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

Technical Note	**LIGO-T070272-00-0**	Date: 12/1/2007

Final Quad Controls Prototype Results

Brett Shapiro, Richard Mittleman

Distribution of this document:

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

This document provides the final results to the testing of the quadruple pendulum controls prototype at LASTI. It serves as an update to T070009-00, which already includes results on system identification, local control, eddy current damping, and thermal testing. Document G070694-00-D, by Lisa Barsotti and Matthew Evans provides results on the electrostatic drive testing. This document updates the system identification with higher resolution data, provides a greater understanding to the disagreement with model predictions, includes results on modal control with state estimation, discusses issues with cabling, and lists some of the changes made on the quad noise prototype.

2 System Identification

Section 2.1 contains transfer functions (TFs) measured from the 6 local control OSEMs on each of the top masses of the main chain and reaction chain. Section 2.2 provides a table comparing the mode (resonant) frequency measurements of each chain to that predicted by the model.

The model in this document is updated from T070009-00 to include off-diagonal moments of inertia. In other words, this model accounts for the fact that the principal axes of rotation of the rectangular masses do not line up exactly with the axes relative to the beam tube. Thus, for example, a pure 'pitch' motion about the top mass's 'y' axis is observed as both pitch and roll in the more global coordinates of the OSEMs, which coincide with the axes of the beam tube. The model exists in Matlab state space form. It was provided by Mark Barton in February 2007.

The discussion of sys-id in this section introduces controllability issues related to weakly coupled modes, the proximity of certain modes to each other, and the effect of off-diagonal moments of inertia.

2.1 Undamped Transfer Functions and Mode Frequencies

Measurements were taken for all 6 diagonal degrees of freedom (DOFs) of each mass, ie X to X, Y to Y, etc. For each DOF, the measurements for the main chain (red) and reaction chain (green) are plotted together with the model predictions (blue) in the form of a bode plot. The model curve is artificially shifted vertically in the magnitude plot so that the shapes of the curves are easier to see. See Figures 1 to 6 below.

The TFs were all measured with a resolution of approximately 1 mHz using a white noise injection from 0 Hz to 10 Hz. No damping of any kind is used in these measurements.

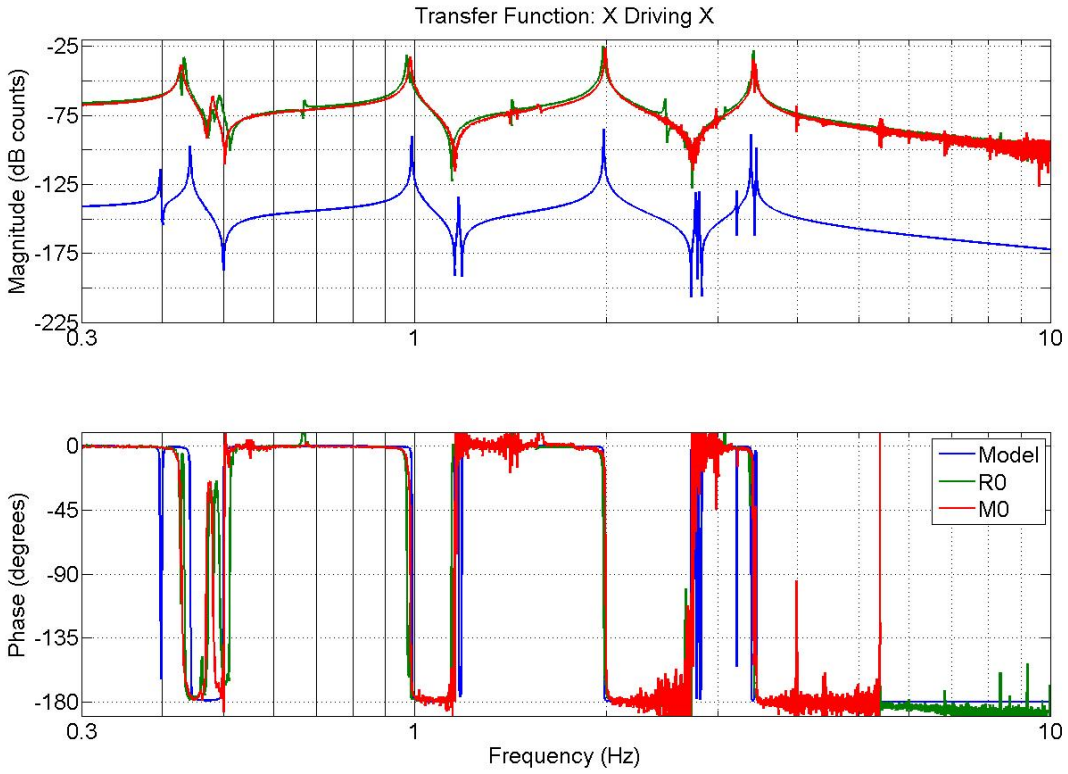

Figure 1: Measured TFs of the X to X (along the beam path) DOF on the main (red) and reaction chains (green). The predicted response from the model (blue) is plotted as well for comparison. The model curve has been artificially shifted vertically for clarity.

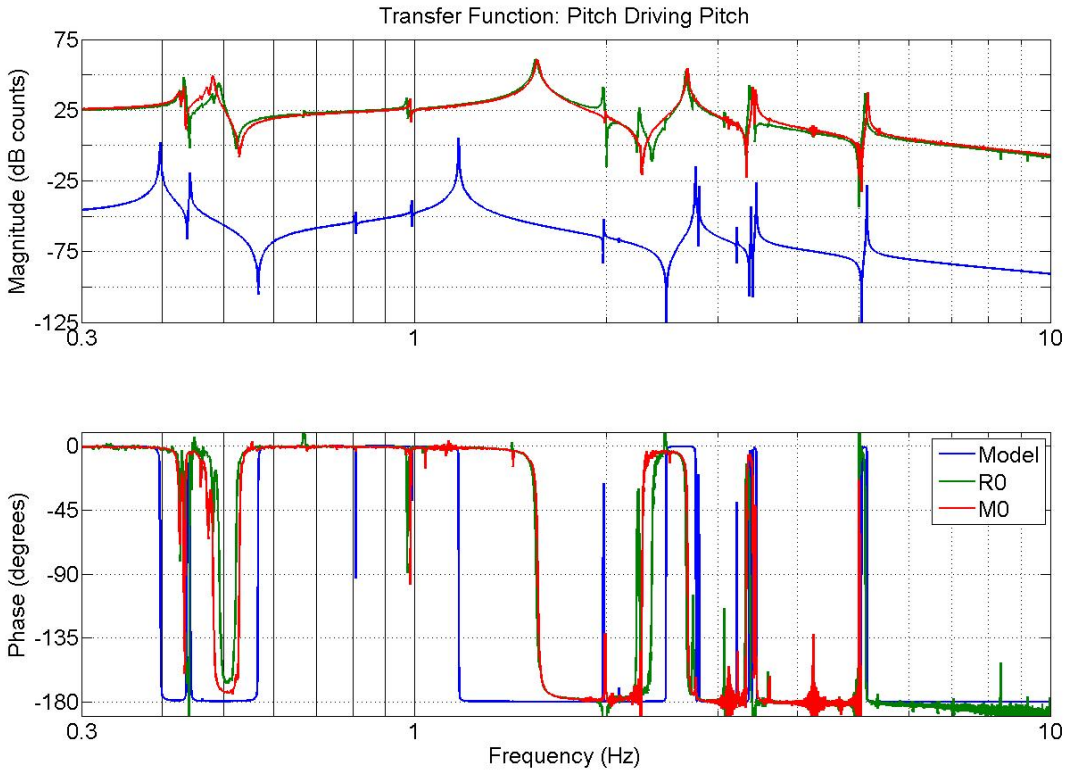

Figure 2: Measured TFs of the pitch to pitch (about y-axis) DOF on the main (red) and reaction chains (green). The predicted response from the model (blue) is plotted as well for comparison. The model curve has been artificially shifted vertically for clarity.

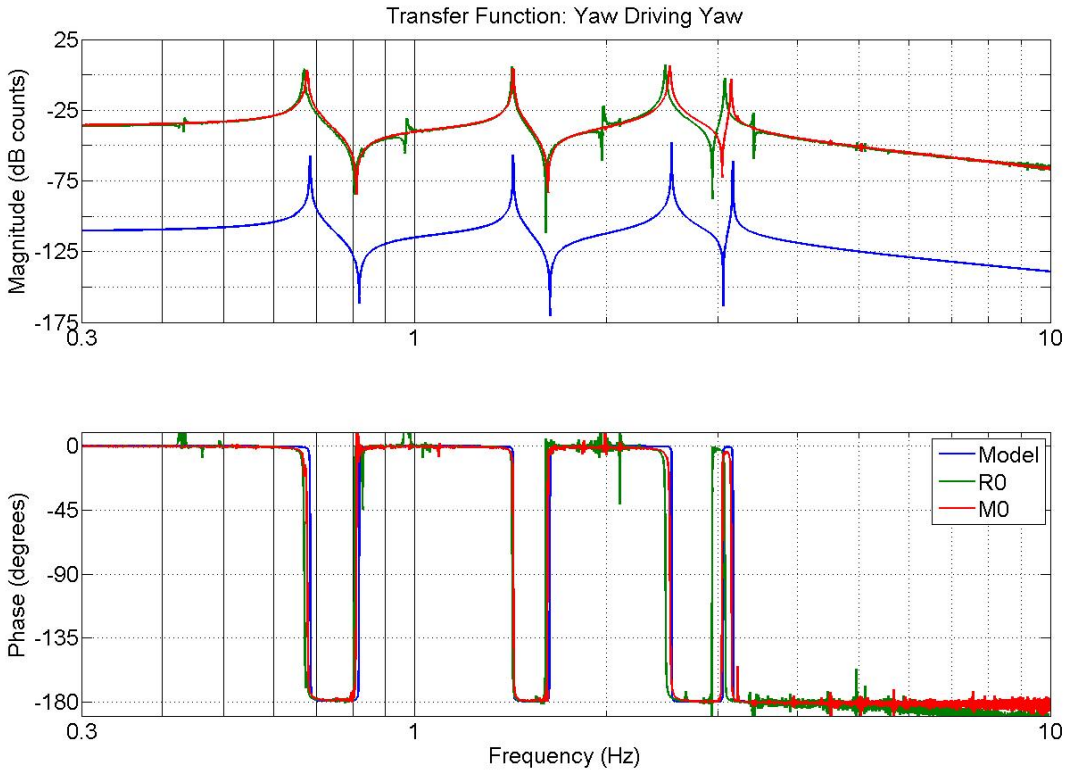

Figure 3: Measured TFs of the yaw to yaw (about z-axis) DOF on the main (red) and reaction chains (green). The predicted response from the model (blue) is plotted as well for comparison. The model curve has been artificially shifted vertically for clarity.

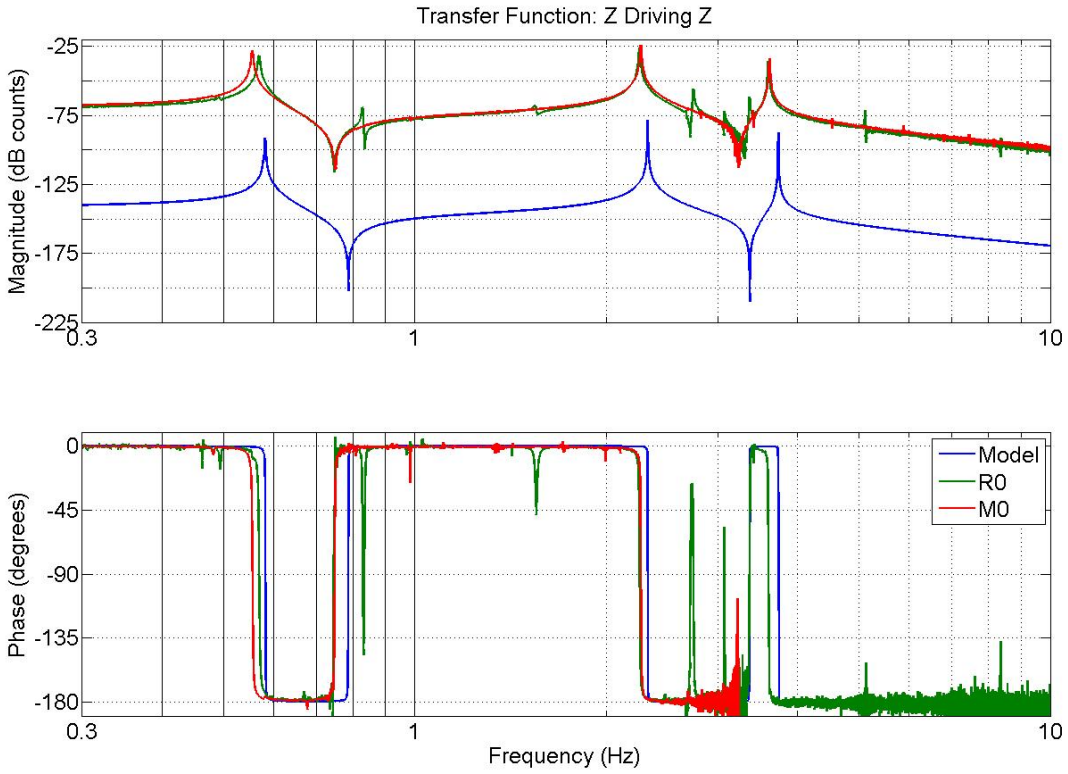

Figure 4: Measured TFs of the Z to Z (vertical) DOF on the main (red) and reaction chains (green). The predicted response from the model (blue) is plotted as well for comparison. The model curve has been artificially shifted vertically for clarity.

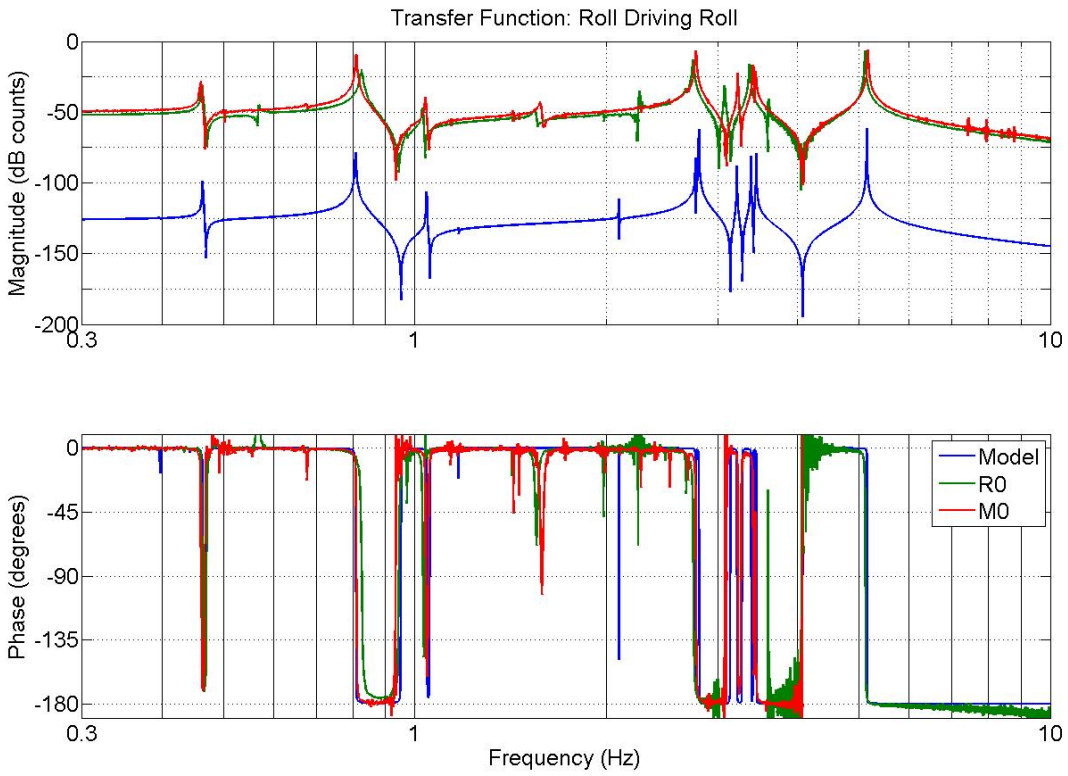

Figure 5: Measured TFs of the roll to roll (about x-axis) DOF on the main (red) and reaction (green) chains. The predicted response from the model (blue) is plotted as well for comparison. The model curve has been artificially shifted vertically for clarity.

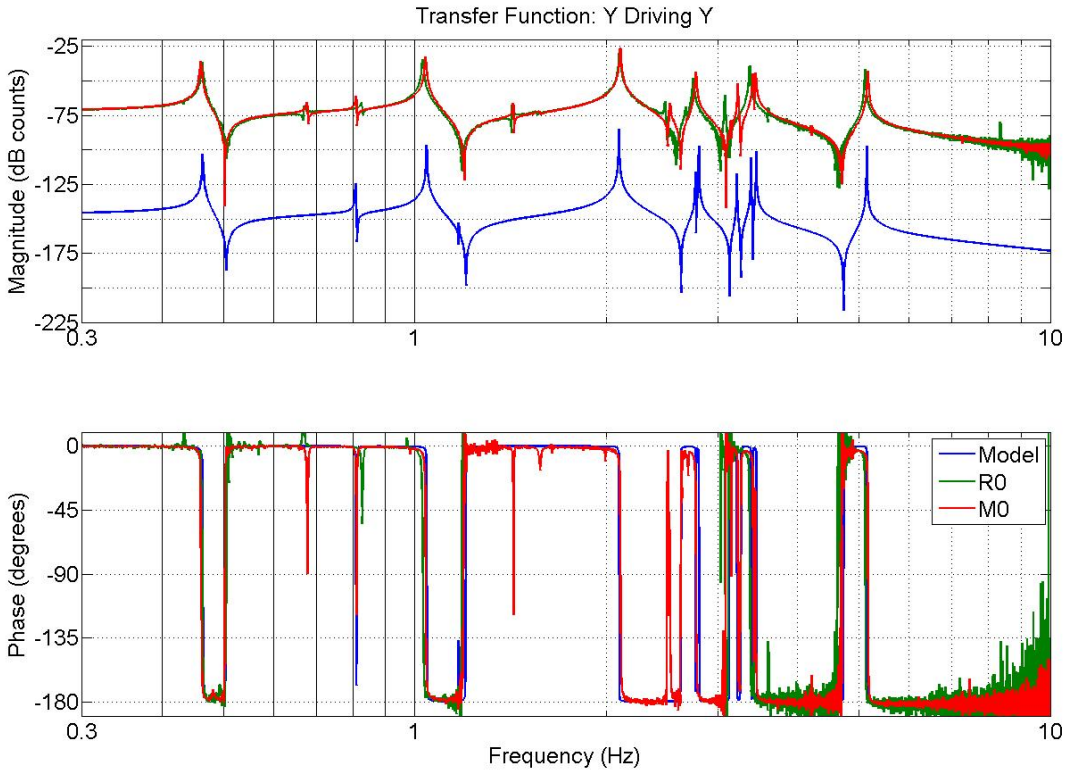

Figure 6: Measured TFs of the Y to Y (transverse to beam) DOF on the main (red) and reaction (green) chains. The predicted response from the model (blue) is plotted as well for comparison. The model curve has been artificially shifted vertically for clarity.

Most of the modes line up fairly well on both chains and the model. The largest discrepancy between the measurements and the model occurs at the first 2 pitch modes, with the second being greater than the model by 25% to 30%. Most of this error in both modes is explainable by a drop in the blade tip heights (parameter dn in the model) in the top masses on the order of 3 mm. An accurate measure of this height was not attainable on the actual suspension, but this number is believed to be accurate within 1 to 2 mm by adjusting various model parameters. Differences in other parameters likely contribute to the rest of the discrepancies such as differences in mass, stiffness, or other wire break off positions. Various attempts to fit the model to the measurements were made with various success. See Section 3 for a brief discussion on model fitting. Section 2.2 has a table that more accurately lists the measured positions of the modes and compares them to the model.

The reaction chain TFs show a lot of cross coupling. Much of this is likely due to effects from the OSEM and electrostatic drive cabling. The bulk and weight of the OSEM cabling, along with the stiffness of the ESD drive cabling proved to have a significant effect on the reliability of the reaction chain. This effect could possibly be mitigated with careful clamping and routing techniques. Section 5 is devoted to a discussion of this issue.

In terms of controllability, one of the modes couples very weakly to the top mass. This is a mode at about 3.2 Hz. It is most visible in the roll plot (Figure 5) but also in the y plot (Figure 6). Most of the mode's energy is in the lower masses, including the test mass. Thus, to effectively reduce the RMS contribution to the optic, a lot of gain needs to be used in the local control, which can introduce sensor noise. Other issues of controllability result from the proximity of some modes to one another. Proximity is mostly a problem for modal control, which is discussed in Section 4.

2.2 Mode Frequency Comparison Table with Model

Here Table 1 shows a more accurate comparison between the mode frequency measurements and the model. This data was collected directly from the TFs above. The TFs and data in the table were measured with a resolution of approximately 1 mHz.

The structure of this table also highlights the effects of the off-diagonal moments of inertia. Originally, like the triple, the dynamics were expected to be dividable into four separate, independent systems, yaw, z (vertical), x-pitch, and y-roll. Accordingly, the table lists the modes under categories of these systems. However, many of the higher frequency x-pitch and y-roll modes, ie above 2.5 Hz, do not dominate in either of these two systems. Thus, there is another category, x-y-pitch-roll, to list these modes that really must be thought of as existing in both systems. The choice between what qualifies a mode for being listed in one category versus another is somewhat arbitrary and open to interpretation, but making these distinctions helps give a more intuitive feel for the behavior of the quad at various frequencies.

It is more obvious in this table then in some of the TFs that some modes are extremely close together. Referencing the main chain, the most notable pair is the couple at 3.407 Hz and 3.437 Hz. In fact, it was not until the resolution of the TFs was taken to mHz that it was possibly to see there were in fact two resonant frequencies right on top of each other.

It is also difficult to be sure which of the model modes these line up with since they are so close, though this probably does not really matter for the same reason. The 2.688 Hz and 2.752 Hz modes have similar issues with proximity.

It is also noteworthy that it is a result of the off-diagonal moments of the inertia that these modes are so close together. The coupling of the x-pitch and y-roll systems creates a 'nest' of modes at the high frequencies. If these systems were completely independent then one mode in the pair would be x-pitch and the other y-roll, and they would not interact. In principle, the fact that they are so close in the coupled system does not really matter since they can still be controlled, until one attempts to do modal control. Theoretically in modal control it does not really matter either. However, modal control requires an estimator with a good model of the system to predict which modes are which so they can be damped effectively and individually. The fact that they are so close means that the estimator's model needs to be that much better to be able to tell them apart. Complicating this issue, model fitting becomes more difficult in the coupled system as well since there are more parameters to consider in the fit. If the systems were uncoupled, the parameters of the x-pitch system could be fit independent of the y-roll system. In the fitting methods tried thus far, the difficulty seems to increase exponentially with the number of parameters. More about model fitting is discussed in the next section.

X-Pitch					**Y-Roll**				
Pred.	Main	Err (%)	React.	Err (%)	Pred.	Main	Err (%)	React.	Err (%)
0.398	0.48	17.19	0.492	23.67	0.464	0.458	-1.29	0.463	-0.24
0.443	0.427	-3.61	0.436	-0.35	0.807	0.803	-0.50	0.825	2.23
0.989	0.980	-0.91	0.972	-1.75	1.044	1.039	-0.48	1.029	-1.44
1.173	1.583	25.93	1.552	32.31	2.096	2.097	0.05	2.102	0.29
1.983	2.000	0.86	1.980	-0.15	25.123	24.346?/25.340		-3.09?/0.86	

Yaw					**Z**				
Pred.	Main	Err (%)	React.	Err (%)	Pred.	Main	Err (%)	React.	Err (%)
0.685	0.674	-1.61	0.670	-2.19	0.581	0.552	-4.99	0.569	-2.07
1.429	1.424	-0.35	1.423	-0.42	2.325	2.278	-2.02	2.257	-2.92
2.536	2.504	-1.26	2.479	-2.25	3.733	3.613	-3.21	3.601	-3.54
3.165	3.135	-0.95	3.076	-2.81	17.332	17.201/17.418		-0.76/0.50	

X-Y-Pitch-Roll				
Pred.	Main	Err (%)	React.	Err (%)
2.765	2.688	-2.87	2.672	-3.36
2.800	2.752	-1.74	2.742	-2.07
3.213	3.210	-0.08	3.076	-4.26
3.381	3.407	0.77	3.365	-0.47
3.447	3.437	-0.28	3.341	-1.07
5.143	5.151	0.15	5.112	-0.603

Table 1: Frequency table showing the error between the measured mode locations compared to those from the model. All units are in Hz. The resolution of the measurements is approximately 1 mHz. The columns labeled 'Pred.' are the mode locations predicted by the model, the column labeled 'Main' contains measurements from the main chain, and the column labeled 'React.' contains the measurements from the reaction chain. The 'Err (%)' columns contain the percent error of their respective chains to the model. The x-y-pitch-roll section is a result of the off-diagonal moments of inertia. The off-diagonal moments of inertia cause the higher frequency x-pitch and y-roll modes to couple together. With purely diagonalized moments, these modes would seperate into x-pitch or y-roll. The predicted 17.2 Hz and 25.1 Hz modes were measured simultanously on the main and reaction chains and thus were not possible to define as coming from one or the other. As a result, the measurements and errors are listed together. The measurement of 24.346 Hz is listed with a question mark since the signal to noise ratio was slightly too low to definitively say that this is actually a resonant frequency.

3 Model Fitting

Attempts were made to fit the model to the observations of the actual suspension with varying degrees of success. The yaw and vertical systems were relatively easy to fit, and were done so successfully using a simple gradient descent method. The parameters of the model likely to have the most error such as mass, moment of inertia, and blade stiffness were altered until the error in the resonant frequencies was minimized. The new best fit model then had resonant frequency error of less than one percent for each mode.

The x-y-pitch-roll system proved much more difficult. The number of parameters with uncertainty and number of modes to fit was too much for a simple gradient descent method to handle. The blade tip heights had uncertainty on the order of millimeters, and some of the masses turned out to be off be hundreds of grams. A Monte Carlo like method was tried as a rough fit method in order to bring the model 'close' to the measured frequencies. In this method, the parameters were simply changed randomly many times within a given range of uncertainty. The combination of parameters that yielded the smallest error became the rough fit. This rough fit could then be used as the starting point for the gradient descent method. The advantage of the Monte Carlo fit is that it is extremely stable and it cannot get stuck in local minima. However, it has no way of intelligently searching for the real solution and must run through many iterations before converging to something reasonable. The gradient descent method can then finish the fit off to bring the error down even further.

The combination of these two methods actually can produce a fit where all the modes line up fairly well with measurements. An example of a pitch TF from a model fit using this hybrid method is shown below.

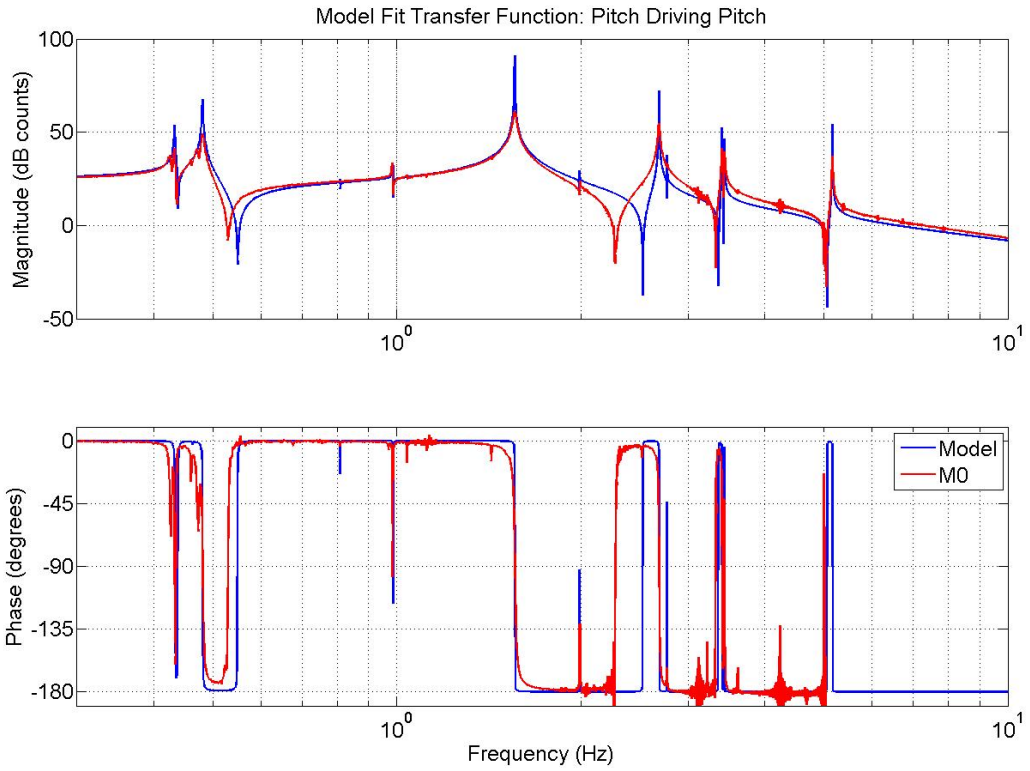

Figure 7: Model pitch TF after fit (blue) compared to the M0 pitch TF (red). The fit appears pretty good, and in fact all the modes line up within 1.5%. However, the fitting code lacks information about the modes shapes and TFs to other stages will not appear the way they should. This fit is not adequate for modal control without correct mode shapes.

Here the fit actually appears pretty good. The model's modes even line up to within 1.5% of the measured modes. So it appears that this fit may actually be a success. However, this model still fails when put up to the test of modal control. The reason it fails is because the fit contains no information about the mode shapes. For this system, the mode shapes are rather complex, since they vibrate in 16 DOFs each. It seems that there are multiple solutions to the model fit that produce the correct resonant frequencies, but yet have differing mode shapes. Measuring the complete mode shapes is not possible with the configuration of the OSEMs on the lower stages. However, even introducing the information of the measurable parts of the mode shapes into the error calculation quickly overflows the fit code with so many parameters that nothing gets adequately fit. Since these are the only fitting methods tried here, it is very possible that other fitting methods will have more success. The yaw and vertical systems appear to be simple enough that simply using the mode frequencies in the error calculation is sufficient to produce a model good enough for modal control. Ideally, in future suspensions the parameters will be well known enough that fancy fitting methods are not necesarry.

4 Modal Control with State Estimation

This document is not long enough to give a good introduction to modal control. For more detail, one should refer to Brett Shapiro's master's thesis (document number P070144-00-Z) for the quadruple pendulum or Laurent Ruet's PhD thesis (document number P070054-00-Z) for the triple.

Modal control allows the pendulum to be mathematically broken down into a series of second order single DOF systems resonant at the pendulum's modes. Then, instead of designing a small set of large complex filters, the pendulum can be damped with a larger set of simple filters to meet both the requirements for damping and sensor noise injection. In fact, the design of the filters can be repeated for all the modes. Thus, the design can be done once for an arbitrary undamped second order oscillator and simply shifted in frequency and gain to meet the needs for each mode.

Another benefit of modal control is that the gains on each modal controller can be tailored to optimize the damping versus noise injection performance. In general, the lower frequency modes contribute the most to settling time whereas the higher frequency modes contribute the most to sensor noise injection. Thus, the gains on the lower modes can be turned up, while the gains on the higher modes turned down, keeping settling time the same but significantly reducing sensor noise.

Still another benefit is introduced with state estimation. State estimation is necessary to reconstruct the DOFs of the pendulum that are not observed by the OSEMs. The estimator uses a model of the pendulum to predict what these DOFs are doing in real time and uses this information to fill in where the sensors are lacking. This lack at first may seem like a draw back, however estimators have the property of further reducing the sensor noise injection if designed properly. They do however, have the disadvantages of requiring a relatively accurate model and a sufficiently powerful computer to run in real time.

The two key equations that make modal control possible are listed below.

$$q = \Phi^{-1} x \tag{1}$$

$$P = (\Phi^T)^{-1} P_m \tag{2}$$

Eq. 1 shows the transform from the 'real' motion, x, sensed by the OSEMs on the top mass. Φ is the basis of eigenvectors for the system. The transformed motion, q, is the modal state. In reality, the full state x is incomplete, so this transform is approximated by the estimator. The control filters then act on this modal signal and generate the modal damping forces P_m. The modal damping forces then need to be transformed back into 'real' damping forces that are useful to the pendulum. This second transformation is done through Eq. 2. A sample schematic of what a modal control layout with state estimation looks like is shown below in Figure 8.

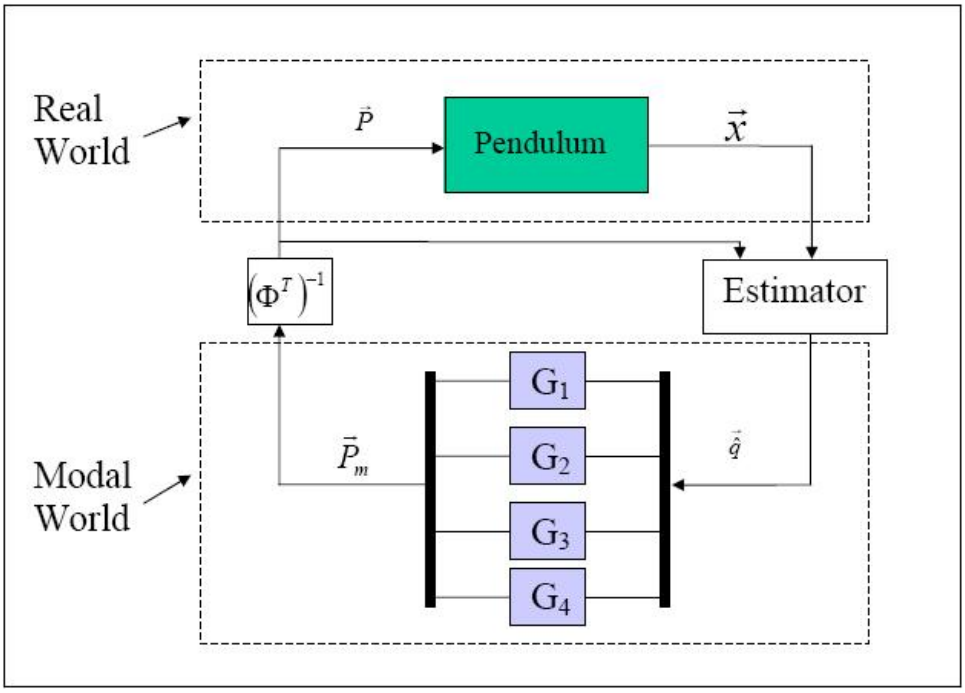

Figure 8: Schematic diagram of a modal control loop. The incomplete sensor signals $\vec{x}$ are reconstructed and simultaneously transformed into modal signals $\vec{q}$ and passed through their respective control filters G_i. The resulting modal forces $\vec{P}_m$ are then transformed into real forces $\vec{P}$ that are applied to both the pendulum and the estimator. The feedback of the modal forces reduces the estimator's error by giving the model more complete information about the closed loop system.

An example of the modal control filter design is shown below in Figure 9. The left part of the figure shows an example of a TF for an idealized 1 DOF undamped oscillator. This is the shape of each mode's TF assumed when designing the modal controller. The resulting design is shown on the right. It is rather simple, consisting of an AC-coupled zero, 2 poles at twice the resonant frequency and a small gain bump (2 more complex poles and zeros) right at the resonant frequency to improve damping.

Actual tests of the damping performance of modal control were conducted on the yaw and vertical systems of the controls prototype. The main limitation preventing the testing of modal control on the x-y-pitch-roll system was the inaccuracy of the model. Figures 10 and 11 show TF and step response experimental results on the yaw system. As shown, the experimental results line up very well with those of simulation, indicating both that the method is (at least partially) applicable to a quadruple pendulum and that the models for these systems are sufficiently close to the real system. Figures 12 and 13 show similar results on the z system.

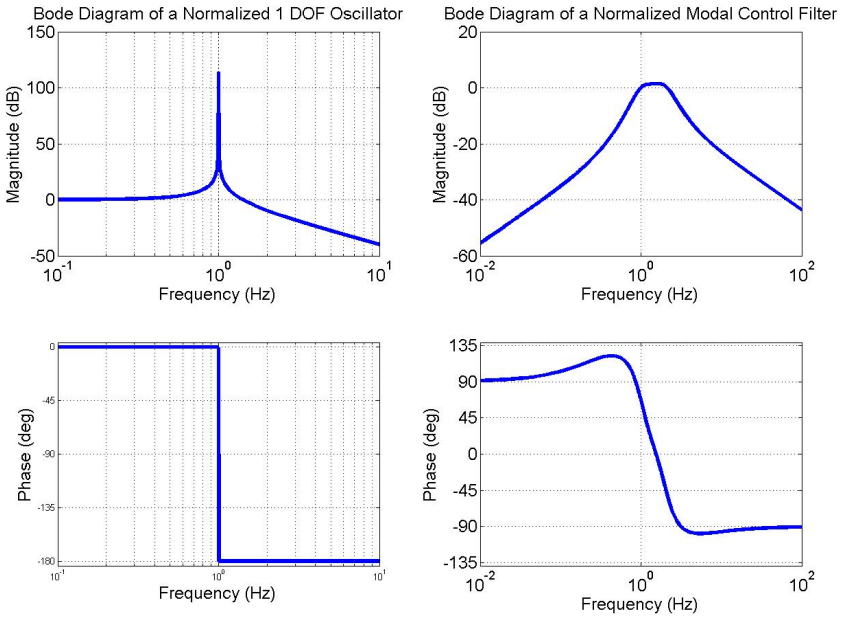

Figure 9: Left, a modal plant as a simple oscillator normalized to 1 Hz. Right, a standard modal control filter normalized to damp at 1 Hz.

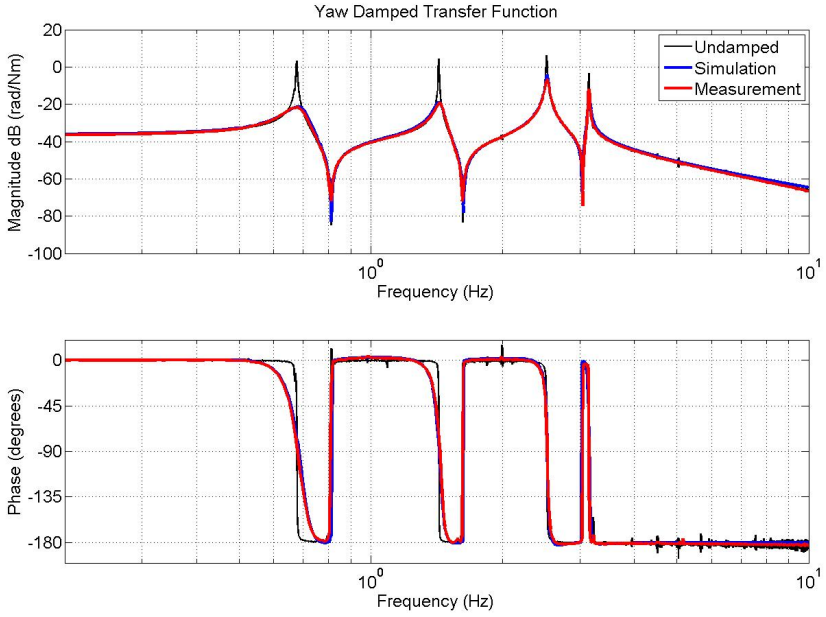

Figure 10: The red curve is the measured yaw TF of the top mass with modal control damping on. The damping gains are set to achieve a settling time of less than 10 s. The blue curve is a simulation of the same test. The black curve, the measured undamped response, is a reference to show peak reduction.

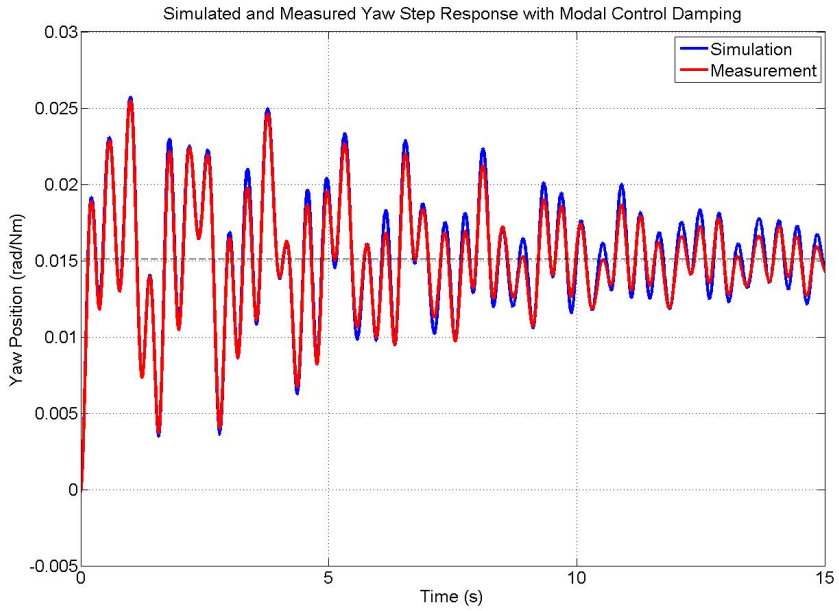

Figure 11: The red curve is the measured yaw step response of the top mass with the modal damping on. The damping gains are set to achieve a settling time of less than 10 s. The blue curve is a simulation of the same test. In this case the settling time is about 8.15 s.

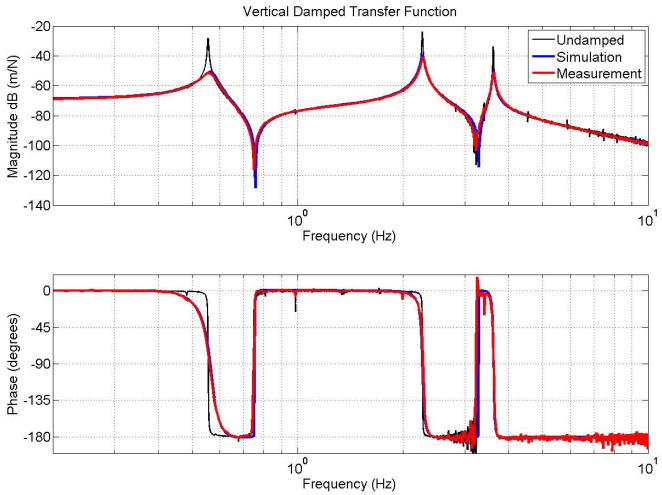

Figure 12: The red curve is the measured vertical TF of the top mass with modal control damping on. The damping gains are set to achieve a settling time of less than 10 s. The blue curve is a simulation of the same test. The black curve, the measured undamped response, is a reference to show peak reduction.

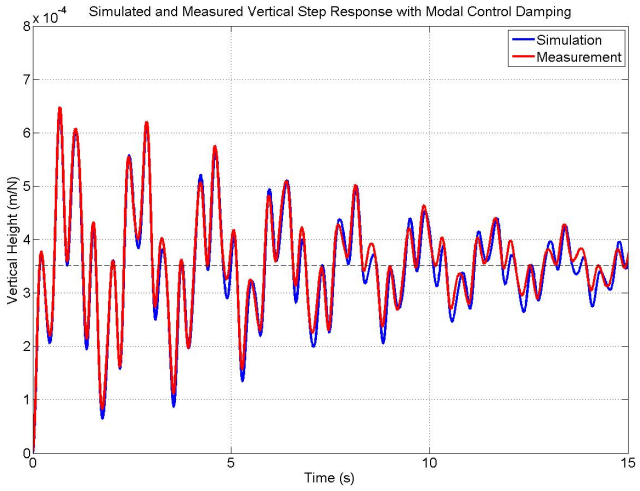

Figure 13: The red curve is the measured vertical step response of the top mass with the modal damping on. The damping gains are set to achieve a settling time of less than 10 s. The blue curve is a simulation of the same test. In this case the settling time is about 8.5 s.

5 Cabling Issues

Most of the cable-related problems experienced on the controls prototype centered around the location where the cables went from the strain relief on the structure to the reaction chain UI mass (note: the cables first entered the suspension at the UI mass, not the top mass). Most of these issues were a result of the large amount of loose slack on top of the reaction UI mass, stiffness of the ESD cables, and insufficient strain relief from the structure. As a result, small shifts in the position of the cables on the UI mass or in the strain relief caused a host of problems such as offsets in pitch, rubbing of cabling along the structure, pulling on the masses, and cross coupling between normally unrelated DOFs. Nearly every vent some OSEMs had to be realigned before closing up. Realigning OSEMs on the UI and penultimate masses is significant work since they are connected directly to suspended mass.

At first a number of causes were considered for the changes in alignment during vents, such as temperature effects in the blade springs. However, further experience with the controls prototype proved the unreliability of large amounts of loose cabling on the reaction chain in the presence of people poking around the suspension during vents, even when all the work was done on the *main* chain. Figure 14 shows photographs of some of these problematic locations. A reduction of weight in the cabling, like what has been done in the noise prototype by removing the shielding, will certainly help these types of issues and make the suspension more user-friendly. However, if massive cables like those on the controls prototype are necessary, then beneficial results may still be attainable through carefully thought out cable routes and strain relief along the structure and suspension.

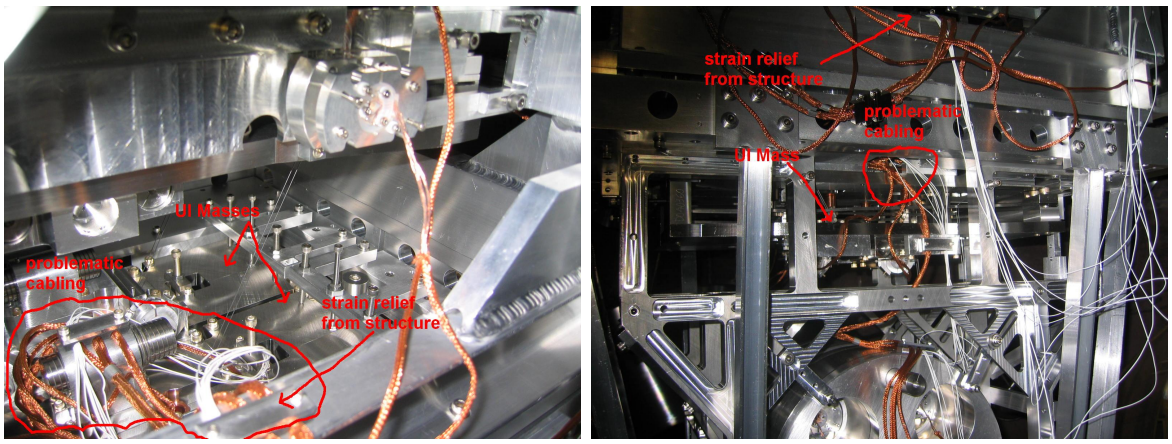

Figure 14: OSEM and ESD cabling between the structure and UI mass. The problematic areas of cabling are circled in red. Most of the issues centered around the loose slack around the UI mass combined with the weight of the OSEM cables and stiffness of the ESD cables. The labeled strain relief clamp in both pictures is where the cabling went from the fixed structure to the reaction chain UI mass.

During one of the lasts vents for the controls prototype, some time was taken to try and rework the cabling around the structure and suspension. Some of the problematic slack around the UI mass was pulled through. Additionally, the strain relief clamp that introduced the cabling to the suspended mass from the structure was relocated. Previously, the cabling was introduced horizontally from the structure to the top of the UI mass, which caused the pitch of the suspension to be very dependent on how much slack there was between the clamp and mass. The new position was more directly above the mass (from the tablecloth). These changes had some success in that the following vent no OSEMs had to be realigned, for one of the first times in its history. However, cross coupling issues are still present, as is shown in the TFs in Section 2.1.

6 Changes from the Controls Prototype to the Noise Prototype

Currently, the quad noise prototype incorporates some changes that should improve upon some of the issues experienced with the controls prototype and make it more user-friendly. First, the shielding from the OSEM cabling has been removed to make them lighter. The cables will still need to be tested to prove that the lack of shielding does not cause other issues. Second, the blade tips have height adjusters and reference points. Not only can these heights be easily measured but they can also be adjusted. Third, the OSEMs and ECDs have been given more lateral adjustments so that they can follow the suspension in pitch if it is adjusted. It is also noteworthy that the ECDs and OSEMs are now connected to the same lateral adjustment plates so that if everything is correctly assembled, aligning the OSEMs also simultaneously aligns the ECDs. Fourth, some eddy current damping has been shifted

from z and roll motion to include x, pitch, and yaw. The number of magnets used is the same, however they are now spread more evenly around the top mass. This new damping layout should improve the stability of modal control since less gain will be needed on the more problematic higher frequency x-y-pitch-roll modes.