A LOCK-IN PRIMER

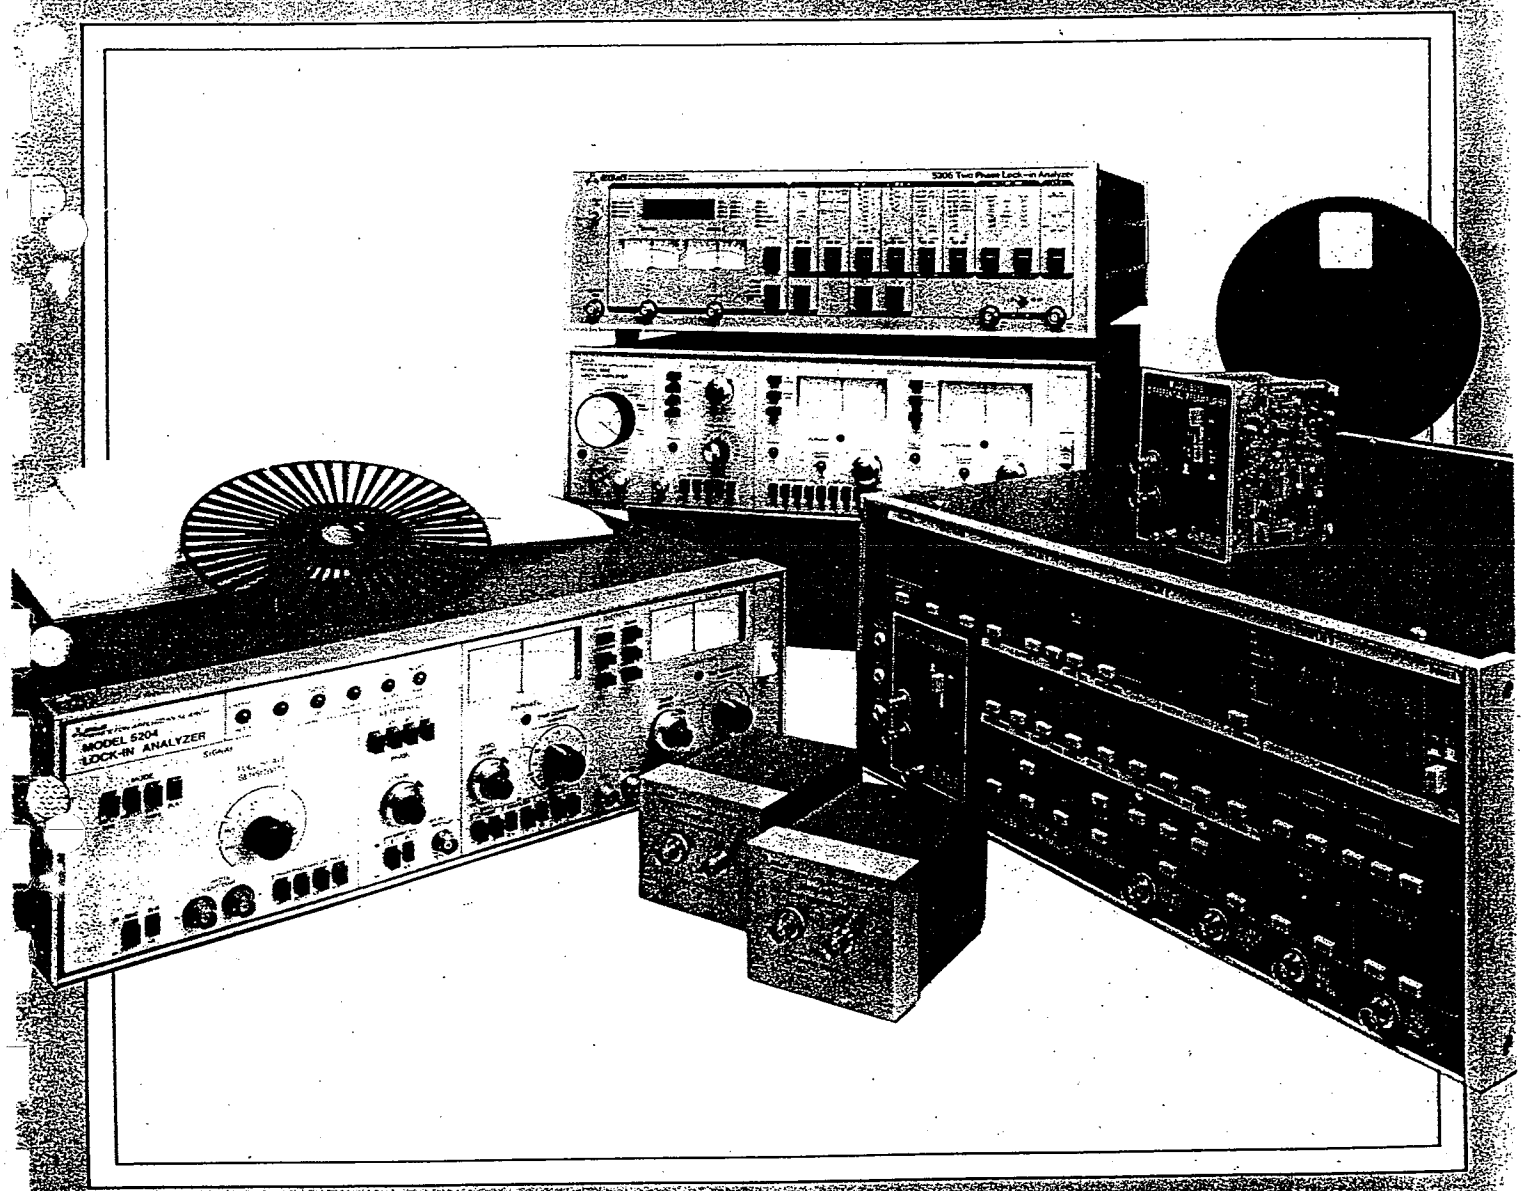

EG&G PRINCETON APPLIED RESEARCH

A Lock-In Primer

In recent years, the variety and complexity of lock-in amplifiers have increased significantly. Selecting the right lock-in for your application or interpreting tricky specifications requires a clear understanding of the different types of instruments available, how they work and their advantages and disadvantages. To better achieve such a clear understanding, we'll start by reviewing some of the basic "building-blocks" used in lock-in amplifiers.

MIXERS

A simple mixer circuit is shown in Figure 1. As shown, this type of mixer is simply a linear multiplier whose output e_3 is the product of its two inputs e_1 and e_2. With two sinusoidal inputs, there are two components in the mixer output, a sum frequency $(f_1 + f_2)$ term and a difference frequency $(f_1 - f_2)$ term. Notice that when the two input signals are *synchronous* (i.e., $f_1 = f_2$), then the difference frequency is *zero* and the difference component is therefore a phase-sensitive dc voltage.

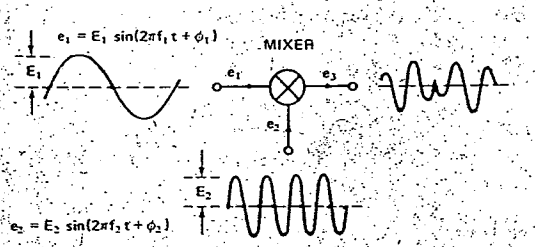

$$e_3 = e_1 \cdot e_2 = E_1 E_2 \sin(2\pi f_1 t + \phi_1) \sin(2\pi f_2 t + \phi_2)$$

$$e_3 = \underbrace{\frac{E_1 E_2}{2} \cos[2\pi(f_1 - f_2)t + (\phi_1 - \phi_2)]}_{\text{difference frequency } (f_1 - f_2) \text{ component}} - \underbrace{\frac{E_1 E_2}{2} \cos[2\pi(f_1 + f_2)t + (\phi_1 + \phi_2)]}_{\text{sum frequency } (f_1 + f_2) \text{ component}}$$

FIGURE 1. Simple mixer operation.

All lock-ins use a phase-sensitive detector (PSD) circuit and all PSD circuits consist of nothing more than a mixer followed by a low-pass filter (see sections 2 and 3). Heterodyning lock-in amplifiers use additional mixers to convert or translate the frequency of the input signal (f_1) to a different or intermediate frequency. Either the sum or the difference frequency of the mixer output may be used as the intermediate frequency (f_i). Such heterodyning mixers are said to "up-convert" when $f_i > f_1$ and to "down-convert" when $f_i < f_1$.

For use in lock-in amplifiers, mixer circuits must be capable of withstanding large amounts of noise (e.g., asynchronous signals, $f_1 \neq f_2$) without overloading. The term *dynamic reserve* is used to specify such noise overload performance. The dynamic reserve of a mixer, or any other circuit for that matter, is defined as the ratio of the overload level (peak value of an asynchronous signal that will just cause significant non-linearity) to the peak value of a full-scale synchronous signal. Dynamic reserve is often confused with input *dynamic range* which is the ratio of the overload level to the minimum detectable signal level.

Due to non-linearity and other problems, the linear multiplier type of mixer shown in Figure 1 cannot provide the dynamic reserve required in a lock-in amplifier. In commercial lock-ins, mixers are invariably of the switching type shown in Figure 2. The square-wave drive or switching signal used in switching mixers contains all odd harmonics of the fundamental frequency of the square wave. The sum and difference outputs of a switching mixer, therefore, are each composed of a large number of frequencies, e.g., $f_1 + f_2$, $f_1 + 3f_2$, $f_1 + 5f_2$, etc. are sum frequencies; $f_1 - f_2$, $f_1 - 3f_2$, $f_1 - 5f_2$, etc. are difference frequencies. Note that for a switching mixer, *synchronous* operation occurs whenever $f_1 = (2n+1)f_2$ where $(2n+1)$ is the harmonic number, compared with the more simple $f_1 = f_2$ condition in a linear multiplier type of mixer. Remember that with synchronous operation, one of the difference frequency components of the mixer output will be at zero frequency or dc. A switching mixer will therefore produce a phase-sensitive dc output whenever $f_1 = (2n+1)f_2$, and for $n > 0$, these outputs are known as the *harmonic responses* of the mixer.

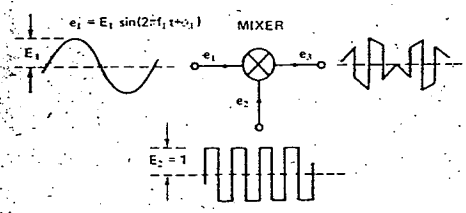

$$e_2 = \frac{4}{\pi} \sum_{n=0}^{n=\infty} \frac{1}{2n+1} \sin[(2n+1)(2\pi f_2 t + \phi_2)]$$

$$e_3 = e_1 \cdot e_2 = \sum_{n=0}^{n=\infty} \frac{4E_1}{(2n+1)\pi} \sin(2\pi f_1 t + \phi_1) \sin[(2n+1)(2\pi f_2 t + \phi_2)]$$

$$\text{therefore } e_3 = \underbrace{\sum_{n=0}^{n=\infty} \frac{2E_1}{(2n+1)\pi} \cos[2\pi(f_1 - (2n+1)f_2)t + (\phi_1 - (2n+1)\phi_2)]}_{\text{difference frequency components}}$$

$$- \underbrace{\sum_{n=0}^{n=\infty} \frac{2E_1}{(2n+1)\pi} \cos[2\pi(f_1 + (2n+1)f_2)t + (\phi_1 + (2n+1)\phi_2)]}_{\text{sum frequency components}}$$

FIGURE 2. Switching mixer operation.

In a switching mixer, the amplitude (E_2) of the square-wave drive (e_2) is relatively unimportant provided that it is sufficient to cause switching. As with the linear multiplier type of mixer shown previously, the output of a switching mixer is the product of its two inputs, i.e., $e_3 = e_1 \cdot e_2$. The effective amplitude of e_2 is now a constant ($E_2 = 1$) however, and only the phase and frequency(s) of this square-wave drive are of consequence. A simple switching mixer employing an inverting amplifier and two complementary semiconductor switches is shown in Figure 3. In the condition shown, the difference frequency is zero (i.e., synchronous condition of $f_1 = f_2$) and outputs are shown for four different phase relationships.

The upper transistor switch shown in Figure 3 turns on or conducts during positive half-cycles of the square-wave drive e_2; the lower switch conducts during negative half-cycles. When $\phi_1 = \pm\phi_2$, the mixer acts simply as a synchronous rectifier. Notice that the mixer dc output can be adjusted to any value from zero to $\pm(2/\pi)E_1$ by varying the phase difference ($\phi_1 - \phi_2$). In a perfect mixer, only synchronous inputs can cause a dc output. In practice, a mixer can produce a dc output with high level noise inputs or even with no (zero) input due to non-linearities associated with the square-wave drive input. Such spurious dc outputs are normally negligible in amplitude. However, at higher frequencies (above 10 kHz typically) the magnitude of such a dc offset and its associated drift may become significant.

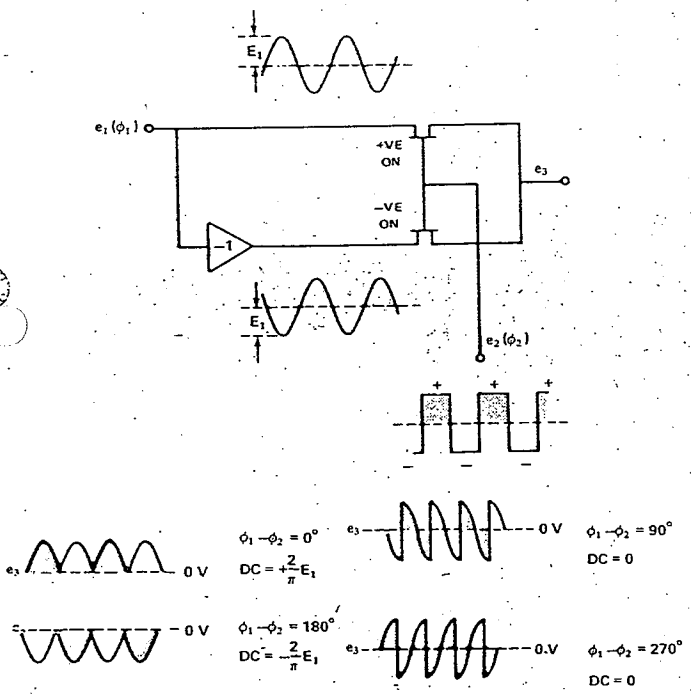

FIGURE 3. Switching mixer circuit (simplified).

LOW-PASS FILTERS

Low-pass filters, such as those shown in Figure 4, are used to pass dc and low-frequency ac signals while severely attenuating higher frequency signals. Commonly, such filters are characterized by a cut-off frequency or *signal bandwidth* (f_c) which is somewhat arbitrarily chosen to be the frequency at which the gain of the filter falls to 70.7% (or −3 dB) of its maximum value (0 dB). When used to reduce noise, it is more meaningful to talk in terms of their *equivalent noise bandwidth* (f_N).

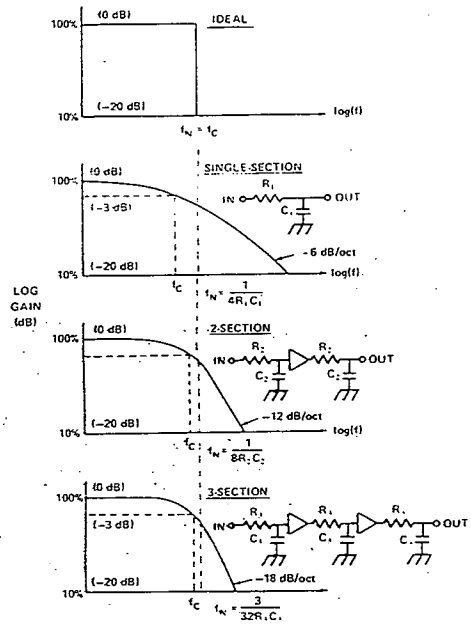

FIGURE 4. Low-pass filter characteristics.

Another characteristic used to compare filters is the rate at which the filter gain *rolls off* after the cut-off frequency (f_c). Note that an octave is a doubling of frequency and that −6 dB, −12 dB and −18 dB correspond respectively to 1/2, 1/4 and 1/8, so that a −6 dB/octave rolloff, for example, means that the filter gain is halved for each doubling of frequency. Simple single-section filters, consisting of one resistor and one capacitor, provide a −6 dB/octave rolloff and are used in inexpensive lock-ins. In more expensive instruments, −12 dB/octave or −18 dB/octave rolloff modes are normally provided, though a −6 dB/octave mode is often also provided to allow stable operation when the lock-in is used in a feedback loop. Each filter section is characterized by its time-constant (RC) value where R (ohms) × C (farads) = RC (seconds). Identical filter sections may be cascaded, using buffer amplifiers for isolation, to provide steeper rolloff rates. Note that, as more sections are added, the filter characteristic begins to approach that of the ideal filter.

The equivalent noise bandwidth of a filter is a theoretical, rectangular shaped frequency response with the same maximum gain as that of the actual frequency response and producing the same rms output noise. All three practical filters shown in Figure 4 have the same noise bandwidth as that shown for the "ideal" filter. Noise bandwidth is inversely related to time constant as shown in Figure 4. Increasing the time constant decreases the noise bandwidth (and hence the output noise) but does so at the expense of increasing the measurement time. Figure 5 shows the time response to a step input for 1, 2 and 3 section filters where the time constant for each filter is different and is chosen so that each filter provides the same noise bandwidth. Notice that the measurement time required by each filter, for its output to reach within 1% (for example) of the final value, decreases with the number of filter sections. By definition, the *ideal* low-pass filter is one offering the shortest measurement time for a given noise bandwidth. In practice, lock-ins offering 2 or 3 section filters come close to this ideal.

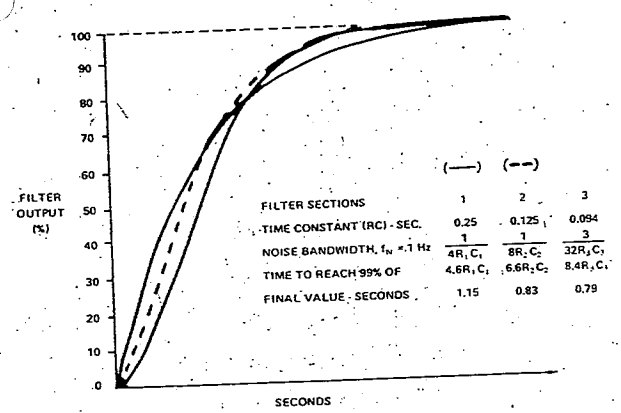

FIGURE 5.

THE PHASE SENSITIVE DETECTOR

The combination of a mixer circuit followed by a low-pass filter is known as a phase-sensitive detector (PSD). As we saw previously in Figure 2, the output of a switching mixer contains a large number of sinusoidal sum and difference frequency components (the number is large rather than infinite since the squareness of the e_2 does not contain all higher odd harmonics). The effect of the low-pass filter (see Figure 6) is to *average* all components of the mixer output which have frequencies beyond the filter cut-off. The average value of a sinusoidal component, i.e., a sine wave, is zero, and when the filter time constant is set correctly so that the filter cut-off frequency (f_c) is less than the fundamental frequency (f2) of the mixer square wave drive, the output of a PSD will contain only those *difference frequency* components having frequencies within the equivalent noise bandwidth of the filter.

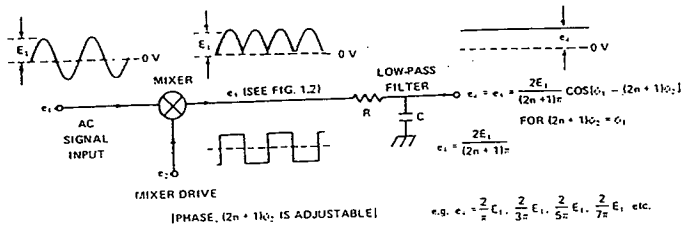

FIGURE 6. PSD operation with synchronous signal.

When the input signal (e1) to a PSD is synchronous, i.e., when $f_1 = (2n+1)f2$, the PSD output will contain a dc signal component as shown in Figure 6. Note that the amplitude of the harmonic response dc outputs are phase-sensitive and are inversely proportional to their harmonic number (n). In a lock-in amplifier, the phase (ϕ_2) of the PSD drive may be adjusted to maximize the dc output signal.

The PSD input signal (e_1) need not be sinusoidal. If e_1 were a synchronous square-wave signal such as that resulting from chopped light experiments, for example, then e_1 would contain a large number of synchronous components, each of which would give rise to an output dc signal from the PSD.

Suppose, as shown in Figure 7, that the PSD input (e1) is an asynchronous (noise) signal of frequency $f_1 = f_2 + \Delta f$. The resulting mixer sum and

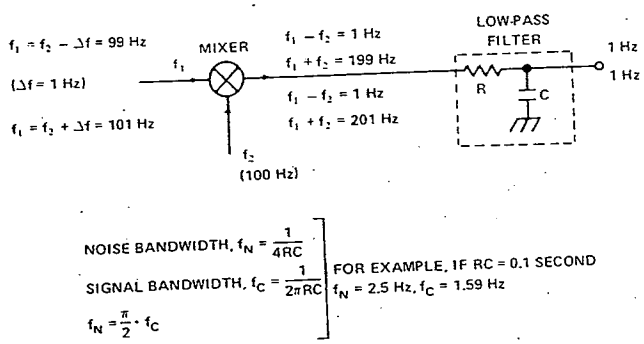

FIGURE 7. PSD operation with asynchronous (noise) signal.

difference frequencies (ignoring harmonics for simplicity) will therefore be $2f_2 + \Delta f$ and Δf respectively. Only the Δf component may be able to pass through the low-pass filter and appear as output noise. Suppose we change the frequency of this input noise to $f_1 = f_2 - \Delta f$. The resulting sum and difference frequencies will respectively be $2f_2 - \Delta f$

and $-\Delta f(=\Delta f)$. Again, only the Δf component can appear as output noise and the low-pass filter "cannot tell" whether its Δf input resulted from a $f_2+\Delta f$ input to the mixer or a $f_2-\Delta f$ input. In addition to its rectifying and phase-sensitive properties, the PSD therefore acts as a *comb-filter*, that is, it provides *bandpass* filter responses automatically centered on all odd harmonics of f_2 (see Figure 8). A PSD is sometimes referred to as a *synchronous filter*. Notice that the bandpass responses are *tracking*—their center-frequencies automatically track changes in the PSD drive frequency f_2.

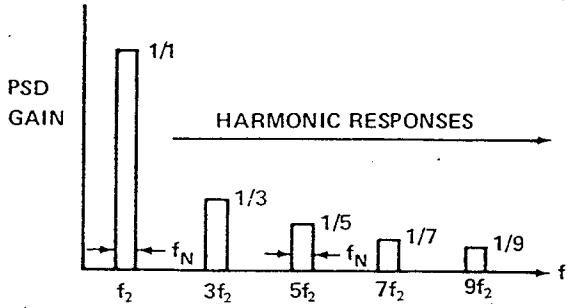

FIGURE 8. PSD frequency response.

Each bandpass response has an equivalent noise bandwidth determined by that of the low-pass filter. If the PSD input consists of white noise, i.e., constant amplitude noise at all frequencies, the effect of the harmonic response ($2n+1=3$, 5, 7, etc.) is to increase the PSD output noise by 11%. For square-wave signal inputs, the additional output noise (11%) caused by the harmonic responses is more than compensated by the increase in signal (23%). If the PSD is used to measure a sinusoidal signal accompanied by white noise, a separate bandpass filter, centered on f_2, may be used in front of the PSD to remove the harmonic responses and thus the additional 11% noise. The improvement in output *signal-to-noise ratio* effected by the use of such front-end filtering is normally insignificant. Front-end filters can be extremely helpful, however, in that by reducing the input noise before it reaches the PSD, the *dynamic reserve* of the lock-in may be improved significantly.

BROADBAND AND TUNED LOCK-INS

Both broadband (flat) and narrowband (tuned) lock-in amplifiers may be represented by the simplified block diagram shown in Figure 9. The reference input waveform to the lock-in may be of almost any waveshape and by definition is at the reference frequency and has zero phase. The output of the phase-locked loop (PLL) circuit is a precise square-wave, locked in phase to the reference input, and at a frequency f_2. Normally, $f_2=f_R$ the reference frequency; most lock-ins also provide a second harmonic mode where $f_2=2f_R$, and this mode is often used for derivative measurements. The phase-shifter circuit provides for precise adjustment of the phase ϕ_2, allowing the phase difference between the two PSD mixer inputs to be zeroed and thereby maximizing the output dc signal. All modern lock-ins have *tracking reference* circuits. That is, the reference circuits will follow or track any change in reference frequency, within the operating range of the instrument, and will maintain the selected value of phase-shift ϕ_2.

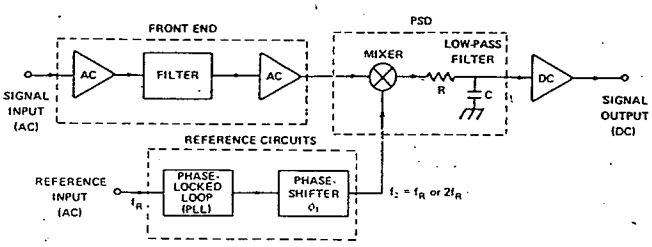

FIGURE 9. Basic lock-in amplifier (simplified).

For simplicity in Figure 9, the *front-end* ac amplifier and filter circuits are shown as separate functions. In practice, these front-end circuits are usually combined—as is normally the case also for the output low-pass filter and dc amplifier. Because of the dc drift of both the mixer and dc amplifier, the gain of the dc amplifier should be minimized to provide optimum *output stability* and ac gain used to provide the overall instrument gain required. Such a gain distribution is practicable and desirable for use with "clean" signals. With noisy signals, however, the ac gain must be reduced to provide increased dynamic reserve and the dc gain increased proportionately. Most high performance instruments provide the controls to allow such a tradeoff between dynamic reserve and output stability (see Figure 10).

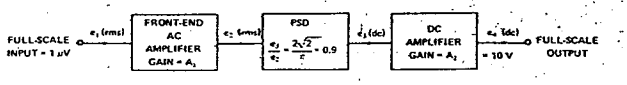

Let PSD overload level = 14.14 V (peak) = 10 V rms, and dc drift of (PSD + DC amplifier) = 10 μV/°C, then

AC GAIN, A_1	e_2 (rms)	e_3 (dc)	DC GAIN, A_2	DYNAMIC RESERVE ($=\frac{10\ V}{e_2}$)	OUTPUT DRIFT (% OF F.S.)
10^4	10 mV	9 mV	1.1×10^3	10^3 (60 dB)	0.11 %/°C
10^3	100 mV	90 mV	1.1×10^2	10^2 (40 dB)	0.011 %/°C
10^2	1 V	900 mV	1.1×10^1	10^1 (20 dB)	0.0011 %/°C

FIGURE 10. Dynamic reserve/output stability trade-off.

The basic difference between broadband and tuned lock-ins is in their front-end filter characteristics—see Figure 11. Broadband instruments are economical and have the advantage that, over a wide frequency range, a change in reference frequency (f_R) will not cause a change in front-end phase-shift. Tuned front-ends are more expensive to manufacture and must be returned if the reference frequency changes significantly. A tuned front-end instrument, however, offers the ultimate dynamic reserve performance. For noisy input signals, the narrowband frequency response effectively "protects" the PSD, reducing noise that would otherwise cause overloading. An approach that yields optimum lock-in performance is, of course, to provide both broadband and tuned modes of operation—the EG&G PARC Models 124A and 5301 provide no less than *five* signal channel modes together with many other unique features.

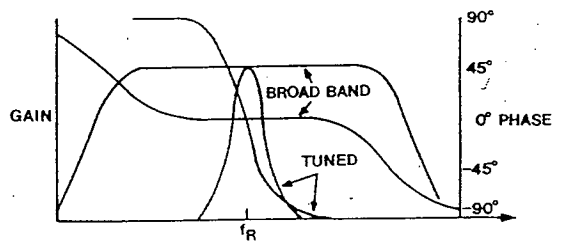

FIGURE 11. Broadband and tuned filter characteristics.

THE HETERODYNING/TRACKING FILTER LOCK-IN

Another method for obtaining a front-end filter uses the long established technique of heterodyning. The signal of interest is up-converted to a higher fixed frequency to provide a front-end bandpass filter (see Figure 12).

With this approach, a fixed frequency filter-amplifier is used to protect the PSD from noise and increase the dynamic reserve of the instrument. Early heterodyning systems divided the operating frequency range into a number of small frequency bands. Each band operated at its own intermediate frequency and required up to 5 circuit boards to be changed for each band. Recent designs eliminated the need to interchange cards since all cards could be installed simultaneously. the main advantage of the heterodyning system is that it approaches the tuned front-end system in dynamic reserve while allowing a signal channel tracking capability. It also removes harmonic response from the system. (See PHASE-SENSITIVE DETECTOR.)

Disadvantages of the heterodyning approach are the relatively poor front-end selectivity (tuned instruments can provide much narrower front-end filtering) and degraded phase stability and noise performance.

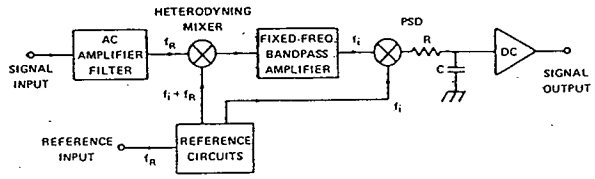

FIGURE 12. Heterodyning "tracking-filter" principle.

THE SYNCHRONOUS HETERODYNING (SYNCHRO-HET®) LOCK-IN

Two of the most important specifications of a lock-in are its dynamic reserve and its output stability. As mentioned previously, the amount of ac gain that can be used in a lock-in is limited by the noise accompanying the signal. The Synchro-Het® principle used in the EG&G PARC Model 186A is a patented technique which effectively allows ac gain to be increased without degradation of the instrument's dynamic reserve. The Model 186A is unique in its ability to *simultaneously* provide both outstanding dynamic reserve and excellent output drift performance.

In a lock-in amplifier, most of the noise accompanying an input signal is not removed until acted upon by the extremely narrowband filtering action of the PSD—the effect of a front-end filter is normally insignificant by comparison. Unfortunately, in conectional instruments, this cleaned-up PSD output is now a dc signal, the stability of which is degraded by any spurious PSD dc output and its associated drift and also by the dc drift of the following output amplifier. As we will see, the Synchro-Het® technique uses a combination of mixers and a unique rotating-capacitor filter to form what is effectively a PSD with *ac* output. This relatively noise-free PSD output may now be ac amplified before final rectification and smoothing takes place.

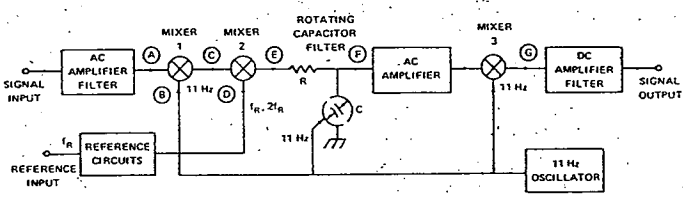

FIGURE 13. The Synchro-Het® principle (Model 186A).

Referring to both Figures 13 and 14, an input signal to the Model 186A is amplified by the front-end ac amplifier and then mixed with an 11 Hz square-wave in mixer 1. The effect of mixer 1 is to phase-chop the signal, i.e., the signal polarity is alternately reversed at an 11 Hz rate. The phase-chopped signal (waveform C) is then mixed with a reference frequency (or $2f_R$ in second harmonic mode) square-wave in mixer 2 to produce waveform E. The combination of mixer 2 and the rotating capacitor filter act to form a PSD with the square-wave ac output shown in waveform F. Note that any spurious dc outputs from mixer 1 or mixer 2 will not be amplified by the ac amplifier following the rotating-capacitor filter. After ac amplification, mixer 3 is used to synchronously rectify the signal (waveform G) and this dc signal is then amplified by the low gain output amplifier-filter. The overall low-pass filter response of the Model 186A rolls off at −18 dB/octave—approaching that of the "ideal" filter (see LOW-PASS FILTERS).

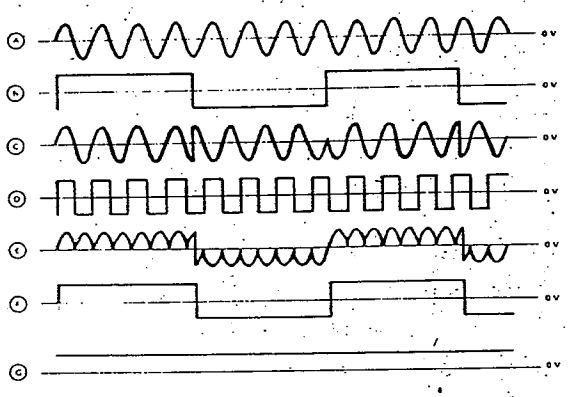

FIGURE 14. Synchro-Het® waveforms (see Figure 13).

TWO-PHASE/VECTOR LOCK-IN AMPLIFIER

This type of lock-in is shown in Figure 15. After front-end ac amplification and filtering, the amplified input signal (e1) is fed to two PSD circuits operating with square-wave drive waveforms which are in quadrature; i.e., 90° out of phase with one another. The low-pass filters shown following the two mixers are combined with their output dc amplifiers so that the time constant in each case is given by R_fC where R_f is the value of the feedback resistor. For two-phase operation, the feedback connection shown in blue from the quadrature output to the reference circuits is disconnected and the instrument simply provides two outputs—one proportional to $e_1\cos(\phi_1-\phi_2)$, the other proportional to $e_1\sin(\phi_1-\phi_2)$. Such an operating mode is particularly useful for two-phase applications such as ac bridge balancing.

In many such applications, it is necessary to measure a small in-phase signal in the presence of a large quadrature signal. For this reason, the

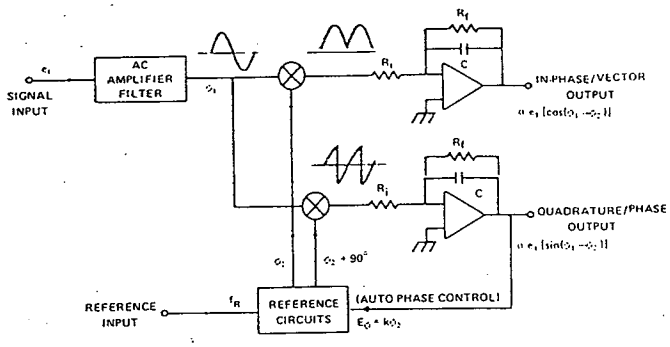

FIGURE 15. Two-phase/vector operation.

accuracy of the 90° phase difference between the two mixer drives, i.e., the *orthogonality*, is very important—as is the phase drift with time and temperature.

For vector operation, the feedback connection shown in blue is connected as shown, and the feedback resistor (R_f) associated with the quadrature output amplifier is disconnected. This output amplifier circuit now becomes an operational *integrator* with very high dc gain and a time constant of R_fC seconds. The integrator output (E_ϕ) is used to control the phase shift (ϕ_2) generated by the reference circuits. The effect of the feedback loop formed by the quadrature mixer, integrator and reference circuit phase-shifter is to automatically force the output of the quadrature mixer toward zero. This will reach zero when $\phi 1=\phi_2$, since $\mathrm{Sin}(\phi_1-\phi_2)=\mathrm{Sin}(0)=0$. If the phase-shift $\phi 2$ is linearly related to the control voltage E_ϕ, the integrator output will be proportional to $\phi 2$ and may be used as a PHASE output. When $(\phi_1-\phi_2)$ is automatically held to zero by the feedback loop, the in-phase channel output becomes proportional to $e_1\cos(\phi 1-\phi_2)=e_1$—i.e., a VECTOR MAGNITUDE output. Such a *phase-insensitive* mode of operation is particularly valuable in applications where the signal phase ($\phi 1$) is constantly changing.

THE MICROPROCESSOR-CONTROLLED LOCK-IN AMPLIFIER

The addition of a microprocessor to the lock-in amplifier creates new computer-controlled measurement possibilities, simplifies and automates manual measurements, and permits special features that increase measurement efficiency.

Full programmability of all instrument parameters via IEEE-488 or RS232 interfaces is a requirement in today's computer-controlled laboratory. Many complicated experiments require a computer to coordinate multiple signal sources, samples, and instruments. Under computer control, microprocessor-based, or *smart*, lock-in amplifiers are now being used for production testing of fiber optics, and for tedious measurements of low-level

light or electrochemical signals. They allow skilled workers to concentrate on the results rather than on data gathering. Built-in features like multiple curve buffers, binary or ASCII curve dump commands, and commands for all front panel functions minimize programming time.

The *smart* lock-in amplifier is valuable in manual use as well. The auto-range, auto-phase, auto-offset, auto-normalize and auto-tune functions simplify lock-in operation, reducing total measurement time. And the auto-measure function automatically finds the signal and adjusts the lock-in to make the complete measurement, even on moderately noisy signals. Moreover, when manually adjusting individual controls, a *smart* lock-in warns the user if an inconsistent setting was selected, thus improving measurement reliability. Additional convenience features include relative phase measurement, logarithm and ratio calculations, and a built-in, crystal-controlled frequency meter.

The marriage of analog and digital electronics brings several more benefits to the art of recovering signals from noise. In the *linear averaging* mode, the microprocessor zeroes in on the signal amplitude within one time constant, compared to the five time constants taken by the normal exponential mode. This is of great benefit in resolving peaks in spectral scans, at only a slight sacrifice in signal-to-noise improvement. The *sliding average* mode, on the other hand, allows measurement resolution to be improved by increasing the number of samples averaged together for each reading.

SIGNAL RECOVERY

One last subject: Just how deep can a signal be buried in noise and still be recovered by a lock-in? Figure 16 shows (in black) three types of noise that are frequently encountered in the output of an experiment. Discrete frequency interference such as line frequency pick-up (50/60 Hz or their harmonics) can usually be eliminated by simply choosing a suitable operating (reference) frequency. In most applications, random 1/f or flicker noise may also be greatly reduced by choosing a relatively high reference frequency. Random white noise with its constant noise power/unit bandwidth presents the most serious problem.

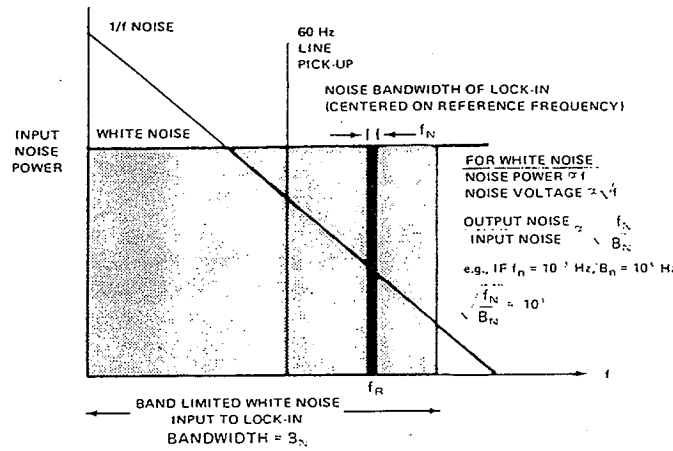

FIGURE 16. Signal/noise improvement.

When white noise is applied to the input of a lock-in, it becomes *band-limited* by the input frequency response of the instrument. Let's assume that this input noise bandwidth (B_N) containing the noise is (for example) 100 kHz or 10^5 Hz. The noise bandwidth (f_N) of the lock-in depends upon the PSD time constant selected. Let's suppose we can afford a time constant of about 100 seconds (i.e., a measurement time of about 500 seconds), so that $f_N = 0.001$ Hz or 10^{-3} Hz. Since the noise reduction or signal recovery effected by the lock-in is given by the *square root* of f_N/B_N, the output signal-to-noise ratio (SNR) will be 10^4 times better than the input SNR. Let's assume that, for use with a chart recorder, the output SNR must be no less than (for example) 10:1. Under these conditions then, the worst input SNR that can be used is simply $(10:1)/10^4 = 1:1000$. Assuming the lock-in has sufficient dynamic reserve so as not to overload on noise peaks (at least 5000 in this example), then a signal may be recovered from noise that is 1000 times larger.

EG&G PRINCETON APPLIED RESEARCH

Scientific Instruments Division

P. O. BOX 2565 • PRINCETON, NJ 08543, USA • 609/452-2111 • TELEX: 843409

T459; 3/86; 2M; B&R