

Application Note (A12)

The Benefits of DSP Lock-in Amplifiers

Revision: A
September 1996



Gooch & Housego

Gooch & Housego 4632 36th Street, Orlando, FL 32811
Tel: 1 407 422 3171 Fax: 1 407 648 5412 Email: sales@goochandhousego.com

If you never heard of or don't understand the term "lock-in amplifier," you're in good company. With the exception of the optics industry where virtually every major laboratory has at least one lock-in amplifier, the majority of people in the fields of science and engineering have never been exposed to this technology.

That majority is rapidly decreasing. Now that DSP (digital signal processing) technology is being applied to lock-in amplification, the improved benefits of this measurement technique are being realized in more applications than ever before.

WHY USE A LOCK-IN AMPLIFIER?

The measurement of low-level electrical signals is a difficult task. Whether due to 50/60 Hz hum pickup from power lines into a signal cable, 1/f noise from high gain preamplifiers, thermal noise from a sensor, dark current drift from a photodetector, or a combination of the above problems, noise sources of all types and frequencies hamper the ability of typical meters to give a fast, accurate reading at low signal levels.

Whether the measurement is AC or DC, AC noise and DC drift cause the reading to be unstable and increase the uncertainty in the accuracy of the reading.

Long time constants can increase the accuracy of the measurement by averaging out AC noise. However, if the meter itself experiences DC drifts during that time, how valid can the measurement be? Besides, who wants to wait for those long time constants?

Fortunately, lock-in amplification provides a technique for rejecting both AC and DC noise sources before the signal is measured. The measured signal can then be averaged to much shorter time constants, allowing faster and more accurate results.

If there is only one measurement to be made of a signal that does not change significantly with time, it may be acceptable to wait as long as it takes for the signal to settle. However, if there are multiple measurements to be made those time constants may slow results to a snail's pace.

Consider, for instance, a typical infrared detector with an r.m.s. noise level of 10mV for a one second integration period, measuring a light level that gives 100mV signal. If the desired accuracy of measurement is 1%, then the signal must be integrated over 100 seconds since the reduction in noise varies with the square root of the integration time. If this measurement were also part of a spectroradiometric scan of 500 data points, the overall scan time would be nearly 14 hours - clearly too long to be of practical use.

When performing the same scan using a lock-in amplifier, most of the noise frequencies are filtered out - reducing the noise by many decades. If, in the above example, the noise was reduced by a factor of 1000, each data point would be obtained at the desired accuracy in 0.1 seconds and the scan time would be less than one minute. This is the basis for using lock-in amplification techniques in most branches of infrared metrology.

Similarly, if the signal to be measured is time varying or coincides with some event, it may be missed entirely or altered in magnitude to the point of irrelevance if long integration times are used. A case in point is the measurement of phosphors. These may be of low intensity and decay over just a few milliseconds. To get an accurate measurement using DC methods may be impossible since noise levels at the short integration times required will generally swamp small signals. Lock-in amplification can reduce noise to levels where measurements are possible.

WHAT IS A LOCK-IN AMPLIFIER?

Lock-in amplifiers have been sometimes referred to as "frequency-selective voltmeters" and "single channel spectrum analyzers." This is because they measure the magnitude of a signal in a very narrow frequency band, while rejecting all components of the signal that are outside that very narrow frequency band.

At first glance, this sounds very simple. All that is required is a bandpass filter inserted between the signal source and a voltmeter, but this rarely achieves the desired result. The noise rejection, speed and accuracy of a good lock-in amplifier surpass what can be done with simple filtering by many orders of magnitude (1).

A lock-in amplifier is capable of measuring small AC signals that are obscured by large amounts of noise. In fact, AC signals can be extracted even when swamped by far greater noise. This ability is the basis for the most common figure of merit used to describe lock-in amplifiers, namely dynamic reserve. Dynamic reserve is usually defined as the maximum noise-to-signal ratio, expressed in dB, that a lock-in amplifier can tolerate while the measurement error is less than 5%. For common analog lock-in amplifiers this can be up to 60 dB (1000), but DSP-based lock-in amplifiers can achieve as high as 100dB (100,000) without pre-filtering (2,3).

In order to accomplish this, the lock in amplifier must be provided with a relatively clean reference signal of the same frequency as the signal to be measured.

If the signal to be measured is DC, it must be modulated by an AC waveform either electrically (as in exciting a strain gauge with an AC voltage), or mechanically (as in passing a light beam through an optical chopper). The signal and its modulation frequency (known as a reference signal) can then be fed to a lock-in amplifier.

HOW A LOCK-IN AMPLIFIER WORKS

In a fashion similar to an AM radio receiver, a lock-in amplifier uses a technique known by

many names including synchronous detection, auto-correlation, signal mixing and frequency shifting. In lock-in amplifier terminology this technique is most commonly known as phase-sensitive detection (PSD).

Phase-sensitive detection consists of multiplying the input signal by a reference signal of a constant, known magnitude, which is phase-locked to the input reference signal.

Some lock-in amplifiers use a square wave reference, others use a sine wave reference. A square wave reference contains odd harmonics of the fundamental that cause noise at those harmonic frequencies to be detected, and this is usually undesirable (4). On the other hand, a perfect sine wave reference will result in only the fundamental being detected.

The equation for the resultant, Φ , of multiplying two sine waves is (5):

$$\begin{aligned}\Phi(t) &= V_1 \sin(\omega_1 t) \cdot V_2 \sin(\omega_2 t + \theta) \\ &= 0.5 \cdot V_1 V_2 \left(\cos((\omega_1 - \omega_2)t + \theta) - \cos((\omega_1 + \omega_2)t + \theta) \right)\end{aligned}$$

Equation 1

where:

V_1 = peak magnitude of the first sine wave

ω_1 = 2π times the frequency of the first sine wave

V_2 = peak magnitude of the second sine wave

ω_2 = 2π times the frequency of the second sine wave

θ = the phase shift of the first sine wave relative to the second.

t = time.

To illustrate the lock-in amplifier process, **Figure 1** shows a signal sine wave multiplied by a phase-locked reference sine wave of the same frequency.

The equation for multiplying these sine waves reduces to:

$$\begin{aligned}\Phi(t) &= 0.5 \cdot V_1 V_2 (\cos(0) - \cos(2\omega t)) \\ &= 0.5 \cdot V_1 V_2 - 0.5 \cdot V_1 V_2 \cos(2\omega t)\end{aligned}$$

Equation 2




Figure 1. Multiplication of two phase-locked sine waves of identical frequency gives a resultant symmetrical about a DC level.




Figure 2. Multiplication of two sine waves of different frequencies gives a resultant symmetrical about zero.

Since the magnitude of one of these sine waves (the reference) is constant in a lock-in amplifier, if this resultant, $\Phi(t)$, is low pass filtered to remove the $2\omega t$ component, the constant DC voltage remaining is directly proportional to the AC signal being measured. Signal components that are not of the exact same frequency as the reference will not give this constant DC voltage, since neither the frequency sum nor the frequency difference equals zero. As seen in **Figure 2**, the resultant for different frequencies is symmetrical about zero so, following low-pass filtering, such components will not contribute to the lock-in amplifier output.

In the case of a DC offset, the lock-in amplifier output component will be the reference frequency times a constant, and will also average out to zero. If the offset drifts slowly, it represents a very low frequency, and will also average out to zero.

Noise components near to the reference frequency will result in low frequency lock-in amplifier output components. This is due, as seen in **Equation 1**, to the fact that the frequency difference term in the equation is near to, but not quite, zero. These low output frequencies, commonly referred to as “beats,” are more difficult to filter out, requiring long integration times and hence slowing measurements. For this reason, lock-in users attempt to use a reference frequency that corresponds to a quiet area of the signal's spectrum.

SIGNAL PHASE

So far the examples only illustrate the case of a signal whose phase, relative to the reference signal, is zero. When the signal phase is not zero, the output of the lock-in amplifier does not directly represent the signal magnitude. Instead, the relationship varies as $\cos(\theta)$. There are two ways to deal with this:

1. *Adjust the phase of the PSD reference vs. the input reference until it matches the signal.*
2. *Perform a dual-phase measurement.*

Contact: sales@goochandhousego.com

www.goochandhousego.com

As part of our policy of continuous product improvement, we reserve the right to change specifications at any time

A dual-phase measurement can be performed with a single-phase lock-in amplifier by taking two consecutive measurements, one with the reference phase equal to that of the signal, and one with the reference phase shifted by 90 degrees from that of the signal.

With the recorded data, both the signal magnitude and phase can be calculated. This technique assumes that the signal does not change between the two measurements.

A dual phase lock-in amplifier performs both measurements simultaneously. In a dual-phase lock-in amplifier, there are actually two PSD sections. One PSD operates on a sine wave reference, and the other operates on a cosine reference (90 degrees from the sine reference). The low-pass filtered DC output of the sine PSD is generally called the X output, whereas that of the cosine PSD is called the Y output.

If the signal is in perfect phase with the reference, the X output is the signal magnitude and the Y output is zero. If the signal is 90 degrees out of phase with the reference, the Y output is the signal magnitude and the X output is zero. But no matter what the phase, the resultant magnitude (V_{out}) of X and Y always represents the signal magnitude.

$$V_{out} = \sqrt{X^2 + Y^2}$$

Equation 3

In addition the phase of the signal can be determined by:

$$\theta = \arctan\left(\frac{Y}{X}\right)$$

Equation 4

WHY IS A DSP LOCK-IN BETTER?

In an analog lock-in amplifier, the PSD is typically implemented by one of two techniques: sine wave demodulation or square wave demodulation.

Analog sine wave demodulation

If sine wave demodulation is used, the resultant is generated using an analog multiplier. This multiplies the input signal by a sine wave generated by an analog oscillator that is phase-locked to the reference signal.

Although many advances have been made in these devices, errors on the order of 1% and harmonic distortion on the order of -60 dB (0.1%) exist even in "precision" commercially available devices.

Analog square wave demodulation

If square wave demodulation is used, an analog switch selects the output of an inverting amplifier stage when the reference is below zero and a non-inverting stage when the reference is above zero.

This technique has the disadvantage of detecting odd harmonics of the reference frequency in the input signal. This means that the user must take great care not only that there is no significant noise in the frequency band near the selected reference frequency, but also that there is no significant noise in the frequency bands near the odd harmonics of the selected reference frequency.

For both analog demodulation techniques, there are gain and offset errors that are not constant with frequency and also drift with temperature and age. There are, in addition, the problems of switching and zero crossing noise.

The output of either PSD is then fed to an analog low-pass filter to remove the AC components. The low-pass filter has gain and offset errors of its own, which also change with frequency and drift with temperature and age. In

addition, the low-pass filter order, time constants

and roll-off slopes, are fixed in hardware and due to cost consideration, a limited selection of settings are available.

DSP lock-in demodulation


The amplified signal is input directly to a high speed, high resolution A/D converter. The reference oscillator, PSD process and low-pass filtering are all done by digital components. A high quality DSP-generated sine wave contains minimal harmonic distortion, and allows the digital PSD to achieve excellent rejection of out-of-band noise. This improved rejection allows a greater dynamic reserve and signal phase angle accuracy - typically 0.001° with DSP as opposed to 1° obtained with analog lock-in amplifiers.

Assuming that the speed and resolution of the A/D and DSP chips are sufficient, the errors in the digital domain will be negligible, and will not drift. Moreover, the DSP low-pass filtering is much more flexible, since it is hardware-independent. So long as the DSP has enough memory and speed, virtually any filter type can be implemented by a change in algorithm and/or filter coefficients. Also, since the PSD is implemented in the digital domain, there is no additional hardware cost involved in providing dual-phase operation.

A COMPARISON OF RESULTS

The relative performance of analog and DSP lock-in amplifiers is easily compared in an electrical testing laboratory. It may be that "top of the range" analog units will give similar results to basic DSP lock-in amplifiers, so in order to provide a like-for-like comparison two units aimed at the same commercial market were tested.

In these tests, sinusoidal signals of 1 mV and 10 mV at various frequencies were introduced with and without a 500 mV, 60 Hz "noise" component. The output of the lock-in amplifiers



with the noise present was divided by that without the noise to provide a measure of noise rejection under difficult, but not unusual, circumstances. Results are shown in **Figure 3**. As can be seen, the DSP lock-in amplifier gives far superior performance to that of the analog amplifier.

Probably the most significant feature of DSP lock-in amplifiers is that this increase in performance does not increase the purchase price! The reason being, in the last decade, the use of DSP has grown exponentially. Driven by commercial uses in CD players, cellular phones, etc., the speed and resolution of DSP and A/D converter ICs are growing rapidly, while the costs are decreasing due to volume production.

These price/performance conditions often make

DSP the first choice of most designers who need to handle any kind of analog signal. In electronic measurement equipment, this evolution has progressed from digital multimeters through digital oscilloscopes and in the last five years, DSP lock-in amplifiers.

OTHER ADVANTAGES TO DSP LOCK-INS

Since DSP lock-in amplifiers avoid the need for some of the analog component of traditional units, they tend to be more compact. This, in turn, leads to smaller units and often more facilities available such as chopper controllers and PMT high voltage supplies being incorporated rather than as separate boxes. Also, fully functional PC boards offering the same, and sometimes superior, facilities of larger traditional stand-alone systems are possible.

DISADVANTAGES TO DSP LOCK-INS

The DSP lock-in method does have a limitation, namely, the sampling rate and resolution of the A/D converter and DSP ICs. To avoid DC offset and sampling anomalies, the maximum frequency of operation is limited to half the sampling rate. However, this limit, typically around 100 kHz, affects only a very small percentage of applications. As the technology progresses, even these very high frequency applications are being accommodated.

CONCLUSION

DSP lock-in amplifiers provide a cost-effective means of rejecting large amounts of signal noise. They are generally more accurate than analog lock-in amplifiers, especially at lower frequencies, and they do a better job of rejecting noise, due primarily to the fact that they produce less harmonic distortion in the PSD section. They provide more flexibility in terms of filter customization and their sampling rate/resolution and general signal architecture make them adaptable to different tasks in addition to lock-in amplification. These applications include FFT frequency spectrum measurements, boxcar averaging, pulse profiling, differential and integrating amplification.

References

1. NASA Tech Briefs, **Technical Support Package GSC-13179**, pp.1
 2. "OL 730D Programmable Lock-in Amplifier," Bulletin 134, **Rev. 9/96**, Optronic Laboratories, Inc.
 3. "OL 4000 DSP Lock-in Amplifier Card," Bulletin 131, **Rev. 6/96**, Optronic Laboratories, Inc.
 4. NASA Tech Briefs, **Technical Support Package GSC-13273**, pp.4
 5. Meade, M.L., "*Lock-in Amplifiers: Principles and Applications*," **1983**, Peter Peregrinus Ltd., chapter 2, pp.16
-