Working Towards Finding An Upper Limit For Crackle in LIGO's Maraging Steel Blade Springs

Ben Levy
Research Adviser: Eric Quintero

Abstract Crackle is a type of nonlinear, high frequency noise generated as individual molecules in a material shift when the entire system is driven at low frequencies. Such unwanted noise may occur in components of LIGO's gravitational wave observatories. Components of specific importance in studying this phenomenon are the steel blade springs that hold the observatories' highly sensitive mirrors in place. These blade springs were the primary focus of this investigation of crackle. For the project, an optical apparatus was created to determine the upper limit for the intensity of crackle that could occur in such springs. Several improvements were made to the apparatus for the purpose of amplifying the crackle, while decreasing the effects of other background noise sources, and a detailed data simulation was created for use in testing the analysis process. Using a demodulation technique to analyze preliminary data, I found that in the worst case scenario, crackle could cause displacement noise of up to $1.84 \times 10^{-14} \frac{m}{\sqrt{Hz}}$ against a background displacement noise of $4.00 \times 10^{-13} \frac{m}{\sqrt{Hz}}$. Longer data collection runs will be needed in order to make more conclusive statements. Ultimately, this study will help LIGO to ascertain whether provisions will be necessary to mitigate the effects of such nonlinear noise in its observatories.

LIGO-T1300839-v1

1 Introduction

1.1 Background

LIGO, or *Laser Interferometer Gravitational-Wave Observatory*, is a large scale experiment run jointly by Caltech and MIT for the purpose of finding evidence for the existence of gravitational waves. LIGO observatories in Hanford, Washington and Livingston, Louisiana are large Michelson interferometers with perpendicular arms each four kilometers in length. Ultimately, the project hopes to observe changes in the apparent length of one arm relative to the other - an indication that a gravitational wave, propagating through the fabric of space-time, has passed through the earth. This can be achieved by observing the interference pattern created at the output of the interferometer. A small change in that pattern could indicate a change in the distance light has traveled to get to the detector[1]. A simple representation of a Michelson is shown in Figure 1. Here, light is emitted from the laser, shown by the heavy black arrow. The beam splitter subsequently splits the beam so that fifty percent of the power travels into each arm of the apparatus. The light then reflects off of the end mirrors and returns to the beam splitter. Fifty percent of the power in each arm is transmitted into the photo detector, while the other half travels back in the direction of the laser. The interference pattern is measured at the photo detector where half of the beam power is from each arm.

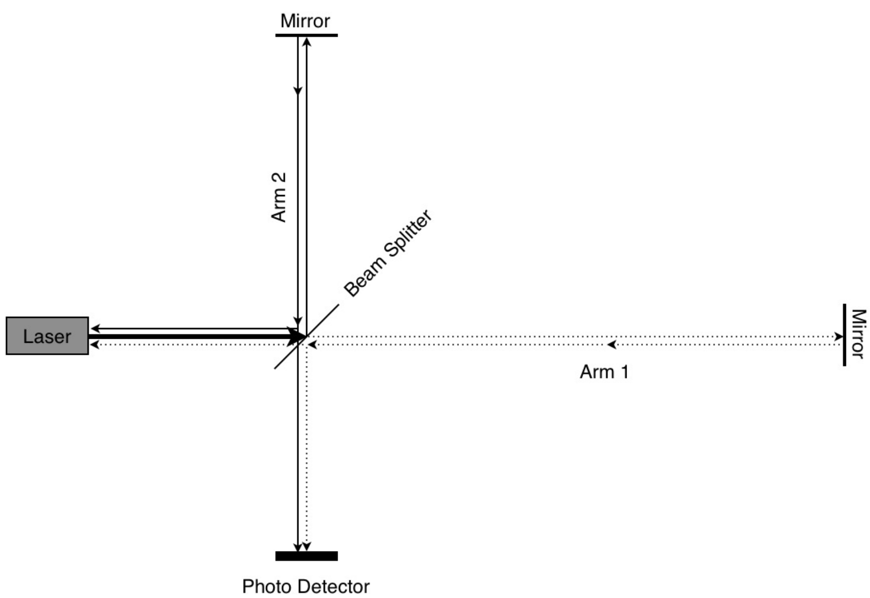

Fig. 1 Schematic drawing of the basic path that light travels through a Michelson interferometer. The beam splitter separates incoming light into two perpendicular beams of equal power. This basic design is used both in the LIGO observatories and in my project.

Despite this clever design, the first iterations of LIGO were not able to detect gravitational waves. Because large sources of gravitational waves (such as colliding binary stars) are so rare, detectable events are likely to be located far from the earth. This means that their waves will be exceedingly weak and difficult to detect against the diverse background of noise sources that constantly bathe our planet. In order to be sure that measurements made at the observatories actually indicate the detection of gravitational waves, LIGO scientists and engineers must isolate the observatories as much as possible from all conceivable sources of noise. For example, they are isolated seismically to reduce the effects of earthquakes and passing trains.[1] Some types of noise can be dealt with in this manner, while others, such as quantum shot noise, are completely unavoidable. This project investigates a type of noise that has not yet been studied thoroughly.

1.2 Crackle Noise and its Importance

One of the many types of noise that may appear in LIGO's signal is "crackle noise." Crackle is a common phenomenon. It occurs when a system responds to slowly changing external conditions (for example, a driving force) through discrete, impulsive events spanning a broad range of sizes.[2] Importantly, crackle is classified as "nonlinear noise" because there is no simple relationship between the slow driving and the discrete events. Though these effects can manifest in systems ranging widely in areas such as superconductors, stock markets and rearranging foam bubbles, one simple example can be demonstrated with paper. It is nearly impossible to bring two ends of a sheet of binder paper together without producing audible crackling sounds. Even if the paper is moving at extremely slow speeds (i.e. slowly changing external conditions), it will be difficult to eliminate these sounds.[2] This type of disturbance is of importance to LIGO because high-frequency noise caused by the slow movement of Earth's tectonic plates could yield crackle in LIGO's suspensions, which would cause them to vibrate and add unwanted noise to the signal. Right now, the effects of crackle on LIGO are not well understood, so it is important that we find an upper limit for this type of noise due to seismic oscillations in our interferometers.[3] Crackle could of course occur in many components, but we have chosen to concentrate on examining its effects in the maraging steel blade springs from which the end mirrors are suspended at the observatories.

1.3 Blade Springs

LIGO's blade springs are a part of the complex system of masses, springs and pendula that work together to isolate the end mirrors from outside noise. They are machined from maraging steel, a highly flexible, light, and crack resistant alloy that can be found in applications from aerospace to fencing blades. As is shown in Figure 2, the mirrors are hung as the fourth pendulum in a series which ultimately is suspended by the blade springs at the very top. In fact, these springs exist at every level of the suspension. Even though crackle may occur in other components of these suspensions, it was most convenient to investigate crackle in the blade springs, so they were chosen for experimentation.[3]

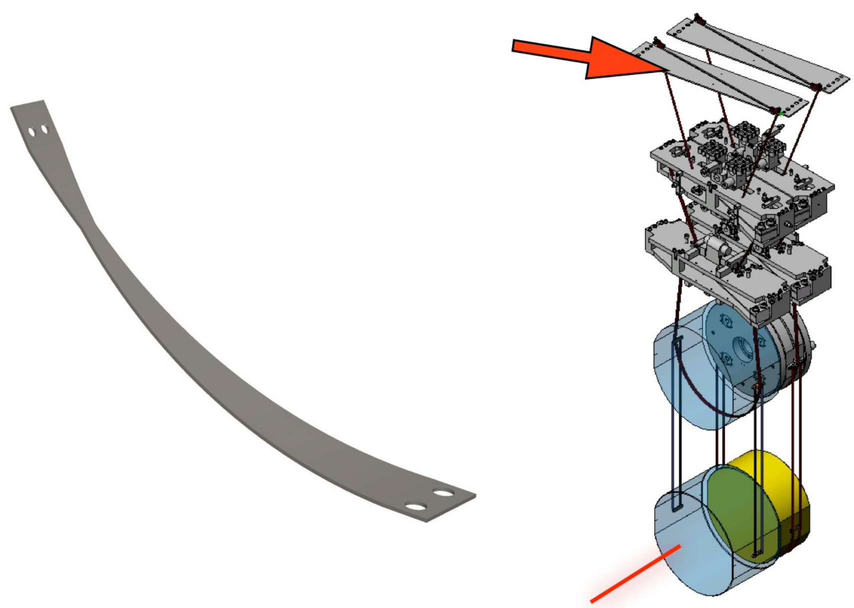

Fig. 2 The narrow metal band on the left is a computer model of one of LIGO's maraging steel blade springs. They are baked into this bent position so that when the end mirrors of the quadruple pendulum suspension are hung from them, they flatten out, and stress is more evenly distributed across their surface. The assembly to the right of the blade shows the suspension with four of the blade springs at the top.

After careful study of these blade springs, a tabletop Michelson interferometer was ultimately created in such a way that it could be driven sinusoidally. It was proposed that by analyzing data taken from the interferometer's photo detector outputs, we would be able to determine the maximum amount of crackle that could be present in two blade springs when they were driven. Undriven test runs would act as a control. Although we were not able to take enough data before the end of the summer session at Caltech to come to any strong conclusions relating to LIGO's observatories, I have successfully calculated a preliminary upper limit for crackle in our signal. This work has also increased our understanding of the analysis pipeline that will be utilized once we have more data.

2 Basic Experimental Overview

Crackle noise is such a low-level phenomenon, that it would be impossible to detect directly without optics on the scale of LIGO's observatories.[3] In order to make such measurements feasible in a lab, we set up an environment in which we knew what our crackle signal should look like, and then we used this information to simplify the detection process.

On a basic level, we measured the differential displacement between two test masses each hung from a blade spring. Crackle is a random phenomenon, so when these two blade springs were driven sinusoidally we hypothesized that discrete random events in the two blades would cause differential displacements which could be measured. In order to actually make this measurement, we used a Michelson interferometer as described in section 1.1 with the two hanging test masses acting as the end mirrors. We set up the optics such that each photo detector (denoted as PD in Figure 3) gave information about one blade spring. When the two signals were subtracted, we had a data stream of the differential displacements between the blades. Because any seismic noise that affected both blades equally was simply subtracted away in this process, along with all other such "common mode" noise, what we were left with was displacement caused by events such as crackle that occurred in just one blade or the other.

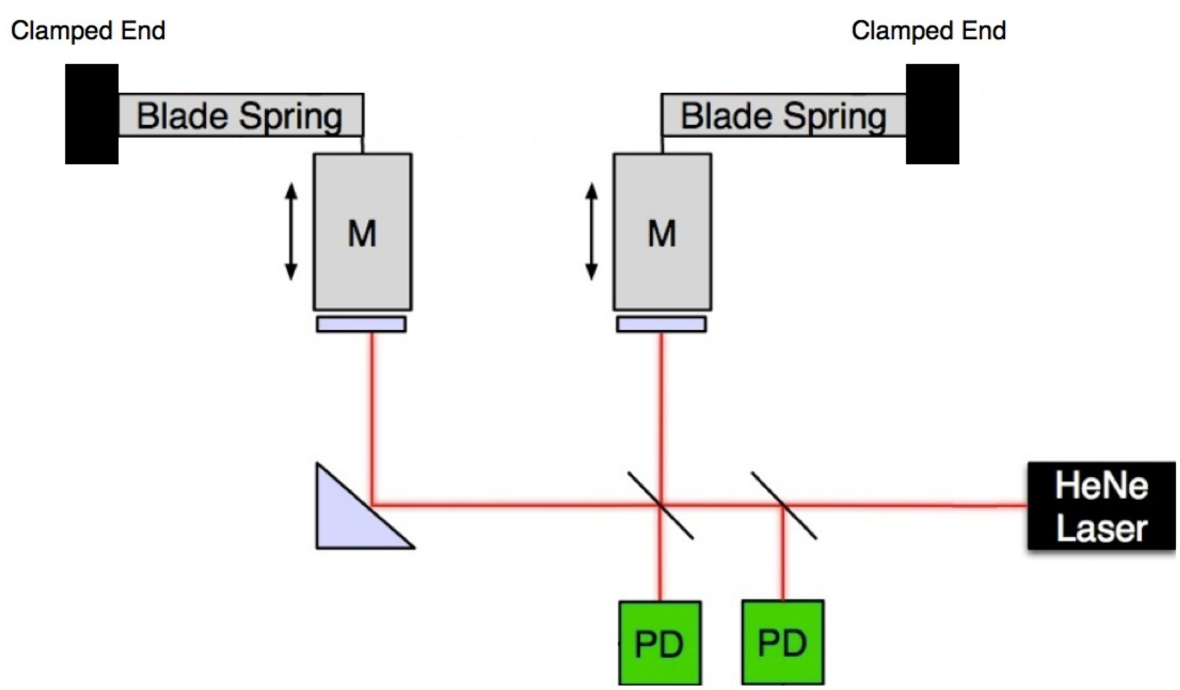

Fig. 3 Basic experimental setup showing two masses, (M), each hanging from the ends of blade springs. The masses are also the end mirrors of a Michelson interferometer, which is used to measure the differential displacement of the masses. Such displacment could be caused by crackle noise.

Crackle is not the only type of random noise that could manifest in just one blade or the other. Other types such as quantum shot noise constantly bathe the photo detector output signals. In order to make sense of this jumble, we use the fact that the amount of crackle present in a system is proportional to the driving force. We drove our table sinusoidally so that the magnitude of the crackle would also oscillate sinusoidally as shown in Figure 4. We then used a demodulation technique to pick out this distinctive signal from the background.

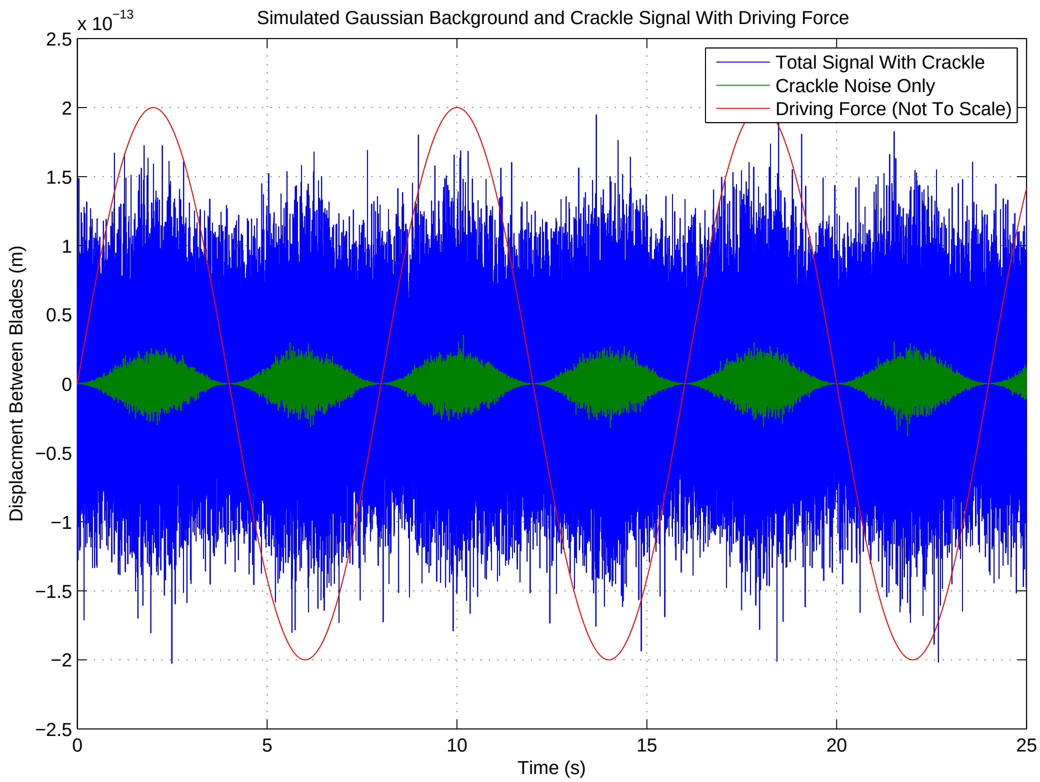

Fig. 4 This plot shows a Gaussian background noise signal including crackle noise similar to what we might detect with our experiment. The sinusoidal red trace is a simulated driving force (shown with a greatly reduced amplitude for clarity), and the green trace is the resulting crackle signal that is usually hidden inside the blue trace. Note how the crackle varies with the driving force.

3 Results

Using the methods outlined above, I developed a data analysis scheme to find the amount of crackle noise in the $400 - 500Hz$ region of our signal. My code utilized the Student's t-test to find a noise spectral density interval. We can be 95% confident that the crackle's noise spectral density lies on this interval. Before the end of the Caltech summer session, we were able to take 2 hours of data with the drive on, and 1 hour with the drive off. The confidence interval for this data is:

$$[-6.84 \times 10^{-15} \frac{m}{\sqrt{Hz}},\ 1.84 \times 10^{-14} \frac{m}{\sqrt{Hz}}]$$

This can be compared with a mean background noise spectral density in the $400 - 500Hz$ region of:

$$4.00 \times 10^{-13} \frac{m}{\sqrt{Hz}}$$

Though the interval is very large (so large, in fact that it encompasses zero, the case where there is no crackle noise at all), it does provide us with an upper limit for the amount of crackle that could possibly be occurring in the blade springs. According to this data we can be 95% sure that no crackle will occur above $1.84 \times 10^{-14} \frac{m}{\sqrt{Hz}}$. Figure 5 shows a plot of the displacement noise spectral density of the total signal with the upper limit of crackle shown as a dashed line.

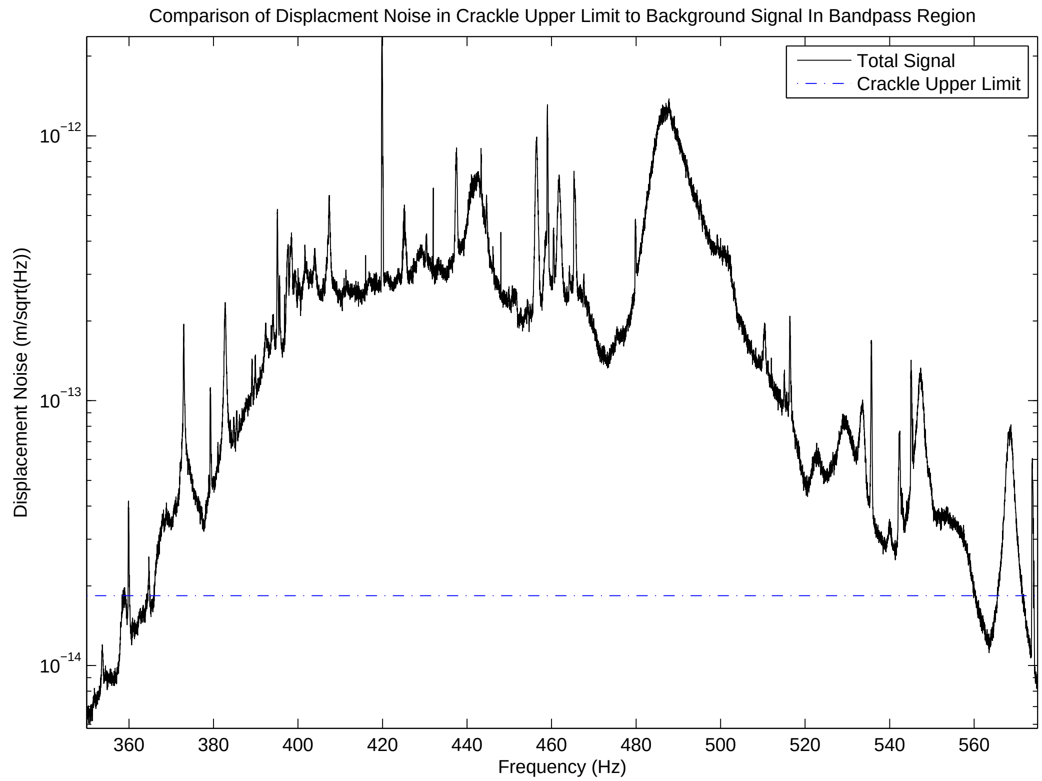

Fig. 5 This plot shows the displacement noise spectral density of the total signal along with the upper threshold for crackle noise that I calculated to be $1.84 \times 10^{-14} \frac{m}{\sqrt{Hz}}$. The data represented has been bandpassed in the $400 - 500Hz$ region.

This experiment is still underway, so there is much more work to be done in the future. Perhaps most importantly, we need to take more data. Currently our confidence bound is on the order of 10^{-14} which is enormous when compared with the actual magnitude of crackle that we expect to see. More data will allow for longer integration times which will in turn allow for much more certain results. However, longer data collection runs will also require more optimization of the systems that stabilize the test masses against common mode vibrations. Currently, vibrations such as heavy footsteps have the potential to “unlock” the masses, ruining the data stream.

Another area that needs investigation is the scaling of our results in this small lab experiment to LIGO’s larger blade springs at the observatories. It is not currently obvious if crackle will occur at the same levels in a spring with a greater volume, and if not, how to appropriately scale quantities so that we can extrapolate to behavior on the full suspensions. Depending on the results of this experiment, it may also be necessary to examine additional components of the quadruple suspension for crackle. The blade springs have the least damped connection to the ground of any part of the suspension, meaning that they would likely be most susceptible to crackle induced by seismic activity. If it were found that crackle levels in the springs could actually be problematic, it would be important to test components that are further down the suspensions and closer to the mirrors. It is crackle at the mirrors - not at the blade springs - that would actually hinder the detection of gravitational waves. The welded connection between the silica fibers and the end mirror could prove to be an area where further study is needed.

4 Methods

In this section I will give a more in-depth description of the experimental design. I will also discuss at length some of the areas of this project on which I spent the most time while working for LIGO.

4.1 Experimental Design and Data Collection

4.1.1 Physical Set-Up

We began with a vacuum chamber on an optics table. In the vacuum chamber we placed two blade springs each with masses hanging from the ends. To the bottoms of the masses, we attached mirrors which acted as the end mirrors for the two arms of a Michelson interferometer similar to, but much smaller than the ones at the heart of each of LIGO's observatories. The input to the interferometer was a 1064*nm* HeNe Laser fed into the chamber by a polarization-maintaining fiber optic cable. As shown in Equation 3, we used the difference between the outputs from the symmetric and asymmetric ports of the interferometer, sampled 4096 times a second with photo detectors, as a control signal.

$$V_{sy} = dx + IntesityNoise + ShotNoise_{sy} \tag{1}$$

$$V_{asy} = -dx + IntesityNoise + ShotNoise_{asy} \tag{2}$$

$$control = V_{sy} - V_{asy} = 2dx + ShotNoise_{sy} - ShotNoise_{asy} \tag{3}$$

In Equations 1 and 2 I have represented the outputs from the symmetric and asymmetric Michelson ports (V_{sy} and V_{asy} respectively) as the sum of any differential displacement between the blades (dx and $-dx$: the crackle signal), intensity noise (one of several common mode noise sources that may be present), and quantum shot noise (an example of a random, unavoidable phenomenon). When we subtract the two outputs, we amplify the crackle signal, dx, while eliminating common mode fluctuations such as intensity noise.

4.1.2 Optimization of Photo Detector Gain for Data Collection

Before actually collecting data, I had to deal with a potentially problematic aspect of the subtraction discussed in the previous section. In reality, minor misalignment of the Michelson could lead to V_{sy} and V_{asy} being scaled by different factors. That is, these two outputs were actually multiplied by unknown gain coefficients (G_{sy} and G_{asy}) that could be different. Equations 1 and 2 could more accurately be written as follows:

$$V_{sy} = G_{sy}(dx + IntesityNoise + ShotNoise_{sy}) \tag{4}$$

$$V_{asy} = G_{asy}(-dx + IntesityNoise + ShotNoise_{asy}) \tag{5}$$

If $G_{sy} \neq G_{asy}$, then taking $V_{sy} - V_{asy}$ would no longer completely eliminate the intensity noise as it was eliminated in Equation 3. To remedy this, I wrote code that would find the best coefficient, C by which to multiply V_{asy} in order to eliminate any gain imbalance. To this end, I made use of mathematical coherence. Ideally, $V_{sy} - V_{asy}$ and $V_{sy} + V_{asy}$ should be as dissimilar (or "incoherent") as possible. The sum should contain no crackle and a great deal of intensity noise, while the difference should contain crackle but no intensity noise. All I had to do was find C in Equation 6 such that coherence between the sum and the difference was minimized.

$$V_{asy} = C \times G_{asy}(-dx + IntesityNoise + ShotNoise_{asy}) \tag{6}$$

Figure 6 shows the mean square coherence functions of an ideal and an imbalanced pair of simulated sum and difference signals.

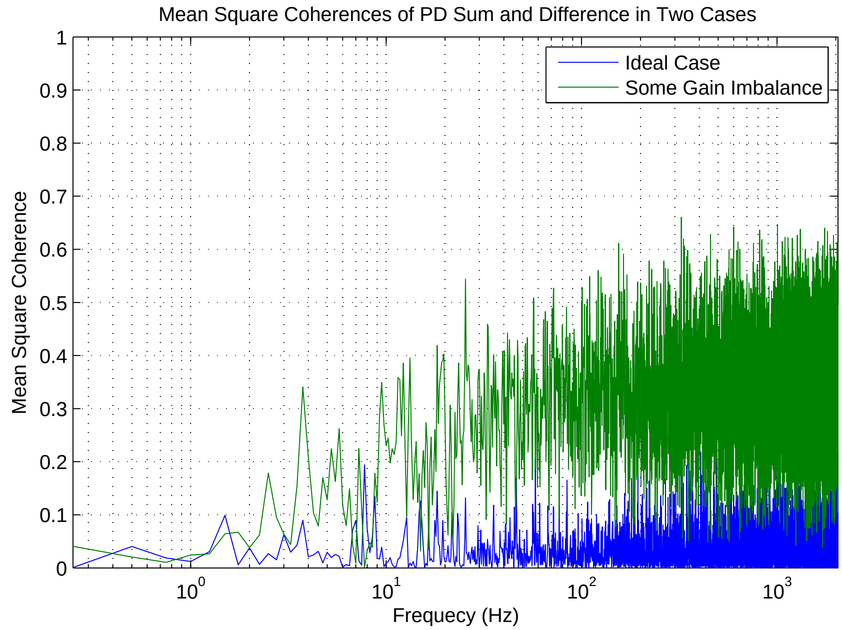

Fig. 6 The top, green trace plots the mean square coherence of the sum and difference signals where V_{sy} and V_{asy} have unbalanced gain coefficients. The blue, bottom trace shows the ideal case where there is no imbalance and the signals are as incoherent as possible. Both signals are from Gaussian distributions.

I then used Newton's root finding method to find the value of C for which the integral of the coherence for the imbalanced signal was least. By applying this coefficient to V_{asy} and re-plotting the coherences of the two signal pairs, I obtained the plot in Figure 7.

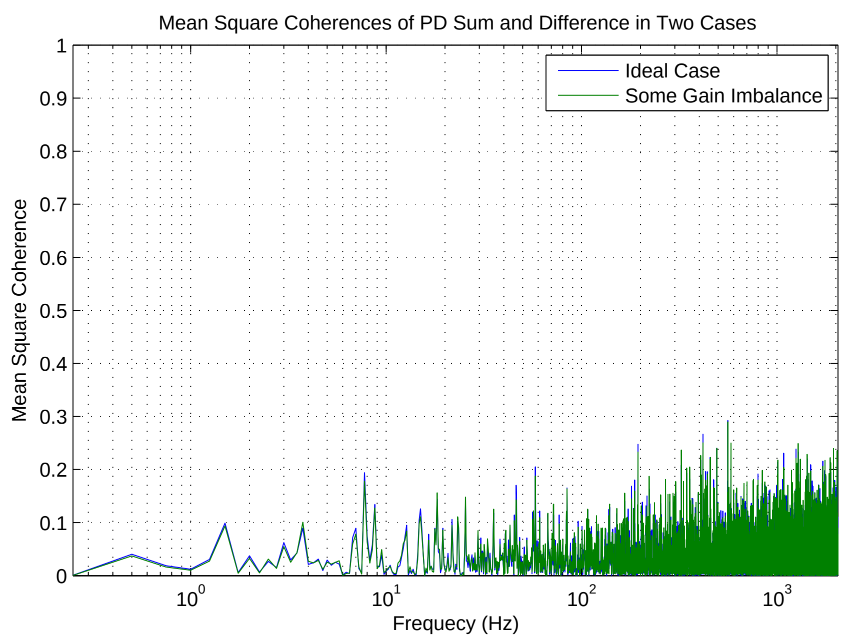

Fig. 7 This is the same plot as in Figure 6, but with V_{asy} for the imbalanced case being multiplied by the optimal coherence-reducing coefficient C.

We took minute-long V_{asy} and V_{sy} data streams and ran them through this code. We then used the coefficient that was returned as a new gain coefficient for the V_{asy} signal in the control software. With this in place we took driven data for 2 hours, and undriven data for 1 hour during my time working on the experiment. Although many days of data will be needed before we can hope to make any strong assertions about the nature of crackle, this was enough data to allow me to begin analysis. The last step before I could begin was to bandpass the $400-500Hz$ region of the signal. This was the region in which we had the most sensitivity and it also represented the most sensitive (and thus most susceptible) region of the LIGO observatories.

4.1.3 Background Noise Damping

Several steps were taken to mitigate the effects of seismic and acoustic vibrations. Except for the laser, the entire experiment was housed within a vacuum chamber which shielded it from air currents and acoustic vibrations, and reduced the likelihood of photons in the laser beam scattering off of air molecules. Additionally, the optics table inside the chamber was mounted on two levels of springs and masses, which reduced the effects of vibrations traveling through the floor.

A multifaceted active damping scheme was also employed. The control signal from the Michelson was used as part of a digital feedback loop in which coil and magnet actuators attached to the ends of each blade spring, locking the masses to a definite differential displacement. This system compensated for most common mode noise while still allowing potentially interesting, isolated events such as crackle noise to create differential displacements between the blades. An additional feedback loop was necessary in order to bring the blades into the vicinity where they could be "locked" by the computer. This system used shadow sensors as input, and actuated using the same coils and magnets as the main feedback loop. Once the Michelson had achieved a lock, this damping was turned off so as not to inject noise from the shadow sensors into the signal. This entire setup is shown as a block diagram in Figure 9, a more complete version of Figure 3. Note the presence of two feedback loops - one using the Michelson output, and the other using the shadow sensor output to drive the coil and magnet actuators.

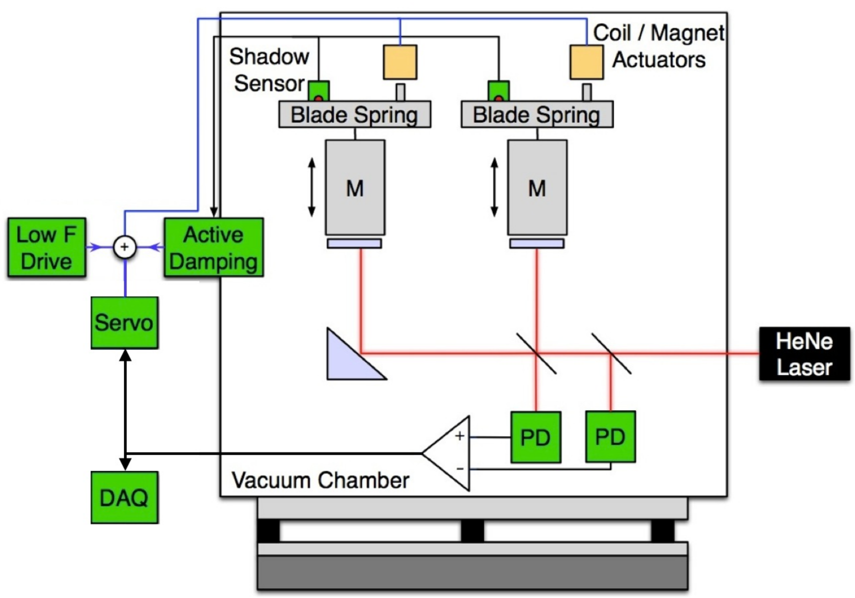

Fig. 8 This block diagram shows the Michelson interferometer at the center of the crackle noise experiment housed inside a vacuum chamber. The signals from the symmetric and asymmetric port photo detectors are subtracted and both used as an input for a digital active feedback loop, and as the data output from the experiment. Shadow sensors provide a coarse level of damping and the signal output provides a finer level, both using the coil and magnet actuators. The low-frequency sinusoidal drive used to induce crackle noise is also applied using these actuators.

4.1.4 Shadow Sensor/Actuator Mount Mechanical Redesign

One area of the experiment that seemed particularly susceptible to vibrations was the mounting mechanism for the shadow sensors and the actuators. When I first began working on the experiment, this sub-assembly was made largely of thin pieces of metal, mesh, and circuit boards. Using this system, we risked inducing inaccuracies that could have hindered the performance of the precision actuators and shadow sensors. Additionally, adjusting the vertical position of the shadow sensor diodes was not easy - a problem because future experimentation would likely call for frequent realignment of the Michelson. I set out to design a new mount that would be both ridged and adjustable.

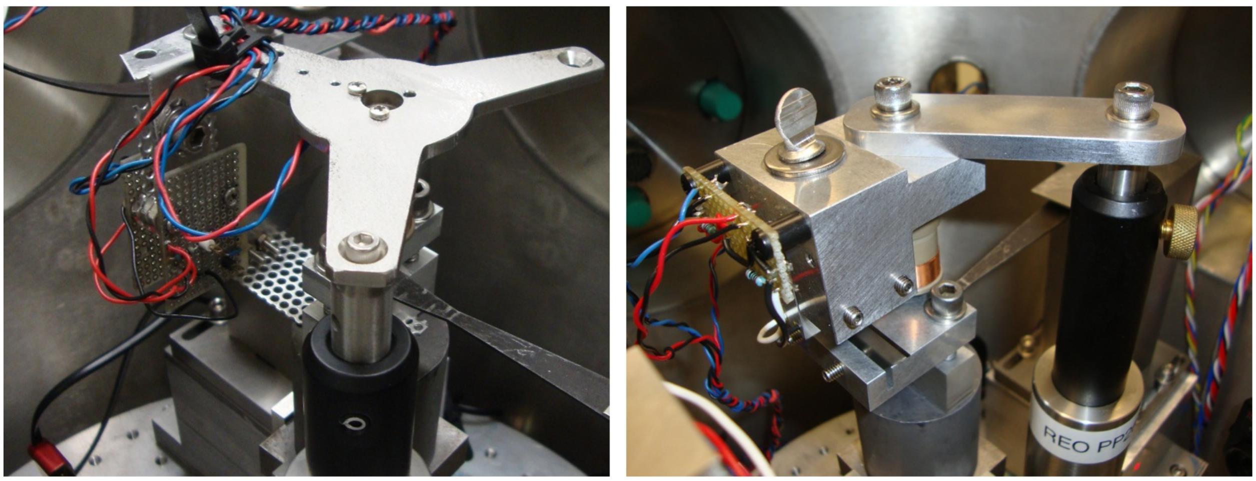

Fig. 9 The lefthand image shows the shadow sensor mount in its previous state with extensive use of aluminum mesh and circuit boards for structural components. The righthand image shows my redesigned mount with stronger parts and a nob for fine adjustments.

I used Solidworks to design a new assembly made from four aluminum parts. After they were machined, I assembled them and soldered the new circuit boards. The product was a much more robust, adjustable mounting apparatus for the coil and magnet actuators and the shadow sensors.

4.2 Data Analysis

4.2.1 Demodulation

I used a demodulation scheme that leveraged our crackle noise's sinusoidal dependence to detect crackle even when it was buried deep within the much larger total noise signal. As demonstrated in Figure 4, the magnitude of crackle is proportional to the driving force. There is also a component of crackle hypothesized to be dependent on the driving jerk, but it is smaller than force-crackle, and I did not include it in my analysis due to time constraints. If we name our displacement due to force-crackle dx_f, and the driving frequency ω_d, we can see that the crackle portion of our control signal ($x(t)$) could be represented by the expression in Equation 7.

$$x(t) \propto 2cos(\omega_d t)dx_f \tag{7}$$

We can square $x(t)$ and, disregarding a DC offset that appears due to the identity $cos^2(x) = \frac{1}{2} + \frac{1}{2}cos(2x)$, we get:

$$x(t)^2 \propto 2cos(2\omega_d t)dx_f^2 \tag{8}$$

At this point we can note that $x(t)^2$ is either positive in the presence of crackle, or zero when no crackle is present. Keeping this fact in mind, we define the quantities Q and I as follows:

$$Q \equiv cos(2\omega_d t)x(t)^2 \tag{9}$$

$$I \equiv sin(2\omega_d t)x(t)^2 \tag{10}$$

Finally, we plug $x(t)^2$ into Equations 9 and 10 and rewrite them using trigonometric identities.

$$Q = dx_f^2 + cos(4\omega_d t)dx_f^2 \tag{11}$$

$$I = 2sin(2\omega_d t)cos(2\omega_d t)dx_f^2 \tag{12}$$

As can be seen in Equations 11 and 12, when there is no crackle signal, dx_f^2 is 0, so both I and Q will equal 0. If however dx_f^2 is not 0, Q will exhibit a DC offset, while I will continue to have an average value of 0. If values of I and Q were recorded over long periods of time both with the drive on and off, we would expect to see the average driven Q value converge to an offset from the average undriven I, Q, and driven I values. The displacement of Q corresponds to the amount of displacement due to crackle in the system (in m^2).

4.2.2 Confidence Bounds

I then used the Student's t-test to find the 95% confidence interval for the average Q value offset. The upper bound of this confidence interval represented the highest level of crackle consistent with our data to 95% certainty. Finally, I converted the confidence bounds of Q to more useful units by first multiplying by the $100Hz$ bandwidth to get power spectral density (in $\frac{m^2}{Hz}$), and then I took the square root to get noise spectral density in $\frac{m}{\sqrt{Hz}}$.

4.2.3 Simulated Data Analysis

In order to check that my demodulation schemes worked, and to gain a better feel for the manner in which they worked, I also developed code to analyze the detection capabilities of the demodulation scheme. To do this, I created a model displacement signal using a Gaussian distribution which could be adjusted to have a noise spectral density from 10^{-15} to $10^{-12}\frac{m}{\sqrt{Hz}}$. To this I added a model Gaussian crackle signal with the crackle represented as variations in the blade spring constants proportional to the driving forces. The crackle signal was multiplied by a coefficient that I called α. To see the full simulated signal, refer to Figure 4. Then, beginning with short integration times of just $600s$, I generated vectors of I and Q values and calculated the minimum α value necessary to generate a level of crackle distinct enough from the background distribution that we could be 95% confident in its detection. As can be seen in Figure 10, detection became more difficult as the background noise spectral density increased. These results matched expectations, and lent credence to the effectiveness of my code.

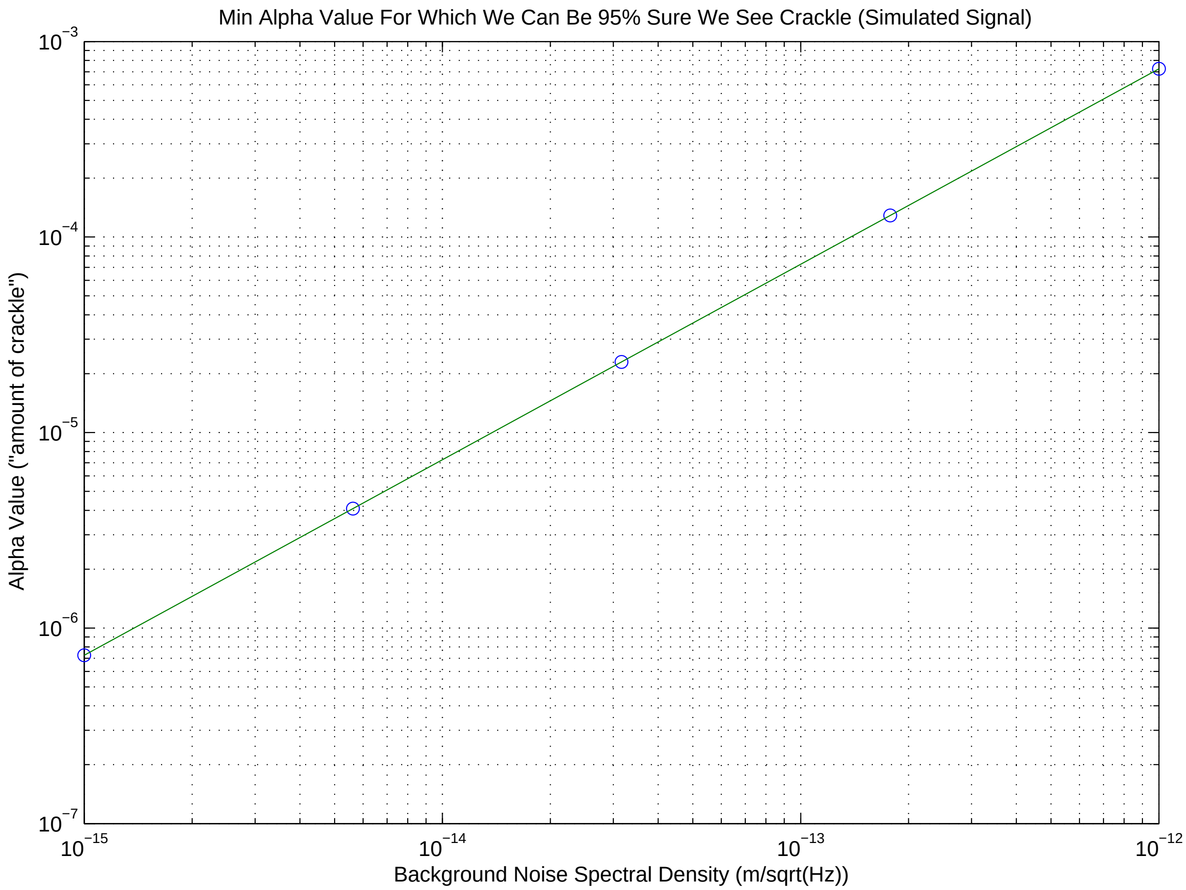

Fig. 10 This plot shows that as the noise spectral density of my background signal increased, it became increasingly difficult to detect crackle noise. A 600*s* integration time was used for this simulated, Gaussian datastream.

5 References

1. "Science of LIGO." Science of LIGO. N.p., n.d. Web. 13 May 2013.
2. Sethna, James P., Karin A. Dahmen, and Christopher R. Myers. "Crackling Noise." Nature 410, 2001: 242-50.
3. Quintero, Eric, Eric Gustafson, and Rana Adhikari. "Experiment to Investigate Crackling Noise in Maraging Steel Blade Springs." Tech. no. LIGO-T1300465v1, 2013.

6 Acknowledgments

I would like to thank my adviser, Eric Quintero for his patient and illuminating support this summer. I would also like to thank Rana Adhikari and Koji Arai for their guidance. Finally, I wish to acknowledge the National Science Foundation and LIGO for giving me this opportunity and for funding my research.