

Phase sensitive detection: the lock-in amplifier

by Dr. G. Bradley Armen
Department of Physics and Astronomy
401 Nielsen Physics Building
The University of Tennessee
Knoxville, Tennessee 37996-1200

*Copyright © April 2008 by George Bradley Armen**

**All rights are reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy, recording, or any information storage or retrieval system, without permission in writing from the author.*

A. Introduction

Phase-sensitive detection is a powerful method for seeing very small signals in the presence of overwhelming noise. Developed in the 1960's, it's become a ubiquitous experimental technique, and the lock-in amplifier¹ is the instrument which makes this method possible.

In this laboratory we'll learn about the method in some generality, and apply it to measure some very small quantities which would be impossible by conventional means. The laboratory consists (like life) of three stages: First we'll look at a synthetic signal in the presence of some well behaved 'noise'. This will give us some insight into the lock-in's behavior in the real world. Next, we'll perform the classic 'light-bulb experiment' which gives a lasting impression of the power of the technique. We'll actually be able to measure the intensity of a flash-light bulb placed at the far end of the 3rd floor hallway, in the presence of all the lighting from the overhead and Coke machines – with a very poor detector². Finally, we'll apply the method to the more serious problem of measuring the

¹ If *Wikipedia* is to be believed: The lock-in amplifier was invented by physicist Robert Dicke of Princeton University, who founded the company Princeton Applied Research (PAR) to market the product.

² Our eyes can do it, but they're very, very good detectors!

Faraday effect. Here, polarized light passing through a dielectric material in the presence of a weak magnetic field rotates slightly. This angle of polarization rotation is very small ($\approx 0.008^\circ$) but we'll be able to measure this easily — and accurately.

The actual purpose of this lab is to leave you with some working knowledge of lock-in amplifiers and what they're good for. At some point in your future career you may very well be designing experiments of your own, and a vague memory of this lab and phase-sensitive capabilities may put you on a fruitful track.

B. Phase sensitive detection

A lock-in, or phase-sensitive, amplifier is simply a fancy AC voltmeter. Along with the input, one supplies it with a periodic reference signal. The amplifier then responds only to the portion of the input signal that occurs at the reference frequency with a fixed phase relationship. By designing experiments that exploit this feature, it's possible to measure quantities that would otherwise be overwhelmed by noise.

The lock-in amplifier operates on a very simple scheme: Consider a sinusoidal input signal

$$V(t) = V_0 \sin(\omega t + \phi).$$

Suppose we also have available a reference signal

$$V_R(t) = \sin(\Omega t).$$

The product of these two gives beats at the sum and difference frequencies

$$V(t)V_R(t) = \frac{V_0}{2} \{ \cos[(\omega - \Omega)t + \phi] - \cos[(\omega + \Omega)t + \phi] \}.$$

When the input signal has a frequency different from the reference frequency Ω , the product oscillates in time with an average value of zero. However, if $\omega = \Omega$ we get a sinusoidal output, offset by a DC (zero frequency) level:

$$V(t)V_R(t) = \frac{V_0}{2} \{ \cos[\phi] - \cos[2\Omega t + \phi] \} \quad \omega = \Omega.$$

If we can extract the DC component of this product, *and* are able to adjust ϕ , we get a direct measure of the signal amplitude V_0 . So, we arrange to have our quantity of interest oscillate at Ω ; any unwanted signals oscillating at different frequencies are rejected. Furthermore, any random³ noise that does oscillate at Ω will also be rejected.

The following figure illustrates how a lock-in works:




Fig. 1 Block diagram of phase-sensitive detection.

The input signal $V(t)$ passes through a capacitor, blocking any pre-existing DC offset, and is then amplified⁴ (A). The reference signal $V_R(t)$ passes through an adjustable phase-shifter (ϕ). These two results are then multiplied, and any resulting DC component is extracted by the low-pass (L.P.) filter.

³ By random, we mean that the noise signal has a time-dependant phase, and upon averaging gives zero.

⁴ The amplifier may be in stages, before or after the multiplier, or both.

The idea is simple enough, but the actual implementation is difficult and Lock-ins tend to be expensive. While this isn't an electronics class, a few of the details are worth noting:

First, the actual reference signal need not be sinusoidal. Lock-ins take the reference signal, pass it through a phase shifter, and then create their own internal reference 'locked' to the phase-shifted external reference. This allows for much greater flexibility (and reliability) in operation. Also, the input need not be sinusoidal, merely periodic with frequency Ω . The lock-in then picks out the fundamental Fourier component of the input waveform.

Second, the low-pass filter is not perfect. In fact, if it were perfect, the instrument would be largely useless. We'll look at this in more detail presently.

Finally, the multiplier (or 'demodulator') is a tricky device to implement in analog form (or used to be). If one first digitizes the input and reference, then of course multiplication is simple. However, there are some disadvantages inherent to the digital approach, primarily one of dynamic range — if you digitize the input with a certain precision (bits) the ability of the instrument to extract small signals is of course limited. Thus, when shopping today one can find both digital and analog models, each having their own merits. Older models (like the one you'll be using, naturally) are analog, but additionally don't realize the exact scheme outlined above: Because the multiplication circuitry is difficult, in older lock-ins the polarity of the input signal is reversed periodically at the reference frequency. This is equivalent to multiplying $V(t)$ with a square-wave reference. The result is the same, however the problem of harmonics is introduced (see App. B).

Regardless of all these kinds of details, let's continue by considering the lock-in as a machine performing the following operation

$$V_{in} = V_0 \sin(\Omega t + \phi) + \text{"noise"} \rightarrow V_{out} \propto V_0 \cos(\phi).$$

Usually, the experimenter adjusts the phase difference $\phi = 0$ so that the signal is a maximum. Lock-ins are calibrated so that this maximum output voltage equals the RMS value of the desired signal, that is $V_{out} = V_0 / \sqrt{2}$.

Now, to apply the lock-in to real world problems we need to understand a slightly subtle, but important (and useful) point: It's seldom the case that we modulate the input *voltage* sinusoidally, but rather some physical parameter. This parameter we modulate causes the physical system we're interested in to respond with frequency Ω , and finally a detector translates this into a voltage! All is well if the system's and detector's responses are linear, however often they are not.

Let's consider this point more carefully. We have a combined system and detector, which creates a voltage which is a function of some stimulus (s) over which we have control: $V = V(s)$. The stimulus might be anything; current, magnetic field, light intensity or wavelength, etc. We then arrange the stimulus to vary sinusoidally around some average value \bar{s} at our reference frequency Ω , with amplitude A :

$$s(t) = \bar{s} + A \sin(\Omega t).$$

The detector output is then a time-varying voltage

$$V(t) = V(s(t)).$$

In general, we can't say much else. The output will still be periodic in time, but will not necessarily be a sine-wave. Our lock-in will then pick out the fundamental Fourier component of this function and report its RMS voltage. However, if we arrange it so that modulation amplitude A is 'suitably' small, then we can approximate $V(s)$ by a Taylor-series expansion about \bar{s} :

$$V(t) = V(\bar{s}) + \left. \frac{dV}{ds} \right|_{\bar{s}} A \sin(\Omega t) + O(A^2)$$

Running this through the lock-in amplifier gives an output

$$V_{\text{out}} \approx \frac{A}{\sqrt{2}} \left. \frac{dV}{ds} \right|_{\bar{s}}.$$

In words: The lock-in's output is proportional, not only to the modulation amplitude A , but also to the derivative of the system's response with respect to the stimulus, evaluated at $s = \bar{s}$. We've also assumed here that the relative phase difference has been pre-adjusted so that $\phi = 0$. If this isn't the case, we must also include the $\cos(\phi)$ factor.

Figure 2 gives a graphical visualization of the above argument. The graph shows some non-linear output function $V(s)$. Two inputs are shown at s_1 and s_2 , each modulated with amplitude A . The output depends not only on the input modulation amplitude A , but also on the slope. Naturally, if the amplitude is not sufficiently small, the outputs become distorted and the outputs are no longer is proportional to dV/ds .

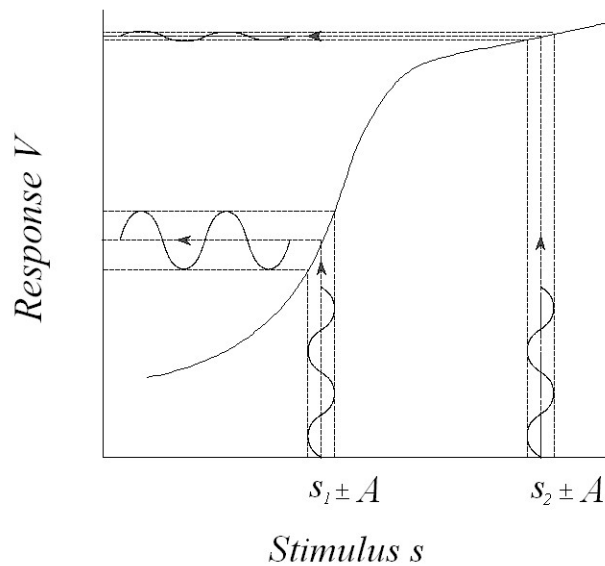


Fig. 2 Effect of a non-linear response on modulated inputs.

The lock-in measures not only the *magnitude* of the response derivative, but also the *sign*. This we've shown mathematically using the Taylor-expansion argument above, but it's also worth seeing intuitively how the lock-in does this. It's obvious that the *amplitude* of the output signal should be proportional to $|dV/ds|$, but it's the phase-sensitive aspect of

the instrument that also allows us to also determine if it's rising or falling. Figure 3 sketches the response to a modulated input for a positive and negative slope. The sign of the derivative is reflected in the *phase* of the output waveform relative to that of the input.

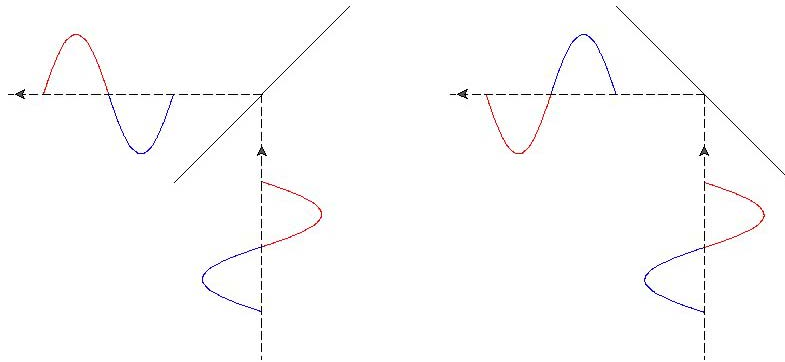


Fig. 3 Relationship between output phase and the sign of the response's slope.

The proportionality of the Lock-in output to dV/ds is actually very useful in a number of situations. An important example arises in spectroscopy where we're looking at the absorption of a sample as a function of incident-light frequency (ν) near some resonance centered at ν_0 . If we vary the frequency slightly around an average value $\bar{\nu}$ sinusoidally we can use our lock-in technique to look at the absorption as a function of $\bar{\nu}$. The following figure sketches the line and its derivative:




Fig. 4 Derivative of an absorption resonance.

The lock-in output (sometimes called a dispersion curve) is bi-polar; positive below resonance and negative above—with a zero crossing at ν_0 . There are a number of advantages to the lock-in's response. In the presence of noise or for weak signals, the dispersion curve proves more accurate in determining the peak center.

Most importantly the dispersion response is really useful for practical purposes in control systems. For frequencies near ν_0 the dispersion curve provides an error signal for regulating ν : if the output is negative you know ν is too large; if positive, it's too small. This error can be used as feedback to stabilize the source.

The following figure outlines an interesting example: stabilizing a diode laser against frequency drift:

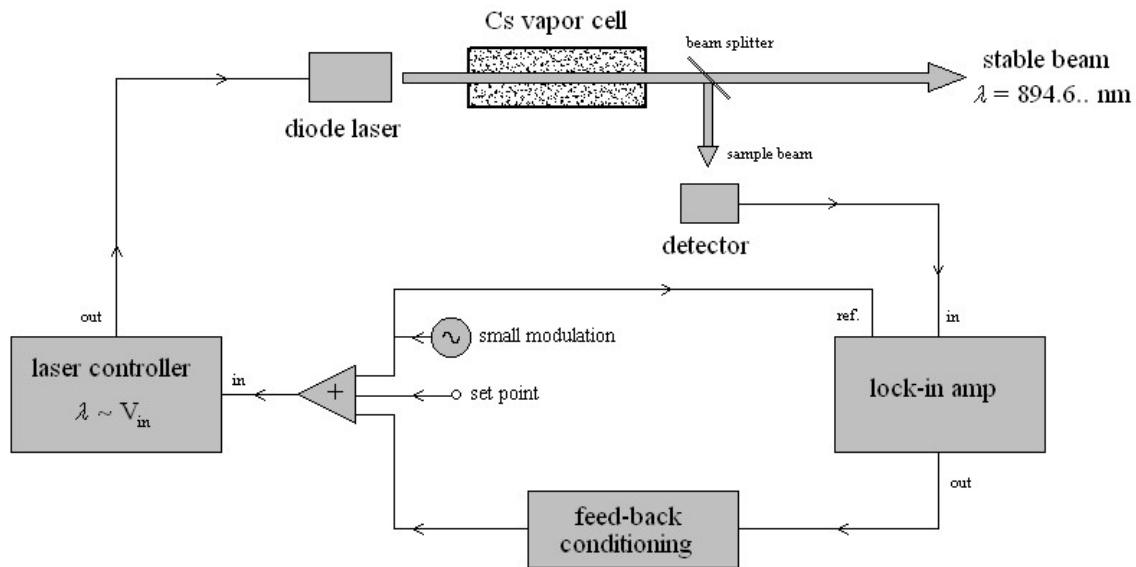


Fig. 5 Simplified feedback system.

Suppose we have a diode laser capable of producing a very monochromatic output beam in the infra-red. The output wavelength is determined by a laser controller, the details of which we won't be concerned with, except that an input voltage can be used to tune the output wavelength λ over a certain range. Unfortunately, the controller isn't perfect and the beam's wavelength will tend to drift away from a set value.

To stabilize the laser, the output beam (or a portion of it) is passed through a Cs vapor cell and the intensity monitored. If the laser wavelength is initially set to the Cs D1 resonance line, a lock-in can be used to maintain this frequency. As the laser wavelength drifts high or low, the lock-in's output provides an error signal which, when conditioned and fed back to the controller, corrects the drift.

With this somewhat lengthy digression out of the way, let's examine one last point which is important in experiment design: Lock-in detection measures the derivative of the response with respect to the stimulus s **which is modulated**. When there's more than one parameter, the choice of *which one* to modulate in a given experiment must be considered carefully.

Let's look at a simple, hopefully familiar, example: Consider determining the Hall constant (R_H) of a metal by measuring the hall voltage

$$V_{hall} = \frac{R_H}{\delta} BI + RI .$$

Here δ is the sample thickness, I the current passing through the sample, B the applied magnetic field, and R is an unknown (and unwanted) offset resistance.

Suppose we want to use our lock-in technique to find R_H . There are two parameters we can modulate, B or I . If we hold B constant and modulate the current $I = I_0 \sin \Omega t$, then measuring V_{hall} with the lock-in gives

$$V_{out} = \frac{I_0}{\sqrt{2}} \left(\frac{R_H B}{\delta} + R \right) .$$

This is a bad experimental design since our output contains the unknown offset resistance which is likely large compared with the term we want. Instead, if we hold I constant and modulate the magnetic field $B = B_0 \sin \Omega t$, we get

$$V_{out} = \frac{B_0}{\sqrt{2}} \left(\frac{R_H I}{\delta} \right) ,$$

with all quantities known except R_H . By choosing to modulate B rather than I , the derivative feature of the lock-in is employed (gainfully) to reject a large nuisance signal.

C. Getting to know the Lock-in

In this section we'll explore some basic lock-in amplifier characteristics. Our laboratory will be based on an *EG&G* model 5101 lock-in amplifier. As mentioned, this model is a bit old (and has a few problems) but is simple, having all the basic features. Figure 3 is a photo of the 5101 front panel:




Fig. 3 Our amplifier.

As you can see, the instrument controls are organized in three sections (neglecting the obvious power switch). Let's go through them:

To the left is the Signal section. It includes the signal input BNC jack and a sensitivity setting. The sensitivity setting is essentially a selectable gain in the amplifier stage (A) of Fig. 1. The sensitivity can be adjusted from 1 μV up to 250 mV full scale in overlapping steps.

The Reference section includes a BNC reference input jack and some controls that manipulate the phase shift applied to the reference signal. These include a continuously adjustable knob that varies from 0° - 90° . To get larger phase shifts, you use this knob in conjunction with the Quadrant buttons, which add an additional 90° and 180° (or both) when pushed. An LED indicator glows red when the instrument's internal circuitry

cannot lock on any recognizable signal. So, all's well if the light is *off* (unless you forgot to turn the power on—there's no power-on indicator!). Finally, the Mode button selects whether to lock onto the input reference frequency f , or at twice this frequency $2f$. This is useful when your response of interest doesn't care about the polarity of the modulation, and so occurs twice as fast. An example of such a case is the AC flicker of light bulbs: they're supplied with 60 Hz current, but since the filaments are indifferent to current direction, light is radiated with a 120 Hz flicker.

The output stage consists of the actual signal output and some other controls. The actual output can be viewed on the analog meter which displays \pm the full-scale sensitivity. Despite the prejudice of the young against anything non-digital, the meter is very useful as a quick guide to what's happening. Along with the meter output there is an analog output (BNC) which can be connected to a DMM or oscilloscope. This output is calibrated to give ± 1.00 V at full scale. So, for example, if this voltage is 0.580 V on the 250 mV range, the signal reading is $0.580 \times 250 = 145 \mu\text{V}$. An over-range light informs you when the sensitivity is set too high.

The Zero knob and offset buttons allow you to subtract an unwanted DC offset — we won't be concerned with these. The offset button should always be set to *off*.

Finally we come to the filter controls. These concern the low-pass filter (L.P. of fig. 1) operation, and we'll investigate these in some detail. The time-constant settings select how sharp the L.P. filter is, *i.e.* how much low-frequency noise is allowed to pass and corrupt our true DC level.

A thoughtful person might wonder just why we want this adjustable — don't we want it as sharp as possible to get the best reading? The answer is a matter of practicality. A Low-Pass filter actually corresponds to a time averaging of the signal (see App. A). If there is a change in the signal, you must wait for a period of about 5 time-constants for the averaging to settle to a good result. An infinitely sharp filter would imply an infinite wait. Thus, there is a trade off between how quickly you want to make a measurement, and how much low-frequency noise you're able to tolerate. Additionally, it's usually the case that you want to make multiple measurements for different values of the stimulus s : How often you vary s (or how fast you can scan s) depends on how much unwanted noise you're willing to accept.

We now need to discuss low-pass filters in a bit more detail. While the following should be familiar to you from your electronics course, it's best to review the material here, in the context of our specific instrument.

As pointed out in appendix A, a first-order low-pass filter with time constant τ has a transfer function

$$T(\omega) = \frac{|V_{out}|}{|V_{in}|} = \frac{1}{\sqrt{1 + (\omega\tau)^2}}.$$

Here, we're ignoring the phase-shift a signal encounters in traversing the filter and concentrating on the attenuation only. Notice that this attenuation depends only on the product $\omega\tau$. Figure 4 displays T as a function of $\omega\tau$.




Fig. 4 First-order low-pass filter.

As expected, signals with frequencies small compared to $1/\tau$ are passed un-attenuated. For higher frequency signals, $\omega\tau \gg 1$, the attenuation goes as $T \rightarrow 1/\omega\tau$ (which engineers like to refer to as “a 6 db per octave rolloff”). The frequency at which the

attenuation drops to $T = 1/2$ is $\omega_{1/2} = \sqrt{3}/\tau$: The larger the time constant, the sharper the filter.

Our *EG&G* lock-in employs two 1st order filters cascaded; a pre- and post filter. The pre-filter has a wide selection of time constants to choose from, ranging from ms up to 30 sec. The post filter only has two choices: 0.1 and 1 sec, or can be removed from the signal path altogether (turned ‘off’). The utility of the post filter is in suppressing noise at larger frequencies. In fact two cascaded 1st order filters make a 2nd order filter with an attenuation $T \rightarrow 1/\omega^2$ at large frequencies. Figure 4 shows a plot of a 2nd order filter’s transfer function that is made from two cascaded 1st order filters.




Fig. 4 2nd order response (solid black line) formed from two cascaded 1st order filters, one with $\tau = 0.1$ s (dashed red) and another with $\tau = 1$ s (green dot-dashed).

The thing to notice here is that the low-frequency transmission of the filter is determined almost exclusively by the stage with the largest time-constant ($\tau = 1$ s). However, the presence of the low- τ stage ($\tau = 0.1$ s) creates a significant reduction of high-frequency transmission.

The general scheme of operation is thus the following: with the post filter off, one first adjusts the pre-filter time-constant to some desired value. Then the post filter is

engaged, choosing its time constant as less than or equal to that of the pre-filter. This increases the high frequency noise rejection. The limited choice of .1 or 1 second for the post filter's τ reflects this secondary role. (Clearly, if you're using pre-filter time constants less than 0.1 s, noise isn't a problem and the post filter is unnecessary.)

Experiment 1:

Let's now experiment in the semi-real world with our lock in. We'll use a simulator box to provide a source of well-controlled signal and *coherent* noise to our lock in. Figures 6 and 7 show the box and its functional workings:

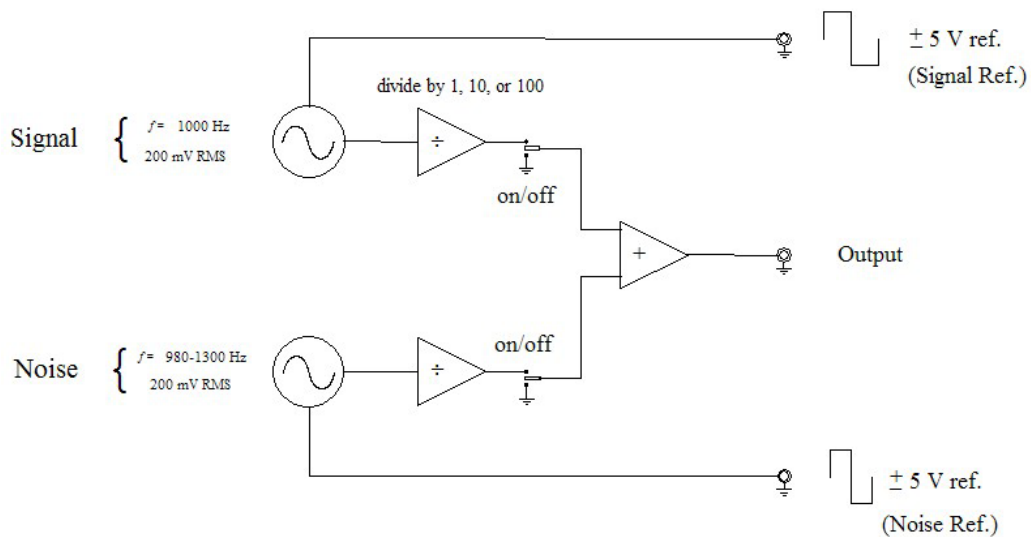


Fig. 6 Block diagram of the simulator box.



Fig. 7 The simulator box.

The output of our simulator box is the sum of a ‘signal’ and a ‘noise’ component, either of which can be turned off. The signal part consists of a 1.00 kHz sine wave whose *RMS* amplitude can be set to 200, 20, or 2 mV. A ‘Signal Reference’, in phase with the signal, is provided for the lock-in’s reference. The noise component is a sine wave of frequency variable over a range of about 0.98 to 1.30 kHz, with a similar range of amplitudes. A reference output for the noise is also provided for accurately measuring the noise frequency.

Step 1: Setup.

We’ll use the simulator box as an input to the lock in, and monitor the lock-in’s response by several means: The lock-in’s panel meter, and the scaled output on both a digital multimeter (set to read **DC** voltage) and an oscilloscope. Using BNC cables, connect the equipment as shown in Fig. 8.

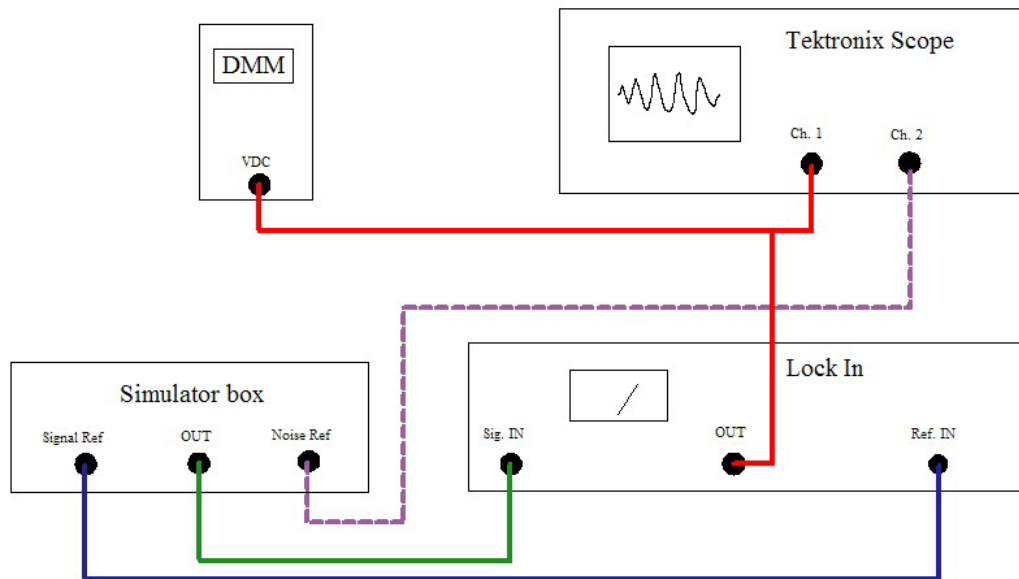


Fig. 8 Setup for experiment 1.

The purple connection from the noise reference to channel 2 of the scope is an optional connection, which can be used later for accurately measuring the noise frequency. Use Channel 1 set to DC coupling.

Step 2: Signal with no noise.

Now we'll look at the lock-in's response to only a pure sine wave at the in sync with the lock-in's reference signal. On the simulator box turn on the signal switch and turn off the noise switch (only the signal LED should be lit). Adjust the signal amplitude to '÷1', which produces a 1.00 kHz signal with an amplitude of 200 mV (RMS).

IMPORTANT NOTE: Despite the cookbook style of this write up, you're *encouraged* to be creative. For instance, at this point you could check the input by patching it directly from the simulator directly into the scope. These sorts of maneuvers are good skills to acquire.

Now, with our known input established look at lock-in's response. The first thing to play with is the phase adjustment. Because the square-wave reference signal is designed to be in phase with the sine wave, you should see a maximum (200 mV) on the *meter* when you adjust to 0°, a minimum (-200 mV) at 180°, and a null reading at 90° and 270°.

Once you're familiar with the phase adjustment, we need to be more precise with our measurements. You will note that the meter doesn't quite show exactly 200mV at maximum. There is an overall calibration error with our lock-in that will be important to us in the last experiment of this lab when absolute measurements are required. We need to establish a correction factor. To do so, use the DMM reading of the output for accuracy. Carefully adjust the phase setting for maximum output and acquire a measurement (recall that you need to multiply the DMM reading by the full-scale sensitivity), and compare this with the 200 mV (RMS). Repeat this using different inputs amplitudes (20.0 mV and 2.00 mV), and using as many possible sensitivity settings as possible. Repeat this procedure with the phase shifted 180°. Hopefully you will find a standard correction factor, independent of the sensitivity. When I did this my average correction factor was

$$V_{TRUE} = 1.12 \times V_{LOCKIN} .$$

Step 3: Noise with no signal.

Now, turn off the signal and turn on the noise. Switch the noise amplitude to the ‘÷1’ scale (200 mV RMS). Noise the noise frequency adjustment clockwise a ways. Now we should have an input whose frequency is higher than the reference frequency. We thus expect the lock-in’s output to be an attenuated sine wave oscillating at the difference frequency. Use the oscilloscope to look at the lock-in’s output. For frequencies too close to the 1kHz reference the scope output is somewhat useless since the difference frequency $\Delta f = |\omega - \Omega|/2\pi$ is so low; however in this situation the panel meter will show you the oscillating behavior. At larger frequencies, the panel meter can’t respond but the scope will let you see the behavior. Play around: What you should see is that as the frequency is increased (1) the output oscillates faster and (2) the amplitude of oscillation decreases.

Now, let’s get a bit more analytical. *Make sure the post filter is OFF.* Our digital Tektronics scopes have a measurement feature: set it to measure the RMS amplitude and frequency of the lock-in’s output (if you don’t know how, get someone to help). Now select a lock-in time constant τ (*i.e.* 0.1 sec). With τ fixed vary the noise frequency and, using the scope, record the amplitude and difference frequency⁵ for a number of frequencies over as wide a range as you can. (You’ll have to keep changing the sweep and range settings). Repeat this for other time constants.

With this data, convince yourself that the lock-in is suppressing the noise in the fashion of a 1st order filter. One way to do this is to note that only the product $x = \tau \Delta f$ enters into the attenuation, and that

$$\left(\frac{1}{V_{out}}\right)^2 = \left(\frac{1}{V_{in}}\right)^2 + \left(\frac{1}{V_{in}}\right)^2 \frac{x^2}{(2\pi)^2}.$$

⁵ And alternate method is to use channel 2, measure the noise reference frequency and subtract 1 kHz. You should at least do this once and compare, so that you really believe we’re seeing the difference frequency.

So, plotting $(1/V_{out})^2$ verses x^2 should give a straight line. Furthermore, performing a linear fit should give you a slope and intercept consistent with our $V_{in} = 200$ mV RMS, corrected for our calibration error.

Next, let's see about the post filter. Select a time constant and noise frequency and record the output amplitude and frequency. Now turn ON the post filter. Record the output amplitude for both the $\tau_p = 0.1$ and 1 sec post-filter settings. Using the post filter should further reduce the amplitude by a factor of $[1 + (2\pi\tau_p\Delta f)^2]^{1/2}$.

Step 3: Signal and noise together.

Finally, turn both the noise and signal channels on. Play around with the settings. Of particular importance is to see the behavior of noise at a frequency very close to the reference. Of course, what you have is the oscillating noise superimposed atop of the DC signal level. In practice it sometimes happens that the beats are so slow you see the meter drift about an average position. If the drift is too slow, you don't notice the oscillation and take a reading in error.

D. Chopping and the light-bulb experiment.

The power of lock-in detection can best be appreciated by considering the simple experiment outlined in figure 3. We'll be performing this as part of the laboratory.

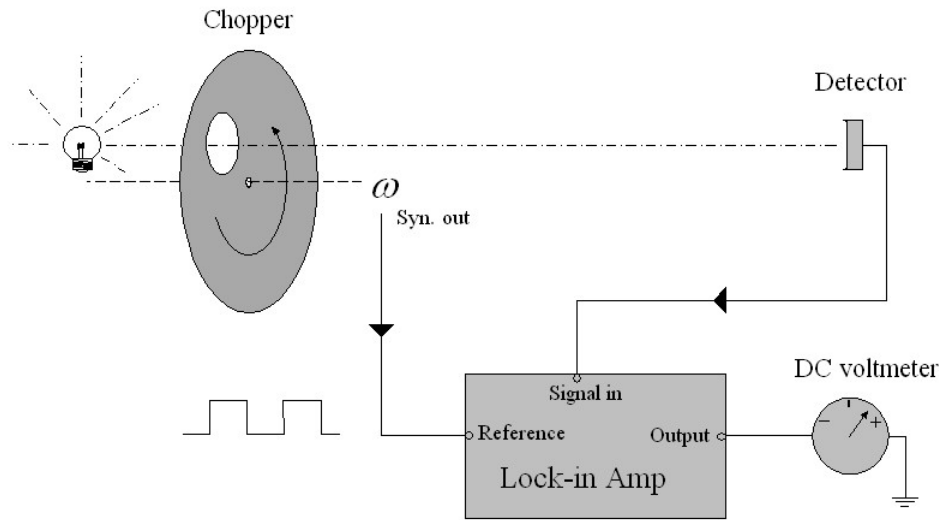


Fig. 9 The Light Bulb experiment.

Here, the light from a small flash-light bulb is chopped by a rotating mask. It's very common in optical experiments to achieve intensity modulation in this way, and optical choppers of high quality are available commercially. Because the signal is periodic, the lock-in detects the signal at the fundamental frequency. Some signal power is lost into higher harmonics of ω as a consequence, but it's a small price to pay for the ease of modulation.

The point of the experiment is to see the amazing ability of lock-in detection to measure the small signal of the bulb in the overwhelming presence of hall lights, coke machines, etc—even at very large distances.

For the moment, we'll assume the voltage from the detector is proportional (constant C) to the light intensity. We might suspect that

$$V_{Det} = \frac{CI_{Bulb}}{r^2},$$

where r is the distance between the bulb and the detector, and I_{Bulb} is the intrinsic light output of the bulb. Since we modulate I_{Bulb} , our lock-in output is

$$V_{out} = \text{some constant} / r^2.$$

In our experiment we will try to verify the inverse-square law. To do so, we will measure V_{out} at a variety of distances down the hallway. The slope of a plot of $\log V_{out}$ versus $\log r$ will give the power law, hopefully somewhere near 2.

DISCLAIMER: Before we get started, let's emphasize that the r^{-2} law would not be valid even in a perfect hallway (unless the walls are painted black). Reflections from the floor, walls, and ceiling that happen to reenter the detector cause a discrepancy that disappears only at fairly large distance (depending on the average reflectivity). To further confound the problem coke machines, doorways, overhead lights, *etc.* all confound the issue by masking reflections at certain points. Look at this example:




Fig. 10 Short-range data.

This plot shows some data I took close to the source: we might expect $r^2 \times V_{out}$ to be constant. The gradual rise is just what would be expected by considering reflections, and the dip occurs where the first soda-pop machine starts to occult these reflections from one wall! So, don't expect perfect results, the actual point is to become familiar with measuring progressively weaker signals as you move down the hall.

Procedure.




Fig. 11: Detector and chopper components.

- 1.) Noise: hook detector to scope and look at DC background plus 120 Hz ripple.
 - 2.) Phase: detector **near** chopped bulb: adjust phase for max – then leave it.
 - 3.) chopper freq: put detector **far** from bulb:
 - Bulb off: try and adjust chopping freq to get min signal (see appendix C)
 - Bulb on and off: can you see it?
 - 4.) take readings down the hall VDC and Vlockin – practice with sensitivity and time constant settings. Beware low freq noise beats: Use a max/min recording DMM, average gives signal and spread gives error.
 - 5.) take good readings near end of hall – more data/distance since this is where we might hope to get a real $1/r^2$ power law.
- Fiddle with chopper speed, hooked to detector minimize noise. Remember phase.

Analysis

- 1.) Plot data, log-log plot to get power law. Perhaps just long-range points.
- 2.) More impressive is to consider signal/noise you are able to extract vs. distance.
- 3.) extra credit: read appendix B and apply corrections for non-linear detector. Make any difference?

E. Measuring the Faraday effect.

As a more serious example of lock-in experimental design we consider the Faraday effect: When a beam of polarized light traverses a material in the direction of an external magnetic field, the angle of polarization is found to rotate slightly. The angle of rotation θ is found to be proportional to the strength of the magnetic field, and to the distance traveled through the sample. The constant of proportionality (\mathcal{V}) is called the Verdat coefficient. We have

$$\theta = \mathcal{V}Bd .$$

While the Verdat coefficient is a property of a specific material (App. 2), in practice is pretty much the same for most substances: $\mathcal{V} \approx$ for solids and liquids. Consider then the rotation of light in our big lab magnet where we can attain $B = 12kG$. The maximum sample length would be the pole-face separation $d = 4.3$ cm. (Perhaps we steer the light in and out with fiber optics?) Under these (extreme) conditions, the light's polarization would rotate by about .

To measure the Verdat coefficient we employ the experimental setup shown in Fig. 4: Our sample is placed in the center of a Helmholtz coil pair (not shown) so that a magnetic field lies along the sample axis. The magnitude of \mathbf{B} is modulated by driving the coils with an AC current, so that $B(t) = B_0 \sin \omega t$. A linearly-polarized laser beam is passed through the sample axis, then through a sheet polarizer, and detected with a photodiode. The Faraday effect rotates the angle of polarization by θ as the beam passes through the sample. For the moment, let's consider the polarizer as rotated by an unspecified angle Ω with respect to the incident polarization direction as shown.

To understand how we can use lock-in detection to measure \mathcal{V} , one needs to know a only small amount of circuit theory. The photodiode works like an ordinary diode except that when a photon strikes it's junction a certain number of electron-hole pairs are produced. These pairs are drawn out of the junction when the device is biased. When connected in reverse bias as shown, only a tiny (dark) current flows in the absence of light. When light of intensity \mathcal{J} is absorbed a photocurrent flows proportional to the

light intensity; we can write this as $i = \beta \mathcal{J}$. This voltage must flow through the resistor, and so a voltage $v = \beta \mathcal{J} R$ develops across R . (We must be careful to measure v with an instrument whose impedance is very much larger than R . In our detector $V_{cc} = 9\text{ V}$ and $R = XXX\ \Omega$.)

Next, consider the light intensity \mathcal{J} striking the photodiode. Supposing that the laser intensity, once it's traversed the sample, is \mathcal{J}_0 Malus's law for the polarizer gives

$$\mathcal{J} = \mathcal{J}_0 \cos^2(\Omega - \theta).$$

Since θ is small, we can expand this in a Taylor series about $\theta = 0$ yielding

$$\mathcal{J} \approx \mathcal{J}_0 (\cos^2(\Omega) - 2 \cos(\Omega) \sin(\Omega) \theta).$$

Now, since B is modulated so is the rotation angle

$$\theta = VdB_0 \sin \omega t \equiv \theta_0 \sin \omega t.$$

The voltage across the resistor thus has an AC and DC component:

$$v = v_{DC} + v_{AC},$$

where

$$v_{DC} = R\beta\mathcal{J}_0 \cos^2(\Omega)$$

and

$$v_{AC} = -2R\beta\mathcal{J}_0\theta_0 \cos(\Omega) \sin(\Omega) \sin \omega t.$$

Now, v_{AC} is what we will measure using lock-in detection with a $\sin \omega t$ reference. To make this quantity as large as possible, we see that we must set $\Omega = 45^\circ$. This is in keeping with our earlier discussion: when modulating θ the lock-in scheme measures $d\mathcal{J} / d\theta|_{\theta=0}$, and we can maximize this signal by adjusting our free parameter Ω .

So, setting the polarizer's pass axis at 45° to the incident polarization, measuring v_{AC} with the lock-in and v_{DC} with a DMM we get

$$v_{DC} = \frac{1}{2} R\beta\mathcal{J}_0 \quad \text{and} \quad v_{lockin} = \frac{1}{\sqrt{2}} R\beta\mathcal{J}_0\theta_0$$

(where the $\sqrt{2}$ comes from the RMS calibration of the lock in). Taking the ratio

$$\frac{v_{lockin}}{v_{DC}} = \sqrt{2} \theta_0$$

eliminates the unknown quantity $\beta\mathcal{J}_0$. Solving for the Verdat coefficient finally yields

$$V = \frac{1}{\sqrt{2} B_0 d} \frac{v_{lockin}}{v_{DC}} = \frac{1}{2 B_{RMS} d} \frac{v_{lockin}}{v_{DC}}.$$

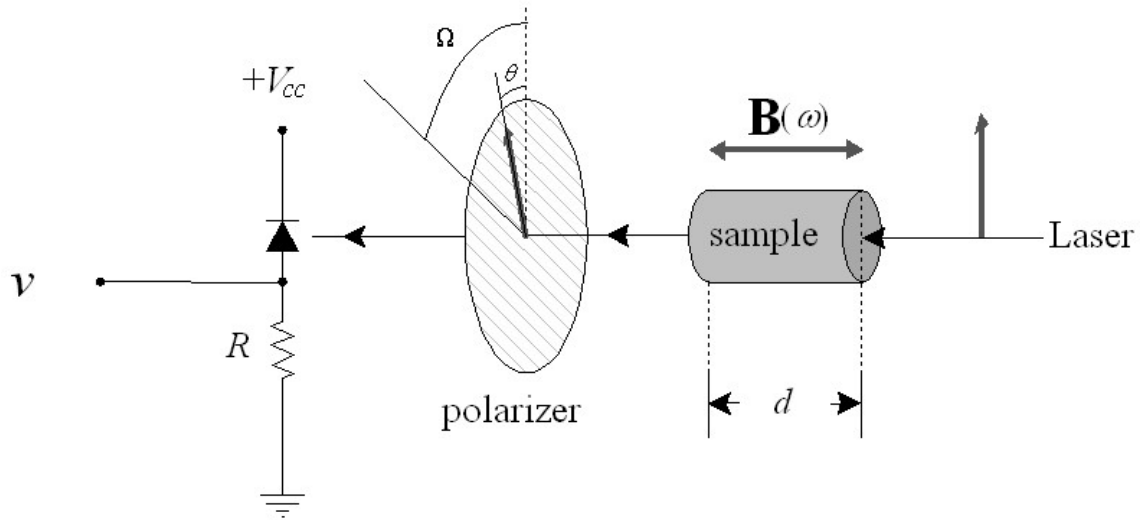


Fig. 4 Schematic experiment for observing Faraday rotation.

Procedure:

Monitor output: V_{rms} (measured by scope) gives measure of B_{rms} of helmholtz coil; using Bell gauss meter in AC mode I measured:

$$B_{RMS} = (0.00780 \text{ G/mV}) \times V_{RMS} .$$

1. Polarization setup: Use DMM and DC output of detector. Adjust θ to find max and min output -- don't saturate the detector, output should be less than 5V. Adjust R in needed. (you can use scale on polarizer to mark 0 deg.) Set polarizer to 45 deg.
2. Put in sample (water and ethanol in culture bottles). Turn on coils and monitor. Measure detector output with lockin referenced to helmholtz syn. out. Take measurements of VDC (DMM) and Vlockin for various settings: Vary B_{RMS} and modulating frequency, both can be measured by monitor output with scope.
3. Calculate angle and Verdet as function of B_0 . Compare with measured values at 640nm.

APPENDIX A: Low-pass filters in the time domain.

We're quite used to the idea of circuit behavior in the frequency domain, and circuit analysis is greatly simplified by Fourier-transforming the time-dependent differential equations into algebraic equations of frequency. This is the method of complex impedances. It's interesting to investigate the behavior of a low pass filter in the time domain. From such a study we can see the connection between frequency filtering and time averaging. Consider the following block diagram:




Fig. A1 Low pass filter showing a representative frequency response.

The *L.P.* box transforms an input signal f into some output function g . Here we can consider f and g as functions of either time or frequency, their connection being the Fourier transform, *e.g.*

$$f(t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(\omega) e^{-i\omega t} d\omega \quad \text{and} \quad f(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(t) e^{+i\omega t} dt ,$$

with similar relations connecting $g(t)$ to $g(\omega)$ and $L(t)$ to $L(\omega)$. The frequency-dependent transfer function $L(\omega)$ connects the input and output. A typical low-pass dependence on ω is sketched in the Fig. A1, the filter passes low frequency components

and blocks those at higher frequencies. We might define some characteristic frequency ω_c (as shown) to roughly define the cutoff region.

The frequency-dependent relation between input and output is simple:

$$g(\omega) = L(\omega)f(\omega).$$

This relation can be transformed into one in time by taking the inverse Fourier transform.

The result is the convolution theorem

$$g(t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} L(\tau)f(t-\tau)d\tau.$$

The simple, algebraic relation in the frequency domain produces a *non-local* relation in time; that is the output g at time t depends on the input f , not only at t , but at all other times as well. For physical systems we expect causal behavior: the output should not depend on the input at future times. Indeed, the differential equations which describe the time evolution of the system are inherently causal. Thus we expect, and demand, that $L(\tau) = 0$ for all $\tau < 0$. Causality puts a severe constraint on the functional form of $L(\omega)$ known as the Kramer-Kronig dispersion relations: the real part of $L(\omega)$ uniquely determines the imaginary part, and vice-versa.

In general we then have the relation that

$$g(t) = \frac{1}{\sqrt{2\pi}} \int_0^{\infty} L(t')f(t-t')dt',$$

where all functions are real valued (and hence $L(\omega)$ *must* be complex). We see that the output is some kind of weighted average of the input over all past times. This holds true for any system with a transfer function $L(\omega)$, the manner of weighting dictated by its Fourier transform $L(t')$.

Let's now consider how the low-pass character of the filter translates into the time scheme. To do so consider the simple 1st order low-pass *RC* filter of Fig. A2.




Fig. A2. 1st order *RC* low-pass filter

To solve for $L(\omega)$ we simply replace the resistance and capacitance with their complex impedances, R and $1/(i\omega C)$ respectively. The transfer function is then obvious as a frequency-dependent voltage divider. After a small amount of rearrangement we have

$$L(\omega) = \frac{1}{1 + i\omega\tau},$$

where we've introduced the familiar *RC* time constant $\tau \equiv RC$. To find the time dependence of the transfer function, we perform the inverse Fourier transform on $L(\omega)$. Using contour integration and residue theory makes the integral quite easy to perform: Since there's only a simple pole ($\omega = i/\tau$) in the upper complex- ω plane, causality is ensured because the integral is zero for times $t < 0$. The result is $L(t) = \sqrt{2\pi} e^{-t/\tau} / \tau$, and substituting this back into the convolution relation gives us

$$g(t) = \frac{1}{\tau} \int_0^{\infty} e^{-t'/\tau} f(t-t') dt'.$$

There are two points worth noting: Firstly, for an input that has been non-zero only recently ($t \ll \tau$) so that $e^{-t'/\tau} \approx 1$, the output is proportional to the integral of the input. The circuit of Fig. A2 is then referred to as an ‘integrator’. Secondly, $g(t)$ is a true average of $f(t)$ in the sense that the weighting function is normalized: *i.e.*

$$\tau^{-1} \int e^{-t'/\tau} dt = 1.$$

Let’s check up on the time-domain convolution result. First, let’s consider a unit sinusoidal input

$$f(t) = \sin(\omega t).$$

We know immediately from our frequency-domain experience that the output will be a phase-shifted sine-wave with an amplitude diminished by a factor of $|L|^{-1}$. Let’s see if we can recover this basic feature using the time-domain description. The integral is a little tedious:

$$g(t) = \frac{1}{\tau} \int_0^{\infty} e^{-t'/\tau} \sin[\omega(t-t')] dt'.$$

Transforming the integration variable to $z = t - t'$ gives

$$g(t) = \frac{e^{-t/\tau}}{\tau} \int_{-\infty}^t e^{z/\tau} \sin[\omega z] dz.$$

This evaluates (after a little algebra) to the expected result

$$g(t) = \frac{1}{1 + (\omega\tau)^2} [\sin(\omega t) - \omega\tau \cos(\omega t)] = \frac{1}{\sqrt{1 + (\omega\tau)^2}} \sin(\omega t - \phi),$$

with the frequency-dependent phase shift

$$\phi = \sin^{-1} \left[\frac{\omega\tau}{\sqrt{1 + (\omega\tau)^2}} \right].$$

Obviously, for sinusoidal inputs it's easier to work in the frequency domain using the phasor representation, but we should be reassured that the time-averaging description of the L.P. filter gives us the results we expect.

Our second example should prove equally reassuring and is a case in which the time-domain description is the easier approach. Consider a constant input that changes discontinuously from f_A to f_B at $t = 0$:

$$f(t) = \begin{cases} f_A & t < 0 \\ f_B & t \geq 0 \end{cases}.$$

Once again transforming the integration variable gives the output

$$g(t) = \frac{e^{-t/\tau}}{\tau} \int_{-\infty}^t e^{z/\tau} f(z) dz.$$

This divides into two cases. First, for $t < 0$ we have

$$g(t) = \frac{e^{-t/\tau}}{\tau} \left[f_A \int_{-\infty}^t e^{z/\tau} dz \right] = \frac{e^{-t/\tau}}{\tau} [f_A \tau e^{t/\tau}] = f_A.$$

This is good: if the input has always been f_A we certainly expect the output to be the same. Next, for $t \geq 0$ we have

$$g(t) = \frac{e^{-t/\tau}}{\tau} \left[f_A \int_{-\infty}^0 e^{z/\tau} dz + f_B \int_0^t e^{z/\tau} dz \right] = \frac{e^{-t/\tau}}{\tau} [f_A \tau + f_B \tau (e^{t/\tau} - 1)]$$

$$= [f_A - f_B] e^{-t/\tau} + f_B$$

This exponential ‘decay’ of an RC circuit is derived in all introductory textbooks (from the differential equation, usually with f_A or f_B zero). As a practical point, how long does it take for the output to accommodate this sudden input change? We can look at the fractional error ε between the output at time t and its asymptotic value:

$$\varepsilon = \left| \frac{g(t) - f_B}{f_B} \right| = \left| \frac{f_A - f_B}{f_B} \right| e^{-t/\tau}.$$

Obviously, the larger the input jump the longer we must wait for the output to settle to a given accuracy. For guiding purposes, let’s presuppose an input jump of order unity (*i.e.* the input decreases by a factor of $1/2$). In this case, our output decays to within 1% of f_B in a time $t = \tau \ln(100) \approx 5\tau$.

This is an important point to remember when using lock-in amplifiers. When the input changes for some reason (perhaps someone walks in front of your source) you need to let the instrument settle long enough to get an accurate reading.

Finally, let’s consider the frequency and time ranges we’re considering — and how they’re related. The time-domain description can be roughly summarized by the time constant τ : the circuit really only averages back over a time interval of a few τ ’s. In the frequency domain we can loosely define the ‘cutoff’ ω_c as the frequency at which $|L(\omega)|$ drops to $1/2$. Thus the circuit can loosely be said to ‘pass’ frequency components of the signal below $\omega_c = \sqrt{3}/\tau$ and ‘reject’ those above.

Remember in the case of a lock-in amplifier that we are trying to extract the DC component of the manipulated signal. Our RC -filter example illustrates the general nature of the dilemma. To get the true DC component, we would need a very good filter,

say $\omega_c \rightarrow 0$, but to do so implies that we must average the signal over a long time $\tau \rightarrow \infty$. No matter the exact nature of the low-pass filter used, we end up with an uncertainty relation⁶ in time and frequency

$$\tau \omega_c \geq \text{some constant.}$$

If our signal of interest varies with time, it is of course not really DC. We must then compromise by selecting a time constant small compared with the time variation of the signal. In doing so, we invariably increase the amount of noise included at non-zero frequencies.

⁶ Uncertainty relations occur between any pair of variables connected by the Fourier transform. In quantum mechanics they arise because of the Fourier connection between energy and time or position and momentum variables.

APPENDIX B: Chopping and harmonics.

Older models of lock-in amplifiers (like ours!) use a simpler method for phase sensitive detection: As mentioned previously, multiplying two voltage signals accurately involves sophisticated circuitry. However, it's relatively easy to change the polarity of a signal, *i.e.* multiply it by ± 1 . This is the scheme used in older, more affordable units.

Suppose we pass our input signal for a time T , invert it for another time T , and keep repeating this process. We are multiplying the input by a square wave (unit amplitude) of period $2T$. The Fourier expansion of this function is a sum of odd harmonics of our reference frequency $\Omega = \pi/T$:

$$f_{ref}(t) = \frac{4}{\pi} \sum_{n=1,3,5,\dots} \frac{1}{n} \sin n\Omega t.$$

Now consider an input signal $V(t)$ that is periodic (but not necessarily sinusoidal) with fundamental frequency ω . Most generally, it can be expanded as

$$V(t) = \sum_{m=1,2,3,\dots} (A_m \sin m\omega t + B_m \cos m\omega t).$$

Consider the product $V(t)f_{ref}(t)$. It will be a sum of terms of two types involving the products (m and n integers, with n odd)

A.) $\sin m\omega t \times \sin n\Omega t$

B.) $\cos m\omega t \times \sin n\Omega t$

We are as usual concerned with any DC (or near DC) results that are passed by the low-pass filter. As already discussed, terms of type A will give a DC contribution if

$$m\omega = n\Omega.$$

Since $\cos m\omega t \times \sin n\Omega t$ is related to the sine of the argument sum and differences, terms of type B will never give a strictly DC contribution, however if

$$m\omega - n\Omega \neq 0 \text{ but small,}$$

then we'll have a slowly oscillating noise contribution that will be passed through the filter.

Obviously if $V(t)$ is an interfering signal we want to reject, we'd better arrange for our reference frequency Ω to be incommensurate with ω . However, in practice what do you do — there is a rational ratio as near to ω/Ω as we care to approximate, and all these terms will pass through the filter? What we need is to choose $\Omega \approx m\omega/n$ so that the *smallest possible* values of m and n are still very large. (This is because our products are weighted by factors of A_m/n or B_m/n , and A_m and B_m both decrease *at least* as fast as $1/m$.) More practically, we need to avoid frequencies where m and n are small! Hence, for a background light signal at 120 Hz, we certainly need to avoid reference frequencies that are low harmonics. Fig. C1 shows the difficulty:

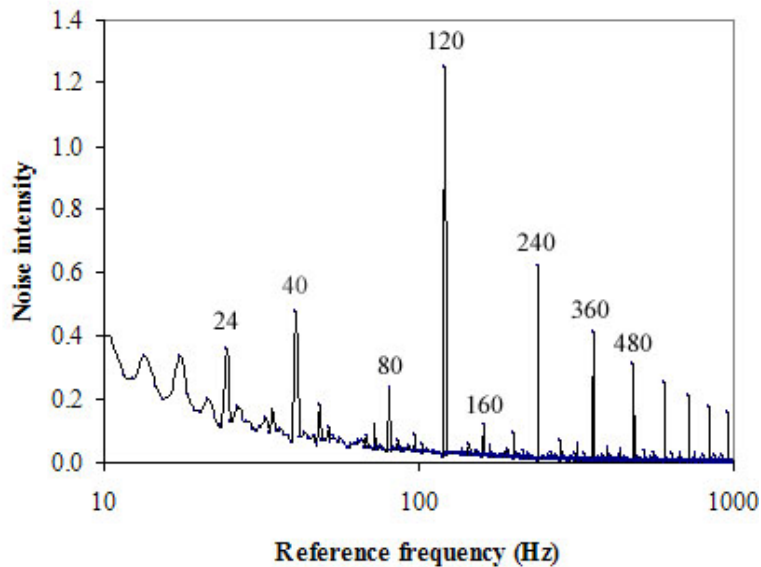


Fig. B1 Potential interference from a 120 Hz noise source.

Figure C1 is a histogram plot of the worst-case noise intensity we might expect from a 'general' 120 Hz noise source. It's been created by weighting all the harmonics for

$m, n \leq 10,000$ by $1/mn$ as discussed, and summing these amplitudes into bins 1 Hz wide. You can see the problem, we want to stay as clear of the huge harmonics as possible else the low-pass filter will let them through—but often some ‘bigish’ peak lies directly between two huge ones. We’re left experimenting to find a useful reference frequency with the noise as low as we can attain.

While we’re at it, there’s one (final) lesson we’re in a position to derive. When $V(t)$ is considered as a general *signal*, with $\omega = \Omega$, the DC component of $V(t)f_{ref}(t)$ is

$$V_{DC} = V(t)f_{ref}(t)|_{DC} = \frac{2}{\pi} \sum_{n=1,3,5,\dots} \frac{A_n}{n}.$$

Suppose for the moment that we have a pure sinusoidal input (only $A_1 = A$ is non-zero).

Then $V_{DC} = 2A/\pi$. The true RMS average of the pure sine is $A/\sqrt{2}$, so we calibrate our

lock-in by scaling it as $V_{Lockin} = (\pi/2\sqrt{2})V_{DC}$. The true RMS value of a general periodic signal is by definition

$$V_{RMS} \equiv \sqrt{\frac{1}{2T} \int_0^{2T} V^2(t) dt} = \sqrt{\frac{1}{2} \sum_{m=1}^{\infty} (A_m^2 + B_m^2)}.$$

Our *calibrated* lock-in responds to the signal as


$$V_{Lockin} = \frac{1}{\sqrt{2}} \sum_{n=1,3,5,\dots} \frac{A_n}{n}.$$

So, the lock-in **doesn’t** necessarily give the true RMS result for *non-sinusoidal* signals!

APPENDIX C: Our ‘bad’ light-bulb detector.

In our light-bulb experiment, we purposely use a simple detector based on a photo-resistor to demonstrate the power of phase-sensitive detection. However, this ‘simple’ detector actually involves us with some complications which are interesting and very instructive.

Let’s look at how it works: A photo-resistor is actually a semiconductor (CdS) device whose resistance R_p depends on ambient light intensity I . Its resistance *decreases* with *increasing* intensity. To get a voltage signal that monitors light intensity, we employ it, along with a fixed resistor R , as part of a voltage divider circuit:



A power supply provides a stable source $V_C = 5\text{ V}$, and the output voltage is simply

$$V_{out} = \frac{R_p}{R + R_p} V_C.$$

Note that even if R_p depended linearly with light intensity (which it doesn’t!) the output voltage would not. We have a nonlinear detector which is the source of our complications: recall that the lock-in

amplifier, under the conditions met in our light bulb experiment, is proportional to the chopped light intensity I_{Bulb} and dV_{out} / dI . Now, the detector output produces a large DC voltage V_{DC} determined by the average hall light (neglect the smaller 120 Hz ripple). Superimposed on this is the small, chopped light-bulb voltage we want to measure with our lock-in. Our derivative term is thusly determined by the DC level

$$\frac{dV_{out}}{dI} \approx \frac{dV_{DC}}{dI}.$$

If the background lighting were constant there would be no trouble, however as we move the detector up and down the hall, we encounter darker and lighter areas and hence V_{DC} and so dV_{out} / dI change from site to site. This is where the non-linear nature of the detector manifests itself. If we're serious about our experiment, we should correct for this.

To do so we need to know the derivative. The easiest way to do this is just to measure it. I've done this by using a super-bright LED and some neutral density filters to produce known (relatively) light levels. Here are my results:




Fig. C1 Measured detector response.

The response is, as expected, pretty non-linear. The power-law fit to my data gives a rough guide to the response is $V_{out} \approx 1.15 \times I^{-0.398}$. Let's write this for the moment as $V_{out} = \alpha I^{-\beta}$. We then have that

$$\frac{dV_{out}}{dI} = \alpha(-\beta)I^{-\beta-1} = \frac{-\beta}{I}V_{out}.$$

We don't know I , but we can invert our fit so that $(V_{out} / \alpha)^{+\beta} = I$. We therefore have

$$\frac{dV_{out}}{dI} = -\beta V_{out} \left(\frac{\alpha}{V_{out}} \right)^{+\beta}.$$

So, to correct our light-bulb data for the variations of ambient light at different points around the hall, we need to divide the lock-in's reading by $|dV_{out} / dI|$ evaluated at V_{DC} :

$$V_{corrected} = \frac{V_{Bulb}}{\beta V_{DC}} \left(\frac{\alpha}{V_{DC}} \right)^{-\beta} = V_{Bulb} \times \frac{\alpha^{-\beta}}{\beta} (V_{DC})^{\beta-1}.$$

This is why we measured V_{DC} . For careful work, these sorts of considerations are important. Of course, for careful work we might have chosen a better detector scheme — such as our photodiode setup used for the Faraday effect which is very linear.

