

DECEMBER 1974

HEWLETT-PACKARD JOURNAL



Improved Accuracy and Convenience in Oscilloscope Timing and Voltage Measurements

Timing measurements are made more easily and accurately with the dual-delayed sweep of a new oscilloscope. An internal microprocessor gives direct readout of time or voltage, greatly simplifying measurement procedures.

by Walter A. Fischer and William B. Risley

MANY ELECTRONIC ENGINEERS would agree that the oscilloscope is the most useful of test instruments. They do not customarily expect a high degree of precision in an oscilloscope, however, and accept the 3 to 5% accuracy that most oscilloscopes provide. The exception has been timing measurements.

Engineers concerned with measurement of very short time intervals such as rise times, propagation delay, clock phasing and other high-speed digital events depend on the oscilloscope for their timing measurements. To get accuracy in these measurements they have had to order instruments with special CRTs and specially linearized sweeps. But even with the best of conventional oscilloscopes, a major source of errors still remains in the measurement technique. The engineer either has to count graticule lines from one point on a waveform to another or, for better accuracy, he has to position the starting point at center screen with the sweep delay control, write down the control setting, position the stopping point at center screen, write down the new control setting, take the difference between the two readings, and multiply the result by the main time base setting to get the answer. Although 1% accuracy can be obtained this way, the procedure obviously has the potential for many errors.

To eliminate this bother and at the same time to improve accuracy, a new technique has been developed for the new HP 1722A Oscilloscope (Fig. 1). This oscilloscope displays two intensified markers on the waveform (Fig. 2). The operator positions the first marker at the point where the time interval measurement is to start and the second marker at the stopping point. A LED digital readout, automatically scaled to the time base setting, then displays the time lapse between the markers directly. The technique is fast and accurate, and it considerably reduces the chance for human error—there is no need to count graticule lines or calculate results from readings.

Voltage Readings Too

Better timing measurements are only one of the new capabilities of this instrument. It also makes frequency measurements and does so quickly by automatically converting a period measurement, made with the use of the markers, to frequency ($f = 1/t$). In addition, it makes measurements on the CRT vertical



Cover: The LED numeric display on this oscilloscope is an essential part of a new way of measuring very short time intervals, such as the propagation delay of a flip-flop breadboarded here on the HP logic lab. A description of the oscilloscope and the new technique begins on this page; the logic lab was described last month.

In this Issue:

- Improved Accuracy and Convenience in Oscilloscope Timing and Voltage Measurements*, by Walter A. Fischer and William B. Risley page 2
- Laboratory Notebook—An Active Loop-Holding Device* page 11
- A Supersystem for BASIC Time-sharing*, by Nealon Mack and Leonard E. Shar page 12
- Deriving and Reporting Chromatograph Data with a Microprocessor-Controlled Integrator*, by Andrew Stefanski page 18
- Adapting a Calculator Microprocessor to Instrumentation*, by Hal Barraclough, page 22.



Fig. 1. Model 1722A Oscilloscope is a high-performance dual-channel instrument with 1.3 ns rise time, 50 Ω or 1 M Ω /11 pF input impedance, sweep times to 1 ns/div, and the dual-delayed sweep that provides higher accuracy, resolution, and convenience in time-interval measurements. Its LED display gives direct readout of time intervals, frequency (1/TIME), and voltage.

axis, presenting a digital reading of the average dc value of the displayed waveform or the voltage difference between any two selected points on the waveform, such as the overshoot on a pulse. It can also derive the percentage of a part of a waveform with respect to the whole, as in measuring modulation on a carrier.

Several developments combined to achieve these capabilities. The first development is the basic oscil-

loscope, which is the same as the laboratory-grade 275-MHz Model 1720A Oscilloscope, an advance in cost-effectiveness described in the September issue of the Hewlett-Packard Journal¹. The second development is a proprietary technique known as "dual-delayed sweep", which gives the capability for more accurate determination of time intervals. The third is the microprocessor used in the HP hand-held calculators², which is built into this instrument to derive answers from the information the instrument provides.




Fig. 2. Two markers are positioned to indicate the start and stop regions of a time-interval measurement and the digital readout shows the time interval between the markers. The example here shows the pulse width to be 18.80 μ s.

Dual-Delayed Sweep

The basics of the dual-delayed sweep are shown in Fig. 3. The delayed sweep circuit itself is conventional but it can be started by either of two comparators. These are enabled alternately such that the delayed sweep starts on one main sweep when the main sweep ramp reaches the E_1 level, and on the next sweep when the ramp reaches a level equal to $E_1 + E_{\Delta t}$.

To make measurements using the dual-delayed sweep, the oscilloscope is operated in the MAIN INTENSIFIED mode in which the main sweep drives the horizontal deflection system and the delayed sweep merely intensifies the trace. The operator sets the delayed sweep to intensify short segments of the main sweep.

E_1 and $E_{\Delta t}$ are adjusted to place the two intensified segments on the points of interest, as shown in Fig. 4a. E_1 is set by the DELAY control and $E_{\Delta t}$ by the




Fig. 3. The dual delayed sweep uses two comparators that are enabled alternately. Comparator A enables the main sweep ramp to trigger the delayed sweep when the ramp reaches the E_1 level. On the next main sweep, comparator B enables the main sweep to trigger the delayed sweep at a later time when it reaches the $E_1 + E_{\Delta t}$ level.

DEC-INC switches (Fig. 5) which, when held to one side or the other, cause $E_{\Delta t}$ to decrease or increase, moving the right-hand marker to the left or right along the waveform. The microprocessor reads the value of $E_{\Delta t}$, converts it to the equivalent time interval scaled according to the time base setting, and displays the result.

Once the segments are positioned, higher accuracy can be obtained by switching the oscilloscope to the DELAYED SWEEP mode, which expands the intensified segments to full screen width, displaying the two segments overlapped as shown in Fig. 4b. The operator can then adjust $E_{\Delta t}$ to superimpose the two waveform segments exactly, as shown in Fig. 4c. The digital readout displays the time interval between the two segments with 4-digit resolution which, on the 20 ns/div sweep time range, can give 20-ps resolution.

Accuracy

The accuracy achievable by the Model 1722A Oscilloscope in time interval measurements is speci-

fied conservatively as $\pm 0.5\%$ of reading $\pm 0.05\%$ of full scale (full scale is 10 CRT divisions) on main time-base settings between 100 ns/div and 20 ms/div. When the time interval is equivalent to less than one CRT division, however, the microprocessor automatically downranges, giving $10\times$ greater resolution in the reading. Accuracy then improves to $\pm 0.5\%$ of reading $\pm 0.02\%$ of full scale (10 divisions). It is in the measurement of very short time intervals that the Model 1722A makes its greatest contribution to measurement accuracy.

Comparisons of the accuracy of the Model 1722A with that of a high-quality conventional oscilloscope are shown in Fig. 6. Whereas the percent error is about the same as a conventional high-quality scope for time intervals approaching the full display width of the CRT, the Model 1722A is superior for very short time intervals.

Measurement accuracy is enhanced by the fact that the start and stop waveform segments are displayed simultaneously. With the segments overlapped as in Fig. 4c, it would immediately become apparent if




Fig. 4. Procedure for making a time-interval measurement with the Model 1722A Oscilloscope. With the instrument set to operate in the MAIN INTENSIFIED mode, the DELAY control is used to position the first brightened segment of the trace to cover the starting point and then the DEC-INC switches are used to place the second segment over the stopping point, as in "a" (main sweep time = $0.5 \mu\text{s}/\text{div}$ and delayed sweep time = $20 \text{ ns}/\text{div}$). The instrument is then switched to the delayed sweep mode (b) and the DEC-INC switches used to superimpose the traces (c). The digital readout then gives the time interval between the pulse leading edges with 4-digit resolution, in this case $1.65 \mu\text{s}$.



Fig. 5. The microprocessor is activated by pressing one of the function buttons (POSN, INPUT, TIME, 1/TIME). If any of the oscilloscope controls are not set appropriately for the measurement selected, the digital readout displays "0".

sweep triggering had been affected by drift in the signal. And, because the operator makes his measurement by superimposing the waveforms rather than by noting where the waveform crosses graticule lines, the CRT serves simply as a null indicator so non-linearities and drift in the vertical and horizontal amplifiers do not affect measurement accuracy.

Accuracy is determined primarily by the sweep ramp generator, which is accurate within 0.02%. The delay potentiometer, which largely determines the accuracy of measurements made by the conventional differential delayed time base technique, does not enter into the measurement. The accuracy with which $E_{\Delta t}$ is derived is better than 0.005% of full scale, so $E_{\Delta t}$ does not introduce significant errors into the measurement.

Other Uses

The dual-delayed sweep gives added measurement flexibility to the oscilloscope by making it possible to view two separate expanded portions of a display simultaneously. In conjunction with the microprocessor, it can also be used as an indicator for adjusting, say, a clock repetition rate to an exact value. In this case, $E_{\Delta t}$ is adjusted to cause the digital readout to display the desired frequency in the 1/TIME




Fig. 6. Specified measurement accuracy of the Model 1722A Oscilloscope compared to a conventional high-quality oscilloscope using the differential delay technique. Plot "a" is for absolute values of time and plot "b" is in terms of horizontal deflection (in a range of 100 ns/div to 20 ms/div). The upward discontinuity at 5 ns is where the accuracy specification changes for sweep times shorter than 100 ns/div. The discontinuity shown at 1 CRT division is where the microprocessor down-ranges to give 10× better resolution.

mode, which is interpreted internally as the desired clock period. The clock repetition rate is then adjusted to cause the two waveform segments to be superimposed.

The dual-delayed sweep can also be used for measurements between points on two waveforms, such as measurements of propagation delay. When the instrument is displaying two waveforms in the ALTERNATE SWEEP mode, the delayed sweep is started by E_1 when channel A is displayed and by $E_1 + E_{\Delta t}$ when channel B is displayed, giving the time interval between the points selected on the two waveforms (Fig. 7).

The phase delay of a two-phase clock can be adjusted, for example, by displaying the master clock on one channel and the delayed clock on the other. $E_{\Delta t}$ is adjusted to cause the readout to display the exact value of phase delay desired. The clock phase delay is then adjusted to align the waveforms.

Measurements on the Vertical Axis

When the button labeled INPUT DC VOLTS (Fig. 5) is pressed, the digital readout displays the average




Fig. 7. In the *ALTERNATE* sweep mode, the oscilloscope measures the time interval between points on two waveforms. Alignment of the two points on the same vertical graticule, in this case the 50% amplitude points as in the photo at right, gives a precise measurement of time interval. (Main sweep: 0.1 μ s/div; delayed sweep: 10ns/div; time interval: 45.3 ns).

value of the input to channel A. The instrument then functions as a $3\frac{1}{2}$ -digit voltmeter with full-scale ranges from 100 mV to 50V. If a 10:1 divider probe is used, a front-panel switch compensates the reading, giving full-scale readings from 1V to 500V.

Pressing the REF SET button stores a reading as a reference. The display will then show the difference between the reference and a new voltage at the channel A input. Normally, the REF SET button is pressed while the input is grounded so subsequent readings give absolute values. Since another voltage may be used as the reference, differential readings are easily made.

The accuracy of dc voltage measurement, specified conservatively as $\pm 0.5\%$ of reading $\pm 0.5\%$ of full scale (full scale corresponds to 10 divisions even though only 6 divisions are displayed), is diagrammed in Fig. 8. As in the case of measurements on the horizontal axis, the operator does not have to count graticule lines nor multiply by range factors. Accuracy is enhanced by the fact that unlike volt-

meters with decade ranges, the vertical deflection factor ranges are in a 1, 2, 5, 10 sequence, which makes it possible to measure most voltages near full scale.

Point-to-Point Voltage Measurements

When the POSN (position) button is pressed, the DVM circuits read the level of the position control voltage. This makes it possible to measure the instantaneous voltage of any part of a waveform through dc substitution. To do this, a reference point on the waveform is selected and brought to a convenient horizontal graticule line (Fig. 9). The REF SET level is pressed to establish this graticule line as the zero level, then the position control is used to bring the point to be measured to the same line. The digital readout then displays the voltage level between this point and the reference.

Since the reference can be set to any level, the technique can be used to measure point-to-point voltages on any part of a waveform. Here again, the CRT serves simply as a null indicator with the reference and measurement point both positioned to the same graticule line, so vertical channel non-linearities, a common source of oscilloscope measurement errors, do not enter into these measurements.




Fig. 8. Curves show specified accuracy of the Model 1722A Oscilloscope in voltage measurements (a and b) as compared to a conventional oscilloscope (c). Curve "a" is for dc voltage measurements. Curve "b" is for point-to-point measurements.




Fig. 9. Double-exposure photo shows how point-to-point voltage measurements are made. The reference point is first brought to a horizontal graticule line, in this case the center line (upper trace), with the vertical position control. The REF SET button is pressed, and the other point is brought to the same line (lower trace). The digital readout displays the voltage difference between the two points.




Fig. 10. Microprocessor related circuits.

Measuring Percent

Measurements of a voltage level as a percent of a waveform are made by switching the channel A attenuator vernier out of the CAL position. The vernier is then used to establish a five-division separation between the desired zero and 100% points of the waveform on the CRT graticule. Next, the zero percent level is positioned to a reference horizontal graticule line, and the REF SET button is pressed. Positioning any other part of the waveform to the reference line then gives a reading of that waveform level in percent.

Besides quickly measuring such quantities as the percent overshoot on a pulse, this technique is also useful for defining percentage levels. For example, it can show exactly where the 50% level is on a pulse for consistent measurements of pulse width, or it can define the 10% and 90% pulse levels for risetime measurements.

Enter the Microprocessor

There are a number of ways that logic may be implemented to perform these various functions. The use of a microprocessor, however, turned out to be the most efficient way in terms of hardware and costs. It also provided a convenient means for broadening the capabilities of the instrument, such as enabling the 1/TIME calculation.

The microprocessor developed for the HP hand-held calculators was an appropriate choice for this instrument, primarily because it already had the means for driving the digital readout. The decimal adder-subtractor lends itself easily to the scaling problem, and the internal flags of the calculator permit separating and controlling the programs. The microprocessor consists of the calculator's arithmetic-and-register and control-and-timing MOS/LSI circuits² working with two ROMs designed expressly for the programs used in this instrument. The two ROMs contain a total of 512 words.

A block diagram of the circuits related to the microprocessor is shown in Fig. 10. The initial problem was to interface the calculator circuits to the oscilloscope controls and to the digital-to-analog converter that derives $E_{\Delta t}$. The front-panel controls serve as the calculator "keyboard" with the controls encoded and multiplexed to appear as keystrokes. As in the hand-held calculators, the microprocessor continuously scans the control settings to see what task is called for (TIME, 1/TIME, DC VOLTS, POSN, %) and what range factors should enter into the calculations. The input interface encodes the appropriate front-panel control settings and these are presented to the microprocessor as particular memory addresses. Programs stored at these addresses perform the indicated functions (compute time, increment, decrement, etc.)

The output interface converts the serial data to parallel data for the digital-to-analog converter (DAC), and retains it temporarily in buffer storage (the microprocessor uses words consisting of 14 BCD digits presented serially on the data bus).

During a time-interval measurement, $E_{\Delta t}$ (Fig. 3) is stored as a digital number in the microprocessor. The DEC-INC switches cause this number to be incremented or decremented, the size of the increment or decrement being determined by which of the three switches is activated. The digital number is converted to the equivalent dc voltage by the DAC.

The scaled value of $E_{\Delta t}$ is presented in units of seconds on the display in scientific notation ($A \times 10^B$) where 10 is implied and only the exponent is given. For example, $3.514 \mu s$ is displayed as $3.514 -6$. However, to simplify interpretation, only the values 9, 6, 3, and 0 are used for the exponent. With this arrangement, 128.6 ms would not be displayed as $1.286 -1$, as it would be in pure scientific notation, but as $128.6 -3$, which is easily interpreted as milliseconds. The same scheme is used for the display of




Fig. 11. Flow diagram of a voltage measurement.

frequency and also for voltage, where the exponent -3 denotes millivolts.

A block diagram of the circuits involved in a time-interval measurement was shown earlier in Fig. 3. Not shown were the interfaces and the entries from front-panel controls other than the time-base setting. These other entries blank the digital readout if the control settings are not appropriate for the measurement. This prevents the display of such ambiguous information as would occur, for example, if the sweep vernier were out of the CAL position, or if MIXED SWEEP had been selected, or if the delayed sweep TRIG LEVEL control were not in the STARTS AFTER DELAY position.

Voltage Measurements

Voltage measurements are made by comparing the input voltage V_{in} to a voltage derived by the microprocessor. The result of the comparison is reported back to the calculator, closing the loop.

The derived voltage is stored as a digital number in the microprocessor and converted to a dc voltage in the DAC. The analog amplifier assembly provides two pieces of information: (1) the polarity of the input voltage; and (2) whether the derived voltage is greater or less than the input. In response to this information, the derived voltage is incremented or decremented until it is within one least significant bit of the input. The value previously stored as the zero reference is then subtracted from this value and the result displayed.

To simplify the program and reduce the number of processing steps, the derived voltage is obtained by a successive approximation procedure. As shown by the logic flow diagram of Fig. 11, at the start of a measurement the most significant digit is set to 5. If the comparison shows this to be greater than the input, the digit is decremented to 4 and the comparison repeated.

This process continues until the comparison shows the most significant digit to be less than the unknown. This digit is retained and now the next most significant digit is set to 5 and comparisons made until the correct value for this digit is found. The process repeats for each digit until finally the derived voltage is within one least significant bit of the input. At most, only 20 iterations are required. The instrument makes about two readings per second.

If in the initial comparison the result shows the most significant digit to be less than the unknown, it is incremented upwards until it exceeds the unknown. It is then decremented one count before the comparison switches to the next most significant digit.

In a percent measurement, the microprocessor is instructed to scale the measurement as 20V/div regardless of the attenuator setting. Thus, a voltage equal to 5-cm vertical deflection is displayed as 100.0. A 5-cm deflection is thus equivalent to 100.0% and all other voltage levels are displayed as a percent of the 5-cm level.

High-Resolution DAC

The digital-to-analog converter obviously is a key element in this system. Since measurement accuracy depends upon its output, it needs superior resolution and stability, but not necessarily fast response. Available DACs that have the requisite resolution and stability are quite fast, and also very expensive. An alternate solution therefore was sought.

The DAC that evolved from this search is built around a "rate multiplier", a device that outputs pulses in proportion to the BCD number at its input.³ For example, if the number were 6, a rate multiplier would output 6 pulses for every 10 input clock pulses.




Fig. 12. Digital-to-analog converter achieves high resolution and stability with inexpensive components.

A block diagram is shown in Fig. 12. This includes the storage buffers that store the parallel data derived from the microprocessor's serial data (in BCD). If the number stored in the buffer were 6432, then, for every 10,000 clock pulses, the rate multiplier for the most significant digit would output 6000 pulses. This multiplier also gates the multiplier for the next most significant digit so this multiplier accepts clock pulses for 1/10 the time, thus outputting 400 pulses for every 10,000 clock pulses. It in turn gates the next

rate multiplier, which outputs 30 pulses. In the same way, the rate multiplier for the least significant digit outputs 2 pulses for every 10,000 clock pulses. The pulses are interleaved such that 6432 discrete pulses are supplied to the integrator, which outputs a dc voltage proportional to the number of pulses. Resolution is 1 part in 10,000.

Careful attention was paid to the design variables that affect stability. For example, it was found that an increase in ambient temperature slowed the pulse transition times while at the same time slightly increasing pulse height. The clock repetition rate was selected so these effects compensate each other, maintaining the area under each pulse constant.


Overall stability of the DAC is 0.005%/°C, eliminating it as a significant source of errors. Total cost of the components, on the other hand, is of a very low order (<\$15).

Acknowledgments

The dual-delayed sweep concept was developed by William Mordan. Product design was by George Blinn, who rearranged the Model 1720A front panel neatly to incorporate the added capabilities of the Model 1722A, and industrial design was by Bill Fischer. The authors also wish to acknowledge the contributions of all those who developed the Model 1720A Oscilloscope¹ used as the basis for the Model 1722A, including CRT designers Henry Ragsdale and Ronald Larson and hybrid circuit production expert Jay Cederleaf, and the many people in engineering, marketing, manufacturing, and quality assurance who contributed valuable suggestions on what form the 1722A should take. 

References

1. P.K. Hardage, S.R. Kushnir, and T.J. Zamborelli, "Optimizing the Design of a High-Performance Oscilloscope", Hewlett-Packard Journal, September 1974.
2. T.M. Whitney, F. Rodé, and C.C. Tung, "The 'Powerful Pocketful': an Electronic Calculator Challenges the Slide Rule", Hewlett-Packard Journal, June 1972.
3. See for example "Operation of the Digital Programmable Frequency Generator", Hewlett-Packard Journal, November 1973, p. 14.



ABRIDGED SPECIFICATIONS

Model 1722A Oscilloscope

Complete specifications available on request

Vertical Display Modes

Channel A; channel B; channels A and B displayed alternately on successive sweeps (ALT) or by switching between channels at 1 MHz rate (CHOP); channel A plus channel B (algebraic addition).

Vertical Amplifiers (2)

BANDWIDTH:

DC-COUPLED: dc to 275 MHz in both 50 ohm and high impedance input modes.

AC-COUPLED: approx 10 Hz to 275 MHz.

RISE TIME: ≤ 1.3 ns

BANDWIDTH LIMIT: limits upper bandwidth to approx 20 MHz.

DEFLECTION FACTOR:

RANGES: 10 mV/div to 5 V/div in 1,2,5 sequence.

VERNIER: continuously variable between all ranges

INPUT RC (selectable)

AC AND DC: 1 megohm shunted by approx 11 pF.

50 OHM: 50 ohms $\pm 2\%$.

MAXIMUM INPUT:

AC AND DC: ± 250 V (dc + peak ac)

50 OHM: 5V rms or ± 250 V peak whichever is less.

A + B OPERATION

Bandwidth and deflection factors are unchanged. Channel B may be inverted for A-B operation.

TRIGGER SOURCE: Selectable from channel A, channel B, or Composite.

Input - DC Volts (Channel A)

DISPLAY: light emitting diodes (LED).

NUMBER OF DIGITS: 3½.

× 1 RANGE: 100 mV to 50 V full scale vertical deflection (10 mV/div to 5 V/div).

× 10 RANGE: 1V to 500V full scale vertical deflection (100 mV/div to 50 V/div with × 10 probe).

ACCURACY: $\pm 0.5\%$ reading $\pm 0.5\%$ full scale (full scale = 10 cm), 20° C to 30° C.

STABILITY: temperature coefficient, $\leq \pm 0.02\%/^{\circ}\text{C}$.

SAMPLE RATE: approx 2/s.

RESPONSE TIME: ≤ 1 s.

REFERENCE SET: voltmeter circuits may be zeroed permitting dc voltage measurements with respect to any voltage within selected range.

OVERRRANGE: flashing display indicates overrange condition.

Position - Volts (Channel A)

(Channel A vernier in CAL detent.) With the following exceptions, specifications are the same as Input - DC Volts.

MEASUREMENT: dc substitution method using channel A position control to determine voltage of any point on displayed waveform using any graticule line as reference.

DYNAMIC RANGE: ± 6 cm from ground referenced to center screen.

REFERENCE SET: meter may be zeroed; permits instantaneous voltage measurements with respect to any voltage within selected range.

ACCURACY: $\pm 1\%$ reading $\pm 0.5\%$ of full scale (10 × the volts/div range) measured at dc.

Position - % (Channel A)

(Channel A vernier out of CAL detent.)

MEASUREMENT: dc substitution method using channel A position control to determine percent of any waveform point with respect to user defined 0 and 100% points.

RANGE: 0 to $\pm 140\%$ (calibrated with vernier so that 100% equals 5 div).

ACCURACY: $\pm 1\%$

REFERENCE SET: voltmeter circuits may be zeroed to permit percent measurements with respect to any waveform point.

Horizontal Display Modes

SWEEP MODES: main, main intensified, mixed, delayed, and × 10.

SWEEP

RANGES: 10 ns/div to 0.5 s/div (24 ranges) in 1,2,5 sequence.

MAGNIFIER: expands all sweeps by a factor of 10, extends fastest sweep to 1 ns/div.

SWEEP TRIGGER MODE

NORMAL: sweep is triggered by internal or external signal.

AUTOMATIC: bright baseline displayed in absence of input signal. Triggering is same as normal above 40 Hz.

SINGLE: in Normal mode, sweep occurs once with same triggering as normal, reset pushbutton arms sweep and lights indicator; in Auto mode, sweep occurs once each time Reset pushbutton is pressed.

TRIGGERING

INTERNAL: dc to 100 MHz on signals causing 0.5 division or more vertical deflection, increasing to 1 division of vertical deflection at 300 MHz in all display modes. Triggering on line frequency is also selectable.

EXTERNAL: dc to 100 MHz on signals of 50 mV p-p or more increasing to 100 mV p-p at 300 MHz.

EXTERNAL INPUT RC: approx 1 megohm shunted by approx 15 pF.

TRIGGER LEVEL AND SLOPE

INTERNAL: at any point on vertical waveform displayed.

EXTERNAL: continuously variable from +1.0V to -1.0V on either slope of trigger signal, +10V to -10V in divide by 10 mode (± 10).

COUPLING: AC, DC, LF REJ, or HF REJ.

AC: attenuates signals below approx 10 Hz.

LF REJ: attenuates signals below approx 15 kHz.

HF REJ: attenuates signals above approx 15 kHz.

TRIGGER HOLDOFF: time between sweeps continuously variable, exceeding one full sweep from 10 ns/div to 50 ms/div.

MAIN INTENSIFIED: intensifies that part of main time base to be expanded to full screen in delayed time base mode. Delay and time interval controls adjust position of intensified portions of sweep.

Delayed Time Base

SWEEP

RANGES: 10 ns/div to 20 ms/div (20 ranges) in 1,2,5 sequence.

MAGNIFIER: (0 to 55°C): same as main time base.

TRIGGERING

INTERNAL: same as main time base except there is no Line Frequency triggering. **STARTS AFTER DELAY:** delayed sweep automatically starts at end of delay period.

TRIGGER: with delayed trigger level control out of detent (Starts After Delay) delayed sweep is triggerable at end of delay period.

TIME INTERVAL MEASUREMENTS: measures time interval between two events on channel A (channel A display); between two events on channel B (channel B display); or between two events starting from an event on channel A and ending with an event on channel B (Alternate display).

ACCURACY

Main Time Base Setting	Accuracy (+20°C to +30°C)
100 ns/div to 20 ms/div	$\pm 0.5\%$ of measurement $\pm 0.02\%$ of full scale for measurements < 1 cm. For measurements > 1 cm, $\pm 0.5\%$ of measurement $\pm 0.05\%$ of full scale.
50 ns/div	$\pm 0.5\%$ of measurement $\pm 0.6\%$ of full scale.
20 ns/div* and 50 ms/div to 0.5 s/div	$\pm 0.5\%$ of measurement $\pm 0.15\%$ of full scale.

*Starting after 3 cm of sweep.

RESOLUTION: intervals < 1 cm, $> 0.01\%$ of full scale; intervals > 1 cm, $> 0.1\%$ of full scale; maximum display resolution, 20 ps.

STABILITY: (0°C to +55°C): short term, $< 0.01\%$. Temperature, $\pm 0.03\%/^{\circ}\text{C}$ deviation from calibration temperature range.

1/TIME FUNCTION: calculates and displays reciprocal of measured time interval.

ACCURACY: same as Time Interval Measurements.

Mixed Time Base

Dual time base in which main time base drives first portion of sweep and delayed time base completes sweep at the faster rate.

Cathode Ray Tube and Controls

TYPE: post accelerator, approx 20.5 kV accelerating potential, aluminized P31 phosphor.

GRATICULE: 6 × 10 div internal graticule, 0.2 subdivision markings on major horizontal and vertical axes, 1 div = 1 cm. Internal flood gun graticule illumination.

INTENSITY MODULATION: +8V, ≥ 50 ns width pulse blanks trace of any intensity, useable to 20 MHz for normal intensity. Input R, 1 k Ω $\pm 10\%$. Maximum input +10V (dc + peak ac).

General

REAR PANEL OUTPUTS: main and delayed gates, vertical output.

CALIBRATOR

TYPE: 1 kHz $\pm 10\%$ square wave.

VOLTAGE: 3V p-p $\pm 1\%$.

RISE TIME: < 0.1 μs .

POWER: 100, 120, 220, 240, -10% +5%; 48 to 440 Hz; 110 VA max.

WEIGHT: 29 lb (13.2 kg).

DIMENSIONS: 13-3/16 W × 7-3/4 H × 20 in. D. (335 × 197 × 508 mm).

PRICE IN U.S.A.: \$4500.

MANUFACTURING DIVISION: COLORADO SPRINGS DIVISION
1900 Garden of the Gods Road
Colorado Springs, Colorado 80907



William B. Risley

A native of Trinidad, Colorado, Bill Risley earned an AB in Physics at Princeton University in 1968. He then went to work at the Army's Fort Monmouth laboratories as a physicist specializing in electronics but left two years later to do graduate work at Colorado State University. On getting his MSEE in 1972, he joined HP's Colorado Springs Division. Spare-time activities include fishing and gardening. Bill and his wife have a one-year-old son.



Walter A. Fischer

Walt Fischer joined the Boonton Radio Corp., then an affiliate of Hewlett-Packard, in 1961 and worked on the 202J FM/AM Signal Generator. He left the next year to fulfill his military obligations but returned in 1964 and contributed to the designs of the 3211A Sweep Oscillator and the 3205A Telemetry Signal Generator. In 1968 he accepted a position as lab manager for an oscilloscope manufacturer but rejoined HP in 1972, this time at the Colorado Springs Division where he is now a group leader. Walt earned a BSEE at the Newark College of Engineering in 1961 and an MSEE at the same institution in 1968. Free time activities include horseback riding and skiing with his wife and two children, ages 13 and 11.

Laboratory Notebook

An Active Loop-Holding Device


For operation on switched telephone circuits, equipment that terminates a line-pair must provide a dc path for the holding current. The usual holding device is an inductor, but for a wideband precision measuring instrument, a prohibitively large value of inductance would be required.

For use on switched networks, an option for the Model 3770A Amplitude/Delay Distortion Analyzer provides for the dc holding current without the use of any inductors (the instrument normally presents an approximate 600Ω resistive impedance to inputs and outputs).

A diagram of the loop-holding device is shown in the drawing. R1 and C1 form a low-pass filter such that only the dc component of the signal can turn on the Darlington pair. When turned on, the Darlington pair can sink up to 100 mA dc.

To the ac component, the Darlington pair in conjunction with impedance Z appears as a current source (Z may be a low-value resistance or, for higher impedance, another active current sink). Thus, ac currents "see" a high impedance (50 kΩ).

R1, R2, and Z were chosen to ensure that the transistors do not begin conduction until the dc voltage across the device is sufficient to allow linear transistor operation with the largest ac signal voltage expected. Thus, when there is no holding current, the transistors are turned off and the device may be



left connected without causing signal distortion. Two circuits connected in parallel opposition enable currents of either polarity to be accommodated.

—David H. Guest
Hewlett-Packard Limited

A Supersystem for BASIC Timesharing

This HP 3000 Computer System is optimized for BASIC-language timesharing, but it also supports concurrent batch processing in BASIC, FORTRAN, COBOL, and SPL.

by Nealon Mack and Leonard E. Shar

THE HP 3000 COMPUTER SYSTEM is a low-cost general-purpose computer system capable of concurrent batch processing and on-line terminal processing. The system can be accessed by many users simultaneously using any of several programming languages and applications library programs. Operation is under the control of the Multiprogramming Executive (MPE/3000).¹

To meet the needs of users who want a computer system primarily for BASIC-language timesharing, MPE has now been modified to emphasize the interactive capabilities of the system. The result, called Multiprogramming Executive for Timesharing (MPET), provides the BASIC/3000 timesharing user with the fastest possible response, yet retains the ability to support concurrent multilingual batch processing.

In its most modest form, MPET supports 16 BASIC users and batch in the background (Fig. 1). Programs written in BASIC, FORTRAN, COBOL, or SPL (HP 3000 Systems Programming Language) can be run in batch mode. Calls to programs or subroutines that have been batch-compiled in FORTRAN, COBOL, or SPL can be included in BASIC user programs, a feature that can greatly increase the speed of execution of BASIC programs. Also unique among BASIC timesharing systems is the new system's ability to store and operate on integer, real, long-precision, and complex numbers in the same program. File systems are identical for timesharing and batch processing, so all files can be made available to any user in either operating mode, as desired by the system or account manager. Other features of MPET and MPE are a simple command language, complete accounting of resources, logging facilities, file backup and security, dynamic resource allocation, and virtual memory.

Two standard hardware configurations capable of running MPET are the HP 3000 Model 100CX and Model 200CX Systems. The HP 3000 Model 100CX includes an HP 3000 Computer with 48K memory, a

line printer, two 4.7-megabyte disc drives, a card reader, and a magnetic tape drive. The HP 3000 Model 200CX consists of an HP 3000 with 64K memory, a larger line printer, a 2-megabyte fixed-head disc drive, a 47-megabyte mass-storage disc drive, a card reader, and a magnetic tape drive.

What Was Done

The MPET project started with performance evaluations of the HP 3000 running under the control of the Multiprogramming Executive. The evaluation was accomplished using special software and hardware measurement aids (these aids will be discussed later) and some purely subjective reasoning. The result of this evaluation indicated that to become an optimal timesharing system, MPE would need improvement in the following areas:

- Log-on or session initiation
- BASIC subsystem access
- BASIC LIST command
- BASIC GET and SAVE commands
- BASIC run-time performance.

The Multiprogramming Executive performs system and user functions as a series of processes. It was found that when a session was initiated at the terminal or the BASIC subsystem loaded, several processes had to be created. These processes would in turn create other processes and transfer control of the system to the newly created processes. The process switching was the main factor in the amount of time required to initiate a session or load the BASIC subsystem.

Fig. 2 illustrates how a typical user process (BASIC) is created under the control of MPE. The processes at or near the root of the process tree in Fig. 2 are high-priority processes. High priority means that the execution of these processes takes precedence over all other system and user code. Thus high-priority processes are executed rather quickly under MPE.

The nodes of the process tree in Fig. 2 that are be-




Fig. 1. MPET, a modification of the Multiprogramming Executive (MPE) for HP 3000 Computer Systems, optimizes the system for BASIC-language timesharing, yet retains the ability to support concurrent multilingual batch processing.

low the dotted line are low-priority processes. This means that these processes are executed on the general process queue with all other system and user code.

When a user initiates a session a unique session main process (SMP) is created for him at high priority. However, the SMP itself executes in the general process queue. This queue is circular and is rotated in a "round-robin" fashion to allow each active process in turn to use the CPU for no more than one time slice. The number of processes on this queue will, of course, be large in a heavily loaded timeshared system.

To improve session-initiation time it is necessary to force the newly created SMP to the head of the general process queue so it can initialize itself immediately without having to wait for its turn. However, the currently executing timeshare process is allowed to complete its time slice to minimize thrashing (excessive moving of code into and out of main memory), which could result from frequent pre-emption.

A queuing analysis was performed on this method of modifying the scheduling algorithm. This study showed that, in the restricted environment of single-language timesharing, average response time could be improved by judicious use of this technique. It was felt that when a user interacts with the system he should immediately get enough CPU time to execute the majority of his requests. To achieve this the relevant process is forced to the head of the general process queue and is given a double-length time slice. Also, when control is transferred between processes on the general process queue, the newly active process is similarly forced to the head of the queue. This improves response to commands that involve the initiation of a new process—in particular, log-on and

entry into the BASIC subsystem.

Normally, initiating BASIC requires a complete process creation with all the necessary linkage editing. In MPET, however, since BASIC is invoked so often, its creation can be speeded by permanently linking it as a part of the operating system. When the first user requests BASIC a "virgin" BASIC process (which is never executed) is created and linked into the operating system's process structure as shown in Fig. 3. Thereafter, when a user requests BASIC the virgin process is merely copied and the copy on which he executes is linked as a son of his SMP.

The LIST Function

On heavily loaded HP 3000 Systems it was observed that, when listing BASIC programs or print-




Fig. 2. In the original MPE process structure, BASIC is a low-priority process that executes in the general process queue. In MPET, the BASIC user's session main process (SMP) is forced to the head of the queue when it is first created. This improves session-initiation time.




Fig. 3. Since BASIC is invoked so often in MPET, it is permanently linked into the operating system. A virgin BASIC process, created when the first user requests BASIC, is simply copied for subsequent requests.

ing data, the output would often be in spurts of four or five lines per time slice. Investigation showed that the MPE terminal buffers could hold 32 ASCII characters each with a maximum of six buffers per terminal, and that only one data transfer was made to these terminal buffers per time slice, after which the user process would be inactive until it went around the general process queue.

It's important to note that terminal buffer I/O continues after the process has lost its time slice. In other words, once the data has been buffered, it becomes a system process to output the data. With this in mind, it was felt that if the terminal buffers were arranged in a circular structure, then a process could continue to fill those buffers that had been emptied during its time slice. This guarantees that a process will have six full terminal buffers to be emptied under the control of the I/O system after it has lost its time slice. With this method it was found that approximately twice as much data could be output during and between a user process's time slices. Fig. 4 shows how this technique works.

Loading and Saving BASIC Programs

The MPE file system is a highly generalized subsystem capable of handling files of practically any type, size, or structure. A certain amount of overhead is the price paid for these conveniences. Although the flexibility of the file system is one of the advantages of MPE (and of MPET), it was obvious that the existing file system was not the most efficient way to load and

save user BASIC program files interactively.

While the MPE file system is very general, BASIC program files are very specific. That is, all BASIC program files have the same record width, 128 words, and are in all other respects identical in structure. It became obvious that a rather simple file system interface could be written to fetch and save BASIC program files. It could be simple because all the options and record sizes handled by the MPE file system would not have to be considered. This specialized interface was implemented, and as a result, most operations on BASIC program files are completed within the user's first time slice. Response times to these commands are improved by a factor of five on a loaded system.

In addition to the file system interface, certain other economies resulted by allowing two extra records for expansion of BASIC program files. Previously the BASIC subsystem created a program file of the exact length of the program to be filed. When a program was modified, it became necessary to purge the old file and create a new one of different length. This is no longer necessary. With the extra space allotted at the creation of the file, the program can




Fig. 4. To speed up the LIST function, MPET terminal buffers are arranged in a circular structure so they can be refilled as data is printed out.

expand and contract within the initial space without having to waste time purging and creating files.

BASIC Interpreter

The BASIC/3000 interpreter presented special problems. It had been well thought out by its designers and coded in a very efficient manner. Yet its run-time performance had to be improved while maintaining all the flexibility that has given this interpreter its outstanding reputation. (We feel that BASIC/3000 is the most powerful BASIC interpreter ever written.)

With the aid of measurement tools, it was determined that the interpreter was spending large amounts of time in a relatively small routine called the expression evaluator. Further investigation showed that a number of procedures were being called by this routine and the amount of time spent executing these procedure calls was significant. These procedures, some of which were very small, proved to be excellent candidates for optimization and/or relocation. Many were placed in-line, thus eliminating the time-consuming procedure calls and providing a substantial increase in run-time performance. Also, computed GOTO statements were rewritten in assembly code and this improved their execution time significantly. These modifications yield a 20% reduction of CPU load for the typical BASIC program.

General Improvements

Certain modifications to MPE and the BASIC/3000 interpreter were aimed at a general improvement of the timesharing environment, especially at increasing the number of simultaneous users of the system. Some of these modifications were initially made to solve particular problems and were later discovered to have significant effects on total system performance.

Among the most important of these modifications was the decision to restrict the types of operations that could be done at a terminal, and to give higher priority to interactive access than to batch access. Allowing only one subsystem, BASIC, to run from the terminals improves system throughput by maximizing the code-sharing capability of MPET and minimizing memory traffic. BASIC/3000 maintains user data areas nicely in that it expands and contracts them as needed, thus leaving more memory available for code. It is also the most popular timesharing language for small machines. In addition to BASIC, all the MPE commands that display system information, manipulate files, and perform general user and operator functions are still allowed from any terminal.

Other general performance modifications include faster, more specialized routines to serve the interac-

tive terminals so that terminal access no longer has to go through the file system, and placing critical sections of code in core when sufficient core is available.

The batch mode of MPE was left unchanged, except that jobs now run on a low-priority subqueue to minimize the effect on the timesharing user. User programs and any of the subsystems supported by MPE may be run in batch mode.

TEPE

MPET was a relatively short project that could not have been successful without certain performance evaluation and measurement tools developed for in-house use at HP.

The Timesharing Event Performance Evaluator (TEPE) is an HP 2100-Computer-based software system that is capable of simulating up to 32 timesharing terminals simultaneously. To run the system, the user provides a script that describes each terminal's conversation with the system under test.

TEPE transmits data from the script file to the system under test and then collects data on response time. This information is written to magnetic tape and later analyzed by an off-line process.

To create realistic models or scripts for the TEPE system and thus obtain reliable information, it was necessary to define what the typical user of a timesharing system does. The literature on this subject is sometimes ambiguous and inconclusive. However, there are a few studies on the subject that have made real contributions.^{2,3,4,5,6} These studies indicate that the typical user loads a timeshare system as follows:

- Approximately 30 to 35% of interactions result in CPU-bound jobs or tasks.
- The user requires an average think time of about 25 seconds between entries (a mixture of getting, running, modifying, and saving programs).

It was felt that if good response times were obtained from TEPE data using these two important quantifiers then there was a good chance that actual system performance would be good. In fact the typical user defined for TEPE is a bit more demanding.

From Fig. 5, a typical TEPE user interaction, it can be seen that the models include most of the operations done at a BASIC timesharing terminal. For example, all models contain operations that are characterized as CPU-bound (e.g., running BASIC programs). However, some are more CPU-bound than others (e.g., shorter programs or more I/O). Fig. 5 also illustrates a typical program file run by the TEPE system. By varying the loop parameter (N) these programs can be made to provide a variable CPU load that is in close agreement with statistics published in the literature. The degree of CPU-boundedness of the various simulated users was chosen to fit the curve shown in Fig. 6, which was derived from the pub-

```

:HELLO TEPE.SHA,R,DATA
SESSION NUMBER = #535
THU, OCT 31, 1974, 11:18 AM
HP32000C.X0.92

: BASIC

BASIC 3.0
>GET PZ075
>10 N=1200
>200 PRINT"XXXXXXXXXXXXXXXXXX"
>201 PRINT"XXXXXXXXXXXXXXXXXX"
>202 PRINT "XXXXXXXXXXXXXXXXXX"
>LIST 1-204
PZ075
1 REM THIS IS PZ075
10 N=1200
20 DIM A(10,10)
30 M=0
40 J=50
50 L=0
51 T=TIM(-1)
52 C=CPU(0)
60 L=L+1
70 IF J<>L THEN 100
80 J=J+50
90 GOTO 199
100 IF L<>N THEN 60
110 A(10,11)=3.1416
111 T=TIM(-1)-T
112 C=CPU(0)-C
113 PRINT "CPU TIME";C,"ELASPE TIME";T
120 IF M=4 THEN 160
130 INPUT K
140 M=M+1
150 GOTO 50
160 STOP
199 GOTO 100
200 PRINT "XXXXXXXXXXXXXXXXXX"
201 PRINT "XXXXXXXXXXXXXXXXXX"
202 PRINT "XXXXXXXXXXXXXXXXXX"
203 PRINT "XXXXXXXXXXXXXXXXXX "
204 PRINT "XXXXXXXXXXXXXXXXXX "

>RUN
PZ075
CPU TIME 4.106          ELASPE TIME 4.156
?5
CPU TIME 4.05          ELASPE TIME 4.059
?6
CPU TIME 4.052          ELASPE TIME 4.059
?7
CPU TIME 4.052          ELASPE TIME 4.059
?8
CPU TIME 4.05          ELASPE TIME 4.059

>NAME SW075DUM
>SAVE
>PURGE SW075DUM
>EXIT

END OF SUBSYSTEM
:BYE

CPU (SEC) = 25
CONNECT (MIN) = 5
THU, OCT 31, 1974, 11:23 AM
END OF SESSION

```

Fig. 5. TEPE, the Timesharing Event Performance Evaluator, was one of the tools used in the MPET project. It simulates up to 32 timesharing users simultaneously. This typical TEPE user interaction includes most of the operations commonly done at a BASIC terminal.

lished statistics.

The TEPE system uses a random think time between one second and 100 seconds. The mean think time is 23 seconds, which again is in close agreement with published statistics. The think time distribution is exponential, as shown in Fig. 7. In general,




Fig. 6. TEPE interactions were adjusted to provide a CPU load in agreement with this curve, which is derived from published statistics.

we feel that TEPE provides a realistic image of the typical user of MPET.

Trace

Another tool, the segment trace system, used a hardware trace facility to collect data pertaining to processes at the time of intersegment transfers. These transfers are the result of procedure calls and exits from procedures in user or system code. The system collects data that, when reduced, reveals the number of segment calls of the traced routines, tells whether the segment called was absent or present, and reveals the time spent in each segment.

Trace is handy for determining resegmentation schemes to minimize segment faults, or absences, for both user and system code.

Sampler

The software sampling system is a useful tool for measuring the relative time spent executing various sections of code. A special external clock interface is used to produce controlled random interrupts. The interrupt receiver for this clock gathers information about the environment prior to the interrupt and dumps this data to magnetic tape. A data reduction program provides reliable histograms of code execution times. Resolution is selectable and can be as fine as single instructions. With this information a programmer can easily determine those sections of code for which optimization will provide the greatest performance improvements. The sampler was especially useful for fine-tuning the BASIC/3000 run-time expression evaluator.

Results

Because the goal of this project was superior interactive performance, that is, fast response times to the user, the results of the modifications as the user sees them are of great importance. On the 16-user HP 3000




Fig. 7. TEPE simulates a random user think time between one second and 100 seconds. The think time distribution is exponential, and the mean value is 23 seconds, which agrees closely with published statistics.

Model 100CX System, and using MPE as a standard for comparison, we find that under MPET the BASIC subsystem can be loaded approximately 14 times faster. All other interactions are from 10% to 450% faster than the same interactions under MPE.

MPET on the HP 3000 Model 200CX System can support 24 or more simulated typical users with approximately these improvements. Although throughput was not measured specifically on MPET, it is evident that it has increased greatly over MPE for a BASIC timesharing load. Of course, actual performance will depend on the system load imposed by the particular user environment.

Acknowledgments

We express our appreciation to the following people for their assistance in the implementation of MPET/3000: Joel Bartlett, Tom Blease, John Dieckman, and John Hawkes (TEPE).

References

1. T.A. Blease and A. Hewer, "Single Operating System Serves All HP 3000 Users," Hewlett-Packard Journal, January 1973.
2. G.E. Bryan "Joss: 20,000 Hours at a Console," Proceedings of the 1967 Fall Joint Computer Conference, AFIPS Press.
3. A.L. Scherr, "Time-Sharing Measurement," Datamation, April 1966.
4. M. Parupudi and J. Winograd, "Interactive Task Behavior in a Time-Sharing Environment," Proceedings of the ACM 1972 Annual Conference, Association for Computing Machinery.
5. H.D. Schwetman and J.R. DeLine, "An Operational Analysis of a Remote Console System," Proceedings of the 1969 Spring Joint Computer Conference, AFIPS Press.
6. S.J. Boies, "User Behavior on an Interactive Computer System," IBM System Journal, Vol. 13 No. 1, 1974.

MPET/3000

PRICES IN U.S.A.:

32010A MPET Operating System (ordered separately), \$5000.
 Complete systems including MPET: HP 3000 Model 50CX, \$99,500.
 HP 3000 Model 100CX, \$129,500.
 HP 3000 Model 200CX, \$171,000.
 HP 3000 Model 300CX, \$203,500.

MANUFACTURING DIVISION: DATA SYSTEMS DIVISION

11000 Wolfe Road
 Cupertino, California 95014 U.S.A.



Nealon Mack (right)

Neal Mack joined HP's Data Systems Division in 1973. He's worked on performance measurement, software quality assurance, and performance and human engineering improvements on MPE/3000. Born in Shreveport, Louisiana, Neal served in the U.S. Air Force from 1963 to 1967, then attended California State University at Long Beach, graduating in 1971 with a BA degree in mathematics. In 1973 he received the MS degree in computer engineering from Stanford University. He also holds community college teaching credentials in electrical engineering, computer science, and mathematics. A resident of Sunnyvale, California, Neal enjoys reading, sports car touring, and the active social life of a bachelor.

Leonard E. Shar (left)

A native of Johannesburg, South Africa, Len Shar received his BSc degree in electrical engineering from the University of the Witwatersrand in 1968. He came to the U.S.A. in 1969 to study computer science at Stanford University and received his MS and PhD degrees in 1970 and 1972. Deciding that he liked the San Francisco bay area so much that he wanted to stay, Len joined HP's Data Systems Division in 1972. Now a project manager there, he's been heavily involved with HP 3000 performance measurements, diagnostics, and interface design. He's a member of IEEE. Bachelor Shar, who lives in Palo Alto, California, and enjoys hiking, reading, and music, is currently "trying" to teach himself to play guitar.

Deriving and Reporting Chromatograph Data with a Microprocessor-Controlled Integrator

Printing retention times next to the peaks while plotting the chromatogram, a new integrator measures the chromatograph peak areas and, at the end of the run, derives concentrations and prints the analysis on the chromatogram.

by Andrew Stefanski

ALTHOUGH GAS AND LIQUID chromatography provides fast and convenient means for analyzing the chemical components of complicated mixtures, identifying and quantifying the raw chromatograph information requires a major effort.

The result of a chromatograph procedure is a chromatogram (Fig. 1), usually made by a conventional strip-chart recorder that monitors the output of the chromatograph's detector. The substance to be analyzed is injected at the input to the chromatograph's column, a long tube packed with particles coated with a particular liquid.* The sample is carried through the tube by a carrier gas or solvent and the chemical components become separated on the basis of the differences in their solubility in the liquid coating. The lighter molecules arrive at the end of the column first, the heavier molecules coming later.

The detector responds to the presence of substances other than the carrier in the emerging stream, tracing a peak on the chromatogram for each chemical component detected. The time of occurrence of each peak corresponds to the travel time through the column and can be used to identify the corresponding chemical component. The area enclosed by the peak corresponds to the concentration of that chemical.

To calibrate the chromatogram, a known amount of a known substance is usually mixed with the sample. Reducing the data then requires the chromatographer to measure the retention times with a ruler, using the known substance as a reference, and to measure areas of the peaks by counting squares, using a planimeter, or cutting out the peaks and weighing the paper. Clearly, a lot of effort can go into reducing the data.

Speeding Data Reduction

This task was eased somewhat by the development

*See "Gas Chromatography", Hewlett-Packard Journal, March 1973, page 4.




Fig. 1. Typical gas chromatogram, this one resulting from a mixture of chlorinated benzenes.




Fig. 2. Integrator printout with the corresponding chromatogram. The lower number of each pair gives the retention time in hundredths of minutes. The upper number gives the area count in μV -seconds with the digit in the right-hand column representing a power of 10 multiplier, e.g. 3311 2 means 3311×10^2 .

of integrators that automatically compute the areas under the peaks and print the area and retention time for each peak (Fig. 2). The early electronic integrators basically were voltage-to-frequency converters that monitored the output of the chromatograph detector and drove a counter activated by rather complex peak-recognition logic. The chromatographer was still required to scale the time base, however, and to compute percentage or absolute concentration from the area counts.

The next step was to derive final results with the aid of a computer working directly from an analog-to-digital converter. The cost of doing it this way, however, was usually justified only by time-sharing the computer with several chromatographs doing repetitive analyses, such as those for checking pesticide residues or drugs.

There was an obvious need for a modestly-priced single-channel instrument that incorporated digital processing. As large-scale integrated-circuit technology advanced during recent years, it was hoped that eventually the cost of digital processing circuits would become low enough to make the computing integrator economically feasible. This hope was realized with the development of the digital processor for the HP pocket calculators.¹

New Concepts

In applying the digital processor to an integrator, the initial goal was to simplify hardware design

while enabling more sophisticated recognition of peaks, but it soon became apparent that the availability of an internal digital processor would present opportunities for new integrator capabilities. We therefore considered means of adding automatic calibration so the integrator could identify the peak belonging to the calibrating sample and then scale results accordingly. The digital processor also provides means for reducing the effects of detector noise, and for letting the instrument select the optimum slope sensitivity automatically so it can be sensitive to small peaks while ignoring noise peaks.

We also considered including the recorder as part of the integrator—combining the numerical data with the graphical data on one piece of paper would make it much easier for the chemist to relate the reduced data to the raw chromatogram. There was one major drawback to this idea: there was no suitably-priced recorder that could print as well as plot. Therefore, we developed our own (see box, page 20).

Communicating with the Processor

Although the presence of a digital processor would allow all integrating parameters to be entered through a calculator-type keyboard, it was realized that the instrument would be easier to operate if certain parameters were entered by means of slide switches. The switch positions are encoded internally and the code is sent to the digital processor. By their setting, the switches provide continuous display of the integrating parameters.

In the final design (Fig. 3), the instrument is operated entirely by the switches when end results are to be printed as "area %" (the percent area that each peak contributes to the total of all peak areas). The keyboard is used only when further computations are to be performed.

At the end of a run, the recorder prints a report on the sample analysis. The analysis identifies each peak by its retention time and gives area count and amount or percent of concentration. It also lists the integration parameters used, such as slope sensitivity and the time between the sample injection and the start of integration (start delay). Thus, the chromatographer has on one chart a complete record that includes the raw chromatogram and the reduced data (Fig. 4).

Processing the Chromatogram

The new instrument's analog-to-digital converter uses integrating digital voltmeter circuits to measure the average amplitude of the chromatograph detector output five times per second. The dual-slope technique² is used to convert the detector output voltage to digital form. The voltmeter output consists of bursts of 10-MHz pulses, the number of pulses in



Fig. 3. Model 3380A Integrator records chromatogram and analysis on the same sheet of paper. Recording and integration are controlled by the slide switches. The keyboard is used only for computations related to calibration procedures.

A Printing Plotter

Instead of a pen, the printer-plotter in the new Model 3380A Integrator uses a thermal print head. The print head, similar to those used in the printers for the HP 9800 series desk-top calculators,* has seven printing elements (heaters) in a row on a ceramic substrate. For normal recorder operation, one element is left on continuously, tracing the chromatogram on heat-sensi-

tive paper. To write characters, the elements are pulsed at appropriate moments as the carriage is moved rapidly across the paper (Fig. A).

When the integrator senses that the output of the chromatograph's detector has crested a peak, it commands the recorder to "steal" a little time from the chromatogram to write out the time of occurrence of the peak, or retention time as it is commonly called. Each peak is thus clearly identified by the retention time printed next to it.




Fig. A. Heating elements are pulsed at appropriate times to write letters and numbers. To trace a chromatogram, one element is turned on continuously.




Fig. B. Thermal print-head mounts on the recorder's carriage.

At the end of a run, the plotter prints the analysis report, identifying each peak by its retention time. The integrator retains the data until the next run so the plotter can be used to print additional copies, or it can print the results of further processing of the data by various calibration methods.

*D.B. Barney and J.R. Dreble, "A Quiet, Low-Cost, High-Speed Line Pointer", Hewlett-Packard Journal, May 1973.



Fig. 4. Typical chromatogram and report generated by the Model 3380A Integrator. The retention time for each peak is printed alongside the peak. The analysis report is printed after the chromatograph run is completed. The final two lines give the settings of the integration controls.

each burst being proportional to the amplitude of the corresponding sample.

To smooth noisy chromatograms, a running average of consecutive samples is calculated by a weighted averaging method. This smooths the high-frequency noise without distorting true peaks.

The system totals the counts obtained on successive samples but it discards the stored count if it decides that it is not measuring a true peak. This decision has always presented a dilemma to integrators. Integrators universally select peaks on the basis of slope. If the slope threshold is set too low, noise on the baseline can trigger integration. If it is set too high, integration starts high on the peak and a significant part of the peak is lost. The digital processor in the new integrator starts integration at the slightest hint of a peak but it discards the count if the peak presence is not confirmed. Thus, even with the slope threshold set high, total peak area is integrated.

Peak Criteria

The digital processor measures slope by continuously comparing each new averaged value to the previous value. If the difference is positive and exceeds a certain minimum for several successive samples, the processor judges that a peak is being detected. It then commands the plotter to mark the chart to indicate that a peak is being integrated, and continues to accumulate counts.

If the sample-to-sample comparison indicates that the slope reverses before the threshold criterion is reached, then it is assumed that a noise peak had been encountered and the total count is discarded.

Once the processor has made the decision that a peak is being detected, a reversal of slope that continues for several consecutive samples indicates that the apex of the peak has been crossed. The processor then commands the plotter to print the time elapsed since the start of the run.

Counts continue to be totaled until the sample-to-sample difference indicates that the detector output has returned to the baseline. At this time the processor stores the accumulated count, commands the plotter to place a mark on the chart indicating the end of integration for that peak, and starts looking for a new peak.

The processor memory is capable of holding counts obtained from 54 peaks in any one chromatograph run. Because there are times when a peak does not return to the starting baseline but returns to a drifting baseline or merges with a following peak, the processor also stores data pertinent to the slope reversal for later evaluation.

Automated Slope Sensitivity

As with earlier integrators, the chromatographer can select the slope criterion (mV/min.). The use of a digital processor, however, provides a new convenience: automatic slope sensitivity selection. To use this feature, the chromatographer depresses the SLOPE SENSITIVITY switch to the TEST position before starting the chromatograph run. This causes the instrument to monitor the detector output for 20 seconds and to store the maximum sample-to-sample difference encountered during that time. This is representative of the maximum noise to be expected, and it becomes the threshold level of the peak-recognition criterion.

The processor compares the beginning and end of each peak to detect baseline drift. It then adjusts the readings to account for drift, if present. This results in more accurate measurements than those made by older instruments that assume a level baseline for each peak.

Another convenience the digital processor offers is operation controlled by an Automatic Sampler. The Sampler starts the Integrator each time a new sample is injected into the Chromatograph. The Integrator's self-timer stops the integration at which time the report for that run is printed. The report includes the identification number of the bottle from which the sample was drawn. Thus long, repetitive analysis may be made with the equipment unattended.

Merged Peaks

A particular problem for integrators is finding the true areas of peaks that overlap or merge on the chromatogram. Two merged peaks are diagrammed in Fig. 5. In the new Integrator, when the sample-to-

Adapting a Calculator Microprocessor to Instrumentation

by Hal Barraclough

While the advantages of digital implementation were being considered for the next generation of chromatograph integrators, HP Labs was developing a microprocessor for a family of hand-held calculators. This microprocessor was being designed specifically for the HP-35 and its descendants, and consequently was severely limited in several characteristics essential for a complete processor structure. The cooperation of the microprocessor development team, in particular, Kenneth Peterson, was therefore most valuable in our effort at applying this microprocessor to an instrumentation problem.

Because the chromatograph integrator receives its input data in a continuous stream, the processor must be capable of real-time operation. The HP-35 is, of course, designed for human use, which signifies two fundamental characteristics: (1) relatively slow speeds, and (2) closed-loop operation through the human's own processor, for its data input rate. The first task we faced, consequently, was to ensure that the microprocessor could keep up with the required rate of data delivered from any gas chromatograph detector. Fortunately that data rate proved slow enough.

The next problem to solve was the instrument's requirement for long-term data storage. This arises in an integrator for several reasons, one of which is that several forms of relatively complex data anomalies, primarily overlapping waveforms, are inherent in the raw measurements delivered to the integrator. Their resolution can be automated accurately only by choosing the most appropriate algorithm, and this choice must be postponed until all data in the vicinity of the anomaly is received; typical examples of this are merged peaks requiring separation, drifting baselines to be distinguished from the onset of a real peak, and digital filtering of $1/f$ noise without loss of small peaks.

Two additional reasons for having large data storage are the normalization of all peaks to percentage values at the end of the run, and the printing of a final report covering the entire run.

Unfortunately, mass data storage is not a capability of calculators. In our chosen microprocessor, for example, mass memory consists of one register, yielding a capacity of one data word. We added 16K bits of data memory, organized as 512 words of 32 bits each. The data format is BCD, for compatibility with the microprocessor; we interfaced to this at the most convenient place, the A/D converter.

Next, moderate study indicated a data word of 8 digits would give us more than sufficient resolution, and 512 words gave us the capability to accept chromatograph runs considerably more complex than our original objective. After a good deal of trial designs for costing purposes, MOS shift registers were rejected for the storage medium, although their initial cost is very

low and they fit nicely into the fully serial architecture of the microprocessor. Magnetic core is not really appropriate for this type of application, and the cost of static MOS was high while yielding no advantage at the cycle times required for this instrument. The optimum choice for our integrator was dynamic MOS RAM. We designed the addressing of our 1K-by-1-bit chips to be counted sequentially through 32 steps, thereby performing the necessary serialization directly.

The most difficult problem remaining was the design of an addressing scheme for all this memory. The choices involved trade-offs between binary and BCD (our microprocessor is strictly BCD), various address computation schemes, the microprocessor's limitation of I/O to only one port, and the implications of all methods upon the execution times of the real-time program loop. Innumerable schemes were created, hardware designed on paper, and the result measured for effectiveness by microcoding the real-time procedures. At the end of all this we selected binary addressing, with the data addresses stored in ROM along with all the machine's instructions, and with a substantial dose of TTL interjected between the microprocessor and its program storage (i.e., ROM). The logic serves to occasionally fool the microprocessor by intercepting pseudo-instructions and treating them as binary data addresses. For a touch of elegance we included an index register, made it conversant with BCD via the I/O port so its utilization by the programmers is easy, and permitted it to be duplexed so it doubles as a general-purpose register. Indexing is a very useful feature for our application because of the episodic nature of much chromatograph data, with repetitive kinds of data points common to most peaks. Finally, code conversion from the index register's BCD to our address register's binary, along with address computation (performed mainly by addition with some concatenation) was mixed with the address counter for serialization. We ultimately eliminated this counter and used the system state counter for *all* timing. One can appreciate the number of timing diagrams that went to the Palo Alto paper recycling center.

The limitation of only one I/O port was solved by assigning a unique storage address to each of the following: (1) the A/D converter, for input, (2) the printer-plotter, for output, (3) the front-panel switches, for human control inputs, and (4) the index register.

The last design phase, performed while accommodating to the continuing changes made to the microprocessor by its development group, was an absolute logic minimization endeavor. All the digital processing was achieved on one 8" × 12" two-layer PC board requiring an 8-package microprocessor, 16 packages of RAM, test capability, and all the TTL and CMOS for the logic.

sample comparison indicates that the slope of the chromatogram changes sense before it reaches the baseline, the processor stores the count accumulated up to that point, starts a new count, and draws a mark on the chromatogram to indicate that a new integration has been started. It also identifies the value of the first sample in the new count for later use.

If the trace returns to the baseline on the next downslope, the two counts obtained are stored as the area counts for the two peaks. This is known as the "dropline" method of merged peak separation. During the final printout, the letter "M" is printed on the line for the second peak to indicate that it was merged with the previous peak and that this method was used.




Fig. 5. Areas of merged peaks (left) are separated by a line dropped from the valley to the baseline. The area of a small peak riding on a tailing peak (right) is computed using the tangent line as a baseline. The area in the trapezoid below the tangent line is allocated to the tailing peak.

A common occurrence is that indicated in the right hand plot of Fig. 5 where a small peak rides on the tail of a larger peak. This calls for separation by a different method, known as the "tangent skim" method.

The digital processor detects the presence of a tailing peak by storing the time elapsed between the start of the peak and its apex, and comparing that to the time from the apex to the end of the peak. A large difference classifies the peak as "tailing".

If the first peak is a tailing peak, and if the ampli-

```

METHOD:  ISTD
%RTW:  5

ID#      RT              AMT
REF  1 - 4 - 8 0        1 0 0 0
      2 - 3 - 0 2        1 0 0 0
      3 - 7 - 9 3        1 0 0 0
      4 - 1 2 - 0 0       1 0 0 0
      5 - 0

LIST
%RTW 5      ISTD
RT  ID#      AMT/AREA
4 00  1      005 319
3 00  2      003 921
7 00  3      002 81
10 00  4      002 332

XF  1.  ISTD AMT  1 000

```

Fig. 6. Typical calibration dialog. When the operator presses the CALIB key, the integrator asks for the method. In this case, the operator responds with ISTD (internal standard). The integrator then asks for the width of the retention time window (%RTW) and the identification of the calibrating peaks, which are identified by their retention times (RT) and by the amount in the sample (AMT). The dialog is ended when the operator presses 0 in response to a request for another reference. When the operator presses the LIST button, the integrator confirms the calibration by listing the parameters and the response factors, which it calculates.

tude of the second peak is less than one-half that of the first peak, then the processor uses the tangent skim method. The end of the second peak is determined by continuously calculating the slope of a line drawn from the start of the second peak to the latest sample, and comparing the line's slope to the slope of the chromatograph curve. When the two slopes coincide, the end of the peak is indicated. The proces-



Andrew Stefanski

Andy Stefanski received his Master's degree in EE from the Warsaw (Poland) Polytechnic Institute in 1962. He worked for a time at the Institute of Telecommunications in Poland, designing TV broadcast equipment but then came to the United States where he worked on an optical print reader and later on advanced consumer electronics while working towards an advanced degree at the University of Pennsylvania. Obtaining his Ph.D. in EE in 1970, he came to work for Hewlett-Packard's Avondale Division. Andy flatly states that he has no hobbies to speak of.



Hal Barraclough

A sometime commercial pilot specializing in helicopters, Hal Barraclough joined HP Labs in 1970 to work on computer architecture but he is presently on a leave of absence to teach computer design in the graduate school at Santa Clara University. Hal earned his BSEE degree at the University of Idaho in 1961 and his MSEE degree at Stanford in the HP Honors Co-op program. His most pleasant spare-time activity is the time spent with his two sons but he also derives satisfaction from designing home electronics and tending a vegetable garden at the Barraclough home in San Jose.

sor then calculates the area. In the final printout, the letter "T" is placed on the line for the second peak. The tangent-skim method can also be invoked manually any time the chromatographer decides a more accurate integration would be achieved.

By similar techniques, the processor derives counts for two or more merged peaks on the tail of a large peak by dropline to the tangent. The letters "TM" will appear in the printout for these peaks.

When a run has been completed, the integrator processes the stored data according to the method selected. It retains the data until a new run is initiated, so the chromatographer can make additional copies of the analysis, or he can process the data again by an-

other method if he so chooses.

When the slide switch labeled AREA%/METHOD is moved to the METHOD position, the digital processor initiates a dialog by way of the printer-plotter. This guides the user through the steps required to establish the calibrating parameters. An example is shown in Fig. 6.

References

1. T.M. Whitney, F. Rodé, and C.T. Tung, "The 'Powerful Pocketful': an Electronic Calculator Challenges the Slide Rule", Hewlett-Packard Journal, June 1972.
2. See, for example, A. Gookin, "Compactness and Versatility in a new Plug-Together Digital Multimeter", Hewlett-Packard Journal, April 1972.

SPECIFICATIONS HP Model 3380A Integrator

Input Characteristics

VOLTAGE INPUT: -0.01 to 1.0V
DYNAMIC RANGE: 10⁶ to 1

Output Characteristics

RESOLUTION: 1 area count = 1 μ V sec.
INPUT EXCEEDED: Warning printed in report
LOGIC MARKS: Peak recognition and termination marks
RETENTION TIME: Printed at peak apex on chromatogram in 0.01 min units, maximum 330 mins
INTEGRATION: Automatic tangent skim on tailing peaks with manual forced tangent skim possible; slope sensitivity may be selected manually or automatically; compensation for up/down drifting baselines is automatic
REPORT: Consists of a chromatogram, calculations, and listing of control settings. All stored peaks are reported, but only those identified as calibrated peaks are calculated to yield amounts.
Report is on 8 1/2 x 11" sheets of Z-fold thermal writing paper. Area% calculation format consists of four columns: retention time (as printed at peak apex); peak type; Area%. Method calculation format consists of five columns; retention time (as printed at peak apex); peak type; area; calibrated peak identifier (ID#); amount.

Controls

ATTENUATOR: 1 to 1024 in binary steps, and log presentation
SLOPE SENSITIVITY: Six settings from 0.01 to 3.0 mV/min, and auto/test positions for automatic selection

CHART: Automatic and ON, OFF positions. OFF position prevents plotting of chromatogram so each run is reported by a calculation only
CHART SPEED: Four settings: 0.5-1-2-4 cm/min
STOP TIMER: Off, and nine settings to 90 mins for automatic termination of run followed by report printout
START DELAY: Off, and nine settings to 64 mins to delay start of integration
METHOD: Selector for Area% or Method calculation, keyboard is deactivated when switch is in Area% position
CALCULATIONS (keyboard controlled): Four are standard: Area%; Normalization; Internal standard; External standard. Latter three use automatically determined or manually entered response factors. Single stored calibration shared by methods permitting any method calculation report for stored run data—no limit to number of report copies, original or modified. Special key for entry of amount of internal standard added to sample and for dilution factor. Up to 54 peaks may be calibrated.
PEAK IDENTIFICATION: Calibrated peaks other than reference peak are automatically identified by relative retention. In ESTD and NORM methods, identification by absolute retention time occurs automatically if reference peak is not found. Analyst may deliberately select this alternate type of identification for all calibrated peaks in ESTD and NORM methods.

General

DIMENSIONS: 20.6 H x 43.5 W x 57.2 D cm (8-1/8 h x 17-1/8 w x 11-1/2 d ins)
WEIGHT: 17 kg (37 lbs)
POWER: 100-120-220-240 (+5, -10%), 50/60 Hz, 150W max.
ENVIRONMENTAL: 10-50°, 0-95% rel. humidity up to 45°C
PRICE IN U.S.A.: \$5200.
MANUFACTURING DIVISION: AVONDALE DIVISION
Route 41 and Starr Road
Avondale, Pennsylvania 19311

Hewlett-Packard Company, 1501 Page Mill Road, Palo Alto, California 94304

HEWLETT-PACKARD JOURNAL DECEMBER 1974 Volume 26 • Number 4

Technical Information from the Laboratories of
Hewlett-Packard Company

Hewlett-Packard S.A., CH-1217 Meyrin 2
Geneva, Switzerland
Yokogawa-Hewlett-Packard Ltd., Shibuya-Ku
Tokyo 151 Japan

Editorial Director • Howard L. Roberts
Managing Editor • Richard P. Dolan
Art Director, Photographer • Arvid A. Danielson
Illustrator • Sue M. Perez
Administrative Services, Typography • Anne S. LoPresti
European Production Manager • Kurt Hungerbühler

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

943046PAGE150PACKAAA 157
MR JOHN L DOE PROGRAMMER
HEWLETT-PACKARD
DATA PROCESSING CENTER
1501 PAGE-MILL ROAD
PALO ALTO CA
94304

CHANGE OF ADDRESS. To change your address or delete your name from our mailing list please send us your old address label (it peels off). Send changes to Hewlett-Packard Journal, 1501 Page Mill Road, Palo Alto, California 94304 U.S.A. Allow 60 days.