T0900083-v1

Modeling Cable Noise in LIGO

Julian E. Freed-Brown*
Harvey Mudd College
(Dated: October 24, 2008)

The Laser Interferometer Gravitational-Wave Observatory (LIGO) has test masses and optics hung as the bottom mass on simple pendula in order to isolate them from the motion of the Earth. In Advanced LIGO, optics will be hung on multistage pendula to increase isolation. However, cables need to be attached to sensors and actuators on some of these pendula, which can potentially degrade the vibrational isolation and change the dynamics of the system. To understand the effects of the cabling, we used *Mathematica* to model the cables as collections of small masses connected by springs. With our model, we can quickly and easily look at how the cables affect the pendula, once we have experimentally determined a few of their physical parameters.

With the upgrade from Initial LIGO to Enhanced and Advanced LIGO, the aim is to increase the sensitivity of the detector by a factor of 10, meaning Advanced LIGO should have a displacement sensitivity of around 10^{-19} m $\mathrm{Hz}^{-1/2}$ at 10 Hz, with better sensitivity at higher frequencies. The fundamental sources of noise in Initial LIGO are shot, seismic, and thermal noise, but Advanced LIGO is designed to be limited at all levels by quantum noise. Because of this ambitious noise floor, much time and effort is put into analyzing and minimizing other noise sources. For example, the Earth is constantly vibrating on the order of a micron. These vibrations need to be removed from the system.

To eliminate this problem, LIGO's test masses and optics are suspended on pendula. A pendulum acts as a natural low-pass filter: the pendulum is affected by low-frequency inputs, particularly at the resonant frequency, but is isolated from high-frequency inputs . LIGO's suspensions take advantage of this effect to isolate the motion of the masses from the motion of the Earth in the frequency range of interest (10 Hz to few kHz). When additional isolation is needed, multistage pendula are used.

In particular, a double pendulum suspends the Output Mode Cleaner (OMC), a piece of equipment that removes noise from light as it leaves the interferometer. The OMC has a suspension designed to isolate it from external inputs (see T060257-03-R[1]). However, electronic cables are needed which attach from the frame of the suspension to the OMC. This has negative consequences because it may short out the suspension at some frequencies, since the cables bypass the pendula and loosely tie the OMC to the ground. As of yet, the magnitude of this problem is unknown. Some previous research has been done to see the effect of the cabling on the OMC, but it has mostly been preliminary work to demonstrate that the cables could have a considerable effect[2, 3].

The goal of our research was to experimentally determine various cable parameters, such as the quality factor and effective spring constant of the cables, and update Mark Barton's *Mathematica* model[4] of the suspensions to include cabling. My Co-SURF, Chihyu Chen, did the experimental work and I focused on the modeling.

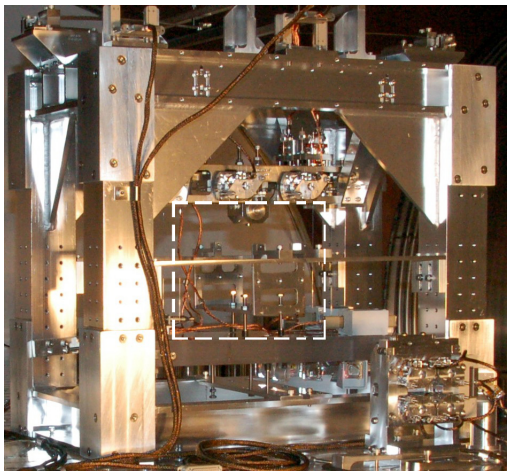

FIG. 1: The OMC suspension installed at the Livingston Observatory. The OMC is mounted on a bench that hangs as the ultimate mass of a double pendulum. From this picture, we can see the cabling that attaches to the bench (boxed).

MODEL CONCEPT

The model starts by taking a series of parameters that define the cable: stiffness, damping, mass, length, and the attachments points. It then takes the cable and "cuts" it into n segments. These segments are approximated as rigid bodies connected by springs. Then, the model calculates a potential function, E_P, and kinetic energy function, E_K, for the cable. By minimizing the potential function, the model finds the equilibrium position of the cable.

Once the kinetic and potential functions have been updated, Dr. Barton's model calculates the vibrational modes and makes transfer function and thermal noise plots using a stiffness matrix, K, where

$$K_{ij} = \frac{\partial^2 E_P}{\partial x_i \partial x_j}.$$

x_i and x_j are the position of the pendulum and cable segments. Consider a classic potential energy equation

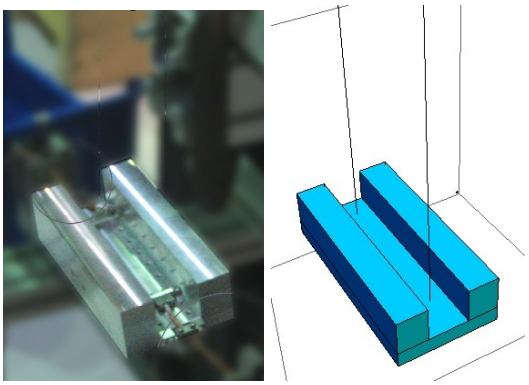

FIG. 2: The test pendulum we have been doing experimental work on (left) and the pendulum given by our model (right).

for a simple spring potential, where

$$E_P = \frac{1}{2}k(x)^2,$$

then there is a one by one stiffness matrix,

$$K = \left[\, k \,\right].$$

This result indicates that the matrix behaves as a matrix of spring constants. Similarly, there is a mass matrix, M, where

$$M_{ij} = \frac{\partial^2 E_K}{\partial \dot{x}_i \partial \dot{x}_j}.$$

These matrices can be used to find the vibrational modes of the system. This is accomplished by solving

$$Ke_i = \omega_i^2 M e_i$$

for the eigenmodes, e_i, and the angular eigenfrequencies, ω_i.

Additionally, the model calculates damping, dissipation dilution, and matrices representing equations of motion and coupling between the different degrees of freedom. From this, it produces transfer function and thermal noise plots.

A more rigorous description of the processes can be found in LIGO-T020205-02-D[4]. Of course, all of these calculations are done automatically in the model, but a general understanding of what the model is doing gives a motivation for the structure of the model for the cables.

This formulation provides critical checks to see if the model is working. For example, the equilibrium position is only based on the attachment points and cable stiffness. Hence, the equilibrium position predicted by our model should have a natural looking arrangement that can be compared to the actual cables. Furthermore, the qualitative behavior of the cable should not change as n is increased. In other words, the approximation with $n = 10$ and $n = 5$ should have differences, but they should exhibit the same general behavior. For the majority of our research, we used the predicted equilibrium position to evaluate the health of the model. Only now that we have a more advanced model of the cables functioning are we starting to compare quantitative results of the model, like resonant frequencies and transfer functions, to experimental data.

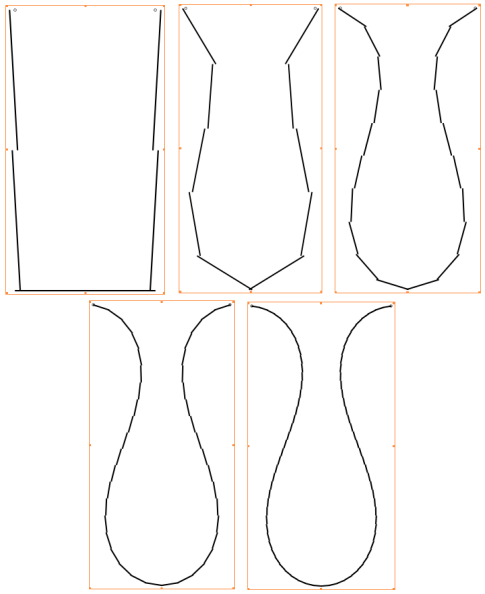

FIG. 3: Sample output of the two dimensional model. n increases by factors of 2 with each picture, starting with 5 (top left) and ending with 80 (bottom right).

TWO DIMENSIONAL MODEL

A two dimensional cable model was used to see if the model concept was valid. A two dimensional model is much easier to work with. In 2D, each cable element has three degrees of freedom: two translational degrees of freedom (an x and y coordinate) and a rotational degree of freedom. This contrasts with three dimensions, where there are six degrees of freedom. Here, there are three translational degrees of freedom (an x, y, and z coordinate) and three rotational degrees of freedom (yaw, pitch, and roll). This makes the 2D model much easier to imagine, so it is a good starting point.

Under this regime, the model still takes the same inputs as the 3D model and gives an equilibrium cable position as output. Using the simplified model, qualitative features of the model could be examined. In particular, did the cables had natural dressing? That is, did the cables generated by our model hang similarly to actual cables?

An example of the outputs of the two-dimensional model can be seen in Fig. 3. It should be noted that this model excludes gravity: the position is completely

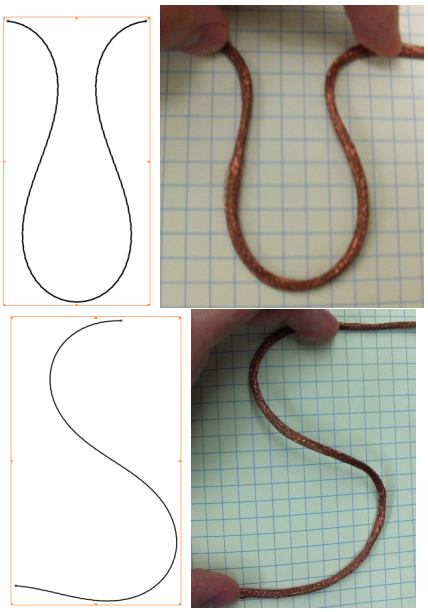

FIG. 4: Comparison between two dimensional model output and actual cables.

determined by springs between the pieces of the cable. The results of the 2D model were very promising. As n increases by factors of 2, the qualitative behavior of the cable does not change. Each cable has the same general shape. This indicates that increasing n only increases the resolution of our model. Hence, the model seems to approximate the same behavior every time, but when the cable is chopped into more pieces, the approximation gets better.

Furthermore, the modeled cables have similar equilibrium positions to cables held with similar fixed points (see Fig. 4). Obviously, they are not exactly the same, but the parameters used in the model were only rough estimates of the cable's length, stiffness and attachments points. However, despite these inexact values, they had very similar shapes and features.

THREE DIMENSIONAL MODEL

Adding Cables

The initial problem of adding cables to Dr. Barton's model reduces to adding a large number of springs to the model. However, there was already a way to add springs, using an object called springlist. To add springs, an entry is added to springlist that defines the springs attachment points, stiffness (using a stiffness matrix to represent six degrees of freedom and their possible coupling), and unstretched length.

To add cables, springlist was adapted into an object called cablelist. With a cable approximated as n segments and length, mass and attachment points specified in the model definition, a function generates cablelist so that there is a spring between every cable segment. When springlist is referenced to make a potential function, cablelist is also referenced so that the potential incorporates the stiffness of the cable. Later, cablelist is referenced again in the potential function so that a gravitational potential term is added for the cable elements. Since we are using the same equation we use for springs, we have that the potential between two cable elements is given by

$$\frac{1}{2}\begin{pmatrix}\Delta x & \Delta y & \Delta z & \Delta\theta_{yaw} & \Delta\theta_{pitch} & \Delta\theta_{roll}\end{pmatrix} K_c \begin{pmatrix}\Delta x\\ \Delta y\\ \Delta z\\ \Delta\theta_{yaw}\\ \Delta\theta_{pitch}\\ \Delta\theta_{roll}\end{pmatrix},$$

where K_c is the stiffness matrix between the cable segments.

To generate K_c, the stiffness of a length of cable is experimentally measured. This measurement is used to find effective spring constants for 5 cm of the cable in the six degrees of freedom. Then, the model uses those values to calculate spring constants for springs between the cable segments. This calculation assumes the springs in the model combine linearly and is based on the length of the cable and n.

It was also useful to adapt springlist into cablelist, because springlist already has damping built into it, which meant that we could add damping to the cable exactly as it was added for springs. Like the calculation of vibrational modes, there is a detailed description of how damping is incorporated into the model in LIGO-T020205-02-D[4].

Once cablelist was in place, the model just needed to be tweaked a little more for it to work. To start, the variables the model solves for needed to be updated. Before the update, the model found the coordinates of the pendulum at equilibrium. With the update, the model also solved for coordinates defining the cable. The kinetic energy function also needed to be changed to include the cable and the dimension of a few matrices needed to be changed. After these changes, the model computed everything it needed to. To finish the update, one last function was added that includes the cabling in the plots of the pendulum (see Fig. 5).

Small Angle Constraints

The improved model worked well for the most part, but did have some serious problems. When it was tested for various values of n, it did not exhibit consistent behavior. With certain values of n, the pendulum would find unnatural equilibria. Often, the cables would be connected

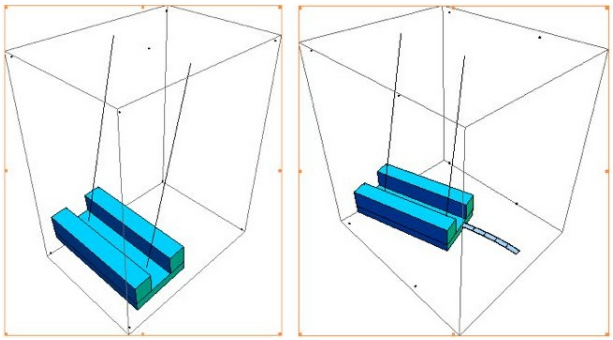

FIG. 5: Sample output of the model without cables (left) and with cables (right). Note that in this case, the cable is noticeably affecting the motion of the pendulum.

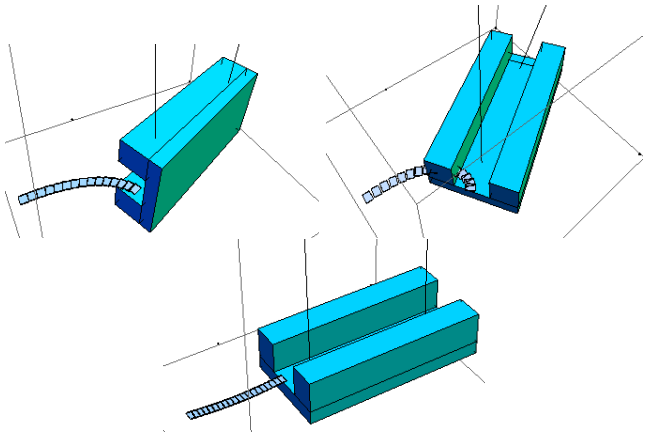

FIG. 6: (Top)Two typical cabling errors. The cables are attached to the pendulum at unnatural angles. (Bottom) The pendulum once the problem was fixed.

to their attachment points at 90 or 180 degree angles (see Fig. 6).

The 90 and 180 degree angles indicated problems with the potential functions for torsional springs. In particular, the potential function contained $\sin\theta$ and $\cos\theta$ in places where we did not think it should belong. This gives unintuitive results because it says the force applied by the spring will decrease if it is twisted more, finding a minimum 180 degrees. This was a serious problem for the cables. Unfortunately, it was not a problem that could be easily disposed of, because attaching three torsional springs simultaneously yields a very strange, nonlinear system which is very difficult to deal with.

Before adding cabling, there was not a problem with the potential because the pendulum calculations used a small angle approximation. Unfortunately, the cables often had large angles, so the potential function failed when it was applied.

To solve this problem, we added a feature where a maximum angle between two cable elements could be specified. This puts us into a sort of medium angle regime. The angle cannot be too large, or it cannot reach the unnatural equilibria at 90 or 180 degrees, but it is broad enough that it will not force the cable into a small angle approximation, which would not apply to our system.

Upon adding this constraint, the problem was immediately fixed.

Number of Modes

The model had a very inconvenient feature. Before cabling was added, all of the modes the model calculated were important, because they all involved the pendulum, and there were not very many of them. There is one mode for every degree of freedom in the system, so the double pendulum would have 12 modes.

When cabling was added, the degrees of freedom of the system increased dramatically. For the final modeling of the OMC, we used two lengths of cable and approximated each as 20 segments. This gave a total of 252 modes for the system. This is a slightly overwhelming number of plots to look at.

However, not all of the modes were significant. Most of the modes involve internal oscillations in the cable and are not actually important in the analysis of the pendulum. Hence, we wanted the model to look at the modes and only give us the ones that involved significant oscillations of the pendulum.

To do this, we made a function that takes each mode and looks at its kinetic energy. It would classify the mode by what elements were contributing a certain percentage of the kinetic energy. For example, if 95% of the energy was coming from the pendulum rotating, it would recognize the mode as a pendulum mode. However, if all but 1% of the kinetic energy is coming from the cabling, then it is probably a cable mode and not a pendulum mode. The function would look at all the modes, report the ones that are pendulum modes and automatically plot the ones we cared about. Also, since the mass of the pendulum is so much greater than the mass of the cable, pendulum motion is given much more weight in this analysis, since smaller motions cause larger kinetic energies. That is, it is unlikely to disregard significant pendulum modes. This is even less likely because the criteria the function uses to classify the modes is adjustable.

With the cabling, there were typically 13 pendulum modes for a double pendulum, regardless of the number of segments we cut the cable into. This is nice, because it was another qualitative check that the model was working well and, physically, it means the cables are not introducing many new modes into the system. Also, the vibrations in the cable tend to be higher frequency than the vibrations of the pendulum. In the double pendulum model, the highest frequency pendulum mode is the nineteenth lowest mode out of 252. The pendulum frequencies ranged from 0.48 Hz to 6.85 Hz whereas the

pendulum frequencies range from 2.06 Hz to 2513 Hz.

Having a function that classified modes for us was very useful. It spend up our analysis and trims the fat off the model. For a less experienced user, it cuts out superfluous information for them and gives them what they are interested in.

Packaging Update

Once the model was working, we wanted to make the update easy to install on other people's computers. Many people use the existing *Mathematica* model. We want people to be able to use the update without completely rewriting the program.

We put the majority of the update in a file called Cables.m. This file contains the majority of the changes needed to add the cables to the model. Now, users only need to add Cables.m to their directory of support files (which exists to let the model run anyway), add two lines of code to the model to reference Cables.m, and add the parameters defining the cables.

Packaging the update makes it significantly easier to install.

Increasing Robustness

Since the update was developed, we have been making it more capable than its first incarnation. In particular, we have changed how many cables the model could handle. At first, we only had a single cable case. Subsequently, we have made a symmetric two cable case, a general two cable case and, finally, an arbitrary number of cables (see Fig. 7). Additionally, we have applied the model to different systems. Specifically, we added cabling to a model of the OMC which mimic the cables in the actual detector (see Fig. 8). Now, users can add and remove cables very easily.

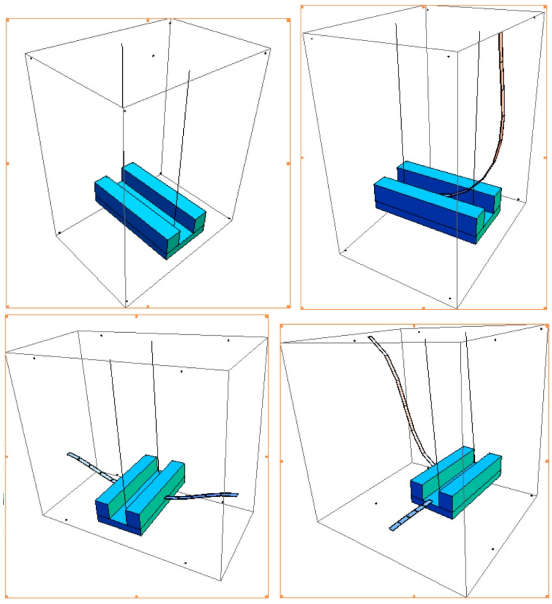

FIG. 7: The evolution of the cabling in the model.

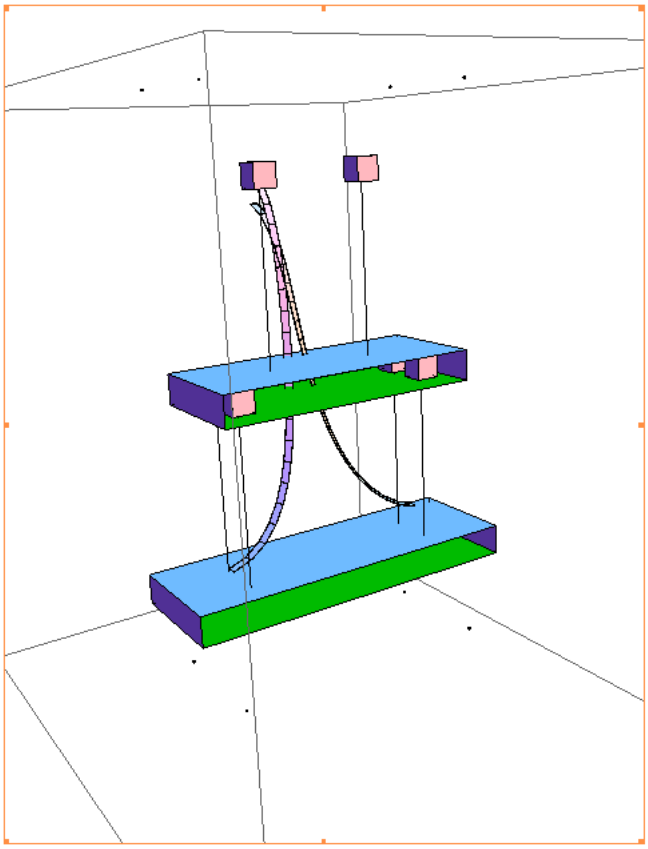

FIG. 8: An accurate model of the OMC. This was the last major step in the project.

CONCLUSIONS

Our project culminated in modeling the actual OMC. In our model, we included two cables going from the frame of the suspension to the optics bench on the OMC. The model assumes that the cables have the same mass per unit length, diameter, stiffness and damping. However, the cables are attached asymmetrically and have different lengths. The attachment points and lengths are taken from the schematics of the OMC suspension. Each cable was approximated 20 rigid segments. The equilibrium position given by the model is shown in Fig. 8.

Using my partner's experimental results, we were able to get measurements for the stiffness of the cable. In the

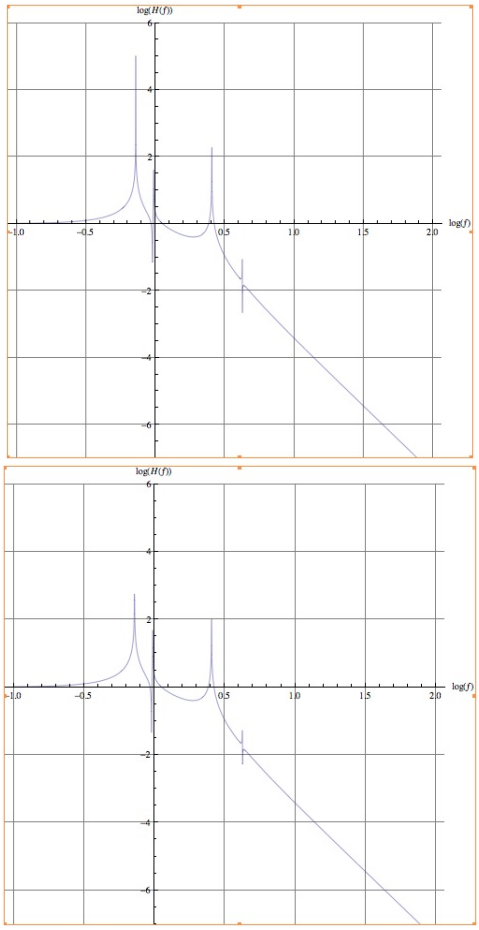

FIG. 9: The transfer function for the x motion of the optic before we added cabling (above) and after we added cabling (below). They are plotted over frequencies from 0.1 Hz to 100 Hz.

model, these measurements are expressed as spring constants for a calibrated length of cable in the six different degrees of freedom. Since cables compress and shear far less than they bend and twist, we made the three spring constants for the translational degrees of freedom arbitrarily high. For the rotational degrees of freedom, we calibrate from the torsional spring constants

$$k_{yaw} = .012 \text{ N m/rad},$$
$$k_{pitch} = .012 \text{ N m/rad, and}$$
$$k_{roll} = .003 \text{ N m/rad}.$$

We calculated k_{yaw}, k_{pitch}, and k_{roll} from my partner's data. For k_{yaw} and k_{pitch}, the spring constants that correspond to bending the wire, we measured deflection of horizontal lengths of cable as different weights were attached to the end. To calculate k_{roll}, we used the cable as a torsional pendulum and measured the frequency. Physically, these values give a magnitude for the stiffness and indicate that the cables are easier to twist than bend. For the damping, we used an improbably high value for the loss angle: $\phi(f) = 0.1$. My partner's experimental data showed that the actual loss angle was approximately an order of magnitude less, which is better for our system.

Using these numbers, we generated transfer functions for the x motion of the optic (Fig. 9). As we can see, the differences in these two transfer functions are nearly imperceptible. The heights of the peaks are different, but since the quality factors are very high in this system, these plots may not be trustworthy around the peaks. They are useful in other ranges, though. When cabling was added, the worry was that the transfer function would fall off at a different rate than it was designed to, which would happen if the cable introduced loss into the system. However, these transfer functions suggest that there is not a substantial change to the system at high frequencies. Therefore, our model indicates that the cable noise should not be a substantial problem in the Enhanced and Advanced LIGO upgrades.

ACKNOWLEDGMENTS

I would like to thank my mentors, Mark Barton and Norna Robertson, for bringing me into this project and guiding me in my research. I would also like to thank Calum Torrie and my partner, Chihyu Chen, for their help.

Finally, I would like to thank Caltech, LSC, SURF, and the NSF for providing the resources and facilities for our research.

---

* Mentors: Mark Barton and Norna Robertson
Co-SURF: Chihyu Chen

[1] Norna Robertson, *Conceptual Design of a Double Pendulum for the Output Modecleaner*. LIGO Science Collaboration, LIGO-T060257-03-R 2006.
[2] Dennis Coyne, *Vibration Transmission to the OMC Bench via Cabling*. LIGO Science Collaboration, 2006.
[3] Norna Robertson, *Vibration Transmission to the OMC Bench via Cabling*. LIGO Science Collaboration, LIGO-T070304-00-R 2007.
[4] Mark Barton, *Models of the Advanced LIGO Suspensions in Mathematica*. LIGO Science Collaboration, LIGO-T020205-02draft-D 2006.