

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

Engineering Note	LIGO-E1000078-v1	Date: 3/12/2010
Electronic Setup and Testing of Advanced LIGO Suspensions		
Brett Shapiro		

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

Contents

1	Introduction	1
1.1	Purpose and Scope	1
1.2	Related Documents and References	2
1.3	Version History	2
2	LIGO Controls Software	2
2.1	Signal Flow Loop	2
2.2	Overview of LIGO Controls Software	3
2.3	MEDM Screens	6
2.3.1	Quad MEDM Screens	7
2.4	Data Viewer	19
2.5	Diagnostic Transfer Function Measurements with DTT	21
2.6	Foton Basics	27
2.7	Brief notes on the Simulink-Frontend Interface	30
3	Transfer Functions	30
3.1	What are Transfer Functions?	30
3.2	Debugging Suspensions with Transfer Functions	32
3.2.1	Quad M0 Reference Transfer Functions	35
3.2.2	UIM Left Blade Spring Clamped	38
3.2.3	Top Mass Right Blade Spring Clamped	43
3.2.4	Top Stage Right Blade Spring Clamped	47
3.2.5	FACE1 OSEM Magnet Rubbing Inside OSEM Bore	51
3.2.6	Top Mass Forward Left Top Stop Clamped	56

1 Introduction

1.1 Purpose and Scope

This document is written as a quickstart guide to setting up the MEDM screens, using DTT, Data Viewer, and Foton; and debugging Advanced LIGO suspension using transfer functions. The upcoming commissioning of the suspensions at LHO and LLO is the motivation for this document.

1.2 Related Documents and References

E1000006: Advanced LIGO Quad Suspension Procedure
 T0900612: AdvLIGO CDS Design Overview
 T080135: AdvLIGO CDS Real-time Code Generator (RCG) Application Developer's User Guide
 G050190: A Brief Overview of LIGO and AdL Controls
 D060403: Quad Top Mass
 D060375: Quad Upper Intermediate Mass
 D080128: Quad ITM Penultimate Mass
 D080117: Quad ETM Penultimate Mass
 T080328: BSC Suspension Test Stand (US) Controls Test Plan
 T080062: UK Satellite Amplifier Pre-Production Prototype Test Plan
 T080014: UK Top Driver Pre-Production Prototype Test Plan
 T080021: UK UIM Driver Pre-Production Prototype Test Plan
 T0900164: UK PUM Driver Pre-Production Prototype Test Plan
 T0900140: AdL Triple Acquisition Driver Board Prototype Test Plan
 T0900190: AdL Triple Low Current Driver Board Prototype Test Plan

1.3 Version History

12 March 2010: Version 1 submitted.

2 LIGO Controls Software

2.1 Signal Flow Loop

The standard signal flow from measuring the pendulum's motion to applying forces to it creates a loop. This loop, when it is complete, is referred to as a feedback loop because the pendulum's motion is sent back to itself to make the pendulum behave in a more desirable way. This loop is illustrated in Figure 1. Note that the figure can be split two ways, top-bottom and left-right. The top and bottom halves are symmetric with each other, where the top half represents the measurement or 'input' side of the loop and the bottom the actuation or 'output' side of the loop. The left side incorporates all the physical analog components while the right has the digital components living inside computer software. All the control is implemented in the computer world on the right side of the loop. All the software discussed in this document interacts with the suspension through this part of the loop.

To understand what happens in this loop let us walk through it from beginning to end. First, the sensors measure the pendulum's motion and convert it into voltage signals. In the case of the OSEMs this happens when the flag blocks part of the LED light entering the photodiode. These signals are then conditioned with nearby electronics by going through some analog filters and amplifiers. They then enter the computer through the Analog-to-Digital-Converter (ADC) which samples the analog signals at some given sampling rate (probably 16384 samples/second, the standard LIGO sampling rate).

Inside the computer we could in theory implement all the control in one step, however it would be much more difficult to design. Instead the control implementation is broken into a series of simpler stages. Once inside the computer the signals first go through input filters which do nothing more than compensate for any relevant analog filtering done by the electronics. This step is important so that the signals we pass to the next stage are nothing more than scaled versions of the sensor signals. Otherwise there might be some undesired frequency dependence. Typically at this step we also compensate for any gain differences between the sensors themselves. For example, if one sensor is weaker than the others, we add a little more gain to its input filter to normalize it with the others. The next step in the loop is the input matrix transformation. This matrix is a way of informing the loop of where the sensors are located on the suspension. In practice this means we add and subtract the sensors to put the signals into an x-y-z coordinate system relative to the mass of the suspension being controlled. Thus, the output of the matrix is a measure of the x, y, z, yaw, pitch, and roll motion of the relevant mass. The signals then go through the controllers (e.g. damping filters). These controllers are nothing more than yet another stage of filtering that is specially designed to make the pendulum do what we want. The output of these controllers is the digital 'forces' we want to apply to the pendulum to control it, still in the x-y-z coordinates. To transform the 'forces' into signals appropriate to each actuator they must go through another matrix. Fortunately, since OSEMs are used, the actuators are in the same place as the sensors. As a result, this output matrix will look a lot like the input matrix (in fact it is the transpose). Before leaving the digital world the signals must pass through another stage of filtering that is exactly analogous to the input filters. They compensate for any relevant filtering the electronics do at the next step and normalize any differences that might exist in the actuators themselves. This step is where the OSEM magnet polarity is taken account of for example.

The signals exit the computer through the Digital-to-Analog-Converter (DAC). This box turns the computer signals back into a real-world voltage. They then go through one last stage of electronics that converts them into a drive useful to the actuators. For the OSEMs this means they go through the 'coil driver' that converts the voltage to the current that goes through the coils. The actuators finally convert the drive into forces and apply them to the pendulum.

2.2 Overview of LIGO Controls Software

The aLIGO suspensions make use of a variety of software, each with a unique role. All interact with the suspension on the digital side of the signal flow loop shown in Figure 1.

The majority of interaction with the suspensions is through MEDM (Motif Editor and Display Manager). MEDM is a graphical interface that incorporates the top level controls through buttons, sliders, meters, text displays/entries, and graphs. For example, this is where the damping loops are enabled/disabled, where the matrix transformations on the signals are entered, and where any gain values are set. When commissioning a new suspension it is likely that some new screens will need to be generated and others adjusted. Editing MEDM is straight forward and easy to learn.

While MEDM does display real-time values of signals it is not so conducive to viewing real-

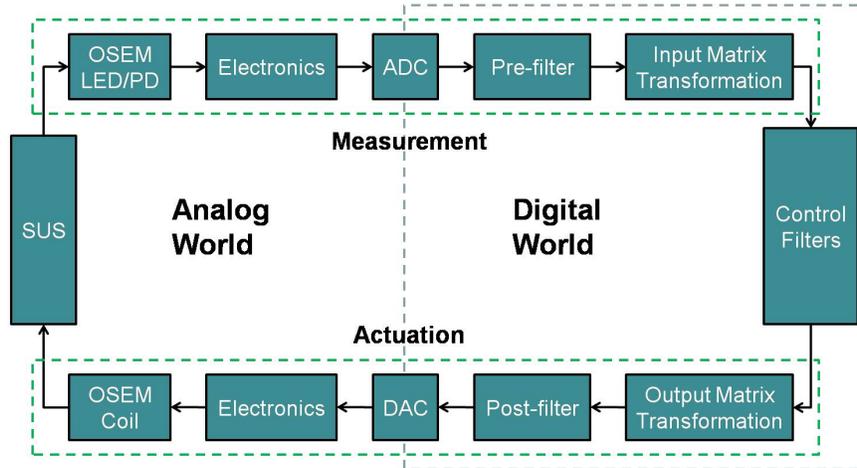


Figure 1: The signal flow loop from the sensor signals to the actuator outputs. The top half represents the measurement side of the flow, the bottom the actuation side. Note that these two parts are symmetric with each other. The figure can also be split left to right, where the left side is the analog ('real') world, and the right is the digital ('virtual') world where the control is implemented.

time plots of data, such as what you might expect from an oscilloscope display for example. For this we have Data Viewer. Data Viewer is fairly self explanatory. It allows you to pull up live plots of any signal, up to 16 simultaneously and with a time scale anywhere from $\frac{1}{16}$ of a second to 16 seconds. It also allows you to look back at the history of signals that are recorded to the frame builder. Data Viewer is also the most passive of these software packages in that it only allows you to view data. It has no ability to influence what forces are applied to the suspension.

The frame builder is a computer that is like a giant archive, amongst other things. For chosen signals it stores some form of their recorded history for all time. More specifically the full data (recorded at some sampling rate, probably 16 kHz or less) is kept on disk for a finite amount of time (days) dependent on the disk size, number of channels recorded, and the sampling rate of those channels. This data is also averaged over second, minute, 10 minute, and hour intervals. This averaged data takes up much less disk space and is kept for a much longer period, if not indefinitely.

DTT (Diagnostic Test Tools) is the general purpose testing software. It allows you to measure transfer functions, amplitude spectrums (generally labeled as power spectrums), and time responses. These tests are all useful for checking the behavior of the suspension.

Foton is a tool used to install control filters (including the damping filters). In control theory terminology it allows one to enter the poles, zeros, and gains of the filters desired. It also includes options for standard filters such as elliptical, notch, bump, butterworth, chebyshev, etc. Once loaded these filters are accessed through the medm screens.

Behind the scenes of all this is the software running on the frontend computer. The frontend is the master computer that runs the suspension. It usually has a convenient name such

as 'sus', 'quad', etc. This software is what defines all the filter modules, channel names, signal flow, etc. Anything that appears on the medm screens is based on something in the frontend software. It also has a convenient interface in Matlab's Simulink environment so editing is fairly straight forward and does not require a trained specialist (for the most part). For example, if you want to reroute a signal, add/remove signals, add/remove filters, this is where to do it. To put it simply, the frontend software runs everything; MEDM, Data Viewer, DTT, and Foton are merely pretty interfaces to interact with it.

2.3 MEDM Screens

MEDM is opened either by clicking on the appropriate icon, or typing `medm` into a terminal window. MEDM can be opened in edit mode either by clicking the edit button in the menu window or by typing `'medm'` (in single quotes) in the terminal window.

Each suspension has a top level overview screen that summarizes everything on the digital half of Figure 1. It has links to more detailed screens for all the filters, matrices, controls, watchdogs, test points, errors, and any other relevant information. It is from this screen that one implements the suspension control and closes the feedback loops. All damping control of the suspension resonances is done from the top mass, while the optical cavity control is split between the stages below.

This section will discuss the MEDM screens following the example of the quad, however all aLIGO suspensions are similar and the discussion applies to all.

All MEDM screens are built around groups of filters and matrices. Every filter has a small screen associated with it as shown in Figure 2. Usually these filter module screens are pasted into a larger screen with a set of common filters. For each filter the input signal flows in from the left and the output goes out through the right. The input signal first passes an ON/OFF switch. The green light below the switch indicates that it is ON and that the signal is allowed to pass through. Red would mean it is off and the signal does not pass. Next it passes a summation point where a user defined excitation signal is added to the input signal. This summed signal then passes another summation point where an offset can be added. The signal then goes through a series of filter sections. There are a total of ten possible sections (5 are above the line and 5 are below). The signal is only filtered by those where the green light shows that they are enabled. Each section has two indicator lights. One light indicates that the user requested the filter to be enabled, the second indicates that it is in fact enabled. For example, if you try to turn on a filter section that has no filter loaded into it, only one light will turn green. In Figure 2, filter sections 2 and 7 are enabled. Sections 1 through 7 have filters loaded into them while 8, 9, and 10 do not. After the filter sections the signal goes through the proportional gain in the black triangle. This gain multiplies the filtered signal by the value indicated. After the gain is a LIMIT switch (aka saturation). When enabled the the LIMIT will not permit any values greater than what is typed in the box to pass. If a signal has a greater value, the value in the LIMIT box will pass instead. The signal, now known as the filter output, will go through another ON/OFF switch. If the switch is enabled the signal will pass onto to the next stage in the digital loop.

Other features in this filter module include the 'test points' indicated by the three white arrows IN1, IN2, and OUT. These are points where the signal can be examined in Data Viewer or DTT. The CLEAR HISTORY button tells the filter sections to 'forget' the history of the input signal. Since the output of a filter depends on the input values that came before in time, the effect of pushing this button can be dramatic for certain types of filters (e.g. integrators). The LOAD COEFFICIENTS button updates the filter sections with any modifications that have been made with Foton. The HOLD OUTPUT button tells the filter module to hold the current value of the output. The 'Ramp Time' option will ramp any changes in the gain over the time period entered into the box, in seconds. Sometimes it is better not to change gains rapidly if the feedback loop is closed. So this box is an alternative

to manually stepping the gain many times.



Figure 2: This is the standard MEDM screen automatically generated for each filter module in the digital loop. Many MEDM screens contain sets of these filters.

2.3.1 Quad MEDM Screens

The quad overview screen is shown below in Figure 3. This screen summarizes all the relevant information and it links to more detailed screens. It has four main areas: M0, main chain top mass; R0, reaction chain top mass; UI, upper-intermediate mass; and PEN, penultimate mass. These four areas represent the OSEMs and control of these stages. The bottom mass has no OSEMs, but there is a link to the electrostatic drive (ESD) at the bottom of the page. Each area graphically follows the input-control-output signal flow shown in Figure 1. The OSEM sensor signals flow in on the left, they go through matrices and control in the middle, and the actuator signals exit on the right. The individual OSEM boxes are also arranged in a configuration similar to how they are mounted to the quad masses, with the exception of the M0 and R0 SIDE OSEMs on the input side. However one could edit this mistake easily enough if desired. Note that the main chain and reaction chain have separate sets of control for the top masses while the stages below do not. This is because each top mass has its own sets of OSEMs while the OSEMs below are literally split between the main chain and reaction chain. On these lower stages the magnets and flags are located on the main chain while the LED, photodiode, and coil are on the reaction chain. Thus, the sensor signals represent the relative motion between the two masses and the control forces are equally applied to both.

The M0 and R0 sections of the screen also have a different appearance from the UI and PEN sections. This difference is in part historical and in part a reflection of the different types of control at each stage. The primary difference is that the top mass control uses only OSEMs while the UI and PEN sections allow the control to choose a combination of the OSEM sensors and the optical cavity signals. In practice the OSEM sensors at these stages are used only for alignment purposes and will probably never appear in the control. Another obvious difference between these sections, which is appearance only, is that the UI and PEN sections have their input and output matrices on the overview screen while in the M0 and R0 sections these matrices are represented by simple black rectangles. The M0 and R0 matrices are mixed in with the control and output filter screens.

At this point it should be said that although there are many wrong ways to setup these screens, there is no unique 'correct' way. Much of what is shown here results from convenient conventions that were developed with the quad prototypes. Some of these conventions reflect efforts to be consistent with the Matlab and Mathematica models of the quad (e.g. consistency with SI units, right handed coordinate systems, and the order of channel names). In the end it does not matter which conventions are adopted. However, it is best to choose a system and stick with it consistently to avoid much confusion later on. It is also beneficial if all aLIGO suspensions adopt the same conventions.

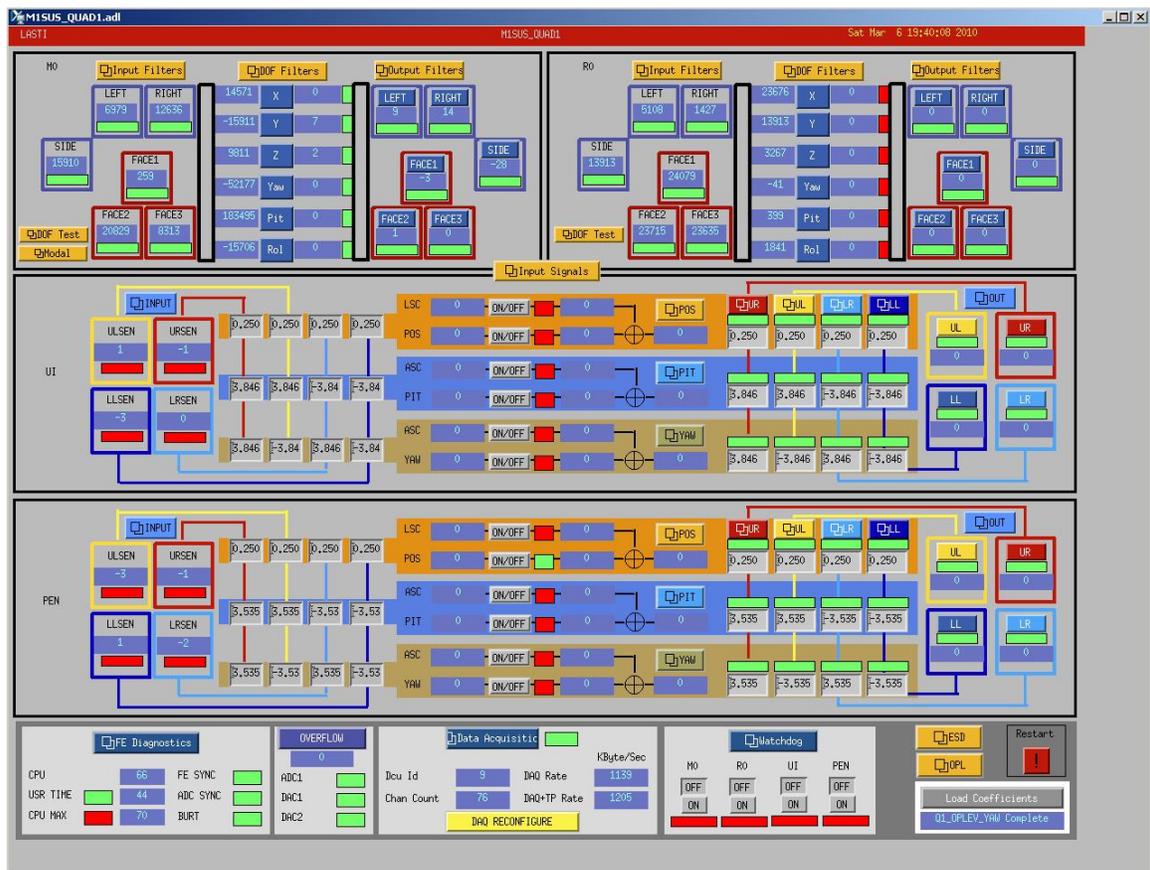


Figure 3: The LASTI quad MEDM overview screen. This is the top level screen and has links to most of the other screens.

The procedure for setting up the matrices and input and output screens is the same everywhere and we will illustrate the process with the M0 screens. The M0_INPUT screen represents the first stage in the digital signal flow loop. It is accessed by clicking on the yellow 'Input Filters' link in the M0 section of the overview screen. Figure 4 has a screen shot of the LASTI screen.

The raw OSEM sensors (post ADC) enter the screen on the left, go through some filtering in the middle, and exit on the right. The filtering component of the screen allows the signal to go through a series of zero to ten individual frequency-dependent filter sections before passing through the proportional gain in the black triangle (the proportional gain is technically part of the filtering). In this case there are no frequency dependent filters loaded or engaged because the previous stage of LASTI electronics has no frequency dependence to be canceled out. Note that the production electronics may NOT be the same. For example the version of this screen used by the LLO training quad does have filters engaged here.

The LASTI screen does use the gains in the black triangles. The reason they are all negative is in part historical and in part a user preference. They can in fact be whatever you wish provided subsequent stages in the loop are adjusted accordingly. The history is that the quad at LASTI originally had a set of electronics that gave us positive raw OSEM signals and the MEDM screens were setup accordingly. Then we received new electronics that were more similar to aLIGO production versions which put out negative raw OSEM signals. Instead of adjusting multiple screens to take account of the polarity shift, the OSEM signals were simply forced back to positive numbers using this M0_INPUT screen. Some people also prefer to think about positive numbers, so this negative gain convention stuck.

The size of the gain in these triangles takes into account the strength of the OSEM sensors. It turns out there is quite a bit a variability in the LED-PD pair between OSEMs. Birmingham has already pre-screened the OSEMs sent to the sites to give us ones that are close in gain, however removing all discrepancies is impractical. So since all the OSEM sensors put out a raw signal close to -30000 counts at full range (counts are the arbitrary units of the signal in the digital part of the loop), the gains here were chosen to normalize each OSEM to +30000 counts. The maximum value for each particular OSEM can be found by removing them completely from the suspension (simply retracting them gives you a similar answer but may differ slightly).

The output of these filters are those shown on the overview screen under the 'Input Filters' link. These values were by default the raw OSEM signals, but this was changed at LASTI to show the normalized positive signals. This change may or may not have propagated to the sites.

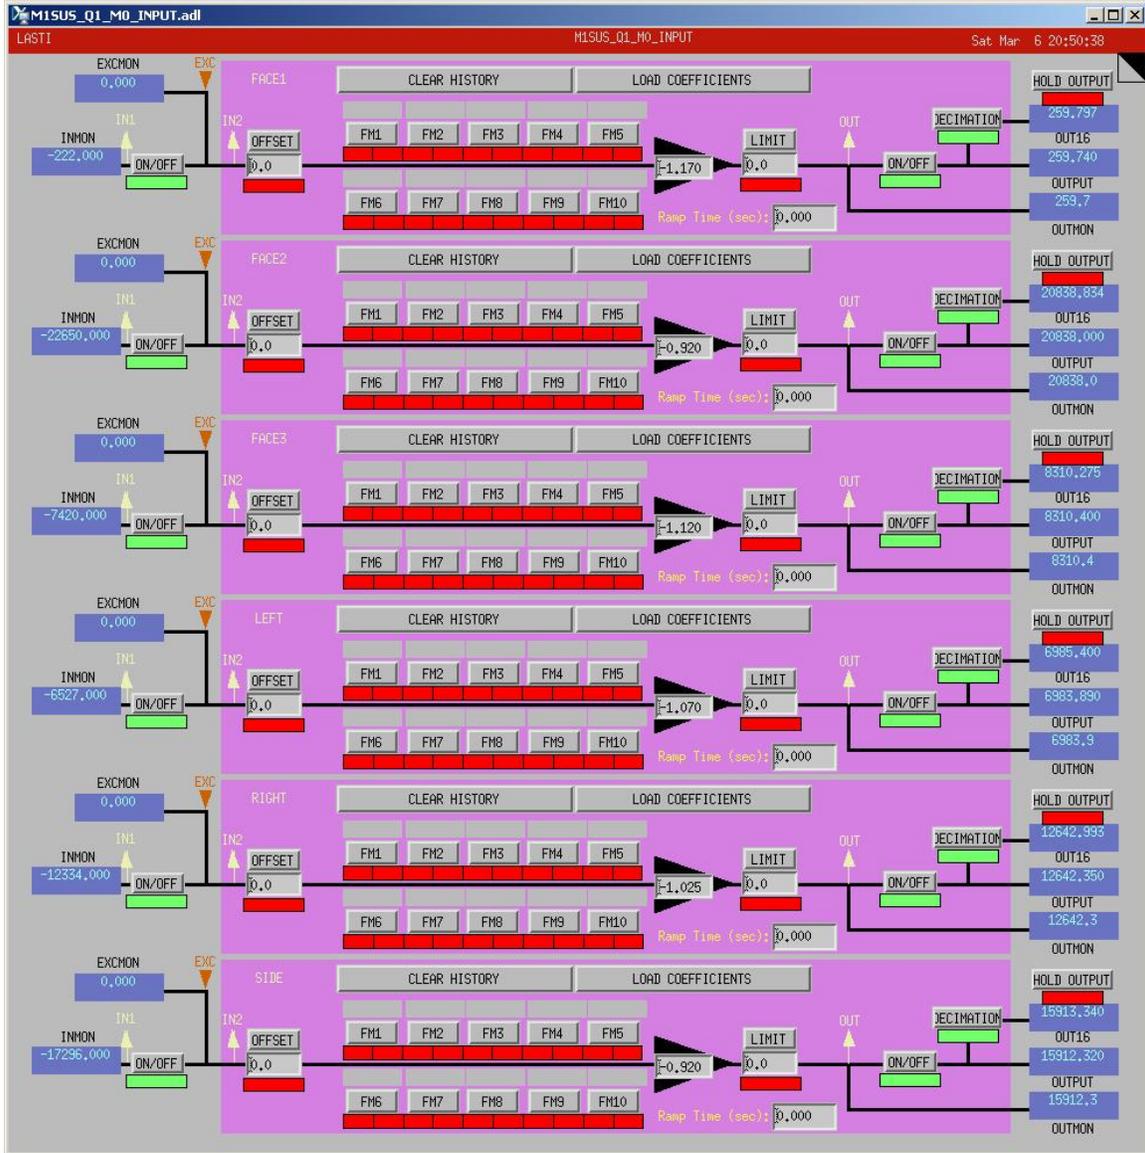


Figure 4: The M0.INPUT screen contains all the pre-filters in the digital signal flow loop. The raw OSEM sensor signals enter on the left, go through a filter in the middle, and leave on the right.

The next screen in the loop is the 'M0_DOF' screen accessed by clicking on the 'DOF Filters' link on the overview screen. Figure 5 has a screen shot. This screen contains both the input matrix and the main chain top mass damping filters. The signals that departed the 'M0_INPUT' screen first go through the matrix on the left side of the screen. It looks like there are actually six matrices, one for each filter, but these are actually subsets of a single larger matrix. The matrix does nothing more than add and subtract OSEM signals. Take the top row of this screen for example. It converts the OSEM signals into an 'X' signal. The matrix says that the X signal is the negative of half the FACE2 signal plus half the FACE3 signal, in other words the negative average of the two.

To understand where these matrix values come from one must know the geometry of the OSEMs around the main chain top mass. Figure 6 illustrates this geometry and the associated input matrix and coordinate system. The desired coordinate system shown on the top right of this figure is defined by the matrix underneath. The coordinate system was chosen in a way that made intuitive sense. The X axis is positive in the direction the mirror points, the Z axis is positive up, and the Y axis is positive to make the system standard right-handed. The rotational DOFs also follow the right-hand standard.

Conveniently, the quad was designed with the OSEMs placed in strategic locations that make use of its symmetry. The FACE2, FACE3, and SIDE OSEMs are placed at the height of the top mass's center of mass (COM). The LEFT and RIGHT OSEMs are placed at the front-back COM position. The FACE1 OSEM is aligned with the COM left-right, but offset vertically. Thus, to generate a positive X signal along the positive X axis we need to sum equal (because they are equidistant from the COM) negative values from FACE2 and FACE3. The reason for the negative is because when the top mass moves closer to the OSEMs, more LED light is blocked so the signal becomes weaker, or in other words the normalized OSEM signal becomes more negative. Since we desire this to generate an increase in the X signal we add minus signs. By convention we do not want the matrix to contribute any gain to the loop so half of each signal is used because there are two signals. Following this convention with Y, we see that SIDE is the only OSEM that directly measures this DOF and give the matrix a -1 in the Y-SIDE element. Z turns out to be the same as X except with the LEFT and RIGHT OSEMs.

The rotational DOFs are calculated by taking the difference of two signals (as opposed to summing them like we did for the translational DOFs above). If we want to maintain the convention of the matrix contributing no gain this is conceptually trickier with rotation. The tricky bit results from the fact that the numerical value of a rotation measurement depends on the lever arm over which it is taken (by analogy, instead of dividing by the number of constituent signals as we did for the translational DOFs, we divide by the distance between the signals). For example, given equal amounts of rotation, subtracting the LEFT and RIGHT OSEMs will produce larger values than subtracting FACE2 and FACE3 because they are farther apart. To compensate for this we divide by the distance between the OSEMs used in each signal, in units of meters. Consequently, a YAW signal is calculated by subtracting FACE2 from FACE3 and dividing the result by the distance between them, 0.24 m. ROLL is found by subtracting RIGHT from LEFT and dividing by 0.36 m. Pitch is more complex because three signals are involved. What we want to do is subtract motion above the COM from motion below. So we take the average of FACE2 and FACE3 (exactly the negative of X) and subtract FACE1 from this and then divide the result by 0.078 meters.

The damping filters are located in the center of the screen. The design of the filters requires some knowledge of control theory and takes into account the dynamics of that particular DOF. In general very simple and robust filters can be implemented just to get things going. For example, if all resonances fall below 10 Hz, a filter with an AC coupled zero and two poles at 20 Hz likely is stable with the appropriate gain value. Usually some type of higher frequency filtering, such as an elliptical filter with a 50 Hz cutoff, is needed to prevent excitation of higher order dynamics with these simple filters.

The negative gains ensure negative feedback. The negative sign here is extremely important. It more or less tells the controller that when the suspension moves one way it should push the other way to try to hold it still. Positive feedback does the reverse and will exponentially amplify the suspension's motion until it either bangs into stops or the actuators can put out no more force. Exceedingly large values of negative gain will do the same. Consequently, a control filter should first be implemented with small magnitude gains close to zero and then incrementally increased until the desired performance is reached. Adjusting these gains should always be done while watching the values of the signals, generally in Data Viewer.

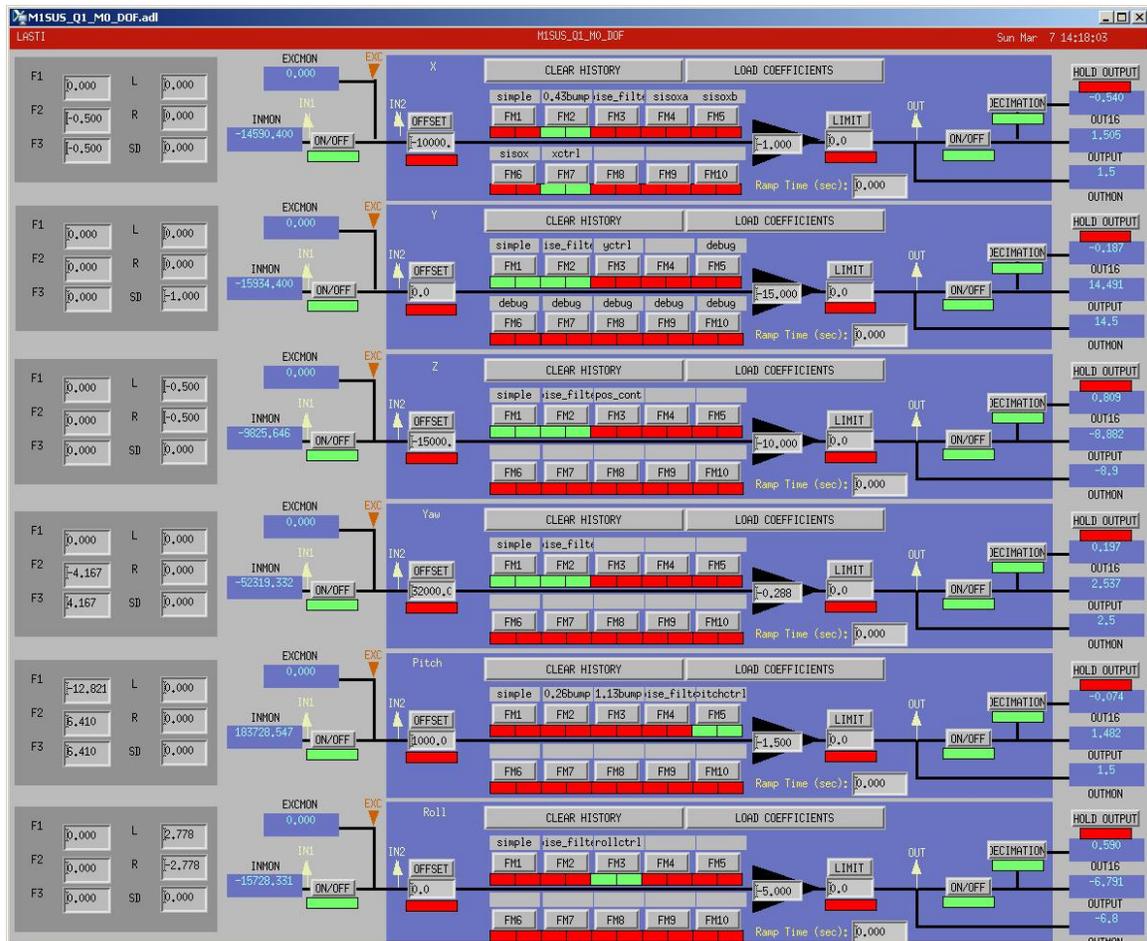


Figure 5: The M0.DOF screen contains both the input matrix and the control filters.

The final screen in the digital signal flow is the output screen. Figure 7 is a screen shot

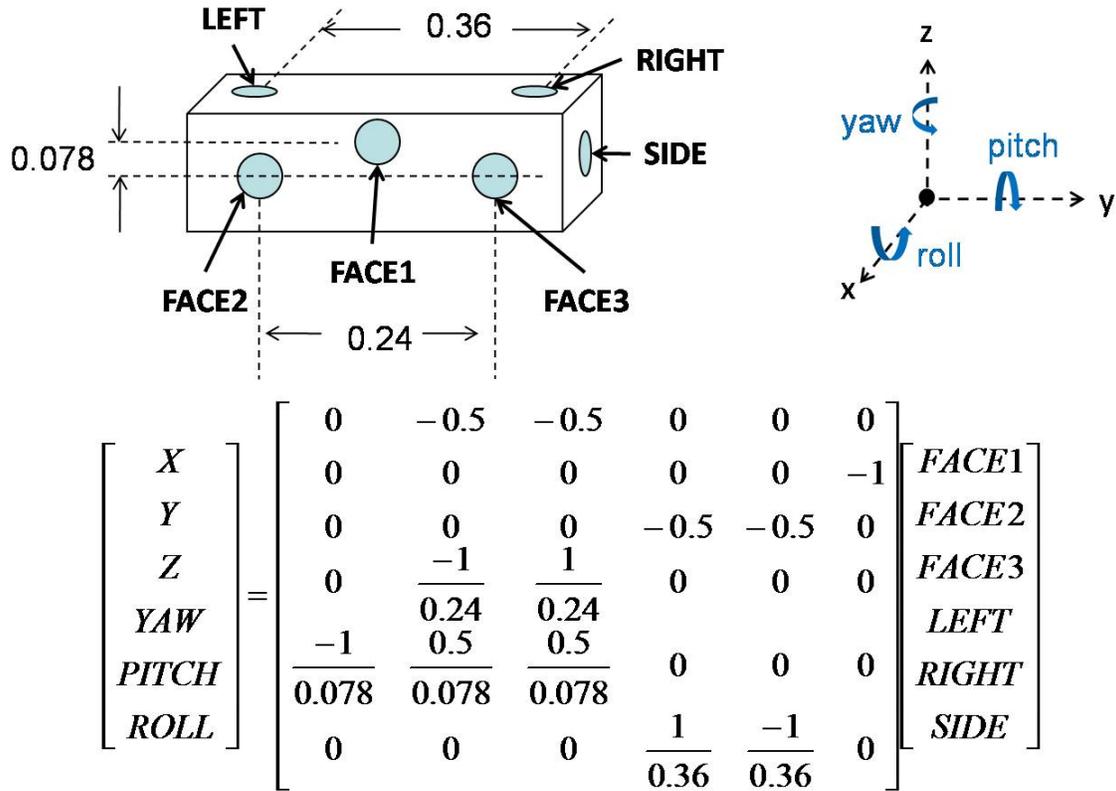


Figure 6: The geometry of the OSEMs around the top mass is shown at the top left. The relevant distances between these OSEMs is also shown in units of meters. The coordinate system defined by the input matrix at the bottom is shown at the top right.

of the LASTI M0_OUPUT screen. It has both the output matrix and the post-filters used to compensate for the magnet polarity and coil driver response. Since we use OSEMs the sensors and actuators are collated. This is good news for the output matrix because the information from the input matrix can be copied in exactly. In fact the output matrix is the transpose of the input matrix. For example, if the input matrix has a -0.5 between FACE2 and X, then the output matrix should have a -0.5 between X and FACE2.

Note that the coil driver electronics (at least for the quad) have different analog filters depending on what mode they are in. The modes are designed to correspond to the states of the interferometer so that you can drive harder, but noisier when the interferometer is out of science mode. See the appropriate electronics documentation to find out what the analog filters are.

To compensate for the magnet polarity you must check each OSEM drive. To do this you apply an offset to one of the output signals in the OFFSET box and observe which way the corresponding normalized sensor signal moves (at the output of the M0_INPUT screen.) The signal should increase for positive offsets. If it does not, switch the sign of the gain on the output screen. If the magnets were installed correctly the FACE2 and FACE3 gains should be opposite and the LEFT and RIGHT gains should be opposite.

The three M0 screens above complete the feedback loop for the main chain top mass. The other stages of the quad are setup the same way.

The R0 OSEMs locations were defined to appear the same as the M0 OSEMs when facing that chain. This means that when you are facing a particular chain LEFT and FACE2 are always on the left side of the mass, and RIGHT, FACE3, and SIDE are always on the right. Consequently, the M0_LEFT OSEM sits in front of R0_RIGHT. Because the layout of the OSEMs 'looks' the same, the input and output matrices can be copied right over from M0. In reality, the R0 coordinate frame will be flipped 180 degrees relative to M0, but this result has a negligible impact on the implementation. Experience at LASTI proved that this flipped coordinate system is the most intuitive while working on the quad because the left hand always deals with 'left' OSEMs and the right always deals with 'right' OSEMs, while all the matrices remain the same.

A side note about good book-keeping with OSEMs: It is VERY helpful to put the OSEMs on the quad in order by serial number. During quad work OSEMs may be repeatedly removed for any number of reasons, and replacing them is much easier if its obvious which serial number goes where. At LASTI this was done according to how they appear in the software. The order of the top mass OSEMs in the software is FACE1, FACE2, FACE3, LEFT, RIGHT, SIDE. So the serials numbers were placed in ascending order following this pattern. The R0 OSEMs were given the same order after the M0 OSEMs.

The LASTI quad has three additional screens related to the top masses. These are auxiliary screens because they are not part of the feedback loop. These screens are the M0 and R0 'DOF_TEST' screens and the 'Input Signals' screen. The M0_DOF_TEST screen is shown in Figure 8 and the Input_Signals screen is shown in Figure 9.

The DOF_TEST screens allow you to add excitations and offsets in the coordinate system of the top mass without going through the damping filters or damping gains. The screen has no inputs, only outputs. The output signals enter the loop between the control filter

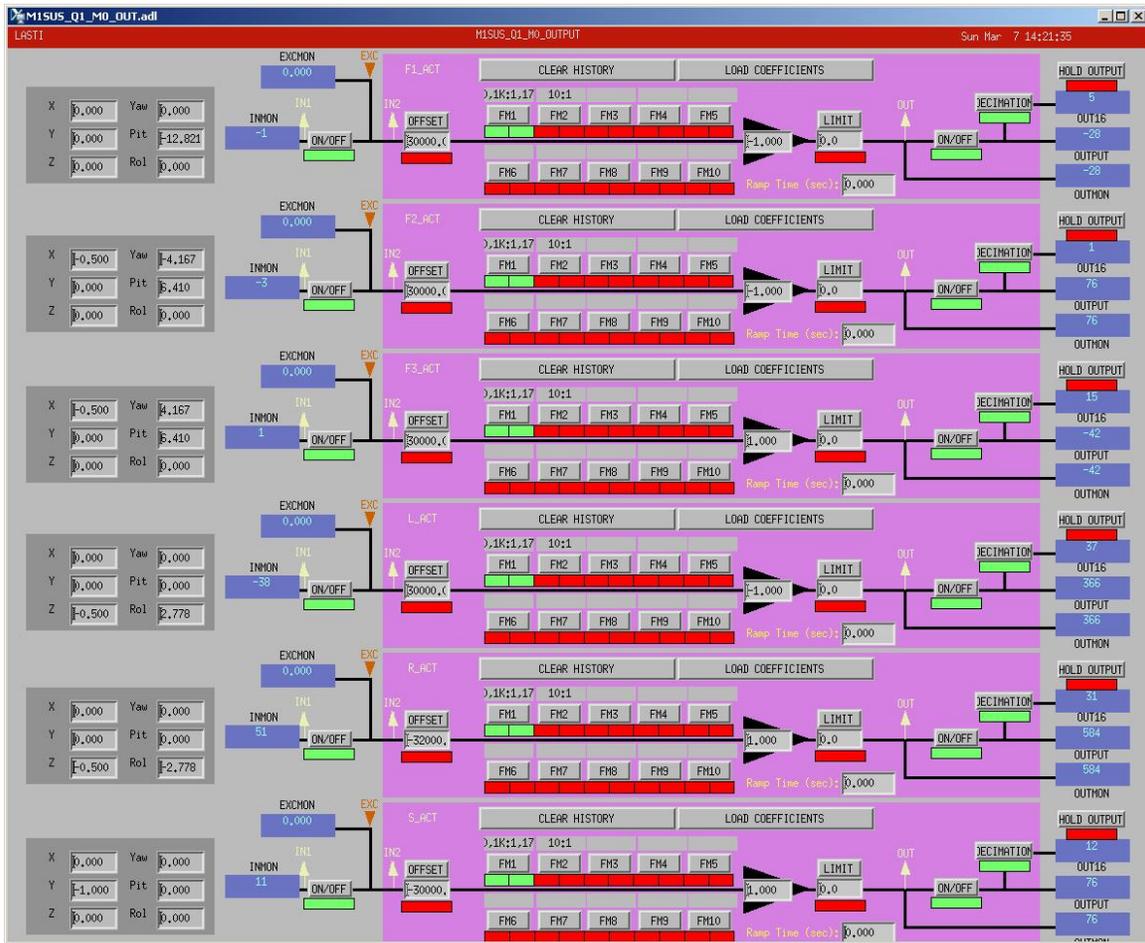


Figure 7: The M0_OUTPUT screen contains both the output matrix and the post-filters in the digital signal flow loop. This is also the screen where magnet polarity is taken into account.

outputs and the output matrix. The only other place to add excitations in the suspension's coordinate frame is in the control filter modules where these excitations are forced to go through the control filters and gains. So this screen becomes helpful if you do not want to worry about how these filters and gains effect your excitation. For example if you are aligning a cavity and you want to apply a PITCH offset, doing so here will bypass these filters and gains (in fact the damping filters will force any static offset to zero if applied on M0_DOF or R0_DOF screens). This screen is also useful for measuring transfer functions in the X-Y-Z coordinate frame. If all your DTT files are set to apply excitations here you will not need to generate separate files for the conditions of damping on or damping off, nor worry about what the damping gain happens to be set to at that time.

Since the DOF_TEST screens are LASTI built they may not have made it to the sites. Adding these screens in requires adding filters to the frontend Simulink diagram. See T080135 for instructions on how to edit the frontend code using Simulink.

The other auxiliary screen is the Input_Signals screen. This screen adds no new information or capabilities to MEDM. Consequently, no editing of the frontend code is required to generate this, only of MEDM itself. The screen is simply a convenient way of monitoring the normalized OSEM sensor signals by displaying them in a large meter format. The original motivation behind this screen was so that you could easily see the OSEM signals from inside the BSC, which is particularly useful when aligning them. The screen is also useful for quickly observing the qualitative behavior of the quad at any given time. When the quad is moving a lot it is much easier to see needles jumping around then numbers on a monitor. Also, since the ideal state of OSEM alignment is when the needles are pointing straight up and down, it is immediately obvious when something is off. Additionally, the screen has a copy of the watchdog for M0 and R0.

The final basic MEDM screen is the watchdog screen. The watchdog's job is to shut off all actuator drives if it detects a possible instability. It is sort of like a digital circuit breaker. The numbers adjacent to each OSEM represent an RMS velocity of the raw sensor signal. Thus, it is sensitive to the magnitude of a signal's oscillation, not the signal itself. Additionally, since the signal represents an average over some period of time there is a short delay between motion of the pendulum and changes in these numbers. The 'SetPoint' number is the limit at which the watchdog will shut off the drives for that particular set of loops. As soon as one RMS velocity signal goes above the SetPoint value, the whole group is shut off. Finding the optimal SetPoint is somewhat of a guess and check, but it is obviously safer for it to be smaller than larger. When the quad is placed in the BSC this watchdog should also be tied into both the ISI and HEPI (may not be by default, so check this) so that the quad watchdog can catch the case when large quad oscillations are caused by them. At LASTI the ISI is even forced off when the quad damping loops are deactivated.

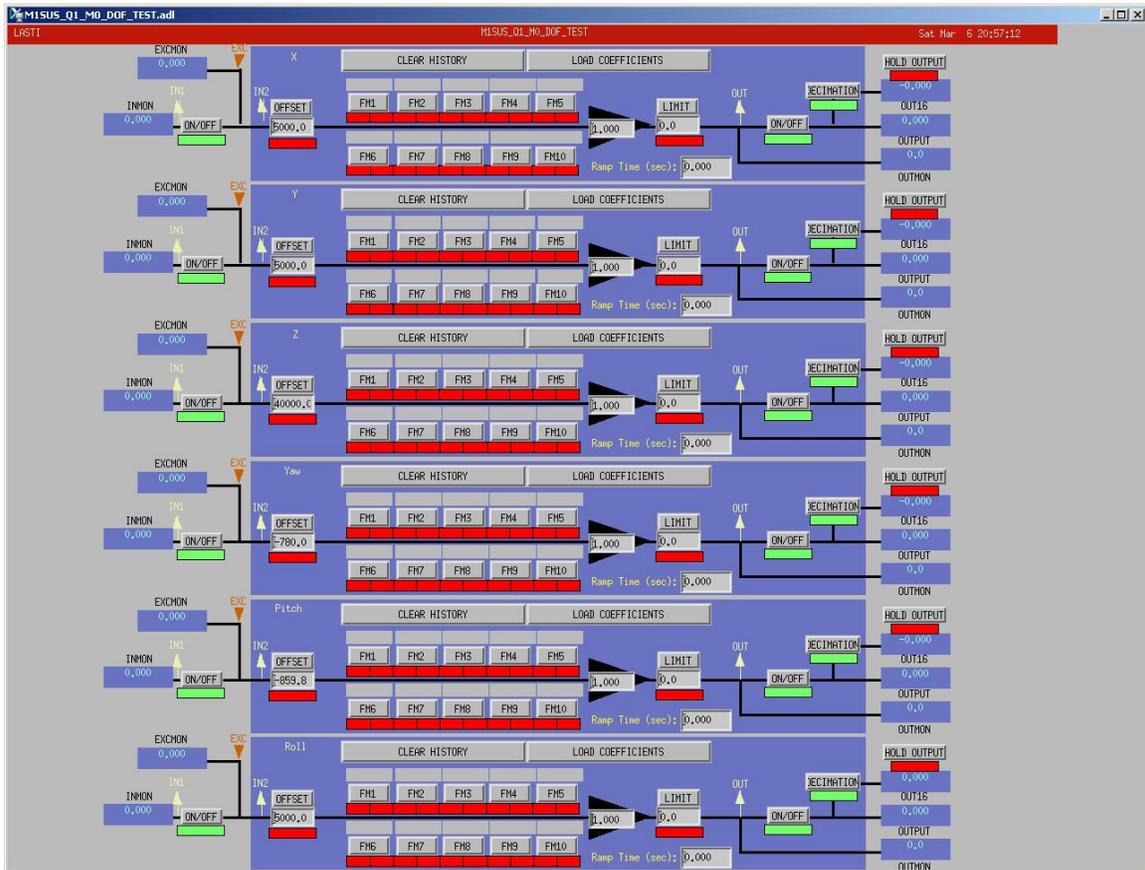


Figure 8: The M0_DOF_TEST screen is used to bypass the damping filters and gains while applying excitations and offsets in the main chain top mass x-y-z coordinate frame.

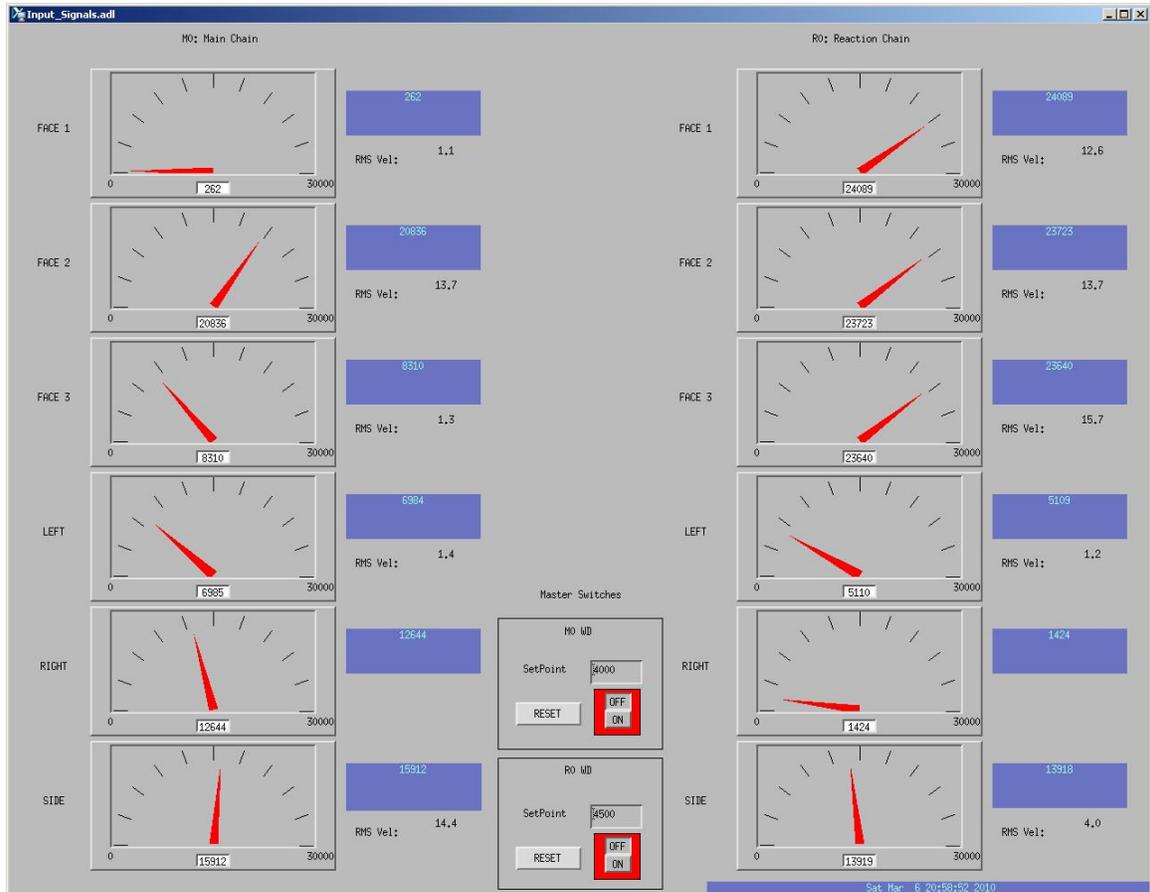


Figure 9: The Input_Signals screen is used to observe both top mass's normalized OSEM sensor signals in a more convenient format. Ideally, when the OSEMs are properly aligned all the needles will point straight up and down. In this screen shot they do not because the lower part of the LASTI quad is outside the BSC being prepared for the monolithic fiber welding.

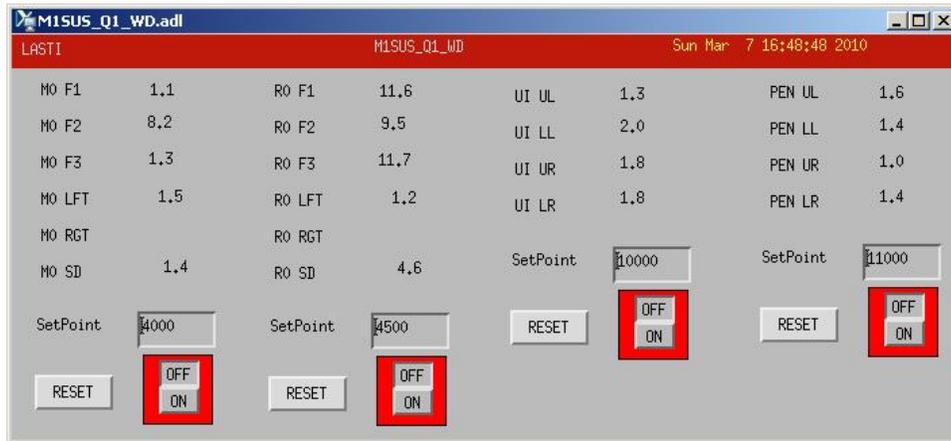


Figure 10: The Watchdog screen prevents the suspension from going unstable by cutting off the drives when the RMS velocity signals pass the SetPoint values. Here all the drives are deactivated because the lower part of the LASTI quad is outside the BSC being prepared for the monolithic fiber welding.

2.4 Data Viewer

Data Viewer is opened either by clicking on the appropriate icon or typing the word `dataviewer` into a terminal.

Data Viewer is a tool to plot and view real time data. Any signal in the MEDM screen can be plotted in real time with Data Viewer. These signals are selected in the 'Signals' tab in the main window. They are then plotted through the 'Realtime' tab. When commissioning a suspension, Data Viewer is often used to help select the sign of the the gains in the output screens that compensate for magnet polarity. As discussed in Section 2.3, the normalized OSEM sensor signal should move in the same direction as an offset in that OSEM's output filter.

During commissioning, the real-time plots are also useful when setting the gains of basic damping filters. One way to do this is by viewing the step responses. Step responses are tested by applying offsets to a DOF in the DOF_TEST screen. The signal of that DOF will then oscillate for some time. If a stable damping loop is closed, this oscillation will be less. Tuning the gain while viewing these oscillations will optimize the value of that gain. Damping ability can also be checked by measuring transfer functions with excitations from the DOF_TEST screen. The reduction in resonance peak height is proportional to the amount of damping for that resonance.

The Data Viewer 'Playback' tab allows one to plot data trends over some time period in the past. This can only be done with data recorded to the framebuilder.

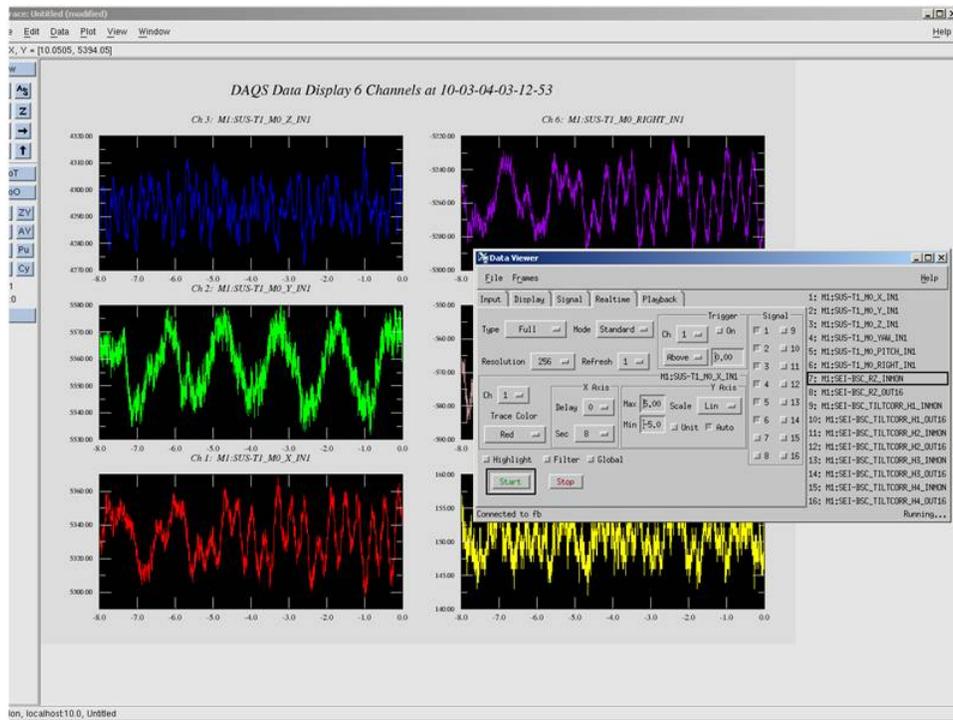


Figure 11: The Data Viewer screens allow you to plot real time data from any signal. It also allows you to plot previously recorded data history for any signals sent to the framebuilder computer.

2.5 Diagnostic Transfer Function Measurements with DTT

DTT is opened either by clicking on the appropriate icon or typing the word `diaggui` into a terminal.

DTT is the typical LIGO software used when running measurements on a system. Measurements can be either passive or active. Passive measurements, those without excitations, are generally power spectrums. Active measurements can also involve power spectrums, but also include transfer functions, and time responses (e.g. step responses amongst others). Transfer function tests are most extensively used to check the overall behavior during balancing, alignment, and commissioning. These checks, known as diagnostic transfer functions, generally involve comparing a recently measured transfer function to a previously measured reference that was known to have a good response.

Naturally, care needs to be taken when establishing these 'good' references. It is helpful to save these references in a convenient directory so they can easily be referred to for future diagnostic checks. A handy feature of DTT is that it also saves all the settings of a measurement. Measuring a diagnostic transfer function is as simple as opening the reference file, checking the data is copied into the 'Reference' field, and pushing the 'Run' button. The LASTI quad has 6 reference files for each suspension chain, one for each of the 6 DOFs in the X-Y-Z coordinate system defined by the matrices. Similarly, the LASTI triple has just 6 reference files since there is only 1 chain. Measurements within this coordinate system takes advantage of the symmetry built into the suspension.

Measurements are set in the 'Measurement' tab of DTT, see Figure 12. The type of measurement is chosen in the 'Measurement' box at the top of this screen. For diagnostic transfer functions there are two options that both have the same end result. 'Fourier Tools' is the choice for quick white noise excitations. These excitations allow one to measure all frequencies simultaneously. 'Swept Sine' does a swept sine test. It will measure one frequency at a time, which is generally more precise but takes longer. Additionally, the swept sine option of DTT does not handle large resonance peaks well, like those of the aLIGO suspensions, because the suspension will still be resonating when measuring nearby frequencies. Ideally, one would want to damp these resonances between each frequency measurement, but DTT does not allow this. My preferred choice is the 'Fourier Tools' option.

The measured signals are listed in the 'Measurement Channels' box. For diagnostic measurements I set 12 signals here, the 6 normalized top mass OSEM signals before the input matrix, and the 6 DOF signals after the input matrix. This way you can chose to check the coupling behavior and symmetry of the responses.

The parameters of the measurement are set in the next box down, also labeled 'Fourier Tools'. DTT will take the measurement in the frequency band from the 'Start' field to the 'Stop' field at the resolution of the 'BW' field. For any aLIGO suspension a good first guess of what to set is 0.1, 10, 0.01, in those fields respectively. This will tell DTT to measure the transfer function from 0.1 Hz to 10 Hz in steps of 0.01 Hz. Note that the smaller the 'BW' field is the longer the measurement will take because DTT must actually take a measurement down to this frequency to get that resolution. If the 'Start' field is above this value, it simply discards any data in between. The 'Averages' field tells DTT the number of times the measurement is repeated and averaged before stopping. The higher this number the better the quality of

the data but the longer it takes. A rough guess of how long the test will take is the number of averages divided by the number in BW. So a test with a BW of 0.01 and 10 averages will take about 1000 seconds. It is actually a bit longer because DTT does some fancier signal processing steps. I usually establish a reference file with 10 averages, and then use 3 averages for diagnostic checks. The other parameters may be left as the default values.

If one chooses to take measurements with the 'Swept Sine' option there will be some other choices to set here. First, the 'BW' box is replaced with the 'Points' box. This is the number of frequency values to be sampled between the start and stop frequencies. Second, is the number of cycles at each frequency set in the 'Measurement Time' field. DTT will measure each frequency with a sine wave of that number of cycles as many times as there are averages. So a rough guess of how long each frequency will take is the number of cycles multiplied by the number of averages divided by the frequency. Third, is whether DTT sweeps up from the start frequency or down from the stop frequency. Choosing to sweep down makes sense because the higher frequencies run faster and mistakes will more quickly be caught. The fourth option is the 'Sweep Type'. The choices are 'Linear' or 'Logarithmic'. This option tells DTT how to space the number of points specified in the 'Points' field within the band established by the 'Start' and 'Stop' fields. The 'Linear' option spaces them evenly, i.e. every 10 Hz segment will have the same number of points. The 'Logarithmic' option spaces them proportionally, i.e. every decade has the same number of points. Thus there will be the same number of points between 0.1 Hz and 1 Hz as there are between 1 Hz and 10 Hz. Since transfer functions are plotted in logarithmic space, it will actually appear on the graph that the points are evenly spaced with the 'Logarithmic' option. If 'Linear' is chosen it will appear that the density of points increases with frequency (transfer functions measured with noise inputs will have linearly spaced points). Finally, the 'Format' option is next. 'Format' is optional. It allows the user to define the amplitude of the excitation wave at each frequency. It will interpolate between the specified frequencies. This option is useful if you know there are tall resonance peaks, where you might want to reduce the drive, and deep zeros where you might want to increase it. If nothing is set in 'Format' then the default is the amplitude specified under the 'Excitation' Tab. The remaining parameters may be left as the default values.

The excitation channel is set in the 'Excitation' tab. I typically set this as one of the DOF_TEST channels to avoid going through damping filters, however the DOF channels can also be used if the filters are off and the gains set to 1. The setup here is again somewhat different depending on what type of measurement one is doing. If using the 'Fourier Tools' option then the 'Waveform' must be chosen. Either of the two 'Noise' options in this drop down menu will do. If using a 'Swept Sine' DTT does not give you a choice other than 'None' because the waveform is already specified. Also specific to 'Fourier Tools' is the 'Filter' option. This option is not required, but getting a good measurement is difficult without it. Since the noise excitation applies drive at all frequencies, you are likely to saturate the actuators before you get any reasonable drive at any of them. Thus, it is necessary to define a bandpass filter between the 'Start' and 'Stop' values set in the 'Measurement' tab. Clicking on the 'Foton' button will open a 'Foton' window where you can set this filter. You can also just type in the parameters if you know the format. A filter that works well is `ellip("BandPass",4,1,40,0.1,10)`, which sets a 4th order filter with a pass band between 0.1 Hz and 10 Hz with a ripple of 1 dB and a reduction of 40 dB outside the band. At LASTI

is was also necessary to add poles as zeros to this filter to cancel the post-filters set in the output medm screens. This was because these post-filters (recall they cancel the dynamics of the coil drivers) amplified the 1 Hz to 10 Hz part of the band enough that the DAC would saturate before enough drive reached the 0.1 Hz to 1 Hz part of the band. Since that part of the spectrum did not need so much drive it was partially suppressed.

The final parameter to set in this tab is the 'Amplitude'. This value determines the strength of the excitation. Setting it is generally guess and check. Start with a small value, take a measurement, and see how it looks. If it is too noisy and has low coherence, increase the value. This process should only occur while the OSEM sensor signals are being monitored through either the Input_Signals screen or Data Viewer to make sure the suspension is not driven too hard. The DAC error light at the bottom of the overview screen also should be monitored to make sure the excitation is not being clipped by the DAC limit. This limit is always set to an output signal of 32000 counts (the largest value any of the actuators can put out). When the DAC saturates the indicator light goes from green to red and the 'Overflow' box above counts up (1 count for every 16384 Hz sample). If using the 'Swept Sine' option it is more obvious when the DAC will saturate because the amplitude of the sine wave is the value entered in this box. One could figure out the amplitude of the OSEM output simply by passing this value through the output matrix (the output filters will change this somewhat depending on the frequency). Note, while setting this amplitude with the guess and check method you can speed up each measurement by increasing the value entered into the 'BW' box in the 'Measurement' tab for noise excitations, or by reducing the number of 'Points' with swept sine excitations.

The test is engaged by pressing the 'Start' button at the bottom. 'Abort' will stop the test.

The 'Results' tab is where all the measurement plots are given. I typically put the transfer function magnitude in the top plot and the coherence in the bottom plot. Make sure to choose the appropriate signals in the 'Traces' tab. The excitation should go in 'Channel A' while the response goes in 'Channel B'. The graph field should be set to 'Transfer Function' on the top plot and 'Coherence' on the bottom. In the 'Units' tab of the transfer function plot, the 'Y' 'Display' should be either 'Magnitude' or 'dB Magnitude'. The 'Range' tab should then be log-spaced for the x-axis. The y-axis should be log-spaced if the 'Magnitude' units were chosen and linear-spaced if 'dB Magnitude' was chosen. For the coherence plot the x-axis should be log-spaced while the y-axis linear-spaced.

If this test is a diagnostic measurement to compare with a reference, the reference will need to be reset before beginning the measurement using the 'Reference' button at the bottom of the screen.

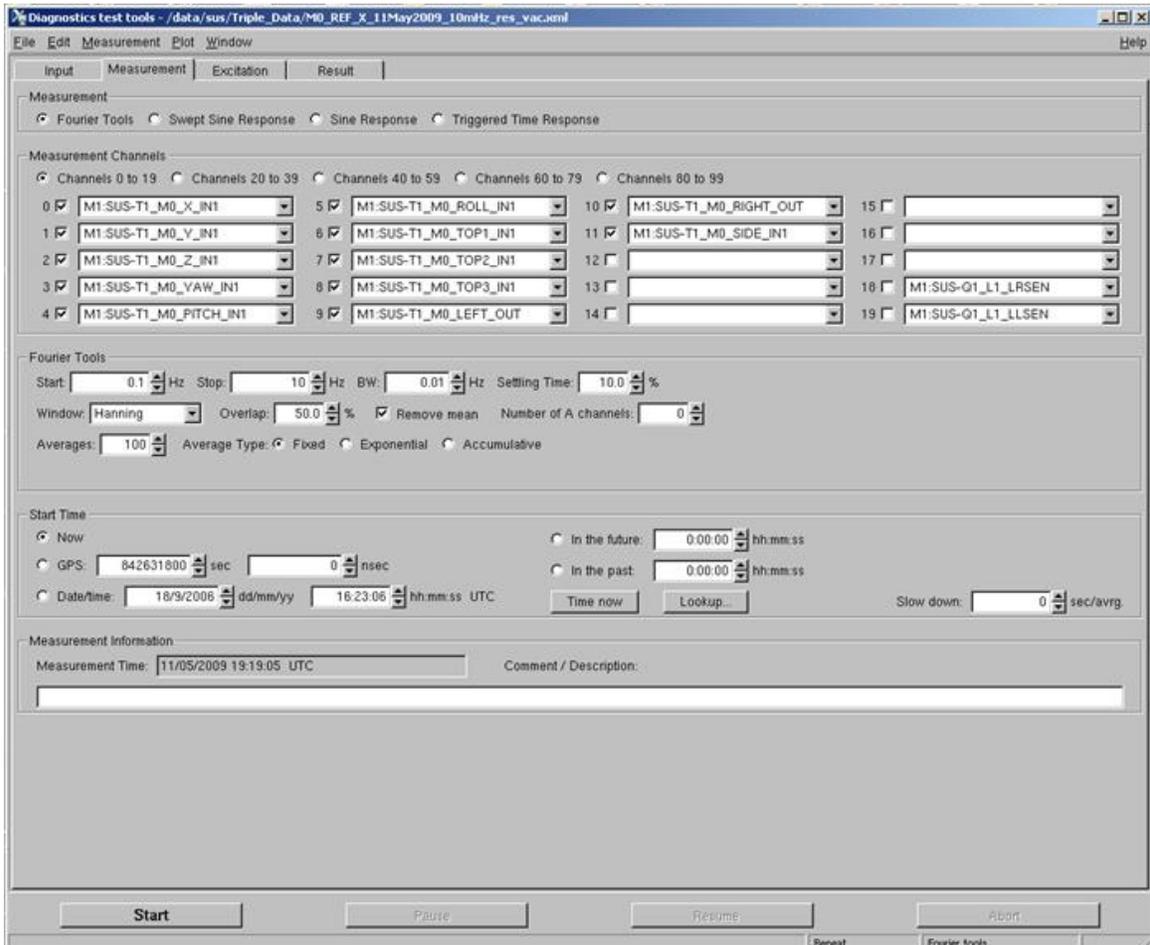


Figure 12: The DTT screen where the measurement is set.

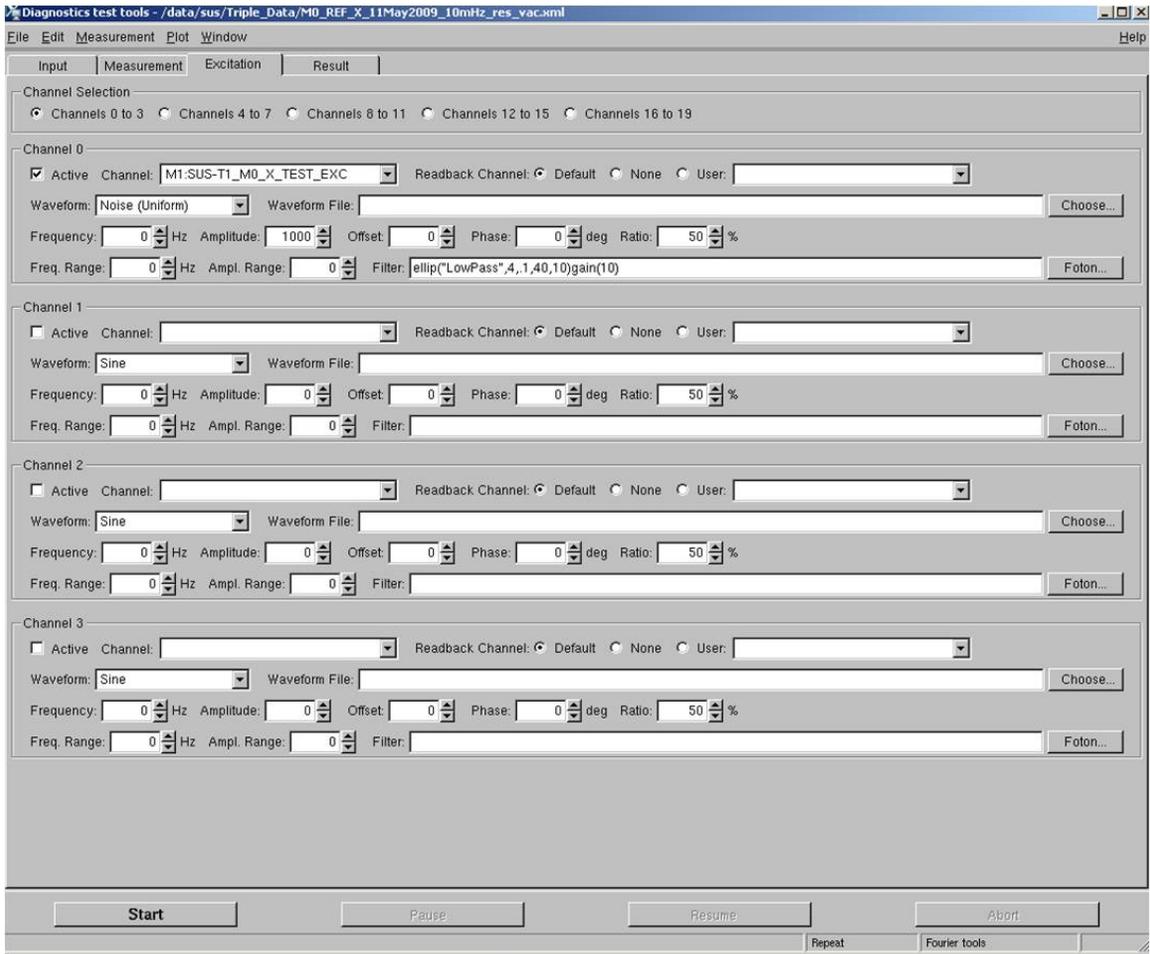


Figure 13: The DTT screen where the excitation is set.

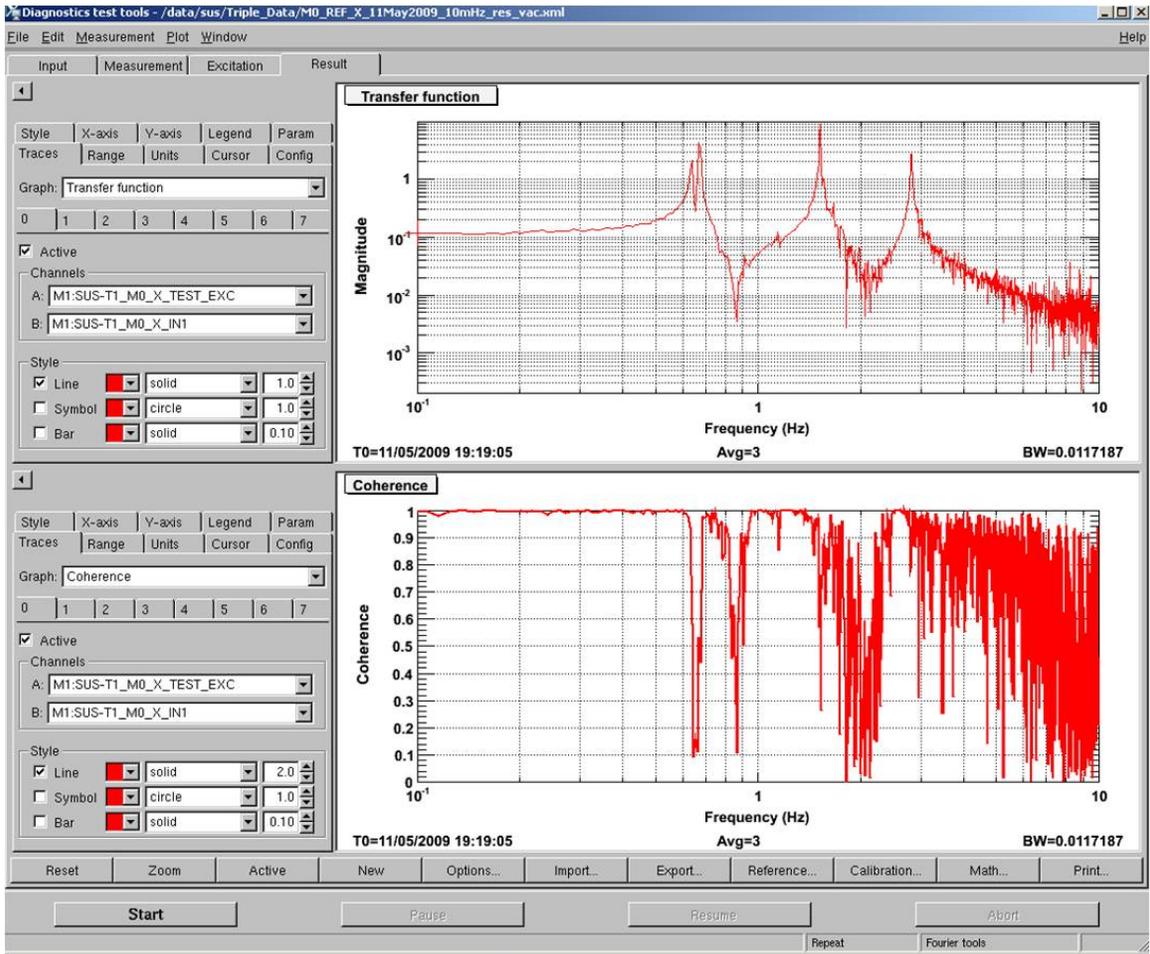


Figure 14: The DTT screen where the results are viewed.

2.6 Foton Basics

Foton is opened either by clicking on the appropriate icon or typing the word foton into a terminal.

Foton is the software used to load filters into the MEDM screens. In Foton each system has a .txt file that stores all the filters sent to MEDM. Foton is a convenient interface to edit this file (the file should NOT be edited manually except under special circumstances!). In Foton the file is loaded under the 'Design' tab in the 'File' field within the 'Module Selection' box. The desired channel is chosen in the 'Module' field. Before anything can be edited it is necessary to un-click the 'Read Only' option in the file menu at the top left.

The 'Sections' box to the right has 10 different filter sections, which correspond to the 10 sections in each MEDM filter module. The names that appear in each section of the MEDM window are given here. The check box establishes which filter section is active in the in the 'Design' box below.

Filters can be directly typed into the 'Command' field but it is usually easier to chose an option in the 'Add' field such as 'ZPK' or 'Ellip'.

IMPORTANT: make sure that the value in 'fSample' is correct. It should be set to the sampling frequency of the system, probably 16384 Hz. The Foton default is not necessarily correct. A wrong value may result in an unstable filter. In fact, this is one of the special circumstances where it is better to edit the text file directly. If all the filters have the same wrong default value it is a simple matter to open the text file in a text editor and do a 'replace all'.

When the filter is entered it can be viewed in the 'Graphics' tab. To display the filter it is necessary to click on the 'Bode Plot' button on the bottom left. Note that the 'Graphics' tab is setup the same as the DTT 'Results' tab.

Press the save button when the filter is ready.

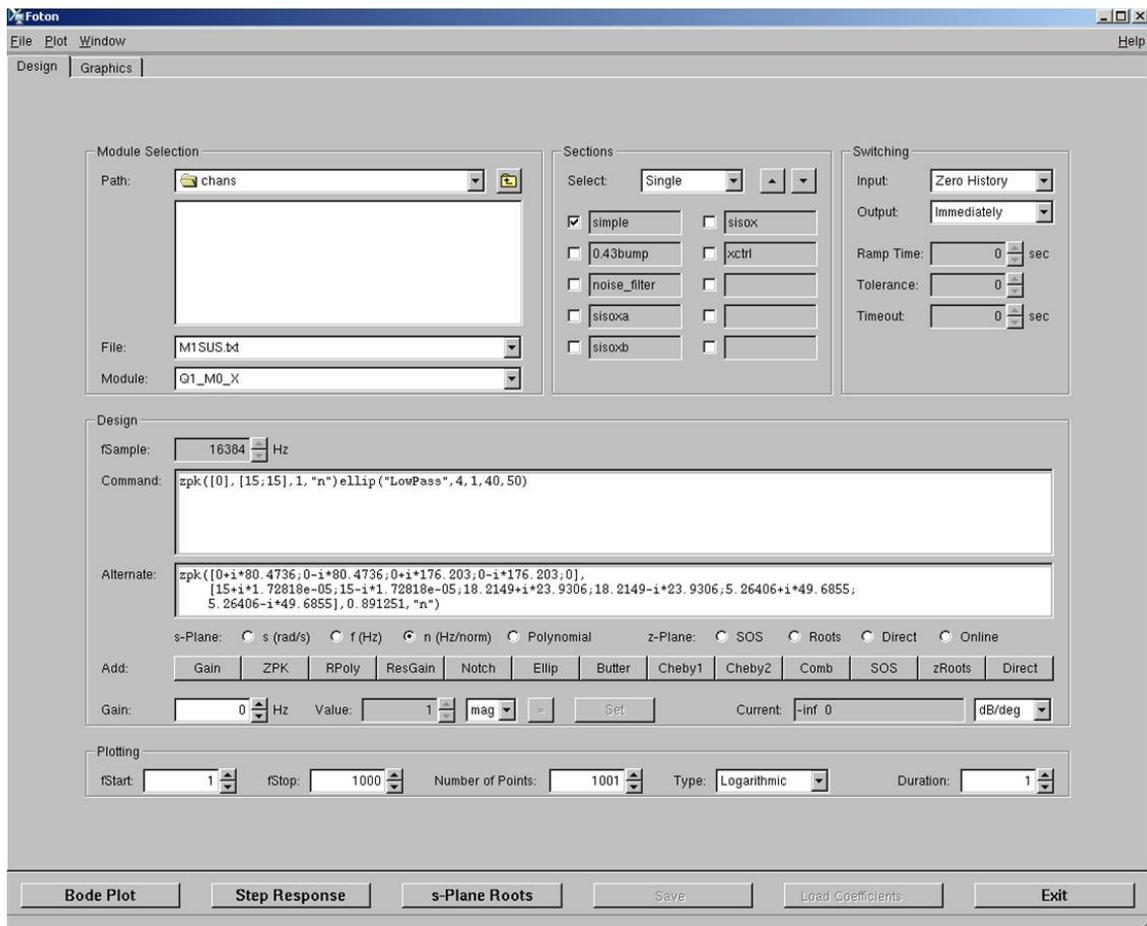


Figure 15: The Foton screen where filters are defined.

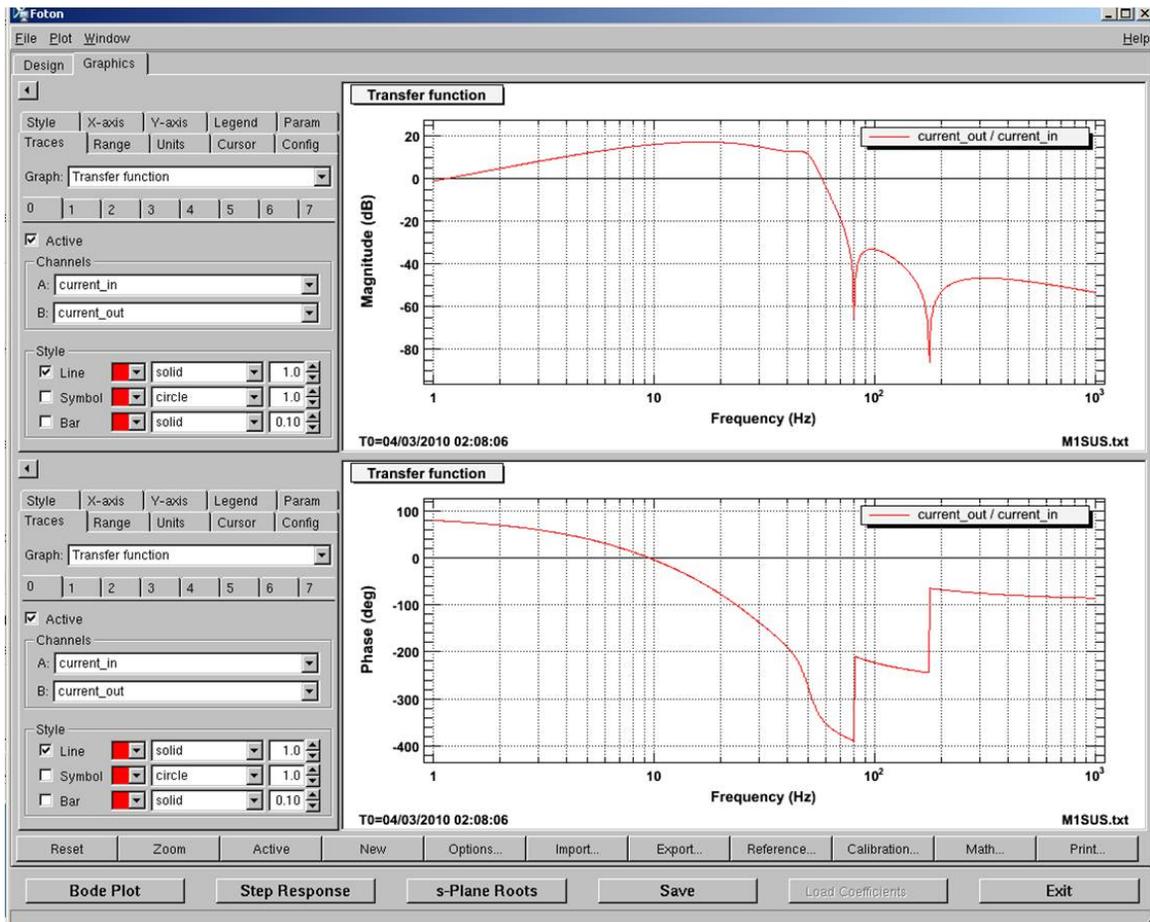


Figure 16: The Foton screen where filters are viewed.

2.7 Brief notes on the Simulink-Frontend Interface

See Rolf's document T080135, "AdvLIGO CDS Real-time Code Generator (RCG) Application Developer's User Guide." There should also be a text file available to the local workstation with a set of instructions on how to compile the Simulink diagram.

3 Transfer Functions

3.1 What are Transfer Functions?

Transfer functions are a method of characterizing the behavior of linear systems. They measure the relationship between the input and output of the system at a variety of frequencies. Since linear systems respond differently at different frequencies, measuring the response over the full spectrum of frequencies will capture all there is to know about the system's behavior. In fact, with a precisely measured transfer function, one could reconstruct the equations of motion of the system and predict exactly how it would respond to any input.

A system is said to be linear when the output from the sum of two inputs is equal to the sum of the outputs from each input. Similarly, scaling the magnitude of an input scales the magnitude of the output equivalently. In equation form this appears as

$$f(x + y) = f(x) + f(y) \quad (1)$$

$$f(ax) = af(x) \quad (2)$$

Pendulums are inherently nonlinear, but for small motions like those permitted by the earthquake stops, we can accurately pretend they are.

There is a property of linear systems that the system responds to a wave (typically a sine wave) at only the frequency of the wave. To be more precise, this statement is mostly true. When a system first receives an input excitation there are usually some temporary transient behaviors, related to motion at the resonance frequencies. This transient response will decay due to the natural damping of the system leaving behind only the forced response at the frequency of the input. The time after the transient response decays is called the steady state. So this property is said to hold at the steady state.

The simplest way to measure a transfer function is to literally follow the property stated above. One injects a force into the suspension in the form of a sine wave at some frequency and amplitude. After some amount of time the suspension responds by moving at this same frequency. Two pieces of information can then be gathered from this response. Its amplitude and phase. Another name for the phase is 'time-delay'. Often, the suspension will not respond immediately to the input so the response sine wave will be shifted out of phase with the input sine wave. This time delay is typically measured in units of degrees, where the length of a full sine wave cycle in time represents 360 degrees. This input-output relationship is shown graphically in the example of Figure 17. In this figure, input and output waveforms of 1 Hz are shown. The transfer function values in this example would be a magnitude of $\frac{1}{2}$ and phase of 180° because the output is half the input and shifted exactly

out of phase. If we were to record these two values over a range of frequencies and plot the results we would have generated a transfer function. Transfer functions can have units as well. If the input was in N and the output in m, then the magnitude plot would have units of m/N. Figure 18 gives an example transfer function.

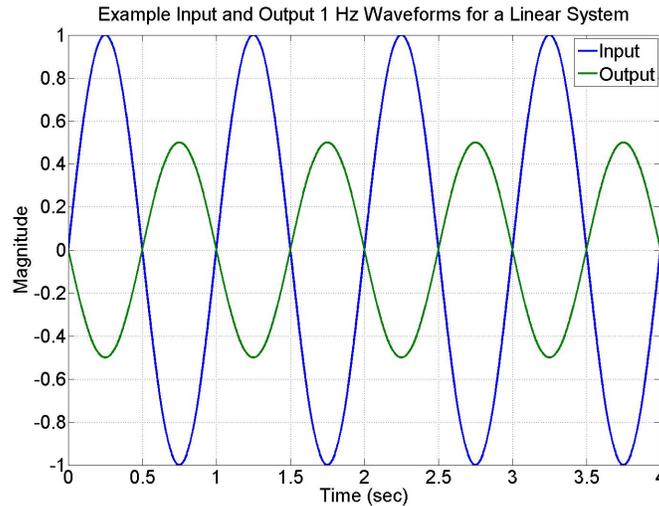


Figure 17: Example input and output waveforms for a linear system. The output amplitude is $\frac{1}{2}$ the input and shifted by 180° .

This method of measuring a transfer function is called a swept-sine or stepped-sine because individual sine waves are tested in increments of frequency. This is not the only way to measure transfer functions. It has the advantage of being very precise, but it can be rather slow. A faster method is to inject an input that has a spectrum of frequencies. The transfer function is then the spectrum of the output divided by the spectrum of the input. Since each signal is spread over an entire space of frequencies the measurement is less precise because each frequency has a weaker signal and any noise present becomes more significant. One can limit the range of frequencies injected to a limited band to improve the accuracy (naturally, the limit as the bandwidth goes to zero reconstructs the swept-sine case). Also, one can repeat this measurement a number of times and then average all the results together to get a more accurate measurement. Then you only need to take enough averages to get the precision you need. DTT can measure transfer functions using both of these methods.

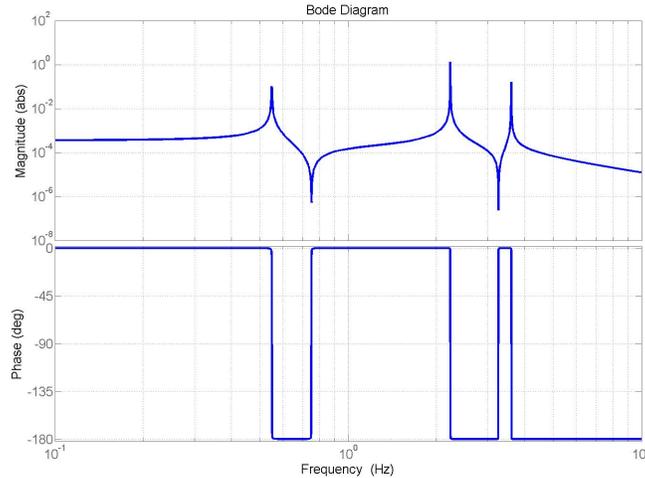


Figure 18: Example transfer function. The top plot shows the ratio of the output and input wave form amplitudes at each frequency (arbitrary units here). The bottom plot shows what the phase difference is between the two wave forms at each frequency. The axes are usually log spaced so that more detail is visible.

3.2 Debugging Suspensions with Transfer Functions

While transfer functions for debugging a suspension in this document are given using the quad suspension as an example, it is applicable to any aLIGO suspension by analogy.

While testing the suspension with transfer functions it is useful to keep in mind the coupling between DOFs one should expect. For example, the suspension can bounce up and down along the Z axis without moving in X, Y, Yaw, PITCH or ROLL. Consequently, the resonances that appear in any Z-Z transfer function should not appear in any other transfer function. Ideally, Z motion is completely decoupled. So if it turns out there is some sharing of Z resonances then there is something less than ideal going on with the suspension. It actually does turn out, at least for the quad, that some coupling between Z and PITCH occurs when the top mass hangs with a small pitch offset. This is likely due to some coupling introduced by the slightly tilted magnets.

Yaw motion is similarly decoupled. X motion is almost entirely decoupled but has a small amount of low frequency coupling to PITCH. The coupling between X and Pitch is minimized due to the close proximity of the wires to the axis of rotation. Y and ROLL are strongly coupled, perhaps more so than any other set of DOFs due to the large distance of the wires from the axis of ROLL rotation.

Perhaps unique to the quad (though I cannot say this for certainty) is some high frequency coupling between PITCH and ROLL, particularly at 5 Hz. This coupling results from the large blade spring clamps in the top mass and UIM. They add a lot of mass in two diagonal corners of these stages. As a result the axis of pitch rotation is slightly slanted off the Y axis defined by the OSEMs. When the top mass tries to pitch it is effectively pointing in the

wrong direction so the LEFT and RIGHT OSEMs see their flags move up and down slightly. Looked at from the actuation side, since the LEFT and RIGHT OSEMs are on opposite sides of the axis of PITCH rotation, when they drive ROLL they also drive a little bit of PITCH. Since ROLL is coupled to Y, PITCH also has some coupling to Y. Basically, the quad PITCH is coupled somehow to everything, with the exception of YAW. In mathematical terms, the mass matrix has some off-diagonal elements. The triple suspensions that have 4 blade springs in each stage, as opposed to two like the quad, may not have this same PITCH-ROLL coupling (the LASTI triple does not).

Figure 19 summarizes into a table many of the possible mechanical bugs and the specific effects they have on the transfer functions. This table can be used to debug the suspension during commissioning. Transfer functions resulting from a number of these cases are plotted in the sections below.

A point about terminology, people often use the term resonance and mode interchangeably. Resonances and modes are manifested in the transfer functions by the large peaks.

The opposite of a resonance is a zero, shown on the transfer functions as the 'valleys'. One way to think of zeros in this context is as a resonance of the stages below the top mass (while measuring the top mass). In fact, if one were to clamp the top mass, the zeros as seen by the top mass would become the resonances of the smaller suspension below. The top mass sees these zeros because the stages below are resonating in such a way that they exactly cancel the force applied by the OSEMs at that frequency.

All the transfer functions shown here were measured with DTT using a noise input that was filtered with a bandpass elliptical filter between 0.1 Hz and 10 Hz. The frequency resolution is 0.01 Hz, and the measurement band is from 0.1 Hz to 10 Hz. Approximately 3 to 10 averages were taken for each measurement. All measurements are from the quad main chain top mass OSEMs (M0). By design, all dynamics of the entire suspension were designed to couple to the top mass. Thus, if anything changes in the stages below it will show up on the top mass transfer functions. For this reason we do not need to worry about measuring transfer functions anywhere else.

TF Debugging Table	Non-sus stop, magnet, or cable interference	Top stage spring locked	Top mass or UIM spring interference (sus)	OSEM/ESD cable tension (sus)	Misaligned OSEM	Response Notes
Drop in magnitude	X	X			X	Most easily observed at low frequencies. Bug is external to suspended chain.
Resonance and zero movement (TF shape)	X	X	X	O		From some form of localized stiffness.
Resonance frequency cross coupling	X	X	X	X	O	If all DOFs are effected something is touching a suspended mass. If some are not affected it is probably a spring or cable tension.
Asymmetric response	O	X	X		X	E.g. left interferences have a greater influence on the left
Visible in all DOFs	X					Giveaway for interferences touching a mass
Noisy appearance	O				X	If it shows up only is DOFs with a common OSEM then check alignment.
Bug notes	Anything touching both structure and suspension. Generally asymmetric response with exception of anything touching the vertical centerline such as FACE1.	If both springs are stopped up the response may be symmetric with no added cross couplings. Asymmetries or cross couplings that do appear are limited to z and roll.	Could be blade stops, UIM blade ECDs, blade tip adjusters, or a bottomed out FACE3 magnet/flag standoff (very hard to seal). The top mass springs produce more asymmetry than the UIM .	Large stiffness or tension will change the TF shape; usually z and pitch. Small stiffness increases may simply add cross couplings	A particular OSEM sensor will have a reduced magnitude. Resonances from other DOFs may poke through if the DOF uses multiple OSEMs, e.g. YAW may show some X if FACE2 or FACE3 is misaligned.	

KEY: sus = suspended, non-sus = non-suspended, X = definite effect, O = possible effect

Figure 19: This table generalizes the mechanical bugs likely to be found on the suspension and what effect they have. It is not exhaustive, but it covers the main categories of possible effects. Keep in mind that multiple bugs may exist simultaneously.

3.2.1 Quad M0 Reference Transfer Functions

Within the X-Y-Z coordinate frame there are a total of 36 possible transfer functions in a 6x6 matrix (e.g. X-X, X-YAW, PITCH-ROLL, etc). The diagonal transfer functions X-X, Y-Y, Z-Z, YAW-YAW, PITCH-PITCH, ROLL-ROLL are almost exclusively used while debugging so they are the only ones plotted here. These reference measurements were taken with the main chain top mass when the quad was in a fully functional, suspended state, in the wire hang configuration (glass optic, metal wires, March 2009 to Feb 2010). They are listed here so the proper responses can clearly be seen before delving into the 'buggy' responses below.

I chose to plot the coherence of the transfer function instead of the phase because the phase information is redundant while the coherence gives an idea of good the measurement is. In fact, one could reconstruct the phase information with only the magnitude plot if one desired. The coherence is more useful because it tells you how well the output waveform agrees with the input at each frequency. A value of 1 means a perfect match, a value of 0 means no match.

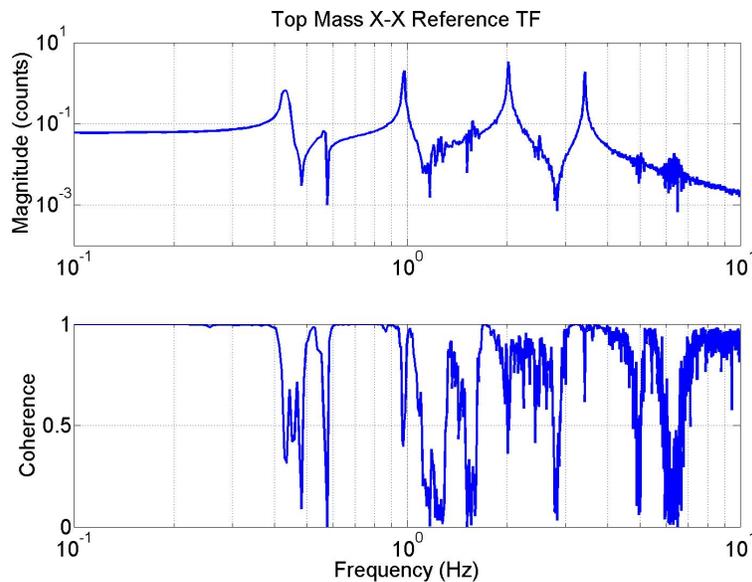


Figure 20: A transfer function measured from an M0_X excitation to an M0_X measurement. The 0.55 Hz pitch resonance is visible.

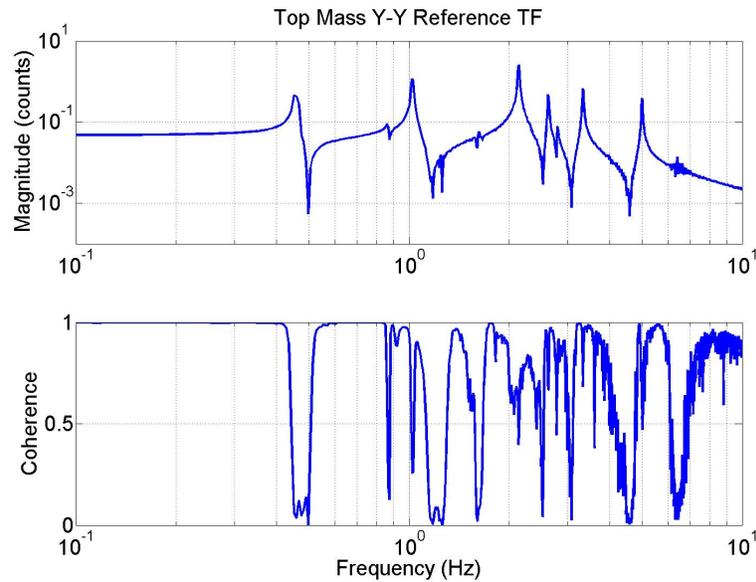


Figure 21: A transfer function measured from an $M0_Y$ excitation to an $M0_Y$ measurement. Many of the resonances in this TF are shared with $ROLL$. Some higher frequencies are also shared with $PITCH$.

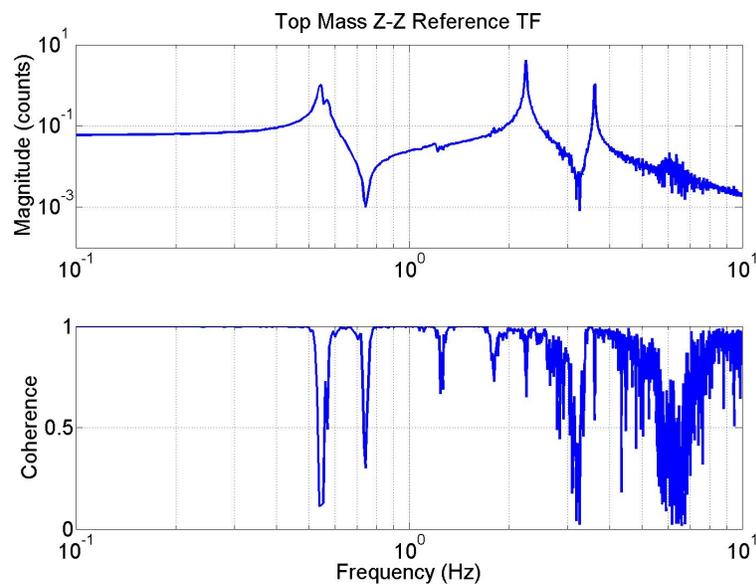


Figure 22: A transfer function measured from an $M0_Z$ excitation to an $M0_Z$ measurement. The first resonance appears doubled because a piece of the 0.55 Hz $PITCH$ resonance is poking through. This coupling is related to $PITCH$ offsets of the top mass.

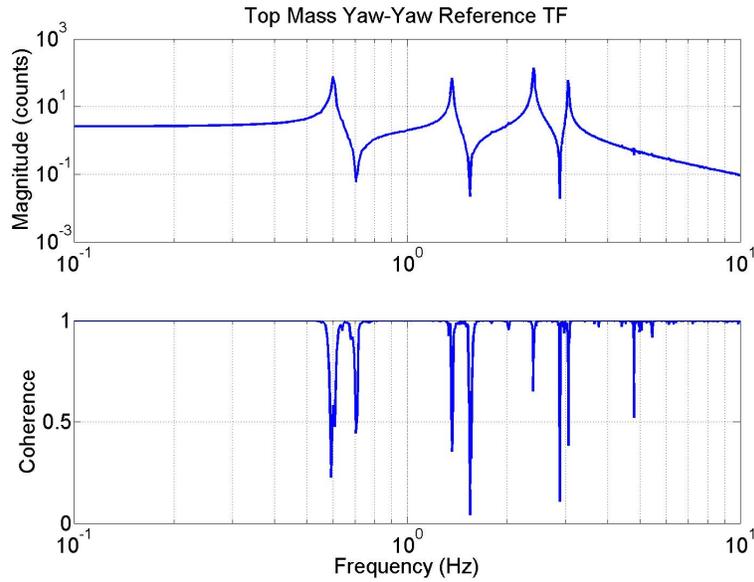


Figure 23: A transfer function measured from an M0_YAW excitation to an M0_YAW measurement. YAW is the most decoupled of all the DOFs. No resonances here are shared in any other DOF. YAW measurements tend to be the 'cleanest' because of the lack of coupling.

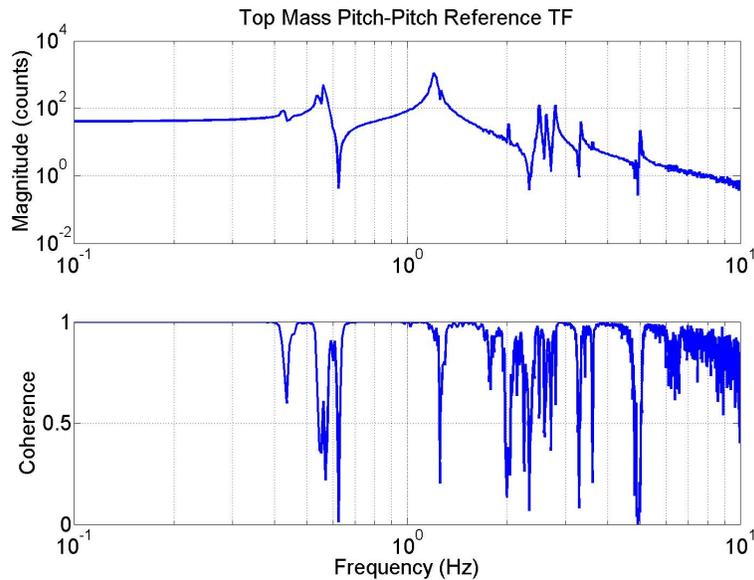


Figure 24: A transfer function measured from an M0_PITCH excitation to an M0_PITCH measurement. PITCH is the most coupled of all the DOFs. A small piece of the 0.42 Hz X resonance shows through, as well as the first Z mode. At higher frequencies many of the resonances are shared with ROLL and Y and cannot be said to belong to either of them.

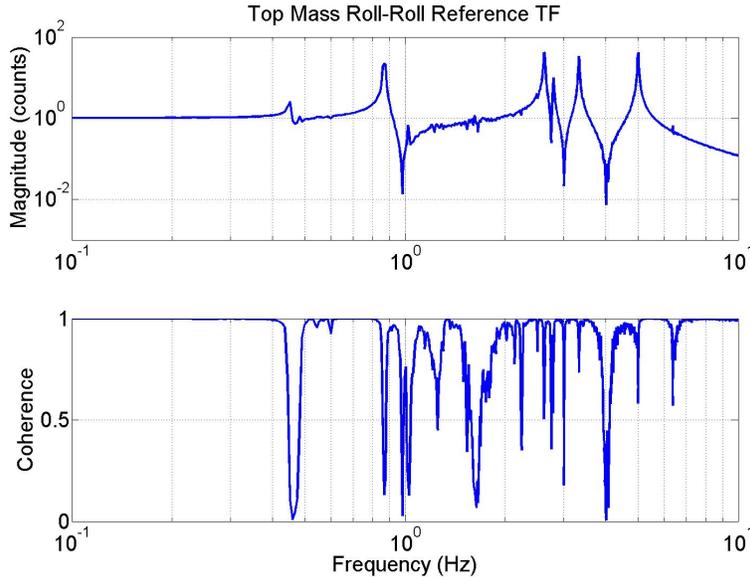


Figure 25: A transfer function measured from an M0_ROLL excitation to an M0_ROLL measurement. ROLL is coupled most strongly to Y, but at higher frequencies the coupling to PITCH is quite strong as well.

3.2.2 UIM Left Blade Spring Clamped

The left spring is defined as the spring supporting the left wire. The right spring supports the right wire (when facing the main chain).

This set of figures show transfer function measurements when the upper intermediate mass left blade is clamped. I chose to plot strictly magnitude plots here in order to zoom in and show more detail. All transfer functions are plotted on top of the references from the previous section to highlight the changes.

The first point to highlight is that the height, or magnitude, of the transfer functions are unchanged as shown by the extreme left ends. The low frequencies fall just on top of their references. This is because a static offset, from the point of view of the top mass OSEMs, does not care what state the springs below them are in. The entire group of four masses moves together. At higher frequencies the masses have more complicated motions and may even move in opposite directions. Thus, the DOFs where springs are most important, such as Z, ROLL, and Y show a significant changes. Some influence is seen at the higher frequency pitch resonances only because of its coupling to ROLL. X and YAW are not influenced by the springs and consequently show no change.

Another point is that the loss of mechanical symmetry can be seen by examining the component OSEM signals making up a measurement. A Z excitation for instance should produce the same response in both LEFT and RIGHT. If these components are different, than one can say that something is disrupting the symmetry of the quad, such as a locked spring. In the Z and ROLL measurements here there is some asymmetry in the region from 1 Hz to 4

Hz.

Note that the eddy current damping magnet on the UIM springs can produce a similar effect if it is rubbing inside the copper damper. Also, a similar effect results when the spring is sagging on the blade spring tip height adjuster.

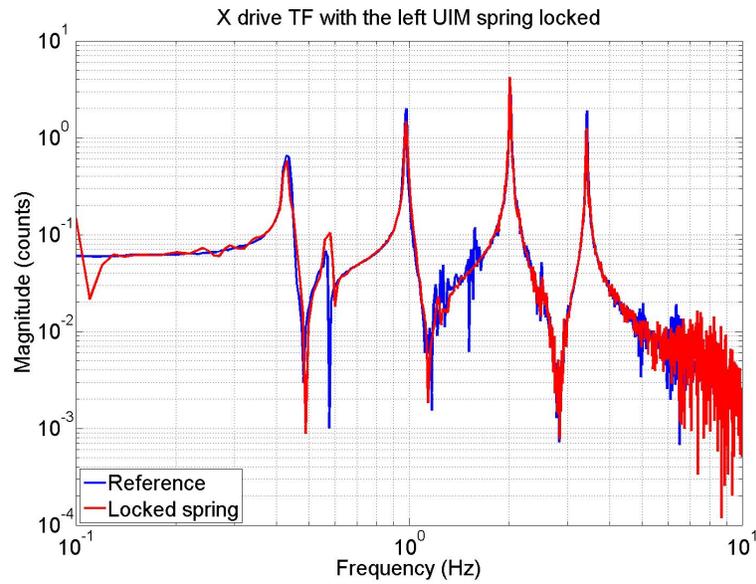


Figure 26: An X-X TF with the UIM left blade spring clamped. There is no dependence on the springs so there is no change.

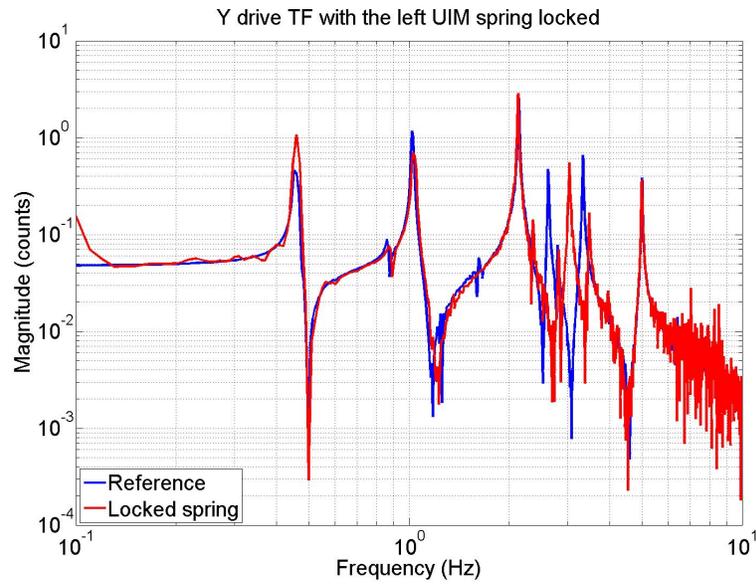


Figure 27: A Y-Y TF with the UIM left blade spring clamped. There is some high frequency modification of the resonances.

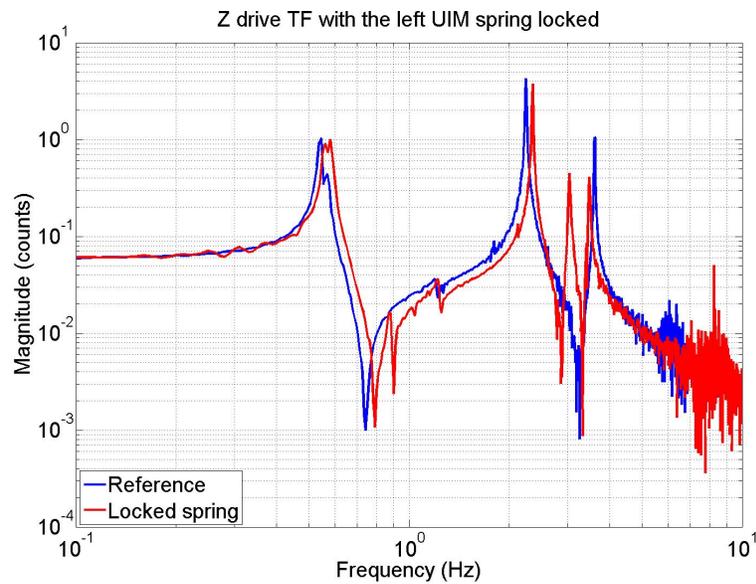


Figure 28: A Z-Z TF with the UIM left blade spring clamped. All the resonances move. A fourth appears because Z now couples to ROLL with the loss of mechanical symmetry.

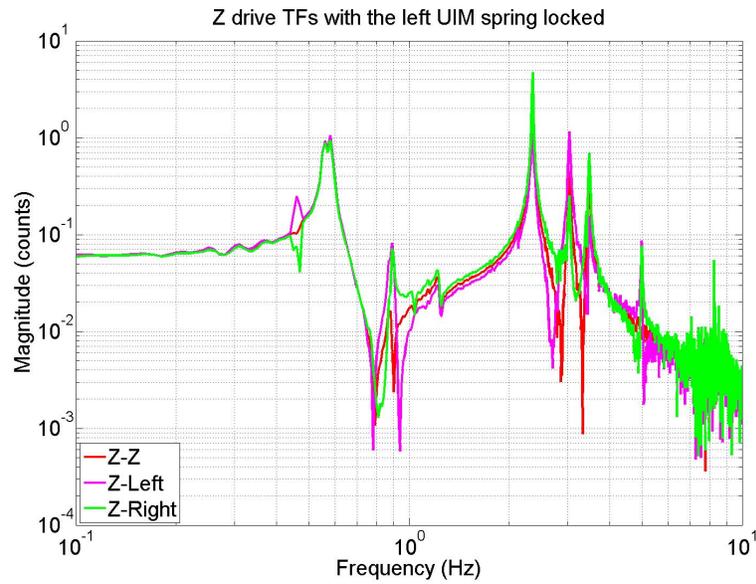


Figure 29: A Z-Z TF with the UIM left blade spring clamped. The Z transfer function is shown with its component signals. The fact that LEFT and RIGHT have a slightly different response between 1 Hz and 4 Hz points to a loss in mechanical symmetry.

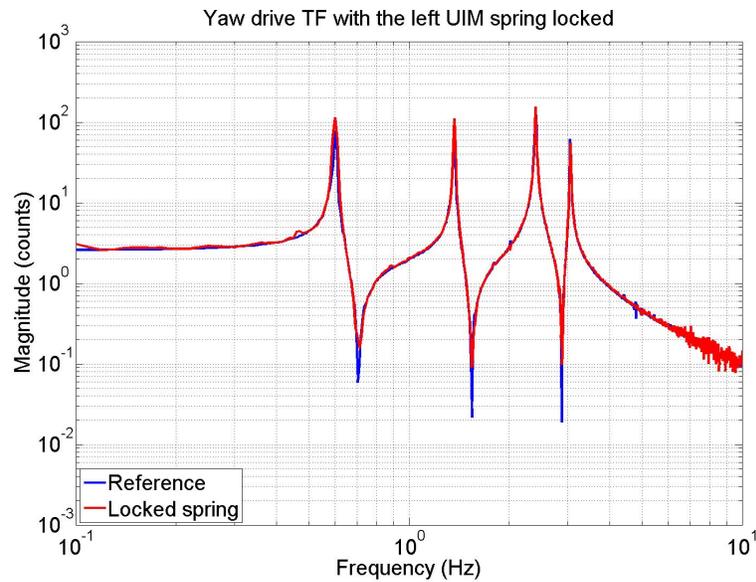


Figure 30: A YAW-YAW TF with the UIM left blade spring clamped. No change, the springs have no coupling to YAW.

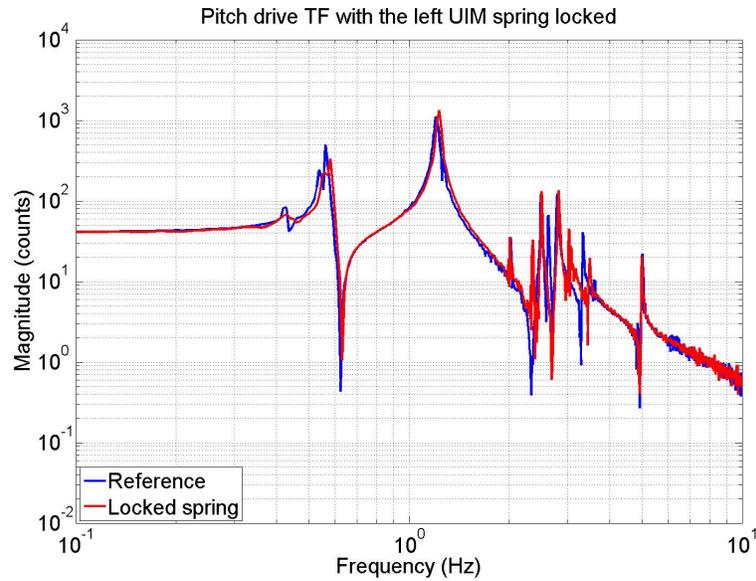


Figure 31: A PITCH-PITCH TF with the UIM left blade spring clamped. Some minor high frequency resonance shifts due to the coupling with ROLL.

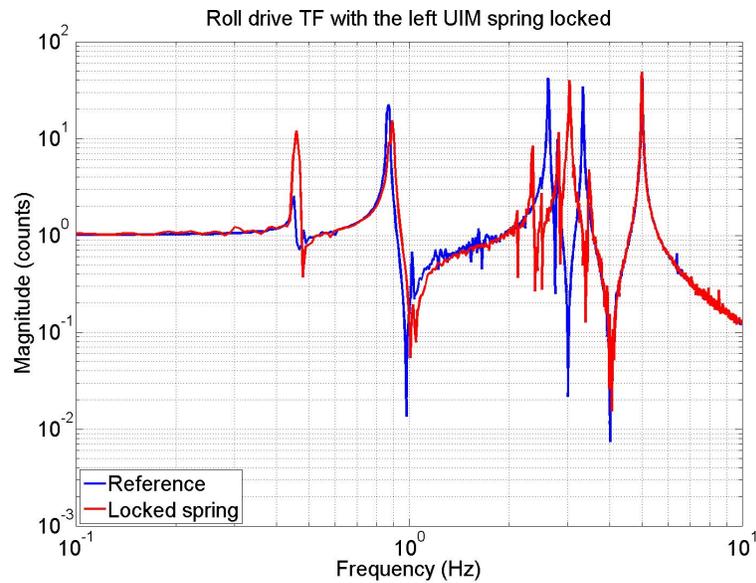


Figure 32: A ROLL-ROLL TF with the UIM left blade spring clamped. There is some high frequency resonance movement.

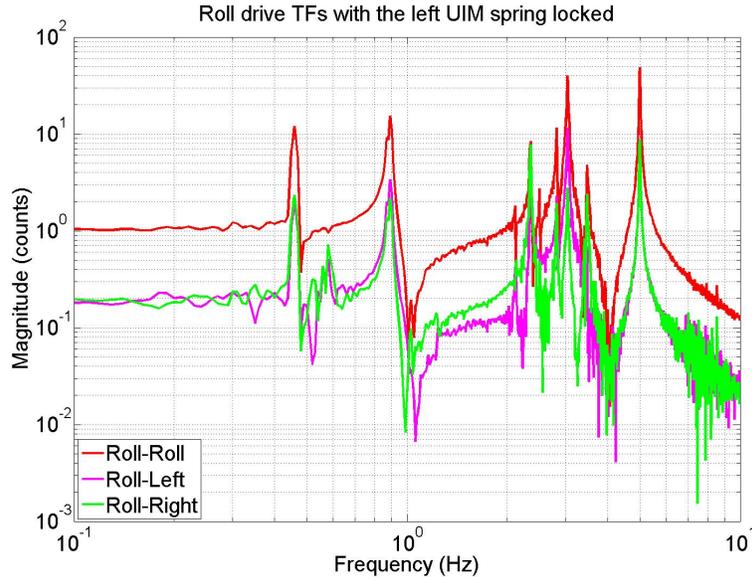


Figure 33: A ROLL-ROLL TF with the UIM left blade spring clamped. The slight difference in the LEFT and RIGHT OSEM response between 1 Hz and 4 Hz points to loss in symmetry. The response is not as pronounced as when a top mass spring is clamped because the UIM springs are less important to top mass motion.

3.2.3 Top Mass Right Blade Spring Clamped

The left spring is defined as the spring supporting the left wire. The right spring supports the right wire.

The next set of figures show transfer function measurements when the top mass right blade is clamped. Since the X-X and YAW-YAW transfer functions are again unchanged due to their lack of dependence on springs, they are not plotted here. The results are largely the same as those with the UIM spring clamped. The transfer functions change in a slightly different way because the top mass spring influences the dynamics in a slightly different way.

More asymmetry between the LEFT and RIGHT responses of the Z and ROLL measurements shows up for the top mass spring than the UIM spring. In particular, notice that at high frequencies above 1 Hz the RIGHT response is lower overall. This information is enough to point to an interference in specifically a right spring. Since the top mass spring stops are difficult to see when the tablecloth is put together, this information can speed the process of finding the issue. Additionally, if one knows beforehand what type of asymmetry to expect in the LEFT and RIGHT responses it is possible to distinguish between a UIM spring and top mass spring.

Note that the $\frac{1}{4}$ – 20 magnet-flag mount behind the FACE3 OSEM will bottom out on the spring anchored behind it if screwed in all the way. This is a VERY hard bug to find if its possibility is not known. The transfer functions will behave as if this spring is partially locked. As with the UIM springs, they can also sag onto the spring tip height adjusters,

producing a similar effect again.

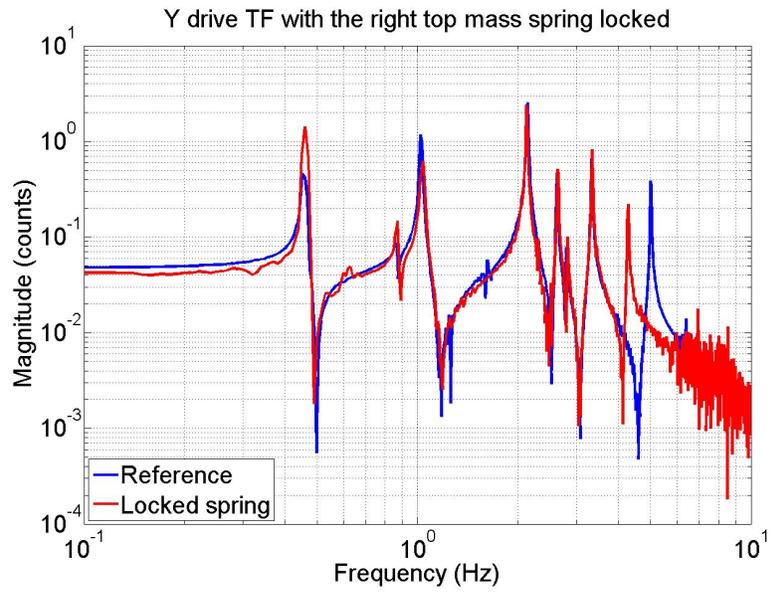


Figure 34: A Y-Y TF with the top mass right blade spring clamped. The highest frequency resonance changes more than the others.

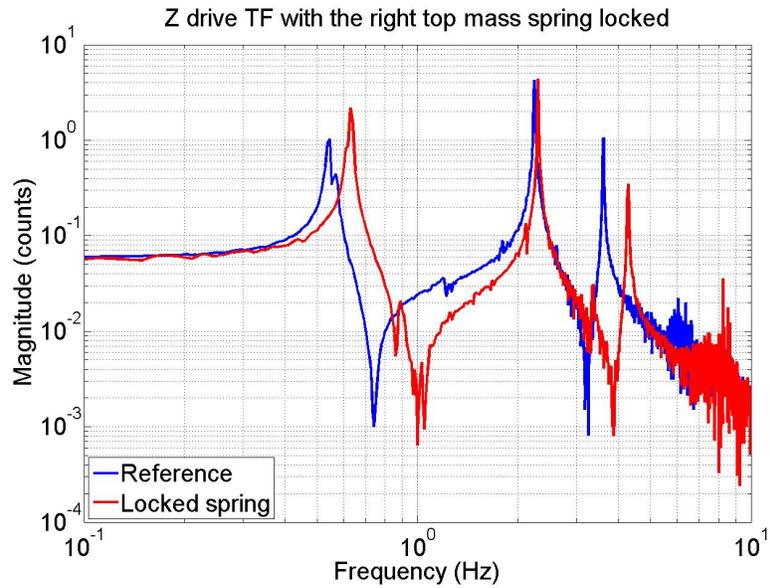


Figure 35: A Z-Z TF with the top mass right blade spring clamped.

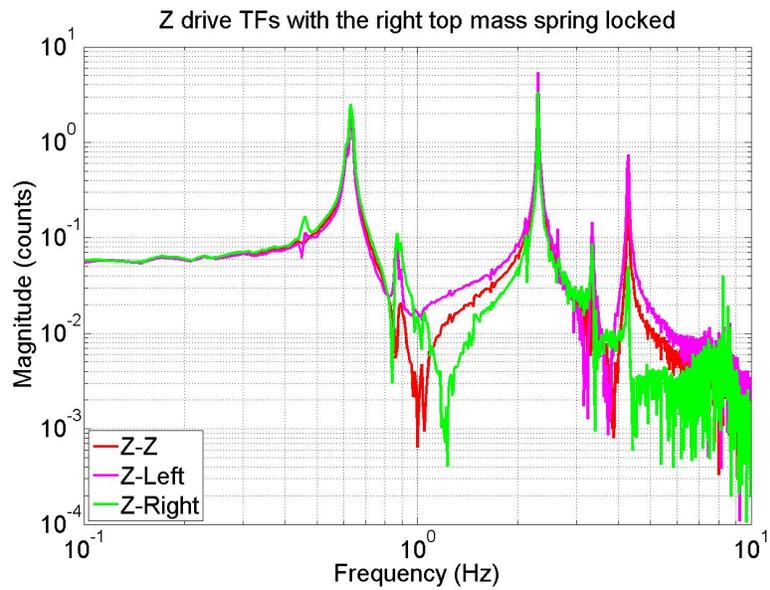


Figure 36: A Z-Z TF with the top mass right blade spring clamped. The fact that the RIGHT OSEM responds less than the LEFT above 1 Hz, plus the fact that the low frequency is unchanged points very strongly to a clamped right spring.

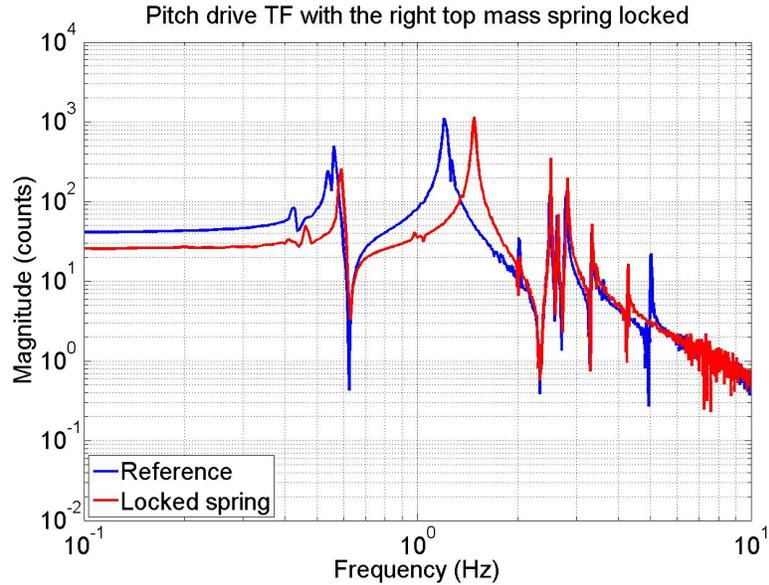


Figure 37: A PITCH-PITCH TF with the top mass right blade spring clamped. This result is different from the PITCH-PITCH transfer function measured with the UIM spring clamped due to a side effect of this measurement. When clamping the top mass spring, the tip was pushed down a significant amount (increasing the d value in the model). Pushing spring tips down increases overall pitch stiffness. This is why the low frequency end of the transfer function has dropped and the first two main peaks moved up in frequency. Referring back to the UIM result, one actually can make out a slight increase in these two frequencies as well for the same reason.

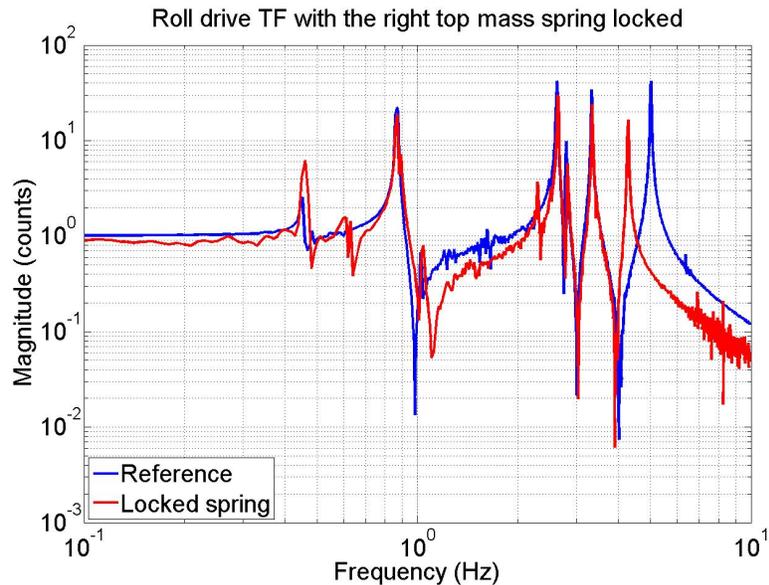


Figure 38: A ROLL-ROLL TF with the top mass right blade spring clamped.

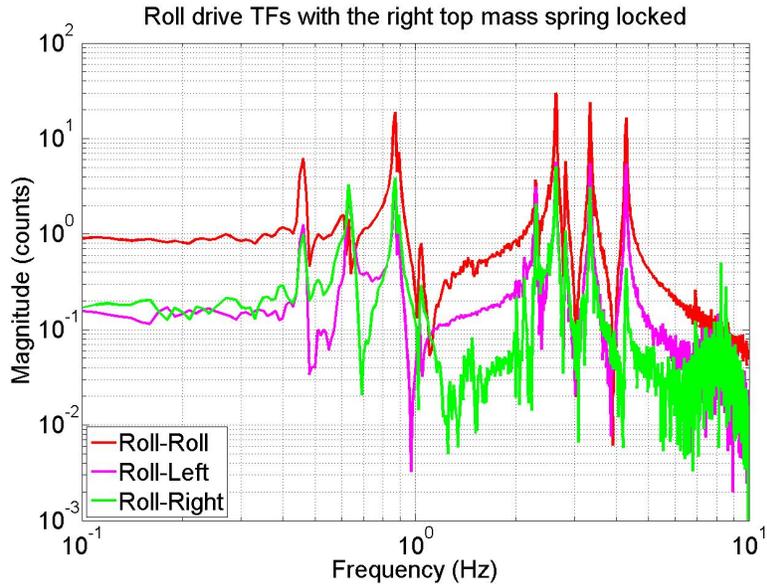


Figure 39: A ROLL-ROLL TF with the top mass right blade spring clamped. The fact that the RIGHT OSEM responds less than the LEFT above 1 Hz, plus the fact that the low frequency is unchanged points very strongly to a clamped right spring in one of the masses.

3.2.4 Top Stage Right Blade Spring Clamped

The next set of figures show transfer function measurements when the top stage right blade is clamped. The X-X and YAW-YAW transfer functions are not influenced and not repeated in this section. While the top stage springs are similar to the springs in the masses there is an important difference. From the point of view of the top mass OSEMs they increase the overall stiffness of the suspension. For this reason a Z-Z transfer will show a low frequency drop in magnitude, indicating the suspension is not fully suspended. This is similar to how Z would respond for an earthquake stop or any other interference external to the suspension. ROLL responds in a similar way, however the low frequency drop in magnitude happens to be less obvious. The difference between the LEFT and RIGHT OSEM signals for these DOFs is immediately obvious, with the LEFT much higher than the right throughout.

The drop in Z-Z magnitude, the asymmetry between the LEFT and RIGHT responses for both Z and ROLL, and the unaffected responses of X and YAW indicate a top stage right spring interference. The shift in the 5 Hz mode in the PITCH TF is further evidence since this mode couples strongly to ROLL.

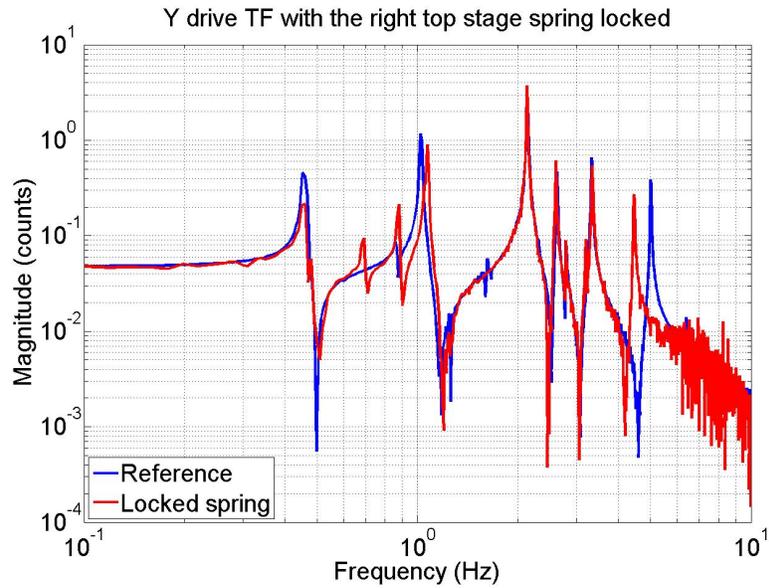


Figure 40: A Y-Y TF with the top stage right blade spring clamped. Some modes have shifted, and coupling to Z is introduced.

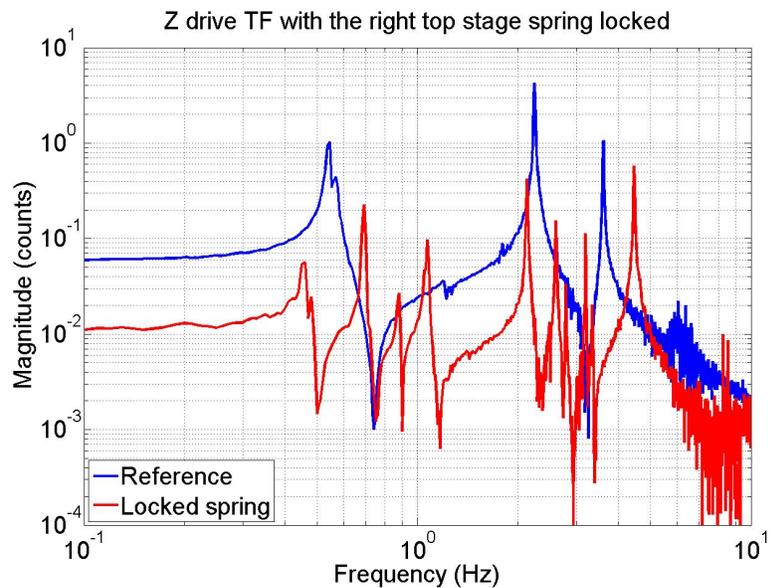


Figure 41: A Z-Z TF with the top stage right blade spring clamped. Note the overall drop in magnitude and the coupling to ROLL and Y.

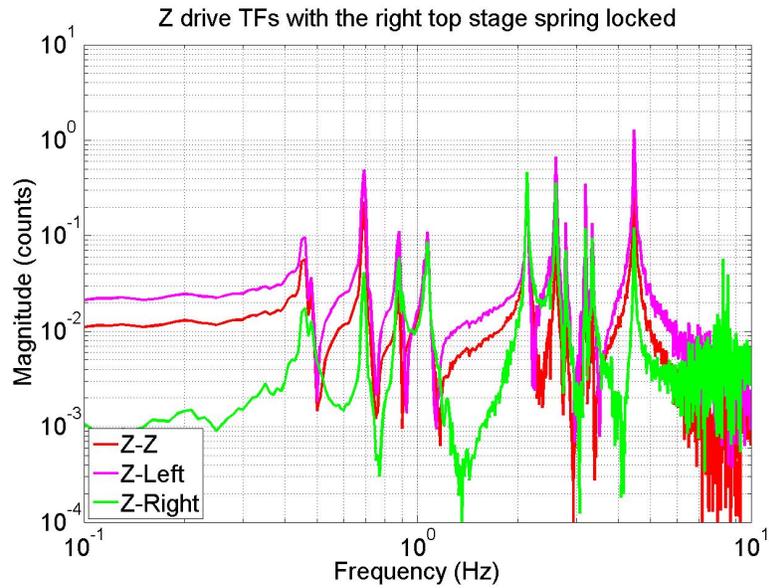


Figure 42: A Z-Z TF with the top stage right blade spring clamped. The RIGHT signal is much smaller than the left.

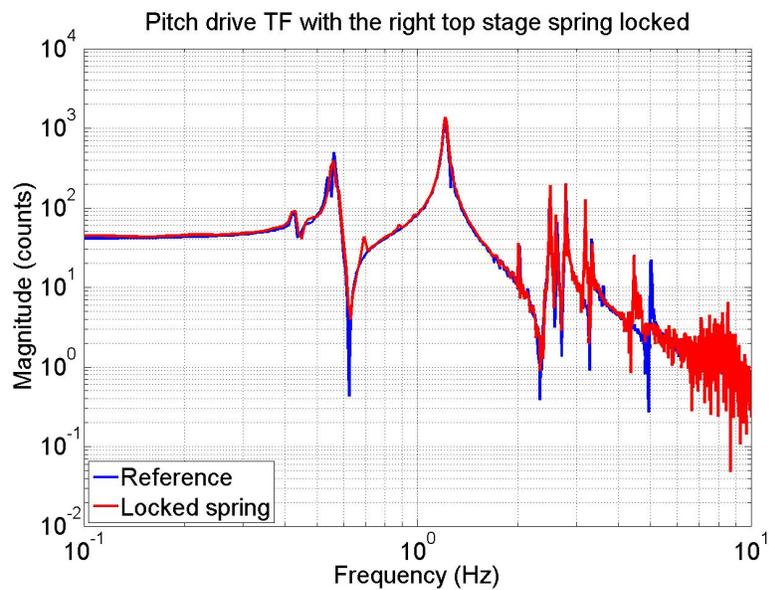


Figure 43: A PITCH-PITCH TF with the top stage right blade spring clamped. It is largely unaffected except for the 5 Hz mode which couples strongly to ROLL.

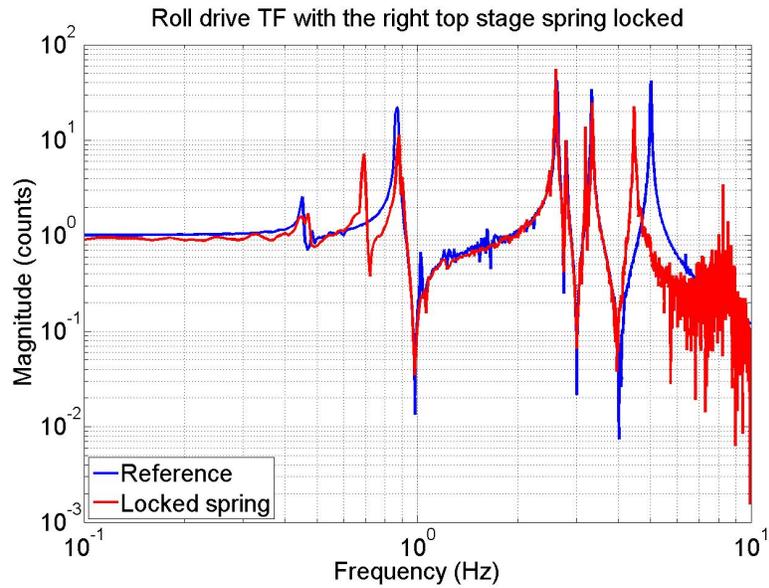


Figure 44: A ROLL-ROLL TF with the top stage right blade spring clamped. The low frequency drop in magnitude is misleadingly small in this case. There is also coupling to Z.

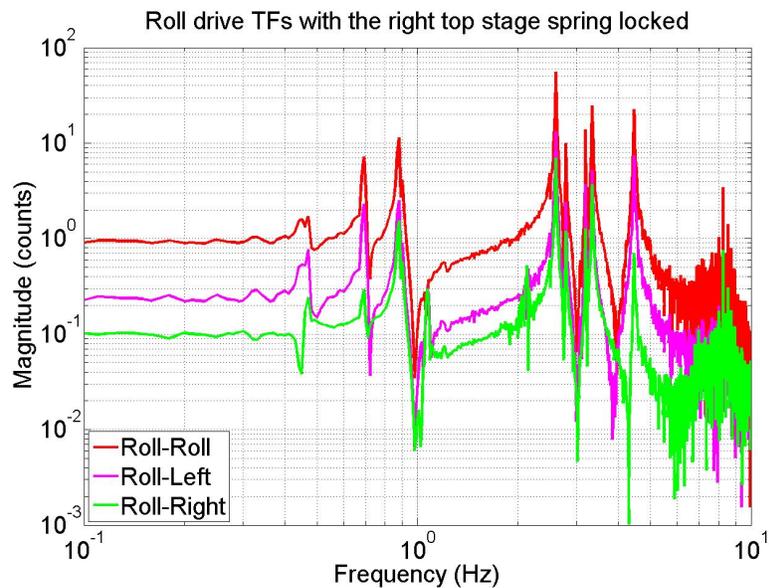


Figure 45: A ROLL-ROLL TF with the top stage right blade spring clamped. The RIGHT signal is much smaller than the left.

3.2.5 FACE1 OSEM Magnet Rubbing Inside OSEM Bore

The next set of figures show transfer function measurements when the FACE1 OSEM magnet is rubbing inside the OSEM magnet bore. In this case the flag did not sit flush on the magnet so it was actually possible to 'align' the flag to the OSEM sweet spot while the magnet was bumped up against the wall inside.

Any external interference to the suspension such as earthquake stops or rubbing magnets (with the notable exception of top stage springs) will impact all of the transfer functions in a similar way. The overall magnitude of all the measurements will drop, resonances will shift, and coupling between DOFs will develop. Typically these interferences introduce asymmetry into the measurements that point to where they are. In this case there is no obvious asymmetry. LEFT and RIGHT respond mostly the same for both Z and ROLL; and FACE2 and FACE3 respond the same for both X and YAW. The nearly symmetric response points to a symmetric interference. The ways of getting such symmetry with these interferences is either through a magnificent stroke of luck where multiple stops are locked down in a perfectly symmetric way (in which case you should play the lottery), or the interference somehow occurs along a line of symmetry. In this case FACE1 happens to be right on the vertical center line of the suspension, so it effects the left and right sides nearly equally.

The other notable way to get a symmetric interference on the quad is from tension on the OSEM and ESD cables going down the reaction chain. Since these cables are internal to the reaction chain they will not effect the low frequency magnitude of any of the transfer functions with the exception of PITCH (PITCH is the only DOF where internal interferences can effect the overall stiffness and thus the low frequency response). Additionally, depending on the particular case of cable tension there may be very little influence on X and YAW.

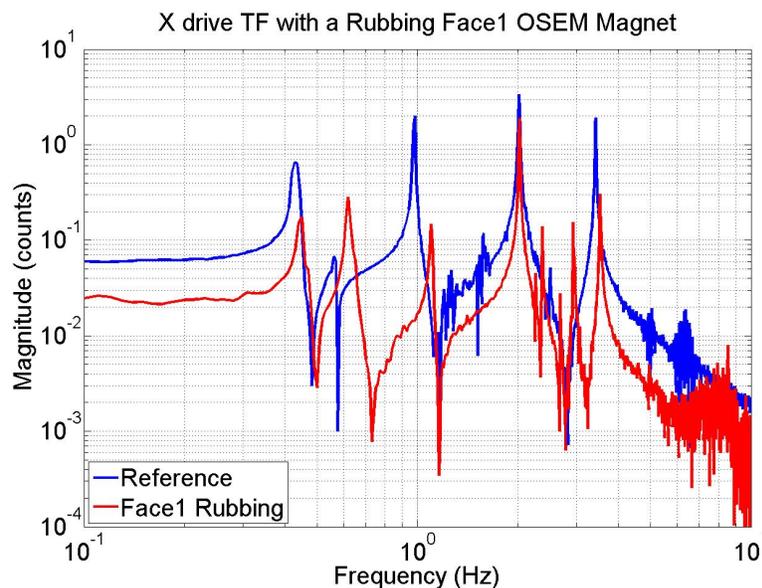


Figure 46: An X-X TF with a rubbing FACE1 OSEM magnet.

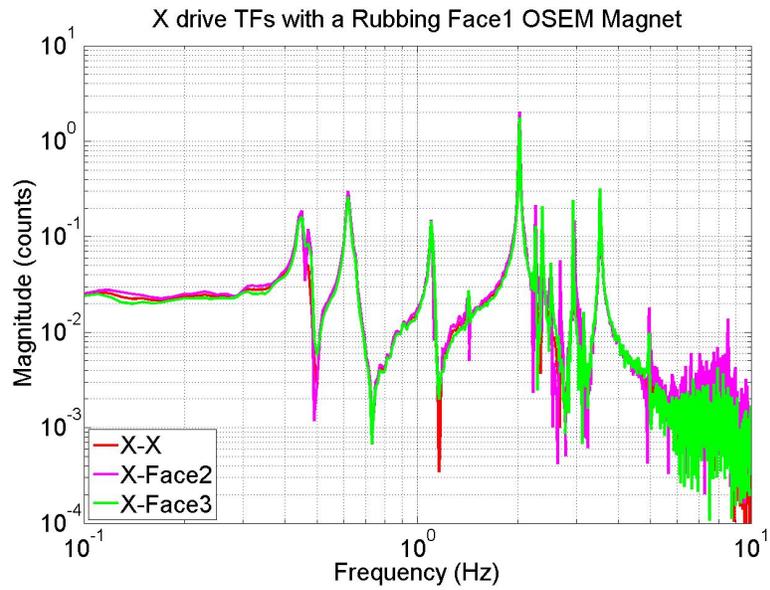


Figure 47: An X-X TF with a rubbing FACE1 OSEM magnet. Nearly symmetric response from FACE2 and FACE3.

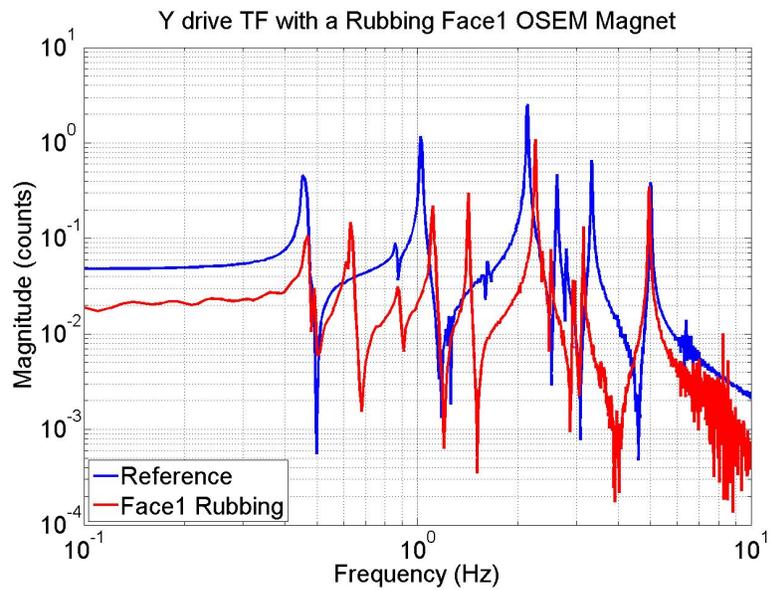


Figure 48: A Y-Y TF with a rubbing FACE1 OSEM magnet.

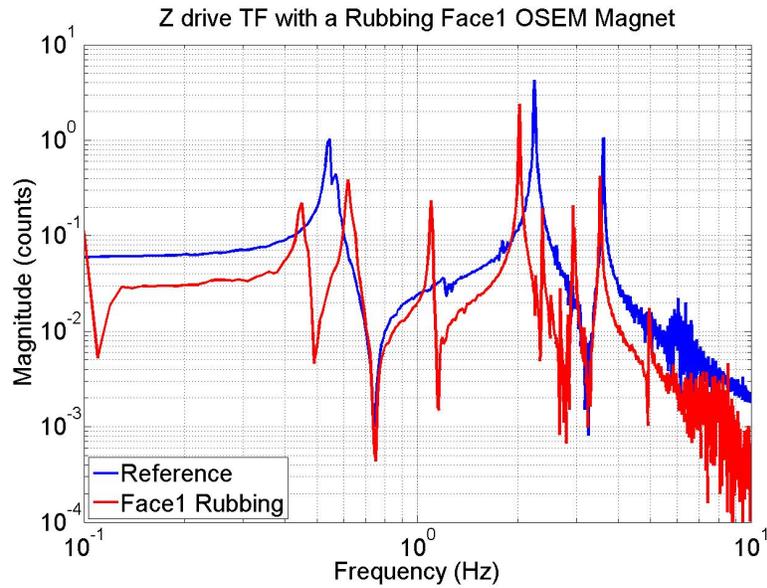


Figure 49: A Z-Z TF with a rubbing FACE1 OSEM magnet.

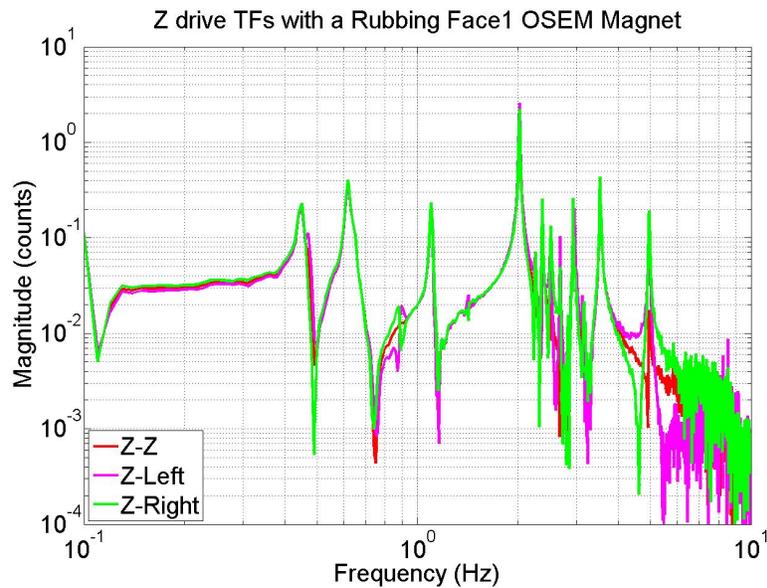


Figure 50: A Z-Z TF with a rubbing FACE1 OSEM magnet. The LEFT and RIGHT responses are mostly symmetric. Probably the magnet is not quite centered in the OSEM bore.

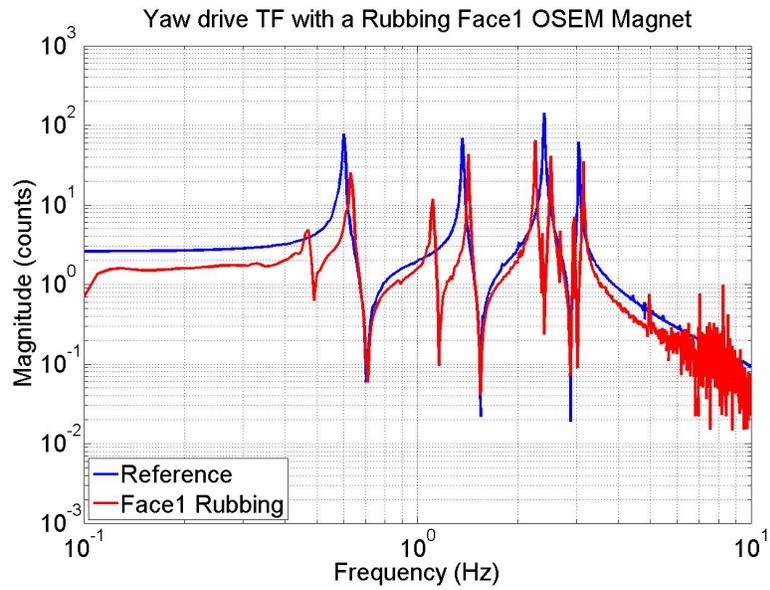


Figure 51: A YAW-YAW TF with a rubbing FACE1 OSEM magnet.

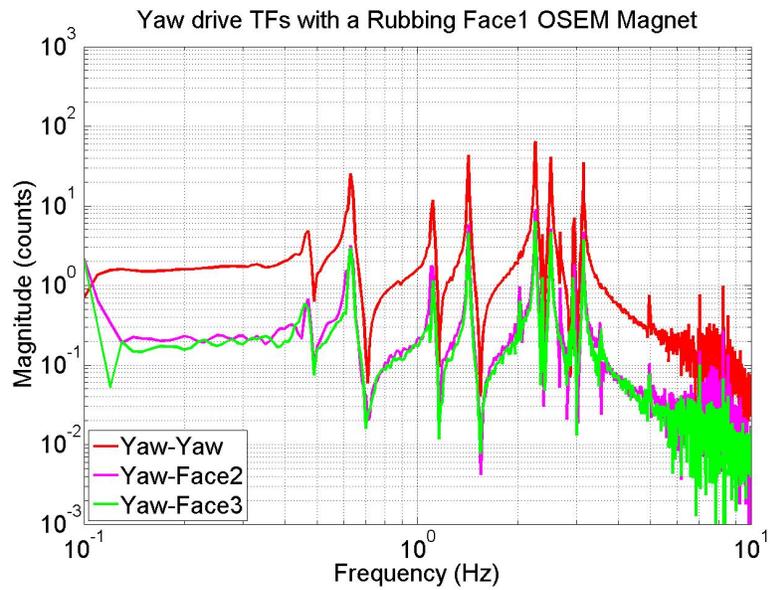


Figure 52: A YAW-YAW TF with a rubbing FACE1 OSEM magnet. A symmetric response from FACE2 and FACE3.

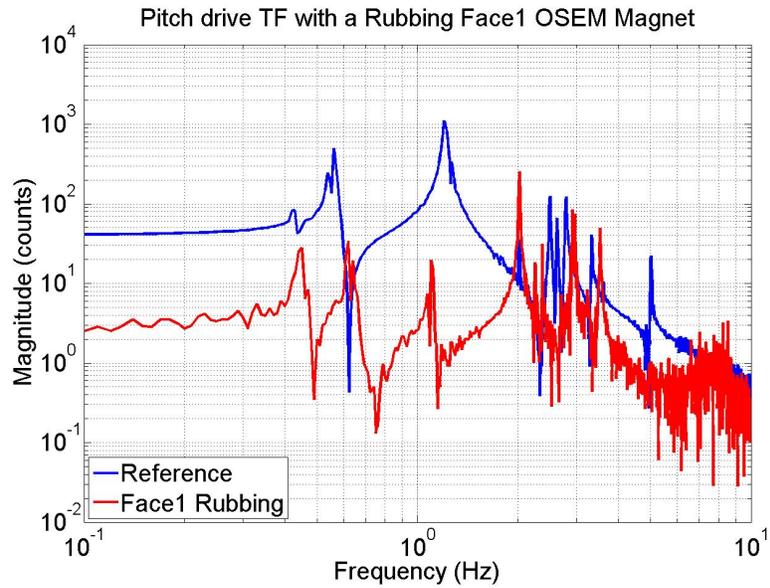


Figure 53: A PITCH-PITCH TF with a rubbing FACE1 OSEM magnet.

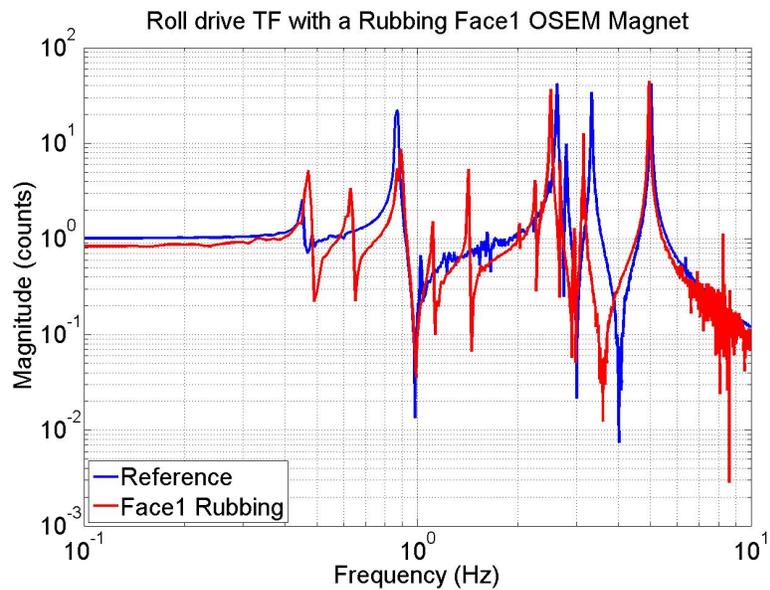


Figure 54: A ROLL-ROLL TF with a rubbing FACE1 OSEM magnet.

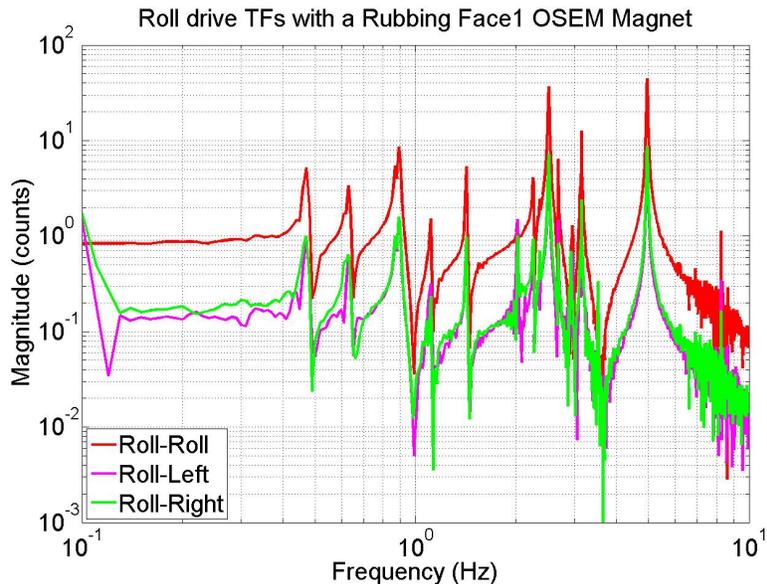


Figure 55: A ROLL-ROLL TF with a rubbing FACE1 OSEM magnet. A symmetric response from LEFT and RIGHT.

3.2.6 Top Mass Forward Left Top Stop Clamped

The next set of figures show transfer function measurements when the top mass was clamped with the forward top left earthquake stop.

The responses here are similar to the case of the rubbing magnet in that all responses have an overall drop in magnitude and resonances and coupling are altered. No DOF is spared. The main difference is the obvious asymmetry between LEFT and RIGHT and FACE2 and FACE3 that points to a left interference. Additionally, many of the measurements look very 'noisy'. All these effects are very characteristic of an external interference such as a locked stop.

Keep in mind that for the quad it is very easy to bottom out the LEFT and RIGHT OSEMs on top of one of the top mass blade spring stops. These stops must be turned down far enough to give sufficient adjustment range for LEFT and RIGHT. If the OSEM does bottom out, the effect on the transfer functions will be very much like what is shown in this section. At the same time one must be careful not to run the stop into the spring when turning it down or else you will end up iterating between this section and Section 3.2.4.

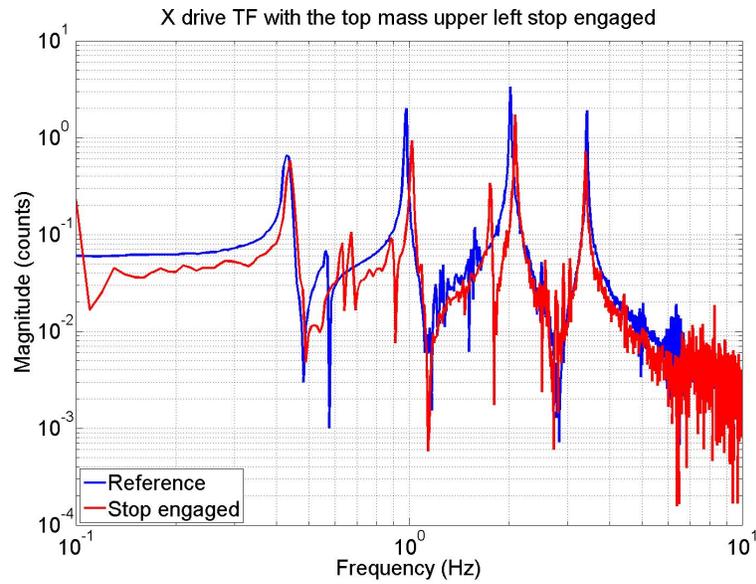


Figure 56: An X-X TF with the top mass forward top left stop clamped.

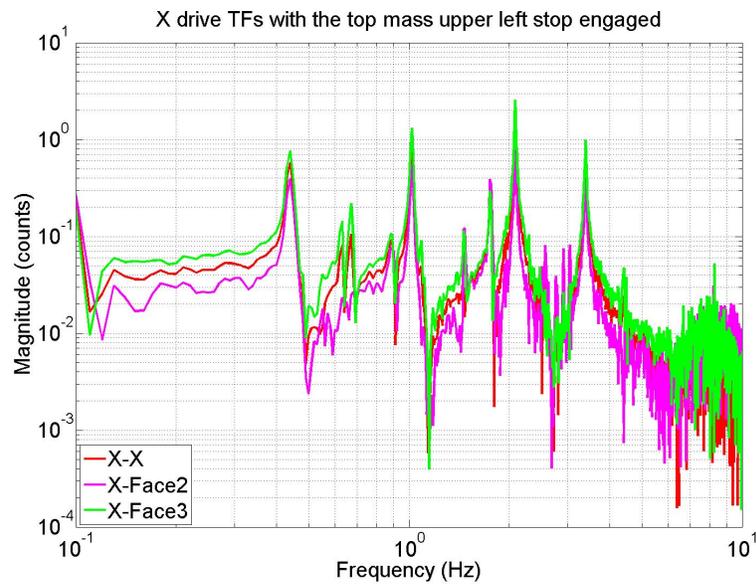


Figure 57: An X-X TF with the top mass forward top left stop clamped. Notice the asymmetry between FACE2 and FACE3.

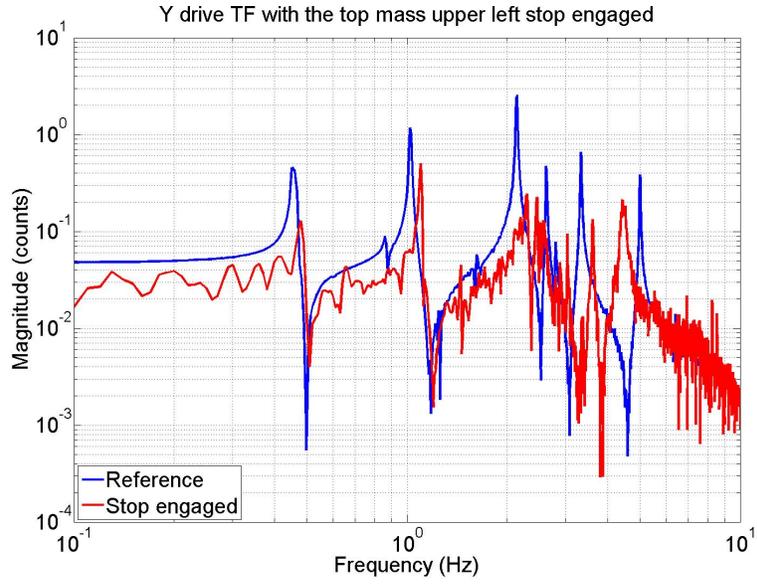


Figure 58: A Y-Y TF with the top mass forward top left stop clamped.

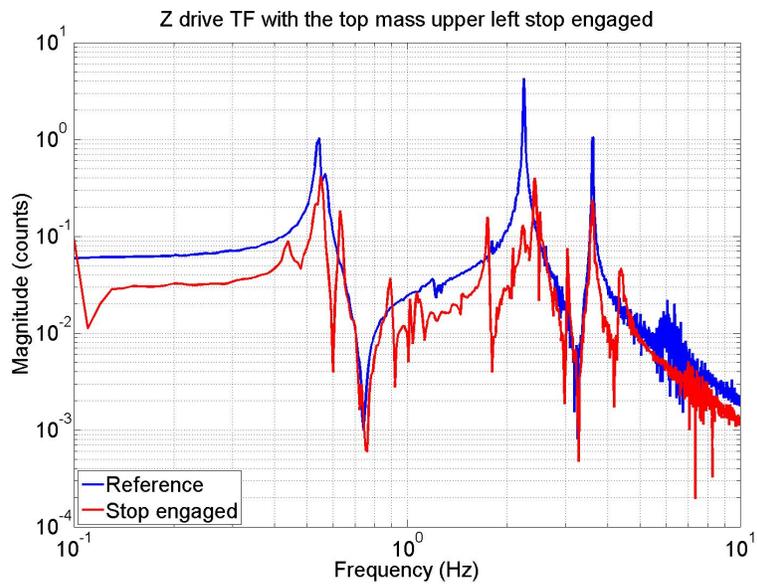


Figure 59: A Z-Z TF with the top mass forward top left stop clamped.

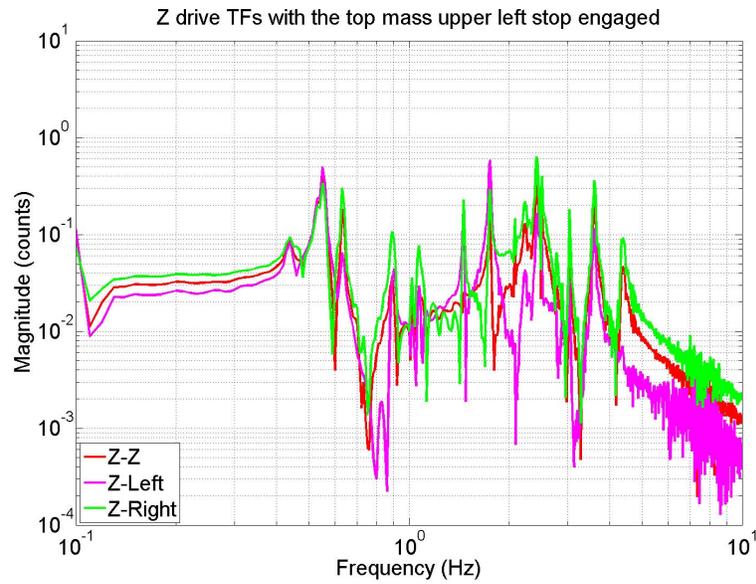


Figure 60: A Z-Z TF with the top mass forward top left stop clamped. Notice the asymmetry between LEFT and RIGHT.

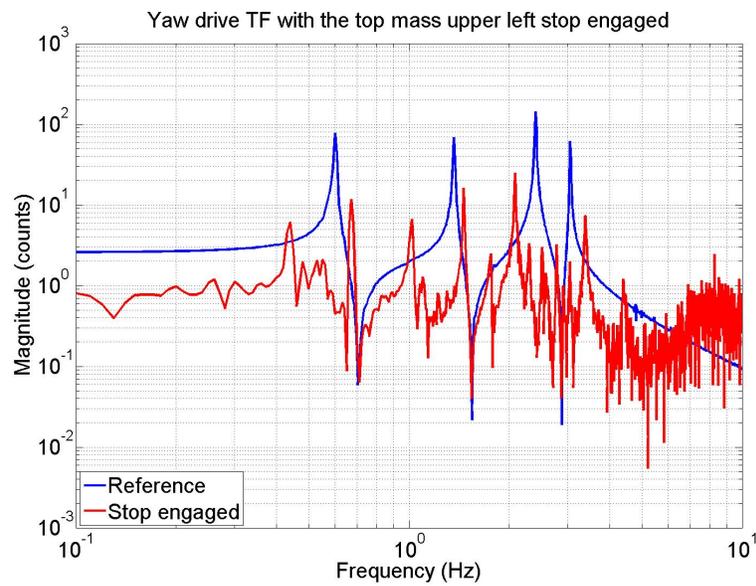


Figure 61: A YAW-YAW TF with the top mass forward top left stop clamped.

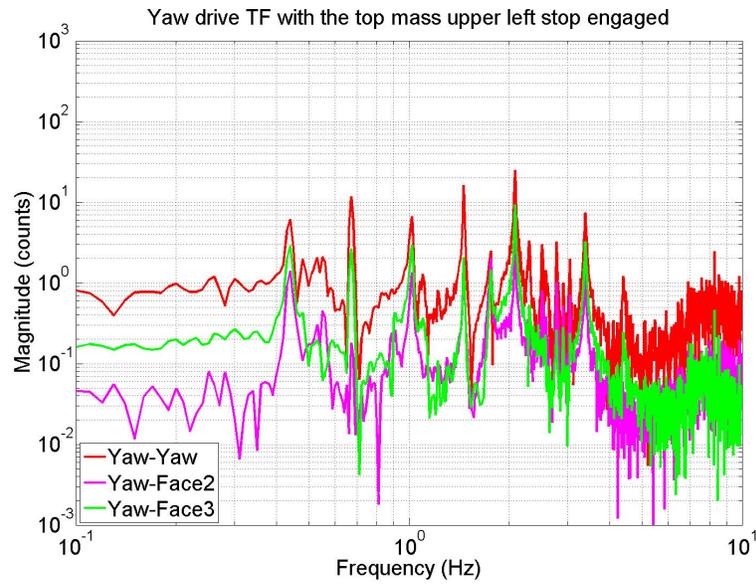


Figure 62: A YAW-YAW TF with the top mass forward top left stop clamped. Notice the asymmetry between FACE2 and FACE3.

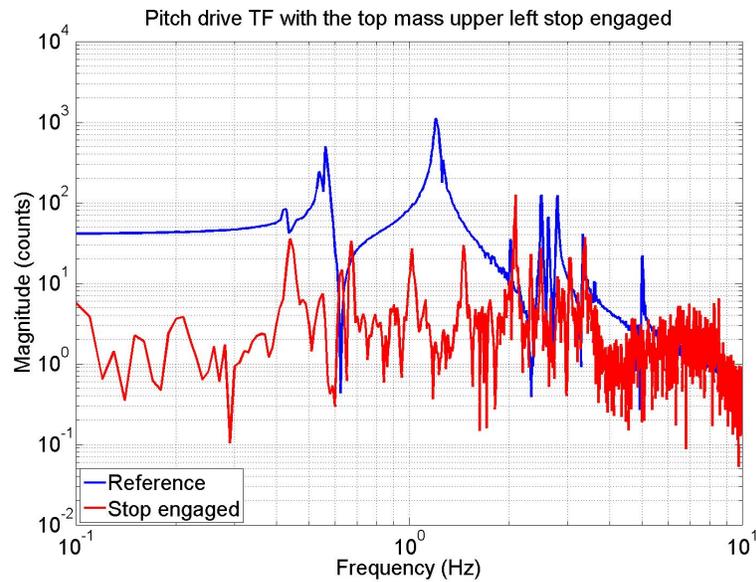


Figure 63: A PITCH-PITCH TF with the top mass forward top left stop clamped.

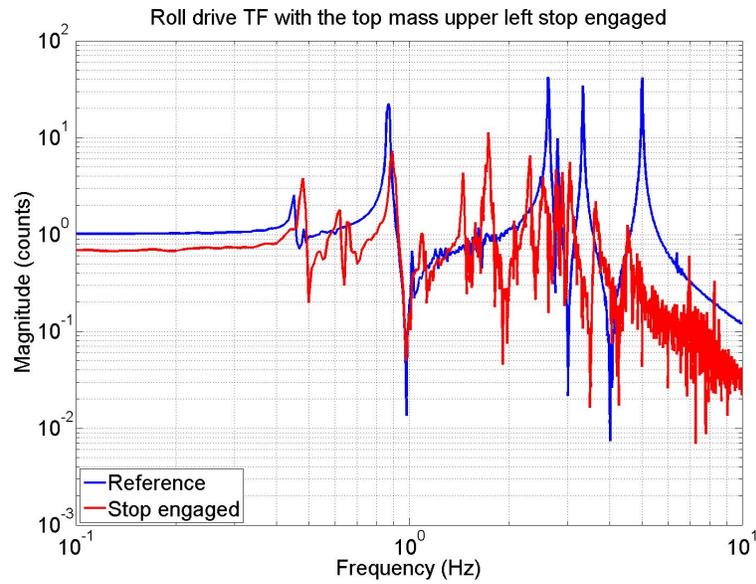


Figure 64: A ROLL-ROLL TF with the top mass forward top left stop clamped.

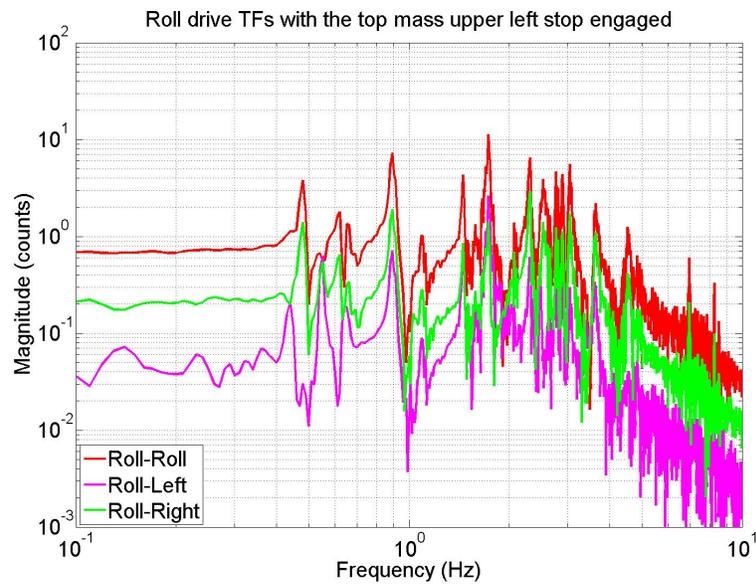


Figure 65: A ROLL-ROLL TF with the top mass forward top left stop clamped. Notice the asymmetry between LEFT and RIGHT.