The Phase Sensitive (Lock-in) Detector

The "lock-in amplifier" is an instrument used in many physics experiments because of its special effectiveness in reducing noise in electrical measurements. But unlike instruments such as oscilloscopes and various types of meters, its operating principle is somewhat subtle. This laboratory will lead the student through the principle of phase sensitive detection—the heart of the lock-in amplifier—and will explore the basic function, features and limitations of the lock-in amplifier—the commercial box built around the phase sensitive detector.

1 Introduction

Many measurements in experimental physics involve the detection of an electrical quantity, either a voltage or a current. Some physical quantities are intrinsically electrical in nature, for example the voltage drop across a diode, or the emission current in a vacuum tube. Other quantities such as temperature, pressure, displacement, or light level can be converted to electrical quantities by means of transducers (sometimes called sensors). Perversely, the electrical quantity (signal) of interest is accompanied by noise, the latter sometimes orders of magnitude greater than the former. Various techniques exist to recover the signal of interest from the composite of signal + noise, and one technique, phase sensitive (lock-in) detection will be explored in this experiment.

Consider the following question: Can one, by being sufficiently careful and clever, reduce the noise content of an electrical signal coming from a transducer to arbitrarily low levels? The answer, in a word, is no. The signal from any transducer with resistive or diode-like characteristics will, from physical principles, have an irreducible amount of (Johnson or shot) noise on its output signal. In addition, some flavor of reducible noise ($1/f$, electromagnetic interference, microphonic, to name a few) is almost always present on transducer output, buttressing the claim that noise will be present at some level on any transducer signal. Finally, just amplifying the composite signal won't help make the signal of interest more distinguishable from noise, since amplifiers boost the level of everything present at the input (noise included), and contribute noise of their own to boot! (For a discussion of the various kinds of electrical noise see Horowitz and Hill, pp. 430-436, as noted in the references at the end of this section).

However, if the experiment in question is one in which measurements are made of the response to a controlled excitation, there may be a way out of the noise quagmire. An example of such an experiment is the measurement of the resistance of a circuit element: apply a known current and measure the corresponding voltage drop. Another example is an optics experiment involving the measurement of the ratio of scattered to incident light intensity. In such situations the experimenter may be able to impart some unique characteristic to the excitation and then measure the response by a method that strongly selects for that characteristic. For instance, in the resistance example, one could use an AC current of a particular frequency f_m and then measure the AC voltage with an instrument tuned to f_m.

The obvious way to accomplish this tuned measurement would be to first modulate the excitation with an AC signal at f_m, place a narrow bandpass filter also tuned to f_m in the response-signal path to filter out all signals but the ones at f_m, and then rectify the result. In many cases the improvement in noise rejection will be sufficient. But what if you want to change the frequency of the excitation? And what if you would like to vary the bandpass width? More sophisticated filter designs might be needed. Moreover, in practice it is difficult to make an extremely narrow bandwidth filter. The bandwidth of a filter may be specified by its Q value, which is defined as the ratio of the center frequency f_0 to the range of frequencies Δf between the points where the filter response falls by one-half: $Q = f_0/\Delta f$. Practically, it is hard to build and use filters with $Q >\approx 50$.

The phase sensitive detector makes an elegant end run around these problems by reversing the order: it first rectifies the signal and then filters it. We'll look at how this is done, and then at why this works so well. Figure 1 shows the basic phase sensitive detector or PSD, greatly simplified. The PSD is composed of two basic circuits: a synchronous switch and a amplifier/filter.

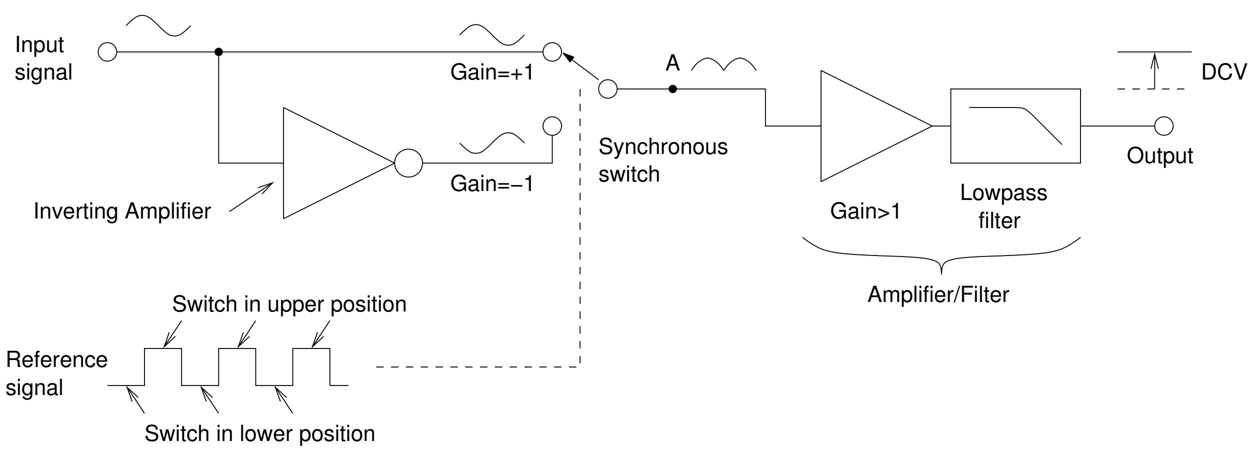

Figure 1: Simplified schematic of the phase sensitive detector.

Consider, first, the action of the switch. It is designed so that it spends an equal amount of time in each position corresponding to the period of a reference signal $T_m = 1/f_m$. In the upper position, it passes the signal unchanged; in the lower position it passes the signal inverted. If we apply a sine wave to the input which has the same frequency as the reference, $V(t) = V_0 \sin(2\pi f_m t + \phi)$, the signal at the output of the switch (at point A in the figure) will depend upon the phase angle ϕ between the reference signal and the input signal.

If the phase angle $\phi = 0$, the two signals are in phase, and the result at point A looks like the lower trace of Fig. 2a: the sine wave has been rectified. If the result is then passed to the amplifier/filter stage, with a time constant τ of the filter sufficiently long, the output signal will be a constant (DC) signal whose value will be proportional to the amplitude V_0. However, if $\phi = 90°$, as in Fig. 2b, the signal at point A will be symmetric about the zero-volt axis, and the output from the amplifier/filter would be zero. Likewise, it is easy to see that if $\phi = 180°$, the output would be proportional to $-V_0$. This dependence on the phase between the input and the reference signals is why this method is called "phase sensitive detection".

Qualitatively, it is now easy to see why this method helps to significantly reduce noise. If one were to apply a signal to the input at a different frequency f than the reference frequency f_m, the phase ϕ would be constantly changing, and the net effect would be to produce a signal that would average to zero.

With this introduction in mind, we will now lead you through a collection of investigations that should make these ideas clearer and more quantitative. The investigations will be in three parts:

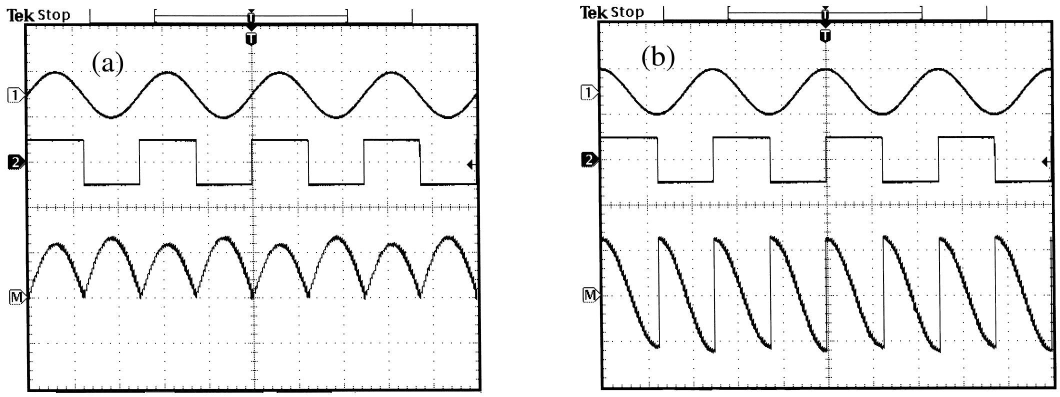

Figure 2: Oscilloscope traces simulating the effect of the synchronous switch section of the phase sensitive detector. (a) Reference (square wave) and signal (sine wave) in phase produce a net positive waveform; (b) Reference and signal 90° out of phase produce a waveform with a zero time average.

1. **Characteristics of the Keithley 822 Phase Sensitive Detector.** First you will study how one particular PSD, the Keithley 822 behaves when presented with a collection of pure signals. You will look at the output of the synchronous switch and the amplifier/filter, and calculate the gain of the amplifier stage. Then you will apply pure signals which are shifted in frequency from the reference by a controlled amount.

2. **Study of a noisy model system.** A box with a photodiode is illuminated with two LEDs. One LED produces a periodic (square-wave) light and the other produces a random light from a digital noise generator. The amplitudes and frequency content of the light from the two LEDs can be controlled. You will use the PSD to measure the signal in the noise for a variety of settings, and study the effects of different filter time-constants.

3. **Quantitative measurement of noise.** You will apply a sine wave of a known amplitude mixed with broadband noise from a white-noise generator to the lock-in, and then measure how the RMS noise amplitude depends on filter time-constant settings. This will allow you to estimate quantitatively how well the PSD can select for a signal buried in a very noisy background.

The following list contains references to short tutorial papers produced by some manufacturers of lock-in amplifiers. These are available for download from the course web site.

References

[1] Ithaco Application Note IAN-35, The Evolution of the Modern Lock-in Amplifier, by Edwin H. Fisher. (Available online).

[2] Ithaco Application Note IAN-47, Introduction to Lock-in Amplifiers, by J. L. Scott. (Available online).

[3] Ithaco Application Note IAN-49, Speed/Accuracy Tradeoff When Using a Lock-in Amplifier To Measure Signals in the Presence of Random Noise, by J. L. Scott. (Available online).

[4] Signal Recovery Technical Note TN-1000, What is a Lock-in Amplifier? (Available online).

[5] P. Horowitz and W. Hill, The Art of Electronics, 2d ed., Cambridge University Press, New York, NY, 1989, pp. 430-436 (noise in general), pp. 1026-1035 (bandwidth narrowing and lock-in detection).

[6] D. Preston and E. Dietz, The Art of Experimental Physics, Wiley, New York, NY, 1991, pp. 367-375 (lock-in detection in general and modulation technique applied to spectroscopy measurements).

2 *Investigation 1: The Keithley 822 PSD*

First, please be aware of several precautions:

1. The Keithley 822 is a venerable, irreplaceable old instrument (built 1970). Before connecting any signal to the SIGNAL input, look at it on the scope and make sure that it does not exceed 3 volts peak to peak. (The OVERLOAD light next to the SIGNAL input will come on if the input is above this level.) A large overload may cause damage; any overload at all will introduce distortion and degrade the performance of the instrument.

2. For certain parts of the following, the analog meter on the Keithley 822 must be turned off (-/OFF/+ switch in center position). If left on the meter will vibrate very vigorously and may be damaged. This warning is repeated at the points in the procedure where the meter needs to be turned off.

3. At all times use DC coupling for the input signals to the scope. Using AC coupling will mask DC level shifts which are an important part of some of the measurements.

Now on to the investigation. To begin, set the switches on the Keithley 822 PSD as follows:

- ZERO: OFF
- METER: 3V and OFF
- TIME CONSTANT: OFF, set knob to .001 seconds (fully CCW).
- FILTER: OFF
- REF: A

Turn the PSD power on to let it warm up.

Turn on the Wavetek 29 function generator, and set the parameters by using the front panel buttons. To adjust the waveform type, frequency and amplitude, follow these steps:

1. Among the FUNCTION keys, select the one marked SINE.

2. Push the FREQ/PER button, and then use the DIGIT buttons to position the cursor on the numerical part of the frequency reading.

3. Use the keypad to enter '1', '0', '0', and then the unit key 'Hz ms dBm' to set the frequency to 100 Hz.

4. Use the FIELD buttons to position the cursor on the numerical part of the amplitude field.

5. Use the numerical keypad to enter '1', and then the unit key 'MHz ns V' to set the amplitude to 1.00 Vpp (volts peak-to-peak).

At various points in the investigations, you will need to adjust the phase between the reference signal from AUX OUT and the main signal from MAIN OUT. This requires editing the trigger parameters:

1. Push the EDIT (dark blue) and then the TRIG buttons.
2. Use the FIELD buttons to move the cursor to the numerical part of the phase field.
3. Use the keypad and ENTER keys to set the phase. For example, push '9', '0', '+/−' and ENTER to set the phase of the AUX OUT signal to $-90°$ relative to the MAIN OUT.
4. To begin, confirm that the phase reads "000" degrees.

Turn on the digital oscilloscope. Connect the equipment together as in Fig. 3. The idea is to supply the PSD reference signal from the logic-level (0-5 volt) square wave output AUX OUT on the function generator, and then to compare the input to the PSD with the output from the PSD. After you have set the amplitude of the signal generator to 1.0 volts peak-to-peak, press the OUTPUT button to turn the signal ON. Set both channels to have 500 mV/div vertical sensitivity, and set up the oscilloscope to trigger on the channel 1 signal.

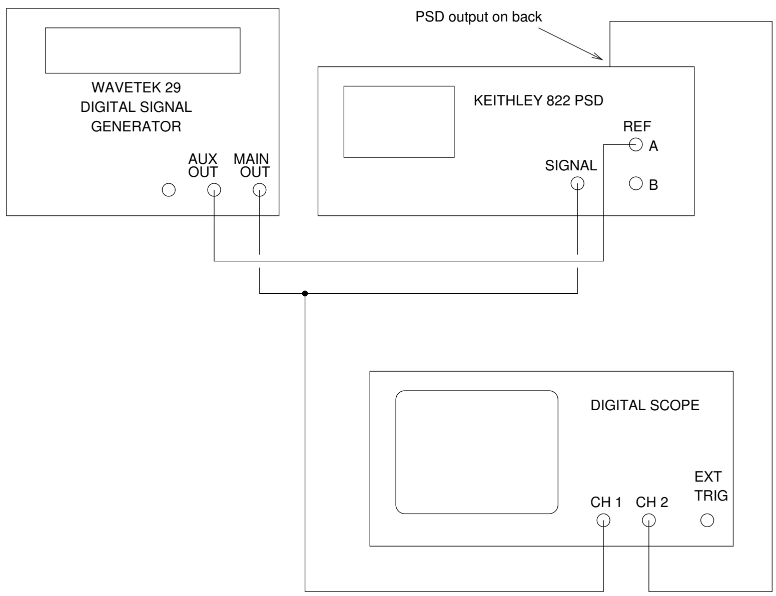

Figure 3: Connection of equipment for the first part of Investigation 1.

If all has been set up correctly, you should see two waveforms on the scope which look like the upper and lower traces of Fig. 2a, but with one notable difference: the lower waveform will be inverted. Can you guess what the reason for this inversion might be? If not, hold your thoughts until you look at the effect of the amplifier/filter in a moment.

Exercise 1 *Use the printing feature of the digital oscilloscope to produce pictures of the input and output waveforms of the Keithley 822 for three (3) different phase relationships between the AUX OUT and MAIN OUT signals from the Wavetek 29:* $0°$, $+90°$, *and* $180°$. *Make sure that the TIME CONSTANT switch on the Keithley is set to OFF, and remember to record the instrument settings for each picture.*

Reset the phase parameter to $0°$, make sure that the TIME CONSTANT knob is set to .001 seconds, and turn the TIME CONSTANT switch to INT (INTernal capacitor). Note what happens to the output waveform. You may see the trace go off screen—this is the effect of the amplifier part of the amplifier/filter.

Adjust the gain of CH 2 on the oscilloscope to bring the output waveform comfortably on screen; about a factor of 10 is needed. Now you should see two differences from the waveform you saw earlier: the output will no longer be inverted, and the output will be smoothed somewhat. Turn the TIME CONSTANT knob to successively higher time constants, and note what happens to the output waveform. Try changing the phase parameter now, and see what happens to the output with various time constants and phases.

Consider the question: what is the gain of the amplifier of the amplifier/filter? This is a little tricky to find, since the waveform is changed so much by the filter. Here's a way: try sending a square wave into the input. Select SQUARE from the FUNCTION keys, and make sure the amplitude is 1.00 volts peak-to-peak as measured on the oscilloscope. Note the output of the PSD; it should be very close to a straight line, even with a .001 second time constant. (But, notice what happens when you set the phase to 90°; the waveform is no longer a straight line!) When the phase is at 0°, the synchronous switch sends a DC signal to the amplifier/filter since it switches at exactly the point where the square wave changes between low and high. The filter does not make a change in a DC waveform, so the result depends only on the gain of the amplifier.

Exercise 2 *Use a square wave input signal to measure the DC gain of the Keithley 822 PSD. To measure the output voltage accurately, use either the Keithley 2000 digital multimeter, or the MEASURE menu of the digital scope. (Press MEASURE and select the "Mean" softkey for CH 2.) To compensate for a small DC offset, make your measurements at both 0° and 180° phase settings, and average the absolute values of the readings. If you get a DC gain of 10, you have made a small error.*

Change the FUNCTION back to SINE, and set the PSD time constant to .1 seconds. Measure the output voltage now. Why is it different for a sine wave than for a square wave? To answer this question remember that the filter *averages* the signal over a period of time that increases with increasing time constant. The average of a time-dependent signal $V(t)$, as measured at point A in Fig. 1, over one cycle T_m of the synchronous switch is given by this formula:

$$\langle V_A \rangle_{T_m} = \frac{1}{T_m}\left(\int_0^{T_m/2} V(t)\,dt - \int_{T_m/2}^{T_m} V(t)\,dt\right) , \tag{1}$$

where the integral has been broken into two intervals in order to account for the effect of the inverting switch. One can see by inspection that a square wave having the same phase as the reference, a period T_m, and a peak-to-peak amplitude of $2V_0$ will average to V_0. If we apply Eq. (1) to a sine wave of the form $V(t) = V_0 \sin(2\pi f_m t + \phi)$, we get

$$\langle V_A \rangle_{T_m} = \left(\frac{2}{\pi}\right) V_0 \cos\phi \ . \tag{2}$$

Note that for $\phi = 0°$, the average is not V_0 as in the case of a square wave, but is a factor of $2/\pi \approx 0.64$ smaller.

Exercise 3 *Prove Eq. (2). Then use your value of the DC gain that you found with a square wave input along with measurements of the output of the PSD given a sine wave input to verify that the factor of $2/\pi$ holds for the measurements.*

Now we will look more closely at the dependence on ϕ as well as the response of the PSD to signals which differ in frequency from the reference frequency f_m. For this we will use the device called the "Digital Doppler Shifter" or DDS. This box generates a 1 volt peak to peak sine wave at 312 Hz, and a second sine wave, also 1 volt peak to peak at 312 Hz+Δf, where Δf can be varied from 0 to ± 9.99 Hz in .01 Hz steps. The DDS also generates a third signal, 3 volts peak to peak at 312 Hz to drive the reference input of the lock-in. With this device, one can see clearly how the lock-in treats signals close in frequency to the reference frequency. (Note: F on the front panel of the DDS means the same as f in this write up.)

With f different from f_m, the phase parameter ϕ will be time dependent. Assume for a moment that we feed the PSD a signal of the form $V(t) = V_0 \sin(2\pi f t)$ where $f \neq f_m$. We can write $f = f_m + \Delta f$ and $V(t) = V_0 \sin(2\pi f_m t + 2\pi \Delta f t)$ and we identify the time dependent phase $\phi(t) = 2\pi \Delta f t$.

First a description of the DDS: the frequency difference Δf is set by the thumbwheel switches. The RESET switch zeroes the phase difference between the two outputs. The BLUE (+) and RED (−) switch positions select the sign of Δf. In the center (HOLD) position, the phase shifting is stopped and the phase difference ϕ is held constant. In this position the frequencies of the two outputs are identical since $\Delta f = (1/2\pi)(d\phi/dt) = 0$. The phase of the 3 volt peak to peak reference signal can be set to either 0° or 180° with respect to the 1 volt signal at frequency f. Finally, a two color LED gives an indication of the phase difference, pure green indicating no phase difference and pure red indicating a 180° phase difference.

2.1 PSD output as a function of phase difference.

By shifting the phase between the two outputs (switch in RED or BLUE position) for an appropriate period of time, and then holding the phase difference constant (switch in HOLD position), the DDS can generate two sine waves of identical frequency and any desired phase difference. Connect the various instruments as shown in Fig. 4. Select a time constant τ of .1 second. Set the frequency difference to some low value (say .01 Hz), put the three position switch in the HOLD position, and zero the phase difference with the RESET switch. Set the oscilloscope trigger to EXT and use the horizontal position knob to park the trigger point (as indicated by the "T" symbol) at the middle of the trace. Adjust the trigger level to make the sine wave input trigger at the positive zero crossing. The output of the PSD should now be maximum positive. Verify that this is so by varying and then holding the phase difference in both directions. Now introduce a phase difference of approximately 90° as observed on the scope, hold this phase difference and verify that the lock-in output is approximately zero. Next introduce a phase difference of 180° as observed on the scope and verify that the lock-in output is at or near its most negative value. Finally, observe the lock-in output for a 270° phase difference. Is there any way to distinguish between a phase difference of 90° and one of 270°?

2.2 PSD output as a function of frequency difference.

Note: Turn off the meter on the lock-in before doing this part. If left on it will vibrate vigorously and may be damaged.

By now, you will have certainly noted that the output of the PSD varies in a periodic way according to the value of Δf chosen on the thumbwheel switches. If the input signal to the PSD is a sine wave shifted in frequency from the reference frequency by an amount Δf, the output of the lock-in will oscillate at a frequency Δf.

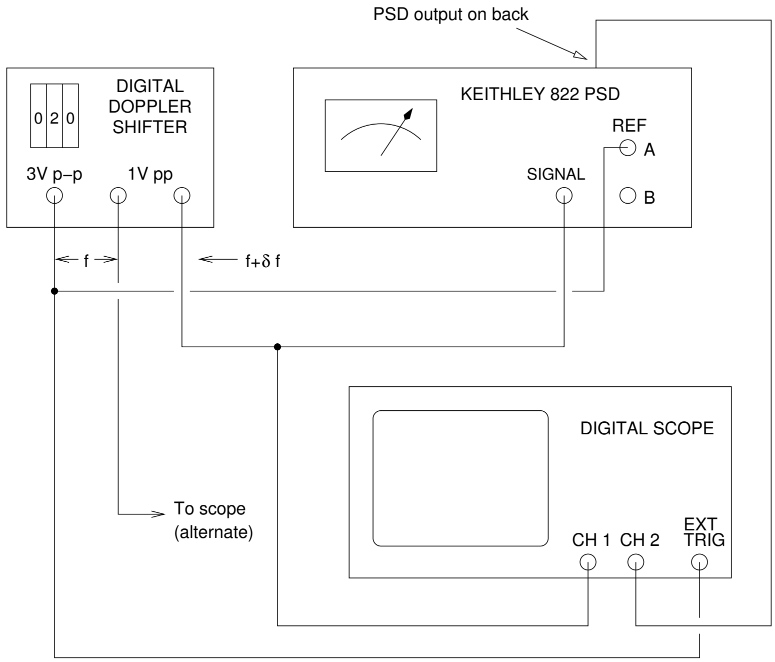

Figure 4: Connections for using the digital Doppler shifter.

Because of the filter stage, as Δf gets larger, the amplitude of the oscillations at the PSD output decreases. With the DDS we can now observe quantitatively how the output depends on Δf with a fixed PSD time constant.

Exercise 4 *Set the time constant to 1.0 seconds, and use the scope to measure the peak to peak output voltage of the PSD for values of Δf between .02 Hz and 6–8 Hz. Plot the results on log-log paper (or use a program like KaleidaGraph). Hints: Use the horizontal cursors to mark the positions of the maximum and minimum output voltages. As the peak to peak voltage decreases turn up the vertical gain on the scope to increase the resolution of the measurement. As Δf increases, increase the sweep speed on the scope. For a good looking graph, space your difference frequencies logarithmically, that is, use .02 Hz, .04 Hz, .08 Hz, etc.*

Your graph should show two distinct behavior regimes—one where increasing Δf by a factor of 2 makes little or no difference in the output, and a second where increasing Δf by a factor of 2 decreases the output by a factor of 2. This is just classic low pass filter behavior, with the response falling off as $1/\Delta f$ for $2\pi\Delta f\tau \gg 1$. The implication of these results is that the Keithley 822 has only one low pass filter on the output. Many lock-ins have multiple low pass filters in series on the output, with the number of filters in the circuit selectable from the front panel. Two sections (filters) is the most common choice as it often proves a good trade-off between increased signal to noise ratio and increased response time.

For pure sinusoidal signals of the form $V(t) = V_0 \sin(2\pi ft + \phi)$ we can write the overall response of the PSD as follows:

$$V_{PSD}(\Delta f, t) = \frac{A_{DC}\left(\frac{2V_0}{\pi}\right)\cos(2\pi\Delta ft + \phi)}{\sqrt{1 + (2\pi\Delta f\tau)^2}} \ , \tag{3}$$

where $\Delta f \equiv f - f_m$ and A_{DC} is the DC gain of the amplifier/filter. The denominator represents the effect of the filter stage in reducing the output for frequencies which differ from the reference frequency f_m.

We can now summarize the effect of the phase sensitive detector: all input signals are shifted downward in frequency by an amount equal to the reference frequency, with the frequency of interest appearing at the output as a DC (0 frequency!) level, and all signals at other frequencies appearing as fluctuations at the difference frequency attenuated by an amount depending on the frequency difference (Δf) and the time constant setting (τ) of the amplifier/filter stage. Simply put, the instrument "locks-in" to the signal at the reference frequency, and rejects, to varying degrees, all others, thus the name "lock-in amplifier".

3 *Investigation 2: Study of a Noisy Model System*

In this investigation, you will use the PSD, which we will now call by its more common name "lock-in", to see how it can extract a signal of interest from a composite of signal plus noise that is dominated by the latter. You will also see that there is a trade-off involved in reducing the noise present at the lock-in output (nothing is free!).

First a brief description of the hardware for this part: inside the LED/Photodiode Box are two LED's (light emitting diodes) and a photodiode detector. The LED's are driven by CMOS logic gates with resistors in series to limit the LED current. One LED (signal) is driven by a periodic square wave whose frequency can be varied between approximately 105 and 145 Hz (by means of the Frequency Adjust knob on the Control Box). The other (noise) LED is driven by a digital noise generator (officially known as a pseudo random bit sequence, or PRBS generator) inside the Control Box. When this LED is on you can see it flicker (noise!). The logic level (0/5 volts) outputs for both signal and noise are available at the Signal Output and Noise Output connectors respectively on the Control Box. (For circuit details and output characteristics of PRBS generators see Horowitz and Hill, pp. 655-664).

Both LED's are aimed at the photodiode detector. The output of this detector (a current) is proportional to the sum of the light from the individual LED's (plus any other light incident on the detector—hence the cover for the box). The photodiode current is converted to a voltage inside the Control Box and this voltage is available at the Photodiode Output connector.

Two switches control the intensity of the signal and noise LED's, and the approximate outputs, given as peak-to-peak voltage measured at the Photodiode Output connector, for the various settings are indicated below:

	Signal LED (mV)	Noise LED (mV)
Lo	25	240
Med	60	600
Hi	140	1400

Two additional switches enable one to turn off either LED completely—useful for observing signal and noise components separately.

The digital noise generator requires a square wave input, as shown in Fig. 5. The frequency of this input determines the noise amplitude at any given frequency. Figure 6 shows traces of the logic level outputs and the Photodiode Output which results from combining these signals; The bottom trace is the Photodiode Output with the Signal LED set on "Hi" and the Noise LED set on "Med";

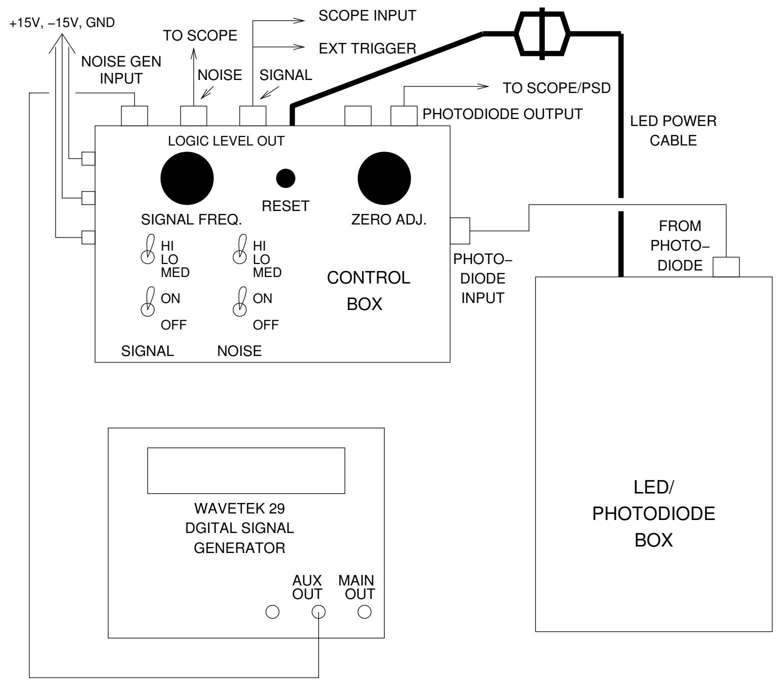

Figure 5: Connections for the LED/Photodiode box investigation.

the middle trace is the logic-level Signal Output with $f_{sig} \approx 114$ Hz, and top trace is the logic-level Noise Output with the input to the noise generator at 1.3 kHz. The noise waveform clearly has frequency components near the signal frequency (how would you justify this claim?).

Note: To shield the photodiode from stray room light, the cover for the LED/Photodiode box should be in place for observations and measurements in parts 3.1, 3.2 and 3.3.

3.1 Observing the signal and noise waveforms.

Connect the various boxes and instruments as shown in Fig. 5. Connect the logic level Signal Output on the Control Box to one scope input and the logic level Noise Output to the other scope input. Turn on the Wavetek 29 function generator and set the frequency to 1 kHz (see Investigation 1 for info on setting the Wavetek 29 controls) and then turn on the power to the Control Box. With the LED On/Off switches in the On position, lift the cover and check that both LED's are illuminated. If the noise LED is not flickering, push the RESET button. Turn the Frequency Adjust knob all the way counterclockwise so that $f_{sig} \approx 105$ Hz. Select External triggering on the scope (push the Trigger Menu button and then the Ext button) and adjust the trigger level (≈ 1 volt works nicely) to achieve a stable display. (Externally triggering the scope on the Signal Output will result in a stable display any time you are looking at anything happening at the signal frequency). By pressing the RUN/STOP button, you can capture and hold one sweep at a time, a useful feature given that the noise signal changes from sweep to sweep. By adjusting the time base and vertical sensitivities you should be able to get a display similar to that in the top two traces of Fig. 6. This display will help you get a feeling for what the signal and noise components look like.

Now disconnect the Noise Output from the scope and connect the Photodiode Output (on the Control Box) to the scope. Turn off the noise LED so that only the signal shows up on the

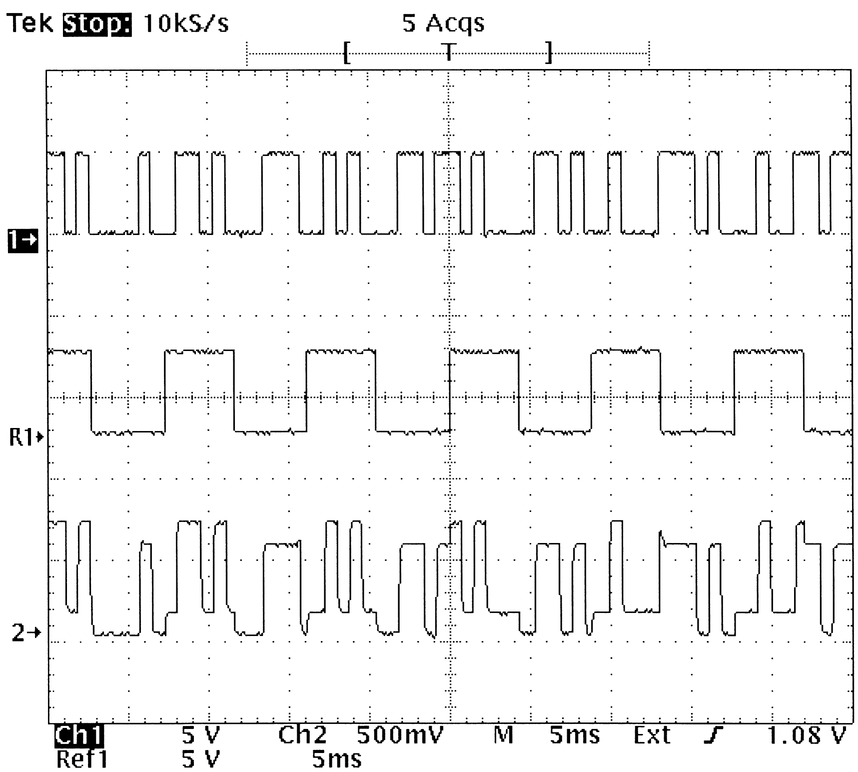

Figure 6: Oscilloscope traces showing output signals from the LED/Photodiode Control Box. Top: Logic-level “noise”, middle: logic-level “signal”, bottom Photodiode Output showing combined noise and signal.

Photodiode Output. Set the signal LED to medium intensity. To get the Photodiode Output positioned sensibly on the scope display it may be necessary to adjust the Zero Offset of this output with the upper right knob on the Control Box; note the position of the waveform with respect to the channel reference marker on the left side of the display screen. You will also need to adjust the channel sensitivity. Note the effect of the various switch settings (Hi, Lo, Med) on the waveform.

Switch the Signal LED off and the noise LED on to "Med" intensity. Use of the RUN/STOP button is essential here since the scope display will be constantly changing and hard to see. Try the different switch settings to see the different amplitudes of the noise waveform. The noise should be about 10 times larger than the signal

Now that you've seen what the signal and noise detected by the photodiode look like separately, turn on both LED's (medium intensity). You should now get a scope pattern like that in lower trace of Fig. 6. By looking at the trace closely you can see the Photodiode Output increase when the signal LED turns on (lower trace goes from low to high) and decrease when it turns off.

Exercise 5 *Use the cursors on the digital oscilloscope to measure the peak-to-peak amplitudes of the photodiode signal for all three settings (Hi, Lo, Med) of the Noise LED and the Signal LED, separately. Compare your values with the values listed in the table above. Then calculate the signal-to-noise ratios of the following switch combinations: (1) Both Signal and Noise set to "Med"; (2) Signal set to "Lo" and Noise set to "Hi". The signal-to-noise ratio, SNR, is defined by the formula* $\mathrm{SNR} = 20\mathrm{dB} \times \log_{10}\left(V_{sig}/V_{noise}\right)$*, where the two voltages must be of the same type—RMS, peak-to-peak, etc.*

3.2 Noise at the lock-in output vs. time constant setting.

Now let us see what the Keithley 822 lock-in does with the signals from the LED box.

1. Set the time constant to 1 second and the associated switch to INT.
2. Set the meter to 1 volt full scale, positive. The FILTER switch should be in the OFF position.
3. Connect the Signal Output on the Control Box to the REF A input on the lock-in and set the REF switch to A.
4. Connect the digital multimeter or DMM (Keithley 2000) and CH 2 on the digital oscilloscope to the lock-in output. The lock-in the REF LEVEL light should come on indicating that the REF signal is at least 3V P-P.
5. Connect the Photodiode Output from the Control Box to the SIGNAL input of the lock-in.
6. Turn both the Signal and the Noise LEDs off; then use the zero knob and its associated switch on the lock-in to zero the output.
7. Now turn on the Signal LED to "Med", and compare the amplitude of the Signal square wave (input to the lock-in) to the output DC level. Is it about right? (Recall what the lock-in does to an in-phase square wave).
8. Turn on the Noise LED to "Med", and notice the lock-in meter, the reading on the DMM and the waveform on the oscilloscope. The lock-in meter should be fluctuating about an average

value of approximately 0.3 volts. Turning off the noise LED should cause the fluctuations to cease. Figure 7 shows oscilloscope screen for the situation where the Noise LED is turned off midway through the trace; notice how the lock-in output varies about the mean with the Noise LED on, and settles to the mean when the Noise is turned off.

9. With both signal and noise LED's on (still medium intensity), vary the time constant setting on the lock-in between .1 and 10 seconds. How does the lock-in output (needle movement and scope trace) change?

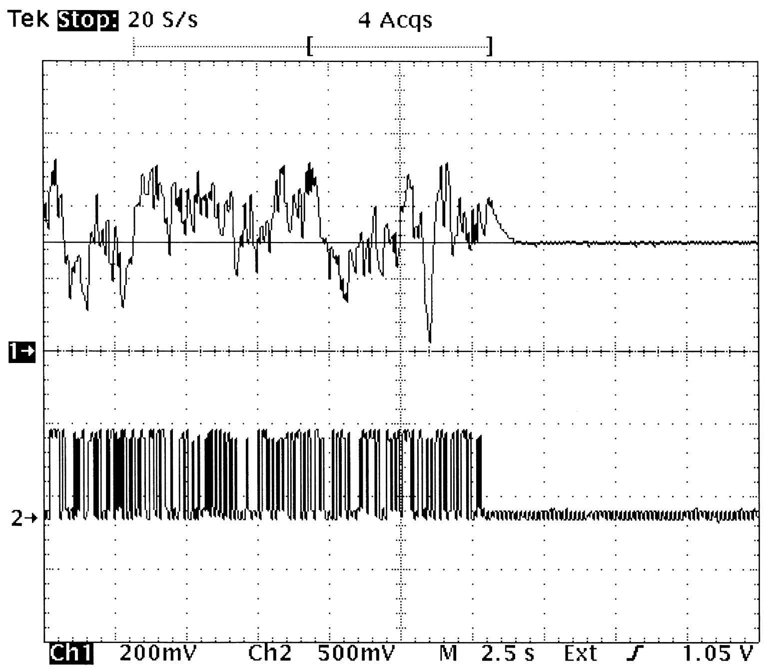

Figure 7: Scope traces for a noisy-changing-to-quiet signal. Lower trace: Photodiode output with noise+signal (both at "Med" setting)switching to signal only at about the 6th division. Upper trace: corresponding lock-in output; time constant = 0.3 seconds.

Exercise 6 *Let's pursue this slightly more quantitatively (the seriously quantitative part comes in the next section). Slow down the horizontal sweep speed so that the time per division on the scope display is approximately equal to the time constant setting on the lock-in. Use the Measure feature of the scope to find the peak to peak (Pk-Pk in Tektronix shorthand) and average (mean) values for the sweep. The peak-to-peak value is an indication of the amount of noise on the lock-in output, and the average is a measure of the amplitude of the signal of interest (the signal LED intensity). Print screen grabs for time constant settings of .1, 1, and 10 seconds, and record the peak-to-peak and mean voltage of the lock-in output as a function of time constant. The trend is obvious. Any prediction as to noise vs. time constant relationship?*

3.3 Lock-in response time vs. time constant setting.

It's clear that you win big in terms of reducing the noise at the lock-in output by increasing the time constant, which may be just what is needed to observe a very small signal. But is there a down side to this? Well, no and yes—it depends on the rate at which the quantity, and associated signal you are measuring, is changing. If the signal amplitude is rock steady over a long period of time, the time constant can be made very long (up to 30000 seconds on some lock-ins!) without sacrificing anything (except time). But what if the signal of interest changes amplitude rapidly and you want to track that change on the lock-in output? Then you have to reduce the time constant.

Thinking about this situation in different terms, the time constant τ of low-pass filter circuit tells you how quickly (or slowly!) the circuit will relax to a new equilibrium when the DC input level is changed. For the lock-in filter, the DC output level corresponds to the magnitude of the signal of interest at the lock-in input. If the latter signal amplitude changes abruptly, the lock-in output will take an amount of time determined by the time constant τ to reach a new equilibrium (about 5τ is required to be within 1% of the new equilibrium value).

Exercise 7 *Now you will take some measurements to see how the output response (relaxation) time varies as a function of time constant setting. Check that the lock-in output is still connected to a scope input. Turn both LEDs off and check that the scope trace shows 0 volts. Turn on the Signal LED (medium intensity) and leave off the Noise LED. Set the time constant to .1 second. With the scope running on "auto" triggering, and a sweep speed of 500 ms/div, switch the Signal LED off, and then press RUN/STOP just before the transition reaches the left end of the screen. Getting a suitable trace requires some timing coordination between turning the LED off and stopping the sweep. The output time constant (time it takes to make 63% of the change between initial and final values) can be estimated from the scope display by using the cursors. Calculate the voltage where the signal drops by 63%, park the cursors at the start of the drop and at the 63%-drop point and read the time difference from the screen. Repeat this procedure for time constant settings up to 10 seconds. Plot the results.*

What if noise is present (which it would be, otherwise you don't need a lock-in)? Can one follow a rapid change in the signal? The next exercise will explore this issue.

Exercise 8 *Set the Signal LED to "Med", the lock-in time constant to .3 seconds, and the oscilloscope trace speed to 500 ms/div. Leave the Noise LED off, and switch the Signal LED between on and off while looking at the screen. Notice how you can clearly see the point where the transition occurs at the lock-in output, even if you change the time constant. Print a screen grab of a representative Signal LED on to off transition.*

Next, turn the Noise LED on to "Lo", and watch the screen as you turn the Signal LED on and off again. Can you see where the transition happens? Try "Med" and "Hi" settings on the Noise LED and repeat the Signal transition. Print a screen grab of a transition with the Noise LED on "Med".

You might think that increasing the time constant would improve matters; try it, and print representative screen grabs. (You may want to increase the time/div setting on the scope.) What happens? Discuss.

The previous exercise shows (or should show, if done correctly) that there are cases where the lock-in can't save the day all by itself. It needs some help in the form of reducing the noise (or resorting to some additional technique like signal averaging the lock-in output).

The question of transient response arises in the Low Temperature Superconductivity experiment in the Advanced Labs. A sample of mercury (inside a glass ampule) sits in a bath of liquid helium and around the sample is a coil which is part of a bridge circuit that is excited at the reference frequency. A slowly varying magnetic field is applied to the sample. As the field changes, the mercury abruptly transitions from normal to superconducting (or vice-versa), its susceptibility changes abruptly, and the input to the lock-in (the lock-in detects the null in the bridge circuit) changes abruptly. If the time constant of the lock-in is set too long, the output will not show the transition clearly. On the other hand, if the time constant is set very short, the noise on the lock-in output gets pretty nasty, and it begins to obscure the transition. This is a real life case where the competing considerations of output noise vs. response time have to be sorted out carefully.

For detailed discussions of the output noise vs. response time issue, see Ithaco Application Notes 47 & 49.

3.4 Why 120 Hz (or 60 Hz) is bad news.

Turn off the meter on the lock-in before doing this part. If left on it will vibrate vigorously and may be damaged.

To wrap up Investigation 2, we will study an example of the effect of power-line frequency noise. In North America, the standard frequency of AC electrical power is 60 Hz, thus the wires in the walls radiate electromagnetic waves at this frequency, and many appliances, such as lights and motors flicker or buzz at 60 Hz and its harmonics.

Let's look at a situation where this can be a problem for lock-in detection. Set the frequency of the signal LED (Frequency Adjust knob) to within 2 or 3 Hz of 120 Hz. Use the Measure feature of the scope to determine the signal frequency. Turn on the signal LED and turn off the noise LED. Check that the lock-in output is still connected to the scope. Set the lock-in time constant to .1 second. Remove the cover to the LED/Photodiode Box so that some (not much) incandescent light from the desk lamp falls on the photodiode. The lock-in output should now have a low frequency component at the difference frequency between the signal frequency (which is the lock-in reference frequency here) and 120 Hz, the frequency at which the light intensity varies (why 120 Hz and not 60 Hz?). Change the lock-in reference frequency and note the effect (amplitude and frequency) on the lock-in output.

Exercise 9 *Print some screen grabs from the scope for different signal frequencies. Explain what you observe.*

It's obvious why one wants to avoid operating at or near 120 Hz (or 60 Hz), and indeed, many lock-in amplifiers have built-in filters to reject, right at the input, any signals at 60 or 120 Hz.

4 *Investigation 3: Quantitative measurement of noise*

This part has much in common with Investigation 2, but a different combination of instruments is used to achieve a more quantitative measure of noise at the lock-in output as a function of time constant setting. A fixed frequency sine wave (this is the signal of interest) is combined with noise and the composite is connected to the lock-in input. The lock-in reference is driven by a square wave at the same frequency as the sine wave. The output of the lock-in is measured by the Keithley 2000 digital multimeter (DMM). This instrument can be set up to find the average (mean) value of a set of readings, the high and low values, and the RMS value of the fluctuations about the mean, which is just the RMS value of the noise component of the lock-in output. Let's get back to the hardware.

Both the sine wave signal and square wave reference are provided by the function generator. The noise component is generated by the General Radio Random-Noise Generator (RNG). Believe it or not, there are instruments whose sole function is to produce noise! This instrument produces a special kind of noise, a noise which has an equal amount of power in any frequency interval, up to some cut-off frequency (as set by the RANGE switch). This kind of noise is called "white" noise, in analogy with the idea that "white" light is light of equal intensity of every wavelength in the visible spectrum. (To be specific in the present situation, if you had a filter that passed only frequencies in a 1 kHz band and connected the Random-Noise Generator to the filter input and an RMS voltmeter to the filter output, the voltmeter would read the same whether the filter passed frequencies between 1 and 2 kHz, between 3 and 4 kHz, and so on). The function generator frequency will be set right in the middle, frequency-wise, of the noise produced by the RNG, and some of this noise will make its way to the lock-in output.

Connect the various boxes and instruments according to the diagram of Fig. 8, except leave the Random-Noise Generator output disconnected from the Summing Amplifier, and the Summing Amplifier output disconnected from the lock-in input for now. Set the lock-in time constant to .01 seconds, the meter to 3 Volts and the meter switch on +. Turn on the various instruments and the power to the Summing Amplifier. Connect the output of the Summing Amplifier to the oscilloscope and adjust the amplitude and frequency of the digital function generator to get a sine wave of 300 mV peak to peak at 10 kHz. Now connect the output of the Summing Amplifier to the lock-in input and check that the the phase between the AUX OUT and MAIN OUT is set to 0°.

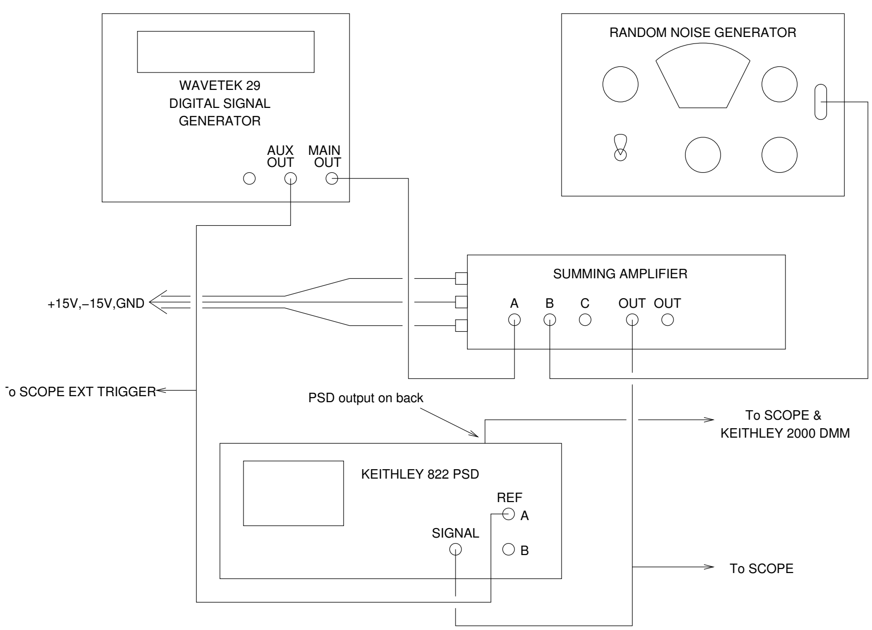

Figure 8: Setup for quantitative measurements of noise vs time-constant setting.

Set the Random Noise Generator to 20 kHz bandwidth, switch to HI, Multiply By knob to .1 and turn the amplitude knob **fully counterclockwise**. Connect the output of the RNG to the Summing Amplifier. Slowly turn the amplitude knob clockwise, watching the lock-in to make sure the INPUT OVERLOAD light does not come on. You should be able to turn the amplitude knob fully clockwise without having the light come on. As you turn the RNG amplitude knob up, you will see on the scope the 10 kHz sine wave become completely obscured by the noise from the RNG. When the knob is fully clockwise you should see a random signal with a peak to peak envelope of approximately 2 volts. With the time constant set to .01 seconds, the lock-in meter needle should be visibly fluctuating.

Configure the Keithley 2000 settings as follows:

- DC volts (push DCV button)
- 10 Volt range (push RANGE △ button)
- 5 digits (display should look like XX.XXX, digits button)
- Filter: Off (push FILTER button)
- Rate: MED (push RATE button)

The DMM reading should be fluctuating, just like the lock-in meter. Disconnect the RNG from the Summing Amplifier and note how the reading on the DMM becomes much more stable. The fluctuations riding on top of this stable DC voltage may be expressed in terms of an AC RMS voltage which is calculated by the formula

$$V_{\text{AC RMS}} = \left[\sum_{i=1}^{N} \frac{(V_i - V_{avg})^2}{N} \right]^{1/2} , \tag{4}$$

which is just the standard deviation of a set of N readings.

The Keithley DMM can make this calculation for you. Here are the steps for acquiring a set of readings, the average and the standard deviation:

1. Press the STORE key.
2. Use the ◄, ►, △, ▽ keys to set the number of readings (◄ and ► keys select the digit to be set; △ and ▽ keys set the value of the selected digit).
3. Press ENTER to initiate the acquisition of the readings. The $*$ annunciator on the front panel will be on while the readings are being taken and stored.

To view the average and standard deviation:

1. Press the RECALL key; the BUFFER annunciator will come on to indicate that stored readings are being displayed.
2. Press the ▽ key twice to bring the buffer to the "Average" location, and then press the ► key to display the average. To display the standard deviation, press the ◄ key, then the △ key and then the ► key.
3. To return to the regular display, push the EXIT key.

Several pages from the Keithley 2000 DMM manual are available on the class web site and in the lab with further information on these operations. You may want to practice these operations on the DMM before taking actual measurements.

Reconnect the RNG output to the Summing Amplifier. Set the rate on the DMM to FAST and collect several sample data sets with the lock-in time constant set to .01 seconds. Note that the AC RMS value changes from dataset to dataset. One can acquire several datasets with the same time constant and then take the mean as a more representative number for the AC RMS value. It is recommended that you acquire a minimum of five datasets of 1000 points each for each value of time constant. (Why not collect one data set of 5000 points? Try it!)

One consideration is the RATE setting on the DMM. This setting should be matched to the time constant of the lock-in, and it is recommended that the RATE be set to FAST for time constants $\leq .03$ seconds, and set to MED for time constants between .1 and 1.0 seconds. Specifications for specific RATE settings are included in the can be found in the Keithley 2000 DMM manual.

Exercise 10 *For time constants spanning at least .01 to 1 seconds, record in your lab notebook, as usual, each data set for the AC RMS value of the lock-in output.*

Plot the AC RMS voltage as a function of time constant. What kind of graph (linear, semi-log, log-log) is appropriate? Graph paper or a plotting program is available in the lab.

The relationship between the time constant and AC RMS on the lock-in output is clear. But how does one think about this formally? The key ideas needed to derive the noise vs. time-constant relationship have already been presented in this write-up. First, recall that the phase sensitive detector shifts all input signals downward in frequency by an amount equal to the reference frequency so that the signal of interest produces a DC (zero frequency) output, and signals differing in frequency from the reference frequency show up at the detector output at the difference frequency Δf, attenuated by the factor $1/\sqrt{1+(2\pi\Delta f\tau)^2}$. Second, recall that the Random Noise Generator output has an equal amount of power in any frequency interval. Since power is proportional to (voltage)^2, the square of the voltage measured over any frequency interval should be proportional to the extent of that interval. Combining these two ideas into a mathematical statement, the square of the RMS noise voltage at the filter output V^2_{RMSN} should be proportional to

$$V^2_{\mathrm{RMSN}} = \int \frac{V^2_{gen}(f)}{1+(2\pi\Delta f\tau)^2}\, d(\Delta f)\ , \tag{5}$$

where the limits of integration are over the bandwidth of the noise generator. The quantity $V^2_{gen}(f)$ has the units of V^2/Hz, that is, it is proportional to the power per unit bandwidth, and since our generator generates white noise, this quantity is a constant (on average), at least up to a frequency where $2\pi\Delta f\tau \gg 1$. Hence we can approximate the integral over the real bandwidth by assuming an infinite noise bandwidth, and let the limits of integration be from 0 to ∞. Thus the integral becomes

$$V^2_{\mathrm{RMSN}} \approx V^2_{gen} \int_0^\infty \frac{1}{1+(2\pi\Delta f\tau)^2}\, d(\Delta f) = \frac{V^2_{gen}}{4\tau}\ , \tag{6}$$

thus

$$V_{\mathrm{RMSN}} \propto \tau^{-1/2}\ . \tag{7}$$

Compare the relationship derived from your data to this result. (Some of the details have been glossed over here, but the basic result still stands. For a detailed treatment of the passband characteristics of the low pass filter see Ithaco Application Note IAN-35).

Let's consider the question of just how deeply the signal of interest can be buried in noise and still be detectable at the lock-in output. First, one has to decide what detectable means. For the situation at hand, let's say we require that at the lock-in output the signal amplitude (the DC, or average component of the output) be at least twice the noise amplitude, i.e., $V_{\mathrm{signal}}/V_{\mathrm{RMSN}} \geq 2$, and that 3 seconds is an acceptable time constant.

Exercise 11 *From the results of graphing V_{RMSN} vs. time constant, determine the noise amplitude at a 3 second time constant. How much can the signal input to the lock-in be reduced so that V_{signal} is about twice the noise at the output? Remarkable, isn't it? If you have time, test the prediction by dialing down the output of the signal generator appropriately.*

5 Final comments and some considerations when using a commercial lock-in amplifier

The Keithley 822 PSD is an old instrument and no longer available. More modern lock-in amplifiers are used in other experiments in the Advanced Labs, and these have features which make for easier and more accurate measurements:

- The 822 has only one output. To get the maximum response from the signal, the phase parameter ϕ needs to be set as close to zero as possible. Most newer lock-ins have two outputs, one which is proportional to $\cos\phi$, like the 822, and one that is proportional to $\sin\phi$. The $\sin\phi$ output is also called the "quadrature" output ("Quadrature" is an engineering term that is thrown around when discussing signals that are out of phase by 90°.) With these two outputs, typically called the X and Y outputs (sometimes called "real" and "imaginary") the experimenter can calculate the magnitude R ($R = \sqrt{X^2 + Y^2}$) and phase ϕ ($\phi = \arctan(Y/X)$) of the signal. In fact, many lock-ins will do this for you; you can choose R and ϕ to show up as voltages by the push of a button.
- Most modern lock-ins also allow the phase between the signal and reference to be adjusted at will. Some will even set the phase to 0° automatically.
- As noted, most modern lock-ins have input notch filters at 60 Hz and 120 Hz to reduce the interference that is so prevalent at these frequencies.
- Finally, most modern lock-ins have additional amplification stages at the input and the output, usually with selectable gains.

The astute reader may have realized that the kind of phase detection scheme implemented in the Keithley 822 would also be sensitive to odd harmonics of the reference frequency. The effect of the switch is to multiply the input by a square wave, and a square wave can be written in terms of a Fourier series by Constant $\times$ $[\sin(\omega t) + 1/3\sin(3\omega t) + 1/5\sin(5\omega t) + \ldots]$ (where $\omega = 2\pi f$. If one multiplies the signal by this series, one would see nonvanishing terms involving f, $3f$, $5f$, etc., if the input signal contains components at these frequencies. Indeed, one can see this effect by recalling that the 822 gave a greater output when presented with a square wave signal than a sine wave signal! Many lock-ins will work around this problem by additional circuitry that causes the effective multiplication to be with a sine wave reference; it is *not* a simple switch!

You could use a lock-in as a spectrum analyzer. A spectrum analyzer is a device which takes a composite signal consisting of individual signals at many different frequencies and tells you the amplitude of the individual signals as a function of frequency (i.e., it gives you the Fourier transform of the input signal). What would happen if you varied the reference frequency over a wide range? If you are curious about this you can try to look at the components of the square wave generated by the signal LED.

Another way to use a lock-in is to modulate, with a small sine wave, some parameter in an experiment, e.g., magnetic field, accelerating potential, temperature or physical displacement, and then see how the parameter of interest responds, by feeding its associated signal to the lock-in input. With this technique the *derivative* of the parameter of interest is sensed, not its absolute value. This technique is used in the Physics 432 lab in the Franck-Hertz experiment, and allows for very precise location of where the derivative of the signal goes to zero. For further reading on the modulation technique see Horowitz and Hill, pp. 1032-1034, and Preston and Dietz, pp. 368-373.

Prepared by J. Stoltenberg, D. Pengra, R. Van Dyck and O. Vilches

`lockin.tex -- Updated 12 October 2005`