

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.

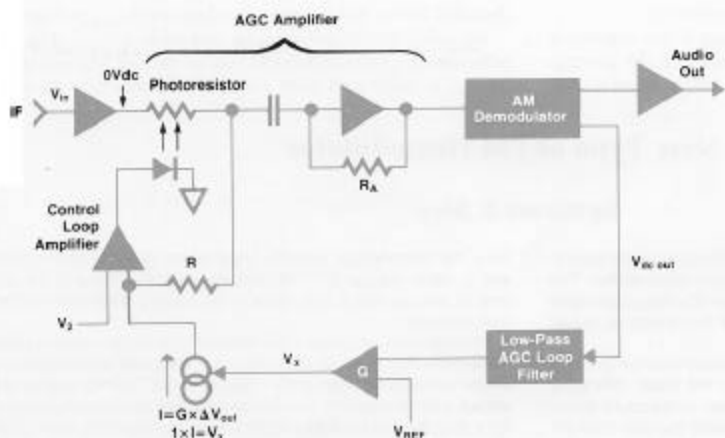


Fig. 11. Automatic gain control (AGC) circuit provides a constant-level IF signal to the AM demodulator. A low-distortion photoresistor is the gain control element.

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 $GK(dV_{out}/dV_c) = GK V_{in}$. 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

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

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

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

An exponential with $V_c \propto e^{K_2 V_x}$ 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 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- μ s 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 μ s is equivalent to a 212-Hz low-pass filter with one pole.)

When FM 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 FM 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. 14. Peak detector is simple but performs well.

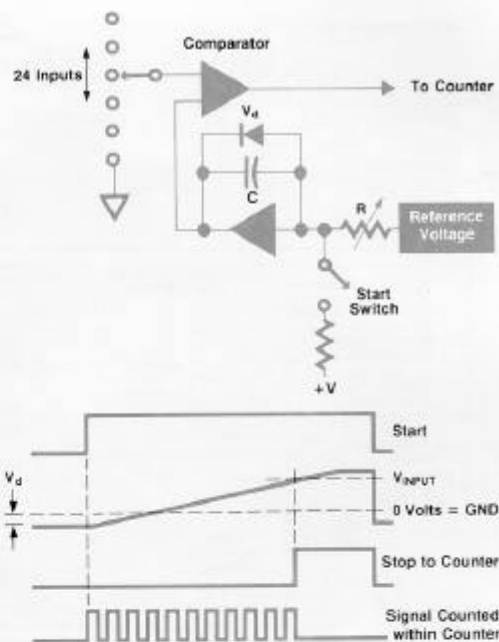


Fig. 15. Analog-to-digital converter is used to measure many voltages in the modulation analyzer. It is a single-slope converter, counting a clock signal while a constant-slope ramp voltage rises to the level of the unknown voltage.

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 (4 μ F). 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

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 I/O 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 analysis³ 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 troubleshoot 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).