HEWLETT-PACKARD
JOURNAL

APRIL 1991

HEWLETT-PACKARD JOURNAL

April 1991 Volume 42 • Number 2

Articles

© Hewlett-Packard Company 1991 Printed in Holland

Departments

In this Issue
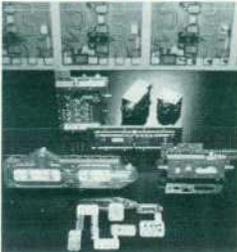
The microwave and millimeter-wave portions of the electromagnetic spectrum encompass frequencies from a few hundred megahertz to about 300 gigahertz. The region is mainly the domain of radar and communications systems. A basic piece of test equipment for such systems and the components that go into them is the sweep oscillator, or sweeper. A sweeper generates a constant-amplitude signal that varies in frequency at a constant rate, that is, it "sweeps" over a user-selected range of frequencies. In the article on page 6, we're introduced to the HP 8360 family of synthesized sweepers—eleven models, some of which generate frequencies as high as 50 GHz, well into the millimeter-wave range. That article includes a brief history of microwave sweepers and the basic problem of combining smooth sweeping with high frequency accuracy. Before the HP 8360, the best solution to this problem was the "lock and roll" technique of the HP 8340 synthesized sweeper, which provides a highly accurate, phase-locked start frequency followed by an unlocked sweep. By adding an end-of sweep calibration to this technique, the HP 8360 improves swept frequency accuracy tenfold. The design of the HP 8360 family is characterized by extensive use of advanced microwave technologies. There are microwave monolithic integrated circuits, thick-film and thin-film hybrids, complex microcircuit assemblies, new connection and packaging techniques, and many contributions in amplifier, filter, mixer, modulator, doubler, and attenuator design. Two complex microcircuits called the modsplitter and the low-band microcircuit exemplify many of the advanced design concepts (see page 36). Other state-of-the-art components include a switched frequency doubler with low spurious outputs (page 31), an optional high-speed pulse modulator that provides two-nanosecond pulse rise and fall times and a 95-dB on-off ratio (page 34), and dc-to-50-GHz programmable step attenuators (page 47). A new automatic level control system (see page 24) offers greatly improved leveling performance over previous generations of microwave sources, along with user flatness correction and self-calibration capabilities. The HP 8360 family has a menu-based user interface, as described in the article on page 6, and many built-in self-test and adjustment features, as described in the article on page 17.

Many radar, communications, and spectroscopic applications already exist at frequencies well above the 50-GHz upper frequency specification of the HP 8360 series. To meet the testing needs of these applications, frequencies up to 110 GHz can be generated by teaming an HP microwave sweeper like the HP 8360 with one of the millimeter-wave source modules described in the article on page 50. The HP 83557A covers V band, 50 to 75 GHz, and the HP 83558A covers W band, 75 to 110 GHz. These modules multiply the microwave sweeper's output frequency by 4 and 6, respectively. After a brief introduction to the millimeter-wave characteristics of the atmosphere, the article tells why it's better to multiply by 2, amplify, and multiply again by 2 or 3 than simply to multiply by 4 or 6 directly. It then describes the design of an R-band amplifier doubler, a V-band amplifier doubler, a W-band amplifier tripler, and a coupler/detector using advanced microcircuit technologies and tools.

Cellular telephone and radio services, according to the article on page 65, have become so popular that the currently allocated frequency bands have become saturated. A new digital cellular standard has been developed that will allow more users in a given portion of the spectrum. Because the new standard requires radios to conform to the old analog specification as well as the new digital system, it is called the North American Dual-Mode Cellular System (yes, there's a different standard in Europe). The new radios will require testing for both operating modes: analog and digital. The analog tests can be done as they are now. Two new HP products will aid the digital tests. The HP 11846A modulation generator (page 65) uses a simple but effective design approach to generate the complex modulation format used by the dual-mode system, while the HP 11847A modulation measurement software (page 73), using digital signal processing techniques, accurately verifies the performance of digital cellular transmitters. The HP 11847A software provides many valuable optional measurements and graphical output modes in addition to those required by the standard.

In software testing, it's important that a test exercise all of the code so that any defects present will have a chance of being exposed. One way to measure the effectiveness of software testing is branch analysis, which seeks to measure what percentage of the branches in a program have executed at least once during a test (branches are the code paths after the decision points in a program). The HP Branch Validator (page 83) is a software execution and test verification tool that instruments programs written in the C language and determines which branches have been executed. It runs on HP 9000 computers, either as a stand-alone application or as a member of the SoftBench environment described in our June 1990 issue.

R.P. Dolan
Editor

Cover:

Representatives of some of the state-of-the-art technologies that are used in the HP 8360 family of synthesized sweepers and the HP 83557A and HP 83558A millimeter-wave sources. In the background is a photomicrograph of a microwave monolithic integrated circuit (MMIC)—a traveling wave amplifier. In the foreground are examples of thick-film and thin-film microwave hybrid microcircuits, waveguide components, and various amplifier, multiplier, and modulator microcircuit assemblies.

What's Ahead

The cover subject of the June 1991 issue will be the HP 48SX scientific expandable calculator. Also featured will be the HP 3588A, a 10-Hz-to-150-MHz spectrum analyzer that uses digital signal processing techniques to speed up swept spectrum measurements by a factor of four or more. There will be three papers from the 1990 HP Software Engineering Productivity Conference—two on software configuration management and one on an integrated computing environment. We'll also have an article on HP GlancePlus, a computer system performance diagnostic tool, and a paper on methods of improving the product development process.

A Family of High-Performance Synthesized Sweepers

Eleven models offer frequency coverage to 50 GHz in coax, extendable to 110 GHz in waveguide with millimeter heads. Swept frequency accuracy is ten times better than previous designs. A menu-based user interface simplifies operation.

by Roger P. Oblad, John R. Regazzi, and James E. Bossaller

MICROWAVE FREQUENCIES came into general interest with the advent of radar at the begining of World War II in 1939. Signal generators were needed to test and develop these early radar systems. The first commercial microwave sources used mechanically tuned klystrons. The frequency accuracy was only as good as the markings on manually operated dials. In 1956, the first HP microwave sweep oscillator, the HP 670A, was produced by coupling a motor to the klystron drive mechanism. These early sweepers brought advances in measurement technology but were very slow and inaccurate. In 1958, sweeper performance improved considerably with the introduction of the electromagnet-tuned backward wave oscillator (BWO). Over the next three decades, the sweeper improved further with the use of YIG oscillators and microprocessor controls to provide many convenient user features.

The first microwave synthesizer was introduced in the early 1970s. With the advent of the synthesizer, it was possible to generate extremely stable and accurate CW signals.[1] Over the next ten years, synthesizers were continually improved in parallel with the improvements made on sweepers.

It still was not possible to have both high frequency accuracy and sweep in the same product. Various schemes were used to achieve better accuracy in sweepers. Sweepers were locked to microwave counters to achieve frequency accuracy equal to that of a counter. Crystal markers marked precise frequency points in otherwise inaccurate sweeps.

Improvements were made to synthesizers by adding automatic step-sweep features. This provided accurate sweeps but with undesirable step discontinuities and slow sweep rates.

The first true marriage of the features and benefits of both the microwave synthesizer and the analog sweeper was the HP 8340A synthesized sweeper, which was introduced in 1982. This was a major advance in microwave source technology. The HP 8340A used a concept called "lock and roll," which was developed for the HP 8566A microwave spectrum analyzer.[2,3] The synthesized sweeper contains a high-performance microwave synthesizer, which is phase-locked to the start frequency of an analog sweep. After the phase-lock operation, the frequency control circuits are unlocked and an analog sample-and-hold circuit maintains the start frequency. Analog sweep control circuits push the frequency over the predetermined sweep range. This mode of operation gives the start frequency synthesizer accuracy. The accuracy of the rest of the sweep is determined by the precision of the analog frequency control circuits.

New Sweeper Family

The next-generation synthesized sweeper is the HP 8360 family, introduced in 1989 (see Fig. 1). This family of products improves the accuracy of the earlier lock and roll scheme by adding an automatic end-of-sweep calibration. The internal synthesizer accurately measures the exact frequency at the end of a sweep. This data is fed back to

Fig. 1. *HP 8360 synthesized sweep oscillators include eleven models that cover the frequency range from 10 MHz to 50 GHz. Millimeter heads extend the frequency range to 110 GHz (Model 83640A shown). High-power models provide 17 dBm at 20 GHz. Analyzer models have special network analyzer synchronization features.*

frequency control circuits, which correct the sweep inaccuracies. This automatic calibration feature has resulted in a tenfold improvement in swept frequency accuracy.

A major design objective for the new sweeper family was orthogonality of the optional features. This means that optional features are independent and can be chosen without compromising other aspects of performance or functionality. This goal required much attention in the design and did result in some added cost and complexity. For example, the modulation features—AM, FM, and pulse—can be used in any combination. These features work in both CW mode and swept mode. This goal also had a significant impact on the user interface, both front-panel and HP-IB, as described later in this article.

From a performance point of view, the three biggest design goals were higher power, larger frequency range, and lower harmonics. Other issues such as frequency switching time, modulation features, power flatness compensation, and an extensive effort to design for quality and low overall cost of ownership rounded out the project objectives.

The HP 8360 family of high-end microwave synthesized sweepers currently includes the following models:

- HP 83620A: 10 MHz to 20 GHz
- HP 83622A: 2 to 20 GHz
- HP 83623A: 10 Mhz to 20 GHz High Power
- HP 83624A: 2 to 20 GHz High Power
- HP 83640A: 10 MHz to 40 GHz
- HP 83642A: 2 to 40 GHz
- HP 83621A: 45 MHz to 20 GHz Network Analyzer Model
- HP 83631A: 45 MHz to 26.5 GHz Network Analyzer Model
- HP 83650A: 10 MHz to 50 GHz
- HP 83630A: 10 MHz to 26.5 GHz
- HP 83651A: 45 MHz to 50 GHz Network Analyzer Model.

All of these models have coaxial outputs. With the addition of the V-band and W-band millimeter-wave heads (see article, page 50), the frequency is extended to 110 GHz in waveguide. A high-power option provides 17 dBm at 20 GHz.

Improvements in the synthesizer and the associated firmware have resulted in <5-ms switching times in step-sweep mode. This, coupled with new synchronization interface hardware, provides significant improvements in network analyzer performance. This new interface hardware is called the "bucket pulse interface." Messages pass between the synthesized sweeper and the network analyzer on two hardware lines. This allows the two instruments to synchronize the exact points where data is taken with the synthesizer's internal sweep calibration and control circuitry. The result is a tenfold improvement in swept frequency accuracy over what was achievable before.

Other new features include deep AM and a high-performance precision leveling system (see article, page 24). A simple internal pulse generator is standard, and there is an optional high-performance internal modulation generator. There is also a high-performance pulse modulator option, discussed in the article on page 34.

Block Diagram

The overall HP 8360 block diagram, Fig. 2, contains seven blocks:
- Microwave components

- Synthesizer
- Modulation, sweep control, and microcircuit control circuits
- Frequency standard
- User interface
- CPU
- Power supply.

The microwave components generate and process all the high-frequency signals. These are all located on the RF deck except the sampler, which is located inside the synthesizer.

The synthesizer block (Fig. 3) implements the HP 8360 indirect synthesis scheme. This scheme uses four key components: a 2-to-7.8-GHz YIG-tuned oscillator (YTO) located on the RF deck, a sampler covering the same frequency range, and two phase-locked loops called the sampler loop and the fractional-N loop. A portion of the microwave output from the YTO feeds the sampler. Inside the sampler microcircuit, this YTO signal mixes with a high-order harmonic of a 200-to-220-MHz VCO. This VCO is part of the sampler phase-locked loop. The resultant mixing product is an IF signal in the 20-to-40-MHz range. A phase detector compares this IF signal with the output of the fractional-N synthesis loop. The phase detector output is summed into the tuning control for the original 2-to-7.8-GHz YIG oscillator. This achieves the overall phase lock for the synthesizer.

The sampler loop can be stepped in 500-kHz steps. The fractional-N loop can be stepped in 0.001-Hz steps. This ultrafine resolution makes it possible for the HP 8360 to maintain its 1-Hz frequency resolution even at 110 GHz, which is achieved by multiplying a 6.11-GHz YIG oscillator frequency by 18. The modulation, sweep control, and microcircuit control circuits reside on several of the card-cage printed circuit boards. These circuits implement the power level and flatness control features and provide all the microcircuit control signals. Sweep control circuits provide the analog sweep feature and compensate for the delay and tracking characteristics of magnetically tuned microcircuits.

Both the sampler and the fractional-N phase-locked loops derive their internal reference signals from a precision oven-stabilized 10-MHz crystal frequency standard.
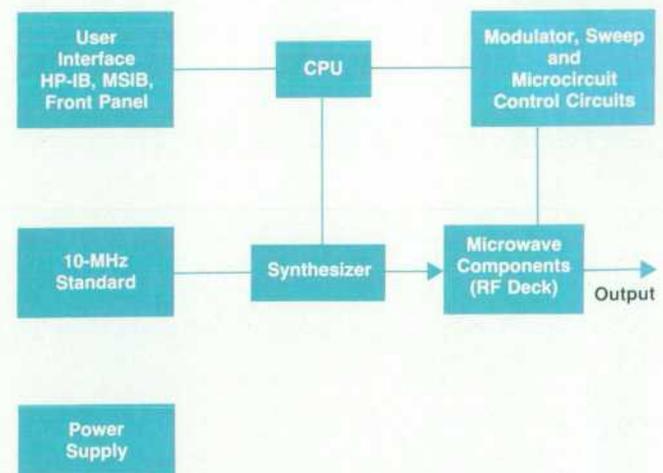
Fig. 2. Basic block diagram of an HP 8360 synthesized sweeper.
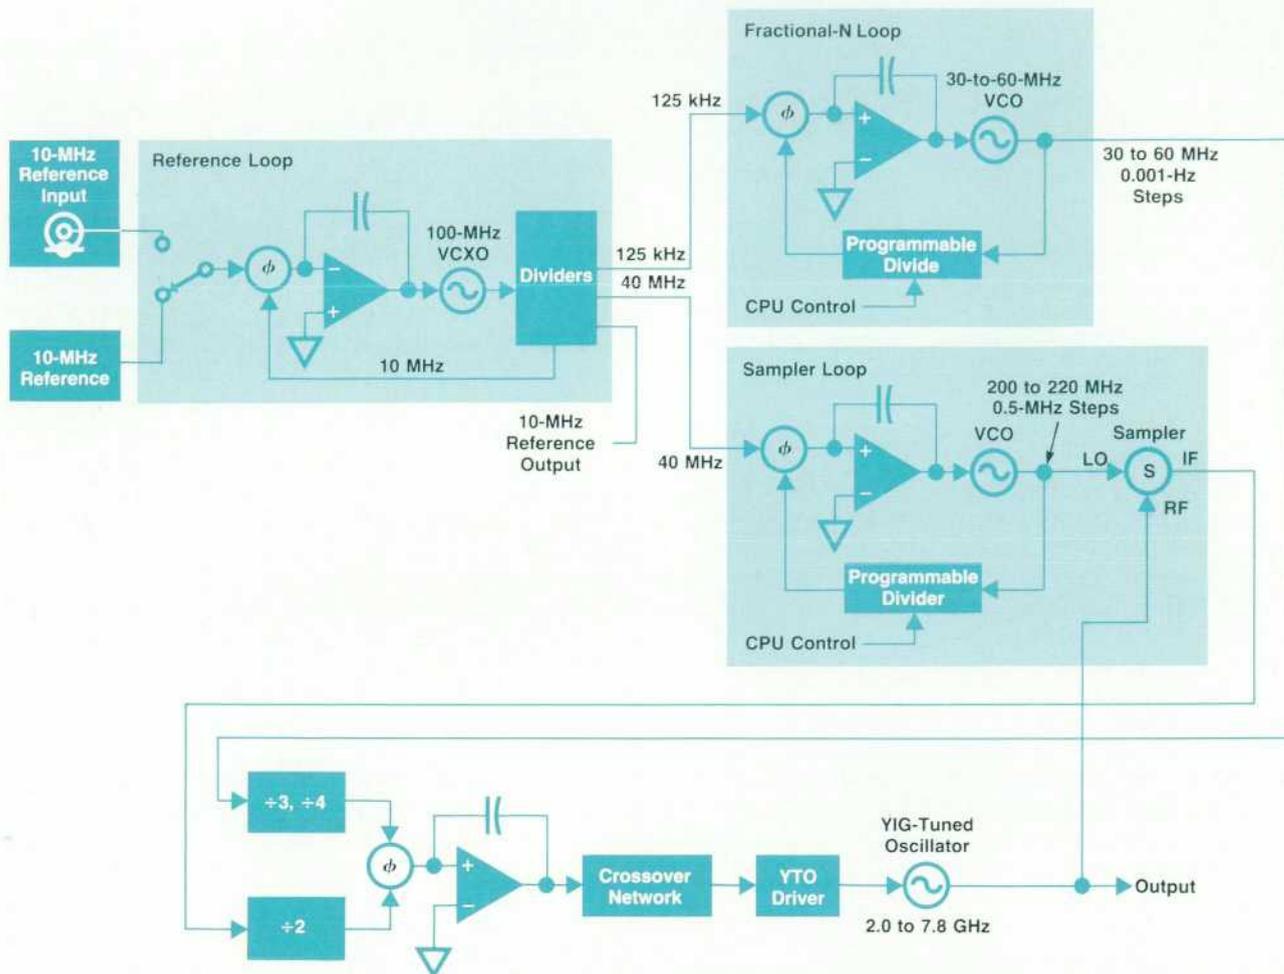
Fig. 3. *HP 8360 frequency synthesizer block diagram. The phase-locked loops are locked at the beginning of a sweep so the start frequency has synthesizer accuracy. The loops are unlocked for continuous sweeping. An end-of-sweep calibration with synthesizer accuracy provides a tenfold improvement in overall swept frequency accuracy over previous designs.*

The user interface includes the front panel, the rear panel, the HP-IB hardware and firmware on the CPU board, and an optional Hewlett-Packard Modular System Interface Bus (HP-MSIB) interface. The MSIB interface is compatible with the HP modular measurement system (MMS). The front panel contains three 8-bit microprocessors of its own to provide the required functionality.

The CPU is a Motorola 68000. It is augmented with 512K bytes of RAM, 512K bytes of EPROM, 32K bytes of EEPROM, a digital voltmeter, a counter timer system, and specialized interface hardware.

The power supply module contains a large switching supply along with a linear standby supply. The supply module consists of two preregulator boards in a shielded enclosure and one postregulator board, which provides the final supply voltages.

RF Deck

The microcircuits in the HP 8360 are laid out on a single RF deck. This was done to allow assembly and testing of the microcircuit assembly separately from the instrument, to provide the flexibility to accommodate different micro-

circuit configurations to meet customer requirements, to shorten the length of the microwave connections to minimize power loss, and to ease service. Several alternatives were considered to achieve these objectives. A side tip-out RF deck was considered but was found unacceptable in rack-mount applications. Two alternatives were considered to reduce power loss caused by excess cable length. One of the alternatives would have resulted in a completely different part layout depending on whether the instrument had the front-panel or the rear-panel output option. This would have reduced cable loss, but would have caused both manufacturing and field support problems. Another approach was to have the microwave output come from the center of the RF deck so both the front-panel and the rear-panel options would be equally penalized.

In the end, the RF deck was configured with a single layout that has the shortest possible path for the high-frequency cables. The conversion from front-panel to rear-panel output is achieved by reversing the RF deck. A special RF deck interface printed circuit board, which has redundant connectors on two sides, makes it possible to install the RF deck "upside down". The entire RF deck can also be
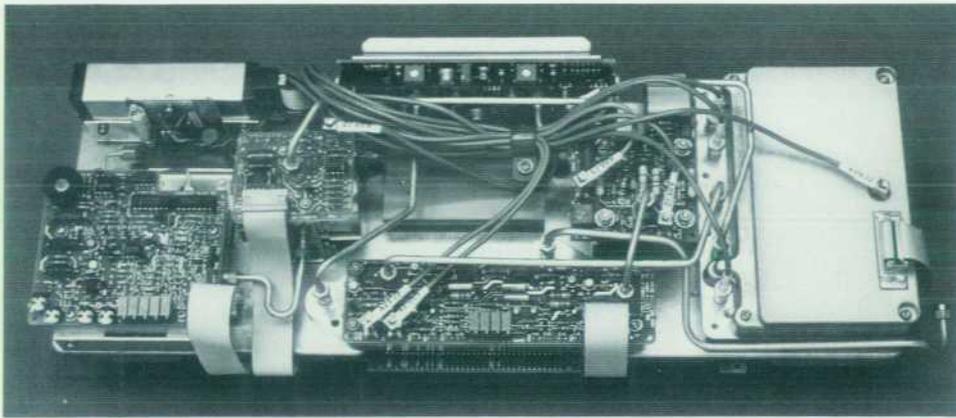
Fig. 4. *The RF deck holds all of the microwave components. For rear-panel input/output options, the deck is reversed to keep microwave path lengths as short as possible.*

mounted on top of the instrument on special mounting points. This facilitates service and production test and allows access to both sides of the deck (see Fig. 4).

Basic Model and Options

The HP 8360 series synthesizers offer a wide range of microwave performance options to the user. By choosing the appropriate model within the HP 8360 family and specifying a few options, one can tailor the microwave performance to match specific measurement needs. The foundation of this new family of synthesizers is the Model HP 83620A. This source uses a proven low-cost approach

to provide 0.01-to-20-GHz frequency coverage and +10 dBm of output power while maintaining harmonics and subharmonics below 50 dBc. Higher output power and broader frequency coverage are available in Models HP 83623A and HP 83650A, respectively. Thus, a user can purchase only as much performance as is necessary.

The basic Model HP 83620A, shown schematically in Fig. 5, uses a single 2.3-to-7.8-GHz YIG oscillator with a combination of heterodyne and multiplier techniques to achieve broad frequency coverage. The low band (0.01 to 2.4 GHz) is generated by mixing the YIG oscillator output with a fixed oscillator at 5.4 GHz (see article, page 36). The
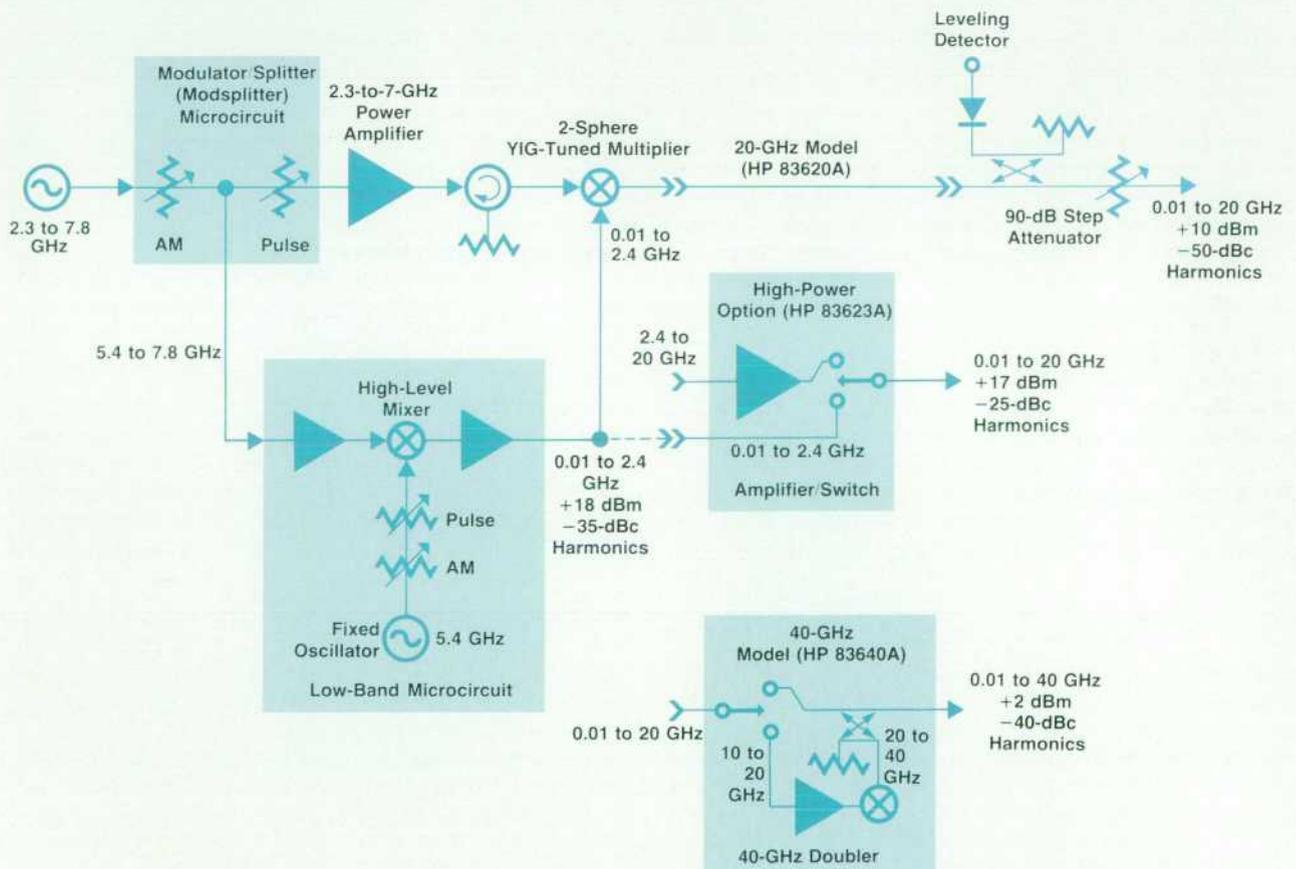
Fig. 5. *HP 8360 microwave block diagram, showing the standard 20-GHz and 40-GHz models and the high-power option.*

Designing for Low Cost of Ownership

An investment in microwave test equipment goes well beyond the original purchase price of the product. In the long term, annual maintenance and repair costs can easily offset any differences in purchase price. When costs related to lost productivity are included, a product's cost of ownership can exceed twice the original purchase price.

Cost of ownership is the sum of three parts: purchase price, repair costs, and calibration costs:

Annual Cost of Ownership = Purchase Price/Useful Life + Annual Repair Costs + Annual Calibration Costs.

The purchase price is spread over the useful life of the product. Repair and calibration costs are usually stated as yearly averages. For repair and calibration, a major cost contributor is lost productivity (for example, loss of shipments in a manufacturing environment). High reliability and extensive service and calibration tools ensure maximum uptime and usefulness.

HP 8360 annual repair and calibration costs are expected to be about 40% of those of similar synthesized sweepers. This is because low cost of ownership was a primary objective of the HP 8360 series design team. Early in the design phase a service and reliability engineer was assigned to help the design team focus on the right issues. In addition, a firmware designer ensured that service-related features were included in early designs.

Designing for Reliability

The key to designing a very reliable product is first to ensure that the objective is clear to everyone involved in the project and then to provide an environment that promotes success. High visibility was provided by assigning a reliability engineer to the design team shortly after design feasibility was proven. The reliability engineer provided each design engineer with the design information and tools necessary for a reliable design, coordinated design reviews, and provided the focal point for all reliability testing.

Component Derating. Strict adherence to HP corporate component selection and derating guidelines in the initial design provided the cornerstone for the highly reliable design. All designers were given a review of the derating concepts for each component type (resistor, capacitor, transistor, etc.). Stress analysis forms were developed to aid in the calculations and ensure common derating practices. The types of stresses analyzed varied with component type, but typically included voltage, current, and temperature. Stress levels were calculated for worst-case operating conditions. Three component derating classifications were defined for each component: good, acceptable, and unacceptable. The reliability engineer worked with the designer when an initial investigation indicated that good components were not available. About 5% of the components used were in the acceptable range. Unacceptable was precisely that—unacceptable. Quite often components selected were specified to operate at more than twice the stress levels encountered in the HP 8360.

Design Reviews. Peer design reviews were held before sending out for first prototype hardware and before any major design changes. All participants were given a package that included a design intent, stress analysis results, a parts list, a block diagram, and schematics. Generating this package gave the designer a good opportunity to double-check the design. Subjects covered by the review team included performance, reliability, manufacturability, and serviceability. Instrument level reviews were held for such items as airflow, thermal profile, and mechanical structure. Key benefits derived from the design reviews were reduced component stress, reduced parts count (30% fewer parts than the HP 8340 synthesized sweeper), fewer adjustable components (75% fewer than the HP 8340), and better self-test and service features. Instrument level design reviews helped ensure a well-distributed airflow and low internal temperature rise while using a quiet fan that meets the growing noise concerns of users. The HP 8360 is quieter than the typical overhead projector.

Reliability Testing. Reliability testing was done at both assembly and instrument levels. The team identified two areas they wanted to test rigorously (power supply and alphanumeric display) and started testing before a complete instrument was available. The alphanumeric display is purchased from an outside vendor, so it was important to qualify the part early and work with the vendor to implement any required design changes (see "Strife Testing the Alphanumeric Display," page 13).

One of the first prototype instruments was dedicated for reliability testing. The instrument was subjected to stresses commonly seen during normal use (vibration, shock, temperature variation). The idea was to start at low stress levels where no failures occur and slowly increment the stress levels until a failure occurred. That failure was then analyzed and the design changed to fix the problem. Retesting at the same stress level was required to verify the new design before the stress level was incremented. This cycle was repeated until stress levels far exceeded commercial standards.

Doing reliability testing as early as possible was a definite advantage for the HP 8360. The team was able to redesign weak areas and not just patch current designs. For example, the initial card cage design used sheet metal. Vibration testing indicated many weak areas that required additional support. This feedback to the mechanical engineers enabled them to justify the tooling costs to build the card cage from cast aluminum parts. This not only resulted in a sturdier instrument, but also in a lower-cost card cage.

Designing for Serviceability

Having a service engineer assigned and located in the lab during early development helped the design team make decisions that promoted serviceability. Close attention to serviceability was carried throughout the design cycle, including printed circuit layout.

Service firmware was designed in conjunction with the hardware, and service tools were often provided to aid in the design. This ensured a design that provided the most service capability without adversely effecting design time or assembly cost. As a result, the extensive self-tests used in conjunction with quick and simple manual diagnostic procedures isolate more than 97% of failures in minutes. The self-tests were developed for fast and easy repair at the assembly level, but also provide a great deal of help if component level repair is required. Firmware-aided adjustments improve adjustment speed and accuracy, reduce test equipment needs, and require less technical skill. For a more detailed description of the service firmware, see the article on page 17.

The RF deck design described in the accompanying article was heavily influenced by the design reviews. The RF deck can be placed in a service position on an extender board for troubleshooting and repair. Performance and configuration flexibility goals were not compromised in the process.

high band (2.4 to 20 GHz) is achieved by multiplying the oscillator output frequency using a YIG-tuned multiplier. The multiplier uses a step-recovery diode to generate second and third harmonics. Although a step-recovery diode is an efficient harmonic generator, a power amplifier is still necessary to boost the signal driving the diode. To produce +13 dBm at 20 GHz the drive signal must be around one-half watt. There was a time when this would have been difficult, but technology improvements in the last few years have made this a relatively simple task in the 2-to-7-GHz range. A two-sphere YIG filter following the step-recovery diode then selects the proper harmonic signal for 7-to-12.5-GHz and 12.5-to-20-GHz coverage.[4] For frequencies in the 2.4-to-7-GHz range, the step-recovery diode is forward-biased and the YIG filter passes the oscillator signal straight through. The two-sphere YIG filter ensures that the unwanted harmonics are suppressed below −50 dBc and provides the HP 83620A with exceptionally low broadband noise. The low-band and high-band frequency ranges are combined at a pin diode switch built into the multiplier for continuous 0.01-to-20-GHz coverage.

Heterodyning and multiplying a single YIG oscillator offers three distinct advantages over other common block diagrams. The first is low cost. YIG oscillators are expensive, and more than one would typically be required to cover the 2-to-20-GHz range. Instruments employing multiple broadband oscillators to achieve full frequency coverage also demand broadband performance in their power amplifiers, level control modulators, and frequency stabilizing circuits. Often, multiple sets of amplifiers and modulators are used because of the difficulty of building broadband components. This further increases cost.

A second advantage of this block diagram is that swept frequency accuracy is superior using a single oscillator. Linearity typically suffers when individual oscillators with differing characteristics are multiplexed, and correcting these differences while sweeping is difficult.

A third advantage is that the bandswitch frequencies are more flexible in a YIG-filtered system. This can be important when measuring narrowband devices with points of interest near the instrument's bandswitch frequency. Depending on the selected sweep width, the instrument can eliminate bandswitch points by using different bands. A multiple-oscillator block diagram usually employs a series of switchable low-pass filters to remove unwanted harmonics. In this case, the switch frequencies cannot be changed without degrading harmonic performance. This is because the corner frequency of each filter is fixed.

High-Power Model

Certain measurement applications require more than +10 dBm of output power from the source. To provide for this case, a broadband 2-to-20-GHz power amplifier with an integral pin diode switch was designed to follow the YIG-tuned multiplier. This extends the output power of the basic source to +17 dBm and is available in the Model HP 83623A. Fig. 5 shows this block diagram. The power amplifier uses four state-of-the-art traveling-wave amplifier ICs to achieve +21 dBm of output power at 20 GHz. The final section of this three-stage amplifier employs two custom HP microwave ICs driven in parallel. This increases the available output power by typically 2.5 dB over that achievable using only a single IC. Also, the harmonics generated are lower since each IC in the last stage is operating at a lower power level. The low-band frequency range is added at the integral pin diode switch to deliver continuous 0.01-to-20-GHz coverage as before.

A clear advantage of offering higher power as a different model in the HP 8360 series is cost. Customers with measurement needs satisfied by the standard product are not burdened by the higher cost of the broadband component.

40-GHz Model

Another performance category receiving considerable at-

tention in the past few years is frequency coverage in coaxial systems. The HP 8360 series offers Model 83640A, which provides 0.01-to-40-GHz coverage in coax. This is achieved by employing a specially designed switchable frequency doubler (see article, page 31). The doubler immediately follows the YIG-tuned multiplier of the standard model. Fig. 5 shows this block diagram. The doubler allows two modes of operation. In the first, the input signals are routed straight through. In the second mode, input frequencies in the range of 10 to 20 GHz are doubled, then coupled to the output. In this way, continuous coverage from 0.01 to 40 GHz is available at one coaxial port. The HP 83640A is available with a 90-dB step attenuator for applications needing very low source power. This step attenuator offers low insertion loss at 40 GHz along with high repeatability of each step (see article, page 47). The doubler and attenuator both use the industry-standard 2.4-mm coaxial connector. These two components have recently been extended to 50 GHz.

Each of the three basic models in the HP 8360 series offers a variety of options that increase the range of choices available to the user. One important option is a high-speed pulse modulator that provides pulse rise and fall times typically less than 2 ns (see article, page 34). This flexibility ensures that users' current measurement needs are met while providing a path for expansion as new applications arise.

Other Features

Each HP 8360 model has a high-capacity switching power supply with enough reserve power to handle the requirements of extra microwave components. The design also provides extra cooling capacity with greatly reduced fan noise. This was achieved by improved fan technology, specialized geometry of the fan grill, and grill placement away from the fan blades. The fan blows air into an open plenum which contains the microcircuits. Air is metered out of this area through cooling ports that conduct air across each of the major printed circuit boards.

Flexible User Interface

The user sees the HP 8360 either from the point of view of the front panel or as a system element that is controlled by a computer.

Front-Panel User Interface

The front-panel interface represents a departure from the traditional way of designing user interfaces for instruments without a CRT. Since the HP 8360 was designed as a platform for synthesizer products for the decade of the 90s, it required an extensible approach to user interfaces. The decision to develop a new front-panel user interface was greatly influenced by our experience with the previous-generation synthesized sweeper, the HP 8340. This product had a large number of keys arranged by functional grouping for all of the major functions of the synthesizer. This method worked well for the functions that could be directly associated with one of the keys, but a large number of minor features had to be relegated to a shift function. A blue shift key was provided to double the number of keys

available. Most of these shift functions had a logical association with a particular set of keys and were allocated from the beginning. As new features were added with later releases of firmware, it was very difficult to find an appropriate key to which to bind some of the new features. The solution to this problem was to develop a menuing system that is not only clear and easy to use but also provides expandability and flexibility.

The first element of the user interface to be selected was the display, since much of the rest of the system had to be designed around this component. The display selection criteria we used were gathered from customer feedback, requirements of a menuing system, our own experience, and environmental and reliability requirements. The display had to be very reliable so that the overall instrument reliability would not be adversely affected, especially since we were competing directly with the excellent reliability of the HP 8340. It had to have a sufficient number of characters and lines to allow a menu structure with menu keys adjacent to the display. The characters had to be easily readable in both bright and dimly lit areas and had to be readable from a distance of several armlengths. We looked at every type of display technology that was available at that time and settled on the six-line-by-40-character vacuum fluorescent display that was available from several manufacturers. The display selected did have some environmental concerns, which were resolved by working closely with a single manufacturer.

Keyboard

The keyboard was specifically designed to address problems in manufacturing, customer complaints, cost, reliability and system requirements. We needed a very reliable key technology that was simple to assemble and aesthetically pleasing. We needed to provide RFI shielding around the front panel and display. Additionally, we needed to block all air flow out of the front to address customers' complaints. Taking all of these considerations into account, we selected a one-piece rubber keypad that could provide all of these functions in a single piece. We were sensitive to concerns that various people had expressed that the feel of a flexible keypad would not be correct. Our industrial design group developed a promising design for an antirock mechanism that would give the rubber keys the type of stability and feel that was necessary (see "Front Panel Designed for Manufacturability," page 15). The final design of the keypad provides "soft touch" keys that don't rock sideways and have a detent-type feel for tactile feedback through the use of the antirock membrane. The keypad seals the area around all of the keys so that no air escapes to irritate the user. Conductive rubber portions of the part reduce the amount of RFI that escapes and prevent air from escaping around the display window and the periphery of the display casting.

Layout and Menus

The layout of the front-panel functions was influenced by the HP 8510 vector network analyzer system, which groups similar functions together with a menu select key that provides access to similar functions on a menu. Our goal was to provide clear single-key access to the most

important and most often used features such as basic sweep control, basic frequency control, basic power control, and necessary system keys such as save and recall. The menu keys are placed under the display so that they are easy to associate with a display field that is seven to eight characters wide by two lines high. Additional menu select keys and menu control keys are placed on both sides of the display (see Fig. 6).

The approach to the menuing scheme was to provide a simple, elegant, easy-to-understand, and easy-to-use interface for the inexperienced user. At the same time, we wanted to provide the power and functionality that would be demanded by the most sophisticated users. The first-level menus are designed to provide all of the features that

are normally used by a large percentage of customers and the less often used (but very powerful) features are placed on the lower-level menus.

A fairly extensive user interface survey was conducted to find problem areas with the menuing system. We used our prototype instruments and invited a cross section of marketing, R&D, production, and field people to try out the menuing system. We gave them minimal instructions so that we could discover how easy-to-use and self-explanatory the system was. We assigned particular areas of the system to certain people depending on their areas of expertise. One of the lab designers was a silent observer and took notes. From this survey a multitude of useful changes and enhancements were discovered. As a result,

---

Strife Testing the Alphanumeric Display

To enhance the user interface of the HP 8360 synthesized sweepers, a multiline alphanumeric display was desired. After a preliminary investigation, we decided to evaluate vacuum fluorescent display technology. A six-line, 40-character alphanumeric display, complete with power supply and drivers, was available from several vendors.

Having no previous experience with vacuum fluorescent display technology, we first had to determine what stress modes and levels would be most beneficial in evaluating the displays. The display is basically a vacuum tube with the plate being the display grid (see Fig. 1). Vibration provided a good check of ruggedness for the internal elements and the printed circuit board. Operating temperature range provided a measure of component electrical derating, while temperature change rate had its biggest effect on the display's vacuum seal. Shock was a good component ruggedness test. Initial testing (before any re-design) was conducted on three displays from each manufacturer.

Initial stress levels chosen were considered minimal and were used as a baseline that indicated the display could survive each stress mode without any damage.

Vibration testing consisted of both random and sine wave tests. A swept sine wave test was run to find any resonant frequencies within the display or printed circuit board. Initial random vibration was conducted at 1g rms. Both tests were run in three axes (x, y, and z). Any resonances were noted for later reference. The display was powered during these tests.

Temperature testing was initially conducted over normal operating conditions (0 to 55°C). The display was subjected to ten temperature cycles and was operated at each extreme for a half hour. The rate of temperature change was limited to 2°C per minute.
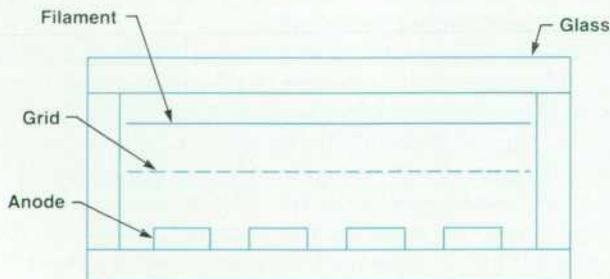
Shock testing was done in six axes with an initial level of 30g for a duration of 11 milliseconds.

As a result of these initial tests, one manufacturer's display was chosen for the HP 8360.

Further Testing

After random vibration testing at 2.1g rms, portions of the chosen display had become dimmer. Closer inspection indicated that filament wires in these areas were a different brightness. Discussion with the manufacturer indicated that the filaments were probably hitting other elements and knocking off their oxide coating. Swept sine measurements confirmed that the filament resonant frequency was within the frequency range of the random vibration test.

This problem was found before the first prototype HP 8360 instrument was complete. This gave us plenty of time to work with the vendor and validate any design changes. The result is a more reliable display. Additional testing uncovered other potential problems that were designed out and never became part of a customer's HP 8360 synthesized sweeper.

At some point, further increases in stress levels do not provide useful feedback because the failures found may not be practical to eliminate or because there is a practical limit on the stress level that can be applied. The maximum stress levels for the alphanumeric display were:

Stress Mode	Maximum Level Passed	Maximum Level Tested
Random Vibration	5.5g rms	6.0g rms (1)
Temperature Range	−40 to +100°C	−40 to +100°C
Temperature Rate	5°C/min	10°C/min (2)
Shock (11 ms)	55g	55g

Notes:
1. Random vibration testing at 6g caused a getter ring inside the display to come loose. Only one display was tested at this level and it still operated.
2. Display lost vacuum (2 out of 3 survived) when tested by itself. Displays installed in instruments survived 10°C/min rates because they were isolated from the chamber temperature.

James R. Stead
Product Support Engineer
Network Measurements Division

Fig. 1. *Basic construction of the vacuum fluorescent display used in the HP 8360 synthesized sweepers.*
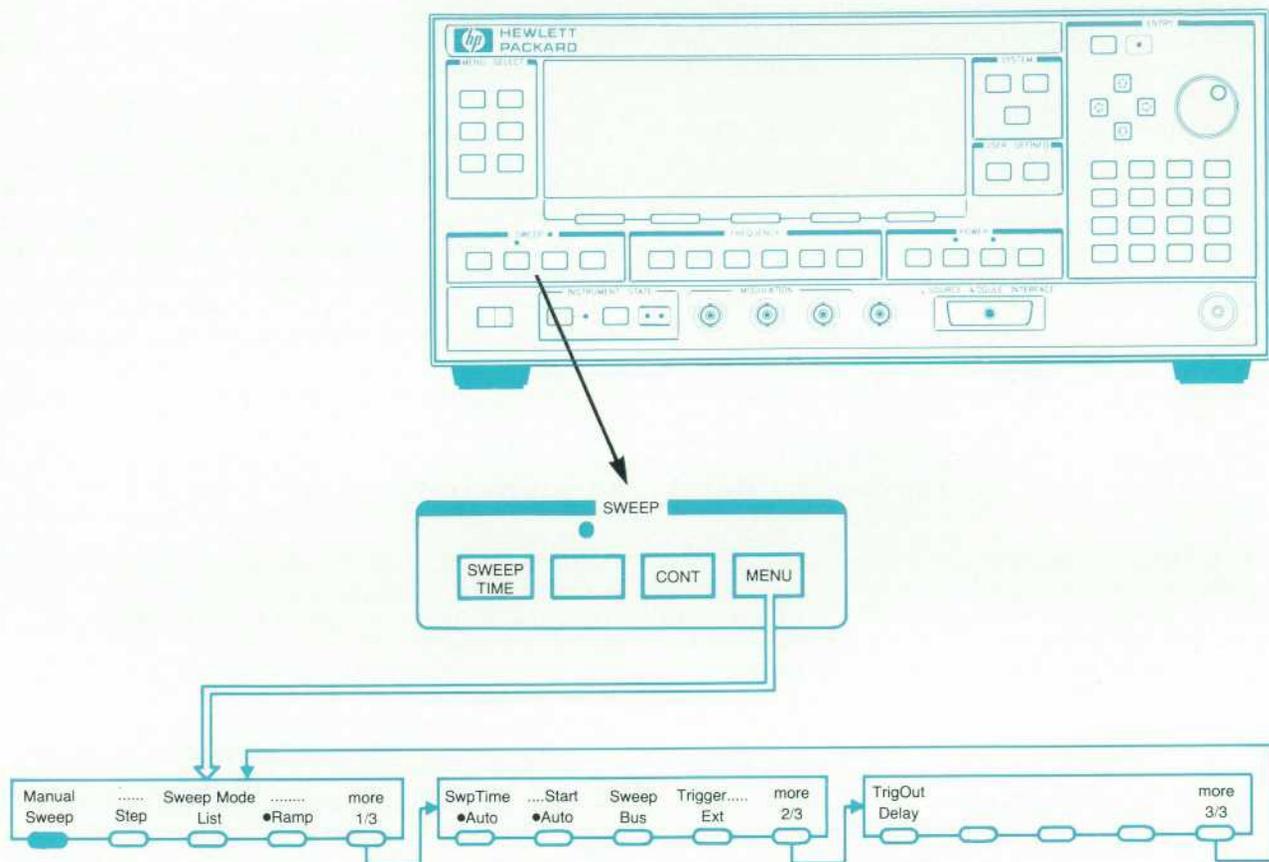
Fig. 6. *HP 8360 menu concept. Similar functions are grouped together, with a menu select key as part of each group. The concept provides clear single-key access to the most often used features and places the less often used functions on the associated menu.*

the character fonts were modified, the information presented on the display was reorganized, many key labels and other text were changed, the user menu feature was changed to make it easier to use, and many menus were reorganized. Obviously, not all suggestions could be implemented because of other conflicting requirements, desires, or limitations of the hardware, but this process did allow a much more thorough discovery of problems that would have otherwise gone unaddressed.

The user-defined menu has turned out to be a very useful way to have often used menu functions at a user's fingertips. The user simply assigns to the user-defined menu those menu items to which immediate access is desired. To use one of these selected functions, the user presses the user menu key to bring up this custom menu and selects the desired function.

The current status of the instrument is communicated in many ways. Dedicated LEDs show sweeping, RF on/off, and bus functions. A dedicated line on the display shows the status of important features such as AM, FM, unleveled, or unlocked. Asterisks near the frequency or power readout signify that an offset is in effect. An asterisk near an on/off menu function signifies that this menu function is on. Thus, the status of the most important features is displayed constantly, and a user is not bothered by the less significant ones.

Many customers requested that certain qualities of the preset key be changed (for instance, power level, power on or off, sweeping or CW). To address these needs, the HP 8360 has a user-defined preset register, which is similar to a save/recall register except that it can be configured to be recalled whenever the preset key is pressed. To avoid confusion, a prominent message is displayed in the entry area of the display when a user-defined preset as opposed to a generic preset is performed. A user can easily put the HP 8360 back in its factory-preset state by pressing a menu key that becomes available after a user-defined preset occurs.

Computer-Control User Interface

All remote computer control of the HP 8360 is either by the IEEE 488.1 HP-IB interface or the optional MSIB interface. The language support required an extensive firmware effort. The HP 8360 follows the very successful HP 8340 microwave synthesizer and therefore it was considered absolutely necessary that the HP 8360 respond to the same programming mnemonics to achieve application software compatibility. This compatibility language was also needed to support the interface required by vector and scalar analyzers. Additionally, a special input parser was needed to implement the U.S. Air Force language CIIL which is

Front Panel Designed for Manufacturablity

Fig. 1 shows the front-panel design of the HP 8360 synthesized sweep oscillators.

A high degree of integration was achieved by basing the front panel design on two key components: a one-piece rubber (flubber) keypad and a casting that incorporates as many mechanical features as possible. These parts plus a printed circuit board, an antirock membrane, and a vacuum fluorescent display meet all of the requirements for the HP 8360 front-panel interface. The total number of parts in the front-panel assembly, not counting the electrical components on the printed circuit board, is 64. The front panel can be assembled in six minutes. The previous-generation front panel contained 243 parts and took 90 minutes to assemble. The front-panel casting alone eliminated the need for 177 individual parts.

The front-panel casting contains all of the standoffs and mounting provisions required by the various components. The casting is painted and silkscreened directly so that no additional dress panel is required. A stippled clear coat covers the front, protecting the surface and giving a long-term blemish-free appearance. Assembly and disassembly are very easy because of the small number of extra assemblies and hardware.

The feel of the keys was important because customers expect a quality keyboard on a relatively expensive product. A patented antirock mechanism was developed which allowed us to achieve this quality feel. The antirock sheets are precisely cut pieces of mylar that grab each key and prevent its flopping over to one side when it is pressed. The antirock sheets are referenced to location pins on the casting. The location pins force each key to be properly aligned with the associated hole in the casting. The overcenter feel is provided by molded details in the rubber keypad. The final design of the keyboard provides "soft touch" keys that don't rock sideways and have a detent-type feel for tactile feedback.

Air is prevented from escaping around any of the keys because the keypad seals the area around the keys. Seals are also provided around the display window and the outer periphery of the display casting. Air blowing in the face of a user, a common complaint in the past, is eliminated with this new design.

To meet certain RFI emission and susceptibility requirements it was necessary to provide an electrical seal for the display window area and the slots around the periphery of the display casting. This was achieved by specifying that conductive rubber be used to mold the keypad in the required areas. In this way a single molded part is able to provide the keypad, the airflow blocking, and the RFI shielding functions.

The display filter is made from two pieces of plexiglass with a stainless steel mesh between them. The conductive portion of the keypad holds the filter in place and provides an electrical path between the conductive filter and the casting.

James E. Bossaller
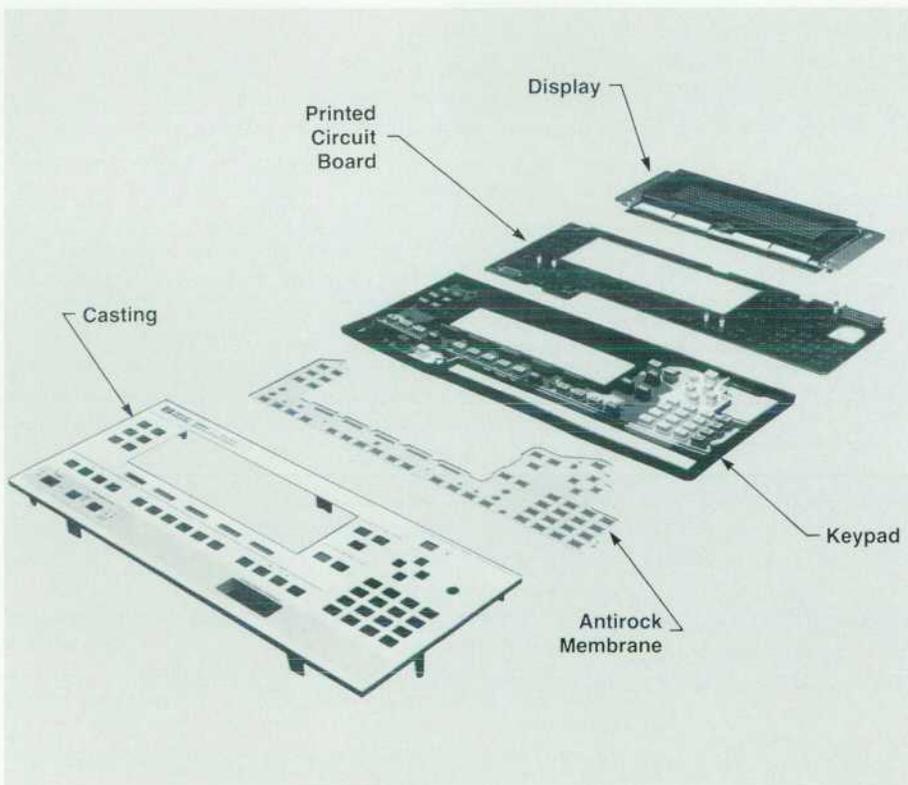Project Manager
Network Measurements Division

Fig. 1. *The rubber keypad and the special casting are key components of the HP 8360 front-panel assembly.*

used by several customers.

SCPI

The largest amount of effort, however, went into a new language that was beginning to be developed by HP during the early part of the HP 8360 development. This language was initially called HP-SL (Hewlett Packard System Language). Later this name was changed to TMSL (Test and Measurement System Language) and was proposed to the industry as a new standard. As this took place, the administration of the language development went to a consortium of instrument manufacturers. The consortium changed the name one last time to SCPI (Standard Commands for Programmable Instruments).

This language is based upon the IEEE 488.2 standard, which defines recommended communication protocol for the IEEE 488.1 HP-IB interface. This standard defines the overall syntax structure, the data structures, and the common system control commands for status, testing, save and recall, identifying, and other functions that are needed in a system regardless of the type of instrument. The SCPI language goes the next step and defines the actual commands that must be used to set, query, or control the features of an instrument. In this approach, all instruments capable of outputting a signal, for example, will program the frequency and other parameters of the signal with the same commands. The SCPI language is a tree-structured language that has been well thought out and provides the capability for the HP 8360 to extend its programming language easily as new features are added to the family.

Some of the major features of the SCPI language are summarized here since the HP 8360 is one of the first major instruments to implement the new language fully.

1. SCPI is a tree-structured language that is extensible to whatever degree is necessary to describe a setup or measurement. SCPI commands are human readable. English names are used for all commands and regular truncation rules are used to shorten the commands if desired. The tree structure allows logical placement of commands.

2. The feature set is orthogonal. In other words, SCPI tries to eliminate interactions between settings of functions that are supposed to be mutually exclusive.

3. SCPI requires order independence within a single message except for immediate execution commands such as synchronization commands and queries. For instance, parameters that are coupled together in some manner, such as FREQ:START, FREQ:STOP, FREQ:CENTER, and FREQ:SPAN, can be sent in pairs in any order with the same result. If the allowed values of a parameter change based upon a mode switch, then the parameter and the mode can be programmed in any order within a single message with the same result.

4. All values and settings can be queried without affecting the current state of the instrument.

5. Capabilities can be set up completely before the function is turned on. For instance, all parameters related to pulse modulation can be specified before pulse modulation is turned on.

6. SCPI contains the synchronization commands *OPC, *OPC? and *WAI. *OPC and *OPC? report when a pending operation or measurement is complete by setting a bit in the standard event status register or by responding to the query with a "1", respectively. With these two commands, the HP-IB bus will not be held off and additional commands can be sent and processed. The *WAI command, on the other hand, suspends the parsing and execution of any further commands until the pending operation is complete.

7. SCPI and IEEE 488.2 encourage the use of an input buffer for improved system throughput. With an input buffer it is not necessary to wait for one instrument to complete an operation before communicating with other instruments in the system tied to the HP-IB.

8. SCPI provides a very complete set of status and service request capabilities. Many standard conditions are required in all instruments and this status system is expandable to any extent required by a particular instrument. The service request control consists of an enable for each bit and both positive-edge and negative-edge transition filters for maximum flexibility.

9. SCPI defines standard conditions for all controllable nodes after a *RST command, ensuring that the system is in a known state.

10. SCPI achieves both vertical and horizontal consistency. Vertical consistency is achieved by having instruments in the same class use the same set of commands so that all synthesizers, for example, look the same when viewed through the HP-IB interface. Horizontal consistency is achieved by having similar functions in different classes of instruments respond to the same commands. An example of this would be FREQ:START <value>, which will have the same effect in both a sweeper and a spectrum analyzer.

The HP-IB language and address can be changed via the front-panel system menu. In addition, rear-panel switches are available for the option that deletes the display.

Acknowledgments

References

1. J.L. Thomason, "Expanding Synthesized Signal Generation to the Microwave Range," *Hewlett-Packard Journal*, Vol. 29, no. 3, November 1977, pp. 2-8.

2. U.S. Patent 4130808, "Phase Lock Stabilized Swept Frequency Signal Source."

3. S.H. Linkwitz, "New Performance Standards in Microwave Spectrum Analysis," *Hewlett-Packard Journal*, Vol. 30, no. 8, August 1979, pp. 3-7.

4. L. Rhymes, "Two-Sphere YIG Multiplier/Filter Ensures Purity," *Microwaves & RF*, April 1988.

Built-in Synthesized Sweeper Self-Test and Adjustments

A combination of hardware features and firmware routines makes it possible to isolate most failures to the assembly level and make many adjustments without external test equipment.

by Michael J. Seibel

N THE DESIGN of the HP 8360 synthesized sweep oscillators, considerable effort was expended on providing built-in self-test and service features. The HP 8360 self-test consists of over 900 analog, digital, and RF measurements taken in approximately 45 seconds. More than 90% of failures are detected and diagnosed to determine the most likely fault and a message is presented directing the operator to a specific entry in the service manual for further isolation. Many of the adjustments can be done without the aid of external test equipment through the use of internal firmware routines. Other firmware routines provide a collection of service tools to manipulate the hardware to facilitate troubleshooting. All of these self-test and service features are easily accessible from the front panel and over the HP-IB (IEEE 488).

This level of serviceability required the full participation of the HP 8360 design team. As each of the printed circuit boards was designed, it was reviewed for serviceability. As a consequence of these reviews, most boards required the addition of small amounts of hardware or connections between pieces of existing hardware. By designing in the service hardware early, the risk of disturbing an otherwise stable design is avoided.

Service firmware was designed in conjunction with the hardware, and service tools were often provided to aid in the design. These service tools perform read/write operations to devices on boards, read voltages in the instrument, control phase-locked loops, and monitor the contents of latches. These became very useful during prototyping because the self-tests provided quick, thorough checks of the hardware and digital control of the boards before the normal instrument firmware had been completed. The service tools allowed the designers to manipulate the hardware to exercise their prototype boards. These tools remain in the HP 8360 and are supported to facilitate troubleshooting to the component level.

Many circuits in the HP 8360 provided the opportunity to use digital-to-analog converters (DACs) in place of potentiometers for adjustments. There was concern about the cost of the DACs on one hand and about the reliability and long-term stability of the potentiometers on the other, so the decision to substitute DACs for potentiometers was made carefully. DACs were substituted for potentiometers only if a single DAC replaced several potentiometers, or if adding the DAC allowed an adjustment to be done using internal

firmware instead of external equipment and procedures.

The hardware changes made for self-test tend to complement those made for adjustment. The adjustment routines make use of the self-test analog bus and peak detectors. The self-tests benefit from the ability to use DAC control of some circuits.

Self-Test

The full instrument self-test is the most often used service feature. This is initiated by a single keystroke from the SERVICE menu. It executes all of the assembly-level self-tests in the instrument and does a diagnosis of failures, if any. A scrolling menu of the self-tests is also provided with the tests organized into assembly-related groups. Any of these groups or a specific test can be run individually by selecting from the menu. The following is an example portion of the self-test menu:

0	Full Instrument Self Test	Passed
1	YO Loop Full Test	Passed
2	YO Loop Digital Interface	Passed
3	YO Loop Phase Detector	Passed
4	YO Loop Integrator	Passed
5	Sampler Loop Full Test	Passed
6	Sampler Phase Amplifier	Passed
7	Sampler IF Amp Bias	Passed

Any test in the list can be looped repeatedly. While looping, a status display is presented showing the cumulative numbers of times the test has passed and failed. Each self-test typically consists of many internal measurements that are checked against high and low limits to determine a single pass/fail result. A single component failure can very likely cause many of the self-tests to fail, so the diagnosis routine examines the pass/fail results in a particular order to determine the most independent test that failed, effectively ignoring the subsequent failures.

There are situations where the self-test measurement data is of value to the user. In the production area it is desirable to examine the data for analysis of trends, comparison to tighter limits, documenting instrument performance, and so on. In these cases, the simple pass/fail result code is insufficient. To make this measurement data available, the self-test firmware is structured to pass all of the measurements through a single procedure to compare them to

specifications. The measurement data, specifications, and a short description are passed to the procedure. If a data logging feature is activated, the measurement data is presented to the user. The following is an example of the measurement data display:

```
213  A14 : Level Correction DAC    --??--
213.007) A14 : Level Correction DAC    :
---data point---     -meas-  -min-  -max-
Bit 5 set [mV]          -68    -102    -42
          ....Options....
Abort          Loop    *Log   Cont
```

The test number (213.007) identifies this measurement data as part of self-test number 213, the A14 : Level Correction DAC test. The data point field describes the individual measurement, which in this case is the DAC output in mV. The measured data (−68) is shown in addition to the limits against which it is tested (−102 and −42). If this were a digital type of test, where data is written to a device and read back, the minimum and maximum specifications would be identical. Thus, the same display and procedure are used for digital and analog self-test measurements.

The user can request all of the self-test measurement data or restrict the display to only failed self-test measurements. The user can also select the destination of the logged self-test measurement data as either the HP 8360 display or an HP-IB printer.

The combination of self-test looping and logging provides some interesting capabilities. For troubleshooting an elusive intermittent failure, one can select the looping option and log failed measurement data to the printer. When the intermittent failure occurs again, the failed self-test data will be captured on the printer, revealing what was wrong and by how much. If the failures are logged to the display, the hardware will be frozen in the state that existed when the failure occurred, while the failed measurement data is being displayed. This allows the circuit to be evaluated in the failed condition.

The ease of self-test execution is of considerable value in Hewlett-Packard's ATE environment. When doing temperature testing of the synthesizers, for example, the self-tests can verify the integrity of the hardware in 45 seconds without test equipment. Since all of the service features are accessible over the HP-IB, they are easy to perform in a test chamber where operator intervention is difficult.

Self-Test Design Philosophy

The goal for the HP 8360 built-in self-tests is to detect hardware failures and quickly diagnose the fault to a replaceable module. The task of detecting failures and diagnosing the probable cause is divided between HP 8360 self-test firmware and repair documentation, with most of the work provided by the firmware. Each of the self-tests performs a series of measurements to exercise a small section of circuitry and return a single pass or fail result. By providing a large number of small tests the tests can be kept more localized to a specific portion of the circuitry. Because of the many small tests, it is quite likely that a single failed component will result in many failed self-tests. The firmware then performs a diagnosis procedure to deduce which of the potentially numerous failures is actually the significant failure. The result of the firmware diagnosis is a display that directs the user to a specific diagnostic manual section for the suspect module. The following is an example of a diagnostic display:

```
The following self-test has failed :
    A5  : Switched Amplifier
Refer to ASSEMBLY-LEVEL REPAIR MANUAL,
paragraph A5.000, for assistance.
          ....Options....
Abort          Loop    Log    Cont
```

The repair manual provides further instructions on how to verify that the diagnosed module is indeed bad. In many cases the manual instructions will call for replacement of the indicated module. In other cases it is necessary to make several simple measurements to be certain of the identity of the failed module. This is particularly true of components in the RF section, where it is difficult to isolate a failure to one of several microcircuits in a chain with a simple dc-level analog bus. In these cases, a power meter measurement may be indicated to identify the offending microcircuit.

Examination of field failure reports indicates that a vast majority of failures are hardware failures, where a part has substantially degraded or failed entirely. Subtle parameter related failures, such as phase noise, spurious responses, and AM flatness, are not the usual cause of failures. These are not intended to be detected by the self-tests and are covered by the repair manual. Since most of the failures to be detected are the result of a significant change in a component, the specifications that the self-tests use to decide the pass/fail result can be rather loose. Loose specifications allow the self-tests to be less susceptible to false failures caused by minor component drift with temperature, humidity, or aging.

Self-Test Hardware Considerations

To isolate failures down to the board level, it was necessary to provide extensive monitoring of the output signals of the assemblies. Fig. 1 shows how this is done.

In Fig. 1, P1, P2, P3, P4, and P5 are monitoring points for self-test. The symbols W1, W2, W3, W4, and W5 are interconnecting cables. If P1 indicates a good signal, then assembly A is functional. If P2 and P3 are then tested, and they indicate good signals, then assembly B is functional. This continues down the circuit chain with each assembly examined only after the detectors supplying its inputs have been verified.

The first failure found is important for diagnosis. Consider a failure on assembly B that causes detector P3 to indicate a bad signal. This would likely cause detector P5 on assembly D to fail. The method of examining the detectors in order, P1 through P5, and stopping the search after the first failure is found causes the search to end at the P3
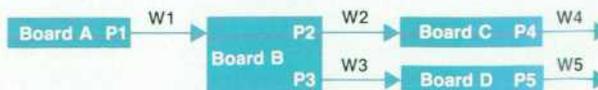

Fig. 1. *An example of failure isolation to an assembly.*

Automatic Frequency Span Calibration

The HP 8360 synthesized sweep oscillators use a "lock-and-roll" method of frequency sweeping similar to that used by the HP 8566A spectrum analyzer and the HP 8340 synthesized sweeper (see article, page 6). The start frequency of the sweep is synthesized by phase-locking the YIG-tuned oscillator (YTO) to a combination of a sampler loop harmonic and a low-frequency oscillator. The difference between the sampler harmonic and the YTO frequency is equal to the frequency of the low-frequency oscillator. This difference signal is referred to as the YTO intermediate frequency (IF). Once the phase-lock at start frequency has been achieved, the YTO phase-locked loop error voltage is maintained on a sample-and-hold circuit, the loop is opened, and a precision current ramp is summed into the YTO tuning coil. In this manner, the start frequency error is eliminated but the remainder of the sweep is subject to small inaccuracies such as gain errors in the tuning ramp, nonlinearities in the YIG tuning characteristics, and dynamic errors in tuning the YIG. The HP 8360 reduces these inaccuracies with a new method of frequency span calibration, which is used in conjunction with the lock-and-roll sweeps.

Fig. 1 is a simple block diagram of the YTO loop showing the IF counter. At the start of a sweep, the sampler loop is needed to lock the YTO to the proper frequency. Once the YTO sample-and-hold circuit has been switched to hold mode to maintain the YTO loop error voltage, the sampler loop can be retuned to a different frequency without affecting the YTO frequency. The sampler loop is reprogrammed in frequency so that a harmonic of its output is equal to about 90% of the YTO frequency span. As the YTO sweeps, it will eventually pass through the frequency of the sampler harmonic and the YTO IF will respond with a low-frequency (10 MHz to 70 MHz) portion of a frequency sweep while the YTO continues sweeping. The sweep timing of the HP 8360 is controlled by digitally generated "buckets" from the sweep generator. There are 1601 of these buckets in each sweep, and events can be set up to occur at certain buckets. One of these events is the enabling or disabling of a gate to a digital counter on the YTO IF path. By turning on the YTO IF counter gate at a precise time in the sweep and leaving the gate open for a precise amount of time, the YTO IF frequency can be calculated by dividing the counter contents by the gate time. Since the YTO IF was sweeping, the resulting frequency will equal the YTO IF frequency at the center of the gate time. This actual IF frequency can be compared to an expected IF frequency and the error in the YTO sweep can be calculated. The YTO sweep error is then corrected on the subsequent sweep by altering the slope of the tuning current into the YTO.

The IF counter is formed by providing a 4-bit TTL counter on the YTO loop board and then sending the counter output to a 16-bit counter/timer on the sweep generator board. These are combined in firmware to produce a 20-bit counter. The gate time, gate position in the sweep, sampler VCO frequency, and gate width are calculated to produce an IF counter reading that will not overflow the counter for any conceivable sweep span errors.
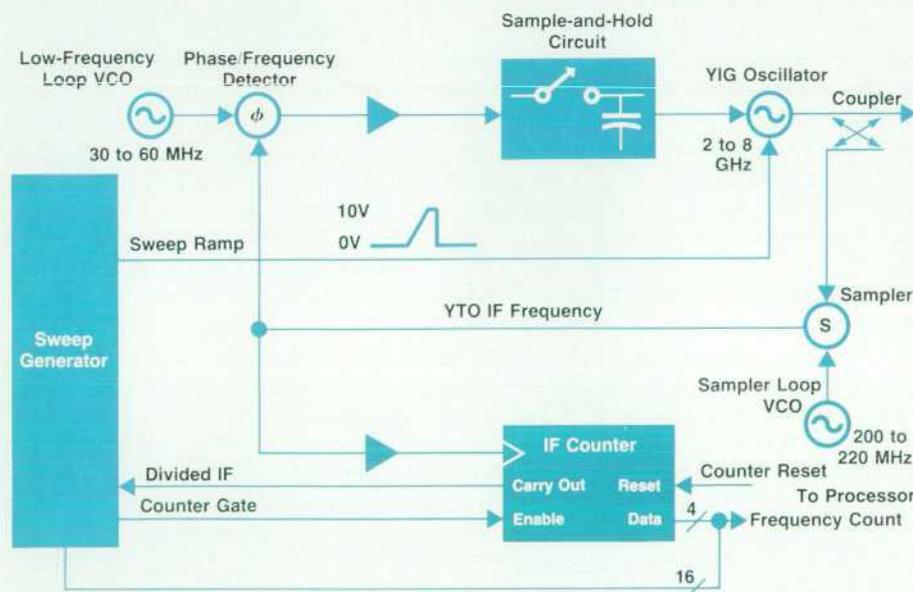
Fig. 1. *Simplified block diagram of the YTO loop showing the IF counter.*

detector. The failure at P5 is of no concern since assembly B is the likely failure, with cables W1 and W3 being possible failures. The failure is isolated to the replaceable assembly (either assembly B, cable W1, or cable W3) by means of simple measurements in the manual-guided procedure.

It is not necessary to provide monitoring points at the assembly inputs since these signals were tested previously on assembly outputs. To provide both input and output monitoring would increase the cost of the self-test hardware and provide little benefit. In fact, monitoring the inputs adds a degree of confusion to the failure isolation. In the unlikely event of a faulty signal detector, output monitoring will still correctly identify the assembly with the bad detector as the most likely failure. With input monitoring, a faulty input detector would cause the most likely failed assembly to be incorrectly identified as the preceding module. The assembly with the faulty detector would be missed until the manual-guided procedure detected it. For these reasons, the HP 8360 self-test hardware is directed towards monitoring the outputs of assemblies. Other monitoring

points are used if they are easily available and provide access to circuitry that is difficult to cover from the output, and to allow possible future fault isolation down to the subassembly or component levels.

The signal detectors in the HP 8360 measure dc voltages, RF signal levels, and digital data. For the analog detection, a 12½-bit analog-to-digital converter (ADC) on the main instrument processor board is used. This ADC is necessary for normal instrument control in addition to the analog self-tests. The ADC is connected through a network of distributed multiplexers to over 100 circuit nodes on the printed circuit cards. This analog bus (ABUS) typically provides eight monitoring points on each of the printed circuit boards. The ADC allows very accurate measurements of dc signals as large as $\pm 13V$.

To monitor RF signal levels on the printed circuit boards, the analog bus is connected to inexpensive diode peak detectors. These provide a simple, low-parts-count, high-reliability way of determining the presence or lack of RF signals. For digital testing, unused bits on data latches and buffers are frequently used in the self-test scheme as stimulus and feedback circuits. In some cases, inexpensive digital buffers were added to the printed circuit boards. The cost of this hardware is minimal.

The analog circuit nodes connected to the ABUS are generally low-impedance. By probing low-impedance nodes, sensitivity to induced interference from the analog bus is reduced. As an added precaution, the node is usually isolated from the bus by a 1-kΩ-to-10-kΩ resistor. When it isn't being used, the analog bus is placed into a carefully chosen idle state that connects all the analog bus lines to very low-impedance sources rather than simply disabling all the multiplexers and allowing the ABUS lines to float. As a result of these precautions, the analog bus has caused no interference problems between printed circuit modules.

Phase-Locked Loop Self-Tests

Phase-locked loops provide a challenge for self-testing and fault isolation. In some cases, the hardware of a phase-locked loop is split between several boards. In the event of a loop that is out of lock, it is difficult to determine which board is at fault. The HP 8360 self-test design addresses this by providing every phase-locked loop with an out-of-lock detector and a method of forcing the loop's phase detector to both plus and minus extremes. This makes it possible to verify the functionality of all portions of the phase-locked loop without depending on the loop's ability to lock. The loop integrator, for example, can be tested by measuring its output when forced to both extremes with the phase detector. If it responds correctly, it is assumed good. The out-of-lock indicator is also verified by forcing an unlocked condition. This simple dc test will detect most failures of the op-amp-based loop integrators since the common failure mode of the amplifiers is to drive the output to one power supply rail.

To verify the functionality of a programmable phase-locked loop without any external frequency counters or meters, the out-of-lock detector is first checked by forcing an unlocked condition and reading the detector. Once the out-of-lock detector has demonstrated its ability to indicate an out-of-lock condition, it can be trusted if it indicates a locked condition. The phase-locked loop is then tuned across its entire range. At each frequency, the out-of-lock detector and the tune voltage to the oscillator are checked. The tune voltage is checked for monotonic (unidirectional) changes as the frequency is stepped across its range. In many cases the output of the VCO is also monitored with a peak detector. This procedure detects output amplitude problems and any subtle output frequency problems caused by VCO tuning range or digital divider problems. If the phase-locked loop can be tuned across its range without indicating an unlock, the tune voltage changes monotonically, and the output level remains above the detector limits, it is unlikely that there is a fault with the phase-locked loop.

The loop out-of-lock detectors are also monitored on a continuous basis. This provides ongoing assurance of the loops' integrity during normal instrument operation.

Example Self-Test Design

A simplified schematic of one of the HP 8360's circuit boards is shown in Fig. 2. The function of this circuit is to phase-lock a 200-to-220-MHz oscillator to a reference signal and to drive a sampler with the oscillator output. The sampler down-converts the microwave signal to an intermediate frequency (IF), which is also amplified and filtered on this board. This board exemplifies many of the self-test techniques used throughout the HP 8360 series.

Testing this board required adding the ability to force the phase detector and integrator to either extreme. This was done by adding the REF ENABLE, DIVIDER ENABLE, and CROSS ENABLE control lines to the input buffers of the phase detector. These control lines allow the processor to shut off either input to the phase detector and to "cross" the reference input over to either phase detector input. Since the reference input is an output from a preceding assembly and has been verified as functional, it is used as a stimulus to test this board. By turning off DIVIDER ENABLE, the phase detector can be checked by changing the CROSS ENABLE and REF ENABLE control lines and measuring the voltage on the PHASE ABUS point.

By forcing the phase detector to its extremes, the UNLOCK digital monitoring point can be examined and its ability to detect an unlocked loop verified.

The loop integrator is tested similarly, by forcing its output to both extremes and measuring the output at the INTEGRATOR ABUS point. The CROSS ENABLE and REF ENABLE lines are used to change the phase detector to provide a stimulus.

The divider in the loop is used to set the frequency of the 200-to-220-MHz VCO. To test it, an inexpensive RC filter was added to the output to provide a simple discriminator. The loop is forced to an extreme using the CROSS ENABLE and REF ENABLE lines as before, and the divide number is varied. As the divide number changes, the frequency of the divider output pulses changes and the dc voltage on the RC filter changes in a predictable fashion.

It is important to verify the IF level since it is an output of this assembly. An inexpensive diode is used as an RF peak detector to measure the IF signal level. The IF LEVEL detector is measured with the ABUS while the microwave input to the sampler is stepped in very small frequency
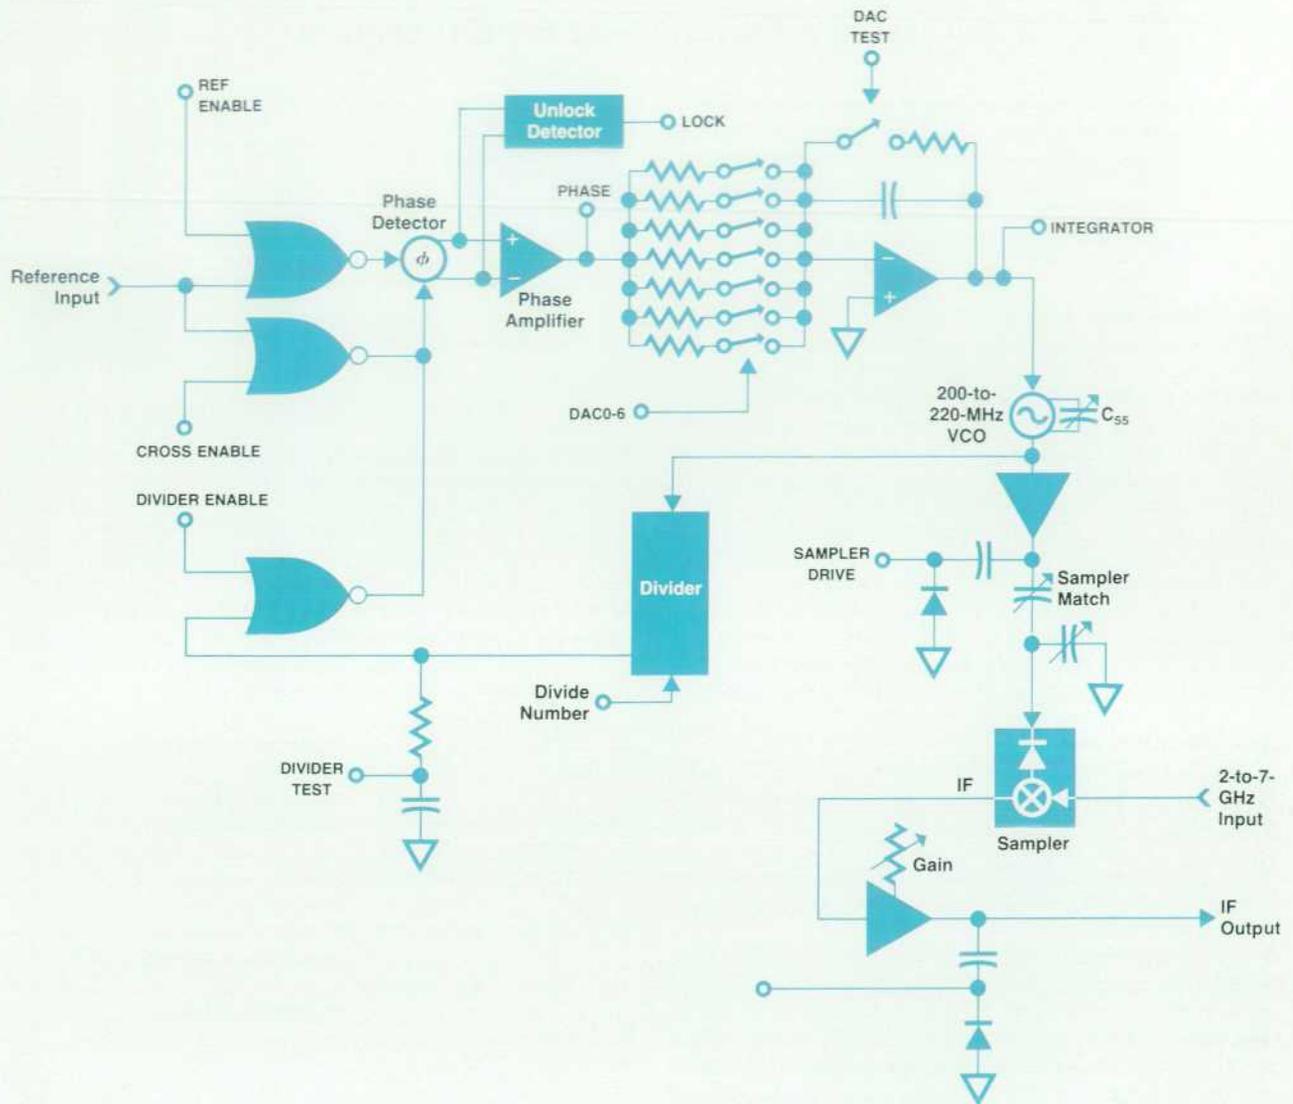
Fig. 2. *Sampler driver phase-locked loop.*

increments. The peak response of the IF LEVEL detector is noted and checked against limits.

Calibration

Although the number of manual adjustments in the HP 8360 has been minimized, further effort was applied to make the remaining adjustments easier to perform. In addition to their use in self-test, the ABUS and the internal ADC assist in making these adjustments. Activating internal firmware procedures allows some of the adjustable components on the printed circuit assemblies to be set without requiring external test equipment. These procedures set up the HP 8360 in test configurations, measure the necessary circuit nodes using the ABUS, and present the operator with a display indicating which direction to adjust the component. The following shows a display that is presented to help the operator adjust the sampler phase-locked loop:

```
*** A6 SAMPLER VCO CAL ***
Adjust C55 to center the 'needle' below:
                              |
---------------->>>  |<<<----------------
        cw->                <-ccw
                                    Done
```

A similar display is used for many adjustments. Some of these adjustments would have required external voltmeters, frequency counters, network analyzers, or spectrum analyzers. Through the use of internal firmware algorithms, these external pieces of equipment have been avoided and replaced with simple dc voltage measurements using the analog bus.

The sampler driver schematic (Fig. 2) is an example of how the same monitoring points are used for both self-test and firmware-aided adjustments. The nominal VCO tuning is provided by an adjustable capacitor, labeled C_{55}. This is set using the "needle" display shown above. The internal service firmware sets the sampler driver to a known frequency and continuously measures the INTEGRATOR ABUS

Accessing a Power Meter for Calibration

Several calibration-related features of the HP 8360 synthesized sweep oscillators require an output power measurement. The results of these power measurements are used by the HP 8360 to update internal power accuracy tables or to calibrate specified user flatness frequencies. The HP 8360 provides two methods of making these power measurements: direct HP-IB control of a power meter and an HP-IB query/response protocol.

Direct Power Meter Control

The HP 8360 can control an HP 437B power meter over the HP-IB. By removing any other controllers from the HP-IB and connecting an HP 437B power meter with an appropriate power sensor (see Fig. 1), the HP 8360 can perform its calibration functions without any operator intervention. The HP 437B was selected because of its ability to retain its own power sensor calibration factors. The HP 8360 is spared the task of editing, storing, and saving calibration factor lists. When the user activates a calibration, the HP 8360 enters a mode that steps through the appropriate operating conditions for the requested calibration. At each condition, the HP 8360 downloads the present output frequency, triggers the power meter over the HP-IB and then reads the corrected power level from the power meter. This process repeats until all the calibration conditions are measured.

Query/Response Method

There are situations in which direct control of a power meter is undesirable or impossible. In some systems it would be considered dangerous to allow an instrument to control the HP-IB if other devices are present. Many systems do not have an HP 437B power meter. In other systems it might be inconvenient to change HP-IB cabling. In these cases, another method of power determination is used.

In the query/response mode, the HP 8360 has the responsibility to configure itself for various conditions of the calibration and notify the user's program when it has settled on each condition. The user's program is notified of the output frequency at each point. The user has the responsibility of providing a power measurement routine. Since this protocol assumes no specific type of measurement equipment, the user is free to use any type of power meter. Once the user's program has measured the HP 8360's output and has corrected the reading using the supplied frequency information, the HP 8360 is notified of the output power.
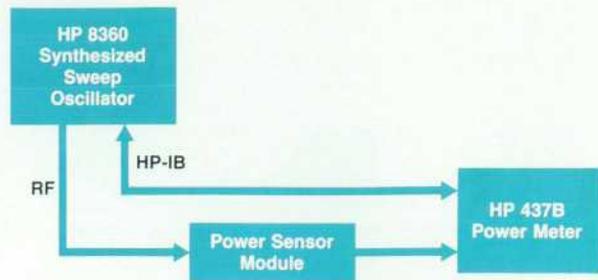

Fig. 1. *The direct control method of power calibration.*

This process repeats for each HP 8360 calibration condition until the HP 8360 determines that it has finished.

Since it is very likely that the user already has a procedure to measure power, it is very simple to control the calibration in the user's program. The following HP BASIC code lines demonstrate how to calibrate the HP 8360's internal ALC detector/logger. It is assumed that the user has already created a function called FNMeasure_power(F) that will return the power in dBm measured at the output of the HP 8360. The frequency (F) is passed to the function to allow calibration factors to be applied.

```
10   OUTPUT 719;"*RST"              ! preset the instrument
20   OUTPUT 719;"CAL:PMET:DET:INIT? IDET"
                                    ! initiate the calibration
30   INPUT 719;Cal_frequency        ! enter the calibration frequency
40   WHILE Cal_frequency>0
50       Power_out = FNMeasure_power(Cal_frequency)
                                    ! measure the power
60       OUTPUT 719;"CAL:PMET:DET:NEXT? ";Power_out
                                    ! send power to HP 8360
70       INPUT 719;Cal_frequency
                                    ! get the next frequency
80   END WHILE
90   END
```

In this fashion, the user can calibrate many of the HP 8360's power features without having to understand any of the details of the calibration process.

point. The error in the measurement is presented to the operator with the needle display as the operator adjusts C_{55}.

The IF output of the sampler driver has an adjustable-gain amplifier to set the output signal level. Since there is an ABUS monitoring point on the output, this can be a firmware-aided adjustment, too. The firmware continuously presents the measured IF output level, using the needle display, as the operator adjusts the gain potentiometer in the IF amplifier.

The loop integrator is an example of how the internal service firmware allows not only improved ease of adjustment, but improved circuit performance as well. The VCO used in the sampler driver is very nonlinear. As the VCO is tuned across its range, the voltage-to-frequency characteristic varies, which changes the loop gain of the phase-locked loop. It is important to keep the loop gain fixed, since it is critical to the HP 8360's phase noise performance. A more traditional approach to this problem would have

been to have many factory-selected diode breakpoints designed into a linearization network. This would have required operator intervention to select and adjust these breakpoints for each VCO. In the HP 8360, this problem was solved with the use of a DAC.

The DAC used to adjust the loop gain of the sampler driver is formed from seven sections of an eight-section analog switch. The main instrument processor writes the correct gain number into the DAC at the time it programs the sampler VCO frequency. By writing the appropriate number into the DAC, the loop gain is kept constant. To determine the correct numbers, a firmware routine is used that steps the VCO across its frequency range and measures the INTEGRATOR ABUS point. The VCO tuning linearity is calculated and then used to determine the correct gain calibration for each frequency. This takes less than a second to complete, and requires no operator intervention or external test equipment.

The DAC uses only seven sections of the analog switch. The remaining unused section is put to use as a DAC-test switch. As shown in Fig. 2, the DAC TEST switch and a resistor turn the integrator into a fixed-gain amplifier. This allows the INTEGRATOR ABUS point to be monitored as the DAC number is changed to perform a self-test on the DAC. This is a good example of how extra unused pieces of hardware are applied to self-test uses where appropriate.

The sampler VCO signal is applied to the input of a microwave sampler through a matching network. To ensure maximum signal transfer to the sampler, the matching network must be adjusted. This was traditionally done with a spectrum analyzer and an external sweep oscillator, but the HP 8360 performs this adjustment using internal firmware and no test equipment. The HP 8360 display prompts the operator to adjust the two matching capacitors while stepping the sampler drive loop across its frequency range. At each frequency, the voltage at the SAMPLER DRIVE ABUS point is measured and presented as a bar graph to the operator. The operator adjusts the matching until the bar graphs for all the frequencies are maximized, indicating a good power transfer across the entire sampler VCO frequency range. The display looks as follows:

```
A6 SAMPLER MATCH: adj C82,C83 for max *s
200 MHz : ***********
207 MHz : **************
213 MHz : ***************
220 MHz : *************
                                     Done
```

Troubleshooting Tools

There are other firmware routines in the HP 8360 to help knowledgeable service personnel perform more detailed troubleshooting. These include:

- Extensive status displays of hardware latch contents, frequency band information, and firmware algorithm fault detection.
- Direct display and control of any phase-locked loop frequency.
- The ability to write directly to any hardware latch or DAC in the instrument in either binary or decimal fashion. Also, a test pattern consisting of the sequence 1, 2, 4, 8, 16, 32, ... can be sent. This is much more powerful than the traditional memory peek and poke utilities in that complex devices such as nibble-loaded DACs can be written to easily with a single keystroke rather than the usual method of four or more memory pokes.
- Direct access to any of the analog bus nodes for reading.
- An extra analog bus node connected to a test point. This acts as a movable DVM probe for those nodes that are not connected to the ABUS.
- An external crystal detector that can be connected to function as a movable, limited-accuracy power meter probe.

Summary

The HP 8360's low cost of ownership results in large part from the use of internal service related firmware. The early integration of service features into the design of the HP 8360 has reduced equipment requirements for repair and calibration, simplified diagnosis of failures, and had minimal impact upon cost and reliability.

Acknowledgments

A High-Performance Sweeper Output Power Leveling System

A feedforward ALC design gives HP 8360 sweepers improved flatness, power accuracy, and modulation performance. Factory calibration techniques minimize measurement errors so as not to degrade the improved specifications.

by Glen M. Baker, Mark N. Davidson, and Lance E. Haag

THE HP 8360 FAMILY of synthesized sweepers offers greatly enhanced leveling performance over previous generations of microwave sources. The improvements are in flatness, leveling accuracy, AM bandwidth, and AM dynamic range. In addition, the ALC design provides the ability to incorporate such features as user flatness correction and self-calibration.

ALC Loop

Automatic level control (ALC) loops using diode detectors to sense RF output power have been used in many microwave sources, such as the HP 8340 synthesized sweeper. The ALC loop ensures accurate output power and AM performance. Fig. 1 is a simplified block diagram of the type of ALC loop used in the HP 8340.

This type of ALC has three main limitations regarding AM modulation performance. First, AM dynamic range (depth) is limited to that of the level detector and associated circuitry. Typically this range is far less than the range of the linear modulator. Second, the AM bandwidth is limited by ALC loop delays. Third, pulse amplitude modulation (simultaneous AM and pulse) further limits the effective AM bandwidth by the duty cycle of the pulse train.

These limitations can be circumvented by using an unleveled mode called power search. In this mode the detector voltage is sampled and held, after which the ALC loop is opened. Open-loop operation allows greater AM bandwidth or depth. However, the power search must be reexe-
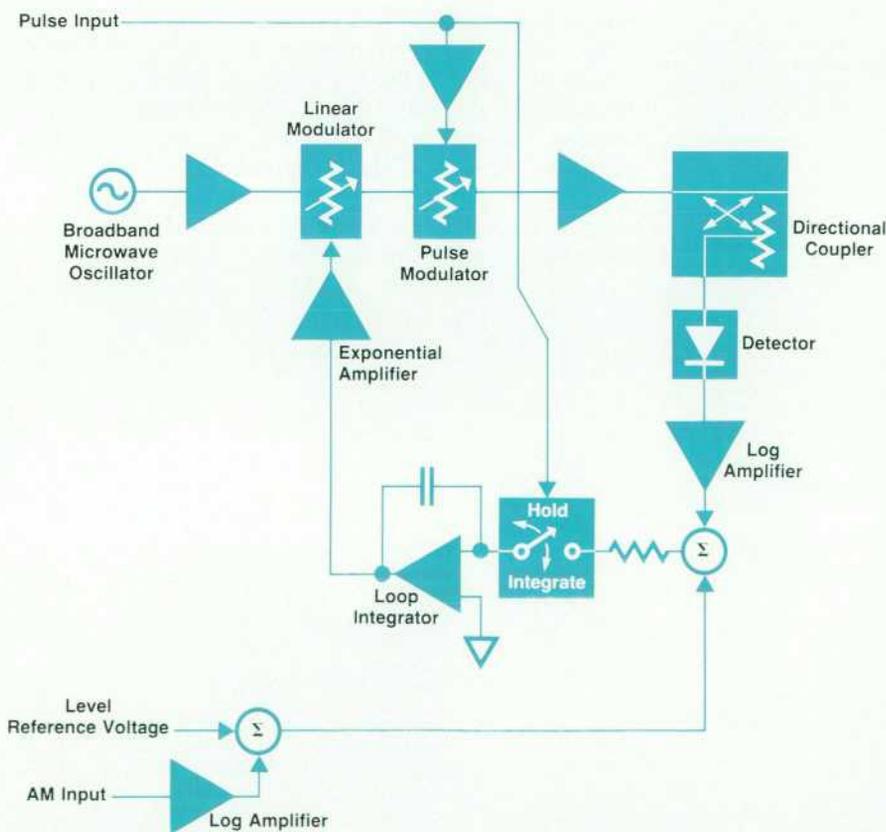

Fig. 1. *The type of ALC loop used in the HP 8340 synthesized sweep oscillator.*

cuted whenever the frequency or the power level is changed, so the response is slow. Without feedback, the power level drifts with time and the AM accuracy is poor.

Feedforward ALC

In contrast to previous designs, the HP 8360 incorporates a feedforward ALC system (Fig. 2). This is an extension of previous designs in which AM and power level information is fed directly to the modulator and to the summing node of the feedback loop, thereby providing wide AM bandwidth (up to 500 kHz) with a relatively low loop bandwidth (10 kHz or 100 kHz, selectable). Analog delay compensates for delays in the modulator and detector portions of the loop. Without this delay a spurious error signal would appear at the loop summing node, causing an undesirable peak in the AM frequency response.

Modulation inputs may cause the RF output power to be reduced to a level that cannot be accurately measured by the diode detector. In this case the range limit comparator opens the feedback loop so that the level control information in the loop integrator is held. When the RF power increases back to a measurable level, the loop is again closed. This technique produces linear AM over a range of up to 80 dB with an accurate, drift-free carrier level. In addition, greater pulse amplitude modulation agility is provided, since the AM bandwidth is independent of pulse duty cycle.

Digital ALC Reference

Frequency related variations in coupling factor, detector performance, and attenuator losses require the ALC reference to vary in a similar manner to achieve flat power at the output port. In previous ALC designs, compensation for these variations was provided by using offset and slope terms derived from the ramp generator. While simple and functional, this approach has obvious limitations in that it is incapable of compensating for nonmonotonic or nonlinear RF losses. Fig. 3 is a block diagram of this type of design.

The HP 8360 ALC loop is significantly more flexible than the previous designs. As shown in Fig. 4, digital circuitry (a DAC and an ADC) has been added to the ramp generator. This circuitry provides 1601 "buckets" for any sweep. ALC correction data, calculated by the microprocessor and stored in on-board RAM, is clocked out by these buckets to the reference DAC, thereby allowing the production of 1601 unique ALC reference voltages. In this manner the ALC reference is continuously updated during analog sweeps to provide flat, accurate output power (Fig. 5).

In the simplest case, the numbers stored in RAM are derived from instrument calibration data (measured in the factory and stored in nonvolatile EEROM). The flexibility of the digital ALC scheme, however, allows additional features such as RF slope, power sweep, and user flatness correction to be implemented by appropriately adjusting the numbers stored in the RAM. In addition, the HP 8360 is capable of reading correction data from the HP 83550 family of millimeter heads via the source module interface, thereby extending the enhanced leveling capabilities to 110 GHz.
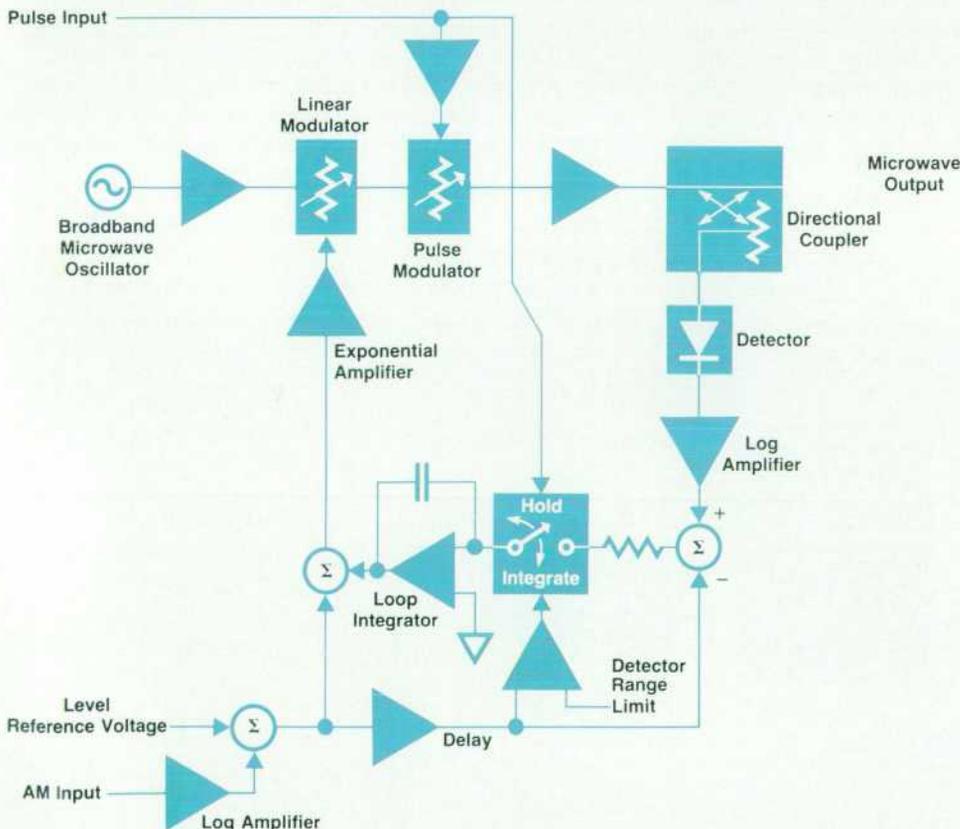

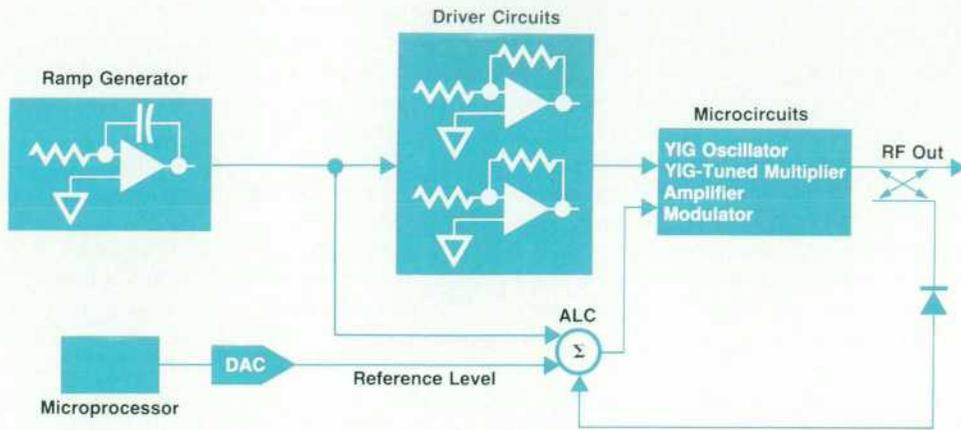**Fig. 2.** *Feedforward ALC loop of the HP 8360 synthesized sweepers.*

Fig. 3. Analog ramp generator used in previous sweeper designs.

User Flatness Correction

In practice the source is rarely connected directly to a DUT. Instead, there are intervening RF components such as cables, connectors, and couplers. Flat RF output power from the source, therefore, will rarely result in flat RF power at the input of the DUT. To achieve flat test-port power the HP 8360 incorporates the user flatness correction feature. The user is able to enter a series of frequency/correction pairs representing the loss (or gain) from the source to the test port. Once this data has been entered and user flatness correction has been enabled, the HP 8360 will automatically interpolate the supplied information and add the results to the ALC reference data in RAM, thereby producing flat power at the test port. As the user changes start or stop frequencies the user data is reinterpolated to yield appropriate corrections for any frequency span chosen.

While the frequency/correction pairs can be entered manually if desired, the HP 8360 also allows the user to specify only the frequencies of interest and measure the correction automatically by using an HP 437B power meter (see "Accessing a Power Meter for Calibration," page 22).

Power Calibration Techniques

The HP 8360's digital ALC correction, stable ALC loop, and repeatable, resonance-free 90-dB step attenuator make very accurate output power possible over the full −110-dBm-to-+20-dBm amplitude range, from 10 MHz to 40 GHz and beyond. Because measurement uncertainty must be taken into account when setting product specifications, the test system for calibrating power becomes a key element for realizing the full capabilities of the HP 8360's power control. The remainder of this article describes the factory calibration techniques and presents an analysis of the measurement errors.

High-Power Measurements. At power levels above −20 dBm, the calibration is done directly with an HP power meter, using one or more of four different HP power sensors. Software automatically selects the optimum power sensor for each power level and frequency range. Because sensor load mismatch is a significant contributor to measurement error,[1] the best sensor for the job is generally the one with the lowest input reflections. The input reflections of various power sensors were measured by an HP vector network analyzer on eight or more samples of each sensor model to determine which sensor was best for each frequency range. As shown in Fig. 6, for frequencies below 1 GHz and power levels below +10 dBm, an HP 8482A 100-kHz-to-4.2-GHz sensor is used. It has superior input match compared to microwave power sensors at the low-frequency end. For HP 8360s with APC-3.5 output connectors, the HP 8485A is used for frequencies above 1 GHz and power levels below +10 dBm. An alternative for these models would be an HP 8487A with an adapter, which may give better input match. For models with 2.4-mm output connectors, the HP 8487A is used above 1 GHz.

Finally, an HP 8485A that is specially calibrated with a 10-dB attenuator permanently attached is used to measure power levels above +10 dBm. At present, only models
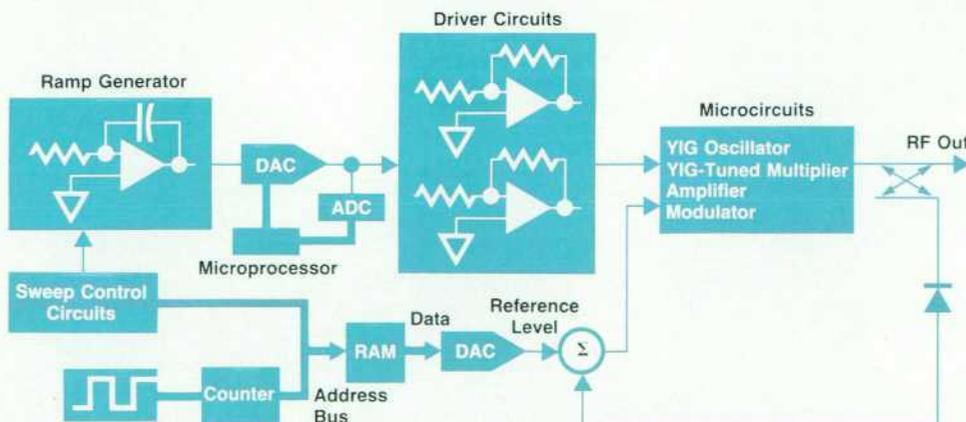

Fig. 4. The HP 8360 ramp generator adds digital control circuitry to provide 1601 ALC reference levels during a sweep.
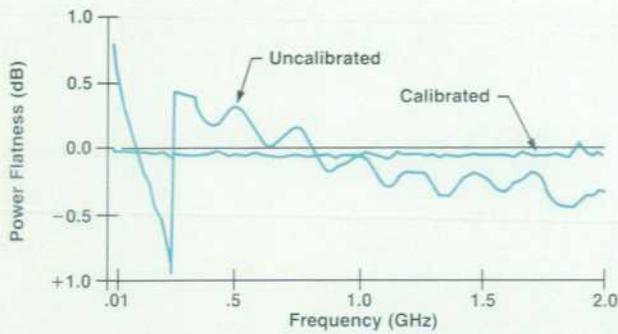
Fig. 5. *Swept power flatness of the HP 8360.*

with 3.5-mm connectors achieve output levels significantly above +10 dBm. There is no reason the same approach can't be taken with the HP 8487A. Measurements show that the sensor with pad has better input match than the unpadded sensor, especially at low frequencies. The useful range of the HP 8485A is extended down to 10 MHz from the specified lower limit of 50 MHz. The main reasons for using the pad are to prevent the sensor's operating in a nonlinear power range and to prevent overload and possible damage from some HP 8360 models, which can produce power levels above +25 dBm at certain frequencies.

Because the HP 8360 has a very stable ALC system, very good incremental power accuracy, and enough correction points to remove virtually all the flatness errors, a unit that is measured with the same set of power sensors that were used to calibrate it will show very small errors, typically <0.1 dB worst-case from −10 to +10 dBm.[2] This level of performance has been measured on HP 8360s after several months of worldwide travel, when returned to the factory and measured on the original test system. When the performance is measured by different meters with different power sensors, the results are significantly worse. This is because the input match and calibration error of the second

meter are different from those of the first. The errors measured are a random combination of the errors of both meters. Another way to look at this is that the accuracy of any power measurement device can be transferred to the HP 8360 with <0.1 dB error.

Low-Power Measurements. For power levels less than −20 dBm, an attenuator calibration system was developed. This system uses an HP 70000 Series modular spectrum analyzer to translate the frequency being measured to a 21.4-MHz IF (see Fig. 7). This IF is then measured by an HP 8902A measuring receiver, which delivers state-of-the-art relative power measurement accuracy.[3,4] The algorithm relies on the HP 8360's power accuracy at high power levels and on the spectrum analyzer's excellent IF linearity to get an accurate measurement. At each frequency to be measured, the HP 8360 is programmed to a power level that has been calibrated with a power meter and is therefore known to have nearly power meter accuracy. The spectrum analyzer is set up with enough input attenuation to keep this reference power within its linear range, and is left in the same state while the HP 8360's output power is set to lower levels. The relative power drop is measured by the HP 8902A at the IF to determine the amount of correction required.

Several measures are taken to minimize error. A 20-dB pad is used at the HP 70908A input to improve input match, thus minimizing errors caused by changes in source match when the attenuator is switched in the HP 8360. For the waveguide mixers, 10-dB pads are used with HP 365 Series isolators to achieve good input match while improving sensitivity. Hardware arrangements and shielding are used to minimize interference of the preattenuated signal with the signal being measured. IF filtering and small offset frequencies are used to prevent interference from residual signals in the spectrum analyzer and other sources in the test system.

The power is not actually measured below −90 dBm
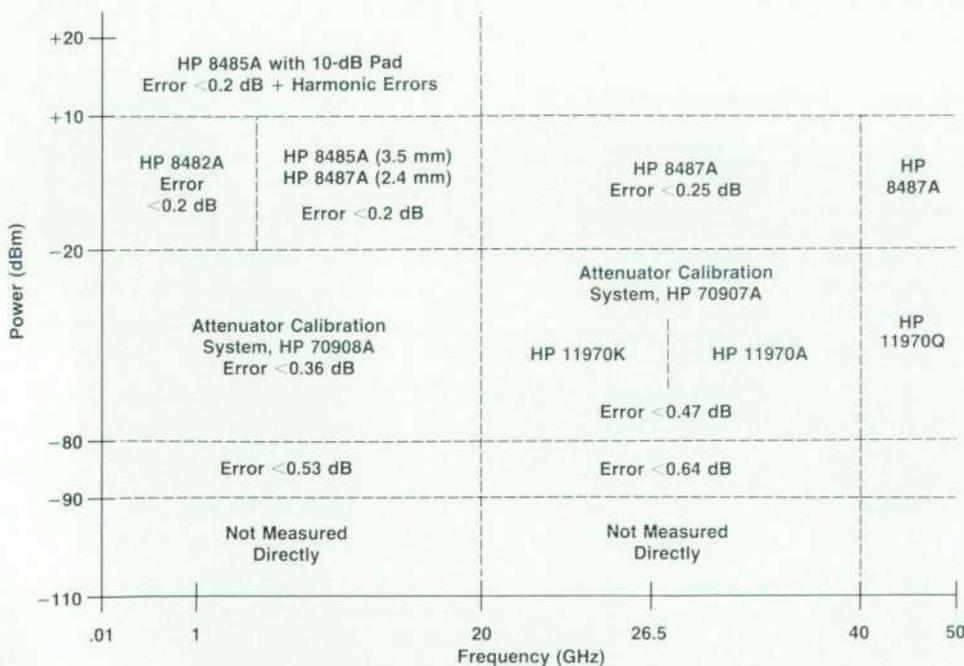
Fig. 6. *A map of the equipment used for factory calibration of the HP 8360 output power and the measurement uncertainty.*
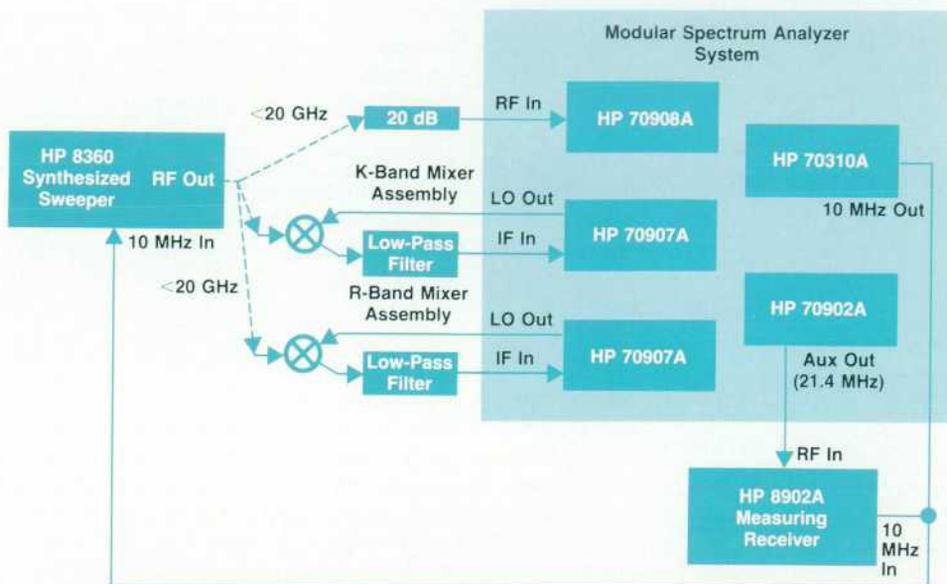
Fig. 7. *Attenuator calibration system for calibrating the HP 8360 output power at low power levels.*

Mismatch Error Calculation for Relative Power Measurements with Changing Source Match

A relative power measurement is a measurement of the ratio of two power levels. In calibrating the HP 8360 synthesized sweeper at low power levels, a reference power measurement is taken with the HP 8360 programmed to an accurately known power level. Then the HP 8360's internal attenuator is switched to reduce the output power and calibrating power measurements are taken at various power levels.

An error is introduced because the source match changes as the attenuator is switched. An analysis was done to determine the magnitude of the error. Fig. 1a shows the conceptual block diagram of the measurement and Fig. 1b is the flow graph. In the following equations, a and b are source and load power, respectively. The s_{ij} are s-parameters and Γ is reflection coefficient.

For a single measurement,

$$\frac{b_L}{a_S} = \frac{s_{21c}s_{21p}}{1-\Delta}$$

$$1-\Delta = A\Gamma_S + B$$

$$A = -s_{11c} - s_{21c}^2 s_{21p}^2 \Gamma_L - s_{21c}^2 s_{11p} + s_{11c}s_{22c}s_{11p}$$
$$+ \Gamma_L s_{11c}s_{22p} - \Gamma_L s_{11c}s_{22c}s_{11p}s_{22p}$$

Variance= 1.05901569097E-5
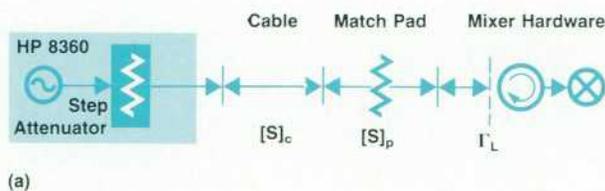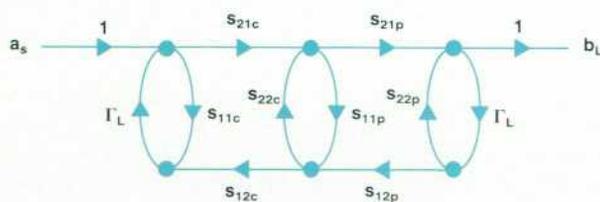
Fig. 1. *(a) Conceptual block diagram of a relative power measurement with changing source match. The source match changes as the attenuator is switched. (b) Flow graph.*
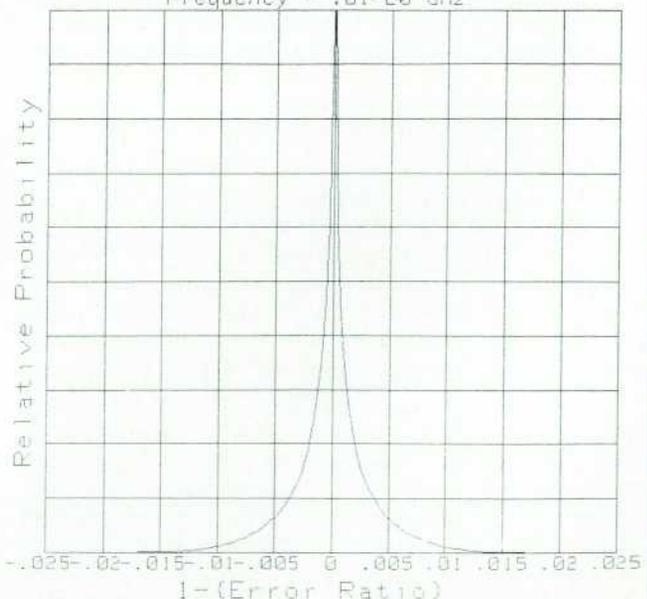
Fig. 2. *Monte Carlo analysis histogram for HP 8360 calibration from 0.01 to 20 GHz. The standard deviation is 0.325% and the error is less than 1.5% at the 99% confidence level.*
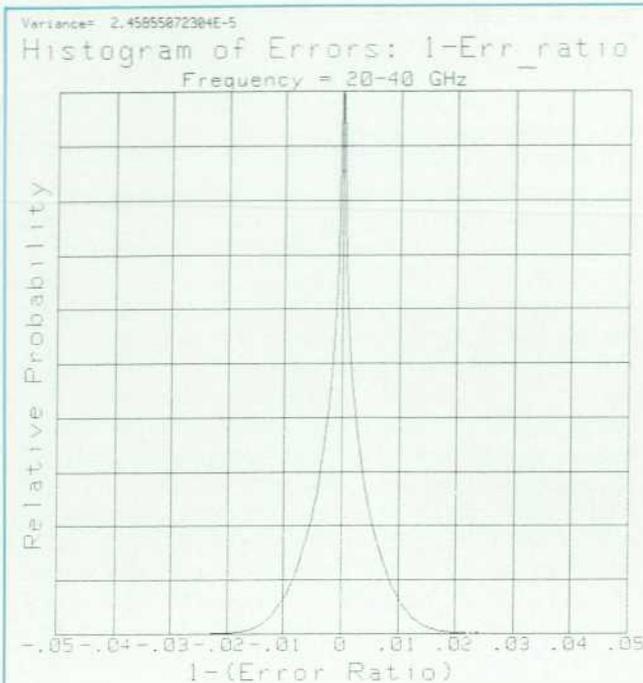
Fig. 3. *Monte Carlo analysis histogram for HP 8360 calibration from 20 to 40 GHz. The standard deviation is 0.496% and the error is less than 2.25% (99% confidence).*

$$B = 1 - s_{22c}s_{11p} - s_{22p}\Gamma_L - s_{21p}{}^2\Gamma_L s_{22c} + \Gamma_L s_{11p}s_{22p}s_{22c}.$$

All of the terms in A and B are constant during a relative power measurement.

For the relative power measurement, the error ratio is

$$\frac{\dfrac{b_{L1}}{b_{L2}}}{\dfrac{a_{S1}}{a_{S2}}} = \frac{(1-\Delta)_2}{(1-\Delta)_1} = \frac{A\Gamma_{S2}+B}{A\Gamma_{S1}+B}$$

in linear terms. To calculate the distribution of errors, Γ_{S1}, Γ_{S2}, s_{11c}, s_{11p}, s_{22c}, s_{22p}, and Γ_S were all modeled as random variables with uniform phase and normal magnitude. The models were developed from measured data and specifications on the components. The transmission parameters were modeled as constants with a nominal value. In the case of a direct connection to a mixer (no cable), the cable reflections were put at zero and the cable transmission at unity.

The Monte-Carlo analysis program picks a value for each random magnitude and phase, calculates the result, and places it into a histogram. This process is repeated several hundred thousand times, until the histogram has a smooth appearance. Figs. 2 and 3 show typical histograms.

Taking the error ratio beyond which the histogram is zero as the maximum error yields an error of about ±1.5% for the 0.01-to-20-GHz frequency range and ±2.25% for the 20-to-40-GHz frequency range. Thus the total power errors used to calculate the measurement uncertainty (see accompanying article) are 3% and 4.5%, respectively.

because of system sensitivity limitations. Attenuator input power of 0 dBm is used, and 90 dB is the maximum attenuation. Power accuracy below −90 dBm is guaranteed by the ALC incremental accuracy. A major advantage of this test system is that the spectrum analyzer is available to measure other aspects of performance such as harmonics and spurious signals using the same setup. Similarly, the other measurements the HP 8902A can make are also available, such as AM and FM performance and noise.

Three identical systems exist in the factory, and low-level accuracy has been cross-verified between the different systems. The systems have been found to disagree by less than 0.1 dB down to −80 dBm, and by less than 0.3 dB down to −90 dBm. These results apply over the full 0.01-to-40-GHz frequency range. An independent verification using a diode power sensor (HP 8487D) down to −60 dBm showed less than 0.3 dB of error, which is mostly attributable to calibration and mismatch differences between the HP 8487D and the power sensor used to calibrate the HP 8360 at high power. These are measured results. Specified performance accounts for the calculated errors of the test system, as well as known environmental performance changes over the specified 0-to-55°C temperature range.[5]

Error Analysis

To determine the measurement uncertainty of the calibration systems, an error analysis was done using statistical models for the error contributors. For errors that add randomly, the standard RSS (root sum of squares) approach was used. Systematic (repeatable but unknown) errors were added to the RSS totals. For mismatch errors, statistical models were developed for the complex reflection coefficients and a computerized Monte Carlo analysis was run to generate the distribution of errors from mismatch. In every case, the phase of the reflection was modeled as a uniformly distributed random variable over 0 to 360 degrees. The magnitude of the reflections was modeled as a normally distributed random variable, with the mean and variance chosen from measured data and specifications. All the variables were assumed to be statistically independent in the Monte Carlo analysis.

High Power. For the power meter error calculation, the independent additive terms that are combined using the RSS method are reference calibration error, instrumentation error, and calibration factor error.[1] The systematic errors added to the RSS total are the mismatch error from the Monte Carlo analysis, and sometimes the nominal loss of the adapter used to connect the power sensor to the HP 8360. The adapter loss's inclusion depends on whether the application is for a male or female connection. If the test connection has a male connector, the HP 8360 can be assumed to be calibrated at the output of the adapter, which is shipped with the HP 8360, and the adapter loss does not need to be added. For female test connections, adapter loss should be included. A good model for adapter loss for HP precision adapters is:

$$\text{Loss in dB} = \frac{5}{300}(1 + 0.1(\text{Frequency in GHz})).$$

The total error is given by:

$$e_{\text{Total}}(\%) = \sqrt{\Delta K_b{}^2 + e_R{}^2 + e_I{}^2} + e_{\text{mismatch}} + \text{Adapter Loss}.$$

The error resulting from this calculation is <0.2 dB below 20 GHz and <0.25 dB below 40 GHz excluding the adapter loss. The reference error is 0.9%, the instrumentation error is 0.5%, the calibration factor error (ΔK_b) is 3.0%, and the mismatch power error is 1.5% below 20 GHz and 2.7% below 40 GHz. These mismatch errors are much smaller than a worst-case analysis would produce, but are beyond the 99th percentile of error calculated based on measured data.

Another error that must be considered when testing a version of the HP 8360 at a power level where significant (> -40 dBc) harmonics exist, is the error in the carrier level caused by inclusion of the harmonic energy in the power measurement. This is significant only in the HP 83623A, the HP 83624A below 12 GHz and above $+10$ dBm, and all versions below 1.8 GHz and above $+10$ dBm. This is reflected in the less stringent HP 8360 power accuracy specifications above $+10$ dBm.

Low Power. For power levels below -20 dBm, the error is larger because of the additional measurement error of the attenuator calibration system. The main contributors to the error in this case are reference power calibration error as calculated above, HP 8360 power error as measured by the calibrating meter, microwave mismatch error caused by a change of source match between the reference and calibrating measurements, and HP 8902A related errors at the IF, including RF mismatch, noise, linearity, range-to-range errors, and others.[3]

The HP 8902A errors are taken directly from the HP 8902A specifications, while HP 8360 power error is taken from production data on power accuracy verifications. The mismatch error distribution was calculated using Monte Carlo simulation, as it was for the simple power meter measurement (see box, page 28).

The total measurement uncertainty is calculated using an RSS summation of all the applicable errors for the given frequency range, attenuator stage, and HP 8902A setting. Most of the errors are systematic in nature, that is, they are very repeatable but unknown. The only exception is HP 8902A noise error. It is assumed that all of the systematic errors are independent. The total systematic error is calculated using the RSS method. The random error from noise is added to the systematic error to form the total error listed in Fig. 6. Notably, the random error applies only for the 90-dB stage. The mismatch error is included in the RSS total in this case because it is composed of two measurements, unlike the mismatch error for the power meter measurement. This causes its error distribution to be closer to normal, with large errors very unlikely.

Summary

There is >99% confidence that the absolute error in a single measurement is less than the amounts shown in Fig. 6 if the assumptions made in the analysis are true. These assumptions are that the HP 8360 has no flaws in its RF output cabling that generate large reflections, that the HP 8360 has been calibrated at high power, and that the microwave test cabling, pads, isolators, and mixers are not generating abnormal reflections.

The HP 8360 power accuracy specifications include two main components: measurement uncertainty and environmental variations over the 0-to-55°C temperature range.[2] The small errors contributed by the measurement system allow the specifications to reflect the true performance of the HP 8360's power control system more closely.

Using the HP 8360's user flatness correction feature, the user can achieve results much better than the HP 8360 specifications at any point where external leveling hardware can be inserted. If the environment is controlled, the source match of the leveling setup is good, and a high-quality power meter is used for calibration, accuracy even better than the results presented here can be achieved. This is because the HP 8360's internally leveled source match is not as good as what can be achieved with external leveling configurations, such as a two-resistor power splitter.[6]

Acknowledgments

References

1. *Fundamentals of RF and Microwave Power Measurements,* Hewlett-Packard Application Note AN 64-1.
2. *HP 8360 Family Technical Data Sheet,* Hewlett-Packard Publication no. 5952-8087.
3. *Accurate Signal Characterization at Millimeter-Wave Frequencies,* Hewlett-Packard Product Note 8902A-2.
4. M. McNamee, "Integration Arms Receiver With Low Noise Floor," *Microwaves & RF,* June 1983.
5. S.L. Read and T.R.C. Read, "Statistical Issues in Setting Product Specifications," *Hewlett-Packard Journal,* Vol. 39, no. 3, June 1988, pp. 6-11.
6. R.A. Johnson, "Understanding Microwave Power Splitters," *Microwave Journal,* December 1975.

A 0.01-to-40-GHz Switched Frequency Doubler

This microcircuit doubler has a passthrough mode for 0.01 to-20-GHz input signals and a doubler mode for 20 to 40 GHz. An integrated RF switch changes modes. Slotline filters reduce spurious outputs to −40 dBc or less.

by James R. Zellers

THERE IS SIGNIFICANT DEMAND for microwave swept sources that span very broad frequency ranges such as 0.01 to 40 GHz. Many of the applications require a relatively spurious-free output signal, but this has been a limitation of traditional doubler structures. To meet such requirements, whether the swept source is used in a stand-alone application or as a stimulus for scalar and vector network analyzers, a switched doubler microcircuit assembly was developed.

This microcircuit extends the frequency range of existing 0.01-to-20-GHz sources up to 40 GHz. Spurious output in the 0.01-to-20-GHz band is limited by the input signal, and is below −40 dBc in the 20-to-40-GHz band. This microcircuit was first used in the HP 83597A 0.01-to-40-GHz plug-in for the HP 8350 sweeper family and scalar network analyzer systems, and in the HP 8516A 0.045-to-40-GHz coaxial test set for the HP 8510 network analyzer family. The latest use is in the HP 8360 family of synthesized sweepers.

Fig. 1 shows the block diagram of the switched doubler microcircuit assembly. In the 0.01-to-20-GHz frequency band, the doubler is disabled and the input signal passes through diode switch S1 to the main output line. In the 20-to-40-GHz doubler band, an input signal in the 10-to-20-GHz range is routed through diode switch S1 to a preamplifier and a power amplifier and is applied to a full-wave frequency doubler. The resulting 20-to-40-GHz signal is then coupled back into the main output line. A broadband 0.01-to-40-GHz sweep results from sequentially sweeping from 0.01 to 20 GHz in the passthrough mode and from 20 to 40 GHz in the doubler mode.

Amplifiers and Doubler

In the 20-to-40-GHz band, the 10-to-20-GHz input signal is first amplified by a preamplifier and a power amplifier to a level sufficient to drive the doubler. The preamplifier consists of three GaAs FET amplifier stages with short sections of low-impedance line used as interstage matching. The gain of this preamp was designed to compensate for the frequency roll-off in the other system components (doubler, switch, power amplifier, etc.), and is approximately 5 dB at 10 GHz and approximately 15 dB at 20 GHz. The power amplifier consists of a two-stage quadrature amplifier driving a second two-stage quadrature amplifier. Each stage of the first quadrature amplifier has two GaAs FET devices, while each stage of the second has a single FET. This power amplifier provides about 9 dB of gain at more than +23 dBm output power, and was originally designed for millimeter-wave source modules.[1]

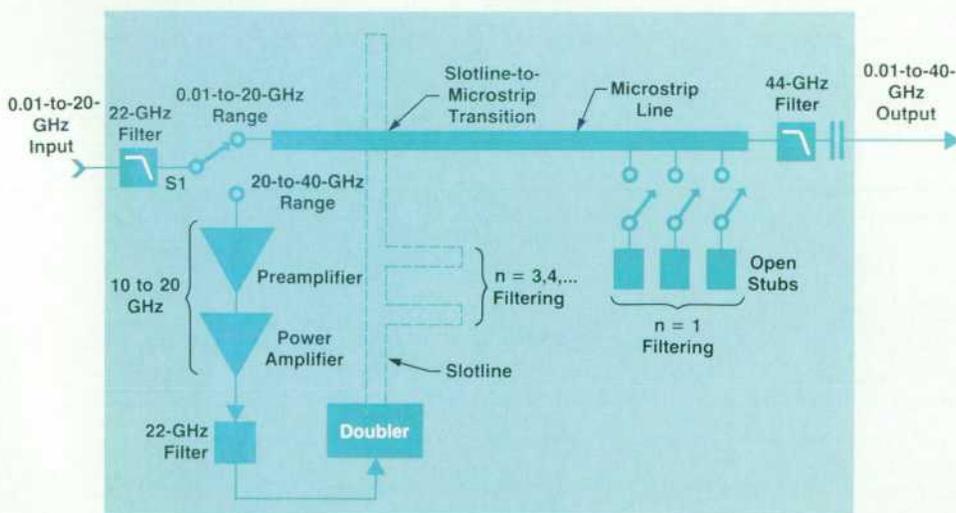The frequency doubler itself consists of two beam-lead diodes mounted across a quarter-wave section of coplanar

Fig. 1. *Block diagram of the switched doubler microcircuit.*
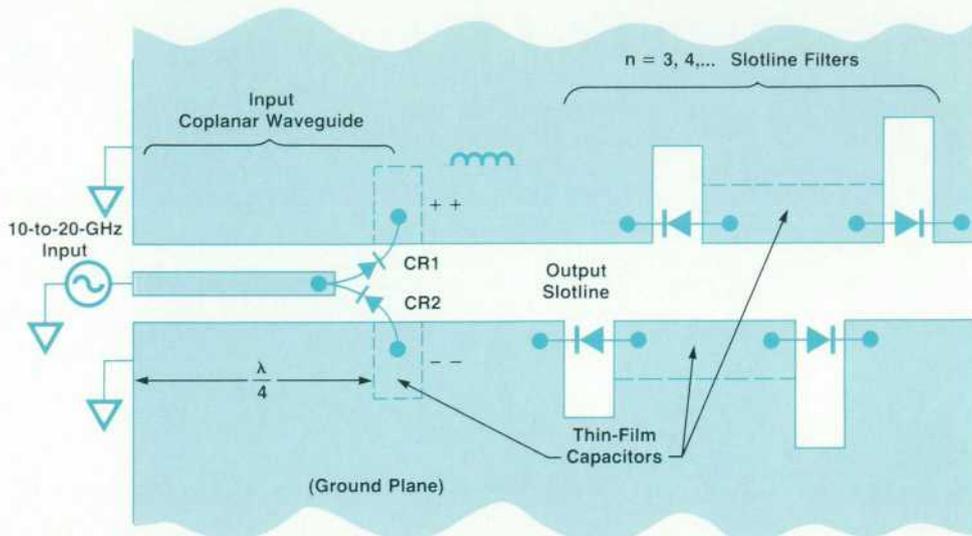
Fig. 2. *Doubler and slotline filter structure.*

waveguide transmission line, as shown in Fig. 2. The input signal drives the diodes through the coplanar waveguide. The output signal is propagated in slotline, a transmission line that is essentially a narrow gap in a ground plane. During one half cycle of the input signal, when the potential of the center conductor of the coplanar waveguide is positive with respect to its ground plane, diode CR1 conducts. This leaves a net positive charge on the upper ground plane with respect to the lower, and therefore a potential difference appears across the output slotline. During the other half cycle, when the potential of the center conductor of the coplanar waveguide is negative with respect to its ground plane, diode CR2 conducts. This leaves a net negative charge on the lower ground plane with respect to the upper, and therefore a potential difference appears across the output slotline that is of the same polarity as during the first half cycle. This results in a signal in the output slotline that is a full-wave rectified version of the input signal, and is therefore doubled in frequency.

The output signal can propagate in two directions in the slotline: back toward the source and out toward the load. The quarter-wave section of coplanar wavequide ensures that the main propagation direction is out toward the load, and provides balance for the fourth harmonic of the input frequency.

A dc bias voltage is applied to the diodes to extend the dynamic range of the doubler and to provide an adjustment for optimum balance. Thin-film bypass capacitors are used to minimize parasitics. A combination of self and fixed bias is used. This extends low-level dynamic range and avoids limiting effects at higher power levels.

Slotline-to-Microstrip Transition

In the 20-to-40-GHz band, the slotline output of the doubler is coupled back into the main output line by a slotline-to-microstrip transition that was described by Gupta and others.[2] An RF switch is incorporated into the transition, as shown in Fig. 3. The slotline crosses the microstrip line at a right angle, extends a quarter wavelength past the crossing point, and is short-circuited at its end. The microstrip line also extends a quarter wavelength past the crossing

point and is open-circuited at its end when diode CR1 is off. This structure allows the 20-to-40-GHz signal from the doubler to be coupled into the microstrip output line when diode CR1 is biased off. In the 20-to-40-GHz band, diode CR2 is biased on to allow the 10-to-20-GHz input signal to drive the preamplifier.

In the 0.01-to-20-GHz band, diode CR1 is biased on and CR2 is biased off, thus routing the input signal directly to the microstrip output line rather than through the amplifiers and doubler. Additional diodes are placed across the slotline where it crosses the microstrip line, and are biased on in this band to short-circuit the slotline and avoid ground-plane discontinuities.

One advantage of this switch design is that the 20-to-40-GHz signal is not passed through any series diodes, which can be very lossy at high frequencies. Another advantage of this technique is that only one switch is required to

Fig. 3. *Slotline-to-microstrip transition. Including the RF switch (CR1 and CR2) minimizes loss.*

switch both input and output, thus reducing signal loss over a conventional double-pole double-throw design.
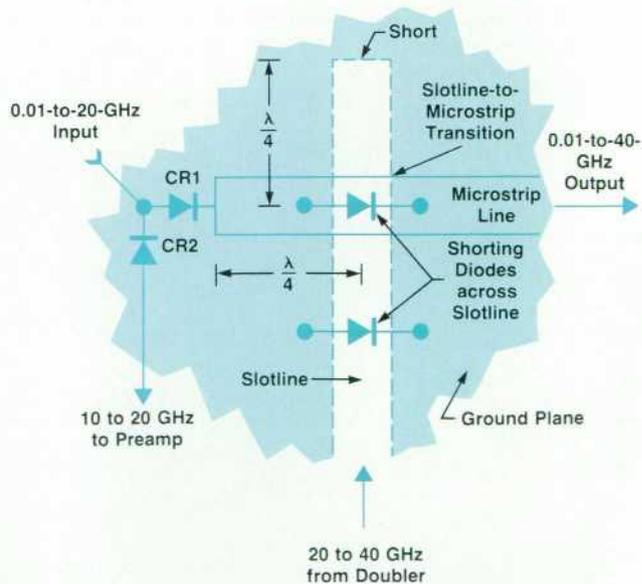
Spurious Reduction

The full-wave rectified signal at the doubler output contains the desired frequency component at twice the input frequency as well as unwanted spurious signals at multiples n = 1, 3, 4, 5, ... of the input frequency. Filtering reduces the level of these spurious signals to less than −40 dBc relative to the desired output. To realize this filtering over an octave frequency range, the 20-to-40-GHz band is divided into three sub-bands: 20 to 25 GHz, 25 to 32 GHz, and 32 to 40 GHz. Diode switching is employed to select filters for each sub-band.

To filter the fundamental feedthrough signal (n = 1), a three-stub notch filter was designed for each sub-band. Each filter consists of three open-circuited microstrip stubs that are switched onto the main output microstrip line, as shown conceptually in Fig. 1. Each stub is a quarter wavelength long at its desired rejection frequency, and a half wavelength long at its desired pass frequency. The exact lengths of the stubs and the spacings between them were computer-optimized for maximum fundamental feedthrough rejection and minimum insertion loss over the sub-band of interest. Two of the stubs are shared between adjacent filters to reduce the total number of stubs required for three filters from 9 to 7.

The higher-frequency spurious products (n = 3, 4, ...) are attenuated by switched stub filters in slotline. In the 20-to-25-GHz sub-band, the filter rejects 30-to-50-GHz signals. In the 25-to-32-GHz sub-band, the filter rejects 37.5-to-50-GHz signals. These filters consist of stubs connected in series with the slotline as shown conceptually in Fig. 2. The stubs are about a half wavelength long at the desired pass frequency. The actual lengths of the stubs and the spacings between them are computer-optimized for maximum spurious rejection and minimum insertion loss over the sub-band of interest. Diodes are used to switch

the filters in and out; when a diode conducts the corresponding stub is bypassed. An advantage of slotline in this application is that since its characteristic impedance is relatively high (about 100 ohms), the diodes contribute negligible loss.

In addition to the switched slotline filters, out-of-band spurious signals above 40 GHz are further attenuated by a fixed 11-element low-pass filter in the output line having a cutoff frequency of approximately 44 GHz.

All the filters described here were modeled, designed, and optimized using a predecessor of the HP 85150 microwave design system linear simulator. The accuracy of the models contributed to the good agreement that was obtained between calculated performance and measured results, thus reducing the time required to complete the final design.

Package Design

To reduce size, line loss, and cost, it was decided to integrate the major system components (switch, amplifiers, doubler, and filters) into one microcircuit package. Fig. 4 shows the general layout of the switched doubler microcircuit assembly. Thin-film circuits are used throughout. Small circuits are epoxied directly to the floor of the machined package. A large sapphire substrate that contains the doubler and the switched filters is clamped in place. Separate carriers hold the preamp and power amplifier circuits. Extensive environmental testing was completed during both the prototype and production-introduction phases to ensure that the design meets all goals for environmental conditions and reliability.

Acknowledgments

I would like to thank the many people who contributed to the design and success of this project. Dave Wilson designed the doubler, breadboarded it with a slotline-to-microstrip transition, and demonstrated feasibility. Dale Albin designed the 10-to-20-GHz power amplifier and contributed the idea for the switched slotline filters. Eric Heyman designed the 10-to-20-GHz preamp and the 44-GHz low-pass filter, developed an automated test system to measure power and harmonics during the entire development process, coordinated prototype assembly and testing, and provided lab support for production introduction. Thanks also to Arlen Dethlefsen and Bob Cannon for their support, and to Joy Flatebo, Gary Gilbreth, Dexter Yamaguchi, Doug Lash, and many other people involved with prototype fabrication, evaluation, and introduction of this microcircuit into production.

References

1. R.D. Albin, "Millimeter-Wave Source Modules," *Hewlett-Packard Journal*, Vol. 39, no. 2, April 1988, pp. 18-25.
2. K.C. Gupta, Ramesh Garg, and I.J. Bahl, *Microstrip Lines and Slotlines*, Aertech House, 1979, p. 235.
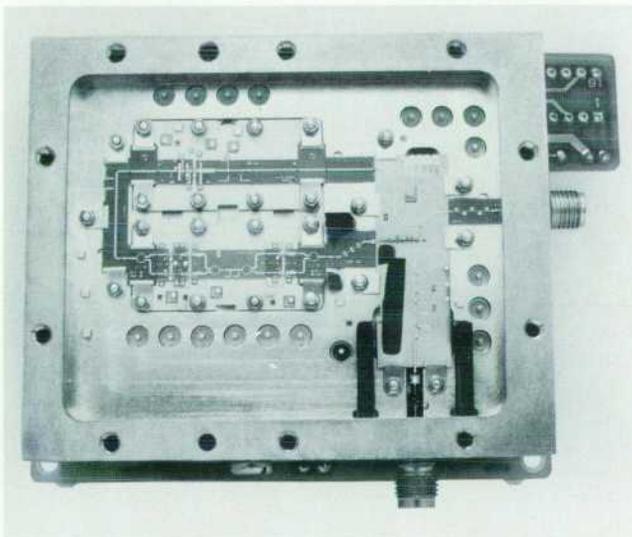
Fig. 4. *Interior of the switched doubler microcircuit. All functions are integrated into one package to minimize cost, size, and transmission loss.*

A High-Speed Microwave Pulse Modulator

This optional fast pulse modulator uses an unequally spaced diode topology to achieve a wide bandwidth and a high on-off ratio without resorting to performance-limiting diode saturation.

by Mary K. Koenig

FAST PULSE MODULATION, Option 006 for the HP 8360 synthesized sweep oscillators, offers improved pulse modulation performance over the standard instrument block diagram. In this configuration, the instrument has an additional pulse modulator at the output of the RF deck. Pulse modulating after the YIG multiplier substantially improves performance. A new wideband modulator is the key to this improved performance.

Standard Instrument

In the standard HP 8360, the modsplitter microcircuit performs the pulse modulation (see article, page 36). A linear modulator and a pulse modulator cover the 2-GHz-to-8-GHz frequency range. These are two-octave modulators. They consist of equally spaced diodes shunted across a microstrip transmission line. The diodes are turned on to attenuate the signal and turned off to pass it. Ideally, these circuits have a two-octave attenuation response centered on the frequency where a quarter wavelength equals the spacing between the diodes.

In the standard HP 8360, pulse modulation occurs before the YIG multiplier, and the bandwidth limitations of the YIG multiplier typically result in pulse rise and fall times of about 15 nanoseconds.

Wideband Modulators

Often applications requiring more than a two-octave bandwidth use the equally spaced diode topology. The spacing between the pin diodes is set for a quarter-wavelength maximum-attenuation response at the high-frequency end of the band. To achieve greater attenuation at the low-frequency end of the band, the diodes are driven into saturation and have less shunt resistance. The disad-

vantage of this scheme is that the rise and fall times of the pulsed waveform depend on the amount of stored charge in the i layer of the pin diodes. As the diodes saturate, the stored charge in the i layer increases and the modulator takes longer to change state.

The new modulator does not depend on diode saturation. Instead, the diodes are spaced unequally so that the minimum attenuation of a given set of diodes is cancelled by other sets.[1]

High-Speed Wideband Modulator

Fig. 1 shows the block diagram of the fast pulse modulation option. At the input and the output of the modulator are high-power amplifiers that provide power, gain, and good reverse isolation characteristics. These amplifiers are monolithic microwave integrated circuits (MMICs) developed by the Hewlett-Packard Microwave Technology Division.

Fig. 2 is a photograph of the fast pulse modulator.

The modulator circuit is constructed on a 0.010-inch-thick alumina substrate with laser-cut holes for precise diode placement. Chip mesa pin diodes are epoxied to the baseplate of the package in these laser-cut holes. Gold mesh is used to connect the diode anode to the circuit pattern. A bias structure and a video feedthrough filter are added to the circuit. The bias circuit extends from the modulator to a 0-pF dc feed. The bias circuit consists of a discrete five-pole low-pass filter with a cutoff frequency of 700 MHz. This bias circuit will pass a fast pulse drive waveform and reject RF energy in the operating bandwidth of the assembly. The video filter is a seven-pole discrete high-pass filter. The series capacitors are fabricated with microstrip technology on sapphire substrate. Shunt inductors are air-
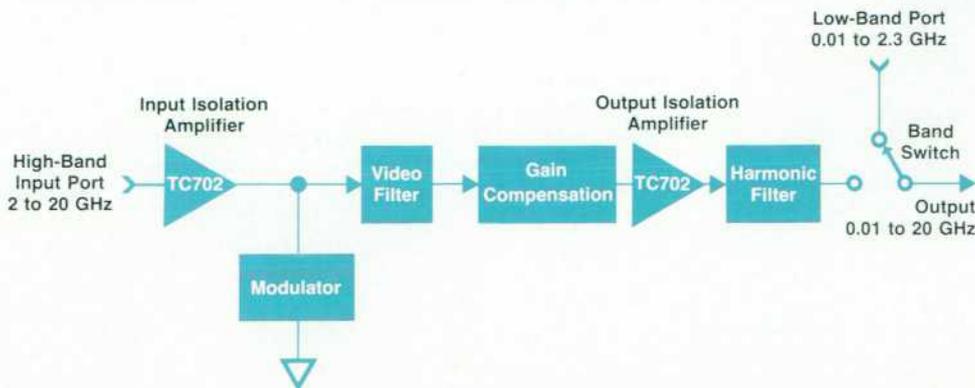
Fig. 1. *Block diagram of the fast pulse modulator.*
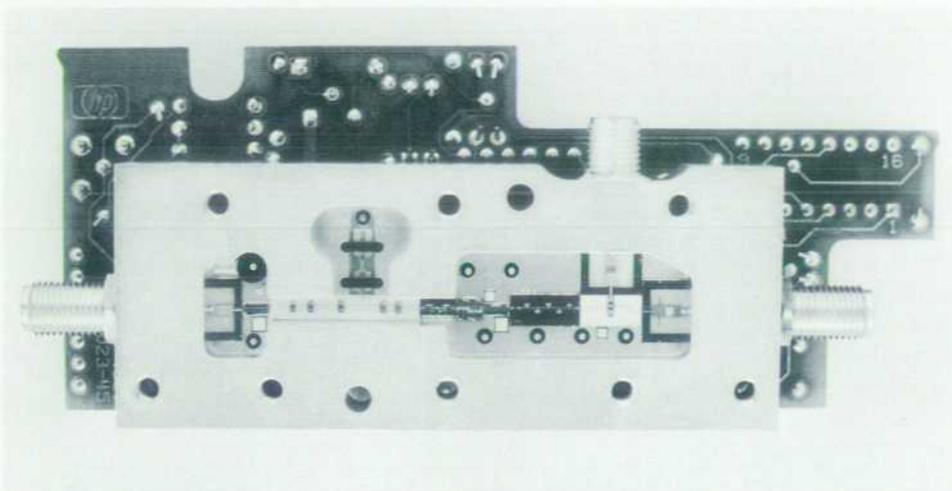
Fig. 2. *Photograph of the fast pulse modulator.*

core coils wound with 0.0015-inch gold wire. A frequency dependent attenuation pad follows the video filter to compensate for the gain slope of the component. Following the output MMIC amplifier, a 22-GHz low-pass filter improves the above-band harmonic response of the component.

An electronic switch at the output of the modulator multiplexes the low-band 10-MHz-to-2.35-GHz frequency range and the high-band 2.35-GHz-to-20-GHz frequency range.

Modulator Performance

The fast pulse modulator's performance was characterized at the microcircuit level with the following results.

On-Off Ratio. Fig. 3 shows the on-off ratio for the pulse modulator. The on-off ratio is 95 dB from 2.0 GHz to 20 GHz.
Output Power. Fig. 4 shows the typical output power of the component.
Gain. Fig. 5 shows the typical gain of the component.
Pulse Performance. Fig. 6 shows the pulsed waveform created by the five-diode pulse modulator. The modulator's rise and fall times are typically better than 2 nanoseconds.

Conclusion

The HP 8360 with Option 006 offers pulse modulation with high power, fast rise and fall times, wide dynamic range, and broad bandwidth. The heart of the pulse modulation system is a new modulator design. The power, gain,

Fig. 3. *Fast pulse modulator on-off ratio.*
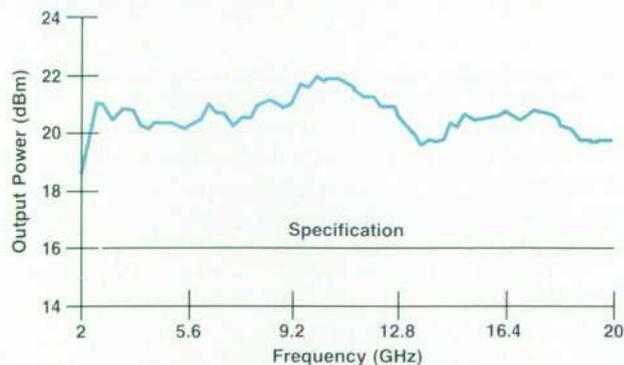
Fig. 4. *Fast pulse modulator power output.*
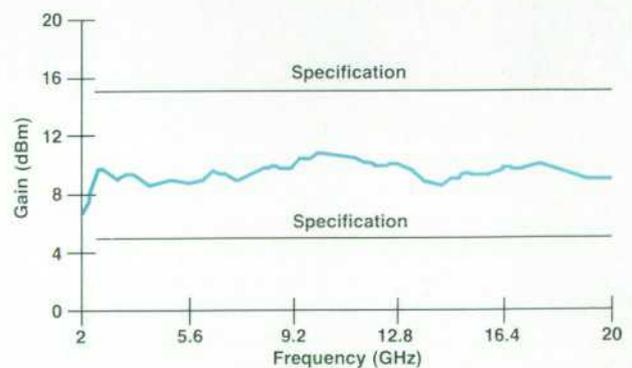
Fig. 5. *Fast pulse modulator gain.*
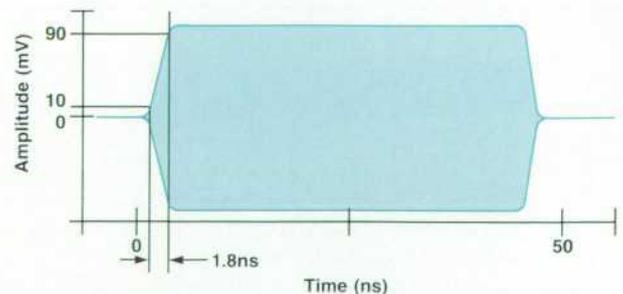
Fig. 6. *Typical pulse envelope.*

and reverse isolation are achieved using HP MMIC traveling wave amplifiers. The fast pulse modulator, at the output of the RF chain of microcircuits, improves the maximum power and pulse performance of the standard HP 8360 instrument.

Reference

1. M. K. Koenig, "A Non-Commensurate Line Length Modulator," *Microwave Journal*, March 1991.

New Technology in Synthesized Sweeper Microcircuits

A new packaging technology using thick-film hybrids and contacts integral to the package simplifies testing and rework and reduces RFI. New circuit designs include a triple balanced mixer and quasi-elliptic low-pass filters. New approaches reduce video feedthrough and harmonic generation.

by Richard S. Bischof, Ronald C. Blanc, and Patrick B. Harper

TWO OF THE NEW MICROCIRCUITS that were developed for the HP 8360 synthesized sweepers are called the low-band microcircuit and the modsplitter (modulator/splitter). These microcircuits were developed with the goals of achieving a high level of integration, improving sweeper performance, and eliminating known system problems. This article will first give a brief overview of the functions of the low-band microcircuit and modsplitter and then discuss how the goals were achieved.

The low-band microcircuit is used to generate the 10-MHz-to-2.3-GHz band in the HP 8360. This band is achieved by mixing the outputs of a YIG-tuned oscillator (YTO) and a fixed-frequency oscillator. The low-band microcircuit receives the YTO input from the modsplitter microcircuit and contains all of the other circuits necessary to provide the low-band signal. These circuits are a 5.4-GHz phase-locked oscillator, an ALC modulator, a pulse modulator, an isolation amplifier, a mixer, an LO amplifier, an IF amplifier, and a bridge detector. The operation of some of these circuits will be discussed below. Fig. 1a shows the low-band microcircuit.

The modsplitter microcircuit routes the YTO signal to several places in the HP 8360 and provides pulse modulation and ALC capabilities. The input to the modsplitter is the YTO output. The modsplitter outputs are:

- Sampler output. This output provides the sampler with a leveled signal by the use of a preleveler, which will

be discussed later.
- Low-band output. This output provides the low-band microcircuit with the required LO frequencies as discussed above.
- High-band output. This output provides the high-band signal, which is amplified and multiplied to provide the 2.3-GHz-to-40-GHz signals.

The modsplitter includes ALC and pulse modulators that have >80 dB of range. Fig. 1b shows a picture of the modspitter.

High Level of Integration

To eliminate as many connections as possible and to address the problems of assembly time and loss of reliability, a high level of integration for both the final package and the circuits was desired. A new packaging technology was developed, making it possible to replace a total of nine separate microcircuits from previous sweeper designs with two microcircuits in the present design. Increases in both functionality and performance were realized. There are two microcircuits and two packages instead of one because the two microcircuits are not both needed in all HP 8360 models. The packaging technology will be discussed in detail later.

At the circuit level, the higher level of integration was achieved largely by a switch from the existing thin-film designs, which offered excellent performance at a high cost,
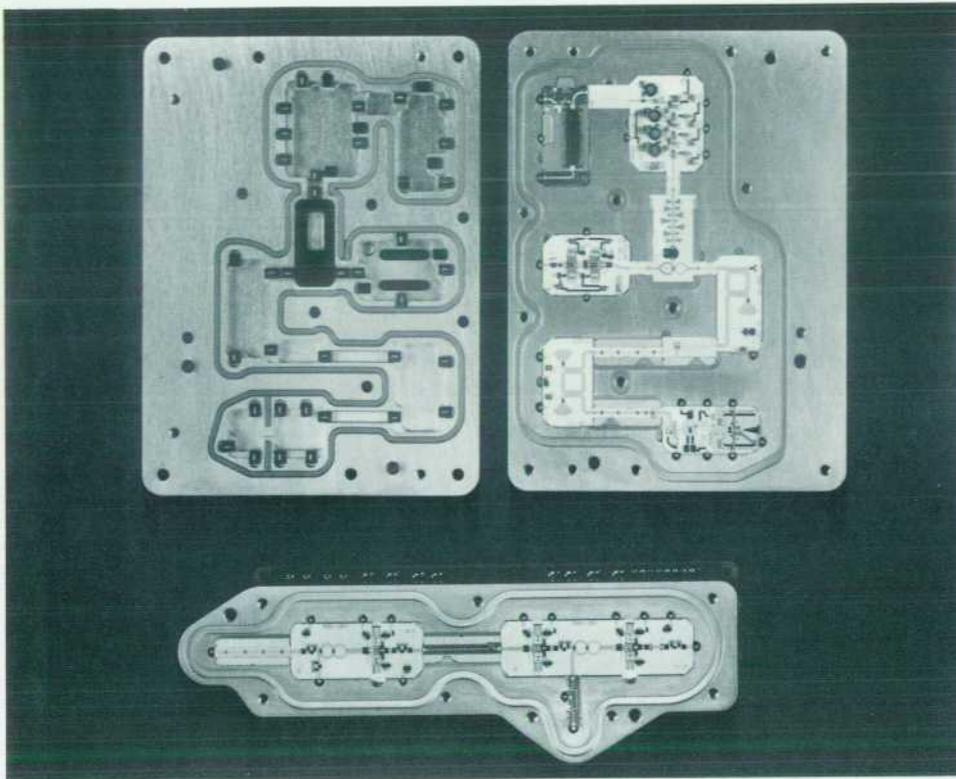
Fig. 1. *(top) Low-band microcircuit with cover. (bottom) Modsplitter microcircuit.*

to thick-film designs, which are able to deliver similar performance at a greatly reduced production cost and at a higher level of integration. For instance, using thick film, it is possible to screen-print the necessary bypass and coupling capacitors, which would otherwise have been added as chip capacitors. Also, where the old design required wire bonds for connections, thick film offers printed dielectric crossovers. Between this and the change to packaged devices, it was possible to eliminate bonding operations totally from some of the circuits.

Using a high-κ dielectric that provides a capacitance-to-area ratio of approximately 10 nF/cm^2, it was possible to design thick-film capacitors with small enough geometries to work well over the 2-GHz-to-8-GHz frequency range wherever tight tolerances were not required. A 10-pF-to-60-pF decoupling capacitor, 0.025 inch wide by 0.047 inch long and built on a 0.025-inch substrate, has excellent VSWR and is used extensively. A bypass capacitor specified at 10 pF to 60 pF and measuring 0.070 inch square including conductive vias and encapsulant works well, but the low-impedance path length to ground must be accounted for in the matching structures. Since these thick-film capacitors are essentially free and do not require epoxying or wire bonding, the cost savings can become significant on circuits where many decoupling and bypass capacitors are used. The LO amplifier of the low-band microcircuit (Fig. 2) is a good example of a circuit with no wire bonds.

A further benefit gained from the switch to thick film is an economy of scale: the thin-film circuits were generally fabricated on substrates with around 2000 mm^2 of usable area, whereas the thick-film circuits are fabricated on larger substrates with about four times as much area. The screen-
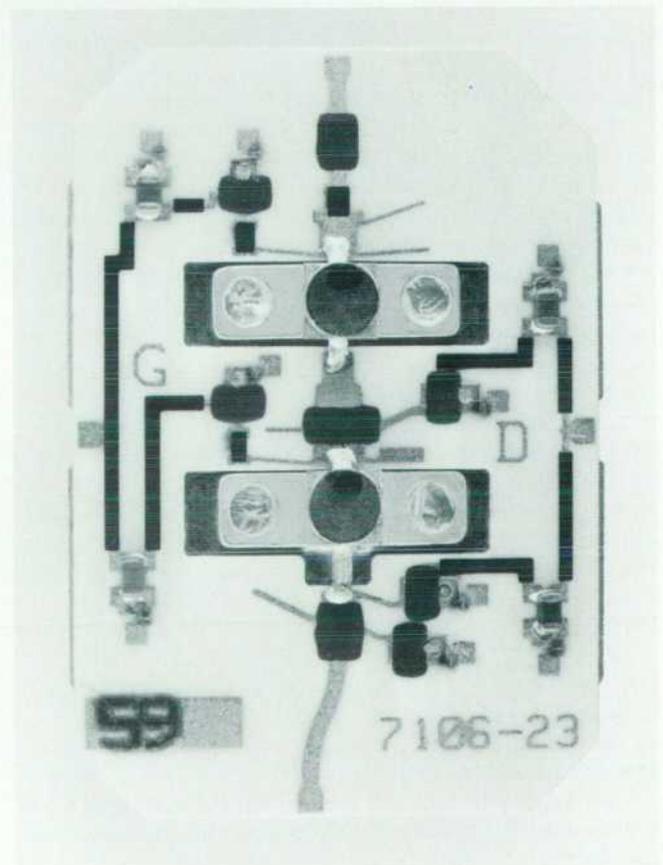
Fig. 2. *The LO amplifier of the low-band microcircuit is an example of a circuit with no bonds. Connections are made by integral contacts on the package lid.*

printing cost is about the same regardless of substrate size, and it is inherently much cheaper than the photolithography required for thin-film fabrication, so it is possible to produce much cheaper circuits with the same functionality and equally good performance.

Still another benefit is that, in general, larger circuits can be produced with thick film, further eliminating connections by placing more circuitry on a single substrate or assembly. In the HP 8360, larger circuits are realized both as single large substrates and, where needed, as collections of smaller substrates on a common carrier. The latter were found to be needed primarily when problems of cracking of the large substrates were encountered and when mixed technology—thin and thick film—was required. In either case, significant reductions in circuit count were achieved. For example, the first stage of the modsplitter consists of several functions (closed-loop preleveler, splitter, amplifier) all on a single thick-film substrate, where previous designs placed each function on its own substrate. Similarly, the modulators used are conventional shunt reflective pin diode modulators with quarter-wavelength microstrip transmission line segments between them, but are realized on a single substrate with laser-drilled holes rather than a series of substrates. Hence they are much easier to assemble, offer more consistent performance, require decreased handling, and are more reliable.

Despite the general switch to thick film, some circuits were left in thin film because of performance requirements. Thin film is capable of generating the fine geometries needed at high frequencies, which thick film cannot match.

New Connection and Packaging Technology

The conventional method of hybrid assembly, which clamps or glues all the circuits into a housing and connects them with wire bonds, was deemed undesirable in this case for two main reasons. First, the process of bonding requires that the entire assembly be heated to 150°C or so, which takes a long time for a packages as large as these and also stresses the circuits. Second, with the goal of a high yield at top-level assembly, it was desirable to be able to test the individual circuits before placing them in the final package. To test conventional circuits under realistic conditions, it would have been necessary to bond them into test packages, which was impractical. To meet the requirements of large-scale integration and economy, and to overcome the limitations of conventional hybrid packaging, a new packaging and connection technology was developed. The design of the new package is more complicated than the traditional package. Grooves are needed for O-ring and elastomer strips, the circuits are aligned by the package itself, and a new type of interconnection called integral contacts is used. The integral contacts are both aligned and located by the package.

Integral Contacts

To connect the circuits, a plated elastomer contact was developed. What was needed was a connector that could be pressed into place and would provide good-quality connections at microwave frequencies. What was developed is effectively a block of silicone rubber with a small (600 micrometer) strip of gold on it. Held in the correct orientation and location by the microcircuit lid, it forms a bridge between two circuits with a gold beam. This connection scheme, referred to as integral contacts, can be seen in Fig. 3.

The integral contacts make it possible to test the circuits before they are enclosed in their final packages. With complex hybrids, such as the low-band microcircuit and modsplitter, and a just-in-time manufacturing process, a high yield of good circuits at the final test system is desired. To this end, pretests are performed on each of the lower-level microcircuits. It is desirable to have the pretest performed by the assemblers who have built the circuits, as soon as possible after the circuit has been completed. In this way the assemblers get instant feedback and the pretest station can be used as a statistical quality control tool. For the line personnel to perform the tests, the systems must be convenient to use and very repeatable. The pretest stations are set up in line with the assembly flow and give a pass/fail indication, so no interpretation is necessary. Because of their ruggedness and repeatability, integral contacts are used in these test fixtures. This also ensures that the test environment for the lower-level circuit is the same as its final operating environment.

The integral contacts also proved to be helpful in the development of the microcircuits. Traditionally, hybrid circuit design consists of several design/build/test cycles. Data from each cycle is used to adjust the subsequent design. The need for large sample sizes must be weighed against the cost of each turnaround. Large sample sizes are needed to determine the range of performance, but this sometimes requires several slightly different versions of a circuit and as many different wafers of active devices as possible. Thus the cost of large sample sizes can be prohibitive, with test package costs being a large part. One method that is used is to rework each package with new circuits and/or devices. The variability that this rework adds to the performance of the circuit can cause extra design cycles. In the HP 8360 low-band and modsplitter microcircuits, integral contacts were used successfully to minimize this problem. The contacts gave repeatable microwave connections and made it
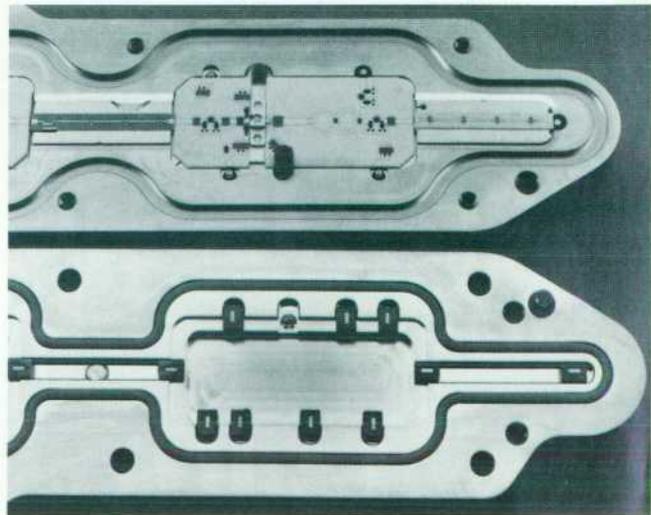

Fig. 3. *The modsplitter microcircuit in an open package, showing the integral contacts.*

possible to test many versions of circuits in the same package and under the same conditions. It should be noted that the mechanical outline of the circuits must be held constant for this approach to work. Significant assembly cost savings are realized because no circuit-to-package assembly is required.

As new versions of the microcircuits are needed or as better devices are available, the designer is faced with the decision of whether to redesign or add to the existing microcircuit. Integral contacts make it relatively easy to configure new versions of the microcircuit as they are needed. The designer only needs to keep the same physical outline as the existing circuits, but can change any or all of the functions. This scheme also allows new devices to be retrofit into the microcircuit as they become available. One assembly line can build any number of versions of the circuits and they can be assembled at the final assembly area into the versions needed.

A complex hybrid such as the low-band microcircuit is inevitably going to require rework. With traditional microcircuit design this rework can be very costly because of the nature of the construction (epoxy lid seal, epoxy circuit attachment, wire bond interconnects). With integral contacts the rework can be done quickly and easily, thereby reducing scrap costs.

Isolation and RFI Concerns

With separately packaged microcircuits, the potential for interference between the various circuits is small, since each package can be well-sealed and the packages can be interconnected with semirigid coax. Once many circuits are placed in a common housing, however, it is much more likely that they will interfere with each other, and careful design is needed to prevent problems. For instance, in the low-band microcircuit, the mixer and the fixed-frequency oscillator are separated by two series modulators with a dynamic range of >150 dB. For the modulators to use their full range, the mixer and the oscillator must be electrically isolated by at least that amount, yet in this case are mounted in a common package less than two centimeters apart. Fortunately, it is possible, by proper use of conductive elastomers, to achieve this degree of isolation.

An O-ring is used for the environmental seal. A strip of conductive elastomer (basically silicone rubber with small pieces of metal embedded in it) is used for all electrical sealing and to clamp some of the circuits in place. In some areas, this seal is accomplished with two strips spaced some distance apart. This provides >150 dB of isolation at frequencies around 5.4 GHz.

A new screw-body dc feed was designed for the low-band and modsplitter microcircuits. The need to do this was driven by the need for a 0.050-inch-diameter pin to accommodate the integral contact, the need for a superior grounding method for RFI, and the need for better microwave rejection. The available knurl-body feeds had two known problems. One was knurl grounding. The conflicting requirements of mechanical strength and good RF grounding often lead to RFI problems. These feeds also showed a large variation in their ability to reject microwave energy. One of the early goals for the low-band microcircuit and modsplitter was superior RFI performance. The screw-body dc

feed contributes to this by providing good mechanical strength and good grounding. Grounding is accomplished by a mechanical mating of the dc feed to the package much the same as a connector is mated to a package. The superior microwave rejection was achieved by building a test fixture to measure the feeds and working with the vendor on process improvements. The 1500-pF dc feed is specified for 60 dB of rejection from 3 GHz to 8 GHz.

Low-Band Microcircuit

The low-band microcircuit is designed to address system-level problems and provide performance improvements. The system issues are integration, which has already been discussed, and video feedthrough. The two key performance improvements obtained are in maximum available power and broadband noise.

Fig. 4 is a block diagram of the low-band microcircuit. One major new design is the RF amplifier, which is a gallium arsenide FET monolithic IC amplifier with a thin-film filter. The need for this amplifier and its function can be explained by examining the phenomenon known as video feedthrough.

In a system with pulse modulation capability, such as the HP 8360, video feedthrough is basically a reproduction of the original pulse at the instrument output, not as a pulse envelope, but as a pulse at the original frequency. For example, if a 1-GHz signal were modulated with a 1-MHz square wave and the result passed through a 100-MHz low-pass filter, ideally nothing should come out. In reality, because of many different mechanisms, there will typically be some component of the original 1-MHz modulation signal left. This is termed video feedthrough.

In many systems, simple filtering will remove video feedthrough. For instance, if pulses occur at a 10-MHz rate and the system only provides signals down to 1 GHz, then a simple high-pass filter can quite effectively remove the 10-MHz signal and leave a clean modulated RF waveform. In the HP 8360, however, outputs are down to 10 MHz (or lower), so simply filtering the output signal to remove the pulses is not practical.

The dominant video feedthrough mechanism in a down-conversion system such as the HP 8360 is interaction between the mixer and pulse modulator. The LO signal (at roughly 20 dBm in this case) leaks across the mixer to some extent because of imperfect mixer balance. This leaves a reasonably strong signal coming out of the RF mixer port and going toward the pulse modulator. This signal is reflected from the pulse modulator (it has a shunt reflective design), travels back to the mixer, and mixes with the original LO signal, generating dc. If the pulse modulator reflection stayed the same, this would not be a problem, since there are several dc blocks in the IF path. However, the reflection changes drastically as the pulse modulator is turned on and off. This leads to two discrete values of dc at the IF output, which form a reconstruction of the original pulse—that is, video feedthrough.

Since the IF must be leveled down to −20 dBm, the video feedthrough could be as bad as 10 dB higher than the desired signal. In this case one would see at the output a large square wave with some RF superimposed—hardly the pulsed RF signal that is desired. A typical good specifi-
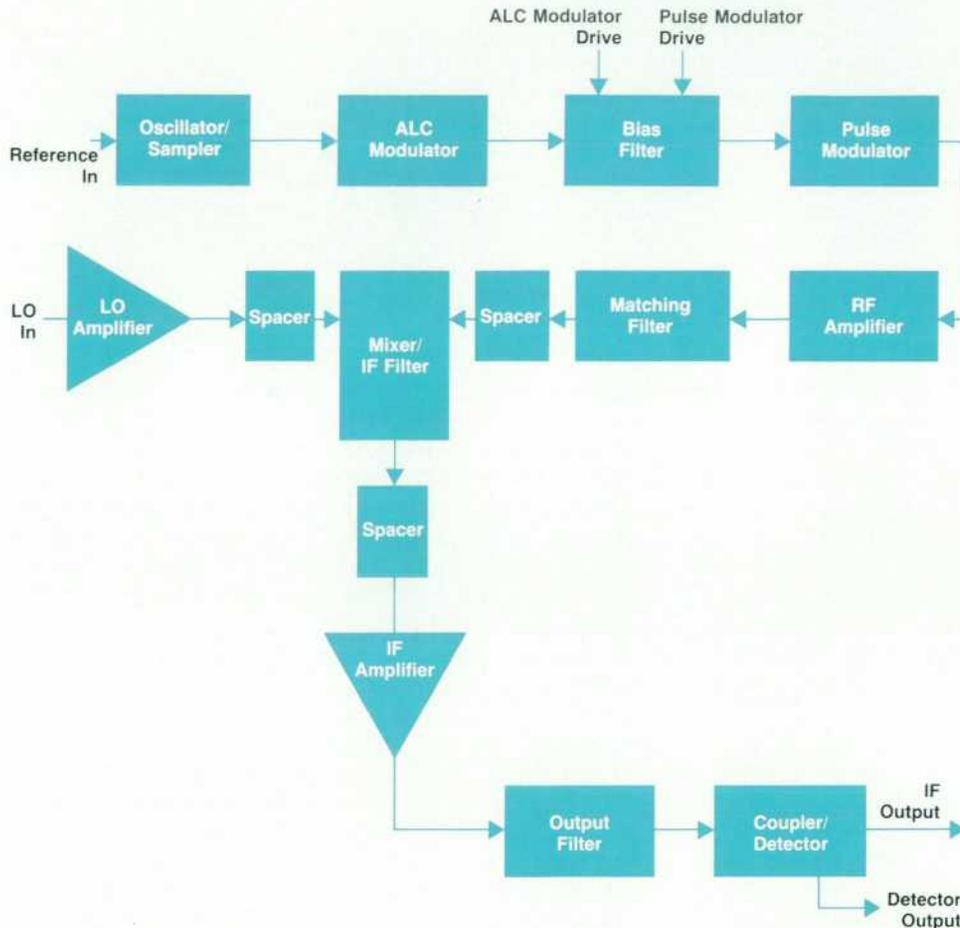
Fig. 4. Functional block diagram of the low-band microcircuit.

cation for video feedthrough is around 1% of the RF signal (−40 dBc). Hence, at least a 50-dB improvement is needed.

To achieve this, the reverse isolation of an amplifier is used. In this case, a two-stage integrated amplifier with an s_{12} of about 60 dB. A magnetic isolator could have been used, as in the HP 8340, but it would have been much too large and too expensive for this application.

While the RF amplifier chip is used primarily for isolation, it is also used to amplify the RF signal slightly. A side benefit is that it will generally sharpen any RF pulse envelopes sent through it.

LO Amplifier

To maintain optimum mixer performance, it is necessary to provide a very flat (±1 dB) LO drive signal, controllable over about a 10-dB range. The source of this signal is the modsplitter, which is not able to meet this flatness requirement. Since it was planned to use the low-band microcircuit in other designs, the flatness control is built in. This is the function of the LO amplifier circuit.

The LO amplifier is a two-stage packaged FET circuit with lossy matching. It is essentially a limiting amplifier, so that the output power is fairly constant and is insensitive to the input power. This circuit has about 15 dB of small-signal gain, operates with a nominal input of around 10 to 15 dBm coming from the modsplitter, and provides a nominal output power of around 20 dBm. Thus it is nominally 5 to 10 dB into saturation and in the worst case can be

fully saturated (no gain at all). Even fully saturated, it puts out a clean signal (−30 dB harmonics) because of the very effective low-pass matching at the output. Output power control is achieved simply by varying the drain bias voltage. This circuit uses interstage matching optimized for minimum sensitivity. A final benefit from this limiting design is that any amplitude noise on the signal feeding the LO amplifier is reduced quite effectively before it hits the mixer—just as it is in limiters used in typical FM systems to eliminate AM.

The LO amplifier is easy to build as designed, since it is insensitive to the devices that make it up.

Triple Balanced Mixer

The triple balanced mixer is the heart of the low-band microcircuit. This design was chosen because it allows a very high-level drive (20 dBm nominal) and provides good spurious performance. The mixer uses a total of eight diodes. The power input is distributed among all eight diodes evenly (four are turned on at a time with 50% duty cycle). With a 20-dBm drive level, only 11 dBm heats any given diode, which is acceptable from a reliability standpoint. The peak power for each diode is twice this level (14 dBm) because of the 50% duty cycle.

The mixer operation is quite simple in principle. The LO signal is used to operate the diodes as switches. The switches connect the RF input to the IF output either directly or 180 degrees out of phase. The RF signal is used

Modular Microwave Breadboard System

One potential disadvantage of the drop-in assembly methods used for the modsplitter and low-band microcircuits is that there is no easy way to test the individual circuits within these hybrids. With the former design method of placing each circuit in its own package, this was not a problem. In the new design, however, the circuits are hidden away inside a final large package and it may not even be possible to gain access to both ports of all of the circuits. A new technique was needed to test the circuits during development and in production.

Fortunately, during development of the microcircuits, the HP 83040 modular microcircuit package was also being developed. This microwave modular breadboard package consists of a series of end blocks with built-in coaxial-to-microstrip launches and center bodies of various lengths for circuits to be tested. Included is all of the hardware needed to connect the various blocks. The end blocks and center bodies in this scheme all share a common mechanical interface, which is genderless so that they can be connected in any order and in any quantity. Fig. 1 shows examples of modular breadboard parts.

Using this system in designing the multistage low-band microcircuit, each separate circuit (or circuit element) was developed and tested in the final microstrip environment. Then the individual blocks were assembled into a full circuit. The entire low-band microcircuit was developed using such a scheme, and the initial fully working breadboard was built using it.

Advantages of the HP 83040 modular microcircuit package include:

- It offers microcircuit breadboard packaging that is very cost-effective compared to custom packaging and saves time because it is available off the shelf.
- All microstrip circuitry is measured directly in microstrip envi-

Fig. 1. *Examples of microwave breadboards using the HP 83040 modular microcircuit package.*

ronment.
- It provides true zero-length through calibration since the end blocks can be directly connected together. Deembedding can be used to make precise microwave measurements.
- Entire microcircuits can be measured and assembled in any combination of subcircuits.

Stan Bischof
Development Engineer
Network Measurements Division

100% of the time, so no power is lost, as it is in a single balanced mixer, which throws away power during half of its cycle.

This type of mixer can be described as a biphase modulator. Ideally, it performs a time-domain multiplication of the RF signal with a square wave at the frequency of the LO drive. In the frequency domain, this yields a convolution of the RF signal spectrum with the spectrum of the square wave drive, which contains the fundamental and odd harmonics of the drive frequency. Of the many terms in the convolution, the principal ones are the $f_{RF} - f_{LO}$ and $f_{RF} + f_{LO}$ terms. The $f_{RF} + f_{LO}$ term is filtered out so that the primary term left is the $f_{RF} - f_{LO}$ term, which is exactly what is wanted for the IF.

The eight diodes are grouped into two quads, which are mounted on the top and bottom of the mixer substrate in a symmetric pattern (Fig. 5).

Fig. 6 shows some of the details of the RF balun. Ideally, as described above, the RF path is switched by the LO drive such that the signal is directed to the IF output in either 0° or 180° phase orientation. Before this happens, the signal is split into two halves, relying on symmetry to ensure that they are equal in magnitude and phase. This is accomplished by two transitions. First, the signal is transformed by a step transition from microstrip line to a coupled suspended microstrip format consisting of top-side and bottom-side electrodes in a symmetric arrange-

ment on the top and bottom of the substrate. Then, on each side, two 100-ohm lines are split off in another step transition. This leaves a total of four leads, two on each side of the substrate. Each lead is connected to a diode switch. The four diodes in each quad are always connected to the IF output in pairs and the IF impedance is 50 ohms, so the RF input port, in principle, is a good 50-ohm match. In practice, of course, none of the transitions is perfect and

Fig. 5. *Diode quads in the triple-balanced mixer.*

the diodes are not ideal switches.

Each balun line is a quarter wavelength long at the lower frequency of 5400 MHz. It is important that all the lines be as close to the same length as possible so that the phases match. At 5400 MHz, 90° of phase shift corresponds to roughly 4400 micrometers in this medium. Thus, to get phase match of, say, ±5°, the paths must match within ±250 μm, or about one circuit thickness.

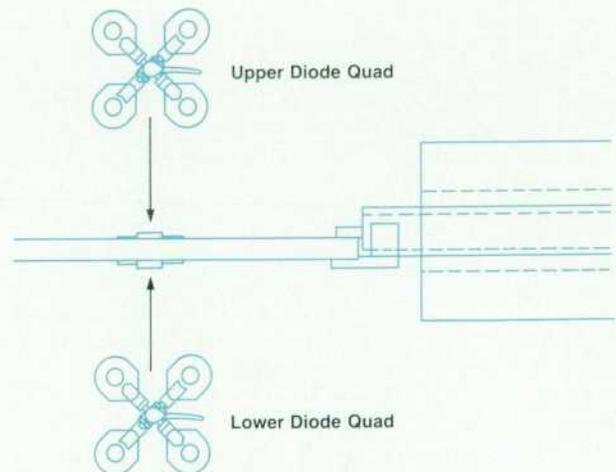Fig. 7 shows the LO balun, which is similar in concept to the RF balun. It, too, is formed of a microstrip-to-suspended transition followed by a split into two paths of suspended balanced microstrip line. In this case, however, the impedances are different. Since the LO provides a large enough signal to turn the diodes on fully, each of the two arms of the balun is effectively terminated in two series diodes. Since each diode has a large-signal impedance of about seven ohms, two diodes in series represent around 15 ohms. For an adequate match, the balun arms are used as quarter-wave transformers to make each arm look like about 100 ohms at its input. To transform 15 ohms to 100 ohms takes $\sqrt{15 \times 100}$ or approximately 40 ohms. Hence each arm is designed to have an impedance of 40 ohms.

Because of symmetry, the RF and IF paths in the LO balun do not, in principle, load down the ends of the balun arms, so they need not be considered in the above analysis. As in the RF path, there are four leads, which are properly terminated. To maintain symmetry, the arms are carefully constructed to have the same length. A balun is used at the IF port to provide a single-sided output to drive a single-sided filter. The high output level of the mixer allows the gain of the IF amplifier to be 10 dB less than in previous designs. This lower gain gives a corresponding improvement in the broadband noise of the low-band microcircuit.

IF Amplifier

The IF amplifier of the low-band microcircuit is a three-stage FET amplifier with 30 dB of gain and 24 dBm of output power. The design is a lower-gain version of the IF amplifier used in the HP 8753A/B/C spectrum analyzers.

The IF amplifier makes possible a +17-dBm instrument power specification, which is needed for the high-power version of the HP 8360, while maintaining the harmonic performance.

Quasi-Elliptic Low-Pass Filters

Throughout the low-band microcircuit, a number of low-pass filters with very high performance are needed. They need sharp cutoffs and very broad stop bands. At the same time they must be low in cost, repeatable, and as small as possible. To do this on microstrip, a new filter was developed. The design uses spiral inductors and radial transmission lines, and has stop-band zeros to extend the stop band. In most cases, the filters are less than a quarter wavelength long at the cutoff frequency, so they are quite small both physically and electrically.

One filter of this type, used in the RF path, is designed to pass 5.4 GHz. Requirements for this filter were that it attenuate the second, third, and fourth harmonics by 40 dB or more, and that it cut off as rapidly as possible above 5.4 GHz. This low-pass filter is a microwave implementation of a simple Chebychev eleventh-order, 0.1-dB ripple, L-input ladder filter, with one major difference, as described below. As such, it consists of the microwave equivalents of six series inductors and five shunt capacitors. The inductors are all realized using spiral coils of high-impedance line. The capacitors are realized using radial transmission lines.

This circuit differs from the conventional Chebychev design in its behavior outside of the passband. While a Chebychev design has all its transmission zeros at infinity, this design places transmission zeros at finite frequencies within the desired stop band. This is accomplished by using the radial transmission lines as shunt resonators. Because these transmission zeros are at finite frequencies, equal ripple is not achieved in the stop band, as it is with a true elliptic filter. Therefore, this filter can be described as a quasi-elliptic low-pass filter.

The reason for using a quasi-elliptic low-pass filter can be seen by looking at the response of a typical distributed
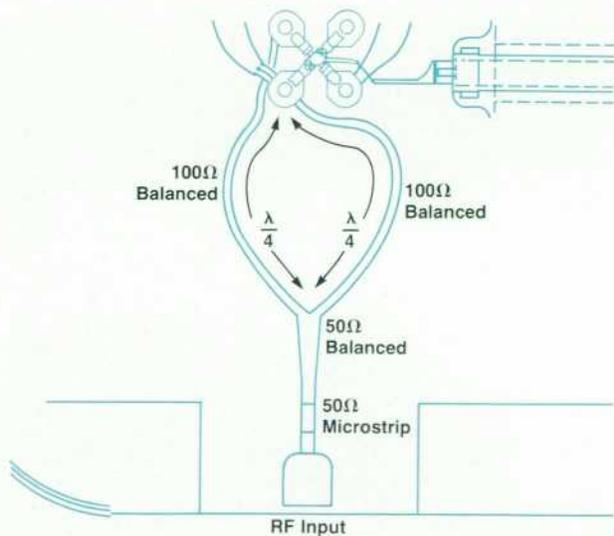

Fig. 6. *Details of the RF balun in the triple-balanced mixer.*
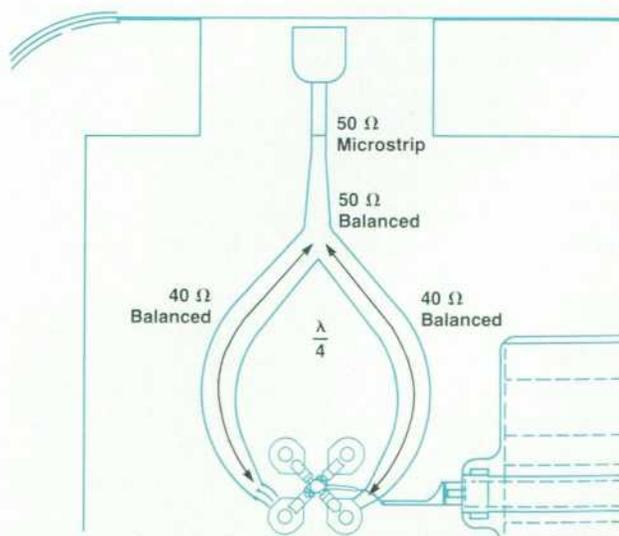

Fig. 7. *Details of the LO balun in the triple-balanced mixer.*

filter. While the passband can be made to look nearly ideal, the distributed elements lead to secondary responses, which create secondary passbands. This effect differs from filter to filter, but typically at around three times the cutoff frequency there is another passband where the response is not much below the response in the main passband. The quasi-elliptic low-pass filter places zeros within and near this secondary passband and is able to keep the response low at much higher frequencies than would otherwise be possible.

The inductor design is shown in Fig. 8. It consists of an inner bonding pad, a length of high-impedance line arranged in a spiral, and a bond wire connecting the inner end of the spiral to another bonding pad.

The ground plane is cut away from beneath the inductor out to a distance of one substrate thickness. This raises the line impedance, or equivalently, removes some of the shunt capacitance of the spiral. This accomplishes two things. First, since the line impedance is higher, the line need not be as long to realize a given inductance. Second, since the shunt capacitance is lower, the self-resonant frequency of the inductor is higher, rendering the filter response useful to higher frequencies.

These spiral inductors are designed using an empirical formula based on the particular geometry and material used. The quasi-elliptic low-pass nature of the filter is not influenced by the inductors, but is rather set by the radial transmission line used as capacitors. Each of the five shunt capacitors is realized by a set of two radial transmission lines of the type shown in Fig. 9.

Like a shunt stub, a radial transmission line resonates at some frequency to transform its open end into a short circuit, thereby forming a transmission zero for the line it is connected to. While the shunt stub transmission line resonates at the frequency where its length is a quarter wavelength, the radial transmission line resonates at a frequency where its length is roughly one sixth of a wavelength. Thus the radial transmission line is more compact. A key advantage of a radial transmission line for this application is that it is pointed at one end, so it forms a simple, well-defined connection, unlike a typical shunt transmission line.

The length of each radial transmission line is chosen for

Fig. 9. *Radial transmission line resonator used in the quasi-elliptic low-pass filters.*

its resonant frequency, and the angle is chosen for the parallel-plate capacitance needed to realize the required shunt capacitance for the Chebychev filter. As Fig. 9 shows, all ten radial transmission lines are different. This provides twice as many zeroes to work with and avoids the situation of two nearly identical resonators being connected together. Using one resonator at each node also works, but this halves the number of usable zeroes.

Since this filter is filtering a narrow-band signal, the transmission zeros are clustered in narrow bands around the harmonics of the RF (i.e., at 10.8, 16.2, and 21.6 GHz). Fig. 10 shows a full quasi-elliptic low-pass filter.

Modsplitter

The modsplitter makes a number of contributions to overall system performance. Fig. 11 shows its block diagram. The modsplitter provides the LO input to the low-band microcircuit. Extending this frequency to 7.8 GHz, higher than in previous designs, allows the low-band microcircuit to improve spurious performance. The higher drive frequency moves the 2/1 spurious response (the $2f_{RF} - f_{LO}$ mixer spur) out of band, where it can be filtered. One of the contributors to low-band video feedthrough is eliminated by locating the modsplitter's pulse modulator after the low-band output. This eliminates the possibility that

Fig. 8. *Details of the spiral inductor used in the quasi-elliptic low-pass filters.*

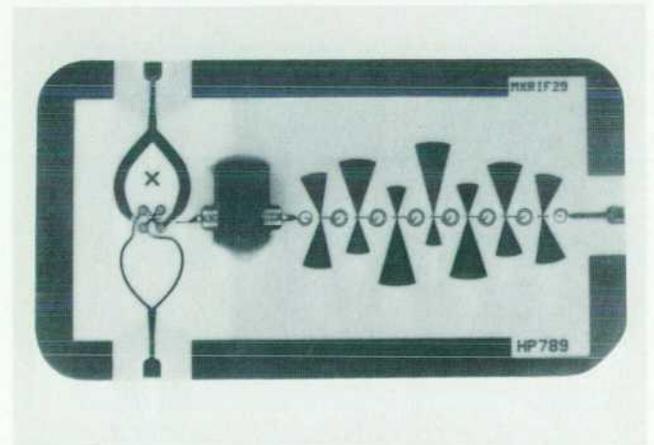**Fig. 10.** *Fifteen-pole quasi-elliptic low-pass filter.*

Quasi-Elliptic Low-Pass Filters

The term quasi-elliptic low-pass filter is used here to describe a particular type of filter designed for use in the low-band microcircuit of the HP 8360 family of synthesized sweepers. It was developed because it offers high performance in a very small area and is realized in microstrip and is therefore compatible with the rest of the circuitry.

This type of filter is distributed, yet achieves both the sharp low-pass cutoff and the broad stop band free of spurious responses that are characteristic of elliptic filters. While it has finite-frequency zeros to enhance the stop band, it is not a true elliptic filter since the zeros are not arranged to provide an equal-ripple stop band.

The elements of this type of filter are round spiral inductors and radial transmission lines. The radial lines function as parallel plate capacitors at low frequencies (within the passband) and as shunt resonators at high frequencies. The length of the resonator is chosen to provide a zero at the desired frequency, and the area (or angular width) is chosen to provide the correct low-frequency capacitance. Fig. 1 shows a typical filter.

This new filter combines the convenience of fabrication of microstrip distributed filters with the spurious-free stop band previously achievable only in lumped filter designs.

Design Procedure

The design procedure is as follows:
1. Determine the needed capacitance and inductance values from basic low-frequency filter tables. For example, a 0.01-dB-ripple Chebychev low-pass prototype (1 ohm, 1 rad/s) requires the following element values (a seven-pole design was chosen based on the required cutoff frequency):

0.796944	Series Inductor
1.39242	Shunt Capacitor
1.74813	Series Inductor
1.63313	Shunt Capacitor
1.74813	Series Inductor
1.39242	Shunt Capacitor
0.796944	Series Inductor.

To get the actual C and L values, denormalize by:
a) Scaling the frequency to 3 GHz (desired cutoff)
 - Multiply L by $1/(2\pi f)$
 - Multiply C by $1/(2\pi f)$
b) Scaling to 50-ohm impedance
 - Multiply L by 50
 - Multiply C by 1/50

The denormalized values are:

L_1	2.114 nH
C_2	1477 fF
L_3	4.637 nH
C_4	1733 fF
L_5	4.637 nH
C_6	1477 fF
L_7	2.114 nH.

2. Compute the needed inductor values from empirical design formulas for the inductor model shown in Fig. 2. If T is the number of turns, L is the inductance in nH, and C is the capacitance in fF,

$$T = (L/1.9)^{0.62}$$
$$C = (90 + 30T).$$
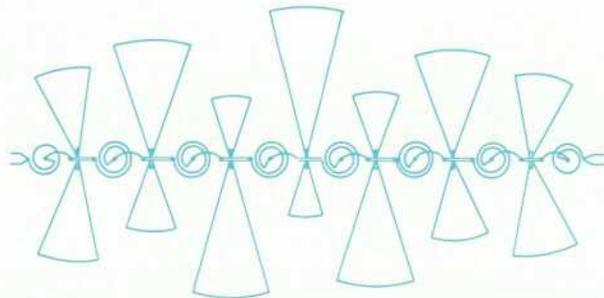
Fig. 1. *Typical fifteen-pole quasi-elliptic low-pass filter.*
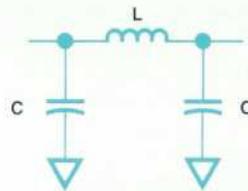
Fig. 2. *Inductor model used in filter design.*

residual signals on the modsplitter pulse modulator bias line will cause residual pulse modulation on the LO signal to the low-band microcircuit.

Excessive harmonics generated in the modsplitter can cause power holes by canceling with the desired harmonics generated by the YIG-tuned step recovery diode multiplier. One of the system contributions of the modsplitter is the addition of a preleveler modulator, which levels the 8 dB of power variation from the YIG-tuned oscillator and keeps the input stage of the modsplitter from saturating and generating high harmonics over much of its operating frequency range. As a result of this preleveler, along with high-power FETs and an overall flat gain response, the modsplitter operates with better than −20 dBc harmonics regardless of the YTO output power or the ALC modulator state.

The high-power and correspondingly high-Q 750-μm FETs made it difficult to achieve a good input match over the entire two-octave bandwidth. The traditional input matching of a shunt lossy gate bias line and a lossless series low-impedance transformation transmission line results in a very low impedance at midband and an impedance slightly higher than 50 ohms at the high end of the band. Adding a low-value series resistor at this point greatly improves the midband match with minimal degradation of the gain or match at the high end of the band.

Because of the gain unflatness contribution of the high-power microcircuits, the HP 8360 requires a wide ALC modulator range. In previous designs, the quarter-wavelength spacing between the four diodes of the ALC modulator was at 8 GHz. This gave maximum attenuation at the high end of the band and effectively added a gain

Thus L1 has 1.05 turns and shunt C of 122 fF

L3	1.75	143
L5	1.75	143
L7	1.05	122.

The inductor design is now finished.

3. Modify the needed capacitors by subtracting the equivalent capacitors from the inductor model values. The shunt Cs in the final filter can be constructed by using the equivalent shunt Cs from the above inductor model. For instance, C_2 is between L_1 and L_3, so part of C_2 can be supplied by 122 fF from L_1 and 143 fF from L_3. Thus the additional capacitance needed to form C_2 is $1477 - 122 - 143 = 1212$ fF.

Capacitance Needed	Amount from Inductors	Additional
C_2 1477 fF	122 + 143 = 265	1212
C_4 1733	143 + 143 = 286	1447
C_6 1477	122 + 143 = 265	1212

4. Determine the lengths of the radial transmission lines by selecting the desired resonant frequencies (zero locations) and getting the lengths from empirical design tables.

Frequency (GHz)	Outer Radius (μm)	Capacitance (fF/degree)
4.5	3750	51.0
5.5	3100	34.9
6.5	2700	26.6
7.5	2300	19.3
9.0	2000	14.7
11.0	1700	10.6

The capacitance shown is pure parallel plate capacitance, which works here because only low-frequency effects need to be considered. The distributed and fringing effects are not seen until well above the passband.

5. Assign radial transmission lines to the various capacitors. Determine the angular widths of the capacitor plates based upon the parallel plate capacitance and the needed total capacitance. C_2 will use two radial transmission lines with 5.5-GHz and 9-GHz resonators and angular sizes of 25 and 23 degrees:

$$C_2 = 1212\,\text{fF} = 3100(34.9\,\text{fF/degree})(25\,\text{degrees}) + 2000(14.7\,\text{fF/degree})(23\,\text{degrees}).$$

Similarly,

$$C_4 = 1447\,\text{fF} = 3750(51.0\,\text{fF/degree})(25\,\text{degrees}) + 1700(10.6\,\text{fF/degree})(16\,\text{degrees}).$$

$$C_6 = 1212\,\text{fF} = 2700(26.6\,\text{fF/degree})(25\,\text{degrees}) + 2300(19.3\,\text{fF/degree})(28\,\text{degrees}).$$

Stan Bischof
Development Engineer
Network Measurements Division

slope to the modulator to provide some filtering of the troublesome second harmonic at 8 GHz. Since the modsplitter preleveler eliminates this harmonic problem, the quarter-wavelength spacing between the ALC diodes is readjusted to the midband for improved modulator range. Further improvement in ALC modulator range is achieved by the use of a lower-resistance pin diode. The midband quarter-wavelength spacing also results in a flatter attenuation-versus-frequency characteristic for the modulator. Modulator flatness is improved by adding padding before and after the ALC modulator to minimize the VSWR interaction between the FETs and the ALC modulator. The flatter modulator contributes to better AM performance in the HP 8360.

RFI leakage from the modsplitter microcircuit can couple back into the RF path after the pulse modulator and cause deterioration in pulse performance. Using 50-pF package dc feeds instead of 0-pF SMB connectors for the ALC and pulse modulator bias lines decreases RFI emissions. When the instrument is leveling at low power, the ALC modulator is in a high-attenuation state and there is less RF energy after the ALC modulator to radiate and escape from the package. Since the 50-pF feeds are not good RFI filters and since the HP 8360's pulse on-off ratio is more sensitive to RFI failures at low leveled power, the ALC modulator bias package dc feed is placed in the package cavity following the ALC modulator to minimize radiation during low leveled power states. Similarly, during pulse modulation,
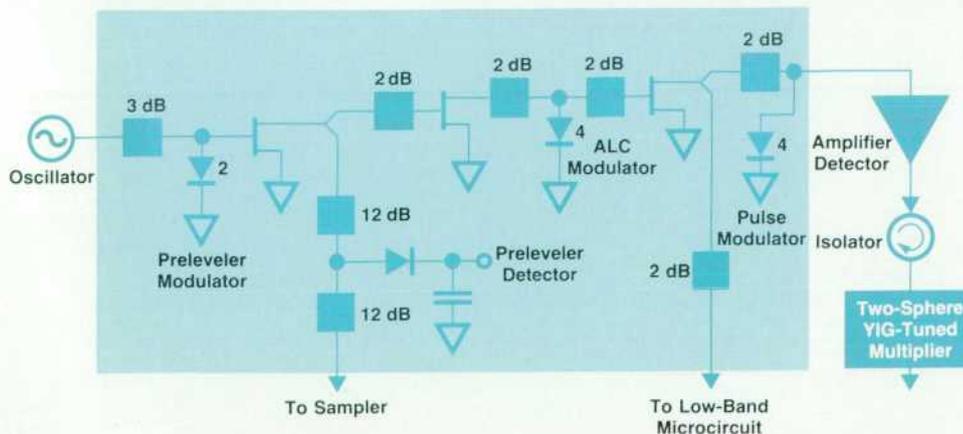
Fig. 11. Functional block diagram of the modsplitter microcircuit in the standard HP 8360 RF section.

the power level following the pulse modulator is extremely low and therefore sensitive to RFI leaking back into the RF path at that point. Therefore, the 50-pF dc feed for the pulse modulator bias is located in the package cavity before the pulse modulator. There are 15,000-pF dc feeds, which would have provided much better RFI filtering and are not location-sensitive, but they cannot be used for fast control signals.

The YIG-tuned oscillator will change its frequency if the load match it is driving changes with time. Therefore, isolation between the YTO and the ALC modulator is important to avoid frequency pulling when the power level is changed and to avoid AM-to-PM conversion. The modsplitter provides two stages of FET amplification before the ALC modulator. This provides high isolation and minimizes frequency pulling problems.

With good harmonic performance under all instrument conditions, extended bandwidth, extended ALC modulator range, controlled gain flatness with ALC modulation, high YTO isolation for minimal AM-to-PM conversion, low RFI radiation and susceptibility, and elimination of residual pulse modulation on the LO drive, the modslpitter reduces or eliminates a number of the system problems inherent in the broadband multiplier RF chain.

Acknowledgments

DC-to-50-GHz Programmable Step Attenuators

Based on HP's proven edgeline technology, these attenuators provide the HP 8360 sweepers with up to 90 dB of attenuation in 10-dB steps.

by David R. Veteran

COAXIAL ATTENUATORS are simple enough in concept but not so easy to design when the required frequency range is dc to 50 GHz. This is especially true when switchability and high repeatability without loss of accuracy and match are required.

A few years ago, type N and SMA were the connectors of choice for coaxial devices. Their maximum upper frequency limit was about 18 Ghz. Then came the 3.5-mm connector and the upper limit moved up to 26.5 GHz. Now we have moved up again with the advent of the 2.4-mm connector and its 50-GHz capability.

The move to ever higher frequencies for broadband measurements has created the need for a whole new generation of coaxial devices. Adapters, pads, terminations, opens and shorts, and air lines came first. They made it possible to make accurate coaxial measurements to 50 GHz, an important first step. Then came broadband sources and test sets which required step attenuators for level control. For the HP 8360 synthesized sweep oscillators, new 50-GHz attenuators were developed. These attenuators are also available separately as the HP 33324/26/27 programmable step attenuators (Fig. 1).

Edgeline Configuration

As a starting point it was decided to base the new attenuators on the edgeline cascaded configuration pioneered by HP 20 years ago. A lot of experience has been gained over the years with these attenuators and they have proven to be very reliable.

In the edgeline system, the center conductor is a flat ribbon suspended between and at right angles to a pair of parallel conducting planes that serve as the outer conductors. The center conductor can be flexed from one contact to another to switch signal paths without changing the transmission line impedance. A big advantage of the edgeline system is that the outer conductor remains fixed, unlike a turret system, in which both inner and outer contacts are broken each time the device is switched.

In the edgeline attenuator the flat flexible center conductor is bent in one direction to make contact to a thin-film attenuator card and in the opposite direction to contact a short length of transmission line that establishes a straight-through path or bypass (Fig. 2). Because electric fields are concentrated along the edges, attachment holes and push-rods can be placed in the center with little negative effect.

To ensure mode-free performance to 50 GHz, the critical internal dimensions of the new attenuators could be no larger than one half that of the HP 26.5-GHz attenuators. Direct scaling of the larger attenuator was not possible because the stiffness of mechanical parts does not scale linearly with size. Also, it was a goal to reduce the residual loss in the new design, and that meant some new thinking would be required on some of the parts. Along with low loss, good environmental performance was a must, including long life. Refinement of these concepts led to the new family of attenuators.

Attenuator Design

The basic structure of the new attenuators is shown in Fig.

Fig. 1. *The programmable coaxial attenuators used in the HP 8360 synthesized sweepers are also available as the HP 33324/26/27 programmable step attenuators. They operate from dc to as high as 50 GHz and provide up to 90 dB of attenuation in steps of 1 dB or 10 dB.*
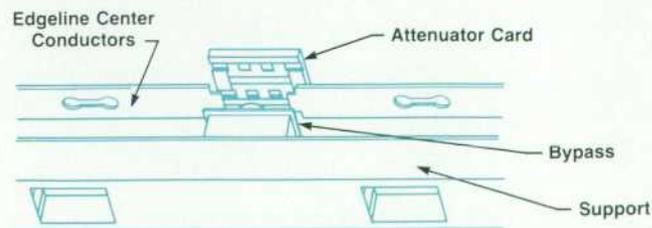

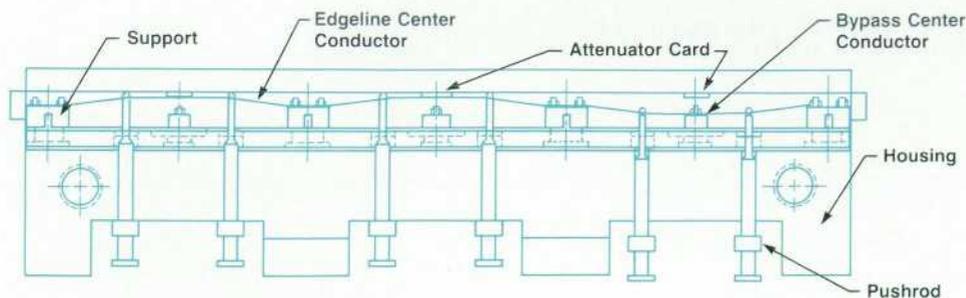**Fig. 2.** *Attenuation section with ground planes removed.*

Fig. 3. Cutaway side view of a three-section attenuator with solenoids removed.

3. Attenuator sections are cascaded and switched in or out in combinations to give the desired attenuation.

The attenuator cards use a single-crystal sapphire substrate with tantalum nitride resistors and plated gold conductors. The resistor pattern is in a π-like configuration. Plated gold center conductor pads on the ends provide contacts for the edgeline switches. The cards are clamped into a common body that forms the outer conductor of the edgeline structure.

The center conductors are mounted on a common support structure, which is molded from two different materials. The towers to which the center conductors are attached are molded from a modified PPO plastic. The base material of the support structure is molded from a low-expansion composite that has the ability to absorb unwanted modes. This center conductor support assembly fits into a T slot in the housing and is spring-loaded for accurate alignment and good shock resistance.

The edgeline conductors are driven by pushrods that attach in the middle where there is little effect on the fields (Fig. 3). The pushrods are made from polyamideimide plastic which is very stable with temperature and very strong. The center conductors are slightly overbent after making contact

to provide a wiping action and ensure a repeatable contact. A finite element analysis program was used to optimize the conductors for low stress with minimum length. HP has recently developed some very useful microwave modeling programs, which were helpful in optimizing the edgeline dimensions for low loss and good match.

The transition from edgeline to coax at the ends is done through an electroformed metal bellows. This provides mechanical isolation between the coax center conductor and the edgeline conductors inside.

The 2.4-mm connectors are mounted on the ends coaxially. They have a flange-type mount with an O-ring RF gasket under the flange. This design makes it easy to mount other types of connectors such as 3.5-mm and 2.9-mm for lower-frequency applications.

The solenoids are mounted in a precision common frame that is keyed to the main housing to ensure accurate alignment. The coil leads are brought out on a circuit board that also has the deenergizing switches mounted on it. The solenoids are a short-stroke design with rubber cushions at the stops for silent operation and longer life.

The main body is machined from brass and is gold plated.
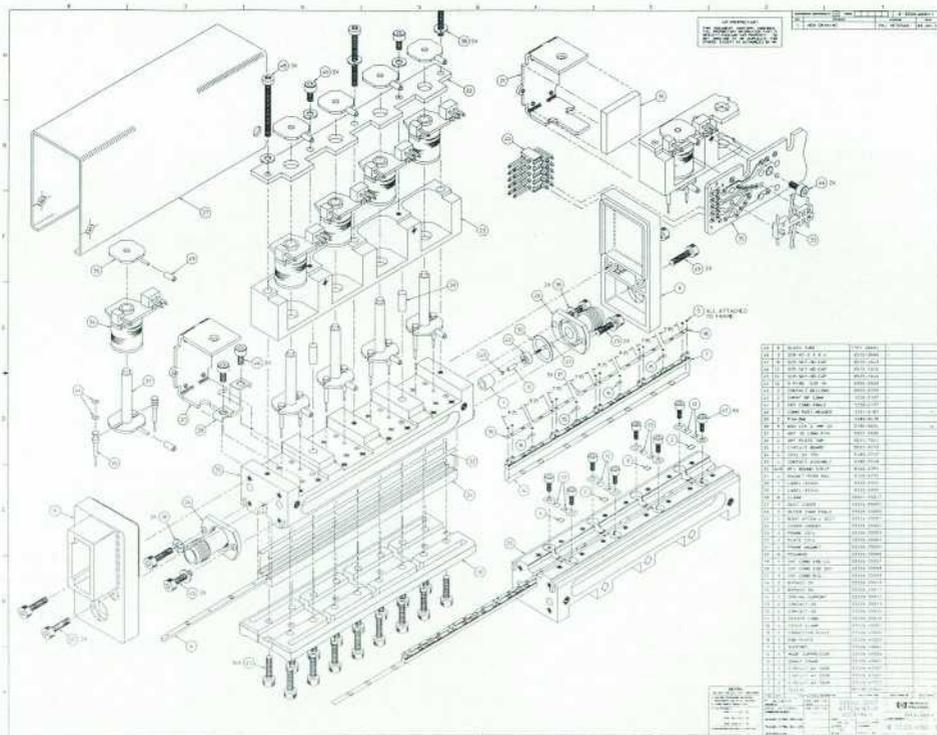
Fig. 4. Exploded view of a four-section attenuator.

This costs a bit more than aluminum but has outstanding long-term resistance to corrosion and contributes to the low residual loss of the device. A circular-cross-section RF gasket material is used between the cover and the body to ensure low leakage. External parts are nonstructural and are used for appearance and protection only. On three sides polycarbonate moldings are used, and the remaining area is enclosed by a sheet-metal dust cover. Mounting holes are provided in the main body with access through holes in the dust cover.

Performance

The performance of this attenuator design has proved to be excellent. The following are typical values:

- Frequency range: dc to 50 GHz
- Attenuation: 90 dB in 10-dB steps
- Residual loss: 1.3 dB at 26.5 Ghz and 2.3 dB at 50 Ghz
- Life: 5 million cycles minimum
- Attenuation accuracy: ±2%
- Repeatability: 0.01 dB
- SWR: 1.3 to 26.5 GHz
- Environmental performance: Exceeds requirements.

The HP 8360 attenuator has four sections with values of 10 dB, 20 dB, 30 dB, and 30 dB, for a total of 90 dB. The same attenuator design is available as a component in these and other values including 1-dB steps. Other connector options are also available for lower-frequency applications.

Fig. 4 shows an exploded view of a four-section attenuator.

Acknowledgments

50-to-110-GHz High-Performance Millimeter-Wave Source Modules

State-of-the-art microcircuit technologies and development tools were employed to produce a W-band amplifier tripler, a V-band amplifier doubler, an R-band amplifier doubler, and a coupler detector for two new frequency multiplier modules.

by Mohamed M. Sayed and Giovonnae F. Anderson

THE MILLIMETER-WAVE PORTION of the electromagnetic spectrum lies between the microwave and far infrared regions. Generally, millimeter-wave frequencies are between 30 and 300 GHz (wavelengths from 10 to 1 mm). The main applications for millimeter-wave systems are in communications, radar, and spectroscopic observation.

Millimeter waves are attenuated by atmospheric constituents and gases at different rates for different frequencies (see Fig. 1). Frequencies where gaseous absorption is a minimum are called *atmospheric windows*. Regions of maximum absorption are called *absorption bands*. The main millimeter-wave atmospheric windows are centered at 35, 94, 140, and 220 GHz, and the main absorption band is around 60 GHz.

The bandwidths of the atmospheric windows are wide—16, 23, and 26 GHz for the windows at 35, 94, and 140 GHz, respectively. In addition, millimeter-wave losses in these windows are low compared to infrared transmission in rain, fog, and smoke. These are advantages of millimeter-wave transmission. Other advantages are that millimeter-wave components are small, and a given antenna aperture yields a narrower beamwidth, higher precision, and better resolution than at microwave frequencies. For example, a 12-cm-diameter antenna produces a 1.8-degree beamwidth at 94 GHz, compared to 18 degrees at 10 GHz. Table I shows the key millimeter-wave characteristics for the windows at 35, 94, and 140 GHz.[1]

Table I
Millimeter-Wave Characteristics for Three Atmospheric Windows

Characteristic	Unit	Frequency (GHz) 35	94	140
Wavelength	mm	8.6	3.2	2.1
Clear Air Attenuation	dB/km	0.12	0.4	1.6
Rain Attenuation	dB/km			
0.25 mm/h		0.07	0.17	0.2
1 mm/h		0.25	0.6	0.7
4 mm/h		1.0	3.0	3.2
16 mm/h		4.0	8.0	9.0
Fog Attenuation	dB/km			
Light 0.01 g/m^3		0.006	0.035	0.007
Thick 0.1 g/m^3		0.06	0.35	0.7
Dense 1.0 g/m^3		0.6	3.5	7.0
Antenna Beamwidth for D = 15.24 cm	Degrees	4.0	1.4	0.98

New Millimeter-Wave Sources

Millimeter-wave sources are essential instruments for developing almost all millimeter-wave systems and for extending the range of microwave systems. The requirements for these sources are high output power, full waveguide band coverage, low output harmonics, and compatibility with other HP microwave sources.

The HP 83557A and HP 83558A millimeter-wave source modules have been developed to complete the waveguide coverage of the millimeter-wave range up to 110 GHz. These source modules take the output of an HP microwave source, such as an HP 8360, and multiply it to millimeter-wave frequencies. They are members of the HP 8355xA series, which was originally introduced in May 1985.[2,3] Fig. 2
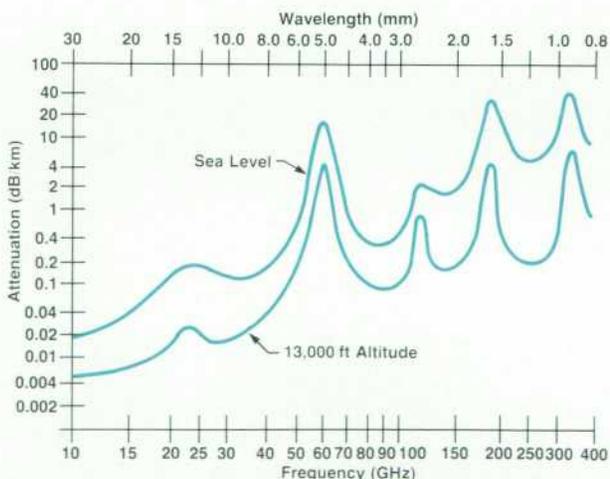

Fig. 1. *Attenuation of millimeter waves for one-way transmission through the atmosphere. There are atmospheric windows at 35, 94, 140, and 220 GHz and absorption bands at 60, 120, and 182 GHz.*
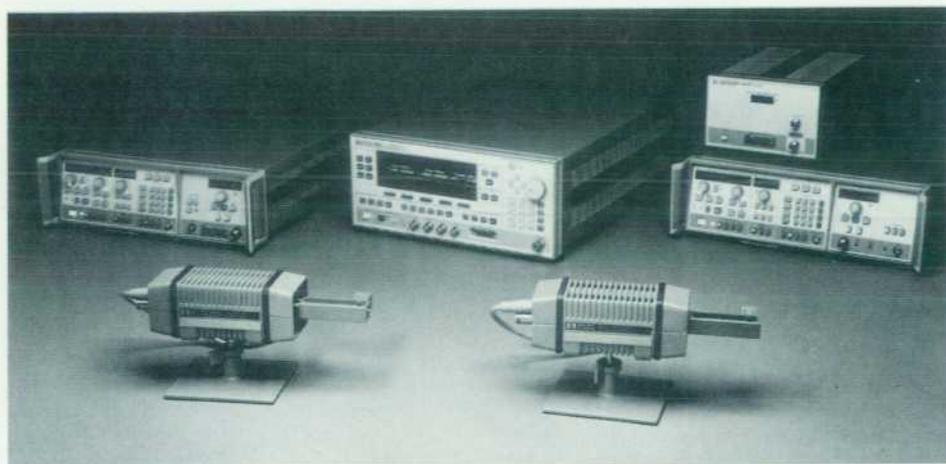
Fig. 2. *The HP 83557A 50-to-75-GHz source module (left foreground) and the HP 83558A 75-to-110-GHz source module (right foreground) are shown with compatible HP microwave sources: HP 83624A, HP 8350B/HP 83550A, and HP 8350B/HP 83595C with HP 8349B (not shown is the HP 8673C/D with HP 8349B). The millimeter-wave source modules receive 12.5-to-18.75-GHz signals from the microwave sources and multiply the frequencies to the millimeter-wave range.*

shows the two new source modules and the microwave sweep oscillators and synthesized signal generators with which they can be used. The HP 83557A covers the V-band range of 50 to 75 GHz and the HP 83558A covers the W-band range of 75 to 110 GHz. These two sources have high output power: +3 dBm for the HP 83557A and 0 dBm for the HP 83558A. Also available is the narrower-bandwidth HP 83558A Option H03, which supplies +3 dBm from 88 to 98 GHz. These new source modules can be used for either scalar or vector network analyzer applications.

The challenge in designing the new source modules was to supply state-of-the-art performance using the latest technologies and the features of the latest microwave sources such as the HP 8360 family (see article, page 6), without sacrificing backward compatibility with HP microwave sources introduced before 1985.

Design Philosophy

Design goals for the two new modules were:
- Higher output power than other sources on the market
- Frequency coverage of an entire waveguide band
- Low output harmonic content for use in scalar applications

- Use of the millimeter-wave source interface already defined for the rest of the HP 8355xA family before 1985
- Use of the same mechanical package as the rest of the HP 8355xA family
- Leveled output power with specifications equal to or better than other HP 8355xA modules
- Use of many of the user interface features offered by the HP 8360 family of sources.

The input frequencies for these two source modules were to be supplied by an HP 83550A millimeter-wave driver or other HP microwave sources along with an HP 8349B microwave power amplifier, which has a bandwidth less than 20 GHz. Thus the multiplier number was determined to be 4 for the HP 83557A and 6 for the HP 83558A, as shown in Fig. 3.

Two alternatives were then available: multiply by 4 or 6 directly or multiply by two and then amplify and multiply by two for V band and by three for W band. Output power, input dc power, and harmonics are very different for the two alternatives. Table II compares these two alternatives for V-band and W-band source modules. For example, to obtain +3 dBm output power in W band by multiplying by 6 directly, an input dc power of over 33 dBm or 2W is
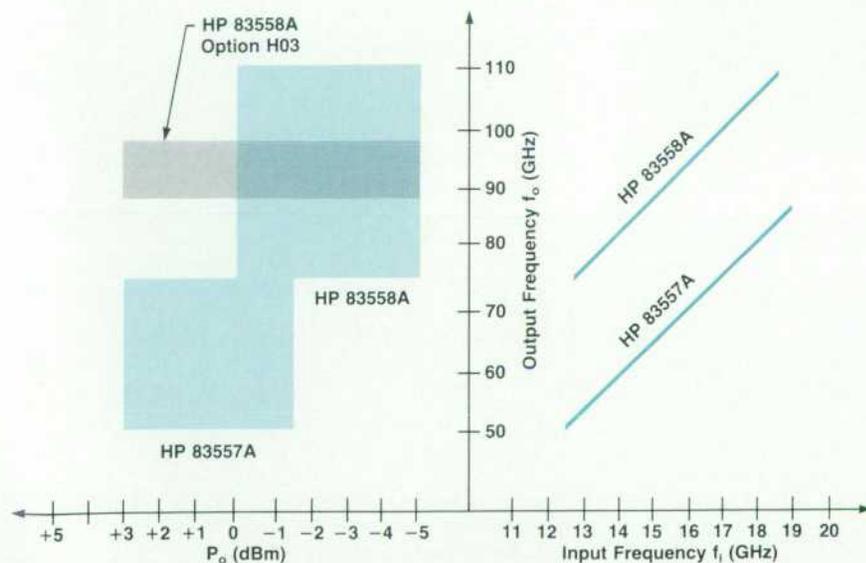

Fig. 3. *Output power and frequency ranges of the new millimeter-wave source modules.*
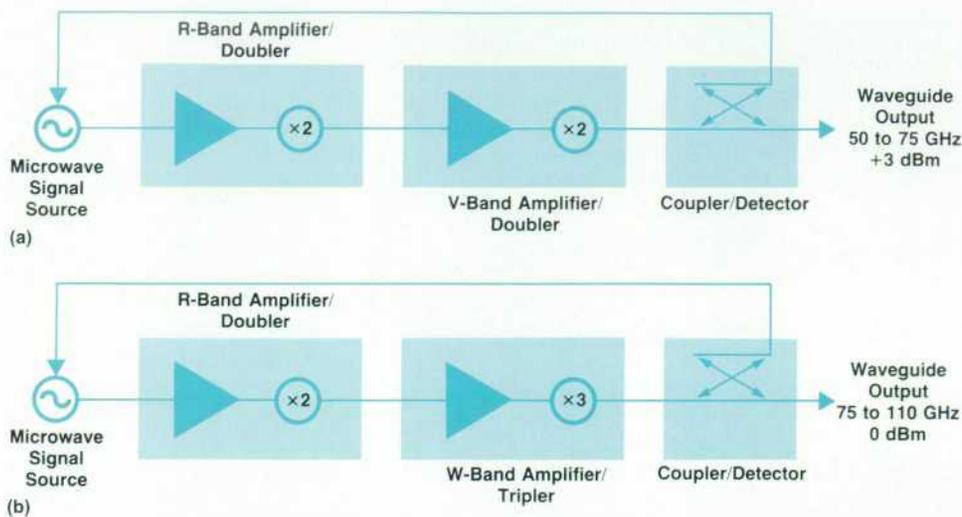
Fig. 4. *Block diagrams of (a) the HP 83557A 50-to-75-GHz source module and (b) the HP 83558A 75-to-110-GHz source module.*

required below 20 GHz. Assuming 20% conversion gain for the microwave amplifier at 20 GHz, over 10W dc is required from the HP 83550A source or HP 8349B. However, the power supply from either the HP 83550A or the HP 8349B to the HP8355xA millimeter-wave source module family is less than 8W. Therefore, the second alternative of multiplying by two, then amplifying and multiplying again was chosen. This alternative is called 2A2 or 2A3 in Table II. The input frequency of 12.5 to 18.75 GHz is first doubled to 25 to 37.5 GHz, then amplified and multiplied again to cover either V-band (with a doubler) or W-band (with a tripler). This method is also efficient in using the available dc power, since for a given amount of dc power, more gain can be obtained at 25 to 37.5 GHz than at 12.5 to 18.75 GHz.

2/3, a V-band amplifier doubler and a W-band amplifier tripler can be used with the same R-band amplifier doubler. This reduced the development time dramatically.

Although the harmonics are the same for direct multiplication by 4 and for the 2A2 method—or 6 and 2A3—the 2A2 and 2A3 methods offer two more degrees of freedom to reduce the mixing harmonics. As explained later, the harmonics of the R-band doubler can control the 5/4 and 7/6 harmonics of the V-band and W-band sources. In addition, the frequency response of the V-band amplifier doubler and the W-band amplifier tripler can control and may eliminate the 3/4 and 5/6 harmonics of the V-band and W-band modules, respectively. The sources used to drive the HP 83557A and HP 83558A have very low harmonic content, better than −50 dBc.

Another challenge to the design team was the physical size of the source module and the available dc power supply from the source. With no increase in the maximum RF input power and dc power supply, the project team had to deliver the same output power at higher frequencies (+3 dBm from 50 to 75 GHz for the HP 83557A). As Fig. 4 shows, this module is actually two sources in one package. These challenges made necessary a high-efficiency amplifier doubler for the V-band module and a high-efficiency amplifier tripler for the W-band module.

An external coupler detector is used to deliver high-

Table II
Comparison of Multiplication Alternatives
for Millimeter-Wave Source Modules

Multiplier Number	Ideal Conversion Loss (dB)	Realistic Conversion Loss (dB)	Band	Frequency (GHz)	Harmonic
2	8	12	R	26.5-40	3/2
	8	14	V	50-75	3/2
4	13	20	V	50-75	3/4,5/4, 6/4
2A2	0-3	3-7	V	50-75	3/4,5/4, 6/4
2	8	12	R	26.5-40	3/2
3	15	19	W	75-110	4/3
6	25	30	W	75-110	5/6,7/6, 8/6
2A3	6-8	8-12	W	75-110	5/6,7/6, 8/6

To make maximum use of the design tasks and components previously developed for earlier HP 8355xA source modules, the HP 83557A and HP 83558A use an R-band amplifier doubler as shown in the block diagram, Fig. 4. This component was designed for the HP 83554A (26.5 to 40 GHz). Since the ratio of the input frequency at the beginning of V band to that at the beginning of W band is exactly
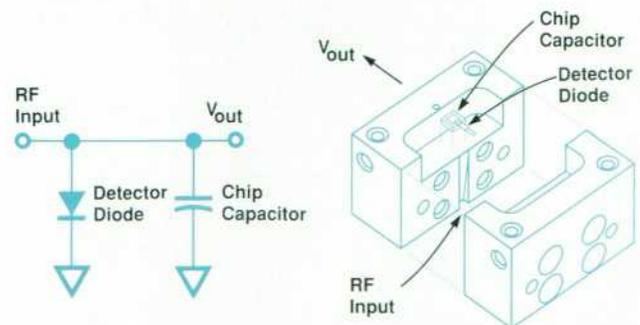

Fig. 5. *Schematic diagram and mechanical layout for the detector portion of the coupler detectors used in the V-band and W-band source modules.*

performance output power leveling. This coupler has superior directivity and therefore a superior source match is achieved. Higher unleveled output power can be obtained by removing the external coupler detector.

For use with the new HP 8360 synthesized sweepers (see article, page 6), a nonvolatile RAM containing calibration data is included in the source modules. This can be read directly by the HP 8360 family to improve the leveling flatness of the system (see "Flatness Correction," page 59).

V-Band and W-Band Coupler Detectors

For maximum flexibility in using the source with or without leveling, the coupler detector was designed to be separate from the V-band or W-band multiplier. (In the earlier HP 83554A, HP 83555A, and HP 83556A source modules, the coupler detectors were integrated with the multipliers). The coupler was designed and is manufactured in-house to obtain superior performance, ruggedness, and mechanical strength. (Special versions of the HP V752C and HP W752C are used). The schematic diagram and the mechanical layout for the V-band and W-band detectors are shown in Fig. 5.

The RF energy travels through a waveguide taper and impinges on a beam-lead diode at the end of the guide. One end of the diode is attached to a 6-pF capacitor, which rectifies the RF signal. The other end of the diode is attached to ground. The device used for this detector is a beam-lead planar doped barrier diode fabricated on a GaAs substrate with a 40-μm^2 active area.

The Use of the HP Microwave Design System in the W-Band Tripler Design

To achieve the maximum conversion efficiency of the tripler described in the accompanying article, it was necessary to provide a very good match to the diode at the tripler input frequencies (25 GHz to 36.66 GHz). An impedance matching filter was designed using the Microwave Nonlinear Simulator (MNS) which is a part of Hewlett-Packard's Microwave Design System (MDS). The design of the filter took place in three stages:

- Developing and simulating a model of the nonlinear device used for multiplication
- Designing and optimizing the filter
- Verifying the design experimentally.

The nonlinear model for the diode is shown in Fig. 1. The diode is a 20-μm^2 beam-lead device fabricated in an n + + LPE epitaxial layer on a GaAs substrate. The model was developed to take into account the parasitic inductances, capacitances, and resistances of the beam-lead diode chip. Values for these parasitic elements were calculated from the geometrical layout of the chip. Values for the diode parameters, such as saturation current and breakdown voltage, were determined by measurement and calculation. Fig. 2 shows the return loss of this device predicted by the nonlinear MDS model without any matching circuit. Because of our interest in the large-signal behavior of the diode, we used a harmonic balance simulation to determine the return loss of the unmatched device at three power levels: 18, 20, and 22 dBm. As one can see, the predicted return loss of around −3 dBm is quite significant. A nine-element Chebyshev filter was designed using radial stubs and microstrip transmission lines. The purpose
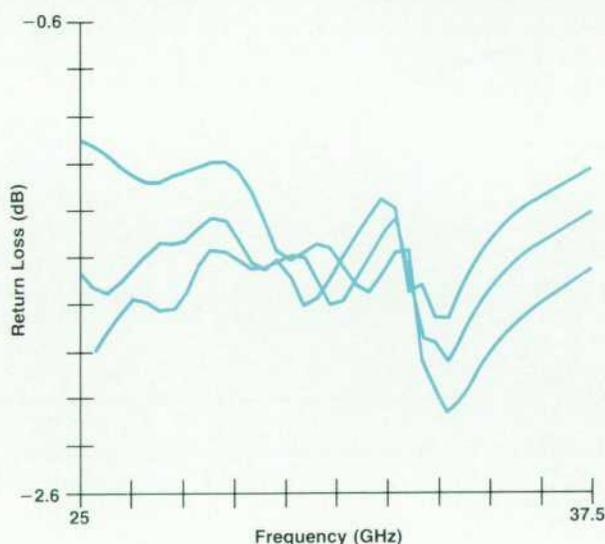
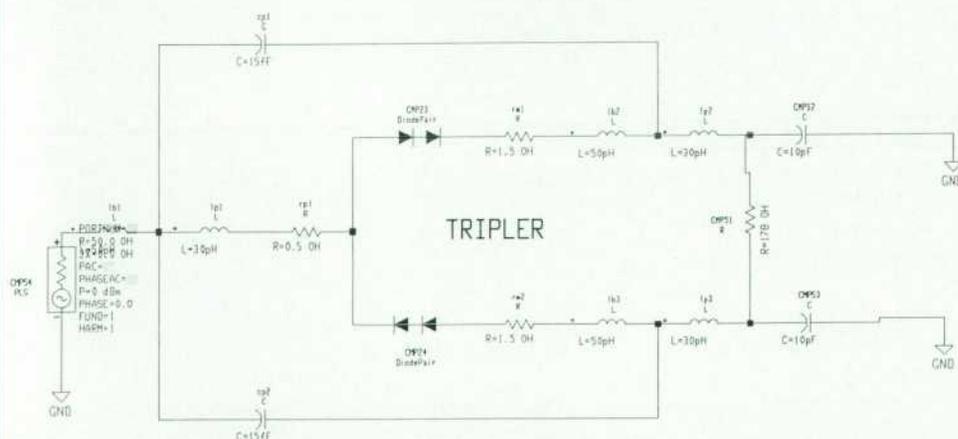**Fig. 2.** *Diode return loss predicted by the HP Microwave Design System at three power levels.*

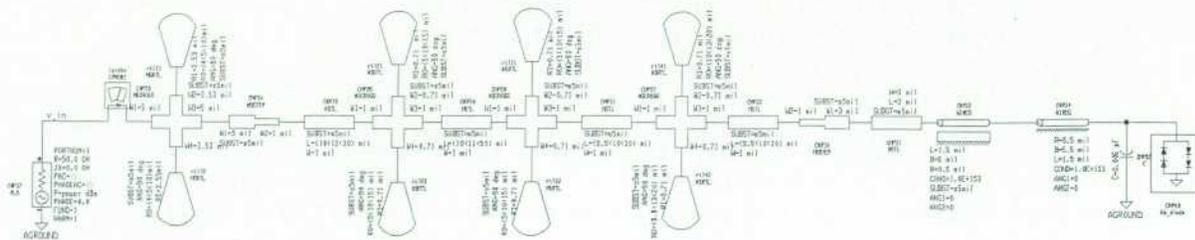**Fig. 1.** *HP Microwave Design System nonlinear tripler diode model.*

Fig. 3. Schematic layout of a nine-element Chebyshev low-pass filter.
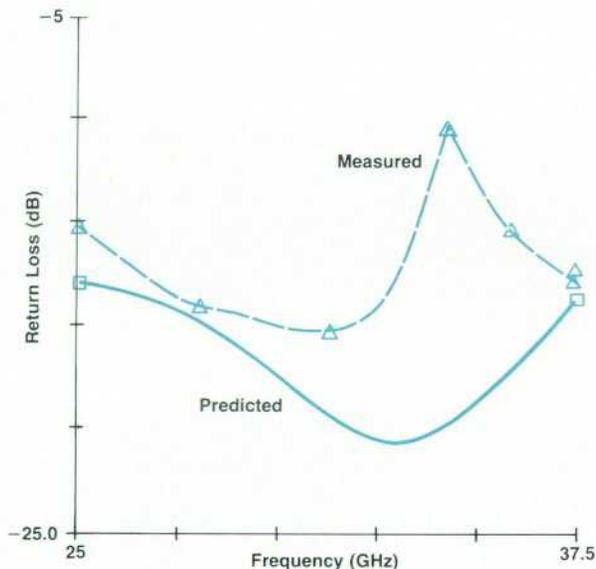
Fig. 4. Predicted optimized return loss of the filter.

of this filter is twofold: to provide an impedance match to the diode and to control the harmonics generated by the diode. The filter has a cutoff frequency of 45 GHz. The last capacitor of the filter is a 0.006-pF chip capacitor. This ensures the presentation of a short circuit to the higher harmonics of the diode. A schematic layout of the filter is shown in Fig. 3. MDS was used to optimize the outer radius of the stubs and the length of the transmission lines. The optimized return loss of the filter circuit, predicted to be between -15 dB and -23 dB, is shown in Fig. 4. The circuit is realized on a 0.005-inch sapphire substrate. Return loss measurements were made using an R-band traveling wave tube amplifier leveled at 22-dBm input power. Measured data (Fig. 4) showed reasonable agreement with the predicted data at the low end of the band but deviated at the high end. The deviation was caused by constraints on line widths and radial stub spacings in the microcircuit substrate fabrication.

Giovonnae Anderson
Development Engineer
Network Measurements Division

The waveguide taper acts as an impedance matching circuit. It reduces the height of the guide from 0.074 inch to 0.006 inch for the V-band detector and from 0.050 inch to 0.005 inch for the W-band detector. The V-band detector employs a linear (straight line) taper, but the W-band detector uses an exponential cosine taper for improved input match.[4] The taper in the W-band detector also contains a high-pass filter, which attenuates below-band harmonics which would have degraded flatness. Fig. 6 shows the V-band and W-band coupler detector output voltages as functions of frequency.

W-Band Amplifier Tripler

Odd harmonics are typically generated by symmetrical distortion of an incoming waveform. For example, if an ideal sinusoidal waveform is clipped symmetrically by a circuit, then the output of that circuit will contain only odd harmonics. The circuit used for the W-band tripler is shown in Fig. 7. The diode used is a 20-μm^2 beam-lead device[5] fabricated on an n++ epitaxial layer on a GaAs substrate.

The filters are designed for proper signal routing. The design of the low-pass filter is crucial to the multiplier design since it provides the impedance match to the diode and a short circuit to the higher harmonics generated in the multiplier.[6] For more discussion on the design of this

Fig. 6. Output voltage of the W-band (a) and V-band (b) coupler detectors for a source module output power of -5 dBm.

Fig. 7. *Diode tripler antiparallel configuration.*

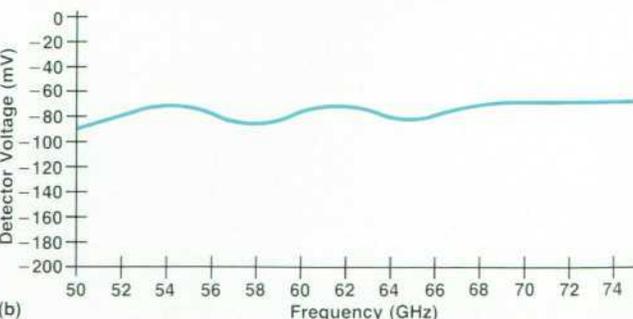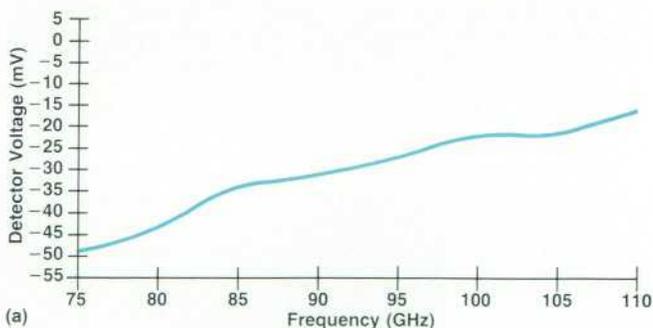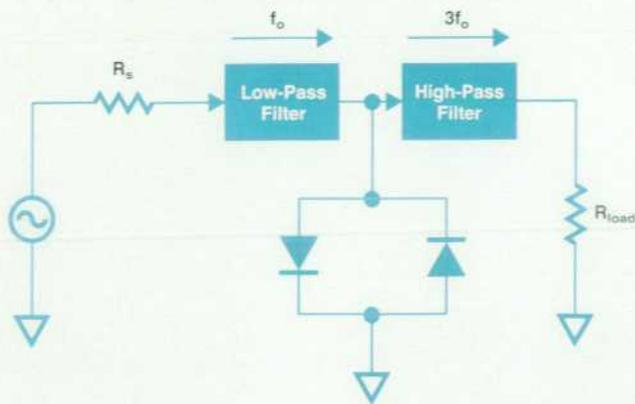filter, see "The Use of the HP Microwave Design System in the W-band Tripler Design," page 53. The final design is realized as a nine-element modified Chebyshev filter with a cutoff frequency of 45 GHz. The last element of the filter is a chip capacitor fabricated in-house with a value of 0.006 pF. This chip capacitor gives the filter excellent out-of-band response because it presents an excellent short circuit to the higher harmonics. The output of the tripler is ac-coupled to the waveguide. The high-pass filter is an exponential cosine taper in the waveguide, similar to the one used in the detector; it provides an open circuit at the plane of the diode to frequencies below 70 GHz. The taper also reduces the height of the guide at the diode end to lower the impedance to 45 ohms to provide a better match to the third harmonic. The physical layout of the multiplier is shown in Fig. 8.

When the source power is too low there is insufficient drive to turn on the diodes and the conversion efficiency is degraded. When the source power is too high, the conversion efficiency is also degraded because the diodes continue to clip at a voltage determined by the barrier height. These effects are shown in Fig. 9 (unbiased curve). Another consequence of increasing power in this circuit is very high diode current which can result in poor device reliability.

The output power from the multiplier is specified at +3 dBm at 110 GHz with a conversion efficiency of better than 19 dB. Self biasing of the multiplier diode is necessary for this goal to be achieved. The bias circuit for the tripler
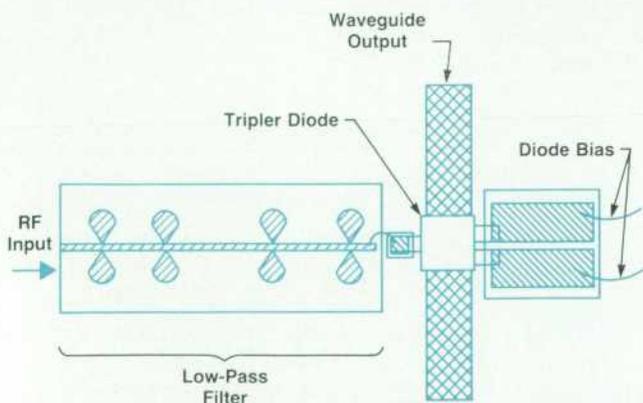
is shown in Fig. 10. The RC time constant is made purposely long in comparison to the RF signal driving the circuit. This results in rectification of the RF and the presence of a positive or negative dc voltage on the internal nodes. The magnitude of the dc voltage depends on the magnitude of the incoming signal. For very high RF drive levels, the dc voltage increases. The diodes will not conduct until the RF voltage exceeds one diode drop above the dc voltage. This means that clipping of the incoming signal occurs at higher voltages and conversion loss is improved. The conversion efficiency is shown in Fig. 9 for biased and unbiased circuits. The conversion efficiency at +18 dBm is about the same as before. For higher power levels, however, the conversion efficiency is improved and remains relatively constant until device failure occurs. The resistor and capacitor values are chosen to allow adequate charge to be bled off the internal node while the diodes are reverse biased by the RF.

Because the diodes now conduct over a smaller portion of the RF cycle, there is a smaller component of direct current flowing in them. This results in less metal migration than in the unbiased tripler. However, because of the presence of the self-bias voltage, the reverse voltage across the diodes is increased. To prevent breakdown, the diodes were redesigned to maintain the same overall capacitance and series resistance.

Bias is also needed to improve the conversion efficiency at low drive levels. For low signal levels, the diode will not turn on as it does in the unbiased tripler. To level the multiplier at −5 dBm a trickle current of about 100 μA is required. The dc voltage across the diodes can be adjusted to improve the harmonic balance of the multiplier. Because of the sensitivity of the input filter circuit to an additional bond wire, there is no center tap to ground for the multiplier. Therefore, the diodes cannot be independently biased.

Fig. 8. *Physical layout of the diode tripler.*
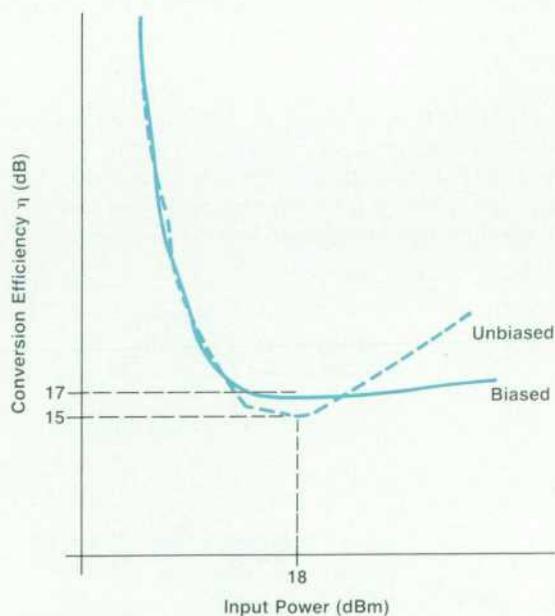
Fig. 9. *Conversion efficiency of unbiased and biased diode triplers.*
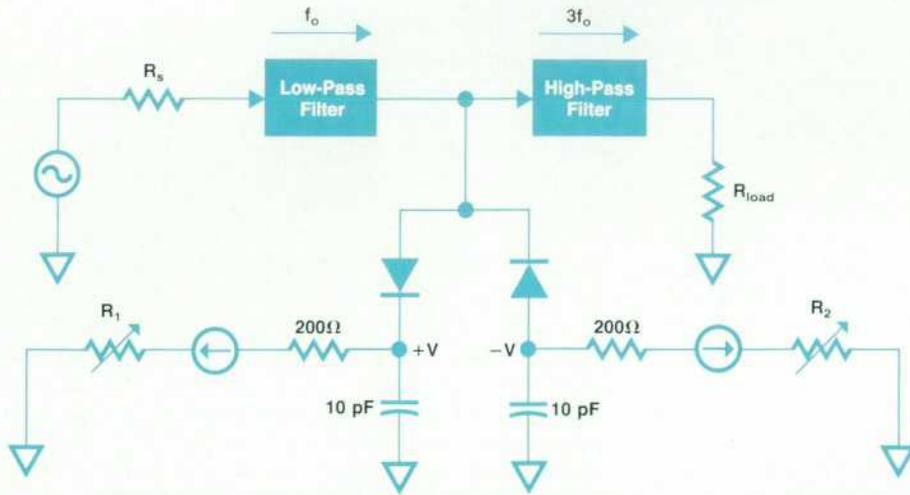
Fig. 10. *W-band tripler with bias circuit.*

The amplifier portion of this component is a thin-film circuit with integrated capacitance fabricated on a 0.005-inch sapphire substrate using a laser milling process. The amplifier has eight stages of amplification to supply the required RF power to the W-band tripler stage. A unit cell of this amplifier is shown in Fig. 11a. It consists of a two-stage GaAs FET amplifier with input, output, and interstage matching networks. This unit cell is used to achieve the required power as shown in Fig. 11b. The same amplifier is also used in the V-band amplifier doubler.[3,7,8]

For a discussion on the design of the W-band tripler package, see "The Use of HP ME 10/30 in the W-Band Tripler Design," page 57.

V-Band Amplifier Doubler

Basic doubling action occurs when a sinusoidal waveform is transformed into a fully rectified output. The output waveform contains only even-order harmonics. This is achieved in both the R-band doubler and the V-band doubler by using a coplanar waveguide-to-slotline structure. The basic structure is shown in Fig. 12.

The fundamental signal is input to the doubler via the coplanar waveguide. The characteristic impedance of the coplanar waveguide input line is 50 ohms at the coax transition, but is transformed down to 30 ohms at the narrowest part of the gap to present a better match to the diode. The coplanar waveguide structure is particularly useful for attaching a beam-lead diode to the circuit. Wire bonds connect the ground planes together, and the unbalanced mode

that exists in coplanar waveguide propagates through these wire bonds unaffected. The fundamental signal is terminated by the doubler diode. Charge flow through the diodes results in the balanced mode across the slotline, giving rise to only even harmonics. The wire bonds are attached to the ground plane approximately one quarter wavelength (doubled frequency) away from the diode, so any second harmonic propagating towards the input on the coplanar waveguide line is reflected in phase with the output. A thin-film capacitor provides RF continuity for the fundamental and doubled signals, while allowing biasing capability for the diode. The second harmonic propagates through the slotline transition and into the next part of the circuit.

The placement of the wire bonds that connect the ground planes is crucial. These bonds are close enough to the diode to appear as an inductive backshort, which resonates with the junction capacitance of the diode. Movement of the diode also affects the 3/2 harmonic of the doubler, since the phasing can be either positive or negative. Movement of the wire bonds closer to the diode will, in general, give a slight improvement of the conversion efficiency at the expense of the unwanted harmonics. The wire bonds extend well up the coplanar line. In the case of the V-band doubler, the wire bonds extend up the entire line. Removal of bonds near the RF pin usually results in reduced conversion efficiency.

Bias is used with both doublers to improve conversion efficiency at higher powers, and to provide leveling loop

(a)

(b)
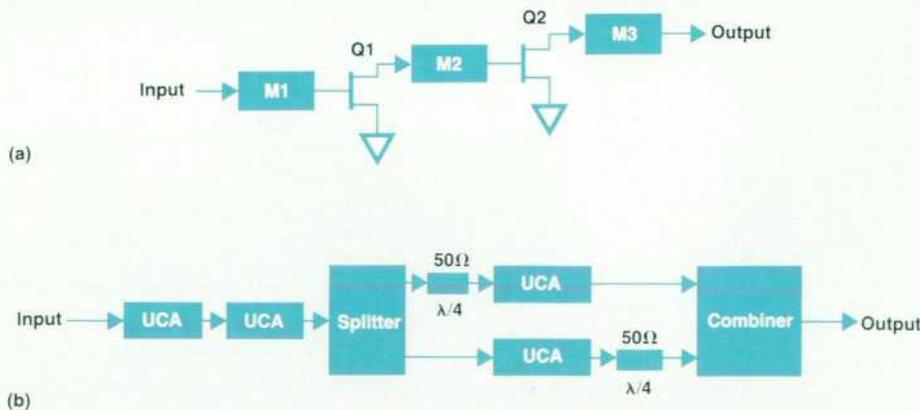
Fig. 11. *(a) Unit cell amplifier (UCA) for the amplifier portion of the W-band amplifier tripler. (b) Amplifier portion of the W-band amplifier tripler.*
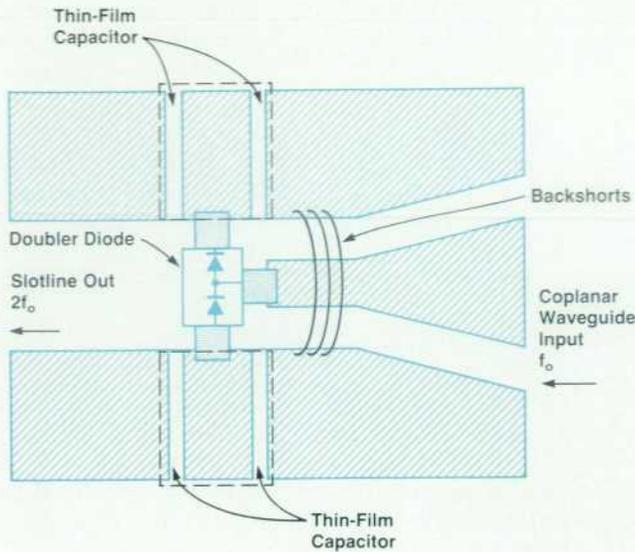
Fig. 12. *Coplanar waveguide input structure.*

stability at lower drive levels. Thin-film capacitors are used as dc blocks for bias currents. Because of the nature of the coplanar structure, a ground return can be added from the center RF line to the ground plane with no degradation of the multiplier. This allows each diode to be balanced separately.

In V band, the doubled signal propagates through a slotline/finline transition and into waveguide, as shown in Fig. 13. The finline circuit dielectrically loads the waveguide, and it is possible to set up surface modes called longitudinal section electric or LSE modes. These modes appear in the output power as sharp, high-Q dips. Three possible
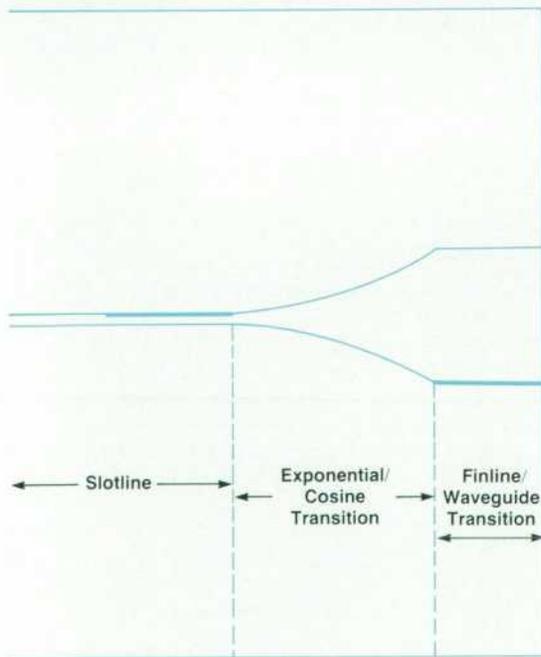
Fig. 13. *Slotline-to-finline transition in the V-band amplifier doubler.*

The Use of HP ME 10/30 in the W-Band Tripler Design

The use of drafting tables, pencils, and paper for design and drafting at HP is rapidly disappearing and being replaced by computers, displays, and keyboards. The primary mechanical design software is HP ME 10 for two-dimensional work and HP ME 30 for solids modeling.

The package for the W-band tripler microcircuit described in the accompanying article was designed using ME 10. The required views were constructed with exact geometry so that other areas involved in documentation and manufacturing would see the correct configuration. This file was then electronically transferred to the model shop for the fabrication of prototype parts. Since this package is relatively complex, consisting of four machined parts, it was translated into an ME 30 solid body for better visualization. The tapered waveguide equation, developed using the HP Microwave Design System (MDS), was also sent to the model shop, routed through a translator, and entered into the ME 30 solid body.

All of the ME 30 information was routed through another translator and then the prototype packages were machined using a computer-controlled machine tool. Once necessary revisions or corrections to the package were implemented, the file was transferred to the drafting area, where a simple command in ME 30 generated all six views of a part and the part was formally dimensioned and formatted for manufacturing specifications. The production machining area was also able to import these files into its own system with very little time for production programming restructuring. Fig. 1 shows the flow chart of the procedure used.

Reference

1. S. Lockhart, "Changing World Markets: the CAID Mandate at HP," *Innovation*, Vol. 8, no. 2, Spring/Summer 1989, pp. 26-29.

Roy Marciulionis
Development Engineer
Network Measurements Division
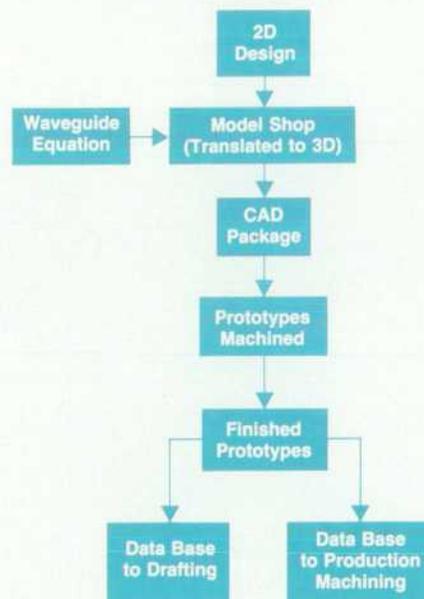
Fig. 1. *Flow chart of the process leading from ME 10 design to final drafting and production.*

methods of suppressing these modes are: changing the substrate to a lower-dielectric material, changing the substrate thickness, and reducing the height of the waveguide to push the cutoff frequency above-band.

This problem is overcome by changing the substrate from 0.010-inch sapphire ($\epsilon_r = 9.6$) to 0.010-inch fused silica ($\epsilon_r = 3.2$) and reducing the height of the waveguide. The choice of material and substrate thickness is a compromise between the need to suppress the LSE modes, the impedance requirements for achieving the proper input match to the diodes, and the feasibility of fabricating the thin-film capacitors necessary for diode biasing. The waveguide height is reduced from 0.074 inch to 0.070 inch to ensure that the cutoff frequency is above-band. The doubling device for both the R-band and the V-band doublers is a 40-μm^2 Schottky diode on an n++ liquid phase epitaxy layer on a GaAs substrate.

Placing the dielectric circuit into the waveguide is mechanically difficult. For this reason, a split block approach is employed. As shown in Fig. 14, the waveguide housing is in two parts, and one part includes machined cavities above and below the waveguide. These cavities are one quarter wavelength (at midband) from the waveguide and filled with polyiron. The polyiron ensures that no modes can exist in the cavities. The cavities themselves present an open circuit, which transforms to a short circuit at the edge of the waveguide.

R-Band Amplifier Doubler

The input signal of 12.5 GHz to 18.75 GHz (maximum input power = 15 dBm) is first amplified by an MMIC (microwave monolithic integrated circuit) amplifier. This amplifier is a reactively matched two-stage MESFET amplifier designed to give 22-dBm output power from 12 GHz to 20 GHz. The amplifier roll-off below 12 GHz provides rejection for the lower harmonics coming from the source. Like many amplifiers, this one had a tendency to oscillate if the input and bias leads were of any significant length. Therefore, several large monoblock capacitors were added around the circuit to keep it from oscillating. The multiplier operates on the previously discussed frequency

Fig. 15. *Slotline-to-microstrip transition of the R-band amplifier doubler with circuit model.*

doubling principle. The circuit is realized on a 0.010-inch sapphire substrate. The frequency-doubled signal propagates down the slotline and through a transition (transformer) into microstrip line. The slotline-to-microstrip transition and its model are shown in Fig. 15. It is formed by overlapping the slotline and the microstrip line.

Amplifier Multiplier Performance

Typical output power performance data for the W-band amplifier tripler, the V-band amplifier doubler, and the R-band amplifier doubler is shown in Fig. 16. The different technologies, devices, and thin-film parameters used in

Fig. 14. *Split block design of the finline-to-waveguide transition of the V-band amplifier doubler.*

Fig. 16. *Typical output power of the W-band amplifier tripler, the V-band amplifier doubler, and the R-band amplifier doubler.*

component development for the HP 83557A and HP 83558A millimeter-wave source modules are shown in Table III.

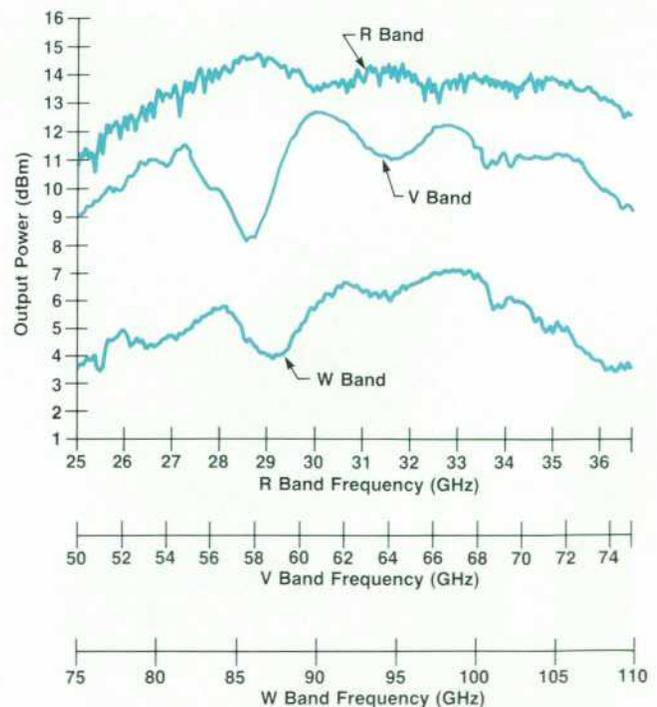### Table III
Technology Used for Components of the HP 83557A and HP 83558A Source Modules

Component	Devices	Thin-Film	CAE
Coupler detector	Beam-lead diode	N/A	ME 10/30 mechanical design
W-band amplifier tripler	Beam-lead diode, GaAs FET	0.005-inch sapphire, integrated capacitors	MDS, ME 10/30 mechanical design
V-band amplifier doubler	Beam-lead diode, GaAs FET	0.005-inch sapphire, integrated capacitors, 0.010-inch fused silica	ME 10/30 mechanical design, MDS
R-band amplifier doubler	Beam-lead diode, MMIC	0.010-inch sapphire	ME 10/30 mechanical design, MDS

Source Module System Evaluation

Once the individual components were developed and met their specifications, the source module was assembled with the component bias board and the source interface board. A few challenging issues were discovered and resolved by the project team for the source and the system to meet specifications.

Dynamic Range. The V-band source module is specified to deliver leveled output power between +3 dBm and −2 dBm. The maximum output power, +3 dBm, determines the required minimum input power to the source module. For the HP 83624A and HP 83550A sources, the minimum
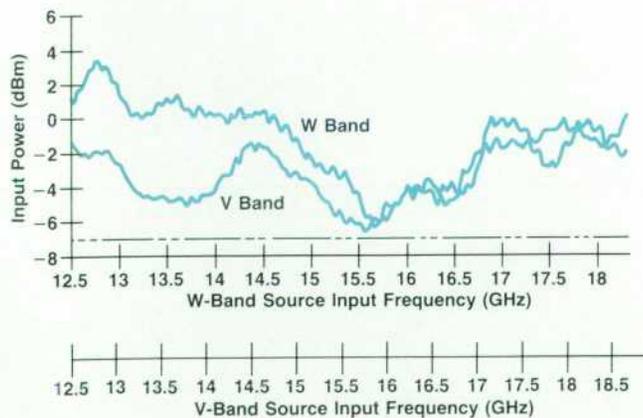
Fig. 17. *Input power for −2-dBm output power for the V-band source module and for −5-dBm output power for the W-band source module.*

Flatness Correction

Among the many features of the HP 8360 family of synthesized sweep oscillators, the flatness correction capability offers the millimeter-wave source module user better output power flatness through the module's own internal flatness correction data and an HP 8360 user-definable flatness calibration routine.

The digital flatness correction for the millimeter-wave source modules works in much the same manner as the internal flatness correction that is standard in all HP 8360 series synthesized sweepers. The flatness correction array is stored in nonvolatile RAM in each millimeter-wave module. It contains offset factors in dB to compensate for the unflatness of the module's coupler and detector. These numbers are stored at frequency intervals corresponding to the module's multiplication factor times 100 MHz (nominal). The V-band module, for example, is a × 4 module and carries offset data every 400 MHz. The W-band module is × 6 and has correction data every 600 MHz. This array is read into the HP 8360's memory whenever a millimeter module is connected and an instrument preset is executed. The module flatness correction then resides in a special array in the instrument. The HP 8360 uses this information to correct the module's coupler and detector flatness at a finite number of frequencies across the millimeter waveguide band. Linear interpolation is performed in the HP 8360 for frequencies between each flatness correction point, resulting in some flatness errors typically on the order of 0.2 dB. Another source of flatness error is the nonlinearity of the detector and logger. The flatness correction data is stored at the specified maximum leveled output power. As the power is lowered, small flatness errors occur, with approximately 0.2 dB error at the minimum specified output power. However, all flatness errors over the power level and frequency ranges are less than 0.5 dB, leaving only the uncertainty of the power measurement in question.

If better flatness than this is required either across the entire waveguide band or over a smaller subset of frequencies one can choose to implement the user flatness correction. The user flatness capability allows the operator to set the user correction array start/stop frequencies with up to 505 frequency points. The user flatness correction is performed at the current output power setting, allowing the operator the choice of output power level. With an HP 437B power meter connected to the HP 8360 over the HP-IB interface and the power sensor connected to the output of the millimeter-wave module, the HP 8360 can measure the power and store the flatness offset data into its user flatness array automatically. Alternatively, a user-written calibration program and any power meter can be used. Once the calibration is complete the operator can choose any subset of frequencies within the user flatness correction array and continue to use that portion of the user correction array.

Lon Dearden
Development Engineer
Network Measurements Division

power is +17 dBm at 18.75 GHz. For other sources that deliver lower power than +17 dBm at 18.75 GHz, the HP 8349B amplifier is recommended. The minimum output power, −2 dBm, determines the required minimum settable power to be delivered by the source. This power is −5 dBm for the HP 83550A, HP 83595C, and HP 83592C without the attenuator option. Therefore, the source mod-

ule is tested with +15-dBm input power to deliver output power of +4.25 dBm from 50 to 75 GHz, taking into consideration the leveling loss and the measurement uncertainty. The source module is also tested for −2-dBm output power with the minimum settable power from the source of −5 dBm from 12.50 to 18.75 GHz. The module's gain and gain flatness are adjusted to meet the specification. The typical input power to the source module to meet its −2-dBm specification is shown in Fig. 17. A similar test is performed for the W-band source module.

Leveling Sensitivity. The V-band source module can be configured with HP 8360, HP 8673C/D, HP 83550A, HP 8340/41B, and HP 8359x sources. The conversion efficiency of the YIG-tuned multiplier of each of these sources is linear in the normal operating output power range. However, at very low output power, it can become nonlinear, which can cause ALC peaking and leveling loop oscillations. Oscillations can also occur if the conversion efficiencies of the R-band amplifier doubler and the V-band amplifier doubler change rapidly with output power variations. Therefore, the V-band source module was also tested with a sample of all applicable sources to verify its ALC stability. The output frequencies were varied and the input power variation was monitored. The worst-case output frequency is used to test every source module to guarantee source system stability. Typical leveling sensitivity of the V-band source module is shown in Fig. 18, which is for 62.5 GHz. The 0.5V/GHz tuning voltage is used to adjust the source module gain and gain shape to meet the required system parameters for stability. A similar test is also performed for the W-band source module.

Harmonic Analysis. The output harmonic content of the V-band source module is generated by the 3/4, 5/4, and 6/4 harmonics as shown in Table II. The desired 6/4 or 3/2 harmonic is generated within the V-band doubler and can be measured and adjusted separately on the subassembly level. The analysis for the 3/4 and 5/4 harmonics is shown in Fig. 19. The desired multiplied signal from the R-band doubler is f_1 and the unwanted signal is $3/2 f_1$ or f_2. Both signals are amplified and doubled again to generate the wanted signal f_o. However, two unwanted signals are generated within the V-band doubler and propagated to the output. The upper portion of f_2 (50 to 56.25 GHz) is $3/4 f_o$.
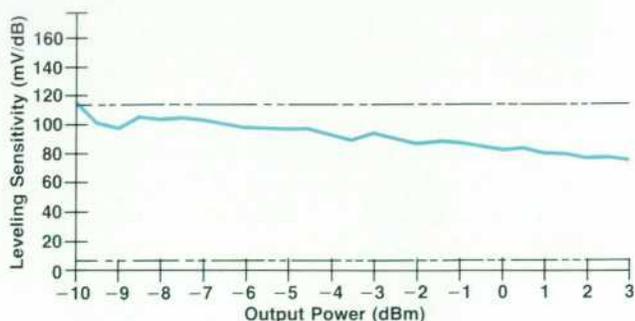
Fig. 18. Leveling sensitivity of the V-band source module.

Since the multiplier is a very efficient mixer, f_1 and f_2 mix with each other and with the wanted output signal to generate $5/4 f_o$. The $3/4 f_o$ product is easily eliminated since the amplifier portion of the V-band amplifier doubler has high attenuation in that frequency range. The $5/4 f_o$ signal is directly related to the third harmonic of the input frequency or f_2 and can also be adjusted by tuning the R-band amplifier doubler. This is a major advantage of the 2A2 scheme of multiplying over the straight ×4 scheme.

For the W-band source module, the unwanted output signals are $5/6 f_o$, $7/6 f_o$, and $8/6 f_o$ as shown in Fig. 20. The $8/6 f_o$ or $4/3 f_o$ product is generated within the W-band tripler and can be measured and adjusted separately at the subassembly level. Again, because the tripler is an efficient mixer, it adds and subtracts the f_1 and f_2 signals to produce $5/6 f_o$ and $7/6 f_o$. The $5/6 f_o$ product is proportional to the gain of the amplifier portion of the W-band amplifier tripler. Since this gain is really an attenuation at 50 to 55 GHz, the $5/6 f_o$ product is less than −40 dBc. However, the $7/6 f_o$ product is proportional to the amplifier gain at 25 and 37.5 GHz, which is high. Thus the R-band amplifier doubler is adjusted to reduce the f_2 signal output, which in turn reduces the $7/6 f_o$ output.

Output Power Level Flatness. The output power can vary by about +3 dB without any correction. Thus an eight-breakpoint scheme is used to correct for the output power flatness. For better than +1-dB flatness, nonvolatile RAM inside the source module contains corrections for level flatness and can be read by the HP 8360 synthesized sweeper
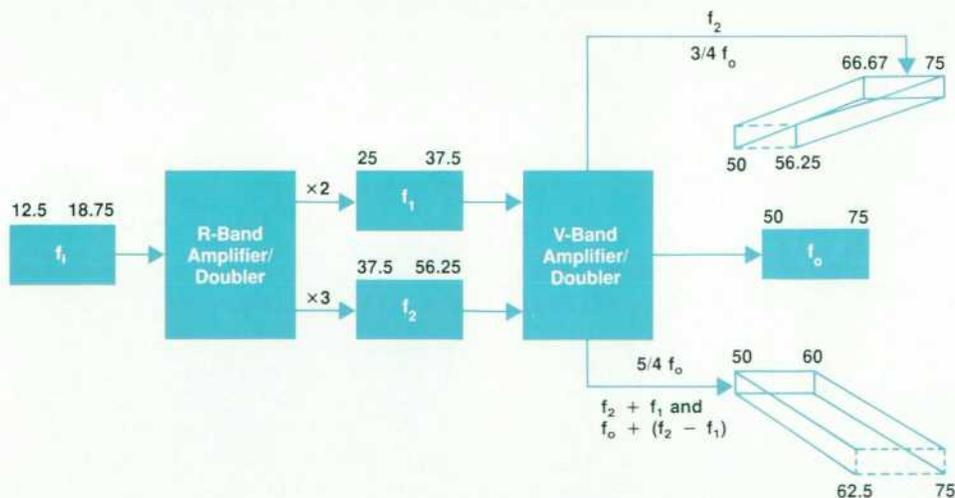
Fig. 19. Harmonic analysis for the V-band source module.
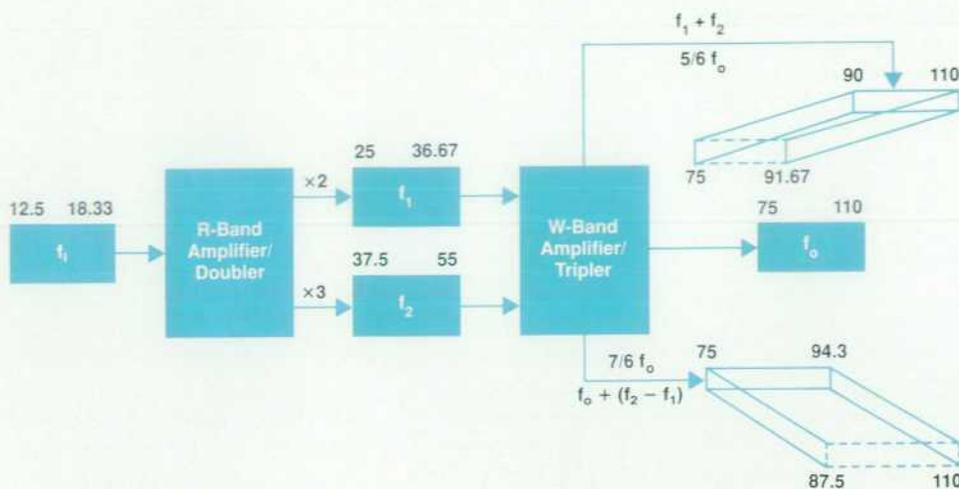
Fig. 20. *Harmonic analysis for the W-band source module.*

family as described on page 6.

Source Module Performance

To test the two new source modules, a production test system was designed. This test system is an extension of the one already used to test the HP 83554A, HP 83555A, and HP 83556A. The challenge was to measure the output power and harmonic level accurately and traceably. Each subassembly of the production test system is calibrated separately. Typical subassemblies include a directional coupler with a load on the straight arm and a power sensor on the coupler arm, a directional coupler with a load on the straight arm and a harmonic mixer on the coupler arm, a switch, and a variable short. Since no primary standards were available from the U.S. National Institute of Standards and Technology at the time,[9] in-house capabilities for secondary standards had to be developed. Typical performance of the V-band and W-band source modules, as measured by the production test system, is shown in Figs. 21 and 22.

Scalar Analysis Application and Performance

A typical W-band scalar network analyzer system is shown in Fig. 23. The performance of a bandpass filter, as measured by this system, is shown in Fig. 24. The traces for ac and dc detection are almost identical, which implies that the error caused by overshoot on the square wave modulation is very small. For devices with a high quality factor (Q), the dc detection scheme is recommended.

Vector Analysis Application and Performance

A typical W-band vector network analyzer system, the HP W85106B, is shown in Fig. 25. The measured performance of 90-GHz device under test is shown in Fig. 26, which shows that the measurement system has a dynamic range of over 95 dB at W band.

Acknowledgments

Every successful project is the result of the extraordinary efforts of many individuals who contribute their expertise, creativity, drive, and vision. We are very grateful for the opportunity to recognize the contributions of the following individuals. Lon Dearden was responsible for the original

---

High-Power W-Band Source Module

The output power of the W-band source module is a function of the sum of the conversion losses of the R-band amplifier doubler, the W-band amplifier tripler, and the coupler detector. Since these conversion or insertion losses are optimum in the middle of the W-band region, high output power can be obtained in this region. The HP 83558A Option H03 source module was created to satisfy the need for high power around the 94-GHz atmospheric window. The output power is maximized by optimizing the W-band amplifier tripler between 88 to 98 GHz and is higher than +3 dBm. The 5/6 and 7/6 harmonics are out-of-band for this model. The flatness is also better for this frequency range. Fig. 1 shows the output power and flatness of the HP 83558A Option H03.

Mohamed Sayed
Project Manager
Network Measurements Division
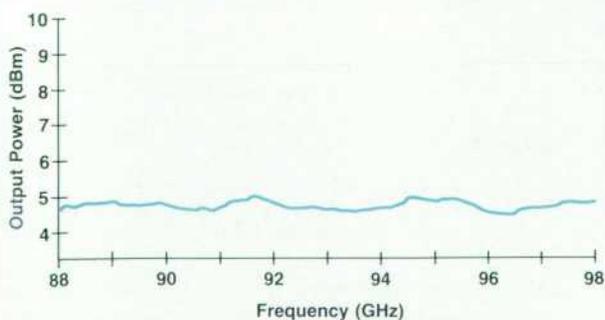
Fig. 1. *Output power and flatness of the HP 83558A Option H03 source module.*

design and technologies for the millimeter-wave amplifier. Steve Goedeke spearheaded the design implementation and the realization of the state-of-the-art performance for this amplifier. Roy Marciulionis designed all of the microcircuit packages, and was also responsible for product design with Ivan Hammer. Jim Grace performed the HP 8510 system evaluation. Gratz Armstrong, Bob Hasenick, and
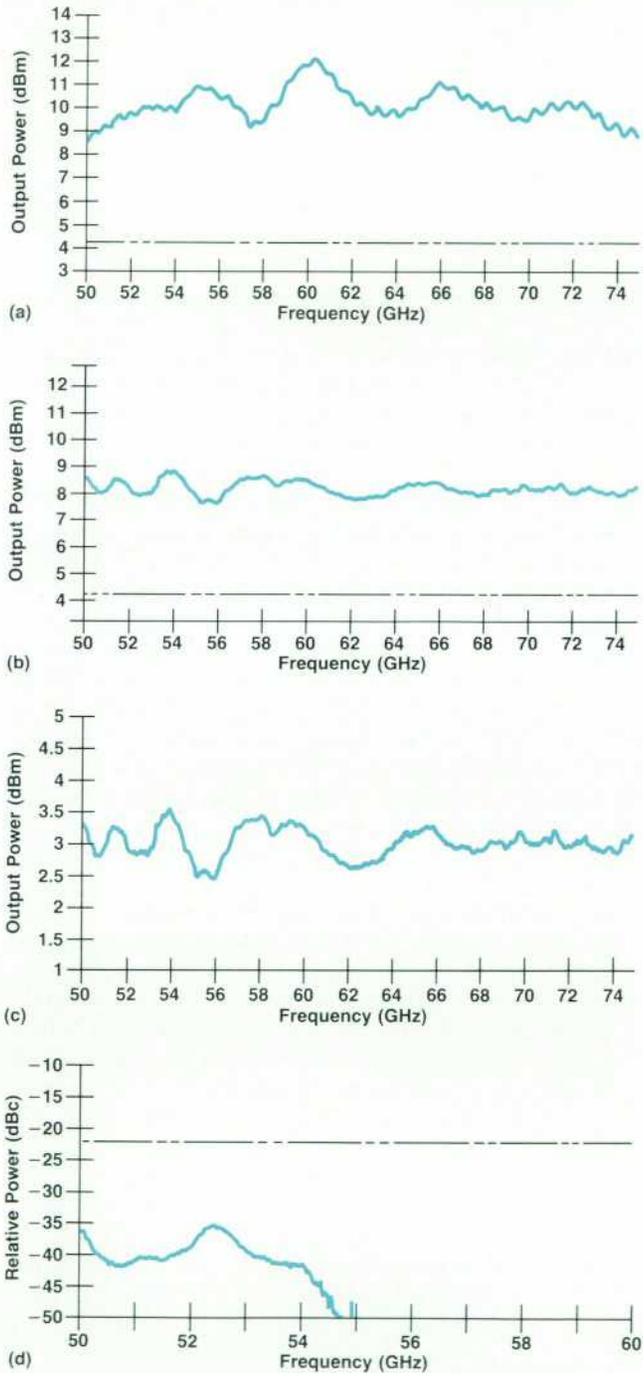
Fig. 21. Typical performance of the HP 83557A V-band source module. (a) Typical unleveled output power. (b) Typical leveled output power. (c) Typical flatness at +3-dBm output power. (d) Typical harmonic performance (5/4 harmonic).
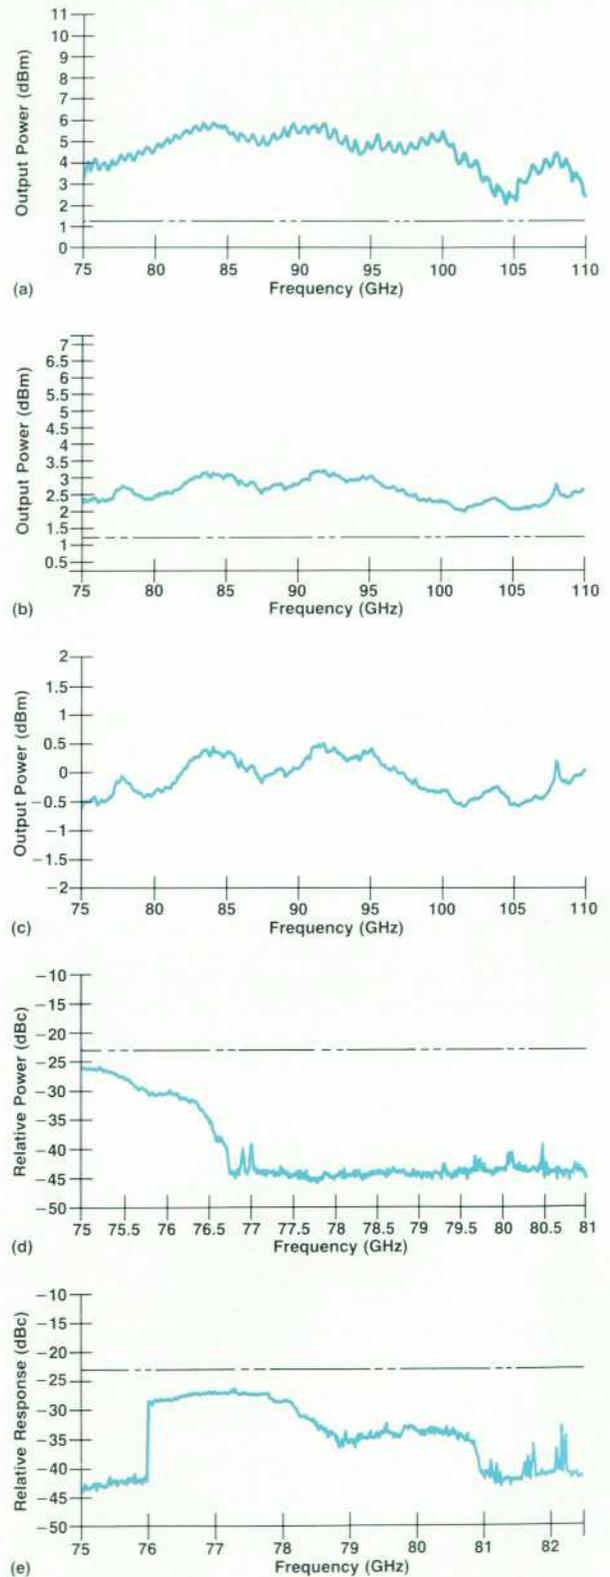
Fig. 22. Typical performance of the HP 83558A W-band source module. (a) Typical unleveled output power. (b) Typical leveled output power. (c) Typical flatness at 0-dBm output power. (d) Typical mixing harmonic (7/6). (e) Typical multiplier harmonic (4/3).
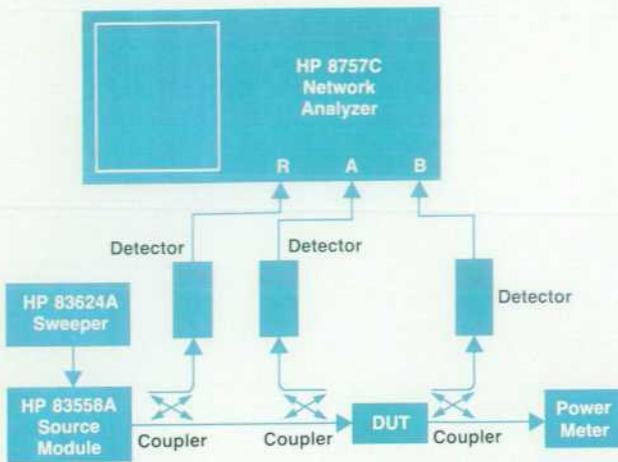
Fig. 23. *Typical scalar network analyzer system.*
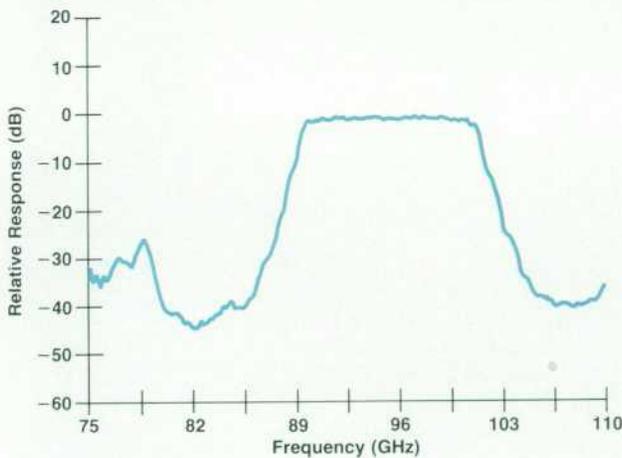
Fig. 24. *Bandpass filter characteristic measured with the scalar analyzer system of Fig. 23.*

Archie Fraser of the microwave test accessory operation designed and developed the detector and the special version of the directional couplers. The thin-film circuit, GaAs diode IC, and MMIC amplifier were fabricated at the Microwave Technology Division. Curtis Bonner was the liaison for the thin-film circuit and Janice Pryst was the liaison for the multiplier diodes. Hiroshi Kondoh designed the MMIC amplifier, and Ron Hogan helped in releasing it to production. John Regazzi, Dale Albin, Frank David, Dave
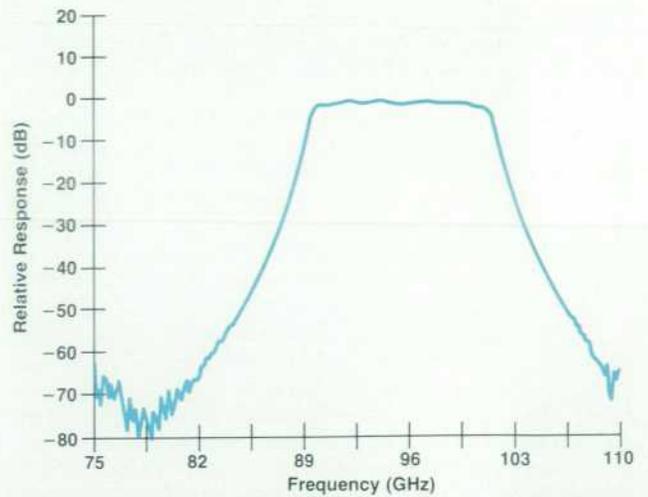
Fig. 26. *Bandpass filter characteristic measured with the HP W85106B network analyzer system.*

Fig. 25. *HP W85106B vector network analyzer system.*

References

1. C.R. Seashore, J.E. Miley and B.A. Kearns, "MM-Wave Radar and Radiometer Sensors for Guidance Systems," *Microwave Journal*, August 1979, pp. 47-58.

2. M.M.. Sayed and J.R. Regazzi, "Millimeter-Wave Sources and Instrumentation," *Hewlett-Packard Journal*, Vol. 39, no. 2, April 1988, pp. 6-11.

3. R.D. Albin, "Millimeter-Wave Source Modules," *ibid*, pp. 18-25.

4. F. David, *Analysis and Synthesis of Exponentially Tapered, Non-Uniform Transmission Line Impedance Transformers*, MSEE Thesis, Oregon State University, 1975.

5. G.F. Anderson, "GaAs Beam-Lead Antiparallel Diodes for MM-Wave Subharmonic Mixers," *International Electron Devices Meeting Digest*, December 1981, pp. 688-691.

6. W.Y. Lau, "Network Analysis Verifies Models in CAD Packages," *Microwaves and RF*, November 1989, pp. 99-110.

7. L. Dearden, G. Miner, and M. Sayed, "Model-Extrapolated S-Parameter Design of MM-Wave GaAs FET Amplifiers," *1986 IEEE MTT-S Digest*, pp. 385-388.

8. L.A. Dearden, *Model-Extrapolated S-Parameter Design of MM-Wave MESFET Amplifiers*, MS Thesis, Brigham Young University, Provo, Utah, 1986.

9. Private Communication, F.E. Marler, Microwave Metrology Group, National Institute of Standards and Technology, Boulder, Colorado.

An Instrument for Testing North American Digital Cellular Radios

The HP 11846A is designed to produce filtered π/4 DQPSK modulated I and Q baseband signals needed to test digital cellular radios.

by David M. Hoover

RAPID GROWTH IN THE DEMAND FOR AMPS (Advanced Mobile Phone Service) cellular radios in North America has caused saturation in current analog cellular frequency bands. This has spurred the development of a new digital cellular standard that will increase the spectrum efficiency of cellular radios to allow more users to share the same frequency spectrum. The Telecommunications Industry Association (TIA) has commissioned a group to define this next-generation cellular radio system. Known as the TR 45.3 committee, this group of industry representatives has issued specifications defining the new digital cellular system. The new system requires radios to conform to the old AMPS analog specification as well as the new digital system, hence the name North American Dual-Mode Cellular System (NADMCS). The dual nature of the radios was deemed necessary to ensure compatibility with current cellular systems. These radios will operate as analog AMPS radios in areas that do not require the increased capacity offered by the digital system, and in the digital mode in high-use areas. The new dual-mode radios will require testing for both the analog and digital operating modes. The analog tests can be made with currently available test equipment. The HP 11846A π/4 DQPSK I-Q generator and the HP 11847A π/4 DQPSK modulation measurement software are two products designed to test the radios' digital mode.

The NADMCS uses a digital modulation format known as π/4 DQPSK. This system is a time-division multiple-access system that allows up to six users access to a given frequency channel concurrently. Because of the backwards compatibility with the current AMPS cellular system, the channel spacing continues to be 30 kHz. The digital information symbol rate is 24.3 kHz, and to control the spectral energy from the digital transmission, a square root raised cosine filter with a roll-off factor of 0.35 is used to smooth the phase transitions. The HP 11846A π/4 DQPSK I-Q generator is designed to produce filtered I-Q baseband signals needed to create the modulation format used by the NADMCS. The HP 11846A approach for implementing this modulation format yields a simple yet effective technique for generating π/4 DQPSK modulation. When used with an I-Q generator, such as the HP 8780A vector signal generator,[1] the HP 11846A can provide accurate π/4 DQPSK modulated signals for testing North American dual-mode digital receivers. An overall block diagram of the HP 11846A is shown in Fig. 1.

The HP 11847A π/4 DQPSK modulation measurement software performs accurate verification of the RF performance of cellular transmitters conforming to the TR 45.3 committee recommendations. This software package uses digital signal processing techniques to demodulate the RF signals, recover the data, and measure the modulation accuracy. The HP 11847A digital signal processing techniques result in excellent modulation measurement accuracy and repeatability.

This article provides some background information about digital modulation format and raised cosine filters, and then covers the implementation of the HP 11846A filtered π/4 DQPSK modulation scheme. The HP 11847A measurement software is described in the article on page 73.
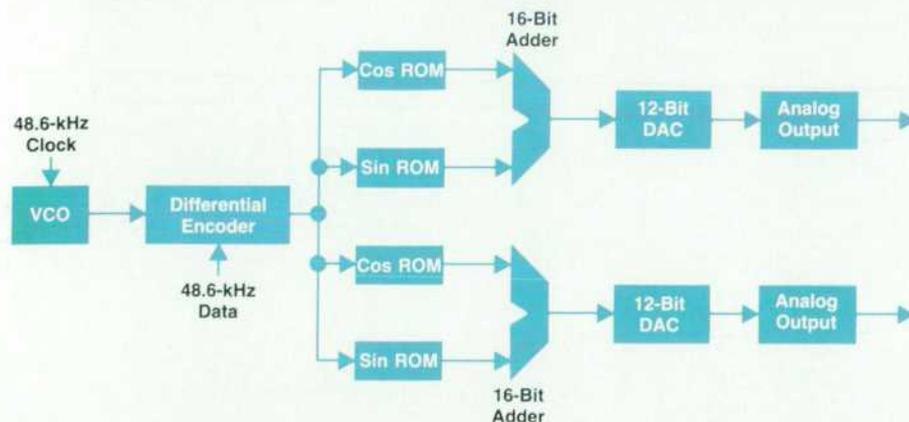
Fig. 1. *HP 11846A π/4 DQPSK generator block diagram.*

(a)

(b)

(c)
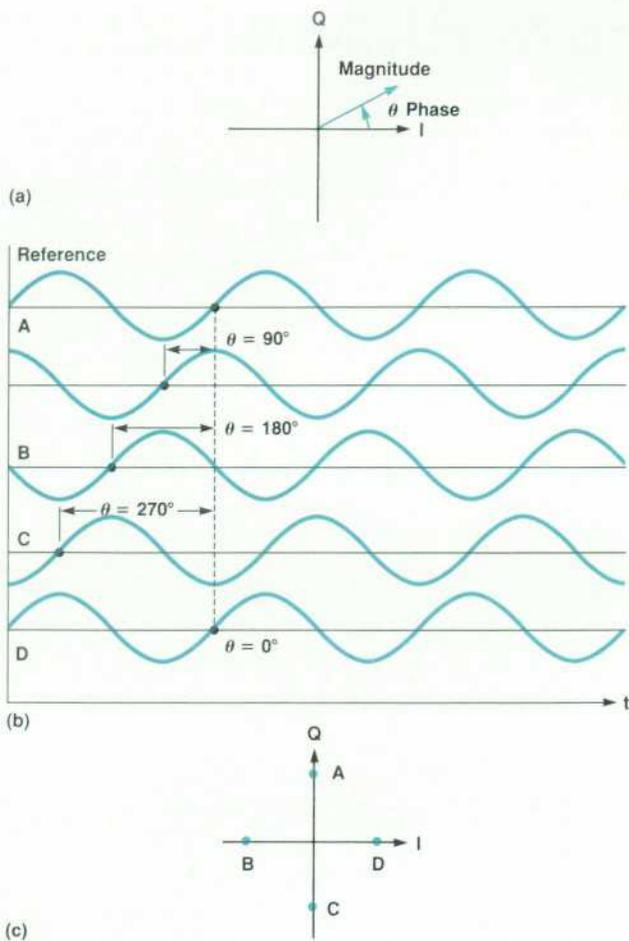
Fig. 2. (a) The I-Q plane. (b) Signals of varying phase. (c) Their I-Q diagram. This type of I-Q diagram is also called a constellation diagram.

Fig. 3. On/off keying I-Q diagram.

(a)

| State 00 | State 01 | State 11 | State 10 |
| $\theta = 0°$ | $\theta = 90°$ | $\theta = 180°$ | $\theta = 270°$ |

(b)

Fig. 4. (a) I-Q diagram of QPSK modulation. (b) QPSK timing diagram.

Digital Modulation

The modulation format chosen for the NADMCS system is $\pi/4$ differential quadrature phase shift keying ($\pi/4$ DQPSK). To understand how this modulation format works, we will look at some basic building blocks for digital modulation.

I-Q Diagrams. I-Q diagrams are frequently used to analyze the performance of digital communication systems. The modulation signal has two components, called in-phase (I) and quadrature (Q). The I axis and the Q axis form a coordinate system that presents the magnitude and phase of the signal being analyzed. The length of the vector from the origin represents the magnitude of the signal, and the angle of the vector referenced from the positive I axis indicates the phase of the signal (Fig. 2). To determine the phase and magnitude of a digital signal, a reference signal must be available. On the transmitting side, the reference signal comes from a local stable oscillator. On the receiving side, the reference is usually derived from the incoming phase-modulated signal. Fig. 2b shows four signals of varying phase and their associated I-Q diagram.

Digital communication systems use a combination of RF signals of specific phase and magnitude to represent specific bit patterns. The simplest method of doing this is

on/off keying. In this method a digital one is transmitted when the radio carrier is on for one data clock period and a digital zero is transmitted when the carrier is off for one period. The I-Q diagram for on/off keying is shown in Fig. 3. Note in Fig. 3 that there are two states, one for each bit of information transmitted.

QPSK Modulation. QPSK, or quadrature phase shift keying, uses one of four phase states to represent a particular data symbol. An I-Q diagram of QPSK is shown in Fig. 4. The reference chosen for zero degrees is arbitrary. In this system, two bits are required to define the possible phase states, because for any data symbol, there are four potential phase states. As shown in Fig. 4b, determining the phase relationships of digital bit patterns requires a known phase reference for comparison.

DQPSK Modulation. DQPSK, or differential quadrature phase shift keying, is a modification of QPSK modulation. This format also has four potential phase states, which require two bits per data symbol. The difference here is that the phase states are defined relative to the last phase state. This means that the absolute phase of the system is not required because the phase of the current symbol is determined from the phase state of the previous symbol. The phase state of a symbol is determined by the phase transition defined for that symbol. For example, consider the phase transitions defined for the following symbols.

Symbol	DQPSK Phase Transition
00	0°
01	90°
10	−90°
11	180°

Based on these definitions, symbols arriving in the following order would have the associated phase states, assuming an initial state of 0°.

Arrival Time	Symbol	New Phase State
t	00	0°
t+1	10	−90°
t+2	10	180°
t+3	00	180°
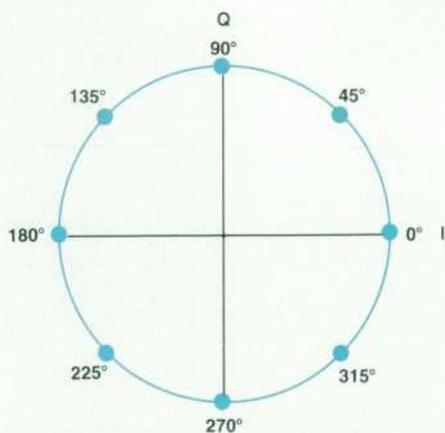
Thus, the phase state of the current symbol is dependent on the phase state of the previous symbol and the phase transition defined for the current symbol.

π/4 DQPSK Modulation. The π/4 DQPSK modulation system is very similar to DQPSK, in that it is a differential quadrature phase shift system. The difference is that in the I-Q plane, the possible phase states rotate 45 degrees for each symbol. Table I compares the phase transitions for DQPSK and π/4 DQPSK.

Table I
Comparison between DQPSK and π/4 DQPSK Modulation Formats

Symbol	DQPSK Phase Transition	π/4 DQPSK Phase Transition
00	0°	45°
01	90°	135°
10	−90°	−45°
11	180°	−135°

The constellation of points on the I-Q diagram now has a total of eight points, but only four are possible for any given symbol. Fig. 5 shows the possible constellation of points for π/4 DQPSK. If the initial phase is at 0 degrees, the possible phase states could be 45, 135, 225, and 315 degrees. For the next symbol, the possible phase states would be 0, 90, 180, and 270 degrees. This pattern continues, allowing a total of 8 phase states for π/4 DQPSK.

The key advantage of π/4 DQPSK modulation is that the spectral energy can be contained to a fairly small bandwidth for a given data rate compared to other modulation formats. In π/4 DQPSK modulation, in which the phase transitions are ±45 degrees and ±135 degrees, spectral energy can be controlled by keeping the phase transition per symbol low, and also by reducing the signal amplitude during large phase transitions. Fig. 6 shows unfiltered phase transition trajectories for π/4 DQPSK. The ±45-degree rotations have small amounts of phase rotation, so the

Fig. 5. I-Q diagram of π/4 DQPSK modulation.
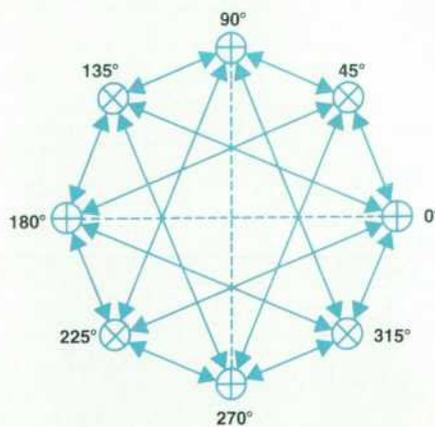
Fig. 6. Unfiltered phase transitions for π/4 DQPSK. The trajectories inside the octagon represent 135° transitions. For example, from 0° there can be a transition to 135° or 225°. The arrows around the outside of the octagon represent 45° transitions.

amplitude remains virtually constant. During the larger ±135-degree transitions, however, the amplitude is significantly reduced, which helps control the spectral energy.

Filtering

The unfiltered phase transitions for $\pi/4$ DQPSK shown in Fig. 6 produce sharp trajectories in the I-Q plane, which result in a large amount of spectral splatter during phase transitions. To control the spectral splatter, filtering is applied to the baseband I-Q signals. The filter used in the NADMCS is a square root raised cosine filter with an α (roll-off factor) of 0.35. The raised cosine filter not only minimizes spectral splatter, but also reduces intersymbol interference. Intersymbol interference is caused by the effects of all the filtering at various locations (transmitter, channel, and receiver) throughout a typical baseband digital system (see Fig. 7). This makes data detection more difficult, and filtering is applied to reduce intersymbol interference.

The theoretical minimum system bandwidth needed to detect 1/T symbols/s without intersymbol interference can be shown to be 1/2T Hz. For this case, a rectangular filter shape in the frequency domain is required. This type of filter is difficult to approximate, but if the filter bandwidth is increased, the approximation task is made much easier. This modification to the filter bandwidth is defined by a term called the filter roll-off factor, or α. Let W_o represent the theoretical minimum bandwidth, 1/2T Hz. Let W represent the bandwidth of the filter. The roll-off factor is defined to be $\alpha = (W - W_o)/W_o$. The roll-off factor specifies the required excess bandwidth divided by the filter's −6-dB bandwidth. Fig. 8a shows the frequency response of a raised cosine filter for several values of α. Notice how the amplitude response is 6 dB down at the theoretical minimum bandwidth point, regardless of the value of α.
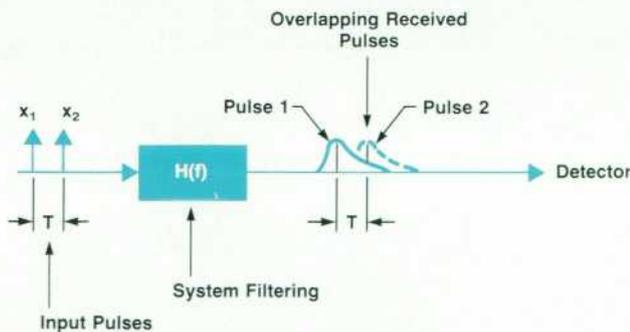
Fig. 9 shows the impulse response of a raised cosine filter with an α of 0.35. By examining this response, one can see that the impulse response crosses through nulls at symbol decision points.* This implies no intersymbol interference when this filter is used.

To achieve optimum signal-to-noise ratio, a matched filter situation must be used. This implies a similar filter in the transmitter and receiver. To accomplish this, the raised cosine filter frequency response is modified by taking the square root of this function. This yields a square root raised cosine filter. The α of the filter is the same as the original raised cosine filter. Fig. 10 shows the impulse response of a square root raised cosine filter with an α of 0.35. Notice that this impulse response does not have nulls at symbol decision points. This indicates that there will be intersymbol interference for this filter characteristic. The intersymbol interference problem is rectified because the receiver has another square root raised cosine filter. When the transmitted data is filtered by the receiver's square root raised cosine filter, the received data should have no intersymbol interference. This is the filter arrangement chosen for the NADMCS. Fig. 11 shows a better comparison of the impulse response of the raised cosine and square root raised cosine filters.

*Decision points are the times when the RF data is demodulated to determine the transmitted data. The integer values along the abscissas in Figs. 8 and 9 represent decision points.
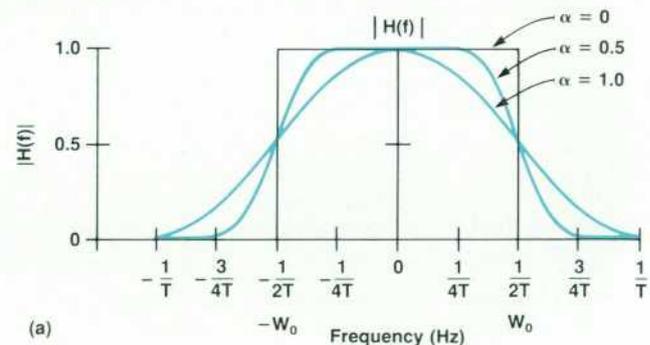
(a)
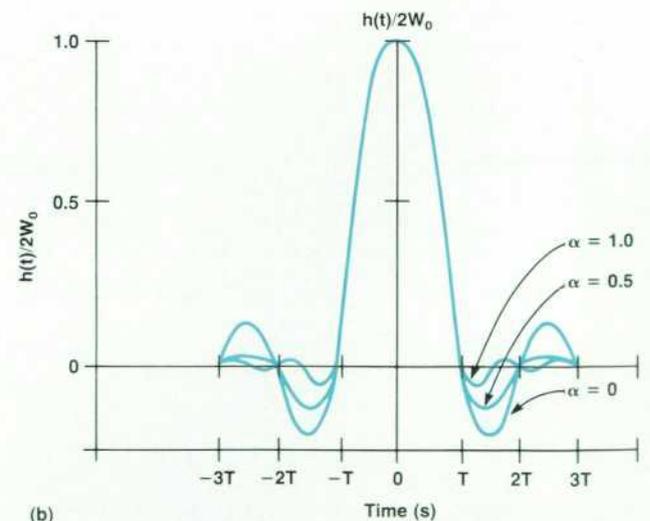
(b)

Fig. 8. *Raised cosine filter characteristics. (a) System transfer function for several values of α. The $\alpha = 0$ roll-off is the minimum bandwidth case. (b) System impulse response.*

Fig. 7. *Intersymbol interference. The incoming pulses overlap, making data detection more difficult. H(f) is the equivalent transfer function for all the filtering effects in the system (i.e., the transmitting filter, the channel filter, and the receiving filter).*
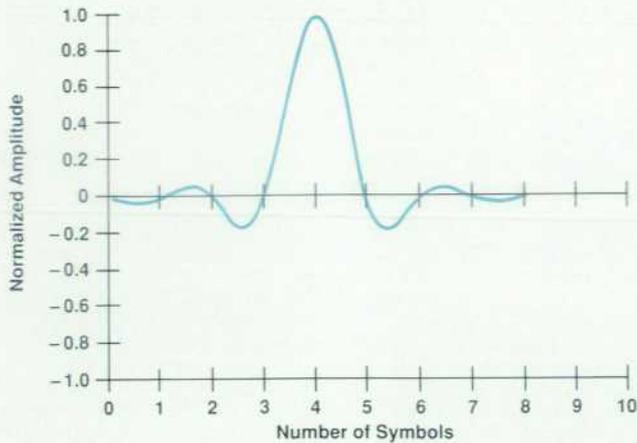
Fig. 9. *Impulse response of a raised cosine filter with an α of 0.35.*

Filtered π/4 DQPSK Modulation

The HP 11846A generates a square root raised cosine filtered I-Q output using an FIR (finite impulse response) type of digital filter. The HP 11846A implementation uses a ROM-based FIR filter, which allows a fairly simple hardware design. The block diagram of a conventional ROM-based filter for π/4 DQPSK is shown in Fig. 12. The incoming data bits are input to the serial-to-parallel converter which separates the first and second bit of each symbol into separate data paths. The symbol data is then differentially encoded and enters a shift register. The length of the shift register is determined by the FIR filter length. In the case of the HP 11846A, the FIR filter length is eight, making the output of the shift registers eight-bit parallel data. Outputs from both shift register banks are applied to the ROM address lines, giving 16 address bits. The ROMs perform a convolution of the input data with the filter's impulse response to produce the I-Q output data. The box on page 71 provides more detail about the ROMs' convolution operation.

The HP 11846A updates the I-Q outputs at a rate 16 times the NADMCS symbol rate. This allows the HP 11846A to
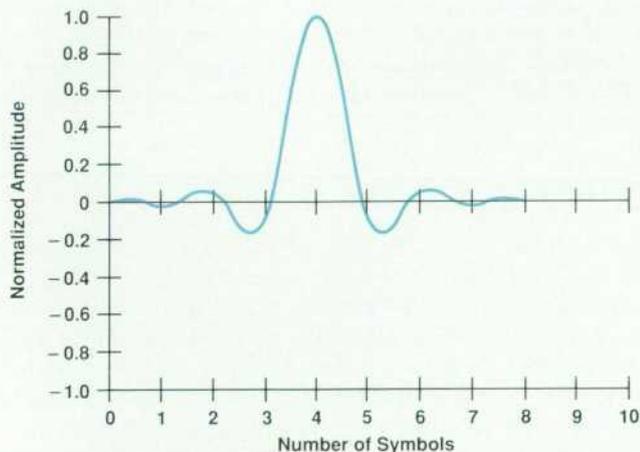
generate a smooth I-Q trajectory between symbol decision points, which is necessary to control unwanted spectral energy. For 16 subintervals per symbol, four address bits are required to be presented to the ROM address line. Adding the 16 address bits mentioned above, this implies a total ROM address space of 20 bits, or 1M bytes of address space. For 16-bit-output ROMs, this implies 16M bytes of storage for both the I and Q ROMs, for a total of 32M bytes of ROM storage. This amount of ROM was impractical, so a different approach to a ROM-based filter was implemented.

Fig. 13 shows the filter block diagram used in the HP 11846A. The key difference is that the I and Q ROMs have been replaced by cosine and sine ROMs. The I and Q outputs are now generated by an addition and subtraction of the sine and cosine ROM data. The key to being able to use this type of approach was recognizing that π/4 DQPSK can be generated from DQPSK by a 45-degree rotation every symbol. By using a coordinate transformation technique, the I and Q outputs are broken down into their cosine and sine components (see the box on page 71). The result is that the address bits required for any of the ROMs in this implementation are:

Number of Bits	Description
8	Bits from shift register, either the first or second bit per symbol.
3	Bit counter. This counts the 45-degree rotation modulo 8.
4	Subinterval counter.

The number of address bits to any given ROM is therefore 15. In this implementation, four ROMs (two ROMs for the I channel and two ROMs for the Q channel) are required, compared to the two ROMs in the conventional filter implementation. The overall result is a reduction by a factor of 16* in the memory required to implement the FIR filter. This reduction of memory not only saves cost, but significant printed circuit board area. In addition, since all of the

*2^{20} bits of address space for the conventional approach versus 2^{16} bits ($2^{15} \times 2$) of address space for the HP approach.

Fig. 10. *Impulse response of a square root raised cosine filter with an α of 0.35.*

Square Root
Raised Cosine
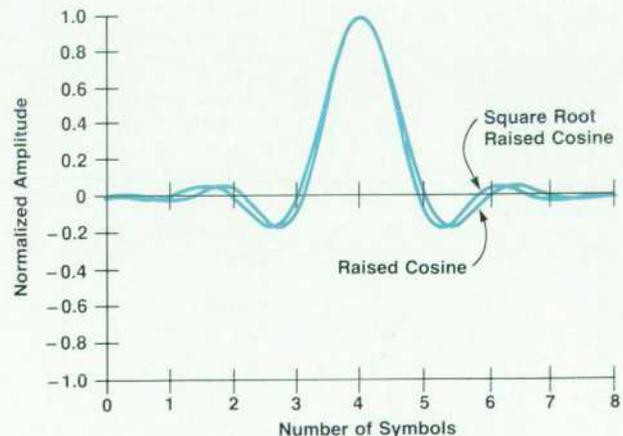
Raised Cosine

Fig. 11. *Comparison of raised cosine and square root raised cosine filter with α = 0.35.*
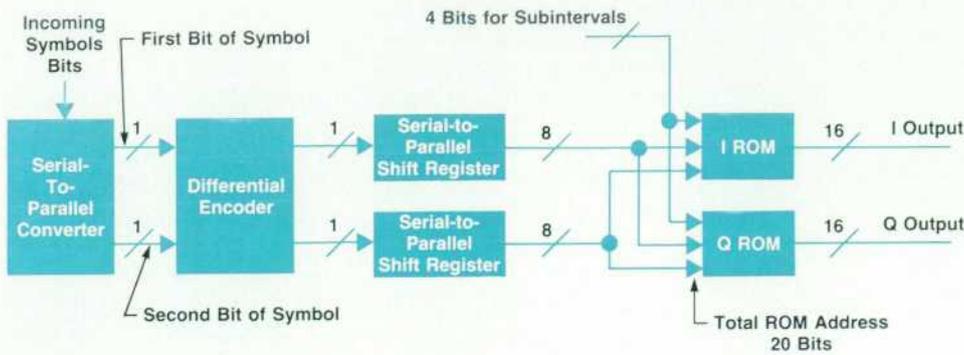
Fig. 12. *Block diagram of a conventional ROM-based FIR filter for π/4 DQPSK.*
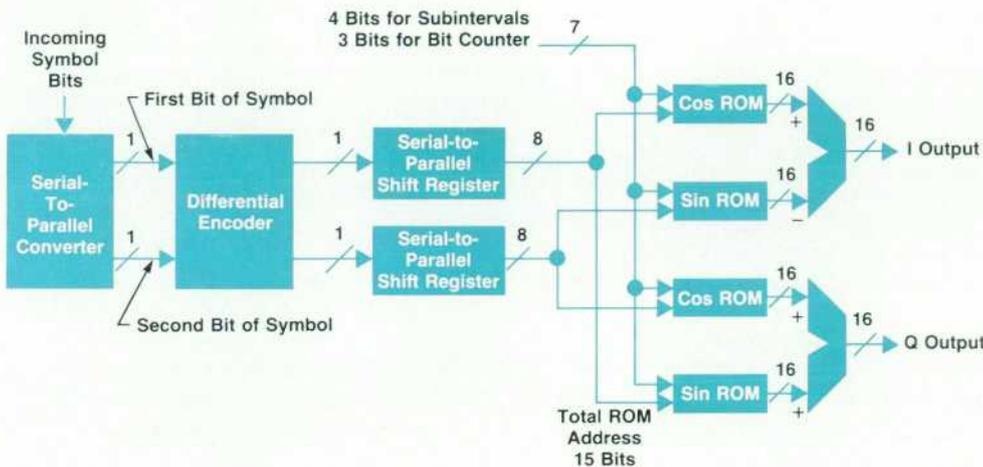
Fig. 13. *HP 11846A ROM-based FIR filter block diagram.*

filter information is stored in ROM, as new communication systems come on line, potentially with different filter characteristics, only a ROM change is required to meet the needs of these systems.

Measurement Specifications

The π/4 DQPSK modulation format required a new method for measuring modulation accuracy. In conventional analog modulation, FM for instance, the figures of merit are distortion and deviation. Therefore, if the deviation, rate, and distortion can be precisely measured, the

Fig. 14. *I-Q diagram showing the parameters used to compute the vector error.*

transmitter is accurately characterized.

In digital modulation systems, such as the NADMCS, phase accuracy at the decision points is the main figure of merit. However, since π/4 DQPSK is not a constant-amplitude system, both the amplitude accuracy and the phase accuracy are important. The TR 45.3 committee has recommended that this modulation be measured in terms of magnitude of vector error. In an I-Q diagram, the vector error is measured by plotting the desired decision point and the measured value at the decision point. The vector error is computed by measuring the vector length between the two points (see Fig. 14). In the NADMCS, the phase of the vector error is not deemed to be important, but only the magnitude of the vector. The HP 11846A has typically less than 1% vector error magnitude for its I-Q outputs. This compares with the NADMCS system modulation accuracy specification of less than 12.5%.

Conclusion

The HP 11846A provides a highly accurate and efficient implementation of π/4 DQPSK modulation. This I-Q source can be used to verify the performance of the new NADMCS radio receivers. The HP 11847A analysis software allows characterization of the NADMCS transmitters for modulation accuracy. Additional test equipment is required to characterize NADMCS radios fully, but these two products test the heart of this new system. The NADMCS system provides almost a six-to-one improvement in capacity over the current analog AMPS cellular system. Because of the increased spectral use, digital modulation formats such as

HP 11846A Filtering Technique

The implementation of the digital filtering scheme for the HP 11846A is shown in the detailed block diagram shown in Fig. 1. The input bit stream a_k for the U.S. cellular system is at 48.6 kbits/s. The $\pi/4$ DQPSK modulation has two bits/symbol, so the serial-to-parallel block separates the first and second bits. The differential encoder block performs phase rotations of the input symbol relative to the current phase state. The differential encoder block is implemented for DQPSK modulation, not $\pi/4$ DQPSK modulation. Transfer to unit circle is a scaling operation that translates the I-Q phase states to a unit circle. Before the phase rotation block the modulation is strictly DQPSK and after phase rotation it becomes $\pi/4$ DQPSK modulation. The final operation is to take the desired filter characteristic, which is a square root raised cosine filter, convolve the input data with the filter impulse response, and generate the desired I and Q outputs.

To help understand the symmetry used for this modulation format, refer to Table I on page 67 which compares the effect on a given input symbol for the two modulation formats. Related to regular DQPSK, the $\pi/4$ shift represents a coordinate shift of $\pi/4$ radians every symbol period. This rotation is described as a modulo 8 characteristic, since after eight symbols the $\pi/4$ shifts have increased the phase 360 degrees.

To generate $\pi/4$ DQPSK, the following coordinate transformation is performed on the DQPSK formatted data (C_k and D_k in Fig. 1).

$$I_k = C_k\cos(k\pi/4) - D_k\sin(k\pi/4)$$
$$Q_k = C_k\sin(k\pi/4) + D_k\cos(k\pi/4),$$

where C_k and D_k represent the I and Q components of the DQPSK

modulation, I_k and Q_k represent the I and Q components of $\pi/4$ DQPSK modulation, and k represents the kth symbol. These two equations show a rotation of $\pi/4$ radians per symbol. The argument $k\pi/4$ repeats every eight symbols, since at this point a rotation of 2π radians has occurred.

Generating ROM Lookup Data

To generate the ROM lookup data and thus the desired output data, the input data must be convolved with the impulse response of the desired filter. A standard equation for this process is:

$$Y(n) = \Sigma_k X(k)h(n-k), \quad -\infty \le k \le \infty,$$

where $X(k)$ is the discrete data input stream, $h(n-k)$ is the impulse response of the desired filter, and $Y(n)$ is the output at the time of the nth sample.

For the following discussion $g(t)$ represents the square root raised cosine impulse response, which is the $\pi/4$ DQPSK filter's impulse response. A plot of this impulse response is shown in Fig. 10 on page 69. Note that for convenience, the impulse response has been shifted in time to go from zero to the filter length. This choice prevents having any negative time representations for the filter impulse response equations. Given this choice the continuous time outputs for the I and Q channels are:

$$I(t) = \Sigma_k I_k g(t-kT)$$
$$Q(t) = \Sigma_k Q_k g(t-kT),$$

where T is the sample time and k is the kth element of I_k or Q_k. Substituting in the expressions for the coordinate transformation
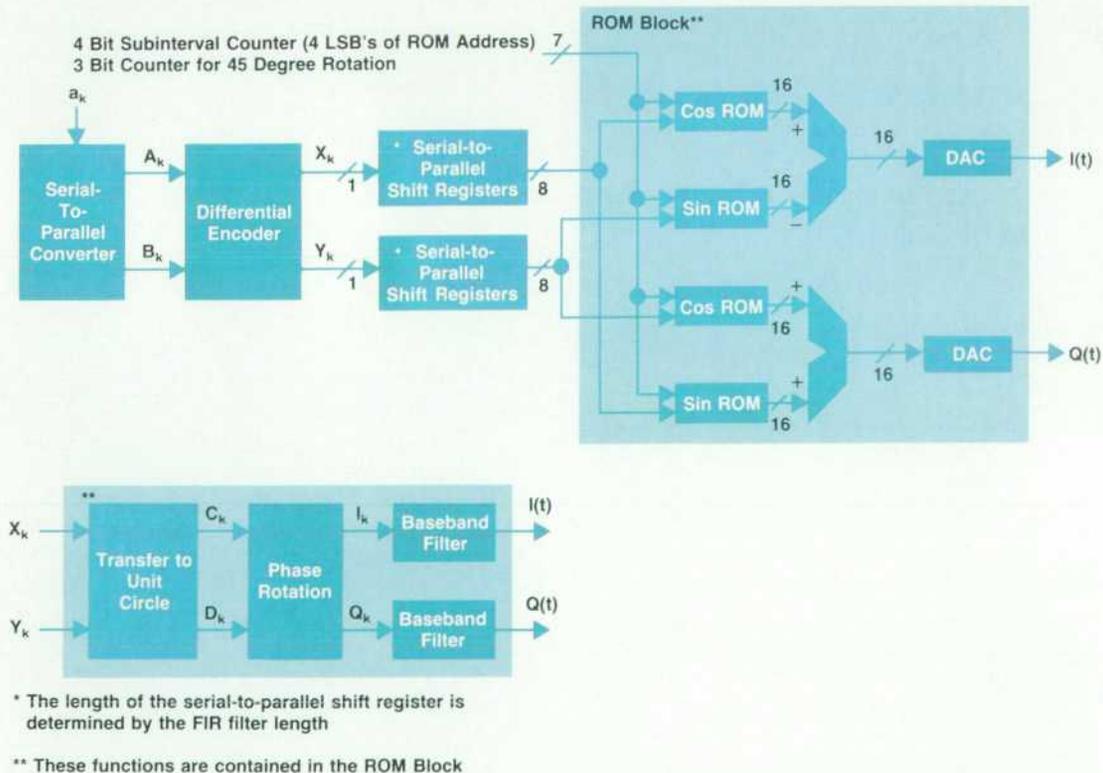
Fig. 1. *HP 11846A detailed block diagram.*

Within the figure:

4 Bit Subinterval Counter (4 LSB's of ROM Address)
3 Bit Counter for 45 Degree Rotation

ROM Block**

a_k → Serial-To-Parallel Converter → A_k → Differential Encoder → X_k → * Serial-to-Parallel Shift Registers → Cos ROM / Sin ROM → DAC → I(t)

B_k → Y_k → * Serial-to-Parallel Shift Registers → Cos ROM / Sin ROM → DAC → Q(t)

X_k → Transfer to Unit Circle → C_k → Phase Rotation → I_k → Baseband Filter → I(t)

Y_k → D_k → Q_k → Baseband Filter → Q(t)

* The length of the serial-to-parallel shift register is determined by the FIR filter length

** These functions are contained in the ROM Block

from DQPSK to $\pi/4$ DQPSK yields:

$$I(t) = \sum_k C_k \cos(k\pi/4)g(t-kT) - \sum_k D_k \sin(k\pi/4)g(t-kT)$$

$$Q(t) = \sum_k C_k \sin(k\pi/4)g(t-kT) + \sum_k D_k \cos(k\pi/4)g(t-kT)$$

where $k = (t/T - L + 1) \leq k \leq (t/T)$ and L is the length of the filter response in symbols.

When these limits of summation are referred to the impulse response then $g(t-kT)$ has limits of $g((L-1)T)$ to $g(0)$. These limits agree with Fig. 10, which has the impulse response defined from 0 to L. From these equations, we can determine the minimum information needed to compute the I and Q filter outputs. The information needed is:

- The time relative to the last data symbol clock time (i.e., the subinterval position with each data bit)
- The past L data bits, $C_n, C_{n-1}, ..., C_{n-L-1}$, and $D_n, D_{n-1}, ..., D_{n-L-1}$
- $k\pi/4$ modulo $(2n)$.

If we let $t = t_1 + nT$, $0 < t_1 < T$, the resolution of t is determined by the desired number of subintervals that should be computed. In the case of the HP 11846A, we chose to implement 16 subintervals per symbol. The benefit of computing subinterval points is that the phase transition between symbol intervals can be smoothed and the spectral energy controlled. Sixteen subintervals has proven to allow a very smooth transition between symbol phase states.

With the above definition for t_1, we can state the following for the first term of I(t).

$$I_1(t_1) = \sum C_{n-m} \cos((n-m)\pi/4)g(t_1 + mT),$$

summed over $m = 0$ to $m = L - 1$.

The information needed to compute this term includes:

- t_1, which is the time from the most recent data symbol occurring at $t = nT$. Since we are using 16 subintervals, four bits are required for the lookup ROM address.
- C_n, which is the most recent input bit. C_{n-L-1} is the input bit the farthest in the past. We are using a total of L values of C_{n-m} in this computation.
- n, which is the input bit counter. Since $\cos((n-m)\pi/4)$ has a period of eight, the input bit counter can be counted modulo 8. This implies that three address bits are required for the ROM address.

The HP 11846A uses 16 subintervals and a filter length of eight, which results in requiring 15 address bits for the lookup ROMs.

From the above discussion, the final design equations for the lookup ROMs are:

$$I(t_1) = \sum_{m=0}^{m=L-1} C_{n-m} \cos((n-m)\pi/4)g(t_1 + mT)$$

$$- \sum_{m=0}^{m=L-1} D_{n-m} \sin((n-m)\pi/4)g(t_1 + mT)$$

$$Q(t_1) = \sum_{m=0}^{m=L-1} C_{n-m} \sin((n-m)\pi/4)g(t_1 + mT)$$

$$+ \sum_{m=0}^{m=L-1} D_{n-m} \cos((n-m)\pi/4)g(t_1 + mT),$$

where $t_1 = t - nT$ and $nT < t < (n+1)T$. T is the symbol period and n is a bit counter.

The sine and cosine ROMs shown in Fig. 1 contain the summation terms for the I and Q equations described above.

the NADMCS system will be the communication systems of choice for the future.

Acknowledgments

The author would like to thank Ray Bergenheier, who developed the mathematical approach used for the HP 11846A ROM lookup data, and Al Tarbutton, who sat in on the TR 45.3 committee and assisted in the design of the HP 11846A digital filter. Thanks also to Ken Thompson and Bob Garner who contributed to writing this article.

Reference

1. *Hewlett-Packard Journal*, Vol. 38, no. 11, December 1988, pp. 4-52.

Measuring the Modulation Accuracy of π/4 DQPSK Signals for Digital Cellular Transmitters

Using digital signal processing techniques, this software accurately verifies the RF performance of digital cellular transmitters conforming to the North American Dual-Mode Cellular System standard.

by Raymond A. Birgenheier

THE HP 11847A π/4 DQPSK modulation measurement software performs accurate verification of the RF performance of cellular transmitters conforming to the standard for the North American Dual-Mode Cellular System (NADMCS). The standard was prepared by the TR 45.3 subcommittee of the Electronic Industries Association and Telecommunications Industries Association (EIA/TIA). Included in the standard are specifications for the maximum carrier frequency error and the modulation accuracy of the transmitter. The modulation accuracy is characterized by the rms error vector magnitude averaged over the detection decision points of a burst. The error vector magnitude is the phasor difference as a function of time between an ideal reference signal and the measured transmitter signal after it has been compensated in timing, amplitude, frequency, phase, and dc offset (see Fig. 1). The rms error over a single burst and the first ten symbols of ten consecutive bursts emitted within a one-minute interval are also specified in the standard.

The HP 11847A software uses digital signal processing techniques to measure carrier frequency error, modulation phase and amplitude error, and error vector magnitude. In addition, it generates valuable graphs of the demodulated signal. These graphs include I-Q diagrams, phase and amplitude graphs, phase- and amplitude-error graphs, and the error vector magnitude graph. Other supplied measurements include measurements of TDMA (time-division multiple-access) amplitude characteristics and FFT measurements of the error vector data.

This article provides an overview of the HP 11847A measurement system and the signal processing techniques used to demodulate RF signals, recover the data, and measure the modulation accuracy of the signal under test.

Measurement Method

The method used in the HP 11847A measurement system to measure the frequency error and error vector magnitude with a measurement error very small compared to the error specification of the transmitter consists of the following steps.

1. An internal reference signal is established by detecting data from the transmitter signal under test.

2. The parameters of the transmitter signal are estimated and used to compensate the transmitter signal.

3. The in-phase (I) and quadrature (Q) components and the magnitude and phase are calculated as functions of time for both the reference signal and the transmitter signal under test.

4. The in-phase and quadrature components of the reference signal and the signal under test are compared to determine the error vector magnitude as a function of time.

This method also compares phase and magnitude functions of the reference signal and the signal under test to determine phase and magnitude errors as a function of time. From the error functions, the rms values of the magnitude error, phase error, and error vector magnitude are calculated.

To generate an estimate of the ideal reference signal, the data sequence over the signal burst must be detected, and to compensate the transmitter signal, the signal parameters of the transmitter must be estimated. The parameters required are the transmitted signal's clock delay, carrier frequency, carrier phase, amplitude scale factor, amplitude droop factor, and I-Q origin offset.* These six parameters

*The amplitude droop factor is the rate of decay in the magnitude of the signal burst in units of dB per symbol. I-Q origin offset is a dc offset caused by carrier feedthrough.
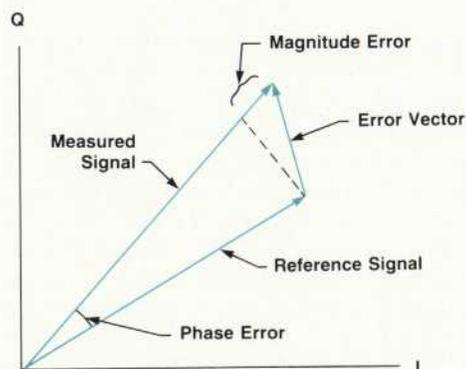

Fig. 1. *A representation of the error vector, which is the phasor difference as a function of time between the measured transmitter signal and an ideal reference signal.*
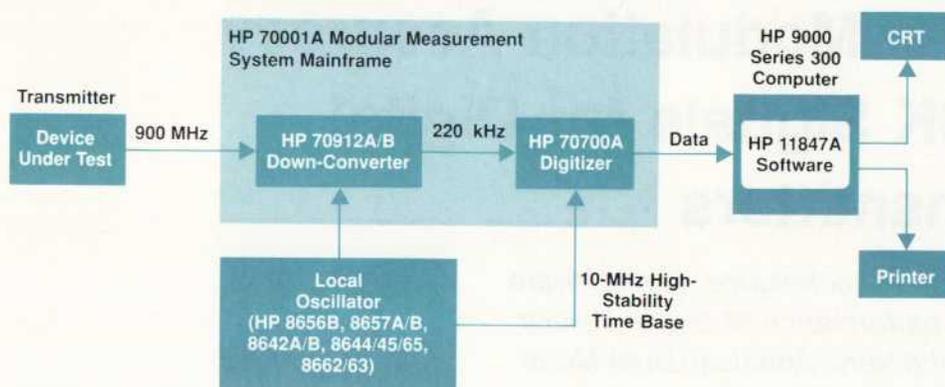
Fig. 2. The modulation measurement system hardware block diagram.

are estimated by finding values that minimize the rms error vector magnitude averaged over the decision points of a burst.

System Overview

The block diagram of the hardware portion of the HP 11847A measurement system is shown in Fig. 2. Shown is the HP 70912A/B down-converter fed by a local oscillator such as an HP 8656B synthesized signal generator, which down-converts the transmitter signal under test to a 220-kHz IF signal. This IF signal is digitized by the HP 70700A digitizer and fed to an HP 9000 Series 300 computer in which the signal processing algorithms described in this paper are executed. The output of the signal processor can be displayed on the CRT monitor of the computer or output to a printer.

The signal processing system flow diagram is shown in Fig. 3. The input to the flow diagram is a discrete-time IF signal derived from the transmitter signal under test. The transmitter signal is down-converted and digitized to produce the discrete-time signal fed through the IF bandpass filter (① in Fig. 3). The purpose of the IF bandpass filter is to eliminate harmonics of the IF signal that may be produced when the transmitter signal is down-converted, and to reduce quantization noise created by the digitization process. A significant reduction in quantization noise is realized by initially sampling the IF signal at a high rate ($\approx$ 41.15 times the symbol rate specified for U.S. digital cellular radios). The sampling rate then is reduced by a factor of four from the input to the output of the IF bandpass filter.

The clock delay, which is the time interval from the first sample from the IF bandpass filter to the leading edge of the first full symbol interval following this first sample, is estimated in the clock-delay estimator ②. In the clock-delay estimator, the modulated IF signal from the output of the IF bandpass filter is squared to produce a signal component at a frequency equal to the symbol clock rate. This squared IF signal is then fed through an FIR bandpass filter that is tuned to the symbol rate. The phase of the signal at the output of this filter provides a very accurate estimate of the symbol clock delay.

After the clock delay, τ, has been estimated, the IF signal is resampled by a Nyquist interpolation filter ③. The purpose of resampling is to provide an integral number of samples per symbol interval and to provide samples pre-

cisely at the detection decision points. To minimize the amount of arithmetic necessary following the resampler, and to minimize signal processing errors that contribute to measurement inaccuracy, a rate of five samples per symbol was chosen for the resampler.

The output of the resampler provides the basic discrete-time IF signal used by the remainder of the signal processor. The uncertainty of the symbol clock phase has been removed so that these samples occur at a rate of five samples per symbol with every fifth sample accurately located at a detection decision point in time.

Three Passes

Following the resampler, the signal processing functions are performed in three passes. To understand the need for multiple passes, consider the signal originating at the transmitter. This signal passes through a square root raised cosine filter in the transmitter and, for a matched receiver, it passes through another square root raised cosine filter in the receiver. The cascading of the two square root raised cosine filters produces the Nyquist frequency response necessary to avoid intersymbol interference. To perform square root raised cosine filtering, the receiver must know the precise frequency of the transmitter. In the signal processor described here, the data is detected and the parameters of the transmitter signal are estimated in the first pass during which noncoherent processing of the IF signal at the output of the resampler is performed. The initial estimates obtained in the first pass are used in the second pass to perform coherent baseband processing.

In the second pass, the IF signal is coherently down-converted to in-phase and quadrature baseband signals which are passed through square root raised cosine filters. The square root raised cosine filtered signals are used again to detect data and obtain refined estimates of the signal parameters. The detected data from the two passes is compared for consistency. A difference implies bit detection errors. The detected bits are used to generate an ideal reference signal, and the refined estimates of signal parameters are used to correct the I-Q origin offset, carrier frequency and phase, and amplitude scale and droop factors of the signal under test. The compensated signal under test and the reference signals are compared to obtain measurements of rms magnitude and phase errors and rms error vector magnitude.

During pass 2 only one sampled value per symbol at the
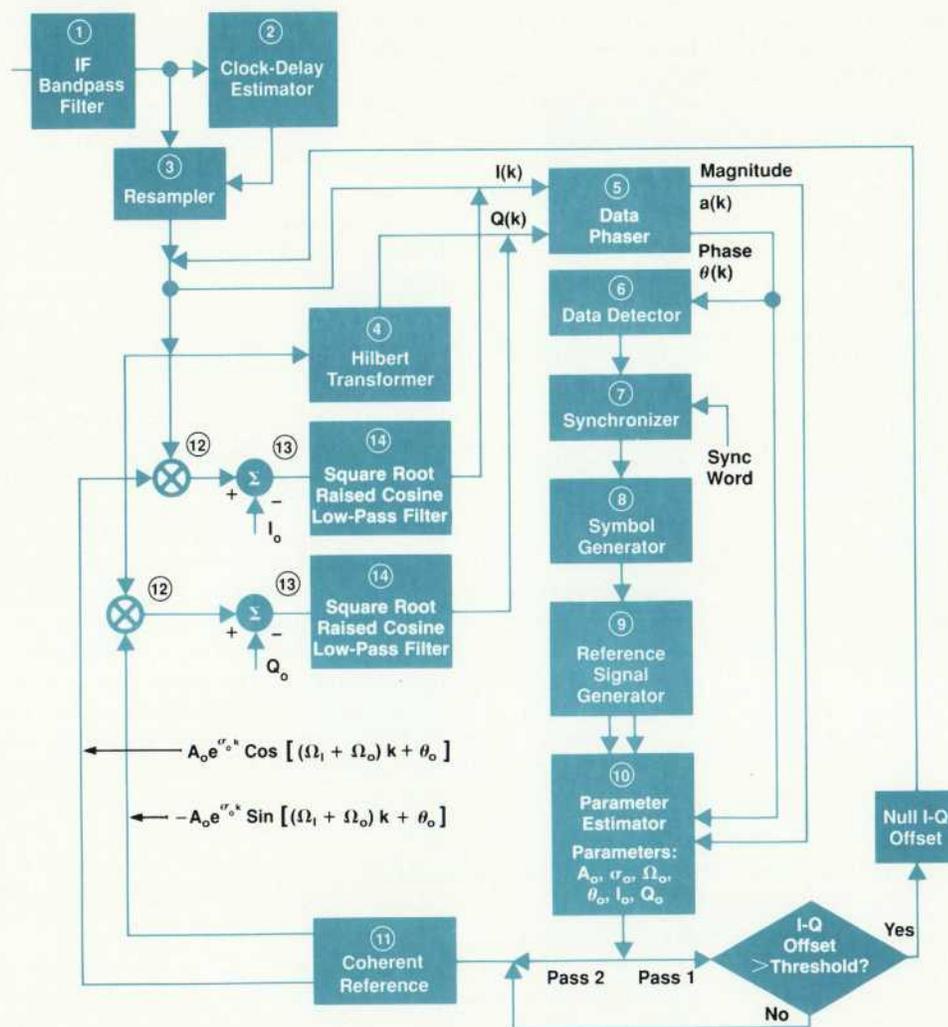
Fig. 3. *The HP 11847A signal processing flow diagram.*

decision points is used. This provides measurements of carrier frequency error and error vector magnitude in a minimum amount of time. Detailed plots of various signals as a function of time over the duration of a burst are calculated in the third pass using up to 20 points per symbol. In addition, Fourier transforms of various signals can be computed in the third pass.

Pass 1. The first pass begins after the signal under test has been resampled to produce an IF signal with five samples per symbol and samples located at the detection decision points. In the discussion that follows, the discrete time variable, k, is the sample number after the clock delay. Relative to the time instant at which the first time sample in taken by the digitizer, time is given in units of symbol intervals as:

$$\text{time} = \text{clock delay} + 0.2k.$$

After resampling is complete, the signal is fed through a Hilbert transformer ④ to provide the quadrature components of the signal. From the in-phase component, $I(k)$, and the quadrature component, $Q(k)$, the magnitude

$$a(k) = \sqrt{I^2(k) + Q^2(k)}$$

and the phase

$$\theta(k) = \tan\left(\frac{Q(k)}{I(k)}\right) - \Omega_I k,$$

where Ω_I is the nominal IF frequency, are computed in the data phaser ⑤. The difference in phase $\theta(k) - \theta(k-5)$ is used in the data detector ⑥ to detect the data symbols as shown in Table I.

Table I
Phase Difference and Detected Symbols

$\theta(k) - \theta(k-5)$			Detected symbols	
0°	to	90°	0	0
90°	to	180°	0	1
−90°	to	0°	1	0
−180°	to	−90°	1	1

After the data bits are detected, they are correlated with a unique 28-bit synchronization word in the synchronizer ⑦ to establish the interval of the burst over which the measurements are to be made. The synchronizer establishes the beginning of the sequence of detected bits to be used

in calculating the reference signals.

After the timing reference has been established and the data bits are detected, symbols are generated in the symbol generator ⑧. The symbol generator performs the serial-to-parallel conversion and differential encoding and mapping to a circle necessary to convert a serial data stream to symbols that drive a quadrature modulator to generate the $\pi/4$ DQPSK modulated reference signal (see the article on page 65 for a description of how a $\pi/4$ DQPSK signal is generated). The symbol generator provides the input to the reference signal generator ⑨ in which the square root raised cosine reference amplitude $a_{ref}(k)$ and reference phase $\theta_{ref}(k)$ are calculated. $a_{ref}(k)$ and $\theta_{ref}(k)$ and the amplitude $a(k)$ and the phase $\theta(k)$ of the signal under test are the inputs to the parameter estimator ⑩, which calculates estimates of the transmitter amplitude scale factor A_o, phase θ_o, frequency error Ω_o, I-component offset I_o, Q-component offset Q_o, and amplitude droop factor σ_o.

The parameter estimator determines the optimal values of A_o, θ_o, Ω_o, I_o, Q_o, and σ_o to minimize the mean-square error vector magnitude, which is

$$\overline{\epsilon^2} = \sum_k |A_o a(k)e^{\sigma_o k - j[\Omega_o k + \theta_o - \theta(k)]} - B_o - a_{ref}(k)e^{j\theta_{ref}(k)}|^2 \tag{1}$$

where

$$B_o = I_o + jQ_o. \tag{2}$$

The optimal values of the parameters can be found to a very good approximation by minimizing:

$$\overline{\epsilon^2} = \sum_k a(k)a_{ref}(k)e^{\sigma_o k}\{[\ln A_o + \sigma_o k + \ln a(k) - \ln a_{ref}(k)]^2$$
$$+ [\theta(k) - \theta_{ref}(k) - \Omega_o k - \theta_o]^2\} \tag{3}$$

with respect to A_o, σ_o, θ_o, and Ω_o. After the parameter values that minimize are found, I_o and Q_o are calculated using the equations:

$$I_o = \frac{1}{N}\sum_k \{A_o a(k)e^{\sigma_o k}\cos[\Omega_o k + \theta_o - \theta(k)] - a_{ref}(k)\cos\theta_{ref}(k)\} \tag{4}$$

and

$$Q_o = \frac{1}{N}\sum_k \{(A_o a(k)e^{\sigma_o k}\sin[\Omega_o k + \theta_o - \theta(k)] - a_{ref}(k)\sin\theta_{ref}(k)\} \tag{5}$$

where N is the number of decision points.

In the event that the I-Q origin offset is large, an offset as given by equations 4 and 5 is removed from the output of the resampler. If this is required, then pass 1 is repeated with the process returning to the input of the Hilbert transformer, repeating the calculations in blocks five through ten of Fig. 2. For very large I-Q origin offsets it may be necessary to repeat pass 1 several times to ensure reliable estimates of the optimal parameter values by minimizing equation 3.

```
HP 11847A MEASUREMENT RESULTS: Single Burst Error Vector Magnitude Test
COMMENTS:
--------------------------------------------------------------------------
Captured peak signal level at digitizer input was approximately -4.94 dBm
Clock delay =.066 symbol intervals from first time sample

DEMODULATED DATA IS:
--------------------------------------------------------------------------

1          11         21         31         41         51

0000000000 0000000001 0001111011 1101100110 0000000000 0000000000

61         71         81         91         101        111

0000000010 1011001100 1010001110 1010100100 0101100111 0010110100

121        131        141        151        161        171

1110101000 0101011110 1000000110 1000100001 1100100001 1111000011

181        191        201        211        221        231

0100011100 1000010100 0000101101 0101000110 0011101111 1100111000

241        251        261        271        281        291

1010111000 0100111011 0011100000 0001010000 1010010100 0010111011

301        311        321

0111011110 1100

Burst Amplitude Droop = -.0001 dB per Symbol
Carrier Frequency Error = -4.794 Hz
I/Q Origin Offset (carrier feedthrough) = -39.76 dB

TEST RESULTS
--------------------------------------------------------------------------
Magnitude error = .792 percent RMS at decision points
Phase error = .444 degrees RMS at decision points
Error Vector magnitude = 1.108 percent RMS at decision points
```

Fig. 4. *A printout of HP 11847A measurement results.*
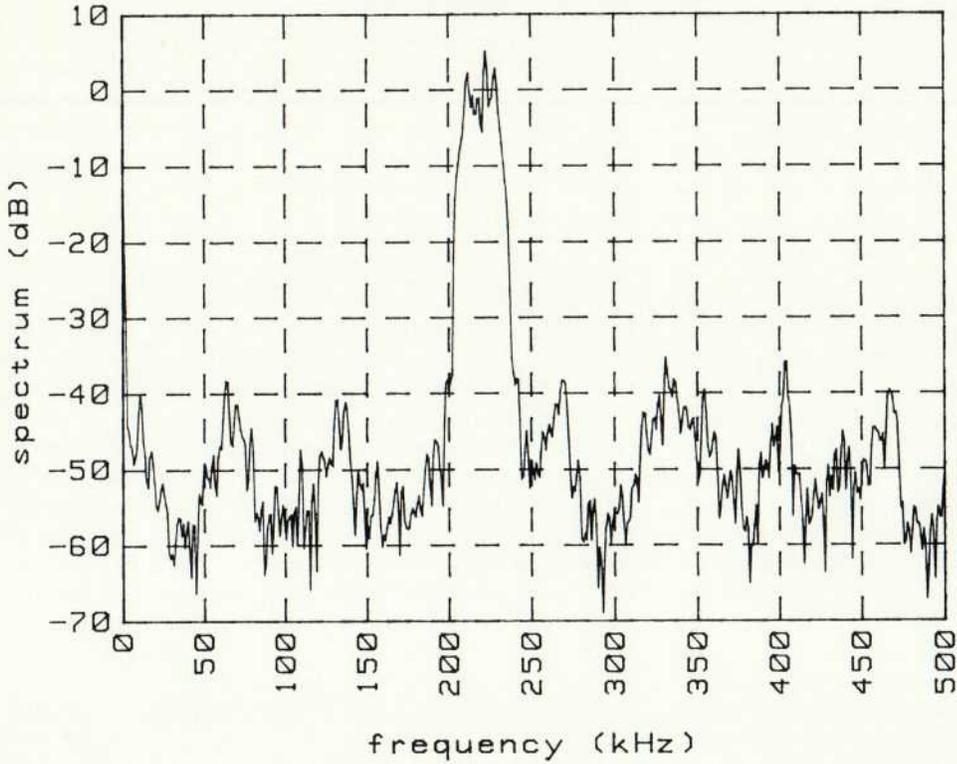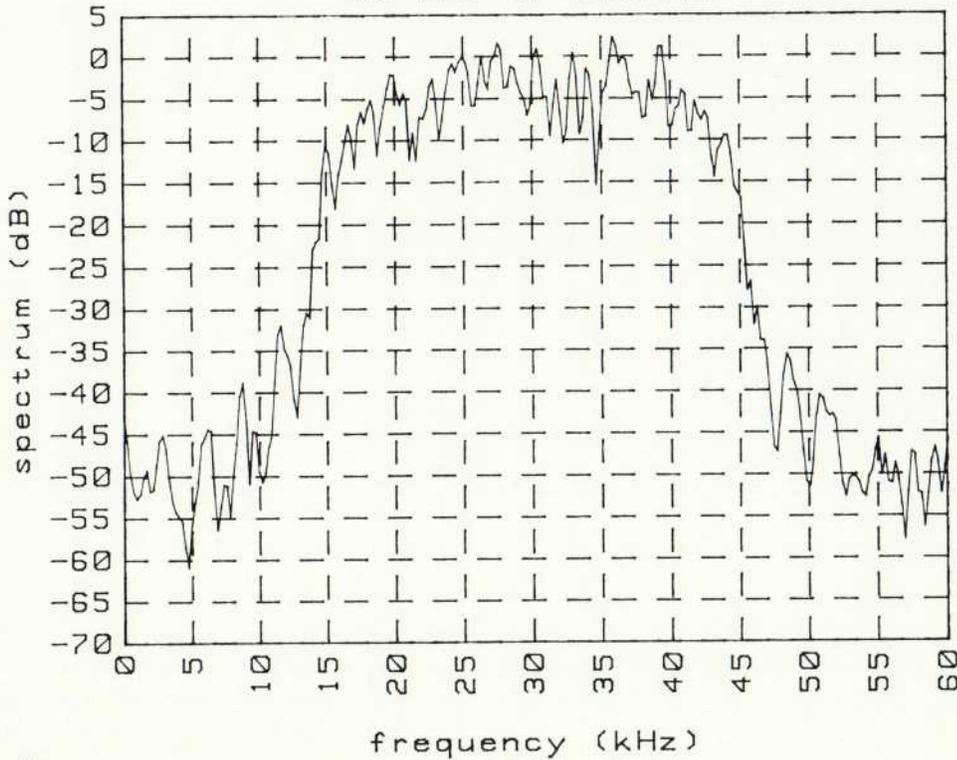
Fig. 5. *HP 11847A plots of (a) a spectrum for a 220-kHz IF signal and (b) a spectrum for a 30-kHz IF signal.*

Pass 2. After optimal signal parameters have been estimated in pass 1, the signal processor proceeds to pass 2. The first step in pass 2 is the baseband conversion of the IF signal to be measured at the output of the resampler. Baseband conversion is performed in blocks eleven through fourteen of Fig. 2 in which coherent local oscillator signals ⑪ given as:

$$A_o e^{\sigma_o k} \cos[\Omega_I + \Omega_o)k + \theta_o] \qquad (6)$$

and

$$-A_o e^{\sigma_o k} \sin[\Omega_I + \Omega_o)k + \theta_o] \qquad (7)$$

are multiplied ⑫ with the output signal from the resampler. I_o and Q_o are then subtracted ⑬ from these products to compensate for the I-Q origin offset. The resulting signals are passed through square root raised cosine filters ⑭ to produce in-phase and quadrature components of the transmitter signal under test. The outputs of the filters are the baseband signal components representing the transmitter signal after it is coherently down-converted to baseband, scaled in amplitude by A_o, compensated for amplitude droop by the factor $e^{\sigma_o k}$, corrected for I-Q origin offset (residual I-Q origin offset remaining after pass 1) by subtracting I_o and Q_o, and then passed through the square root raised cosine filters. At this point, $I(k)$ and $Q(k)$ are the baseband components of the compensated transmitter signal with full Nyquist filtering. The full Nyquist filtering is a result of cascading a square root raised cosine filter in the transmitter with the square root raised cosine filter ⑭ in the receiver as mentioned before.

The next step in pass 2 is the calculation of the magnitude and phase of the baseband signal at the detection decision points in the data phaser. The binary data is again detected from the phase difference $\theta(k) - \theta(k-5)$ for the baseband signal according to Table I, and then compared to the binary data detected during pass 1. If there are no differences in the detected bits from the first and second passes, then the next step is the generation of Nyquist reference signals in the reference signal generator ⑨. In the event there are differences in the detected bits, synchronization and symbol generation are repeated in blocks ⑦ and ⑧ before the Nyquist reference signals are generated.

During pass 2 the reference signal generator calculates the amplitude $a_{ref}(k)$, phase $\theta_{ref}(k)$, in-phase component $I_{ref}(k)$, and quadrature component $Q_{ref}(k)$ of the Nyquist reference signal. $a_{ref}(k)$, $\theta_{ref}(k)$ and the magnitude and phase, $a(k)$ and $\theta(k)$, of the Nyquist baseband signal under test provide the input to the parameter estimator ⑩ to calculate refined estimates of the transmitter signal parameters as given in equations 1 to 5. The refined estimates of the signal parameters are used in the baseband converter (blocks ⑪ to ⑭) to recalculate the in-phase and quadrature components, $I(k)$ and $Q(k)$, of the baseband signal. The signal components obtained from pass 2 are then used to calculate the error vector magnitude as a function of time using:

$$\epsilon(k) = \sqrt{[I(k) - I_{ref}(k)]^2 + [Q(k) - Q_{ref}(k)]^2} \qquad (8)$$

and the rms value of the error vector magnitude averaged over the decision points is computed using:
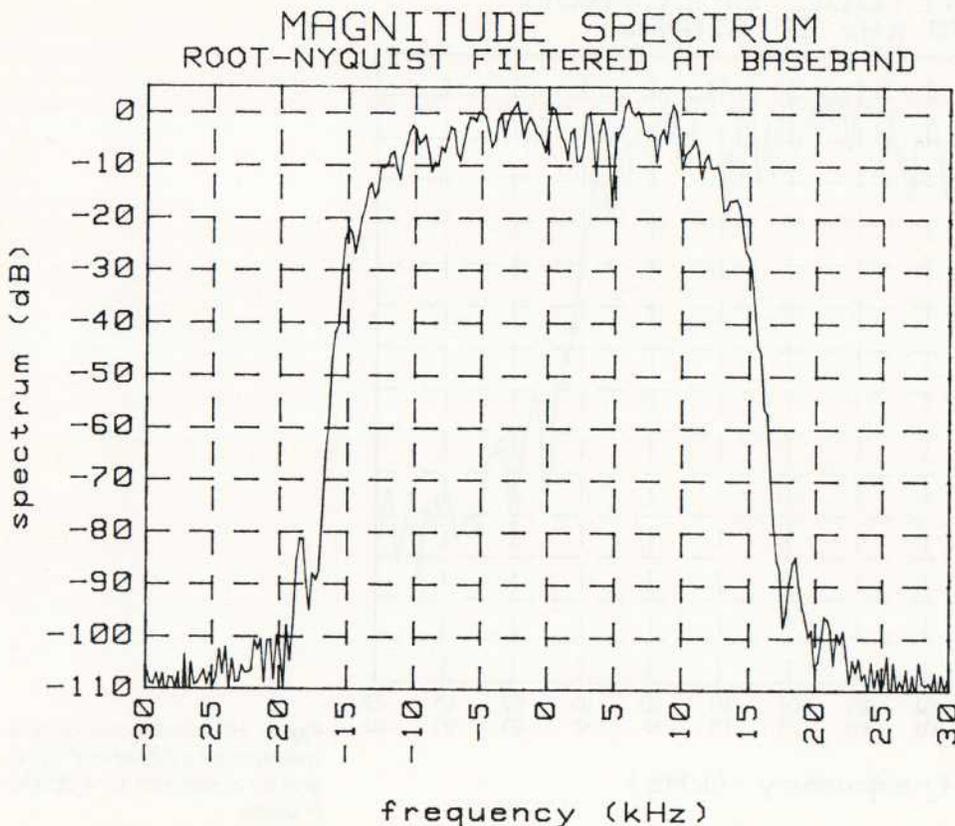
Fig. 6. A typical HP 11847A plot of the spectrum of a signal that has been down-converted and passed through a square root raised cosine filter.

$$\epsilon_{rms} = \sqrt{(1/N)\sum_k \epsilon^2(k)} \qquad (9)$$

The magnitude and phase errors and the rms values of these errors are computed using equations similar to equations 8 and 9.

Pass 3. The signal processing may conclude with the completion of pass 2 when the rms errors are the only desired outputs. When waveforms as a function of time or magnitude spectra of waveforms are desired, the processor proceeds to pass 3. In pass 3 the operator selects the number of points per symbol (1, 2, 4, 5, 10, or 20 points) at which the waveforms will be calculated. The processor then proceeds to compute signal components for the baseband signal under test and the reference signal for the number of points selected. From these signal components it is possible to obtain plots of eye diagrams, I-Q diagrams, magnitude, magnitude error and droop in dB, phase, phase and frequency error, magnitude error, and error vector magnitude. In addition the processor has an FFT routine that allows the computation and display of the spectra of the IF signals, the square root raised cosine filtered baseband signal, and the vector error.

Typical Output Plots

Figs. 4 to 9 show typical outputs from the HP 11847A measurement system. Fig. 4 presents a summary of the measurement results for a single burst of 157 symbols. Listed are the peak signal level at the digitizer input and the clock delay relative to the first time sample. The other measurement results summarized in Fig. 4 are the detected data bits, burst amplitude droop rate, carrier frequency error, I-Q origin offset and rms values of the magnitude error, phase error, and error vector magnitude averaged over the detection decision points of the burst.

Fig. 5 presents typical spectra of the IF signals. Fig. 5a shows the spectrum of the 220-kHz IF signal which is the signal fed from the digitizer to the signal processor shown in Fig. 3. This signal is sampled at a rate of 1 megasample per second so the entire spectrum up to the folding frequency is presented. The 220-kHz IF signal is fed through a bandpass filter (① in Fig. 3) and resampled at a rate of 250 kilosamples per second to produce a 30-kHz IF signal. The 30-kHz signal is resampled again by the resampler (③ in Fig. 3) at a sampling rate of five samples per symbol, which is a rate of 121.5 kilosamples per second for U.S. digital cellular radios. A typical plot of the spectrum of a 30-kHz IF signal is presented in Fig. 5.

After a coherent local oscillator has been established by
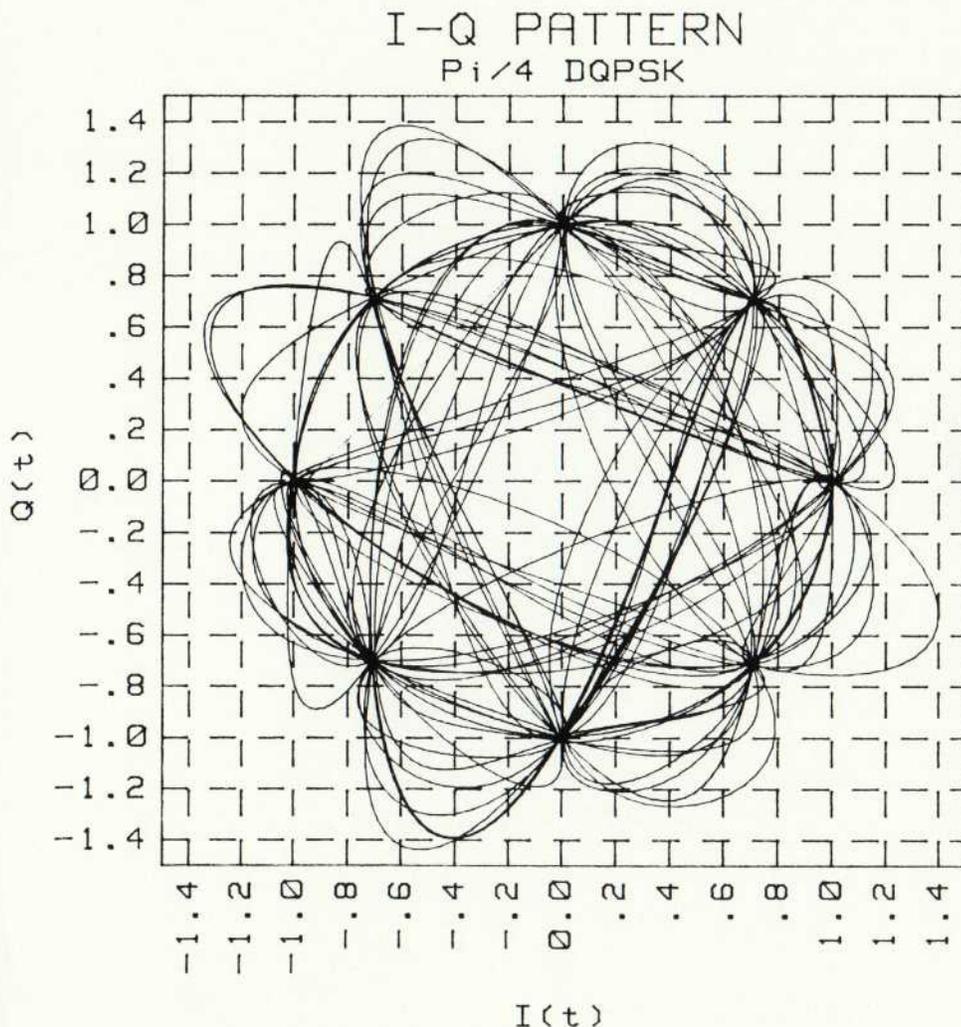
Fig. 7. *A typical π/4 DQPSK I-Q diagram.*

the signal processor, the 30-kHz IF signal is down-converted to baseband and passed through square root raised cosine low-pass filters in the baseband converter (⑪ through ⑭ in Fig. 3). A typical plot of the spectrum of the square root raised cosine filtered baseband signal is presented in Fig. 6. The tight tolerances on the square root raised cosine filter are evident from the low sidelobes seen by comparing Figure 6 to Fig. 5b.

Fig. 7 shows a typical plot of the quadrature component, Q(t), of the transmitter signal plotted against the in-phase component, I(t). For the π/4 DQPSK signal there are eight constellation points equally spaced 45 degrees apart. Ideally, the transmitter signal should pass precisely through the constellation points. If it does not, then this is an indication of intersymbol interference introduced by modulation errors of the transmitter.

Another indication of transmitter errors that would give rise to intersymbol interference can be seen from the eye diagrams. Fig. 8 shows a typical I-channel eye diagram. At the detection decision points, which are integer values on the time axis of Fig. 8, the I-channel signal should have precisely one of the five values seen on the diagram. Any deviation from this condition is indicative of intersymbol interference introduced by modulation errors of the transmitter.

Figs. 9a to 9d show magnitude error, phase error, and error vector magnitude as a function of time. The vertical bars which are most obvious on the error vector magnitude plot in Fig. 9c show the errors at the detection decision points. It is only the errors at these decision points that are used in the calculation of the rms errors presented in Fig. 4. A feature of the software allows the user to zoom in on any of the time plots to examine a particular anomaly more carefully, as shown in Fig. 9d.

Acknowledgments
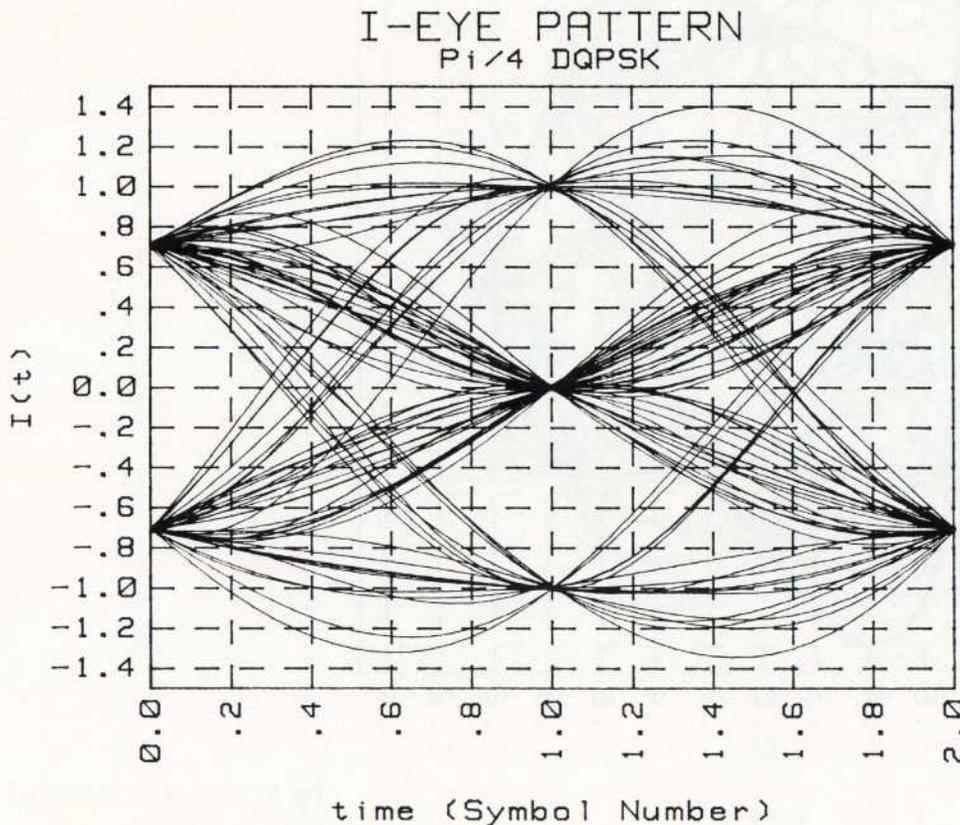
Fig. 8. *A typical I-channel eye diagram.*

MAGNITUDE ERROR
Pi/4 DQPSK

error (percent)

time (Symbol Number)

(a)

PHASE ERROR
Pi/4 DQPSK

error (degrees)

time (Symbol Number)

(b)
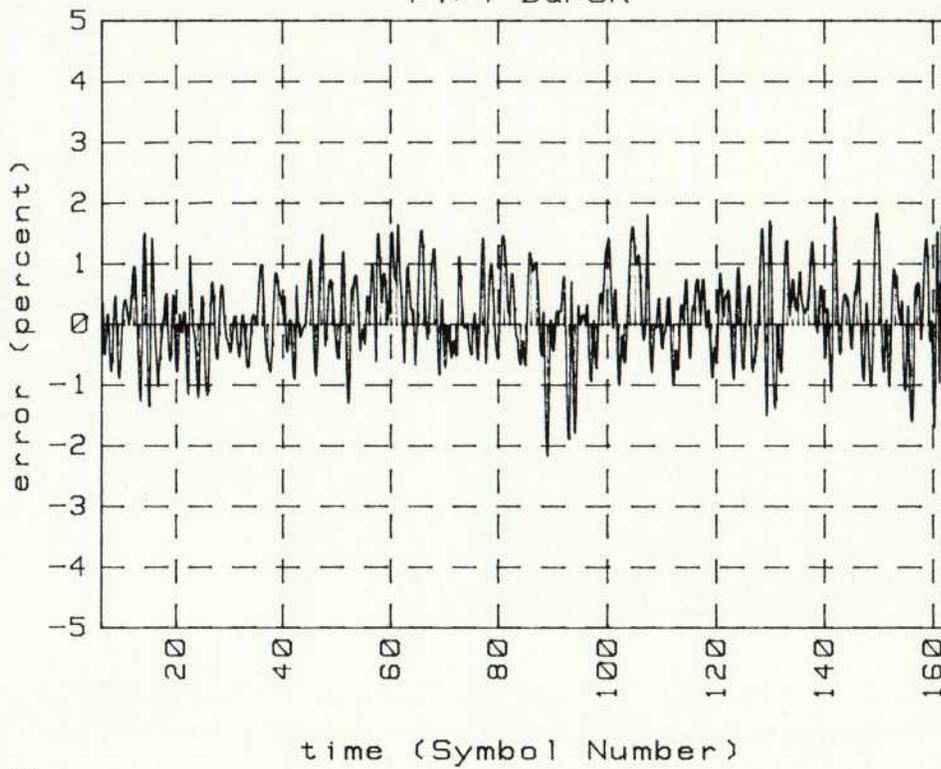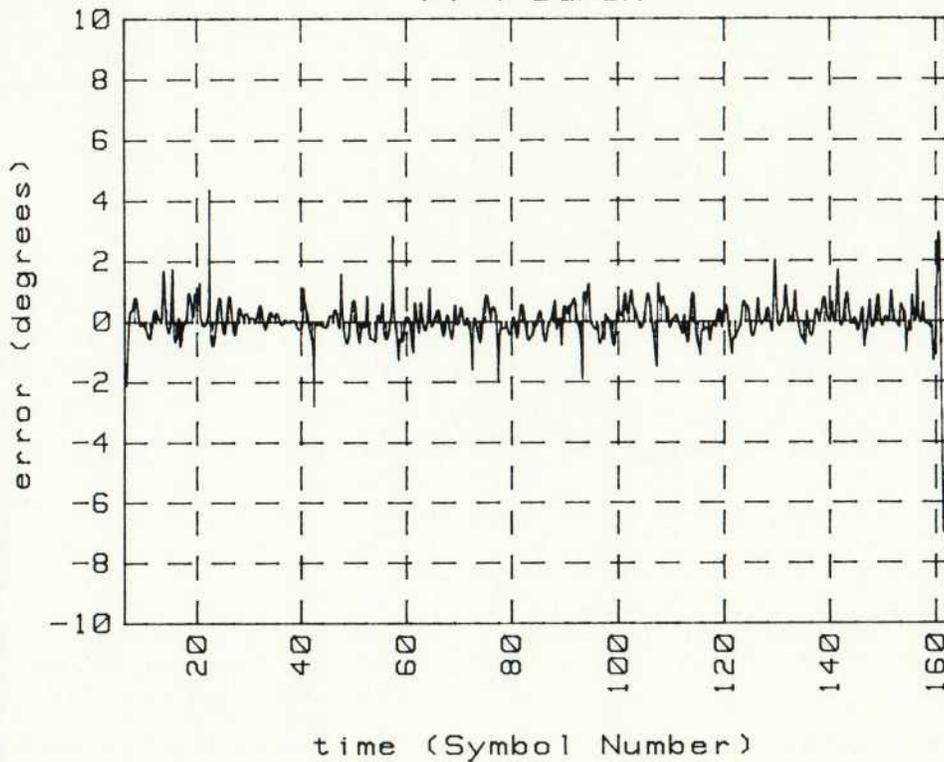
Fig. 9. *HP 11847A plots (a) magnitude error and (b) phase error. (Fig. 9 continued on next page.)*

ERROR VECTOR MAGNITUDE
Pi/4 DQPSK

(c)
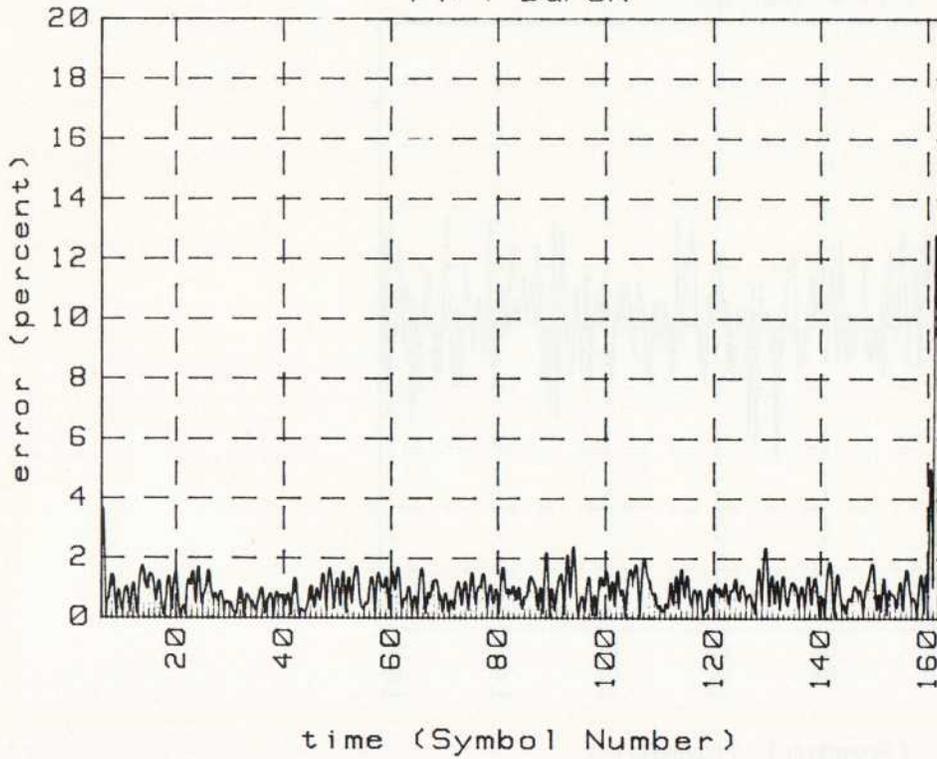
ERROR VECTOR MAGNITUDE
Pi/4 DQPSK
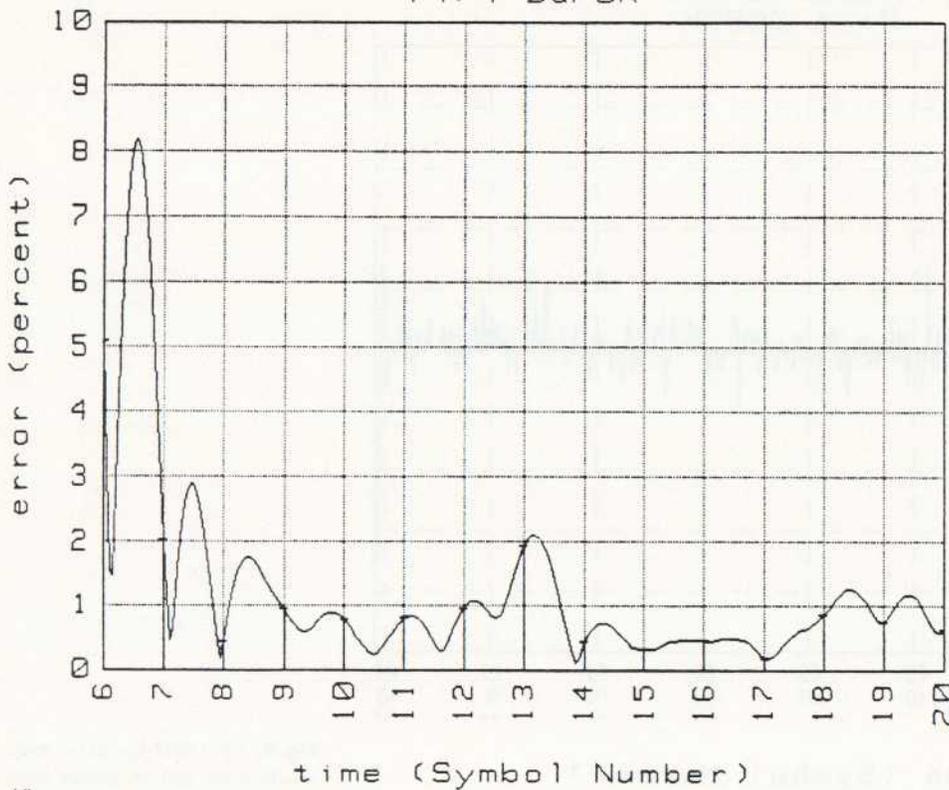
(d)

Fig. 9 (continued). *HP 11847A plots of (c) error vector magnitude and (d) an expanded view of the first 20 symbols of the error vector magnitude shown in (c).*

A Test Verification Tool for C and C++ Programs

The HP Branch Validator provides an automated tool that enables software developers to test and verify the branch coverage of their modules as they are created.

by David L. Neuder

SOFTWARE PRODUCT TESTING takes different forms throughout the software development cycle. Fig. 1 shows the opportunities for testing. The project schedule, the type of software product, and other factors determine whether testing is performed at every phase and how intense it is. Executable modules are usually tested, either individually (module-based testing) or collectively (integration testing), regardless of the type of software product. Module-based tests are written at the same time the code is written and as such they tend to be more complete because of the developer's immediate familiarity with the code. Properly constructed module tests can be very effective in reducing undetected defects and side effects resulting from code changes. The problem with module-based testing is that without some automated assistance, it is time-consuming for developers to write the code to test every module, and it is difficult to measure the effectiveness of such testing.[1]

The HP Branch Validator is a software execution and test verification tool that uses branch analysis to assist and increase the effectiveness of module-based testing. It instruments C programs and validates, in a Boolean manner, which branches of a software application have been executed. Branches are defined as the paths after the decision points in the code. Each of the various paths after a decision point becomes a branch of the application. The HP Branch Validator determines if these branches have executed at least once. Fig. 2 illustrates branches after decision points of program flow.

The HP Branch Validator makes use of the Encapsulator[1,2] to provide a friendly window interface which speeds up the process of analyzing branch coverage results. This interface allows the HP Branch Validator to be run as a stand-alone product or as an integrated member of the Softbench environment.[2] In addition, the HP Branch Validator can be used in a command-line interface to produce branch coverage reports similar to the SoftBench reports. This second interface is valuable for the customer who desires to run the HP Branch Validator on a terminal.

The native and embedded versions of the HP Branch Validator are hosted on HP 9000 Series 300, 400, and 800 computer systems. The embedded version supports AxLS C language for embedded microprocessor development applications. The native version supports both HP-UX C and C++.

The HP Branch Validator is designed to work on C pro-grams. For it to work with C++ programs, the C++ files must be compiled to C files and then the HP Branch Validator can be run on the C files. Choosing the appropriate options with the C++ compiler allows this compile process from C++ file to object files to occur in one step. The results obtained will reflect back to the original C++ source files.

The HP Branch Validator consists of three components: The HP Branch Validator preprocessor, the HP Branch Val-
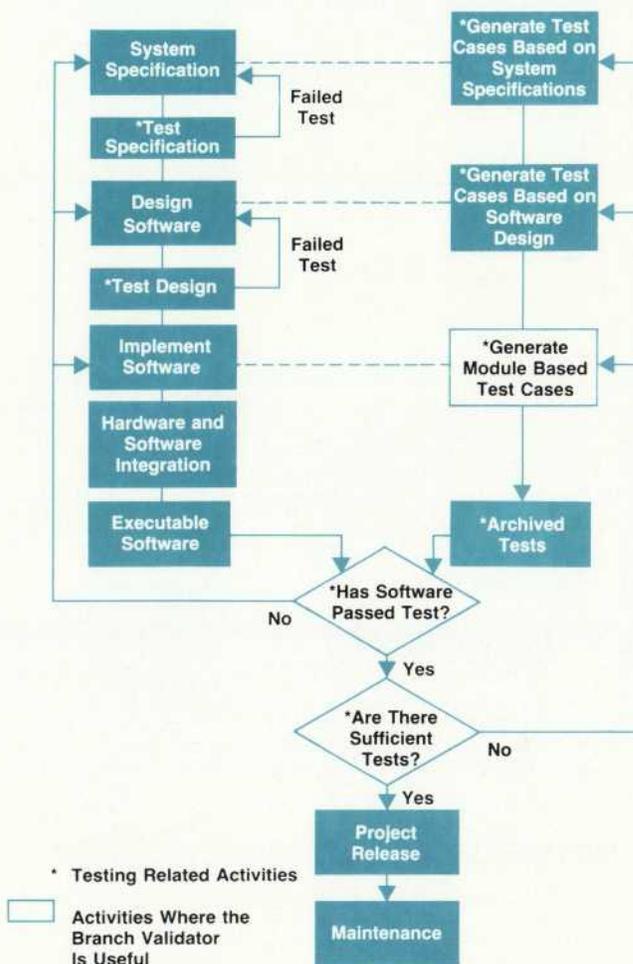
Fig. 1. *Opportunities for testing during the software development process.*
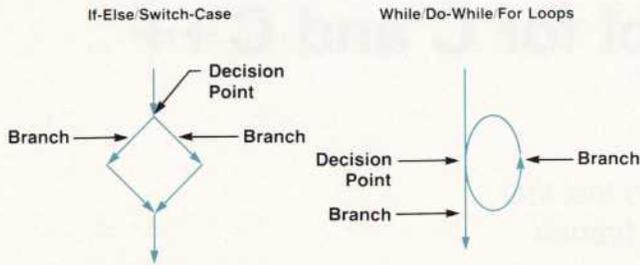
Fig. 2. *Decision points and branches in program flow.*

idator report generator, and the HP Branch Validator SoftBench interface. Fig. 3 illustrates the HP Branch Validator SoftBench interface.

The Preprocessor

The HP Branch Validator preprocessor replaces the C preprocessor (ANSI or non-ANSI) to instrument C or C++ code for branch analysis. The preprocessor inserts array definitions and statements into the source file to record which branches are executed in the file. In addition, the preprocessor generates a map file that contains information necessary to map a specific branch to an associated line in the source file. As an example of the operation of the preprocessor consider the following small source file.

```
extern int goodbye;

yousay(hello)
int hello;
{
    if (hello = = goodbye)
        goodbye = hello + 3;
    else
        goodbye = hello - 5;
}
```

Instrumenting this file with the HP Branch Validator will yield the following intermediate file which is then compiled.

```
1   static unsigned char __bA_array[3                    ] = {0};
2   extern int goodbye;
3
4   yousay(hello)
5   int hello;
6   {
7       __bA_array[0] = 1;if (hello = = goodbye)
8         {__bA_array[1] = 1;goodbye = hello + 3;}
9       else
10        {__bA_array[2] = 1;goodbye = hello - 5;}
11  }
12
13  struct __bA_probe_struct_         {
14              unsigned char insertprotocol;
15              char mapsuffix;
16              char sourcemodtime[9];
17              unsigned char options[4];
18              unsigned int numentries;
19              unsigned char *dataarray;
20              char *sourcepath;
21      };
22  struct __bA_probe_struct_ __bA_eybdoog_2omed_crs_
        abb_tsetws_ =
23  /*   Branch Validator internal version */
24  {  6,
25  /*   Map file extension */
26     'M',
27  /*   Date of compilation (encoded) */
28     {'B', '{', 'j', 'p', ' ', 'g', 'M', 'b', 'p'},
29  /*   Branch Validator compile options */
30     {0xa7, 0x05, 0x02, 0x00},
31  /*   Number of branches */
```
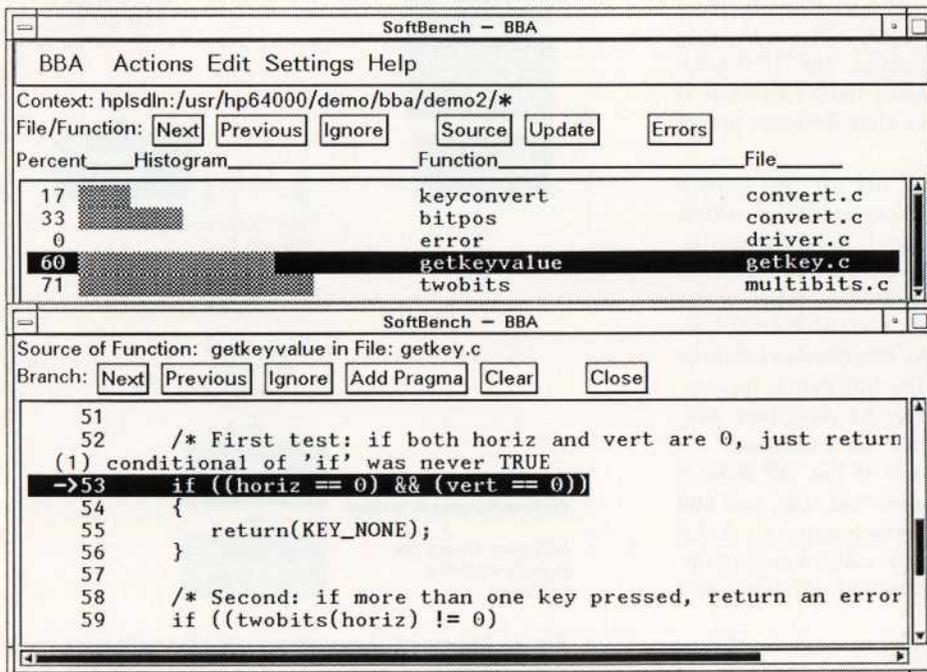

Fig. 3. *The HP Branch Validator in a SoftBench interface.*

```
32      3,
33   /*   Pointer to Branch Array */
34      _bA_array,
35   /*   Path to this file */
36      "/swtest/bba/src/demo2/goodbye.c",
37   };
```

This intermediate file shows the additional statements that the HP Branch Validator preprocessor has added to the original source. These additional statements include a branch array (char_bA_array [3] = {0}; defined on line 1), branch execution statements (_bA_array [x] = 1; defined on lines 7, 8, and 10) and a branch validator structure. The branch array on line 1 records the branches that have been executed. An array element is set to 1 when the corresponding branch is executed. The statements _bA_array[x] = 1; on lines 7, 8 and 10 set the appropriate elements of the branch array to 1 when they execute. Note that the branch statements are inserted so that they do not change the logic of the original program. The first statement on line 7 is used to record if this function is ever called. The statement on line 8 records whether the then part of the if condition executes. The statement on line 10 records whether the else part of the if condition executes. The preprocessor automatically includes the { and } symbols to maintain the logic of the original program. The use of these simple assignment statements and the character branch array keeps the performance time and memory space intrusion per branch to a minimum (typically the time and memory increases will be less than 20%). Finally on lines 13 to 36, the preprocessor adds a structure definition and a data structure for use by the HP Branch Validator report generator. This data structure records the date, filename, compile options, and a pointer to the branch array. This data structure will be dumped into a file at the appropriate time for processing by the report generator. The map file generated from this file is shown in Fig. 4.

This map file contains the information needed to map branch array elements to the source text of the branch. Each branch array element is designated by a line denoted :probe. The numbers associated with each probe completely describe for each branch the type and scope of the branch, source line and column ranges of the branch, and source line and column ranges of the statements the branch controls. Lastly, each branch statement in this file contains the text that is used as the controlling statement for the branch. This text is used to determine if a branch should be ignored.

The Two Preprocessors

The HP Branch Validator preprocessor is actually two preprocessors, which are called serially. The first preprocessor is essentially an ANSI C preprocessor with a special option to improve source referencing. The special source referencing option provides information about how macros are expanded to the second preprocessor. As an example, consider the following macro definition and macro call:

```
#define IF_GREATER_5_ SQUARE(a) if (a > 5) { a = a * a;}
IF_GREATER_5_SQUARE(c);
```

This results in the following code segment:

```
^Aif (c > 5) { c = c * c;} ^BIF_GREATER_5_ SQUARE(c) ^C;
```

The ^A, ^B and ^C (control-A, B, and C) are used to determine column and line positions of the original macro source when a macro is called. This information is used to map the branches associated with a macro to the original macro text. The first preprocessor also adds line number information to help match the line numbers of the processed text to the original text.

The second preprocessor parses the results of the first preprocessor. Using yacc and lex utilities, this second preprocessor interprets the file produced by the first preprocessor and adds the appropriate branch statements and array definitions. In addition, this second preprocessor generates the map file which contains the information necessary to map a specific branch to an associated source file. The following source program is used to show how the preprocessors operate.

```
1    /* Sample function showing many types of branches */
2
3    #define IF_GREATER_5_SQUARE(a) if (a > 5) { a = a * a; }
4
5    main()
6    {
7       int a = 0;
8       int b = 0;
9       int c = 0;
10
11      /* If statement with an else */
12      if (a = = 0)
13         b++;
14      else
15         c++;
16
17      /* If statement with no else */
18      if (b = = 0)
19         b++;
20
21      /* Switch statement with empty case, a case with fall-through
22         execution, and a default */
23      switch(a)
24      {
25      case 0:
26      case 1: b++;
```

```
:id Basis Branch Analysis Source Mapping File
:protocol 6
:options 205A7 :cppver @(A.01.00 03Oct90)@
:source  0 /swtest/bba/src/demo2/goodbye.c :modtime B(jp'gMbp
:probe  1 1 0 0 2 7 5   7 24 :ctl    0 6 2 6 22 @if (hello == goodbye)@
:probe  2 2 0 0 2 9 5   9 24 :ctl    0 6 2 6 22 @if (hello == goodbye)@
:probe  0 0 0 0 1 6 2 10   0 :fname  0 3 0 3 12 @yousay@
```

Fig. 4. *An example of an HP Branch Validator map file.*

```
27       break;
28    default: b++ ;
29    }
30
31    /* Switch statement with no default */
32    switch(a)
33    {
34    case 0: b++;      /* non-empty, no fall-through */
35       break;
36                      /* no default */
37    }
38
39       /* While statement */
40       while (b != 0)
41          b--;
42
43       /* For loop */
44       for (c = 0; c < b; c++)
45          b--;
46
47       /* Conditional assignment */
48       c = (a == 0) ? 123 : 321;
49
50       /* Do-while loop */
51       do
52          b++;
53       while (b != 0);
54
55       /* Macro conditional */
56       IF_GREATER_5_SQUARE(c);
57    }
```

The results of calling the first preprocessor would be:

```
1    #pragma CURRENT_DATE ,1990,06,12,12,45,54
2    #1 "sample.c",1990,06,12,12,41,10
3    /* Sample function showing many types of branches */
4
5
6
7    main()
8    {
9       int a = 0;
10      int b = 0;
.
.
.
```

Lines 11 through 55 are identical to lines 9 through 53 of the source file

```
.
56
57    /* Macro conditional */
58    ^Aif (c > 5) { c = c*c; } ^Bif_GREATER_5_SQUARE(c) ^C;
59    }
```

Note that except for ^A, ^B and ^C, this would be similar to the results of calling the standard C preprocessor. The results of calling the second preprocessor (to instrument the source file for branch analysis) would be:

```
1    # 1 "sample.c"
2    static unsigned char __bA_array[20         ] = { 0 };
3    static unsigned char __bA_temp_[1           ] = { 0 };
4    # 1 "sample.c"
5    /* Sample function showing many types of branches */
6
7
8
9    main()
10   {
11      int a = 0;
12      int b = 0;
13      int c = 0;
14
15      /* If statement with an else */
16      __bA_array[0] = 1;if(a == 0)
17      {__bA_array[1] = 1;b++;}
18      else
19      {__bA_array[2] = 1;c++;}
20
21      /* If statement with no else */
22      if (b == 0)
23      {__bA_array[3] = 1;b++;else __bA_array[4] = 1;
24
25      /* Switch statement with empty case, a case with fall-through
26         execution, and a default */
27      switch(a)
28      {
29      case 0:
30      __bA_array[5] = 1; goto _ bA_ switch_label_1_1;case 1: \
         __bA_array[6] = 1;__bA_switch_label_ 1_1:b++;
31         break;
32      default: __bA_array[7] = 1;__bA_ switch_label_1_2:b++;
33      }
34
35      /* Switch statement with no default */
36      switch(a)
37      {
38      case 0: __bA_array[8] = 1;_ bA_ switch_label_1_3:b++; \
                  /* non-empty, no fall-through */
39         break;
40                  /* no default */
41      default:_bA_array[9] = 1;}
42
43      /* While statement */
44      __bA_temp_[0] = 0;while (b != 0)
45         {__bA_array[10] = 1;_ bA_temp_[0] = 1;b--;}
         if(__bA_temp_[0] == 0)_ bA_ array[11] = 1 ;
46
47      /* For loop */
48      for (c = 0; c < b; _bA_array[12] = 1,c++)
49         {__bA_array[13] = 1;b--;}
50
51      /* Conditional assignment */
52      c = (a == 0) ? (__bA_array[14] = 1,123) :
                       (__bA_array[15] = 1,321);
53
54      /* Do-while loop */
55      do
56         b++;
57      while ((b != 0)?(__bA_array[16] = 1,1):(__bA_array[17] = 1,0));
```

```
58
59      /* Macro conditional */
60      if (c > 5) {__bA__array[18] = 1; c = c * c; else __bA__array[19] = 1;;
61      __bA__label: __bA__dump();
62      {struct __bA__probe__struct__
63              unsigned char insertprotocol;
64              char mapsuffix;
65              char sourcemodtime[9];
66              unsigned char options[4];
67              unsigned int numentries;
68              unsigned char *dataarray;
69              char *sourcepath;
70      };
71      struct __bA__probe__struct__
        __bA__elpmas_2omed_abb_omed_00046ph_rsu_2__csid__        =
72      /*  Branch Validator internal version  */
73      {   6,
74      /*  Map file extension  */
75          'M',
76      /*  Date of compilation (encoded)  */
77          {'B', '{', 'm', ' ', 'h', 'F', 'I', 'Y', ' ' },
78          {0xff, 0xff, 0x03, 0x00},
79      /*  Branch Validator compile options  */
80          20,
81      /*  Number of branches  */
82          __bA__array,
83      /*  Pointer to Branch Array  */
84          "/disc.37.2/usr/hp64000/demo/bba/demo2/sample.c",
85      /*  Path to this file  */
86      };
```

The second HP Branch Validator preprocessor does the following:

- Creates an array with one entry for each branch. Each entry in the array is initialized to 0. Line 2 shows that 20 branch entries are initialized to 0.
- For each branch that is located:
  □ An array entry is assigned to the branch.
  □ An assignment statement is inserted as the first executable statement within the branch. The assignment statement sets the associated array entry to 1, showing that it was executed. Also, the assignment statement is inserted on the same line as the first executable statement so that debuggers can report the correct line numbers of the source statements. Line 17 shows that array element 1 is set when the true condition of the if (a = = 0) branch is executed. Line 19 shows that array element 2 is set when the false condition of the if (a = = 0) branch is executed.
  □ An entry into a map file is created for that array entry. The entry into the map file specifies what type of branch it is (e.g., an if statement or a for loop) and the associated line number and column numbers of the branch.
- At the end of the newly created source file an environment data area is created on lines 62 through 85. This data prevents mismatching the map files, and allows changes in the source file to be monitored. The data area specifies the date of compilation, number of branches, and other parameters.
- Finally, the second preprocessor inserts calls to a branch

analysis dump routine. The dump routine (line 61) dumps the data contained in the data arrays to a dump file when the routine exits. An additional feature allows this dump routine to be called in response to an HP-UX system signal or before a specified routine.

The second preprocessor will add branch analysis assignments for a variety of conditions, including:

- Functions. An array assignment is added as the first executable statement of a function. Line 16 shows that array element 0 is set to 1 when the function is executed.
- If Statements. An array assignment is added as the first executable statement of the then part of any if statement.
- Else Statements. An array assignment is added as the first executable statement of the else part of any if statement. Optionally, else conditions and array assignments can be added in the locations where else statements have not been written.
- While Statements. An array assignment is added as the first executable statement within any while loop. Optional tests can be performed to determine if a while loop always executes (never skipping the loop) or if a do-while loop is executed more than once. Lines 44 to 45 and 55 to 57 illustrate the complete testing of the while loop. Note that the while loop on lines 44 to 45 makes use of a temporary variable to verify that the loop is sometimes skipped.
- For Statements. An array assignment is added as the first executable statement within any for loop. Optional tests can be performed to determine if the third statement of a for loop is ever executed. The third statement of a for loop may be skipped if the for loop contains a break or a return.
- Case and Default Statements. An array assignment is added as the first executable statement following a case or default statement within a switch statement. Optional tests can be added to detect if the code of a given case or default statement is executed because of fall-through from a higher-level case statement or a direct call to the specified case. In addition, cases with no defaults can detect the execution of a default by adding a default and observing if the default is ever executed (line 41).
- Conditional Assignments. Two array assignments can be optionally added to detect when conditional assignments evaluate TRUE and FALSE.

Additional features of the preprocessor allow the user to exclude or ignore specified branches from analysis, or to alert the user if a particular branch is ever executed. The map file contains this information and will direct the report generator to take the appropriate actions when either of these conditions is found. All of these different measurements can be selected at compile time using command line options.

After the HP Branch Validator's two preprocessors finish, the code is ready to be compiled with the standard C or ANSI C compiler. After compilation and linking in the HP Branch Validator dump utility (which knows how to dump the HP Branch Validator data of each file), the resulting executable is executed and run through a series of tests. When the executable exits or is instructed to dump its data, it will generate a dump file. This file is interpreted by the report generator to generate summary and source level reports about the branch coverage results.

The Report Generator

The HP Branch Validator report generator produces summary reports that show the number and percent of branches executed or hit in each of the functions instrumented by the preprocessors, and source reports that provide information about the individual branches that are executed. The report generator uses the information stored in the branch arrays during program execution combined with the preprocessor-generated map files to produce its reports. A typical summary report looks like:

hit	total	%	IA	function	file
8 /	12	(66.67)		keyconvert	convert.c
6 /	9	(66.67)		bitpos	convert.c
0 /	1	(0.00)		error	driver.c
3 /	5	(60.00)		getkeyvalue	getkey.c
5 /	7	(71.43)		twobits	multibits.c

22 out of 34 retained branches executed (64.71%)
[23 branches were ignored]

In this example, the function keyconvert in file convert.c had 66.67 percent of its branches executed or 8 out of 12. The total branch coverage was 64.71 percent or 22 out of 34 branches for all five of these functions. Twenty-three branches were ignored during this analysis.

The source report portion of the report generator presents information about the individual branches of a function that are not executed. The information presented includes the line number of the branch, the text of the branch, and information describing what part of the branch did not execute. A typical source report looks like:

```
keyconvert      convert.c
    (1) conditional of 'if' was never TRUE
36 ->       if (keycode > KEY_9)      /* special case for 0 key */

    (1) conditional of 'if' was never FALSE (no 'else' statement)
44 ->           if (horiz = = 1)

    (1) conditional of 'do while' was never TRUE
49 ->       while (horiz != 0);

    (1) conditional of 'do while' was never FALSE
49 ->       while (horiz != 0);
```

8 out of 12 branches executed (66.67%)

In this example, the function keyconvert in file convert.c had four branches that were not executed. First, the branch on line 36 was never TRUE (the value of keycode was never greater than KEY_9). Second, the branch on line 44 was never FALSE (the value of horiz always equaled 1). The statement no 'else' statement is an informational message stating that there was not an else condition in the original code. Finally, the two branches on line 49 are the same branch. The fact that the branch was never TRUE and never FALSE is an implicit way of saying that this branch was never executed at all. A return statement must have occurred before this branch.

If an entire group of branches is controlled by a single, higher-level branch, the controlling branch will be the only one indicated as not executed. The number in parentheses with the description indicates the number of hidden branches that are not executed because of the controlling statement. For example:

```
    (3) 'then' part of 'if' was never executed
52          /* First test: if both horiz and vert are 0 */
53          if ((horiz = = 0) && (vert = = 0))
54          {
55 ->           is_error = check_error();
56              if (is_error)
57                  return(KEY_ERROR);
58              else
59                  return(KEY_NONE);
60          }
```

The fact that this branch did not execute implies that the three underlying branches did not execute. There are three branches, one for the then part of the if ((horiz = = 0) && (vert = = 0)) branch, one for the then part of the if (is_error) branch, and one for the else part of the if (is_error) branch. Observe also that this example shows that the source report can present more lines of text about the branches of interest if desired.

Each execution of an HP Branch Validator instrumented program will append the data stored in the branch arrays along with the environment data to the dump file.

The report generator then examines the dump file and the map files of each source file to find the information for generating summary and source line reports. In particular the report generator cycles through each branch array of each of the instrumented source files. Each element of the array is mapped back into the source mapping file, which defines the type of branch and the expected value of the array element for the branch to be executed. If a particular branch has not executed, the type of branch and the reason it did not execute will be reported. The text surrounding the branch is determined from the line numbers of the branch saved in the mapping file. The text of the branch is then read back from the original source file. As an example:

```
    (4) conditional of 'if' was never TRUE
53  - > if (value = = 0) if (value > 100)
         ^^^^^^^^^^^^^
```

The then part of the branch if (value = = 0) may have been assigned array element 35. Because this branch did not execute, the value of array element 35 was zero. When the report generator found array element 35, it mapped it back to the then part of an if conditional statement. The mapping also provided the report generator with the line number for this branch, the column range of the text for the if conditional, and the number of branches (4) affected because this branch did not execute.

The objective of using the HP Branch Validator is to achieve a branch coverage percent that is as high as possible. Adding additional test suites is the primary method of creating a higher coverage percent. But some branches do not lend themselves to being easy to test or perhaps they have already been fully tested by another user. In either case it may become important to ignore these
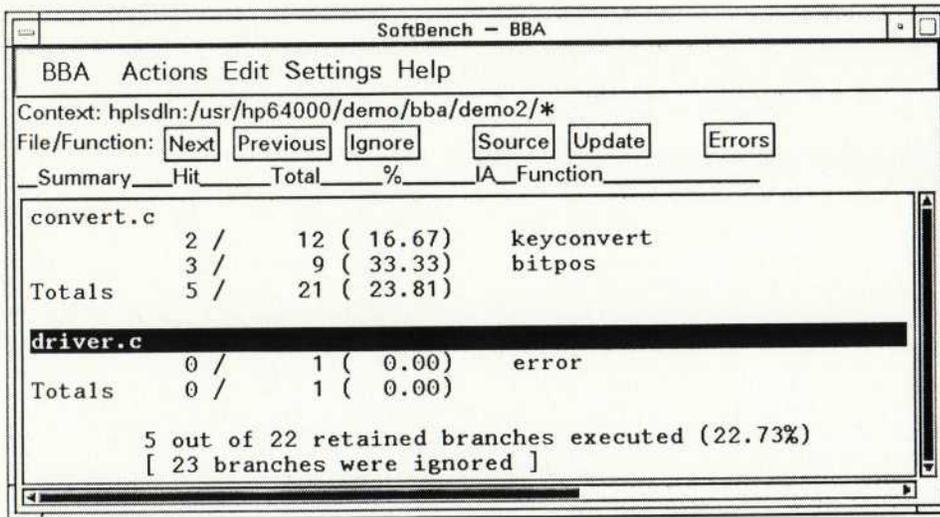
Fig. 5. A file and function summary report.

```
                        SoftBench — BBA

 BBA   Actions Edit Settings Help
Context: hplsdln:/usr/hp64000/demo/bba/demo2/*
File/Function: [Next] [Previous] [Ignore]    [Source] [Update]    [Errors]
_Summary___Hit_____Total_____%_____IA_Function_____

convert.c
             2 /      12 ( 16.67)      keyconvert
             3 /       9 ( 33.33)      bitpos
 Totals      5 /      21 ( 23.81)

driver.c
             0 /       1 (  0.00)      error
 Totals      0 /       1 (  0.00)

         5 out of 22 retained branches executed (22.73%)
         [ 23 branches were ignored ]
```

branches or exclude entire functions from branch coverage percents. A mechanism exists in the HP Branch Validator that enables the user to tell the report generator to ignore files, functions, and branches while it is creating reports. To use this capability the user enters the names of files, functions, and branches to be ignored into an *ignore file*. When the report generator is preparing a report it will examine the ignore file and remove from the report any of the specified items (files, functions, and branch statements) listed. The report generator does this by entering each of the items in the ignore file into an ignore hash table. When each file, function, and branch is sequentially examined it will first be checked to see if it is in the ignore hash table. If the item exists in the ignore hash table it will be ignored (or skipped). To ignore a branch the user must enter a predefined sequence of characters that represents the controlling statement of the branch. To ignore branches, the report generator matches the controlling statements for each branch contained in the map file with the controlling statements in the ignore file. For example, if the ignore file contains the information:

    /users/cc/driver.c
    /users/cc/convert.c : keyconvert
    /users/cc/convert.c : bitpos : case 32:

the following branches will be ignored:
■ All of the branches in the file /users/cc/driver.c

■ All of the branches in the function keyconvert in the file /users/cc/convert.c
■ The branches associated with case 32: in the function bitpos in the file /users/cc/convert.c.

The SoftBench Interface

The SoftBench interface is a friendlier user interface than the command-line interface of the HP Branch Validator. The SoftBench interface allows the user to examine the results of the branch coverage interactively, find unexecuted branches quickly, and take appropriate actions to improve the branch coverage results. By simply clicking the mouse button, the user is able to perform operations such as ignoring a branch, editing a source file containing a missing branch, or positioning the sprite on the next or previous unexecuted branch.

The SoftBench interface provides three primary reports: a file and function summary report, a function histogram report, and a source report. The file and function summary report is similar to the standard report generator report except that it provides summary information about the number of branches executed on a file level as well as function level coverage information. Fig. 5 illustrates the file and function summary report. The function histogram presents the same function summary information in a histogram format (see Fig. 6). The function histogram report helps the user quickly determine the functions that do not meet specified branch coverage criteria.


```
                        SoftBench — BBA

 BBA   Actions Edit Settings Help
Context: hplsdln:/usr/hp64000/demo/bba/demo2/*
File/Function: [Next] [Previous] [Ignore]    [Source] [Update]    [Errors]
Percent____Histogram_____Function_____File_____

   17  ▓▓▓▓▓                          keyconvert             convert.c
   33  ▓▓▓▓▓▓▓▓▓                      bitpos                 convert.c
    0                                 error                  driver.c
   60  ▓▓▓▓▓▓▓▓▓▓▓▓▓▓▓▓               getkeyvalue            getkey.c
   71  ▓▓▓▓▓▓▓▓▓▓▓▓▓▓▓▓▓▓             twobits                multibits.c
```
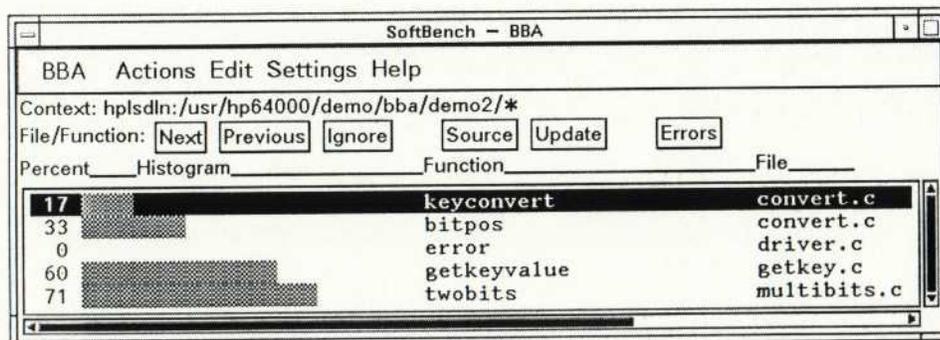
Fig. 6. A function histogram report.

```
┌─────────────────────────────────────────────────────────────────┐
│ ─                   SoftBench — BBA                        □  ▢   │
│ Source of Function: keyconvert in File: convert.c                 │
│ Branch: [Next] [Previous] [Ignore] [Add Pragma] [Clear]  [Close]  │
│ ┌───────────────────────────────────────────────────────────┐▲  │
│ │    (1) switch never went to 'case'                        │▒  │
│ │ ███──>28███████case 2:██/*██pressing█second█column█of█keys█*/│▒ │
│ │     29           keycode = bitpos(horiz) + KEY_F;          │   │
│ │     30           break;                                    │   │
│ │    (1) switch never went to 'case'                        │   │
│ │  ──>31           case 4:  /* Pressing third column of keys */│▒│
│ │     32           keycode = bitpos(horiz) + KEY_1;          │▒  │
│ │     33           break;                                    │▒  │
│ │    (3) switch never went to 'case'                        │   │
│ │  ──>34           case 8:  /* pressing fourth column of keys */│ │
│ │     35           keycode = bitpos(horiz) + KEY_6;          │▼  │
│ │ ◄────────────────────────────────────────────────────►   │   │
│ └───────────────────────────────────────────────────────────┘   │
└─────────────────────────────────────────────────────────────────┘
```

Fig. 7. *A source report for the function* keyconvert.

The source report combines the descriptions of the missed branches and the complete text of the selected source file or function into a source listing. This source listing allows the user to see the missed branches in the context of the complete file or function more clearly. Fig. 7 shows the selected source report and the missed branches in the context of the selected function.

The SoftBench interface uses the report generator and various filters to produce these summary, histogram, and source reports. When the user selects a summary or histogram report, the report generator is run to produce a summary report which is then passed through a filter to produce the corresponding summary or histogram report. When the user selects a source report, the report generator is run to produce a source report. A second filter program combines the statements about the missed branches with the source file. The location of the source file and the range of text for a function are read from the map file.

The Encapsulator was used to build the Softbench interface. The Encapsulator provided many benefits, perhaps the foremost being that we could produce a mouse-driven OSF/Motif[3] style interface quickly for an existing command-line-driven product. We were able to demonstrate many of the features of the HP Branch Validator six weeks after we started using the Encapsulator. A second benefit of the Encapsulator is the SoftBench message server.[4] The message server allowed us to simplify our interactions with the SoftBench edit and builder tools. We could simply send an edit command and not be concerned about the type of editor the user might want to use and the appropriate command to start the editor.

A third benefit we found in using of the Encapsulator is the Encapsulator description language. The Encapsulator description language provides a tool developer with the capability of defining menu buttons and their associated operations very easily. This feature enabled us to try various features and develop the proper look and feel for our product. As an example, the following text defines a pull-down menu to show the HP Branch Validator histogram.

```
make_object(menu, "showHistogram", MenuButton, "Show Histo-
gram", NULL, make_event(User, "Select", bba_histogram()));
```

Observe that even without explaining the details of the Encapsulator description language or how menu buttons work, the meaning of the above command is almost intuitive. This text defines a menu button with the text "Show Histogram". The object will be named "showHistogram", and it will call the function bba_histogram() when the user selects this menu item.

In addition, the Encapsulator makes it easy to define display windows where text is to be displayed. For example, the following text describes the window to display the summary text.

```
results = make_object(pane, "summaryList", List, NULL,
              (XOFFSET : 0 WIDTH : 678 HEIGHT : 255),
              make_event(User, "Select", line_select()));
```

This text declares an object to be a list widget with the specified width and height in pixels. The object is to be named summaryList and will call the function line_select() when the user selects a line. The list widget contains the methods (functions) that know how to scroll and do all of the appropriate operations for a listing display. When it is necessary to write information to the list widget, a file is loaded and the widget handles the rest. Invoking the load command (i.e., load_file(results, tmp_summary)), results in loading the contents of the file tmp_summary into the object defined as results (i.e., the list widget).

A fourth benefit of the Encapsulator is the capability for a process to send commands to a hidden shell or subprocess and get called when the command completes. This feature is used to send report generator and filter commands to a hidden shell. This is perhaps not the most technical feature, but the fact that the commands can execute quickly and that an application can be informed at the completion of a command was another factor in saving development time in producing an HP SoftBench interface to the HP Branch Validator.

Along with the benefits of using the Encapsulator there are a few drawbacks. The primary drawback is the limited set of widgets that can be used and the fact that they must be used in fairly rigid ways. One of the problems associated with the list widget (the widget that allows the user to display textual information) is that it is hard to determine

exactly which line a user has selected because the text of the line is returned but not the line number. If there are two identical lines of text on the screen, it is impossible to know which line the user selected. To work around this problem, we added a line number to each line. The numbers are appended to the end of each line of text in an area of the list widget that is typically off-screen. The user can observe these numbers if the list widget is scrolled to the right, but in typical situations they are not shown. When the line of text that a user selects is read in, the line number is also read to determine which line the user selected and which line to go to next. A second problem with the Encapsulator is not being able to specify a multiple-word label directly as the name of the application. The HP Branch Validator cannot be named the "Branch Validator" on the display because the Encapsulator will only accept a contiguous character string as a name. In addition, the name should be about 9 characters or less. To get around this problem, we elected to use the name BBA which reflects the underlying technology of the HP Branch Validator, specifically basis branch analysis.

In general, the Encapsulator provided the HP Branch Validator with an easy-to-implement graphic interface. This dramatically reduced our time to market while allowing us to maximize the number of features. In addition, porting the HP Branch Validator SoftBench interface to other platforms has required only a simple recompilation of the Encapsulator file. However, porting the HP Branch Validator preprocessor and the HP Branch Validator report generator has not been quite as simple.

C++ Branches

The HP Branch Validator can also provide information about C++ branches that are not executed. The way to get C++ to work with the HP Branch Validator is to compile the C++ into C files and then let the HP Branch Validator compile the C files. Using the C++ +i compiler option in coordination with other options allows the compilation of C++ source files to object files to occur in one step. During the process of compiling C++ to C files, line and file directives about the C++ file are put into the standard C file. The HP Branch Validator preprocessor will pick up these line and file directives and reference a branch in the C file back to the original C++ source line that caused the branch.

The results of C++ branch analysis can be seen via the command-line or SoftBench interface. Since the HP Branch Validator is primarily a C tool and the C++ function names that it sees are encrypted, the command-line version of the HP Branch Validator must be run through the c++filt(1) utility to produce more readable function calls. The following is an example of a C++ summary report.

hit	_total_	_%_	_IA_	_function_	_file_
1/	1	(100.00)		BoardRec::SetPiece(PieceRec*)	board.c
1/	1	(100.00)		BoardRec::Clear()	board.c
1/	1	(100.00)		BoardRec::Repaint(ScrnRec*)	board.c
3/	3	(100.00)		BoardRec::Repaint()	board.c
2/	3	(66.67)		DoReplace(void*)	board.c
3/	3	(100.00)		DoPaintMe(void*)	board.c
19/	21	(90.48)		Paint(int,int,Direction,Direction,int)	flow.c
10/	11	(90.91)		TripFlow(void*)	flow.c
6/	7	(85.71)		FlowRefigureDeltas()	flow.c
4/	5	(80.00)		FlowRepaint()	flow.c
2/	3	(66.67)		FlowFast()	flow.c
1/	1	(100.00)		FlowStart()	flow.c
5/	5	(100.00)		InitLevel()	init.c
2/	3	(66.67)		PieceRec::Paint(ScrnRec*,int,int)	piece.c
7/	9	(77.78)		ScrnRec::RefigureSizes()	screen.c
1/	1	(100.00)		ScrnRec::GetPieceCache(PieceRec*)	screen.c

A portion of the HP Branch Validator source report for the above file would look like:

```
DoReplace(void*) board.C
    (1) conditional of 'if' was never FALSE (no 'else' statement)
        [control lines]
56
57
58        static void DoReplace(void *){
59 ->         if (newpiece) {
60                newboard->SetPiece(newpiece);
61                newpiece = NULL;
62            }
```

```
Paint(int,int,Direction,Direction,int) flow.C
    (1) 'case' code was never executed
74                XFillRectangle(dpy, window, gc, x0 + ssize
                    - ipix,
75                             y0 + lssize, ipix, flowwidth);
76            break;
77        case BadDir:
78 ->          Punt("Bad in direction in Paint!");
79            }
80        }
81        opix = numpix - ipix;

    (1) 'case' code was never executed
97                XFillRectangle(dpy, window, gc, x0 + hssize, y0
                    + issize,
98                             opix, flowwidth);
99            break;
100        case BadDir:
101 ->          Punt("Bad out direction in Paint!");
102            }
103        }
104    }
```

Note that the only inconsistency in the command-line interface of the HP Branch Validator is in the filenames reported in the summary reports. The filenames reported are the C filenames and not the C++ filenames. Simply replace the C extension with the C++ extension to determine the C++ filename. In this example replace .c with .C or board.c with board.C.

The command-line interface has one additional problem that appears when trying to ignore branches. To ignore a specific branch the C filename plus the entire C++ encrypted function name and encrypted branch must be entered into the ignore file. For example, to ignore the case statements on lines 78 and 100 the sequence:

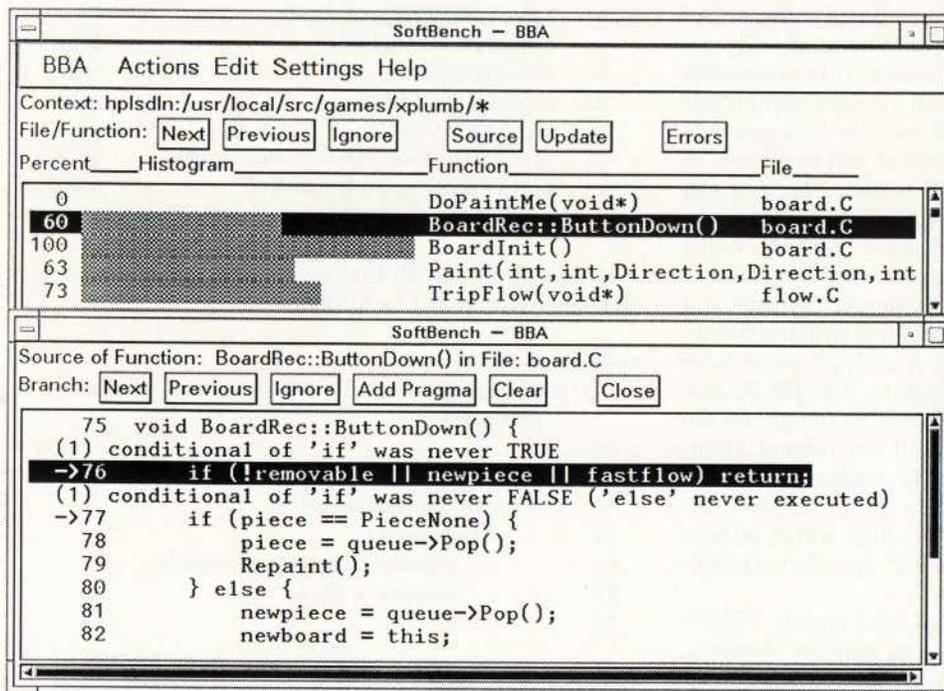/disc.14.4/games/xplumb/flow.c: Paint FiTl9DirectionT3T1:case −1:

Fig. 8. *A portion of a C++ program as it would appear in an HP Branch Validator Softbench interface.*

must be entered into the ignore file. This becomes a little difficult because so many characters have to be typed into the ignore file and because the file flow.c must be examined to determine the branch function name and statement as the HP Branch Validator sees it.

The SoftBench interface overcomes these deficiencies by using the map files and various filters to produce the correct C++ results. All C files are mapped back to C++ files if they exist. In addition, the effort to ignore a branch is simply the click of a mouse button. The SoftBench interface reads the map files and determines the proper filename, function name, and branch string of functions and branches to ignore. Fig. 8 shows an example of a portion of a C++ program being examined using the HP Branch Validator SoftBench interface. This shows the benefits of using the Encapsulator to overcome deficiencies and greatly enhance the capabilities of a command-line interface.

Conclusion

The HP Branch Validator provides a mechanism to quantify the level of testing that has occurred on a software product. This quantification lets the user rapidly define software test sites that provide an acceptable level of software testing. In addition, the HP Branch Validator indicates "logically dead" code (branches that are never executed) and can be a useful tool in the reverification of an existing software product.

The HP Branch Validator can operate through a Soft-Bench encapsulated interface or through a command-line interface. The SoftBench interface will work independently of SoftBench or in conjunction with SoftBench version A.01.00 or greater. The HP Branch Validator works well in 8M bytes of memory when it is used in stand-alone mode with the SoftBench encapsulated interface. However, 12M bytes of memory are required when the HP Branch Val-

idator is used with SoftBench. The command-line interface does not have any major memory constraint. The command-line interface provides the same branch coverage information as the SoftBench interface but without some of the convenience provided by an OSF/Motif graphical user interface.

Acknowledgments

The HP Branch Validator is based on the HP Basis Branch Analyzer,[5] which was originally designed and introduced by Bruce Erickson. Grant Grovenburg made extensive changes to the core technology of the HP Branch Validator to better adapt it to the native and SoftBench environment. Thanks also to Eric Kuzara and Dave Wallman for many hours of management and marketing support.

References

1. B. D. Fromme, "HP Encapsulator: Bridging the Generation Gap," *Hewlett-Packard Journal*, Vol. 41, no. 3, June 1990, pp. 59-68.
2. M. R. Cagan, "The HP SoftBench Environment: An Architecture for a New Generation of Software Tools," *Hewlett-Packard Journal*, Vol. 41, no. 3, June 1990, pp. 36-47.
3. *Hewlett-Packard Journal*, Vol. 41, no. 3, June 1990, pp. 6-35.
4. C. Gerety, "A New Generation of Software Development Tools," *Hewlett-Packard Journal*, Vol. 41, no. 3, June, pp. 48-58.
5. B. Erickson, "Testing Strategies for High Quality Software Include Basis Branch Analysis," *HP Design Center Magazine*, no. 3, 1988, pp. 28-34.

Authors

April 1991

6 ══ Synthesized Sweeper ══

Roger P. Oblad

Roger Oblad was the program manager during development of the HP 8360 family of high-end microwave synthesized sweepers. He joined HP's Microwave Division in 1972, shortly after earning a BSEE degree from the University of Utah. He earned an MSEE degree in 1975 from Stanford University. Roger has been a development engineer for the HP 8505 network analyzer (analog circuit design), the HP 8501 storage normalizer (digital hardware design and firmware), and the HP 8340 synthesized sweeper (firmware). He later became project manager for the HP 8340. He is named as an inventor in two patents. A six-year veteran of the U.S. Air National Guard, he was born in Salt Lake City, Utah,
lives in Santa Rosa, California, is married and has five children.

James E. Bossaller

An enthusiastic flyer, Jim Bossaller recently received his private pilot's license. He also piloted the HP 8360 microwave synthesized sweeper family as the project manager responsible for development of mechanical systems and firmware. He joined HP in 1969 in Colorado Springs as a cathode ray tube test technician. Jim worked full-time and attended college until 1975, when he earned a BSEE degree from the University of Colorado. In 1975, he was transferred to Sunnyvale, California, where he served as a development engineer on the HP 9411 and HP 9413 matrix switching system products. In

1978, he moved to HP's Network Measurements Division as a materials engineer and member of the HP 8340 microwave synthesizer development team. He served in the U.S. Air Force from 1964 to 1968. Born in Inglewood, California, he lives in Rohnert Park, is married, and has five children.

John R. Regazzi

John Regazzi joined HP's Network Measurements Division in 1979. He earned a BSEE degree in 1976 from Rutgers University and an MSEE degree in 1979 from Lehigh University. John was the system testing project leader on the systems' evaluation team for the HP 8360 family of microwave synthesized sweepers. He also served as an R&D engineer and production engineer for the HP 8350 microwave sweeper and millimeter modules. He is now serving as a project manager. He is named as an inventor in a patent for an analog circuit design. Born in Madison, New Jersey, he lives in Santa Rosa, California. He is married, has two children, and is active in youth baseball.

17 ⎓ Sweeper Self-Test Design

Michael J. Seibel

Development engineer Mike Seibel designed the self-test and service firmware and some of the analog and digital circuitry for the HP 8360 synthesized sweepers. He also was a member of the HP committee that created the now industry-accepted SCPI (formerly TMSL) instrument control language. Mike worked summers at HP in 1976 and 1977 while earning his BSEE degree from the California Polytechnic State University in San Luis Obispo in 1979. He joined HP's Santa Rosa Division in 1979 as a production engineer on the HP 8554, HP 8555, and HP 8566 spectrum analyzers. He then developed the test process for the HP 8340 synthesizers, and served as production engineer for the product. Born in San Francisco, Mike lives in Windsor, California. He is married, has two children, and enjoys racquetball, boardsailing, car rallying, audio engineering, and playing drums in local jazz and rock bands.

24 ⎓ Sweeper Power Leveling

Lance E. Haag

High-frequency analog electronics, signal theory, and electromagnetic theory are the professional interests of Lance Haag. Born in St. Paul, Minnesota, he earned a BSEE degree (1985) at the University of Minnesota and an MSEE degree (1988) from Stan-

ford University. Lance joined HP's Network Measurements Division in 1985. He has done production engineering for HP 8350, HP 8620, and plug-in sweeper products, and he developed a production test process and did systems engineering for the HP 8360 synthesized sweeper family. Before joining HP, he was an R&D engineer with Interlaken Technology in Minnesota, developing force, pressure, and temperature sensors. Lance lives in Santa Rosa, California, and enjoys cross-country skiing, bicycle racing, and cooking.

Mark N. Davidson

Born in Atlanta, Georgia, Mark Davidson received his BSEE degree (1979) and MSEE degree (1984) from the Georgia Institute of Technology. He joined HP's Network Measurements Division in 1981. As a product R&D engineer, Mark designed the level control and frequency synthesis hardware for the HP 8360 synthesized sweeper family. He has also worked as a production engineer for microwave sweeper products. Before joining HP, Mark was an R&D engineer with Tektronix Corporation. He lives in Santa Rosa, California, where he enjoys running, backpacking, fast motorcycles, and music.

Glen M. Baker
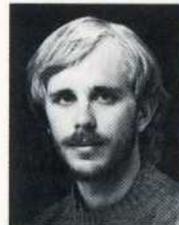
A member of the Squaw Valley Ski Patrol in California, Glen Baker used his keen eye for sweeping slopes while designing firmware for the HP 8360 synthesized sweeper family. Previously, he was a production engineer for the HP 8340 and HP 8350 synthesizers and several microcircuits. He joined HP's Network Measurements Division in 1983, shortly after earning a BSEE degree from the University of Washington. Born in Seattle, Washington, he lives in Santa Rosa, California, and enjoys skiing, boardsailing, acrobatic flying, hang gliding, mountain biking, competitive sailing, and music.

31 ⎓ Switched Doubler

James R. Zellers

Jim Zellers designed the package and switched filtering for the HP 0.01-to-40-GHz switched frequency doubler, and served as a project manager for the HP 83597A 40-GHz plug-in in the Network Measurements Division R&D lab. After joining HP's Santa

Clara Division in 1968, he worked on swept synthesizer products, digital displays, and RF network analyzers. Before coming to HP, Jim was an engineer at Hughes Aircraft Company. He earned a BSEE degree in 1966 from the University of California at Berkeley and an MSEE degree in 1968 from the University of Michigan. Born in Los Angeles, Jim lives in Santa Rosa, is married, and has a teenage daughter and son. He enjoys swimming, hi-fi music, and tinkering with computers.

34 ⎓ High-Speed Pulse Modulator

Mary K. Koenig

The fast pulse modulator to improve performance of the HP 8360 synthesized sweeper was designed by development engineer Mary Koenig, who is named an inventor in a patent for the modulator design. She also developed the HP 8349A amplifier after joining HP's Network Measurements Division in 1982. Mary received a BS degree in 1981 in electrical engineering from the University of Washington in Seattle. Born in Rochester, New York, she lives in Santa Rosa, California, and enjoys white-water kayaking and rafting.

36 ⎓ New Microcircuit Technology

Ronald C. Blanc

Development engineer Ron Blanc developed the mod-splitter microcircuit used in the HP 8360 synthesized sweeper. He joined HP's Network Measurements Division as a microcircuit production engineer in 1975, shortly after graduating from Harvey Mudd College in Claremont, California with a Master's degree in engineering. Ron was born in San Mateo, California, and lives in Santa Rosa. He is married, and enjoys hiking and tennis.

Richard S. Bischof

Project leader for development of two new microcircuits for the HP 8360 synthesized sweepers, Stan Bischof joined HP's Network Measurements Division in 1979. He developed microcircuits for the HP 8340 sweeper and is now working on productivity engineering, microwave design tool development, and support. Stan's professional interests include RF/microwave design, controls, filters, and electronic tools. He received his BSEE degree (1975) and his MSEE degree (1978) from the University of California at Santa Barbara, where he

was born. He now lives in Santa Rosa, is married, has two children, and enjoys badminton and gardening.

Patrick B. Harper

Microcircuit introduction project leader Pat Harper helped develop the test processes and test systems for the HP 8360 microcircuits. As a microcircuit technician after joining HP's Stanford Park Division in 1980, he helped introduce microcircuits for the HP 8340 sweeper and the HP 8753 network analyzer. Born in San Francisco, Pat lives in Santa Rosa, is married, and has one child. He enjoys skin diving and gardening.

David R. Veteran

A 39-year veteran with HP, Dave Veteran joined HP's Microwave Division in 1952 when the company reached its 13th year. As a project engineer, he has designed families of HP microwave components, including pads, terminations, switches, and step attenuators—and the dc-to-50-GHz programmable step attenuators for the HP 8360 synthesized sweep oscillators. Dave, who attended DeAnza and Foothill Colleges in California, is named as an inventor in four patents, with a fifth one pending, on coaxial microwave devices. A native of San Jose, California, Dave lives in Santa Rosa, is married, and has two daughters. He is a pilot, airplane owner, amateur winemaker, and 10-acre landowner.

Mohamed M. Sayed

Mohamed Sayed was the project manager for the HP 83557A V-band and the HP 83558A W-band millimeter-wave source modules. Previously, he helped develop the HP 85106A/B/C millimeter wave network analyzer system, the HP 85109A/B on-wafer network analyzer system, the HP 8349B microwave amplifier, the HP 8355x and HP 85100V/W source modules, the HP 70907A external mixer interface module, the HP 71300A millimeter spectrum analyzer, the HP 5350A microwave frequency counter, and the HP 5356A/B/C microwave frequency converters. He is now project manager for microwave and millimeter-wave amplifier products. A member of the IEEE and author of over 15 articles related to his work, Mohamed's professional interests include microwave and millimeter-wave applications. He earned a BSEE degree (1964) and

an MSEE degree (1968) from Cairo University, and a PhD degree (1973) from the Johns Hopkins University. Before joining HP's Stanford Park Division in 1973, Mohamed was a postdoctoral fellow in solar cell studies at the University of Delaware (1972 to 1973). He was a lecturer in electrical engineering at San Jose State University from 1978 to 1982. Born in Cairo, Egypt, he is married and has two children. He is president of the Santa Rosa Youth Central Soccer Club, coaches and referees youth soccer, and enjoys traveling with his family.

Giovonnae F. Anderson

Codeveloper of the HP 83557A and HP 83558A millimeter-wave source modules at HP's Network Measurements Division, Giovonnae Anderson joined HP's Microwave Technology Division in 1979 as a development engineer. She has also designed beam-lead diodes, and is now working on pulse measurement systems. Giovonnae received a BS degree in physics from the Hampton Institute in 1970, an MSEE degree from Cornell University in 1972, and a PhD degree in electrical engineering from the University of California at Davis in 1979. Before joining HP, she worked on magnetic materials at IBM in San Jose. Giovonnae is the author of two articles published in the IEEE Journal of Solid State Devices, and is a member of the IEEE and the Society of Women Engineers. Her professional interests include semiconductor materials and devices. Born in Richmond, Virginia, Giovonnae lives in Santa Rosa, California, is married, has two children, and enjoys bicycling, camping, and sewing.

David M. Hoover

Dave Hoover likes to build things, including his own house in Spokane, Washington. He also builds cellular telephone testing instruments for HP. His latest project, the HP 11846A DQPSK I-Q signal generator, tests North American dual-mode digital cellular telephone receivers. After joining HP's Spokane Division in 1981, he worked on the HP 8645A signal generator and developed a frequency doubler for the HP 8644A, HP 8645A, and HP 8657B signal generators. He also worked on the HP 8657 (Option 22) Groupe Special Mobile (GSM) European digital cellular system, developed the HP 11835 data buffer, and contributed to development of the HP 11847A software package. Dave is named as an inventor on a patent for a synthesis technique used in the HP 8645A signal generator, and on a patent pending for the HP 11846A generator's filter technique. Dave's professional interests include RF communications and digital signal processing. He received his BSEE degree in 1980 and his MSEE degree in 1981 from the University of Washington. Born in Seattle, Washington,

he lives in Spokane, is married, and has one son. Dave enjoys downhill and cross-country skiing, scuba diving, backpacking, golf, carpentry, and working with computers.

Raymond A. Birgenheier

Ray Birgenheier has been a consultant and development engineer in digital signal processing and digital communications at HP's Spokane Division since 1981. He is named as an inventor of the signal processing techniques used in the HP 11836A and HP 11847A products. He developed the software for the HP 11847A measurement systems that verify the RF performance of digital cellular transmitters. He received a BSEE degree (1963) from Montana State University, an MSEE degree (1965) from the University of Southern California, and a PhD degree (1972) in electrical engineering from the University of California at Los Angeles. Ray worked for Hughes Aircraft Company's Radar Systems Division from 1963 to 1980, where he became a senior scientist in 1976. Since 1980, he has served as a professor and chairman of the Department of Electrical Engineering at Gonzaga University. He has authored numerous articles on adaptive antennas and signal processing techniques, and is named as an inventor in three patents on adaptive antennas, one on premodulation filters, and on three pending patents on premodulation filters, modulation measurement techniques, and apparatus. Ray is a member of the IEEE societies on communications systems and engineering education. His professional interests also include communication theory and systems, and digital processing of random signals. A U.S. Navy veteran, Ray was born in Billings, Montana. He is married, lives in Spokane, Washington, and has seven children and five grandchildren. He is active in his church and enjoys hiking, hunting, and fishing.

David L. Neuder

As the principal software designer for the HP Branch Validator, Dave Neuder developed an automated tool that enables software designers to test and verify the branch coverage of modules. He has also worked on development of HP Teamwork and the HP 64610 timing analyzer. Dave received a BSEE degree (1977) and an MSEE degree (1979) from Michigan State University, and joined HP's Logic Systems Division in 1979. He is named as an inventor on a patent on marking technology in timing analyzers, and is a member of the IEEE. Born in Wyandotte, Michigan, Dave lives in Colorado Springs, is married, and enjoys singing in choirs, skiing, backpacking, cycling, and climbing Colorado's 14,000-foot peaks.

```
FR: DICK DOLAN/97LDC          00053803
                                   446
TO: LEWIS, KAREN
    CORPORATE HEADQUARTERS
    DDIV  0000   20BR
```

**HEWLETT
PACKARD**

5091-1154E