

Advanced LIGO ITM/ETM suspension violin modes, operation and control T050267-01-K

K.A. Strain, G. Cagnoli

March 9, 2006

1 Introduction and references

The primary topic is the violin mode resonances in the quad suspensions. This document touches on several aspects of the interaction between the ETM and ITM suspensions and the rest of the interferometer. A toy global control loop is included to evaluate violin mode damping. It is not intended for other purposes such as development of whitening filters etc.

Reader beware: most of the graphs shown are from the MATLAB model, which assumes velocity damping, and for round fibres. The information necessary to extrapolate from these results to the case of ribbons with structural damping is given in the text.

Main references LIGO-T020205-01: description of the Mathematica suspension modelling toolkit see also

<http://www.ligo.caltech.edu/~mbarton/SUSmodels/asus4XLV5/index.html>

LIGO-T040110-01-K: includes statements of OSEM noise requirements and introduction to hybrid damping.

Yinglei Huang, PhD thesis, Syracuse 1996: describes the physics of violin modes including fibre stiffness.

Stefan Gossler, PhD thesis, Hannover, 2004: presents information that allows the passive of violin mode damping to be evaluated.

The present document completes the work outlined in LIGO-T050108-00-K and effectively replaces LIGO-T050107-01-K.

1.1 Revision -01 changes

There was an error in section 8 regarding thermal noise driven amplitude of the mirror motion at the violin mode resonances. The material in that section has been modified to remove the error.

2 Model

The model employed here is the Mathematica Quad suspension model and toolkit by Barton. A few parameters were changed to represent the latest quad suspension design. These changes, and other aspects of model validation are discussed in section 3 below.

To assist control system design, a MATLAB/Simulink model was built around the state-space matrices generated in, and then exported from the Mathematica toolkit. The Mathematica

toolkit allows state-space models to be exported with the damping changed to velocity damping. This change is required to allow straightforward use of MATLAB (structural damping is imperfectly represented in numerical ‘LTI’ models). The state-space model was exported with velocity damping set to match the calculated structural damping at the frequency of the first violin mode. This frequency was chosen as that is the mode whose Q is most significant in any assessment of stability (it can therefore be expected to be correct in this respect in the MATLAB generated plots). The Q s of the higher modes are wrong in the MATLAB model.

There were frequent numerical problems with the model and consistency checks were required, more significantly in MATLAB.

The Mathematica model includes a comprehensive representation of the 4 pendulum stages, including four fibres in the lowest stage each composed of 5 ‘beads’ connected by springs. The fibres therefore have mass, and stiffness is also included at the end sections. The model represents round fibres and the application of the results to ribbons is discussed in section 8. The model was modified to produce about 1% spread in the frequencies of the violin modes among the 4 fibres. This was achieved by adjusting the mass of the top two beads on each fibre. The degree of spread represents approximately the worst case for fibres that otherwise meet the requirements presented at the monolithic suspension PDR (see T050215-00-K and related documents).

In a model with only 5 beads per wire the mode frequencies of the higher-order modes are not expected to be correct. Results for higher modes should be viewed with some caution. Real ribbons will have some taper and so their modes will be non-trivial to calculate in any case.

3 Model validation and adjustment

Version 3.2 of the Mathematica model was employed.

The following changes were made before recalculating the model. Parameter $sl = 0.015$ – the value in the latest suspension models. The thermoelastic damping parameters $taufibre$ and $deltafibre$ were set to the correct values for silica, as was $phisilica$. The Young’s modulus $Y\beta$ of the final stage was set to the value for silica. Suspension d parameters for the last stage were set to the correct value (1 mm). The values of all of the material parameters used are those from *bench*, as distributed with the Mathematica model.

With these changes the damping at the test mass (TM) as a function of frequency was found to match that predicted by other models and the resulting suspension thermal noise prediction matches that from *bench* with the equivalent parameters, over the important frequency range. Given that the modes of the system appear at the correct frequencies, and the damping is correct as a function of frequency, it seems very likely that the modes will also have the correct Q and other properties.

The model simulates round fibres. The results are useful as they allow a direct comparison with many previous experiments and Huang’s thesis. As is shown below the match is very good and it is probably adequate to use the (simply calculated) dilution factor to model the change from fibres to ribbons.

The Q values resulting from the violin modes (of order 10^9 , for the first mode, up to $\sim 2 \times 10^{10}$ for the highest mode) are as expected, based on results of other models, and they are all above the experimental lower limits for comparable test systems (Cagnoli, Heptonstall, personal communications). The model does not contradict any prior knowledge of which we are aware.

The Mathematica warning *LinearSolve::luc* was frequently encountered during Q calculation and graph plotting. The graphs were smooth and the errors are thought to be minor.

Towards the end of the work the mass of top two beads on each fibre was adjusted to produce a spread in violin mode frequency. This was done by inserting multiplicative factors (0.99 etc.) to

the mass definition in the ‘kinetic’ matrix. The consequences of this change were not validated in detail.

4 Initial results – violin mode transfer functions, round fibres

Transfer functions from force applied at the penultimate mass (PM) or test mass, in the longitudinal direction (x in the model), to the x displacement of the beads on the fibres and to displacement of the PM and TM were calculated. By comparing the magnitudes of these near resonance it was possible to assess the effect of fibre stiffness. The observed ratio of fibre to mass motion was compared to the a model neglecting stiffness (which we call a ‘reaction’ model). The effect of the fibre stiffness was seen to enhance the motion of the mirror by about 50 times.

In the reaction model the effective mass of the fibre falls off with increasing mode number, whereas in the stiff-fibre model the stiffness of the end of the fibre should exaggerate the amplitude of the higher modes as seen at the masses. The Mathematica model was employed to find peak mode-height as a function of frequency. This was done by plotting the PM to TM transfer function near each resonance. The results were obtained using the ‘findQ’ utility function supplied in the toolkit. (Tip: this often works better if the initial guess for Q is slightly too high.) The slope of the resulting graph, which was straight on a log-log plot, was close to $1/f^{-3.5}$. Since the underlying PM-TM transfer function falls off as $1/f^4$, the progressively higher order modes are actually being enhanced by a larger factor at higher frequency. This can be understood as a result of the increasingly steeper gradient near the ends of the fibre in higher order modes. The stiffness of the end section leads to increased coupling to the mass for a given amplitude of excitation. With a 5-bead wire it is hard to be sure how well the model represents reality in this respect.

The same calculation repeated in MATLAB gives a steeper slope (near $1/f^4$, but not constant), due to the incorrect damping model, and numerical errors. It is only correct at 450 Hz where the results agree quite well. **It is important to remember that the higher order violin modes are significantly higher than indicated in the MATLAB results given below.** Note also that closed-loop MATLAB plots do not necessarily show the peaks of the modes, even though the frequency vector was made to include all pole frequencies in the open-loop system. The MATLAB and Mathematica results are compared in figure 1.

5 The MATLAB/Simulink model

This section is mainly of interest to someone who wants to use the model (i.e. very likely nobody). The model was constructed using MATLAB R14 SP3, which has only trivial changes from R14. The state-space system exported from Mathematica consists of the A matrix (with damping set to be correct at 450 Hz, as explained above), together with a B matrix representing 70 inputs. The inputs are the 6 support point coordinates, plus forces on each of the remaining 64 degrees of freedom (6×4 for the masses plus 2×20 for the beads – the latter being constrained to move in x-y planes). The C matrix provides 64 outputs in the same order as the inputs on the B matrix. The D matrix is the null matrix of the correct size.

A two-level Simulink diagram was made to assist investigation of the x -direction transfer functions. The inner part (figure 2) includes 5 degrees of freedom eddy-current damping (ECD) – i.e. excluding longitudinal, and the same 5 degrees of freedom of local control applied at the top mass. The dampers were set up approximately as described in T040110-01. This was to ensure reasonably correct behaviour of the longitudinal controllers. Nothing was learned beyond that stated in the previous document, and work still needs to be done in this area. The inner diagram

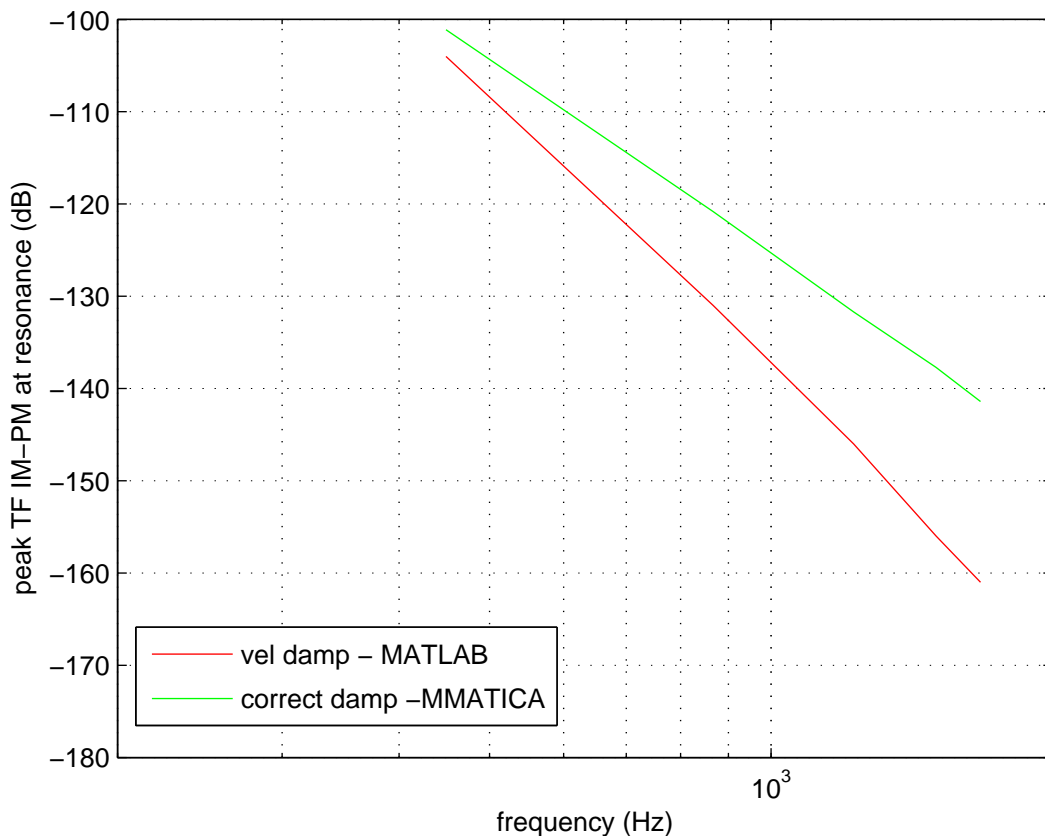


Figure 1: Plots of peak (violin mode) PM-TM transfer function amplitude for the 5 modes in each model. Only valid at resonance – 449.04, 867.22, 1225.94, 1500.86 and 1673.50 Hz. Note that the Mathematica results were checked for stability against small perturbations. The MATLAB curve is included only to show that the MATLAB results underestimate the peak values. This is likely to be both because the damping model in MATLAB is simplified and due to numerical imprecision. The violin modes are degenerate among the 4 fibres in these results.

also includes termination of the unused i/o etc. Its main purpose is to hide all this detail in normal use.

The outer diagram (figure 3) enables wiring of local and global control blocks, eddy current damping (ECD) and violin mode damping (passive or active). The simulations reported here were all of longitudinal transfer functions and their associated poles and zeros. The outer diagram also includes a representation of a filtered seismic noise spectrum, for use in evaluating feedback forces (approximated from the curves in the SEI conceptual design document).

The diagram holds a number of LTI objects (blocks), these are generated by running a script (generate_simulink) before opening the model. The script also defines a useful frequency vector that includes all of the open loop poles of the system as points (to help to show peaks on graphs correctly).

The results were obtained using the new ‘Control and Estimation Tools Manager’ to linearise the system and pass the result to the LTIviewer.

The LTI blocks contain poles and zeros defined by a number of s-plane convenience functions. Since these are mentioned at various places in the text below, they are stated here for the interested reader. The lowpass and highpass (freq, gain) are obvious. The ‘transdif’ and ‘transint’




Figure 2: The inner block of the simulink diagram.

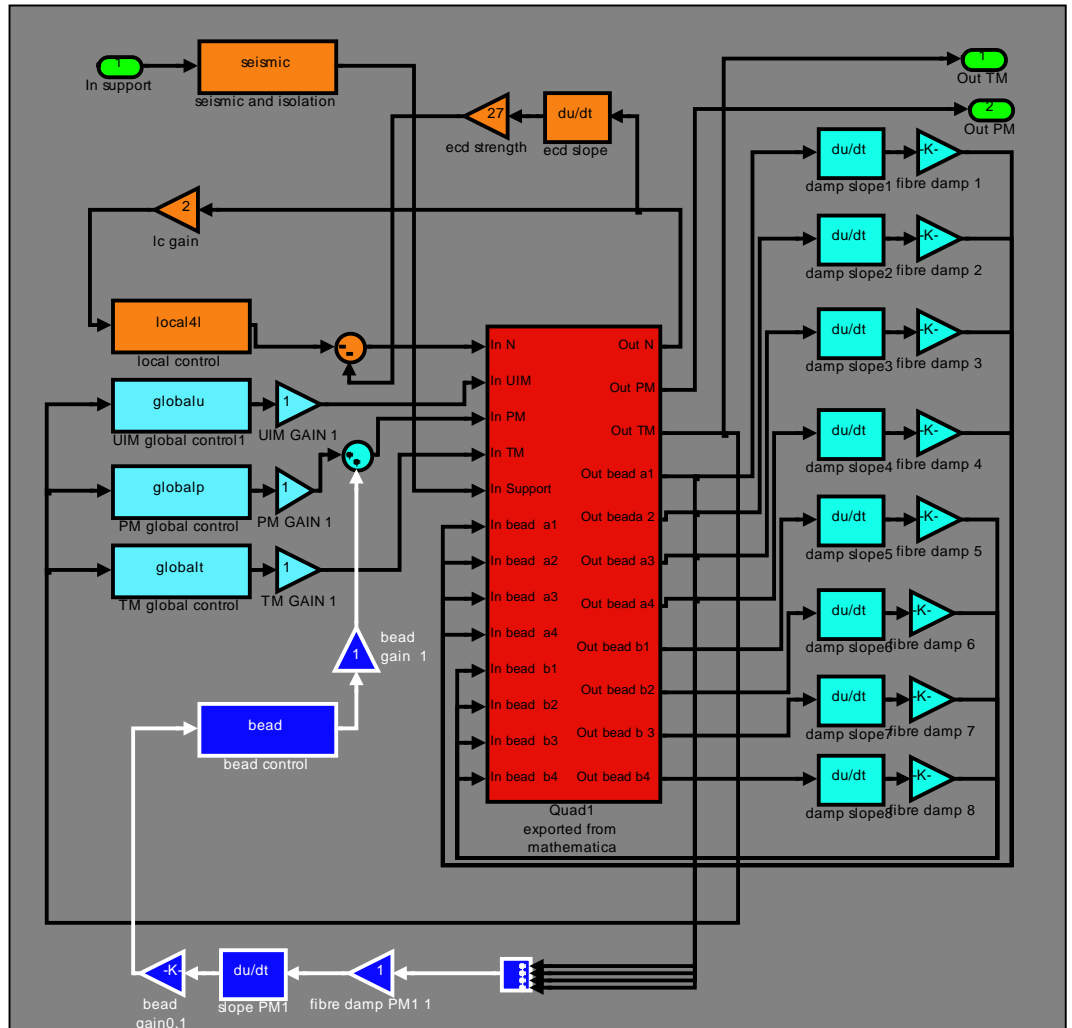


Figure 3: The simulink diagram for the system.

functions are simple real zero pole and real pole zero pairs (parameters are: start freq., stop freq., gain below first freq.). The other functions used: ‘sculte’ and ‘transintqd’, both produce a pair of complex poles and zeros (freq. poles, q poles, freq. zeros, q zeros, gain), but with different scaling of the gain. Frequencies are in hertz.

6 Violin mode damping

There are several approaches to violin mode damping. These are discussed along with relevant results from the model and elsewhere.

6.1 Goals and requirements

The primary driver, as concluded in T050108-00-K, is settling time after a disturbance. The maximum decay time-constant should be much less than an hour to allow recovery from even violent disturbance within a few hours. This implies a $Q \ll 10^7$, so a reasonable goal is $Q \sim 10^6$. This is much lower than is required for reliable global control (see section 7). Note that if an active solution to violin mode damping is chosen, it may be desirable to make it ‘adapt’ to the mode amplitude. In this way the damping time can be minimised for a given noise level in science mode. For the purpose of evaluation, the target Q is taken to be $\approx 10^6$, although a solution that just failed to meet this goal may still be considered acceptable if it has other merits (i.e. the requirement is not yet clear).

In general, additional damping leads to increased noise – either thermal or electronic for passive or active methods, respectively. The noise must be low enough to meet the ‘factor of 10’ requirement on technical noise. In addition noise could increase the violin mode amplitude above the thermal level (which we believe to be acceptable with respect to ADC range, etc.). The goal is, of course, to obtain a similar or lower level of excitation of the violin mode in science mode compared to the undamped case, and to achieve this in the shortest time after a disturbance.

The mechanisms and materials involved must be light enough to be supported by the suspension and/or its support structure, and not too complex to house in vacuum.

6.2 Passive methods

Two passive methods are considered. First is direct application of a damping compound (like the amorphous Teflon employed in GEO 600) to the ribbons. Second is the addition of tuned dampers on the PM.

Taking the latter first, this is theoretically possible, but practically very difficult. To obtain reasonable coupling requires rather high Q tuned dampers (at least 10^5). It is probably too difficult to maintain these in a suitably tuned condition (against temperature drifts etc.). Also about 12 such dampers would be required to cover the modes up to around 1.2 kHz (per suspension). The advantage of this method would be simplicity of wiring, and avoidance of risk to the ribbons.

Going back to the GEO 600 method, the main disadvantages are potential vacuum incompatibility, potentially excessive vertical-mode damping (see below), and possibly reduced rate of success in ribbon production. The vacuum compatibility issue has been looked at (monolithic review and referenced documents), and the increased failure rate was not significant in GEO (probably a few % of fibres were lost at the coating stage). That leaves the vertical damping to be evaluated. This method would require development of a semi-automated process to assist application of the desired distribution of damping material on to ribbons.

The advantages are that essentially all modes are damped at least to some extent (not true with the active systems), and simplicity of design/interfaces (no cables etc.)

6.2.1 Passive ribbon damping: desired damping vs. undesired damping

Teflon coating has been suggested as a solution to limit the Q of the violin modes of the suspension ribbons in Advanced LIGO. Although the experience from in GEO 600 was positive, the parameters and the performance of Advanced LIGO are quite different, so much so that a dedicated study is necessary. The results from GEO 600 are found in Stefan Gossler's PhD thesis, where they are analysed in detail. We have checked these results and applied them to the Advanced LIGO case.

From the GEO 600 tests the imaginary part of the Young's modulus of Teflon was obtained. The result is: $E_T \Phi_T = 3 \times 10^8$ Pa. Here we use E for Young's modulus and Φ for loss angle, subscript T for Teflon. From this figure it is possible to work out the loss angle of the vertical bounce mode of the last stage of the suspension. The total spring constant of a fibre or ribbon coated with Teflon is given by the following relation:

$$k_t = \frac{1}{\frac{1}{k_1} + \frac{1}{k_2 + k_c}},$$

where

$$k_1 = \frac{ES}{L - x} \quad (1)$$

$$k_2 = \frac{ES}{x} \quad (2)$$

$$k_c = \frac{E_T S_T}{x}, \quad (3)$$

are the spring constants of uncoated silica fibre, coated silica fibre and Teflon, respectively. Further, for our case, $E = 7 \times 10^{10}$ Pa, the cross-section $S = 1.28 \times 10^{-7}$ m², the fibre length $L = 0.6$ m, of which length x is coated with a coating of cross section $S = 2.4 \times 10^{-7}$ m², i.e. a layer thickness of about 100 μ m all round the fibre or ribbon.

This results in a loss angle associated with the spring k_t of $\approx 1.3 \times 10^{-3}$ if a 100 mm length of the ribbon is coated. This produces about 300 times the tolerable thermal noise around 10 Hz, and about 30 times the tolerable noise at around 100 Hz. The detail is shown in figure 4 where it should be remembered that the added noise should be, for example, $\leq 10^{-20}$ m/ $\sqrt{\text{Hz}}$ at 10 Hz.

Now, we concentrate our interest on the damping performance. Assuming a total coating length of 100 mm distributed as 80 mm across the middle of the fibre and 20 mm across the point at a quarter of the total fibre length, the dilution factor is given by the following graph at frequencies ranging from 100 Hz to 1600 Hz (figure 5).

The first three modes are at frequencies of 493 Hz, 985 Hz and 1477 Hz ¹ The effect of the additional mass of the coating has not been taken into account on the mode frequency calculation. The numbers on the y-axis have to be multiplied by $E_t/E = 3 \times 10^8 / 7 \times 10^{10} \approx 0.0043$ to have the actual Φ of the suspension in the longitudinal direction. For the first violin mode, the value of the loss angle is $1/Q \approx 2.2e - 8$. This value corresponds to a decay time of about 8 hours.

Although the effect on the vertical bouncing mode was seen to be huge, making the overall thermal noise figure over 10 times higher than the target value, the damping effect is completely inadequate.

¹These differ from the frequencies given by the approximate numerical model as presented elsewhere in this document.




Figure 4: Thermal noise arising from Teflon coating applied to 10 cm (upper, red curve) and 1 cm (middle, blue curve) of the ITM and ETM suspensions, compared to the total thermal noise without any added damping (lowest, magenta curve). The additional thermal noise is seen to be far beyond what is acceptable.

There are, of course, some degrees of freedom to explore, like the thickness of the coating. The damping effect depends more strongly on the coating thickness than does the loss angle of the vertical mode. Again, the $100\ \mu\text{m}$ thickness coating is already considered to be quite thick. Even a thickness increase of 10 times, rather unfeasible, does not produce a much better result.

The reason why the passive damping is much harder in Advanced LIGO than in GEO 600 is because the elastic energy density stored in the LIGO fibre is very much higher. The tolerance for damping of the bounce mode is also very much smaller.

6.2.2 Conclusion on passive damping

Due to the properties of the ribbons, and especially the very high tension upon them, passive damping does not reduce the damping time sufficiently (i.e. to less than of order days) without also introducing unacceptable amounts of thermal noise in the bounce mode.

6.3 Active methods

Two active methods are considered. The first employs direct sensing at the ribbons, with feedback to the PM, the second has sensing at the PM, with co-located feedback. (The third possibility involving feedback to the ribbons does not seem to bring any significant advantage except, perhaps, unconditional stability for very high order modes – modes that we are unlikely to need to damp in any case. The disadvantage of actuators acting on the ribbon is the lack of a conveniently located, quiet platform from which to mount them.)

Taking the latter method first, this has the disadvantage of requiring a sensor which is $\sim 10^5$




Figure 5: Effect of chosen distribution of Teflon expressed as the dilution factor as a function of frequency.

times better than direct sensing of the ribbons. It is also likely to be hard to make this method stable when the modes of the four ribbons are spread too widely. Consider that the modes may be spread by up to 1% in frequency. There is no point in the loop where the signals for the 4 ribbons become naturally separated for individual filtering, so a filter would need to be very complicated.

Direct sensing of the ribbons has the advantages of moderate sensing noise requirements and simplicity of loop design. The disadvantages are that 4 sensors are needed for each suspension and not all modes can be damped. The sensing should be near the top of the ribbon, perhaps 20% of the way down to allow damping at least 4 modes per ribbon - this is seen to be possible in the '5 bead' model.

To estimate the required sensor performance, we assume that the desired $Q \sim 10^6$. There are two ways in which sensor noise could limit system performance. First the residual *rms* motion of the ribbon could lead to ADC-range problems. To evaluate this we note that the motion will be dominated by the component within the resonance bandwidth. For the first mode that is $\sim 4 \times 10^{-4}$ Hz, when damped to the desired Q .

It is very likely that any practical filter has lower Q than the (damped) violin mode, and sensor noise will tend to drive the system at frequencies just outside the bandwidth of the mode. The advantage of 4 separate sensors is that each can be equipped with filters to damp one set of modes (those for one ribbon). The individual filters can be of much higher Q . The limit to the filter Q is probably set by thermal drift of the violin mode frequencies and so is probably of order 10^5 (thermal change in Young's modulus being the main effect). The sensor noise can, therefore, be suppressed by a factor of up to $\sim 10^5$ outside a band about $\sim 4 \times 10^{-3}$ Hz around the fundamental violin mode (for example).

The MATLAB model was used to test the effect of sensor noise in the circumstance described

above. First with simple $Q = 100$ filters covering each band of violin modes (as a very crude first step). In this case there is filtering of at least 170 dB (to the TM) except in a narrow band around the violin modes. To ensure that the resulting TM displacement noise does not exceed $8 \times 10^{-22} \text{m}\sqrt{\text{Hz}}$ the sensor noise should not exceed $\sim 3 \times 10^{-12} \text{m}\sqrt{\text{Hz}}$. In that case the violin mode amplitude is about 10^{-13}m *rms* on the ribbons and $\sim 10^{-18} \text{m}$ *rms* on the TM, this does not seem to be a problem.

The situation can be further improved by implementing separate higher Q filters for each ribbon. Increasing the filter Q to $\sim 10^4$ allows the servo gain to be reduced by another 40 dB outside the narrow bands of the filters. The sensor could then be about 100 times noisier, when the violin modes would still be below their free (thermal noise) level observed at the TM. Clearly though, maximum flexibility would be obtained by using the best possible sensor.

A sensing range of about 1 mm in two dimensions will be required to cover all possible suspension alignments within the structure. It would be wise to specify a design where the minimum clearance to the ribbon is about 1 mm.

There are two technologies appropriate to this sensing application: optical sensing and capacitive (electrostatic) sensing. A sensor design study will be required if the active solution is pursued.

A low Q bead feedback loop covering the first 3 modes of all 4 ribbons with non-degenerate mode frequencies is defined by

```
ag = 0; bg = 0; cg = 0; dg = 1; [a,b,c,d] =
transintqd(449,1e2,449,6,1);
[ag,bg,cg,dg]=series(ag,bg,cg,dg,a,b,c,d); [a,b,c,d] =
transintqd(867,1e2,867,3,1);
[ag,bg,cg,dg]=series(ag,bg,cg,dg,a,b,c,d); [a,b,c,d] =
transintqd(1226,1e2,1226,2,1);
[ag,bg,cg,dg]=series(ag,bg,cg,dg,a,b,c,d);
bead = ss(ag,bg,cg,dg);
```

(In examples below all the ‘series’ commands are omitted for clarity.) Note that overall phase lead is generated in a separate filter shown in the model diagram (see figure 3).

A high- Q bead feedback loop covering the first 4 modes of one ribbon is defined by the series connection of

```
ag = 0; bg = 0; cg = 0; dg = 1e-5;
[a,b,c,d] = transintqd(449.037,1e5,449.037,14,1);
[a,b,c,d] = transintqd(867.218,1e5,867.218,3,1);
[a,b,c,d] = transintqd(1225.941,1e5,1225.941,1.4,1);
[a,b,c,d] = transintqd(1500.864,1e5,1500.864,1,1);
```

Both of these work, the latter with much better sensing noise rejection as explained above. Clearly it would be a little more effort to determine the values for the poles in the latter filters.

The former loop was implemented for the results presented in the following sections. The resulting violin mode Q s are clustered around 10^6 as is indicated in figure 6.

6.3.1 Conclusion on active damping

7 Global control with violin mode damping

This section is intended to show that the violin modes, damped as above, create no great challenge to global control design. The design presented here is by no means optimised, but

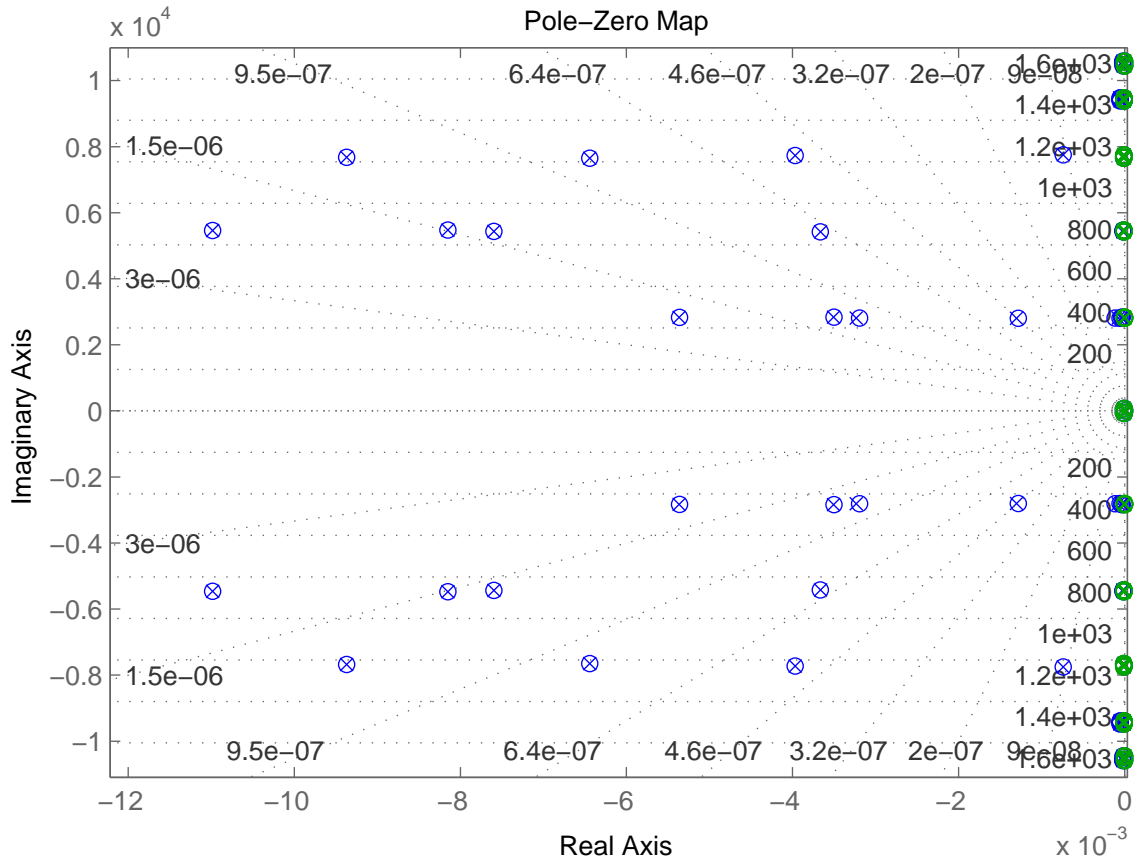


Figure 6: Poles associated with violin modes damped (blue, typical $Q \sim 10^6$) and undamped (green, typical $Q \sim 10^{10}$).

should be adequate to meet most requirements. The results replace those in T050107-00. ²

In all cases presented here at least 30 dB margin was included between the peak of a violin mode and unity gain, to permit enhanced transmission with ribbons (23 dB) as predicted in section 8, and at least 7 dB margin. As stated above, the simplified damping model and numerical errors in MATLAB suggest that this margin should increase with mode-frequency (by a few dB per mode to be safe). It is observed that this is not particularly problematic as only the first couple of modes play a part in servo design if low-pass filters are fitted as suggested.

In the TM and PM feedback paths the design employs a relatively simple, low-pass filter above the 100 Hz target UGP (arbitrary choice, above the optical spring frequency, but in line with other suggestions), and below the first set of ribbon modes. The filter offers a simple way to kill all violin modes without the need for complex tuning. The phase shift at 100 Hz is slightly worse than the many notch filters otherwise required (30 degrees phase margin). The filtering provided is much stronger than necessary and should even suffice if the violin mode damping is turned off. A weaker filter could be used with slightly better phase margin at the UGP. A similar arrangement suffices for PM feedback (here the crossover is set at 20 Hz when the phase margin is 30 degrees). ³

²Although it was originally intended that other aspects of global control be discussed in this document, work is underway within ISC to consider these other aspects. A separate document will follow to review the design of the actuators. This will be written when ISC requirements are clarified.

³The modelling suggested that under some circumstances a little filtering was also needed in the UIM path,

Performance of these paths is shown in figure 7, note that the violin modes are shown accurately on the plot (their peak values are plotted).

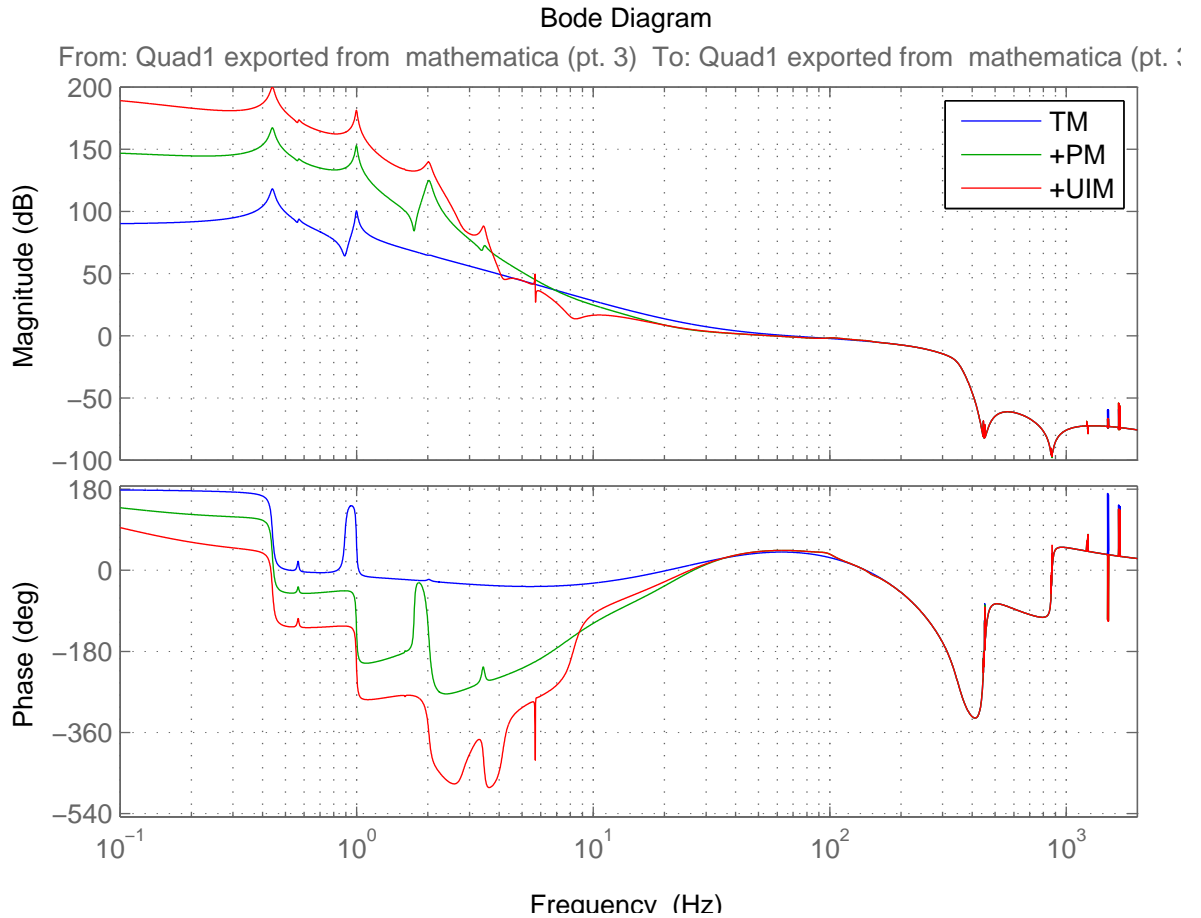


Figure 7: Open loop global control transfer function showing the effect of the 3 feedback paths with the violin mode damping on. The UIM path is of poor design especially around 8 Hz. The peaks of the violin modes are correctly shown to be negligible – their peak amplitudes are shown on the graph but that barely modifies the underlying curve.

7.1 TM feedback

The presence of the violin modes affects the response of the TM to forces in the relevant frequency range.

There is no difficulty arising from the spread of modes among the various fibres (although the number of notches becomes large with 8 ribbons per interferometer – two ETMs).

It is, however, worth considering how many such modes require to be notched and whether there could be a simpler approach. The simplest path seems to be to include a low-pass filter between the unity gain frequency and the first violin mode. This filter reduces the phase margin, but probably by an acceptable amount. Lower Q notches can then be employed at all relevant modes. Either approach should work, a toy low-pass model is included.

The ‘globalt’ block is the series connection of the following:-

but this seems incredible, perhaps numeric precision is stretched too far. Anyway it is trivial to add if really needed.

```

ag = 0; bg = 0; cg = 0; dg = -12e5;
[a,b,c,d] = transdif(30,190,1);
[a,b,c,d] = transdif(30,190,1);
[a,b,c,d] = transint(3,30,10);
[a,b,c,d] = sculte(270,0.7,447,20,1);
[a,b,c,d] = sculte(310,1.4,450,20,1);
[a,b,c,d] = sculte(350,5,867,50,1);

```

As can be seen in figure 8, the notching may even allow stable operation with the (active) violin mode damping turned off, as the necessary 30 dB margin is just available at the fundamental mode.

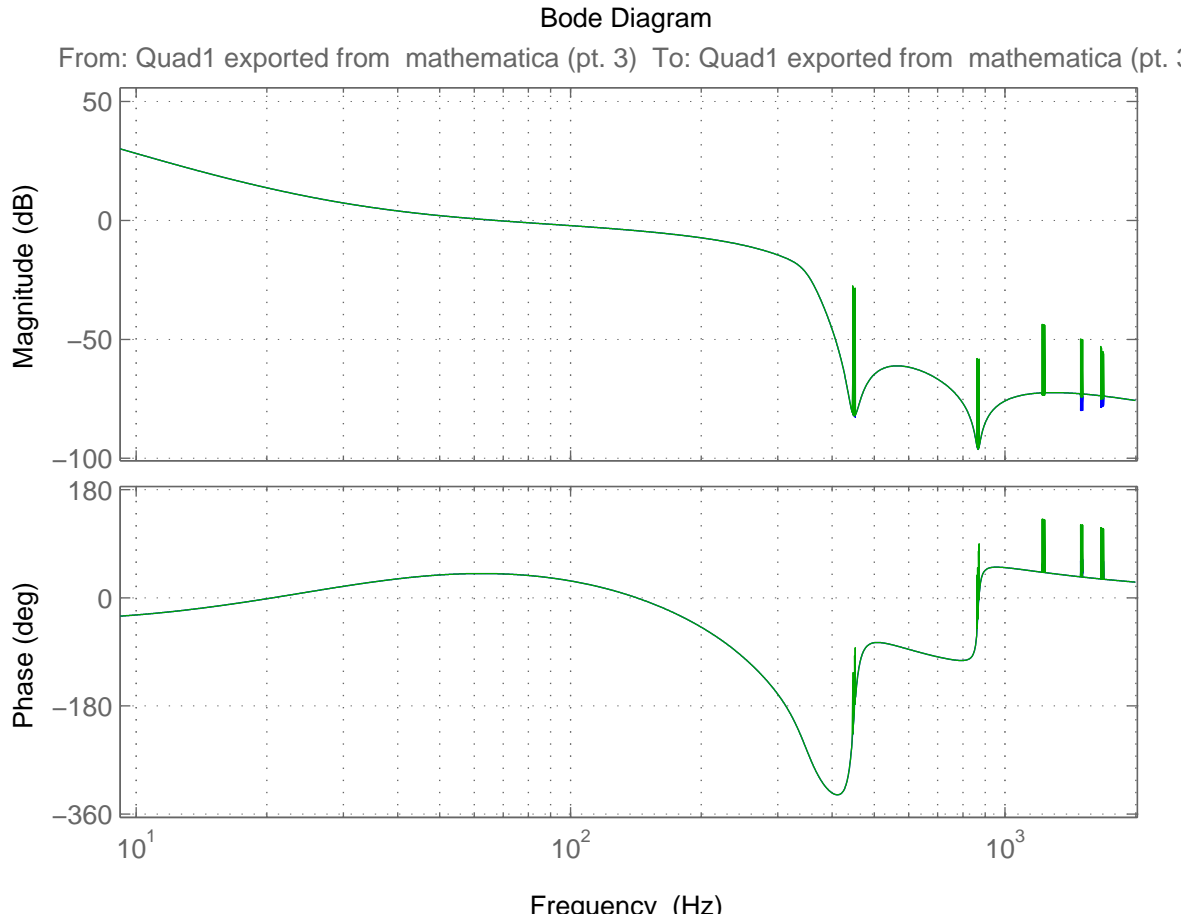


Figure 8: Open loop global control transfer function showing only the TM feedback with violin mode damping off (green). When the damping is on the peaks are reduced by around 60 dB (essentially invisible blue curve, but see figure 7 above).

7.2 PM feedback

The presence of the violin modes has a strong effect on the PM feedback path, especially if its crossover to TM feedback is above about 20 Hz. The main problem is the fundamental mode. The need to have relatively strong phase lead at the crossover tends to ‘boost’ the violin mode unless a low-pass filter is included and/or yet more notches which of course may be shared with the TM path if that is desired.

The ‘globalp’ block is the series connection of the following:-

```

ag = 0; bg = 0; cg = 0; dg = -14e7;
[a,b,c,d] = transdif(4,40,1);
[a,b,c,d] = transint(0.1,20,200);
[a,b,c,d] = transintqd(2,10,2,1,1);
[a,b,c,d] = transdif(10,30,1);
[a,b,c,d] = sculte(100,5,447,100,1);
[a,b,c,d] = sculte(150,10,451,100,1);
[a,b,c,d] = sculte(200,20,867,100,1);

```

7.3 UIM feedback

This loop requires to be improved, but plays little role in the study of violin modes.

8 From fibres to ribbons

The model represents round fibres, the actual suspensions should have ribbons, the main difference is a change in the stiffness of the end sections. This can produce a difference in the peak transmissibility at the violin modes. Here are a few comments on the link with theory.

In the notation adopted by Huang, the energy of the n^{th} mode of a fibre is

$$W \approx \pi^2 m f_1^2 (n A_n)^2,$$

where A_n is the amplitude of the mode, m the fibre mass f_1 the fundamental frequency, and the approximation neglects elastic energy in favour of the greater tension energy (larger by the dilution factor).

The corresponding motion of the TM is predicted to be

$$\frac{A_M}{A_n} = n\pi \sqrt{\frac{EI}{TL^2}},$$

where A_M is the TM motion, and the factor with the square root is half of the dilution factor. The dilution factor is, in turn, the ratio of twice the bending length ($\sqrt{EI/T}$) to the total length L . Here Young's modulus E is $7 \times 10^{10} \text{N/m}^2$ and tension T is about 100 N. For the round fibres in the model $I = \pi r^4/4 \approx 1.3 \times 10^{-15} \text{m}^4$. For the ribbons the corresponding figure is $ab^3/12 \approx 10^{-16} \text{m}^4$. Under the appropriate circumstances the ribbon modes stick up ≈ 23 dB more than the fibre modes.

Here 'appropriate circumstances' means where the 'driving force' is applied to a mass, when the ribbons take whatever energy they need without significantly affecting the motion of the mass. Application of a driving force on to the ribbons themselves does not lead to any change from the reaction model because the ribbons appear to be stiffer due to the tension (so are harder to drive in the first place). As an example, the thermal noise force in fibres or ribbons, ideal or otherwise, leads to – approximately – the same TM motion for both cases.

For our parameters, the observed factor of ~ 50 increase in the fundamental mode transmission seen in Mathematica due to fibre stiffness compares very well to a prediction of about 63 times from Huang's result. The lack of equally good agreement for the higher modes is expected given the limitations of the 5-bead model. In the worst case the higher modes would stick up a little more than predicted here.

9 Summary

Due to the high Q of the violin modes in the ITM and ETM quad suspensions, it will be necessary to damp the modes to some extent to obtain sufficiently rapid recovery from an excitation. The goal would be to damp the lower-frequency modes to $Q \sim 10^6$. With this Q , or even much higher, the effects on global control are quite mild, requiring only simple filtering to suppress completely.

Active damping of the violin modes should work, but a design study for a sensor is required to find the best approach. Passive damping does not seem to offer a practical solution for damping of the most important, lower order modes.