Using a Lock-In Amplifier

Jared M. Maxson
Lehigh University Department of Physics
Physics 262
(Dated: February 14, 2008)

In this lab, we explore the principles of the use of a lock-in amplifier, using a He-Ne laser, light chopper, and a photodiode as the lock-in input. We explain and use the concepts of phase sensitive detection, dynamic reserve, and time constants to determine both the range of linearity of the photodiode in use, as well as the frequency response of the photodiode. Furthermore, we determine whether the frequency response of the photodiode is independent of intensity. Photodiode output voltage data as measured on the lock-in (varying both frequency and intensity) are presented.

I. INTRODUCTION AND MOTIVATION

For nearly all modern physical experimentation, the reduction of the signal-to-noise ratio is of paramount importance. In laser spectroscopic experiments, for example, the desired signal can be several orders of magnitude less than the signal from myriad noise sources, such as stray light, dark current, or inherent device noise. In most cases, signal filtration and subsequent amplification can still yield an unsatisfactory ratio of signal to noise, given the random distribution of noise signal frequencies and intensities. The lock-in amplifier is a device created to surmount this problem, by modulating the input signal by a reference signal ω_R (created by a light chopper and laser, in our case), and upon signal detection, measuring only the voltage input modulated by ω_R.

II. LOCK-IN THEORY

A. General Overview

The input for the lock-in amplifier for this experiment was a He-Ne laser run through a light chopper, incident upon a photodiode. In this case, the light chopper was attached to the lock-in to provide the reference frequency, but in general any voltage source can be used, with any modulation source (a function generator, for example). On the bottom surface of the light chopper is an LED and photodetector pair, sends the square wave ”chopping” signal to the lock in as reference. The lock-in then converts this square wave input into sinusoidal form, for use in processing the input signal.

Noise sources, in general, are continuous in their frequency distribution; however, we may roughly approximate the input signal as a sum of the desired signal and discrete noise sources and frequency:

$$V_{sig} = V_d \cos(\omega_R t + \phi_d) + \sum_n V_n \cos(\omega_n t + \phi_n) \tag{1}$$

where there are n sources of noise, with different intensities, frequencies and phases, and the input signal has the reference frequency. It should be noted that depending on the chopper blade and beam geometry, the time dependence of V_d will not be perfectly sinusoidal, but in this illustrative model we assume perfect cosine dependence. Upon detection by the lock in, V_{sig} is then multiplied by a known reference signal in the form of $V_{ref}\cos(\omega_R + \phi_{ref})$, which, using a trigonometric identity, yields:

$$V_{mult} = \frac{1}{2}V_d V_{ref}\cos((\omega_R + \omega_R)t + \phi_R + \phi_d) + \frac{1}{2}V_d V_{ref}\cos(\phi_R - \phi_d) + V_{ref}\cos(\omega_R + \phi_{ref})\sum_n V_n \cos(\omega_n t + \phi_n) \tag{2}$$

This multiplied signal, V_{mult}, is then passed through a low-pass filter, such that only the DC component of the signal is retained. Thus, the only remaining signal is the second term above, plus a possible noise term having frequency ω_0, from the expansion of the sum.

$$V_{filt} = \frac{1}{2}V_d V_{ref}\cos(\phi_R - \phi_d) + \frac{1}{2}V_N V_{ref}\cos(\phi_R - \phi_N) \tag{3}$$

It should be noted that second term need not necessarily appear: if there is no noise signal having the reference frequency, this signal will be filtered out, and will not enter into V_{filt}. Thus, using this method in an ideal case, all noise except that possessing ω_R, can be eliminated. Also, V_{filt} is dependent on the cosine of the reference and signal phase, $\cos(\phi_R - \phi_d)$. The reference phase may be adjusted so as to maximize V_{filt}, corresponding to phase matching. Thus, this method of detection is termed "Phase-Sensitive Detection". After the initial filtration, the signal is then amplified, to facilitate measurement. If the reference and signal phase are matched, then an accurate measurement of the V_d (rms) is then displayed.

B. Time Constants, Dynamic Reserve, and Experiment Goals

Two concepts that play an integral role in the operation of a lock-in amplifier are the time constant and dynamic reserve. Every lock-in has at least one low-pass filter. The time constant display on the front panel is the time constant of this filter, and is a controllable parameter. For very low time constants, say 10 μs, the lock-in display voltage responds quickly to real changes in the input signal, but the voltage reading is rarely stable enough for proper measurement. For high time constants on the order of seconds, the voltage reading is quite stable, but can take many time constants for the reading to respond to a real change in input. The time constant should be adjusted to provide an optimal balance of responsiveness and stability, depending on if you are changing input or taking measurements.

The dynamic reserve is also a controllable parameter of the lock-in amplifier. It is defined as the largest tolerable ratio of noise to signal. Often the actual value of dynamic reserve is not specifically stated on the panel, but can be adjusted from high to low, for instance. As a rule, the dynamic reserve should be kept as low as possible without overloading, in that for analog lock-in amplifiers a high dynamic reserve can affect the linearity of the phase sensitive detection, and for digital lock-ins, high dynamic reserve produces more output noise from the analog to digital converter. However, for high-noise experimentation, an increased dynamic reserve can enable measurement of a signal that a lower setting would read as an overload.

In this experiment, we are first to determine whether or not our photodiode is operating in the linear range. If the photodiode is operating in the linear range, a fractional decrease in incident light intensity will correspond to the same fractional decrease of output voltage. If the incident intensity is too strong, a decrease in intensity may produce a disproportionately smaller decrease in the voltage; this is termed "saturation" of the photodiode. Next, we are to determine the frequency response of the photodiode. The output of the photodiode is dependent on chopper frequency in that the photodiode takes a finite time to respond to incident light, and if the chopper frequency is too high, not all of the light intensity is capable of being read before the next blade of the chopper covers the beam, and so on. Thus, we expect for low chopper frequencies a steady lock-in reading, but at a certain frequency, the time the beam in steadily on the face of the photodiode is shorter than the response time of the photodiode, and thus output voltages should decrease. Finally, we are to determine whether the frequency response of the photodiode is independent of laser intensity.

III. PROCEDURE

The experiment was conducted as follows:

A. Setup

Both the laser and the photodiode were attached securely to a rail, allowing sufficient room for accessories to be attached in between them. The laser beam was centered on the photodiode surface, and the chopper was placed so that it's blades could obstruct the beam. The photodiode output was connected to the A input of the lock-in amplifier, and the chopper reference output was attached to lock-in reference input. A neutral density filter-holder was placed in the path of the beam, and a piece of frosted scotch tape, which acts as a light diffuser, was placed on the surface of the photodiode cover, in the path of the beam. In general during measurement, a time constant of 300 μs and a normal dynamic reserve was used. Also, after every equipment or chopper frequency change, and at the start of every session, the the phase was adjusted to yield zero signal, and then shifted by 90 degrees to yield maximal signal.

B. Photodiode Linearity Determination

To determine whether the photodiode was operating in the linear range, a number of neutral density filters and a power meter were employed. First, an NG 9 filter was placed in the path of the beam, and using the power meter, the power output of the filtered beam was determined. Then the meter was removed, and the beam allowed to hit the photodiode, whereafter the voltage was read off of the lock-in. This was repeated for a number of filter configurations: NG9 + NG5, NG9 + NG5 + NG5, NG9 + NG5 + NG5 + NG11, NG9 + NG9. The power and voltage output of each configuration was measured. In each successive run, the power was reduced such that the power in the final run was just barely discernable on the power meter. A linear range was determined, and a filter corresponding to the largest linear configuration was chosen. In this case, only the NG9 was used.

C. Frequency Response of the Photodiode

To determine the frequency response of the photodiode, the initial frequency of the chopper was taken down to the lowest stable chopper frequency. At chopper frequencies below this point (in this case, a frequency of approximately 50 Hz), the imperfections of the chopper blade cause the reference signal to "unlock". In this range the photodiode output voltage was sampled at intervals of 10 or 20 Hz. Above 200 Hz, this interval was increased steadily, as the frequency roll off had already become apparent. This data was plotted as output voltage as a function of frequency.

D. Frequency Response and Intensity

To determine whether the frequency response of the photodiode was independent of laser intensity, we utilized a different (polarized) laser, an SR830 lock-in amplifier, power meter, and a Labview (version 7.1) program to both control and take data from the SR830. Frequency dependence was taken at three different power levels: 0.063 mW, 0.498 mW, and 1.110 mW. These values were chosen because the upper value was believed to saturate the photodiode, whereas the first has been demonstrated to be in the linear range. The power was adjusted by changing the angle on the polarizer, and the power was read on the power meter. The power meter was then removed, and with the laser incident on the photodiode (outputting to the lock-in), the Labview program began to take frequency response data. The program was written to take 5 data points at 15 different frequencies between 25 and 1000 Hz. For these measurements, the time constant was set at 1 second. At a given frequency, output voltage data was taken once every two seconds, to allow time for the display to change. A delay of 10 seconds between different frequency readings was built in to provide time for the chopper to maintain the desired frequency. Furthermore, the frequency was scanned from lowest to highest, to allow the chopper to build up to 1000 Hz gradually.

IV. EXPERIMENTAL RESULTS

A. Photodiode Linearity

A table of different light intensities and photodiode output are shown in FIG 1. There are two sets of measurements tabulated: those numbered, and those given a letter. The numbered runs were done using a probably malfunctioning 6-bladed chopper, which produced strangely low lock-in voltages yet reasonable relative changes in lock-in voltage with changing intensity. In FIG 2, we see that the RUN 2 (marked in red) output voltage seems show that this power range is neither linear nor saturated. In this run, the power was reduced by half, but the output voltage was reduced by well over half. Saturation would be indicated by a reduction by less than half, and linearity would be indicated by nearly half. Thus, the validity of this measurement is doubted. Furthermore, using 30-blade chopper and different chopper frequency (measurements A and B), we find that in that range, the behavior is indeed linear: halving the power halves the output voltage. Thus, it is apparent from Fig 2 and runs A and B that all measurements charted (but RUN 2) display approximately linear behavior, and a filter of NG9 was chosen for further experimentation.

B. Photodiode Frequency Response

A plot of the photodiode voltage versus chopper frequency is found in Figure 3. A smooth, decreasing curve is defined by the data points; however, as the curve's profile suggests, the minimum chopper frequency of 50 Hz was

RUN	FILTERS(NG#)	POWER (mW)	PD VOLTAGE (mV)
1	9	1.36	0.509
2	9,5	0.68	0.177
3	9, 5, 5	0.34	0.129
4	9, 5, 5, 11 (2 mm)	0.16	0.0518
5	9, 9	0.05	0.0178
A	9	1.35	9.21
B	9, 5	0.68	4.53

FIG. 1: Photodiode output voltage and laser power data. Numbered runs were made with a possibly malfunctioning 6-blade chopper, which accounts for the uncharacteristically low lock-in voltages. The lettered runs were done at a different frequency with a normally functioning 30-blade chopper.

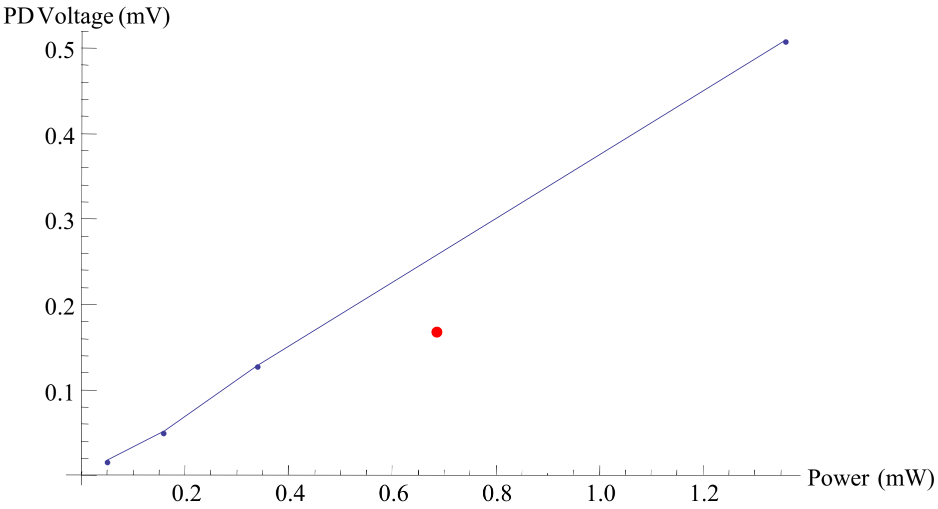

FIG. 2: Photodiode output plotted as a function of laser power, using the 6-bladed chopper. From this plot as well as runs A and B, we see that RUN 2 (marked in red) is thus most probably an erroneous measurement.

not slow enough to display the beginning of the voltage roll off. From this graph we may only determine that the region of constant output voltage as a function of frequency is below 50 Hz, and that using a frequency above 50 Hz would result in an inaccurate measurement of the input signal, V_d.

C. Frequency Response and Intensity

In FIG 4, the frequency response of the photodiode is plotted at three different laser intensities, 0.063 mW (red), 0.498 mW (purple), and 1.110 mW (blue), with corresponding error bars (standard deviation of the mean) at each data point. FIG 5 is logarithmic plot of the same data (omitting error bars), allowing a more qualitative comparison between the curves. The measured error for many runs was zero or nearly zero, which is why the error bars in FIG 4 are often barely visible. This suggests a strong lack of independence between measurements, and thus that the 2 second wait between measurements of the same frequency is too short to provide complete independence. However, independence between measurements of different frequencies is relatively certain given the 10 second wait between them. In FIG 5 we see that for all intensities, though the amplitude of the response is naturally different, the curve profiles are nearly identical. Thus, in the frequency range of 25 to 1000 Hz, we find that the frequency response (except of course for overall amplitude), is independent of intensity. Furthermore, this data also shows that for the photodiode in use, the rolloff frequency is even smaller than 25 Hz, as the beginning of this rolloff is not noted in the used frequency range.

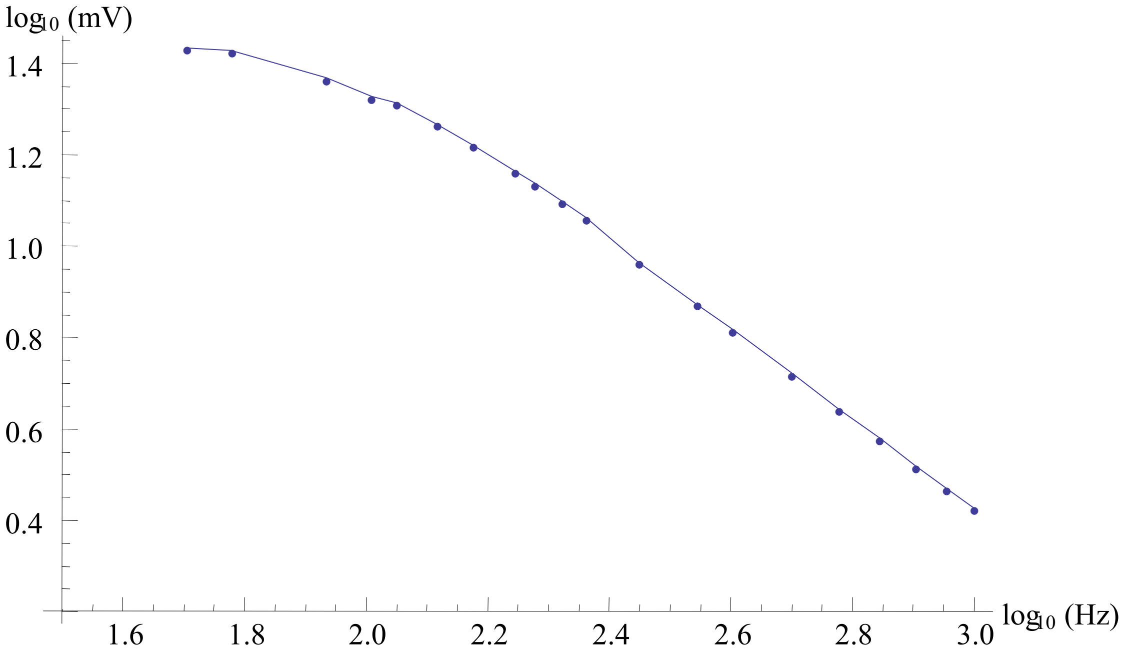

FIG. 3: A plot of the logarithm of photodiode output voltage as a function of the logarithm of chopper frequency, using the 30-blade chopper. The chopper would unlock at frequencies below 50 Hz, and thus the first three points were slightly unstable.

V. CONCLUSIONS

In this experiment, I was able to learn the function and many of the common features of lock-in amplifiers, and to successfully employ them to determine the frequency response of a photodiode. As a preliminary step, I was also able to determine a portion of the linear range of the photodiode in use. The portion of the linear range that was determined was for incident laser light of 0.05-0.136 mW sent though a scotch tape diffuser. Operating at 0.136 mW to maximize the photodiode output, I was able to determine the frequency response of the photodiode in the frequency range allowed by the chopper. From these data, a smooth frequency dependence was determined, and from these data it is apparent that the frequency independent region lies below the lowest chopper frequency allowed. It has also been demonstrated that the frequency response of the photodiode is quite independent of intensity for intensities between 0.063-1.110 mW. This apparent independence may be due to the fact that the rolloff frequencies were not able to be observed with the available choppers, and it is in the rolloff region where intensity dependent effects may be observed. Thus, we may only conclude that the frequency dependence is independent of intensity between 25-1000 Hz. Furthermore, a chopper capable of producing frequencies below 25 Hz would be required to measure the actual signal voltage and the frequency roll off; it has been demonstrated that chopper frequencies above that value produce output voltages less than would be the actual value.

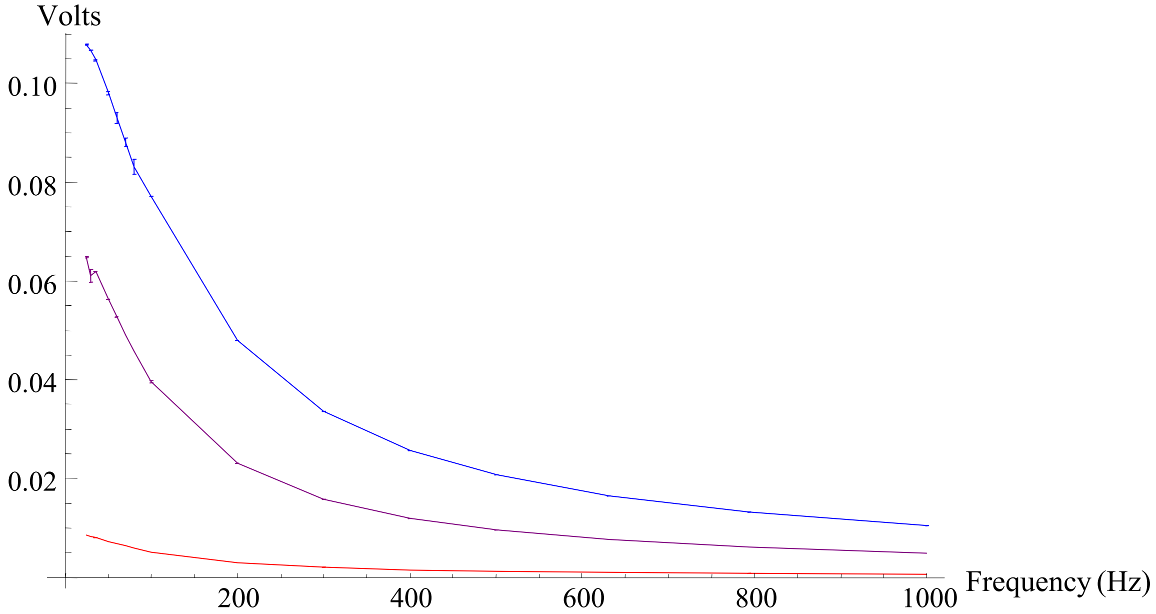

FIG. 4: A plot of the output voltage as a function of frequency for different intensities, with error bars for the standard deviation of the mean for each data point. The three different laser intensities are 0.063 mW (red), 0.498 mW (purple), and 1.110 mW (blue). Given many nearly zero error bars, the independence of measurements at a given frequency is doubted.

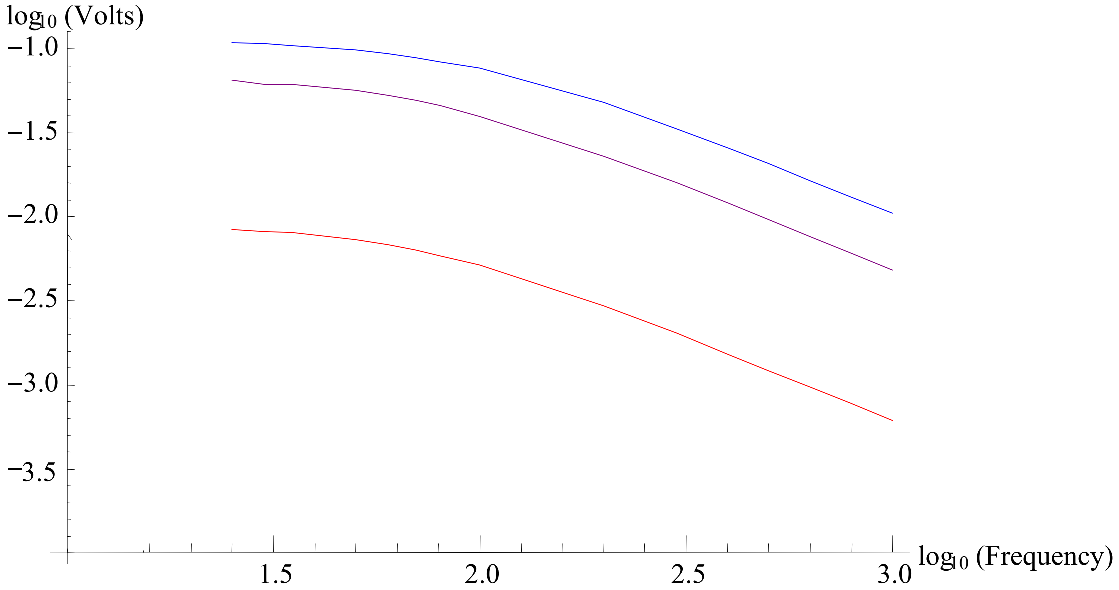

FIG. 5: A log-log plot of the data from FIG 4. From this plot it is apparent that for the given frequency range, the frequency response is almost entirely independent of intensity (between 0.063-1.110mW).