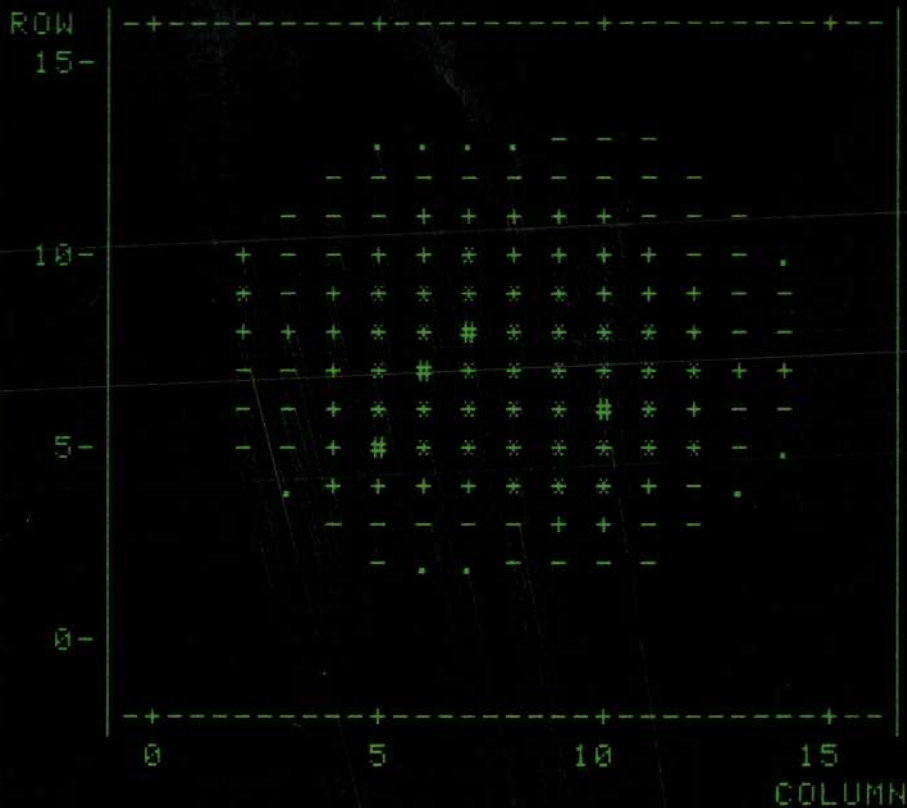



HEWLETT-PACKARD JOURNAL

JUNE 1984

PROG=TC10 ID=LAM LOT#HMC06 WAFER#3
PPG FILE"STEP:INTER"



VAR NAME | VTH T2 |
CHIPS/WAFER = 132
**** LIMIT ****
>= 0.560E+00
* >= 0.555E+00
+ >= 0.550E+00
- >= 0.540E+00
.< 0.540E+00
N/A



Contents:

3 A Parametric Test System for Accurate Measurement of Wafer-Stage ICs, by Yoh Narimatsu and Keiki Kanafuji *Moving this system's 48-pin switching matrix close to the device under test simplifies test cabling and allows accurate current measurements down to 1 pA.*

9 Powerful Test System Software Provides Extensive Parametric Measurement Capability, by Takuo Banno *Setting up wafer tests and probe patterns for the HP 4062A Test System is made easier by this standard set of test instructions and probing utilities.*

12 A High-Speed 1-MHz Capacitance/Conductance Meter for Measuring Semiconductor Parameters, by Tomoyuki Akiyama and Kenzo Ishiguro *This instrument can perform high-resolution C-t and C-V measurements by itself or as part of a large test system.*

24 Authors

26 An Electronic Tool for Analyzing Software Performance, by Gail E. Hamilton, Andrew J. Blasciak, Joseph A. Hawk, and Brett K. Carver *A plug-in option for the HP 64000 Logic Development System, this subsystem lets the software designer measure module execution time, monitor program activity, and collect data about transfers between modules.*

33 Counter Module Simplifies Measurements on Complex Waveforms, by Donald J. Smith, Johnnie L. Hancock, and Thomas K. Bohley *Expanding the feature set of the HP 1980A/B Oscilloscope Measurement System, this plug-in gated counter measures frequencies, periods, and time intervals and counts events.*

In this Issue:

Integrated circuit chips are formed dozens or hundreds at a time on thin discs of silicon called wafers. After processing, the individual chips are cut apart. Those that work (there are always some that don't) are packaged and sold, representing revenue for the manufacturer. Those that don't work are thrown away, representing costs but no revenue. The percentage of chips that work is called the yield. Naturally, increasing production yields is high on every semiconductor manufacturer's priority list. To this end, test element groups consisting of diodes, transistors, resistors, and capacitors are formed at strategic locations on IC wafers. Measurements on these groups provide valuable information about the production process, information that can lead to improvements in materials, layout, and processing, and ultimately to higher yields. The subject of the article on page 3 is a system designed to make and process these measurements. Our cover design this month shows a typical IC wafer and the processed results of some measurements made on it by the HP 4062A Semiconductor Parametric Test System. Among the contributions of the 4062A are its low-level current and capacitance measurement capabilities, its test head design, and its software. The system has one-picoampere sensitivity (that's 0.000000000001 ampere), no small engineering accomplishment when you consider that currents this small can be found running around on the surface of many circuit boards without regard to the metal traces that currents are supposed to follow. The test head design minimizes errors by putting the switching matrix for the measurement pins as close as possible to the wafer, and the software (see page 9) helps the customer write test programs and process data easily.

An important element of the 4062A is its capacitance measurement subsystem, the 4280A C Meter/C-V Plotter, another HP instrument (page 12). The 4280A contributes low-level and very fast capacitance and conductance measurements. It's able to plot capacitance transients that show imperfection concentrations and other characteristics of semiconductor devices.

On pages 26 and 33 are articles about enhancements to two products we've featured in previous issues—the HP 64000 Logic Development System and the HP 1980A/B Oscilloscope Measurement System. The HP 64310A Software Performance Analyzer helps improve the performance of microcomputer-based products by revealing where software redesign might eliminate bottlenecks. The HP 1965A Gated Universal Counter automates universal counter and gated timing measurements, thereby turning the 1980A/B into a much more comprehensive measurement system.

-R. P. Dolan

A Parametric Test System for Accurate Measurement of Wafer-Stage ICs

Special test instruction software and a hardware system fully characterized up to the measurement pins of the test head make this system easy to use for accurate parametric evaluations.

by Yoh Narimatsu and Keiki Kanafuji

THE CONSTANTLY GROWING FIELD of semiconductor technology has created an increased demand for more capable measurement systems to characterize the electrical performance of semiconductors more accurately and conveniently.¹ To keep pace with the industry's needs, the HP 4062A Semiconductor Parametric Test System (Fig. 1) was developed to provide greater measurement flexibility and to eliminate inconsistencies associated with the measurement of semiconductor wafers, chips, and packaged devices.

The 4062A's design is based on HP's state-of-the-art technologies in low-current and capacitance measurements. Overall system measurement performance (test instruments, cables, and the switching matrix) is completely specified, electrically and mechanically, to obtain dependable measurement results. The 4062A's dc and ac measurement characteristics are fully specified up to the measure-

ment pins of the switching matrix. Since the switching matrix (acting as the test head of the wafer prober) is placed close to the device under test, residual parasitics are very small and are readily predictable.

The software for the 4062A's controller, the HP 9000 Model 236 Computer (formerly designated the HP 9836), is written in HP's powerful BASIC language. This language system, in conjunction with the newly developed test instruction set (TIS), provides the user with a maximum degree of flexibility for writing measurement procedures. TIS is a set of subprograms that provides control of the measurement, the switching matrix, and the wafer prober. Each TIS subprogram can be used as if it were a BASIC statement.

Another special feature is the virtual front panel (VFP). This feature allows measurement results to be displayed on the Model 236 Computer's screen in the VFP monitor mode, and the VFP control mode allows system settings




Fig. 1. The HP 4062A Semiconductor Parametric Test System makes it easier to measure dc and ac parameters of semiconductor devices more accurately, both in the wafer state and after packaging.

to be modified manually via the keyboard of the controller. This is a great help for program debugging and for one-time measurements, because programming is not required.

Because parametric test systems are used in environments where new device technologies, new circuit implementations, and new process approaches are tried from day to day and from user to user, HP has chosen a fundamental software policy of not offering specific application software packages to satisfy the needs of only a few customers, but to offer helpful programming tools so that any customer can simply and easily construct test programs.

Features of the 4062A include:

- A wide range of dc voltages and currents can be either sourced or monitored at any of the 48 measurement pins. Currents of less than ± 1 pA up to ± 100 mA and voltages of ± 1 mV up to ± 100 V can be sourced or monitored on an automatic wafer prober.
- Accurate capacitance and conductance measurements at 1 MHz between any two of the 48 measurement pins. Parasitic effects caused by cables, switches, etc. (except for the residual capacitance between the two selected measurement pins) are fully compensated to obtain a basic accuracy of $\pm 0.5\%$ at the 1-MHz test frequency.
- A fully guarded personality board to ensure low-leakage interfacing between the test head and the user's probe card
- A variety of switching matrix test fixtures for measuring packaged devices and discrete devices, and for other special applications
- A test instruction set to provide easy control of the measuring instruments and wafer prober, and to enable measurement programs to be written easily when used along with the controller's friendly, yet powerful BASIC language system
- A virtual front panel to provide the user with a convenient, flexible man-to-machine interface
- Utility subprograms to measure some of the more frequently required parameters such as threshold voltage for MOS transistors, breakdown voltage, and h_{FE}
- Data manipulation utility subprograms to generate wafer

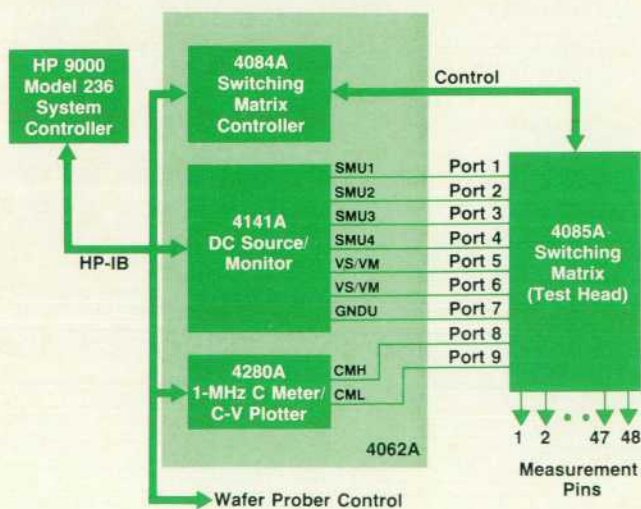


Fig. 2. Block diagram of the 4062A.

maps, histograms, and trend charts

- Complete system diagnostics, including isolation of faulty relays and system components, and verification of system performance.

Hardware Architecture

The block diagram of the 4062A is shown in Fig. 2. The HP 4141A DC Source/Monitor (dc measurement subsystem) and the HP 4280A 1-MHz C Meter/C-V Plotter (capacitance measurement subsystem) are the key instruments, taking care of dc current/voltage measurements and ac capacitance/conductance measurements, respectively. These two instruments are equal to a total of nine instrumentation ports. To multiplex these instrumentation ports into the 48 measurement pins, the HP 4085A Switching Matrix (switching matrix subsystem) is provided. The relays in the 4085A are controlled by the system controller via the 4084A Switching Matrix Controller. The 4084A converts the HP-IB (IEEE 488) relay control codes from the system controller into machine-oriented serial signals. The dc power necessary to actuate the relays in the 4085A is also supplied by the 4084A. All system components, including the wafer prober, are controlled by the 4062A's system controller via the HP-IB.

A great deal of engineering effort was concentrated on the design of the 4085A Switching Matrix. To provide the best available performance for both low-current and capacitance measurements, the 4085A's design avoids using com-

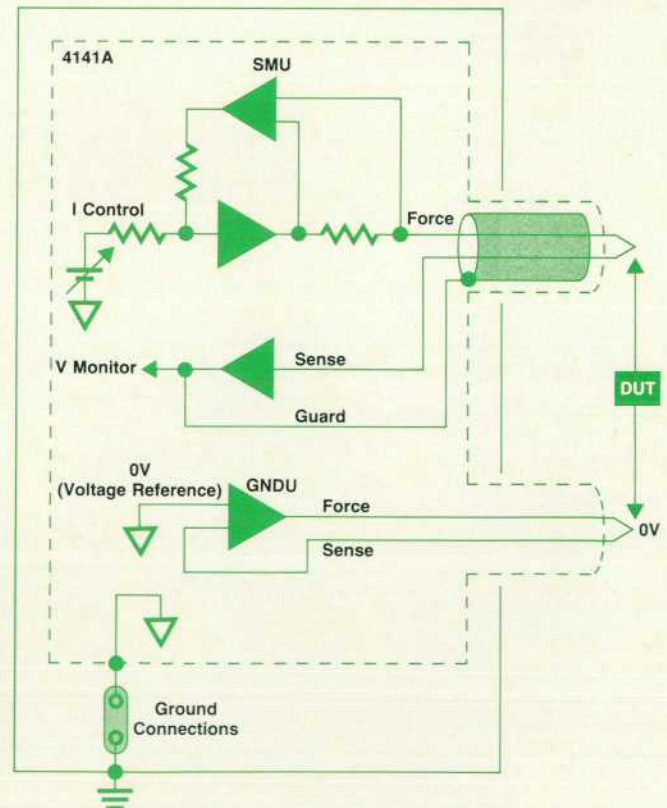


Fig. 3. Connection of the 4141A DC Source/Monitor for measuring the voltage across a DUT driven by a constant current.

Some Examples of 4062A Applications

Measurement data obtained with the 4062A can be arranged into many forms. Shown in Fig. 1 through Fig. 4 are the results of several basic semiconductor measurements.

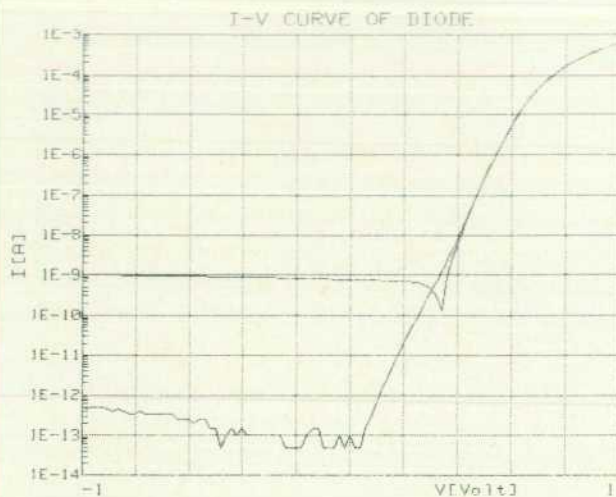


Fig. 1. The current-versus-voltage characteristics of a pn junction diode fabricated on a silicon wafer can be directly measured on the wafer by using the 4062A and a wafer prober. Absolute values of current are used to permit logarithmic scaling of the measurement results. With the pn junction reverse-biased, leakage current increases by several nanoamperes (upper curve) when the microscope lamp provided with the wafer prober's microscope is turned on. Current in the range of picoamperes (lower curve) was measured with the microscope lamp off, but with the room lights still on.

DEVICE	V_t	V_t	V_t	V_t
	$V_{bg}=0$	$V_{bg}=-1$	$V_{bg}=-2$	$V_{bg}=-9$
5/5 DRIVER	1.18	1.80	1.99	2.98
5/5 LOAD	-7.35	-4.24	-3.80	-2.65
3/4 LOAD	-7.80	-3.76	-3.28	-2.10
5/2 DRIVER	1.20	1.35	1.80	2.30
DELTA W LOAD _s =	2.54 μ m			
DELTA W DRIVER _s =	1.59 μ m			
DELTA L DRIVER _s =	0.12 μ m			
DELTA L LOAD _s =	1.25 μ m			
BV _{dss} =	8.75V			

Fig. 2. By measuring various test patterns, various semiconductor process parameters can be monitored and analyzed. The above list shows one example of process parameter monitoring.




Fig. 3. The capacitance-versus-voltage characteristics of an MOS capacitor on a silicon wafer are shown here. Slight variations in capacitance can be accurately detected.




Fig. 4. These plots of MOS transistor gate voltage versus drain current were obtained while varying the backgate bias voltages applied to the test device. The test device was an n-channel MOSFET on a silicon wafer. Currents of several picoamperes can be clearly resolved.

mon signal paths for both dc and ac (capacitance/conductance) measurements, although configuration flexibility is somewhat degraded. In the 4085A, each signal path serves only one function and is assigned to only one of the instrumentation ports so that its electrical characteristics can be clearly characterized and defined. This is essential to guarantee the total performance of the system. The guarding technique and the routing of signal paths to allow measurement of currents as small as 1 pA at any desired pin required a lot of new ideas and skills.

Designing new signal cables between the instruments

and the switching matrix (test head) was another important job. Ordinary twisted-pair cables introduce a significant amount of current noise caused by frictional electricity generated between the inner insulator and the guard conductor when the cable is flexed. To eliminate this effect and to reduce the capacitance between the signal conductor and the guard conductor, which also introduces noise, a special cable was developed in cooperation with a cable manufacturer. The cable between the capacitance meter and the test head is also special, because cable characteristics such as inductance, capacitance, and resistance per

unit length play a large part in determining the accuracy of capacitance measurements. Consequently, these parameters must be well controlled. Such design requirements resulted in relatively heavier-gauge and more expensive signal cables. Because 48 of these cables would have been impractical to use for connection between the switching matrix and the usual tiny probe card, the test head approach where the switching matrix is moved close to the DUT was a much better choice. Some of the many advantages of this test head approach are:

- Different types of cables for dc measurements and for capacitance measurements can be properly used as described above.
- The section of cabling where the dc signals and ac signals (capacitance measurement signals) must share the same signal path (which cannot be practically designed to satisfy all requirements) can be kept to a minimum.
- The switching circuits, including the relays and the motherboard in the test head, which are subject to the effects of external vibration and electromagnetic interference when making measurements at extremely small signal levels, can be placed apart from the system instruments and computer—likely sources of such external noise.
- The test head acts as an electrostatic shield and as a fairly good light barrier between room lights and the device under test when it is mounted on a wafer prober. Measurement results on the order of 1 pA for pn junction reverse-leakage current (see example in the box on page 5) were obtained for a wafer-stage junction diode measured on an automatic wafer prober. The data was taken in an environment where fluorescent lamps provide illumination. The current increases several nanoamperes if the wafer prober's microscope illuminator lamp is turned on.

The test head is designed so that it can be easily mounted onto most wafer probers.

DC Measurement Subsystem

The HP 4141A DC Source/Monitor, referred to as the dc

measurement subsystem (DCS), takes care of the 4062A's dc measurements and requirements. The 4141A is a modified version of HP's 4145A Semiconductor Parameter Analyzer.² The 4145A's CRT, keyboard, and flexible disc drive were removed, keeping four SMUs (stimulus/measurement units), two voltage monitors, and two voltage sources. Using this as a base, the following features were added to the 4141A:

- Kelvin (four-terminal) connection capability for the SMUs. This allows the unwanted voltage drops caused by current flowing through the measurement cables to be eliminated.
- Ground unit (GNDU) with Kelvin connection capability. A true ground (common) potential can be defined at the DUT (not at the instrument), and the GNDU's 500-mA source/sink capability ensures a ground potential of 0V even if the four SMUs source or sink dc currents of 100 mA each (total of 400 mA) through the test device.
- Firmware suitable for system operation. The internal firmware controlling the 4141A is designed so that the best performance can be obtained when the 4141A is controlled by an external system controller via the HP-IB. Since the outermost conductor of the triaxial output connector of each SMU is at guard potential and may possibly exceed $\pm 42V$, a special cable assembly with a connection box at each end is provided. This cable assembly not only prevents a user from touching hazardous voltages, but also provides a convenient method of connection to the 4085A Switching Matrix and eliminates the possibility of improper cable connection.

Quadraxial cables (twisted-pair center conductors surrounded by two shield conductors) must be used to obtain the best performance from the SMUs. The insulation of the two center conductors is coated with a special conducting material to ensure noise-free measurements. The two center conductors are surrounded by the guard, which is essential to achieve low-leakage, high-speed measurements. The outermost conductor of each quadraxial cable is kept at the ground potential (chassis potential) of the system so that measurement results will be immune to any externally gen-


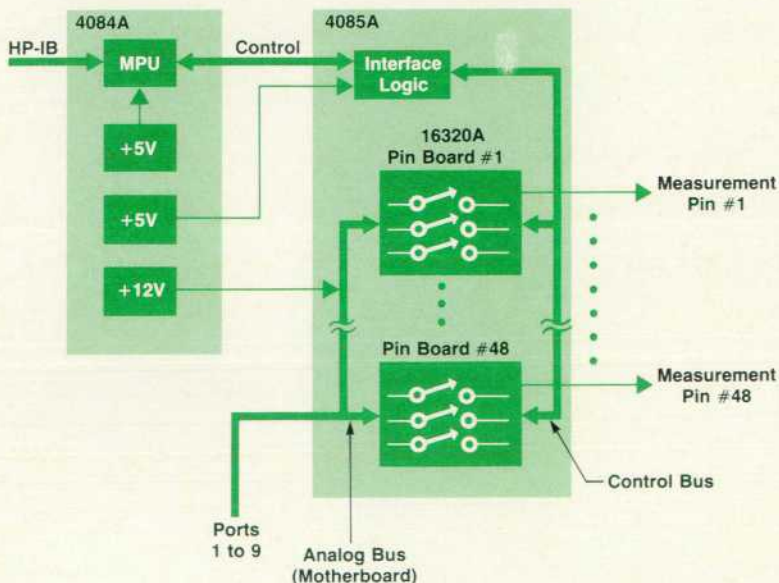


Fig. 4. Simplified block diagram of the switching matrix subsystem consisting of the 4085A Switching Matrix (photo) and the 4084A Switching Matrix Controller.

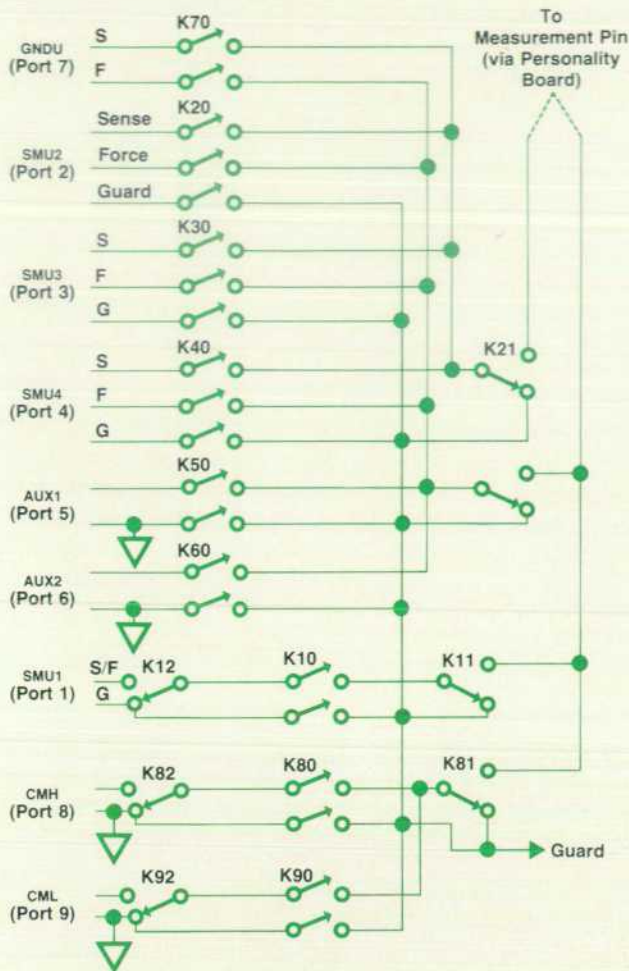


Fig. 5. Relay connections for the HP 16320A Pin Board.

erated noise, and to prevent the system from radiating any electromagnetic interference.

Switching Matrix Subsystem

Fig. 4 shows a simplified block diagram of the 4084A Switching Matrix Controller and the 4085A Switching Matrix, referred to as the switching matrix subsystem (SWM). The 4084A provides the power for the relays and relay control circuitry of the 4085A. The output noise of the switching regulators in the 4084A is kept to a minimum so that low-level dc and ac measurements are not affected.

Relay control statements sent from the 4062A's system controller are applied to the 4084A via the HP-IB. In the 4084A, each statement is interpreted into a logic signal to control a specified relay in the 4085A. The signal is then serially fed to the 4085A, where it is routed to the pin board (HP 16320A) containing the specified relay. The relay control signal contains pin board address information, and a custom gate array IC on each pin board compares the information with its own board address (which corresponds to the measurement pin number). Each pin board's address is determined by the physical location of the pin boards in the switching matrix, because address data for each pin board position is provided at each slot on the motherboard. Since all pin boards are identical, the pin board configuration can be rearranged whenever necessary. Optoisolators between the 4084A and 4085A ensure that no logic noise from the 4084A interferes with the 4085A.

One of the key parts of the 4085A is the motherboard, because a majority of the analog bus lines pass through it. Although the motherboard is of ordinary glass-epoxy construction, various investigations were made and experiments were performed to determine the best possible construction for the motherboard to keep its insulation resistance at the highest possible value. An insulation resistance of more than 10^{10} ohms—even at an ambient operating temperature of 40°C and a relative humidity of 70%—has been achieved by employing guarding and using a special coating. Thus, currents for which the allowable leakage is

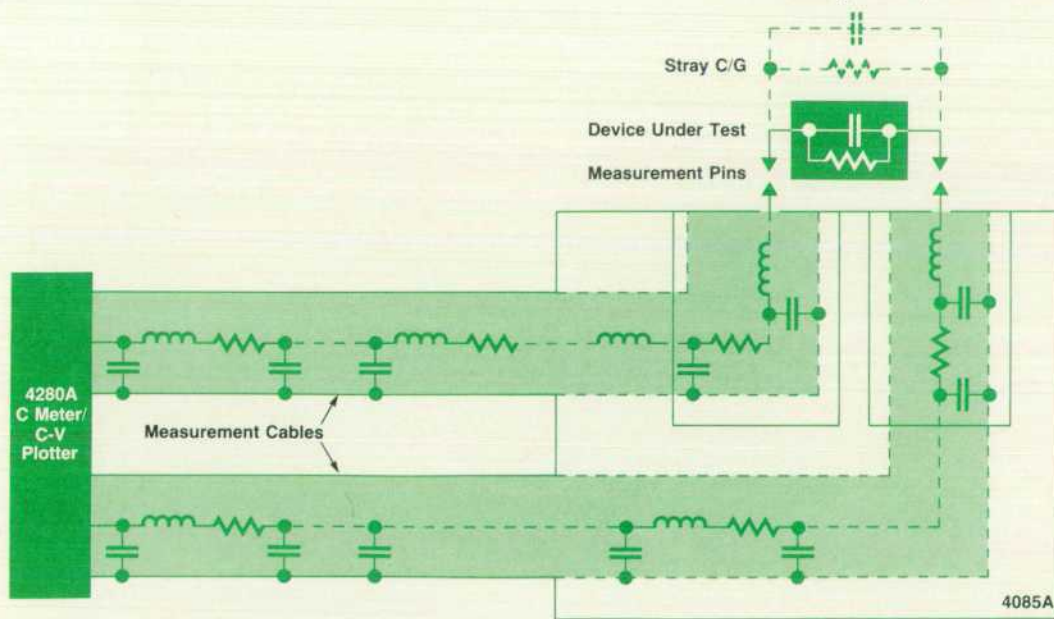


Fig. 6. Between the capacitance measurement subsystem (HP 4280A C Meter/C-V Plotter) of the 4062A and the DUT, a number of parasitic components exist. The test software of the 4062A can be used to correct for the presence of these values.

less than 10 nA when the applied voltage is $\pm 100\text{V}$ can be routed through the motherboard (ports 2, 3, and 4).

The allowable leakage specification for port 1, on the other hand, is 2 pA at 100V. This requires an insulation resistance of more than 5×10^{13} ohms. Because the ordinary glass-epoxy material cannot be used to achieve this order of insulation resistance, a special bus line insulated with Teflon™ was designed. A spring-loaded pin is used to provide electrical contact between the bus line and each pin board. The spring-loaded measurement pins of the pin boards are also insulated with Teflon.

Pin Boards

Fig. 5 shows the relay connections on the HP 16320A Pin Boards. Not only the sense and the force signal lines, but also the guard lines are switched by 3A-type relays to realize a guarded Kelvin (four-terminal) connection up to the measurement pins. An insulation resistance on the order of 10^{13} ohms is provided. To obtain such a high insulation resistance, selected reed relays are used and a special coating is applied to them by the relay manufacturer.

For port 1, where allowable leakage current is less than 1 pA, additional transfer relays are incorporated. That is, when port 1 is off, relay K12 guards the signal line. When port 1 is on (K12, K10 and K11 are closed), relays K21 and K81 provide the guard potential around the signal line. Thus, the signal line concerned always has the proper guard potential, which essentially ensures an extremely small leakage current, independent of the applied signal voltage.

To ensure maximum relay life, "dry" switching of the reed relays is used in the 4085A. A more detailed discussion of this technique is given in the article on page 9. Experimental results indicate that the life of reed relays is 100 to 1000 times longer if the dry switching method is used instead of the ordinary "wet" switching method.

Capacitance Measurement Subsystem

The HP 4280A C Meter/C-V Plotter (see article on page 12) acts as the capacitance measurement subsystem (CMS) of the 4062A.

Because long measurement cables and the switching matrix are inserted between the 4280A and the test sample, various parasitic impedances and admittances will be introduced when a capacitance or conductance measurement is made. These parasitic components cannot be neglected—especially at 1 MHz—and consequently affect measurement results. Among the parasitics are series inductance

and resistance and parallel capacitance and conductance of the connection cables between the 4280A and the 4085A, and of the motherboard and the pin boards of the 4085A (Fig. 6). The 4280A automatically measures and stores these values when the START program of the 4062A is executed, and the TIS subprograms automatically compensate for their effects when a test device is measured.

Note that the stray capacitance and conductance between measurement pins is not eliminated when the START program is executed, because the 4062A's system controller does not know which measurement pins will be used for capacitance measurements, and also because the test device must be manually removed from the socket when a test fixture is used. Therefore, a user must first measure the stray capacitance and conductance between the specified measurement pins without a test sample connected, then measure the capacitance and conductance of the test sample plus the parasitics, and finally subtract the values of the parasitic components from the total measured value. This is the way individual capacitance measurements are commonly performed. When an automatic wafer prober is used for capacitance measurements, the aforementioned procedure can be performed automatically by using a simple program and the help of the TIS.

Acknowledgments

The 4062A is a result of real team effort. Design engineers who deserve to be acknowledged are Takeshi Kyo, Akira Fujita, Teruo Takeda, Hirokazu Sugawara, Hideyuki Hasegawa, Hitoshi Imaizumi, Katsuhiko Ishiguchi, Yoshiyuki Kudo, Keita Gunji, and Naomi Miyata. Mechanical and industrial design was provided by the mechanical engineering team: Hiroshi Shiratori, Hideo Kishimoto, Tsuneji Nakayasu, Yoshimasa Shibata, and Akihiko Goto. Toshio Ichino contributed greatly to the smooth transfer of the system to production. Ken Abiko provided some very helpful leadership in the early investigation stages of the project. Special thanks are due to Mitsutoshi Mori, director, Shiro Kito, lab manager, and Haruo Ito, section manager, for their valuable advice and continuous encouragement.

References

1. G.P. Carver, et al, "Use of Microelectronic Test Structures to Characterize IC Materials, Processes, and Processing Equipment," *Solid-State Technology*, Vol. 23, no. 9, September 1980.
2. K. Maeda, et al, "Intelligent Instrument Streamlines dc Semiconductor Parameter Measurements," *Hewlett-Packard Journal*, Vol. 33, no. 10, October 1982.

Powerful Test System Software Provides Extensive Parametric Measurement Capability

An easy-to-use set of test instructions, "dry" switching of test relays, and a utility for specifying wafer probing patterns provide powerful support for users of HP's semiconductor test system.

by Takuo Banno

PARAMETRIC TESTING has become widely adopted as a method for monitoring semiconductor processing, and many test method approaches have been considered. Therefore, the programming language for a parametric test system must be able to incorporate a diversity of methods and be able to accommodate new techniques.

To meet these requirements, the BASIC language used by the HP 9000 Model 236 Computer (formerly designated the HP 9836) was chosen as the test programming language for the HP 4062A Semiconductor Parametric Test System. This version of BASIC is not only powerful, but also very easy to use. Its on-line editing and debugging capability makes it easier for users to create, modify, and try their own test programs.

Test Instruction Set

The 4062A's test instruction set (TIS) is an extension of BASIC, and adds the capabilities necessary to perform measurements and to control the 4062A's hardware subsystems. TIS statements are provided in the form of compiled subprograms and can be called from a simple BASIC language program. Using these compiled subprograms results in optimum test program execution speeds.

If a user makes a mistake in a test program (e.g., an invalid pass parameter value is specified), the 4062A's language

system (BASIC plus TIS) will report an error message along with the calling line number as shown in the example below. On the other hand, a user subprogram written using only BASIC will report only the line number of the erroneous part of the subprogram.

```
•
•
110 Force_v(FNSmu(2), 120) ! Voltage value must be less than
                             ±100V
•
•
ERROR 392 IN 110           ! 110 = Line number
                             ! 392 = Error number
```

The 4085A Switching Matrix for the 4062A contains many relays. To ensure maximum relay life, a "dry" switching technique has been incorporated into the TIS subprograms. This means that the output of a 4062A source (an SMU of the 4141A or the dc bias of the 4280A) is set to 0V with a current limit of 10 μ A before the switch is actuated (opened or closed). TIS ensures that dry switching will occur, and therefore, frees the user from having to worry about writing this into a program. If a current source is already set to output less than 10 μ A, or if a voltage

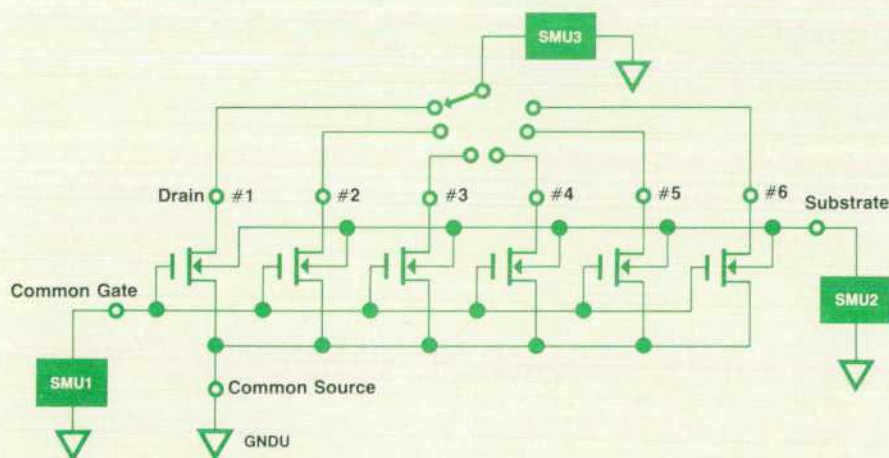


Fig. 1. MOS transistor array for test program example in text.

source's current compliance setting is already less than $10 \mu\text{A}$, the output of the source remains unchanged during switching. Instead of turning off all sources, only the source to which the specified relay is connected and whose output may possibly exceed $10 \mu\text{A}$ is set to 0V with a current limit (compliance) of $10 \mu\text{A}$. Experiments have shown that $10 \mu\text{A}$ is not enough current to cause damage to the relay contacts. Dry switching in this way minimizes the number of source turn-on/turn-off cycles and provides protection for the relays.

The example test circuit shown in Fig. 1 is an array of MOS transistors in which all six transistors share a common gate and a common source. The ground unit (GNDU) is connected to the common source terminal, SMU1 is connected to the common gate terminal, SMU2 is connected to the substrate, and the drain current of each transistor is measured by SMU3. A partial listing of the test program to make these measurements is:

```

210 Connect(FNGnd, Source)
220 Connect(FNSmu(2), Substrate)
230 Connect(FNSmu(1), Gate)
240 Force_v(Substrate, -2)
250 Force_v(Gate, 2)
260 FOR I=1 TO 6
270 Connect(FNSmu(3), Drain(I))
280 Force_v(Drain(I), 5, 0, 0.1)
290 Measure_i(Drain(I), Idrain(I))
300 Connect(FNSmu(3))
310 NEXT I

```

At line 270, SMU3 is automatically turned off before the specified relay is activated. All the other SMUs stay unchanged during the execution of the FOR...NEXT loop (lines 260 through 300). Although not shown in the program listing, dry switching is provided by TIS, relieving the user of such a cumbersome job.

The measurement of test chips on a semiconductor wafer is one of the major applications for the 4062A. To measure the many test chips on a wafer, an automatic wafer prober is required. TIS contains prober control statements that allow the prober's chuck to position a wafer by using X-Y

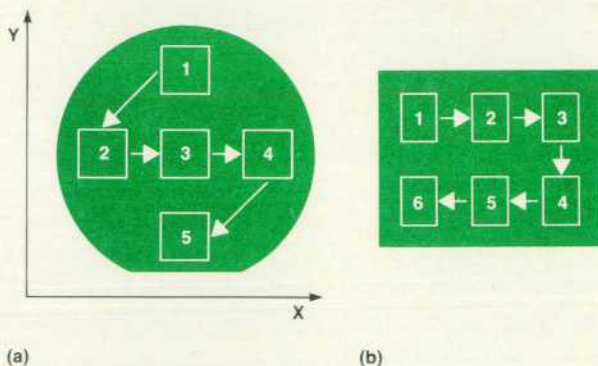


Fig. 2. In this wafer probe pattern example, a user wants to probe five selected chips (a), and each chip contains six modules (subchips) to be probed (b). The PPG utility allows the probing sequence, chip and module addresses, etc. to be determined by the user.

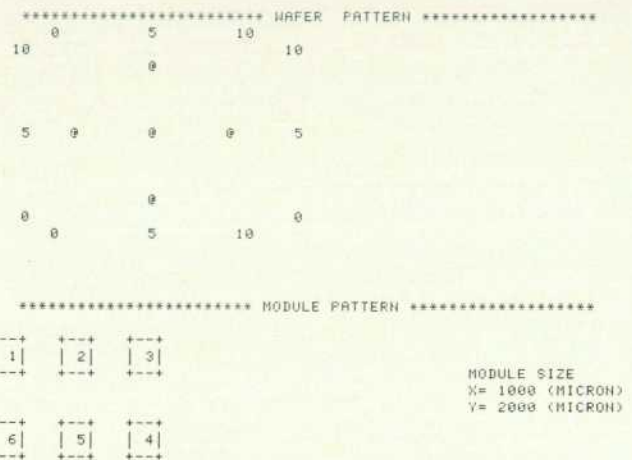


Fig. 3. Display on the CRT of the 4062A's controller for probing example shown in Fig. 2.

coordinate addresses. These TIS statements are similar to those for X-Y plotters:

```

500 Porig                ! Specifies the position of origin.
510 Pscale(1000,2000)   ! Sets the scale for the X axis (1000 μm)
                        ! and for the Y axis (2000 μm).
520 Pmove(2,5)          ! Moves the probe to address (2,5). This
                        ! is 2 mm to the left and 5 mm above the
                        ! position of origin.

```

A parameter-oriented utility library named PARA is also contained in the system software. This library contains a set of subprograms that make it easy to obtain some of the more commonly required dc parameters such as V_{th} , h_{FE} , and others. All PARA subprograms are written in BASIC with TIS, and are easy to understand and modify. PARA is like a "sample" program and can be easily adjusted to suit the requirements of a user.

Virtual Front Panel

The virtual front panel (VFP) is a friendly and interactive front-panel emulator that allows the user to control the 4062A's subsystems and to make measurements without any programming effort. VFP not only makes it possible to control the system as if it were a stand-alone instrument, but also allows a user to monitor the status of the hardware subsystems while the test program is running. At any time during test program execution, the measurement program can be paused, and the 4062A can be manually controlled by VFP in addition to its monitor capability. In this way, VFP is extremely convenient for test program debugging.

VFP is implemented by a BASIC subprogram. A time interval interrupt is required for VFP to be executed concurrently with a test program, and this interrupt is generated by the 16321A VFP Timer installed in the rear of the Model 236 controller.

Probing Pattern Generator

The probing pattern generator (PPG) is an interactive utility program used to set up the probing pattern on a wafer and to create a probing pattern data file. An example of a probing pattern is shown in Fig. 2. A user can assign an

		Observations				
		O ₁	O ₂	O ₃	• • •	O _k
Variables	V ₁	Data	Data	•	•	•
	V ₂	Data	Data	•	•	•
	V ₃	Data	•	•	•	•
	•	•	•	•	•	•
	V _l	•	•	•	•	•

(a)

Record # 1	████████████████████
# 2	████████████████████
# 3	████████████████████
	•
	•
# m	████████████████████

(b)

Fig. 4. (a) 4062A data structure (b) 4062A file structure.

address to a chip either by using the chip's X-Y coordinates or by first assigning a number to each chip and then referring to a chip by its number. The display on the CRT of the 4062A's controller for the example in Fig. 2 is shown in Fig. 3.

A prober control library (PCL) is also available as a utility library to provide control for a wafer prober by using probing pattern data files.

Data File

The structure of a 4062A data file is shown in Fig. 4. This file is an $l \times k$ matrix array whose rows correspond to variables and whose columns correspond to observations. Variables are assumed to be parameters such as breakdown voltage and gate capacitance. Observations are assumed to be chip number, wafer number, etc. The sizes l and k are




Fig. 5. Wafer map produced by the HP 4062A Parameter Test System.

variable. Title, number of variables, number of observations, and variable names are stored in the first record of the data file. The second record and succeeding records contain numeric data. This structure allows any amount of data to be stored into mass storage efficiently.

This data file structure is compatible with that of the HP 98820A Statistical Library. The 98820A Library provides general statistics, statistical graphics, regression analysis, and much more, and can be used for further reduction of a user's collected measurement data.

A wafer map, as shown in Fig. 5, can be generated by using the WAFER MAP program portion of the 4062A's software. With this program, wafer maps based on data contained in the measurement data file and the probing pattern data file can be easily produced.

Acknowledgments

The author would like to express special thanks to Ales Fiala and Steve Schink at HP's Fort Collins Systems Division for their help in developing the system software.

Reference

1. K.Y. Kwinn, R.M. Hallissy, and R.E. Ison, "The 9826A/9836A Language Systems," *Hewlett-Packard Journal*, Vol. 33, no. 5, May 1982.

A High-Speed 1-MHz Capacitance/Conductance Meter for Measuring Semiconductor Parameters

This fast, high-resolution instrument is equipped with a built-in timer, a sweepable dc bias source, and a pulse generator for high-speed C-t and C-V measurements.

By Tomoyuki Akiyama and Kenzo Ishiguro

PRECISE MEASUREMENT of the physical parameters of semiconductor devices in a reasonably short time and with a simple measurement setup is important both in the semiconductor research and development field and in the process control area of semiconductor production. This capability is essential to develop new devices, develop and improve semiconductor process technology, and monitor and control the semiconductor fabrication process properly.

By measuring the capacitance and conductance characteristics of a semiconductor device—such as a pn junction or an MOS structure—we can obtain many of the basic parameters used to explain semiconductor behavior. Capacitance-versus-voltage (C-V) and capacitance-versus-time (C-t) measurements are two of the more common methods of measuring semiconductor parameters.

The HP 4280A 1-MHz C Meter/C-V Plotter (Fig. 1) was developed to satisfy the requirements of semiconductor device/material characterization. Besides being a traditional high-frequency capacitance/conductance meter, the 4280A can also perform C-V and C-t measurements using its internal dc bias source and timer, and its analog output capability allows results to be plotted on an X-Y recorder. The C-t measurement capability combined with the other functions in one box contributes to improved measurement accuracy, speed, and sweep range, ease of operation, and performance/cost ratio. All bias parameters required for a C-V measurement and all time and bias parameters required for a C-t measurement can be remotely controlled via the HP-IB (IEEE 488).

One difference between the 4280A and other HP capacitance and impedance meters is that the 4280A is designed




Fig. 1. The HP 4280A 1-MHz C Meter/C-V Plotter is the first stand-alone instrument capable of both C-V and C-t measurements. If equipped with the optional **C-HIGH RESOLN** function, the 4280A can measure capacitance in 50 ms with 5½-digit resolution.

to satisfy the requirements of a system component while maintaining the features of a stand-alone instrument. The system component concept is based on the assumption that one instrument able to make simple measurements and provide raw data on the device under test will not be able to analyze the results, because analysis of semiconductor devices usually requires both ac and dc parameters, and the conversion of C-V, C-t, and dc parameters to actual semiconductor parameters requires extensive calculations. In addition, the acquisition of the raw data itself sometimes requires a complicated system setup. The C measurement and C-V plot functions of the 4280A are mainly those of a stand-alone instrument, whereas the C-t measurement function is mainly for systems applications. All functional modes are designed for interfacing with a system controller, with other supporting instruments, and with the device under test (DUT).

Stand-Alone Features

Some of the stand-alone features of the 4280A are:

- High accuracy. The 4280A uses a simple two-terminal-pair measurement port and can measure both capacitance and conductance with a basic accuracy of 0.1% over three ranges—10 pF, 100 pF, and 1 nF.
- High resolution and high speed. The 4280A has a maximum 5½-digit display resolution with the **C-HIGH RESOLN** option. The implementation of the automatic C-offset function makes 5½-digit resolution possible without a reduction of measurement speed. The minimum time required for measurements in this mode is only 50 ms.
- Cable compensation and parasitic compensation. The error-causing stray capacitance and conductance of the test fixture and the admittance of connection cables can be cancelled using an internal correction algorithm. Maximum cable length is five meters.
- Floating and grounded connection modes. The 4280A can measure both grounded and floating devices, with a choice of fourteen connection modes.
- C-V measurement. The 4280A has an internal dc bias source that can be programmed from its front panel or via the HP-IB to function as a pulse generator, sweepable (staircase) bias source, or constant dc bias source. Maximum output is $\pm 100\text{V}$ with three-digit resolution over three ranges. The best resolution is 1 mV on the 1V range. Measurements with two independent dc bias sources (the internal dc bias source and one external dc bias source, or two external dc bias sources) connected to the **LOW** and **HIGH UNKNOWN** terminals can be performed.
- C-t measurement. Measurement intervals from 10 ms to 32 s can be set. If an external pulse generator is used, measurement intervals as short as 10 μs can be set. Real-time burst measurements can be used for devices that have slow response. For devices that have fast transient capacitance characteristics, a special sampling technique can be used.
- X-Y recorder outputs. Two Y-axis outputs, one corresponding to the C display and one corresponding to the

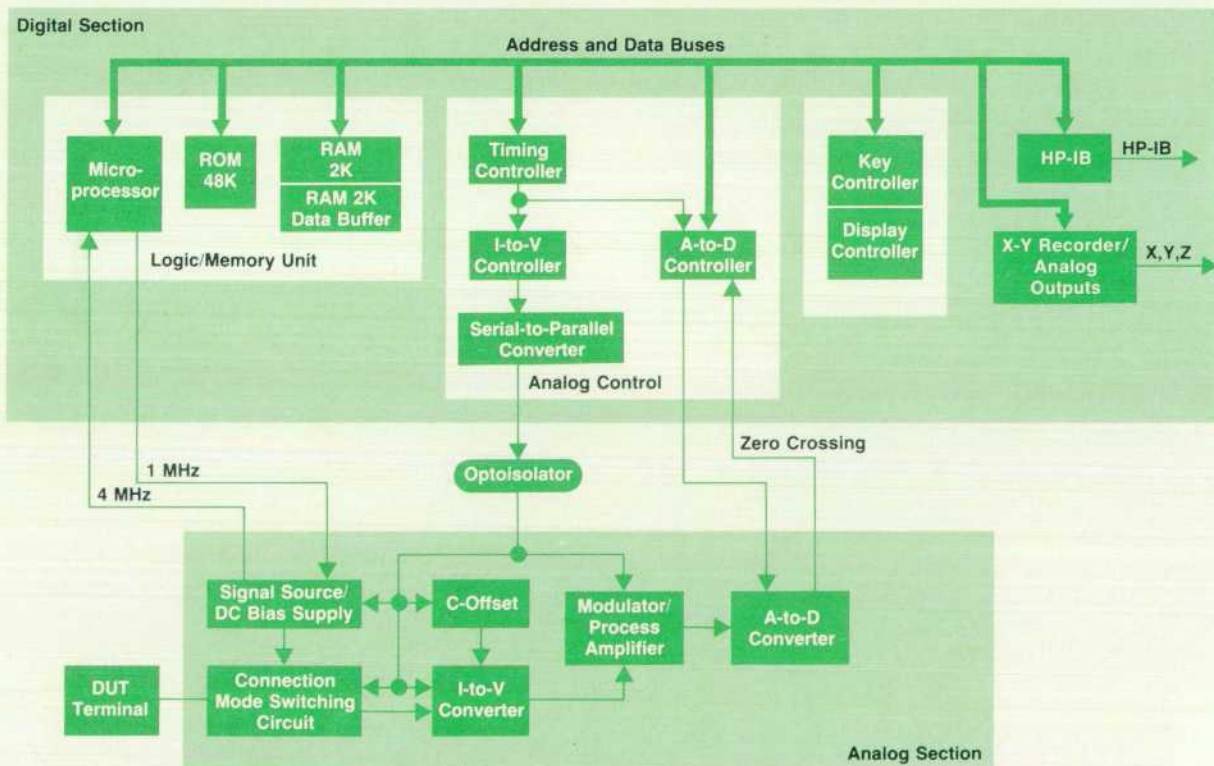


Fig. 2. The block diagram of the 4280A consists of a grounded digital section and a floating analog measurement section. Communication between the two sections is through optoisolators to ensure proper ground isolation.




Fig. 3. The 4280A uses a two-terminal pair configuration for connections to the DUT.

G display, and one X-axis output corresponding to swept voltage, time, or number of triggers are provided. Full-scale output voltage is $\pm 10V$.

- Self-test. An automatic built-in functional test verifies proper operation of the 4280A's analog and digital circuits.
- Operation error codes. About sixty error codes related to operator errors can be displayed.

System Features

Features of the 4280A for systems applications include:

- ASCII or binary output data formats
- Standard or block data transfer modes. The data buffer is large enough to hold the results of 680 C-V and C-t measurements.
- An isolation mode to shut off the test signal source and isolate the internal measurement circuits electrically from the UNKNOWN terminals.

Fig. 2 shows a block diagram of the 4280A. The digital section has five functional blocks: logic/memory unit, analog control, key/display control, HP-IB control, and X-Y recorder/analog output. In the analog section, there are six functional blocks: connection mode switching circuit, signal source/dc bias supply, current-to-voltage (I-to-V) converter, C offset circuit, modulator/process amplifier, and analog-to-digital converter (ADC).

The functions of the analog section are to (1) transduce the admittance of the DUT into a vector voltage referenced to the test signal (this is accomplished by the signal source, I-to-V converter, modulation amplifier, and C-offset circuit), (2) determine the ratios between the vector voltages (process amplifier and ADC), (3) supply the required dc bias voltage to the DUT, and (4) set the proper connection configuration of the dc bias supply, signal source, and I-to-V converter (connection mode switching circuit and internal dc bias source).

The signal source outputs a 4-MHz clock to the 4280A's microprocessor and the microprocessor returns a 1-MHz square-wave signal to the signal source. The signal source

reclocks this 1-MHz signal with the 4-MHz signal and passes this signal to the 1-MHz low-pass filter to generate the test signal. Most of the circuits in the analog section are floating to ensure accurate measurements.

The serial-to-parallel converter circuits process control signals and serial data sent from the digital section to the analog section. These signals are sent via optoisolators to keep noise from the digital section out of the analog section and to maintain proper ground isolation.

The A-to-D controller provides the signals necessary to control the dual-slope ADC in the analog section. It also detects the zero crossing signal output from the ADC integrator at the completion of each integration cycle. Because the ADC requires very fast control, the control lines from the A-to-D controller are tied directly to the ADC; optoisolators are not used.

The display controller controls two groups of large seven-segment displays, which display C and G measurement results, and one group of small seven-segment displays, which displays voltage and time parameters and error codes. The display controller also controls all other front-panel indicators.

The key controller controls all the keys and switches on the front panel. The HP-IB block controls all HP-IB functions.

The X-Y recorder/analog outputs provide dc voltages proportional to the displayed values, and they can be connected directly to an X-Y recorder to make hard copies. The capacitance and/or conductance output and voltage/time output are internally synchronized to maintain the relationship between the C and/or G data and the dc bias or time setting.

The logic/memory unit contains a 6802 microprocessor, 48K bytes of program ROM, and 2K bytes of RAM. In addition, there is a 2K-byte data buffer for block mode measurements. In block mode (only for C-t and C-V measurements), the 4280A executes only a measurement and a data storage operation, and skips display refresh, correction calculation, deviation calculation, and HP-IB data output. For C-t burst measurements in block mode, the 4280A executes a measurement and stores the results in the data buffer within 10 ms. As a result, the 4280A can be used for very fast real-time C-t measurements.

The analog control block has four circuits—timing controller, current-to-voltage controller, serial-to-parallel converter, and analog-to-digital controller. The timing controller generates the signals necessary to control the measurement sequence. This circuit has seven 16-bit programmable

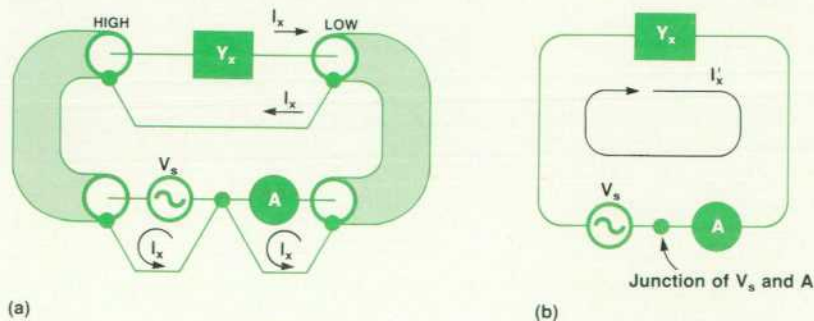


Fig. 4. The actual two-terminal pair configuration of the 4280A is as shown in (a), but the current I_x flowing as shown in (b) is very small. Since negligible current flows through the junction of V_s and the ammeter, configuration (a) is equivalent to Fig. 3.

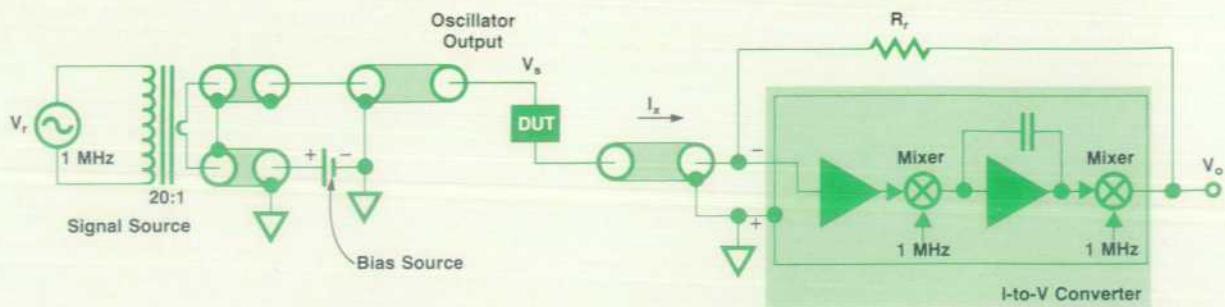


Fig. 5. The measurement transducer consists of a signal source and an I-to-V converter.

timers to control measurements. The time base generator, one of the timers, generates a 0.5-ms clock for C-t measurements and a 10-ms clock for C-V measurements, based on the 1-MHz system clock. This means that the 4280A has the same time interval accuracy as the system clock. These three clock sources (1 μ s, 0.5 ms, and 10 ms) provide a wide range of time intervals. For C-t measurements, time intervals from 10 μ s to 32 s can be selected. For C-V measurements, the range is from 3 ms to 650 s.

The I-to-V controller receives timing signals from the timing controller and produces the actual control signals for the I-to-V converter in the analog section. For sampling

mode measurements used for fast C-t measurements, precise timing signals are required by the I-to-V converter. These circuits generate signals with 1- μ s resolution.

Measurement Techniques

The 4280A uses a two-terminal-pair configuration (Fig. 3) because it is less sensitive to measurement cable parameters. The advantage of the two-terminal-pair configuration over a four-terminal-pair configuration is its ability to handle the narrow pulses (less than 10- μ s duration) required for high-speed C-t measurements. This is possible because the two-terminal-pair configuration uses 50 Ω coaxial ca-

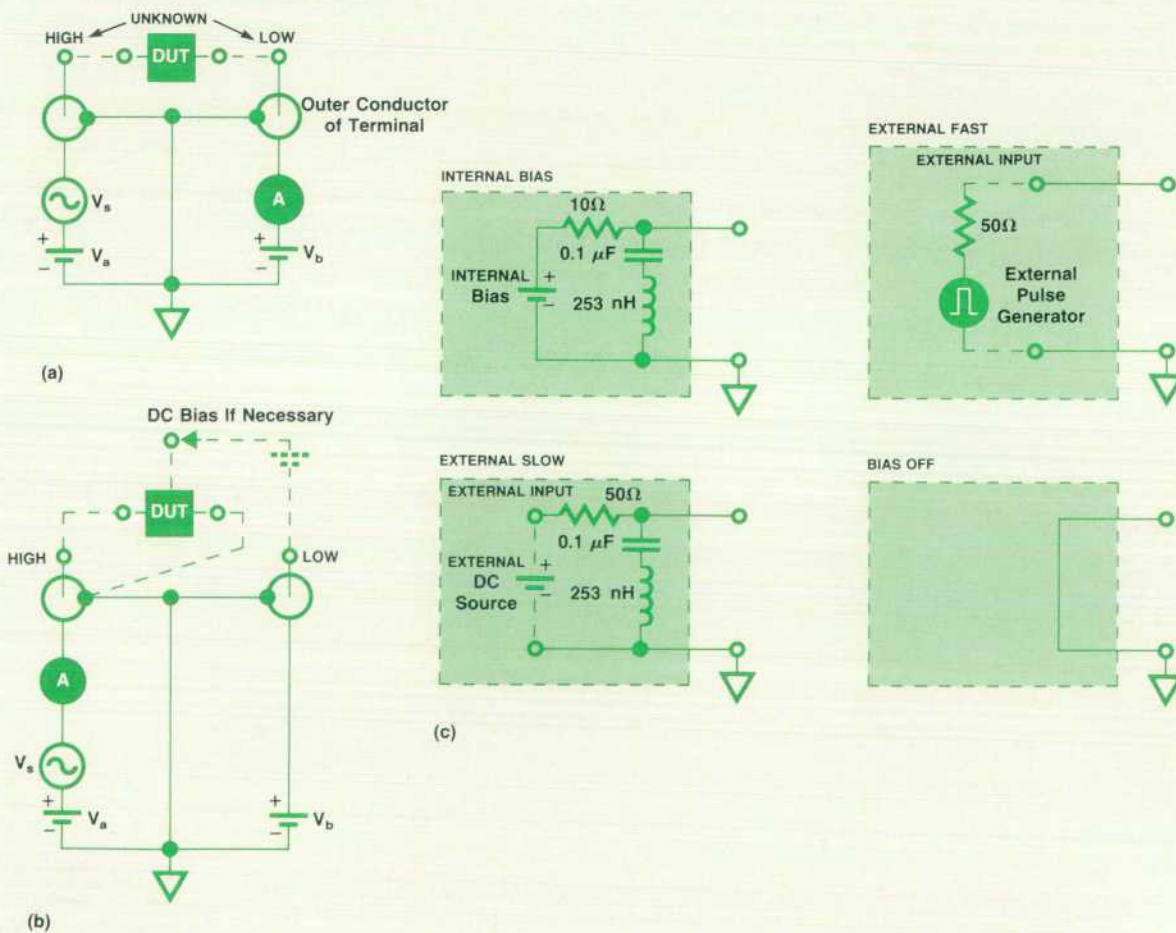


Fig. 6. Basic DUT connection modes, floating (a) and grounded (b), and four dc bias source configurations (c).




Fig. 7. External fast pulse at the end of a 1-m-long test cable. The sine wave superimposed on the pulse is the 4280A's 1-MHz test signal (signal level is 30 mV rms).

bles. Thus, if the output impedance of the pulse generator is 50Ω , the pulses are distortion-free. The two-terminal-pair configuration also has the same advantage as the four-terminal-pair configuration; that is, it eliminates measurement errors caused by mutual inductance between **HIGH** and **LOW** measurement cables since the outer conductor of each cable provides a return path for the test signal current. No magnetic fields are generated around the test cable because the magnetic fields produced by the currents through the inner and outer conductors cancel each other. The same current that flows through the center conductor also flows through the outer conductor, but in the opposite direction. This is true even if the outer shields of the signal source and the current meter are connected together as shown in Fig. 4a. The mutual inductance current I'_x that flows as

shown in Fig. 4b is very small. This is known as the balun effect of a coaxial cable at high frequencies.

The transducer (Fig. 5) consists of a signal source, which applies a complex voltage to the DUT, and an I-to-V converter, which converts the current flowing through the DUT into another complex voltage. The signal source is essentially a 1-MHz oscillator coupled to the front-panel terminals through a 20:1 step-down transformer. The output impedance is nearly zero. The I-to-V converter is basically a high-gain (typically 140 dB at 1 MHz), inverting feedback amplifier with narrow bandwidth characteristics obtained using a modulation technique.¹ The range resistor R_r in the feedback loop of the I-to-V converter converts the input current I_x to a voltage V_o , because the gain of the amplifier is almost infinite and the current input port is at virtual ground. Hence, the input impedance of the I-to-V converter is almost zero. The relationship among V_s , I_x , and Y_x is

$$I_x = Y_x V_s = -V_o/R_r$$

Therefore, the complex admittance Y_x of the DUT can be expressed as:

$$Y_x = -(1/R_r)(V_o/V_s)$$

Since the 4280A measures capacitance and conductance, Y_x is more properly expressed in rectangular format as:

$$Y_x = j\omega C + G$$

Thus,

$$C = \frac{\text{Im } Y_x}{j\omega} = -\frac{I}{j\omega R_r} \text{Im} \left(\frac{V_o}{V_s} \right)$$

$$G = \text{Re } Y_x = -\frac{1}{R_r} \text{Re} \left(\frac{V_o}{V_s} \right)$$

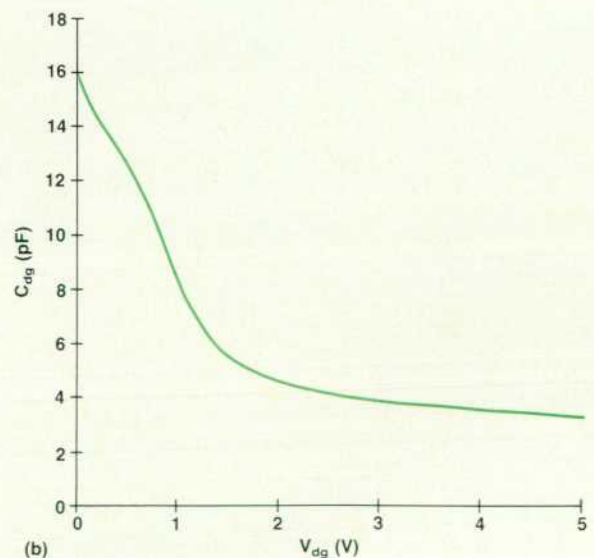
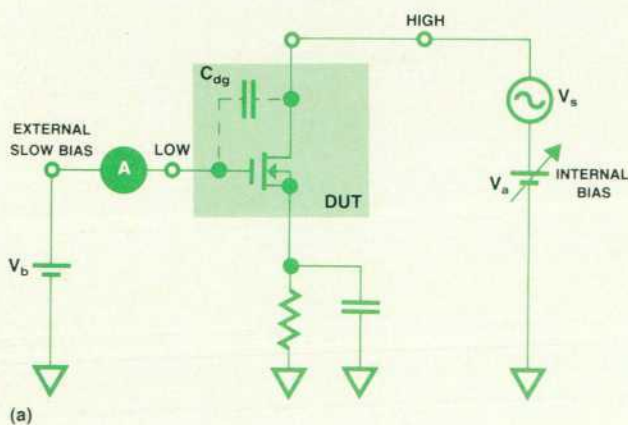


Fig. 8. Test configuration for a MOSFET drain-to-gate capacitance measurement (a) and the plot of the measurement results (b).

The vector ratio detector consists of a process amplifier and an ADC that convert the C and G counts directly; no calculations are required.

Connection Modes

The connection mode switching circuit acts as an interface between the DUT, the transducer, and the dc or pulse bias source. The 4280A provides fourteen connection configurations of these elements. These connection configurations are divided into two groups, floating and grounded, as shown in Fig. 6a and 6b. All other connection modes are variations on the two basic connection modes. That is, the signal source, I-to-V converter, and two bias sources, V_a and V_b , are reconfigured for each variation. Four bias source configurations are shown in Fig. 6c—internal bias, external slow, external fast, and bias off.

The internal dc bias source is essentially a stable power amplifier driven by a 12-bit digital-to-analog converter (DAC). It functions as a staircase bias generator in C-V measurement mode and as a pulse bias generator in C-t measurement mode. Maximum output is $\pm 100V$. Current output is limited to $\pm 5\text{ mA}$. Switching from voltage mode to current-limited mode or vice versa is well controlled. The output is always kept from overshooting to ensure accurate bias voltage at any time. This internal dc bias source has a special safety feature—when the front-panel **INTERNAL BIAS** limit switch is set to the **$\pm 42V\text{ MAX}$** position, the bias source is automatically shut down if the output voltage exceeds 50V for some reason.

External slow bias is fed through a 50 Ω resistor and then applied across a series LC circuit, which is in series with the 4280A's test signal source. The LC circuit resonates at 1 MHz to pass the 1-MHz test signal while filtering any noise that may be generated by the external dc source, usually a pulse generator. The impedance of this LC circuit is unaffected by the output impedance of the pulse generator, and is low enough to ensure measurement accuracy. The input resis-

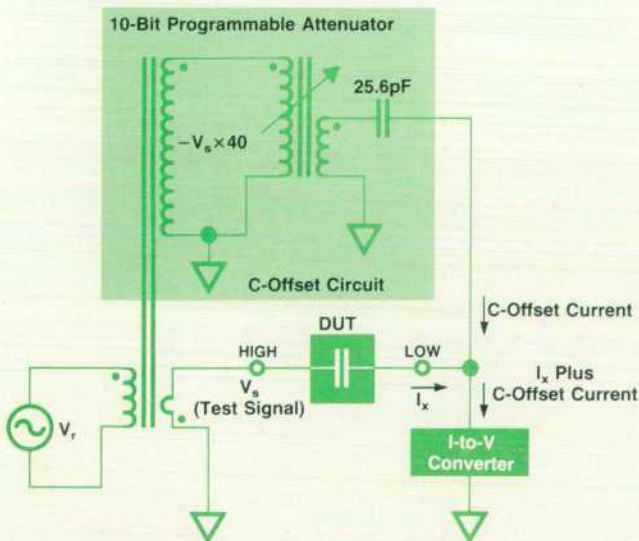


Fig. 9. Block diagram of the C-HIGH-RESOLN function. The C-offset current cancels a portion of the DUT current I_x , allowing the I-to-V converter to measure the difference current on a more sensitive range.

tor, 0.1- μF capacitor, and 50 Ω output impedance of the pulse generator limit the rise time to 10 μs .

The external fast bias input is intended to pass narrow pulses from an external pulse generator that has a 50 Ω output impedance and is used for fast C-t measurements. A bias pulse at the end of a one-meter-long test cable connected to a 100-pF DUT is shown in Fig. 7. The sine wave superimposed on the pulse is the 4280A's 1-MHz test signal.

The bias off mode eliminates residual offset voltage at the **UNKNOWN** terminals and allows higher currents than the $\pm 5\text{ mA}$ allowed by the internal bias current limiter. The internal dc bias source, the external slow bias source, and bias off mode can be selected for dc bias sources V_a and V_b shown in Fig. 6. The signal source and the I-to-V converter are isolated from measurement circuit common to allow these different measurement modes. An example of a floating-mode measurement of drain-to-gate capacitance versus drain voltage on a MOSFET and the resulting plot are shown in Fig. 8.

C-HIGH RESOLN

The 4280A has an optional high-resolution capacitance measurement function that increases display resolution by one digit on the 100-pF and 1-nF ranges, making 5½-digit resolution possible without a reduction in measurement speed. Fig. 9 shows the block diagram of the C-HIGH RESOLN function. The C-offset current, which is proportional to a user-programmed offset capacitance value, is generated in the C-offset circuit from a 1-MHz sine wave whose amplitude is forty times that of the signal applied to the DUT and whose phase is 180° ahead of the test signal. A 1-MHz programmable 10-bit attenuator sets the signal level and a 25.6-pF standard capacitor converts the voltage to a

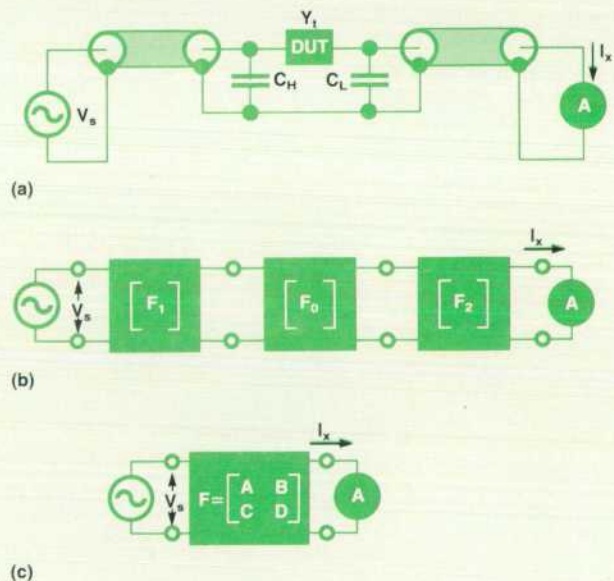


Fig. 10. (a) Measurement of the DUT admittance Y_1 is influenced by the test cable parameters and the stray capacitances C_H and C_L . (b) Mathematical expression of (a) using a chain-matrix. (c) The matrix of (b) expressed as one matrix, where element B is a function of Y_1 , cable parameters, C_H , and C_L .

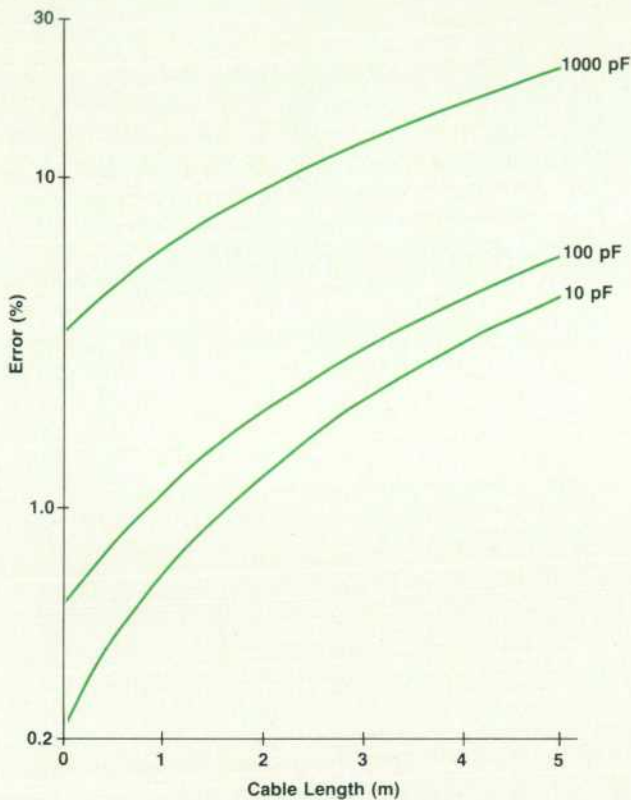


Fig. 11. Error caused by external test cables. The 0-m reference is the front-panel UNKNOWN terminals.

current. The attenuator consists of three step-down transformers. Each transformer has step-down ratios of 2:1, 4:1, and 8:1 to provide 9 bits of attenuation and the 10th bit is obtained by bypassing the transformers. Linearity deviation is less than one LSB (least-significant bit) and phase shift is typically less than one degree for any attenuation setting. The LSB is equivalent to 1 pF because the CV product is

$$25.6 \text{ pF} \times (-V_s) \times 40 \times (1/1024) = (-1 \times V_s) \text{ pF}$$




Fig. 12. The error caused by stray capacitances C_H and C_L . This error remains after internal error correction but can be corrected by an external HP-IB controller if desired.

where 1/1024 is the resolution of the attenuator. The maximum offset capacitance value that can be programmed is 1023 pF.

To maintain continuity between normal mode measurements and high-resolution measurements, the high-resolution capacitance measurement routine contains the following steps:

1. Turn off the error correction function, and then measure the total capacitance (C_o) and conductance (G_o) of the DUT and any test fixturing being used. Store the results in the 4280A's internal memory.
2. Set the offset capacitance to the values measured in step 1, and then downrange and measure the difference (C_r) between the actual capacitance (C_o) of the DUT and the offset capacitance output from the C-offset circuit. Since this small difference is measured on a lower range, we effectively get one more digit of resolution. For example, if the measurement in step 1 is made on the 1-nF range (maximum sensitivity is 0.1 pF), the difference value will be measured on the 100-pF range (maximum sensitivity is 0.01 pF).
3. Calculate the actual offset value C_{off} as

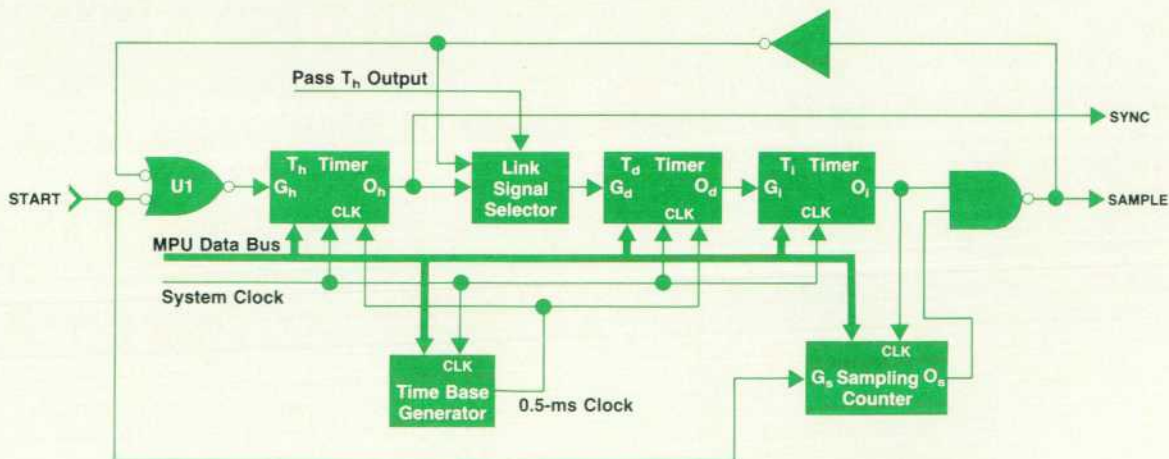


Fig. 13. A simplified diagram of the timing control circuitry used for C-t measurements.

$$C_{\text{off}} = C_o - C_r$$

C_{off} is the offset value actually output from the C-offset circuit, and will differ slightly from C_o . The C-offset circuit tries to output an offset capacitance value equal to C_o , but can only come close because of the resolution limitations (1 pF) of the 10-bit attenuator in the C-offset circuit and the inaccuracies inherent in the circuit components.

4. Make another measurement to obtain new difference data (C'_r), perform error correction, and then calculate true capacitance C .

$$C = C_{\text{off}} + C'_r = C_o - C_r + C'_r$$

As long as the **C-HIGH-RESOLN** function is turned on, step 4 will be repeated each time the 4280A is triggered (e.g., each step of a bias sweep). If the test sample's capacitance and conductance values do not change from those obtained in step 1, then C'_r in step 4 will be equal to C_r , and the values displayed in step 4 will be equal to the step 1 value, C_o . The only difference will be one additional digit of display resolution. This extra digit, which represents 0.001 pF on the 10-pF range, allows observation of minute changes in capacitance and conductance as functions of voltage or time.

The **C-HIGH RESOLN** mode is especially useful for C-t measurements such as those used in deep-level transient spectroscopy² applications that require detection of very small capacitance changes caused by the combination-recombination process of minority carriers in semiconductor devices.

Error Correction and Cable Compensation

As shown in Fig. 10, one pair of measurement cables connects the device under test to the signal source and current meter. The stray capacitances, C_H and C_L , that exist between the DUT and the outer conductors of the cables are sources of error when measuring the true admittance Y_t of the DUT. The cables also have electrical characteristics that cause a measurement error. Hence, the measured admittance Y_m is different from the true admittance Y_t . Y_m is expressed as:

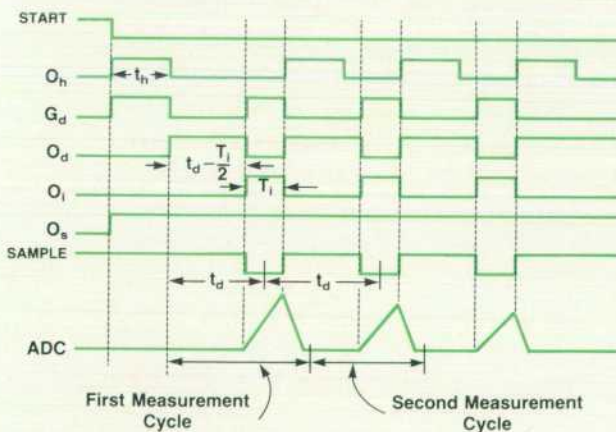


Fig. 14. Timing diagram for burst mode C-t measurements.

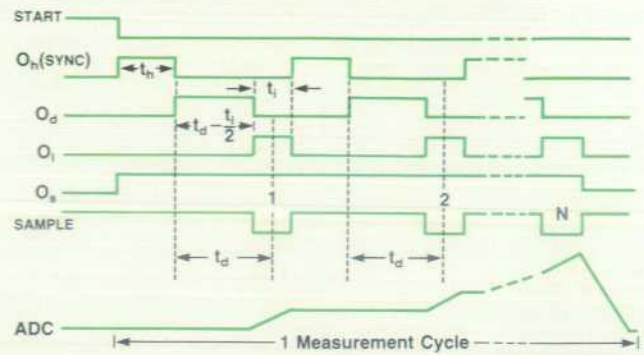


Fig. 15. Timing diagram for sampling mode C-t measurements.

$$Y_m = I_x/V_s = j\omega \times C_m + G_m$$

where C_m is measured capacitance, G_m is measured conductance, and ω is the angular frequency at 1 MHz.

Fig. 11 shows the percent error as a function of cable length when the DUT is a 10-pF, 100-pF, or 1000-pF capacitor. This error can be corrected by the 4280A's error correction function. The equivalent circuit, including test cables, can be expressed using a mathematical expression as shown in Fig. 10b. The chain matrix components **F1** and **F2** are determined by the cable length L . Because the parameters of the test cables are strictly controlled, this makes the conversion to the matrix component easy. However, a measurement using test cables of an arbitrary length is possible because the cable length calibration function of the 4280A can measure the capacitance between the inner and outer conductors of the test cable and convert the measured value into the cable length automatically.

Y_t and stray capacitances C_H and C_L are converted into the **F0** matrix. The three matrices can be expressed as a single matrix **F**, as shown in Fig. 10c, where the relation between the parameters can be expressed as:

$$Y_m = I_x/V_s = 1/B$$

where B is an element of matrix **F** and is a function of Y_t , the cable parameters, and C_H and C_L . The above equation can be rewritten as:

$$Y_t = f(Y_m, \text{cable parameters}, C_H, C_L)$$

The internal correction algorithm calculates Y_t using the above equation. However, it ignores C_H and C_L in the process, that is, assumes them to be zero. Therefore, if C_H and C_L exist, an uncorrected error remains. The residual error caused by C_H and C_L is shown in Fig. 12. This error is considered negligible for most applications. For measurements that require higher precision, this C_H/C_L -caused error can be corrected automatically with the help of an external computer.

C-t Measurement

The 4280A measures transient capacitance (C-t) and conductance (G-t) after a bias pulse has been applied to the

device under test. Timing control for the bias pulse, delay time, and I-to-V converter operation is the most important function of the 4280A. Fig. 13 is a simplified diagram of the timing control circuitry for C-t measurements. The T_h timer determines hold time (or more precisely, the width of the bias pulse). During sampling mode measurements, the output (SYNC) from this timer is used as a synchronization signal for the external pulse generator. The T_d timer determines delay time, which is defined as the time from the end of the bias pulse to the center of the first measurement period. From this definition, the data programmed into the T_d timer is $(t_d - T_i/2)$, where T_i is the integrator charge time determined by the T_i timer and t_d is the selected delay time. Both the T_h and the T_d timer interval parameters are programmed through the 4280A's microprocessor data bus. Table I lists time setting range, resolution, and accuracy.

The time base generator outputs a 0.5-ms pulse. This signal is sent to the T_h and T_d timers and is used as the time base for counting long time intervals (65.5 ms to 32 s).

The 4280A has two methods for C-t measurements. If the delay time is long (at least 10 ms), then the 4280A uses burst mode. If the delay time is short, the 4280A uses sampling mode. During burst mode measurements, only one bias pulse, with a width of t_h , is applied to the DUT for each measurement cycle, and measurements are made every delay time period. During the delay time periods, the 4280A's microprocessor controls the measurement, measurement setup, A-to-D conversion, data storage, display refresh, data output, and so on. The integrator charge time T_i is set to 1 ms or 10 ms by the microprocessor depending on the selected measurement speed. Fig. 14 shows the timing diagram for burst mode measurements.

During sampling mode measurements, the integrator

Table I

Hold Time/Delay Time

Range	Resolution	Accuracy
10 μ s to 65 ms	10 μ s	$\pm(0.02\% + 100 \text{ ns})$
65.5 ms to 1 s	0.5 ms	
1 s to 10 s	1 ms	$\pm(0.02\% + 0.5 \text{ ms})$
10 s to 32 s	10 ms	

charge time T_i is divided into small chunks $t_i = t_d/5$ and the T_i timer is reset to t_i . A sequence of t_h , t_d , and t_i periods is repeated until the sum becomes close enough to the overall integrator charge time T_i . If the sum does not exactly equal T_i , the fraction will be calculated by the microprocessor. The timing diagram for sampling measurements is shown in Fig. 15. The sampling counter counts the number of integrator charge periods. In burst mode measurements, this number is set to 1. In sampling mode measurements, this number is determined by the equation, $N = \text{integral part of } (T_i/t_i)$ for $N \geq 1$.

At the beginning of a C-t measurement, all timers except the time base generator are set to the appropriate conditions, and all outputs are forced low. When the START signal goes low, the T_h timer and sampling counter are activated. The sampling counter output O_s goes high. The T_h timer output O_h goes high initially, and then goes low after hold

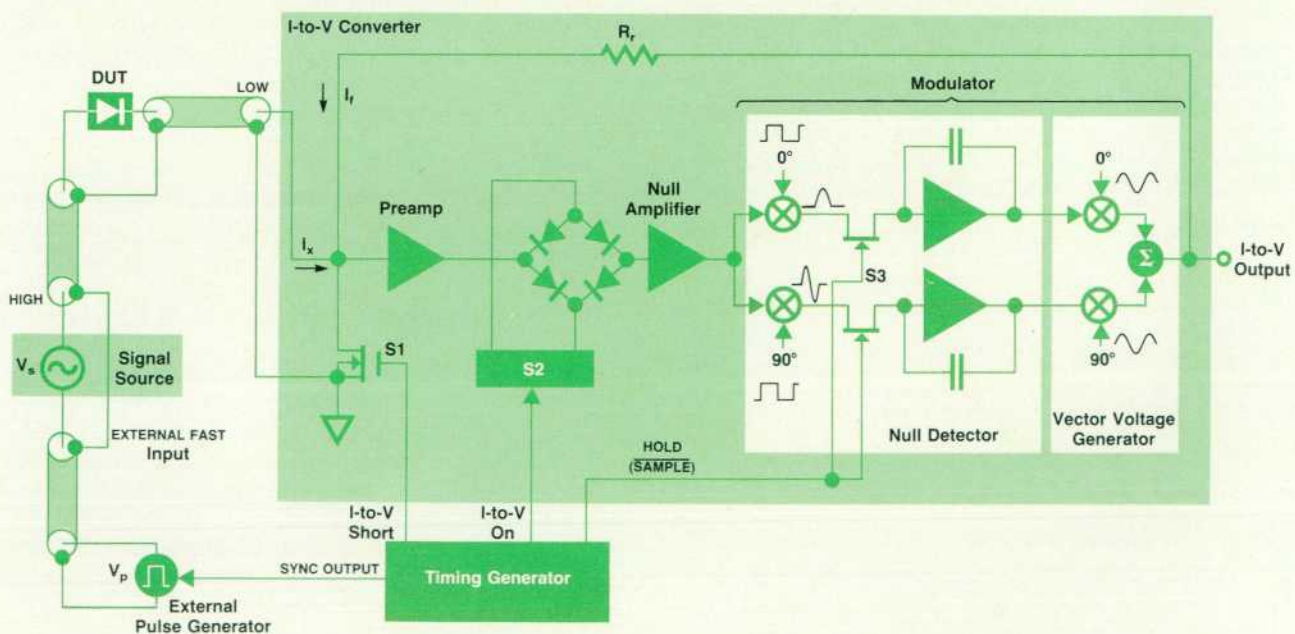


Fig. 16. Simplified block diagram of the I-to-V converter. The hardware configuration used during C-t sampling mode measurements is shown.

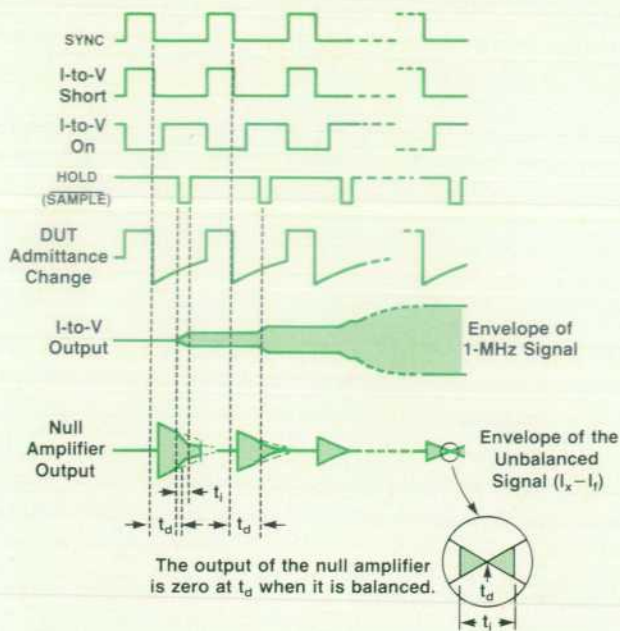


Fig. 17. I-to-V converter timing diagram for C-t sampling mode measurements.

time interval t_h . This negative-going signal at the T_h timer output activates the T_d timer. The T_d timer output O_d remains high until the $t_d - T_i/2$ (burst mode) or $t_d - t_i/2$ (sampling mode) interval has elapsed. When the O_d goes low, the T_i timer is activated. After the T_i timer is activated, the T_i timer output O_i remains high until the time interval T_i (burst mode) or t_i (sampling mode) has elapsed. O_i is gated by the sampling counter output O_s . This gated output (SAMPLE) is sent to the measurement section to enable A-to-D conversion. Also, the SAMPLE signal is sent back to the T_h timer through U1 and to the T_d timer through the T_d link-signal selector. The T_d link-signal selector selects either O_h or SAMPLE as the T_d timer input G_d . The output from the T_d link-signal selector then activates the T_d timer to restart the timer linking operation. For sampling mode measurements, SAMPLE activates the T_h timer through U1. The T_d link-signal selector selects O_h as input G_d and a sequence (t_h , t_d , t_i) continues until the sum of t_i periods equals T_i . (For burst mode measurements, just one bias pulse is necessary for the measurement.) After finishing the first measurement, the T_d link-signal selector passes the SAMPLE signal to the G_d input, thus linking the T_d and T_i timers to each other. This linking operation continues until the specific number of sampling mode measurements is completed.




Fig. 18. (a) A transient (time constant τ_1). (b) Amplifier (time constant τ_2). Shaded area (c) is the error that results when transient (a) is passed through amplifier (b).

I-to-V Converter Operation

Fig. 16 shows a block diagram of the I-to-V converter, detailing C-t sampling mode control hardware. For fast C-t measurements (delay time t_d less than 10 ms), the bias pulse must be provided by an external pulse generator. When the measurement begins, the 4280A outputs a synchronizing pulse (SYNC OUTPUT) that can be used to control the width of bias pulses output from an external pulse generator. The width of the pulse is set by the user-specified hold time t_h . If the pulse generator has an external pulse width control function (i.e., the pulse generator outputs a pulse whose width is exactly the same as the width of a signal applied to its trigger input), then the only parameter that must be set on the pulse generator is the pulse output level. For pulse generators that do not have this capability, the pulse control parameters must be set so that the bias pulse and the 4280A's SYNC OUTPUT pulse end at exactly the same time to maintain delay time (t_d) accuracy. The duty factor of the synchronizing pulse, which is defined as t_h/t_d , can vary from 0.00001 to 500.

The 1-MHz test signal is superimposed on the external pulse signal, and both are applied to the DUT. The pulse current that flows through the DUT is routed to circuit common by FET switch S1 at the input to the preamplifier. This preamp is a low-noise, high-speed amplifier, and is capable of recovering from any saturation condition within a few microseconds. Diode switch S2 between the preamp and the null amplifier is turned off while the bias pulse is applied to the DUT to keep the large bias pulse from entering the high-gain null amplifier. This is necessary, because if the null amplifier becomes saturated, it takes a long time to recover, resulting in measurement error. Switch S2 is turned on just after the bias pulse is removed. The incoming signal is then amplified by about 80 dB and fed to the null detector.

The null detector operates only while FET sampling switch S3 is on. This switch is turned on for a period that is approximately 20% of the specified delay time t_d . For example, if t_d is 10 μ s, this switch will turn on for 2 μ s approximately 9 μ s after the t_d period begins.




Fig. 19. Example of temperature-biased C-t measurement on a GaAs Schottky diode. The shift of the C-t curve with temperature (indicated by arrows) can be clearly seen.

The unbalanced signal output from the null amplifier is separated into its orthogonal components by the 0° and 90° phase detectors. The dc outputs from the integrators contain phase information and are amplified by about 100 dB. These dc outputs remain constant while the sampling switch is off and are fed to the vector generator to modulate 0° and 90° 1-MHz sine waves. The modulated 0° and 90° signals are summed and then fed back to the I-to-V converter's input through range resistor R_r . Since the null loop closes only when the sampling switch is on, the output from the null amplifier can balance in response to the transient admittance of the DUT only when this switch is on, which occurs at the end of the t_d period (Fig. 17). The I-to-V output signal, which is sent to the vector ratio detector, settles at a steady 1-MHz sine wave after the loop is balanced. The voltage ratio detector begins operation only after the I-to-V converter is balanced.

The time required for the loop to balance is about 50 μ s. One way to express the ability to catch the transient response of the device under test is the response time of the measurement instrument. For example, if a transient with a time constant τ_1 (Fig. 18a) is passed through an amplifier with a response time of τ_2 (Fig. 18b), the result will be as shown in Fig. 18c. The shaded area represents an error. The error related to changes of the DUT's admittance is expressed as

$$\text{Error} = \frac{1}{1 - \tau_2/\tau_1} \left[\exp\left(\frac{\tau_2/\tau_1 - 1}{\tau_2} \times t_d\right) - \frac{\tau_2}{\tau_1} \right]$$

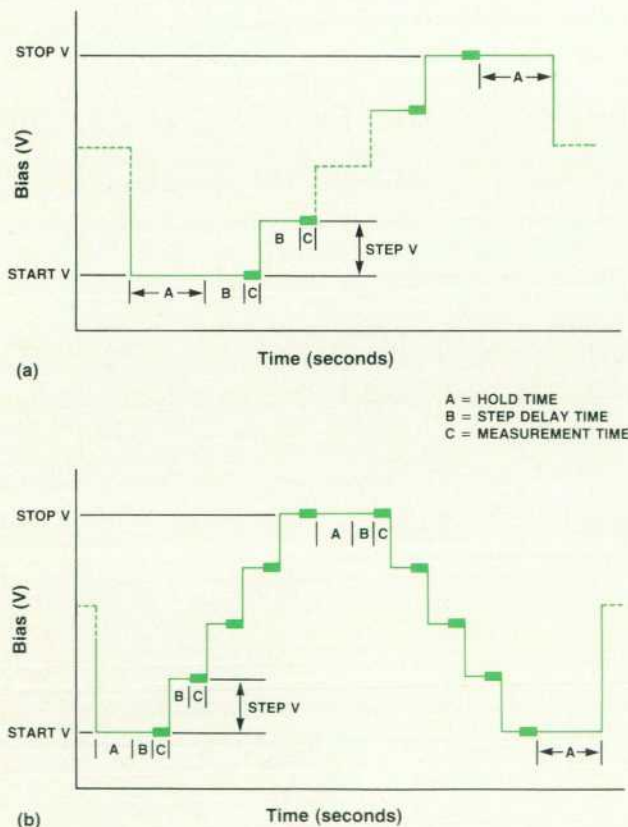


Fig. 20. Single staircase (a) and double staircase (b) dc bias sweep parameters.

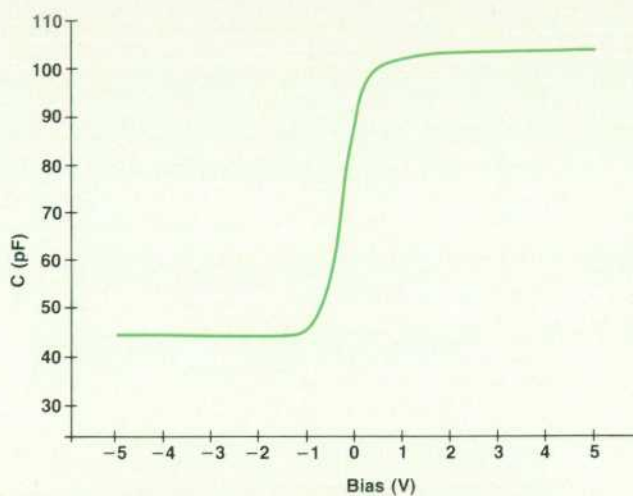


Fig. 21. Reproduction of an X-Y recorder plot of a swept C-V measurement. The DUT is an MOS capacitor. Measurement parameters were **START V** = -5V, **STOP V** = 5V, **STEP V** = 0.1V, **HOLD TIME** = 60 s, and **STEP DELAY TIME** = 1 s.

If we assume that $\tau_2 \ll \tau_1 < t_d$, a realistic assumption, then the above equation can be simplified to

$$\text{Error} = -\tau_2/\tau_1$$

The above equation shows that the error is directly proportional to the response time of the amplifier and inversely proportional to the response time of the DUT. The response time of a sampling circuit, such as the I-to-V converter discussed here, is determined by the response time of the null detector, which in the case of the 4280A is about 1 μ s.

For example, assume we're trying detect a 1-pF change in capacitance with a time constant of 100 μ s out of a 100-pF bulk capacitance. The maximum error is 0.01 pF if $t_d = 10 \mu$ s, the minimum time delay of the 4280A. This error reduces to about 0.001 pF, the minimum resolution on the 100-pF range with the **C-HIGH RESOLN** option, if $t_d = 200 \mu$ s. The 1- μ s response time, in addition to accuracy and stability, would have been impossible to attain without this sampling and modulation technique.

Fig. 19 shows a sample plot of C-t measurements on a GaAs Schottky diode, where t_d is varied from 10 μ s to 100 s, and temperature varies from -70°C to +65°C (203K to 378K).

C-V Sweep Mode Operation

In C-V mode, the 4280A measures capacitance C and conductance G as functions of swept dc bias. Output from the internal dc bias source can be swept in either a single-staircase or double-staircase manner. One of the advantages of C-V measurements with the 4280A is fast measurement speed. For example, a 30-point swept measurement can be done in less than one second (C-only, fast measurement speed, block data transfer mode). This is possible because the internal dc bias source is synchronized with the measurement section to eliminate the time required to set the dc voltage at each measurement point, as would be the case in a multi-instrument C-V measurement system. To

perform a swept bias measurement in C-V mode, five sweep parameters—**START V**, **STOP V**, **STEP V**, **HOLD TIME**, and **STEP DELAY TIME**—must be set (Fig. 20). Table II lists the parameter setting range and resolution. Fig. 21 shows a reproduction of a C-V plot that was made using the X-Y recorder output.

Software

Fig. 22 diagrams the major modules of the 4280A's software. There are three categories: Foreground modules, background modules, and utilities.

There are five foreground modules. The diagnostic module executes the self-test of the 4280A just after power-on. The background executive module controls all background tasks, scheduling them in accordance with the control settings and executing them in proper order. The nonmaskable interrupt (NMI) controller module is closely related to measurement timing. Timing control is the most important function of the 4280A. All timeout conditions from the timers are sent to the 4280A's microprocessor over an NMI hardware line. The processor detects the NMI and determines which timer is requesting service. Once an NMI occurs, this module executes both hardware and software settings immediately. To ensure quick responses to NMIs, these modules are kept as short as possible. The IRQ (interrupt request) controller module functions as an interface between the 4280A and the operator. Measurement condition changes entered from the 4280A's front panel or via the HP-IB are detected by this module. This module also sets up all software flags for the background tasks.

When the program starts, the background executive calls the swap controller, background task scheduler, and background task controller. In the swap controller, software flags for the measurement are swapped if a special measurement is required (e.g., zero/open, cable, and L-R measurements). The task scheduler sets the task bit of the BGT TASK

Table II

Hold Time/Delay Time		Start/Stop Voltage	
Range	Resolution	Range (V)	Resolution
3 ms to 65 ms	1 ms	±(0 to 1.999)	1 mV
0.07 s to 99.99 s	10 ms	±(2 to 19.99)	10 mV
100 s to 650 s	100 ms	±(20 to 100)	100 mV

FLAG according to the software measurement flags. BGT TASK FLAG has eight bits and each bit is assigned to one of the eight background tasks: bit0, self-test; bit1, hardware setup; bit2, measurement control; bit3, correction CAL; bit4, deviation CAL; bit5, display/recorder output; bit6, HP-IB output; bit7, postprocess. When measurement start is requested, the task scheduler sets bit1 and bit2. If the error correction setting is on, then this scheduler also sets bit3, and so on, depending on the existing measurement conditions. The background task controller executes the background tasks using the BGT TASK FLAG bits. After executing a background task, this controller clears the corresponding task bits, and continues this sequence until all bits are cleared. After all background tasks are executed, program control is returned to the background executive and a new sequence starts.

There are eight background modules. The self-test controller module executes the self-test only when the self-test function is on. The hardware setup controller sets up both digital and analog hardware conditions for the measurement. The measurement controller executes the actual measurement (C, C-V, C-t, C-offset, special measurement). The

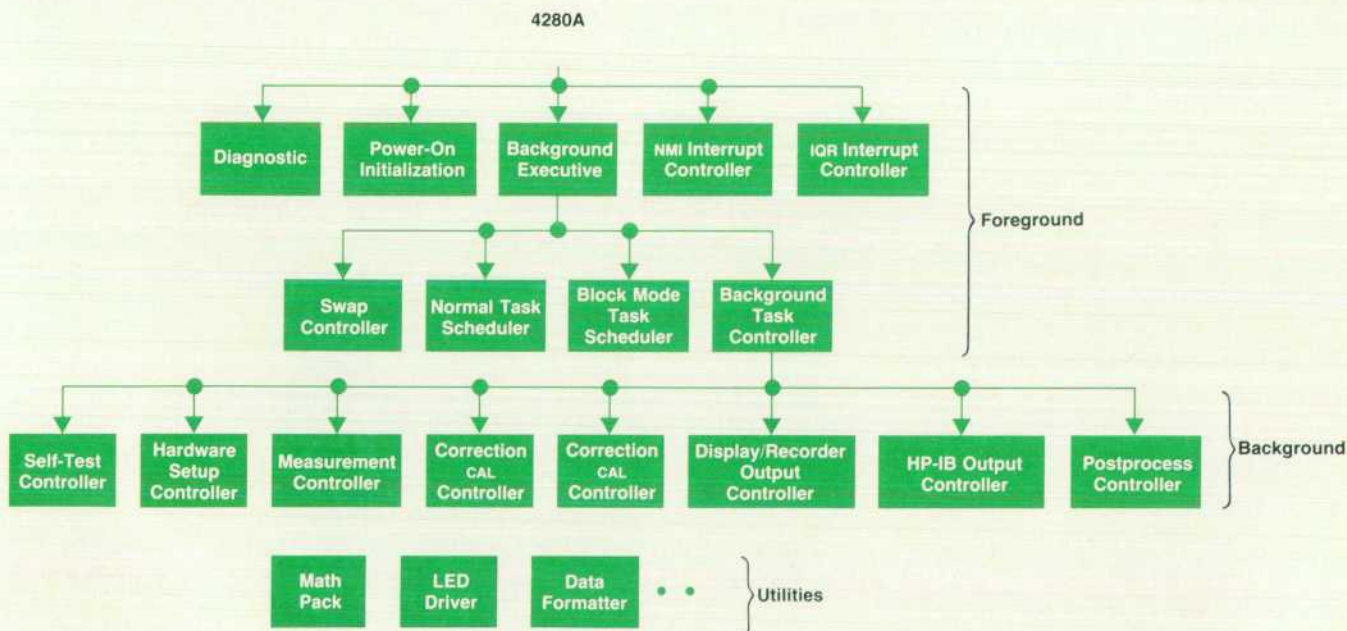


Fig. 22. Major software modules of the 4280A.

correction CAL controller executes the error correction routine if the error correction function is on. The deviation CAL controller calculates deviation or percentage if one of the **MATH** functions is on. The display/recorder output controller updates the measurement data for display and X-Y recorder output. If the 4280A is in block mode or binary data output mode, this task is not scheduled. The HP-IB output controller stores HP-IB output data in the output buffer. The postprocessor clears all hardware/software settings for the next measurement.

The utilities software has many modules that are used with both foreground and background modules. Examples are a math pack, LED driver, data formatter, and so on.

Acknowledgments

Hitoshi Noguchi was the project leader. Hisao Yoshino,

Yasunori Hiratsuka, Hiroshi Kanamori, Koh Murata, and Yoshiaki Okuyama developed the logic and software. Hideo Ohgawara, Hideo Akama, and Norio Sone designed the analog section. Akira Uchiyama and Akihiko Goto handled mechanical design and the test fixtures. Tsuneji Nakayasu did the industrial design. Special thanks go to Haruo Ito, our section manager, and to Takuo Banno, who gave us much useful advise, and to many other people who made special contributions to the project.

References

1. K. Maeda, "An Automatic, Precision 1-MHz Digital LCR Meter," *Hewlett-Packard Journal*, Vol. 25, no. 7, March 1974.
2. L. Forbes and U. Kaempf, "Capacitance and Conductance Deep-Level Transient Spectroscopy Using HP-IB Instruments and a Desktop Computer," *Hewlett-Packard Journal*, Vol. 30, no. 4, April 1979.

Authors

June 1984

3 Parametric Test System

Yoh Narimatsu



Yoh Narimatsu joined Yokogawa-Hewlett-Packard in 1971 as a development engineer. After several years, Yoh transferred to HP's Santa Clara Division where he was involved in the 5342A Microwave Frequency Counter project. After returning to Japan, he

helped develop the 4275A LCR Meter and the 4192A LF Impedance Meter before he was given project manager responsibilities for the 4062A Semiconductor Parametric Test System. Recently, Yoh was placed in charge of semiconductor test instrument development. Yoh holds a BSEE degree from Kyoto University and an MSEE from Stanford University. He and his wife have two children, a son and a daughter. Yoh enjoys woodworking in his spare time.

Keiki Kanafuji



Keiki Kanafuji earned his BS degree in electrical engineering from the Tokyo Institute of Technology in 1972 and then began work at Yokogawa-Hewlett-Packard. Before joining the 4062A project, Keiki designed several instruments, including the 4262A LCR Meter and the 4140B pA Meter/DC Voltage Source. He designed the analog section of the

4062A system and especially contributed to the implementation of the system's low-current measurement capability through the 4085A Switching Matrix. During his leisure hours, you'll probably find Keiki playing contract bridge.

9 Parametric Test Software

Takuo Banno



Takuo Banno received his BSEE degree from the University of Electro-Communications in 1973. He then joined YHP's Tokyo Sales office as a field engineer. In 1977, Takuo came to the YHP Instrument Division as a development engineer. Before

starting the 4062A project, he was the project leader for the 4061A Semiconductor/Component Test System. Takuo's hobbies include sailing, board sailing, skiing, running, and playing bridge.

12 C-V/C-t Meter

Tomoyuki Akiyama



Tomoyuki Akiyama joined Yokogawa-Hewlett-Packard in 1973 after receiving a BSEE degree from Sophia University, Tokyo. After a few years as a development engineer, he transferred to HP's Loveland Instrument Division in 1977 and was involved in the development of the 3582A Spectrum Analyzer. During his stay in the U.S.A., he received an MSEE degree from Colorado State University. He designed the I-to-V converter for the 4280A C Meter/C-V Plotter. He is married, has two daughters, and enjoys classical music, woodworking, and skiing.

Kenzo Ishiguro



Kenzo Ishiguro holds BSEE and MSEE degrees from Shinshu University, Japan, and an MSEE degree from Stanford University. He joined Yokogawa-Hewlett-Packard in 1974. After several years as a development engineer, he transferred to HP's Santa Clara

Division in 1979 and was involved in the development of the 5180A Waveform Recorder. He headed the digital/software design team of the 4280A C Meter/C-V Plotter. He enjoys playing tennis and golf. He is married and has two daughters.

26 Software Analyzer

Joseph A. Hawk



A native of Victor, Colorado, Joe Hawk received a BS degree in electrical engineering and computer science from the University of Colorado at Colorado Springs in 1980. Joining HP in 1979, he did performance verification for the 64600 Timing Analyzer and software for the 64310A Analyzer before his current work on high-level software analysis. Joe is married, has four children, and lives in Colorado Springs, Colorado. When not spending time with his family and working on a new home, Joe's leisure

activities include doing stained glass with his wife, playing darts and pool, camping, playing guitar and piano, and panning for gold.

Andrew J. Blasciak



Andy Blasciak joined HP in 1979 as a production engineer after receiving a BSEE degree from Clemson University. Now an R&D engineer at HP's Logic Systems Division, he is single and enjoys skiing, bicycling, and camping in the Colorado mountains.

Born at Travis Air Force Base in California, he lives in Colorado Springs, Colorado.

Gail E. Hamilton



A project manager for software analysis tools at HP's Logic Systems Division, Gail Hamilton managed the development of the 64310A Analyzer and earlier contributed to the development of the 8085 and 6502 personality modules for the 1611A Logic Analyzer. She is the

author of two papers on software performance analysis and has conducted parts of HP's Digital Seminars. Born in Fort Sill, Oklahoma, Gail joined HP in 1971 and is a graduate of the University of Colorado (BSEE/CS 1975). Active in many civic affairs, she is on the board of directors for the Colorado Opera Festival, a member of the Educational Task Force for the local Chamber of Commerce, a delegate for the Republican Party, and a part of the Citizen's Goals Leadership program. Gail lives in Colorado Springs, Colorado, is married, and has two sons. When she has time to relax, she enjoys gardening, golf, skiing, politics, aerobic exercise, the performing arts, and studying Japanese.

Brett K. Carver



A native of Monterey, California, Brett Carver moved to Colorado Springs, Colorado, when he joined HP in 1981. He currently is working on high-level software analysis. Brett holds BS (1979) and MS (1981) degrees in computer science

from California Polytechnic State University. He is married and interested in photography, firearms, skiing, and dancing.

33 Gated Counter Module

Thomas K. Bohley



Tom Bohley is an R&D engineer at HP's Colorado Springs Division. He's been with HP for 18 years and has contributed to the design of many oscilloscope and display products. Two patents on oscilloscope circuits have resulted from his work. A native of St. Louis,

Missouri, Tom received his BSEE degree in 1961 and his MSEE in 1966 from the University of Missouri. Between degrees, he served as a lieutenant in the U.S. Army for two years. His major interest is horses—he owns three, participates in fox hunts, and competes in dressage and combined training events. He's married and lives in Colorado Springs.

Donald J. Smith



Don Smith received his BSEE degree in 1975 and his MSEE in 1977 from the University of Texas at Austin. He joined HP's oscilloscope laboratory in 1977, contributed to the design of the 1950A Two-Channel Module, and served as project manager

for the 1980A/B Oscilloscope Measurement System and the 1965A Gated Universal Counter. He's now managing a new oscilloscope development project. Don was born in Wichita Falls, Texas and lives in Colorado Springs, Colorado. He enjoys skiing, tennis, soccer, and backpacking.

Johnnie L. Hancock



Johnnie Hancock has been with HP since 1979. He's done analog and digital hardware design for the 1950A Two-Channel Module and the 1965A Gated Universal Counter, and was a project leader for the 1965A. Born in Coral Gables, Florida, he served

four years in the U.S. Marine Corps, attaining the rank of sergeant. He received his BSE degree in electrical engineering from the University of South Florida in 1979. He's married, has two children, lives in Colorado Springs, Colorado, and enjoys skiing, spear fishing, and refinishing antiques.

An Electronic Tool for Analyzing Software Performance

Improving software performance requires measurement of program activity and duration under different conditions. This subsystem for the HP 64000 Logic Development System makes it easy to obtain such data.

by Gail E. Hamilton, Andrew J. Blasciak, Joseph A. Hawk, and Brett K. Carver

TODAY, SOFTWARE COSTS are often greater than hardware costs in developing microprocessor-based products. This growth in demand for software creates a tremendous challenge for software engineering. The challenge is threefold: to increase software development productivity significantly, to increase the efficiency of software maintenance, and to increase the reliability of the software. The problem is magnified by product designs needing high performance within tight memory constraints. One method for resolving this problem that has not been fully used until recently is software performance analysis. By providing a nonintrusive view of software as it executes in real time, performance analysis helps designers locate bottlenecks, characterize software performance, and improve software efficiency much earlier in the development cycle.

The importance of performance analysis lies not only in determining whether a system meets certain objectives; it is also valuable in understanding if and how system performance can be improved. The goal of any quantitative software measurement is information that will help in making better design decisions. Performance measurements are valuable in that they help detect anomalies and predict what may be seen in future designs. They can assist the designer in making decisions about where to spend more time improving the product. One of the primary benefits of performance analysis is that it allows the software designer to see a problem from more than one perspective.

Performance measurements are cost-effective only when they are based on acquired information, easy to evaluate, and available early enough to help. Software that does not perform adequately near the end of a project can cause major delays in getting the product to market. One technique for writing software is to generate a total solution in as high a level of abstraction as possible (i.e., with a high-level programming language such as Pascal). Then, by measuring the performance of various code modules against design specifications, detailed optimization efforts can be directed only where absolutely necessary. The characterization of software is critical to understanding the complex interrelationships and gaining insight for product enhancements.

Many software performance measurement concepts are not new; these measurement needs and benefits evolved through early experience with mainframe computers and

later, minicomputers. In the beginning, the emphasis in computer software measurements was on measuring system resource use. Many of these systems used a hardware monitoring approach with extensive probe-point libraries. They were often expensive and usually difficult to use. Other measurement tools were implemented in software; they were intrusive and not transparent to the software being measured.

Complementing the HP 1630A/D Logic Analyzer and the HP 64620A State Analyzer¹ subsystem of the HP 64000 Logic Development System,^{2,3} the new HP 64310A Software Performance Analyzer (Fig. 1) allows software performance measurements to become a major part of the basic tool set that every software engineer needs. With the 64310A, the user can measure the activity within defined areas of memory, the time spent in a code module, the time spent going between modules, and the number of transfers from one module to another. The 64310A's friendly user interface, symbolic event specification, automatic configuration capability, and data display choices provide the




Fig. 1. The HP 64310A Software Performance Analyzer is a subsystem designed to work with the emulation capability of the HP 64000 Logic Development System. It allows a software designer to study code module execution times, program activity, and other software characteristics in a nonintrusive manner.

software engineer with an effective and accurate electronic tool for generating efficient and reliable code in an economical time frame.

User Interface

For an electronic product to be easy to use, the user interface must be simple and very clear. We all are familiar with products that can make very complex measurements, but are so difficult to use that often one can't do the required task. The 64310A Analyzer is capable of performing very complex and powerful measurements. The design team wanted an easily accessible tool, so a major design goal was to create a user interface that would be as simple and easy to use as possible.

Several guidelines were followed to reach this goal. One makes the commands read as much like a spoken language as possible. The 64310A commands use meaningful words and the syntax of the commands reads like a sentence (e.g., copy display to printer). Thus, the command explains what it does, making it unnecessary to refer to a manual. Another benefit of making the commands read well is that to some extent, they become self-documenting.

Another guideline that makes the 64310A easy to use is restricting the grammar required for a decision point to one level of softkeys whenever possible. This means that no matter where in the command syntax the user might be, there are at most eight choices for the next keyword or item. This eliminates having to search through several groups of softkeys to select the next step. By making smaller steps for generating a command, there are fewer choices at each step, and less confusion.

Two other important softkey factors are the order of the softkeys and keyword abbreviations. Order is very important; it can reduce errors of omission, and it affects the "feel" of the instrument. Abbreviations for the softkeys were selected carefully to reduce the potential for confusion.

Two final areas that affect the friendliness of a user interface are the error/status messages and the high-level interface. Users will make mistakes. It is important for the instrument to be able to identify these mistakes and tell the user what each mistake is and where it is. The 64310A has a comprehensive set of error and status messages to aid in building command sequences. In addition, when an error occurs, the display cursor is placed under the word in the command string that caused the error. The high-level interface allows the user to refer to program addresses by name or the source code line number. This allows the user to think in terms of the C or Pascal program, and to not worry about absolute addresses.

Suggestions from users and HP engineers were considered during the design of the user interface for the 64310A Analyzer and many were implemented. Functionally, the interface has two major components: setup grammar and execution-mode grammar. This division was made since the commands needed in each area are, for the most part, mutually exclusive.

Measurement setup uses eight commands on one softkey level. The order of the commands on the softkeys from left to right is, for the most part, the order of use in defining and making measurements. The first is the `define` command. With this command, the user defines the address or time

events that will be used later by other commands. The high-level interface is helpful here; it allows users to define events in terms of high-level constructs. Events can be numbered, named, and grouped together for easy reference by later commands. The second command is the `setup` command. This command is used to establish measurement termination, event periods, windows, measurement enable and disable conditions, trigger enable conditions, and the absolute file name used by the high-level interface. The third command is `measure`. Here, the user selects the measurement to be made: memory activity, program activity, module duration, module usage, intermodule duration, or intermodule linkage. The other five softkeys are for the `copy`, `show`, `configuration save/load`, `execute`, and `end` commands.

Execution mode also has eight commands. The `display` command controls the information being displayed. The options are occurrences or time, relative or absolute, and histogram or event data list. The `rescale` command controls the histogram display by controlling the axis limits. In addition, there are commands to compute statistics or not, to suspend or resume display updating, to copy the display, and to restart, halt, and end the measurement.

Symbology

As more and more microprocessor-based products appear, better ways of interfacing with the user are needed. A good solution is the symbolic interface provided in the 64310A Analyzer. This interface allows the user to specify measurement events using the same procedures and functions used in high-level Pascal or C programs, as well as line numbers from compiled listings and labels in either high-level or assembly language programs.

To accomplish this, the symbol tables generated by the 64000's linker, compilers, and assemblers were modified to include the additional information required by the 64310A. The linker symbol table provides a mapping between the physical load address of each file and the various files linked together to form the program. It also includes physical addresses of any global symbols specified. Symbol files generated by the compilers and assemblers provide a mapping between the various symbols, line numbers, and procedure/function beginning and ending offsets. With this information, the 64310A can compute the physical address of symbolic information required by its hardware to detect occurrences of a specified event.

For the user, the first step is specifying the name of the absolute file before entering any symbolic specifications. This provides the 64310A with the name of the `link_sym` file described above. The user can now specify global symbols and these can be looked up in the linker symbol table. To specify any symbolic information local to a particular file, the user must further specify the path by providing the name of the file containing the desired symbol. This becomes the default path name and the user is not required to specify it subsequently, unless the user later specifies some other file that would overwrite the old default file. Since a procedure or a function normally consists of a range of addresses in the memory space, the compiler defines the beginning symbol of the range by the name of the procedure or function in question. The endpoint of the range

is provided by the compiler by appending an R to the front of the name. For example, if a procedure named Hewlett-Packard is compiled, the compiler generates an entry in the symbol table with the name Hewlett-Packard and an offset associated with the beginning of that procedure. In addition, the compiler generates another entry in the symbol table that contains the name RHewlett-Packard and the offset associated with the return instruction from that procedure. These offsets are added to the load address found in the linker symbol table for the file that contains the procedure to create the physical address range of the procedure. The information is then entered into the appropriate data structure within the 64310A. Once an event is specified in this manner, the user can then define measurements on the event by specifying the event either by name (the event name takes on the name of the procedure) or number (as events are specified, they are numbered either automatically by the 64310A or as specified by the user).

This symbolic interface within the 64310A Software Performance Analyzer allows the user to specify a measurement at the level that is most familiar, that is, the source code and the procedures and functions found within it. This eliminates the tedious task of having to find offsets and load addresses and provide the physical addresses in some numeric entry format such as hexadecimal. In addition, when events are specified using the symbolic interface, the displayed results contain the symbols specified as well, thereby correlating the measured results to the appropriate procedure or function.

Automatic Configuration

The 64310A is used in conjunction with an emulator. Signals provided on the 64000's emulation bus vary from emulator to emulator and terms used to describe the signals are processor-specific. This presented a design challenge, because the need existed for the 64310A, in some fashion, to obtain the information it requires to operate. The solution chosen was to provide the information for each emulator independently by creating a special file format. This file supplies the 64310A with all of the specific information required to operate with each emulator.

This technique provides many benefits. A design benefit is that by packaging this information with the emulation software, the 64310A Software Performance Analyzer then becomes processor-independent. Changes in emulation, or the introduction of other emulators after the introduction of the 64310A, can be reflected in the emulation software. This minimizes the need to upgrade the hardware of the 64310A to keep the products compatible. In addition, the performance analyzer software doesn't have to decipher the emulator/analyzer interaction, because the software performance analyzer hardware is automatically configured to the emulation status and address bus. As mentioned earlier, a key feature of the 64000 System is its simple-to-use softkeys. One of the functions of automatic configuration is setting up and labeling softkeys for the specific processor (such as the status and protection levels). User interactions are as simple and high-level as possible, resulting in major time savings in setting up measurements.

To match automatic configuration files to particular emulators, the board identifier (ID) of each emulator is

appended to the end of the file name. The 64310A reads the board IDs of the cards found in the 64000's development station and creates the names for configuration files for each emulator present in the station's card cage. If the file for a particular emulator exists on the disc memory, then that emulator is supported by the 64310A Software Performance Analyzer. If more than one supported emulator is present in the development station, the user simply indicates which one is to be used for the current measurement session.

Information in the configuration file includes:

- Special flags for indicating various special cases
- Pointers to the softkey menus
- Pointers to the bus state conditions associated with the various softkey menus
- Pointers to the name field describing the associated emulator
- Pointers to the opcode fetch status
- Bit maps for configuring the address and status lines on the emulation bus
- Size list for the size of the various menus
- Pointer to the don't-care status
- The softkey label menus
- The bus state conditions associated with each softkey
- Name field for displaying the description of the emulator.

Once this file is loaded into memory, it becomes an integral part of the 64310A's software, allowing the analyzer to tailor its personality to the processor being monitored.

How It Works

The 64310A Software Performance Analyzer operates with the emulation subsystem of the 64000 Logic Development System. It is contained on a single card that monitors the 24 address and eight status lines of the 64000's emulation bus. Access to the 64310A is through the 64000's measurement system. Up to three 64310A modules can operate with a single emulator.

Data is collected by a sampling technique called sweeping. The 64310A hardware monitors a specified range (either time or address) for a specified interval, counting all events (either 1- μ s counts or number of bus cycles) that happen. The hardware is then reprogrammed for the next range. After all ranges are monitored once, a scan is completed. This measurement cycle continues until a specified termination condition is met.

The measurement is controlled by an on-board microprocessor system. This system consists of a 68B09 microprocessor, 6K bytes of RAM, and 2K bytes of ROM. Programs for controlling each measurement are loaded into the RAM from the 64000's disc memory as each measurement is called up. The processor also unloads the hardware values and stores the measurement data for later transmission to the 64000 Development Station. The ROM also contains performance verification code and a program that interacts with the 64000 to load disc-based measurement programs.

Data acquisition is performed by three real-time processes. The first process is address/status recognition. Address recognition is performed by two 24-bit ranging comparators. Each comparator generates three signals: in range, equal to lower bound, and equal to upper bound. Status

recognition is accomplished with a 256×4-bit RAM. By using a RAM, status patterns can be ORed as required by the measurement.

The second process is the state machine. The state machine performs context recognition on the incoming information and provides controlling signals for acquiring the data. The action depends on the measurement being run. The machine is constructed of programmable logic arrays, which allow a very complex set of states to be implemented with a minimum number of components.

The third process consists of programmable counters. These counters are programmed to record the events being measured. Events can be time, states, or transitions between address ranges. The counters are programmed differently for each measurement by the 64310A's microprocessor. The counters also send an interrupt signal to the microprocessor at the end of each period, which then shuts down the measurement. The microprocessor unloads and resets the counters after the interrupt is received.

The 64310A Analyzer interacts with the 64000 host processor according to a parallel processing scheme with various functions assigned as listed in Table I. The scheme allows the 64310A's local microprocessor system to run measurements independently of the host processor. The host processor controls the user interface, data display, and statistical calculations. Transmission of data to the host processor only occurs when the host has finished processing data from the previous transmission. When this happens, the host processor requests a "dump" from the 64310A. The analyzer pauses at the end of the current scan and signals the host that it is ready to transmit its measurement data. The host and the 64310A then interact, transferring the data to the host. After the host has received all the data, the 64310A resumes the current measurement.

Table I

Host Functions:
Graphics display, user interface, statistics calculations
64310A Functions:
Data acquisition, measurement control, termination of measurement on specified condition
Shared Functions:
Data transfer

Measurements

The measurement set of the 64310A allows the user to characterize code from a global point of view. Three basic measurement types are provided: activity, time distribution, and linkage. Each of these measurements has several modes of operation.

Activity. Activity measurements show frequency of activity occurring within defined areas of memory. The user can select up to 12 areas of interest by specifying address ranges, single address values, module names, or program symbols. Activity is recorded in two ways—occurrences and time. Occurrences are the number of bus cycles for the emulated processor. Time is the number of 1- μ s intervals.

Counts representing how much activity happens within each range and counts of total system activity are gathered.

Memory activity allows the incoming addresses to be qualified by the type of access they represent. These accesses can be memory reads, memory writes, I/O operations, etc. The types available depend on the status information supplied by the emulated microprocessor.

Program activity measures the amount of activity caused by executing instructions. For example, an opcode that causes a stack push may reflect only one cycle of activity within the range being monitored, but result in multiple write operations to the stack area. If the memory address of the stack push opcode is within a range being measured by a memory activity measurement, only the opcode fetch is counted. However, if the same range is being measured by a program activity measurement, the write operations to the stack are also counted (Fig. 2). By measuring areas of code using a memory activity measurement followed by a program activity measurement, the designer can get an idea of how much activity in other areas is caused by the code being analyzed.

Time Distribution. Time distribution measurements provide a best-case/worst-case characterization of code execution time. These measurements record code execution times that fall within the time ranges specified. They are best suited to structured code (modules with single entry and exit points, such as Pascal procedures). Three time distribution measurements are provided: module duration, module usage, and intermodule duration. Up to 12 time events with minimum lower limits of 1 μ s (depending on the bus cycle period of the emulated processor) and maximum upper limits of 671 s can be specified.

A module duration measurement characterizes the execution time of a module from entry to exit. Each time the module is executed, timers determine whether the measured execution time falls within one of the user-specified time ranges. When this occurs, a counter for that range is incremented. Another counter is also incremented whenever the module is executed, regardless of its execution time. The measurement can be set up to include or exclude time spent in other modules called by the module being measured. The other modules can be called by a function, by invocation, by interrupt, or by other methods, but control must eventually return to the target module.

Where possible, the module duration measurements include algorithms that correct for prefetch conditions. Since the 64310A is a general-purpose instrument that works in conjunction with all 64000 microprocessor emulators, a general approach is necessary when performing this correction. For a general approach to work, some assumptions must be made about the program module being measured. First, it is assumed that the module has a single entry point, and second, that the module also has a single exit point. Third, the module must be at least four opcodes in length (a hardware restriction). With these assumptions, two prefetch conditions that have a large effect on measurement accuracy can be corrected. The first case is prefetching beyond the exit point of the module. To correct for this case, the following algorithm is implemented on the hardware state machine:

■ If the acquired address is equal to the exit point of the

procedure, prepare to pause timers.

- If the acquired address is outside the address range of the specified procedure, pause timers.
- If the acquired address is back within the address range of the specified procedure without going past the entry point, resume timing.
- If the acquired address is the entry point to the specified procedure, reset timers and count occurrence.

The second case is prefetching into a procedure. To compensate for this condition, the following algorithm is used:

- If the acquired address is equal to the entry point of the specified procedure, reset timers and count previous occurrence if present.
- If the specified procedure is left by any point other than its exit point, stop time if excluding calls.
- If the specified procedure is reentered via its entry point, clear and reset timers.

This approach to correcting for an entry prefetch does cause problems when measuring a recursive procedure. The analyzer will interpret a recursive call as an entry prefetch and clear the timers, causing only the last recursive call to be timed.

Module usage and intermodule duration measurements provide the same type of information as module duration measurements. However, the area being timed is defined differently. Module usage is the inverse of module duration in that the time is measured from the point of leaving the module to the point of reentering the module via its entry point. That is, the time spent outside the module is measured.

Intermodule duration measures the time between leaving one module and entering another by their respective exit and entry points. These measurements do not accommodate the including or excluding of module calls.

Linkage. The linkage measurement records the immediate transfers from within one module to other modules. This measurement is useful for watching software traffic pat-

terns. For example, suppose module A calls modules B, C, and D. These calls are conditional, based on incoming data. By doing a linkage measurement on A to B, A to C, and A to D, the user can see which module is being called most often. Each transfer out of the "from" module to the "to" module is counted by the hardware, as well as all transfers out of the "from" module's address range.

Measurement Enable/Disable

At times it may be useful to be able to start/stop a measurement when a particular section of code is being executed. The 64310A provides the ability to start a measurement on user-selected address/status events. Up to two events in sequence can be defined. Measurements (only activity, usage, and duration) can also be disabled by an event. An enable and a disable can be combined to form a measurement window for data acquisition. These enable/disable functions become important when a user is performing measurements on a bank-switched or overlay system.

Measurements can also be enabled/disabled by other 64000 subsystems using the 64000's intermodule bus (IMB). The 64310A can drive or receive a trigger signal to or from other 64000 subsystems. For example, the 64302A Logic Analyzer can be told to send a trigger signal to the 64310A when the 64302A recognizes its trigger point. The measurement system functions allow up to three 64310As to be connected via the IMB, greatly enhancing measurement capabilities.

Display of Data

Two displays summarize the acquired data: a histogram and a table. Both displays present results in absolute or relative form. In addition, activity measurements can display results based on time or occurrence information. When the measurement is displayed in relative mode, the software calculates the displayed percentages based only on



Fig. 2. Typical program activity display provided by the 64310A Analyzer.

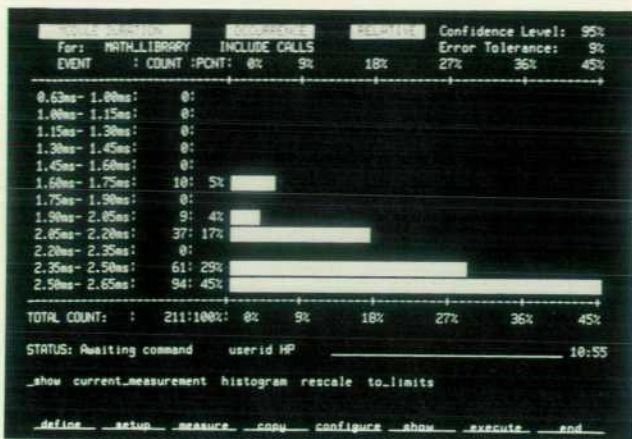


Fig. 3. Histogram display of a module duration measurement. Two automatic scaling modes are available to allow emphasis of selected parameters.

the displayed events. If the displayed data represents only a small portion of the total activity, a relative display may be preferable. If the 64310A is instructed to display the data in absolute form, the percentages are then calculated based on all activity. Interpretation of an absolute display depends on the measurement being executed. For example, if the total percentage for all displayed events for a specified activity measurement is 75%, the missing 25% means that 25% of the total activity occurs outside the events being measured. However, if the user is running a duration measurement, the missing 25% indicates that 25% of the module executions are not included within any of the time events chosen by the user. Finally, if the user is running a linkage measurement, the missing 25% represents transfers out of the "from" module to modules not specified in the measurement.

The histogram display (Fig. 3) can be rescaled any way the user chooses. This allows the user to bring out the detail in a histogram if desired. Two automatic scaling modes are supplied. One mode rescales the graph using the smallest percentage value as the lower graph limit and the largest percentage value as the upper graph limit. The other rescale method constantly adjusts the upper graph limit to the greatest percentage reading.

The tabular display (Fig. 4) shows all counts associated with the events. The counts are displayed in integer form up to the maximum count possible. The histogram display converts counts into scientific notation after 524,288 because of display memory size limitations. The table includes statistics related to each event. All data displayed on the screen can be copied to a file for later reference or additional processing.

Sampling Theory and 64310A Operation

The 64310A acquires data by sampling address/time ranges. Each range in the measurement scan is monitored for a time period called an event period. While a range is being monitored, all data is collected for that range. Both qualified and unqualified information is collected (unqualified data is used for the absolute display). After the event period for the monitored range is complete, the next range

in the scan is monitored. After all ranges in a measurement scan have been monitored, the first range to be monitored for the next scan is selected randomly to ensure statistical independence.

The event period can be based on time or occurrences, depending on which the user chooses. Default values were determined from data collected through computer simulations done in the early phases of developing the 64310A Software Performance Analyzer. The default value for activity and linkage measurements is 800 μ s. In the simulations, this value resulted in convergence to an acceptable level of accuracy in a minimum amount of time. The value chosen for duration measurements is 25 occurrences, that is, the module being measured must execute 25 times for each time range in the measurement scan. Thus, if the module executes infrequently, it may take a long time for a scan to complete. However, by specifying occurrences instead of time, the sampling period is independent of procedure execution time. As an example, consider the case where the module being measured has an execution time in the range of 10 ms to 20 ms. If the sampling period is set to 15 ms, clearly some of the execution times will be missed, since their duration exceeds the period during which the module is being measured. If an event period based on time is used and the default time is overridden, the selected time should be at least twice as long as the longest module execution time. Experimentally, it has been determined that selecting a time interval approximately twice the expected execution time plus the expected time between module calls reduces the convergence time of the measurement.

Sampling assumes that the program being measured is periodic in nature. The longer the period of the program, the longer it takes the measurement to converge. The 64310A can accumulate large amounts of data because it stores only ongoing counts for each event. The actual time a measurement runs before exceeding the resources of the analyzer is dependent on the measurement, the number of events, and the sample period. However, activity measurements usually run for four to five hours while duration and linkage measurements may run for up to eight days.



Fig. 4. Tabular display showing all counts associated with the measured events.

Statistics

To understand the statistical information provided by the 64310A Software Performance Analyzer, one needs to consider the nature of the information acquired by the instrument. First, since the analyzer measures the performance of a computer program, the program is assumed to be a stochastic process, that is, a process that is a function of time with values that are random variables. The 64310A provides information about that stochastic process by estimating values from observed data gathered by the sampling technique described above. Therefore, the results provided by the performance analyzer are not expected to be exact. The role of statistical inference is invaluable. It allows us to make generalizations about the process from which the samples were obtained, thereby providing a link between the real world and the probability models assumed in the analysis. Statistical information provides a measure of the degree of uncertainty, or conversely, the degree of confidence with which the measured results reflect the actual parameters of interest within a specified error tolerance.

The probability model used by the 64310A is based on the Student's *t*-distribution, which lends itself well to situations where observations are independent of the process. To ensure that the observations are indeed independent, extra efforts were made to create the pseudorandom number generator that determines a different starting event for each sweep. In addition, the user can define sampling event periods or occurrence counts over a wide range to ensure independent observations.

The Student's *t*-distribution used in the 64310A is based on a table that was generated with an HP 9845B Computer. This table is stored on disc, and only the values associated with the specified level of confidence are brought into host memory. These values, along with the sample mean and standard deviation computed from the observed data, are used to compute the error tolerance. This error tolerance

is displayed in the upper right portion of the display of measurement results.

An important design goal for the 64310A Analyzer was to produce sampling measurements that approximate actual target system activity as closely as possible. For example, it is necessary to use some measurement time to reset counters and transmit the information gathered. New data cannot be obtained during these steps. This measurement dead time has been minimized to avoid distorting the results. Another consideration was to reduce the number of tasks assigned to the 64310A's local microprocessor. Sums and sums of squares are computed by this processor and then transmitted to the 64000 host processor where the statistical computations are completed. The host processor calculates and displays the mean, standard deviation, percentage, and error tolerance using an IEEE floating-point software package with a modified square-root algorithm. The square-root algorithm is based on a hardware routine, and has proved to be a vast improvement in speed over the Newton-Raphson approximation technique used in the early prototypes.

Acknowledgments

Thanks go to Geraldine Paxton for production engineering and Mark Barrett for writing the performance verification software.

References

1. R.A. Nygaard, Jr., et al, "A Modular Analyzer for Software Analysis in the 64000 System," *Hewlett-Packard Journal*, Vol. 34, no. 3, March 1983.
2. M.W. Davis, J.A. Scharrer, and R.G. Wickliff, Jr., "Extensive Logic Development and Support Capability in One Convenient System," *ibid.*
3. T.A. Saponas and B.W. Kerr, "Logic Development System Accelerates Microcomputer System Design," *Hewlett-Packard Journal*, Vol. 31, no. 10, October 1980.

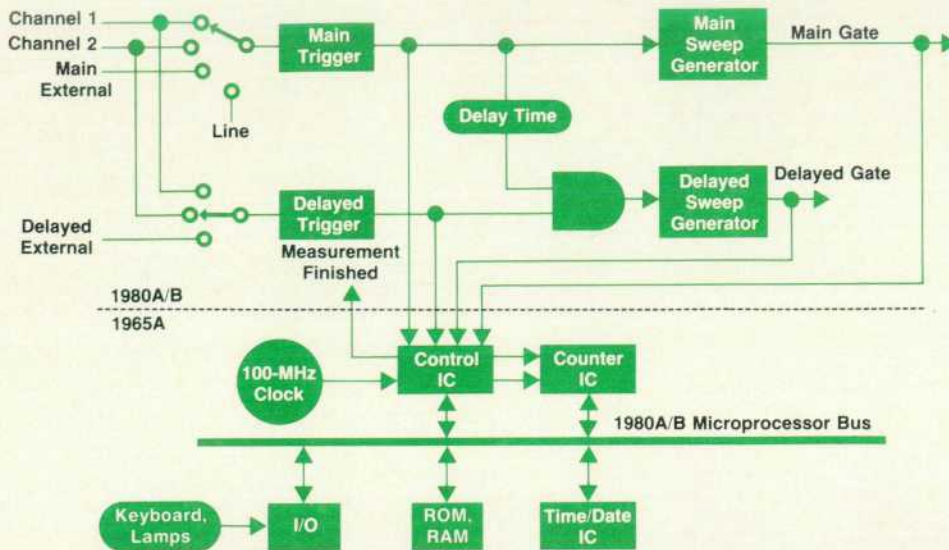


Fig. 2. Block diagram of the combined 1980A/B and 1965A. The gated universal counter uses only the 1980A/B's main and delayed gates and triggers to window specific intervals and perform all of its measurement functions.

two triggers, main trigger and delayed trigger. These triggers are routed to the 1965A and can be defined independently as the counter's A and B trigger sources. The main and delayed triggers can originate from the following 1980A/B inputs:

Main Trigger	Delayed Trigger
Channel 1	Channel 1
Channel 2	Channel 2
Main External	Delayed External
Line (60 Hz)	

The 1965A's 100-MHz clock is derived by phase modulation and frequency multiplication of either an internal 10-MHz temperature-compensated crystal oscillator or an external 10-MHz clock. Random phase modulation is used to prevent phase coherence between the counter's clock and repetitive time intervals being measured by time interval averaging (see box, page 39). Such coherence decreases resolution during averaging.

Custom Integrated Circuits

Incorporating a universal counter into an oscilloscope environment produced an extensive measurement set. Because of the power and size constraints on the expansion module, two custom integrated circuits were developed to perform all of the signal control, gating, and counting functions of the 1965A. These two ICs are configured to minimize the burden placed on the 1980A/B's processor.

The processor determines which measurement is to be made and then starts the 1965A. The measurement, which may take microseconds or hours, proceeds by itself to completion, and the 1980A/B's processor is informed of the completion. The processor reads and processes the accumulated data and displays the answer on the 1980A/B's CRT.

The simpler of the two chips is the counter IC (Fig. 3). It consists of two 31-bit binary counters capable of counting rates greater than 100 MHz. The counters use emitter-follower logic (EFL) with emitter-coupled logic (ECL) inputs. Also included are a 64-to-8 multiplexer with TTL outputs

and an interrupt circuit. The long (31-bit) counters allow hardware accumulation of counts for time interval averaging, relieving the 1980A/B's processor of this task.

The smaller but more complicated IC is the control chip (Fig. 4). It receives measurement instructions from the 1980A/B over an 8-bit data port, and it receives control signals from the 1980A/B time base and a 100-MHz reference clock from the 1965A time base. This chip then causes appropriate signals or events to be counted by the counter chip.

Depending on the particular measurement being made, the E logic block and the T logic block may select any input as an enable, the same or another signal as a disable, and the main trigger, delayed trigger, or 100-MHz reference as a clock, all under control of the control lines and reset/measurement control circuit.

For example, in time interval measurements, the 100-MHz reference is the clock. It is enabled in both the E and T logic circuits by the main gate. It may be disabled by the main trigger after the beginning of the delayed gate in the E counter and by the delayed trigger after the end of the delayed gate in the T counter. Meanwhile, the C counter

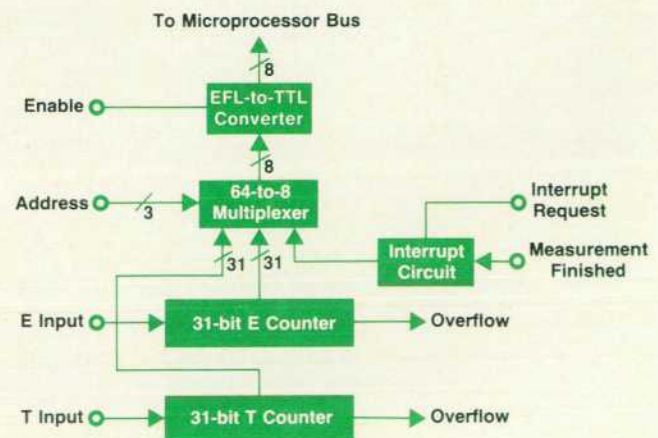


Fig. 3. Block diagram of the counter chip, one of two custom ICs in the gated universal counter module.

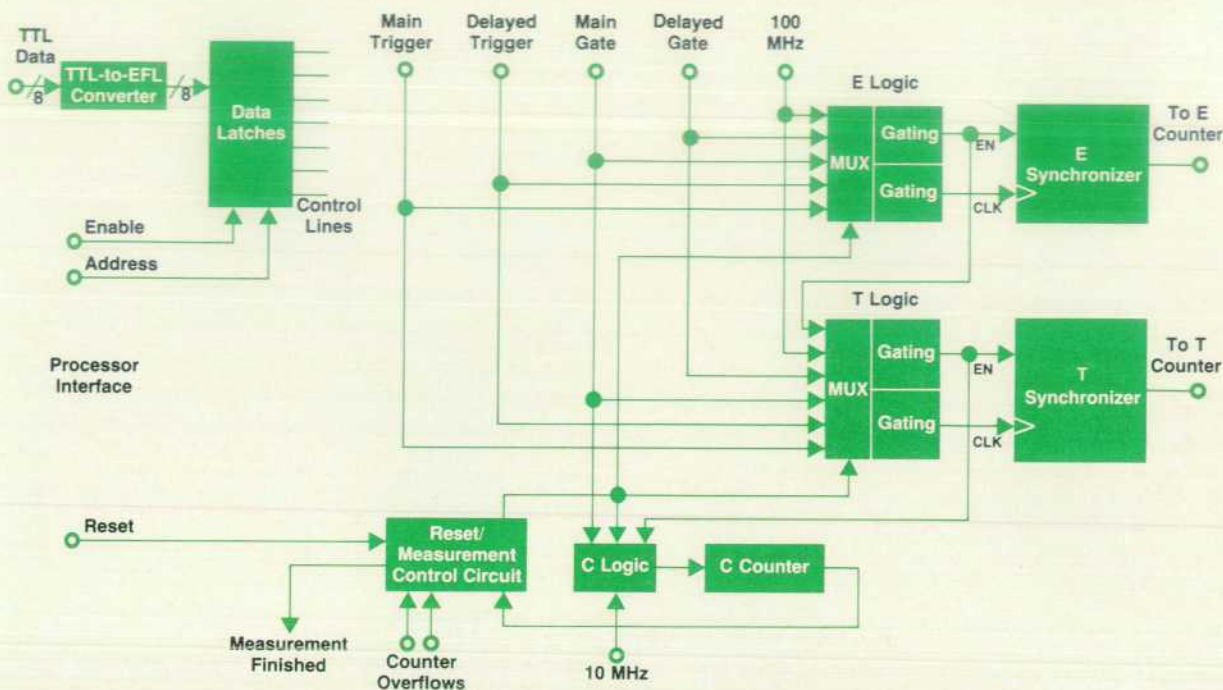


Fig. 4. The control chip, a custom IC, receives instructions, control signals, and the 100-MHz reference clock and causes the counter chip to measure the appropriate signals or events.

is counting the main-gate events to determine when the proper number of measurements have been averaged.

For a gated frequency measurement, the E logic enables the main or delayed trigger as a clock for the duration of the delayed gate. The T logic enables the 100-MHz reference under control of the E synchronizer enable circuit, while the C counter accumulates the total measurement time.

Frequency and Period Measurements

The 1965A uses the reciprocal counting technique to measure frequency and period. During both of these measurement functions, one of the 31-bit counters accumulates 100-MHz reference oscillator clocks and the input signal triggers are accumulated in the other 31-bit counter. The control IC stops the counting on the first input trigger after the sample time is satisfied. This makes measurement resolution independent of frequency, resulting in a maximum resolution of one part in 10^8 . The contents of the two counters are read by the microprocessor. Depending on the function selected, one counter is divided into the other and the result is displayed as frequency or period.

Fig. 5 shows an example of a burst frequency measurement using the gated mode. The upper waveform shows the actual counting interval. The beginning of the delayed sweep (i.e., the intensified portion of the lower waveform) arms the measurement to begin counting on the next trigger event. The end of the delayed sweep disarms the measurement, stopping the counting on the next trigger event. Measurement data is collected during the gated measurement interval on each sweep cycle of the oscilloscope. The mea-

surement continues until the cumulative duration of the gated measurement intervals equals the selected sample time.

Frequency and period can also be measured with the 1965A in its unarmed and armed modes. The unarmed mode can be used for conventional asynchronous measurements. The armed mode is used to hold off the beginning




Fig. 5. A gated frequency measurement on a burst is performed using the intensified portion (i.e., the delayed sweep gate) to window the measurement.

How Computer Control of the Oscilloscope Measurement System Makes Complex Measurements Easy

Running on HP Series 80 Personal Computers or HP 9000 Series 200 Computers, the HP 19800A/B or 19801A/B/C Waveform Measurement Library is a software package that can automatically set up waveforms, make time-domain measurements, and perform waveform comparisons using the 1980A/B Oscilloscope Measurement System and two of its options, the 19860A Digital Waveform Storage and the 1965A Gated Universal Counter. The library is composed of 48 subprograms that can be configured within a main starter program (included in the package) to provide measurement control and results.

A typical task that the 1965A can perform easily using the waveform measurement library is parametric measurements on critically underdamped waveforms. The usual difficulty in setting up this type of measurement is that of determining the 0% and 100% levels. Some counters include automatic peak detection that may passively filter the waveform to detect steady-state levels. This technique puts limits on the bandwidth and waveshape of the waveforms to be measured.

By combining several subprograms from the library, steady-state levels (top and base) can be quickly measured with the 19860A Digital Waveform Storage. Trigger levels can be set to any desired percent level. Fig. 1 shows a time interval measurement of the rise time of a peaked waveform. The following is a partial listing of the BASIC program for setting the trigger levels at 20% and 80% and measuring the rise time.

```
200 CALL Fndpct("M",0,Mr0,Mr100,Mf0,Mf100)
210 CALL Fndpct("D",0,Dr0,Dr100,Df0,Df100)
220 CALL Pcttrg("M",20,Mr0,Mr100)
230 CALL Pcttrg("D",80,Dr0,Dr100)
240 CALL Set65("T","A",0)
250 CALL Setsrc("A","M")
260 CALL Setsrc("B","D")
270 CALL Read65(Rt)
```

Lines 200 and 210 execute the subprogram Fndpct (find percent). Fndpct digitizes the waveform using the 19860A and finds the 0% and 100% steady-state trigger levels. Automatic compensation for trigger hysteresis is included. Lines 220 and 230 exe-




Fig. 1. A time interval measurement of the rise time of a peaked waveform.




Fig. 2. Measurement of the pulse width of the seventh pulse of a train of pulses.

ecute the subprogram Pcttrg (percent trigger). Trigger levels of the main and delayed triggers can be set to any percent value based on the 0% and 100% levels determined in lines 200 and 210. In this program, main is set to 20% and delayed is set to 80%. Lines 240 to 270 set up the 1965A counter to perform the appropriate rise time measurement and store the result in Rt.

Another difficult measurement task that this system can easily perform is parametric measurements on complex pulse trains. With the aid of the library, just a few lines of code can be written to program the system to measure any time parameter of any particular pulse in a train of pulses. Fig. 2 shows a photo of one such measurement. The following is a partial listing of the BASIC program that sets up and measures the pulse width of the seventh pulse.

```
200 CALL Fndpct ("D",0,Dr0,Dr100,Df0,Df100)
210 OUTPUT @Scope;"dt + 1,0;dm3;dd7"
220 CALL Pcttrg("D",50,Dr0,Dr100)
230 CALL Set65("T","U",0)
240 CALL Read65(Ref)
250 CALL Setref(Ref)
260 OUTPUT @Scope;"dt - 1,0"
270 CALL Pcttrg("D",50,Df0,Df100)
280 CALL Read65(Pw)
```

This program performs two time interval measurements. The first measurement is from the beginning of the main sweep (externally triggered) to the 50% point on the rising edge of the seventh pulse. The second measurement is from the beginning of the main sweep to the 50% point on the falling edge of the seventh pulse. The 1980A/B has a feature called digital delay that causes the delayed sweep to begin on any specified trigger event after the beginning of the main sweep. In line 210, dm3 is the HP-IB command calling for the digital delay mode, and dd7 specifies the seventh triggered event. Using the reference feature of the 1965A, the answer (Pw in line 280) is computed as the difference of the two time interval measurements.

-Johnnie Hancock

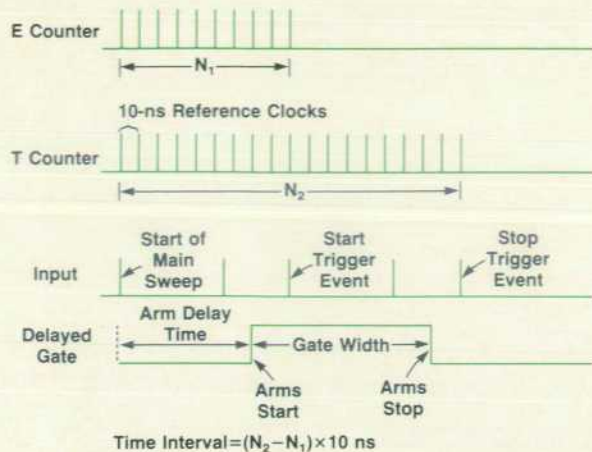


Fig. 6. The two-counter technique for measuring time intervals allows negative as well as positive results. With averaging and random phase modulation of the clock, very small intervals approaching 10 ps can also be measured. By setting the arm delay time and gate width, a time interval between specific trigger events can be measured.

of a measurement, perhaps to allow a signal to settle before it is measured. An example of this is measuring the steady-state frequency of a startable oscillator.

Time Interval Measurements

The time interval function measures the time between two trigger events. Specific trigger events are windowed, or selected for measurement, by setting the delayed sweep mode, arm delay, and gate width.

The 1965A uses two counters to measure time intervals. Both 31-bit counters are started at the beginning of the 1980A/B's main sweep gate and accumulate 100-MHz reference clocks. As shown in Fig. 6, the start trigger event halts the first counter and the stop trigger event halts the second counter. The 1980A/B displays the result after the contents of the counters are subtracted and multiplied by 10 ns. If the stop trigger event occurs before the start trigger event, the answer is negative, resulting in a time interval measurement range of -10 seconds to +10 seconds. Resolution in the time interval mode is increased or decreased by selecting 1, 10^2 , 10^4 , or 10^6 averages. Averaging is performed by hardware in the control and counter ICs, resulting in increased measurement rates compared to averaging in software. During averaging, the reference clock is randomly phase modulated to prevent coherence between the input signal and the reference clock. The combination of phase modulation and averaging makes it possible to measure repetitive time intervals as small as 10 ps.

Fig. 7 shows a gated time interval measurement on a complex digital waveform. The interval measured is from the second rising edge to the fifth rising edge. In the gated mode, time is measured from the first start trigger event after the beginning of the delayed sweep gate to the first stop trigger event after the end of the delayed sweep gate. Thus, unwanted trigger events are ignored by the 1965A, and the measurement is taken only between specified pulses. In the example of Fig. 7, the third and fourth trigger events are ignored.




Fig. 7. Measuring a time interval on a data word demonstrates the 1965A's ability to delay and measure anywhere on a waveform. The top trace is the count view waveform, which indicates the measurement interval.

When the 1965A is in the armed mode, both the start trigger event and the stop trigger event are enabled at the beginning of the delayed sweep gate. Typical applications include measuring rise time, fall time, pulse width, and propagation delay. The unarmed mode of the 1965A's time interval mode uses only one of the counters in the counter IC and measures the time from the beginning of the main sweep gate to the beginning of the delayed sweep gate. With the delayed sweep in triggered or digital delay mode, the measured time is referenced to the main and delayed trigger events. This allows a time interval measurement relative to the first main trigger event on-screen. The unarmed mode can be used with digital delay and reference




Fig. 8. Glitch detection is possible using the 1965A's events A (gated) function. The region tested is indicated by the intensified portion. The result should always be an integral number—in this case it is 3. If the result is nonintegral, in this case 3.125, a glitch occurred.

offset to perform measurements on complex waveforms in a two-step process.

Events

For counting trigger occurrences where the time parameter is unimportant, the 1965A provides six different event measurement functions: events A (gated), events A during B, ratio A/B, totalize A, totalize A+B, and totalize A-B.

During the events function, the two 31-bit counters in the counter IC are connected to the 1980A/B trigger inputs. The counters accumulate triggers depending on the event function selected. Totalize A, events A (gated), and events A during B use only one 31-bit counter. The control IC permits trigger events to be counted as a function of the start/stop key for totalize A, the delayed sweep gate for events A (gated), and input B for A during B. Ratio A/B, totalize A+B, and totalize A-B use both 31-bit counters. The microprocessor reads the contents of the counters and adds, subtracts, or divides depending on the selected function.

An example of using the events function to detect glitches is shown in Fig. 8. This example uses the events A (gated) mode and counts the number of events of the signal that occur during the delayed sweep. The results are averaged and should produce an integer answer. If an intermittent glitch (not always visible) occurs, the answer will contain a fractional count. The events A (gated) function and the other event functions are very useful in analyzing digital pulse trains.

Automatic Parametric Measurements

The 1965A takes advantage of the 1980A/B's calibrated trigger levels and voltage sensitivity to scale and measure six pulse timing parameters: rise time, fall time, pulse width, duty cycle, propagation delay, and phase shift.

The accuracy of these automatic parametric measurements depends on how accurately the trigger levels can be set. The 1965A can set trigger levels with an accuracy of one percent. This accuracy includes trigger level timing (hysteresis), trigger level offset, and trigger level nonlinearities. The one-percent level of accuracy is obtained by the 1965A because it can expand the waveform and compensate for hysteresis on trigger levels.

When a parametric measurement is selected, the input waveform is expanded to approximately 17 divisions peak-to-peak for maximum trigger resolution. The waveform's absolute peaks are then located on both main and delayed triggers. This is accomplished by an internal algorithm in which trigger levels are incremented and decremented. The counter reading is used to detect when triggering is lost, with the last trigger level that gave a stable reading being used as the absolute maximum or minimum.

To enhance the accuracy of peak detection, the 1965A automatically compensates measured peaks for trigger hysteresis. Fig. 9 shows how hysteresis can affect the trigger search. While searching for the absolute maximum, the 1965A detects a loss of triggering when the fire point (the actual trigger point) is incremented beyond the 100% point. The trigger level in divisions is recorded as the relative 100% point. Similarly, decrementing the trigger level yields the absolute minimum. However, when the 1965A detects

a loss of triggering, the arm point—not the fire point—is at the 0% point. The trigger level associated with this loss of trigger is offset from the actual 0% point, and the actual trigger point or fire point is above this level by the amount of trigger hysteresis. Hysteresis values in divisions are measured and stored during the 1965A's front-panel calibration. The 1965A accesses the appropriate value during the peak search and compensates the 0% point when searching on a rising slope. If the peak search is performed on a falling slope (fall-time measurement), the algorithm compensates the 100% point.

After the peaks have been determined for the main and delayed trigger sources, the selected autoperimeter is measured using time interval techniques in the armed mode with trigger levels set to 10%, 50%, or 90%.

System Considerations

Additional features in the 1965A provide flexibility and compatibility in a system environment.

- Propagation delays caused by cabling and probing produce significant measurement errors. The 1965A's reference offset feature is very useful in canceling systematic errors, measuring drift, and making measurements on complex waveforms.
- An extremely noisy measurement environment can cause false triggering. To provide extra noise immunity, the 1980A/B's trigger hysteresis can be increased to approximately twice the normal hysteresis.
- An external oscillator input allows synchronization of the 1965A time base with an external 10-MHz reference standard.
- For time-related measurements, the 1965A's four-year nonvolatile real-time clock/timer lets the user program up to 50 sequential alarm times. These alarms produce a service request (SRQ) over the HP-IB, a TTL pulse at the front-panel **TIMER BNC**, and/or a key sequence operation if the HP 19811A Plot/Sequence ROM is installed.
- The 1965A's self-diagnostics allow the operator to confirm proper RAM and ROM operation and give the operator confidence that the internal hardware is functional.

In systems, software is often as important as hardware. Successful software is software that is easy to use. The 1965A's functions and features are all programmable. In addition, the 1980A/B and the 1965A can operate with an HP Series 80 Personal Computer or an HP 9000 Series 200 Computer running the 19800A/B or 19801A/B/C Waveform Measurement Library (see box, page 36). The library in-




Fig. 9. Effect of hysteresis on the trigger point.

Random Phase Modulation Breaks Coherence for High-Resolution Averaging

For statistical time interval averaging to be effective, the phase relationship of the internal reference clock must be totally random relative to the input signal(s) under measurement. If this random phase relationship does not exist, then coherence is present and must be broken. One common technique for breaking coherence is by phase modulating (i.e., dithering) the counter's reference oscillator (Fig. 1). The 1965A Gated Universal Counter incorpo-




Fig. 1. The 1965A's reference oscillator produces a 100-MHz clock by frequency multiplying a 10-MHz signal from the internal temperature-compensated crystal oscillator (TCXO) or an external source. Phase modulation is added to break coherence with the signal being counted.

rates phase modulation in a rather unusual yet simple way.¹

The phase modulator section of the oscillator consists of a 10-MHz tank circuit, a varactor diode, and a digital noise source applied to an RC timing circuit as shown in Fig. 2. The digital noise source is an 8-pin IC used in electronic organs. Made by

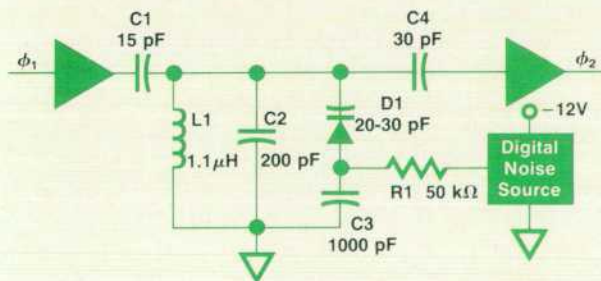


Fig. 2. For a source of random noise, the phase modulator section of the 1965A reference oscillator uses an MM5837 pseudorandom sequence generator, a part used in electronic organs.

National Semiconductor, this part (MM5837) is an MOS/MSI pseudorandom sequence generator designed to produce a broadband white noise signal for audio applications. It produces an output that switches between 0V and -12V at "random" times.

The output of the noise source is applied to the R1-C3 network. Since the pulse widths of the digital noise are always less than one RC time constant, the potential at the anode of the variable-capacitance diode, D1, is a virtually random ramping of the reverse bias voltage on the diode. The instantaneous voltage has an approximately Gaussian distribution, because the digital noise applied to R1 is a broadband white noise signal that has a Gaussian distribution of pulse widths. Fig. 3 is a photo of the digital noise and the resultant virtually random ramping signal.

The change in capacitance of D1 has a nearly one-to-one relationship to the change in its reverse bias voltage. Changes in the capacitance of D1 detune the 10-MHz tank circuit slightly, which results in a change of phase. This phase change also has




Fig. 3. The bottom trace is the digital noise source output, and the top trace is the virtually random ramp voltage at the anode of varactor diode D1 in Fig. 2.




Fig. 4. The phase modulated 10-MHz reference signal (top), an expanded view of one of its leading edges (middle), and a histogram showing the Gaussian distribution of the relative phase shift at the 50% point on the leading edge.

an approximately Gaussian distribution. Fig. 4 is a photo of the phase modulated 10-MHz reference signal. The lower waveform is an expansion of a leading edge, and the digitized waveform superimposed on it shows the relative phase distribution at the 50% point on the leading edge.

After the 10-MHz clock has been phase modulated, the frequency of the reference oscillator is stepped up to 100 MHz. The change in phase is also stepped up by a factor of 10, so that the absolute peak-to-peak phase shift of the 100-MHz clock exceeds 360 degrees relative to itself. When the peak-to-peak phase shift exceeds 360 degrees, the Gaussian distribution of phase overlaps into the preceding and succeeding 10-ns win-




Fig. 5. Unmodulated 10-MHz and 100-MHz signals.

dows, resulting in overlapping distributions. When these distributions are summed, the resulting distribution is approximately flat, thereby producing virtually random occurrences of clock edges and allowing measurement of time intervals as short as 10 ps. Figs. 5 and 6 show photos of the 10-MHz signal and the output 100-MHz signal with and without Gaussian phase modulation.

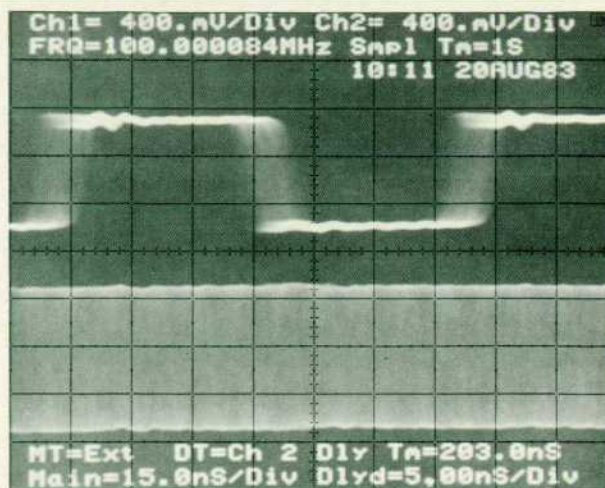


Fig. 6. 10-MHz and 100-MHz signals with random phase modulation.

Reference

1. J. Hancock, "Innovative design techniques yield optimum counter oscillator," *EDN*, October 13, 1983, p. 209.

-Johnnie Hancock

cludes first-day measurement programs (two work with the 1965A), library subprograms (13 are primarily for the 1965A), and program development aids. The library combines the setup, operation, and measurement modes of the 1965A, the 1980A/B, and the 19860A Digital Waveform Storage module into linkable, BASIC-callable subprograms.

Acknowledgments

The authors would like to acknowledge the successful efforts of the people who worked on the 1965A. Special recognition goes to Warren Tustin for his development

work on the firmware, and to Lynne Camp, who was responsible for the product design. Don Skarke and George Blinn contributed to the latch and front-bezel design. The 1965A software library extensions were written by Jonathan P. Mahaffy. Thanks also to Bill Risley and Roy Wheeler for their early product definition work.

Reference

1. W.B. Risley, "Oscilloscope Measurement System Is Programmable and Autoranging," *Hewlett-Packard Journal*, Vol. 33, no. 9, September 1982.

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

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

HEWLETT-PACKARD JOURNAL

JUNE 1984 Volume 35 • Number 6

Technical Information from the Laboratories of
Hewlett-Packard Company

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

Hewlett-Packard Central Mailing Department
Van Heuven Goedhartlaan 121

1181 KK Amstelveen, The Netherlands

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

Hewlett-Packard (Canada) Ltd.

6877 Goreway Drive, Mississauga, Ontario L4V 1M8 Canada

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