

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

Technical Note	LIGO-T0900426-v1-0	Date: 9/1/2009
<h1>Quad Noise Prototype Results Summary</h1>		
Brett Shapiro and John Miller		

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	Applicable Documents	1
2	Introduction	2
3	Eddy Current Damping (ECD)	3
4	Ring Heater Testing	3
5	Cavity Work	6
5.1	Parametric Instability	6
5.2	Charging Measurements	7
6	Mechanical Impact of Electrical Cabling and Clamping Issues	8
7	Model Fitting	13
8	Modal Damping	17
9	Procedural Notes	19
10	Ongoing work	19

1 Applicable Documents

1. Summary of ECD Results on the Quad Noise Prototype - T090245-00
2. Quad Noise Prototype Cable Routing - T0900380-v1
3. Charge Measurement and Mitigation for the Main Test Masses of the GEO 600 Gravitational Wave Observatory - published 2007 in Classical and Quantum Gravity by M Hewitson et. al. Abstract at <http://adsabs.harvard.edu/abs/2007CQGra..24.6379H>

2 Introduction

This document summarizes the results obtained from the quad noise prototype at LASTI as of September 1, 2009. These results include eddy current damping, ring heater testing, cavity work, the mechanical impact of the electrical cabling, model fitting, and modal damping. A note about procedures developed during mechanical work is also included towards the end of the document. The monolithic effort is not included since that work is being led by the Glasgow group.

The list of 'Applicable Documents' are technical notes with more detail about some of the topics covered in the sections below. T090245-00 'Summary of ECD Results on the Quad Noise Prototype' is the primary source for all the ECD results, T0900380-v1 'Quad Noise Prototype Cable Routing' provides a list of photographs characterizing the OSEM and ESD cable routing from the ISI table to the test reaction mass.

3 Eddy Current Damping (ECD)

Document 1 in the applicable documents list, "Summary of ECD Results on the Quad Noise Prototype - T090245-00" includes all the details of the eddy current damping tests. The basic conclusion was that the dampers were giving on the order of two times more damping than was desired. The over damping was due to the effect of both longitudinal and transverse motion of the magnet inside the copper block, the latter of which was previously unmodeled.

Another issue with the dampers was that they were difficult to install because of the small clearance of the magnets inside the copper block holes. Exacerbating the clearance issue was the fact that the magnet is mounted inside of an aluminum cup with a certain wall thickness that decreases the clearance further. Consequently, the Suspensions Group decided to solve the clearance issue and at least partially the damping issue simultaneously by opening the diameter of the copper holes from 12mm to 13mm. The extra millimeter turns out to have a minimal effect on damping. However, since the copper blocks are not suspended there is freedom to reduce the damping further by either retracting the blocks, or pulling some out altogether.

4 Ring Heater Testing

The ring heater was tested to measure the thermal response of the quad structure. Measurements were done at two different power levels, 20W and 40W for about 24 hours each. The response of the structure was measured with $10k\Omega$ thermistors placed in seven locations from top to bottom. The locations of these thermistors is shown in Figure 1 below. Figures 2 and 3 show the results from the 20W and 40W tests respectively. The closest thermistor to the ring heater, located at the very bottom of the structure, warms by about $0.8C/10W$.

During the first measurement the vacuum pressure increased briefly by about 10ntorr. No RGA was available at the time to measure the components of this outgassing. However, a measurement done during a similar controls prototype test showed primarily water outgassing. Additional tests using a theoretical upper limit on the power are planned. This upper limit still needs to be determined, but is believed that work done in Ryan Lawrence's thesis will provide insight. The upper limit will be determined by the spectrum emitted by the heater. Higher powers tend to be dumped into visible light, which does not deposit thermal energy onto the optics.




Figure 1: The locations of the thermistors on the quad structure used to measure the response of the ring heater.

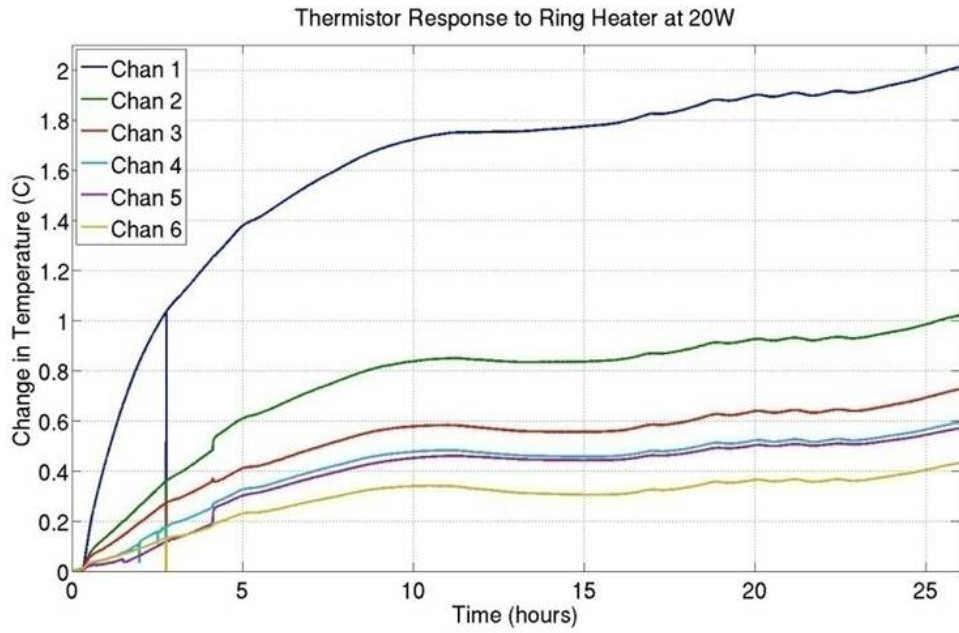


Figure 2: The response of the quad structure to the ring heater at 20W.

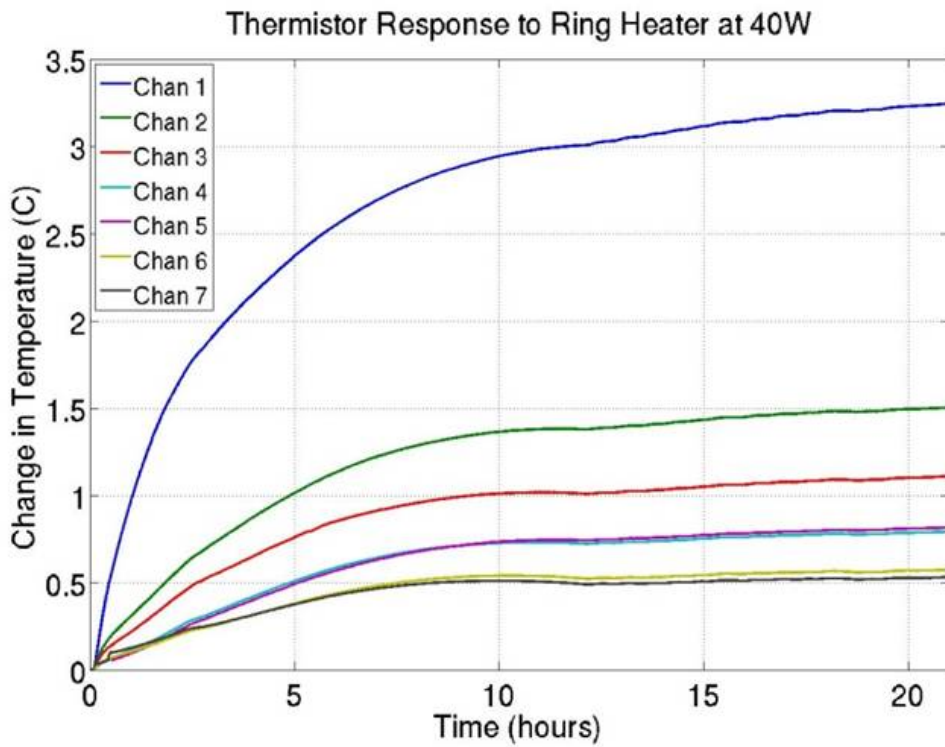


Figure 3: The response of the quad structure to the ring heater at 40W.

5 Cavity Work

The measurements with the quad-triple cavity at LASTI were done by John Miller with input from Lisa Barsotti and Matt Evans.

5.1 Parametric Instability

The high mechanical Q test masses coupled with the high resonant power in the Fabry-Perot cavities of advanced LIGO risks ringing up the test mass vibrational (acoustic) modes. See Figure 4. Simulations and measurements at LASTI show that these modes can be sufficiently damped using the ESD. See Figure 5.

This work has also shown that damping performance can be improved by redesigning the ESD gold pattern with slight asymmetries (often described as looking like 'mistakes'). A perfectly circular ESD pattern has symmetry similar to that of the shape of some of the vibrational modes of the test mass. Consequently the ESD may miss some of these modes by overlapping only with their nodes. These asymmetries break with the typical shape of these modes, reducing the required damping voltage.

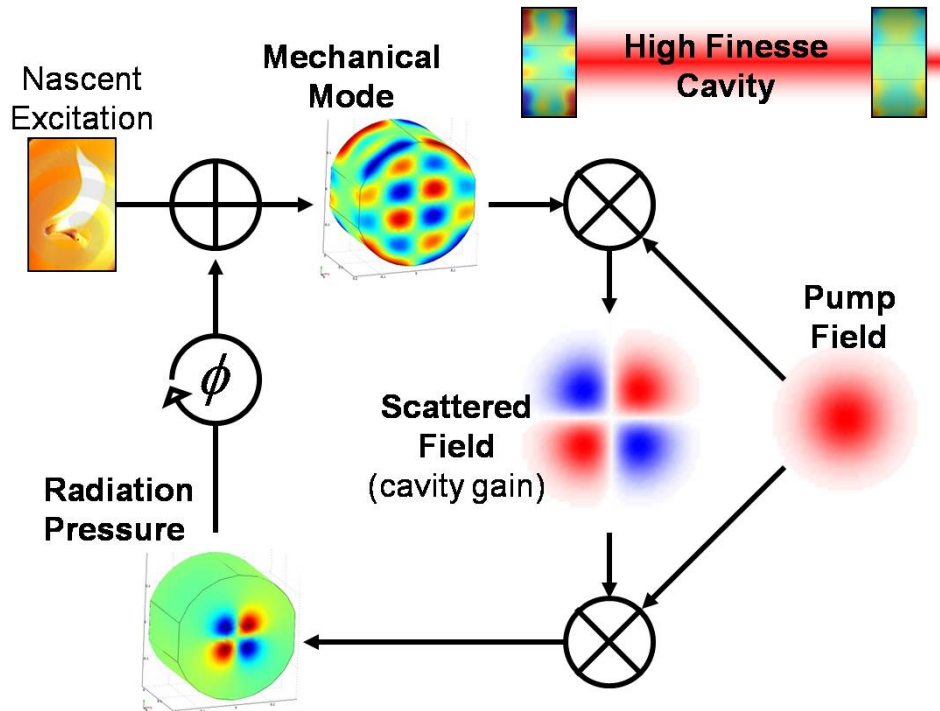


Figure 4: Diagram of the parametric instabilities possible in Advanced LIGO detectors.




Figure 5: A simulation showing that all modes can be damped with a single quadrant of the ESD using less than 1% of its capacity.

5.2 Charging Measurements

Approximations of the test mass charge were made by measuring the cavity response to ESD excitations. With no bias applied the ESD excitation should have no effect on the test mass motion. However, a charge on the test mass will act like an effective bias. See the work done by GEO in item 3 of the Applicable Documents section. Thus, the charge can be estimated by measuring the zero crossing of the test mass response to varying bias voltages, i.e. observing what bias offset is needed to cancel the effective bias produced by the charge. Figure 6 plots the response of the test mass, measured with the quad triple cavity, when an 11Hz drive is applied to the ESD using varying amounts of bias voltage. That particular measurement showed a bias offset of 11V is needed to cancel the effect of the charge. After a following vent of LASTI a repeat test measured an offset of 28.7V.

Since the response of the test mass is a function of not just the size of the charge but also its distribution, it is difficult to infer the size of the charge from this measurement. An FEA model estimates that the 28.7V offset corresponds to -3nC if it is assumed uniformly distributed on the surface. However, we know that the charge is not uniform because excitations of different quadrants of the ESD yield different offsets. Assuming the charge is localized to a 2mm^2 square patch yields anything from -0.3nC to 20nC depending on where it is on the surface.

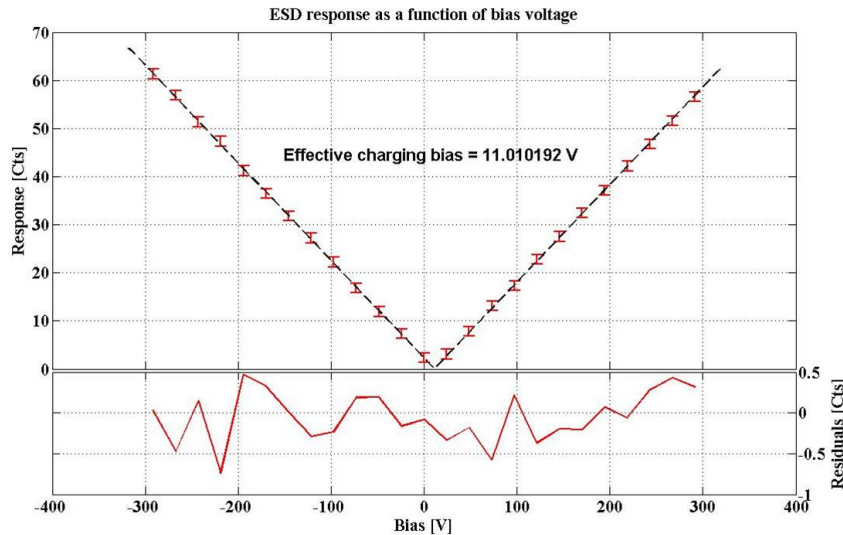


Figure 6: Response of the cavity to ESD drives at 11Hz using a range of bias voltages. This measurement shows a bias offset of 11V is needed to cancel the effect of the charge.

6 Mechanical Impact of Electrical Cabling and Clamping Issues

See T0900380-v1 'Quad Noise Prototype Cable Routing' for more information.

The quad noise prototype work proved the need for all around better clamping of the OSEM and ESD cables on the reaction chain. It is recommended that the production quads be assembled with optimized clamps with the unique needs of each clamping location carefully considered. The available space and geometry between the masses vary from point to point on the suspension. Additionally, all clamps should have protective padding to prevent damage to the delicate cables, particularly those of the ESD. Additionally, each clamp should take into account the vastly different diameters of the ESD and OSEM cables. RAL is currently taking on the task of redesigning the cable clamps with input from LASTI.

On the quad noise prototype it was possible to measure the mechanical effect of the OSEM D-connectors between the top and UI masses. The left photograph in Figure 7 shows the cables in this location. Figure 8 shows transfer function measurements of the response of the reaction chain to the cable connectors in this location. The OSEM cables and ESD cables have connectors in this location so that the lower and upper structures can be separated without impediment from the electrical cabling. The effect of cable stiffness and dangling connectors must be considered when clamping the cables in this location. A cable's stiffness is proportional to the modulus of its materials and cross sectional area, but much more than inversely proportional to its length. Additionally, a cable connector is effectively rigid compared to the floppiness of the cable itself. Thus, considering the relatively short gap between the top and UI masses, it makes sense not to consume more of this gap with bulky connectors.

When the relatively large OSEM D-connectors were left to dangle between the top and UI masses there was effectively a very short length of cable slack separating the two stages. The added stiffness caused pitch and z modes to couple to each other and also raised the frequency of at least one of the z modes. This extra stiffness was minimized by clamping the D-connectors directly at either stage. Measurements of transfer functions indicated that the effect of cable stiffness disappeared with clamped connectors. The word 'disappeared' in this case means that the response of the reaction chain was identical to the response when these connectors were completely disconnected, meaning zero cable stiffness. The ESD cable connectors were left to dangle, as shown in Figure 7, simply because there was no room to clamp them next to the OSEM D-connectors. Fortunately, the ESD cables and connectors added no measurable stiffness to the transfer functions taken. However, the performance of the reaction chain can only be improved by clamping the ESD connectors directly at the masses as well.

If the decision to use thicker ESD cables in the production quads goes through, clamping these ESD cables will become that much more important. It should also be noted that the production OSEM cables will be stiffer as well due to the added shielding.

Stiffer cables do pose the risk of creating issues in other areas. The gap between the UI and penultimate masses is also rather short and, although this gap posed no problem at LASTI, it may simply be insufficient for the bulkier cables planned for the production quads. Preliminary results from the LLO training quad at the October 5-9, 2009 training session showed that this location in addition to the gap between the top and UI masses may contribute cable stiffness to the suspension when OSEM cables with multiple layers of shielding are used. It should be noted however that there was insufficient time to fully study the effect of routing these cables during the training session. If the shielding on the production cables is at least reduced in certain places, such as in the gaps between masses, then the issue will be reduced.




Figure 7: The picture on the left shows the reaction chain OSEM and ESD cables and connectors between the top mass and the UI mass. The picture on the right shows the ESD and OSEM cables entering the suspension from the optical table at the reaction chain top mass.

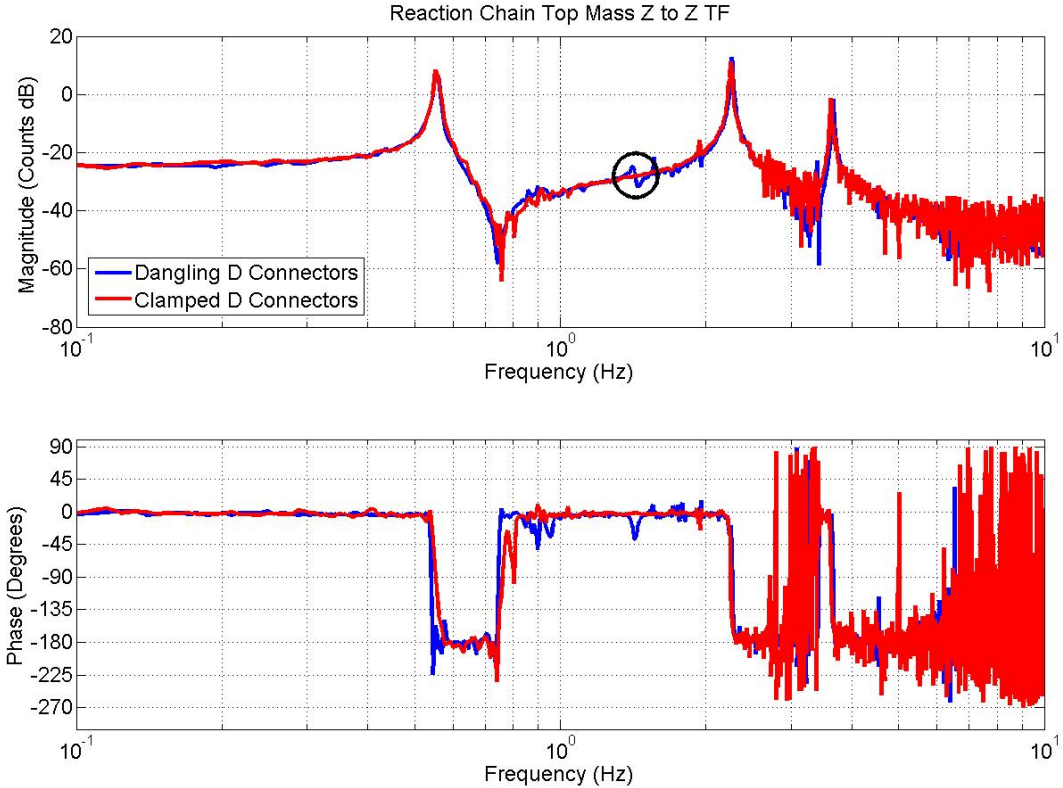


Figure 8: The OSEM cable connectors between the top and UI masses have a measurable impact on the TFs of the reaction chain when they are left to dangle in the gap between the two masses. When they are clamped firmly to the UI and top masses, as shown in the left photograph of Figure 7, their impact on the TFs goes away, i.e. the measurement is identical to the case when the connectors are unplugged. The blue curve shows a top mass to top mass vertical bounce (z) TF when the OSEM connectors are dangling. The red curve shows the same TF when they are clamped. Dangling connectors increase coupling to pitch, as the circled mode indicates. Less obvious is a slight increase in the 3.5 Hz mode frequency in the blue curve.

Other issues with the cable clamps on the noise prototype were related to a lack of padding and a lack of clamps. All of the clamps compressed aluminum directly on the cables. To add protection to the cables, Viton pads were placed in all the clamps.

Generally all the masses on the quad have a sufficient number of clamps, with the exception of the penultimate mass. The only available clamp is located at the bottom of the mass. A second at the top should be sufficient. An improvised piece of Viton was implemented to act as a guide to keep the cables together. See Figure 9.




Figure 9: The picture shows the improvised cable 'clamp' of washers and Viton used to support the cables at the top of the penultimate mass. The production quads should have a real clamp in this location.




Figure 10: The ESD and OSEM cables clamped to the ISI table just before entering the quad upper structure. The clamp here is also an improvised piece of Viton.

7 Model Fitting

Effort was put into fitting the parameters of Mark Barton's Matlab state space model to the main chain dynamics. It is believed an efficient technique is identified for doing this. Because the complexity of fitting a model to data relates very strongly to the dimension of the parameter space, the approach was chosen to reduce the number of parameters as much as possible. As Figure 11 illustrates, a complete set of transfer function measurements is taken with the top mass OSEMs when only the top mass is free. The parameters relating to the top mass are thus isolated from the rest of the suspension. The uncertain parameters of this stage, namely the rotational inertias, and the blade spring dependent d parameters (height of wire attachments relative to center of mass) are run through a Matlab fitting code with these measurements. When the code outputs the matched parameter values, the UI stage is then released. In this state all the parameter uncertainty is localized to the UI mass and another set of transfer functions is measured. The top mass is used again for the measurements since it alone has a full complement of OSEMs. Further, the quad is fully controllable and observable from the top mass in any partially locked configuration (provided the top mass is free of course). The UI mass parameters and the corresponding set of measurements are sent through the fitting code which then outputs the appropriately modified UI mass parameters. This process is continued all the way down to the test mass to fit the entire suspension.

During this process care must be taken to lock the masses down as close to their suspended position as possible. The primary concern is that shifted blade spring tip heights will alter the d 's enough to change the suspension response. Less of a concern, but still an issue is the pitch balance. Large pitch tilts of the top mass artificially introduce extra coupling between measurements that should not exist.

The fit code itself was written in Matlab. It is primarily a simple gradient descent approach. A parameter is moved a small amount in one direction. The error is computed and compared to that of the previous state. If the error decreases the parameter is kept at the new position. If the error increases the parameter is moved in the reverse direction. This simple algorithm is repeated many times over all the parameters in the fit until a local error minimum is found. The code is written to detect when it has fallen into a local minimum. When a detection is made, the code engages a Monte Carlo algorithm where it shifts all the parameters randomly within the uncertainty bounds to see if there are any other points in the space below this local minimum. If one is detected it starts the gradient descent again from the new point. If none are detected it calls the local minimum global and outputs the new parameter values.

The error is calculated by summing the magnitude of the difference between each point in the measured and modeled transfer functions from about 0.1 Hz to 10 Hz, at some chosen resolution. Since the shape of the transfer functions encompass all the relevant dynamic information, i.e. through the resonances, zeros, and gain, the result is more precise and robust than calculating the difference in resonance frequencies alone. Since y , pitch, and roll are highly coupled they were run through the fit code simultaneously. The others, x , z , and yaw are sufficiently decoupled to fit independently.

In general the uncertain parameters are the variable parameters in the quad, i.e. the rotational inertias effected by adjustable mass and the d 's set by the blade spring tip heights.

Unfortunately, like the quad controls prototype, there is additional uncertainty on otherwise static parameters due to the enormous book keeping task during the prototyping of something as complex as a quad. Thus, the precise values of some wire lengths and wire spacings (the model 'n' values) are still in question. Consequently, there remains some error in the model even after this parameter fitting process because it was unknown at the time the number and extent of the uncertain parameters. However, the error is small enough to permit modal damping, which is a control technique that imposes rather strict limits on error. RAL is currently assisting the process of tracking down the as built values of these static parameters for the noise prototype.

Overall the x , z , yaw, and pitch TFs ended with basically zero error. The roll and y TFs ended with error of a few percent on a resonance or zero because of the uncertainty in some of these as built parameters that were not considered in the fit code. See the results of pitch and roll in Figures 12 and 13 below.

It is expected that the production quads will not have the same issues with parameter uncertainty as the prototypes because the as built values will have long since been finalized.




Figure 11: The model fitting approach conceptually broke the quad into smaller sections. The top mass was first left free on its own and transfer functions were measured. The model parameters relating to the top mass were then fit with the Matlab script. Transfer functions were then measured when both the top and UI masses were free, and the fit code was repeated for the UI mass parameters. This process was repeated until the entire quad was suspended.

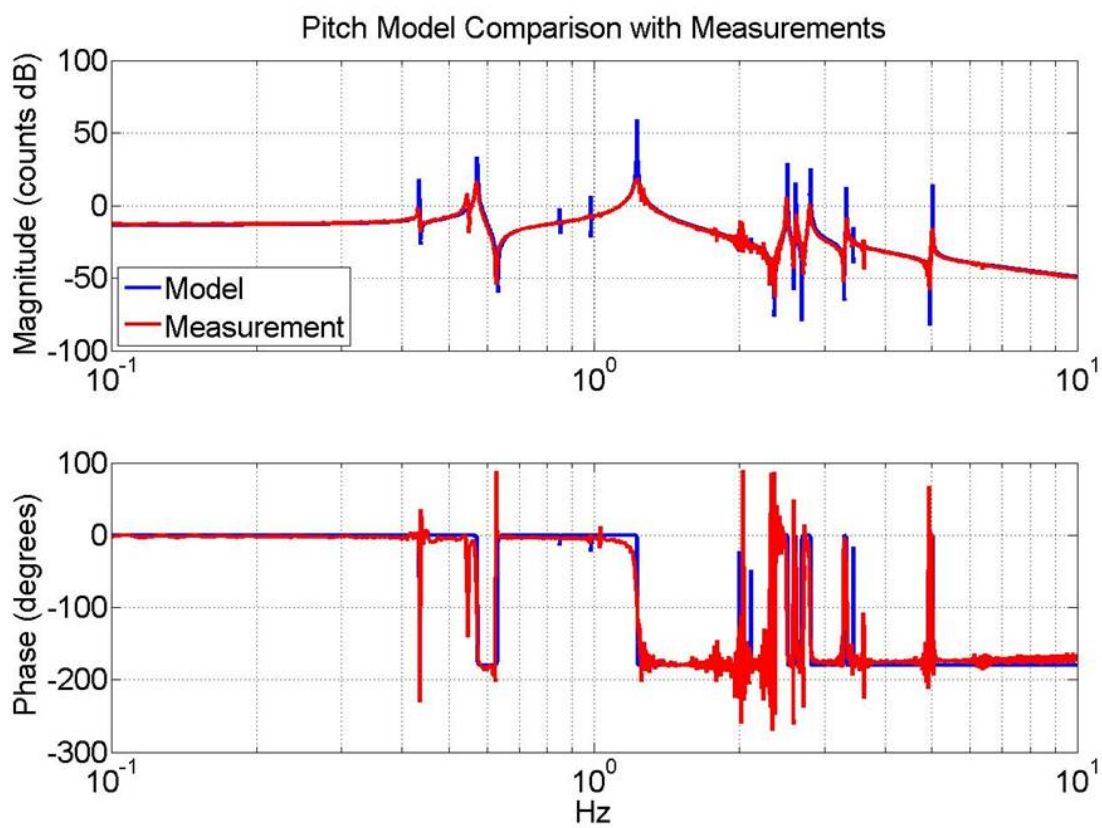


Figure 12: A comparison of the model to a measured top mass to top mass transfer function of pitch after all the stages of the pendulum were fit. In this case pitch is nearly a perfect match.

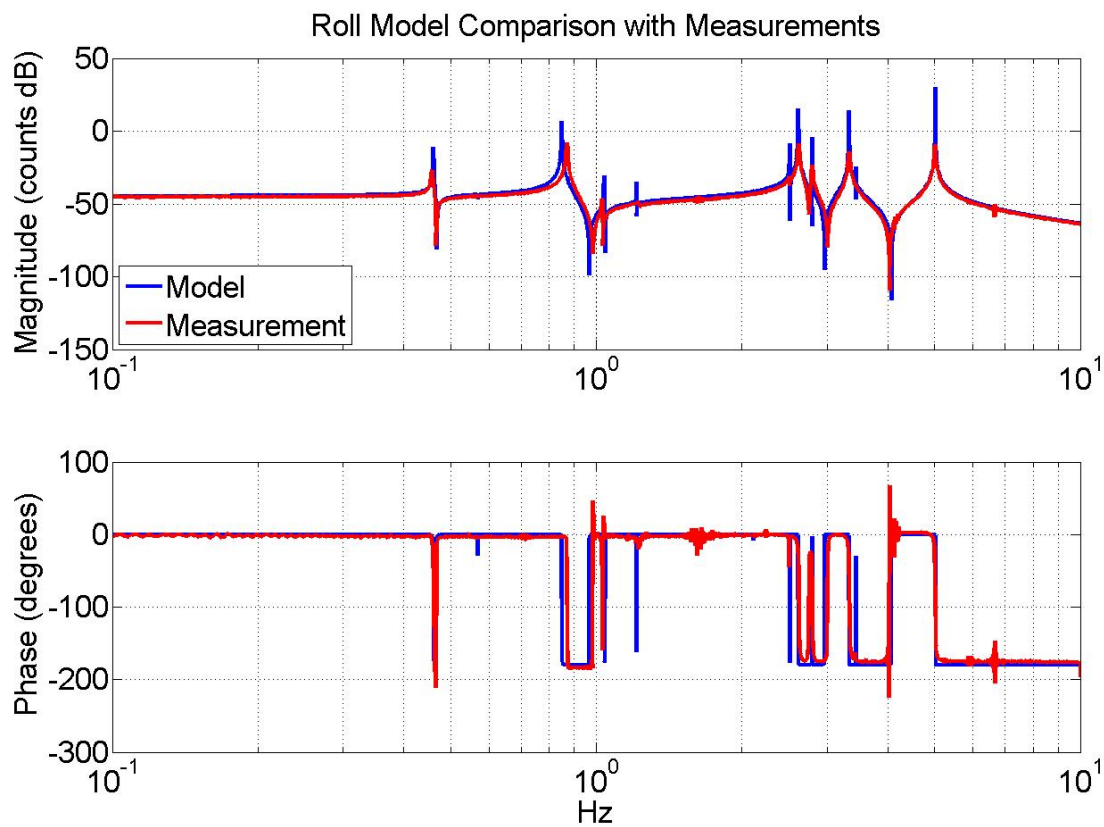


Figure 13: A comparison of the model to a measured top mass to top mass transfer function of roll after all the stages of the pendulum were fit. In this case the uncertainty in certain as built parameters prevented the error from converging completely to zero as indicated by the mismatch in the 0.85Hz mode.

8 Modal Damping

A detailed description of the modal damping technique is found in Laurent Ruet's thesis 'Active Control and Sensor Noise Filtering Duality Application to Advanced LIGO Suspensions' (P070054-00-Z), or Brett Shapiro's masters thesis 'Modal Control with State Estimation for Advanced LIGO Quadruple Suspensions' (P070144-00-Z).

Full modal damping was successfully implemented on the quad main chain. Modal damping involves transforming the measured x , y , z , yaw, pitch, and roll coordinates of the quad into modal coordinates where each modal degree of freedom conceptually represents an independent two degree of freedom (DOF) oscillator resonant at one of the quad's resonance frequencies. Thus, instead of designing optimal control loops for the complicated quad transfer functions, simple loops can be designed in the modal basis. In fact, since each modal response is identical in shape, only shifted in frequency and gain, the same design can be applied to all the modes of the quad. Thus, a few very complex control loops are replaced with many simple ones. The noise performance can then be optimized by tuning the damping gain of each mode individually. The caveat to modal damping is that the modal transformation requires all 24 DOFs of the quad to be measured. Since we have only the 6 top mass measurements from the local control OSEMs, an estimator is employed that predicts the remaining DOFs using a digital model of the quad running in parallel (hence the need for a good model).

Previously only modal damping was successful on the x , z , and yaw DOFs. Conceptually and computationally, these DOFs are much simpler because their dynamics are independent and they can be controlled individually as SISO systems. Note, that there is a some coupling between x and pitch, but it is negligible enough to ignore it. The pitch, roll, and y DOFs are highly coupled and must be considered as a single MIMO system. Further, the digital model of this MIMO system is more difficult to fit since it depends on more parameters with (conveniently) the most significant uncertainty.

Recently though, the combined efforts of computationally fitting the uncertain parameters and debugging previously unknown errors in the frontend software controlling the quad enabled the successful implementation of modal damping on the remaining coupled DOFs. Figure 14 shows a measured modally damped pitch transfer function compared with the expected simulation results.

These results are still preliminary however. No work has yet been done to optimize the design of the estimator or find appropriate gains. Additionally, a sensitivity analysis of the model error on the stability of the system will provide more insight on how good the model needs to be.

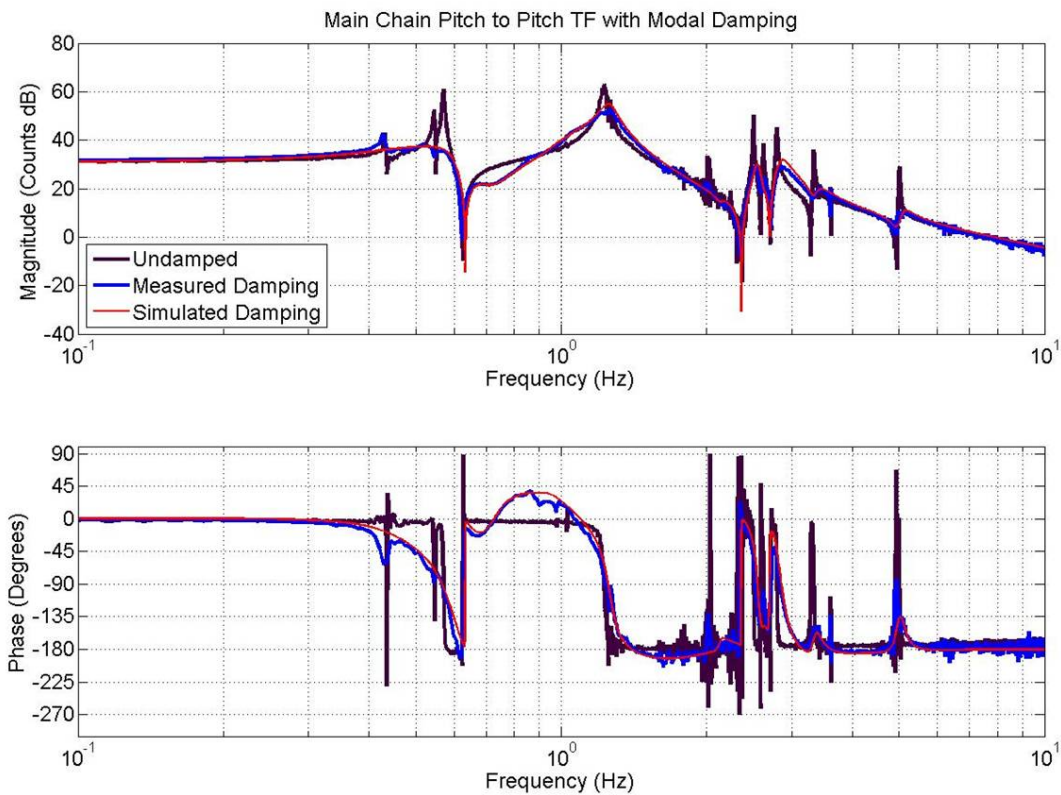


Figure 14: A top mass to top mass transfer function of pitch with modal damping running. The blue curve is the measured response, the red curve is the simulated response. The black curve is a reference showing the measured undamped response. The first visible mode at about 0.42 Hz is not predicted by the simulation because it represents an X mode. The modal damping estimator was simplified by removing all x, yaw, and z dynamics from the model.

9 Procedural Notes

A number of procedures have come out of the work on the quad noise prototype this year. The hanging of the glass optics on metal wires in January of 2009 resulted in T0900055-v1, 'Rehang of the Quad NP with glass masses but metal wires'. The document describes the steps taken at LASTI to replace the quad in the BSC with glass optics as a precursor to the monolithic installation. It also includes the ESD wiring and soldering procedure developed from the extensive work at LASTI. T0900518, 'Quad electronic and optical aligning procedure', is an extension to T080165-00 'Metal Quad Noise Prototype Balancing and Alignment Procedure' and includes steps to go from a roughly balanced quad to the fine tuned configuration.

10 Ongoing work

More testing with the ring heater is expected to examine the thermal and mechanical responses of the systems inside the BSC when full power is applied. Different global control techniques for maintaining lock in a cavity will also be tried including the standard hierarchical methods as well as various new ones. Eventually the seismic isolation of the triple pendulum will be augmented with the installation of a HAM ISI next year. The monolithic upgrade of the quad is expected to begin as soon as the end of this year. This work is expected to last a couple months altogether. Consequently, much of the global control testing will wait until after the monolithic installation, and possibly until after the HAM ISI installation depending on the timing of LASTI's schedule.