1

LIGO: The Laser Interferometer Gravitational-Wave Observatory

P. Fritschel on behalf of the LIGO Scientific Collaboration

Abstract

This chapter features the USA based LIGO, the Laser Interferometer Gravitational-Wave Observatory - the first of three case studies covering different worldwide interferometric gravitational wave detectors. In addition to describing the basic interferometer operation and its various components, we discuss the technological challenges that have been overcome for its successful operation.

1.1 Introduction

The prediction of gravitational waves (GWs), oscillations in the space-time metric that propagate at the speed of light, is one of the most profound differences between Einstein's general theory of relativity and the Newtonian theory of gravity that it replaced. As discussed in the first chapter, GWs remained a theoretical prediction for more than 50 years until the first observational evidence for their existence came with the discovery and subsequent observations of the binary pulsar PSR 1913+16, by Russell Hulse and Joseph Taylor (Weisberg and Taylor, 2005). In about 300 million years, the PSR 1913+16 orbit will decrease to the point where the pair coalesces into a single compact object, a process that will produce directly detectable gravitational waves. In the meantime, the direct detection of GWs will require similarly strong sources – extremely large masses moving with large accelerations in strong gravitational fields. The goal of LIGO, the Laser Interferometer Gravitational-Wave Observatory (Abramovici *et al.*, 1992) is just that: to detect and study GWs of astrophysical origin. Achieving this goal will mark the opening of a new window on the universe, with the promise of new physics and astrophysics. In physics, GW detection could provide information about strong-field gravitation, the untested domain of strongly

curved space where Newtonian gravitation is no longer even a poor approximation. In astrophysics, the sources of GWs that LIGO might detect include binary neutron stars (like PSR 1913+16 but much later in their evolution); binary systems where a black hole replaces one or both of the neutron stars; a stellar core collapse which triggers a Type II supernova; rapidly rotating, non-axisymmetric neutron stars; and possibly processes in the early universe that produce a stochastic background of GWs (Cutler and Thorne, 2002).

In the past few years the field has reached a milestone, with decades-old plans to build and operate large interferometric GW detectors now realized in several locations worldwide. This chapter, based on a previously published article, Abbott *et al.* (2009), focuses on LIGO, which operates the most sensitive detectors yet built. We aim to describe the LIGO detectors and how they operate, explain how they have achieved their remarkable sensitivity, and review how their data can be used to learn about a variety of astrophysical phenomena.

1.2 The LIGO detectors

The oscillating quadrupolar strain pattern of a GW is well matched by a Michelson interferometer, which makes a very sensitive comparison of the lengths of its two orthogonal arms. LIGO utilizes three specialized Michelson interferometers, located at two sites (see Fig. 1.1): an observatory on the Hanford site in Washington houses two interferometers, the 4 km-long H1 and 2 km-long H2 detectors; and an observatory in Livingston, Louisiana, houses the 4 km-long L1 detector. Other than the shorter length of H2, the three interferometers are essentially identical. Multiple detectors at separated sites are crucial for rejecting instrumental and environmental artifacts in the data, by requiring coincident detections in the analysis. Also, because the antenna pattern of an interferometer is quite wide, source localization requires triangulation using three separated detectors.

The initial LIGO detectors were designed to be sensitive to GWs in the frequency band $40-7000\,\mathrm{Hz}$, and capable of detecting a GW strain amplitude as small as 10^{-21} over a 1ś integration time (Abramovici *et al.*, 1992). With funding from the National Science Foundation, the LIGO sites and detectors were designed by scientists and engineers from the California Institute of Technology, the Massachusetts Institute of Technology and the University of Florida, constructed in the late 1990s, and

Figure 1.1 Aerial photograph of the LIGO observatories at Hanford, Washington (top) and Livingston, Louisiana (bottom). The lasers and optics are contained in the white and blue buildings. From the large corner building, evacuated beam tubes extend at right angles for 4 km in each direction (the full length of only one of the arms is seen in each photo); the tubes are covered by the arched, concrete enclosures seen here.

commissioned over the first 5 years of this decade. From November 2005 through September 2007, they operated at their design sensitivity in a continuous data-taking mode. The data from this science run, known as S5, were analyzed for a variety of GW signals by a group of researchers known as the LIGO Scientific Collaboration[1] (LSC). At the most sensitive frequencies, the instrument root-mean-square (rms) strain noise reached an unprecedented level of 3×10^{-22} in a 100 Hz band. Table 1.1 lists the main design parameters of the LIGO interferometers.

[1] Homepage of the LIGO Scientific Collaboration, http://www.ligo.org

1.3 Detector description

As discussed in the earlier chapters, to measure a GW strain using a Michelson interferometer, the challenge is to make the instrument sufficiently sensitive: at the targeted strain sensitivity of 10^{-21}, the resulting arm length change is only $\sim 10^{-18}$ m, a thousand times smaller than the diameter of a proton. Meeting this challenge involves the use of special interferometry techniques, state-of-the-art optics, highly stable lasers, and multiple layers of vibration isolation, all of which are described in the sections that follow. And of course a key feature of the detectors is simply their scale: the arms are made as long as practically possible to increase the signal due to a GW strain. See Table 1.1 for a list of the main design parameters of the LIGO interferometers.

1.3.1 Interferometer Configuration

The LIGO detectors are Michelson interferometers whose mirrors also serve as gravitational test masses. A passing gravitational wave will impress a phase modulation on the light in each arm of the Michelson, with a relative phase shift of 180 degrees between the arms. When the Michelson arm lengths are set such that the un-modulated light interferes destructively at the antisymmetric (AS) port – the dark fringe condition – the phase modulated sideband light will interfere constructively, with an amplitude proportional to GW strain and the input power. With dark fringe operation, the full power incident on the beamsplitter is returned to the laser at the symmetric port. Only differential motion of the arms appears at the AS port; common mode signals are returned to the laser with the carrier light.

Two modifications to a basic Michelson, shown in Fig. 1.2, increase the carrier power in the arms and hence the GW sensitivity. First, each arm contains a resonant Fabry-Perot optical cavity made up of a partially transmitting input mirror and a high reflecting end mirror. The cavities cause the light to effectively bounce back and forth multiple times in the arms, increasing the carrier power and phase shift for a given strain amplitude. In the LIGO detectors the Fabry-Perot cavities multiply the signal by a factor of 100 for a 100 Hz GW. Second, a partially-reflecting mirror is placed between the laser and beamsplitter to implement power recycling (Meers, 1988). In this technique, an optical cavity is formed between the power recycling mirror and the Michelson symmetric port. By matching the transmission of the recycling mirror to the optical losses

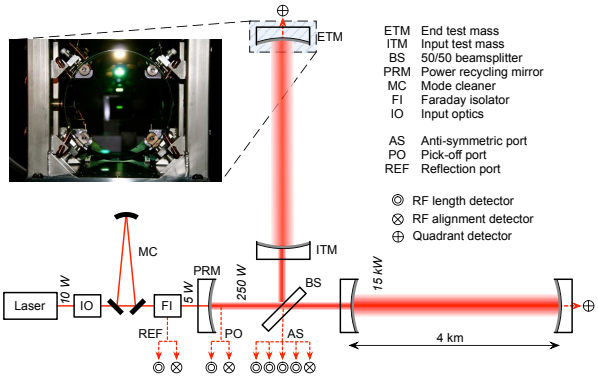

Figure 1.2 Optical and sensing configuration of the LIGO 4 km interferometers (the laser power numbers here are generic; specific power levels are given in Table 1). The IO block includes laser frequency and amplitude stabilization, and electro-optic phase modulators. The power recycling cavity is formed between the PRM and the two ITMs, and contains the BS. The inset photo shows an input test mass mirror in its pendulum suspension. The near face has a highly reflective coating for the infrared laser light, but transmits visible light. Through it one can see mirror actuators arranged in a square pattern near the mirror perimeter.

in the Michelson, and resonating this recycling cavity, the laser power stored in the interferometer can be significantly increased. In this configuration, known as a power recycled Fabry-Perot Michelson, the LIGO interferometers increase the power in the arms by a factor of $\approx 8,000$ with respect to a simple Michelson.

1.3.2 Laser and Optics

The laser source is a diode-pumped, Nd:YAG master oscillator and power amplifier system, and emits 10 W in a single frequency at 1064 nm (Savage et al., 1998). The laser power and frequency are actively stabilized, and passively filtered with a transmissive ring cavity (pre-mode cleaner, PMC). The laser power stabilization is implemented by directing a sample of the beam to a photodetector, filtering its signal and feeding it back to the power amplifier; this servo stabilizes the relative power fluctuations of the beam to $\sim 10^{-7}/\sqrt{\text{Hz}}$ at 100 Hz (Abbott and King, 2001). The laser frequency stabilization is done in multiple stages that are more fully described in later sections. The first, or pre-stabilization stage uses the traditional technique of servo locking the laser frequency

to an isolated reference cavity using the Pound-Drever-Hall (PDH) technique (Drever et al., 1983), in this case via feedback to frequency actuators on the master oscillator and to an electro-optic phase modulator. The servo bandwith is 500 kHz, and the pre-stabilization achieves a stability level of $\sim 10^{-2}\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ at 100 Hz. The PMC transmits the pre-stabilized beam, filtering out both any light not in the fundamental Gaussian spatial mode and laser noise at frequencies above a few MHz (Willke et al., 1998). The PMC output beam is weakly phase-modulated with two radio-frequency (RF) sine waves, producing, to first-order, two pairs of sideband fields around the carrier field; these RF sideband fields are used in a heterodyne detection system described below.

After phase modulation, the beam passes into the LIGO vacuum system. All the main interferometer optical components and beam paths are enclosed in the ultra-high vacuum system ($10^{-8} - 10^{-9}$ torr) for acoustical isolation and to reduce phase fluctuations from light scattering off residual gas (Zucker and Whitcomb, 1996). The long beam tubes are particularly noteworthy components of the LIGO vacuum system. These 1.2 m diameter, 4 km long stainless steel tubes were designed to have low-outgassing so that the required vacuum could be attained by pumping only from the ends of the tubes. This was achieved by special processing of the steel to remove hydrogen, followed by an *in-situ* bakeout of the spiral-welded tubes, for approximately 20 days at 160 C.

The in-vacuum beam first passes through the mode cleaner (MC), a 12 m long, vibrationally isolated transmissive ring cavity. The MC provides a stable, diffraction-limited beam with additional filtering of laser noise above several kilohertz (Skeldon et al., 1996; Yoshida *et al.*, 2000), and it serves as an intermediate reference for frequency stabilization. The MC length and modulation frequencies are matched so that the main carrier field and the modulation sideband fields all pass through the MC. After the MC is a Faraday isolator and a reflective 3-mirror telescope that expands the beam and matches it to the arm cavity mode.

The interferometer optics, including the test masses, are fused-silica substrates with multilayer dielectric coatings, manufactured to have extremely low scatter and low absorption. The test mass substrates are polished so that the surface deviation from a spherical figure, over the central 80 mm diameter, is typically 5 angstroms or smaller, and the surface microroughness is typically less than 2 angstroms (Walsh et al., 1999). The mirror coatings are made using ion-beam sputtering, a technique known for producing ultralow-loss mirrors (Wei, 1989; Rempe et al., 1992). The absorption level in the coatings is generally a few parts-per-

	H1	L1	H2
Laser type and wavelength	Nd:YAG, $\lambda = 1064$ nm		
Arm cavity finesse	220		
Arm length	3995 m	3995 m	2009 m
Arm cavity storage time, τ_s	0.95 ms	0.95 ms	0.475 ms
Input power at recycling mirror	4.5 W	4.5 W	2.0 W
Power Recycling gain	60	45	70
Arm cavity stored power	20 kW	15 kW	10 kW
Test mass size & mass	$\phi\,25\,\text{cm} \times 10\,\text{cm}$, 10.7 kg		
Beam radius ITM/ETM ($1/e^2$ power)	3.6 / 4.5 cm	3.9 / 4.5 cm	3.3 / 3.5 cm
Test mass pendulum frequency	0.76 Hz		

Table 1.1 *Parameters of the LIGO interferometers. H1 and H2 refer to the interferometers at Hanford, Washington, and L1 is the interferometer at Livingston, Louisiana.*

million (ppm) or less (Ottaway et al., 2006), and the total scattering loss from a mirror surface is estimated to be 60 – 70 ppm.

In addition to being a source of optical loss, scattered light can be a problematic noise source, if it is allowed to reflect or scatter from a vibrating surface (such as a vacuum system wall) and recombine with the main beam (Vinet et al., 1996). Since the vibrating, re-scattering surface may be moving by $\sim$10 orders of magnitude more than the test masses, very small levels of scattered light can contaminate the output. To control this, various baffles are employed within the vacuum system to trap scattered light (Vinet et al., 1996; Vinet *et al.*, 1997). Each 4 km long beam tube contains approximately two hundred baffles to trap light scattered at small angles from the test masses. These baffles are stainless steel truncated cones, with serrated inner edges, distributed so as to completely hide the beam tube from the line of sight of any arm cavity mirror. Additional baffles within the vacuum chambers prevent light outside the mirror apertures from hitting the vacuum chamber walls.

1.3.3 Suspensions and Vibration Isolation

Starting with the MC, each mirror in the beam line is suspended as a pendulum by a loop of steel wire. The pendulum provides f^{-2} vibration isolation above its eigenfrequencies, allowing free movement of a test

mass in the GW frequency band. Along the beam direction, a test mass pendulum isolates by a factor of nearly 2×10^4 at 100 Hz. The position and orientation of a suspended optic is controlled by electromagnetic actuators: small magnets are bonded to the optic and coils are mounted to the suspension support structure, positioned to maximize the magnetic force and minimize ground noise coupling. The actuator assemblies also contain optical sensors that measure the position of the suspended optic with respect to its support structure. These signals are used to actively damp eigenmodes of the suspension.

The bulk of the vibration isolation in the GW band is provided by four-layer mass-spring isolation stacks, to which the pendulums are mounted. These stacks provide approximately f^{-8} isolation above ~ 10 Hz (Giaime et al., 1996), giving an isolation factor of about 10^8 at 100 Hz. In addition, the L1 detector, subject to higher environmental ground motion than the Hanford detectors, employs seismic pre-isolators between the ground and the isolation stacks. These active isolators employ a collection of motion sensors, hydraulic actuators, and servo controls; the pre-isolators actively suppress vibrations in the band $0.1 - 10$ Hz, by as much as a factor of 10 in the middle of the band (Abbott *et al.*, 2004).

1.3.4 Sensing and Controls

The two Fabry-Perot arms and power recycling cavities are essential to achieving the LIGO sensitivity goal, but they require an active feedback system to maintain the interferometer at the proper operating point (Fritschel et al., 2001). The round trip length of each cavity must be held to an integer multiple of the laser wavelength so that newly introduced carrier light interferes constructively with light from previous round trips. Under these conditions the light inside the cavities builds up and they are said to be on resonance. In addition to the three cavity lengths, the Michelson phase must be controlled to ensure that the AS port remains on the dark fringe.

The four lengths are sensed with a variation of the PDH reflection scheme (Regehr et al., 1995). In standard PDH, an error signal is generated through heterodyne detection of the light reflected from a cavity. The RF phase modulation sidebands are directly reflected from the cavity input mirror and serve as a local oscillator to mix with the carrier field. The carrier experiences a phase-shift in reflection, turning the RF phase modulation into RF amplitude modulation, linear in amplitude for small deviations from resonance. This concept is extended to the

full interferometer as follows. At the operating point, the carrier light is resonant in the arm and recycling cavities and on a Michelson dark fringe. The RF sideband fields resonate differently. One pair of RF sidebands (from phase modulation at 62.5 MHz) is not resonant and simply reflects from the recycling mirror. The other pair (25 MHz phase modulation) is resonant in the recycling cavity but not in the arm cavities.[2] The Michelson mirrors are positioned to make one arm 30 cm longer than the other so that these RF sidebands are not on a Michelson dark fringe. By design this Michelson asymmetry is chosen so that most of the resonating RF sideband power is coupled to the AS port.

In this configuration, heterodyne error signals for the four length degrees-of-freedom are extracted from the three output ports shown in Fig. 1.2. The AS port is heterodyned at the resonating RF frequency and gives an error signal proportional to differential arm length changes, including those due to a GW. The PO port is a sample of the recycling cavity beam, and is detected at the resonating RF frequency to give error signals for the recycling cavity length and the Michelson phase (using both RF quadratures). The REF port is detected at the non-resonating RF frequency and gives a standard PDH signal proportional to deviations in the laser frequency relative to the average length of the two arms.

Feedback controls derived from these errors signals are applied to the two end mirrors to stabilize the differential arm length, to the beamsplitter to control the Michelson phase, and to the recycling mirror to control the recycling cavity length. The feedback signals are applied directly to the mirrors through their coil-magnet actuators, with slow corrections for the differential arm length applied with longer-range actuators that move the whole isolation stack.

The common arm length signal from the REF port detection is used in the final level of laser frequency stabilization (Adhikari, 2004) pictured schematically in Fig. 1.3. The hierarchical frequency control starts with the reference cavity pre-stabilization mentioned in Sec. 1.3.2. The pre-stabilization path includes an Acousto-Optic Modulator (AOM)driven by a voltage-controlled oscillator, through which fast corrections to the pre-stabilized frequency can be made. The MC servo uses this correction path to stabilize the laser frequency to the MC length, with a servo bandwidth close to 100 kHz. The most stable frequency reference in the GW band is naturally the average length of the two arm cavities,

[2] These are approximate modulation frequencies for H1 and L1; those for H2 are about 10% higher.

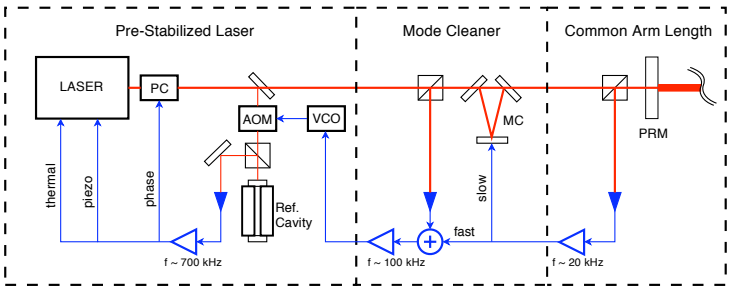

Figure 1.3 Schematic layout of the frequency stabilization servo. The laser is locked to a fixed-length reference cavity through an AOM. The AOM frequency is generated by a Voltage Controlled Oscillator (VCO) driven by the MC, which is in turn driven by the common mode arm length signal from the REF port. The laser frequency is actuated by a combination of a Pockels Cell (PC), piezo actuator, and thermal control.

therefore the common arm length error signal provides the final level of frequency correction. This is accomplished with feedback to the MC, directly to the MC length at low frequencies and to the error point of the MC servo at high frequencies, with an overall bandwidth of 20 kHz. The MC servo then impresses the corrections onto the laser frequency. The three cascaded frequency loops – the reference cavity pre-stabilization; the MC loop; and the common arm length loop – together provide 160 dB of frequency noise reduction at 100 Hz, and achieve a frequency stability of 5 μHz rms in a 100 Hz bandwidth with respect to the common mode of the arm cavities.

The photodetectors are all located outside the vacuum system, mounted on optical tables. Telescopes inside the vacuum reduce the beam size by a factor of ~ 10, and the small beams exit the vacuum through high-quality windows. To reduce noise from scattered light and beam clipping, the optical tables are housed in acoustical enclosures, and the more critical tables are mounted on passive vibration isolators. Any back-scattered light along the AS port path is further mitigated with a Faraday isolator mounted in the vacuum system.

The total AS port power is typically 200 – 250 mW, and is a mixture of RF sideband local oscillator power and carrier light resulting from spatially imperfect interference at the beamsplitter. The light is divided equally between four length photodetectors, keeping the power on each at a detectable level of 50 – 60 mW. The four length detector signals are summed and filtered, and the feedback control signal is applied differ-

entially to the end test masses. This differential-arm servo loop has a unity-gain bandwidth of approximately 200 Hz, suppressing fluctuations in the arm lengths to a residual level of $\sim 10^{-14}$ m rms. An additional servo is implemented on these AS port detectors to cancel signals in the RF-phase orthogonal to the differential-arm channel; this servo injects RF current at each photodetector to suppress signals that would otherwise saturate the detectors. About 1% of the beam is directed to an alignment detector that controls the differential alignment of the ETMs.

Maximal power buildup in the interferometer also depends on maintaining stringent alignment levels. Sixteen alignment degrees-of-freedom – pitch and yaw for each of the 6 interferometer mirrors and the input beam direction – are controlled by a hierarchy of feedback loops. First, alignment motion at the pendulum and isolation stack eigenfrequencies are suppressed locally at each optic using optical lever angle sensors. Second, global alignment is established with four RF quadrant photodetectors at the three output ports as shown in Fig. 1.2. These RF alignment detectors measure wavefront misalignments between the carrier and sideband fields in a spatial version of PDH detection (Morrison et al., 1994; Hefetz et al., 1997). Together the four detectors provide 5 linearly independent combinations of the angular deviations from optimal global alignment (Fritschel et al., 1998). These error signals feed a multiple-input multiple-output control scheme to maintain the alignment within $\sim 10^{-8}$ radians rms of the optimal point, using bandwidths between ~ 0.5 Hz and ~ 5 Hz depending on the channel. Finally, slower servos hold the beam centered on the optics. The beam positions are sensed at the arm ends using DC quadrant detectors that receive the weak beam transmitted through the ETMs, and at the corner by imaging the beam spot scattered from the beamsplitter face with a CCD camera.

The length and alignment feedback controls are all implemented digitally, with a real-time sampling rate of 16384 samples/sec for the length controls and 2048 samples/sec for the alignment controls. The digital control system provides the flexibility required to implement the multiple input, multiple output feedback controls described above. The digital controls also allow complex filter shapes to be easily realized, lend the ability to make dynamic changes in filtering, and make it simple to blend sensor and control signals. As an example, optical gain changes are compensated to first order to keep the loop gains constant in time by making real-time feed-forward corrections to the digital gain based on cavity power levels.

The digital controls are also essential to implementing the interferom-

eter *lock acquisition* algorithm. So far this section has described how the interferometer is maintained at the operating point. The other function of the control system is to acquire lock: to initially stabilize the relative optical positions to establish the resonance conditions and bring them within the linear regions of the error signals. Before lock the suspended optics are only damped within their suspension structures; ground motion and the equivalent effect of input-light frequency fluctuations cause the four (real or apparent) lengths to fluctuate by $0.1-1$ μm rms over time scales of $0.5-10$ sec. The probability of all four degrees-of-freedom simultaneously falling within the ~ 1 nm linear region of the resonance points is thus extremely small and a controlled approach is required. The basic approach of the lock acquisition scheme, described in detail in reference (Evans et al., 2002), is to control the degrees-of-freedom in sequence: first the power-recycled Michelson is controlled, then a resonance of one arm cavity is captured, and finally a resonance of the other arm cavity is captured to achieve full power buildup. A key element of this scheme is the real-time, dynamic calculation of a sensor transformation matrix to form appropriate length error signals throughout the sequence. The interferometers are kept in lock typically for many hours at a time, until lock is lost due to environmental disturbances, instrument malfunction or operator command.

1.3.5 Thermal Effects

At full power operation, a total of $20-60$ mW of light is absorbed in the substrate and in the mirror surface of each ITM, depending on their specific absorption levels. Through the thermo-optic coefficient of fused silica, this creates a weak, though not insignificant thermal lens in the ITM substrates (Winkler et al., 1991). Thermo-elastic distortion of the test mass reflecting surface is not significant at these absorption levels. While the ITM thermal lens has little effect on the carrier mode, which is determined by the arm cavity radii of curvature, it does affect the RF sideband mode supported by the recycling cavity. This in turn affects the power buildup and mode shape of the RF sidebands in the recycling cavity, and consequently the sensitivity of the heterodyne detection signals (D'Ambrosio and Kells, 2006; Gretarsson et al., 2007). Achieving maximum interferometer sensitivity thus depends critically on optimizing the thermal lens and thereby the mode shape, a condition which occurs at a specific level of absorption in each ITM (approximately 50 mW). To achieve this optimum mode over the range of ITM absorption and stored

power levels, each ITM thermal lens is actively controlled by directing additional heating beams, generated from CO_2 lasers, onto each ITM (Ballmer *et al.*, 2005). The power and shape of the heating beams are controlled to maximize the interferometer optical gain and sensitivity. The shape can be selected to have either a Gaussian radial profile to provide more lensing, or an annular radial profile to compensate for excess lensing.

1.3.6 Interferometer Response and Calibration

The GW channel is the digital error point of the differential-arm servo loop. In principle the GW channel could be derived from any point within this loop. The error point is chosen because the dynamic range of this signal is relatively small, since the large low-frequency fluctuations are suppressed by the feedback loop. To calibrate the error point in strain, the effect of the feedback loop is divided out, and the interferometer response to a differential arm strain is factored in (Landry and the LIGO Scientific Collaboration, 2005); this process can be done either in the frequency domain or directly in the time domain. The absolute length scale is established using the laser wavelength, by measuring the mirror drive signal required to move through an interference fringe. The calibration is tracked during operation with sine waves injected into the differential-arm loop. The uncertainty in the amplitude calibration is approximately ±5%. Timing of the GW channel is derived from the Global Positioning System; the absolute timing accuracy of each interferometer is better than $\pm 10\,\mu$sec.

The response of the interferometer output as a function of GW frequency is calculated in detail in references (Meers, 1989; Fabbro and Montelatici, 1995; Rakhmanov et al., 2008). In the long-wavelength approximation, where the wavelength of the GW is much longer than the size of the detector, the response R of a Michelson-Fabry-Perot interferometer is approximated by a single-pole transfer function:

$$R(f) \propto \frac{1}{1 + if/f_p}\,, \tag{1.1}$$

where the pole frequency is related to the storage time by $f_p = 1/4\pi\tau_s$. Above the pole frequency ($f_p = 85$ Hz for the LIGO 4 km interferometers), the amplitude response drops off as $1/f$. As discussed below, the measurement noise above the pole frequency has a white (flat) spectrum, and so the strain sensitivity decreases proportionally to frequency in this

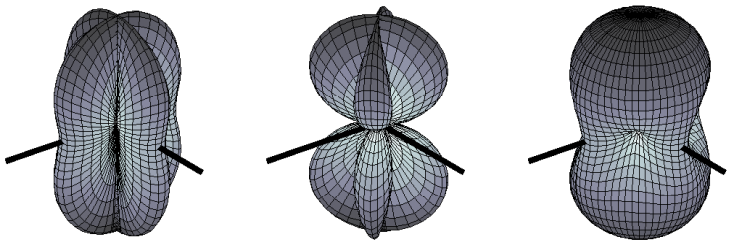

Figure 1.4 Antenna response pattern for a LIGO gravitational wave detector, in the long-wavelength approximation. The interferometer beamsplitter is located at the center of each pattern, and the thick black lines indicate the orientation of the interferometer arms. The distance from a point of the plot surface to the center of the pattern is a measure of the gravitational wave sensitivity in this direction. The pattern on the left is for + polarization, the middle pattern is for × polarization, and the right-most one is for unpolarized waves.

region. The single-pole approximation is quite accurate, differing from the exact response by less than a percent up to ~ 1 kHz (Rakhmanov et al., 2008).

In the long-wavelength approximation, the interferometer directional response is maximal for GWs propagating orthogonally to the plane of the interferometer arms, and linearly polarized along the arms. Other angles of incidence or polarizations give a reduced response, as depicted by the antenna patterns shown in Fig. 1.4. A single detector has blind spots on the sky for linearly polarized gravitational waves.

1.3.7 Environmental Monitors

To complete a LIGO detector, the interferometers described above are supplemented with a set of sensors to monitor the local environment. Seismometers and accelerometers measure vibrations of the ground and various interferometer components; microphones monitor acoustic noise at critical locations; magnetometers monitor fields that could couple to the test masses or electronics; radio receivers monitor RF power around the modulation frequencies. These sensors are used to detect environmental disturbances that can couple to the GW channel.

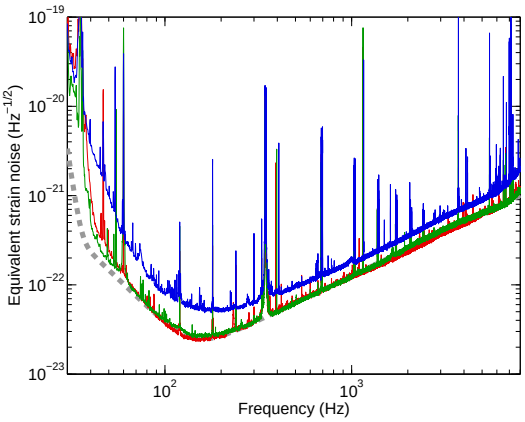

Figure 1.5 Strain sensitivities, expressed as amplitude spectral densities of detector noise converted to equivalent GW strain. The vertical axis denotes the rms strain noise in 1 Hz of bandwidth. Shown are typical high sensitivity spectra for each of the three interferometers (darkest solid curve H2; lighter solid curves H1 and L1), along with the design goal for the 4-km detectors (dashed curve).

1.4 Instrument performance

1.4.1 Strain Noise Spectra

During the commissioning period, as the interferometer sensitivity was improved, several short science runs were carried out, culminating with the fifth science run (S5) at design sensitivity. The S5 run collected a full year of triple-detector coincident interferometer data during the period from November 2005 through September 2007. Since the interferometers detect GW strain amplitude, their performance is typically characterized by an amplitude spectral density of detector noise (the square root of the power spectrum), expressed in equivalent GW strain. Typical high-sensitivity strain noise spectra are shown in Fig. 1.5. Over the course of S5 the strain sensitivity of each interferometer was improved, by up to 40% compared to the beginning of the run through a series of incremental improvements to the instruments.

The primary noise sources contributing to the H1 strain noise spectrum are shown in Fig. 1.6. Understanding and controlling these instrumental noise components has been the major technical challenge in the development of the detectors. The noise terms can be broadly divided into two classes: displacement noise and sensing noise. Displacement

noises cause motions of the test masses or their mirrored surfaces. Sensing noises, on the other hand, are phenomena that limit the ability to measure those motions; they are present even in the absence of test mass motion. The strain noises shown in Fig. 1.5 consists of spectral lines superimposed on a continuous broadband noise spectrum. The majority of the lines are due to power lines ($60, 120, 180, ...$Hz), "violin mode" mechanical resonances ($340, 680, ...$Hz) and calibration lines ($55, 400$, and 1100 Hz). These high Q lines are easily excluded from analysis; the broadband noise dominates the instrument sensitivity.

1.4.2 Sensing Noise Sources

Sensing noises are shown in the lower panel of Fig. 1.6. By design, the dominant noise source above 100 Hz is shot noise, as determined by the Poisson statistics of photon detection. The ideal shot-noise limited strain noise density, $\tilde{h}(f)$, for this type of interferometer is (Meers, 1988):

$$\widetilde{h}(f) = \sqrt{\frac{\pi\hbar\lambda}{\eta P_{\rm BS} c}} \frac{\sqrt{1+(4\pi f\tau_{\rm s})^2}}{4\pi\tau_{\rm s}}, \tag{1.2}$$

where λ is the laser wavelength, $\hbar$ is the reduced Planck constant, c is the speed of light, $\tau_{\rm s}$ is the arm cavity storage time, f is the GW frequency, $P_{\rm BS}$ is the power incident on the beamsplitter, and η is the photodetector quantum efficiency. For the estimated effective power of $\eta P_{\rm BS} = 0.9 \cdot 250$ W, the ideal shot-noise limit is $\widetilde{h} = 1.0 \times 10^{-23}/\sqrt{\rm Hz}$ at 100 Hz. The shot-noise estimate in Fig. 1.6 is based on measured photocurrents in the AS port detectors and the measured interferometer response. The resulting estimate, $\widetilde{h}(100{\rm Hz}) = 1.3 \times 10^{-23}/\sqrt{\rm Hz}$, is higher than the ideal limit due to several inefficiencies in the heterodyne detection process: imperfect interference at the beamsplitter increases the shot noise; imperfect modal overlap between the carrier and RF sideband fields decreases the signal; and the fact that the AS port power is modulated at twice the RF phase modulation frequency leads to an increase in the time-averaged shot noise (Niebauer et al., 1991).

Many noise contributions are estimated using stimulus-response tests, where a sine-wave or broadband noise is injected into an auxiliary channel to measure its coupling to the GW channel. This method is used for the laser frequency and amplitude noise estimates, the RF oscillator phase noise contribution, and also for the angular control and auxiliary length noise terms described below. Although laser noise is nominally common-mode, it couples to the GW channel through small, unavoidable

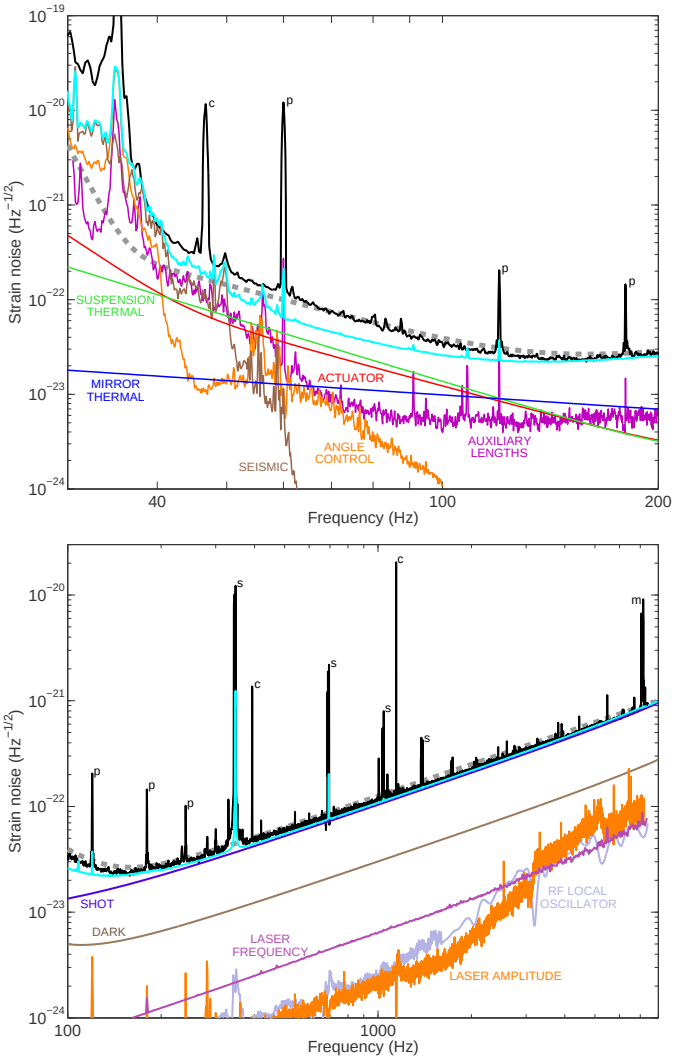

Figure 1.6 Primary known contributors to the H1 detector noise spectrum. The upper panel shows the displacement noise components, while the lower panel shows sensing noises (note the different frequency scales). In both panels, the darkest curve is the measured strain noise (same spectrum as in Fig. 1.5), the light solid curve is the root-square-sum of all known contributors (both sensing and displacement noises) and the light dashed curve is the design goal. The labelled component curves are described in the text. The known noise sources explain the observed noise very well at frequencies above 150 Hz, and to within a factor of 2 in the 40 – 100 Hz band. Spectral peaks are identified as follows: c, calibration line; p, power line harmonic; s, suspension wire vibrational mode; m, mirror (test mass) vibrational mode.

differences in the arm cavity mirrors (Sigg, 1997; Camp et al., 2000). Frequency noise is expected to couple most strongly through a difference in the resonant reflectivity of the two arms. This causes carrier light to leak out the AS port, which interferes with frequency noise on the RF sidebands to create a noise signal. The stimulus-response measurements indicate the coupling is due to a resonant reflectivity difference of about 0.5%, arising from a loss difference of tens of ppm between the arms. Laser amplitude noise can couple through an offset from the carrier dark fringe. The measured coupling is linear, indicating an effective static offset of ~ 1 picometer, believed to be due to mode shape differences between the arms.

1.4.3 Seismic and Thermal Noise

Displacement noises are shown in the upper panel of Fig. 1.6. At the lowest frequencies the largest such noise is seismic noise – motions of the earth's surface driven by wind, ocean waves, human activity, and low-level earthquakes – filtered by the isolation stacks and pendulums. The seismic contribution is estimated using accelerometers to measure the vibration at the isolation stack support points, and propagating this motion to the test masses using modeled transfer functions of the stack and pendulum. The seismic wall frequency, below which seismic noise dominates, is approximately 45 Hz, a bit higher than the goal of 40 Hz, as the actual environmental vibrations around these frequencies are ~ 10 times higher than was estimated in the design.

Mechanical thermal noise is a more fundamental effect, arising from finite losses present in all mechanical systems, and is governed by the fluctuation-dissipation theorem (Saulson, 1990; Levin, 1998). It causes arm length noise through thermal excitation of the test mass pendulums (*suspension thermal noise*) (Gonzalez, 2000), and thermal acoustic waves that perturb the test mass mirror surface (*test mass thermal noise*) (Harry et al., 2002). Most of the thermal energy is concentrated at the resonant frequencies, which are designed (as much as possible) to be outside the detection band. Away from the resonances, the level of thermal motion is proportional to the mechanical dissipation associated with the motion. Designing the mirror and its pendulum to have very low mechanical dissipation reduces the detection-band thermal noise. It is difficult, however, to accurately and unambiguously establish the level of broadband thermal noise *in-situ*; instead, the thermal noise curves in Fig. 1.6 are calculated from models of the suspension and test masses,

with mechanical loss parameters taken from independent characterizations of the materials.

For the pendulum mode, the mechanical dissipation occurs near the ends of the suspension wire, where the wire flexes. Since the elastic energy in the flexing regions depends on the wire radius to the fourth power, it helps to make the wire as thin as possible to limit thermal noise. The pendulums are thus made with steel wire for its strength; with a diameter of 300 μm the wires are loaded to 30% of their breaking stress. The thermal noise in the pendulum mode of the test masses is estimated assuming a frequency-independent mechanical loss angle in the suspension wire of 3×10^{-4} (Gillespie and Raab, 1994). This is a relatively small loss for a metal wire (Cagnoli et al., 1999).

Thermal noise of the test mass surface is associated with mechanical damping within the test mass. The fused-silica test mass substrate material has very low mechanical loss, of order 10^{-7} or smaller (Penn et al., 2006). On the other hand, the thin-film, dielectric coatings that provide the required optical reflectivity – alternating layers of silicon dioxide and tantalum pentoxide – have relatively high mechanical loss. Even though the coatings are only a few microns thick, they are the dominant source of the relevant mechanical loss, due to their level of dissipation and the fact that it is concentrated on the test mass face probed by the laser beam (Levin, 1998). The test mass thermal noise estimate is calculated by modeling the coatings as having a frequency-independent mechanical dissipation of 4×10^{-4} (Harry et al., 2002).

1.4.4 Auxiliary Degree-of-freedom Noise

The auxiliary length noise term refers to noise in the Michelson and power recycling cavity servo loops which couple to the GW channel. The former couples directly to the GW channel while the latter couples in a manner similar to frequency noise. Above $\sim$50 Hz the sensing noise in these loops is dominated by shot noise; since loop bandwidths of $\sim$ 100 Hz are needed to adequately stabilize these degrees of freedom, shot noise is effectively added onto their motion. Their noise infiltration to the GW channel, however, is mitigated by appropriately filtering and scaling their digital control signals and adding them to the differential-arm control signal as a type of feed-forward noise suppression (Fritschel et al., 2001). These correction paths reduce the coupling to the GW channel by 10 – 40 dB.

We illustrate this more concretely with the Michelson loop. The shot-

noise-limited sensitivity for the Michelson is $\sim 10^{-16}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$. Around 100 Hz, the Michelson servo impresses this sensing noise onto the Michelson degree-of-freedom (specifically, onto the beamsplitter). Displacement noise in the Michelson couples to displacement noise in the GW channel by a factor of $\pi/(\sqrt{2}F) = 1/100$, where F is the arm cavity finesse. The Michelson sensing noise would thus produce $\sim 10^{-18}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ of GW channel noise around 100 Hz, if uncorrected. The digital correction path subtracts the Michelson noise from the GW channel with an efficiency of 95% or more. This reduces the Michelson noise component down to $\sim 10^{-20}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ in the GW channel, 5 – 10 times below the GW channel noise floor.

Angular control noise arises from noise in the alignment sensors (both optical levers and wavefront sensors), propagating to the test masses through the alignment control servos. Though these feedback signals affect primarily the test mass orientation, there is always some coupling to the GW degree-of-freedom because the laser beam is not perfectly aligned to the center-of-rotation of the test mass surface (Kawamura and Zucker, 1994). Angular control noise is minimized by a combination of filtering and parameter tuning. Angle control bandwidths are 10 Hz or less and strong low-pass filtering is applied in the GW band. In addition, the angular coupling to the GW channel is minimized by tuning the center-of-rotation, using the four actuators on each optic, down to typical residual coupling levels of $10^{-3} - 10^{-4}$ m/rad.

1.4.5 Actuation Noise

The actuator noise term includes the electronics that produce the coil currents keeping the interferometer locked and aligned, starting with the digital-to-analog converters (DACs). The actuation electronics chain has extremely demanding dynamic range requirements. At low frequencies, control currents of ~ 10 mA are required to provide $\sim 5\,\mu$m of position control, and tens of mA are required to provide ~ 0.5 mrad of alignment bias. Yet the current noise through the coils must be kept below a couple of $\mathrm{pA}/\sqrt{\mathrm{Hz}}$ above 40 Hz. The relatively limited dynamic range of the DACs is managed with a combination of digital and analog filtering: the higher frequency components of the control signals are digitally emphasized before being sent to the DACs, and then de-emphasized following the DACs with complementary analog filters. The dominant coil current noise comes instead from the circuits that provide the alignment

bias currents, followed closely by the circuits that provide the length feedback currents.

1.4.6 Additional Noise Sources

In the 50 – 100 Hz band, the known noise sources typically do not fully explain the measured noise. Additional noise mechanisms have been identified, though not quantitatively established. Two potentially significant candidates are nonlinear conversion of low frequency actuator coil currents to broadband noise (upconversion), and electric charge build-up on the test masses. A variety of experiments have shown that the upconversion occurs in the magnets (neodymium iron boron) of the coil-magnet actuators, and produces a broadband force noise, with a f^{-2} spectral slope; this is the phenomenon known as Barkhausen noise (Cote and Meisel, 1991). The nonlinearity is small but not negligible given the dynamic range involved: 0.1 mN of low-frequency (below a few Hertz) actuator force upconverts of order 10^{-11} N rms of force noise in the 40 – 80 Hz octave. This noise mechanism is significant primarily below 80 Hz, and varies in amplitude with the level of ground motion at the observatories.

Regarding electric charge, mechanical contact of a test mass with its nearby limit-stops, as happens during a large earthquake, can build up charge between the two objects. Such charge distributions are not stationary; they tend to redistribute on the surface to reduce local charge density. This produces a fluctuating force on the test mass, with an expected f^{-1} spectral slope. Although the level at which this mechanism occurs in the interferometers is not well-known, evidence for its potential significance comes from a fortuitous event with L1. Following a vacuum vent and pump-out cycle partway through the S5 science run, the strain noise in the 50 – 100 Hz band went down by about 20%; this was attributed to charge reduction on one of the test masses.

In addition to these broadband noises, there are a variety of periodic or quasi-periodic processes that produce lines or narrow features in the spectrum. The largest of these spectral peaks are identified in Fig. 1.6. The groups of lines around 350 Hz, 700 Hz, *et cetera* are vibrational modes of the wires that suspend the test masses, thermally excited with kT of energy in each mode. The power line harmonics, at 60 Hz, 120 Hz, 180 Hz, *et cetera* infiltrate the interferometer in a variety of ways. The 60 Hz line, for example, is primarily due to the power line's magnetic field coupling directly to the test mass magnets. As all these lines are

narrow and fairly stable in frequency, they occupy only a small fraction of the instrument spectral bandwidth.

1.5 Future directions

From its inception, LIGO was envisioned not as a single experiment, but as an on-going observatory. The facilities and infrastructure construction were specified, as much as possible, to accommodate detectors with much higher sensitivity. There are a set of relatively minor improvements to the first generation instruments (Adhikari et al., 2006) that yield a factor of 2 increase in strain sensitivity and a corresponding factor of 8 increase in the probed volume of the universe. The two most significant enhancements are higher laser power and a new, more efficient readout technique for the GW channel. Higher power is delivered by a new master oscillator-power amplifier system, emitting 35 W of single frequency 1064 nm light (Frede et al., 2007), 3.5 times more power than the initial LIGO lasers. For the readout, a small mode-cleaner cavity is inserted in the AS beam path, between the Faraday isolator and the length photodetectors. This cavity filters out RF sidebands and the higher-order mode content of the AS port light, reducing the shot-noise power. Instead of RF heterodyning, signal detection is done by slightly offsetting the differential arm length from the dark fringe, and using the resulting carrier field as the local oscillator in a DC homodyne detection scheme. These improvements (known collectively as Enhanced LIGO) were implemented and commissioned on H1 and L1.

Significantly greater sensitivity improvements are possible with more extensive upgrades. Advanced LIGO's significantly improved technology will achieve a factor of at least 10 in sensitivity over the initial LIGO interferometers and will lower the seismic wall frequency down to 10 Hz (Fritschel, 2003; team, 2006). Advanced LIGO has been funded by the National Science Foundation, begining in April 2008. Installation of the Advanced LIGO interferometers is planned to start in early-2011.

The Advanced LIGO interferometers are configured like initial LIGO – a power-recycled Fabry-Perot Michelson – with the addition of a *signal recycling* mirror at the anti-symmetric output. Signal recycling gives the ability to tune the interferometer frequency response, so that the point of maximum response can be shifted away from zero frequency (Meers, 1988). The laser wavelength stays at 1064 nm, but an additional high-power stage brings the laser power up to 200 W (Willke et al., 2008).

The test masses will be significantly larger – 40 kg – in order to reduce radiation pressure noise and to allow larger beam sizes. Larger beams and better dielectric mirror coatings combine to reduce the test mass thermal noise by a factor of 5 compared to initial LIGO (Harry et al., 2007).

The test mass suspensions become significantly more intricate to provide much better performance. They incorporate four cascaded stages of passive isolation, instead of just one, including vertical isolation comparable to the horizontal isolation at all stages except one (Robertson et al., 2002). The test mass is suspended at the final stage with fused silica fibers rather than steel wires; these fibers have extremely low mechanical loss and will reduce suspension thermal noise nearly a hundred-fold (Heptonstall et al., 2006). The current passive seismic isolation stacks that support the suspensions are replaced with two-stage active isolation platforms (Abbott et al., 2002). These stages are designed to actively reduce the ground vibration by a factor of $\sim$1000 in the $1-10$ Hz band, and provide passive isolation at higher frequencies. The combination of the isolation platforms and the suspensions will reduce seismic noise to negligible levels above approximately 10 Hz.

The successful operation of Advanced LIGO is expected to transform the field from GW detection to GW astrophysics. We illustrate the potential using compact binary coalescences. Detection rate estimates for CBCs can be made using a combination of extrapolations from observed binary pulsars, stellar birth rate estimates, and population synthesis models. There are large uncertainties inherent in all of these methods, however, leading to rate estimates that are uncertain by several orders of magnitude. We therefore quote a range of rates, spanning plausible pessimistic and optimistic estimates, as well as a likely rate. The rate estimates for Advanced LIGO for CBCs involving NSs of typically $1.4M_{\odot}$ and/or BHs of up to $10M_{\odot}$ are: $0.4-400\,yr^{-1}$, with a likely rate of $40\,yr^{-1}$ for NS-NS binaries; $0.2-300\,yr^{-1}$, with a likely rate of $10\,yr^{-1}$ for NS-BH binaries; $2-4000\,yr^{-1}$, with a likely rate of $30\,yr^{-1}$ for BH-BH binaries.

References

Abbott, B. P. et al. 2009. *Reports on Progress in Physics*, **72**(7), 076901.
Abbott, R., et al. 2002. *Classical and Quantum Gravity*, **19**(7), 1591–1597.
Abbott, R. et al. 2004. *Class. Quant. Grav.*, **21**, S915–S921.

Abbott, R. S. and King, P. J. 2001. *Review of Scientific Instruments*, **72**(2), 1346–1349.

Abramovici, A. et al. 1992. *Science*, **256**, 325–333.

Adhikari, R. 2004. *Sensitivity and Noise Analysis of 4 km Laser Interferometric Gravitational Wave Antennae*. Ph.D. thesis, MIT.

Adhikari, R., Fritschel, P., and Waldman, S. 2006. Tech. rept. LIGO-T060156-01. LIGO Project.

Ballmer, S. et al. 2005. Tech. rept. LIGO-T050064-00. LIGO Project.

Cagnoli, G., et al. 1999. *Physics Letters A*, **255**(4-6), 230 – 235.

Camp, J. B., et al. 2000. *Journal of the Optical Society of America A*, **17**, 120–128.

Cote, P. J. and Meisel, L. V. 1991. *Phys. Rev. Lett.*, **67**, 1334. (c) 1991: The American Physical Society.

Cutler, C. and Thorne, K. S. 2002. *ArXiv General Relativity and Quantum Cosmology e-prints*.

D'Ambrosio, E. and Kells, W. 2006. *Phys. Rev.*, **D73**, 122002.

Drever, R. W. P., et al. 1983. *Applied Physics B: Lasers and Optics*, **31**, 97–105.

Evans, M., et al. 2002. *Opt. Lett.*, **27**(8), 598–600.

Fabbro, R. D. and Montelatici, V. 1995. *Appl. Opt.*, **34**(21), 4380–4396.

Frede, M., et al. 2007. *Opt. Express*, **15**(2), 459–465.

Fritschel, P. 2003. Pages 282–291 of: Cruise, M. and Saulson, P. (eds), *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference, vol. 4856.

Fritschel, P., et al. 1998. *Appl. Opt.*, **37**(28), 6734–6747.

Fritschel, P., et al. 2001. *Appl. Opt.*, **40**(28), 4988–4998.

Giaime, J., et al. 1996. *Review of Scientific Instruments*, **67**(1), 208–214.

Gillespie, A. and Raab, F. J. 1994. *Phys. Lett.*, **A190**, 213–220.

Gonzalez, G. 2000. *Classical and Quantum Gravity*, **17**(21), 4409–4435.

Gretarsson, A. M., et al. 2007. *J. Opt. Soc. Am. B*, **24**(11), 2821–2828.

Harry, G. M., et al. 2002. *Classical and Quantum Gravity*, **19**(5), 897–917.

Harry, G. M., et al. 2007. *Classical and Quantum Gravity*, **24**(2), 405–415.

Hefetz, Y., Mavalvala, N., and Sigg, D. 1997. *Journal of the Optical Society of America B Optical Physics*, **14**, 1597–1605.

Heptonstall, A., et al. 2006. *Physics Letters A*, **354**, 353–359.

Kawamura, S. and Zucker, M. E. 1994. *Appl. Opt.*, **33**(18), 3912–3918.

Landry, M. and the LIGO Scientific Collaboration. 2005. *Classical and Quantum Gravity*, **22**, 985–+.

Levin, Y. 1998. *Phys. Rev.*, **D57**, 659–663.

Meers, B. J. 1988. *Phys. Rev.*, **D38**, 2317–2326.

Meers, B. J. 1989. *Phys. Lett.*, **A142**, 465–470.

Morrison, E., et al. 1994. *Applied Optics*, **33**, 5041–5049.

Niebauer, T. M., et al. 1991. *Phys. Rev. A*, **43**(9), 5022–5029.

Ottaway, D., et al. 2006. *Opt. Lett.*, **31**(4), 450–452.

Penn, S. D., et al. 2006. *Physics Letters A*, **352**(1-2), 3 – 6.

Rakhmanov, M., Romano, J. D., and Whelan, J. T. 2008. *Classical and Quantum Gravity*, **25**(18), 184017 (13pp).

Regehr, M. W., Raab, F. J., and Whitcomb, S. E. 1995. *Opt. Lett.*, **20**(13), 1507–1509.

Rempe, G., et al. 1992. *Opt. Lett.*, **17**(5), 363–365.

Robertson, N. A., et al. 2002. *Classical and Quantum Gravity*, **19**, 4043–4058.

Saulson, P. R. 1990. *Phys. Rev.*, **D42**, 2437–2445.

Savage, R., King, P., and Seel, S. 1998. *LASER PHYSICS*, **8**(3), 679 – 685.

Sigg, D. 1997. Tech. rept. LIGO-T970084-00. LIGO Project.

Skeldon, K. D., et al. 1996. *Review of Scientific Instruments*, **67**(7), 2443–2448.

team, A. L. 2006. Tech. rept. LIGO-M060056-08. LIGO Project.

Vinet, J.-Y., Brisson, V., and Braccini, S. 1996. *Phys. Rev.*, **D54**, 1276–1286.

Vinet, J.-Y. et al. 1997. *Phys. Rev.*, **D56**, 6085–6095.

Walsh, C. J., et al. 1999. *Appl. Opt.*, **38**(13), 2870–2879.

Wei, D. T. 1989. *Appl. Opt.*, **28**(14), 2813–2816.

Weisberg, J. M. and Taylor, J. H. 2005. Pages 25–+ of: Rasio, F. A. and Stairs, I. H. (eds), *Binary Radio Pulsars*. Astronomical Society of the Pacific Conference Series, vol. 328.

Willke, B., et al. 1998. *Opt. Lett.*, **23**(21), 1704–1706.

Willke, B., et al. 2008. *Classical and Quantum Gravity*, **25**(11), 114040–+.

Winkler, W., et al. 1991. *Phys. Rev. A*, **44**(11), 7022–7036.

Yoshida, S. et al. 2000. *Gravitational Wave Detection II*. Universal Academy Press, Tokyo. Chap. Recent Development in the LIGO Input Optics.

Zucker, M. E. and Whitcomb, S. E. 1996. Pages 1434–+ of: Jantzen, R. T., Mac Keiser, G., and Ruffini, R. (eds), *Proceedings of the Seventh Marcel Grossman Meeting on recent developments in theoretical and experimental general relativity, gravitation, and relativistic field theories*.