HOMODYNE DETECTION IN A LASER LOCKING SYSTEM

by

Aaron Bennett

A senior thesis submitted to the faculty of

Brigham Young University

in partial fulfillment of the requirements for the degree of

Bachelor of Science

Department of Physics and Astronomy

Brigham Young University

August 2010

Copyright © 2010 Aaron Bennett

All Rights Reserved

BRIGHAM YOUNG UNIVERSITY

DEPARTMENT APPROVAL

of a senior thesis submitted by

Aaron Bennett

This thesis has been reviewed by the research advisor, research coordinator, and department chair and has been found to be satisfactory.

Date	Dallin Durfee, Advisor
Date	Eric Hintz, Research Coordinator
Date	Ross L. Spencer, Chair

ABSTRACT

HOMODYNE DETECTION IN A LASER LOCKING SYSTEM

Aaron Bennett

Department of Physics and Astronomy

Bachelor of Science

I discuss a high speed, low noise homodyne photo-detector. This detector will be used to better implement laser locking techniques such as the Pound Drever Hall method or saturated absorption with lock-in amplification. I present a basic explanation of these methods and their benefits. I discuss aspects of the detector which allow it to operate with low noise over a high bandwidth.

ACKNOWLEDGMENTS

Most of all, I would like to acknowledge my wife whose constant support and unwavering loyalty makes all the difference in my life. I would also like to acknowledge my advisor who is possibly the only professor at BYU who would have the patience to advise me for over two years.

Contents

List of Figures

Chapter 1

Introduction

In Professor Durfee's lab, we are building a Ramsey-Borde atom interferometer and an ion interferometer. In these interferometers, lasers are used to drive narrow transitions in calcium and strontium atoms. Driving these transitions requires ultra-stable lasers locked to a frequency reference such as a high finesse optical cavity or an atomic absorption line. To keep a laser locked despite fluctuations in the laser, continuous adjustments to the laser cavity and the current driving the laser are necessary. The homodyne detector is the primary element in a feedback loop that dictates these adjustments. Without this feedback, the laser will eventually drift off resonance, rendering the laser useless in our experiments. This thesis discusses the homodyne detector that I designed and built.

In this chapter, I discuss the basics of lock-in detection, including two different ways that it can be used to lock a laser to a frequency reference.

1.1 Lock-In Detection

Lock-in detection is a method used to extract information about a signal even in situations where there is a low signal to noise ratio. To understand this principle, let's examine a standard lock-in detection experiment. Say that a photodiode is placed across a room from an LED and we want to be able to tell if something is preventing the light from the LED from getting to the photodiode. If we simply turn the LED on and leave it on, the photodiode would detect both the background noise in the room and the light from the LED, making it very difficult to determine whether the light from the LED makes it to the photodiode (especially if we have a low signal to noise ratio). Let's say that instead of simply turning the LED on and leaving it on, we modulate the voltage powering the LED with a square wave such that it switches back and forth between on and off. Because the background noise in the room does not switch on and off like the LED, the signal from the photodiode is much more useful now. We send the on/off square wave to the photodiode (to identify which portions of the photodiode signal occurred when the the LED was on and when it was off) and subtract the photodiode signal when the LED was off away from the photodiode signal when the LED was on. We perform this process many times and average the results. This effectively averages away the noise (because the noise is random) while not averaging away the signal (which is not random). If a significant signal survives this averaging, we know that the light from the LED is not blocked on its way to the photodiode.

In the previous example, we were able to extract a signal by modulating with a square wave and removing all the unmodulated portion. In practice, to perform an experiment like this, we usually don't modulate with a square wave; rather, we modulate with a sine wave. Furthermore, we don't actually subtract out the unmodulated portion. Instead, we extract information about the signal by mixing it with another signal. The advantages of doing it this way are best illustrated by example. Let's return to the previous example where we are determining if light from an LED is unblocked on its way to a photodiode. However, this time, instead on turning the LED on and off, we modulate the amplitude of the LED light with a sine wave generated by a local oscillator. The equation for the local oscillator (y_{LO}) is:

$$y_{LO} = A_{LO} sin(\omega_{LO} t - \phi_{LO})$$

If the LED light is unblocked, the photodiode output has three main components. First, it has a DC offset because the light intensity oscillates about a positive value. Second, on top of the DC bias, there is a signal oscillating at the modulation frequency (ω_{LO}) that is in phase with the modulation and has some amplitude A_m. Third, on top of these other two components, the light in the room adds noise at various frequencies that is larger in amplitude than A_m. This noise term with all its frequency components will be denoted by the letter n. Mathematically then, the photodiode output (y_{pd}) when the light is unblocked is:

$$y_{pd} = A_{pd}(1 + A_m sin(\omega_{LO} t - \phi_{LO})) + n.$$

On the other hand, when the light is blocked, the DC offset and modulated light will not be present, leaving us with:

$$y_{pd} = n.$$

To discover whether the light is blocked, we multiply y_{pd} by y_{LO} with a mixer. Let's examine the case where the LED light is unblocked. In this case, the mixer output (M_{out}) is:

$$M_{out} = y_{pd} y_{LO} = A_{pd} A_{LO} sin(\omega_{LO} t - \phi_{LO})[(1 + A_m sin(\omega_{LO} t - \phi_{LO})) + n]$$

which simplifies to:

$$M_{out} = A_{LO}sin(\omega_{LO}t - \phi_{LO})(A_{pd} + n) + A_{pd}A_{LO}A_m sin(\omega_{LO}t - \phi_{LO})sin(\omega_{LO}t - \phi_{LO}).$$

I intentionally did not write the sine functions in the last terms as one sine squared function to illustrate that we can use a product to sum identity to rewrite M_{out} in this way:

$$M_{out} = A_{LO}sin(\omega_{LO}t - \phi_{LO})(A_{pd} + n)$$
$$+\frac{A_{pd}A_{LO}A_m}{2}[cos((\omega_{LO}t - \phi_{LO}) - (\omega_{LO}t - \phi_{LO})) + cos(2(\omega_{LO}t - \phi_{LO}))]$$

which simplifies to:

$$M_{out} = A_{LO}sin(\omega_{LO}t - \phi_{LO})(A_{pd} + n) + \frac{A_{pd}A_{LO}A_m}{2}[1 - cos(2(\omega_{LO}t - \phi_{LO}))].$$

This result highlights three important concepts. First, the second term of M_{out} has a DC offset because the two signals that were mixed together were at the same frequency. Second, the magnitude of the DC offset is proportional to the amplitude of the modulation on the photodiode signal. Finally, the relative phase of the two signals defines the sign of the offset and also affects the magnitude of the offset. In this example, the signals were exactly in phase so the product-to-sum identity left us taking the cosine of zero. If this phase difference is between 0 and 90 degrees, the offset is positive; if the phase difference is between 90 and 180 degrees, the offset is negative (except at 90 degrees where the offset is zero). Furthermore, the closer to 90 degrees the phase difference between the two signals is, the smaller the magnitude of the DC offset.

This information about the second term from M_{out} helps us see that the first term in M_{out} will only have a DC offset if part of the noise oscillates at the modulation frequency.

The DC component of the mixer output is the most valuable piece of the signal because it tells if any of the light hitting the photodiode is oscillating at frequency ω_{LO}. As long as we chose a high modulation frequency, we can be fairly certain that almost none of the noise in the room oscillates near this frequency. Therefore, the mixer output will have a DC bias if the LED light is unblocked and essentially no DC bias if it is blocked. Since all we really care about in this instance is the DC portion of the mixer output and all other portions may be distracting, we often low-pass filter the mixer output to isolate the portion of the signal that contains the pertinent information.

One of the biggest advantages to using this method is that all the terms arising from noise at frequencies not close to the modulation frequency become irrelevant because they are filtered out. Therefore, the final output after the low-pass filter can have a high signal to noise ratio even if y_{pd} does not. This is extremely useful because we can make a very quiet output signal by simply choosing a modulation frequency where there is only a small amount of noise. It is a common practice to choose a fast modulation frequency because many physical systems tend have less noise at higher frequencies.

This process of modulating a signal and mixing it with a signal at the same frequency is known as lock-in detection because we are able to detect (or lock onto) the portion of the signal that is at the modulation frequency while ignoring all of the other contributions. It is also called homodyne detection because it involves detecting a signal by mixing it with a signal at the same frequency.

1.2 Implementation

While this is not a paper about laser locking techniques, it is important to at least briefly discuss the way that my detector will be used in these methods. There are two types of locking systems that I will discuss in this section: locking to a reference cavity and locking to an atomic absorption line. While these methods are different in nature, they both use the principles of lock-in detection.

1.2.1 Locking to a Cavity

An optical cavity is essentially two semi-transparent mirrors placed some distance apart from each other. When light is incident on a cavity, a maximum amount of light transmits when the cavity's length is an integer number of half wavelengths of the light. The frequencies associated with these wavelengths are known as resonant frequencies. When the laser drifts off resonance, some light still transmits through the cavity (but not as much as when the light is resonant). The transmitted light intensity continues to decrease the further the frequency drifts from resonance. This behavior can be seen in Fig. 1.1.

Because the amount of light transmitted through the cavity is dependent on the laser's frequency, optical cavities can be used to lock a laser's frequency. One method for doing so is shown in Fig. 1.2. Here, light from a laser is directed through an optical cavity onto a photodiode on the opposite side of the cavity. When the light is on resonance, transmission through the cavity is maximized, causing a maximum photodiode output. When the laser drifts off resonance, transmission through the cavity decreases, causing the photodiode output to decrease. One could assume that an effective way to lock to the cavity's resonance is to send the photodiode output through some feedback loop that will adjust the laser's frequency to maximize the

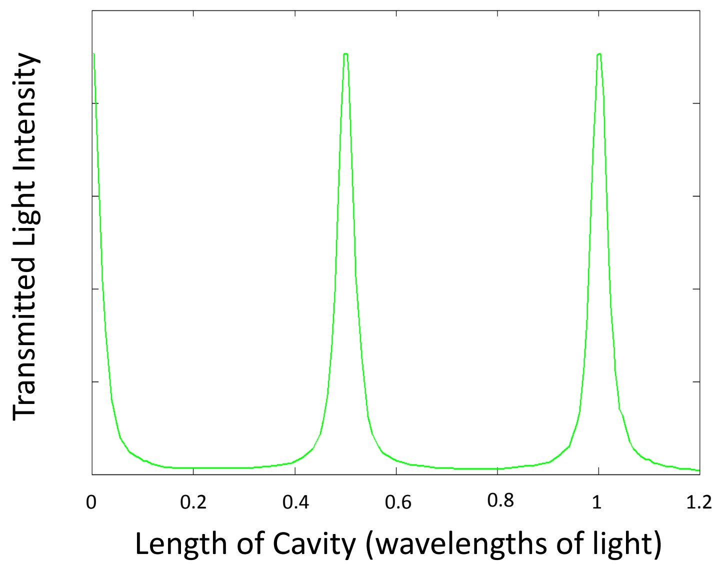

Figure 1.1 The transmission through an optical cavity of varying length. When the length of the cavity is an integer number of half wavelengths of light, the transmission through the cavity is maximum. This figure was reproduced from [1].

photodiode's output. Unfortunately, locking this way does not work because the curve is symmetric about the resonance. If the frequency drifts, it is impossible to tell which direction it drifted because less light transmits whether the frequency drifts high or low.

We could overcome this problem by using the same setup but instead locking to a point that is on the side of the resonance curve. Then, as long as we know which side of the curve we are locking to, it is easy to determine the direction of frequency drifts just by looking at the photodiode's output. While this does work, there are some limitations when locking this way. First, the lock point is dependent on light intensity. For example, if the light intensity drops while the laser's frequency remains constant, the photodiode output will also drop, triggering a change in frequency when no change is necessary. Second, locking to the side of the transmission resonance curve

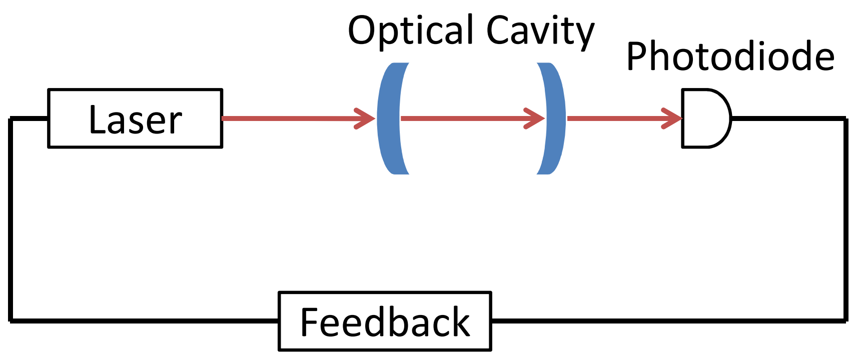

Figure 1.2 An example of a poor method for locking a laser to a cavity.

results in a slow lock because it will take some time for the light in the cavity to ring down if the frequency drifts. Furthermore, if the goal is getting a lot of light into the cavity (perhaps for a frequency doubling application) locking to the side of the resonance curve greatly reduces the amount of light coupled to the cavity.

Fortunately, this is not the only way to lock a laser to a reference cavity. A far superior method, known as the Pound-Drever-Hall technique, utilizes the principles of lock-in detection to create a fast, precise lock. The basic setup for this method is shown in Fig. 1.3. Directly out of the laser, the beam is split and most of the light is directed to the experiment. The remaining light is phase modulated by an electro-optical modulator (EOM) at a frequency dictated by a radio frequency (RF) local oscillator. After modulation, the light is directed to the optical cavity. Instead of examining the transmission through the cavity as in the previous example, here we look at the reflections off of the cavity with a photodiode. The resulting signal from the photodiode is amplified, filtered, mixed with a signal at the modulation frequency, and low-pass filtered. The signal after the filter is the error signal that is fed into

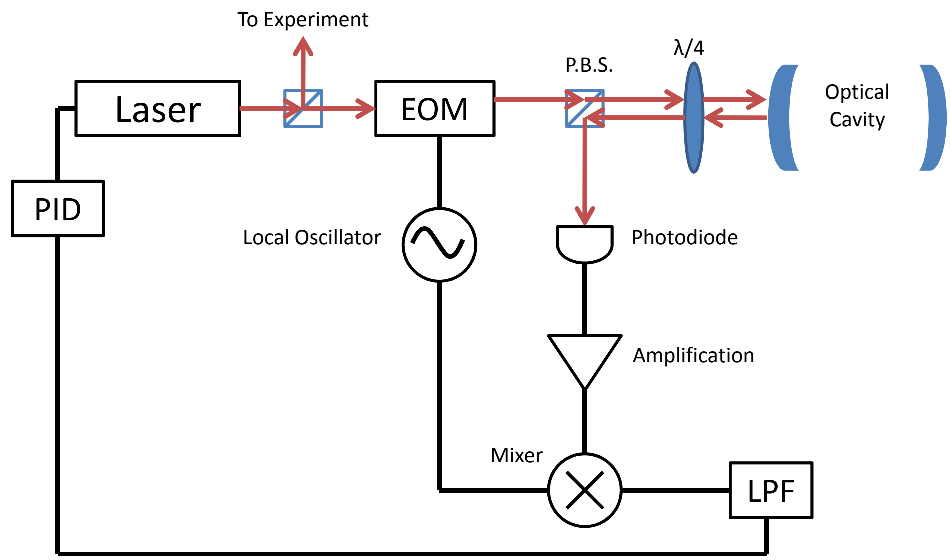

Figure 1.3 The Pound Drever Hall method for locking a laser.

a mechanism that automatically adjusts the laser to keep it on resonance with the cavity.

Because this thesis is not focused on laser locking techniques, I forgo an in depth quantitative derivation of the error signal and instead provide a qualitative explanation with an occasional quantitative portion. Before doing so, let's assume that the modulation frequency is smaller than the linewidth of the cavity. Using similar techniques, one can also examine situations where the modulation frequency is larger than the linewidth of the cavity.

To understand the error signal, we need to start by examining the effect of phase modulation on the laser. Before modulation, the electric field of the laser can be written as:

$$E = E_0 e^{i\omega t}.$$

After passing through the EOM, the phase is now modulated, such that:

$$E = E_0\, e^{(i\omega t + Asin(\Omega t))}$$

(where Ω is the modulation frequency). A Fourier transform of this modulated signal reveals that, in frequency space, this phase modulation adds frequency sidebands to the carrier frequency of the laser (with the side bands located at $\omega - \Omega$ and $\omega + \Omega$). This result becomes clearer when we approximate the modulated electric field in the following way [2]:

$$E \approx E_0(J_0(A) + 2iJ_1(A)sin(\Omega t))e^{i\omega t} = E_0(J_0(A)e^{i\omega t} + J_1(A)e^{i(\omega+\Omega)t} - J_1(A)e^{i(\omega-\Omega)t}).$$

In the above equation for E, it is important to note that modulating the phase has no effect on the amplitude.

Because the phase modulation adds frequency sidebands to the carrier frequency of the light, there is beating between the carrier and the side bands at the modulation frequency. The beat-note from the carrier frequency beating with the lower sideband (the sideband with the lower frequency) is out of phase with the beat note from the carrier beating with the upper side band (even though they are out of phase, they still beat at the same frequency). Now, treating the reflected signal as the sum of the incident light directly reflected and the light leaking from the cavity, let's qualitatively look at the reflected light. When on resonance with the cavity, equal amounts of both sidebands reflect. Consequently, there is no beating between the carrier and sidebands when on resonance (because the two beats have the same magnitude and are summed together out of phase). Off resonance though, different amounts of each sideband reflect. Because one of the beat notes is larger than the other, this time they don't cancel. Furthermore, the beat either has or does not have a π phase shift (depending on which beat-note is larger). As the laser gets farther from resonance, more of the carrier reflects off of the cavity. This causes the beat note to grow in amplitude. Therefore, the photodiode detecting this light produces a signal oscillating at the modulation frequency whose amplitude tells how far the carrier is from the cavity's resonance and whose phase tells which direction from resonance the carrier drifted. Consequently, when we mix this signal from the photodiode with the signal from the local oscillator and low-pass filter the output, the result is a DC signal whose amplitude tells how far the carrier frequency is from resonance (because the amplitude of the signal after the mixer will be proportional to the amplitude of the signal from the photodiode) and whose sign (which is dependent on the phase of the signal from the photodiode) tells which way from resonance the laser drifted. This filtered output is the error signal that dictates the adjustments needed to bring the laser back on resonance with the cavity.

This method for locking lasers is far superior to the method previously discussed for several reasons. First, we do not have to wait for the light in the cavity to ring down because we use reflections off the cavity. Second, drifts in intensity do not matter when on resonance because we are locking to a zero point (i.e. if the intensity drifts but the frequency doesn't, we will still have a 0 V DC signal when on resonance). Then, when the frequency does drift, the signal used to lock the laser tells both the direction and amplitude of the drift. Finally, because this method implicitly uses lock-in detection, the error signal should have a high signal-to-noise ratio. Overall, this method provides a very fast, ultra stable lock.

1.2.2 Locking to an Atomic Absorption Line

In many of the experiments performed in atomic, molecular, and optical physics, a laser is used to drive a transition in an atom to an excited state. For these experiments to work, the laser must be at and stay at a specific frequency (called the atomic resonance frequency) for the duration of the experiment. To stay locked on resonance, we need to be able to determine the relationship between the laser frequency and the resonance frequency and then make proper adjustments.

An excellent way to determine this relationship is to create a vapor cloud of the same type of atoms used in the experiment (e.g., our ion interferometer uses strontium atoms, so the vapor cloud should also be strontium atoms), direct a portion of the laser through this cloud, and then look at the intensity of the light that transmits through the cloud with a photodiode. Interpreting the output of the photodiode requires an understanding of the how the light intensity transmitted through the cloud is dependent on the frequency of the light. This relationship is addressed briefly below with a more thorough explanation included in Appendix A.

In the atoms that make up the vapor, the energy difference between the ground and excited state may be denoted by ΔE. The frequency of light needed to transition the atom between the ground and excited states is related to ΔE through Planck's law: $\Delta E = \hbar\omega_0$ (where $\omega_0 = 2\pi\nu_0$, or $\nu_0 = \frac{\Delta E}{2\pi\hbar}$). If a laser is at the atomic resonance frequency (ν_0), the atom will have a maximum probability of absorbing some of the laser's intensity as it is struck by the beam. As the laser frequency drifts away from ν_0, the probability that the atom will absorb some of the light intensity decreases. Appendix A discusses in depth how the laser's intensity as it propogates through a cloud of atoms is affected by it's frequency. Here, it is sufficient to say that if light at frequency ν propagates a distance x through a cloud of atoms with an atomic resonance frequency ν_0, a resonance linewidth Γ, an effective cross-sectional area of σ, and a number density n_0, the output intensity (I) will be related to the initial intensity (I_0) as follows:

$$I = I_0 e^{\sigma n_0 x}$$

where:

$$\sigma = \sigma_0 L(P_0 - P_1)$$

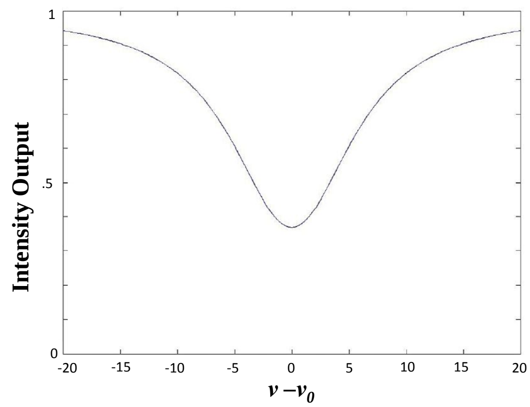

Figure 1.4 Graph depicting reduction in light intensity for a laser traveling through a vapor cloud with the output intensity normalized to I_0. The x-axis has arbitrary frequency units.

and:

$$L = \frac{1}{1 + \frac{4(\nu - \nu_0)^2}{\Gamma^2}}.$$

In these expressions, P_0 and P_1 are the probability that the atoms will be in either the ground or excited state (respectively) and σ_0 relates to the cross sectional area of the atom when directly on resonance with the transition. A graph relating the laser frequency to the intensity of light transmitted through the cell (not absorbed by atoms) is shown in Fig. 1.4.

Perhaps the easiest way to lock to the atomic resonance may appear to be having a feedback loop adjust the laser frequency to minimize the amount of light on the photodiode placed after the vapor cloud. However, this does not work because the resonance curve is symmetric about the resonance (making it impossible to determine the direction of frequency drifts).

We could overcome this particular problem by using the same setup but this time choosing a lock point that is on the side of the curve. This makes it is easy to determine the direction of frequency drifts (e.g. if we are locking to a point on the right side of the curve, we know that the frequency went up if the photodiode output increases and we know that the frequency went down if the photodiode output decreases). The problem with this method is that the lock point is dependent on the density of atoms in the vapor cloud and the initial laser intensity. For example, if the density of atoms in the laser's path decreases while its frequency remains constant, the photodiode output will increase, triggering a change in frequency when no change is necessary.

Overcoming these problems requires a somewhat unintuitive solution that draws on the power of lock-in detection. In this solution, the beam is split directly out of the laser and most of the light is directed to the experiment. The remaining light is frequency modulated with an acousto-optic modulator (AOM) at a frequency determined by a local oscillator. This modulation does not affect the main beam sent to the experiment. After passing through the AOM, we direct this light through a vapor cloud onto a photodiode. Because the light intensity transmitted through the vapor cell is a function of the laser's frequency, modulating the laser frequency causes the photodiode output to oscillate. This output is amplified, mixed with the signal used to modulate the laser's frequency, and low-pass filtered (removing any high frequency components). This filtered signal is a perfect error signal because it contains information about both the magnitude and direction of any frequency shifts.

Let's examine how this error signal comes to contain this valuable information. To do so, I will examine four different cases. Those four cases are when the laser's frequency is: (1) on resonance, (2) well above resonance, (3) well below resonance, and (4) only slightly above resonance. The graphs in Fig. 1.5 should help in this process by tracking three different pertinent values for each case as time evolves. In these graphs, ν and ν_0 have the same meaning as they did in the discussion of Beer's Law.

Fig. 1.5(a) shows this process when the laser frequency oscillates about resonance.

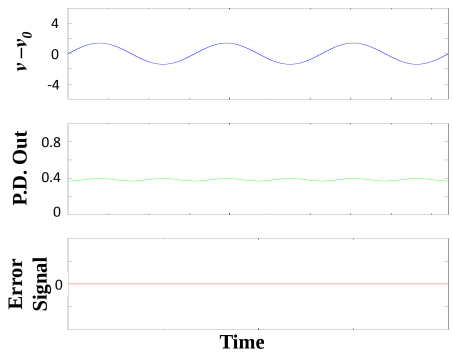

(a) The laser frequency centered on resonance.

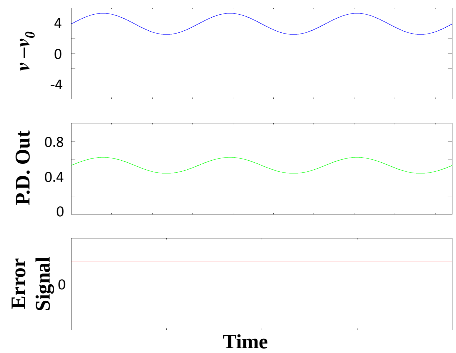

(b) The laser frequency centered well above resonance.

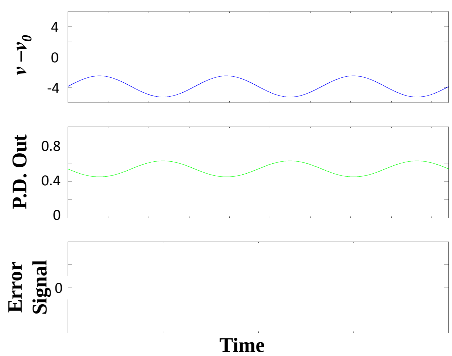

(c) The laser frequency centered well below resonance.

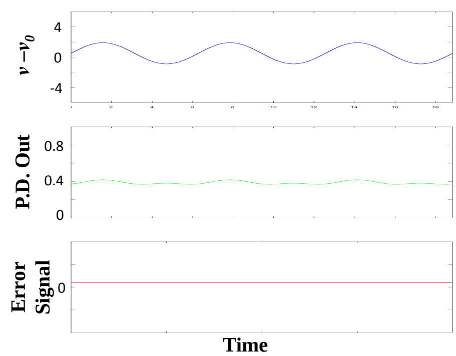

(d) The laser frequency centered just above resonance.

Figure 1.5 These graphs track the frequency of the laser, the photodiode output, and the error signal as time evolves for four different situations. In the situations where the frequency is off resonance, the lock would correct the frequency and bring it toward resonance over time. These graphs then show a situation where an error signal is being produced but not being used to correct the laser frequency. The values of $\nu - \nu_0$ in this graph directly relate to the values shown in Fig. 1.4.

This figure begins with the laser frequency on resonance. When on resonance, the photodiode output is minimal. As the modulation drives the frequency above resonance, the photodiode output increases until the modulation begins driving the frequency downward and the photodiode output then begins to drop. The photodiode output continues to decrease until the frequency passes through the resonance. As the frequency continues to decrease past the resonance, the photodiode output then increases until the laser frequency its minimum value. As the frequency increases from this point, the photodiode output decreases again until it passes through the resonance, after which the photodiode output increases again. This completes one period of the frequency modulation. In this time, the laser frequency maximizes only once while the photodiode output maximizes twice. Therefore, the photodiode output oscillates at twice the modulation frequency (as can be seen in the figure). Consequently, the filtered output from the mixer will be a 0 V DC signal because the photodiode output has no components that oscillate at or near the modulation frequency.

Fig. 1.5(b) shows this process when the laser frequency oscillates about a point well above resonance. This instance is much simpler than the last. When the modulation increases the laser frequency, the photodiode output increases and when the modulation decreases the laser frequency, the photodiode output decreases. The amplitude of the photodiode output oscillations when centered above resonance is much larger than the amplitude of the oscillations when they are centered on resonance because the resonance curve is much steeper away from resonance (see Fig. 1.4). In this situation, the photodiode output oscillations are at the same frequency as and in phase with the modulation signal. So, when they are mixed together and low-pass filtered, the result is a positive DC signal whose amplitude tells how far off resonance the laser is.

Fig. 1.5(c) shows this process when the laser frequency oscillates about a point well below resonance. Because we are on the opposite side of the resonance curve as the previous example, the photodiode output in this situation is also exactly opposite. Now, when the modulation increases the laser frequency, the photodiode output decreases and when the modulation decreases the laser frequency, the photodiode output increases. The amplitude of the oscillation in this example is the same as the previous example because the resonance curve is symmetric. In this situation, the photodiode output oscillations are at the same frequency as the modulation signal but the signals are 180 degrees out of phase. Therefore, when they are mixed together and low-pass filtered, the result is a negative DC signal (negative because the photodiode output and modulation are completely out of phase) whose amplitude tells how far off resonance the laser is.

Fig. 1.5(d) illustrates what happens if the laser frequency oscillates about a point that is just above resonance. Here, the error signal is not as large as the well above resonance error signals for two reasons. First, the resonance curve is not as steep closer to the resonance. Second, the photodiode output in this case has two main frequency components. One of them oscillates at the modulation frequency in phase with the modulation and the other oscillates at twice the modulation frequency. When the photodiode output is mixed with the modulation signal and low-pass filtered, the frequency component that oscillates at twice the modulation frequency doe not contribute to the error signal because it is filtered out. Therefore, as the laser frequency approaches resonance, the error signal decreases because the steepness of the curve decreases and the amplitude of the frequency component that oscillates at the modulation frequency decreases.

To recap, we are trying to lock to an atomic resonance. We determine how far from resonance the laser frequency is and whether we are above or below resonance by modulating the laser frequency, directing the light through a vapor cloud onto a photodiode, amplifying and mixing the photodiode output with the modulation signal, and low-pass filtering the mixer output. This filtered signal is our error signal. The sign of the error signal tells the direction of the frequency drift and the amplitude tells how far away from resonance the laser drifted. If the error signal is a 0 V DC signal, the laser is on resonance and no adjustment is necessary.

This process has several advantages. First, the error tells both the direction and magnitude of any frequency drifts. Second, drifts in light intensity or atom density do not matter when we are on resonance because we are locking to a zero point. Finally, because this method uses lock-in detection, the error signal should have a high signal to noise ratio.

While this method does work, there is one problem with locking this way. Because the atoms in the vapor cloud are in constant motion, the resonance curve will be Doppler-broadened. As a result, the width value of the resonance curve will be much larger than the actual linewidth of the atomic transition. A process known as saturated absorption (also call Doppler-free spectroscopy) was developed to compensate for the Doppler-broadening of an atomic resonance. Since this process is not central to my thesis but may be of some interest to the reader, I have included information in Appendix B that briefly explains it.

As a side note, one way to make the vapor cloud discussed in this section is with a vapor cell. A vapor cell is basically a cylinder with two window ports on either end. To get a cloud of atoms inside the cell, a sample of the desired material is placed inside. The cell is then sealed, pumped down to vacuum, and heated. As the cell heats, the sample vaporizes. As long as the cell stays within an appropriate temperature range, there will be a fairly even distribution of vaporized atoms throughout the cell.

1.2.3 My Contribution to the Lock

Both of these locking processes involve detecting modulated light, amplifying the oscillating signal, and mixing it with the local oscillator. For the mixer and everything before it, the high frequency part of the signal is important. Consequently, I put all these components together onto one board. We call this circuit the "homodyne detector." This is the circuit that I have constructed. It is possible to make a circuit similar to mine with commercially available parts, but doing so has two main downsides. First, the performance of a commercially available alternatives is lacking compared to my circuit because it would involve purchasing all of the parts separately and linking them together with BNC cables (i.e., buying a high-speed photodiode and linking that to an amplifier which is connected to a mixer, etc.). Doing so would introduce significantly more noise than my circuit does because the components in my circuit are all placed close together on the same board. Not only would the noise be negatively affected, but linking through BNC cables to different components may also negatively affect the speed of the circuit. Overall, my circuit should out perform a commercial one because mine is quieter and runs at comparable speeds. The second downside to a commercial alternative is cost. My board is relatively inexpensive to produce while commercial parts may be very expensive (especially considering the high bandwidth these parts would need). The only downside to my board is that it will take some time to assemble a board once the parts have been acquired, but assembling and testing a board should not take long.

For my detector, I had two main goals. First, I wanted my detector to be shot-noise limited for a milliwatt laser. Second, I wanted to be able to modulate laser light at a high bandwidth, ideally up to 100 MHz. While these are ambitious goals, I have been able to make significant advances toward achieving them. Before presenting the results, I will first take an in depth look at my detector.

Chapter 2

Detector Layout

Fig. 2.1 shows a schematic of my detector. To help understand the schematic, I have included Fig. 2.2 (which shows a simplified block diagram of the circuit). In my circuit, light is detected by a photodiode which produces a current. That current is converted into a voltage using a transimpedance amplifier. The voltage signal is then filtered, amplified, filtered again, amplified again, mixed with a local oscillator, low-pass filtered, and sent to a PID controller that adjusts the laser frequency. In this chapter, I discuss each of these stages more thoroughly, though I save an in depth look at the op amps for chapter 3.

2.1 Photodiode

A photodiode is a device that converts optical power into electric current following the equation $I = AP$. In this equation, I stands for current and has units of Amperes. P stands for optical power and has units of watts. The constant of proportionality between the two, A, is known as the responsivity. Using dimensional analysis, one can see that this constant must have units of $\frac{amps}{watt}$. For any given diode, the responsivity

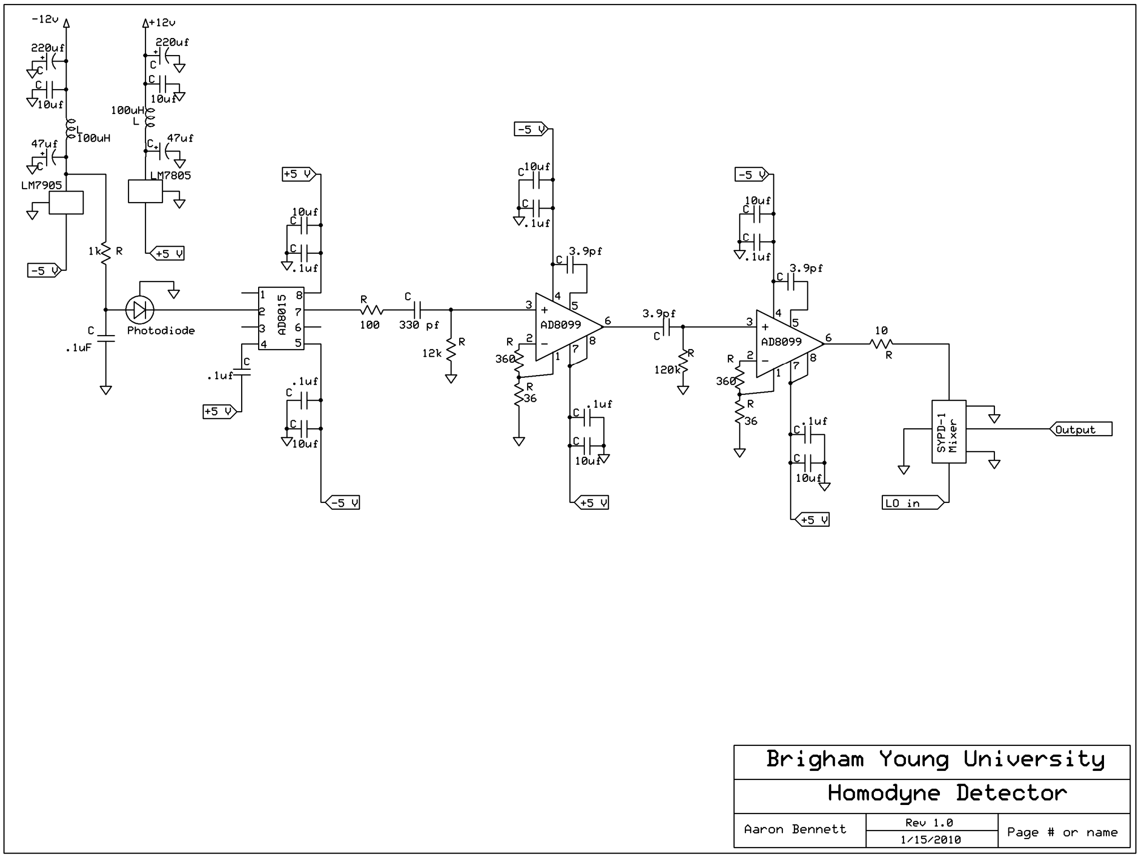

Figure 2.1 The schematic for my detector.

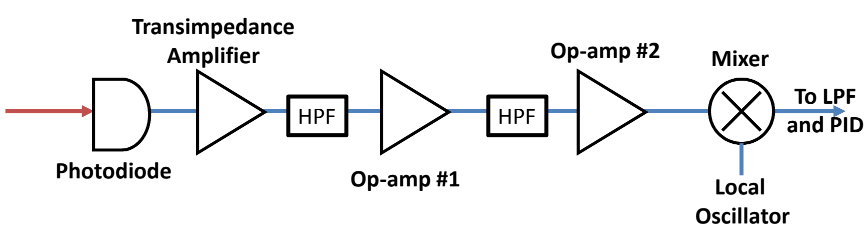

Figure 2.2 A simplified block diagram of the detector. In this figure, HPF stands for high-pass filter, LPF stands for low-pass filter, and PID refers to the PID controller that the error signal is sent to.

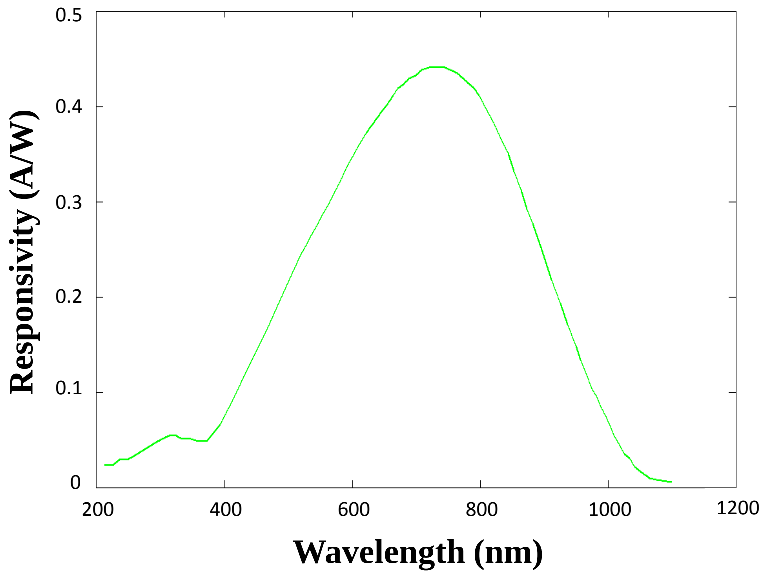

Figure 2.3 The responsivity curve for the photodiode used in my detector. This figure was reproduced from [6].

varies depending on the wavelength of light being detected. Most photodiode data sheets have a responsivity curve, showing the value of A at any given wavelength. Fig. 2.3 shows a a responsivity curve for the particular photodiode used in my detector.

A common method for determining a photodiode's output current is to put that current across a resistor and measure the voltage drop across the resistor. Then, using Ohm's law, one can calculate the current. Unfortunately, this method does not always work because there is a limit to the voltage a photodiode can output because it is a diode. Consequently, a photodiode cannot always output a current proportional to the power incident on the diode. Typically, a photodiode starts saturating when outputting between .2 and .4 V and will totally saturate somewhere around half a volt. In the circuit I built, I overcame this problem by using a transimpedance amplifier (discussed in the next section).

One particularly challenging aspect of working with photodiodes is their response time. The depletion region of a diode acts as a capacitor (because there are charges built up on either side of it). The capacitance of a photodiode in connection with whatever resistance is in series with the output forms a low-pass filter with a time constant of $\tau = RC$ (where R is the value of the resistor and C is the value of the capacitor). This time constant corresponds to a 3 dB bandwidth of $\frac{1}{2\pi RC}$. Based on this equation, there are two ways to increase the bandwidth of the circuit. First, we can decrease the value of R. Unfortunately, this has the negative effect of also decreasing the voltage drop across the resistor. Second, we could decrease the capacitance of the photodiode. Recalling that the photodiode acts as a parallel plate capacitor and that the capacitance of such a capacitor is $C = A\frac{\epsilon}{d}$ (where C is capacitance, A is the area of the capacitor, d is the distance between plates, and ϵ is the permitivity of the material between the plates), we see that the capacitance of the photodiode can be reduced in two ways: decreasing the area of the photodiode or increasing the depletion region (the effective area between the plates).

The first of these solutions is only effective to a point. If the area of the photodiode becomes so small that it is difficult to hit it with a beam, it is not really effective. For my project, I used a photodiode with an area of one millimeter by one millimeter. I chose this size because it has a relatively small area yet it is still reasonably easy to focus light on it.

The second way to reduce the capacitance (increasing the depletion region) is accomplished by a process know as reverse biasing. The schematic of a reverse biased diode is shown in Fig. 2.4. In my circuit, I reverse biased my diode with a voltage of -12 V. When biasing a diode, it is important to pay careful attention to the sign of the voltage. If one were to forward bias the diode (applying a positive voltage instead of a negative one), current will flow almost completely uninhibited through the diode, most likely destroying it.

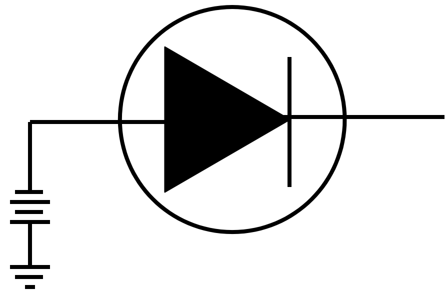

Figure 2.4 The schematic for a reverse biased photodiode. In my detector, I reverse biased the photodiode I used with -12 V.

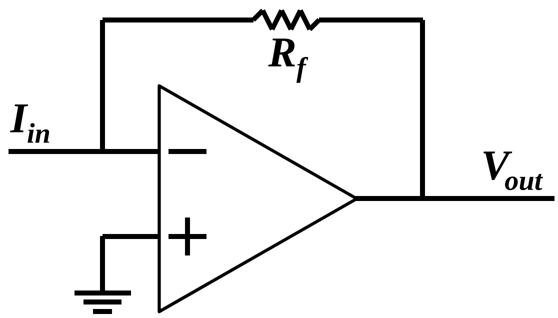

Figure 2.5 The transimpedance amplifier.

2.2 Transimpedance amplifier

In the previous section, I discussed two problems with using a resistor to transform the current from the photodiode into a voltage with a resistor. The first problem was associated with saturating the photodiode. Because the photodiode can only output up to a certain voltage, we cannot choose a resistor that is too large. The second problem dealt with speed. Making the resistor small increases the speed but also results in a lose of signal. We overcome both these problems by using an op amp wired as shown in Fig. 2.5.

A simple evaluation of the circuit shows that the output voltage is proportional to the input current as follows: $V_{out} = -I_{in} R_f$. Because the current is converted to a voltage across a resistor using an amplifier, this circuit is called a transimpedance amplifier. As evident in the equation for V_{out}, the size of the output voltage is directly proportional to the size of the feedback resistor.

The transimpedance amplifier solves the problem with saturation because the photodiode is attached to a virtual ground, so its output voltage stays at zero even as the current through the diode increases (until the amplifier saturates, which is usually at a much higher voltage than the photodiode saturation point). Consequently, when using a transimpedance amplifier, the linearity of the photodiode does not present a problem.

Connecting the photodiode to the virtual ground of the transimpedance amplifier also solves the speed problem. When connected to the transimpedance amplifier, the photodiode essentially behaves as though it is connected to zero resistance (because this is the resistance necessary to keep the voltage zero for arbitrary currents). Previously, I mentioned that the 3 dB bandwidth for the photodiode was $\frac{1}{2\pi RC}$. If R goes to zero, the photodiode should theoretically have infinite bandwidth. Unfortunately, the speed of the circuit is not infinite (meaning the response time is not 0) because it is limited by the capabilities of the op amp I am using, but it does perform much better than just using a resistor to transform the current into a voltage.

In my circuit, I use the AD8015 for my transimpedance amplifier. The block diagram for this chip is shown in Fig. 2.6. Including the whole circuit in one chip removes the loop around the op amp that would be necessary had I chosen to just use a high speed op amp wired as a transimpedance amplifier. This loop would add stray inductance to my circuit, which would significantly decrease its speed.

Another advantage to this chip is that there is both an inverted and non-inverted

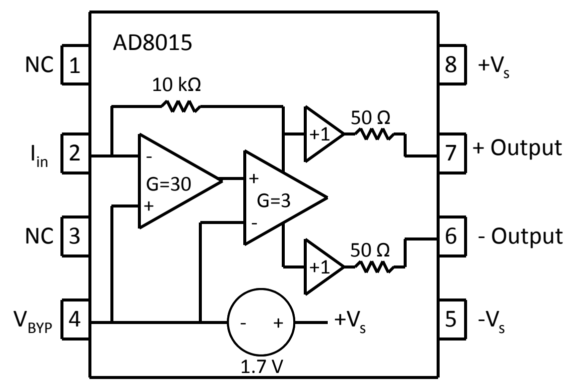

Figure 2.6 The AD8015 block diagram. I reproduced this image from [5].

output pin (meaning that the chip has two outputs that are 180 degrees out of phase with each other). The only major disadvantages to the AD8015 are a somewhat limited voltage swing and a DC offset on the ouput pin (both of which will be discussed more later in my thesis). For now it will suffice to say that I avoid saturating the output and I put a high-pass filter after the AD8015 to remove any DC bias.

2.3 Gain Stage

Because a majority of the initial light intensity is directed toward the experiment, the amount of light used for the lock is often very small, sometimes even on the order of a microwatt. As a result, the signal after the transimpedance amplifier will also be very small. In order to mix well, the signal needs to be amplified many times (the mixer will be discussed in greater detail later). It would be ideal to perform this amplification all in one stage, but this is limited by the gain-bandwidth product of available op amps (a property discussed in chapter 3). To achieve the desired gain and preserve the speed of the circuit, I use two op amps that each have a gain of eleven. The final output from these op amps should produce a large enough signal to

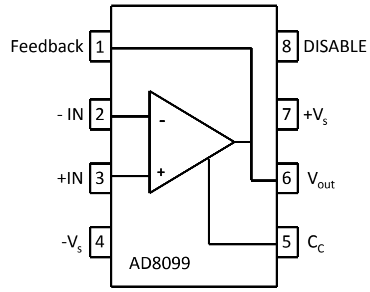

Figure 2.7 The AD8099 block diagram. I reproduced this image from [4].

satisfy the mixer.

The op amp that I am using for my detector is the AD8099. The block diagram for the AD8099 is shown in Fig. 2.7. The three main advantages to this op amp are a high gain-bandwidth product, very low noise, and the special feedback pin (pin 1). This pin makes it possible to lower stray inductance in the circuit. The two main disadvantages to this chip are that it adds a large DC bias and that it can only drive a very small capacitive load.

2.4 Mixer

The final stage of the detector is a mixer that multiplies the output signal from the last AD8099 with the signal from the local oscillator used to modulate the laser. As previously discussed, the amplitude and sign of the DC part of this signal will determine both the magnitude and direction of any adjustments that need to be made to the laser. This signal is sent through a low-pass filter to a PID controller that adjusts the laser.

The mixer I use in my circuit is the SYPD-1 from Mini-Circuits. According to the data sheet, the SYPD-1 produces the largest output when both input signals are at 7 dBm. The input pins for the signals are inductors to ground that act as 50 Ω resistors when receiving a signal between 1 and 100 MHz. Because the input pins are inductors, signals that are at or close to DC experience a short to ground. This is a problem because if there is a DC portion to the signal going to the mixer, the short will draw large amounts of current from the op amp (potentially killing both the mixer and the op amp).

Chapter 3

Non-idealities in op amps

The purpose of this chapter is to discuss op amps at a deeper level than what is usually discussed in a basic electronics course. I also discuss the characteristics of the specific amplifiers that I use in my homodyne detector. To ensure that the reader is familiar with basic op amp concepts and terminology, I first revisit the basic principles of op amps. Then, I discuss some of the non-idealities in op amps, particularly the ones that made the design of this circuit difficult.

3.1 Back to the Basics

An op amp has two input pins (one inverting and one non-inverting) and one output pin. Both input pins ideally have infinite resistance (therefore drawing no current) and no capacitance. The op amp's job is to make the signal at the two input pins identical by outputting a voltage which is connected to the inverting input pin through a feedback loop. To further illustrate this point, I will discuss two common op amp setups, a non-inverting and an inverting amplifier.

A schematic of a non-inverting op amp is shown in Fig. 3.1. In this figure, V_1 is

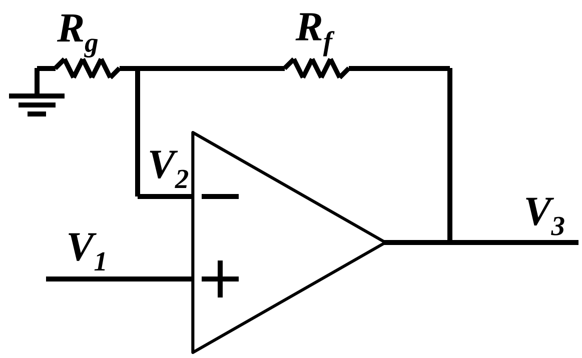

Figure 3.1 An op amp in a non-inverting setup.

the incoming voltage on the non-inverting input pin, V_2 is the voltage on the inverting input pin, and V_3 is the output voltage. In this instance, the purpose of the feedback loop is to divide down (reduce) V_3 (through a standard voltage divider) so that V_1 and V_2 are equivalent.

Let's figure out how much an amplifier like this amplifies a signal. The goal is to find some equation relating V_3 to V_1 in this form:

$$V_3 = GV_1,$$

where G is the gain of the op amp. Because the op amp makes the input pins equal, let's start by equating V_2 and V_1:

$$V_1 = V_2.$$

Next, let's find V_2 in terms of V_3:

$$V_2 = \frac{R_g}{R_g + R_f} V_3.$$

Because $V_1 = V_2$, we can substitute V_1 for V_2 in the previous equation and rearrange terms to find that:

$$V_3 = \frac{R_g + R_f}{R_g} V_1.$$

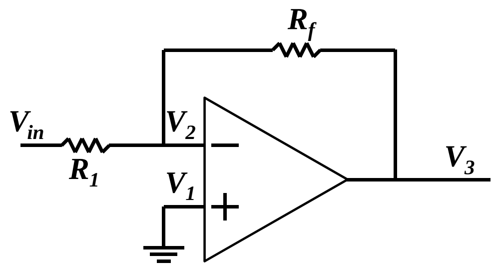

Figure 3.2 An op amp in an inverting setup.

Therefore, the gain G for a non-inverting op amp is: $G = \frac{R_g+R_f}{R_g}$. This setup earns its name because the output signal has the same sign as the input signal.

One common application of a non-inverting op amp is called a follower (also known as a buffer). In a follower, the gain is set to one by simply hooking the output pin straight into the inverting input pin. A follower has the advantage that everything after the op amp is isolated from everything before the op amp while the input and output signals are identical.

The next setup I will investigate is an inverting op amp (shown in Fig. 3.2). I have labeled the pins in this figure the same as for the non-inverting setup: V_1 is the non-inverting input, V_2 is the inverting input, and V_3 is the output. In this setup, neither one of the input pins is the voltage being amplified, so I labeled the voltage being amplified in this figure as V_{in}. To determine how much this circuit amplifies V_{in}, we want to relate V_3 to V_{in} by some factor G. To begin, because V_1 is grounded, V_2 will be a virtual ground. V_2 is also a superposition of V_3 and V_{in} according to the following equation:

$$V_{in} = \frac{V_3 R_1}{R_1 + R_f} + \frac{V_{in} R_f}{R_f + R_1} = 0.$$

Rearranging terms and solving for V_3, we find that:

$$V_3 = -\frac{R_f}{R_1} V_{in}.$$

So, for the inverting op amp, the gain is $G = -\frac{R_f}{R_1}$. This setup earns its name because the output and input signals have opposite signs.

This concludes the basic review of op amps. Unfortunately, an op amp's performance is not solely based on just picking the right resistors and deciding on an inverted or non-inverted signal. There are many other things that affect its performance, particularly when amplifying an AC signal. The rest of this chapter is devoted to discussing various properties of op amps that may cause them to operate non-ideally. While this chapter is a good guide, the data sheet should always be the first source of information for any op amp. Reading the entire data sheet may be long and tedious but often saves immense amounts of time later in a project.

3.2 Gain Bandwidth Product

3.2.1 What to be Aware of

When amplifying AC signals, it is important to consider the speed of the amplifier (whether it will be able to react as quickly as the signal oscillates). The speed of an amplifier is commonly characterized by its gain bandwidth product (GBWP). An op amp's GBWP tells the bandwidth over which an amplifier will behave linearly when amplifying by some gain. For example, if an op amp has a GBWP of 4 GHz, it should be able to amplify a signal by a gain of 10 for speeds up to 400 MHz. It could also amplify by a gain of five for speeds up to 800 MHz, etc. When the product of the gain and bandwidth exceed the GBWP, the output from the op amp dies off exponentially as the frequency increases. This can be seen in Fig. 3.4 (shown later in the section) which shows the frequency response of the op amp used in my detector. In that figure, the gain increases some amount before dying off. This phenomenon is known as gain peaking and is fairly common in op amps.

When considering the GBWPs of an op amp, it is important to note several things. First, the GBWP is only an estimate and doesn't always accurately predict how the op amp will perform. Second, many op amps are only stable over a specific range of gains (information found in the data sheet). For instance, the OPA-657 (produced by Texas Instruments) is only stable for gains larger than seven. The OPA-656, on the other hand, is only stable between unity gain and a gain of about five. Finally, the GBWP is specified for small signal gains (i.e., when V_{out} is less than about .2 V). For larger signal outputs, the bandwidth will be smaller than for small signal outputs. This is sometimes due to the slew rate of the op amp, a property discussed in the next section.

3.2.2 Specific to my circuit

There are two different areas in my circuit where I need to worry about the bandwidth of the amplifiers: the transimpedance amplifier and the gain stages. In general, the GBWP cannot be directly applied to a transimpedance amplifier because it performs a different function from most op amp applications (the transimpedance amplifier converts a current into a voltage while a standard inverting or non-inverting op amp simply amplifies a voltage). Although the GBWP has a different meaning for transimpedance amplifiers, I didn't have to worry about how it would affect my circuit because the AD8015 is hard-wired with a particular transimpedance gain. According to its data sheet, the AD8015 has a 3 dB bandwidth of 100 MHz. Fig. 3.3 shows the frequency response of the AD8015.

The AD8099 has the remarkably high GBWP of 3.8 GHz. Unfortunately, when outputting larger signals, the bandwidth deteriorates. Fig. 3.4 shows the frequency response of the AD8099 in both small and large signal situations. Clearly, the large signal response is not as fast as the small signal response. It is also important to

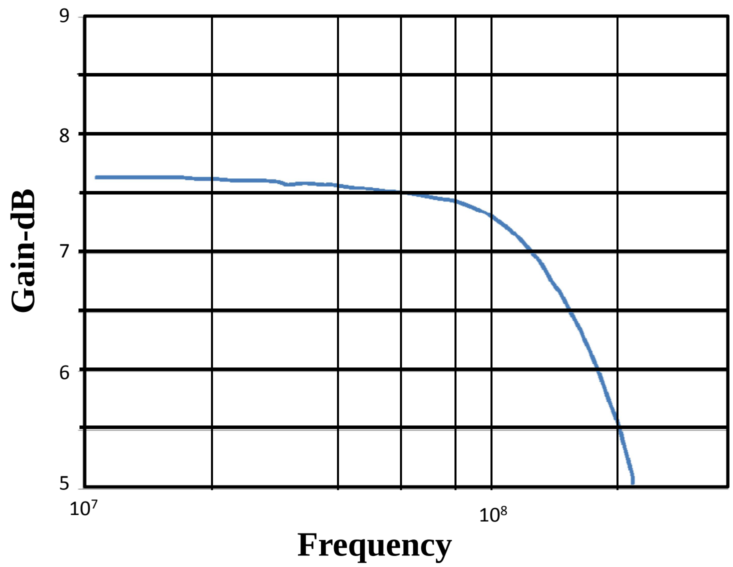

Figure 3.3 The frequency response of the AD8015. I reproduced this image from [5].

note that the AD8099 is stable for gains between 2 and 20 for both the inverting and non-inverting arrangements.

3.3 Slew Rate

The slew rate of an op amp defines the maximum rate of change for the output voltage signal. A common unit for slew rate is volts per microseconds. If the output of an op amp is a sine wave whose equation is:

$$V = V_0 sin(\omega t)$$

it's maximum rate of change will be:

$$max|\frac{dV}{dt}| = max|V_0 \omega cos(\omega t)| = V_0 \omega.$$

As long as $V_0 \omega \leq SR$ (where SR is the slew rate), the op amp will be able to respond fast enough to amplify the signal without any problems. On the other hand, if

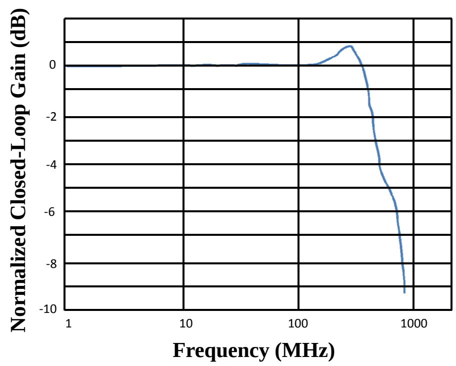

(a) AD8099 small signal frequency response.

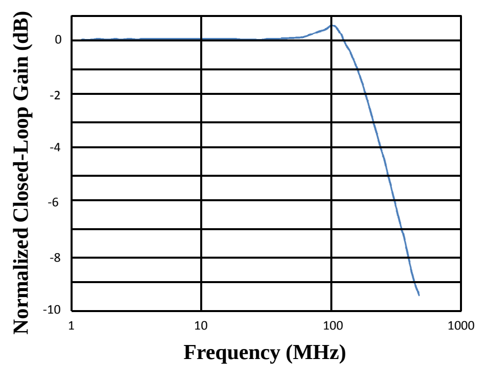

(b) AD8099 large signal frequency response when outputting 2 V peak-to-peak.

Figure 3.4 Frequency response of the AD8099 when amplifying by a gain of 10. I reproduced these images from [4].

$V_0\omega > SR$, the op amp will not be fast enough and the output signal will be distorted.

In the previous section on GBWP, I mentioned that there was a difference in bandwidth for small and large signals. In some situations, the slew rate of the op amp is primarily responsible for this difference.

In my circuit, the AD8099 has a slew rate of 1350 $\frac{V}{\mu s}$ when amplifying by a gain of ten. The fastest signals I would ever amplify would be somewhere around 100 MHz and the largest signal I would ever be driving at that frequency would be around 1.4 V peak-to-peak (about 7 dBm, the ideal input for the mixer). In this instance, $V_0\omega = 70\frac{V}{\mu s}$, which is much smaller than the slew rate of the AD8099. Therefore, the slew rate should not cause any problems in the operation of my detector, even when outputting 7 dBm.

3.4 Load Considerations

An op amp's response is also determined by the load it is driving. Resistive loads that are too small draw more current than the op amp can supply, causing the op amp to be unstable. In my circuit, the data sheet for the AD8015 states that it can drive resistive loads as small as 50 Ω. In practice, I found that it can drive 50 Ω but the output signal is significantly more stable if the load is 100 Ω instead. Since driving this larger load does not negatively affect my detector's performance in any way, I chose to use a 100 Ω load on the output of the AD8015. The data sheet for the AD8099 never specifically gives a minimum output impedance, but it does show several diagrams where the load resistance is 1000 Ω. In my detector, the load resistance for the first AD8099 is simply the input impedance of the next AD8099 combined in parallel with the resistance used to create the high-pass filter between the two chips. The second AD8099 simply drives the 50 Ω input impedance of the mixer plus a 10 Ω resistor that I put between this chip and the mixer (the 10 Ω resistor is so that any DC portion of the signal going to the mixer is not shorted to ground). Driving this small of a load does not appear to negatively affect the performance of this op amp.

While driving resistances that are too small makes an op amp unstable, capacitive loads generally cause an op amp problems when they are too large. Driving a large capacitive load will usually not kill an op amp but it will destroy the performance of most high-speed op amps (often causing the circuit to oscillate). The data sheet for the AD8015 does not mention a specific maximum capacitive load that it can drive. In my detector, I have the AD8015 driving a 100 Ω resistor followed by a 330 pf capacitor. I tested to see if this capacitive load affected the output of the detector and found that the response curves of the AD8015 both with and without the capacitor are essentially identical. From this I concluded that this capacitive load does not negatively affect the operation of the AD8015. The graphs for the AD8099 in its data sheet suggest that it has trouble driving anything over 5 pf, although the data sheet never gives a specific value for a maximum capacitive load. In practice, I found that it does behave poorly with a load much larger than that, so I chose to use a 3.9 pf capacitor in the filter after the first AD8099. For reasons discussed in a later section, I do not have a capacitor after the second AD8099.

Because the AD8099 can only drive very small capacitive loads, it is very difficult to probe. In doing tests of this circuit, I found that even though my scope probe has a capacitance of only 12 pf, this extra capacitance caused the AD8099 to oscillate in unpredictable ways. The inability to probe anything after the first AD8099 makes this a very difficult circuit to debug.

3.5 Input Impedance and Capacitance

In the review of basic op amps at the beginning of the chapter, I stated that an ideal op amp has infinite input resistance and no input capacitance. Unfortunately, in real life, op amps don't meet this condition. Every op amp has some finite resistance and some non-zero capacitance that must be considered when evaluating a circuit, especially when amplifying high frequency signals. When listing input resistances (and possibly other characteristics of the op amp), a data sheet will often refer to common mode and differential mode. Common mode refers to when the two signals on the input pins are the same and differential mode refers to when these signals are different. In this analysis, I assume that the signals on the input pins are close enough that any effect from differential mode impedances is negligible.

To include the input resistance and capacitance in an analysis of a circuit, one can

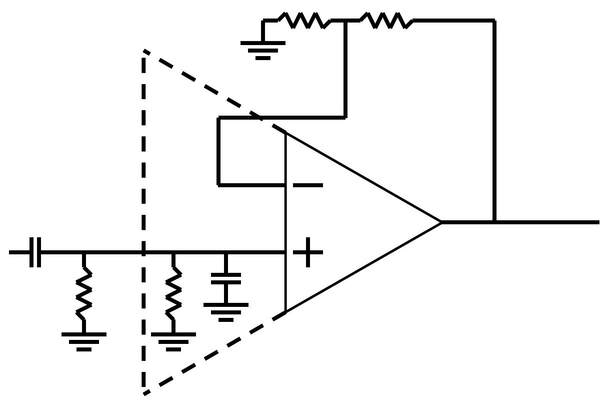

Figure 3.5 The resistor and capacitor inside the dotted lines represent the internal resistance and capacitance of the AD8099.

model the common mode input values as some impedance to ground directly before an ideal op amp. To better illustrate this point, let's examine the part of the circuit between the two AD8099s. Fig. 3.5 shows how to include the internal resistance and capacitance in the analysis of the circuit. On the actual board, we have placed a 3.9 pf capacitor in series with the output followed by a 100 KΩ resistor to ground. The AD8099 has a common mode input resistance of 10 MΩ and input capacitance of 2 pf. We treat the impedance of the 100 KΩ resistor and our two input impedances as one impedance (the parallel combination of all three) to ground. Then, we simply treat the circuit as a voltage divider (with imaginary impedances) between two ideal op amps. Fig. 3.6 shows the frequency response of this voltage divider with and without considering the input impedance and capacitance. Clearly, including these values makes a difference.

3.6 Voltage Swing

The voltage swing of an op amp tells what maximum and minimum voltage the op amp will be able to output (or, in the case of the AD8015, it tells the maximum

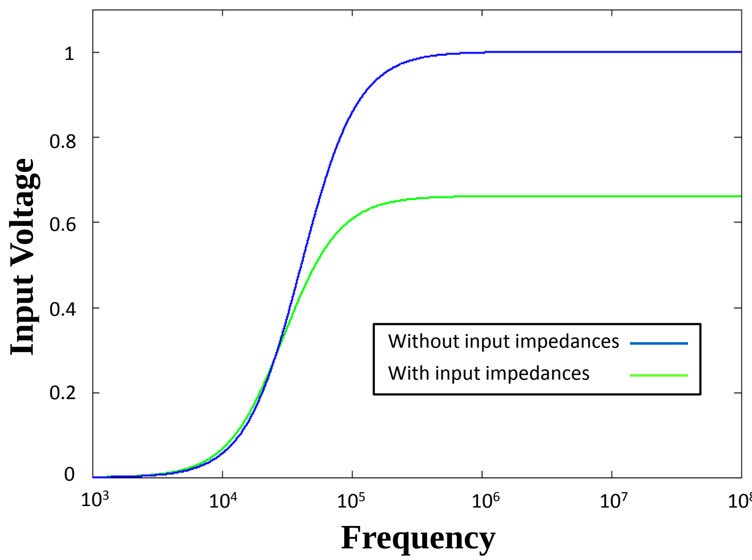

Figure 3.6 This graph shows the importance of including input resistance and capacitance in the complex analysis of a circuit. The blue line represents the predicted normalized output of the amplifier when ignoring input resistance and capacitance while the green line represents the predicted output when including these impedances. In this instance, the input capacitance of the op amp is the primary cause of the decreased output.

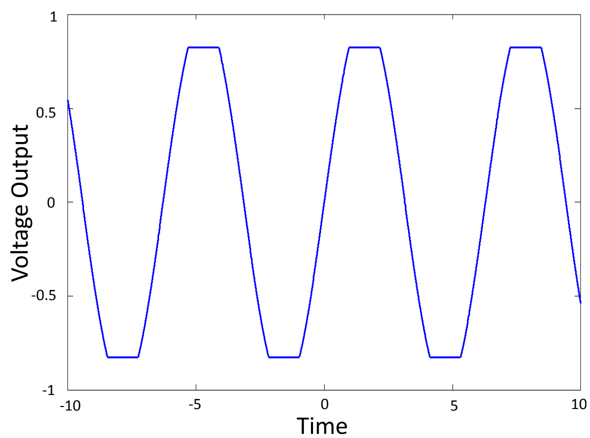

Figure 3.7 In this figure, the voltage swing would be about 1.64 V peak-to-peak.

peak-to-peak voltage it can output). Because the op amp physically cannot output voltages outside of its voltage swing, it will instead cut off the top and bottom of the signal to keep it within the voltage range. Fig. 3.7 depicts a signal that has exceeded the voltage swing of the op amp that is outputting the signal.

For my circuit, the AD8015 has a voltage swing that is dependent on the resistive load it is driving. For an infinite load, it has a voltage swing of 1 V peak-to-peak. For a load of only 50 Ω, it's voltage swing is 0.6 V. Because my AD8015 is driving 100 Ω, I expect my voltage swing to be around 0.6 V. Just to be safe though, I kept the output signal much less than 0.6 V. The AD8099 has a voltage swing of roughly -3.5 V to 3.5 V, with slight changes depending on the output resistive load.

3.7 DC Bias

There are many sources of DC bias in my circuit. First, the signal from the photodiode has some bias because the laser signal is the sum of a constant light level plus an oscillating light level. Next, the AD8015 intentionally adds a bias to its output signal (which is needed in some of its applications). Further biases occur because the inputs of the op amps are not ideal infinite impedances and therefore draw a small bias current. Finally, the AD8099's have intrinsic voltage offsets.

In my circuit, the AC part of the signal carries all of the information necessary to produce the error signal. Not only does the DC portion of the signal carry no pertinent information, it actually harms the performance of my circuit for two main reasons. First, the DC bias may saturate the output of my op amps. For example, the AD8015 adds a 1.3 V bias. Both AD8099's in my circuit amplify by a gain of eleven and can only output voltages between -3.5 V and 3.5 V. Therefore, trying to amplify the bias alone from the AD8015 saturates the first AD8099. The second reason that these biases present a problem is that the input of the mixer acts as a short to ground when it receives a DC signal because the input pin is just an inductor to ground (the inductor acts as a 50 Ω resistor for an AC signal above 1 MHz). This short to ground draws large amounts of current out of the op amp, potentially killing both the op amp and the mixer.

Since I only use the AC part of the signal and the DC portion actually harms the circuit, I remove biases throughout the circuit using high-pass filters. A high-pass filter can be made with a resistor and capacitor as shown in Fig. 3.8. To analyze the frequency response of the filter, we treat this circuit as a voltage divider with the capacitor having complex impedance of $Z_1 = \frac{1}{i\omega C}$. Treating it this way, we find that:

$$V_{out} = \frac{R}{Z_1 + R} V_{in} = \frac{R}{R + \frac{1}{i\omega C}} V_{in}.$$

Fig. 3.9 shows a standard frequency response curve for V_{out}. This circuit is called a high-pass filter because the high frequency signals pass through unaffected while low frequency signals do not survive. This filter has a 3 dB point (the point when the power output has been filtered to half of it's maximum output) of $\frac{1}{2\pi RC}$.

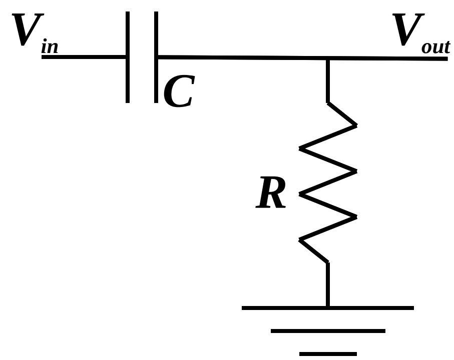

Figure 3.8 This is one way to make a high-pass filter.

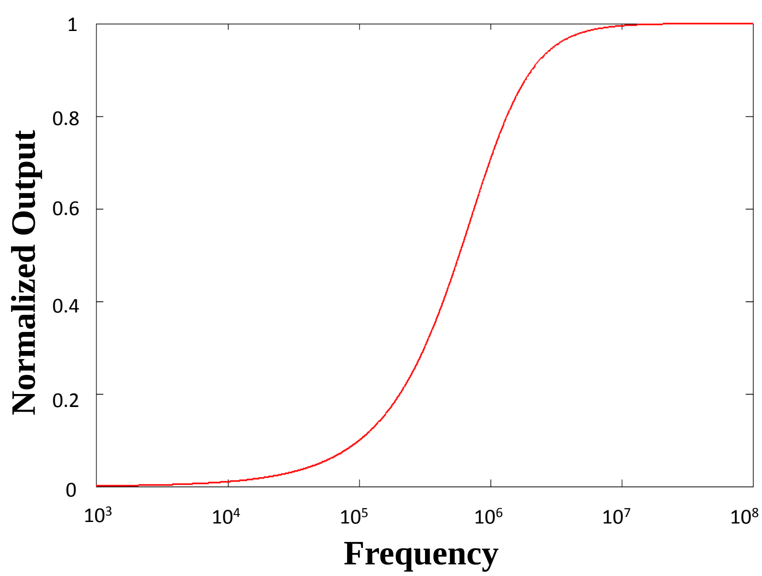

Figure 3.9 The normalized frequency response of a high-pass filter with a 3 dB point at 1 MHz.

To remove the bias after the AD8015, I used a high-pass filter with a 330 pf capacitor and a 12000 Ω resistor. This gives a 3 dB point around 40 KHz. This point is high enough that any DC signal is definitely filtered out and low enough that it won't filter out any of our modulated signal (we generally modulate above 5 or 10 MHz).

The DC bias for the AD8099 is dependent both on the intrinsic input voltage offset and the input bias current. The AD8099's input voltage offset is typically 0.1 mV. Because both of the AD8099's have gains of eleven, each op amp will add an offset of 1.1 mV to the output signal. The input bias current adds a DC offset to the signal on top of the 1.1 mV. The size of this offset is dependent on the state of pin 8 (referred to as the disable pin in the data sheet). If the disable pin is floating, the input bias current will be sixty times larger than if the disable pin is connected to the positive power pin. Consequently, in my circuit, I have chosen to connect these two pins. The only negative repercussion to doing so is that this introduces slightly more noise. With this smaller current, the offset voltage produced as a result of the input bias current will be on the order of a mV as well (this voltage value depends in part on the components that come before the chip).

Even though each individual AD8099 does not add that large of an offset, it is important to remember that the offset on the output of the first AD8099 will be multiplied by the gain of the second. Consequently, we put a high-pass filter in between them. Because the mixer wants only an AC signal, it would also be advantageous to be able to place a high-pass filter after the second AD8099. I was not able to do this, however, because the AD8099 cannot drive much capacitance and the input resistance of the mixer is only 50 Ω. Even if we put a 10 pf capacitor in our filter (which seems to be the largest capacitor the AD8099 could possibly handle), the 3 dB point of this high-pass filter would be well above 100 MHz (meaning that nothing below 100 MHz would get through the filter).

Not being able to filter the signal between the last AD8099 and the mixer introduced some problems for this circuit. Because the DC bias that the AD8099 adds is present any time the op amp is powered up, the AD8099 will always be trying to drive some DC signal into the mixer. This is a problem because the input of the mixer acts as a short to ground when it receives a DC signal (the input pin is just an inductor to ground). The data sheet for the AD8099 says that it will output somewhere near 140 mA when shorted to ground. However, the maximum current input for the mixer is 20 mA. So, it is necessary to put some resistor after the last AD8099 to limit the current it outputs to the mixer. The value of this resistor should be large enough to ensure the safety of the mixer but small enough so that it won't make too large of a voltage divider when handling an AC signal. Based on the DC input bias current on the AD8099 and the resistance to ground directly before the op amp, I calculated that the maximum DC offset after the last op amp may be as high 130 mV (though it will likely be much smaller than this). A 10 Ω resistor should then be large enough to protect the mixer. This value seems ideal because it is also fairly small compared to the 50 Ω input of the mixer when receiving an AC signal (so the RF signal will not be divided down by very much).

3.8 Parasitics

When discussing parasitics in this thesis, I specifically mean stray capacitance and stray inductance introduced by the layout of the board. Remembering that a capacitor is simply two metal sheets in some parallel orientation to each other and that an inductor is simply a loop (or series of loops) of wire, it is not surprising that a circuit with many elements like mine ends up with parasitics. Because my circuit is dealing with signals at such high speeds, it is important to design the board in such a way that reduces parasitics as much as possible.

One of the biggest things that I did from the beginning that helped me in this aspect was to use surface mount components. Because surface mount components are smaller, more compact, and sit closer to the board, they are far superior to other packaging types for removing parasitics. Although surface mount components help, there are still other necessary things to consider when looking to reduce parasitics.

To model stray capacitance, I wrote a Matlab script that determines the stray capacitance of a wire some distance away from a ground plane. This script mainly helped develop intuition for how the orientation of two metal planes affects the stray capacitance of the setup. Furthermore, it provided an estimate as to the order of magnitude for typical stray capacitance. Details about this script are included in Appendix C. After running the simulation many times, I found that the two factors that affect stray capacitance the most are the proximity of the trace (the line that the signal travels along) to the ground plane and the length of the trace. The closer together the planes are and the longer the trace, the larger the stray capacitance (this was the result that I was expecting since the equation for the capacitance of a parallel plate capacitor is $C = \frac{A\epsilon}{d}$ where A is the shared area and d is the separation between the plates). So, to stop stray capacitance from affecting my circuit, I cleared the ground planes on both the top and bottom of the board that were in close proximity at all to any line carrying a high frequency signal.

The inductance of a loop of wire is $L = \frac{N\Phi}{I}$ where L is inductance, N is number of turns, Φ is magnetic flux through the loop, and I is current in the loop. To reduce stray inductance I decided to first remove any unnecessary loops from the line carrying my signal. In one of the earlier versions of my circuit, I used a different op amp whose data sheet suggested placing the feedback loops for the op amps on the back of the board. Because of the suggestion, I designed my board with the loop on the back, having the signal pass through two vias (small holes in the board used to transport the signal from the top of the board to the bottom) to complete the loop. Upon testing, I found that this version of the board was extremely slow. I concluded that the vias were likely adding inductance which slowed the entire circuit. In later versions of the board, I removed these vias. I also switched from using a high speed op amp for my transimpedance amplifier to using the AD8015 because it required no wire loops around the exterior of the chip. After making these changes, the line carrying the high frequency signal does not pass through any unnecessary loops.

Unfortunately, it is impossible to remove all the wire loops from my circuit (both op amps in the gain stage need feedback loops). Although loops are necessary, reducing the area of the loop will decrease inductance (because the magnetic flux is dependent on this area). The design of the AD8099 helps in this aspect. It is designed with the ouput pin connected to a special feedback pin that is on the same side of the chip as the input pins (see Fig. 2.7). This feedback pin allows us to make a feedback loop with an extremely small area. Consequently, the only two loops the high frequency signal does pass through have very little inductance.

3.9 Noise Considerations

In this section, I examine the noise that various electronic elements add to a circuit. All of the data, including figures and terms, that I use for the analysis of the noise in op amps comes from reference [3]. This section follows that reference closely, although I do skip some of the finer details and jump straight to the conclusions. During this analysis, I use the letter N with various subscripts to denote noise power.

The first type of noise that I will discuss is noise inherent in resistors called

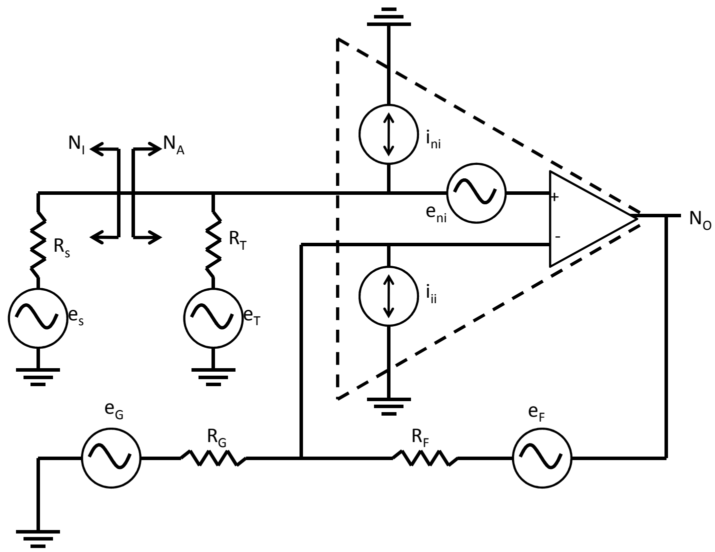

Figure 3.10 The noise sources for a non-inverting op amp. For the AD8099, the current noise on both inputs is the same. I reproduced this image from reference [3].

Johnson-Nyquist noise (or just Johnson noise, also sometimes called thermal noise). Johnson noise exists because of the quantum nature of electrons. The Johnson noise for a resistor with resistance R at temperature T is $N = 4k_bTR$ (where k_b is Boltzmann's constant). For instance, a 1 MΩ resistor at room temperature will have noise of 1.55 x $10^{-14}\frac{V^2}{Hz}$.

The second type of noise that I will discuss is noise introduced from op amps. There are two dominant types of noise that op amps generate: input voltage noise and input current noise. The resistors in the feedback loop also add Johnson noise. Fig. 3.10 shows the sources of noise for a non-inverting op amp. In this analysis, e with various subscripts denotes voltage noise and i with various subscripts denotes current noise. N_I is the noise present due to all of the components prior to the op amp. N_I has two contributions: the noise on the incoming signal (denoted by e_s) and the Johnson noise of the resistor R_s. N_A is all of the noise that the amplifying circuit introduces (this noise is amplified by the gain of the op amp). Three of the noise sources in the figure (e_T, e_G and e_F) are a result of Johnson noise. The other noise sources, e_{ni}, i_{ni}, and i_{ii} are the input voltage noise, input current noise on the non-inverting pin, and input current noise on the inverting pin. Because the data sheet for the AD8099 does not specify a separate input current noise value for the two different input pins, I will assume that both these pins have the input current noise listed in the data sheet. N_O is the total noise on the output signal. When our op amp has a voltage gain of G, $N_0 = (N_I + N_A)G^2$.

With respect to all of the values in this circuit:

$$N_A = c_1 e_{ni}^2 + c_2 i_{ni}^2 + c_3 i_{ii}^2 + c_4 e_T^2 + c_5 e_G^2 + c_6 e_F^2$$

where the c coefficient in each of these terms is defined below.

The input voltage noise (e_{ni}) adds in quadrature directly to N_A, so $c_1 = 1$. At the non-inverting pin, i_{ni} translates to a voltage through the parallel combination of R_S and R_T such that:

$$c_2 = (\frac{R_S R_T}{R_S + R_T})^2.$$

At the inverting input, i_{ii} translates to a voltage through the parallel combination of R_F and R_G such that:

$$c_3 = (\frac{R_F R_G}{R_F + R_G})^2.$$

The noise sources e_T and e_G result from resistors R_T and R_G. The magnitude of the noise from these resistors is divided down by resistors in the circuit. Therefore, the contributions to N_A due to the Johnson noise of R_T and R_G are:

$$c_4 e_T^2 = 4 k_b T R_T (\frac{R_S}{R_S + R_T})^2$$

and:

$$c_5 e_G^2 = 4 k_b T R_G (\frac{R_F}{R_F + R_G})^2.$$

Unlike any of the other noise sources for this circuit, e_f (the Johnson noise from R_F) adds noise to the output signal. Since all the other noise terms add noise to the input, I will write the noise on R_F as if it appeared on the input signal as well. I do this by reducing the voltage noise from this resistor by a factor of G (in this setup, $G = \frac{R_G}{R_F+R_G}$). The contribution to N_A due to the Johnson noise of R_F is therefore:

$$c_6 e_F^2 = 4k_b T R_F (\frac{R_G}{R_F + R_G})^2.$$

Because my circuit does not have any inverting op amps, I will not include a noise analysis of an inverting op amp. However, reference [3] contains this information for those who may be interested.

We also need to do a noise analysis of the transimpedance amplifier. The two main sources of noise here are the input voltage noise and input current noise, e_{trans} and i_{trans} respectively. The Johnson noise of the feedback resistor will also contribute to the noise of the circuit, but this contribution is small compared to the other input voltage and current noise so I am omitting it from this analysis. Referring to the resistor in the feedback loop as R_{trans}, the noise produced by a transimpedance amplifier is:

$$N_A = e_{trans}^2 + (R_{trans} i_{trans})^2.$$

In the results section, I present both the theoretical and measured noise for my circuit.

3.10 Harmonic Distortion

When amplifying a signal that oscillates at frequency ν, an op amp may introduce oscillations that are harmonics of the original signal. These added oscillations are known as harmonic distortion. For most op amps, only the second and third harmonic distortions are specified because these two are typically much larger than distortions from higher order harmonics. In some instances, only even or only odd harmonics are introduced (which is one reason both the second and third harmonic distortions are defined). In my circuit, the second and third harmonic distortion are small enough that they never cause any significant problems.

Chapter 4

Results

For my detector, I had two main goals: low noise and high bandwidth. In this chapter, I will discuss how my detector performs in both of these areas.

4.1 Noise

Specifically, I wanted my detector to be shot noise limited for a milliwatt laser. To calculate this noise limit, let's start by setting up an equation for the variation in energy, ΔE, as light passes by a point over an interval of time Δt. ΔE is related to the energy of each photon, $\hbar\omega$, and the number of photons, N, by the following relationship:

$$\Delta E = \sqrt{N}\hbar\omega.$$

If the light has power P, we can write N as the total energy to arrive during the time interval divided by the energy of each photon:

$$N = \frac{P\Delta t}{\hbar\omega}.$$

Plugging in this value of N into the equation for ΔE we get that:

$$\Delta E = \sqrt{\frac{P\Delta t}{\hbar\omega}}\hbar\omega = \sqrt{P\Delta t\hbar\omega}.$$

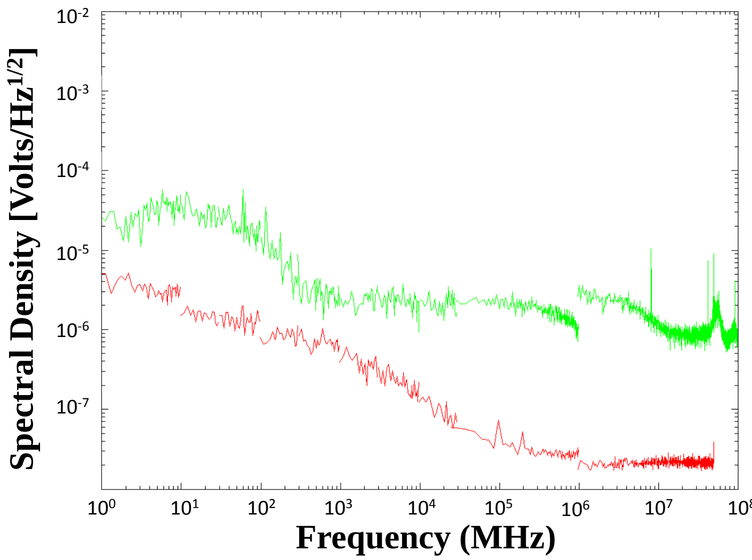

Figure 4.1 The noise spectrum of the homodyne detector measured after the mixer.

If the light has wavelength of 600 nm, then ΔE over 1 second will be 1.816 x 10^{-11} $\frac{Watts}{Hz}$. This relates to a voltage noise of 4.326 x 10^{-6} $\frac{V}{\sqrt{Hz}}$. So, if the noise of my detector is less than this limit, my detector is shot noise limited for a milliwatt laser.

Using the noise analysis presented in the previous chapter, I calculated that the theoretical total noise on the signal going into the mixer is 3.6 x 10^{-6} $\frac{V}{\sqrt{Hz}}$. The noise from the transimpedance amplifier dominates over all of the other noise introduced from the other elements in the circuit because this noise is amplified by both op amps. As long as the mixer in my detector does not add significant noise, my detector should be about shot noise limited for a milliwatt laser.

In practice, the noise spectrum measured after the mixer looks like Fig. 4.1. In this figure, there are two lines. The green one represents the noise floor of the oscilloscope used to take this data. The solid line is the noise spectrum when there is light on the photodiode that is modulated at about 50 MHz (likely explaining the small peak in the spectrum right around 50 MHz). This data shows that the noise for my circuit at frequencies above 1 KHz is actually lower than the predicted 3.6 x 10^{-6} $\frac{V}{\sqrt{Hz}}$. This means that the error signal from my detector will be very quiet, providing a very stable lock.

4.2 Bandwidth

In chapter three, I examined the bandwidth for each of the amplifiers in my circuit. The AD8015 has a 3 dB bandwidth of 100 MHz. In both the small and large signal regime, the AD8099 has a 3 db bandwidth above 100 MHz when amplifying by a gain of 11 (although the small signal response is much faster). The complex analysis of the circuit (presented in Appendix D) shows that the various resistors and capacitors in the circuit should not cause any significant additional frequency roll off even out to 100 MHz. Therefore, the only thing that will cause frequency roll-off in my detector is the bandwidth limitations of the amplifiers. Because these limitations do not affect signals at low frequencies, the detector is working properly if the output at low frequencies is consistent with the complex analysis of the circuit and if the output at high frequencies has not rolled off by much more than 3 dB.

Unfortunately, I was unable to make any conclusive findings about the bandwidth of my detector. Fig. 4.2 shows the normalized detector output as a function of frequency (normalized to the maximum output). There are two things that I have yet to resolve. First, the detector output at all frequencies is about three times smaller than theoretically calculated. This was true for all different light intensities. Second, I have been unable to determine whether the unusually large output at about 7 MHz is a result of a resonance in the circuit. If it is a resonance, the output rolls off to 3 dB around 50 MHz and has decreased by about 9 dB around 100 MHz. If it is not a resonance, the circuit rolls off much faster than expected (dropping off about 12 dB

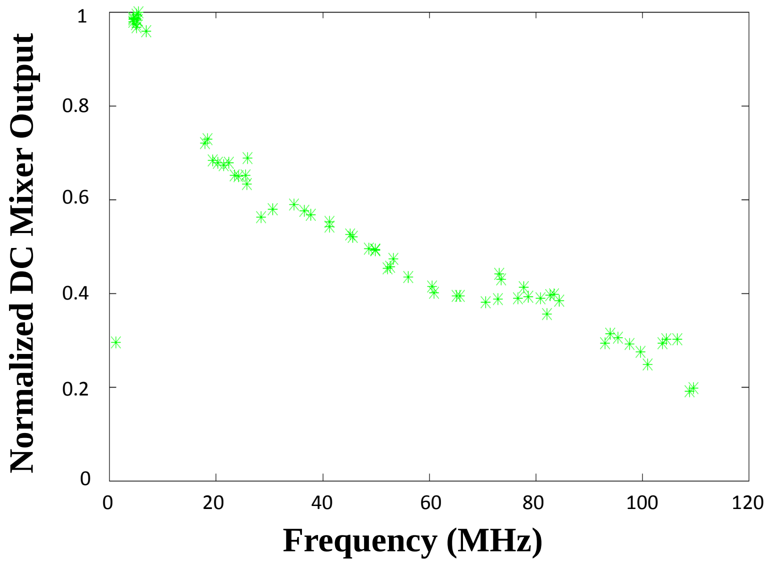

Figure 4.2 The frequency response of the homodyne detector.

by 100 MHz) for some unknown reason.

The first of these two problems may have a fairly simple explanation. The section of the circuit between the two AD8099's is particularly sensitive to stray capacitance. My complex analysis predicts that the larger the stray capacitance on the second op amp's input pin, the smaller the signal will be. Even as little as 10 pf of stray capacitance (which, according to the Matlab script on stray capacitance I wrote, is a very reasonable number) on the second AD8099's input pin will cause the signal at all frequencies to be reduced by a factor of three. Unfortunately though, if this is the cause of the reduction, there is not an easy solution. I have already designed the board as well as I know how to reduce stray capacitance. It is possible then that this overall reduction in the size of the error signal will not be fixable.

While this does not ruin the circuit, it does make it less sensitive to frequency drifts. Recalling from Chapter 1, the size of the error signal is directly proportional to how far the frequency has drifted. Therefore, if the error signal is much smaller than expected, the detector will not be sensitive to small frequency drifts. It will still be possible to lock the laser, the lock just won't be as stable as it would be if the error signal was not reduced in this way.

While the first of my problems in analyzing the detector's bandwidth may have a fairly simple explanation, the second of these problems, not being able to accurately characterize the bandwidth because of the possible resonance, cannot be easily explained. Until some explanation and solution is found to resolving this issue, it would be unwise to use this detector when modulating the laser much higher than 40 or 50 MHz. It is possible to lock when modulating at higher frequencies but the lock becomes increasingly insensitive to small drifts as the modulation frequency increases.

4.3 Conclusions

My goal was to create a homodyne detector to be used in an ultra stable laser lock. I planned to make a detector with lower noise and a larger bandwidth than a commercially available alternative. The detector that I created was particularly quiet. At frequencies above 1 KHz, the noise on my detector's output is less than 3 x 10^{-6} $\frac{V}{\sqrt{Hz}}$. Because my circuit is this quiet, the error signal it produces will also be very quiet (which is key in creating and maintaining an ultra stable laser lock). I was unable to make any conclusive findings about the bandwidth of my detector. This detector should not be used in a laser lock until the problems related to characterizing the bandwidth of the detector are resolved.

Bibliography

[1] J. Peatross, M. Ware, *Principles of Light and Optics.* (Brigham Young University, Provo, 2010).

[2] E. D. Black, "An introduction to PoundDreverHall laser frequency stabilization," Am. J. Phys. **69,** 79-87 (2001).

[3] J.Karki, "Calculating noise figure in op amps," Analog Applications Journal, **4Q,** 31-37 (2003).

[4] AD8099 Analog Devices Datasheet (2004).

[5] AD8015 Analog Devices Datasheet (1996).

[6] FDS010 Thor Labs Datasheet (2010).

Appendix A

Beer's Law

To derive Beer's Law, let's first imagine a thin slab of atoms of thickness dx and number density n_0. To begin, we can assume that all of the atoms in our slab are in the ground state and that our laser is exactly at the frequency necessary to drive this transition (we will later correct our model for when there is some superposition of ground state and excited state atoms and when the laser is not exactly on resonance). Let's say that each atom has a certain area where, if the laser passes through, the atom will absorb some intensity. It turns out that this area is proportional to the wavelength of light squared in this way: $\sigma_0 = \frac{3}{2\pi}\lambda_0^2$. Because there are many atoms and each will absorb some optical power if the light is incident on them, we write the change of intensity after passing through our slab as:

$$dI = -\frac{I}{A}\sum_j \sigma_0 = -\frac{I}{A}\sigma_0 j$$

where A is the area of the slab and j is the number of atoms in the slab. Since we have the number density of atom, n_0, and we have the both the area, A, and the thickness of our slab, dx, we can rewrite j in the following way: $j = n_0 A dx$. Plugging this into our equation for dI, we get that:

$$dI = -\frac{I}{A}\sigma_0 n_0 A dx = -I\sigma_0 n_0 dx.$$

Solving this differential equation gives us:

$$I = I_0 e^{-\sigma_0 n_0 x}$$

where I_0 represents the original intensity of light before entering the cloud of atoms and x is the distance that the light propagates through the atoms. This is Beer's law.

Now we need to extend this equation to situations where our light may not be exactly on resonance and where our atoms may be in a superposition of the ground and excited states. In this instance, the only thing to change from our previous model will be σ (because the number density of atoms and the distance the light has propagated will remain the same). First, to compensate for the fact that the light may not be on resonance, we add a Lorentzian to σ such that $\sigma = \sigma_0 L$. Here L is dependent on the frequency of light, ν, the resonant frequency of the transition, ν_0, and the natural linewidth of the transition, Γ, as follows

$$L = \frac{1}{1 + \frac{4(\nu-\nu_0)^2}{\Gamma^2}}.$$

By adding this factor, we have now compensated for our light not being exactly on resonance with the transition.

Next, we need to compensate for the fact that our atoms may be in the ground or excited states. If the atom is in the ground state and it absorbs a photon, the total light intensity drops. If it is in the excited state and it absorbs a photon, it undergoes stimulated emission and the light intensity increases. We can represent this mathematically by further changing σ in this way: $\sigma = \sigma_0 L(P_0 - P_1)$, where P_0 and P_1 represent the probability that the atom will be in the ground or the excited state, respectively. We can now plug in our modified σ into the equation we previously derived for I, finding that:

$$I = I_0 e^{\sigma n_0 x}$$

where $\sigma = \sigma_0 L(P_0 - P_1)$.

Appendix B

Saturated Absorption

When locking to an atomic resonance line, we compensate for Doppler broadening through a process called saturated absorption. The setup for saturated absorption is diagrammed in Fig. B.1. Here, two beams counter-propagate through a vapor cell. One of the beams is significantly more intense than the other beam. The more intense beam is called the pump beam and the less intense beam is the probe beam. As the pump beam goes through the cloud of atoms, the beam drives the atoms into an equal superposition of the ground and excited state (because absorption and stimulated emission are equally likely to occur). We then overlap the probe beam on top of the pump beam and look at its intensity as it exits the cell. Because the atoms that it propagates through aren't all in the ground state as they were before, the curve relating light frequency and output intensity now looks like the graphs in Fig. B.2 (with a hole burned into the response curve at the desired lock position). The width of this curve is much closer to the linewidth of the resonance we are trying to drive.

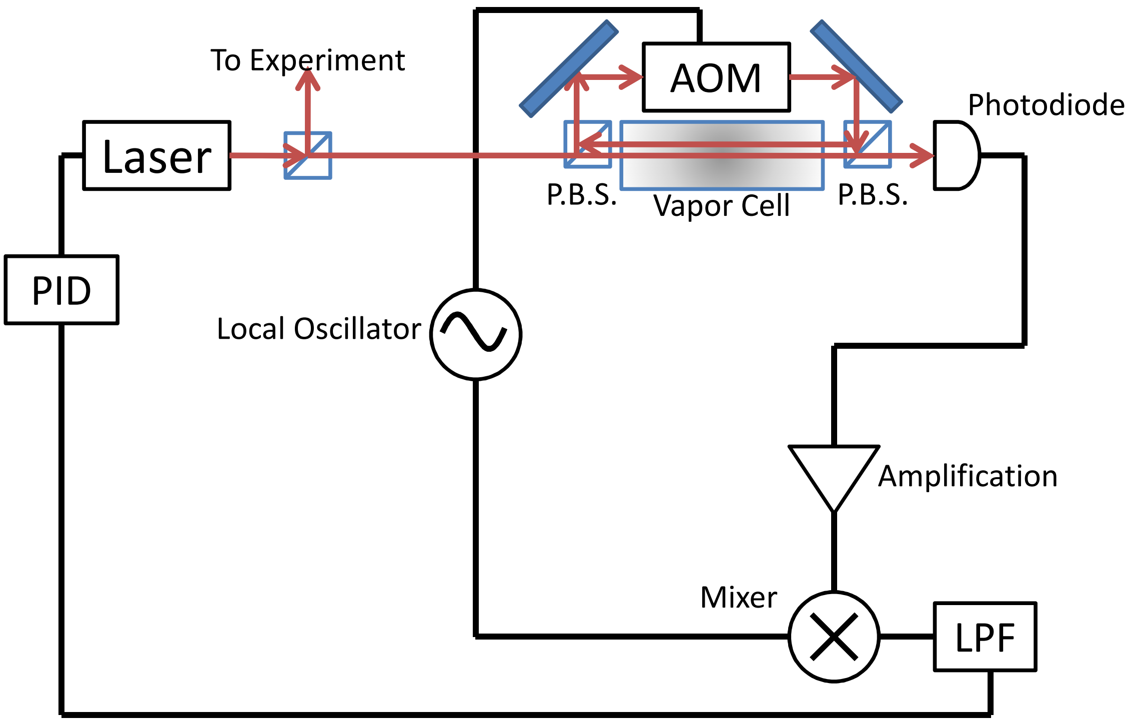

Figure B.1 A standard saturated absorption setup. In this figure, the pump beam is the beam that travels through the vapor going left. The AOM that the pump beam passes through determines where the hole in the resonance curve appears. P.B.S. stands for polarizing beam splitter.

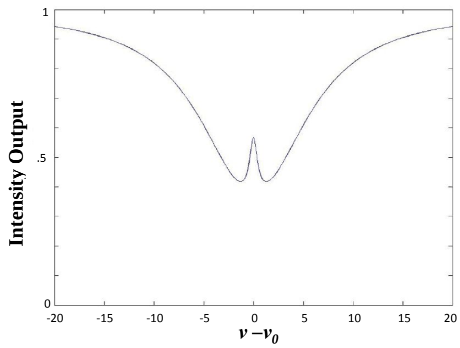

(a) The pump beam used to generate this curve is exactly at the atomic resonance.

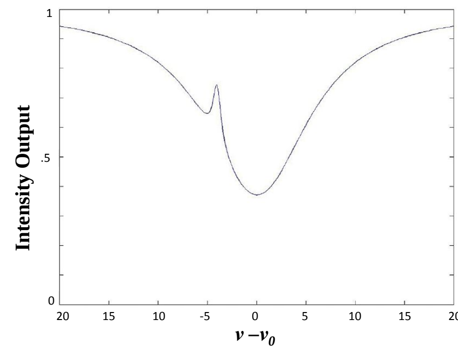

(b) The pump beam used to generate this curve is red-detuned with respect to the atomic resonance.

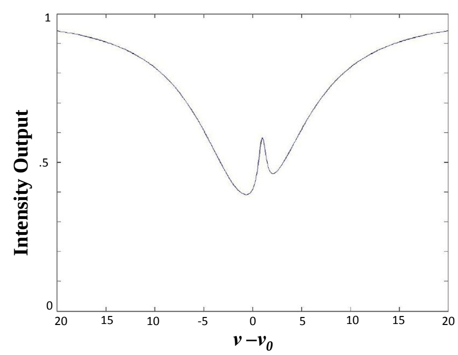

(c) The pump beam used to generate this curve is blue-detuned with respect to the atomic resonance.

Figure B.2 These graphs show the output intensity as a function of frequency for the probe beam in a saturated absorption application. These curves are calculated and not the result of an actual measurement. In the last two curves, the "hole" in the curve is different distances away from resonance simply to demonstrate that the lock point can be set anywhere and is not limited to just one specific distance away from resonance.

Appendix C

Computing Stray Capacitance

Because stray capacitance can cause problems when dealing with high frequencies, I made a mathematical model to understand this phenomenon. The model that I made is more to help understand what features contribute to stray capacitance than to calculate an exact stray capacitance value for some setup. The main things that I wanted to discover were how proximity of the trace to a ground plane as well as the width of the trace affect stray capacitance and what order of magnitude stray capacitance would have.

I accomplished this by first solving LaPlace's equation in two dimensions for a trace with some voltage that is some distance from a ground plane (with the ground plane being on either side of the trace). I used the process of Successive Over Relaxation to solve for the potential at all points outside of the ground plane and the trace. Once I solved for the voltage everywhere, I used that voltage and LaPlace's equation to solve for the charge built up at all points around the trace. I then added up these charges and divided by the voltage on the trace to find the total stray capacitance of the setup.

The code solves for the capacitance on a trace that is three millimeters long under varying circumstances. Each time the code runs, the user defines the width of the trace and the proximity of the trace to the ground plane by setting the left and right hand limits for the trace and the the ground plane. Fig. C.1 shows how the stray capacitance is affected by these two parameters. In Fig. C.1(a), the width of the trace is a constant 3 mm and the distance between the trace and the ground varies. In Fig. C.1(b), the distance between the ground plane and the trace is a constant 3 mm and the width of the trace varies. While the values for stray capacitance on these graphs may not be exact, they do show that stray capacitance can be on the order of tens of picofarads.

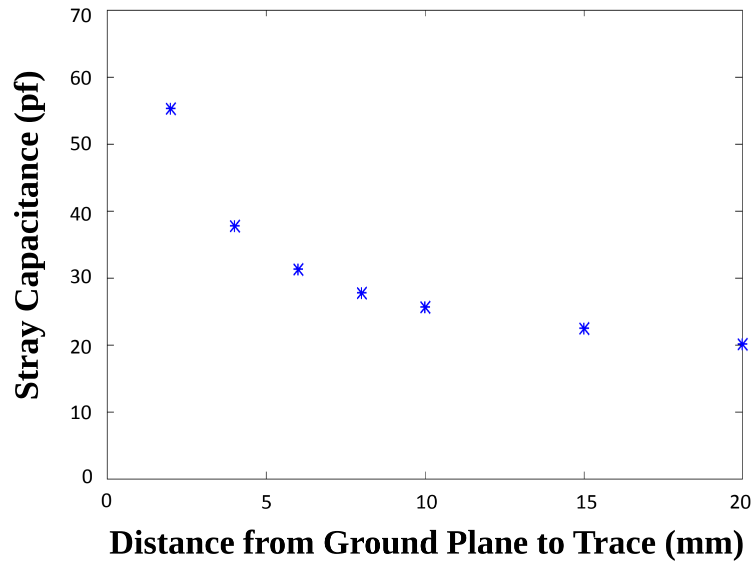

(a) In this figure, the width of the trace was held constant at 3 mm and the distance between the trace and the ground plane was adjusted.

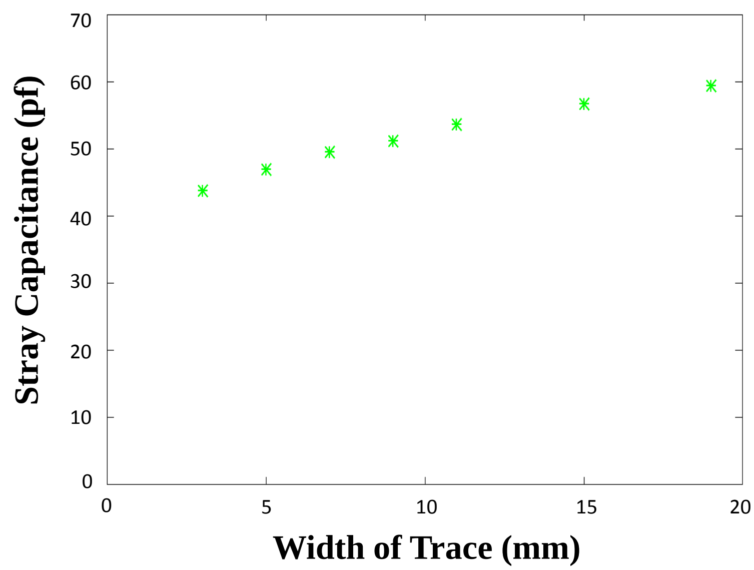

(b) In this figure, the distance between the ground plane and the trace was held constant at 3 mm and the width of the trace was adjusted.

Figure C.1 These graphs show the stray capacitance of a trace in close proximity to a ground plane.

Appendix D

Complex Analysis

Because my detector has many elements that interact in different ways with high frequency signals, the bandwidth of my detector will not just be determined by the bandwidth of each of my amplifiers. To discover how each of these elements contribute to the bandwidth of the circuit, I performed a complex analysis on my circuit . This complex analysis does not take into account the GBWP of the op amps. Instead, it focuses primarily on how the elements other than the amplifiers affect the bandwidth of the entire circuit. This section also provides the tools to be able to predict the maximum detector output for a given amplitude of modulated light.

Fig. D.1 shows each of the elements that will contribute to the circuit bandwidth. The elements inside the black dashed lines represent the common mode input capacitance and resistance of the op amp or the mixer. I modeled the sections of the circuit between the amplifiers and after the last amplifier as voltage dividers with complex impedances. Doing so, I was able to model any bandwidth limitations that may be caused by the elements in the circuit other than the op amps. The results, shown Fig. D.2, are normalized to what the amplitude of the voltage before the mixer would be if the signal from the transimpedance amplifier was simply multiplied by the gain

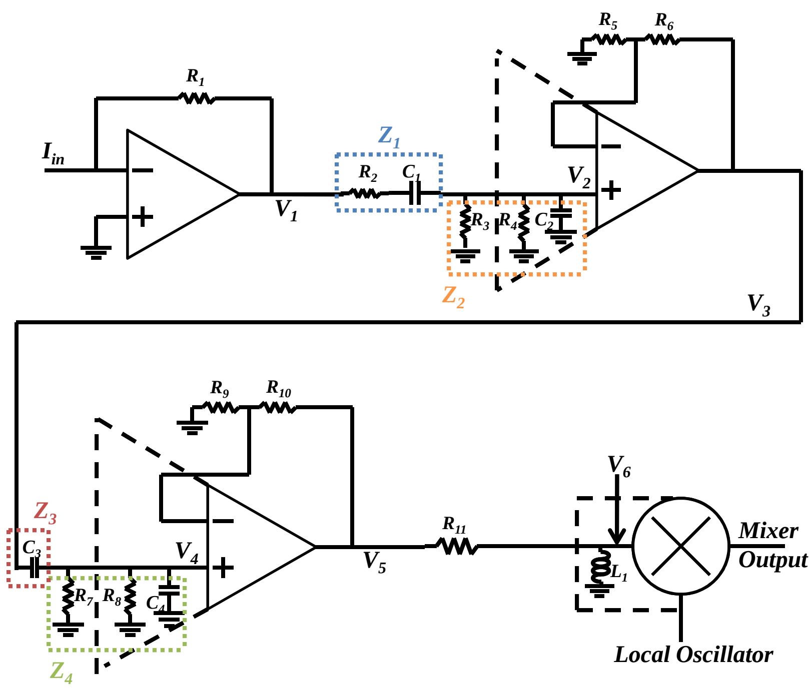

Figure D.1 A schematic showing all the elements that will affect the bandwidth of my detector. Because I treat the elements surrounding the amplifiers as voltage dividers, I grouped impedances together with the dotted colored lines. For instance, Z_1 is the combination of R_2 and the imaginary impedance of C_1 in series while Z_2 is the parallel combination of R_3, R_4, and the imaginary impedance of C_2. Using this notation, $V_2 = \frac{Z_2}{Z_1+Z_2}V_1$, $V_3 = \frac{R_6+R_5}{R_5}V_2$, etc. In my analysis, I treat L_1 (the inductor that provides the input impedance of the mixer) as a 50 Ω resistor.

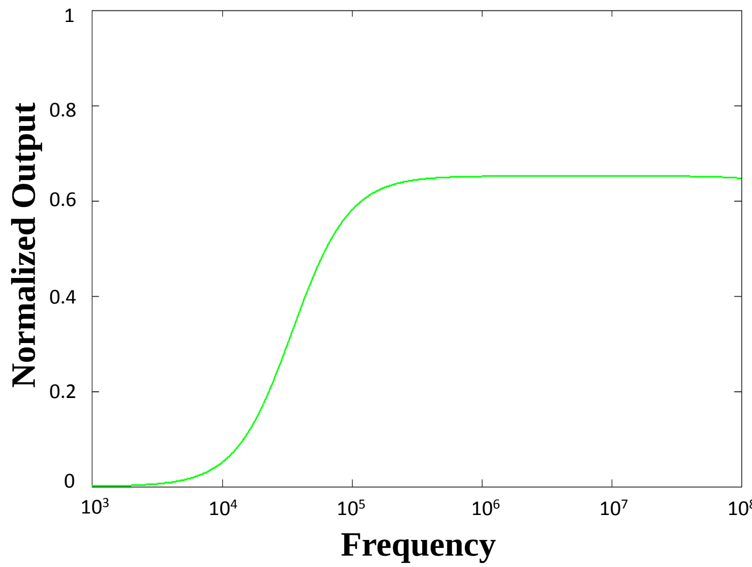

Figure D.2 The theoretical frequency response of the the detector up to the mixer (not including the response of the amplifiers). The y-axis is normalized to what the detector would be if the voltage out of the transimpedance amplifier was simply multiplied the gain of both op amps. In other words, the y-axis is normalized so that 1 represents the voltage output of the transimpedance amplifier multiplied by 121 (because each op amp has a gain of 11).

of both op amps (i.e., if the combinations of capacitors and resistors did not affect the signal at all, the curve would reach 1 on the y-axis).

This figure highlights two main things about the detector's frequency response. First, it shows that the combinations of resistors and capacitors form a high-pass filter. This is intended to remove any DC biases or noise at low frequencies. Second, including these resistors and capacitors cuts the signal by almost 40 percent at high frequencies. A more thorough investigation shows that the filter between the two AD8099s (this filter was analyzed in section 3.5) is primarily responsible for this cut-off.

Since the goal of this analysis is to predict the size of the error signal for a given amount of modulated laser light incident on the photodiode, it is necessary to an-

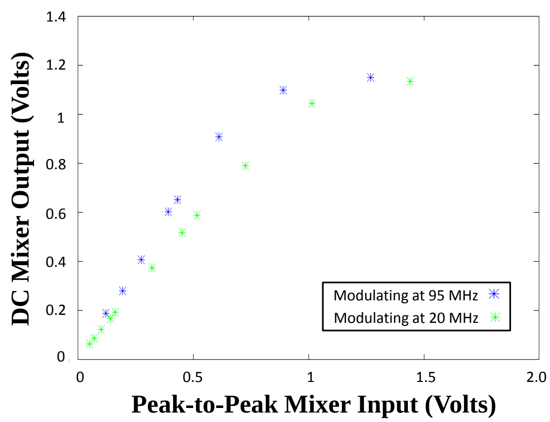

Figure D.3 To take this data, I held one of the mixer inputs constant at 7 dBm (the value suggested by the data-sheet). I varied the other input (plotted on the x-axis) and recorded the mixer output's DC offset (plotted on the y-axis). The two data sets are taken when the inputs were oscillating at 20 and 95 MHz.

alyze the mixer response to various-sized input signals. Modeling the mixer output was somewhat complicated because the data sheet provides no information about the mixer output when the inputs are anything other than 7 dBm. At first, I tried modeling this output as a linear response: the mixer output decreasing proportionately with the amplitude of the input signal (assuming that the one input stays at a constant 7 dBm). Testing this model proved that it was incomplete, so I decided to experimentally determine how the input amplitude affected the magnitude of the output. Fig. D.3 shows the results. In this figure, the x-axis is the peak-to-peak amplitude of the input signal (7 dBm is about 1.41 V peat-to-peak) and the y-axis is the maximum mixer output (when the two signals are either in phase or completely out of phase). Because the data sheet suggests that the maximum output will vary with the frequency of the two signals mixed together, I mapped out the output magnitude vs. the input amplitude at two frequencies that were at either end my detector's frequency range (one data set was taken at about 20 MHz and the other at about 95

MHz).

Testing the mixer proved very valuable for several reasons. First, I learned that the mixer's maximum output was significantly higher than the data sheet suggested (around 1.14 V as opposed to .932 V). Second, I learned that the maximum mixer output was roughly the same for different modulation frequencies (also contrary to the data sheet). Finally, I learned that the detector saturates even when the input is well below 7 dBm. This information about the mixer coupled with the information about the frequency response of the resistors and capacitors throughout the circuit allowed me to precisely calculate the theoretical detector output for a given amount of modulated light incident on the photodiode.