

About Lock-In Amplifiers

Application Note #3

Lock-in amplifiers are used to detect and measure very small AC signals—all the way down to a few nanovolts. Accurate measurements may be made even when the small signal is obscured by noise sources many thousands of times larger. Lock-in amplifiers use a technique known as phase-sensitive detection to single out the component of the signal at a specific reference frequency and phase. Noise signals, at frequencies other than the reference frequency, are rejected and do not affect the measurement.

Why Use a Lock-In?

Let's consider an example. Suppose the signal is a 10 nV sine wave at 10 kHz. Clearly some amplification is required to bring the signal above the noise. A good low-noise amplifier will have about 5 nV/√Hz of input noise. If the amplifier bandwidth is 100 kHz and the gain is 1000, we can expect our output to be 10 μV of signal (10 nV × 1000) and 1.6 mV of broadband noise (5 nV/√Hz × √100 kHz × 1000). We won't have much luck measuring the output signal unless we single out the frequency of interest.


If we follow the amplifier with a band pass filter with a Q=100 (a VERY good filter) centered at 10 kHz, any signal in a 100 Hz bandwidth will be detected (10 kHz/Q). The noise in the filter pass band will be 50 μV (5 nV/√Hz × √100 Hz × 1000), and the signal will still be 10 μV. The output noise is much greater than the signal, and an accurate measurement can not be made. Further gain will not help the signal-to-noise problem.

Now try following the amplifier with a phase-sensitive detector (PSD). The PSD can detect the signal at 10 kHz with a bandwidth as narrow as 0.01 Hz! In this case, the noise in the detection bandwidth will be 0.5 μV (5 nV/√Hz × √0.01 Hz × 1000), while the signal is still 10 μV. The signal-to-noise ratio is now 20, and an accurate measurement of the signal is possible.

What is Phase-Sensitive Detection?

Lock-in measurements require a frequency reference. Typically, an experiment is excited at a fixed frequency (from an oscillator or function generator), and the lock-in detects the response from the experiment at the reference frequency. In the following diagram, the reference signal is a square wave at frequency ω_r . This might be the sync output from a function generator. If the sine output from the function generator is used to excite the experiment, the response might be the signal waveform shown below. The signal is $V_{sig}\sin(\omega_s t + \theta_{sig})$ where V_{sig} is the signal amplitude, ω_s is the signal frequency, and θ_{sig} is the signal's phase.

Lock-in amplifiers generate their own internal reference signal usually by a phase-locked-loop locked to the external reference. In the diagram, the external reference, the lock-in's reference, and the signal are all shown. The internal reference is $V_L\sin(\omega_L t + \theta_{ref})$.



The lock-in amplifies the signal and then multiplies it by the lock-in reference using a phase-sensitive detector or multiplier. The output of the PSD is simply the product of two sine waves.

$$V_{psd} = V_{sig} V_L \sin(\omega_s t + \theta_{sig}) \sin(\omega_L t + \theta_{ref})$$

$$= \frac{1}{2} V_{sig} V_L \cos([\omega_r - \omega_L]t + \theta_{sig} - \theta_{ref}) - \frac{1}{2} V_{sig} V_L \cos([\omega_r + \omega_L]t + \theta_{sig} + \theta_{ref})$$

The PSD output is two AC signals, one at the difference frequency ($\omega_r - \omega_L$) and the other at the sum frequency ($\omega_r + \omega_L$).

If the PSD output is passed through a low pass filter, the AC signals are removed. What will be left? In the general case, nothing. However, if ω_r equals ω_L , the difference frequency component will be a DC signal. In this case, the filtered PSD output will be:

$$V_{psd} = \frac{1}{2} V_{sig} V_L \cos(\theta_{sig} - \theta_{ref})$$

This is a very nice signal—it is a DC signal proportional to the signal amplitude.

It's important to consider the physical nature of this multiplication and filtering process in different types of lock-ins. In traditional analog lock-ins, the signal and reference are analog voltage signals. The signal and reference are multiplied in an analog multiplier, and the result is filtered with one or more stages of RC filters. In a digital lock-in, such as the SR830 or SR850, the signal and reference are represented by sequences of numbers. Multiplication and filtering are performed mathematically by a digital signal processing (DSP) chip. We'll discuss this in more detail later.

Narrow Band Detection

Let's return to our generic lock-in example. Suppose that instead of being a pure sine wave, the input is made up of signal plus noise. The PSD and low pass filter only detect

signals whose frequencies are very close to the lock-in reference frequency. Noise signals, at frequencies far from the reference, are attenuated at the PSD output by the low pass filter (neither $\omega_{\text{noise}} - \omega_{\text{ref}}$ nor $\omega_{\text{noise}} + \omega_{\text{ref}}$ are close to DC). Noise at frequencies very close to the reference frequency will result in very low frequency AC outputs from the PSD ($|\omega_{\text{noise}} - \omega_{\text{ref}}|$ is small). Their attenuation depends upon the low pass filter bandwidth and rolloff. A narrower bandwidth will remove noise sources very close to the reference frequency; a wider bandwidth allows these signals to pass. The low pass filter bandwidth determines the bandwidth of detection. Only the signal at the reference frequency will result in a true DC output and be unaffected by the low pass filter. This is the signal we want to measure.

Where Does the Lock-In Reference Come From?

We need to make the lock-in reference the same as the signal frequency, i.e. $\omega_r = \omega_L$. Not only do the frequencies have to be the same, the phase between the signals can not change with time. Otherwise, $\cos(\theta_{\text{sig}} - \theta_{\text{ref}})$ will change and V_{psd} will not be a DC signal. In other words, the lock-in reference needs to be phase-locked to the signal reference.

Lock-in amplifiers use a phase-locked loop (PLL) to generate the reference signal. An external reference signal (in this case, the reference square wave) is provided to the lock-in. The PLL in the lock-in amplifier locks the internal reference oscillator to this external reference, resulting in a reference sine wave at ω_r with a fixed phase shift of θ_{ref} . Since the PLL actively tracks the external reference, changes in the external reference frequency do not affect the measurement.

Internal Reference Sources

In the case just discussed, the reference is provided by the excitation source (the function generator). This is called an external reference source. In many situations the lock-in's internal oscillator may be used instead. The internal oscillator is just like a function generator (with variable sine output and a TTL sync) which is always phase-locked to the reference oscillator.

Magnitude and Phase

Remember that the PSD output is proportional to $V_{\text{sig}} \cos \theta$, where $\theta = (\theta_{\text{sig}} - \theta_{\text{ref}})$. θ is the phase difference between the signal and the lock-in reference oscillator. By adjusting θ_{ref} we can make θ equal to zero. In which case we can measure $V_{\text{sig}} (\cos \theta = 1)$. Conversely, if θ is 90° , there will be no output at all. A lock-in with a single PSD is called a single-phase lock-in and its output is $V_{\text{sig}} \cos \theta$.

This phase dependency can be eliminated by adding a second PSD. If the second PSD multiplies the signal with the reference oscillator shifted by 90° , i.e. $V_L \sin(\omega_L t + \theta_{\text{ref}} + 90^\circ)$, its low pass filtered output will be:

$$V_{\text{psd2}} = \frac{1}{2} V_{\text{sig}} V_L \sin(\theta_{\text{sig}} - \theta_{\text{ref}})$$

$$V_{\text{psd2}} \sim V_{\text{sig}} \sin \theta$$

Now we have two outputs: one proportional to $\cos \theta$ and the other proportional to $\sin \theta$. If we call the first output X and the second Y,

$$X = V_{\text{sig}} \cos \theta \quad Y = V_{\text{sig}} \sin \theta$$

these two quantities represent the signal as a vector relative to the lock-in reference oscillator. X is called the 'in-phase' component and Y the 'quadrature' component. This is because when $\theta = 0$, X measures the signal while Y is zero.

By computing the magnitude (R) of the signal vector, the phase dependency is removed.

$$R = (X^2 + Y^2)^{1/2} = V_{\text{sig}}$$

R measures the signal amplitude and does not depend upon the phase between the signal and lock-in reference.

A dual-phase lock-in has two PSDs with reference oscillators 90° apart, and can measure X, Y and R directly. In addition, the phase (θ) between the signal and lock-in is defined as:

$$\theta = \tan^{-1}(Y/X)$$

Digital PSD vs. Analog PSD

We mentioned earlier that the implementation of a PSD is different for analog and digital lock-ins. A digital lock-in, such as the SR830, multiplies the signal with the reference sine waves digitally. The amplified signal is converted to digital form using a 16-bit A/D converter sampling at 256 kHz. The A/D converter is preceded by a 102 kHz anti-aliasing filter to prevent higher frequency inputs from aliasing below 102 kHz.

This input data stream is multiplied, a point at a time, with the computed reference sine waves described previously. Every 4 μs the input signal is sampled, and the result is multiplied by both reference sine waves (90° apart).

The phase sensitive detectors (PSDs) in the digital lock-in act as linear multipliers; that is, they multiply the signal with a reference sine wave. Analog PSDs (both square wave and linear) have many problems associated with them. The main problems are harmonic rejection, output offsets, limited dynamic reserve, and gain error.

The digital PSD multiplies the digitized signal with a digitally computed reference sine wave. Because the reference sine waves are computed to 20 bits of accuracy, they have very low harmonic content. In fact, the harmonics are at the -120 dB level! This means that the signal is multiplied by a single reference sine wave (instead of a reference and its many harmonics), and only the signal at this single reference frequency is detected. The SR810, SR830 and SR850 digital lock-ins are completely insensitive to signals at harmonics of the reference. In contrast, a square wave multiplying lock-in will detect at all of the odd harmonics of the reference (a square wave contains many large odd harmonics).

Output offset is a problem because the signal of interest is a DC output from the PSD, and an output offset contributes to error and zero drift. The offset problems of analog PSDs are eliminated using the digital multiplier. There are no erroneous DC output offsets from the digital multiplication of the signal and reference. In fact, the actual multiplication is virtually error free.

The dynamic reserve of an analog PSD is limited to about 60 dB. When there is a large noise signal present, 1000 times (or 60 dB) greater than the full-scale signal, the analog PSD measures the signal with an error. The error is caused by non-linearity in the multiplication (the error at the output depends upon the amplitude of the input). This error can be quite large (10 % of full scale) and depends upon the noise amplitude, frequency and waveform. Since noise generally varies quite a bit in these parameters, the PSD error causes a lot of output uncertainty.

In the digital lock-in, dynamic reserve is limited by the quality of the A/D conversion. Once the input signal is digitized, no further errors are introduced. Certainly, the accuracy of the multiplication does not depend on the size of the numbers. The A/D converter used in the SR810, SR830 and SR850 is extremely linear, meaning that the presence of large noise signals does not impair its ability to correctly digitize a small signal. In fact, the dynamic reserve of these lock-ins can exceed 100 dB without any problems. We'll talk more about dynamic reserve a little later.

A linear, analog PSD multiplies the signal by an analog reference sine wave. Any amplitude variation in the reference amplitude shows up directly as a variation in the overall gain. Analog sine-wave generators are susceptible to amplitude drift: especially as a function of temperature. The digital reference sine wave has a precise amplitude and never changes. This avoids a major source of gain error common to analog lock-ins.

The overall performance of a lock-in amplifier is largely determined by the performance of its phase sensitive detectors. In virtually all respects, the digital PSD outperforms its analog counterparts.

What Does a Lock-In Measure?

So what exactly does the lock-in measure? Fourier's theorem basically states that any input signal can be represented as the sum of many sine waves of differing amplitudes, frequencies and phases. This is generally considered as representing the signal in the "frequency domain". Normal oscilloscopes display the signal in the "time domain". Except in the case of clean sine waves, the time domain representation does not convey very much information about the various frequencies which make up the signal.

A lock-in multiplies the signal by a pure sine wave at the reference frequency. All components of the input signal are multiplied by the reference simultaneously. Mathematically speaking, sine waves of differing frequencies are orthogonal, i.e. the average of the product of two sine waves is zero unless

the frequencies are EXACTLY the same. The product of this multiplication yields a DC output signal proportional to the component of the signal whose frequency is exactly locked to the reference frequency. The low pass filter (which follows the multiplier) provides the averaging which removes the products of the reference with components at all other frequencies.

A lock-in amplifier, because it multiplies the signal with a pure sine wave, measures the single Fourier (sine) component of the signal at the reference frequency. Let's take a look at an example. Suppose the input signal is a simple square wave at frequency f . The square wave is actually composed of many sine waves at multiples of f with carefully related amplitudes and phases. A 2 Vpp square wave can be expressed as:

$$S(t) = 1.273\sin(\omega t) + 0.4244\sin(3\omega t) + 0.2546\sin(5\omega t) + \dots$$

where $\omega = 2\pi f$. The lock-in, locked to f , will single out the first component. The measured signal will be $1.273\sin(\omega t)$, not the 2 Vpp that you'd measure on a scope.

In the general case, the input consists of signal plus noise. Noise is represented as varying signals at all frequencies. The ideal lock-in only responds to noise at the reference frequency. Noise at other frequencies is removed by the low pass filter following the multiplier. This "bandwidth narrowing" is the primary advantage that a lock-in amplifier provides. Only inputs with frequencies at the reference frequency result in an output.

RMS or Peak?

Lock-in amplifiers, as a general rule, display the input signal in volts rms. When a lock-in displays a magnitude of 1 V (rms), the component of the input signal (at the reference frequency) is a sine wave with an amplitude of 1 Vrms, or 2.8 Vpp.

Thus, in the previous example with a 2 Vpp square wave input, the lock-in would detect the first sine component, $1.273\sin(\omega t)$. The measured and displayed magnitude would be 0.90 Vrms (or $1.273/\sqrt{2}$).

Degrees or Radians?

In this discussion, frequencies have been referred to as f (Hz) and ω ($2\pi f$ radians/s). This is because people measure frequencies in cycles per second, and math works best in radians. For purposes of measurement, frequencies as measured in a lock-in amplifier are in Hz. The equations used to explain the actual calculations are sometimes written using ω to simplify the expressions.

Phase is always reported in degrees. Once again, this is more by custom than by choice. Equations written as $\sin(\omega t + \theta)$ are written as if θ is in radians, mostly for simplicity. Lock-in amplifiers always manipulate and measure phase in degrees.

Dynamic Reserve

The term "dynamic reserve" comes up frequently in discussions about lock-in amplifiers. It's time to discuss this

term in a little more detail. Assume the lock-in input consists of a full-scale signal at f_{ref} plus noise at some other frequency. The traditional definition of dynamic reserve is the ratio of the largest tolerable noise signal to the full-scale signal, expressed in dB. For example, if full scale is 1 μ V, then a dynamic reserve of 60 dB means noise as large as 1 mV (60 dB greater than full scale) can be tolerated at the input without overload.


The problem with this definition is the word "tolerable". Clearly, the noise at the dynamic reserve limit should not cause an overload anywhere in the instrument—not in the input signal amplifier, PSD, low pass filter or DC amplifier. This is accomplished by adjusting the distribution of the gain. To achieve high reserve, the input signal gain is set very low so the noise is not likely to overload. This means that the signal at the PSD is also very small. The low pass filter removes the large noise components from the PSD output which allows the remaining DC component to be amplified (a lot) to reach 10 V full scale. There is no problem running the input amplifier at low gain. However, as we have discussed previously, analog lock-ins have a problem with high reserve because of the linearity of the PSD and the DC offsets of the PSD and DC amplifier. In an analog lock-in, large noise signals almost always disturb the measurement in some way.

The most common problem is a DC output error caused by the noise signal. This can appear as an offset or as a gain error. Since both effects are dependent upon the noise amplitude and frequency, they can not be offset to zero in all cases and will limit the measurement accuracy. Because the errors are DC in nature, increasing the time constant does not help. Most lock-ins define tolerable noise as levels which do not affect the output more than a few percent of full scale. This is more severe than simply not overloading.

Another effect of high dynamic reserve is to generate noise and drift at the output. This comes about because the DC output amplifier is running at very high gain, and low-frequency noise and offset drift at the PSD output or the DC amplifier input will be amplified and appear large at the output. The noise is more tolerable than the DC drift errors since increasing the time constant will attenuate the noise. The DC drift in an analog lock-in is usually on the order of 1000 ppm/ $^{\circ}$ C when using 60 dB of dynamic reserve. This means that the zero point moves 1 % of full scale over 10 $^{\circ}$ C temperature change. This is generally considered the limit of tolerable.

Lastly, dynamic reserve depends on the noise frequency. Clearly noise at the reference frequency will make its way to the output without attenuation. So the dynamic reserve at f_{ref} is 0 dB. As the noise frequency moves away from the reference frequency, the dynamic reserve increases. Why? Because the low pass filter after the PSD attenuates the noise components. Remember, the PSD outputs are at a frequency of $|f_{noise} - f_{ref}|$. The rate at which the reserve increases depends upon the low pass filter time constant and rolloff. The reserve increases at the rate at which the filter rolls off. This is why 24 dB/oct filters are better than 6 or 12 dB/oct filters. When the noise frequency is far away, the reserve is limited by the

gain distribution and overload level of each gain element. This reserve level is the dynamic reserve referred to in the specifications.



The above graph shows the actual reserve vs. the frequency of the noise. In some instruments, the signal input attenuates frequencies far outside the lock-in's operating range ($f_{noise} \gg 100$ kHz). In these cases, the reserve can be higher at these frequencies than within the operating range. While this creates a nice specification, removing noise at frequencies very far from the reference does not require a lock-in amplifier. Lock-ins are used when there is noise at frequencies near the signal. Thus, the dynamic reserve for noise within the operating range is more important.

Dynamic Reserve in Digital Lock-Ins

The SR810, SR830 and SR850, with their digital phase sensitive detectors, do not suffer from DC output errors caused by large noise signals. The dynamic reserve can be increased to above 100 dB without measurement error. Large noise signals do not cause output errors from the PSD. The large DC gain does not result in increased output drift.

In fact, the only drawback to using ultra-high dynamic reserves (>60 dB) is the increased output noise due to the noise of the A/D converter. This increase in output noise is only present when the dynamic reserve is increased above 60 dB and above the minimum reserve. (If the minimum reserve is 80 dB, then increasing to 90 dB may increase the noise. As we'll discuss next, the minimum reserve does not have increased output noise: no matter how large it is.)

To set a scale, the digital lock-in's output noise at 100 dB dynamic reserve is only measurable when the signal input is grounded. Let's do a simple experiment. If the lock-in reference is at 1 kHz, and a large signal is applied at 9.5 kHz, what will the lock-in output be? If the signal is increased to the dynamic reserve limit (100 dB greater than full scale), the output will reflect the noise of the signal at 1 kHz. The spectrum of any pure sine generator always has a noise floor, i.e. there is some noise at all frequencies. So even though the

applied signal is at 9.5 kHz, there will be noise at all other frequencies, including the 1 kHz lock-in reference. This noise will be detected by the lock-in and appear as noise at the output. This output noise will typically be greater than the lock-in's own output noise. In fact, virtually all signal sources will have a noise floor which will dominate the lock-in output noise. Of course, noise signals are generally much noisier than pure sine generators and will have much higher broadband noise floors.

If the noise does not reach the reserve limit, the digital lock-in's own output noise may become detectable at ultra-high reserves. In this case, simply lower the dynamic reserve and the DC gain will decrease, and the output noise will decrease also. In general, do not run with more reserve than necessary. Certainly don't use ultra-high reserve when there is virtually no noise at all.

The frequency dependence of dynamic reserve is inherent in the lock-in detection technique. The SR810, SR830 and SR850, by providing more low-pass filter stages, can increase the dynamic reserve close to the reference frequency. The specified reserve applies to noise signals within the operating range of the lock-in, i.e. frequencies below 100 kHz. The reserve at higher frequencies is actually greater but is generally not that useful.

Minimum Dynamic Reserve

The SR810, SR830 and SR850 always have a minimum amount of dynamic reserve. This minimum reserve changes with the sensitivity (gain) of the instrument. At high gains (full-scale sensitivity of 50 μV and below), the minimum dynamic reserve increases from 37 dB at the same rate as the sensitivity increases. For example, the minimum reserve at 5 μV sensitivity is 57 dB. In many analog lock-ins, the reserve can be lower. Why can't the digital lock-ins run with lower reserve at this sensitivity?

The answer to this question is: "Why would you want lower reserve?" In an analog lock-in, lower reserve means less output error and drift. In the SR800 series lock-ins, more reserve does not increase the output error or drift. But, more reserve can increase the output noise. However, if the analog signal gain before the A/D converter is high enough, the 5 $\text{nV}/\sqrt{\text{Hz}}$ noise of the signal input will be amplified to a level greater than the input noise of the A/D converter. At this point, the detected noise will reflect the actual noise at the signal input and not the A/D converter's noise. Increasing the analog gain (decreasing the reserve) will not decrease the output noise. Thus, there is no reason to decrease the reserve. At a sensitivity of 5 μV , the analog gain is sufficiently high so that A/D converter noise is not a problem. Sensitivities below 5 μV do not require any more gain since the signal-to-noise ratio will not be improved (the front-end noise dominates). The SR800 series lock-ins do not increase their gain below the 5 μV sensitivity. Instead, the minimum reserve increases. Of course, the input gain can be decreased and the reserve increased; in which case, the A/D converter noise might be detected in the absence of any signal input.

Dynamic Reserve in Analog Lock-Ins

Because of the limitations of their PSDs, analog lock-in amplifiers must use different techniques to improve their dynamic reserve. The most common of these is the use of analog prefilters. The SR510 and SR530 have tunable, band-pass filters at their inputs. The filters are designed to automatically track the reference frequency. If an interfering signal is attenuated by a filter before it reaches the lock-in input, the dynamic reserve of the lock-in will be increased by that amount. For the SR510 and SR530, a dynamic reserve increase of up to 20 dB can be realized using the input band pass filter. Of course, such filters add their own noise and contribute to phase error: so they should only be used if necessary.

A lock-in can measure signals as small as a few nanovolts. A low-noise signal amplifier is required to boost the signal to a level where the A/D converter can digitize the signal without degrading the signal-to-noise. The analog gain in the SR850 ranges from roughly 7 to 1000. As discussed previously, higher gains do not improve signal-to-noise and are not necessary.

The overall gain (AC and DC) is determined by the sensitivity. The distribution of the gain (AC versus DC) is set by the dynamic reserve.

Input Noise

The input noise of the SR810, SR830 or SR850 signal amplifier is about 5 $\text{nVrms}/\sqrt{\text{Hz}}$. The SR530 and SR510 lock-ins have 7 $\text{nVrms}/\sqrt{\text{Hz}}$ of input noise. What does this noise figure mean? Let's set up an experiment. If an amplifier has 5 $\text{nVrms}/\sqrt{\text{Hz}}$ of input noise and a gain of 1000, then the output will have 5 $\mu\text{Vrms}/\sqrt{\text{Hz}}$ of noise. Suppose the amplifier output is low-pass filtered with a single RC filter (6 dB/oct rolloff) with a time constant of 100 ms. What will be the noise at the filter output?

Amplifier input noise and Johnson noise of resistors are Gaussian in nature. That is, the amount of noise is proportional to the square root of the bandwidth in which the noise is measured. A single stage RC filter has an equivalent noise bandwidth (ENBW) of $1/4T$, where T is the time constant ($R \times C$). This means that Gaussian noise at the filter input is filtered with an effective bandwidth equal to the ENBW. In this example, the filter sees 5 $\mu\text{Vrms}/\sqrt{\text{Hz}}$ of noise at its input. It has an ENBW of $1/(4 \times 100 \text{ ms})$ or 2.5 Hz. The voltage noise at the filter output will be 5 $\mu\text{Vrms}/\sqrt{\text{Hz}} \times \sqrt{2.5 \text{ Hz}}$, or 7.9 μVrms . For Gaussian noise, the peak-to-peak noise is about 5 times the rms noise. Thus, the output will have about 40 μVpp of noise.

Input noise for a lock-in works the same way. For sensitivities below about 5 μV full scale, the input noise will determine the output noise (at minimum reserve). The amount of noise at the output is determined by the ENBW of the low pass filter. The ENBW depends upon the time constant and filter rolloff. For example, suppose the lock-in is set to 5 μV full scale, with a 100 ms time constant, and 6 dB/oct of filter rolloff. The lock-in

will measure the input noise with an ENBW of 2.5 Hz. This translates to 7.9 nVrms at the input. At the output, this represents about 0.16 % of full scale (7.9 nV/5 μ V). The peak-to-peak noise will be about 0.8 % of full scale.

All of this assumes that the signal input is being driven from a low impedance source. Remember resistors have Johnson noise equal to $0.13 \times \sqrt{R}$ nVrms/ $\sqrt{\text{Hz}}$. Even a 50 Ω resistor has almost 1 nVrms/ $\sqrt{\text{Hz}}$ of noise! A signal source impedance of 2 k Ω will have a Johnson noise greater than the lock-in's input noise. To determine the overall noise of multiple noise sources, take the square root of the sum of the squares of the individual noise figures. For example, if a 2 k Ω source impedance is used, the Johnson noise will be 5.8 nVrms/ $\sqrt{\text{Hz}}$. The overall noise at the lock-in's input will be $[5^2 + 5.8^2]^{1/2}$, or 7.7 nVrms/ $\sqrt{\text{Hz}}$.

Noise Sources

What is the origin of the noise we've been discussing? There are two types of noise we have to worry about in laboratory situations: intrinsic noise and external noise. Intrinsic noise sources, like Johnson noise and shot noise, are inherent to all physical processes. Though we cannot get rid of intrinsic noise sources, by being aware of their nature we can minimize their effects. External noise sources are found in the environment—such as power line noise and broadcast stations. The effect of these noise sources can be minimized by careful attention to grounding, shielding, and other aspects of experimental design. We will first discuss some of the sources of intrinsic noise.

Johnson Noise

Every resistor generates a noise voltage across its terminals due to thermal fluctuations in the electron density within the resistor itself. These fluctuations give rise to an open-circuit noise voltage:

$$V_{\text{noise}}(\text{rms}) = (4kTR\Delta f)^{1/2}$$

where k =Boltzmann's constant (1.38×10^{-23} J/ $^{\circ}$ K), T is the temperature in Kelvin (typically 300 $^{\circ}$ K), R is the resistance in ohms, and Δf is the bandwidth of the measurement in Hz.

Since the input signal amplifier in a lock-in typically has a bandwidth of approximately 300 kHz, the effective noise at the amplifier input is $V_{\text{noise}} = 70\sqrt{R}$ nVrms, or $350\sqrt{R}$ nVpp. This noise is broadband. So if the source impedance is large, it can determine the amount of dynamic reserve required.

The amount of noise measured by the lock-in is determined by the measurement bandwidth. Remember, the lock-in does not narrow its detection bandwidth until after the phase sensitive detectors. In a lock-in, the equivalent noise bandwidth (ENBW) of the low pass filter (time constant) sets the detection bandwidth. In this case, the measured noise of a resistor at the lock-in input, typically the source impedance of the signal, is simply:

$$V_{\text{noise}}(\text{rms}) = 0.13\sqrt{R}\sqrt{\text{ENBW}} \text{ nV}$$

Shot Noise

Electric current has noise due to the finite nature of the charge carriers. There is always some non-uniformity in the electron flow which generates noise in the current. This noise is called "shot noise". This can appear as voltage noise when current is passed through a resistor, or as noise in a current measurement. The shot noise, or current noise, is given by:

$$I_{\text{noise}}(\text{rms}) = (2qI\Delta f)^{1/2}$$

where q is the electron charge (1.6×10^{-19} Coulomb), I is the rms AC current or DC current depending upon the circuit, and Δf is the bandwidth.

When the current input of a lock-in is used to measure an AC signal current, the bandwidth is typically so small that shot noise is not important.

1/f Noise

Every 10 Ω resistor, no matter what it is made of, has the same Johnson noise. However, there is excess noise in addition to Johnson noise which arises from fluctuations in resistance due to the current flowing through the resistor. For carbon composition resistors, this is typically 0.1 μ V to 3 μ V of rms noise per volt applied across the resistor. Metal film and wire-wound resistors have about 10 times less noise. This noise has a 1/f spectrum and makes measurements at low frequencies more difficult. Other sources of 1/f noise include noise found in vacuum tubes and semiconductors.

Total Noise

All of these noise sources are incoherent. The total random noise is the square root of the sum of the squares of all the incoherent noise sources.

External Noise Sources


In addition to the intrinsic noise sources discussed previously, there are a variety of external noise sources within the laboratory. Most of these noise sources are asynchronous, i.e. they are not related to the reference, and do not occur at the reference frequency or its harmonics. Examples include lighting fixtures, motors, cooling units, radios, computer screens, etc. These noise sources affect the measurement by increasing the required dynamic reserve or lengthening the time constant.

Some noise sources, however, are related to the reference, and if picked up in the signal, will add or subtract from the actual signal and cause errors in the measurement. Typical sources of synchronous noise are ground loops between the experiment, detector and lock-in; and electronic pick up from the reference oscillator or experimental apparatus.

Many of these noise sources can be minimized with good laboratory practice and experiment design. There are several ways in which noise sources are coupled into the signal path.

Capacitive Coupling

An AC voltage from a nearby piece of apparatus can couple to a detector via stray capacitance. Although C_{stray} may be very small, the coupled noise may still be larger than a weak experimental signal. This is especially damaging if the coupled noise is synchronous (at the reference frequency).



We can estimate the noise current caused by a stray capacitance by:

$$i = C_{\text{stray}} \frac{dV}{dt} = \omega C_{\text{stray}} V_{\text{noise}}$$

where ω is 2π times the noise frequency, V_{noise} is the noise amplitude, and C_{stray} is the stray capacitance.

For example, if the noise source is a power circuit, then $f = 60$ Hz and $V_{\text{noise}} = 120$ V. C_{stray} can be estimated using a parallel plate equivalent capacitor. If the capacitance is roughly an area of 1 cm^2 separated by 10 cm , then C_{stray} is 0.009 pF . The resulting noise current will be 400 pA (at 60 Hz). This small noise current can be thousands of times larger than the signal current. If the noise source is at a higher frequency, the coupled noise will be even greater.

If the noise source is at the reference frequency, the problem is much worse. The lock-in rejects noise at other frequencies, but pick-up at the reference frequency appears as signal!


Cures for capacitive noise coupling include:

- 1) Removing or turning off the noise source.
- 2) Keeping the noise source far from the experiment (reducing C_{stray}). Do not bring the signal cables close to the noise source.
- 3) Designing the experiment to measure voltages with low impedance (noise current generates very little voltage).
- 4) Installing capacitive shielding by placing both the experiment and detector in a metal box.

Inductive Coupling

An AC current in a nearby piece of apparatus can couple to the experiment via a magnetic field. A changing current in a nearby circuit gives rise to a changing magnetic field which

induces an emf ($d\Phi_B/dt$) in the loop connecting the detector to the experiment. This is like a transformer with the experiment-detector loop as the secondary winding.




Cures for inductively coupled noise include:

- 1) Removing or turning off the interfering noise source.
- 2) Reduce the area of the pick-up loop by using twisted pairs or coaxial cables, or even twisting the two coaxial cables used in differential connections.
- 3) Using magnetic shielding to prevent the magnetic field from crossing the area of the experiment.
- 4) Measuring currents (not voltages) from high-impedance detectors.

Resistive Coupling or Ground Loops

Currents flowing through the ground connections can give rise to noise voltages. This is especially a problem with reference frequency ground currents.



In this illustration, the detector is measuring the signal relative to a ground far from the rest of the experiment. The experiment senses the detector signal as well as the voltage from the noise source's ground return current, which passes through the finite resistance of the ground between the experiment and the detector. The detector and the experiment are grounded at different places which, in this case, are at different potentials.

Cures for ground loop problems include:

- 1) Grounding everything to the same physical point.

- 2) Using a heavy ground bus to reduce the resistance of ground connections.
- 3) Removing sources of large ground currents from the ground bus used for small signals.

Microphonics

Not all sources of noise are electrical in origin. Mechanical noise can be translated into electrical noise by microphonic effects. Physical changes in the experiment or cables (due to vibrations for example) can result in electrical noise over the entire frequency range of the lock-in.

For example, consider a coaxial cable connecting a detector to a lock-in. The capacitance of the cable is a function of its geometry. Mechanical vibrations in the cable translate into a capacitance that varies in time—typically at the vibration frequency. Since the cable is governed by $Q=CV$, taking the derivative yields:

$$C \frac{dV}{dt} + V \frac{dC}{dt} = \frac{dQ}{dt} = i$$

Mechanical vibrations in the cable which cause a dC/dt will give rise to a current in the cable. This current affects the detector and the measured signal.

Some ways to minimize microphonic signals are:

- 1) Eliminate mechanical vibrations near the experiment.
- 2) Tie down cables carrying sensitive signals so they do not move.
- 3) Use a low noise cable that is designed to reduce microphonic effects.

Thermocouple Effects

The emf created by junctions between dissimilar metals can give rise to many microvolts of slowly varying potentials. This source of noise is typically at very low frequency since the temperature of the detector and experiment generally changes slowly. This effect is large on the scale of many detector outputs, and can be a problem for low frequency measurements: especially in the mHz range. Some ways to minimize thermocouple effects are:

- 1) Hold the temperature of the experiment or detector constant.
- 2) Use a compensation junction, i.e. a second junction in reverse polarity which generates an emf to cancel the thermal potential of the first junction. This second junction should be held at the same temperature as the first junction.

Input Connections


In order to achieve the best accuracy for a given measurement,

care must be taken to minimize the various noise sources which can be found in the laboratory. With intrinsic noise (Johnson noise, $1/f$ noise or input noise), the experiment or detector must be designed with these noise sources in mind. These noise sources are present regardless of the input connections. The effect of noise sources in the laboratory (such as motors, signal generators, etc.), and the problem of differential grounds between the detector and the lock-in, can be minimized by careful input connections.

There are two basic methods for connecting a voltage signal to the lock-in amplifier; the single-ended connection is more convenient while the differential connection eliminates spurious pick-up more effectively.

Single-Ended Voltage Connection (A)

In the first method, the lock-in uses the A input in a single-ended mode. The lock-in detects the signal as the voltage between the center and outer conductors of the A input only. The lock-in does not force the shield of the A cable to ground. Rather, it is internally connected to the lock-in's ground via a resistor. The value of this resistor is typically between 10 Ω and 1 k Ω . The SR810, SR830 and SR850 let you choose the value of this resistor. This avoids ground loop problems between the experiment and the lock-in due to differing ground potentials. The lock-in lets the shield 'quasi-float' in order to sense the experiment ground. However, noise pickup on the shield will appear as noise to the lock-in. This is bad since the lock-in cannot reject this noise. Common mode noise, which appears on both the center and shield, is rejected by the 100 dB CMRR of the lock-in input, but noise on only the shield is not rejected at all.




Grounds may be at different potentials

Differential Voltage Connection (A–B)

The second method of connection is the differential mode. The lock-in measures the voltage difference between the center conductors of the A and B inputs. Both of the signal connections are shielded from spurious pick-up. Noise pickup on the shields does not translate into signal noise since the shields are ignored.

When using two cables, it is important that both cables travel the same path between the experiment and the lock-in. Specifically, there should not be a large loop area enclosed by

the two cables. Large loop areas are susceptible to magnetic pickup.



Common Mode Signals

Common mode signals are those signals which appear equally on both center and shield (A) or both A and B (A-B). With either connection scheme, it is important to minimize both the common mode noise and the common mode signal. Notice that the signal source is held near ground potential in both illustrations above. If the signal source floats at a nonzero potential, the signal which appears on both the A and B inputs will not be perfectly cancelled. The common mode rejection ratio (CMRR) specifies the degree of cancellation. For low frequencies, the CMRR of 100 dB indicates that the common mode signal is canceled to 1 part in 10^5 . Even with a CMRR of 100 dB, a 100 mV common mode signal behaves like a 1 μ V differential signal! This is especially bad if the common mode signal is at the reference frequency (this happens a lot due to ground loops). The CMRR decreases by about 6 dB/oct (20 dB/decade) starting at around 1 kHz.

The Lock-In as a Noise Measurement Device

Lock-in amplifiers can be used to measure noise. Noise measurements are generally used to characterize components and detectors. Remember that the lock-in detects signals close to the reference frequency. How close? Input signals within the detection bandwidth set by the low-pass-filter time constant and rolloff appear at the output at a frequency $f = f_{\text{sig}} - f_{\text{ref}}$. Input noise near f_{ref} appears as noise at the output with a bandwidth of DC to the detection bandwidth.

The noise is simply the standard deviation (root of the mean of the squared deviations) of the measured X, Y or R. You can measure this noise exactly by recording a sequence of output values and then calculating the standard deviation directly. The noise, in volts/ $\sqrt{\text{Hz}}$, is simply the standard deviation divided by the square root of the equivalent noise bandwidth of the time constant.

For Gaussian noise, the equivalent noise bandwidth (ENBW) of a low pass filter is the bandwidth of the perfect rectangular filter which passes the same amount of noise as the real filter.

Noise Estimation

The above technique, while mathematically sound, can not provide a real-time output or an analog output proportional to the measured noise. Lock-ins (such as the SR510, SR530, SR810, SR830 and SR850) do provide these features, however. The quantity X_{noise} is computed from the measured values of X using the following algorithm. The moving average of X is computed. This is the mean value of X over some past history. The present mean value of X is subtracted from the present value of X to find the deviation of X from the mean. Finally, the moving average of the absolute value of the deviations is calculated. This calculation is called the mean average deviation, or MAD. This is not the same as an rms calculation. However, if the noise is Gaussian in nature, the rms noise and the MAD noise are related by a constant factor.

SRS lock-in amplifiers use the MAD method to estimate the rms noise quantities X_n , Y_n and R_n . The advantage of this technique is its numerical simplicity and speed. For most applications, noise estimation and standard deviation calculations yield the same answer. Which method you use depends upon the requirements of the experiment.