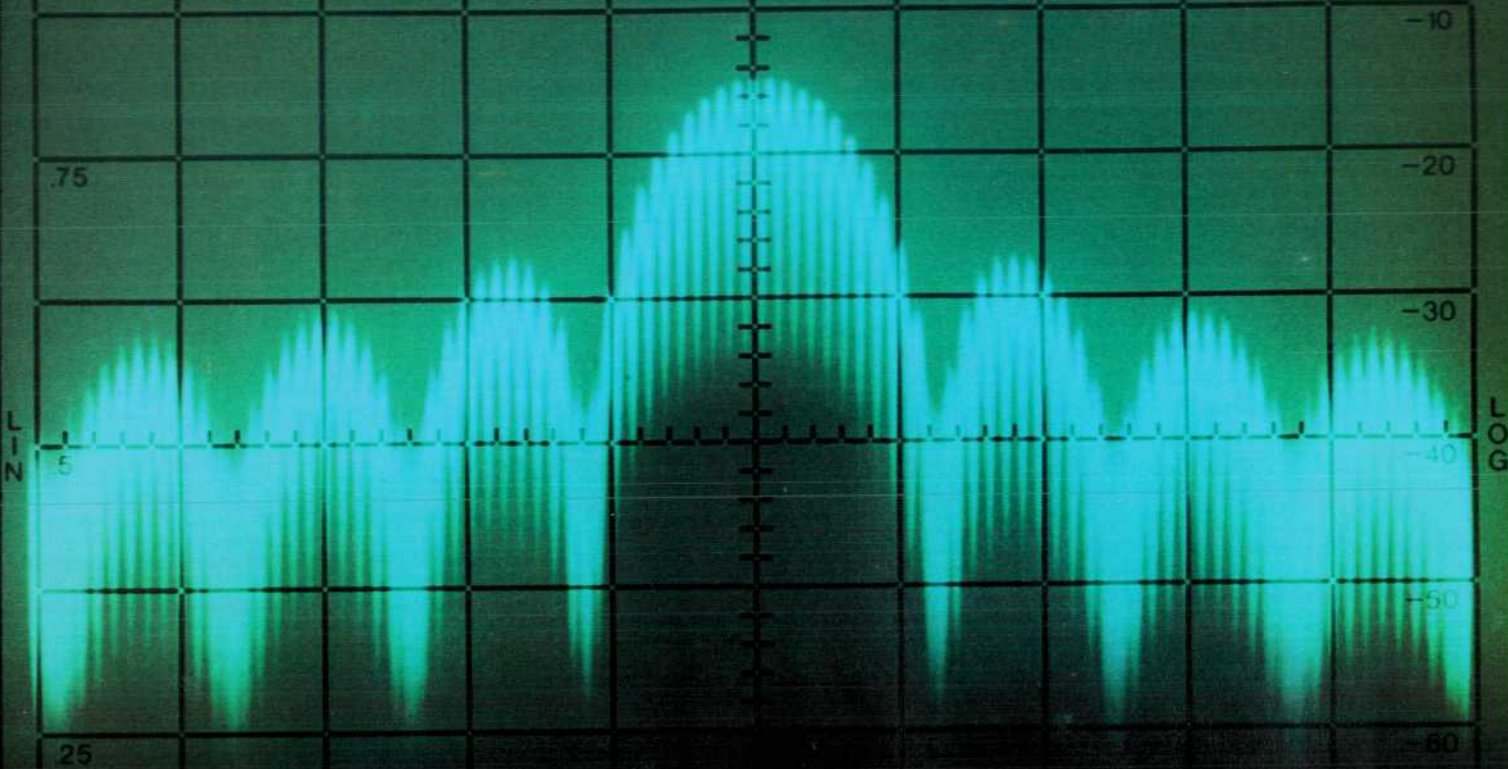


HEWLETT-PACKARD JOURNAL



Contents:

- 3 2-to-26.5-GHz Synthesized Signal Generator Has Internally Leveled Pulse Modulation**, by William W. Heinz and Paul A. Zander *Other features are microcomputer control, front-panel data entry, digital sweep, and internal diagnostic capability.*
- 7 Sample-and-Hold Leveling System**, by Ronald K. Larson *A logarithmic amplifier in the feedback loop reduces the effects of loop-gain variations.*
- 10 A Wideband YIG-Tuned Multiplier and Pulsed Signal Generation System**, by Ronald K. Larson and Lawrence A. Stark *This system enhances output power and frequency range and reduces pulse rise time for HP's latest synthesized signal generator.*
- 12 Autopeaking**, by Paul A. Zander *A small amount of hardware and some microprocessor code adjusts a YIG-tuned multiplier to the center of its passband.*
- 17 Compact Digital Cassette Drive for Low-Cost Mass Storage**, by William A. Buskirk, Charles W. Gilson, and David J. Shelley *This portable, battery-operated unit provides low-cost data and program storage for HP-IL systems.*
- 25 Scientific Pocket Calculator Extends Range of Built-In Functions**, by Eric A. Evett, Paul J. McClellan, and Joseph P. Tanzini *Complex number and matrix computations, integration, and equation solving are only a few of this programmable calculator's capabilities.*
- 36 A Pocket Calculator for Computer Science Professionals**, by Eric A. Evett *This new programmable calculator works with Boolean algebra, converts number bases, and manipulates bits. It also does arithmetic.*

In this Issue:



For the design, production, and maintenance of radar and communications systems, sources of pulsed high-frequency signals are always in demand. The more versatile the source, the more jobs it can do and the more cost-effective it is (they're never cheap). Desirable attributes include broad frequency coverage, a good mix of modulation capabilities for simulating real-world signals, accurate output power controllable over a wide range, low noise, programmability for use in automatic systems, and of course, reliability.

This month's cover subject, the HP 8673A Synthesized Signal Generator, is designed to satisfy the needs of some pretty demanding applications. In radar receiver testing, for example, it can simulate the radar pulses picked up by an antenna, even mimicking the variation in the strength of the pulses as the antenna rotates past a target. Its capabilities come partly from microprocessor control and partly from advanced microwave circuit design. The full design story is on pages 3 to 16. The cover shows the classic pulse-train spectrum of the pulsed microwave output from the 8673A.

A few issues ago, the subject of several of our articles was the Hewlett-Packard Interface Loop, or HP-IL, a method of interconnecting battery-powered computers, peripherals, and instruments to form portable systems that can go into the field to make measurements and record data. Currently, there's just one battery-powered HP-IL peripheral for recording data, and that's the HP 82161A Digital Cassette Drive. This book-sized device can store and retrieve 131,000 characters of data using a matchbox-sized magnetic tape cartridge. Because of its intended use, it has to be light and rugged and use little power. In the article on page 17, its designers tell us how they met these goals.

The articles on pages 25 and 36 describe a pair of unusually capable programmable pocket calculators. One of them, the HP-15C, is for engineers, scientists, and mathematicians. It has a truly remarkable repertoire of mathematical functions that one would normally expect to need a computer for. Now, if you need matrix arithmetic, complex-number operations, equation solving, or numerical integration, you can just reach in your pocket. It's a vivid demonstration, I think, of what integrated circuit and computer technologies are doing for us.

The other calculator, the HP-16C, is for computer programmers and designers of digital circuits. Besides the usual arithmetic functions, it performs many of the basic operations of digital circuits and computers. For example, it can shift, rotate, mask, set, clear, test, and count the bits in a digital word. It has the fundamental combinational logic functions (OR, AND, NOT, and XOR) and it manipulates numbers in four different bases—hexadecimal, decimal, octal, and binary. With this little marvel, the professional can simulate a computer algorithm or logic circuit to see if it works—another computer-sized job cut down to pocket size.

-R. P. Dolan

2-to-26.5-GHz Synthesized Signal Generator Has Internally Leveled Pulse Modulation

This second-generation instrument features microprocessor control, sophisticated sweep capabilities, programmability, and enhanced serviceability.

by William W. Heinz and Paul A. Zander

BROADBAND, synthesized microwave signal generators offer the stability, frequency accuracy, and spectral purity of a synthesizer together with the level accuracy and AM and FM capabilities of a signal generator. They have found numerous applications in communications and radar testing and simulation. Programmability has generated widespread use of these instruments in automatic test systems.

Since the introduction of the HP 8672A Synthesized Signal Generator in 1976,^{1,2,3} increasing user sophistication and demands for performance enhancements have led to the next generation in this instrument family. The new 8673A Synthesized Signal Generator (Fig. 1) covers the 2-to-26.5-GHz frequency range and features internally leveled pulse modulation capability. The addition of mi-

crocomputer control provides keyboard operation, data entry from the front panel, digital sweep capability, and many other features. Calibrated output levels from -100 dBm to +8 dBm are available over the 2-to-18-GHz frequency range. Maximum power is +4 dBm to 22 GHz and 0 dBm to 26 GHz.

Failure rates are of major concern to users of modern sophisticated instruments. The excellent reliability of the 8672A has not been compromised in the design of the 8673A. A number of service features and diagnostics have been incorporated to reduce repair time.

System Operation

The organization of the 8673A is similar to that of the 8672A. The digital control unit (DCU), containing a mi-



Fig. 1. The HP 8673A Synthesized Signal Generator features metered AM, low-distortion FM, and high-performance pulse modulation. Its frequency range is 2 to 26.5 GHz.




Fig. 2. Block diagram of the 8673A RF output section showing the major microwave components and the automatic level control (ALC) circuitry.

microprocessor, directs the internal oscillators, switches, attenuators and digital-to-analog converters (DACs) to appropriate operating points to produce the frequency, level, and modulations requested by the user either via the front panel or remotely via the HP-IB (IEEE 488). The frequency synthesis section is similar to that of the 8672A. This section includes a 2-to-6.625-GHz YIG-tuned * oscillator (YTO) which is phase-locked to signals derived from a stable 10-MHz reference oscillator. The RF output section contains a YIG-tuned multiplier (YTM) which either passes the amplified YTO signal or multiplies it by 2, 3, or 4 to cover the 2-to-26.5-GHz frequency range (see article, page 10). A block diagram of the RF output section is shown in Fig. 2. The signal from the YTO is amplified and passed through a pin diode modulator for amplitude modulation and level control (ALC). Pulse modulation is performed in the subsequent series/shunt diode modulator before power amplification and multiplication in the YTM.

A number of benefits accrue from this configuration, in which pulse modulation is performed ahead of the multiplier. Adding a modulator after the YTM would absorb valuable power, particularly at the higher frequencies. The design of such a component to 26 GHz would be difficult, especially if the desired pulse on/off ratio of 80 dB is to be maintained. Another advantage of modulation ahead of the multiplier is the virtual elimination of pulse video feed-through by the filtering action of the YIG filter in the YTM. A disadvantage of this system is the deterioration of pulse rise time through the YTM. The approach used in solving this problem is discussed in the article on page 10.

The multiplier is driven by a broadband GaAs FET power amplifier which produces a typical output power in excess of +27 dBm.⁴ After multiplication in the YTM, the signal passes through a broadband directional coupler that has a leveling detector at the coupled port. The detected dc voltage is fed through a logarithmic amplifier and summed into the ALC loop. In pulse mode, the ALC loop error signal is generated by sampling during the on time of the pulse and holding between pulses. To achieve accurate leveling of pulses down to 100-ns pulse widths, the sampling gate window must include only the flat top of the sampled pulse, yet be as broad as possible to maximize effective duty cycle.

*The frequency-determining element is a sphere of yttrium-iron-garnet in a magnetic field. Its resonant frequency is tuned by varying the magnetic field strength.

Thus the rise time of the detected pulse must not be degraded through the detector and logarithmic amplifier. This requires a low bypass capacitance (3 pF) for the detector and sufficient bandwidth for the logarithmic amplifier. The actual gain-bandwidth product exceeds 500 MHz.

After the leveling coupler, the signal passes through a step attenuator which provides a maximum of 90 dB of attenuation in 10-dB steps. Continuous control of power levels between steps is provided by the ALC loop reference voltage which is adjustable by means of the vernier knob on the front panel. In remote operation, a DAC provides the reference voltage in 0.1-dB steps.

Performance

The instrument performance specifications are published only after characterization of a sufficient number of instruments to provide meaningful statistical data. The specification allows for measurement uncertainties (usually calculated to achieve 95% confidence levels), temperature and environmental variations, and any potential drift that may occur over time. For these reasons, measured room-temperature performance may be better than a specification might indicate.

Maximum power output of a signal generator is of major concern to users interested in using the instrument as a local oscillator or where measurement system losses are high. Fig. 3 is a plot of maximum output power obtained from ten production 8673A instruments together with specified power as a function of frequency. The graph shows the margin between the specifications and the power available at room temperature.

The level accuracy of the instrument is important to those measuring receiver sensitivities or the transmission characteristics of amplifiers or other devices. Fig. 4 shows level accuracy plots for ten production instruments for several ranges down to -100 dBm. In view of the excellent results obtained, it was decided that microprocessor correction of output level was not needed, eliminating the disadvantages of requiring new ROMs when components after the YTM are replaced.

Fig. 5 shows typical pulse performance at 26 GHz. Specified maximum rise time is 35 ns and maximum overshoot is 20%. Leveling accuracy down to 100-ns pulse widths is ± 1 dB relative to CW. To achieve 80-dB on/off




Fig. 3. Maximum output power plots for ten production 8673As at 25°C.

ratio, careful shielding of components is required. Leakage or radiation from components ahead of the pulse modulator is kept low enough not to leak back and be amplified by the power amplifier's 25 dB of small-signal gain.

Programmability

The 8673A is fully programmable via the HP-IB from an external controller. Extensive design effort went into making the remote programming as user-friendly as possible. One HP-IB innovation is the master-slave sweep. This is

useful in testing of receivers and mixers where it is necessary to have two signal generators sweeping with a fixed offset between their output frequencies. One 8673A is designated as the master unit and sends out HP-IB commands to one or more slave 8673As. Master-slave sweep can be performed without a computer to control the system. The slave holds off the HP-IB handshake until its output has settled. The master looks for release of the handshake before proceeding to the next frequency. This ensures that the two synthesizers track each other.




Fig. 4. Level accuracy range for ten production 8673As at 0 dBm, -70 dBm, and -100 dBm output levels.




Fig. 5. A typical 8673A RF pulse at a carrier frequency of 26 GHz. The horizontal scale is 20 ns/div.

Of particular interest to system programmers trying to achieve maximum performance is the ready bit in the HP-IB status byte. The 8673A may take different times to settle after a frequency change. By sensing when the 8673A has actually settled instead of always waiting for a fixed worst-case delay time, the test system can run faster. The ready bit indicates that the 8673A has phase-locked and finished the YTM autopeak routine (see page 12) at a new frequency. The ready bit can be configured by a user program to generate a service request interrupt. Using this feature, the computer can be performing other useful work while the 8673A is busy.

Rear-panel output connectors provide sweep and blanking voltages for sweep displays on recorders, oscilloscopes, or network analyzers such as the HP 8755C or the HP 8410B/C. Stop-sweep and trigger outputs required for operation with network analyzers are available from a 14-pin connector on the rear panel of the 8673A.

Other interfaces provided at this connector permit useful functions without the need for an HP-IB controller. These include remote frequency incrementing and decrementing, a trigger-sweep input, an end-sweep output, blanking of the frequency display, and sequential storage register recall.

Digital Control Unit

One of the design objectives of the 8673A was to replace the combinational logic control section of the 8672A with a microprocessor. By using circuit boards and software designed for the 8662A 0.01-to-1280-MHz Signal Generator, it was possible to make a first breadboard controller capable of keyboard entry of microwave frequencies in a matter of weeks. With a base of proven circuits as a starting point, the development effort concentrated on user friendliness and enhanced performance of the analog circuitry and serviceability to try to match the expected applications.

One of the features of the 8673A is digital sweep. The digital control unit (DCU) can completely synthesize a sweep from a series of discrete frequencies. The sweep range can be entered as either start-stop or center

frequency- Δf (span sweep). The DCU has the necessary firmware to calculate either pair of values from the other. For example, if the user enters a start frequency of 10 GHz and a stop frequency of 5 GHz, the 8673A will calculate a center frequency of 7.5 GHz and a Δf of -5 GHz.

One area in which the 8673A departs from conventional analog sweepers is in the control of sweep rate. The 8673A is primarily a microwave synthesizer, so sweeps must be divided into discrete steps. The user can enter the number of steps or the step size. Obviously, the more steps to be generated, the slower a sweep will be. The other aspect of sweep rate is the time per step. In a synthesized signal generator, this time is the sum of the transition time between frequencies and the dwell time on each frequency. The 8673A can achieve a faster overall sweep rate by allowing the user to specify the dwell time. The DCU automatically checks the loops for phase lock after each step and controls the timing and Z-axis blanking accordingly.

For automatic sweep with dwell times shorter than 5 milliseconds, the DCU does not wait for a complete phase-lock of all four loops. A heuristic algorithm considers which phase-locked loops are required to change on a particular step and estimates the required delay to be "close enough" that the frequency error will not be significant on most swept displays. As a result, the overall sweep time can be reduced. This makes the 8673A useful for such applications as aligning a circuit while watching a real-time swept display. At the same time, the synthesized nature of the 8673A eliminates the effects of drift commonly associated with analog sweepers on narrow sweeps. If the application requires a complete phase lock at each step, the user can simply enter a dwell time of at least 5 ms and the 8673A will give a complete phase lock.

The microprocessor-based controller of the 8673A is used in several ways to improve performance over that of the earlier 8672A. The most striking of these is the autopeak function. The microprocessor adjusts the tuning of the YIG-tuned multiplier to track the output frequency for maximum power and best modulation performance. This is described in detail in the box on page 12.

A number of internal features add measurably to overall ease of use and performance. For example, the 8673A has a number of functions that need timing. Instead of programming the microprocessor to go into a delay loop counting down numbers until the necessary time has elapsed, which would restrict the 8673A to performing one function at a time, an LSI multiple-timer IC and the basics of real-time, multitasking programming are built in. This allows the 8673A to do several things simultaneously.

Another internal feature is the hardware divider circuit. Whenever the 8673A output frequency is above 6.6 GHz, the frequency of the YIG-tuned oscillator must be multiplied. From the perspective of the DCU, the desired frequency must be divided to calculate the YTO frequency. For output frequencies from 6.6 to 12.3 GHz, the YIG-tuned multiplier multiplies by 2, so the DCU needs to divide by 2. Division by 2 is fairly easy. The binary number is simply shifted one place to the right. For output frequencies above 18.6 GHz, the YTM multiplies by 4, so the DCU must divide by 4. Division by 4 can be done by two divisions by 2. However, from 12.3 to 18.6 GHz, the YTM multiplies by 3

Sample-and-Hold Leveling System

by Ronald K. Larson

Design goals for the 8673A's automatic level control (ALC) system were as follows:

- Broadband power leveling from 2.0-26.5 GHz in pulse, CW, AM, and FM modes
- Good AM performance (10 Hz to 200 kHz, 0 to 90% depth) with a stable leveling loop
- Temperature-stable output power from -10 to +10 dBm
- External leveling capability using a power meter or a diode detector
- Pulse duty cycle as low as 0.0001
- Minimum power transients when changing frequency or power
- Full HP-IB control.

The leveling loop is shown in Fig. 2 on page 4.

One of the most difficult problems to solve in a very broadband leveling system is caused by the variation in gain with frequency of the components in the microwave signal path. In a conventional or linear leveling loop the voltage-gain variations directly produce loop-gain variations. This can cause loop stability problems and degrade AM performance. Also, it is usually desirable for a microwave signal generator to control and meter output power in dBm rather than volts. With a linear loop, the system reference voltage must be a nonlinear function of power out in dBm, and therefore cannot be used directly to control a meter calibrated in these units.

The use of a logarithmic amplifier in the feedback path of the loop has a valuable effect—the microwave gain of each part of the signal path in the leveling loop does not affect loop gain. Instead, the gain factor for each microwave component operating linearly is 1 dB/dB. The YTM is a nonlinear voltage-gain device but exhibits a nearly constant gain in dB/dB, when properly biased, for each multiplying band and over the full range of power output. Typical gain factors for the YTM are 1.0, 1.4, 1.6, and 1.8 dB/dB on bands 1, 2, 3, and 4, respectively. A factory-set loop-gain adjustment in the leveling system for each multiplying band compensates for the YTM gain changes.

A detector operating in its square-law region will produce a large voltage-gain variation as power output varies. This would produce another source of loop-gain variation in a linear loop. In the logarithmic loop the gain factor is simply 2 dB/dB, greatly reducing loop-gain variations.

Thus the log amplifier reduces total loop-gain variations to a few dB over the entire range of power and frequency. Loop gain is also independent of variations in small-signal gain that can occur in the microwave amplifiers. The result is a loop with nearly constant AM bandwidth and excellent stability.

When the log amplifier is used, its voltage output varies linearly with power output in dB. Thus the reference voltage into the summing junction is linear with RF power out in dBm. This simplifies the control and metering of output power since the reference voltage need not be shaped and can be used directly to control the deflection of a meter calibrated in linear dBm units.

The exponentiator following the integrator lets the integrator voltage control the modulator output power linearly in dB. The

exponentiator gain factor is 0.9 decade of current per volt input. The pin diode modulator has the property that any decade of input current produces the same dB change in output power. The exponentiator/modulator cascade has a sensitivity of 36 dB/volt at the integrator output.

To produce amplitude modulation in the loop, the modulation voltage is summed into the summing junction after being logarithmically shaped. Shaping is necessary to produce RF envelope voltage variations that are linear with the AM input voltage. The resultant AM typically has less than 5% distortion at 90% depth and 100 kHz rates. YTM linearity in dB/dB is a definite factor in achieving this kind of AM performance. The YTM using self-bias rather than fixed bias can maintain bias stability, freedom from parametric phenomena, and a nearly constant gain factor over its full dynamic range. The actual dynamic range needed for power output from -10 to +10 dBm in CW mode with 90% AM depth is 40 dB. Microwave amplifier compression can add another 10 to 15 dB. Therefore, the total ALC loop dynamic range is 50 to 55 dB.

The voltage output of the detector varies with temperature. The temperature coefficient (TC) varies with power level. To correct for this varying TC, the logging amplifier has a thermistor in a resistor network to correct for detector drift at -4 dBm power output. This leaves a residual drift term at all other power levels which must be corrected. This term is proportional to power level in dB. It is corrected by using a linear-TC resistor in the reference voltage circuit. The result is typically less than 0.1 dB of drift over the specified temperature range of 15 to 35°C and over the full range of power and frequency.

Operation of the leveling loop in the pulse mode is identical to the CW mode since the sampling switch is closed only when RF is on. When the switch is open the integrator capacitor holds its charge, thus maintaining constant output voltage. The current into the pin modulator is constant when the RF is off. When RF again turns on, the switch closes and any charge that may have leaked off the capacitor is replaced. This design—where the loop actually samples the error voltage—eliminates the requirement seen in some peak-leveling systems to slew the hold capacitor voltage to the full value of the pulse on each RF pulse. The integrator capacitor also provides the system's dominant pole.

The ALC loop has two selectable bandwidths. The wide bandwidth allows high rates of amplitude modulation on a carrier and fast transient response. The narrow bandwidth is used for CW signals and reduces AM noise on the carrier, but has slow transient response. The 8673A's digital control unit automatically selects the wide bandwidth whenever the frequency is switched and then switches to the narrower bandwidth when appropriate. This allows the ALC circuit to recover more rapidly after a frequency change. Switching the ALC bandwidth as a function of frequency switching, modulation, and sweep mode would be impractical without a microprocessor.

and so the DCU must divide by 3. A general-purpose division routine for the microprocessor would take more than 20 milliseconds for this division. This is more than the specified worst-case frequency switching time! It would be unacceptable to require that much DCU processing time for certain output frequencies.

The design used for division by 3 in the 8673A includes a special circuit, Fig. 6, to speed up this process. The microprocessor starts the division cycle by clearing the latch of any possible remainder from a previous calculation. Then it gets the first digit of the frequency from memory and writes it to the latch. For example, for a frequency of 12,345,678 kHz, this is a 1. The latch stores the 1, and in turn drives the ROM with 01. On the next instruction, the microprocessor reads the ROM output through the data buffer. In our example, it will read a partial quotient of 0 and a remainder of 1. When the microprocessor writes the next digit (2 in the example) to the latch, the remainder from the previous digit is automatically put into the latch at the same time. A dividend of 2 and a remainder of 1 from the previous digit combine to make 12. This will divide by 3 to give 4 with 0 remainder. By freeing the microprocessor from manipulating the remainder between digits and from calculating the absolute table address for each digit, the division routine for an eight-digit number can be accomplished in less than 0.2 millisecond. This is 100 times faster than the general-purpose software routine, and almost as fast as division by 2.

Electromagnetic Compatibility

One of the classical problems in digital system design is that digital circuits tend to generate electromagnetic interference. Preventing the signals generated in the 8673A's digital control unit from coupling to other circuits required several measures. The first and primary measure was to design the hardware and firmware so that the microprocessor spends most of the time in a wait-for-interrupt state. When this happens, almost all of the logic lines in the DCU are quiet. As a second precaution, the buses carrying control signals to the front panel, phase-locked loops, and output section are all driven by latches. The latch outputs only change when necessary. This further confines the gen-

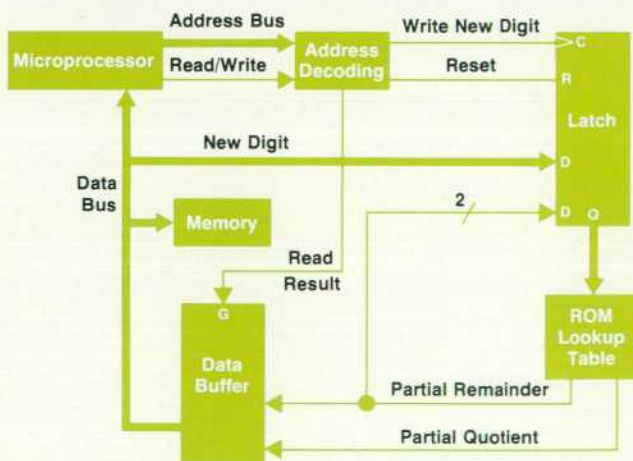


Fig. 6. This circuit reduces the time it takes to divide by 3 by a factor of five over a software method.

eration of noise to the DCU. These precautions plus good engineering practice with grounds and bypass capacitors reduce the coupling and radiation of digitally generated signals to a level well below the specified limits. The 8673A passed the radiated-electromagnetic-interference test the first time a prototype was put in the screen room. Just the same, extensive type testing was performed to verify that nothing had been overlooked.

Reliability

With the increasing complexity of today's instruments reliability becomes a major concern. Considerable effort was devoted to thermal design, component ratings, vendor history and stress analysis during the design of the 8673A.

Statistics acquired over the last five years since the introduction of the 8672A indicate a warranty failure rate well below the original goal. Actual warranty rate is about 32% per year, or an MTBF of 6700 hours assuming an operating time of 2000 hours per year for the instrument's 3100 components.

The 8673A design followed the successful approach of the 8672A. Thermal profiles of the instrument were done to measure local temperature rises and to check that average rise was less than 10°C above ambient. The increased power dissipation of the 8673A led to a more massive heat sink and supplemental individual heat sinks for the series-pass power supply transistors. A stress analysis computer program was used to evaluate each component in its operating environment to verify that internally generated derating guidelines were not exceeded. Careful failure analyses were performed on failed components to gain an understanding of the failure mechanisms and to provide vendors with appropriate information to rectify the problem. The failure rate analysis program, from which the 8672A failure rate has been accurately computed, predicts an 8673A MTBF greater than 5400 hours for 3320 components.

Serviceability

A number of features allow fast and easy troubleshooting if a failure does occur. First, every time the power is turned on, the DCU does a self-check of RAM and ROM. If a failure is detected, a code indicating the suspect IC is displayed on the front panel. During operation, if the DCU detects an abnormal condition, it displays a message on the front panel. For example, an output power unlevelled condition causes the **ALC UNLEVELLED** annunciator to be lit. This could be an indication of a malfunction, or simply that the user is trying to get more than the specified power at that particular frequency. Less likely failures are indicated by message numbers that are explained on the pull-out card. The DCU always attempts to continue operation. Despite a problem, the 8673A may still be useful for the particular measurement and service can be scheduled for a more convenient time.

A special function key accessible when the top cover is open (or via the rear-panel programming connectors) makes it possible to use the controller section to simplify servicing the other portions of the instrument. For example, in the 8672A it is necessary to connect a low-frequency function generator to the YTM drive circuit as part of the procedure to measure the passband and align the circuit. The 8673A can simply sweep the YTM fine-tuning DAC to generate the

ramp. The push of a couple of buttons eliminates the need for a piece of test equipment.

The HP 11726A Service Support Kit includes some special active extender boards. The primary purpose of these boards is to allow the service technician to tell quickly whether a problem is in the controller or elsewhere. For example, if the front-panel frequency entry and display appear to function but the output frequency is wrong, the problem could be in the controller or in a phase-locked loop. The numeric display on the extender board shows the DCU frequency output in decimal numbers. This makes it quick and easy to verify that the DCU is functioning properly.

If the digital circuitry is malfunctioning, a number of tests facilitate component-level troubleshooting. A key part of these tests is a 2K-byte ROM which is used only for troubleshooting purposes. Normally, its data outputs are not connected to the rest of the controller circuitry. The debug ROM contains a test program that tests not only the other ROMs but also itself. If the debug ROM does not pass the test, the next step is to put the microprocessor into a free-run mode and use signature analysis to verify the address decoders and the debug ROM data. If this doesn't find the problem, it is time to check the clock waveforms and the power supply voltages.

A dozen other routines are included in the debug ROM. Each routine is designed to test a specific part of the 8673A. Many of these tests also use signature analysis. For example, one routine exercises all of the digital logic in the output section. Each of the control latches is cycled through all valid data settings and each of the three DACs is stepped through a voltage ramp. If the DAC output waveform is not a simple staircase ramp, digital signatures can be taken at the DAC input with the same test setup.

Acknowledgments

The capable design team responsible for the 8673A had many contributors besides the authors in this issue. Credit goes to Ron Kmetovicz for the original concept and for his leadership, to Charles Cook for the product design, and to John Hasen for the design of the YTM driver, the DAC circuitry, and the test system. Dave Veteran was responsible for the improved step attenuator design. Jerry Ainsworth and Derrick Kikuchi worked on the pulse circuitry. Randy White was the reliability engineer. Don Pierce was responsible for the microcircuit production engineering and John Sims and Ray Yetman were the instrument production engineers. Eric McHenry contributed to the test system software. Donn Mulder was the product marketing engineer.

William W. Heinz



Bill Heinz has been with HP for seventeen years. He's been a diode applications engineer, a designer of microwave components and oscillators, and a project manager for synthesized signal generators, the latest being the 8673A. A senior member of the IEEE, he has coauthored several published letters on ferrite devices and a paper on a microwave amplifier. His work has resulted in a patent on a microstrip resonator. Bill was born in Havana, Cuba and attended the Polytechnic Institute of Brooklyn, receiving his BSEE degree in 1960 and his MSEE in 1962. Before

coming to HP he worked with masers, paramps, cryogenic ferrite components, and superconducting magnets for five years. He lives in Palo Alto, California, has two sons, and likes to travel abroad, hike, listen to music, and attend the theater.

Paul A. Zander



Paul Zander participated in a co-op study program with General Motors Institute in Flint, Michigan and received the MS degree from Purdue University before joining HP in 1970. He has been responsible for a number of analog and digital design tasks, most recently the hardware and software for the digital control unit in the 8673A Synthesized Signal Generator. Two patent applications have resulted from his work on the 8673A. A native of Cleveland, Ohio, and a member of the IEEE, Paul is married, has two children, and lives in Los Altos, California. He is building a home computer, enjoys photography and running, holds the amateur radio call sign AA6PZ, and has had several magazine articles published about his station.

References

1. J.L. Thomason, "Expanding Synthesized Signal Generation to the Microwave Range," *Hewlett-Packard Journal*, Vol. 29, no. 3, November 1977.
2. K.L. Astrof, "Frequency Synthesis in a Microwave Signal Generator," *Hewlett-Packard Journal*, Vol. 29, no. 3, November 1977.
3. B.C. Stribling, "Signal Generator Features for a Microwave Synthesizer," *Hewlett-Packard Journal*, Vol. 29, no. 3, November 1977.
4. M. Furukawa, "A Broadband 2-to-7-GHz Power Amplifier," *Hewlett-Packard Journal*, Vol. 33, no. 2, February 1982, p. 20.

A Wideband YIG-Tuned Multiplier and Pulsed Signal Generation System

by Ronald K. Larson and Lawrence A. Stark

THE KEY TO THE 2.0-to-26.5-GHz frequency range available at the output of the 8673A Synthesized Signal Generator (see article on page 3) is a broadband YIG-tuned multiplier (YTM), which is shown schematically in Fig. 1. YIG stands for yttrium-iron-garnet, a ferrite material. When a YIG sphere is placed in a magnetic field, it exhibits a sharp resonance at a frequency that is a function of the magnetic field strength.

The operating frequency range of the YTM is divided into four bands, which correspond to frequency multiplication ratios of 1, 2, 3, and 4. In band 1, step recovery diode D2 in Fig. 1 is forward-biased to a low impedance and no significant harmonic generation occurs. The four bands and the corresponding frequency ranges are listed below.

Band Number	Output Frequency Range (GHz)	Input Frequency Range (GHz)
1	2.0 - 6.6	2.0 - 6.6
2	6.6 - 12.3	3.3 - 6.15
3	12.3 - 18.6	4.1 - 6.2
4	18.6 - 26.5	4.65 - 6.625

The YTM consists of a standard step recovery diode multiplier which generates a comb of harmonics of the input frequency. The input frequency is tunable over a broad range and the multiplication ratio is varied by tuning a YIG filter to select a single harmonic component. The multiplier is inherently broadband in that the comb spectrum generated by the diode extends from the input frequency to an upper limit greater than 30 GHz. By tuning the YIG filter at

the output of the multiplier to one particular harmonic, all unwanted output signals are suppressed and the desired frequency is delivered to the output of the device. The input low-pass filter prevents the output signals from returning to the input and prevents harmonics of the power amplifier feeding the YTM from interfering with the multiplied signal.

In the multiplying bands, the step recovery diode is biased to act as a charge controlled switch which produces a narrow voltage impulse when the diode switches from forward to reverse bias. The impulse width is determined by the circuit inductance and the diode capacitance, assuming that the transition time of the diode is negligible. Impulse widths of 40 picoseconds are necessary to obtain high conversion efficiency at 26 GHz, and since the diode transition time should be a small fraction of the pulse width, it is very important to obtain diodes with low capacitance and short transition times.

The proper timing of the switching action is controlled by the dc self-bias voltage. The ideal timing point for switching to low capacitance occurs when the diode current has reached its maximum value. The resulting diode voltage impulse goes negative initially and then positive, at which point the diode switches back to the conducting state.

The input loop inductance of the YIG filter and the diode capacitance form the resonant circuit for the impulse generator. The input low-pass filter provides an impedance match at the input frequency and a short circuit to harmonics.

Mechanical Design

The complete multiplier is constructed on a single sapphire substrate. The design goal is to provide the closest possible physical proximity between the step recovery diode and the YIG output filter. This increases the broadband capabilities of the circuit by reducing the path length of the unwanted harmonics that are reflected from the YIG filter back to the diode.

Fig. 2 shows a photograph of the region around the YIG filter and step recovery diode. The diode is die-attached to the top surface of the chip capacitor which is the final element of the input low-pass filter. The capacitor is epoxied to the heat sink, which provides an extension of the ground plane in the region of the rectangular hole. This opening is laser-cut, while the circular hole that holds the 0.023-inch-diameter YIG sphere is drilled ultrasonically.

The assembled substrate is held in a magnetic package consisting of a center body and two magnets. The design of this package provides differential expansion of the magnets and the center body so that the gap between the magnet pole faces remains constant when the temperature of the struc-

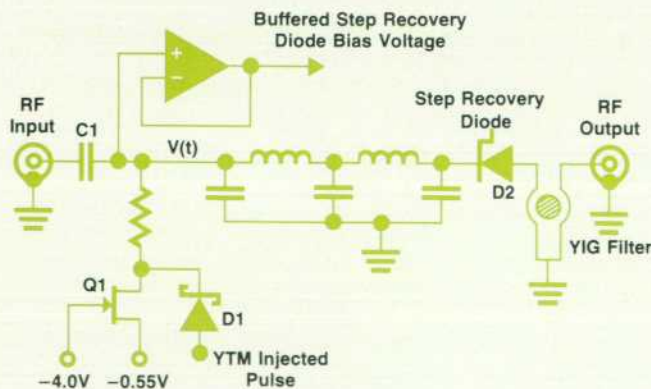


Fig. 1. YIG-tuned multiplier schematic. At output power levels greater than approximately 0 dBm, the steady-state value of the step recovery diode bias voltage $V(t)$ is directly proportional to the RF input voltage. The FET resistance is controlled to give the highest conversion efficiency consistent with stable parametric-free RF output.




Fig. 2. (a) Photograph of the YTM circuit showing the components and filter. (b) Photograph of the microcircuit showing the ground plane, YIG sphere, and output coupling loop.

ture changes. This keeps the YIG filter tuned precisely to the output frequency.

Input Low-Pass Filter

Optimizing the input low-pass filter for a resonance-free stopband up to 26.5 GHz was a major effort. The transverse resonances in the distributed transmission line filter elements forced a modification of the basic design, which was derived from standard filter element values. It was evident that the final capacitive element in the low-pass filter could not be realized on the 0.01-inch-thick sapphire substrate. A thinner dielectric was needed to push the transverse resonances past the upper frequency limit of 26.5 GHz. A 0.004-inch-thick single-dielectric capacitor with a nominal value of 2.0 pF was chosen empirically on the basis of performance in actual circuit tests. This was significantly lower than the prototype value of 2.7 pF.

Step Recovery Diode Characteristics

Two different step recovery diodes (SRD) have been used successfully in these YTM's. Both were developed for in-house use at Hewlett-Packard. One diode is an outgrowth of HP's standard SRD product line and the other was developed by HP's Santa Rosa Technology Center as a second source. A major effort in the development of this device was to keep the doping profiles as abrupt as possible. Coupled

with intrinsic-layer thicknesses less than 1 μm , this was seen as the key to obtaining the shortest possible transition time. Two other diode parameters of importance are the reverse bias capacitance and the recombination time. As the diodes are made physically smaller, the recombination time begins to drop because of the influence of the sides and contacts of the device. Many different diodes were evaluated to select the optimum junction capacitance for maximizing the conversion efficiency at 26 GHz. The final diode dimensions are a compromise among aspect ratio, capacitance, breakdown voltage, and transition time.

Bias Control

Optimum RF multiplication requires that the appropriate dc conditions be established for the diode. This is done through a separate bias circuit. A blocking capacitor prevents the dc from flowing in the microwave input circuit. Fig. 3 shows how the SRD's dc operating point is dependent on the microwave signal amplitude. The circuit that supplies these dc conditions is called a self-bias circuit because the dc operating point is established by the microwave signal input amplitude. The advantage of the self-bias circuit is that a stable operating point is obtained easily over the full dynamic range of the multiplier because the dc conditions follow the RF power level smoothly. The correct operating point is determined by the bias resistance. If the bias resistance is set lower than the optimum value, the multiplication will be stable, but the power output will be low. As the bias resistance is raised, an optimum point is found where the power is high and the operation is still stable. At higher resistance, either the output power drops, or the operation becomes unstable.

It has been found empirically that the optimum self-bias resistance is a function of frequency within a single multiplying band. To provide for this variation, a single JFET chip is used as a voltage-controlled resistor on the multiplier substrate. The resistance is controlled by adjusting the gate voltage.

Another way of establishing the proper dc bias is with a

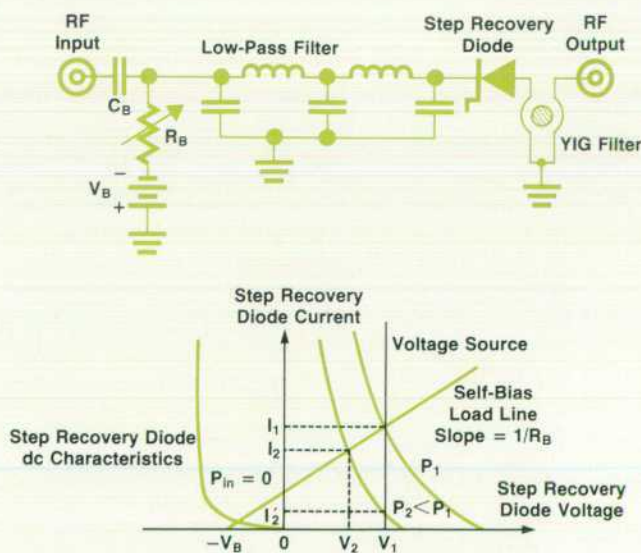


Fig. 3. Self-bias circuit operation.

Autopeaking

by Paul A. Zander

One aspect of the YTM's performance is that it acts as a very narrow bandpass filter that tracks the source output frequency. If the tuning is off by as little as 10 MHz, maximum available power is reduced, the pulse modulation waveform is distorted, and the frequency modulation sidebands are filtered asymmetrically.

Avoiding these problems requires tuning accuracy better than 0.1%. Because of nonlinearities, hysteresis, and temperature sensitivities in the YTM, this kind of tuning accuracy is impractical to achieve in a straightforward open-loop tuning system. Autopeak is a combination of a small amount of hardware and 1500 bytes of microprocessor code that adjust the YTM tuning to the exact center of the passband (see Fig. 4 on page 13).

Because the peaking process introduces some perturbations on the output which could affect certain measurements, peaking is never allowed to occur spontaneously. Peaking can be prevented by a front-panel pushbutton or an HP-IB command.

Peaking is performed in conjunction with a number of operator inputs. It occurs whenever peaking is switched on, when the RF output is switched on or the FM deviation range is changed, and after every frequency change that results in an output frequency more than 50 MHz away from the last one where peaking was performed. Because peaking is critical to achieving good pulse performance, peaking is automatically enabled and performed each time the pulse modulation function is turned on. The time required for peaking depends on the YTM passband shape and the amount of correction required. To maximize the measurement rate in automatic systems, a bit in the HP-IB status byte is set to indicate that peaking has been completed. This can be used by the HP-IB system as an interrupt to initiate the next part of an measurement cycle, thus preventing the taking of false data during the peaking routine.

The Peaking Algorithm

The peaking process consists of four phases: setup, tuning the YTM, measuring the YTM self-bias (for pulse injection), and restoring normal operation of the 8673A. During the setup phase, the digital control unit (DCU) suspends normal operation of the output section. All modulation is switched off. The internal diode ALC detector is selected. If a very large frequency step has just occurred, some delay is used to allow the YTM frequency to settle completely before using the peaking as a fine tuning adjustment. Then the ALC circuit is put into hold mode using circuits already included for the sample-and-hold operation during pulse modulation. Control flags are set in software so that the 8673A can accept front-panel and remote commands to turn on modulation, but not execute the commands until peaking has been completed.

In the peaking process, the DCU uses an eight-bit digital-to-analog converter to tune the YTM ± 200 MHz in steps of 1.6 MHz. The peak sensing circuit, Fig. 1, is used by the DCU to direct the search for the center of the passband. Buffer amplifier A1 amplifies the output of the detector logarithmic amplifier in the ALC circuit. Because of the ALC log amp, a change in the YTM output of 1 dB causes a change of 300 mV at the output of A1 regardless of the absolute power. A2, C1, and S1 act as a sample-and-hold circuit which is controlled by the DCU. A3 compares the output of the sample-and-hold circuit, which represents a previous YTM tuning setting, with the present output of A1 and sends the result to the DCU. The entire circuit consists of a quad FET operational amplifier IC and a few discrete parts.

An offset corresponding to 3 dB can be switched in before A3. Without the offset, the output of A3 will change when the present YTM output is less than when the sample was taken. With the offset, the comparator output will not change until the YTM output is 3 dB less than the sample. The timing of the sample-and-hold operation, and whether the offset is on or off, depends upon which of two algorithms the DCU is using.

In the coarse tuning algorithm, the offset is not used. The DCU closes S1 long enough to charge C1 and then opens S1. Thus a voltage proportional to the RF power is held on C1. Next the YTM is tuned 10 MHz. After a delay of 200 microseconds to allow the tuning coil time to respond, the DCU checks the output of A3. If the microwave output power is greater at the new tuning setting, the output of A3 will be high, otherwise it will be low. If the power is higher, the DCU repeats this process of sampling the output power on C1, tuning the YTM to a new setting and checking to see if the YTM output is greater or less. With suitable refinements to ignore minor resonances in the YTM passband shape, this algorithm does a good job of locating the main peak of the passband. However, it does not do a repeatable job of finding the exact peak. Instead, it tends to stop at a point a fraction of one dB past the true center of the YTM passband. For achieving maximum power output, it is adequate and represents a major improvement over the older HP 8672A, which has no automatic peaking. However, the pulse and frequency modulated signals show some distortion with even a small amount of mistuning, so a second algorithm is built into the DCU.

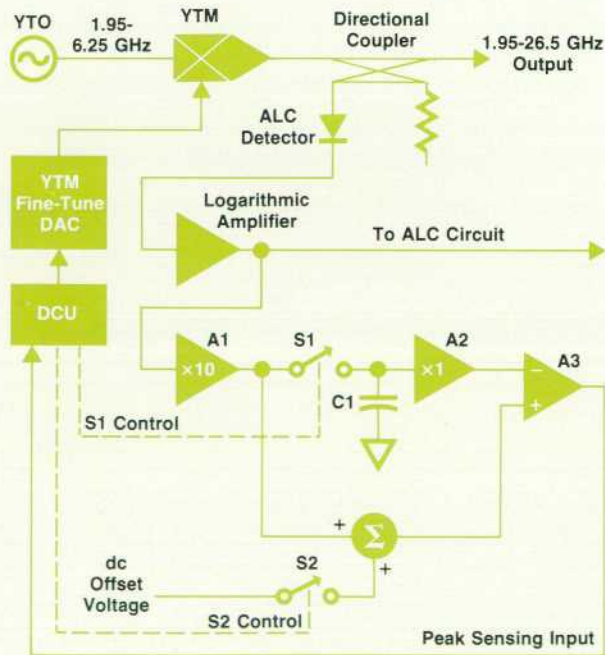


Fig. 1. Block diagram of the YTM peak sensing circuit.

The basic problem with the coarse tuning algorithm is that the power output changes very slowly as the YTM tunes through the middle of the passband. Finding the exact middle becomes an example of the classical measurement problem of accurately finding the small difference between two large values. The solution was to develop a centering algorithm.

The centering algorithm is based on the fact that the YTM passband shape is reasonably symmetrical in the region from 3 to 15 dB below the peak. The DCU performs the centering algorithm by first assuming that the initial tuning point is close to the peak. It samples and holds a voltage on C1 which corresponds to the initial power output. Then it switches in the offset voltage corresponding to 3 dB and starts stepping the YTM until the comparator output changes. This point corresponds to an output power level 3 dB less than the original point. Next the DCU returns to the original point and starts stepping the YTM tuning in the other direction until the other 3-dB point has been located. The center of the passband is calculated to be halfway between the 3-dB points. Because the YTM passband is symmetrical, it makes no difference whether the original point was at the exact center or slightly below the center.

Normally, the open-loop tuning is accurate enough that the centering algorithm alone is adequate. For small frequency steps, the time to search is further reduced by starting at the tuning correction for the previous frequency. Only when the centering algorithm encounters a difficulty, such as when the YTM tuning has been grossly pulled by a reactive load impedance, is the

coarse tuning algorithm used.

Bias Sample

Once the YTM frequency has been tuned, the self-bias must be measured for use with pulse injection. This is only done in pulse mode. The DAC output that controls the pulse injection is compared with the actual YTM bias voltage for CW output at that particular microwave frequency. The DCU performs a straightforward successive approximation algorithm to measure the bias voltage. For a fixed frequency and output power level, that would be the end of it. However, the injection must change as the output power level is changed. If only the power is changed, the DCU skips the YTM tuning procedures and simply measures the bias for the new power level. The DCU eventually builds up a table in memory of bias versus power level every 0.4 dB at that frequency. Once the needed data is in the table, the DCU simply looks it up, thereby avoiding the possible user inconvenience of having the DCU continuously switching to CW so that the bias injection can be measured. Of course, when the YTM tuning is changed (for example, when the frequency is changed) the table of bias values is erased, and the data must be remeasured.

After the YTM tuning has been corrected and the bias measured, the DCU restores normal modulation and operation to the output section. The time it takes to perform this process varies somewhat with the amount of tuning needed, but is typically 5 to 10 milliseconds. In return for this brief delay, the 8673A delivers more output power and better modulation.

voltage source. This works well at constant power, but requires external adjustment if the power level is varied. As shown in Fig. 3, as the power level drops, the operating voltage should drop too, so a voltage-source bias network needs to be adjusted as the input power to the multiplier is changed.

At low power levels there is not enough diode current to produce the required bias. The fixed bias voltage shown in Fig. 3 provides the needed bias on the knee of the SRD's dc characteristic.

YIG Filter

The YIG filter is central to the operation of the overall YTM. The operating characteristics that had to be dealt with were bandwidth, crossing modes, magnet alignment, out-of-band rejection, acoustic squegging, and loop inductance.

The shape of the coupling wires affects nearly all of these

factors. The design uses tight coupling to maximize bandwidth and minimize loop inductance for narrow pulses. This leads to more severe crossing modes which have been minimized experimentally by optimizing the sphere orientation. As Fig. 4 shows, the YIG filter passbands are free of spurious modes and the 3-dB bandwidths are typically 60 MHz or more. A 3-dB bandwidth less than 40 MHz adversely affects the rise time of pulsed RF signals.

The parameters of the YIG sphere are:

- Saturation magnetization $4\pi M_S = 800$ gauss
 - Diameter = 0.58 mm (23 mils)
 - Operating temperature = 80°C
 - Bias field axis = 30° off $\langle 100 \rangle$ towards $\langle 110 \rangle$
- The YIG filter is temperature stabilized with a heater and thermistor feedback control which maintains the sphere at




Fig. 4. YIG filter passband characteristics. Horizontal scale: 40 MHz/div. Vertical scale: 3 dB/div.

80°C, independent of external temperature changes.

Nearly all of the key specifications of the 8673A depend on whether the YIG filter is centered at the output frequency. An autotune circuit (see page 12) controlled by the microprocessor ensures that the YIG filter is operating at the center of its passband. It works by making small corrections to the magnet current and monitoring the output of the ALC detector to find the maximum output power as a function of magnet tuning. Thus, output power is maximized and optimum pulse shape is maintained in the face of magnet tuning hysteresis and nonlinear tuning characteristics caused by reactive circuit pulling of the filter center frequency.

Pulse Modulation System

For the YTM to frequency-multiply microwave input pulses without producing excessively long rise times, the diode bias voltage $V(t)$ must reach its steady-state value in a time at least as short as the rise time of the input microwave pulse. Using the self-bias scheme described earlier, the dc bias conditions are created by the rectification inherent in the operation of the diode. Because of charge storage in the step recovery diode, the rise time of the bias voltage $V(t)$ is about 100 to 300 ns, producing microwave pulse rise times of approximately the same length.

The rise time of the output microwave pulse is longest at the high end of each multiplying band (e.g., at 12.3, 18.6, and 26.5 GHz). This occurs because reverse-recovery current of the SRD flows for a higher fraction of each input frequency cycle at the high end of each band. Hence it takes longer for the SRD bias voltage $V(t)$ to reach its steady-state value and the output pulse rise time is correspondingly longer.

The method used to eliminate the rise time degradation of the YTM is to charge the capacitor $C1$ in Fig. 1 to the required final value before the arrival of the microwave pulse. In this way, $V(t)$ is at the correct voltage when the pulse arrives, and the RF pulse passes through the YTM undistorted. $C1$ is charged by applying a short positive pulse to diode $D1$ before the RF pulse arrives. This pulse is called the YTM injected pulse. Using this technique, the microwave pulse rise time is limited only by the YIG filter bandwidth.

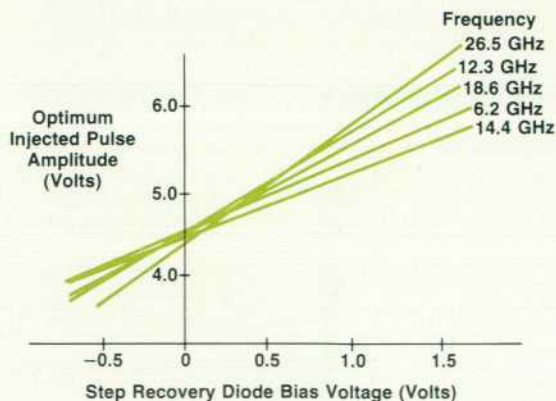


Fig. 5. The optimum injected pulse amplitude to eliminate rise time degradation is a linear function of the steady-state step recovery diode bias voltage.

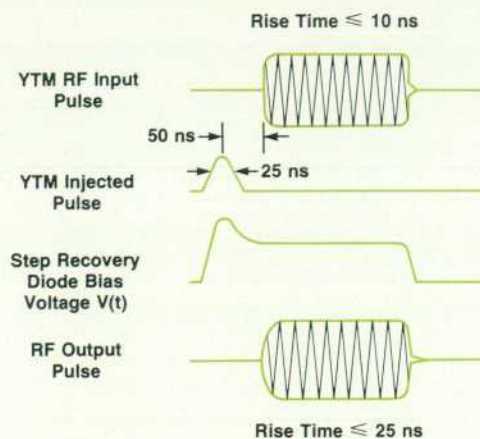


Fig. 6. Pulse injection forces the step recovery diode bias voltage to the correct value before the RF pulse arrives. Output rise time is then a function only of the YIG filter bandwidth.

The successful implementation of this system rests on the ability to predict the correct amplitude of the YTM injected pulse at each frequency and power level. Since the purpose of the injected pulse is to raise the bias voltage $V(t)$ to the value it would have in normal operation, the first step is to measure the steady-state value of the diode voltage while the YTM is in CW operation. Then this voltage is used to control the YTM injected pulse amplitude.

One requirement for this system to work effectively is that the optimum injected pulse amplitude must be a linear function of steady-state SRD bias voltage. The proportionality factor and offset must be relatively constant over any multiplying band. A simple model of the circuit predicts and experimental data confirms that the required linear relationship exists and that the variation in slope and offset with frequency can be accommodated by adjusting the gain and offset of the circuit that controls the injected pulse amplitude. This adjustment is made by observing pulse shape on each multiplying band while adjusting the injected pulse amplitude.

Fig. 5 shows a graph of the optimum injection pulse amplitude as a function of the steady-state value of step recovery diode bias voltage $V(t)$. Different values of bias voltage correspond to different microwave power levels at the input of the YTM. It can be seen that different slope and offset values are obtained depending on the operating frequency, although the total variation is not large. Within one multiplying band the variations are small enough that the behavior over the entire band can be adequately approximated by a single straight line.

Since the straight line is only an approximation, the actual RF pulse at a particular frequency can have either overshoot or lengthened rise time depending on whether the injected pulse amplitude is larger or smaller than the optimum value. The slope and offset parameters of the pulse control circuit that generated the curves in Fig. 5 are factory-set so that the total variation of rise time and overshoot is within the specified limits of 20% maximum overshoot and 35 ns maximum rise time. Typical performance is 10% maximum overshoot and 25 ns maximum rise time.

The pulse control system waveforms are shown in Fig. 6.




Fig. 7. An 18.6-GHz pulse (a) without and (b) with pulse injection. Horizontal scale: 50 ns/div. A 26.5-GHz pulse (c) without and (d) with pulse injection. Horizontal scale: 100 ns/div.

The timing of the application of the YTM injected pulse is critical and is fixed. The pulse arrives approximately 50 ns before the RF input pulse to allow ringing transients to die out of the SRD bias voltage $V(t)$. When these transients die out, the SRD voltage is at the correct value and the output microwave pulse rise time is limited only by the bandwidth of the YIG filter.

Fig. 7 shows pictures of RF pulse shapes that demonstrate the dramatic improvement in rise time as a result of the pulse injection.

The overall pulse modulation control system, Fig. 8, operates as follows. Whenever frequency changes by 50 MHz or more, or power changes by 0.4 dB or more, the input RF

voltage to the YTM can change significantly. Thus the SRD bias voltage $V(t)$ may change, which will change the required YTM injected pulse amplitude. To compensate for these changes the microprocessor switches the leveling system into the CW mode for about 200 μ s. During this time the microprocessor changes the DAC (digital-to-analog converter) output until it equals the steady-state value of bias voltage $V(t)$. Pulse mode is then enabled and the injected pulse amplitude is again the correct value to produce short-rise-time pulses.

The DAC output voltage drives an amplifier, which provides a gain and an offset adjustment for each band. The amplifier output voltage controls the injected pulse




Fig. 8. 8673A pulse modulation system.

amplitude.

Another requirement in the reproduction of high-quality output pulses is that the RF pulse into the YTM must be free of ringing and overshoot. This is accomplished by using a series-shunt pulse modulator before the power amplifier.¹ This minimizes mismatch reflections between the preamp output and the power amplifier input. The series pulse driver provides a pulse to the series diode while the shunt pulse driver provides a pulse to the shunt pin diodes. This method achieves the low reflections of a series-shunt modulator and retains the short rise time of a shunt pin diode modulator.

Also included in the pulse modulation system is a pulse width detector which turns on the front-panel unlevelled indicator when the input pulse width is less than 100 ns. The specified level accuracy at 100-ns pulse width is ± 1 dB relative to the CW level. Typically, specified level accuracy is maintained down to 80-ns pulse width. Pulse widths less than 80 ns are available but level accuracy is degraded. Maximum pulse repetition frequency for the specified level accuracy is 1 MHz. Typically, specified level accuracy is maintained to 5 MHz, as shown in the table below.

Pulse Performance Summary

Parameter	Specified Performance	Typical Performance
Level Accuracy (relative to CW)	± 1.0 dB	+0.4 dB ²
Minimum Pulse Width (for specified accuracy)	100 ns	80 ns
Minimum Duty Cycle (for specified accuracy)	0.0001	0.00003
Maximum PRF (for specified accuracy)	1 MHz	5 MHz
Pulse On/Off Ratio	>80 dB	>90 dB
Overshoot	20% (25%, 6.6-6.7 GHz)	10%
Rise Time	35 ns	20 ns ¹

1. Typical system performance gives rise time less than 25 ns on the multiplying bands and 15 ns on the 1.95-6.6 GHz band.

2. Typical level accuracy relative to CW is +0.4 dB at 100-ns pulse width.

Acknowledgments

We are grateful for the contributions of Don Pierce, Ray Yetman, John Sims, Bob Skinner, Sandy Smith, Pauline Prather, and Klaus Model. Thanks to Steve Sparks for his original development of the ALC loop.

Reference

1. R. Larson, "A Fast 2-to-8-GHz Pulse Modulator," Hewlett-Packard Journal, Vol. 29, no. 3, November 1977, p. 6.

Lawrence A. Stark



Larry Stark received his BSc degree in electrical engineering from the Massachusetts Institute of Technology in 1965 and his MSc and PhD degrees from Cornell University, the latter in 1971. He came to HP in 1979 after two years as an associate professor at Carleton University, Ottawa and six years in microwave components R&D. At HP's Stanford Park Division, he was a development engineer for the 8673A Synthesized Signal Generator and is now a product marketing engineer. A member of the IEEE, he has coauthored seven papers on microwave devices and systems. Larry was born in Newton, Kansas and now lives in Palo Alto, California. He's married, has a daughter, and enjoys tennis, skiing, playing piano, and photography. He also enjoys planning landscaping projects which he'll get around to "one of these days."

Ronald K. Larson



With HP since 1972, Ron Larson has designed the 11720A Pulse Modulator and the output sections of several microwave synthesizers, including the 8673A. His pulse injection system for the 8673A is the subject of a pending patent. A native of Minneapolis, he served ten years as an electronic technician in the U.S. Navy, attaining the rank of Chief Petty Officer, then worked as an electronic technician for two more years before returning to school. He received his BSEE degree from the University of Minnesota in 1972. Ron is a guitar player, a karate student, a small boat sailor, and a backpacker. He's married, has a son, and lives in San Jose, California.

Compact Digital Cassette Drive for Low-Cost Mass Storage

This portable battery-operated unit uses minicassettes to store programs and data inexpensively for HP-IL systems.

by William A. Buskirk, Charles W. Gilson, and David J. Shelley

THE HP 82161A Digital Cassette Drive (Fig. 1) is a portable, programmable, mass storage peripheral for the Hewlett-Packard Interface Loop (HP-IL).¹ The storage medium is a removable minicassette that can store up to 128K bytes of information. Portability is achieved by the use of a four-cell nickel-cadmium battery pack, recharger, and power supply system similar to that used in other portable HP products. The 82161A is styled to fit in a family of compact peripheral devices such as the 82143A and 82162A Printer/Plotters, and to fit nicely in a system controlled by an HP-41 Handheld Computer or an HP-75 Portable Computer. The 82161A makes use of much of the package design of the 82143A Printer/Plotter,² producing a unit 178 mm wide by 133 mm deep by 57 mm high. Replacing the 82143A's printer mechanism on the top right side is a transport mechanism with a **REWIND** key and a door **OPEN** key located in front. To the left of these keys is the power switch and indicators **POWER**, **LOW BATTERY**, and **BUSY**. The top left side of the package offers a compartment to store two minicassettes. The two HP-IL cables and the re-

charger cable are connected to the 82161A via plug receptacles on its back panel.

Electronic System

Fig. 2 is a block diagram of the electronic system of the 82161A. An internal microcomputer controls the head and motor drive electronics for the transport assembly and interacts with the HP-IL interface logic and data buffers.

The criteria for microcomputer selection for the 82161A included low cost, ready availability, low power consumption, and adequate I/O. To limit the number of electrical parts in the 82161A, a microcomputer that also contained ROM, RAM, and a timer, and could generate the encoded bit timing during a write operation was needed. A 3870 microcomputer with 2K bytes of ROM and 64 bytes of RAM was selected.

The logical interface of the 82161A is a generalized mass storage driver that provides the capability to execute operations such as initializing the tape, seeking a record, reading or writing a record, and rewinding the tape (see Table I).



Fig. 1. The HP 82161A Digital Cassette Drive is a compact battery-operated mass-storage unit designed for use in portable HP-IL systems.

Table I
HP 82161A Digital Cassette Drive Commands

DDL0	Write buffer 0	DDT0	Read buffer 0
DDL1	Write buffer 1	DDT1	Read buffer 2
DDL2	Write	DDT2	Read
DDL3	Set byte pointer	DDT3	Read address
DDL4	Seek	DDT4	Exchange buffers
DDL5	Format	DDT5	Transfer buffer 0→1
DDL6	Partial write		
DDL7	Rewind		
DDL8	Close record		
DDL9	Transfer buffer 0→1		
DDL10	Exchange buffers		

Buffer space for two 256-byte records of data is provided. Buffer 0 is used for data transfers between the HP-IL and the minicassette tape, and buffer 1 can be used by the HP-IL controller as virtual memory. The intent is to provide space to store a page of the tape directory and thereby reduce the number of seeks to the directory at the beginning of the tape. The DDL3 (set byte pointer), DDL8 (close record), and DDL6 (partial write) commands allow a memory-limited controller such as the HP-41 Handheld Computer to modify parts of a record without having to buffer the entire record in its mainframe. The record is read into buffer 0, modified, and written back to the tape with only the modification information passing around the HP-IL.

The ability to use the 82161A for extended remote data gathering tasks has been enhanced by the addition of the power-up/down commands. When the power switch on the front-panel keyboard is in the **STANDBY** position and a loop-power-down (PWRDN) command is received, the drive's power supply is turned off. When the HP-IL controller requires the loop to be active again, it sends a string of identify message frames, which turns the drive's power supply back on.

Software

The 2K bytes of machine code in the ROM of the microcomputer can be divided into three major areas: the

power-on idle routine, the HP-IL routine, and the device control routines. The power-on idle routine, which uses approximately 160 bytes of code, sets up the initial state of the 82161A at power on and then alternates between testing for a cassette to be inserted into the drive, the **REWIND** key to be pressed, and calling the HP-IL routine. This routine also executes the device-clear and power-down functions. If either command is received, the HP-IL routine flags that fact and the drive responds after it finishes its latest task and returns to the idle routine loop.

The HP-IL routine, which takes approximately 460 bytes, provides the 82161A with basic talker and listener capabilities. This routine takes care of all communication with the HP-IL interface chip and passes all necessary information to the device control routines, primarily through a set of flag registers and one data register. This polled solution to HP-IL, in contrast to an interrupt-driven solution, is required because most of the device control routines need exclusive use of the microcomputer and can only give up control at specific times.

The device control routines, which take the remaining 1420 bytes of ROM, can be further divided. One part is the command decode portion. The device control is done with device-dependent commands (DDCs). When the HP-IL routine receives a DDC that it decides is of interest to the 82161A, it passes the DDC on to the command decode routine. Either the command is executed immediately, as in the case of a read or exchange buffer operation, or flags are set to control future actions, such as write and set byte pointer where the flags control where data bytes are put. A one-byte command buffer is used to hold a DDC received when the drive is busy. The HP-IL ready-for-command (RFC) message frame following this command is not retransmitted until the present task has been completed and the new command has been decoded.

The DDL5 (format) command initializes the record positions on the tape by recording all 512 records on both tracks. Each record contains a sync byte, a byte for the record number, a second sync byte, 256 bytes of data (each data byte initialized to 255), a checksum, and a final sync byte. Only during initialization is the first sync byte and the record number written. In all following write operations, the record number is read before the remaining part of the




Fig. 2. Block diagram of electronic system of the 82161A Digital Cassette Drive.




Fig. 3. Timing diagram for signal lines and one-shot multivibrator states used to decode bit values stored on the 82161A tape cassette.

record is written. This serves two purposes. The first is to verify that the proper record is being written and the second is to fix the record position on the tape so that it does not move along the tape when it is overwritten.

The major criterion in selecting an encoding method for the 82161A was reliability. The tape drive system requires that the method have a large speed-variation tolerance and use a microcomputer to generate the encoded bit stream during a write operation. The tape lengths required to record a one and a zero should be the same so that the length of a record does not depend on the ratio of ones to zeros within the record. Also, the code should be self-clocking for easy decoding.

The method best qualified is the biphas-level or Manchester code. The rules of this code are 1) there is always a transition in a bit cell center, and its direction specifies the value of the bit, and 2) there is a transition on a bit cell edge only when the two bits on either side have the same value (see Fig. 3).

In a write operation, the time between transitions, bit cell midpoint to bit cell edge, is $64 \mu s$. During this time the transport status (stall, cassette present, and end of tape) is checked, the next nibble is read from the buffer and added to the checksum, and the next transition is calculated. The bit stream generated is sent to the sense amplifier on the DIO line.

There are two signals used in a read operation, DIN (data in) and DRDY (data ready). DRDY is the extracted clock and DIN is the latched data derived from the signal read from the tape. The microcomputer reads DIN and DRDY simultaneously and checks for DRDY to change state. When it does, the value of DIN is shifted into the register building the data. While it is in the read loop, the microcomputer also checks the transport status, stores the complete nibbles in the buffer, adds them to the checksum, and maintains a counter to detect when signal dropouts occur.

When the read/write routine is entered, the motors are turned on, the record number is read and verified, and the data portion of the record is then read or written. If a record number error and/or (in the case of a read) a checksum error is detected, the drive attempts the read or write operation a second time. If the microcomputer still detects an error, it stops the drive and reports the error to the HP-IL controller.

Seek operations are always attempted in a relative manner first. When the new record number is received, it is checked to see if it is in range (i.e., <512), and the difference between the present position and the desired position is calculated. The transport is turned on to move in the appropriate direction, and by watching the DRDY line, the microcomputer counts interrecord gaps until the transport reaches the record immediately before the desired record. This record is read, and the record number is checked. If it is correct, the transport is stopped with the desired record next. If an error is detected, the tape is rewound, and the seek is attempted again, but this time from the beginning of the tape. This gives four chances of reading the record correctly and ensures accurate seeks.

Data Storage and Retrieval

The microcomputer handles digital information to and from the read/write electronics on a bit-by-bit basis using three data-related lines (DIN, DIO, and DRDY, see Fig. 3) and two control lines (REC and TRK). DIO is a bidirectional data line whose transfer direction is controlled by the state of the REC line. In the read mode (REC low), DIO is driven by the sense amplifier, while in the write mode (REC high), the sense amplifier goes into a high-impedance state and DIO is driven directly by the microcomputer. Both DIN and DRDY are generated by the decoder circuitry and are derived solely from DIO level changes. The TRK line is driven by the microcomputer to select which tape track (0 or 1) is read from or written to.

The sense amplifier is a custom bipolar integrated circuit. It contains the signal conditioning and logic circuits to drive the magnetic head during a write operation and to translate the low-level analog signals at the head to time-related digital signals at DIO during a read operation.

Writing to the tape is accomplished by controlling the current flowing through the windings of the magnetic head. These currents produce a magnetic field across the gap at the front of the head. Three wires (two ends and a center tap) are attached to each winding (track) on the head. During a write operation, the center tap is connected to a constant-current sink, and each end of the winding is alternately driven high to control the direction of the current and thus the polarity of the magnetic field at the gap. The use of a current sink allows maximum rate of change of current, yet limits the peak direct current to 150% of that required to completely magnetize the tape.

During a read operation, the voltage at the terminals of the head is proportional to the rate of change of the magnetic flux across its gap and reaches a peak value when the gap is directly opposite a flux reversal on the tape. To decode recorded information properly, a digital signal with level changes corresponding in time to these voltage peaks must be generated. The sense amplifier generates this signal by amplifying and then differentiating the analog signal from the head. A zero crossing at the output of the differentiator corresponds to a peak of the amplified signal and is used to clock DIO level changes. The DIO level (high or low) is related to the polarity of the amplified signal at clock time and indicates the direction of the flux transition. Hysteresis is included to provide protection from unwanted transitions caused by electrical noise.

Data content is encoded by the direction of the DIO level transition at the midpoint of a bit cell. Transitions at bit cell edges are used only as required to set up DIO for the proper change at the next midpoint (see Fig. 3). The decoder hardware ignores these edge transitions and provides the microcomputer with two signals—DRDY and DIN. A change at DRDY notifies the microcomputer that the signal at DIN represents valid data.

For every DIO transition a 100-ns pulse is generated and appears at the trigger input of a nonretriggerable one-shot multivibrator. The timing period of the one-shot multivibrator is set so that, if triggered by a midcell transition, the next edge transition, when it exists, will occur during the cycle and thus be ignored. When the timing cycle expires, the level of DIO, which corresponds to the encoded bit value, is latched into the DIN flip-flop. Approximately 2 μ s later, the output of the DRDY flip-flop changes, notifying the microcomputer that data is valid (see Fig. 3).

To ensure that the one-shot multivibrator is triggered only by midcell transitions and that it ignores any cell edge transitions, a sync byte is included at the beginning of each record. This special bit pattern maps to a stream of level changes that includes only midcell transitions. Thus, the one-shot multivibrator is set up to be triggered only at midcell, and, if speed stays within allowable limits, synchronization is maintained throughout the entire record.

The encoder hardware can tolerate timing variations in DIO of $\pm 30\%$. Electronic jitter, aliasing, phase shift in the amplifiers, external electromagnetic noise, and true speed variations all contribute to the total timing variance at DIO. Fortunately, the wide acceptability range of the decoder easily overcomes these factors and ensures good unit-to-unit compatibility.

Transport Mechanism

The 82161A's mechanical design uses an $8 \times 34 \times 56$ mm minicassette designed especially for digital applications. The cassette contains nominally 24 meters of usable tape 3.81 mm wide, allowing two tracks of data 1.45 mm wide. This tape is hub-driven as opposed to capstan-driven, meaning that the only way the tape can be moved is by turning the appropriate stack of tape directly.

The selection of a hub-driven cassette was the key step in a "simplicity" approach to the design of the 82161A

mechanism. It allows the use of a two-motor drive (one per hub) and eliminates additional motors or controlled actuator devices that would be required by capstan mechanisms or single-motor drives. Another key to simplicity is the use of a two-track magnetic head. This eliminates having to move the cassette, either by the user or by a mechanism, to access both tracks. Other factors contributing to a straightforward design are the extensive use of injection-molded thermoplastics, cost-effective fasteners such as adhesives and press fits, and a low part count.

Aside from simplicity, another goal of this mechanism design was modularity. This produced a drive system module that can be removed from the 82161A and designed into other products with a minimum of change, electrical or mechanical.

The primary part of this device is the headframe assembly (Fig. 4). This assembly consists of a molded plastic frame into which a magnetic head is aligned and glued, and an optoelectronic device which forms half of the beginning-of-tape/end-of-tape (BOT/EOT) sensor. The single-gap, two-track head's coil winding parameters were chosen for both read and write functions. The headframe is molded from glass-filled polycarbonate, a very stable compound that allows some key dimensions to be held within a tolerance of 0.025 mm. Two posts on the headframe position the cassette housing relative to the head and two tape guides on the frame guide the tape relative to the head. These features are molded into the headframe.

The headframe assembly is joined with a door, window, and cassette pressure springs to form a door assembly very similar to that used in many conventional cassette tape recorders. The cassette is loaded into a slot in the door and the action of closing the door positions the cassette in the mechanism. When the door is released by pressing the **OPEN** key, a spring pulls it open so that the cassette can be easily removed.

When the door containing a cassette is shut, it pushes a pin, closing a switch spring located directly on the drive printed circuit assembly and informing the electronics that a cassette is present. A 45-degree mirror located within each cassette completes the optical BOT/EOT detection path. The optoelectronic device and the head in the headframe are connected to the drive's printed circuit assembly by a flexible ribbon cable.




Fig. 4. Exploded view of the 82161A headframe assembly.

The backbone of the transport is the mainframe. It is a precision part made of glass-filled polycarbonate formed by injection molding. All of the key subassemblies, including the door/headframe assembly, the motors and gear systems, the door latch, and the printed circuit assembly are fastened to the mainframe to form a complete modular transport mechanism.

Two identical drive motors are used. One drives the left stack of tape and is called the forward motor. The other drives the right stack and is called the reverse motor. The use of two motors coupled directly to the tape stacks in this fashion makes possible a simple speed control scheme for reading and writing data. The motor selection involved a tradeoff between motor performance (hence cost) and product capability (primarily data capacity). Ironless-rotor motors with their low inertia/torque ratio were selected. The software definition requires interrecord gaps long enough to allow the tape velocity to change between zero and read/write speed between records. Selecting low-inertia motors allows record length to interrecord gap ratios of around 3:1. If higher-inertia motors had been used, this ratio, as well as data capacity, would have been lower. Another important characteristic of ironless-rotor motors in this application is their linear tachogenerator feature. The construction of the motors is such that their EMF, with the commutation ripple filtered out, can be used to detect motor speed changes typically within 2 percent. Perhaps the most important result of the selection of low-inertia motors is the reduction of dynamic tape tension. With the two-motor, hub-drive technique, the driving motor must pull the tape against the other motor's inertia. When accelerating, the trailing motor's inertia is multiplied by the square of the gear reduction as it is reflected into the tape. Hence, peak tape tensions are very sensitive to motor inertia. Tape life was found to be dominated by dynamic tape tensions and so the long tape life achieved in the 82161A is strongly related to the selection of low-inertia motors.

Motors could not be found that would go slow enough while maintaining the torque required to drive the tape hubs directly. Therefore, a speed reducer is required. Many methods were considered, beginning with the logical choice of motor/gearbox combinations, but these proved to be too expensive. O-ring and toothed-belt drive designs were tried, but both exhibited a common problem of requiring increased belt tension to avoid slipping. The higher belt tension produced higher shaft friction, which in turn, led to increased tape tension and reduced tape and bearing life.

The 82161A uses a custom gear drive (Fig. 5) consisting of a pinion and drive gear for each motor with a ratio of 1:4 (15 teeth to 60 teeth). Both gears are injection-molded at a custom gear-molding house and exceed AGMA* quality No. 7. The diametral pitch is 96 and the gear material is lubricated acetal resin. The pinions are pressed on the motor shafts and the drive gears run free on ground stainless-steel shafts pressed into the mainframe (Fig. 5a). The motors are positioned by concentric collars fastened also to the mainframe. Precision molding of the mainframe allows the motor-to-drive-gear-shaft dimension to be held to within 0.025 mm. Also running on the drive gear shafts,

and axially coupled to the drive gears, are the drive splines (see Fig. 5b). These splines fit into the cassette hubs and transmit torque to them from the driven gears. In the event that a cassette hub does not align with the spline when the door is closed, the spline is spring-loaded and can be pushed down, allowing the door to shut. When that particular side of the drive moves, spline relative to hub, the spring pushes the spline up to engage it with the hub. These spline springs also produce a braking action or a drag friction. This drag was optimized for system speed performance by adjusting the spring constant and preload.

Head Alignment

Accurate alignment of the magnetic head in the head-frame is crucial to give unit-to-unit read/write compatibility for systems using multiple transports. This alignment is done electrically, with the head actually reading signals





Fig. 5. The tape drive system in the 82161A uses two identical motors, each driving a spring-loaded spline. (a) Close-up photograph of drive gear and motor assembly. (b) Exploded drawing of motor and drive gear system.

*American Gear Manufacturers' Association

from two tracks simultaneously, rather than aligning optically as has been done by HP in the past.^{3,4} The head aligner tool consists of an endless-loop tape deck, an oscilloscope, and an ac voltmeter. A master head on the tape deck is used to write perfectly phased signals on both tracks of a tape loop. This loop is then read by a head requiring alignment. The head is held inside a headframe by a tooling fixture that sets penetration and varies azimuth (perpendicularity of the head gap relative to tape motion) and tracking (vertical position of the head poles relative to tape position). The two signals read are observed on the oscilloscope and the azimuth is adjusted until both tracks are in phase and the combination of signals with maximum amplitudes has been found. This ensures that the azimuth has been "perfectly" aligned in the fixture. Next, the tracking is adjusted by moving the head up and down until the sum of the signals from both tracks is a maximum as measured by the voltmeter. This ensures that the pole spacing on the head optimally matches the track positions on the alignment tape. The two adjustments are practically decoupled on the alignment fixtures, so convergence is not necessary. After the head is aligned, the assembly and tooling fixture is removed, and the head is glued in place with a fast-curing acrylic adhesive. The headframe assembly can then be removed from its fixture and the alignment rechecked to observe any movement caused by glue cure. This procedure produces azimuth alignment better than ± 5 arc-minutes, and tracking alignment, relative to the headframe, better than ± 0.05 mm. The master head is periodically used to monitor the accuracy of the alignment tape. To check the master head, a Möbius tape is installed on the tape deck. This tape alternately presents front and back sides to the master head. A signal is written on the front side and read from the back side. The amplitude is reduced, but only the track-to-track phase is important. If the master head is "perfectly" aligned in azimuth, no phase difference will occur. An iterative process is used to align the master head if necessary.

System Modeling

The electromechanical system of the 82161A tape transport was modeled to allow studies of tape velocities and the effects thereon of motor parameters and mechanism inertias and frictions. The basic model is shown in Fig. 6 and the basic equations of motion are

$$\begin{aligned} I_1 \ddot{\theta}_1 &= R_1 [K_2 (R_2 \ddot{\theta}_2 - R_1 \ddot{\theta}_1)] \\ I_2 \ddot{\theta}_2 &= -R_2 [K_2 (R_2 \ddot{\theta}_2 - R_1 \ddot{\theta}_1)] + R_3 [K_1 (R_4 \ddot{\theta}_3 - R_3 \ddot{\theta}_2)] \\ I_3 \ddot{\theta}_3 &= -R_4 [K_1 (R_4 \ddot{\theta}_3 - R_3 \ddot{\theta}_2)] + R_2 [K_2 (R_1 \ddot{\theta}_4 - R_2 \ddot{\theta}_3)] \\ I_4 \ddot{\theta}_4 &= -R_1 [K_2 (R_1 \ddot{\theta}_4 - R_2 \ddot{\theta}_3)] \end{aligned}$$

This multi-degree-of-freedom system was studied by identifying different motor-to-tape-stack and motor-to-motor modes whose frequencies depend on the reduction method (belts or gears) and tape stack ratios. Many of the natural frequencies found by this model can be satisfactorily filtered out by altering the servo design, but one mode consistently showed up in the analysis that cannot. This was identified as a "tape mode" or the opposing oscillation of both hub stacks with the spring being the length of tape




Fig. 6. Diagram defining tape motor drive parameters used to model 82161A transport system behavior.

between the two stacks. During read/write operation, when the servo is controlling tape speed, this tape mode, if excited, is superimposed on the steady-state tape velocity and becomes what was found to be the major cause of tape speed jitter. In all of the configurations studied, this jitter frequency is near 500 Hz. This tape mode cannot be effectively corrected by the servo because it is totally isolated from the motors and hence is "invisible" to the servo. This tape mode creates oscillations in tape tension, but the magnitude of the oscillations was found to be less than the steady-state tape tension. Thus, the treatment of the tape as a linear spring was not negated by the fact that the tape cannot be put into compression. This phenomenon will be discussed further in reference to the slow-start circuitry in the motor drive.

The tape mode oscillations would not be a major difficulty if there was no 500-Hz excitation in the mechanism to get them going. Unfortunately, because of the size of the transport, the gear pitch selection was limited, and 500-Hz perturbations caused by gear backlash are unavoidable. Hence, two methods are used to damp the tape mode. First, friction is added by using the spline spring as described previously, and second, viscous rubber snubbers are added between the drive gear and the spline (see Fig. 5). In this position, the snubbers serially provide a viscoelastic element between the perturbation (gear backlash) and the elements (the cassette stacks) producing the tape mode oscillations. The combination of these two damping schemes reduces speed jitter by approximately 50% to a range of 10% of average tape speed.

Motor Drive

The 82161A mechanism combines software, electronics, and mechanics to control both the position and the velocity of the tape. TTL-compatible inputs to the motor drive circuitry allow the microcomputer to select any of five possible modes of operation.

The fast forward and rewind modes move the tape at 76 to 152 cm/s, during which time the microcomputer counts interrecord gaps to determine tape position (record number). Once the desired position has been reached, the slow forward mode is activated for a data read/write operation. Forward and reverse braking is accomplished by using the back EMF of the trailing motor to generate a reverse torque to decelerate the system.

The fast forward, rewind, and slow forward modes use the leading motor as the actuator and the trailing motor is "pulled" by the tape. The no-load friction of the trailing motor and its associated gears provides tape tension to aid speed control and help keep the tape in contact with the magnetic head. The forward and reverse braking modes use the trailing motor as the actuator and the tape as the mechanical link to decelerate the leading motor.

The heart of the motor drive electronics is the velocity control circuitry (Fig. 7). To ensure read/write compatibility, linear tape velocity past the magnetic head must be a controlled, repeatable function of tape position. Although holding the angular velocity of one motor constant would satisfy this objective, tape capacity would be severely limited because the linear tape speed would vary over a wide range as the radii of the takeup and supply reels, respectively, increase and decrease. However, holding the sum of the angular velocities of both motors constant not only satisfies the above requirements, but dramatically increases data capacity by maintaining a more uniform linear tape velocity.

The input to the servo is a controllable reference voltage. The servo acts to hold the sum of the back EMFs of the two motors equal to this reference. As shown in Fig. 7, the forward transfer path consists of an error amplifier, a power stage, and the mechanical system. The back-EMF summer forms the feedback path. All necessary frequency compensation is implemented in the error amplifier and includes a

pole at the origin to integrate out dc errors, a low-frequency zero at 4 Hz to compensate a pole of the mechanical system, and a second pole at ≈ 40 Hz to filter out unwanted motor commutation noise that appears in the feedback signal. At frequencies within the range of interest (40 Hz), the open-loop transfer function of the system, including compensation, consists of a single pole at the origin. Local feedback for the error amplifier is derived from the output of the power stage to minimize crossover distortion. As discussed earlier, the transfer function of the mechanical system is quite complex and includes several oscillatory modes. Fortunately, these modes are either at frequencies well outside the bandwidth of the servo or are invisible to the servo so that no serious electronic stability problems arise.

A novel feature of this servo is the speed sensor which sums the back EMF from each motor. Since no current flows through the trailing motor, the back EMF is simply its terminal voltage and is readily available to determine motor speed. However, the current required to produce drive torque generates a voltage across the rotor resistance of the leading motor which is superimposed on its back EMF. In the past, this speed measurement problem has been avoided by using either pulse-width modulators, which sample back EMF by momentarily removing power, or transducers which do not rely on back EMF. For this application, low-frequency pulse-width modulators would dissipate additional power in the motor and generate electrical and mechanical noise caused by their switching transients. Transducers are too expensive, too large, and require too much additional hardware.

The chosen scheme dynamically sums the terminal voltage of each motor and subtracts the voltage caused by the drive currents in the leading motor. Referring to Fig. 7,

$$V_0 = EMF_L + I_M R_M - [(EMF_L + I_M R_M + I_M R_S) - (EMF_L + I_M R_M)](R_3/R_2) + [(EMF_L + I_M R_M) - (EMF_L + I_M R_M - EMF_T)](R_3/R_1)$$

If $R_1 = R_3$, and canceling terms,

$$V_0 = EMF_L + I_M R_M - I_M R_S (R_3/R_2) + EMF_T$$

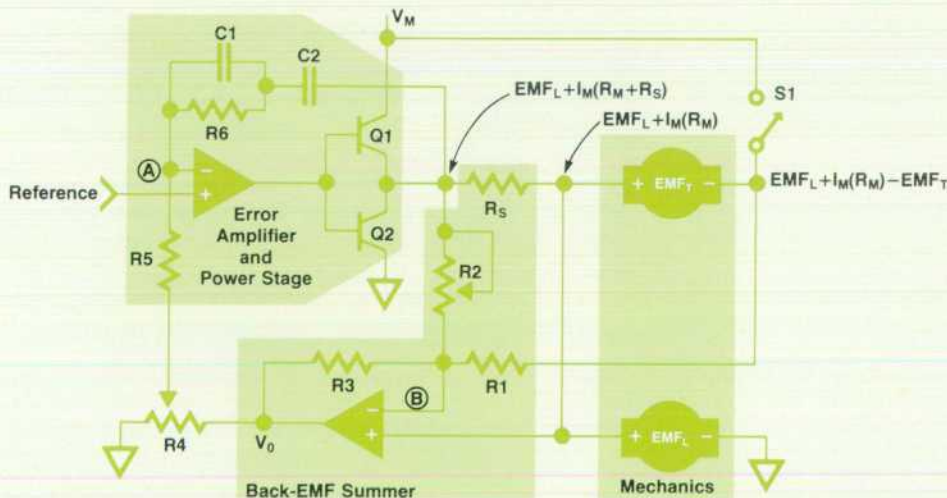


Fig. 7. Simplified schematic of velocity control servo.

Then, if R_2 is adjusted such that $R_M/R_S=R_3/R_2$,

$$V_0=EMF_L+EMF_T$$

As can be seen from the derivation above, if resistance matching is done (using potentiometer R_2), the output of the feedback amplifier is the true sum of the back EMF of the motors. R_S is specified as a copper wirewound resistor so that its temperature coefficient of resistivity will match that of the motor's ironless rotor, thus holding the ratio of R_M to R_S constant and ensuring consistent speed control over the full operating temperature range.

The fast forward and rewind modes are implemented by "fooling" the servo. Grounding point A in Fig. 7 eliminates the feedback and causes the output of the error amplifier to go high and drive the motor at high forward speed. Forcing point B low causes V_0 to be high, thus forcing the output of the error amplifier low. This, in combination with closing switch S1 and lifting the ground on the leading motor (using transistor switches), results in the rewind mode.

To use the feedback from both motors to control speed, it is essential that the motors be mechanically linked in a predictable, linear fashion. In the 82161A, this link is the tape. Because the tape cannot support compressive forces, slack in the tape can occur and totally uncouple the leading and trailing motors. The typical result is a "bang-bang" servo action. The leading motor is driven until its back EMF equals the reference value. Suddenly the tape slack is taken up and the trailing motor begins to move and injects a step function into the feedback signal. The error amplifier responds by slowing the leading motor, which allows the trailing motor to spool up and form another loop. This starts the process all over again.

Once the 82161A has attained stable slow forward operation, this problem is prevented by the tape tension generated by system friction. However, when the slow-forward mode is initiated, there is always some amount of slack in the tape, and this slack must be eliminated first before accelerating to full speed. In addition, since overshoot can cause the same problem, the rate of change of the speed reference voltage must be slowed to the point where the servo can keep up.

The slow-start circuit performs these functions by controlling the reference voltage to the servo. When the slow forward mode is selected, the reference is held to a low value for approximately 130 ms, during which time the slack is removed from the tape. Then, the reference voltage rises exponentially towards an asymptotic value, allowing a smooth acceleration to read/write speed without overshoot.

Acknowledgments

The authors would like to thank Roger Quick and Tom Braun for their direction and leadership in the development of the 82161A, and George Custer for the product package design, head sourcing, and innumerable other design details. Also greatly appreciated is the work of Mark Matsler in developing the head alignment process, and the hard work of everyone else who helped bring the 82161A Digital Cassette Drive to production.

References

1. R.D. Quick and S. L. Harper, "HP-IL: A Low-Cost Digital Interface for Portable Applications," Hewlett-Packard Journal, Vol. 34,

no. 1, January 1983.

2. R.D. Quick and D.L. Morris, "Evolutionary Printer Provides Significantly Better Performance," Hewlett-Packard Journal, Vol. 31, no. 3, March 1980.

3. D.J. Collins and B.G. Spreadbury, "A Compact Tape Transport Subassembly Design for Reliability and Low Cost," Hewlett-Packard Journal, Vol. 31, no. 7, July 1980.

4. R.B. Taggart, "Designing a Tiny Magnetic Card Reader," Hewlett-Packard Journal, Vol. 25, no. 5, May 1974.



William A. Buskirk

Bill 'Buzzy' Buskirk joined HP in 1977 after receiving a BSEE degree from the University of Colorado. He worked in production engineering on various calculators before moving to R&D in 1978. Bill worked on the card reader for the HP-41 Handheld Computer and the 82161A Cassette Drive before assuming his current responsibility as a project manager. He was born in Bloomington, Indiana and during his undergraduate studies, worked for the National Oceanic and Atmospheric Administration. Bill is married, has a son, and lives in Albany, Oregon. Outside of work, he is kept busy remodeling his home, reading, camping, working with home computers, and helping his wife run their gift shop.



Charles W. Gilson

Charlie Gilson graduated from California Polytechnic State University with a BSME degree in 1973. He worked three years on computer modeling of missile shock isolation and air-launch systems before joining HP in 1976. He worked on the mechanical design of various calculators and the 82161A Cassette Drive's transport mechanism before moving to production engineering to do cost-reduction design. Born in San Francisco, California, he now lives in King's Valley, Oregon. He is married and has two children, a girl and a boy. His interests include raising sheep and slowly remodeling an old farmhouse.



David J. Shelley

Dave Shelley was born in Seattle, Washington and attended the nearby University of Washington, earning a BSEE degree in 1973. He then joined HP's San Diego Division and contributed to the electrical design for the 7245A Plotter/Printer and the 9872A Graphics Plotter. In 1977, he transferred to HP's Corvallis Division and worked on the electrical design for the 82143A Printer and the 82161A Cassette Drive. Dave is currently a project manager for cost reduction of portable computers. He is named coinventor for a patent related to the thermal head design for the 7245A. Dave is married, has two sons, and lives in Corvallis, Oregon. He enjoys bicycling, camping, and playing racquetball.

Scientific Pocket Calculator Extends Range of Built-In Functions

Matrix operations, complex number functions, integration, and equation solving are only some of the numerous preprogrammed capabilities of HP's latest scientific calculator, the HP-15C.

by Eric A. Evett, Paul J. McClellan, and Joseph P. Tanzini

THE NEW HP-15C Scientific Programmable Calculator (Fig. 1) has the largest number of preprogrammed mathematical functions of any handheld calculator designed by Hewlett-Packard. For the first time in an HP calculator, all arithmetic, logarithmic, exponential, trigonometric, and hyperbolic functions operate on complex numbers as well as real numbers. Also, built-in matrix operations are provided, including addition, subtraction, multiplication, system solution, inversion, transposition, and norms.

The HP-15C also performs the **SOLVE** and \int_y^x functions, which are very useful tools in many applications. The **SOLVE** operator numerically locates the zeros of a function programmed into the calculator by the user.¹ The \int_y^x operator numerically approximates the definite integral of a user-programmed function.²

Design Objectives

The HP-15C was designed with the following goals in mind:

- Provide all functions of the HP-11C and HP-34C Calculators in the same slim-line package used for the HP-11C
- Provide additional convenient, built-in advanced mathematical functions which are widely used in many technical disciplines.

Achieving the first objective posed a keyboard layout problem. The nomenclature for the HP-11C functions filled every position on the keyboard. Since the display is limited to seven-segment characters, functions could not be removed from the keyboard and accessed by typing the function name as is done on the HP-41 Handheld Computers. Therefore, to free some space on the keyboard, only the two most common conditional tests are placed on the keyboard, $x=0$ and $x \leq y$. A **TEST** prefix is added to access the other ten




Fig. 1. The HP-15C is an advanced programmable calculator with special functions that enable the user to solve problems involving matrices, integrals, complex arithmetic, and roots of equations. Its slim-line design fits easily in a shirt pocket.

tests by executing **TEST 0**, **TEST 1**, . . . , **TEST 9**. A table on the back of the calculator indicates the correspondence between the digits and tests. This frees enough positions on the keyboard to add the **SOLVE** and \int_y^x functions, plus a few more.

In striving to attain the second objective, we noted that nearly every text covering advanced mathematics for science and engineering includes chapters on complex analytic functions and matrix algebra. They are fundamental tools used in many disciplines. Furthermore, the complex functions and many of the matrix operations can be viewed as extensions of the functions already on the keyboard. This is an important consideration because of the limited number of key positions available. Thus, our goal was to extend the domain of some of the built-in functions to complex numbers and matrices in a natural way without altering how those functions operate on real numbers.

Complex Mode

A complex mode was introduced in which another register stack for imaginary numbers is allocated parallel to the traditional register stack for real numbers (Fig. 2). Together they form what is referred to as the complex RPN* stack.

The real X register is always displayed. A complex number $a+ib$ is placed in the X register by executing **a**, **ENTER**, **b**, **i**. The user may display the contents of the imaginary X register by executing **Re** \leq **Im** to exchange the contents of the real and imaginary X registers. Or the user may hold down the **(i)** key to view the imaginary part without performing an exchange.

ENTER, **R** \downarrow , **R** \uparrow , **x** \leq **y**, and **LST x** all operate on the complex stack, but **CLx** and **CHS** operate only on the real X register so that one part of a complex number can be altered or complemented without affecting the other. For example, the complex conjugate is performed by executing **Re** \leq **Im**, **CHS**, **Re** \leq **Im**.

The following functions include complex numbers in their domain: **+**, **-**, **x**, **÷**, **1/x**, \sqrt{x} , **x²**, **ABS** (magnitude), **LN**, **e^x**, **LOG**, **10^x**, **y^x**, **SIN**, **COS**, **TAN**, **SIN⁻¹**, **COS⁻¹**, **TAN⁻¹**, **SINH**, **COSH**, **TANH**, **SINH⁻¹**, **COSH⁻¹**, and **TANH⁻¹**. These functions assume the complex inputs are in the rectangular form, $a+ib$.

Often complex numbers are expressed in polar form: $re^{i\theta} = r(\cos\theta + isin\theta)$. In complex mode, the polar-to-rectangular conversion functions **→P** and **→R** provide a

*Reverse Polish notation, a logic system that eliminates the need for parentheses and "equals" keystrokes in calculator operations.





Fig. 2. To handle complex numbers, the HP-15C uses another register stack in parallel with the traditional RPN stack. Only the contents of the X register in the real stack are displayed.



$$Z_1 = 3 - j4$$

$$Z_2 = 10$$

$$Z_{eq} = 1 / (1/Z_1 + 1/Z_2)$$

Fig. 3. The complex arithmetic capabilities of the HP-15C make it easy to solve for the equivalent impedance of this parallel circuit (see text).

convenient means for converting between the polar and rectangular forms of a complex number.

Complex numbers are used extensively in electrical engineering. For example, to find the equivalent impedance in the parallel circuit shown in Fig. 3, perform the following steps on the HP-15C:

Keystrokes	Calculation
3 ENTER 4 CHS i	Z_1
1/x	$1/Z_1$
1 0	Z_2
1/x	$1/Z_2$
+	$1/Z_1 + 1/Z_2$
1/x	$Z_{eq} = 1 / (1/Z_1 + 1/Z_2)$ Hold down (i) key to view imaginary part. $Z_{eq} = 2.9730 - 2.1622i$
→P	Convert to phasor form. $Z_{eq} = 3.6761 \angle -36.0274^\circ$

This example is a very elementary application of the built-in complex function capability. Since complex operations can be used in conjunction with the **SOLVE** and \int_y^x functions, the HP-15C can be programmed to carry out some sophisticated calculations such as computing complex line integrals and solving complex potentials to determine equipotential lines and streamlines.³

Matrix Descriptors

No set of matrix operations is complete without addition, subtraction, multiplication, system solution, and inversion. To provide these operations on the HP-15C, it seemed natural to extend the domains of the **+**, **-**, **x**, **÷**, and **1/x** functions to include matrix arguments. Since these functions operate on the stack contents, a means of placing a matrix name (descriptor) on the stack is essential. The set of alpha characters that can be represented in a seven-segment font is limited, but the letters A, B, C, D, and E have reason-



Fig. 4. The seven-segment font used in the HP-15C's liquid-crystal display allows representations of the alphabetic characters A, B, C, D, and E as shown above for use in labeling matrices.

able representations (Fig. 4).

Thus the decision was made to allow up to five matrices to reside in memory simultaneously, named **A**, **B**, **C**, **D**, and **E**. Their descriptors are recalled into the X register by the sequence **RCL MATRIX** followed by the appropriate letter. When the X register contains a matrix descriptor, the matrix name and dimensions are displayed. Matrix descriptors may be manipulated by stack operations and stored in registers just like real numbers, and certain functions accept matrix descriptors as valid inputs. For example, suppose **C** and **D** are 2-by-3 and 3-by-4 matrices, respectively, which are already stored in memory. To compute the matrix product **CD** and place the result in matrix **A**, the user parallels the steps required for real number multiplication, except that the result destination must be specified:

Keystrokes	HP-15C Display
RCL MATRIX C	C 2 3
RCL MATRIX D	d 3 4
RESULT A	d 3 4

At this point, the HP-15C's RPN stack contains the information shown in Fig. 5a. The matrix operands are in the stack, and the result matrix is specified. The user now executes **x** to compute the matrix product. When **x** is executed, the presence of the matrix descriptions in the Y and X registers is detected, the matrices are checked for compatible dimensions, the result matrix **A** is automatically dimensioned to a 2-by-4 matrix, the product is computed, and the matrix descriptor of the result is placed in the X register and displayed (Fig. 5b).

The operators **+** and **-** work similarly, and **÷** performs the matrix operation $X^{-1}Y$ if the X and Y registers contain matrix descriptors. This is useful for linear system solution, since the solution to the matrix equation $XR=Y$ is $R=X^{-1}Y$. The **1/x** function key performs matrix inversion.

Other important matrix operations that are not natural extensions of functions on the keyboard are accessed by the prefix **MATRIX** followed by a digit. These include transpose, determinant, and matrix norms. A table on the back of the calculator indicates the correspondence between the digits and matrix operations.

Internal Format of Descriptors

Normal floating-point numbers are internally rep-

resented in the HP-15C by using a 14-digit (56-bit) binary-coded-decimal (BCD) format (Fig. 6).

The exponent *e* is given by *XX* if *XS*=0, and by $-(100-XX)$ if *XS*=9. The value of the number is interpreted as $(-1)^S(M.MMMMMMMMM) \times 10^e$. For example,

01234000000002 represents 1.234×10^2

and

91234000000994 represents -1.234×10^{-6} .

Matrix descriptors, on the other hand, are distinguished by a 1 in the mantissa sign digit and a hexadecimal digit corresponding to the matrix name in the most significant digit of the mantissa field. For example, the matrix descriptor **C** is represented internally as 1C000000000000.

When a matrix descriptor is detected in the X register, the matrix name is displayed, and the current dimensions of that matrix are fetched from a system memory location and also displayed.

Creating Matrices and Accessing Individual Elements

A matrix is dimensioned by entering the row and column dimensions in the Y and X registers of the stack, respectively, and then executing the **DIM** prefix followed by the matrix name. Individual matrix elements are accessed by executing the **STO** or **RCL** prefixes followed by the matrix name. The element accessed is determined by the row and column indexes stored in registers **R0** and **R1**, respectively.

Matrix data is usually entered or reviewed from left to

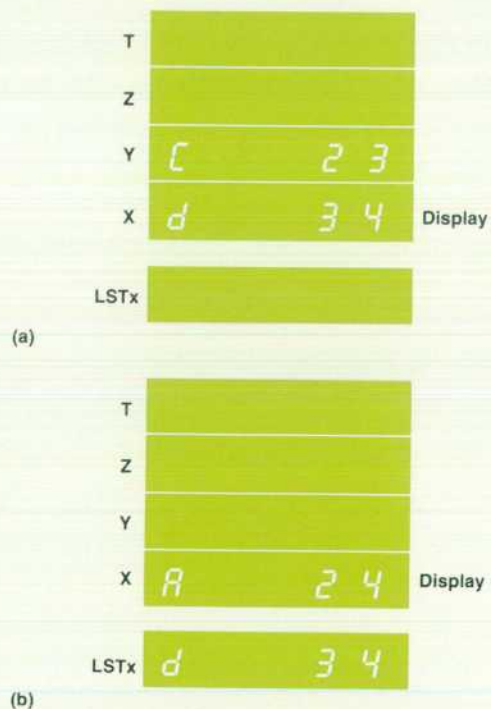


Fig. 5. Before multiplying matrices **C** and **D**, the information in the RPN stack is located as shown in (a). After multiplication, the result, matrix **A**, is located as shown in (b) and the **LSTx** register contains the information for matrix **D**.

right along each row and from the first row to the last. To facilitate this process, a user mode is provided in which the indexes are automatically advanced along rows after each **STO** or **RCL** matrix access operation. After the last element of the matrix has been accessed, the indexes wrap around to 1,1. As an added convenience, executing **MATRIX 1** initializes the indexes to 1,1.

The following example illustrates some of these features by solving the following matrix equation for C:

$$\begin{bmatrix} 5 & -2 \\ 4 & 6 \end{bmatrix} \begin{bmatrix} c(1,1) & c(1,2) \\ c(2,1) & c(2,2) \end{bmatrix} = \begin{bmatrix} 8 & 3 \\ 2 & -6 \end{bmatrix}$$

Keystrokes	Display	Comments
USER		Select USER mode.
MATRIX 1		Initialize indexes in registers R0 and R1 to 1.
2 ENTER DIM A	2.0000	Dimension matrix A to 2 by 2.
DIM B	2.0000	Dimension matrix B to 2 by 2.
5 STO A	5.0000	a(1,1)
2 CHS STO A	-2.0000	a(1,2)
4 STO A	4.0000	a(2,1)
6 STO A	6.0000	a(2,2). Indexes wrap around to 1,1.
8 STO B	8.0000	b(1,1)
3 STO B	3.0000	b(1,2)
2 STO B	2.0000	b(2,1)
6 CHS STO B	-6.0000	b(2,2). Indexes wrap around to 1,1.
RCL MATRIX B	b 2 2	Recall right-hand side
RCL MATRIX A	A 2 2	Recall coefficient matrix
RESULT C	A 2 2	Specify matrix C as result.
÷	C 2 2	Compute $C=A^{-1}B$.
RCL C	1.3684	c(1,1)
RCL C	0.1579	c(1,2)
RCL C	-0.5789	c(2,1)
RCL C	-1.1053	c(2,2)



Fig. 6. The internal representation for floating-point numbers in the HP-15C uses a 14-digit (56-bit), binary-coded-decimal format.

Available Matrix Memory, Speed

A maximum of 64 matrix elements can be distributed among the five matrices. Since the HP-15C can invert matrices in place, up to an 8-by-8 matrix can be inverted. There is also enough memory to solve a 7-by-7 linear system of equations. Table I specifies the approximate time required to perform certain matrix operations.

Table I
Time in Seconds for Selected Matrix Operations

Order of Matrix	Determinant	Solving a System	Matrix Inversion
1	0.5	0.5	0.5
2	1.3	2.0	1.8
3	2.8	4.2	5.3
4	5.3	7.6	12
5	9.1	12	22
6	14	19	36
7	21	28	55
8	30	—	80

Designing the Complex Function Algorithms

After deciding to extend the real-valued functions and the RPN stack to the complex domain, our next step was to design the algorithms for complex arithmetic. Although their defining formulas are very simple, some disturbing examples made us question what accuracy should be achieved to parallel the high quality of the real-valued functions.

The real functions are generally computed with a small relative error (less than 6×10^{-10}) except at particular points of certain functions, where it is too costly in execution time or ROM space for the result to be computed that accurately.³

The relative difference $R(x,y)$ between two numbers x and y is given by

$$R(x,y) = \frac{|x-y|}{|y|}$$

When X is an approximation of x , then we say $R(X,x)$ is the relative error of the approximation X . Notice that the size of the relative error is related to the number of digits that are accurate. More precisely, $R(X,x) < 0.5 \times 10^{-n}$ implies that X

is an approximation to x that is accurate to n significant digits.

If we always wish to obtain small relative errors in each component of a complex result, then the outcome of the following example is very disappointing. For simplicity we will use four-digit arithmetic, instead of the 13 digits used internally to calculate the 10-digit results delivered to the X register of the calculator.

Example 1: Using the definition for complex multiplication,

$$(a + ib)(c + id) = (ac - bd) + (ad + bc)i,$$

consider the four-digit calculation of $Z \times W$, where $Z = 37.1 + 37.3i$ and $W = 37.5 + 37.3i$. We get,

$$\begin{aligned} Z \times W &= (1391 - 1391) + (1384 + 1399)i \\ &= 0 + 2783i \end{aligned}$$

Since the exact answer is $-0.04 + 2782.58i$, it is clear that accurate components are not always achieved by a simple application of this formula. The difference $a \times c - b \times d$ has been rounded off to result in a loss of all significant digits of the real part. The loss can be eliminated, but the calculation time would increase roughly by a factor of 4. Is it really worth this higher cost in execution time? For comparison we will consider an alternative definition of accurate complex results.

Complex Relative Error

As with real approximations we often want our errors small relative to the magnitude of the true answer. That is to say, we want $|(approximate\ value) - (true\ value)| / |(true\ value)|$ to be small enough for our purposes. So relative error may be extended to the complex plane by $R(Z, z) = |Z - z| / |z|$. This extension may be applied to vectors in any normed space. A simple geometric interpretation is illustrated in Fig. 7. Approximations Z of z will satisfy $R(Z, z) < \delta$ if and only if the points Z lie inside the circle of radius $\delta|z|$ centered at z . This condition for complex relative accuracy is weaker than that for component accuracy. If the errors in each component are small, then the complex error is small. To show this, assume that $R(X, x) < \delta$ and $R(Y, y) < \delta$ where $z = x + iy$. Then,

$$\begin{aligned} R(Z, z) &= \frac{|(X - x) + i(Y - y)|}{|z|} \\ &\leq \frac{|X - x|}{|z|} + \frac{|Y - y|}{|z|} \\ &\leq R(X, x) + R(Y, y) \\ &< 2\delta \end{aligned}$$

Actually, $R(Z, z)$ is less than δ , but this is slightly more difficult to show. On the other hand, however, a small complex error does not imply small component errors. Referring back to Example 1, we see that $R(ZW, zw) = 0.0002$, which is respectably small for four-digit precision, even though the real component has no correct digits.

It is not unusual for only one component to be inaccurate when the result is computed accurately in the sense of complex relative error. In fact, because the error is relative to the size $|z|$, and because this is never greatly different from the size of the larger component, only the smaller component can be inaccurate.




Fig. 7. A simple geometric representation of complex relative error $R(Z, z) < \delta$.

To show this we shall assume, without loss of generality, that $|x|$ is the larger component. Then

$$|z|/|x| = \sqrt{1 + |y/x|^2},$$

which implies that $1 \leq |z/x| \leq \sqrt{2}$, since $|y| \leq |x|$ by assumption. Thus $|x|$ and $|z|$ do not differ greatly. The important part is that $|x| \geq |z|/\sqrt{2}$. This gives

$$|X - x|/|x| \leq |Z - z|/|x| \leq \sqrt{2} |Z - z|/|z| < \sqrt{2} R(Z, z)$$

So the relative error of the larger component (assumed to be x here) is very nearly as small as the complex relative error bound $R(Z, z)$. It also follows that the smaller component is accurate relative to the larger component's size (i.e., $|Y - y|/|x| \leq |Z - z|/|x| \leq \sqrt{2} R(Z, z)$).

This provides a quick way to determine which digits of a calculated value can possibly be incorrect when it is known that the calculated value has a certain complex error. By representing the smaller component with the exponent of the larger component, the complex error indicates the number of correct digits in each component.

For instance, in Example 1 we obtained the approximation $Z = 0 + 2783i$ of the true answer $-0.04 + 2782.58i$. Since the larger component is 2.783×10^3 we will represent the first component with the same exponent (0.000×10^3) to obtain $Z = 0.0000 + 2783i$. These components must be accurate to nearly four digits since $R(Z, z) = 0.0002$.

Perhaps the zero component of Z confuses the issue here, so another example may be appropriate. First, let

$$Z = 1.234567890 \times 10^{-10} + 2.222222222 \times 10^{-3}i$$

Then think of Z as

$$Z = 0.0000001234567890 \times 10^{-3} + 2.222222222 \times 10^{-3}i$$

If the complex relative error indicates 10-digit accuracy, i.e., $R(Z, z) < 0.5 \times 10^{-10}$, then this implies that the first 10 digits are correct, that is,

$$Z = 0.000000123 \times 10^{-3} + 2.222222222 \times 10^{-3}i$$

Error Propagation

We have seen that computing the product of two complex numbers in the straightforward manner does not necessarily result in a small error in each component (Example 1). However it can be shown that the product does have a complex relative error bound of roughly 10^{-n} whenever n digits of precision are used in the calculation. Moreover, small relative errors in the input values give rise to relative errors nearly as small in the output values. This is not true for small component errors. One acceptable rounding error in an input value may produce an inaccurate component, even when the multiplication is exact. This is illustrated by the following example.

Example 2: Let $z = (1 + 1/300) + i$ and $w = 1 + i$, then using four-digit precision we have

$$Z = 1.003 + 1.000i \\ \text{and } W = 1.000 + 1.000i$$

Therefore,

$$ZW = 0.003 + 2.003i \\ = 3.000 \times 10^{-3} + 2.003i$$

exactly, yet

$$zw = 3.333 \times 10^{-3} + 2.003i$$

to four digits. The single rounding error of $1 + 1/300 \rightarrow 1.003$ in the component of the input Z was magnified from a relative error of 0.0003 to 0.1.

So, in general, computing accurate components will not improve the result of a chain calculation because intermediate input values are often inexact (this is the idea of backward error analysis and is explained more fully in reference 3). It is important to realize that this is not, in itself, a good reason to forsake accurate results based on the assumption that the input values are not exact. For example, if we assume that X has an error in its eleventh digit and thus decide that $\sin(X)$ for $X > 10^5$ degrees, say, need not be computed accurately, then we would have failed to provide a useful result for those special cases where we know that the input value is exact.

As a simple illustration consider accurately calculating the value $\sin(1,234,567,899.1234567890)$ where the argument is in degrees. Using

$$\sin(1,234,567,899) = 0.9876883406$$

is grossly inaccurate. Instead, let $x = 1.234567899 \times 10^9$ and $y = 0.123456789$, then evaluate

$$\sin(x+y) = \sin(x)\cos(y) + \cos(x)\sin(y).$$

Here we know x is exact, and since $\sin(x)$ and $\cos(x)$ are computed accurately by the HP-15C, the final result $\sin(x+y) = 0.9873489744$ is very accurate.

The point here is that clean results (in particular accurate components) are desirable, but in our estimation the cost of adding ROM and increasing execution time was too high on this machine to provide complex arithmetic that is accurate in each component. However, accurate components are delivered in those functions where it is more practical. This is

discussed further in the following section.

In general, the HP-15C delivers complex results that satisfy $R(Z,z) < 6 \times 10^{-10}$, except where functions involving trigonometric calculations (in radians) are evaluated at very large arguments or near transcendental zeros such as multiples of π . This inaccuracy is embedded in the real-valued functions and is an example of an error that is too costly to correct completely.^{3,4}

Some Specific Complex Functions

For complex arithmetic we obtained accurate results (i.e., small complex relative errors) from the standard formulas used to define each operation. But, in general, defining formulas are usually not accurate for computers. In this section we will single out two particular functions, $\sin(z)$ and \sqrt{z} , and very briefly focus on some difficulties that arise.

■ **Sin(z).** A typical defining formula for the complex sine function is given by

$$\sin(z) = \frac{e^{iz} - e^{-iz}}{2i} \quad (1)$$

If this is used to compute $\sin(z)$ for small $|z|$, the two exponential terms will be nearly equal and thus cause a loss of accuracy. This will result in a large complex relative error even though each step of the calculation is very accurate. If equation (1) is replaced by

$$\sin z = \sin(x) \cosh(y) + i \cos(x) \sinh(y) \quad (2)$$

where $z = x + iy$, then the relative error problem for small $|z|$ will be solved, and furthermore the components will become accurate (except for the trigonometric difficulty with large angles mentioned earlier). To observe the striking difference in results, we calculate

$$w = \sin(1.234567 \times 10^{-5} + 9.876543 \times 10^{-5}i)$$

for each formula. The outcome is represented below.

Eqn.	W (10-digit calculation of w)	R(W,w)
(1)	$1.234567006 \times 10^{-5} + 9.876530000 \times 10^{-5}i$	10^{-6}
(2)	$1.234567006 \times 10^{-5} + 9.876543015 \times 10^{-5}i$	10^{-10}

The HP-15C's internal calculation is based on equation (2), with minor modifications that exploit the relationships between the real functions to eliminate redundant computation.

■ **\sqrt{z} .** The most common definition of the principal square root is

$$\sqrt{z} = \sqrt{|z|} e^{i\theta/2} \quad (3)$$

where θ is the Arg (z), satisfying $-\pi < \theta \leq \pi$.

This formula is accurate with respect to complex relative error, but not accurate in each component. This can be seen by working through the calculation of $\sqrt{-1}$, where $a = -1 + (-1 \times 10^{-15}i)$, with 10-digit precision. Here $\theta/2$ rounds to precisely 90 degrees, thus causing $\sqrt{a} \rightarrow 0 - i$, while the true

value rounds to $5 \times 10^{-16} - i$. The complex error is small but certain information in the real component is lost. The fact that \sqrt{a} lies on the right side of the imaginary axis can be critical when computing near discontinuities called branch cuts. For example, $\ln(-i\sqrt{a}) \approx -i\pi/2$, but the inaccurate component of \sqrt{a} will cause it to evaluate to $i\pi/2$ since $-i\sqrt{a}$ is near the branch cut of $\ln(z)$. More will be said about branch cuts in the next section.

It turns out that \sqrt{z} can be computed with accurate components and without loss in execution time. This function, along with the inverse trigonometric and hyperbolic functions, is computed on the HP-15C with accurate components. Their algorithms are not described by a simple formula as with $\sin(z)$ in equation (2), but rather are described in terms of their components. These accurate components are achieved by recognizing and eliminating errors such as those described above.

Principal Branches

The function \sqrt{z} is an inverse function of $f(z) = z^2$. As is often the case with defining inverses, we must select from more than one solution to define the principal branch of the inverse. This is done for the real function by selecting the non-negative solution of $x^2 = a$ and denoting it by \sqrt{a} . Because of the branch point at 0, any branch for \sqrt{a} must have a discontinuity along some slit (branch cut). In equation (3) above, it is along the negative real axis. Notice in Fig. 8 that values below the negative real axis map to values near the negative imaginary axis, while above the slit, values map near the positive imaginary axis. Since it is traditional to have $i = \sqrt{-1}$ we must attach the slit (negative real axis) to the upper half plane, making it continuous from above and not from below, that is, $-\pi < \theta \leq \pi$. One will occasionally see \sqrt{z} defined for $0 \leq \theta < 2\pi$, which places the discontinuity along the positive real axis. We have avoided doing something like this in the branches of all of the complex inverse functions so that each will be analytic in a region about its real domain. This is important since complex computation is often performed in a region about the real domain in which the function's values are defined by the analytic continuation from the real axis.

The placement of the branch cuts and the function values along the slit are fairly standard for \sqrt{z} and $\ln(z)$, but the inverse trigonometric and hyperbolic functions have not, as yet, become standardized. However, by following a few reasonable rules there is not much room for variation.

The first rule, analyticity about the real domain, has already been mentioned. Secondly, we have tried to preserve fundamental relationships such as the oddness or evenness of functions (e.g., $\sin(-z) = -\sin(z)$) and the computational formulas relating functions to the standard principal branches of $\ln(z)$ and \sqrt{z} (e.g., $\pi/2 - \sin(z) = g(z) \sqrt{1-z}$ where $g(z)$ is analytic at 1, that is, a power series in $z-1$).

The determination of formulas involving a choice of branches is often quite complicated. W.M. Kahan has presented a very enlightening discussion⁵ of branch cuts and has pointed out to us that the HP-15C branch cuts should satisfy certain simple formulas relating them to the principal branch of $\ln(z)$. These formulas are satisfied and are reproduced below.

$$\ln(z) = \ln(|z|) + i \operatorname{Arg}(z)$$

and

$$\sqrt{z} = \exp(\ln(z)/2)$$

with $-\pi < \operatorname{Arg}(z) \leq \pi$ and $\sqrt{0} = 0$

$$\operatorname{arctanh}(z) = [\ln(1+z) - \ln(1-z)]/2$$

$$= -\operatorname{arctanh}(-z)$$

$$\operatorname{arctan}(z) = -i \operatorname{arctanh}(iz)$$

$$= -\operatorname{arctan}(-z)$$

$$\operatorname{arcsinh}(z) = \ln(z + \sqrt{1+z^2})$$

$$= -\operatorname{arcsinh}(-z)$$

$$\operatorname{arcsin}(z) = -i \operatorname{arcsinh}(iz)$$

$$= -\operatorname{arcsin}(-z)$$

$$\operatorname{arccos}(z) = \pi/2 - \operatorname{arcsin}(z)$$

$$\operatorname{arccosh}(z) = 2 \ln[\sqrt{(z+1)/2} + \sqrt{(z-1)/2}]$$

These are not intended as algorithms for computation, but as relations defining precisely the principal branch of each function.

Matrix Calculations

As mentioned earlier, the HP-15C can perform matrix addition, subtraction, and multiplication. It can also calculate determinants, invert square matrices, and solve systems of linear equations. In performing these last three operations, the HP-15C transforms a square matrix into a computationally convenient and mathematically equivalent form called the LU decomposition of that matrix.

LU Decomposition

The LU decomposition procedure factors a square matrix, say **A**, into a matrix product **LU**. **L** is a lower-triangular square matrix with 1s on its diagonal and with subdiagonal elements having values between -1 and 1, inclusive. **U** is an upper-triangular square matrix. The rows of matrix **A** may be permuted in the decomposition procedure. The possibly row-permuted matrix can be represented as the matrix




Fig. 8. The complex function $Z = w^2$, shown in (a) has an inverse function $w = \sqrt{Z}$, shown in (b), which maps the Z plane onto the right half plane of w with a branch cut along the negative real axis of the Z plane.

product \mathbf{PA} for some invertible matrix \mathbf{P} . The LU decomposition can then be represented by the matrix equation $\mathbf{PA} = \mathbf{LU}$ or $\mathbf{A} = \mathbf{P}^{-1}\mathbf{LU}$.

The HP-15C uses the Doolittle method with partial pivoting to construct the LU decomposition. It constructs the decomposition entirely within the result matrix. The upper-triangular part of \mathbf{U} and the subdiagonal part of \mathbf{L} are stored in the corresponding parts of the result matrix. It is not necessary to save the diagonal elements of \mathbf{L} since they are always equal to 1.

Partial pivoting is a strategy of row interchanging to reduce rounding errors in the decomposition. The row interchanges are recorded in the otherwise underused XS format fields of the result matrix's diagonal elements. The recorded row interchanges identify the result matrix as containing an LU decomposition and the result matrix's descriptor includes two dashes when displayed.

The determinant of the decomposed matrix \mathbf{A} is just $(-1)^r$ times the product of the diagonal elements of \mathbf{U} , where r is the number of row interchanges represented by \mathbf{P} . The HP-15C computes the signed product after decomposing the argument matrix \mathbf{A} into the result matrix.

The HP-15C calculates the inverse of the decomposed matrix using the relationship

$$\mathbf{A}^{-1} = [\mathbf{P}^{-1}\mathbf{LU}]^{-1} = \mathbf{U}^{-1}\mathbf{L}^{-1}\mathbf{P}$$

It does this by inverting both \mathbf{U} and \mathbf{L} , computing the product of their inverses, and then interchanging the columns of the product in the reverse order of the row interchanges of \mathbf{A} . This is all done within the result matrix.

Solving a system $\mathbf{AX}=\mathbf{B}$ for \mathbf{X} is equivalent to solving $\mathbf{LUX}=\mathbf{PB}$ for \mathbf{X} , where $\mathbf{PA}=\mathbf{LU}$ denotes the LU decomposition of \mathbf{A} . To solve this system, the HP-15C first decomposes the matrix \mathbf{A} in place. The calculator then solves the matrix equation $\mathbf{LY}=\mathbf{PB}$ for matrix \mathbf{Y} (forward substitution) and finally $\mathbf{UX}=\mathbf{Y}$ for matrix \mathbf{X} (backward substitution), placing the solution \mathbf{X} into the result matrix.

The LU decomposition is returned by a determinant or system solution calculation and can be used instead of the original matrix as the input to subsequent determinant, matrix inverse, or system solution calculations.

Norms and the Condition Number

A norm of a matrix \mathbf{A} , denoted by $\|\mathbf{A}\|$, is a matrix generalization of the absolute value of a real number or the magnitude of a complex number. Any norm satisfies the following properties:

- $\|\mathbf{A}\| \geq 0$ for any matrix and $\|\mathbf{A}\| = 0$ if and only if $\mathbf{A} = \mathbf{0}$
- $\|a\mathbf{A}\| = |a| \times \|\mathbf{A}\|$ for any number a and matrix \mathbf{A}
- $\|\mathbf{A} + \mathbf{B}\| \leq \|\mathbf{A}\| + \|\mathbf{B}\|$ for any matrices \mathbf{A} and \mathbf{B}
- $\|\mathbf{A}\mathbf{B}\| \leq \|\mathbf{A}\| \times \|\mathbf{B}\|$ for any matrices \mathbf{A} and \mathbf{B} .

One measure of the distance between two matrices \mathbf{A} and \mathbf{B} is the norm of their difference, $\|\mathbf{A}-\mathbf{B}\|$. A norm can also be used to define a condition number of a square matrix, which measures the sensitivity of matrix calculations to perturbations in the elements of that matrix.

The HP-15C provides three norms. The Frobenius norm of a matrix \mathbf{A} , denoted $\|\mathbf{A}\|_F$, is the square root of the sum of the squares of the matrix elements. This is a matrix generalization of the Euclidean length of a vector.

The HP-15C also provides the row (or row-sum) norm. The row norm of an m -by- n matrix \mathbf{A} is the largest row sum of absolute values of its elements and is denoted by $\|\mathbf{A}\|_R$:

$$\|\mathbf{A}\|_R = \max_{1 \leq i \leq m} \sum_{j=1}^n |a_{ij}|$$

The column (or column-sum) norm of a matrix \mathbf{A} is denoted by $\|\mathbf{A}\|_C$ and is the largest column sum of absolute values of its elements. It can be computed as the row norm of the transpose of the matrix \mathbf{A} .

For any choice of norm, a condition number $K(\mathbf{A})$ of a square matrix \mathbf{A} can be defined by

$$K(\mathbf{A}) = \|\mathbf{A}\| \times \|\mathbf{A}^{-1}\|$$

Then $K(\mathbf{A}) = \|\mathbf{A}\| \times \|\mathbf{A}^{-1}\| \geq \|\mathbf{AA}^{-1}\| = \|\mathbf{I}\| \geq 1$ for any norm. The following discussion assumes the condition number defined by the row norm. Similar statements can be made for the other norms.

If rounding or other errors are present in matrix elements, these errors will propagate through subsequent matrix calculations. They can be magnified significantly. Consider, for example, the matrix product \mathbf{AB} where \mathbf{A} is a square matrix. Suppose that \mathbf{A} is perturbed by the matrix $\Delta\mathbf{A}$. The relative size of this perturbation can be measured as $\|\Delta\mathbf{A}\| / \|\mathbf{A}\|$. The relative size of the resulting perturbation in the product is then

$$\begin{aligned} \|\Delta\mathbf{A}\mathbf{B}\| / \|\mathbf{A}\mathbf{B}\| &= \|\Delta\mathbf{A}\mathbf{A}^{-1}\mathbf{A}\mathbf{B}\| / \|\mathbf{A}\mathbf{B}\| \\ &\leq \|\Delta\mathbf{A}\mathbf{A}^{-1}\| \\ &\leq \|\Delta\mathbf{A}\| \times \|\mathbf{A}^{-1}\| \\ &= K(\mathbf{A}) \|\Delta\mathbf{A}\| / \|\mathbf{A}\| \end{aligned}$$

with equality for some choices of \mathbf{A} , \mathbf{B} , and $\Delta\mathbf{A}$. Hence $K(\mathbf{A})$ measures how much the relative uncertainty of a matrix can be magnified when propagated into a matrix product.

Uncertainties in the square system matrix \mathbf{A} or the matrix \mathbf{B} of the system of equations $\mathbf{AX}=\mathbf{B}$ will also propagate into the solution \mathbf{X} . For small relative uncertainties $\Delta\mathbf{A}$ in \mathbf{A} , say $\|\Delta\mathbf{A}\| / \|\mathbf{A}\| \ll 1/K(\mathbf{A})$, the condition number is a close approximation to how much the relative uncertainty in \mathbf{A} or \mathbf{B} can be magnified in the solution \mathbf{X} .⁶

A matrix is said to be ill-conditioned if its condition number is very large. We have seen that errors in the data—sometimes very small relative errors—can cause the solution of an ill-conditioned system to be quite different from the solution of the original system. In the same way, the inverse of a perturbed ill-conditioned matrix can be quite different from the inverse of the unperturbed matrix. But both differences are bounded by the condition number; they can be relatively large only if the condition number is large.

Singular and Nearly Singular Matrices

A large condition number also indicates that a matrix is relatively close to a singular matrix (determinant = 0). Suppose that A is a nonsingular matrix.

$$1/K(A) = \min (\|A-S\| / \|A\|)$$

$$\text{and } 1/\|A^{-1}\| = \min (\|A-S\|),$$

where each minimum is taken over all singular matrices S .⁶ $1/\|A^{-1}\|$ is the distance from A to the nearest singular matrix. $1/K(A)$ is this distance divided by the norm of A .

For example, if

$$A = \begin{bmatrix} 1 & 1 \\ 1 & 0.9999999999 \end{bmatrix}$$

then

$$A^{-1} = \begin{bmatrix} -9,999,999,999 & 10^{10} \\ 10^{10} & -10^{10} \end{bmatrix}$$

and $\|A^{-1}\| = 2 \times 10^{10}$. Therefore, there should exist a perturbation matrix ΔA with $\|\Delta A\| = 5 \times 10^{-11}$ that makes $A + \Delta A$ singular. Indeed,

$$\Delta A = \begin{bmatrix} 0 & -5 \times 10^{-11} \\ 0 & 5 \times 10^{-11} \end{bmatrix}$$

has $\|\Delta A\| = 5 \times 10^{-11}$, and

$$A + \Delta A = \begin{bmatrix} 1 & 0.99999999995 \\ 1 & 0.99999999995 \end{bmatrix}$$

is singular.

In principle, because the HP-15C's matrices are bounded in size, exact arithmetic and exact internal storage could be used to ensure 10-digit accuracy in matrix calculations. This was considered prohibitively expensive, however. Instead, the HP-15C is designed to perform arithmetic and store intermediate calculated values using a fixed number of digits.

Numerical determinant, matrix inversion, and system solution calculations using a fixed number of digits introduce rounding errors in their results. These rounding errors can be conceptually passed back to the input data and the calculated results interpreted as exact results for perturbed input data $A + \Delta A$. If the norm of the conceptual perturbation ΔA is comparable to $1/\|A^{-1}\|$, the original nonsingular input matrix A may be numerically indistinguishable from a singular matrix.

For example, a square matrix is singular if and only if at least one of the diagonal elements of U , the upper triangular matrix in the LU decomposition of A , is zero. But because the HP-15C performs calculations with only a finite number of digits, some singular and nearly singular matrices cannot be distinguished in this way.

The matrix

$$B = \begin{bmatrix} 3 & 3 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 1/3 & 1 \end{bmatrix} \begin{bmatrix} 3 & 3 \\ 0 & 0 \end{bmatrix} = LU$$

is singular. Using 10-digit accuracy, the calculated LU decomposition is

$$LU = \begin{bmatrix} 1 & 0 \\ 0.3333333333 & 1 \end{bmatrix} \begin{bmatrix} 3 & 3 \\ 0 & 10^{-10} \end{bmatrix}$$

which is the decomposition of the nonsingular matrix

$$D = \begin{bmatrix} 3 & 3 \\ 0.9999999999 & 1 \end{bmatrix}$$

Hence the calculated determinants of B and D are identical. On the other hand, the matrix

$$A = \begin{bmatrix} 3 & 3 \\ 1 & 0.9999999999 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 1/3 & 1 \end{bmatrix} \begin{bmatrix} 3 & 3 \\ 0 & -10^{-10} \end{bmatrix} = LU$$

is nonsingular. Using 10-digit accuracy, the calculated LU decomposition is

$$LU = \begin{bmatrix} 1 & 0 \\ 0.3333333333 & 1 \end{bmatrix} \begin{bmatrix} 3 & 3 \\ 0 & 0 \end{bmatrix}$$

which is the LU decomposition of the singular matrix

$$C = \begin{bmatrix} 3 & 3 \\ 0.9999999999 & 0.9999999999 \end{bmatrix}$$

The calculated determinants of A and C are also identical.

Because the calculated LU decompositions of some singular and nonsingular matrices are identical, any test for singularity based upon a calculated decomposition would be unreliable. Some singular matrices would fail the test and some nonsingular ones would pass it. Therefore, no such test is built into the HP-15C.

Instead, if a calculated diagonal element of U , which we call a pivot, is found to be zero during the LU decomposi-

tion, rather than aborting the matrix calculation and reporting the input matrix to be singular, the HP-15C replaces the zero pivot by a small positive number and continues with the calculation. This number is usually small compared to the rounding errors in the calculations. Specifically, it will be about 10^{-10} times the largest absolute value of any element in that column of the original matrix. If every element in that column of the original matrix has an absolute value less than 10^{-89} , the value 10^{-99} is used instead.

An advantage of replacing zero pivots by nonzero pivots is that matrix inversion and system solution calculations will not be interrupted by zero pivots. This is especially useful in applications such as calculating eigenvectors using the method of inverse iteration. Example programs calculating eigenvalues and eigenvectors can be found in reference 3.

The effect of rounding errors and possible intentional perturbations causes the calculated decomposition to have all nonzero pivots and to correspond to a nonsingular matrix usually identical to or negligibly different from the original matrix.

Complex Matrix Calculations

The HP-15C only operates on real matrices, that is, matrices with real elements. However, it is possible to represent complex matrices as real matrices and to perform matrix addition, subtraction, multiplication, and inversion of complex matrices and to solve complex systems of equations using these real representations.

Let $Z = X + iY$ denote a complex matrix with real part X and imaginary part Y , both real matrices. One way to represent Z as a real matrix is as the partitioned matrix

$$Z^P = \begin{bmatrix} X & \\ & Y \end{bmatrix}$$

having twice the number of rows but the same number of columns as Z . Complex matrices can be added or subtracted by adding and subtracting such real representations.

Another computationally useful real representation for Z is

$$\tilde{Z} = \begin{bmatrix} X & -Y \\ Y & X \end{bmatrix}$$

having twice the number of rows, and columns as Z . The HP-15C's built-in matrix operation **MATRIX 2** performs the transformation

$$Z^P \rightarrow \tilde{Z}$$

The operation **MATRIX 3** performs the inverse transformation

$$\tilde{Z} \rightarrow Z^P$$

Suppose A , B , and C are complex matrices and A is

invertible. Then complex matrix multiplication, inversion, and system solution can be performed with real matrices and built-in HP-15C operations using the relationships:

$$(AB)^P = \tilde{A}B^P,$$

$$(\tilde{A}^{-1}) = (\tilde{A})^{-1},$$

$$AC = B \rightarrow C^P = (\tilde{A})^{-1} B^P.$$

These procedures are illustrated in the *HP-15C Owner's Handbook*.

Matrix Transpose

The operations **MATRIX 2** and **MATRIX 3** perform their transformations using a matrix transpose routine. The rows and columns of a matrix are interchanged to form the transpose of that matrix. The transformation is performed in place, replacing the original matrix by its transpose. This routine is available to the user as **MATRIX 4**. Consider the following example:

$$\begin{bmatrix} a & b & c \\ d & e & f \end{bmatrix} \rightarrow \begin{bmatrix} a & d \\ b & e \\ c & f \end{bmatrix}$$

Here the elements of the matrices have been displayed in a two-dimensional format. However, they are stored in a one-dimensional sequence within the calculator's memory. For this example, the transpose operation changes the ordering of the elements within the calculator memory as

$$a b c d e f \rightarrow a d b e c f.$$

The **MATRIX 4** operation moves the elements according to



These movements form disjoint loops. The first value in the sequence is the first candidate for moving. As a value is copied into its destination, that destination is tagged in its XS field. The previous value at that location is the next candidate for moving. Movement along a loop continues until a destination is encountered that is already tagged. The content of the tagged destination is not changed and the current loop is terminated. The value in the location immediately following that tagged destination is the next candidate for moving.

This operation continues moving values along loops until the sequence is exhausted, at which point all destination tags are removed. Finally, the recorded dimensions of the matrix are switched.

Accuracy of Matrix Calculations

Accuracy specifications for all matrix operations are given in reference 3. These specifications are stated in terms of both backward and forward error analysis. Reference 3 includes a general rule of thumb for the number of significant digits in a calculated matrix inverse or system solution. It also includes descriptions of techniques to improve upon the accuracy of calculated system solutions and to reduce the ill-conditioning of systems of equations.

Acknowledgments

Numerous individuals made valuable contributions to the HP-15C software effort. As the software project manager, Rich Carone helped formulate some of the original design concepts and kept the software effort on track. Diana Roy, Robert Barkan, and Hank Schroeder wrote the *HP-15C Owner's Handbook*. We would like to give special thanks to Professor William Kahan, who contributed many design ideas, provided strong guidance in developing the mathematical algorithms, and wrote a portion of the *HP-15C Advanced Functions Handbook*. His unbounded enthusiasm for the product helped keep us going, especially when we still had features to implement and no ROM space left.

References

1. W.M. Kahan, "Personal Calculator Has Key to Solve Any Equation $f(x)=0$," *Hewlett-Packard Journal*, Vol. 30, no. 12, December 1979.
2. W.M. Kahan, "Handheld Calculator Evaluates Integrals," *Hewlett-Packard Journal*, Vol. 31, no. 8, August 1980.
3. *HP-15C Advanced Functions Handbook*.
4. D.W. Harms, "The New Accuracy: Making $2^3=8$," *Hewlett-Packard Journal*, Vol. 28, no. 3, November 1976, p. 16.
5. W.M. Kahan, "Branch Cuts for Complex Elementary Functions," working document for IEEE Floating-Point Standards Subcommittee, September 9, 1982.
6. E. Atkinson, *An Introduction to Numerical Analysis*, John Wiley & Sons, 1978, pp. 461-463.



Paul J. McClellan

Paul McClellan started working part-time at HP in 1979 while he was a graduate student at Oregon State University. He has a BS degree in physics and mathematics awarded by the University of Oregon in 1974 and an MS degree in statistics from Oregon State. He plans to complete the requirements for the PhD degree in statistics this year. Paul began full-time work at HP in late 1982 and worked on the **SOLVE** and matrix computation routines for the HP-15C. He is a member of the American Statistical Association and coauthor of the *HP-15C Advanced*

Functions Handbook. When not busy with work or his studies, Paul enjoys rock climbing, mountaineering, cross-country skiing, and visiting Portland, Oregon. He is single and lives in Corvallis, Oregon.



Joseph P. Tanzini

Joe Tanzini was born in Trenton, New Jersey and attended Trenton State College, receiving a BA degree in mathematics in 1973. He continued his mathematics studies at Lehigh University and earned the MS and PhD degrees in 1975 and 1982. Joe started at HP as a software engineer in 1980 and wrote firmware for calculators, including coding the integrate algorithm and complex functions for the HP-15C. He also coauthored the *HP-15C Advanced Functions Handbook*. Joe enjoys bicycling and walking during his leisure time. He is married, has two daughters, and lives in Corvallis, Oregon.

A Pocket Calculator for Computer Science Professionals

This compact, yet powerful pocket calculator is designed for technical professionals working in computer science and digital electronics. Boolean operations and bit manipulation are some of its capabilities.

by Eric A. Evett

LOGIC DESIGN and computer programming require mathematical operations not ordinarily provided by small calculators. A large amount of tedious paperwork is often required to convert among number bases, perform logic operations, shift and rotate bits in a word, or check processor instruction flow. To simplify such work, Hewlett-Packard recently introduced a programmable pocket calculator especially designed for people who deal with bits. The HP-16C (Fig. 1), like other HP calculators, uses a reverse-Polish-notation (RPN) system and provides standard floating-point decimal arithmetic (including square root). Its novel capabilities become apparent, however, when the HP-16C is switched into the integer mode. Only integers are allowed in this mode, and they can be keyed in and displayed in either hexadecimal, octal, binary, or decimal format. In this mode, number base conversion, integer arithmetic, logical operations, and bit manipulations can be done.

Integer Mode

In the integer mode, all numbers are represented internally in binary form. The word size is selected by the user and can range from 1 to 64 bits. The user also can select whether the numbers are to be interpreted as one's complement, two's complement, or unsigned integers. In the unsigned integer mode with a 64-bit word size, numbers up to $2^{64}-1$ (18,446,744,073,709,551,615) can be represented. Although the HP-16C normally displays the eight least-significant digits of a number, a scrolling capability is provided to display higher-order digits.

Programming

In addition to the four-register RPN stack, 203 bytes of user memory are available for storing program steps and use as storage registers. When the program memory is cleared, all 203 bytes are allocated to storage registers. The number of storage registers available depends on the selected word



Fig. 1. The HP-16C Programmable Calculator is designed for computer science and digital electronics applications. Besides the normal four-function calculator features, it has a number of capabilities for setting number bases and word sizes, performing Boolean operations, and manipulating bits.

Real (Floating-Point) Format

Real numbers are represented in the HP 3000 memory by 32 bits (two consecutive 16-bit words) separated into three fields. These fields are the sign, the exponent, and the mantissa. The format is known as excess 256. Thus, a real number consists of (see Fig. 1):

- Sign (S), bit 0 of the first 16-bit word. Positive=0, negative=1. A value X and its negative $-X$ differ only in the value of the sign bit.
- Exponent (E), bits 1 through 9 of the first 16-bit word. The exponent ranges from 0 to 777 octal (511 decimal). This number represents a binary exponent biased by 400 octal (256 decimal). The true exponent, therefore is $E - 256$; it ranges from -256 to $+255$.
- Fraction (F), a binary number of the form $1.xxx$, where xxx represents 22 bits stored in bits 10 through 15 of the first 16-bit word and all bits of the second 16-bit word. Note that the 1 is not actually stored, there is an assumed 1 to the left of the

binary point. Floating-point zero is the only exception. It is represented by all 32 bits being zero.

The range of nonzero real values for this format is from 0.863617×10^{-77} to 0.1157920×10^{78} . The formula for computing the decimal value of a floating-point representation is: Decimal value = $(-1)^S \times 2^{E-256} \times F$.




Fig. 1. Diagram of real (floating-point) format used in the HP 3000.

LBL	A	MASKL	Create mask of 23 bits, left-justified.
2'S	Set two's complement mode.	AND	Extract upper 23 bits.
OCT	Convert to octal integer mode, and return integers y and x such that $2^x y = \text{original input}$.	F?	4 Did round cause a carry-out of most significant bit?
SF	3 Leading zeros will be displayed.	ISZ	If yes, increment exponent.
X=Y	Was input 0?	SL	Shift off implied 1 bit.
GTO	1 If yes, then branch to Label 1.	RCL	1 Recall biased exponent.
4		OR	Concatenate exponent to fraction part.
3		F?	0 Is mantissa sign to be positive?
7		GTO	1 If yes, branch to Label 1.
+	Bias exponent. $287 = 256 + 31$	1	
STO	1 Store biased exponent in index register.	1	
X≥Y	Swap exponent and mantissa.	SB	Set the sign bit.
SF	0 Set flag 0.	LBL	1
X<0	Mantissa negative?	1	
CF	0 If yes, clear flag 0.	2	
ABS	Absolute value of mantissa.	RRn	Rotate sign, exponent, and fraction to proper position.
4		STO	0 Store the 32-bit result in register 0.
0		2	
WSIZE	Set word size to 32.	0	
4		WSIZE	Change word size to 16 bits.
0		RCL	1 Recall 16-bit word 1.
0		RCL	0 Recall 16-bit word 2.
+	Round mantissa to 23 bits.	RTN	
2			
7			

Fig. 2. Outline of HP-16C subroutine to convert numbers given in the HP 3000 Computer's real format to decimal floating-point format.

Using the HP-16C

Listing the features of a programmable calculator rarely provides a complete picture of its capabilities. Examples of the application of the calculator's features are often required to demonstrate to the user what can be done and why a particular feature is useful. Hence, several examples of the use of the HP-16C are given below.

Add with Carry

The HP-16C can be programmed to simulate instructions commonly found in commercial processors. The following subroutine performs an add with carry ($Y+X+C \rightarrow X$). It adds the numbers in registers X and Y along with the carry bit indicated by the state of flag 4 and returns the result in the X register. The carry flag is set (indicated by the C annunciator in the display) if there is a carry-out of the most significant bit of the result.

001	LBL A	Labels subroutine
002	0	
003	RLC	Generates 0 or 1 depending on carry flag
004	+	Adds carry to second operand
005	CF 0	
006	F? 4	Copies carry flag 4 to flag 0
007	SF 0	
008	+	Adds first operand to the total
009	F? 0	
010	SF 4	Sets carry flag if first add carried
011	RTN	

To use this routine, enter the two operands in registers Y and X, and press **GSB A**.

Example:

2'S	Set two's complement mode
HEX	Set number base to hexadecimal
8	
WSIZE	Set word size to 8
CF 4	Clear carry flag
FE	First operand
ENTER	Enter first operand into Y register
72	Second operand
GSB A	Displays 70 (FE+72+0) with carry set (C annunciator on)

Bit Extraction

The following subroutine extracts a field from a bit pattern. The field is specified by the bit numbers of the pattern corresponding to the lowest-order and highest-order bits of the field. The least-significant bit of the bit pattern is bit number 0. Hence, the result in the X register is the bits of the pattern in the Z register from the bit number in the Y register to the bit number in the X register, inclusive.

001	LBL B	Labels subroutine
002	R↓	Bring down value in Y register
003	RRn	Right-justifies field
004	R↑	Raise stack
005	LST x	Recall Y value
006	-	Subtract Y from X
007	1	
008	+	Computes number of bits in field
009	MASKR	Creates mask same width as field
010	AND	Extracts field
011	HEX	Exits in the hexadecimal mode
012	RTN	

size. When the word size is eight bits, 203 registers are available; a 16-bit word size results in 101 available registers, and so on. Each programmable instruction takes one byte of memory. As program steps are inserted, the number of available storage registers decreases. A program can have up to 203 steps if no storage registers are required.

Editing capabilities to make program development easier include insert, delete, back-step, single-step, and go-to-line-number operations. The user may single-step through program execution to help debug programs. Other programming features include label addressing (sixteen labels), subroutines (up to four levels deep), conditional tests, branching, and six user flags.

These flags can be set, cleared, and tested under program control. Three of the flags are special. Leading zero digits in a word are suppressed in the display unless flag 3 is set. Flag 4 is the carry flag, and flag 5 is the overflow flag. Two annunciators in the display (C for carry and G for > largest representable number) give a visual indication of the state of flags 4 and 5, respectively. The overflow flag is set if the true result of an operation cannot be represented in the selected word size and complement mode. The carry flag is set under various conditions, depending on the operation. For example, addition sets the carry flag if there is a carry-out of the most significant bit; otherwise the carry is cleared (see box above for examples). The shift-left instruction sets the carry if a 1 bit is shifted

off the left end of the word; otherwise the carry is cleared.

Logic Operations

The rich selection of bit manipulation and logical operations, along with user-selectable complement mode and word size, make the HP-16C a flexible logic and program design tool. Programs can be written to simulate individual instructions commonly found on commercial processors, to extract a field from a bit pattern, or to convert from one numeric format to another.

A common problem is the conversion between the internal binary floating-point format of a particular machine and decimal floating-point format. The HP-16C provides a feature that can be used to great advantage in programs designed to perform such conversions. This feature provides a mode for performing standard decimal floating-point calculations. Upon switching from the integer mode to decimal floating-point mode (by using the **FLOAT** function), the integers y and x in the Y and X stack registers are converted to the floating-point equivalent of $2^x y$, which is then placed in the X register and displayed. Converting back to integer mode (by pressing the **HEX**, **DEC**, **OCT**, or **BIN** keys), causes the contents of the X register to be converted to a pair of integers y and x such that y is a 32-bit integer ($2^{31} \leq |y| < 2^{32}$ unless y=0) and $2^x y$ is equal to the value in the X register before mode conversion. The integers y and x are then placed in the Y and X registers.

To use this routine, the user places the bit pattern in register Z, the number of the lowest-order bit in the field in register Y, and the number of the highest-order bit in the field in register X. The user then presses **GSB B**.

Example: Extract bits 2 through 5 from 39_{16} (00111001).

```

8
WSIZE      Set wordsize to 8
HEX        Set hexadecimal mode
39         Bit pattern
ENTER
2          Lowest-order bit
ENTER
5          Highest-order bit
GSB B     Displays E (1110) as result.
  
```

Conversion Between Binary and Gray Code

Gray code has the property that only one bit changes between the representations of any two adjacent numbers. If the word size is n bits, then binary-to-Gray-code conversion is given by

$$\begin{aligned}
 G_0 &= B_0 \text{ XOR } B_1 \\
 G_1 &= B_1 \text{ XOR } B_2 \\
 &\vdots \\
 G_{n-2} &= B_{n-2} \text{ XOR } B_{n-1} \\
 G_{n-1} &= B_{n-1}
 \end{aligned}$$

where G is the Gray code number, B is the binary number, and subscript 0 indicates the least-significant bit of G and B , subscript 1 indicates the next least-significant bit, and so forth.

The Gray-code-to-binary conversion is given by

$$\begin{aligned}
 B_0 &= G_0 \text{ XOR } G_1 \text{ XOR } \dots \text{ XOR } G_{n-1} \\
 B_1 &= G_1 \text{ XOR } G_2 \text{ XOR } \dots \text{ XOR } G_{n-1} \\
 &\vdots \\
 B_{n-2} &= G_{n-2} \text{ XOR } G_{n-1} \\
 B_{n-1} &= G_{n-1}
 \end{aligned}$$

Binary-to-Gray-code subroutine:

```

001  LBL C
002  ENTER  Copies binary number to Y register
003  SR     Shifts binary number in X register to
           the right
004  XOR    Computes Gray-code equivalent
005  RTN
  
```

Gray-to-binary-code subroutine:

```

001  LBL D
002  ENTER  Copies Gray code number to Y register
003  LBL 2
004  SR     Shift Gray code number in X register to
           the right
005  XOR    Exclusive OR operation
006  LST x  Recall previous number
007  X≠0    Loop until Gray code number is 0
008  GTO 2
009  R↓
010  RTN
  
```

To use these routines, the user sets the HP-16C to the binary mode by pressing **BIN**, places the number in the X register and presses **GSB C** for binary-to-Gray or **GSB D** for Gray-to-binary-code conversions.

LBL B		XOR	Extract biased exponent part.
HEX	Set hexadecimal base mode.	LST x	Recover fraction part.
2'S	Set two's complement mode.	1	
2		7	
0		SB	Set bit 23 (the implied 1-bit). Bit 0 is the least significant.
WSIZE	Set word size to 32 bits.	F? 4	Test carry flag. Was sign bit negative?
X↔Y	Swap word 1 and word 2.	CHS	If yes, complement mantissa.
1		ASR	Arithmetic shift right mantissa 1 bit.
0		X≥Y	Swap mantissa and exponent part.
RLn	Shift word 1 16 bits to left.	9	
OR	Concatenate word 2 to word 1.	RLn	Rotate exponent part 9 bits left, placing it at right end.
SL	Shift sign bit left into carry flag.	1	
ENTER		1	
X=0	Is input 0?	6	
GTO 0	If yes, branch to Label 0.	-	Unbias exponent. $E-278=E-256-32$
1		LBL 0	
7		FLOAT	Compute 2^y
MASKR	Create mask of 23 bits, right-justified.	RTN	
AND	Extract fraction part.		

Fig. 3. Outline of HP-16C subroutine to convert decimal floating-point numbers into the format used by HP 3000 Computers.

The subroutines listed in Fig. 2 and Fig. 3 convert between the HP 3000 Computer's FORTRAN real (floating-point) format² and decimal floating-point format (see box on page 37). Because the HP-16C views bit 0 as the least significant bit of a word and the HP 3000 views it as the most significant bit, some of the steps listed in Fig. 2 and Fig. 3 are used to convert between these two opposing views.

To use these programs after they are entered in the HP-16C, a user performs the following steps.

- HP 3000 to decimal:
 1. Select octal base (**OCT**).
 2. Enter word 1 in the Y register and word 2 in the X register.
 3. Execute **GSB B**. Answer is displayed.
 4. Repeat steps 1, 2, and 3 for each new conversion.
- Decimal to HP 3000:
 1. Select decimal floating-point mode (**FLOAT 4**).
 2. Enter number in the X register.
 3. Execute **GSB A**. Word 1 is placed in the Y register, word 2 in the X register.
 4. Repeat steps 1, 2, and 3 for each new conversion.

Acknowledgments

Several people made significant contributions to the HP-16C software effort. John Van Boxtel developed many of the fundamental design concepts. Rich Carone also contributed to the software design and coded the complex

routines that format and build the display. Stan Blascow did a large portion of the software testing and Diana Roy wrote the owner's handbook.

References

1. M.E. Sloan, *Computer Hardware and Organization*, Science Research Associates, Inc., 1976, pp 95-97.
2. *FORTRAN Reference Manual*, HP 3000 Computer Systems, Hewlett-Packard, Santa Clara, 1978, Section II.

Eric A. Evett



Eric Evett received the BS and MS degrees in mathematics from the University of Arizona in 1970 and 1972. He taught mathematics there until 1975 and then spent three years developing software for calculators before joining HP in late 1978. Eric wrote portions of the microcode for the HP-11C, HP-15C, and HP-16C Calculators and now is an R&D software project manager at HP's Corvallis, Oregon facility. He was born in Colfax, Washington, and now lives in Corvallis. He is married and has two sons. His outside interests include jogging, hiking, basketball, reading, movies, and tennis (he played on his college's tennis team which was then ranked 7th in the U.S.A.).

CORRECTION

In the March issue, the Pascal statements at the top of the back page were printed in the wrong order. Here is the correct version.

```
buffer[w]:=getch;
c:=c+a[buffer[w]];
w:=w+1;
```

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

HEWLETT-PACKARD JOURNAL

MAY 1983 Volume 34 • Number 5

Technical information from the Laboratories of
Hewlett-Packard Company

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

Hewlett-Packard Central Mailing Department
Van Heuven Goedhartlaan 121

1181 KK Amstelveen, The Netherlands

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

Hewlett-Packard (Canada) Ltd.

6877 Goreway Drive, Mississauga, Ontario L4V 1M8 Canada

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

02000207078 &&BLAC&CA00
MR C A BLACKBURN
JOHN HOPKINS UNIV
APPLIED PHYSICS LAB
JOHNS HOPKINS RD
LAUREL MD 20707

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