-tone. Frequency and amplitude levels are directly read by placing cursors on the peaks. A two-tone spectrum at 1476.56250 Hz and 701.171875 Hz identifies the third key pressed as a 3.

Conversion to frequency domain is accomplished using a fast Fourier transform. The implementation is a radix-4, in-place, real-in/complex-out algorithm¹ that executes on the 68000 microprocessor in 450 ms for 1024 source samples. Conversion of the complex output of the FFT to either a magnitude (in dBV or dBm) or a phase spectrum requires Cartesian-to-polar coordinate transformation. Rather than compute this transformation directly from the equations:

$$\text{Magnitude}_i = \sqrt{(\text{real}_i)^2 + (\text{imaginary}_i)^2}$$

$$\text{Phase angle}_i = \tan^{-1} \left(\frac{\text{imaginary}_i}{\text{real}_i} \right)$$

which requires many slow transcendental function computations, a method known as the CORDIC computing technique (COordinate Rotation DIgital Computer) is used.² Developed originally as a special-purpose computer for real-time airborne computation to solve navigation equations, the CORDIC technique is well-suited for providing coordinate transformations of the FFT output data.

The selection of a windowing function for a frequency-domain measurement is dependent upon input signal characteristics and the type of frequency domain information that is being sought. For example, if amplitude accuracy is the critical measurement, the FLATTOP window is the best choice. The data in Fig. 11a was generated using

source waveforms at random frequencies, all having rms power levels of 0 dBm. The spectrum in dBm was computed for each source. The results, which show a maximum variation of less than 0.1 dB across the spectrum, agree closely with the theoretical prediction and show that the accuracy of the result is optimum when the FLATTOP window is used. Fig. 11b shows a similar plot for the HANNING window. The HANNING window doesn't provide the amplitude accuracy of the FLATTOP window, but has superior frequency resolving properties.

The phase spectrum analysis function is especially useful for measuring time differences between multiple inputs. A typical example of the device under test would be an amplifier or a filter whose phase response needs to be determined. The signal source is stepped through a succession of test frequencies and the phase difference between the inputs is measured at each frequency. The data is directly available as delta-phase versus frequency.

This same technique can be used to calibrate systematic timing characteristics between measurement channels when using any of the digitizing oscilloscopes. Fig. 12 illustrates the results of using this technique on an HP 5185T Digitizing Oscilloscope. The upper waveform of Fig. 12 was captured first, then the cables were reversed and the new result was the lower waveform. These results are mirror images about zero degrees, indicating that the primary timing difference in the measurement channels is contributed by the cables (to the eye, the cables are the same length). The delta-phase results directly obtainable from the HP 51089A can be converted to timing differences from the expression

$$\Delta \text{ time} = \frac{\Delta \text{ phase}}{\text{input frequency} \times 360}$$

Selecting 75 MHz as a sample frequency, the expression yields

$$\frac{3.5^\circ}{75 \text{ MHz} \times 360^\circ} \approx 130 \text{ picoseconds}$$

Resolution for delta-phase measurements is typically better than 0.1 degree.

Average Frequency

This analysis function is a software implementation of a synchronous gating reciprocal frequency counter.³ The result is average frequency, computed from the formula:

$$f_{av} = (\text{number of input periods})/(\text{time interval}).$$

Because the time interval value always spans an integral number of periods, no error in the average frequency result is caused by inclusion of partial periods.

Using techniques that enhance the accuracy of the time interval portion of the average frequency measurement produces frequency counting capability comparable to instruments dedicated to measuring average frequency. However, unlike counters, a waveform image is available, so it is apparent what is being measured.

Dedicated frequency counters operate on live data so there is only one chance to obtain the result. The HP 5180T, HP 5183T, and HP 5185T operate on captured waveform data. The delimitation feature of analysis provides the equivalent of a built-in adjustable external gate for the average frequency measurement. By combining these capabilities, a single-shot, zero-dead-time frequency profile of a captured waveform can be obtained. This is illustrated in Fig. 13, which shows a captured pulsed RF sequence, its overall

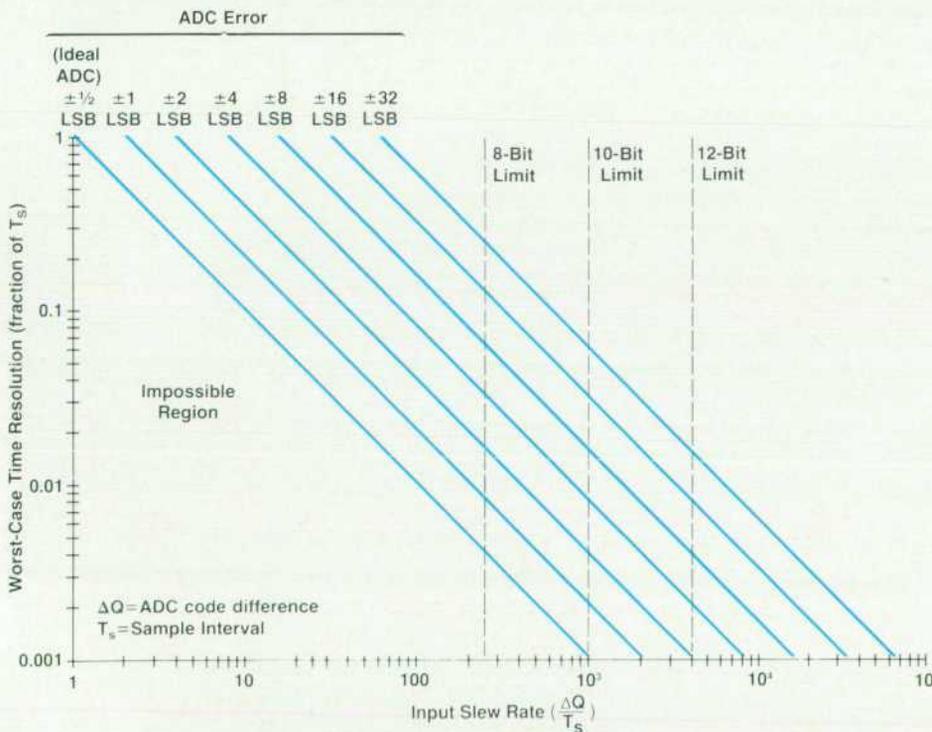


Fig. 14. Timing resolution can be greater than indicated by the sample rate. This plot shows worst-case time resolution as a function of digitizer error and input slew rate.



Fig. 15. Using the event crossing function at different threshold levels, a single marginal pulse can be detected automatically.

magnitude spectrum, and the recovered frequency modulation, obtained with delimited average frequency analysis. The recovered modulation waveform clearly shows what the other two cannot—the modulation source was a triangular signal freely running with respect to the pulse repetition frequency.

Waveform Reconstruction for Time Estimation

A traditional frequency counter minimizes time error by accumulating clock pulses in the time register at a high rate. The HP 5345A Universal Counter uses a 500-MHz sample clock, resulting in fundamental timing resolution of 2 ns. The digitized data being evaluated in the HP 5180T, HP 5183T, or HP 5185T will not have been sampled at rates as high as the clock rates in a frequency counter, so

it might appear that the timing error would be greater. For example, evaluating data captured on an HP 5180T at 20 MHz (50 ns/sample), one might expect worst-case total timing error of ± 50 ns for an interval measurement. However, the digitized data sequence contains information that can be used for timing resolution improvement in a way not possible with a frequency counter. Signal level information on either side of the threshold level can be used to determine crossing times more precisely. Simple linear interpolation between samples leads to significant reduction of timing error. Additional waveform reconstruction algorithms are used in the HP 51089A to obtain a more precise timing measurement.

As the input slew rate increases, the significance of ADC amplitude error diminishes in the timing estimate. Fig. 14

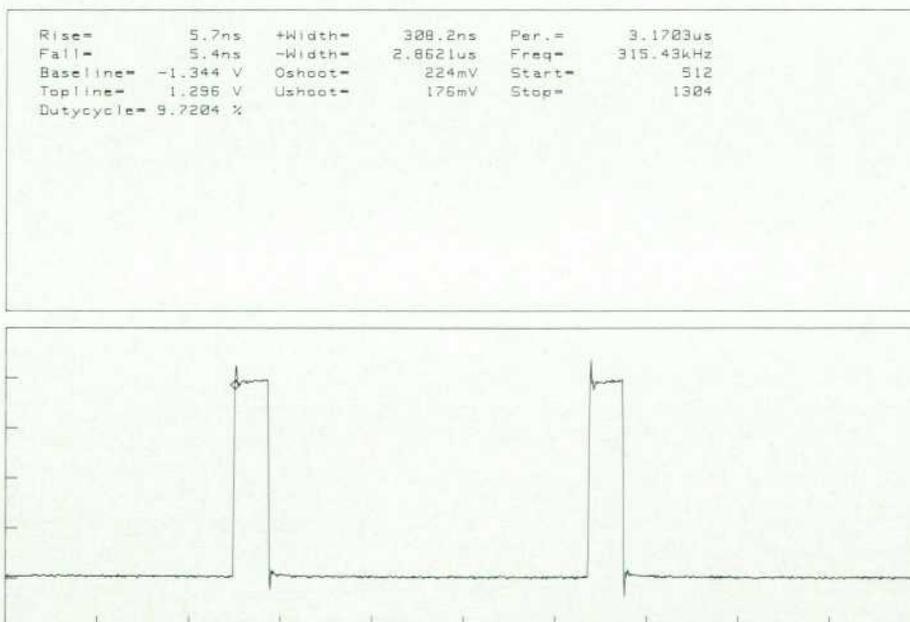


Fig. 16. Pulse analysis result.

Handling of Significant Digits

Typically, measurement instruments provide more digits than are significant, relying on the user's ability and patience to determine what is significant. A major weakness of this approach is that the number of significant digits is determined by the largest possible value for the measurement, resulting in overstated resolution for values that aren't the largest. For example, consider computing the peak-to-peak voltage of a waveform captured on the 5V range of the HP 5183A. The ADC quantization step is the resolution of this measurement, which turns out to be 2.5 mV. The largest possible value for peak-to-peak voltage on this range is about 10V. To support the largest value to correct resolution requires six digits (e.g., 10.0025V). With a fixed six digits, smaller results are vastly overstated. For example, 27.5 mV will be stated as 27.5000 mV, a value that seems to suggest resolution on the order of 100 nanovolts! To avoid this kind of overstatement, a significant digits handling algorithm was developed and implemented that dynamically adjusts the number of significant digits based on the magnitude and resolution of the result. The expression is:

$$D = L[\log(x)] - L[\log(F[r])] + 1$$

where D is the number of mantissa digits required to represent the value properly, L is a function returning the closest integer value that is equal to or lower than the number (e.g., $L[2.3] = 2$, $L[-6.8] = -7$), x is the number being represented, F is a function returning a number containing only the least-significant nonzero digit in the input number (e.g., $F[2.35] = 0.05$, $F[1.00000025] = 5 \times 10^{-8}$), and r is the resolution appropriate for x.

Using the example above, $x = 10.0025$ would produce $L[\log(10.0025)] = 1$ and $F[0.0025] = 0.0005$, so $L[\log(0.0005)] = -4$ and $D = 1 - (-4) + 1 = 6$ digits. The value $x = 27.5$ mV would produce $L[\log(0.0275)] = -2$, so $D = -2 - (-4) + 1 = 3$ digits. In each case, the result is expressed to the proper number of significant digits.

One special problem in significant digit handling occurs when the resolution term is not an integer power of ten. In the preceding example, the resolution is 2.5 mV, which is not a power of 10. A single result expressed with the proper number of significant digits, such as 4.3425 volts, gives the incorrect impression that 4.3424 and 4.3426 are possible results, when in fact the nearest possible results are 4.3450 and 4.3400 volts. The user has imagined 100- μ V resolution when the actual resolution is only 2.5 mV. All five digits need to be displayed in the original result since the last digit is truly meaningful (not displaying it creates bias error). It is unfortunate that, in such cases, providing a single result to its proper resolution doesn't identify the magnitude of the resolution.

A more dramatic example of this dilemma occurs when computing frequency domain results. The frequency bin size of an FFT is f_s/N , where f_s is the sampling frequency and N is the number of samples in the original time-domain source data. For preset conditions on an HP 51089A with an HP 5183A (i.e., an

HP 5183T) this gives a frequency resolution of 3.90625 kHz. The frequency value obtained by using a cursor to measure the 200th sample of the magnitude spectrum is 777.34375 kHz. There is a tendency to look at a value like this and conclude that the resolution is around 10 mHz, when in fact it is around 4 kHz! Implementations that always provide a fixed number of digits get into serious trouble in a situation like this, because the worst-case measurement here is a magnitude of 2 MHz with 3.90625 kHz resolution, which requires nine digits. When a small result is computed, such as 7.81250 kHz, the fixed implementation would represent this as 7.81250000 kHz, suggesting resolution to 10 μ Hz!

The typical alternative implementation, arbitrarily limiting the output to five digits (perhaps rounding as well) in an attempt to keep the results looking good on the display and avoid overstatement, is really the worst approach of all. Not only does overstatement of results continue to be a problem (for results near the resolution), but now understatement of results often occurs as well. Using an example from the paragraph above, 777.34 kHz may look good, but it is needlessly biased away from the true result by 3.75 Hz.

The best solution is to provide the resolution value as well as the result. With this information, results like 777.34375 kHz are not apt to be interpreted as having such fine resolution because the message **RESOLUTION: 3.90625 kHz** would appear next to the result. The practical reason for not doing this is that it doubles the number of numerical values on the display, which complicates the display appearance and makes extracting information tedious.

The following list summarizes the alternatives. In any event, the present implementation in the HP 51089A always provides the correct number of digits for each measurement.

Implementation	Comments
Fixed digits, large enough to cover worst case	Extreme overstatement of resolution. Easy to implement for given measurement. Hard or impossible to find one value for all cases.
Fixed digits, limited to look good on display	Rarely will provide result to correct resolution. Often understates and biases results. With small results, will overstate resolution.
Implementation using formula	Result always stated correctly. Resolution may appear overstated when resolution step has many digits (above methods have this problem also).
Implementation using formula with resolution value indicated to user	This is ideal if you can afford the space and additional computation time.

illustrates the interrelationships between input slew rate at the threshold crossing, the amount of timing enhancement as a fraction of the sample interval, and the resolving capability of the digitizer. The plot considers the digitizer error to be the only source of error in timing estimation; for many measurements, this is the dominant error source.

Using an HP 5185T, single-shot timing resolution better than 200 ps is achieved for inputs ranging from 1 MHz to 100 MHz, digitizing at 250 megasamples per second.

Event Crossings

This analysis function is much like the totalize function

* : -12.064E-6 s; 48E-3 V
o : 11.856E-6 s; -16E-3 V

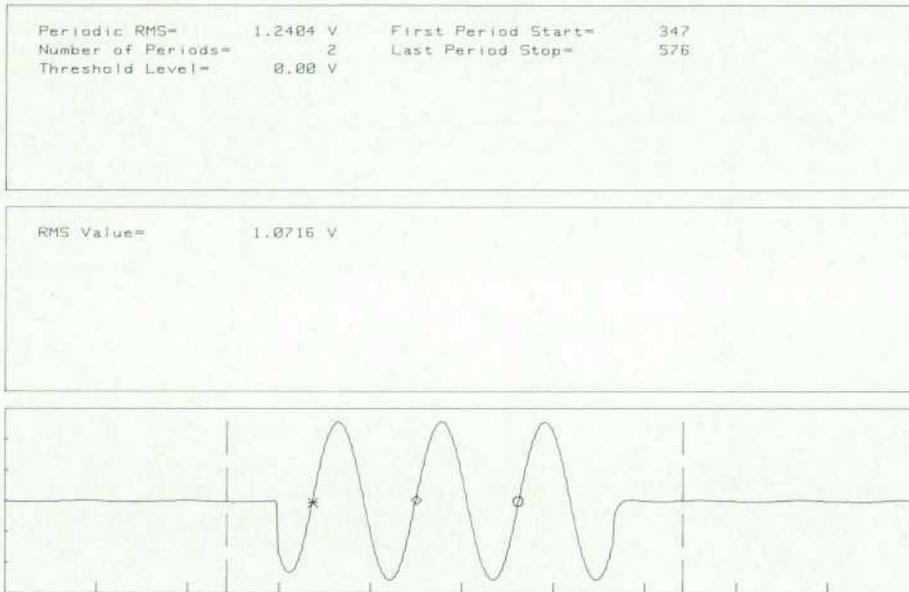


Fig. 17. The periodic rms function can avoid errors caused by inclusion of partial periods in the rms result.

found on some frequency counters. Events in each direction (+, -) are computed independently. Whereas a traditional frequency counter uses analog hardware to determine when events occur, the implementation in the analysis, display, and I/O module uses software for event determination.

The event crossing function is useful for analyzing complex pulse trains. Measurements easily performed using this function include locating a missing or extra pulse and identifying pulses that are close to noise margins (Fig. 15).

Pulse Analysis

The pulse analysis package is based upon IEEE standards.⁴ An example of its use is shown in Fig. 16.

Periodic RMS

This analysis function computes the same result as the basic rms function with the exception that it prequalifies the X-axis range for the computation to span an integral number of periods. As with the average frequency function, the threshold level can be adjusted to control the Y-axis level used for period determination. The advantage of periodic rms over rms is realized in the evaluation of periodic signals, where inclusion of partial periods in the rms result can cause a large departure from the expected result (Fig. 17).

Hold Min/Max

This group of analysis functions provides a convenient method for monitoring long-term variations of waveforms. The hold minimum function maintains a waveform of the smallest value occurring since the function was initiated. The hold maximum function is identical to hold minimum except that it maintains a time-ordered waveform of largest values. By using these functions together, a pair of waveforms is generated that brackets the extreme values that the source has exhibited.

Acknowledgments

The HP 51089A is the result of the persistence and dedication of many talented individuals. Those involved with the firmware and hardware development were John Fenwick, Carolyn Beck, Scott Titus, Lee Cosart, Don Shremp, Phil Scott, Al Foster, Tom Burris, Ralph Smith, Andy Gong, Nancy Nelson, Nina DeLu, Dick Fowles, Chris Szeto, Jason Cotton, David George, Dana Stoffers, Celia Vigil, and Phil Vitale. Mechanical Engineering was provided by Mike Detro, Ron Keeley, Jim Ammon, and Russ Zandbergen. At the forefront of the firmware testing team were Alice Kwei and Nathan Brown. Production engineers Janet Anvick and Al Scalesi worked closely with the R&D team and did a great job throughout. Special thanks to Jack Folchi, who first proposed the HP 51089A and served as project manager and friend throughout the development of the HP 51089A.

References

1. G.D. Bergland and M.T. Dolan, "Fast Fourier Transform Algorithms," *Programs for Digital Signal Processing*, IEEE Press, 1979, pp. 1.2.1-1.2.18.
2. J.E. Volder, "The CORDIC Trigonometric Computing Technique," *IRE Transactions on Electronic Computers*, September 1959.
3. J.L. Sorden, "A New Generation in Frequency and Time Measurements," *Hewlett-Packard Journal*, Vol 25, no. 10, June 1974, pp. 2-8.
4. *IEEE Standard of Pulse Measurement and Analysis by Objective Techniques*, ANSI/IEEE Standard 181-1977.

Developing a Printed Circuit Board Design System

HP's Printed Circuit Board Design System (HP PCDS) was developed to meet the needs of electronic designers dealing with ever-increasing complexity and density.

by Elaine C. Regelson

TODAY'S MICROMINIATURIZATION of parts and assemblies has made possible products of a complexity and sophistication virtually unimaginable thirty years ago. These products range from Christmas cards animated by tiny chips playing carols to fully functional computers that can be held in one hand. As the parts have increased in complexity and density, so have the boards that hold them. Modern computers can contain boards with more than 20 layers of tiny conductive traces—fine lines of copper etched down to widths of 0.003 inch. Production of such complex boards has required development of vastly more sophisticated manufacturing processes and software tools to assist with design and layout.

In 1982 Hewlett-Packard addressed the computer-aided printed circuit board layout problem with HP EGS (HP's Engineering Graphics System), a product that provides design capture (via an electrical schematic) and general physical design capabilities. In 1986 HP introduced its first fully automated printed circuit board layout product (Fig. 1) as a member of the HP DesignCenter family. This HP Printed Circuit Design System (HP PCDS) is a fully functional computer-aided design application that couples printed circuit board layout to electrical engineering design, manufacturing, and testing. In this issue are a number of articles about HP PCDS, including detailed discussions of some of its features and the development required to bring a large, high-quality product to market. This article provides an overview of the product and the environment in which it runs, and includes a discussion of HP PCDS' Design Module.

Electronic Product Development Steps

There are a number of steps from the conception of an idea for an electronic circuit to its realization as a working product. The engineer sketches out the basic blocks of the circuit and then creates a functional model to test the design concept. Next the engineer refines the design, describing it in the form of a schematic diagram, and perhaps runs a series of simulations to verify that the design is logically correct. More extensive simulations could perform timing and fault analysis.

When the design is complete, a physical layout of the printed circuit board with the components and the connections between them is built. This physical design layout process includes definition of the board outline, assignment of circuit logic to physical devices, placing parts on

the board, electrically connecting devices to one another as dictated by the schematic, and modifying the board design to improve manufacturability. All this must be accomplished in the context of manufacturing constraints and maintenance of design integrity. Once the board has been laid out (Fig. 2a), the instructions to drive the manufacturing machines that build and test the board must be generated. Finally, the board can be built, loaded with parts (Fig. 2b), and tested.

During the development process, the design data must be maintained and archived to provide version control in the event multiple versions of the circuit board are developed, and to secure the design data to ensure that changes are made only by appropriate people, and by only one person at a time. Furthermore, a component library containing accurate and consistent definitions of parts available for use in the products must be maintained.

HP's Electronic Design System (HP EDS) provides the electrical engineer with solutions for schematic capture and digital logic simulation. This product has tight bidirectional data links with HP PCDS and both products share a synchronized parts library. HP PCDS provides a complete collection of tools for the layout process—from accepting the initial logic design all the way through generating manufacturing outputs and links to an HP 3065 Board Test System. The link between HP EDS and HP PCDS allows back annotation and bidirectional engineering change, eliminates redundant data entry, and reduces chances for logical and physical design inconsistencies.

History

HP PCDS is based on source code licensed from Northern Telecom Inc. and its subsidiary, Bell Northern Research, in late 1985. This product had been evolving for over ten years in a mainframe computing environment and was chosen as an excellent foundation for developing a family of Hewlett-Packard products. It provided good functionality for most areas of printed circuit board design and was highly tuned in a number of areas.

The source code was ported to the HP 9000 Series 300 technical workstations to provide single-workstation and networked operation on HP's HP-UX operating system, which is derived from the UNIX® operating system and adheres to AT&T's *System V Interface Definition*, Issue I. In addition to the port, extensive revisions were made to

UNIX is a registered trademark of AT&T in the U.S.A. and other countries.

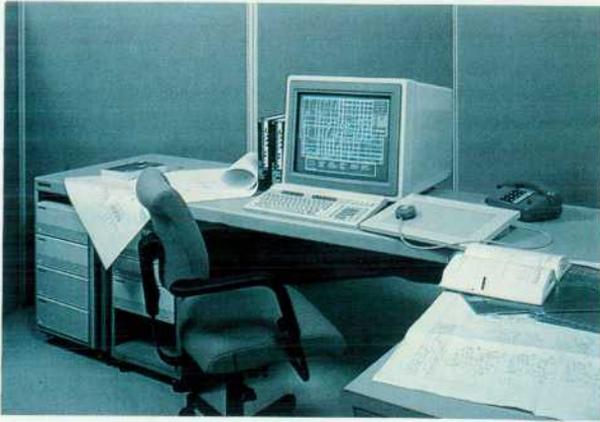


Fig. 1. The HP Printed Circuit Design System (HP PCDS) is a full-function computer-aided design system for the design of printed circuit boards. Links to the HP Electronic Design System (HP EDS) and manufacturing machinery are some of the features.

improve the product, particularly in the areas of design data management, and links to the HP EDS front end and HP 3065 Board Test System back end. In addition to functionality extensions, extensive rework was done to provide greatly increased quality and generality to the product.

In April 1987, HP PCDS was updated to provide new capabilities to address the rapidly changing market and add even more generality for the user community. New workstation platforms broaden the price/performance range of options available to customers. A schematic net editor (SNE) was released which provides very easy entry of netlists for a large group of users who wish to enter their designs into HP PCDS from paper schematics instead of electronically. Significant functionality improvements were made throughout the product, including such areas as design verification and automatic routing.

HP PCDS

HP PCDS is made up of three modules: the Design Module, the Autorouter Module, and the Library Module. These modules reside in the Design System Manager (DSM) environment. DSM is a product that is tuned to manage data and programs for CAD and general-purpose applications. The article on page 71 provides an overview of DSM's capabilities. The article on page 80 is a detailed discussion of the approach used to implement this functionality. In addition to data management, the DSM environment provides for remote program execution, peripheral spooling, spooled remote process management, and network access. The article on page 77 focuses on the special problems of providing appropriate spooling capabilities for a CAD environment.

The Design Module discussed later in this article is the heart of the printed circuit board design system. The Design Module provides board definition, packing, placing, interactive design editing, clean-up, and manufacturing output generation functions. The Design Module contains an extensive set of specific information about how printed circuit boards are manufactured, which it uses throughout the development of the board.

The Autorouter Module automatically routes the board, providing the electrical connections needed for the design. The Autorouter Module and the automatic placement capability of the Design Module are described in the article on page 68.

Board components are maintained in a parts library and managed by the Library Module. The Parts Library provides a data base of more than 8500 standard components from which customers can draw to begin their printed circuit board designs. The Library Module allows maintenance—creation, modification, and validation—of parts in the component data base. The parts library and the Library Module are discussed in the box on page 82.

Product Realization

Porting a software product the size of HP PCDS and bringing it to market quickly and efficiently were challenging processes. Automated tools were developed and special teams formed to complete the job. Special software tools automatically converted the original product language specifics into the HP-UX workstation languages. System dependent foundations of the system, notably the graphics library and file-handling utilities, were redesigned to use the graphics and file capabilities provided by the workstation environment. In addition to language differences, there were occasionally specific problems caused by different uses of machine architecture, such as discrepancies in machine memory allocation and use. Both manual and automated processes were developed to find and correct these discrepancies.

After the product was ported to the HP-UX environment, it was evaluated for opportunities to make it better by using features provided in the new environment. Networking and graphics environments proved to be especially fruitful areas for change. Networked workstations, for instance, provided the opportunity to offer distributed data bases and required additional engineering to solve the special problems this capability poses. An example of the new environment's impact on graphics was the reworking of parts of the interface to take full advantage of the multiple window display capabilities provided by HP Windows/9000 instead of requiring an extra terminal for message display and program control.

To ensure and enhance product quality, small teams were formed that specialized in each of the major functional areas. They were responsible for learning in depth how the area was supposed to work and making sure it still did so after the port. A special test team provided concentrated reliability testing of HP PCDS, particularly focusing on verifying consistency between the software and the manuals describing it. Printed circuit board design specialists also joined the project team and used the product to design a variety of real boards. Their experience enhanced the reliability and usability of the product and has significantly contributed to the list of functional enhancements planned for it.

The article on page 84 discusses in detail many of the special quality assurance teams, tools, and processes used to build and verify HP PCDS.

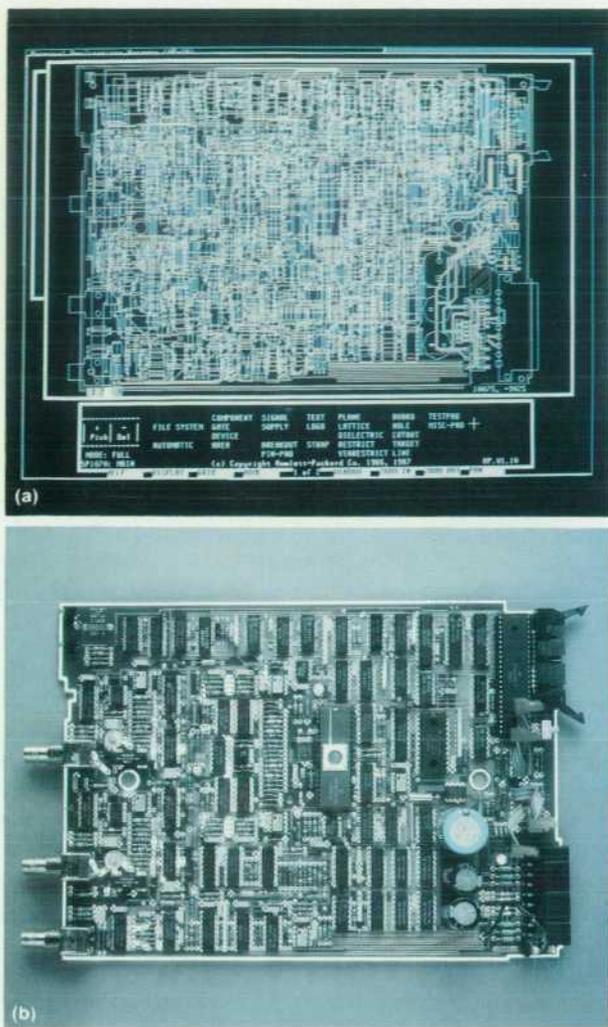


Fig. 2. (a) Layout of board using HP PCDS. (b) Actual board for (a).

Design Module

Most designer interaction with HP PCDS is done through the Design Module. The Design Module provides a set of powerful features to aid the designer in all steps of the printed circuit board and hybrid layout process. These features are accessed through a hierarchically structured user interface.

The user interface is generally object-oriented and easy to use. Picking an object from the display or menu brings up a menu of operations relevant to that object type. Additional menu options allow simple access to such features as zooming in on specific display areas, associating colors with object types, or accessing the file system to save or recover designs.

The first step in creation of a physical design is definition of the board blank shape and the manufacturing technology to be used in fabricating that board. A technology file specifies such items as number and use of layers, default trace widths and spacings, pad definitions, and photoplotter apertures. Technology file information is used throughout the design process to warn of design violations and to ensure that the board will be manufacturable. HP PCDS is

delivered with a number of technology files that serve as useful examples. Many sites, however, may choose to modify these to reflect their own manufacturing processes.

When the board technology and shape have been defined, the logical design information is loaded in the form of a packed, unpacked, or partially packed netlist. HP PCDS contains an automatic packer that associates logical circuits from the engineer's design to physical parts, but still has the flexibility to allow the designer interactive control over the packing of critical parts.

The next step is to place the parts onto the board. The Design Module has a number of features to help the designer interactively place critical parts. These include automatically sorting them for placement in a sequence based on logical connectivity, for instance, and allowing easy movement or rotation of parts. The automatic place and improve functions described in the article on page 68 automatically place parts in their near-optimal positions and swap placed devices and gates to minimize logic length (straight-line distance between connection points) and improve routability.

Interactive routing functions allow electrical connection of parts on the board. The Design Module assists interactive routing by providing pattern repetition capabilities, automatic insertion of vias (feedthroughs) when changing layers, and the ability to add lattices with automatic avoidance and addition of thermal relief pads. Warnings, such as "aperture size not available," notify the designer when design violations or possible manufacturing problems occur. In general, the designer routes critical traces first, then sends the board to the Autorouter Module described in the next article.

A final clean-up, or tidy, function is used to improve manufacturability of the board. Tidy functions make final adjustments to such things as trace spacing, layering, and widths, eliminate unnecessary or unused elements such as traces, breakouts, and vias, automatically spread traces, add teardrops, and widen traces for specified signals.

Before manufacturing the board it is important to reverify the design to ensure that design violations that may have been ignored when they occurred have been taken care of. Design rules check for violations such as too little clearance between board objects (e.g., board edge and elements, vias and pads, trace segments, etc.).

The Design Module generates a variety of manufacturing outputs, including a complete set of printed circuit and hybrid masks, reports for manufacturing and purchasing, and links to photoplotters, numerically controlled (NC) drills, and the HP 3065 Board Test System. In addition, a set of design data access routines (DDAR) is provided. These routines allow custom data extraction to enable customer sites to develop their own links to local manufacturing systems.

Acknowledgments

I thank Jean Jasinski, Jenifer Morrissey, and Rick Stahlin for their patient and careful review and input to this article.

Bibliography

1. HP Printed Circuit Design System Technical Specifications.

Automating the Printed Circuit Board Design Process

Using a computer to place components on a printed circuit board and route most connections greatly simplifies the design process. To accomplish this, careful selection of the algorithms is essential.

by Gary Jackoway

TWO MAJOR REASONS for using a computer-aided design system to design printed circuit boards are to shorten the design cycle and to free the designer from time-consuming, routine tasks. Design automation can play an important role in achieving both of these goals by automatically placing the parts on the board and determining routes for electrical connections and signal paths.

The standard scenario for placing and routing a printed circuit board using HP's Printed Circuit Design System (HP PCDS) is that the user places critical parts manually and then has the autoplacer place the rest. If for some reason the autoplacer fails to place all parts, the user completes the task. Once parts are placed, it is common to try to improve the placement. This can be done manually or automatically at the pin, gate, or part level. When the placement is complete, electrical connections between pins are made by adding traces and vias* to connect all pins carrying the same signal. Routing can be done automatically or manually. Again, it is common for the user to make critical connections manually and then run the autorouter to make the rest. If the autorouter cannot make all of the remaining connections, the user manually completes the board. Routing must take into account the trace width, via size, and spacing requirements of the board manufacturing system.

The placement of parts on the board and gates within parts greatly affects the routability of a printed circuit board. The next section describes the placement method used in HP PCDS. Then placement improvement methods are discussed. The final section describes the autorouter, its parameters and its method of operation.

Autoplacement

The autoplacer determines part locations on the printed circuit board so that the autorouter can completely route the board within the given design constraints. Part placement must consider a number of practical constraints such as board obstructions and part sizes. In determining how routable a board is the standard norm used is logic length—the minimum wire length needed to make all the connections if wires were allowed to cross. Logic length underestimates the final trace length that will be used on the board, because connections must avoid other signals. Since it is difficult or impossible to predict where the traces will

actually go, logic length is used as a predictor of the difficulty of making the trace interconnections. The goal of an autoplacer is to minimize total logic length.

In general, it is not acceptable for an autoplacer to place parts in arbitrary locations on the board. Designers often want the parts aligned to a placement grid because it simplifies automatic insertion of the parts on the board, and assists routing by providing natural channels between parts for connecting traces. Furthermore, if parts are aligned to a grid, there may be a pattern that routes all the connections in one section of the board, as is common for memory arrays.

HP PCDS uses an autoplacer whose basic method is force-directed. In force-directed placement, a mechanical analogy is applied. In Fig. 1, each part to be placed is represented by a block and each interconnection to be made is represented by a spring connecting the two blocks with that interconnection. If all these springs are stretched and

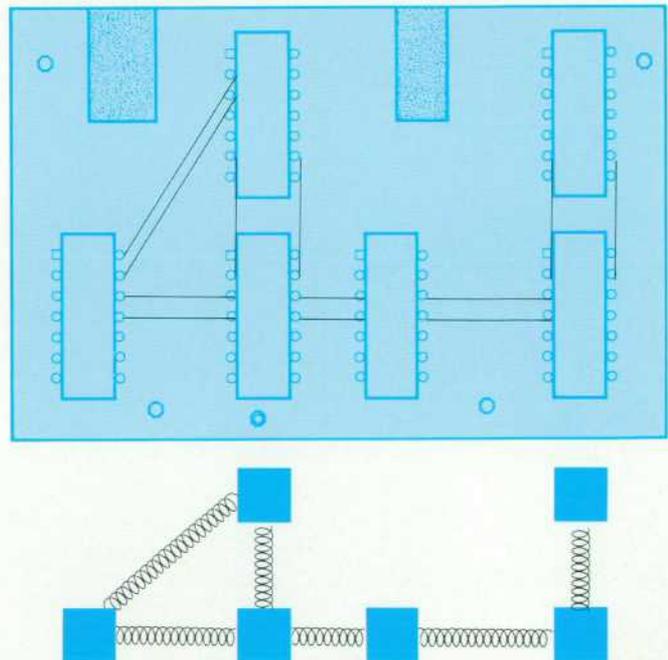


Fig. 1. Force-directed placement is based on a mechanical analogy. (Top) Small printed circuit board. (Bottom) Mechanical representation of board.

*A printed circuit board is often made up of several separate layers of connections. A via is a plated hole in the board that makes connections between layers.

then released, the springs will pull the system together. Now, add repulsion to the system so that all the blocks are not pulled to the center. The autoplacer finds the equilibrium position of blocks in this block/spring problem, and uses the block locations as the part locations on the printed circuit board. Once the placer has these starting locations, it moves each part to a nearby placement grid location in such a way that no parts overlap. When this is completed, the board has been placed. Note that the block/spring problem has been solved optimally, but the placement problem has only been approximated. Finding the minimum total logic length would take much too long, so some simplification must be made.

Improvement

Placement gives a first cut at determining locations for all parts. Since the minimum total logic length in general has not been achieved, there is usually room for improvement. In the improvement phase, the placement is refined. The different types of improvement that can be attempted are:

- **Parts improvement.** Logic length may be shortened when one part swaps location with another. If logic length goes down, the swap is kept, otherwise the parts are returned to their original locations.
- **Gate improvement.** Many parts contain more than one gate. A 74LS00 package, for example, contains four NAND gates. If there are many 74LS00 parts on a board, the gates in all of these parts can be relocated to improve logic length. The user can also specify that improvement is to be tried only within parts. That is, no gate is to be moved to another part, but the four NAND gates within a given package can be reassigned.
- **Pin improvement.** A particular gate may have pins that are equivalent and can be swapped. The two input pins on a NAND gate, for instance, can be swapped with no change in functionality. Pin improvement is more local in nature, but still can ease the routing task by relocating pins on the side of a part where they can be connected more easily.

Pin and gate improvements are straightforward because the small number of possible alternatives allows a thorough search. Part improvement, however, can be a bottleneck. In HP PCDS, part improvement is achieved using a force-directed method. Each part has a zero force location—a location at which the springs pull the part equally in all directions. This location can be considered as the best location for that part. Thus, instead of trying to swap a part with all other parts, the search is concentrated in the neighborhood around the zero force location.

The force-directed method has been shown to be one of the best available. Hanan and his coworkers¹ compared many popular methods on several real-world problems and concluded: "The force-directed pairwise relaxation algorithm . . . yielded the minimum or near minimum result on all six problems. In addition, it accomplished this in the minimum time for the two large problems."

Routing

Routing connections is the most time-intensive step in creating a printed circuit board design. For this reason,

automatic routing is often seen as the critical element in a printed circuit board CAD system. Sales literature abounds with autorouting claims, the most prominent of which is the 100% autorouter—an autorouter that makes all of the connections necessary on the board. It is easy to demonstrate that no practical autorouter lays out 100% of the connections all of the time. An autorouter that tries all possibilities could take years to solve large boards. And, if an autorouter fails to consider all alternatives, it can miss the alternative critical to a specific board. The best autorouter is the one that achieves 100% completion on most boards and gets close to 100% completion on the boards for which it does not achieve 100%, while always maintaining user-specified design rules.

Another attribute of autorouters that has gained much publicity is the choice of internal data structure. Autorouters are divided into two classes: gridded and gridless. A gridded autorouter uses an array as a representation of the board, whereas a gridless autorouter does not. Why would anyone care about the internal data structure of the product they are buying? Because that structure has implications concerning how accurately a board can be routed. Gridded data structures have been seen as inaccurate because object placement has to be rounded to the nearest grid location. A fine grid is possible, but to attain 0.001-inch resolution for a 10-inch-by-10-inch board one would need 100 million grid cells for each layer. Gridless approaches have potential resolutions much finer than 0.001 inch, but so far no purely gridless autorouter has established itself in the marketplace.

The key for the user is not the data structure being used, but the accuracy of the routes that are made. This accuracy has two parts: board space must be used efficiently, but enough space must be left between traces to satisfy the design rules required for that board. The objectives of an autorouter, then, are to achieve the highest completion rate possible in a reasonable amount of time without violating design rules and while using the board space as efficiently as possible.

HP PCDS Autorouter

To achieve high completion rates, the HP PCDS Autorouter Module uses a combination of custom grid specifications, intelligent and flexible connection sorting methods, and user parameter control. The HP PCDS autorouter is a gridded autorouter. It has been carefully constructed, however, to get as much out of the grid structure as possible. Objects are represented accurately from the perspective of traces lying on the tracking grid. That is, the autorouter only routes traces on the tracking grid, so the critical concern is whether these traces can or cannot go to a specific location on that tracking grid. Thus, HP PCDS has the advantage of the proven technology of gridded autorouters without the disadvantage of inaccurate representation. The only limitation placed on the user is that the traces can fall only on grid cells. This limitation is minimized by giving the user control over the grid.

HP PCDS includes custom grid specifications which include parts of the router tuned for specific situations and matched with carefully selected sets of parameters to complement that custom grid. Custom grids are available for

Finding Paths in a Gridded Data Structure

The reason for the popularity of gridded autorouters over the years is that there is a simple algorithm that yields powerful results. The Lee, or wavefront, algorithm involves filling the grid cells with values (flooding) in an attempt to find the shortest path between two pins. The advantages of this algorithm are that if there is a path through grid cells, this algorithm will find it and it will be the shortest one available. Of course, the shortest path uses the least board space and thus leaves the most space available for other paths.

The wavefront algorithm begins by marking the starting cell (source) and the ending cell (target). All grid cells that have already been used are marked as obstructions. In Fig. 1a, source is marked as S, target is marked as T, and obstructions are indicated by X.

Starting from the source cell, mark the empty cells immediately surrounding it with a 1 (surrounding in this example means orthogonally adjacent, not diagonally). The cells marked 1 are exactly one cell away from the source. Now mark with a 2 all unmarked cells surrounding all cells marked with a 1. These cells are exactly two cells away from the source cell. If we continue in this manner until the target cell is reached, the grid looks like Fig. 1b.

At this point we see that the target would be numbered 12, and thus is 12 cells away from the source. Now we find a length-12 path from the target back to the source by following the cell numbers backwards from 12. Note that there are several options along the way. Even in the first step we could go either left or down. The length of the path is not affected by our choice, but the quality could be. A heuristic approach may be applied for selecting a path. For example, start by going down, and then always choose the straightest possible path. This yields the path shown in Fig. 1c.

The wavefront algorithm can be extended in many ways. First, it is possible to add more directions. For a multilayer printed circuit board, vertical directions need to be considered. If diagonal routing is allowed, another four directions are added. Second, there is no need to set the increment (cost) for going one step in each direction to the same value. Instead the path can be controlled by increasing the cost in less-desired directions. Vias can be made more expensive by increasing the cost for moving between layers. Printed circuit boards are designed with layers having a preferred orientation—one layer might be horizontal while another is vertical. The strength of this preference can be controlled by the assigned costs. When costs are different for different directions, the wavefront algorithm no longer yields the shortest path from source to target. Instead it gives the least-

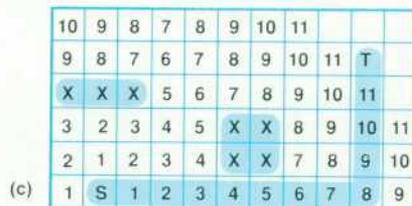
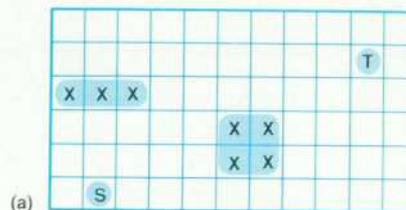


Fig. 1. Wavefront algorithm for autorouting. (a) Grid showing source point S, ending target T, and obstructions X. (b) Marking closest cells with a 1, next closest cells with a 2, and so forth. (c) Choosing the straightest path between S and T.

expensive path, which is more desirable. Finally, there is no requirement that the source or the target be a single cell. Instead, a large network of interconnected pins and traces can serve as source and/or target for the flood. The salient feature of the wavefront algorithm will still be observed: the least-expensive path between the two structures will be found.

It should be noted that the wavefront algorithm gives the least-expensive path only for the one connection it is making. There is no guarantee that by continually making least-expensive connections the board as a whole will achieve the best results. However, in practice the wavefront algorithm combined with carefully selected costs and ordering of the connections results in high completion rates.

achieving one trace between IC pads, for two traces between IC pads, and for boards with surface mount parts (which have a smaller spacing between pads than leaded ICs). The custom grids are optimized to these situations so that the connections made take into account where other traces will need to go. There is also a general grid for use in other circumstances including three traces between IC pads, and simultaneous two traces between IC pads and one trace between surface mount pads. The design parameters control how close traces run to other objects on the board. The two-between grid, for example, will only place two traces between pads if there really is enough room. The autorouter

does not violate design rules, thus saving the user from a clean-up step.

The order of connection plays a major role in how well the board space is used, and thus on the autorouter's completion rate. In Fig. 2, for instance, connecting the pads marked A blocks the connection of the B pads. This simple example suggests that expected connection length is a good ordering of connections. As well as length, sorts are available starting with traces nearest the top, bottom, left, or right of the board. These sorts are critical for aiding the autorouter in locating patterns within the board. A special sort is available that prefers connections that must traverse

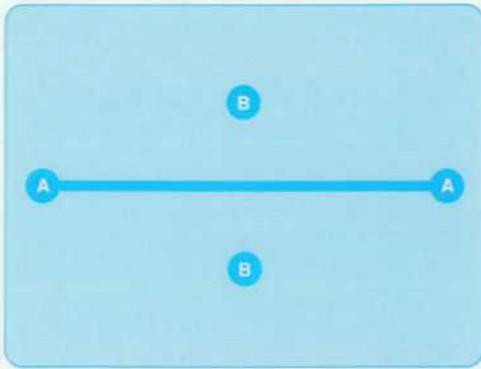


Fig. 2. The order of connection is important in autorouting. Here, connecting the path marked by A first makes it impossible to complete the path marked by B.

congested areas of the board. By routing these first, other traces that do not need to use these critical areas will not inadvertently take these spots.

The user has many other controls over the routing process. A single run of the autorouter is broken into passes which can use different parameter values. For instance, it is common to limit the autorouter's use of vias in early passes. This is helpful because vias tend to be larger than traces and thus block those areas from use by other traces. Also, the user can control how far out of its way the autorouter will go in trying to make a connect. Thus, the auto-

router makes all of the easy connects before adding meandering connects that turn easy connects into hard ones.

The user's control goes to the very heart of the autorouter in setting the cost parameters. The cost parameters tell the autorouter the relative costs to go one step in each direction (including adding vias). If the user has a board where vias are very expensive or if the user prefers not to have any 45° traces, this information can be factored directly into the router's inner loop. Of course, HP PCDS comes with a complete set of default values for strategies to handle most situations so that the user is not required to become an instant expert on autorouters to use the product. Users with demands outside the supplied set can successfully use the control available to them.

Summary

The automatic tools provided to the customer by HP PCDS dramatically affect the total design cycle time by completing in hours tasks that could take a designer weeks. The automatic placement tools can be used individually or as a group to achieve a high-quality placement. The automatic router efficiently routes the board while maintaining the design rules required for manufacturability.

Reference

1. Hanan M., et al, "A Study of Placement Techniques," *Journal of Design Automation and Fault Tolerant Computing*, Vol. 1, no. 1, October 1976, p. 60.

Managing HP PCDS with the Design System Manager

Engineering and design organizations must effectively manage design information to reap the productivity benefits of CAE/CAD systems. The Design System Manager addresses the information management needs in the HP PCDS design environment.

by Paul S. Reese and Mark E. Mayotte

AUTOMATED CAD SYSTEMS like HP's Printed Circuit Design System (HP PCDS) pay for themselves by speeding up the design cycle and thus reducing the time to get a new product to market. But the flood of data produced by automated design tools increases the project management workload and can drag down the net productivity gain from the CAD investment. Engineers and board designers must spend time away from their designs, coordinating efforts and doing administrative tasks that

support their work. As a result, the efficiency of the design environment goes down.

Users also need their CAD solutions to work in a patchwork environment of hardware platforms, software applications, and network links. This further complicates the data management picture.

Engineers and designers use a variety of CAE/CAD tools to do their work. Frequently these tools are linked through a local area network (LAN) to facilitate data sharing. Be-

cause users need access to the same information, they may generate multiple copies of the same files. Or they may move the files they need to their local machines, changing the network location of the files and requiring the next user to seek them out. In such an environment, the design tools must be supported by other tools that record access permissions and the names and network locations of multiple file copies, and provide easy access to applications and peripherals, either locally or across the network.

The Design System Manager (DSM) is an integral part of HP PCDS and provides a flexible solution to its design management needs. DSM addresses four main areas of design management: application integration, network support, file security, and control of multiple versions of files. In addition, DSM includes a spooler system for configuring and accessing peripheral devices (see article on page 77).

The primary contribution of DSM in a design environment is supporting the way a user runs the application programs available on a stand-alone computer or on a network. DSM contains a definition of supported applications that includes all the parameters and operating system commands necessary to initialize and run the application software. The user only has to press a button on a DSM menu.

DSM provides a file-grouping mechanism for the data associated with each application. The groups of files are called filesets. Fileset definitions are unique to each application, and permit users to manipulate a set of files as a logical unit rather than requiring that each file operation be performed individually on each file. This is one way that DSM manages the explosive growth of data created by automated systems.

To make application programs easier to use, DSM manages network locations so that users do not need to know where application files or data bases reside. To an application user, a networked environment appears the same as a single workstation. This is because DSM keeps a list of workstations that will be used for design tasks. This list is cross-referenced in the DSM data base with design information or peripheral devices residing on each workstation. When the user enters a command that requires access to a remote workstation, DSM performs the network connections and file transfers without the user's knowledge and without undue changes in system response.

Project managers may not want users to have equal access to applications or to design information. Security needs will differ depending on company policies and the nature of different designs. Therefore, the security features of DSM are designed to ensure that different sites can establish the appropriate access permissions with a minimum amount of configuration. To accomplish this, DSM allows permissions to be set for different categories of design information. Such permissions apply to all users unless a user is given a specific alternate permission. To configure such permissions, a site administrator chooses the appropriate default, or categorical permission, and then specifies exceptions to these permissions for particular users.

Design History and Version Control

DSM manages files in the context of the design process. During the design and manufacture of a product, three levels of file management are relevant: control and syn-

chronization of designs across manufacturing releases, control of all working files within or before each manufacturing release, and control of working files for each user.

To track files at each of these levels, DSM attaches a version string to each copy of a file or fileset. The version string is incremented with each recorded change or update to the fileset. Therefore, at each level of control a succession of recorded, annotated changes exists as a version stream for the design. The position of a version in the version stream indicates which version is current, which is the previous version, and so forth. If necessary, previous or older versions can be brought forward in the version stream if the current version becomes corrupted.

In addition to a unique version stream, DSM also associates each fileset within a project, that is, it forms a design hierarchy. This structure is used in many engineering organizations. The association ensures that within any project and design the current version of a fileset is known. Only one user at a time is given update access to the fileset.

The project-design-fileset structure and the version string and network information allow DSM to record the location and status of each piece of design information. In this way DSM can provide access across the network to any version of any fileset, provided that the user has the correct access permissions.

Using DSM, therefore, amounts to creating the definitions in the DSM data base for such things as workstations involved in design tasks, applications and related filesets, users and their access permissions, and the names of projects and designs. Once these definitions exist, DSM commands query or update the information as users manipulate information in the design environment. Users can also generate reports on any information contained in DSM.

Like HP PCDS, DSM runs under the HP-UX operating system on HP 9000 Series 300 engineering workstations and Series 800 Precision Architecture computers. DSM currently uses HP Network Services remote file access capability for its network functions, and will soon support the ARPA/Berkeley functionality of the HP-UX system. Although users do not need HP-UX commands to run applications or use DSM functions, DSM does not restrict knowledgeable users from taking advantage of HP-UX communication and system management tools.

DSM Architecture

DSM keeps track of the design environment by recording all relevant information in a set of system tables. The tables are a relational data base of design information. They record the workstations and applications involved in design tasks, user accounts and access permissions, the file structure where information resides, and the network location and revision history for each piece of design information.

Users build, query, and update the system tables by using DSM commands. The same commands that define DSM users or the DSM file structure in the system tables also create the necessary directories in the HP-UX file system and update the necessary operating system files. Only privileged DSM users can create definitions in the system tables, change the file structure, and define access permissions for other users. These special classes of users along with the unique ownership of the DSM file structure pro-

Use of Filesets in HP PCDS

The fileset mechanism provides an important function for data-intensive applications like those used for CAE/CAD. For example, the number of manufacturing files generated by HP PCDS for a printed circuit board is approximately a function of the number of layers in the board multiplied by the types of machinery used to fabricate the board. Thus each layer may have a file for a particular brand of photoplotter, drill tape and drill plot files for a particular numerically controlled drilling machine, and necessary solder mask files.

For a four-layer board, the system produces about 20 manufacturing files in addition to management reports and general documentation such as material lists and checkplots of photoplotter output. Because DSM groups all these files into a single manufacturing fileset with subfilesets for photoplotting, drilling, silk-screening, and other outputs, users can operate on large collections of data with a single command.

The alternative to the fileset mechanism would be less usable and less desirable. Users would need to perform one command for each file, as well as remembering the names of all the files. The fileset mechanism thus bridges the gap between application specific file-naming conventions and the user's view of the data needed to do a particular design task. Filesets allow DSM to keep track of input and output data from each application so that users can concentrate on doing their job.

vide the framework for secure data storage.

Using DSM amounts to defining the information for DSM to manage, then querying and updating that information as design work progresses. Fig. 1 shows the basic structure of DSM.

DSM Networks

In a networked environment, one workstation acts as the central node for DSM information. The system tables on that node contain all the information for the portion of the LAN known to DSM. As users add more workstations to the DSM network, DSM copies the system tables to each new workstation/node. Each of these remote sets of tables contains the definitions of what exists on the network, plus historical and location records for design information that resides on that particular node.

The system tables do not need to be distributed in this manner to support access from other nodes in a network. As long as each node can access the central node, the information is available for DSM operations. But there is a catch—the central node may not always be accessible. Network failures, however rare, prevent users from accessing data, and all design work ceases. One of the original design objectives for DSM was to implement a robust network strategy that would minimize the effects of a network failure. By distributing the system tables, work on each node can continue if the network goes down. Then after full network operation returns, DSM queries the remote tables for changes that occurred while the network was down. If the central tables are different from the remote tables, the user can synchronize the network by specifying which node contains the correct information as each discrepancy is found. This comparison and correction process is largely

automated in a DSM maintenance command.

Interface

DSM commands are HP-UX routines written in the C programming language. Users interact with the commands through Personal Application Manager (PAM) menus running in HP Windows 9000. Each PAM menu lists the available CAE/CAD applications programs, design information that resides at that level of the file structure, and shell scripts that call the lower-level DSM commands. In addition, users can select the shell scripts from the menu, or type them or other HP-UX shell commands on a command line.

Rather than restricting access to the HP-UX system, DSM eases the role of the system administrator by combining system table operations and HP-UX tasks in the same DSM command. This helps the engineer or designer who has the responsibility for setting up and maintaining the system. DSM provides an administrator menu for creating user accounts, adding workstations or peripherals to the network, and creating the file structure. The shell scripts on this menu prompt the administrator to supply information needed by DSM and HP-UX.

Fig. 2 shows how the system creates a user account. The system prompts for the user name, home directory, and home node in a makeuser script and passes this data as parameters to the DSM command that adds the user to DSM. Executing this command writes the necessary entries in the HP-UX `/etc/passwd` and `/etc/group` files, creates a working directory for the user's files, copies the HP-UX environment files into this working directory, and creates an entry in the DSM user definition table. These definitions are repeated on every node in the DSM network. DSM then recognizes the new user, and an HP-UX login exists for the user on every workstation in the DSM network.

File Structure

The DSM file structure provides the basis for file security and for keeping track of multiple copies of design files. DSM partitions the design environment into projects, de-

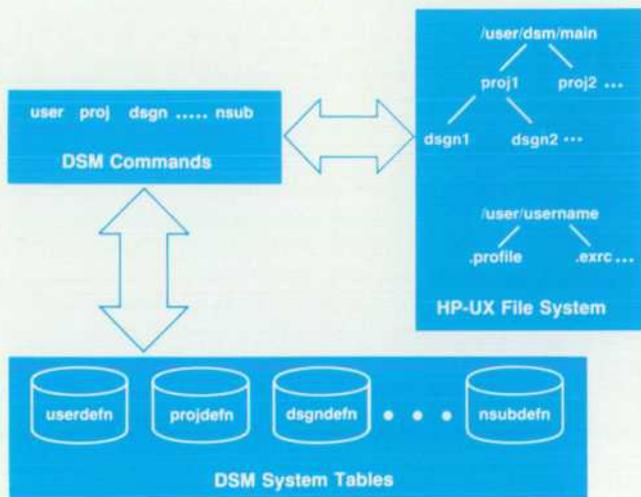


Fig. 1. The Design System Manager (DSM) for HP PCDS consists of system tables, file structure, and DSM commands.

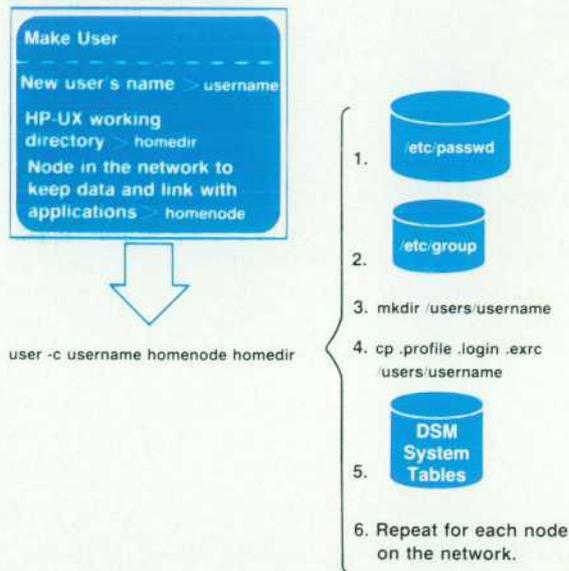


Fig. 2. Creating a user account through DSM.

signs, and filesets (logical groups of HP-UX files whose relation is defined by a particular application). The project level of the hierarchy contains information related by end product, a technology, or common project management. In HP PCDS, each design also contains all the information to design, test, and manufacture a printed circuit board.

The project-design-fileset hierarchy was chosen over other alternatives because it is more easily understood by customers in many engineering organizations and design shops. Before the implementation of DSM, many paper-based data management systems used a project/design hierarchy. The fileset concept was included to add value and functionality to DSM beyond that provided by manual systems. Fileset definitions permit users to manipulate a set of files as a logical unit. The grouping of files into filesets and the names of the filesets are unique to each application. Because of this, filesets can be thought of as chunks of data associated with a subprocess of a CAE/CAD application.

For HP PCDS, fileset definitions correspond to the phases of printed circuit board design. The main fileset groups include a physical fileset for all the connectivity and layout information and a manufacturing fileset for all manufacturing data and supporting documentation. The filesets also contain a hierarchy of subfilesets which allows control of smaller sets of design information. Hence, users can work on different subsets of manufacturing information in parallel. In addition, a security-conscious site can restrict access to whatever subsets of information their needs require. For example, one user may have update access to photoplotter tapes but no access to the silk-screen masks.

The project-design-fileset hierarchy allows users to have an intuitive view of information stored in the system. Project and design names are chosen by the user, and fileset names can reflect specific functions performed by an application. When a user needs access to a piece of information, the user only has to enter the relevant project and design in the DSM file structure and request access to the fileset

that reflects the work to be done. For example, an HP PCDS user working on a board layout can log in, specify the project and design, access the layout fileset, and then run the HP PCDS Design Module program. The designer does not need to know the whereabouts or names of the actual data files in the layout fileset either before or after running the application program.

DSM can support this intuitive user view of design information because it records both the logical relations of projects, designs, and filesets and the locations of filesets on the network. The file structure definitions in the DSM tables are the same on all workstations, so users can log in anywhere on the network without changing either their intuitive view or DSM's logical view of information. When the user requests a copy of a fileset, the network location tables allow DSM to fill the user's request without requiring the user to know where the data physically resides. In fact, for most DSM operations a networked environment appears the same to a user as a stand-alone workstation.

Design Access

DSM includes a special login for an administrative user. This administrator sets up and maintains the design environment and uses this login to create or remove user accounts, network nodes, and the project/design file structure. The administrator also "owns" all design data files and application programs by virtue of HP-UX file permissions.

DSM commands that manipulate design information require a project-design-fileset path to locate any data under DSM control. The project and design of interest are set once as defaults and then referenced as each command is executed. DSM prompts the user for a fileset to update or view. In the HP-UX file system, the project-design-fileset structure logically resides in the home directory of the DSM administrator user. Users cannot view this file structure and have no access to design information kept within.

To update a piece of design information, users must go through a checkout/check-in process. Checking out a fileset from DSM is similar to checking a book out of the library when the reader cannot find the book on the shelves. The reader consults the librarian, who determines if the book is already checked out, or if anyone else has reserved access to the book. If the book is available, and the reader has a library card, the librarian locates the book and allows the reader to take it from the library. While the book is out of the library, its condition is the responsibility of the person who checked it out. For example, a public-spirited reader may "update" a book by mending a page torn by a previous borrower.

When a designer requests access to a fileset, DSM queries a status table to see if the fileset is already checked out. If it is not, and the requesting user has the correct access permissions, DSM places a copy of the fileset (a copy of the data files represented by the fileset name) in the user's working directory, and records in the status table that the fileset is checked out from DSM and is therefore unavailable for update by other users.

The user's working directory provides a common location for DSM to place requested filesets and for application programs to find input files and return output files. Because

Version Strings

DSM attaches a unique version string to each copy of a fileset. The checkout/check-in process increments the version string to reflect new manufacturing releases, new versions within the manufacturing release, and different working copies of a fileset. The following steps illustrate the changes to a version string associated with a typical design cycle:

1. Initialize the design for a new manufacturing release (A.00.00).
2. Check out the design to begin layout (A.00.01).
3. Save a working copy (A.00.02) before making experimental changes.
4. Experiment (A.00.02) did not work out, get saved version as the new current version (A.00.03), and complete work.
5. Design complete, check in current version (A.01.00) and lock to prevent further changes—all temporary versions are automatically removed.
6. Verification shows minor design flaw, unlock and check out for correction.
7. Correction complete, check in and lock new current version (A.02.00).
8. Verification passed, synchronize to last A-level version number (A.**.**) and prepare for next manufacturing revision (B.00.00).

the working directory is known throughout the DSM network, users can run applications from any network location. This is one way that DSM integrates the working environment for different application programs while letting users run their applications from a common menu.

When the user is finished working on the fileset, it must be checked in. This returns the fileset to DSM control and updates the central copy kept by DSM (unlike a library, DSM always keeps a copy of a checked-out fileset, and updates this copy only at the user's request). DSM then removes the fileset from the user's working directory, places it in the DSM file structure as the current version, and changes the status table to indicate that the fileset is again available for updating.

Every time a user checks out a fileset to work on it or checks it back in, a new entry is made in the status table for that fileset. As a result, DSM knows when someone is updating a fileset. During such a time, other users can get a read-only copy of the same fileset for reference. They can change this reference copy, but DSM will not update the central copy to reflect these changes unless the user performs commands that require a more intimate knowledge of how DSM works. In this way DSM keeps designers and engineers from stepping on each other's work—they cannot simultaneously update the same fileset in the same design. On the other hand, more advanced users can choose to ignore this protection scheme. The detailed interaction with DSM required of these users ensures that they are aware of the consequences of their actions.

The status table not only records accesses to a fileset to prevent simultaneous updates, but also serves as a log of all accesses made to a design. Each version created by check-in and checkout operations creates a new entry that contains the name of the user that performed the operation, the node in the network on which the files reside, the

version string that uniquely identifies the files operated upon, a comment specified by the user, and a time stamp. This information acts to keep a complete history of design progress, and is readily available through reports to managers and others with permission to view the data.

File Security

Like the design information managed by DSM, the security mechanism is defined in the system tables and is tied to the file structure. The file structure involves a hierarchy of projects and designs, and filesets within designs that point to HP-UX files. One advantage of this structure is that user access permissions can be specified at any level of the hierarchy. Also, permissions set at higher levels of the hierarchy govern access to lower levels. Default permissions can be set for the entire network of workstations defined in DSM, and also for each project and each design. In addition, specific permissions that act as exceptions to the default permissions can be set for individual users. The allowed permissions are update access, read-only access, or no access.

Thus, users have access to design information according to the file structure permissions and to HP-UX file permissions. The HP-UX permissions provide unlimited access for the DSM administrator user who owns all files in the DSM file structure. This user can define the DSM network, set the network-wide access permission, create user accounts, and designate other users as project managers. Project managers have unlimited access to commands and data within the projects they manage. They also can define the access permissions for other users for all designs and filesets within their projects. Other users have access to design information according to the permissions set by the administrator and project manager users.

The goal of the DSM security scheme is to be flexible enough to meet the security needs of a variety of design environments, and to allow each site to perform a minimum amount of security setup. For example, the network access permission can be set to update, read-only, or no access. A global update access may be appropriate in an open design environment where most of the designers work on most of the designs. In this case, access permissions at lower levels of the file structure need not be set. In a more restricted environment, design information may be classified. The appropriate global permission may deny access to all data and the lower-level permissions may specify update access for a particular user who has clearance to work on a particular design or fileset within a design. DSM thus provides access permissions that are appropriate to the company environment or the nature of the designs.

Version Control and the Design Cycle

DSM manages files in the context of the design process. Three levels of file management are relevant in the process of designing and manufacturing a product: control and synchronization of designs across manufacturing releases, control of all working files before manufacturing release, and control of working files for each engineer or designer.

For each level of control, the check-in/checkout process establishes different file versions. When initializing the DSM structure for a new design, a project manager in effect

checks out the files in the design and makes them available for work on a new manufacturing release. This establishes a unique version string that will be associated with all filesets in the design. Within the manufacturing release, designers check out and work on individual filesets. Because the designer can go through several iterations of an automated design task before getting satisfactory results, the designer may want to save working copies of the fileset. Each of these working copies is a temporary backup copy of the checked-out fileset. Such backup copies are useful while the designer is experimenting with alternate strategies. However, when the designer gets the desired results and checks the fileset back into DSM control, DSM purges the temporary copies from the system.

The control of design versions across manufacturing releases, within manufacturing releases, and for temporary working copies allows for saving or locking the current state of the design at major milestones in the design cycle. The largest milestone is established at each manufacturing release. Because DSM prohibits changes to versions of files already released to manufacturing, these released files should be archived elsewhere to save file space in the design system.

In addition, the filesets for each manufactured design should be available for reference in the next release. DSM provides this capability by synchronizing all filesets with a unique version string in the process of initializing the next manufacturing release. The version string identifies the current version of all information in the design, that is, the latest version sent to manufacturing.

Within a manufacturing release, milestones can occur that require verification and approval before continuing to the next step in the design process. For example, after completing a printed circuit board, a site may require that a pilot production run be made and tested before the board is released for volume manufacture. The designer checks in the board design and the project manager locks the fileset to prevent unauthorized change to the layout while the pilot run is completed. The locked design is still available for reference (to permit the creation of documentation and manufacturing files), but changes are not permitted.

Working copies of a design become the designer's personal version stream. Milestones at this level are what the designer considers a meaningful point to save work in progress, and may also depend on special application features. For example, a designer could choose to save working copies of the layout after each iteration of the placement-improve function. The designer could then choose which of the working copies contains the best placement strategy.

DSM handles design versioning, so the user does not need to crowd the working directory with multiple design copies, or worry about overwriting a previous version of the design. DSM automatically creates a new directory for each fileset version within the file structure. The user has the option of keeping these versions on-line, archiving specific versions, or physically removing unwanted versions.

Archive schedules depend on the amount of available disc space and the amount of data the user will risk losing in the event of disc or file system corruption. Users can archive specific versions by using the unique version string as an identifier, or by using relative identifiers that refer-

ence the current version, the previous version, or the oldest version that has yet to be archived. To save work in progress, users can also archive temporary working copies of checked-out filesets. A design history log shows the status and current location of filesets, including backup copies, archived versions, and, if requested, network locations of data files.

Spooler System

Because a CAE/CAD environment includes special processors and peripheral devices, DSM provides a spooler system to support devices other than the line printers supported by HP-UX. DSM records information about peripheral devices in the system tables and secures spooler system files in a special project in the file structure. As with other applications, users access the supported devices through a consistent menu-driven interface. Default choices in the menu scripts make it possible to use peripherals without knowing their network location or command languages.

Users can customize the spooler system to include devices for particular applications (for example, simulators, routers, or paper tape punch machines). The existing devices can also be modified as necessary. For example, some photoplotters only accept 9-track magnetic tape files written in a particular format.

Customization and Usability

Other DSM functions also provide the opportunity for customization. The commands that call each application or check out and check in designs include a parameter that will run a user-defined script for application or site-specific preprocessing or postprocessing. The application integration capability is discussed in the article on page 80.

The goal of DSM is to manage the design environment and let users do their work unencumbered by administrative details. A CAE/CAD system should not require users to know operating system or network commands, and a design management system should not require extensive additional training for application users or get in the way of application use. DSM runs in a separate window from HP PCDS and other applications. DSM functions such as generating working copies of filesets or archiving designs are available by shuffling windows. Moreover, when accessing DSM, users do not lose their place or need to retrace a menu structure to return to the application they were using. Because of the way HP-UX allows concurrent processes, the user can perform DSM functions while an HP PCDS process continues as a background operation in its own window.

Acknowledgments

The authors would like to acknowledge the contributions of Jon Gustafson, Colin Clement, Debbie Lienhart, and Eric Widholm in the design and implementation of DSM. Jeff Alley, Dave Martin, Jafar Nabkel, Jenifer Morrissey, Tom Rideout, and Rick Roeling contributed to the verification of the functionality and usability of the system. George Williams, Harvey Clawson, and Gary Jackoway supplied ideas and, when required, code for some parts of the system.

A Multidevice Spooler for Technical Applications

The variety and complexity of shared devices for CAD systems such as one used for printed circuit design require a flexible spooler with a common access method for all applications.

by Deborah A. Lienhart

TECHNICAL SOFTWARE APPLICATIONS produce files that are sent to peripherals, including printers and plotters. As a physical design system, the HP Printed Circuit Design System (HP PCDS) produces output that goes to a wide variety of physical devices. This output includes manufacturing reports to be printed, check plots to be plotted, photoplot files to be stored on a nine-track magnetic tape, and numerically controlled (NC) drill tapes generated by a tape punch. HP PCDS's Design System Manager (DSM) also uses a cartridge tape drive to archive design files. In addition to sending files to physical devices, it is necessary to send files to the Autorouter Module (the printed circuit board router), which is a batch process.

The target machine for HP PCDS, an HP 9000 Series 300 technical workstation running under the HP-UX operating system, provides a spooler for line printers and several methods of writing to tapes, but it does not provide direct support for plotters or tape punch machines, nor for the spooling of processes. Since HP PCDS users are printed circuit board designers, not necessarily computer programmers, we decided to provide a spooler that would support and manage each of the devices necessary for HP PCDS, using a common access method.

Requirements

Some requirements are common to all of the output devices. There must be a common user interface for sending jobs to and getting the status of each of the devices. A common user interface lets the user first learn how to interact with one device and then use that knowledge to access other devices. The spooler must be able to recover from device problems, such as the plotter being turned off or a tape write error. Sending jobs can take a long time if the files are large, which is especially infuriating if a job must be sent again because of a device error.

Users must be able to add devices to the spooler and/or customize the existing devices. The new and/or modified devices may be like devices already supported, such as another plotter, or they can be completely different. The user must be able to send a job to and get the status of a device on any node in the network subset.

Further investigation reveals that each of these devices has special requirements to be considered. Because a line printer spooler (called `lp`) is shipped with the HP-UX operating system, a new spooler must be compatible with `lp so`

that both spoolers do not send files to the printer at the same time.

The HP-UX commands to save data on tape store the name of the file on the tape along with the contents of the files. With some formats, directory hierarchies and file ownership information are also saved. Although it is often easier to keep track of the files if they can be given new names meaningful to the spooler, the spooler must maintain both the user's file names and the user's directory structures.

Some devices, such as tape drives and some plotters, require time between jobs for the user to change the tape or plotter paper. A spooler must supply a mechanism for the user to let the spooler know when the device is ready for the next job.

Interface programs had to be written for the plotter and tape punch since none are provided by HP-UX. Finally, the HP PCDS Autorouter Module must be able to send the results of the route process back to the user.

Spooler Processes

The DSM spooler system consists of three types of processes: the daemon (an HP-UX term for a noninteractive background process), the drivers (processes that manage devices), and the scripts (an HP-UX term for files containing operating system commands).

There is one daemon in the DSM spooler system. At regular, user-defined intervals the daemon checks to see if new jobs have arrived. When a new job arrives for a driver, the daemon wakes up the driver and lets it know that the job is waiting. The daemon also processes commands for the spooler system issued by the user.

There is one driver for each device. The driver maintains the job queue and status information for the device. When the daemon tells the driver that a new job has arrived, the driver adds the job to its queue. If the device is not busy, the driver executes the script for the device; otherwise, it waits for the device to become available.

The script actually sends the job files to the device. The scripts for the physical devices are Bourne shell scripts that use HP-UX commands such as `tcio`, `dd`, or `lp` to send the output to the device, or use the interface programs for the plotter or tape punch. When the driver in the Autorouter Module has a job to run, it executes the printed circuit board router.

The spooler processes communicate with each other through the use of mailboxes. There is a different mailbox for each pair of communicating processes in each direction of communication. Fig. 1 shows the spooler processes, their communication paths, and the mailboxes involved. In general, communication flows one way between processes. The only exception is the communication between the router driver and the router. In this case, two-way communication is used to negotiate the conclusion of an Autorouter Module session. This ensures that the router can produce a results file that is complete and ready to be sent back to the user.

The mailboxes are implemented with HP-UX shared memory, semaphores, and signals. Shared memory is a section of memory that can be used by multiple processes. In the spooler processes, the shared memory contains the command, which is the contents of the mailbox. A semaphore is similar to shared memory in that it can be accessed by multiple processes, but the semaphore operations guarantee that only one of the processes accesses the semaphore at a time. In the spooler, a semaphore is used to indicate whether the information in the shared memory is valid, or whether the mailbox is full or empty.

Signals are system-supplied mechanisms that allow one process to interrupt another process. In the spooler, a signal is sent to tell a process to check the mailbox for incoming mail.

The protocol for the mailboxes is the same for all communication paths. For example, when the user issues a spooler command, the command process first checks to make sure that the daemon's incoming mailbox is empty. When the mailbox is empty, the command process puts the command information in the mailbox, marks the mailbox as full, and signals the daemon that there is mail waiting. When the daemon gets the signal, it checks to make sure that the mailbox is full, takes the command information from the mailbox, and then marks the mailbox as empty. For communication paths where there is only one possible message, only the signal part of the mailbox and protocol needs to be used. This is the case when a script process needs to tell the respective driver that it is finished.

Spooler Job Flow

Fig. 2 shows the job flow through the spooler system. A spooler job (#1 in Fig. 2) is a directory that contains the data files for the job and the driver command file. The driver command file contains the date, the user's name, the directory and node from which the files came (usually

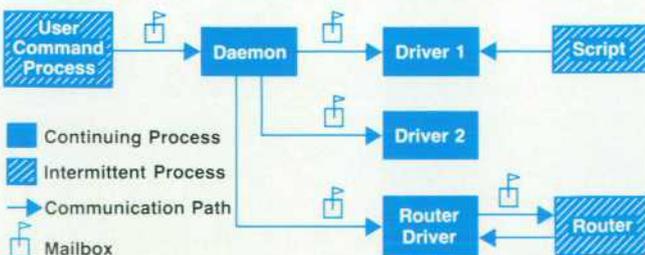


Fig. 1. Spooler processes, communication paths, and mailboxes.

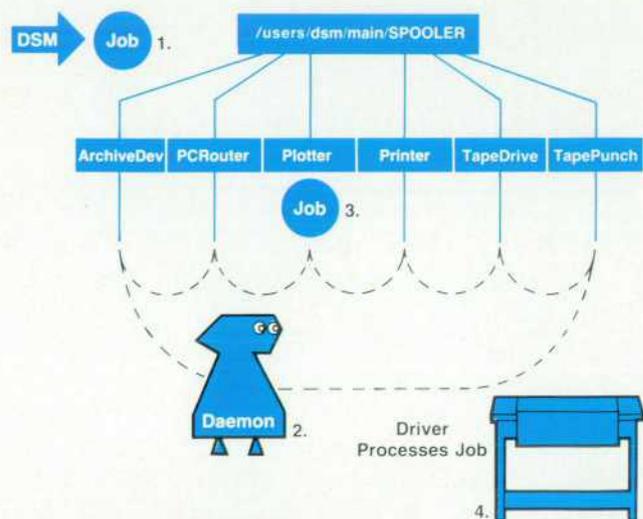


Fig. 2. Job flow through the DSM spooler system.

the user's home node and home directory), the design name, and a message string. Since each job is contained within its own directory, the data files can keep their original names and directory structure without worrying about naming conflicts with other jobs.

When the job is in the process of being created, its HP-UX directory permissions are set to a state that the spooler sees as locked. When the job is ready for the spooler, its HP-UX directory permissions are set to an unlocked state. The daemon checks each driver's directory to see if there are new jobs (#2 in Fig. 2). The spooler sees the new job and wakes up the responsible driver (#3 in Fig. 2). The driver changes the name of the job directory to show the job's position in the job queue, creates a string from the driver command file that will be used to identify the job, and then adds the job to the status file. When the device is available, the first job in the queue is executed (#4 in Fig. 2).

Several commands can be issued by a user to interact with the spooler system. Some of these commands change the configuration of the spooler system while the spooler is running. These commands are starting or stopping a driver, and changing the polling interval. Besides the configuration commands, there are commands that tell the spooler to do something specific. With these specific commands the user can cancel a job, tell the driver that the device is ready for the next job, reset the driver after a device error, and tell the daemon to poll once for each of the drivers.

Customizing the Spooler System

The initial configuration of the spooler is determined by the daemon configuration file. This file is a normal text file so that it can be easily modified. The configuration file contains the polling interval and information for each driver. There is a spooler command that allows users to add drivers to the spooler without editing this file directly.

The configuration information for each driver includes the name of the driver's directory, the path to the driver code, the self-start bit, and the active bit. The self-start bit specifies whether the device can start jobs without operator

setup. If the device cannot start jobs without operator setup, the driver puts a message in the driver status file that says that the device needs attention and waits for the operator to issue the `device_ready` command. When the driver gets this command it runs the job. If the device does not require operator setup, the driver starts the job immediately.

The active bit says whether the driver should be started when the daemon starts. The daemon can start drivers that are listed in the daemon configuration file, so if the user does not always want to run a particular driver, the active bit can be set so that the daemon does not start that driver. The information will still be in the file in case the user wants to start the driver later. For example, if a plotter is not usually attached to a node, but there is one occasionally, this bit can be off.

The driver scripts are Bourne shell scripts which can be customized by editing them. For example, there is a line in the `TapeDrive` script file that writes files to the nine-track tape drive. The user can change this line of the file to change the block size or format of the tape. If the user wants to add another plotter to the system, the driver script for the new plotter can be created easily by modifying a copy of the script file for the existing plotter. To add a completely new device such as a numerically controlled drill with an RS-232-C/V.24 interface to the spooler, only a new driver script and possibly an interface program would need to be written.

Spooler and DSM

The Design System Manager (DSM) manages design information for CAE/CAD applications and is included with HP PCDS. DSM maintains the design information in a hierarchy of projects and designs. The menus used by the 1.0 and 1.1 versions of HP PCDS, DSM, and the spooler are also based on this hierarchy. Fig. 3 shows a DSM hierarchy of projects and designs with the spooler command menus.

The spooler system has its own special DSM project called SPOOLER. The spooler commands can be found in the SPOOLER project directory. Some of the spooler commands can also be found in the Spooler menu in each DSM

design directory.

There is a DSM design for each driver in the SPOOLER project. The driver's job queue and status files are maintained in the driver design directory. The `send_job` and `driver_status` commands for each driver are found in that driver's design directory. These commands are also found in menus with the same name as the driver under the design directory for each DSM design.

The `send_job` and `driver_status` commands provide the interface to the DSM spool command. The spool command can create spooler jobs on any node in the network subset and can copy files from any node in the network subset. The spool command uses the DSM system tables to gather information about the user who is sending the job, the files that are included in the job, the location of the files, and the location of the device. For example, if the user is sending a job to the Autorouter Module on the router workstation, DSM can supply the names and locations of the input files, and the name and location of the router workstation. DSM also opens any necessary network connections and copies the files to the router's driver design directory on the destination node. If the routing is to be done on the local node, the spooler is notified that a new job has arrived. All of this is done just by picking the `send_job` command in the PCRouter submenu of the menu for the current design and choosing all of the defaults. When the job is finished, the router driver uses DSM to send the results file back to the user. Of course, the job can be processed on a different node or use different input files if the default conditions are not chosen.

Summary

The DSM spooler system provides a method for accessing each of the peripherals and the Autorouter Module through a common user interface. This lets HP PCDS users concentrate on designing printed circuit boards instead of interacting with each device individually. The DSM spooler can be easily customized. This allows customer sites to configure the system so that it works with their peripherals and manufacturing processes.

Because the spooler is integrated within DSM, the user

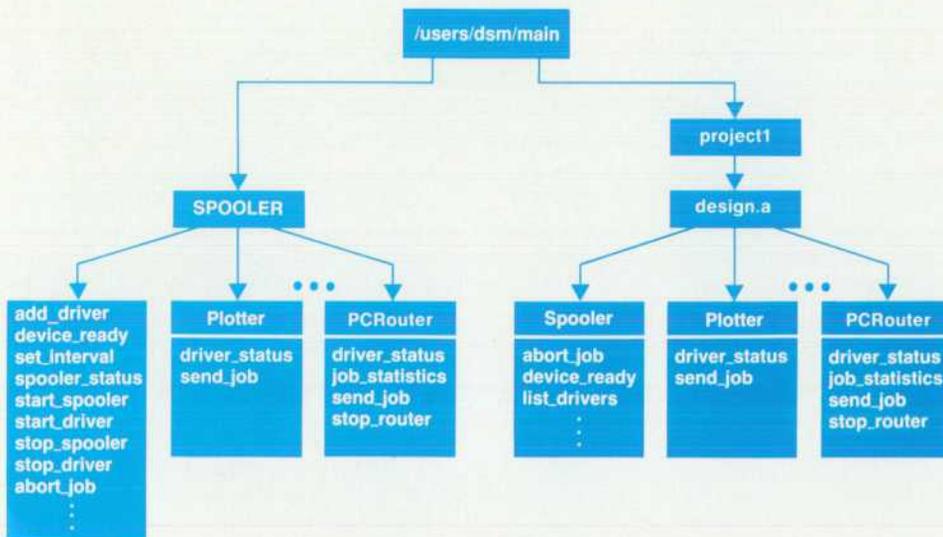


Fig. 3. DSM project and design hierarchy with spooler menus.

does not need to specify information that is already known by DSM, such as the names and locations of files. In addition, DSM can be used to transfer job files and status information across the network.

Acknowledgments

The plot program was written by Don Hoffman, the tape punch program was written by Colin Clement, and the router interface was written by John Kessenich. Dave Martin, Mark Mayotte, and Gary Jackoway provided inspiration during design and moral support during implementation.

Integrating Applications in a Design Management System

The many data files and applications involved in a CAD system require a sophisticated management system to preserve data integrity, prevent conflicts, and maintain design documentation.

by Mark E. Mayotte

A DESIGN MANAGEMENT SYSTEM that manages versions of and access to design information satisfies the basic requirement of keeping a company's data organized and protected. If the applications that create and maintain the data are excluded from the management system's control, data integrity is jeopardized. HP's Design System Manager (DSM) exceeds these usual file management tasks by providing application integration functions. These capabilities include automatic invocation and environment initialization and serve as the primary mechanism for user customization and expansion of DSM.

Design System Manager was developed in conjunction with HP's Printed Circuit Design System (HP PCDS) to allow an organization to control the access to and the distribution of their CAD data. The basic functions of a design manager are to provide the file management tasks of versioning, access control, archival storage, and reporting. If applications are not integrated into a design manager, the link between it and the huge amounts of data produced during the design process must be the user.

DSM provides an open, intrinsically customizable application management capability that integrates the applications used (in the case of HP PCDS, schematic capture and simulation, circuit board design, and generation of manufacturing outputs) into its management environment. DSM is a general-purpose data management environment and is not limited to supporting applications that currently use its capabilities. Any data-intensive application implemented on the HP-UX system can benefit from integration into DSM. Two types of integration are supported: environment initialization and automatic execution.

Environment initialization uses information in a data base maintained by DSM and optional preprocessing and

postprocessing information. The stored definitions are used to interface the CAD applications with the operating system and information kept in the design data base. The benefits of this form of integration are ease of use and controlled access to applications and their data.

An example of this feature's use by HP PCDS is the PCDesign printed circuit design layout editor. To execute the editor, 16 parameters must be specified. These parameters include the locations of the editor's system files, working files, and component data bases. The data bases and working files are maintained by DSM, and the locations vary according to the design in progress, the user performing the work, and the network configuration. The environment initialization capability enables DSM to supply this information to the application.

Automatic execution allows applications to be executed as a side effect of routine data management operations. The benefits of this second form of integration are the enforcement of a site's design process rules, ease of use, and customization of the design manager. This customization ability also permits DSM to be less restrictive in its default configuration. A site need not add any automatic processing if a relaxed development process is desired, but more-rigid systems can be supported.

The automatic execution feature is provided primarily for DSM users and is not currently used by HP PCDS. A possible use is the mailing of a notification to a project manager when a designer checks out a design for initial work. A more powerful use is the automatic execution of a design rule checker when the design is checked in—if the design does not meet site manufacturing standards it cannot be checked in until the errors are corrected.

Environment Initialization

Large complicated applications, especially data-intensive ones such as CAD applications, require nontrivial initialization, invocation, and postprocessing just to get them to work. Add these operating system requirements to the rules that a site requires and the result is a complicated set of incantations that the user must learn before any useful work can begin. If all of the HP-UX operating system commands and commands necessary to follow conventions used at the user's site can be captured in one place, the user is freed from having to learn more than what is necessary. A user knowledgeable in the operating system can use the mechanism that DSM provides to capture this information and manage its distribution on the network. These rules are stored in DSM's data base and customize its operation to the needs of the site. By including this capability, DSM frees each site to concentrate on what must be customized instead of how to manage the result.

Initialization and cleanup are accomplished through the use of preprocessing and postprocessing Bourne shell scripts (collections of HP-UX system commands and simple control constructs). These scripts are executed before and after the application they surround, which provides points of control where the required incantations can be performed. Separate scripts are used to avoid the need for users to alter files installed as part of a supported product.

These preprocessing and postprocessing scripts are called triggers. The name is taken from the manner in which they are automatically run when the user requests execution of the related application, thus triggering their execution. DSM provides for the association and disassociation of triggers with applications and the distribution and management of trigger files across a network. A proprietary data base stores the triggers associated with each application.

In addition to providing for site customizable preprocessing and postprocessing via triggers, a link to the information maintained in DSM's data base is necessary to tighten the link to an application further, such as between HP PCDS and DSM. The HP-UX operating system has the ability to pass parameters to a program from the command interpreter. This parameter passing capability is used in conjunction with a simple query mechanism (see page 82) to communicate information from the DSM data base applications. Most of the information passed between DSM and HP PCDS applications concerns the location of files under DSM's control. These files are component data bases and user-specific working areas. It is not enough to supply just the locations of the files since they can reside anywhere on the network. If a file does not reside on the local file system, DSM automatically initializes a network connection between the local system and the node on which the files reside. It is through this ability that a degree of network transparency is implemented; HP PCDS need not be network smart to make use of the distributed file system provided by HP's NS/9000 remote file access capability.

Automatic Execution

Once the definition of an application is available to DSM, it can be used at the request of the user, as outlined above, or automatically. By associating the application with a project or design in the DSM structure, operations on that DSM

structure will trigger execution of the application. Event-driven automatic execution of this kind enables a site to enforce the design discipline dictated by its needs.

Optional preprocessing and postprocessing triggers can be associated with a project or design. Whenever a check-in, checkout, or other operation that manipulates versions and provides access to data is requested, any existing triggers are executed. For example, consider a printed circuit design group that uses a formal sign-off procedure when releasing a design to manufacturing. An application has been written that records whether or not each approval has been obtained. When the designers are finished, their manager locks the applicable files to prevent modification during the verification and final sign-off process. (The lock operation could have attached a postprocessing trigger that would mail a notification to each party whose approval is needed.) A second trigger is placed at the front of the manufacturing release command.* This trigger checks the sign-off record and fails if any of the required approvals is missing, thus aborting the attempt to release the design. The sign-off checker accesses a file that is maintained by DSM. Since the file can exist anywhere on the network, the first type of application integration described above is used to package all of the commands necessary to execute the sign-off checker. Pseudocode for the preprocessing trigger to the manufacturing release command is shown below.

Both types of integration are shown. The sign-off checker program was defined in DSM, thus eliminating the need to perform operating system incantations in the preprocessing trigger, and the automatic execution of the checker prevents unapproved designs from entering the manufacturing process.

- If a ReleaseDesign operation: Use `dsmappl` (discussed later) to execute the sign-off checker storing the success indicator in the variable status with any messages saved in a file called results.
- If status is failed: Send mail to the appropriate manager indicating the failure with the results file attached for details, abort the ReleaseDesign operation.

The Bourne shell script is:

```
# The init_version command calls the DSM command ckinit. The
# -r option to the ckinit command indicates a manufacturing release
# level initialization. The test checks these two values.
# The script assumes that the proj_mgr user is responsible
# for acting on failed manufacturing release attempts
```

```
if [ $cmd = ckinit -a $1 = -r ]
then
  dsmappl sign_off_checker ) /tmp/results
  if [ $? != 0 ]
  then
    echo "FAILED SIGN_OFF_CHECKER:" ) /tmp/header
    cat /tmp/header /tmp/results | mail proj_mgr
    rm /tmp/header /tmp/results # some cleanup
    exit 1 # exiting the trigger effectively aborts command
  fi
fi
# end of preprocessing trigger, ReleaseDesign operation follows
```

*DSM has a command called `ReleaseDesign` that is used to submit a design for manufacturing release. This command initializes a new stream of versions separate from the earlier stream and prevents any changes to the earlier stream.

DSM Query Facility

The DSM command that handles the application interface is called `dsmapl`. When creating an application reference the customizer specifies:

- The name that DSM will use to refer to the application
- The program to be executed
- The preprocessing trigger
- The postprocessing trigger
- A definition for each positional parameter to be resolved from the data base and passed to the application.

The `dsmapl` process gathers the positional parameters and constructs the argument vector used to execute the application as a child process. If preprocessing and/or post-processing triggers have been attached to the application, an environment must be set up that executes the application and the triggers in one process. Certain HP-UX environment initializations are not inherited across sequentially executed processes. To avoid this loss of environment, a

new command must be constructed.

The new command is begun by first writing the preprocessing trigger to a temporary file. A command is constructed that calls the application using the information supplied in the definition and this command is appended to the end of the temporary file. The postprocessing trigger is then appended to the end of the temporary file. The result is a Bourne shell script that can be executed as the child process of the `dsmapl` call.

The constructed command is executed using the HP-UX `fork/exec` procedures and `dsmapl` waits for its child, the application that the user requested, to complete. `dsmapl` stays around to maintain the environment that it initialized while gathering the positional parameters, the most important component being the initialized network connections.

The types of parameters that can be specified with the applications definition in the data base are:

- `project`, the name of the current project

HP PCDS Library Module

The Library Module (LM) is the data base management system for HP's Printed Circuit Design System (HP PCDS). The major contributions of the Library Module are its speed and its large data base containing over 8500 component definitions. The Library Module provides the capability to enter, modify, and store component information required by HP PCDS for printed circuit board design. This information is organized and stored in three data bases: `USER`, `PROJ`, and `CORP`. The three data bases are searched in hierarchical order for component information, starting with data supplied by the user/designer, progressing to project data, and finally moving to corporate (or master) data.

Component information required in the design process can be extracted from data bases on the local node or remotely over the local area network (LAN). Data base access across the network is made transparent to the user through the Library Module's tight interface with the HP PCDS Design System Manager (DSM). The Library Module does not limit technologies (e.g., surface mount devices or thick-film hybrids) that can be used in HP PCDS. The Library Module data bases are in a multiple-reader, one-writer environment that allows multiple applications to read the same data base at the same time and allows one application to write to the data base while other applications are reading.

A forms editor lets the user enter or modify component data. HP's logical layout system, DCS, can be used to enter logical-to-physical mapping information for pin, gate, and group swapping. This information can be defined in the Library Module or defined in DCS and transferred to the Library Module data bases using a supplied component-transfer program. A macro command language provides links for user-customized use of the Library Module interface. Graphical component display and input are provided by the Library Module macros. On-the-fly component verification is provided in the forms editor and the verification rules can be modified by the user to meet site specifications.

Data Definition

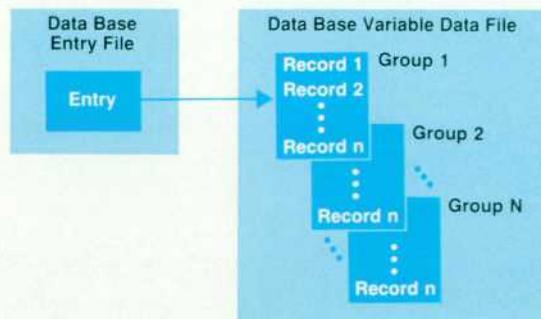
Data stored in the Library Module is broken into two pieces: master data and variable data. The master data contains entry names, and the variable data contains the data used to define the entry. The entry types that HP PCDS currently supports are:



Fig. 1. The composition of a component definition allows sharing of entries to reduce data redundancy.

- `EQVC`, which defines the logical-to-physical mapping information for a component.
- `PHYS`, which defines physical component information such as pin layout, component outline, autoplacement and autoroute parameters, and silk-screen information.
- `COMP`, which is the main component entry type and contains logical pointers to `EQVC` and `PHYS` entries along with component descriptions, component versioning, and information for an HP 3065 Board Test System. The name typically given to a `COMP` entry is the component name, such as 74LS04.
- `GRAF`, for entries used to define board logos.
- `MATL`, which allows a designer to define thick-film hybrid components.

The separation of the mapping, physical information, and component identification allows a design group's librarian to reduce



data redundancy by sharing entries across component main entries. For example, hex inverters generally have the same logical mapping information, but do not always have the same package. Therefore, they can share the same EQVC entry, but have different PHYS entries (Fig. 1). Similarly, a dual four-input NOR gate and a quad analog switch may share the same PHYS entry, but have different EQVC entries.

Each entry is made up of several groups. A group is a table of information similar to that found in a relational data base model, and each group is made up of one or more fields and one or more records (tuples). The entry name is used to produce a hash table index key into the master file. The index into the master file contains a logical pointer to the definition of that entry in the variable file (Fig. 2). This scheme provides very fast access to the contents in the data base, allowing designers to access any of thousands of components at interactive speeds. The speed in data access is apparent both locally and while accessing the data base remotely using LAN remote file access. This is a result of the simple storage mechanism and file structure.

Data Environment

Three data bases can be accessed simultaneously for component definitions:

- USER, a user's or designer's own component entry data
- PROJ, component entry data specific to a group of users
- CORP, component entry data central to an organization or corporation.

Entries that make up a complete component description need not be in the same data base; they can be distributed between data bases (see Fig. 3). The HP PCDS Design Module sequentially searches through the data bases for entries in the order USER-PROJ-CORP (Fig. 4 shows the data base hierarchy). The COMP main entry must be at or below the PHYS and EQVC subentries in the search hierarchy. Therefore, if a COMP entry is found in the PROJ data bases, the subentry definitions must be in the PROJ or CORP data bases. The Library Module allows a user to build a subentry in any one of the three data bases, but only

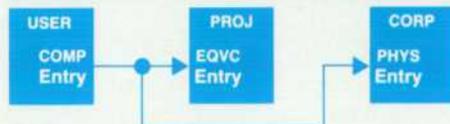


Fig. 3. The sharing of information across data bases and the search order to resolve entry definitions.

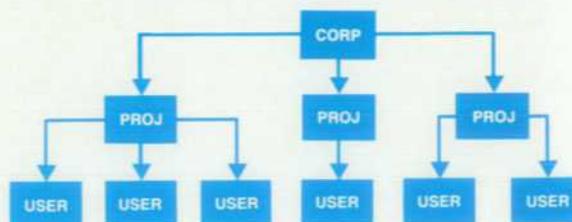


Fig. 4. The data base hierarchy.

allows a COMP entry to be built that accesses a PHYS or EQVC entry in the USER-PROJ-CORP search order.

Data base access across a network is supported through LAN remote file access capabilities. The Library Module uses the PCDS Design System Manager as a sort of data dictionary to inform it of the locations of the data bases. If a data base is located on a remote node, DSM opens the network path to that remote node before invoking the Library Module.

Macro Language

The Library Module has a built-in macro command language that allows users to write their own commands and manipulate different sets of graphically defined data. The macro language is an interpretive language. It is extensive enough to allow a data base administrator to define menus, read from or write to files, and modify entry field data.

Customizability

The data base schema can be modified by the design team as desired. The system allows customers to define new groups and fields within the entries. The new groups defined in the schema can be manipulated in the full screen editor using macros the design team writes. The new groups are also carried along with the component into a Design Module design file. This new group information can then be accessed programmatically through the Design Module design data access routines (DDAR).

The group and field descriptions are also stored in the schema and can be modified. The Library Module uses these descriptions for user help. Field and entry validation rules can also be modified or added to the schema to ensure compliance with design team guidelines and specifications when building a component.

John M. Agosta
Development Engineer
Electronic Design Division

- design, the name of the current design
- project.design.file, the location of the current version of a piece of DSM-maintained data. By using DSM to manage the data and requesting the location from DSM, the application is assured of using the most up-to-date version. DSM will also initialize any required network links.
- home, the current user's working directory. In HP-UX each user has a home directory. DSM extends this convention by adding a home node in the network. Resolving the home parameter returns the path to the standard location for working files on the network and initializes any required network connections.
- ?prompt_string?. The prompt string is displayed and the user must enter a response which then becomes the value of the parameter.
- literal_string. The literal string is interpreted by the HP-UX Bourne shell and then passed to the application. The

experienced HP-UX user can define literal strings that include wild-card characters and variable references that provide access to the rich set of capabilities supported for HP-UX shell programming.

To understand the purpose of each parameter type it is important to remember that DSM maintains all data in a simple hierarchy of projects and designs within projects. Versions of data files are kept within each design. DSM also provides for the maintenance of user logins and environments on an HP-UX system. If an application wants to make use of DSM-maintained data, it must be able to query the project/design hierarchy, where each user wishes to keep work in progress, and the locations of files in DSM's control.

Software Quality Assurance on the HP Printed Circuit Design System Project

by David E. Martin

EFFORTS TO ENSURE SOFTWARE QUALITY must be planned for and kept visible throughout the entire project life cycle. Such efforts were particularly important during the development of Hewlett-Packard's Printed Circuit Design System (HP PCDS). Focusing on quality alone has the inherent danger of never releasing the product because it is not perfect. Developing methods that result in a product that meets high quality standards without sacrificing aggressive goals for a product release date is a monumental challenge for HP as well as the rest of the industry. Two key ingredients in addressing this challenge are quality assurance and engineering productivity.

QA Plan

Often a quality assurance plan is viewed as a necessary nuisance at best and busy work at worst. For a QA plan to contribute to bettering the quality of the software, engineers and management should view it as a contract among themselves detailing their quality efforts.

The QA plan includes the release criteria for the project, which are the basis for the rest of the document. Although release criteria are usually custom tailored for each project, there are some items that are fairly common, such as a clear downward trend of the defect rate, execution of all of the individual test plans, no known critical defects remaining, etc. The rest of the plan describes how it is anticipated that the release criteria will be met.

An optimal QA plan requires extensive effort from everyone. The management team provides the framework for specifying what areas need coverage and ensuring that the coverage is adequate. The details about how to perform the testing are left to the engineering team. It is vital that everyone knows the plan, believes it, and is committed to fulfilling it. To that end, the QA plan is not a static document, but a dynamic one. It is impossible to foresee all events that might cause certain sections of the document to become infeasible, inadequate, or even unnecessary. Thus the plan becomes much more valuable to a project team if it is kept current, reflecting the real intentions of the team. As the testing progresses, results are entered. Performance measurement results are especially important for comparing later revisions of the product. The results give the management team invaluable information for deciding when the formal QA phase should end. A side benefit of keeping the QA plan current is that it automatically becomes a final document for the product archives.

Defect Tracking

Software defects come in many forms. Many are very innocuous, such as the misspelling of a word or unclear error messages. Others are of such a critical nature that the application software aborts unexpectedly and data is lost.

Throughout the life cycle of any software project a large number of defects are encountered and fixed. Attempting to keep track of thousands of defects in a large software project using a manual method is untenable and especially abhorrent, given the fact that HP is in the computer business. Having an automated defect tracking system is viewed as a necessity for helping the management team balance resources and provide the engineers with data on locating and fixing defects. The two major defects tracking systems used in HP are STARS and DTS.¹ STARS (Software Tracking And Reporting System) runs on HP 3000 Computers under the MPE operating system and is used by marketing and field sales organizations. DTS (Defect Tracking System) runs on HP 9000 Computers under the HP-UX operating system.

HP PCDS Development

Since HP's Printed Circuit Design System was developed in the HP-UX environment, it was felt that the engineers would be more productive using DTS. With other systems for tracking defects, it is not uncommon for the R&D lab to be unaware of defects because the procedure for reporting them is too cumbersome. However, with DTS the ease of entering defect reports facilitated the logging of the vast majority of defects encountered. Furthermore, because DTS was installed on every engineer's workstation, engineers found it very convenient to use DTS as a lab notebook. Notes about what caused a defect and ideas on how to fix it were stored directly with the defect data packet. Reminders to enter copyright notices, say, were entered as "defects." Ideas for new features were entered as enhancement requests.

Throughout the HP PCDS project, the management team was kept up to date with accurate statistics that were used to measure the quality of the software. For example, one measurement that was closely monitored was the number of defects reported versus time. A clear downward trend of this curve is evidence that QA efforts are improving the quality of the software.

Although there were many advantages for the lab in using DTS, the needs of the field, sales, and marketing organizations could not be ignored. Before manufacturing release, the outstanding defects from DTS were copied into the STARS data base. Defects from the field and sales organizations are entered in STARS and then copied over to DTS for the lab engineers. There is now software that automatically provides a link between DTS and STARS which keeps the two data bases synchronized. Although supporting two defect tracking systems simultaneously is certainly not easy, the advantages make the effort worthwhile.

Environmental Issues

It is unreasonable to expect anyone to produce quality results in an atmosphere that is not conducive to the task being performed. In the ideal case, the physical environment for a software project should be completely unobtrusive, allowing the engineers to focus their efforts on software development, rather than on manipulating that environment. With each engineer equipped with a workstation, the environment immediately becomes somewhat less than ideal because the numerous workstations create a need for some level of systems administration.

The use of a local area network (LAN) and the addition of a full-time systems administrator can turn this potential liability into an asset. For example, the systems administrator provides expertise in installing and updating systems, troubleshooting hardware and operating system failures, and maintaining the LAN hardware and software. The LAN provides many services that improve the productivity of engineers. The use of electronic mail over the LAN promotes communication with little effort on the part of the engineers. The systems administrator facilitates this activity by establishing and maintaining electronic mailboxes for all the LAN users. Fig. 1 illustrates a method that minimizes the overhead generated by maintenance. Using this method, changes can be handled easily by the systems administrator, since only the files on the central system need to be modified. The addition and deletion of user mailboxes is very straightforward for the rest of the lab.

The LAN also plays an important part in the engineer's development environment. Typically, a team of engineers works on the same software module. This requires that each engineer have access to all of the module's files and that there be an easy way to prevent inadvertent simultaneous updates to the same file by more than one person. The use of SCCS (source code control system), a standard toolset provided by HP-UX, with the LAN provides a foundation on which to build a stable working environment. SCCS requires that files be checked out for the purpose of making modifications, and uses a semaphore to allow only one user to check out a file at a time. While a file is being edited

by one user, other users can still get a read-only version of the file. SCCS also keeps version information so that, if need be, an earlier revision of a file can be restored.

In a distributed environment, there may be a tendency to have each engineer responsible for a defined set of files, with the SCCS files resident on the engineer's local workstation. This approach unfortunately creates extra maintenance for everyone when files shift responsibility or if there is joint ownership of files. Keeping the SCCS files on one workstation designated as a team master system avoids this confusion, yet the LAN still allows the necessary flexibility. Although the SCCS files reside on a remote machine over the LAN, an engineer still has access to them by remote file access (RFA). Makefiles can be set up to perform all commands necessary to build the software module, referencing the location of the SCCS files relative to the local machine. By doing this, each engineer has a local copy of the module which is built from the latest versions of the SCCS files. While an engineer has a file checked out for editing, the other members of the team still have copies of the file before the checkout. Only after the engineer has made appropriate tests is the revised file checked back in and made available to all. In this way, the entire team has access to the latest software, but is protected from the experimentation of individuals.

Additional benefits are derived from the physical environment chosen. Since HP PCDS was developed on the product's target hardware, testing was more convenient for the engineers. The LAN was extensively used when defects were encountered. The engineer could view data files on a remote system or copy them to a local workstation using the LAN. Then debugging activities could take place at the engineer's desk. In other cases, a special version of the code that could be used with the debugger was copied to the test system and the debugging activities took place on the test system. The LAN reduced the overall need for physical media such as magnetic tape for the transferring of files.

System Integration

At some point, all modules must be integrated and the product as a whole must be built on a single workstation. The process must be highly reliable, because at the end of the project cycle, time cannot be wasted trying to figure out how to build the system, or chasing down the reason that some file won't compile. To have continuing confidence in the process, it was decided to rebuild the HP PCDS product every night. A workstation serving as a product master system was dedicated to this task of system integration. This system provided a central repository for files shared by the team masters. Every night each team would build files required by other teams. These would be copied to the master system before a prescribed time. Teams would then copy the files required from other teams from the master system and proceed to build their respective modules. This distributed method of building the product ensured that interteam dependencies were checked continually. This method was incremental in nature because only the files that had changed were recompiled.

On a periodic basis, it was necessary to verify that the whole product could be built entirely from scratch on a

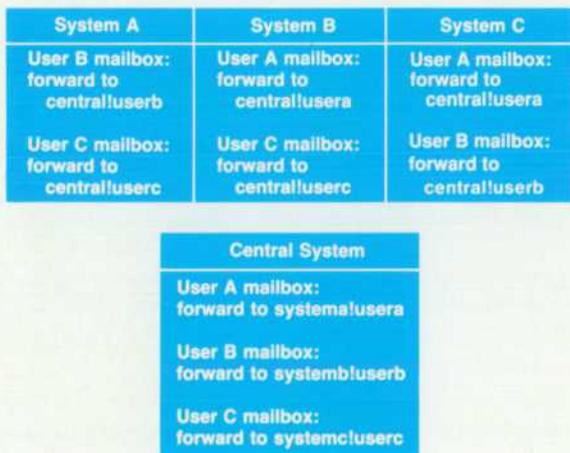


Fig. 1. Method of establishing mailboxes that reduces the overhead required for LAN maintenance by the systems administrator.

single workstation. This had the advantage of uncovering any implicit assumptions that might not be valid. This was done by copying all of the SCCS files from the team masters to the master system. The product was then built, and an installation tape was produced for verification and testing. By establishing this build-and-system-integration process early in the HP PCDS project, valuable time was not consumed at the end of the project, the most critical phase, attempting to get the process under control.

Formal QA Activities

Before entering the QA phase, some QA activities have already been accomplished. Design and code reviews certainly contribute to the quality of the software, yet must be done during the design and implementation phases of the project. During the QA phase of the project, the QA plan becomes the guiding force. To complete the project and ensure that the product has high quality, the plan must be executed. The confidence in the product's quality is in direct proportion to the amount of testing that has been performed on the product. However, testing hours alone do not tell the whole story. There must be a variety of testing, with emphasis placed on testing the right things. Engineers are very good at deriving "white box" tests. That is, since they have developed the software, they know how it functions, and are best able to develop tests that stress boundary conditions, limits, and other structural tests.

A group of engineers from a different section of the lab formed a small test team. This test team had the charter to test the HP PCDS product against its manuals. Nothing is more frustrating to a customer than having the behavior of the software be completely different from what is described in the documentation. The test team members were able to bring their engineering skills to bear on the problem without being biased by intimate knowledge of the internal workings of the software.

Another opportunity for testing came in the form of support engineer and field engineer training classes. Getting prerelease software stable enough to support training classes during the early phase of the project helped establish installation and system integration processes and illuminate deficiencies and weaknesses in the product. The training classroom had the serendipitous advantage of providing ideal resources for hardware configuration testing during idle times.

Another whole category of testing is more important to customers. This is fitness for use, that is, can the customer use the product in its intended application? Software engineers rarely have strong backgrounds in the use environments for the applications they are developing. The addition of an application resource (AR) team composed of typical end users fills this gap. In the case of HP PCDS, several experienced printed circuit board designers were assembled as the AR team. Their job was to use HP PCDS to design actual printed circuit boards. It was extremely advantageous that the AR team was in the same physical location as the software lab. During the process of designing a board, the AR team would give instant feedback to the lab engineers concerning problems, suggestions, etc.

The AR team was particularly qualified to evaluate the results produced by HP PCDS. With an experienced eye, they could quickly spot a problem that a software engineer would never notice. It is typical for prospective buyers of a printed circuit board design system to request that a printed circuit board representative of their requirements be designed on the system before purchase. This benchmarking activity was very professionally accomplished by the AR team.

The AR team also coordinated and gave support to the alpha test sites within Hewlett-Packard. Alpha sites were chosen such that HP PCDS would have a wide exposure to different printed circuit board technologies. There was a strong partner relationship between the lab and the alpha sites. The alpha sites benefited from their participation by receiving the software much earlier than the rest of HP and by being able to provide input regarding desired product refinements. Naturally, the design team benefited from the extra end-user testing. The AR team and the alpha sites had the experience to evaluate the quality of HP PCDS properly from the customer's perspective. They were the "voice of reality."

Acknowledgments

I would like to express my sincere appreciation to the following people. Elaine Regelson provided me with guidance and enthusiastic support, which encouraged me to do a better job than I thought I was capable of. Rick Roeling joined the project at a time when I was desperate for some help. I am still amazed at how quickly he came up to speed and how he was able to take over the responsibility of administering DTS. Many of the benefits derived from using DTS were a direct result of Rick's efforts. Pati Porter served as our systems administrator and developed the methods for keeping the electronic mail system and the LAN running smoothly. Also, I am indebted to Jon Gustafson, Debbie Lienhart, Mark Mayotte, Elaine Regelson, and Denese Van Dyne for their efforts in helping me prepare this article.

Reference

1. S.R. Blair, "A Defect Tracking System for the UNIX Environment," *Hewlett-Packard Journal*, Vol. 37, no. 3, March 1986, pp. 15-18.

Silicon-on-Insulator MOS Devices for Integrated Circuit Applications

Several techniques for fabricating regions of crystalline silicon on insulating substrates are available. These methods are described briefly and device design considerations introduced by the use of SOI are discussed.

by Jean-Pierre Colinge

SILICON-ON-INSULATOR (SOI) technologies are a growing class of techniques for fabricating integrated circuits. Unlike conventional devices fabricated in bulk silicon, SOI devices are made in a thin silicon layer deposited on an insulator, which can be either amorphous or crystalline. When a crystalline insulator is used the silicon layer can be grown epitaxially, provided the lattice parameters of the crystalline insulator are not too different from those of silicon.

Silicon-on-sapphire (SOS) is the best-known technology making use of silicon epitaxy on an insulator. SOS is a technically successful technology (many satellites contain SOS circuits), but the high cost of the single-crystal sapphire wafers renders this technology impractical for most commercial applications. Other crystalline materials, such as calcium fluoride, are being investigated as a replacement for sapphire.

Another alternative is the use of an amorphous insulator. Silicon dioxide (SiO_2) is the most widely used insulator since it is fully compatible with existing integrated circuit fabrication techniques. The most mature methods for forming a silicon-on-oxide structure involve either implanting oxygen into a silicon wafer or recrystallizing polysilicon deposited on an SiO_2 layer. Both approaches are being investigated at Hewlett-Packard.^{1,2}

The oxygen implantation method is called SIMOX (separation by implantation of oxygen). During this process, a very high dose of oxygen (2×10^{18} atoms/cm²) is implanted into a silicon wafer to form a buried layer of stoichiometric SiO_2 . Implantation energy and temperature are carefully chosen to leave a thin layer of single-crystal silicon on top of the silicon dioxide layer.

The other method, recrystallization of polycrystalline silicon deposited on an oxide layer, can be achieved by various techniques. All of these are miniature versions of the zone-melting recrystallization (ZMR) technique used to produce bulk silicon single crystals. In micro-ZMR techniques, a thin polysilicon film is melted using either an electron beam, a focused laser beam, or a focused incoherent light source. The thermal profile of the silicon film is controlled in such a way that upon cooling, large single-crystal areas are obtained in which devices and circuits can be formed. Micro-ZMR also opens the door to three-dimensional (3D) integration where two or more active layers are formed on a silicon wafer. Using a laser, for instance,

it is possible to recrystallize a silicon layer deposited on an existing integrated circuit to realize another active layer. The Japanese Ministry of Trade and Industry (MITI) is fostering research in 3D integration, and 3D prototype circuits such as camera photosensors and memory chips have been produced.³

In the U.S.A., most SOI research is oriented towards producing radiation-hardened circuits and providing a cheaper alternative to SOS technology. Good radiation hardness is not the only advantage presented by SOI structures. The inherent dielectric insulation of SOI devices and the relative thinness of the silicon film in which devices are made make SOI an ideal choice for high-speed, very-large-scale integrated circuit fabrication.

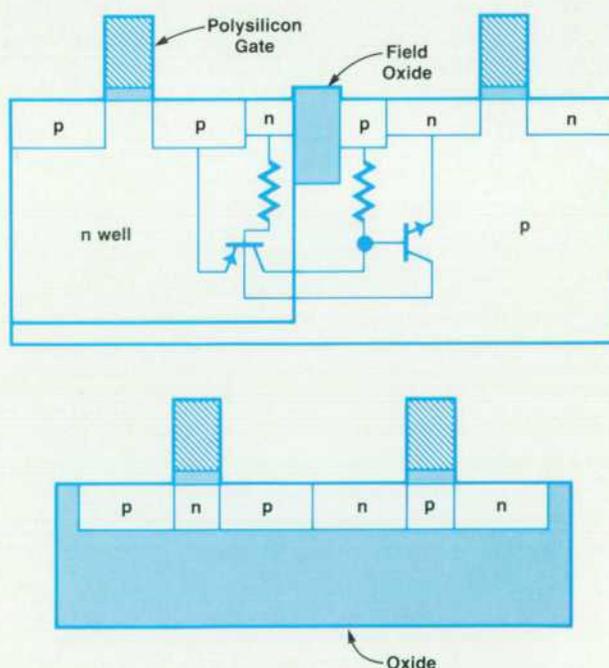


Fig. 1. Cross-sections of CMOS inverters made using bulk-silicon (top) and SOI (bottom) technologies. Resistors and bipolar transistors in the bulk device symbolize the nnp latch-up path.

General Properties of SOI Devices

Although different types of transistors (bipolar, JFET, DMOS, etc.) can be fabricated in SOI films, MOS and more particularly CMOS devices seem to take best advantage of SOI technology. These advantages are absence of latch-up, process simplicity, radiation resistance, and reduction of parasitic capacitance.

Absence of Latch-Up. The latch-up phenomenon is the unwanted triggering of the parasitic npnp (thyristor) structure present in a standard CMOS circuit. The thyristor can be schematically represented by two bipolar transistors (Fig. 1, top). Triggering can be induced by light, alpha particles, and even thermal noise. In an SOI CMOS structure (Fig. 1, bottom), there is no such npnp structure since the transistors are dielectrically insulated from one another by silicon dioxide.

Process Simplicity. Fabrication of CMOS structures is much simpler in SOI than in bulk silicon. Acceptor or donor impurity doping of the silicon islands can be achieved during the implantation step used to control threshold voltage. Therefore, there is no need for creating diffused wells, an operation that consumes both process time and silicon real estate in a bulk-silicon CMOS process. Bulk-silicon technologies use techniques such as LOCOS (local oxidation of silicon) or etched trenches backfilled with an insulator to isolate devices from one another. In SOI, a patterned silicon etch can be used readily to form islands. Relatively thin LOCOS oxides can also be used to isolate the silicon islands. In this case, the resulting oxide thickness is approximately twice that of the silicon film. SOI LOCOS field oxides, therefore, have thicknesses in the 200-to-400-nm range, compared to a typical thickness of 800 nm for bulk-silicon devices. The use of such modest oxide thicknesses limits the undesirable LOCOS "bird's beak," the lateral growth of the oxide into the adjacent device areas. This allows increased packing density.

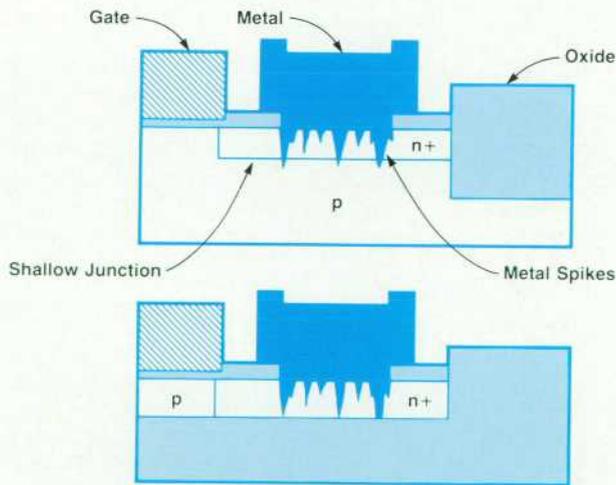


Fig. 2. Illustration of metal spikes punching through a shallow junction in a bulk transistor (top). In the case of an SOI device (bottom) where the source and drain regions extend throughout the entire film thickness, no junction punchthrough is possible.

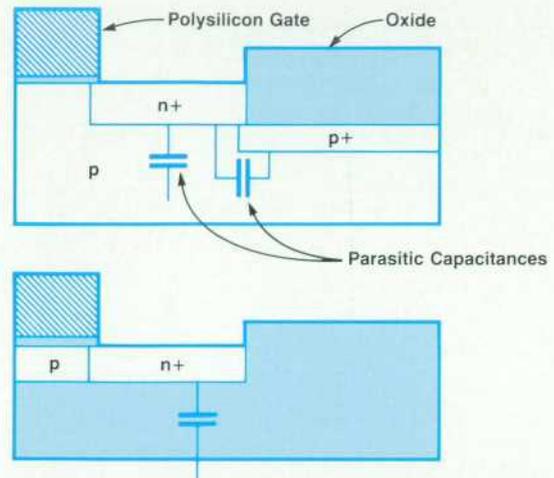


Fig. 3. Drain-to-bulk and drain-to-field-implant capacitances in a bulk transistor (top). In an SOI device (bottom), the drain-to-substrate capacitance is reduced because of the insulating oxide underneath the device.

Packing density can be increased further because SOI provides the ability to merge n and p diffusions. The bottom half of Fig. 1 represents the cross-section of an SOI CMOS inverter. Unlike an inverter fabricated in bulk silicon, the drains of the p-channel and the n-channel transistors can touch one another. This feature allows SOI CMOS circuits to reach the density levels of NMOS circuits.

Because the source and drain junctions cannot extend deeper than the silicon film thickness, realization of shallow junctions is automatically achieved when thin (<200 nm) SOI films are used. Similarly, it is impossible for the contact metal system to spike down through a junction when the source and drain diffusions extend through the entire silicon film thickness (Fig. 2). This simplifies metallization processes and reduces leakage current problems.

Radiation Resistance. Because they require a smaller volume of silicon, SOI devices are more resistant to radiation damage than bulk devices. SOS has always been a technology of choice for space and military applications. SOI devices, which can be made in thinner films than SOS devices, are therefore good candidates for applications where radiation hardness is a key issue.

Reduced Parasitic Capacitance. Source and drain junctions of MOS transistors have parasitic capacitance, the value of which is roughly inversely proportional to the thickness of the space charge region below the junctions. Modern VLSI MOS circuits are made in relatively heavily doped substrates. Higher substrate doping means thinner space charge regions, and hence increased parasitic capacitances. SOI devices are, by definition, sitting on top of an insulator, usually SiO_2 . The parasitic capacitance of a source or a drain now depends on the thickness of the insulating oxide. This thickness is typically $0.35 \mu\text{m}$ in the case of SIMOX material, and $1.0 \mu\text{m}$ in the case of beam-recrystallized silicon. Taking into consideration that the dielectric constant of silicon dioxide is three times smaller than that of silicon, substantial reduction of parasitic capacitances can be obtained in SOI circuits. In addition,

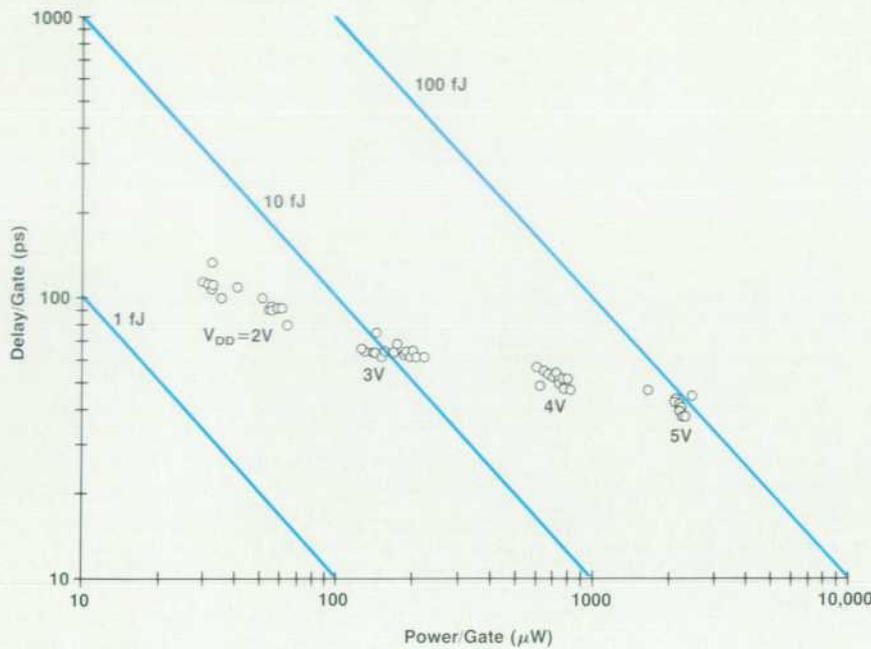


Fig. 4. Scatter plot of delay per gate versus power dissipation per gate in CMOS SOI ring oscillators. The straight lines represent different delay-power products. The effective channel lengths are 0.25 and 0.47 μm for the n-channel and p-channel devices, respectively.

parasitic capacitances between junctions and field implants are essentially nonexistent in SOI devices (Fig. 3).

Reduced parasitic capacitances, increased packing density, and easier processing, even at submicrometer dimensions, permit the realization of high-speed circuits in SOI films. Fig. 4 shows the performance of CMOS ring oscillators realized in SIMOX material. The effective gate length is 0.25 μm for the n-channel transistors and 0.47 μm for the p-channel devices. Gate delays down to 40 picoseconds (at $V_{DD} = 5\text{V}$) have been obtained, which is the lowest value obtained in silicon-based CMOS circuits as of February 1987. More recently (October 1987), CMOS frequency dividers have been made at HP, using a flip-flop type of logic. These circuits, in which the minimum gate length is 0.65 μm , operate at clock rates of 2 GHz.

Physics of an SOI MOS Transistor

The major difference between an SOI MOSFET and a bulk-silicon MOSFET is that the SOI device is made in a silicon film of finite thickness. The thickness is usually comparable in magnitude to the depth of the space charge region created by the gate. A few years ago, SOI devices were fabricated in silicon layers having a thickness of 0.5 μm . Recent years have seen film thicknesses reduced to 0.2 μm or less for high-speed applications, and 0.3 μm for radiation-hard devices. A diagram of energy band curvature helps demonstrate how such thick-film and thin-film SOI transistors differ from bulk devices. For simplicity, only the n-channel device is described.

In a bulk transistor that is being turned on, the energy bands are horizontal deep in the bulk of the silicon material (the doping profile is assumed to be constant). The bands are bent in the space-charge region (depletion region) until a surface potential of $2\phi_F$ is reached and an inversion layer is created (Fig. 5a). The width of the depletion region has a maximum value W_{max} equal to $\sqrt{4\epsilon_{Si}\phi_F/qN_A}$, where N_A is the channel dopant concentration, q is the electron

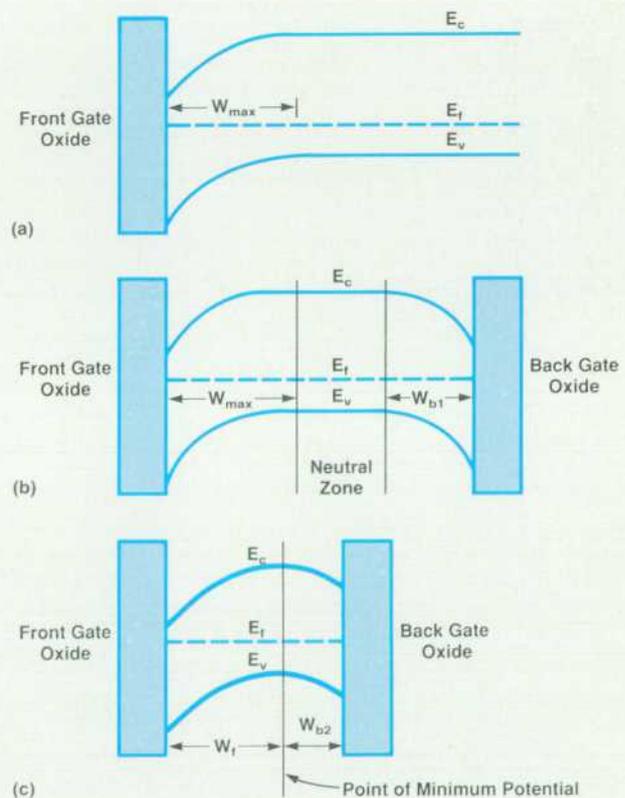


Fig. 5. Band diagram for n-channel transistors under inversion. (a) Bulk device. (b) Thick SOI device. (c) Thin SOI device. W_{max} is the maximum depletion depth. W_l , W_{b1} , and W_{b2} are the depletion depths related to the front gate, the back gate in a partially depleted film, and the back gate in a fully depleted film, respectively. The following relationships apply: $W_{max} > W_l$ and $W_{b1} > W_{b2}$.

charge, ϵ_{Si} is the silicon dielectric constant, and ϕ_F is the Fermi potential.

In a thick-film SOI MOSFET, a similar band curvature is observed near the gate-oxide/silicon interface. Since the device is sitting on a silicon dioxide film on top of a silicon wafer, a second oxide/silicon interface is present at the bottom of the silicon film (Fig. 5b). The underlying oxide film can be viewed as a second gate oxide (back gate oxide) with the silicon wafer acting as a backside gate. Band curvature at the back interface depends on the bias of the substrate and on the interface states present at the back Si/SiO₂ interface. In general, the back interface can be under accumulation, depletion, or inversion (in which case leakage current flows from source to drain.)

In a thin-film SOI MOS device, the silicon film thickness is less than twice the maximum depletion width. Thus, the silicon film can be fully depleted. In this case, no neutral silicon is left between the front and backside depletion zones⁴ (Fig. 5c). MOS transistors made in such films have remarkable properties. For instance, a thin-film SOI MOSFET can be switched from a fully depleted state to a partially depleted state by applying a bias to the back gate, which accumulates holes at the bottom silicon interface. **Threshold Voltage.** Under similar channel doping conditions, bulk-silicon and thick-film transistors have the same threshold voltages. Since the front-gate and back-gate depletion zones do not merge, the front and backside threshold voltages are decoupled. Hence, both front and back threshold voltages are given by the classical expression:⁵

$$V_T = \phi_{MS} + 2\phi_F - q(N_{ss}/C_{ox}) - (Q_b/C_{ox}) \quad (1)$$

where ϕ_{MS} is the difference in work function between the gate material (i.e., front gate or back gate) and the silicon film, N_{ss} is the surface state density at the Si/SiO₂ interface, and C_{ox} is the capacitance of either the gate or insulating oxide layer. Q_b is the depletion charge, equal to $qW_{max}N_A$, where N_A is the dopant concentration.

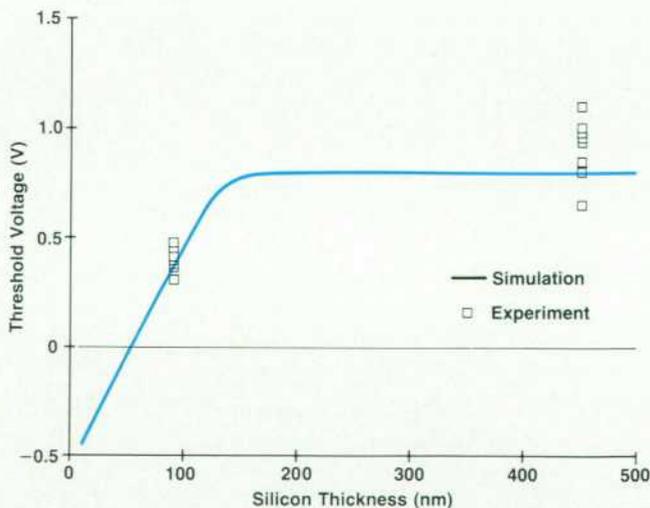


Fig. 6. Threshold voltage as a function of film thickness in *n*-channel SOI MOSFETs. The boron doping level is constant ($8 \times 10^{16}/\text{cm}^3$), and the gate oxide thickness is 25 nm.

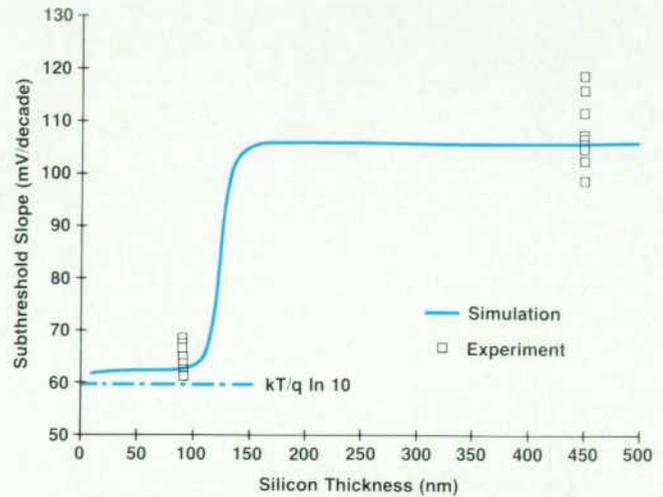


Fig. 7. Subthreshold slope as a function of film thickness in *n*-channel SOI MOSFETs. Same parameters as for Fig. 6.

Things are quite different in a thin-film device. Where the silicon film is fully depleted, there is no neutral silicon. The front-gate depletion zone extends from the gate oxide interface to the point of minimum potential. At that point, the electric field changes orientation, since the bottom part of the depletion zone is related to the backside gate. If the backside surface potential is increased, the back-gate depletion zone extends to a larger portion of the film at the expense of the front-gate space-charge region, thereby modifying the front-gate threshold voltage. Similarly, a change in front surface potential affects backside threshold voltage.

Threshold voltage modulation by a back-gate potential is not acceptable in a commercial circuit where substrate potential is usually kept grounded, but it provides a unique way of investigating the device physics and deriving the properties of very thin devices where the film thickness t_{Si} is smaller than W_{max} . Equation 1 can be used to derive the threshold voltage of thin devices, provided that Q_b is replaced by an appropriate depletion charge value located between $qN_A t_{Si}$ and $qN_A t_{Si}/2$ and depending on the surface state density at the back interface and the back-gate voltage.

Fig. 6 shows the threshold voltage of an SOI *n*-channel transistor as a function of silicon film thickness. If we start with a thick film (500 nm) and the film thickness t_{Si} is gradually decreased, Q_b and V_T remain constant until the film becomes fully depleted. At that point the back depletion zone encroaches on the front space-charge territory. This gives rise to a reduction of the front depletion depth from W_{max} to W_f (Figs. 5b and 5c). As a consequence, Q_b and V_T are linearly decreased as film thickness is decreased (Fig. 6).

Subthreshold Slope. The subthreshold slope (sometimes called inverse subthreshold slope or subthreshold swing) of an MOS device is the slope of the drain current I_D versus gate voltage V_G below the threshold. The steeper the slope, the smaller the gate voltage swing required to turn the device on and off. In modern CMOS circuits, high-speed operation requires the lowest admissible threshold voltage. Since the off current of the device must be very low to minimize standby power consumption, the lower bound

of threshold voltage is set by the subthreshold slope. Typical subthreshold slopes are 90 mV per decade. This means that a 90-mV increase in gate voltage is needed to increase drain current by a factor of 10 in the subthreshold region.

The subthreshold slope S is classically given by:⁶

$$S = \{ (kT/q) \ln(10) \} [1 + (C_d + C_{it})/C_{ox}] \quad (2)$$

where k is Boltzmann's constant, T is temperature, and C_d and C_{it} are the depletion and interface trap capacitances, respectively. The interface trap capacitance is equal to qD_{it} , where D_{it} is the interface trap density. Equation 2 basically expresses that the gate charge is distributed in a capacitive divider. The smaller the depletion capacitance, the steeper the subthreshold slope. In a thick-film SOI transistor, the depletion capacitance, and hence the subthreshold slope, is a function of the doping level N_A just as it is for bulk devices. In a thin-film transistor, on the other hand, the silicon film thickness can be smaller than W_{max} . Once the film is fully depleted, any further increase in gate voltage does not bring about an increase of depletion charge ($C_d = \partial Q_b / \partial V_G = 0$), but instead increases the electron concentration in the forming inversion layer. Therefore, the value of S is predicted to be smaller in thin-film transistors and should eventually reach the theoretical limit of $(kT/q) \ln(10) = 60$ mV/decade in the case of a very thin film with no traps at the Si/SiO₂ interfaces at room temperature.

Fig. 7 presents experimental values of S for n-channel MOS devices made in two different SOI film thicknesses compared with the theoretical dependence of S on film thickness obtained by numerical simulation. The doping profile in the channel is constant and equal to 8×10^{16} boron atoms per cm³. A sharp transition occurs when the

film thickness is equal to W_{max} . This transition thickness depends on dopant concentration and marks the limit between thin-film and thick-film behavior.

Effects Related to the Floating Substrate. In bulk MOS transistors, the channel region below the gate is part of an "infinitely thick" substrate that is grounded in most applications. In an SOI transistor, the substrate in which the device is made has a finite thickness (the thickness of the silicon film). Furthermore, this substrate is electrically floating and can have a potential ranging from 0V to a few hundred millivolts for an n-channel device.

When an MOS device operates in the saturation region, impact ionization can occur near the drain. This creates electron-hole pairs in the channel region. The electrons are attracted by the gate potential and merge into the channel current flowing from source to drain. Holes flow towards the region of minimum potential, which is the floating substrate (neutral zone) in the case of a partially depleted SOI film (Fig. 5b). When no holes are injected in the floating substrate, its potential is 0V (this occurs at low drain currents and low drain voltages). At higher drain potentials, hole injection into the floating substrate occurs. As holes accumulate in the floating substrate its potential becomes positive and current flows through the forward-biased source-substrate diode. The substrate potential is set by an equilibrium condition between the hole current generated by ionization and the current in the forward-biased source-substrate diode.

Any increase of the floating potential brings about a decrease of the device threshold voltage, which shows up in the form of a kink in the output characteristics of the transistor (Fig. 8). For a thin, fully depleted film, the point of minimum potential is located near the middle of the film

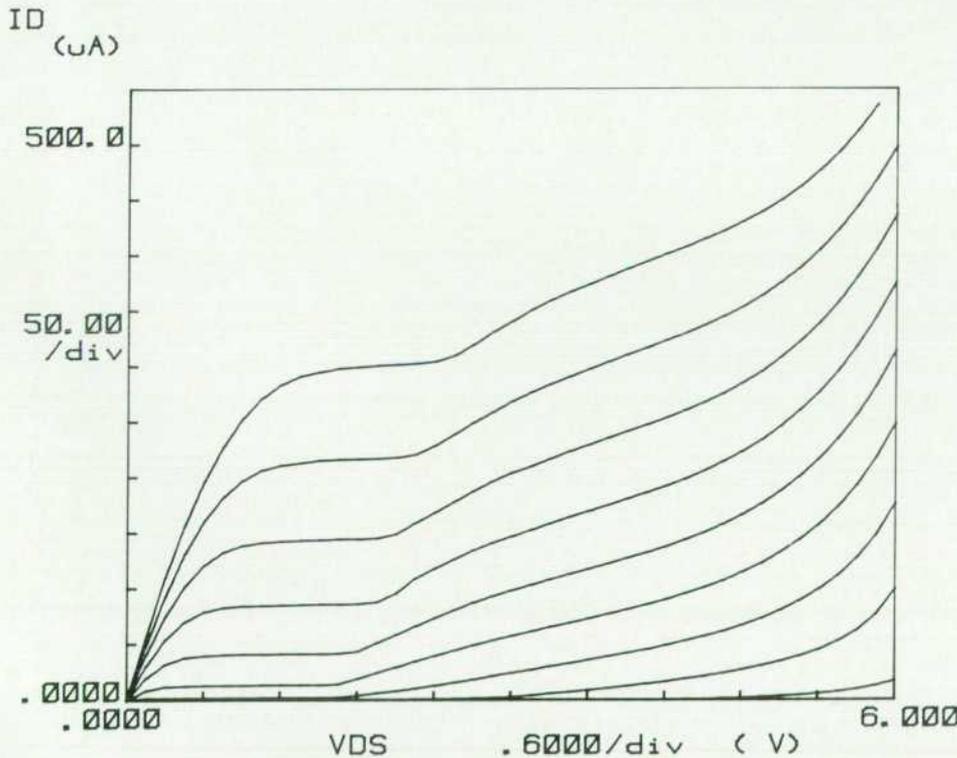


Fig. 8. Output characteristics of an n-channel MOSFET made in a 100-nm-thick SOI film. The back-side gate bias is set to -20V to accumulate the bottom of the silicon film and create a floating substrate in the device channel region. Gate voltage ranges from 0 to 2.5V in 0.25V steps.

(Fig. 5c). This potential is positive, even in the absence of impact ionization. Since there is no neutral silicon in the film, holes cannot accumulate in a floating substrate. Holes generated by impact ionization migrate to the point where the potential is lowest, but this potential is high enough to prevent them from accumulating. Therefore, the holes flow to the source through the forward-biased source-substrate diode, the potential distribution in the film is hardly modified by hole injection, and no kink effect is observed.

The dependence of the kink effect on the level of depletion in the silicon film can be easily demonstrated. Fig. 9 presents the output characteristics of an SOI transistor made in a 100-nm-thick silicon film. The channel doping concentration N_A is $1.3 \times 10^{17}/\text{cm}^3$, corresponding to a W_{max} of 91 nm. When a potential of 0V is applied to the back gate (silicon mechanical substrate), the channel region is fully depleted ($t_{\text{Si}} < 2W_{\text{max}}$), and no kink can be seen in the output characteristics. When a potential of -20V is applied to the back gate (the back gate oxide is 800 nm thick), holes accumulate at the bottom of the device and a floating substrate is created. Then, a kink in the output characteristics can readily be observed (Fig. 8). It is also worth noting that the threshold voltage increases about 0.5V when accumulation is created at the back interface. This is because the front-gate-related depletion depth is increased from approximately $t_{\text{Si}}/2$ to W_{max} when negative back bias is applied. It can also be observed that transconductance ($\partial I_D / \partial V_G$) is larger in the fully depleted device. Indeed, the vertical electric field E_s near the silicon/gate-oxide interface is smaller. In the case of a fully depleted device, E_s is proportional to the front-gate depletion depth. A smaller surface electric field means higher channel electron mobility and hence, larger transconductance.

Threshold Voltage Dependence on Gate Length. Threshold voltage variation with gate length is observed in all types of MOS devices and is often called the short-channel effect. The threshold voltage of a long-channel n-channel MOSFET is traditionally given by equation 1. When the gate length is decreased to submicrometer values, electric field lines near the source and drain terminate on the source and drain instead of the gate. This reduces the effective depletion charge Q_b , controlled by the gate and, therefore, decreases V_T . The situation is similar in devices made in thick, fully depleted SOI films. In thin SOI films, the close proximity of the backside depletion region reduces the influence of the source and drain field lines on the front depletion zone. As a result, long-channel behavior is preserved at shorter gate lengths than in bulk devices.⁷ Fig. 10 presents the calculated threshold voltage in a bulk device and a thin-film SOI device as a function of gate length. The SOI device has a lower threshold voltage since the silicon film is fully depleted. The short-channel effect is observed in the bulk transistor for gate lengths shorter than $0.5 \mu\text{m}$, while threshold voltage remains stable down to gate lengths as short as $0.25 \mu\text{m}$ in the SOI device. When a thicker SOI film is used, short-channel effect improvement is less dramatic, but still clearly observable.

Fig. 11 presents the short-channel behavior of bulk and SOI devices made in the same lot. Although the SOI film is 200 nm thick and not fully depleted, it can be observed that the SOI devices are less affected by short-channel effects than bulk devices. Almost as important as the value of threshold voltage at a given gate length is its variation across a device lot. Indeed, gate length can only be controlled within a given accuracy across a wafer. In short-channel bulk devices, any variations of gate length caused

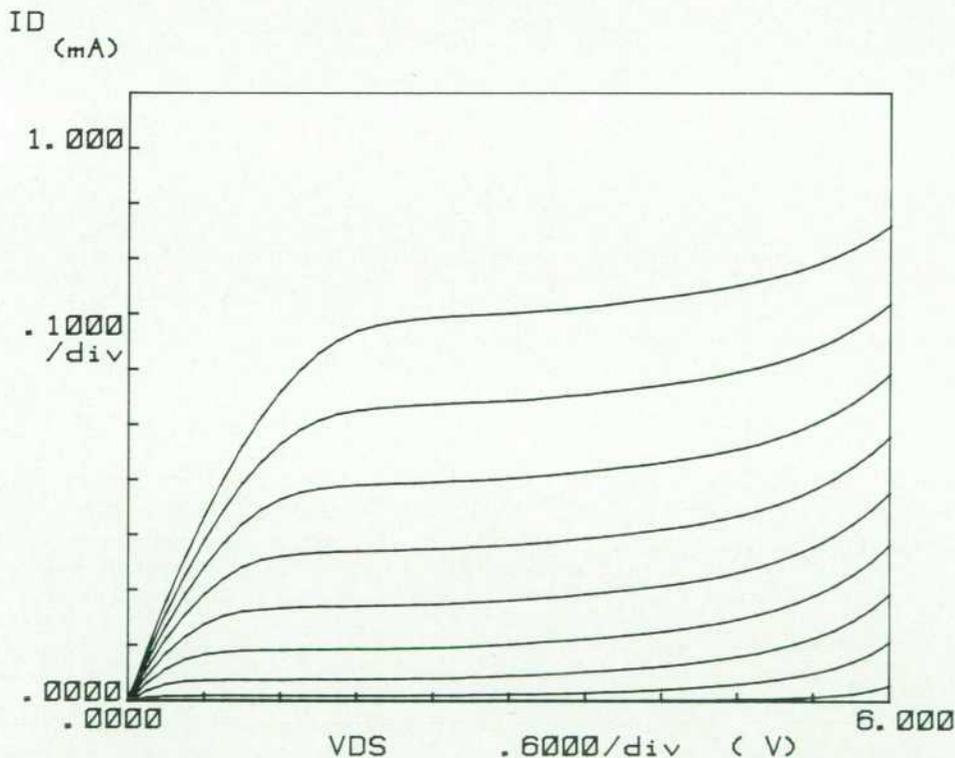


Fig. 9. Same device as in Fig. 8, but the backside gate voltage is now set to 0V. The kink in the output characteristics has disappeared since the device is now fully depleted. Gate voltage ranges from 0 to 2.5V in 0.25V steps.

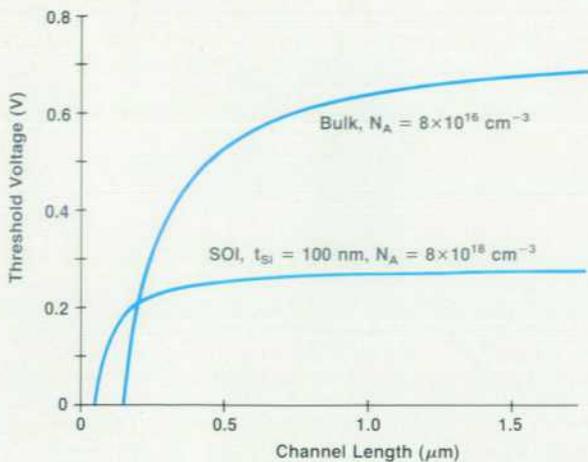


Fig. 10. Calculated threshold voltage in *n*-channel transistors as a function of channel length. The top curve represents a bulk device, while the bottom curve represents an SOI device. Both transistors have the same constant channel boron doping profile ($8 \times 10^{16}/\text{cm}^3$). The SOI transistor is made in a 100-nm-thick silicon film.

by process control variations give rise to a large range of threshold voltages. In SOI devices where threshold voltage is less sensitive to gate length, similar differences in gate lengths only give rise to a moderate range of V_T values. This can be observed in Fig. 11, where V_T varies over 400 mV in the bulk transistors, and only over 250 mV in the SOI devices (at a gate length of 0.25 μm).

Conclusions and Perspectives

Silicon-on-insulator technologies provide the ability to realize high-speed submicrometer CMOS circuits much more easily than conventional bulk-silicon technologies do. Unlike SOS technology, which has been successful only in some niche markets, it is expected that SOI technol-

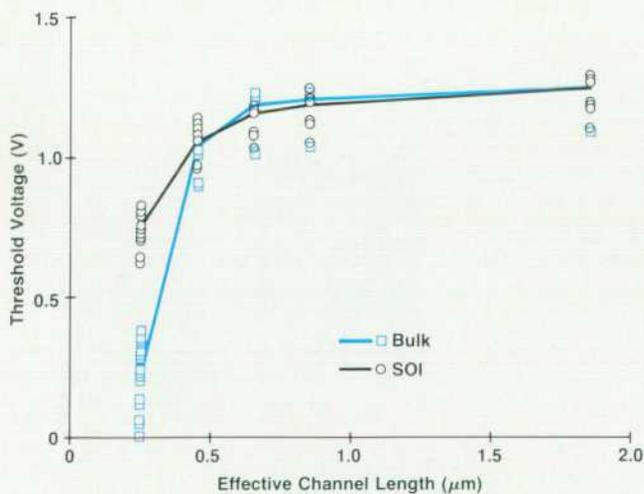


Fig. 11. Measured threshold voltages in bulk and SOI *n*-channel MOSFETs as a function of effective channel length. The devices were made in the same lot and received the same threshold implants. The SOI devices were made in a 200-nm-thick SOI film.

ogy will gain more and more interest in all domains of CMOS circuit applications. The reasons are quite simple. SOS circuits must be processed in special process lines, mainly because the sapphire material contaminates diffusion furnaces. SOI wafers, on the other hand, contain only silicon and SiO₂ and can be processed in conventional process lines together with bulk-silicon products. SOS material requires sophisticated amorphization and regrowth processes to render its quality suitable for submicrometer applications but SOI films of suitable quality can be obtained readily by merely performing implantation and thermal annealing steps.

As bulk-silicon CMOS technologies progress towards the submicrometer level, process complexity increases dramatically, and costly techniques must be developed to create dense field isolation and shallow junctions. In SOI circuits, these problems can be solved in a more straightforward manner.

Finally, the high cost of oxygen-implanted wafers is starting to drop because of the advent of commercially available high-current oxygen implanters that have adequate throughputs for high-volume production.

Acknowledgments

Many individuals are contributing to the success of HP's SOI project. The author's appreciation goes to Ted Kamins for his work on improving the SOI material quality, to Jan Turner who performed the electrical measurements, and Paul Rissman and Marsha Long for electron-beam lithography. The author also wishes to thank the engineers and technicians of the silicon process laboratory and the electron beam department in HP Laboratories for their help and support in wafer processing, as well as Shang-Yi Chiang and John Moll for their enthusiastic support.

References

1. J.P. Colinge, et al, "High-Speed, Low-Power, Implanted-Buried-Oxide CMOS Circuits," *IEEE Electron Device Letters*, Vol. EDL-7, no. 5, 1986, p. 279.
2. J.P. Colinge, H.K. Hu, and S. Peng, "Fabrication of Thin SOI Films Using Laser Recrystallization," *Electronics Letters*, Vol. 21, no. 23, 1985, p. 1102.
3. Y. Akasaka and T. Nishimura, "Concept and Basic Technologies for 3-D IC Structure," *Technical Digest, 1986 IEEE International Electron Devices Meeting*, p. 488.
4. J.P. Colinge, "Transconductance of SOI MOSFETs," *IEEE Electron Device Letters*, Vol. EDL-6, no. 11, 1985, p. 573.
5. A.S. Grove, *Physics and Technology of Semiconductor Devices*, John Wiley & Sons, 1967.
6. S.M. Sze, *Physics of Semiconductor Devices*, second edition, John Wiley & Sons, 1981.
7. T. Sekigawa and Y. Hayashi, "Calculated Threshold-Voltage Characteristics of an XMOS Having an Additional Bottom Gate," *Solid-State Electronics*, Vol. 27, no. 8/9, 1984, p. 827.

Authors

February 1988

15 == 12-Bit Waveform Recorder

Albert Gee



With HP from 1979 to 1986, Al Gee joined the digital signal analyzer R&D section at the Santa Clara Division after receiving his BE degree in electrical engineering and computer science from the University of California at Berkeley. He also holds a 1982 MSEE from Stanford University. His early R&D activities included the design of digital signal processing algorithms and analog-to-digital converter systems. More recently, he designed the HP 5183A analog-to-digital system. Born in Sacramento, California, Al is married and enjoys Tai Chi, football, basketball, softball, and skiing.

Nancy W. Nelson



With HP since 1980, Nancy Nelson is currently a Santa Clara Division project manager. Her earlier projects include firmware for the HP 5182A Waveform Recorder/Generator and analog circuits for frequency counter instrumentation. Nancy grew up in New Mexico and received a BS degree in chemistry and psychology from the University of New Mexico in 1974. She then began work on a BSEE at the University of New Mexico and completed her studies at California State University, San Jose. Past president of the Santa Clara County Chapter of the Society of Women Engineers, she is currently working on plans for the 1989 national convention in Santa Clara County. Nancy is married and has one son. Outside of work she enjoys camping, reading, gardening and Cub Scouts activities.

Ronald W. Young



An HP R&D engineer from 1979 to 1987, Ron Young worked on the HP 5180A and 5183A Waveform Recorders. His work on the HP 5183A included contributions to the topology of the measurement system and the design of the input amplifier and the reference oscillator. Ron received a BSEE degree in 1978 from Stanford University and an MSEE degree in 1981 from the University of California at Berkeley. He was born in Bakersfield, California, and his interests include volleyball, windsurfing, tennis, and mountain activities. He also enjoys piano music and reading.

26 == Waveform Reconstruction

Allen S. Foster



Allen Foster was born in San Francisco and has worked for HP for 23 years. His BSEE degree is from the University of California at Berkeley (1963) and his MSEE is from Stanford University (1971). His outside interests include Porsches, computer hardware design, UNIX system software, and remodeling of his mountain home in Saratoga, California.

Richard W. Page

Author's biography appears elsewhere in this section.

6 == Waveform Recorders/Oscilloscopes

James L. Sorden



With HP from 1964 to 1987, Jim Sorden was in R&D management for most of his HP career. He was the program manager and authored earlier HP Journal articles on the HP 5345A Frequency Counter and the HP 5180A Waveform Recorder, and was section manager for the HP 5180T/U and 5183T/U Digitizing Oscilloscopes. After serving in the U.S. Army, he attended the Universities of Georgia and Wisconsin, and Stanford University. His BSEE degree is from Wisconsin and he did graduate study in electronic engineering at Wisconsin and Stanford. Jim is married and has two teenage children. His hobbies include swimming, skiing, house remodeling, and real estate and stock investing.

23 == Adaptive Sample Rate

Richard W. Page



A native of Los Angeles, Richard Page was with HP for 14 years and is now a senior scientist with another firm, working in the areas of adaptive signal processing and spectrum estimation. He was responsible for the adaptive sample rate algorithm for the HP 5183 Waveform Recorder. He received his BSEE degree from California State Polytechnic University at Pomona in 1973 and his MSEE from Stanford University in 1978. Rich and his wife live in Los Altos, California. His hobbies include sailing and music, and he is working on his PhD degree in signal processing.

32 == 250-MHz Waveform Recorder

Patrick D. Deane



Pat Deane received his BSEE degree from Washington University of St. Louis in 1976 and his MSEE degree from Stanford University in 1978. With HP's Santa Clara Division since 1978, he has designed the input amplifier of the HP 5180A Waveform Recorder, installed CAD systems for hybrid and printed circuit layout, contributed to the design of the preamp and ADC hybrids for the HP 5185A Waveform Recorder, and served as project leader for HP 5185A printed circuit hardware. He's a member of the IEEE and a registered professional engineer, and specializes in analog, high-speed digital, and hybrid circuit design. Raised in Omaha, Nebraska, Pat is married, enjoys basketball and golf, and is learning to play piano.

Rayman W. Pon



Ray Pon has been a development engineer with HP's Santa Clara Division since 1983 and is now working for his MSEE degree at Stanford University through the HP resident fellowship program. He contributed front-panel, I/O, external trigger, memory,

and control circuitry to the HP 5185A Waveform Recorder. He is a member of the IEEE and holds a 1983 BS degree in electrical engineering and computer science from the University of California at Berkeley. Ray was born in Hong Kong. His leisure activities include piano, tennis, volleyball, and skiing.

Mark A. Unkrich



Mark Unkrich is a high-frequency analog circuit designer. With HP's Santa Clara Division since 1981, he has designed the input amplifier ICs and preamp hybrid for the HP 5185A Waveform Recorder, developed software for instrument control and testing,

and authored an IEEE paper on crystal oscillator startup. A native of Cincinnati, Ohio, he received his BS degree in electrical engineering from the Massachusetts Institute of Technology in 1979 and his MS in electrical engineering from the University of California at Berkeley in 1981. For relaxation, he likes skiing, windsurfing, bicycling, dancing, and traveling.

Simcoe Walmsley, Jr.



Simcoe Walmsley is a substrate process engineer in the Santa Clara Division hybrid department. He was responsible for the development of the hybrid substrate processes for the HP 5185A Waveform Recorder. Previously, with the Optoelectronics Division,

he was a production supervisor for gallium arsenide and gallium phosphide crystals. He studied chemistry at San Francisco City College and worked in thin-film development for Singer Research for four years before joining HP in 1971. A native of Phoenix, Arizona, Simcoe is married, has two sons, and enjoys bicycling, camping, and traveling.

39 — Design for Dynamic Performance —

Steven C. Bird



Now a development engineer with HP's Information Technology Group, until last October Steve Bird was with the Santa Clara Division, where he was responsible for the design of the high-speed digital system of the HP 5185A

Waveform Recorder, including the 250-MHz SAW oscillator, time base, trigger, and real-time control. He joined HP in 1979 after receiving his BSEE degree from California State Polytechnic University at Pomona. A member of the IEEE, he specializes in high-speed digital design, low-phase-noise oscillator design, and printed circuit technology. He's named an inventor on a patent on the dropout trigger for the HP 5185A. Steve was born in Cut Bank, Montana. He's married and has two children. Raquetball and church and community activities occupy much of his time.

Brian J. Frohring



Brian Frohring joined HP's Santa Clara Division in 1978 after receiving his BS degree in electrical engineering from the Massachusetts Institute of Technology. He has done ADC research and development, helped develop

HP 5185A Waveform Recorder test system software, and served as project leader for HP 5185A front end and data deceleration IC development. He's a member of the IEEE. Brian is a native of Cleveland, Ohio. He likes hiking and camping, plays piano, and studies French "for fun."

53 — Waveform Analysis, Display, and I/O —

Douglas C. Nichols



Doug Nichols joined HP's Santa Clara Division in 1979 as a product marketing engineer. Three years later, he moved to the R&D lab, and is now temporarily back in marketing, supporting the precision digitizing oscilloscope product line. In the lab, he served as project

leader and wrote the analysis firmware for the HP 51089A Analysis, Display, and Input/Output Module. His special interest is algorithm development and implementation. A native of Ann Arbor, Michigan, Doug attended the University of Michigan, graduating in 1979 with a BSEE degree. He is married, has two children, and plays golf, guitar, and piano.

Bruce E. Peetz



Bruce Peetz is a project manager with HP's Santa Clara Division. Since joining HP in 1977, he has contributed to the design of the HP 5180A Waveform Recorder, been responsible for quantizer design for the HP 5185A Waveform Recorder, and served as project

manager for the HP 5185T Oscilloscope. He received his BS degree in electrical engineering from the Massachusetts Institute of Technology in 1973. Before coming to HP, he was with Hughes Aircraft Company, involved in radar receiver design. He specializes in data conversion and analog design, and has authored several articles on analog-to-digital conversion and testing. Bruce comes from Santa Monica, California, is married, and has three daughters. He plays piano and guitar and is interested in European history and cartography.

49 — Packaging a 250-MHz ADC —

Patrick D. Deane

Author's biography appears elsewhere in this section.

Farid Dibachi



Farid Dibachi holds a 1981 BS degree in mechanical engineering from the University of Tennessee, a 1982 MSME from Stanford University, and a 1986 Master of Engineering degree in electrical engineering from Cornell University earned on HP's resident

fellowship program. He joined HP in 1982 as an IC process engineer. With the Santa Clara Division R&D lab since 1984, he has done hybrid development and microwave sampler design and was responsible for the mechanical and thermal design of the hybrids in the HP 5185A Waveform Recorder. He's a native of Tehran, Iran.

65 — Developing HP PCDS —

Elaine C. Regelson



Currently a project manager of the HP PCDS architecture and library team at HP's Electronic Design Division, Elaine Regelson has contributed to the development of a number of technical-computer-aided work tools, among them HP

TechWriter and various parts of HP PCDS. Before joining HP in 1981, she was a project manager for advanced speech recognition and generation training systems for air traffic controllers. She holds a BA degree in biology awarded in 1972 by the University of California at San Diego and is the coauthor of an earlier HP Journal article on HP TechWriter. Elaine was born in Santa Monica, California and now lives in Fort Collins, Colorado with her husband, three step-children, and five-month-old son. She is interested in folk and traditional country dancing and enjoys storytelling, reciting poetry, singing, gardening, and playing the concertina. She serves on the advisory board of the Colorado State Science Fair and as a volunteer at a community crisis and information center.

68 Automating Circuit Board Design

Gary Jackoway



Interested in algorithms and artificial intelligence, Gary Jackoway is a project manager for design automation tools for HP PCDS. Among other products, he has worked on HP Spice for the HP 1000 Computers and the Autorouter Module for HP PCDS. With HP since

1979, he holds a BS degree in mathematical sciences (1979) from Stanford University and an MA degree in computer science (1984) from Duke University. Gary is also coauthor of an article on a new autorouting technique. Born in St. Louis, Missouri, he now lives in Fort Collins, Colorado with his wife. He enjoys bridge and is currently Unit President for the American Contract Bridge League in northern Colorado.

71 Design System Manager

Mark E. Mayotte

Author's biography appears elsewhere in this section.

Paul S. Reese



Before joining HP in 1985, Paul Reese was an assistant editor for the *Journal of Technical Writing and Communication*. Interested in print and video-based promotional materials, he was responsible for the user documentation and training materials for the HP

PCDS Design System Manager. Paul is currently responsible for sales literature related to HP's electronic design automation products. A native of Burlington, Vermont, he studied English at the University of Vermont (BA 1984) and technical writing at Rensselaer Polytechnic Institute (MS 1985). He

also wrote an earlier article on HP PCDS for HP's *DesignCenter* magazine. He is married, has two daughters, and lives in Fort Collins, Colorado. In his spare time, he enjoys playing with his kids, restoring his 1890's home, and strumming on a guitar.

77 Multidevice Spooler

Deborah A. Lienhart



Born in Berkeley, California, Debbie Lienhart studied geography with an emphasis on cartography at Humboldt State University (BA 1977) and computer science at Colorado State University (MS 1985). She worked as a technical illustrator for five years before joining HP in 1983. At HP Debbie has worked on BASIC applications software and contributed to the development of the Design System Manager and spooler for HP PCDS. She is a member of the ACM special interest groups on design automation and computer and human interfaces. Debbie also serves as an HP Visiting Scientist at a local elementary school. Married to another HP employee, she lives in Fort Collins, Colorado, and enjoys weaving, quilting, raising llamas, and playing the flute at her church.

80 Integrating Applications

Mark E. Mayotte



Born in Sacramento, California, Mark Mayotte went to Boulder, Colorado to study applied mathematics with computer science at the University of Colorado (BS 1980). He returned to California after graduation to join HP and worked on RAPID/3000, a transaction processing system. More recently, Mark was the principal developer of the Design System Manager for HP PCDS. He is author of an article on data management systems and a member of the IEEE. Now living in Fort Collins, Colorado, Mark is married, has two sons, and enjoys playing volleyball, skiing, woodworking, and improving his golf game.

84 Software QA

David E. Martin



Dave Martin received a BS degree in secondary education from New Mexico State University in 1976 and taught high-school chemistry for two years before returning to earn an MS degree in computer science awarded in 1980. He then joined HP's Disc Mem-

ory Division and administered alpha-site testing of the predecessor to HP PCDS. He developed data base utilities for that project and then became QA manager for HP PCDS and contributed to the user interface for the Design System Manager and the installation of HP PCDS. Born in Mesa, Arizona, he now lives in Fort Collins, Colorado with his wife and two children. Outside of work Dave enjoys volleyball, golf, skiing, and teaching Bible class at his church.

87 SOI

Jean-Pierre Colinge



Jean-Pierre Colinge is the author or coauthor of more than 45 papers concerned with SOI and 3D integrated circuit technologies. Interested in IC manufacturing techniques, particularly in SOI, his work has resulted in three French patents and two pending U.S. patents. Jean-Pierre holds an EE degree (1980), a bachelor's degree in philosophy (1980), and a PhD in applied sciences (1984), all awarded by the Université Catholique de Louvain. He joined the technical staff of HP Laboratories in 1985 and has worked on bipolar-CMOS processing and SOI and 3D materials and devices. He is a member of the IEEE Technical Program Committee on SOS and SOI technology and the Scientific Committee of the European SOI Workshop. Born in Brussels, Belgium, Jean-Pierre now lives in Palo Alto, California with his wife and son.

Hewlett-Packard Company, 3200 Hillview
Avenue, Palo Alto, California 94304

HEWLETT-PACKARD JOURNAL

February 1988 Volume 39 • Number 1

Technical Information from the Laboratories of Hewlett-Packard Company

Hewlett-Packard Company, 3200 Hillview Avenue
Palo Alto, California 94304 U.S.A.
Hewlett-Packard Central Mailing Department
P.O. Box 529, Stairbaan 16
1180 AM Amstelveen, The Netherlands
Yokogawa-Hewlett-Packard Ltd., Sugnam-Ku, Tokyo 168 Japan
Hewlett-Packard (Canada) Ltd.
6877 Goreway Drive, Mississauga, Ontario L4V 1M8 Canada

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

CHANGE OF ADDRESS: To subscribe, change your address, or delete your name from our mailing list, send your request to Hewlett-Packard Journal, 3200 Hillview Avenue, Palo Alto, CA 94304 U.S.A. Include your old address label, if any. Allow 60 days.