Cover Feature

Integration arms receiver with low noise floor

A measurement receiver that combines the features of a sensitive power meter, a modulation analyzer, and an audio analyzer performs power measurements down to −127 dBm. Some new tricks were needed to keep the noise at bay.

Mike McNamee, Development Engineer, Hewlett-Packard Co., 1620 Signal Dr., TAF C-34, Spokane, WA 99220

The model 8902A Measuring Receiver offers an alternative to the existing avenues of low-level power measurement. It handles input signals from 150 kHz to 1300 MHz and combines the features of a wide-dynamic-range (157-dB) power meter, a tuned modulation analyzer, and an audio analyzer.

Such an instrument is useful in a variety of applications. The 8902A's primary purpose is to calibrate signal generators, but its combinaton of functions allows it to replace larger, more complex multicomponent systems. It can also serve as an attenuation measurement center with a 127-dB dynamic range. Its relative power measurement accuracy is not limited by the need to make an absolute reference calibration. And in transmitter testing, the model 8902A can replace many of the components typically used to measure power, modulation, and audio distortion.

Perhaps the most valuable feature of the 8902A is its tuned RF level measurement capability, which allows power readings to be taken down to −127 dBm (0.1 μV into 50 ohms), with linearity better than 0.5 dB. The fundamental measurement accuracy derives from a model 436A type power meter which is integrated into the circuitry of the measuring receiver.

A model 8901A modulation analyzer has also been integrated into the 8902A. It measures AM depths to 99 percent at rates as high as 100 kHz. The modulation analyzer can also measure FM deviations as high as 400 kHz at rates up to 200 kHz, both with an accuracy of better than ±1 percent. The third type of modulation handled by the modulation analyzer section is phase, at peak deviations up to 400 radians and rates up to 20 kHz, with an accuracy of ±3 percent. Demodulated signals are supplied to a front-panel RF connector for further processing.

Audio measurements of either the demodulated RF signal or an external input

1. The 8902A Measuring Receiver *works in conjunction with the 11722A sensor module to make power measurements down to −127 dBm.*

can be made with the audio analyzer section. This section features audio frequency counting of signals at better than 50-kHz rates, with six-digit resolution. It provides for true RMS level measurements of signals up to 50 kHz. In addition, it has the capability of analyzing distortion levels down to −54 dB at the standard calibration frequencies of 400 and 1000 Hz.

Single-conversion receiver

The 8902A is basically a single-conversion receiver with detectors at several key locations (Fig. 2). In most cases, a measurement is made by connecting a model 11722A sensor module to the signal under test. The sensor module contains a thermocouple sensitive to relative power levels; an input signal can be routed into this thermocouple or, through a specially designed RF cable, into the input section of the 8902A. Once in the measuring receiver, the signal is attenuated or amplified to an optimum level for downconversion to an IF of either 455 kHz or 1.5 MHz.

Several measurements are made in the receiver's IF section. A signal is simultaneously amplitude and frequency (or phase) demodulated, and the recovered audio signal is made available at front and rear panel output ports. High-pass and low-pass filters, as well as selectable de-emphasis filters, can be switched in when signal processing is required. And depending on the measurement, peak, RMS, and average detectors are available.

Another component in the IF section

is a new tracking synchronous detector. It provides very linear relative level measurements under extremely noisy conditions. The detector has a wide dynamic range, using 60 dB of variable IF gain as well as the available RF gain and attenuation to achieve that range.

The 455-kHz frequency of the 8902A was chosen to allow the use of common components and electronic switching. As a result, a high degree of accuracy in the IF gain steps was possible, and it was achieved without expensive, elaborate attenuation standards.

All measurements performed by the 8902A are ultimately made by the frequency counter section. The RF frequency is indirectly read by counting the LO and IF signals and determining the difference frequency, LO − IF. The frequency counter, together with a voltage-to-time converter, doubles as a voltmeter, measuring the outputs of the internal detectors. The voltages are converted into a result that is displayed, with the appropriate units of measurement, on the 8902A's front panel.

The dynamic range of the 8902A is broken down into several measurement ranges. Two measurement modes are used: RF power (which relies on the power sensor in the model 11722A) and tuned RF level (which depends on the IF tracking synchronous detector). The calibration factors for the sensor in the 11722A module are printed on the back of the device. These can be entered into the memory of the 8902A for automatic correction of power sensor frequency response.

The internal power meter is set to the zero point and calibrated by connecting the 11722A to a 0-dBm power reference output from the 8902A's front panel. When an operator presses the "zero" key, the 11722A disconnects the power sensor and the 8902A subtracts any offset in the power metering circuitry under zero-input conditions. When the operator presses the "calibrate" key, the power sensor is switched in and the reference output is turned on. When the "save cal" key is engaged, the output voltage from the power metering circuitry is read by the measuring receiver's internal voltmeter, and the instrument controller equates that voltage to a power level of 0 dBm. From that point on, all readings will be compared to this reference calibration point, and all subsequent measurements can be traced to this initial calibration. (This accuracy is essentially equal to that of the model 436A power meter.)

Ranges rely on TSD

The tuned RF level ranges rely on the tracking synchronous detector (TSD). Each range is obtained by extending the

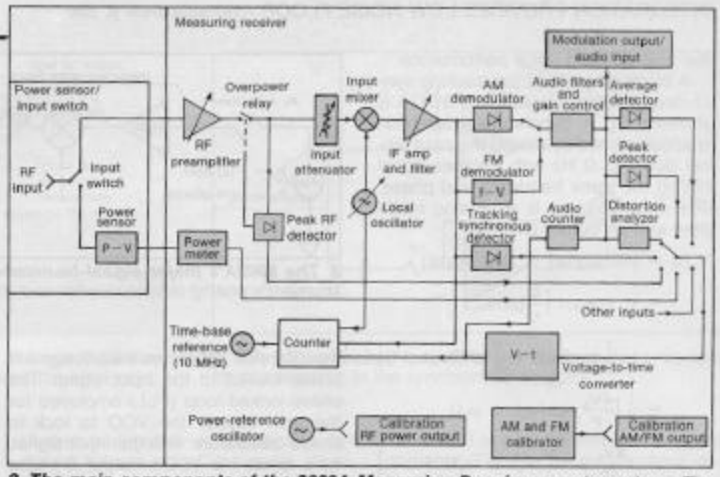

2. The main components of the 8902A Measuring Receiver *are shown here. The newly developed tracking synchronous detector appears in the center.*

TSD's basic 15-dB range via 60 dB of IF gain in 10-dB steps. The three overlapping ranges result from three different gain/attenuation configurations in the RF section. The highest range corresponds to 50 dB of RF attenuation, the middle range to 10 dB of attenuation, and the lowest range represents 24 dB of RF gain. This gain is achieved with a broadband, thick-film microcircuit amplifier specifically designed for low input SWR or high input return loss and minimum noise figure.

A high degree of relative accuracy is maintained within each range. The nonlinearity over the 15-dB range of the TSD is only about 0.003 dB, while the 10-dB IF gain step accuracy is approximately ± 0.02 dB.

The overlapping of the RF and tuned RF bands allows for cross-calibration between the ranges. This calibration information is used to achieve absolute calibration to the 0-dBm reference.

In order to calibrate the tuned RF level ranges for measurements to −127 dBm, the following procedure is practiced. First, the RF power measurement range is calibrated with the 0-dBm reference output. Then, the source under test, perhaps a signal generator, is set to the frequency of interest at a level lying in the overlap region between the RF power range and the first tuned RF range. Next, the 11722A sensor module is connected to the source under test, and the tuned RF level mode is selected. The 8902A will automatically tune to the signal. The operator at this point presses the "calibrate" key, which switches the 8902A to the RF power mode. It takes a calibrated level reading using the power sensor, switches back to the tuned RF level mode, and takes another level reading using the tracking synchronous detector. (Through all this, the source under test should be maintained at a constant output power level.) The 8902A then defines the tracking detector output to correspond to the power level just read by the power sensor. Now, the first tuned RF level range is absolutely calibrated. All absolute RF and IF losses are effectively calibrated out at that frequency. In addition, any mismatch errors due to complex impedance differences between the two internal configurations of the power sensor module are also calibrated out.

The other tuned RF ranges are calibrated similarly. Each time a lower range is calibrated, the source under test is simply set to an output level that corresponds to the overlap region between the two ranges of interest.

The calibration data that are obtained by this method are valid for signals within 5 percent of the calibration frequency.

Considering instrument noise

Attempts to accurately measure signal levels as low as −127 dBm are invariably complicated because of the degradation of the predetection signal-to-noise ratio (SNR) by thermal noise. Figure 3 illustrates the major noise-influencing components of the 8902A receiver. The basic level of kTB noise power at −174 dBm/Hz results in an input SNR of 47 dB in a 1-Hz bandwidth. This SNR decreases by the 8902A's double-sideband noise figure of 10 dB, which includes noise from the front end as well as image noise from the single-balanced mixer. An equivalent noise bandwidth of 30 kHz in the IF further decreases the SNR by a value of 10log(30 kHz) or 45 dB. As a result, the SNR in the IF is −8 dB.

Because the SNR is so low, standard signal detection methods are meaningless. For this reason, the synchronous detector unit was developed to provide

(continued on next page)

the best possible noise performance.

A block diagram of the tracking synchronous detector used in the 8902A is shown in Fig. 4. Synchronous detection is accomplished by mixing the input signal down to 0 Hz with another signal having the same frequency and phase. The mixer involved is an analog multiplier with an output given by

$$V_o(t) = (\text{RF signal}) \times (\text{LO signal})$$

$$= V_s \cos\omega_c t \left[\frac{4}{\pi}(\cos\omega_c t - \frac{1}{3}\cos 3\omega_c t + \frac{1}{5}\cos 5\omega_c t \ldots)\right]$$

$$= \frac{1}{2}\left[\frac{4V_s}{\pi}\cos(\omega_c t \pm \omega_c t) - \frac{4V_s}{3\pi}\cos(3\omega_c t \pm \omega_c t) + \ldots\right]$$

$$= \frac{2V_s}{\pi} + \frac{2V_s}{\pi}\cos 2\omega_c t - \frac{2V_s}{3\pi}\cos 2\omega_c t$$

$$+ \text{higher frequency terms} \quad (1)$$

This output is processed through a low-pass filter, leaving $2V_s/\pi$, a DC component proportional to the input signal level. If the input signal happens to be amplitude modulated, the demodulated audio will also appear at the output of the synchronous detector.

At high input SNRs, detection linearity is determined by the switching multiplier. As a result, nonlinearities as small as approximately 0.003 dB for each 10-dB input change are possible. Complications arise, however, when the SNR is reduced; in this case, significant nonlinearities can result.

Most of the TSD's circuitry is used to generate the high-level LO input signal required by the mixer. This signal must be in phase with the test input signal and must be capable of tracking an input that may be drifting over several kilohertz. In order to accomplish this, the VCO that generates the LO signal is phase locked to the input signal. The phase-locked loop (PLL) employed for the task causes the VCO to lock in phase quadrature with the input signal. As a result, the VCO's output is phase shifted by 90 degrees in order to create the proper zero-degree phase relationship between the LO and the input signal. The synchronous detector error, as well as the overall loop response, is determined by the PLL filter.

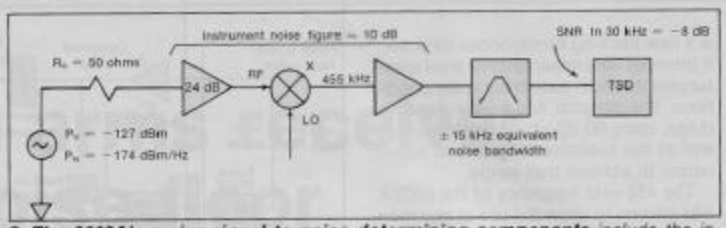

3. The 8902A's major signal-to-noise-determining components *include the instrument's analog downconverter and the tracking synchronous detector (TSD).*

Synchronous detection offers considerable improvement in noise performance over conventional envelope detection techniques. For example, assume an input signal consists of an unmodulated carrier at 455 kHz surrounded by ±15 kHz additive noise. The noise can be separated into two independent in-phase and quadrature components: $\eta_I(t)$ and $\eta_Q(t)$. These components are uncorrelated stochastic signals, limited in bandwidth to 15 kHz. As a result, the composite input signal looks like:

$$V_i(t) = A\cos\omega_c t + \eta_I(t)\cos\omega_c t - \eta_Q(t)\sin\omega_c t \quad (2)$$

To simplify this form, consider an ideal LO signal, $2\cos\omega_c t$, and perform basic linear analog multiplication. Equation 2 then becomes:

$$V_o(t) = 2\cos\omega_c t\left[A\cos\omega_c t + \eta_I(t)\cos\omega_c t - \eta_Q(t)\sin\omega_c t\right]$$

$$= A\cos(\omega_c t \pm \omega_c t) + \eta_I(t)\cos(\omega_c t \pm \omega_c t) - \eta_Q(t)\sin(\omega_c t \pm \omega_c t)$$

$$= A + \eta_I(t) + \left[A + \eta_I(t)\right]\cos 2\omega_c t - \eta_Q(t)\sin 2\omega_c t \quad (3)$$

This result highlights two advantages to the use of synchronous detection. For one thing, the noise is additive at the output. It isn't multiplied onto the input signal as in envelope detection. As a result, the desired signal can always be recovered. Output noise can be reduced whenever necessary simply by adding low-pass filtering to the output. It should be noted, however, that filtering adds delay to the response time and should be used only when the noise level absolutely must be lowered. The second advantage in using synchronous detection is that only the in-phase noise component is translated to DC; the quadrature component is rejected. When downconverted, the noise sidebands on each side of the input carrier "fold over" and add in power level, since the two sidebands are uncorrelated. In contrast, the peak value of the input signal is translated directly to the output, due to the synchronization of the LO with the carrier. The power in the desired output component is related to the peak of the

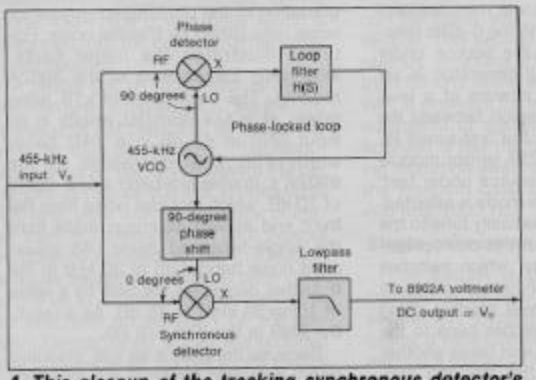

4. This closeup of the tracking synchronous detector's hardware *reveals how the VCO locks in phase quadrature to the 455-kHz input signal.*

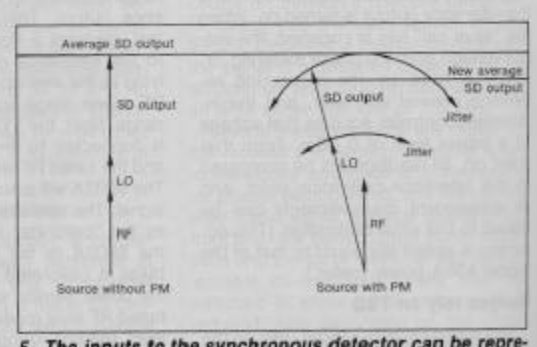

5. The inputs to the synchronous detector can be represented as phasors. *The ideal case (left), a source with no phase modulation, is compared to an actual case (right).*

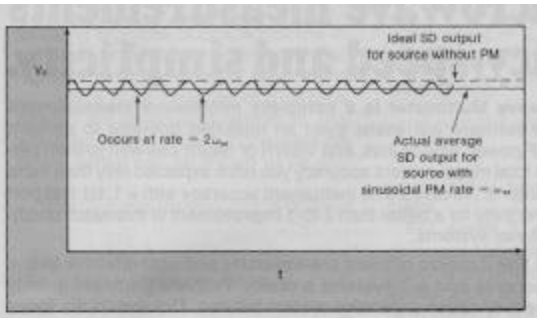

6. Phase modulation on the output of the synchronous detector *causes a level that fluctuates below the ideal case.*

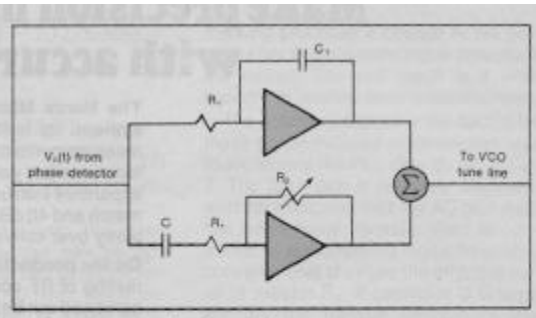

7. This phase-locked loop filter *was implemented to correct for phase errors in the synchronous detector.*

input signal voltage. The undesired noise output power is related to the RMS noise input voltage.

Unfortunately, these two advantages hold only for ideal synchronous detection, where the LO signal is exactly in phase with the input signal. Most problems associated with implementing synchronous detectors stem from the inability of the PLL to provide an LO signal with exactly the right phase.

One of the realities faced by designers of synchronous detectors is the unavoidable existence of residual phase modulation on the input signal. The synchronous detector is specifically designed for use with CW signals. When there are fluctuating phase conditions at the input signal, the PLL will attempt to follow and track out the phase modulation. If the bandwidth of the PLL were large enough to accept all the phase-modulated sidebands contained with the input signal, the VCO would track the incoming signal's phase perfectly, and the two inputs of the synchronous detector would be exactly in phase. However, due to noise constraints, the bandwidth of the PLL must be narrow.

The synchronous detector inputs can be viewed as phasors that are multiplied together (Fig. 5). The synchronous detector's output is proportional to the cosine projection of the input phasor on the LO phasor. The maximum output occurs when the two phasors line up exactly. The slight PLL mistracking caused by residual input phase modulation produces the relative phase jitter between the inputs. This jitter causes the average output level to fluctuate below the ideal maximum output level.

The effect of having the actual synchronous detector output at a lower than ideal level can be illustrated by the

(continued on p. 107)

INTEGRATION PROVIDES LOW NOISE FLOOR *(continued from p. 105)*

following example. Suppose that the input signal has a small amount of sinusoidal pulse modulation at a rate of ω_m, with a peak deviation of β radians. Relying on an analysis similar to that which led to Eq. 1, the synchronous detector output can be shown to be

$$v_o(t) = \frac{2V_s}{\pi} \cos\left[\beta\cos\omega_m t - \beta'\cos(\omega_m t + \theta)\right] + \text{higher frequency terms} \quad (4)$$

where β is the equivalent peak LO phase deviation in radians, and θ is the phase error between the two inputs. If ω_m is much less than the PLL bandwidth, the VCO will track the input signal almost perfectly. This ideal case produces a synchronous detector output of $2V_s/\pi$.

When ω_m is much greater than the PLL bandwidth, the VCO will track very little of the incoming phase modulation, and the detector output can be approximated by:

$$v_o(t) \simeq \frac{2V_s}{\pi} \cos(\beta\cos\omega_m t) \quad (5)$$

This can be expanded with the Bessel series representation of

$$\cos(x\cos\phi) = J_0(x) - 2J_2(x)\cos 2\phi + 2J_4(x)\cos 4\phi \ldots \quad (6)$$

which yields

$$v_o(t) = \frac{2V_s}{\pi}\left[J_0(\beta) - 2J_2(\beta)\cos 2\omega_m t + \ldots\right] \quad (7)$$

where the J terms are Bessel coefficients.

For small β deviations, the synchronous detector output can be approximated by the first two terms shown above, a DC and an AC term. Since the term $J_0(\beta)$ is less than or equal to 1, which indicates that suppression of the DC term is taking place, there will be errors for inputs with phase modulation, as shown in Fig. 6.

If the phase errors as a percentage of DC output were constant, they would cancel in any relative measurement. Unfortunately, the percentage of error changes with the input level. Because of the design of the PLL phase detector, the equivalent loop bandwidth is proportional to the input level. This changing bandwidth causes different amounts of input phase modulation to be tracked out by the VCO as the input level changes. Assuming that the input signal's spectral purity remains constant as the signal level varies, the change in PLL tracking produces a change in the percent error in the synchronous detector's DC output. The end result is a nonlinearity in relative level measurements.

The solution adopted in the 8902A for these phase-induced nonlinearities was to implement the PLL filter shown in Fig. 7. The approach is relatively standard, with the exception that the AC gain path has been made variable. Gain is controlled by a multiplying digital-to-analog converter that changes the effective value of resistor R_2. If capacitor C is large enough, the transfer function of this loop filter will be:

$$H(s) = -\left[\frac{R_2}{R_1} + \frac{1}{sC_1R_1}\right] \quad (8)$$

This loop filter allows the PLL bandwidth to be held constant, regardless of input signal level, thereby eliminating the PM-induced error. The second major problem that must be overcome when using a tracking synchronous detector is additive noise at the input. This error will occur regardless of the spectral purity of the input carrier. It depends only on the input additive noise level. The error increases as the input SNR decreases.

Additive noise errors are caused by phase modulation of the VCO. In a

(continued on next page)

mathematical analysis, assume that the input with the additive noise looks like

$$v_i(t) = V_s\cos\omega_c t + \eta_I(t)\cos\omega_c t - \eta_Q(t)\sin\omega_c t \quad (9)$$

while the VCO signal appears as

$$2\sin(\omega_c t + \theta_o(t))$$

Since the phase detector is essentially a switching mixer like the synchronous detector, the phase detector output can be given by:

$$v_o(t) = V_s\sin\theta_o(t) + \eta_I(t)\sin\theta_o(t) - \eta_Q(t)\cos\theta_o(t) \quad (10)$$

which does not include higher-frequency terms removed by the loop filter. In order to find the error caused by the VCO's phase modulation, the term $\theta_o(t)$ must be determined. Unfortunately, due to intermodulation (IM) in the PLL between the VCO phase modulation quadrature sidebands, and in-phase and quadrature components on the incoming additive noise, this determination is not straightforward. The best solution lies in a statistical representation of the magnitude of $\theta_o(t)$. Previous analyses have shown that a relationship exists between $\overline{\theta_o(t)^2}$ and the SNR in the loop[1]:

$$\overline{\theta_o(t)^2} \simeq \frac{1}{2SNR} \quad (11)$$

where SNR is defined as $\frac{P_s}{\eta 2B_L}$. P_s is the signal power, η is the noise spectral density, and B_L is the equivalent PLL bandwidth.

In order to understand additive input noise, assume that the VCO phase modulation is represented by $\theta_o(t)$. By definition, this phase modulation is in phase quadrature with the VCO carrier signal. The phase-modulated signal mixes in the synchronous detector with the incoming signal and additive noise. If the VCO noise were correlated with the incoming quadrature noise component, a DC component at the synchronous detector's output would result. This component would add to the DC output caused by the VCO carrier mixing with the incoming carrier, causing a positive DC error.

In practice, however, almost the opposite occurs. The synchronous detector's DC output is actually lower than that for an incoming signal of the same level with no additive noise. Due to the IM that occurs in the PLL, the VCO phase modulation is uncorrelated with the incoming quadrature noise, and the synchronous detector's input phasors appear much as they do in the right-hand side of Fig. 5. Due to the relative jitter between the phases of the two synchronous detector inputs, the average detector DC output is suppressed.

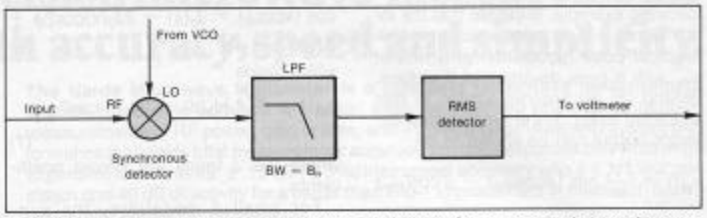

8. This is one technique for correcting additive noise errors *in the synchronous detector. It relies on using an RMS detector as an analog computer.*

Since the detector's DC output is proportional to the average cosine projection of the input phasor on the VCO phasor, the detector output error can be calculated from the SNR in the loop. The procedure is as follows:

$$v_o(t) \propto \overline{\cos\theta_o(t)} \simeq \cos\theta_{oRMS} \quad (12)$$

$$\theta_{oRMS} = \sqrt{\overline{\theta_o{}^2}} \simeq \sqrt{\frac{1}{2SNR}} \quad (13)$$

$$\text{Error (dB)} = 20\log\cos\theta_{oRMS} = 20\log\cos\left[\frac{1}{\sqrt{2SNR}}\right] \quad (14)$$

If this error were a constant percentage of the input signal level, it would cancel, since the level measurements are relative. However, the error is a function of the SNR, which varies with changes in the input level, resulting in a synchronous detector output nonlinearity. If this SNR could be determined, it would then be possible (in the software) to correct the error via Eq. 14.

The error can also be corrected by using an RMS detector in an analog computer, as shown in Fig. 8. This detector is based on the known synchronous detector output of

$$V_o \propto \overline{\hat{V}_s\cos\theta_o(t)} \simeq \hat{V}_s\cos\theta_{oRMS}$$

$$\cos\theta_{oRMS} \simeq 1 - \frac{\theta_{oRMS}{}^2}{2} \quad (15)$$

It is possible to make the following substitution:

$$\overline{\hat{V}_s\cos\theta_o(t)} \simeq \hat{V}_s\left(1 - \tfrac{1}{2}\theta_{oRMS}{}^2\right) = \hat{V}_s\left(1 - \tfrac{1}{2}\overline{\theta_o(t)^2}\right) \quad (16)$$

But it must be remembered that

$$\overline{\theta_o(t)^2} = \frac{1}{2SNR}$$

thus,

$$V_o = \hat{V}_s\left[1 - \tfrac{1}{2}\left(\frac{1}{2SNR}\right)\right] \quad (17)$$

By substituting

$$SNR = \frac{P_s}{2\eta B_L}$$

the following equivalency is obtained:

$$V_o = \hat{V}_s\left[1 - \tfrac{1}{2}\left(\frac{1}{\frac{2P_s}{\eta 2B_L}}\right)\right] = \hat{V}_s\left(1 - \frac{1}{4}\frac{2\eta B_L}{P_s}\right) \quad (18)$$

Building upon this equation, the RMS detector input (V_i) can be written as:

$$V_i = V_{DC} + V_{AC} - V_s\left(1 - \tfrac{1}{4}\frac{2\eta B_L}{P_s}\right) + \hat{V}_s\frac{V_o}{\hat{V}_s}\,(\text{in } B_N);$$

$$\overline{V_i^2} = \overline{(DC + AC)^2} = \overline{DC^2} + \overline{AC^2} + \overline{2DC\cdot AC} = \overline{DC^2} + \overline{AC^2} \quad (19)$$

where the synchronous detector AC input (the noise voltage) is represented by V_n. By substituting the AC and DC voltage terms into this last equation, the following form results:

$$\overline{V_i^2} = \hat{V}_s{}^2\left(1 - \tfrac{1}{4}\frac{2\eta B_L}{P_s}\right)^2 + \hat{V}_s{}^2\left[\frac{V_o(\text{in } B_o)}{\hat{V}_s}\right]^2 = \hat{V}_s{}^2\left[1 - \tfrac{1}{2}\frac{2\eta B_L}{P_s} + \frac{1}{16}\left(\frac{2\eta B_L}{P_s}\right)^2\right] + \hat{V}_s{}^2\frac{V_n{}^2(\text{in } B_N)}{\hat{V}_n{}^2}$$

$$\simeq \hat{V}_s{}^2\left(1 - \tfrac{1}{2}\frac{2\eta B_L}{P_s}\right) + \hat{V}_s{}^2\left[\frac{P(\text{in } B_N)}{\hat{V}_n{}^2}\right] \quad (20)$$

where P is the average power.

Since

$$\frac{P_n(\text{in } B_N)}{\hat{V}_s{}^2} = \frac{\text{Power In Noise Out Of SD}}{2 \times P_s \text{ In The IF}} = \frac{2\eta B_N}{2P_s} \quad (21)$$

by substituting into Eq. 20, the final result appears as:

$$\overline{V_i^2} = \hat{V}_s{}^2\left(1 - \tfrac{1}{2}\frac{2\eta B_L}{P_s}\right) + \hat{V}_s{}^2\frac{2\eta B_N}{\hat{V}_s{}^2} = \hat{V}_s{}^2\left(1 - \tfrac{1}{2}\frac{2\eta B_L}{P_s} + \tfrac{1}{2}\frac{2\eta B_N}{P_s}\right) \quad (22)$$

This last equation shows that by setting $B_N = B_L$, the last two terms will cancel, leaving $\overline{V_i^2} = \hat{V}_s{}^2$, so that the error due to additive noise has been effectively cancelled. Of course, this technique relies on a fixed value of B_L. Still, this has been already guaranteed by the technique used to correct for phase modulation on the source.

(concluded on p. 119)

INTEGRATION PROVIDES LOW NOISE FLOOR (*continued from p. 108*)

In viewing this technique qualitatively, note that the VCO power is constant. When the PLL transfers phase noise onto the VCO, the power in the VCO carrier is reduced by the amount of power that goes into the phase noise sidebands. Since the synchronous detector output is proportional only to the VCO carrier, the output will be suppressed by this amount. If $B_N = B_L$, the RMS detector will add a voltage based on the same amount of noise power that was taken from the VCO carrier line. P&A: 8902A, $21,000; 11722A, $1900; 90 days. **Inquiries Manager, Hewlett-Packard Co., 1507 Page Mill Rd., Palo Alto, CA 94304. (415) 857-1507. CIRCLE NO. 158**

Acknowledgments

The author would like to acknowledge the enormous contributions of Stuart Carp and Andrew Naegeli of Hewlett-Packard Co. for their work on the theoretical and hardware techniques described in the article.

References

1. F.M. Gardner, *Phaselock Techniques*, 2nd ed. (Wiley, New York, 1979), p. 32.
2. A.B. Carlson, *Communication Systems*, 2nd ed. (McGraw-Hill, New York, 1975).