HEWLETT-PACKARD
JOURNAL

APRIL 1989

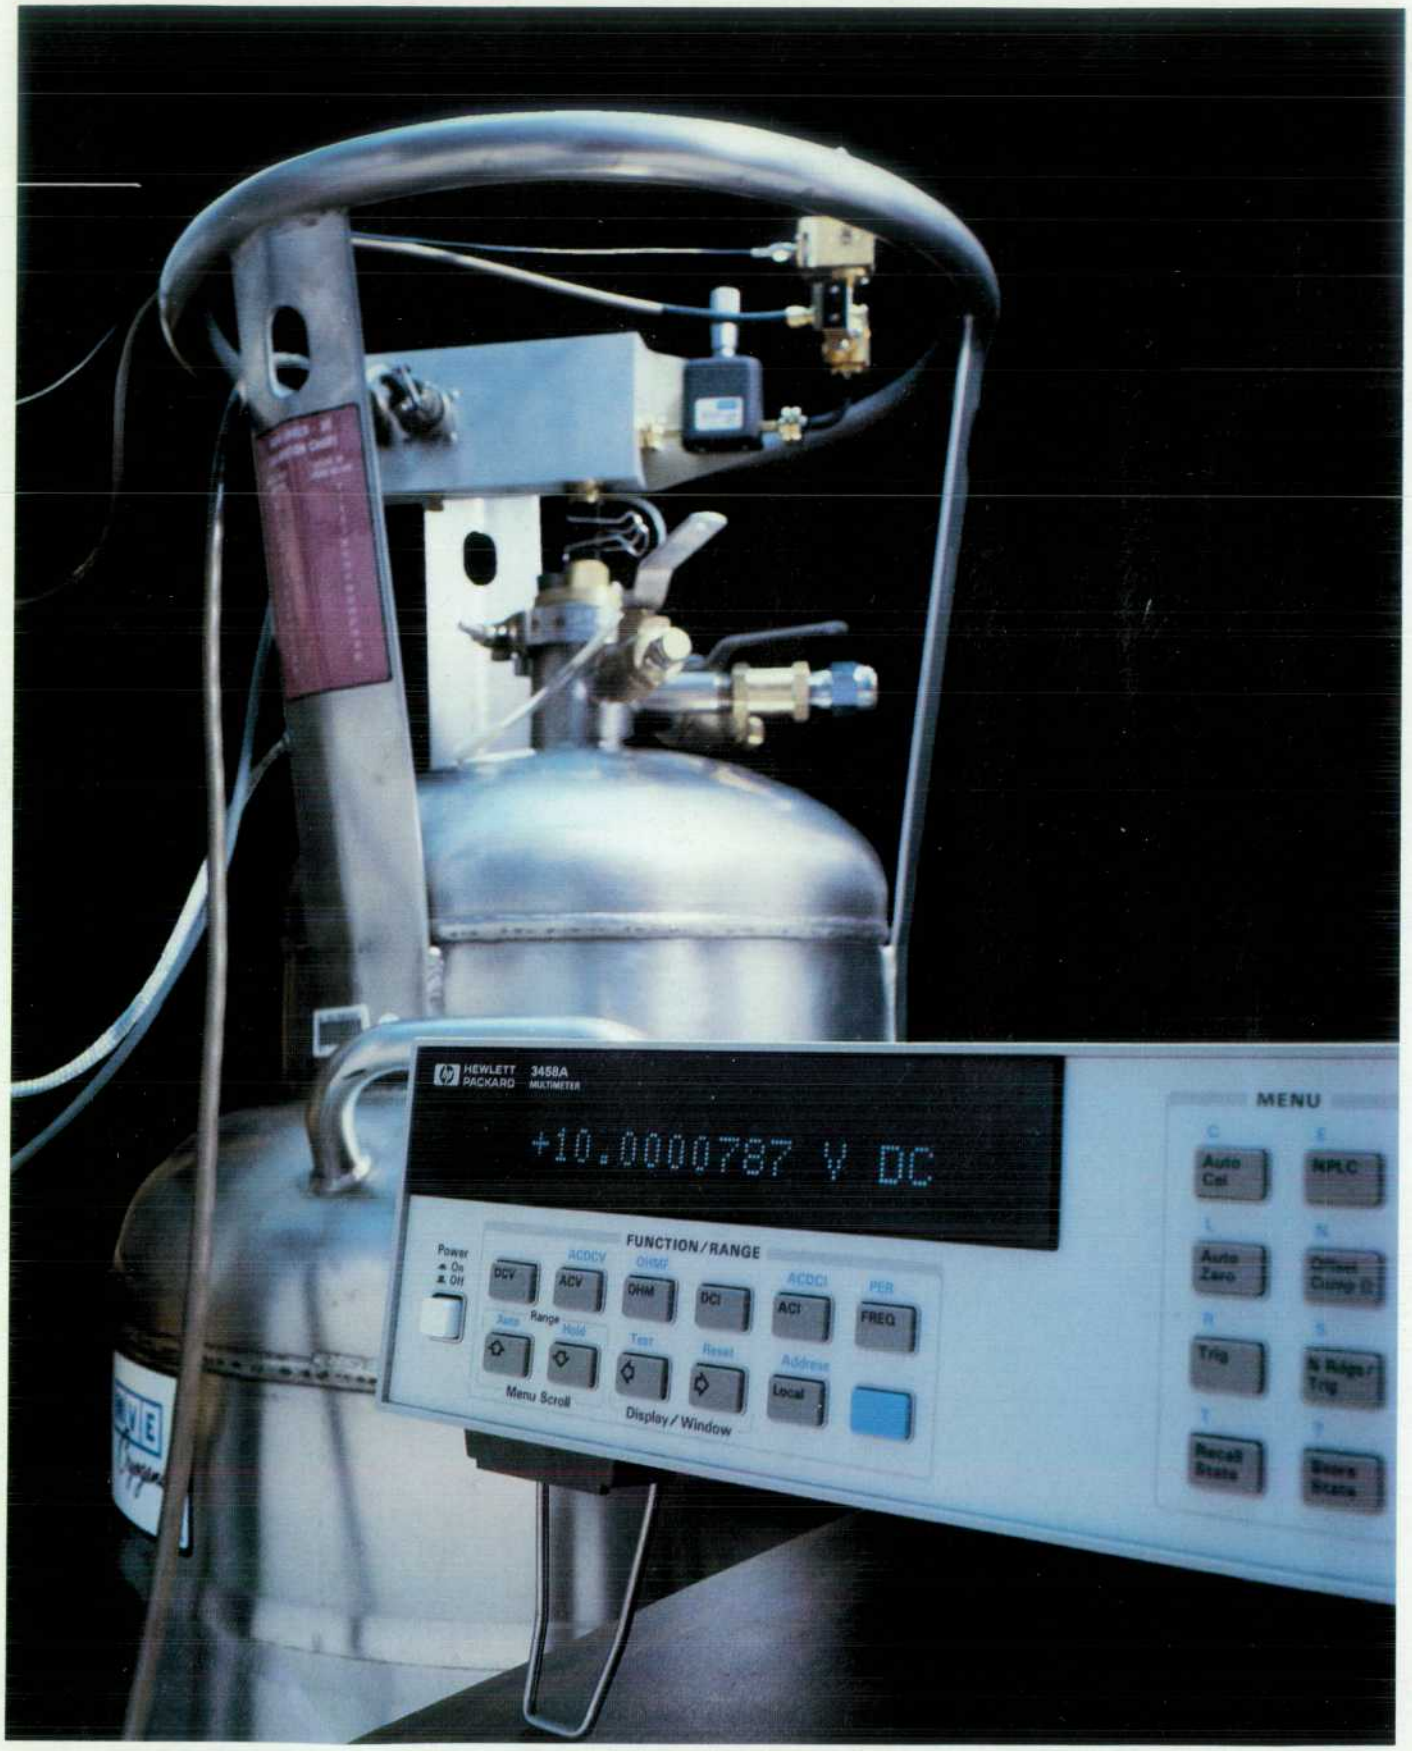

HEWLETT PACKARD

HEWLETT-PACKARD JOURNAL

April 1989 Volume 40 • Number 2

Articles

Departments

The **Hewlett-Packard Journal** is published bimonthly by the Hewlett-Packard Company to recognize technical contributions made by Hewlett-Packard (HP) personnel. While the information found in this publication is believed to be accurate, the Hewlett-Packard Company makes no warranties, express or implied, as to the accuracy or reliability of such information. The Hewlett-Packard Company disclaims all warranties of merchantability and fitness for a particular purpose and all obligations and liabilities for damages, including but not limited to indirect, special, or consequential damages, attorney's and expert's fees, and court costs, arising out of or in connection with this publication.

Subscriptions: The Hewlett-Packard Journal is distributed free of charge to HP research, design, and manufacturing engineering personnel, as well as to qualified non-HP individuals, libraries, and educational institutions. Please address subscription or change of address requests on printed letterhead (or include a business card) to the HP address on the back cover that is closest to you. When submitting a change of address, please include your zip or postal code and a copy of your old label.

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

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

In this Issue

If you thought that voltmeters, those ancient and fundamental instruments, had evolved about as far as possible and that there couldn't be much more to say about them, you'd be wrong. Today, they're usually called digital multimeters or DMMs rather than voltmeters. You can find highly accurate ones in calibration laboratories, very fast ones in automated test systems, and a whole spectrum of performance levels and applications in between these extremes. Generally, more speed means less resolution—that is, fewer digits in the measurement result. Conversely, a higher-resolution measurement generally takes longer. Some DMMs are capable of a range of speeds and resolutions and allow the user to trade one for the other. The HP 3458A Digital Multimeter does this to an unprecedented degree. It can make 100,000 4½-digit measurements per second or six 8½-digit measurements per second, and allows the user an almost continuous selection of speed-versus-resolution trade-offs between these limits. You'll find an introduction to the HP 3458A on page 6. The basis of its performance is a state-of-the-art integrating analog-to-digital converter (ADC) that uses both multislope runup and multislope rundown along with a two-input structure to achieve both high speed and high precision (page 8). So precise is this ADC that it can function as a ratio transfer device for calibration purposes. With the ADC and a trio of built-in transfer standards, all of the functions and ranges of the HP 3458A can be calibrated using only two external traceable standards—10V and 10 kΩ. The article on page 22 explains how this is possible. At the high end of its speed range, the ADC allows the HP 3458A to function as a high-speed digitizer, an unusual role for a DMM (page 39). In fact, ac voltage measurements are made by digitizing the input signal and computing its rms value, eliminating the analog rms-to-dc converters of older designs (page 15). Finally, moving data fast enough to keep up with the ADC was a design challenge in itself. How it was met with a combination of high-speed circuits and efficient firmware is detailed in the article on page 31.

The seven papers on pages 50 to 90 are from the 1988 HP Software Engineering Productivity Conference and should be of interest to software engineers and users concerned with software defect prevention. Collectively, the papers spotlight areas where vigorous software engineering activity is occurring today, namely in structured and object-oriented analysis, design, and testing, and in the development of reliable metrics with which to measure software quality. ▶ In the paper on page 50, engineers from Yokogawa Hewlett-Packard and Tokyo University describe a joint effort to find the flaws in design procedures that increase the likelihood of human errors that result in program faults. Working backwards from faults to human errors to flawed procedures, they propose various structured analysis and design solutions to eliminate the flaws. ▶ The paper on page 57 urges expansion of the software defect data collection process so that project managers can not only determine how best to prevent future defects, but also build a case for making the necessary changes in procedures. The time required to collect and analyze the additional data is shown to be minimal. ▶ That complexity leads to defects is well-established, so monitoring the complexity of software modules during implementation should point out modules that will be defect prone. The paper on page 64 tells how HP's Waltham Division is taking this approach to improve the software development process, using McCabe's cyclomatic complexity metric to measure complexity. ▶ Object-oriented programming, originally conceived for artificial intelligence applications, is now finding wider acceptance. The paper on page 69 reports on problems and methods associated with testing software modules developed with an object-oriented language, C++, for a clinical information system. ▶ In the paper on page 75, Greg Kruger updates his June 1988 paper on the use of a software reliability growth model at HP's Lake Stevens Instrument Division.

The model has generally been successful, but there are pitfalls to be avoided in applying it. ▶ On a software project at HP's Logic Systems Division, some of the engineers used structured methods and some used unstructured methods. Was structured design better? According to the paper on page 80, the results were mixed, but the structured methods offered enough benefits to justify their continued use. ▶ In software development, system analysis precedes design. The paper on page 86 describes a new object-oriented approach suitable for analyzing today's increasingly larger and more complex systems. The authors believe that designs based on object-oriented specifications can be procedure-oriented or object-oriented with equal success.

Modular measurement systems consist of instruments on cards that plug into a card cage or mainframe. A user can tailor a system to an application by plugging the right modules into the mainframe. The VXIbus is a new industry standard for such systems. Modules and mainframes conforming to the VXIbus standard are compatible no matter what company manufactured them. The articles on pages 91 and 96 introduce the VXIbus and some new HP products that help manufacturers develop VXIbus modules more quickly. HP's own modular measurement system architecture conforms to the VXIbus standard where applicable. However, for high-performance RF and microwave instrumentation, HP has used a proprietary modular system interface bus (HP-MSIB). Patent rights to the HP-MSIB have now been assigned to the public so that this architecture can be used by everyone as the high-frequency counterpart of the VXIbus.

R.P. Dolan
Editor

Cover

So precise is the 3458A Digital Multimeter that verifying some aspects of its performance is beyond the limits of conventional standards. In the HP Loveland Instrument Division Standards Laboratory, the HP 3458A's linearity is measured using a 10-volt Josephson junction array developed by the U.S. National Institute of Standards and Technology. The array is in a specially magnetically shielded cryoprobe in the center of a liquid-helium-filled dewar (the tank with the protective "steering wheel.") On top of the dewar are a Gunn-diode signal source (72 GHz) and various microwave components. A waveguide and voltage and sense leads connect the array to the external components. For more details see page 24.

What's Ahead

Subjects to be covered in the June issue include:

- The HP 9000 Model 835 and HP 3000 Series 935 Midrange HP Precision Architecture Computers
- Programming with neurons
- A new 2D simulation model for electromigration in thin metal films
- Data compression and blocking in the HP 7980XC Tape Drive
- Design and applications of HP 8702A Lightwave Component Analyzer systems
- A data base for real-time applications and environments
- A hardware/software tool to automate testing of software for the HP Vectra Personal Computer.

An 8½-Digit Digital Multimeter Capable of 100,000 Readings per Second and Two-Source Calibration

A highly linear and extremely flexible analog-to-digital converter and a state-of-the-art design give this DMM new performance and measurement capabilities for automated test, calibration laboratory, or R&D applications.

by Scott D. Stever

THE DIGITAL MULTIMETER OR DMM is among the most common and most versatile instruments available for low-frequency and dc measurements in automated test, calibration laboratory, and bench R&D applications. The use of general-purpose instrumentation in automated measurement systems has steadily grown over the past decade. While early users of programmable instruments were elated to be able to automate costly, tedious, error-prone measurements or characterization processes, the sophistication and needs of today's users are orders of magnitude greater. The computing power of instrument controllers has increased manyfold since the mid-1970s and so have user expectations for the performance of measurement systems. Test efficiency in many applications is no longer limited by the device under test or the instrument controller's horsepower. Often either the configuration speed or the measurement speed of the test instrumentation has become the limiting factor for achieving greater test throughput. In many systems, the DMM is required to perform hundreds of measurements and be capable of multiple functions with various resolutions and accuracies.

In some applications, several DMMs may be required to characterize a single device. For example, measurements requiring high precision may need a slower DMM with calibration laboratory performance. Usually, the majority of measurements can be satisfied by the faster, moderate-resolution capabilities of a traditional system DMM. In extreme cases, where speed or sample timing are critical to the application, a lower-resolution high-speed DMM may be required. A single digital multimeter capable of fulfilling this broad range of measurement capabilities can reduce system complexity and development costs. If it also provides shorter reconfiguration time and increased measurement speed, test throughput can also be improved for automated test applications.

The HP 3458A Digital Multimeter (Fig. 1) was developed to address the increasing requirements for flexible, accurate, and cost-effective solutions in today's automated test applications. The product concept centers upon the synergistic application of state-of-the-art technologies to meet these needs. While it is tuned for high throughput in computer-aided testing, the HP 3458A also offers calibration laboratory accuracy in dc volts, ac volts, and resistance. Owners can trade speed for resolution, from 100,000 measurements per second with 4½-digit (16-bit) resolution to six measurements per second with 8½-digit resolution. At 5½-digit resolution, the DMM achieves 50,000 readings per second. To maximize the measurement speed for the resolution selected, the integration time is selectable from 500 nanoseconds to one second in 100-ns steps. The effect is an almost continuous range of speed-versus-resolution trade-offs.

Fig. 1. *The HP 3458A Digital Multimeter can make 100,000 4½-digit readings per second for high-speed automated test applications. For calibration laboratory applications, it can make six 8½-digit readings per second. Fine control of the integration aperture allows a nearly continuous range of speed-versus-resoultion trade-offs.*

Measurement Capabilities

Measurement ranges for the HP 3458A's functions are:

- Voltage: 10 nV to 1000V dc
  <1 mV to 700V rms ac
- Current: 1 pA to 1A dc
  100 pA to 1A rms ac
- Resistance: 10 μΩ to 1 GΩ, 2-wire or 4-wire
- Frequency: 1 Hz to 10 MHz
- Period: 100 ns to 1 s
- 16-bit digitizing at effective sample rates to 100 megasamples/second.

The ac voltage bandwidth is 1 Hz to 10 MHz, either ac or dc coupled.

To increase uptime, the HP 3458A is capable of two-source electronic calibration and self-verifying autocalibration. Autocalibration enhances accuracy by eliminating drift errors with time and temperature. The dc voltage stability is specified at eight parts per million over one year, or 4 ppm with the high-stability option. Linearity is specified at 0.1 ppm, transfer accuracy at 0.1 ppm, and rms internal noise at 0.01 ppm. Maximum accuracies are 0.5 ppm for 24 hours in dc volts and 100 ppm in ac volts. Midrange resistance and direct current accuracies are 3 ppm and 10 ppm, respectively.

The HP 3458A can transfer 16-bit readings to an HP 9000 Series 200 or 300 Computer via the HP-IB (IEEE 488, IEC 625) at 100,000 readings per second. It can change functions or ranges and deliver a measurement 200 times per second (over 300/s from the internal program memory), about four times faster than any earlier HP multimeter.

The following five articles describe the technologies required to achieve this performance and the benefits that result. In the first paper, the development of a single analog-to-digital converter capable of both high resolution and high speed is discussed. The second paper describes the development of a technique for the precise measurement of rms ac voltages. The application of these technologies to provide improved measurement accuracy over extended operating conditions and to provide complete calibration of the DMM from only two external traceable sources (10V dc, 10 kΩ) is discussed in the third article. Hardware and firmware design to achieve increased measurement throughput is the topic of the fourth paper. The final paper discusses several applications for the HP 3458A's ability to perform high-resolution, high-speed digitizing.

An 8½-Digit Integrating Analog-to-Digital Converter with 16-Bit, 100,000-Sample-per-Second Performance

This integrating-type ADC uses multislope runup, multislope rundown, and a two-input structure to achieve the required speed, resolution, and linearity.

by Wayne C. Goeke

THE ANALOG-TO-DIGITAL CONVERTER (ADC) design for the HP 3458A Digital Multimeter was driven by the state-of-the-art requirements for the system design. For example, autocalibration required an ADC with 8½-digit (28-bit) resolution and 7½-digit (25-bit) integral linearity, and the digital ac technique (see article, page 22) required an ADC capable of making 50,000 readings per second with 18-bit resolution.

Integrating ADCs have always been known for their ability to make high-resolution measurements, but tend to be relatively slow. When the HP 3458A's design was started, the fastest integrating ADC known was in the HP 3456A DMM. This ADC uses a technique known as multislope and is capable of making 330 readings per second. The HP 3458A's ADC uses an enhanced implementation of the same multislope technique to achieve a range of speeds and resolutions never before achieved—from 16-bit resolution at 100,000 readings per second to 28-bit resolution at six readings per second. In addition to high resolution, the ADC has high integral linearity—deviations are less than 0.1 ppm (parts per million) of input.

Multislope is a versatile ADC technique, allowing speed to be traded off for resolution within a single circuit. It is easier to understand multislope by first understanding its predecessor, dual-slope.

Basic Dual-Slope Theory

Dual-slope is a simple integrating-type ADC algorithm. Fig. 1 shows a simple circuit for implementing a dual-slope ADC.

The algorithm starts with the integrator at zero volts. (This is achieved by shorting the integrator capacitor, C.) At time 0 the unknown input voltage V_{in} is applied to the resistor R by closing switch SW1 for a fixed length of time t_u. This portion of the algorithm, in which the unknown input is being integrated, is known as runup. At the end of runup (i.e., when SW1 is opened), the output of the integrator, V_o, can be shown to be

$$V_o(t_u) = -(1/RC)\int_0^{t_u} V_{in}(t)dt$$

or, when V_{in} is time invariant,

$$V_o(t_u) = -(1/RC)\, V_{in} t_u.$$

Next a known reference voltage V_{ref} with polarity opposite to that of V_{in} is connected to the same resistor R by closing SW2. A counter is started at this time and is stopped when the output of the integrator crosses through zero volts. This portion of the algorithm is known as rundown. The counter contents can be shown to be proportional to the unknown input.

$$V_o(t_2) = V_o(t_u) - (1/RC)V_{ref}t_d = 0,$$

where t_d is the time required to complete rundown (i.e., $t_d = t_2 - t_u$). Solving for V_{in},

$$V_{in} = -V_{ref}(t_d/t_u).$$

Letting N_u be the number of clock periods (T_{ck}) during

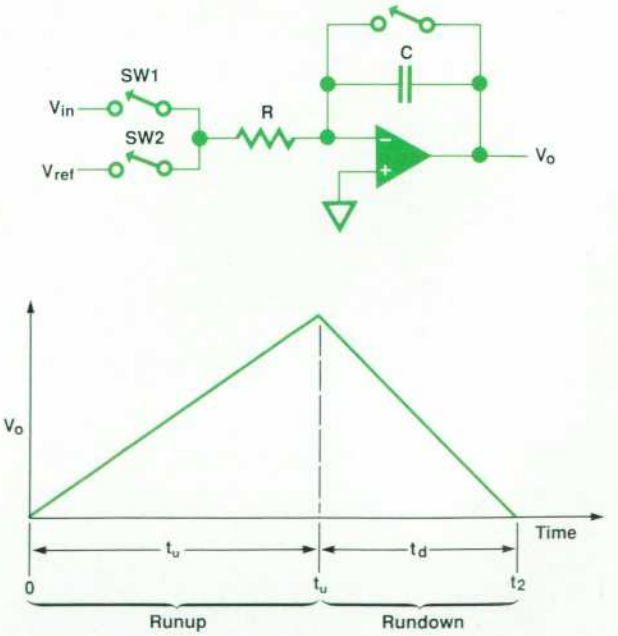

Fig. 1. *Dual-slope integrating ADC (analog-to-digital converter) circuit and a typical waveform.*

runup and N_d the number of clock periods during rundown, time cancels and

$$V_{in} = - V_{ref}(N_d/N_u).$$

The beauty of the dual-slope ADC technique is its insensitivity to the value of most of the circuit parameters. The values of R, C, and T_{ck} all cancel from the final equation. Another advantage of the dual-slope ADC is that a single circuit can be designed to trade speed for resolution. If the runup time is shortened, the resolution will be reduced, but so will the time required to make the measurement.

The problem with the dual-slope algorithm is that its resolution and speed are limited. The time T_m for a dual-slope ADC to make a measurement is determined by:

$$T_m = 2T_{ck}M,$$

where T_m is the minimum theoretical time to make a full-scale measurement, T_{ck} is the period of the ADC clock, and M is the number of counts of resolution in a full-scale measurement. Even with a clock frequency of 20 MHz, to measure a signal with a resolution of 10,000 counts requires at least 1 millisecond.

The resolution of the dual-slope ADC is limited by the wideband circuit noise and the maximum voltage swing of the integrator, about ±10 volts. The wideband circuit noise limits how precisely the zero crossing can be determined. Determining the zero crossing to much better than a millivolt becomes very difficult. Thus, dual-slope is limited in a practical sense to four or five digits of resolution (i.e., 10V/1 mV).

Enhanced Dual-Slope

The speed of the dual-slope ADC can be nearly doubled simply by using a pair of resistors, one for runup and the other for rundown, as shown in Fig. 2.

The unknown voltage, V_{in}, is connected to resistor R_u, which is much smaller than resistor R_d, which is used during rundown. This allows the runup time to be shortened by the ratio of the two resistors while maintaining the same resolution during rundown. The cost of the added speed is an additional resistor and a sensitivity to the ratio of the two resistors:

$$V_{in} = - V_{ref}(N_d/N_u)(R_u/R_d).$$

Because resistor networks can be produced with excellent ratio tracking characteristics, the enhancement is very feasible.

Multislope Rundown

Enhanced dual-slope reduces the time to perform runup. Multislope rundown reduces the time to perform rundown. Instead of using a single resistor (i.e., a single slope) to seek zero, multislope uses several resistors (i.e., multiple slopes) and seeks zero several times, each time more precisely. The ratio of one slope to another is a power of some number base, such as base 2 or base 10.

Fig. 3 shows a multislope circuit using base 10. Four slopes are used in this circuit, with weights of 1000, 100, 10, and 1. Each slope is given a name denoting its weight

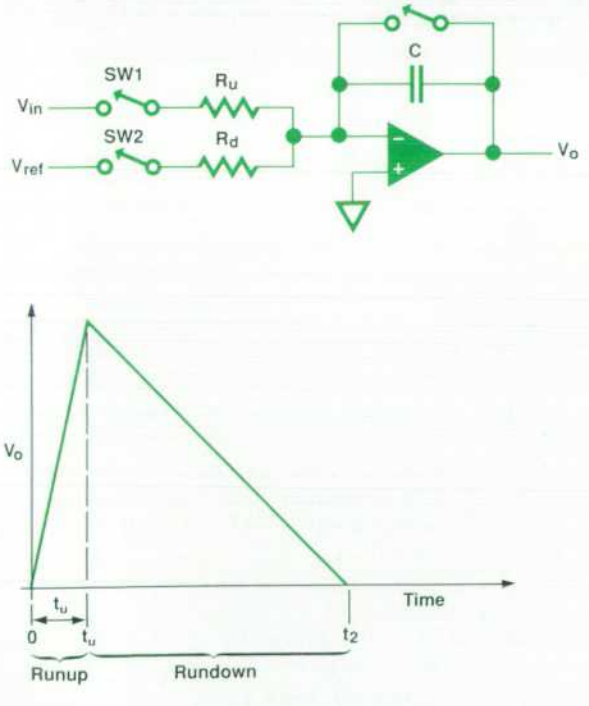

Fig. 2. *Enhanced dual-slope ADC circuit uses two resistors, one for runup and one for rundown.*

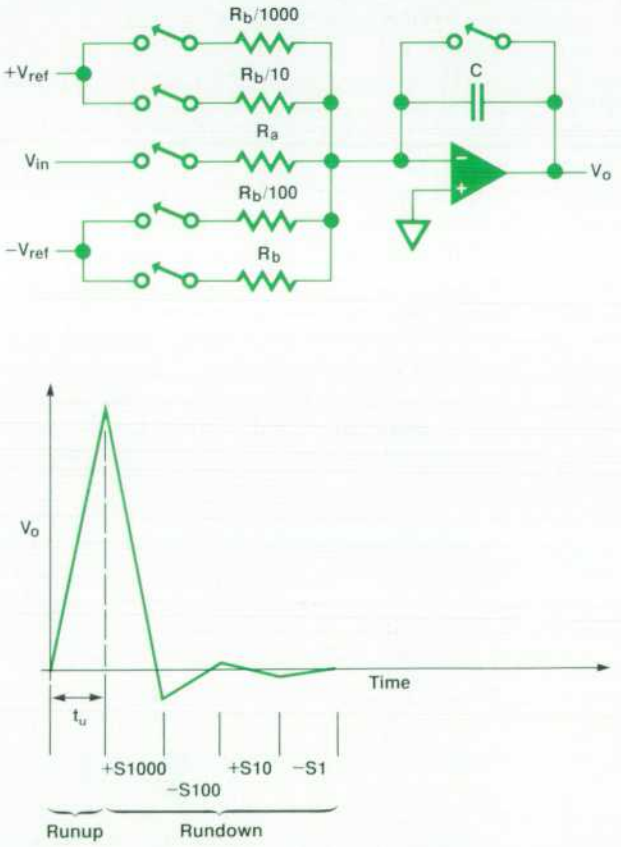

Fig. 3. *Base-10 multislope rundown circuit.*

and polarity. For example, S1000 is a positive slope worth 1000 counts per clock period and −S100 is a negative slope worth −100 counts per clock period. A slope is considered to be positive if it transfers charge into the integrator. This may be confusing because the integrator (an inverting circuit) actually moves in a negative direction during a positive slope and vice versa.

The multislope rundown begins by switching on the steepest slope, S1000. This slope remains on until the integrator output crosses zero, at which time it is turned off and the next slope, −S100, is turned on until the output crosses back through zero. The S10 slope follows next, and finally, the −S1 slope. Each slope determines the integrator's zero crossing ten times more precisely than the previous slope. This can be viewed as a process in which each slope adds another digit of resolution to the rundown.

If each slope is turned off within one clock period of crossing zero, then each subsequent slope should take ten or fewer clock periods to cross zero. Theoretically, then, the time t_d to complete a multislope rundown is:

$$t_d < NBT_{ck},$$

where N is the number of slopes and B is the number base of the slope ratios. In practice, the time to complete rundown is higher, because it isn't always possible to to turn off each slope within a clock period of its zero crossing. Delays in detecting the zero crossings and delays in responding by turning off the slopes cause the actual time to be:

$$t_d < kNBT_{ck},$$

where k is a factor greater than one. The delay in turning off a slope results in the integrator output's overshooting zero. For each clock period of overshoot, the following slope must take B clock periods to overcome the overshoot. Typical values of k range from two to four. The multislope rundown shown in Fig. 3 completes a measurement yielding 10,000 counts of resolution in 4.0 μs assuming a 20-MHz clock and k = 2. This is 125 times faster than the equivalent dual-slope rundown.

Multislope can be optimized for even faster measurements by choosing the optimum base. Noting that the number of slopes, N, can be written as $\log_B(M)$, where M is the number of counts of resolution required from rundown,

$$t_d < kB\log_B(M)T_{ck}.$$

This yields base e as the optimum base regardless of the required resolution. Using base e in the above example results in a rundown time of 2.5 μs. This is a 60% increase in multislope rundown speed as a result of using base e instead of base 10.

There is a cost associated with implementing multislope rundown. A resistor network must be produced with several resistors that have precise ratios. The tightest ratio tolerance is the reciprocal of the weight of the steepest slope and must be maintained to ensure linear ADC operation. If the ratio tolerances are no tighter than 0.05%, then this requirement is feasible. Multislope also requires a more complex circuit to control and accumulate the measurement, but with the reduced cost and increased density of digital circuits, this is also feasible.

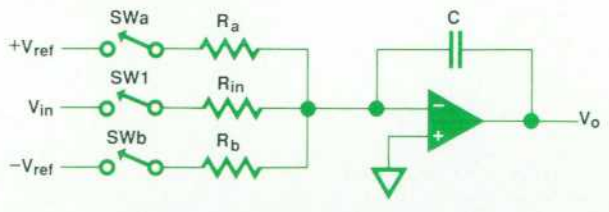

Fig. 4. *Multislope runup circuit.*

Multislope Runup

Multislope runup is a modification of dual-slope runup with the purpose of increasing the resolution of the ADC. As mentioned earlier, the dual-slope technique's resolution is limited by the maximum voltage swing of the integrator and the wideband circuit noise. Multislope runup allows the ADC to have an effective voltage swing much larger than the physical limitations of the integrator circuit hardware.

The technique involves periodically adding and subtracting reference charge to or from the integrator during runup such that the charge from the unknown input plus the total reference charge is never large enough to saturate the integrator. By accounting for the total amount of reference charge transferred to the integrator during runup and adding this number to the result of rundown, a measurement can be made with much higher resolution. Fig. 4 shows a circuit for implementing multislope runup.

A precise amount of reference charge is generated by applying either a positive reference voltage to resistor R_a or a negative reference voltage to resistor R_b for a fixed amount of time. The following table shows the four possible runup reference currents using this circuit.

Slope Name	SWa	SWb	Integrator Direction	Current
S_+	$+V_{ref}$	0	↘	+I
S_{+0}	0	0	–	0
S_-	0	$-V_{ref}$	↗	−I
S_{-0}	$+V_{ref}$	$-V_{ref}$	–	0

Notice that, like multislope rundown, S_+ adds charge to the integrator and S_- subtracts charge from the integrator. If we design the S_+ and S_- currents to have equal magnitudes that are slightly greater than that of the current generated by a full-scale input signal, then the reference currents will always be able to remove the charge accumulating from the input signal. Therefore, the integrator can be kept from being saturated by periodically sensing the polarity of the integrator output and turning on either S_+ or S_- such that the integrator output is forced to move towards or across zero.

Fig. 5 shows a typical multislope runup waveform. The dashed line shows the effective voltage swing, that is, the voltage swing without reference charge being put into the integrator. The integrator output is staying within the limits of the circuit while the effective voltage swing ramps far

beyond the limit. The HP 3458A has an effective voltage swing of ±120,000 volts when making 8½-digit readings, which means the rundown needs to resolve a millivolt to achieve an 8½-digit reading (i.e., 120,000V/0.001V = 120,000,000 counts).

The multislope runup algorithm has two advantages over dual-slope runup: (1) the runup can be continued for any length of time without saturating the integrator, and (2) resolution can be achieved during runup as well as during rundown. The HP 3458A resolves the first 4½ digits during runup and the final 4 digits during rundown to achieve an 8½-digit reading.

An important requirement for any ADC is that it be linear. With the algorithm described above, multislope runup would not be linear. This is because each switch transition transfers an unpredictable amount of charge into the integrator during the rise and fall times. Fig. 6 shows two waveforms that should result in the same amount of charge transferred to the integrator, but because of the different number of switch transitions, do not.

This problem can be overcome if each switch is operated a constant number of times for each reading, regardless of the input signal. If this is done, the charge transferred during the transitions will result in an offset in all readings. Offsets can be easily removed by applying a zero input periodically and subtracting the result from all subsequent readings. The zero measurement must be repeated periodically because the rise and fall times of the switches drift with temperature and thus the offset will drift.

Multislope runup algorithms can be implemented with constant numbers of switch transitions by alternately placing an S_{+0} and an S_{-0} between each runup slope. Fig. 7 shows the four possible slope patterns between any two S_{+0} slopes. Varying input voltages will cause the algorithm to change between these four patterns, but regardless of which pattern is chosen, each switch makes one and only one transition between the first S_{+0} slope and the S_{-0} slope, and the opposite transition between the S_{-0} slope and the second S_{+0} slope.

The cost of multislope runup is relatively small. The runup slopes can have the same weight as the first slope of multislope rundown. Therefore, only the opposite-polarity slope has to be added, along with the logic to implement the algorithm.

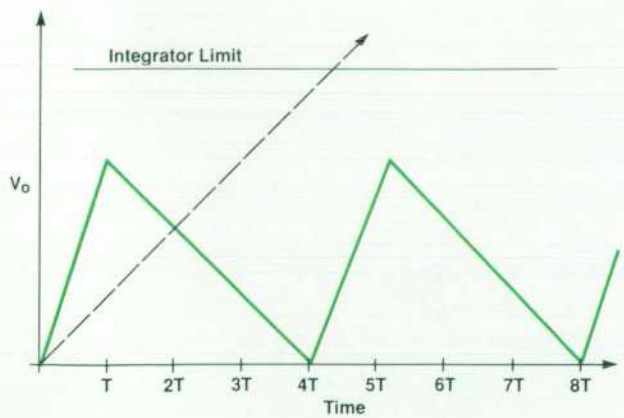

Fig. 5. *Integrator output waveform for multislope runup. The dashed line shows the effective integrator output voltage swing.*

HP 3458A ADC Design

The design of the HP 3458A's ADC is based on these theories for a multislope ADC. Decisions had to be made on how to control the ADC, what number base to use, how fast the integrator can be slewed and remain linear, how much current to force into the integrator (i.e., the size of the resistors), and many other questions. The decisions were affected by both the high-speed goals and the high-resolution goals. For example, very steep slopes are needed to achieve high speed, but steep slopes cause integrator circuits to behave too nonlinearly for high-resolution measurement performance.

One of the easier decisions was to choose a number base for the ADC's multislope rundown. Base e is the optimum to achieve the highest speed, but the task of accumulating an answer is difficult, requiring a conversion to binary. Base 2 and base 4 are both well-suited for binary systems and are close to base e. Base 2 and base 4 are actually equally fast, about 6% slower than base e, but base 2 uses twice as many slopes to achieve the same resolution. Therefore, base 4 was chosen to achieve the required speed with minimum hardware cost.

Microprocessors have always been used to control multislope ADCs, but the speed goals for the HP 3458A quickly

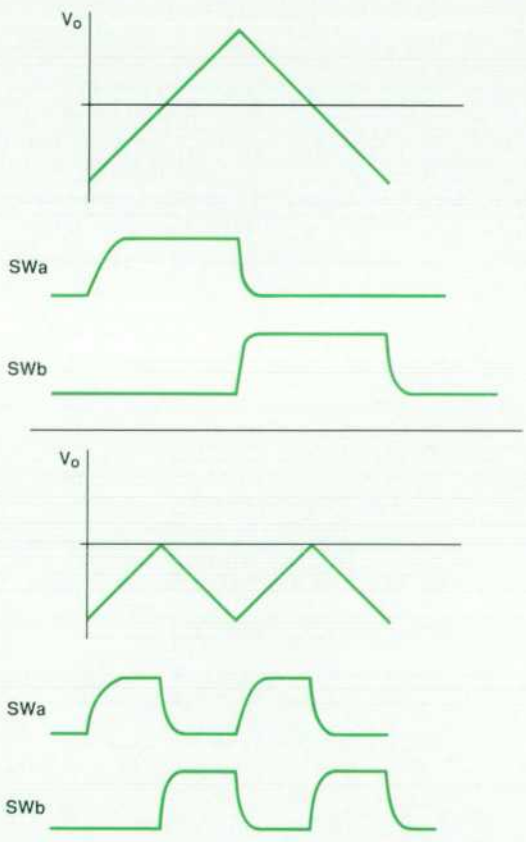

Fig. 6. *Ideally, these two waveforms would transfer equal charge into the integrator, but because of the different number of switch transitions, they do not.*

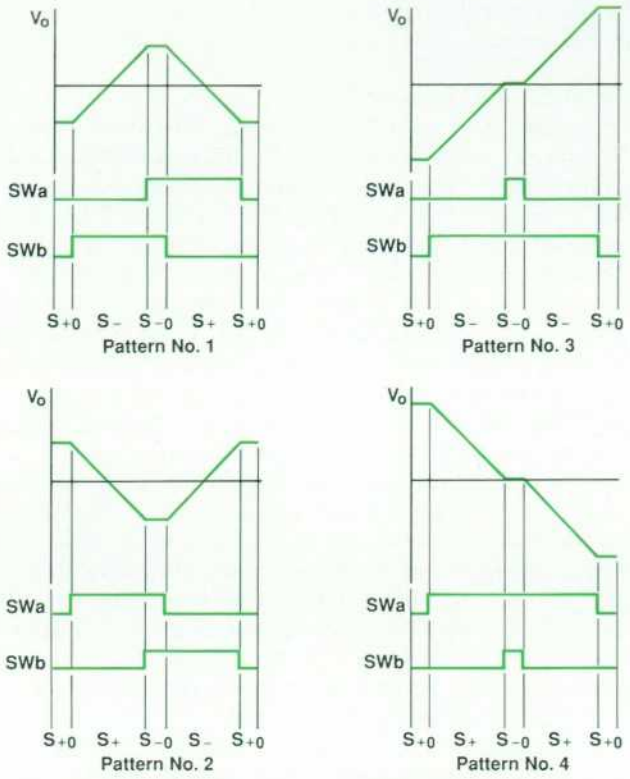

Fig. 7. *Multislope runup patterns for an algorithm that keeps the number of switch transitions constant.*

eliminated the possibility of using a microprocessor to control the ADC algorithm. It was anticipated that the ADC clock frequency would have to be between 10 and 20 MHz, and making decisions at these rates requires dedicated hardware. Therefore, a gate array was chosen to implement state machines running at 20 MHz for ADC control. The ADC control and accumulator functions consume approximately half of a 6000-gate CMOS gate array. The other half of the gate array is devoted to the timing and counting of triggers and a UART (universal asynchronous receiver/transmitter) to transfer data and commands through a 2-Mbit/s fiber optic link to and from the ground-referenced logic (see article, page 31).

The number of slopes and the magnitude of the currents for each slope are more subtle decisions. If the slope currents get too large, they stress the output stage of the integrator's operational amplifier, which can cause nonlinear behavior. If the currents are too small, switch and amplifier leakage currents can become larger than the smallest slope current, and the slope current would not be able to converge the integrator toward zero. A minimum of a microampere for the smallest slope was set to avoid leakage current problems. Also, it was believed that the integrator could handle several milliamperes of input current and remain linear over five or six digits, but that less than a milliampere of input current would be required to achieve linearity over seven or eight digits. On the other hand, greater than a milliampere of current was needed to achieve the high-speed reading rate goal. Therefore, a two-input ADC structure was chosen.

As shown in Fig. 8, when making high-speed measurements, the input voltage is applied through a 10-kΩ resistor, and the ADC's largest slopes, having currents greater than a milliampere, are used. When making high-resolution measurements, the input voltage is applied through a 50-kΩ resistor and the largest slopes used have less than a milliampere of current. The largest slope was chosen to be S1024, having 1.2 μA of current. This led to a total of six slopes (S1024, S256, S64, S16, S4, and S1) with S1 having about 1.2 μA of current. S1024 and S256 are both used during multislope runup; therefore, both polarities exist for both slopes. The ±S256 slopes (0.3 mA) are used when the 50-kohm input is used and both the ±S1024 and the ±S256 slopes (1.5 mA total) are used in parallel when the 10-kΩ input is used. The S256 slope is 25% stronger than a full-scale input to the 50-kΩ resistor, which allows it to keep the integrator from saturating. The 10-kΩ input is five times

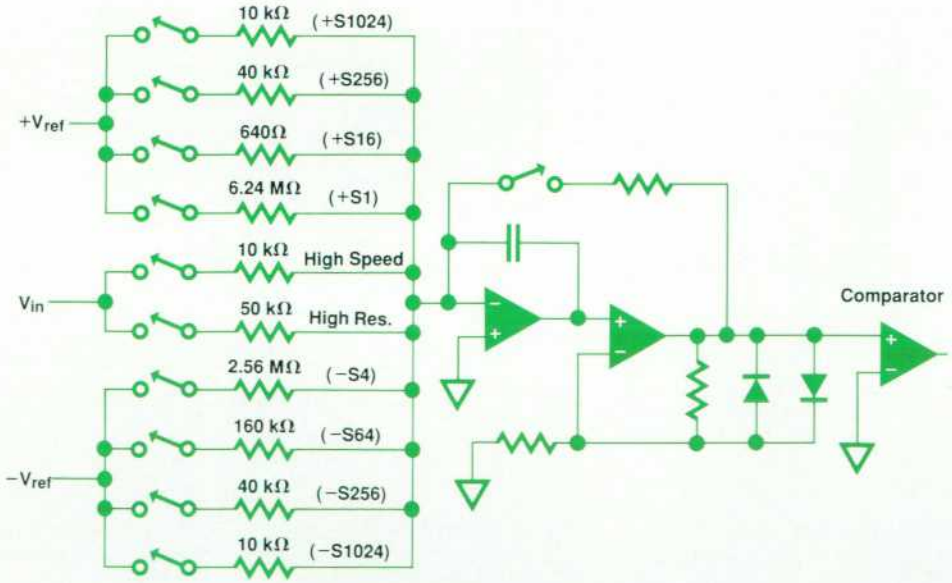

Fig. 8. *Simplified HP 3458A ADC circuit.*

stronger than the 50-kohm input; thus, by using both S1024 and S256, 25% stronger reference slopes can be maintained during high-speed measurements.

Integrator

The integrator's operational amplifier behaves more nonlinearly as the integrator slew rate (i.e., the steepest slope) approaches the slew rate of the amplifier. Two factors determine the integrator slew rate: the total current into the integrator and the size of the integrator capacitor. Wanting to keep the integrator slew rate less than 10V/μs led to an integrator capacitor of 330 pF. This capacitor must have a very small amount of dielectric absorption since 50 fC of charge is one count.

The integrator circuit has to respond to a change in reference current and settle to near 0.01% before the next possible switch transition (about 200 ns). It also has to have low voltage and current noise, about 100V/μs slew rate, a dc gain of at least 25,000, an offset voltage less than 5 mV, and a bias current of less than 10 nA. A custom amplifier design was necessary to achieve all the specifications.

Resistor Network

The resistor network has several requirements. The most stringent is to obtain the lowest ratio-tracking temperature coefficient possible. It is important to keep this coefficient low because the gain of the ADC is dependent on the ratio of the input resistor to the runup reference slope resistors. An overall temperature coefficient of 0.4 ppm/°C was achieved for the ADC. Even at this level, a temperature change of 0.1°C results in a five-count change in a full-scale 8½-digit measurement. (Autocalibration increases the gain stability to greater than 0.15 ppm/°C.)

Another requirement for the resistor network is to have a low enough absolute temperature coefficient that nonlinearities are not introduced by the self-heating of the resistors. For example, the 50-kΩ input resistor has a input voltage that ranges from +12V to −12V. There is a 2.88-milliwatt power difference between a 0V input and a 12V input. If this power difference causes the resistor to change its value, the result is a nonlinearity in the ADC. A 0.01°C temperature change in a resistor that has an absolute temperature coefficient of 1 ppm/°C causes a one-count error in an 8½-digit measurement. The network used in the HP 3458A's ADC shows no measurable self-heating nonlinearities.

The final requirement of the resistor network is that it maintain the ratios of the six slopes throughout the life of the HP3458A. The tightest ratio tolerance is approximately 0.1% and is required to maintain linearity of the high-speed measurements. This is a relatively easy requirement. To maintain the ADC's 8½-digit differential linearity at less than 0.02 ppm requires ratio tolerances of only 3%.

Switches

A last major concern for the ADC design was the switches required to control the inputs and the slopes. Because the switches are in series with the resistors, they can add to the temperature coefficient of the ADC. A custom chip design was chosen so that each switch could be scaled to the size of the resistor to which it is connected. This allows the ADC to be sensitive to the ratio-tracking temperature coefficient of the switches and not to the absolute temperature coefficient. Another advantage of the custom design is that it allows the control signals to be latched just before the drives to the switches. This resynchronizes the signal with the clock and reduces the timing jitter in the switch transitions. The result is a reduction in the noise of the ADC.

Performance

The performance of an ADC is limited by several nonideal behaviors. Often the stated resolution of an ADC is limited by differential linearity or noise even though the number of counts generated would indicate much finer resolution. For example, the HP 3458A's ADC generates more than 9½ digits of counts but is only rated at 8½ digits because the ninth digit is very noisy and the differential linearity is about one eight-digit count. Therefore, when stating an ADC's speed and resolution, it is important to specify under what conditions the parameters are valid. Fig. 9 shows the speed-versus-resolution relationship of the HP 3458A ADC assuming less than one count of rms noise.

Given a noise level, there is a theoretical limit to the resolution of an ADC for a given speed. It can be shown that the white noise bandwidth of a signal that is the output of an integration over time T is

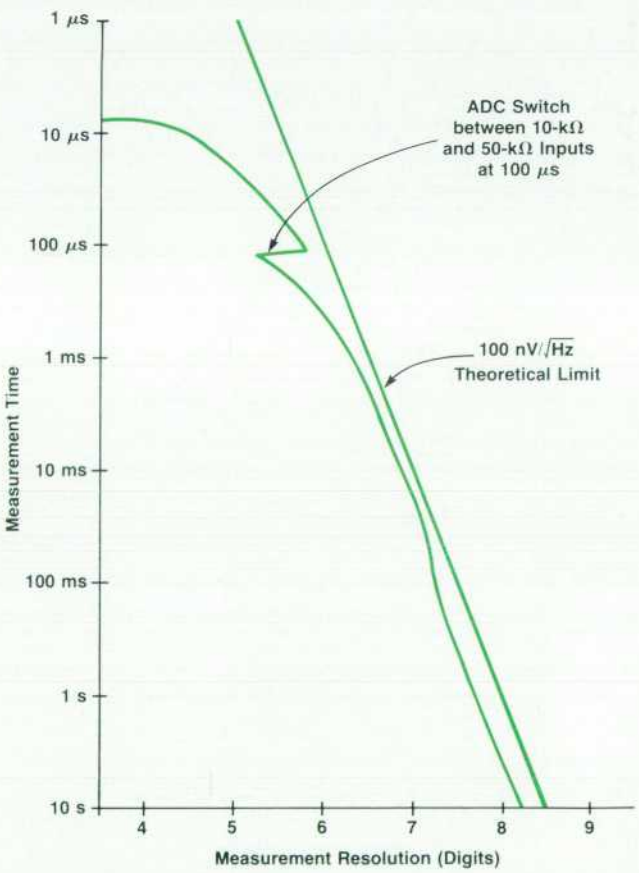

Fig. 9. *HP 3458A ADC speed versus resolution for one count of rms noise.*

$$BW = 1/2T.$$

If rundown took zero time then an integrating ADC could sample once every T seconds. At this rate. the counts of resolution, M, of an ADC are noise-limited to

$$M = (V_{fs}\sqrt{2T})/V_n,$$

where V_{fs} is the full-scale input voltage to the ADC and V_n is the white noise of the ADC in $V/\sqrt{Hz}$. Fig. 9 shows the best theoretical resolution for an ADC with rms noise of 100 $nV/\sqrt{Hz}$ and a full-scale input of 10 volts. The HP 3458A comes very close to the theoretical limit of an ADC with a white noise of 130 $nV/\sqrt{Hz}$ near the 7-digit resolution region. At lower resolutions the ADC's rundown time becomes a more significant portion of the overall measurement time and therefore pulls the ADC away from the theoretical limit. At higher resolutions the 1/f noise of the ADC forces a measurement of zero several times within the measurement cycle to reduce the noise to the 8½-digit level. This also reduces the measurement speed.

Another way of viewing the ADC's performance is to plot resolution versus aperture. The aperture is the integration time, that is, the length of runup. This is shown in Fig. 10 along with the 100-$nV/\sqrt{Hz}$ noise limit and the ADC's resolution without regard to noise. At smaller apertures, the HP 3458A's resolution is less than the theoretical noise limit because it is limited by noise in detecting the final zero of rundown. That is, the algorithm does not have enough resolution to achieve the theoretical resolution.

Linearity

High-resolution linearity was one of the major challenges of the ADC design. The autocalibration technique requires an integral linearity of 0.1 ppm and an differential linearity of 0.02 ppm. One of the more significant problems was verifying the integral linearity. The most linear commercially available device we could find was a Kelvin-Varley divider, and its best specification was 0.1 ppm of input. Fig. 11 compares this with the ADC's requirements, showing that it is not adequate.

Using low-thermal-EMF switches, any even-ordered deviations from an ideal straight line can be detected by doing a turnover test. A turnover test consists of three steps: (1) measure and remove any offset, (2) measure a voltage, and (3) switch the polarity of the voltage (i.e., turn the voltage over) and remeasure it. Any even-order errors will produce a difference in the magnitude of the two nonzero voltages measured. Measurements of this type can be made to within 0.01 ppm of a 10V signal. This left us with only the odd-order errors to detect. Fortunately, the U.S. National Bureau of Standards had developed a Josephson junction array capable of generating voltages from −10V to +10V. Using a 10V array we were able to measure both even-order and odd-order errors with confidence to a few hundredths of a ppm. Fig. 4a on page 23 shows the integral linearity error of an HP 3458A measured using a Josephson junction array.

The differential linearity can be best seen by looking at a small interval about zero volts. Here a variable source need only be linear within 1 ppm on its 100-mV range to produce an output that is within 0.01 ppm of 10 volts. Fig. 4b on page 23 shows the differential linearity of an HP 3458A.

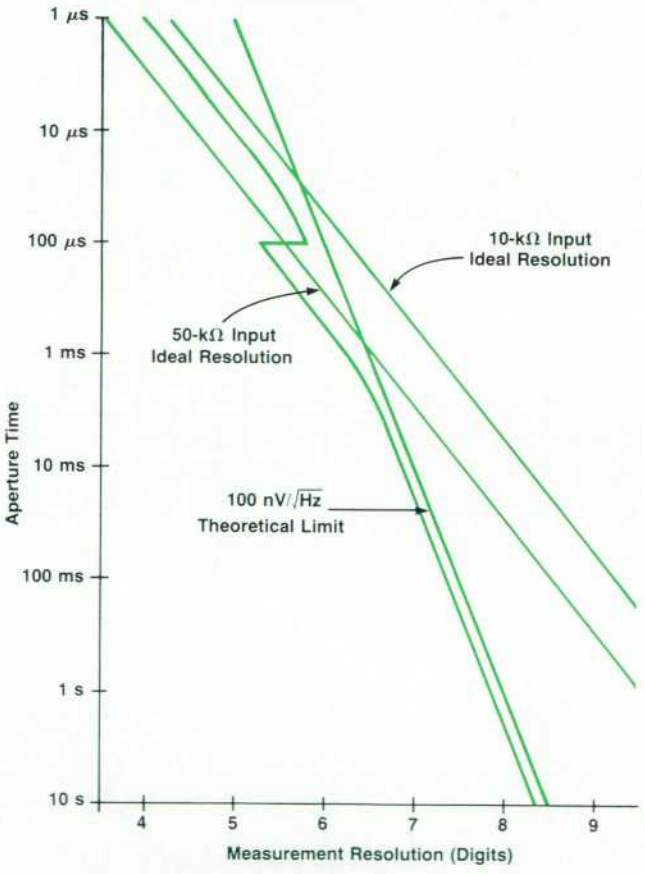

Fig. 10. *HP 3458A ADC aperture (runup time) versus resolution.*

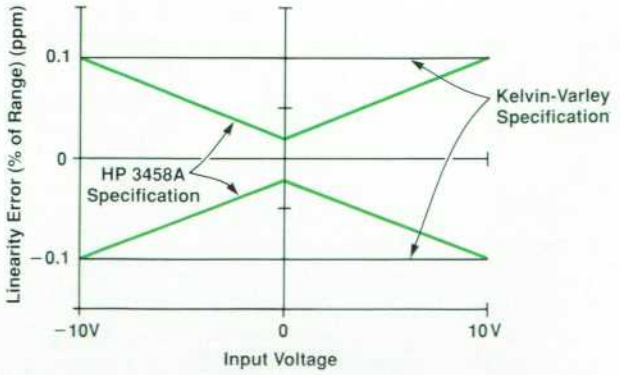

Fig. 11. *HP 3458A linearity specification compared with a Kelvin-Varley divider.*

Acknowledgments

The evolution of multislope ADC technology started with the development of the HP 3455A DMM over 15 years ago and has been refined by many designers working on several projects. Through the years Al Gookin, Joe Marriott, Larry Jones, James Ressmeyer, and Larry DesJardin have made significant contributions to multislope ADC concepts. I also want to thank John Roe of Hewlett-Packard's Colorado Integrated Circuits Division for his efforts in the design of the custom switch chip, David Czenkusch and David Rustici for the gate array development, and Steve Venzke for his valuable input to the integrator design. I would like to commend Clark Hamilton and Dick Harris of the U.S. National Bureau of Standards for developing the Josephson junction array voltage system that contributed to the refinement of the ADC technology developed for the HP 3458A.

Precision AC Voltage Measurements Using Digital Sampling Techniques

Instead of traditional DMM techniques such as thermal conversion or analog computation, the HP 3458A DMM measures rms ac voltages by sampling the input signal and computing the rms value digitally in real time. Track-and-hold circuit performance is critical to the accuracy of the method.

by Ronald L. Swerlein

THE HP 3458A DIGITAL MULTIMETER implements a digital method for the precise measurement of rms ac voltages. A technique similar to that of a modern digitizing oscilloscope is used to sample the input voltage waveform. The rms value of the data is computed in real time to produce the final measurement result. The HP 3458A objectives for high-precision digital ac measurements required the development of both new measurement algorithms and a track-and-hold circuit capable of fulfilling these needs.

Limitations of Conventional Techniques

All methods for making ac rms measurements tend to have various performance limitations. Depending on the needs of the measurement, these limitations take on different levels of importance.

Perhaps the most basic specification of performance is accuracy. For ac measurements, accuracy has to be specified over a frequency band. Usually, the best accuracy is for sine waves at midband frequencies (typically 1 kHz to 20 kHz). Low-frequency accuracy usually refers to frequencies below 200 Hz (some techniques can work down to 1 Hz). Bandwidth is a measure of the technique's performance at higher frequencies.

Linearity is another measure of accuracy. Linearity is a measure of how much the measurement accuracy changes when the applied signal voltage changes. In general, linearity is a function of frequency just as accuracy is, and can be included in the accuracy specifications. For instance, the accuracy at 1 kHz may be specified as 0.02% of reading + 0.01% of range while the accuracy at 100 kHz may be specified as 0.1% of reading + 0.1% of range. The percent-of-range part of the specification is where most of the linearity error is found.

If a nonsinusoid is being measured, most ac rms measurement techniques exhibit additional error. Crest-factor error is one way to characterize this performance. Crest factor is defined as the ratio of the peak value of a waveform to its rms value. For example, a sine wave has a crest factor of 1.4 and a pulse train with a duty cycle of 1/25 has a crest factor of 5. Even when crest factor error is specified, one should use caution when applying this additional error if it is not given as a function of frequency. A signal with a moderately high crest factor may have significant frequency components at 40,000 times the fundamental frequency. Thus crest factor error should be coupled with bandwidth information in estimating the accuracy of a measurement. In some ac voltmeters, crest factor specifications mean only that the voltmeter's internal amplifiers will remain unsaturated with this signal, and the accuracy for nonsinusoids may actually be unspecified.

Some of the secondary performance specifications for rms measurements are short-term reading stability, settling time, and reading rate. These parameters may have primary importance, however, depending on the need of the measurement. Short-term stability is self-explanatory, but the

difference between settling time and reading rate is sometimes confusing. Settling time is usually specified as the time that one should wait after a full-scale signal amplitude change before accepting a reading as having full accuracy. Reading rate is the rate at which readings can be taken. Its possible for an ac voltmeter that has a one-second settling time to be able to take more than 300 readings per second. Of course, after a full-scale signal swing, the next 299 readings would have degraded accuracy. But if the input signal swing is smaller than full-scale, the settling time to specified accuracy is faster. Therefore, in some situations, the 300 readings/second capability is actually useful even though the settling time is one second.

The traditional methods for measuring ac rms voltage are thermal conversion and analog computation. The basis of thermal conversion is that the heat generated in a resistive element is proportional to the square of the rms voltage applied to the element. A thermocouple is used to measure this generated heat. Thermal conversion can be highly accurate with both sine waves and waveforms of higher crest factor. Indeed, this accuracy is part of the reason that the U.S. National Institute of Standards and Technology (formerly the National Bureau of Standards) uses this method to supply ac voltage traceability. It can also be used at very high frequencies (in the hundreds of MHz). But thermal conversion tends to be slow (near one minute per reading) and tends to exhibit degraded performance at low frequencies (below 20 Hz). The other major limitation of thermal conversion is dynamic range. Low output voltage, low thermal coupling to the ambient environment, and other factors limit this technique to a dynamic range of around 10 dB. This compares to the greater than 20 dB range typical of other techniques.

Analog computation is the other common technology used for rms measurements. Essentially, the analog converter uses logging and antilogging circuitry to implement an analog computer that calculates the squares and square roots involved in an rms measurement. Since the rms averaging is implemented using electronic filters (instead of the physical thermal mass of the thermal converter), analog computation is very flexible in terms of reading rate. This flexibility is the reason that this technique is offered in the HP 3458A Multimeter as an ATE-optimized ac measurement function (SETACV ANA). Switchable filters offer settling times as fast as 0.01 second for frequencies above 10 kHz. With such a filter, reading rates up to 1000 readings/second may be useful.

Analog computation does have some severe accuracy drawbacks, however. It can be very accurate in the midband audio range, but both its accuracy and its linearity tend to suffer severe degradations at higher frequencies. Also, the emitter resistances of the transistors commonly used to implement the logging and antilogging functions tend to cause errors that are crest-factor dependent.

Digital AC Technique

Digital ac is another way to measure the rms value of a signal. The signal is sampled by an analog-to-digital converter (ADC) at greater than the signal's Nyquist rate to avoid aliasing errors. A digital computer is then used to compute the rms value of the applied signal. Digital ac can exhibit excellent linearity that doesn't degrade at high frequencies as analog ac computation does. Accuracy with all waveforms is comparable to thermal techniques without their long settling times. It is possible to measure low frequencies faster and with better accuracy than other methods using digital ac measurements. Also, the technique lends itself to autocalibration of both gain and frequency response errors using only an external dc voltage standard (see article, page 22).

In its basic form, a digital rms voltmeter would sample the input waveform with an ADC at a fast enough rate to avoid aliasing errors. The sampled voltage points (in the form of digital data) would then be operated on by an rms estimation algorithm. One example is shown below:

Num	= number of digital samples
Sum	= sum of digital data
Sumsq	= sum of squares of digital data
ac rms	= SQR((Sumsq−Sum*Sum/Num)/Num)

Conceptually, digital rms estimation has many potential advantages that can be exploited in a digital multimeter (DMM). Accuracy, linearity over the measurement range, frequency response, short-term reading stability, and crest factor performance can all be excellent and limited only by the errors of the ADC. The performance limitations of digital ac are unknown at the present time since ADC technology is continually improving.

Reading rates can be as fast as theoretically possible because ideal averaging filters can be implemented in firmware. Low-frequency settling time can be improved by measuring the period of the input waveform and sampling only over integral numbers of periods. This would allow a 1-Hz waveform to be measured in only two seconds—one second to measure the period and one second to sample the waveform.

Synchronous Subsampling

A thermal converter can measure ac voltages in the frequency band of 20 Hz to 10 MHz with state-of-the-art accuracy. Sampling rates near 50 MHz are required to measure these same frequencies digitally, but present ADCs that can sample at this rate have far less linearity and stability than is required for state-of-the-art accuracy in the audio band. If the restriction is made that the signal being measured must be repetitive, however, a track-and-hold circuit can be used ahead of a slower ADC with higher stability to create an ADC that can effectively sample at a much higher rate. The terms "effective time sampling," "equivalent time

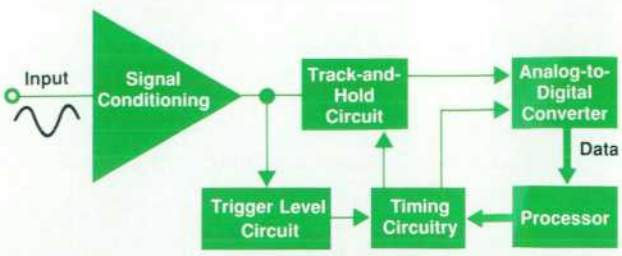

Fig. 1. *Simplified block diagram of a subsampling ac voltmeter.*

sampling," and "subsampling" are used interchangeably to describe this technique.

The concept of subsampling is used by digitizing oscilloscopes to extend their sample rate to well beyond the intrinsic speed of the ADC. The concept is to use a trigger level circuit to establish a time reference relative to a point on a repetitive input signal. A timing circuit, or time base, is used to select sample delays from this reference point in increments determined by the required effective sampling rate. For example, moving the sampling point in 10-ns increments corresponds to an effective sampling rate of 100 MHz. A block diagram of a subsampling ac voltmeter is shown in Fig. 1.

Fig. 2 is a simple graphic example of subsampling. Here we have an ADC that can sample at the rate of five samples for one period of the waveform being measured. We want to sample one period of this waveform at an effective rate of 20 samples per period. First, the timing circuit waits for a positive zero crossing and then takes a burst of five readings at its fastest sample rate. This is shown as "First Pass" in Fig. 2. On a subsequent positive slope, the time base delays an amount of time equal to one fourth of the ADC's minimum time between samples and again takes a burst of five readings. This is shown as "Second Pass." This continues through the fourth pass, at which time the applied repetitive waveform has been equivalent time sampled as if the ADC could acquire data at a rate four times faster than it actually can.

The digital ac measurement technique of the HP 3458A is optimized for precision calibration laboratory measurements. Short-term measurement stability better than 1 ppm has been demonstrated. Absolute accuracy better than 100 ppm has been shown. This accuracy is achieved by automatic internal adjustment relative to an external 10V dc standard. No ac source is required (see article, page 22). The internal adjustments have the added benefit of providing a quick, independent check of the voltage ratios and transfers that are typically performed in a standards laboratory every day. Fast, accurate 1-Hz measurements and superb performance with nonsinusoids allow calibration laboratories to make measurements easily that were previously very difficult.

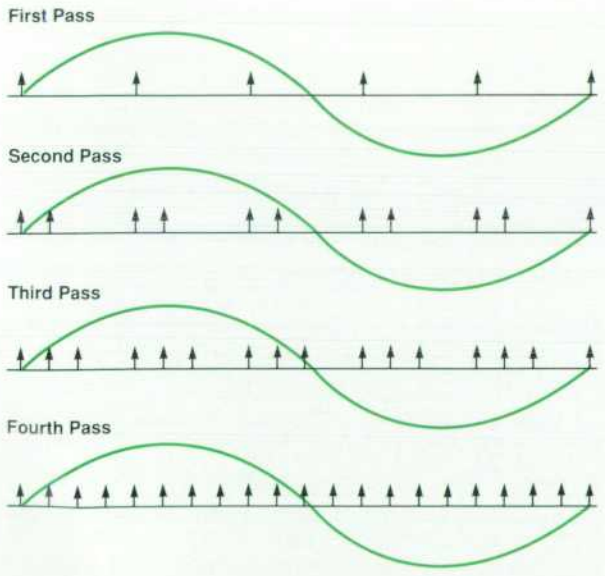

Fig. 2. *An example of subsampling.*

The HP 3458A enters into the synchronously subsampled ac mode through the command SETACV SYNC. For optimal sampling of the input signal, one must determine the period of the signal, the number of samples required, and the signal bandwidth. The measurement resolution desired and the potential bandwidth of the input waveform are described using the commands RES and ACBAND. The period of the input signal is measured by the instrument. The more the HP 3458A knows about the bandwidth of the input and the required measurement resolution, the better the job it can do of optimizing accuracy and reading rate. Default values are assumed if the user chooses not to enter more complete information. An ac measurement using the SYNC mode appears to function almost exactly the same to the user as one made using the more conventional analog mode.

Subsampled AC Algorithm

The algorithm applied internally by the HP 3458A during each subsampled ac measurement is totally invisible to the user. The first part of a subsampled ac measurement is autolevel. The input waveform is randomly sampled for a period of time long enough to get an idea of its minimum and maximum voltage points. This time is at least one cycle of the lowest expected frequency value (the low-frequency value of ACBAND). The trigger level is then set to a point midway between the minimum and maximum voltages, a good triggering point for most waveforms. In the unlikely event that this triggering point does not generate a reliable trigger, provision is made for the user to generate a trigger signal and apply it to an external trigger input. An example of such a waveform is a video signal. Even though video signals can be repetitive, they are difficult to trigger on correctly with just a standard trigger level.

With the trigger level determined, the period of the input waveform is measured. The measured period is used along with the global parameter RES to determine subsampling parameters. These parameters are used by the timing circuitry in the HP 3458A to select the effective sample rate, the number of samples, and the order in which these samples are to be taken. In general, the HP 3458A tries to sample at the highest effective sample rate consistent with meeting the twin constraints of subsampling over an integral number of input waveform periods and restricting the total number of samples to a minimum value large enough to meet the specified resolution. This pushes the frequency where aliasing may occur as high as possible and also performs the best rms measurement of arbitrary waveforms of high crest factor. The number of samples taken will lie somewhere between 4/RES and 8/RES depending on the measured period of the input waveform.

The final step is to acquire samples. As samples are taken, the data is processed in real time at a rate of up to 50,000 samples per second to compute a sum of the readings squared and a sum of the readings. After all the samples are taken, the two sum registers are used to determine standard deviation (ACV function), or rms value (ACDCV

function). For example, suppose a 1-kHz waveform is being measured and the specified measurement resolution is 0.001%. When triggered, the HP 3458A will take 400,000 samples using an effective sample rate of 100 MHz. The timing circuit waits for a positive-slope trigger level. Then, after a small fixed delay, it takes a burst of 200 readings spaced 20 μs apart. It waits for another trigger, and when this occurs the timing circuit adds 10 ns to the previous delay before starting another burst of 200 readings. This is repeated 2,000 times, generating 400,000 samples. Effectively, four periods of the 1-kHz signal are sampled with samples placed every 10 ns.

Sources of Error

Internal time base jitter and trigger jitter during subsampling contribute measurement uncertainty to the rms measurement. The magnitude of this uncertainty depends on the magnitude of these timing errors. The internal time base jitter in the HP 3458A is less than 100 ps rms. Trigger jitter is dependent on the input signal's amplitude and frequency, since internal noise will create greater time uncertainties for slow-slew-rate signals than for faster ones. A readily achievable trigger jitter is 100 ps rms for a 1-MHz input. Fig. 3 is a plot generated by mathematical modeling of the performance of a 400,000-sample ac measurement using the HP 3458A's subsampling algorithm (RES = 0.001%) in the presence of 100-ps time base and trigger jitters. The modeled errors suggest the possibility of stable and accurate ac measurements with better than 6-digit accuracy.

Errors other than time jitter and trigger jitter limit the typical absolute accuracy of the HP 3458A to about 50 ppm, but there is reason to believe that short-term stability is better than 1 ppm. Many five-minute stability tests using a Datron 4200 AC Calibrator show reading-to-reading standard deviations between 0.7 ppm and 3 ppm. Other measurements independently show the Datron 4200 to have similar short-term stability. More recently, tests performed using a Fluke 5700 Calibrator, which uses a theoretically quieter leveling loop, show standard deviations under 0.6 ppm.

The above algorithm tries to sample the applied signal over an integral number of periods. To do this, the period of the signal must first be measured. Errors in measuring the period of the input waveform will cause the subsequent sampling to cover more or less than an integral number of periods. Thus, the accuracy of the subsampled ac rms measurement is directly related to how accurately the period of the input waveform is known relative to the internal sample time base clock.

Period measurements in the HP 3458A are performed using reciprocal frequency counting techniques. This method allows accuracy to be traded off for measurement speed by selecting different gate times. A shorter gate time contributes to a faster measurement, but the lower accuracy of the period determination contributes to a less accurate ac measurement. Fig. 4 is a graph of the error introduced into the rms measurement by various gate times. At high frequencies, this error is a constant dependent on the resolution of the frequency counter for a given gate time. At lower frequencies, trigger time jitter increases, causing increased error, because random noise has a larger effect on slower signals. At still lower frequencies, where the period being measured is longer than the selected gate time, this error becomes constant again. This is because the gate time is always at least one period in length, and as the frequency is lowered, the gate time increases just fast enough to cancel the effect of increasing trigger jitter.

Any violation of the restriction that the input waveform be repetitive will also lead to errors. A common condition is amplitude and frequency modulation of the input. If this modulation is of a fairly small magnitude and is fast compared to the total measurement time this violation of the repetitive requirement will probably be negligible. At most, reading-to-reading variation might increase slightly. If these modulations become large, however, subsampled ac accuracy can be seriously compromised. The signal sources typically present on a lab bench or in a calibration laboratory work quite well with the subsampling algorithm of the HP 3458A.

Random noise spikes superimposed on an input can make an otherwise repetitive input waveform appear nonrepetitive. Induced current caused by motors and electrical devices turning on and off is just one of many ways to generate such spikes. Large test systems tend to generate more of this than bench and calibration laboratory environments. Low-voltage input signals (below 100 mV) at low

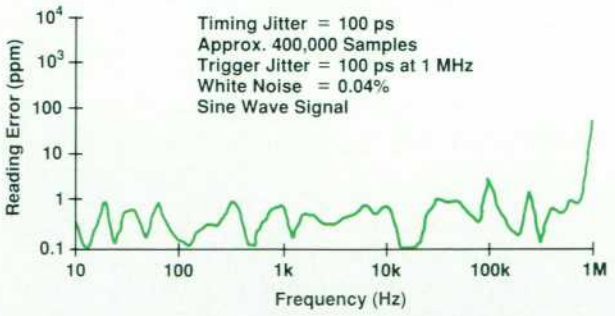

Fig. 3. *Subsampling errors resulting from timing uncertainties.*

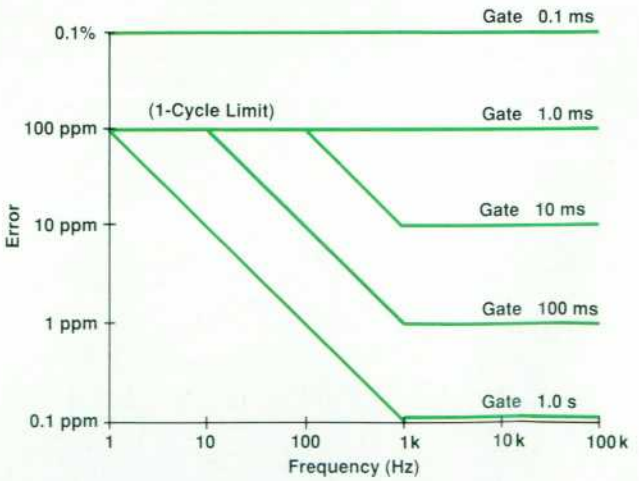

Fig. 4. *Subsampling error as a function of gate time.*

frequencies are the signals most susceptible to these errors.

Two ways are provided by the HP 3458A to deal with these potential errors. The first is to use the internal 80-kHz low-pass trigger filter to reduce high-frequency trigger noise (LFILTER ON). If this is not enough, provision is made for accepting external synchronization pulses. In principle, getting subsampled ac to work in a noisy environment is no more difficult than getting a frequency counter to work in the same environment.

If nonsinusoidal signals are being measured, the subsampling algorithm has some additional random errors that become greater for signals of large crest factor. All nonsinusoidal repetitive signals have some of their spectral energy at frequencies higher than their fundamental frequency. Signals of high crest factor generally have more of this high-frequency energy than those of lower crest factor. Random timing jitter, which tends to affect higher frequencies the most, will create measurement errors that are greater for large-crest-factor signals. These random errors can be reduced by specifying a higher-resolution measurement, which forces more samples per reading to be acquired. The additional measurement error induced by a signal crest factor of 5 can be as low as 50 ppm in the HP 3458A.

Track-and-Hold Circuit

Track-and-hold performance is critical to the accuracy of digital ac measurements. Track-and-hold linearity, bandwidth, frequency flatness, and aperture jitter all affect the error in a sampled measurement. To meet the HP 3458A performance objectives, track-and-hold frequency response flatness of ±0.0015% (15 ppm) was required from dc to 50 kHz, along with a 3-dB bandwidth of 15 MHz. In addition, 16-bit linearity below 50 kHz and low aperture jitter were needed. A custom track-and-hold amplifier was developed to meet these requirements.

The most basic implementation of a track-and-hold circuit—a switch and a capacitor—is shown in Fig. 5. If the assumption is made that the switch is perfect (when open it has infinite impedance and when closed it has zero impedance) and if it assumed that the capacitor is perfect (no dielectric absorption), then this is a perfect track-and-hold circuit. The voltage on the capacitor will track the input signal perfectly in track mode, and when the switch is opened, the capacitor will hold its value until the switch is closed. Also, as long as the buffer amplifier's input impedance is high and well-behaved, its bandwidth can be much lower than the bandwidth of the signal being sampled. When the switch is opened, the buffer amplifier's output might not have been keeping up with the input signal, but since the voltage at the input of the amplifier is now static, the buffer will eventually settle out to the hold capacitor's voltage.

The problem with building Fig. 5 is that it is impossible at the present time to build a perfect switch. When the switch is opened it is not truly turned off; it has some residual leakage capacitance and resistance. In hold mode, there is some residual coupling to the input signal because of this leakage capacitance. This error term is commonly called feedthrough. Another error term is pedestal voltage. The process of turning real-world switches off induces a charge transfer that causes the hold capacitor to experience a fixed voltage step (a pedestal) when entering hold mode.

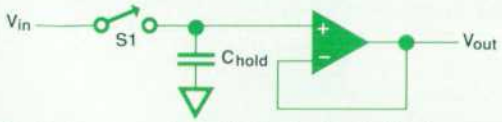

Fig. 5. *Basic track-and-hold circuit with ideal components.*

Another problem with Fig. 5 is that it is impossible in the real world to build a perfect capacitor. Real-world capacitors have nonideal behaviors because of dielectric absorption and other factors. This dielectric absorption will manifest itself as a pedestal that is different for different input-voltage slew rates. Even if the capacitor is refined until it is "perfect enough," the switch and the buffer amplifier may contribute enough capacitance in parallel with C_{hold} that the resultant capacitance has dielectric absorption problems.

Fig. 6 is an implementation of Fig. 5 using real-world components. The switch is implemented with a p-channel MOS FET. When the drive voltage is −15V, the circuit is in track mode. If the FET has an on resistance of R, then the 3-dB bandwidth of the circuit is $1/(2\pi RC_{hold})$. C_{dg} (the drain-to-gate capacitance) is always in parallel with C_{hold}, so even if C_{hold} and the buffer amplifier have low dielectric absorption, the dielectric absorption associated with C_{dg} will cause this circuit to exhibit pedestal changes with different input signal slew rates.

When the drive voltage is changed to +15V, the FET turns off and puts the circuit into hold mode. The drain-to-source capacitance (C_{ds}) contributes feedthrough error equal to C_{ds}/C_{hold}. If the drive voltage changes infinitely fast, the pedestal error is $(30V)(C_{dg}/C_{hold})$. If the drive voltage changes at a slower rate, the pedestal error will be less, but a gain error term will now appear. Assume that the drive voltage changes slowly relative to the bandwidth of the track-and-hold circuit $(1/(2\pi RC_{hold}))$. Assume also that the FET is on until the drive voltage is equal to V_{in} and that it is off when the drive voltage is greater than V_{in}. The process of going into hold mode begins with the drive voltage changing from −15V to +15V. As the voltage changes from −15V to V_{in}, C_{hold} experiences very little pedestal error since the current $C_{dg}(dv/dt)$ mostly flows into the FET, which is on. When the drive voltage reaches V_{in}, the FET turns off and all of the $C_{dg}(dv/dt)$ current flows into C_{hold}. The pedestal in this case is $(15V - V_{in})(C_{dg}/C_{hold})$. Notice that this is a smaller pedestal than in the previous case where the drive voltage changed infinitely fast. Also notice that there is a V_{in} term in the pedestal equation. This is a gain error.

Pedestal errors are easy to deal with in the real world. There are a number of easy ways to remove offset errors.

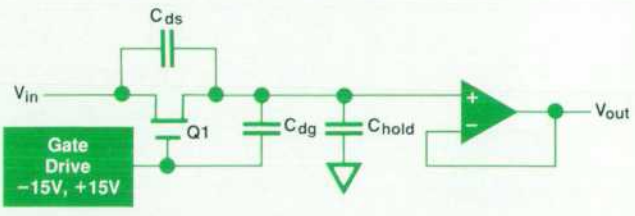

Fig. 6. *Basic track-and-hold circuit with real components.*

Gain errors are not necessarily bad either, since ideal gain errors can be corrected with a compensating gain stage. But because C_{dg} is a semiconductor capacitance, it tends to change value with the applied voltage. This leads to a form of error called nonlinearity. In general, gain errors that are caused by semiconductor capacitances (like that calculated in the above paragraph) lead to nonlinearity errors. A track-and-hold circuit application that is affected by nonlinearity errors is sampling a signal and calculating its Fourier transform. Feedthrough and dielectric absorption errors also are hard to deal with. Commonly, a different track-and-hold architecture is used to achieve better linearity, feedthrough, and dielectric absorption performance.

Fig. 7 is a diagram of the track-and-hold architecture used most often to achieve 16-bit or better resolution along with 2-MHz bandwidths. In track mode the drive voltage is −15V, turning Q1 on. The output voltage is the inverse of the input voltage. The inverse of the input voltage is impressed across C_{hold} during track mode. When the drive voltage is changed to +15V, Q1 turns off and V_{out} is held.

Fig. 7 has several advantages over Fig. 6. Since the switch (Q1) is at a virtual ground point, the pedestal voltage is constant with V_{in} and equal to $(30V)(C_{dg}/C_{hold})$. This is because the drain and source are always at zero so that when Q1 is turned off the same amount of charge is always transferred to C_{hold}. Also, since no point on Q1 moves with V_{in}, the FET does not contribute any dielectric absorption error terms.

Fig. 7 does have feedthrough error. It is equal to $\frac{1}{2}(C_{ds}/C_{hold})$. Theoretically this error could be substantially eliminated if a second switch could be turned on after entering hold mode to ground the junction of the two resistors. However, a real drawback of this circuit is that the op amp U1 has to have the same bandwidth and slew rate capabilities as the signal being sampled. In the descriptions of Figs. 5 and 6 it was mentioned that the buffer amplifier need not have the same bandwidth as the signal being sampled. So in summary, Fig. 7 eliminates some of the errors of the previous circuits but introduces at least one new limitation.

HP 3458A Track-and-Hold Architecture

Fig. 8 is a modification of Fig. 6 that has most of the advantages and very few of the disadvantages of the previous circuits. Here the switch is implemented with two n-channel JFETs and one p-channel MOS FET. In track mode the JFETs Q1 and Q2 are on and the MOS FET Q3 is off. Q1 and Q2 are on because their gate-to-source voltages are zero, since their gates track V_{in}. Their gates track V_{in} because in track mode point B is an open circuit and CR1 and CR2 act like resistances of about 1 kΩ. CR1 and CR2 are current regulator diodes, which are simply JFETs with their gates wired to their sources. In hold mode, Q1 and Q2 are off and Q3 is on. Q1 is now off because point B is now at −15V and thus the gate of Q1 is at −15V. CR2 now appears as a current source of high resistance and the gate of Q2 is clamped at about 7V below V_{out}, turning off Q2. Q3 is on because its gate (point A) is at −15V.

In hold mode, feedthrough error is very low, since the feedthrough caused by C_{ds1} is shunted into the ac ground created by Q3's being on. Also, the pedestal error caused by C_{dg2} is constant for all V_{in}, since the gate of Q2 is clamped at 7V below V_{out}. Since V_{out} is tracking V_{in} during track mode (or will settle out to V_{in} after hold mode is entered), the pedestal error caused by C_{dg2} is $(-7V)(C_{dg2}/C_{hold})$ and has no V_{in} dependent terms. Therefore it makes no difference to the linearity errors of the track-and-hold circuit whether C_{dg2} is nonlinear with bias voltage.

It is not so obvious that C_{ds2} contributes almost nothing to the pedestal errors and the nonlinearity errors of the circuit. In addition to being a T-switch that reduces feedthrough errors in hold mode, Q1, Q2, and Q3 when switched in the correct sequence act to remove almost all of the pedestal errors caused by C_{ds2}. This is very important, since C_{ds2} is nonlinear, and if its pedestal errors remained, the linearity of the circuit would be no better than that of Fig. 6. Q1 is selected such that its pinchoff voltage (V_{gsoff}) is greater than that of Q2. Thus, as point B is driven to −15V, Q2 turns off before Q1. Once Q2 is off, the only coupling path to C_{hold} is through the capacitance C_{ds2}.

Fig. 9 shows the various waveforms present in the circuit. When Q1 is finally turned off, the voltage on C1 has a pedestal error of $(V_{in} - 15V)(C_{dg1}/C_1)$. This pedestal couples into C_{hold} through C_{ds2}. The magnitude is $(V_{in} - 15V)(C_{dg1}/C_1)(C_{ds2}/C_{hold})$. Since C_{dg1} is nonlinear and the coupling has a V_{in} dependent term, the pedestal on C_{hold} now has a nonlinear component. But after Q1 and Q2 are off, point A is driven to −15V, turning Q3 on. C_1 is now connected to V_{out} through the on resistance of Q3 and approaches the voltage V_{out}. This voltage movement, which

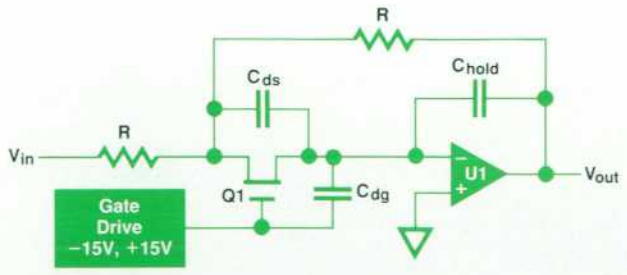

Fig. 7. *Conventional track-and-hold architecture.*

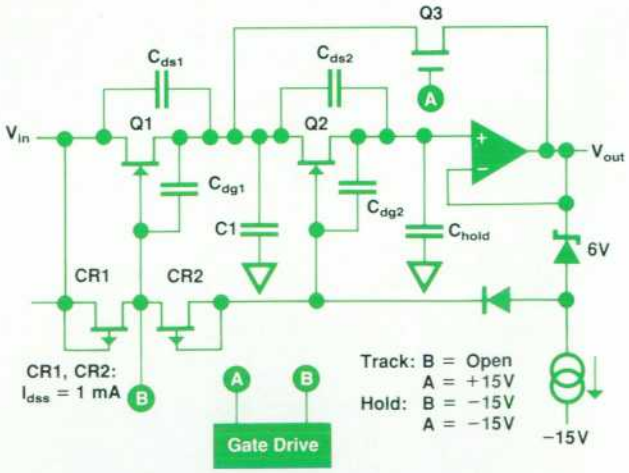

Fig. 8. *HP 3458A track-and-hold architecture.*

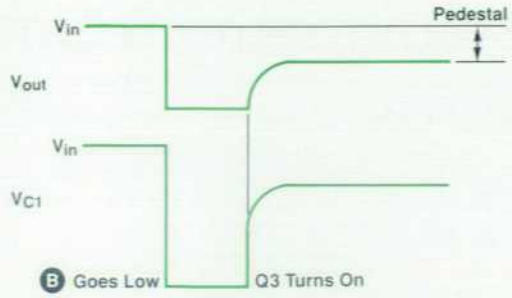

Fig. 9. *Waveforms in the circuit of Fig. 8.*

is of the same magnitude as C_1's previous change but in the opposite direction, couples into C_{hold} through C_{ds2} and totally removes the pedestal error previously coupled into C_{hold} through C_{ds2}.

Another point that also might not be so obvious is that Q1, Q2, and Q3 do not contribute any dielectric absorption errors to the track-and-hold circuit. Since in track mode the drains, sources, and gates of Q1 and Q2 are at the same potential (V_{in}), none of the FET capacitances has charge on it before it is put in hold mode. Therefore, the charge transferred to C_{hold} through the FET capacitances when hold mode is entered is the same for any value or slew rate of V_{in}, so it doesn't matter whether the FET capacitances have high dielectric absorption.

Summary

The performance of the HP 3458A with sinusoidal and nonsinusoidal inputs is known to be very good. The DMM was tested against a synthesized arbitrary waveform generator under development at the U.S. National Bureau of Standards which was capable of generating sine waves and ANSI-standard distorted sine waves with an absolute uncertainty of 10 ppm. The HP 3458A measured all of the various test waveforms with errors ranging from 5 ppm to 50 ppm for 7V rms inputs from 100 Hz to 10 kHz.

The digital ac measurement capability of the HP 3458A combines the best features of the traditional thermal and analog computational ac rms techniques in addition to adding several advantages of its own. Measurement accuracies for digital ac are comparable to thermal techniques for both sinusoidal (crest factor 1.4) and large-crest-factor nonsinusoidal waveforms. Like analog computation, digital ac reading rates are reasonably fast compared to thermal rms techniques. The major advantages of digital ac include linearity superior to traditional analog rms detection methods and significantly faster low-frequency rms ac measurements (less than six seconds for a 1-Hz input). Short-term reading stability is excellent, allowing previously difficult characterizations to be performed easily.

Acknowledgments

Credit should be given to Larry DesJardin for his direction and support for the development of the digital ac techniques implemented in the HP 3458A. I would also like to thank Barry Bell and Nile Oldham of the National Bureau of Standards for allowing access to their synthesized ac source to help validate the digital ac concept.

Calibration of an 8½-Digit Multimeter from Only Two External Standards

Internal transfer standards and autocalibration simplify external calibration and extend the period between external calibrations to two years.

by Wayne C. Goeke, Ronald L. Swerlein, Stephen B. Venzke, and Scott D. Stever

ONE OF THE EARLIEST PRODUCT CONCEPTS for the HP 3458A Digital Multimeter was to develop a means for calibrating its measurement accuracies from only two external reference standards. This is not possible with the traditional design for a DMM, which requires independent adjustment of the full-scale gain and zero offset for each measurement range and function.

Calibration is a process in which individual gain and offset values are adjusted, manually or electronically, to yield minimum error relative to an applied input, as shown in Fig. 1. Gain and offset calibration values are generally determined using precision ratio transfer measurements relative to a smaller set of working standards whose errors are directly traceable to national standards. In the United States, standards are kept by the National Institute of Standards and Technology (NIST), formerly the National Bureau of Standards (NBS). Dc voltages are often derived from a 1.018-volt saturated electrochemical cell known as a Weston standard cell. The output voltage is divided, or otherwise ratioed, to yield other traceable values. For example, the output would be divided by 10.18 to produce 0.1V. The ratio transfer process is, in general, different for each calibration value. It is prone to both random and systematic errors, which may propagate undetected into instrumentation through the calibration process. This calibration (or verification) uncertainty will produce a "floor" measurement error sometimes equal to or greater than the uncertainty of the instrument alone.

The objectives for two-source calibration are to reduce this floor uncertainty and to provide an independent method to increase confidence in the overall calibration process. The HP 3458A uses a highly linear analog-to-digital converter (ADC) to measure the ratio between a traceable reference and its divided output. The ADC performs the function of the precise ratio transfer device.

Sources of Error

The errors in any ratio measurement can be divided into two general types: differential errors (D) and integral errors (I). A differential error is a constant percent of full scale and is independent of the input. These errors are handled like dc offsets. An integral error is a function of the input, and the relationship is usually nonlinear. These errors are generally thought of as gain errors. The maximum total error can be expressed as:

$$E_1(x) = I(x/100\%) + D,$$

where x is the input to the ratio device and $E_1(x)$ is the error, both expressed as a percent of full scale. The general form of the error bound is shown in Fig. 2.

What is of concern is the error in the output or measured value expressed as a percent of that value. Expressed in this form, the maximum error is:

$$E_2(x) = I + D(100\%/x),$$

Where $E_2(x)$ is the total error in the output or measured value expressed as a percent of x. The general form of this error bound is shown in Fig. 3. For ratios less than one, the total error is dominated by the differential errors of the ratio transfer device. Since the differential error term is equal to the differential linearity error multiplied by one over the divider ratio, this error grows to infinity as the divider ratio grows smaller.

HP 3458A Uncertainty

The design goal for the HP 3458A DMM was for internal ratio transfer errors to be equal to or lower than those achievable with commercially available external ratio dividers. This set the total ratio measurement error (linearity) requirement for the ADC for a 10:1 transfer to approximately 0.5 ppm of output or 0.05 ppm of input.

Fig. 4 illustrates the integral and differential linearity achieved with the HP 3458A ADC design. The test data was generated using a Josephson junction array intrinsic voltage standard (see "Josephson Junction Arrays," page

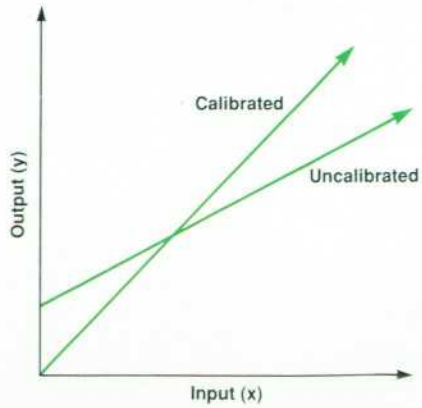

Fig. 1. *Calibrated and uncalibrated gain and offset in a measurement.*

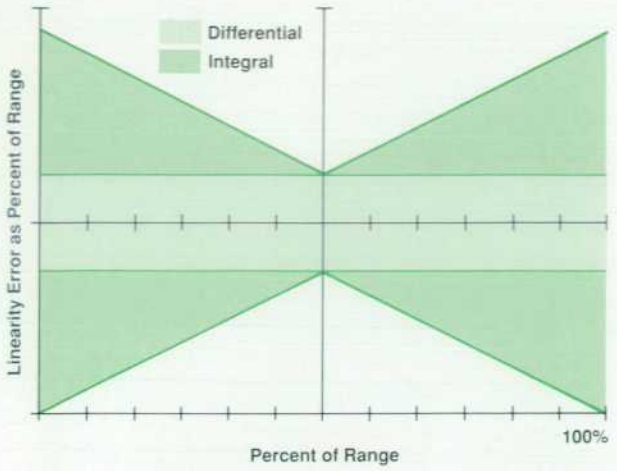

Fig. 2. *Linearity error as a percent of range.*

24). Fig. 4a shows typical deviation from a straight line for the input voltage range from minus full scale to plus full scale expressed in ppm of full scale. This expresses the test data in a form dominated by the integral linearity effects. Integral error less than 0.1 ppm of full scale was achieved. Fig. 4b shows typical test data expressed as ppm of reading (output). This data indicates differential linearity error less than 0.02 ppm of reading. For a 10:1 ratio transfer the predicted error would be approximately I + 10D or 0.3 ppm. Fig. 4c shows measured data, again using a Josephson junction array standard to characterize the error at 1/10 of full scale relative to a full-scale measured value. The data indicates a 10:1 ratio error of 0.01 ppm of the input or 0.1 ppm of the measured (output) value. This represents typical results; the specified 3σ ratio transfer error is greater than 0.3 ppm. Measurement noise contributes additional error, which can be combined in a root-sum-of-squares manner with the linearity errors.

Offset Errors

Linear measurement errors in a DMM are of two general types, offset errors and gain errors. Offset error sources include amplifier offset voltages, leakage current effects (IR), and thermocouple effects generated by dissimilar metals used in component construction or interconnection. Fig. 5 shows a simplified schematic of the dc measurement function. Switches S1 and S2 are used to provide a zero reference during each measurement cycle. Offset errors common to both measurement paths, for example the offset voltage introduced by amplifier A1, are sampled and subtracted during each measurement sequence. This is referred to as the autozero process.

Correction of the remaining offset error is achieved by

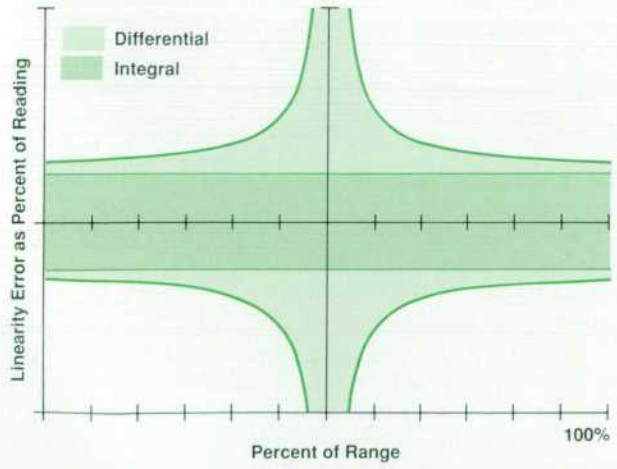

Fig. 3. *Linearity error as a percent of reading.*

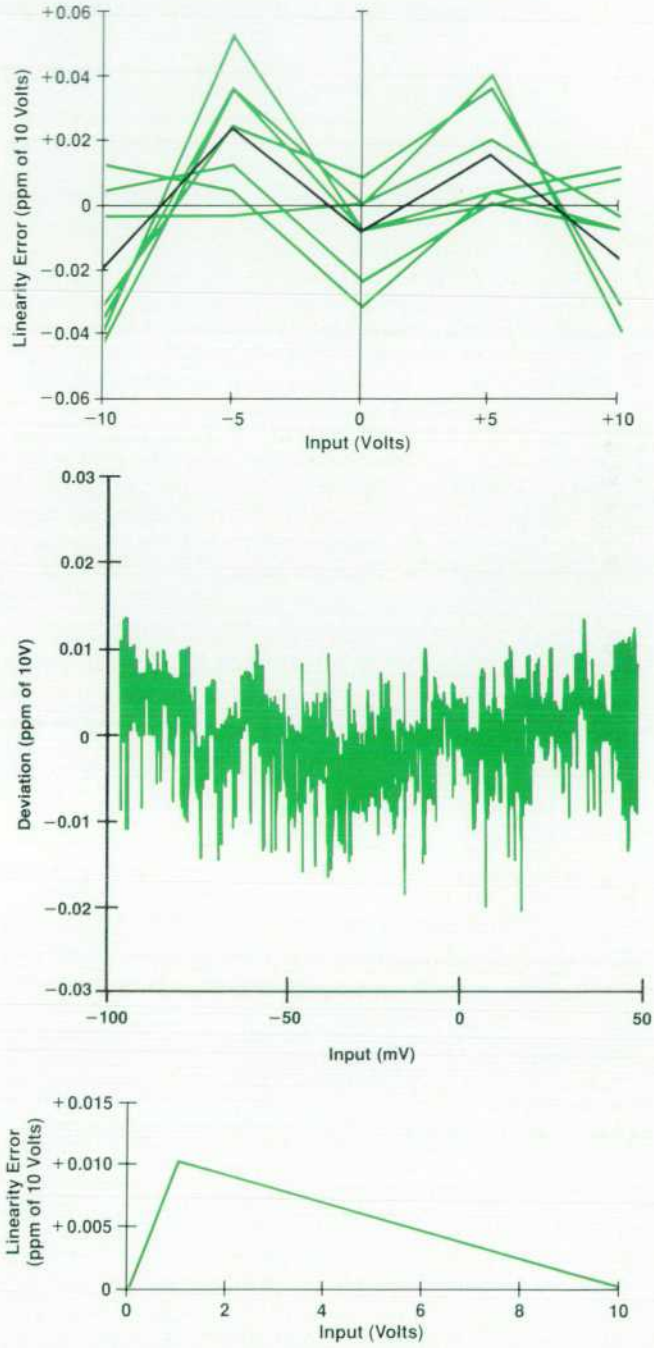

Fig. 4. *Results of HP 3458A linearity tests using a Josephson junction array. (a) Seven passes and the average result for linearity error characterization. (b) Differential linearity characteristic. (c) Linearity error for an internal 10:1 ratio transfer.*

Josephson Junction Arrays

A Josephson junction is formed by two superconductors separated by a thin insulating barrier. When cooled to liquid helium temperatures (4.2K), these devices exhibit very complex nonlinear behavior that has led to a wide range of applications in analog and digital electronics. A quantum mechanical analysis shows that these junctions generate an ac current whose frequency is related to the junction voltage by the relation $f = 2eV/h$ where e is the electron charge and h is Planck's constant. When the junction is driven by an ac current the effect operates in reverse. The junction oscillation phase locks to the applied ac current and the junction voltage locks to a value $V = hf/2e$. This phase locking can also occur between harmonics of the applied ac current and the Josephson oscillation. Thus, the junction I-V curve displays a set of constant-voltage steps (Fig. 1) at the voltages $V = nhf/2e$, where n is an integer. The Josephson junction thereby provides a means of translating the inherent accuracy of the frequency scale to voltage measurements.

In July of 1972 the Josephson effect was adopted as the definition of the U.S. legal volt. For the purpose of this definition the quantity 2e/h was assigned the value 483593.42 GHz/V. Since then, tests of the Josephson voltage-to-frequency relation have verified its precision and independence of experimental conditions to the level of a few parts in 10^{17}.[1]

The Josephson voltage standards of 1972 had only one or two junctions and could generate voltages only up to about 10 mV. This low voltage required the use of a complex voltage divider to calibrate the 1.018V standard cells used by most standards laboratories. To overcome the limitations of these low voltages,

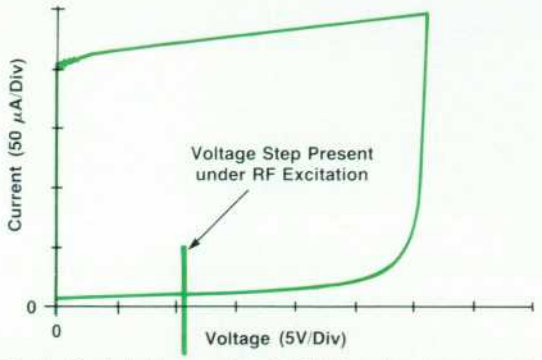

Fig. 1. *Partial I-V curve of an 18,992-junction Josephson junction array without RF excitation. Also shown is a typical I-V curve under 75-GHz excitation, which is a constant-voltage step at a voltage V = nhf/2e. The voltage V is between −12V and +12V, and is determined by controlling the bias current and source impedance to select the value of n.*

researchers at the U.S. National Institute of Standards and Technology (formerly the National Bureau of Standards), and PTB in West Germany have developed superconducting integrated circuits that combine the voltages of several thousand junctions. The most complex of these chips uses 18,992 junctions to generate 150,000 constant-voltage steps spanning the range from −12V to +12V (Fig. 2). The chip uses a finline to collect 75-GHz

providing a copper short across the input terminals. A reference measurement is taken and the measured offset is stored. Values are determined for each measurement function and range configuration. The offset is subtracted from all subsequent measurements. The HP 3458A performs all zero offset corrections by automatically sequencing through each of the required configurations and storing the appropriate offset correction during the external calibration process. These offsets are the b term in the linear equation $y = mx + b$, where y is the calibrated output result and x is the internal uncalibrated measurement. These calibrated offsets can be made small and stable through careful printed circuit layout and component selection.

Gain Errors

Gain errors in a DMM result from changes in amplifier gains, divider ratios, or internal reference voltages. Each gain term exhibits a temperature coefficient and some finite aging rate, and can potentially change value following exposure to high-humidity environments or severe shock or vibration. Periodically, known values close to the full scale of each measurement function and range are applied to the DMM to calibrate the gain ratio m such that $y = mx + b$ is precisely equal to the known input value, y. However, even after gain calibration, a DMM can easily be exposed to conditions that may introduce new errors. The HP 3458A DMM implements a special method for self-adjusting all instrument gain errors and many offset errors relative to its own internal references.

DC Calibration

Calibration of the dc function begins by establishing traceability of the internal voltage reference. The internal 7V Zener reference (see "A High-Stability Voltage Reference," page 28) is measured relative to an externally applied traceable standard. A traceable value for this internal reference is stored in secure calibration memory until the next external calibration is performed. Next, the gain of the 10V range is determined by measuring the internal 7V reference on this range. The gain value is stored in secure autocalibration memory. This gain value can be recomputed at any time by simply remeasuring the internal 7V reference. The stability, temperature coefficient, and time drift errors of the internal 7V reference are sufficiently small (and specified) compared with other gain errors that remeasurement or autocalibration of these gains will yield smaller measurement errors in all cases. Adjustment of the full-scale gain values of all other ranges relies on the precise ratio measurement capabilities of the HP 3458A ADC as demonstrated in Fig. 4c. For the 1V-range gain adjustment, the traceable internal 7V reference is divided to produce a nominal 1V output. The exact value of this nominal 1V is measured on the previously adjusted 10V measurement range at approximately 1/10 of full scale. The measured value, a ratio transfer from the internal 7V reference, is used to adjust the gain of the 1V range of the dc voltage function. This gain value is again stored in secure autocalibration memory. Neither the precise value nor the long-term stability of the nominal 1V internal source is important. The internal 1V

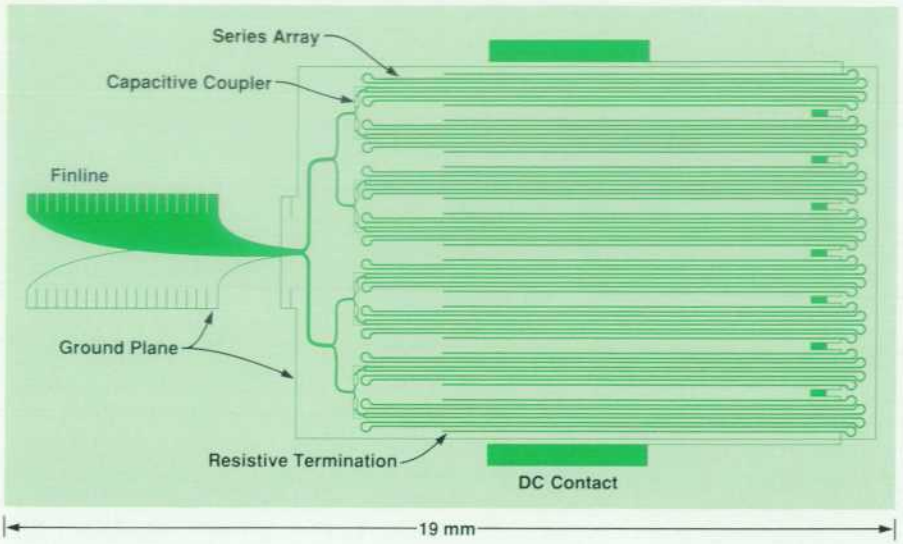

Fig. 2. *The layout for an 18,992-junction voltage standard array capable of generating voltage steps in the range of −12V to +12V. The horizontal lines represent 16 striplines, each of which passes through 1187 junctions. The junctions are too small to be distinguished on this drawing.*

power from a waveguide and direct it through a set of power splitters to 16 striplines, each of which passes through 1187 junctions. A network of high-pass and low-pass filters allows the microwave power to be applied in parallel while the dc voltages add in series.[2]

In operation, the array is cooled to 4.2K in a liquid-helium dewar. A Gunn-diode source at room temperature provides the required 40 mW of 75-GHz power. It is possible to select any one of the 150,000 constant-voltage steps by controlling the bias current level and source impedance. A continuous voltage scale can be obtained by fine-tuning the frequency. The accuracy of the voltage at the array terminals is equal to the accuracy of the time standard used to stabilize the Gunn-diode source. Actual calibrations, however, are limited by noise and thermal voltages to an accuracy of a few parts in 10^9.

The ability to generate exactly known voltages between −12V and +12V can eliminate the problems and uncertainties of potentiometry from many standards laboratory functions. For example, Josephson array standards make it possible to perform absolute calibration of voltmeters at levels between 0.1V and 10V without the uncertainty of a resistor ratio transfer from standard cells. Another application is the measurement of voltmeter linearity with an accuracy higher than ever before possible.

References

1. R.L. Kautz and F.L. Lloyd, "Precision of Series-Array Josephson Voltage Standards," *Applied Physics Letters*, Vol. 51, no. 24, December 1987, pp. 2043-2045.
2. F.L. Lloyd, C.A. Hamilton, K. Chieh, and W. Goeke, "A 10-V Josephson Voltage Standard," *1988 Conference on Precision Electromagnetic Measurements*, June 1988, Tokyo.

John Giem
Development Engineer
Calibration Laboratory
Loveland Instrument Division

source must only be stable for the short time required to perform the two measurements of the transfer.

Each of the remaining dc voltage ranges is automatically gain adjusted relative to the internal 7V reference through a similar sequence of full-scale-to-1/10-full-scale transfer measurements. All gain errors can then be readjusted relative to the internal reference to remove measurement errors at any later time. The only gain error that cannot be adjusted during autocalibration is the time and temperature drift of the internal 7V reference.

Ohms and DC Current Calibration

Calibration of the ohms functions is similar to that of the dc voltage function. Traceability for the internal 40-kΩ reference resistor is established first. The internal reference resistor is measured relative to an externally applied traceable 10-kΩ standard resistor. The traceable value for this internal reference is stored in secure calibration memory until the next external calibration is performed. Resistance measurements are made by driving a known current I through an unknown resistance R and measuring the resultant voltage V. The unknown resistance value R is computed from Ohm's law, R = V/I. Since the dc voltage measurement function has been previously traceably adjusted, only the values of the ohms current sources (I) need be determined to establish calibration.

Adjustment of the ohms current source values begins by applying the nominal 100-microampere current source (10-kΩ range) to the traceable 40-kΩ internal resistance standard. The value of the current source is computed from the traceable measurements and stored in secure autocalibration memory. The 100-μA current source can be remeasured (autocalibrated) at any time to correct for changes in its value. Residual errors in this autocalibrated measurement are reduced to those of the internal reference resistor and the autocalibrated error of the 10V dc voltage range—essentially the drift of the internal voltage reference. For resistance measurements, only drift in the internal resistance reference will affect measurement accuracies. The gains of the voltage measurements V and the current sources I, which are derived from the internal voltage reference, will also change as this reference drifts, but the computed value for R is not affected since the V/I ratio remains unchanged.

The known 100-μA current, its value determined in the previous step, is next applied to an internal 5.2-kΩ resistor

(an internal 10-to-1 ratio transfer measurement). The value of this resistor is determined, again from Ohm's law. This new resistor R is computed (R = V/I) from the 100-μA current previously determined relative to known traceable standards and the previously calibrated dc voltage function. The value of this resistor is stored in autocalibration memory. This resistor is actually the 10-μA dc current function shunt resistor. With the shunt resistor R traceably determined, traceable dc current measurements can be computed from Ohm's law, I = V/R.

Now that the 5.2-kΩ internal shunt resistor is known, the 1-mA ohms current source (1-kΩ range) is applied and its value computed as a ratio relative to the 100-μA current source value. The 1-mA current source value is stored in autocalibration memory. This combined ohms current source and dc current shunt resistor ratio transfer process continues until all six currents and all eight shunt resistors are known relative to the two external standards.

As we set out to show, all gain errors for dc voltage, ohms, and dc current measurements have been traceably adjusted relative to only two external standard values: 10V dc and 10 kΩ. Table I summarizes the HP 3458A errors for the internal ratio transfer measurements described so far.

Table I
Internal Ratio Transfer Errors

External 10V dc	→ Internal 7V	0.03 ppm
Internal 7V	→ 10V dc	0.02 ppm
10V dc	→ 1V dc	0.33 ppm
External 10 kΩ	→ Internal 40 kΩ	0.30 ppm
Internal 40 kΩ	→ 100 μA	0.15 ppm
Internal 40 kΩ	→ Internal 5.2 kΩ	0.50 ppm
100 μA	→ 1 mA	0.50 ppm

Additional Errors

Gain and offset variations are the dominant sources of measurement error in a DMM, but they are by no means the only sources of error. Measurement errors are also produced by changes in leakage currents in the input signal path. These may be dynamic or quasistatic leakages. A more complete schematic of the input circuit of the HP 3458A is shown in Fig. 6. Recall that switches S1 and S2 are used to null the dc offsets of amplifier A1 and its input bias current. However, the capacitance C1 causes an error current I_{err} to flow when S1 is turned on. This current, sourced by the input, generates an exponentially decaying error voltage $I_{err}(R + R_i)$. If R_i is large, as it is for ohms measurements, significant measurement errors can result.

These errors can be reduced by providing a substitute source (shown in the shaded section of Fig. 6) to provide the charging current for the parasitic capacitance C1. Amplifier A2 follows the input voltage so that when switch S3 is turned on between the S2 and S1 measurement periods, C1 will be precharged to the input voltage. Second-order dynamic currents flow because of the gate-to-drain and gate-to-source capacitances of the switches, which are FETs. The HP 3458A performs complementary switching to minimize these effects. During an autocalibration, the offset of buffer amplifier A2 is nulled and the gain of the complementary switching loop is adjusted to reduce errors further.

High ohms measurements are particularly sensitive to parasitic leakage currents. For example, 10 ppm of error in the measurement of a 10-MΩ resistor will result from a change of 5 pA in the 500-nA current source used for the measurement. Over the 0°C-to-55°C operating temperature range a 5-pA change can easily occur. During autocalibration, which can be performed at any operating temperature, several internal measurements are performed with various hardware configurations. The results are used to solve simultaneous equations for leakage current sources. Knowing these leakage currents allows precise calculation of the ohms current source value for enhanced measurement accuracy.

Many other errors are also recomputed during autocalibration. Autocalibration can be performed in its entirety or in pieces (dc, ohms, or ac) optimized for particular measurement functions. The dc voltage autocalibration, for example, executes in approximately two minutes. The autocalibration process for the ohms functions, which also calibrates the dc current function, takes about eight minutes to complete. If the user is only concerned with correcting errors for dc or ac measurements, the ohms autocalibration

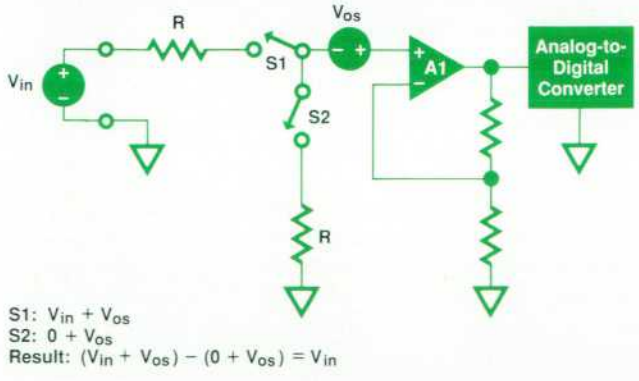

Fig. 5. *Simplified schematic of the dc voltage measurement function.*

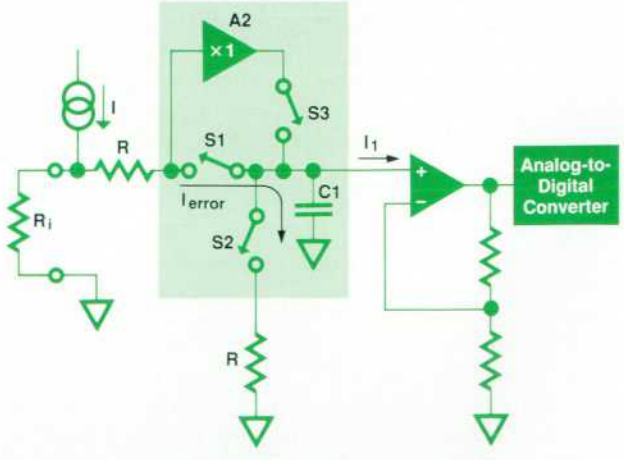

Fig. 6. *A more complete schematic of the HP 3458A input circuit.*

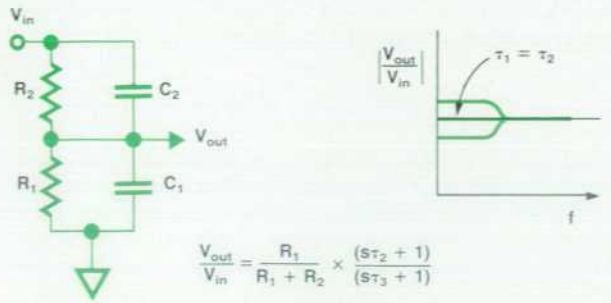

$$\frac{V_{out}}{V_{in}} = \frac{R_1}{R_1 + R_2} \times \frac{(s\tau_2 + 1)}{(s\tau_3 + 1)}$$

$$\tau_1 = R_1C_1 \qquad \tau_2 = R_2C_2$$

$$\tau_3 = \left(\frac{R_1R_2}{R_1 + R_2}\right)\left(C_1 + C_2\right)$$

Fig. 7. *RC attenuator gain-versus-frequency characteristic.*

sequence can be omitted to save time.

AC Frequency Response Calibration

The goals for self-calibration of the HP 3458A extended beyond the dc measurement functions. Just as the concept of sampling a signal and digitally computing its true-rms value goes against traditional DMM methods, so does the idea of adjusting the frequency response and gain of an ac voltmeter without applying external ac calibration sources. Normally, the first step in the calibration of an ac voltmeter would be to adjust the instrument for constant gain at all frequencies. This frequency flatness correction is generally performed by manually adjusting either resistive or capacitive circuit components. Resistive components usually determine gains at lower frequencies and capacitive components usually determine gains at higher frequencies. The frequency response characteristic of the HP 3458A ac measurement function is dominated by five compensated RC divider networks, which are used to condition the input signal for each measurement range. The gain-versus-frequency characteristic of an RC attenuator circuit is shown in Fig. 7. When the attenuator is properly compensated ($\tau_1 = \tau_2$), the resulting divide ratio is a frequency independent constant determined solely by the resistive elements.

It can be shown using Fourier transforms that if the input to a linear circuit is a perfect voltage step and the output of the same circuit is also a perfect voltage step, then the circuit transfer function is constant with frequency. The hardware used to implement the digital ac measurement technique of the HP 3458A is also used to sample a step output of the RC attenuator. The sampled data is used to compensate the internal RC divider networks for flat gain versus frequency without external inputs.

A simplified schematic for the 10V ac measurement range is shown in Fig. 8. The active compensation of the divider network is achieved by generating a "virtual trimmer" circuit element to allow the adjustment of the divider time constants. The trimmer is a programmable-gain bootstrap amplifier connected across resistor R1. The variable-gain amplifier allows control of the voltage across R1, effectively varying R1's value. The resistive divider ratio can be electronically servoed to match the fixed capacitive divider ratio given a measurable error function. The servo error signal is generated by applying an extremely square voltage step to the network. The step output is sampled at least twice. An amplitude difference between samples indicates the presence of an exponential component resulting from miscompensation of the attenuator. The digitally controlled loop servos the difference signal to adjust the virtual trimmer to achieve precise cancellation of frequency dependent errors. Sample times can be optimized for maximum sensitivity to the attenuator time constant RC, thus improving servo-loop rejection of second-order time constants resulting from capacitor dielectric absorption or other parasitic effects.

Sampling of the voltage step uses the same internal tools required to perform the digital ac measurement function. The flatness autocalibration voltage step is sampled with the integrating ADC configured for 18-bit measurement resolution at 50,000 conversions per second. An internal precision sampling time base is used to place samples with 100-ns resolution and less than 100-ps time jitter. Fig. 9 shows the range of attenuator output waveforms present during frequency flatness autocalibration. When the attenuator is compensated correctly, the output waveform will closely resemble an ideal voltage step as shown. Test data has shown that the automated compensation yields less than 50 ppm of frequency response error from dc to 30 kHz. Autocalibration of the frequency response will correct for component changes caused by temperature, humidity, aging, and other drift mechanisms. Correction

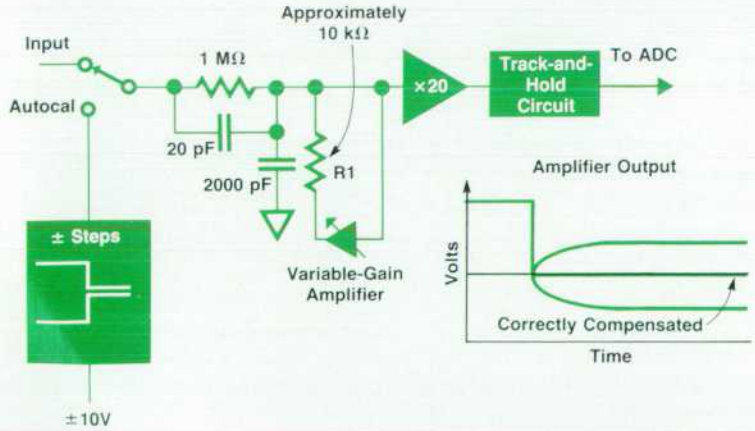

Fig. 8. *Simplified schematic of the 10V ac measurement range.*

A High-Stability Voltage Reference

Autocalibration in the HP 3458A Digital Multimeter is a process of transferring the gain accuracy of a single voltage reference to all measurement gains. The design goal for the internal reference of the HP 3458A was to provide long-term stability and temperature stability comparable to external standards that would normally be used to calibrate an 8½-digit multimeter. These goals were achieved by using a temperature-stabilized solid-state Zener reference. Without temperature stabilization, the Zener's voltage drift with temperature is approximately 50 ppm/°C. A proportional temperature control loop senses the chip temperature of the reference device and reduces this drift to less than 0.15 ppm/°C.

The long-term drift of each voltage reference assembly is measured by an automated drift monitoring and screening process. Reference assemblies, including the temperature controller, are monitored until the aging rate is shown to be less than the 8 ppm/yr stability requirement of the HP 3458A. Summarized test data for a number of 8 ppm/yr reference assemblies is shown in Fig. 1. Monitoring the references for additional time allows the selection of assemblies that exhibit aging rates less than 4 ppm/yr for the high-stability option.

David E. Smith
Development Engineer
Loveland Instrument Division

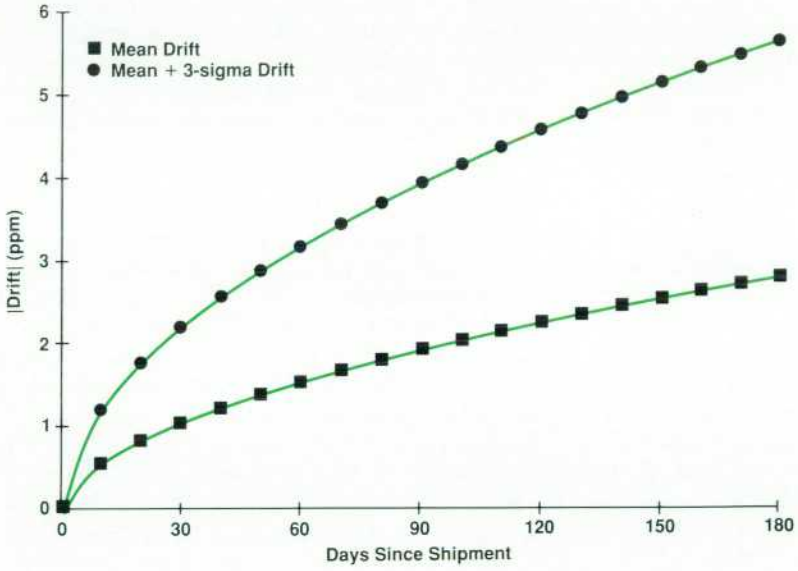

Fig. 1. *HP 3458A internal voltage reference drift distribution.*

of these errors allows a single specification to apply for extended operating conditions.

AC Gain Calibration

Once the frequency flatness characteristics are adjusted, the second step of calibration can be completed. Gain correction for the measurement must still be achieved. In Fig. 7 it can be seen that when frequency compensation is achieved, the attenuator gain can be established equally well at any frequency as long as the calibration signal amplitude is precisely known. Adjustment of the circuit gain using a dc signal is convenient since a traceably calibrated dc voltage reference and a dc voltage measurement function are available. Gain adjustment of the ac measurement function using known dc voltages allows complete autocalibration of ac measurement accuracy in much the same manner as the dc voltage measurement function.

Several mechanisms can limit the accuracy of a dc gain adjustment. Dc offsets or turnover errors can be minimized by performing gain adjustment calculations using known positive and negative voltages. Errors caused by white noise are reduced by averaging 40,000 samples for each voltage measurement made through the wide-bandwidth track-and-hold circuit. Low-frequency 1/f noise is minimized by chopping these 40,000 readings into groups of 1000, each group sampling alternating polarities of the known internal

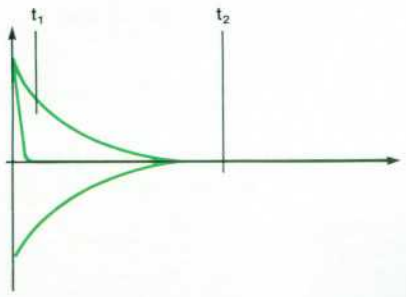

Fig. 9. *The range of attenuator output waveforms present during frequency flatness compensation. The output waveform closely resembles an ideal voltage step when compensation is correct.*

dc calibration voltages. This voltage chop is performed at a fast enough rate to achieve maximum cancellation of the 1/f noise voltage. A final error mechanism results from aliasing of internal spurious signals. The internal 10-MHz clock signal tends to be present in small amounts everywhere. The ac signal path and the track-and-hold circuit (2-ns sample aperture) each have sufficient bandwidth to couple the internal clock into measurements. If the sample spacing is a multiple of the 100-ns clock period, the internal spurious clock will be aliased or mixed down to contribute a dc offset in the measurement. A 100-μV-peak spurious clock signal can lead directly to a 100-μV error in measuring the internal dc calibration signal as shown in Fig. 10. The HP 3458A uses a random sampling time base mode during this calibration sequence. The time base generates randomly spaced sample intervals with a resolution of 10 ns. The chopped groups of random samples, 40,000 in all, are averaged together to obtain the net gain of the divider. Errors caused by dc offsets, white noise, 1/f noise, and clock aliasing are reduced using this internal calibration algorithm. Gain calibration of the ac measurement function relative to the internal dc reference is accomplished with less than 10 ppm error for intervals extending to two years. The residual dc gain calibration error will limit the absolute measurement accuracy for low-frequency inputs.

Additional Errors

Besides adjusting the frequency response and gain of each ac measurement range, other corrections are performed during autocalibration. Offset voltage corrections are determined for each ac amplifier configuration. The offset of the analog true-rms-to-dc converter is determined. The offset of the analog trigger level circuit is nulled. Internal gain adjustments for various measurement paths are performed. For example, the track-and-hold amplifier gain is precisely determined by applying a known dc voltage and measuring the output in track mode using 7½-digit internal dc measurements. A gain ratio is computed using this measurement and the hold mode gain is determined by averaging 40,000 samples using the 6-μs, 50,000-reading-per-second, 18-bit conversion mode of the integrating ADC. This gain is critical to the accuracy of the digitally computed rms ac voltage function and to the wideband sampling functions. Ac current measurements use the same shunt resistors as the dc currents. A differential amplifier is used to sample the voltage across the shunt resistors for ac current measurements, and the gain of this amplifier is computed during autocalibration.

As a result of autocalibration, the ac measurement accuracy of the HP 3458A is unchanged for temperatures from 0°C to 55°C, for humidity to 95% at 40°C, and for a period of two years following external calibration. Execution of only the ac portion of the autocalibration process is completed in approximately one minute.

Summary

Two-source calibration of a state-of-the-art digital multimeter provides several benefits:

- Increased process control within the standards labora-

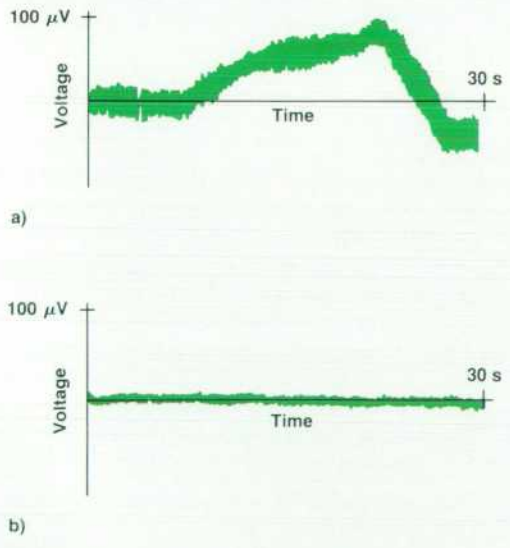

Fig. 10. *(a) Using the clock-derived time base, a 100-μV spurious clock signal can lead directly to a 100-μV error in measuring the internal dc calibration signal. (b) The HP 3458A uses a random sampling time base mode to eliminate this error source.*

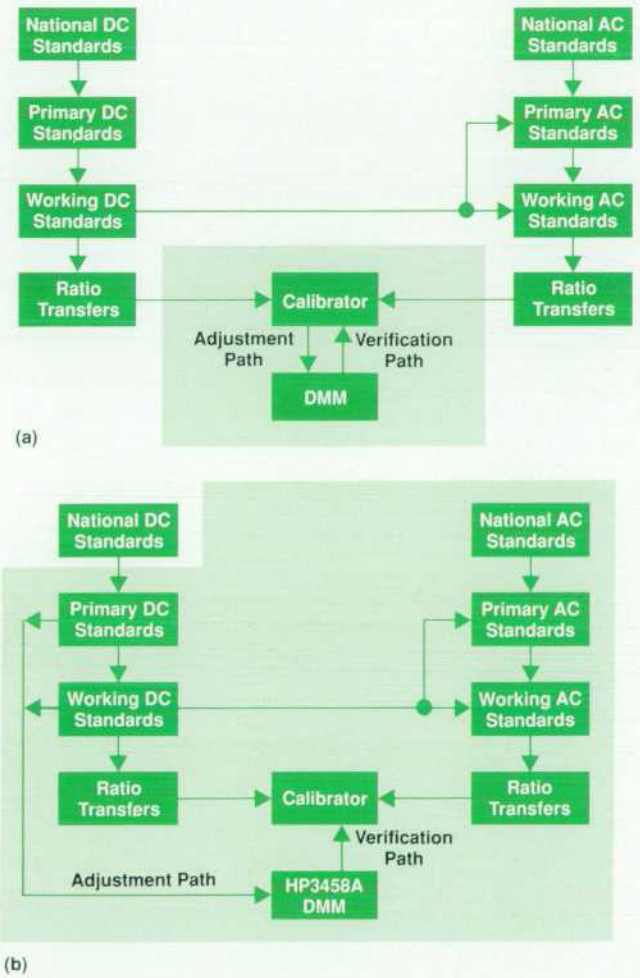

Fig. 11. *(a) Traditional calibration chain for dc and ac voltage. (b) HP 3458A calibration chain, showing the increased verification confidence that results from internal calibration.*

tory through the independent ratio transfers of the DMM

- Reduced calibration time
- Increased measurement accuracies in real environments
- Increased confidence through complete self-testing.

The greatest benefit of two-source calibration is seen not by the instrument end user but by the calibration facility supporting those instruments. Fig. 11 shows the normal instrument and two-source calibrated instrument traceability chain. When verifying the results of the two-source calibration process, the metrologist now has the independent checks of the HP 3458A to catch inadvertent human errors in the normal process. Technique, cabling, and other instruments used in the generation of calibration values are no longer open-loop errors that may propagate through a calibration laboratory. Two-source calibration can identify errors anywhere within the traceability chain, from primary standards to final values.

The HP 3458A autocalibration procedures are also performed during the instrument self-test, which takes about one minute. The only difference is reduced averaging of the internal results for faster execution. Also, the results are not retained in memory afterward. The self-test procedures perform highly accurate measurements on each range of each function, thereby providing a comprehensive analog and digital confidence test of the system.

Acknowledgments

Verification of the two-source calibration performance of the HP 3458A was not a particularly easy process. Identifying and quantifying errors, normally a fairly straightforward proposition, was made difficult by both the high accuracy and the independent nature of two-source calibration. Many people have contributed to the success of these efforts. Particular thanks are given to Bill Bruce, the Loveland Instrument Division standards lab manager, and his staff. Their assistance and guidance were of immeasurable help. They were disbelieving and skeptical initially, but in the end, their ardent support was greatly appreciated. Kudos go to our production engineering test development group of Bert Kolts and Randy Hanson for developing a fully automated system for performance verification testing of the HP 3458A. Last and most important, special thanks to the staff of the U.S. National Bureau of Standards: to Dick Harris and Clark Hamilton at NBS Boulder, Colorado for allowing us access to their Josephson junction array during our product development and for helping us set up our own Josephson array voltage standard system. Thanks also go to Neil Oldham of NBS Washington, D.C. for making available his transportable digital ac reference standard, also for performance verification testing.

Design for High Throughput in a System Digital Multimeter

High-speed custom gate arrays, microprocessors, and supporting hardware and a substantial investment in firmware design contributed to the design of the HP 3458A DMM as a system for moving data efficiently.

by Gary A. Ceely and David J. Rustici

MANUFACTURERS OF ELECTRONIC and other types of products have learned that high test system throughput is vital to maintaining production capacity. As a primary component of automated test and data acquisition systems, the system digital multimeter (DMM) has become a major factor in determining system throughput. A DMM must not only be able to take and transfer high-speed bursts of readings, but must also have the ability to reconfigure itself quickly when measuring several different parameters in rapid succession.

Historically, DMM performance has been hindered by a number of factors, such as relay switching times, ADC conversion delays, and the limited processing power of early-generation microprocessors. In addition to controlling the ADC hardware, taking and transferring readings, and parsing commands, the microprocessor has been saddled with scanning the front-panel keyboard, updating the display, and polling various peripheral ICs to monitor and update status information. Increasing demands on the capabilities of firmware written for these machines have only compounded the problem. Adoption of more English-like programming languages has added greatly to both bus overhead (because of the length of these commands) and parsing time, which formerly was a minor factor.

Another performance limitation in system DMMs has resulted from the need to make floating measurements, that is, measurements referenced to the **LO** terminal instead of earth ground. Since the **LO** terminal may be raised to a potential several hundred volts above ground, the ADC hardware must also float with this voltage. The problem here is that the HP-IB (IEEE 488, IEC 625), and therefore the hardware that interfaces to it, is earth-referenced, requiring that the ADC hardware be isolated from the controlling microprocessor. In many cases, the ADC hardware is designed around a second microprocessor which communicates with the main microprocessor via an isolated serial link, forming a bottleneck in high-speed ADC programming and data transfers.

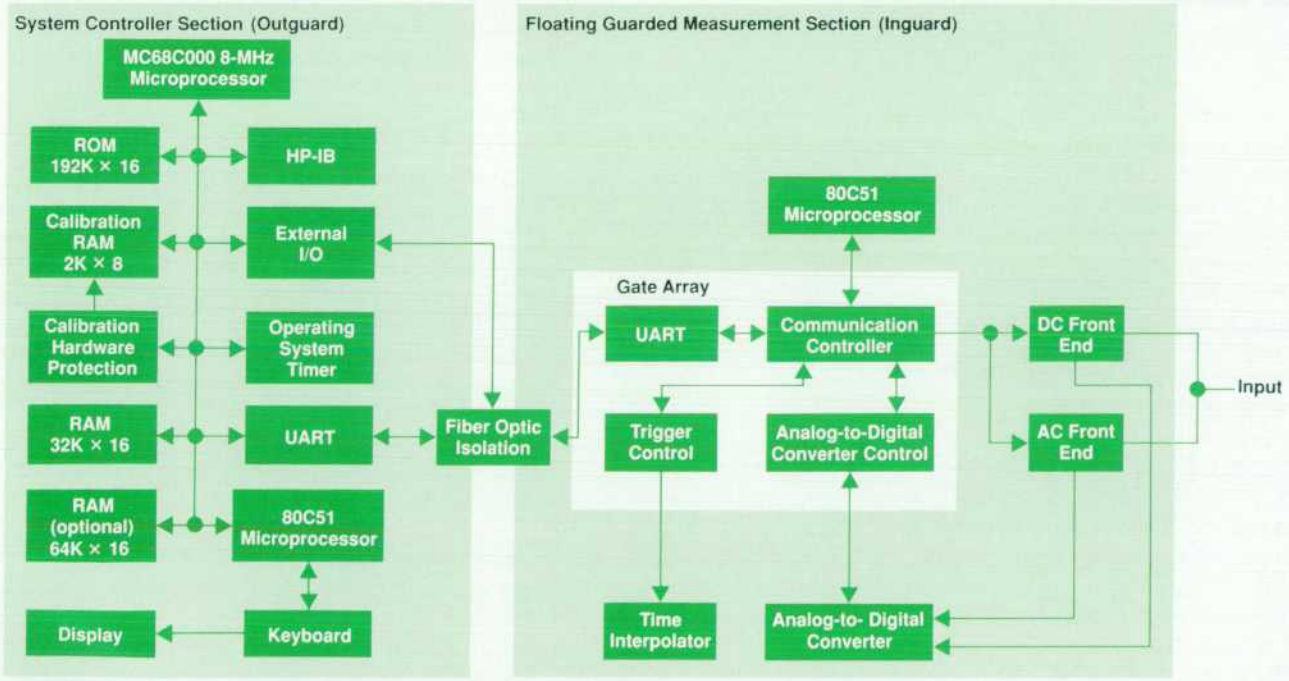

Fig. 1. *HP 3458A Digital Multimeter system block diagram.*

Considering the history of DMM performance, it becomes obvious that the design of the instrument as a system in itself is critical to the performance of the surrounding automatic test system as well. Two key design goals for the HP 3458A were that it be able to reconfigure itself and take a reading 200 times per second, and that it be able to take and transfer readings (or store them internally) at a burst rate of 100,000/s. To achieve these goals, system design for the HP 3458A focused on expediting the flow of data through the instrument, both in the hardware and in the firmware.

Design Overview

A simplified block diagram of the HP 3458A is shown in Fig. 1. Like previous designs, the DMM is divided into two sections, inguard and outguard, which correspond to the hardware inside and outside of the guarded (isolated) section of the DMM. In this design, however, the bottleneck of the serial interface between the two sections is overcome by the use of a high-speed (5 Mbits/s) fiber optic data link and custom gate arrays on each end to decode and buffer received data.

Performance features on the outguard side include an 8-MHz MC68C000 main processor, high-speed RAM and ROM (requiring no wait states from the processor), a separate 80C51 microprocessor to control the front-panel interface, and a programmable timer used as an operating system clock. This represents a significant upgrade in the outguard hardware over previous 6800-based designs, and not only yields faster execution of instructions, but also frees the main processor from polling peripherals, since all I/O and interprocessor communications are now interrupt-driven. Additional gains are realized through the use of a double-buffered HP-IB input scheme (the parser reads data from one buffer while an interrupt service routine fills the other) and a hardware HP-IB output buffer, which allows the main processor to write data to the HP-IB in words (16 bits) instead of bytes (8 bits).

Outguard RAM is divided into three sections: an EEPROM for storing calibration constants, standard RAM (nonvolatile), and optional RAM (volatile). Calibration RAM is distinct from the rest of RAM because it is protected from accidental overwrites by a hardware mechanism that also makes writing to it rather slow. Standard RAM is divided into program memory, reading memory (10K 16-bit readings), state storage, and system overhead (stacks, buffers, etc.). Nonvolatile RAM is used here to protect stored instrument states, subroutines, and user key definitions. Optional RAM is available only as additional reading storage (64K readings).

Inguard hardware is also under microprocessor control (an 80C51, in this case), but the heart of the inguard section is a 6000-gate, 20-MHz CMOS gate array. Functions performed by the gate array include communications with the outguard section through a custom UART, trigger logic control, analog-to-digital conversion, and communications between the UART and other parts of the inguard section. Shift registers are incorporated to minimize the number of interconnections between the gate array and other inguard circuits (the ADC, the ac and dc front ends, and the trigger control logic). Five shift registers containing 460 bits of information reduce the number of interface lines to just three per circuit. Communications are directed by the processor, which also interprets messages sent from the outguard section and generates response messages (see "Custom UART Design," page 36.

Firmware Structure

The division of tasks between the inguard and outguard processors is based on the need to minimize the flow of messages between them. Inguard firmware is responsible for controlling the ADC measurement sequence, controlling the trigger logic during measurements, and directing configuration data to the other inguard circuits. Outguard firmware responsibilities are as shown in Fig. 2. Primary functions, such as parsing, command execution, display updating, and keyboard input are performed by separate tasks under operating system control. Other functions, such as HP-IB I/O and interprocessor communications, are interrupt-driven, are coded in assembly language for maximum speed, and communicate with the primary tasks via signals and message exchanges. High firmware throughput is achieved by focusing on optimization of time-intensive tasks, such as data transfer and manipulation, parsing and execution of commands, task switching overhead, and the measurements themselves.

Fig. 3 shows the flow of data through the HP 3458A. Data flow is divided into two main paths: the input path for messages received from the controller, and the output path for measurements generated by the instrument. When a controller sends a command such as DCV 10, the data flow is from the controller to the HP 3458A through the HP-IB. The HP-IB handler accepts incoming data and passes it on to the outguard processor's parser, which interprets the command and then passes control to an execution routine. After determining the necessary actions, the execution routine sends state change data to RAM and inguard-bound messages to the UART. Messages sent to the inguard section are of two types: measurement messages, which control the type of measurement (e.g., dc voltage or ac voltage), and configuration messages, which define the state of the front ends and the ADC and timer control circuits. Data is received by the inguard UART and passed to the inguard processor, which parses the message and either acts upon it or directs it through the communication controller to one of the other inguard circuits. Once the configuration phase is complete, the ADC is ready to take a reading, and throughput becomes a matter of getting the reading out of the instrument quickly. Referring again to Fig. 3, the output data path is from the ADC to the inguard UART, through the fiber optic link, and on to the outguard processor. The processor performs any required math and formatting operations, and then directs the data either to reading storage or to the HP-IB.

Data Input, Configuration, and Measurements

Programming commands coming in over the HP-IB are received and buffered by an interrupt service routine, which in turn signals the HP-IB parser/execution task. The interrupt code is designed to continue reading characters from the HP-IB chip as long as they continue to come in at a rate of 100 μs/character or faster. In this manner, an

Firmware Development System

Firmware for the HP 3458A DMM was developed on four HP 9000 Computers (Models 320 and 350) under the HP 64000-UX microprocessor development environment. Each system was fully equipped to operate as an independent development station, and the systems were networked to facilitate transfer of code revisions (see Fig. 1). A fifth station was used for consolidating code modifications to be tested using a prototype HP 3458A and the HP 3458A production test system. After passing an extensive battery of tests, code was released in EPROM form for other prototype instruments.

Firmware tasks were divided along lines intended to minimize interdependence between the designers. The areas of responsibility were (1) measurements and calibration, (2) digitizing, (3) data processing, formatting, and storage, and (4) parsing, I/O, and operating system overhead. Fig. 2 shows a breakdown of the amount of object code generated by various modules. Al-

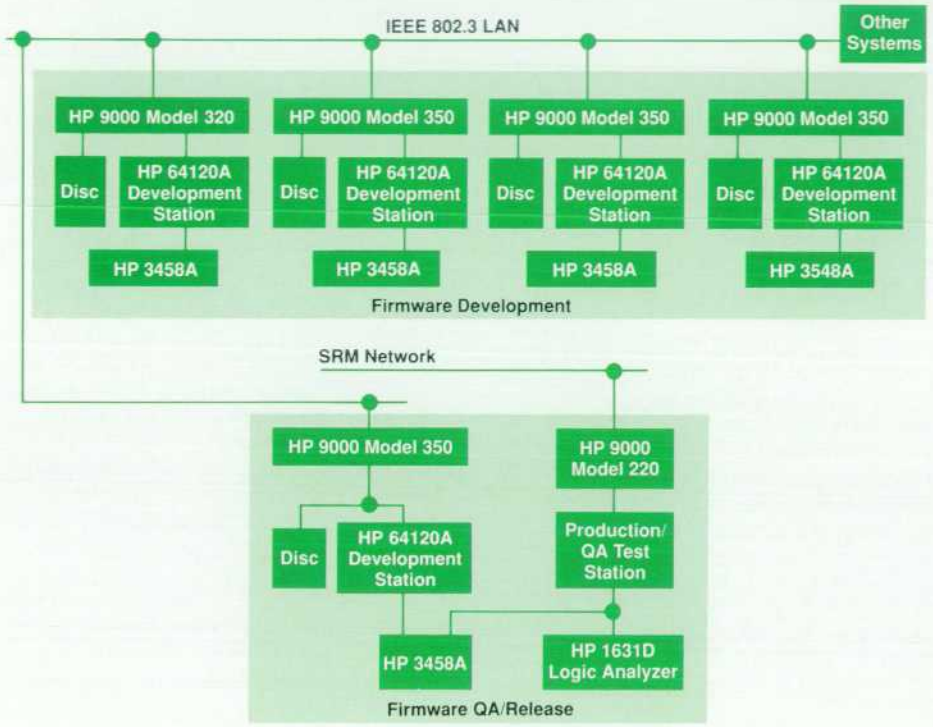

Fig. 1. *HP 3458A firmware development and QA regression test systems.*

entire command or string of commands can be read in during a single invocation of the interrupt routine, thereby generating only one signal to the parser task. In reality, two input buffers are used: one that is filled by the interrupt routine, and another that is read by the parser task. After the interrupt routine signals the parser that data is present in one buffer, that buffer belongs to the parser task, and the other buffer is used for the next command that comes in. When the parser empties a buffer, that buffer is freed for later use by the interrupt routine. Using two buffers simplifies pointer manipulation so that data can be read in and passed to the parser quickly.

To maximize the flow of data to the HP-IB parser/execution task, the instrument must first be programmed to an idle state (e.g., using TARM HOLD). This allows the operating system to keep the HP-IB parser task active so that no task switching is necessary when an HP-IB command is received. The parser is a table-driven SLR (simple left-right) design, with all critical components coded in assembly language. Simple commands can be parsed in as little as 1 ms; longer commands take as much as 3 ms. For a further increase in system throughput, command sequences can be stored as subprograms, in which case they are first compiled into assembly language by the parser/code generator. Executing command sequences in this fashion eliminates most of the overhead of bus I/O and parsing and allows the HP 3458A to perform reconfiguration and trigger operations almost twice as fast as the same sequence with individual commands (340/s instead of 180/s).

In many situations, the HP 3458A will be reconfigured for a different measurement setup with each test, which may include only one measurement. The setup changes in these cases may take more time than the measurement, so the configuration time must be minimized. To perform 180 reconfiguration and trigger operations per second, the instrument must be able to transfer, parse, and execute a command in slightly over 5 ms. Of this total, several hundred microseconds are spent in bus transfer and system overhead, and up to 3 ms may be spent parsing the command. Given that an additional several hundred microsec-

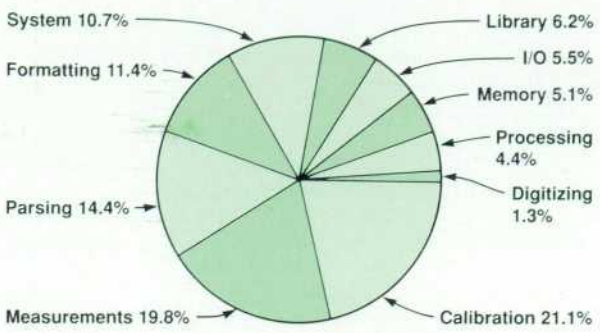

Fig. 2. *Outguard firmware modules.*

together, over 28,000 lines of C code were written, representing roughly 80% of the 356K bytes of object code generated. The remainder (12,000 lines) was written in 68000 assembly language.

During the most intense period of firmware development, code revisions were released on a weekly basis. To relieve the firmware team of the time-consuming task of generating and testing code revisions, a fifth team member was given this responsibility. Firmware designers uploaded source code weekly to the fifth system, where it was compiled, linked, and downloaded to an emulator. Having source code available on this system made it possible to trace and analyze defects using a dedicated QA system to reproduce them. The fifth development system was also used for archiving firmware revisions using RCS (UNIX revision control system). To reduce duplication of effort, the test system used for firmware development was a replica of the HP 3458A production test system, which had been developed earlier in the project cycle to be used in environmental testing and prototype characterization.

As the firmware construction phase neared completion, two engineers were added to the project so that test software could be developed in parallel with the firmware effort. To save test writers the trouble of learning the details of test system operation, drivers and utilities were written that allowed each new test to be written as an isolated subroutine. The test system executive simply loaded and ran each test as it was needed, thereby providing an efficient mechanism for adding new tests throughout the construction and test phases. Both hardware and firmware designers wrote tests for the test suite. Each was assigned a specific area of functionality to be tested, using both white-box and black-box approaches.

In addition to the tests written specifically to verify firmware operation, each revision of code was subjected to the production test software (which mainly tested the analog hardware for measurement accuracy). Additional test coverage included the entire HP 3458A user's manual, with emphasis on the command reference, example programs, and randomly generated combinations of valid and invalid syntax. As defects were found, they were fixed and the test code run again for verification. Following a successful run through the test suite, code was released and source code was saved using RCS. Saving old code revisions enabled the firmware team to recreate earlier code revisions to help track down defects that may not have been reproducible on a newer code revision. When a new defect was found, tests were written and added to the test suite to ensure that the defect would not recur. By the end of the project, the test suite had grown to where 12 hours were required to run all tests. To assess testing progress and effectiveness, defects were submitted to HP's DTS (defect tracking system). Metric reports were generated and analyzed on a weekly basis to help assess the firmware status.

Victoria K. Sweetser
Development Engineer
Loveland Instrument Division

onds will be spent taking and transferring the reading, only about 1 ms is left for the execution of the command. In this millisecond, the execution routine must range-check parameters, calculate the gain and offset values, and configure the trigger controller, the ADC, the front-end hardware, and the inguard processor. In the worst case, performing these operations takes considerably longer than a millisecond. A complete configuration of all the inguard sections takes 1.4 ms, and settling time for the front-end relays adds another 1.3 ms. In addition, a function command may require as many as six floating-point calculations, each taking 0.3 ms. This all adds up to well over 4 ms; therefore, a number of optimizations have been incorporated to reduce configuration time.

The first step is to avoid reconfiguring the instrument or a section of inguard if there is no change. For example, if the present function is ac volts and the new command is ACV, only the range is configured (if it changes), not the function. The ADC configuration is the same for dc volts, ohms, and dc current, so the ADC section is not reconfigured for changes between these functions. The trigger configuration changes only for digital ac voltage or frequency measurements, so a new configuration is sent only when entering or leaving these functions. In general, reconfiguration occurs only to the extent required by a given command.

Each combination of function and range uses different gain and offset values for the ADC readings. The gain and offset values are scaled by the ADC's aperture, so if the aperture increases by 2, the gain and offset are scaled by 2. An execution routine retrieves the gain and offset values from calibration memory and scales them by the aperture. Then the 120%- and 10%-of-full-scale points are calculated for overload detection and autoranging. The autoranging algorithm uses a different ADC aperture and has a separate set of 120% and 10% points. These two calculations were removed from the execution routine, and are done at calibration time since the autoranging algorithm always uses the same ADC aperture. To reduce the effect of the other four calculations, a data structure is used that saves the gain and offset for each function and range as it is needed. If the aperture of the ADC is changed, the data structure is cleared, and as function and ranges are revisited, the data structure is filled in. This eliminates recalculation of values that are constant for a given aperture.

An operation that is not always necessary but takes considerable time during a range or function change is a special

sequence of relay closures in the front-end circuitry. This sequence protects the relays from damage when high voltage is on the input terminals during range changes, but is not needed when measuring low voltages or if high voltage is only present when the instrument is set to a high voltage range. Therefore, the HP 3458A provides an HP-IB programmable command to defeat the protection scheme, speeding up the relay sequence by a factor of five. If an overvoltage condition occurs while protection is inactive, an interrupt is generated and the relay sequence is reversed, thereby protecting the relays from damage. A delay of 0.4 second is then inserted to prevent a rapid recurrence of the overload condition, and the instrument reverts to normal (protective) relay sequencing thereafter.

Another technique used to reduce the configuration time is to defer computations until the last possible moment. The scale factor used in the format conversion of ADC readings from integer format to real or ASCII format is an example of this technique. Many commands cause the scale factor to change, so instead of each command computing the scale factor, a flag is set and the calculation is performed when the scale factor is first used. This eliminates wasted time from unnecessary calculations when many intermediate configuration changes are sent to the instrument, and reduces the time spent responding to even a single HP-IB command.

Data flow between the outguard and inguard sections has the potential to be a bottleneck, because the UART and the inguard processor can only accept configuration data at a rate of 20,000 words/s. Furthermore, commands to change relays can take a millisecond for the inguard processor to execute. To relieve the outguard processor of the need to wait on the inguard processor, a buffer was added to store messages bound for the UART. This buffer is deep enough to hold an entire configuration change—128 commands. This allows the outguard processor to overlap its activities with the inguard processor's. If the buffer is empty and the UART is not busy sending data, the 68000 will send a command directly to the UART, avoiding the overhead of the buffer. If the UART is busy, data is written to the buffer instead. In this case, the UART generates an interrupt when it is ready to accept the next word, which is then retrieved from the buffer and sent.

In addition to fast reconfiguration, system throughput depends on the time required to make a measurement. Fig. 4 shows the steps an ADC reading goes through before it is sent to the HP-IB. The first step is autoranging: if a reading is less than 10% of the range or greater than 120%, the instrument switches to the next range, changes the ADC's aperture for a fast measurement, and takes a reading. This procedure is repeated until the correct range is found, and then the final measurement is made with the ADC's original aperture. Although this algorithm is very fast (typically 8 milliseconds), it usually requires that the ADC take several

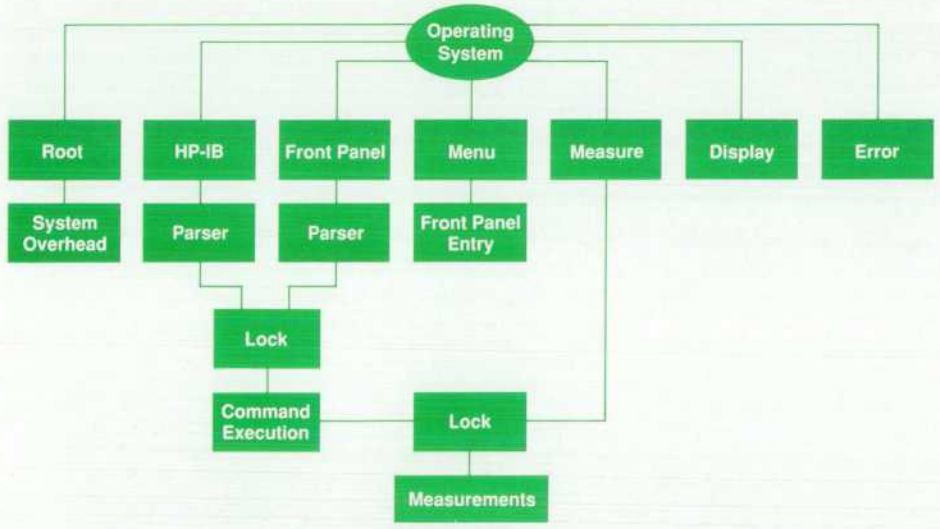

(a)

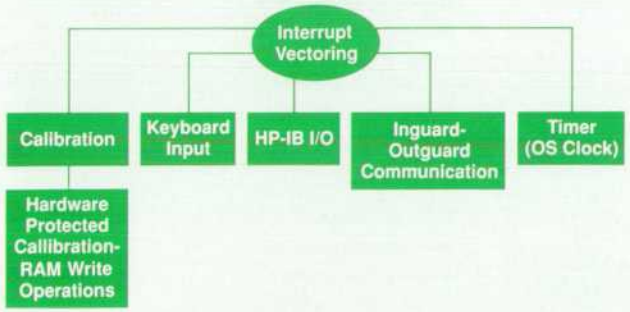

(b)

Fig. 2. *HP 3458A firmware structure. (a) Tasks under operating system control. (b) Interrupt service routines.*

Custom UART Design

At the center of the communications link between the inguard and outguard sections of the HP 3458A DMM is the custom UART (universal asynchronous receiver/transmitter). A serial interface was chosen because an isolated parallel interface would have been prohibitively expensive. Unfortunately, conventional UARTs are too slow to meet the HP 3458A's required data rate of 200 kbytes/s, which corresponds to a baud rate of 2 Mbits/s, counting start and stop bits. Therefore, fiber optic couplers were chosen, which also provide the benefit of infinite isolation resistance.

Conventional UARTs require a clock rate that is 16 times the baud rate; thus, to generate the 3458A's required baud rate, the clock rate would have to be 32 MHz. The 16× clock rate is needed to compensate for mismatched clock frequencies and waveform distortion. These two factors can be controlled within the HP 3458A, so a clock rate of three times the baud rate is used. The UART design is implemented as part of a CMOS gate array driven by a 10-MHz clock. This clock rate yields a baud rate of 3.3 Mbits/s, which meets the design goal with some margin.

The data format for the UART is shown in Fig. 1. The first bit is the start bit and indicates the beginning of a message. The next bit is the handshake bit. If this bit is high, a data/command message will follow immediately. If the bit is low, the message is a handshake and the next bit will be a stop bit. A handshake message is sent each time a data message is read by the processor, ensuring that a new message will not be sent until the previous message has been read. The next-to-last bit is the interrupt or command bit, used to indicate whether the preceding message was data or a command. A command message from the inguard section could be an ADC conversion failure, an end of sequence message, or a change in the front or rear terminals. Command messages generate interrupts, eliminating the need for software to check the data from the UART to determine the message type. The middle 16 bits of the message represent the data or command, and the last bit is the stop bit.

Fig. 2 shows a block diagram of the UART and the communication controller. When the decoding state machine detects that a start bit has been received, it waits three cycles to decide whether the message is a handshake. If so, the state machine returns to its initial state. If the message is data, the next 16 bits are clocked into the input shift register. The state machine then examines the next bit (the command/data bit). If the message is a command, an interrupt is generated.

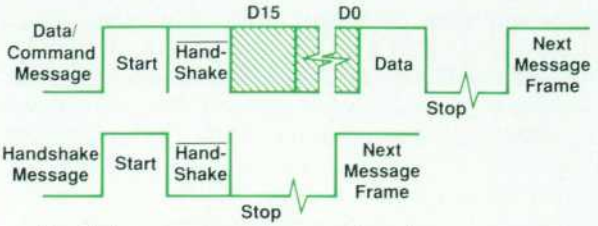

Fig. 1. *Interprocessor message formats.*

For transmitted messages, the encode machine first generates a start bit. If the message is a handshake, the next bit is set high; otherwise (if the message is data), the next bit is set low. The 16 bits of data are sent next (if required), and if the message is a command, the last bit is set high.

Buffers in the UART are used both for received data and data to be transmitted. This allows the ADC to leave data in the buffer while starting the next measurement, thus maximizing the overlap between outguard and inguard. Once the buffer has been emptied, the handshake message is sent and an interrupt can be generated. The interrupt can be used as a request for more data to be sent. The buffer queues requests from four sources: the ADC's error detection circuitry, the ADC's output register, the trigger controller messages, and the inguard processor.

The input buffer also has a direct output mode to the shift registers. When data is sent to the inguard section, the processor is interrupted, the data is parsed, and, if the message is a configuration message, the direct output mode is selected in the communication controller. This mode allows the next message to be sent to both the processor and the shift register, thereby sending the configuration data directly to the appropriate section. In this case, the processor receives the message but does not act upon it, thereby eliminating the overhead of processor intervention in the configuration process.

Although the use of microprocessors has enabled instruments to offer greatly enhanced measurement capability, a severe speed penalty may be incurred if firmware is burdened with tasks that are best left to hardware. The HP 3458A's use of a custom UART coupled directly to the measurement hardware optimizes performance by balancing the workload between hardware and firmware.

David J. Rustici
Development Engineer
Loveland Instrument Division

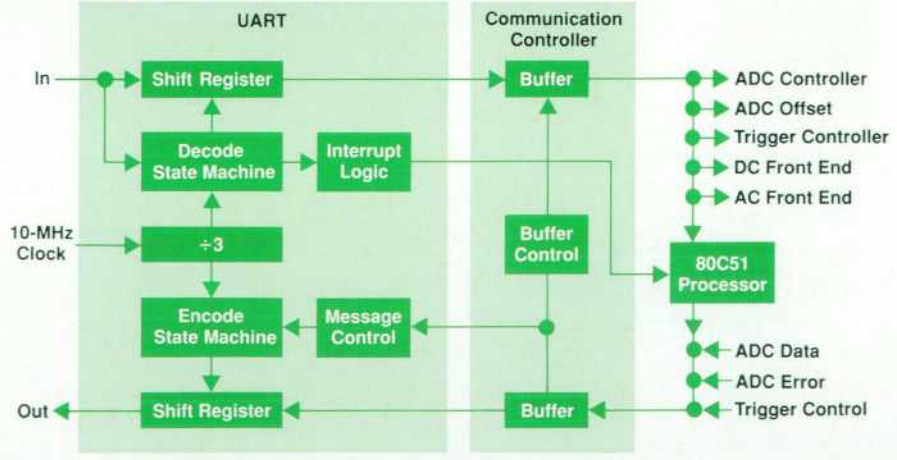

Fig. 2. *Block diagram of the UART and data communication portions of the inguard gate array.*

samples to generate one reading. Therefore, a faster measurement will be made if autoranging is turned off.

Throughput is also enhanced by minimizing operating system overhead. In cases where high throughput is not an issue (e.g., long integration times), measurements are handled by a background task, which runs whenever the instrument is not actively executing commands. This task simply monitors the trigger and trigger arm states to see if a measurement should be taken. When throughput is an issue, however, measurements are initiated directly by the HP-IB command parser/execution task. In this case, the overhead of task switching (approximately 250 μs) is eliminated, leaving only the overhead of communication between the interrupt service routine and the HP-IB task. Another speed enhancement is the use of preprogrammed states, which fall into two categories: predefined states (activated using the PRESET command), and user-defined states (stored using the SSTATE command and activated using the RSTATE command). Since these commands cause an extensive reconfiguration, their primary benefit is in putting the instrument in a known desired state. However, they can also save time when the alternative is to send long strings of commands to program the instrument to the same state.

Output Data Path

Once the instrument has been configured and triggered, a measurement is taken by the ADC and transmitted through the fiber optic link to the outguard processor. The format for this reading is either a 16-bit or a 32-bit two's complement result with the range offset subtracted. The next step is to convert the readings into volts, ohms, or amperes by multiplying by the gain of the range. If a math operation is active, it is initiated using a procedure variable that points to the math subroutine. At this point, the reading is in a 64-bit floating-point format, and a format conversion is required for an integer, ASCII, or short real format. The last step is to display the result and send it to memory or the HP-IB. Some steps can be eliminated using the appropriate HP-IB command; for example, the display operation is deleted using the DISP OFF command.

If autoranging, math, and the display are turned off and the output format matches the ADC's internal format, the measurement can be sent directly to the HP-IB or memory. Special assembly language routines were written to handle these high-speed modes. The time allowed to read the measurement and send it out is 10 μs (given a maximum reading rate of 100,000 per second). There are two data paths: one that sends readings to memory and one that sends them to the HP-IB.

Reading Storage. The memory structure dictated by HP's multimeter language is a general circular buffer in which readings may be added or removed at any time. This buffer can be used in either of two modes: FIFO (first in, first out) or LIFO (last in, first out), the main distinction being that the LIFO mode will overwrite the oldest readings when memory fills, whereas the FIFO mode will terminate when memory fills, thus preserving the oldest samples. A general program loop for receiving readings from the ADC and writing them into memory is as follows:

- Wait until the ADC has taken a reading.
- Write the reading into the current fill location and increment the fill pointer.
- Has the fill pointer reached the top of memory (buffer pointer wrap-around)?
- If memory is full and the memory mode is FIFO, stop.
- Terminate the loop when the end of sequence is sent.

Within 10 μs, the 68000 will allow only about three decisions to be made. Even using hand-optimized assembly

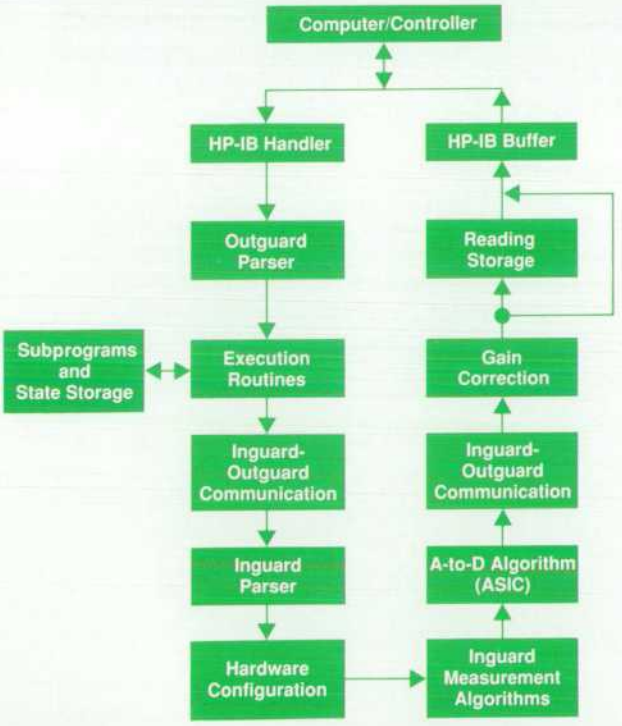

Fig. 3. *Input and output data flow paths.*

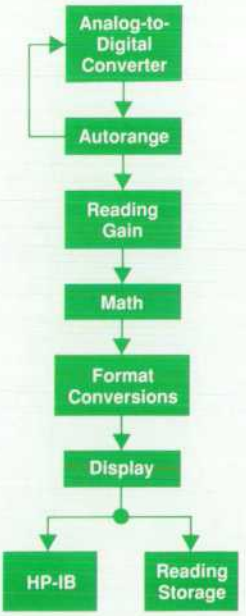

Fig. 4. *Processing of readings.*

language, a single program loop could not be written to implement the general memory model in the allotted time. The solution uses the fact that if enough decisions are made before the start of the burst, the number of on-the-fly decisions can be reduced. Before the start of a burst of samples, it is known how many readings can be added before the buffer pointers wrap around, and how much room is left before the circular buffer fills. The problem is divided into a set of special cases. For example, assume that 1000 readings are expected from the ADC. Memory fill and empty pointers indicate space for 2000 readings, but the fill pointer is only 100 samples from buffer wraparound. Under these conditions, the memory fill algorithm can be stated as follows:

- Fill memory with samples until the buffer fill pointer reaches the top of memory.
- Wrap around the fill pointer to the bottom of memory.
- Fill memory with samples until the sequence is complete.
- Exit the routine.

Any memory scenario can be expressed as a combination of the following special-case loops:

- Fill memory with samples until the fill pointer reaches the top of memory, then wrap around the fill pointer to the bottom of memory.
- Fill memory with samples until memory is full (fill pointer = empty pointer).
- Fill memory with samples until the sequence is complete.

Four factors influence the algorithm used: memory mode, number of readings expected, total available memory, and number of samples before wraparound. All possible combinations of these factors can be accommodated using only ten special-case combinations. Any particular special case can be built out of one to four of the routines listed above. Routines are linked together by pushing their addresses onto the stack in the reverse of the order in which they are to be executed (the address of the exit routine is pushed first), and the first routine is called. In the example above, the first routine is called to fill memory until it detects buffer wraparound. It then loads the fill pointer with the address of the bottom of memory and executes an RTS (return from subroutine) instruction, which pops the address of the next routine from the stack and jumps to it. The next routine continues filling memory until the burst is complete, then terminates in another RTS instruction, which pops the address of the exit routine. The exit routine performs some minor cleanup (restoring pointers, setting flags, etc.) and leaves.

HP-IB Output. The high-speed output routine for the HP-IB uses some of the same concepts as the memory routines. In this case, the algorithm is as follows:

- Initialize pointers.
- Wait until the ADC has taken a reading, then enter the readings.
- Wait until the HP-IB buffer is ready to accept more data.
- Transfer the reading to the HP-IB buffer.
- Terminate the loop when the end-of-sequence command is sent.

The HP-IB buffer accepts a 16-bit word from the processor and sends the lower eight bits to the HP-IB interface chip. Once this byte has been transmitted, the HP-IB chip signals the buffer, and the buffer then sends the upper eight bits without intervention from the processor. Use of a buffer relieves a congestion point in the output data flow that would occur if the processor wrote directly to the HP-IB chip, since the HP-IB is an eight-bit bus while all other internal data paths are 16 bits wide. Using this scheme, the HP 3458A is able to offer complete memory and HP-IB functionality at the full speed of 100,000 16-bit dc voltage readings per second.

Summary

Achieving high throughput in a system DMM is a matter of designing the instrument as a system for moving data efficiently. Hardware and firmware must be designed as integral elements of this system, not as isolated entities. In the design of the HP 3458A, experience with DMM performance limitations provided invaluable insight into key areas of concern. As a result, significant improvements in throughput were achieved through the development of high-speed custom gate arrays for ADC control and interprocessor communications. Use of high-performance microprocessors and supporting hardware also contributed greatly to meeting design goals, as did the substantial investment in firmware design and development that was necessary to translate increased hardware performance into increased system performance.

Acknowledgments

We would like to thank Greg Wale and Dave Czenkusch, who were responsible for a significant part of the firmware development, and Vicky Sweetser, who was responsible for coordinating, testing and releasing firmware revisions. We also thank Bill Lutton and Jerry Metz of the HP 3235A design team for contributing parts of their firmware to be reused in the HP 3458A.

High-Resolution Digitizing Techniques with an Integrating Digital Multimeter

Capabilities and limitations of the HP 3458A Digital Multimeter as a high-resolution digitizer are summarized. Performance data is presented for selected applications.

by David A. Czenkusch

WITH ITS INTEGRATING analog-to-digital converter (ADC) capable of making 100,000 conversions per second, the HP 3458A Digital Multimeter (DMM) raises the possibility that, for the first time, a voltmeter can satisfy many requirements for high-resolution digitizing.

What are the characteristics of a high-resolution digitizer? Digitizing requires a combination of fast, accurate sampling and precise timing. It also needs a flexible triggering capability. The HP 3458A allows sampling through two different signal paths, each optimized for particular applications.

Converting a signal using the dc volts function (which does not use a sample-and-hold circuit, but depends on the short integration time of the ADC) provides the highest resolution and noise rejection. The direct sampling and subsampling functions, which use a fast-sampling track-and-hold circuit, provide higher signal bandwidth and more precise timing.

High-Resolution Digitizer Requirements

As the block diagram in Fig. 1 illustrates, a digitizer consists of an analog input signal conditioner followed by a sampling circuit. A trigger circuit and time base generator controls the timing of samples. The output of the sampling circuit is converted to a number by an analog-to-digital converter (ADC). Once converted to a number, the sample data can be processed digitally and displayed to the user.

Many types of instruments fit this definition of a digitizer, including digital oscilloscopes, dynamic signal analyzers, and digital multimeters (DMMs). Digitizing products can be roughly differentiated by four characteristics: analog signal bandwidth, sample rate, signal-to-noise ratio (which can be expressed as effective bits of resolution), and type of data displayed (time, frequency, etc.). In general, digital oscilloscopes tend to have high bandwidth and sample rate and relatively low resolution, while DMMs and dynamic signal analyzers tend to have much higher resolution and correspondingly lower bandwidth and sample rate.

Digital oscilloscopes are known for their high bandwidth, typically 100 MHz or greater, and their digitizing rates of 50 megasamples to one gigasample per second, making them useful for capturing very fast, single-shot events. Their resolution of five to eight effective bits is well-suited for displaying waveforms on a CRT, since one part in 200 is perfectly adequate for the human eye.

Dynamic signal analyzers, on the other hand, are used in applications that call for higher resolution—typically 10 to 14 bits. Examples include dynamic digital-to-analog converter testing, telecommunications, SONAR, and seismic or mechanical measurements that require digital signal processing. These applications require higher resolution and typically involve frequency-domain analysis. Therefore, to judge the attributes of a high-resolution digitizer, we should also examine the characteristics of discrete Fourier transforms (DFTs) performed on the digitizer's output data.

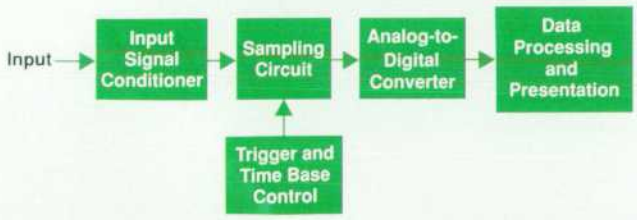

Fig. 1. *Generalized block diagram of a digitizer.*

Digitizer Spectral Attributes

"Effective bits" is a measure of the resolution of an ADC. Essentially, it is a measure of the signal-to-noise ratio in a

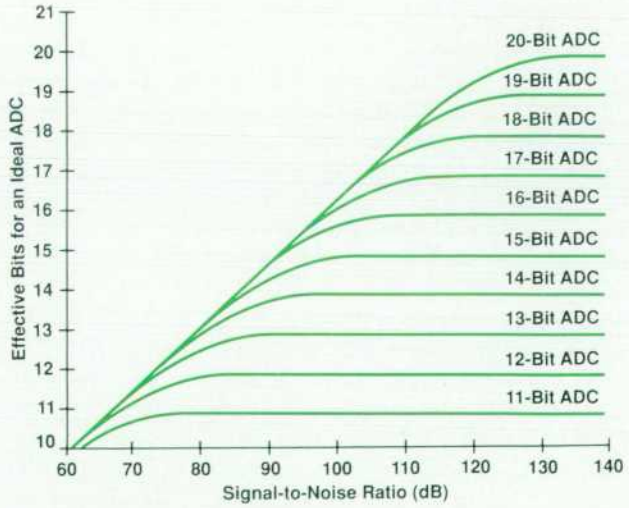

Fig. 2. *Analog-to-digital converter (ADC) effective bit limitation because of excess ADC noise and noise present in the ADC input signal.*

digitizing system expressed as a power of two. This can be expressed mathematically as:

$$N(\text{effective}) = (S/(N + D) - 1.8)/6.02$$

where S/(N + D) is the ratio of the signal power of a full-scale input to the total power of noise plus distortion, expressed in dB. Notice that the effective bits rating and the signal-to-noise ratio expressed in dB are both logarithmic scales related by the constant 6.02. This means that increasing the resolution of a measurement by one effective bit results in a 6-dB improvement in the signal-to-noise ratio. The system noise term, N + D, is the rms result of the power contributions of harmonic distortion and noise from various sources. For an otherwise noise-free, distortion free-system, there is minimum noise component because of the fundamental quantization error of the ADC. If this is the only source of error, the number of effective bits approaches the basic resolution of the ADC. Fig. 2 shows how the number of effective bits decreases as errors from other sources increase.

Other types of errors will appear as random noise. These include noise in the input signal, noise in the analog input circuits, random jitter in the timing of samples, and noise and differential nonlinearity in the ADC.

Linearity error is a measure of the deviation of the output of an ADC from the ideal straight-line relationship it should have with the input voltage. Fig. 3 shows a graph of the linearity error of a typical ADC as a function of input voltage. Integral linearity error is the large-scale bow in the total linearity error plot. This deviation from a straight line can often be described by a second-order or third-order function. Differential linearity error, on the other hand, has no large-scale structure, so it looks very much like noise.

If the noise in a digitizer is truly random, then a point-by-point average of many independent ensembles of waveform data taken with the same input signal will reduce this noise by the square root of the number of ensembles, provided the different ensembles of data have the same phase relationship to the input signal. Analog noise in the input amplifier and ADC and noise caused by random timing errors tend to be uncorrelated with the input signal, and so can be reduced by waveform averaging. On the other hand, differential linearity error in the ADC and systematic timing errors, while appearing like random noise in a single pass of data, are repeatable from pass to pass, and so are correlated with the input and cannot be reduced by averaging. This provides a way of determining if the signal-to-noise ratio of a given digitizing system is dominated by input noise or by differential linearity error.

Effective Bits from the DFT

One way to characterize the signal-to-noise ratio of a digitizer is to sample a quiet (low-noise) and spectrally pure full-scale sine wave and perform a discrete Fourier transform (DFT) on the resulting data. The dynamic range (in dB) from the peak of the fundamental to the noise floor of the DFT gives an idea of the low-level signals that can be resolved. The level of the noise floor depends on the number of frequency points (bins) in the DFT, and hence on the number of samples taken, since if the same noise power is spread over more frequency bins, there will be less noise power per bin.

The DFT spectrum can be used to produce an estimate of the signal-to-noise ratio of a digitizer by performing essentially the same measurement digitally that a distortion analyzer performs electronically. A distortion analyzer supplies a low-distortion sine wave as the input to a circuit under test. A notch filter is used to remove the fundamental frequency from the output signal. The power in the filtered signal is measured and a ratio is formed with the total output power of the circuit under test. A distortion analyzer measurement assumes that the power in the filtered output signal is dominated by harmonic terms generated by distortion in the circuit under test. In practice, however, the analyzer is unable to separate this power from the power contribution of wideband noise, and hence is actually measuring the signal-to-noise ratio of the output signal.

An analogous operation can be performed on the DFT spectrum of a digitized pure sine wave. A certain number of frequency bins on either side of the fundamental peak are removed from the DFT data. The data in each of the other frequency bins is squared (to yield a power term) and summed with similar results from the other frequency bins to calculate the total noise power. The data within the narrow band around the fundamental is squared and summed to give the total signal power. The ratio of these two terms,

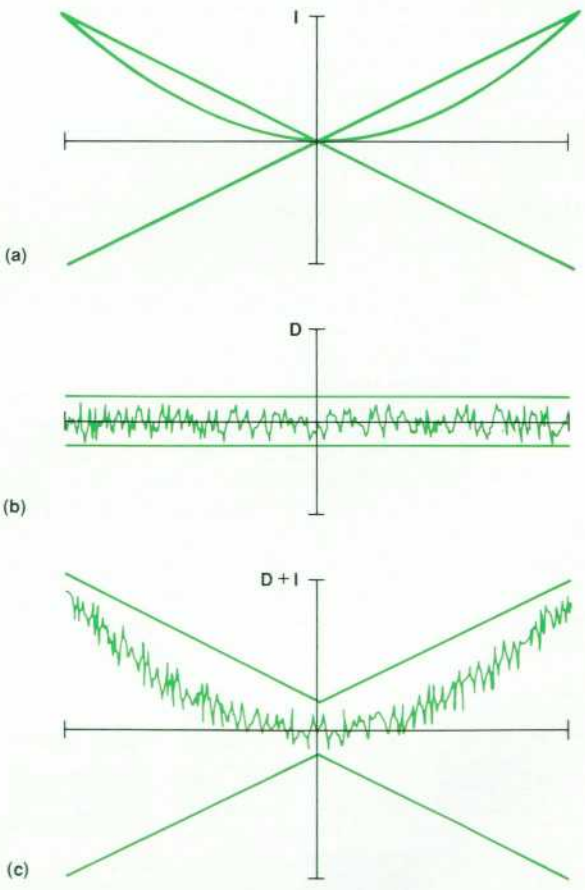

Fig. 3. *Linearity errors in an ADC. (a) Integral linearity error. (b) Differential linearity error. (c) Total linearity error.*

expressed in dB, can be used to compute the number of effective bits of resolution of the digitizer.

Calculations of effective bits from DFT spectra will show variations if the test is performed repeatedly. This variation can be reduced if the spectral values from many independent trials are averaged point by point (as opposed to averaging the time-domain data). Spectral averaging will not reduce the level of the noise floor in the DFT data, but only the amount it varies. Therefore, if enough ensembles of spectral data are averaged, the number of effective bits calculated will converge to a single number.

Fig. 4 shows the DFT for 4096 samples of a mathematically generated ideal sine wave quantized to 16 bits (±32,767 counts). From this, we see that a perfect 16-bit digitizer will show a noise floor of about −127 dB when quantization error is the only source of noise. If the signal-to-noise ratio is calculated using the method described above, the result is 97.0 dB, or 16.0 effective bits, which is what we would expect.

Other types of digitizer errors can show up on a DFT plot. Distortion reveals itself as harmonic components at multiples of the fundamental input frequency. This can be distortion in the input signal, harmonic distortion in the input amplifier, or integral nonlinearity in the ADC. As mentioned before, integral linearity error can be approximated by a second-order or third-order term in the transfer function of the ADC. These higher-order terms generate spurious harmonic components in the DFT spectrum.

Other spurious signals can show up in the DFT spectrum besides harmonic distortion. Internal clock signals can produce unwanted signal components (spurs) either by direct cross talk or through intermodulation with the input signal. These effects are commonly grouped together into a single specification of spurious DFT signals.

Effect of Sample Aperture

Another aspect of digitizers that should be considered is the effect of the finite acquisition time of the sampling circuit that provides the input to the ADC. This is typically some type of sample-and-hold or track-and-hold circuit. For maximal time certainty, an ideal track-and-hold circuit would acquire a voltage instantaneously when triggered to take a sample. In reality, of course, all sampling circuits require some finite time to acquire a sample. This sampling function can be approximated by a rectangular window of time T over which the input signal is sampled.

The Fourier transform of a square pulse defined over the interval $-T/2 \leq t \leq T/2$ in the time domain has the form $[\sin(\pi fT)]/\pi fT$, which is the familiar function sinc(fT). This means that sampling a signal for a time T is equivalent in the frequency domain to multiplying the input spectrum by the function sinc(fT). Fig. 5 shows that the spectral envelope of the sinc function approximates a single-pole low-pass filter with a 3-dB corner frequency of $f_c \approx 0.45/T$.

From this analysis we can conclude that making the sample time as short as possible produces the flattest possible response because it maximizes the aperture roll-off corner frequency. A less desirable trade-off, however, is that this also increases the equivalent white noise bandwidth of the sampler, thereby increasing its sensitivity to noise. Therefore, in applications where noise is a greater problem than frequency roll-off, it would be desirable to have a wider sample aperture to reduce the noise bandwidth.

The transform above was defined for a square pulse extending from $-T/2$ to $T/2$. Since a real sampler cannot anticipate its input, the sample must actually occur over the interval $0 \leq t \leq T$. This implies that any sampler that acquires a signal over a nonzero time interval T will introduce an apparent time delay equal to T/2 to the output. In most real applications, however, this distinction is not significant.

Another characteristic of the sinc function that can be useful is that its transfer function goes to zero at all frequencies that are multiples of 1/T. This means the sampler will reject all harmonics of a signal whose fundamental period is equal to the sample aperture. Therefore, a selectable aperture allows the rejection of specific interference frequencies that may be present in the measurement environment.

HP 3458A Digitizing Characteristics

Many of the same design characteristics required to make

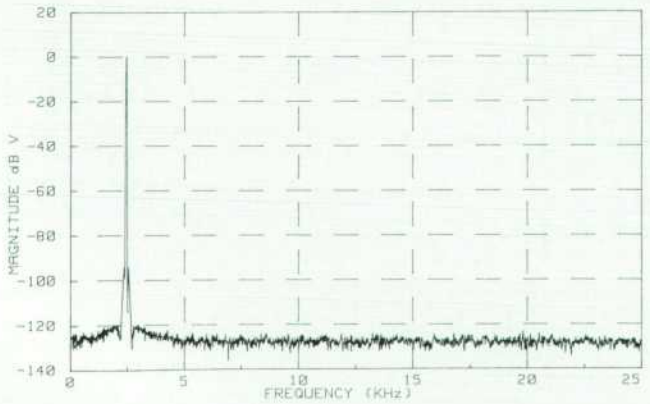

Fig. 4. *Discrete Fourier transform of an ideal sine wave sampled with an ideal 16-bit ADC.*

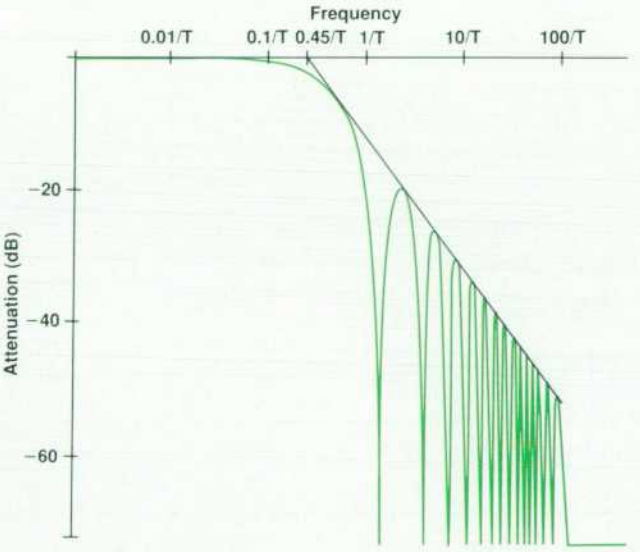

Fig. 5. *Attenuation of the input signal as a function of frequency resulting from sampling with an aperture of width T.*

Time Interpolation

To implement the subsampling (effective time sampling) required for the HP 3458A DMM's digital ac measurement technique, some means of synchronization with the input signal was necessary. To minimize errors caused by aliasing of the sampled data, a time base with 10-ns resolution was desired. However, the internal 10-MHz clock would only allow a sample resolution of 100 ns relative to a synchronizing trigger event. These design requirements dictated the development of the time interpolation circuit of the HP 3458A.

The instrument's 10-MHz clock is used to generate sample timing pulses of variable period in 100-ns (10-MHz) steps. The time interpolator extends the resolution of the time base from 100-ns steps to 10-ns steps for initial burst delays (the delay from a trigger event to the start of sampling). This enables the HP 3458A to digitize signals with spectral content up to 50 MHz without introducing aliasing errors.

The time interpolator, Fig. 1, uses analog techniques to convert time to stored charge on a capacitor. Before an input trigger, the interpolator is reset by shorting both capacitors (S1 and S2 closed) with the current source shorted to ground (S3 and S4 in position B). An asynchronous input trigger, generated either by the ac path's trigger level circuit or by an external trigger input, initiates charge accumulation on C1 by opening S1 and setting S3 and S4 to position A. This charge accumulation process continues until the next positive edge of the 10-MHz clock occurs.

On this edge S3 and S4 switch to position B, forcing the accumulated charge to be held on C1. This charge, Q_1, is directly proportional to the elapsed time (T_{var1}) between the input trigger and the next 10-MHz clock edge. Likewise, the voltage across C1 (V_{var1}) is also proportional to T_{var1}, which varies between 50 ns and 150 ns depending on the timing of the asynchronous input trigger relative to the internal 10-MHz clock.

The interpolator remains in this "hold" state for an integral number of clock cycles, T_{delay}. The next positive-going clock edge after T_{delay} initiates the second charge accumulation process. At this time, S2 opens and S3 and S4 are switched to position A. During this time, the same charge, Q_2, is accumulated on C1 and C2. This process continues until the voltage on C1, V_{ramp}, crosses the programmable comparator threshold V_t. This transition generates an output trigger that signals the track-and-hold circuit in the ac section to enter hold mode, thus acquiring a sample for subsequent ADC conversion. By programming V_t to various values, the system can alter this delay in increments of 10 ns, allowing precise timing of a burst of samples relative to an asynchronous starting event.

The output trigger also switches S3 and S4 to position B. This not only turns off the current source, but also creates a loop between C2, R1, and the buffer amplifier's input and output. Feedback forces a current through C2, removing its accumulated charge, Q_2. The resulting current flows through both C1 and C2, removing the charge Q_2 from capacitor C1. The process completes with C1 holding the original charge, Q_1, which is proportional to the delay between the first trigger and the rising edge of the internal 10-MHz clock. During the ADC conversion, (a

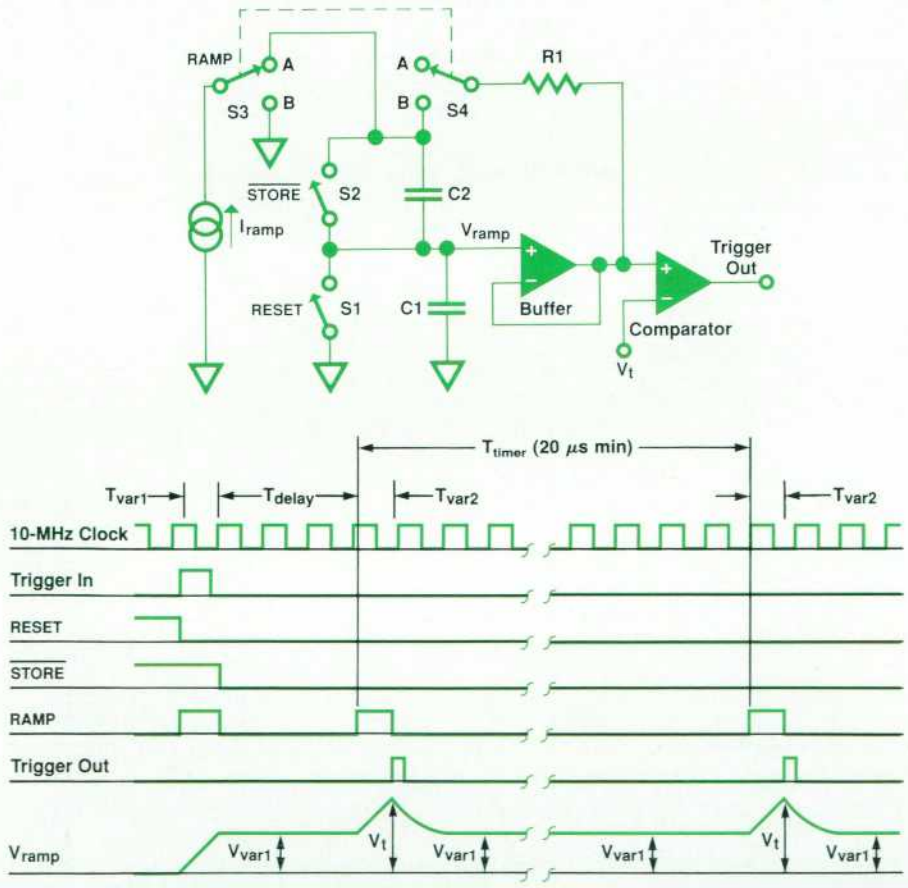

Fig. 1. *HP 3458A DMM time interpolator block and timing diagrams.*

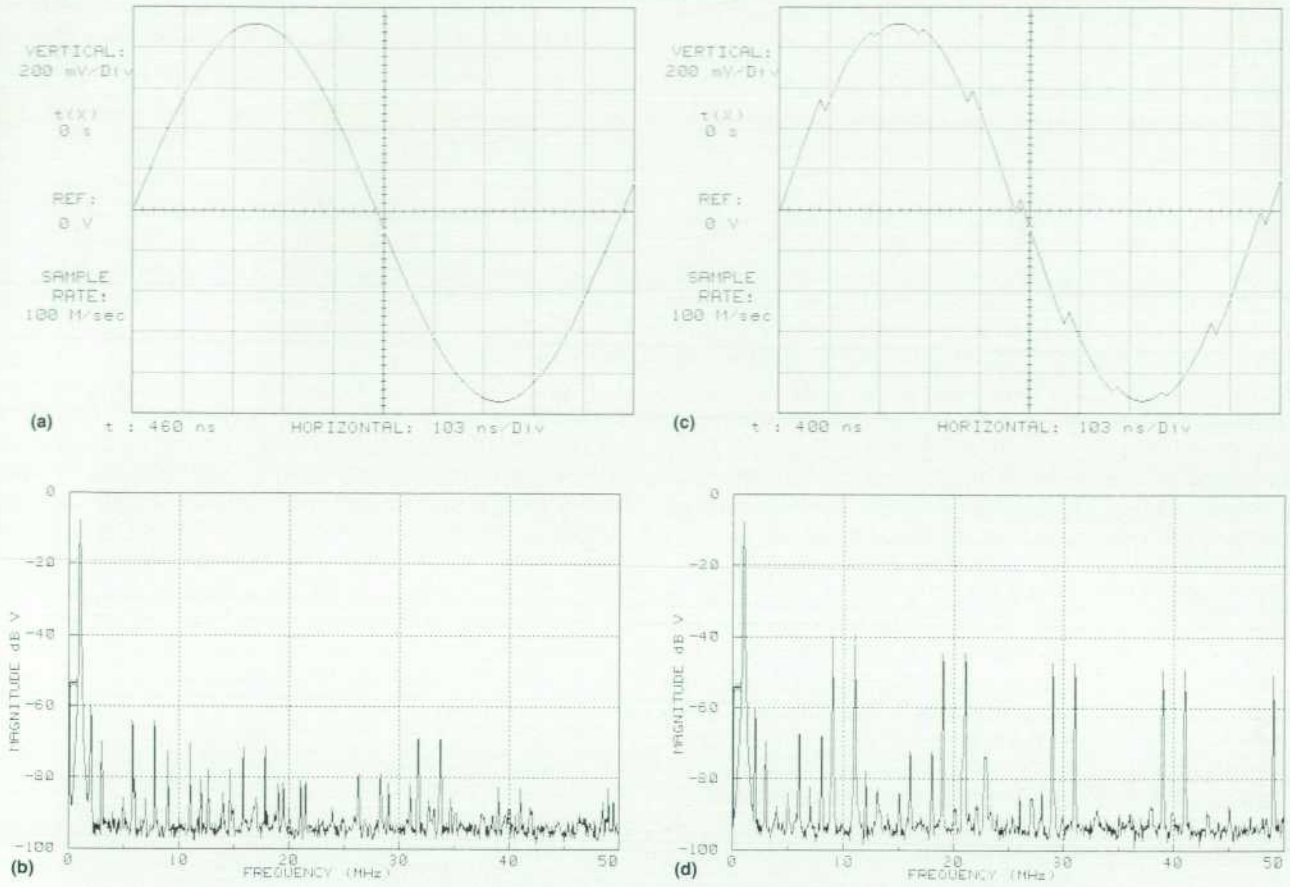

Fig. 2. *The time interpolator's accuracy is adjusted by calibrating I_{ramp}. (a) A digitized 1-MHz waveform after I_{ramp} calibration. (b) Fourier transform of (a), showing a noise floor 80 dB below the fundamental and spurious signals below −55 dB. (c) A digitized 1-MHz sine wave with I_{ramp} misadjusted. (d) Fourier transform of (c).*

minimum of 20 μs for subsampling) the time base circuit waits an interval T_{timer} before repeating the charge/discharge cycle.

The accuracy of the 10-ns increments is ensured by calibration of the circuit gain. Since the time interpolator's absolute delay is a function of I_{ramp}, C1, and V_t, many variables can prevent the 10-ns increments from being exactly one tenth of a 100-ns time base step. Interpolation for ten 10-ns intervals must precisely equal one 100-ns clock period (10 MHz) to minimize sampling errors. By adjusting I_{ramp} (Fig. 2), the slew rate and threshold errors are adjusted to yield 10-ns steps within ±50 ps. Time jitter is held to less than 100 ps rms. Low temperature coefficients for C1 and the DAC that generates V_t ensure interpolator accuracy over the operating temperature range. The time interpolator is adjusted by applying a 2-MHz sine wave to the input and executing a calibration routine which alternately programs 100-ns delays into either the time base or the time interpolator. By adjusting the DAC that controls I_{ramp}, the routine converges the two delays. This time base performance contributes to the a noise floor 80 dB below the fundamantal and spurious signals below −55 dB.

The design of the time interpolator circuit was refined using analog simulation methods on an HP 9000 Model 320 Computer. Computer-aided engineering provided timely feedback during development, allowing rapid evaluation of alternative circuit topologies. Critical design characterizations, difficult to achieve by traditional means, were performed accurately and simply using CAE simulations. The resulting circuit performance exceeded our original design goals.

David E. Smith
Development Engineer
Loveland Instrument Division

high-accuracy ac and dc measurements also allow the HP 3458A to perform well as a high-resolution digitizer. For instance, because it makes true-rms ac measurements using digital techniques, it has a scope-like trigger level circuit for waveform synchronization. A precise trigger timing circuit allows sample intervals to be specified to a resolution of 100 nanoseconds and initial delays from a trigger event to the first sample of a burst can be specified to a resolution of 10 nanoseconds using an analog time interpolator.

As the block diagram in Fig. 6 shows, the HP 3458A provides two distinct input paths for digitizing, corresponding to the two amplifiers used for the dc volts and ac volts functions. Each path has advantages and disadvantages. The dc input path should be used when maximum resolution and noise rejection are required and the bandwidth of the input signal is relatively low. Because it

uses a track-and-hold circuit, the ac input path can be used on signals of higher bandwidth or when the signal must be sampled at a very precise point in time.

High-Resolution DC Input Path

The dc input path allows higher-resolution sampling as well as a higher single-shot measurement speed, providing 16-bit samples at up to 100,000 samples per second. The bandwidth of this amplifier varies from 50 kHz to 150 kHz, depending on the range selected. The widest bandwidth is available on the 10V range, when the amplifier is operating at unity gain. In this path, the sampling function is performed by the ADC itself with its selectable integration time (sample aperture). Historically, digital multimeters with integrating ADCs have allowed only a few discrete values for integration time. These values were chosen to be multiples of the power-line frequency—the most common signal to interfere with a high-resolution voltage measurement. In the HP 3458A, integration times can be specified from 500 ns to 1 s in increments of 100 ns. This allows the rejection of an interference signal of arbitrary frequency that may be present in the input, and provides attenuation of other frequencies above the sample rate by the approximate single-pole roll-off characteristic of the sample aperture's sinc function. The longer the integration aperture specified, the more resolution is provided by the ADC. Fig. 7 shows the resolution that can be obtained for a given aperture.

Because the dc input path is designed for extremely low noise, low offset, and part-per-million (ppm) accuracy, the DFT spectra produced in this mode are quite good. In fact, it is difficult to determine whether the harmonic distortion and noise floor measurements are dominated by the HP 3458A or by the input signal.

Fig. 8a shows the DFT calculated on 4096 samples of a 1-kHz waveform aquired at a rate of 50,000 samples/s with an integration time of 10 microseconds. The noise floor and spurious DFT signals are below −120 dB, and harmonic spurs are below −106 dB. If the signal-to-noise ratio is computed from the spectral data, the result is approximately 93.9 dB, yielding 15.3 effective bits.

The input signal for this test was provided by the oscillator output of an HP 339A Distortion Measurement Set, whose distortion at this frequency is specified to be less than −96 dB at the first harmonic. It is unclear whether the first-harmonic term at −107 dB is caused by distortion in the source signal or distortion in the HP 3458A at this sample rate. However, tests performed at slower sample rates (and greater resolution) also exhibit this second-harmonic term.

The averaging effect of the relatively wide sample aperture (10 μs vs. 2 ns) reduces random noise contributions to the DFT noise floor to a level comparable to those of systematic nonlinearities. Because of this, waveform averaging only provides an extra 4.4 dB improvement in the signal-to-noise ratio, yielding an extra 0.7 effective bit. Fig. 8b shows the DFT spectrum that results if 64 waveforms are averaged.

A striking example of the high-resolution digitizing capability of the dc volts sampling mode involves measuring an ultralow-distortion signal source used to characterize the performance of seismic measurement systems. The output of the source is a 0.03-Hz sine wave whose noise and harmonic distortion products are guaranteed by design to be at least 140 dB below the level of the fundamental. Superimposed on this is a 1-Hz sine wave whose amplitude is 120 dB below the level of the 0.03-Hz signal. The goal of the measurement system two-tone test is to be able to see the 1-Hz tone clearly in the presence of the full-scale

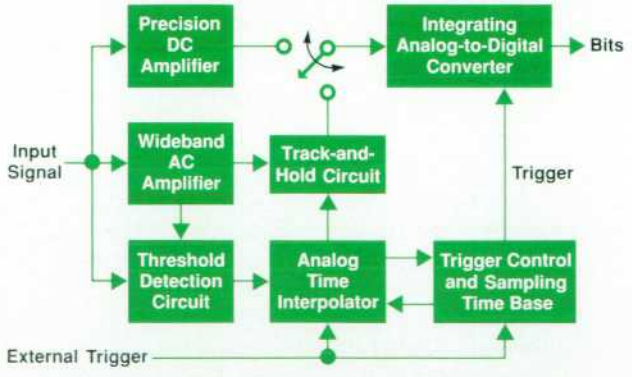

Fig. 6. *HP 3458A block diagram, showing the two measurement paths.*

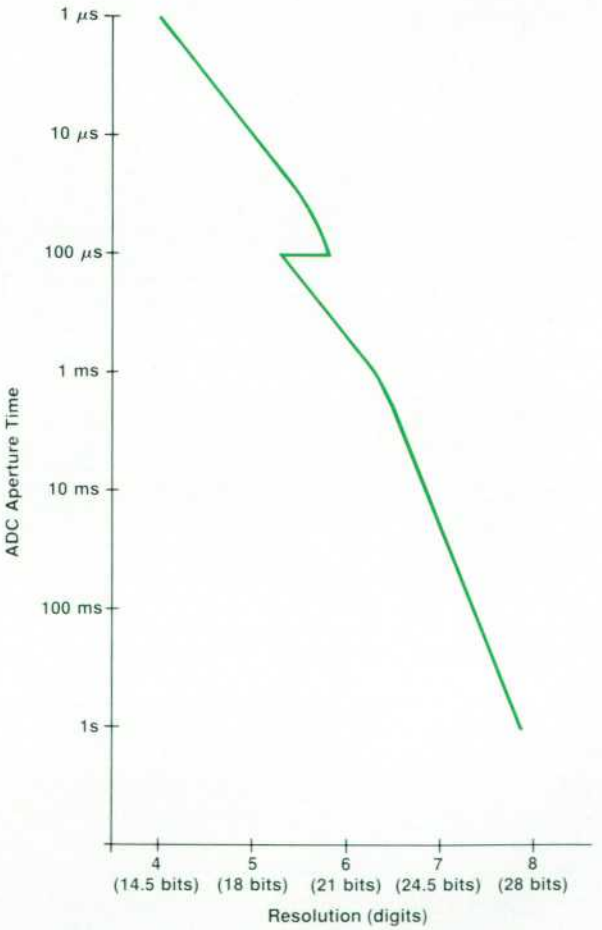

Fig. 7. *HP 3458A measurement resolution as a function of aperture time (speed).*

(10V peak) input without introducing extraneous distortion products. The HP 3458A was used to acquire 4096 samples with an ADC aperture of 20 milliseconds and a sample interval of 50 milliseconds, resulting in a resolution of 24 bits (7½ digits).

The DFT plot in Fig. 9 shows the result of this test. Only a portion of the full 10-Hz bandwidth is shown to make the component at 0.03 Hz more apparent. The 1-Hz spike at −120 dB is clearly visible above a noise floor of −150 dB. If the 1-Hz component is notched out along with the 0.03-Hz fundamental, and the remaining power is considered noise, a signal-to-noise calculation yields 19.6 effective bits. As before, it is not clear whether the DFT noise floor in this measurement is dominated by noise in the input signal or noise in the HP 3458A. If the rms noise of the HP 3458A is characterized with the same ADC aperture (20 ms) and a quiet dc source is substituted as input, measurements demonstrate a performance of 22 effective bits. The HP 3458A is clearly capable of verifying the performance of this source to the levels guaranteed by its designers. We are told that earlier measurements had never been able to achieve these low levels of noise and distortion.

In the dc volts mode, the input signal is sampled directly by the ADC. The sampling is synchronous with the instrument's internal 10-MHz clock. This leads to a 100-nanosecond peak uncertainty in the time latency of a sample or group of samples relative to an external or level trigger event. While a time uncertainty of 100 nanoseconds from an asynchronous trigger event is perfectly adequate for most applications, other applications require more precise sample timing.

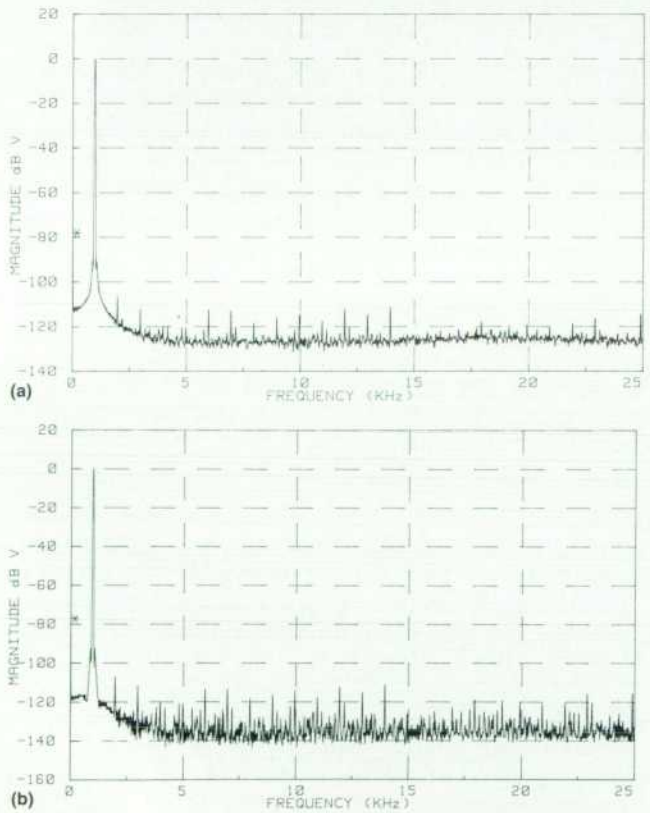

Fig. 8. *(a) Single-shot, 4096-time-sample discrete Fourier transform (DFT) of a 1-kHz input signal. (b) DFT for 64 averaged acquisitions of the 1-kHz input signal. Effective bits are 15.3 for (a) and 16.0 for (b).*

Digital AC Input Path

The ac input path provides a wider analog bandwidth and more precise timing than the dc path. The bandwidth of the ac amplifier is 12 MHz on all ranges except the 10-mV and 1000V ranges, where the bandwidth is 2 MHz. Autocalibration guarantees a frequency response flatter than 0.01% (0.001 dB) throughout the frequency band from 200 Hz to 20 kHz, making this path ideal for characterizing frequency response in the audio band. While the maximum single-shot sample rate of 50,000 samples per second is somewhat lower than the dc input path because of the additional settling time required by the track-and-hold circuit, a precise timing circuit allows effective time sampling (subsampling) of repetitive input signals with effective sample intervals as short as 10 ns.

Achieving true-rms ac measurements with 100-ppm accuracy using digital techniques requires an extremely linear track-and-hold circuit. This same track-and-hold circuit provides 16-bit linearity in digitizing applications. A sample acquisition time of approximately 2 ns results in a 3-dB aperture roll-off frequency of at least 225 MHz. This means that amplitude errors caused by the sample aperture are insignificant through the entire measurement band. The timing of the track-and-hold circuit is controlled by an analog ramp interpolator circuit which operates asynchronously with the internal 10-MHz clock, giving a burst-to-burst timing repeatability error less than 100 picoseconds. The time interpolator allows programming of delays from an external or internal trigger with a resolution of 10 ns, allowing single samples to be timed very precisely.

While the greater equivalent noise bandwidth of the input amplifier and track-and-hold circuit results in fewer effective bits of resolution in a single-shot measurement than the dc input path, the DFT performance for this path is still quite good. Fig. 10a shows a typical 2048-point DFT plot for a 1-kHz sine wave sampled at the single-shot limit of 50,000 samples per second. A signal-to-noise ratio calculation on this data yields 10.4 effective bits. The ac input path has a greater equivalent noise bandwidth than the dc input path, so random noise dominates the signal-to-noise

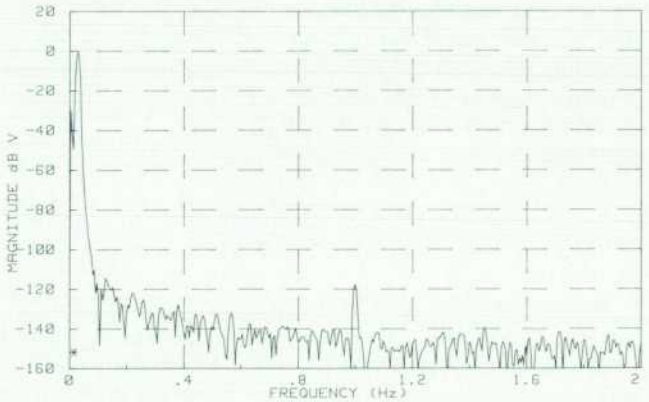

Fig. 9. *DFT for a 4096-sample HP 3458A acquisition of a 0.3-Hz sine wave with a 1-Hz, −120-dB sine wave superimposed. The effective bits rating is 19.6.*

Measurement of Capacitor Dissipation Factor Using Digitizing

No capacitor outside of a textbook exhibits the theoretical current-to-voltage phase lag of 90 degrees. This is another way of saying that in the real world all capacitors are lossy to some extent. These losses are caused by a number of factors, such as lead resistance and dielectric hysteresis.

At a given frequency, the dissipation factor of a capacitor is defined to be the ratio of the equivalent series resistance (ESR) and the capacitive reactance. Dissipation factor is important for many applications. At high power levels, capacitors with poor dissipation factor can overheat. The precision of capacitively compensated attenuators can be compromised by dissipation factor. Also, the capabilities of track-and-hold circuits are degraded by the dissipation factors of their hold capacitors.

There are two common ways to measure dissipation factor. In the first method, the impedance of the capacitor under test (CUT) is measured at a given frequency and the deviation in phase angle from the ideal 90 degrees is used to calculate the dissipation factor. Bridges are another method used to measure dissipation factor. In essence, the CUT is in a bridge with three other capacitors, one of which is adjustable in both C and ESR. When the bridge is nulled, the values of the adjustable C and its ESR determine the dissipation factor of the CUT.

The ac attenuator in the HP 3458A uses a 20-pF capacitor that has a dissipation factor requirement of 0.0001 (0.01%) at 10 kHz. Commercially available automated equipment exhibits reading-to-reading noise of 0.01% and dissipation factor accuracies of 0.04%. This is inadequate to screen this capacitor reliably. High-quality manual bridges can do this job, but their operation is not well-suited to a production environment.

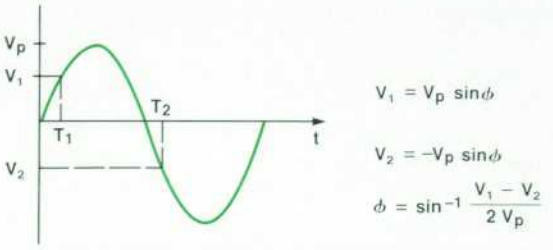

Fig. 1. *Measuring phase shift in a sine wave.*

By making use of the high-resolution digitizing and precision ac measurement capabilities of the HP 3458A, it is possible to construct an automated dissipation factor meter that is capable of making accurate and stable 0.001% dissipation factor measurements and capacitance measurements that are stable to 0.001 pF.

Circuit Description

In Fig. 1, a method of determining the phase shift of a sine wave relative to an external timing pulse occurring at the sine wave's zero crossing is shown. Theoretically, only V_1 is needed to determine this phase shift. The advantage of using a second sample (V_2) spaced one half cycle later in time is that ($V_1 - V_2$)

measurement to a much greater extent. Because of this, the noise floor can be lowered another 20.6 dB by waveform averaging, producing 13.8 effective bits as shown in Fig. 10b.

The ac input path supports two digitizing functions: direct sampling and subsampling, which is also referred to as effective time sampling. The article on page 15 describes the subsampling technique. Subsampling allows the sampling of repetitive waveforms with effective sample intervals as short as 10 ns, thus allowing the user to take full advantage of the 12-MHz analog input bandwidth. The subsampling parameters are somewhat complex to calculate for an arbitrary effective interval and number of samples, but the user need not understand the details of the algorithm. All that need be specified is the desired effective sample interval and number of samples, and the HP 3458A will compute the number of passes, the number of samples per pass, the delay increment per pass, and the ADC sample rate required to complete the task most efficiently. Furthermore, if the samples are directed to the instrument's internal memory, they will be sorted into the correct time order on the fly.

If the number of samples required for a subsampled measurement exceeds the size of the instrument's internal memory, the samples can be sent directly from the ADC to a computer via the HP-IB. Since the HP 3458A cannot sort the data in this mode, the samples recieved by the computer generally will not be in the correct time order. If this is the case, the waveform can be reconstructed in the computer's memory using an algorithm requiring three sorting parameters supplied by the HP 3458A.

Subsampling is essentially the same as direct sampling when the effective sample rate is less than or equal to 50,000 samples per second. Why, then, is direct sampling even offered? The answer is that the subsampling technique only allows sampling based on the internal time base, whereas the direct sampling function includes all the same trigger modes as the dc volts function. This means that the user can supply an external time base via the external trigger input to allow sampling at odd frequencies that cannot be realized with the 100-ns quantization of the internal time base. An example would be the 44.1-kHz sample rate required by many digital audio applications. Direct sampling is also useful for acquiring single samples with minimum time uncertainty. Samples can be precisely placed with 10-ns delay resolution relative to an external trigger event and with 2-ns rms time jitter. "Measurement of Capacitor Dissipation Factor Using Digitizing" on this page shows an example of these measurement capabilities of the HP 3458A.

HP 3458A Limitations

Since the HP 3458A must be a voltmeter first and a digitizer second, it is not surprising that it has some limitations as a digitizer. Perhaps the most significant is the lack of an anti-aliasing filter. Because no single filter could be included to cover all possible sample rates, and because it would degrade the analog performance, the design team

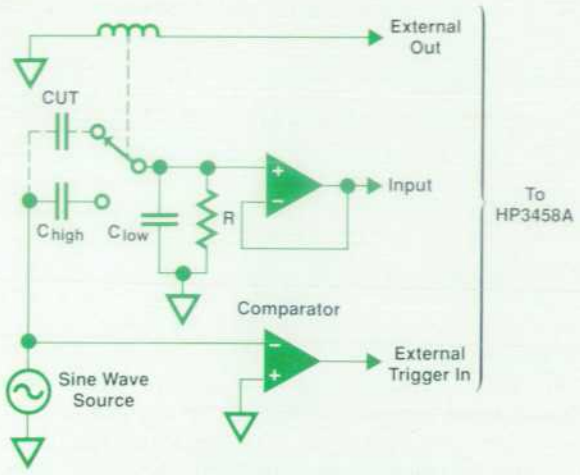

Fig. 2. *Circuit to measure dissipation factor.*

is insensitive to voltage offsets on the sine wave.

Fig. 2 shows a circuit using the technique of Fig. 1. A sine wave is applied to one of two capacitive dividers, one formed by the CUT and C_{low} and the other formed by C_{high} and C_{low} (R provides the dc bias for the buffer amplifier). This sine wave is also applied to a comparator that detects zero crossings. The output of the comparator is routed to the external trigger input of the HP 3458A and the output of the buffer amplifier is applied to the input of the HP 3458A. The HP 3458A can use its ac section to measure the rms value of this output waveform and thus V_p in Fig. 1 can be determined very precisely. The HP 3458A can also measure the period of the output waveform and set up sample timing parameters to sample the output sine wave relative to the external trigger signal as shown in Fig. 1. Thus all the information is present to determine the phase shift of the sine wave through the capacitor divider network.

The absolute phase shift of one side of the capacitor divider is not the information desired, however. What is desired is the phase shift caused by the dissipation factor of the CUT in the divider formed by the CUT and C_{low}. This will provide the information needed to determine the dissipation factor of the CUT.

Computing the difference between the absolute phase shift of the reference divider (C_{high} and C_{low}) and the input divider (CUT and C_{low}) is the first step towards determining the phase shift in the input divider resulting from the dissipation factor of the CUT. The HP 3458A's **EXT OUT** output is used to select either the reference divider or the input divider. Taking the phase difference between the reference and input measurements removes errors caused by the buffer amplifier and the comparator. If C_{high} had zero dissipation factor, CUT had the same capacitance value as C_{high}, and the switching relay was perfect, this phase difference would be entirely a result of the dissipation factor of the CUT. If this phase difference is ϕ, the dissipation factor of the CUT is:

$$DF = \tan(\phi)\ \frac{(CUT + C_{low})}{C_{low}}.$$

In general, the CUT will not be the same size as C_{high}, C_{high} will not have zero dissipation factor, and the switching relay will not be perfect. However, these conditions are easily controlled. The feedthrough capacitance of the relay in Fig. 2 can be reduced by implementing the relay as a T-switch. If the CUT is different in magnitude from C_{high}, a phase difference will be measured even if the CUT has zero dissipation factor. This is because the phase shift of the parallel combination of R and C_{high} and C_{low} is different from that of the combination of R and the CUT and C_{low}. This error can be removed by appropriate correction factors implemented in software. Also, in general, the dissipation factor of the CUT will not be zero. A zero calibration against a reference capacitor can remove this error.

Summary

The precision digitizing capabilities of the HP 3458A DMM have been applied to make a traditionally difficult measurement of capacitor dissipation factor. Test results show measurement accuracies approaching 0.001%. This corresponds to a phase error of 0.0005 degree or a time error of 150 ps at 10 kHz. Also, since the capacitance of the CUT is computed as part of the dissipation factor calculation, accurate capacitance measurements are also generated that are stable to 0.001 pF.

Ronald L. Swerlein
Development Engineer
Loveland Instrument Division

decided it would be impractical to include one.

Another limitation is the latency from an external or internal trigger to the start of sampling. The ramp time of the analog time interpolator produces a minimum delay of at least 400 ns. This means that if the input frequency is greater than about 500 kHz, the signal edge that is used to synchronize the waveform in a subsampled measurement will not even show up in the output data. Oscilloscopes typically include some form of analog delay to match the timing of the signal path to the trigger circuit, but again this was not compatible with the requirements of a high-precision DMM.

Another effect inherent in the design of the analog time interpolator is voltage droop. Essentially, the phase of the input signal relative to the internal 10-MHz clock is represented by a voltage stored on a hold capacitor, which is captured at the beginning of a measurement burst and held throughout the burst. Since there will always be some leakage in the circuits attached to this node, the voltage on this capacitor will slowly leak off, causing an apparent lengthening in the time between samples. This produces an apparent frequency modulation in the output data, which continues until the charge leaks off completely, at which time the sample interval will again be stable. This droop rate gets worse as leakage increases with higher temperature. Measurements on a typical unit at room temperature show a droop rate of about 500 ns/s, which persists for about 140 ms. In other words, during the first 140 ms of a reading burst, a sample interval of 20 μs will be lengthened by about 10 ps per sample.

Waveform Analysis Software

One factor limiting the effectiveness of the HP 3458A as a stand-alone digitizer is the lack of a built-in CRT for waveform display. This shortcoming has been addressed with a software library that turns an HP 3458A and a computer into a real-time single-channel digital oscilloscope and DFT analyzer.

The optional waveform analysis library allows a user with an HP 9000 Series 200 or 300 workstation or an IBM PC/AT-compatible computer with HP BASIC Language Processor to display waveforms in real time. In addition, routines are included to perform parametric analysis, waveform comparisons, and FFT spectral calculations and to store and recall waveforms from mass storage.

The real-time oscilloscope subprogram, Scope58, began as a means to demonstrate how quickly waveforms could be acquired by the HP 3458A and displayed. It soon became an indispensable tool in the development of the ADC and high-speed firmware. Since the program had proven so valuable during development, we decided it should be included in the waveform analysis library. A user interface was added to give the look and feel of a digital oscilloscope, including horizontal and vertical ranging, voltage and time markers, and an FFT display mode. The program can acquire and plot waveforms at a rate of approximately 10 updates per second—fast enough to provide a real-time feel.

The heart of the Scope58 subprogram is a set of specialized compiled subroutines for fast plotting, averaging, and interpolation of waveforms. Since speed was the overriding design consideration for these routines, most of these subroutines were written in MC68000 assembly language rather than a higher-level language like Pascal or BASIC. The fast plotting routine, in particular, required certain design compromises to achieve its high speed. It uses a simplified plotting algorithm which requires that there be one sample per horizontal display pixel, which means that the only way to change the horizontal scale is to change the sample rate unless the waveform data is interpolated to increase its time resolution before plotting. Also, the plotting routine bypasses the machine independent graphics routines and writes directly to the bit-mapped frame buffer of the graphics screen. This makes the routine fast, but it complicates the programming task, since a special version of the routine must be written for every supported display interface.

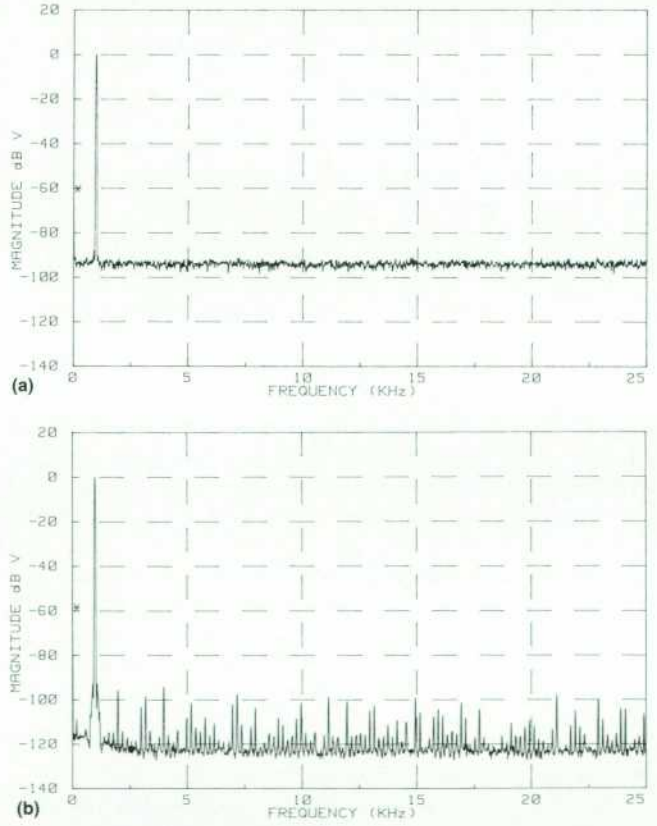

Fig. 10. *(a) Typical DFT for 2048 samples of a 1-kHz sine wave sampled at the HP 3458A ac path's single-shot limit of 50,000 samples per second. Effective bits are 10.4. (b) Effective bits can be increased to 13.8 by averaging data for several acquisitions.*

In addition to the Scope58 subprogram, the waveform analysis library includes routines that help with waveform acquisition, analysis, and storage. Since the HP 3458A is capable of synchronizing with external switching instruments like a normal DMM, it can be switched to acquire a waveform per channel in a multichannel data acquisition system. This feature, combined with the waveform analysis library, can be used to make many complex measurements in automated test applications.

The library's analysis capabilities include routines to extract parametric data such as rise time, pulse width, overshoot, and peak-to-peak voltage, and routines to compare waveforms against high and low limit arrays. There is also a compiled utility for calculating Fourier and inverse Fourier transforms. This routine can compute a 2048-time-point-to-1024-frequency-point transform in as little as 1.2 s if the computer's CPU includes a 68881 floating-point coprocessor. Finally, routines are provided for the interpolation of waveforms using the time convolution property of the sinc(x) function. This technique is common in digital oscilloscopes, and allows the accurate reconstruction of waveforms with frequency components approaching the Nyquist limit of half the sampling frequency.

The precision digitizing characteristics of the HP 3458A and the display capabilities of the waveform analysis library combine to form a powerful waveform analysis tool in R&D or automated test applications. For instance, an HP 3458A together with a digital pattern generator can be used to test digital-to-analog converters (DACs). The waveform comparison capability of the waveform analysis library can be used to provide a pass/fail indication. Assuming a DAC settling time of 10 μs and an HP 3458A measurement time of 20 μs (only 10 μs of which is spent integrating the input signal), all codes of a 14-bit DAC (16,384 levels) can be acquired in approximately 328 ms.* The dynamic characteristics of the DAC can be tested using the FFT library routine. The DAC can be programmed to output a sine wave, which the HP 3458A can digitize. A DFT on the resulting data can be analyzed to characterize the DAC for noise floor and total harmonic distortion (THD).

Summary

The capabilities of a high-resolution digitizer can best be characterized by examining its performance in the frequency domain. To be able to resolve very low-level phenomena, it must have a wide dynamic range and very low levels of distortion and spurious signals. The excep-

*If you multiply 16,384 by 30 μs, the result is actually 492 ms. However, for at least 10 μs of each ADC conversion, the HP 3458A is not measuring the input, and provides a TTL signal indicating this fact. This time can be overlapped with the DAC's settling time, thereby reducing the total acquisition time.

tional DFT performance of the HP 3458A results from its combination of precise timing and the nearly ideal noise rejection capability of an integrating ADC. Also, its high-resolution track-and-hold circuit allows very fast sampling with maximal time certainty. These features, combined with the display capabilities of a host computer, are all that is needed to implement a high-resolution single-channel oscilloscope or DFT analyzer.

Acknowledgments

I would especially like to mention the contributions of Dave Rustici who implemented the original version of the real-time oscilloscope program, Evan Whitney of Santa Clara Division who wrote the sinc interpolation software, and Ron Swerlein who wrote the original compiled DFT routine. Also, I'd like to thank Brad Waite, who provided programming support on a prototype of the waveform analysis library, Greg Wale, who helped with software testing, and Brian Berry, whose inputs helped greatly to refine the definition of the waveform analysis library.

References

1. D.R. Flach, "Characterization of Waveform Recorders," *Digital Methods in Waveform Metrology*, NBS Special Publication 707, October 1985, pp 31-52.
2. B. Allen Montijo, "Digital Filtering in a High-Speed Digitizing Oscilloscope," *Hewlett-Packard Journal*, Vol. 39, no. 3, June 1988.