LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
– LIGO –

LIGO Laboratory / LIGO Scientific Collaboration

LIGO-T020034-01-D *ADVANCED LIGO* 19 March 2002

Low-frequency Cutoff for Advanced LIGO

P Fritschel, D Shoemaker, D Coyne

Distribution of this draft:

California Institute of Technology
LIGO Project - MS 51-33
Pasadena CA 91125
Phone (818) 395-2129
Fax (818) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project - NW17-161
Cambridge, MA 01239
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

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

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

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

1 INTRODUCTION

This note addresses issues involved with the choice and impact of the low-frequency cutoff for the Advanced LIGO interferometers. To date, the low-frequency cutoff has been specified as 10 Hz. We begin by reviewing the potential impact of variations around this frequency on astrophysical source detection. Within the interferometer subsystems, the targeted cutoff most strongly affects the test mass suspension design. Specifically, the challenge is in achieving a sub-10 Hz eigenfrequency for the vertical mode of the last stage of the suspension (though this may not be strictly required to achieve a 10 Hz cutoff frequency; see section 4 below). We examine the trade-offs involved in this design issue.

2 IMPACT ON ASTROPHYSICAL SOURCE DETECTION

The effect of cutoff frequency on the detection range of neutron star binary inspirals and on the sensitivity to a stochastic background is readily calculated with the BENCH program. The result is shown in Figure 1, which indicates that neither source is significantly affected for cutoff frequencies in the range of 10-15 Hz.

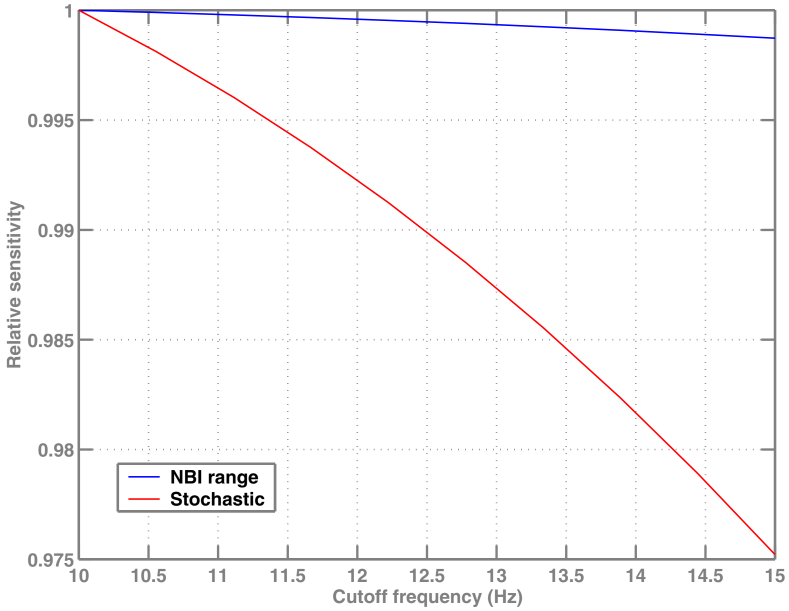

Figure 1. The sensitivity for NBI and stochastic background detection (assuming a frequency independent $\Omega(f)$), relative to the sensitivity level at 10 Hz, for the baseline Advanced LIGO

The gravitational wave theorist community was asked to comment on the relative merits of low-frequency sensitivity, and estimate the strength and likelihood of other potential signals in the ~10-20 Hz band. Input was received from K Thorne, P Brady, V Kalogera, J Centrella, S Hughes, B Schutz, E Flannagan and J Romano. The general conclusion is that there is insufficient information to rule out or to 'rule in' any sources, and the likelihood that a type of source would be missed due to, e.g., an increase in a low-frequency cutoff of 10 Hz to 13 Hz is impossible to evaluate. No anticipated source classes would be excluded by excluding the 10-13 Hz band; none would be included by pushing the band down to 8 Hz. In general, there are no estimates of rates for these sources (for pulsars, there are counts of nearby visible sources). If one is constrained to looking above and below a significant suspension resonance, only coherent (fixed-frequency) sources near the narrow resonance would be significantly impacted; other source signals could be ‘knitted together’ across the inaccessible frequency band.

Core collapse of massive stars. The most up-to-date review article on the various signals resulting from core collapse of massive stars is from Fryer, Holz, and Hughes (astro-pf/0106113 v2). One scenario that may lead to low-frequency signals involves the fragmentation of a fairly large object (~300 M_S). In the model of Fryer et alia, both strain and frequency are boosted if the core splits into two pieces, which then fall into a Keplerian orbit, conserving angular momentum. If this instability occurs and the pieces orbit coherently for ~10 cycles, the waves may be detectable at redshifts $z \sim 5$: for these parameters the model predicts the gravitational wave strain to be above Advanced LIGO’s strain curve in the band 13-70 Hz (with $h_{rms} \approx 4 \times 10^{-22}$). The signal spectrum has a ‘knee’ at above 30 Hz–below 30 Hz the strain amplitude is fairly constant with frequency, but it drops off quickly above 30 Hz. Given this, a variation in the interferometer cutoff frequency around 10 Hz would make no difference to the source detectability.

Low frequency pulsars. There has been little work on end-of-life pulsar gw emission since a collection of articles from Brady et al. (PRD, 57, 2102). The best model for the pattern of evolution (and gw emission) as a pulsar ages is a spin-down from a rotation rate of several hundred Hz to the 1-10 Hz region, where it ‘sticks’. The upper limit of gw emission is calculated from the measured spin-down rate, and assuming that all of the decrease in rotational energy is due to gw radiation; our best information indicates that most of the energy is lost as electromagnetic radiation. Figure 2 from Schutz is based on the Brady et alia paper, and represents our best estimates of the upper limits from electromagnetically-visible pulsars. Most of the sources actually lie just under 10 Hz, a result which could argue for positioning an obscuring interferometer resonance just above 10 Hz.

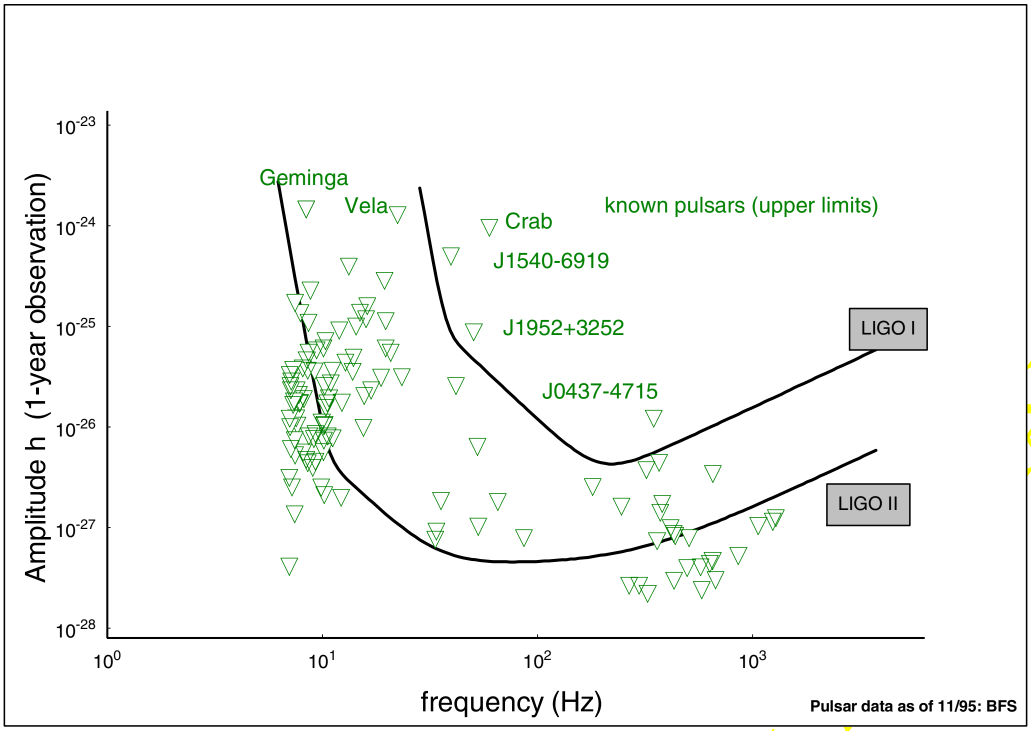

Figure 2. Pulsar CW sources (Schutz; ‘LIGO II’ curve slightly modified to resemble Advanced LIGO).

Massive binary inspirals. In the region of 10 Hz, inspirals of ~1000M_s BH-BH binaries are a potential source of signals for Advanced LIGO. A lower cutoff frequency means that more massive systems could be detected, with an inverse relationship: $M_{\text{binary}} \sim 2000 M_{\text{sun}}(10\ \text{Hz}/f_c)$. The implications of this have been addressed by Thorne (LIGO-T000146-00) and Kalogera. They conclude that 20-30% changes in the cutoff frequency around 10 Hz would not be significant for this source. As Kalogera states, ‘detecting systems with [individual] BH masses of 800, 1000, or 1200 M_{sun} does not advance our physical understanding in any different way, because the realization that a 800M_{sun} BH exists has the same implications that the realization of a 1000M_{sun} BH has.’ Thorne and Kalogera estimated how the search volume and event rate could scale with the low-frequency cutoff (with Kalogera making some refinements to Thorne’s initial estimates). The bottom line is that moving from 10 Hz to 13 Hz, e.g., might reduce the rate by a factor of roughly 2.

3 INTERFEROMETER DESIGN ISSUES

Here we examine the limitations to low frequency performance, within the basic advanced LIGO interferometer design concepts. The fixed components of the design are:

- active seismic isolation system, with performance as described in LIGO-E990303-03-D.
- quadruple test mass suspensions, with fused silica fibers for the last stage; operated at room temperature
- interferometer configuration: signal- and power-recycled, Fabry-Perot arm Michelson; mirror transmissions fixed and optimized for inspiral detection, but signal recycling mirror phase is tunable
- sapphire test masses; 40 kg mass
- laser power: up to 125 W at interferometer input

3.1. Limiting noise levels

The limiting low-frequency noise levels for the advanced LIGO design are shown in Figure 3. Given the basic design choices listed above, the limiting noise sources are those we have little or no ability to reduce further:

- *Ground noise.* Seismic mounting platform is assumed to have a displacement noise level of $2\times10^{-13}\,\text{m}/\sqrt{\text{Hz}}$ at 10 Hz (see SEI DRD, LIGO-E990303-03-D). Isolation of the quad pendulum is calculated (with damping small enough not to limit isolation) for last stage pendulum lengths of 60 cm and 90 cm (line thickness in the figure indicates the resulting noise spread).
- *Gravitational gradient background.* Yellow band: estimate from Hughes and Thorne, *Phys. Rev. D,* **58**, 122002; Grey band: estimate from R. Schofield, using H&T formulation, but with measured Hanford ground noise, anisotropy ratios and propagation velocities (ground noise measured at a relatively noisy time).
- *Test mass thermal noise:* Test mass aspect ratio and beam size are optimized, as described in the LIGO Systems Design, LIGO-T010075-00.
- *Quantum noise:* Laser power is 30 W into interferometer, signal recycling mirror has no detuning. At this point, quantum noise is smaller than or equal to thermal noise at all fre-

quencies below ~50 Hz.

- *Suspension thermal noise:* Horizontal thermal noise range: lower edge of the range is for circular fibers, radius chosen to minimize damping (r = 400 micron); upper end of the range is for the baseline ribbon design, 1.15×0.115 mm cross section. In each case the fiber length is 60 cm, and a surface loss of 180 micron dissipation depth is included.

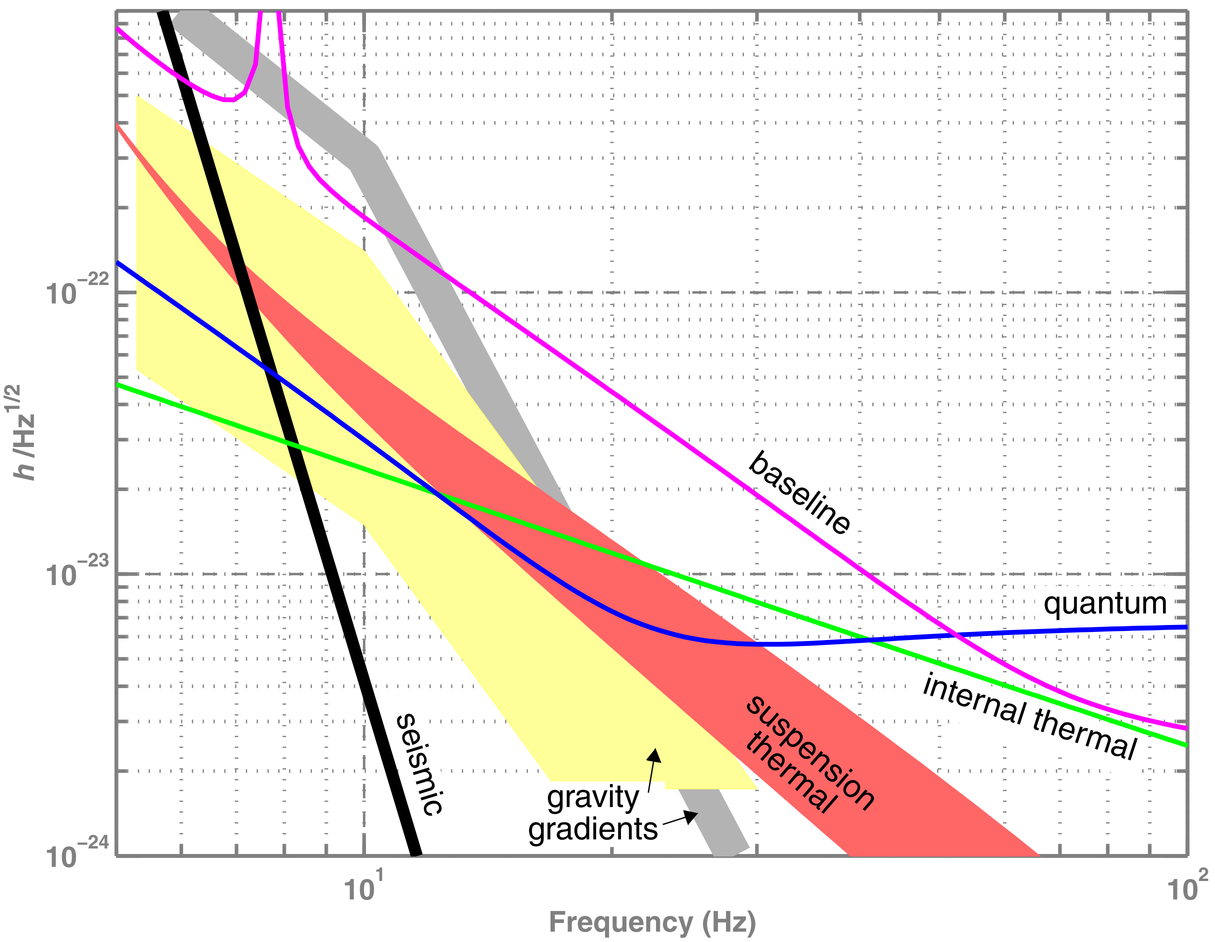

Figure 3. Levels of the limiting noise sources at low frequencies. Also shown is the total strain noise for the baseline interferometer design, optimized for binary inspiral detection. See text for specifics on the individual curves. Parameters for the ‘baseline’ curve can be found in the Advanced LIGO Systems Requirement Document, LIGO-T010075. The baseline design is optimized for binary inspiral detection, which leaves the ~10-50 Hz band dominated by quantum noise; for better low-frequency sensitivity the input laser power can be reduced and the signal recycling retuned.

- Vertical thermal and seismic noise are not included in Figure 3, as their levels depend more strongly on remaining variables in the suspension design. In particular, vertical thermal noise is significant in a small band around the highest-frequency vertical mode of the suspension (involving the stretching of the test mass suspension fibers), and the specific eigenfrequency of this mode involves various trade-offs, as discussed below. In addition, the degree to which vertical motion couples to the test mass readout can in principle be minimized with careful engineering (even conceiving of compensating the 3×10^{-4} coupling due to earth curvature with a designed-in mechanical coupling).

3.2. Suspension design trade-offs

It clearly would be advantageous to have the highest vertical mode lie outside (below) the gravitational wave band (ideally all mechanical resonances in the interferometer would be made to lie outside the gravitational wave band!), however there are other performance and practical issues to consider with the parameters the affect this mode frequency. The trade-offs involved with the three parameters at our disposal – mass of the penultimate stage, fiber length, fiber cross section – are indicated in Figure 4.

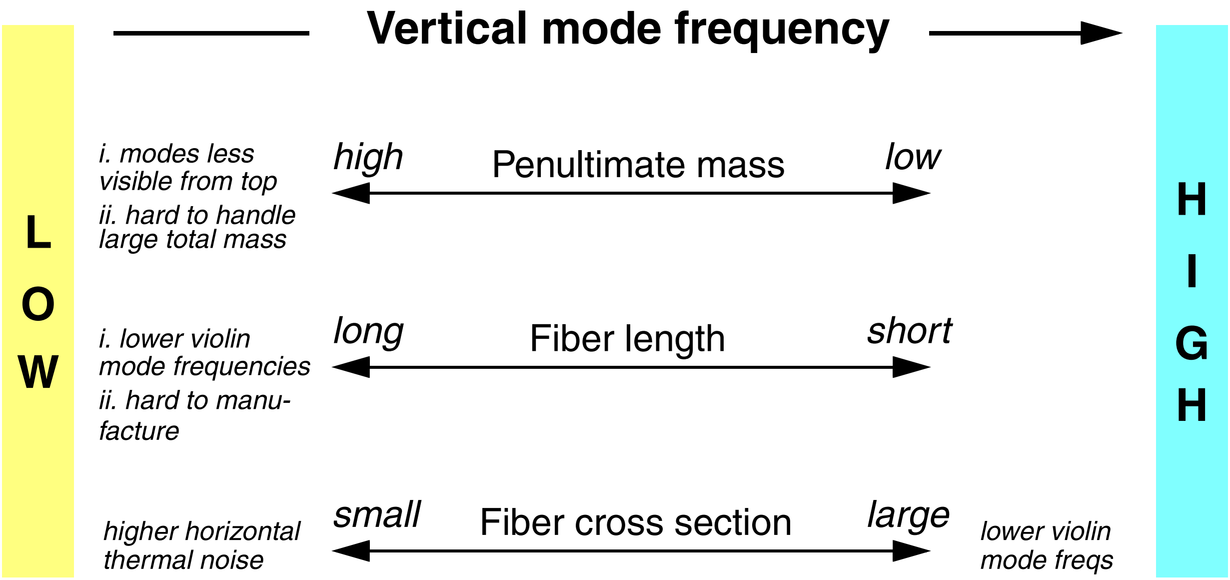

Figure 4. Trade-offs involved with the parameters affecting the highest vertical mode frequency in the suspension. For each parameter, any *disadvantages* of pushing it to one extreme or the other are listed.

To further clarify the disadvantages given in Figure 4:

- *Modes less visible from the top:* If the penultimate mass were made much larger than the test mass, motion of the test mass would hardly show up at the top of the suspension, and the local damping would not be entirely effective. Certainly the penultimate mass can be made larger than the test mass and still achieve adequate damping, but this general tendency should be kept in mind as the penultimate mass is increased.
- *Hard to manufacture and handle large total mass:* This encompasses several concerns: difficulty in finding vendors to supply and work with dense materials; difficulty in assembling and installing suspensions having larger masses; load limitations of the seismic platform (before a significant change to the seismic design would need to be made).
- *Lower violin mode frequencies:* The fundamental violin mode can cause an instability for the global control feedback to the test masses; dealing with these resonances (i.e., by filtering them out) becomes more difficult the lower their frequencies.
- *Harder to manufacture (longer fibers):* Producing uniform, strong fibers of the right cross-

section is more difficult the longer they are.

- *Higher horizontal thermal noise:* This refers to the ‘nonlinear’ thermo-elastic damping in the stressed fiber; there is a stress level at which the temperature dependence of the elastic modulus cancels the thermoelastic damping arising from thermal expansion. This occurs at a relatively small fraction (~10%) of the breaking stress; thus a higher stress design actually suffers from higher thermal noise of the pendulum longitudinal mode.

The vertical (f_{v}) and fundamental violin mode (f_1) frequencies, and the nonlinear thermo-elastic damping are linked by the formulas:

$$f_{\mathrm{v}} \approx \frac{1}{2\pi}\sqrt{\frac{gY\mu}{L\sigma}} \qquad f_1 = \frac{1}{2L}\sqrt{\frac{\sigma}{\rho}} \qquad \sigma_{\mathrm{opt}} = \frac{\alpha Y}{\beta} \tag{1}$$

where L is the fiber length, σ is the stress in the fiber, g is the gravitational acceleration, and $\mu = (m_{\mathrm{tm}} + m_{\mathrm{pm}})/m_{\mathrm{pm}}$, with m_{tm} and m_{pm} the masses of the test mass and penultimate mass; α is the thermal expansion coefficient, and Y and β are the Young’s modulus and its temperature derivative for the fused silica fiber material; σ_{opt} is the stress at which the thermo-elastic damping is zero (where the linear thermal expansion is cancelled by the Young’s modulus change). There is clearly a competition between choosing parameters to simultaneously achieve a low vertical mode frequency, a high violin mode fundamental, and low damping. For example, take the upper and lower edges of the suspension horizontal thermal noise band shown in Figure 3, and assume the penultimate mass is twice the test mass: for the thick circular fiber (lower edge), with constant diameter along its length, the vertical eigenfrequency would be about 15 Hz; for the ribbon (upper edge), it lies just below 8 Hz.

3.3. Variable fiber cross section

Another option, recently proposed by P Willems, is to make the last stage suspension fiber have a variable cross section: at the ends, the cross sectional area is made relatively large, such that the linear thermal expansion is cancelled by the Young’s modulus temperature dependence to null out thermo-elastic damping in the regions that affect the pendulum thermal noise; the middle of the fiber is made with a small area to get as large a stress as possible, thereby lowering the vertical mode frequency and increasing the violin mode frequencies. This design approach could be applied to either circular or ribbon fibers.

Willems has calculated the thermal noise associated with a 60 cm long circular fiber, where the 10 cm length at each end is thick (767 micron diameter) and the middle 40 cm is thin (380 micron diameter); see LIGO-T020003. The vertical eigenfrequency is just under 8 Hz, the thermal noise at 10 Hz (horizontal + vertical) is $1\times10^{-19}\ \mathrm{m}/\sqrt{\mathrm{Hz}}$, and above ~15 Hz it essentially follows the lower edge of the suspension thermal noise band of Figure 3. The first violin mode frequency is approximately 480 Hz, versus about 490 Hz for the baseline ribbon geometry.

While this seems like an ideal solution, two things remain to be proven: that the cancellation of thermal expansion by the Young’s modulus change really works and is controllable; that fibers with a varying cross section can be reliably produced and have sufficient strength. The latter task seems no more (and potentially less) challenging than producing strong ribbons with twists.

3.4. Technical noise

Achieving the noise limits shown in Figure 3 requires of course that technical noise sources be sufficiently controlled, a task which can become increasingly difficult at low frequencies. Thus it is imperative to consider if any technical noise sources might present serious obstacles to reaching the more fundamental noise levels. Three such sources that traditionally have been difficult to control at low frequencies are technical radiation pressure, frequency noise, scattered light and suspension local sensor noise.

Technical radiation pressure (power fluctuations of the input light producing differential arm length fluctuations through unbalanced power levels in the arms) is difficult because it has such a steep frequency dependence: $x(f) \propto \delta P(f)/f^3$, with the power fluctuations $\delta P(f)$ typically going inversely with frequency. The Advanced LIGO Systems design calls for a level of power stability that would produce, with an input power of 125 W, an equivalent strain noise 10× below the baseline curve. At 30 W, the same relative power stability would produce a strain noise of $h = 5\times10^{-24}/\sqrt{\mathrm{Hz}}$ at 10 Hz – below the pendulum thermal noise by a factor of 6-8. This is based on a power imbalance of 1% in the arms, and a relative input power stability of $2\times10^{-9}/\sqrt{\mathrm{Hz}}$ at 10 Hz. Both levels will be challenging to achieve, but are not out of reach.

Good frequency stability at low frequency has traditionally been difficult to achieve because the interferometer error signal used in the final stabilization loop is also fed back to the arm length to damp the test mass pendulum modes, limiting the frequency noise suppression factor. In AdLIGO, it should be possible to leave the test masses locally damped and avoid common mode feedback to the arms; this would be possible with eddy current damping of the suspension, or possibly with light active damping. In this scenario the loop gain limitations to frequency stability would be removed; the remaining challenge of achieving adequate sensing performance should be no more difficult than the gw channel sensing. Thus it is likely that frequency noise can be brought below the levels of Figure 3.

Scattered light can produce upconverted noise if the scattering path is modulated by more than a wavelength. The maximum frequency of the upconverted signal is: $f_{\mathrm{max}} = 2v_{\mathrm{max}}/\lambda$, where v_{max} is the maximum relative velocity of the scattering path; thus we would require $v_{\mathrm{max}} < 5\lambda/\mathrm{sec}$ for $f_{\mathrm{max}} < 10\ \mathrm{Hz}$. This will inevitably require some isolation of the interferometer tables or components external to the vacuum system, as the ground motion can easily exceed 5 microns/sec in velocity. It will be a difficult noise source to control, but it does not appear to present a threshold in the 10-15 Hz region.

On the test mass suspensions, all local damping–if active–can be turned off during operation except for the vertical motion damping (there are interferometer signals for the longitudinal, pitch and yaw motions, but not vertical motion). Local sensor noise will induce a displacement signal according to the noise floor of the sensors, the vertical-horizontal coupling factor, and the level of active damping; assuming velocity proportional feedback, the equivalent test mass motion is:

$$x(f) \approx 3\times10^{-17}\frac{\mathrm{m}}{\sqrt{\mathrm{Hz}}}\left(\frac{10}{Q}\right)\left(\frac{x_{\mathrm{sensor}}}{10^{-10}\mathrm{m}/\sqrt{\mathrm{Hz}}}\right)\left(\frac{\text{v-h coupling}}{10^{-3}}\right)\left(\frac{10\mathrm{Hz}}{f}\right)^{7}$$

Clearly for active vertical damping to be feasible, we need some combination of low damping (the Q above corresponds to that of the lowest vertical mode), lower sensor noise, and electronic (digital) filtering of the local sensor feedback signal. Even then, given the steep slope to the coupling, getting the local sensor noise out of the way significantly below 10 Hz may not be feasible. A much more promising alternative from the noise standpoint is eddy current damping of the suspension modes. The ‘equivalent sensing noise’ of eddy current dampers is so small that even low Q’s would be allowed (if the damping force could be made large enough).

3.5. Sensitivity trade-offs

We should also ask whether pushing the interferometer performance at low-frequencies might otherwise compromise its sensitivity. The most significant concern would be stressing the final stage suspension fiber so much (to get a low vertical mode frequency) that excess mechanical noise is introduced (non-gaussian, stress-release events). The nominal design for a sub-10 Hz vertical mode uses a fiber loaded at ~1/3 of the breaking stress. Unfortunately there is little or no data on non-gaussian events in stressed fused silica fibers.

Another concern might be increasing the density (in frequency-space) of violin modes if the suspension fiber length were increased. Over the range being considered, however, this isn’t a significant sensitivity driver, since the fraction of the gw band occupied by the violin modes would still be very small.

4 SIGNAL PROCESSING OPTIONS

It is worth considering what signal processing might be able to achieve in removing low frequency components from the data stream. Kalman filtering will undoubtedly be applied to remove the suspensions’ violin modes from the data–why not do the same for the vertical mode? If this could be done successfully, the vertical mode frequency then would not represent a minimum low-frequency cutoff, and we would be free to choose its location based on technical trade-offs. Several factors make this line removal a more difficult task than what has been accomplished to date, however. Kalman filtering of the violin modes, applied to data from the Caltech 40m prototype, has achieved a ~40 dB suppression of a dozen or so modes, each over approximately a 1 Hz bandwidth[1]; the suppression was limited by non-linear couplings not included in the model.

For a suspension design with a vertical mode slightly above 10 Hz, the thermally driven vertical motion would be $1-2\times10^{-13}$ m-rms, the seismically driven motion around 10^{-15} m-rms. Assuming a vertical-horizontal coupling of 0.001, vertical mode motion would exceed horizontal thermal motion over a 2-3 Hz bandwidth; the vertical mode peak would need to be suppressed by 60dB to bring it to the level of the broadband horizontal motion. Compared to what has been

---

[1] S. Mukherjee and L.S. Finn, “Removing Instrumental Artifacts: Suspension Violin Modes”, Gravitational Waves, Third Edoardo Amaldi Conference, ed. S. Meshkov, AIP, 1999.

achieved with this algorithm, an additional factor of 10 suppression is needed (which may force one to deal with the non-linear couplings), and it must be applied to a much lower frequency line.

5 UPGRADE OPTIONS

Consider the scenario where it is concluded that the technical risks associated with a design that achieves a sub-10 Hz vertical mode are too great for implementation in AdLIGO, and at some point the performance of AdLIGO becomes limited by the vertical mode noise. It is worth considering what initial design choices would better lend themselves to later upgrades aimed at eliminating the performance limitation. These would be:

- Replacing the last stage suspension fiber with one of different cross-sectional geometry, but same length (e.g., replacing a relatively thick round fiber with a variable cross-section fiber). This is the most conceivable upgrade, in that it would have the least impact on the rest of the system (the suspension design must include some strategy for replacing/repairing fibers in any case).
- Replacing the last stage suspension fiber with a longer fiber, possibly also of different cross-sectional geometry. This is still conceivable, but would have to anticipated by incorporating a spacer between the initial suspension and the seismic platform, which would be removed (or shortened) when the longer fibers are installed (to keep the interferometer optic axis at the same height).
- Replacing the penultimate mass with a heavier mass (possibly in conjunction with the above changes). This would be the most disruptive modification, since the suspension wires and blade springs in the upper stages of the suspension would have to be changed as well to accommodate the increased load; essentially the each (test mass) suspension would have to be rebuilt. The seismic platform would have to be designed from the start for the larger load; if a heavier suspension were installed, ballast would be removed from the platform to maintain a constant total load.

6 RECOMMENDATIONS

While we strive to design for the best interferometer performance possible, there is not a strong case for demanding a sub-10 Hz suspension vertical mode: there appears to be no significant impact for source detection; furthermore, it appears quite promising that the vertical mode signal could be effectively removed from the data with Kalman filtering; finally, Figure 3 indicates that the band below ~15 Hz may be obscured by gravitational gradients in any case[1]. Based on the considerations discussed above, we make the following recommendations:

- Vertical mode frequency: 12 Hz or lower. Given the width of the vertical mode, this would establish a region of no observations between 11 and 13 Hz. Line removal of the resonance may allow this frequency band to be recovered. A sub-10 Hz vertical mode is still desirable, if it can be designed to also meet the additional conditions below (and without

---

[1] It is conceivable that the gravitational gradients could be measured independently with an array of seismometers, allowing their effect to be subtracted from the gw data; this is not planned for Advanced LIGO, but could be part of a later upgrade.

introducing increased potential for excess noise); this probably requires a ribbon or a variable cross section circular fiber.

- Horizontal thermal noise: continues to be specified as 10^{-19} m/$\sqrt{\mathrm{Hz}}$ or lower at 10 Hz (per test mass; or 5×10^{-23} /$\sqrt{\mathrm{Hz}}$ in strain for the four test masses).
- Violin mode fundamental frequency: no lower than 400 Hz.
- Technical noise sources to be held at levels to allow observation down to 10 Hz; in addition to the noise sources discussed in section 3.4., this also applies to the suspension's local damping noise.