

Ph 77 – ADVANCED PHYSICS LABORATORY

Lab 2 - Small-Signal Detection Using the Lock-In Amplifier

I. BACKGROUND

Modern physics research often involves observing small signals buried in noise. Consider an experiment in which we wish to measure the response of a system to some stimulus, and we know that the response is very weak. Furthermore the output of the system, which is typically converted to an electronic signal, contains a lot of electronic noise we can't get rid of. Put in terms of voltages, our output signal looks like

$$V(t) = V_0 + V_{noise}(t)$$

where V_0 is the response of the system we want to measure and $V_{noise}(t)$ is noise.

Given this situation, one approach is simply to average the output signal with respect to time. If the noise is well behaved and has zero expectation value, $\langle V_{noise} \rangle = 0$, then by averaging we will obtain $\langle V(t) \rangle = \langle V_0 + V_{noise}(t) \rangle = \langle V_0 \rangle + \langle V_{noise}(t) \rangle \rightarrow V_0$, so we'll get the answer we seek if we wait long enough. Of course, if $\langle V_{noise} \rangle \neq 0$, then we have systematic errors in our experiment in addition to random errors, and no amount of averaging will allow us to extract V_0 without first removing or otherwise measuring the systematic errors.




Figure 1. White noise versus $1/f$ noise.

The averaging procedure described in the preceding paragraph works great if the noise we're faced with is pure **white noise**, sometimes also called **Gaussian noise**. In this type of noise V_{noise} fluctuates randomly and rapidly with time, the statistics of the noise are independent of time, and each value of $V_{noise}(t)$ is completely uncorrelated with $V_{noise}(t')$, provided $|t - t'|$ is greater than some correlation time τ . A picture of white noise as a function of time is shown in Figure 1. If you look at the power spectrum of this kind of $V_{noise}(t)$ (where the power spectrum is equal to $|\tilde{V}_{noise}(\omega)|^2$, the absolute square of the Fourier transform of $V_{noise}(t)$), you will find that the power

is independent of frequency up to frequencies $f = \omega/2\pi \approx \tau^{-1}$. At higher frequencies the power typically goes to zero.

In the real world, however, noise is seldom white. Amplifiers and other noise-generating elements often drift slowly with time. Over short periods the output of a typical amplifier will drift a small amount, and over long times it will drift a larger amount. Some of this comes from environmental effects – temperature drifts in the lab, creep in materials, etc., but some is intrinsic to the electronics or other hardware. Detailed studies of many different types of electronic noise have shown that the frequency spectrum of these kinds of noise sources often goes approximately like $1/f$, where f is frequency, and hence the noise is called “ $1/f$ ” noise. Many electronic devices exhibit $1/f$ noise out to very low frequencies (even to periods of a month). After much study there is still considerable debate over just what causes $1/f$ noise in many devices, but it certainly exists and is almost ubiquitous in amplifiers and most pieces of analog electronics. Figure 1 shows some $1/f$ noise in comparison with white noise.

In addition to white noise and $1/f$ noise, most detectors also include some unwanted DC offsets to go along with whatever signal one would like to measure. Including these terms, we see that a somewhat more realistic picture for our signal voltage might be

$$V(t) = V_0 + V_{offset} + V_{white-noise}(t) + V_{1/f-noise}(t)$$

Even if we could remove V_{offset} , at low enough frequencies the noise will still be dominated by $1/f$ noise, and then a simple time average is not going to work well. A long average will nicely beat down the high-frequency noise contributions, but the longer we average the more the low-frequency drifts inherent in $1/f$ noise will contaminate the signal. Once we’re dominated by $1/f$ noise or voltage offsets, signal averaging will not be an effective way to improve our determination of V_0 .

What can often save us in this situation is that *we control the signal*, and with that control we can apply some tricks. A particularly powerful trick is to chop the signal on and off at some high frequency and take the difference $V_{on} - V_{off}$. It’s easy to see that low-frequency drifts disappear from this difference voltage, and we’re just left with the higher-frequency noise contributions. If we then average the difference signal over time we’ll be left with *only* a noise contribution at the chopping frequency. Chopping the signal is a way to completely eliminate all low-frequency noise, so we can then average over long times again to beat down the noise.

If our signal voltage is being digitized on the fly, then it’s a simple matter to perform the difference-and-average procedure in software; one just has to include a flag to keep track of when the signal is on and when it’s off. But another route is to use a *lock-in amplifier*, which is nothing more than a box of electronics that does the task for you. A lock-in is a general purpose, all-around piece of laboratory equipment, that can quickly be adapted to different experiments. As such it’s like a fancy voltmeter, useful for lots of different things.

The basic idea of how a lock-in amplifier works is shown in Figure 2. The first stage of a lock-in is where the input signal is amplified and filtered to remove noise above and below the reference frequency (since the real signal is all at the chopping frequency and maybe a few higher harmonics if the chopping produces a square wave signal). One of the things you get for your \$5,000.00 when you buy a lock-in is a robust, well-behaved, low-noise preamplifier and filter to feed your signal into. Sometimes the signal filtering is done with analog electronic filters, sometimes it’s done digitally; it

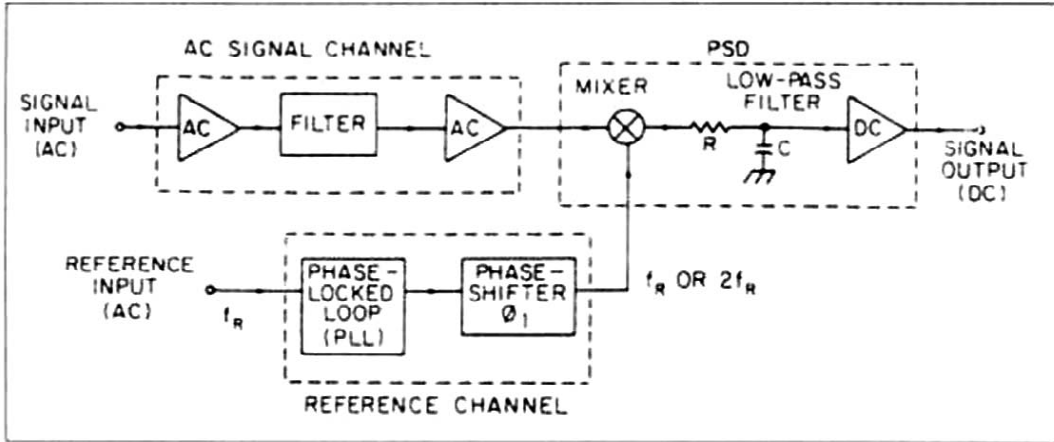


Figure 2. The inner workings of a typical lock-in amplifier.

depends on the particular lock-in, but is usually transparent to the user.

In addition to the signal input, one also needs to provide a reference input, which contains some signal with a strong frequency component at the frequency you’re modulating your signal. The lock-in electronics then “locks” onto this reference signal, which then sets the frequency. A good lock-in can often pull a decent reference out of a pretty weak reference input, but more typically one feeds the lock-in a hulking big several volt square wave which provides an absolutely unmistakable reference. Note that the reference signal contains not only a frequency but also a phase, and the latter is often relevant to one’s experiment, as we’ll see below. (The reference amplitude is irrelevant, of course.)

Many lock-ins provide a feature where the reference frequency is doubled before further processing. This is useful if you modulate your system at a frequency f , and a signal comes out at $2f$ (for example, a light bulb – 60 Hz goes in, and the light intensity coming out is modulated at 120 Hz). Presumably you could provide a $2f$ reference signal yourself, but it’s a common enough thing that some manufacturers give you a button to switch between f and $2f$.

With a locked reference and a pre-processed input signal, the lock-in then combines the two signals in the *mixer*. Mixers are nonlinear devices that come in various forms, some analog and some digital, but again the details are not usually important to the user. In essence the reference signal is converted into a sine wave with some adjustable phase, and this is multiplied by the signal input and averaged to form

$$V_{out} = \langle V_{sig} \cos(\omega t + \varphi) \rangle \quad (1)$$

where the average is a running time average. This procedure gives about the same result as the difference-and-average procedure described above, provided the phase angle is chosen properly, and you should convince yourself that the result is that $V_{out} = \alpha V_0$, where α is some proportionality constant and V_0 is the signal we’re trying to measure. (The exact prescription for mixing varies among different lock-ins, and although the proportionality constant α is always of order unity, it will sometimes include additional factors of $\sqrt{2}$ or π or some such thing. You can find all the details in

the lock-in manual, or you can send in a known signal and just measure α directly. If the experiment is really important it would be wise to do both.)

A *dual-phase lock-in*, like the one you will be using in the lab, produces two outputs, the *in-phase* and *quadrature* outputs, given by

$$V_X = \langle V_{sig} \cos(\omega t + \varphi) \rangle$$

$$V_Y = \langle V_{sig} \sin(\omega t + \varphi) \rangle$$

where again φ is a parameter you set from the lock-in front panel. The lock-in can also convert these signals digitally to amplitude and phase,

$$V_R = (V_X^2 + V_Y^2)^{1/2}$$

$$\Phi = \tan^{-1}(V_Y/V_X)$$

Both representations are useful, of course, depending on what kind of signal is being examined.




Figure 3. A typical application of the lock-in amplifier – measuring fluorescence from a sample illuminated by a laser.

To see what a flexible instrument the lock-in amplifier is, consider the experiment shown in Figure 3. Here the goal is to observe fluorescence from a sample illuminated by a laser. The laser beam is chopped using a mechanical chopping wheel, which in turns chops the fluorescence on and off. The first thing you gain by using a lock-in in this experiment is simply flexibility. The lock-in has a low-noise input amplifier, and the sensitivity can be adjusted over many orders of magnitude. Thus just about any input signal can be stuffed into the box – this is great when you start changing samples, or the laser parameters, etc.

The second thing you gain is freedom from worrying about ambient light getting into your detector. With the chopper working at 1 kHz or so, the lock-in can completely reject the DC

and 60/120-Hz signals coming from ambient lights, provided these sources don't swamp either the detector or the lock-in (both will have limited dynamic range). You could always carefully shield your detector from ambient light, of course, but in practice it's usually much less effort to use a lock-in and do less shielding. Finally, problems with voltage offsets and $1/f$ noise in the detector are also taken care of using a lock-in.

One thing the lock-in won't fix in this experiment is the problem of scattered laser light. Any stray laser light hitting the detector will look just like the fluorescence signal, and this will cause systematic errors in the measurement. If the laser beam must be bright, while the fluorescence itself is weak, then scattered light will be a real problem in this experiment, with or without the lock-in. If the fluorescence signal changes with laser frequency, then another trick you might use in this situation is to chop the laser frequency instead of the intensity. In this case the lock-in would also mostly get rid of scattered light problems, since the scattered light probably doesn't change much with laser frequency.

Noise Spectral Density. For any physically meaningful noise we can define a new “smoothed” noise function $V_{noise,\tau}(t)$ that is a running time average of the noise

$$V_{noise,\tau}(t) = \frac{1}{\tau} \int_{t' = t}^{t+\tau} V_{noise}(t') dt'$$

For the equivalent of white noise we then have that $\langle V_{noise,\tau}(t) \rangle = 0$, and $\langle V_{noise,\tau}(t)^2 \rangle^{1/2} \equiv \sigma_{noise,\tau}$ is some constant (measured in volts). Physically this is a reasonable definition because we never actually measure the noise voltage at an instantaneous time, but rather we're always averaging over some short time τ . What we see plotted in Figure 1 is not $V_{noise}(t)$, but rather $V_{noise,\tau}(t)$. Naturally $V_{noise,\tau}(t)$ depends on our averaging time τ ; the longer the average, the smaller $\sigma_{noise,\tau}$ and $V_{noise,\tau}(t)$ will be, and for white noise $\sigma_{noise,\tau} \sim \tau^{-1/2}$. We see that the averaging time τ is effectively equivalent to the noise correlation time mentioned above.

In the same way that measuring noise in the time domain always involves some averaging time τ , measuring the noise power spectrum always involves an average over a range of frequencies, called the measurement *bandwidth* (typically stated in Hz). If we compute the noise power spectrum using measurements over a finite time T_{ave} , then for white noise we find

$$\begin{aligned} \tilde{P}_{noise,B}(f) &= \left| \frac{1}{T_{ave}} \int V_{noise}(t) e^{i2\pi ft} dt \right|^2 \\ \langle \tilde{P}_{noise,B}(f) \rangle &= \frac{1}{T_{ave}^2} \left| \sum_{j=1}^{j=N=T_{ave}/\tau} \sigma_{noise,\tau}(t_j) e^{i2\pi ft} \right|^2 \\ &= \frac{1}{T_{ave}^2} N \sigma_{noise,\tau}^2 \tau^2 \\ &= \frac{\sigma_{noise,\tau}^2 \tau}{T_{ave}} \\ &= \sigma_{noise,\tau}^2 \tau B \end{aligned} \tag{2}$$

where $B = 1/T_{ave}$ is the measurement bandwidth. The sum was evaluated knowing that the noise is uncorrelated over times greater than τ . Put another way, $V_{noise}(t)$ exhibits no long-range correlations, hence on long timescales the integral random-walks and has a mean-squared value

proportional to T_{ave} .

Since $\sigma_{noise,\tau} \sim \tau^{-1/2}$ we see that $\langle \tilde{P}_{noise,B}(f) \rangle$ is equal to some constant times the bandwidth B . We therefore define the bandwidth-independent *power spectral density*

$$\begin{aligned} S(f) &= \lim_{T_{ave} \rightarrow \infty} \frac{1}{T_{ave}} \left| \int V_{noise}(t) e^{i2\pi ft} dt \right|^2 \\ &\approx \tilde{P}_{noise,B}(f)/B \end{aligned}$$

We see that $S(f)$ is a well-defined function that depends on the intrinsic noise in the system as a function of frequency. $S(f)^{1/2}$ has the dimensions of volts/ $\sqrt{\text{Hz}}$ (called “volts per root Hertz”), and this function is typically all you need to know about the random noise in a signal. For pure white noise $S(f)^{1/2}$ is equal to a constant, while for $1/f$ noise $S(f)^{1/2} \sim 1/f$.

By comparing Eqn.1 and Eqn.2 we see that $\tilde{P}_{noise,B}(f)^{1/2}$ is precisely what the lock-in amplifier measures in the absence of any signal. Thus the fluctuations in the lock-in output V_{out} will be proportional to the power spectral density at the reference frequency, $S(f)^{1/2}$, times $B^{1/2}$, the square root of the bandwidth of the measurement. Note that $B = \beta/T_{filter}$ is the equivalent noise bandwidth when using a lock-in, where T_{filter} is the averaging time indicated by the lock-in, and β depends on how the signal averaging is performed (which is described in detail in the manual). Thus we see that the noise in a lock-in measurement will go like $T_{filter}^{-1/2}$, as we would expect.

One source of measurement noise in our measurement is the input amplifiers inside the lock-in, which the manufacturer has worked hard to reduce. For example, our lock-in manual states the input noise is no more than 6 nV/ $\sqrt{\text{Hz}}$ at 1 kHz. Thus if our reference is at 1 kHz and we short the lock-in input with some small resistor (so the input signal is zero), then the signal we measure would have an effective noise of 6 nV/ $\sqrt{\text{Hz}}$. Thus if we integrate for 10 seconds, the output display should display RMS fluctuations of about $6/\sqrt{10}$ nV ≈ 2 nV.

Thermal Noise. Resistors can introduce noise into a measurement that comes from thermal fluctuations. This is called resistor *thermal noise*, or *Johnson noise*, and the RMS amplitude of the thermal noise voltage is given by

$$V_{RMS, thermal} = \sqrt{4kTRB} \quad (\text{volts})$$

or in terms of spectral density

$$\begin{aligned} V_{RMS, thermal}(\text{spectral density}) &= \sqrt{4kTR} \quad (\text{volts}/\sqrt{\text{Hz}}) \\ &= 0.128 \sqrt{\frac{R}{1 \Omega}} \quad \text{nV}/\sqrt{\text{Hz}} \end{aligned}$$

where R is the resistance value and B is the bandwidth of measurement.

II. LABORATORY EXERCISES.

To get a feeling for how the lock-in amplifier works in the lab we will first create what might be called an artificial experiment, namely measuring the resistance of a small length of brass wire. If this were a real experiment you can imagine that the wire would actually be 20 nanometers in diameter and 100 nanometers long, chilled to 50 milliKelvin at the bottom of some expensive cryostat. Of course, a wire like that is terribly easy to burn up, and maybe some weeks of effort to repair, so you can see why we would like to begin with something a bit more robust. Our emphasis will be on trying to make an accurate measurement using the lock-in, and looking at some of the

random and systematic errors that limit what we can do.

Problem 1. Given that our wire is made of brass (some random alloy) and has a length of 17 cm and a diameter of 0.4 mm, estimate it's resistance (Hint: check the web for relevant data).

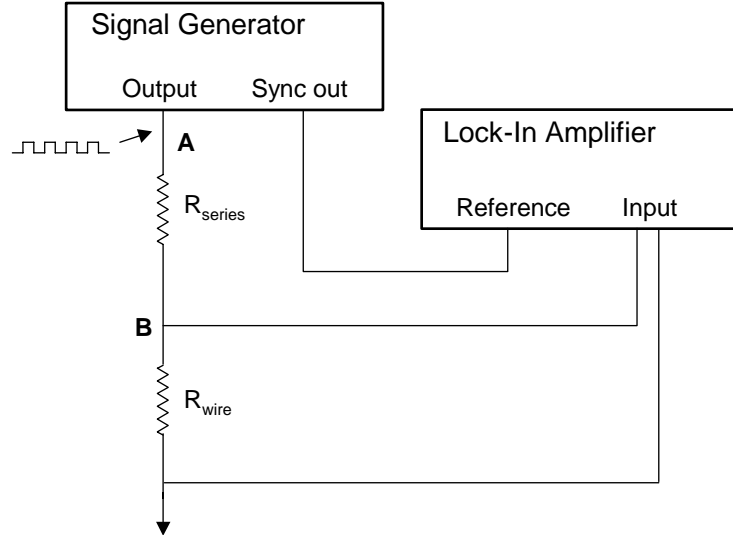


Figure 4. Layout of hardware for measuring the resistance of a short length of brass wire. Note the lock-in input is a single coaxial cable.

There are essentially two ways to measure the electrical resistance of any device – you can either send a known current through it and measure the resulting voltage, or you can apply a known voltage and measure the resulting current. In most cases, including our case, the former option will be technically easier. Thus we will use the basic resistor divider circuit shown in Figure 4, with $R_{series} = 1 \text{ k}\Omega$. The current in this case is $I = V_A / (R_{series} + R_{wire})$, and the voltages we need to measure are V_A and V_B . The wire resistance can then be computed as

$$R_{wire} = \frac{V_B}{V_A} R_{series} \left(1 + \frac{R_{wire}}{R_{series}} \right)$$

$$\approx \frac{V_B}{V_A} R_{series}$$

where in our case the approximation is accurate to a part in 10^4 . We see that our measurement of R_{wire} can be no more accurate than our knowledge of R_{series} , but since R_{series} is approximately $1 \text{ k}\Omega$ we can measure it fairly accurately with a simple digital multimeter.

Given the general truth that large voltages are easier to measure than small voltages, you can immediately conclude that the best way to measure R_{wire} is to send a lot of current through it. But for the real nanoscale experiment it's pretty obvious that large currents are a bad idea. Thus you should keep in mind that although we may send 10 milliamps through the wire today, you might be limited to much smaller currents if you're working with nanowires. (Given a maximum allowable current density of, say, 1 ampere/cm², what might you really be allowed in the nanoscale experiment?)

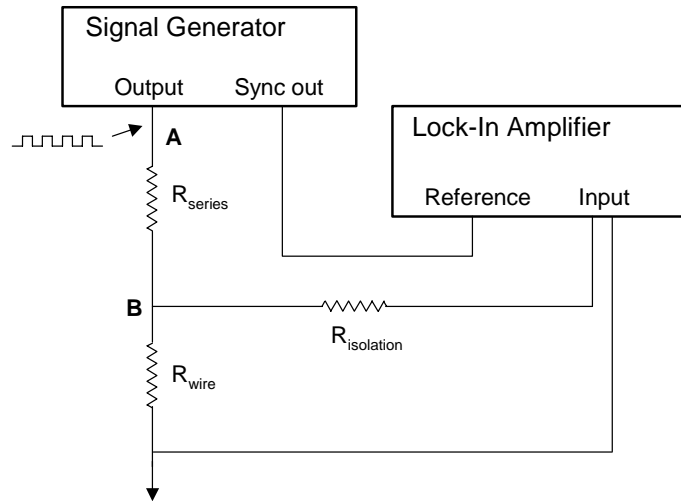


Figure 5. Another test set-up, with an isolation resistor between the wire being measured and the lock-in.

Exercise 1. Set up the circuit in Figure 4 using the small electronics box provided. Set the signal generator to a 1-volt sine wave and measure R_{wire} as a function of the signal generator frequency. Measure V_A , $V_{B,X}$ and $V_{B,R}$ using the lock-in, from which you can derive R_{wire} . (Don't just assume V_A will be 1 volt! – measure it using the lock-in.) Use a 3-second time constant on the lock-in. Note you should have $V_{B,X} = V_{B,R}$ if the circuit in Figure 4 corresponds with reality (which it does up to a point). Record measurements at frequencies of $\nu = 10, 30, 100, 330, 1000, 3000, 10000, 30000$ Hz. (Try 300 Hz when V_A is small; you will find V_B acts strangely because 300 Hz is a harmonic of 60 Hz, and there is a lot of 60-Hz noise everywhere.) Make a nice plot of R_{wire} in $m\Omega$ versus ν in Hz. Ask your TA if you don't understand what you're doing.

You should find that your measurement of R_{wire} is fine at low frequencies, but behaves strangely above a few kHz. This is because of stray capacitances in your circuit. Since the impedance of a capacitor is purely imaginary, stray capacitances will affect $V_{B,R}$ more than $V_{B,X}$. Thus your measurements of R_{wire} should be better using $V_{B,X}$ than using $V_{B,R}$. At high enough frequencies the capacitive effects strongly perturb both measurements.

Exercise 2. Set the signal generator to a 10 mV sine wave and measure the *noise* in your R_{wire} measurements as a function of frequency. Use $V_{B,X}$ and remember to measure V_A using the lock-in as well. Use a 300 msec time constant, and make ten measurements at each frequency (separated by a few seconds, which is many time constants). Compute σ_R (in $m\Omega$) from the standard deviation of the ten measurements. Use the same frequencies as before. Make a nice plot of σ_R (in $m\Omega$) versus ν (in Hz).

This time you should see the measurements are much noisier at low frequencies. This is essentially because of $1/f$ noise; more specifically the internal noise in the lock-in preamplifier is higher at low frequencies. There is a lot of detail on this point in the lock-in manual, since people using lock-ins

usually care a lot about noise.

Exercise 3. Set the signal generator frequency to 1 kHz, and measure R_{wire} as a function of V_A . Set the signal generator to a sine wave with amplitudes of $V_{A,set} = 4, 10, 30, 100, 300, 1000$ mV. Use a 3-second time constant, and derive R_{wire} from measurements of V_A and $V_{B,X}$. Make a nice plot of R_{wire} versus $V_{A,set}$. You should see some small systematic errors creeping in when V_A is small.

The moral of this story is that stray couplings and stray signals happen, and you won't always know they're there. You typically wouldn't have predicted them, and they often don't make much sense, at least in the beginning. This is really why measuring small signals is hard. Of course these things are particularly nasty when you're trying to do something no one has ever done before. A major weapon you have against systematic errors is to measure what you're trying to measure in different ways. In this experiment you begin to notice problems when you change the lock-in frequency and/or the signal amplitude. If you just measured R_{wire} at a single V_A and ν , you wouldn't have any real impression of the systematic errors.

Exercise 4. To see a different kind of random noise, set up the circuit in Figure 5 by putting a resistor between the small electronics box and the lock-in. Measure the noise in V_B for different $R_{isolation}$, and compare with what you expect from Johnson noise in $R_{isolation}$.