LIGO-T880005-00-R

REPORT ON SEISMIC ISOLATION AND SUSPENSION SYSTEMS

Introduction to the Issues and Status of Research at MIT and Caltech

Yekta Gürsel Fred Raab Peter Saulson
Robert Spero

28 October, 1988

1 Introduction

1.1 Goals of the Working Group

The working group will address the immediate and long-term issues in the design and construction of suspension and seismic isolation systems. For the short-term, the science teams need to plan the initial system to be installed in the MIT prototype, as well as the next upgrade to the Caltech prototype. We shall make recommendations as to how this work should proceed in the next several months. Of comparable urgency are aspects of the seismic isolation that relate to LIGO designs. We plan to provide guidance to the engineering team on design issues related to seismic isolation, and to survey the isolation options for LIGO detectors.

Table 1: Background Seismic Noise and LIGO Target Sensitivity Levels. Units: meters/$\sqrt{Hz}$.

f (Hz)	Background	Early LIGO	2/88 Target	Later LIGO
10	10^{-9}	—	10^{-16}	10^{-20}
100	10^{-11}	$3 \cdot 10^{-17}$	$3 \cdot 10^{-19}$	10^{-21}
1000	10^{-13}	10^{-18}	10^{-19}	10^{-21}

1.2 Order-of-Magnitude Requirements of Suspension System

The test masses in all interferometers built or planned are roughly "free" in that the final stage of the suspension system—usually one or more thin wires supporting a mass—is a pendulum oscillator with a fundamental frequency (around 1 Hz) much lower than the band of gravity waves being searched. Though free in this sense, great care is required to set up test masses that serve as inertial platforms for the detection of gravitational waves: the ultimate suspension point connects to the ground, which, on the required scale of sensitivity, is a poor approximation to an inertial frame.

Table 1 shows order-of-magnitude estimates for the seismic noise background and the amount of that noise allowed to reach the test masses for three different levels of detector sensitivity. As a rule, the low-frequency sensitivity levels are set by seismic noise, and the high-frequency entries are set by shot noise. The "Early LIGO" and "Later LIGO" entries are from the December 1987 proposal, Figures A-4a and A-4b. The "2/88 Target" entries are the middle-ground sensitivity levels that were presented to the NSF during the February 1988 site visit.

We emphasize that the uncertainties are large—the Background column is based on some rough measurements made at Edwards, combined with a guess of the frequency-dependence, and the roughness of the sensitivity estimates is indicated by their span (three or four orders of magnitude at each frequency between the extremes). The "Later LIGO" entries, which assume the shot noise can be driven down by periodic recycling, impose the most stringent requirements on seismic isolation.

1.3 The Low-frequency Cutoff

As the frequency increases, the background seismic noise falls rapidly and, for a given isolation system, the fraction of that background penetrating the isolation to the test masses also falls as a high power of the frequency. These effects combine to make a sharp frequency cutoff, below which signals are unobservable. At just higher frequencies, other noise sources (such as shot noise) dominate. The cutoff frequency is a figure of merit for seismic isolation systems. How low should it be?

There is no practical low-frequency cutoff to the spectrum of gravitational waves. Spacecraft experiments tend to work best for periods of several minutes to many days, and pulsar timing measurements are sensitive to waves of still much lower frequency. The lower the frequency, the stronger the signals tend, so at any stage of LIGO development we stand to gain by improving the seismic isolation until some other low-frequency source of noise dominates.

1.4 Competition to Seismic Noise: How Much Isolation is Enough?

Potential noise sources competing with seismic noise at low frequency are gravity gradients, the standard quantum limit, and thermal noise. Noise from fluctuating gravitational gradients sets a natural low-frequency cutoff around 1 Hz, but with moderate precautions gravitational gradients can be reduced to insignificant levels above 10 Hz. Noise from the standard quantum limit is perhaps evadable, though a practical scheme that works for broad-band detectors has not yet been imagined. Nevertheless, with masses on the order of one ton SQL noise is unlikely to dominate except in the most advanced detectors.

A more serious competitor to seismic noise at low frequency is thermal noise. The mechanical Q required to reduce thermal noise level to the amplitude spectral density displacement sensitivity $\tilde{x}(f)$, at frequency ω is $Q \sim 4kT\omega_0/m\omega^4\tilde{x}^2(f)$. At 10 Hz, the corresponding "Later LIGO" requirement is $Q > 10^{11}$. To maintain this extraordinarily high Q requires a vacuum level of approximately 10^{-8} torr to reduce the fluctuations in the Brownian motion of one-ton test masses with density of 10 gm/cm^3.

More imposing is the tolerable loss due to bending of the suspension wires. A simple order of magnitude shows that the inherent Q of the suspension wire can be lower than the pendulum Q by $\sim 10^4$, assuming a wire of at least 1 meter in length. Even so, a material with room-temperature Q of 10^7 may be hard to come by. Our estimates of likely achievable thermal noise limits, to be provided in the sequel, will be an important parameter in the specification of seismic isolation.

1.5 Vertical Isolation

Input seismic noise is approximately equal for vertical and horizontal motions. Vertical motion of the test masses does not directly couple to the gravity wave signal, but this is compensated by the anisotropy of pendulum suspensions, which provide good horizontal isolation and no vertical isolation below their violin resonance of several hundred Hz. An optimized design has roughly equal contributions to the detector noise from vertical and horizontal seismic motion.

What is the ratio of vertical to horizontal motion of the test mass for equal noise contributions? The answer depends on the mechanism that couples vertical noise into the interferometer; an upper limit probably comes from the departure from level of the terrain. A site might have, say, a 5-meter difference in height relative to the horizontal for test masses separated by 5 km; this implies that the vertical motion should be less than 1000 times the horizontal motion if it is not to dominate. Other couplings, such as optical effects of vertical motion, might make the acceptable vertical motion smaller still.

1.6 Distribution of Isolation

Five distinct types of isolation are under consideration

- Isolation external to the vacuum chambers.
- Low-Q passive systems, such as stages based on rubber
- High-Q passive systems, such as from coil springs

- Active systems, such as magnetic levitation or auxiliary interferometers
- Pendulum suspensions of test masses

Proposed designs use some or all of these methods together. A full analysis is necessary to determine the best distribution of the different types of isolation, including the order for different stages for best overall performance.

1.7 How Much Seismic Isolation is Achievable?

The most difficult isolation requirement in Table 1 is at 10 Hz, where an isolation factor of 10^{11} is called for in the most sensitive detectors. Giazotto and company (RSI **59**, p 292, 1988) claim this factor is achievable, and they even describe a system they believe will do the job. Although there are no experimental demonstrations of such isolation at low frequencies, there are fairly solid data at high frequencies. The seismic noise in the Caltech detector above 500 Hz has been measured to be less than 10^{-19} meters/$\sqrt{\text{Hz}}$, and the same suspension at a remote site would probably exhibit considerably less seismic noise. We're comfortably below the Early LIGO level, and even approaching the Later LIGO level at 1 kHz. The region between 10 Hz and 500 Hz is unexplored territory, and that is where we need concentrate our efforts.

2 Recent Vibration Isolation Research at MIT

2.1 Nested Double Pendulum

The largest effort in suspension and vibration isolation research at MIT has centered on development of the nested double pendulum test mass suspension. The benefits of this scheme over a single pendulum are isolation proportional to the negative fourth power of frequency, and the provision of a quiet platform from which to apply the feedback forces to the test mass. The cost comes in extra mechanical complexity, in particular in doubling the number of mechanical degrees of freedom which need servo damping.

We have built a test model of a nested double pendulum. It is completely instrumented for damping and control. We have used the equations of motion of the system to compute a number of mechanical transfer functions, and verified the transfer functions by actual measurements on the system. We have investigated different schemes for damping the suspension, and have found that all of the normal modes are controllable in stable loops.

The very next priority for the nested double pendulum is to perform a series of tests in vacuum. We will learn, first of all, whether in fact all of the components function satisfactorily in vacuum (as we expect they will). We will also learn about the total outgassing rate, and make sure there are no unforeseen electrical interference problems. The most interesting tests will be measurements of the vibration transfer function (i.e. seismic isolation), and reconfirmation of the damping servo operation under conditions where we expect that the mechanical Q's of all the resonances will be high. Measurement of Q's is also important for predicting the thermal noise on the test mass. Finally, we will test the electrostatic system for applying feedback to the test mass, in particular investigating schemes for minimizing the excitation of mirror tilt by the feedback forces.

It is only when we have a suspended interferometer operating with good performance that the suspensions can be fully tested. There is no substitute for transfer function measurements performed with the interferometer. Furthermore, measurement of other noise terms besides transmitted vibration are only possible at position sensitivity levels comparable to LIGO sensitivity. For this last reason, it will be important to use whatever combination of input laser power, recycling, and storage time is required to measure the position noise spectrum.

2.2 Magnetic suspension

A parallel effort has gone into the development of a magnetic spring vibration isolation system to supplement a pendulum suspension. A magnetic spring should be able to give the isolation of a very soft massless spring, or in other words have a very low lumped body resonant frequency without compromise of the isolation by internal resonances. The cost is, again, the complexity of a servo system, since a magnetic suspension is always unstable in one (or more) degrees of freedom, in our case the vertical.

We have built a full-size test model magnetic suspension capable of supporting one of our in-vacuum optical tables and the attached components.

The floating part has been successfully suspended, and some mechanical tests have been made.

The magnetic suspension is not yet in a state where it should be a very good isolator. The vertical dynamic range is small, due to a large negative spring constant in the Alnico magnets used. We are in the process of replacing the magnets with new ones which use Samarium Cobalt. The hope is that this will allow lower gain in the servo loop which controls the height of the platform, yielding a softer, less-damped spring.

Once the new magnets are installed and some preliminary tests are performed, we plan to put this line of research on hold. This isn't because we don't believe in the long term promise of this scheme, but is simply a matter of allocating limited resources to highest priority items (tests with the stationary interferometer, along with suspensions and other parts for the suspended 5-meter interferometer.)

2.3 Coil Springs

Even with a double pendulum, additional stages of isolation would be extremely helpful. With a single pendulum, extra stages are more important still. In the 40-meter interferometer, this role is played by rubber springs, but their outgassing rate is too high for LIGO. In parallel with the study of the use of bellows to solve this problem, we are investigating the use of standard coil springs as isolators. The attraction is that they are completely passive and already engineered. The drawback is that their damping is low, which is especially a problem considering that high Q characterizes their internal resonances as well, compromising the high frequency isolation.

Nevertheless, the simplicity of metal springs is appealing. The current test model of the nested double pendulum uses coil springs for vertical isolation of the outer stage. We recently measured (in air) the transfer function of a mass hung from one of these springs. It looks like a good (better than 55 dB) isolator in this simple test. A number of internal resonances are visible at 370 Hz intervals, but their Q declines substantially with frequency. Although by no means a clean bill of health for these simple springs, it is encouraging enough for further tests in vacuum. We will test them as part of the nested double pendulum, and also make use of them in a quick and dirty isolation system for the table on which the stationary interferometer is built.

3 Vibration Isolation Research at Caltech

3.1 Introduction

The 40 meter prototype Laser Interferometric Gravitational Wave Detector at Caltech relies on multiple stacks of rubber and lead for its first stage of seismic isolation. An additional stage of isolation is obtained by suspending the test masses from the top plate of the rubber-lead stacks. This arrangement works very well for the prototype. However, the same design can not be used in LIGO receivers because of the stringent requirements on the level of outgassing that is allowed in the chambers. The purpose of this section is to give an overview of the past experiences in the laboratory in using rubber and metal springs as vibration isolators and to describe planned tests of the newer designs which are candidates for vibration isolators in LIGO receivers.

3.2 Stacks with Rubber Exposed to Vacuum

The current vibration isolation stacks in the 40 meter prototype use rubber erasers that are exposed to vacuum. These were chosen for their low spring constants and their "acceptable" outgassing rates. Tests were performed on them to determine their outgassing rate by an undergraduate student, and he found that their outgassing rate was better than other kinds of rubber which were as soft as these. The erasers were then adapted as the spring part of the vibration isolation stacks.

The performance of the stacks is still about the same after being in vacuum for approximately five years. However, some unexpected things surfaced after the stacks had been exposed to vacuum for an extended period of time. A viscous fluid started to ooze out of the erasers and collected on the supports and on the central beam splitter mass in small pools. We guess that this fluid is a remnant of the solvent that was used in casting the erasers. This is certainly unacceptable for any decent vacuum system.

One might ask the question why the elastic materials that are known for their "good" behavior in vacuum were not used initially. The answer is that such materials tend to have a rather large spring constant which makes them useless as vibration isolators. Recently we started experimenting with solvent free rubbers such as two-part RTVs. It is possible to make very soft rubber compounds by adjusting the ratio of the two parts that are mixed together. Since the silicone rubber compounds are resistant to degradation at relatively high temperatures, it is possible to pre-bake these in order to reduce their outgassing rate and to take out any remaining uncured rubber compound.

We installed a vibration isolation stack using one of these rubber compounds as a vibration isolator for our passive mode cleaning cavity. After being exposed to vacuum for several months the rubber shows no sign of anything oozing out of it. However, the outgassing rate is still too high to meet the stringent requirements for the LIGO receivers.

Much more research is needed to determine whether an exposed rubber system is suitable as a part of the vibration isolation system in the receiver design for LIGO. The properties of elastic materials which are soft enough to make a good vibration isolator have to be examined after they are exposed to vacuum for an extended period of time. Because of this, such research is likely to take a very long time and the results from it may not be available in time to be considered as a part of the first receiver design.

3.3 Encapsulated Rubber Springs

A solution to the problems mentioned above is to encapsulate the rubber in some "material" which prevents the vapors and the liquids escaping from the rubber compound from contaminating the high-vacuum chambers. There are several proposed methods to accomplish this feat, although none of these methods have been tested extensively and none are "proven" to work. In the following sections, We will examine some of these methods.

3.4 "Vacuum" Encapsulated Rubber

This idea is first introduced by Ernie J. Franzgrote of LIGO staff. It proposes to put the rubber compounds in separate chambers which are not completely sealed from the main high-vacuum chambers. The tops and the bottoms of the chambers holding the rubber compound are attached to the masses that form the massive part of the stack while the rubber is squeezed in between to provide the spring action. The tops and the bottoms however, are detached from each other by a small gap which allows the rubber to be compressed. The idea is to pump out the small chambers that are housing the rubber compounds separately from the main chambers at a higher pumping speed. This prevents the vapor escaping from the rubber from entering the main chambers.

Although it is quite possible to arrange the pumping lines for a multi-stage stack in such a way that they do not short-circuit the isolation provided by the rubber, this method suffers from the small dynamic range of the stack which is caused by the necessity of having a small gap in order to achieve a high pumping speed with a reasonably fast pump and reasonably sized pumping lines. Also, in the case of a pump failure, the material escaping from the rubber can potentially contaminate the main chambers. There are currently no plans to test this method, but this may change in the future.

3.5 Sealed Bellows Encapsulated Rubber

Another scheme to stop the rubber from contaminating high-vacuum chambers involves encapsulating the rubber in "soft" metal bellows. This is first introduced by Ron Drever.

The idea is that the metal bellows can be made much softer than the rubber to minimize the possibility of the metal short-circuiting the isolation provided by the rubber. The encapsulation is done in the following manner: The rubber is placed inside the metal bellows, the bellows is coated on the inside with a compound that is designed to damp the vibration of the bellows. The bellows is then sealed and evacuated. The evacuation may be necessary to prevent the air inside the bellows from short-circuiting the isolation provided by the rubber against relatively high-frequency sound waves. This whole assembly is then used as the spring in a vibration isolation stack.

There are several problems with this approach: First of all, the bellows can not be arbitrarily soft since it has to withstand a differential pressure of one atmosphere between inside and outside. The sound waves can travel through the metal of the bellows in the longitudinal mode which will not be damped by the coating inside the bellows very well. Since the bellows is sealed, the height of the stack will change as the vacuum chamber is brought back to atmospheric pressure.

The softest metal bellows that can withstand the pressure difference of one atmosphere are welded bellows which are made out of stainless steel. Because of the method of construction these bellows are much more expensive than the ordinary corrugated bellows. They are also very stiff in the "shear" direction. In order to get good isolation in the horizontal direction the bellows must be made to tilt as the mass they are carrying moves in the horizontal direction. This can be accomplished by installing flexure mounts between the load and the top of the bellows.

We have a welded, stainless steel bellows which has the appropriate dimensions to be used in the vibration isolation stacks for our 6 foot chamber and we may test this way of encapsulating rubber. The testing procedure involves constructing the bellows encapsulated rubber spring as described above, and actively measuring its transfer function under load in vacuum. It seems at first sight that it is difficult to load a single bellows at the loading rates required without it becoming unstable. The solution to this problem is to use a loading platform that is supported by three exposed rubber springs and measure its transfer function. Then, one repeats the measurement with one of the rubber springs replaced by the bellows encapsulated rubber spring.

3.6 Vented Bellows Encapsulated Rubber

One way to overcome some of the problems associated with using sealed bellows encapsulated rubber is to use a "vented" bellows scheme in which the bellows do not have to withstand a pressure differential as large as an atmosphere. This can be arranged by pumping the bellows separately and synchronously with the main chamber with the aid of an active control system. The control system monitors the pressures in the chamber and in the bellows and adjusts the pumping (or venting) rate of the bellows to hold the pressure differential below a certain value.

Since the pressure differential is not large, the bellows can be made as thin as possible. This will improve the isolation provided by the bellows. This method also guarantees that the height of the stack will stay the same whether the chamber in vacuum or at atmospheric pressure. In a multi-stage system it is possible to arrange the pumping lines so that all of the bellows can be pumped (or vented) by a single line with an automatic valve outside the vacuum chamber without short-circuiting the isolation provided by the rubber inside. Since the bellows can be made out of very thin metal sheets, corrugated bellows can be used without compromising the low spring constant requirement. This will reduce the cost of the bellows significantly. Also, the bellows can be made very small without being very stiff. This may be needed for secondary isolation stacks on top of the primary ones.

The problem with this method is that it is very fragile. If the active control system fails, the bellows will burst contaminating the high-vacuum chambers in the process. Since bellows are very thin, they can easily be damaged by accidental mishandling during installation or modification of the receiver test masses. We may test this kind of encapsulation in the near future.

3.7 Compensated Bellows Encapsulated Rubber

Another way of sealing rubber which overcomes almost all of the problems mentioned above is to encase the rubber in a small chamber which has two relatively small bellows in a compensated arrangement. This idea is also due to Ron Drever.

A sufficiently thick, soft rubber disk with a hole at the center is placed in a chamber which has similar size holes at the top and at the bottom. A metal plate which has the same size as the disk is placed on top of the rubber disk. An adequately thick metal rod is welded to the plate at the center. The rod goes through the hole at the center of the rubber disk and protrudes through the holes at the top and at the bottom of the chamber. Open ends of two small, soft metal bellows are welded to the holes on the chamber respectively. The other ends of the bellows are welded to the rod.

The closed chamber thus formed isolates the bare rubber from the high-vacuum environment. The chamber sits on its bottom on a large mass with a hole bored into the metal so that the rod that protrudes from the bottom and the bellows attached to it are free to move in the hole. Another large mass hangs from a flexure mount which is attached to the rod that protrudes from the top of the chamber. These masses form the massive part of an isolation stack. The rubber encased in the sealed chamber between the masses is the spring part of the isolation stack.

Since the bellows are relatively small, it is possible to get bellows that are made out of thin metal sheets that can withstand a pressure difference of one atmosphere. These bellows are much softer than stainless steel bellows, and they are usually made of brass. The height of the stack does not change with varying external pressure since the bellows are attached in a compensated arrangement. The short length of metal rod gives some mechanical advantage to make the spring constant of the system lower in the horizontal direction.

The major problem with this system is that the length of the metal rod can not be too long if one wants to make the system stable against small perturbations while carrying a reasonable range of weight. One can make the rubber harder to prevent instability, but this increases the vertical resonance frequency of the system. Another way of making the system stable is to increase the diameter of the rubber disk. This may take up too much horizontal space in the vacuum chambers.

We have tested a bench-top system demonstrating the fact that one can get horizontal resonance frequencies as low as 1 Hertz with an associated vertical frequency of 5 Hertz without making the system unstable and without taking too much horizontal space. A further decrease in the horizontal frequency seems to be impractical because of stability requirements. This is our prime candidate for an encapsulated rubber spring at the moment and we are planning to test other parts of this system. One crucial measurement that is needed is to plot the transfer function of the bellows as a function of the frequency since they are the "weak link" in the encapsulating scheme which can short the isolation provided by the rubber.

3.8 Metallized Plastic Pouch Encapsulated Rubber

Another way of sealing rubber is to follow what the vacuum packers of food and other items do. They use a special material that constitutes a good vapor barrier which can also be sealed well. This material is made out of two laminated layers of plastic: Polyethylene and mylar. The polyethylene layer gives the property of making good seals since that plastic is well known for its stickyness. The mylar layer forms the vapor barrier. The object to be sealed is placed in a bag that is made out of this material with the polyethylene layer facing inside. Then the entire assembly is put in a vacuum chamber and the air is pumped out. The bag is then heat-sealed under vacuum.

If one uses this kind of a bag to encapsulate rubber, one is still putting material in the vacuum chambers that can potentially outgas and contaminate the high-vacuum chambers. One way to overcome this difficulty is to use the same material with a third layer made out of aluminum. There are two ways of putting an aluminum layer on this material. One of them is vapor depositing which gives a very thin layer of aluminum which may have a lot of pinholes. The other way is to laminate an aluminum foil over the mylar layer. The foil is substantially thicker than the vapor deposited layer. This may seem a good way of sealing rubber, but it is not free of problems by any means.

First of all, the bags have to withstand atmospheric pressure with a good vacuum inside. Since the walls of the bag is relatively thin, they will crease and make sharp corners around the edges of the object that is inside the bag. This substantially weakens the walls of the bag at those points. It will almost certainly crack the metal layer along those edges. If the pressure outside the pouch is cycled between a good vacuum and the atmospheric pressure, then the pouch will expand and contract and will eventually burst. A way around this weakness is to put yet another layer of heat resistant nylon or polyester paper outside the aluminum layer. Unfortunately, this brings one back to putting exposed plastics in vacuum.

Another serious problem is the noise generated by the laminate as it is bent and relaxed by the moving mass carried by the rubber it encapsulates. This is the familiar crackling sound that is heard when one squeezes mylar balloons which are made from a similar laminated material. This noise can leak through the seismic isolation system and reduce the sensitivity of the interferometer.

We will be getting samples of this material and we may test them. It is clear that much research is needed in this area before one can find a satisfactory solution. The material described above is an industry standard and it is readily available. It may be possible to custom make such a laminate from other metals and plastics which will have very good outgassing characteristics. However, the research is likely to take a long time and it may be too costly to pursue.

3.9 Metal Springs

The metal springs are used in the 40 meter prototype as the last stage of a vibration isolation stack for the passive mode cleaner cavity. This stack performs exceptionally well: The seismic noise is attenuated to a level where it is possible to observe the thermal excitations of the normal modes of the quartz tube that forms the cavity.

The reason for such a performance is that all the high frequency noise that can potentially travel through the spring gets attenuated by the rubber spring stack, the metal spring forms the last stage which gives a very low frequency isolation. We have no plans at this time to test metal spring vibration isolators.

4 Goals of Future Work

Our next report will describe a coordinated suspension research plan. The issues that need to be addressed include the following:

- We need to better characterize the candidate isolation systems. Which are likely to perform best over what frequency range? Some systems may suppress mechanical modes such as torsion, transverse and longitudinal excitations better than others. Similarly a system with very good low frequency isolation may have high acoustic transmission.

- We need to develop simple yet sufficiently accurate models, backed up by experimental data.

- We must identify what data are needed on currently open projects within the two science groups. We must also consider which other systems developed elsewhere should receive attention. One observation is that many groups which have working prototypes (both bars and interferometers) have rubber in them somewhere. If rubber is important, we ought to find out why it is special.

It will be worthwhile to consider planning for a succession of LIGO suspensions. The earliest is likely to be a very simple and robust system relying mainly on passive isolation. Improved performance might be sought later, when we will use more advanced techniques. Active systems are currently under development; some may be mature enough to use in the first LIGO receivers.