Thermal Noise in the Initial LIGO Interferometers

Thesis by

Aaron D. Gillespie

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

1995

(submitted April 26, 1995)

Acknowledgment

I thank my advisor, Dr. Frederick J. Raab, for making my studies at Caltech both rewarding and enjoyable. I am indebted to Fred for teaching me the experimental attitude required to solve problems, that a combination of persistence, thoroughness, and a good sense of humor are required to learn nature's secrets.

I thank the entire LIGO team for their encouragement. The 40-meter interferometer, where many of the measurements of this thesis were made, could not be maintained and operated without the combined efforts of LIGO's many scientists, engineers and technicians. In particular I thank Stan Whitcomb, who provided insights into just about every issue that I investigated, and Torrey Lyons, who has been both a good friend and also the unfortunate officemate who suffered through a first presentation of most of my bad ideas. I am also deeply indebted to Robbie Vogt, who protected me from all of the turmoil that surrounded the LIGO project during my tenure and allowed me to concentrate on science without worrying about the problems of politics, money, and management.

Finally, I thank my wife Ruth for providing the emotional support necessary to navigate the stresses of graduate school and for making me a very happy man.

This work was supported by the National Science Foundation cooperative agreement number PHY-9210038.

Abstract

Gravitational wave detectors capable of detecting broadband gravitational wave bursts with a strain amplitude sensitivity near 10^{-21} at frequencies around 100 Hz are currently under construction by the LIGO (Laser Interferometer Gravitational-wave Observatory) and VIRGO groups. One challenge facing these groups is how to detect the motion of the center of an inertial mass to a precision of 10^{-18} m when the mass consists of atoms each of which individually moves much more than that due to thermal energy. The uncertainty in the interferometer's measurement due to these thermal motions is called thermal noise. This thesis describes the thermal noise of the initial LIGO detectors.

The thermal noise was analyzed by modelling the normal modes of the test mass suspension system as harmonic oscillators with dissipation and applying the fluctuation dissipation theorem. The dissipation of all modes which contribute significant thermal noise to the interferometer was measured and from these measurements the total thermal noise was estimated.

The frequency dependence of the dissipation of the pendulum mode was characterized from measurements of the violin modes. A steel music wire suspension system was found to meet the goals of the initial LIGO detectors.

A mathematical technique was developed which relates the energy in each vibrational mode to the motion of the mirror surface measured by the interferometer. Modes with acoustic wavelengths greater than the laser beam spot size can contribute significant thermal noise to the interferometer measurements.

The dissipation of the test masses of LIGO's 40-m interferometer at Caltech was

investigated, and a technique for suspending and controlling the test masses which lowered the dissipation and met the thermal noise goals of the initial LIGO detector was developed. New test masses were installed in the 40-m interferometer resulting in improved noise performance.

The implications of thermal noise to detecting gravitational waves from inspiralling compact binaries was investigated. An optimal pendulum length for detecting these signals was found. It was shown that the narrow band thermally excited violin resonances could be efficiently filtered from the broadband gravitational wave signal.

Table of Contents

List of Tables

List of Figures

Chapter 1: Introduction

1.1 What are Gravitational Waves?

General relativity predicts that whenever a mass is accelerated, it emits energy in the form of gravitational waves. This process is analogous to the emission of electromagnetic waves by accelerated charges. Gravitational waves take the form of strains in space; that is if two inertial observers were at some distance from each other when a gravitational wave passed, they would observe the distance between them to stretch or shrink, and the magnitude of the stretching would be proportional to the distance between the observers.

The lowest order moment of a mass distribution that can change in an oscillatory way is the quadrupole moment. Oscillatory changes in the monopole moment are forbidden due to conservation of mass, and changes in the dipole moment are forbidden due to conservation of momentum. This is in contrast to electromagnetic radiation, which has a dipole as its lowest order term. Hence gravitational waves are quadrupolar in nature. If there were three inertial observers in a plane they would observe distances along one direction stretch while distances along the orthogonal direction would shrink and visa versa in response to a passing gravitational wave.

The magnitude of emitted gravitational waves can be estimated by the quadrupole formula:

$$h \approx \frac{G}{c^4 D} \frac{d^2 Q}{dt^2}. \qquad 1\text{-}1$$

h is the strain at the observer, and G, c, D, and Q are the gravitational constant, speed of light, distance between the source of the waves and the observer, and quadrupole moment.

In SI units, G/c^4 is 9×10^{-45} sec^2 kg^{-1} m^{-1}, so everyday motions (kilogram bodies moving meters on second time scales) would emit strains of order 10^{-44}. Clearly gravitational waves are not easily observable. To emit appreciable gravitational waves requires solar size masses being accelerated to a significant fraction of the speed of light in a fraction of a second. These types of violent events are not believed to occur everyday or even every few years in our galaxy; observing them on human time scales requires monitoring many galaxies. The expected strains from such events at intergalactic distances are less than of order 10^{-21}. The LIGO [1] (Laser Interferometer Gravitational-wave Observatory) and VIRGO [2] projects are building detectors to measure such strains. Their strategy is to monitor changes in the separation of test masses several kilometers apart to better than 10^{-18} m, or less than a thousandth of the diameter of a proton, in a broad bandwidth around 100 Hz. One of many challenges facing these groups is to figure out how to measure the position of the center of an inertial test mass to better than 10^{-18} m when the mass is made up of many atoms, each of which is moving much more than that due simply to its thermal energy. This challenge is the subject of this thesis.

1.2 Gravitational Wave Detectors

Three types of gravitational wave detectors have been proposed and are being significantly developed. The first to be developed was the resonant bar detector [3][4][5]. The bar detector, shown schematically in Figure 1-1, is simply a cylinder of aluminum suspended in a cryogenic vacuum system. All external disturbances are reduced as much as possible. A sensor monitors the amount of energy in the quadrupole vibrational mode(s) of

the bar. Any sudden change in the energy or phase of the bar could be attributed to a gravitational wave.

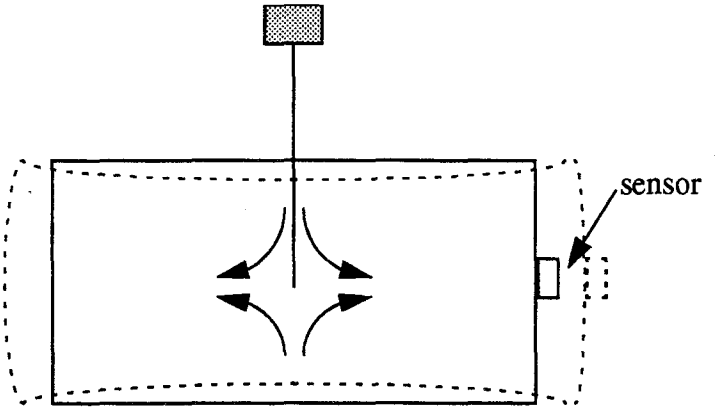

Figure 1-1: Schematic view of a resonant bar detector with a gravitational wave passing through.

The primary advantage of bar detectors is simplicity. Simplicity is an advantage that cannot be overestimated in a gravitational wave detector; a simpler system has fewer mechanisms by which external forces can affect the measurement. Another advantage is cost; bar detectors are cheaper than other types of gravitational wave detectors. An important feature of the bar detector is its resonant response, which increases the detector's sensitivity relative to a broadband detector, but gives little information about the gravitational wave beyond the amplitude of the wave in a particular bandwidth. The primary disadvantage is that increasing the dimensions of the bar changes its resonant response, so bar detectors cannot take advantage of long baselines to measure strain.

The type of gravitational wave detector that is the focus of this thesis is the laser interferometer detector. Such a detector is shown schematically in Figure 1-2. Four inertial test masses are suspended along orthogonal axes, so that a passing gravitational wave would cause the distance between the masses on one axis to increase and the distance along the

other axis to decrease, and visa versa. To extract information, the test masses have mirrored surfaces, and a laser interferometer is constructed around them so that relative motions of the mirrors are measured as a phase shift in the laser light at the beam splitter.

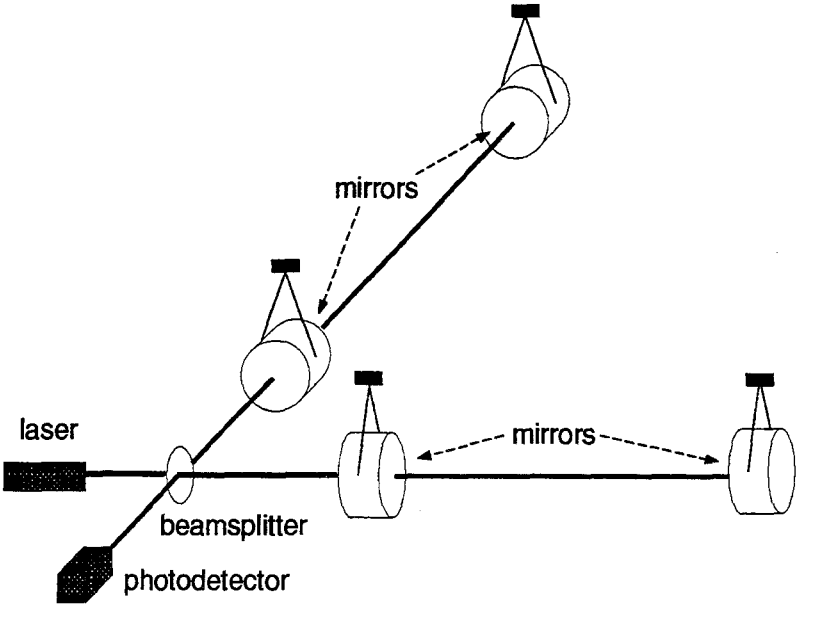

Figure 1-2: Schematic view of a laser interferometer gravitational wave detector.

The primary advantage of the laser interferometer is its inherently broadband response, allowing the extraction of the full waveform of the gravitational wave. Also, the interferometers can take advantage of long baselines to measure strain because as the distance between the mirrors, L, is increased, the interferometer becomes more sensitive to a given gravitational wave strain ($h = \Delta L/L$). The primary disadvantage of the laser interferometer detector is its complexity; many potential external sources of disturbance must be simultaneously considered in its design.

The last type of gravitational wave detector being pursued is a space based version of the laser interferometer, proposed by the LISA project [6]. The distances between three satellites in a heliocentric orbit would be carefully monitored. The advantages of a space

based system are that the baseline can be made tremendously long, leading to increased sensitivity, and that the satellites would be in continual free fall, so they would not have to be held by anything and hence would be less susceptible to external disturbances. The strong disadvantage of a space based detector is the tremendous cost of putting a such system into space.

1.3 Noise in Laser Interferometers

Laser interferometer gravitational wave detectors have many noise sources which could limit the sensitivity of the detector. The three primary noise sources are seismic noise, shot noise, and thermal noise, the focus of this thesis.

Seismic noise includes all noise which is mechanically transmitted from the external world to the mirrors. These vibrations include natural ground motion, acoustic vibrations of the vacuum system, and man made disturbances from objects as diverse as vacuum pumps and freight trains. To reduce the effects of seismic noise, vibration isolation systems are built between the vacuum system and the mirrors. Seismic noise is expected to limit the interferometer performance at low frequencies (below ~50 Hz).

Photon shot noise is noise due to the quantum uncertainty in the phase measurement of the light. In a simplistic model shot noise results from a finite number of photons hitting the photodetector in the measurement time. The effects of shot noise are reduced by using optical schemes designed to extract the maximum amount of information from each available photon. Also, to maximize the number of available photons, high power lasers and low loss mirrors are used. Shot noise is expected to limit the interferometer sensitivity at frequencies above 200 Hz

Thermal noise refers to all noise which is a consequence of the thermal motion of the atoms which make up the masses and their suspensions. In principle thermal noise could be reduced simply by cooling the masses; however technical difficulties make that option prohibitively expensive. In practice, thermal noise is controlled by careful experimental design. How this is done is described in detail in this thesis. Thermal noise is expected to limit the interferometer sensitivity in a bandwidth between about 50 and 200 Hz.

In addition to seismic, thermal, and shot noise there are many other potential sources of noise. These include everything from noise in the orientation of the mirrors to frequency noise in the laser to electronics noise in the control systems and many more. Every interferometer design choice must account for all of these noise sources. Although this thesis concentrates only on thermal noise, the design goals and choices are affected by other noise sources.

1.4 The LIGO Project

The LIGO project is currently building two 4 km long laser interferometer gravitational-wave detectors with the support of the U.S. National Science Foundation. In addition to designing and constructing the full scale interferometers, LIGO operates a 40 meter long interferometer on the campus of the California Institute of Technology. The 40-m interferometer serves as a test bed for ideas and designs that may be applied to the full scale interferometers. The thermal noise investigations of this thesis is one part of the research effort by the LIGO team of scientists and engineers to develop the technology to design the full scale interferometers. The 40-m interferometer is one of the primary research tools used for this work.

Most of the capital construction cost of the LIGO interferometers goes into building the vacuum system, and not into the interferometer itself (i.e., the laser, the vibration isolation system, the optics, etc.) For this reason the philosophy of the LIGO project has been to build a vacuum system which is as flexible as possible so that as technology improves, the interferometers can be upgraded to higher sensitivities. Therefore two types of research are carried out within the gravitational wave community. Work on the "initial" detectors is concerned with integrating presently available technology into a functioning interferometer with relatively modest performance goals. "Advanced" detector research investigates more aggressive techniques which have the promise of better performance. Once the initial interferometers are functioning properly, experience gained in operating them will be integrated with the advanced detector research programs to design more sensitive interferometers to be installed in the vacuum system. The work described in this thesis deals primarily with the initial detector, so it focuses more on building a working system with low thermal noise than developing techniques to achieve the lowest possible thermal noise. Where appropriate, how aspects of this research can be applied to the advanced detector research programs will be discussed.

1.5 Goals of this Work

When this work was begun in 1991, the field of thermal noise in laser interferometers gravitational wave detectors was in its infancy. Mechanical thermal noise was familiar to the gravitational wave community, for it had long since been detected in resonant bar detectors [7], and it was well known that thermal noise would ultimately limit laser interferometers. But most of the early research (and most of the present research, for that mat-

ter) on laser interferometers concentrated on understanding their subtle complexities and making them work. As the sensitivity of the small scale interferometers increased, so did interest in thermal noise. In 1990 Peter Saulson published "Thermal Noise in Mechanical Experiments," which laid the foundation for studies of laser interferometer thermal noise [8]. By 1991 several experimental groups were investigating various aspects of thermal noise [9][10]. However, no detailed systematic study of the thermal noise in an interferometer had been undertaken.

Two conditions existed at Caltech from 1991-1994 which together made the work of this thesis unique. The first condition was that thermal noise was measured in the 40-m interferometer at Caltech [11]. The second condition was that funding for the full scale LIGO interferometers had been approved and the interferometers were scheduled to be constructed on 5 to 7 year time scales.

These conditions made the goal of this work to be to design a suspension system that simultaneously met the thermal noise goals of the initial interferometer and could be fit into existing concepts of how the initial interferometers would be built and operated. To demonstrate that this goal had been accomplished, the candidate suspension system was installed in the 40-m interferometer and tested by observing the noise performance of that instrument. Meeting all of the thermal noise goals in one suspension system required developing an experimental program that systematically investigated every relevant mode of the suspension system. The goal was not to achieve the lowest possible noise in any one mode, but rather to achieve noise which was low enough in all modes for the initial LIGO interferometers. Specific attention was paid both to the details of how the suspension system would be assembled and installed in both the 40-m interferometer and the LIGO inter-

ferometer and to the problems of how control forces would be applied to the test masses.

The result of this work is a suspension system that has demonstrated acceptably low noise in a working interferometer. The remainder of this thesis describes the theory and experimental program behind the development of a candidate suspension system that meets the thermal noise goals of the initial LIGO interferometer.

Chapter 2: Mechanical Thermal Noise

When referring to a mechanical system, thermal noise generally includes any motion that results from the system being at a finite temperature. The temperature of a system is a measure of its thermal energy, i.e., the random energy contained in each degree of freedom. When the spacing of the quantum energy levels of a system is much less than the average energy of the modes of the system, the temperature, T, is

$$T = \frac{2\bar{E}}{k_B}. \qquad 2\text{-}1$$

$\bar{E}$ is the mean energy in each degree of freedom and k_B is Boltzman's constant. In SI units k_B is 1.38×10^{-23} joules/Kelvin.

Eq. 2-1 is commonly used to describe gases where the degrees of freedom are very easy to interpret. For example, in a monatomic gas each atom has three (translational) degrees of freedom, from which one can derive the heat capacity of the gas and the ideal gas law. Molecular gases have more degrees of freedom in vibrational and rotational modes, but can be analyzed in a analogous way to a monatomic gas.

Solids are likewise made up of atoms and their mean energy per mode can also be described by Eq. 2-1. However, due to the complicated interactions that occur between atoms in a solid, it is virtually impossible to analyze the degrees of freedom on an atom by atom basis for a macroscopic ($10^{>20}$ atoms) body. Instead, a statistical approach is used. The theory that is used is the theory of elasticity. Here the solid is modelled not as a collection of atoms of individual masses and binding strengths, but rather as a continuous medium of uniform density and some average distributed spring constant (per unit length).

The disadvantage of this simplification is that it breaks down on length scales over which the medium is no longer uniform and hence cannot predict all of the degrees of freedom. In order to use this simplification, it must be shown that the modes which are missed are unimportant.

The general strategy for analyzing the thermal noise of a mechanical system is to find all relevant degrees of freedom, or normal modes, and assign $\frac{1}{2}k_{\mathrm{B}}T$ of energy to each mode. $\frac{1}{2}k_{\mathrm{B}}T$ corresponds to much more motion in the interferometers than one might expect from a gravitational wave, so in order to have any hope of detecting gravitational waves, the energy must be packaged in frequency bands where gravitational waves will not be detected. The subject of this thesis is a study of how to predict the thermal energy distribution of the gravitational-wave detector's suspension systems and test masses, and how to build the suspension systems so as to minimize the impact of thermal noise on the detection of gravitational waves.

2.1 The Normal Modes

The first step in predicting the thermal noise is to figure out what the normal modes of the interferometer are and how they will affect the gravitational-wave measurement. The test masses of the interferometers are suspended as pendulums. It is convenient to divide the modes of the suspension system into two classes: suspension modes, which treat the test mass itself as an ideal rigid body and deal primarily with the motion of the center of mass of the mass as well as the motion of the suspending fibers, and internal vibrational modes, which generally treat the test mass as a free elastic body in which the center of

mass does not move, but all vibrations of the mass itself are included.

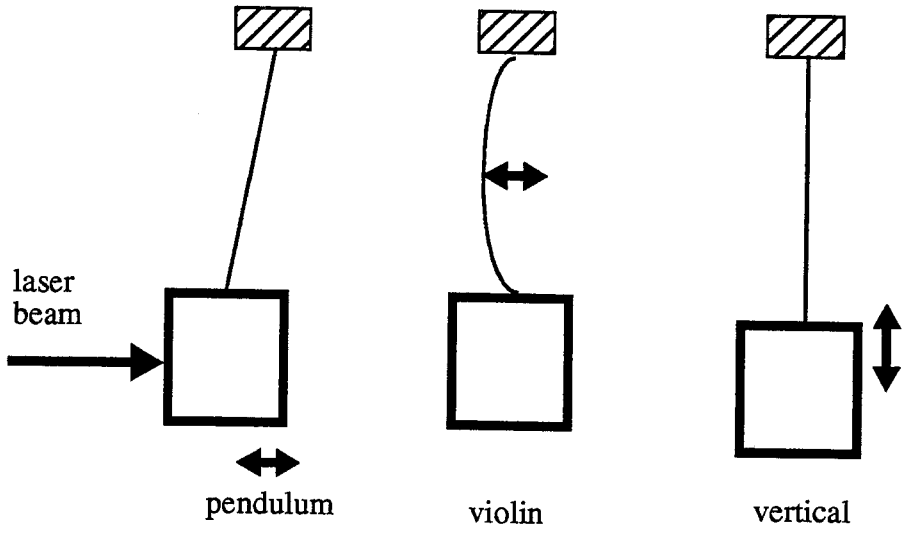

Figure 2-1: Schematic view of the suspension modes.

There are three important types of suspension modes, shown in Figure 2-1. The most important is the pendulum mode, in which the test mass swings back and forth. All of the motion of the pendulum mode is coupled directly to the interferometer. The second type of mode is the violin mode. Here the wire vibrates like a violin string. The test mass recoils only slightly from the wire due to the large mass difference (a factor of $\sim 10^5$) between the test mass and the wire, however, the geometry of the system puts much of the energy of the first violin mode in a bandwidth near several hundred Hz, the middle of the gravitational wave observing bandwidth. The violin modes form a harmonic series; only the first harmonic is shown in Figure 2-1. The last important (from a thermal noise standpoint) mode of the suspension system is the vertical spring mode in which the wire acts as a spring and the test mass bobs up and down. This mode couples to the interferometer signal only if the light propagating between the test masses is not level or the interferometer arms

are so long that the curvature of the earth makes the local vertical frames of the test mass not completely parallel (as will be the case for LIGO, which will have a coupling factor of about 10^{-3}). However, as will be seen in Chapter 4, the factors affecting the vertical mode's thermal noise are less favorable than those affecting the pendulum mode's thermal noise, so the vertical mode must be studied. In addition to these modes there are also torsional modes, rocking modes, and pendulum and violin modes perpendicular to the laser beam. These modes do not contribute significant thermal noise to the interferometer and will not be discussed here.

In order to find the internal vibrational modes of the test mass, the time dependent equations of elasticity must be explicitly solved [12]. Due to the three dimensional nature of the test mass and the complicated interactions between shear and bulk motions, there are quite a rich variety of mode shapes. Due to the variety of shapes, the interactions between the vibrational modes and the interferometer signals are rather complicated. These issues are discussed in detail in Chapter 5.

2.2 Energy in the Modes

The energy in each mode can be found using Eq. 2-1. The root-mean-squared motion, x_{rms} can be estimated by modelling each mode as a simple harmonic oscillator:

$$x_{\mathrm{rms}} \approx \sqrt{\frac{k_B T}{m\omega_0^2}}. \qquad 2\text{-}2$$

m is the mass of the oscillator and ω_0 is the angular resonant frequency. For the parameters of the 40-m interferometer (1.6 kg mass, 1 Hz resonant frequency, temperature of 300 K),

the rms motion of the pendulum is 8×10^{-12} m. For the violin modes (4×10^{-6} kg effective mass, 600 Hz resonant frequency) the rms motion of the center of the wire is also 8×10^{-12} m.[1] The recoil of the mass against the wire is 2×10^{-17} m.

Recall that the target sensitivity of the initial interferometers is an rms noise background of about 10^{-18} m for broadband bursts around 100 Hz. This motion is much less than the expected thermal motion of either the pendulum or the violin modes. Clearly, even a small fraction of the total thermal noise in the gravitational-wave bandwidth could affect the sensitivity of the interferometer. In order to predict the spectral density of the thermal energy and determine what parameters are important for minimizing the noise in a particular bandwidth, a more sophisticated model of the thermal energy must be used. That model is the fluctuation dissipation theorem.

2.3 Fluctuation Dissipation Theorem

The fluctuation dissipation theorem is a quantum-mechanically derived theorem describing the equilibrium level of any system with an external heat bath. The theorem depends on the principle of microscopic reversibility, that is, if some process can occur, then a corresponding time-reversed process also has a finite probability of occurring. If there is some mechanism for energy in a particular mode to leak out of that mode in the form of heat, then, by the same mechanism, energy can leak into that mode from the surrounding heat. The fluctuation dissipation theorem describes the equilibrium level of this

1. The agreement here is no coincidence. They are the same because the same gravitational force applies the restoring force for both the pendulum and violin modes. The fact that the 40-m interferometer has four wires gives the appropriate number of factors of 2 to make the thermal noise amplitude identical for both the violin wire in its first harmonic and the pendulum mode.

energy flow and hence the average energy in the system at thermal equilibrium.

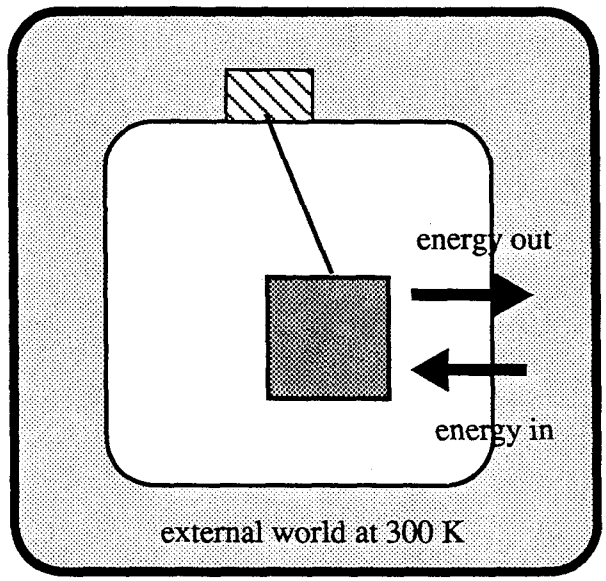

Figure 2-2: Schematic picture of the fluctuation dissipation theorem. The fluctuation dissipation theorem describes the equilibrium level between energy leaking out of a system in the form of heat and heat leaking into a system.

Mathematically, the fluctuation-dissipation theorem states [13]:

$$S_F(f) = 4k_B T \Re(Z) . \qquad 2\text{-}3$$

$S_F(f)$ is the spectral density of the generalized force acting on a system, and $\Re(Z)$ is the real part of the impedance of the system, Z. For a mechanical system, the generalized force is simply the mechanical force on the system and the impedance is the force divided by the velocity ($Z = F/v$). For an electrical system, the generalized force is the voltage and the impedance is the voltage divided by the current ($Z = V/I$). The most familiar form of thermal noise is the Johnson noise of a resistor with resistance R ($= V/I$):

$$S_V(f) = 4k_B TR . \qquad 2\text{-}4$$

The thermal noise discussed here is simply the mechanical analog of Johnson noise in a

resistor.

2.4 Dissipation of a Harmonic Oscillator

In order to use the fluctuation dissipation theorem to describe a simple harmonic oscillator, the basic model of the simple harmonic oscillator,

$$-m\omega^2\tilde{x} + k\tilde{x} = \tilde{F}, \qquad 2\text{-}5$$

must be modified to include dissipation. ω, k, $\tilde{x}$, and $\tilde{F}$ are the angular frequency, spring constant, Fourier transformed displacement, and Fourier transformed force. A general way to include dissipation is by allowing the spring constant to have a small imaginary part [14]:

$$k \rightarrow k\,[1 + i\varphi(\omega)]\,. \qquad 2\text{-}6$$

Then the equation of motion, Eq. 2-5, becomes

$$-m\omega^2\tilde{x} + k\,[1 + i\varphi(\omega)]\,\tilde{x} = \tilde{F}. \qquad 2\text{-}7$$

This new equation of motion has dissipation because it includes a component of the restoring force which is out of phase with the displacement.

$\varphi(\omega)$ is called the *loss function*. It is the fractional loss of energy of the system per radian of oscillation at a frequency ω. In order to model a large variety of loss mechanisms, $\varphi(\omega)$ is allowed to be an arbitrary function of ω. The specific nature of a particular loss mechanism is contained in the frequency dependence of $\varphi(\omega)$. In general, if the spring constant has a frequency dependent imaginary part, then, in order to be physically realizable, its real part must also have a frequency dependence according to the Kramers Kronig

relation. The magnitude of the deviations of the real part of the spring constant from the average value are of order the size of the imaginary part of the spring constant. For parameters relevant to gravitational-wave detectors ($\varphi(\omega) < 10^{-5}$), the frequency dependence of the real part of the spring constant is negligible, and will be ignored for the remainder of this thesis.

From Eq. 2-7 the impedance of the harmonic oscillator can be derived:

$$Z = \frac{k\varphi(\omega)}{\omega} + i\left(m\omega - \frac{k}{\omega}\right). \qquad 2\text{-}8$$

Combining Eq. 2-3, Eq. 2-7, and Eq. 2-8, the spectral density of the displacement of the mass due to thermal noise (*what the interferometer measures*!) can be derived [8]:

$$S_x(f) = \frac{4k_B T}{m\omega} \frac{\omega_0^2 \varphi(\omega)}{\left(\omega_0^2 - \omega^2\right)^2 + \omega^4 \varphi^2(\omega)}. \qquad 2\text{-}9$$

Eq. 2-9 is THE equation for predicting the thermal noise as seen by the interferometer. Here, the spectral density of the displacement is written as a function of the frequency, f, and is in the units (SI) of m^2/Hz. However, the function is written explicitly in terms of the angular frequency, ω. This notation, which is standard in the field, is perhaps initially confusing but simplifies the equation in that it eliminates many 2π's. The rms motion of the mass in the bandwidth $[f_1, f_2]$ is then

$$x_{\text{rms}} = \sqrt{\int_{f_1}^{f_2} S_x(f)\, df} = \sqrt{\int_{2\pi f_1}^{2\pi f_2} S_x(f) \frac{d\omega}{2\pi}}. \qquad 2\text{-}10$$

Most of the parameters of Eq. 2-9 are straightforward to measure or calculate. The temperature and the resonant frequency can be easily measured. For the pendulum mode

the mass can also be easily measured. For the violin modes and the vibrational modes of the test mass, the moving mass relevant to the interferometer is not quite so simple; however those masses can be calculated (see Section 4.3 for the violin modes and Section 5.1 for the vibrational modes). The remaining (and most challenging) parameter to be determined is $\varphi(\omega)$.

Chapter 3: The Loss Function

The primary challenge in understanding and predicting thermal noise is determining what the loss function, $\varphi(\omega)$, is. Once $\varphi(\omega)$ is known, it can simply be plugged into 2-9 to determine the thermal noise.

3.1 Measuring $\varphi(\omega)$

Measuring the loss function at the resonant frequency, $\varphi(\omega_0)$, is a simple matter. In the limit that $\varphi(\omega_0)$ is very small, the equation of the free oscillation becomes (from the equation of motion, Eq. 2-7):

$$x(t) \approx x_0 e^{\pm i\omega_0 t} \exp\left(-\frac{\varphi(\omega_0)\,\omega_0 t}{2}\right). \qquad 3\text{-}1$$

From this equation, it can be seen that a simple way to measure $\varphi(\omega_0)$ is to excite the resonance, turn off the excitation to create the conditions of a free oscillation, and measure the amplitude decay time, τ_A, of the free oscillation. $\varphi(\omega_0)$ can then be computed by

$$\varphi(\omega_0) = \frac{2}{\omega_0 \tau_A} = \frac{1}{Q}. \qquad 3\text{-}2$$

$\varphi(\omega_0)$ is related to the mechanical quality factor, or Q, of the resonance as shown.

Measuring $\varphi(\omega_0)$ gives little information about the general frequency dependence of $\varphi(\omega)$. To directly determine $\varphi(\omega)$, one must measure the phase shift, θ, between a driving force on the oscillator and the amplitude of the oscillation at the frequency of interest, ω. $\varphi(\omega)$ can then be calculated using Eq. 2-7:

$$\varphi(\omega) = \theta(\omega) \qquad \omega \ll \omega_0$$
$$\varphi(\omega) = \pi - \frac{\omega^2}{\omega_0^2}\theta(\omega) \qquad \omega \gg \omega_0 \qquad 3\text{-}3$$

Such a technique can been used to measure $\varphi(\omega)$ below the resonant frequency for a relatively high loss ($\varphi(\omega) \sim 10^{-3}$) system [15], however, the expected loss functions relevant to a gravitational wave detector are much lower. For example, the expected loss function of the pendulum mode near 100 Hz is of order 10^{-5} (to be discussed in Section 4.7). With a resonant frequency of 1 Hz, the phase shift that would have to be measured to determine φ(100 Hz) directly would be (π - 10^{-9}). Such a measurement does not appear to be possible.

For a multimode system, sometimes the losses for one mode can be predicted from the losses of other modes. This can be done if a relationship can be found between the damping forces of two or more modes. By measuring the losses on resonance, $\varphi(\omega_0)$, for one mode, the loss function at that frequency can be calculated for the other modes. Hence the loss function for an individual mode can be estimated by measuring the $\varphi(\omega_0)$ of all the other modes, computing the loss function at those frequencies for the particular mode, and then interpolating between the resonances to determine the general frequency dependence of $\varphi(\omega)$. Such a technique is used to estimate $\varphi(\omega)$ of the pendulum mode and the violin modes (Section 4.2).

3.2 Sources of $\varphi(\omega)$

There are certain types of loss mechanisms affecting gravitational wave detectors

from which $\varphi(\omega)$ can be predicted by measuring other macroscopic quantities of the system. Usually the general frequency dependence of $\varphi(\omega)$ can be determined simply by knowing the loss mechanism. Also, when the loss mechanism and the parameters that affect it are known, careful experimental design can be used to reduce the effects of those loss mechanisms. Here four loss mechanisms which can affect the LIGO interferometers are discussed.

Residual Gas Damping

Whenever a mass moves through a gas, it exchanges energy with the gas with the net result being to heat the gas and reduce the energy of the mass. The effects of this damping come in two very different forms depending on the number density of the gas and the distances moved by the mass. In the fluid regime, the mass moves distances which are much greater than the mean free path of the gas molecules. In this case, the gas must be treated as a fluid and the formulations of fluid dynamics must be used to solve for the damping. This regime applies to pendulums moving macroscopic distances in a gas at atmospheric pressure; however it does not apply to the pendulums of gravitational wave detectors which move microscopic distances in a vacuum. In this case, each molecule can be treated as individually making elastic collisions with the mass. The damping is found by calculating the average difference in momentum transfer between molecules hitting the front of the moving mass and molecules hitting the back. The loss function of the mass is then given by:

$$\varphi(\omega) \approx \frac{2A_{\perp}P}{m}\sqrt{\frac{m_{\mathrm{gas}}}{k_{\mathrm{B}}T}}\frac{\omega}{\omega_0^2}. \qquad 3\text{-}4$$

$A_\perp$, P, and m_{gas} are the surface area of the mass perpendicular to the direction of motion, the pressure of the gas, and the mass of the gas molecules. The magnitude of this damping mechanism can be controlled simply by adjusting the pressure of the vacuum system (i.e. by controlling what materials are allowed in the vacuum system and by buying vacuum pumps with sufficient pump speeds. For parameters relevant to the initial LIGO interferometer, $\varphi(\omega) \approx 10^{-10}\omega \times (P/10^{-6}$ Torr), putting moderate restrictions on the vacuum system. Optical cleanliness and pressure fluctuations in the beam path put more stringent requirements on the vacuum system. The frequency dependence of the loss function, $\varphi(\omega) \propto \omega$, is sometimes referred to as "viscous damping."

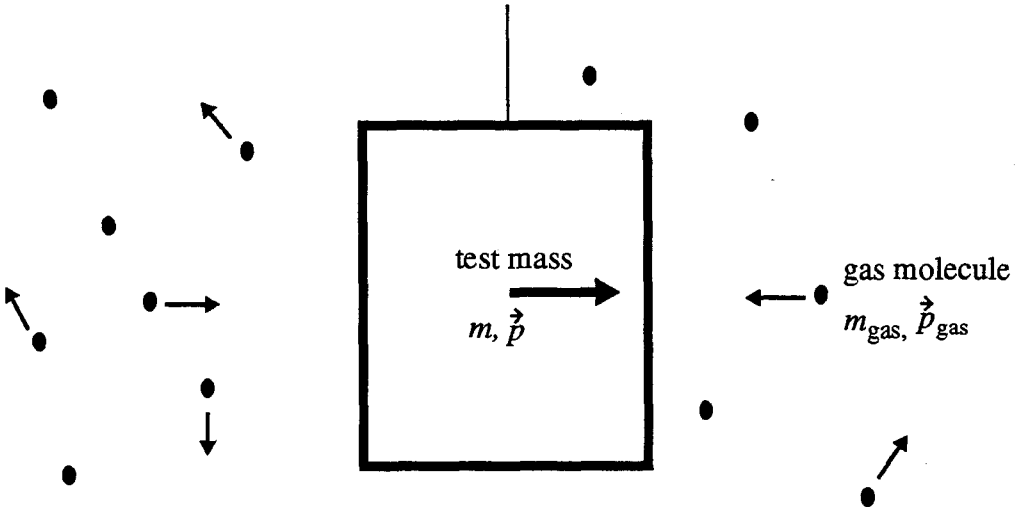

Figure 3-1: Schematic view of the test mass with residual gas to illustrate mechanism of residual gas damping. There is a larger momentum difference between the test mass and molecules striking the front of the test mass than between the test mass and molecules striking the back of the test mass resulting in a net damping force.

Eddy current damping

The test masses of the initial LIGO and the 40-m interferometers have magnets attached to them to exert control forces. When those magnets move, they induce currents

in any surrounding conductors. If those conductors have a resistance, energy from the motion of the magnets is dissipated in the conductors. This loss mechanism is referred to as eddy current damping.

The effects of eddy current damping can be seen most easily for the simple case of the damping in the circuitry that applies the control forces to the test mass. If the output circuitry is purely resistive, then the losses are given by:

$$\varphi(\omega) = \frac{\gamma^2 \omega}{R\omega_0^2}. \qquad 3\text{-}5$$

R is the output impedance of the circuit and γ is the ratio of the force on the test mass to the current in the coil. R and γ are parameters which depend on the design of the circuit, the number of turns in the coil and the size of the magnet. They must be chosen as a compromise between thermal (and electrical) noise requirements and the servo dynamic range. Like residual gas damping, eddy current damping in a purely resistive medium is viscous in nature.

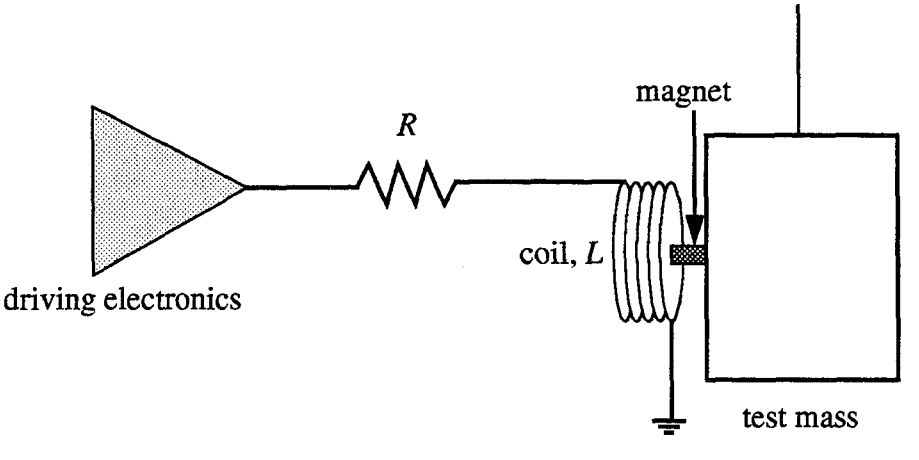

Figure 3-2: Schematic view of test mass with magnet and driving electronics to illustrate the relevant parameters of eddy current damping. Motion of the magnet results in a current through the coil which dissipates energy in the resistor.

The output impedance of the driving electronics need not be purely resistive. In fact the coil will have some inductance, and additional inductance can be designed into the circuit. If the output inductance of the circuit is L, the loss function becomes:

$$\varphi(\omega) = \frac{\gamma^2 R\omega}{\omega_0^2\left[R^2 + \omega^2 L^2\right]}. \qquad 3\text{-}6$$

The frequency dependence of $\varphi(\omega)$ is not viscous, but rather has a peaked frequency dependence with maximum of the damping function at $\omega = R/L$. The frequency dependence of Eq. 3-6 is quite common for damping mechanisms which have a characteristic time scale, τ, and a characteristic strength, Δ. In the general case, Eq. 3-6 can be written as

$$\varphi(\omega) = \Delta\frac{\tau\omega}{1 + (\tau\omega)^2}. \qquad 3\text{-}7$$

The shape of Eq. 3-7 is given in Figure 3-3. For the specific case of eddy current damping through an inductive load, $\Delta = \gamma^2/\omega_0{}^2 L$ and $\tau = L/R$. At low frequencies ($\omega << 1/\tau$), very little electromotive force is induced by the moving magnets, so there is only a small amount of current and dissipation. At large frequencies ($\omega >> 1/\tau$), the inductive load effectively blocks the current from flowing so there is also little dissipation.

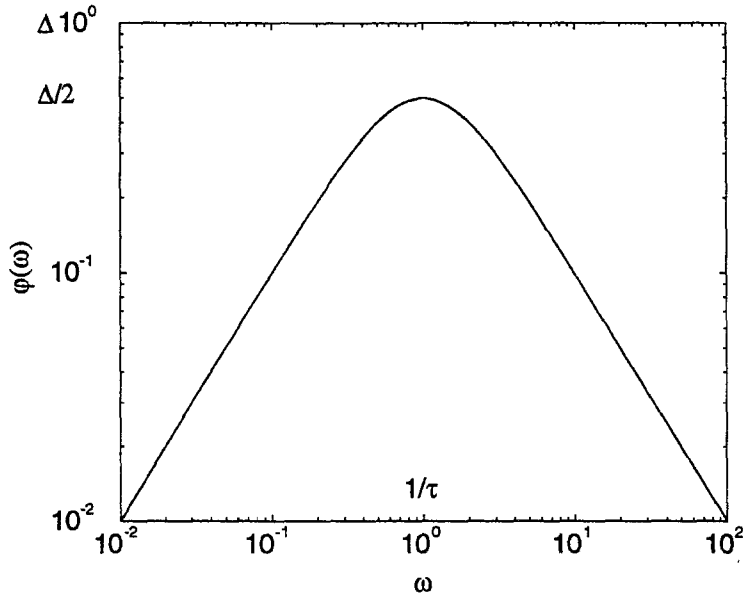

Figure 3-3: Plot of Eq. 3-7, with Δ and τ both set equal to 1.

One scenario (of many possibilities--the initial LIGO control system has not yet been designed) for the initial LIGO control system is for the coil driver to have an output resistance of 100 Ω, an output inductance of 4 H, and a γ of 0.05 N/Amp. In this case the maximum of the loss function would be $\varphi(4\text{ Hz}) = 5\times10^{-7}$ and $\varphi(100\text{ Hz}) = 4\times10^{-8}$. Eddy current damping in the output impedance of the coil driving electronics is more important than residual gas damping for the design of the initial interferometers, but the initial interferometer can be designed to reduce this damping to acceptable levels. As one looks forward to more advanced interferometers, eddy current damping becomes a more challenging problem. For this reason (and several others, e.g., see Chapter 6), advanced LIGO interferometers probably will not have magnets attached directly to the test masses.

In addition to eddy currents in the drive electronics, eddy currents will be induced in all conductors surrounding a moving magnet. To control this source of damping, the

amount of metal near the magnet must be minimized and the geometry of those conductors must be carefully considered.

Recoil Damping Losses

The pendulum suspensions of the LIGO interferometers are attached to soft, heavily damped seismic isolation stacks. When the pendulum swings, the seismic isolation stack recoils against it, and some energy is coupled into the stack itself. The magnitude of this damping mechanism can be estimated by considering the simplified model of two coupled harmonic oscillators shown in Figure 3-4. One oscillator (the pendulum) is idealized as having no loss if it were isolated. The other oscillator (the support structure) is lossy. By calculating the real part of the mechanical impedance between the pendulum and the (fixed) ground, an effective loss function of the pendulum due to the support structure can be estimated. In the approximation that the support structure is much more massive than the pendulum,

$$\varphi_p(\omega) \approx \frac{(m_p/m_s)\,\omega_p^2\omega_s^2\varphi_s(\omega)}{\left(\omega_s^2-\omega^2\right)^2+\omega_s^4\varphi_s^2(\omega)}. \qquad 3\text{-}8$$

The subscripts s and p represent the parameters of the support structure and pendulum, respectively. This equation is plotted in Figure 3-5 with parameters relevant to the 40-m interferometer (parameters are in Table 3-1). The loss function of the support structure, $\varphi_s(\omega)$, is taken to be a constant to emphasize the frequency dependence due to the resonant system and not the frequency dependence of the losses of the individual springs. In reality, $\varphi_s(\omega)$ is not necessarily a constant. At low frequencies ($\omega<<\omega_s$) the damping of the pendulum is proportional to the ratio of stiffnesses of the pendulum and the support

structure, whereas at higher frequencies, the pendulum and the support structure become decoupled.

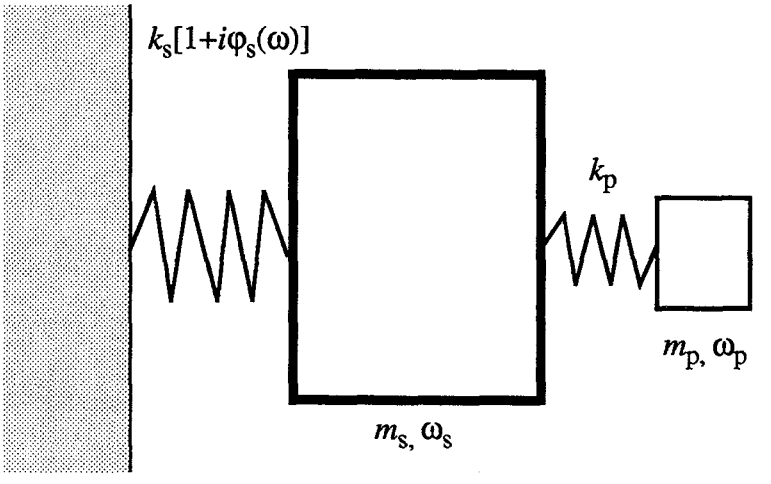

Figure 3-4: Schematic of damping of two coupled harmonic oscillators.

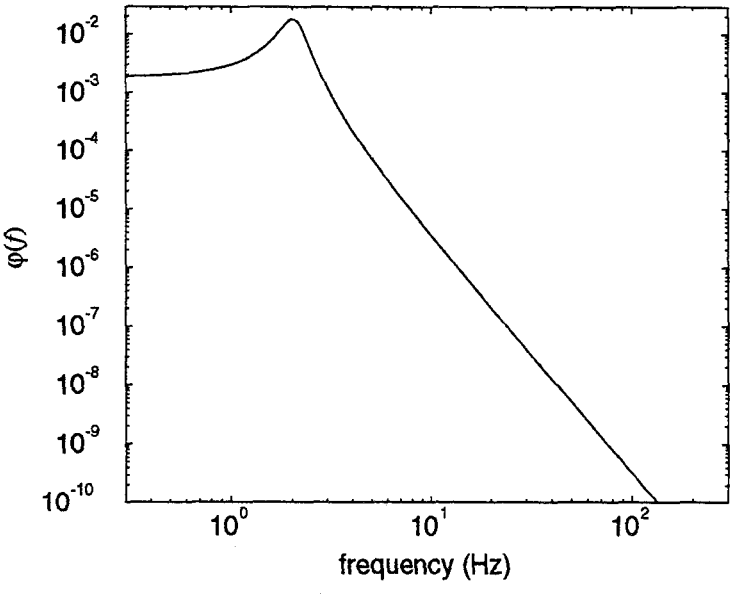

Figure 3-5: Approximate loss function of the pendulum due to the recoil of the support structure. The parameters used are from the 40-m interferometer (Table 3-1).

Table 3-1: Parameters of 40-m interferometer seismic isolation stack.

m_p	1.6 kg
m_s	200. kg
f_p	1. Hz
f_s	2. Hz
$\varphi_s(\omega)$	0.33

The model of Eq. 3-8 and Figure 3-5 is a simplification of the actual system. The seismic isolation stack of the 40-m interferometer has not one stage of isolation, but rather 3. In addition, each stage has six degrees of freedom and the coupling between the various degrees of freedom is not negligible. To get a more accurate estimate of the damping, a finite element analysis of all the modes and their couplings must be used. Nevertheless, Figure 3-5 does give a general idea of the shape of the loss function: there is an enhancement in the damping at frequencies near the resonant frequencies of the support structure, and at frequencies above those resonant frequencies, the loss function falls off as f^{-4}.

The design of the seismic isolation stack for the initial LIGO interferometer has not yet been finalized. However, its parameters will likely be similar to those of the 40-m interferometer. From Figure 3-5, it can be seen that contribution to the loss function from the recoil of the support will be much less than eddy current damping at frequencies near 100 Hz, and hence recoil losses will be negligible for the initial LIGO interferometer. However, advanced detectors will push the observing bandwidth down to 10 Hz, making recoil damping much more important. One obvious way to reduce recoil damping is to put a low loss stage between the heavily damped seismic isolation stack and the pendulum. An example of such a configuration would be a double pendulum [16].

Thermoelastic Damping

When a fiber is bent, one side is compressed and heats up, and the other side is stretched and cools off. The change in temperature across the wire results in an irreversible heat flow which is a source of dissipation. The form of the dissipation is given by Eq. 3-7 [17]. At frequencies corresponding to time scales which are large compared to the time scale for heat to flow across the wire, the wire is always near thermal equilibrium and there is very little heat flow, and at frequencies corresponding to time scales which are small compared to the heat flow time scale, the heat does not have time to flow. The characteristic frequency ($f_0 = 1/2\pi\tau$) for a round wire is

$$f_0 = 2.16\frac{\kappa}{Cd^2}, \qquad 3\text{-}9$$

and the strength is

$$\Delta = \frac{\mathcal{E}\alpha_{th}^2 T}{C}. \qquad 3\text{-}10$$

κ, C, d, $\mathcal{E}$, and α_{th} are the thermal conductivity, heat capacity per unit volume, diameter of the wire, Young's modulus, and coefficient of thermal expansion, respectively. When the wire is under tension as is the case for a pendulum wire, the strength is more complicated. The implications of thermoelastic damping for a pendulum suspension will be discussed in more detail in Section 4.7.

3.3 Other Sources of Loss

The four loss mechanisms described above are the only mechanisms known to be

within several orders of magnitude of limiting the interferometer performance. For each of those loss mechanisms, the parameters affecting them (i.e. the pressure of the vacuum system, the output resistance of the coil driver, etc.) can be adjusted using careful experimental design. When all of these loss mechanisms are sufficiently reduced, loss mechanisms which are not well understood become important. Since the nature of the loss mechanisms is uncertain, the frequency dependences of the loss mechanisms are also highly uncertain.

Many solid materials for which the specific nature of the loss mechanism is unknown exhibit losses which are independent of frequency [18]. A speculative explanation for the cause of this behavior is that the losses are not dominated by a single mechanism, but rather that there are a great many of microscopic structures, consisting of anything from grain boundaries to defects to impurities, each of which individually responds differently when the sample is put under stress. Typically, each loss point may be thought of as having a loss function something like Eq. 3-7 with its own strength and relaxation time. However, when summed over a large number of these points, the net loss function can appear to be nearly independent of frequency over large frequency bands.

It should be emphasized that this model for frequency independent losses is speculative. However, even though the source of the loss mechanism is unknown, this loss mechanism is very important for gravitational wave detectors. In both Chapter 4 and Chapter 6 frequency independent losses are found to be the dominant loss mechanism affecting both the pendulum and the internal vibrations of the test mass.

Chapter 4: Suspension Thermal Noise

The test masses of the LIGO detectors will be suspended as pendulums. The pendulum configuration is chosen because it is a very simple way to achieve a relatively low resonant frequency (below 1 Hz). At frequencies above the resonant frequency, the pendulum's mass behaves as if it were a free mass. A very low resonant frequency is desirable because the pendulum acts as a two-pole low pass filter for the purposes of attenuating external vibrations due to seismic or acoustic noise to the test mass. The low resonant frequency also puts the resonance well below the desired observing bandwidth of the detector, thus putting the measurement bandwidth far out on the high frequency tail of the pendulum thermal noise peak and reducing the thermal noise.

There is an additional thermal noise benefit to the choice of the pendulum as the means to hold the test mass. When a pendulum is excited, most of its energy is in either potential energy of the mass in the earth's gravitational field or kinetic energy. The earth's gravitational field is lossless, and if the pendulum is put in a modest vacuum, the viscous damping of the kinetic energy due to residual gas can be made very small. Therefore the pendulum's loss function, $\varphi(\omega)$, can be dominated by much smaller effects. If the recoil losses of the support structure and the losses due to clamping could be eliminated, then the loss function of the pendulum would be dominated by losses in the bending of the suspension fiber at the endpoint. In this case, the energy in the deformation of the wire can be computed simply from the radius of the wire, r, and the Young's modulus of the wire material, E. The ratio of elastic to gravitational energy, and hence the ratio of the loss function of the wire material to the pendulum loss function is given by:

$$\frac{E_{elast}}{E_{grav}} = \frac{\varphi_{pend}(\omega)}{\varphi_{material}(\omega)} = \beta \frac{1}{l} \sqrt{\frac{2EI}{\mathcal{T}}}. \qquad 4\text{-}1$$

Here, $\mathcal{T}$ and l are the tension and length of the wire. I is the area moment of inertia ($\pi r^4/4$ for a cylindrical wire). β is a factor of order unity that depends on the geometry of the pendulum, specifically on to what extent the mass rotates as the pendulum swings. This ratio can be of order 10^{-3} for pendulum configurations with metal wires. In addition, the loss function of many metals is of order 10^{-3}, so pendulum loss functions due to dissipation in the wire can be of order 10^{-6} for typical metal wire suspensions. The potential for very low dissipation along with a low resonant frequency make pendulum suspensions appear to be the most desirable way to support a test mass.

4.1 Early Measurements

One of the early attempts to make a very low dissipation pendulum achieved a Q of 8×10^6 [19]. This pendulum used an aluminum flexure hinge rather than a wire so that the area moment of inertia could be made very small in one direction. Based on this data, a pendulum Q of 10^7 was believed to be achievable for the initial interferometers. At that time a loss function of $\varphi(\omega) = Q^{-1}(\omega/\omega_p)$ was assumed, resulting in a target $\varphi(100\text{ Hz}) = 10^{-5}$ for a 1 Hz pendulum. This value gave a thermal noise level which, when compared to other projected noise sources and the signal strengths that were expected, was determined to be a reasonable and attainable goal for the initial interferometers. Reaching this level of dissipation was set as the goal for my work.

Kovalik and Saulson did an additional set of measurements to determine what a rea-

sonable choice of $\varphi(\omega)$ would be for several candidate suspension materials [9]. They made measurements on unstressed wires so that they could easily investigate a wide range of frequencies. They assumed that Eq. 4-1 could be used to predict the pendulum loss function from these measurements on unstressed wires. For their four candidate wire materials, tungsten, silicon, sapphire, and fused silica, $\varphi(\omega)$ was either limited by thermoelastic damping or independent of frequency.

The implications of $\varphi(\omega)$ being independent of frequency are very important. If $\varphi(\omega)$ were indeed independent of frequency, then to achieve a $\varphi(100\ \text{Hz})$ of 10^{-5} would require a pendulum Q of only 10^5, two orders of magnitude lower than was achieved at Glasgow and was believed to be needed for the initial interferometers. This difference illustrates the need to develop an experimental method of determining $\varphi(\omega)$ for a pendulum without relying solely on the pendulum Q.

4.2 Predicting Pendulum Losses from the Violin Modes

Not only does the pendulum mode have much lower losses than the material out of which it is constructed, but so do the violin modes of the suspension system. This is because most of the potential energy of the violin mode goes into work done against the tension of the wire due to the gravitational field and relatively little energy goes into deforming the wire. Saulson alluded to the idea that in a similar manner in which Eq. 4-1 is derived, one could predict the losses of the violin modes from the losses in the material [8]. Several years later it was shown that [20]:

$$\varphi_v(\omega) \approx \frac{2}{l}\sqrt{\frac{EI}{T}}\left[1+\frac{1}{2l}\sqrt{\frac{EI}{T}}(n\pi)^2\right]\varphi_{material}(\omega)\ . \qquad 4\text{-}2$$

Here n is the harmonic number of the mode. From Eq. 4-1 and Eq. 4-2, if the losses of the pendulum mode and the violin modes are limited by the material losses, then their loss functions are interrelated. Thus measurements of the loss function of the pendulum mode could be used to predict the loss function of the violin mode and visa versa.

This line of reasoning can be generalized to include not only dissipation internal to the wire material but in fact all loss mechanisms that occur near the endpoints of the wires. Recall that for the pendulum mode the wire is straight except for a bend very near the connection point, and explicitly calculating the loss from the energy stored in the bend results in Eq. 4-1. For the violin mode there is a similarly sharp bend near the end and also a more gentle bend along the length of the wire. The relative importance of the bend at the endpoints and the bend along the length of the wire can be seen by examining Eq. 4-2; the first term indicates the contribution to the loss function due to the bend at the ends, and the second term indicates the losses due to bending along the length of the wire. For parameters relevant to the LIGO detectors, the bending along the length of the wire is small compared to the bending at the ends of the wire for n less than about 10. Hence for the LIGO interferometers essentially all the material losses in the pendulum mode and the lower few violin modes occur near the endpoints of the wires.

To generalize the argument to include all loss mechanisms which occur near the endpoints of the wire, including losses due to imperfect clamping, consider the energy not in terms of the ratio of bending energy to total energy, but instead as a function of the bending angle of the wire near the end [11]. The calculation assumes that there is some loss of energy, $\Delta E(\theta)$, whenever the wire is bent through some angle θ. The strategy is to calculate the total energy of both the pendulum (E_p) and the violin (E_v) modes as functions of θ,

and to deduce the relationship between their respective damping by comparing $E_p(\theta)/\Delta E(\theta)$ with $E_v(\theta)/\Delta E(\theta)$.

The energy in the pendulum mode is, to second order in θ,

$$E_p \approx \frac{1}{2} mgl\theta^2 . \qquad 4\text{-}3$$

g is the acceleration due to gravity. If one assumes that the there is no energy stored in the bending along the length of the wire, then it follows that the shape of the violin modes is sinusoidal. In this case, the energy in the violin modes is, to second order in θ,

$$E_v \approx \frac{1}{4} Tl\theta^2 . \qquad 4\text{-}4$$

This equation is independent of which harmonic is chosen. The pendulum has the same amount of energy at a given angle regardless of the number of wires. Since each wire in a four wire system supports one quarter of the mass, the tension is $\frac{1}{4}mg$, and $E_v \approx (1/16)mgl\theta^2$. Therefore for a given angle θ the pendulum mode has eight times as much energy as a violin mode, $E_p = 8E_v$.

For a given mode, the loss of energy per cycle, ΔE, is given by

$$\Delta E(\theta) = 2\pi E(\theta)\varphi(\omega) . \qquad 4\text{-}5$$

If the losses for both the violin modes and the pendulum mode are primarily at the end-points, then for a given angle of bend the losses in the pendulum would equal the losses in the four wires,

$$\Delta E_p(\theta) = \Delta E_{v1}(\theta) + \Delta E_{v2}(\theta) + \Delta E_{v3}(\theta) + \Delta E_{v4}(\theta) . \qquad 4\text{-}6$$

Here the mass is assumed to be constrained by the four wires so that it does not rotate and the wire bends at both ends in the pendulum mode. The damping of the pendulum mode, in terms of the damping of the violin modes, is then

$$\varphi_p = \frac{1}{8}(\varphi_{v1} + \varphi_{v2} + \varphi_{v3} + \varphi_{v4}) . \qquad 4\text{-}7$$

A consequence of the losses being concentrated at the endpoints is that the losses are a function not only of frequency, but also of the length of the suspension wires. The explicit length dependence can be obtained by combining Eq. 4-4 and Eq. 4-5:

$$\varphi(\omega, l) = \frac{l_0}{l}\varphi(\omega, l_0) . \qquad 4\text{-}8$$

l_0 is an arbitrary reference length. This relationship is an important test of the hypothesis that the losses occur near the endpoints.

This model for the losses serves as the foundation to study the thermal noise of a pendulum suspension. It is consistent with the models of Gonzalez and Saulson [20] and of Logan et al. [21]. Eq. 4-7 is the primary means to estimate the frequency dependence of the loss function of the pendulum, and Eq. 4-8 is the test of the validity of Eq. 4-7. However, this model is not the complete picture of the damping of the pendulum mode, for it does not include forces which act directly on the mass such as residual gas damping or eddy current damping. These other loss mechanisms must be independently investigated.

4.3 Predicting the Suspension Thermal Noise

Once the loss function $\varphi(\omega)$ is found, it is straightforward to predict the thermal noise from the general thermal noise lineshape (Eq. 2-9). The thermal noise of a single violin

mode becomes:

$$S_x(f) = \frac{4k_BT}{\omega\mu_n}\frac{\omega_n^2\varphi_n(\omega)}{\left(\omega_n^2-\omega^2\right)^2+\omega_n^4\varphi_n^2(\omega)}. \qquad 4\text{-}9$$

ω_n is the angular frequency of the nth mode; μ_n is the corresponding reduced mass.

$$\mu_n = \frac{q}{2}m(\omega_n/\omega_p)^2 \qquad 4\text{-}10$$

for the violin modes ($n > 0$) [8][22]; ω_p is the pendulum angular resonant frequency and q is the number of wires. μ_n accounts for the small recoil of the test mass against the thermal motion of the wire. To determine how many violin modes are required to determine the low frequency thermal noise accurately, rewrite Eq. 4-9 in the low frequency limit with $\omega_n = n\omega_1$ (ω_1 is the resonant frequency of the first harmonic) and sum over all the violin modes:

$$S_x(f) \propto \sum_n \frac{\omega_p^2\varphi_n(\omega)}{\omega(n\omega_1)^4}. \qquad 4\text{-}11$$

The $1/n^4$ term makes the first few violin modes dominate S_x for any reasonable choice of $\varphi(\omega)$; therefore in order to predict the thermal noise due to the suspension system only the violin modes with resonant frequencies at or below the frequency of interest need to be included.

4.4 Measurements on the 40-m Interferometer

The Q's of individual suspension wires in the 40-m interferometer were determined

by exciting the violin resonance, turning off the excitation, and measuring the decay times ($Q_n = \frac{1}{2}\omega_n \tau_A$; τ_A is the measured amplitude decay time). A schematic view of the 40-m interferometer suspension is shown in Figure 4-1. The details of the interferometer length measurement process are described in the Appendix. Magnets and actuator coils were attached to the end test masses for applying calibration signals and to maintain resonance between the optical cavities and the light. On all masses a piezoelectric transducer was normally used to damp the residual pendulum motion. For measurements on the end masses the driving signal was applied to the actuator coils. The vertex masses were driven using the damping transducers on their control blocks. To verify that neither the pendulum damping transducers nor the interferometer control servomechanisms damped the wire resonances, the decay times were measured both with and without servomechanism signals applied to the actuators.

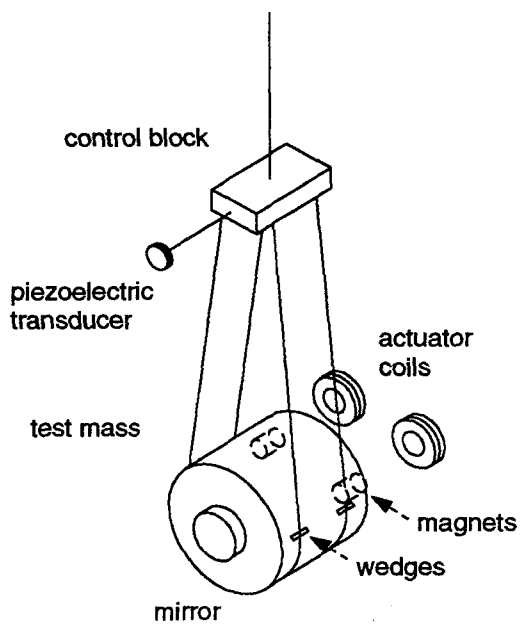

Figure 4-1: Schematic view of the 40-m interferometer suspension.

The frequencies and the Q's of the fundamental violin modes of the suspension system of the 40-m interferometer measured in the Spring of 1992 are given in Table 4-1. The uncertainties in the Q measurements are approximately 10%. The differences in frequencies among the different masses are explained by differences in the suspensions (which arose through various modifications to the interferometer). The east end mass used 150 μm diameter wire. The south end mass used 100 μm diameter wire. Both vertex masses used 75 μm diameter wire. Steel music wire was used for all masses; the length of all the suspension wires was 25 cm, and each test mass was 1.5 kg.

The differences in the Q's of the violin modes for the different masses might partly depend on the wire diameters, which affect both the tensile strength and the stiffness of the wires. However there were also significant variations in the connections of the wires to the masses. In all cases a small glass wedge was used to define the point where the wire connects to the test mass, and a similar aluminum wedge was used where the wire meets the control block. Unfortunately fine details (such as how the wedges and the wires were bonded to the mass) varied among the wires. Furthermore, the end masses and the vertex masses had different types of control blocks, for historical reasons.

The Q's of the second harmonics of five wires, two wires from the east end mass, two wires from south end mass including both a high and a low Q wire, and a high Q wire from the east vertex mass, were measured and are also shown in Table 4-1. Two additional higher order harmonics of the east vertex wire were measured. The Q of all harmonics were the same as the corresponding fundamental resonance Q within 25%. This result is consistent with a loss function which is independent of frequency in the range of several hundred hertz. This damping model was adopted to describe the thermal noise in the 40-m

interferometer.

Table 4-1: The measured Q's of the violin resonances.

Test Mass	Frequency (Hz)	Q (10^3)	Q (10^3) of *n*th harmonic; *n*
East End	319.65	13	
	324.90	16	19; 2
	326.08	19	
	328.45	15	16; 2
East Vertex	594.35	240	
	596.68	280	260; 2
			220; 3
			220; 4
	598.15	43	
	605.02	110	
South End	505.85	66	
	506.88	120	110; 2
	512.85	23	23; 2
	514.90	16	
South Vertex	592.70	295	
	592.80	295	
	596.42	356	
	600.22	163	

4.5 Comparison of Estimated Suspension Thermal Noise with Experiment

The lineshapes of the violin resonances were compared with the thermal noise prediction using Eq. 4-9. This comparison is shown in Figure 4-2 for frequencies near the first harmonics of the east end mass. The solid line is the measured displacement noise spectrum of the 40-m interferometer. The dashed line is formed simply by plugging the values in Table 4-1 into Eq. 4-9, and by adding the noise of the four wires together. It is not a fit

to the data! Next to two of the resonances are smaller peaks which are believed to be modes of the wires with polarizations which are nearly but not completely orthogonal to the optical axis. A finer comparison between the theory and experiment is given in Figure 4-3 where the data points (for the middle resonances in Figure 4-2) are plotted with their associated error bars. The uncertainties are statistical errors due to the limited number of averages in the power spectrum. The systematic uncertainty of the calibration is 10%. The uncertainty in the theoretical curve due to the uncertainties in the Q's is smaller. The agreement between the measured noise and the predicted noise is quite good. The corresponding rms fluctuations in the test mass due to a single violin resonance is about 0.06 fm.

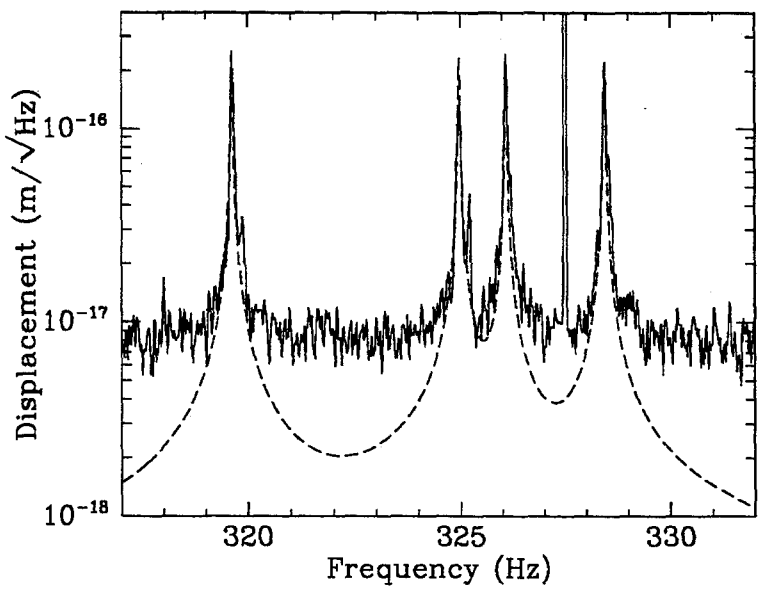

Figure 4-2: Comparison of predicted thermal noise (dashed line) to the measured noise (solid line) for the violin resonances of the east end test mass. The large peak at 327.5 Hz is a calibration signal.

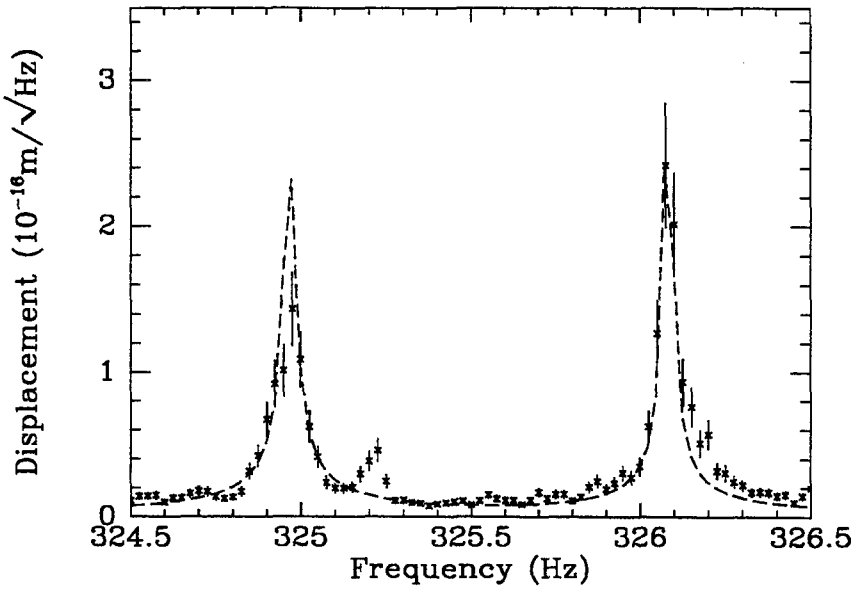

Figure 4-3: Fine comparison of thermal noise prediction (dashed line) and two of the violin resonances of the east end test mass (×'s).

Using the measured Q and assuming that the pendulum mode is damped as derived above, the total thermal noise of the wire suspensions can be predicted by summing over all modes:

$$S_x(f) = \sum_{4\ \text{masses}} \left(S_{\text{pendulum}}(f) + \sum_{4\ \text{wires}} S_{\text{violin}}(f) \right). \qquad 4\text{-}12$$

Figure 4-4 shows this thermal noise prediction and compares it to the measured interferometer noise of June 1992. As can be seen from the figure, thermal noise contributed to the 40-m interferometer by the suspension wires was much smaller than other noise sources except in narrow frequency bands centered on the resonant frequencies of the violin modes. The increase in the predicted suspension noise below 200 Hz is due to the pendulum mode. The resonances near 320 Hz are excited above thermal noise by a factor of two by excess noise in the pendulum damping control electronics used at that time; when

that control system was turned off (which degraded the overall performance of the interferometer) those peaks were at thermal noise (see Figure 4-2). The electronics have been fixed since then. The peaks in the noise between 1800 and 2000 Hz which appear to coincide with the violin resonances are due to residual frequency noise from the laser, not due to the violin resonances.

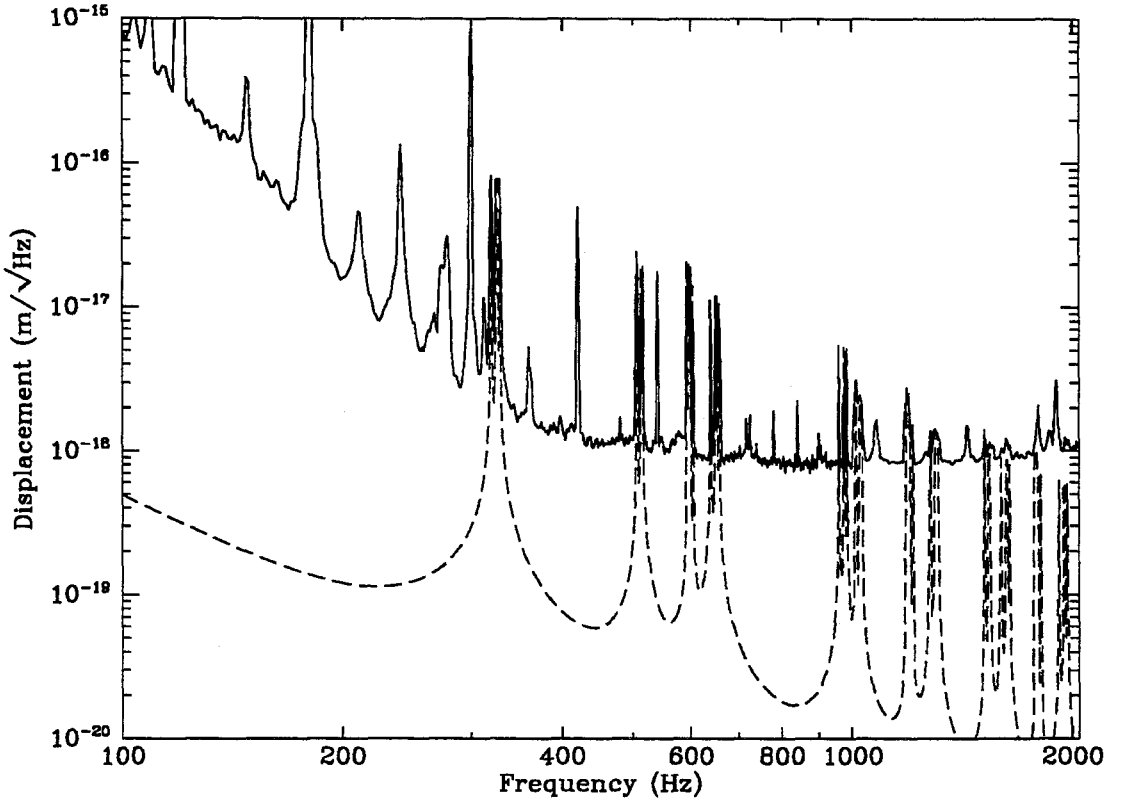

Figure 4-4: The solid line is the interferometer noise; the dashed line is the thermal noise prediction. The data were taken at two different bandwidths: 1.25 Hz from 100 Hz to 1000 Hz and 6.25 Hz from 1000 Hz to 2000 Hz. The thermal noise prediction was generated by averaging the lineshape over the appropriate bandwidth. Due to this averaging the violin resonances appear broadened.

The remarkable agreement between the theory and the noise of the 40-m interferometer shows that the fluctuation dissipation theorem can be used to predict the thermal noise.

However, those data are not a good test of the model used to predict $\varphi(\omega)$. The 40-m interferometer was limited by thermal noise only in very narrow bandwidths around the resonant frequencies of the violin modes. In the limit that the Q of a resonance becomes very large, all thermal noise lineshapes (in which the loss function varies smoothly over the bandwidth ($\Delta f = Q^{-1} f_0$) of the resonance) converge to the same shape near the resonant frequency. The significance of those data is that they give confidence in the ability of the fluctuation dissipation theorem to predict the thermal noise in a mechanical system, and, since the fluctuation dissipation theorem has a quantum mechanical derivation, these data can be interpreted as yet another of many examples showing the remarkable power of quantum mechanics to predict the behavior of physical systems that are macroscopic in size.

4.6 Independent Tests of the Damping Model

In order to test the damping model as expressed in Eq. 4-7 and Eq. 4-8, an independent experimental apparatus was developed which enabled more rapid tests than would be possible with the 40-m interferometer. A schematic of the apparatus is shown in Figure 4-5. The suspension system consisted of a 1.6 kg cylinder of fused silica (with a diameter of 10 cm and a length of 8.8 cm) suspended by two loops of 75 μm diameter steel music wire. The violin resonances were excited by a piezoelectric transducer glued to the support bar near the clamping point at the top of the suspension. The motion was measured by focusing a HeNe laser onto the wire and monitoring the diffraction pattern on a split photodiode. To damp the seismic motion of the mass while the violin mode Q's were being measured, magnets were attached directly to the mass and current was fed back from an

active damping circuit to nearby coils. The pendulum mode was excited by physically pushing the mass just prior to pumping the vacuum system. The test mass motion was monitored by an edge sensor (an LED which cast a shadow of the corner of the test mass on a photodiode).

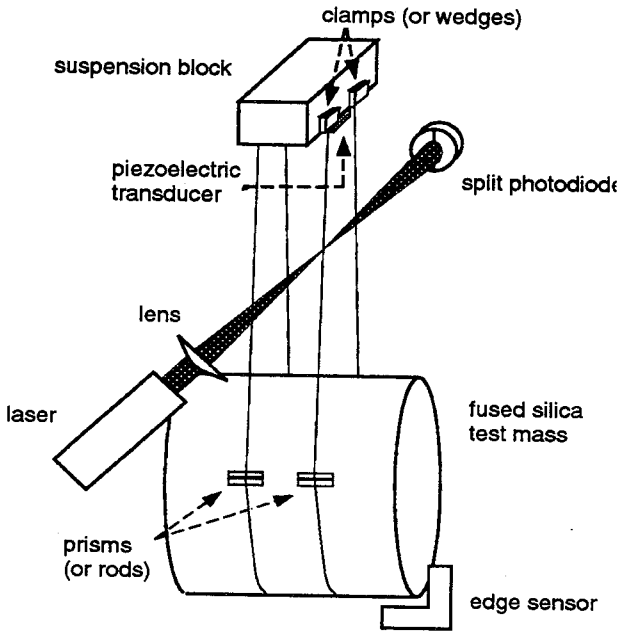

Figure 4-5: Schematic view of the experimental apparatus.

For each measurement, new wire was taken directly from the spool and wiped with acetone to remove the residual oils; no annealing or preloading treatments were used. The wire tension was approximately one half of its breaking strength. Two different types of clamps were used at the top. In one type the wire was simply clamped between two aluminum bars; the corners of the bars had approximate radii of curvature of 100 μm and the faces of the bars where the wires left were flush to approximately that level. For the other clamping arrangement, the wire left a 45° wedge with a 25 μm radius edge. No difference in the Q's was found between the two clamping methods. At the test mass the wires were

looped over a standoff to maintain a significant pressure on the wire and thus avoid rubbing at the point of contact. Either fused silica prisms (13 mm long by 2 mm equilateral triangles) or rods (13 mm long by 1 mm in diameter) were used as standoffs. The standoffs were held in place by one of three methods: held simply by the pressure of the wire with no glue, glued on the cylinder with a cyanoacrylate based glue, or attached to the cylinder with a vacuum sealant epoxy. There was no observable difference between the Q's of either type of standoff even though the sharp edge of the prism presumably exerted more contact pressure on the wire than the gentler curve of the rod, and the means of holding the standoffs in place had no effect. However, when no standoffs were used, the Q's of the violin mode resonances were degraded by a factor of approximately 30.

The experiment was designed so that losses that directly affected the pendulum were minimized (as discussed in Section 3.2). The apparatus was placed in a vacuum system with an operating pressure of 10^{-4} Torr, limiting the pendulum losses to a minimum of $\varphi \approx 10^{-7}$. When the pendulum losses were measured, the magnets were removed from the mass, making eddy current damping negligible. The support structure was designed such that the recoil losses would limit the pendulum losses at a level of $\varphi \approx 10^{-7}$.

4.7 Suspension System Losses

The measured losses of the violin modes are plotted in Figure 4-6. For these measurements a suspension wire length of 35 cm was used. The solid squares represent the average of measurements made on several repeated suspensions of the test mass. The error bars indicate the statistical variation of the losses from wire to wire. The value at 0.8 Hz was derived from measurements of the pendulum mode, but for comparison the value was

expressed as an effective violin mode loss at that frequency using Eq. 4-7. The open circles represent the measured losses of different harmonics of a single wire; the reproducibility of the measurements on that single wire is comparable to the size of the circles. These data show a variation in the measured losses of the violin modes of only a factor of four over four decades of frequency.

The rise in $\varphi(\omega)$ at high frequencies can be explained by thermoelastic damping (see Section 3.2). The thermodynamic properties of steel can vary substantially depending on both the specific alloy and the treatment; extreme values give a possible range of characteristic frequencies of 4 to 11 kHz, and a possible range of Δ of (1-3) $\times$ 10^{-3} [23]. The exact composition of the wire used was unknown; for illustrative purposes, f_0 was taken to be 7 kHz and Δ was taken to be 2.5 $\times$ 10^{-3}. When the wire is under tension, the thermoelastic losses (Eq. 3-7) must be modified according to Eq. 4-2. The factor due to the tension of the wire, $\frac{l}{2}\sqrt{\frac{\mathcal{T}}{\mathcal{E}I}}$, is 620. The second term of Eq. 4-2 adds some frequency dependence to the thermoelastic loss curve due to the curvature of the wires becoming increasingly important at higher order harmonics. For the parameters of this system, the second term of Eq. 4-2 becomes equal to unity at $n \approx 16$ (frequency of 6.8 kHz), almost cancelling the peaked shape of Figure 3-3. Calculated thermoelastic losses are shown as the dashed line in Figure 4-6.

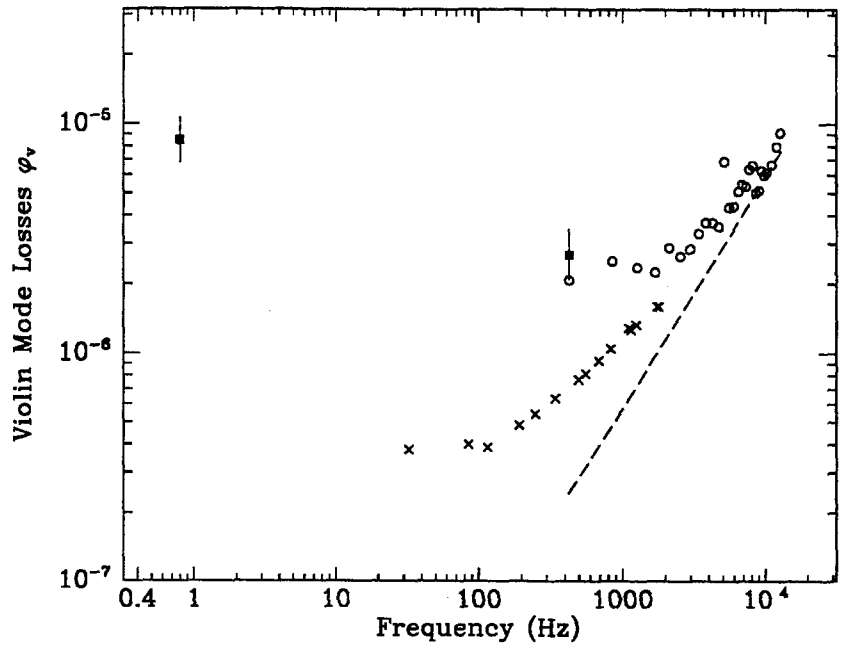

Figure 4-6: Frequency dependence of the violin mode losses. The solid squares represent the average losses of several wires at both the fundamental violin mode resonant frequency and at the pendulum resonant frequency; the open circles represent measurements on different harmonics of a single wire. The ×'s are the predicted losses from the wire material loss measurements discussed in Figure 4-8. The dashed line is the predicted loss due to thermoelastic damping.

The data of the violin mode losses in Figure 4-6 can be divided into two parts, a slowly varying section between 1 Hz and 2 kHz, and a section above 2 kHz which is consistent with thermoelastic damping. Based on this data, a model in which the losses are independent of frequency is adopted as a simple approximation to the slowly varying lower frequency data. The uncertainty in the measurements from wire to wire and the lack of data points between 1 Hz and 400 Hz make a more detailed model unwarranted. Losses which are independent of frequency are a common characteristic of many solid materials [18], and have been observed in steel at lower frequencies [15].

Frequency independent losses that occur near the endpoints of the wire result from

Eq. 4-8 in a simple relation between the Q of the violin resonances and the wire length,

$$Q \propto l. \qquad 4\text{-}13$$

By measuring the Q of the violin resonances as a function of length, the fundamental assumption of this model, that the losses occur near the endpoints of the wire, can be verified.

The results of such measurements are given in Figure 4-7. Each × represents a measurement of the Q of a different wire. As in Figure 4-6 the reproducibility of a measurement made on a single wire was approximately 5%; the spread in the points at a given length indicates the variation in the Q's from wire to wire. The line in the figure represents the model of Eq. 4-13. The data support the assumption that the losses occur near the endpoints. Note that due to different wire lengths, the data were taken at different frequencies. The observed weak frequency dependence of the losses can give systematic departures from the model of Eq. 4-13 of the order of 10% over this frequency band (180-800 Hz).

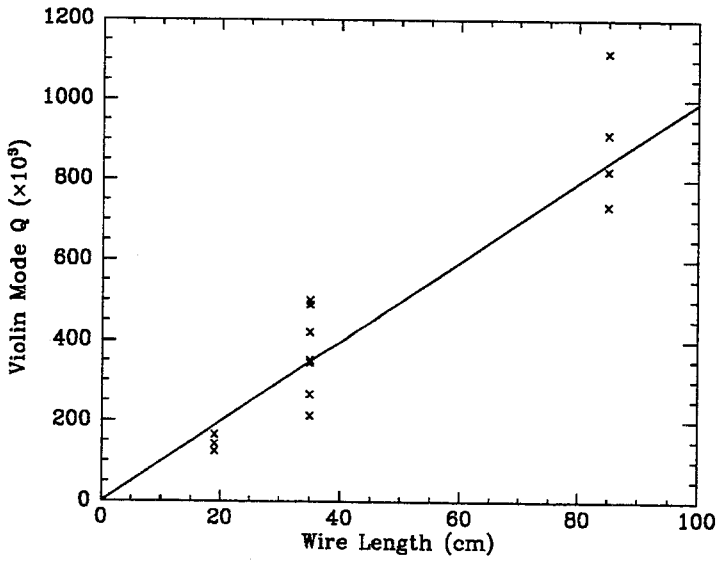

Figure 4-7: Measured Q's as a function of wire length.

Two possible sources of losses at the endpoints have been identified. One is damping in the clamps connecting the wire to either the test mass or the support structure. Since modelling the losses due to imperfect clamps is difficult, two different types of clamps were used at both the top and the bottom of the wire (described in Section 4.6). No difference in the measured Q's was found for any clamp combination. This result is interpreted as indicating that the losses were probably not due to the clamps. The remaining possibility is that the losses were in the wire material itself.

4.8 Wire Material Losses

To compare the losses of the violin modes to the intrinsic losses of the wire material, the wire material Q's were measured using a method analogous to that of Kovalik and Saulson [9]. Pieces of wire were clamped at one end using the same aluminum bar which was used for the suspension system, and the other end was left free. The modes were excited using the piezoelectric transducer glued to that clamp, and the motion was measured in the same manner as for the violin modes.

The frequency dependence of the losses of the wire material is shown in Figure 4-8. Three different wire lengths, 4, 5, and 6 cm, were used to arrive at the density of points. In each case, new wire was taken directly from the spool and wiped with acetone. No attempt was made to straighten the short lengths of wire used, resulting in a gentle bend over the length of the wire (about a 1 cm deflection over a 6 cm length). The error bars indicates the statistical reproducibility of the measurements on a single wire; at lower frequencies the measurement was limited by seismic and acoustic vibrations. The smoothness of the data gives an indication of the reproducibility from wire to wire. The dashed line on the

figure is the loss due to thermoelastic damping.

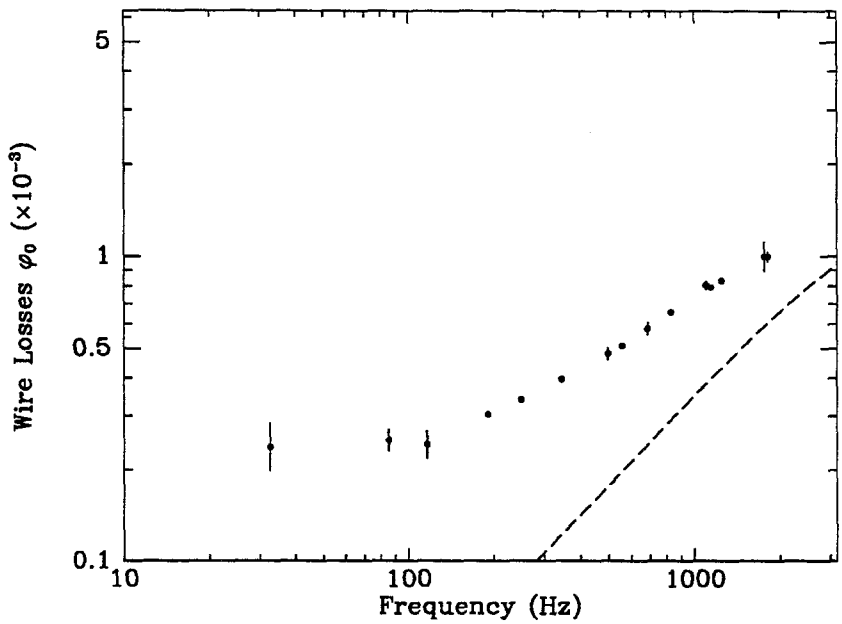

Figure 4-8: Wire material losses. The dashed line is the predicted loss due to thermoelastic damping. The error bars for the points in the mid frequency range are comparable to the size of the points.

A prediction of the losses in the violin modes of the test mass suspension due to the wire losses, using the model of Eq. 4-2, is represented by the ×'s in Figure 4-6. Note that although the results of Section 4.7 suggest that the losses may be due to the wire material itself, the measured test mass suspension losses exceed the losses predicted from the measurements on unloaded wires. One possible explanation for this discrepancy is that the losses in the wire material are a function of the stress in the material. Careful studies of the stress dependence of the losses in beryllium/copper [24] and in stainless steel and tungsten [25] have not shown stress dependent dissipation in these materials.

4.9 Vertical Resonance

In addition to the pendulum and violin mode resonances, the vertical spring mode of the suspension system must also be included in a complete thermal noise estimate. Since the full scale LIGO interferometers will be built on the curved earth, the local gravitational forces on test masses at opposite ends of the 4 km long arms will not be parallel to each other. Therefore some component of the local vertical frames will lie along the direction of the laser beam and couple into the gravitational-wave signal.

To judge the importance of this coupling, the thermal noise due to the vertical spring mode above its resonant frequency is compared to the thermal noise of the pendulum mode above its resonant frequency. The ratio of the two is:

$$\frac{S_{vert}(f)}{S_p(f)} \approx \Theta^2 \frac{\omega_{vert}^2 \varphi_{vert}(\omega)}{\omega_p^2 \varphi_p(\omega)} \approx 0.02 . \qquad 4\text{-}14$$

Θ is the angle subtended on the earth by the two test masses and is about 0.6 mrad for a 4 km arm length. For steel wires, ω_{vert}/ω_p is about 15. The losses of the vertical spring mode are dominated by the material losses since all of the potential energy is stored in the wire; φ_{vert}(15 Hz) was measured to be 3×10^{-4}, consistent with the loss measurements made on the unstressed wire (see Figure 4-8). From Figure 4-8 it is reasonable to assume that, at frequencies of interest for the suspension thermal noise, (below 200 Hz) that φ_{vert} is a constant independent of frequency, like φ_p. With all of these factors considered, the ratio of the vertical to pendulum thermal noise power is about 0.02.

This factor of 50 indicates that vertical thermal noise will not be an important factor

in the design of the initial LIGO interferometer. However, this factor is not overwhelming, and while looking forward to advanced interferometer, one must verify that the vertical thermal noise is still at a low enough level. In particular, there are several instances where vertical thermal noise may become significant. One way to dramatically decrease the pendulum losses is to use alternative geometries to the cylindrical wire suspension that reduce the area moment of inertia of the suspension (I in Eq. 4-1) [19][26]. This technique does not improve the losses in the vertical mode and hence by itself cannot produce a large improvement in the thermal noise. Another problem as one looks toward advanced interferometers is that the target frequency band for observation extends down to 10 Hz. This bandwidth includes the resonant frequency of the vertical suspension mode near 15 Hz; that mode would effectively foil any observation at nearby frequencies. The VIRGO project is investigating techniques to lower the vertical resonant frequency by attaching the suspension wire to a soft spring [27]. This method has the disadvantage that it creates another weakly coupled higher frequency resonance in the observing bandwidth.

Chapter 5: Test Mass Vibrational Modes

The vibrational modes of the test mass are more complicated to deal with than the suspension modes. For the suspension modes, the wire vibrates and the test mass can be treated as an ideal point mass without any loss of information, whereas for the test mass modes, the test mass vibrates in a complex three dimensional shape with different parts of the mirror moving different amounts. Therefore for the test mass modes, the motion of all parts of the mirror sampled by the laser beam must be considered. In addition, the varying light intensity across the optical spot must also be considered. However, the only information about the mirror motion that the interferometer extracts when the light is recombined at the beamsplitter is the overall phase shift that the light picks up when it reflects from the mirrors. Therefore, all of the information regarding the geometry of the laser beam and the mirror that is relevant to the interferometer can be condensed into a single parameter. My choice of parametrization is in terms of an effective mass coefficient. This chapter describes the coupling between the interferometer signal and the test mass vibrational modes and the consequences of test mass vibrational thermal noise.

5.1 Effects of Mirror Vibrations on Optical Modes

To find the displacements detected by the optical system when the mirror surfaces are vibrating, the interaction between two different types of modes must be considered: the mechanical modes of the vibrating test masses, and the optical modes of the electromagnetic field resonating in the Fabry-Perot cavities.

The vibrational eigenmodes and eigenfrequencies of a free right solid cylinder can be

found by solving the equations of elasticity using an analytic series solution [12]. From the results of that solution, the amplitude of the displacement of the mirror surface at each point, $\vec{u}_n(\rho, \theta)$, can be calculated for the nth mode normalized to a fixed energy, E. Implicit in the mode shape, $\vec{u}_n$, is a time dependence $e^{i\omega_n t}$.

The optical modes of the interferometer can be described by Hermite-Gaussian functions, ψ_{lm} [28]. The interferometer typically operates with the TEM_{00} mode, ψ_{00}, on resonance. To avoid interference between modes, the length of the interferometer and the curvature of the mirrors are chosen such that the TEM_{00} mode and other higher order modes cannot resonate simultaneously [29].

The optical mode experiences a phase shift upon reflection from a mirror excited in a particular vibrational mode:

$$\psi_{00}(\rho, \theta, z) \rightarrow \psi_{00}(\rho, \theta, z)\, e^{i2\vec{k}\cdot\vec{u}_n(\rho,\theta,z)} \approx \psi_{00}\left[1 + i2\vec{k}\cdot\vec{u}_n - 2\left(\vec{k}\cdot\vec{u}_n\right)^2\right]. \qquad 5\text{-}1$$

$\vec{k}$ is the wave vector, and $\left|\vec{k}\cdot\vec{u}_n\right|$ is taken to be much less than unity for all points on the mirror surface. This new perturbed mode in the interferometer can be described in terms of the unperturbed modes:

$$\psi = \sum_i \sum_j c_{ij}\psi_{ij} \qquad 5\text{-}2$$

$$c_{ij} = \int_S \psi_{ij}^{*}\psi_{00} e^{i2\vec{k}\cdot\vec{u}_n(\rho,\theta,z)}\, d\sigma. \qquad 5\text{-}3$$

The integral is an area integral over the mirror surface, S. Since only the TEM_{00} component of the perturbed light resonates, the new resonating mode can be written as

$$\psi \approx \psi_{00}\left[1 + i2\int_S \psi_{00}^{*}\psi_{00}\vec{k}\cdot\vec{u}_n d\sigma - 2\int_S \psi_{00}^{*}\psi_{00}\left(\vec{k}\cdot\vec{u}_n\right)^2 d\sigma\right]. \qquad 5\text{-}4$$

Each term of this expression can be easily interpreted. The imaginary term contains the phase shift from which the apparent length change Δl_n can be determined,

$$\Delta l_n = \left|\vec{k}\right|^{-1}\int_S \psi_{00}^{*}\psi_{00}\vec{k}\cdot\vec{u}_n d\sigma. \qquad 5\text{-}5$$

The second integral term describes the light which is scattered out of the TEM_{00} mode. The apparent motion, Δl_n, is generally different for different vibrational modes of the mirror with the same energy.

To remove the dependence of the amplitude of the displacements, $\vec{u}_n$, on the energy normalization, it is convenient to parametrize the coupling of each mode in terms of an effective mass coefficient, α_n, defined as

$$\alpha_n \equiv \frac{E}{\frac{1}{2}m\omega_n^2\Delta l_n^2}. \qquad 5\text{-}6$$

m and ω_n are the mass of the mirror and the angular resonant frequency of the vibrational mode, respectively. With this parametrization, the apparent motion of the illuminated mirror surface oscillating in a particular vibrational mode can be modeled as if it were a point mass of magnitude $\alpha_n m$ vibrating with a resonant frequency ω_n measured by an ideal one-dimensional laser beam. The result of this mathematical procedure is shown schematically in Figure 5-1.

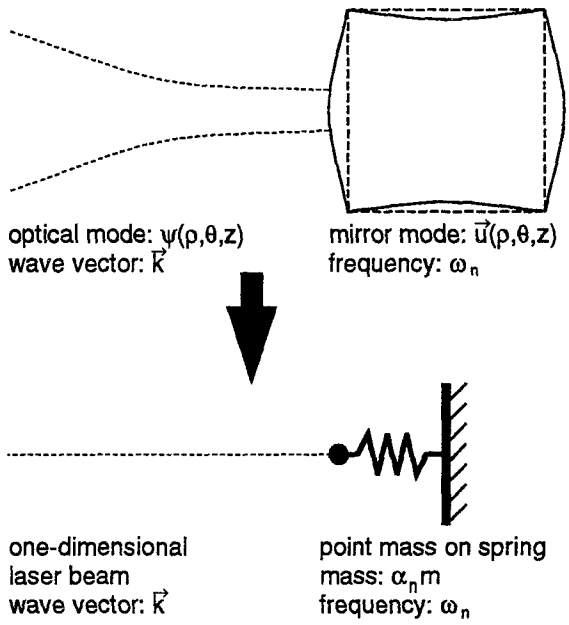

Figure 5-1: Schematic representation of the effective mass coefficient. The effective mass coefficient, α, is the parameter used to model the interactions between the optical mode of the interferometer and a vibrational mode of the mirror as a one dimensional mass-spring system.

As an example, the mode shapes, resonant frequencies, and effective mass coefficients of the first six axisymmetric modes of a 10-cm mirror are shown in Figure 5-2. A 10-cm mirror (used in the 40-m interferometer) is a fused silica cylinder with a diameter of 10 cm, a length of 8.8 cm, and a mass of 1.6 kg. The effective mass coefficients depend not only on the parameters of the mirror, but also on the geometry of the optical mode on the mirror surface. For all the calculations of this chapter, except in Section 5.4 where it is explicitly stated otherwise, the laser parameters used with the 10-cm mirror are those relevant to the 40-m interferometer, which has a spot size (the radius at which the intensity is $1/e^2$ of its maximum) of 0.22 cm. The beam is assumed to be centered on the mirror and in the TEM_{00} optical mode, so that non-axisymmetric vibrational modes do not contribute

apparent motion to the mirror surface. The benefit of this choice of beam location is that the center of the mirror is a node for all non-axisymmetric modes, making calculations much simpler by decreasing the number of modes involved. The implications of moving the beam spot off center are discussed in Section 5.5.

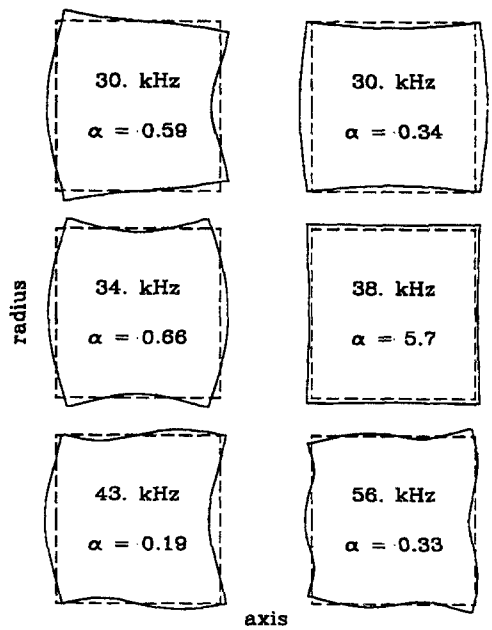

Figure 5-2: Mode shapes, resonant frequencies, and effective mass coefficients of a 10-cm mirror.

The effective mass coefficients of the first 100 axisymmetric modes of a 10-cm mirror are plotted against their respective resonant frequencies in Figure 5-3. The effective mass coefficients vary by several orders of magnitude, reflecting the wide variety of mode shapes. There is, however, a general trend toward lower effective masses at larger resonant frequencies. To draw attention to this trend, a dashed line representing $\alpha_n \propto f_n^{-1}$ is drawn. This line is approximately the median effective mass coefficient in a given bandwidth. The significance of this line will become apparent in Section 5.4. This general trend arises

from the cylindrical symmetry of the axisymmetric modes, which dictates that the largest antinode of motion in the axial direction is at the center of the mirror, the position being sampled by the laser. When the acoustic wavelength is shorter, this antinode will have a narrower width and thus a larger amplitude for a given energy. Hence the apparent motion of the mirror surface sampled by the laser light can be relatively larger for the higher frequency modes, and the effective mass can be correspondingly smaller. To illustrate this point, the shapes of the mirror surface for four modes with small effective masses are shown in Figure 5-4 (these modes are shown as solid circles in Figure 5-3); notice that a relatively small portion of the mirror around the center of the mirror surface moves much more than any other spot.

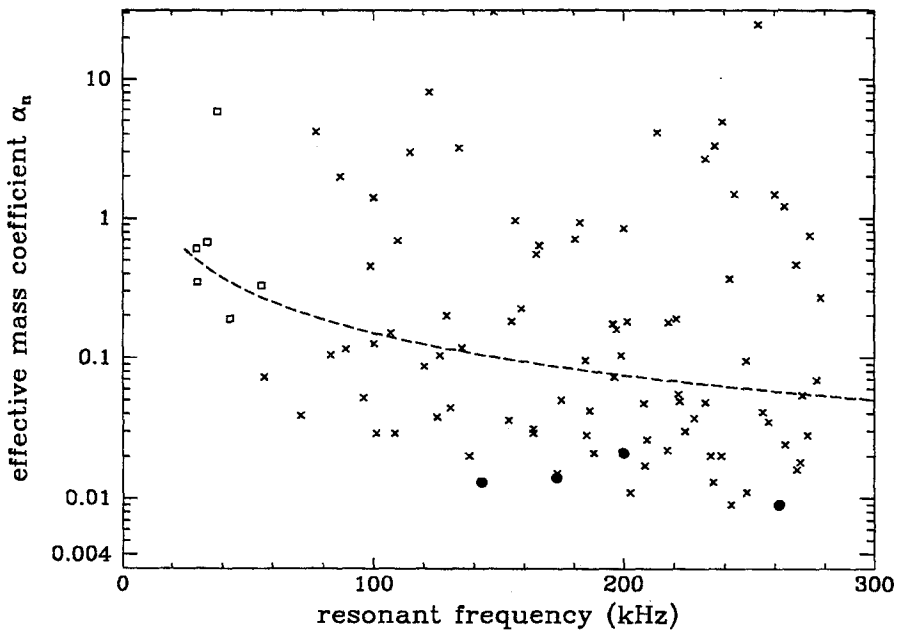

Figure 5-3: The effective mass coefficients as a function of mode resonant frequency for a 10-cm mirror. The modes shown in Figure 5-2 are plotted as open squares; the modes shown in Section Figure 5-4 are plotted as solid circles. The dashed line corresponds to $1/f$.

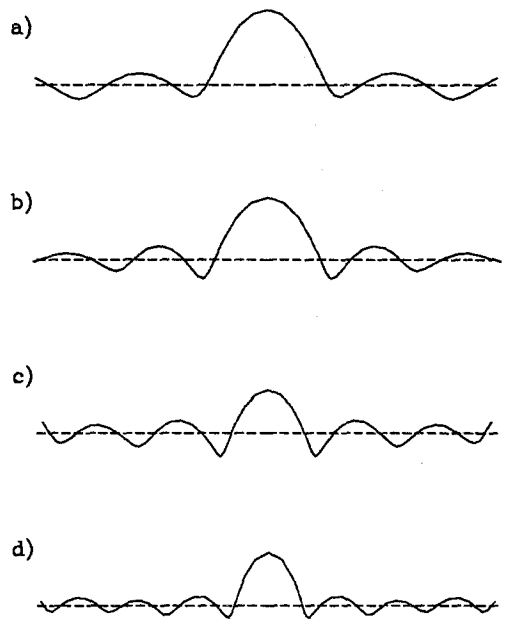

Figure 5-4: The shape of the mirror surface for four higher frequency modes with small effective masses. a) $f = 143$ kHz, $\alpha = 0.013$; b) $f = 173$ kHz, $\alpha = 0.014$; c) $f = 200$ kHz, $\alpha = 0.021$; d) $f = 262$ Hz, $\alpha = 0.009$.

Previous estimates (see, for example, [8]) of the coupling between the vibrational modes of the mirror and the optical path length of the interferometer considered only purely longitudinal motion. In this case the mode shapes can be described simply by

$$u_n = u_0 \sin\left(\frac{n\pi z}{2h}\right), \qquad 5\text{-}7$$

where h is the thickness of the mirror. Such a model predicts an effective mass coefficient of 0.5 for all modes. This approximation gives a reasonable estimate of the effective mass coefficients of some low frequency modes (e.g. the first few modes of Figure 5-2). However, such a model fails to predict the much lower effective masses of some of the higher frequency modes that are shown in Figure 5-3, and therefore this model underestimates

the total noise contributed by mirror vibrations.

To check the self-consistency of this three dimensional model, the approximations that went into its formulation must be examined. The first assumption is that $\left|\vec{k} \cdot \vec{u}_n\right| \ll 1$. This is approximately the same as requiring that $\left|\vec{k}\right| \cdot \Delta l_n \ll 1$. If the vibrations are thermally excited, it follows from the equipartition theorem that

$$\Delta l_n \approx \sqrt{\frac{k_B T}{\alpha_n m \omega_n^2}}. \qquad 5\text{-}8$$

Hence at a temperature of 300 K, $\left|\vec{k}\right| \cdot \Delta l_n$ is of order 10^{-10} for the lower frequency modes shown in Figure 5-2, clearly satisfying the requirement $\left|\vec{k} \cdot \vec{u}_n\right| \ll 1$.

The second integral term of Eq. 5-4, which describes the scattered light, is of order 10^{-20} of the light incident on the mirror for the modes of Figure 5-2. This is small compared to the constant scattering due to microroughness and figure errors in the mirror, which will be of order 10^{-4}-10^{-6}. However, the amplitude of the light scattered by thermally excited mirror vibrations will be time dependent, and so to estimate its importance, its magnitude is compared to another fundamental noise source in the amplitude of the light, the quantum uncertainty of the light power or photon shot noise. An order of magnitude estimate of the frequency bandwidth of the mirror modes (the full width at half maximum of the mirror resonances) is 10^{-3} Hz, and the laser light at the input of the Fabry-Perot cavity is of order 1 W. The corresponding shot noise in the light power at the input mirror is of order 10^{-11} W, whereas the fluctuations in this light due to scattering from mirror vibrations are of order 10^{-20} W. Thus the scattering term and all higher order terms of

Eq. 5-4 can safely be ignored.

Another assumption is that the TEM_{00} component of the distorted optical mode reflected from the vibrating mirror surface still resonates in the Fabry-Perot cavity, and that the light scattered into other modes does not. This is true if the change in the resonant frequencies of the optical modes of the cavity due to the vibrations of the mirror is less than the linewidth of the cavity. The frequency shift, $\Delta f = f\Delta l_n/l$ is of the order of 10^{-3} Hz for the modes of Figure 5-2, and the linewidth of the optical cavity under consideration is of order 10^2 Hz, so the light continues to resonate in the primary mode, and modes which were previously separate from the primary mode do not resonate.

5.2 Verification of the Effective Mass Coefficients

The numerical code used to calculate the effective mass coefficients was subjected to a number of consistency checks. The resonant frequencies were calculated using a FORTRAN code largely provided by J.R. Hutchinson. The calculated eigenfrequencies agreed with the theoretical work of Hutchinson [12] and also with the experimental work of McMahon [30] for the appropriate cylinder materials and dimensions. The mode shapes were checked for self consistency by comparing the elastic energy of deformation in the mode to the kinetic energy in the mode one quarter of a cycle later:

$$\int_{\mathcal{M}} \left[\frac{1}{2} K u_{ll}^2 + \mu \left(u_{ik} - \frac{1}{3}\delta_{ik} u_{ll} \right)^2 \right] dV = \int_{\mathcal{M}} \frac{1}{2} \rho \omega_n^2 \vec{u} \cdot \vec{u} dV. \qquad 5\text{-}9$$

K, μ, δ_{kj}, and ρ are the bulk modulus, shear modulus, Kronecker delta, and the density of fused silica. The integral is over the mirror volume, $\mathcal{M}$. u_{ik} is the strain tensor, defined in

terms of the displacement vector as

$$u_{ik} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_k} + \frac{\partial u_k}{\partial x_i}\right). \qquad 5\text{-}10$$

Repeated indices are summed. The effective mass coefficients can also be checked against analytical calculations in the one dimensional limit where the mirror is made very long and thin (Eq. 5-7). The effective mass coefficients of all modes in this approximation are 0.5. The numerical code passed these basic checks.

A simple experiment was performed to verify directly the calculation of the effective mass coefficients for modes which have acoustic wavelengths much larger than the beam spot size. In this case the laser spot can be approximated by a point. These modes can be driven by gluing a small magnet (with dimensions much smaller than the acoustic wavelength of the mode) to the center of the back of the mirror, which by symmetry has the same effective mass coefficient as the front. Using an interferometer to measure the mirror's response to forces applied by a current in a coil near the magnet, one can experimentally infer the effective mass of the mode. The apparent motion on resonance is

$$\bar{x} = \frac{Q_n}{\alpha_n m \omega_n^2} \bar{F} \qquad 5\text{-}11$$

where $\bar{x}$ and $\bar{F}$ are the root mean squared displacement and force. The force is proportional to the current, I, through the drive coils; therefore

$$\alpha_n \propto \frac{Q_n}{m\omega_n^2} \frac{\bar{I}}{\bar{x}}. \qquad 5\text{-}12$$

Such an experiment was carried out using one of the 10-cm mirrors, investigating the

first five modes of Figure 5-2. The resonant frequencies were found to agree with the calculation within 2%; the mirror had a wedged shape for optical reasons which made the thickness of the mirror ill-defined at the 1% level. Figure 5-5 shows a comparison of the measurement of the effective mass coefficients for these five modes to the values of α_n calculated using Eq. 5-12. The line on the figure indicates a fit to a direct proportionality (the current to force ratio was not calibrated independently). Figure 5-5 indicates that the effective mass coefficients are an accurate way to model the vibrational modes of a mirror, and that the numerical code used in the calculations was functioning properly.

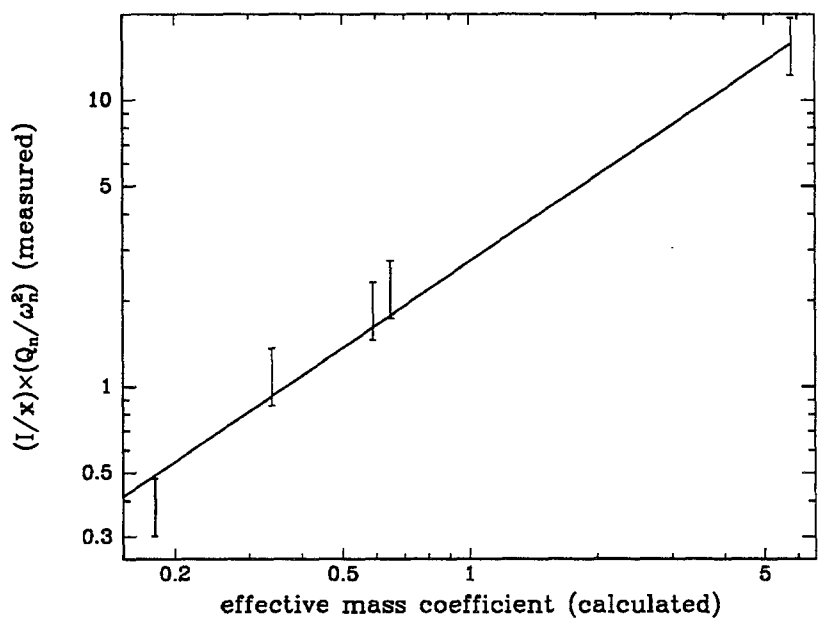

Figure 5-5: Comparison between calculated effective mass coefficients and the measured coefficients. The solid line is a direct proportionality.

5.3 Spectral Density of the Thermally Excited Motion

To estimate the thermal noise due to the vibrational modes of the mirrors, one must again return to the general thermal noise lineshape (Eq. 2-9), except the mass, *m*, of the

oscillator is replaced with the effective mass of each mode, $\alpha_n m$. The frequency band of interest for earth based gravitational-wave detectors (10 Hz to 10 kHz) is generally below the resonant frequency of the lowest vibrational mode of the mirror, in which case the total thermal noise can be approximated by

$$S_x(f) \approx \sum_n \frac{4k_B T}{\alpha_n m \omega_n^2} \frac{\varphi_n(\omega)}{\omega}. \qquad 5\text{-}13$$

All of the parameters of Eq. 5-13 can be readily measured or calculated except the loss function. An approximate value of $\varphi_n(\omega)$ can be inferred from the results of several recent measurements. Q values of order 10^7 have been measured for the vibrational resonances of the 10-cm mirror (see Section 6.2). Furthermore fused silica oscillators with Q's of order 10^7 and resonant frequencies ranging from 1 Hz to 10 kHz have been built [31]. These results all give a $\varphi_n(\omega_n)$ which is of order 10^{-7} and approximately independent of the resonant frequency. If one assumes that the loss function is the same for all of these modes, then these results are consistent with a loss function of fused silica of $\varphi(\omega) \approx 10^{-7}$, independent of frequency. This value of $\varphi(\omega)$ is adopted for the purpose of estimating thermal noise for the remainder of this chapter.

This estimate of φ should not be taken as either the fundamental level or frequency dependence of the dissipation of fused silica, which is unknown at frequencies near 100 Hz and at room temperature, but rather a summary of what has been observed thus far. The dissipation depends on both the purity and the preparation of the sample, and there is no strong reason to believe that the dissipation could not be reduced if sufficient care were taken.

5.4 Addition of Modes

The total thermally excited motion of the mirror surface at a given frequency can be predicted from the sum of Eq. 5-13. This raises the question of how many modes must be included in the sum to achieve an accurate estimate of the thermal noise. Figure 5-6 shows the cumulative contribution to the thermal motion at 100 Hz from the first one hundred axisymmetric modes for the 10-cm mirror, taking φ(100 Hz) to be 10^{-7} for all modes. The line indicates the cumulative contribution of all modes with resonant frequency less than *f*, and the ×'s indicate the individual contribution of each mode. The contribution of each individual mode decreases with resonant frequency due to the ω_n^2 term in the denominator of Eq. 5-13, but the mode density increases linearly with frequency (the axisymmetric modes form a two dimensional system). Also, there is a general decrease in the α_n terms in the denominator in Eq. 5-13. The result is a cumulative contribution which increases almost linearly with the maximum resonant frequency included. The approximate dependence of the effective mass coefficients on the resonant frequency needed to give this relation is 1/*f*, plotted as the dashed line in Figure 5-3. More than 100 modes were not included because of numerical precision errors in the series solving the equations of elasticity when more terms were added.

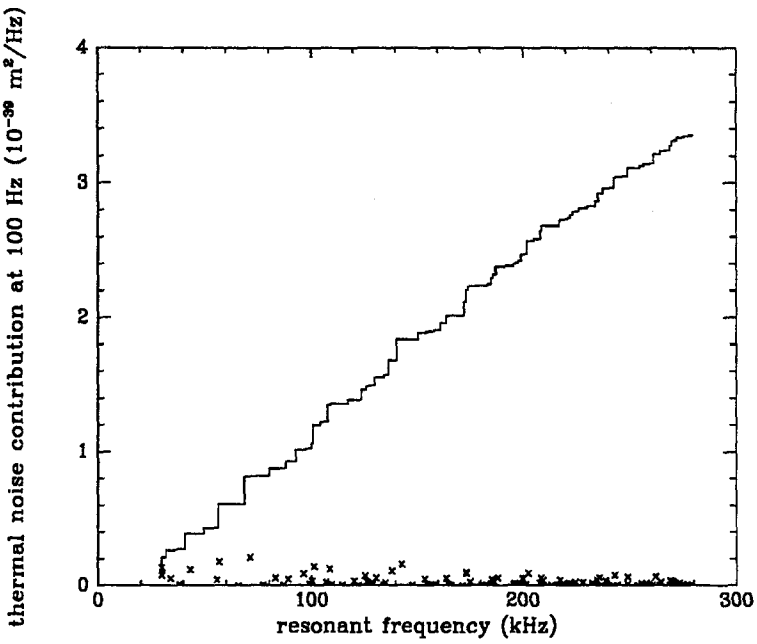

Figure 5-6: The contribution to the low frequency thermal noise of the modes of the 10-cm mirror with a 0.22 cm beam spot size. The line is the cumulative contribution and the ×'s are the individual contributions of the modes.

The thermal motion of the mirrors does not actually diverge, but convergence of the series depends upon the beam spot size. As an example, the initial LIGO interferometers will use mirrors with a diameter of 25 cm, a thickness of 10 cm, a mass of 11 kg, and a spot size of 2.2 cm on the vertex mirrors (the mirrors closest to the beam splitter). Figure 5-7 shows the convergence for an initial LIGO vertex mirror at 100 Hz with the same dissipation as in Figure 5-6. Figure 5-7 also shows the convergence for a 10-cm mirror with the same 2.2 cm size spot. The dot-dashed vertical lines on the figure indicate the frequencies at which half an acoustic wavelength becomes equal to the diameter at which the laser beam intensity is $1/e$ of its maximum. At this frequency the displacement of the modes, $\vec{u}(\rho, \theta)$, begins to cancel in the phase integral of Eq. 5-4, resulting in larger effective masses and smaller contributions to the thermal noise. The first vertical line applies to

transverse acoustic waves, and the second applies to longitudinal acoustic waves. Each mode is a combination of both transverse and longitudinal motion, and the sum levels off in the region between those two frequencies.

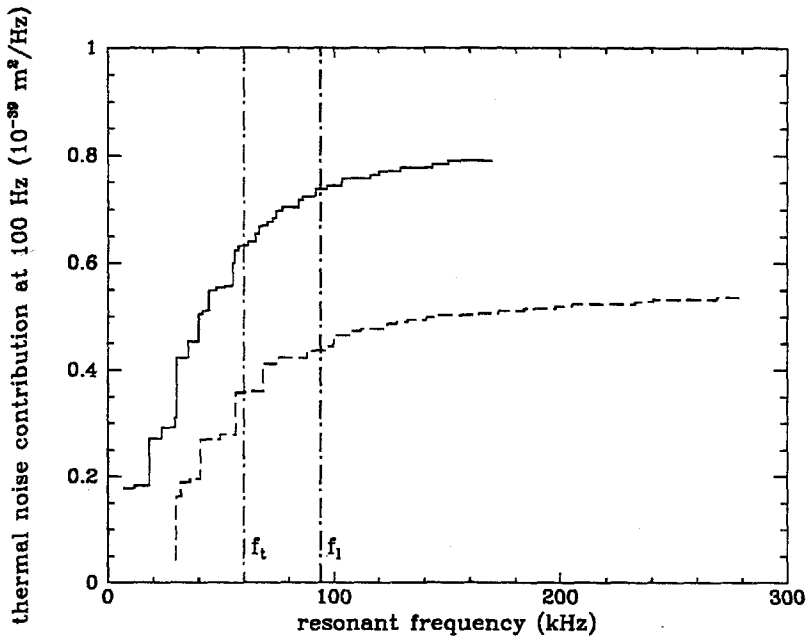

Figure 5-7: The cumulative contribution to the low frequency thermal noise of both the 10-cm (dotted line) and the LIGO 25-cm (solid line) mirrors with a 2.2 cm beam spot. The vertical dot-dashed lines indicate the frequencies at which an acoustic wavelength is of order the beam spot size for both transverse (left line) and longitudinal (right line) waves.

The thermal noise has the following pattern of convergence: the contribution increases approximately linearly with the largest resonant frequency included until the laser beam spot diameter becomes comparable to half an acoustic wave length. At that frequency the thermal noise sum levels off. In all, of order one hundred or more (depending on the geometry of the mirror) modes must be included to describe the thermal noise accurately; this number is far more than the few modes used in previous simpler models. Once this pattern of convergence was established, it was verified numerically for several mirror

and spot size combinations.

Figure 5-7 also indicates that mirrors of different geometries can have significantly different thermal noise levels. Typically more modes can contribute to thermal noise for larger mirrors and fewer modes can contribute for larger spot sizes. One should be cautioned, however, that the 10-cm mirror would be inappropriate for use in a full scale (km length) interferometer. Since the radius of the 10-cm mirror is only twice the spot radius, the optical loss at the edge of the mirror would be too large. Also, in determining the shape of the mirrors one must consider that the light passes through the vertex mirrors; therefore the optical loss in fused silica could restrict the allowed thickness of the mirrors. Optical requirements thus give additional constraints on the geometry of the mirrors.

The convergence pattern of the sum in Eq. 5-13 is independent of the loss function as long as all of the modes have the same loss function. This is because the relevant parameter for predicting the thermal noise is the loss function evaluated at the frequency of interest, in our case 100 Hz, and not the Q of the resonance, which could depend on the resonant frequency for an arbitrary loss function. The linear increase in the low frequency thermal noise as more modes are included does depend somewhat on the geometry of the mirror in that the axisymmetric modes are a two dimensional system. To meet this criterion, both the mirror diameter and its thickness must be greater than the laser spot size. This criterion is met in most realistic applications.

5.5 Effects of Moving the Beam Off Center

The numerical calculations described in the previous sections assumed that the laser was in the TEM_{00} mode and was centered on the mirror. This configuration puts the reso-

nant spot at an antinode of all axisymmetric modes and a node of all non-axisymmetric modes. By moving the beam off center, the contribution of the axisymmetric modes is decreased and the contribution of the non-axisymmetric modes is increased. Figure 5-8 shows the effect of moving the beam off center on the LIGO 25-cm mirror. 420 modes with circumferential order (number of nodal lines across the center; axisymmetric modes are order 0) of up to six were included, 60 modes each of circumferential order 0, 1, 2, 3, 4, 5, and 6. These modes included most of the modes with resonant frequencies below 100 kHz (approximately the maximum frequency at which one half an acoustic wavelength is equal to the spot diameter). The thermal noise is computed using the same parameters as in Section 5.4 (the laser beam spot size was 2.2 cm). Numerical precision errors prevented more modes from being included; only 60 axisymmetric modes were included to avoid giving those modes too much weight.

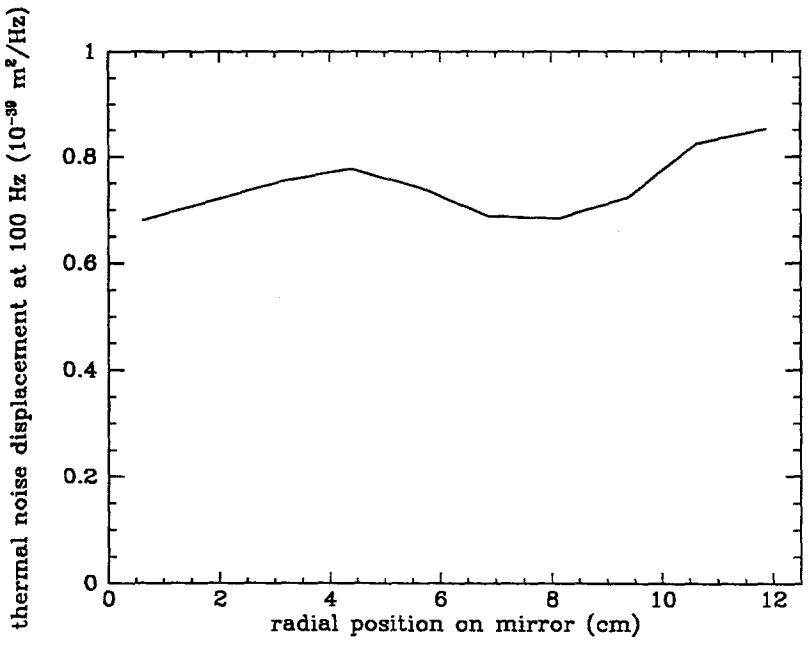

Figure 5-8: The thermal noise of a LIGO 25-cm mirror as a function of position on the mirror with 420 modes included.

The figure shows that except for spots quite near the edge, the thermal noise estimate is relatively insensitive to the exact location of the beam. The thermal noise power varies by 13% across the face of the mirror (an amplitude variation of only 6%). Since the motion of any spot on the mirror surface contains different contributions from each of the mechanical modes, the motion sampled by the light appears to be insensitive to exactly which combination of modes is sampled. It is not clear whether the small variation seen across the mirror surface is real or whether it is an effect of including only a finite number of modes. If all modes were included, then once the laser beam is more than a few spot diameters (or a few acoustic wavelengths of the highest frequency relevant modes) from the edge of the mirror, the mirror surface may effectively appear to the laser to be an infinite plane with all points having similar amounts of motion. Other interferometer details, such as noise in the alignment of the mirror [32], are likely to put much more stringent requirements on the positioning of the beam on the mirror.

5.6 Applicability of Results to Real Mirrors

Actual interferometer mirrors differ from ideal right solid cylinders. As mentioned in Section 5.2, the mirrors are wedge shaped for optical reasons. The surface might also be slightly curved. As described in Chapter 4, the mirrors are not completely free but are suspended as pendulums by fibers. To define the point of contact between the fiber and the mirror and to prevent rubbing, there may be attachments glued to the mirror. Furthermore, in order to apply forces to control the orientation of the mirror or to keep the TEM_{00} optical mode on resonance in the interferometer, there may be magnets attached to the mirrors.

These differences can be described as perturbations to shape and boundary conditions

of a right solid cylinder. Such perturbations can cause mixing of the modes [10], making the effective mass coefficients difficult to calculate. This problem becomes particularly severe for modes with high resonant frequencies, where the mode density is large (the mode density increases as the resonant frequency squared for a three dimensional system). However, the results of Section 5.5 indicate that the total thermal noise at low frequencies is relatively insensitive to the exact position of the laser on the mirror, or, put another way, the noise power is insensitive to exactly which combination of modes is sampled. Therefore it is reasonable to expect that the effects of such perturbations cancel when summed over a large number of modes so that the unperturbed calculation is expected to give a good approximation to the total noise power at frequencies far from the mechanical resonances.

Not only will the mechanical modes of the mirror differ from those calculated here assuming an ideal right cylinder, but also the optical modes of the Fabry-Perot cavity will differ from ideal Hermite-Gaussian optical modes. These deviations will be primarily caused by imperfections of the mirror surface. Such perturbed optical modes can be described in terms of their own orthonormal basis or can be expanded in terms of the Hermite-Gaussian orthonormal basis, and then the formalism outlined here can be used to estimate the thermal noise. The effect of a distorted optical mode would be to adjust the effective mass coefficients of the axisymmetric mechanical modes and, if the distortion were not cylindrical, to allow other mechanical modes to couple to the optical path length. Once again, since these results are insensitive to which modes are sampled, there would not be a strong dependence of the thermal noise power on the optical mode shape.

As will be seen in Chapter 6, attachments can also add dissipation to the system and

may decrease the Q's of the mirror by large factors. This additional dissipation may affect each mode differently and can be frequency dependent. An additional example of such a dissipation mechanism which has been studied in detail is resonant coupling between a suspended cylinder and its suspension wires [33]. Any damping mechanism which affects each mode differently in some systematic manner will alter the convergence pattern discussed in Section 5.4. The additional dissipation from attachments can likely be minimized by using careful experimental technique. These and other surface losses will generally cause less additional dissipation in mirrors which have larger ratios of volume to surface (or contact) area, such as those to be used in the full scale interferometers.

5.7 Significance of these Calculations

This estimate of the thermal noise level indicates that vibrational thermal noise may affect the sensitivity of advanced laser interferometer gravitational-wave detectors. The noise level shown in Figure 5-7 for a LIGO 25-cm mirror, corresponding to a displacement noise of 5×10^{-20} m/$\sqrt{\text{Hz}}$ or a strain noise of 1.3×10^{-23} 1/$\sqrt{\text{Hz}}$ at 100 Hz with all four mirrors included, lies between the "initial" and "advanced" detector noise levels of reference [1]. Vibrational thermal noise was not included in reference [1] because it was previously believed that this noise could be made smaller than other noise sources in the detectors.

This method of estimating the noise can be compared to previous commonly used methods which did not account for the geometry of the mirror or the laser beam. For direct comparison, take Saulson's example [8], which considered a mirror 10 cm in radius and 16 cm thick. This mirror has approximately the same mass as the LIGO 25-cm mirror.

Saulson chose the length and radius to minimize the thermal noise from the two lowest frequency modes. $\varphi(\omega)$ was assumed to be 2.5×10^{-7} and independent of frequency. Saulson's thermal noise estimate, using only the two lowest frequency modes and assuming an effective mass coefficient of 0.5, was (using his units) $S_x(f) = (2.5\times10^{-34}\ \mathrm{cm}^2)\times(1/f)$ [34]. Applying the three dimensional method of calculation and using Saulson's parameters, along with the LIGO laser spot size of 2.2 cm, $S_x(f) = (1.5\times10^{-33}\ \mathrm{cm}^2)\times(1/f)$. (This thermal noise level, when adjusted for the differing loss functions, is within 25% of the thermal noise estimate for the LIGO 25-cm mirror; in addition to thermal noise requirements, optical and fabrication considerations determined the geometry of the LIGO mirror.)

By considering the full mirror geometry, the thermal noise power is increased by about a factor of 6. The significance of this factor can be realized by comparing the internal vibrational thermal noise levels to other noise sources and by determining the change in the expected event rate. Saulson's estimate for the internal vibrational thermal noise was slightly above the published "advanced" LIGO noise goals. However, if one assumes a more optimistic loss function which has been achieved in some samples (3×10^{-8}) [35], then the resulting vibrational thermal noise level would be approximately the same as the "advanced" LIGO goals. The extra factor of six in noise power that results from taking into account the full geometrical shape of the vibrational mode thus changes the level of vibrational thermal noise from being near the "advanced" detector goals to being significantly above the "advanced" detector goals. If vibrational thermal noise were the limiting noise source to detecting a particular gravitational wave source, and this calculation indicates that it may be, then a factor of six in noise power would result in a reduction in the event rate by a factor of fifteen.

The noise levels discussed here should not be taken as the final estimate of thermal noise, but they do indicate that the noise goals of the "advanced" interferometers likely require test masses with lower dissipation than is available with present technology. The estimate of the dissipation, $\varphi(\omega)$, of 10^{-7} at 100 Hz was based on several recent experiments and reflects the levels of dissipation which have already been achieved. Data obtained so far probably do not represent the fundamental level of dissipation in fused silica, which is unknown. Meeting the thermal noise goals of advanced gravitational-wave detectors will require a better understanding of the dissipation in fused silica (or finding a better mirror material) and an improvement in the loss over what has been achieved thus far.

Chapter 6: Installation of the New Test Masses

6.1 Background

In late 1991 and early 1992 two major improvements were made in the 40-m interferometer's performance. First, new orientation control systems were installed which greatly reduced the noise below 1 kHz [32], and second, the laser power at the input of the interferometer was increased, reducing the shot noise and thus improving the interferometer performance at higher frequencies. These two improvements left a large region between 500 Hz and 1.5 kHz in the noise spectrum where the noise was slowly varying and of unknown origin (see Figure 6-1).

In attempting to explain this region of the noise spectrum, many potential noise sources were ruled out. Any sort of mechanical excitation above thermal noise transmitted through the suspension wires could be ruled out because the violin resonances were at thermally excited levels [11]. Frequency noise and amplitude noise on the light and electronics noise in the readout and coil driver were measured and eliminated as potential noise sources. Noise due to scattered light was a possibility because there was much scattered light in the restricted vacuum system of the 40-m interferometer at that time, but there were no models describing how the scattered light would couple to the interferometer noise at this level. In 1993, the Mark II version of the interferometer with its expanded vacuum system became operational. The change made a qualitatively radical change in the scattered light in the vacuum system, but the interferometer noise spectrum remained virtually unchanged above 500 Hz. This eliminated scattered light as a likely noise source. In addition, the improvements in the seismic isolation of the Mark II resulted in a reduction

of the seismic noise, extending the unexplained and slowly varying region of the noise spectrum down to 200 Hz (Figure 6-1).

One remaining possible source of the noise was thermal noise in the internal vibrations of the test mass. The test masses in the 40-m interferometer at that time were 8.8 cm long by 10 cm diameter fused silica cylinders with a 25 cm diameter hole bored through the center. The faces of the test mass were polished; the sides were not. A 38 mm diameter mirror was optically contacted to one side of the test mass. The mechanical quality factors of the test masses were measured and found to be lower than expected from measurements done before the test masses were installed. A summary of the measurements of the Q's of the axisymmetric modes of the old test masses are in Table 6-1. The Q measurements were made by sweeping a driving force through the resonance and measuring the FWHM of the response of the interferometer output. For the end masses the driving forces were applied by the driving coils and magnets attached to the masses; for the vertex masses the driving forces were electrostatic. These modes were identified as the axisymmetric modes by the following procedure. The approximate resonant frequencies of the axisymmetric modes were found using finite element analysis. The most strongly excited modes near the appropriate frequency were identified as the axisymmetric modes.

Table 6-1: Q's of old test masses as installed in 40-m interferometer.

mirror	resonant frequency (Hz)	Q
East End	27416 27673 34748	2700 5100 33100
East Vertex	27240 27580	800 2000
South Ver-tex	27395 27576	30000 1700
South End	27436 27996 34748	54900 14400 41100

In order to estimate the off resonance thermal noise from these Q's using the theory of Chapter 5, several assumptions must be made. First, the loss function was assumed to be independent of frequency. This assumption enables the loss function to be defined as

$$\varphi_n(\omega) = \frac{1}{Q_n}. \qquad 6\text{-}1$$

This assumption was driven primarily by a desire to get a thermal noise prediction that was close to the measured noise spectrum. Any weak frequency dependence could give the appropriate noise level. Second, the effective mass coefficients must be determined. Since these masses were not solid right circular cylinders, calculation of the coefficients is difficult. Effective mass coefficients of 0.5 were assumed for all modes. Lastly, one must decide how many modes to consider. In principle, all modes with acoustic wavelengths larger than the laser spot size could contribute to the low frequency thermal noise. In this case, all the measured modes had different Q's, and determining the Q's of the higher

order modes was not possible. The noise was assumed to be dominated by the few lower order low Q modes that were measured and the higher order modes were not included.

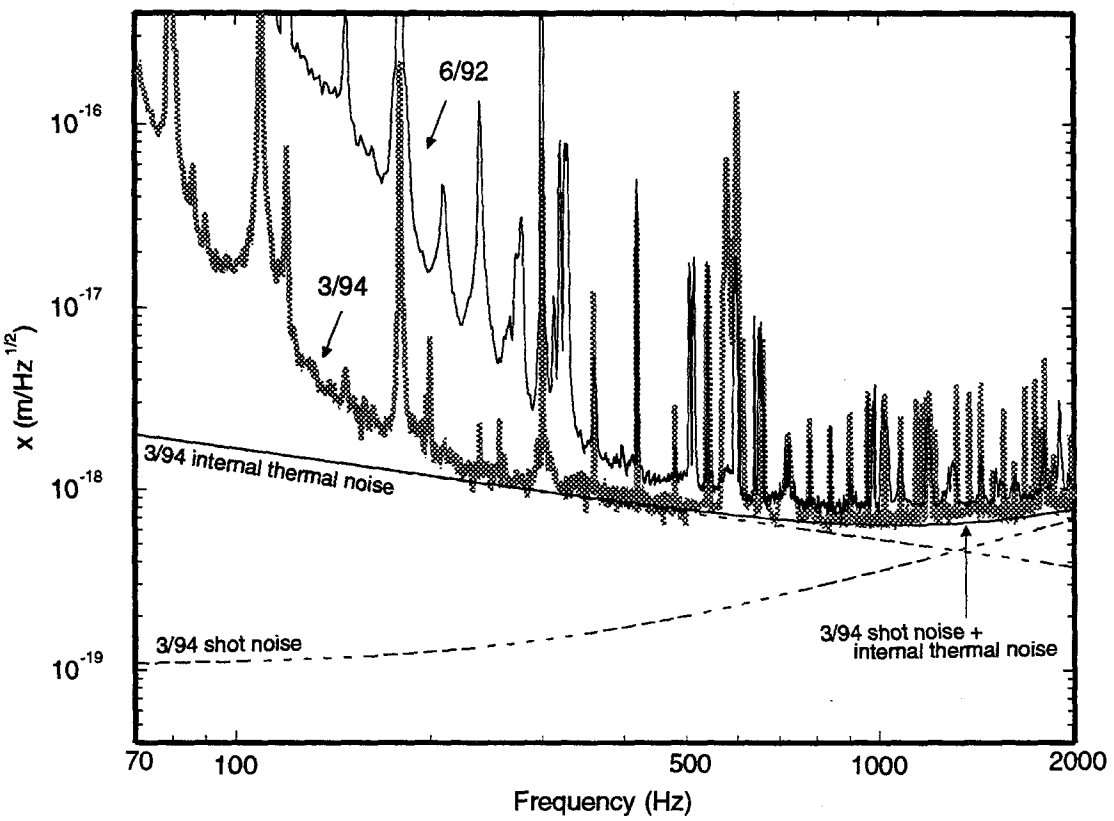

Figure 6-1: 40-m interferometer noise spectra in June 1992 and March 1994. Also shown are the internal vibrational thermal noise estimate and the shot noise.

Given these assumptions, the predicted noise level at 500 Hz is 8×10^{-19} m/$\sqrt{\text{Hz}}$, consistent with the observed noise level. The prediction of the internal thermal noise, along with the shot noise prediction and the measured noise spectrum are shown in Figure 6-1. The absolute level of the thermal noise was then adjusted (by 10%), keeping the shape fixed, to give the best agreement between the noise prediction of internal thermal noise plus shot noise and the observed noise spectrum. Such an adjustment was justified given the uncertainties of the prediction, which were estimated to be a factor of 3. The thermal

noise plus shot noise model gave excellent agreement to the noise spectrum above 250 Hz. With this prediction, reduction of internal thermal noise was made a priority after the shakedown of the Mark II interferometer.

6.2 Initial Experiments

An experimental program to improve the Q's of the test masses was begun in early 1992. The work was begun with the assumption that the old test masses would be replaced with new test masses that had been ordered. The new test masses were solid right fused silica cylinders, and 10 cm in diameter by 8.8 cm long. The test masses were polished on all sides, and the mirror coating was applied directly on the face of the mass so that there were no optical contacts. The new mirrors were of higher optical quality than the existing mirrors.

The suspension development apparatus, shown schematically in Figure 6-2, was used to evaluate the Q's of various test mass suspension configurations. The Q's were determined by exciting a particular resonance, turning off the excitation, and measuring the decay time of the oscillation. For configurations with magnets, a magnetic driver was used to excite the resonance; for configurations with no magnets, an electrostatic driver [36] was used. For each suspension configuration, the Q's of the five lowest frequency axisymmetric modes, shown in Figure 5-2, were evaluated.

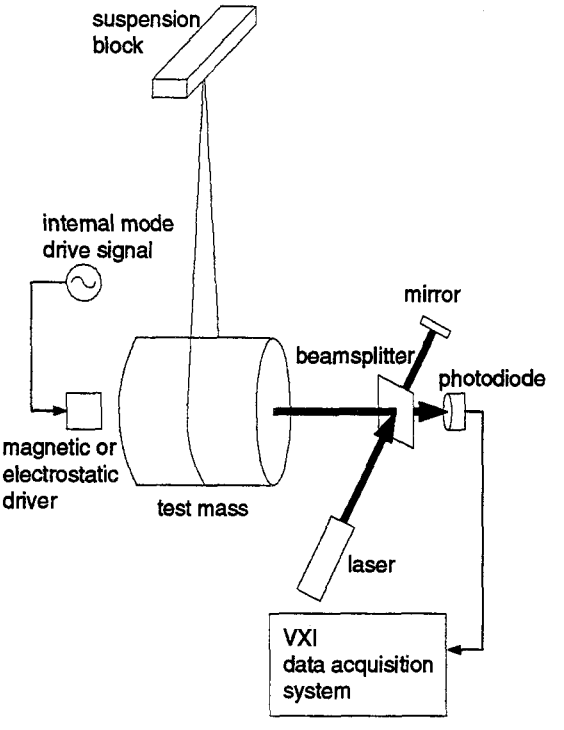

Figure 6-2: Schematic view of the suspension development apparatus.

The suspension configuration in which the least additional loss above the intrinsic loss in the fused silica might be expected, is to suspend the test mass in a single wire loop with no magnets attached. The measured test mass Q's of this configuration are given in column 1 of Table 6-2. The Q's of all modes except the 30.2 kHz mode are above 5×10^6. The 30.2 kHz mode is excessively damped by the suspension wire. From a close examination of the modes in Figure 5-2 it can be seen that the 30.2 kHz mode (the left 30. kHz mode) is the only mode with substantial axial motion at the center where the wire is situated. With no well defined attachment point, the wire is free to rub against the side of the test mass as the test mass vibrates in that mode.

For the next two configurations the suspension wires were moved to the positions which they would occupy in the interferometer. Two wire loops were placed 25 mm apart.

In one case the test mass simply sat in the wire loops, in another the points at which the wires leave the mass were defined by epoxying small (1 mm diameter by 3 mm long) fused silica rods on the test mass. Results from these two configurations are shown in the second and third columns of Table 6-2. Notice that the Q's in some modes are slightly degraded from the single wire loop configuration, but with the wire positions better defined, the 30.2 kHz mode's Q was actually improved. The configuration with the fused silica rods was adopted for the vertex masses in the 40-m interferometer (see Table 6-5).

The end test masses of for the 40-m interferometer required magnets be attached to them for control purposes. Column 4 of Table 6-2 shows the Q's when two 6 mm diameter by 6 mm long neodymium-iron-boron magnets were glued near the edges of the mass. The Q's were degraded by three orders of magnitude for some modes and varied by several orders of magnitude among the modes. To test whether this degradation was a property of the magnetic field, the magnets' dipole moments were reduced to 0.5% of their original values by baking. The data for the baked magnets is shown in column 5 of Table 6-2. The magnetic field strength had no effect on the Q's. In addition, to test for the possibility that the glue was responsible for the additional loss, two different types of glue were used: a cyano-acrylate glue and a vacuum sealant epoxy. There was no difference in the Q's depending on the type of glue. These results indicate that the dominant loss mechanism was due to the magnets themselves and not due to their magnetic fields or the glue joints.

Table 6-2: Test mass Q's: effects of attachments.

mode resonant frequency (kHz)	simple single loop Q (10^3)	2 wire loops Q (10^3)	with breakaways Q (10^3)	with magnets Q (10^3)	demagnetized magnets Q (10^3)
30.2	230	3,200	3,600	9.8	8.6
30.7	9,100	7,100	1,800	890	830
34.7	5,700	4,500	3,100	3.3	3.0
39.0	6,900		940		20.
44.0	5,500	8,100	7,300		180.

6.3 Reducing the Losses Due to the Magnets

The magnitude of the losses in the new masses due to the magnets were of the same order as the losses of the old existing masses in the 40-m interferometer (see Table 6-1). Therefore in order to make a substantial improvement in the interferometer's noise performance, the exact nature of the loss mechanism was investigated. There are three simple ways in which mechanical energy could be coupled into the magnets from the test mass. One is through motion perpendicular to the attachment joint, or axial motion; another is through motion parallel to the connection joint, or radial motion; and the last is through expansion across the connection joint, or strain motion. The amounts of each of these motions at the point of the magnet attachment was computed for each mode, and quantities that were inversely proportional to the energy were constructed. These constructed quantities--(1/axial motion^2), (1/radial motion^2), and (1/strain^2)--would be proportional to the Q if that particular type of energy transfer were the dominant loss mechanism. The

measured Q's were plotted against these constructed quantities to look for proportionalities.

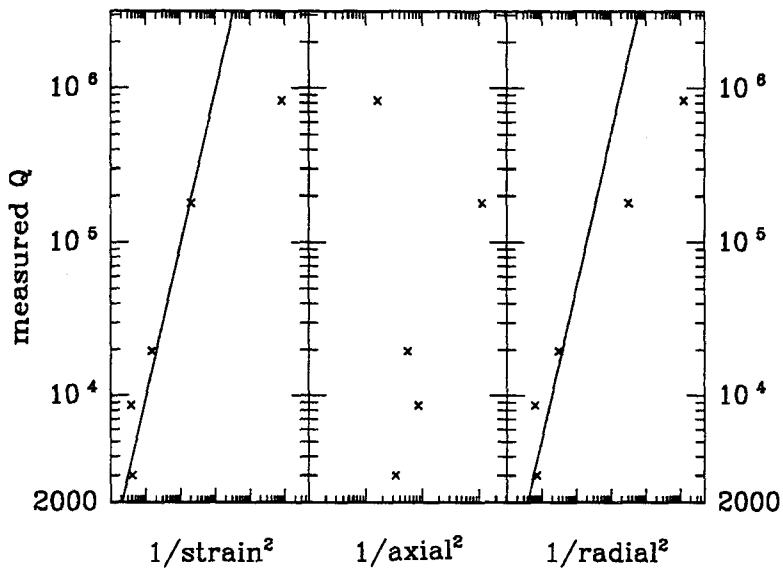

Figure 6-3: Comparison of measured Q with 3 possible loss mechanisms: couplings of strain energy, axial motion energy, and radial motion energy into the magnets.

Such a plot is shown in Figure 6-3. The demagnetized magnet data were used since those data were most complete.[1] At the factor of three level, four points in the strain damping model and three points in the radial motion damping model lie on a line of direct proportionality. The axial damping motion can be ruled out. For both the strain damping model and the radial damping model, the points that do not fit on the line are the higher Q modes, and in each case these modes have additional damping. There could be some additional damping mechanism that dominates the weakly damped modes but not the strongly

1. Chronologically, the measurement with the magnets attached was the earliest measurement made. At that time the interferometer was not sensitive enough to measure the Q's of the two higher frequency modes. Since the demagnetized magnet Q's were equivalent to the magnetized magnet Q's for three modes, they were assumed to be equivalent for all modes, and the magnetized magnet measurement was never repeated with the more sensitive interferometer.

damped modes. If the points had fallen on the other side of the proportionality line, or were insufficiently damped, then one could have ruled out the model.

To distinguish between the two models, one must use an additional information which is not on the plot. For the four lowest order modes, the face of the test mass undergoes nearly uniform expansion, so that the strain at a particular point is proportional to the radial motion (see Figure 5-2). Of the modes considered, only the 44 kHz mode does not have uniform expansion across the face of the mass. This fact is why the radial and strain damping model plots look similar except for at one point. That point has the Q of 180,000. For the strain damping model that point lies on the direct proportionality line and for the radial damping model it does not. For this reason, the strain damping model was the strongest candidate for the cause of the damping.

Strain energy is coupled across a surface area. Therefore, the most logical step to minimize the damping would be to minimize the area of contact between the magnet and the test mass. To verify that this worked, smaller (3 mm diameter by 5 mm long) magnets were attached to the mass and the Q's were measured. These results are shown in column 2 of Table 6-3. The Q's were increased by factors of approximately four for all modes. This improvement was consistent with the model that the damping was due to coupling of strain into the magnets, but was not sufficient to meet the goal of reducing convincingly the noise in the 40-m interferometer. It was also apparent that one could probably not use magnets which were small enough to meet the goal.

To further reduce the damping, the magnets were not attached directly to the test mass, but rather a spacer with a small contact area was inserted between the magnet and the test mass to attenuate the strain reaching the magnet. Aluminum (for speed and ease of

construction) spacers 1.5 mm in diameter and 3 mm long were designed. A schematic view of the spacer design as it was finally adopted for use in the 40-m interferometer can be seen in Figure 6-6. The Q's were measured with the spacers attached to the test mass and with no magnets (column 3 of Table 6-3). The results were very promising, and the large magnets were attached to the ends of the spacers. Those Q's are given in column 4 of Table 6-3.

The results were initially disappointing. The results also could not be explained with the simple axial, radial, and strain damping models. However, progress was made when the resonances of the aluminum spacer were considered.

Table 6-3: Test mass Q's: effects of the magnets and spacers.

resonant frequency (kHz)	with large magnets Q (10^3)	with small magnets Q (10^3)	with Al spacers only Q (10^3)	spacers and big magnets Q (10^3)	effective low frequency Q (10^3)
30.2	8.6	44.	1,600	86	190
30.7	830	4,000	3,600	250	4,500
34.7	3.0	9.1	970	53	200
39.0	20.		2,600	350	940
44.0	180	380	1,300	400	1,500

The observed losses could be explained by a model including a combination of damping due to strain transmitted through the aluminum spacer to the magnet and axial motion coupled into the longitudinal spring mode of the aluminum spacer. The energy coupled into the longitudinal spring mode as a function of frequency is

$$E(\omega) \propto \frac{\omega^4}{\left(\omega_0^2 - \omega^2\right)^2 + \omega^4 \varphi_{spacer}^2(\omega)} x_{axial}^2 \qquad 6\text{-}2$$

(the damping is proportional to the energy coupled into the spacer). A plot of such a function is given in Figure 6-4. Well below the resonant frequency, the spacer acts as a rigid body. It is not compressed and no energy is coupled into it. Well above the resonant frequency the spacer is compressed and stretched by the amount of the axial motion and the energy coupling becomes independent of frequency. Near the resonant frequency, the spacer resonance is excited and the damping is enhanced.

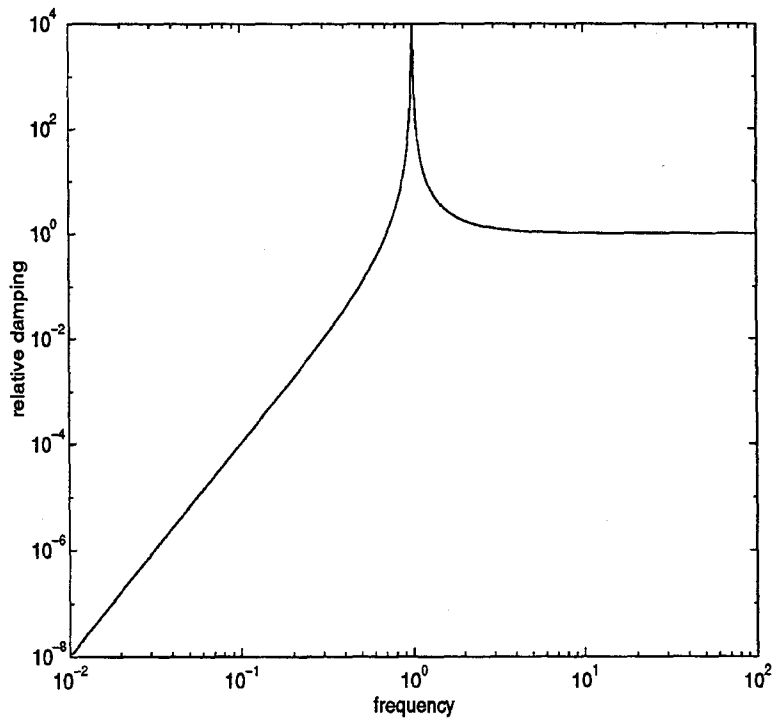

Figure 6-4: Damping function of test mass due longitudinal spacer mode. On this plot $\omega_0 = 1.$ and $\varphi_{spacer}(\omega) = 0.01.$

The measured losses are compared to the model of strain damping plus longitudinal

spring mode damping in Figure 6-5. The spacer resonant frequency was about 25 kHz. The ten points come from two sets of measurements on the five modes. One set of measurements was done with the magnets and spacers attached 13 mm from the edge of the mass and the other was done with the magnets at the edge of the test mass. These two sets included different samplings of the strain and axial motion of the test mass and gave the additional number of points necessary to do a multi-parameter fit to the damping model. All of the points agree with the model within a factor of two except a few high Q points which show excess damping.

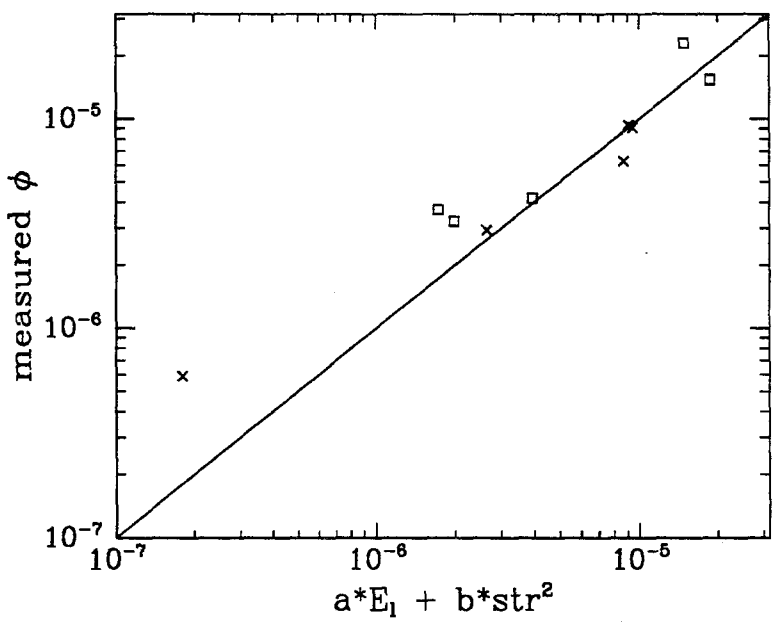

Figure 6-5: Comparison of measured losses with the model of a combination of strain damping and longitudinal spacer resonance damping. The crosses and the squares indicate measurements taken with the magnets at different locations on the test mass.

Since the frequencies of interest for the 40-m and LIGO interferometers (100 Hz-1kHz) are well below the spacer resonant frequency and since the spacer damping falls off quickly with decreasing frequency (ω^4), the spacer spring mode portion of the damping

can be subtracted for the purpose of predicting the low frequency thermal noise. Then one is left with the effective Q's shown in column 5 of Table 6-3. These effective Q's would give a significant drop in the internal thermal noise in the 40-m interferometer.

6.4 Choice of Magnet Size

The spacer-magnet assembly had one significant weakness. The assembly had a transverse flexural resonance in the spacer which was at 1.7 kHz, a significantly lower frequency than the longitudinal resonance. This resonance was near the unity gain point of the interferometer length stabilization servo and would make the servo unstable. Estimates of how well one could reasonably center the magnet on the spacer and in the coil indicated that for the parameters of this spacer, a gain bump of about a factor of three and an additional phase shift of about 90° should be expected. Measured transfer functions were consistent with these estimates. In order to push the transverse resonance to a sufficiently large frequency and maintain the improvement in Q, smaller magnets were used.

The smaller magnets (on the same size spacers) had a transverse resonant frequency at 5 kHz and a longitudinal resonant frequency at 42 kHz. The measured Q values and the effective low frequency Q values are given in Table 6-4 (this table gives the Q's of the test masses which were installed in the 40-m interferometer). The effective Q's for this system were better than for the larger magnet system, and it was adopted as the configuration to install in the 40-m interferometer.[1]

1. There was some risk in the decision to change the magnet size in that we were far from certain that the interferometer would acquire lock with the smaller magnets. In addition, even if the interferometer did acquire lock, the smaller magnets would require that in order to obtain optimal performance, adjustments in the servo would have to be made, thus coupling servo changes with test mass changes and complicating the project.

Several alternative spacer designs were investigated to increase the transverse resonant frequency and maintain the larger magnets. These included designs using hollow cylinders, designs with shorter spacers, and designs using three small rods widely spaced on the magnet. Steel and fused silica spacer material were also investigated. None of these alternatives gave as good results as the aluminum spacer with the small magnets.

6.5 Test Mass Design Criteria

Having achieved good results with a trial and error experimental program, the important aspects for future test mass design can be identified. The most important aspect of the design is the magnets. To achieve good Q results one should carefully consider how to minimize the magnet size and, if possible, consider alternative designs that do without magnets altogether. The smaller the magnet is, the smaller the surface area contact can be, and the smaller the mass on the spacer is, the higher the spacer resonant frequencies can be.

Once the magnet size is chosen, one must choose a magnet and spacer geometry. There are two driving design criteria. To maximize the Q, one must minimize the strain energy coupled to the magnet. Longer, thinner spacers will have less strain energy coupling. There is the additional constraint that the transverse resonant frequency must be large enough to avoid servo problems. Shorter, thicker spacers will have larger resonant frequencies. With these criteria the relative merits of various spacer aspect ratios can be explored using a finite element analysis.

In addition to these basic design criteria, quality control must also be considered. With small surface area contacts, one must worry about the quality of the glue joints. For

this spacer size, approximately 5% of the joints gave abnormally low Q's. With two glue joints per spacer, four magnets per mass, and four masses per interferometer, there is a high probability of getting a bad glue joint. Each set of glue joints was tested before installation and one was redone. If the surface area were made smaller, the reproducibility of the glue joint might become worse.

Another quality issue is that the long, thin spacer design is fragile. Installation of a test mass requires considerable handling through a glue joint test, a vacuum bake, an optical test, and the physical installation procedure. The success rate without breaking off a magnet was only 50%. Any new test mass design should make provisions for careful handling throughout these procedures.

Of lesser importance than the magnets are the wire attachments. The data show that by using relatively small fused silica wire attachment points the wires can be made to contribute relatively little additional loss. Probably this can be made even smaller if one were to use even smaller attachment points, but then the difficulties in handling and placing such small pieces must be considered.

6.6 Installation in the 40-m Interferometer

The new test masses were installed in the 40-m interferometer. The end masses, shown schematically in Figure 6-6, had the small magnets attached to them with the spacer design. The vertex masses had no magnets attached to them. The end masses had their wires placed 76 mm apart and the vertex masses had their wires placed 25 mm apart (the difference was due to differences in the existing control blocks of the two types of suspensions). In addition, the vertex masses had a 0.5° wedge and the end masses had a 2°

wedge. The difference in wedge angle made the resonant frequencies slightly different.

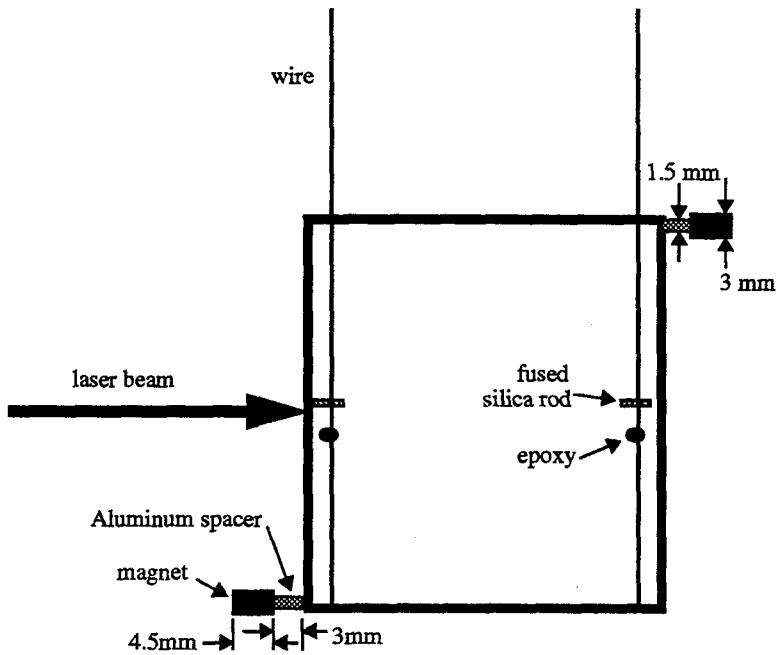

Figure 6-6: Schematic view of an end test mass as suspended in the 40-m interferometer. Vertex masses are similar except they lack magnets and the wires are closer together.

The measured Q's of all of the masses and the effective Q's of the end masses are shown in Table 6-4. The modes are the five lowest modes shown in Figure 5-2. The fifth mode is not listed for the end masses because the large wedge angle caused that mode to mix with a non-axisymmetric mode, making identification of that mode ambiguous.

Table 6-4: Q's of new test masses as installed in 40-m interferometer.

test mass	resonant frequency (Hz)	measured Q (10^3)	effective low frequency Q (10^3)
East End	30371 31178 35125 39369	200 680 43 79	510 >5,000 530 2,500
East Vertex	30196 30751 34776 38967 44035	3,900 8,100 1,700 3,300 5,800	
South Vertex	30192 30719 34693 38942 43997	5,900 6,100 1,700 2,200 4,600	
South End	30102 31104 34995 39190	190 850 40 60	430 >5,000 460 2,200

These Q's gave a dramatic decrease in the thermal noise. Figure 6-7 shows the noise spectrum of the 40-m interferometer after the installation of the new test masses and the adjustment of the servo. Also plotted on the figure is the internal vibrational thermal noise prediction for the new masses. To make this estimate, the effective low frequency Q's were used, and the loss function was assumed to be independent of frequency. There is no convincing reason to believe that the loss function should be independent of frequency for the new test masses, it was simply chosen because it explained the data from the old test masses. In addition to the 18 modes in Table 6-4, all other modes with acoustic wave-

lengths greater than the laser spot size were included. The loss function for these modes was assumed to be the average loss function for the measured modes of each particular mass. Also, the thermal noise due to the longitudinal resonances of the spacers were included. The spacer was heavily damped with a Q of about 25, which, with the factor of 5000 mass difference between the magnet and the test mass, made the spacer mode's thermal noise be of the same order of magnitude as the test mass modes' thermal noise.

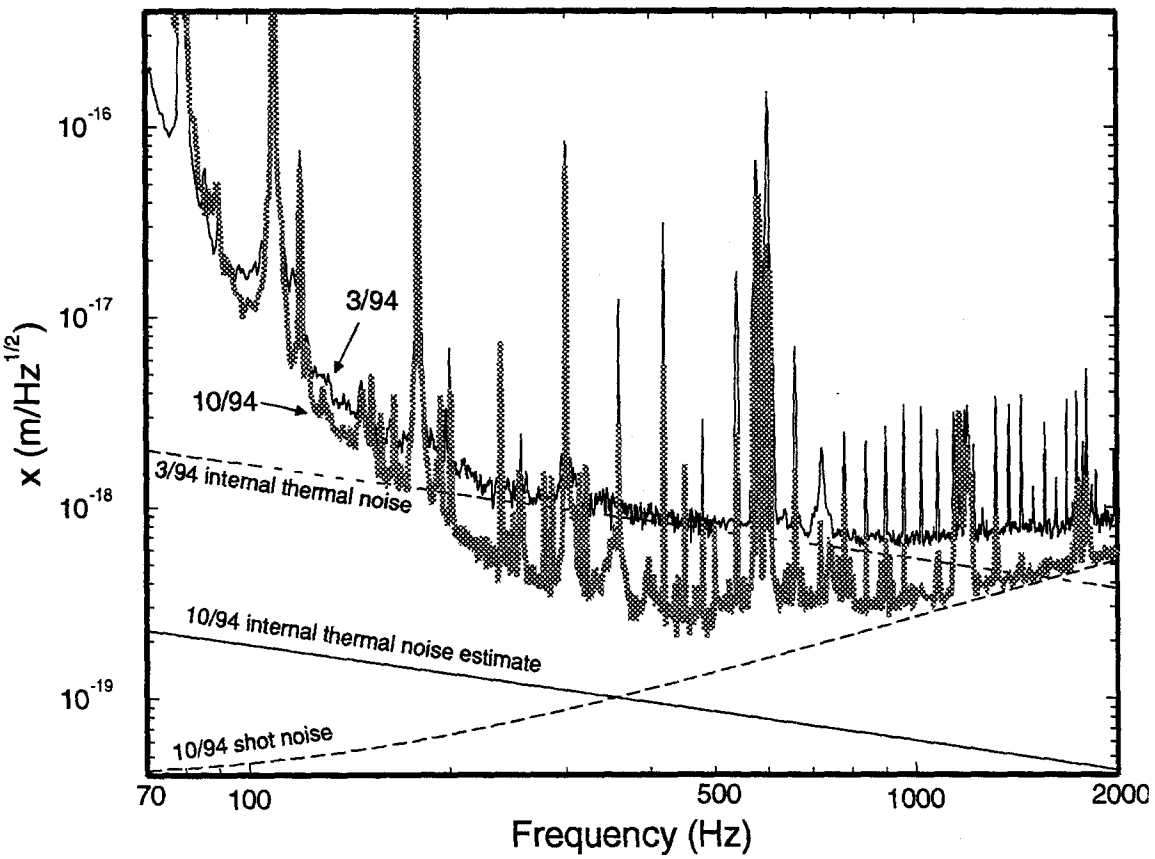

Figure 6-7: 40-m interferometer displacement noise in October 1994, along with the 10/94 shot noise and internal vibrational thermal noise estimates. For comparison the 3/94 noise spectrum and internal vibrational thermal noise estimates are also plotted.

For comparison the figure also includes the 3/94 interferometer noise and thermal noise estimate, as well as the new shot noise prediction. The shot noise prediction and the

improvement in the high frequency portion of the spectrum indicate the improved optical quality of the new mirrors.

6.7 Total Thermal Noise

In addition to replacing the test masses, the suspension wires were also replaced. The new violin mode Q's are given in Table 6-5. The pendulum mode loss function can be partially predicted from the violin mode Q's as described in Chapter 4. Recall that the violin mode losses do not reflect damping forces that act directly on the test mass. Such a force that is important in the 40-m interferometer is the eddy current damping of the pendulum due to the magnets attached to the mass and the magnet driving coils (see Section 3.2). For the 40-m interferometer, the force to current ratio was 2.025×10^{-2} newtons/ampere and the output impedance of the coil drivers was 70 ohms for the east end and 220 ohms for the south end (the differences are due to historical servo evolution). The noise spectrum along with all three thermal noise sources--internal vibrational thermal noise, suspension thermal noise in the wires, and eddy current damping thermal noise in the pendulum--are plotted in Figure 6-8. This plot is the complete thermal noise prediction of the 40-m interferometer as of October 1994.

Table 6-5: Violin mode Q's.

test mass	violin mode resonant frequency (Hz)	Q (10^3)
east end	571.60	80
	578.40	360
	594.25	50
	599.00	140
east vertex	581.10	94
	584.00	81
	595.98	117
	597.68	100
south vertex	585.02	70
	593.98	190
	597.72	140
	598.50	150
south end	578.72	70
	582.42	88
	597.98	78
	605.42	31

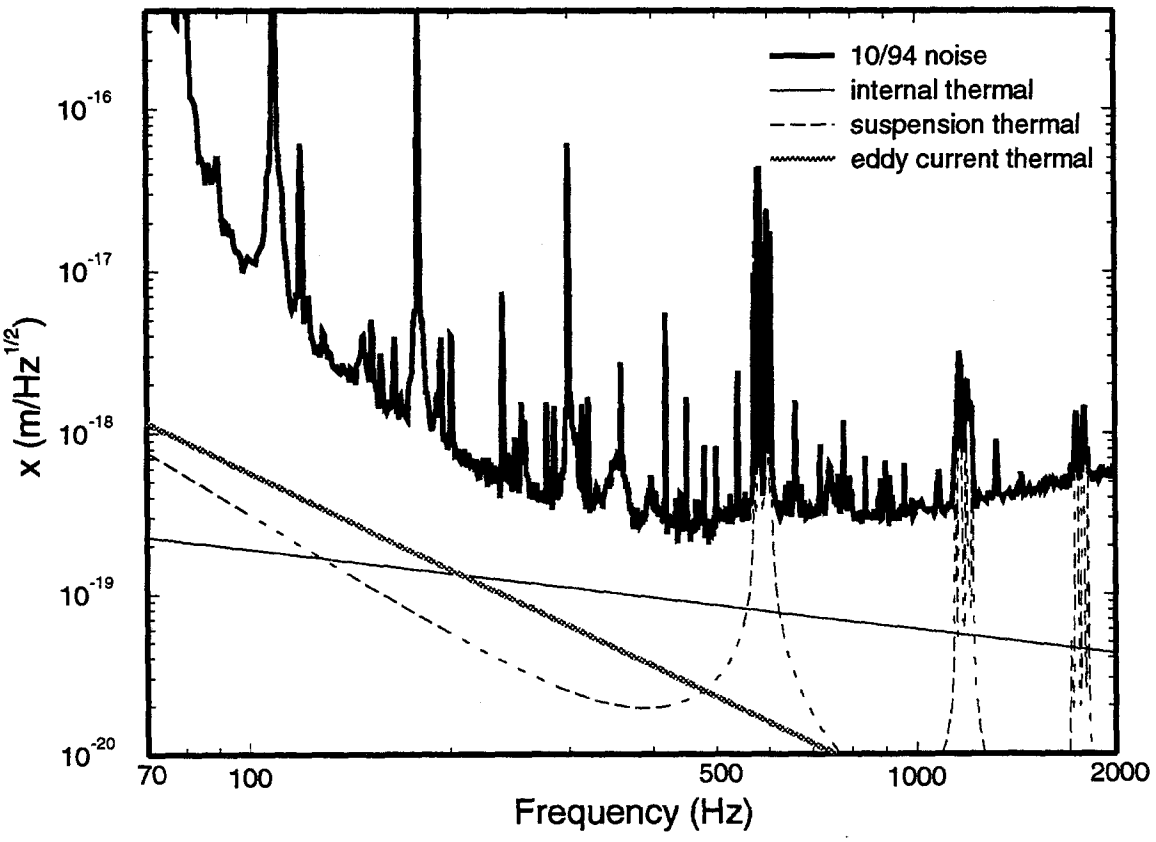

Figure 6-8: Total thermal noise of 40-m interferometer as of October 1994.

Chapter 7: Gravitational Waves from Inspiralling Compact Binary Stars

The best understood source of gravitational waves in the LIGO bandwidth is the final inspiral to coalescence of a binary neutron star system. The basic inspiral waveform can be reasonably well predicted from general relativity. One neutron star binary system, PSR1913+16, has been observed for a sufficiently long time to measure the orbital phase shift due to the emission of energy in the form of gravitational waves [37]. The measurement agrees with the predicted phase shift to 0.4%. Three neutron star binary systems which will inspiral in the lifetime of the universe have been observed, and from the statistics of radio pulsar searches in the solar neighborhood, the density of these star systems can be estimated, resulting in an estimate of the expected event rate of final inspirals as a function of the volume of space surveyed [38][39]. Since inspiralling compact binary stars are the only gravitational wave source for which both the waveform and the event rate are reasonably well known, LIGO's detectors are designed with those waves in mind. This chapter investigates how suspension thermal noise in particular can affect the sensitivity of an interferometer to inspiralling binary waveforms, and a simple search for inspiralling waveforms in the 40-m interferometer noise is described.

7.1 Inspiralling Binary Waveform

The quadrupole approximation of the waveform of an inspiralling compact binary star system is given by (see, for example, [40])

$$h(t) = \frac{Q\mu M}{Dr(t)}\cos[\phi(t)]$$
$$r(t) = \left(\frac{256}{5}\mu M^2\right)^{1/4}(t_c - t)^{1/4} \quad . \qquad 7\text{-}1$$
$$\phi(t) = -2\left[\frac{1}{5}\mathcal{M}^{-1}(t_c - t)\right]^{5/8} + \phi_c$$

h is the strain waveform, r is the separation between the two stars, and ϕ is the phase. Q, μ, M, and D are, respectively, a geometrical factor representing the relative orientation of the plane of the star system and the detector, the reduced mass ($m_1m_2/(m_1+m_2)$), the total mass (m_1+m_2), and the distance to the star system. $\mathcal{M}$ is the chirp mass, defined as $\mu^{3/5}M^{2/5}$. t_c and ϕ_c are the time of coalescence and the phase at coalescence. These equations set $G = c = 1$; i.e., 1 m = 3.336×10^{-9} s and 1 kg = 2.477×10^{-36} s.

The quadrupole approximation treats the masses as spinless point particles. All general relativistic effects are ignored except that gravitational radiation is allowed from an accelerating quadrupole mass distribution. The waveform's amplitude and frequency increase with time. The fourier transform of the wave form is given by

$$\tilde{h}(f) = \frac{Q}{D}\mathcal{M}^{5/6}f^{-7/6}\exp[i\psi(f)]$$
$$\psi(f) = 2\pi f t_c - \phi_c - \frac{\pi}{4} + \frac{3}{4}(8\pi\mathcal{M}f)^{-5/3} \quad . \qquad 7\text{-}2$$

The waveform has been computed to precision much greater than the quadrupole approximation, and there is an ongoing effort to calculate the waveform with still greater precision. The result of adding higher order (post-newtonian and above) corrections to the quadrupole formula is to significantly change the phase evolution of the system (especially when integrated over a large number of cycles), but not to change the amplitude

evolution much [41]. In addition, the higher order terms contain information about the individual masses and spins of the stars.

7.2 Optimal Filtering

In order to determine whether a signal exists at the output of the interferometer, $o(t)$, the output is passed through a filter. Any linear filter can be expressed in terms of a correlation integral:

$$c(t) = \int_{-\infty}^{\infty} o(t')\, q(t'-t)\, dt' = \int_{-\infty}^{\infty} \tilde{o}(f)\, \tilde{q}^{*}(f)\, e^{-2\pi i f t} df. \qquad 7\text{-}3$$

$q(t)$ is the filter function and $c(t)$ is the correlation. A signal is determined to be present if $c(t)$ is above some predetermined level. The simplest filter is a threshold detector, which looks for the output to cross some threshold. In this case $q(t)$ would be a Dirac delta function. The goal of optimal filtering is to choose the appropriate $q(t)$ so that the signal to noise ratio (SNR) is maximized.

If the form of the signal that is being measured, $h(t)$, is well known, then the filter which gives the maximum SNR is

$$\tilde{q}(f) = \frac{k\tilde{h}(f)}{S(f)}. \qquad 7\text{-}4$$

$S(f)$ is the spectral density of the noise and k is a constant chosen to give the desired normalization. The optimal SNR is then

$$(\mathrm{SNR})^2 = 4\int_0^{\infty} \frac{|\tilde{h}(f)|^2}{S(f)} df. \qquad 7\text{-}5$$

The SNR measured with an optimal filter will be the optimal SNR with a variance of 1.

From Eq. 7-5 it follows that the phase evolution of the signal is not important in terms of calculating the expected optimal SNR. Therefore for the purpose of analyzing an interferometer's noise performance in terms of predicting probable SNR's of inspiralling binary signals, the quadrupole approximation is adequate. However, for detecting signals and estimating the parameters of the compact binary star systems using Eq. 7-3, the phase evolution of the signal is very important and post-newtonian terms will need to be added to Eq. 7-1 [41].

Eq. 7-5 is a powerful tool for evaluating the detector noise in terms of signal detectability. In addition to giving the expected SNR for a given signal, it can also yield the SNR in a given frequency bandwidth. With this information one can determine what noise sources are important for detecting a particular signal.

As an example, the projected contribution, per unit bandwidth, to the SNR^2 (the integrand of Eq. 7-5) of an inspiralling binary star for an initial LIGO interferometer is plotted in Figure 7-1. The bold line indicates the region of the spectrum where the noise is dominated by suspension thermal noise (including both the pendulum thermal noise around 100 Hz and the violin mode thermal noise around 450 Hz. In the region between 130 and 400 Hz, shot noise is the largest noise source, however suspension thermal noise is within a factor of 2 of the total noise throughout that region. Virtually all of the signal of an inspiralling binary star system comes from regions of the noise spectrum where suspension thermal noise is an important noise source. Therefore, when the interferometer is designed, parameters of the suspension system which can be freely adjusted should be carefully considered.

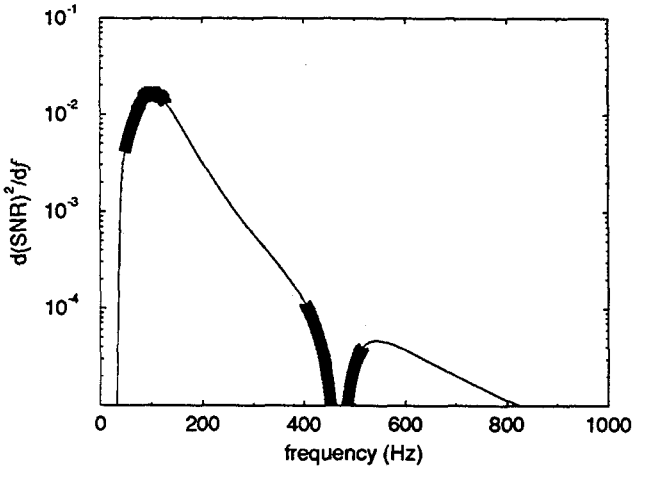

Figure 7-1: Contribution to the SNR per unit bandwidth of an inspiralling compact binary star gravitational-wave signal in a projected initial LIGO detector. The star system is taken to be two 1.4 solar mass neutron stars distance of 280 Mpc and optimally aligned with respect to the interferometer. The bold line indicates the region where the detector noise is dominated by suspension thermal noise.

7.3 Pendulum Length

The length of the pendulum of the LIGO interferometer can be chosen subject solely to the constraints of the vacuum system. From Eq. 4-13 and Figure 4-7 it is apparent that making the pendulum longer decreases its loss function and therefore decreases its thermal noise. Increasing the pendulum's length also decreases its resonant frequency, further decreasing its thermal noise as well as improving its vibration isolation factor. The disadvantage of increasing the pendulum length is that it lowers the resonant frequencies of the violin modes, thus putting more violin mode harmonics in the observing bandwidth. Decreasing the violin mode resonant frequency also decreases the reduced mass of the violin modes (Eq. 4-10), thus increasing the coupling of the violin mode to the interferometer. The situation is further complicated by the loss function of the violin modes being

decreased by an increase in the pendulum length. To illustrate all of these points, the total noise of the LIGO interferometer for three different pendulum lengths is shown in Figure 7-2. In the figure, the level of the shot noise is held fixed, and the seismic noise is varied according to the change in isolation with changing wire length. The level of the thermal noise are taken from Chapter 4 and the level of the shot and seismic noise is from [1].

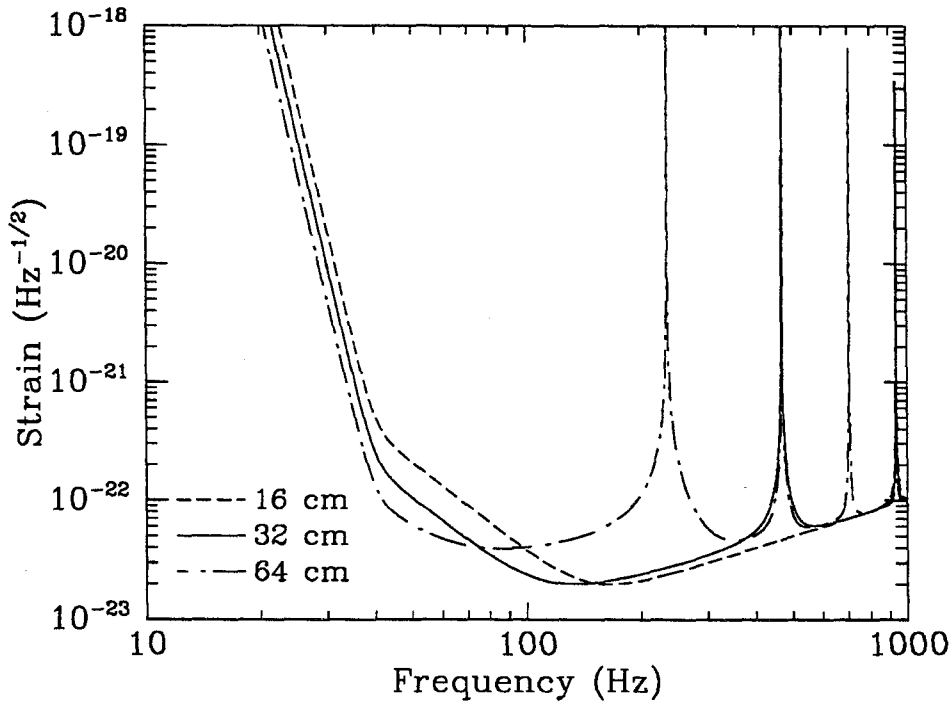

Figure 7-2: Projected initial LIGO thermal noise for pendulum lengths of 16, 32, and 64 cm.

In order to determine the optimal pendulum length, the suspension thermal noise as a function of pendulum length must be quantitatively compared to both the other noise sources (seismic noise and shot noise) and the signal that is to be detected. The obvious choice of signal is the wave from of an inspiralling neutron star binary since that is the best understood waveform. The optimal signal to noise ratio, computed using Eq. 7-5, as a

function of pendulum length is shown in Figure 7-3. The suspension parameters used are from the results of Chapter 4. The optimal pendulum length for the initial LIGO interferometer is 32 cm. The corresponding noise spectrum is shown as the middle curve in Figure 7-2. This optimal pendulum length depends on the gravitational wave to be measured and the other noise sources. A lower frequency gravitational wave would have a longer optimal pendulum length. Improving the seismic isolation would also make the optimal length longer. Decreasing the shot noise would shorten the optimal length. Finding a lower loss suspension system would favor using a longer pendulum.

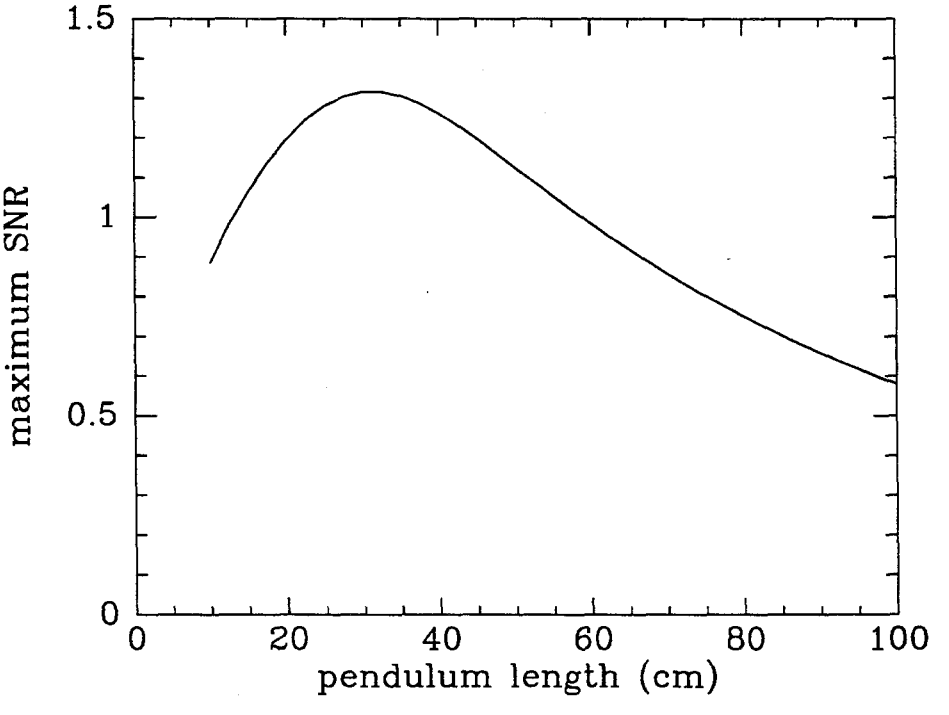

Figure 7-3: The maximum signal to noise ratio of the initial LIGO interferometer detecting two inspiralling 1.4 solar mass neutron stars at a distance of 280 Mpc with the binary system optimally aligned with the interferometer.

7.4 Filtering the Violin Thermal Noise

Once the optimal filter is built using Eq. 7-4, in principle it is straightforward to calculate the probability that a signal has arrived at a particular time using Eq. 7-3. Of the two mathematically equivalent choices of filtering in Eq. 7-3, time domain filtering (left integral) and frequency domain filtering (right integral), frequency domain filtering is the most computationally efficient for determining the probability that a signal has arrived at all times. To find the probability at all times using time domain filtering, $q(t'-t)$ must be evaluated for each time of arrival, t, and the integral of Eq. 7-3 must be redone for each time. With frequency domain filtering, the probability at all times can be found by doing just one integral since both $\tilde{o}(f)$ and $\tilde{q}(f)$ are independent of the arrival time. The price for this simplification is that two Fourier transforms must be performed on the data, one to put the data into the frequency domain and another to return the data to the time domain to determine the arrival times.

For very long data sets, a Fourier transform is computationally cumbersome. A fast Fourier transform requires approximately $N\log N$ multiplications (N is the number of points) and for efficiency, all of the points must be held in memory. Therefore it is more efficient to split very long data sets into shorter subsets. The price for this division is that in order to avoid splitting a signal between two subsets of data, each data subset must overlap the adjacent subsets by the length of the filter, n. Hence there is a trade-off between minimizing the size of the subsets and minimizing the fraction of the data that overlaps adjacent sets. This issue is discussed by Schutz [42]; the recommended overlap is approximately 25%, with the optimum being rather broad. The data set length is thus $N \approx$

$4n$.

An inspiralling neutron star binary will be in the initial LIGO detector's observing bandwidth (above about 70 Hz) for about 6 s. Hence the optimal length of each data set would be approximately 24 s. That time scale is shorter than the coherence time scale of the violin resonances ($\tau_{coh} = Q/2\pi f_0$). This raises the issue of how well the violin resonances can be filtered out of the data.

The simplest way to filter the data, that is to naively apply Eq. 7-3, has problems when sharp spectral features such as the violin modes exist. Figure 7-4 shows the noise spectrum of the violin resonance thermal noise (solid line). Also shown is the noise spectrum of a perfect sine wave with an rms amplitude equivalent to the rms amplitude of the violin resonance as computed when Fourier transformed in this naive fashion (often referred to as box windowing or flattop windowing). If the data sets were infinitely long, the sine wave would Fourier transform as a delta function. Each data set here is 40 s long. The effect of the box window is to smear the sine wave over a frequency band which is wider than the natural width of a violin resonance. Therefore, if this particular type of window were used, the thermal noise of the violin resonance would be spread out over a frequency band which is greater than its natural width, and the interferometer sensitivity would be compromised.

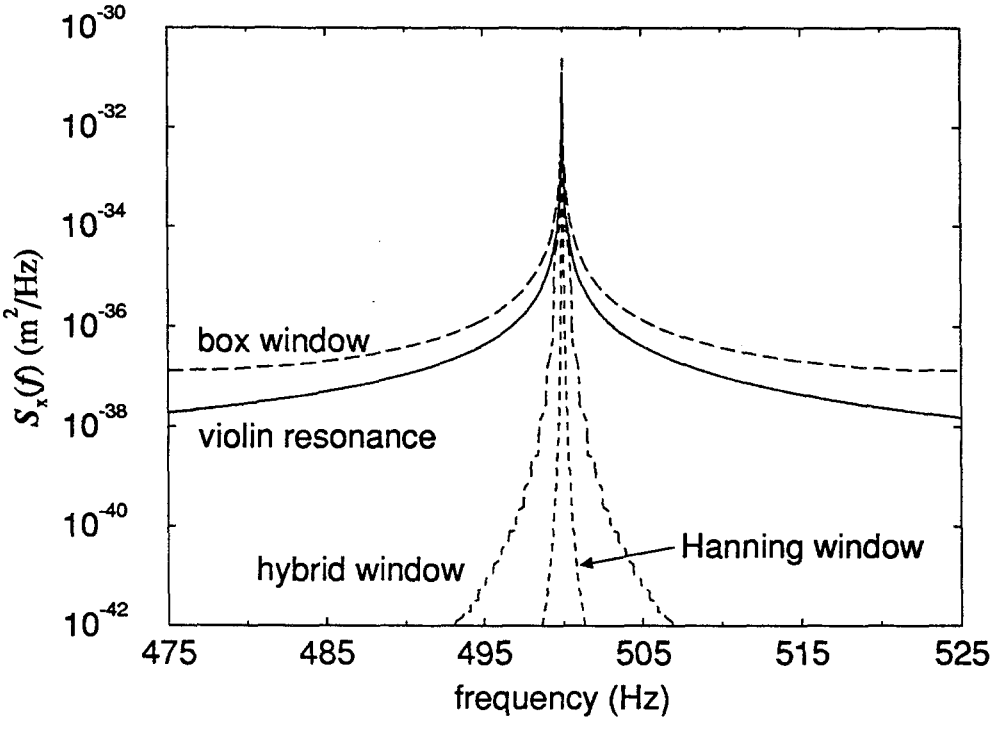

Figure 7-4: Comparison of the resolution of the box (long dashed line), Hanning (short dash), and hybrid (dot-dash) windows with the power spectrum of a thermally excited violin resonance (solid line). The sample length is 40 s. The violin resonance was taken to have a resonant frequency of 500 Hz and a loss function of 10^{-5}.

The mathematics of this smearing works as follows: Fourier transforming a finite length of data is equivalent to multiplying an infinite length of data by a window function and Fourier transforming the product. In this case the window is a box, represented by the function

$$\begin{aligned} w(t) &= 0 \quad t < t_1 \\ w(t) &= 1 \quad t_1 < t < t_2 \\ w(t) &= 0 \quad t_2 < t \end{aligned} \quad , \qquad 7\text{-}6$$

where (t_1,t_2) is the interval of the data. This function is plotted as the solid line in Figure 7-5. The Fourier transform of the output data is

$$\tilde{o}(f) = \int_{-\infty}^{\infty} o_{\text{ideal}}(t)\, w(t)\, e^{2\pi i f t} dt = \int_{-\infty}^{\infty} \tilde{o}_{\text{ideal}}(f')\, \tilde{w}(f-f')\, df'. \qquad 7\text{-}7$$

o_{ideal} is an idealized infinitely long data stream. Thus in the frequency domain, the data are convolved with the window function,

$$\tilde{w}(f) = \frac{1}{N^2}\left\{\frac{\sin[2\pi(t_2-t_1)f]}{\sin[2\pi(t_2-t_1)f/N]}\right\}^2. \qquad 7\text{-}8$$

This function is oscillatory, and the envelope falls off as $\tilde{w}(f) \approx 1/[2\pi(t_2-t_1)f]^2$. Only the envelope is plotted in Figure 7-4. When the sharpest possible spectral feature, an idealized sine wave of frequency f_0 in the time domain or a Dirac delta function in the frequency domain, is convolved with a window, the result is $\tilde{w}(f-f_0)$.

The problem with the box window is that it has very sharp edges in the time domain which transform into rather broad features in the frequency domain. The general solution to this problem is to chose a window which has smoother features. A window which is commonly used when the design criterium is to minimize the leakage of spectral information into the tails of the feature is the Hanning window [43]:

$$w(t) = \frac{1}{2}\left[1 - \cos\left(\frac{2\pi t}{t_2-t_1}\right)\right]. \qquad 7\text{-}9$$

This window is plotted in Figure 7-5, and its convolution with a perfect sine with the rms motion of a thermally excited violin resonance is shown in Figure 7-4. The Hanning window has significantly less leakage than the box window, and it has a resolution which is much finer than the violin resonance. Therefore, a Hanning window could be used to filter a violin resonance efficiently with negligible loss of information.

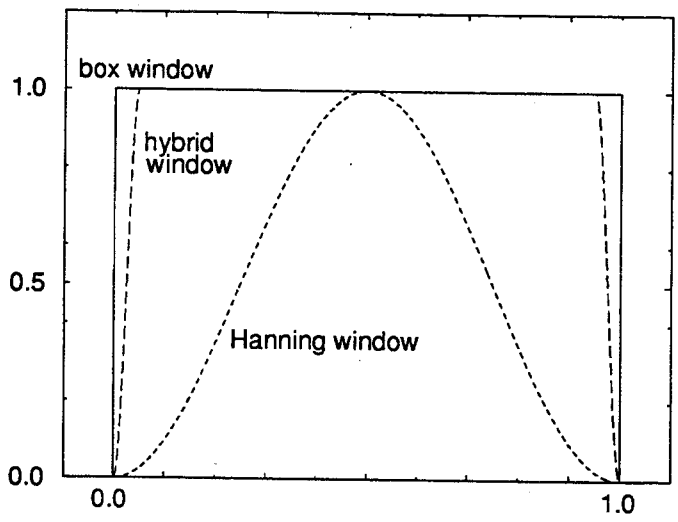

Figure 7-5: The shapes of the box (solid line), Hanning (dotted), and hybrid (dashed) windows. The time is scaled so that $t_1 = 0$ and $t_2 = 1$.

The Hanning window has one complication. It gives different weights to parts of the data at different times, making the results more difficult to interpret, whereas the box window weights the entire data sample evenly. One solution to this problem is to find a function $v(t)$ which is multiplied by $c(t)$ to correct for the effects of the window:

$$v(t) = \mathrm{rms}\left[\frac{\int o(t')\, q(t'-t)\, dt'}{\int o(t')\, w(t')\, q(t'-t)\, dt'}\right]. \qquad 7\text{-}10$$

This solution requires computing another function for each filter, holding it in memory, and doing an additional set of multiplications.

Another solution to the window problem is to find a compromise between the even weighting of the box window and the low leakage of the Hanning window. The window labeled as the hybrid window in Figure 7-4 and Figure 7-5 is formed by making a Hanning window which is one tenth of the length of the data sample. The window is split in the middle and filled with ones to create a window of the desired length. This window has

smooth edges and weights the middle 90% of the data evenly. Like the box window, this hybrid window's Fourier transform has oscillatory tails; the envelope is shown in Figure 7-4. This window has more leakage than the Hanning window, but the resolution is still much narrower than the violin resonances. Also, in order to ignore the 10% of the data set which is not evenly weighted and not lose any information, the data sets must be overlapped by 10%, in addition to the overlap for the length of the filter, for the entire filter length must be in an evenly weighted section. The overlap of the data sections is shown schematically in Figure 7-6. In order to use the hybrid window, longer data sets must be used, thus increasing the computational requirements. For example, the 6 s filter length of inspiralling neutron stars in the initial LIGO detector detected with a 10% Hanning-90% box hybrid window would require a data set length of 40 s (the data set length used in Figure 7-4) as opposed to a 24 s data set length in the ideal case. However, if the length of the Hanning part of the window is shorter than the filter, this hybrid window is more efficient than using a Hanning window and correcting the filter output using Eq. 7-10. This hybrid window is used in determining the data statistics that follow.

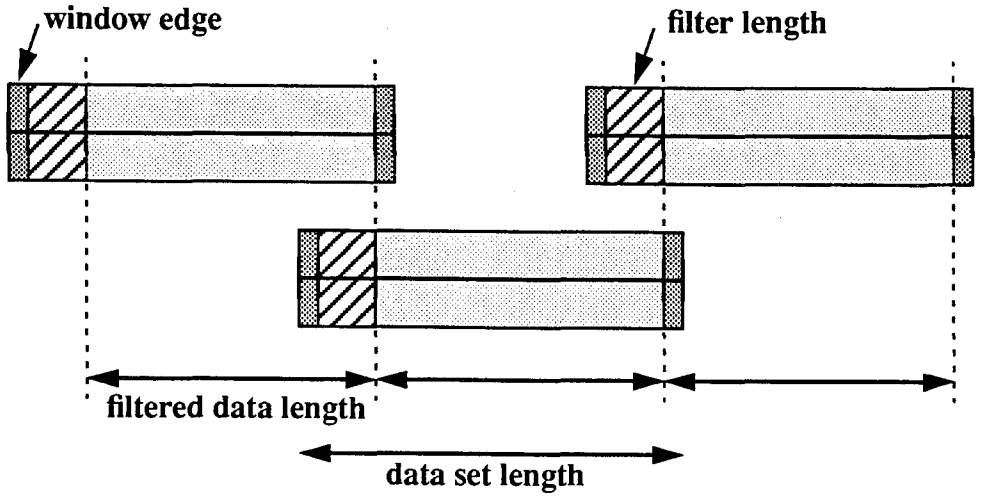

Figure 7-6: Schematic view of the temporal overlap of the data sets. Each data set must

overlap the adjacent set by the filter length and the length of time lost by the windowing.

7.5 Gravitational Wave Search with the 40-m Interferometer

Given the dramatic improvements that have been made in the 40-m interferometer noise performance over the past several years leading up to and including the replacement of the test masses (Chapter 6), a search for gravitational waves from inspiralling compact binaries was scheduled after the replacement of the test masses. Such a search was last done with the 40-m interferometer in 1987 by Smith [44]. The noise characteristics of the interferometer have changed dramatically since then. The amplitude of the noise near 500 Hz has been reduced by about a factor of 100, and the lower edge of the effective bandwidth of the interferometer has been reduced from about 300 Hz to about 100 Hz.

The method used by Smith took advantage of the fact that for the quadrupole approximation of an inspiralling binary waveform (Eq. 7-1), the phase evolution depends only on the chirp mass, $\mathcal{M}$, and the time [45]. By resampling the data at the appropriate variable time rate, the chirp mass can by mapped into evenly spaced frequencies. Then by doing a single Fourier transform, all chirp masses can be simultaneously investigated. Smith's algorithm is not appropriate for searching for signals with LIGO's interferometers. The method's primary flaw is that it cannot be adjusted to account for post-newtonian corrections which will be necessary for detecting signals and extracting parameters from LIGO's noise. For the purposes of the data analysis done here, the method used is to apply Eq. 7-3 in a straightforward manner after applying a hybrid window.

The data were taken over the course of the week of November 14-21, 1994. The sampling rate was 9868 Hz (a convenient multiple of the computer clock rate), and the data

were divided into sections approximately 4 hours long (corresponding to the amount of data that could be stored on an exabyte tape). Unfortunately, over the course of the week allocated for the data collection, the computer power supply failed, the laser logic controller failed, and one of the test mass mirrors suffered from a sudden increase in its optical losses, resulting in a limited amount of interesting data. For the purposes of this data analysis, the five best time sections were used for this analysis. The primary cause of lost data was not time that the interferometer was out of lock, i.e., both arms of the interferometer were not resonating simultaneously, but rather the time required for the violin resonances to equilibrate (about 10 e-folding times or 5 minutes) after being excited to high levels by the impulsive forces required to acquire lock (which took about 30 seconds). Good data sections had long consecutive lengths of times in which the interferometer was locked, so that data could be collected. After subtracting the time that the interferometer was out of lock and 5 minutes at the beginning of each locked section, there were 13.5 hours of data on the 5 tapes which were analyzed, representing 71% of the total time recorded by those tapes. In addition to the gravitational wave signal, several other environmental and control signals were recorded (see the Appendix).

49 filters were chosen to make the search. The filters are labelled by the mass of one of the stars in the binary system in solar masses, with the assumption that the stars were of equal mass. There was a finely spaced set of 36 filters between 1.2 and 1.5 solar masses, designed to search for signals of neutron star binaries in the galaxy with a maximum signal loss of 10%. There was also a set of 14 filters which were more coarsely spaced spanning, 0.12 to 10.17 solar masses as well as a pulse filter (i.e., an infinite mass filter). These filters were chosen to investigate the effects of the filter length on the noise and they

spanned the same range of numbers of cycles in the observational bandwidth (up to 10,000 cycles) as expected to be seen by the advanced LIGO detectors.

The outputs of nearby filters were nearly identical, so the discussion of the data will concentrate on only a few filters. Figure 7-7 shows the event rate histograms for four filters corresponding to masses of 4.4 (which had several cylces in the observing bandwidth), 0.97 (several hundred cycles), and 0.18 (several thousand cycles) solar masses and a pulse filter. An event is defined as a time when the filtered output crosses a threshold. In order to avoid double counting events, an event contained all of the time between when the filter output initially crossed the threshold and when there was 20 ms in which there were no threshold crossings. 20 ms was chosen because it was twice the period of the effective lower cutoff frequency of 100 Hz. The threshold was chosen to be a SNR of 3.5 so that there was a low probability that there would be a threshold crossing in any one 20 ms time sample. The SNR was computed by taking the rms of the entire output stream of the filter. The peak value of the event was then divided by this rms to get the SNR. The data were histogrammed, and the number of events in each bin was divided by the total length of the sample to get an event rate in a particular bin. This value is plotted. The SNR of events is histogrammed up to a SNR of 20; all of the overflow points were put in the 20 bin. The SNR's above 20 were scattered in amplitude; typical maxima for individual filters were near 100.

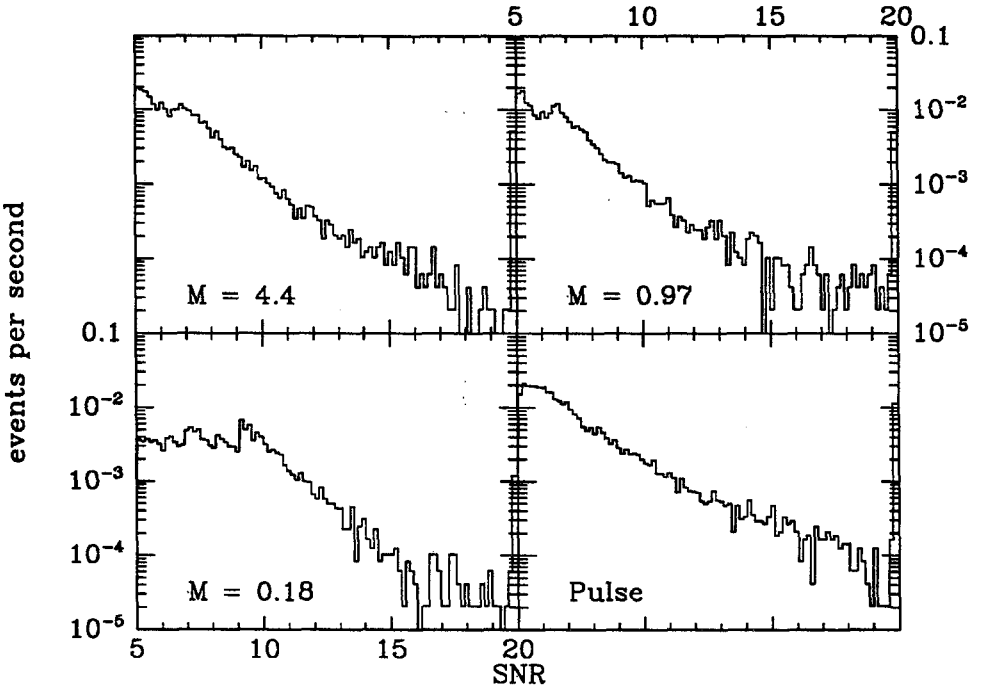

Figure 7-7: Histograms of event rates of four filters.

The histograms of Figure 7-7 indicate that the data are definitely NOT gaussian in nature. The probability of randomly getting an event with an SNR of greater than about 6 in this data set, assuming event lengths of 20 ms and that the data were purely gaussian noise, is approximately 0.01 (below the bottom edge of the histograms). Several thousand such noise events were found for each filter.

From histograms such as these, the probability of accidental coincidences in two or more detectors can be determined, and thus a threshold for detection can be set. As an example, assume that there are two interferometers which are aligned with each other and 3000 km apart. Suppose that the tolerable false alarm rate is 1 every ten years. Let each signal have a duration of 20 ms; the coincidence window at 3000 km is also 20 ms. Since the interferometers are aligned with each other, the signals should have identical ampli-

tudes, so assume that in addition to the signals being coincident in time, they must also have the same amplitude to within 2σ. Then the allowed event rate in any 2σ bin would be 4×10^{-4} per second. For the particular binning in Figure 7-7, there are 13 bins in a 2σ width, so therefore, on the scale of Figure 7-7, the threshold for signal detection for an error rate of 1 every ten years would be at 3×10^{-5}. This level is reached at a SNR of approximately 17 for M = 4.4, 15 for M = 0.97 and 0.18, and at 19 for the pulse filter. For this amount of data, the threshold level corresponds to about 2 events per bin, so the level of the threshold is correspondingly uncertain.

The absolute sensitivity of this search can be determined from the normalization of Eq. 7-4. A SNR of 1 in the 1.4 solar mass filter corresponded to a gravitational wave signal from an optimally aligned inspiralling binary neutron star at a distance of 400 kpc. If there existed two interferometers with the sensitivity of the 40-m interferometer which were aligned with each other, then with a threshold of 15, inspiralling binary stars could be detected with confidence at distances of up to 27 kpc. Not all star systems are optimally aligned with respect to the detector; considering all alignments, signals could be detected at an expected root mean squared distance of 11 kpc. Actually measuring the distance (as opposed to an average expected distance) requires at least 3 detectors sampling both polarizations of the waveform. For comparison, 8 kpc is the distance to the center of the galaxy. However, at the time of the search, there was only 1 detector running, and the largest event detected by the 1.4 solar mass filter had a SNR of 101, corresponding to a randomly oriented event at an expected root mean squared distance of 1.5 kpc.

The threshold SNR of 15 to 19 can be compared to the threshold expected from gaussian noise, which is approximately 4.5 for the above example. However, when analyzing

thresholds for the LIGO detectors where as many as several thousand filter may be used, the appropriate allowed error rate might not be one every ten years per filter, but rather one every ten years for all filters. This criterium would make the threshold higher, up to 5.5 or 6 (depending on the exact number of filters and the allowed false signal rate) for gaussian noise.

7.6 Time Variability of the Event Rate

The event rate changed dramatically from one time to another. As an example, Figure 7-8 shows the event rate of the output of the M = 0.97 filter. M = 0.97 is a convenient filter to study because it has as many cycles (about 600) above 100 Hz, the approximate lower cutoff frequency of the 40-m interferometer, as a typical 1.4 solar mass neutron star would have above 70 Hz, the approximate lower cutoff frequency of the initial LIGO interferometer. The figure show a time section (a single exabyte tape with approximately 3 hours of useful data) when the event rate was particularly low (14NOV94.2), a time section when the event rate was particularly large (19NOV94.4) and a time section when the event rate was about average (20NOV94.1). For the 19NOV94.4 time section, the fall off in events at SNR below 7 is because the estimator of the rms made from the first few data blocks was inaccurate and lower SNR events were not recorded. As more data were analyzed from the time section, the rms estimator approached its the final value, however the first few blocks missed the lower SNR events. The data show that there are some times when there are considerably fewer large SNR events than at other times. For example, the threshold for an error rate of one per ten years as described above would be at an SNR of 7 to 8 for the 14NOV94.2 data as opposed to 15 for all of the data. This factor of 2 differ-

ence in amplitude would result in a factor of 8 difference in the volume of space that could be surveyed by the gravitational-wave detector.

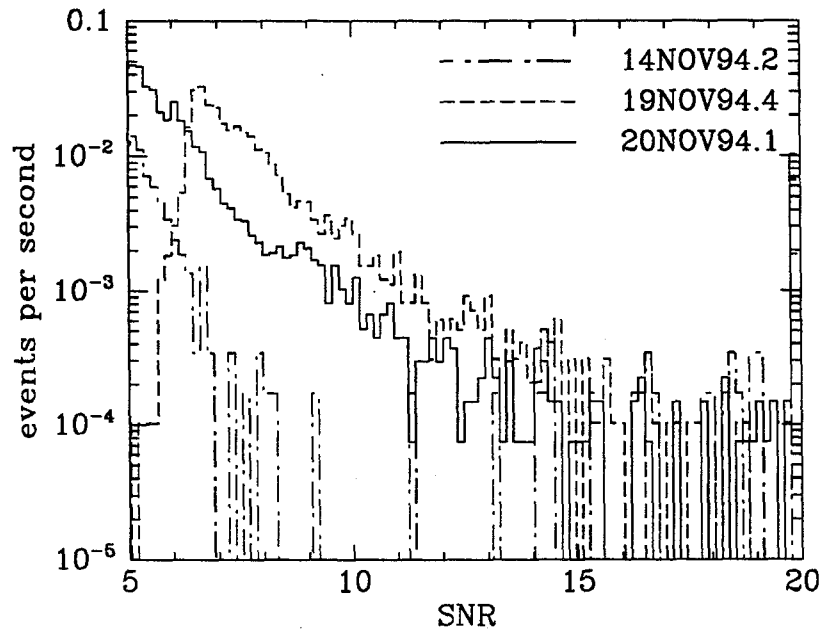

Figure 7-8: Comparison of event rate histograms for three different time sections.

The time scales over which the interferometer was in a high event rate state can be determined by examining the autocorrelation function of the output of the detector, shown in Figure 7-9. The figure shows that for two of the time sections, the interferometer was in states of more frequent noise events that lasted several hundred seconds. For the 19NOV94.4 time section, which had the most events, there is no sharp corner in the autocorrelation function indicating that the interferometer was nearly always noisy. For comparison, other time scales of the data are the 1 s length of the filter, the 5.1 s lengths of the data blocks on the tapes, the 26 s length of the block of data that is Fourier transformed, and the approximately 1000 s time scale that the interferometer is in lock. The autocorrelation functions of the other filters were similar to the M = 0.97 filter.

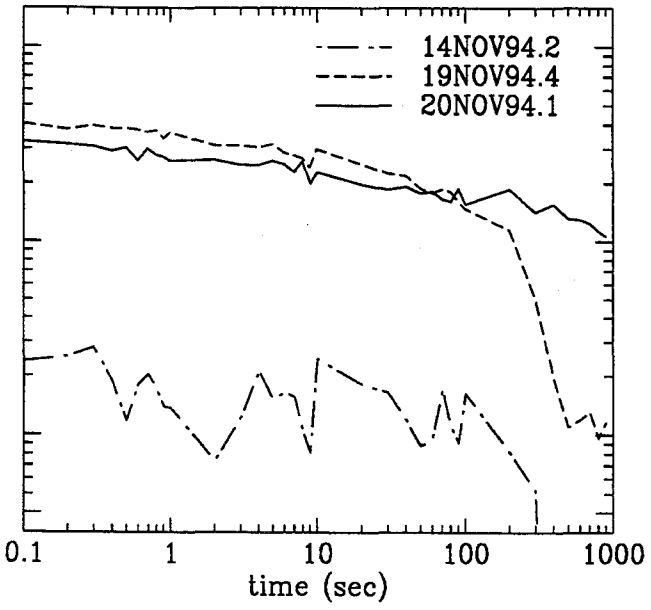

Figure 7-9: Autocorrelation function of the output of the M = 0.97 filter for 3 of the time sections. The vertical scale is not normalized.

Since the autocorrelation functions indicated that many of the noise events were clustered within time scales of a few hundred seconds, in order to determine if the times of the events were the same for all of the filters, a correlation can be made between the outputs of different filters. As an example, the correlation functions between the M = 0.97 filter and the pulse filter for three time sections are shown in Figure 7-10. The sign of the time shift is set so that the time of the pulse filter output precedes the time of the M = 0.97 filter output. These correlation functions are very similar to the autocorrelation functions of Figure 7-9. Figure 7-10 is typical of the correlation functions between different filter outputs. These data show that there exist periods in the data lasting several hundred seconds in which there are an abnormally large number of large SNR noise events.

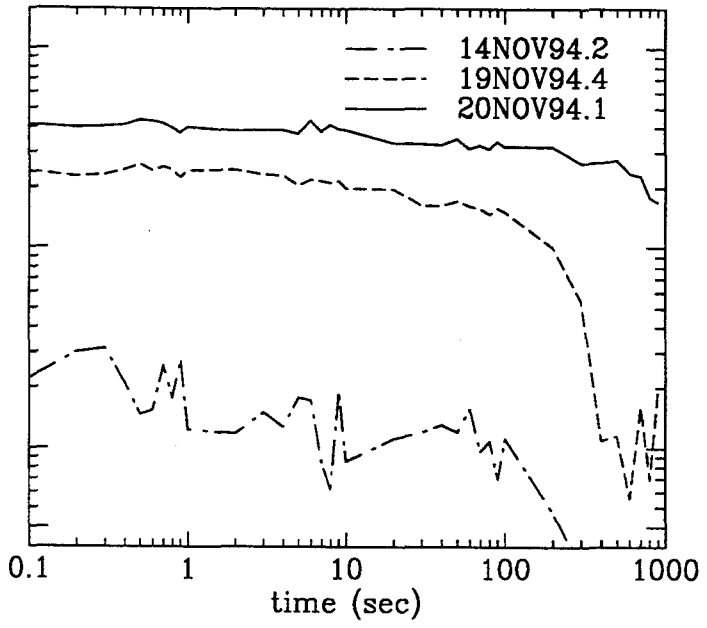

Figure 7-10: Correlation of the outputs of the M = 0.97 filter and the pulse filter.

The most surprising result of this study is that the use of optimal filters to search for inspiralling binary gravitational wave forms did not dramatically reduce the amount of excess noise events. The expectation was that since the optimal filters integrated over a large number of cycles, the effects of impulsive noise sources would be dramatically reduced. However, Figure 7-7 shows that the number of events above the expected gaussian background was only slightly reduced as the number of cycles of the filter was increased (smaller masses correspond to more cycles). The number of noise events was not dramatically suppressed by the optimal filters as compared to the pulse filter because many of the events in the noise filter were in fact not best described as impulsive events, but rather were events that lasted for of order 100 ms and had of order 10 cycles with characteristic frequencies in the range of 100 to 300 Hz. The inspiralling binary filters had most of their signal in the 100 to 300 Hz bandwidth, and the process of optimal filtering

would select the portion of the noise burst which best matched the phase of the filter and return that as the SNR. The noise bursts had the right frequency for the optimal filters, and were long enough that there was some increase of the signal due to the integration, yet short enough that the accumulated phase errors do not significantly reduce the integral.

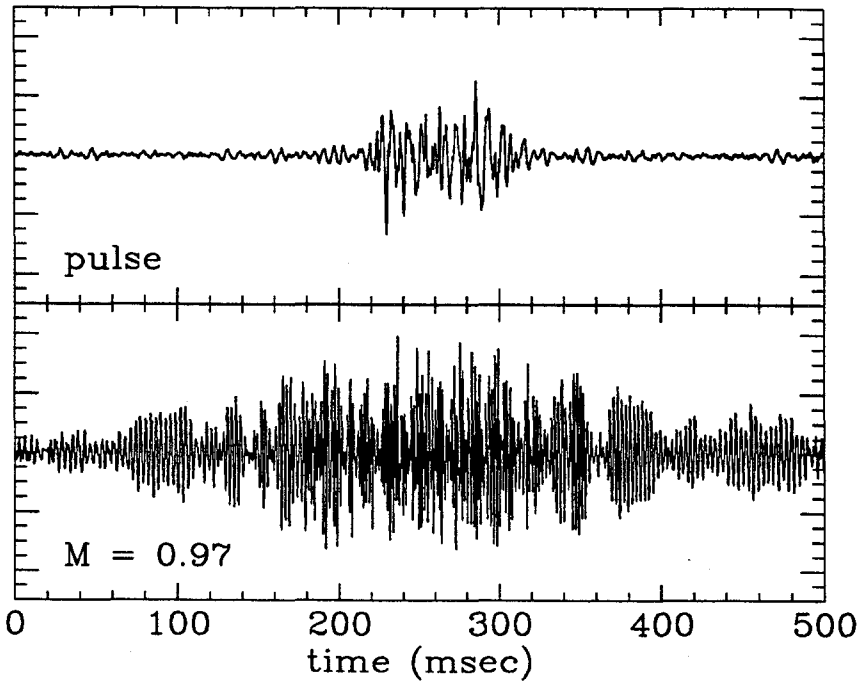

Figure 7-11: Example of a noise event as it is filtered through the pulse filter and the M = 0.97 filter.

As an example, a noise event is shown after it has passed through both the pulse filter and the M = 0.97 filter in Figure 7-11. The pulse has a length of about 100 ms and characteristic frequencies of several hundred Hz. When correlated with a finite length filter, the pulse gets spread over a significantly longer time span. This sort of noise event was fairly typical of the other events.

An attempt was made to find the cause of these noise events by examining the auxiliary signals that were also recorded (described in the Appendix). These signals were

examined for anomalously large noise that correlated with the noise events in the gravitational wave signal. Also, the direct current levels of the signals were investigated to determine if the noisy periods of time corresponded to times when control signals were near the edge of their dynamic ranges. No correlations between the noise events and the auxiliary channels were discovered; the cause of the noise events is unknown.

7.7 Summary

Many large (compared to the gaussian level of the noise) inspiralling binary events were found. These noise events were correlated with noise bursts in the interferometer that lasted of order 100 ms and had characteristic frequencies of 100-300 Hz, so, aside from the astrophysical implications, it is unlikely that the events seen at the detector output were actually due to inspiralling binaries. Also, the occurrence of such events were clustered in time over time scales of several hundred seconds. If two interferometers which were aligned with each other and had similar noise characteristics, then a threshold of about 15 σ could be used for detection to get no more than 1 false alarm every 10 years. Presumably, if the cause of the 100 ms noise bursts could be found and eliminated, then the sensitivity of a pair interferometers to inspiralling binaries could be improved by up to a factor of 3. The largest event in the inspiralling neutron star filters over the 13.5 hour search was 101 σ, corresponding to a randomly oriented inspiralling pair of neutron stars at an expected rms distance of 1.5 kpc. For comparison, the largest event of the previous search by the 40-m interferometer [44] was 22 σ, corresponding to a distance of 20 pc, over a search time of 36 minutes.

Chapter 8: Thermal Noise in the LIGO Detectors

Up to this point I have concentrated primarily on aspects of the thermal noise relevant to the 40-m interferometer. However, the primary goal of this work was to develop a suspension system that meets the thermal noise goals of the initial LIGO detector. Here I will summarize the results of the previous chapters and place them in the context of the LIGO interferometers.

8.1 Suspension Thermal Noise

From Figure 4-6 and Figure 4-7 the loss function of the steel music wire pendulum suspension at 100 Hz can be estimated:

$$\varphi_p(100\text{Hz}) \approx 1.5\times10^{-6}\left(\frac{35\text{cm}}{l}\right). \qquad 8\text{-}1$$

However, in order to scale to the LIGO interferometer, the change in tension and thickness of the wires must also be considered. If the losses in the pendulum mode are due to losses in the wire material, then the effects of these changes can determined from Eq. 4-1:

$$\varphi(\omega) \propto \sqrt[4]{\frac{r}{\mathcal{T}}}, \qquad 8\text{-}2$$

where r is the wire radius and $\mathcal{T}$ is the wire tension. In order to have as low a φ as possible, it is desirable to minimize the radius of the wire. Therefore the wire is always stressed to some fixed large fraction of its yield stress (typically near 50%) so that the wire radius becomes a function of the tension ($r \propto \mathcal{T}^{1/2}$). Therefore, only the additional tension dependence must be added to Eq. 8-1. The loss function of a pendulum suspended by steel

music wire then becomes:

$$\varphi_p(100\text{Hz}) \approx 1.5\times10^{-6}\left(\frac{35\text{cm}}{l}\right)\left(\frac{\mathcal{T}}{4\text{N}}\right)^{1/2}. \qquad 8\text{-}3$$

Recall that the results in Figure 4-8 indicate that the loss function for the pendulum mode may not be due to the losses of the steel. The other option was that the losses were due to imperfect clamping. If this were the case then Eq. 8-3 may not be applicable. However, the LIGO suspensions will require thicker wires which could make the details of the clamping less important. Therefore, there is reason to speculate that if the losses are due to clamping details, then the loss function may improve with thicker wires. However, the more conservative estimate of Eq. 8-3 will be adopted.

The scaling of the interferometer noise with length is rather complicated and has been discussed in Section 7.3, with the conclusion that for the noise parameters of the initial LIGO interferometer and the signal of an inspiralling compact binary star system, the optimal pendulum length is 32 cm. The mass of the initial LIGO mirror will be approximately 11 kg. The initial LIGO test mass will also be suspended in a single wire loop suspension, so the tension in the wire is 54 N. Hence the loss function of the pendulum in the initial LIGO suspension is estimated as

$$\varphi_p(100\text{Hz}) \approx 6\times10^{-6}. \qquad 8\text{-}4$$

In the region around 100 Hz, the loss function is approximately independent of frequency. This value of the loss function is therefore consistent with the goal of the initial LIGO interferometer which is equivalent to $\varphi_p(\omega) \approx 10^{-5}(f/100\text{Hz})$ in the region between 70 Hz and 150 Hz [1]. It follows from this reasoning that for the violin modes, $\varphi_v(\omega) \approx 1.2\times10^{-5}$.

In the frequency band of interest. Hence a steel wire suspension system is sufficient for the initial LIGO interferometers.

8.2 Internal Vibrational Thermal Noise

The scaling of internal thermal noise with the geometry of the test mass is rather complicated; it was described in detail in Chapter 5. For scaling from the 40-m interferometer to the LIGO interferometer, the most important point to recall is that due to the factor of 100 increase in length of the LIGO interferometer, the spot size of the resonant light will be increased by a factor of 10. Therefore, significantly fewer modes must be included in estimates of the LIGO internal thermal noise than were included in the estimates of the 40-m internal thermal noise, resulting in less noise in the LIGO interferometers.

Another complication in scaling the thermal noise from the 40-m case to the LIGO case is that the dominant loss function, the attachment of the magnets to the mass, was extremely sensitive to the particular geometry of each mode. Also, the control systems of the LIGO interferometer have not yet been finalized, so the necessary magnet size has not yet been determined. On one hand, the LIGO masses are 7 times larger than the 10-cm masses installed in the 40-m interferometer, requiring correspondingly more force for the same amount of control; on the other hand the remote sites have about an order of magnitude less background seismic vibration, so smaller control forces will be required to overcome this noise. The LIGO magnets will be assumed to be identical to the 40-m magnets.

For a given amount of uniform strain, the amount of energy in the mode scales as the volume of the body. Therefore for a given amount of energy, $\frac{1}{2}k_B T$, the LIGO masses will

have on average one seventh as much strain energy at the point of the magnet attachments. Therefore, on average the LIGO mass will be damped by a factor of 7 less than the 10-cm masses were damped due to the magnets. However, the control scheme of the LIGO interferometer requires 4 magnets on the face of each mass including the vertex masses, as opposed to 2 magnets on the end masses only for the 40-m interferometer. The damping of the LIGO end masses will thus be reduced to 2/7 of the damping of the 10-cm end masses, and the vertex masses (which contributed very little thermal noise in the 40-m interferometer) will contribute more noise. I will also (conservatively) assume that no mode will have a Q of greater than 10^7 since no Q higher than that was observed with the 10-cm masses even with no magnets attached. With all of these corrections, and adding the (unproven) assumption that the loss function is independent of frequency, then the expected average loss function of the LIGO mirrors is

$$\varphi(\omega) \approx 4\times10^{-7}. \qquad 8\text{-}5$$

The assumption that the losses are independent of frequency is conservative; the losses near 500 Hz cannot be much larger, for then excess thermal noise would be observed in the 40-m interferometer. However, the losses near 500 Hz could be smaller.

With this value of the loss function, the thermal noise can be estimated using the process of Chapter 5. In addition, the thermal noise of the magnet spacer longitudinal resonance must be included. The result is that the thermal noise due to the vibrations of the test mass approximately 1.0×10^{-19} m/$\sqrt{\text{Hz}}$ at 100 Hz. This noise level is comparable to other projected noise sources in the interferometer, and hence meets the goals of the initial interferometer.

8.3 Initial LIGO Thermal Noise

With these predictions for the thermal noise and the level of seismic noise and shot noise projected for the initial LIGO interferometer in reference [1], the total noise for an initial LIGO interferometer, shown in Figure 8-1, can be predicted. Also shown in the figure is the noise goal described in reference [1]. The current projected noise is slightly below the noise goal for frequencies below 120 Hz and slightly above the noise goal above 120 Hz. The current projected noise is larger than the noise goal at higher frequencies primarily because the noise goal did not include the thermal noise of the violin modes, whereas the current more complete projection does include the thermal noise of the violin modes. Also the vibrational thermal noise makes an additional contribution to the overall noise which was not accounted for in the noise goal.

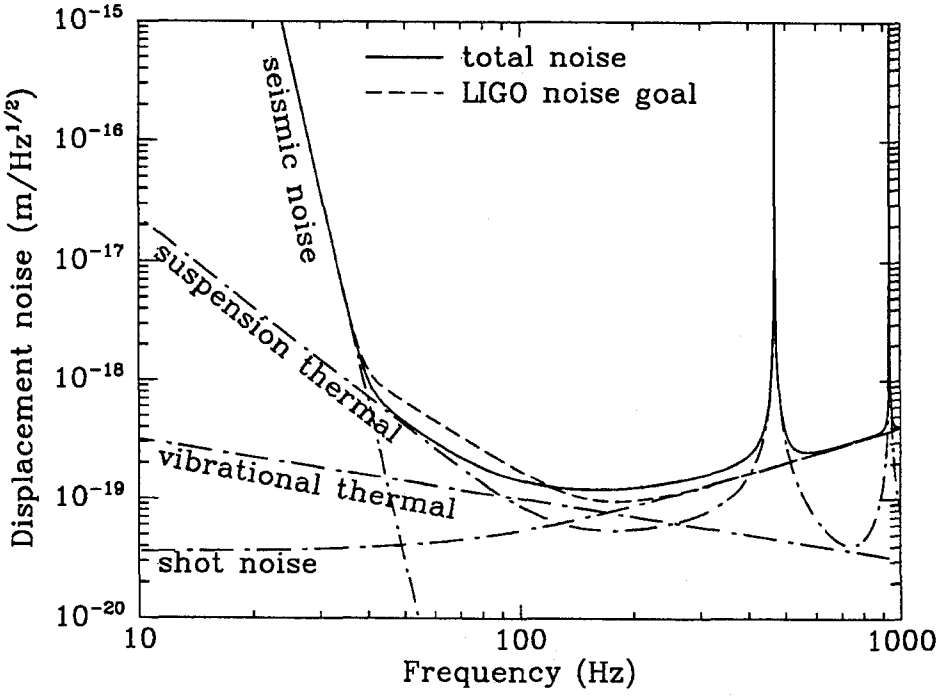

Figure 8-1: Noise of the initial LIGO interferometer. The dot-dashed lines are the noise

contributions of the primary noise sources. The solid line is the total projected noise. The dashed line represents the noise goal of the initial LIGO detector [1].

8.4 Conclusion

This thesis has been an investigation of thermal noise in the initial LIGO interferometer. Thermal noise has been observed both on resonance in the violin modes and far from resonance in the vibrational modes of the test mass. Although the fluctuation dissipation theorem and the basic quantum mechanics from which it was derived were not in doubt, it is nevertheless reassuring to find that the simple models used here to predict the thermal noise agree with the experimental data. The observation of thermal noise in the 40-m interferometer is yet another example of the remarkable ability of quantum mechanics to predict the behavior of macroscopic objects. Besides the simple observation of thermal noise, several important developments have been described in this thesis.

A steel music wire suspension system has been characterized and found to be suitable for the initial LIGO detector. Probably there is not much of a future in steel music wire suspensions beyond the initial detectors; preliminary work on fused silica suspension systems show that fused silica has very good prospects for reaching advanced LIGO goals [46]. However, the techniques described here for characterizing fully all of the modes of the suspension system in order to predict the full frequency dependence of the loss function will be useful well into the future.

A formalism has been developed for determining the coupling between the vibrational modes of the test mass and the phase measurement of the interferometer. The biggest surprise was that some high frequency modes are extremely well coupled to the

interferometer signal. The result of these strongly coupled modes is a significant contribution to the total thermal noise power from modes with resonant frequencies up to the frequency at which an acoustic wavelength corresponds to the laser spot size.

A technique was found to attach mechanically lossy magnets to the test masses without degrading the losses of the test masses too much. Using that technique, the Q's of the 40-m interferometer test masses were dramatically improved and the noise level of the 40-m interferometer was dramatically reduced. When scaled to the initial LIGO interferometer, the projected thermal noise from the internal vibrations of the test masses with the magnet attachments is less than other noise sources. However, these results indicate that there is not much of a future in using magnets beyond the initial interferometers, and hopefully will spur serious development of control systems which do not require magnets attached directly to the test masses.

The implications of suspension thermal noise for the detection of gravitational waves from inspiralling compact binary star systems was investigated. The noise of the 40-m interferometer was characterized after being optimally filtered to detect inspiralling binaries. The choice of pendulum length was shown to affect the sensitivity of the interferometer to such signals and that the optimal length is a convenient size for the LIGO vacuum system. It was also shown that the narrow bandwidth violin resonances can be effectively filtered out of the interferometer signal with a rather small cost in computing power.

The most important result of this work is the end prediction of the initial LIGO thermal noise shown in Figure 8-1. The primary goal of this research was to demonstrate that the thermal noise goals of the initial LIGO detector could be met. The suspension systems described in this thesis meet the initial LIGO noise goals.

Appendix: 40-m Interferometer

Many of the measurements of this thesis were made on the LIGO 40-m interferometer at Caltech. This appendix describes the configuration of the interferometer at the time of these measurements (1991-1994) and discusses how displacement measurements were made. Finally, the control signals that were monitored during the data run are described.

A.1 Interferometer Topology

A schematic of LIGO's 40-m interferometer is shown in Figure A-1. The laser light is delivered by the laser prestabilization system where the amplitude and frequency of the light are stabilized. The light passes through the mode cleaner which acts as a passive spatial filter and a length reference for the laser prestabilization system. Next a Pockels Cell applies 12 MHz phase modulation to the light so that RF reflection locking can be used [47]. After the beamsplitter, half of the light goes to the east arm and the other half goes to the south arm. On its way to a vertex mass (the masses closest to the beamsplitter), the light passes through a circulator, made up of a polarizing beamsplitter cube and a quarter wave plate, so that light (of the proper polarization) from the beamsplitter goes to the vertex masses but light reflected from the vertex masses is directed to the photodiodes. The light is then incident on the vertex mass and is coupled into the optical cavity formed by the vertex and end masses. The light does not return to the beamsplitter and is not interfered with light from the other arm.[1]

1. As this thesis is being written, there is currently a project under way in the 40-m interferometer to recombine the reflected beams at the beamsplitter.

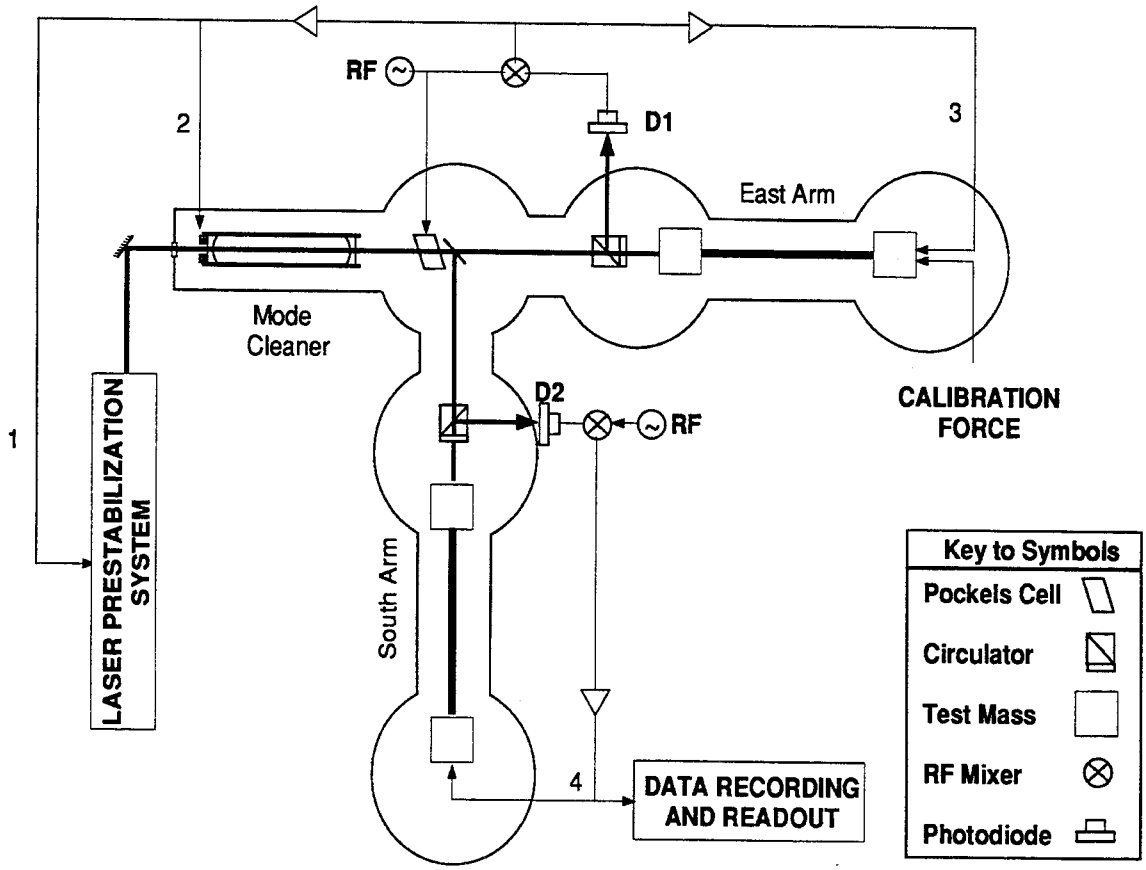

Figure A-1: Schematic of the control layout of the 40-m interferometer.

The interferometer control signals were extracted using the Pound-Drever RF reflection locking technique [47]. For the east arm, a low frequency (below 5 Hz) control signal (labelled as signal 3 on the figure) was fed back to the east end test mass to overcome the seismically excited low frequency motions of the test masses. The force was applied through currents in a coil which was near a magnet attached to the test mass. Control signals from about 5 Hz to 1 MHz (including the gravitational wave band) were fed back to a length control on the mode cleaner (labelled as signal 2) and a frequency control in the laser prestabilization system (labelled as signal 1). Therefore, at frequencies above 5 Hz, the laser frequency followed the length of the east arm.

The south arm was controlled independently of the east arm. The control signal was fed back solely to a coil-magnet actuator at the south end (control signal 4 on the figure). The control loop of the south arm had a frequency bandwidth of about 1 kHz. Therefore, the length of the south arm was forced to follow the length of the east arm (as encoded on the laser frequency) for frequencies between about 5 Hz and 1 kHz. The south and east arms were interchangeable, i.e., by rewiring the control systems the east arm could be forced to follow the south arm.

In addition to the control loops described here, there were many auxiliary control systems, controlling parameters such as the orientations and residual motions of all optics after the mode cleaner that can couple noise in the interferometer. The only auxiliary control system which had a direct bearing on the work of this thesis was the pendulum damping controller of the east end test mass, described in Section 4.5, which excited those violin resonances above thermal noise.

A.2 Extracting the Gravitational Wave Signal

The length measurement was made by interfering the light incident on the south vertex mass, which was controlled so that its frequency contained information about the length of the east arm, with the light resonant in the south arm, which contained information about the length of the south am, at photodiode D2. Since the control system forced the south arm length to follow the east arm length, differences in the lengths of the two arms were extracted by measuring the force applied to the south end mass. Therefore, the feedback voltage applied to south arm coil (signal 4) was a measure of the length difference between the two arms.

However, signal 4 was not directly proportional to the length difference. The signal was a measure of force, not length, and the signal also contained the response function of the control systems of the interferometer. The interferometer displacement, $\tilde{x}(f)$ was determined as follows:

$$\tilde{x}(f) = -k\frac{\tilde{o}(f)}{\omega^2\tilde{r}(f)}. \qquad \text{A-1}$$

Signal 4, $\tilde{o}(f)$, was divided by the interferometer response function, $\tilde{r}(f)$. $\tilde{r}(f)$ was determined by applying a known signal at a given frequency to the east end coil (the calibration force in Figure A-1) and measuring the amplitude and phase of the response in signal 4. The calibration signal was made much larger than the noise so that the difference in the arm lengths was due solely to the calibration force. This measurement was repeated at several different frequencies, and the response function was an interpolation of the results of those measurements. (The response function was a smooth function of frequency.) The resulting function was then divided by $-\omega^2$ to transform the calibrating force into a displacement. Lastly, the output signal was multiplied by a calibration constant, k, which transformed the arbitrary units into meters. The calibration constant was most easily determined by measuring the phase shift of the light in the east arm in response to the calibrating force applied to the east end test mass. Eq. A-1 was used for the measuring the noise spectra in Chapter 4 and Chapter 6.

For frequencies well above 1 kHz, the control system for the south arm had no gain. Therefore there was a larger measurement signal to noise ratio before the control amplifier than after it. However, Eq. A-1 could still be used with a different response function which was independently determined. This signal was used to measure the test mass Q's in the

40-m interferometer (described in Chapter 6).

When doing the optimal filtering described in Chapter 7, applying Eq. A-1 to every subset of data would be cumbersome. Therefore, rather than calibrating the output of the interferometer, the filters were simply transformed from displacement space to output voltage space:

$$\tilde{h}(f) \rightarrow -\frac{\omega^2 \tilde{r}(f)}{k} \tilde{h}(f) . \qquad \text{A-2}$$

Also the voltage power spectrum, $S_V(f)$, rather that the displacement power spectrum, $S_x(f)$, was used to build the optimal filters.

A.3 Other Signals

In addition to the gravitational wave signal, other signals can give useful information about the overall performance of the interferometer. During the data run described in Chapter 7, the following signals were also monitored: signals 1, 2, and 3 of Figure A-1; a version of signal 4 which was filtered to emphasize frequencies below 5 Hz, reflecting the low frequency motions of south masses rather than the differential motion of the two arms in the gravitational wave band; the laser power stabilization feedback signal from within the laser stabilization system; a seismometer signal; a microphone signal measuring background acoustic noise in the lab; a magnetometer signal; and signals measuring the efficiency with which the laser light was coupled into the optical cavities (the visibility of the cavities).

Bibliography

[1] A. Abramovici, W.E. Althouse, R.W.P. Drever, Y. Gursel, S. Kawamura, F.J. Raab, D. Shoemaker, L. Sievers, R.E. Spero, K.S. Thorne, R.E. Vogt, R. Weiss, S.E. Whitcomb, and M.E. Zucker, Science 256 (1992) 325.

[2] C. Bradaschia, R. Del Fabbro, A. Di Virgilio, A. Giazotto, H. Kauzky, V. Montelatici, D. Passuello, A. Brillet, O. Cregut, P. Hello, C.N. Man, P.T. Manh, A. Marroud, D. Shoemaker, J.Y. Vinet, F. Barone, L. di Fiore, L. Milano, G. Russo, J.M. Aguirregabiria, H. Bel, J.P. Duruisseau, G. le Denmat, Ph. Tourrenc, M. Capozzi, M. Longo, M. Lops, I. Pinto, G. Rotoli, T. Damour, S. Bonazzolla, J.A. Marck, Y. Gourghoulon, L.E. Holloway, F. Fuligni, V. Iafolla, and G. Natale, Nucl. Instrum. Methods A 289 (1990) 518.

[3] J. Weber, Phys. Rev. 117 (1960) 306.

[4] Peter F. Michelson, John C. Price, and Robert C. Taber, Science 237 (1987) 150.

[5] W.O. Hamliton, in *Proccedings of the Sixth Marcel Grossman Meeting on General Relativity*, ed. H. Sato (World Scientific, Singapore, 1993).

[6] K. Danzmann, private communication (1994).

[7] J. Weber, Phys. Rev. Lett. 17 (1966) 1224.

[8] .P.R. Saulson, Phys. Rev. D 42 (1990) 528

[9] J. Kovalik and P.R. Saulson, Rev. Sci. Instrum. 65 (1994) 182.

[10] J.E. Logan, N.A. Robertson, J. Hough, and P.J. Veitch, Phys. Lett. A 161 (1991) 101.

[11] A. Gillespie and F. Raab, Phys. Lett. A 178 (1993) 357.

[12] J.R. Hutchinson, J. Appl. Mech. 47 (1980) 901.

[13] H.B. Callen and T.A. Welton, Phys. Rev. 83 (1951) 34.

[14] A.S. Nowick and B.S. Berry, *Anelastic Relaxation in Crystalline Solids* (Academic Press, New York, 1972).

[15] P.R. Saulson, R.T. Stebbins, F.D. Dumont, and S.E. Mock, Rev. Sci. Instum. 65 (1994) 182.

[16] .M. Stephens, P. Saulson, and J. Kovalik, Rev. Sci. Instrum. 62 (1991) 924.

[17] C. Zener, Phys. Rev. 52 (1937) 230; 53 (1938) 90.

[18] A.L. Kimball and D.E. Lovell, Phys. Rev. 30 (1927) 948.

[19] W. Martin, *Experiments and Techniques for the Detection of Gravitational and Pulsed Electromagnetic Radiation from Astrophysical Sources,* Ph.D. Thesis, University of Glasgow, 1978.

[20] G. Gonzalez and P.R. Saulson, J. Acoust. Soc. Am. 96 (1994) 207.

[21] J.E. Logan, J. Hough, and N.A. Robertson, Phys. Lett. A 183 (1993) 145.

[22] N. Mio, Japan J. Appl. Phys. 31 (1992) 1243.

[23] E. Washburn, ed., *International Critical Tables of Numerical Data, Physics, Chemistry, and Technology* (McGraw-Hill, New York, 1929).

[24] T.J. Quinn, Anelasticity of Cu-Be as a function of stress up to the yield stress, presentation to the First international workshop on thermal noise in laser interferometer gravitational-wave detectors, Caltech (1994).

[25] Y. Huang, Tests of tension dependence of violin mode quality factors, presentation to the First international workshop on thermal noise in laser interferometer gravitational-wave detectors, Caltech (1994).

[26] D.G. Blair, L. Ju, and M. Notcutt, Rev. Sci. Instrum. 64 (1993) 1899.

[27] R. Flaminio, Thermal noise due to vertical-horizontal coupling in the suspension,

presentation to the First international workshop on thermal noise in laser interferometer gravitational-wave detectors, Caltech (1994).

[28] H. Kogelnik and T. Li, Appl. Opt. 5 (1966) 1550.

[29] Similar considerations are used in choosing the mirror curvature as were used in: A. Rüdiger, R. Schilling, L. Schnupp, W. Wrinkler, H. Billing, and K. Maischberger, Opt. Acta 28 (1980) 641.

[30] G.W. McMahon, J. Acoust. Soc. Am. 36 (1964) 85.

[31] V.B. Braginsky, V.P. Mitrofanov, and O.A. Okhrimento, JETP Lett. 55 (1992) 432, and references therein.

[32] S. Kawamura and M.E. Zucker, Appl. Opt. 83 (1994) 3912.

[33] J.E. Logan, N.A. Robertson, and J. Hough, Phys. Lett. A 170 (1992) 352.

[34] This value is corrected from the stated value in [8] which was in error. P. Saulson, private communication (1995).

[35] V.B. Braginsky, V.P. Mitrofanov, and S.P. Vyatchanin, Rev. Sci. Instrum. 65 (1994) 3771.

[36] A. Cadez and A. Abramovici, J. Phys E: Sci. Instrum. 31 (1988) 453.

[37] Joseph H. Taylor, Rev. of Mod. Phys. 66 (1994) 711.

[38] R. Narayan, T. Piran, A. Shemi, Ap.J. 379 (1991) L17.

[39] E.S. Phinney, Ap. J. 380 (1991) L17.

[40] K.S. Thorne, in *300 Years of Gravitation*, ed. S.W. Hawking and W. Israel (Cambridge University Press, Cambridge, 1987).

[41] C. Cutler, T.A. Apostolatos, L. Bildsten, L.S. Finn, E. Flanagan, D. Kennefick, D. Markovic, A. Ori, E. Poisson, G.J. Sussman, and K.S. Thorne, Phys Rev Lett 70