PRINCETON APPLIED RESEARCH

A Lock-in Amplifier Primer

By Donald M. Munroe*

*Don Munroe is Division Manager, Scientific Instrument Division

I'm sure you've all heard of the astounding event that took place recently. Dr. D. P. Freeze, a well known experimental physicist, who had been lost and presumed dead during an Antarctic expedition some years ago, was discovered entombed in a huge block of ice.

To the amazement of his rescuers, Freeze was not dead and after thawing out was able to return to his post at the University.

A few days after being given his first assignment, Freeze visited his supervisor Dr. W. I. Thit, radiating gloom and despondency. "Bill," said Freeze, "you've given me an impossible task. This measurement you want me to make needs an incredibly sophisticated electronic system. Just look at the specifications we need."

"First of all, the system must amplify a 1 nV ac signal and turn it into a 10 volt dc signal. Bill, do you realize that's a gain of 10^{10} or 200 dB? Think of the shielding we'll need!"

"The huge gain wouldn't be so bad if we had a clean signal, but look at the input signal to noise ratio we can expect. I calculate that our 1 nV signal will be drowned by an interfering signal that's bigger by five orders of magnitude."

"For God's sake, Bill, don't you realize that means an input dynamic reserve of at least 10^5 if the system is not to overload. And look at the dynamic range that implies. We need to resolve our signal to 10 pV or one part in a hundred, and that means a minimum input dynamic range of 10^7 or 140 dB!"

By this time, Freeze was pacing backwards and forwards, his face pink with passion.

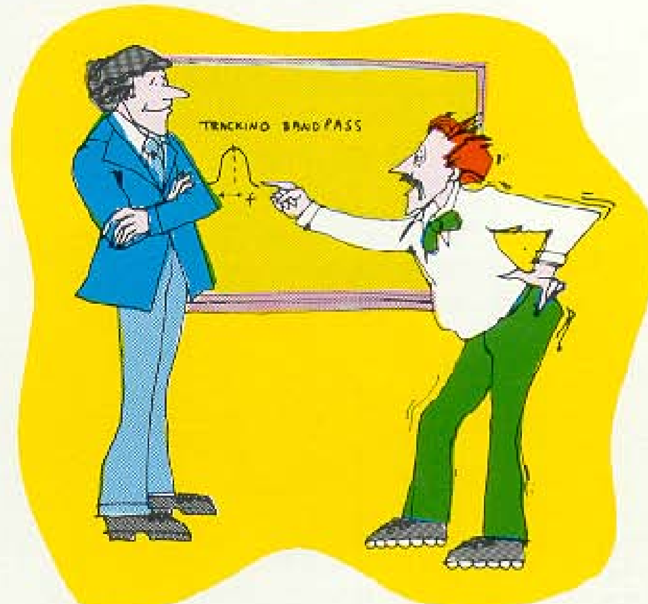

"And, Bill," he said in a choked voice, "the worst is yet to come - look at the filtering performance we need."

"Our system must lock-in, in both frequency and phase, to a reference signal and not only can the waveshape of our reference be sinusoidal, square, triangular, narrow pulses, or anything in between, but its frequency need not be constant - it can change continuously over a 10^5 range of frequency."

At this point, Freeze took a grip on himself and continued more calmly. "You see, Bill," he said, "I don't think you fully appreciate the problem."

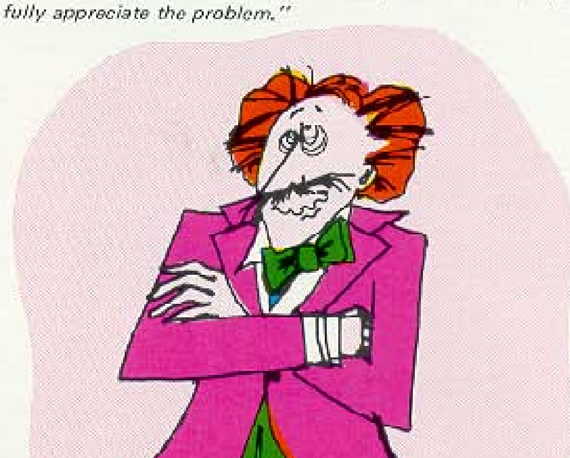

"Our system must act as a selective or tuned amplifier and amplify only a selectable narrow band of frequencies centered on the frequency of the reference signal. In other words, Bill, our system needs to be a frequency tracking bandpass filter with enormous gain and with selectable bandwidth or Q."

"If nothing else will convince you," said Freeze, "just look at the Q requirements. We both know what Q is - the filter center-frequency divided by the bandwidth, right? You remember when we were students learning circuit design - with either tuned L-C circuits or operational amplifier R-C circuits, a Q of 100 is about as high as you can go and still have acceptable frequency and amplitude stability, not to mention phase stability."

"All right," said Freeze, taking a deep breath, "this system you're asking for needs a bandwidth of 0.001 Hz at a center-frequency of 100 kHz. That implies a Q of 10^8 or 100 million - now do you believe it's impossible!" He slumped tiredly into a chair and waited for his supervisor to reply.

"It's understandable, of course," said W. I. Thit with a smile, "in view of your fifteen years of hibernation.

"What you've been describing is called a Lock-In Amplifier and there's a company called Princeton Applied Research who for the past 13 years has specialized in their design, manufacture and applications. Here's their latest catalog - take it with you and choose the model you need."

"Hello, P.A.R.C. applications engineer, I have this problem. I have your new lock-in amplifier catalog and..."

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.

1. Mixers

   A simple mixer circuit is shown in **Figure** 1.1. 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.

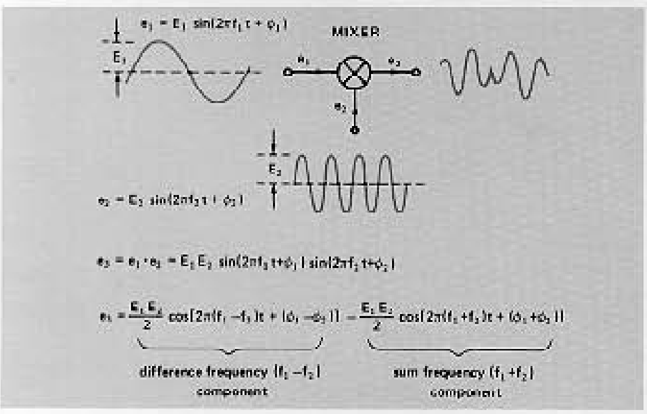

FIGURE 1.1 — SIMPLE MIXER OPERATION

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 (i.e. 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 1.1 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** 1.2. 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, i.e., $f_1 + f_2$, $f_1 + 3f_2$, $f_1 + 5f_2$, 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.

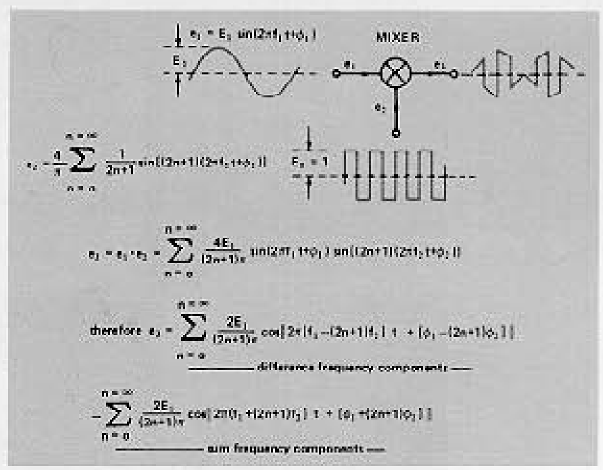

FIGURE 1.2 — SWITCHING MIXER OPERATION

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 1.3. 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 1.3 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.

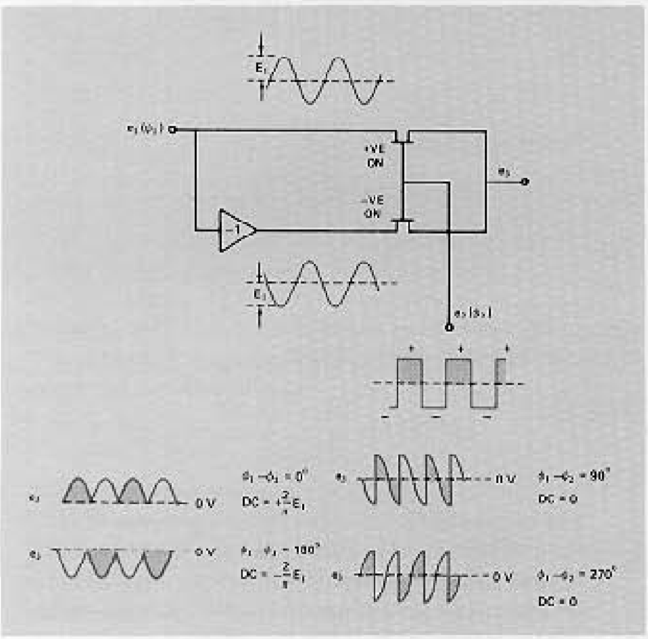

FIGURE 1.3 – SWITCHING MIXER CIRCUIT (SIMPLIFIED)

2. Low Pass Filters

Low-pass filters, such as those shown in Figure 2.1, 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).

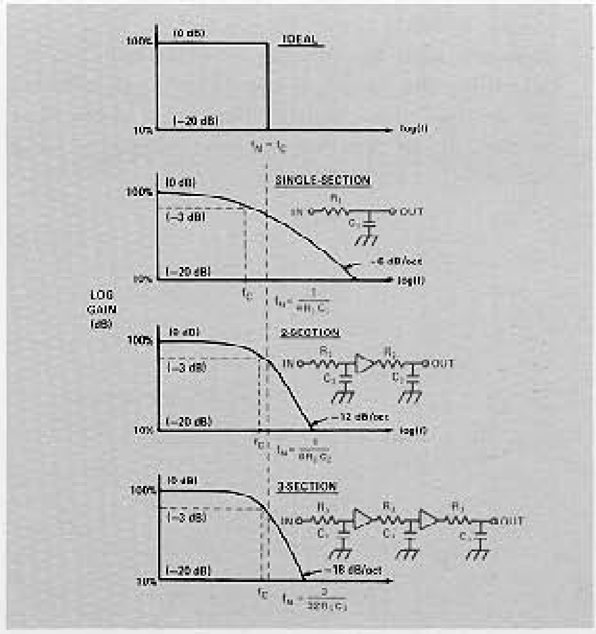

FIGURE 2.1 – LOW-PASS FILTER CHARACTERISTICS

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) x 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 2.1 have the same noise bandwidth as that shown for the "ideal" filter. Noise bandwidth is inversely related to time-constant as shown in Figure 2.1. 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 2.2 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% (say) 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.

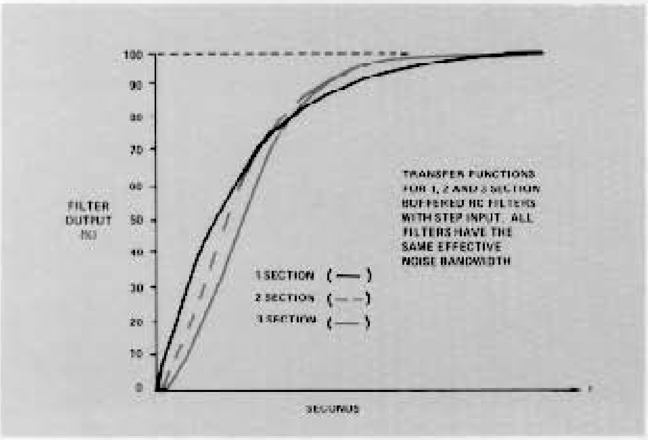

FIGURE 2.2

3. 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 1.2, 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 drive signal is not perfect and e_2 does not contain all higher odd harmonics). The effect of the low-pass filter (see Figure 3.1) 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 sinewave, 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 (f_2) 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.

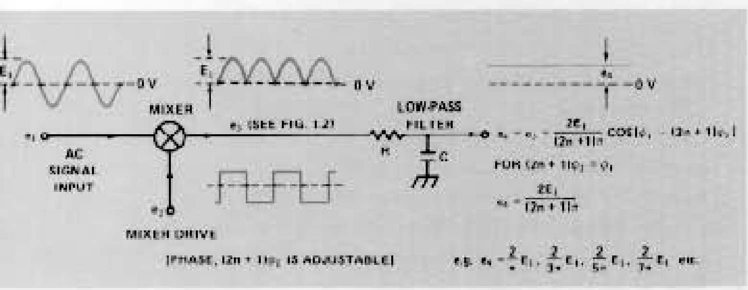

FIGURE 3.1 – PSD OPERATION WITH SYNCHRONOUS SIGNAL

When the input signal (e_1) to a PSD is synchronous, i.e., when $f_1 = (2n + 1)f_2$, the PSD output will contain a dc signal component as shown in Figure 3.1. 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 for example, such as that resulting from chopped light experiments, 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 3.2, that the PSD input (e_1) is an asynchronous (noise) signal of frequency $f_1 = f_2 + \Delta f$. The resulting mixer sum

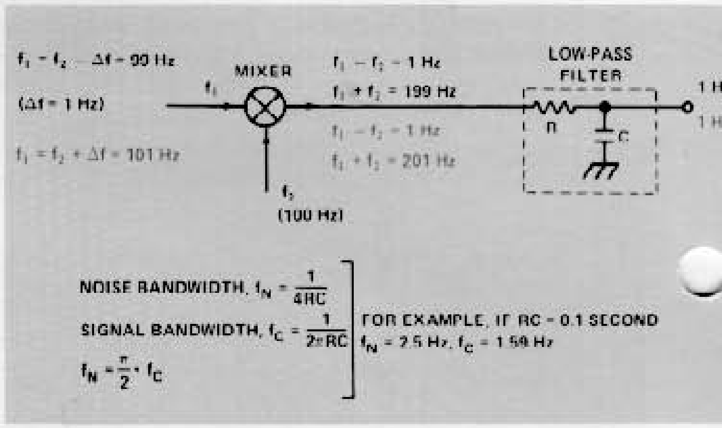

FIGURE 3.2 – PSD OPERATION WITH ASYNCHRONOUS (NOISE) SIGNAL

and 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 3.3. 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.

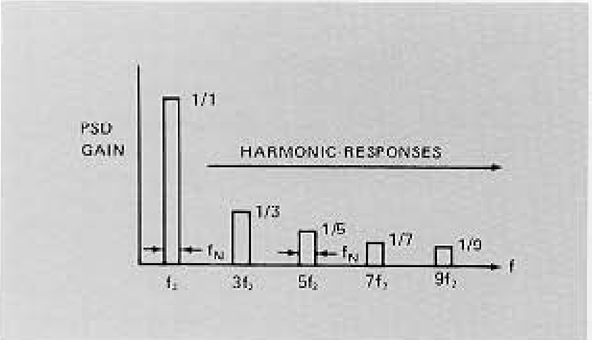

FIGURE 3.3 – PSD FREQUENCY RESPONSE

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 responses ($2n + 1 = 3, 5, 7$ 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.

4. Broad-Band and Tuned Lock-Ins

Both broad-band (flat) and narrow-band (tuned) lock-in amplifiers may be represented by the simplified block diagram shown in Figure 4.1. 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.

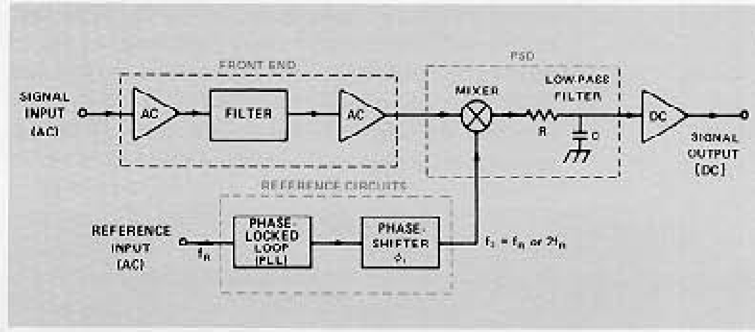

FIGURE 4.1 – BASIC LOCK-IN AMPLIFIER (SIMPLIFIED)

For simplicity in Figure 4.1, 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 trade-off between dynamic reserve and output stability - see Figure 4.2.

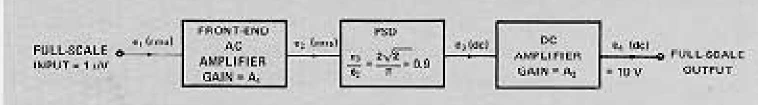

Let PSD overload level = 14.14 V (peak) = 10 V rms, and dc drift of (PSD + DC amplifier) = 10 µV/°C, then

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
10^5	100 mV	90 mV	1.1×10^2	10^2 (40 dB)	0.011%/°C
10^6	1 V	900 mV	1.1×10^1	10^1 (20 dB)	0.0011%/°C

FIGURE 4.2 – DYNAMIC RESERVE/OUTPUT STABILITY TRADE-OFF

The basic difference between broad-band and tuned lock-ins is in their front-end filter characteristics - see Figure 4.3. Broad-band 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 retuned if the reference frequency changes significantly. A tuned front-end instrument, however, offers the ultimate dynamic reserve performance. For noisy input signals, the narrow-band 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 broad-band and tuned modes of operation - the P.A.R.C. Model 124A provides no less than *five* signal channel modes together with many other unique features.

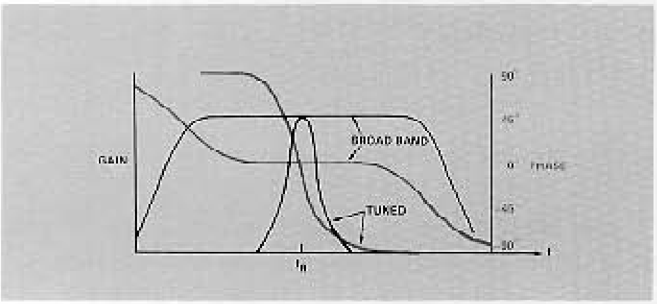

FIGURE 4.3 – BROAD-BAND AND TUNED FILTER CHARACTERISTICS

5. The Heterodyning/Tracking Filter Lock-In

As mentioned previously, the PSD circuit used in all lock-ins is a tracking filter. A competitive instrument, not made by P.A.R.C. but included here for completeness, uses an up-conversion heterodyning scheme to provide an additional front-end bandpass filter - see Figure 5.1.

With this approach, a fixed frequency filter-amplifier is used to protect the PSD from input noise and increase the dynamic reserve of the instrument. In order to use such a fixed-frequency filter, the operating frequency range of the instrument is divided into smaller frequency bands; for each band, the input signal frequency is heterodyned up to the center-frequency of the filter. The advantage of this approach is that, over a limited frequency range, the instrument offers dynamic reserve approaching that of a tuned front-end instrument, without requiring manual tuning. Another feature emphasized for this competitive instrument is that the front-end heterodyning removes harmonic responses (see section 3).

Disadvantages of this approach are the relatively poor front-end selectivity (tuned instruments can provide much narrower front-end filtering), degraded output stability at higher frequencies caused by PSD operation at a high intermediate frequency rather than at the reference frequency (see section 1) and the necessity of changing a large number of components in order to change frequency bands.

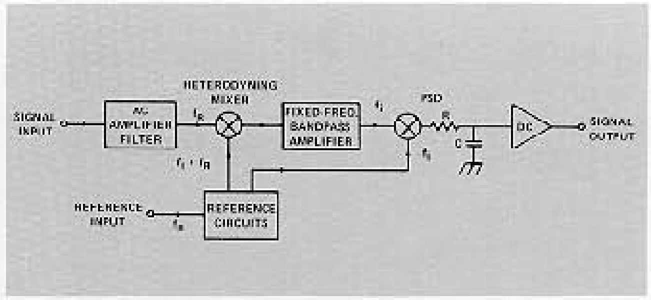

FIGURE 5.1 – HETERODYNING "TRACKING-FILTER" PRINCIPLE

6. 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 P.A.R.C. Model 186A is a recently patented technique which effectively allows ac gain to be increased without degradation of the instrument's dynamic reserve. The 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 narrow-band filtering action of the PSD - the effect of a front-end filter is normally insignificant by comparison. Unfortunately, in conventional 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.

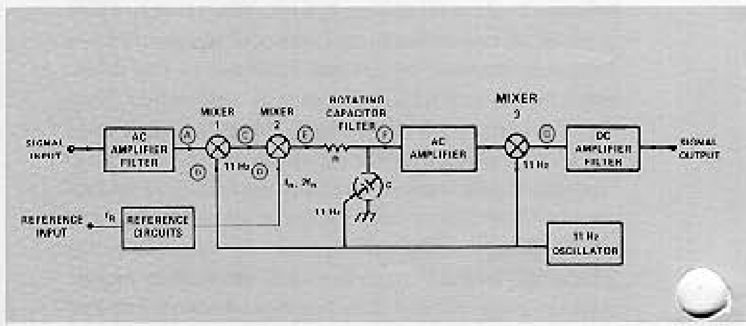

FIGURE 6.1 – THE SYNCHRO-HET® PRINCIPLE (MODEL 186A)

Referring to both Figures 6.1 and 6.2, 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 186A rolls off at −18 dB/octave - approaching that of the "ideal" filter (see section 2).

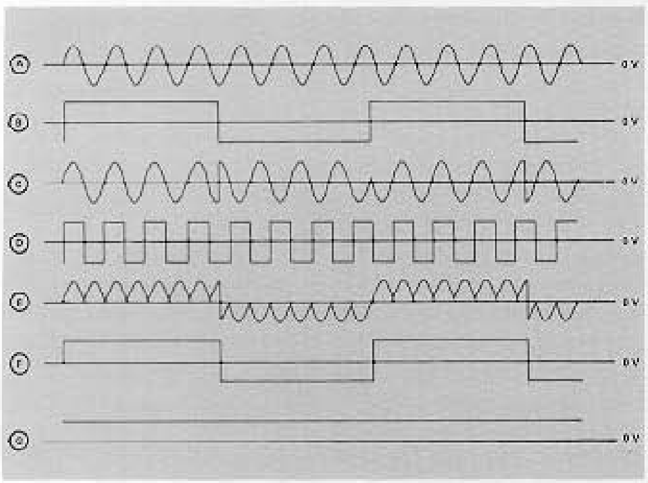

FIGURE 6.2 – SYNCHRO-HET® WAVEFORMS (SEE FIGURE 6.1)

7. Two-Phase/Vector Lock-In Amplifier

This type of lock-in is shown in Figure 7.1. After front-end ac amplification and filtering, the amplified input signal (e_1) 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.

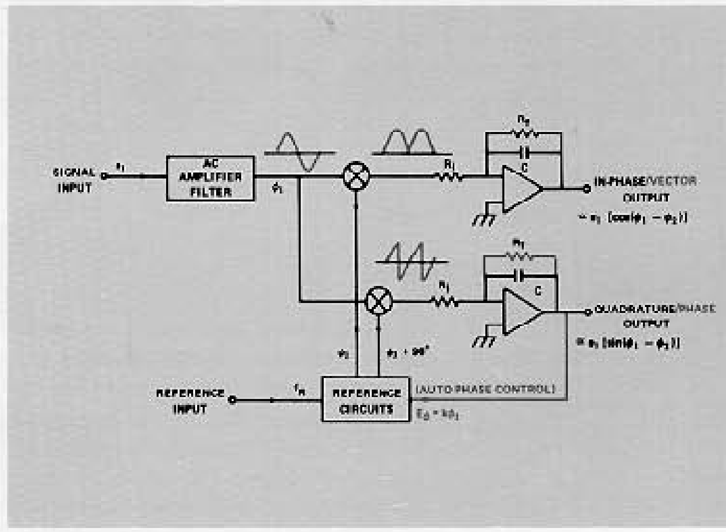

FIGURE 7.1 – TWO-PHASE/VECTOR OPERATION

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 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_iC 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 $\text{Sin}(\phi_1 - \phi_2) = \text{Sin}(0) = 0$. If the phase-shift ϕ_2 is linearly related to the control voltage E_ϕ, the integrator output will be proportional to ϕ_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 (ϕ_1) is constantly changing.

8. Signal Recovery

One last subject. Just how deep can a signal be buried in noise and still be recovered by a lock-in? Figure 8.1 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.

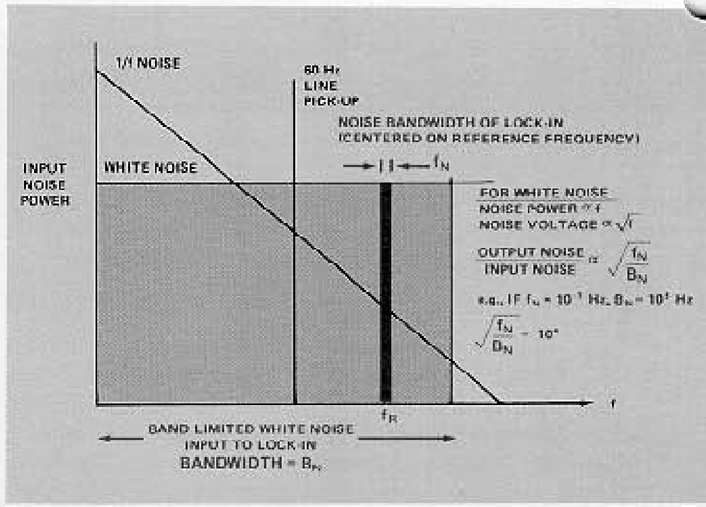

FIGURE 8.1 – SIGNAL/NOISE IMPROVEMENT

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 (say) 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 (say) 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.

First Class
Permit No. 285
Princeton, New Jersey

BUSINESS REPLY CARD
No postage stamp necessary if mailed in the United States

Postage will be paid by

PRINCETON APPLIED RESEARCH CORPORATION
P.O. Box 2565, Princeton, New Jersey 08540

PRINCETON APPLIED RESEARCH

CORPORATE HEADQUARTERS
Princeton Applied Research Corporation
Post Office Box 2565,
Princeton, New Jersey 08540
Telephone: 609/452-2111
Telex: 84-3409 Cable: PARCO

T351-15M-8/75-CP Printed in U.S.A.