

THz Sensitivity Characterization of a LiTaO₃ Pyroelectric Detector

Background

The THz field has advanced dramatically during the past 10 years largely in large part because of the development of useful solid-state devices, components, and systems operating at room temperature. Among the successful components are Schottky diode mixers, rectifiers, and varactor multiplier; ultrafast photoconductive (aka “Auston”) switches and photomixers; millimeter-wave low-noise and power amplifiers; and room-temperature thermal detectors (bolometers and pyroelectrics). Among the successful components and systems are varactor multiplier chains, Schottky-diode-based modularized heterodyne receiver front ends, scalar network analyzers, time-domain spectrometers, and photomixing sweep oscillators.

The goal of this laboratory project is to give the students “hands-on” experience with one of the most useful and affordable of these devices – the pyroelectric thermal detector. Measuring power – be it thermal or coherent radiation – is one of the most common and important THz tasks, and it has long been somewhat controversial. Researchers developing new THz oscillators have often measured the output power with cryogenic bolometers that are poorly understood and characterized by spectral variations and saturation effects. So they may likely underestimate or overestimate the output of their oscillator, leading to disagreement with other researchers that impedes the progress of the field as a whole. This is a primary reason for the “standardization” that has occurred in nearly every RF band below the THz region over the past 50 years. Unfortunately, the THz field is still a bit too primitive to create such standardization, although there are certain efforts in this direction.

If sensitive enough and spectrally uniform, the pyroelectric detector has a good chance of becoming a THz standard in the next few years. The sensitivity requirements are roughly $NEP' \sim 0.1 \text{ nW/Hz}^{1/2}$, and $NE\Delta T' \sim 1.0 \text{ K/Hz}^{1/2}$. Both are comparable to the THz sensitivity of the best Golay cells. But the pyroelectric detectors are much more rugged, cost much less (< 10%), and are amenable to fabrication in one- or two-dim arrays. Judging from how important uncooled thermal detector arrays (so-called “focal plane” arrays) have been in the mid-IR bands, a room-temperature THz array is an exciting possibility.

Laboratory Procedure

The experiments will all be carried out in the Laboratory of Dr. E.R. Brown, 2209 Engineering Sciences Building. Each student will be assigned their own pyroelectric detector, a Model SPH-45 (manufactured by Spectrum Detector, Inc.) mounted in a TO-5 can. The detector is a single-crystal LiTaO₃ capacitor with a front-side chrome electrode that doubles as the absorbing layer. The pyroelectric detector is coupled to a transimpedance amplifier (TIA) that is packaged on the same TO-5 header. The pyroelectric capacitor and TIA are biased by plugging the TO-5 pins into a socket in the Spectrum Detector Bias box (SPH-CM-test). Make sure to first turn the bias voltage off (switch on side) and wear an ESD-mitigating wrist strap during installation of the TO-5 pins. Also make sure the TO-5 pins are aligned properly with respect to the socket on the PCB of the Bias Box (the tab on the TO-5 should be directly above the tab printed on the PCB) before applying force. Once the TO-5 is properly installed, re-assemble the SPH-CM bias box, making sure the two pieces of black polyethylene are placed in

the optical window to block background IR radiation from the pyroelectric detector. Then mount the SPH-CM on the Gimbal mount (NEP set-up) or the optical bench (NEAT set-up). Photographs of these two set-ups are provided in Figs. 1 and 2, respectively.

Note that two “grey-body” sources are provided in the NEAT set-up. They both consist of tap water in a polyethylene wash bottle which is heated internally by wire-wound resistors and a separate power supply. One of the wash bottles is maintained at a “hot” temperature T_2 and the other at a “warm” temperature T_1 . Measurement of the response of the pyroelectric detector at these two temperatures will be used below to determine the NEAT. Please do not change the settings on the power-supplies as it will cause a change in the water temperature that takes a long time to reach steady state. Also, we collect the results from all the pyroelectric detectors at the end of the project to assess their uniformity. Changing the temperature will add a systematic error and render the uniformity results less useful.

A. Determination of the Optical NEP

A.1. Measurement of the Optical Responsivity

Dr. Brown’s lab is fortunate to have a small collection of “Frequency Extension Modules” (FEMs) manufactured by Virginia Diodes, Inc. These have been used in the past to characterize some of the most sensitive room-temperature millimeter-wave and THz detectors in the world. The one we will use is tunable between 600 and 650 GHz. We will use it a single “spot” frequency of 600 GHz where the FEM is known to put out ~ 1.0 mW (this number may change so please be ready to record the latest estimate during your session). The output of the detector is measured with a Stanford Research Model SR-830 digital lock-in amplifier (LIA).

After inserting your pyroelectric detector in the socket of the Bias Box, mount the bias box on the gimbal mount shown in Fig. 1. The gimbal has been pre-aligned with respect to the FEM feedhorn so that the pyroelectric detector should be located close to the “bore-sight” position. Then turn on the bias to the detector and TIA (switch on side of bias box). Make sure Channel 1 of the LIA is set for magnitude (R), and that a healthy reading appears Channel 1 of the LIA (a few mV). Using the MATLAB Script provided, rotate the pyroelectric detector in azimuth and elevation over 90 degrees for each. The two-dim antenna pattern is displayed as the rotation occurs. After the rotation is done, save your data to a .txt file. Then either by inspection or using a program (e.g., Excel), find the maximum pyroelectric output signal, making sure your maximum value is not a “glitch.” Then convert your value to an optical responsivity (Question#1 or Write-Up).

A.2. Measurement of the Electrical Noise

One of the most mysterious yet essential tasks of THz sensor characterization, or any sensor characterization for that matter, is measuring the electrical noise. We are fortunate to be using the SR-830 because it measures the noise automatically using the MAD method (mean absolute deviation). This method is not as crazy as it may sound. And because it is done by sampling and digital computation, the results are independent of the noise statistics and of the bandwidth (i.e., time constant) of the LIA. The noise readings are computed for both the in-phase (I) and quadrature (Q) components of the signal, consistent with the fact that electrical noise generally has two degrees-of-freedom [amplitude and phase, or $X (=I)$ and $Y (=Q)$ in an Argand representation].

To take the noise readings, block the FEM feedhorn with a metal plate. Set the SR830 LIA to Xnoise on Channel 1 and Ynoise on Channel 2. Wait for the displays to reach steady state (this may take a few minutes). Then record both readings. You will use these along with the optical responsivity value above to derive the optical NEP in the Laboratory Write-Up

A minor drawback of the MAD method is that there is a memory effect so that any changes in the experiment (or in the sensitivity setting of the LIA) will require some waiting before the X and Y noise readings settle. Note, however, that the time constant of the LIA does not change the readings. More detail on the MAD method can be found in the last few pages of the attached SR830 documentation,¹ and some of these issues are addressed in the Laboratory Write-Up

B. Determination of the Optical NEAT

B.1. Setting up the THz Water-Based Grey Body

At the beginning of the course we saw how water has very interesting dielectric properties in the THz region in that the real part of the dielectric function drops rapidly and the imaginary part increases with increasing frequency. So the absorption coefficient increases, but the reflection coefficient drops compared to low frequency values, making the emissivity close to unity (recall, according to Kirchoff's relation, $e = 1 - R$). Therefore, water can be used as calibrated source of THz radiation just by heating it well above ambient and holding it in a transparent container such as a styrofoam cup or a polyethylene wash bottle. We will assume the emissivity is less than unity, but independent of wavelength, which simplifies the task of estimating the incident power. This is commonly called the "grey body" estimation in infrared science and technology.²

The first step in the NEAT experiment is to insert your pyroelectric detector into the SPH-CM bias box in the laboratory set-up. Again, practice precaution by wearing a wrist-strap during the installation, and make sure to carefully align the pins on the TO-5 header with the socket receptacles in the bias box before pushing the TO-5 can into place. Once the installation is done, re-assemble the SPH-CM bias box, making sure the two pieces of black polyethylene are placed in the optical window to block background IR radiation from striking the pyroelectric detector. Then re-mount on the optical bench top as shown in Fig. 2.

B.2. Measuring the Grey-Body Responsivity and Noise

Start by turning on the bias to the SPH-CM and the controller for the optical chopper to 10 Hz. Make sure the chopper blade is as close as possible to the SPH-CM but doesn't touch it. Re-locate the THz turning mirror at approximately 45 degrees with respect to the chopper blade, and locate the hot water bottle as shown in Fig. 2. Put the thermocouple midway into the hot water bottle but not too close to the heating resistors. Record the temperature in your notebook. Take the reading from the SR-830 LIA with the Channel 1 display set to magnitude (R). Adjust the position of the turning mirror manually to maximize the reading. Record this in your notebook as V_H .

Next turn off the chopper and block the pyroelectric detector with a solid sector of the blade. Change the settings on the LIA to Xnoise on the Channel 1 display and Ynoise on the

¹ an even better reference is http://en.wikipedia.org/wiki/Mean_absolute_deviation

² See, for example, http://www.efunda.com/formulae/heat_transfer/radiation/blackbody.cfm

Channel 2 display. Wait for the noise levels to settle (this may take a few minutes). Then record these two readings in your notebook.

Finally swap out the hot water bottle for the warm water bottle in the same location. Re-set the LIA to read R on Channel 1, re-optimize the signal and record V_W . Then repeat the temperature and noise measurements as above. You will use the V_H , V_W and noise readings in the Write-Up to derive the $NE\Delta T$.




Fig. 1.




Fig. 2.

WHAT IS A LOCK-IN AMPLIFIER?

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. A good low noise amplifier will have about 5 nV/ $\sqrt{\text{Hz}}$ of input noise. If the amplifier bandwidth is 100 kHz and the gain is 1000, then we can expect our output to be 10 μV of signal (10 nV x 1000) and 1.6 mV of broadband noise (5 nV/ $\sqrt{\text{Hz}}$ x $\sqrt{100 \text{ kHz}}$ x 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/ $\sqrt{\text{Hz}}$ x $\sqrt{100 \text{ Hz}}$ x 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 only 0.5 μV (5 nV/ $\sqrt{\text{Hz}}$ x $\sqrt{0.01 \text{ Hz}}$ x 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 diagram below, 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 $V_{\text{sig}} \sin(\omega_r t + \theta_{\text{sig}})$ where V_{sig} is the signal amplitude.

The SR830 generates its own sine wave, shown as the lock-in reference below. The lock-in reference is $V_L \sin(\omega_L t + \theta_{\text{ref}})$.



The SR830 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.

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

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

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

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

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

Narrow band detection

Now suppose 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 roll-off. 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 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.

All lock-in measurements require a reference signal.

In this case, the reference is provided by the excitation source (the function generator). This is called an external reference source. In many situations, the SR830'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_{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

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

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

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, such as the SR830, has two PSD's, with reference oscillators 90° apart, and can measure X, Y and R directly. In addition, the phase θ between the signal and lock-in reference, can be measured according to

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

WHAT DOES A LOCK-IN MEASURE?

So what exactly does the SR830 measure? Fourier's theorem basically states that any input signal can be represented as the sum of many, 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.

What does the SR830 measure?

The SR830 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. In the SR830, 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.

The SR830, 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 2V pk-pk 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 SR830, locked to f will single out the first component. The measured signal will be $1.273\sin(\omega t)$, not the 2V pk-pk 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 at 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 the SR830 displays a magnitude of 1V (rms), the component of the input signal at the reference frequency is a sine wave with an amplitude of 1 Vrms or 2.8 V pk-pk.

Thus, in the previous example with a 2 V pk-pk square wave input, the SR830 would detect the first sine component, $1.273\sin(\omega t)$. The measured and displayed magnitude would be 0.90 V (rms) ($1/\sqrt{2} \times 1.273$).


Degrees or Radians?

In this discussion, frequencies have been referred to as f (Hz) and ω ($2\pi f$ radians/sec). 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.

THE FUNCTIONAL SR830

The functional block diagram of the SR830 DSP Lock-In Amplifier is shown below. The functions in the gray area are handled by the digital signal processor (DSP). We'll discuss the DSP aspects of the SR830 as they come up in each functional block description.



THE PHASE SENSITIVE DETECTORS (PSD's)

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. The signal amplifier and filters will be discussed later.

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

Digital PSD vs Analog PSD

The phase sensitive detectors (PSD's) in the SR830 act as linear multipliers, that is, they multiply the signal with a reference sine wave. Analog PSD's (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 SR830 is 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 PSD's 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 totally free from errors.

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 quite a bit of output uncertainty.

In the digital lock-in, the 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 SR830 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 the SR830 can exceed 100 dB without any problems. We'll talk more about dynamic reserve a little later.

An analog linear 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 eliminates a major source of gain error in a linear analog lock-in.

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.

We've discussed how the digital signal processor in the SR830 computes the internal oscillator and two reference sine waves and handles both phase sensitive detectors. In the next section, we'll see the same DSP perform the low pass filtering and DC amplification required at the output of the PSD's. Here again, the digital technique eliminates many of the problems associated with analog lock-in amplifiers.

TIME CONSTANTS and DC GAIN

Remember, the output of the PSD contains many signals. Most of the output signals have frequencies which are either the sum or difference between an input signal frequency and the reference frequency. Only the component of the input signal whose frequency is exactly equal to the reference frequency will result in a DC output.

The low pass filter at the PSD output removes all of the unwanted AC signals, both the 2F (sum of the signal and the reference) and the noise components. This filter is what makes the lock-in such a narrow band detector.

Time Constants

Lock-in amplifiers have traditionally set the low pass filter bandwidth by setting the time constant. The time constant is simply $1/2\pi f$ where f is the -3 dB frequency of the filter. The low pass filters are simple 6 dB/oct roll off, RC type filters. A 1 second time constant referred to a filter whose -3 dB point occurred at 0.16 Hz and rolled off at 6 dB/oct beyond 0.16 Hz. Typically, there are two successive filters so that the overall filter can roll off at either 6 dB or 12 dB per octave. The time constant referred to the -3 dB point of each filter alone (not the combined filter).

The notion of time constant arises from the fact that the actual output is supposed to be a DC signal. In fact, when there is noise at the input, there is noise on the output. By increasing the time constant, the output becomes more steady and easier to measure reliably. The trade off comes when real changes in the input signal take many time constants to be reflected at the output. This is because a single RC filter requires about 5 time constants to settle to its final value. The time constant reflects how slowly the output responds, and thus the degree of output smoothing.

The time constant also determines the equivalent noise bandwidth (ENBW) for noise measurements. The ENBW is NOT the filter -3 dB pole, it is the effective bandwidth for Gaussian noise. More about this later.

Digital Filters vs Analog Filters

The SR830 improves on analog filters in many ways. First, analog lock-ins provide at most, two

stages of filtering with a maximum roll off of 12 dB/oct. This limitation is usually due to space and expense. Each filter needs to have many different time constant settings. The different settings require different components and switches to select them, all of which is costly and space consuming.

The digital signal processor in the SR830 handles all of the low pass filtering. Each PSD can be followed by up to four filter stages for up to 24 dB/oct of roll off. Since the filters are digital, the SR830 is not limited to just two stages of filtering.

Why is the increased roll off desirable? Consider an example where the reference is at 1 kHz and a large noise signal is at 1.05 kHz. The PSD noise outputs are at 50 Hz (difference) and 2.05 kHz (sum). Clearly the 50 Hz component is the more difficult to low pass filter. If the noise signal is 80 dB above the full scale signal and we would like to measure the signal to 1% (-40 dB), then the 50 Hz component needs to be reduced by 120 dB. To do this in two stages would require a time constant of at least 3 seconds. To accomplish the same attenuation in four stages only requires 100 ms of time constant. In the second case, the output will respond 30 times faster and the experiment will take less time.

Synchronous Filters

Another advantage of digital filtering is the ability to do synchronous filtering. Even if the input signal has no noise, the PSD output always contains a component at 2F (sum frequency of signal and reference) whose amplitude equals or exceeds the desired DC output depending upon the phase. At low frequencies, the time constant required to attenuate the 2F component can be quite long. For example, at 1 Hz, the 2F output is at 2 Hz and to attenuate the 2 Hz by 60 dB in two stages requires a time constant of 3 seconds.

A synchronous filter, on the other hand, operates totally differently. The PSD output is averaged over a complete cycle of the reference frequency. The result is that all components at multiples of the reference (2F included) are notched out completely. In the case of a clean signal, almost no additional filtering would be required. This is

increasingly useful the lower the reference frequency. Imagine what the time constant would need to be at 0.001 Hz!

In the SR830, synchronous filters are available at detection frequencies below 200 Hz. At higher frequencies, the filters are not required (2F is easily removed without using long time constants). Below 200 Hz, the synchronous filter follows either one or two stages of normal filters. The output of the synchronous filter is followed by two more stages of normal filters. This combination of filters notches all multiples of the reference frequency and provides overall noise attenuation as well.

Long Time Constants

Time constants above 100 seconds are difficult to accomplish using analog filters. This is simply because the capacitor required for the RC filter is prohibitively large (in value and in size!). Why would you use such a long time constant? Sometimes you have no choice. If the reference is well below 1 Hz and there is a lot of low frequency noise, then the PSD output contains many very low frequency components. The synchronous filter only notches multiples of the reference frequency, the noise is filtered by the normal filters.

The SR830 can provide time constants as long as 30000 seconds at reference frequencies below 200 Hz. Obviously you don't use long time constants unless absolutely necessary, but they're available.

DC Output Gain

How big is the DC output from the PSD? It depends on the dynamic reserve. With 60 dB of dynamic reserve, a noise signal can be 1000 times (60 dB) greater than a full scale signal. At the PSD, the noise can not exceed the PSD's input range. In an analog lock-in, the PSD input range might be 5V. With 60 dB of dynamic reserve, the signal will be only 5 mV at the PSD input. The PSD typically has no gain so the DC output from the PSD will only be a few millivolts! Even if the PSD had no DC output errors, amplifying this millivolt signal up to 10 V is error prone. The DC output gain needs to be about the same as the dynamic reserve (1000 in this case) to provide a 10 V output for a full scale input signal. An offset as small as 1 mV will appear as 1 V at the output! In fact, the PSD output offset plus the input offset of the DC amplifier needs to be on the order of 10 μ V in order to not affect the measurement. If

the dynamic reserve is increased to 80dB, then this offset needs to be 10 times smaller still. This is one of the reasons why analog lock-ins do not perform well at very high dynamic reserve.

The digital lock-in does not have an analog DC amplifier. The output gain is yet another function handled by the digital signal processor. We already know that the digital PSD has no DC output offset. Likewise, the digital DC amplifier has no input offset. Amplification is simply taking input numbers and multiplying by the gain. This allows the SR830 to operate with 100 dB of dynamic reserve without any output offset or zero drift.

What about resolution?

Just like the analog lock-in where the noise can not exceed the input range of the PSD, in the digital lock-in, the noise can not exceed the input range of the A/D converter. With a 16 bit A/D converter, a dynamic reserve of 60 dB means that while the noise has a range of the full 16 bits, the full scale signal only uses 6 bits. With a dynamic reserve of 80 dB, the full scale signal uses only 2.5 bits. And with 100 dB dynamic reserve, the signal is below a single bit! Clearly multiplying these numbers by a large gain is not going to result in a sensible output. Where does the output resolution come from?

The answer is filtering. The low pass filters effectively combine many data samples together. For example, at a 1 second time constant, the output is the result of averaging data over the previous 4 or 5 seconds. At a sample rate of 256 kHz, this means each output point is the exponential average of over a million data points. (A new output point is computed every 4 μ s and is a moving exponential average). What happens when you average a million points? To first order, the resulting average has more resolution than the incoming data points by a factor of million. This represents a gain of 20 bits in resolution over the raw data. A 1 bit input data stream is converted to 20 bits of output resolution.

The compromise here is that with high dynamic reserve (large DC gains), some filtering is required. The shortest time constants are not available when the dynamic reserve is very high. This is not really a limitation since presumably there is noise which is requiring the high dynamic reserve and thus substantial output filtering will also be required.

DC OUTPUTS and SCALING

The SR830 has X and Y outputs on the rear panel and Channel 1 and 2 (CH1 and CH2) outputs on the front panel.

X and Y Rear Panel Outputs

The X and Y rear panel outputs are the outputs from the two phase sensitive detectors with low pass filtering, offset and expand. These outputs are the traditional outputs of an analog lock-in. The X and Y outputs have an output bandwidth of 100 kHz.

CH1 and CH2 Front Panel Outputs

The two front panel outputs can be configured to output voltages proportional to the CH1 and CH2 displays or X and Y.

If the outputs are set to X or Y, these outputs duplicate the rear panel outputs.

If they are set to Display, the output is updated at 512 Hz. The CH1 display can be defined as X, R, X Noise, Aux Input 1 or 2, or any of these quantities divided by Aux Input 1 or 2. The CH2 display can be defined as Y, θ , Y Noise, Aux Input 3 or 4, or any of these quantities divided by Aux Input 3 or 4. If a display is defined as simply X or Y, this display, when output through the CH1 or CH2 output BNC, will only update at 512 Hz. It is better in this case to set output to X or Y directly, rather than the display.

X, Y, R and θ Output scales

The sensitivity of the lock-in is the rms amplitude of an input sine (at the reference frequency) which results in a full scale DC output. Traditionally, full scale means 10 VDC at the X, Y or R BNC output. The overall gain (input to output) of the amplifier is then 10 V/sensitivity. This gain is distributed between AC gain before the PSD and DC gain following the PSD. Changing the dynamic reserve at a given sensitivity changes the gain distribution while keeping the overall gain constant.

The SR830 considers 10 V to be full scale for any output proportional to simply X, Y or R. This is the output scale for the X and Y rear panel outputs as well as the CH1 and CH2 outputs when configured to output X or Y. When the CH1 or CH2 outputs are proportional to a display which is simply

defined as X, Y or R, the output scale is also 10 V full scale.

Lock-in amplifiers are designed to measure the RMS value of the AC input signal. All sensitivities and X, Y and R outputs and displays are RMS values.

Phase is a quantity which ranges from -180° to $+180^\circ$ regardless of the sensitivity. When CH2 outputs a voltage proportional to θ , the output scale is $18^\circ/\text{Volt}$ or $180^\circ=10\text{V}$.

X, Y and R Output Offset and Expand

The SR830 has the ability to offset the X, Y and R outputs. This is useful when measuring deviations in the signal around some nominal value. The offset can be set so that the output is offset to zero. Changes in the output can then be read directly from the display or output voltages. The offset is specified as a percentage of full scale and the percentage does not change when the sensitivity is changed. Offsets up to $\pm 105\%$ can be programmed.

The X, Y and R outputs may also be expanded. This simply takes the output (minus its offset) and multiplies by an expansion factor. Thus, a signal which is only 10% of full scale can be expanded to provide 10 V of output rather than only 1 V. The normal use for expand is to expand the measurement resolution around some value which is not zero. For example, suppose a signal has a nominal value of 0.9 mV and we want to measure small deviations, say 10 μV or so, in the signal. The sensitivity of the lock-in needs to be 1 mV to accommodate the nominal signal. If the offset is set to 90% of full scale, then the nominal 0.9 mV signal will result in a zero output. The 10 μV deviations in the signal only provide 100 mV of DC output. If the output is expanded by 10, these small deviations are magnified by 10 and provide outputs of 1 VDC.

The SR830 can expand the output by 10 or 100 provided the expanded output does not exceed full scale. In the above example, the 10 μV deviations can be expanded by 100 times before they exceed full scale (at 1 mV sensitivity).

The analog output with offset and expand is

$$\text{Output} = (\text{signal/sensitivity} - \text{offset}) \times \text{Expand} \times 10\text{V}$$

where offset is a fraction of 1 (50%=0.5), expand is 1, 10 or 100, and the output can not exceed 10 V. In the above example,

$$\text{Output} = (0.91\text{mV}/1\text{mV} - 0.9) \times 10 \times 10\text{V} = 1\text{V}$$

for a signal which is 10 μV greater than the 0.9 mV nominal. (Offset = 0.9 and expand =10).

The X and Y offset and expand functions in the SR830 are output functions, they do NOT affect the calculation of R or θ . R has its own output offset and expand.

CH1 and CH2 Displays

The CH1 display can show X, R, X Noise, Aux Input 1 or 2, or any of these quantities divided by Aux Input 1 or 2. The CH2 display can show Y, θ , Y Noise, Aux Input 3 or 4, or any of these quantities divided by Aux Input 3 or 4.

Output offsets ARE reflected in the displays. For example, if CH1 is displaying X, it is affected by the X offset. When the X output is offset to zero, the displayed value will drop to zero also. Any display which is showing a quantity which is affected by a non-zero offset will display a highlighted **Offset** indicator below the display.

Output expands do NOT increase the displayed values of X, Y or R. Expand increases the resolution of the X, Y or R value used to calculate the displayed value. For example, CH1 when displaying X does not increase its displayed value when X is expanded. This is because the expand function increases the resolution with which the signal is measured, not the size of the input signal. The displayed value will show an increased resolution but will continue to display the original value of X minus the X offset. Any display which is showing a quantity which is affected by a non-unity expand will display a highlighted **Expand** indicator below the display.

Ratio displays are displayed as percentages. The displayed percentage for X/Aux 1 would be

$$\text{Display \%} = \frac{(\text{signal/sensitivity}-\text{offset}) \times \text{Expand} \times 100}{\text{Aux In 1 (in Volts)}}$$

where offset is a fraction of 1 (50%=0.5), expand is 1, 10 or 100, and the display can not exceed 100%.

For example, if the sensitivity is 1V and CH1 display is showing X/Aux 1. If X= 500 mV and Aux 1= 2.34 V, then the display value is $(0.5/1.0) \times 100 / 2.34$ or 21.37%. This value is affected by the sensitivity, offset and X expand.

In the case of θ , the full scale sensitivity is always 180°.

The **Ratio** indicator below the display is on whenever a display is showing a ratio quantity.

Display output scaling

What about CH1 or CH2 outputs proportional to ratio displays? The output voltage will simply be the displayed percentage times 10V full scale.

In the above example, the displayed ratio of 21.37% will output 2.137V from the CH1 output.

DYNAMIC RESERVE

We've mentioned dynamic reserve quite a bit in the preceding discussions. It's time to clarify dynamic reserve a bit.

What is dynamic reserve really?

Suppose 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 then 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 noise 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 1000ppm/ $^{\circ}$ 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 0dB. 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 roll off. 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 ($f_{\text{noise}} \gg 100 \text{ kHz}$). In these cases, the reserve can be higher at these frequencies than within the operating range. While this may be 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 the SR830

The SR830, with its digital phase sensitive detectors, does 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 ($>60 \text{ dB}$) 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 above 60 dB AND set to High Reserve or Normal. However, the Low Noise reserve can be very high as we'll see shortly.

To set a scale, the SR830'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 SR830'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 SR830'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 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 SR830, 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 higher but is generally not that useful.

Minimum dynamic reserve (Low Noise)

The SR830 always has a minimum amount of dynamic reserve. This minimum reserve is the Low Noise reserve setting. The 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 SR830 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 SR830, more reserve does not increase the output error or drift. More reserve can increase the output noise though. 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 SR830 does not increase the 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.

SIGNAL INPUT AMPLIFIER and FILTERS

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 SR830 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 plus 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 SR830 signal amplifier is about 5 nVrms/ $\sqrt{\text{Hz}}$. What does this noise figure mean? Let's set up an experiment. If an amplifier has 5 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 roll off) 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 μV pk-pk 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. See the discussion of noise later in this section for more information on ENBW. The ENBW depends upon the time constant and filter roll off. For example, suppose the SR830 is set to 5 μV full scale

with a 100 ms time constant and 6 dB/oct of filter roll off. The ENBW of a 100 ms, 6 dB/oct filter is 2.5 Hz. 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 μV). 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 2k Ω will have a Johnson noise greater than the SR830'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 2k Ω source impedance is used, the Johnson noise will be 5.8 nVrms/ $\sqrt{\text{Hz}}$. The overall noise at the SR830 input will be $[5^2 + 5.8^2]^{1/2}$ or 7.7 nVrms/ $\sqrt{\text{Hz}}$.

We'll talk more about noise sources later in this section.

At lower gains (sensitivities above 50 μV), there is not enough gain at high reserve to amplify the input noise to a level greater than the noise of the A/D converter. In these cases, the output noise is determined by the A/D noise. Fortunately, at these sensitivities, the DC gain is low and the noise at the output is negligible.

Notch filters

The SR830 has two notch filters in the signal amplifier chain. These are pre-tuned to the line frequency (50 or 60 Hz) and twice the line frequency (100 or 120 Hz). In circumstances where the largest noise signals are at the power line frequencies, these filters can be engaged to remove noise signals at these frequencies. Removing the largest noise signals before the final gain stage can reduce the amount of dynamic reserve required to perform a measurement. To the extent that these filters reduce the required reserve to either 60 dB or the minimum reserve (whichever is higher), then some improvement might be gained. If the required reserve without these notch filters is below 60 dB or if the minimum reserve is sufficient, then these filters do not significantly improve

Current Input (I)

The current input on the SR830 uses the A input BNC. The current input has a 1 k Ω input impedance and a current gain of either 10^6 or 10^8 Volts/Amp. Currents from 1 μ A down to 2 fA full scale can be measured.

The impedance of the signal source is the most important factor to consider in deciding between voltage and current measurements.

For high source impedances, greater than 1 M Ω (10^6 gain) or 100 M Ω (10^8 gain), and small currents, use the current input. Its relatively low impedance greatly reduces the amplitude and phase errors caused by the cable capacitance-source impedance time constant. The cable capacitance should still be kept small to minimize the high frequency noise gain of the current preamplifier.

For moderate to low source impedances, or larger currents, the voltage input is preferred. A small value resistor may be used to shunt the signal current and generate a voltage signal. The lock-in then measures the voltage across the shunt resistor. Select the resistor value to keep the shunt voltage small (so it does not affect the source current) while providing enough signal for the lock-in to measure.

Which current gain should you use? The current gain determines the input current noise of the lock-in as well as its measurement bandwidth. Signals far above the input bandwidth are attenuated by 6 dB/oct. The noise and bandwidth are listed below.

<u>Gain</u>	<u>Noise</u>	<u>Bandwidth</u>
10^6	130 fA/ $\sqrt{\text{Hz}}$	70 kHz
10^8	13 fA/ $\sqrt{\text{Hz}}$	700 Hz

AC vs DC Coupling

The signal input can be either AC or DC coupled. The AC coupling high pass filter passes signals above 160 mHz (0.16 Hz) and attenuates signals at lower frequencies. AC coupling should be used at frequencies above 160 mHz whenever possible. At lower frequencies, DC coupling is required.

A DC signal, if not removed by the AC coupling filter, will multiply with the reference sine wave and produce an output at the reference frequency. This signal is not normally present and needs to be removed by the low pass filter. If the DC component of the signal is large, then this output will be large and require a long time constant to remove. AC coupling removes the DC component of the signal without any sacrifice in signal as long as the frequency is above 160 mHz.

The current input current to voltage preamplifier is always DC coupled. AC coupling can be selected following the current preamplifier to remove any DC current signal.

INTRINSIC (RANDOM) NOISE SOURCES

Random noise finds its way into experiments in a variety of ways. Good experimental design can reduce these noise sources and improve the measurement stability and accuracy.

There are a variety of intrinsic noise sources which are present in all electronic signals. These sources are physical in origin.

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,

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

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

Since the input signal amplifier in the SR830 has a bandwidth of approximately 300 kHz, the effective noise at the amplifier input is $V_{\text{noise}} = 70\sqrt{R}$ nVrms or $350\sqrt{R}$ nV pk-pk. This noise is broadband and if the source impedance of the signal is large, 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

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

The ENBW is determined by the time constant and slope as shown in the following table. Wait time is the time required to reach 99% of its final value.

T= Time Constant

Slope	ENBW	Wait Time
6 dB/oct	1/(4T)	5T
12 dB/oct	1/(8T)	7T
18 dB/oct	3/(32T)	9T
24 dB/oct	5/(64T)	10T

The signal amplifier bandwidth determines the amount of broadband noise that will be amplified. This affects the dynamic reserve. The time constant sets the amount of noise which will be measured at the reference frequency. See the SIGNAL INPUT AMPLIFIER discussion for more information about Johnson noise.

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

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

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 μV -3 μV of rms noise per Volt of 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 in the 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 a 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 $f = 60 \text{ Hz}$ and $V_{\text{noise}} = 120 \text{ V}$. C_{stray} can be estimated using a parallel plate equivalent capacitor. If the capacitance is roughly an area of 1 cm^2 at a 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, then 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 2 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 plus the voltage due to the noise source's ground return current passing 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 $Q=CV$,

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

taking the derivative, we have

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.

NOISE MEASUREMENTS

Lock-in amplifiers can be used to measure noise. Noise measurements are generally used to characterize components and detectors.

The SR830 measures input signal noise AT the reference frequency. Many noise sources have a frequency dependence which the lock-in can measure.

How does a lock-in measure noise?

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 roll-off 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.

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.

The ENBW is determined by the time constant and slope as shown below. Wait time is the time required to reach 99% of its final value.

T= Time Constant

Slope	ENBW	Wait Time
6 dB/oct	1/(4T)	5T
12 dB/oct	1/(8T)	7T
18 dB/oct	3/(32T)	9T
24 dB/oct	5/(64T)	10T

Noise estimation

The noise is simply the standard deviation (root of the mean of the squared deviations) of the measured X, Y or R.

The above technique, while mathematically sound, can not provide a real time output or an analog output proportional to the measured noise. For these measurements, the SR830 estimates the X or Y noise directly.

To display the noise of X, for example, simply set the CH1 display to X noise. 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, then the RMS noise and the MAD noise are related by a constant factor.

The SR830 uses the MAD method to estimate the RMS noise of X and Y. The advantage of this technique is its numerical simplicity and speed.

The noise calculations for X and Y occur at 512 Hz. At each sample, the mean and moving average of the absolute value of the deviations is calculated. The averaging time (for the mean and average deviation) depends upon the time constant. The averaging time is selected by the SR830 and ranges from 10 to 80 times the time constant. Shorter averaging times yield a very poor estimate of the noise (the mean varies rapidly and the deviations are not averaged well). Longer averaging times, while yielding better results, take a long time to settle to a steady answer.

To change the settling time, change the time constant. Remember, shorter settling times use smaller time constants (higher noise bandwidths) and yield noisier noise estimates.

X and Y noise are displayed in units of Volts/ $\sqrt{\text{Hz}}$. The ENBW of the time constant is already factored into the calculation. Thus, the mean displayed value of the noise should not depend upon the time constant.

The SR830 performs the noise calculations all of the time, whether or not X or Y noise are being displayed. Thus, as soon as X noise is displayed, the value shown is up to date and no settling time is required. If the sensitivity is changed, then the noise estimate will need to settle to the correct value.