LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Technical Note	LIGO-T1900242–v2	2019/09/27

Active Reduction of Residual Amplitude Modulation in Free Space EOMs

S.M.Aronson, A.Gupta, R.Adhikari, C.Cahillane

California Institute of Technology
LIGO Project, MS 18-34
Pasadena, CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project, Room NW22-295
Cambridge, MA 02139
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

LIGO Hanford Observatory
Route 10, Mile Marker 2
Richland, WA 99352
Phone (509) 372-8106
Fax (509) 372-8137
E-mail: info@ligo.caltech.edu

LIGO Livingston Observatory
19100 LIGO Lane
Livingston, LA 70754
Phone (225) 686-3100
Fax (225) 686-7189
E-mail: info@ligo.caltech.edu

http://www.ligo.caltech.edu/

Contents

1 Introduction

1.1 LIGO

The Laser Interferometer Gravitational-wave Observatory, also known as LIGO, is an earth-based network of two gravitational wave detectors located in the United States. The design is similar to a Michelson Interferometer, but with the addition of Fabry-Perot cavities along both arms. Due to the high finesse cavities in LIGO, it requires ultra-stable frequency lasers to make the precise measurements needed for gravitational wave detections. Utilizing an array of detectors, it is possible to better find the location of an event in the sky and study it.

1.2 EOMs

An Electro-Optic Modulator (or EOM for short) is one type of optical component used within LIGO. For our purpose, we use it as a phase modulator in a Pound-Drever-Hall (PDH) set-up for laser frequency stabilization. An EOM consists of an electro-optic crystal, which means the crystal's birefringence depends on an applied electric field, and a set of electrodes on the top and bottom of the crystal. This crystal has two principal axes that can have different refractive indices dependent on the electric field in that direction. The voltage difference of the electrodes creates an electric field across that axis of the crystal which then changes the index of refraction for the crystal in the respective axis. With two different indices of refraction for the principal axes, we essentially have a tunable waveplate. As we change the birefringence of the crystal, the light's phase is shifted in turn. Due to polarization misalignments, etalons, and other factors, EOM's in use as phase modulators also introduce residual amplitude modulation (RAM).

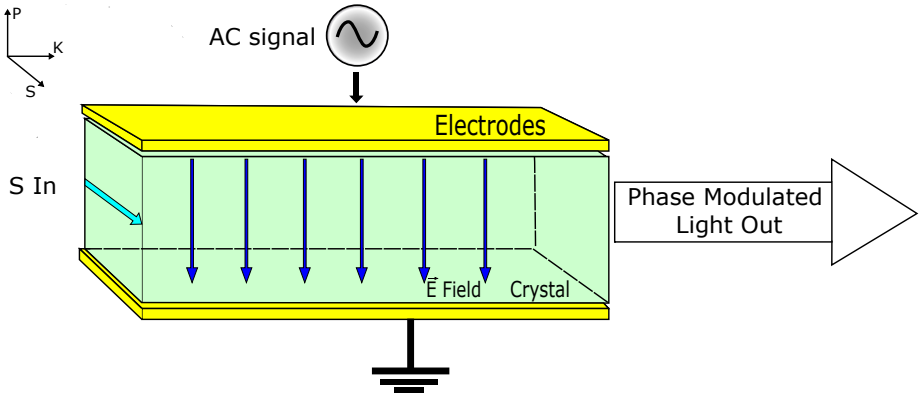

Figure 1: Experimental Set-up

1.3 Pound-Drever-Hall and Frequency Stabilization

The purpose of PDH is to lock the frequency of a laser to a cavity, or vice versa. This is usually done by measuring the power being reflected from the cavity and generating a

feedback signal to control the frequency of the laser or the length of the cavity. For light to resonate in the cavity its frequency must be an integer multiple of the cavity's free spectral range $\Delta\nu_{fsr} = c/2L$, where c is the speed of light and L is the length of the cavity. When the laser frequency is near resonance, the amount of light being reflected can be approximated as a linear function of the laser's frequency[1]. Phase modulation creates upper and lower sidebands, which when far from resonance can be treated as seeing the maximum reflection from the cavity. The photodetector sees the power from the superposition of these sidebands and the reflected carrier from the cavity. This power is converted to a voltage and then used as an error signal for locking.

Any RAM introduced by an EOM is fed through the experiment and is picked up on the photodetector. After the photodetector and the local oscillator's signals are mixed, this RAM introduces an offset on the signal being fed to the locking servo. This unwanted offset means the system is not at the frequency it expects. So instead of locking the laser's frequency to resonance, we unintentionally lock off-resonance reducing the effectiveness of the system. Any changes in RAM over time would change this offset uncontrollably leading to unpredictable behavior from the system.

2 Objective

Our goal is to use an EOM as a phase modulator and create an ultra-fast feedback mechanism by controlling the temperature and DC bias of the EOM to actively suppress RAM to the level of 1×10^{-5}. This will be done by using a free space EOM to phase modulate a beam, which will be partially picked off and sampled by a photodetector for use as our feedback signal to the temperature and DC bias control. The DC bias component influences the in-phase response 20 times more than the quadrature response, but the temperature control influences the quadrature response 5 times more than the in-phase response[2]. Working together, it should be feasible to have good suppression of both components over an extended period.

3 Approach

3.1 Mathematical Motivation

From our photodetector the amplitude modulated current for a given modulation frequency ω is given by

$$I(\omega_m) = -sin(2\beta)sin(2\gamma)|\epsilon_0|^2 J_1(M) \times sin(\omega t)sin(\Delta\phi + \Delta\phi_{DC})[2].$$

Where β andγ represent polarization misalignments of the half-wave plates from the crystal's axes, ϵ_0 is the amplitude of the incoming laser field, $J_1(M)$ is the first order Bessel function with M being the difference in modulation depth between the ordinary and extraordinary axis, and $\Delta\phi + \Delta\phi_{DC}$ represents the crystals natural phase shift summed with the phase shift induced by the bias voltage. We can see that when this sum is zero or an integer multiple of π the *sin* term will become zero yielding no RAM.

3.2 One Possible Set-Up for Experiment

The set-up involves using feedback mechanisms for temperature and DC_{bias} suppression of RAM in an EOM being used as a phase modulator. As reported by [2] the in-phase RAM is much more effective at suppressing through the DC_{bias} voltage, and the quadrature component is better suppressed by the temperature control. The PD acts as our RAM signal which we then decompose into the in-phase component to feedback to the DC_{bias} port of the EOM, as well as a quadrature component which is used to control a heater in thermal contact with the EOM. The DC Servo's output is fed through a high voltage amplifier to reach the DC bias needed for RAM suppression in free space EOM's, which is usually a couple of hundred volts.

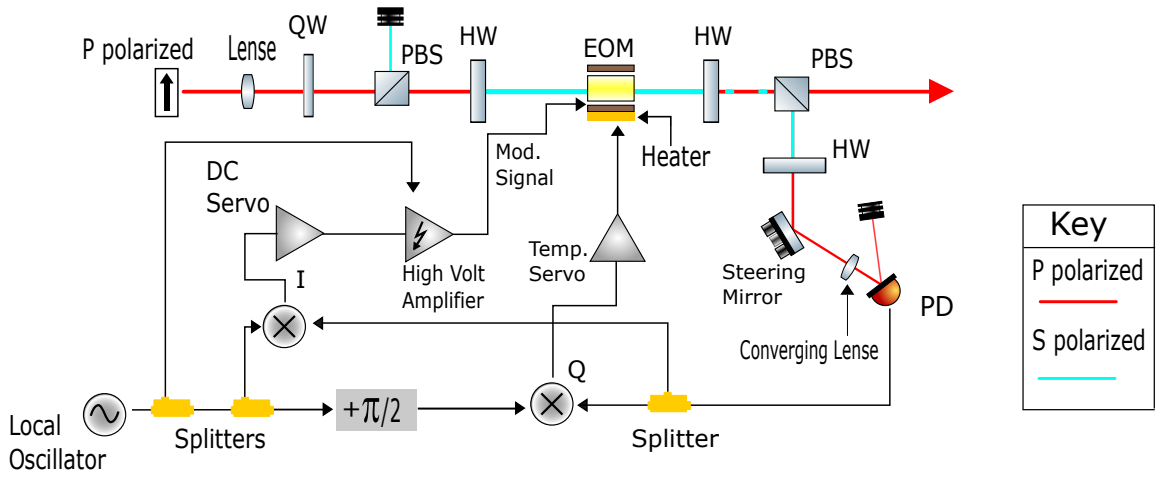

Figure 2: Experimental Set-up

4 DC Bias-Related Measurments

4.1 Bias Tee Characterization

The first step in being able to introduce DC bias feedback was to characterize the bias tee used to see the attenuation of the DC and RF signals into each port of the device. This was done by using the network analyzer to measure the S matrix for the device. Because the DC port of the device did not have any connector attached, I soldered a female BNC connector to the attached to the network analyzer as shown in the picture.

4.2 Fitting to Simplified Model

I made 4 plots from this data, three showing the data for each port individually and one showing the transfer functions that will apply when using the bias as expected. The first three plots gave the signals read from the two alternate ports when the signal is injected

Bias Tee with Soldered BNC connector

into the respective port. The last plot is the coupling between the RF and RF+DC ports as well as between the DC and RF+DC ports, which characterizes the transfer functions we expect to see in practice. While the device has multiple capacitors and inductors to perform the biasing, I used the S32 and S31 transfer functions to fit the electrical schematic as one capacitor and one inductor, shown below.

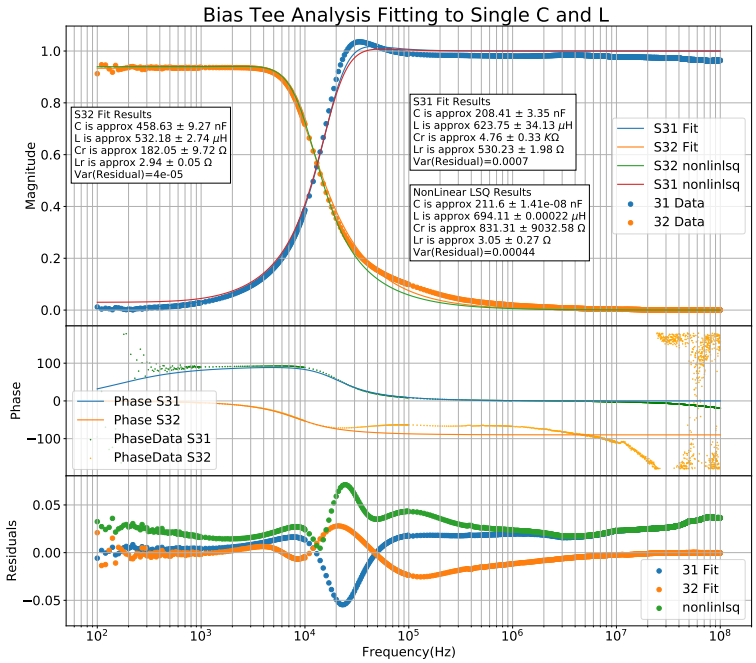

Fit of Data to Single Capacitor and Inductor

4.3 Modeling

Once we had realistic values for the capacitor and inductor, I used LTspice to add a model of the bias tee to the EOM driver and see what value of a tunable inductor was needed to maintain the resonant circuit at 37 MHz and 36 MHz for two different drivers. After confirming we had the correct inductor I soldered a new EOM driver board for this purpose. Here we see the left driver is complete, while the right one has no inductor or connectors.

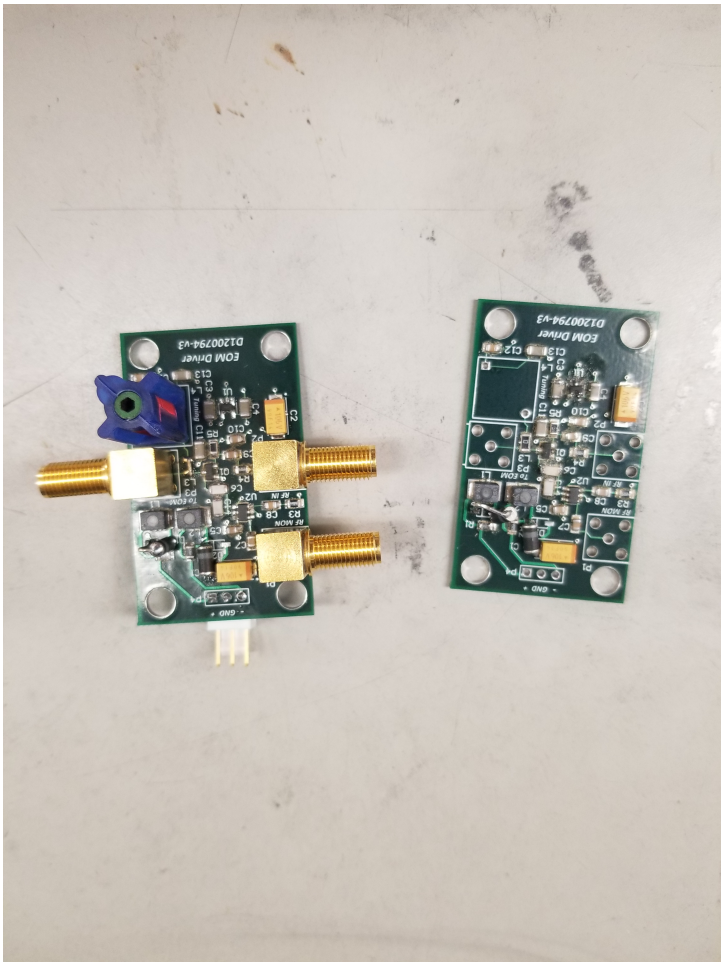

EOM Drivers

4.4 EOM Driver Resonant Circuit Update

There was a good possibility the tunable inductor I found would not be in the correct range to have a resonant circuit around 37 MHz. This turned out to be correct, as a replaced the 891nH - 1270nH inductor with a 460nH-788nH inductor to finally achieve the correct peak location. To take the measurement I attached a 50-ohm termination to the bias tee's DC port, and directly connected a 20pF 'Dummy EOM' to the RF+DC port which matches the EOM expected capacitance. I used a network analyzer to measure the transfer function through the bias tee and EOM system while tuning the inductor to move the peak to 37 MHz. This resonant circuit will ensure no higher-order harmonics of our modulation frequency will create unwanted modulation at higher frequencies. The transfer function from the RF Input of the EOM Driver to the EOM has a max gain of around 10, this gain will be enough to get a 40 Vpp signal which equates to a 0.3 modulation depth.

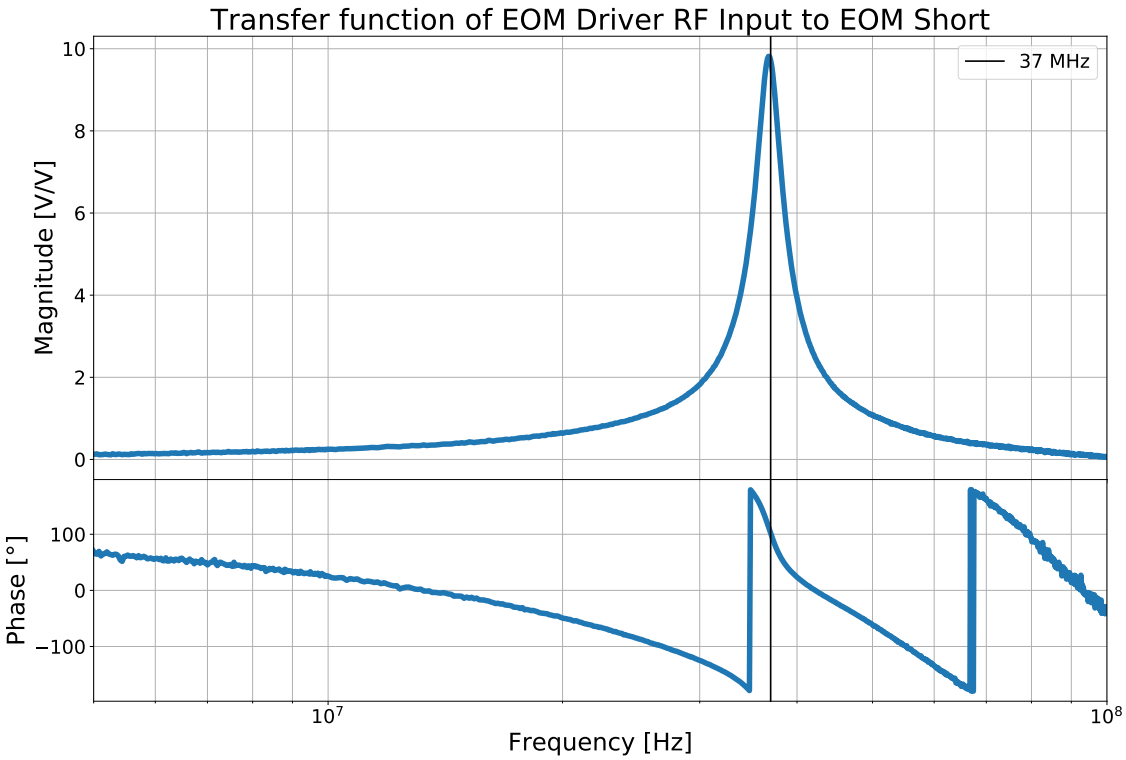

Resonant Circuit Transfer Function

4.5 RAM Measurement Set-up

The eventual goal of the project to actively suppress the RAM in EOM's. The first step in this process on the optical table is to set up some sort of pick-off by which we can measure and correct for the RAM present. We used the PBS within the faraday isolator (FI) as a pickoff, and rotate the HW plate at the input to control the power reflected. I then used a mirror and a focusing lens to steer the picked-off beam onto the active area of an 1811 new focus photodetector. Using the DC coupled output I maximized the output on an o-scope to get the best possible alignment. I then used the AC output port to measure the spectrum of the signal. I saw a peak at 14.75 MHz from a previous EOM, as well as its second harmonic at 29.5 MHz. The third peak is the RAM due to the 37 MHz EOM which we focus on.

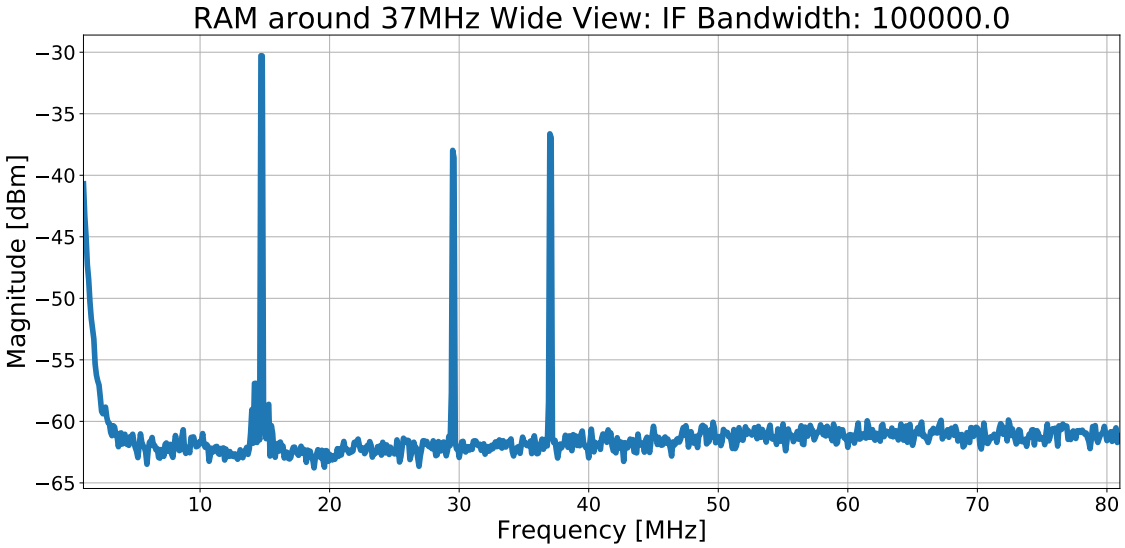

Spectral Density of signal from Photo-Detector

4.6 Reduction of Gain in EOM Driver

After the implementation of the EOM driver with the bias tee onto the input of the EOM, we did some preliminary analysis of the amplitude of the AC voltage reaching the EOM. For our set-up, we wanted to maintain a modulation depth of 0.3 which requires an amplitude of 20 V. Without the introduction of the bias tee this amplitude was attainable using the EOM driver present in the set-up and described before.

With the introduction of the bias tee, we noticed that the value of the monitor port at 37 MHz was around -3.49 dBm which corresponds to a signal with an amplitude of 6.4122 V. This was converted to the signal seen at the EOM by changing the signal at RF monitor to voltage, multiplying by sqrt(2) and dividing by the value of the simulated TF of the RF monitor port to the EOM short at 37MHz which is about 0.0328. This back of the envelope calculation showed the amplitude is less than the expected 20 V from a 16.5 dBm, and not enough for our desired modulation depth.

This motivated us to take a more robust measurement of the voltage seen by the EOM as a function of the input power. A power sweep from -5 dBm to 15dbm reveals the gain limitations of the EOM driver with the Bias Tee in series. As we can see, at 15 dBm we only see an 8 V signal at the EOM.

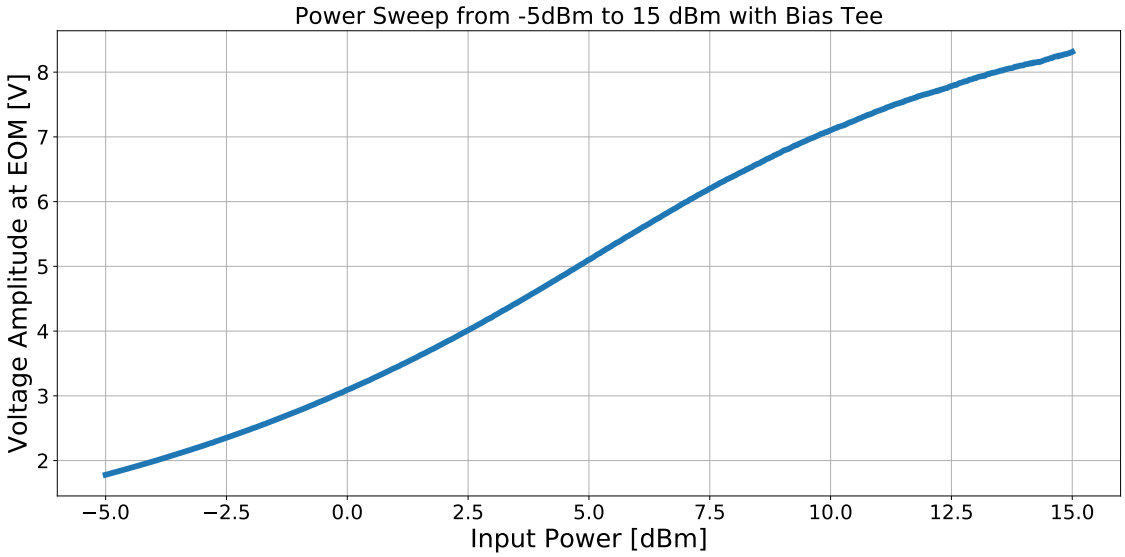

Power Sweep of Driver with Bias Tee and Dummy EOM

4.7 First Stage of Preliminary EOM Driver

Due to the power-bandwidth limitations of the current EOM driver, we set out to build a new driver for use with the bias tee. Fig.9 below is an LTspice diagram of a preliminary design for the first stage of the driver.

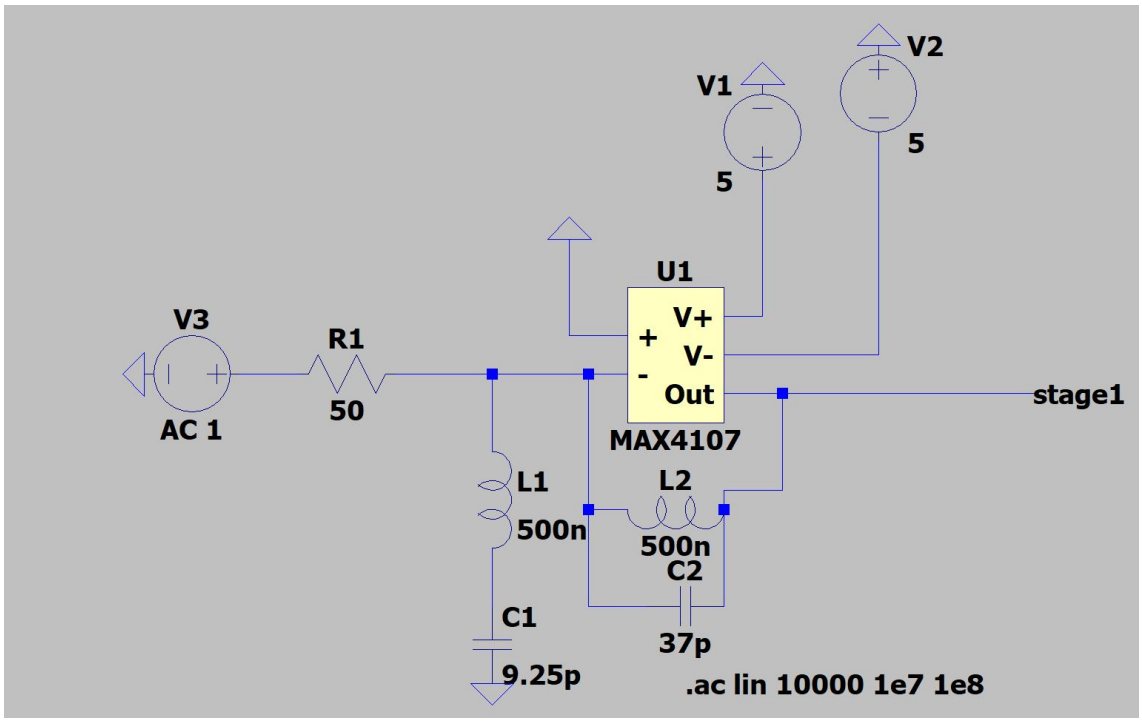

Figure 3: Preliminary EOM Driver 1st Stage

This stage acts as a resonant circuit creating a peak at the modulation frequency of 37

MHz. We designed a valley at 74 MHz to suppress the second-order harmonic of 37MHz. This is only the first stage as the EOM driver needs to be able to reach peaks of 20 V. The MAX4107 has a slew rate of 500 V/μs and delivers a 300MHz -3dB bandwidth. The purpose of using an op-amp in the first stage of the driver was to isolate the rest of the driver from any changes in impedance due to the bias tee introduction. We planned to find a transistor gain stage to place after the MAX4107 output which has a range of $\pm$ 200 V with a gain of around 8.

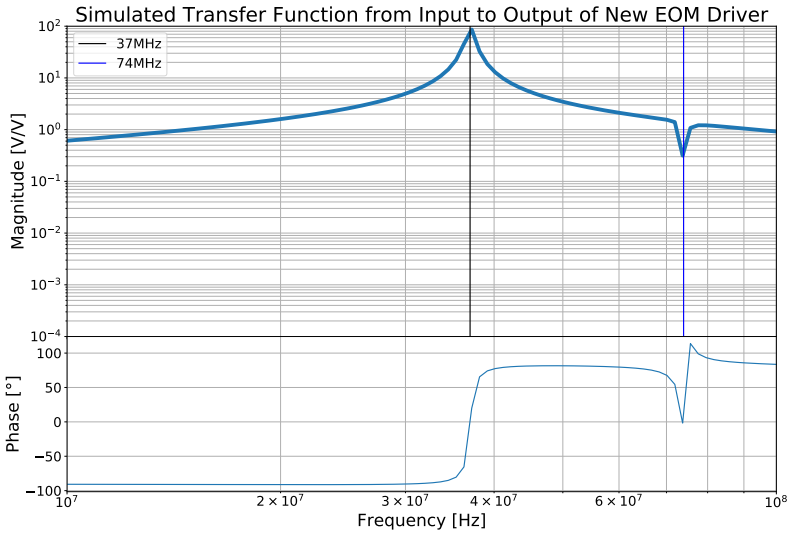

Figure 4: 1st Stage Zero Simulated Transfer function

5 Temperature Sensor Related Measurments

5.1 Reflow Solding

While soldering I learned of a new form called reflow soldering. This involves using a paste that has solder beads suspended in a flux, which when heated evaporates and fuses the solder. To spread awareness of how this process is done, I made a little step by step tutorial within the ELOG and linked a youtube video of the fusing process. After learning reflow soldering I used this technique to solder 3 channels to the AD590 temperature sensor board below, changing one resistor to bring it more in the room temperature range.

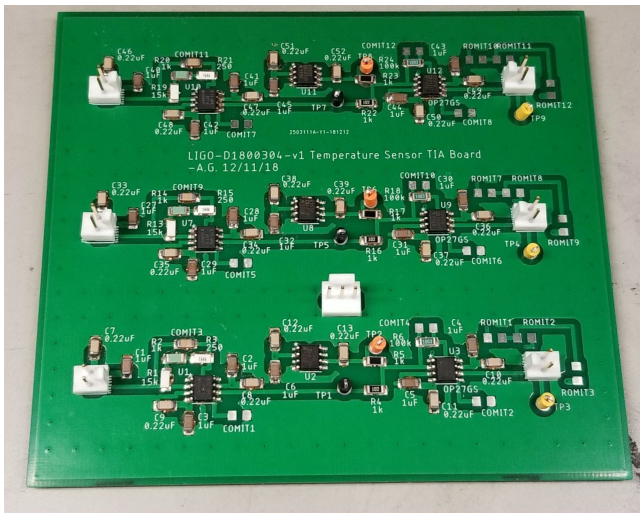

AD590 Temperature Sensing Board

5.2 AD590 Temperature Sensing Results

The first step in understanding the noise from AD590 temperature sensors was to use Zero to model all the known noise contributions present at the output. I was able to estimate the noise in terms of the input-referred current noise density by dividing the output noise by the overall transfer function. Then knowing the AD590s conversion from temperature change to current output, I found the simulated input-referred temperature noise density of the circuit as a whole from the sensor to computer acquisition.

From there I set up an experiment that measures the time series of three AD590's in thermal contact with a large metal block. The large metal block has a high heat capacity so it should be insensitive to small local temperature fluctuations, yielding the lowest noise power spectral density as possible due to the environment. I measured the input-referred noise of the ADC I was using by shorting the input and measuring the output. Because the ADC has a gain of one this output noise is the input-referred noise. I also measured the noise due to the 15V regulated power supply used to power the circuit. I used the power noise times the power supply rejection ratio of the op-amps in the circuit to estimate the input-referred noise of the power supply as well. I took the difference of the time series for each pair of AD590

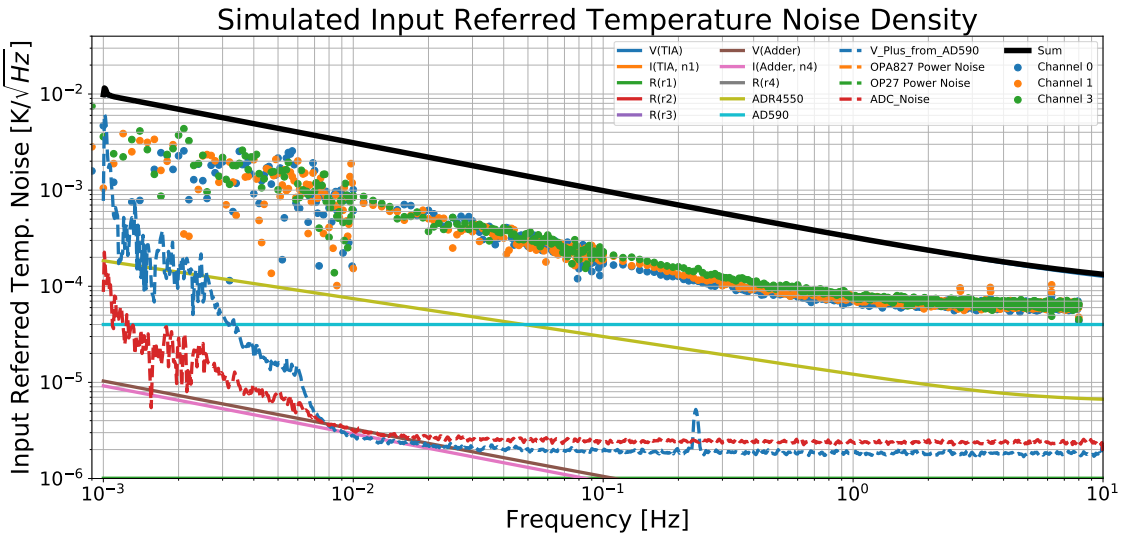

AD590 Input Referred Temperature Noise Density

sensors to get differential data, which should not contain any common temperature-induced fluctuations.

My first assumption was that the signal received from each detector was of the form $S(t) = T(t)+n+S_0$ where S is the signal, t is time, T is temperature, and n is the inherent stochastic noise of the device. This assumes all channels have the same transfer function but different offsets, so when taking the PSD of the difference of channels, it has the sum of the noise of both devices in quadrature. The difference PSDs, for instance, the difference between channels 0 and 1, then takes the form

$$N_{01} = n_0^2 + n_1^2$$

By adding difference PSDs with the channels we want and subtracting the other we get the noise of a single device squared. Taking the square root we can pull out the individual noise spectral density for each channel, this is also known as the three corner hat method. Plotting the estimated noise sum and the measured individual noise of each device on one graph we were able to directly compare our expected noise with reality. I found the measured noise was below the simulated noise for all frequencies between a millihertz and ten hertz. By integrating this real noise over the bandwidth of a millihertz to ten hertz, we found the total noise of the device is always under one millikelvin. This means the circuit is sensitive enough to resolve this level of temperature change up to ten times per second, which matches our goal for the circuit.

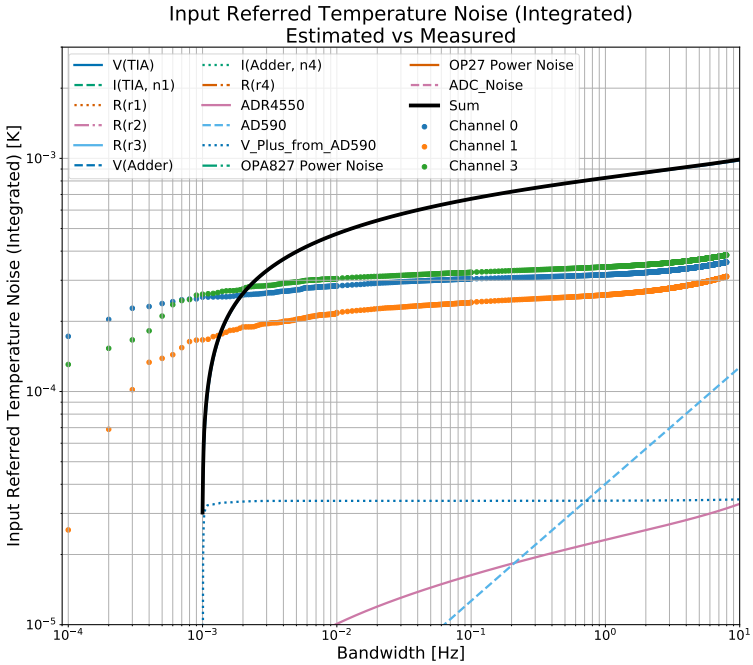

AD590 Integrated Temperature Noise Sum

5.3 Optimal Op-amp for Temperature Sensing

To determine what would be the best op-amp to use in the transimpedance amplifier of the AD590 temperature sensing circuit I ran a Zero simulation that calculated the noise sum of the circuit as a whole for different op-amp choices. A major consideration is we do not want the circuit's response to change with temperature in a temperature sensing circuit. Also since this is the location where the current information of the temperature given by the AD590 is converted into a voltage response, we want any noise introduced after this transimpedance stage to be much less than this stage. After the analysis, I found the Op-amp we were using (OPA827) had the lowest temperature drift, but AD8671 or LT1128 had an all-around lower noise density. Because we found the noise level of the circuit was a millikelvin sensitive, I thought OPA827 was a good choice to reduce the temperature dependence as much as possible.

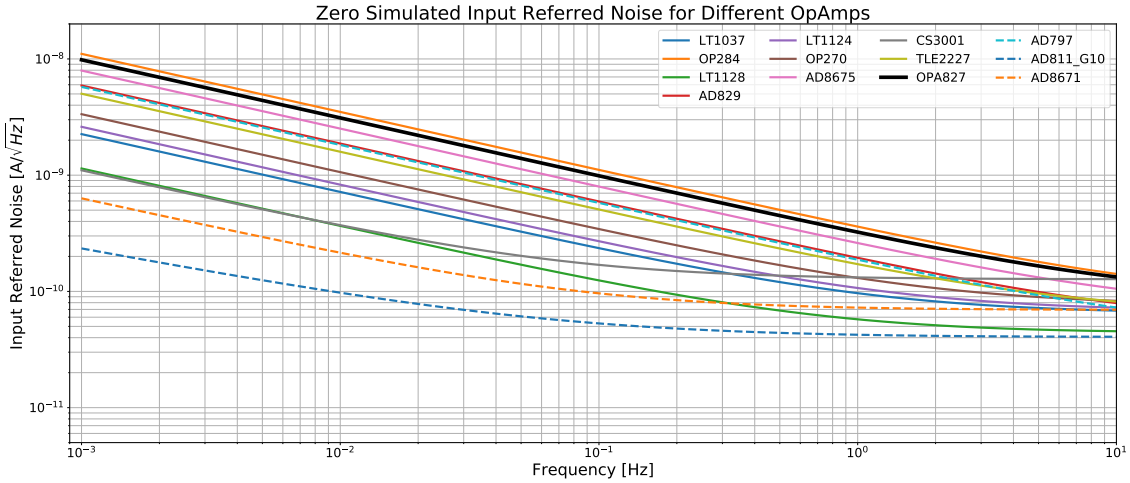

Noise Density of Different Op-Amps

5.4 Temperature Sensing with the board inside and outside of insulated box

A previous measurement showed a substantial difference in the noise level of the temperature sensing circuit when it was in the isolated box vs when it was not insolated (Fig.5.4). These measurements were taken on different days and showed the noise of the non-insolated configuration was lower than the insolated one which made little sense.

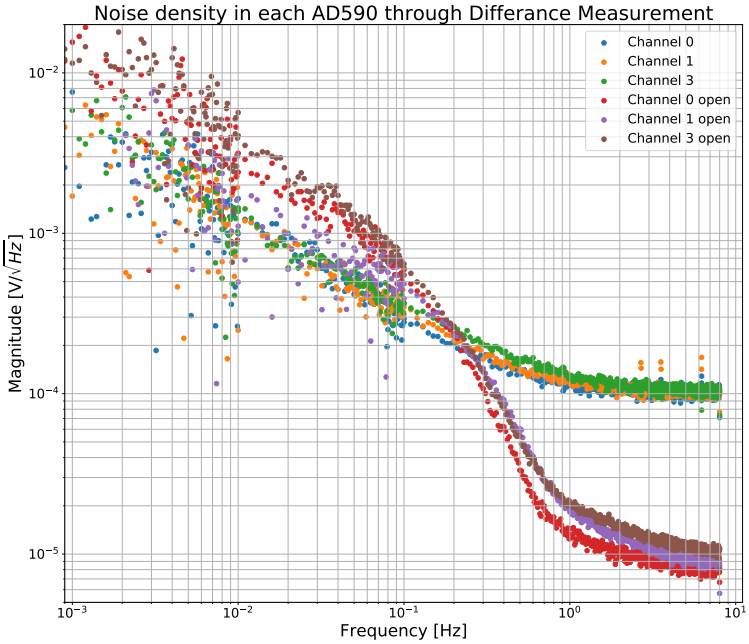

Open vs Closed ASD

To confirm or refute this discrepancy I re-ran the experiment with both configurations over one day. Fig.5.4 conclusively indicates the previous measurement had something else going on we were not aware of. We now see there is no obvious difference in the amplitude spectral noise density of the open configuration vs the insolated configuration.

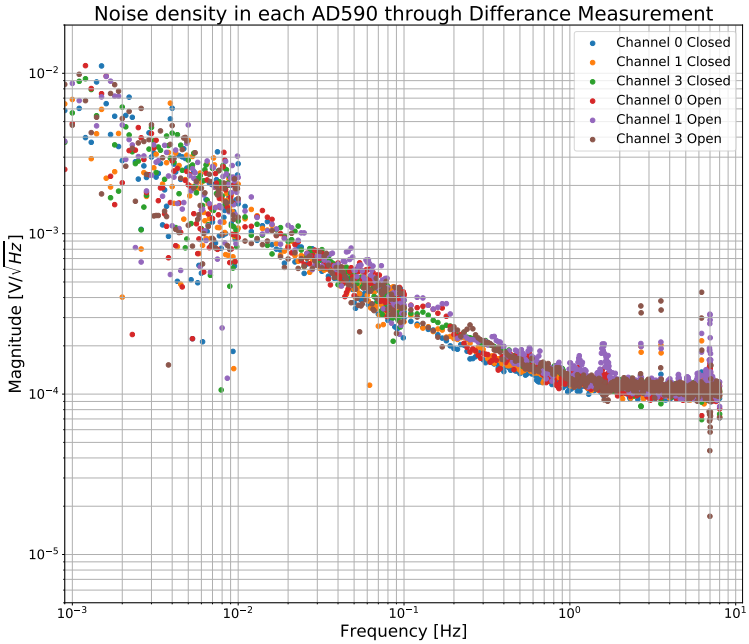

Open vs Closed ASD Pt.2

6 Conclusion

The next steps are to finish the design of and build the DC bias servo which will take the signal from the photodetector, downmix it with a local oscillator at the modulation frequency and then low pass filter it to yield a DC signal which is proportional to the amount of RAM seen. This DC value will then be fed through a high voltage amplifier to get into the value of hundreds of volts, which is the order of the V_π of the EOM. For the high voltage amplifier, the plan is to begin at a gain of 10 and investigate the open-loop gain to determine stability in feedback. After the DC Servo is complete we will need to build the temperature servo, which will require a much lower bandwidth, [2] had 0.1 Hz. We also need to work on finishing the new EOM driver for use with the bias tee. This will entail determining what transistor stage will be optimized for the high voltage gain as well as if the MAX4107 is the best op-amp available for the transfer function shaping.

References

[1] Black, E. D. An introduction to Pound-Drever-Hall laser frequency stabilization. Am. J. Phys. 69, 79 - 87 (2001).

[2] W. Zhang, M. J. Martin, C. Benko, J. L. Hall, J. Ye, C. Hagemann, T. Legero, U. Sterr, F. Riehle, G. D. Cole, and M. Aspelmeyer, "Reduction of residual amplitude modulation to 1 x 10-6 for frequency modulation and laser stabilization," Opt. Lett. 39, 1980-1983 (2014)