

EXPERIMENTS WITH A LASER
INTERFEROMETRIC GRAVITATIONAL WAVE
ANTENNA

Thesis by
Michael Edward Zucker

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Pasadena, California
1989

(Submitted July 1, 1988)

Acknowledgements

Sincere thanks are due the many people who have assisted me through these years. Among them, my advisor, Ron Drever, as well as Stan Whitcomb, Kip Thorne and most recently Robbie Vogt have helped to shape my research and have supported my efforts. The laughter, sound advice and friendship shared with Mark Hereld, Dana Anderson, Carl Caves, Steve Clark, Sheri Smith, Harry Ward and Bob Spero have rendered this difficult and demanding work enjoyable and even (at rare moments) relaxing. I should also acknowledge the utter indifference of our four cats, past and present. Finally, I am most grateful to my dear wife Gail, for her steadfast support and deepest affection.

Abstract

Sources and effects of astrophysical gravitational radiation are explained briefly to motivate discussion of the Caltech 40 meter antenna, which employs laser interferometry to monitor proper distances between inertial test masses. Practical considerations in construction of the apparatus are described. Redesign of test mass systems has resulted in a reduction of noise from internal mass vibrations by up to two orders of magnitude at some frequencies. A laser frequency stabilization system was developed which corrects the frequency of an argon ion laser to a residual fluctuation level bounded by the spectral density $\sqrt{S_\nu(f)} \leq 60 \mu\text{Hz}/\sqrt{\text{Hz}}$, at fluctuation frequencies near 1.2 kHz. These and other improvements have contributed to reducing the spectral density of equivalent gravitational wave strain noise to $\sqrt{S_h(f)} \approx 10^{-19}/\sqrt{\text{Hz}}$ at these frequencies.

Finally, observations made with the antenna in February and March of 1987 are described. KiloHertz-band gravitational waves produced by the remnant of the recent supernova are shown to be theoretically unlikely at the strength required for confident detection in this antenna (then operating at poorer sensitivity than that quoted above). A search for periodic waves in the recorded data, comprising Fourier analysis of four 105-second samples of the antenna strain signal, was used to place new upper limits on periodic gravitational radiation at frequencies between 305 Hz and 5 kHz. In particular, continuous waves of any polarization are ruled out above strain amplitudes of 1.2×10^{-18} R.M.S. for waves emanating from the direction of the supernova, and 6.2×10^{-19} R.M.S. for

waves emanating from the galactic center, between 1.5 and 4 kilohertz. Between 305 Hz and 5 kHz no strains greater than 1.2×10^{-17} R.M.S. were detected from either direction. Limitations of the analysis and potential improvements are discussed, as are prospects for future searches.

Preface

The Caltech 40 meter antenna is in fact a prototype for much larger antennas being designed for the planned Laser Interferometer Gravitational Observatory (LIGO), scheduled to begin observations early in the next decade. In this capacity our antenna is most often used as a testbed for the challenging new technology needed to achieve the sensitivity goals of the observatory project; as we shall see, its sensitivity currently falls short of the gravitational waves strengths predicted for most plausible astrophysical mechanisms. Its own role as an observing instrument is thus mostly confined to placing upper limits on anomalous radiation.

A further consequence of the prototypical nature of our apparatus is that it is constantly in a state of flux. In trying to describe the technical work in a logical way I must frequently jump between different times in the development of the antenna, and I often assess the state of knowledge on certain issues with considerable reference to historical context. At the very least, the numerous changes to the apparatus which are *not* described here (many people participate actively in our research) must often be mentioned insofar as they affect the contributions under consideration. While I believe this style presents the findings most realistically, readers of early drafts have pointed out that the text reads like a novel with too many flashbacks. I hope successive revisions have made it less objectionable.

Contents

Acknowledgements	ii
Abstract	iii
Preface	v
1 Ideas and Objectives	1
1.1 Motivations	1
1.2 Generation of Gravitational Waves	4
1.3 Effects of Plane Gravitational Waves	7
1.4 Notation; Power Spectral Densities	13
1.5 Interferometric Antennas	15
1.5.1 Practical Considerations	15
1.5.2 “Fundamental” Sensitivity Limits	17
1.5.3 Fabry-Perot Interferometers	20
1.5.4 Less Fundamental Sensitivity Limits	24
2 The Test Masses	27
2.1 Seismic Isolation	27
2.2 Acoustic Isolation	30
2.3 Suspension and Control	30
2.4 Shark Detectors	34

2.5	Thermal Noise	38
2.6	Mass Engineering	40
3	The Lasers	48
3.1	Laser Intensity Noise	49
3.1.1	Electrooptic Feedback	51
3.1.2	Acoustooptic Feedback	58
3.2	Laser Frequency Noise	65
3.2.1	Frequency Noise in Argon Lasers	67
3.2.2	The Servo	69
3.2.3	The Electronics	75
3.2.4	Performance	81
3.3	Locking the Secondary Cavity	86
4	The Observations	91
4.1	State of the Interferometer	91
4.2	Data Runs	92
4.3	Potentialities and Limitations of the Data	94
4.4	Periodic Gravitational Waves	95
4.5	Data Analysis	102
4.5.1	General Method	102
4.5.2	Transformation	103
4.5.3	Windowing	104
4.5.4	Finding the Peaks	107
4.5.5	Statistical Expectations	109
4.6	Results	114
4.6.1	First Impressions	114
4.6.2	Microphonics	119

4.6.3	Aliased Line Multiples	121
4.6.4	Remaining Peaks	122
4.7	Conclusions	128
4.7.1	Limits on Strain in the Antenna	128
4.7.2	Interpretation; Strains in Space	129
4.7.3	Comments	136
A	Test Mass Testing	139
B	Principle of the Optical Phase Detection	144
C	Burst and Coalescing Binary Sources	149
C.1	Short Bursts of Gravitational Radiation	149
C.2	Coalescing Compact Binaries	152
D	Data Acquisition	156
D.1	Signal Preparation	156
D.2	Timing	159
D.3	Calibration	160

List of Tables

1	Analyzed Tape Segments	115
2	Effect of Local Interference Vetoes on Peak Lists	123
3	Coordinates of Antenna and Two Possible Sources	134
4	Antenna Beam Factors for the LMC and our Galactic Center . .	135
5	Power Line Frequencies During the Four Analyzed Segments . .	138

List of Figures

1	Three Test Particles and a Gravitational Wave	8
2	Wave, Antenna and Planet Coordinates	10
3	Coordinates on the Celestial Sphere	11
4	Fabry-Perot Interferometer	21
5	Phasor Diagram of Reflected Optical Field	23
6	Strain Spectrum of the Caltech Antenna: October 1987	26
7	A Test Mass and Its Control Systems	33
8	Schematic of Shark Detector Circuitry	36
9	A Pre-1984 Test Mass	41
10	Interferometer Spectrum with Original Masses	42
11	Strain Spectrum Before and After Rebuilding the Test Masses	45
12	Setup of Intensity Servo	53
13	Schematic of Intensity Servo Electronics	54
14	Bode Plot for the Electrooptic Intensity Servo	55
15	Improvement in Intensity Servo Performance with Lens System	59
16	Schematic of the Acoustooptic Intensity Servo	62
17	Bode Plot of the Acoustooptic Intensity Servo	63
18	Laser Frequency Fluctuations vs. Cooling Water Flow	68
19	Laser Frequency Fluctuations with a Rebuilt Resonator	70
20	Schematic Diagram of the Laser Stabilization System	73

21	Topology of the Servo Amplifier	77
22	Laser Stabilization Servo Bode Plot	78
23	RF Photodetector Circuitry	79
24	Laser Stabilization Servo Amplifier Circuit	80
25	Servo Layout with a Mode Cleaner Inside the Loop	84
26	Performance of the Stabilized Laser	85
27	Secondary Cavity Servo Loop Diagram	87
28	Schematic of the Antenna System	89
29	Interferometer Noise Spectrum; February 1987	90
30	Kilohertz-Band Periodic Sources and Experimental Upper Limits	100
31	Peak Height Histogram for Random Noise	111
32	Peak Height Histogram From Some Real Data	117
33	Coarse Power Spectrum of Segment 45(a)	118
34	Suspected Wire Resonances	125
35	Number of Peaks vs. Location in Spectrum	126
36	Number of Peaks vs. Location in Spectrum	127
37	Upper Limit for Periodic Strain in Antenna: Tape 45(a)	130
38	Upper Limit for Periodic Strain in Antenna: Tape 45(b)	131
39	Upper Limit for Periodic Strain in Antenna: Tape 49	132
40	Upper Limit for Periodic Strain in Antenna: Tape 56	133
41	Composite Limits on Continuous Periodic Gravitational Waves .	137
42	Resonances of a Test Mass	141
43	Some Eigenmodes of an Old Mass	142
44	Phasor Diagram of the Phase Modulation	146
45	A Cavity Resonance	148

46 Published Burst Upper Limits 153

Chapter 1

Ideas and Objectives

1.1 Motivations

Building instrumentation to detect the weak influence of gravitational waves bathing the Earth has proven to be a difficult and challenging task. Since the first antennas were operated in the 1960's, persistent worldwide efforts by many researchers have improved the sensitivity of successive generations by several orders of magnitude. Still more sensitivity will be necessary before regular observations of gravitational waves can begin to yield astrophysical insight. These intense programs of technological development and observation are motivated by the exceptional properties predicted for gravitational waves and in the unique information they can bring us about their sources.

Our perception of the universe has been revolutionized by astronomical observations in unexplored electromagnetic bands. The processes which typically produce radio waves, x-rays, infrared, ultraviolet and visible light are often so physically different that entire classes of objects and events went totally unobserved in some wavelengths while dominating the sky in others. Nevertheless, these radiations are nearly always generated in regions of low matter density and weak gravity; they share the common weakness that they are obscured and

trapped by the very matter and gravity we would like to study [1].

As the observation of neutrinos from the stellar collapse associated with Supernova 1987a [2,3] demonstrates, non-electromagnetic probes can penetrate to the very core of the most dynamic and dense regions. The subnuclear processes at the heart of the “superficial” optical display are thus revealed. Gravitational waves convey a further complementary picture, one emphasizing the physics of mass and of gravity itself. They penetrate, as do neutrinos; unlike neutrinos, however, they are not “triggered” by microcosmic reactions, but are evolved by the dynamical object as a whole, defying localization.

Indirect but compelling verification that gravitational radiation exists, and indeed that much of gravitation is well described by General Relativity, has arisen from observations of the radiation–reaction drag on the binary pulsar PSR 1913+16 [4]. Much remains to be learned about the waves themselves; for example, whether they travel at precisely the same speed as light could be checked if an optical display, such as a supernova, were identified with a burst of gravitational waves. Waves from a coalescence of black holes (or stellar collapse to a black hole) would provide incontrovertible evidence for the existence of these objects. This observation could also provide the most stringent possible test of General Relativity in the extreme limits of strong gravity and rapid motion. Neutron stars are likewise expected to reveal much of their structure, dynamics and evolution by the gravitational waves they emit at their birth and in their interactions with neighboring stars. The independent determination of the Hubble constant from observations of extragalactic coalescences of compact binaries is a further tantalizing benefit. The highly predictable gravitational waveforms these stellar systems emit as they inspiral toward collision are an ideal “standard candle,” and the identification of such an event’s host galaxy, or statistical conclusions from a number of detections, could allow determination of

the redshift/distance relation with unprecedented accuracy [6]. Numerous compact binary coalescences are predicted within the range of planned gravitational wave observatories [7].

It is important, however, to reemphasize the impact of serendipitous discoveries uncovered by new modes of observation. We may quite reasonably expect unforeseen objects and phenomena to appear with the opening of gravitational wave observatories, just as radio telescopes first showed us quasars, galactic jets, and pulsars.

1.2 Generation of Gravitational Waves

Gravitational radiation is predicted to arise from acceleration of massive systems, much as acceleration of electric charges generates electromagnetic radiation. An important difference, however, is that there exists only one kind of “charge” (i.e., mass) for gravitational interactions. As a result, the equivalent of electric dipole radiation is prohibited in gravitation by the conservation of linear momentum; furthermore, the equivalent of magnetic dipole radiation is prohibited by conservation of angular momentum¹.

However, nonvanishing quadrupole radiation is possible in gravity. The power radiated by an electric quadrupole is proportional to the square of the third time derivative of the electric quadrupole moment; similarly, an massive system will radiate gravitational wave power proportional to the square of the third time derivative of its mass quadrupole moment [8, Chapter 36]. Heuristically, one can envision a pair of masses in orbit (a changing mass quadrupole). To first approximation, the differential contributions of the two masses to the curvature of spacetime at a great distance will cancel one another. But if allowance is made for the fact that one mass will typically be farther away from a distant observer than the other, a change in their orbital phase will appear in the contribution made by the nearer mass an instant before the change in the contribution from the farther mass arrives, since news of the altered position must travel from each mass at a finite speed ($= c$ in General Relativity and most other gravitation theories).

The distant observer can carefully monitor changes in the curvature of spacetime by planting inertial test particles at various positions and monitoring the

¹A useful analogy: in a radio broadcast whip antenna opposite charges are given opposite acceleration, and in a loop antenna opposite charges are given opposite torques. Gravitational equivalents cannot exist because the total linear and angular momenta of the “antenna” are constants of the motion.

proper lengths of geodesics joining them using clocks and photons. Geodesics parallel to the plane of the oscillating source’s motion and those in a perpendicular plane will expand and contract in antiphase; geodesics oriented at 45° to these planes will, to first order, remain unaffected. This defines two orthogonal (in the physical sense) states of plane polarization a wave may adopt, referred to as “plus” (+) and “cross” (\times) states. In general the waves will contain an admixture of polarization states, often (as in our orbiting–mass example) with fixed relative phases and amplitudes, leading to a natural description in terms of circular or elliptical polarization states somewhat analogous to their electromagnetic counterparts. The single antenna to be described here responds to the projection of an arbitrary incident gravitational wave onto a single linear polarization state corresponding to its own orientation, so we shall from now on consider only linear polarization; the combination of several such antennas at different orientations can in principle give a complete picture of both polarization components of a wave [9].

The magnitude of a linearly polarized plane gravitational wave is quantified by its “strain amplitudes” h_+ and h_\times , defined as the fractional difference in length between orthogonal, initially equal geodesics in the plane perpendicular to the waves’ propagation direction, for each of the two polarization states. Technically, these amplitudes are components of the difference between the metric tensor in the absence of the wave and the time–varying metric perturbations attributable to the wave, i.e. $h_{jk}(t) = g_{jk}(t) - g_{jk}(0)$ if it is assumed that no waves were present at $t = 0$. In order of magnitude, the strain components at the position of the detector (distance r from the source) will be

$$|h_{jk}| \approx \frac{2G}{c^4} \frac{\ddot{I}_{jk}}{r} \sim \frac{G}{c^4} \frac{E_{ns}}{r}, \quad (1)$$

where I_{jk} is the mass quadrupole moment tensor of the source and E_{ns} is, roughly,

the “nonspherical component” of the source’s internal kinetic energy [1]. This result yields disappointingly small amplitudes, almost certainly undetectable, for constructible laboratory sources of gravitational waves. We are thus unable to test the properties of the waves in a controlled Hertz–type experiment on Earth; we must look to astrophysical sources to find enough quadrupole moment (i.e. enough aspherical mass) and rapid enough acceleration for visible effects. As an example, a nonaxisymmetric neutron star with with an ellipticity ϵ of one percent (\sim the fractional disparity between equatorial components of its quadrupole tensor) rotating with a period of 2 milliseconds could produce a strain $|h|$ of order one part in 10^{21} at a distance of 10 kiloparsecs. More detailed estimates of the waves emitted by various sources will be presented in Chapter 4 and Appendix C; the latest theoretical calculations are summarized in [1].

It is conventional to classify sources according to the time evolution of their waveforms $h(t)$. “Bursts,” as might be produced by aspherical stellar collapse and collisions or perturbations of black holes, last only a few cycles and then die off. Periodic sources, like the rotating neutron star mentioned above, are expected to emit continuous radiation at a fixed or nearly fixed frequency (or set of frequencies), maintaining coherence over many cycles. Somewhere in between lie the signals expected from compact binary star systems which, through gravitational radiation, lose orbital energy and spiral together more and more rapidly, emitting a “chirp” whose frequency rises with time until the stars coalesce in a final burst. Finally, relict waves from the formation of the universe, or the “confusion limit” arising from the random combination of numerous overlapping signals, can be classified as stochastic waves. The various classes typically require different strategies for optimal detection in the presence of noise, but the antennas to be described are in principle sensitive to all of them.

Earthbound antennas are plagued by seismic interference, which will probably prevent sensitive observations of waves having frequency components below 10 Hz (the seismic limit is more like 100 Hz in current prototypes, but we anticipate much improvement). Astrophysical considerations imply that waves at frequencies above a few tens of kilohertz should be relatively weak; as a result, we will henceforth focus on the problems associated with detecting signals in the audio frequency range.

1.3 Effects of Plane Gravitational Waves

Gravitational waves impinging on systems of free inertial test masses are easily characterized in a comoving coordinate system where inertial test particles maintain fixed coordinate positions. If external forces are absent (or fluctuate with vastly different characteristic timescales than the gravitational waves under consideration), this system has the convenience that the coordinate positions of test particles remain constant.

To define such a system, imagine a set of three free-falling test particles arranged in the xy plane at positions $(0, 0, 0)$, $(0, -l, 0)$, and $(l, 0, 0)$. To detect any change in the metric tensor, as from a gravitational wave, one can measure the integral of

$$0 = (ds)^2 = g_{\alpha\beta} dx^\alpha dx^\beta \quad (2)$$

along two orthogonal geodesics by clocking the proper time taken by light pulses emitted simultaneously at the $(0, 0, 0)$ particle to return after reflection from the $(0, -l, 0)$ and $(l, 0, 0)$ particles (Figure 1).

Specifically, consider a metric tensor containing a plane gravitational wave superimposed on a flat background

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu},$$

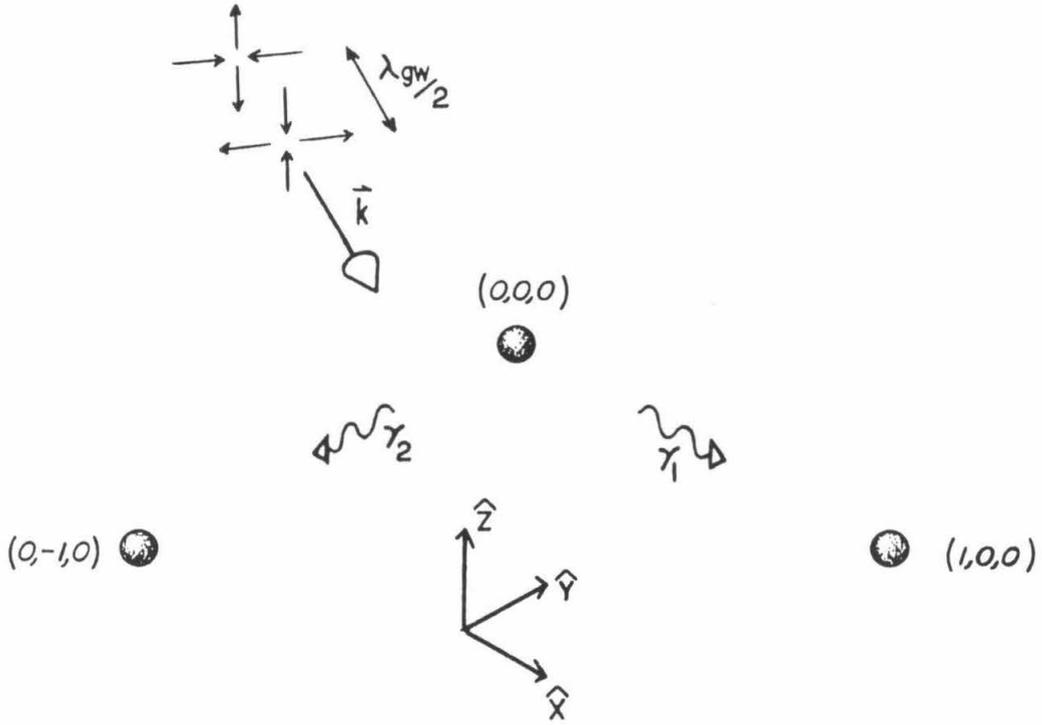


Figure 1: Three inertial test particles falling freely through space encounter a plane gravitational wave with wavevector \vec{k} and wavelength λ_{gw} . Photons γ_1 and γ_2 are sent out simultaneously by the vertex particle to the outlying particles, and are reflected back. The difference in their round-trip times, measured by a clock carried on the vertex particle, is proportional to the gravitational wave amplitude h ; the experiment is repeated continuously to determine the waveform $h(t)$. Two orthogonal geodesics are used to cancel variations in the clock rate over the mean roundtrip interval $\tau_{rt} = 2l/c$; the clock only needs to keep good time for the short difference $\delta\tau_{rt} = h\tau_{rt}$. If a perfect clock were available, one baseline would suffice.

where η is the Minkowski metric. The metric perturbation $h_{\mu\nu}$ can be rendered transverse (i.e., only spatial components are nonzero, and these are divergence-free) and traceless by a suitable gauge transformation [8, 946–950]. Then the relevant components of Equation 2 determining the proper lengths of the geodesics traversed by our light pulses will be h_{xx} and h_{yy} , giving

$$(dt)_{(x)}^2 = (1 + h_{xx})(dx)^2 \quad (3)$$

and

$$(dt)_{(y)}^2 = (1 + h_{yy})(dy)^2. \quad (4)$$

respectively. A plane gravitational wave impinging on the system with wavevector \vec{k} , strain amplitude $h_0(t)$ (in the polarization state for which it is a maximum) and polarization angle ψ with respect to the kz plane will have

$$\begin{aligned} h_{xx} &= h_0(t)[\cos 2\psi(\cos^2 \phi - \sin^2 \phi \cos^2 \theta) - \sin 2\psi \sin^2 \phi \cos \theta] \\ &\equiv h_0(t)\alpha_{xx}(\theta, \phi, \psi) \end{aligned} \quad (5)$$

and

$$\begin{aligned} h_{yy} &= h_0(t)[\cos 2\psi(\sin^2 \phi - \cos^2 \phi \cos^2 \theta) + \sin 2\psi \sin^2 \phi \cos \theta] \\ &\equiv h_0(t)\alpha_{yy}(\theta, \phi, \psi). \end{aligned} \quad (6)$$

Here \vec{k} can be written as

$$\vec{k} = \frac{2\pi}{\lambda_{GW}}(\sin \theta \sin \phi \hat{x} - \sin \theta \cos \phi \hat{y} + \cos \theta \hat{z}) \quad (7)$$

with λ_{GW} the wavelength of the gravitational radiation and θ and ϕ the co-altitude and azimuth with respect to the \hat{z} and \hat{x} axes (Figures 2 and 3) [10,11].

The path integral of the geodesic equation (2) is then evaluated for the two light pulses sent out; the elapsed proper times for the two orthogonal roundtrips are subtracted, cancelling out to lowest order possible errors in the rate of the

clock. This gives the round trip time delay $\delta\tau_{rt}(t)$ induced by the wave as a function of proper time. In the usual limit $|h_0| \ll 1$, and for $l \ll \lambda_{GW}$, this time delay is given by

$$\begin{aligned} \delta\tau_{rt}(t) &\simeq h_0(t) \frac{2l}{c} [\alpha_{xx}(\theta, \phi, \psi) - \alpha_{yy}(\theta, \phi, \psi)] \\ &\equiv \frac{2l}{c} h(t), \end{aligned} \tag{8}$$

defining $h(t)$ as the projection of the astrophysical strain registered in the antenna [10]. The “antenna beam” factor $\alpha_{xx}(\theta, \phi, \psi) - \alpha_{yy}(\theta, \phi, \psi)$ is more conveniently evaluated if the transformation $\psi \rightarrow \psi'$ is made redefining the polarization state with respect to the fixed celestial meridian of the source, rather than the rotating detector axes (Figure 3). This convention will be adopted subsequently; a wave having “plus” (+) linear polarization state, corresponding to $\psi' = 0, \pi/2, \pi, \dots$, is hence defined as one which alters geodesics laid out along North–South and East–West celestial coordinate directions. The “cross” (\times) state, of course, is rotated by $\pi/4$ with respect to this choice.

The extensions of Equation 8 to the cases where the length l is comparable to or greater than the gravitational wavelength λ_{GW} , and/or the light traverses each path many times, have been derived in [10,12,13]. The magnitude of the roundtrip time difference remains proportional to the astrophysical strain h_0 , but has a different characteristic form in each regime.

1.4 Notation; Power Spectral Densities

Most of the signals and measurements to be described are presented in the frequency domain. This convention arises naturally in the contexts of servo system design, mechanical resonances, and periodic gravitational wave searches to be discussed. Noise or other predominantly stochastic processes will be characterized by their power spectral densities, or PSD’s, denoted by $S_{\cdot}(f)$ where the blank is filled in by the name of the time-dependent variable being described. For example, the power spectral density of the linear displacement $x(t)$ of some moving object is defined as

$$S_x(f) = \lim_{T \rightarrow \infty} \frac{1}{T} \left| \int_{-T}^T x(t) e^{-2\pi i f t} dt \right|^2 \quad (9)$$

where f is the cyclic frequency in Hertz, and has units (meters)²/Hz [14]. For a random process $x(t)$ this should strictly be replaced by its expectation value over an ensemble of identical random systems, but for convenience we shall keep the same symbol. Of course, real measurements can only involve estimates of this quantity based on finite observation times T , and the ensemble average is always replaced by an average over samples of the same process taken successively (equivalent if the process is approximately stationary). The frequency f is usually in the audio range, corresponding to the frequencies of gravitational wave signals we expect to look for; we’ll call the optical frequency of the laser light ν for distinction. The square root of the power spectral density at a frequency f is often quoted for comparison with some expected signal amplitude. This “root PSD,” or RPSD, should not be confused with the Fourier transform of a deterministic signal, however. We denote the Fourier transform of a waveform $y(t)$ by $\tilde{y}(f)$, which (again for the example that y represents a linear displacement) has units meters/Hz. The RPSD has the property that a signal of R.M.S. amplitude a , which lasts for a duration τ , will be detectable in noise characterized by a RPSD

\sqrt{S} if $a > \sqrt{S\tau}$. Thus for a continuous signal (e.g. $h(t) = h(0) \cos(2\pi f_{GW}t)$) buried in noise having a constant mean RPSD, the signal-to-noise ratio will improve as the square root of the integration time of the measurement.

1.5 Interferometric Antennas

The first instruments built to detect the weak effects of gravitational waves did not spring directly from the above description involving free masses. Instead, bar antennas (as they are widely known) are arranged to register very small changes in the state of a harmonic oscillator, usually a large crystal or bar of metal. A gravitational wave passing through such an oscillator would deposit some energy in it (or remove some, depending on the bar’s initial state and the phase of the wave). For more discussion of bar antennas see [15,16,17]. Recent variants of these instruments have achieved sensitivities approaching $h \approx 10^{-18}$ for bursts of millisecond duration.

The use of inertial masses monitored by light beams, while perhaps as old a concept as the bar (see [1] and references therein for historical context), was not tried experimentally until relatively late. I will confine this discussion to certain practical realizations of this method. The first free-mass interferometric instrument was constructed and operated by Forward at Hughes Corporation in the early 1970’s and achieved a strain RPSD of $\sqrt{S_h(f)} \simeq 3 \times 10^{-16}/\sqrt{\text{Hz}}$ between 1 and 3 kilohertz [11]. Since then more sophisticated interferometers have been built at the University of Glasgow in Scotland, Max Planck Institute in Germany, Weizmann Institute in Israel, and at both MIT and Caltech in the U.S.

1.5.1 Practical Considerations

The test masses in the simple description of Section 1.3 were assumed to be in free-fall, uninfluenced by any external forces. A realistic experiment done on the Earth’s surface can only approximate these conditions. One simplification is that the masses need only be untethered in the plane of the monitored distances.

Another is that, provided the instrument responds with sufficient linearity to external forces, those which fluctuate at frequencies well away from expected signal frequencies can be filtered out or ignored. These considerations suggest hanging the test masses on long wires, leaving them free in the horizontal plane at frequencies far above the resonant frequency of the pendulum mode. Current gravitational wave interferometers have their test masses suspended on wires about 0.3 to 1.5 meters in length; the masses are thus ‘inertial’ over timescales significantly shorter than 1 second. Fluctuating forces from seismic and artificial ground motion will tend to mask the tiny displacements sought; the pendulum suspension serves to isolate the mass from these influences. Some additional antiseismic and acoustic isolation between the outside world and the pendulum suspension point is generally needed as well. Since hanging the mass from wires will typically allow it to rotate freely, some angular control must be applied to keep the light reflector aligned, or a special angle-insensitive optical arrangement must be used.

Having ensured the freedom of the test masses from external influences, we are left with the problem of accurately monitoring the relative delay between the geodesics joining masses along the two arms. In current instruments continuous beams of coherent light are used, rather than discrete pulses. The imprint of proper time is derived from the phase of each returning light wave, determined by interfering them together or with the common source.

1.5.2 “Fundamental” Sensitivity Limits

The Quantum Limit

Applying a “Heisenberg’s Microscope”-type argument to the measurement of the position of a free mass yields an uncertainty principle limit

$$\Delta x \gtrsim \sqrt{\frac{2\hbar\tau}{m}} \quad (10)$$

where m is the mass and we have assumed we are looking for a transient signal of duration τ (\approx the inverse bandwidth of the measurement) [18]. Unless some scheme for circumventing this limit for free masses is discovered [19], the smallest detectable motion is limited by the equivalent displacement spectral density

$$\sqrt{S_x(f)} \geq 2 \times 10^{-21} \left(\frac{1 \text{ kg}}{m}\right)^{\frac{1}{2}} \left(\frac{1 \text{ kHz}}{f}\right) \frac{\text{m}}{\sqrt{\text{Hz}}} \quad (11)$$

or, for a detector of arm length l , a strain spectral density

$$\sqrt{S_h(f)} \geq 1 \times 10^{-22} \left(\frac{1 \text{ kg}}{m}\right)^{\frac{1}{2}} \left(\frac{40 \text{ m}}{l}\right) \left(\frac{1 \text{ kHz}}{f}\right) \text{Hz}^{-\frac{1}{2}} \quad (12)$$

where we have used the rigorous calculation of quantum limit spectral density (assuming an optimum measurement strategy) provided by Caves [21].

This limit is obtained by increasing the light power until the effect of fluctuating radiation pressure on the masses (due to quantum fluctuation of the number of photons hitting the mirror) becomes equal to the uncertainty in the measured phase (due to quantum fluctuation of the number of photons detected). To actually encounter this quantum limit in any current antenna would require much more light power than suitable present-day lasers can provide (of order one kilowatt for the Caltech instrument).

Shot Noise

With light sources available today, a more pressing concern is the statistical fluctuation in the finite number of photons that can be mustered for the phase measurement. This fluctuation is commonly referred to as the “shot noise” in the light. For most ordinary light sources individual photons are independent, so their arrivals at a detector are Poisson–distributed in time. This leads to a fluctuation $\Delta N = \sqrt{\bar{N}}$ in the number of photons N arriving in an interval over which \bar{N} are typically expected. If one measures the relative phase between two beams of light, for example by interfering them and estimating the resultant intensity as a fraction of the maximum (in–phase) intensity, it is possible to interpolate to precision $\Delta\phi \gtrsim \pi/\sqrt{\bar{N}}$ by integrating for a period in which \bar{N} photons are typically expected, implying a strain sensitivity limit of order $h \gtrsim \lambda_{opt}/2\pi l\sqrt{\bar{N}}$ where λ_{opt} is the wavelength of the light used.

The Question of Length

Many noise sources which might interfere with the strain measurement contribute a fixed amount of displacement noise to each mass, irrespective of the readout system used. Among these are the “quantum limit” noise mentioned above and the thermal Brownian motion of the test masses. To make these sources contribute as little noise as possible to the strain measurement, it is desirable to separate the masses as much as possible. The largest reasonable length is set by the finite speed of the light beams; when $2l \simeq \lambda_{GW}$ the wave–induced phase shift reverses sign before a given photon has completed its journey, and the effect will tend to average away over successive half cycles². For kilohertz waves there is little practical advantage to making l any more than 150 kilometers.

²Some plausible gravitational waveforms $h(t)$ have “memory”; a nonzero strain remains indefinitely after the wave passes, and the phase shift can continue to accumulate [22].

A detector this physical length would be expensive, so a further compromise is based on evaluating the local noise sources more carefully. Fortunately, from the point of view of minimizing the contribution of shot noise, it is not the physical length but the *optical* length which matters. Several schemes have been proposed to actually fold the optical path into a shorter arrangement³. The gravitational wave—induced phase shift is then accumulated over b successive reflections of the beam from each test mass, yielding the same optical phase shift as a single-bounce detector b times as long. The total time a given photon spends retracing each path before its phase is measured is, of course, subject to the same constraint imposed by the finite period of gravitational waves under consideration. Two major folded optical schemes are under investigation, the optical delay line [24] and the Fabry-Perot cavity [25].

An optical delay line comprises two concave spherical mirrors arranged so that a light beam penetrating one mirror through a hole near the periphery is made to execute a sequence of discrete reflections between the two before reemerging through the same hole [26]. The number of bounces on each mirror is determined by their separation and curvature and the condition that the exit beam fall on the entry hole. This system has several attractive properties, one being that the optical path is relatively insensitive to the angles of the mirror surfaces. Gravitational wave detectors employing this technique have been developed extensively by several groups [27,28,29].

The Fabry-Perot cavity configuration of gravitational wave interferometer has been chosen for study at Caltech. For a review of many alternative optical, mechanical, and other detectors of gravitational radiation, please see [1] and its

³Folding the optical path to make the instrument fit into a practical laboratory was a key innovation of Michelson's interferometers [23]. The techniques described here greatly enhance the possible number of reflections while adding relatively little complication to the mirror structure, another important consideration to be discussed in Chapter 2.

references.

1.5.3 Fabry-Perot Interferometers

Fabry-Perot interferometers have been applied to precise measurements of wavelength and refractive index for many years [30,31]. In essence, the simplest such instrument consists of two plane, parallel, slightly transmissive mirrors. Successive internal reflections of collimated monochromatic light entering through one mirror will superimpose constructively if the mirrors are an integral number of half-wavelengths apart, leading to a buildup of light power inside; a very small change in the separation of the reflectors (or the light wavelength) will scramble the relative phases of successive reflections and the buildup will disappear. To see how the principle can be applied to accumulating the optical phase shift induced by a gravitational wave it may be instructive to follow a wave train through the instrument.

Consider Figure 4. In each arm of a standard Michelson interferometer we have interposed a partially transmitting mirror (M_1 and M_3) of reflectivity R_1 . These inner mirrors may be connected rigidly to the beamsplitter or may be on separate test masses nearby. The mirrors on the distant test masses have been made concave to match the spherical phase fronts of the diffracting laser beams; the systems M_1M_2 and M_3M_4 are thus stable optical resonators [32, Chapter 8]. The phase of every light wave at a given axial position in such a resonator (often called a “cavity”) has the same dependence on cylindrical radius ρ , which is thus ignorable.

The beamsplitter divides the laser power equally between the two arms. In each, a certain fraction E_0 is reflected promptly at the near mirror while the rest proceeds to the far test mass, whereupon it is sent back bearing the imprint of the geodesic proper length in its phase $\delta\phi$. On return a fraction is transmitted

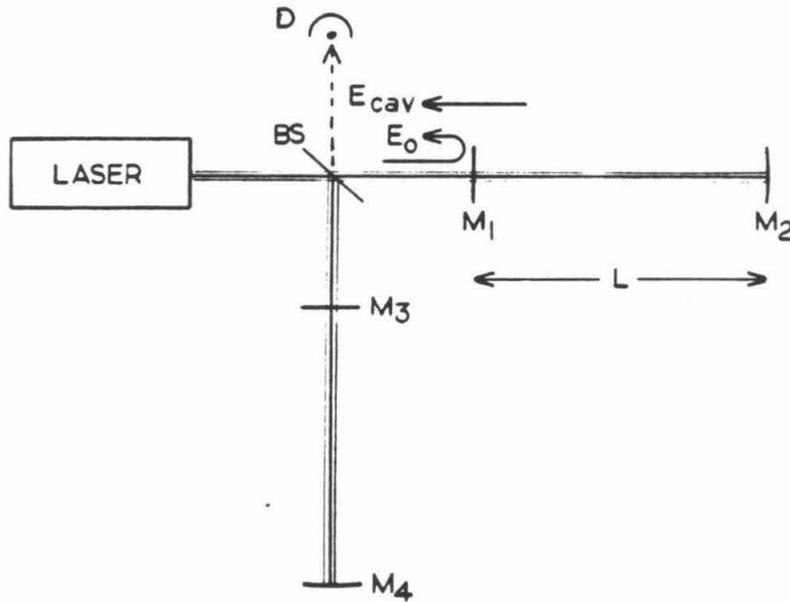


Figure 4: A Fabry-Perot based gravitational wave detector. M_1 and M_3 are partially transmissive, while M_2 and M_4 have the highest reflectivity possible. Both distances M_1M_2 and M_3M_4 are held by servo systems (not shown) to roughly an integral number of half-wavelengths of the laser light. This causes the light to build up in each arm by resonance, and enhances small wave-induced phase shifts in the light leaking back out through the input mirror. This light field E_{cav} interferes with the incident laser field E_0 reflected promptly by the input mirror; the resultants from the two arms may be diverted and examined separately for the two cavities, or (as shown) may be interfered with each other back at the beamsplitter and the final resultant monitored with photodetector D.

back toward the beamsplitter while the rest is reflected to repeat the journey, each successive trip accumulating phase shift $\delta\phi$ and experiencing a small loss to the outside at the input mirror. The experimenter arranges (by methods to be discussed in Section 3.2) to keep $\delta\phi \approx 2n\pi$ on average; electric field components from successive round trips thus add up almost coherently, and the light inside each cavity is built up by resonance to an intensity $I_c \approx I_0 \times (1 - \sqrt{R_1 R_2})^{-1}$ where I_0 is the laser intensity incident on M_1 (Figure 5).

The light leaking backward out of each cavity E_{cav} comprises the sum of electric fields from each of the many round trips, and this resultant field interferes with the promptly reflected input E_0 from the laser; the sum $E_{cav} + E_0$ may itself be interfered at the beamsplitter with the result of the symmetric process returning from the other arm to find the phase difference. In current practice the returning light from the two arms is usually diverted by directional isolators (not shown) and is not recombined. The relative phase between E_{cav} and E_0 is measured separately for each cavity and the results are effectively subtracted to give the gravitational wave strain.

The induced optical phase difference between the outputs of the two cavities for a sinusoidal gravitational wave strain of amplitude h and frequency f_{gw} is given by

$$\Delta\phi(f_{gw}) = \frac{4\pi^2 c}{\lambda} \left(\frac{1}{4\pi\tau_e} - if_{gw} \right)^{-1} h(f_{gw}) \quad (13)$$

for $f_{gw} \ll c/l$ [10]. Here $\tau_e \equiv l/[c(1 - \sqrt{R_1 R_2})]$ is the cavity storage time; a given photon bouncing between the mirrors has probability $1/e$ of remaining after this amount of time. The number of bounces a typical photon survives without escaping is therefore $b \sim c\tau_e/l$, making the connection with discrete-bounce systems clear. Equation 13 shows the fall in optical phase response for a given amplitude h as the gravitational wave frequency exceeds the characteristic frequency of the cavity $f_c \equiv 1/4\pi\tau_e$. In the Caltech prototype this frequency has

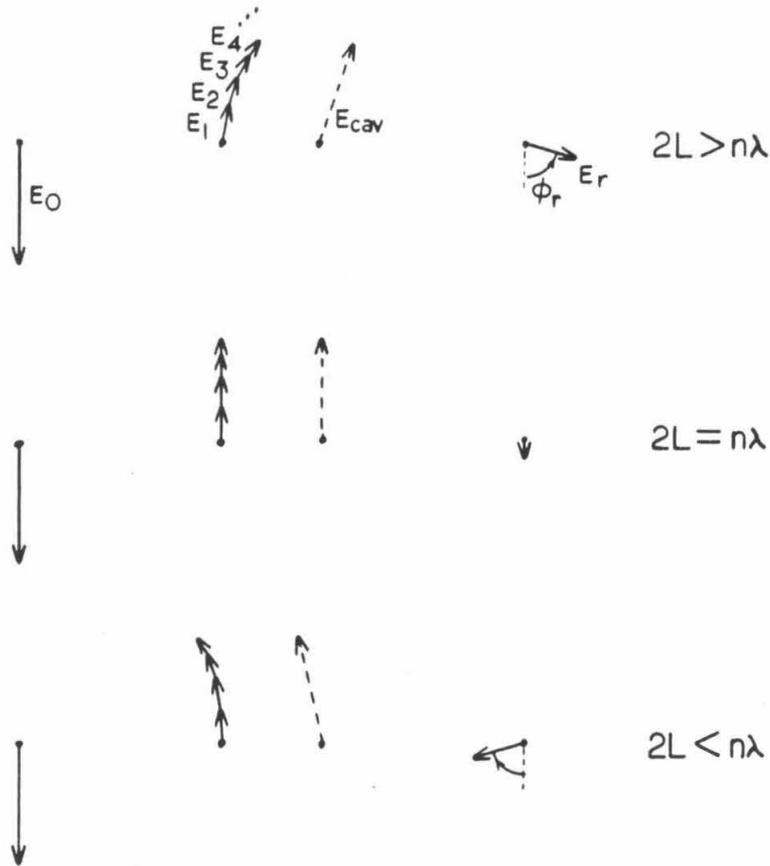


Figure 5: Schematic representation of the optical electric field reflected by the Fabry-Perot cavities in the interferometer of Figure 4, as seen in a frame which rotates with the laser field at $\Omega_{opt} \approx (2\pi) \cdot 6 \times 10^{14}$ Hz. The promptly reflected laser field E_0 acquires π phase shift on reflection with respect to the portion transmitted into the cavity. This transmitted portion bounces back and forth, leaving a fraction E_1, E_2, \dots back out through M_1 on each return bearing the imprint of geodesic $M_1M_2M_1$ in its phase. The sum of all these fields E_{cav} , and the resultant E_r after combining with E_0 , are shown for several values of the proper distance.

been varied between 125 and 450 Hz, depending on the transmission selected for the cavity input mirrors. The long storage times implied by these figures are achieved by using modern ultra-low-loss mirrors, which lose as little as 50 parts per million of incident power per reflection.

The “shot noise limit” to measuring ϕ , when expressed as a RPSD of equivalent strain, will be independent of frequency up to this breakpoint and will rise as f above it, the expected effect of averaging the optical phase shift over more than half a period of the gravitational waves. The actual limit depends on the exact method used to measure the optical phase and on losses and mismatching in the cavities; this has been calculated rigorously by Whitcomb [33].

1.5.4 Less Fundamental Sensitivity Limits

The Caltech interferometer has rarely been truly limited by photon shot noise over more than a narrow band of frequencies. Much of the work to be described was and remains concerned with unanticipated or underestimated spurious noise sources. Many of these were only discoverable in the instrument itself, leading to the inevitable subservience of its role as an observing instrument to that of a technology testbed. Each noise source is identified as it appears and comes to dominate the spectrum and subsequently eliminated, revealing the next obstacle.

One can divide the spurious noise sources identified to date among those associated with the test masses themselves, and those related to the lasers and external optics. Some potentially important noise sources, such as light scattering in the interferometer cavities and fluctuations in the refractive index of residual gas, fit into neither category; the reader is referred to [34], [35] and [36] for more information.

The remainder of this thesis is concerned with investigations leading to significant improvements of two key elements of the antenna, the test masses and the

laser light. Finally, some observations made with the antenna are discussed and analyzed for evidence of periodic gravitational radiation; this should be viewed not as a culmination of the improvements discussed, but as an interlude in the continuing process of refinement. Indeed, many of the technical innovations to be described have transpired since the observations. As one consequence of recent improvements, the antenna's strain sensitivity is approximately three times better than it was for the observing run (Figure 6). Even excluding the effects of local periodic interference, which hampered those observations and have now been greatly reduced, the antenna could then detect a given source over only one twenty-seventh the volume it can cover now. The observations are nevertheless important for two reasons; first, they were recorded soon after the very close supernova in the Large Magellanic Cloud, an unusual astrophysical event. Second, they place an upper limit on the overall flux of periodic gravitational waves bathing the Earth, improving on previous experiments by a considerable margin.

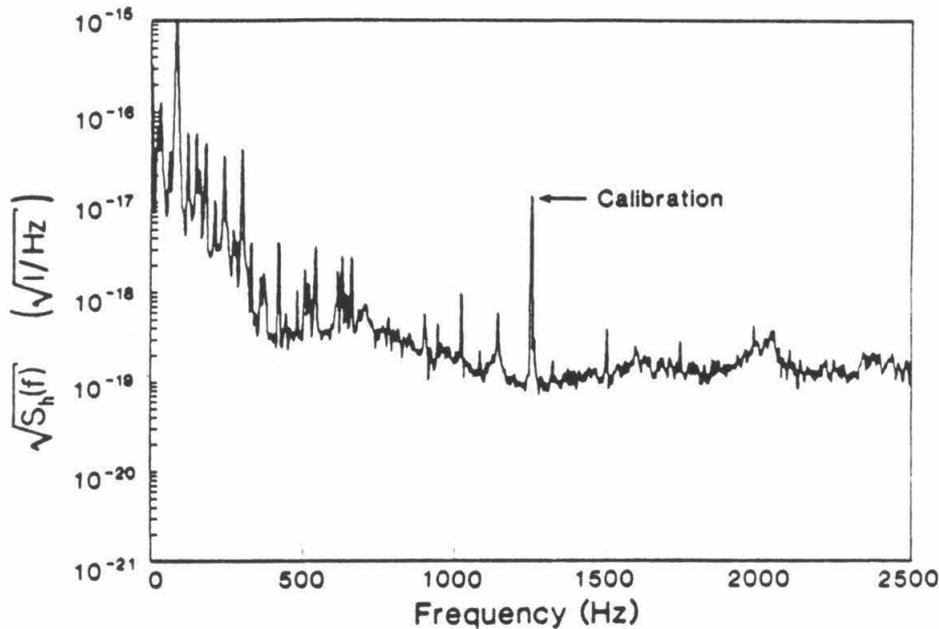


Figure 6: Strain RPSD of the Caltech 40 meter antenna in late 1987, showing its lowest noise level achieved to date. The calibration peak at 1.2 kHz was induced by driving one test mass with a sine wave at 10^{-15} meter RMS amplitude; the bandwidth of each bin in the spectrum is 4.7 Hz. Among the most recent improvements contributing to this performance are all-quartz optically contacted test bodies, magnetic cavity length servo systems, penetration of the single-mode optical fiber into the vacuum system and isolation of its output coupler, and the inclusion of a narrowband “mode cleaner” optical filter in series with the laser to reduce high-frequency laser fluctuations. Some of this work will be described here; the optical fiber improvements are mentioned in [37]. Above 1.3 kHz the spectrum is within a factor of two of the calculated shot noise [33] for the modest light power used, about 25 mW. The bump at 2 kHz may be an artifact of fluctuating secondary-cavity servo loop gain (Appendix D). At low frequencies the many peaks may be caused by resonances of the beamsplitter mass, which is of the obsolescent design described in Chapter 2, and acoustic vibration of the laser, fiber and other optics outside the vacuum enclosure. Since this performance was achieved, efforts have focussed on further reducing the bandwidth and increasing the power throughput of the mode cleaning cavity.

Chapter 2

The Test Masses

Making gravitational test masses behave like the pointlike inertial bodies discussed above has turned out to be a fairly difficult task. From a laboratory practitioner’s standpoint, this may not be surprising. The strongest gravitational waves might change the separation of the test masses in our 40 meter antenna by only 10^{-18} meter, one hundredth the diameter of a proton. Extraordinary freedom from external forces and internal distortions is thus required.

2.1 Seismic Isolation

Between earthquakes, the Caltech campus shakes in the horizontal plane with a displacement RPSD of order $10^{-13} \text{ m}/\sqrt{\text{Hz}}$ at 1 kHz. The spectrum rises very steeply at lower frequencies, roughly as f^{-3} [38]. Since the seismic vibrations are likely to be uncorrelated over distances much greater than the acoustic wavelength in the ground, one cannot rely on cancellation of “common mode” motion in an apparatus 40 meters in size. A factor of order 10^6 isolation between each test mass and the lab floor is thus needed at 1 kHz to reach the strain RPSD of $\sqrt{S_h(f)} = 10^{-19}/\sqrt{\text{Hz}}$ shown in Figure 6; considerably more isolation is required at 300 Hz.

The test masses of the 40 meter interferometer are isolated in several stages. The outermost shield is provided by separating the laboratory floor into four segments. Three small segments, at each end and in the corner area, are supported on concrete piles which extend approximately 5 meters downward to a relatively solid stratum. The optical tables and vacuum tanks are all supported by these three separate pieces of floor. Vacuum pumps, electronics, and people all occupy the remaining area which is built in the standard way on the surface. Walking on this ordinary floor does not affect the interferometer, but even treading lightly on the isolated segments noticeably disturbs the masses.

On each pad sits a Newport optical bench, supported by air legs. The air legs are not usually inflated; their off-the-shelf height servos are insufficiently stable to maintain the angle of the tabletop to the required precision. The tables are advertised to strongly damp internal vibrations. Vacuum tanks are strapped to each table, and an aluminum tower is bolted to the inside of each tank to support the next isolating stage. These are connected to 40 meter long, 20 cm diameter vacuum pipes via soft stainless steel bellows, arranged in pairs to compensate the force of atmospheric pressure; in this way the tanks (and thus the mass suspensions) remain largely decoupled from the pipes.

The next stage of isolation is provided by a stack of two lead weights separated by compliant toy rubber cars. The lower weight rests on more rubber cars placed on the tower; a massive platform which houses the mass suspension and its control transducers sits on still more rubber cars atop the upper lead weight. The intended effect is that of three coupled three-dimensional harmonic oscillators, damped by the rubber's viscosity. The fully loaded stack assembly has a resonant frequency of approximately 5 Hz in the horizontal plane, and 10 Hz vertically; the quality factor Q is less than 10 [38].

The final and most critical isolation factor is provided by the pendulum suspension of each test mass. In the Caltech instrument each mass is actually hung from a small intermediate mass, called a “control block,” which is constrained in translations but allowed to rotate in response to torque transducers; the test mass itself is hung by four wires in such a way that it follows the angular orientation of the control block (Figure 7). These wires are of hard carbon steel piano wire, 30 cm long and just thick enough not to break under the weight of the test mass¹. The length is chosen to place the pendulum frequency near one Hertz, and the wire mass (i.e., thickness) is then made as small as possible to maximize the frequency of the fundamental “violin” mode of the wire and the spacing between higher modes. In our detector the lowest violin resonances occur near 630 Hz.

For frequencies well above the pendulum frequency the result of a displacement δx_s of the suspension point is an acceleration of the mass $\ddot{x}_m = g\delta x_s/L$ where L is the wire length, so the displacement transfer function falls as f^{-2} . This behavior is interrupted by the wire resonances, which provide narrow transmission “windows,” and by the viscous damping of the pendulum which compromises its isolation at frequencies above the “quality factor” Q times its resonant frequency. The pendulum Q of each mass in the Caltech instrument is at least several thousand in vacuum, so this latter departure is not significant in the kilohertz band. Excitation at the fundamental and first harmonic wire resonances has been observed in the gravitational wave signal, but with some care to filter out signals applied to the suspension point transducers at these frequencies it has usually been suppressed.

¹A factor of two in breaking stress is reserved for safety. For two of the masses this left inadequate margin; the Richter 5.9 earthquake which occurred on 1 October 1987 at a distance of about 5 kilometers (and one Richter 5.1 aftershock) broke some wires. The masses were caught with minimal damage by safety stops included for this purpose.

Measurements of the linear transfer function of the lead/rubber stack assembly [39] and of the entire structure including the pendulum [40] indicate that the attenuation is more than adequate to block ambient seismic noise at frequencies above 100 Hz, although the consequences of nonlinear and nonorthogonal coupling have not been investigated in detail.

2.2 Acoustic Isolation

The low pressure in the experimental chambers ($\lesssim 10^{-5}$ torr) greatly reduces the risk of direct acoustic interaction with the test masses. Coupling of sound to the wire suspension point via some mechanical path from the air outside might be more serious. The transmission path for acoustic noise is thus largely indistinguishable from that for seismic noise. Attempts to measure direct acoustic transmission have revealed that optics outside the vacuum, especially the laser, provide the dominant path of entry for airborne acoustic interference. There is no explicit acoustic shielding on the vacuum system, although at one time an effective shield was built around the laser [41]. Acoustic interference complicated the data analysis presented in Chapter 4, by adding narrowband noise at a few discrete frequencies; however, it is not a major contributor to the broadband noise.

2.3 Suspension and Control

To avoid exciting high-order spatial modes of the optical cavities, the mirrors must be aligned within 10^{-5} radian [38, page 70]. The second-order effect of small angle fluctuations on the perceived length of a cavity with perfect average alignment gives an effect $\Delta l \simeq \rho l (\Delta\theta)^2$ where ρ is a dimensionless factor of

order unity which depends on the curvature of the cavity mirrors and $\Delta\theta$ is the angular deviation of the mirror [42]. Predicting the spectral density of the equivalent displacement induced by this effect requires assuming some form for the spectral density of the angle fluctuations at all frequencies above and below the frequency of interest. This has been calculated for the case where the angle fluctuations are due purely to the thermal excitation of bending modes of the mirror; for this “best case” limit, at least, the noise seems inconsequential [43].

Letting the optic axis stray a small amount d from the center of mass introduces a first-order coupling $\Delta l \simeq d\Delta\theta$ which can be far more serious. Other plausible mechanisms for mirror angle affecting the perceived cavity length abound, ranging from the trivial, e.g., applying torques about axes other than the principal inertial axes, to the baroque, e.g., sweeping the beam across small figure distortions or scattering centers. This latter mechanism would also upconvert low-frequency angle noise to high frequency path length jitter. Motion of the input laser beam may also multiply up an otherwise insignificant slow mirror angle fluctuation to higher frequencies [44].

In response to these potential problems, some of which are discussed quantitatively in [37], we have attempted to minimize angle deviations. The test masses are all servoed by 40 meter optical levers to reference photodiodes fixed on each optical table. A helium-neon laser beam, suitably expanded to endure diffraction, is sent through the evacuated pipe to each cavity mirror and reflected back to a quadrant photodiode near the laser. The photocurrents from the four quadrants are suitably summed and differenced to provide altitude and azimuth correction signals. These are amplified, filtered and fed to torque transducers on the mass control block. A secondary “local” control loop using 1 meter optical levers (for greater dynamic range) is provided for each mass to enable starting up from a state of total misalignment.

The orientation servo loops have unity-gain bandwidths of from 10 to 30 Hz; above this range the signals to the torque transducers are filtered strongly. The moment of inertia of the test mass is relied upon to resist AC torques at higher frequencies. Special care is taken to bring the axes of mass rotation (defined by the pivot points of the wires on the mass and the control block) as close as possible to the principal inertial axes. This reduces the coupling of angular acceleration to linear acceleration (and vice versa) and is particularly important in view of the intrinsic anisotropy of the antiseismic system.

The current form of torque transducer is depicted in Figure 7. The control block is hung from a single length of wire (or a pair of wires in an inverted "V" arrangement) which allows it to rotate and tilt about its geometric center. Additional constraint wires define the horizontal position of the control block without restraining its rotations. Small samarium-cobalt magnets are glued near the ends of the control block, and four copper coils are connected to high-current amplifiers in such a way as to induce torques in the block about the vertical axis and one horizontal axis. The other horizontal axis (the one nominally parallel with the optic axis of the cavity) is not servoed. The equilibrium distance of each magnet from its coil is adjusted to bring the force per unit current near its maximum value, so the transducer is linear to first order in angle. The four wires extending downward to the mass are arranged to transmit the applied torques to the mass and keep it parallel to the control block.

Translation of the control block along the optic axis can be effected by use of a piezoelectric stack translator affixed to the suspension platform. A stiff wire anchored to the free end of this translator passes through the center of the control block, which is locally clamped to it, and is held in tension by a spring on its outboard end. The "wire pushing" piezo effectively moves the suspension point of the mass along the optic axis by translating the control block; it is

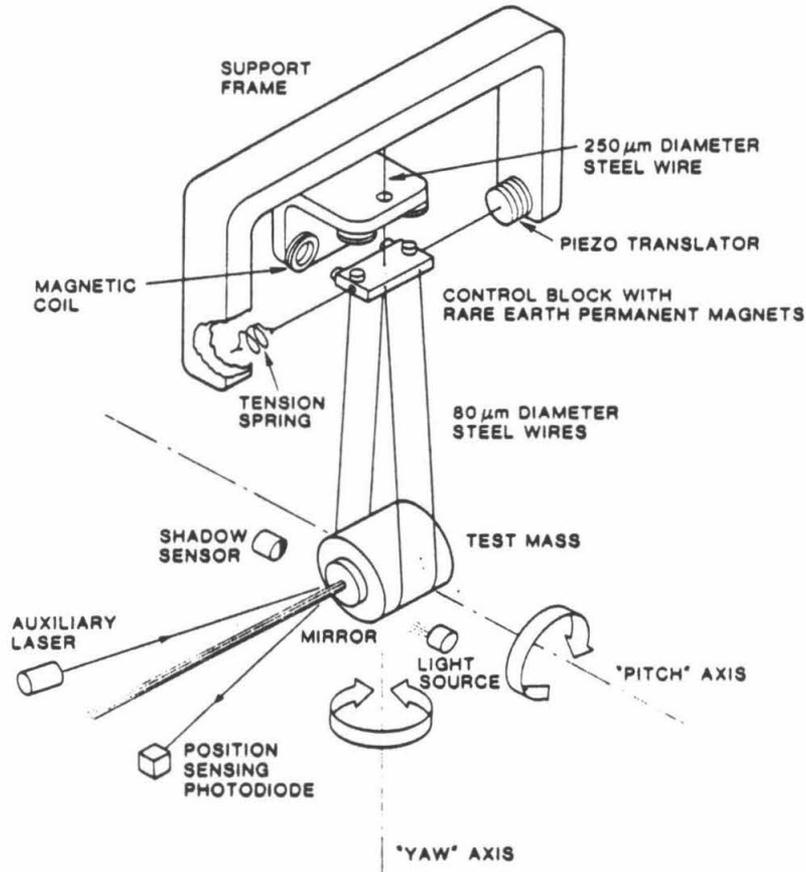


Figure 7: Semi-schematic depiction of one of the cylindrical test masses used since 1984, with transducers and sensors used to control its angle and position. The auxiliary laser and position sensing photodiode sense the angle of the mirror surface; the magnetic coil/permanent magnet system servo the pitch and yaw of the control block to keep the test mass aligned with this reference. The separation between the coils and their magnets is exaggerated for clarity. Also shown are the light source (an LED) and shadow sensor (a silicon photodiode) which comprise the “shark detector” sensor (§ 2.4), a device which senses the velocity of the mass and feeds corrections to the piezo translator at top in order to damp its seismically-excited swinging. The support frame above is supported inside the vacuum tank by the lead and rubber stack assembly described in the text. In the current arrangement four masses like this one carry the critical cavity mirrors, while a fifth more complex mass (like the one shown in Figure 9) carries the beamsplitter and other beam-handling optics.

driven to compensate drift of the cavity lengths at very low frequencies (below .3 Hz) and also to damp the seismically-excited pendulum motion (from .3 Hz up to 10 Hz). Above 100 Hz the piezo signal is strongly filtered to keep electronic noise from driving the wires and mass at interesting frequencies. Wire pushing signals are derived from two sources. When the cavity in question is servoed to an interference fringe, the signal from the servo can be low-pass filtered and fed to both wirepushers on that arm to help keep the main actuators within their dynamic range. However, if the cavity falls off the fringe, this signal disappears (or worse, turns into nonsense), and the masses may be seismically excited to huge amplitudes before the servo can reacquire. The amplitude can reach hundreds of fringes peak to peak at one Hertz, and the fringe servo may be unable to stop the mass rapidly enough as one flies by. The shark detector was developed primarily to deal with this annoyance.

2.4 Shark Detectors

The shark detector² was devised to sense the velocity of a swinging mass and generate a damping signal to be applied to its wirepushing piezo. The sensor simultaneously needs a dynamic range of some millimeters to be usable in serious “pendulum emergencies” (as often arise when people step on the isolated floor pads to make adjustments) and yet needs sufficient sensitivity to damp the mass motion to less than a fringe or so (about .25 micron peak). An auxiliary Michelson interferometer equipped with a fringe composing device had been used successfully in previous experiments (see [38, page 87]) but was deemed inappropriate because of its complexity and sensitivity to alignment, as well as the inaccessibility of some of the masses. A “shadow meter” device was settled on

²So named by Mack Hereld.

for compactness and simplicity. A light-emitting diode (LED) is focussed across a gap onto a 4 mm \times 4 mm square silicon photodiode. The edge of the test mass is interposed in the gap and simply casts a shadow on the photodiode (Figure 7). The LED beam is set tangent to the arc described by the edge as the mass rotates about its center of mass. This decouples the mass orientation from the sensor to first order.

Concerns that stray laser beams, room lights and scattered light from the resonating cavity would interfere with the sensor led to a design in which the LED is amplitude-modulated at a high frequency (100 kHz) and the light is coherently demodulated after detection. Although the device was only needed to act in a narrow band around the 1 Hz resonant frequency, it was speculated that modulated LED light might contaminate other signals and appear in noise spectra if the frequency was lower.

This demodulated shadow signal is differentiated over a limited bandwidth, amplified, and fed to the wire pushing piezo. Computer analysis of the circuitry [45] has indicated that with sufficient gain the filtering shown in Figure 8 will indeed afford critical damping of the pendulum motion. In a brief experiment with a powerful magnetic mass drive system, the shark detector signal was used to damp the mass from some millimeters amplitude to a standstill in less than one cycle (indicating overdamping) without instability. Because of the sharp narrowband filtering, the damping contributes no thermal Nyquist force to the pendulum at higher frequencies; the technique has thus been called “cold damping.”

The equivalent displacement RPSD of the sensors and circuit is approximately $100 \text{ nm}/\sqrt{\text{Hz}}$ near 1 Hz. By decreasing the gap, reducing the size of the beam and/or photodiode (sacrificing dynamic range) and increasing the photodiode luminosity, as well as using a diode preamp with lower input current noise, it is

expected that equivalent displacement noise levels below one $\text{nm}/\sqrt{\text{Hz}}$ can be achieved; for practical reasons the sensors were anchored to the vacuum tank, which is such a noisy platform that the sensor noise is not significant. At present the shark detector signals are applied to the wire pushing transducers with limited gain, damping the pendulum only to an effective Q of 30 or so. The masses are not critically damped because this would have the effect of locking them to the noisy vacuum tank, and there is currently no provision for cutting out the shark signals when the cavity acquires a fringe. In the future rather simple switching circuitry and an isolated, low-noise platform on which to mount the sensors should allow full use of the shark detectors' rapid damping capability without degrading the performance of the interferometer.

2.5 Thermal Noise

The $\frac{1}{2}k_B T$ of thermal energy populating each degree of freedom of an apparatus requires careful consideration when the instrument must register subnuclear displacements. A simple one-dimensional harmonic oscillator of mass m , resonant frequency ω_0 and quality factor Q maintains thermal equilibrium with its environment by experiencing an apparent force with spectral density $S_F(f) = 4k_B T \cdot m\omega_0/Q$. This “white noise” force results in a fluctuating displacement having power spectral density

$$S_x(\hat{f}) = \frac{4k_B T}{m\omega_0^3 Q} \left[(1 - \hat{f}^2)^2 + \frac{\hat{f}^2}{Q^2} \right]^{-1} \quad (14)$$

where $\hat{f} \equiv 2\pi f/\omega_0$ is the frequency normalized to units of the resonant frequency [38]. Rather than solve for the normal mode spectrum of the entire apparatus it seems a reasonable approximation (provided that Q 's are high) to separately evaluate the direct contribution of each mode to the thermal noise spectrum of the mirror displacement x as though it were an isolated simple harmonic oscillator. Perhaps the most transparent example is the pendulum motion of the mass as a whole; internal mechanical modes of the test masses themselves, however, have been more troublesome in practice.

Equation 14 reduces in the limit $\hat{f} \gg 1$ to $S_x(\hat{f}) = 4k_B T / (m\omega_0^3 Q \hat{f}^4)$ for high Q , so for a one-Hertz pendulum at room temperature the displacement RPSD at audio frequencies due to thermal excitation is

$$\sqrt{S_x(f)} \simeq 2.6 \times 10^{-19} \left(\frac{10^3}{Q} \right)^{\frac{1}{2}} \left(\frac{1\text{kg}}{m} \right)^{\frac{1}{2}} \left(\frac{1\text{kHz}}{f} \right)^2 \frac{\text{m}}{\sqrt{\text{Hz}}} \quad (15)$$

where the parameters chosen are roughly those of the current Caltech test mass suspensions. This may become a serious limit for the Caltech interferometer in the future, but larger masses ($\gtrsim 100$ kg) and greater Q 's ($Q \gtrsim 10^7$ has been demonstrated [47]) could forestall the eventuality for some orders of magnitude.

Troublesome thermal fluctuations can arise in internal modes of the test masses, which drive the mirror surfaces with respect to the centers of mass. The normal modes of each mass can have resonant frequencies in the frequency band of interest and/or at higher frequencies, invoking different regimes of Equation 14.

Take $m_e \lesssim m$ and y respectively as the effective mass and normal coordinate of the mode under consideration, and call the projection of this normal coordinate on the distance of the mirror surface from the center of mass $\alpha = x/y$ ($0 \leq |\alpha| \leq 1$). These parameters can be calculated rigorously, but for simple modes they can be estimated³. Each mode is thus described by a standard equation of motion

$$m_e \ddot{y} + \frac{m_e \omega_0}{Q} \dot{y} + m_e \omega_0^2 y = F_{\text{th}}(T) \quad (16)$$

where $F_{\text{th}}(T)$ is the effective thermal force discussed above, so the resulting mirror displacement is again given by Equation 14 with the substitutions $x \rightarrow y$ and $m \rightarrow m_e$. If the resonant frequency $f_0 \equiv \omega_0/2\pi$ of the mode in question is in the band where we are looking for gravitational waves, the height of the resonant peak in the spectral density may be of interest; taking $\hat{f} = 1$ gives $S_x(f_0) = 4k_B T Q \alpha^2 / m_e \omega_0^3$. The width $\Delta f = \omega_0/2\pi Q$ of the resonance must be compared with the frequency spread of expected signals; for some applications a few known, well-catalogued resonances are not terribly damaging. However, in general it is better to place most resonant modes at significantly higher frequencies than expected signals. If this can be managed, the appropriate limit $\hat{f} \ll 1$ of Equation 14 gives the frequency-independent PSD $S_x(f) = 4k_B T \alpha^2 / m_e \omega_0^3 Q$. Although there are a very large number of such modes (of order the number of molecules in the mass) whose effects in this pseudo-independent treatment should be summed in quadrature, only a few of the lowest ones will have a significant

³For example, the first longitudinal vibrational mode of a long cylinder with a massless mirror mounted on the end face will have $m_e \simeq m/2$ and $\alpha \simeq 1$; the first flexural mode of the same cylinder might have $m_e \simeq m/3$ but $\alpha = 0$.

coupling α .

The state of the art in test mass design has advanced considerably since work began on these interferometers, and it has taken some effort to make a workable system employing masses with specifications like those quoted in Equation 15. Some interesting work characterizing the original Caltech test masses and how subsequent mass designs evolved are described next.

2.6 Mass Engineering

Before 1984 the Caltech test masses, named Huey (northeast), Dewey (vertex), and Louie (southwest)⁴ each comprised two 36 cm aluminum disks joined in five places by 11 cm long posts (Figure 9). The upper disk was machined away in several places for access and was made correspondingly thicker to place the center of mass of the assembly midway between the disks. The lower disk was 2.5 cm thick and acted as a small optical table on which the various optical components were arranged. Cavity mirrors were placed in standard optical mounts glued down near the edge of each mass; the Dewey mass also held a beamsplitter, directional isolators, and various steering mirrors needed to snake the beam around the central axis (obstructed by the suspension wire attachment post). All three masses carried separate auxiliary mirrors in standard adjustable mounts for the local and global orientation servos.

To see if the complicated structure could be contributing to the excess of noise afflicting the interferometer at the time, a dummy mass was instrumented with strain transducers and several readout systems to measure the frequencies and quality factors of some of the normal modes, and if possible to estimate their effect on the motion of the cavity mirror. The investigation is summarized in

⁴After Donald Duck's three mischievous nephews [48].

coupling α .

The state of the art in test mass design has advanced considerably since work began on these interferometers, and it has taken some effort to make a workable system employing masses with specifications like those quoted in Equation 15. Some interesting work characterizing the original Caltech test masses and how subsequent mass designs evolved are described next.

2.6 Mass Engineering

Before 1984 the Caltech test masses, named Huey (northeast), Dewey (vertex), and Louie (southwest)⁴ each comprised two 36 cm aluminum disks joined in five places by 11 cm long posts (Figure 9). The upper disk was machined away in several places for access and was made correspondingly thicker to place the center of mass of the assembly midway between the disks. The lower disk was 2.5 cm thick and acted as a small optical table on which the various optical components were arranged. Cavity mirrors were placed in standard optical mounts glued down near the edge of each mass; the Dewey mass also held a beamsplitter, directional isolators, and various steering mirrors needed to snake the beam around the central axis (obstructed by the suspension wire attachment post). All three masses carried separate auxiliary mirrors in standard adjustable mounts for the local and global orientation servos.

To see if the complicated structure could be contributing to the excess of noise afflicting the interferometer at the time, a dummy mass was instrumented with strain transducers and several readout systems to measure the frequencies and quality factors of some of the normal modes, and if possible to estimate their effect on the motion of the cavity mirror. The investigation is summarized in

⁴After Donald Duck's three mischievous nephews [48].

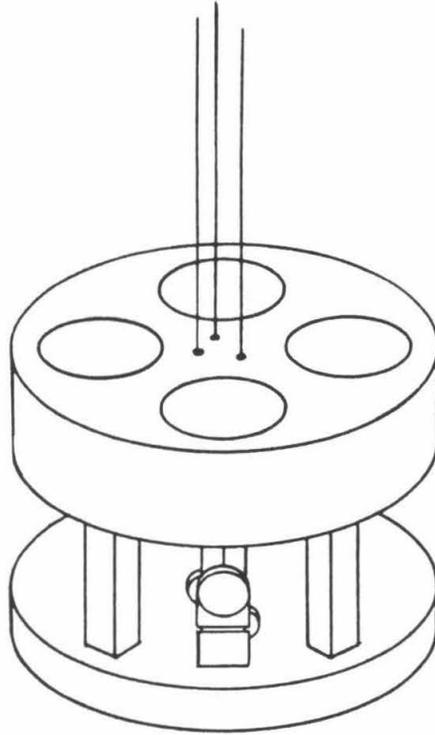


Figure 9: One of the three test masses used before 1984. It is suspended by three wires attached to the center post near the midplane; the optic axis is arranged to intersect the center of gravity, which also lies approximately in this plane. Adjustable mirror mounts (foreground) were attached near the edge with glue.

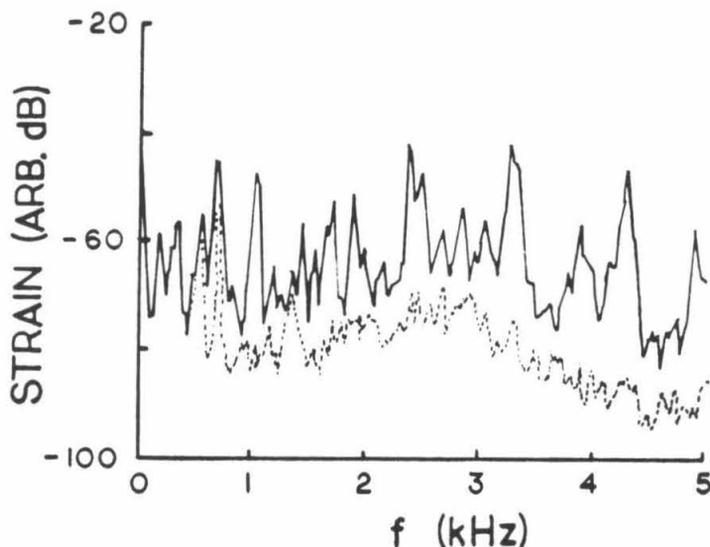


Figure 10: An uncalibrated interferometer strain spectrum from 1984 with (solid) and without (broken) pseudorandom white noise applied to a piezoelectric exciter bonded onto Huey.

Appendix A. Briefly, it was concluded that many resonances with Q 's in the low hundreds coupled strongly to the motion of the cavity mirrors. This is shown in Figure 10, in which artificial mechanical excitation was applied to one of the test masses while the interferometer was running to gauge the effect of internal vibration on the strain signal. Many resonances were enhanced which matched features in the strain spectrum; modes of the other two masses might well have contributed the remaining features. The excitation of thermal noise in these resonances could account for the heights of many peaks in the interferometer noise spectrum, even without any transmission of acoustic energy from the outside.

At around the same time the German and Scottish groups were discovering

similar results, and the decision was taken to rebuild the test masses to simplify the mode structure, increase the Q 's, and push all significant resonances which might affect the mirror position to much higher frequencies. The new configuration adopted at Caltech was shown in Figure 7. The end masses were replaced by right circular cylinders, 10 cm long and 10 cm in diameter, of brass (Huey) or aluminum (Louie) with the cavity mirrors attached to piezoelectric washers glued on each cylinder. The two piezoelectric washers stacked between Louie and its mirror were used to servo the length of that cavity to the laser light wavelength (Section 3.3); the single piezo washer between Huey and its mirror was used to calibrate the interferometer by inducing a known displacement. Each cylinder was suspended with two loops of hard steel piano wire, kinematically fixed by grooved pegs press-fit into the mass.

At the vertex Dewey presented some problems. It was desirable to retain all the beamsteering and isolation functions provided by optics on the existing big mass and just remove the cavity mirrors onto simplified masses of their own. However, the disks occupied the entire vacuum chamber area save a 5 cm annulus at the edge. The conflict was resolved by clearing off the regions of Dewey where the old cavity mirror mounts had resided and allowing new cylindrical masses to intrude between the disks, as well as into the mouth of the 40 m pipe. New suspension frames with torque transducers and especially compact control blocks were hung off the edges of the existing support structure, just clearing the inside of the vacuum enclosure. The two wire loops cradling each cylinder were placed just 2.5 cm apart to avoid touching Dewey or the inside of the tank. Cavity mirrors were attached directly to the cylinders, which were both made of brass.

While the work pushed on some dummy cylinders without mirrors were instrumented with small piezo strain transducers; the mode spectra were measured

by exciting one piezo (glued near the equator) and sensing the voltage induced on another piezo glued nearby. The cylinders were hung by loops of wire to simulate actual conditions. The lowest mode of the brass test bar was at 15.2 kHz, and the Q 's of all the low frequency modes were close to 10,000. For such resonances occurring well above frequencies of interest, the white thermal noise contribution to the displacement RPSD will be

$$\sqrt{S_x(f)} \approx 10^{-19} \alpha \left(\frac{1 \text{ kg}}{m_e} \right)^{\frac{1}{2}} \left(\frac{10 \text{ kHz}}{f_0} \right)^{\frac{3}{2}} \left(\frac{10^4}{Q} \right)^{\frac{1}{2}} \frac{\text{m}}{\sqrt{\text{Hz}}}. \quad (17)$$

These findings suggested that the new masses should solve the problem, providing that the mounting of the mirrors and piezos introduced no complications.

As evidenced by the calibrated noise spectrum in Figure 11 the new masses really improved not just the level of the noise but also the usable low-frequency signal limit of the interferometer, particularly important in view of the low-frequency bias of many expected astrophysical sources.

Modest improvements to other aspects of the interferometer over the next two years brought the performance to an apparently white strain RPSD of approximately $2.5 \times 10^{-19}/\sqrt{\text{Hz}}$ over a broad band extending from 1 to 3 kHz. Shot noise and other obvious noise sources were ruled out by direct tests. At this time the Glasgow detector also seemed limited at about this same displacement sensitivity, while the delay line detector under development in Garching, West Germany was about a factor of three better. At Glasgow a serious effort was put forth concentrating on the differences between the Garching test masses and the masses used there and at Caltech. Of the many differences one that stood out was the fact that using a delay-line system had allowed the Garching group to employ masses that were essentially one piece (the mirror itself) with no piezoelectric transducers. The possibility of noise or noise-inducing nonlinearity in piezoelectric materials was investigated carefully at Glasgow [49] and at

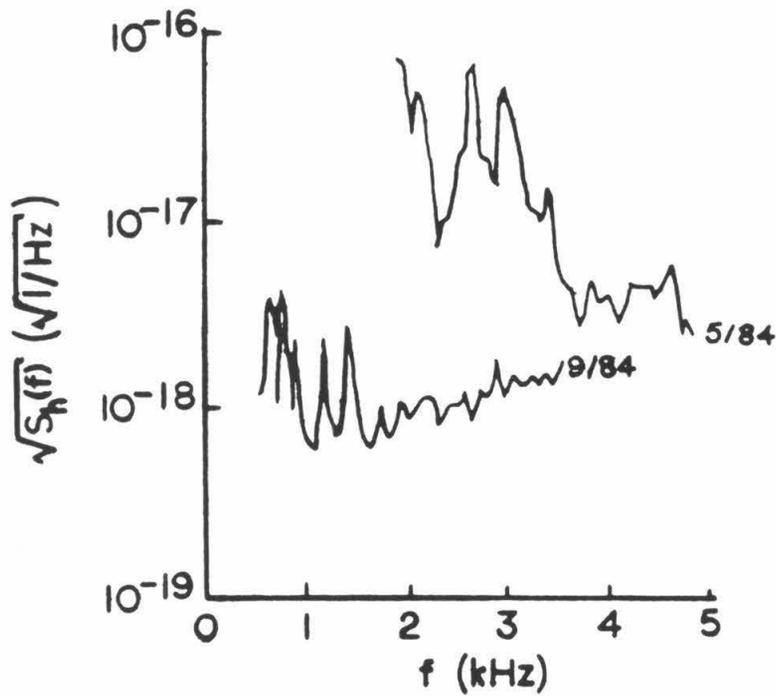


Figure 11: The strain RPSD is shown before (5/84) and after (9/84) replacement of the three test masses with five mechanically simplified, separate test bodies. Much of the remaining structure below 1.5 kHz is actually due to line frequency interference and excitation of the suspension wire violin resonances, features which are not resolved in this low-resolution spectrum. They were subsequently reduced. The noise above 2 kHz is very close to the calculated contribution of photon shot noise for the low laser power used (less than five milliwatts).

Caltech [51] with negative results; in the course of their measurements, however, the Glasgow group uncovered a source of noise in the joints between the piezos and the mirrors and masses.

Mirrors were then mounted by sticking them to a thin layer of viscous vacuum grease (e.g., Apiezon ‘L’). It was believed that at kilohertz frequencies the viscous grease would resist flowing and respond more like a solid, while allowing disassembly and differential thermal expansion. Tests revealed that the mirror was actually only weakly coupled to the mass at high frequencies [43]. Thermal excitation of the low Q , low-effective-mass “differential” mode of the mirror/mass system might explain the noise level. Grease joints were eliminated from the masses; mirrors and piezos were bonded to the masses and to each other using a type of glue known to have extremely low viscous losses [52]. This glue was applied by heating it and the pieces to be joined to 100°C and pressing the pieces together with weights to squeeze the glue very thin. The complete assembly was then cooled slowly to room temperature.

Unfortunately the rigid bond between the fused silica mirror and the metal mass (or the piezoelectric ceramic washers) now resisted differential thermal expansion, and the mirrors were warped. The interferometer was usable, but the cavity shape was distorted such that the optical performance was poor and fluctuated rapidly. Indeed, the light could often be seen resonating in two nominally orthogonal spatial modes at once, implying new susceptibilities to angle fluctuation. This was the state of the interferometer for the observing to be described later; during this run, the noise dropped temporarily (possibly due to fortuitous thermal conditions) to approximately $\sqrt{S_h(f)} \approx 1.8 \times 10^{-19}/\sqrt{\text{Hz}}$ near 1.5 kHz. This suggested that despite its impracticability the new bonding method had reduced the noise level significantly.

The current generation of test bodies solves the thermal distortion problem

by making the mass of the same material as the mirror. In the latest Glasgow and Caltech implementations the mirrors, coated on fused silica substrates, have been optically contacted to ground and polished cylinders of fused silica. No piezo transducers are used. The molecular bond achieved through optical contacting eliminates glue altogether, and the matched expansion coefficients not only prevent figure distortion but also reduce internal stresses (potential energy sources for “lock-and-slip” noise generation). Small pieces of ground quartz tubing bonded to the cylinder define the pivot points for the two wire loops supporting each mass. The Q of a bare 9 cm long, 10 cm diameter fused silica cylinder was measured to be over 80,000 in its first longitudinal resonance, at approximately 30 kHz [53]. This would imply a thermal displacement RPSD below $10^{-20}\text{m}/\sqrt{\text{Hz}}$ at 1 kHz. Similar measurements on a smaller optically contacted composite system indicate that such joints may not appreciably degrade the mechanical quality factor over that of a single piece [53].

A significant obstacle to introducing this clean, almost-joint-free technology was posed by the lack of a ready replacement for the piezo transducers, previously used to servo the optical cavities and to provide direct calibration. A magnetic force transducer system was developed for this purpose. Further improvements in sensitivity have resulted since the move to all-silica masses, although many noise sources not directly linked to mass construction are suspected in this regime and it is difficult to assess their relative contributions. Future test bodies may comprise single pieces of fused silica or sapphire, with mirror coatings applied directly to their polished surfaces.

Chapter 3

The Lasers

An ideal laser beam in free space is a simple object; it has a fixed wavelength, unvarying direction, constant amplitude, and a perfect Gaussian profile. Of course, real laser beams only roughly approximate these qualities. We have found that to use lasers for the ultraprecise measurements gravitational wave astronomy demands, even tiny departures from these ideals may be devastating. Each degree of freedom must typically be controlled, carefully decoupled from the phase measurements, or both. Some qualities, for example the laser wavelength, must be stabilized to the ultimate precision allowable by the shot noise in the light itself, requiring carefully optimized control systems.

Two of the laser's more unruly degrees of freedom are output power and wavelength. Both quantities fluctuate considerably in the argon ion lasers chosen for this work. The wavelength stability has actually been the more serious obstacle for us, although it is likely that power stability will become a problem at some future sensitivity. Geometric fluctuations in laser beams and cavities, and their effects on precise phase measurements, have been explored by Sheryl Smith and others [37,42,44,54]; geometric effects greatly complicated much of the work described here, but will only be mentioned briefly.

3.1 Laser Intensity Noise

Commercial argon ion lasers typically exhibit intensity fluctuations far in excess of their photon shot noise at low frequencies, and sometimes well above one megahertz. This noise results partly from mechanical vibration of the laser resonator which influences the cavity alignment, partly from fluctuation of the plasma tube excitation current, and partly from turbulence in the lasing medium. At least the first two noise sources, which appear to dominate at low frequencies, could be improved enormously with modest effort by laser manufacturers.

In the interferometer configuration we have chosen, the measurement is arranged to be insensitive to the laser intensity to first order. This is done by holding both cavities at interference minima; then, the intensity varies only quadratically with the phase. The interference condition is monitored by dithering the laser phase at a radio frequency (12.335 MHz) and looking for a component in the reflected cavity light fluctuating at this frequency (see § 3.2 and Appendix B). This “transformation” of the low frequency measurement up to a radio frequency band allows only the intensity noise components near the RF carrier frequency to contaminate the phase measurement; although the laser is not guaranteed to be quiet here, it is sufficient that shot noise dominates in the sample of the light actually hitting the photodetector. In experiments so far, using less than 50 milliwatts of the several watts available from the laser on each detector, this has indeed been the case; the RF photocurrent detected on a photodiode is dominated by shot noise.

However, many parametric and nonlinear couplings are still available to intensity noise. Some important mechanisms occur if the fringe servo doesn’t quite enforce the null condition and holds the optical phase somewhere off the true interference minimum. For example, since the servo loops have finite loop gain,

and the buffeting of the phase by various noise sources is only reduced by this finite factor, the cavities will typically be somewhere slightly off the minimum. As another example, there may be some D.C. offset in the servo loop; such things result from offsets in electronic amplifiers, or from RF leakage producing a fake phase error signal. The servo loop will compensate for the offset by introducing an equal and opposing phase error to cancel it out; the optical phase is locked somewhere off the true interference minimum. Now, the coefficient linking optical phase to detected photocurrent is clearly proportional to the light intensity (just the optical gain); so if the intensity decreases, for example, the phase offset will have to increase to maintain cancellation of the fixed voltage offset.

Intensity coupling can be hard to evaluate experimentally. Just putting a sinusoidal test modulation onto the laser intensity and looking for the resulting peak in the interferometer signal's spectrum can be deceptive, as the phase fluctuation inducing the coupling can easily be reversing the sign of the effect many times during the spectrum accumulation. Putting a stronger modulation on (to try to see the effect in the time domain), or using incoherent random noise that won't be averaged away, may strain the limits of linearity of the electronics or otherwise extrapolate poorly to the level of the natural excitation. One reassuring way of evaluating the contribution of intensity noise (and also the way to reduce it if necessary) is to directly reduce the laser fluctuations as much as possible and see if other things change.

Commercial intensity noise reduction devices are available, but are often not sufficiently effective or introduce added complications. They generally work by using a quiet reference photodiode to sense changes in the laser output and servo the transmission of some variable optical element to hold the output constant. By carefully considering the requirements for our particular application it has been possible to design and build considerably better intensity servos. Two such

systems were implemented at different times in our apparatus; in both cases the operation of the servo clearly reduced the intensity noise but not the overall strain noise, indicating that all sources of laser intensity coupling were exceeded by other noise contributions at those times. The observations to be described in Chapter 4, and most other tests, were made with the installed servo inactive.

3.1.1 Electrooptic Feedback

One commercial intensity servo device (a Coherent Model 307 “Noise-Eater”) was based on a four-crystal electrooptic cell arranged to produce rotation of the incident optical field polarization in response to applied voltage. The light was then analyzed by a fixed polarizer, turning the rotation into intensity modulation (Figure 12). The transmission of the cell/polarizer unit thus varied as the \cos^2 of the applied voltage; for small deviations about some fixed mean transmission, the extinction curve was sufficiently linear for servo applications. A beamsplitter directed a sample of the output onto a fast photodiode, whose signal was subtracted from a stable D.C. bias, amplified, and fed back to the electrooptic cell. The unit was specified to have a ten megahertz unity-gain bandwidth and up to 40 dB gain. Two seemingly contradictory problems arose; first, that this gain didn’t seem high enough given the potential bandwidth, and second, that the unity-gain frequency was uncomfortably near the frequency of the phase modulation used to derive fringe signals in the interferometer. Even if the device could perform flawlessly at this frequency by reducing the fractional intensity fluctuation to the level of the shot noise in the small diverted sample, this fractional fluctuation would still be far in excess of the intrinsic shot noise fluctuation in the remaining output light. As it happened, electronic noise and incipient unity-gain oscillations conspired to make the noise considerably higher than shot noise at ten megahertz.

A new electronic system was designed to *decrease* the bandwidth of the servo so that it would not add noise to the light at any frequency, and at the same time to increase the gain at low frequencies as much as possible given that limited bandwidth. To sharpen the point, it would make no sense to servo the intensity to a reference noisier than the laser light already is at high frequencies, where the laser noise falls off and is eventually dominated by shot noise. At low frequencies, on the other hand, the laser noise is so high that more than 40 dB of loop gain is necessary to suppress it fully.

For these reasons the design was chosen to have a unity-gain bandwidth of one megahertz. The light sampled was to be about one milliwatt (the laser then in use was giving about 100 milliwatts total power). The circuit is shown in Figure 13. Its electronic noise level was intended to be less than shot noise at one milliwatt of diverted sample power, but in practice was measured to be about equal to the shot noise. The loop gain Bode plot (Figure 14) shows a $1/f^2$ slope below 1 kHz and $1/f$ slope above. The low-frequency gain might be increased still further by moving the breakpoint closer to the unity-gain frequency, or by using a still greater slope at low frequencies. However, the dependence of the effective overall gain on the electrooptic cell voltage (as the slope of the $\cos^2 \theta$ extinction curve changed) sparked concern that a narrow conditionally stable gain region might limit the usable dynamic range, or even “lock out” the servo from entering stable operation at all.

The electrooptic cell has a nominal “half-wave voltage” (i.e., that voltage required to rotate the incident polarization by π radians) of 380 volts. The commercial servo had a full dynamic range of twice this, so that it would be possible to select from two stable operating points on the cyclic extinction curve. This was intended to allow for the thermal drift of the cell’s “zero-voltage” rotation value which could shift the polarization corresponding to a given voltage by as

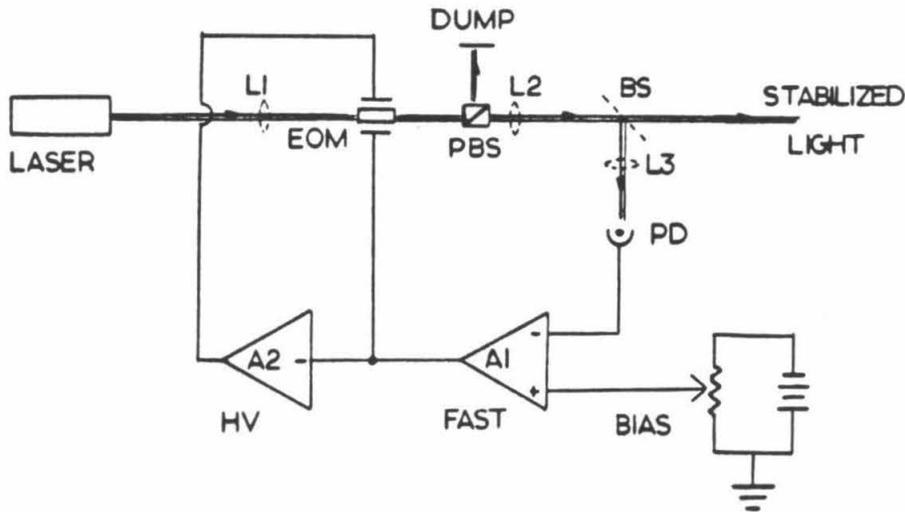


Figure 12: A semi-schematic layout of the electrooptic intensity servo. A small sample of the light is deflected by beamsplitter BS to photodiode PD. The photocurrent is subtracted from a stable reference bias, amplified, and fed back to the electrooptic modulator EOM, which rotates the polarization of its output light in proportion to the applied voltage. The polarizing beamsplitter analyzes the light and turns the rotation into amplitude modulation, stabilizing the light falling on the detector (and thus the output light) to a fixed intensity. The reduction in intensity fluctuations was apparently limited by beam geometry interactions; focussing the beam down to a smaller diameter inside the electrooptic modulator using lens system L_1L_2 , and simultaneously placing the sensing photodiode at the focus of L_3 , dramatically improved the servo's performance.

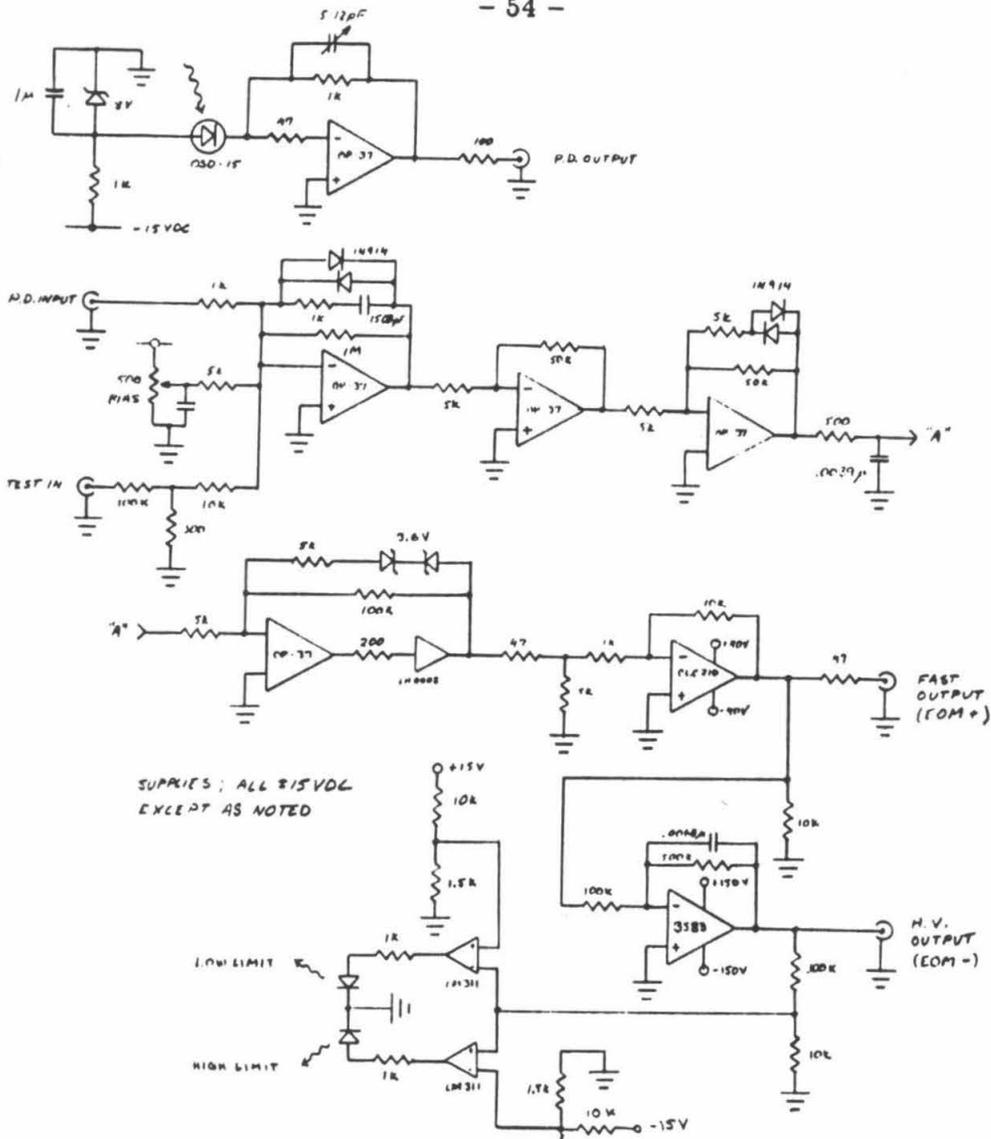


Figure 13: Electronic circuit of the servo amplifier built for the electrooptic intensity servo. The "Fast Output" and "HV Output" signals are applied directly to opposite plates of the electrooptic cell. An LED (lower left) glows when the correction voltage nears either limit, warning the operator that the laser or electrooptic cell has drifted out of range. The 500 Ω potentiometer is used to adjust the DC light reference. The remote photodiode/preamp circuit (top) was housed in a separate shielded enclosure, and joined by shielded signal and power cables to the remaining circuitry. The low-ripple power supply required to obtain good suppression of line frequency harmonics is not shown.

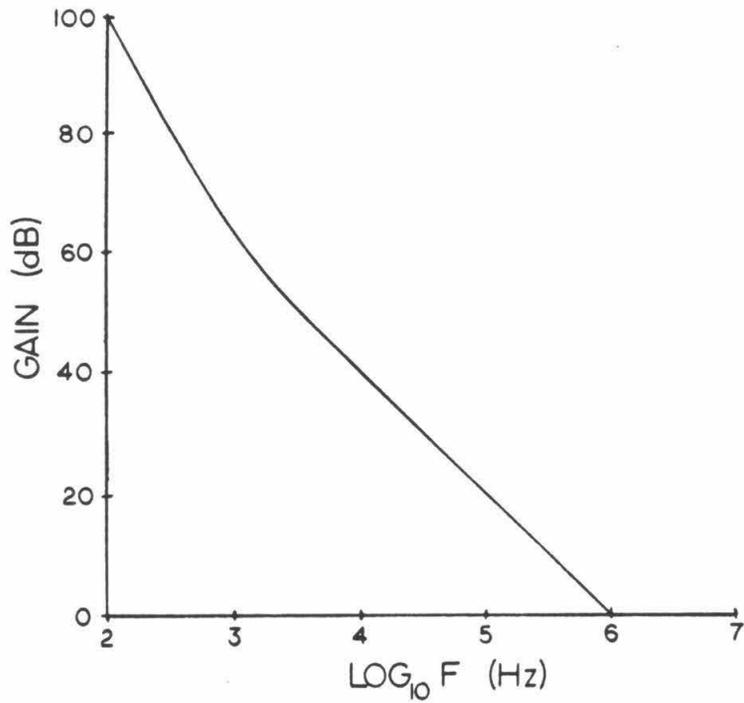


Figure 14: Estimated gain of the electrooptic intensity servo, shown for a diverted sample of 1 mW striking the sensor. The servo is designed to remain stable over a 100 : 1 range of overall gain, to allow for the variable slope of the extinction vs. voltage curve.

much as $\pi/2$. The commercial system also had provision for jumping to the alternate operating point when the voltage approached its limit. These considerations were ignored in the homebuilt system for simplicity. The high frequency feedback is applied to one plate of the cell by a medium-voltage high speed operational amplifier (Comlinear CLC210AI), with the low frequencies handled by a high-voltage operational amplifier (Burr-Brown 3583) driving the opposite plate. This gives the servo a total dynamic range of about 360 volts at D.C., sufficient to cope with thermal drifts over periods up to an hour when set up carefully. In practice a more serious operational difficulty is the propensity for slipping “over the top” on the extinction curve during transient dips in the laser intensity, hanging the servo up on the wrong slope. Frequent resetting was needed unless the system was set up with an average extinction near 50%, very wasteful of laser power (especially when added to the electrooptic cell’s high insertion loss of 10% and the intrinsic losses in the polarizers).

Results with the system were at first disappointing. The servo reduced the laser fluctuations, but not nearly by the anticipated factor. The whole seemed worse than the sum of its parts; open-loop tests of the each component showed up no forgotten noise terms, and the photocurrent in the servo sensing diode was suppressed to near the photon shot noise (of course, independent photodiodes outside the servo loop were used to evaluate the loop performance). One answer seems to lie in a conspiracy of two parasitic effects, one in the electrooptic cell and one in the photodiode. Electrooptic intensity modulators rotate the light polarization by inducing a voltage-dependent birefringence in each crystal. The birefringence causes one polarization component to follow a different optical path through the material. In these instruments the four crystals are arranged to cancel the beam “walk-off” as well as possible, but because of inevitable manufacturing tolerances, some voltage-dependent deflection persists. Also, the

walk-off (as well as the intended polarization effect) depends to some extent on spatial position inside the cell. While attempting to control the intensity, correction signals also altered the beam direction. Photodiodes (including the servo sensor and any independent monitors) have slight nonuniformities in their photosensitivity; moving the beam to a less sensitive spot on the servo diode will decrease the detected photocurrent just as well as actually cutting the beam intensity. Of course it is impossible to null all these degrees of freedom with just one sensor and actuator. Further complications might be induced by wiggle in the original laser beam itself. The cumulative effect is a "floor" below which the servo cannot decrease any intensity fluctuations without increasing the parasitic fluctuations equally. In typical tests the floor was at about 50 dB attenuation near 1 kHz (depending on the relative positions of servo and monitoring detectors), not enough to remove the biggest features in the intensity spectrum and far shy of the 70 dB expected on the basis of the calculated gain.

To reduce the parasitic effects, the laser beam was focussed down inside the electrooptic cell from its initial diameter (about 2 mm) to about 0.2 mm with a lens system. This improved the performance somewhat, presumably by reducing the interaction with nonuniformities in the cell, but when the servo sensing photodiode was also placed at the focus of another lens the improvement was dramatic (see Figure 15). The high measured attenuation at low frequencies actually exceeded that needed to bring the natural laser noise down to the shot noise in the sample, so the laser noise was artificially enhanced to measure the transfer function. The result is consistent at kilohertz frequencies with the calculated loop gain. The slight worsening of the attenuation above 10 kHz with the lens system, however, is not well understood. When later installed in the interferometer, the servo was closed around the single-mode optical fiber used to "clean" the laser mode for the main interferometer; this should have had an

even stronger effect of isolating the photodiode from geometric fluctuations in the laser and electrooptic cell, but the complexity of the interferometer’s optical system (much of which was enclosed in the loop) hampered efforts to measure the performance repeatably.

These findings closely parallel those from Glasgow [57], where a very similar system was constructed independently around the same time. In retrospect, the “bottoming out” of noise reduction despite increasing gain may be the reason for the modest gain designed into the commercial unit. Unpublished measurements from Glasgow [49] made after the inclusion of a single-mode optical fiber within the feedback loop of their intensity servo indeed verify that this step further improves the immunity from parasitic couplings, and allows better ultimate intensity servo performance.

With the lens systems in place (and artificial noise turned off), the servo was able to reduce the fluctuations in the laser light to nearly the limit set by shot noise in the diverted sample (actually, to the sum of shot noise and electronic noise in the diode preamplifier). Much of the noise power in the laser light was at multiples of the 60 Hz line frequency; careful shielding and grounding, and construction of a low-ripple DC power supply for the servo electronics, were required to remove interference at these frequencies.

At various times after installation in the main system, the servo was activated to test the influence of intensity noise on interferometer performance. Usually, there was no effect, and the servo was switched off. Eventually, the excessive losses in the electrooptic cell precipitated its removal.

3.1.2 Acoustooptic Feedback

The need for a ready means for reducing laser intensity fluctuations came up again later, at an opportune time. An acoustooptic modulator (AOM) had been

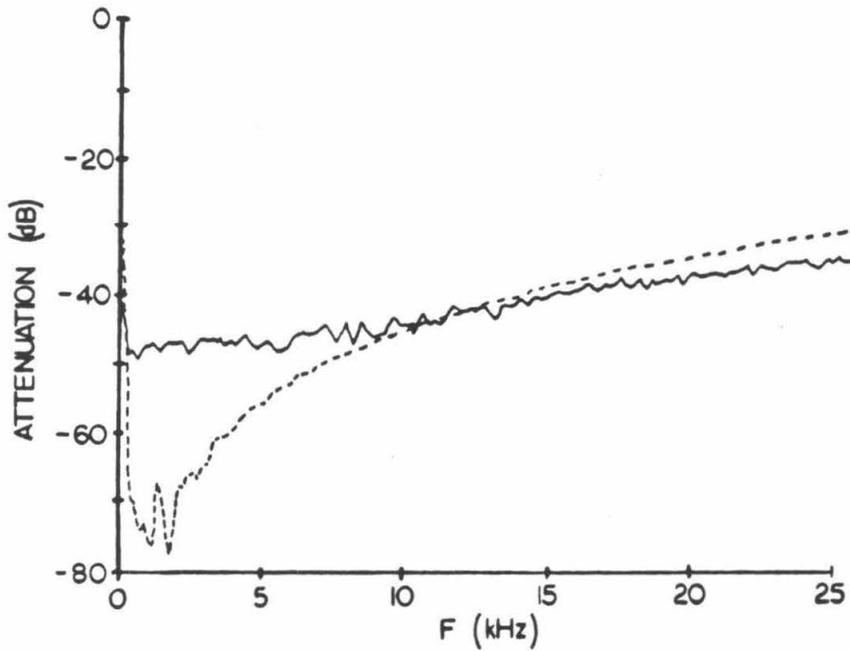


Figure 15: The measured attenuation of artificially enhanced laser intensity noise, before (solid) and after (broken) adding lenses to reduce the beam diameter inside the electrooptic cell and to focus the sample beam on the sensor. The degradation in performance above 10 kHz was not investigated; with the lenses, the attenuation measured at one kilohertz is consistent with the calculated servo loop gain (Figure 14).

installed in the system to act as an optical isolator between the laser and other optical components. Briefly, these devices work by setting up an acoustic travelling wave in an elasto-optic material (one whose index of refraction depends on mechanical strain). The wave sets up a moving diffraction grating in the material; light entering the sound field at a glancing incidence is diffracted by the periodic index fluctuations through a small angle. The motion of the grating (at the speed of sound in the medium) Doppler-shifts the diffracted beam with respect to the input by an amount equal to the acoustic wave frequency, typically 30 to 300 MHz. This allows the AOM to isolate lasers from subsequent optics, since light leaking backwards through the system is at a slightly different wavelength than the laser light, and therefore is not resonant with the laser cavity. Here the useful quality is the ability to change the efficiency of the diffraction grating by changing the intensity of the sound field. The acoustic wave is generated by electrically exciting a piezoelectric transducer bonded to the elasto-optic medium; by modulating the RF power to the piezo, one can change the grating depth and hence the transmission of power into the diffracted beam. The undiffracted power is left in the original beam, which warms a beam dump.

The bandwidth such a system can achieve as a servo actuator is limited in principle by the causal delay introduced as any corrective change in the acoustic wave intensity propagates at the speed of sound toward the interaction region. By contrast, in an electro-optic cell the adjustments propagate at the speed of light. It is thus desirable in this application to pass the beam as close as possible to the transducer (minimizing the mean delay) and to make the beam as small as possible (minimizing the delay gradient across the beam). Although it is theoretically possible to focus the beam to only a few optical wavelengths diameter, and to bring it almost this close to the transducer, the medium typically cannot handle too much optical power density without destruction (leading to

a bandwidth/power tradeoff) and the acoustic field is often not uniform close to the transducer (thus a bandwidth/efficiency tradeoff). Nevertheless, real-time bandwidths of several megahertz have been achieved in acoustooptic servo systems [58,59].

The AOM then in use (an IntraAction ADM-40) has a flint glass interaction medium, and operates at 40 MHz nominal acoustic frequency. Glass is not the fastest medium used in these modulators, but it is quite efficient; with careful adjustment the modulator could transfer 90% of the incident power into the diffracted beam, including optical insertion losses¹. The calculated delay $\tau \approx d/c_s$ is about a microsecond for $d \approx 3$ mm \approx the distance from the transducer to the beam aperture, and $c_s \approx 3$ km/sec \approx the speed of sound in the glass. This is equivalent to a phase shift of $\pi/2$ radians at a frequency of 250 kHz; additional delays and phase shifts in the electronics cut the unity-gain bandwidth actually achieved to just over 100 kHz. The modulator was provided with its own oscillator/driver (IntraAction DE-40BM), which was provided with external amplitude modulation capability. For convenience, the original sensing photodiode and preamp (Figure 13) were recycled, and a commercial low-noise amplifier (Ithaco 1201) provided the gain. The DC reference voltage was provided by a battery. The amplified control signal was fed into the amplitude modulation input on the AOM driver, after suitable filtering and limiting (to avoid destruction of the driver). The schematic is shown in Figure 16.

While this servo cannot have the loop gain or bandwidth to match the performance of the electrooptic system, it attenuates the noise enough (at least 30 dB at 1 kHz) to show a marked change when switched on if the interferometer

¹However, the output beam is somewhat distorted, so probably more of this power was eventually lost in coupling the beam into the single-mode fiber than might have been without the AOM.

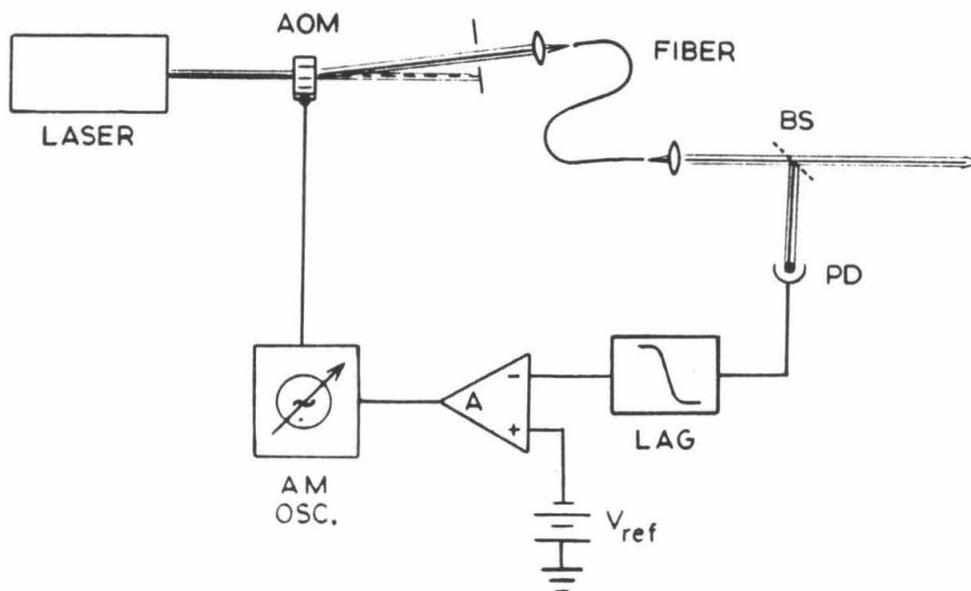


Figure 16: Optical and electronic semi-schematic of the acousto-optic intensity servo system. Light from the laser is diffracted by the acoustic field in the glass in proportion to the RF power provided by the amplitude-modulated oscillator. The AM control is provided by the amplified, filtered difference between the photocurrent falling on a detector and a stable DC bias (in this case a battery). While not designed for the high performance of the electrooptic system, it is extremely simple to operate and can be set up and torn down in under an hour. It is thus well suited to the diagnostic role usually assigned it.

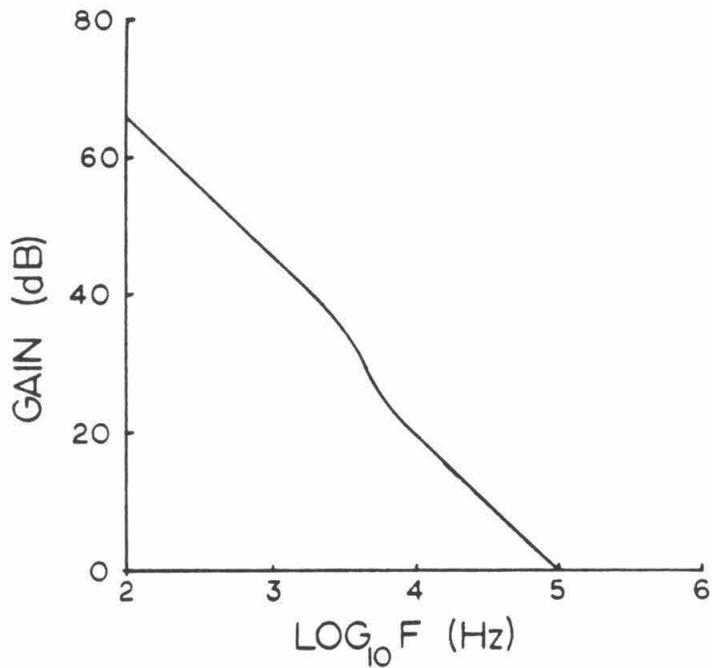


Figure 17: Calculated gain of the acoustooptic intensity servo as it was typically employed. The intensity fluctuations were attenuated by up 35 dB at 1 kHz, less than predicted on the basis of this gain; the discrepancy was not pursued, but may be due to noise in the commercial AOM driver.

noise is affected by intensity effects. It is also trivial to set up and employ, needing nothing but a standard laboratory preamp, a photodiode, and a battery (the AOM and driver are usually already installed to perform their isolation function). These features suit it ideally to quick diagnostic tests.

The AOM system is generally less prone to accidentally unlocking than the other servo; the extinction vs. voltage characteristic is *almost* monotonic, so it won't often find itself suddenly in a region with the wrong servo sign and pin itself. Unfortunately, near the maximum drive power the diffraction pattern becomes overmodulated, and higher orders of diffraction (seen as a fanlike array of beams leaving the AOM) rob power from the intended first order. If the servo accidentally enters this region, say during a momentary drop in the laser output, the slope will indeed just turn over and the servo will latch up at full RF power to the transducer. Rapid attention can readily fix this; if left for more than a few seconds, however, the high RF power will warm the AOM glass, which expands and gradually deflects the beam out of alignment with the optical fiber (further exacerbating the dearth of light on the sensor). After this, there is no recourse but to shut off the RF and wait up to 5 minutes for the AOM to cool. Efforts to sharply limit the RF power below the critical value were not entirely successful, partly because the commercial driver unit drifts considerably. Less frequent loss of servo lock is thus almost compensated by the inconvenience of recovering. The servo is usually set up with mean transmission between 50% and 75%, depending on how attentive one is prepared to be and how unstable the laser is. When inactive, the RF level is restored to the optimum transmission level, giving 90% throughput; while this is not much better than the electrooptic cell's optimum throughput, the AOM performs another function for the same price.

3.2 Laser Frequency Noise

In principle, the frequency ν of the laser light should cancel out in this type of differential measurement; a fluctuation in laser frequency produces the same phase signal in both cavities, and ought to be subtracted when the two phase measurements are subtracted. However, this would demand that the phase sensitivity of the two cavities be matched to an unrealistically high precision. Inevitable differences in the absolute length of the two arms and in the reflectivities of the mirrors severely limit the practicality of using unstabilized laser light.

Another main reason for narrowing the bandwidth of the laser light is the method chosen to multiply the number of times a photon measures the proper distance. In a Fabry–Perot cavity, the light must maintain coherence over the entire length of a photon’s lifetime in the cavity, since it is effectively interfered with incoming “fresh” light at every returning bounce on the input mirror. Recall that in looking for kilohertz–band gravitational waves one should store the light in each arm for a millisecond or so. The cavity will then have a optical bandwidth of order a kilohertz; just to get the light to enter the cavities efficiently, its spectrum must be at least this narrow².

For simplicity, in the 40 meter interferometer the light returning from the two arms is not interfered. Instead, the light returning after storage in each cavity is separately interfered with the incident laser light. One cavity (called the “primary” cavity because it usually has to be in resonance before the other one—the “secondary”—can resonate) acts as a precision reference. The laser frequency is servoed as tightly as possible to the resonance of the primary cavity. Then, to the extent that the laser’s own deviations are truly suppressed by the servo action, the light falling on the secondary cavity contains the imprint of

²The green light from an unstabilized commercial argon ion laser, at $\lambda = 514.5$ nm or $\nu = 5.83 \times 10^{14}$ Hz, typically has a spectral linewidth of several megahertz.

the primary's length. The difference between this light and the secondary cavity resonance will then reflect the difference in lengths of the two cavities. In this way, although the light from the cavities is never directly compared, the length of the primary is transferred through a "yardstick," the laser wavelength, which then measures the secondary independently.

In this arrangement the servo should suppress the laser's own frequency fluctuations to a residual level $\delta\nu/\nu \lesssim h$, since here there is no direct subtraction of the phases in the two arms. If the servo is not quite up to it, it is possible to get rid of a small residual amount of frequency noise by electronically subtracting the common laser-induced phase error signals from the two cavities. The laser error, which is common to the two cavities, should cancel out if the subtraction (commonly termed "phase subtraction") is done very carefully. Another way to get rid of a small persistent amount of frequency noise is to let the entire interferometer "breathe" with it; dilate or contract both arms in unison to match the dilation or contraction of the laser. Again, this "common mode" action (here equivalent to a uniform scaling of the entire instrument) will have no effect on the *differential* measurement being pursued. Both techniques were proposed [47] as ways to patch up the situation if the laser's frequency noise cannot be totally brought under servo control. Both are limited; again, difficulties arise in the need to match the signals from the two arms precisely. Phase subtraction has been used with success, but cannot in practice be relied upon for more laser noise suppression than a factor of 30; the second technique may be able to do better. In any case it remains desirable to make the laser frequency servo do the whole job if possible.

3.2.1 Frequency Noise in Argon Lasers

Two different models of laser have been used in the Caltech prototype; the small frame Lexel 95 (4 Watts nominal, all lines) and, more recently, the Coherent Innova 100–20 (20 Watts nominal, all lines). Both are used in single line, single mode operation with temperature-regulated solid intracavity etalons; under these conditions they can give about 1 watt and 6 watts, respectively. As with intensity noise, their frequency noise is unacceptable as they come from the factory. We will characterize the spectrum of frequency fluctuations with the quantities $\sqrt{S_\nu(f)}$, the RPSD of fluctuations in laser frequency ν (with confusing units Hz/ $\sqrt{\text{Hz}}$), and $\sqrt{S_{\delta\nu/\nu}(f)}$, the RPSD of fractional frequency fluctuations $\delta\nu/\nu$ (with units $\sqrt{1/\text{Hz}}$).

At low frequencies, the dominant contribution to this spectrum is mechanical vibration of the laser resonator. Much of the vibration is typically generated by the rush of cooling water around the plasma tube, which must remove about 99.98% of the energy supplied to the laser as waste heat (up to 30 kW for the Innova). This is demonstrated by the frequency noise spectra in Figure 18, taken at two different rates of water flow in the coolant loop. The vibration is apparently transmitted to the laser mirrors and etalon by the rigid connection of the tube to the resonator. (The separation of the laser mirrors defines the frequency of the laser.)

Many experimenters have broken this connection and placed the mirrors (and sometimes the etalon as well) in a separate, isolated resonator structure. This can reduce the frequency noise considerably, often with the added bonus of reducing output beam wiggle as well. Our Innova laser was rebuilt very rapidly, in an effort mainly directed at gaining more work space inside the laser cavity. The laser reflectors were removed and placed in standard optical mounts on blocks

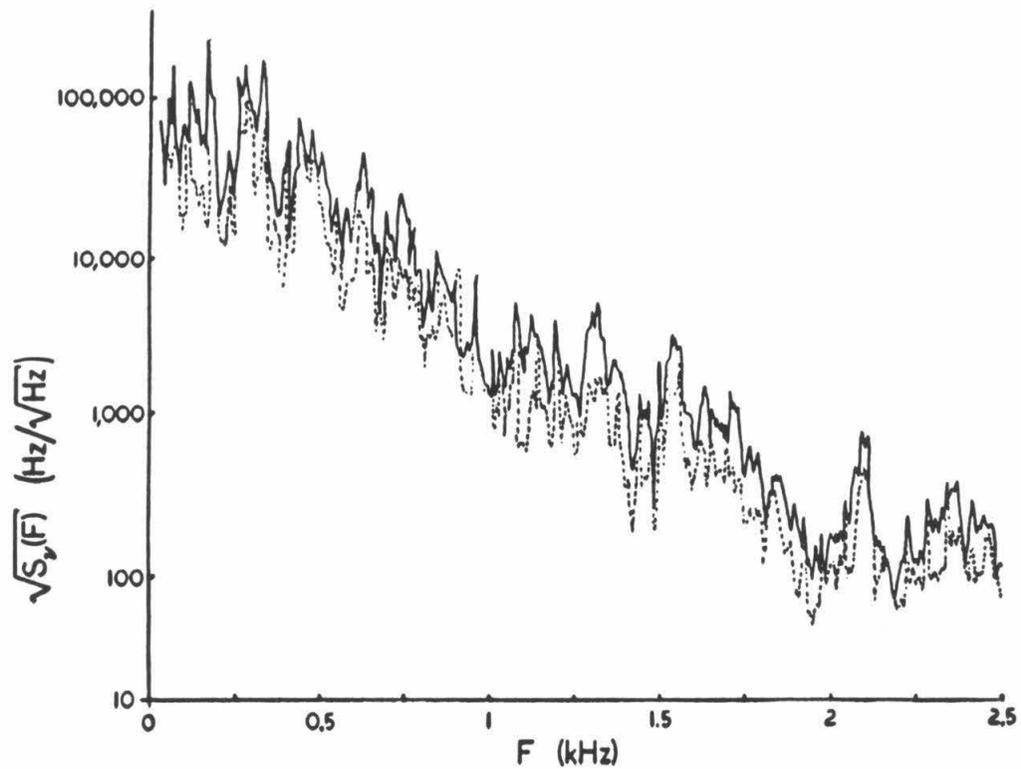


Figure 18: Frequency fluctuations of a Coherent Innova 100-20 argon ion laser used in the Caltech experiment. The laser was run at low current so that the coolant flow around the plasma tube could be temporarily reduced without overheating. The solid curve was measured with normal coolant flow; the broken curve, with the flow reduced by about 20%.

clamped to the optical bench, with the original frame (still holding the plasma tube) supported on a stack of metal plates and rubber. This almost certainly hurt the laser’s thermal stability, as the optical bench is ordinary stainless steel (the original resonator was made of super Invar, an ultra-low expansion steel). However, the rebuilding significantly reduced the laser’s frequency fluctuations (Figure 19). The thermal expansion of the laser with the optical table has not presented any serious difficulties; it appears that the table has enough thermal mass, and the laboratory air conditioning enough stability, to prevent excessive drift.

The improvement scored by rebuilding the laser is useful, but still leaves a great deal of frequency noise to contend with. For example, the goal of $\sqrt{S_h(f)} \lesssim 10^{-19}/\sqrt{\text{Hz}}$ requires a suppression of the noise at 1 kHz by a factor of order 10^7 , and considerably more at lower frequencies, below the improved level shown in Figure 19. That suppression factor is equal to the loop gain the laser frequency servo must achieve to clear this sensitivity goal.

3.2.2 The Servo

Most of the work to be described relates to the electronic amplifiers and the servo actuators used in the system. The remaining critical piece of the servo, the sensor, employs a technique developed specifically for this job [60]. It has proved so suitable that it has been widely adopted by researchers in other fields who require an extreme degree of laser stabilization. A heuristic explanation of the method is provided in Appendix B. Briefly, the phase of the laser is modulated at a radio frequency. Because of the storage effect of the optical cavities, the phase modulation is stripped off the light entering the cavities; it remains only on that which is promptly reflected by the input mirror. Thus, the relative phase of the input laser light and the stored light returning from the cavity is

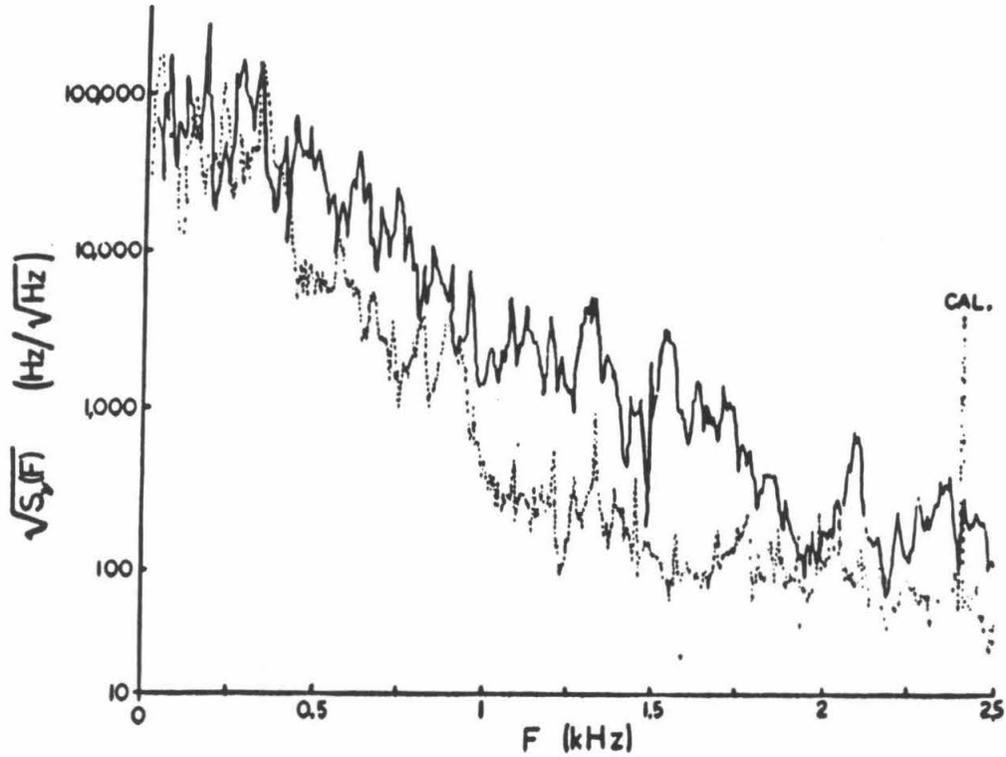


Figure 19: Frequency fluctuations of the Coherent Innova 100-20 argon ion laser used in the Caltech experiment, before (solid) and after (broken) the laser cavity mirrors were removed from the frame and mounted on the optical table. The water-cooled plasma tube (still mounted to the original laser frame, along with the intracavity etalon) was supported on a stack of metal plates and rubber to cut vibrations reaching the table.

modulated. By examining the component which varies synchronously with the imposed modulation in the interfered resultant of these two fields, the sign and magnitude of the mean phase difference can be derived. The photocurrent from a photodiode receiving the interfered light is multiplied by the phase modulating signal (usually in a double-balanced mixer) and integrated, yielding a pseudo-linear voltage representation of the optical phase error between the laser and the cavity. The cavity's storage of the light divides the behavior of the sensor into two regimes. For deviations which are rapid with respect to the storage time τ_e the cavity cannot maintain equilibrium with the laser, and the sensor reports the instantaneous *phase* error; for gradual deviations the cavity light can follow the laser, and the sensor instead reports the time derivative of the instantaneous phase error, i.e., the *frequency* error. The transition between the regimes is smooth, and has the simple character of an RC integrator with $RC = 2\tau_e$.

The fringe signal above has to be processed electronically and made to influence the laser frequency. The actuator most often used at Caltech has been a Pockels cell placed inside the laser cavity. In this application, the cell's refractive index changes in response to the electric field across its crystals; thus, the effective optical length of the laser cavity can be altered with an applied voltage. The principal advantage of this kind of actuator is that it is fast; providing one can supply the current to quickly charge the capacitance of the crystals, the effect on the light is virtually instantaneous. It unfortunately is only modestly effective; the Pockels cells we use (Gsänger PM-25) change their optical length by half a wavelength with an applied voltage of 1150 volts. In a laser two meters long, this gives a shift of 75 MHz in the optical frequency ($\delta\nu/\nu = -\delta l/l$, providing the laser keeps oscillating in the same longitudinal mode). Disadvantages include a loss of power over the “bare cavity” laser, nonlinear saturation effects

which limit the power still more, and a propensity for fatal damage to the crystals in the harsh intralaser environment (probably due to UV emission from the plasma discharge). This latter is especially serious given the high cost of Pockels cells. Alternative actuator types have been developed at Caltech and elsewhere [58,59,61,62].

The amplifier to drive the Pockels cell needs several typically incompatible characteristics. A brief summary of the requirements follows:

- Extremely low noise. The dominant noise contribution must be from the shot noise in the detected RF photocurrent.
- Large dynamic range. While acquiring a fringe the amplifier may handle transient signals 10^9 times larger than quiescent signals seen in normal operation; when locked, out-of-band noise at high frequencies can be 10^5 times larger than critical kilohertz-band signal components.
- High linearity. This also arises from the large range of signal amplitudes to be handled.
- High voltage capability. At low frequencies the system needs to provide high voltage to the Pockels cell to correct large excursions of the laser.
- High current capability. At high frequencies, considerable current is needed to charge the capacitance of the Pockels cell rapidly.
- High bandwidth. To achieve the extremely high loop gain desired at audio frequencies, the servo needs to apply corrections up to the megahertz band in order to maintain stability.

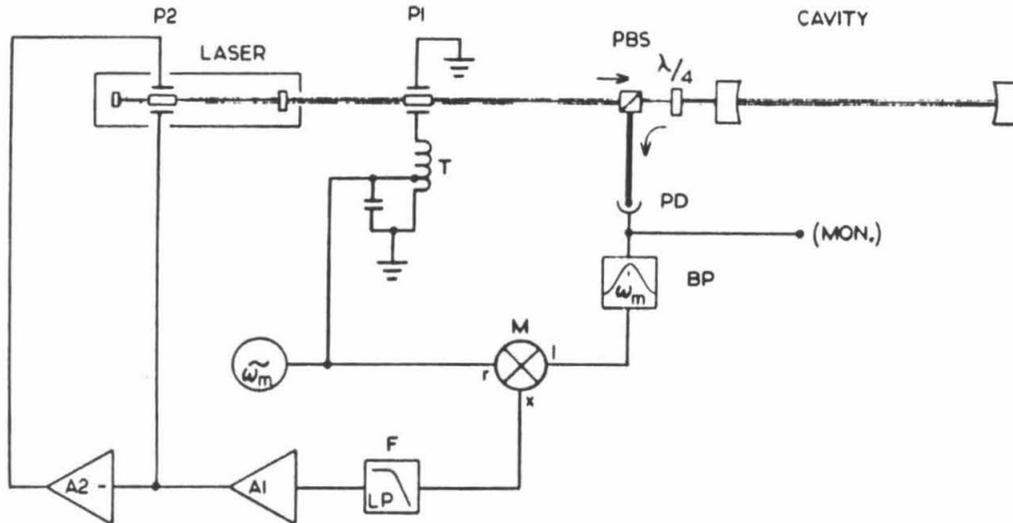


Figure 20: The laser frequency stabilization system (following [60]). Pockels cell P1, driven with approximately one watt of RF power through resonant step-up autotransformer T, phase modulates the laser light at RF frequency ω_m . The light returning from the cavity is diverted, by a directional isolator comprising a polarizing beamsplitter and quarter-wave plate, to a photodiode. The RF photocurrent is bandpass-filtered and demodulated by double-balanced mixer M against the modulating signal and lowpass filtered, leaving a signal which (for rapid fluctuations) is proportional to the phase error between the laser light and the light stored in the cavity (Appendix B). This signal is amplified by fast servo amplifier A1 and fed to one plate of intralaser Pockels cell P2, which corrects the laser frequency by adjusting the effective length of the laser resonator. At low frequencies the servo is assisted by high-voltage amplifier A2, which copes with large frequency excursions. Additional stages (not shown) assist at still lower frequencies by adjusting a piezoelectric translator on one of the laser mirrors, and by adjusting the cavity length to correct very slow drifts and pendulum motion. The DC component of the photocurrent is also monitored to check and adjust the alignment of the cavity to the laser beam; the “fringe contrast” $K \equiv (I_{max} - I_{min})/I_{max}$, where I_{min} and I_{max} are the detected photocurrents when the cavity is on and off resonance, respectively, is a maximum for optimal coupling of the laser mode to that defined by the cavity.

- Low propagation delay. High bandwidth in the sense of passing signals at all frequencies is not sufficient; the corrections must reach the laser immediately .

The unity-gain bandwidth was targeted to be at least 1 MHz. To achieve this goal, the total delay around the loop needed to be kept under 250 ns; this was just attainable with available integrated circuits after the amplifier delay is added to delays from 20 odd meters of beam path and coaxial cable, plus delays in photodiodes, preamps, etc.

In a simple, unconditionally stable servo (one which is stable irrespective of the overall frequency-independent gain) this unity gain frequency would give a gain of only a thousand at 1 kHz. Constraining the gain to be 10^7 at 1 kHz and ≤ 1 at 1 MHz requires some precipitous filtering in between. This filtering must invariably be done at the cost of phase shift³; the phase shift is “unwound” by differentiating the signal a suitable number of times just below the unity gain frequency to satisfy the stability criteria, but this locks the unity gain frequency into the design. Reduced gain levels, as from variations in laser power or twiddled knobs, can bring the unity-gain frequency into a regime where the phase shift is greater than π , just as increased gain will. The servo is conditionally stable. For safety, the servo is designed to remain stable over a 10 dB range in overall gain, to allow for misadjustment or fluctuations in cavity alignment and laser power.

Since the intralaser Pockels cell adjusts the laser *frequency* and the cavity provides a signal proportional to the laser *phase* (for fast corrections) the system already implicitly contains the single integration needed for the simple (unconditional) servo. Measurements of the “optical” transfer function (the ratio of mixer

³It is true that filters can be constructed with sharp amplitude response and no effect on phase, but to do so they must be able to peek at the future development of the signal (i.e., they are acausal). To get this peek the entire signal must be delayed, prohibiting use in a real-time servo system.

response to laser Pockels cell excitation) under more or less typical conditions of light power, fringe contrast, modulation depth, etc. yield the value

$$\frac{\partial V_{mixer}}{\partial V_{pc}} \approx 150 \left(1 + \frac{f}{500 \text{ Hz}} \right)^{-1} \frac{\text{V}}{\text{V}} \quad (18)$$

when the laser is locked to the center of a fringe. Thus the electronics must provide an amplification of 13 or so at 1 MHz to make the overall loop gain (the product of optical and electronic contributions) unity at that frequency. This value, and the requirements listed above, define the servo amplifier's transfer function almost entirely (Figure 22).

3.2.3 The Electronics

Implementing that transfer function in a manner consistent with the formidable performance constraints we listed is difficult. The design was evolved from servo amplifiers developed for dye laser stabilization by Hall et al. [63]. A key innovation of those designs satisfied the diverse requirements by splitting the signal path internally into several complementary signal paths, each of which is optimized for a different job, and then adding the effects together. One principal advantage over an ordinary cascade of gain stages (for the high gain required at low frequencies) is that the propagation delays of successive stages do not add up; high-frequency components are "bypassed" around the sluggish low-frequency gain elements, leading to the name "bypass topology" for this kind of circuit. Further bypassing is used to satisfy the input noise requirements; a low-noise, but slow, preamp can be made to dominate the overall noise figure at low frequencies, but fade out above a few kilohertz. Still more bypassing allows use of a slow output stage with high voltage capability on one plate of the Pockels cell while a faster, high current driver (with an inverted sign) drives the opposite plate. The electric field in the cell, and thus its refractive index, depends only

upon the difference between the plate voltages.

A semi-schematic diagram showing our specific variant of the bypass topology, along with the characteristics of the various stages, is presented in Figure 21. It is important to note that the overall gain is the sum of gains of various elements, each of which dominates the others in a different frequency regime. The “crossover” frequencies, at which two different signal paths have roughly equal effect on the output, are particularly troublesome. If the relative phase between two added signals is too large (more than about $3\pi/4$) at the crossover, the transition will create a virtual antiresonance as the vector sum dips below the magnitude of either component taken alone. Like a true resonance this can have dire consequences for the servo’s delicate balance of stability. As a result the crossover frequencies and relative phases of the component signal paths have to be chosen carefully.

Another problem is associated not with the steady-state performance of the amplifier, but with what it does at startup, when first attempting to acquire a cavity fringe. The D.C. gain of the circuit is quite large—around 10^{11} —so the output is effectively always saturated hard when the feedback loop is open (just the thermal noise at the input would produce many kilovolts at the output, if the circuit could provide it). The behavior in the saturated state is important when one considers that in this near-dead condition the servo must somehow lock to a cavity fringe as it whizzes by, *before* the amplifier can enter linear operation and start doing its job. This property is arranged largely by trial and error, but there is a plan. Strategically placed diode limiters [63] are used to “design” the saturated state; the fastest element of the servo comes out of saturation first and jumps in with a crude, but stable, loop gain characteristic, just enough to catch a fringe briefly; the other stages then successively come on line to fill out the gain curve.

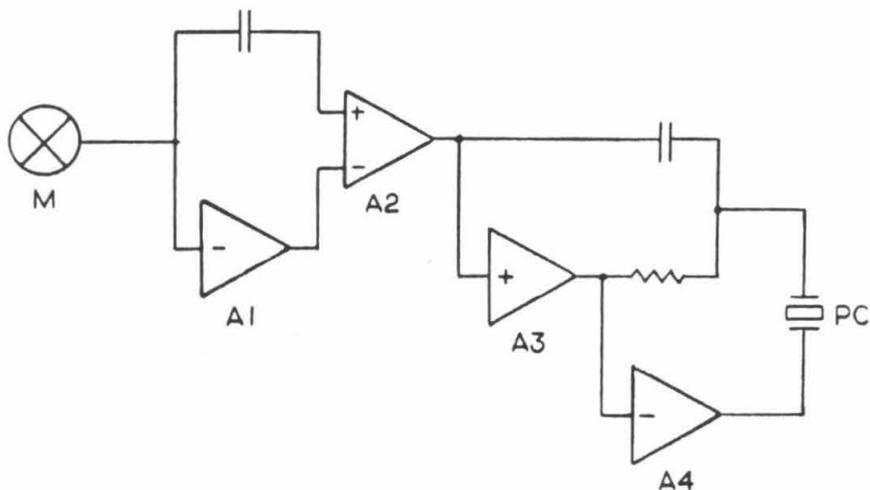


Figure 21: The composite laser servo amplifier is broken down into specialized subunits, each of which handles particular tasks best. Complimentary signal paths are blended together to generate the overall transfer function of the unit, shown in Figure 22, while still meeting stringent noise, time delay, linearity, dynamic range, and output drive requirements. The fastest correction signals are capacitively bypassed to high speed, high current, low delay amplifier A2 from the output of mixer M. They are then applied directly to the upper plate of the Pockels cell PC. Low noise amplifier A1 also receives the mixer output; its contribution to A2's input dominates below 30 kHz, suppressing the noise of A2 and providing gain. A3 provides additional gain and current drive to the upper plate of the cell between 30 and 300 kHz, and also boosts the gain of signals fed to high-voltage amplifier A4, which dominates the drive across the pockels cell below 30 kHz. Not shown are additional bypass stages which adjust a piezo mirror on the laser (dominating below 300 Hz) and drive the wire pushing piezos on the cavity test mass suspensions (effective below 10 Hz).

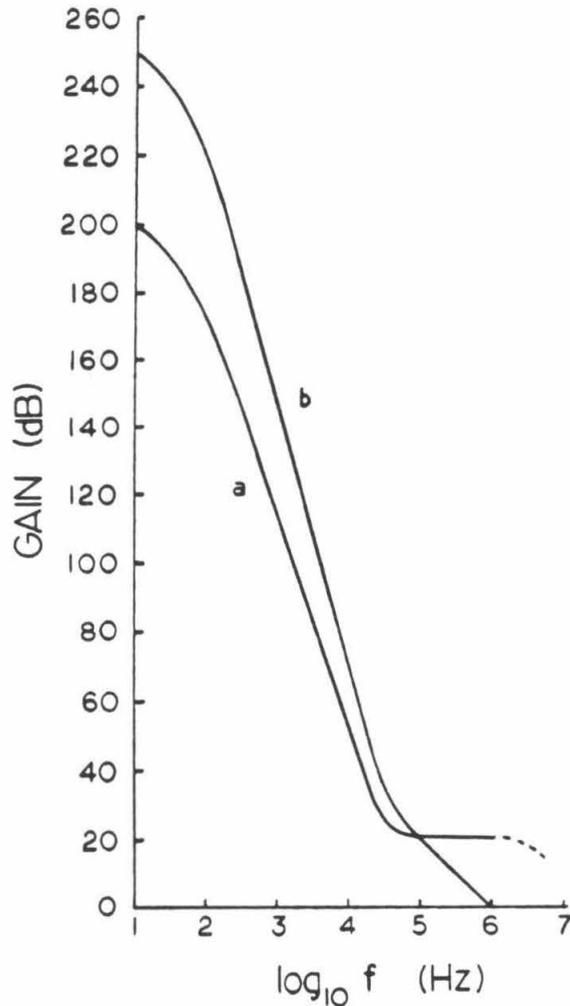


Figure 22: Calculated Bode plot for the stabilization servo loop. The gain of the electronic amplifier alone (a) is shown along with the total combined loop gain of the amplifier, RF and optical systems (b). The curve is highly constrained by the requirements of high gain at 1 kHz ($\gtrsim 10^7$) and stable phase ($\leq \pi$) at the unity-gain frequency of 1 MHz. Practical difficulties hamper efforts to directly measure such large gain figures; however, the ratio of unstabilized laser frequency fluctuations (Figures 19) to the measured upper limit on stabilized fluctuations (Figure 26) is consistent with this gain and inconsistent with 10 dB less.

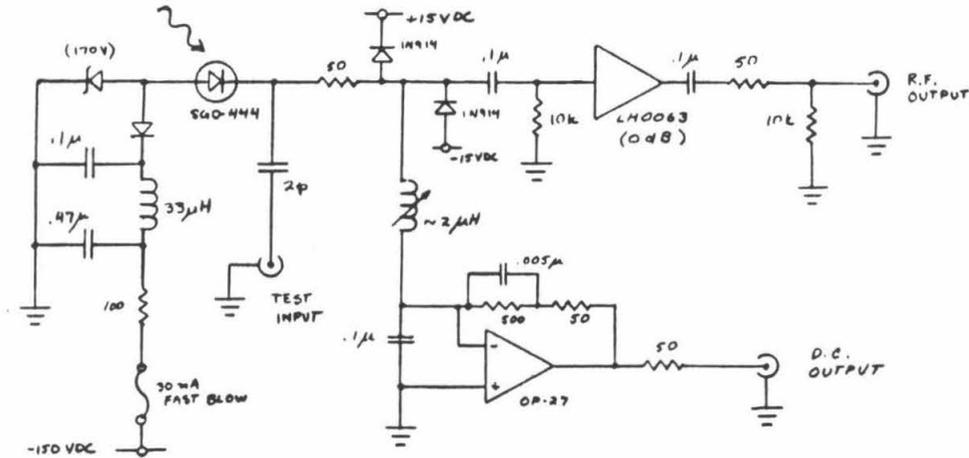


Figure 23: Circuit diagram of the RF photodiode units. Used as “front ends” for both primary and secondary cavities, they detect and separate the component in the reflected cavity light containing the optical phase information, which oscillates at 12 MHz. The SGD-444 silicon photodiode is reverse-biased at -150 volts DC. Its junction capacitance is made to resonate with the adjustable inductor at the 12 MHz modulation frequency; the Q is deliberately diminished by the 50Ω series resistor to maintain high servo bandwidth. The voltage induced by the RF photocurrent is buffered by an LH0063 current amplifier and passed to the mixer. Additional RF gain stages may be added between the R.F. output and the mixer if needed. The RF noise of the circuit is exceeded by photon shot noise in as little as $500\ \mu\text{W}$ of incident power. The quantum efficiency of the photodetector (including window reflections) is about 45% for the 514.5 nm wavelength. The low frequency current-to-voltage converter (bottom) monitors the D.C. photocurrent and provides a readout of the fringe contrast, used to optimize cavity alignment.

The circuit, as it existed in late 1987, is shown in Figure 24. Additional circuitry, not shown, is used to compensate slow drifts of the laser and/or cavities by adjusting a piezoelectric stack on one laser mirror and by adjusting the wire suspension points of the test masses (via the same wirepushing piezos driven by the shark detectors). These measures typically allow the laser to stay locked to a cavity fringe for tens of minutes without interruption, depending on environmental factors; if it loses lock, usually from excessive drifts⁴ exceeding the dynamic range of those adjustments, it finds a new fringe and relocks in a fraction of a second.

3.2.4 Performance

The servo has evolved considerably since its introduction in May 1986, in the effort to keep laser frequency noise one step ahead of the system’s other limitations. The effort has been largely successful; tests for frequency noise (for example, deliberate enhancement of the frequency noise by a known factor) have *usually* vindicated it. When such tests did show something, it was often anomalous and therefore interesting. One challenging problem was a noise source in the servo amplifier which only appeared when the loop was closed, and disappeared at other times. It was traced⁵ to asymmetrical slew rate limiting on stray RF signals reaching the low noise OP-37 input preamp; carefully isolating the inputs of this stage from out-of-band signals, and readjusting the rest of the circuit to maintain the same frequency response, cured the problem.

Another noise source, which is still under intensive scrutiny, seems to have arisen not in the amplifier but in the photodiodes. The laser produces excess frequency and amplitude noise at around 1 MHz which ordinarily should have

⁴Or due to a faulty laser autofill circuit, which at inopportune moments opens a gas valve to fill the tube, and then changes its mind and closes it again.

⁵With guidance from Jeff Harman.

no effect. A viable hypothesis is that these signals may interact with nonlinearities in the photodiodes themselves to produce noise at the 12 MHz modulation frequency, contaminating the phase measurements in both cavities. This could explain the repeated failure of increased light power to reduce the noise level, as would have been expected of photon shot noise. The frequency and intensity servos have insufficient gain at high frequencies to suppress these fluctuations, so a third smaller optical cavity was introduced in series with the laser to act as a passive optical–frequency bandpass filter. Again, because of its storage of the light, the cavity averages away rapid phase or intensity fluctuations. The initial “mode cleaning” cavity, which was 40 cm long and was equipped with supermirrors, had an optical bandwidth of about 80 kHz. Getting the light to resonate simultaneously in this cavity and the 40 meter cavity was not a trivial task. Briefly, the mode cleaner had to resonate to transmit light to the interferometer and permit the laser to be stabilized, but the light needed to be pre-stabilized to resonate in the mode cleaner, leading to a “chicken and egg” problem at startup. It was solved⁶ by inventing an intermediate state in which the laser was simultaneously locked to both mode cleaner and 40 meter cavities. A further problem was imposed by the narrowness of the mode cleaner bandwidth itself; since this cavity was inside the feedback loop, its filtering action prevented the laser servo from operating with the required 1 MHz bandwidth. The solution was to divert the fastest frequency corrections applied to the laser, integrate them, and apply them as *phase* corrections to an additional Pockels cell placed after the mode cleaner (Figure 25).

This system allowed the mode cleaner to work within the high-gain stabilization loop. The result was a 7 dB improvement in the interferometer noise spectrum, which is shown expressed in terms of an equivalent upper limit to

⁶With the indispensable help of Harry Ward, who was then visiting from Glasgow.

laser frequency noise in Figure 26. Work now in progress centers on reducing the bandwidth of the mode cleaning cavity to see if better filtering of the light will further reduce the noise. This has required drastic alterations in the stabilization system (including the removal of the unappealing "chicken and egg" problem). It should be noted that the upper limit shown in Figure 26 implies the laser light is exhibiting fractional frequency fluctuations of at most 4 parts in 10^{18} over millisecond intervals.

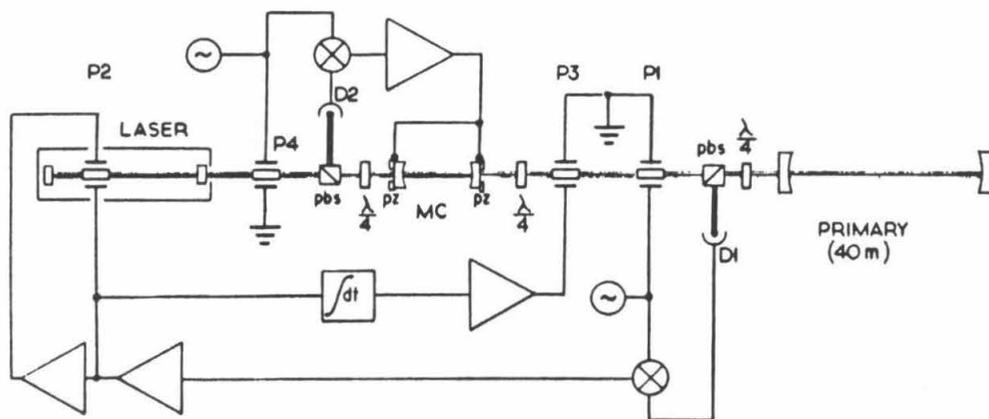


Figure 25: Schematic arrangement of the laser stabilization loop after modification to incorporate mode cleaning cavity MC as a passive optical narrowband filter. A secondary independent servo loop (using modulating pockels cell P4 and detector D2) matches the mode cleaner's resonance to the laser light by use of piezoelectric mirror translators (pz). This keeps the light resonant and allows transmission to the primary cavity. Cell P1 re-applies the phase modulation (which has been stripped off by the mode cleaner's filtering action); the laser servo operates as in Figure 20, except that fast laser corrections, which are also attenuated by the mode cleaner, are diverted, integrated and applied to post-mode cleaner Pockels cell P3 as phase corrections. This restores the original servo bandwidth and gain. Additional apparatus required for startup has been omitted for clarity.

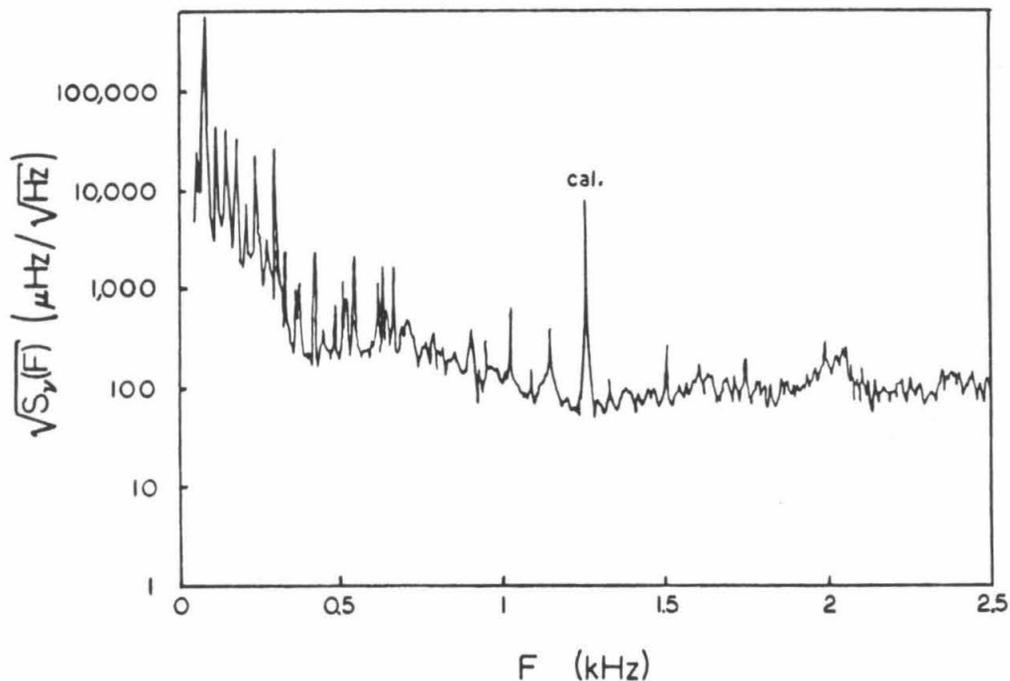


Figure 26: Measured upper limit to the RPSD of frequency fluctuations in the stabilized laser, after inclusion of the mode cleaning cavity. Compare with the laser's intrinsic noise, Figure 19. Dividing by the mean laser frequency $\nu = 5.8 \times 10^{14}$ Hz gives the spectrum in terms of fractional laser frequency fluctuations; since no electronic phase subtraction has been used, this will equal the RPSD of the antenna strain h (Figure 6).

3.3 Locking the Secondary Cavity

With the laser locked to the primary cavity, we need to resonate the light with the secondary in order to obtain the difference in their lengths. The relative optical phase is sensed by an identical system to that used on the primary, in fact using the same imposed modulation. The secondary servo system, however, is otherwise quite different, since in this case it is the cavity which must be servoed to the laser light. For the observations described in Chapter 4, the cavity length was adjusted by a piezoelectric transducer sandwiched between the test mass and its mirror. The transducer was driven by a high-voltage amplifier through a network which allowed a unity-gain bandwidth of approximately 5 kHz (Figure 27). This bandwidth was limited by resonances in the masses, which (for the brass cylinders then in use) appeared between 13 and 16 kHz. The piezo had several advantages, one of which was the ability to provide large motion (up to a micron) in less than a millisecond. Although the laser was highly stabilized by the primary loop, the low-frequency swinging of the masses often made the fringes fly by quickly. Swinging velocities of order 5 microns/second are typical, but the fringe width is only about 0.1 angstrom; a rapid actuator was helpful in stopping the mirror surface inside this range.

Unfortunately, the piezos had to be abandoned when the move was made to optically-contacted fused silica masses. In their place, small rare-earth permanent magnets were bonded onto each end mass, and copper coils were arranged to push and pull the mass as a whole. Since the new system actuates a force on the mass rather than a direct displacement of the mirror, it is difficult to provide enough impulsive current to halt the entire mass within a single fringe width. However, the electromagnet's superior dynamic range at low frequencies makes the magnetic servo more robust once locked. The new masses have much higher

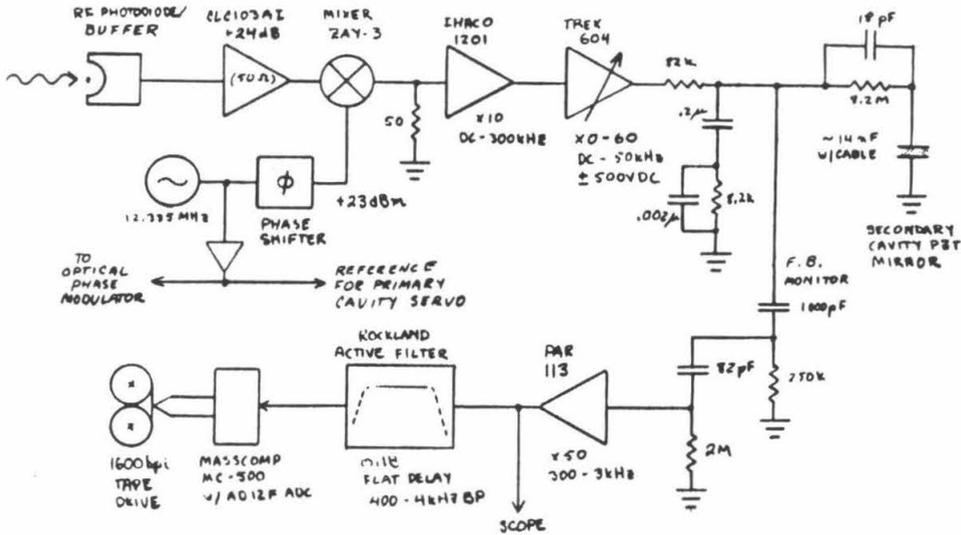


Figure 27: Schematic diagram of the secondary cavity servo loop, showing the high-voltage amplifier and filter network used to drive the piezoelectric mirror translator which keeps the cavity in resonance. Also shown is the high-impedance feedback monitor; when locked at high servo gain, the feedback voltage represents the correction applied to the secondary cavity, and thus the strain in the antenna. The signal is strongly filtered to remove large low-frequency components and is analyzed and/or recorded.

intrinsic Q , and this limits the usable servo bandwidth to about 3 kHz; if fed signals at higher frequencies, the internal vibration of the mass at its longitudinal resonant frequency of 30 kHz is excited and grows in amplitude.

As shown in Figure 28, the feedback signal applied to correct the length of the secondary cavity, whether piezo voltage or magnetic coil current, represents the difference in the cavity lengths, and is therefore proportional to the strain h . This signal is the "output," and is recorded and/or analyzed. A spectrum of this signal is presented in Figure 29; this was the antenna noise level as the observing run described in the next chapter commenced.

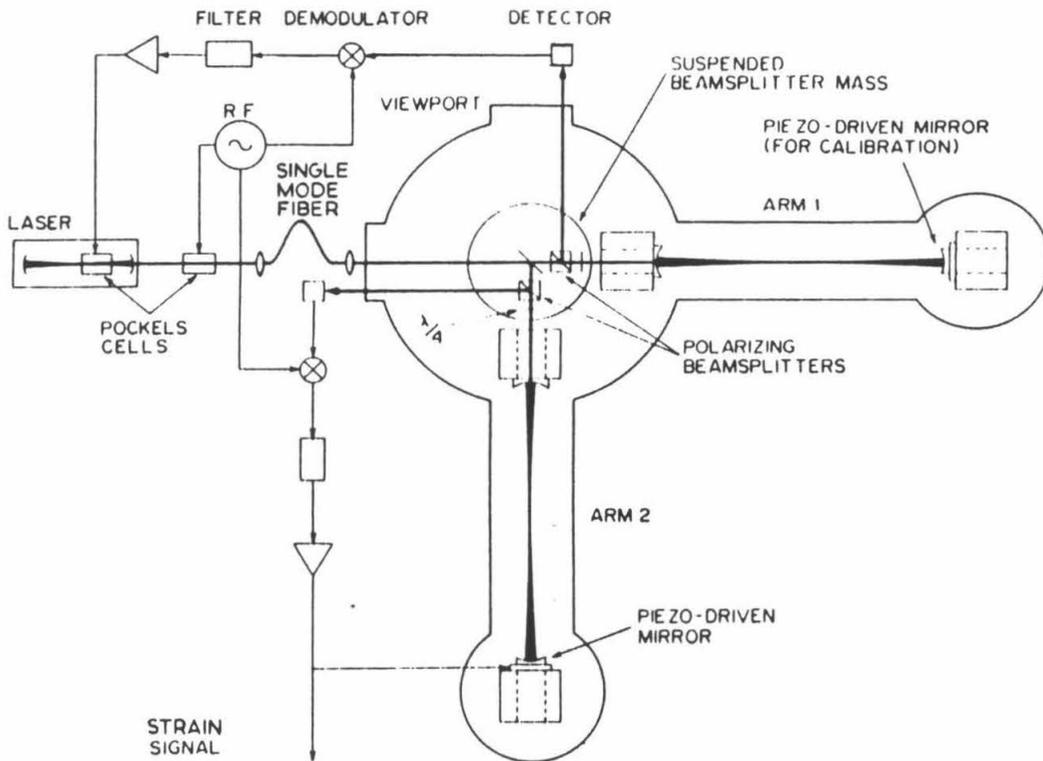


Figure 28: Operational schematic of the antenna with both arms servoed. The secondary cavity resonance is sensed in the same way as the primary; the resulting signal is fed back to the end test mass to keep the secondary locked, using either a piezoelectric mirror transducer (pre-May 1987; as shown) or a pair of electromagnetic coils forcing small permanent magnets bonded to the mass. The secondary feedback is proportional to the difference in the cavity lengths, i.e. the strain. A symmetric transducer is provided on the primary cavity end mass to allow calibration of the entire instrument. The single-mode optical fiber is used to suppress geometric fluctuations of the laser beam, which have been shown to interfere with the phase measurements. The mode-cleaning cavity now in use (Figure 25) has a similar mitigating effect on geometric fluctuations, so the fiber will probably be eliminated in the future.

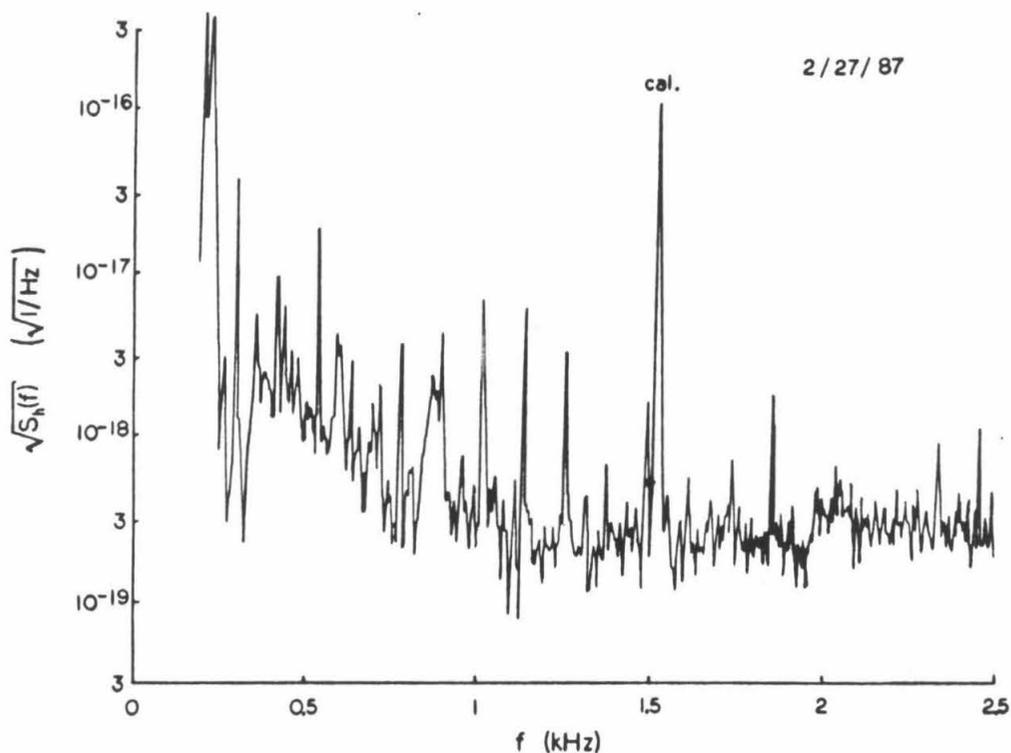


Figure 29: Calibrated spectrum derived from the feedback to the secondary cavity piezoelectric mirror, taken just prior to the observing run. The feedback voltage has been converted to strain units by use of a series of calibration signals applied to the primary cavity piezo at various frequencies. One such calibration signal is shown at 1.5 kHz. The numerous narrow spikes are caused by line frequency interference. It is now believed that this interference may have been picked up directly on the wires leading to the piezo mirrors; nearly all the spikes disappeared with the switch to a magnetic transducer system (cf. Figure 6). Some peaks at lower frequencies and the broad bump just below 1 kHz may be due to resonances in the beamsplitter (Dewey) mass. The very large features near 200 Hz arise from a rotational resonance of the Huey control block, excited by the orientation servos. The noise significantly exceeds the measured contribution of the shot noise (at the effective light power of 10 mW) everywhere but in the narrow band 1.5—2 kHz, where it is consistent with shot noise.

Chapter 4

The Observations

4.1 State of the Interferometer

In late February 1987 the Caltech interferometer was open to air undergoing adjustments. The persistent problem of mirror distortion induced by rigid bonding of the mirror substrates to the metal masses was forcing operation with very poor alignment, contrived to discourage the appearance of high-order modes in the cavities. This deliberately poor alignment required some redirection of global orientation servo beams. The noise spectrum fluctuated and was often higher than normal, but occasionally dipped to approximately $4 \times 10^{-19}/\sqrt{\text{Hz}}$ near 1.5 kHz.

The discovery of Supernova 1987a in the Large Magellanic Cloud was announced on 24 February. No gravitational wave interferometer was operating at the projected time of the stellar collapse; furthermore, all of the cryogenic bar detectors were also off the air and far from operating temperature. Rough estimates of the possible gravitational wave flux emitted by the collapse [65] came in just below the sensitivity of the best current instruments, anyway. Likely models for the remnant object (presumably a neutron star or black hole) also

seemed unlikely to produce a detectable flux. Nevertheless, the Caltech prototype interferometer and the interferometer in Glasgow, beset as they were with technical problems, were the among most sensitive broadband antennas to date; the closeness of the supernova (about 50 kiloparsecs) led both groups, along with the MIT interferometer group, to press the antennas into observing service. All three instruments were at the time in various stages of disassembly; the Garching interferometer, which had been the most sensitive interferometer up to this time, was unfortunately so far disassembled that the German group was unable to join the effort. At MIT, Glasgow and Caltech, the work began to restore the detectors to operation as soon as possible, with the fear that any source energetic enough to be detectable might fade rapidly.

Our prototype was quickly closed up and brought to an operable state; the noise level (a selected “good” sample of which was presented in Figure 29) and the fringe contrast fluctuated, and the cavities were more prone than usual to falling out of lock (as well as more resistant to reacquiring), but the performance was deemed the best possible without unacceptable delays. This work was completed in only three days; unfortunately, the other groups were beset with major reconstruction tasks and took about a week longer. Simultaneously an improvised data-acquisition timing system (to be described in Appendix D) was assembled, and the data acquisition computer was programmed with a rudimentary algorithm to digitize signals from the interferometer and four auxiliary systems and write the results on magnetic tape.

4.2 Data Runs

Six tapes were hastily recorded at Caltech starting at 2 a.m. Greenwich time on 27 February, 1987. Problems with the tape drive led us to experiment with

sampling rates and buffer sizes, and about one hour of tape was recorded at varying bandwidths and with varying complements of auxiliary veto signals over a five hour period. Efforts to fix the tape drive, improve the recording process and stabilize the interferometer alignment consumed the next 41 hours, after which about 1.5 hours of data spanning an 8 hour baseline were taken. Sometime around 1 hour into this second run the clock was accidentally offset by an unknown number of milliseconds, breaking the data into two continuous-phase segments (before and after the mishap). Should future analysis require it, the offset may be recoverable from the WWV radio broadcast track recorded on the tapes.

On 5 March the Glasgow 10 meter interferometer was ready and an initial run was made in two-way coincidence; both interferometers were unusually prone to falling out of lock, so after about half an hour of low-duty-cycle observing the experiment was deferred until 8 March. At this time the MIT 1.5 meter interferometer came on line, and we took data first in two-way coincidence with MIT and then in three-way coincidence with both MIT and Glasgow on 8 and 9 March. On 11 March all three groups once again recorded data simultaneously.

This concluded the first coincident observing run by interferometric gravitational wave antennas. The three interferometers were on the air a comparatively short time; the 70 tapes recorded at Caltech represent only about 8 hours of data, of which the interferometer was locked for 5 or so, and the actual three-way coincidence measurements lasted only 2.5 hours (the actual "live" time with all three instruments locked and actively recording has not yet been computed, but this is of course an upper limit). Attention has therefore been focused on looking for gravitational wave sources which might be directly associated with the supernova's aftermath, rather than sources expected to occur more or less uniformly in space and time.

4.3 Potentialities and Limitations of the Data

Since these observations were recorded, we were enabled, in a sense, to “plan our experiment,” to consider what astrophysical phenomena we might look for or seek to rule out. The general-purpose nature of the broadband antenna and the flexibility of wideband digital recording allows one to be fairly nonspecific about what one wishes to look for. Within well-defined limits, all the information is there on the tapes. Future sophistication and specialization of interferometric receivers, as well as the crush of continuous volumes of recorded data, will force more real-time interaction, but for now it is sufficient to peruse this hasty snapshot of the universe at our leisure.

A useful framework for deciding more or less objectively what kinds of gravitational waves the data set is best suited to finding or ruling out is organized by posing the following questions:

- Can the source considered be associated in time and space with the recent supernova?
- If *not* associated with the supernova, then under the assumption of approximate ergodicity in the universe, would a detection by this search be consistent (in a statistical sense) with null results from previous observations ?
- Can a source detectable with high confidence in these observations be reconciled with widely held beliefs about astrophysical generation of gravitational waves?

These questions were addressed for various potential sources in an attempt to establish rough guidelines for the analysis. The considerable gains achievable by combining worldwide data from the three detectors were also estimated, although

at this writing the other groups have not yet analyzed their data far enough to predict coincidence sensitivities rigorously.

An examination of burst and compact binary coalescence gravitational wave sources, described in Appendix C, arrived at the following conclusions. First, it seems unlikely that such signals would be associated with the supernova. Second, since the Stanford bar detector had observed for a much longer period in 1982 [66] with null results at equal or somewhat better sensitivity to these classes of signals, anything observable in our short data set would have to arise from a statistically improbable large event. Finally, the detectability threshold of the interferometer for the coalescing binary signal implies so small a range (of order 20 parsecs for pairs of neutron stars [37]) that it is improbable for such an event to have happened inside this volume during the few hours spanned by our observations. Although such definite predicted strengths are unavailable for short bursts, a detectable burst in this short period would also need to be an improbable statistical accident to be consistent with cherished beliefs about the energy reserves of our galaxy [1].

4.4 Periodic Gravitational Waves

We therefore turn to the prospect of looking for periodic gravitational waves in the data. A lack of previous searches for continuous periodic signals at the strain sensitivity achieved during these runs was brought to our attention by Bernard Schutz [67]. In addition, the possibility that the supernova formed a neutron star (confirmed by the detection of a pulse of neutrinos from the event [3]) and that the new star may be rotating rapidly (still just speculation) may provide a mechanism for the remnant to emit strong gravitational waves. For example, the newborn might not at first be perfectly axisymmetric. Here we might reasonably

expect some gravitational radiation to be emitted directly by the remnant of the supernova. The question remains; could such radiation be detected?

In looking for a truly monochromatic source radiating at some unknown frequency, an optimal search procedure will improve the signal-to-noise ratio as the square root of the observing time, providing the detector noise is itself random and aperiodic. Practical difficulties with coherently piecing together data having large gaps in the time record (discussed by Jeff Livas in [54]) make it considerably more tractable to consider using something less than the entire baseline of the data run or, as Livas did, using an incoherent combination of shorter coherent spectra to get something in between in sensitivity. An added complication is the Doppler shift due to the velocity of the detector as the earth rotates and revolves around the sun, as well as the amplitude modulation due to the rotation of the antenna pattern (Equation 8). For integrations longer than about half an hour the Doppler shift of a kilohertz-band wave can smear the signal power out over enough frequencies to make it undetectable. As previous periodicity-seekers will attest, even with the necessary computing and data handling power, removing these motion effects can be a formidable task.

Furthermore, a fair guess which will later be justified is that the assumption of random aperiodic noise in the detectors is very poor, and the smallest signal detectable with confidence is therefore not so much limited by the broadband random noise as by the presence of local *periodic* interference. Although powerful methods for rejecting local periodicities in a single detector have been developed [54, § 4.3], comparison with the concurrent data from the other detectors would be considerably more general in discriminating against such interference. This justifies restraint in devoting resources to extensive single-detector analysis.

A reasonable initial compromise is to analyze the data from single tapes or fractions of tapes, segments without major time gaps and (at seven minutes per

tape) having negligible Doppler and A.M. effects. The noise spectrum of our interferometer was not white, but a sizable fraction of the spectrum was near $\sqrt{S_h(f)} \approx 3 \times 10^{-19} / \sqrt{\text{Hz}}$. This implies a sinusoidal strain of 3×10^{-20} in the antenna would just equal the noise level after coherent averaging for 100 seconds. A further statistical confidence factor of from three to five, if the noise were *truly* random and aperiodic, as well as allowance for the non-optimal orientation of the antenna, are required to translate this mean noise level into a minimum detectable astrophysical strain. With these factors included the antenna might be able to detect with some confidence monochromatic waves coming from the supernova’s direction of amplitude $h_{min} \gtrsim 2 \times 10^{-19}$ at one kilohertz.

The strain amplitude expected from a spinning neutron star at distance r with rotational period¹ $2/f_{GW}$ and having a “gravitational ellipticity” ϵ , after averaging over the possible orientations of the spin axis in space, is given by Thorne as

$$h(f_{GW}) \approx 1.5 \times 10^{-20} \epsilon \left(\frac{I}{10^{45} \text{ g cm}^2} \right) \left(\frac{f_{GW}}{1 \text{ kHz}} \right)^2 \left(\frac{50 \text{ kpc}}{r} \right) \quad (19)$$

with ϵ being the disparity between the projections of the star’s quadrupole moment along each of the equatorial principal axes, normalized by the star’s moment of inertia I [1]. This assumes that the rotation axis is also a principal axis; if it is not, one can also obtain radiation at or near the rotation frequency as well as at twice the rotation frequency [68]. In general the radiation can also have higher harmonics, determined by the complexity of the deviations from symmetry.

The size of ϵ is open to broad speculation. Upper limits can be formed for those neutron stars which have been observed as pulsars, based on attributing the observed slowing of the pulse rates to gravitational radiation drag. Unless the remnant of 1987a becomes visible as a pulsar, we are unable to estimate its

¹The waves have half this period, as the quadrupole moment takes on the same values twice in each rotation.

ellipticity in this way.

Another upper limit is placed by assuming the neutron star was born with infinite (or very large) angular momentum and spun down entirely by emission of gravitational radiation since, with no further energy input. This condition implies the relation $\dot{P}/P \approx 1/\tau_{ssn}$ where τ_{ssn} is the elapsed time since the supernova collapse and P is the star’s rotational period. That “maximum” rate of power emission gives

$$\begin{aligned} \dot{E}_{GW} &= -\dot{E}_{rot} = -\frac{E_{rot}}{\tau_{ssn}} \\ &= -\frac{1}{2\tau_{ssn}} I \Omega_{rot}^2 \\ &\approx 8.2 \times 10^{45} \left(\frac{I}{10^{45} \text{ g cm}^2} \right) \left(\frac{f_{GW}}{1 \text{ kHz}} \right)^2 \left(\frac{1 \text{ week}}{\tau_{ssn}} \right) \frac{\text{ergs}}{\text{second}}. \end{aligned} \quad (20)$$

The strain corresponding to this maximum power, when averaged over possible orientations of the rotation axis, will be

$$h \approx 1.0 \times 10^{-24} \left(\frac{I}{10^{45} \text{ g cm}^2} \right)^{\frac{1}{2}} \left(\frac{1 \text{ week}}{\tau_{ssn}} \right)^{\frac{1}{2}} \quad (21)$$

independent of the frequency. The emission frequency would continually decrease, since $\dot{f}_{GW}/f_{GW} = -1/\tau_{ssn}$; this has important consequences for some proposed detection strategies [69]. At the one-week mark, kilohertz waves would slip in phase by half a cycle in only 25 seconds with this power output. Special techniques would be needed to follow the slewing signal frequency if coherent integrations longer than this are to be used to pick it out of noise.

The above estimates, while disappointing in view of the sensitivity figures quoted above, presume that there is no power source other than the neutron star’s rotational energy. Mass transfer from nearby objects has been invoked as a likely mechanism by which the very rapid millisecond pulsars acquired their huge angular momenta. A companion star or other “battery” driving our candidate

source could provide more radiative gravitational wave power. The energy flux at earth corresponding to a sinusoidal strain h at frequency f_{GW} is given by [1, Equation 11]

$$\begin{aligned} \frac{dE_{GW}}{d^2x dt} &\approx \frac{\pi c^3}{4G} f_{GW}^2 \overline{h^2} \\ &\approx 3.2 \times 10^6 \left(\frac{f_{GW}}{1 \text{ kHz}} \right)^2 \frac{\overline{h^2}}{(10^{-19})^2} \frac{\text{ergs}}{\text{cm}^2 \text{ sec}} \end{aligned} \quad (22)$$

where $\overline{h^2}$ is the total mean square strain (including both polarization components). Thus at a radius r , and assuming “typical” orientation of the source, the power output of a system producing this flux would have to be

$$\begin{aligned} \frac{dE_{GW}}{dt} &\approx 4\pi r^2 \frac{dE_{GW}}{d^2x dt} \\ &\approx 9.6 \times 10^{53} \left(\frac{f_{GW}}{1 \text{ kHz}} \right)^2 \left(\frac{r}{50 \text{ kpc}} \right)^2 \frac{\overline{h^2}}{(10^{-19})^2} \frac{\text{ergs}}{\text{sec}} \\ &\approx 0.53 \left(\frac{f_{GW}}{1 \text{ kHz}} \right)^2 \left(\frac{r}{50 \text{ kpc}} \right)^2 \frac{\overline{h^2}}{(10^{-19})^2} \frac{M_{\odot} c^2}{\text{sec}}. \end{aligned} \quad (23)$$

The last form of this expression makes it clear that a hypothetical source detectable in our 1987 observations would have not only to be extraordinarily efficient at producing gravitational waves, but would also need an enormous reservoir of raw material to keep it going (of order 10^5 solar masses per day). Detecting a signal from the 1987a remnant in our data would thus have revolutionary implications for astrophysics.

What about the universe in general? The center of our own galaxy, for example, is five times closer than the Large Magellanic Cloud, and contains considerably more “raw material” than the neighborhood of our supernova remnant. Previous kilohertz-band searches for coherent gravitational radiation have either been restricted to very narrow frequency bands [38,70] or have attained lower sensitivity [54,72] than is available from our observations. The results are summarized in Figure 30.

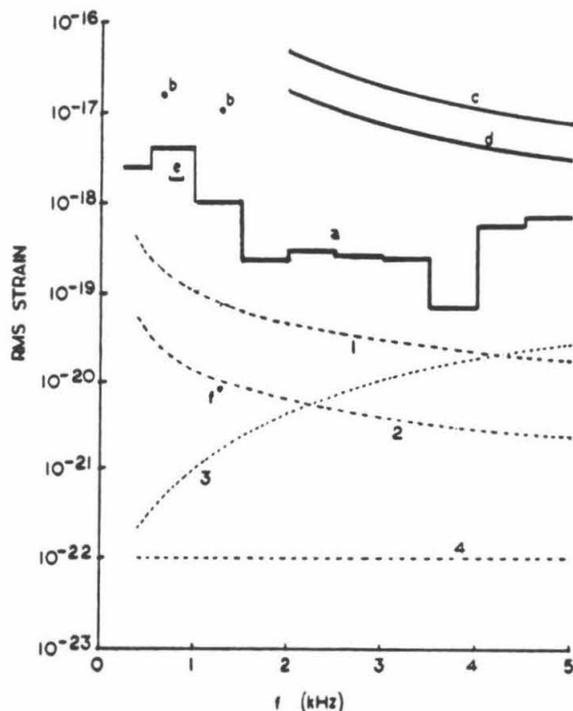


Figure 30: Experimental upper limits on continuous kilohertz-band gravitational radiation are presented with some hypothetical source models. Curve (a) is our upper limit on “(+)”-polarized periodic strains emanating from the galactic center (§ 4.7). Points (b) are limits determined by Hereld [38] on periodic radiation from the millisecond pulsar PSR1937+214, at 642 Hz (the radio pulse rate) and 1284 Hz ($2f_{pulse}$). Point (f) is an upper limit on the same source at 1284 Hz reported by Hough et al. [70]. Curves (c) and (d) are the limits set by Livas [54] on radiation from arbitrary directions and from the direction of the galactic center, respectively. Curve (e) is the upper limit reported by Hough et al. [72] on continuous radiation from all directions. Theoretical estimates (broken lines) are as follows; (1) would result from 1 solar mass per minute converted entirely into sinusoidal gravitational waves and radiated isotropically from the galactic center; (2) represents 1 solar mass per hour radiated similarly. Curve (3) represents radiation from a neutron star with an ellipticity ϵ of 10^{-2} located at the galactic center; and curve (4) is the upper limit to radiation from the supernova remnant based on a radiative spindown time of one week (for average orientation of the spin axis; the peak strain at optimum source orientation is about twice as large). No allowance has been made for the different confidence levels of the various experimental limits. The reader is referred to the cited publications for details.

In summary, detectable periodic gravitational radiation from the remnant of the supernova seems theoretically improbable. On the other hand, the sensitivity of even a modest search for periodic strains from *any* source would improve on previously published surveys by a large factor over a considerable bandwidth.

4.5 Data Analysis

Several segments of interferometer data were scanned for evidence of periodic gravitational waves using standard techniques. As expected, a disproportionate amount of time was spent recognizing and rejecting candidate periodicities which turned out to be local in origin; before giving details, it may be worth reminding the reader that a simple comparison of these data with those recorded at MIT and Glasgow could serve to reject many local disturbances immediately, so these procedures should only be viewed as a coarse “first pass.”

4.5.1 General Method

An estimate of the power spectrum of the signal derived from the secondary cavity servo was formed by a standard Fourier transform computer algorithm. First, the chart recorder records (Appendix D) and detailed tables of “good” and “bad” tape segments prepared by Sheryl Smith [37] were scanned to find contiguous stretches of 2^{20} samples (105 seconds) within a tape during which both cavities remained locked. Momentary dropouts and electronic saturations were by and large ignored; these were fairly common (again, likely due to the distortion of the cavities and the resulting instability against optical mode perturbations). No overwhelming difference was seen between spectra formed from data sets without any dropouts and those with a few, indicating that ones of short duration contributed a negligible amount of energy to the spectrum. As a rule, dropouts long enough to register clearly on the chart record of the secondary cavity reflected light (about .1 second or longer) disqualified the segment as “discontinuous.”

4.5.2 Transformation

Each selected segment was analyzed for periodicities by calculating its discrete Fourier transform, or DFT. Briefly, this technique forms a discrete approximation \tilde{X}_k to the continuous, infinite Fourier integral transform $\tilde{X}(f)$ of the time record $x(t)$. The computer took N discrete samples $x_j = x(t = t_0 + Tj/N)$ of the interferometer strain signal. Each of the N elements of the discrete transform is then given by the sum

$$\begin{aligned} \tilde{X}_k &= \sum_{j=0}^{N-1} x_j e^{2\pi i \frac{jk}{N}} \\ &\approx \frac{N}{T} \tilde{X}(f_k = \frac{2\pi k}{T}). \end{aligned} \tag{24}$$

In our case, the input data are real numbers; the number of independent Fourier components \tilde{X}_k is thus reduced from N to $N/2$ by the symmetry of the transform [73]. These complex components represent equally-spaced frequencies from $k = 0$ (DC) up to $k = N/2$ (the Nyquist frequency, which is half the sampling rate). The discrete Fourier series formed with these coefficients (which reproduces the original sample vector) is formally equivalent to a least-squares fit of the original continuous signal to a sum of sines and cosines at integrally-related basis frequencies $f_k = kN/T$, $k = 0, 1, 2, \dots, N$. The method is thus optimal (in the least squares sense) for discovering sinusoidal signals in random noise.

One key motivation for using this method is the Fast Fourier Transform or FFT, a class of algorithms by which the N^2 terms needed to evaluate the DFT (Equation 25) at all N frequencies is drastically reduced to a number of order $N \log_2 N$ by virtue of the intrinsic symmetry of the exponential function inside the sum. The idea, which is most readily applicable where N is an integral power of 2, is nicely explained in [74]. The particular computer routine used, called MSFFT, is designed for optimum application in the unfortunate case that

the entire data vector will not fit into the memory of the computer [75]. Since our Masscomp computer has only 1 Mbyte of usable random-access memory, at most 262,144 four-byte numbers could be stored in memory at one time, requiring considerable amounts of comparatively slow random access to the mass storage device (a magnetic disk) during the computation of each million-point transform used here. Without any effort to optimize execution time beyond that implicit in the subroutines used, the transformation of each 2^{20} -point data vector considered here took approximately 1.2 hours of real time if the machine was not loaded with other tasks. The excessive time consumption (over 4,000 seconds per 105-second data segment) is largely due to the memory limitation, making comparison with other methods and hardware difficult. For some useful benchmark data on state-of-the-art equipment, see [54, Appendix F].

After forming the DFT of the data vector, subsequent procedures concentrated on a quantity proportional to the single-sided power spectrum,

$$S_k = |\tilde{X}_k|^2, \quad (25)$$

for each complex Fourier component \tilde{X}_k . The phase of each component was discarded in this analysis; however, phase information was useful in discovering related classes of local disturbances, and can always be addressed by retransforming the original data vector.

4.5.3 Windowing

Selected segments were read from the original data tape into a disk file on the computer, after conversion from the recorded short integer format to 4-byte real format (for compatibility with existing computer programs). Each stored data vector was windowed in-place by a simple algorithm. The purpose of this step is to reduce the errors inherent in estimating the power spectrum of a continuous

signal from only a finite number of discrete samples; if simply “chopped off” abruptly at the beginning and end of the segment, the sharp edges will add considerable spurious power to the spectral estimate, seen as a long “tail” of sidebands appended to any narrowband feature². Since this search turned up many locally generated narrowband interference signals, such leakage from the large interference spikes would tend to mask smaller features nearby.

A von Hann (or “raised cosine”) window was adopted for its high leakage suppression [76, §9-5]; the window function is simply

$$w_j = \frac{1}{2} \left(1 - \cos \left(\frac{2\pi j}{N-1} \right) \right) \quad (26)$$

where N is the number of samples in the data vector. Each element in the original vector x_j is multiplied by w_j to form a new N -element vector x'_j , which is substituted for the original in subsequent procedures. Here, the most important consequences of this procedure are:

- A sharp spectral feature is effectively broadened into the three adjacent frequency bins. For example, a true sinusoidal signal at a basis frequency will show up with its power redistributed among that bin and its two nearest neighbors in the approximate ratio $\frac{1}{4} : \frac{1}{2} : \frac{1}{4}$. This is due to the effective amplitude modulation of any signal component by the window function (Equation 26), whose period equals the record length (producing “modulation sidebands” one resolution bin away on each side of a peak).
- Leakage into more distant sidebands is markedly reduced, since the window function itself has very small Fourier components at higher multiples of the basis frequency.

²Unless the period coincides with a “basis period” of the transform, an exact integral divisor of the length of the data vector. This is an atypical occurrence.

- The actual numerical value of a power spectrum feature is proportional to the area under the square of the window function [74, page 424]. (This factor comes to 3/8 for the von Hann window, and unity for the rectangular or “uniform” window.)
- The procedure affects coherent signals differently than incoherent random noise; an unresolved spectral feature, as noted above, is diluted over adjacent channels, while broadband noise is simultaneously present in all nearby frequency bins. As a result, the “significance” of such an unresolved peak with respect to the background is degraded. The effect is similar to that of increasing the effective noise bandwidth by a factor of 1.5 over the resolution bandwidth f_{samp}/N [54,73,74].
- A spectral line falling midway between basis frequencies will be shown at only 84% of the estimated amplitude (70% of the power) of an otherwise identical signal which happened to fall right on a basis frequency. While in principle one can easily correct for this, it is not significant in view of other estimation errors. Other window functions can increase amplitude resolution at the expense of frequency resolution and *vice versa* [77, § 5].

The windowing routine and the Fourier transform program were invoked in sequence by a single command; the windowing added another 15 minutes or so to the total transform time. This time could be spared by windowing the data as it is transferred from tape (the modest computations are insignificant compared to the tape reading and buffering), but it was deemed useful to sneak a look at the unwindowed raw data once on the disk to check its integrity and to save the first few thousand samples as a “fingerprint.” This helped verify that future references indeed addressed the same segment.

4.5.4 Finding the Peaks

Even such modest-length spectra (half a million bins for the 105-second segments) need some kind of automation for finding and examining candidate peaks. A simple plot of the results, allowing an eye-straining 0.1 mm per bin, would stretch over 50 meters. The procedure adopted closely parallels methods used by Livas [54] and Armstrong [78] for similar searches.

A computer program stepped through the disk file containing the Fourier transform, stored in eight-byte complex binary format. The power spectrum of the recorded signal is far from white; it was therefore necessary to form local estimates of the mean power spectrum in order to decide if any given bin contained an unusually large amount of energy. The bandwidth defining “local” was fixed despite considerable variation in the frequency scales over which the background noise fluctuated. At low frequencies, mechanical resonances and sharp filtering modified the power spectrum significantly between frequencies separated by only a few Hertz, while at higher frequencies gentler trends affected the background more gradually (with some exceptions due to local interference). A bandwidth of 256 resolution bins, about 2.44 Hz, was settled on after some experimentation. This provided a reasonably small statistical variance for the local estimated background, while allowing the estimate to follow most of the structure in the noise closely.

The Fourier components were read in records of this length from the disk file, and each record was treated as an isolated subspectrum. First, the squared magnitudes of elements in each record were averaged to form the “local mean power” for that record. Then, each of the 256 squared magnitudes was compared with this local mean. Every bin whose power exceeded the mean by a factor greater than a preset threshold was tagged as a candidate peak. Optionally, a separate subroutine could then cull out candidates which fell within a preset

“dead band” around the frequency of each known local periodic disturbance. The primary culprits, and the only ones pernicious enough to demand automatic vetoing, were related to harmonic multiples of the 60 Hz power mains frequency. Any peaks above local threshold passing these secondary cuts were reported to a data file, noting the frequency, absolute size (in squared digital units), and relative size with respect to the local mean. After checking all 256 bins the value of the local mean and the boundaries of the record were also recorded in another file for reference. The next record was then read in and the process repeated.

Thresholds were typically set to report at most a few hundred outstanding peaks, making it possible to peruse the list manually (that being the idea of automation). Plots of the number of peaks as a function of their height (a “peak height histogram”) or of their approximate frequency (a “peak distribution histogram”) were derived from these lists, made at a variety of thresholds. The naive peakfinding algorithm was inefficient; each peak list took approximately 25 minutes to generate, depending on the threshold and the vetoes enabled.

When added to the time required for reading the tape onto disk (about 15 minutes), windowing, and transforming the data vector, each 105-second segment took over two hours of real time to crunch. This burden significantly reduced the flexibility of the search, especially when combined with another limitation; the fixed disk drive on the Masscomp (a 70 megabyte Fujitsu) was all but full, even after sacrificing other users’ space and removing instruction manuals and utilities to backup tapes. As a result, at any given time only one million-point segment or its transform could be stored, hampering comparison of data from different times and dates. A used 690 megabyte drive was purchased, but the necessary controllers and cables arrived too late to influence the analysis. This and the time-consuming code are the primary reasons for not analyzing more of the 70 tapes recorded.

The peaks and local mean background levels were stored in “raw” transformed tape units. They could then be expressed as strains in the antenna by use of a cubic spline interpolation to coarse calibration spectra, produced by driving one mirror of the interferometer with pseudorandom noise (as described in § D.3). This calibration was only accurate above about 300 Hz; the search was restricted to this range. The same spline parameters were used for all segments analyzed without adjustment, since the various calibration tests were mutually consistent.

4.5.5 Statistical Expectations

Deciding statistically whether a peak in the spectrum could represent an astrophysical signal or not is severely hampered by the presence of colored noise and periodic interference. We have tried to deal with the coloration of the noise by estimating its mean power in local bands. Temporarily neglecting the periodic interference, we can at least get a rough idea of how many peaks to expect just from random fluctuations. This baseline will drive the hunt for departures from the theoretical ideal; more relevant, it will indicate the potential benefits of expunging local interference, e.g., by employing data from other detectors.

Under the temporary assumption that the noise is random and aperiodic, we can look upon the set of $N/2$ complex Fourier amplitudes as samples from N distinct random number generators (one each for real and imaginary parts). Furthermore, each generator will produce a roughly Gaussian distribution with zero mean (since each x_j in the sum of Equation 25 itself comes independently from an assumed random sequence). It follows that the spectral power $S_k = |\bar{X}_k|^2 = |Re(\bar{X}_k)|^2 + |Im(\bar{X}_k)|^2$ will be distributed as a χ^2 random variable having two degrees of freedom. If we normalize the S_k by the local mean power $\bar{S}(f_k)$ (an abstract property of the physical process generating the noise, to be estimated by our 256–point local averages) the resulting random variable $u \equiv (S_k/\bar{S}(f_k))$

will have the simple probability density distribution

$$p(u) = e^{-u}; \quad (27)$$

that is, the probability that u will be found between a and b ($b > a$) is just

$$P(a < u < b) = \int_a^b p(u) du = e^{-a} - e^{-b}. \quad (28)$$

These properties were verified by substituting the band-limited, amplified thermal noise of a carefully screened 1 M Ω resistor held at room temperature for the gravitational wave antenna readout at the input of the ADC, and recording and processing this “true noise” as described above. The normalized powers u for 102,400 bins of the 524,240-bin spectrum are histogrammed in Figure 31, showing agreement with Equation 27.

For an imaginary ensemble of independent spectra produced by identical copies of this ideal process, we can derive the probability distribution of the discrete variable $n_m(v)$, defined as the number of bins with normalized power u exceeding some given threshold v in a spectrum of m bins. This is simply the probability that, say, the first n_m bins exceed v while the rest do not, times the number of possible ways to choose these n_m special bins, i.e. a binomial distribution

$$P[n_m(v)] = (e^{-v})^{n_m} (1 - e^{-v})^{m-n_m} \binom{m}{n_m} \quad (29)$$

Over the ensemble, the average number of bins exceeding threshold v will clearly be me^{-v} , and the variance $\sigma_{n_m(v)}^2 = me^{-v}(1 - e^{-v})$. Applying the results for this hypothetical ensemble of spectra produced by identical physical processes to the set of spectra taken from our one apparatus at different times requires assuming the noise is ergodic; this assumption is also violated in the real world, but relatively weakly, so we’ll ignore that restriction as well.

For example, if each spectrum of 524,288 bins was truly produced by the above “idealized” random process, and if our estimates of the local mean power

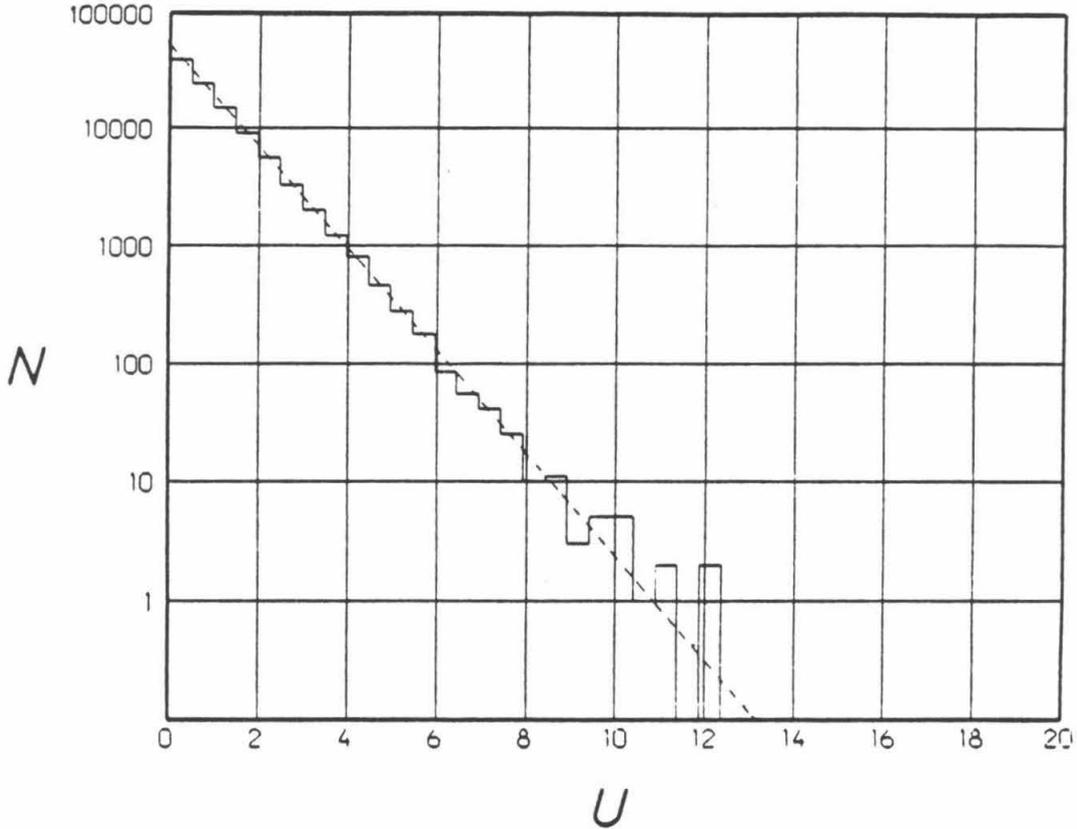


Figure 31: Thermal noise from a resistor was amplified and band-limited, then recorded and processed in the same manner as the actual data. $M = 102,400$ bins from the normalized power spectrum are histogrammed here; the theoretical prediction of Equation 27 is also shown. Both curves are normalized to unit probability (i.e., enclose area M). The two peaks found with $u \geq 11.8$ are not exceptional; the theoretical distribution should produce two peaks this size or bigger in about 1 out of 7 trials (cf. Equation 29)

are accurate (meaning the noise power spectrum is smooth on the scale of our averaging bandwidth), we would expect an average spectrum to have about 24 bins with $u > 10$. Furthermore, 9 out of 10 similarly prepared spectra would each be expected to contain between 16 and 32 such bins.

The task of deciding whether or not a peak is important proceeds by assigning a different probability distribution to the normalized power u reported for a bin which contains an astrophysical signal added to the noise, and then determining the likelihood that a given bin belongs to this distribution and not to the “pure noise” distribution above. For example, a waveform like

$$x(t) = a \cos(2\pi f_{GW}t) + x_n(t), \quad (30)$$

a sinusoid of amplitude a added to the original random noise $x_n(t)$, will produce power spectra having $(N/2) - 1$ bins distributed as above and one bin³ nearest the frequency f_{GW} distributed with the Rice-squared probability density [78, Equation A2]

$$p_a(u) = \exp[-(u + a^2)] \mathbf{I}_0(\sqrt{4a^2u}) \quad (31)$$

where \mathbf{I}_0 is the zeroth hyperbolic Bessel function. For a peak distinguished by its low probability of arising from noise alone, this distribution function would allow us to estimate the amplitude a ; if nothing were found, it would allow us to estimate the probability of having accidentally missed something actually present at a given a .

Most of the work with the actual data was devoted to trying to get rid of the effects of local periodic interference, which rendered the above results academic. The initial lists of “statistically significant” peaks were at first a factor of ten longer than would be produced by the “simple” noise process we considered.

³Here we neglect the windowing effects discussed above.

Surely these were not all gravitational waves! After some detective work, however, the biggest offenders were eliminated, and the pared-down peak list more nearly approached the estimates mentioned in Section 4.4.

4.6 Results

Eight segments of seven different tapes were selected from the 70-tape library accumulated during the 1987 observations. All were recorded at 10 (kilosamples/second)/channel with the 5 channel assignments described in Appendix D. Of these, two were short (less than 30 second) calibration tests in which pseudo-random white noise of known amplitude had been applied to the Huey (primary cavity) mirror. One 105-second segment was abandoned after initial analysis revealed anomalously high broadband noise⁴; another segment was analyzed before spurious peak rejection procedures had been fully developed, and was not subsequently reanalyzed. The four remaining segments showed reasonably uniform characteristics. They were chosen from three tapes, two from the coincidence run of 9 March and one from 11 March, 1987. All eight segments are summarized in Table 1.

4.6.1 First Impressions

Before starting it was known (from routine diagnostic spectra) that 60 Hz line multiples contaminated the noise spectrum at low frequencies. It was not known if these “line spikes” would appear up to the 83rd harmonic (4980 Hz, just below our Nyquist limit) since the diagnostic spectra typically had low resolution. As a result the peakfinding algorithm was equipped with a means of rejecting any peak found within a preset deadband of ± 100 bins (about 0.95 Hz) either way from any line frequency multiple. The mean line frequency was determined by examining the full-resolution spectrum itself, forming an average of the frequencies of the

⁴Recall that the noise fluctuated occasionally due to the time-varying mode structure of the distorted cavities; in this recorded segment the chart recorder trace showed no significant effect in the fringe contrast, but the average contrast was so poor ($\sim 40\%$) that this is not a sensitive gauge of the mode quality.

<i>Tape (seg)</i>	<i>Start Time</i>	<i>Length</i>	<i>Comments</i>
27	067:14:26:23.36	~0.41 s	white noise calibration
34	067:15:49:10.00	104.86 s	partial analysis
45 (a)	068:01:10:36.76	104.86 s	included
45 (b)	068:01:14:17.74	104.86 s	included
49	068:01:51:43.32	104.86 s	included
52	070:06:08:00.00	~0.50 s	white noise calibration
53	070:06:18:19.91	104.86 s	high broadband noise
56	070:06:53:30.00	104.86 s	included

Table 1: The four tape segments studied most intensively, listed along with the two short calibration recordings and two partially analyzed segments not included in the conclusions (noted in case someone is particularly interested in those times). The starting times are in the format [day of 1987]:[hour]:[minute]:[seconds].[hundredths], universal time (UT).

largest obvious line spikes (each divided by the appropriate integer). The 200-bin zone width was chosen to allow for the width of the broadest line spike, and was applied up to 5 kHz. This results in the loss of about $83 \times 1.9 \text{ Hz} = 158 \text{ Hz}$, or 3% of the search bandwidth; later developments (§ 4.6.3) doubled the loss. More refined methods could probably reduce the total bandwidth ignored in avoiding line spikes; for future observations, of course, they should just be eliminated⁵. A further practical matter, the absence of accurate calibration at very low frequencies, limited the serious peak search to frequencies above 305.18 Hz (bins #32,000 to #524,288).

After the exclusion of suspected line spike regions, the three spectra computed from tape segments 45(a), 45(b) and 49 yielded 167, 64 and 158 bins, respectively, above threshold $v = 10$. Recalling the predictions of Equation 27, we would only anticipate about 22 of the 475,720 “live” bins to exceed $v = 10$ if the noise were

⁵The 40m antenna has indeed been purged of most line interference since the 1987 observations.

aperiodic and random; 64 or more would be expected by chance in only one out of 10^{13} spectra. Figure 32 shows the situation for a particular region of the spectrum from Tape 45(b). Conclusion; the antenna output is *not* random, aperiodic noise.

As expected, the peaks were most numerous in the low-frequency portion of the spectrum, below 2 kHz (cf. Figure 35). Resonances in the apparatus tend to be most numerous and least susceptible to isolation measures at lower frequencies. A look at the coarse power spectrum (Figure 33) also shows that the noise spectrum undergoes sharp changes over short frequency intervals, confounding our smoothed estimate of the broadband noise. Some peaks reported as “significant” were thus compared with an inaccurate (low) estimate of the local background. On the other hand, some smaller peaks were also probably masked by high background estimates, formed in averaging bands which sported a large spurious peak (e.g., a line spike).

Under the assumption that few if any of the reported peaks were caused by gravitational waves, mechanisms of local interference were examined to try and explain at least the largest ones. As little prejudice as possible was included about the nature of true astrophysical signals; thus, for example, peaks with resolved broadband structure (and hence a lower “ Q ” than the astrophysical processes guessed at above) were not discounted simply on the basis of their shape.

Several potential sources of spurious peaks were investigated. Two were found, both effects of poor isolation of the apparatus and recording system, which allowed rejection of substantial numbers of peaks. A third suspected mechanism, nonlinearity in the recording system, yielded inconclusive results and merits more study.

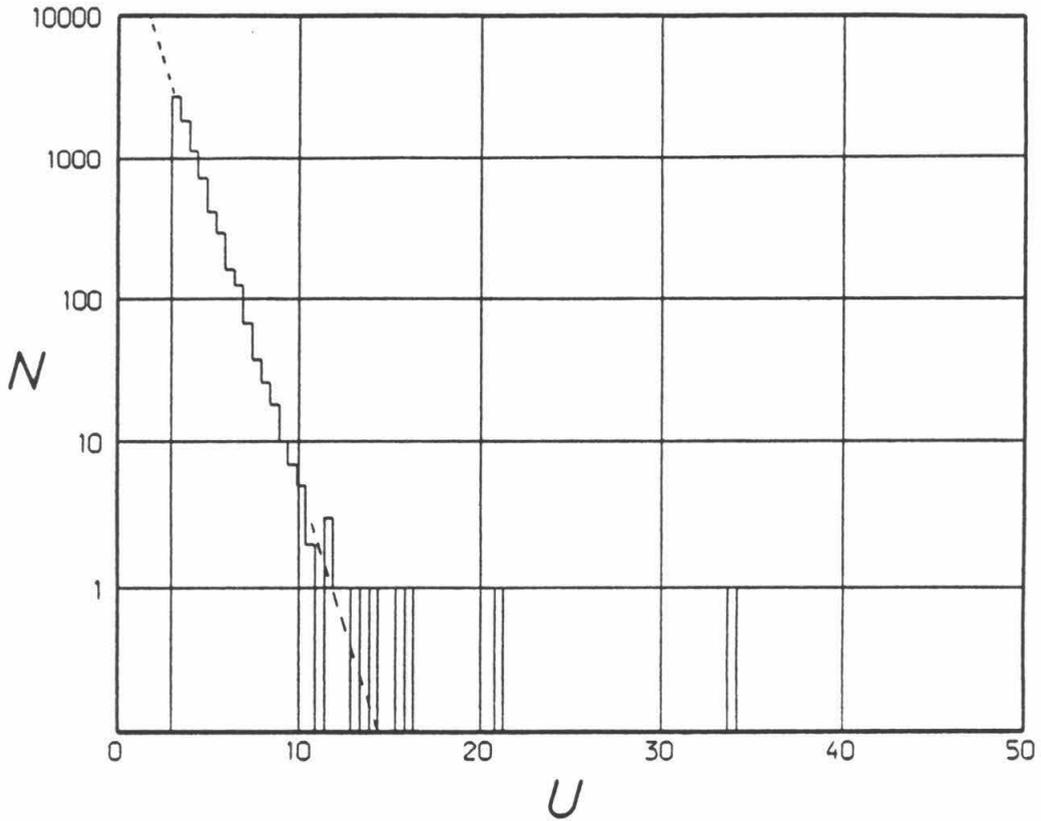


Figure 32: Histogram of bin heights in the spectrum of Tape 45(b) between 1 kHz and 1.5 kHz (52,428 bins). A threshold of $u = 4$ was used; below this, the distribution follows the theoretical curve for random noise (sketched in) closely. An excess of high-power bins is evident (cf. Figure 31, Equation 27). This graph is just an example; some regions of the spectrum showed fewer departures from the theoretical curve, while others showed more.

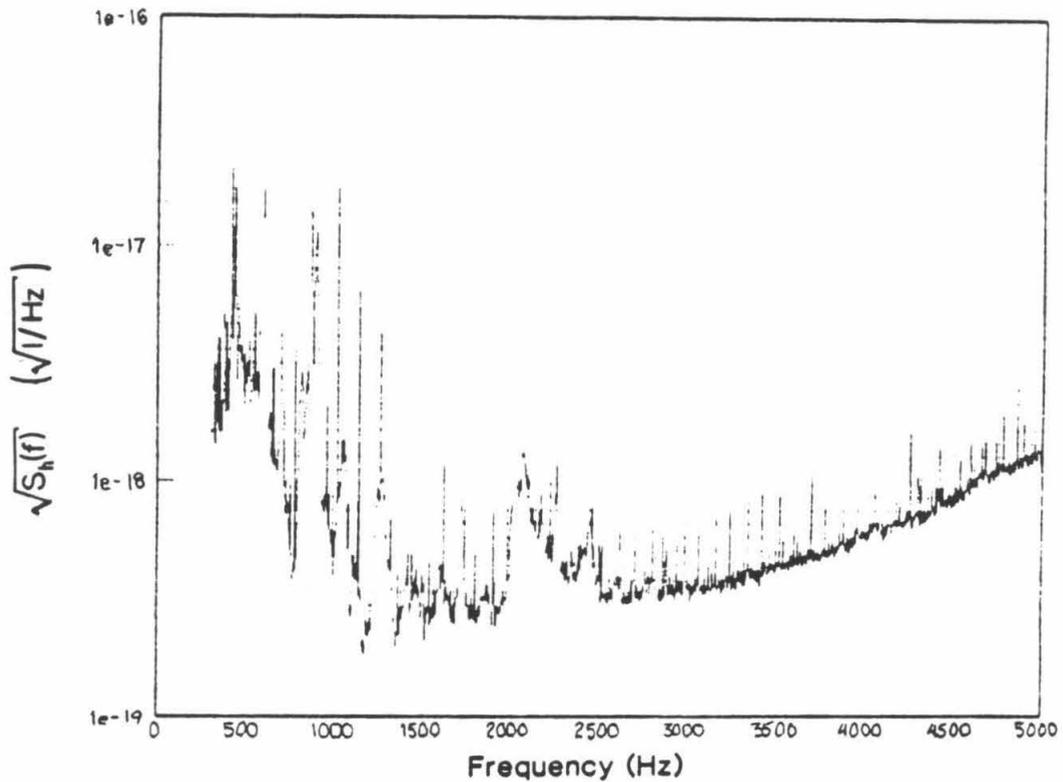


Figure 33: Local mean power per bin calculated for each 256-bin region of the transform of tape segment 45(a), replotted as a coarse strain spectrum $\sqrt{S_h(f)}$. Note that averaging 256 adjacent bins of the fine-resolution spectrum is equivalent to averaging 256 successive lower-resolution spectra, the more usual practice for power spectrum estimation.

4.6.2 Microphonics

Acoustic excitation of the laser and optical fiber was known to induce broadband noise in the apparatus, but it was not realized until these data were analyzed that the ambient acoustic spectrum in the laboratory also contains numerous fine, high- Q features invisible to previous lower-resolution tests. The microphone track recorded on tape along with the strain signal was transformed and analyzed in exactly the same way, except that the microphone was not calibrated. The spectrum revealed many “excess” (non-statistical) peaks, as did the strain signal; moreover, a large number appeared at exactly the same frequencies in both spectra, even after excluding line multiples. This was fairly concrete circumstantial evidence that those peaks in the strain spectrum were caused by periodic sound interference. To be sure there was no coincidence, several of the suspects were scrutinized by playing a tape back through the computer DAC into an HP3562A Spectrum Analyzer and looking for coherence in the time-averaged cross spectrum of the two signals⁶. For all the coincident peak pairs examined in this way, the microphone and strain signals maintained phase coherence over 240 second intervals; further support came from the striking similarity of the shapes of those lines that were broad enough to be resolved in both spectra.

Two natural questions arose: How big a peak in the microphone spectrum was sufficient grounds for disqualifying a nearby peak in the strain spectrum; and how close did they have to be in frequency to justify the assumption of cause and effect? If the coupling transfer function is roughly linear, and doesn’t vary much over frequencies of order the width of the peaks (i.e., the transfer function has lower Q than the offending peaks), then (a) the ratio of the peak height to the

⁶This trick saved a lot of programming, since the commercial instrument is designed to do such involved calculations trivially, but it required considerable faith in the accuracy of the DAC. Independent tests showed that the DAC was indeed rather good, in fact much better than the ADC.

nearby background ought to be larger in the spectrum of the “cause” than the “effect” (or at least be equal), and (b) the frequencies should match within the half-width of the feature. Spot checks of the two spectra implied that property (a) was always satisfied; all suspected pairs were also well aligned in frequency (often in the exact same bins), satisfying (b) as well. Hence it was sufficient to scan the microphone spectrum with the same threshold v used to find the peaks in the strain spectrum, and to look only for close coincidence (within 10 bins) between reported groups common to the two peak lists. These were then simply crossed off the candidate list by hand.

The satisfaction of condition (a) above, in addition to pointing to a very simple kind of acoustic interaction, could also allow us to attribute some fraction of the broadband noise in the detector to acoustic infiltration. Some peaks whose “significance” with respect to the local background was nearly equal in both spectra show that the broadband noise in those regions was probably also acoustic in nature. However, except for very narrow features (less than 0.1 Hz wide) little coherence was seen between the microphone and strain spectra. It is possible that a time-varying or nonlinear coupling mechanism would allow acoustic energy to infiltrate incoherently⁷.

Many of the microphonic peaks were fully resolved (i.e., significantly broader than 10 mHz) and occurred at frequencies below 1 kHz. The only exception was a very large ($u \sim 50$) unresolved peak at 4997.63 Hz, appearing at the same place (within 1 bin) in all four spectra. The source for this (and the other “normal” acoustic peaks) remains a mystery. One possibility is that the blades of cooling fans in electronic lab equipment produced periodic sound, at arbitrary but

⁷In this case one might still reveal a cause/effect relationship by cross-correlating the *squares* of the two signals, as in an intensity interferometer [79], but this method (suggested to me by Ron Drever) is severely limited in cases like this where the signals have pronounced periodicities and strongly colored spectra to begin with.

nearly steady frequencies set by the motor speeds. The peak amplitudes are all consistent with their having been inaudible over the incoherent background noise (according to [11], the human ear and brain have an equivalent noise bandwidth of about 120 Hertz).

4.6.3 Aliased Line Multiples

The microphonic peak veto cleaned up some low-frequency peaks, but a large excess remained. Line frequency harmonics were visible all the way up to #83 (just below the 5 kHz Nyquist frequency); furthermore, a simple peak pairing program revealed that even after rejecting direct multiples of the line frequency, a significant number of pairs of large peaks were still *separated* by multiples of that frequency. An obvious explanation was that these were in fact very high frequency line multiples which somehow got past the anti-aliasing filters and were aliased into the band of interest. This speculation turned out to be right; beyond spike #83, the series continued back *downward*, as though reflected by the 5 kHz Nyquist limit.

In retrospect, it seems understandable that line harmonics above 5 kHz could sneak into the ADC inputs around the anti-aliasing filters. Transient pulses at a repetition rate of 180 Hz, with frequency components up to 20 kHz, were often seen superimposed on sensitive signals in the laboratory; they were probably produced by the switching of rectifiers in the laser power supply. These spikes were both radiated and directly added to the power distribution lines in the lab. Interferometer signals were routinely amplified and filtered by AC-powered equipment in a uniformly grounded rack system inside the laboratory, near the vertex optical table, while the computer system was located in a nearby control room; the rapid setup for the observing run consisted of just running some ten-meter cables through the door to the ADC inputs. This could have caused

ground loops and/or enormous antennas, which bypassed the filters and went undiscovered in the rush to bring the antenna on line.

The peakfinding algorithm was therefore modified to skip over a new series of peaks, starting with #84 (at 5040 Hz, but aliased down to $10,000 - 5040 = 4960$ Hz) and proceeding downward toward DC. At these high frequencies, the line spikes were much broader in absolute terms; this effect would be expected if the peaks were broadened by phase noise resulting from changing reactive loads on the power grid. It was thus necessary to determine the mean line frequency during each tape segment very accurately to avoid letting the “tails” of the highest harmonics evade the veto deadband. As it was, the deadband of 1.9 Hz was barely enough to cover them. Aliased spikes were seen up to #147 (8,820 Hz); the program was automated to delete them up to #166 (just below the 10 kHz sample rate) for simplicity. It is quite possible that members of the set existed in yet another aliased series proceeding up from #167, which would appear aliased to 20 Hz, but none were detected.

All told, about 316 Hz of the 4,695 Hz search bandwidth (6.7 %) was intentionally given up to line frequency interference. However, the actual presence of a line spike tended to greatly inflate the local mean power estimate used to calculate the peak thresholds; therefore, each spike could well have interfered with the discovery of other legitimate peaks, potentially over a region as broad as 4.88 Hz (if the line spike straddled two local averaging bands and poisoned both). This effect was not fully accounted for, depending as it does on the strengths of the spikes and on the unknown distribution of peak heights.

4.6.4 Remaining Peaks

The results of applying the added vetoes on the four peak lists are summarized in Table 2; although the number of peaks above threshold $u = 10$ remained

<i>Tape(seg)</i>	<i>Threshold</i>	<i>L</i>	<i>L,A</i>	<i>L,A,M</i>
45(a)	10.0	167	147	103
	16.0	48	48	27
	25.0	21	21	9
45(b)	10.0	64	57	48
	16.0	16	16	9
	25.0	10	10	6
49	10.0	158	135	90
	16.0	36	32	21
	25.0	11	10	6
56	10.0	n/a	61	48
	16.0	n/a	27	18
	25.0	n/a	13	9

Table 2: The number of bins reported exceeding each of three thresholds, after applying three successive vetoes; first, peaks within a preset interval of a line multiple were excluded (column marked “*L*”); then, peaks close to aliased line multiples were excluded (column marked “*L,A*”); and finally, peaks common to the microphone spectrum were deleted (column marked “*L,A,M*”). Since these are numbers of *bins*, large features which affect several bins are reported several times; in most cases just a few persistent large peaks are responsible for all the remaining counts.

excessive, it was considerably reduced. No other obvious local interference was exposed. Many of the remaining peaks are probably attributable to high- Q resonances in the apparatus. Some are very likely to be violin-mode resonances of the wires suspending one or more test masses. Most of the larger peaks reported actually comprised several close bins, leading to multiple reporting of many features; recall that even a monochromatic feature will typically be smeared out into three bins by the windowing. Thus the tabulated numbers include many correlated bins, and don’t accurately represent the statistical distribution of the peak heights.

An obvious additional test, to mechanically excite each mass and determine

the frequencies and Q 's of major resonances, was unfortunately not done before these masses were torn out and replaced by more modern ones. The wire resonances, in particular, are thus at unknown frequencies, although from previous experience we know they can appear in the spectrum. The best we can do is estimate their number and approximate frequencies and try to guess which features they are. A grouping of four large, narrow peaks in the range 630 to 640 Hz, seen to some degree in all the spectra, are almost certainly the slightly nondegenerate resonances of the four wires on one test mass. These are shown in Figure 34. The three peaks reported (one fell near a line harmonic and got vetoed) account for 8 of the 18 bins above $u = 16$ and 6 of the 9 bins above $u = 25$ in this spectrum from Tape 56.

More large peaks were still evident at lower frequencies, even after the veto procedures. For example, the distribution of bins above a uniform threshold $u = 10$ vs. frequency has been plotted for Tape 45(a) both before and after application of all the interference vetoes (Figure 35). Another example is presented from Tape 56 in Figure 36. The height of the biggest "outlier" peaks in any given stretch of spectrum depended strongly on frequency as well. These two qualities motivated breaking the spectrum up somewhat arbitrarily into several distinct regions, in a sense coarsely recognizing the different statistical nature of the noise (or interference) at different frequencies.

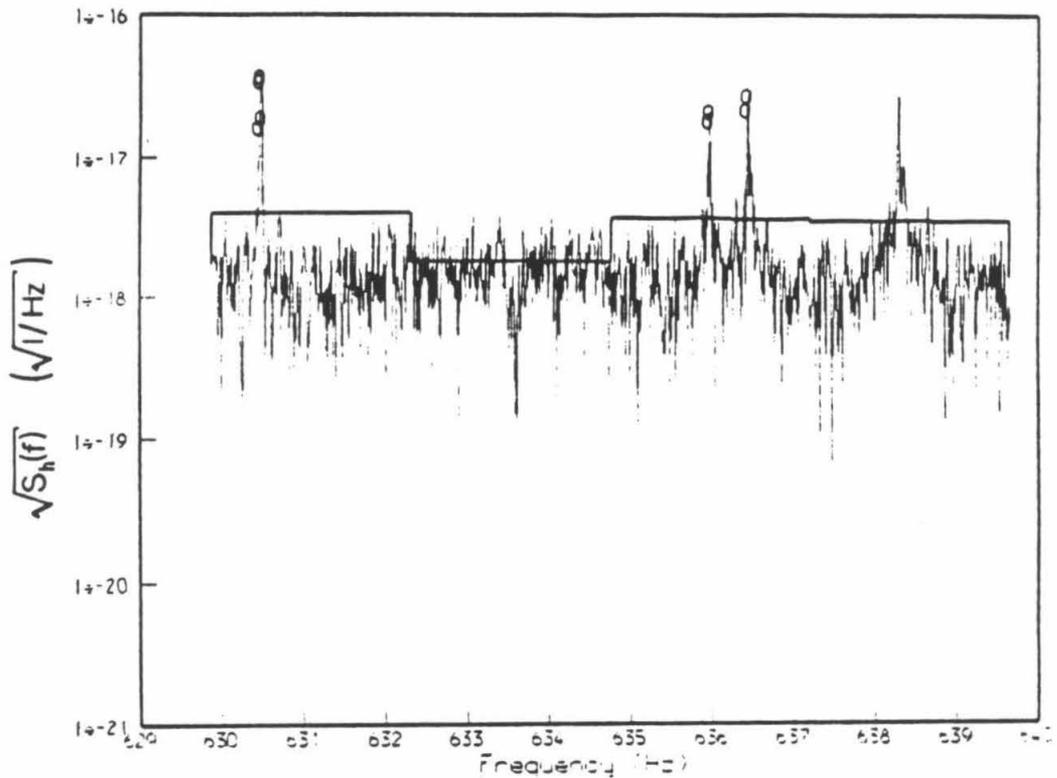


Figure 34: These peaks, which appeared in all four spectra, are probably the violin-mode resonances of the suspension wires on one of the test masses. They may have been excited by electronic noise on the suspension transducers passed by faulty filters. The widths indicate Q 's in excess of 10,000. No other group of features having high Q 's in such close proximity were found, indicating that efforts to quell wire resonances were *usually* successful. This plot is from Tape 56; the rightmost peak fell within the veto band of a nearby line spike, and was not reported in the peak list.

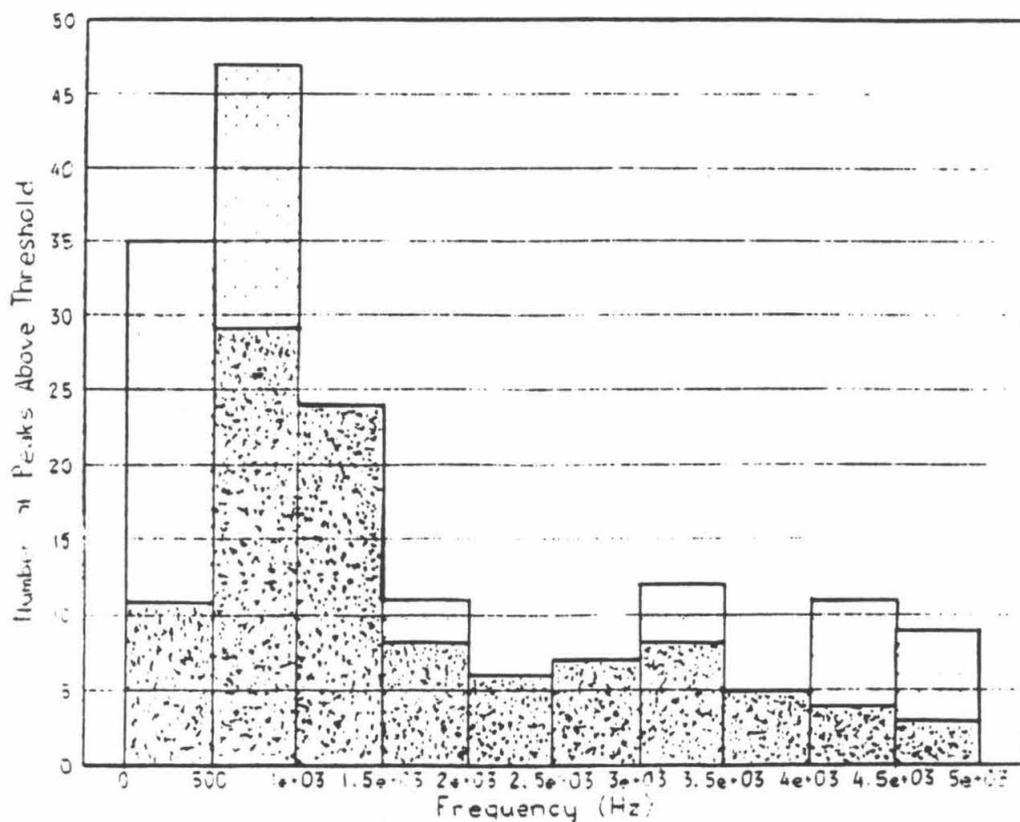


Figure 35: The number of bins above $u = 10$ found in 500-Hz portions of the spectrum taken from Tape 45(a), before (light shading) and after (dark shading) microphone and aliased line frequency vetoes had been applied. The lowest bin really extends only from 500 Hz down to 305 Hz, since the search was cut off below this. An excess of peaks at lower frequencies is still evident after applying the vetoes.

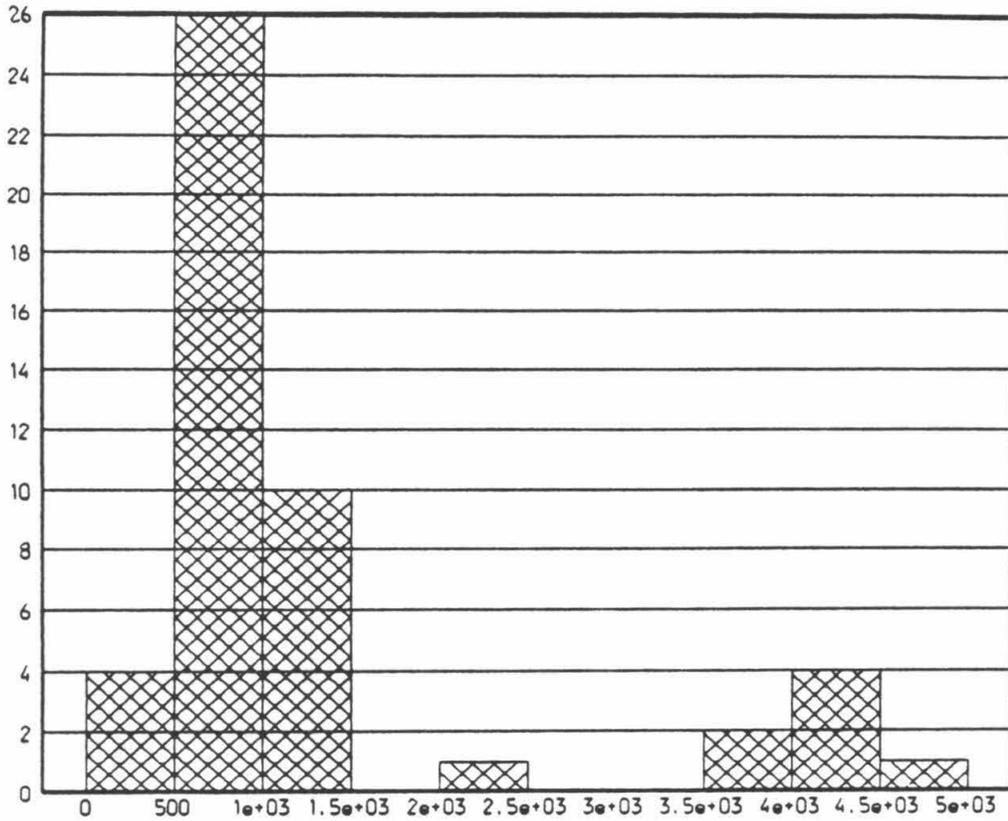


Figure 36: The number of bins above $u = 10$ found in 500-Hz portions of the spectrum taken from Tape 56, after all spurious peak vetoes had been applied. The lowest histogram bar really extends only from 500 Hz down to 305 Hz, since the search was cut off below this.

4.7 Conclusions

Since there remains an excess of peaks found at a “statistically significant” level, we present the data in a straightforward way. Based on our knowledge of the apparatus and the local nature of those peaks we did manage to identify, it seems unlikely that all the peaks are due to gravitational waves. Therefore, we proceed by assuming that *none* of them are, and attempt to place an upper limit on the size of any periodic gravitational waves based on the fact that no bigger peaks were seen. For reasons explained above, it seems logical to split up the spectrum into regions according to qualitative features such as the density and size of peaks.

4.7.1 Limits on Strain in the Antenna

The final peak lists were plotted out for each of the four spectra considered, after translation of the peak values into equivalent R.M.S. strain for a monochromatic signal at the nearest bin center frequency. A horizontal line was drawn through the largest peak in each 500-Hz band of each plot, giving separate and independent upper limits for 305–500 Hz, 500–1,000 Hz, 1,000–1500 Hz, etc. In regions where no peaks were seen above $u = 10$, the line was placed a factor of 10 (in power) above the highest level of the averaged background in that region (excluding values inflated by line spikes, etc.). The completed boundary was then uniformly raised by a factor of 2.5 (in power); the four resulting plots are presented in Figures 37 through 40.

These rather conservative (and slightly baroque) conditions would ordinarily imply a low probability, considerably less than ten percent, of having missed a genuine signal bigger than that final line. It seems inappropriate, however, to

invest much in estimating the “confidence level” of this limit based on some idealized peak height distribution; as we have seen, our limited knowledge of the process producing these spectra makes it impossible to extrapolate the probability distribution so far out on the “tail” without considerably more experimental evidence. Future analyses in conjunction with the spectra from Glasgow and MIT are virtually guaranteed more predictable aggregate statistics.

4.7.2 Interpretation; Strains in Space

The ratio of antenna strain limits to the minimum detectable amplitudes of actual gravitational waves emitted by an astrophysical source depends on the source polarization and on its position angles relative to the antenna arms, according to Equation 8. The maximum value of this factor, 1.0, obtains for the source directly overhead or directly below, with “plus” polarization; two other interesting directions are toward the galactic center and toward the Supernova 1987a remnant. The relevant angles pertinent to the Caltech antenna site and the two sky positions are summarized in Table 3. The corresponding “antenna beam” factors are tabulated in Table 4, for the two fixed polarization states defined with respect the celestial meridian of each source (as opposed to the detector axes, which rotate). A third number often stated in discussion of future detectors is the integrated average value of this factor over 4π steradians and all polarizations; this is inappropriate here, since the universe is extremely lumpy on the scale of our maximum detection range (almost exclusively sub-galactic).

Finally, the antenna strain limits calculated above were divided by the modulus of the beam factor for each tape, and the most stringent results were assembled for each polarization and for each selected source direction. The final

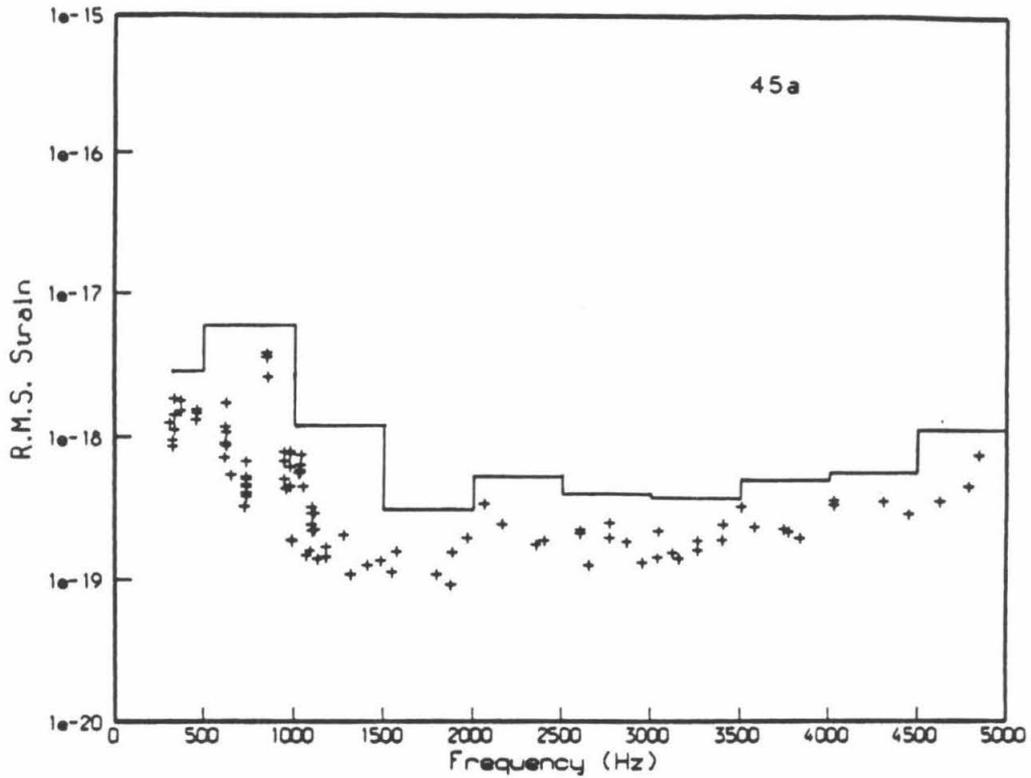


Figure 37: Experimental upper limit on periodic gravitational wave strains in the Caltech antenna for Tape 45(a). Peaks found in the spectrum above $u = 10$ (ten times the local mean power spectral density) are marked with "+" signs. The ordinate is labeled with equivalent R.M.S. strain for a monochromatic (spectrally unresolved) signal; the resolution bandwidth of these spectra is 9.54 mHz. Line spike regions and other bands vetoed from the analysis have not been shown explicitly; the total bandwidth excluded is approximately 320 Hz.

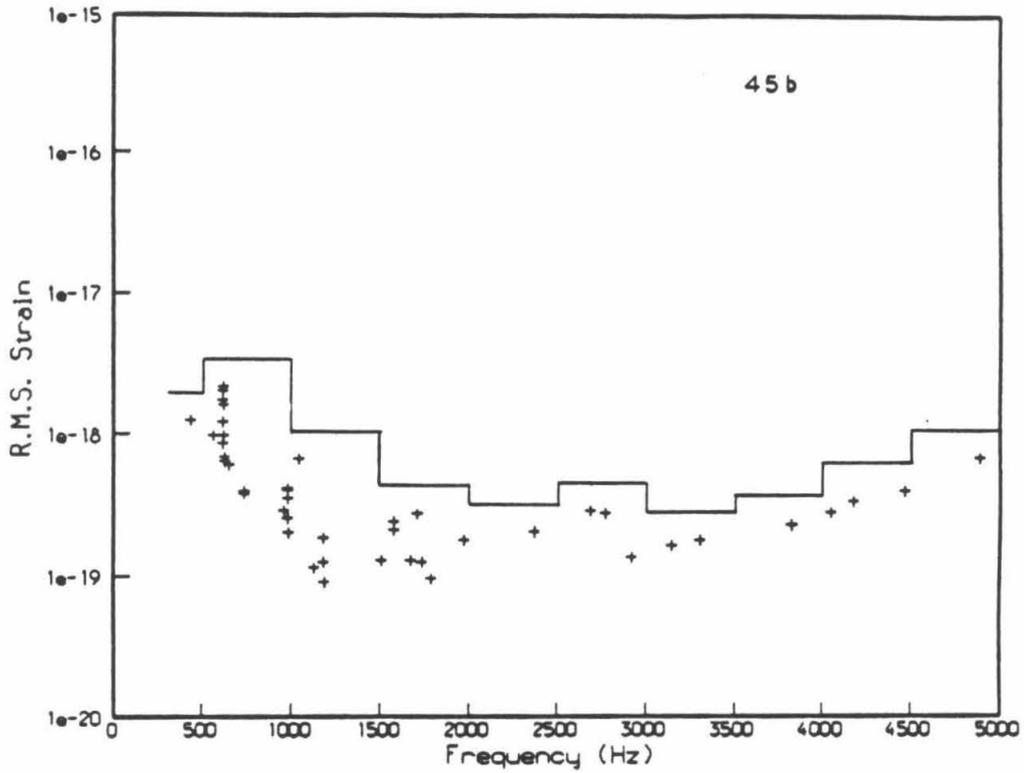


Figure 38: Experimental upper limit on periodic gravitational wave strains in the Caltech antenna for Tape 45(b), with locations and strength of peaks found in the spectrum above $u = 10$.

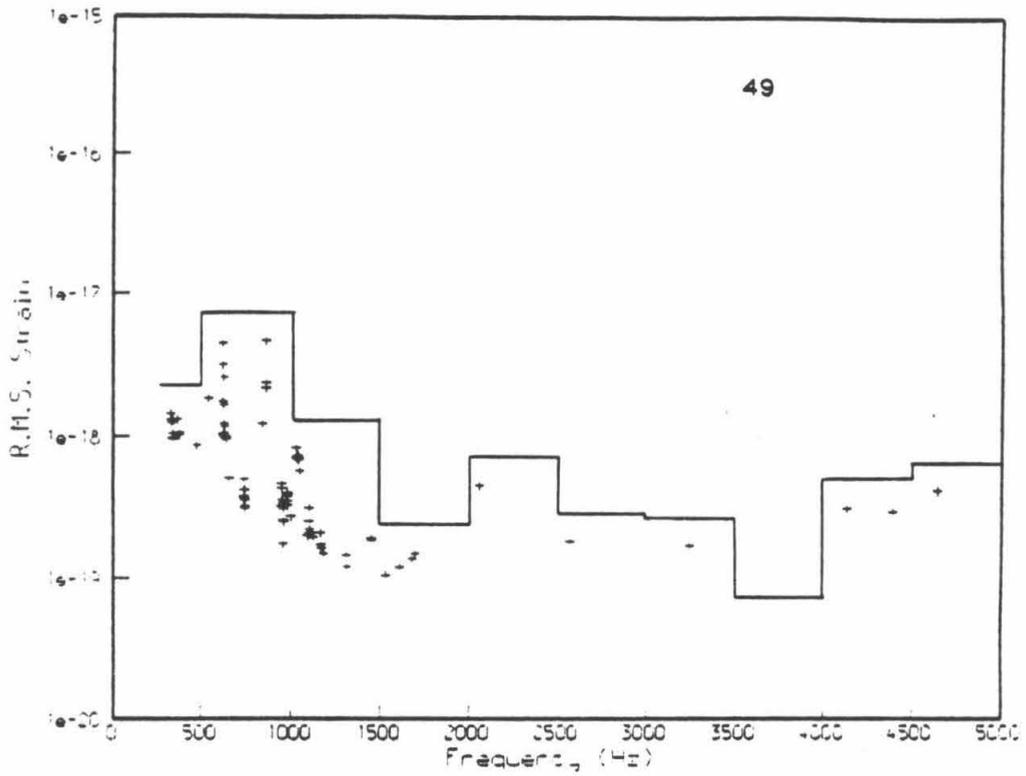


Figure 39: Experimental upper limit on periodic gravitational wave strains in the Caltech antenna for Tape 49, with locations and strength of peaks found in the spectrum above $u = 10$.

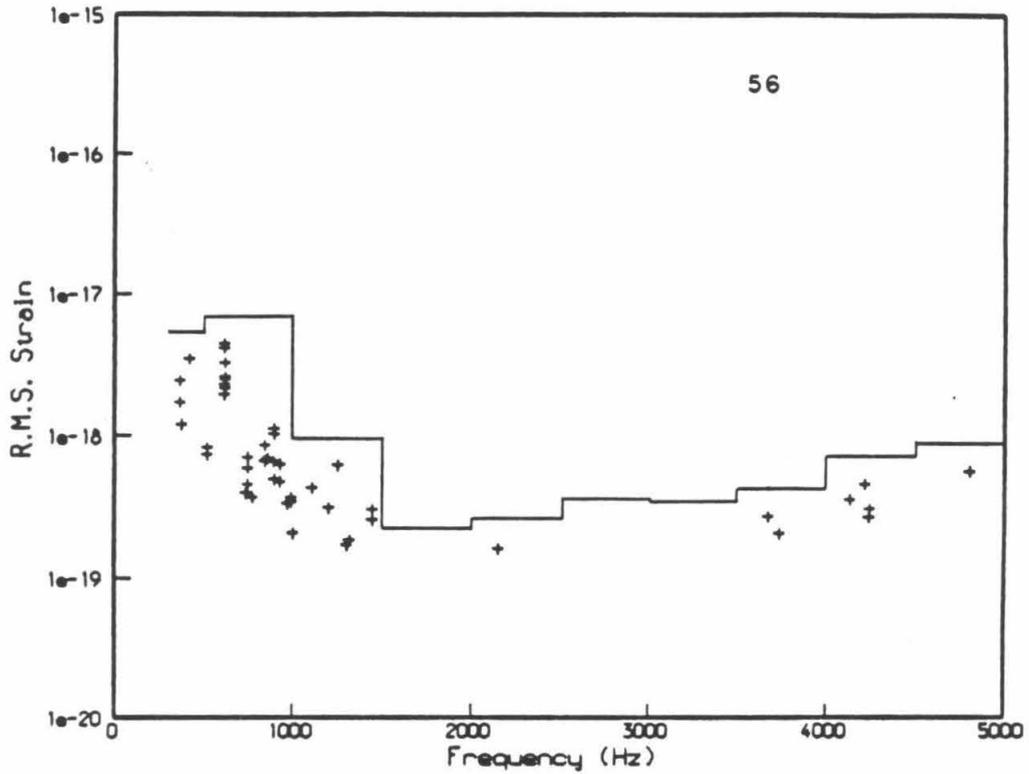


Figure 40: Experimental upper limit on periodic gravitational wave strains in the Caltech antenna for Tape 56, with locations and strength of peaks found in the spectrum above $u = 10$.

The Caltech Antenna	
latitude	$34^{\circ} 8'$
longitude	$118^{\circ} 7'$
orientation	North (Louie arm)

SN1987a (LMC)	
right ascension	$5^{\text{h}} 26^{\text{m}}$
declination	-69°
distance	~ 50 kpc

Galactic Center	
right ascension	$17^{\text{h}} 42^{\text{m}}$
declination	-29°
distance	~ 10 kpc

Table 3: The location and orientation of the Caltech 40 meter antenna, and the position angles of two points in the sky for which the beam pattern factor was calculated as a function of time during the observing run.

SN1987a (LMC)		
Tape (seg)	(+)	(×)
45(a)	0.45	-0.30
45(b)	0.47	-0.26
49	0.53	-0.11
56	-0.42	0.30

Galactic Center		
Tape (seg)	(+)	(×)
45(a)	0.92	-0.33
45(b)	0.93	-0.30
49	0.98	-0.17
56	-0.27	0.64

Table 4: Ratio of Caltech antenna strain h to gravitational wave strain amplitude h_0 , for waves emanating from the direction of the supernova remnant (top) or the center of our galaxy (below), or their diametric opposites. The polarization state “+” refers to alignment with the celestial meridian of the source; the “×” polarization state is rotated by 45° clockwise, looking skyward.

composite limits (comprising results from all four data segments) on linearly polarized astrophysical strains emanating from the two directions considered are presented graphically in Figure 41. Note that this composition implicitly assumes no change in source strength occurred between the first and last tape (a span of 53 hours). Limits free of this assumption, and for other source directions, can be derived from the beam pattern factors (Equation 8) and the “raw” antenna strain limits (Figures 37 through 40).

4.7.3 Comments

The antenna strain limits are obviously degraded by the presence of large spurious peaks in the spectrum, especially below 1500 Hz. In particular, the region between 500 Hz and 1 kHz is totally dominated in all four spectra by the suspected wire resonances mentioned previously (at 630.47, 635.96, 636.45, and 638.29 Hz) and/or a set of close peaks between 870 and 911 Hz likely to be resonances of the beamsplitter mass. If these could be shown to be spurious, as we strongly suspect they are, the strain thresholds could be lowered by at least a factor of 5 in those bands. This is an extreme example. In other regions of the spectrum the limits are also dominated by large peaks which, with sufficient diligence, could possibly be traced to local interference. However, in three of the four spectra no peaks above 1 kHz exceeded the local background by more than a factor of six in strain; this strongly implies that even if they could be eliminated the statistical background would not allow even a factor of two better sensitivity.

In addition, most such obstacles could be easily dispatched by comparison with one or more spectra from the other antennas, and in particular the Glasgow machine, which had essentially identical strain performance over much of the

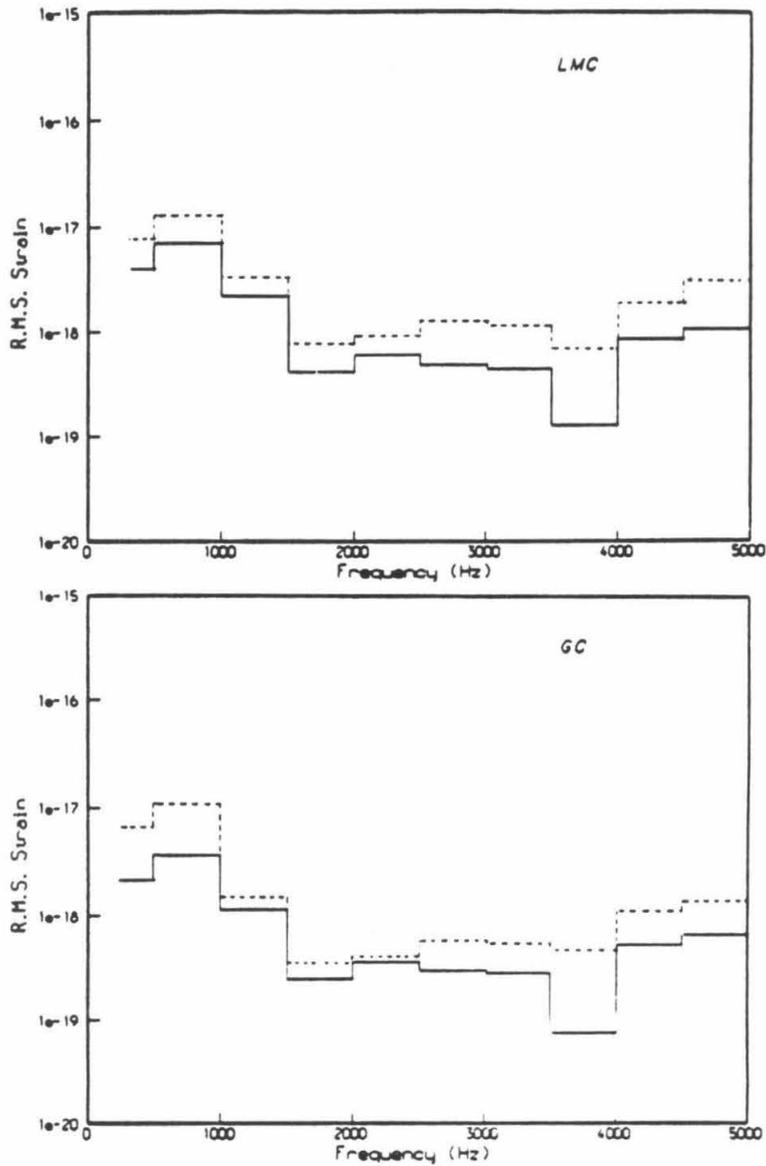


Figure 41: Composite limits on continuous gravitational waves emanating from the directions of Supernova 1987a (top) and the center of our galaxy (below). Results are shown for “+” polarization (i.e. strain aligned with celestial meridian), shown as the solid line, and “x” polarization, the broken line.

Tape (seg)	f_l (Hz)
45(a)	60.000
45(b)	60.006
49	59.995
56	60.011

Table 5: Frequencies of powerline interference in each of the four analyzed tape segments, estimated from an ensemble of line harmonic features in each spectrum with an estimated uncertainty of ± 0.005 Hz. Peaks in each spectrum within 0.95 Hz of any integral multiple (up to $n=166$) of this frequency were rejected.

spectrum. Glasgow’s latitude gives a somewhat better view of the Large Magellanic Cloud, and a worse view of the Galactic Center, but this is a small effect. Short lists of peaks above a given threshold (say, $v = 10$) are extremely unlikely to contain a coincident pair unless produced by an astrophysical signal, since the antennas are both of highly unique construction.

Another problem is the loss of a considerable fraction of the search bandwidth to line multiple interference. It may be possible to reduce the loss by including frequencies closer to each line multiple, but any peak reported too close to a multiple would have to remain suspicious. The only real solution, as mentioned above, is to eliminate line interference; in future runs this should be a priority. In the meantime, it must be stressed that our quoted limits are only valid for frequencies f satisfying $f_l n + 0.95 \text{ Hz} < f < f_l(n+1) - 0.95 \text{ Hz}$ and $10,000.95 \text{ Hz} - f_l n > f > 9,999.05 \text{ Hz} - f_l(n+1)$ for all integers $6 \leq n \leq 166$, where f_l is the line frequency. For the data tapes considered, f_l is given in Table 5.

Appendix A

Test Mass Testing

In this appendix, I describe some exploratory measurements undertaken to better understand the array of modes found in the test masses used in the interferometer before May 1984 (Figure 9). Some of the more pronounced resonances depicted in Figure 10 were identified in these tests. The mapping of their eigenmodes, in addition to providing an answer to the immediate question of their coupling to the mirror displacement, also instilled considerable intuition about mechanical structures which helped to guide subsequent design tasks.

An unused mass assembled from rejected pieces was set on four of the toy rubber cars atop an optical table. Piezoelectric disks were glued flat to its horizontal plates, employing the dilation of the disk in a plane perpendicular to the applied electric field to induce strain in the metal. Several locations were chosen on both top and bottom plates, to minimize the chance of missing any big resonances by a fortuitous placement of the driver at an antinode. The disks were driven by an audio power amplifier with up to $80 V_{p-p}$ of pseudorandom white noise (generated by an HP3582A spectrum analyzer) in order to pinpoint interesting spectral features, and then with a sine wave to investigate details of each feature.

Several methods were used to detect the vibration of the mass. A mirror

mount was attached to the lower plate holding a dummy cavity mirror. A helium–neon laser beam, reflected from this mirror onto a quadrant photodiode two meters away, provided a readout indicating the angles of altitude and azimuth. A Michelson interferometer set up to measure the linear displacement of the mirror proved unwieldy because of excess low–frequency motion of the mass on the rubber, and was abandoned. The voltages developed across the other piezos glued to the mass were also monitored. A spectrum obtained in this way is shown in Figure 42, with the measured quality factors of the more prominent modes indicated. However, only crude information could be derived about the shape of these modes and their effect on the cavity length from these measurements.

A very instructive readout was obtained with a tiny commercial accelerometer (made by Vibrametrics) which could be moved around on the mass and stuck down with doublesided adhesive tape. Its mass was small enough (< 10 grams) not to excessively perturb the modes. By driving the mass on a resonance and using a lock–in analyzer to detect the amplitude and phase of the acceleration at numerous locations, one could very quickly map out the eigenfunction of each mode. Some of the more interesting ones are depicted schematically in Figure 43. These examples should all couple strongly to the displacement of a mirror, and most also produced large signals on the quadrant diode mirror angle readout.

The dummy mass was clearly not identical to the actual interferometer masses; for example, the real masses were suspended by wires at the center post, and had several mirrors and counterweights. Two piezo disks were therefore glued onto the Huey test mass *in situ* and one was excited with pseudorandom noise while running the interferometer; the more prominent peaks in the resulting interferometer output (Figure 10) were later looked at again with the second piezo, used as a strain gauge, to measure the Q 's of the resonances. Some of these peaks were close to those of the test mass, and within the estimated uncertainty arising

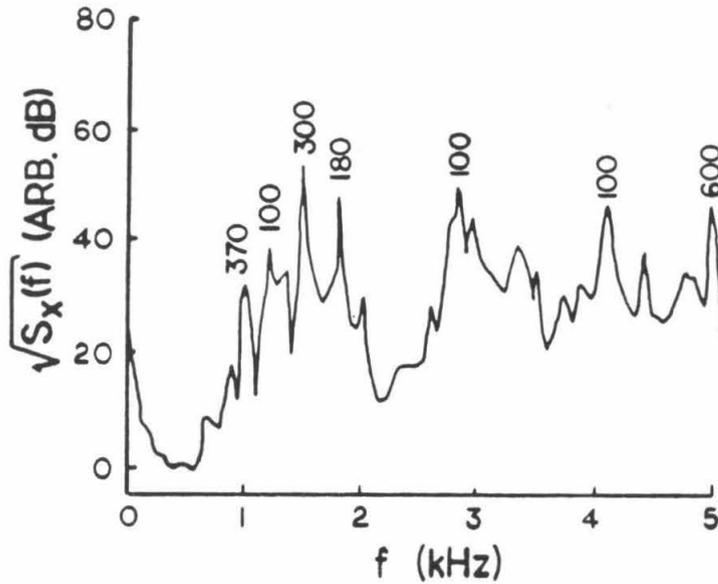


Figure 42: This transfer function between two piezoelectric strain transducers bonded to an instrumented dummy test mass like the one depicted in Figure 9 shows a large number of mechanical resonances in the band of interest. The quality factors Q of the most prominent modes were measured by exciting them individually and timing their ringdown after cutting off the excitation.

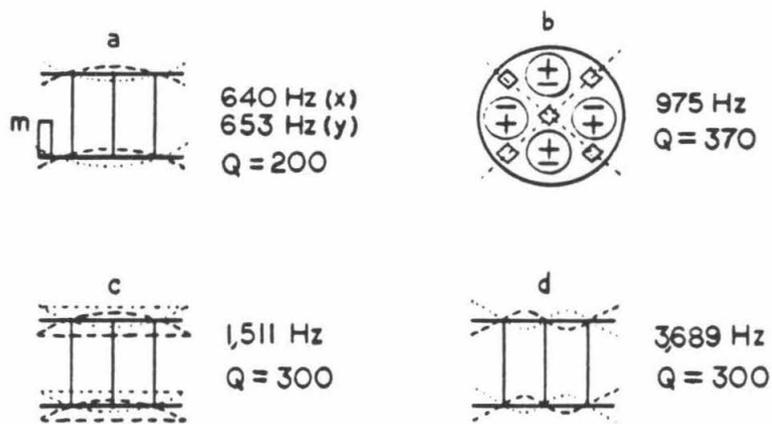


Figure 43: Each vibrational mode depicted here was mapped by exploring a dummy test mass with a small accelerometer (see text). The frequencies and Q 's of (a) and (c) would be altered for a mass suspended by wires at the center rather than sitting on soft rubber as in the tests. The effective mass and coupling to mirror displacement can be estimated from such diagrams.

from the different suspension, the lack of air damping, and the different number and arrangement of optical mounts change the frequencies and Q 's somewhat.

To estimate the contribution of thermal noise in the modes we identified we took the measured Q and frequency along with estimates of the parameters α and m_e from the accelerometer maps. In units appropriate to the 3.2 kHz resonance measured on Huey, the RPSD of strain at the peak due to thermal noise (Equation 14 with $\hat{f} = 1$) should be

$$\sqrt{S_h(f_0)} \approx 3 \times 10^{-17} \left(\frac{Q}{260}\right)^{\frac{1}{2}} \left(\frac{\alpha}{.3}\right) \left(\frac{3\text{kg}}{m_e}\right)^{\frac{1}{2}} \left(\frac{3.2\text{kHz}}{f_0}\right)^{\frac{3}{2}} \frac{1}{\sqrt{\text{Hz}}}, \quad (32)$$

matching heights of peaks near this frequency in the strain spectrum (Figure 11, upper curve) within a factor of two.

Appendix B

Principle of the Optical Phase

Detection

In this appendix we present a heuristic explanation of the modulation technique developed by Drever et al: [60] for determining the relative phase or frequency error between a Fabry–Perot optical cavity and an incident laser beam.

Returning to the pictorial language of Figure 5, recall that the light reflected from a Fabry–Perot cavity whose length is somewhere near an integral number of half–wavelengths has two dominant contributions to its electric field. One is the field promptly reflected from the input cavity mirror, \vec{E}_0 ; this is of course just a sample of the laser output at the time. The second piece, \vec{E}_1 , is light which has entered the cavity in the past. This will itself be the vector sum of components which entered $2l/c$ ago and bounced once, $4l/c$ ago and bounced three times (once at the input mirror), $6l/c$ and bounced five times, . . . , leaking away some energy with every bounce.

We want to measure the phase difference between the two components. To do this, we tag the laser field \vec{E}_0 with a distinctive signature. We impose on it a phase modulation at radio frequency ω_m . The original field is passed through an electrooptic (Pockels) cell driven by a sinusoidal voltage, effectively introducing a

sinusoidally varying time delay with period $2\pi/\omega_m$. This is depicted in Figure 44, where we have kept our frame of reference locked to the laser field *before* the modulation is imposed (i.e. our frame still rotates at a constant rate equal to the mean frequency of the laser light). Now, it is true that light entering the cavity will also have this modulation on it; but unless by lack of foresight the chosen modulating period is an integral multiple of the round-trip light travel time $2l/c$, the average effect of the phase dithering on the sum \vec{E}_1 will wash out to zero. The cavity “cleans off” the phase modulation. Looking at it briefly in the frequency domain, the phase modulation puts sidebands on the optical carrier; these sidebands are spaced ω_m away from the carrier, so that if the carrier is resonant with an optical cavity of bandwidth $\Delta\omega \ll \omega_m$, the sidebands are not resonant and are reflected. The frequency ω_m , in our case $2\pi \times 12.335$ MHz, is chosen to satisfy this condition.

With \vec{E}_1 stripped of the modulation, whatever the average phase difference is between the laser light and the cavity leakage field, a component of that phase difference will increase when our imposed “tagging” phase increases and decrease when the imposed phase decreases. Again referring to Figure 44, this means that a component in the reflected light intensity (the quantity $|\vec{E}_0 + \vec{E}_1|^2$) will fluctuate at ω_m if the mean phase difference is not precisely zero. In fact, the intensity will fluctuate in sympathy with our modulation if the mean relative phase is negative, say, and against the modulation if the mean relative phase is positive. Providing we keep track of the modulation imposed by the Pockels cell, we can decide which way and by how much the laser phase differs from the phase of the light stored in the cavity by the sign and magnitude of the fluctuations seen in the reflected light at the modulation frequency.

The reflected light is shone on a photodiode; the photocurrent is bandpass-filtered to transmit only components near ω_m . Then, the filtered current is

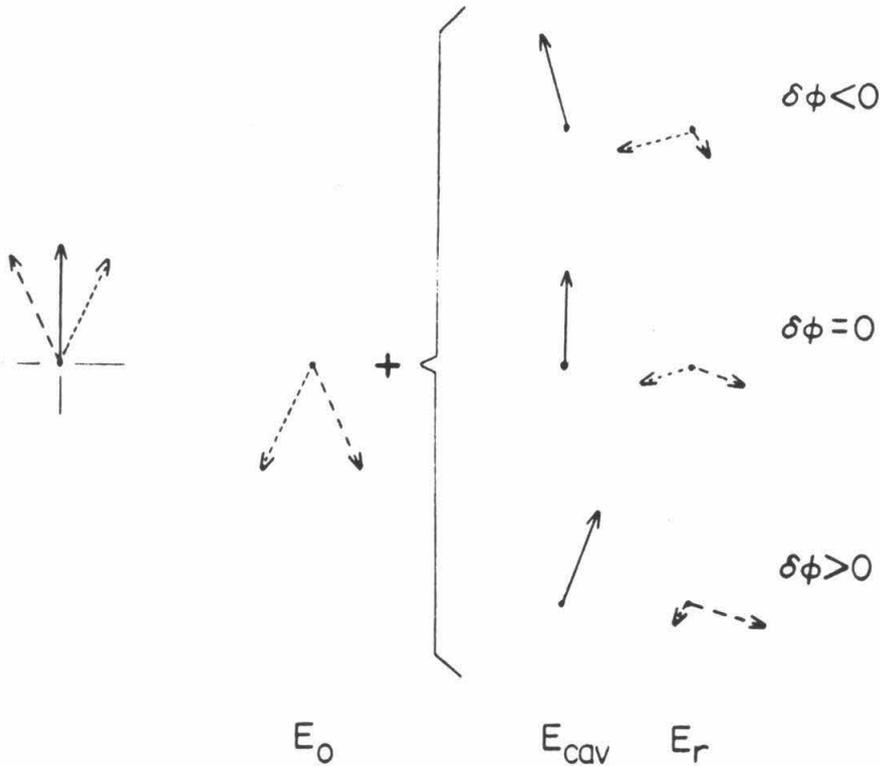


Figure 44: A phasor representation of the phase modulation fringe discriminator. The laser field (solid, at left) is phase modulated at a radio frequency between the two extremes shown (dotted and dashed to denote opposite swings of the modulating waveform). On reflection from the cavity input mirror the laser phase is shifted by π radians. This E_0 is added to returning cavity light E_{cav} ; if E_{cav} bears a retarded phase ($\delta\phi < 0$) with respect to the mean laser phase, the resultant E_r will be longer at the dotted portion of the modulation cycle than at the dashed portion (top), and the light power, which is proportional to $|E_r|^2$, will be greater in the dotted half cycle than in the dashed half cycle. If the relative optical phase is instead advanced, E_r (and the reflected power) will be larger during the dashed half cycle of the modulation (bottom); at zero relative optical phase the reflected power will be the same during symmetric swings of the modulation waveform (center), and no net fluctuation at the modulating frequency will appear. Multiplying the reflected power (detected on a photodiode) by the modulating waveform and integrating the result over several periods gives a signal proportional to the component fluctuating at the modulating frequency, and therefore proportional to the optical phase error $\delta\phi$.

multiplied (usually in a diode bridge balanced mixer) by the modulating signal itself. This picks out the piece which is coherent with the modulation, and after low-pass filtering the product provides a pseudolinear representation of the desired optical phase difference. The result is presented in Figure 45, showing the reflected light intensity and the mixer output voltage obtained by scanning a laser through the resonance of a small cavity.

A final, but critical subtlety; because the cavity stores light from the past, there is a difference in its response to changes of laser phase depending on how rapidly they occur. For changes that are sudden with respect to the storage time, the cavity will not reach equilibrium with the laser; but for gradual changes (like the scan which produced the picture in Figure 45), the light inside the cavity will tend to catch up in phase. The sensing system works differently in the two regimes. For short durations, the demodulated signal follows the instantaneous phase difference; over longer durations, the signal reflects the *rate of change* of the phase difference, i.e., the *frequency* difference. The effect is the same as that of a simple RC lowpass network, with time constant $RC \approx 2\tau_e$, where τ_e is the cavity storage time (see Equation 13).

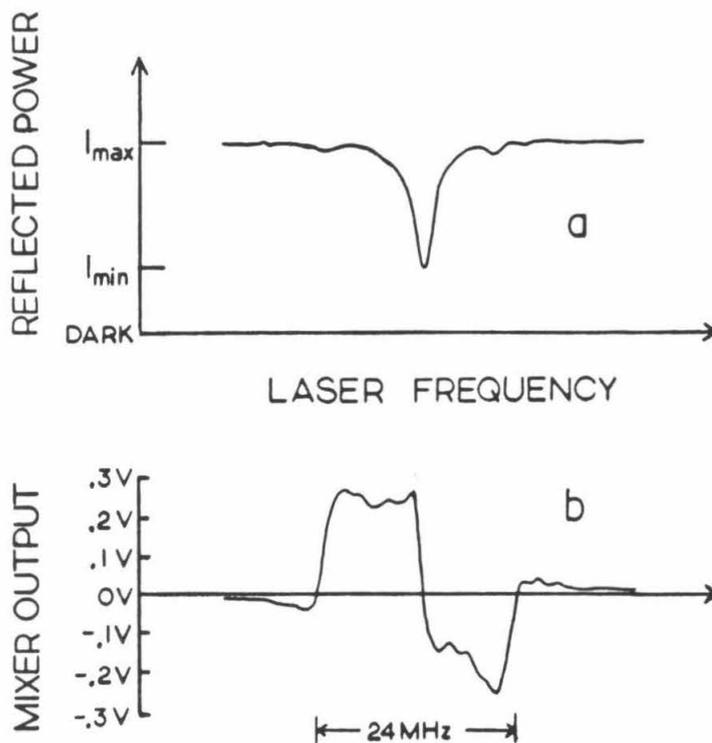


Figure 45: Picture captured as the laser was scanned through the resonance of a 30 cm cavity. Trace (a) is the total photocurrent detected from light reflected from the cavity, showing the dip in reflection when the laser passes through resonance. The ratio $K = (I_{max} - I_{min})/I_{max}$ is the fringe contrast. Trace (b) is the filtered output of the RF mixer; near the resonance the voltage is roughly proportional to the frequency difference between the cavity and the laser. A tiny dip in the total light on either side of the main resonance is just visible, with corresponding phase signals in the mixer output. These are the resonances of the cavity with the modulation sidebands themselves, which are 12 MHz away from the optical carrier (setting the frequency scale of the picture). In the 40 meter cavities, which are equipped with much better mirrors, the resonances are 3,000 times narrower. Note that the phase signal has an opposite slope at the sidebands; this keeps a servo from accidentally locking onto a sideband instead of the carrier.

Appendix C

Burst and Coalescing Binary Sources

C.1 Short Bursts of Gravitational Radiation

Burst sources of gravitational waves have been sought most often by experiments to date, and the early history of the field was largely driven by an explicit or implicit assumption that bursts from supernovae were most likely to provide the first confirmed detection. We can define as a burst any waveform $h(t)$ which is zero before some initial instant, has appreciable amplitude for only a few cycles, and then dies off. Fitting this category are the wide variety of waveforms postulated to emanate from stellar collapse to a neutron star or black hole [1,80,81], or infall of material into a black hole [82]. At present there is no firm consensus regarding how strong these waves are or how frequently they might be generated in a given volume of space.

As far as we know, no detailed mechanism has been proposed by which a supernova remnant might emit a large burst of gravitational radiation days or weeks *after* its collapse. While this ignorance could well be interpreted as an incentive to look for such events, we will for now assume that detectable bursts belong to the class of signals not associated with the supernova, and thus having no particular *a priori* reason to have occurred just when we were looking. Thus,

we address chances for detecting signals at strengths and/or event rates not ruled out statistically by previous experiments.

The best published limit to date on gravitational wave bursts has been set by observations with the Stanford University cryogenic bar antenna [66]. In comparing their limit, which is expressed in terms of energy deposited in the bar’s fundamental mode of vibration, to the detectability threshold of a broadband interferometer it is necessary to specify the waveform under consideration. This is of course not generally possible, so investigators have proposed model “canonical forms” for bursts, varying a finite number of model parameters (e.g. amplitude, number of half-cycles, and characteristic frequency) to span an infinity of believable theoretical waveforms [84,86,87]. The approach has its pitfalls. The energy deposited in a narrowband bar¹ is essentially determined by the value of the impinging signal’s power spectrum evaluated at the bar frequency. Sharp discontinuities in the canonical waveform or its derivatives, taken for simplicity of calculation, can serve to redistribute the pulse energy over a broader spectral band than the comparatively smooth natural waveform being approximated. This causes an overestimation of the bar’s response to bursts having characteristic frequencies far from the bar frequency, analogous to the spectral leakage seen in Fourier analysis of data which have been too sharply “windowed”.

A further subtlety not addressed in some comparisons of bars and interferometers is also related to the bandwidth issue. When translating the quoted mean noise level of an instrument into a statistically measurable rate of astrophysical events, it is important to define the time interval over which the detector output has some acceptably small probability of exceeding a certain threshold in the absence of any signal. Narrowband bar detectors, when filtered optimally,

¹The cryogenic resonant bar detectors operated so far have had bandwidths of order 1% or less of their resonant frequency [88,89].

are interrogated about once every tenth of a second to one second, almost independent of the expected signal duration; the quoted mean noise level is on average exceeded once in every e such intervals. By contrast, an interferometer has time resolution of order $\Delta t \approx \tau_c/\pi\text{SNR}$ where τ_c is the characteristic length of the burst, and SNR is its detected signal-to-noise ratio. The broadband experimenter pays for knowing the exact pulse arrival time by having to quote a mean noise level that is actually exceeded once in e *milliseconds* (for bursts of that length). As Thorne [1] points out, if the detector noise is perfectly Gaussian or if one combines an ensemble of independent detectors such that the aggregate noise approaches a Gaussian distribution, this merely requires setting a slightly higher threshold for the interferometer to reach the desired “accidental detection” rate, say three false alarms per year. Single detectors tend to have significant excesses of large events over the Gaussian model [37,66,86], with typical spurious event rates falling much more slowly than an exponential with increasing event energy. This can leave a longer way to go for the “broadbander” to achieve an acceptable false alarm rate.

With these caveats in mind a look at the Stanford results shows a tough mark to match. I used one of the better-behaved short pulse waveforms considered by Dewey² [86], and varied the characteristic timescale of the burst (not adhering to any particular model, but once again as a convenient parametrization) between .25 and 2 milliseconds. The published bar event rate limits have been replotted against the approximate R.M.S. signal amplitude h_0 they represent for this parametric family of burst waveforms (Figure 46). Also plotted is the projected event rate seen by a hypothetical broadband detector whose noise is white between 500 Hz and 4 kHz (and very high elsewhere), with a Gaussian probability distribution,

²The waveform is derived from a numerical simulation of a particle falling into a rotating black hole performed by Kojima and Nakamura [82].

and with a mean noise level given by $S_h(f)^{1/2} \approx 5 \times 10^{-19} / \sqrt{\text{Hz}}$. Nearly Gaussian noise at this level might conceivably be achieved in a coincidence analysis of data from the Glasgow and Caltech interferometers as they were operating in March of 1987, but the curve should not be taken too seriously. It really represents a “best case” scenario.

The calculations indicate that the interferometer data would not place any new limits on the rate of occurrence of bursts having this particular form. Had the observing run been extended for some weeks, it might have been possible with coincidence analysis to explore some new territory.

Choosing a different waveform with more cycles at the characteristic frequency, i.e. a higher Q , affects the analysis by narrowing the spectrum of the signal. Bursts like this could be more likely to have “missed” exciting the Stanford bar if at frequencies far from the bar frequency, while registering in a more broadband detector. At this point, however, the limited event rate floor imposed by the short coincidence time seems discouraging.

C.2 Coalescing Compact Binaries

The only type of gravitational wave signal with a highly predictable waveform, a known strength, and a fairly predictable event rate is the “chirp” generated by the inspiral and coalescence of a compact binary system comprising black holes and/or neutron stars [1]. The unique form of the signal may also allow accurate and independent determination of the Hubble constant, if the source can be identified with an optical display or if a large sample of waveforms can be examined statistically [6]. Planned large-scale gravitational wave observatories [90] may be able to see such events occurring anywhere in the observable universe.

The waveform, which is described extensively elsewhere [8,37,87,91], begins

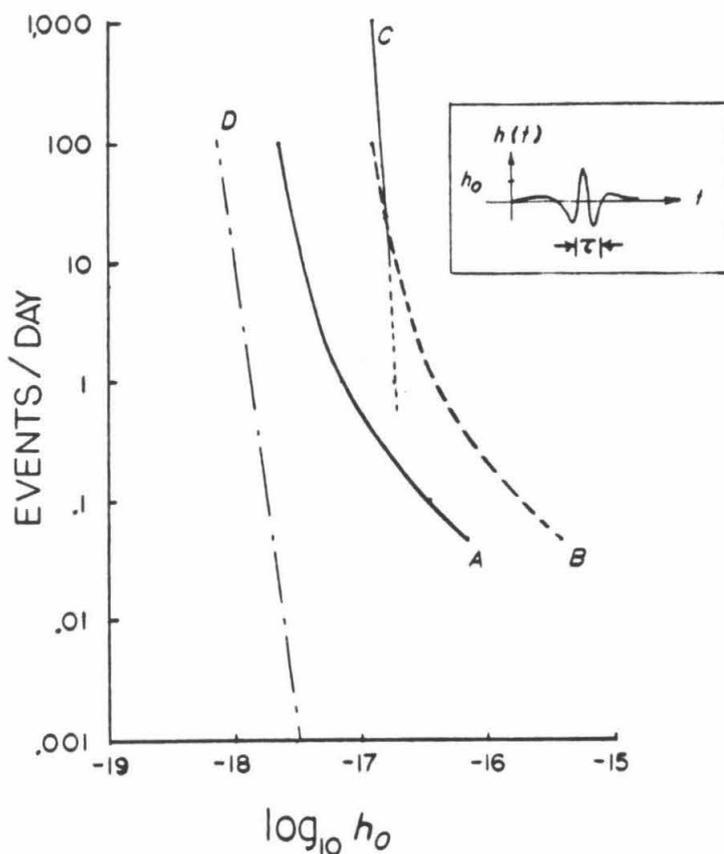


Figure 46: Published upper limits on the rate of gravitational wave bursts reported by Boughn et al. [66] for the Stanford bar experiment are here reinterpreted as R.M.S. strain for the class of waveforms $h(t)$ depicted (really a waveform calculated by Kojima and Nakamura [82] for a specific burst mechanism, pressed into service as a general-purpose short pulse). The timescale τ of the burst is varied from .25 to 2 milliseconds; at both extremes the Stanford results place a limit on burst event rates as a function of h_0 shown in curve B. The bar is optimally sensitive to bursts with timescale equal to the bar's period (1.2 ms); the limit for such bursts is shown as curve A. Line C represents a hypothetical limit which could conceivably result from coincidence analysis of the Caltech and Glasgow data from the 1987 observations (it is nearly independent of the burst timescale); however, since the coincident observations spanned at most a few hours, event rates below 10 per day (dotted portion of curve C) are inaccessible. For fairness, the limit quoted in [66] for a coincidence search with two similar bars is also interpreted for $\tau = 1.2$ ms (D); recent coincident observations (now being analyzed) between comparable bars at LSU, Stanford and CERN could probe this range.

as a sinusoid at the frequency of the nearly classical initial orbit. Gravitational radiation reaction removes orbital energy and angular momentum, and the amplitude and frequency increase with time as the objects spiral closer together and revolve faster. The final burst, when the objects are tidally disrupted, is likely to be very complicated for neutron stars, but for black holes numerical simulations should be able to calculate the waveform. From the detection standpoint, however, relatively little of the total emitted energy is left at the coalescence; the lion's share will have been radiated away during the long inspiral. Strategies are therefore being developed [37,92,93] which concentrate on filtering the outputs of broadband detectors to follow the rising frequency and integrate over many cycles, taking advantage of the pseudo-coherent release of orbital binding energy over a long period. The time variation of the signal depends on a single “mass parameter,” which is the reduced mass of the binary system multiplied by the total mass raised to the $2/3$ power. This parametrization, as well as some physical insight as to the maximum gravitational wave frequency achieved before the system breaks up (e.g., about 500 Hz for two neutron stars or 10 solar mass black holes [1]), allows filtering for all possible mass combinations simultaneously in real time [93].

Once again, there seems to be no compelling reason to link this specific mechanism with the aftermath of a supernova; indeed, the known size of this kind of signal essentially puts the location of the supernova out of range of our apparatus³. The Stanford search of 1982 [66] places the strictest published limit on the rate of these signals as well as for short bursts. Following Dewey's estimates of bar and interferometer response to these signals, we can compare the bar and interferometer signal-to-noise ratios for a given source distance and

³The range of the search performed, about 25 parsecs for neutron star pairs, is limited by factors discussed in [37]. It encompasses a volume highly unlikely to contain enough such objects to expect a coalescence during the observations.

mass parameter. The low-frequency limit of the interferometer is important; as pointed out above, the bulk of the energy is emitted at low frequencies (the signal power spectrum falls as $f^{-7/3}$). Reducing the low-frequency cutoff frequency to 300 Hz instead of 500 Hz would increase the collected signal energy by almost a factor of two, allowing detections $\sqrt{2}$ farther away and encompassing $2^{3/2} \approx 2.8$ times more volume in the “detectability envelope.” In practice, of course, the seismic noise which forces the cutoff has a complicated spectrum, and the filter should be given an optimal shape to account for this. For this observing run, the coalescing binary search performed by Sheryl Smith included frequencies as low as 300 Hz, although this was not optimal for all mass parameters considered [94].

For very massive black hole pairs ($M_1 \sim M_2 \gtrsim 10M_\odot$) there might be negligible excitation of an 850 Hz bar detector if the coalescence itself doesn’t produce a large high-frequency burst, while a broadband detector sensitive at 500 Hz might still pick something up. This is one possibility not ruled out by the general-purpose 1982 Stanford search, although it is admittedly a narrow one. Once again, however, the very limited length of the observing run combined with the small volume of the universe observable at this sensitivity make it unlikely that such a binary might go off within range in such a short time. Even assuming a distribution of compact binary masses “designed” to evade the 850 Hz bar frequency, one or more coalescences per few hours within our own galaxy is irreconcilable with the observed galactic matter density [1].

Appendix D

Data Acquisition

Some technical details of the data acquisition and timing systems are presented in this Appendix. Relatively little work could be spent on refining the hardware and software, since (as we mentioned in Chapter 4) the system was assembled in a mad hurry. This probably makes it all the more important to describe it fully, as things were not generally done in a way someone with more time to think would have done them.

D.1 Signal Preparation

The data acquisition computer, a Masscomp MC-500 equipped with 12-bit multiplexed ADC's and a single 1600 bpi magnetic tape drive, was programmed to digitize analog signals on each of five input channels sequentially; the five-sample sets were recorded into memory buffers which were sequentially dumped to the tape drive in raw 16-bit binary integer form. No headings, markers, or other frills were recorded. Simultaneously a wider array of auxiliary signals was recorded on an eight-channel analog chart recorder, using "pulse stretching" circuitry to insure registration of impulsive events in the broadband signals.

The actual strain signal recorded was a heavily filtered version of the feedback

voltage applied to the secondary cavity piezoelectric mirror. This signal reflects any differential length change between the two cavities in that it is an applied correction, and is relatively independent of various electronic and optical transfer function characteristics providing the gain of the secondary cavity loop is kept high (a condition satisfied for much of the frequency spectrum to be considered). The actual signal used for routine diagnostic monitoring of the interferometer, derived from an Ithaco 1201 FET-input preamp sensing a high-impedance tap on the HV mirror drive, was further filtered by a commercial Rockland 452 active filter set to pass frequencies between 400 and 4000 Hz with the equivalent of two cascaded 4-pole Bessel filters. The schematic is shown in Figure 27 . With our typical sampling rate of 10k (samples/sec)/channel, the Nyquist rate of 5 kHz is rather close to the corner frequency of the lowpass filter section, and so some aliasing can be expected to have occurred at the high end of the frequency spectrum. This potential ambiguity was accepted in order to cover as much bandwidth as possible.

Four other analog signals were also digitized and recorded concurrently on the computer tape. The high-voltage laser feedback from the primary cavity/laser servo loop served as a monitor of laser misbehavior, mode hops, mechanical or acoustic disturbances, or temporary overloading "glitches" as the primary servo approached its dynamic range limits. The output of a Ranger seismometer, placed underneath the vertex optical table and oriented to sense vertical accelerations, was amplified and filtered; this low-frequency signal was added to the high-pass filtered amplified output from a microphone, hung by its cord above the laser table, and the sum was recorded on a third channel. A filtered version of the output of a WWV VLF receiver (HP117A) was recorded on the fourth channel, as a coarse absolute timing reference. Finally, the light power falling on the secondary cavity photodiode was recorded, monitoring the fringe contrast on

this cavity and thus telling when both cavities were locked to fringes¹.

All these signals were digitized at equal intervals and at the same rate. All but the strain signal and perhaps the microphone diagnostic signal require minimal bandwidth for faithful reproduction, and would have served equally well had they been recorded at much lower rates. This would in turn have allowed recording other signals which might be used as added diagnostics or event vetoes while still maintaining the overall data rate below the limits set by our hardware, or alternatively, might have reduced the outrageous speed with which we filled up tapes, but the programming for such an efficient algorithm was too complex to complete and debug in time.

To partly satisfy the need for additional monitoring of the interferometer performance, and also for overall redundancy, a multichannel analog chart recorder running at 1 cm/minute was supplied with all the signals going to the computer (except the WWV timing) and several others. The laser intensity falling on an auxiliary photodiode served to normalize the fringe contrast on both secondary and primary cavities (also recorded). A differently filtered version of the secondary cavity mirror feedback, emphasizing low-frequency drifts and trends, provided insight into the overall expansion and contraction of the laser table and the interferometer arms. The microphone and strain signals were processed by special pulse-stretching circuits² intended to make sure that rapid transients would still show up on the trace in rough proportion to their peak amplitude.

¹Losing lock on the primary cavity automatically throws the secondary out, as it is unable to track the much larger fluctuations in the unstabilized laser wavelength.

²Essentially a full-wave rectifier driving a fast peak hold circuit, which is arranged to have a droop time longer than the pen response time of the chart recorder. The unit was designed and built by colleagues in Glasgow.

D.2 Timing

A complete clock and WWV receiver system was put together very quickly for us by Robert Carlson, with some equipment borrowed from the Jet Propulsion Laboratory and a handbuilt counting circuit. The ADC sampling was synchronized by an HP 105B oven-stabilized quartz oscillator. The 1 MHz output of the oscillator was counted down by the counting circuit which gave out a TTL pulse every n oscillator periods, with n adjustable. The circuit was externally gated by one line of the BCD time code output from a Systron-Donner 8120 Time Code Generator, which was itself clocked by the same quartz oscillator. The gating was rigged so that at 20-second intervals, every odd ten seconds, the countdown circuit was enabled and (if cleared and armed by the operator in the last 20 seconds) started giving out clocking pulses at the selected intervals. Each pulse triggered the ADC to digitize one sample on one channel, and the channels were addressed cyclically, starting with the gravitational wave strain channel. During most of the observing run the countdown circuit was set to $n = 20$, giving an aggregate sample rate of 50,000 samples/second, or 10,000 (samples/second)/channel.

The time code generator was manually set to within 50 milliseconds of Universal Time by use of a WWV receiver tuned to the 10 or 15 MHz WWV broadcasts. After the run the absolute time and the rate of the quartz oscillator were checked more accurately using a portable cesium beam standard; during the run the drift of the oscillator was also constantly checked by long-term phase comparison with the 60 kHz VLF WWV broadcast signal, using the HP 117A VLF Phase Comparator. The adjustments indicated by these comparisons are far below the level of precision required by any of the analyses discussed here, but for

future searches these data can probably be corrected to 100 microsecond absolute timing accuracy if needed.

The starting time (time of the first sample on the first channel) of each tape was recorded by the operator from the display on the time code generator; the known sampling interval kept time from that point onward for that tape. The system was briefly checked for missed samples and other obvious errors. The WWV VLF signal recorded on one of the tape channels served as an independent check for gross timing errors. A spot check of one tape showed this signal was continuous and in agreement with the sampling interval.

At the inefficient data rate quoted above, each reel of tape lasted only 7 minutes 13 seconds. Although after some difficulty we were able to borrow a second tape drive, our computer persistently rejected the implant and we were unable to get two drives working in complementary fashion. Hence we were forced to accept some “dead time” while rewinding each tape and mounting a fresh one. After some initial training, we were able to do this in about 135 seconds (largely limited by the tape drive rewind time) for an overall duty cycle of 70% . It should be mentioned that the other major source of dead time in this run, loss of fringe lock in one or both cavities of the interferometer, occurred during 20 to 40% of the observing time, depending on a large number of environmental factors. Both losses are associated with unfortunate, though far from inevitable, peculiarities of the apparatus, and can be made insignificant with modest attention in the future.

D.3 Calibration

Uncertainties in the secondary cavity servo loop gain may introduce complications in the straightforward interpretation of the mirror feedback signal as a strain. A major cause of fluctuating servo gain was the changing fringe contrast

in the cavity, caused by alignment variations and exacerbated by the distorted condition of the mirrors. During this observing run, the servo unity-gain frequency probably fluctuated between 3 and 4 kHz, leading to some variation in the calibration above 2 kHz. The actual transfer function was occasionally checked by exciting the calibration piezo transducer on Huey (end mass of the primary cavity) to induce an artificial strain. The piezoelectric constant of this piezo had been measured to about 1 dB accuracy before the mass was installed in 1984, and several independent checks have since verified that this value was stable over time. A further concern, that the piezo might not respond linearly at the extremely small displacements (of order 1 Fermi R.M.S.) used to calibrate the antenna, was answered directly by crosschecks against the magnetic force transducers installed subsequently. At intervals during the observing run several different test waveforms were applied to this piezo, including sine waves at various frequencies, step functions, and broadband pseudorandom noise. This allowed accurate compensation for the frequency response of the readout.

References

- [1] K.S. Thorne, “Gravitational Radiation” in *300 Years of Gravitation*, eds. S.W. Hawking and W. Israel. Cambridge University Press, Cambridge (1987).
- [2] David Helfand, “Bang: The Supernova of 1987,” *Physics Today* **40** (8), p. 24 (1987).
- [3] Masa-Toshi Koshiha, “Observational Neutrino Astrophysics,” *Physics Today* **40** (12), p. 38 (1987).
- [4] J.M. Weisberg and J.H. Taylor, “Observations of Post-Newtonian Timing Effects in the Binary Pulsar PSR 1913+16,” *Phys. Rev. Lett.* **52** (15), 1348 (1984).
- [5] J.H. Taylor and J.M. Weisberg, “A New Test of General Relativity: Gravitational Radiation and the Binary Pulsar PSR 1913+16,” *Ap. J.* **253**:908–920 (1982).
- [6] Bernard F. Schutz, “Determining the Hubble constant from gravitational wave observations,” *Nature* **323** (25), p. 310 (1986).
- [7] Rochus E. Vogt, Ronald W.P. Drever, Kip S. Thorne, and Rainer Weiss, *Caltech/MIT Project for a Laser Interferometer Gravitational Wave Observatory*, Renewal Proposal to the National Science Foundation. California

- Institute of Technology and Massachusetts Institute of Technology (1987).
- [8] C.W. Misner, K.S. Thorne, and J.A. Wheeler, *Gravitation*. W.H. Freeman, San Francisco (1973).
 - [9] Massimo Tinto, in preparation (1988).
 - [10] Y. Gursel, P. Linsay, P. Saulson, R. Spero, R. Weiss and S. Whitcomb, “The Response of a Free Mass Gravitational Wave Antenna,” unpublished communication (1983).
 - [11] R.L. Forward, “Wideband laser–interferometer graviational[sic]–radiation experiment,” *Phys. Rev. D* **17**, p. 379 (1978).
 - [12] F.B. Estabrook and H.D. Wahlquist, *General Relativity and Gravitation* **6**, p. 439 (1975).
 - [13] B.F. Schutz and M. Tinto, “Antenna patterns of interferometric detectors of gravitational waves—I. Linearly polarized waves,” *Mon. Not. R. Astr. Soc.* **224**, p. 131 (1987).
 - [14] J.H. Laning and R.H. Batten, *Random Processes in Automatic Control*. McGraw Hill, New York (1956).
 - [15] K. Maischberger, “The Gravitational Wave Detection Experiment in Frascati,” in *Colloques Internationaux du C.N.R.S. No. 220, Ondes et Radiations Gravitationelles*. C.N.R.S., Paris (1973).
 - [16] William O. Hamilton, “Near Zero: Force, Force Gradient, Temperature. Cryogenic Bar Detectors of Gravitational Radiation (*A Primer for Non-Specialists*), preprint (1982).

- [17] David G. Blair, “Resonant Bar Detectors for Gravitational Waves,” in *Gravitational Radiation*, eds. N. Deruelle and T. Piran. North Holland, Amsterdam (1983).
- [18] R.W.P. Drever, “Interferometric Detectors of Gravitational Radiation,” *Workshop on Gravitational Radiation*. Ecole d’Ete de Theoretique, Les Houches, France (1982).
- [19] Carlton M. Caves, “Quantum–Mechanical Radiation–Pressure Fluctuations in an Interferometer,” *Phys. Rev. Lett.* **45** (2), p. 75 (1980).
- [20] Carlton M. Caves, “Defense of the Standard Quantum Limit for Free–Mass Position,” *Phys. Rev. Lett.* **54** (23), p. 2465 (1985)
- [21] C.M. Caves, in *Quantum Measurement and Chaos*, ed. E.R. Pike. Plenum, New York (1987).
- [22] V.B. Braginsky and K.S. Thorne, “Gravitational–wave bursts with memory and experimental prospects,” *Nature* **327** (6118), p.123 (1987).
- [23] A.A. Michelson and E.W. Morley, “On the Relative Motion of the Earth and the Luminiferous Ether,” *American Journal of Science* **XXXIV** (203), p. 333 (1887).
- [24] R. Weiss, “Electromagnetically Coupled Broadband Gravitational Antenna,” *RLE Quarterly Progress Report 105*. Massachusetts Institute of Technology, Cambridge, p. 54 (1972).
- [25] R.W.P. Drever, G.M. Ford, J. Hough, I.M. Kerr, A.J. Munley, J.R. Pugh, N.A. Robertson and H. Ward, “A Gravity–Wave Detector Using Optical Cavity Sensing,” 9th International Conference on General Relativity and Gravitation at Jena, GDR (1980).

- [26] D.R. Herriott and H.J. Schulte, “Folded Optical Delay Lines,” *Applied Optics* 4 (8), p. 883 (1965).
- [27] R. Benford, M. Burka, N. Christensen, M. Eisgruber, P. Fritschel, A. Jeffries, J. Kovalik, P. Linsay, J. Livas, P. Saulson and R. Weiss, “Interferometric Gravitational Wave Detection at MIT,” 13th *Texas Symposium on Relativistic Astrophysics*, ed. M. Ulmer. World Scientific, Singapore (1986).
- [28] D. Shoemaker, R. Schilling, L. Schnupp, W. Winkler, K. Maischberger, and A. Rüdiger, “Noise behavior of the Garching 30 meter prototype gravitational wave detector,” Max-Planck-Institut für Quantenoptik manuscript MPQ 130 (1987).
- [29] R.W.P. Drever, J. Hough, W.A. Edelstein, J.R. Pugh, and W. Martin, “On Gravitational Radiation Detectors Using Optical Sensing Techniques” in *Gravitazione Sperimentale*, ed. B. Bertotti. Academia Nazionale dei Lincei, Rome (1977).
- [30] M. Born and E. Wolf, *Principles of Modern Optics*. Pergamon, New York (1980).
- [31] V.B. Braginsky and A.B. Manukin, *Measurement of Weak Forces in Physics Experiments*, ed. David H. Douglass. University of Chicago Press, Chicago (1977).
- [32] A.E. Siegman, *An Introduction to Lasers and Masers*. McGraw-Hill, New York (1971).
- [33] Stanley E. Whitcomb, “Shot Noise in the Caltech Gravitational Wave Detector—The Mid-1984 Configuration,” unpublished communication (1984).

- [34] H. Billing, W. Winkler, R. Schilling, A. Rüdiger, K. Maischberger, and L. Schnupp, “The Munich Gravitational Wave Detector Using Laser Interferometry,” in *Quantum Optics, Experimental Gravity, and Measurement Theory*, eds. Pierre Meystre and Marlan O. Scully. Plenum, New York (1983).
- [35] Andrej Čadež, “Internal Scattering in Fabry–Perot Interferometers,” in preparation (1987).
- [36] S.E. Whitcomb, “Optical Pathlength Fluctuations in an Interferometer due to Residual Gas,” unpublished communication (1984).
- [37] Sheryl L. Smith, *A Search for Gravitational Waves from Coalescing Binary Stars Using the Caltech 40 Meter Gravity Wave Detector*. Ph.D. thesis, California Institute of Technology (1988).
- [38] Mark Hereld, *A Search for Gravitational Radiation from PSR 1937+214*. Ph.D. thesis, California Institute of Technology (1983).
- [39] Ralph Wolf, *Transfer Characteristics of Mechanical Filters Used in Caltech’s Gravity Wave Antenna*, undergraduate research thesis. California Institute of Technology (1987).
- [40] Gary M. Gutt, “Report on Vibration Isolation Measurements,” unpublished research report (1985).
- [41] James M. Kaufman, private communication (1985).
- [42] Joseph Clapp, “Geometrical Fluctuations in a Fabry–Perot Interferometer,” unpublished communication (1980).
- [43] Andrej Čadež, private communication (1987).

- [44] A. Rüdiger, R. Schilling, L. Schnupp, W. Winkler, H. Billing and K. Maischberger, "A mode selector to suppress fluctuations in laser beam geometry," *Optica Acta* **28** 5, p. 641 (1981).
- [45] Jefferson Harman, private communication (1987).
- [46] EG&G Princeton Applied Research Corp., *Model 5204 Lock-In Analyzer Operating and Service Manual*.
- [47] R.W.P. Drever, private communication.
- [48] Walt Disney, "Huey, Dewey and Louie; Junior Woodchucks." Whitman Comics.
- [49] Work done by scientists, students and staff of the University of Glasgow Gravitational Physics group; reported to the author verbally by Henry Ward and/or James Hough.
- [50] Henry Ward, private communication (1986).
- [51] Robert Carlson, private communication (1987).
- [52] R.W.P. Drever and J. Hough, private communication.
- [53] Andrej Čadež and Alex Abramovici, "Measuring High Mechanical Quality Factors of Bodies Made of Bare Insulating Materials," Caltech Goldenrod Preprint GRP-154 (accepted for publication), (1987).
- [54] Jeffrey C. Livas, *Upper Limits for Gravitational Radiation From Some Astrophysical Sources*. Ph.D. thesis, Massachusetts Institute of Technology (1987).

- [55] Jeffrey Livas, “Broadband Search for Periodic Sources of Gravitational Radiation,” in *NATO Advanced Research Workshop: Gravitational Wave Data Analysis*, ed. B.F. Schutz. Reidel, Dordrecht (1987).
- [56] Coherent Associates, Inc., *Model 307 Noise-Eater Operation and Maintenance Manual*.
- [57] N.A. Robertson, S. Hoggan, J.B. Mangan, and J. Hough, “Intensity Stabilization of an Argon Laser using an Electro-Optic Modulator—Performance and Limitations,” *Appl. Phys. B* **39**, p. 149 (1986).
- [58] G. Camy, D. Pinaud, N. Courtier and Hu Chi Chuan, “Recent developments in high resolution saturation spectroscopy obtained by use of acousto-optic modulators,” *Revue Phys. Appl.* **17**, p. 357 (1982).
- [59] J.L. Hall and T.W. Hansch, “External dye-laser frequency stabilizer,” *Optics Letters* **9** (11), p. 502 (1984).
- [60] R.W.P. Drever, J.L. Hall, F.V. Kowalski, J. Hough, G.M. Ford, A.J. Munley, and H. Ward, “Laser Phase and Frequency Stabilization Using an Optical Resonator,” *Appl. Phys. B* **31**, p. 97 (1983).
- [61] G.A. Kerr, N.A. Robertson, J. Hough and C.N. Man, “The Fast Frequency Stabilization of an Argon Laser to an Optical Resonator Using an Extra-Cavity Electro-Optic Modulator,” *Appl. Phys. B* **37**, p. 11 (1985).
- [62] C. Man and P. Cérez, “Measurement of frequency response of a piezoelectric transducer using an acoustooptic modulator,” *J. Phys. E: Sci. Instrum.* **13**, p. 157 (1979).
- [63] J. Helmcke, S.A. Lee and J.L. Hall, “Dye laser spectrometer for ultrahigh spectral resolution: design and performance,” *Applied Optics* **21** (9), p.

1686 (1982)

- [64] Roland Schilling, private communication (1987).
- [65] Kip Thorne, private communication (1987).
- [66] S.P. Boughn, W.M. Fairbank, R.P. Giffard, J.N. Hollenhorst, E.R. Mapoles, M.S. McAshan, P.F. Michelson, H.J. Paik, and R.C. Taber, “Observations with a Low-Temperature, Resonant-Mass, Gravitational Radiation Detector,” *Ap. J.* **261**:L19–L22 (1982).
- [67] Bernard F. Schutz, private communication (1987).
- [68] M. Zimmerman and E. Szedenits, Jr., “Gravitational waves from rotating and precessing rigid bodies: I. Simple models and applications to pulsars,” *Phys. Rev. D* **20**, p. 351 (1979).
- [69] Tsvi Piran and Takashi Nakamura, “Can Current Detectors Detect Gravitational Radiation from Newborn Pulsars?,” preprint KUNS 901 (submitted for publication), (1987).
- [70] J. Hough, R.W.P. Drever, H. Ward, A.J. Munley, G.P. Newton, B.J. Meers, S. Hoggan and G.A. Kerr, “Direct observational upper limit to gravitational radiation from millisecond pulsar PSR1937+214,” *Nature* **303**, p. 216 (1983).
- [71] R.W.P. Drever, “Contribution on the Millisecond Pulsar,” *Meeting of the Royal Astronomical Society*, London (1982).
- [72] J. Hough, J.R. Pugh, R. Bland and R.W.P. Drever, “Search for continuous gravitational radiation,” *Nature* **254**, p. 498 (1975).

- [73] Lawrence R. Rabiner and Bernard Gold, *Theory and Application of Digital Signal Processing*. Prentice–Hall, Englewood Cliffs (1975).
- [74] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William Vetterling, *Numerical Recipes*. Cambridge University Press, Cambridge (1986).
- [75] Donald Fraser, “Optimized Mass Storage FFT Program,” in *Programs for Digital Signal Processing*, eds. Digital Signal Processing Committee, IEEE Acoustics, Speech and Signal Processing Society. IEEE Press, New York (1979).
- [76] E. Oran Brigham, *The Fast Fourier Transform*. Prentice–Hall, Englewood Cliffs (1974).
- [77] *The Fundamentals of Signal Analysis*, Hewlett Packard Application Note 243 (1985).
- [78] J.W. Armstrong, F.B. Estabrook, and H.D. Wahlquist, “A Search for Sinusoidal Gravitational Radiation in the Period Range 30–2000 Seconds,” *Ap. J.* **318**:536–541 (1987).
- [79] R. Hanbury Brown, *The Intensity Interferometer*. Taylor & Francis, London (1974).
- [80] R.A. Saenz and S.L. Shapiro, “Gravitational Radiation from Stellar Core Collapse: Ellipsoidal Models,” *Ap. J.* **221**:286 (1978).
- [81] R.F. Stark and T. Piran, “Gravitational–Wave Emission from Rotating Gravitational Collapse,” *Phys. Rev. Lett.* **55**, p. 891 (1985).
- [82] Y. Kojima and T. Nakamura, “Gravitational Radiation from a Particle with Zero Orbital Angular Momentum Plunging into a Kerr Black Hole,” *Physics Letters* **96A**, p. 335 (1983).

- [83] R.W.P. Drever, J. Hough, R. Bland and G.W. Lessnoff, "Experiments and Observations with Wide Band Gravity Wave Detectors," in *Colloques Internationaux du C.N.R.S. No. 220, Ondes et Radiations Gravitationelles*. C.N.R.S., Paris (1973).
- [84] D.G. Maeder, "Proposed Standardization of Gravitational Burst Detector Signals," in *Colloques Internationaux du C.N.R.S. No. 220, Ondes et Radiations Gravitationelles*. C.N.R.S., Paris (1973).
- [85] R.W.P. Drever, J. Hough, R. Bland and G.W. Lessnoff, "Search for Short Bursts of Gravitational Radiation," *Nature* **246** (5432), p. 340 (1973).
- [86] Daniel Dewey, *A Search for Astronomical Gravitational Radiation with an Interferometric Broad Band Antenna*. Ph.D. thesis, Massachusetts Institute of Technology (1986).
- [87] Daniel Dewey, "Comparison of Bars and Interferometers: Detection of Transient Gravitational Radiation," in *NATO Advanced Research Workshop: Gravitational Wave Data Analysis*, ed. B.F. Schutz. Reidel, Dordrecht (1987).
- [88] P.F. Michelson and R.C. Taber, "Sensitivity analysis of a resonant mass gravitational wave antenna with resonant transducer," *J. Appl. Phys* **52**, p. 4313 (1981).
- [89] M.F. Bocko, M.W. Cromar, D.H. Douglass, R.Q. Gram, W.W. Johnson, M. Karim, C.C. Lam, D. Macaluso, J.R. Marsden, B. Muhlfelder, L. Narici, and M. Zucker, "The Rochester Gravitational Wave Detector," *J. Phys. E: Sci. Instrum.* **17**, p. 604 (1984).

- [90] Andrew D. Jeffries, Peter R. Saulson, Robert E. Spero, and Michael E. Zucker, "Gravitational Wave Observatories," *Scientific American* **256** (6), p. 50 (1987).
- [91] P.C. Peters and J. Matthews, "Gravitational Radiation from Point Masses in Keplerian Orbit," *Phys. Rev.* **131** (1), p. 435 (1963).
- [92] Bernard F. Schutz, "Sources of Gravitational Radiation," in *NATO Advanced Research Workshop: Gravitational Wave Data Analysis*, ed. B.F. Schutz. Reidel, Dordrecht (1987).
- [93] Sheryl L. Smith, "An Algorithm to Search for Gravitational Radiation from Coalescing Binaries," *Phys. Rev. D* **36**, p. 2901 (1987).
- [94] Sheryl L. Smith, private communication (1988).
- [95] Douglas M. Eardley, "Theoretical Models for Sources of Gravitational Waves," in *Gravitational Radiation*, eds. N. Deruelle and T. Piran. North Holland, Amsterdam (1983).
- [96] John Paul Adrian Clark and Douglas M. Eardley, "Evolution of Close Neutron Star Binaries," *Ap. J.* **215**:311-322 (1977).
- [97] Robert V. Wagoner, "Gravitational Radiation from Accreting Neutron Stars," *Ap. J.* **278**:345-348 (1984).
- [98] Martin J. Rees, "Extragalactic Sources of Gravitational Waves," in *Gravitational Radiation*, eds. N. Dereulle and T. Piran. North Holland, Amsterdam (1983).
- [99] W.M. Smart, *Textbook on Spherical Astronomy*. Cambridge University Press, Cambridge (1977).