New Folder Name Low Frequency Noise

LOW FREQUENCY NOISE IN GRAVITATIONAL WAVE INTERFEROMETERS

PETER R. SAULSON*
Joint Institute for Laboratory Astrophysics
National Institute of Standards and Technology and University of Colorado
Boulder, CO, USA

submitted to: Gravitational Astronomy - Instrument Design and Astrophysical Prospects
eds D McClelland and H.-A Bachor, Singapore: World Scientific

ABSTRACT

The ultimate low frequency performance of terrestrial gravitational wave detectors is limited by the uncertainty principle and by fluctuating gravitational fields from the noisy environment. In any device likely to be built in the near future, however, seismic noise and Brownian motion of the test masses will almost certainly dominate the noise budget below 100 Hz. I present the expected amplitudes of these noise sources. The predictions of the thermal noise spectrum are based on new models using measured loss properties of the materials that might be used for pendulum wires. For seismic noise, I describe active systems optimized for vibration isolation between 1 and 10 Hz. The factors that could limit performance, including available servo bandwidth, sensor noise and thermal noise, are discussed.

1. Introduction

In the quest for the most sensitive interferometric gravitational wave detector, high frequency performance will be the province of the optics experts. The fundamental limit should be shot noise in the laser light, and improvements will come from the development of more powerful lasers, optical components which can withstand them, and clever optical schemes to make more efficient use of the light.

At low frequencies, the mechanical aspect of the interferometers will probably dominate the performance. What is important there is all of the ways that the mirrored test masses depart from the idealization of a freely-falling test body. Fundamental limits come from quantum mechanics and thermodynamics.

* Address starting January 1, 1991: Department of Physics, Syracuse University, Syracuse, NY 13244-1130, USA

There are also important effects from mechanical and gravitational coupling to the noisy terrestrial environment.

In this paper, I will give a brief review of the expected levels of these low-frequency noise sources. I will give the greatest emphasis to the results of a design study for active seismic isolation in the 1-10 Hz band. My goal is to assess the prospects of reducing the mechanically filtered seismic noise of the test mass to the level of the irreducible noise limit in that band, gravitational noise driven by the external seismic spectrum.

2. The Quantum Limit

The quantum limit to sensitivity of an interferometer can be understood by thinking about it as an oversized version of the "Heisenberg Microscope". There is an inevitable competition, at the quantum level, between the precision of a measurement and the disturbance that the measurement makes to the system to be measured. The quantum limit to measurement precision *per se* is the shot noise in the intensity at the output port of the interferometer, mentioned above. Disturbance of the system comes from shot noise in the difference in the radiation pressure in the two arms. A clear quantum mechanical explanation of the existence of these two effects in an interferometer was first given by Caves.[1] (See the discussion in this volume by Bachor.)

The measurement precision is improved by increasing the number of photons N_τ detected in the measurement interval τ ($x_{shot}(f) \propto 1/\sqrt{N_\tau}$). But the radiation pressure noise increases as $x_{rad'n}(f) \propto \sqrt{N_\tau}$. Thus there is an optimum power, itself a function of frequency, at which the noise due to combination of the two effects is minimized (that is, where both effects have the same magnitude). The locus of this minimum noise as a function of frequency is given by[2]

$$x_{q.l.}(f) = \sqrt{\frac{2\hbar}{m}} \frac{1}{2\pi f}$$

and is shown as the dashed line in Figure 1. Note that any particular interferometer (specific combination of arm length, mirror mass, and circulating optical power) can only reach the quantum limit at a single frequency.

3. Gravitational Noise

A gravitational wave detector is designed to be exquisitely sensitive to weak gravitational effects, and insensitive to other influences. Our hope is that the gravitational effects we will see will be due to something interesting, namely

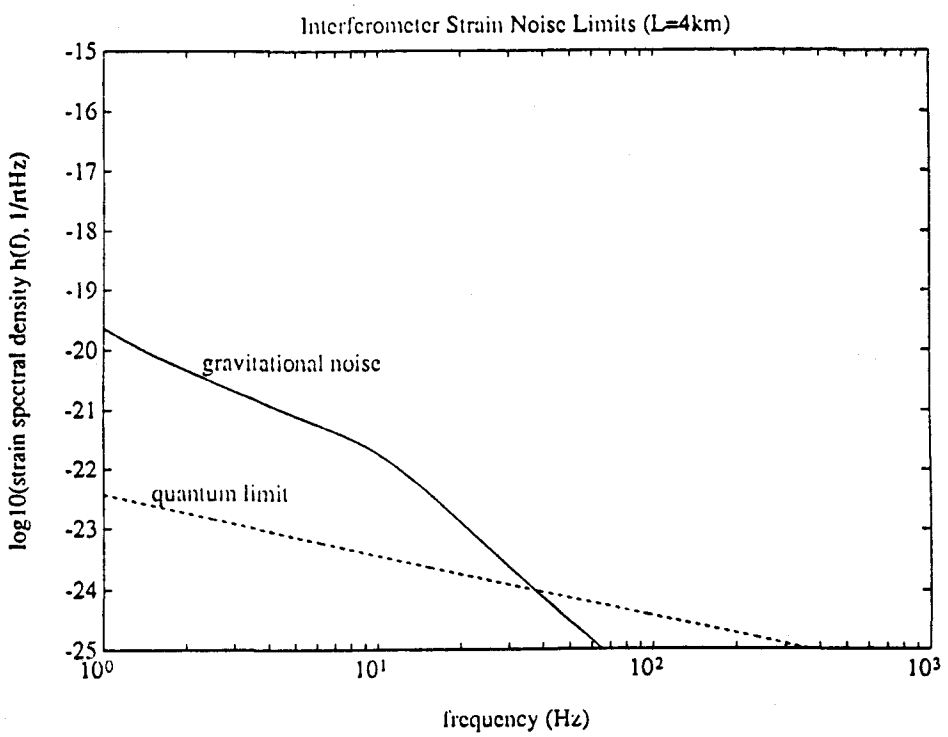

Figure 1: Low frequency limits to the sensitivity of interferometric gravitational wave detectors. The gravitational noise was calculated assuming the seismic spectrum of Section 5.1. The quantum limit line is the locus of minima of the set of fundamental optical noise spectra (combined shot noise in fringe read-out and differential radiation pressure noise). The frequency at which the actual noise spectrum would reach the quantum limit could be adjusted by, for example, varying the injected light power.

gravitational waves from extraterrestrial sources. But as a matter of principle, we could not isolate the apparatus from any other gravitational effects that might fall in our signal band.

If the distribution of matter in the vicinity of the interferometer is time-dependent, then the gravitational field can fluctuate as well. This will cause the (nearly freely-falling) test masses to move, generating a spurious gravitational response in the interferometer. Motion of macroscopic objects (such as vehicles, animals, or scientists) in the vicinity of the test masses can cause detectably large signals in some cases.[3]

Easier to quantify, and harder to prevent, is the gravitational noise caused by the seismic background.[4] The compressional waves in the ground cause density gradients to be established in the vicinity of each test mass. Transverse surface waves add and subtract thin layers of material in the horizontal plane. The $1/r^2$ character of gravity means that the effects are dominated by what happens to

the ground in the nearest (reduced) acoustic wavelength. The coupling is, up to a small factor of order unity, the same as if the test masses were attached to the ground with a spring that gave a resonant frequency of $\omega_{grav} = \sqrt{G\rho}$, where ρ is the density of the earth in the vicinity of the test masses.

A graph of that spectrum, using an updated model of the seismic background spectrum, is shown in Figure 1. (By coincidence, the analogous contribution from infrasonic density fluctuations in the atmosphere has roughly the same magnitude.) Note that it is larger than the quantum limit for all frequencies below about 30 Hz.

4. Thermal Noise

The test masses of a gravitational wave interferometer are suspended as pendulums with very small damping. Special care is taken to keep the damping low, because associated with mechanical loss is so-called thermal noise, whose force power spectrum is given by the fluctuation-dissipation theorem:[5]

$$F^2(f) = 4k_B T R(f).$$

Here $R(f)$ is the real part of the mechanical impedance at the test mass.

Usually, the losses in a pendulum have been modeled as if they were produced by a dashpot with $F_{damp} = bv$. Then $R(f) = b$, and the force power spectrum is white. Applied to the pendulum, this gives a response $x_{th}(f) \propto f^{-2}$ above the resonant frequency. But the losses in the flexing pendulum wires (the dominant source of damping at high vacuum) often show a quite different frequency dependence. Damping in materials can be represented by a Young's modulus with a small fractional imaginary part $\phi(f)$. Although sometimes $\phi(f) \propto f$ (which would mimic velocity damping), it is more often constant, and sometimes departs even more widely from dashpot form. Whatever the functional form of the loss angle $\phi(f)$, the fluctuation-dissipation theorem predicts[6]

$$x_{th}^2(f) = \frac{4k_B T k\phi}{2\pi f((k - m(2\pi f)^2)^2 + k^2\phi^2)},$$

or $x_{th}(f) \propto f^{-2.5}$ for frequencies above the resonance of the pendulum (if $\phi(f)$ is constant).

Losses in several possible pendulum wire materials were measured by Kovalik.[7] Tungsten has a loss factor ϕ of about 7×10^{-4}, roughly independent of frequency between 30 Hz and 1 kHz, while fused quartz showed about 2×10^{-6}, also nearly independent of frequency. Sapphire and etched silicon gave values in between those of tungsten and fused quartz.

A pendulum made of these materials should show substantially lower losses than a conventional mass-spring oscillator. This is because the material loss factor only acts on that fraction of the pendulum's energy stored in the flexing wire. The fraction is about 3×10^{-4} for tungsten wires 30 cm long, in a 4-wire (2 loop) pendulum supporting a 10 kg mass at half the breaking stress. For fused quartz, the energy fraction is about 1×10^{-2}.

Note that the thermal noise will not be much reduced by increasing the size of the test masses. The explicit $m^{-1/2}$ dependence is partly counteracted by an increase in the fraction of energy stored in the wires as the mass is increased (assuming the number of wires is kept fixed but their cross section is scaled with the mass). The net dependence is $x(f) \propto m^{-1/4}$ (Ref. 8).

5. Seismic Noise

5.1. Passive Isolation

The spectrum of seismic noise at a typical remote site[9] can be approximated by $10^{-7}\text{cm}/\sqrt{\text{Hz}}$ from 1 to 10 Hz, and by $10^{-5}\text{cm}/\sqrt{\text{Hz}}(1\text{Hz}/f)^2$ for frequencies above 10 Hz. A brief calculation shows that the attenuation of a pendulum with a 1 Hz resonant frequency is insufficient to reduce the motion of the test mass below a negligible level, say $10^{-17}\text{cm}/\sqrt{\text{Hz}}$, at frequencies below about 1 kHz. For this reason, the pendulum's isolation is typically supplemented by a multi-element vibration isolator, supplying many additional low-pass poles in the vicinity, usually, of around 10 Hz.

Why not make an isolation system with extremely low resonant frequencies, and thus very good isolation? A completely passive isolator must be long in order to be soft – a Hooke's Law spring will sag under load by an amount $\delta z = mg/k$. A long spring exhibits internal resonances at much lower frequencies than a short stiff spring, shrinking the frequency range over which one can construct a useful isolator.

The most ambitious design for a passive isolator is the Super Attenuator of the Pisa group.[10] It consists of seven 400 kg masses, each hung from the last by a specially-designed air spring that gives nearly isotropic compliance soft enough for the fundamental resonant frequency to be around 1.4 Hz. Such a multi-stage isolator is expected to give isolation of $(1.4\text{Hz}/f)^{14}$, for frequencies large compared to the resonant frequencies. This isolation would be sufficient to make seismic noise negligible above 10 Hz. Measurements have confirmed isolation of at least 150 dB at 10 Hz. It is important to note that this system partly evades the difficulties described in the previous paragraph because it is not a truly passive system. The lengths of the springs are adjusted by careful setting and control of the pressure in the air springs.

Application of such a degree of vibration isolation in a gravitational wave interferometer could yield substantial dividends, assuming the other noise sources were small. But even with such heroic efforts, a wall of seismic noise prevents useful observations in the decade from 1 to 10 Hz.

5.2. Suspension Point Interferometer

A technique proposed by Drever[11] should be able to reduce the effect of seismic noise on an interferometer, over a broad band starting just above the pendulum resonance. A second interferometer, called a suspension point interferometer or seismic monitor interferometer, measures the difference in separations of the tops of the test mass pendulum wires. When this interferometer is locked (by feedback to the tops of one or more of the pendulums), differential length changes of the two arms are reduced. This component of the seismic noise input is exactly what causes spurious interferometer output that can mimic a gravitational wave.

The reduction in noise can be susbstantial, but will be limited by one of the following factors. 1) Servo sensitivity and bandwidth must, as always, be maximized, but are unlikely to limit the performance below 10 Hz. 2) The reduction factor will not be greater than the finesse of the main cavity arms, since the suspension point interferometer is conceived as a one-bounce system, and thus won't properly account for noise from motion of the beamsplitter. 3) The transfer functions of the four pendulums must also be well matched, since the method depends on their responding the same to common-mode inputs (which are not attenuated by the system). 4) Finally, the system will respond incorrectly to vertical seismic noise. The interferometer has a (reduced) sensitivity to vertical noise because the optic axes will not be precisely level, either because the interferometer is installed on a slope, or at least because the curvature of the earth forbids a straight line to be precisely level at more than one point. This last effect limits the maximum reduction factor of the suspension point interferometer to be about 3000.

For specificity in our models below, we'll take the improvement from this technique to be a factor of 1000.

5.3. Active Isolation

To get around the technical limitations to passive vibration discussed above, one can consider using feedback to give active isolation. An active vibration isolation system for one degree of freedom starts with an inertial proof mass resonantly suspended from a movable platform. (See Figure 2.) A position

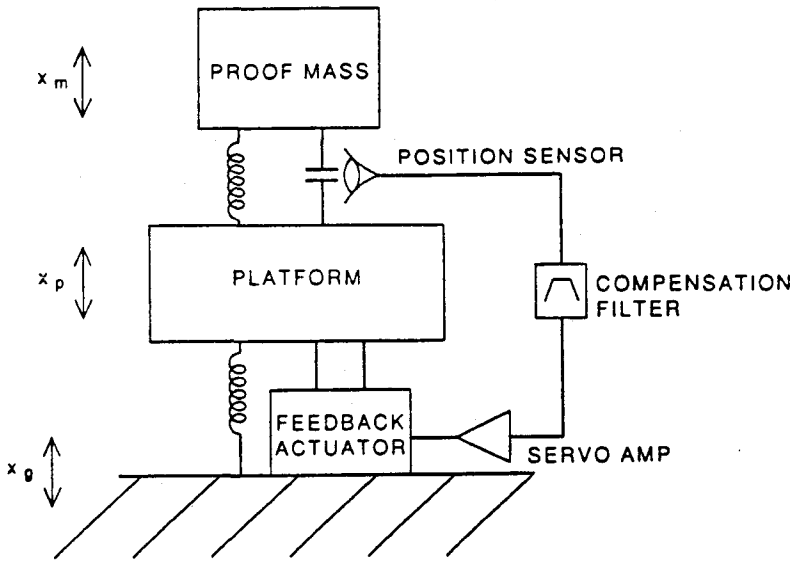

Figure 2: Schematic diagram of a one-dimensional active vibration isolation system. The position sensor measures the relative displacement between the platform and proof mass. The feedback actuator applies a force to the platform that is proportional to the (suitably filtered) position sensor output.

sensor measures the displacement between the platform and the proof mass. The difference between the actual displacement and the nominal DC spacing is the error signal that is amplified, filtered, and used to drive a mechanical actuator that can move the platform. The feedback, by nulling the error signal, causes the platform to track any motion of the proof mass.

Why does this provide vibration isolation? One way to see this is to note that external vibration only moves the proof mass if, by moving the platform, the spring that connects the platform and proof mass is stretched, thus applying to the proof mass a Hooke's Law force $F = -k\delta x$. If the servo works to make the platform follow the proof mass, then that spring's length is being held approximately fixed, nulling the force that would have been applied to the proof mass. And in bootstrap fashion, since the platform's position is being servoed to that of a well-isolated proof mass, it in turn responds less to external vibration than in the open-loop state.

Single-axis active vibration isolation systems have been built by me and others working on gravitational wave detectors, and in geophysics.[12–14] For the past several years, the JILA group of Bender, Faller, and Stebbins have been studying the design of multi-axis active isolation systems appropriate to long-baseline interferometric gravitational wave detectors.[15]

5.4. Ideal Single Degree of Freedom System

The open loop system has a transfer function from seismic vibration to motion of the platform like that graphed in Figure 3. It is in most respects like the transfer function of a single stage of seismic isolation, with unit value at low frequencies and transmission proportional to $(f_0/f)^2$ at high frequencies. The null and second resonant peak are due to the attached proof mass.

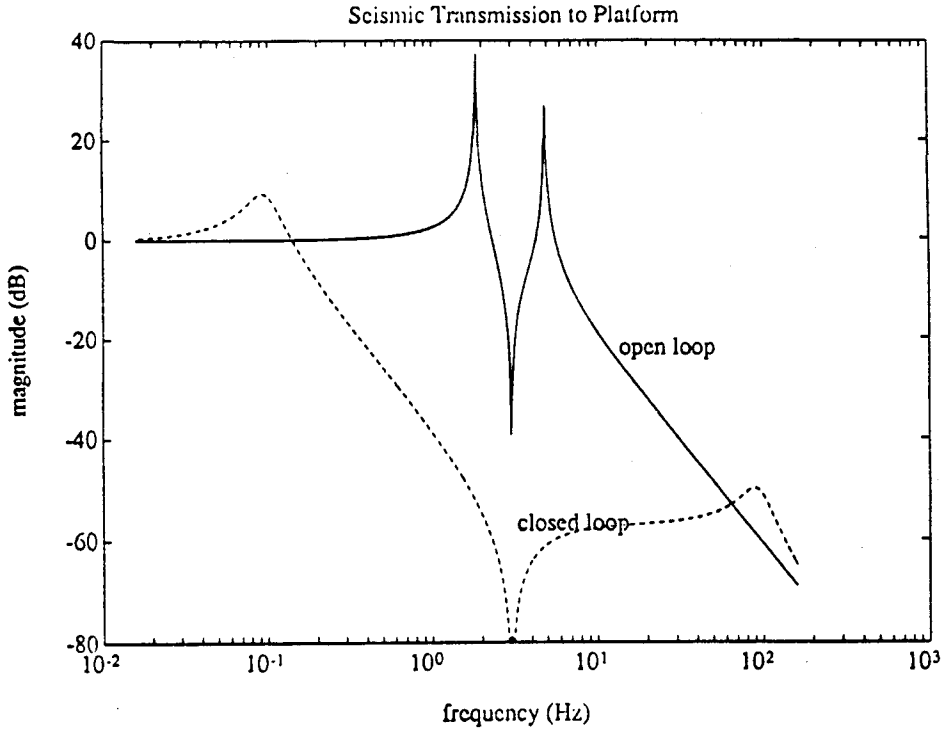

Figure 3: Seismic transfer functions for a one-dimensional active vibration isolation system, defined as the magnitude of the response of the platform to a unit vibration amplitude of the floor on which the system rests. The solid line shows the passive (or open loop) response. The dashed line shows the response when the loop is closed with a gain that gives a bandwidth of approximately 100 Hz. The Q of the closed loop resonances is determined by the lead and lag filters used for compensation.

The vibration isolation servo is based on sensing the displacement difference between the proof mass and the platform, and applying a feedback force to the platform to null the error signal. The loop transfer function for such a servo is presented as a Bode plot in Figure 4. At high frequencies, it falls proportionally to f^{-2}, since the proof mass is approximately fixed in inertial space, while the response of the platform to an applied force is governed by its inertia according to Newton's Second Law, $a = F/m$, or $x = F/m\omega^2$. The two normal modes of

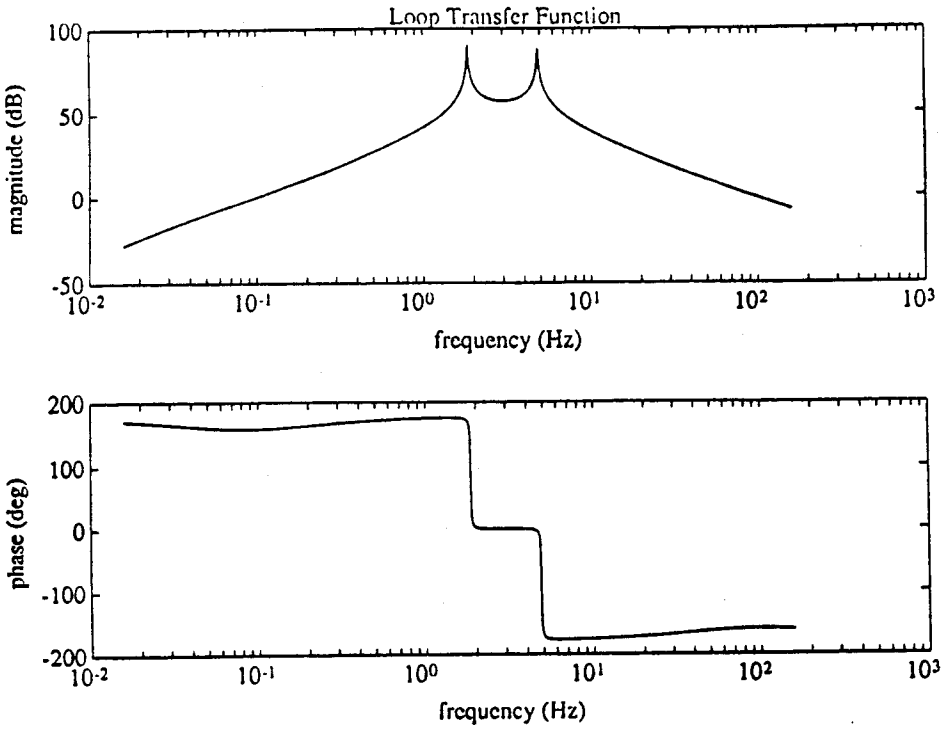

Figure 4: Loop transfer function for a one-dimensional active vibration isolation system, presented as a Bode plot showing the magnitude in decibels and the phase in degrees. A compensation filter consisting of a lag at low frequencies and a lead at high frequencies has been included in the system to reduce the Q of the closed loop resonances.

vibration appear as resonances in the loop transfer function, each accompanied by a 180 degree phase shift. At low frequencies, the loop transfer function is proportional to f^2 – the platform motion is governed by Hooke's Law, $x = F/k$, but the difference in positions of the platform and proof mass is that of an accelerometer, since we are below the proof mass's resonant frequency.

This loop transfer function has several interesting features that are worth pondering. Perhaps most unusual is the fact that the gain goes to zero at DC, instead of the more common case of large finite DC gain. As I will show below, this feature is reflected in the way the closed-loop system mimics, at low frequencies, a passive vibration isolator of lower resonant frequency than the open loop system. Another unusual (and related) feature is that the phase of the loop transfer function makes a total excursion of -360 degrees from low to high frequency. To demonstrate that such a servo loop can be stable, we can examine the Nyquist plot of the loop transfer function, shown in Figure 5. The f^2 and f^{-2} sections of the transfer function appear as sections that asymptotically approach the negative real axis as the gain approaches zero. Because there are

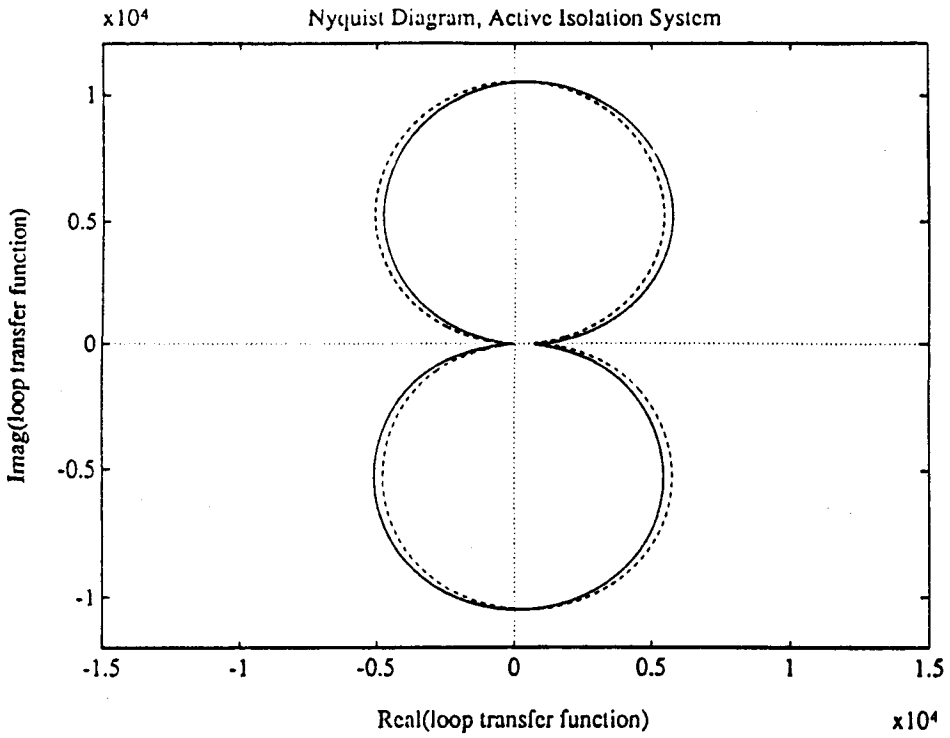

Figure 5: Nyquist diagram, or polar plot of the loop transfer function, for a one-dimensional active vibration isolation system. The solid line shows the positive frequency section of the locus, the dashed line the negative frequency section. Note that the locus is asymptotic to, but never crosses, the negative real axis. This demonstrates the stability of the (ideal) system.

no crossings of the negative real axis, there are no encirclements of the point -1, and thus such a servo will be stable for arbitrarily high gains, albeit with vanishing phase margin. (The phase margins can be adjusted by a compensation filter.) I will discuss below the most important departures of real systems from this ideal model.

When the loop is closed, the transfer function for platform motion due to ground motion will look like the dashed line in Figure 3. Instead of peaks at the passive system's resonant frequencies, the closed loop system has two resonances, one at each of the two frequencies at which the loop transfer function has unit magnitude. Just above the lower closed-loop resonance, the vibration transfer function is proportional to f^{-2}, exactly as if the active system had a mechanical resonance at such a low frequency. A way to see graphically that this is the expected behavior goes as follows. The ratio of any open loop transfer function to the equivalent closed loop transfer function is given by $1/(1+G(f))$, where $G(f)$ is the loop transfer function. The denominator becomes significantly greater than unity above the low-frequency unity-gain point, growing in magnitude like

than unity above the low-frequency unity-gain point, growing in magnitude like f^2. So the closed loop seismic transfer function should fall below the open loop transfer function (which is constant in this frequency range) by a factor f^{-2}, as the detailed calculation shows.

The high frequency behavior is also easily understood by this line of reasoning. Between the second open-loop resonance and the high-frequency unity-gain point, the factor $1/(1+G(f))$ is proportional to f^{-2}, so the improvement that the closed loop system makes over the open loop case falls with increasing frequency. But the isolation of the passive system (defined as the inverse of the seismic transfer function) rises with increasing frequency as f^2. The product of the two factors gives the constant response of the closed loop system in this frequency range.

5.5. Limits to Performance: Internal Resonances

The preceding discussion explains how an ideal active isolation system works. Any real version will depart from the model in several respects, some of which will limit the performance. The most important effects to understand are the limits to the useful bandwidth, and the response of the system to noise inputs other than vibration of the ground.

One of the assumptions that we made in deriving the equations of motion was that we could treat the masses as perfectly rigid, and the springs as perfectly massless. Actually, neither assumption can give a good representation at high frequencies. The lumped-mass or discrete model is only an approximation to a distributed system that has high-frequency normal modes in the neglected degrees of freedom. Even with good mechanical design, it is hard to build a system of reasonable size that has no additional resonances below 1 kHz. All of these resonances cause some modification of the vibration transfer functions, appearing as transmission peaks or regions where the isolation is poor. A subset of these will have an even more deleterious effect, adding extra 180 degree phase shifts to the loop transfer function. Such additional large phase shifts will add to the Nyquist plot an extra loop that crosses the negative real axis, enabling it to encircle the point -1 at a finite gain. For the system to be stable, those crossing points must have a magnitude of $G(f)$ less than unity. Thus, internal resonances can set a limit to the useful gain and thus to the improvement in isolation.

Not all internal resonances will add additional 180 degree phase shifts to the loop transfer function. The only ones that do so are those whose mode shape adds a node between the location of the actuator and the part of the platform to which the platform side of the position sensor is mounted.[16] Thus this problem at least can be minimized by placing the actuator and sensor in close proximity, so that only very high frequency modes can limit the useful bandwidth.

The other important feature we need to add to our models is the response of the system to disturbances other than vibration of the ground below. Some of these, other forces that act directly on the platform, are attenuated by the servo to the same degree as external vibration, and are therefore not worrisome. Disturbances applied directly to the proof mass are a different matter. The action of the servo system is to make the platform track the proof mass, on the assumption that it is more nearly fixed in inertial space. To the extent that the proof mass is itself noisy, the platform (in the high gain limit) will also be noisy. One unavoidable noise term of this sort is the thermal noise of the proof mass. We will evaluate its magnitude in the models below. Another noise that enters in the same way, although it is not properly a noisy force, is noise in the position sensor. This causes a spurious error signal, unrelated to platform motion, just as if the proof mass were actually moving.

The response of a feedback system to noise is most clearly analyzed by reference to a block diagram, such as the one shown in Figure 6. The transfer function from noise injected at any point to the response at the output is given by $H_{fio}/(1+G)$, where H_{fio} is the gain in the forward direction from the injection

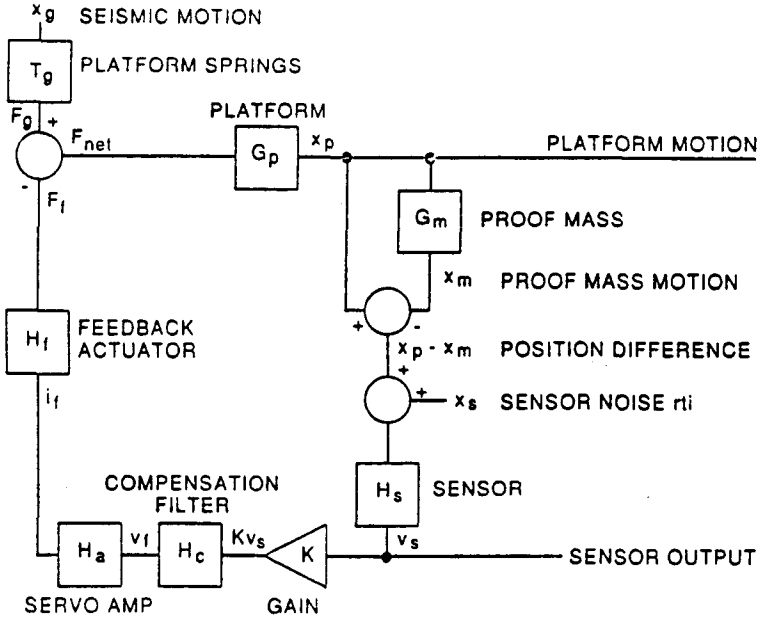

Figure 6: Block diagram of a one-dimensional active vibration isolation system. Each block represents the frequency-dependent transfer function of an individual element in the servo. The loop transfer function is given by $G(f) = KG_p(1-G_m)H_sH_cH_aH_f$. The seismic transfer function is $T_gG_p/(1+G(f))$.

point to the output.[17] Typically, the sensor is the component whose noise is most important in a feedback system, since the path from the sensor to the output can contain most or all of the loop. This means that the feedback denominator is cancelled by the numerator in the expression above. In other words, increasing the feedback gain would not reduce the effect of sensor noise (in contrast to the effect of feedback on external vibration).

In this particular system, H_{fio} contains all of the elements whose gains are multiplied to form G, with the exception of the block G_m and the following difference junction. These elements represent the transfer function from platform motion to the difference in motions between platform and proof mass (which is what the position sensor actually looks at). Thus, in the band where $G >> 1$, the noise transfer function is approximately

$$\frac{x_p}{x_p - x_m} = \frac{\omega^2 - \omega_0^2}{\omega^2}.$$

Above the proof mass resonant frequency ω_0, this transfer function is nearly equal to unity, equivalent to the usual servo case. Below the resonance, the magnitude of the noise transfer function is approximately ω_0^2/ω^2. This amplification of the sensor noise is caused by the attenuation of the response of the system at low frequencies, reflecting the fact that the proof mass is not well anchored to inertial space below its resonance, but instead tends to track the platform motion. The strong low frequency response of this system to sensor noise must be an important consideration in the design process.

5.7. Multi-Degree-of-Freedom System Design

So far our discussion has been limited to systems that provide active isolation in one dimension only. In most applications, we will want to provide approximately isotropic isolation. Gravitational wave interferometers are clearly most sensitive to vibration along the optic axis of each arm, but pendulums are such anisotropic isolators themselves that any additional isolation must almost certainly be close to isotropic. In addition, the suspension point interferometer is a completely anisotropic device. If we want to include the full improvement factor from it, any additional noise reduction must be isotropic.

Many design considerations carry over, of course, from a one-dimensional system to a multi-axis version. In the remainder of this section, I will discuss several alternative designs for six-degree-of-freedom systems, and will analyze one in detail to indicate what the chief problems are likely to be and what performance we can expect to achieve.

A related device that represents a kind of idealization of the problem is a "drag-free" satellite.[18] The key idea is to place a freely-falling proof mass inside

a housing, which is servoed to track the motion of the proof mass in all degrees of freedom. The housing protects the proof mass from external forces (such as atmospheric drag or solar radiation pressure), and by the action of the servo is itself made less sensitive to such effects. This system has a manifest symmetry between the different degrees of freedom, making it plausible that each could be treated as a one-dimensional system uncoupled to the others. Then the design could proceed as above. Note, though, that the drag-free satellite, because it is freely falling, feels no force of gravity. Thus, the proof mass needs no springs to hold it up.

The non-zero gravitational acceleration in a laboratory isolation system is responsible for two important complications. One is the need to support the proof mass with a spring, leading to the existence of a finite resonant frequency for the proof mass, with its resulting attenuation of low-frequency signals and consequent amplification of low-frequency noise. The other is an unavoidable confusion between horizontal accelerations of the platform and tilts about horizontal axes, both of which yield non-zero error signals in horizontal inertial sensors.

To actively isolate a platform in all six of its rigid body degrees of freedom, we need to provide a means of inertially sensing all of the translations and rotations, and a means of applying appropriate forces and torques to null the measured error signals. One can imagine several ways of doing this. A scheme that is a close analogy of a drag-free satellite is to use a single proof mass in the center of a hollow platform. Six position sensors measuring displacements between the proof mass and the platform at various locations are enough to allow all of the degrees of freedom to be controlled. One possible arrangement is shown in Figure 7. Actuators may be placed immediately adjacent to the sensors, on the opposite side of the platform wall. Magnetic linear motors are suitable for this purpose, providing a force in one direction without constraining motion in orthogonal directions.

Another set of schemes replaces the single proof mass with six proof masses. Each proof mass is suspended so that it is free to move in one degree of freedom only. In one version, six linear seismometers are attached to the platform, oriented the same way as the position sensors in the single-proof-mass scheme described in the previous paragraph. Another version uses three linear sensors and three angular sensors – these can be arranged with one linear sensor and one angular sensor aligned with each of three orthogonal axes.

What are the advantages and disadvantages of these distinct schemes for multi-axis vibration isolation? The single proof mass design has the virtue of mechanical simplicity, having the fewest parts. This may be important in achieving the highest useful bandwidth free of phase shifts from high frequency resonances. The single proof mass also means that, for a given total system mass, thermal noise should be lower, as the moving mass for each degree of freedom

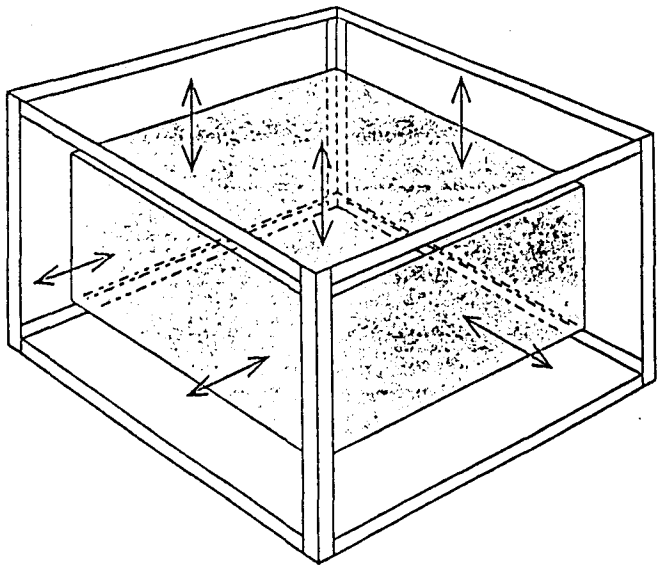

Figure 7: Schematic perspective view of a single proof mass inside a hollow platform, for use in a six degree-of-freedom active vibration isolation system. The arrows indicate the location and orientation of six one-dimensional position sensors that together could be used to control all of the degrees of freedom. If six separate seismometer movements were used, they could be placed and oriented as indicated by the arrows.

is the whole proof mass. On the other hand, the mechanical elegance may be purchased at the cost of requiring careful balancing of the arrangement of springs with respect to the center of the proof mass, to avoid introducing more than the minimal cross-coupling of the degrees of freedom. And, what may prove to be the chief disadvantage of this scheme is that in such a quasi-isotropic suspension of a mass it is difficult to achieve resonant frequencies much below 3 Hz. Since sensor noise is amplified for the whole frequency band below the proof-mass resonance, this will put strong limits on the allowable sensor noise.

If either of these disadvantages prove to be important limits to performance, a system with six linear sensors may be a better solution. It is likely that the sensors will be easier to align and null. An additional advantage is that commercial seismometer movements are available with resonant frequencies as low as 0.2 Hz.

A feature that both of these schemes share is an ability to construct quasi-independent feedback loops based on a local pairing of each sensor with an immediately adjacent actuator. This means that the feedback loops have the least possible sensitivity to the existence of resonant modes in the platform structure. I will show below that loops constructed in this fashion have loop transfer functions

similar to that of a single degree of freedom system, with only minor changes due to the existence of additional degrees of freedom.

5.8. Analysis of Model System

To make these considerations quantitative, I present below the results of calculations of a model of a multi-axis vibration isolation system. The system I chose to model is of the single proof mass type. A schematic diagram is shown in Figure 8. Each mass is free to move in three degrees of freedom, a horizontal direction (x), the vertical (z), and θ, rotation about the y-axis (i.e. rotation in the x-z plane). This should include all of the essential features of a true six degree-of-freedom system, as it exhibits horizontal-tilt coupling as well as control of tilt and vertical motion using a pair of separated vertical loops. The calculations were performed using the program MATLAB (from The MathWorks, Inc.) on a Macintosh II computer.

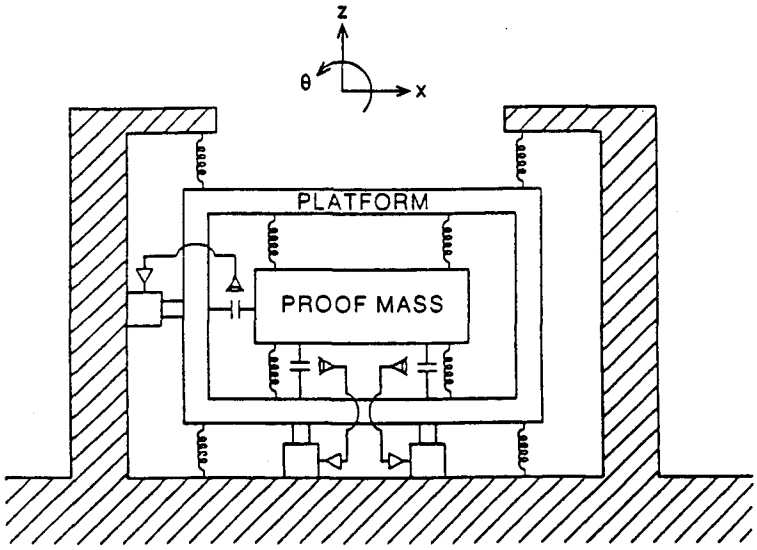

Figure 8: A schematic diagram of the three-degree-of-freedom active vibration isolation system whose performance was modelled in detail. The platform and proof mass are each free to move in the three degrees of freedom (x, z, and θ) confined to the plane of the paper. Three loops are used to provide the active isolation. One loop, using a horizontal position sensor and horizontal actuator at the height of and aligned with the centers of mass of the platform and proof mass, controls the x motion. Two vertical sensor-actuator pairs are offset on either side of the center of the sytem. Even though the loops are closed locally only, as shown in the diagram, well-behaved active isolation is achieved for both the z and θ coordinates.

The model system contains a proof mass of 200 kg, with horizontal dimensions of 80 cm. The platform is 100 cm across, and has a mass of 400 kg. The springs give resonant frequencies in the neighborhood of 3 Hz, with Q's of 10^3.

For each loop that we will want to close in the system, we can calculate the loop transfer function. These are given in Figure 9. The horizontal loop has a transfer function nearly the same as a one-dimensional system, since we have assumed, in effect, that the sensor and actuator are aligned with the center of mass of a symmetrically suspended system. The vertical loops show more interesting structure, since they are offset from the center of mass to allow them to control tilt in addition to vertical displacement. More than two normal modes of the system enter, yet the overall shape of the transfer function is quite similar, in both magnitude and phase, to that of a one-dimensional system.

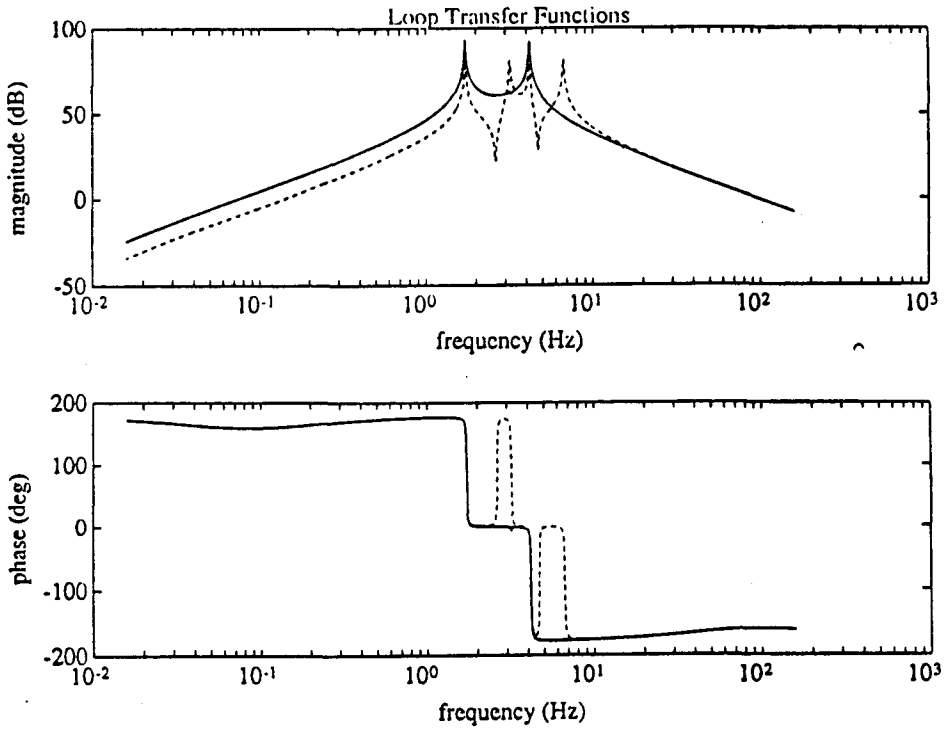

Figure 9: Bode plots of the loop transfer functions for the system shown in Figure 8. The solid lines show the loop tranfer function for the horizontal loop. The dashed lines show the (identical) loop transfer functions for the two vertical loops. The horizontal transfer function is almost identical to that of a one-dimensional system, while the vertical loops exhibit extra structure due to the coupling of vertical and tilt degrees of freedom. As in the one-dimensional system, compensation filters are used to control the Q of the closed loop resonances.

The design of multi-degree-of-freedom control systems has spawned much literature but little intuition, compared with the theory of single-input-single-output systems.[19] One of the few intuitive precepts is that, if a complicated system looks as if it can be broken down into a set of single-input-single-output systems, that is usually the most fruitful control strategy to try first. We follow that strategy here, closing three loops, each of which pairs a sensor with its adjacent actuator. We can verify by direct calculation that the system is stable and has the isolation characteristics that we expect from our one-dimensional experience. Stability of the closed loop system is checked by calculation of its resonant frequencies. All have negative real parts. Thus, cross-coupling between different degrees of freedom does not prevent the system from operating stably.

The closed loop transfer functions, shown in Figure 10, confirm that the performance of the multi-dimensional system has the same basic features as the one-dimensional analog. To make a definite prediction, I assumed that a bandwidth of 100 Hz was available. This is typical of what has been achieved by previous work. The model shows approximately equal isolation against horizontal, vertical, and tilt inputs.

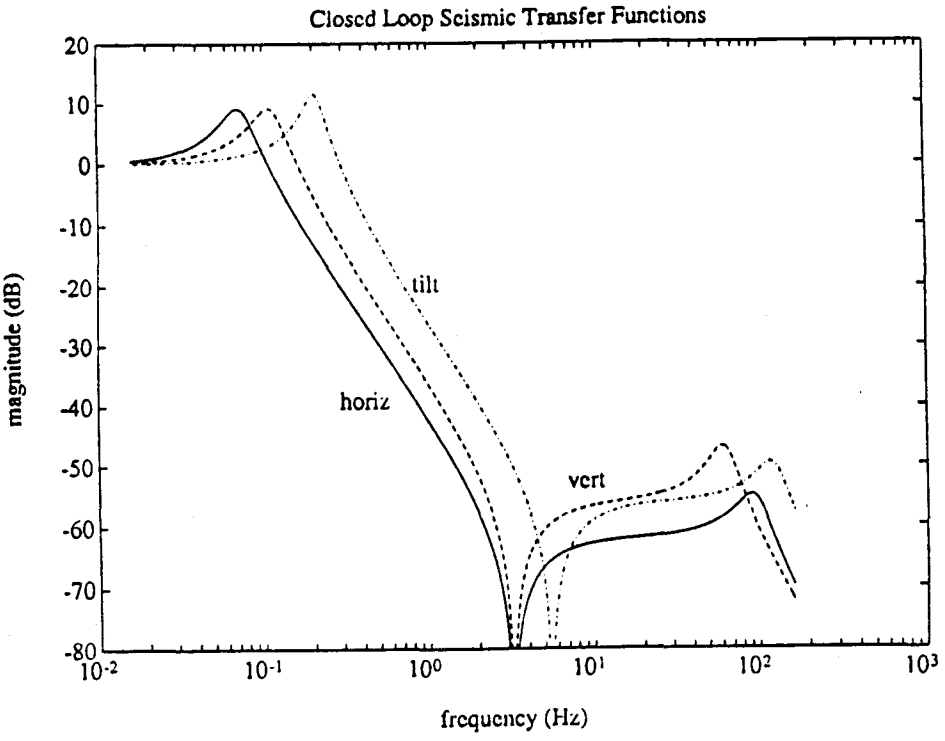

Figure 10: Seismic transfer functions for the three degrees of freedom. The figure shows the platform response when all three loops are closed with gain set to give about 100 Hz bandwidth.

To predict the actual vibration spectrum at the platform, we need to multiply the closed-loop transfer functions by the expected input spectra. In Figure 11, the dashed line is based on the model noise spectrum estimated for the LIGO interferometer sites. I have extended the model to lower frequencies to represent the microseismic peak around 0.1-0.2 Hz. This spectrum should be typical of the noise level at specially constructed remote installations.

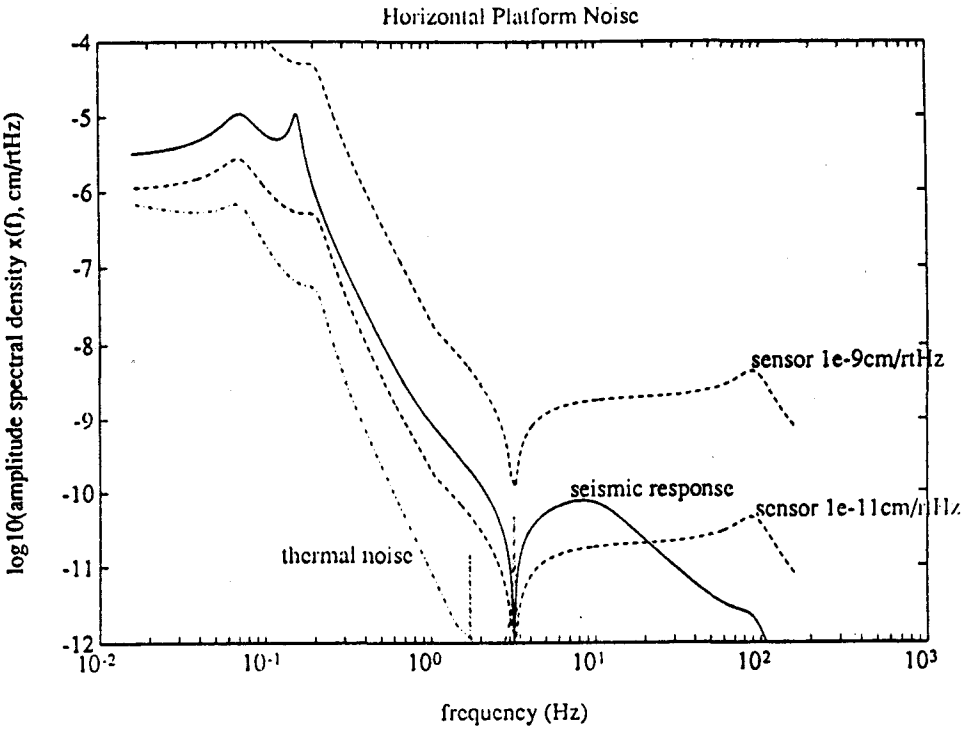

Figure 11: Noise power spectra for the horizontal degree of freedom of the modelled system. The solid curve shows the seismic response to the model input spectrum. The dashed lines show the horizontal platform noise due to sensor noise. The upper curve represents the noise from good split-diode position sensors, while the lower one shows the level caused by modest interferometric sensors. The dash-dot line shows the amount of horizontal platform motion caused by the Brownian motion of the proof mass.

To fully characterize the model, two additional assumptions are needed. One is that the spectra in all linear degrees of freedom are of the same magnitude. This has been typically found to be a good approximation.[20] The other assumption is that an upper limit on the tilt spectrum can be derived by assuming that vertical noise is due only to transverse waves, with no contribution from compressional

waves. Then, the tilt noise spectrum is given by

$$\theta(f) = \frac{2\pi}{\lambda_{acoustic}} x(f).$$

This is almost certainly too large, but probably not by a large factor.

The predicted horizontal vibration spectrum for the actively isolated platform is shown by the solid line in Figure 11. The other degrees of freedom have similar spectra.

We must verify that other additive noise sources will not dominate the isolated seismic noise at the platform. The dash-dot line in Figure 11 shows the calcuated closed loop noise due to Brownian motion of the test mass. The calculation was based on a frequency-independent model of the internal losses in the springs, using the methods described above. Comparison with the other curves in the figure shows that thermal noise will not dominate the noise level of this system.

Sensor noise will be an important concern. Figure 11 shows the closed loop vibration of the platform due to the spurious response of the loop to sensor noise of two different levels. The upper dashed line shows the noise spectrum that results if the sensor has a noise spectrum of 10^{-9} cm/$\sqrt{\text{Hz}}$, representative of the shot noise in good split-diode sensors. The lower dashed line gives the platform vibration that comes from a sensor with a noise level of 10^{-11} cm/$\sqrt{\text{Hz}}$, the shot noise level in low-power interferometric sensors. As discussed above, it is important to compare sensor-driven noise to attenuated seismic noise both in the servo's active frequency band, and in the amplified low-frequency region. The better sensor yields noise lower than transmitted vibration at nearly all frequencies.

6. Interferometer Noise Spectrum

Finally, I summarize the low frequency noise budget in Figure 12. The quantum limit and gravitational noise lines are the same as those in Figure 1. The line labelled "pendulum thermal noise" corresponds to a 1000 kg mass hung from tungsten wires. The seismic noise curve represents the LIGO model spectrum, filtered by an active isolation system like the one described in the previous section, attenuated by a factor of 10^{-3} from a suspension point interferometer. I assumed that the active isolator uses position sensors with sensitivity slightly better than 1×10^{-11}cm/$\sqrt{\text{Hz}}$. The seismic noise is still the dominant noise term. Note that additional provisions for passive seismic isolation for frequencies above 10 Hz have *not* been included in the model.

What would it take to make the seismic noise negligible between 1 and 10 Hz? If a second active isolation system, using sensors with noise below

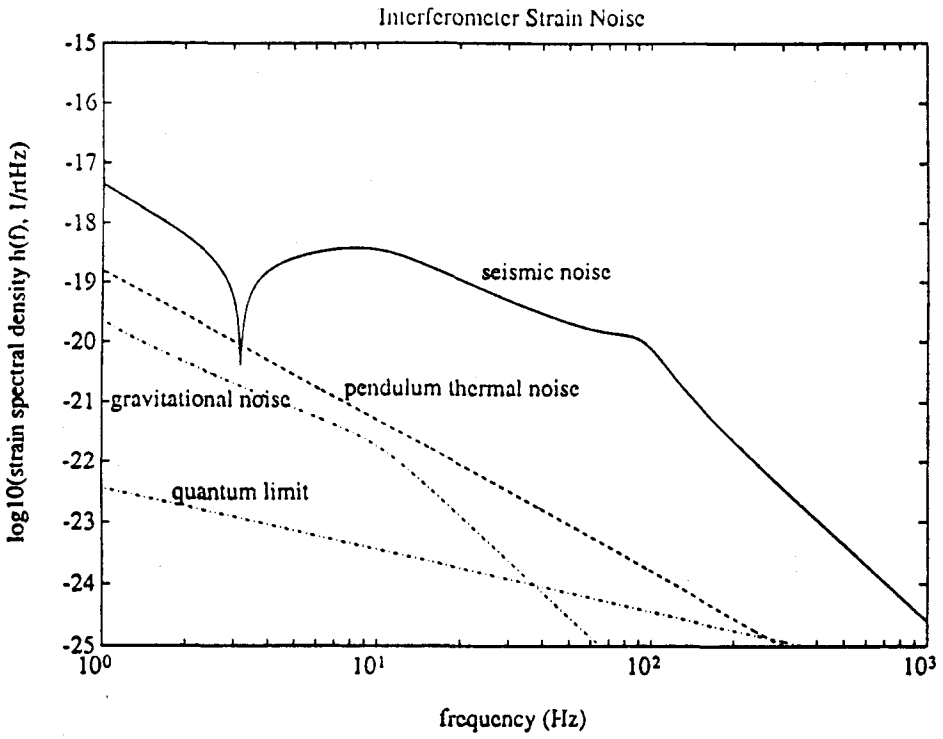

Figure 12: Low frequency noise budget for the modelled system. The seismic noise is attenuated by one active isolation stage of the sort described in the text, and by a suspension point interferometer giving an attenuation of 10^3. No other isolation is assumed.

$1 \times 10^{-13}\mathrm{cm}/\sqrt{\mathrm{Hz}}$, were cascaded from the first,[15] the noise budget would be that shown in Figure 13. The noise from the seismic isolation system is dominated by the thermal noise of the 200 kg proof mass in the second stage (attenuated by the suspension point interferometer). This noise level is comparable to the gravitational noise on the test masses. Both are about an order of magnitude below the predicted thermal noise from the test mass suspension.

Acknowledgements

I thank the Joint Institute for Laboratory Astrophysics for the award of a Visiting Fellowship, during which much of this work was done. Peter Bender, James Faller, and Robin Stebbins proposed the strategy for the application of active isolation to an interferometric antenna. The detailed analysis was carried out in consultation with Peter Nelson and David Newell.

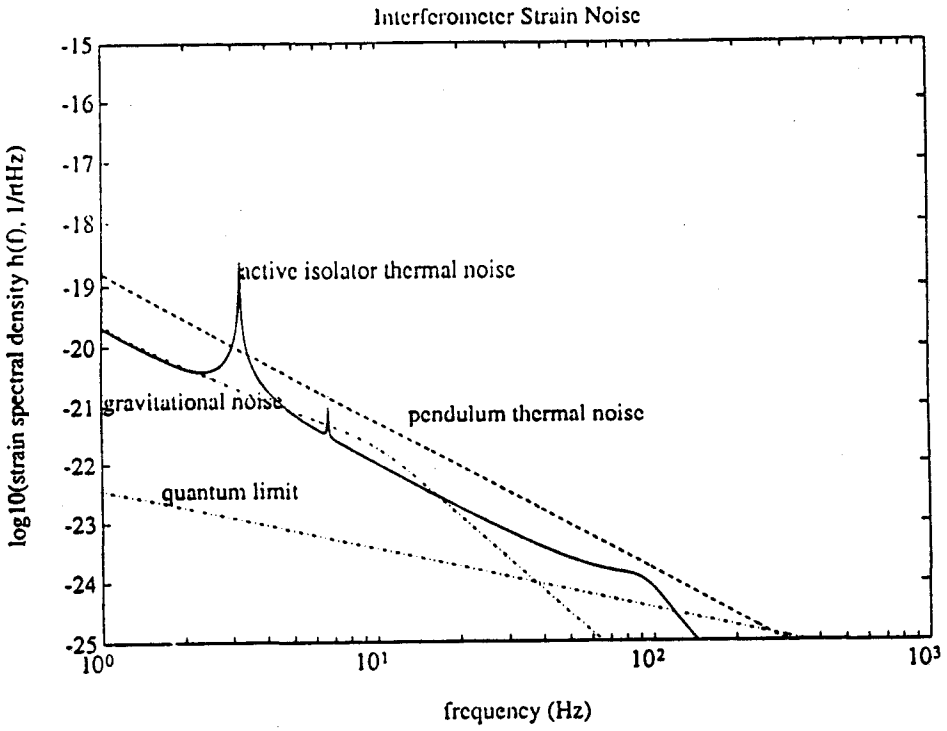

Figure 13: Low frequency noise budget, if two stages of active isolation were used. The pendulum thermal noise could be reduced about an order of magnitude if fused quartz fibers were used instead of tungsten wires.

References

1. C.M. Caves, *Phys. Rev. D* **23**, 1693 (1981).
2. C.M. Caves, in *Quantum Measurement and Chaos*, ed. E.R. Pike and S. Sarkar, (Plenum, New York, 1987), p. 195.
3. R. Spero, in *Science Underground*, Proceedings of the Los Alamos Conference, 1982, ed. M. M. Nieto *et al.* (AIP, New York, 1983), p.347.
4. P.R. Saulson, *Phys. Rev. D* **30**, 732 (1984).
5. H.B. Callen and R.F. Greene, *Phys. Rev.* **86**, 702 (1952); H.B. Callen and T.A. Welton, *ibid.* **83**, 34 (1951).
6. P.R. Saulson, *Phys. Rev. D* **42**, 2437 (1990).
7. J. Kovalik and P.R. Saulson, paper in preparation.
8. R. Weiss, private communication. For derivation, see Ref. 6.
9. R. Weiss, private communication.
10. R. Del Fabbro, A. di Virgilio, A. Giazotto, H. Kautzky, V. Montelatici, and D. Passuello, *Phys. Lett. A* **132**, 237 (1988).

11. R.W.P. Drever, J. Hough, A.J. Munley, S.-A. Lee, R. Spero, S.E. Whitcomb, J. Pugh, G. Newton, B. Meers, E. Brooks III, and Y. Gursel, in *Quantum Optics, Experimental Gravity, and Measurement Theory*, ed. P. Meystere and M.O. Scully, (Plenum, New York, 1983), p. 503.
12. J.E. Faller and R.L. Rinker, *NBS Dimensions* **63**, 25 (1979).
13. N.A. Robertson, R.W.P. Drever, I. Kerr, and J. Hough, *J. Phys. E: Sci. Instr.* **15**, 1101 (1982).
14. P.R. Saulson, *Rev. Sci. Instr.* **55**, 1315 (1984).
15. R.T. Stebbins, P.L. Bender, J.E. Faller, D.B. Newell, and C.C. Speake, talk presented at the 12th International Conference on General Relativity and Gravitation, Boulder, CO USA (1989).
16. N.A. Robertson, Ph.D. thesis, University of Glasgow, 1981.
17. See, e.g., R.C. Dorf, *Modern Control Systems* (Addison-Wesley, Reading, MA, 1980).
18. B. Lange, *AIAA Journal* **2**, 1590 (1964).
19. See, e.g., J.M. Maciejowski, *Multivariable Feedback Design* (Addison-Wesley, Reading, MA, 1989).
20. See, e.g., J.E. Fix, *Bull. Seis. Soc. Am.* **62**, 1753 (1972).

"Low Frequency Noise in Gravitational Wave Interferometers," by Peter Saulson
(Submitted to *Gravitational Astronomy—Instrument Design and Astrophysical Prospects*)

Circulated by R. Vogt 1/9/91

PLEASE PASS ALONG PROMPTLY

L. Sievers
R. Spero ✓
→ M. Zucker ✓ 1/28
A. Abramovici ✓ 1/29
Y. Gürsel
F. Raab
R. Drever

C. Akutagawa (Science Library) -last