FEBRUARY 1975

HEWLETT-PACKARD JOURNAL

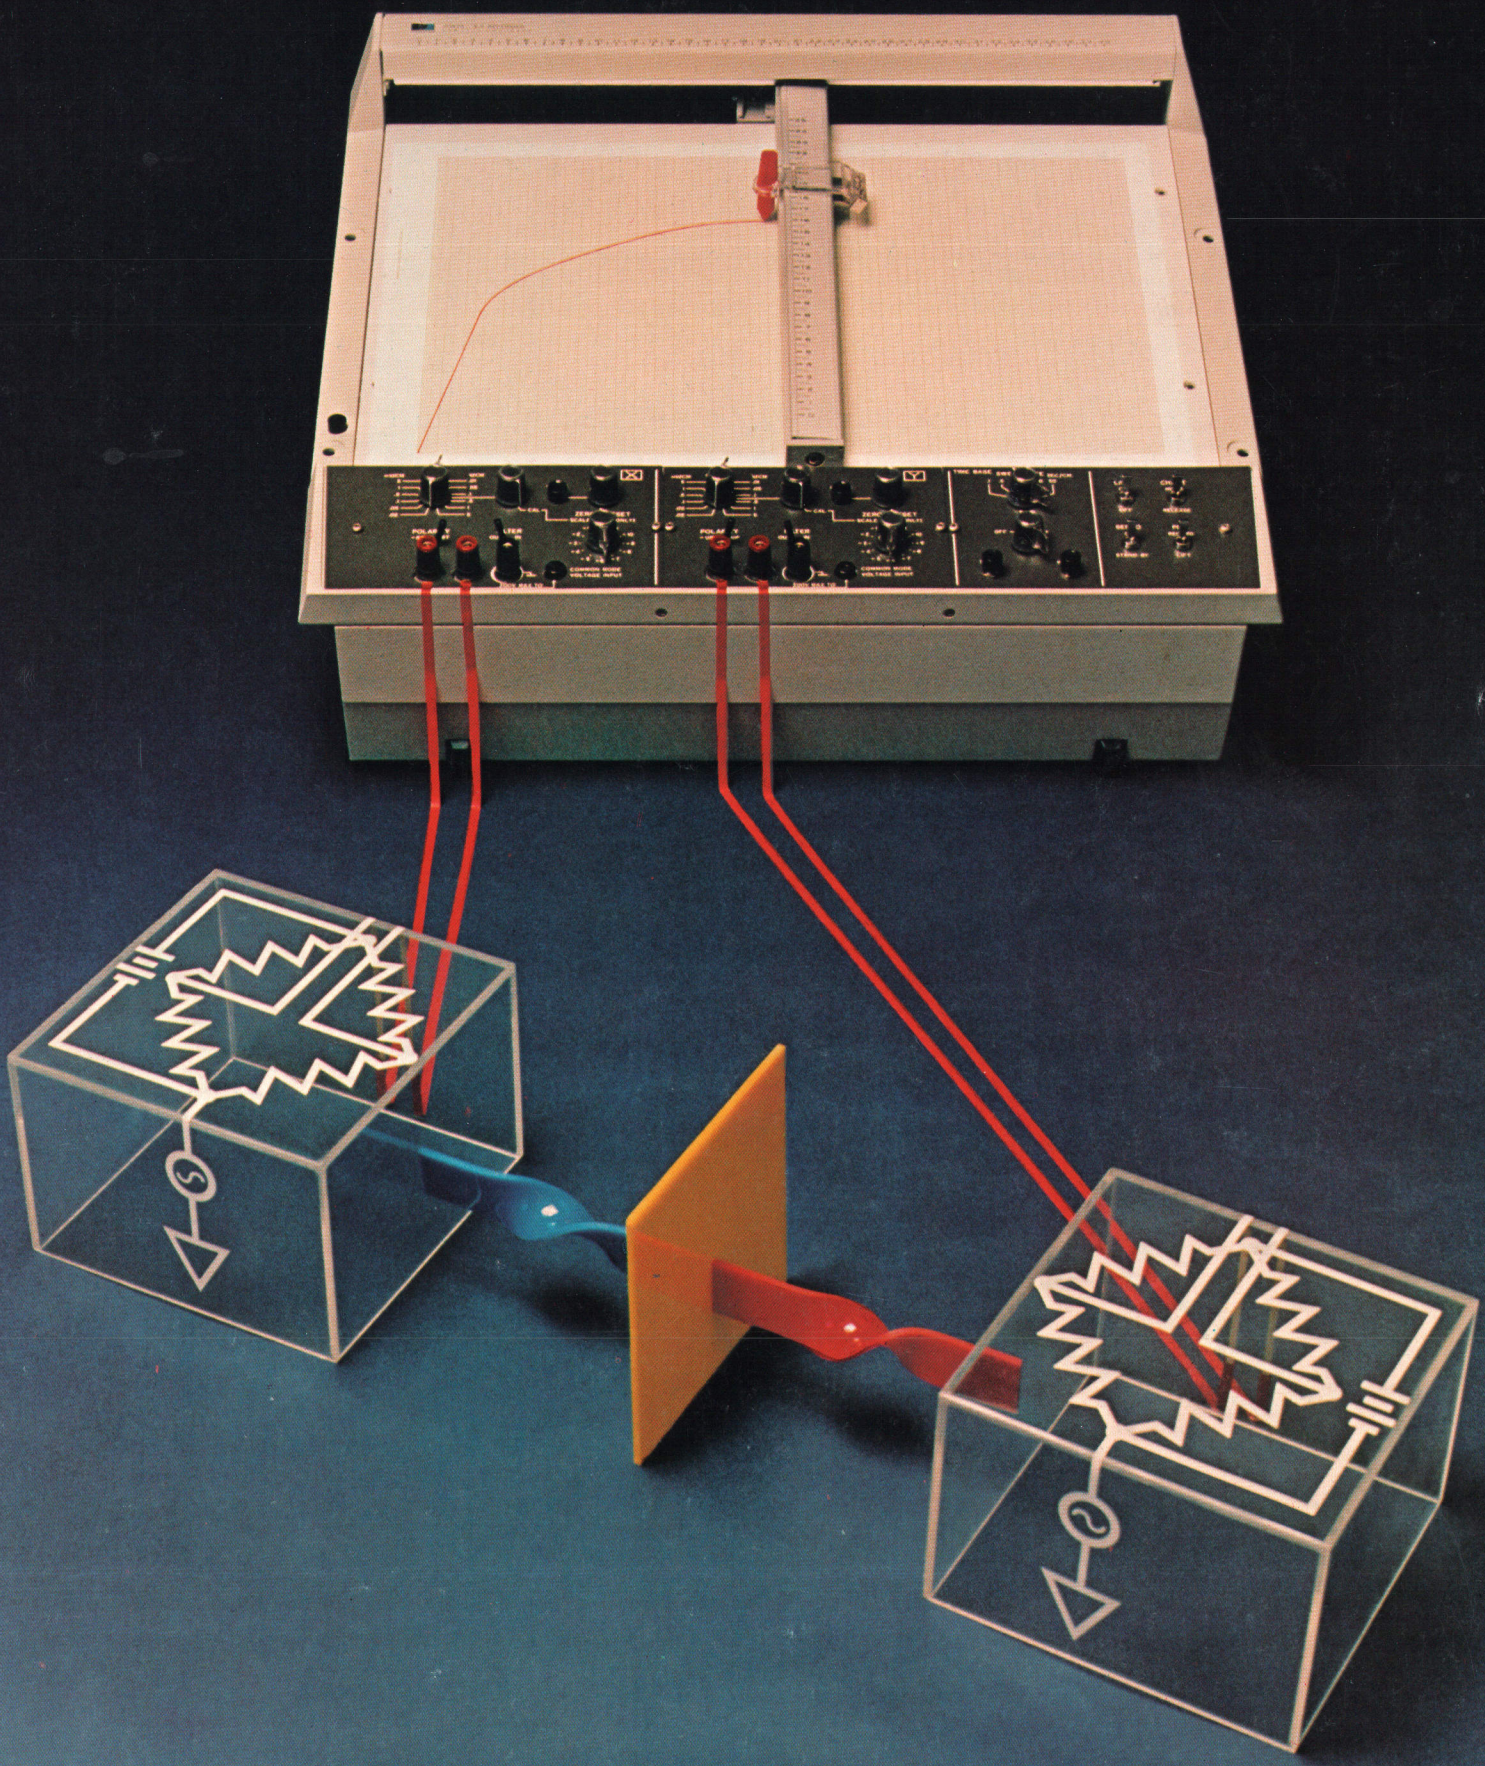

High-Sensitivity X-Y Recorder Has Few Input Restrictions

To match its high sensitivity, this new recorder provides 130 dB of common-mode rejection for virtually any input configuration. No external "guard" connection is needed unless the common-mode voltage exceeds ten volts peak.

by Donald W. Huff, Daniel E. Johnson, and John M. Wade

THE NEW MODEL 7047A RECORDER is the fastest, most sensitive X-Y recorder ever built by Hewlett-Packard. This rugged, reliable, one-pen laboratory recorder (Fig. 1) has maximum input sensitivity of 0.02 millivolts per centimetre on both axes. It is capable of slewing at more than 76 cm/s and of accelerating at rates exceeding 5080 and 7620 cm/s on the X and Y axes, respectively.

Making the high sensitivity usable is a special common driver amplifier circuit that, acting as an internally driven guard, provides 130 dB common-mode rejection with virtually any input configuration. This is a significant difference from other recorders, which require that the LO input terminal always be connected to the low side of the source if high common-mode rejection is needed. No external "guard" connection is required by the 7047A unless the common-mode voltage exceeds ten volts peak.

Standard features of the 7047A Recorder are twelve calibrated dc input ranges on each axis, a time base that provides six sweep speeds from 0.1 s/cm to 50 s/cm, calibrated zero offset with zero control, switchable input filters, polarity reversal switches, and remote control of many functions by TTL signals or contact closures. The servo-motors are continuous-duty aluminum-framed dc motors that do not wear when the pen is driven off scale. Chart paper up to 11 by 17 inches or DIN size A3 is held in place electrostatically. Pens are disposable, and four colors of ink are available.

Options include metric or English calibration and an event marker.

7047A Design

The 7047A Recorder is a member of the 7040 Series and shares many parts with other members of that series. The basic aluminum mainframe casting is the same in all members of the series. The mechanics, motors, and servo electronics of the 7047A are the same as those of other high-performance members of the series.

The principal differences between the 7047A and other members of the 7040 Series are in the preamplifiers that condition the input signals and drive the servo electronics. The dc amplifier used in other members of the series does not have the very low

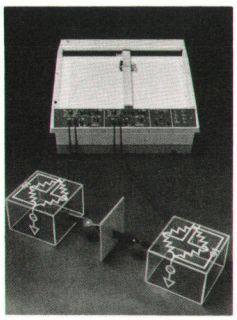

Cover: *Our artist's representation of a two-bridge measurement that's easy for the 7047A X-Y Recorder, difficult or impossible for others. Stress-versus-strain for an unknown sample is measured by applying equal forces to the unknown and to a sample of known characteristics, using the 7047A to plot one deflection against the other. One arm of each bridge is a strain gauge. The ac generators represent common-mode voltages frequently encountered in such measurements.*

In this Issue:

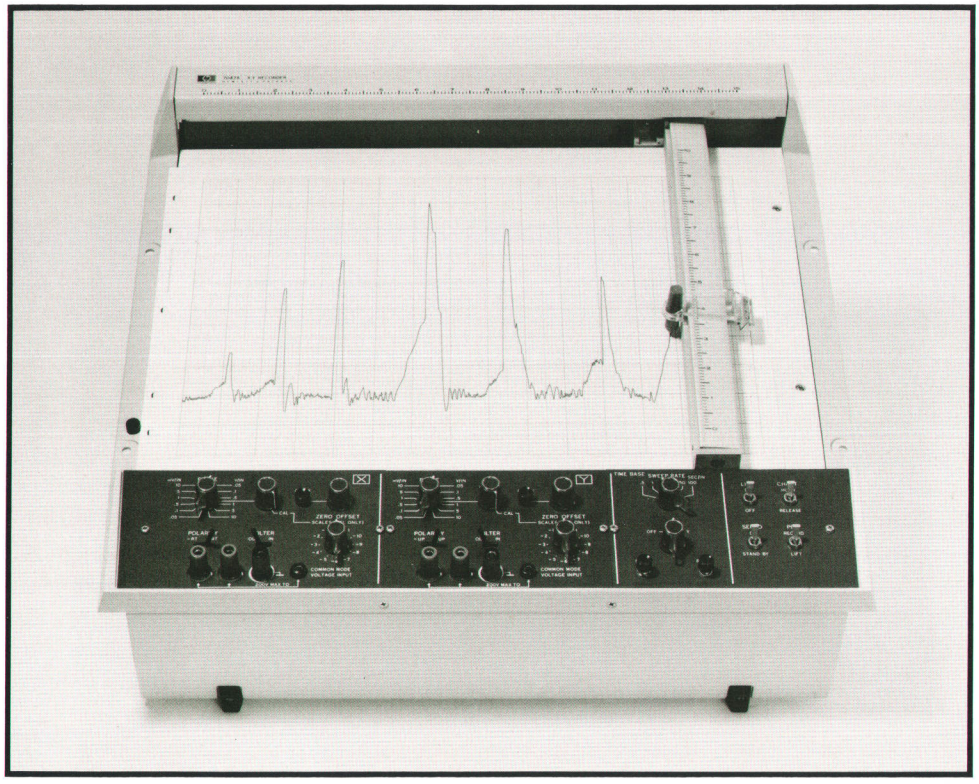

Fig. 1. *Model 7047A X-Y Recorder has maximum input sensitivity of 0.02 mV/cm (or 0.05 mV/in with English scaling) on both axes. Common-mode rejection is greater than 130 dB with 1-kΩ impedance in series with either input terminal (or both). Six sweep speeds, calibrated zero offset, polarity switches, and input filters are standard features.*

noise and drift required in a high-sensitivity recorder. Therefore a new amplifier was designed, using a chopper configuration (Fig. 2). It provides the necessary gain (1200 on the most sensitive range) without any significant noise or drift problems.

Also because of the new recorder's high sensitivity, greater common-mode rejection was required for the 7047A than for other 7040-Series recorders. Its specified common-mode rejection ratio (CMRR) of 130 dB is not only 20 dB greater than that of others in the ser-

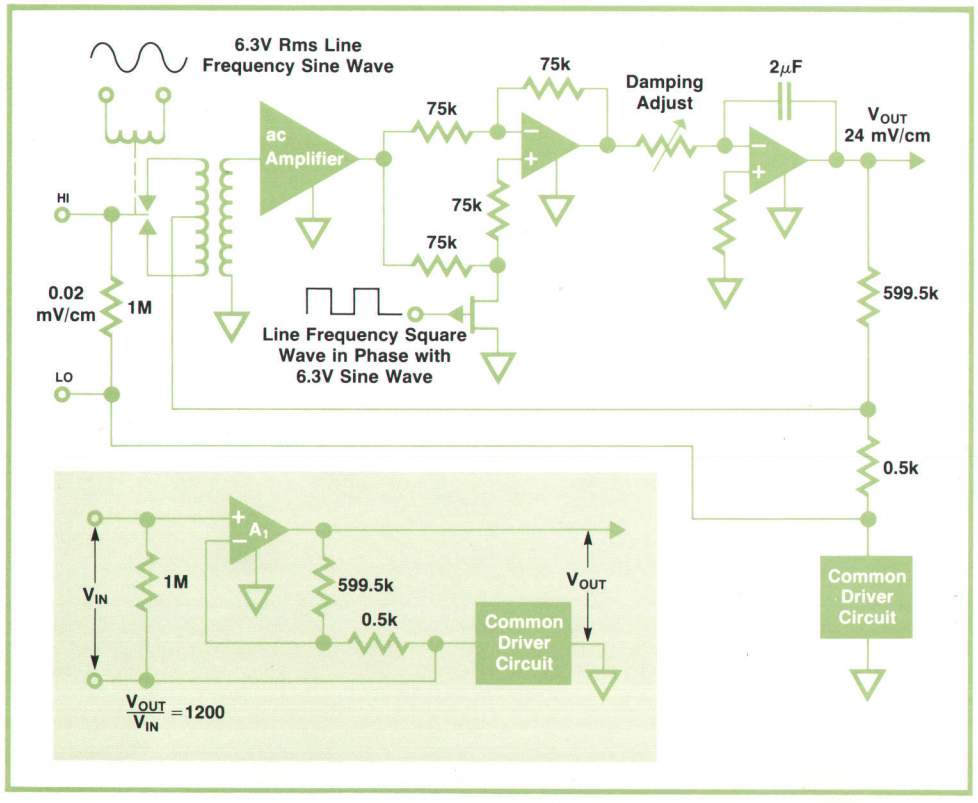

Fig. 2. *High sensitivity of the 7047A Recorder comes from its chopper stabilized input preamplifier. Shown here are a simplified schematic of the preamp and an equivalent circuit (inset) for the most sensitive range. The preamp acts as a dc operational amplifier, A_1.*

ies, but guaranteed as well for a 1-kΩ impedance in either the HI or the LO input lead. The input configuration is virtually unlimited, as shown in Fig. 3. Most recorders cannot operate on their most sensitive ranges with the HI terminal grounded and a 1-kΩ impedance at the LO terminal (Fig. 3b), because a noticeable line-frequency buzz appears at the pen tip. This problem is caused by the unbalance capacitance to ground at the secondary of the power transformer, which creates a noise voltage and causes a current to flow out of the common input or LO terminal, through the unbalance and/or source resistance and then to ground. The resulting voltage drop across the resistance appears as an input signal to the recorder. The high CMRR requirement plus the need to eliminate this noise pump-out current from the LO terminal led to the adoption of the "common driver" scheme. This circuit, which is described in detail later in this article, represents the biggest single contribution to the 7047A's ability to make the wide variety of measurements it does make.

The packaging that houses the somewhat sophisticated front-end electronics of the 7047A (input terminals, zero-check switch, filter switch, range switch, chopper preamp, and common driver) as well as various other functions (polarity switch, calibrated zero offset switch, zero potentiometer and vernier potentiometer) consists of two shielded, totally enclosed, virtually identical boxes, one for each axis (Fig. 4). A protruding cable with multipin connector connects each preamplifier to the appropriate servo electronics inside the mainframe.

The casting used to mount the boxes, the time-base switches, and power controls is the same as that used to mount the input terminals, various switches, and power controls for the older 7046A X-Y Recorder. Hence, no new castings were required for the 7047A, and except for that associated with the input boxes, no new tooling was required. Thus the development cost was kept relatively low, and most of the engineering effort could be directed toward the design of the chopper amplifier and the common driver circuit.

CMR and Noise Problems

Recorder amplifier circuits usually have a common input terminal, LO, and a high input terminal, HI. The input signal V_i and source resistances R_H and R_L are connected in series between these terminals, while a common-mode voltage V_{CM} is applied to both. See Fig. 5.

The secondary winding of the transformer in the power supply of the instrument usually has an unbalance capacitance to ground, C_{UB}. This creates a noise voltage V_N and causes a current I_N to flow out of the common input terminal, through the source resistance R_L and then to ground. The resulting voltage drop, $I_N R_L$, across the source resistance appears as an input signal to the circuit. Hence the true input signal is distorted.

Another problem arises when a common-mode voltage is applied as shown in Fig. 5. The common-mode voltage causes a current I_{CM} to flow through the leakage impedance* $R_{CG} || C_{CG}$ to ground and back to the common mode voltage source. This again causes a voltage drop across the source resistance R_L, which appears as an input to the circuit and distorts the true input.

Several methods have been used to reduce the ex-

*C_{UB} is neglected here because $C_{UB} << C_{CG}$.

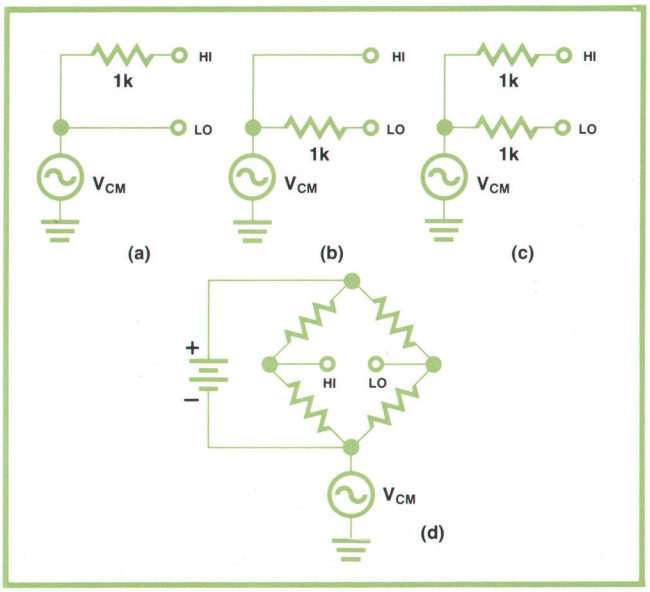

Fig. 3. *7047A meets its CMR specifications with any of these input configurations.*

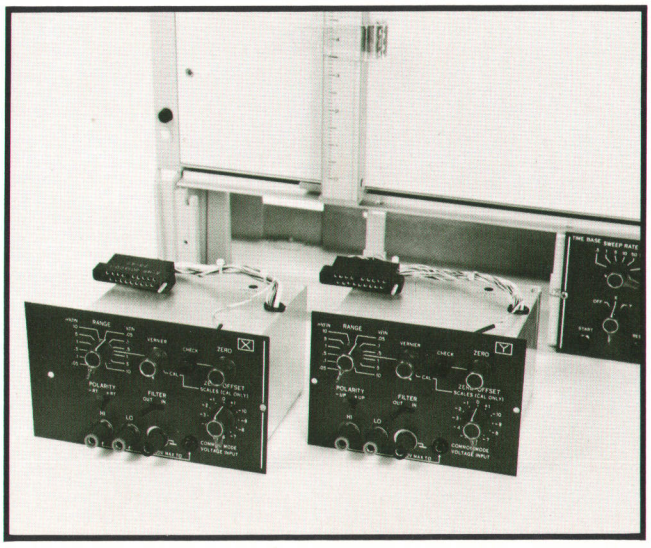

Fig. 4. *Two shielded boxes house the sophisticated 7047A front-end electronics. The boxes are removable for servicing.*

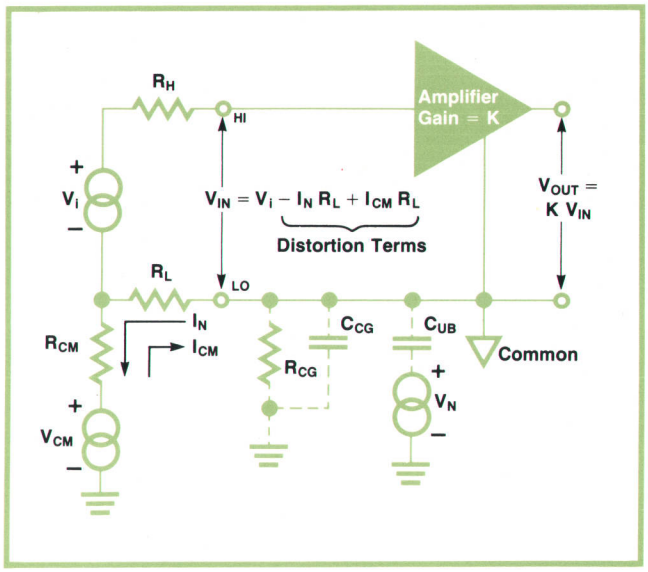

Fig. 5. *Conventional input amplifiers have problems because of pump-out and common-mode currents flowing through the source resistance, R_L.*

traneous currents superimposed on the input signal. One method is to enclose the secondary windings of the power transformer in a box shield. This eliminates most of the stray capacitance to ground. However, since there must be an opening for the transformer wires, the capacitance shield is not complete. This method can reduce the stray capacitance by an order of magnitude, but is very expensive.

Another possible solution is to prohibit any input resistance between ground and the common input terminal. This is unsatisfactory because in many measurement applications a source resistance cannot be avoided. Fig. 3d shows one such situation.

A third method applies when only one side of the input signal has a source resistance. It provides a polarity switch that reverses the polarity of the preamplifier output for the same polarity of input signal. By reversing the leads to the input terminals and using this polarity switch, the source resistance can always be placed in series with the HI input terminal. This removes the input resistance from the extraneous current path and greatly reduces the distortion of the input signal. This solution, however, is inadequate when there is an impedance in both input terminals simultaneously (e.g., bridge measurements), or when an impedance in series with one terminal has to be switched from one terminal to the other without adjustments.

Another method uses separate amplifiers for each input terminal followed by a differential amplifier to drive the output. The high impedance of the input amplifiers prevents any currents from flowing out of the input terminals. However, resistor matching problems limit the expected common-mode rejection to about 80 dB, far below the 130 dB desired for the 7047A. Another problem with this configuration is that the voltage drift of the two input amplifiers has an additive effect on the output voltage. Therefore, to achieve a given stability both amplifiers must have no more than half the drift of the original single preamplifier. Also, the total noise voltage is approximately 40% more than the original noise voltage.

A widely used method is guarding, or placing an electrical shield called a guard around the circuit to isolate the circuit from ground. This achieves two things. First, the internally generated noise current flowing out of the input terminals is greatly reduced. Second, high common-mode rejection, especially at line frequency, can be realized by driving the guard terminal at the common-mode voltage. This diverts the common-mode current around the source impedance in the LO terminal. Guarding has some drawbacks, such as additional cost, and the burden of driving a guard terminal properly.[1]

The Common Driver Amplifier

The common driver amplifier used in the 7047A permits measurements to be made with better than 130 dB common-mode rejection at dc and line frequency with a 1-kΩ resistance in either the HI or the LO input lead or both. This holds over the full Hewlett-Packard Class B environmental range (0° to 55°C, up to 95% relative humidity at 40°C). The effects of unbalance capacitance from the power transformers are virtually eliminated and a "guard" terminal does not have to be driven provided the common-mode voltage is less than the output voltage range of an operational amplifier in the circuit. Since roughly 90% of all measurements and recordings taken involve less than 10V, this voltage was chosen as a reasonable limiting value. When this voltage is exceeded, an auxiliary input terminal (COMMON MODE VOLTAGE INPUT) may be driven to circumvent the saturation limitation of the amplifier. The only shielding required is to isolate the HI and LO input terminals from chassis ground. This shield is connected to circuit common, which is now an inaccessible point inside the recorder.

How It Works

Fig. 6 is a simplified schematic diagram of the complete 7047A preamplifier, including the common driver amplifier. The noise generator V_N and unbalance capacitor C_{UB} represent the current source from the power transformer. The leakage impedance ($R_{CG}||C_{CG}$) is due to many factors; for example, the preamplifier shield, the servo amplifier, and the servo motor, all referenced to circuit common, are in close proximity to chassis ground. The input signal to be amplified is V_i and the source resistances in the

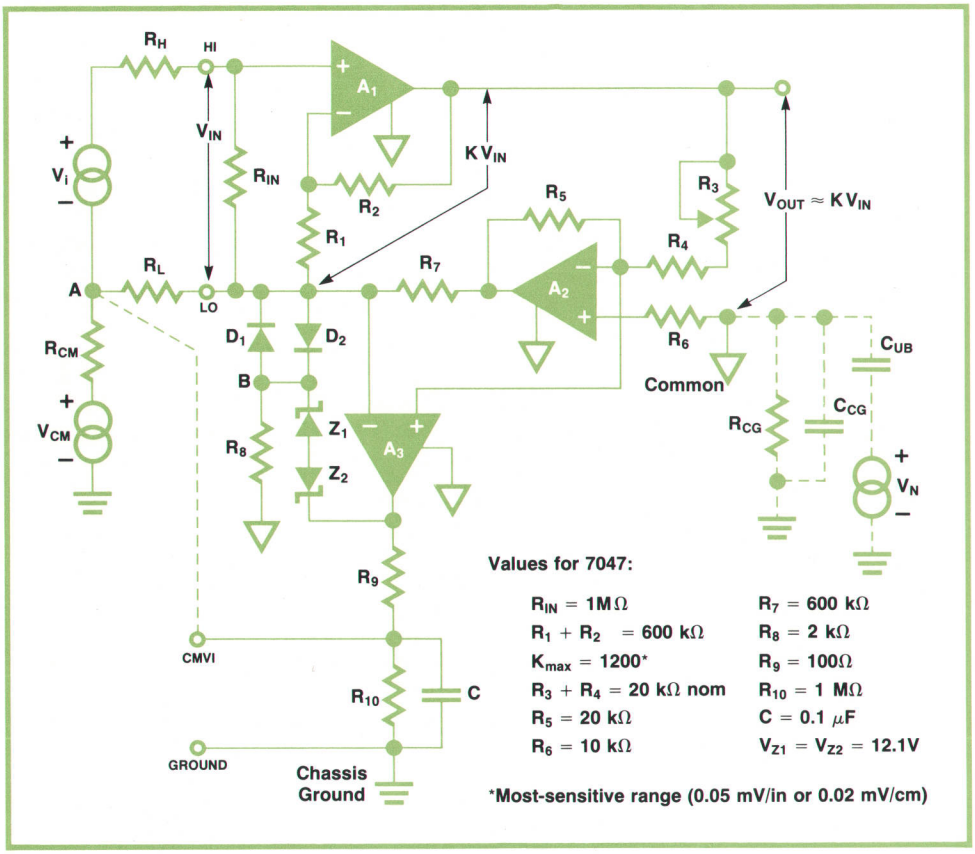

Fig. 6. *Simplified schematic of the 7047A input circuit. A_1 is the equivalent of the input preamplifier and A_3 is the common driver amplifier. The COMMON MODE VOLTAGE INPUT (CMVI) does not have to be connected unless the common-mode voltage exceeds ±10 V.*

HI and LO terminal leads are shown as R_H and R_L, respectively. The common-mode source V_{CM} is shown in series with its source resistance R_{CM}.

The main preamplifier is shown in Fig. 6 as operational amplifier A_1. Although it actually consists of a chopper, an ac amplifier, a demodulator, and an integrator, its overall effect is that of a dc operational amplifier. A_1 provides a non-inverting gain of

$$K = 1 + R_2/R_1$$

for the voltage V_{IN}, which appears across the input HI and LO terminals.

The operational amplifier A_2 inverts the output voltage V_{OUT} so that, ideally, all of the current flowing through R_2 and R_1 because of V_{OUT} has a return path through R_7 and into the output of A_2. The potentiometer R_3 allows a precise adjustment of this current sinking so virtually no current flows out of the LO terminal.

The operational amplifier A_3, the common driver amplifier, has its negative terminal connected to the LO input terminal. It serves to drive the circuit common to the same voltage as the voltage applied to the LO input terminal. (Note that all three operational amplifiers A_1, A_2, and A_3 have their power supplies, and hence their outputs, referenced to circuit common.)

Resistors R_9 and R_{10} and capacitor C are connected in a series-parallel combination between the output of A_3 and chassis ground to provide a shunting path to ground for the transformer noise coupled into the circuit common from the power supply. Thus chassis ground forms part of a negative feedback path for amplifier A_3 through the elements R_9, R_{10}, and C, the common mode voltage source V_{CM} and R_{CM}, and the impedance R_L.

The output of A_3 is also connected to its negative input terminal through an alternative feedback path consisting of the zener diodes Z_1 and Z_2, and the diodes D_1 and D_2. The point B is connected through resistor R_8 to circuit common. This diode-resistor network activates the alternative feedback path before amplifier A_3 reaches the saturation point. The auxiliary common-mode voltage input terminal (CMVI) is connected to the output of A_3 through resistor R_9 to provide an alternate reference voltage (instead of chassis ground) for the output of A_3 when the common-mode voltage exceeds the range of A_3.

In normal operation, the common-mode voltage is less than the voltage range of A_3 and the CMVI terminal is not externally connected. The voltage V_{IN} is amplified by A_1 and its feedback resistors R_1 and R_2. The voltage V_{OUT} at the output of A_1, which is referenced to circuit common, drives the unity-gain amplifier A_2, which inverts this voltage. Because R_7 equals

R_1+R_2, virtually all of the current through R_1 and R_2 is drawn through resistor R_7, and not out of the LO terminal.

Amplifier A_3 serves as a voltage follower and keeps the voltage potential on circuit common equal to the voltage on the LO input terminal. The result of this is that the currents through R_L resulting from V_{CM} and V_N are greatly reduced compared to the case shown in Fig. 5. Consequently, the CMRR with a resistance in the LO terminal is greatly improved and the pump-out noise current is virtually eliminated (see Appendix I).

The current through R_H because of V_{CM} is also considerably reduced because circuit common is moving with respect to ground along with this common-mode voltage. Therefore the CMRR with a resistance unbalance in the HI terminal is improved significantly (for the 7047A, at least an order of magnitude over the earlier 7045A, which is conventionally floated).

The main limitations on CMRR are imperfect shielding between the HI and LO terminals and chassis ground, and finite gain in amplifier A_3. The input imperfections in the operational amplifiers (offset voltages and bias and offset currents) are not serious with appropriate design and component selection (see Appendix II).

When V_{CM} exceeds the range of A_3 the alternative diode-resistor feedback path will be activated, unless a voltage within ± 10V of V_{CM} is applied at the auxiliary CMVI terminal. For the straightforward situation of Fig. 6, the best connection is shown by the dashed line to point A. In this case A_3 need only supply the voltage difference between the common-mode voltage applied through the LO input terminal and the voltage applied to the CMVI terminal.

If no signal is applied to the CMVI terminal under the above conditions and the alternative feedback path is activated, the circuit still functions but it loses its ability to shunt to ground the current generated by the noise voltage V_N, and this flows out of the LO input terminal as in Fig. 5. Also, the CMRR is reduced, especially for the case of the 1-kΩ resistance in series with the LO terminal.

The alternative feedback path for A_3 is especially necessary to keep it from saturating when the signal source V_i is floating (e.g., a battery) and there is no common-mode voltage source. The absence of V_{CM} (with $R_{CM} = \infty$) disrupts the normal feedback path through chassis ground. Without a negative feedback path amplifier A_3 would become an open-loop amplifier and would lose its ability to reduce the voltage difference between the LO terminal and circuit common to near zero. However, before A_3 reaches its saturation point, its output voltage reaches the turn-on voltage of zener diodes Z_1 and Z_2. The output current of A_3 produces a voltage difference across R_8, which in turn produces a voltage difference across the signal diodes D_1 and D_2 sufficient to forward bias one of them, depending on the polarity of the biasing voltage. Thus the alternative negative feedback path is activated to keep amplifier A_3 from saturating.

Other 7047A Features

The front-panel POLARITY switches, one for each axis, determine whether positive signals drive the recorder right-to-left or left-to-right on the X-axis, and down-to-up or up-to-down on the Y-axis. Thus the user can drive the recorder in a predetermined direction regardless of the polarity of the input signal and without having to exchange input leads. If he cannot drive the CMVI terminal with a common-mode voltage exceeding 10V, he can use the POLARITY switch along with exchanging the input leads to maximize CMR, provided there is no resistance in series with the LO terminal. The CMR performance will then approximate that of the 7045A Recorder (i.e., 110 dB at dc and 90 dB at line frequency for a 1-kΩ resistance in series with the HI terminal).

The front-panel FILTER switches, one for each ax-

SPECIFICATIONS

HP Model 7047A High Sensitivity X-Y Recorder

TYPE OF INPUT: Front input only. Floating, guarded. Employs a unique common driver circuit that eliminates need for connecting CMV to recorder, if CMV ≤10V peak.

INPUT RANGES: 0.05, 0.1, 0.5, 1, 5, 10 mV/in.; 0.05, 0.1, 0.5, 1, 5, 10 V/in. (metric calibration available in 0.02, 0.05, 0.1, 0.5, 1, 5 mV/cm; 0.01, 0.05, 0.1, 0.5, 1, 5 V/cm). Continuous vernier between ranges.

INPUT RESISTANCE: 1 megohm constant on all ranges.

SOURCE RESISTANCE: 10k ohm maximum on all ranges except .05 mV/in. and .1 mV/in. (.02 mV/cm, .05 mV/cm, and .1 mV/cm) where it is 2k ohm maximum.

ACCURACY: ±0.2% of full scale (includes linearity and deadband) at 25°C. Temp coefficient ±0.01% per °C.

RANGE ACCURACY: ±0.2% of full scale ±0.2% of deflection (includes linearity and deadband) at 25°C. Temp coefficient ±0.01% per °C.

DEADBAND: 0.1% of full scale.

COMMON MODE REJECTION: 130 dB dc and 130 dB ac with 1k ohm imbalance in either the high or low terminal (exceeds 150 dB under normal lab conditions). CMR decreases 20 dB per decade step in attenuation from the most sensitive range.

NORMAL MODE REJECTION: 30 dB minimum at line frequency with FILTER IN. (50 dB typical at 60 Hz and 40 dB typical at 50 Hz.)

SLEWING SPEED: 30 in/sec (76 cm/sec) minimum. 38 in/sec (97 cm/sec) typical under normal lab conditions.

ACCELERATION PEAK:
Y-AXIS: 3000 in/sec², (7620 cm/sec²)
X-AXIS: 2000 in/sec², (5080 cm/sec²)

OVERSHOOT: 1% of full scale maximum.

ZERO OFFSET: Eleven calibrated scales of zero offset in both axes. Switchable in steps of one full scale from +1 to −10.

OFFSET ACCURACY AT 25°C (APPLIES TO CALIBRATED ONLY): ±0.1% of full scale times N where N = number of scales of offset.

OFFSET TEMPERATURE COEFFICIENT: ±0.004% of full scale times N per °C.

TIME BASE: Standard: 6 speeds; 0.5, 1, 5, 10, 50, 100 sec/in (metric calibration .1, .5, 1, 5, 10, 50 sec/cm) switchable into X or Y axis.

TIME BASE ACCURACY: 1.0% at 25°C. Temp coefficient ±0.1%/°C.

POWER: 115 or 230 Vac, ±10%, 48 to 66 Hz, 180 VA

PEN LIFT: Electric (remote via TTL level)

WRITING AREA: 10 × 15 in (25 cm × 38 cm)

WEIGHT: Net 41 lb (18,6 kg); shipping 53 lb (24 kg).

PRICE IN U.S.A.: $2850. Event Marker Option 002, $85.

MANUFACTURING DIVISION: SAN DIEGO DIVISION
16399 West Bernardo Drive
San Diego, California 92127 U.S.A.

APPENDIX I
Analytical Expressions For Common-Mode Rejection Ratio and Noise Pump-Out Current

The common driver amplifier, A_3 in Fig. 6, is shown in the diagram below with gain A. Also shown are the sources of the unwanted current I_L, the common-mode voltage V_{CM}, and V'_N, which is the Thévenin equivalent of the noise source V_N in Fig. 6, that is,

$$V'_N = \frac{j\omega C_{UB} Z_{CG}}{1+j\omega C_{UB} Z_{CG}} V_N \tag{1}$$

where

$$Z_{CG} = \frac{R_{CG}}{1+j\omega C_{CG} R_{CG}} \tag{2}$$

also from Fig. 6. The Thévenin equivalent impedance is

$$Z'_{CG} = \frac{R_{CG}}{1+j\omega C'_{CG} R_{CG}} \tag{3}$$

where

$$C'_{CG} = C_{CG} + C_{UB} \tag{4}$$

The output impedance Z_o is the parallel combination of R_o and C_o, which are, respectively, R_{10} and C in Fig. 6. The input resistance R_i is equivalent to the series combination of R_1 and R_2 in parallel with R_7. All components obviously insignificant to the analysis are omitted.

We can now write the loop equations

$$V_{CM} + A(V_C - V_L) - I_L(R_{CM} + R_i + Z_o) - I_N Z_o = 0 \tag{5}$$

and

$$V'_N - I_N(Z'_{CG} + Z_o) + A(V_C - V_L) - I_L Z_o = 0 \tag{6}$$

Since

$$V_C - V_L = -I_L R_i \tag{7}$$

equations (5) and (6) become

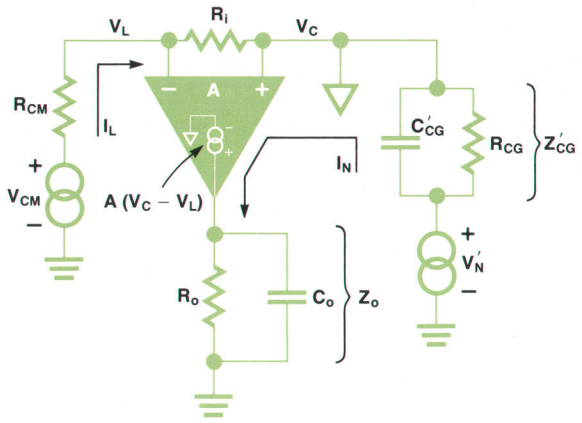

$$V_{CM} - I_L[R_{CM} + (1+A)R_i + Z_o] - I_N Z_o = 0 \tag{8}$$

and

$$V'_N - I_N(Z'_{CG} + Z_o) - I_L(Z_o + AR_i) = 0 \tag{9}$$

Solving for I_N in equation (9) we have

$$I_N = \frac{V'_N - I_L(Z_o + AR_i)}{Z'_{CG} + Z_o} \tag{10}$$

Substituting (10) into (8) we obtain

$$I_L = \frac{V_{CM} - V'_N \frac{Z_o}{Z'_{CG}+Z_o}}{R_{CM} + Z_o + (1+A)R_i - \frac{Z_o(Z_o + AR_i)}{Z'_{CG}+Z_o}} \tag{11}$$

Under all normal operating conditions it can be assumed that

$$|Z'_{CG}| >> |Z_o| \tag{12}$$

$$|A| >> 1 \tag{13}$$

$$|AR_i| >> |R_{CM} + Z_o| \tag{14}$$

Hence

$$I_L \approx \frac{V_{CM}}{AR_i} - \frac{V'_N Z_o}{AR_i Z'_{CG}} \tag{15}$$

Since from (1), (2), (3), and (4)

$$\frac{V'_N}{Z'_{CG}} = j\omega C_{UB} V_N \tag{16}$$

we have

$$I_L \approx \frac{V_{CM}}{AR_i} - j\omega C_{UB} V_N \frac{Z_o}{AR_i} \tag{17}$$

The first term in equation (17) represents the current resulting from the common-mode voltage V_{CM}. Therefore, for a resistance R_L in series with the LO terminal,

$$CMRR \approx \frac{V_{CM}}{\frac{V_{CM} R_L}{|A|R_i}} = \frac{|A|R_i}{R_L} \tag{18}$$

For R_L = 1 kΩ and R_i = 300 kΩ,

$$CMRR \approx 300|A| \tag{19}$$

Using design values from the 7047A we get the following approximate values for CMRR:

2.0×10^9 or 186 dB, dc nominal
1.6×10^8 or 164 dB, dc minimum
6.0×10^7 or 155 dB, 60 Hz nominal
4.7×10^6 or 133 dB, 60 Hz minimum.

These values agree closely with measured values for the 7047A.

The second term in equation (17) represents the current resulting from the power transformer noise current. This is reduced by a factor of AR_i/Z_o, which for the 7047A is nominally 2×10^6, absolute value. With a 10-kΩ source resistance and the HI terminal connected to ground, this source of pump-out current was not measurable on the 7047A.

is, provide more than 20 dB additional normal-mode rejection at line frequency when the filter is IN. Some degradation of rise time occurs.

Acknowledgments

Hal Beach performed much of the testing and proved the feasibility for the common driver circuit when it was initially developed. He also was instrumental in designing the various protection schemes for the common driver amplifier. His insights on low-level preamplifiers were particularly valuable for the development of the 7047A.

Tom Daniels, group leader for X-Y Recorders, was instrumental in seeing that the common driver scheme was used on the 7047A, and had the insight to recognize its significance early in its development.

Finally, Tom Werts performed some very inven-

tive product design in developing the input boxes to house the preamplifier, common driver, and all those terminals, switches, and potentiometers. He did this while minimizing manufacturing cost in various ingenious ways and allowing for relatively easy testing and maintenance of the circuitry.

APPENDIX II
Effects of Amplifier Offsets

The common driver amplifier A_3 from Fig. 6 is shown in the diagram below with its associated resistors and an external source resistance of R_L. The HI terminal is grounded, since this represents a more significant case than when the LO terminal is grounded. The offset voltage of A_3 is shown as a variable battery v_{03} (an internal potentiometer is used to adjust its equivalent value). The bias and offset currents of A_3 may be neglected because a low-leakage dual-input FET is used. The sum of the offset voltage and the product of offset current and drift compensation resistance R_6 for amplifier A_2 in Fig. 6 is shown as the fixed battery v_{02}. v_{01} is the equivalent offset referred to the input of the chopper amplifier A_1. V_0 is the sum of the series batteries v_{02} and v_{03}, and $-V_0$ is the voltage with respect to circuit common at the negative input of A_3. V_1 and V_2 are the voltages with respect to circuit common at the outputs of A_1 and A_2, respectively (see Fig. 6). Since a dc analysis is being done, all capacitors are ignored.

Assuming A_3 to now be an ideal operational amplifier (i.e., infinite gain, infinite input impedance, and zero offset), we can use superposition to write

$$\begin{aligned} I &= V_0/300k\Omega + V_1/600k\Omega + V_2/600k\Omega \\ &= V_0/300k\Omega + V_1/600k\Omega - V_1/600k\Omega - v_{02}/300k\Omega \\ &= V_0/300k\Omega - v_{02}/300k\Omega \\ &= (v_{02} + v_{03})/300k\Omega - v_{02}/300k\Omega \\ &= v_{03}/300k\Omega \end{aligned}$$

Hence, by adjusting only the offset voltage v_{03} of the common driver amplifier A_3 to zero, the pump-out current I due to offsets can be eliminated. The effect of drift in amplifier A_2 on pump-out current I is eliminated by the algebraic cancellation of the terms containing v_{02}. This is the result of connecting the + terminal of amplifier A_3 to the − terminal of amplifier A_2 (see Fig. 6), a connection whose advantages are not obvious at first glance.

The offset voltage at the output of the preamplifier is given by

$$V_1 = -Kv_{01} - v_{02} - v_{03}$$

For the 7047A design the Kv_{01} term is only significant on the most sensitive ranges when K is large. Hence, only $v_{02} + v_{03}$ is compensated by adding a series equivalent low-impedance fixed voltage on a less-sensitive range. Kv_{01}, which comes into play only on the most sensitive ranges, is tolerated, and typically is on the order of a few millivolts or an offset of about 0.1 cm at the pen tip on the 0.02 mV/cm range (the sensitivity at the output of the preamplifier is 24 mV/cm on all ranges).

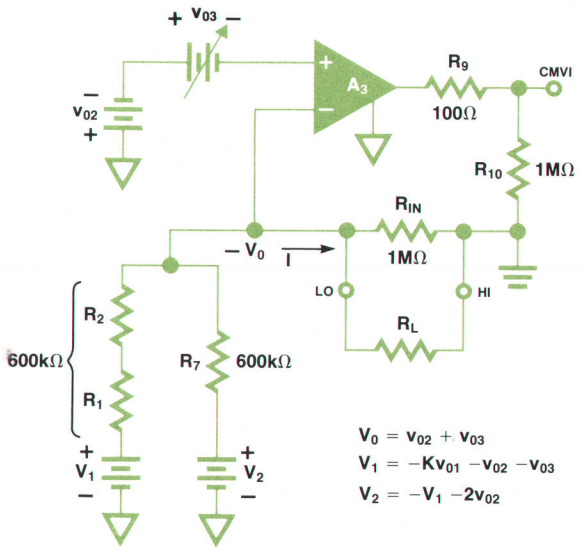

Reference

1. "Floating Measurements and Guarding," Application Note 123, Hewlett-Packard Co., 1970.

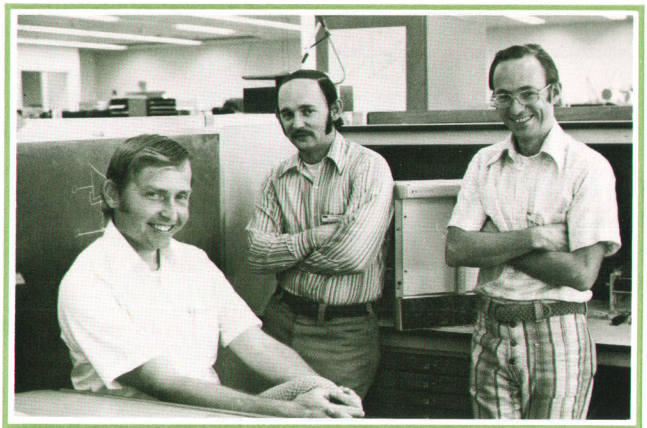

John M. Wade (left)
John Wade has been with HP since 1961, the year he graduated from the University of California at Los Angeles with a BSEE degree. He's served as design engineer and project leader for recorders and accessories, group leader for X-Y and strip chart recorders, and San Diego Division quality assurance manager, his present position. He's co-author of the patent application on the common driver circuit used in the 7047A Recorder. Born in Des Moines, Iowa, John is married, has three children, and lives in Poway, California, near San Diego. He's active in church work and relaxes by restoring antique automobiles.

Donald W. Huff (right)
Don Huff received his BSEE degree from California State Polytechnic College at Pomona in 1965, then joined HP's Loveland, Colorado Division, where he did product and circuit design. In 1969 he received his MSEE degree from Colorado State University, and moved to San Diego, where he subsequently became 7130A Strip Chart Recorder project leader. Named group leader for strip chart recorders in 1973, he's now group leader for instrumentation tape recorders. Don is a co-author of the patent application on the common driver circuit used in the 7047A Recorder. Originally from Massachusetts, he's married, has two children, and lives in Poway, California. He's active in his church and lists amateur radio, photography, and reading as leisure-time activities.

Daniel E. Johnson (center)
Dan Johnson, project leader for the 7047A High-Sensitivity X-Y Recorder, joined HP's San Diego Division in 1970 with five years' experience in circuit design and analysis. He helped design the 7046A Recorder and was project leader for the 7041A Recorder before taking on the 7047A. His BSEE degree is from Lafayette College (1965) and his MSEE is from the Polytechnic Institute of New York (1967). Dan is a member of IEEE. Among his non-electronic interests languages and sports predominate. He's fluent in Spanish and English, has some knowledge of German, and is studying Portugese. He plays softball and is a fan of baseball, football, basketball, horse racing, and jai-alai. He's also found time to gain a junior-college teaching certificate in electronics and be active in mental health volunteer work. He and his wife, a native of Mexico, are expecting their first child early in 1975.

Digital High-Capacitance Measurements to One Farad

Here's a new high-C meter capable of measuring today's wide-value electrolytic and tantalum capacitors. It's for production testing, incoming inspection, and the laboratory.

by Kunihisa Osada and Jun-ichi Suehiro

HEWLETT-PACKARD'S MODEL 4350A/B High-Capacitance Meter[1], first produced in 1971, had a wide measurement range (1μF to 300 mF) and simultaneously measured capacitance and dissipation factor, features that were considered advanced at the time. Since 1971, however, the capacitor industry has made considerable progress. Manufacturers of high-value electrolytic capacitors, used in computer power supplies and many other applications, now routinely produce capacitors as large as one farad. Tantalum capacitors are now produced in an expanded range of values (0.1μF to 1000μF) with smaller dissipation factors.

The new Model 4282A Digital High-Capacitance Meter, Fig. 1, has a measurement range wide enough and accuracy high enough to satisfy the present measurement needs of capacitor makers and users, and then some. Its four-digit readout, digital output, and remote programmability are expected to make it useful in production test and incoming inspection systems, as well as in the laboratory.

Model 4282A measures capacitance in nine ranges from 10 nanofarads to one farad. Resolution as fine as one picofarad is possible. Basic accuracy in capacitance measurements is 0.4% to 2.5%, depending on range. Dissipation factors as high as 10.00 can be measured. Resolution is 0.001 and basic accuracy is 1.5%.

The new high-C meter can also measure the product of capacitance and equivalent series resistance and display the result in ohm-farads. Capacitance is read on the LED four-digit display; dissipation and ohm-farads are read using three full digits with a fourth for overranging. The reading rate is continuously variable from 0.3 to 2 seconds, and capacitance and dissipation factor (or ohm-farad) measurements can be displayed alternately if desired, thus eliminating the need for resetting panel switches or using two instruments.

Four internal test frequencies are available: 50, 60, 100 and 120 Hz. An internal dc bias supply is continuously variable from 0 to 10 volts.

Fig. 1. *Model 4282A Digital High-Capacitance Meter measures capacitance to one farad, dissipation factor to 10.00, the product of capacitance and equivalent series resistance (ΩF), and dc voltages to 600V. An option adds leakage current measurements.*

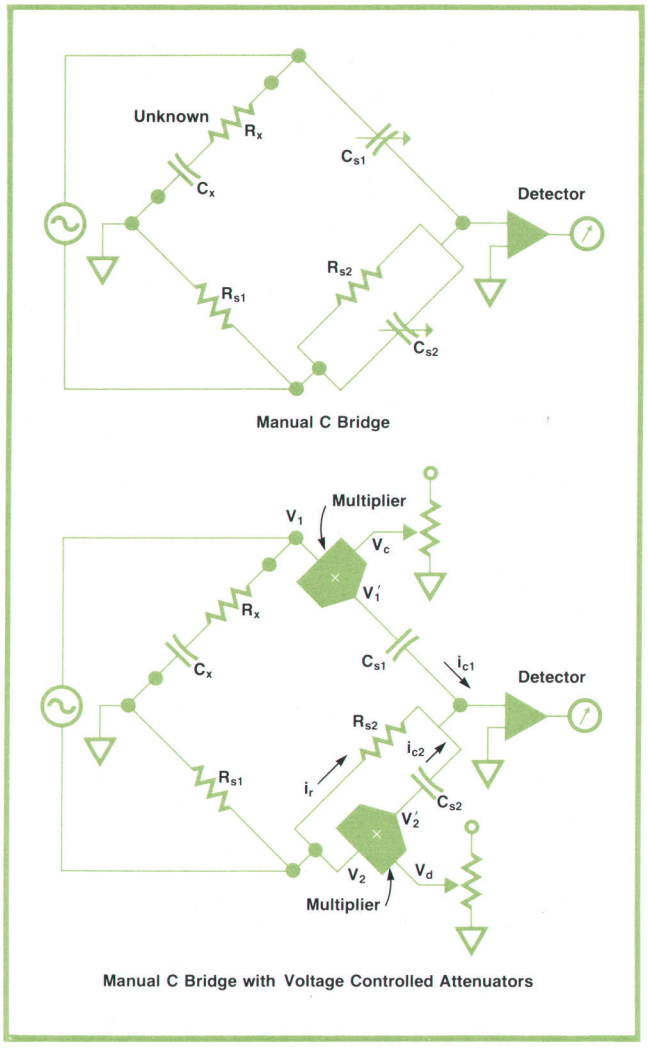

Fig. 2. *A capacitance bridge may be balanced by adjusting standard capacitors, as shown in (a), or by adjusting the gain of a voltage-controlled attenuator, or multiplier, as shown in (b). Model 4282A uses the latter approach, adding a feedback loop so balance is achieved automatically.*

Leakage-current measurement ranges from 1 microampere to 10 milliamperes are added with Option 001. Results are displayed on three digits with 18 percent overranging. Dc bias from 0 to 100 volts and provision for bias from an external source are added to the standard instrument with this option.

Digital and analog output are available on the rear panel. Function, range, and reset are independently remotely programmable by contact closure to ground or TTL/DTL logic levels. I_L with Option 001 is also programmable.

In addition to these high-C-meter capabilities, Model 4282A can be used as a three-digit dc digital voltmeter to measure up to 600 volts.

Bridge Measurement

Because the bridge method is the most accurate way to measure impedance, the 4282A makes bridge measurements. However, the bridge is balanced automatically, so measurements are much faster than with a manual bridge.

Fig. 2a shows a manual bridge for measuring the capacitance and dissipation factor of an unknown. A series equivalent circuit is assumed for the unknown. The bridge is balanced by adjusting standard capacitors C_{s1} and C_{s2}. At balance,

$$C_x = \frac{R_{s2}}{R_{s1}} C_{s1}.$$

For the series equivalent circuit, dissipation factor is defined as

$$D_x = \omega C_x R_x.$$

At balance,

$$D_x = \omega C_{s2} R_{s2}.$$

Shown in Fig. 2b is the same bridge, but with analog multipliers acting as voltage-controlled attenuators replacing the variable standard capacitors. Balance is obtained by varying V_c and V_d. At balance,

$$i_r + i_{c1} + i_{c2} = 0.$$

How the New Meter Works

Fig. 3, the simplified block diagram of the 4282A High-Capacitance Meter, has a basic configuration similar to that of Fig. 2b. The major difference is that in the 4282A the bridge is balanced automatically by means of a negative feedback loop consisting of multipliers, a summing amplifier, synchronous detectors, and integrators.

The test signal from the meter's internal oscillator is applied to the unknown through an AGC amplifier and a power amplifier. A transformer makes it possible to apply dc bias to the unknown.

The voltage across the unknown capacitor is amplified by A_1 in Fig. 3. The voltage across the range resistor R_{s1} is amplified by A_2. The output of amplifier A_1 goes to one of the multipliers, which in turn provides current i_{c1} to the summing amplifier through standard capacitor C_{s1}. The output of A_2 goes to the other multiplier, which sends current i_{c2} to the summing amplifier through standard capacitor C_{s2}. Reference current i_r comes to the summing amplifier directly from the output of A_2 through standard resistor R_{s2}.

The output of the summing amplifier is divided into two parts by the synchronous detectors, which produce dc voltages proportional to the real (in-

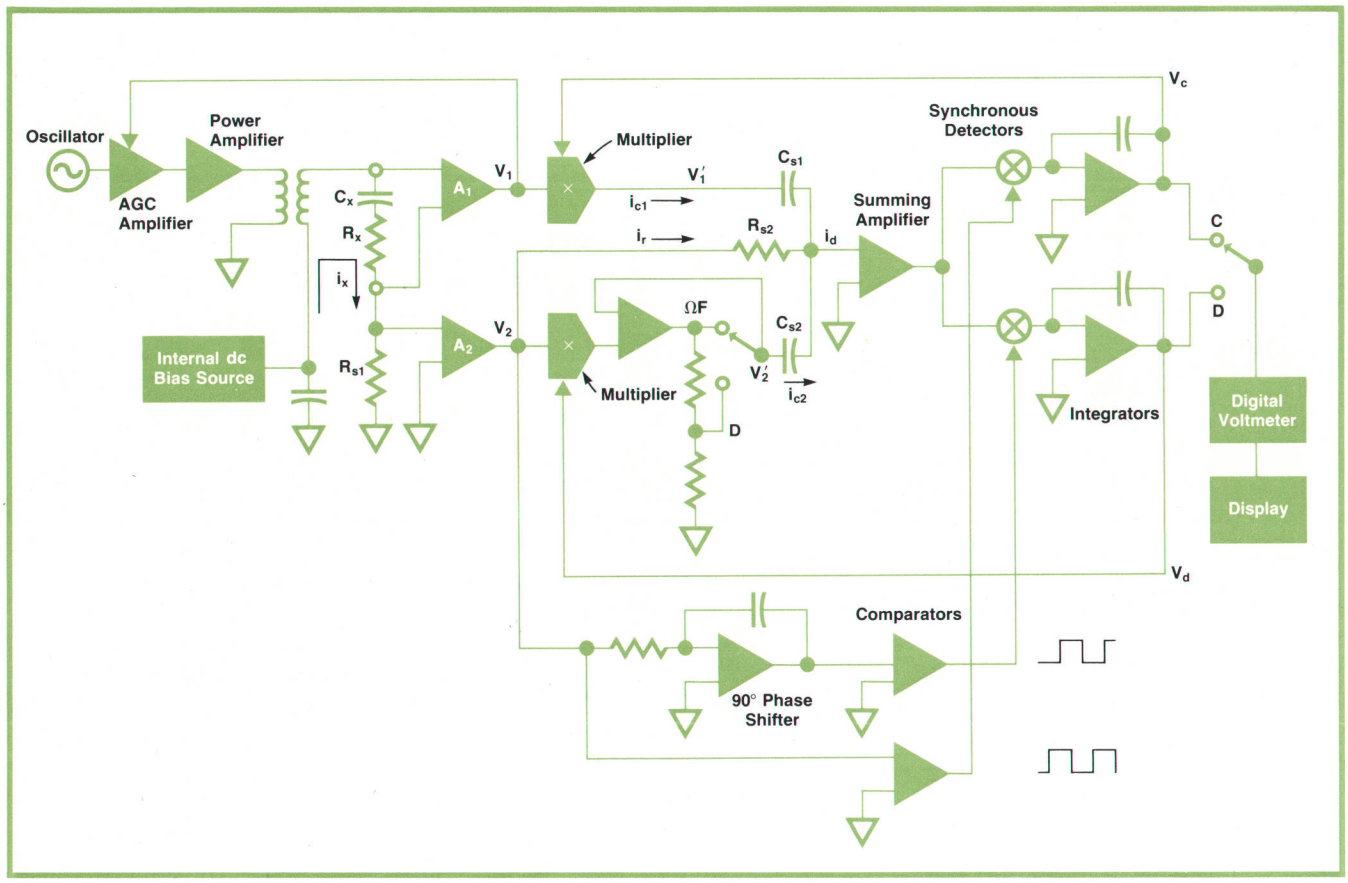

Fig. 3. *Simplified circuit diagram of the 4282A Digital High-Capacitance Meter. The transformer makes it possible to apply dc bias to the unknown along with the ac test signal at 50, 60, 100, or 120 Hz.*

phase) and imaginary (quadrature) parts of the current i_x through the unknown. These voltages are integrated and the integrator outputs are the second inputs to the two multipliers.

The feedback loop acts to set $i_d = 0$. At balance, the input current to both integrators is zero, and the integrator output voltages are V_c and V_d, which are proportional to the unknown capacitance and dissipation factor, as we will now show.

The multiplier input voltages in Fig. 3 are

$$V_1 = -i_x\left(R_x + \frac{1}{j\omega C_x}\right)A_1$$

and

$$V_2 = i_x R_{s1} A_2.$$

The output of each multiplier is proportional to the product of its two inputs:

$$V_1' = V_1 V_c / K_1$$

and

$$V_2' = V_2 V_d / K_2.$$

The currents into the summing amplifier are

$$i_{c1} = j\omega C_{s1} V_1'$$

$$i_{c2} = j\omega C_{s2} V_2'$$

and

$$i_r = V_2 / R_{s2} = i_x R_{s1} A_2 / R_{s2}.$$

When the bridge is balanced,

$$i_d = i_{c1} + i_{c2} + i_r = 0.$$

Substituting the values above into this equation, setting the sums of the real and imaginary parts equal to zero, and solving for the unknowns gives the following relations:

$$C_x = \frac{C_{s1} R_{s2} A_1}{K_1 R_{s1} A_2} V_c$$

and

$$D_x = \omega C_x R_x = \frac{\omega C_{s2} R_{s2}}{K_2} V_d.$$

The built-in digital voltmeter measures V_c and V_d and displays C_x and D_x directly. Ohm-farads are obtained by multiplying the dissipation factor by $1/\omega$.

Four-Terminal Measurement

Model 4282A has nine decade capacitance ranges, from 10 nF full-scale to 1 F. Thus the unknown impedance may range from milliohms to megohms, and the 4282A must be able to cope with the problems of both low-impedance and high-impedance measurements.

A four-terminal configuration, as shown in Fig. 4, is used to eliminate low-impedance measurement errors caused by test-lead resistances (R_l in Fig. 4). Two terminals are used to inject current into the unknown and two other terminals are used to detect the voltage across the unknown. Because lead resistance between the unknown and the range resistor R_{s1} may cause an error, two different terminals are used to detect the voltage across the range resistor. Very-high-input-impedance voltage followers at the voltage-detection terminals make the four-terminal configuration practical for bridge measurements, where the unknown voltage and the reference voltage must be detected simultaneously while keeping the current in the voltage leads negligibly small compared to the current in the current leads.

Stray Capacitance Compensation

Capacitance measurement resolution may be as fine as 1 pF on the lowest range, so stray capacitance and amplifier input capacitance must be considered a source of error. Of the stray capacitances shown in Fig. 5, C_1 causes mainly capacitance errors and C_2 causes mainly dissipation factor errors. C_3 does not affect measurements. Without compensation, capacitance errors might be as large as 20% of full scale on the lowest range.

To eliminate these errors, Model 4282A has the stray-C compensation circuit shown in Fig. 5. Amplifiers A_4 and A_5 send compensation currents i_4 and i_5 through capacitors C_4 and C_5. The result is that

$$i_1 + i_2 + i_4 + i_5 = 0$$

so the current that flows through the unknown capacitor is equal to the current flowing through the range resistor.

For high-voltage protection a pair of zener diodes is connected to each input voltage follower as shown in Fig. 6. In normal operation there is no ac voltage across the diodes that connect the amplifier input terminal to the zener diodes, so these diodes introduce no distortion and no additional stray capacitance.

The Multipliers

The accuracy of the new high-C meter in capacitance and dissipation factor measurements is mainly a function of the linearity of the two multipliers shown in Fig. 3. The multipliers, a kind of voltage-controlled attenuator, are linear within 0.2% or better from 1/10 of full scale to full scale for both inputs.

The multiplier circuit, Fig. 7, consists of a pulse-

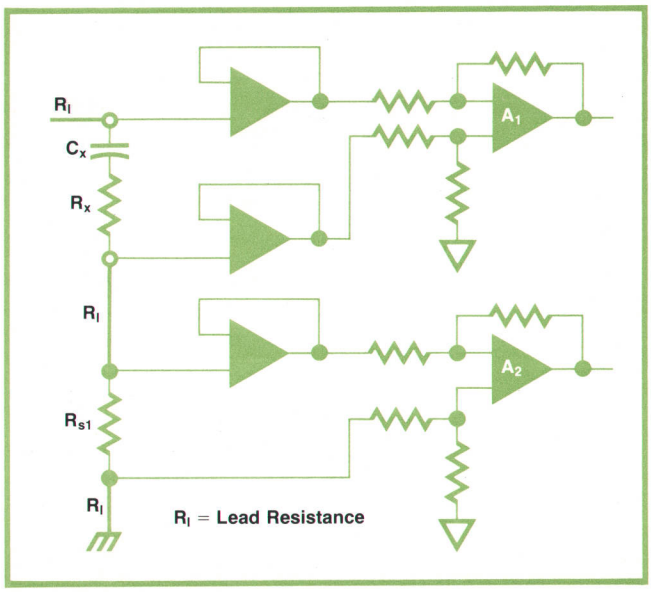

Fig. 4. *A four-terminal configuration with very-high-input-impedance voltage followers in the voltage-sensing leads eliminates errors caused by lead resistances.*

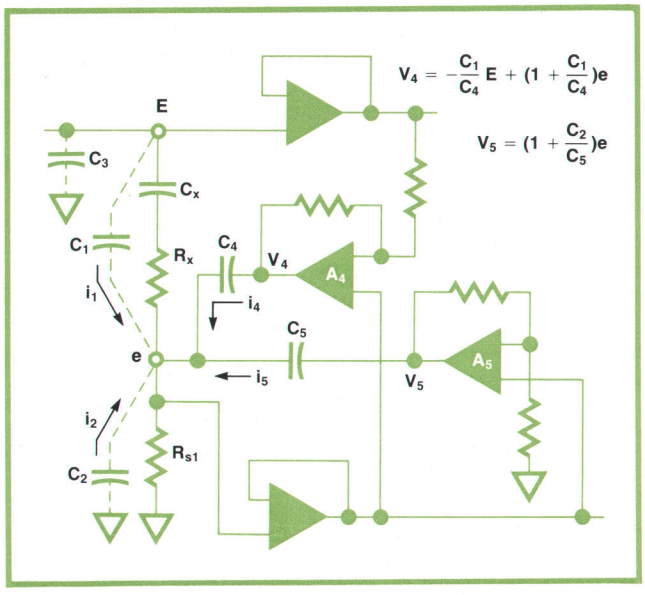

Fig. 5. *Stray-C compensation circuit makes the current through the unknown equal to the current through the range resistor R_{s1}.*

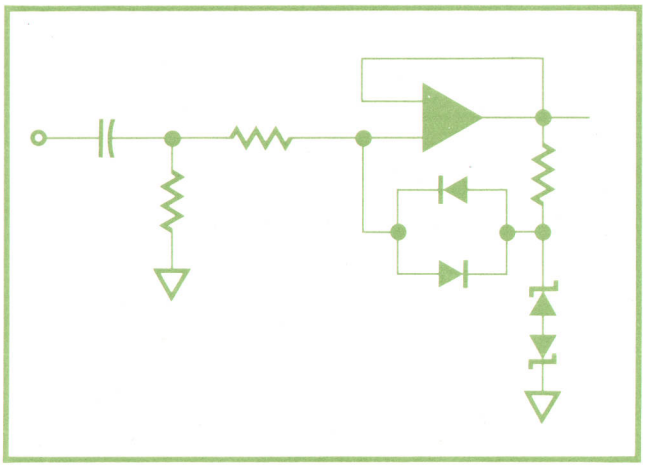

Fig. 6. *For high voltage protection, zener diodes are connected to each input voltage follower. Circuit is designed so that, in normal operation, no ac voltage appears across the diodes, so they introduce no distortion or additional stray capacitance.*

width modulator followed by a demodulator. The duty cycle of the modulator output pulse is proportional to one of the multiplier input voltages, and the gain of the demodulator is proportional to the other multiplier input voltage.

The clock generator in the modulator circuit of Fig. 7 produces a sawtooth wave, which fixes the period of the modulated pulse. The clock frequency is 200 kHz, far higher than the input signal frequency. The input voltage to the zero detector is half the sum of the clock generator output voltage and the integrator output voltage.

Inputs to the integrator come from the variable multiplier input voltage E_a, which corresponds to V_1 or V_2 in Fig. 3, and the fixed reference voltage $\pm E_r$. During time T_1, the zero detector input voltage is positive and switch S_1 connects $+E_r$ to the integrator. During T_2, the zero detector input is negative and S_1 switches to $-E_r$. T_1+T_2 is equal to the clock generator period. At steady state the sum of the

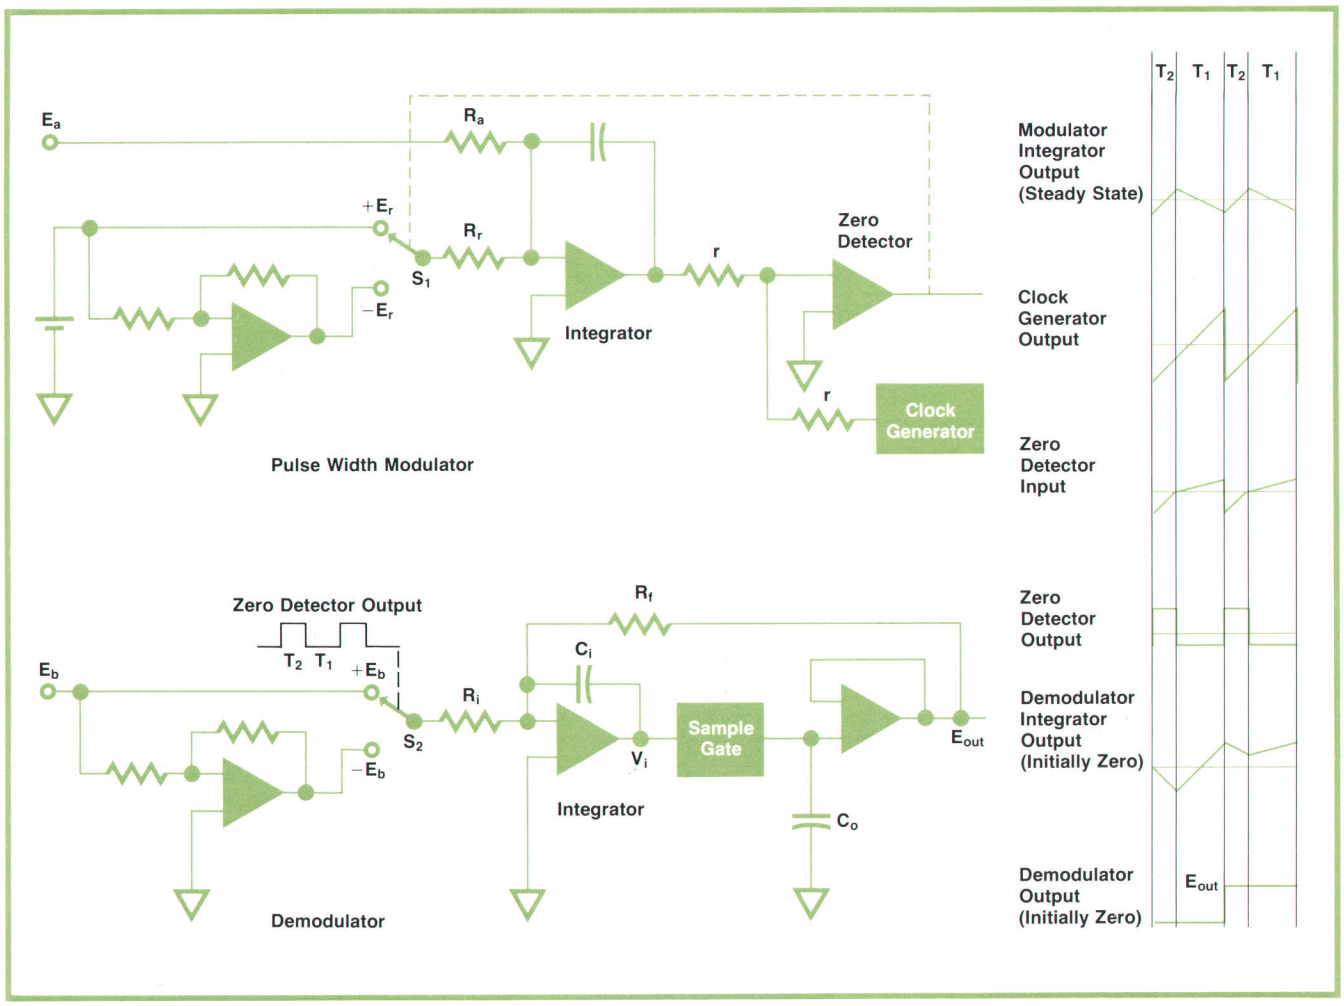

Fig. 7. *Multiplier consists of a pulse width modulator and a demodulator. One multiplier input varies the duty cycle of the modulator output pulse, and the other varies the demodulator gain. Multiplier linearity is better than 0.2% for both inputs.*

integrator input currents over the complete period is zero. Hence

$$\frac{E_a}{R_a}(T_1+T_2) + \frac{E_r}{R_r}T_1 - \frac{E_r}{R_r}T_2 = 0$$

and

$$T_1-T_2 = -\frac{E_aR_r}{E_rR_a}(T_1+T_2).$$

Thus T_1-T_2 is proportional to the input voltage E_a. In this closed-loop modulator, the hysteresis of the zero detector causes no error.

In the demodulator circuit of Fig. 7, switch S_2 is controlled by the zero detector output pulse. During time T_2, S_2 connects the voltage $-E_b$ to the integrator, where E_b is the other multiplier input voltage, corresponding to V_c or V_d in Fig. 3. During time T_1, $+E_b$ is connected to the integrator. The sample gate samples the integrator output voltage V_i at the end of each period (T_1+T_2) and capacitor C_o holds the sampled voltage.

Assuming that the integrator output is zero and the demodulator output E_{out} is also zero at the beginning of the first period, the integrator output at the end of the first period is

$$V_{i1} = -\frac{E_bT_1}{R_iC_i} + \frac{E_bT_2}{R_iC_i}.$$

V_{i1} is sampled and held on capacitor C_o. At the end of the next clock period the integrator output is

$$V_{i2} = V_{i1} - \frac{E_bT_1}{R_iC_i} + \frac{E_bT_2}{R_iC_i} - V_{i1}(T_1+T_2)\frac{1}{R_fC_i}.$$

The values of C_i, R_f, and the clock period are chosen so that

$$1 - (T_1+T_2)\frac{1}{R_fC_i} = 0.$$

Therefore

$$V_{i1} = V_{i2}.$$

This means that if T_1-T_2 is constant, demodulation is completed in one clock period and E_{out} remains constant thereafter, as shown by the waveforms in Fig. 7. The output voltage is

$$E_{out} = -\frac{E_b}{R_iC_i}(T_1-T_2) = \frac{E_aE_bR_r}{R_iC_iR_aE_r}(T_1+T_2)$$

which is proportional to the product of the two multiplier input voltages E_a and E_b.

A simple low-pass filter could have been used instead of the integrator and sample-and-hold circuit of the demodulator, but there would have been time delay and ripple problems. The 4282A approach eliminates these.

Kunihisa Osada

Kunihisa Osada received his BSEE degree from the University of Electronic Communications, Tokyo, in 1968. At Yokogawa-Hewlett-Packard since 1968, he has designed several circuits for the 4350A/B HI-C Meter, the 4332A LCR Meter, and the 4265A Universal Bridge. He designed the analog circuits of the 4282A. Kunihisa is currently serving YHP as a manufacturing engineer. His favorite leisure-time activities include taking 8mm movies and skiing with his wife.

Jun-ichi Suehiro

Jun-ichi Suehiro graduated from Kagoshima Technical College in 1970. At Yokogawa-Hewlett-Packard since 1970, he has designed several circuits for the 4350A High Capacitance Meter and the 4332A LCR Meter. He developed the measurement section circuitry for the 4282A Hi-C Meter. He is a member of the IEE of Japan. In his spare time, Jun-ichi likes to swim and to listen to classical music.

SPECIFICATIONS

HP Model 4282A Digital High-Capacitance Meter

MEASURING RANGES:

CAPACITANCE: 10.000nF to 1.0000F, four full digits. 9 ranges in decade steps, manual selection, 18% overranging.

DISSIPATION FACTOR: 1.000 to 10.00, three full digits, 2 ranges, auto selection, 18% overranging.

OHM FARAD: 1.000ΩmF to 10.00ΩmF three full digits, 2 ranges, auto selection, 18% overranging.

DC VOLTAGE: 10.00V to 1.000kV, three full digits. 3 ranges, in decade steps, manual selection (maximum voltage is 600V), 18% overranging.

MEASURING CIRCUIT: Series equivalent circuit using four-terminal method.

MEASURING FREQUENCIES: 50Hz, 60Hz, 100Hz and 120Hz (50Hz and 60Hz synchronized by line frequency). Accuracy: ±1.5%.

MEASURING VOLTAGES: 10nF to 10mF ranges: <1Vrms.; 100mF range: <0.1Vrms.; 1F range: <10mVrms.

ACCURACY (+23°C ±5°C after half hour warm up): ±(% of reading +% of full scale).

CAPACITANCE:

C Range	% of reading	% of full-scale
10nF	1.0 + 0.9 · D rdg	0.2
100nF	0.5 + 0.5 · D rdg	0.1
1μF to 1mF	0.4 + 0.5 · D rdg	0.05
10mF	1.0 + 0.5 · D rdg	0.05
100mF	1.5 + 0.5 · D rdg	0.5
1F	2.5 + 0.5 · D rdg	1.0

DISSIPATION FACTOR:

C Range	% of reading	% of full-scale
10nF	1.5 + 0.5 · D rdg	0.2 · Cfs/C rdg + 0.3
100nF to 1mF	1.5 + 0.2 · D rdg	0.2 · Cfs/C rdg + 0.3
10mF	1.5 + 0.2 · D rdg	0.2 · Cfs/C rdg + 0.5
100mF, 1F	1.5 + 0.2 · D rdg	0.2 · Cfs/C rdg + 3

OHM-FARAD:

C Range	% of reading	% of full-scale
10nF	1.0 + 0.5 · ΩF rdg	0.2 · Cfs/C rdg + 0.3
100nF to 1mF	1.0 + 0.2 · ΩF rdg	0.2 · Cfs/C rdg + 0.3
10mF	1.0 + 0.2 · ΩF rdg	0.2 · Cfs/C rdg + 0.5
100mF, 1F	1.0 + 0.2 · ΩF rdg	0.2 · Cfs/C rdg + 3

DC VOLTAGE:

10V range: ±(0.05% of reading + 0.1% of full-scale).

100V and 1kV ranges: ±(0.2% of reading + 0.1% of full-scale).

TEMPERATURE COEFFICIENT: (referred to +23°C, and temperature range of 0°C to 50°C):

C: ±0.02% of reading/°C

D, ΩF: ±0.03% of reading/°C

V: ±0.01% of reading/°C

Option 001 Leakage Current Measurement

RANGE: 1.000μA to 10.00mA, 5 ranges, three full digits.

OVERRANGING: 18%

ACCURACY:

1μA range: ±(2% of reading +2.0% of full-scale).

10μA to 10mA ranges: ±(2% of reading +0.3% of full-scale).

BIAS VOLTAGES:

INTERNAL SOURCE: 0 to 10V, 0 to 100V, 2 ranges, continuously variable over each range.

MAXIMUM CURRENT: 100mA for 10V range and 60mA (for 1 minute) for 100V range.

EXTERNAL SOURCE: Useable up to 600Vdc across EXT BIAS terminals on rear panel.

PROTECTIVE RESISTOR: 1kΩ for 100V range and for external bias, 1Ω for 10V range.

General

DC BIAS VOLTAGE: 0 to 10V, continuously adjustable with DC BIAS control. Maximum charging current is 100mA.

BALANCING TIME: Normally one second (when measuring on C ranges of 10nF through 10mF, capacitance value near full-scale, dissipation factor less than one and without dc bias).

READING RATE: Continuously variable from 0.3 seconds to 2 seconds with RATE control.

RESET: Initiates one reading by depressing RESET INT pushbutton or contact closure to ground or TTL low level at RESET EXT line.

DIGITAL OUTPUT:

OUTPUT SIGNALS: BCD +1-2-4-8, data parallel, decimal point, function and unit, overload and unbalance, and polarity.

LEVEL: Low: 0.3V ± 0.3V, Max. sink current 15mA. High: 3.9V ± 1.5V, Max. load current 300μA.

PRINT COMMAND OUTPUT: Negative going TTL pulse of approx. 1ms.

PRINTER HOLD INPUT: TTL low level or contact closure to ground.

REMOTE PROGRAMMING: All "FUNCTIONS"; "C Range"; "I_L Range" (option 001) and Reset by TTL low level or contact closure to ground.

ANALOG OUTPUT: Dc output of 1V full-scale in proportion to displayed value. ACCURACY: Add ±0.5% of reading to accuracy specification.

OPERATING ENVIRONMENT: 0°C to +50°C, <90% RH.

POWER REQUIREMENTS: 100V, 120V, 220V or 240V ±10%, 50Hz or 60Hz, approx. 70VA.

DIMENSIONS: 425 W × 88 H × 467 D mm.

WEIGHT: Net 8.8 kg.

PRICES IN U.S.A.:

HP 4282A, $3500

HP 4282 with Option 001, $3750

MANUFACTURING DIVISION: YOKOGAWA-HEWLETT-PACKARD LTD.
9-1, Takakura-cho
Hachioji-shi
Tokyo 192 Japan

Acknowledgments

The authors wish to express their thanks to the many people who participated in the development of the Model 4282A Digital High Capacitance Meter. To Yoshihisa Kameoka, who was the first project leader of the 4282A, we owe special gratitude for his formative work. Hisao Nakae and Shinichi Sano, both electronic engineers, made significant contributions. Shiro Kito, our section manager, provided day-to-day encouragement and gave us many helpful suggestions. Our other project members included Yohichi Matsuzaki, a mechanical engineer, who solved mechanical problems and did the layout, and Motoji Suzuki and Hiroshi Sakayori, electronic engineers who designed the digital circuit and the leakage-current option.

Computer Performance Improvement by Measurement and Microprogramming

The speed of a computer-based Fourier analyzer was increased by a factor of ten by creating several new machine language instructions using firmware. Areas to be microprogrammed were selected on the basis of performance measurements.

by David C. Snyder

MICROPROGRAMMING, or microcode, is a comparatively old technique for computer design. It was first described by Wilkes in 1951, when most logic gates were vacuum tubes and the transistor was only a three-year-old infant. The original objective was "to provide a systematic alternative to the somewhat *ad hoc* procedure used for designing the control system" of a computer.[1,2]

The control system of a computer implements its basic machine-language instruction set. In many computers (e.g., the 1967 vintage HP 2116), this system is a combination of random logic, complicated clocking, and special-purpose data paths. In a microcoded computer, a set of primitive operations (gate register R3 to bus, gate bus to register R1, issue ADD command, etc.) is first identified. One then builds a primitive machine and microprograms the primitive machine to emulate the desired machine.

Microcoding introduces a degree of freedom into the design process by breaking one problem, defining the desired machine, into two relatively noninteractive ones: defining a set of primitives (using hardware), and microprogramming these primitives to emulate the desired machine.

As the years passed, memory costs (both ROM and RAM) dropped, memory speeds rose, and more designs used microcode. Today one finds microcode used in large machines (e.g., most models of the IBM 360 and 370), minicomputers (e.g., Hewlett-Packard 2100 and 21MX, Microdata Model 1600), microcomputers (e.g., National GPCP and Intel MCS-4, but not Intel MCS-8), and some calculators (e.g., HP 35, HP 9830). The benefits to the designer include reduced design time and increased hardware product life. The benefits to the user include the ability to increase performance greatly through the use of low-cost, microcoded instructions tailored to the user's own problems.

Where Microcode Fits

Machine language is the lowest user-programmable level in most computers. FORTRAN or ALGOL programs must be translated into this language by FORTRAN or ALGOL compilers before the programs can run. BASIC and other interpretive languages are implemented by machine language programs. Assembly language is translated into machine language by an assembler, which replaces instruction mnemonics with numerical opcodes. If one were to stop a computer and look in its memory, one would find machine language programs and their data (numbers, tables, lists, etc.), nothing else.

All stored-program computers work in the following manner: a program address register (P) contains the address of the next machine language instruction. That instruction is fetched, P is incremented to the next instruction address, and the fetched instruction is executed. This pattern then repeats. In a machine that is not microcoded, these phases are implemented via complicated clocking, discretionary wiring, and control gates arranged in a seemingly random pattern. In a microcoded machine, the execute phase—and sometimes fetch, interrupt, and direct memory access—is implemented by a microprogram that runs in a simpler, more regular microprocessor. The advantage of such a machine hierarchy—implementing the "2116 language machine" by a microprogram stored in a "2100 micromachine"—is that the hardware is simpler, hence more reliable, and considerably more flexible: changing the visible machine requires microcode changes or additions, not rewiring, layout, new parts, and so on.

In appearance, the HP 2100 is a rather conventional minicomputer (Clark Kent appeared to be a rather conventional reporter). It may have up to 32K words of nonvolatile memory and many I/O devices. Its machine language is identical to that of the earlier

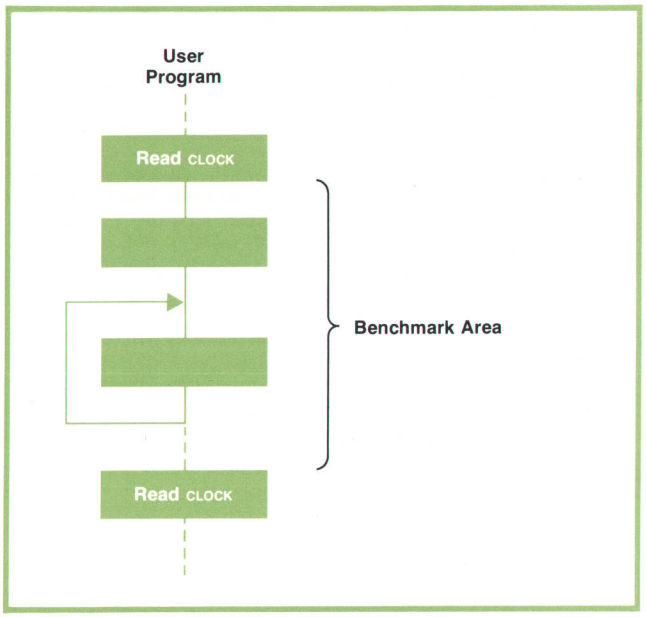

Fig. 1. *Analyzing system performance is the first step in deciding where to use microcode. A common method is to use a software-programmable clock to time various operations. This method requires considerable intuition.*

HP 2116 and its operation speeds are typical of most minicomputers. What sets it apart from most other minicomputers, however, is that it is a microcoded machine, and the user is encouraged to take advantage of this microcode capability.[3,4,5]

The importance of microcode to the user is that it makes possible significant performance increases in many situations. Performance may mean speed to one user, accuracy to another, and a special data format to a third. For example, the HP 2100's Fast FORTRAN Package (FFP), a microcode package implemented in read-only memory (ROM), yields a speed increase of more than 30 times in extended-precision floating-point operations. Along with this speed increase, the user gets a significant increase (approximately 900 words) in available memory because operations previously performed by machine language subroutines are replaced by microcoded instructions implemented in ROM. This kind of architectural extension requires no hardware extensions such as extended arithmetic units or special-purpose processors.

Microprograms are stored in some form of relatively permanent memory. ROMs, PROMs, and RAMs are used, in addition to the somewhat less structured programmable logic arrays (PLAs). When RAMs are used, they are sometimes called writable control store (WCS). This special memory is generally several times faster than the computer's main memory, where machine language programs and data reside. It is also usually wider (some machines use 200-bit-wide microinstructions) to allow for substantial parallelism of operations.

In the HP 2100, microcode is organized into four 256-word modules. Each word is 24 bits wide. Module 0 defines the base instruction set. It is implemented in ROM and is always present. Modules 1, 2, and 3 are for extensions to the base set. They may be absent, or may be implemented in ROM, PROM, or WCS. Although all run at 196 nanoseconds per microinstruction, WCS has a speed advantage since its microprograms may be dynamically tuned for speed. However, it is volatile and therefore requires copies of its microprograms in the computer's main memory.

The speed advantage of microcode in the 2100 comes from four sources:

The fast microprocessor: 196 ns/microinstruction is five million microinstructions per second.

Storage of microcoded programs outside of main memory. This eliminates competition of programs and data for main memory cycles and thus doubles the speed of most instructions.

Provision for parallel operation of the computing resources—for example, calculation during a memory access.

The ability to create special-purpose assembly language instructions and thereby overcome performance compromises that may have been made when the basic instruction set was designed.

Analyzing Where to Microcode

It is easy to be misled by intuition in the area of tim-

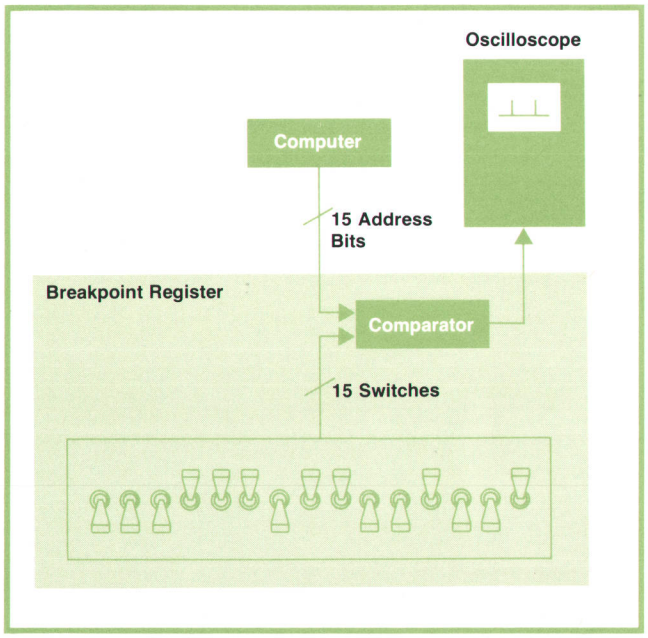

Fig. 2. *A useful tool for timing analysis is a hardware breakpoint register, which emits a pulse whenever the computer program accesses the memory location set into the switch register. The comparator output may be displayed on an oscilloscope or be fed to a computer, even the one being measured.*

ing. Quite often our feelings about where computer programs spend their time are wrong, primarily because we ask the wrong questions. It is important to know if a speedup by a factor of five in a particular subroutine will effect a speed increase of 5, or 1.5, or 1.0005 at the system level.

Fortunately, there are techniques that allow us to measure system performance, localize areas with highest system performance leverage, and predict reliably what system performance increment to expect from microcode (or assembly language, or simple program reorganization, for that matter). System performance, in this context, means basically the timing associated with user-visible operations. For example, in the eyes of a user processing SONAR echoes in real time, the per-echo processing time is a system characteristic; the timing of subroutines is relevant only in relation to this.

By far the most common technique to measure system timing is to use a hand-held watch. Unfortunately, this method collapses all information about timing into a single number, time. A second common method is to use a software-programmable clock to time various operations (see Fig. 1). Many computers contain programmable clocks with millisecond or microsecond resolution. These require considerable intuition in deciding what to measure and what to change. Also, there is no easy way to predict what the effect of a change will be. For example, the user may recode a large matrix inversion and eigenvalue package much more efficiently, only to find that the system timing really depended on a tiny vector-dot-product routine used by the matrix inversion routine.

An extremely useful tool for timing analysis is a hardware breakpoint register, such as the HP 10676A or the American Asian BPR-2100 (see Figs. 2 and 3). This device emits a pulse every time the address set into its 16 switches matches the M (memory address) register. One can determine the time in microseconds spent in any subroutine by setting its entry address into the switches and timing the pulse separation on any scope. The entry address is accessed via the JSB instruction when the routine is entered and via the JMP,I instruction when the routine is exited. Thus the breakpoint register emits a pulse at the beginning and end of the subroutine. The breakpoint register can also answer such questions as "Is this section of code used?," "Did the DMA completion interrupt happen before or after some external event?," and so on.

Perhaps the most useful tool for timing analysis is the activity profile (see Fig. 4). This is a plot showing the percentage of time spent in each area of memory. It not only localizes timing hot-spots, it measures their intensity.

One of the simplest ways to generate an activity profile is simply to push the HALT switch ten or twenty times while the system is operating and write down the program address register contents each time the machine is halted. Normally some range of P-values will stand out. This, along with a printout of the subroutines loaded into various areas of memory, points to the programs that need attention. It also suggests what speed improvement is possible. For example, if seven out of twenty halts found P in the .FLUN floating-point unpacking subroutine, then there would be speed increase of about 7/20 or 35% if the .FLUN routine were made infinitely fast.

A fancier version of this procedure uses the computer to measure itself. I have used this technique on

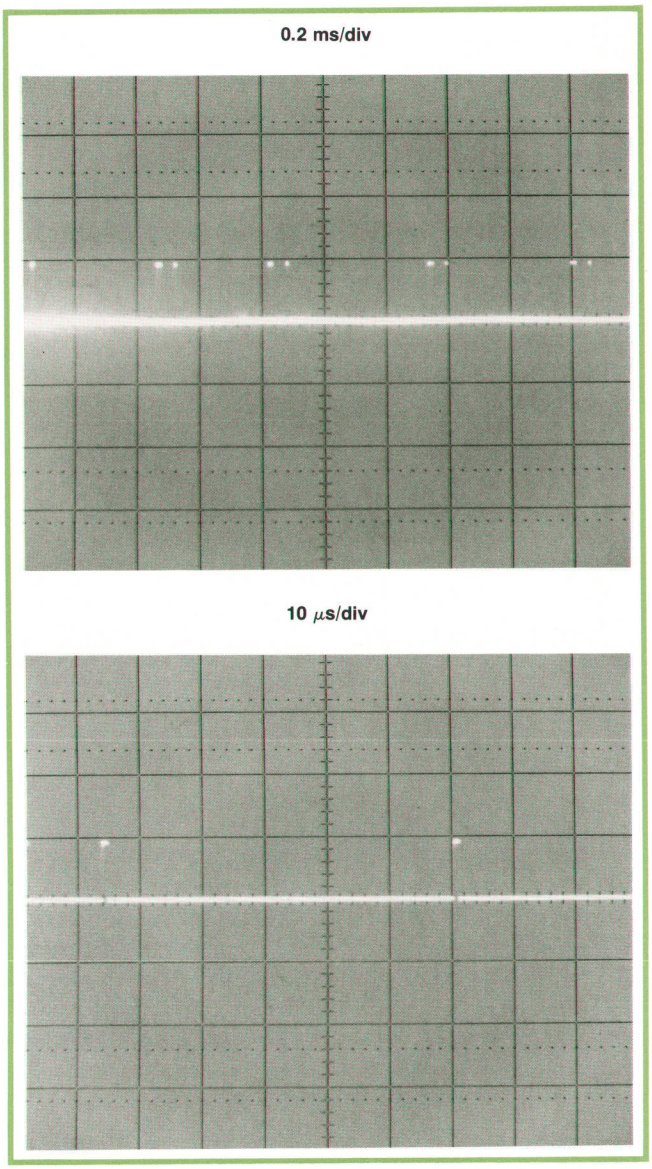

Fig. 3. *Setting a subroutine's entry address into the switches of a breakpoint register makes it possible to see how often that subroutine is used and what its execution time is. Top photo shows this subroutine is invoked about every 0.5 millisecond. Bottom photo shows subroutine is about 70 microseconds long; pulses are entry, parameter pickup, and exit.*

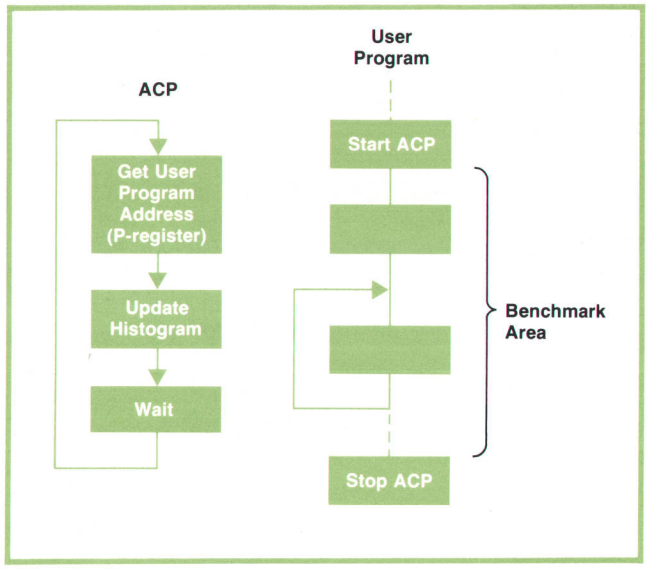

Fig. 4. *The activity profile, a plot showing the percentage of time spent in each area of memory, is an excellent tool for timing analysis. Generating it requires that the program labeled ACP run at the same time as the user's program. Any computer's interrupt system provides this capability.*

several different machines. The IBM 360/67 required creation of a high-priority asynchronous task running off the time-of-day clock. The CDC 6600 used an interrupt processor resident in one of its peripheral processing units (PPUs), and the CDC 3800 used a conventional interrupt processor. In the HP 2100, a short assembly language subroutine (about 50 words long) that used the computer's interrupt system to read an endless loop of paper tape was written. As each interrupt occurred (approximately every three milliseconds using a 300-cps photoreader), the user program would pause, the interrupt subroutine would read the P-register and increment a histogram stored in a 256-word buffer, and then the user program would continue.

Fig. 5 shows this technique applied to a benchmark portion of a FORTRAN program. The CALL to ACP defines the range of P-register values to be included in the histogram and begins the test. The CALL to ACPO ends the test and produces the plots shown in Figs. 6, 7, and 8.

```
      CALL ACP(IPMIN,IRHO,IBUF,256,10B)
      ILOOP=1
      DO 30 I=1,ILOOP
      DO 25 J=1,10
      DO 25 K=1,256
   25 JRPP(K)=ALOG(FLOAT(IRPP(K,J)))*100.
   30 CONTINUE
      CALL ACPO(SUM)
```

Fig. 5. *A portion of a user's FORTRAN program with CALL statements to start and stop ACP processes.*

Fig. 6 shows that the only area in the entire machine that affects system timing is located near memory location 32000. This area is expanded in Fig. 7. Note that both the location and the intensity of timing hot spots stand out. Fig. 8 shows the effect of adding the floating point ROMs.

Upgrading a Computer-Based Signal Processor

The HP 5451B (Fig. 9) is a computer-based digital signal processor. Its calculator-like keyboard allows the user to observe, synthesize, and manipulate waveforms in both the time domain and the frequency domain. Internally, all waveforms are represented digitally to provide large dynamic range (16 bits, equivalent to 96 dB) and the ability to perform precise manipulations.

Fourier analyzers are widely used to solve problems in mechanical vibration, signature analysis, modal analysis, impact testing, digital shaker control, noise source detection, underwater acoustics, servo testing, communications, filter design, and many other areas.

An objective of the 5451B project was to use the microcode capability of the 2100 computer to provide a factor of two or three improvement in speed over the already optimized 5451A software. By using factory programmed PROMs to create several new machine language instructions, and then using these new instructions in some of our time-critical software modules, we hoped to be able to pass on the speed improvement to the user at zero additional cost.

The first engineering task was to develop a method to measure processing performance in a way that would allow comparison of various processing techniques. The methods developed were described in the preceding section. Applied to the 5451A, they produced results like Figs. 10 and 11.

At the same time we assembled the necessary hardware: a writable control store (WCS) for development of microprograms, a PROM writer for later "burning" the debugged microprograms onto PROMs, and an HP 2100 Computer.

After a time we began to get an intuitive feel for what we could do with microcode. During the calculation of a power spectrum, one has to convert an array of complex numbers (A+jB) into an array of magnitude-squares (A^2+B^2). The 2100 machine language MPY instruction takes about 11 microseconds, but we were able to do the entire conversion, including multiplication, double-precision addition, memory access, and array subscripting at only 6.5 μs per word. Only 30 microinstructions were required.

A complicated in-place sorting operation is required in the fast Fourier transform (FFT). Each word (typically there are 1024 words) is moved to its bit-

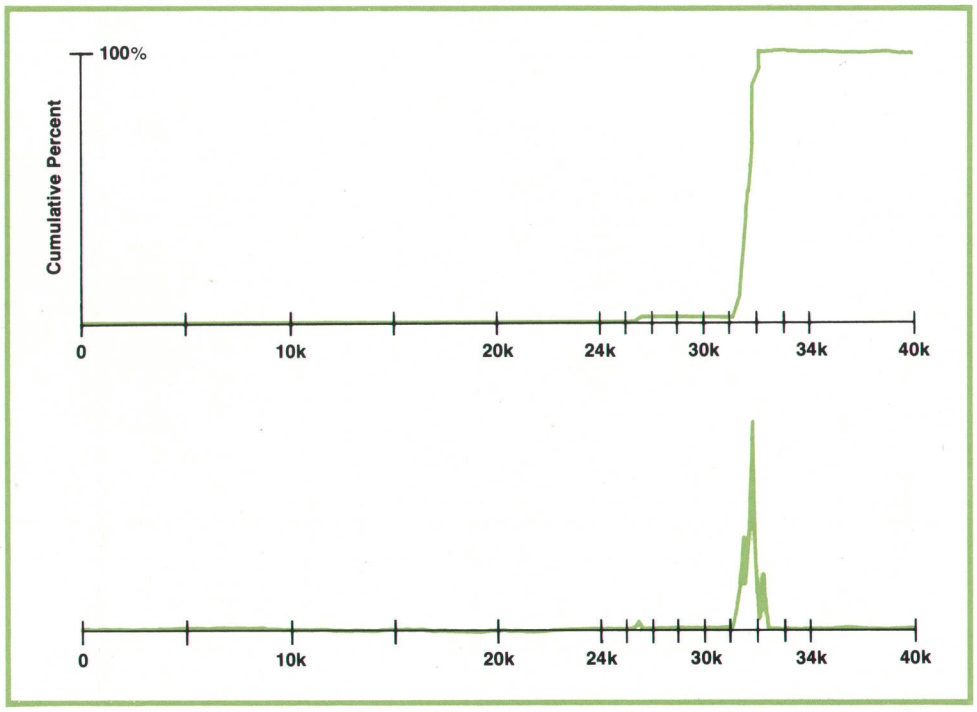

Fig. 6. *Activity profile generated by ACP. The horizontal axis shows memory locations 0 to 40000$_8$ (16384$_{10}$). The subroutines located between 31200$_8$ and 32200$_8$, approximately, are the only ones that influence the execution time of this user's benchmark program. The little bump at 25400$_8$ is the benchmark program itself.*

reversed image; thus the contents of address 0001_2 moves to address 1000_2, the contents of address 0010_2 to 0100_2, the contents of address 0011_2 to 1100_2, etc. The machine language program that performed this in-place sort took about 80 μs/word and required about 60 instructions. The new machine language program takes 4 μs/word and requires only ten instructions. One new machine language instruction was created with a 30-microinstruction microprogram which is not only faster but also smaller than its machine language counterpart.

Fig. 12 compares the speed improvements and incremental costs of firmware and special-purpose hardware methods of improving the speed of the Fourier analyzer.

As we saw more uses for microcode, we realized that ROMs were too inflexible in our application. We did not have a static problem to solve, such as emulating the 2116 instruction set. Furthermore, standard programming techniques for speed enhancement are much easier to apply with WCS than with ROM. As a result, our approach changed. To provide a structure that would allow any future software development effort (by factory, field, or user) to comprise both a ma-

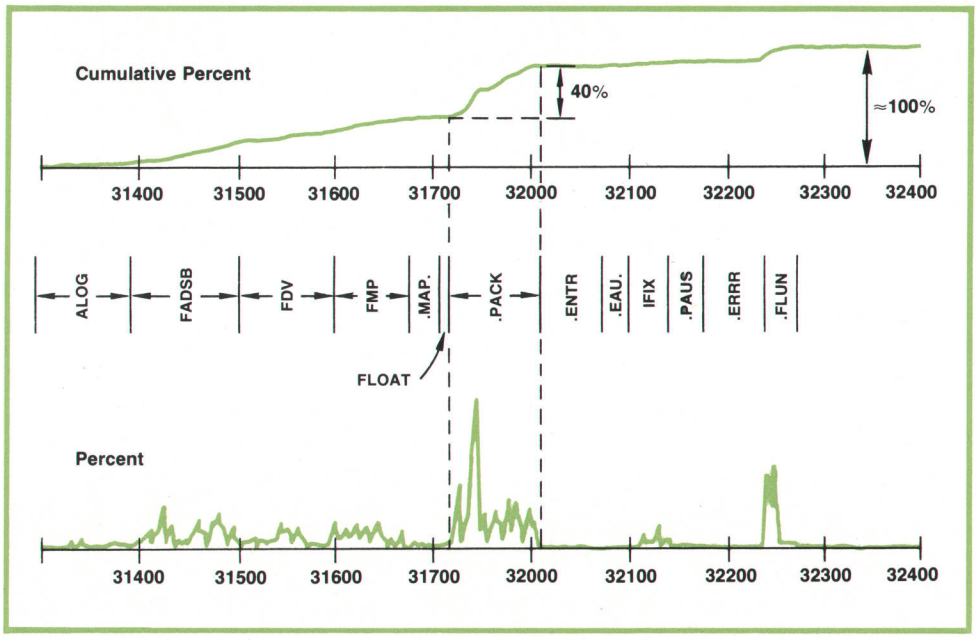

Fig. 7. *An expanded view of the area around 32000$_8$ in Fig. 6. Resolution is two words, compared to 100 words in Fig. 6. Notice how the areas with system performance leverage stand out. The .PACK subroutine takes about 40% of the time, but the .MAP. subroutine, which resolves subscripts, has essentially zero leverage for this benchmark.*

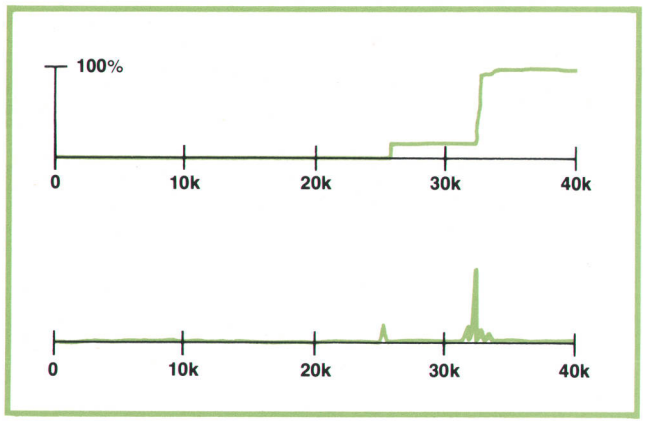

Fig. 8. *Activity profile for the benchmark of Figs. 6 and 7 with floating-point firmware replacing the FADSB, FMP, FLOAT, FIX, and .FLUN routines and eliminating .PACK entirely. The benchmark runs six times faster, an improvement that could have been predicted from Fig. 7. Subroutine ALOG now takes 70% of the time, most of it in a 28-word section, .MAP. takes about 10% of the time, and the user's program about 20%. Thus there is still room for improvement.*

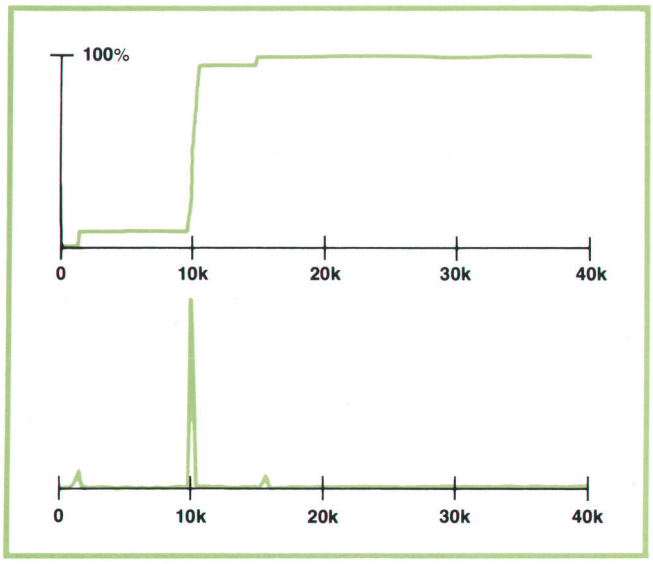

Fig. 10. *Activity profile for a benchmark process used in the 5451A shows that 80% of the time is spent in subroutines near memory location 10000₈.*

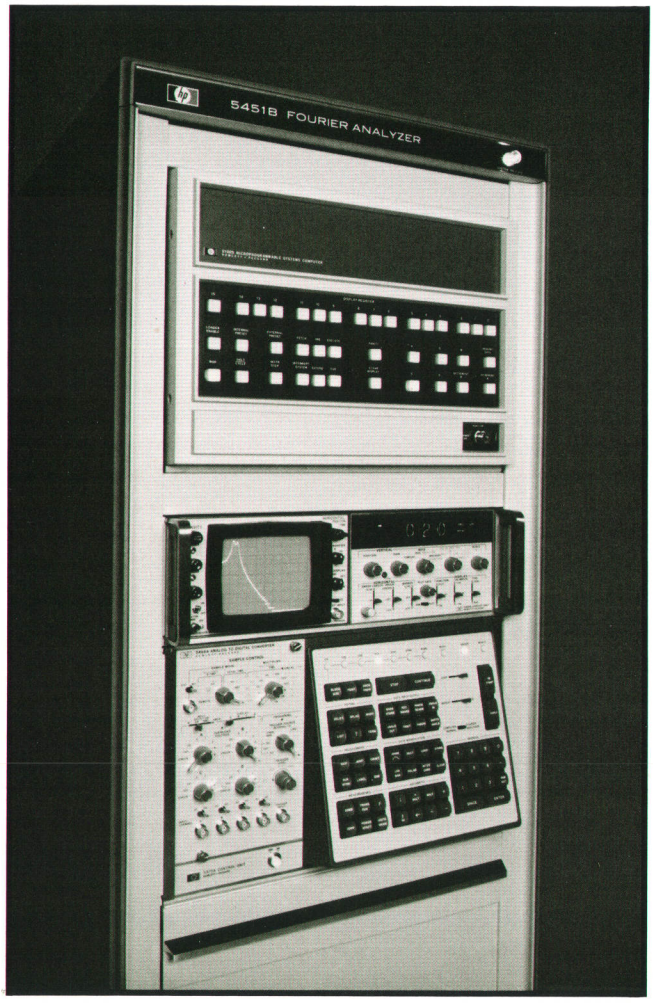

Fig. 9. *Thanks to the microcode capability of the HP2100S computer, the 5451B Fourier Analyzer is ten times faster than the 5451A.*

chine language portion and a microcode portion, and to capitalize on the intrinsic flexibility and higher speed potential of writable control store, WCS was included in the 5451B hardware configuration.

A structure that allows the sharing of WCS can pose coordination problems. Assume that a program uses special machine language instructions defined by microcode residing in WCS. Should the program be responsible for loading WCS? What if it calls a subroutine that loads WCS with something else? Should subroutines be deprived of the use of microcode? If a machine interrupt occurs, should an interrupt processor be allowed to use WCS? If some of the WCS microcode is committed to PROM to eliminate WCS loading time, should the programs using that microcode be changed at all?

The solution chosen in the 5451B involves the use of both PROM and WCS along with the concept of virtual microcode. Users of special microcoded machine language instructions need not be aware of the storage medium or the state of WCS. Any assembly language routine in the 5451B that uses microcode assumes that WCS is loaded with the appropriate microprograms. If that assumption is true, that microcode is executed with zero linkage overhead. If that assumption is false, an illegal opcode trap occurs, and the trap interrupts into an assembly language system subroutine that quickly loads WCS and causes the no-longer-illegal opcode to be re-executed. The dynamic character of WCS is invisible to the user, who may simply assume that all of the microcode is available all of the time.

The virtual microcode concept required several features not present in the normal 2100 architecture:

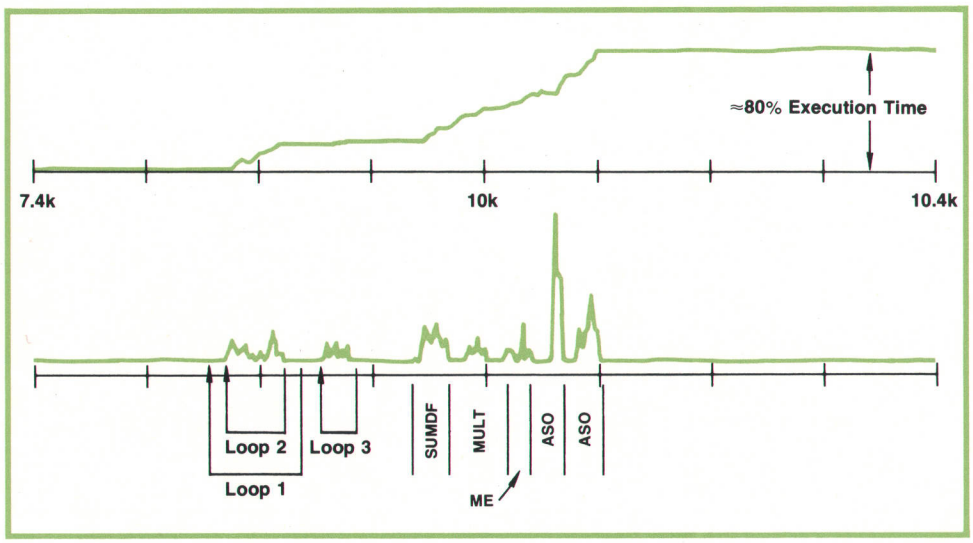

Fig. 11. *Expanded view of the area around 10000₈ in Fig. 10. Elementary operations that determine system speed include complex addition, subtraction, and multiplication, double-precision add-test-round, and trigonometric interpolation.*

illegal instruction trapping, interrupts not involving the I/O structure, a third programmable DMA channel, and others. PROMs were microprogrammed to fill this need. They contain micro-utilities for parameter passing (giving the machine language programmer access to invisible registers and the microprogrammer access to global parameters like the location of the operand stack), illegal-opcode trapping, pseudo-interrupts (allowing microcode to call assembly language routines), pseudo-DMA (500 kHz rate without requiring any special hardware), and other capabilities.

This new set of capabilities allows the user to ignore the operation of loading WCS as well as the time "lost" during the loading. For example, it is possible for the execution time of a subroutine as short as ten milliseconds to be cut by a factor of five or ten, even including the loading of WCS. In fact, during the most time-critical operations in the 5451B Fourier Analyzer, WCS may be reloaded as often as every 10 to 50 milliseconds. The slight cost of loading is more than offset by the additional speed available and by the flexibility that allows different microprogramming efforts (application groups, users, the factory) to work together on a problem solution.

	5451A Software	5451A/ 5471A Hardware	5451B Microcoded	5451A/ 5470A Hardware
Fast Fourier Transform, 1024 Points	1000-1100 ms	160 ms	90-110 ms	16 ms
Power Spectral Average	1400 ms	240 ms	150 ms	30 ms
Cost	$0	$4.5K	$1.5K	$17K
Hardware		4 I/O Cards	1 WCS Card	20 Cards
User Programmable				
Software	Yes	Yes	Yes	Yes
Firmware		No	Yes	No

Fig. 12. *Comparing special-purpose hardware with firmware for speeding Fourier analyzer operation.*

Acknowledgments

I would like to thank Chuck Weyland for his capable help in microcoding, Norm Rogers for his display de-flicker software, Mellba Lindgren for her computer room support, and Peter Roth, whose good judgment allowed us to build a product capitalizing on the flexibility of WCS rather than the (apparent) economy of PROMs.

References

1. R.F. Rosin, "Contemporary Concepts of Microprogramming and Emulation", Computing Surveys, Vol. 1, No. 4, December 1969, pp. 197-212. Describes various methods for realizing the control function of machines; examples of microprogram control in both large and small contemporary (1969) systems are given. Research areas are described, e.g., high level language support, uses of writable control storage, emulation, microlanguages, etc.
2. S. Husson, "Microprogramming Principles and Practices", Prentice-Hall, 1970. This book emphasizes large machines like IBM System 360, Honeywell H4200, and the RCA Spectra 70/45. These machines typically require a very large amount of microcode (e.g., one or two thousand 100-bit microinstructions).
3. C.T. Leis, "Microprogramming, ROMs, Firmware, and All That", Hewlett-Packard Journal, October 1971. Describes the benefits of a "computer within a computer." The system design tradeoffs in the 2100 are discussed. Describes customer engineer's "briefcase" ROM simulator which, using microcoded machine diagnostics, can exercise all hardware logic in the 2100 in only a few minutes. This briefcase contains a papertape reader, a small WCS, a control panel (with built-in breakpoint register), and a cable which plugs into the 2100.

4. "Microprogramming Guide for HP 2100 Computer", Hewlett-Packard publication number 5951-3028. This book describes all the microprogramming characteristics of the 2100. It includes a microcode listing of the 256-word module 0. This module emulates the entire HP 2116 instruction set and then some.
5. "Microprogramming Software for the 2100 Computer", Hewlett-Packard publication number 02100-90133. This shows how to use the microassembler, what columns to type in, etc. It shows how to use the PROM burner (HP 12909A) to convert microcode on paper tape into a permanent pattern on PROMs. Some microprogramming examples are included.
6. "ACP, Activity Profile Generator", Hewlett-Packard Contributed Program Library, No. 22646. This describes the use of the ACP program for the 2100 computer. Examples are given.
7. A.M. Abd-Alla and D.C. Karlgaard, "The Heuristic Synthesis of Applications-Oriented Microcode", the 6th Annual Workshop on Microprogramming, Sept. 1973, ACM MICROSIC, University of Maryland, pp. 83-90. This paper is concerned with heuristically developing specialized instruction sets. A trace file (sequential list of all memory accesses while a benchmark program runs) is used as an analysis technique. Examples of performance improvement are shown for the IBM 370/145 computer.
8. "Microcode Support for the 5451B Fourier Analyzer System," Hewlett-Packard publication number 10637-81003. Describes the internal structure of the 5451B relevant to a user who will be writing his own microcode. Gives examples of analyzing where to microcode as well as examples of microcode.

David C. Snyder

Dave Snyder was project manager for the 5451B Fourier Analyzer. He joined HP in 1970 with five years experience in software development in the fields of high-energy physics, astrodynamics, and airlift management. His HP projects include work on the 5406B Multiparameter Analyzer and the 5407A Scintigraphic Data Analyzer. He's a member of ACM. A 1965 graduate of the University of California at Berkeley, Dave has a BS degree in engineering physics and has done graduate work at three universities in systems, numerical analysis, digital design, and computer science. He's married, has three children, and is planning to build a home someday in the Santa Cruz mountains. For now, he keeps busy with photography, electronic tinkering, and raising Irish setters.

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

HEWLETT-PACKARD JOURNAL
FEBRUARY 1975 Volume 26 • Number 6

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 • Michel Foglia

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