Contents:

In this Issue:

If our cover photograph tells you that this month's featured instrument tests the kinds of transmitters that are used by fire, police, and airline personnel, we've been successful. Model 8901A Modulation Analyzer is a new kind of analyzer that makes it possible to do a better job of testing AM and FM transmitters, signal generators, crystal oscillators, and many other devices that operate in the radio frequency range, specifically 150 kilohertz to 1300 megahertz.

What makes this new analyzer different from other radio frequency analyzers is that it is designed especially for measuring the characteristics of *modulated* signals—like those that carry voice communications. It measures amplitude modulation, frequency modulation, phase modulation, or combinations of AM and FM or ΦM. An important feature is its ability to measure very small amounts of one kind of modulation in the presence of very large amounts of another, which means it can help discover unintentional AM on an FM signal or vice versa.

The new modulation analyzer measures modulation much more accurately than anything that's been available before—so accurately that it was necessary to develop new standards to compare it with (page 22) and a special signal generator to test it (page 30). A built-in microcomputer makes it nearly automatic (page 26). For production testing, metrology, research, and development, this new concept in RF analyzers will solve many long-standing problems.

-R. P. Dolan

Precise, Convenient Analysis of Modulated Signals

This new modulation analyzer measures a signal's frequency, power, and modulation components with unprecedented accuracy in many measurements. It adds up to a new concept in RF analyzers.

by Allen P. Edwards

HEWLETT-PACKARD'S new Model 8901A Modulation Analyzer, Fig. 1, is a complete measurement system for accurately characterizing signals in the 150-kilohertz-to-1300-megahertz frequency range. The name modulation analyzer was chosen instead of modulation meter because the 8901A does so much more than just meter the modulation. Combining the capabilities of three separate instruments, the new modulation analyzer can measure carrier frequency and RF peak power (often eliminating the need for a separate power meter), and it can accurately measure modulation and recover the modulating signal. It makes those measurements most commonly needed to characterize a signal completely, including a signal's frequency, frequency drift, peak power level, amplitude modulation (AM), frequency modulation (FM), phase modulation (ΦM), and AM and FM noise components. It recovers the modulating signal with very low added distortion for further analysis. The article on page 19 describes some of its uses.

The 8901A measures the amount of modulation on a signal whether or not the type of modulation being measured is the dominant modulation type on the signal. This means, for example, that one can make accurate FM measurements on an AM signal, and FM measurements are now possible on a signal that is not being intentionally modulated at all. Since the 8901A's internally generated FM noise is typically less than 1 Hz (rms deviation in a 50-Hz-to-3-kHz bandwidth), direct noise measurements can be made, for example, on a high-quality 500-MHz transmitter or signal generator. The analyzer's low noise level also allows accurate measurements of low-peak-deviation modulation.

Another modulation analyzer benefit is its ability to measure a signal's carrier frequency directly and accurately.

Fig. 1. *Model 8901A Modulation Analyzer measures a signal's frequency, drift, power, AM, FM, and residual noise. It also recovers the modulating signal. Its carrier frequency range is 150 kHz to 1300 MHz and its modulation depth or deviation measurements are generally accurate within one percent of reading.*
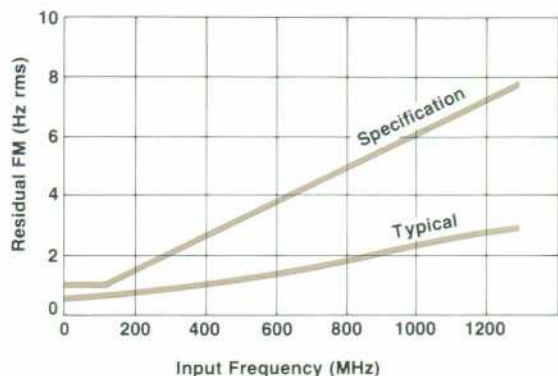
Fig. 2. *8901A Modulation Analyzer's performance exceeds that of nearly all signal generators. Residual FM in a 50-Hz-to-3-kHz bandwidth is typically less than 1 Hz up to 500 MHz, as shown here.*

Signals can be counted in the presence of combinations of AM and FM that might confuse a conventional counter. There is also an indication of carrier power with sufficient accuracy for many applications.

But the major contribution of the 8901A to signal analysis is that it makes extremely precise modulation measurements quickly and easily. Modulation depth or deviation accuracy is generally better than one percent of reading. Specified residual noise in a 50-Hz-to-3-kHz bandwidth is less than 0.01 percent for AM and less than 8 Hz for FM at 1300 MHz carrier frequencies, decreasing to less than 1 Hz below 100 MHz (Fig. 2). As an example of how easy the analyzer is to use, assume that you have a signal and that

you have no knowledge of its frequency or power or modulation content. Applying the signal to the 8901A almost instantly causes a change in the display from two dashes to four dashes, indicating that an input has been sensed. Approximately one second after the signal is applied, a settled reading appears on the display, indicating the measured value of the quantity that was selected. After that, about as fast as the buttons can be pressed, measurements of AM, FM, analog phase modulation, frequency, or power can be made. At no time does the user have to tune the instrument or set up input attenuators. Although the instrument can be controlled entirely manually if desired, it is much faster to let the modulation analyzer adjust itself.

The accuracy claimed by the modulation analyzer is unprecedented. To achieve this accuracy, standards had to be developed beyond what existed. Two self-calibrating modulation standards were developed, one for FM and one for AM, and are available as options. The article on page 22 describes the theory and implementation of these calibrators.

It was also necessary to develop a high-quality signal source to test the modulation analyzer, since the analyzer's performance exceeds that of nearly all signal generators. This special-purpose signal source is described in the article on page 30.

Modulation analyzer operations can be controlled and all measurement results can be transferred via the Hewlett-Packard Interface Bus (HP-IB), Hewlett-Packard's implementation of ANSI/IEEE Standard 488-1978.

Instrument Control Philosophy

The front-panel control philosophy of the new modula-
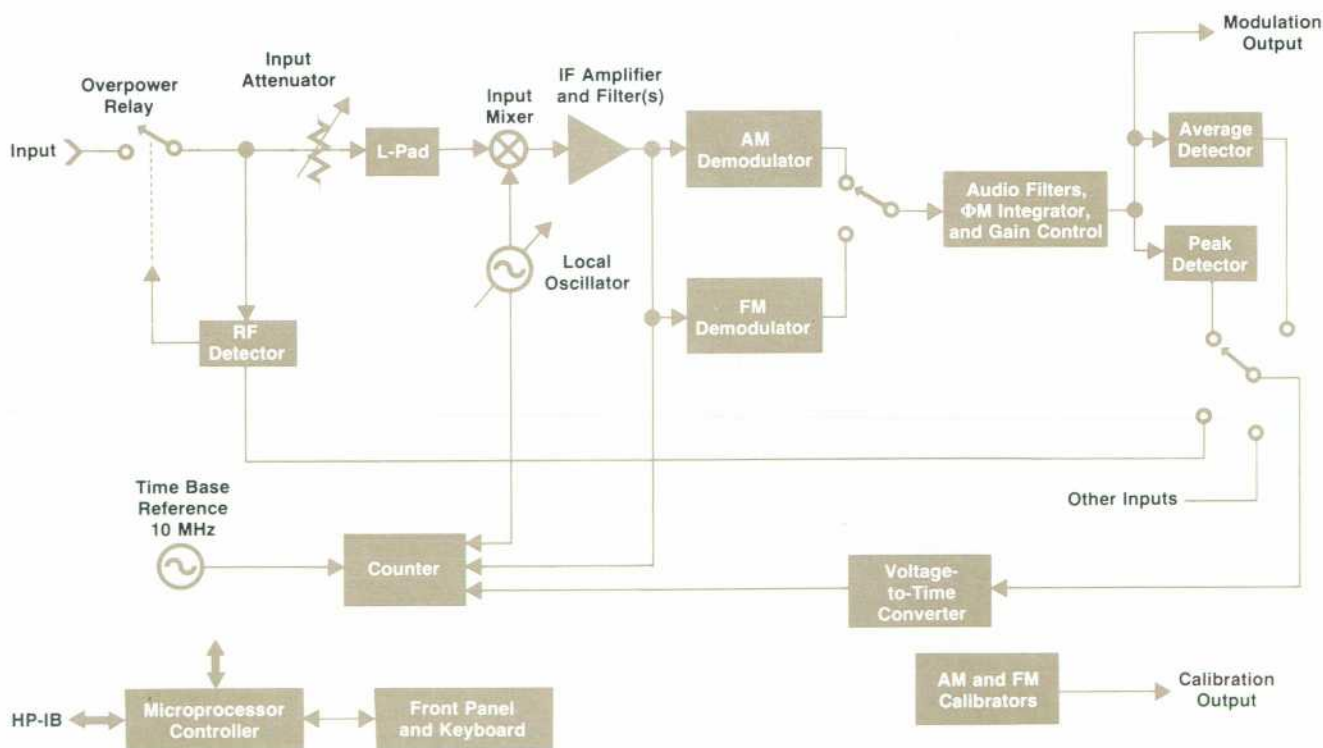
Fig. 3. *Simplified block diagram of Model 8901A Modulation Analyzer. Basically a calibrated receiver, the analyzer down-converts the input signal and then recovers either the amplitude or frequency modulation regardless of which is the dominant type.*

tion analyzer combines the simple with the powerful. The panel is intentionally kept simple and can be understood with almost no training. Even the most casual user can make all of the normal measurements with the maximum specified accuracy. Usually a single keystroke is all that is necessary to select a new measurement. The user can select the bandwidth of the measurement using audio filters, and can select whether to read the deviation as peak or average. On the right side of the front panel is a block of keys labeled MANUAL OPERATION. As long as the automatic mode of operation has been selected, the user never has to use this keypad.

This simplicity of operation has been achieved without sacrificing total user flexibility. From the front panel the more advanced user can override all automatic settings or select measurements not available in automatic mode. The most obvious thing that can be done is to enter the input frequency desired directly. Once entered, the frequency can be changed either by entering a new one or by stepping up or down a settable number of kilohertz.

Other functions are available through the special function key, labeled SPCL. Special functions are numbered 1 to 39 for user special functions and 40 to 99 and 0 for service special functions. The first eight special functions represent instrument settings, such as input attenuation, IF frequency, peak detector response time, and so on. To program a special function the function number is keyed in, followed by a decimal point and then the desired value relating to that function. For example, to command the input attenuator to 10 dB attenuation the user keys in 1.2 and presses the SPCL key. All of this is explained on a pullout card. Pushing the SPCL key without entering a number causes the display to show the programmed status of the first eight special functions. Pressing the SPCL button twice without entering a number causes the display to show the actual settings of the first eight functions. Thus, pushing SPCL once may show that a function is set to autorange (default), and pressing SPCL again will show what range was selected by the controller.

A Calibrated Receiver

The block diagram of the 8901A Modulation Analyzer is shown in Fig. 3. The diagram is that of a highly calibrated receiver. After being attenuated to the desired drive level, the incoming signal is down-converted by a low-residual-noise local oscillator. Down-conversion was chosen to convert the incoming signal to an intermediate frequency (IF) range where low-noise and high-linearity modulation detectors can be built.

The IF signal is amplified and fed to AM and FM detectors, either of which may be selected for further processing by the audio section and presentation to the modulation output connector and to the audio detectors. Both peak and average responding audio detectors are provided. The detector outputs are digitized and displayed when called for.

The requirement for continuous coverage of the RF range all the way down to 150 kHz along with the ability to handle FM rates of 200 kHz and deviations of 400 kHz at other input frequencies is responsible for a major feature of the block diagram. The incoming frequencies are divided into three regions. Frequencies above 10 MHz are converted down to 1.5 MHz. The IF at 1.5 MHz is broadband, about 3 MHz

Frequency Range? Which One?

So many of the terms describing a modulated signal are in units of frequency (hertz) that the 8901A Modulation Analyzer actually has many operating frequency ranges. The input carrier frequency operating range is from 150 kHz to 1300 MHz. This carrier can be either amplitude or frequency modulated. For an FM signal the *amount* of modulation can be between 0 and 400 kHz peak deviation, which means that the instantaneous frequency can go higher or lower than the average carrier frequency by up to 400 kHz. The *rate* of modulation, which is the frequency of the audio tone modulating the carrier, can be between 20 Hz and 200 kHz. The specifications for phase modulated (ΦM) signals are derived from the FM specifications by observing that the amount of phase deviation in radians is the ratio of the frequency deviation to the modulation rate. In the case of amplitude modulation the *amount* of modulation, which is also called the modulation *depth* or modulation *index*, can be between 0 and 99%. The *rate* of modulation can be between 20 Hz and 100 kHz.

8901A accuracy is specified as percent of reading instead of percent modulation. For example, the basic uncertainty in reading a 30% AM signal is specified as 1% of reading, or ±0.3% AM in this example.

wide. It is amplitude and phase compensated for good incidental AM and low FM distortion (see "IF Filters," page 10). Above 2.5 MHz the input is down-converted to 455 kHz. The AM and FM detectors are broadband and work well at IF or carrier frequencies between 150 kHz and 2.5 MHz, so this change in IF is not a problem. The 455-kHz IF is about 200 kHz wide (455 kHz ±100 kHz), broad enough to handle voice-grade bandwidths. This IF offers a higher degree of selectability; in fact, the 8901A can be set up to receive the commercial FM band (useful when working late in the lab). In the region between 150 kHz and 2.5 MHz no down-conversion takes place, and detection is done at the carrier frequency.

The input uses a single-balanced mixer. This type of mixer allows the incoming signal to pass through into the IF without down-conversion, a useful feature of a single-balanced mixer. Other advantages are the mixer's low distortion and high-level operation, both extremely important for low-distortion AM operation.

When the 1.5-MHz IF is being used, a 4-MHz low-pass filter is switched into the system just past the mixer and in front of the IF amplifier. This filter is relatively insensitive to the source impedance variations that are present at the mixer's output, and it improves the conversion loss of the mixer by matching the IF port. This filter also prevents the high-level RF from causing distortion in the IF amplifier. A 3-MHz low-pass filter is placed after the IF amplifier to define the 1.5-MHz IF passband and to eliminate the high-frequency noise contribution of the IF amplifier. This filter is phase compensated for an overall flat group delay, and is isolated from mixer impedance variations by the IF amplifier. This filter arrangement helps maintain the low-distortion and low-residual-noise performance over the multi-decade input frequency range.

When the 455-kHz IF is being used, the 455-kHz bandpass filter that defines this IF replaces the 4-MHz low-pass filter ahead of the IF amplifier.

Since a single-balanced mixer has no rejection of signals at the RF port, there could be a problem if signals at the

intermediate frequency are presented to the input along with the desired signal. This can happen if AM broadcast stations are picked up in the device under test. To alleviate this problem a pre-mixer high-pass filter is available to the user; it rejects the offending frequencies when operating with the wide IF.

Overload Protection

Another requirement that influences the 8901A's block diagram is the ability to withstand large power overloads. The design goal called for protection from signals as large as 100 watts from a 50Ω source. The actual damage level is a function of source SWR (standing wave ratio) so that the actual specification is written as 25 watts from a 4:1 mismatch. Under normal operation the instrument can make measurements on signals up to 1 watt. The input attenuators had to be designed to absorb over 1 watt continuously without damage and up to 100 watts for short periods of time.

An RF relay in front of the input attenuators opens in the event of an overpower condition, thus protecting the instrument from damage. The input power is detected by the RF level detector, which is designed for safe operation with over 100 volts applied. For the short period of time until the relay opens, the mixer is protected by a pair of clamp diodes. The clamping diodes come after an impedance transformation L-pad that converts the 50Ω system to 25Ω in anticipation of the mixer's impedance-doubling characteristic. The L-pad isolates the clamp that protects the mixer from input overpower. Without this isolation the clamp could limit the signal applied to the detector and the overpower condition would not be detected. Also, at 25Ω, the parasitic reactances of the clamp are less significant. All of the components of this RF power and overpower circuit are matched to 50Ω by modifying the impedance of the transmission lines to form a low-pass filter structure for each component. This allows operation at 1300 MHz without affecting the input SWR specification.

The input power detector has a tendency to load the line and cause inaccuracies at high RF frequencies. To eliminate this, the RF level detector is biased off the line when the analyzer is being used as a receiver. It is only used when needed, such as during attenuator setting or during RF level measurements. The bias level is set below the overpower trip point so that overpower will still be sensed and cause the input relay to open.

Automatic Operation

The goal of making the analyzer autoranging, autotuning, and generally easy to use has also had a major impact on the block diagram.

The ease-of-use and autoranging requirements dictated a system under the control of an intelligent controller. The 8901A is controlled by a microprocessor with 18K bytes of program memory. Of course, the decisions the controller makes are only as good as the information fed to it. Therefore, many things are monitored in the instrument in the course of its operation. Decisions as to the proper operating point, for example, are based on input RF level and frequency, peak IF level, average level out of the AM detector, and AGC control current. Depending on the instrument

setup, invalid levels at these detectors or invalid relationships between levels will cause autoranging, retuning, or user warnings.

Many of these detectors are measurement detectors, but this does not mean that they are examined only during specific measurements. Most detectors are examined during every measurement cycle to insure proper setup and valid readings. Some detectors are installed primarily to insure proper setup, but this doesn't mean that they can't be used to take readings. One example of this is the IF level detector used primarily for autotuning. This detector measures the peak level in the IF. The level at this detector is measured and a second-order curve fit is made to linearize it. Then a frequency response calibration factor is applied, and finally, by adjusting for input attenuator setting, the power in the IF referred to the input can be displayed. This is called tuned RF level, and represents a secondary measurement mode. It is quite useful and can be used to measure RF harmonics within about −50 dBc to an accuracy of a few dB. It can also be used as a high-sensitivity power indicator with maximum resolution of 1 nW. Another example is the average level out of the AM detector, which is displayed as IF level. Normally the AGC keeps this at 100%, but with the AGC turned off, this can also be a useful tuned detector, although it is most useful in manual setup to indicate the presence of an IF signal.

Another source of inputs for the controller is the counter. The counter is implemented partly in hardware and partly in the microprocessor. The microprocessor determines the input and count time and does the appropriate scaling. The counter has eight inputs, but the most important two during normal operation are the local oscillator frequency (LO) and the intermediate frequency (IF).

The controller monitors the IF to insure that the instrument is properly set up, and if it is not, causes the appropriate response to take place. Also, it is the combination of knowing the LO frequency and the IF that allows the input frequency to be known. The relationship that applies is $IF = n(LO) \pm m(RF)$, where RF is the input frequency. An elaborate procedure of searching that avoids harmonic mixing and checks for input harmonics and images gives considerable confidence that $m = 1$ and $n = 1$, so that the RF frequency
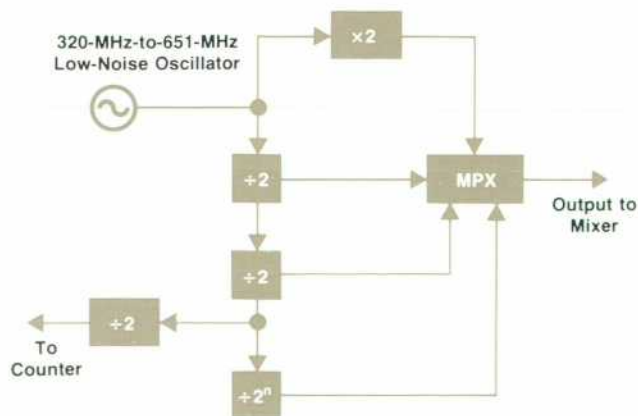
Fig. 4. Low-residual-FM local oscillator generates a signal between 1.25 MHz and 1302 MHz by dividing or multiplying the output of a low-noise octave oscillator.

is in fact (LO−IF).

To make the 8901A both autoranging and controllable, all of the internal switching is under software control. The front-panel readouts and annunciators are also under software control. The keyboard provides input requests to the controller, which examines each request and decides what response is appropriate.

The I/O system for the microprocessor is organized as a nibble serial system. Each four-bit nibble is addressed by four address lines and one of eight clock lines to create a 512-bit I/O space. Serializing the I/O greatly reduced the number of interconnections required, while sending four bits in parallel allowed available MSI to be used efficiently so that remote latching was practical. The latches for the functions are located remotely from the microprocessor and near the functions being controlled. This allows the I/O bus to be contained on a 16-line ribbon cable that connects the various functional modules. Within each module, ten lines of output and one to four lines of input represent the entire I/O bus. This made it possible to have all the required interconnections on one side of a motherboard, and a ground plane for shielding on the other side. Thus it was possible to build the 8901A without any costly multilayer printed circuit boards.

Local Oscillator Design

The principal design goal for the local oscillator used in the modulation analyzer was low residual FM. Specifically we wanted a local oscillator that could down-convert a high-quality narrow-band FM signal without degrading the signal-to-noise ratio. After discussions with users, we established a design goal of 1 Hz total FM noise (70 dB below 3 kHz deviation) in a 50-Hz-to-3-kHz demodulation bandwidth at a carrier frequency of 500 MHz. This level of performance exceeds that of many frequency synthesizers. We rejected the conventional solution of a separate crystal for every desired channel and set out to synthesize a signal economically. The approach that was eventually selected meets our design goal, yet does so at a very low cost, thus representing good value to our customers. The cost of the local oscillator is only a fraction of the cost of a conventional synthesizer.

Two important considerations differentiate a local oscillator application from a signal generator. Since we are driving a mixer, which is rich in harmonic generation, harmonics out of the local oscillator are not a problem. In fact, we found that square waves represented a superior signal for driving our mixer. Also, since we only wanted to generate a CW signal, we were free to consider designs that did not lend themselves to being modulated.

Fig. 4 is a block diagram of the local oscillator. The LO can generate a signal anywhere within the range of 1.25 MHz to 1302 MHz. All outputs start with an octave oscillator operating between 320 and 651 MHz. That fundamental octave is either divided down to the desired frequency or, in the case of the top band, multiplied up. The appropriate output is selected by a multiplexer and presented to the mixer. Because we use square waves, much of our multiplexing is done with ECL (emitter-coupled logic), with the high-speed and final multiplexing done with high-speed differential pair transistors manufactured by HP.

The doubler is a custom thin-film microcircuit packaged in a TO-8 package. The doubled output is multiplexed into the mixer drive with PIN diodes to maintain the required bandwidth.

The first two divide-by-two circuits divide the 320-to-651-MHz signal down to 80-162.75 MHz. These dividers, which are also manufactured by HP, are always left on, providing a signal to the microprocessor-based counter through an ECL divide by-two circuit. Thus the microprocessor can always count the local oscillator. This information is used in tuning the local oscillator to its desired frequency, and when determining the frequency of the input signal. The other dividers are gated off when not in use to minimize problems caused by subharmonics.

Octave Oscillator

Before developing the block diagram for the octave oscillator, we translated our design goal of 1 Hz residual FM to more familiar units. The amount of residual FM within a band can be expressed as:[1]

$$\Delta f(res) = \sqrt{2} \int_a^b \Lambda(f_m) \; f_m^2 df_m$$

where $\Lambda(f_m)$ is defined as:[2]

$$\left[\frac{\text{Single-Sideband Noise Power in a 1-Hz Bandwidth,} \; f_m \text{ away from the Carrier}}{\text{Signal Power}} \right]$$
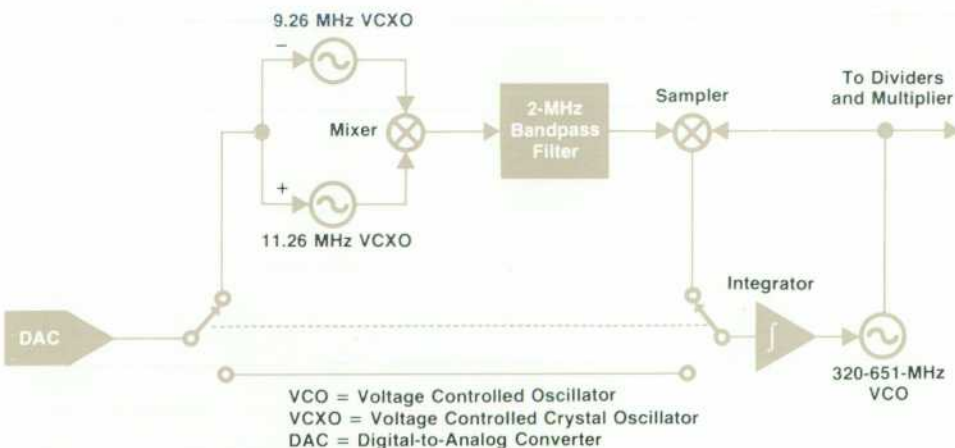

Fig. 5. Octave oscillator is phase-locked to a harmonic of a tunable crystal reference oscillator to stabilize the 320-651 MHz output. The VCO is tuned approximately by the DAC, then locked to the VCXO.

VCO = Voltage Controlled Oscillator
VCXO = Voltage Controlled Crystal Oscillator
DAC = Digital-to-Analog Converter

Solving this for various noise shapes produced an interesting result. We found that if the phase noise in a 1-Hz bandwidth 1 kHz away from the carrier was −100 dBc, then we were very close to our 1-Hz residual FM goal. This was true whether the noise was flat or falling off at 9 dB/octave in the band of interest (flat yields 1.34 Hz, 6 dB/octave yields 0.77 Hz and 9 dB/octave yields 0.90 Hz). This gave us a convenient guideline by which to evaluate designs.

A very important factor in local oscillator design is the requirement to have continuous frequency coverage. It would be simple to lock the output to a crystal oscillator and achieve good noise performance, but that would produce only one frequency per octave. If the output is locked to a lower-frequency reference and the harmonic of lock is selectable, we can get one frequency per harmonic number, but still do not have continuous coverage. To fill in the gaps created by harmonic mixing there are basically two options. One possibility is to add in another signal that can be tuned over the required amount and fill in the difference between adjacent harmonics. This signal must also have low noise, but this is a simpler task because a much lower frequency can be used. This approach requires two high-frequency oscillators and two phase-locked loops (one phase-locked loop is required as an output tracking loop), and was rejected for cost reasons.

The other possibility, the one we use to get continuous frequency coverage, is to use a variable reference frequency. Because the noise of the reference is multiplied by the harmonic number, the reference must have very low noise. The condition of continuous coverage states that the reference must be able to pull the output frequency at least as far as the next frequency available without pulling. We can express this requirement as:

$$f_{out} = nf_{ref} + n\Delta f_{ref} = (n+1)f_{ref}.$$

Rearranging this gives:

$$\Delta f_{ref} = f_{ref}/n = f_{out}/(n^2+n).$$

Thus for a given f_{out}, the choice of a low f_{ref} and a large n reduces the requirements on Δf_{ref} significantly.

Assume, for example, that f_{out} = 300 MHz and f_{ref} = 2 MHz (n=150). Then Δf_{ref} (max) = $300 \times 10^6 \div (150^2+150)$ = 13.2 kHz = ±6.6 kHz. For n=250 (f_{out} = 500 MHz) the noise requirement for the reference is 48 dB better than the requirement for the output. Allowing something for sampling noise, the reference noise must be <−150 dBc at 1 kHz.

Clearly this is an application for a crystal oscillator. However, to do this with a 2-MHz crystal oscillator would require ±3000 ppm pulling, and this cannot be done and maintain the noise properties of the crystal. Instead, we took two higher-frequency crystals around 10 MHz and mixed them together to form a 2-MHz reference. Since the crystal frequencies are now five times higher than the reference and there are two VCXOs (voltage controlled crystal oscillators) that can be pulled against each other, the pulling requirement on each VCXO is only ±300 ppm. This is a formidable amount but one that was achieved.

Fig. 5 is a diagram of the local oscillator for the 8901A. The 9.26-MHz and 11.26-MHz VCXOs are mixed to produce a 2-MHz reference frequency. The reference is mixed in a sampler with the output of the high-frequency VCO, which is then phase-locked to the reference. Thus we have met our design goal with a single-loop design, which represents a considerable savings in complexity and cost.

The good residual FM performance can only be maintained if the spurious mixing products on the output are kept to −80 dBc. This translates to −130 dBc at the reference signal for a multiplication of 50 dB. With 9.26 MHz and 11.26 MHz for the two references, the lowest-order spurious product to cross the output signal is 82nd order. The next closest spurious output is at 130 kHz offset and is 51st order. With some care we were able to keep the actual spurious levels below detectability.

Care has also been taken to reduce the microphonic susceptibility of the local oscillator by shock mounting it. In an informal test, we have had the 8901A operating next to a large punch press without adverse effect.

Tuning the Oscillator

The first step in tuning the local oscillator to a desired frequency is selection of the octave of interest. Once the desired octave is determined and output to the multiplexer, the job of getting the fundamental oscillator to the required frequency begins. Before the phase-locked loop is closed, the VCO is tuned to the desired frequency under the control of a digital-to-analog converter (DAC). This positions the output frequency approximately on the required harmonic of the reference. The positioning is accomplished by counting the output frequency with the microprocessor-based counter and adjusting the control word to the DAC until the input to the counter is within tolerance.

The DAC control is arranged so that the phase-locked-loop integrator is charged to the correct operating voltage. The DAC can now be disconnected without the VCO changing frequency. This is done, and the DAC is then connected to the tune line of the reference VCXOs and set to a value that causes the reference frequency to be 10% above its lowest value. Finally, the phase-locked loop is closed.

At this point there are two possibilities. If the loop is within its capture range (about 200 kHz), the loop will pull in and acquire lock. More often than not, however, the loop will be outside the capture range, so an offset is injected into the loop so that the resulting error signal will tune the VCO down in frequency until lock is achieved with the closest lower harmonic. When lock is achieved, the output is somewhere between 200 kHz above the desired frequency and 1.8 MHz below the desired frequency. The DAC controlling the reference frequency, which was set for 10%
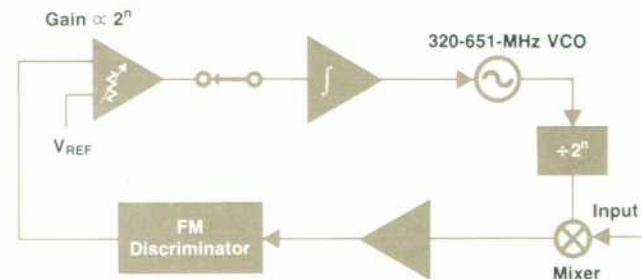

Fig. 6. *To track an input signal, the octave oscillator is phase-locked to the input signal by the track loop shown here.*

more than the lowest reference frequency, can then be used to tune the VCXOs until the desired output frequency is obtained.

Examining the numbers involved, it can be seen that the reference can be moved enough to cover the space between n numbers (2 MHz). Because of its initial setting, the DAC can move the output down 10% (at least 200 kHz) or up 90% (at least 1.8 MHz). This is typically exactly what is required. However, there are conditions where the initial locking to the VCXO harmonic might be off by one n number. If the desired output frequency is not obtained, then a special routine is performed that steps the phase-locked loop one harmonic number and another attempt to fine-tune the VCXOs is made. This process may be repeated as often as necessary, so the correct harmonic number can be obtained regardless of how far off the initial lock was. This represents a certain amount of redundancy in tuning and a certain tolerance to some types of hardware failure.

Tracking an Input Signal

Another mode for setting the output frequency of the local oscillator is called track mode. In track mode the reference that is used to set the local oscillator output frequency is the frequency of the RF input. The circuit is similar to an AFC (automatic frequency control) circuit except that a signal can be followed continuously anywhere between 10 MHz and 1300 MHz. The lock loop for track mode uses the input mixer to generate an intermediate frequency that is converted by the FM discriminator to a voltage proportional to the IF frequency. The discriminator output is compared to the desired voltage and used as an error voltage to tune the VCO (see Fig. 6). The LO signal is fed to the mixer through a divide-by-2^n circuit to generate an output within the desired octave. The rate of change of the local oscillator output frequency as a function of the VCO tuning voltage is a function of the division factor. This is compensated within the loop by a resistive network that changes the input resistor to the loop integrator by powers of two. With constant loop gain, the loop bandwidth can be kept constant. There is some variation in loop bandwidth due to changes in the sensitivity of the VCO over any octave, but this has been kept to a minimum by design. The loop bandwidth is optimized to produce less than 1% error for 1-kHz FM rates. Higher rates are less affected, of course, and very low FM rates are removed by the loop.

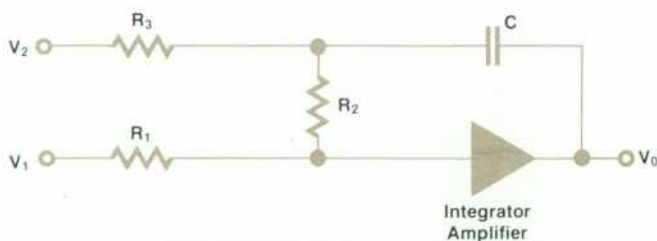When the input moves in frequency enough that the VCO approaches one of its octave band edges, the microprocessor takes over. The track loop is opened and the DAC is connected to the VCO. The local oscillator frequency is adjusted to the same point but in the next octave. Then the DAC is disconnected and the track loop closed. This process takes only about 50 milliseconds so that lock appears unbroken for slowly moving signals such as that generated by manually tuning a signal generator. The band was made slightly larger than an octave to allow for some hysteresis.

It is interesting to note that two different types of inputs are required of the loop integrator. In the case of the phase-locked loop (Fig. 5), the integrator pole must be removed with a zero before loop crossover for a stable system, because the phase detector (sampler) output voltage is proportional to phase difference and phase is the integral of frequency. This gives the loop a double-integrator response. The second integrator pole is removed by putting a resistor in series with the integrator capacitor. For the track loop (Fig. 6), a true integrator response is desired because frequency discrimination rather than phase discrimination is used. The requirement for a true integrator response is even more stringent for the sweep circuit that is discussed next. Examination of Fig. 7 and the following equations will show how these two requirements were met:

Lock Loop: $V_0/V_1 = -(1+sR_2C)/(sR_1C)$
Track Loop: $V_0/V_2 = -1/(sR_3C)$.

Finding an Input Signal

The modulation analyzer is capable of finding a signal that is applied to its input. The algorithm that accomplishes this is complex (see article, page 26). It is capable of pre-

VCO Overshoot = Δf
Loop Turnoff Delay = D
For Compensation, $\Delta f = DI/C = IR$
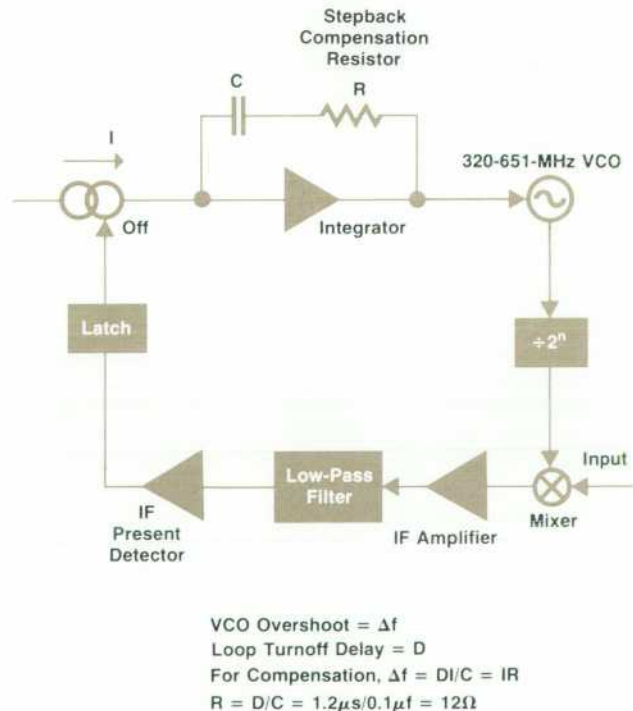$R = D/C = 1.2\mu s/0.1\mu f = 12\Omega$

Fig. 8. Sweep loop, controlled by the microprocessor, tunes the octave oscillator rapidly across its band, looking for an input signal. When an IF signal is detected, the sweep stops. The stepback resistor compensates for the oscillator's overshooting the input frequency during the time it takes to stop.

Fig. 7. Loop integrator has two types of inputs, one for the lock loop (Fig. 5) and one for the track loop (Fig. 6).

IF Filters for the 8901A Modulation Analyzer

by Andrew H. Naegeli

The 8901A Modulation Analyzer has two alternate intermediate frequencies. The primary IF is 1.5 MHz, chosen to accommodate an FM signal with 400-kHz deviation at a 200-kHz rate. The secondary IF is 455 kHz, used to avoid spurious conversion problems in the mixer for input frequencies below 10 MHz. This IF is limited to rates of 10 kHz and deviations of 40 kHz.

Low-Pass Filter and Phase Compensation

The primary IF uses a low-pass filter with a bandwidth of 2.5 MHz. The low-pass filter is used because of the very wide bandwidth requirements of the worst-case FM signal described above. The filter must reject IF signals of 6 MHz and higher by at least 60 dB, and must have a very flat amplitude response to preserve the FM-rejection performance of the AM detector. These requirements led to a 6-pole Cauer-Chebyshev filter design with two zeros of transmission at 6.1 and 8.4 MHz. This filter was synthesized as a passive network using standard inductors and capacitors as the filter elements.

To prevent distortion of FM signals, a phase compensation filter is cascaded with the 2.5-MHz low-pass filter. For an FM signal with 100-kHz deviation at a 100-kHz rate, the phase compensation filter reduces the distortion from 2% to less than 0.1%. This filter is an all-pass network with four poles and four zeros. It is realized as an active filter, as shown in Fig. 1.
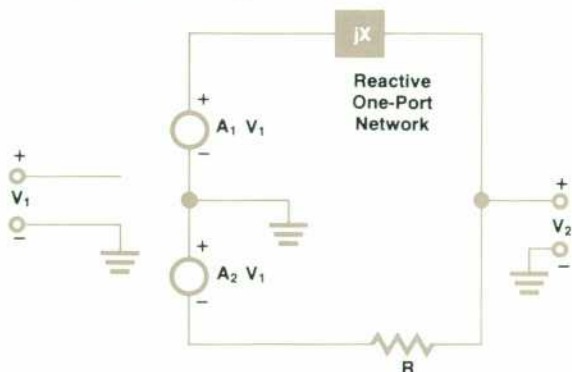

Fig. 1. All-pass network acts as a phase compensation filter to prevent distortion of FM signals.

The two voltage sources are models for two amplifiers that deliver the input signal to two points, 180 degrees out of phase. Assume for the moment that A1 = A2 = 1, that is, both channels have the same unity gain. The circuit may be analyzed using superposition. First, with the lower voltage source shorted, the output voltage is:

$$V1 \frac{R}{R + jX} \tag{1}$$

Next the contribution from the lower voltage source is found by shorting the upper voltage source:

$$-V1 \frac{jX}{R + jX} \tag{2}$$

The sum of (1) and (2) is the output voltage, V2. Therefore, the transfer function is:

$$\frac{V2}{V1} = \frac{R - jX}{R + jX} \tag{3}$$

This transfer function is an all-pass function because the zeros of transmission (the roots of the numerator) are a mirror image of the poles (the roots of the denominator) across the frequency axis of the s-plane.

The amplitude and phase responses of this network must be adjusted to give optimum performance. Component variations in the reactive elements can change the locations of the poles and zeros, leading to phase changes. The phase response is adjusted to the desired shape by adjusting the value of R.

The amplitude response of the network is theoretically flat, but losses in the reactive elements can affect the amplitude flatness. For example, a small resistance, r, in series with the reactive network changes the transfer function to:

$$\frac{V1}{V2} = \frac{(R - r) - jX}{(R + r) + jX} \tag{4}$$

Now, if the reactance term, jX, is zero at some frequency, the amplitude is no longer = 1, but is less than 1. This causes ripples that depend on the location of the poles and zeros of the reactance, jX. The reactance network for this filter has a zero at 2.2 MHz and poles at 1.0 and 4.3 MHz. These can cause substantial ripples in the amplitude response. To adjust for this problem, the two gains are adjusted such that A1 is not equal to A2. This changes the amplitude response primarily, so the two adjustments are only marginally interactive.

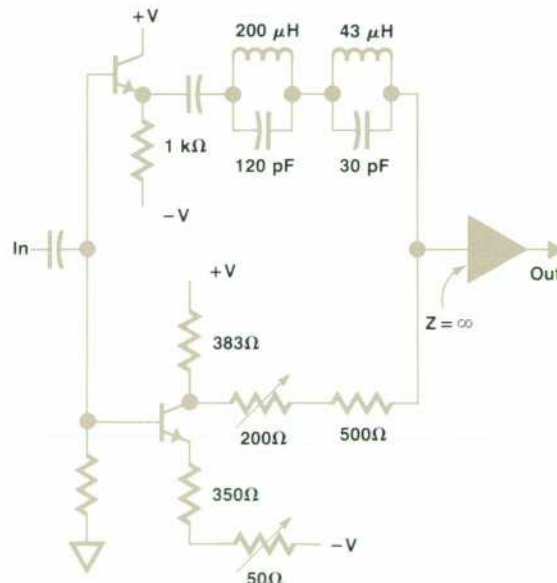The transistor circuit realization of this filter is shown in Fig. 2.

Fig. 2. All-pass network is a four-pole, four-zero active filter.

Bandpass Filters

The secondary IF uses a bandpass filter at a 455-kHz center frequency. The 3-dB bandwidth of this filter is 200 kHz. The wide bandwidth is necessary because of the wide deviation of permissible FM signals. The phase response of this filter is very important, because of FM distortion. The filter was designed to add less than 0.03% distortion to an FM signal with 10-kHz deviation at a 10-kHz rate. To

accomplish this, the phase response was made very linear, while sacrificing a sharp cutoff in the amplitude response. This required the filter to be designed as an arithmetically symmetrical bandpass filter (same response at $f_0 - \Delta f$ as at $f_0 + \Delta f$, where f_0 is center frequency), as opposed to the standard geometrically symmetrical design (same response at $\frac{1}{2}f_0$ as at $2f_0$).

In standard bandpass filter design, a low-pass prototype filter is transformed geometrically up to the desired center frequency. Unfortunately, this transformation distorts the phase response in the filter passband. To avoid this problem, the desired bandpass function was approximated directly by analyzing the response of pole-zero patterns using a computer. A transitional Butterworth-Thompson design was used as the low-pass prototype, to make the group delay response flat across most of the passband. When the desired constellation of poles and zeros was achieved, the seven-pole bandpass filter was synthesized directly using standard insertion loss techniques.[1]

The filter bandwidth is wide enough so that standard inductors and capacitors are used as the filter elements. The amplitude and phase responses are adjusted using two slightly interactive adjustments, adjusting the value of two inductors. The variable elements were chosen from the results of a sensitivity analysis performed with a computer-aided-design program.

The secondary IF filter may also be used at input frequencies above 10 MHz to make the 8901A more frequency-selective. This can be useful when measuring the level of signals at various frequencies in the input spectrum, using the TUNED RF LEVEL function.

References
1. E.A. Guillemin, "Synthesis of Passive Networks," Wiley, 1957.

Andrew H. Naegeli
Born in Arcadia, California, Andy Naegeli received his BSEE degree from Stanford University in 1975 and his MSEE in 1979, also from Stanford. An HP employee since 1975, Andy worked on the RF input circuits, IF filters and microprocessor software for the 8901A Modulation Analyzer. He is currently a project manager with the HP Stanford Park Division. Married to a professional musician and living in Menlo Park, Andy spends much of his spare time fixing up his 30-year-old house. A reborn Christian, he plays string bass with a folk singing group at his church. He also works with stained glass and enjoys camping, water-skiing and volleyball.

cisely locating signals that are moving (FM), only there part of the time (AM), and that have harmonics only 10 dB down (or 30 dB down if higher than third). In the process of finding the fundamental of the incoming signal the octave oscillator is swept rapidly across its band about 20 times. To do this, a current source is fed into the integrator input to the VCO (Fig. 8). The output of the integrator is a ramp that causes the VCO to move down in frequency. Another current source retraces the VCO at the end of a band if no signal is found.

The microprocessor sweeps each octave until a signal is found. A fast detector in the IF chain senses the presence of a signal in the IF and switches off the current source driving the integrator. When a sweep is successful in generating an IF signal, the local oscillator is counted to determine where it was when it stopped.

The sweep rate is very fast, up to 1 MHz/μs. This presents the problem that any time involved in recognizing the presence of an IF signal and turning the sweep source off causes the VCO to go too far. To compensate for this, a step-back resistor is put in the integrator, such that removing the sweep current removes $V = IR$ from the output voltage of the integrator. R is chosen so that the frequency represented by the sweep current times the resistor is equal to the amount the output frequency moved during the turnoff delay. This allows the output of the VCO to be counted to determine where it was when the signal first entered the IF.

Demodulators

The modulation analyzer's two demodulators operate on the IF signal and produce a pair of audio signals. One demodulator extracts the frequency modulation components of the IF signal and the other extracts the amplitude modulation components. The two demodulators are designed for low susceptibility to incidental modulation, or conversion from one type of modulation to the other. They are also designed not to add significantly to the modulation noise output so that they do not degrade the system residual modulation characteristics. They are also both very linear and have flat frequency responses so that the overall 1% accuracy specification can be maintained. Both detectors do a fine job of meeting their design goals.

The job of removing the AM from the IF signal before it is sent to the FM demodulator is done by a series of three limiter stages. Considerable care was taken to minimize the effect of amplitude variations on the phase of the limiter output (minimum AM to ΦM conversion). It is most important that both the positive half cycle and negative half cycle be treated the same and always symmetrically. It is obvious that if the limiting does not take place exactly at the zero crossing, changes in amplitude will affect the zero crossing and, in turn, the FM output. (However, the 8901A's full-wave FM demodulator averages the zero crossings so that some symmetrical movement can be tolerated.) To minimize this effect, we treat the signal within the limiter in a dual-ended differential arrangement (Fig. 9).

Each limiter stage consists of a differential pair with the emitters tied together and to a current source. The collector loads are individual resistors with emitter followers to take the signal differentially to the next stage or to the output. By maintaining strict symmetry, both phases of the signal are treated the same. Since the zero crossing information is carried in the difference between the two outputs any asymmetry that might exist at one output will not exist in the difference between the two. Symmetry was further assured by having all of the transistors for a given limiter on the same transistor array integrated circuit chip.

Perhaps the most important consideration involves maintaining all of the operating conditions independent of input level. The current source on the emitters keeps the operating current and consequently the output level independent of base voltage. The differential arrangement keeps the transistors out of saturation (a necessity) and the emitter output stages buffer the loads from changes in the input
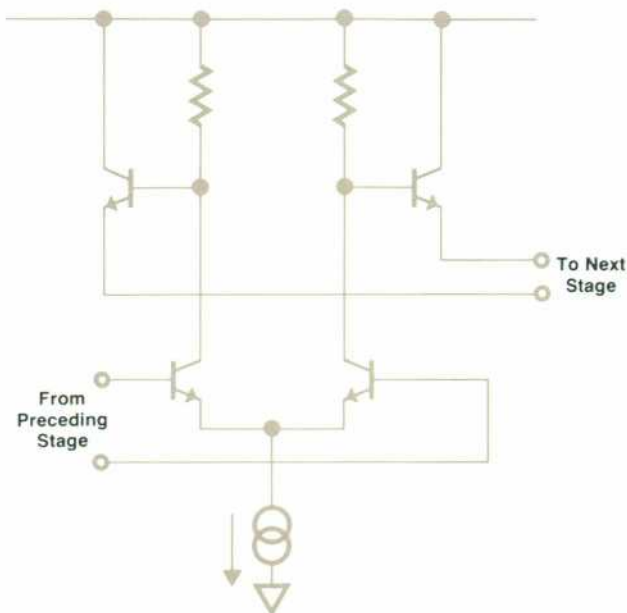
Fig. 9. *Limiters (one stage shown here) remove the AM from the IF signal before it goes to the FM demodulator. The design minimizes AM-to-ΦM conversion.*

impedance of the following stage as it begins to limit.

Some care was taken to make the stage delay independent of input level. Normally the delay through a differential stage gets smaller as the stage is driven harder. In this design, the small-signal bandwidth is extended by applying negative feedback. (When the stage is limiting, the feedback is not effective, of course.) The feedback equalizes the delay through the stage so that the small-signal delay is equal to the large-signal delay.

With all of these steps, phase shift at the limiter output is typically less than ½ degree for a 3-to-1 change in input level.

FM Demodulator

The most difficult task in selecting an FM demodulator was to achieve both the noise specification and the distortion specification we needed for our various potential users. To satisfy the distortion requirements of broadcast FM applications we are providing FM distortion that is typically 70 dB down at rates and deviations up to 100 kHz. This is very good, and suggests a pulse-count type of discrimination. On the other end of our target applications are the high-quality, low-deviation FM mobile radio applications. For these we are providing discriminator noise in a 3-kHz bandwidth that is small compared with 1 Hz (around 1/3 Hz, typically). This is also very good, and suggests a tuned-circuit type of discriminator. The problem is that a pulse-count discriminator is too noisy and a tuned discriminator is not linear enough. We were able to develop a new type of discriminator that is neither pulse-count or tuned-circuit, yet meets all our requirements. The beauty of this circuit is that it is inherently linear and yet does not have a significant noise mechanism. It is described in detail on page 13.
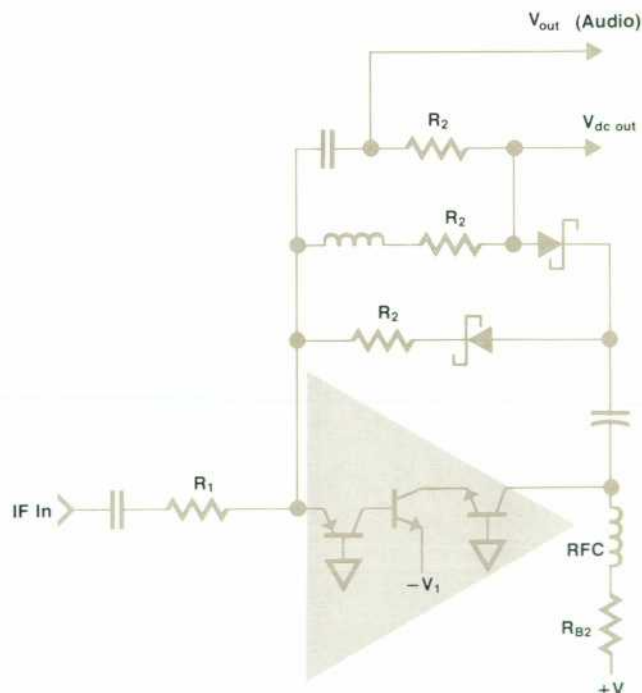
AM Demodulator and AGC

The AM demodulator detects the amplitude modulation

on the IF signal and produces an audio-frequency output. This detector has a very wide IF bandwidth and therefore very low conversion of FM to AM. At most operating frequencies, conversion of FM to AM is negligible, approaching the published specification only at the lowest input frequencies when the RF input circuit and the mixer begin to roll off slightly.

The demodulator is a half-wave average-responding rectifier with sufficient bandwidth to work well with an IF above 2.5 MHz. It is similar to a circuit sometimes used at lower frequencies, with diodes in the feedback loop of an amplifier. The difference is that the amplifier is a very high-gain amplifier capable of the high slew rates required to turn the diodes on during the IF signal's transition through zero.

A simplified schematic is shown in Fig. 10. (Some biasing and compensation elements have been omitted.) Gain is provided by a cascade of three transistors arranged for a high gain-bandwidth product. The current source (grounded base) output allows the fast voltage changes required to switch the diodes as the sign of the current reverses. The audio output is taken from one of the half-wave outputs by way of a constant-resistance network. The constant-R network filters the IF signal so that very little of it reaches the buffer amplifier that follows the demodulator. Without the constant-R filter, the detection diode would be back-biased at each IF zero crossing by the average value of the half-wave-rectified IF signal, and this would degrade performance because the amplifier would have to slew farther.

The IF gain of the circuit is R_2/R_1. To have reasonable gain and yet keep R_2 small enough to get the required bandwidth, R_1 has to be small. Its value is actually 100Ω, so

Fig. 10. *AM demodulator has a wide IF bandwidth and therefore very low FM-to-AM conversion. It detects the modulating signal (V_{out}) and the average carrier level ($V_{dc\ out}$).*

it requires a good stout IF amplifier to drive it.

The amplitude modulation percentage is the amplitude of the AM detector audio output divided by the average carrier level times 100%. The average carrier level is also detected by the AM detector; it is labeled $V_{dc\,out}$ in Fig. 10.

An AGC circuit ahead of the AM demodulator (Fig. 11) provides a constant-level IF signal to the demodulator. $V_{dc\,out}$ is compared with a stable reference and an error

A New Type of FM Demodulator

by Russell B. Riley

We have called the wideband low-noise FM demodulator used in the 8901A Modulation Analyzer a charge-count discriminator. The basic idea is to form pulses of constant charge at a rate proportional to frequency and then average these pulses to produce an output voltage proportional to the input frequency.

The operation of the charge-count discriminator is similar to that of the more familiar pulse-count discriminator, the basic difference being the pulse shapes involved. The typical pulse-count circuit forms a pulse of constant amplitude and constant duration once per cycle of the signal to be demodulated. Note that both the amplitude and the duration have to be controlled with great accuracy and stability for linear, low-noise performance. In practice it is usually jitter

Thus the demodulator linearity depends on passive components and a stable voltage ΔV. The principal source of noise is the op-amp U, with resistor R and noise on ΔV making somewhat smaller contributions.

The circuit actually used in the modulation analyzer includes some refinements. For example, transistor Q1 is also used, driving circuitry similar to that connected to the collector of Q2. The two outputs are added, with the result that the demodulator transfer function (sensitivity) is doubled and the output ripple frequency is doubled, making the filtering job easier. A damped inductor is added in series with C1 to introduce a controlled amount of overshoot so that the steering diodes cut off cleanly.
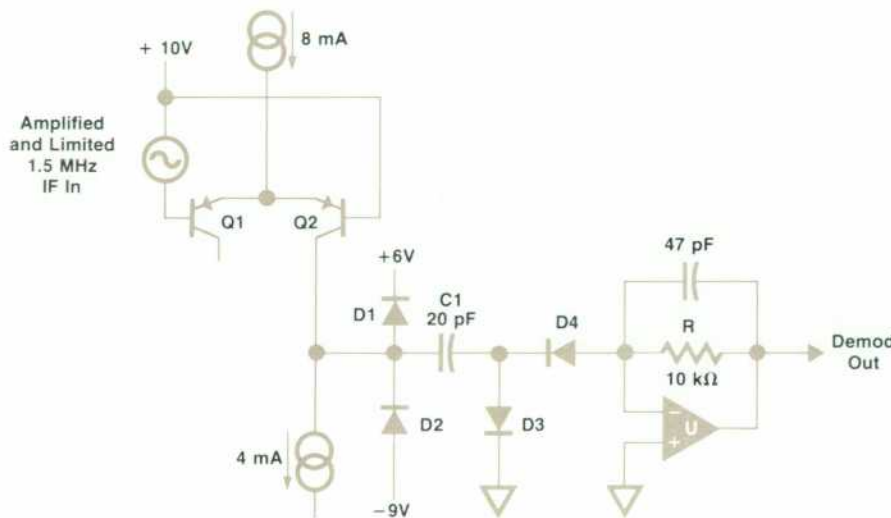

Fig. 1. *Simplified schematic of the charge-count discriminator.*

in the pulse duration that limits noise performance. In contrast, the charge-count circuit requires only that a dc voltage be accurate and stable.

In the circuit diagram (Fig. 1), diodes D1 and D2 clamp the left end of capacitor C1 to a voltage swing of 15V plus two diode drops. Charge-steering diodes D3 and D4 limit the right end of C1 to a voltage excursion of two diode drops. Thus the charge that flows back and forth in C1 is well defined and is given by

$$Q = C1\,\Delta V$$

where ΔV is 15V in this example. Since current is given by the rate of change of charge with respect to time the average current I_{avg} flowing in resistor R, through the action of the charge-steering diodes and the operational amplifier U, is given by

$$I_{avg} = Qf = C1f\,\Delta V$$

where f is the number of cycles per second (the frequency of the input signal). The average output voltage V_{avg} is simply

$$V_{avg} = RI_{avg} = RC1f\,\Delta V$$
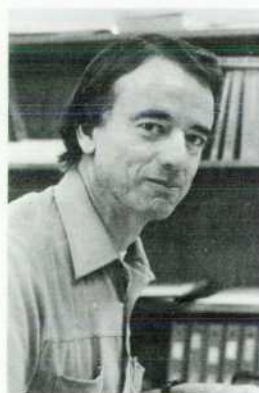
Russell B. Riley
A 1959 graduate of the University of Colorado, Russ Riley completed his studies for the PhD degree at Stanford University in 1961. His responsibilities at HP have included the 938A and 940A Frequency Doubler Sets, waveguide thermistor mounts, the 423A and 424A Crystal Detectors, the 415E SWR Meter, the 432A Power Meter, and parts of the 8558A Spectrum Analyzer. He is named as inventor on several patents on these products. Most recently he contributed to the design of the 8901A Modulation Analyzer, especially the FM limiter-discriminator; he's now with the optoelectronics section of HP Laboratories. Born in Kansas City, Missouri, Russ is married and lives in Portola Valley. His daughter attends high school, and his two sons are in college. He spends some of his leisure time gardening and singing in his church choir, and he played clarinet in the Peninsula Symphony for three years.
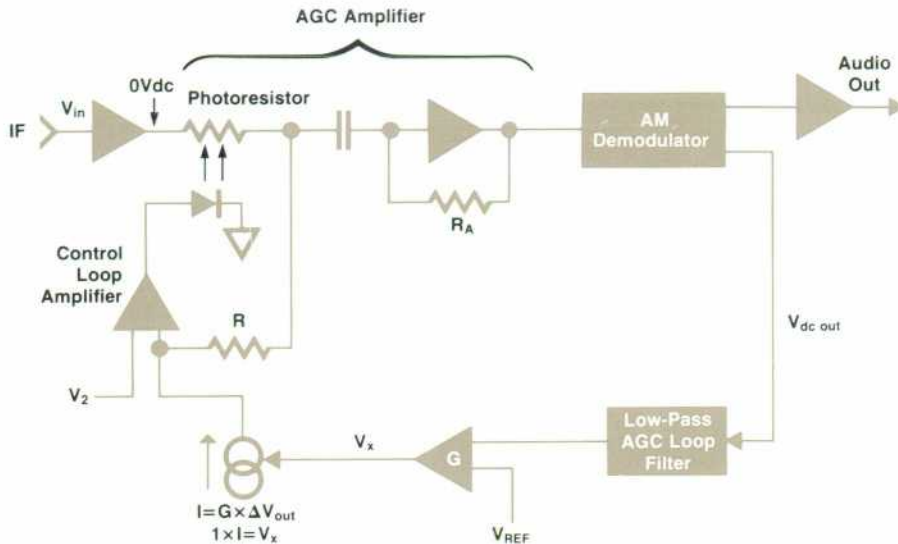
Fig. 11. *Automatic gain control (AGC) circuit provides a constant-level IF signal to the AM demodulator. A low-distortion photoresistor is the gain control element.*

voltage is developed to control the gain of the AGC amplifier. The limiting factor in the design of the AGC amplifier is the AM distortion goal of −60 dB harmonic distortion. A special low-distortion photoresistor is used as the gain control element in a feedback arrangement. The approach works well, but is not without pitfalls. The photoresistor is a slow device with a low-frequency pole that tends to move with the operating point. It also has a gain factor that is not very stable. In general, it is not the kind of element a designer likes to have determining the dynamics of a loop.

To avoid these problems, the photoresistor is placed in a control loop to set its resistance (Fig. 11). A small dc current is put through the photoresistor and an operational amplifier senses the voltage across the photoresistor. The op-amp changes the current through the light-emitting diode that illuminates the photoresistor until the voltage across the photoresistor equals a predetermined value. In this way the dc current controls the resistance of the photoresistor and consequently the gain of the AGC amplifier. The pole of the photoresistor as well as its gain variations are removed by the action of this local control loop. Since the AM signal going through the photoresistor is at IF and the resistance-setting signal is at dc, these signals do not interfere with each other.

The AGC loop can be modeled as a feedback loop around a multiplier (see Fig. 12). The multiplication factor V_c is the gain of the AGC amplifier, so that $V_{out} = V_{in} \times V_c$. The loop gain is $GK(dV_{out}/dV_c) = GKV_{in}$. Thus the loop gain is proportional to the input level. Although this is characteristic of an AGC loop, it is undesirable because the loop bandwidth and response time are functions of the loop gain. If the loop bandwidth is too high, the circuit will begin to remove some of the amplitude information from the signal, and ringing will occur for step RF inputs. If the loop bandwidth is too small the response time will be very long. Unless something is done to compensate for the changing loop gain, one or the other of these effects can be expected to occur.

Our solution is to make the gain factor K inversely proportional to V_{in}. Since $K = dV_c/dV_x$, we will have constant

bandwidth if

$$\frac{dV_c}{dV_x} \propto \frac{1}{V_{in}} \quad \text{or} \quad \frac{dV_c}{dV_x} \propto \frac{V_c}{V_{out}}$$

Because the circuit is an AGC circuit, $V_{out} = V_{ref} =$ constant. Therefore we need

$$\frac{dV_c}{dV_x} \propto V_c \, .$$

An exponential with $V_c \propto e^{K_2 V_c}$ meets this requirement. In the 8901A the exponential characteristic is approximated by the series combination of R and R_P in Fig. 11. Current I develops a voltage $I(R+R_P)$ across this series combination. The control loop amplifier drives the photoresistor to make this voltage equal to the reference voltage V_2. The AGC amplifier gain is

$$\text{Gain} = V_c = \frac{R_A}{R_P} = \frac{IR_A}{IR+V_2} \approx \frac{dV_c}{dV_x}$$

For the values used, the maximum gain error of this approximation over a 24-dB IF input range is 2 dB. This 12-to-1 improvement allows us to maintain both accuracy and response time over the operating range of the circuit.
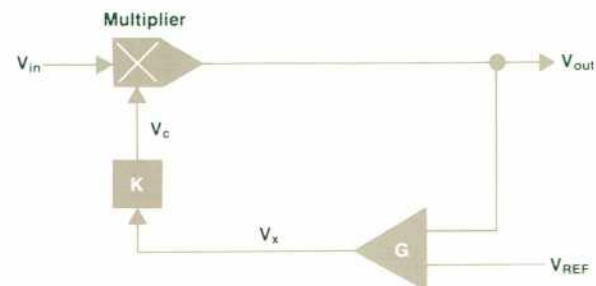
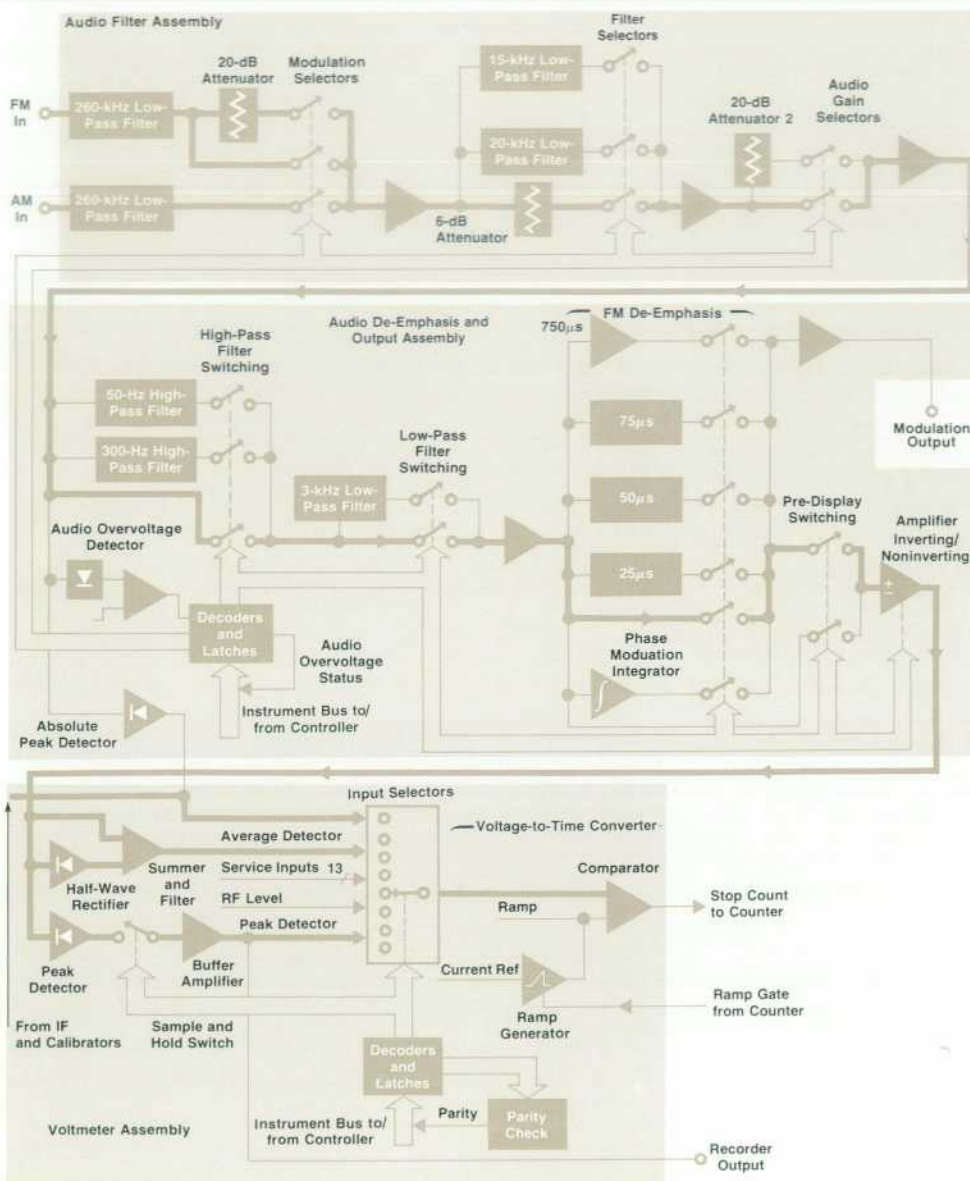**Fig. 12.** *AGC loop can be modeled as a feedback loop around a multiplier.*

Fig. 13. *Audio section processes and measures the modulation signal.*

Audio Section

The audio section (Fig. 13) selects the output of either the AM demodulator or the FM demodulator, generates a digital result proportional to the amount of modulation, and outputs the demodulated signal to the front panel. In the process, the signal is filtered, amplified, de-emphasized, detected and converted to digital format.

Three low-pass filter bandwidths and two high-pass filter bandwidths are provided, along with a wideband filter representing the unfiltered case. The wide filter's 3-dB bandwidth is a few tenths of one hertz at the low end to about 250 kHz at the high end. The low-frequency cutoff had to be kept very low to preserve FM stereo information at low frequencies. For good stereo separation only small phase shifts between the audio tone and the 38-kHz subcarrier are permissible. The lower cutoff frequency of 0.3 Hz limits the maximum measurable separation at 50 Hz to 50 dB.

This filter is required to be flat to 200 kHz for FM and to 100 kHz for AM. Separate filters are used for the FM and AM inputs to minimize the problems associated with switching between modulators. The actual switching throughout the audio section is done at high-impedance points so that switch impedance variations do not lead to gain variations. The wide filters are designed to attenuate the 1.5-MHz IF signal and its sidebands.

The two high-pass filters, 50 Hz and 300 Hz, represent standard frequencies a user might need.

The three low-pass filters are 3kHz, 15kHz, and >20 kHz. The 15-kHz filter is a five-pole Butterworth filter. It is designed to attenuate the carrier when operating with a 150-kHz or greater input frequency. The 3-kHz filter is an active five-pole filter used for noise measurements. This filter is designed with many poles so that noise that might be rising with audio frequency (typical in FM systems) will be filtered enough to give meaningful results. The >20 kHz filter has a 3-dB bandwidth of about 90 kHz but has a Bessel shape, so that is is down about 2% at 20 kHz. It is optimized

for minimum overshoot to a square wave input. This filter can also attenuate the IF when operating at a 455-kHz IF. Either this filter or the 15-kHz filter may be used with the 455-kHz IF.

When FM is selected there is a choice of de-emphasis networks. These are required to reconstruct the modulation signal, since an FM transmitter typically has its frequency response pre-emphasized. When a de-emphasis network is selected it is always switched in ahead of the demodulated output. The display, however, can represent the modulation either before or after de-emphasis. This allows either deviation (without de-emphasis) or flatness (including de-emphasis) measurements to be made while maintaining the correct frequency response at the demodulated output. When the 750-μs de-emphasis network is selected a gain of 10 is automatically inserted. This gives the added resolution often needed with this de-emphasis network because it reduces the received noise so much (750 μs is equivalent to a 212-Hz low-pass filter with one pole.)

When ΦM is selected the FM signal is integrated to form $\Phi = \int f dt$. The phase modulation capability assumes wide-deviation phase modulation, with readings to 400 radians. Modulation schemes that depend upon zero radians being the same as 2π radians will cause erroneous readings, and phase modulation with incidental AM much greater than 99% may also cause incorrect readings. The ΦM demodulator works very well with analog phase modulation and is useful with some types of digital PSK (phase shift keying).

Average and Peak Detectors

Two detectors are provided in the audio section. The average-responding detector, similar to that found in high-quality voltmeters, is included mainly to measure noise. The peak detector is used mainly to measure maximum modulation. An amplifier that can be programmed to a gain of +1 or -1 precedes the peak detector so that either +pk or -pk readings can be made. The performance of the peak detector is key to the operation of the modulation measurements and much effort went into its development.

As simple as the peak detector is (see Fig. 14), it has very good performance. Its maximum allowable error for 100-kHz audio rates is 0.15% on our production line. A sine wave is within 0.15% of peak value for only 1.7% of the time, or only 170 ns at 100 kHz. This requires a very fast circuit. For this reason, a high-speed comparator is used as
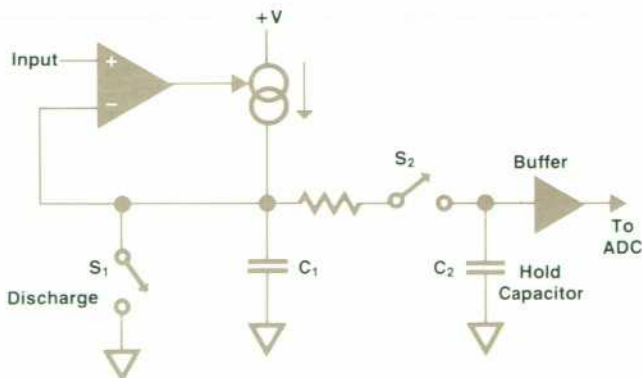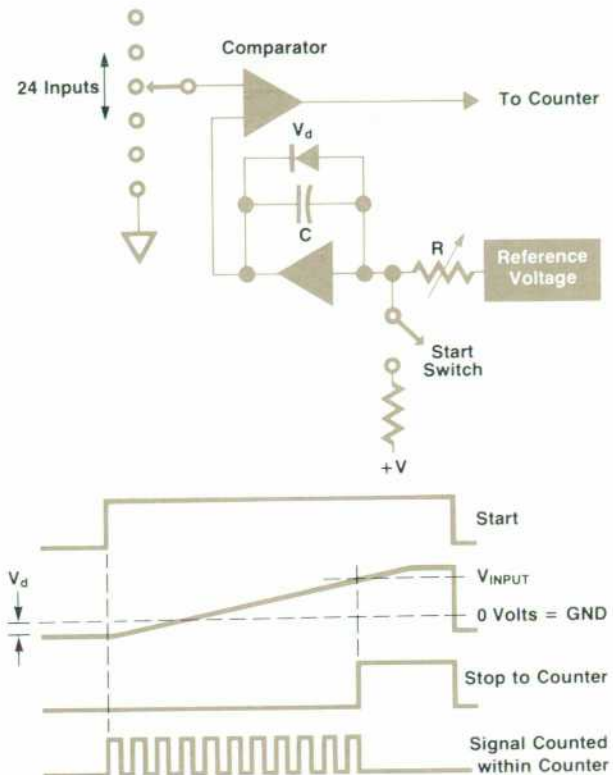

Fig. 15. *Analog-to-digital converter is used to measure many voltages in the modulation analyzer. It is a single-slope converter, counting a clock signal while a constant-slope ramp voltage rises to the level of the unknown voltage.*

the basis of the circuit. At the other end of the spectrum, for a 50-Hz input the sample capacitor must retain its charge for 20 ms. With a maximum of 0.15% discharge allowed, this represents a time constant of 13 seconds. To maintain this long time constant given the levels and input bias current in the comparator, the sampling capacitor is large (4μF). This long time constant also requires that the circuit be actively discharged to maintain reasonable measurement speed. The circuit is cycled automatically to measure a new peak value 10 times a second.

The current that charges the sample capacitor is large so that the detector is able to respond to complex signals, such as two simultaneous tones. In this case, which is typical of avionics signals, peaks occur only at the difference frequency between the two tones, yet the widths of the peaks are proportional to the highest frequency. The problem here is to charge the sample capacitor fully while the two inputs are summing to a maximum. On the other hand, the higher the charge current, the greater the potential overshoot. The charge current level is chosen so that this overshoot is less than one least significant digit.

After the peak sample has been taken, the result is transferred to a hold capacitor and buffered for the analog-to-digital converter. If S_2 in Fig. 14 is held closed for a sufficient length of time, then nearly all the voltage on C_1 is transferred to the smaller C_2. However, if some post-peak-detection filtering is required, then S_2 is closed for a shorter time. In this way a fraction of the voltage difference between C_1 and C_2 is transferred for each time period. The result is

Fig. 14. *Peak detector is simple but performs well.*

a time-sampled approximation to an ideal one-pole post-detection filter. The advantage is that it is programmable and it does not add much hardware—just two resistors plus the latch and transistor that change the timing to the type-555 timer that sequences the peak detector.

If the timer is turned off and S_2 is left closed, then any signal that comes along will be captured. This is called peak hold; it is implemented with digital storage to provide indefinite retention of maximum levels of peak modulation.

The analog-to-digital converter is used to read many voltages inside the modulation analyzer. Readings are made to determine the result and validity of each measurement. The ADC can make measurements as fast as 10 ms, has resolution of 1 part in 40,000, and has only one adjustment and inherently no offset.

The converter is a single-slope converter (Fig. 15). When it receives a start transition, it generates a ramp by integrating a precision current. When the ramp exceeds the input voltage, a comparator changes state. A clock is counted between the start transition and the comparator state change. The number of counts is proportional to the analog voltage.

Instead of starting at ground, the ramp starts one diode drop below ground, and goes up to about 5 volts. Once per measurement cycle a reading of a grounded input is taken and stored. The input voltage is then computed as the signal reading minus the ground reading. Offset is not a problem because ground is actually measured. Temperature drift in the integration capacitor is compensated by a temperature coefficient introduced into the voltage reference for an overall effect of less than 30 ppm/°C.

Serviceability

Several features were included in the design of the 8901A to make servicing easier. These start with the power-up sequence and with four dedicated troubleshooting I/O pins on one of the microprocessor chips. Four lights on the top of the controller board are associated with these four pins. If all is well these lights quickly go through a flashing sequence on power-up, then all stay on. If any other pattern of these lights is displayed it means that one of the power-up checks has failed. After the power-up sequence is complete, two of the lights are used to indicate interrupts (keyboard and HP-IB), and one light flashes to indicate that the controller is running.

Many items are checked during power-up. First a checksum is made on each of the nine read-only-memory chips to insure that the program is all right. The checksum is compared against a value stored in one of the ROMs. The code to run this program is located entirely in the first ROM, so that a bad ROM will tend to show up before it has any effect on the validity of what is being done. Next all of the read/write memory (RAM) is checked by writing all zeros, then all ones, and verifying that the same patterns are read back.

An exclusive-OR gate tied to the output of one of the I/O data latches forms a parity check on that latch and is read back to check that the instrument I/O is working. The keyboard is scanned to make sure a key is not stuck down and an error light is turned on if a stuck key is detected. Finally, a test is run on the local oscillator to verify that it is working properly. Tested are its frequency range, loop gains, and lock stability. The LO tests are done in sequence to help isolate any problem. A special function command from the keyboard produces a readout that indicates what test number failed. At the beginning of the power-up tests, a scan of the four service pins already mentioned is made to see if one is grounded. A grounded pin indicates that the user wants the instrument to go into one of its signature analysis[3] routines.

To troubleshoot the microprocessor, a special extender board is used that allows "no-operation" instructions to be forced onto the processor control bus. If all is going well the processor will increment its program counter and step through all of memory. The signature analysis probe will recognize characteristic signatures in address lines, memory decoders, and memory chips, allowing easy identification of problem parts.

The keyboard assembly is also troubleshot using signature analysis. The routine scans each key and outputs a 1 or 0 for each key to one of the four troubleshooting pins on the microprocessor board. There is a signature for no key down and a unique signature for each key. There is also a signature for each pair of keys, and so on. If a fault is found, the board itself can be probed until the problem is isolated to a particular component. Also, by correlating patterns of key failures with the organization of the keyboard matrix, many clues can be obtained even without going onto the board.

The software-controlled counter circuits also have a signature analysis routine to exercise them. The various latches are located and cleared and the counters allowed to count. By comparing signatures with the expected signatures on a schematic, one can isolate a fault to a particular part.

The remaining digital circuits and the various I/O latches are not tested by signature analysis. The controller part of the I/O circuitry is tested automatically on power-up. To test a particular latch controlling some function on a board another set of special functions is included. These functions allow the service person to write any desired data to any latch in the instrument. The controller then presents stable data and addressing lines to the particular latch and continuously strobes its clock or latch input. By examining the outputs and comparing them with the inputs, the fault can be isolated. This approach offers some advantage over signature analysis. First, if one is looking at a particular latch there is probably reason to believe that at least the function controlled by that latch isn't working. Also, because a latch is, in general, controlling analog hardware, probing need not stop when the TTL levels stop. Thus one can set up a command word and follow it onto a board, into a latch, into a level translator, and to an FET switch. The problem may be well out of the interface hardware, and this function allows one to set up a stable state of the analog hardware, which is sometimes necessary to isolate the problem.

Some of the inputs to the voltmeter and counter are there for service. The voltmeter is able to measure the instrument's five power supply voltages to verify their operation. The counter is able to count the internal and external time bases and thus it can be verified which time base is actually being used (the switchover is automatic to an external time base when one is connected).

References

1. D. Scherer, "Design Principles and Test Methods for Low Phase Noise RF and Microwave Sources," HP Microwave Symposium, 1979.
2. "Frequency Stability Specification and Measurement, High-Frequency and Microwave Signals," U.S. National Bureau of Standards Technical Note 632.
3. R.A. Frohwerk, "Signature Analysis: A New Digital Field Service Method," Hewlett-Packard Journal, May 1977.

Allen P. Edwards

A native Californian, Allen Edwards earned his BSEE and MSEE degrees at Stanford University in 1971 and joined HP the same year. He worked as a design engineer on the 8558A Spectrum Analyzer, the 435A Analog Power Meter, and the 8481A Coaxial Power Sensor, and served as project manager for the 436A Digital Power Meter and the 8901A Modulation Analyzer. He is named as inventor on two patents. A resident of Palo Alto, Allen spends part of his spare time backpacking and bicycling; he also enjoys photography.

SPECIFICATIONS
HP Model 8901A Modulation Analyzer

RF Input

FREQUENCY RANGE: 150 kHz to 1300 MHz.
OPERATING LEVEL:
150 kHz - 650 MHz: 12 mV$_{rms}$ (−25 dBm) to 7 V$_{rms}$ (1 W$_{peak}$).
650 MHz - 1300 MHz: 22 mV$_{rms}$ (−20 dBm) to 7 V$_{rms}$ (1 W$_{peak}$).

Frequency Modulation

RATES: 150 kHz - 10 MHz: 20 Hz to 10 kHz.
10 MHz - 1300 MHz: 20 Hz to 200 kHz.[1]
DEVIATIONS:
150 kHz - 10 MHz: 40 kHz$_{peak}$ maximum.
10 MHz - 1300 MHz: 400 kHz$_{peak}$ maximum.[1]
ACCURACY:[2]
250 kHz - 10 MHz: ±2% of reading ±1 digit, 20 Hz to 10 kHz rates.
10 MHz - 1300 MHz: ±1% of reading ±1 digit, 50 Hz to 100 kHz rates. ±5% of reading ±1 digit, 20 Hz to 200 kHz rates.
DEMODULATED OUTPUT DISTORTION:[3]
400 kHz - 10 MHz: <0.1% THD, deviations <10 kHz.
10 MHz - 1300 MHz: <0.1% THD, rates and deviations <100 kHz.
AM REJECTION (for 50 % AM at 400 Hz and 1 kHz rates):[2] <20 Hz peak deviation measured in a 50 Hz to 3 kHz BW.
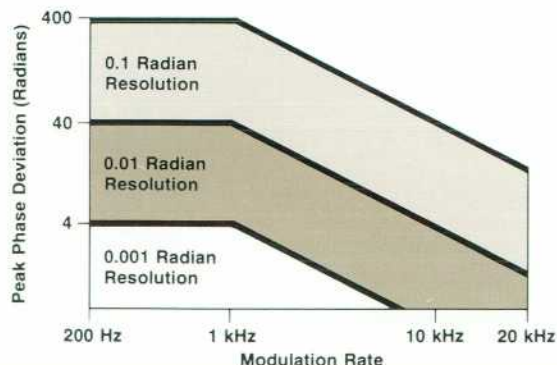**RESIDUAL FM** (50 Hz to 3 kHz BW): <8 Hz$_{rms}$ at 1300 MHz, decreasing linearly with frequency to <1 Hz$_{rms}$ for 100 MHz and below.

Phase Modulation

CARRIER FREQUENCY: 10 MHz to 1300 MHz.
RATES: 200 Hz to 20 kHz.
DEVIATION AND MAXIMUM RESOLUTION:

ACCURACY:[2] ±3% of reading ±1 digit.
DEMODULATED OUTPUT DISTORTION: <0.1% THD.
AM REJECTION (for 50% AM at 1 kHz rate):[2] <0.03 radian peak in a 50-Hz-to-3-kHz BW.

Amplitude Modulation

RATES: 150 kHz - 10 MHz: 20 Hz to 10 kHz.
10 MHz - 1300 MHz: 20 Hz to 100 kHz.
DEPTH: to 99%.
ACCURACY:[2,4]
150 kHz to 10 MHz: ±2% of reading ±1 digit, 50 Hz to 10 kHz rates, >5% depth. ±3% of reading ±1 digit, 20 Hz to 10 kHz rates.
10 MHz to 1300 MHz: ±1% of reading ±1 digit, 50 Hz to 50 kHz rates, >5% depth. ±3% of reading ±1 digit, 20 Hz to 100 kHz rates.
FLATNESS (variation in indicated AM depth for constant depth on input signal):
10 MHz to 1300 MHz: ±0.3% of reading ±1 digit, 90 Hz to 10 kHz rates, 20 to 80% depth.
DEMODULATED OUTPUT DISTORTION: <0.3% THD for ≤50% depth, <0.6% THD for ≤95% depth.
FM REJECTION (at 400 Hz and 1 kHz rates, 50 Hz to 3 kHz BW):[2]
250 kHz to 10 MHz: <0.2% AM for <5 kHz$_{peak}$ deviation.
10 MHz to 1300 MHz: <0.2% AM for <50 kHz$_{peak}$ deviation.
RESIDUAL AM (50 Hz to 3 kHz BW): <0.01% rms.

Frequency Counter

RANGE: 150 kHz - 1300 MHz.
SENSITIVITY:
150 kHz - 650 MHz: 12 mV$_{rms}$ (−25 dBm).
650 MHz - 1300 MHz: 22 mV$_{rms}$ (−20 dBm).
ACCURACY: Reference accuracy ±3 counts of least significant digit.
INTERNAL REFERENCE:
FREQUENCY: 10 MHz.
AGING RATE: <1×10^{-6}/month.
(Optional: 1×10^{-9}/day after 30-day warmup).

	Standard	Option 002
Aging Rate	<1 × 10^{-6}/mo.	<1 × 10^{-9}/day
Temperature Effects	<2 × 10^{-7}/°C	<2 × 10^{-10}/°C
Line Voltage Effects (+5%, −10% line voltage change)	<1 × 10^{-6}	<6 × 10^{-10}
Short term stability		1 × 10^{-9} for 1s average

Modulation Analyzer Applications

by Allen P. Edwards

THE 8901A MODULATION ANALYZER is a useful tool for analyzing many types of signals. Often it can provide needed information that has been difficult to obtain, such as incidental FM or residual FM. It can replace large, complex test systems, and speed and simplify measurements.

The modulation analyzer is well suited for measuring mobile communications and other transmitters. This single instrument can be used in making most of the measurements normally made on transmitters, such as carrier power, carrier frequency and stability, AM depth, FM de-

viation, hum and noise, incidental AM or FM, modulation limiting (instantaneous and steady-state), and audio frequency response.

For avionics applications the 8901A can be very useful in measuring navigation signals. In testing ILS transmitters the analyzer can be used to measure depth of modulation very accurately. For broadcast AM and FM it can be used to measure AM depth or FM deviation, and it can accurately recover the modulation for making measurements such as stereo separation and distortion.

Its accuracy makes the modulation analyzer an excellent
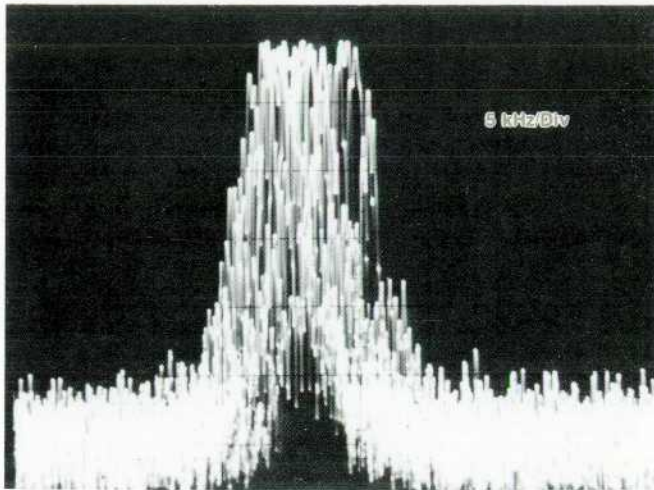
Fig. 1. *Spectrum analyzer display of a supposedly amplitude modulated signal is more typical of broadband FM than single-frequency AM.*
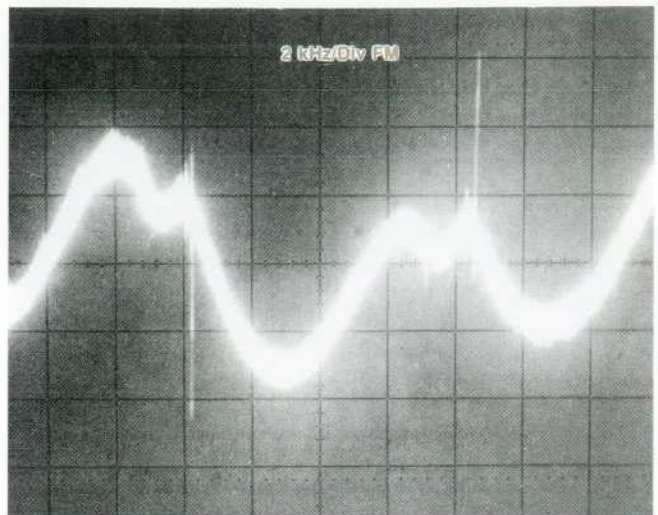
Fig. 3. *FM recovered by the modulation analyzer from the signal of Fig. 1 shows sharp spikes causing the broadband FM appearance of the signal's spectrum.*
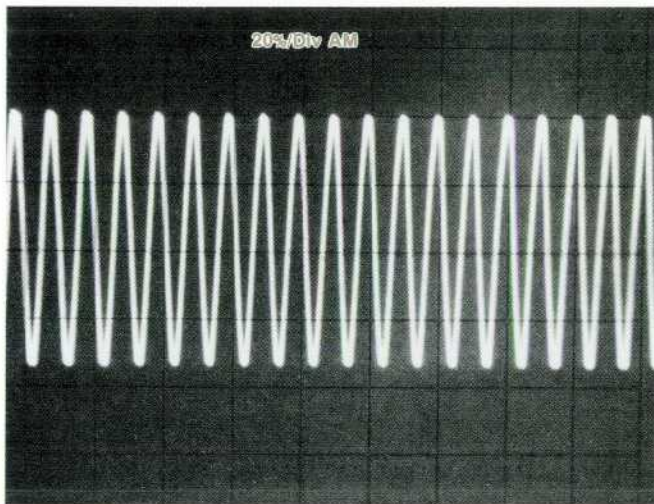
Fig. 2. *AM recovered by the 8901A modulation analyzer from the signal of Fig. 1 is a single-frequency sine wave.*

characterize RF and IF designs, evaluate modulators, and test individual ICs or modules.

In the following sections, three specific applications are described that demonstrate the modulation analyzer's capabilities.

Solving a Signal Generator Problem

The 8901A's ability to separate amplitude modulation and frequency modulation is demonstrated by this example. Fig. 1 is a spectrum analyzer display of a signal coming from a high-power signal generator that was being amplitude modulated at 1 kHz. The spectrum is more typical of broadband FM than of single-frequency AM.

To see what was happening, the signal was applied to an 8901A. Fig. 2 shows the AM detected by the modulation analyzer, revealing that the signal is indeed being amplitude modulated at 1 kHz. However, the analyzer's FM output, Fig. 3, shows that the generator is also being frequency modulated with line-related interference and that the sharp spikes are causing the broadband FM we saw in the spectrum analyzer display. With these pictures, the

addition to a metrology laboratory. An example of its usefulness is in calibrating signal generators, especially high-performance signal generators such as the HP 8640B. The modulation analyzer's capabilities exceed those required to verify many signal generator specifications. Besides improving the accuracy of these measurements, it greatly reduces the time involved. Also, the optional calibrators provide a new level of modulation standard accuracy and help ensure accurate measurements.

Because the modulation analyzer is capable of characterizing all types of signals, it is useful in research and development laboratories for characterizing VCOs, measuring residual noise on crystal oscillators, measuring incidental modulation, measuring the frequency of low-level signals, and so on. When used with a signal source it can
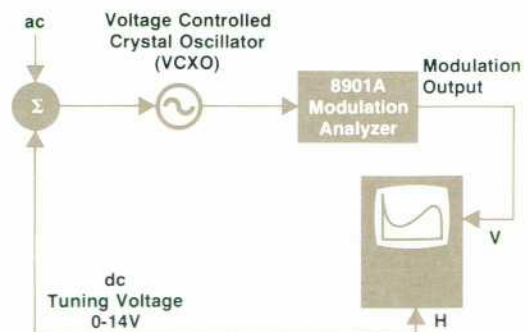
Fig. 4. *Using the modulation analyzer to characterize voltage-tunable oscillators. The low noise of the 8901A's local oscillator makes this application possible.*

53 Hz
= −4.29 dB

85 Hz = 0 dB
68 Hz

279.9927 MHz   279.9985 MHz
0V   14V

Fig. 5. *Oscilloscope display from the setup of Fig. 4. Ac input to the VCXO was adjusted to give 100-Hz deviation at the low end of the VCXO's tuning range. The display shows how the actual deviation varies as the VCXO is tuned across its band.*
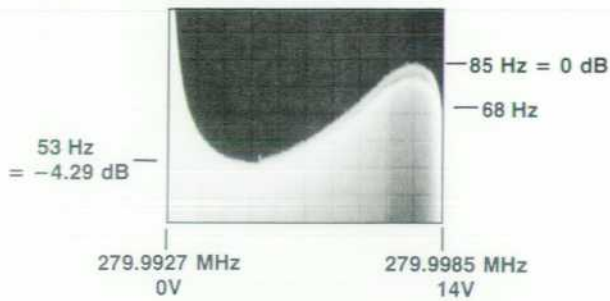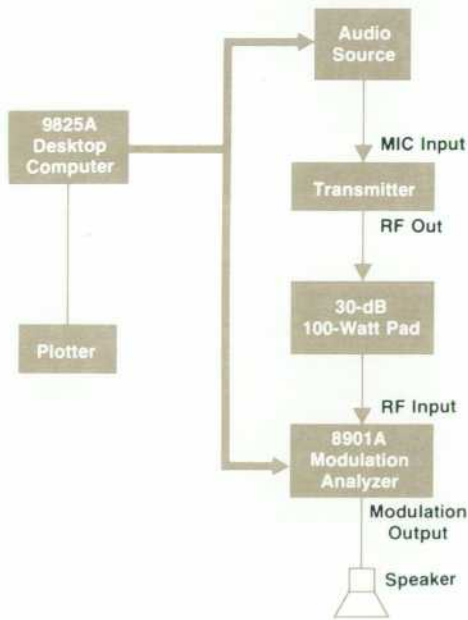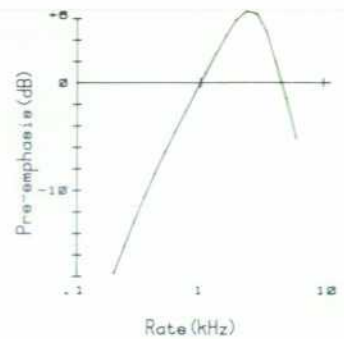

Fig. 6. *An automatic system for testing mobile transmitters. A desktop computer controls the modulation analyzer via the HP Interface Bus.*

FM Mobile
Transmitter
Model XYZ

Transmitter Tests

Power=	48.5 Watts
Frequency=	132.00008 MHz
Freq. Error=	80 Hz
Mic. Sens. =	123 mV
Mod. Limiting=	4.27 kHz
Residual FM=	2 Hz
Incidental AM=	0.01 %
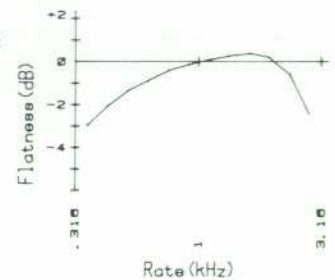
Fig. 7. *Output plot from the system of Fig. 6.*

manufacturer of this product might have been able to cure this problem by a little attention to the power supply rectifiers, which, one would suspect, are turning off very fast and causing the spikes.

Characterizing Voltage-Tunable Oscillators

The next application demonstrates how the 8901A is useful for characterizing VCOs and VCXOs. Fig. 4 shows a VCXO being measured by the 8901A. The VCXO is operating at 280 MHz and can be pulled 5800 Hz. The idea is to vary the dc tuning voltage from 0V to 14V, causing the VCXO frequency to change by 5800 Hz, while a small ac test signal of constant amplitude frequency-modulates the VCXO. The ratio of the ac frequency deviation to the ac test signal ($\Delta f/\Delta v$) is the gain of the oscillator at the dc operating point. With the ac test signal held constant, the ac frequency deviation is proportional to gain. An oscilloscope is used to plot the dc tuning voltage against the frequency modulation on the VCXO output, as detected by the 8901A. The en-

velope of the oscilloscope display shows the peak FM deviation of the VCXO as a function of center frequency.

Fig. 5 shows the oscilloscope display. Notice that the peak deviation of the VCXO output varies with center frequency, even though the FM input signal is constant. With this setup the oscilloscope display of gain allows the designer to optimize the circuit and get immediate results when changes are made.

So as not to disturb the measurement, the peak deviation of the FM at 0V is adjusted to only 100 Hz. This is done by adjusting the ac modulating signal while reading the peak deviation with the 8901A. The 8901A's 300-Hz to 3-kHz audio filters are used to limit the noise. The front-panel dB button is pushed to make the display read in dB. Resolution is better than 0.2 dB, which is more than enough for this application. The 8901A's post-peak-detector filter is selected using the analyzer's special function capability and a completely stable display is produced. For this setup, accuracy is limited by the resolution of 0.2 dB.

Assuming that everything was stable this measurement could be made with a counter by painstakingly plotting f_{out} versus tuning voltage and differentiating the plot. Obviously the modulation analyzer produces a more convenient method of analyzing the gain of this system. However, without the extremely low noise of the modulation analyzer's local oscillator, the deviation required to make this type of measurement might perturb the oscillator and limit the resolution of the measurement.

Automated Tests

Fig. 6 shows an automatic system for testing mobile transmitters, taking advantage of the programmability of the 8901A Modulation Analyzer. An HP 9800 Series Desktop Computer controls the system. Fig. 7 shows a typical output plot.

Assuring Accuracy in Modulation Measurements

by Leslie E. Brubaker

O NE OF THE MOST DIFFICULT PROBLEMS in making very accurate measurements of AM depth or FM deviation is calibrating the measuring instrument. A precisely modulated signal must be generated to serve as a calibration standard. Option 010 for the 8901A Modulation Analyzer solves this problem for the user by providing built-in precise AM and FM modulation standards. The calibrators built into an 8901A that has Option 010 can also be used to calibrate other 8901As that do not have this option.

The AM standard is generated by summing two identical 10.1-MHz signals. When one of the signals is switched on and off at a 10-kHz rate, the result is 33.33 percent AM depth. By internally measuring* any slight difference in the levels of the 10.1-MHz signals the analyzer is able to determine the actual depth to ±0.1 percent accuracy. To further improve the modulation envelope, the rise and fall transitions are smoothed to eliminate ringing that might other-

wise occur in the 8901A's audio filters when this signal is measured.

The FM standard is generated by square-wave modulating a VCO with a nominal 33-kHz peak deviation. By using the internal counter to measure* the upper and lower frequency of this signal, the actual peak deviation is determined to ±0.1 percent accuracy. To prevent ringing, the square wave is modified to a round-edged trapezoid.

Once the actual AM depth and FM deviation have been calculated and stored, the modulation analyzer can be calibrated. With the calibration signal applied to the analyzer's RF input port, the internal demodulators measure the signal being applied, and a pair of percentage

*The 8901A interacts with the AM and FM calibrators to calculate the actual amount of generated modulation, but the AM and FM detectors, which will later be calibrated using the signals generated by the AM and FM calibrators, are not used in the calculation. Because of this indirect method, any counter could determine the FM deviation and any digital voltmeter could measure the two RF levels for the AM depth. The 8901A is simply a convenient instrument to use to find the magnitudes of the calibration signals.
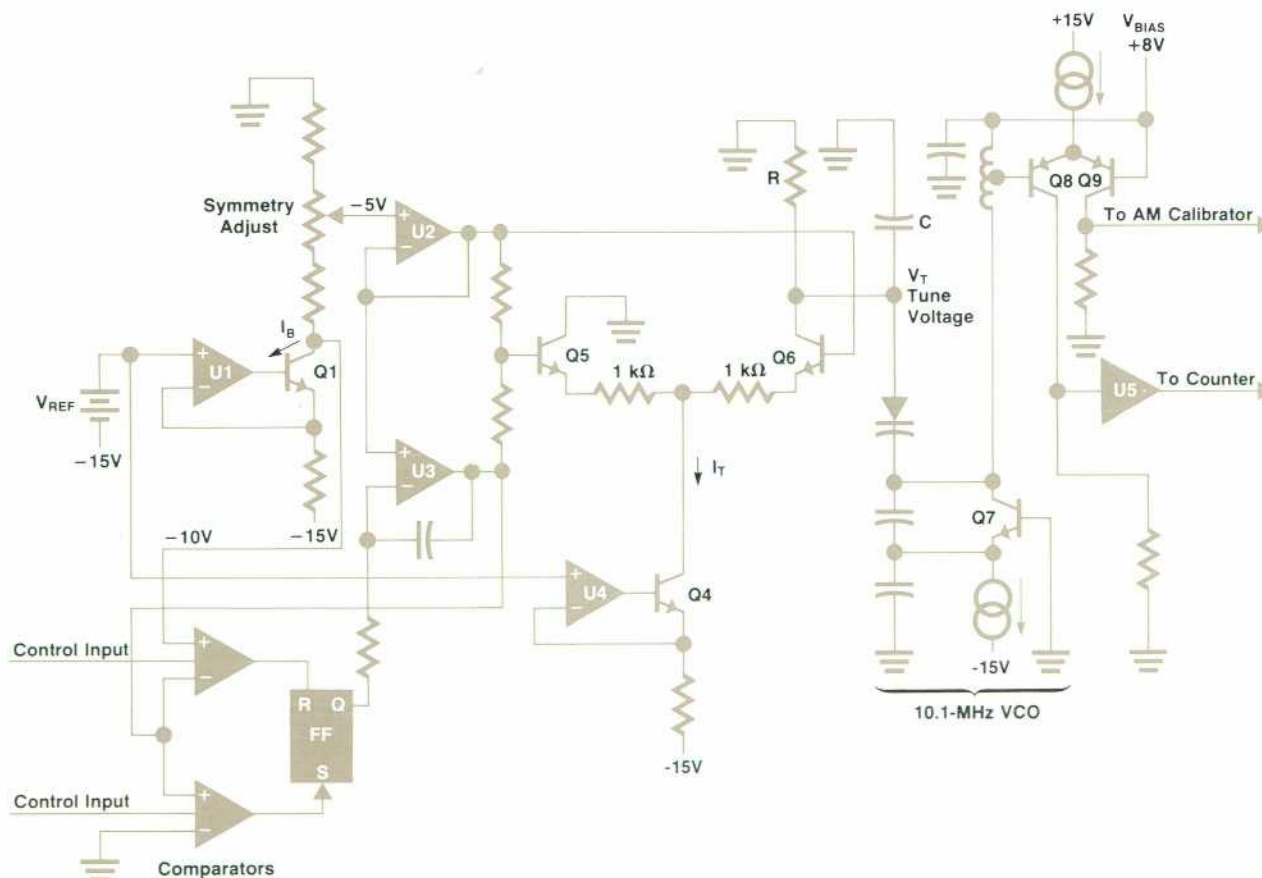

Fig. 1. *FM calibrator circuit diagram. 10.1-MHz oscillator Q7 is switched between two frequencies by the output of differential amplifier Q5-Q6.*

accuracies are determined using the calculated values of AM and FM as references. For example, if the measured value of AM or FM deviates from the calculated value by −0.1%, the display shows 99.90%. The two percentage accuracies, one for FM and one for AM, may be stored and used as scale factors to correct subsequent readings. If commanded to do so by means of a special function command, the analyzer will correct its readings automatically.

Because the modulation standards are internal to the analyzer, there is little need for metrology laboratories to purchase separate calibration standards. Also, because of the technique used, it is easy to verify that the calibrators are operating properly. The two levels in AM and the two frequencies in FM are very accurately determined in a nonmodulating mode. By analyzing the circuit performance, it is possible to show that the amount of modulation generated is the same as that indicated by knowing the nonmodulated states. The error analysis will follow the circuit description.

FM Calibrator

The FM calibrator circuit diagram appears in Fig. 1. The main portions of the circuit are the 10.1-MHz VCO (Q7) and the differential amplifier made up of Q5 and Q6 with its emitter current source U4 and Q4. The oscillator is a varactor-tuned Colpitts type with a tapped inductor for signal extraction. The 10.1-MHz signal is removed by the differential amplifier made up of Q8 and Q9. The collector of Q9 drives the input to the AM calibrator, which looks like a through path to the front-panel output when in the FM calibration mode. The collector of Q8 drives a line receiver with TTL output. This stage acts as a level converter for driving the 8901A's internal counter and provides reverse isolation that keeps counter noise from modulating the VCO. The collector bias of Q7 also supplies base bias for the differential amplifier and reverse bias for the varactor diode. This bias partially temperature compensates drift of the varactor capacitance. Low drift is required to improve accuracy when calculating the deviation of the FM calibration signal, since any drift will directly add to or subtract from the difference measurement.

The oscillator tuning voltage is generated by switching a constant current on and off with Q5 and Q6. Suppose that Q6 is fully on and Q5 is fully off. All the current from Q4 flows through Q6 and none through Q5. This makes $V_T = I_T R$. Since this causes greater reverse bias on the varactor, the frequency is increased to f_U, which is counted and stored. If Q5 then is turned on and Q6 is off, all of I_T flows
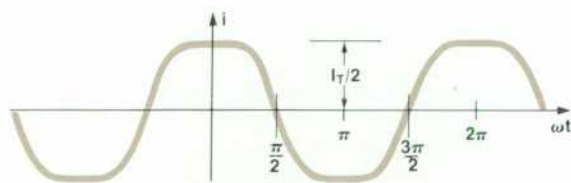
Fig. 2 Current waveform in the FM calibrator differential amplifier. The FM calibration frequency is proportional to this current. The rounded corners prevent ringing in the 8901A's audio filters.
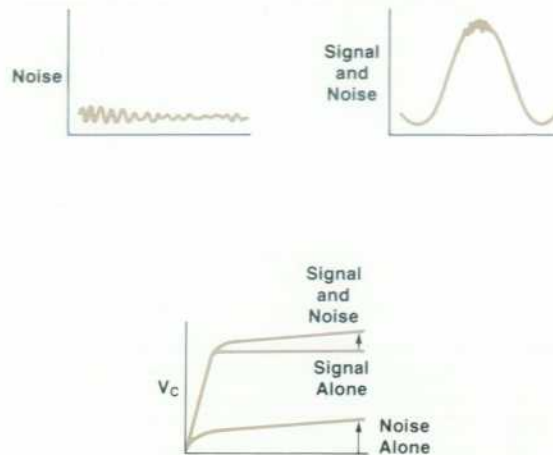
Fig. 3. Effective noise during peak detection.

through Q5 and none through Q6. Then $V_T = 0V$, and a lower frequency f_L is generated and counted. The peak FM deviation is $\frac{1}{2}(f_U - f_L)$. If Q5 and Q6 are allowed to switch at a 10-kHz rate, we then have an RF signal modulated at a 10-kHz rate with a precisely known peak deviation. This is the signal that is used to calibrate the 8901A FM demodulator.

The remaining portion of the calibrator consists of control and reference circuits. A very stable voltage source is used as a reference for the tuning current and a bias current (I_B). U3 integrates the output voltage of a CMOS flip-flop and causes Q5 to be switched on and off. The output of the integrator is sensed by two comparators that control the state of the flip-flop. When in the modulating mode, this circuit causes a triangle waveform at the U3 output. The comparison points are 0V and −10V. I_B flows through three resistors to generate the −10V. There is a voltage tap at half of the −10V that sets the bias of Q5 and Q6 and the reference for the integrator. Thus Q5 and Q6 are always biased halfway between the comparison points even if I_B drifts. This is necessary because otherwise the area above and below the midway point would change, making the indicated plus peak deviation different from the minus peak deviation. The peak-to-peak deviation would still be correct, but it would be more difficult to set the demodulator gain accurately.

To prevent ringing in the audio filters (ringing would result in an improper indication from the 8901A's peak detector), it was necessary to modify the transfer characteristic of the Q5-Q6 differential amplifier. The large (1 kΩ) emitter resistors give the stage a gain of 1.3 and cause the transistors to turn on and off with rounded, rather than sharp corners (see Fig. 2). The rounding also causes a reduction in high-frequency components that would otherwise have to be handled by the system.

FM Error Analysis

Among the errors in the FM calibrator, several sources are negligible: the drift of the VCO with time, the change of the VCO tuning constant with temperature, and the change of α in the current source and differential amplifier transistors.

The counter introduces uncertainty in the measured frequency difference: a resolution of 2 Hz and accuracy of ±1 count gives 0.01% error for 30-kHz deviation. The varactor diodes require RF bypassing at the collector of Q6, giving rise to an RC time constant that must be short compared to the modulation period. Eight time constants are allowed for settling to the final value, giving 0.03% error. Finally, the residual FM of the unmodulated carrier causes an uncertainty in the peak indication (see Fig. 3). When random noise alone is applied to a peak detector, the output voltage will build up to a particular level. If signal alone is applied to the detector, the detector output will rise to the peak of the input. When signal and noise are applied together, some amplitudes of the noise will not be large enough to put additional charge onto the holding capacitor, so the effective amount of noise is reduced. Analysis of the 8901A circuits yields a noise reduction factor of 0.7, and we know that factor with a certainty that will introduce less than 0.025% error. Adding 0.7 times the peak residual FM to the calculated value of the FM from the calibrator gives the correct amount of deviation that the 8901A should indicate when measuring the calibration signal with its FM demodulator. Some of these errors are random while others are always in one direction. By appropriately combining errors, we can arrive at an uncertainty of 0.06% in determining the actual modulation. This is well within the 0.1% specified accuracy.

AM Calibrator

Fig. 4 is the block diagram of the AM calibrator. The unmodulated 10.1-MHz signal from the FM calibrator goes through a limiting stage and then is split into two identical paths containing a buffer, an RF modulator and a common-base stage. The two signals are joined again by summing the collector currents of the common-base stages. From there the signal goes to the front-panel output as the AM calibration signal and to an RF level detector. The output of the detector is monitored by the microprocessor both directly and after $10\times$ amplification for increased resolution on the AM calculation.

The technique of generating AM is similar to that of generating FM. Two known states are precisely determined in a nonmodulating mode, and we then switch between those two states at an audio rate. Assume that I_B is off and that I_A is on. Also, assume that V1 is large enough to turn on D1 and turn off D2. All of I_A will flow through D1 causing $V_A = 0$. When V1 causes D1 to be off and D2 to be on, I_A will cause a voltage (V_A) at the emitter of Q1. In this fashion a square wave of voltage at the RF rate will appear at the emitter of Q1 with amplitude V_A. Notice that this amplitude depends only on the impedance at node A and the value of the dc current source I_A, not on V1 or the diodes (as long as the diodes are being switched on and off by V1). The output of Q1 is amplified and detected. An identical operation occurs for the B modulator when I_B is on and I_A is
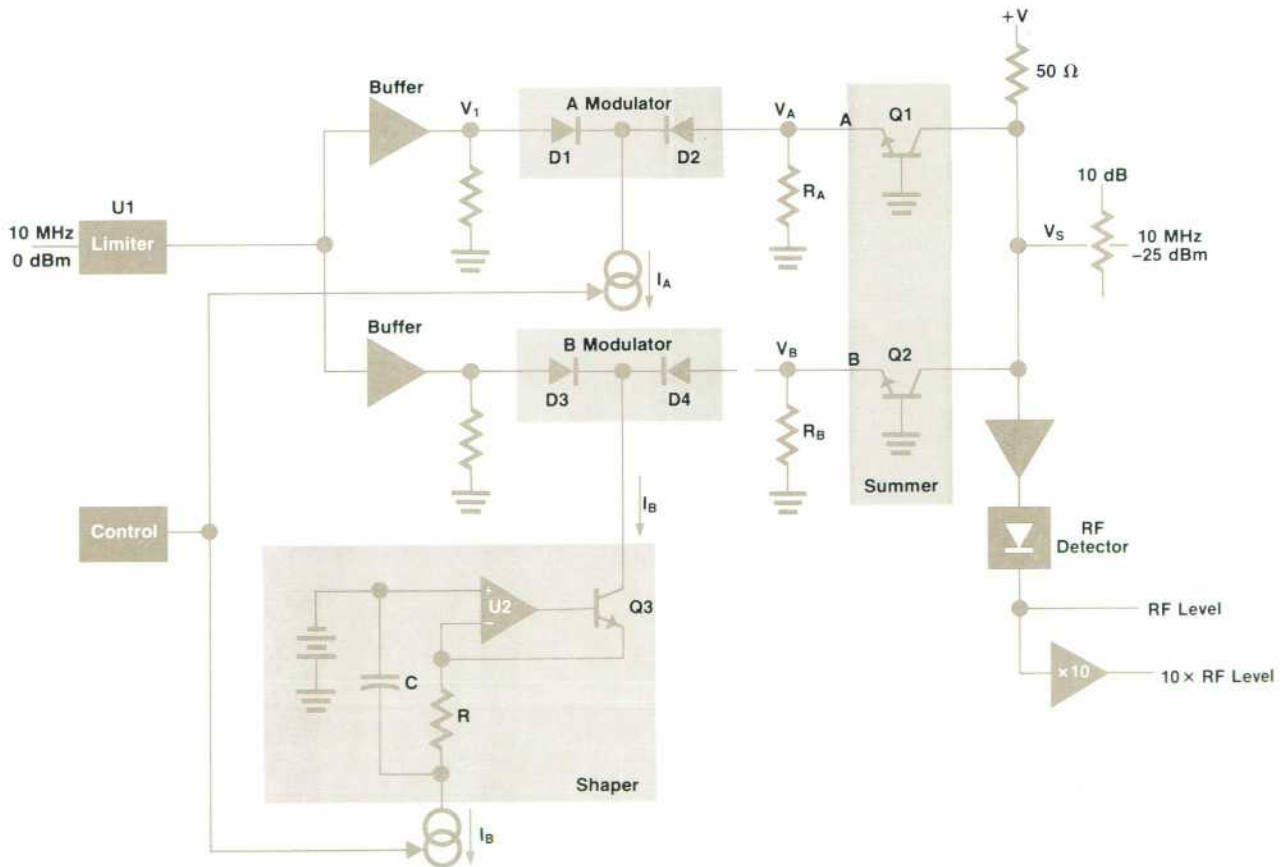

Fig. 4. *AM calibrator circuit diagram. The 10.1-MHz signal from the FM calibrator is split into two identical paths. The signal in one path is switched on and off, and the two paths are then summed.*
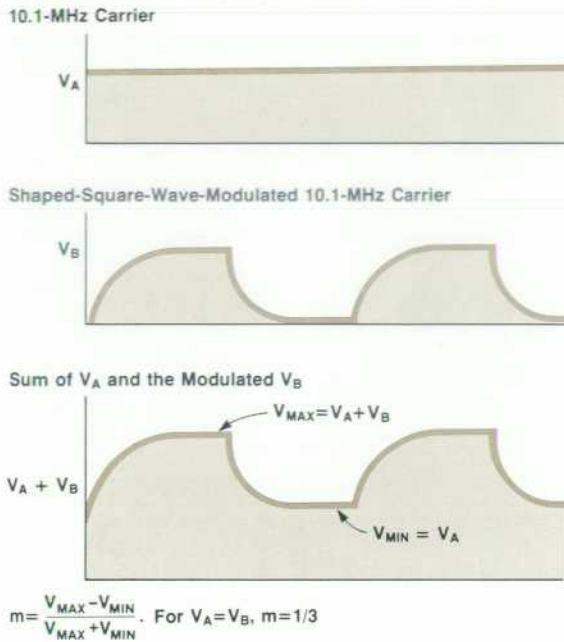
10.1-MHz Carrier

V_A

Shaped-Square-Wave-Modulated 10.1-MHz Carrier

V_B

Sum of V_A and the Modulated V_B

$V_{MAX} = V_A + V_B$

$V_A + V_B$

$V_{MIN} = V_A$

$m = \dfrac{V_{MAX} - V_{MIN}}{V_{MAX} + V_{MIN}}$. For $V_A = V_B$, $m = 1/3$.

Fig. 5. *Generation of the AM calibration signal by summing.*

off, thus generating V_B. If both I_A and I_B are on, then the sum of the two voltages appears at V_S. AM is generated by turning on I_A continuously but switching I_B on and off at an audio rate (see Fig. 5). For one-half period, the output is V_A and for the other half period it is $V_A + V_B$. For $V_A = V_B = V$, the modulation index is $m = 1/3$. The expression for m is:

$$m = \frac{V_{max} - V_{min}}{V_{max} + V_{min}}$$

$$m = \frac{(V_A + V_B) - V_A}{(V_A + V_B) + V_A} = \frac{1}{3 + 2(V_A - V_B)/V_B}$$

We have added the 10× amplifier for V_A and V_B to get increased resolution on the $V_A - V_B$ measurement. This expression shows that, as V_A and V_B become equal, any uncertainty in measuring V_A and V_B becomes insignificant. This removes stringent requirements for the RF detector to be highly linear or to make highly accurate absolute measurements.

The shaping of the I_B current drive to the B modulator is accomplished by driving a current through a parallel RC network using U2 (see Fig. 4). Assume that I_B has been off and that there is no charge on C. There is also no current in R. Now let I_B make a step increase to its maximum value. All of I_B flows through C and starts to build up charge. Since the voltage between the + and − terminals of the operational amplifier must be zero, a current is forced through R, developing a voltage that will equal the voltage on C. This current is the emitter current of Q3, and it reduces the current through C. This process continues until all of I_B flows in Q3. The current in Q3 follows a simple exponential rise. When I_B is turned off, C supplies the current to R giving an exponential decay. Since the size of the square waves of RF developed at node B are directly proportional to I_B, the RF waveform also has an exponential rise and decay.

AM Error Analysis

Once again it is necessary to verify that the amount of modulation generated in the switching mode is the same as that calculated from measurements in the static mode. The first errors that need to be considered are in determining the RF level. The factors involved are the nonlinearity and offset voltage of the RF detector and the internal digital voltmeter, the resolution of the digital voltmeter, and the gain uncertainty of the 10× amplifier. A list of magnitudes and their influence on the uncertainty in knowing m is given below:

A. Nonlinearity of RF detector ≈ 0.07% for 3% change of RF level
B. Offset voltage of RF detector ≈ 15 mV = V_0
C. Gain uncertainty of 10× amplifier ≈ 0.5%
D. Nonlinearity of DVM <0.1%
E. Offset voltage of DVM <1 mV (0.05% for 2V input)
F. Resolution of DVM = 1 Bit (0.05% for 2V input).

Let

X_1 = % error in knowing V_0 = 100%
X_2 = % error in knowing $V_A - V_B$
   = 0.07 + 0.5 + 0.1 + 2(0.05) = 0.77%
X_3 = % error in knowing $V_B = \dfrac{0.015V}{V_B} \times 100 + 2(0.05)$
   = 0.85%

Then the uncertainty in knowing m because of the detector errors is:

$$\frac{\Delta m}{m} = \frac{2}{3} \frac{V_A - V_B}{V_B} \left(\left| \frac{V_0}{V_B} X_1 \right| + |X_2| + |X_3| \right).$$

For example, let V_A be within 3% of V_B. Then:

$$\frac{V_A - V_B}{V_B} \leq 0.03$$

and

$$\frac{\Delta m}{m} = \frac{2}{3}(0.03)(0.75 + 0.77 + 0.85) = 0.047\%$$

We test the value of $(V_A - V_B)/V_B$ for less than 3% difference during each calibration cycle and halt if the test fails. This insures that there will never be more than 0.047% error because of the detectors.

Since the AM calibration technique depends on adding two signals to get a peak value that is the sum of the two signals, the peaks of the waveform must be coincident. Therefore the differential phase shift between the two modulator paths must be small. A typical phase difference of ½° results in approximately 0.001% error.

Since the modulators are not ideal and have stray capacitive coupling around the diodes, they do not totally shut off the signal to the summing amplifier when I = 0. The level of this feedthrough is about 60 dB down from the desired RF level, but since it is at an angle of 90° it has negligible effect on the magnitude of the resultant signal.

Another source of error comes from generating an exponential leading and trailing edge on the modulating current source. There must be enough time for the exponential to reach the same final value as was measured statically. Approximately nine time constants occur from the start to the end of each half period, resulting in 0.012% error.

The change of α in the current source transistors also affects the AM calculation, but in this case it is not negligible because of higher power dissipation. When measuring V_B, Q3 dissipates 35 mW, but when modulating, Q3 has only a 50% current duty cycle, so the dissipation is only 17.5 mW. For the transistor used this results in a 6°C temperature change and a change in α of 0.03%.

Finally, the residual AM measurement must be modified to include the effects of peak-detecting a signal plus noise. The effective peak noise is 0.7 of the measured peak noise, as in the FM case. There is an additional factor to be considered in the AM case because of AGC action in the 8901A's AM demodulator. The AGC loop adjusts the gain of an amplifier to maintain the same average IF level into the demodulator. When V_B alone is applied the gain is 1.5 times higher than when the modulated signal is applied. Notice that the modulated signal goes between V_B and $V_A + V_B$, resulting in a higher average RF level. The reduction of gain while modulating also reduces the noise applied to the demodulator by a factor of 0.8, so the actual amount of noise affecting the calibration is 0.56 times the peak noise

Leslie E. Brubaker
Les Brubaker completed his BSEE studies in 1972 at the University of New Mexico, having previously worked as a technician for several years in Albuquerque. Two years later he received his MSEE degree from Stanford University. An HP employee since 1972, Les has contributed to the 8505A Network Analyzer and the 8482A, 8483A and 8484A power sensors. He designed the calibrators for the 8901A Modulation Analyzer, and designed the 11715A AM/FM Test Source. He's now a production engineering supervisor. Les is an avid tennis player and enjoys radio-controlled flying. He land his wife have a daughter and live in Sunnyvale, California.

measured with V_B only. We know the factor of 0.56 to an accuracy that gives 0.025% uncertainty to the calibration. During calibration, the 8901A should then indicate the value that is calculated in the AM calibrator plus 0.56 times the peak noise.

Combining the above errors results in 0.07% uncertainty in the AM calibration. This is within the 0.1% specification.

Interactive Modulation Analyzer Control

by Paul J. Lingane

THE MICROPROCESSOR-BASED CONTROLLER in the 8901A Modulation Analyzer interacts closely with the hardware subsections of the instrument to create a unified system. The controller also interacts with the user, both through the front panel and by responding to the RF signal applied to the input. Control in the instrument takes place at several levels.

At the functional block level the controller uses various circuits as separate instruments: the RF level detector as a power meter, the local oscillator as a signal generator, and the analog-to-digital converter as a voltmeter. The frequency counter, IF section, and audio section are blocks that can be manipulated in various ways to tune to a signal and make a measurement.

The microprocessor is also used as a circuit component.

Half of the hardware registers and associated circuits in the frequency counter were replaced by software registers. The local oscillator relies on the speed of the microprocessor to control its phase-locked loops; manual control is impossible with the block diagram that was selected.

The controller makes it possible to transfer some complexity of hardware into complexity of software. This tends to improve reliability and flexibility in the use of circuits, allowing more to be done with the same hardware. The local oscillator was designed with that in mind. Also, software for linearization of nonlinear detectors and predistortion of control voltages to the VCO and VCXO allow reductions in the cost and complexity of that hardware.

This close interaction between the controller and the hardware impacts the service strategy, as traditional service
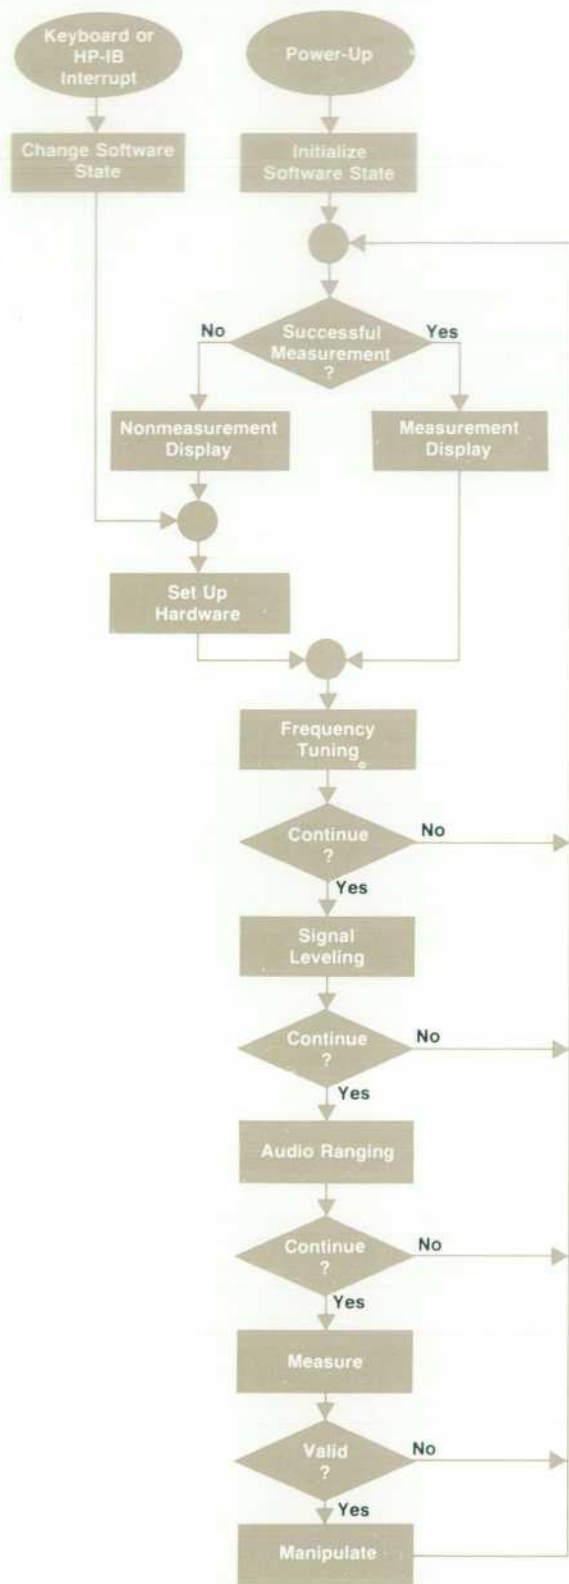
Fig. 1. *8901A software supervisory loop is continuously traversed by the microprocessor-based controller. A measurement is made near the bottom only if all preceding tests are passed in immediate succession.*

procedures are not always effective at identifying problems. The action of the controller can multiply the effect of a circuit failure so it may appear as a collection of symptoms that superficially seem unrelated. I will say more about this later in this article.

Software Structure

An important maxim in the design philosophy of the modulation analyzer was, "Don't mislead the user." Both the software structure and a system of error messages were designed to instill confidence in the measurement.

The 8901A software is structured in a form that we have dubbed the supervisor (see Fig. 1). It is a loop that is continuously traversed, with the measurement made near the end, after checks for proper frequency tuning, proper RF and IF level, and correct audio range. Arithmetic manipulation (e.g. for the ratio function) follows the measurement, and the program then loops back up to display.

The frequency, level, and audio blocks verify that the 8901A is adjusted to make an accurate measurement. A measurement is not made until all of the tests are passed in immediate succession. Let us say, for example, that the RF input level has changed. The signal leveling block therefore reduces the RF attenuation to boost the IF level to an acceptable value. The decision after this block forces the program back to the top of the supervisor, bypassing the measurement for that loop. This reduces interaction between the blocks and aids identification of the nature of the input signal change.

In the example above, perhaps the real situation was that an octave bandswitch was changed on the RF signal generator under test, so now the second harmonic of the desired signal has been tuned into the IF of the 8901A. Since it is lower in level than the fundamental the leveling block boosts the IF level. But the real problem is that the LO is now improperly tuned to the fundamental. Looping back before a new measurement is made gives the frequency tuning block a chance to correct the tuning.

The software interface with the hardware makes use of two concepts called software state and hardware state. The software state is located in 22 bytes of RAM and totally describes the state of the instrument. On power-up, the initialization procedure loads the software state from ROM. Keyboard and HP-IB entry routines modify only the software state and do not affect hardware immediately. The setup block in the supervisor is where the hardware state is made to conform with the software state. Setup is not the only place where hardware is affected; the frequency tuning, leveling, audio ranging, and measurement blocks manipulate the hardware as well.

The keyboard and HP-IB can be thought of as a medium through which the user requests a certain instrument setup. In fact, the actual instrument setup is guaranteed to conform to the keyboard request only at the moment a measurement is taken. The controller may change the instrument hardware at other times to optimize its tuning, leveling, and ranging functions.

In a normal stable measurement cycle the program takes the measurement display branch at the top of the supervisor and so avoids the time overhead associated with the setup block. However, if the program loops back before taking a

measurement, or if an error condition exists, the nonmeasurement display branch will be traversed, thus lighting an appropriate display and going through the setup block. This refreshing of the hardware state when error conditions are sensed contributes to the robustness of the software structure.

Tuning to a Signal

As a measuring instrument the 8901A must be able to cope with input signals whose nature is unknown, except that they must be within the specified frequency, power, and modulation ranges. The controller manipulates the various subsections of the instrument to first understand the nature of the input signal and then to respond effectively to it by changing the setup to make the measurement. An example is the frequency tuning block of the supervisor.

The user simply wants to apply a signal and make a measurement. From the point of view of the modulation analyzer, however, it may be presented with a complex RF spectrum consisting of an RF signal with its harmonics, perhaps drifting in frequency and being modulated with a complex waveform at high FM deviations and large amounts of AM. The modulation analyzer must accurately tune the local oscillator to the signal, range the RF and audio levels, and make a fully settled and stable measurement within a second or two after a signal is applied.

The 8901A is basically an accurately calibrated radio receiver that tunes to an RF signal by mixing a single local oscillator (LO) signal with the RF to produce an intermediate frequency (IF) signal that is demodulated and analyzed. To locate a signal automatically, the local oscillator is quickly swept throughout the entire frequency range of the instrument, looking for a signal to appear in the IF section. Since the LO is a square wave, rich with its own harmonics, it is swept down from high to low frequencies so

that the fundamental of the LO will be the first LO signal to mix with the RF signal.

If there is AM present on the RF signal it is possible that the fundamental of the LO will sweep past the input signal while the RF is deep in the AM trough, so that a harmonic of the LO (following the LO down a bit later in time) might catch the input signal as it comes out of the trough. This would make it appear that the RF signal is at a lower frequency, since the fundamental of the LO would be at a low frequency. So the LO is swept several times, and each time it finds a new signal at a higher frequency it throws out the previously found frequencies. This continues until a signal is found at the same frequency during five separate sweeps. This effectively eliminates the possiblity of LO harmonic mixing. Note that this also has the effect of favoring the highest-frequency RF signal if more than one is present.

To avoid the possibility that the LO's sweep repetition rate might be synchronized with the amplitude modulation rate, causing the RF signal to be consistently missed, a variable delay of up to 10 ms is used after sweeps that fail to find any signal. This introduces jitter into the sweep repetition rate to make it uncorrelated with the input signal AM rate.

The LO is swept down with a ramp signal generated by a current source driving an integrator. When an IF level monitoring circuit senses the presence of a signal in the IF, it instantly turns off the current source thus freezing the LO where it is roughly tuned to the RF signal. The LO is quickly counted with the frequency counter to determine the approximate frequency of the RF signal. However, the 8901A has no image rejection, so the LO can be either above or below the RF signal frequency to produce a signal in the IF. Knowing the LO and IF frequencies will not uniquely identify the frequency of an RF signal unless it is known which image is being used. To resolve this ambiguity the LO is

Fig. 2. *8901A local oscillator was developed using the controller as an essential element.*
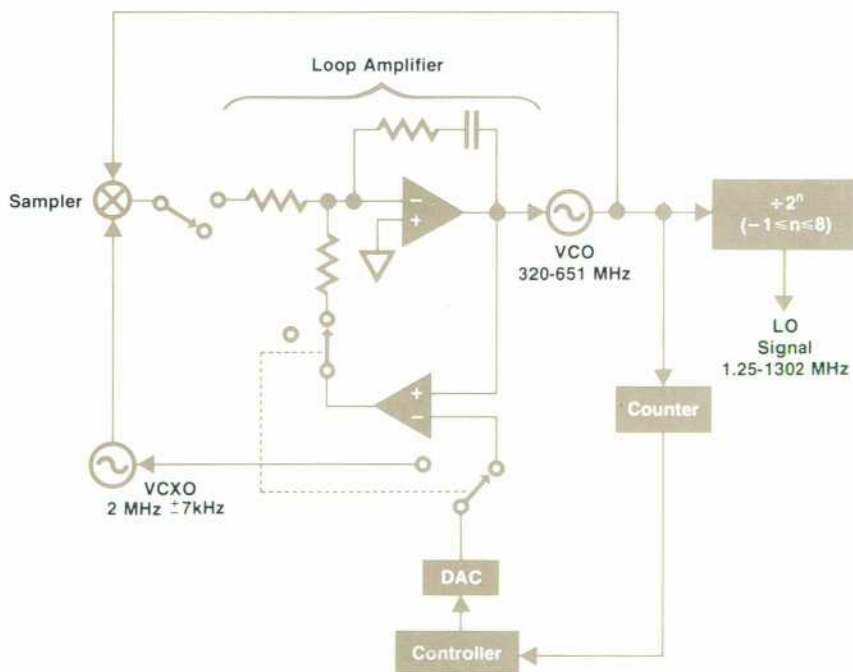
Switches shown connected to pretune the VCO with the DAC

tuned to the frequency where an IF signal appeared, then it is stepped up in frequency while monitoring the IF level until it sees the IF signal level moving down the skirt of the IF low-pass filter. The LO at this point must be above the RF signal frequency, thus identifying the image.

The RF signal frequency is known by accurately counting the LO and IF. This count lasts over 100 ms to gain sufficient resolution, so a drifting RF signal may move out of the IF passband during the count. To avoid this the LO is configured for track mode during the count, so it will always remain tuned to the signal.

Now that a signal has been found and accurately identified as to frequency, call it f, it is still possible that it is actually a harmonic of the desired signal. So the controller directs the LO to tune to f/3 and then looks for the presence of a signal that is higher in amplitude than the signal at f. If this happens then the signal at f is assumed to be the third harmonic of the applied input signal, so the modulation analyzer will tune to f/3 as the fundamental. If the signal at f is not the third harmonic, a similar check is made at f/2 to see if the second harmonic is involved.

Once the fundamental of the input signal is identified, the LO can be properly tuned and the rest of the instrument initialized and set up to take the desired measurement.

This simplified description of one algorithm in the 8901A illustrates the close interaction of the controller with the various subsections of the instrument at the functional block level.

Control of the Local Oscillator

The controller is also closely tied with the hardware at a circuit level. As an example the block diagram of the local oscillator was developed with the controller as an essential element (see Fig. 2).

When configured for low-noise, fixed-frequency operation, the heart of the local oscillator consists of a high-frequency voltage-controlled octave oscillator (VCO) that is locked to a harmonic of a low-noise voltage-controlled crystal oscillator (VCXO). An internal frequency counter is available to monitor the frequency of the VCO. The specification of the two oscillators calls for only a modest degree of linearity between the tune voltage input and the output frequency. This is because the controller can tune the two oscillators by adjusting the voltage control until exactly the desired frequency is attained.

The first step in the tuning procedure is to pretune the high-frequency VCO to approximately the desired frequency. Then the tuning DAC (digital-to-analog converter) is disconnected from the loop amplifier and connected to the low-frequency VCXO, which is then adjusted near the bottom of its range. Meanwhile the VCO has remained where it was because of the sample-and-hold nature of the integrator in the loop amplifier. Now the sampler can be connected to the loop amplifier, and a dc offset voltage at the output of the sampler drives the VCO down in frequency until it acquires lock on the nearest harmonic of the VCXO. The VCXO can then be adjusted to fine-tune the VCO to the proper frequency. Binary dividers and a doubler convert the VCO signal to the desired octave in the 1.25-to-1302-MHz range.

Frequency Counter

The frequency counter chain is partitioned into two parts, roughly four decades of higher-speed counters in hardware and about five decades of lower-speed counters emulated by the controller. A 10-MHz crystal oscillator is divided down to 6.25 kHz to provide the basic timing signal monitored by the microprocessor. The controller times a count by keeping track of this 6.25-kHz time base while accumulating frequency counts at rates up to 3 kHz. When sufficient time has passed to attain the frequency resolution desired, the controller turns off the counter gate, and reads the contents of the hardware decades. With this count and the count accumulated in the software registers, the microprocessor then calculates the frequency of the signal.

Implications for Service

The close interaction between hardware and software in the 8901A has definite implications for the use and servicing of the instrument. Service special functions were developed that can be used as service tools to diagnose problems. Some self-tests are available to check for certain failure modes. Since it is not possible to predict all types of failures, general-purpose software routines were developed that can be used in a variety of situations. One of these is a direct control special function that is called from the keyboard. It can read or write data to any instrument hardware location that is controlled by the microprocessor. A service person can use this to operate instrument circuits manually to check their performance.

A special class of error displays called service errors is available. When enabled by a service technician with a special function, they can reveal information about the internal operation of the instrument. Thus, for example, if the demodulated output is periodically being blanked, enabling service errors may result in an "Error 72" display. This would indicate that the audio overload protection circuit is being tripped.

The local oscillator section required particular attention. To lock the low-noise loop a series of switch openings and closings and DAC movements must be performed within an 8-ms interval. Since a service technician couldn't hope to duplicate this procedure manually, a series of self-tests concentrating on the LO is of value in testing gross operation of the circuit. These can be called from the keyboard,

Paul J. Lingane

Born in Boston, Paul Lingane earned his bachelor's degree at Harvard in engineering and applied physics. Three years later, in 1977, he received a master's degree in electrical engineering and computer science from the University of California at Berkeley. With HP since 1975, Paul was responsible for hardware and software design of the controller in the 8901A Modulation Analyzer. Married and living in Palo Alto, he spends his leisure time cooking, gardening, wine-tasting, bicycle touring, enjoying music, and brewing beer.

with results displayed on the front panel.

Some checks are automatically performed at power-up, with results flashed on four LEDs located within the instrument on the controller board. These include the LO tests just mentioned, as well as checks on the operation of the ROM and RAM of the controller, the internal instrument bus, the front panel, the keyboard, and the frequency counter. Signature analysis is also used in troubleshooting much of the digital hardware.

A clear line is drawn between service special functions and normal instrument operation. When most service special functions are used, normal instrument functions are suspended. When the service special function mode is left to resume normal measurements, all effects of the service functions on hardware are lost. As an example, a direct control special function can be used to activate a particular RF attenuator to check its operation. But once normal measurements are resumed the attenuator setting will revert back to what it was before the service special function was invoked.

Special Signal Source Tests Modulation Analyzer

by Leslie E. Brubaker

BECAUSE THE 8901A MODULATION ANALYZER (see article, page 3) is designed to test the highest-quality transmitters and signal generators, it has to be more stable, noise-free, and distortion-free than these signal sources. Where, then, does one find a signal source good enough to test the 8901A?

The answer is a special-purpose signal source, Model 11715A AM/FM Test Source (Fig. 1). This easy-to-use, low-distortion AM and FM source has extremely flat modulation characteristics. Although intended for verifying 8901A performance, it is versatile enough to be useful for other purposes if its restricted frequency coverage matches a particular application. It provides especially high-quality FM in the 88-to-108-MHz frequency range.

While the 8901A's accuracy is determined by its optional built-in AM and FM calibrators, this calibration represents only a single-point performance test. For a complete test, it is also necessary to determine how the 8901A responds when changing the modulation rate (flatness) and the modulation level (linearity). The function of the 11715A is to provide AM and FM signals of sufficient flatness and linearity to characterize the 8901A. In addition, a signal with extremely low residual FM is provided to verify the residual FM of the 8901A at a particular frequency.

The AM output of the 11715A is designed to have very low noise so that it can be used to check the residual AM of the 8901A. The 11715A is also designed to have low distortion, so that the distortion of the 8901A's AM and FM demodulators and audio section can be checked. Finally, the AM source is designed to have low incidental ΦM and the FM source is designed to have low incidental AM.

What's Inside

The block diagram of the 11715A is shown in Fig. 2. The

Fig. 1. Model 11715A AM/FM Test Source provides AM and FM signals of sufficient flatness and linearity to characterize the 8901A Modulation Analyzer.

test source contains two oscillators. A crystal oscillator generates the low-residual-FM output. This oscillator is followed by a full-wave frequency doubler and appropriate filtering. The output residual FM is much less than 0.1 Hz in a 50-Hz-to-3-kHz bandwidth.

The other oscillator, which generates the AM and FM outputs, is a varactor-tuned voltage-controlled oscillator (VCO) with a frequency range of 352 MHz to 432 MHz. The VCO output is buffered, filtered and brought to the front panel. This front-panel output can be modulated at up to 400-kHz peak deviation and 200-kHz rate. Its residual FM at 400 MHz is typically 15 Hz in a 50-Hz-to-15-kHz bandwidth. The VCO output is also applied to an ECL (emitter coupled logic) divide-by-four counter whose output is filtered and brought to the front panel. This output covers the FM broadcast band and allows up to 100-kHz peak deviation. After another division by eight, the VCO output becomes the local oscillator for an amplitude modulator. The output of the modulator can supply both FM and AM carriers from 11 MHz to 13.5 MHz. The FM deviation from this output can be as high as 12.5 kHz at a rate of 10 kHz. The AM can be as high as 99% at rates from 20 Hz to 100 kHz.

Since the FM÷4 and FM÷32 outputs are generated by digital division, the FM linearity between the two outputs is ideal. This property can be used to check the range-to-range accuracy of the audio section of the 8901A Modulation Analyzer. Setting a particular deviation on one range with the FM÷32 output, then changing to the FM÷4 output, should cause the 8901A to show exactly eight times the modulation.

The center frequency control varies the bias on the varactors and thus sets the VCO RF frequency. The audio input is applied to the VCO through a relay and a 30-dB attenuator. The audio is applied to the mixer through the relay, a dc blocking capacitor, and a 20-dB attenuator. The relay and blocking capacitor are positioned within the series branch
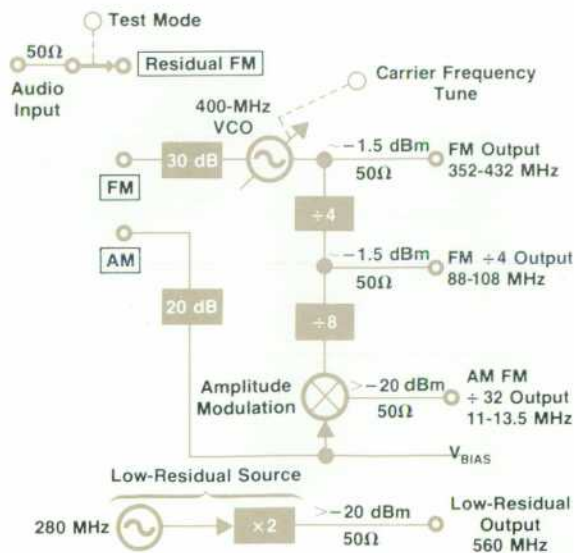

Fig. 2. *Model 11715A AM/FM Test Source has a low-residual-FM crystal oscillator, a VCO to generate the FM outputs, and a modulator (a double-balanced mixer) to generate the AM output.*

of the attenuator, rather than at the input, to minimize the effects of their impedance variations as functions of frequency. At the input the relay would switch a 50Ω source to a 50Ω load, while within the attenuator, it switches between a 50Ω source and a 300Ω load. The effects of loss variations due to relay contact impedance are reduced by about a factor of 3. The FM input is dc-coupled and therefore the 11715A can be phase-locked to an external reference. The dc coupling and extreme flatness of audio response imply very good stereo separation throughout the FM audio range. With a stereo test signal applied to the 11715A's audio input, stereo separation of >60 dB at 1 kHz was measured at 100 MHz using the 8901A Modulation Analyzer.

The power supply uses simple three-terminal regulators mounted on a separate printed-circuit board from the oscillator circuits. This was done to eliminate the effect of the power supply diodes modulating the RF and appearing as 120-Hz sidebands on the VCO output.

Except for the power supply, the entire circuit is mounted within an enclosure that is shock mounted to reduce microphonic noise. (The enclosure also acts as a shield to prevent electromagnetic interference.)

Obtaining VCO Linearity

The design of the 11715A VCO took advantage of one of the features of the 8901A: the RF track mode of operation. Since the 8901A can automatically acquire and track a sweeping RF signal, it is possible to make a real-time plot of VCO FM sensitivity versus VCO center frequency, or differential linearity. The technique is to apply a constant amount of audio to the 11715A's audio input, and a slowly varying signal to the 11715A's frequency tune input and the horizontal input of an oscilloscope. The 11715A RF output is applied to the 8901A and demodulated. The demodulated audio is applied to the vertical input of the oscilloscope. A change in the amplitude of the resulting plot is a change in sensitivity of the VCO. A description of the technique and a typical oscilloscope display appears in the article on page 19 (see Figs. 4 and 5 of that article).

By adjusting component values, a nearly flat envelope, indicating little change in sensitivity, was achieved quite readily. Operating at the flattest portion of the curve results in minimum FM distortion. At 400 MHz the 11715A has less than 0.025% total harmonic distortion (THD) for 400 kHz peak deviation at a 100 kHz rate.

AM Generator

The modulator used to generate AM is a double-balanced mixer. A constant bias is applied to the mixer IF port and the divided VCO signal is applied to the LO port. This up-converts the dc to the LO frequency, causing an RF level at the RF port that is directly proportional to the dc bias. If an audio signal is summed with the dc, the total bias changes at the audio rate. This causes the RF level to vary at the audio rate, and the result is AM. A peak audio equal to the dc bias gives 100% AM.

Since it isn't feasible to build an AM detector and prove its distortion and linearity with a sufficiently high level of confidence to call it a standard detector, it was necessary to verify the 11715A's AM performance with measurements on the AM generating system itself. AM distortion is related

to mixer compression and LO-to-RF isolation. If LO-to-RF isolation is not ideal, at high AM percentages the minimum value of RF is greater than desired. For example, if the LO is 40 dB below the carrier, the minimum value of RF is ideally 0.01 of the carrier level. At 99% AM the desired level is 0.01, but the resulting level is $\sqrt{2}$ (0.01) because of LO feedthrough, which is at 90° to the desired RF. This causes clipping, resulting in distortion and nonlinearity. Further, since the resultant RF is at an angle other than 0°, incidental phase modulation is generated.

Mixer compression causes the peak of the AM to be clipped. A conservative approximation to the relationship between RF level and compression is that compression is reduced as the cube of the RF level reduction. By measuring the compression at a level considerably higher than the operating level, it is possible to predict a worst-case distortion.

For mixer compression of less than 0.05 dB at a −5 dBm RF output level, LO-to-RF isolation of 68 dB and an unmodulated RF level of −17 dBm, the AM distortion is less than 0.05% and the AM linearity is better than 0.08% to 95% AM. Also, the incidental phase modulation is less than 0.005 radian for 50% AM.

Verifying Flatness

FM flatness is determined by the audio circuits between the input and the varactors, while AM flatness is determined by the components between the audio input and the IF port of the mixer. This assumes no frequency response variation in the varactors for FM or in the mixer for AM. By characterizing the components, including the relay and coupling capacitor, the frequency response can be calculated. The result of this model indicates flatness much better than we can measure. Therefore, the 11715A AM flatness specification is closer to our measuring system performance than to the calculated value. An additional measurement is made at 10 MHz to ensure that none of the components is unusually lossy.

Acknowledgments

The author wishes to thank Stuart Carp, who did the original design work on the VCO and crystal oscillators, and Bob DeVries for his excellent product design. Also, thanks to Russ Riley for his assistance with the performance analysis.

SPECIFICATIONS
HP Model 11715A AM/FM Test Source

FM OUTPUTS
FREQUENCY RANGE: 11 to 13.5 MHz at AM FM ÷ 32 output.
    88 to 108 MHz at FM ÷ 4 output.
    352 to 432 MHz at FM output.
FM PEAK DEVIATION: >12.5 kHz, 11 to 13.5 MHz carrier.
    >100 kHz, 88 to 108 MHz carrier.
    >400 kHz, 352 to 432 MHz carrier.
FM DISTORTION <0.025% THD (<−72 dB):

Carrier Frequency	Peak Deviation	Modulation Rate
12.5 MHz	12.5 kHz	<10 kHz
100 MHz	100 kHz	<100 kHz
400 MHz	400 kHz	<100 kHz

FM FLATNESS: ±0.1%, dc to 100 kHz rates. ±0.25%, dc to 200 kHz rates.
INCIDENTAL AM: <0.08% (100 MHz carrier; <50 kHz peak deviation; 1 kHz rate; 50 Hz to 3 kHz bandwidth).
AM OUTPUT
FREQUENCY RANGE: 11 to 13.5 MHz at AM FM ÷ 32 output.
AM DEPTH: to 99%.
AM DISTORTION: <0.05% THD (<−66 dB) for 50% AM; 20 Hz to 100 kHz rates.
    <0.1% THD (<−60 dB) for 95% AM; 20 Hz to 100 kHz rates.
AM FLATNESS: ±0.1%, 50 Hz to 50 kHz rates. ±0.25%, 20 Hz to 100 kHz rates.
INCIDENTAL ΦM: <0.008 rad peak (12.5 MHz carrier; 50% AM; 1 kHz rate; 50 Hz to 3 kHz bandwidth).
AM LINEARITY: ±0.1%, ≤95% AM. ±0.2%, ≤99% AM.
RESIDUAL AM: <0.01% rms, 50 Hz to 3 kHz bandwidth.
LOW RESIDUAL OUTPUT
RESIDUAL FM: <3 Hz rms, 50 Hz to 3 kHz bandwidth.

General
LINE VOLTAGE: 100 or 120 Vac +5% −10% 48 to 440 Hz. 220 or 240 Vac +5% −10% 48 to 66 Hz.
POWER DISSIPATION: 40 VA maximum.
CONDUCTED AND RADIATED ELECTROMAGNETIC INTERFERENCE: Conducted and radiated interference is within the requirements of methods CE03 and RE02 of MIL STD 461A, VDE 0871 (Level B), and CISPR publication 11.
CONDUCTED AND RADIATED ELECTROMAGNETIC SUSCEPTIBILITY: Meets the requirements of methods CS01, CS02, and RS03 (1 volt/metre) of MIL STD 461A dated 1968.
NET WEIGHT: 4.4 kg (9.5 lb) nominal.
DIMENSIONS: 102 mm H × 212 mm W × 444 mm D (4.0 × 8.4 × 17.5 in) nominal.
TEMPERATURE:
    OPERATING: 0 to 55°C.
    STORAGE: −55 to 75°C.
PRICE IN U.S.A.: $1550.
MANUFACTURING DIVISION: STANFORD PARK DIVISION
    1501 Page Mill Road
    Palo Alto, California 94304 U.S.A.

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