LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY

LIGO Laboratory / LIGO Scientific Collaboration

LIGO- T070009-00 *ADVANCED LIGO* 16 January 2007

Quadruple Pendulum Controls Prototype Data Summary

Brett Shapiro, Richard Mittleman

Distribution of this document:
LIGO Science Collaboration

This is an internal working note
of the LIGO Project.

California Institute of Technology
LIGO Project – MS 18-34
1200 E. California Blvd.
Pasadena, CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project – NW17-161
175 Albany St
Cambridge, MA 02139
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

LIGO Hanford Observatory
P.O. Box 1970
Mail Stop S9-02
Richland WA 99352
Phone 509-372-8106
Fax 509-372-8137

LIGO Livingston Observatory
P.O. Box 940
Livingston, LA 70754
Phone 225-686-3100
Fax 225-686-7189

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

1 Introduction

This document summarizes the data collected from the quadruple pendulum controls prototype at LASTI as of December 2006. It also serves as an update to the document T060134-00 written last June. It includes the transfer functions (TFs) of the undamped pendulum and their resonance frequencies as well as damping control data such as the control filters and impulse responses. Other results included in this document are the thermal heat load test, eddy current damping (ECD) test, and problems encountered during the testing procedures.

2 Undamped Local Transfer Functions and Resonance Frequencies

2.1 Local Transfer Functions

These six figures are plots of transfer functions at the top mass of each pendulum. Each transfer function is measured from a swept sine excitation at one of the degrees of freedom (DOFs) and measured from the same DOF, i.e. longitudinal to longitudinal. The red curve represents the prediction from the model, the blue the measured response from the reaction chain, and the black the main chain. The red curves include some velocity damping in the state space model to add some realism and increase their proximity to the measured data. Specifically I added -0.1 velocity damping to each mode in the state space. There is no other damping in these plots. The model curves are shifted above the others simply for clarity. Some of the reaction chain data is not plotted on these figures because in the original data sets the OSEMs are saturating and the resolution was too low. The undamped data was not retaken because the eddy current dampers were then installed in the reaction chain. It was decided to leave these in for the time being instead of retaking the undamped sysid because the location of the modes can still be determined from the eddy current damped TFs. In addition, the installation of the ECD proved extremely difficult, which is discussed further in Section 7 Mechanical/Electrical Difficulties. The original data sets can still be referenced in the previous document T060134.

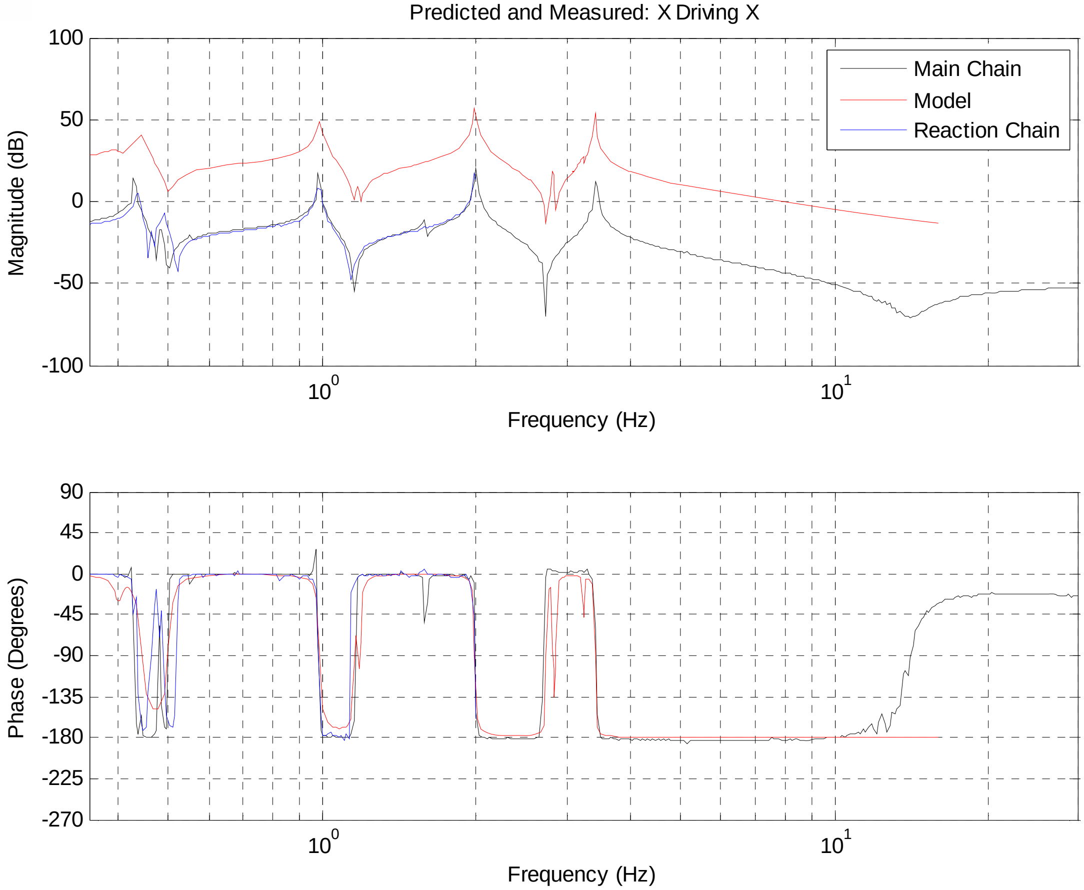

Predicted and Measured: X Driving X
Main Chain
Model
Reaction Chain
Magnitude (dB)
100
50
0
-50
-100
10^0
10^1
Frequency (Hz)
Phase (Degrees)
90
45
0
-45
-90
-135
-180
-225
-270
10^0
10^1
Frequency (Hz)

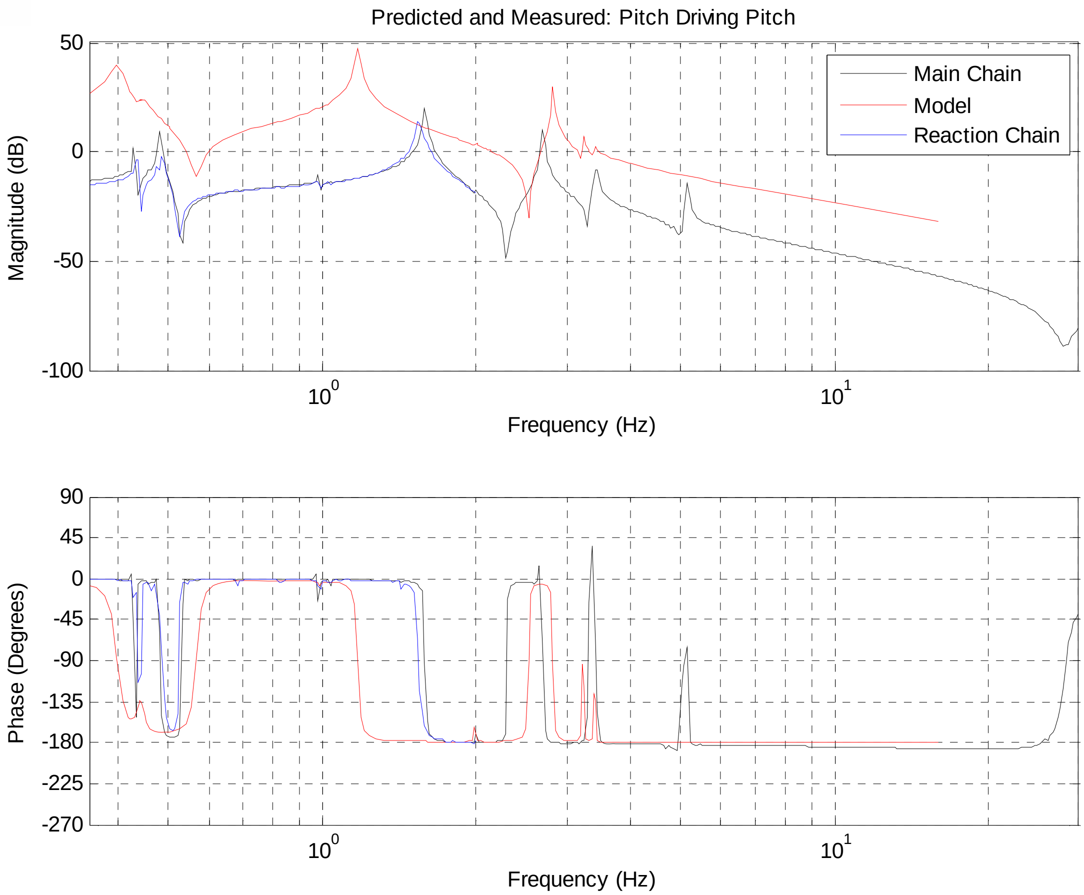
Predicted and Measured: Pitch Driving Pitch
Main Chain
Model
Reaction Chain
Magnitude (dB)
50
0
-50
-100
10^0
10^1
Frequency (Hz)
Phase (Degrees)
90
45
0
-45
-90
-135
-180
-225
-270
10^0
10^1
Frequency (Hz)

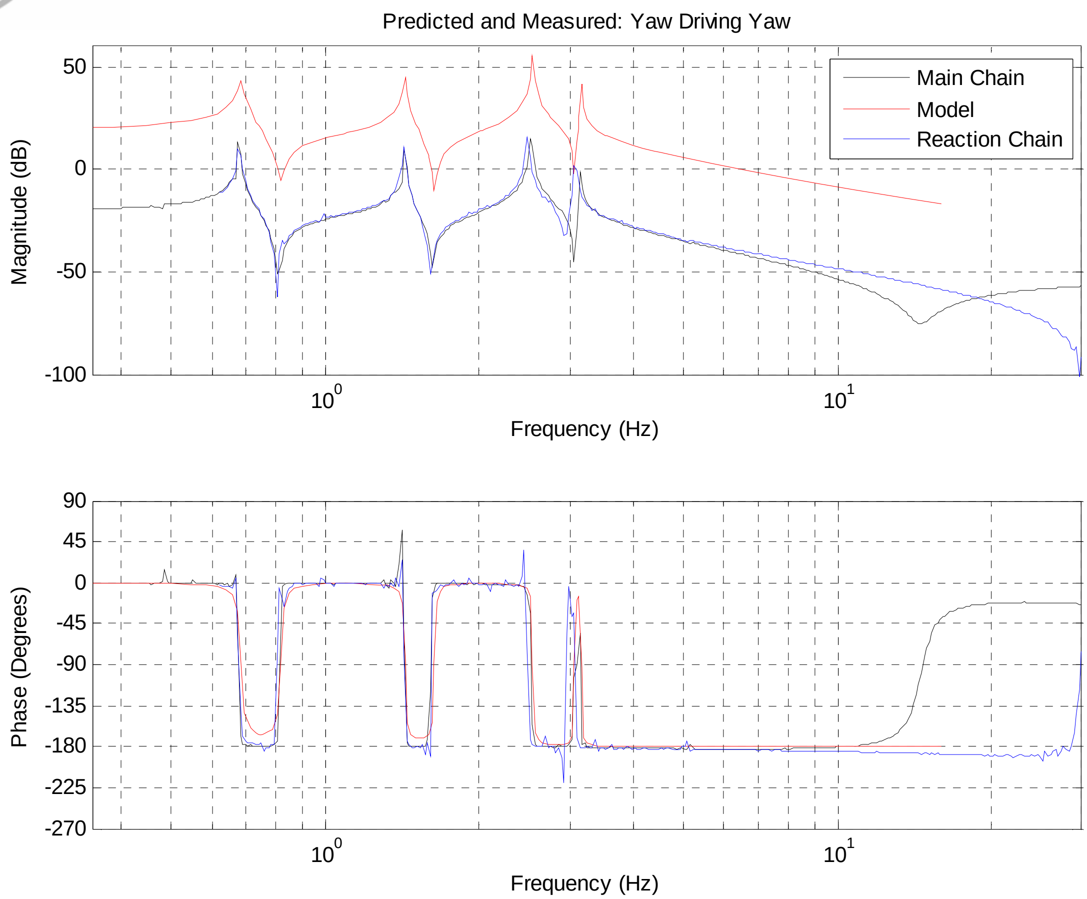
Predicted and Measured: Yaw Driving Yaw
Magnitude (dB)
50
0
-50
-100
10^0
10^1
Frequency (Hz)
Main Chain
Model
Reaction Chain
Phase (Degrees)
90
45
0
-45
-90
-135
-180
-225
-270
10^0
10^1
Frequency (Hz)

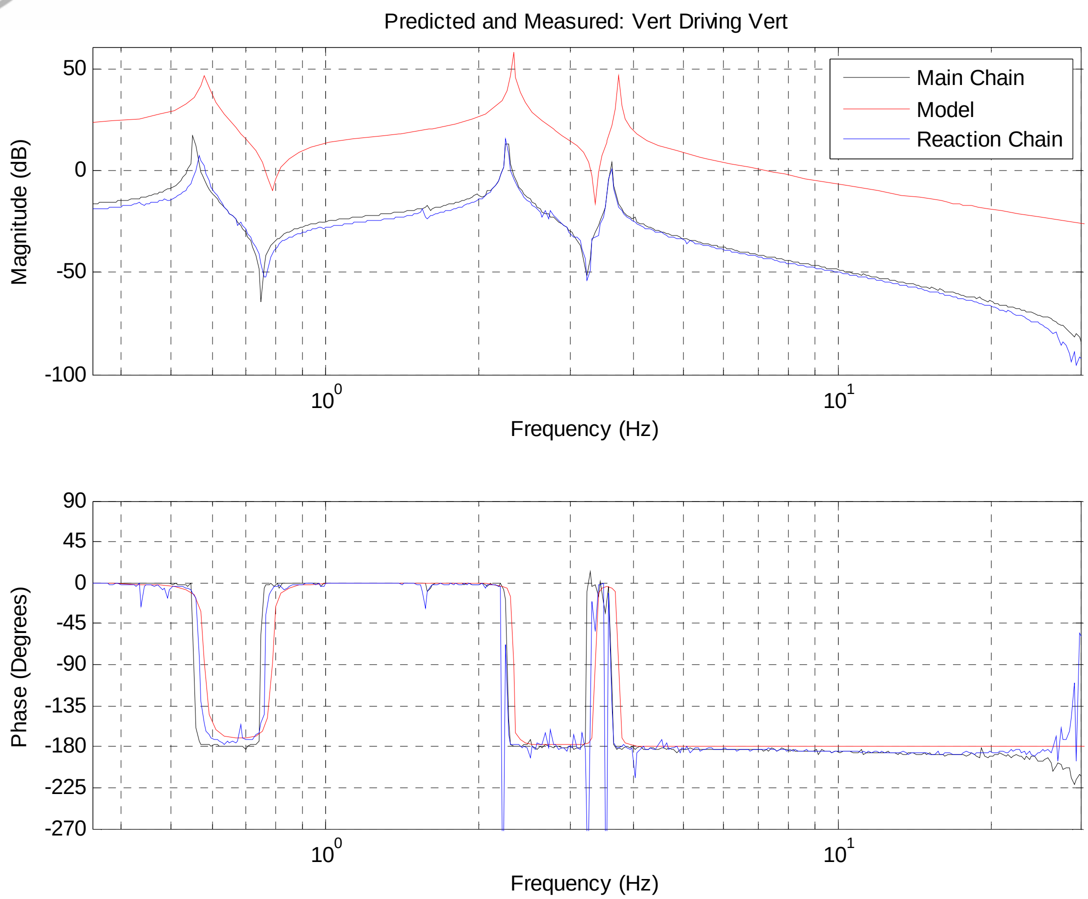
Predicted and Measured: Vert Driving Vert
Main Chain
Model
Reaction Chain
Magnitude (dB)
50
0
-50
-100
10^0
10^1
Frequency (Hz)
Phase (Degrees)
90
45
0
-45
-90
-135
-180
-225
-270
10^0
10^1
Frequency (Hz)

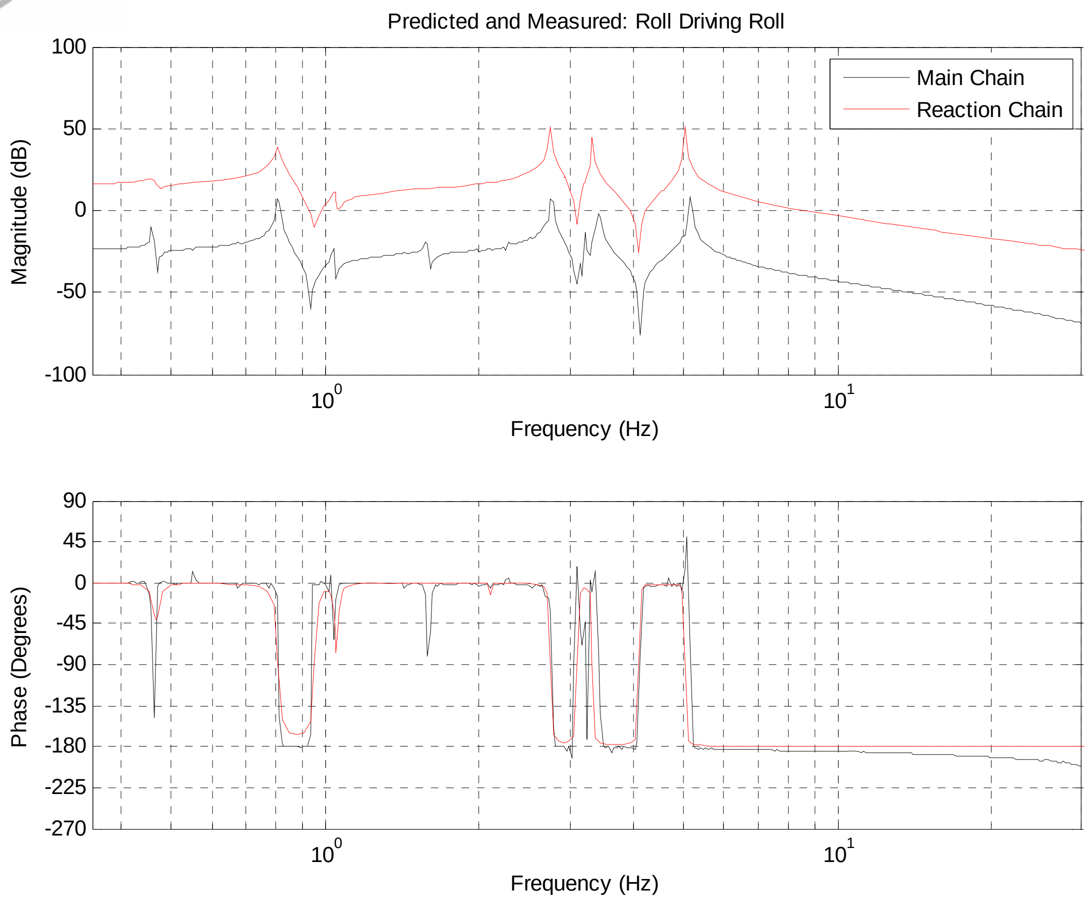
Predicted and Measured: Roll Driving Roll
Main Chain
Reaction Chain
Magnitude (dB)
100
50
0
-50
-100
10^0
10^1
Frequency (Hz)
Phase (Degrees)
90
45
0
-45
-90
-135
-180
-225
-270
10^0
10^1
Frequency (Hz)

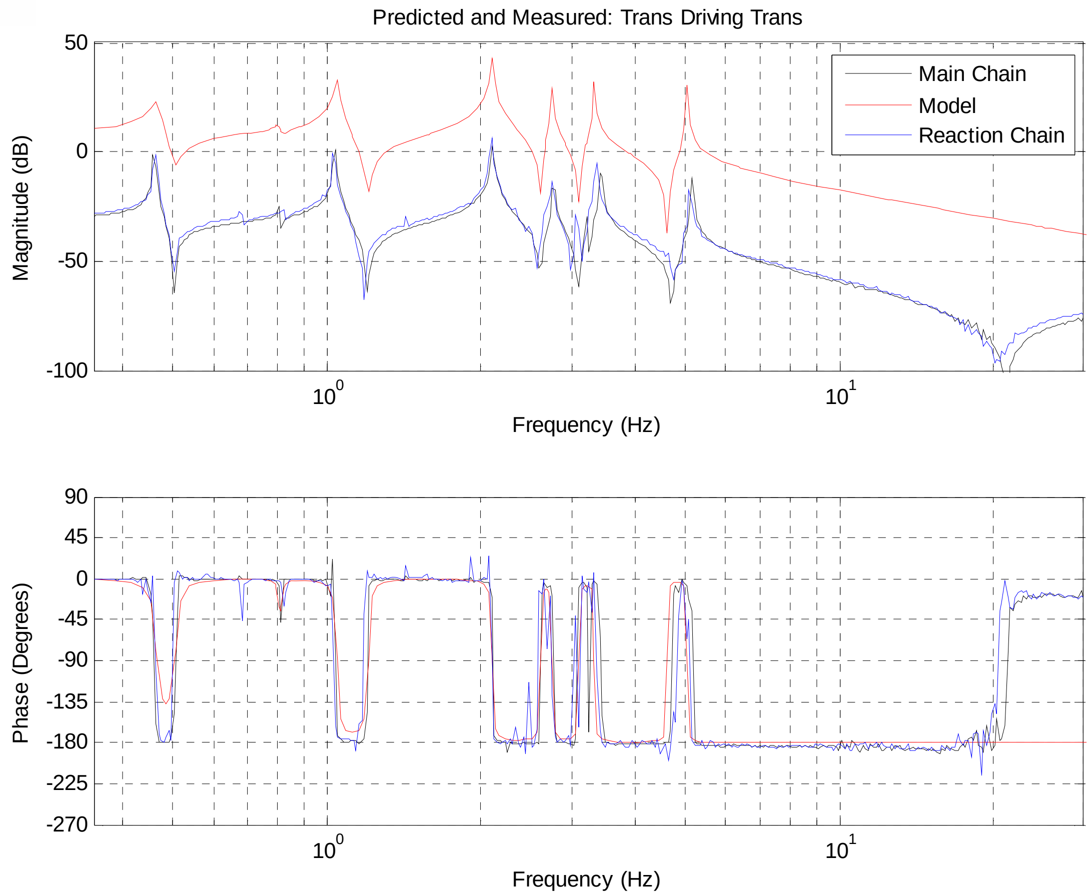

Both the main chain and reaction chain match each other very well in most instances and generally to the predictions as well. The greatest exception is pitch. The first two pitch modes are greater than the predictions. The first mode is greater by roughly 20% and the second by 30-35%. Most of this error can be reduced to a maximum of 10% in the model by increasing the wire break off points in the top mass (the dn values) by about 4 mm. However, attempts to directly measure this value on the quad itself have failed due to the limited access inside the masses. There is also a 5.1 Hz mode, which is coupling to roll due to asymmetrical trim mass. Likewise the 1.58 Hz pitch mode is visible in the roll measurement.

Some TFs appear to start rolling up at the high frequency end. This is an electronic issue which is discussed in further in the Mechanical/Electrical Difficulties section.

2.2 Local Resonance Frequencies

Longitudinal-Pitch						Transverse-Roll				
Pred.	Main	%err+-1	React.	%err+-1		Pred.	Main	%err+-1	React.	%err+-1
0.397	0.480	20.91	0.4856	22.32		0.464	0.458	-1.29	0.4634	-0.13
0.443	0.427	-3.61	0.4371	-1.33		0.808	0.803	-0.62	0.8221	1.75
0.989	0.980	-0.91	0.9798	-0.93		1.044	1.039	-0.48	1.027	-1.63
1.171	1.583	35.18	1.528	30.49		2.096	2.097	0.05	2.097	0.05
1.983	2.000	0.86	1.977	-0.30		2.752	2.752	0	2.752	0
2.819	2.688	-4.65	2.656	-5.78		3.325	3.210	-3.46	3.365	1.20
3.232	3.437	5.96	3.365	4.12		5.046	5.151	2.08	5.09	0.87
3.405	3.407	0.06	3.405	0		25.123	24.346?/25.340		-3.09?/0.86	

Yaw						Vertical				
Pred.	Main	%err+-1	React.	%err+-1		Pred.	Main	%err+-1	React.	%err+-1
0.685	0.674	-1.61	0.674	-1.61		0.581	0.552	-4.99	0.5654	-2.69
1.429	1.424	-0.35	1.425	-0.28		2.325	2.278	-2.02	2.251	-3.18
2.536	2.504	-1.26	2.474	-2.44		3.733	3.613	-3.21	3.613	-3.21
3.165	3.135	-0.95	3.061	-3.29		17.332	17.201/17.418		-0.76/0.50	

This table lists all the predicted modes next to all the main and reaction chain modes. The reaction chain modes not shown in the undamped transfer functions above were located using the eddy current damped transfer functions below. The highest roll and vertical modes are too weak to be seen from the top mass so they were measured by driving the penultimate mass OSEMs. Since these OSEMs are between the two chains, both are excited simultaneously and it is not possible to distinguish a main chain mode from a reaction chain mode. For this reason there is no distinction of which measurement goes to which chain in the table. The question mark by the 24.346 Hz vertical measurement is because it is not clear whether this is in fact a mode or noise.

3 Eddy Current Damping Results

The eddy current dampers were designed to damp the vertical and roll modes. Thus, measured TFs of these degrees of freedom are presented here along with predictions from the model. Yaw is also presented for reasons discussed below.

All these measurements were taken on the reaction chain since installing dampers on the main chain would have required realigning the whole pendulum.

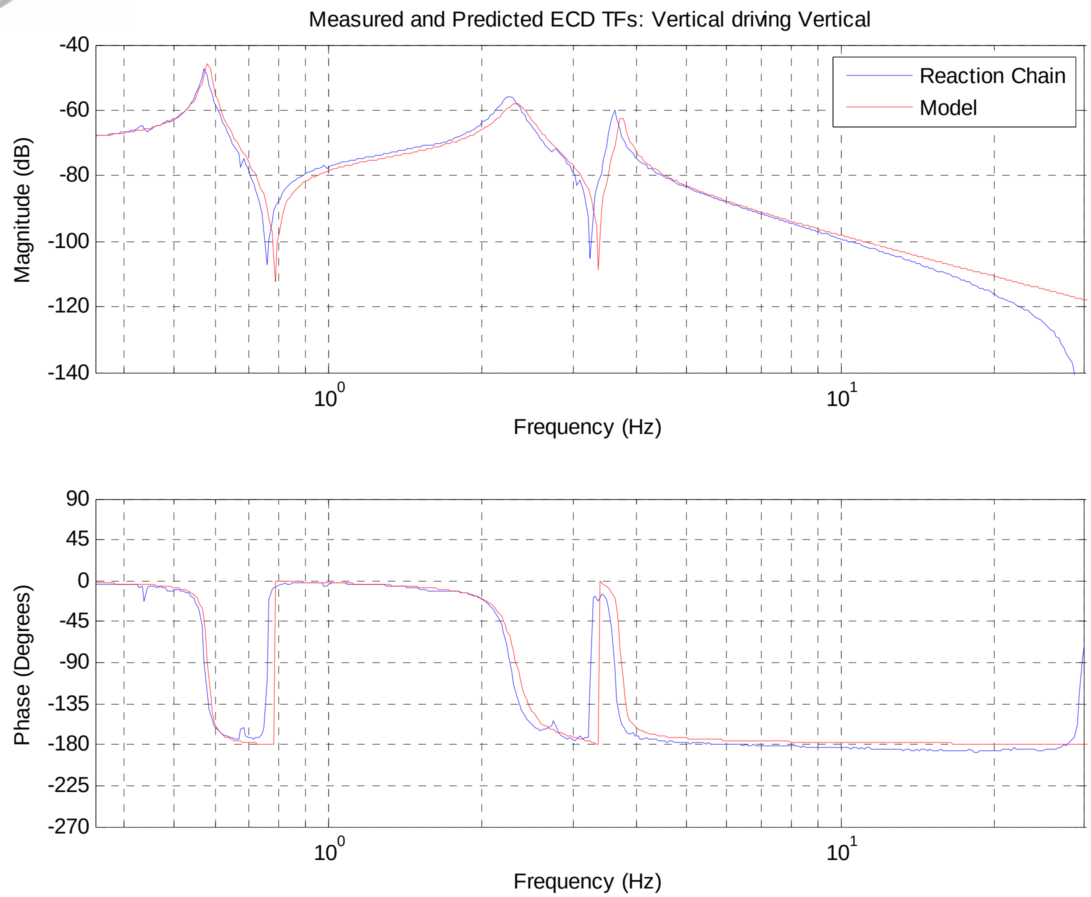
Measured and Predicted ECD TFs: Vertical driving Vertical
Magnitude (dB)
-40
-60
-80
-100
-120
-140
Reaction Chain
Model
10^0
10^1
Frequency (Hz)
Phase (Degrees)
90
45
0
-45
-90
-135
-180
-225
-270
10^0
10^1
Frequency (Hz)

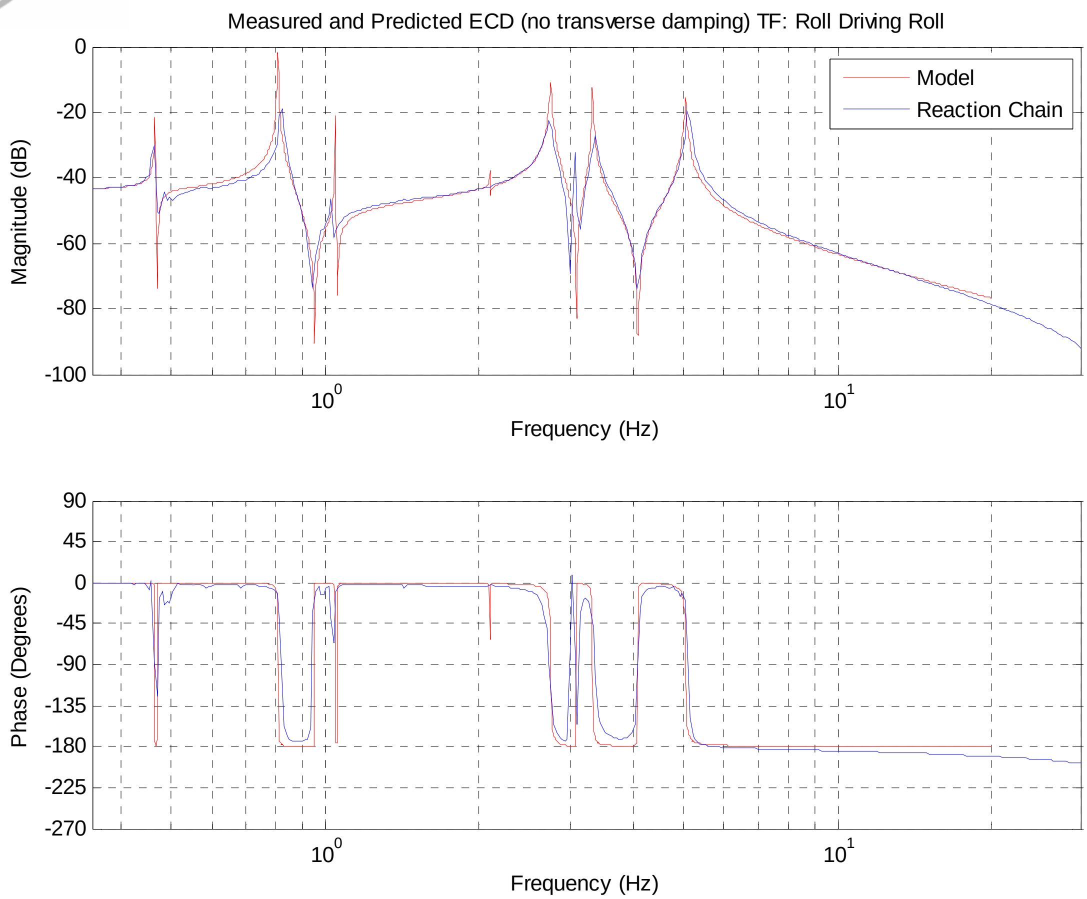
Measured and Predicted ECD (no transverse damping) TF: Roll Driving Roll
Model
Reaction Chain
Magnitude (dB)
0
-20
-40
-60
-80
-100
10^0
10^1
Frequency (Hz)
Phase (Degrees)
90
45
0
-45
-90
-135
-180
-225
-270
10^0
10^1
Frequency (Hz)

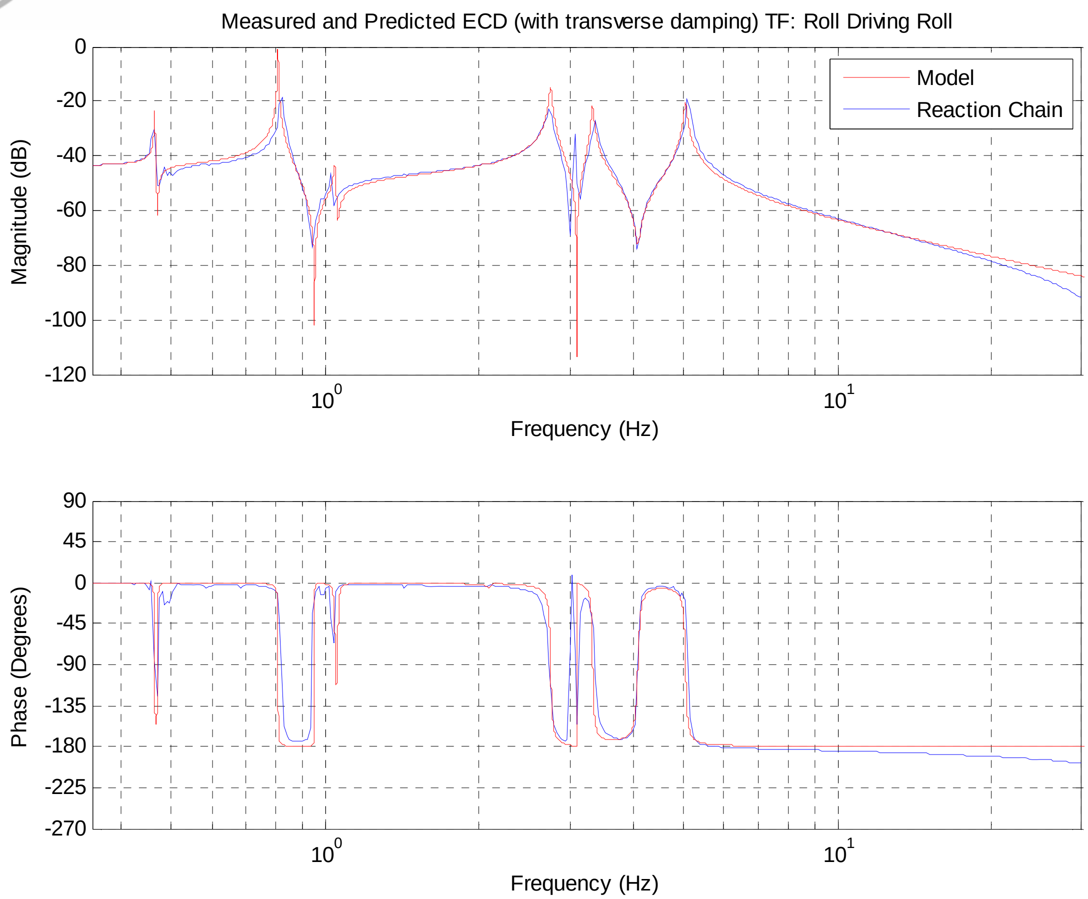
Measured and Predicted ECD (with transverse damping) TF: Roll Driving Roll
Model
Reaction Chain
Magnitude (dB)
0
-20
-40
-60
-80
-100
-120
10^0
10^1
Frequency (Hz)
Phase (Degrees)
90
45
0
-45
-90
-135
-180
-225
-270
10^0
10^1
Frequency (Hz)

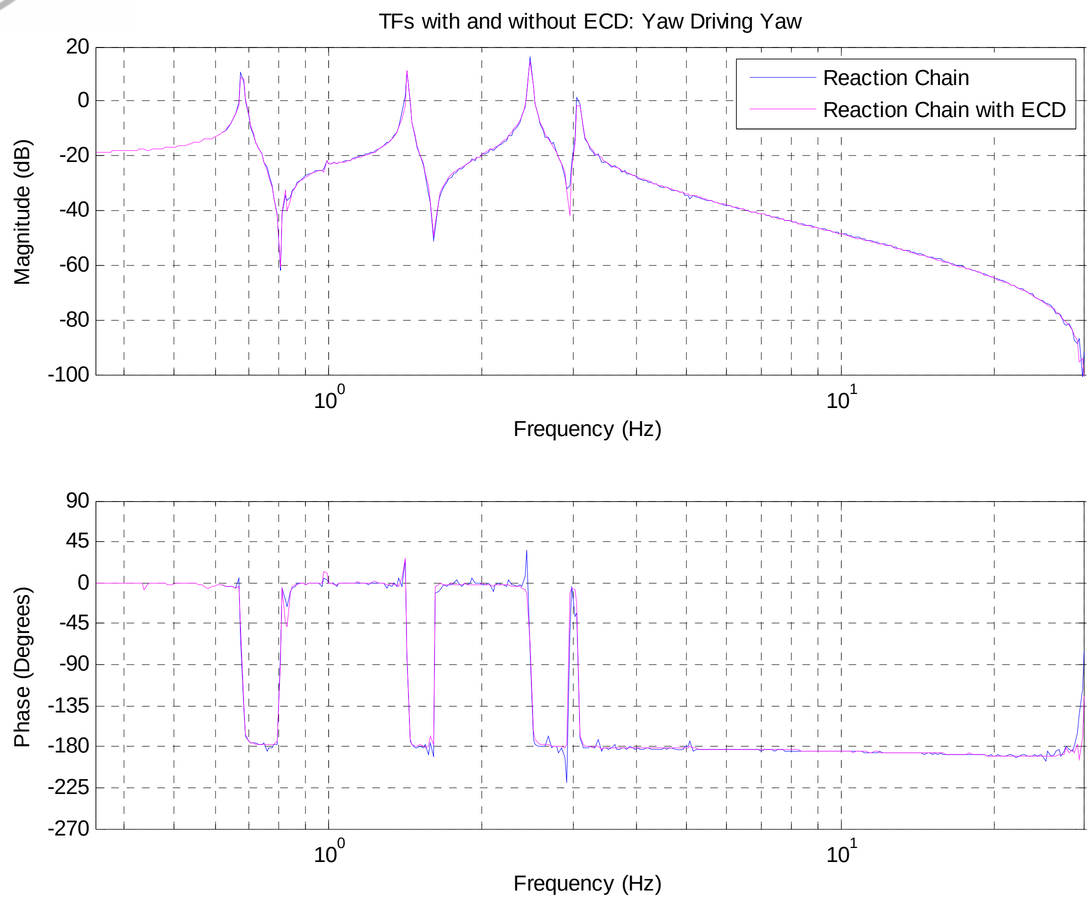

The first figure in this section show the damped vertical to vertical measured TF along with the prediction from the model. Here the two are never more than 2.5 dB off from each other, so the agreement is rather good.

Roll does not agree quite as well since it shows more damping than predicted. To test an idea that there is extra coupling to the transverse degree of freedom, some damping was explicitly added into the model for motion in that direction. This brings the results closer to observation suggesting that extra coupling may exist. To rule out the possibility that extra damping exists from the horizontal motion of the ECD magnets in the copper (damping should only exist for vertical motion) yaw was also looked at. Since yaw is the only DOF with no theoretical coupling to vertical motion of the ECD magnets than there should be virtually no additional damping unless horizontal motion does in fact produce damping as well. Cross coupling from yaw to other DOFs has already proven to be minimal. The last figure in this section shows yaw before and after the ECD were installed. Since the two measurements in this figure are virtually identical then we can conclude all the damping is due to vertical motion and that extra cross coupling from transverse to roll is likely the culprit for this extra damping. Examining the measurements of the transverse DOF also shows extra damping.

Even though the ECD is designed to damp vertical and roll motion, they also couple directly to pitch due to the finite size of the dampers in this direction. The rotation of the top mass in pitch causes the ECD magnets to move vertically in and out of the copper array resulting in some damping. There is also some indirect damping of the x DOF due to the coupling of x to pitch.

4 Cross Coupling

Two modes cross couple between the DOFs stronger than the others. One is the 5.1 Hz roll which couples very strongly to pitch on both chains. The other is the 1.5 Hz pitch coupling to roll, though much more so on the main chain. Extra coupling between roll and pitch is expected for unbalanced trim mass and is likely, at least partially, to blame in both cases.

The next most significant coupling appears in the vertical motion of the reaction chain. Here 3 pitch modes, an x, and a roll are visible even if just barely. All these modes are very apparent in pitch and are likely coupling through that DOF. A possible explanation is the electrostatic drive cables which drape over the top of the UI mass before connecting to the structure. These cables have caused difficulties in the past since it is not always obvious if they are pulling on the mass. Even a small amount of tension might weigh one side of the mass down as it tries to move vertically which would without a doubt excite pitch.

Other observed coupling is rather weak. On the main chain the1.5 Hz pitch is visible in the vertical TF as well as the 5.1 Hz roll in x. On the reaction chain the 0.44 and 0.99 Hz x are just barely visible in yaw. Both yaw OSEMs are also both of the x OSEMs which could explain this observation if there is a slight imbalance in force, position, or mass.

Observing cross-coupled TFs directly, i.e. driving yaw and observing vertical, was also attempted. This proved difficult however since the signals were so weak. Electronic issues also impeded these observations.

In all cases the coupled modes do not appear to change the shape of the TFs other than showing up as a peak in other DOFs.

5 Local Damping

All the damping data reported here are from the reaction chain since it has experienced greater luck with its sensors and has more complete data. The main chain, with one exception, uses the same damping filters and behaves very similarly, due to the symmetry between the two.

5.1 Local Damping Filters

There are currently two types of damping filters on the quadruple pendulum. The first is a simple velocity damper, titled as ‘Standard Damping’ in the figure, which damps all the degrees of freedom except transverse. Due to an intermittent instability of the coupled 5.1 Hz roll mode a notch was added to add a more robust phase margin. It was then necessary to add an elliptical filter to roll the high frequency gain off at a sufficient rate. If the gain is too large at high frequencies instability develops around 90 Hz. The interesting thing is that this instability is coupled between the chains and only exists while both are simultaneously damped.

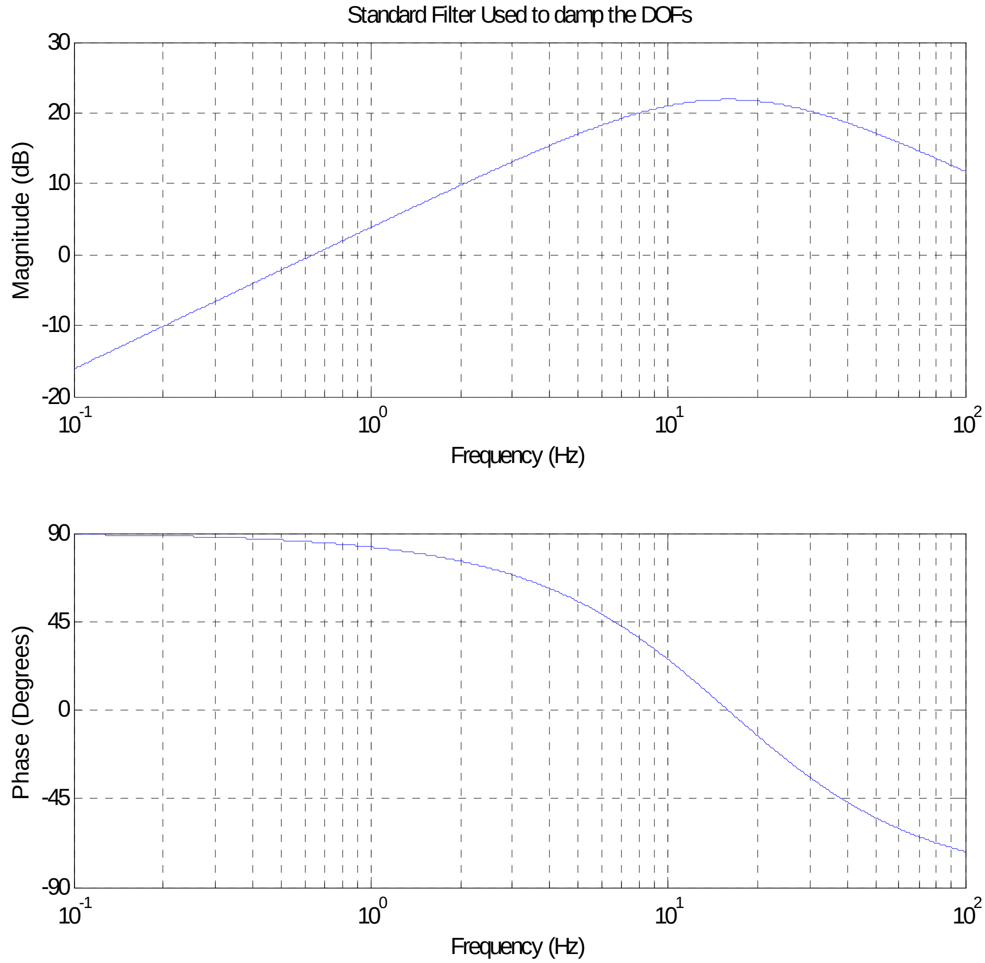

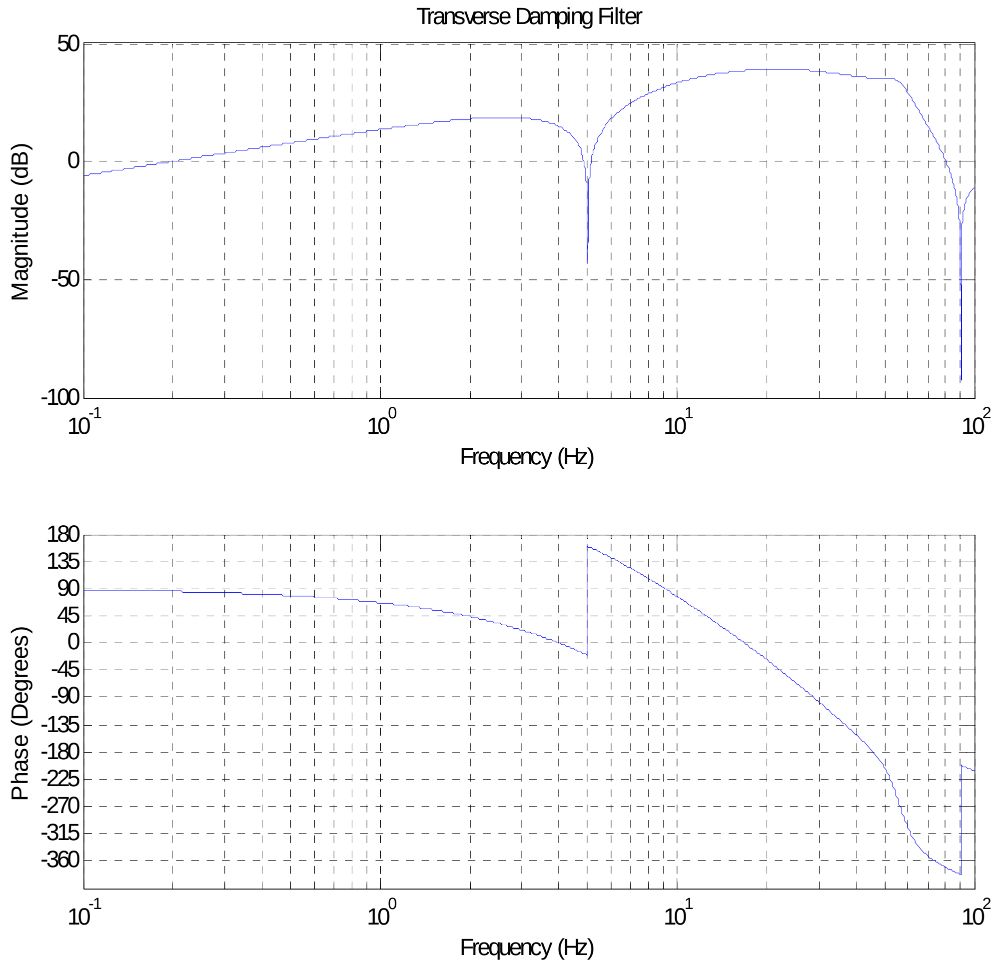

The standard filter uses a zero at 0 Hz and two poles at 16 Hz. Some of the DOFs have an additional gain between 0.5 and 5 in order to improve damping or stability. The transverse filter is similar but with an elliptical filter that begins rolling off the gain at 56 Hz. The latter filter is not used on the main chain.

5.2 Local Damping Transfer Functions

The results in this section are unchanged from those in the previous document T060134. The filters here prove the feasibility of damping but are not optimized. As mentioned in the Future Work section, the optimization of the damping will be the focus of my master's thesis.

These plots show the reaction chain undamped local transfer functions (same locations as section 2) together with the damped transfer functions. The damped transfer functions are measured from the excitation to the response and thus have the form of PC/(1+PC) where P is the plant and C the controller.

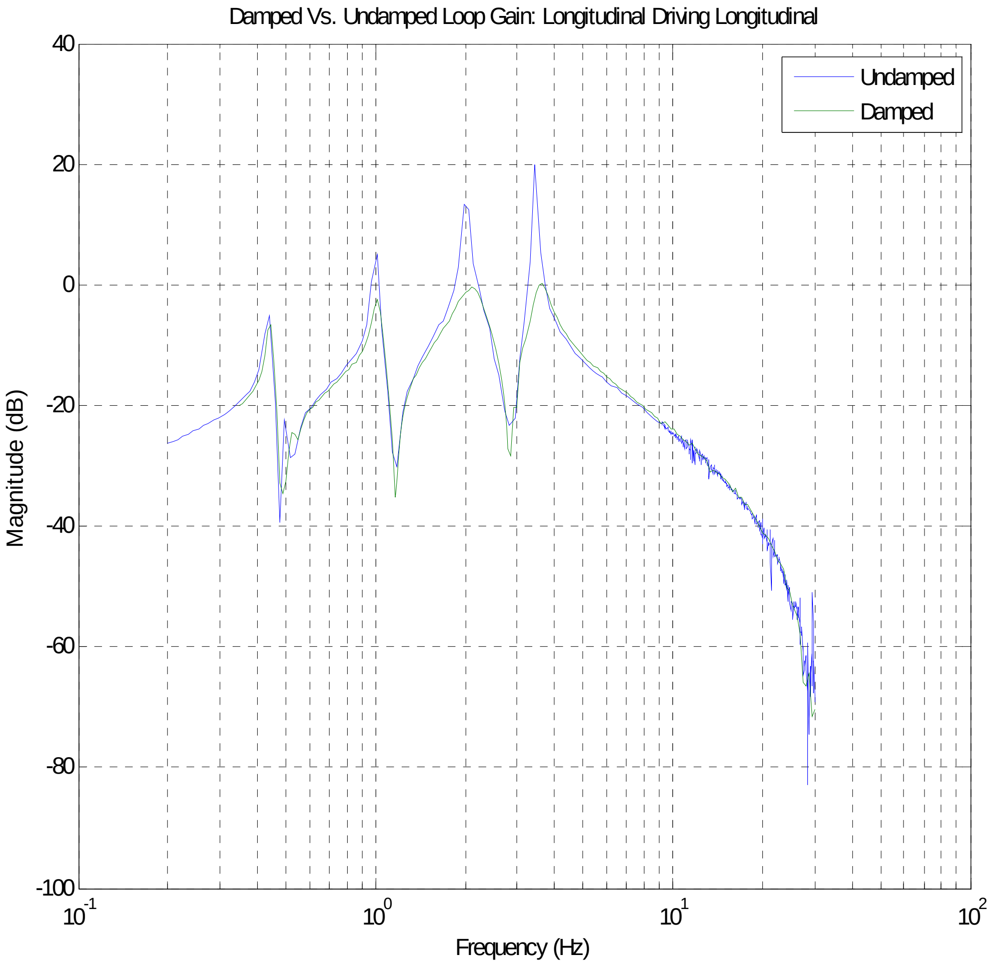
Damped Vs. Undamped Loop Gain: Longitudinal Driving Longitudinal
Undamped
Damped
Magnitude (dB)
Frequency (Hz)
40
20
0
-20
-40
-60
-80
-100
10^{-1}
10^{0}
10^{1}
10^{2}

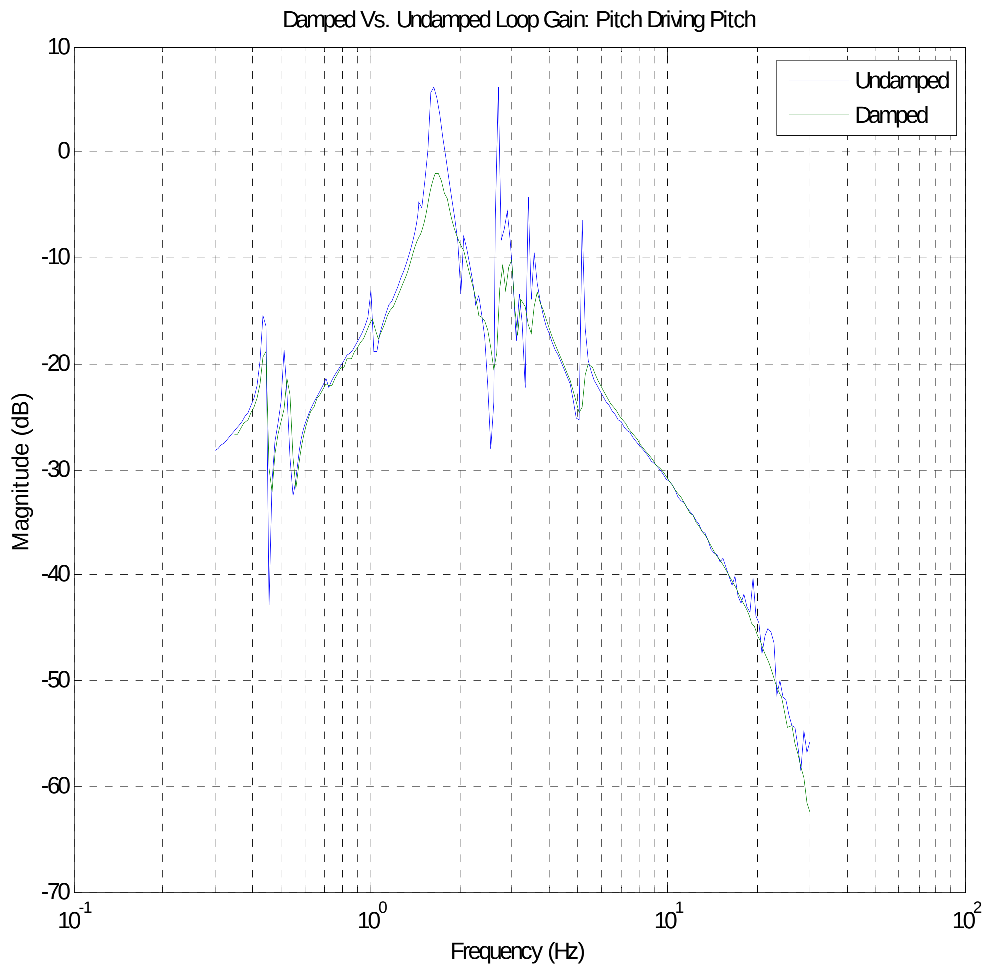

Damped Vs. Undamped Loop Gain: Pitch Driving Pitch
Undamped
Damped
Magnitude (dB)
Frequency (Hz)
10
0
-10
-20
-30
-40
-50
-60
-70
10^{-1}
10^{0}
10^{1}
10^{2}

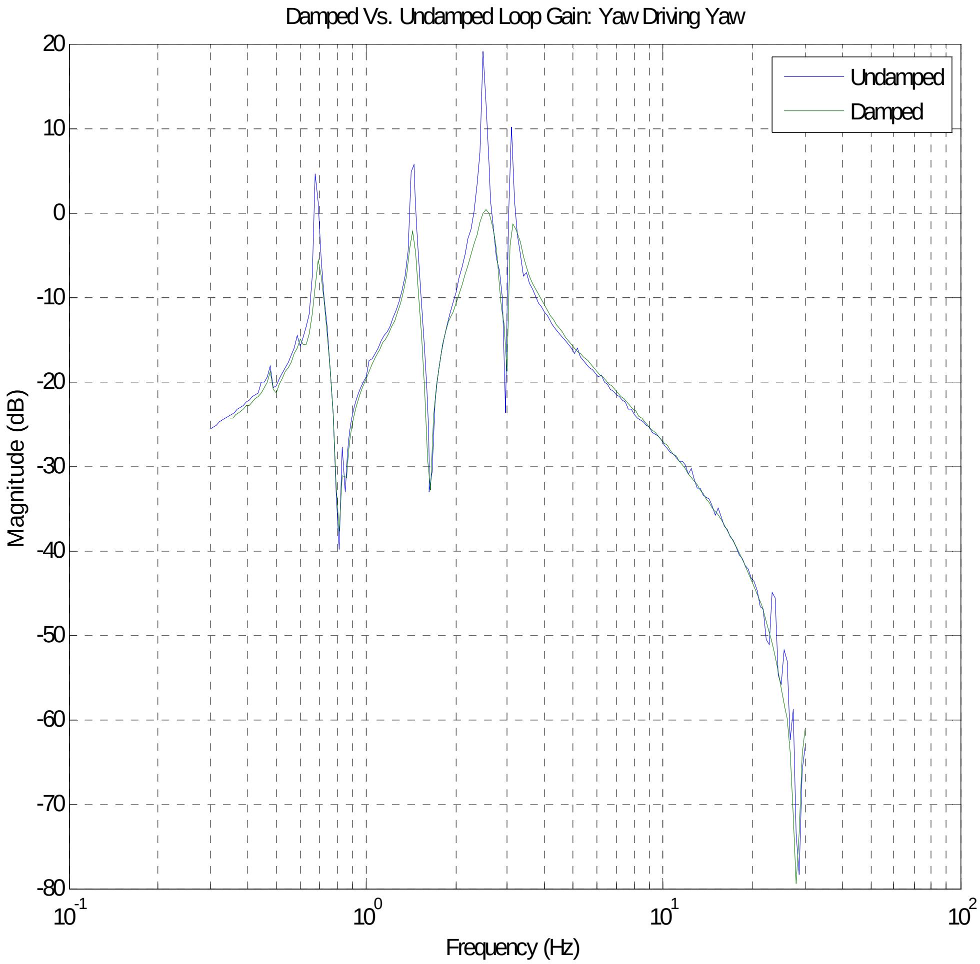
Damped Vs. Undamped Loop Gain: Yaw Driving Yaw
Undamped
Damped
Magnitude (dB)
Frequency (Hz)
20
10
0
-10
-20
-30
-40
-50
-60
-70
-80
10^{-1}
10^{0}
10^{1}
10^{2}

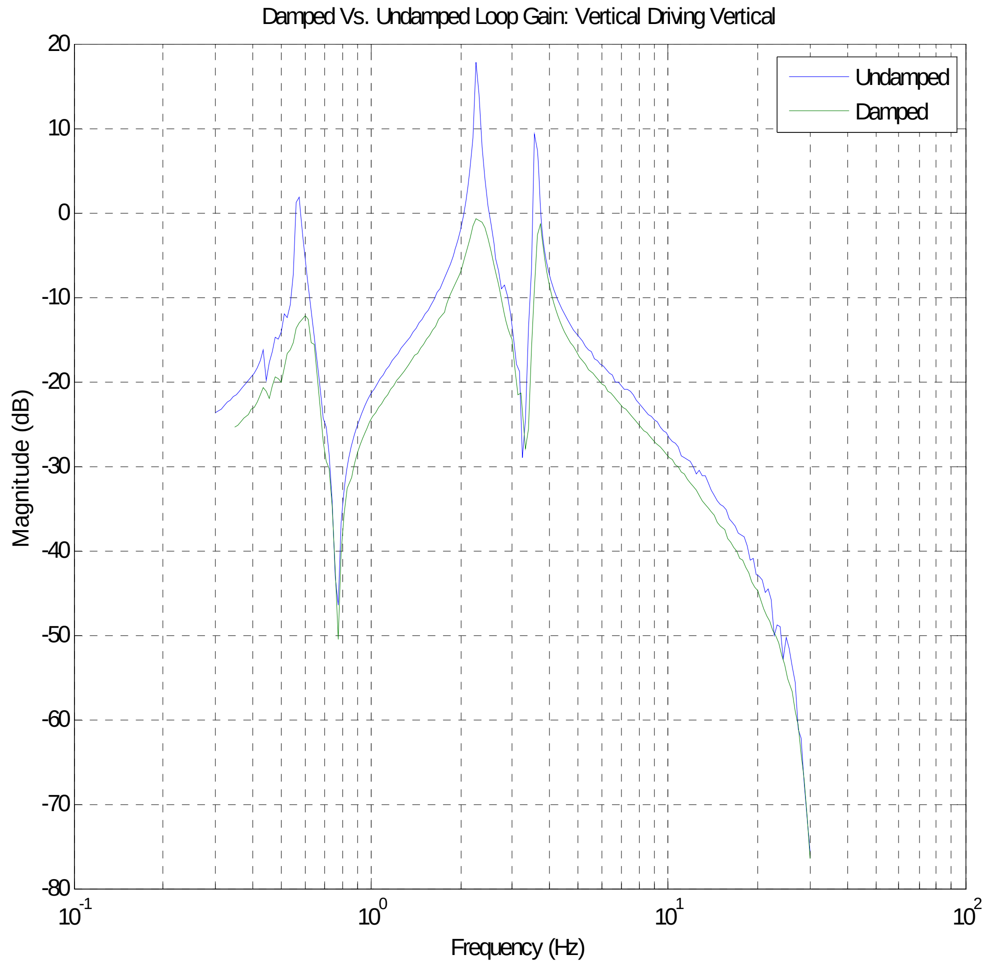
Damped Vs. Undamped Loop Gain: Vertical Driving Vertical
Undamped
Damped
Magnitude (dB)
Frequency (Hz)
20
10
0
-10
-20
-30
-40
-50
-60
-70
-80
10^{-1}
10^{0}
10^{1}
10^{2}

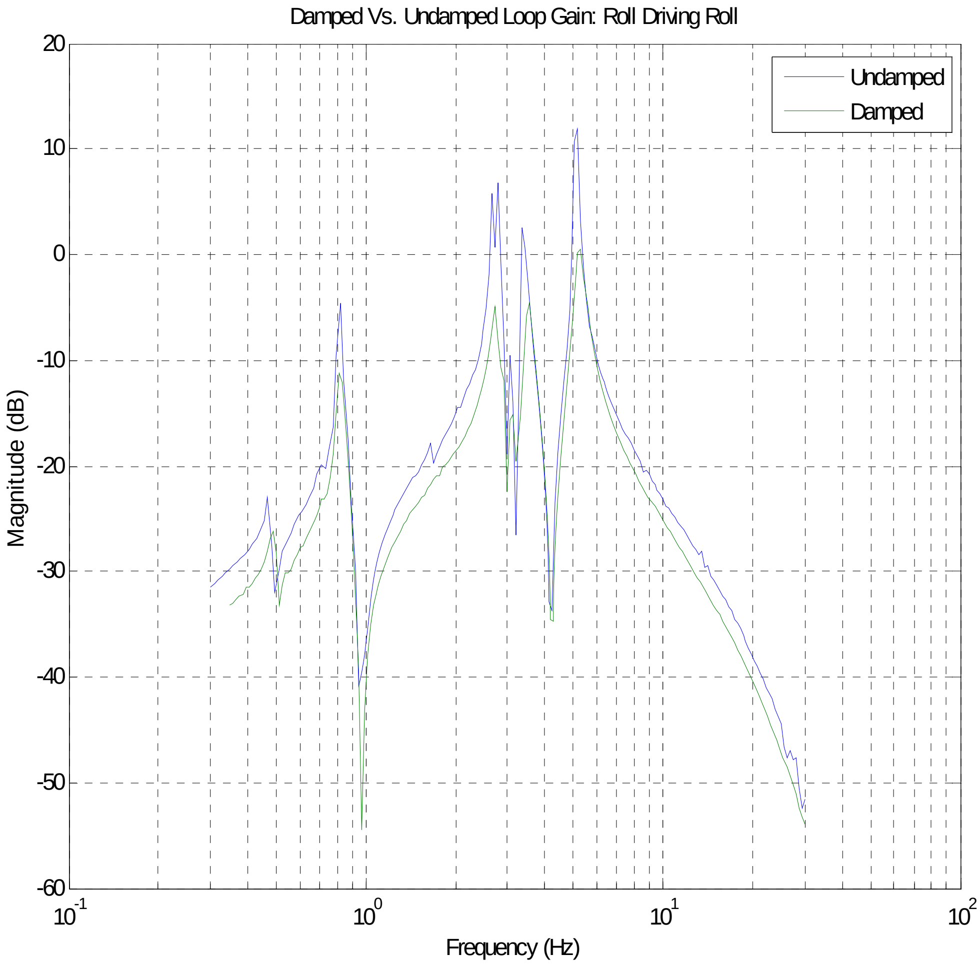

Damped Vs. Undamped Loop Gain: Roll Driving Roll
Undamped
Damped
Magnitude (dB)
Frequency (Hz)
20
10
0
-10
-20
-30
-40
-50
-60
10^{-1}
10^{0}
10^{1}
10^{2}

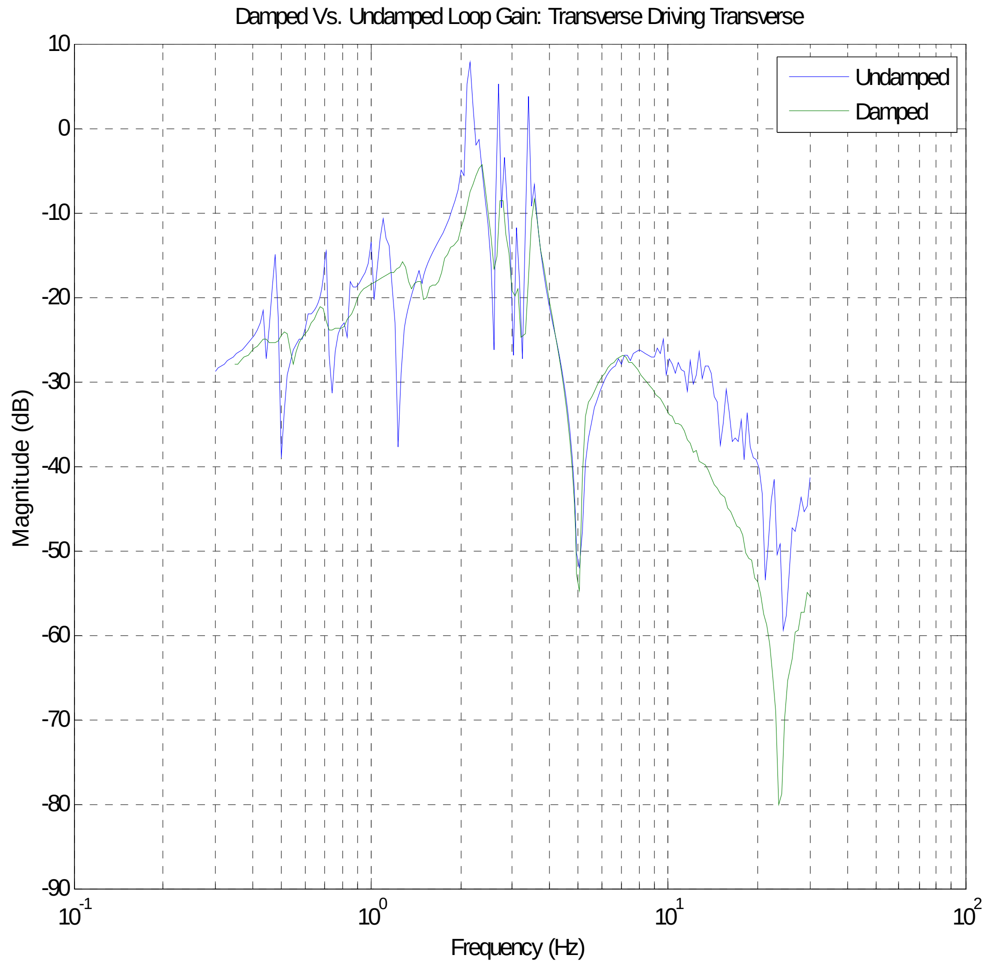

The lowest longitudinal mode does not show much damping here but the controller gain has been adjusted since this data was taken. There was not enough time to include the new data but damping is now much more significant. The roll damping has been adjusted as well and also has more damping then shown. Simple optimization techniques such as adding a bump in the gain at the lower modes also significantly improve damping of these modes.

5.3 Local Damping Impulse Response

These impulse responses are measured at the same local DOFs as above. The requirement on the impulse response is to damp to 1/e (≈37%) within ten seconds. Here the red lines represent damping to 10%. As shown, all the DOFs damp within 10% of the initial maximum in less than ten seconds and more than meet the requirements. This data was collected with some of the improved damping filters mentioned earlier.

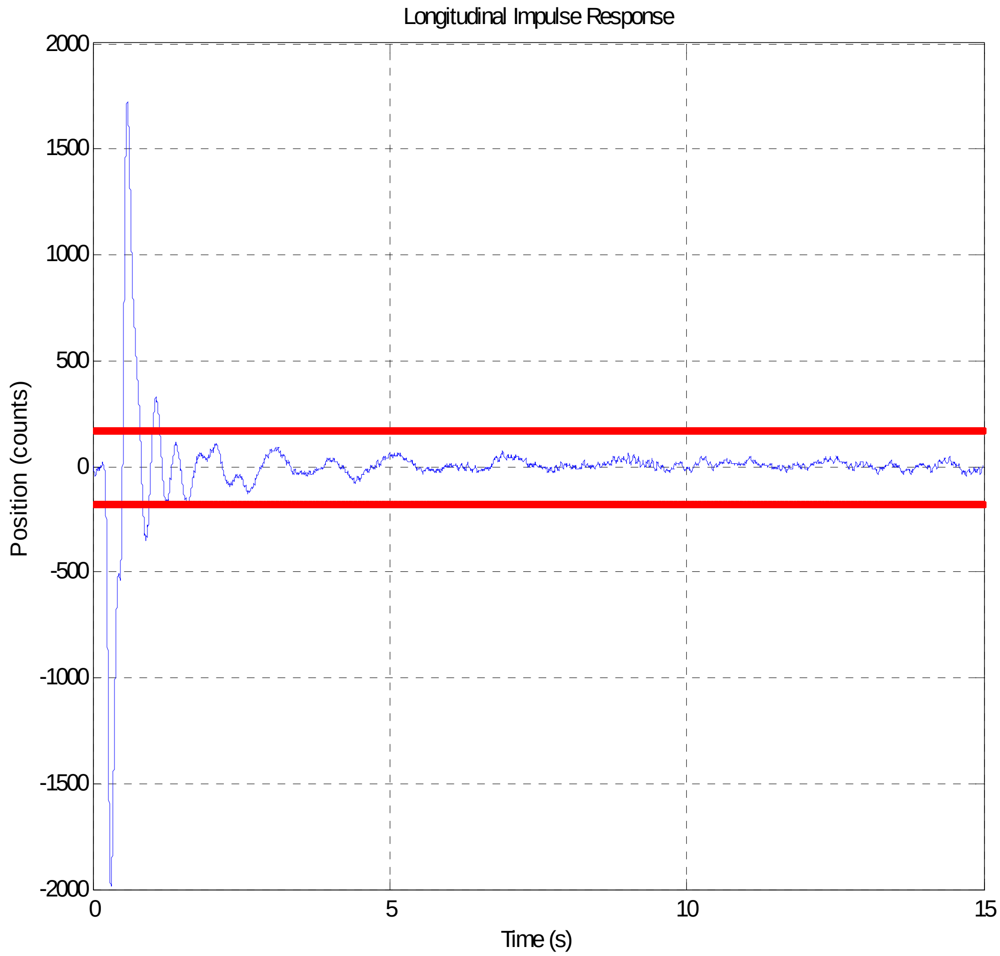
Longitudinal Impulse Response
2000
1500
1000
500
0
-500
-1000
-1500
-2000
Position (counts)
0
5
10
15
Time (s)

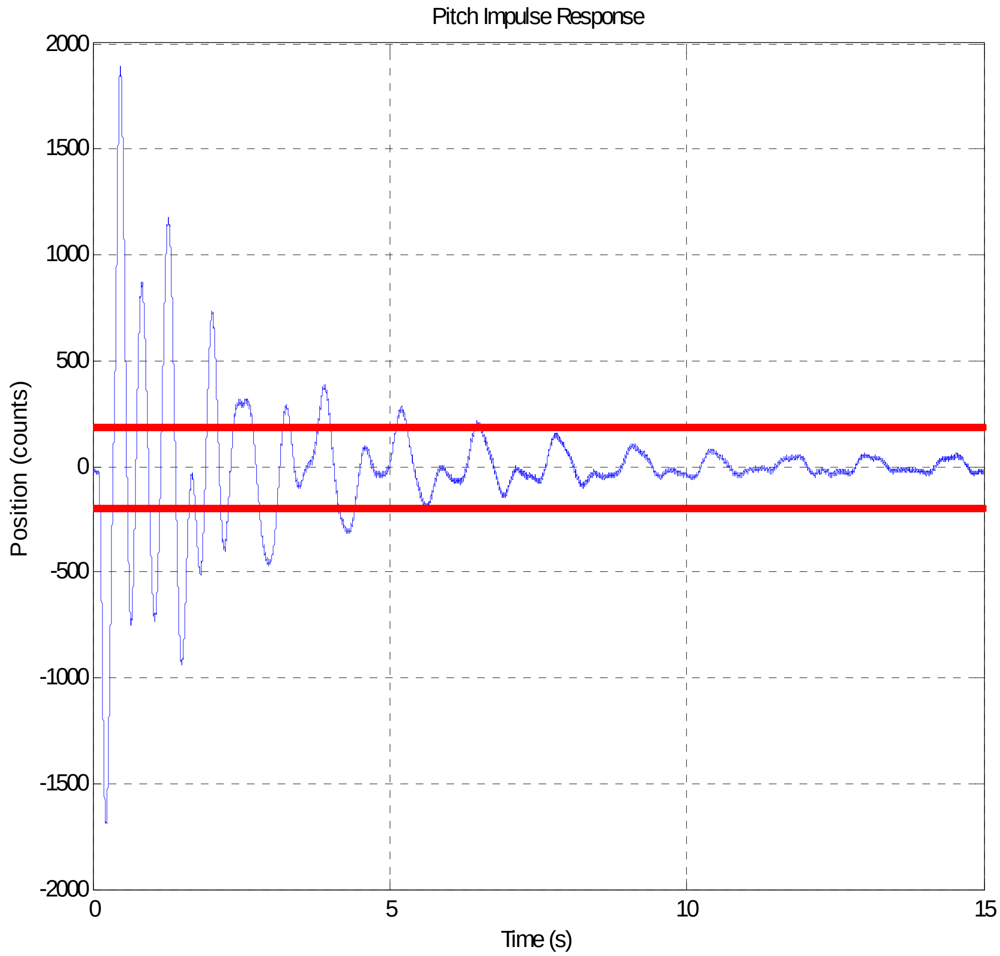

Pitch Impulse Response
Position (counts)
Time (s)
2000
1500
1000
500
0
-500
-1000
-1500
-2000
0
5
10
15

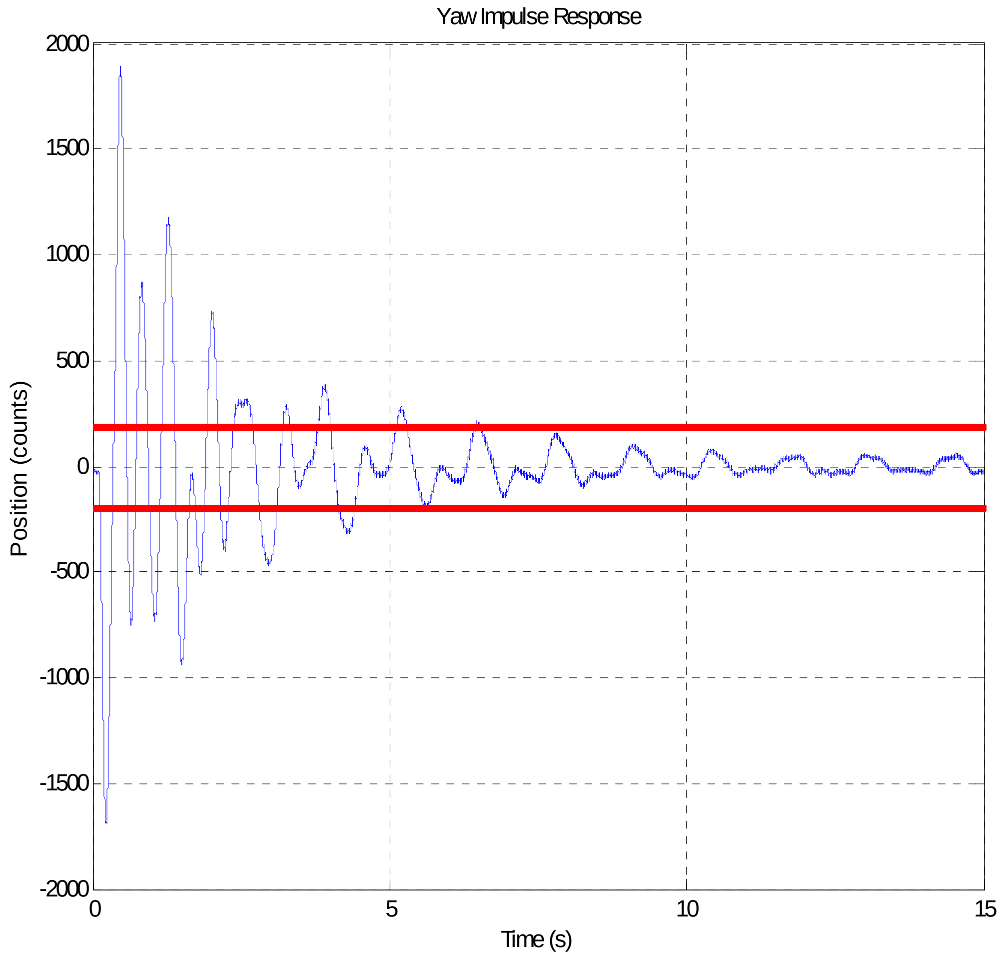
Yaw Impulse Response
2000
1500
1000
500
0
-500
-1000
-1500
-2000
Position (counts)
0
5
10
15
Time (s)

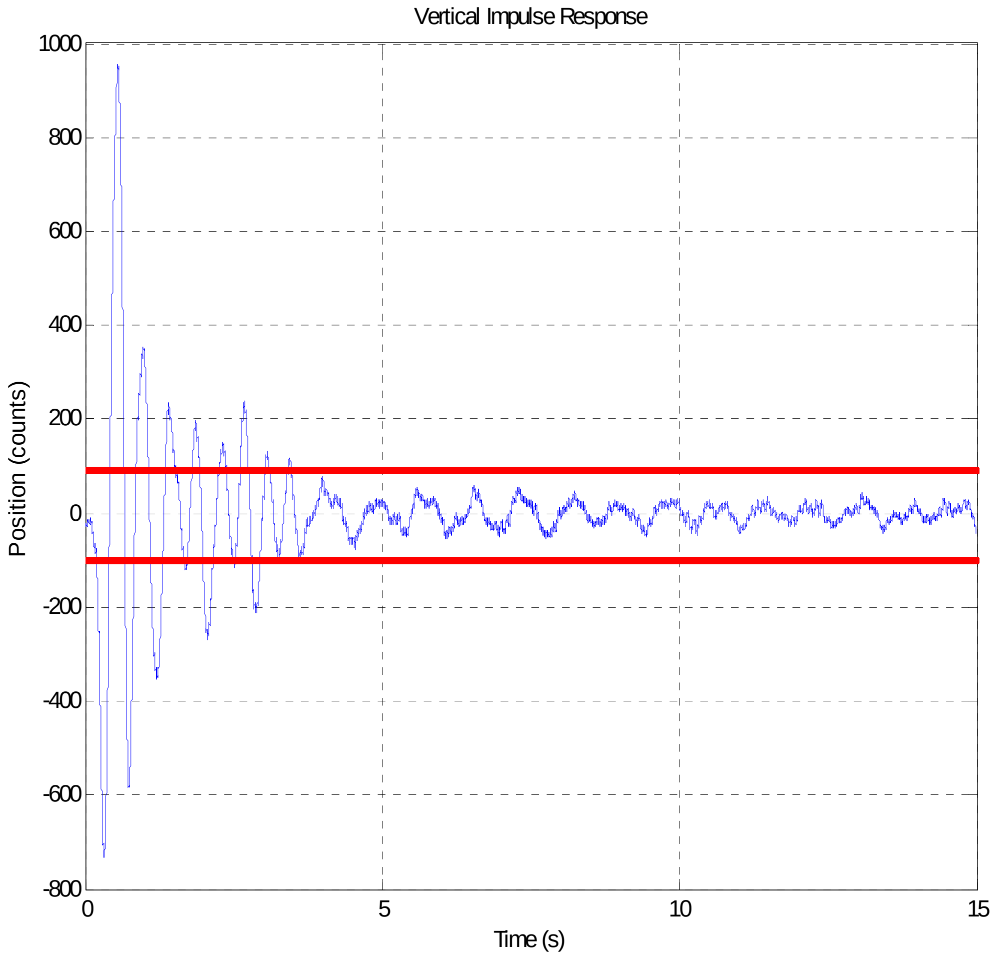

Vertical Impulse Response
1000
800
600
400
200
0
-200
-400
-600
-800
Position (counts)
0
5
10
15
Time (s)

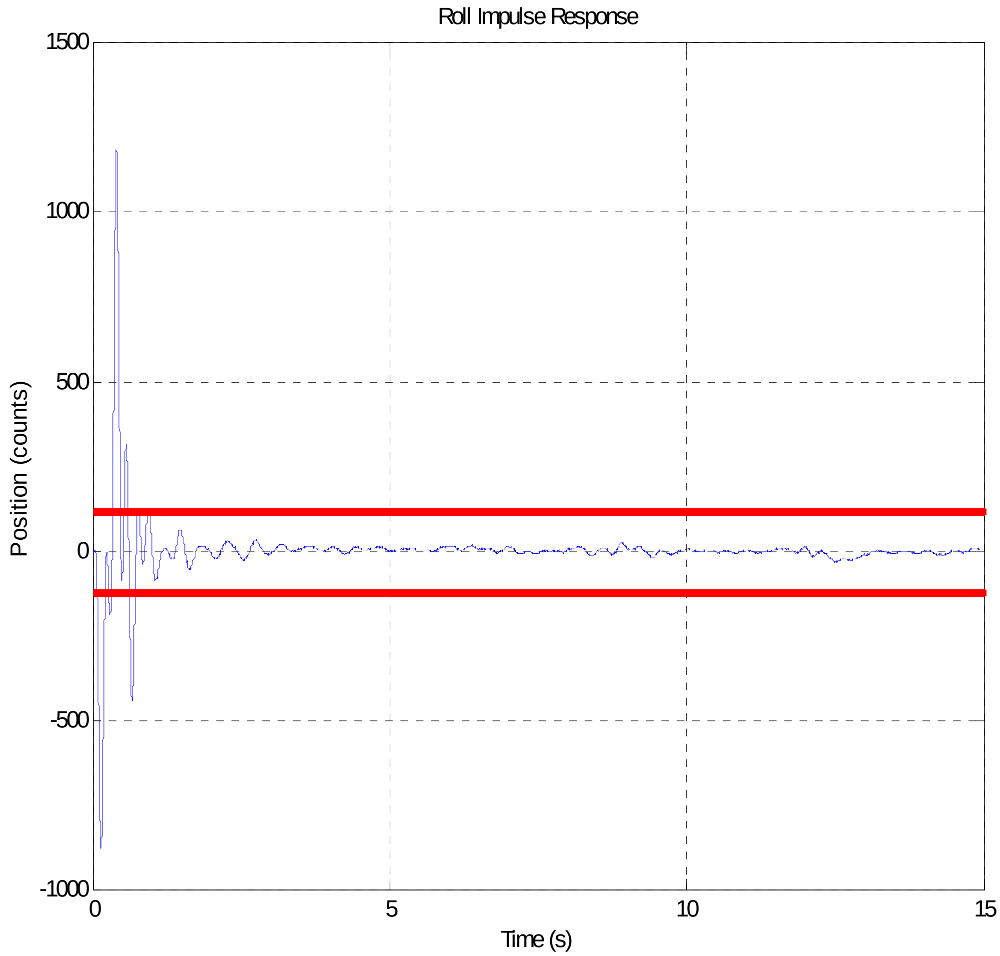
Roll Impulse Response
Position (counts)
1500
1000
500
0
-500
-1000
0
5
10
15
Time (s)

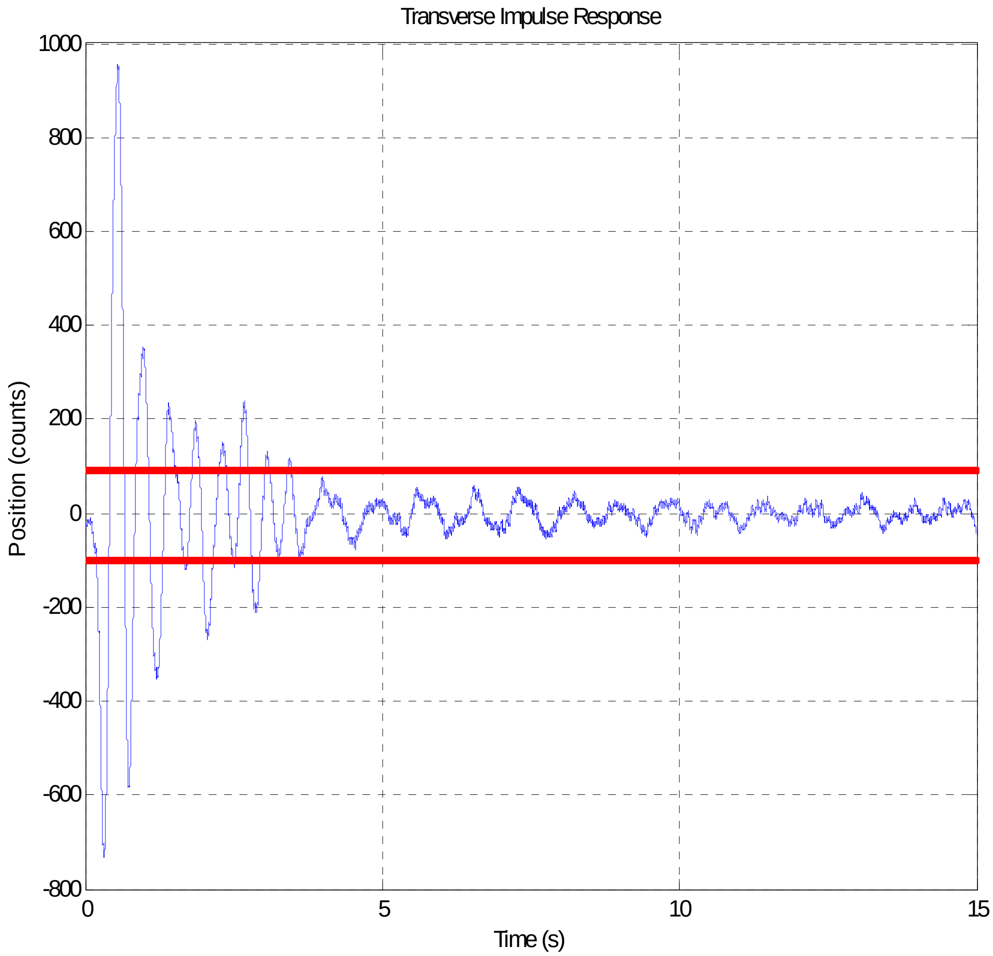

6 Thermal Load Test

This test was done to simulate the effect of the ring heater on the quad. 20 W were injected around the main chain test mass for 24 hours with 6 heating resistors at the locations illustrated below. The photograph shows the locations of where the temperature was measured, along with the corresponding channel numbers which are necessary to interpret the test results on the third figure.

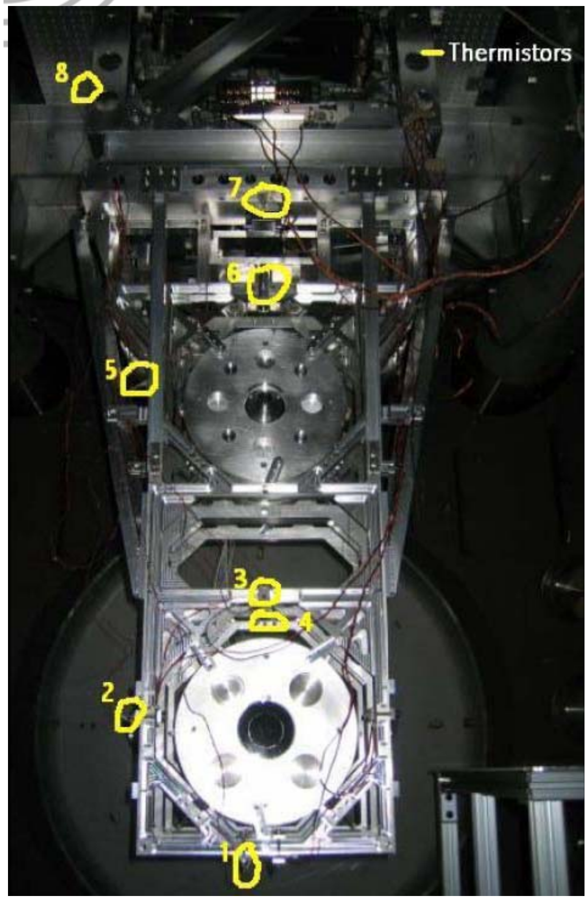
Thermistors
8
7
6
5
3
4
2
1

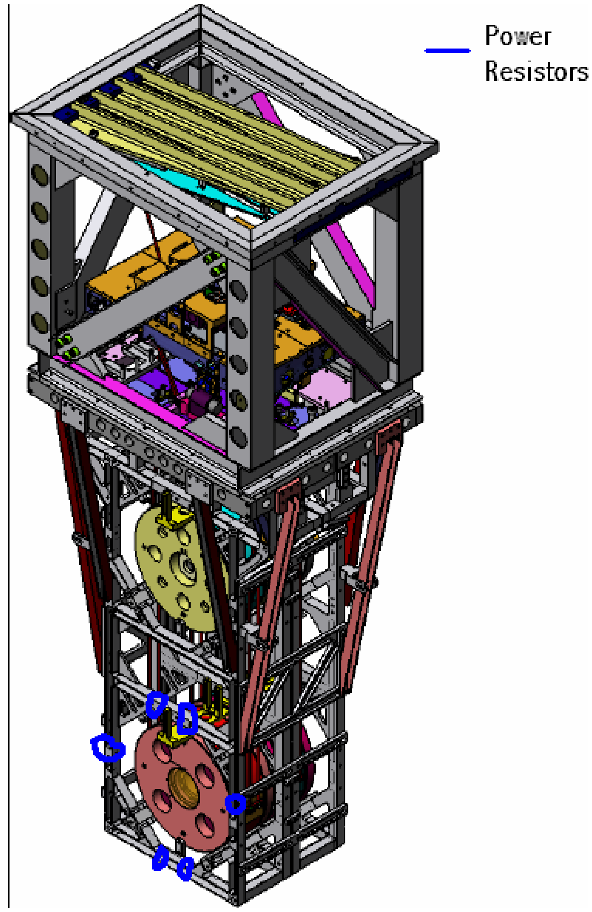
Power
Resistors

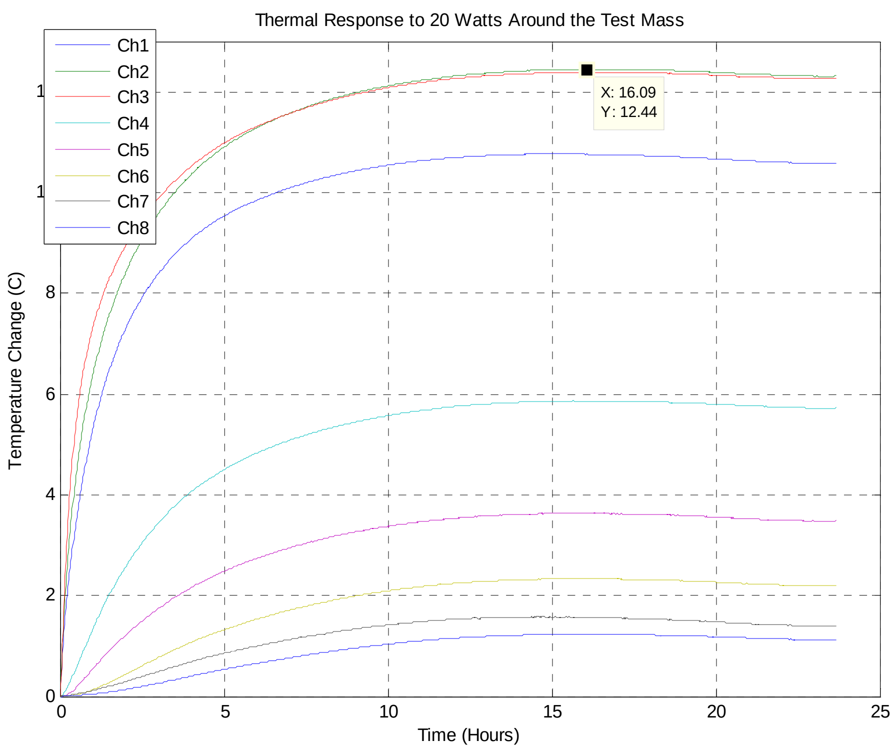
Thermal Response to 20 Watts Around the Test Mass
Ch1
Ch2
Ch3
Ch4
Ch5
Ch6
Ch7
Ch8
X: 16.09
Y: 12.44
Temperature Change (C)
0
2
4
6
8
0
5
10
15
20
25
Time (Hours)

The maximum temperature rise is 12.4 C near the bottom of the structure. The slight drop towards the end of the experiment is likely due to the large temperature fluctuations experienced at LASTI. A point 6 inches from the bottom of the structure was observed in order to measure the structure's elongation. The maximum drop of this point was 0.007 +- 0.001 in (0.18 +- 0.025 mm).

It was predicted that the maximum temperature rise would be 26 C (see T060243-00). However, it is believed these measurements are accurate since all power resistors and thermistors made good contact being almost entirely surrounded by the structure and metal in contact with the structure. That being said, these results are encouraging since the injected heat has less effect on the quad than previously thought.

7 Mechanical/Electrical Difficulties

There were a few issues that came up during the testing of the controls prototype. The first is the electrostatic drive cables. As mentioned earlier in the cross-coupling section it is not obvious if the cables have tension and are pulling on the UI mass of the reaction chain. They enter the structure right in front of this mass near to where a lot of work takes place. Thus, it was very common for them to get bumped and then slip through the strain relief on the structure causing them to pull on the UI mass. An easy fix to this problem would be to have them leave the structure at a less vulnerable location or to provide sufficient cable so that they could be shifted around as needed. In this case the cables were simply unplugged from the feedthrough since they weren't in use anyhow.

A second problem was the installation of the Eddy Current Dampers. The first issue encountered with those was that some of the installation tooling has no function. Specifically a piece of the tooling appeared to be designed for horizontal alignment of the ECD array, however the only available horizontal motion was given in the set screws that hold the array in place. In fact the only way to make adjustments was to loosen the set screws and reposition every DOF of the array simultaneously. Since visibility was very poor it was not possible to be sure that the array was aligned correctly without testing. Further complicating this issue is that when you tighten the set screws to lock it in place the screws end up turning the array as well. The final positioning of the ECDs was achieved through a series of iterations of repositioning and testing until damping appeared to reach a maximum. The design of the ECDs on the noise prototype appears to address all these issues.

Some trouble occurred with the OSEMs as well. The biggest problem seems to be from cracking insulation on the wires inside the OSEMs, causing them to short to the structure. There were various instances of multiple OSEMs simultaneously shorting to the structure and thus themselves. Many LEDs were lost this way. The problem was solved by adding extra spacers inside the shorting OSEMs. The extra spacers provided room to lift the wires off the aluminum surface and also reduced the risk of further cracking. Another problem with these OSEMs is the cracking and shedding of the ceramic washers. Ceramic fragments have added a lot of dust and debris to quad and the inside of the BSC. The most vulnerable and difficult to clean sites are the OSEM pockets inside the penultimate mass, especially since fully locking this mass involves installing additional hardware. The design of the noise prototype OSEMs is different and should not have these problems. Refer to LIGO document T050111-01-K for a description of these new OSEMs.

Another notable issue was electronic limitations which are likely dominated by the coupling of the cables outside the BSC. Because of this problem measuring signals beyond 10 Hz became tricky. This is illustrated in many of the TFs included above. Beyond the highest mode the TFs should roll off at $1/f^2$ however, in many they appear to include a zero above 10 Hz and then start leveling off around 30 Hz. Clearly this is not part of the dynamics of the pendulum. The problem seemed to limit the measurements of the cross-coupled TFs because the OSEM signals were weaker in those cases.

8 Future Work

The next steps for the controls prototype will be to optimize the local control filters and then to apply the modal control with state estimation method of control developed by Laurent Ruet for the triple pendulum. Other testing will involve forming an optical cavity between the quad and one of the triples to directly observe the motion of the quad's bottom mass. The implementation of the modal control with state estimation will be the focus of my master's thesis.