

3 LIGO: The Basic Idea

LIGO uses interference of light waves to detect gravitational waves. This is accomplished using a device called an interferometer. LIGO actually uses several interferometers for different tasks, but we will begin this discussion by describing how interferometry can be used to detect the passage of gravitational waves.

3.1 What is an Interferometer?

An interferometer is a device used to interfere one light beam with another. LIGO uses a Michelson interferometer with Fabry-Perot resonant cavities. The first is used to measure the stretching and compressing of spacetime by gravitational waves. The second tunes and amplifies the laser beam inside the Michelson interferometer. In this section we consider Michelson interferometers. Fabry-Perot will be saved until the section about noise mitigation.

The Michelson interferometer was devised by A. A. Michelson in the 19th century. Its design is simple, as shown schematically in Figure 3.1. The basic elements are a beam splitter, which is a partially-reflecting mirror, and two fully-reflecting mirrors, one placed at the end of each arm of the interferometer. The beam splitter is arranged so that part of an incoming light beam is reflected at 90° (b_R), while the remaining light is transmitted along the original path (b_T). The interferometer, thus, has the shape of an “L.” Ideally, both beams travel an identical distance before being reflected by the mirrors at the end of the arms. They then return to the beam splitter along their path of approach. When they encounter the beam splitter, part of the light is again transmitted and part reflected, with some of each going to a photo detector, shown at the bottom of Figure 3.1.

Notice that the beam splitter introduces an asymmetry between the two beams. If its mirrored surface is on the side toward the input laser, then the reflection of the laser upward to form b_R introduces a phase shift of π : b_R is exactly out of phase with the input beam. However, the reflection of b_T upon its return to the beam splitter produces no phase shift because the reflection happens inside the beam splitter at the glass/vacuum boundary. Both beams undergo a π phase shift at the ends of their respective arms, and so the two beams are exactly out of phase upon their arrival at the photo detector, and no light is observed there. This is true as long as the arm lengths are both an integral number of wavelengths long.

On the other hand, if the paths traveled by b_T and b_R differ, then the cancellation will in general not be exact. Some light might be observed at the photo detector. We can, of course, have cancellation with different path lengths. In general we will have total cancellation whenever the path difference is an integer number of wavelengths, including zero.

The exact manner in which the recombined beams interact depends on the difference in

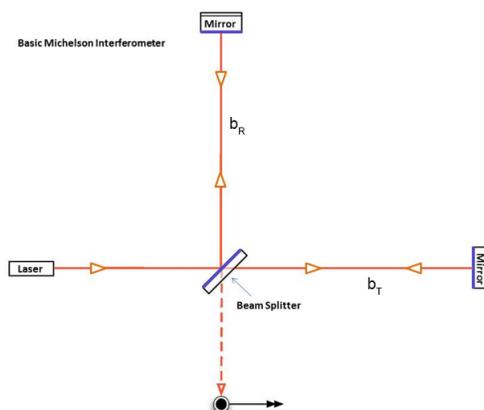


Figure 3.1: This diagram illustrates the basic idea of a Michelson interferometer. Light, in this case from a laser, enters from the left. The beam is split into two perpendicular paths, which are then reflected back toward the beam splitter using mirrors. The beam is then split again, and part of it is sent to a detector, shown at the very bottom of the diagram. Credit: Caltech/MIT/LIGO Lab

path length between the arms. As noted, if the path lengths are the same (or differ by an integral number of wavelengths), then the beams arrive π out of phase⁷. By adjusting the length of the arms we can achieve whatever relative phase we desire for the two beams, and thus whatever intensity we desire at the photo detector. Conversely, for motions that are small compared to the wavelength of the laser light, we can infer the difference in the path length from changes in light intensity at the photo detector. This interference phenomenon is exactly what LIGO uses to detect gravitational waves.

Use the [Michelson Interferometer](#) applet at GWO to explore the way the output signal from the device depends on the relative lengths of its arms. Try to make the output signals (there are two of them in the applet) go to zero. Are the conditions needed to make each signal vanish the same? Try to explain what is happening as each signal goes to zero.

Another interesting GWO applet is the [Cavity Calculator](#). In this applet you can tune a simulated Fabry-Perot cavity within a Michelson interferometer and see how your tuning changes the power output of the device. One of the cavity mirrors is curved, while the other is flat. Three spatial modes are shown for the laser beam inside the cavity. Each has a different distribution of light intensity in the direction perpendicular to light propagation. One mode, (TEM_{00}), is the fundamental. The other two are higher order modes (denoted

⁷The observant reader will note that b_T traverses the glass of the beam splitter twice, whereas b_R traverses it only once. This difference introduces a small change in phase as well. To compensate, some Michelson interferometers (but not LIGO) place a non-reflecting glass plate in the arm with b_R to compensate.

TEM_{10} and TEM_{20}). These designations refer to the order of Laguerre polynomial used to represent them, with the first subscript being the radial index of the mode (which gives the number of nodes of the light intensity) and the second being the azimuthal index. The coordinate system used is plane-polar with the z -coordinate running down the axis of the cavity. You will note that the modes respond differently as the resonant cavity tuning is changed, a result of their different phases inside the cavity. The applet allows adjustments to the position and curvature of the end mirror, and the reflectivity of both mirrors.



Figure 3.2: The LIGO facility in Hanford, WA. In this aerial shot, the arms of the interferometer can be seen extending 4 km from the central building that houses the input laser and the photodetector. Credit: Caltech/MIT/LIGO Lab

3.2 How Does LIGO Detect Gravitational Waves?

As a gravitational wave passes by, the wave will alternately stretch and compress each arm of the interferometer. What's more, because of the way gravitational waves are structured (discussed below), compression in one side will occur when the other side is being stretched. Since the length of each arm has changed, the travel time of the light from the beam splitter to the end mirror (and back) has been changed. When the beams are recombined at the beam splitter, the differences in light travel time are realized as a phase difference, and this causes a change in light intensity at the photo detector. That is the signal recorded by the interferometer. It can be used to infer how each arm's length changed as the gravitational wave passed by. A simplified view of this process is provided by [an animation](#) on the LIGO website.

As we saw in Section 2, the amount of stretching and compression of each arm is generally described in terms of the fractional change of the arm's length, called the strain.

Strain is denoted by h and is defined mathematically in Equation 2.1

Strain is the natural way to think about gravitational waves because their mode of oscillation is to stretch and compress space along axes perpendicular to one another and to the direction of propagation of the waves. Typical modes of oscillation are shown in Figure 3.3. The two fundamental polarization modes for gravitational waves are depicted. In this perspective, the waves approach the viewer head-on, and the dots show the effect of passing waves on a set of freely-falling objects, called *test masses*, as the waves pass by. When no distortion is present the objects are arrayed in a circle.

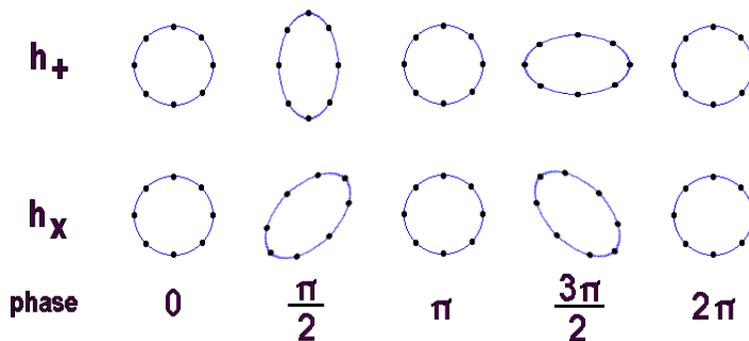


Figure 3.3: This shows the h_+ and h_\times polarizations for gravitational waves. The waves propagation is into or out of the page. They displace “test masses” represented by the small dots. Notice how the waves alternately stretch and then compress the space. Image: <http://www.johnstonsarchive.net/relativity/pictures.html>.

The amplitude of the motion in Figure 3.3 is greatly exaggerated. In reality, the strain caused by gravitational waves is expected to be so tiny that it would not be close to being visible. The typical smallest strain LIGO can detect is of the order $h \sim 10^{-21}$. That is why the interferometer arms have been made so long; at 4 km they are expected to be barely long enough to allow detection of the strongest astrophysical sources.

In Figure 3.3 the test masses represented by the dots are assumed to be freely falling. In the language of general relativity, they should be in inertial frames of reference. Gravitational waves can be viewed as small perturbations to the flat spacetime metric that describes inertial frames in special relativity, and such perturbations will disturb freely falling objects in those frames⁸. In any case, it is clear that the LIGO interferometers do

⁸See the materials from the first LIGO course (Geometry and Gravity in Weak Fields) if you are unsure what that means and would like to know more.

not occupy inertial reference frames and are not freely falling. They are sitting on Earth’s surface, not following their respective geodesics. However, it is not strictly necessary for the interferometers to be in inertial frames in order to detect the effect of passing gravitational waves. It is only necessary to approximate such frames along the dimension of the interferometer that is sensitive to small displacements, and that is what a Michelson interferometer is ideal at doing.

Since the device is only sensitive to displacements along the lengths of its arms, the experiment need only allow “free fall” in that dimension, and only for a distance comparable to the expected displacement caused by a passing gravitational wave. From the exercises in Section 1 you should be convinced that the allowable motions are not very large. You should have also seen that the motion of the test masses is larger for longer interferometer arms, but is still very small, even for the 4 km arms in LIGO.

LIGO creates a sufficient region of free-fall by making the mirror at the end of each of the arms the bob of a pendulum. Each pendulum lets its mirror move freely, for very small displacements, along the axis of that arm. If a gravitational wave passes by, the mirrors can move in and out slightly, thereby causing the length of the arms to oscillate with the wave. This motion creates a changing intensity at the photo detector. For this reason, the mirrors are generally referred to as *test masses*. LIGO uses many mirrors, but only the ones at the ends of the arms are test masses. We will adopt this nomenclature from now on to avoid confusion. The test masses used for LIGO are each 40 kg.

3.3 The LIGO Interferometer - Details

That is the basic idea of LIGO⁹. In practice, things are more complicated. A device sensitive enough to measure strains of 10^{-21} will show many effects beyond gravitational waves. For example, LIGO is sitting on the ground. As a result, tiny seismic vibrations will displace the test masses, and by far larger amounts than any gravitational wave is expected to do. LIGO is, in effect, the most sensitive seismometer ever built. Even the Brownian motion due to impacts with air molecules will have a measurable effect. Means must be employed to remove or mitigate the noise created by these effects, and others.

Other details are also important. One is related to the properties of light itself. For such tiny shifts in interferometer arm length, the phase shifts of the light will also be tiny. In order to measure them, we need to be able to see minute changes in the light intensity. But because the photons that make up the light beam are randomly distributed - like the sound of raindrops on a roof - there is uncertainty in the amount of light received at the photodetector. The *Poisson statistics* that determine those fluctuations become fractionally smaller with higher light intensity (or average number of photons), so we need high-power lasers and Fabry-Perot cavities to boost the number of photons, and thus lower

⁹Much of this section is based on material in [LSCollaboration \(2015\)](#)

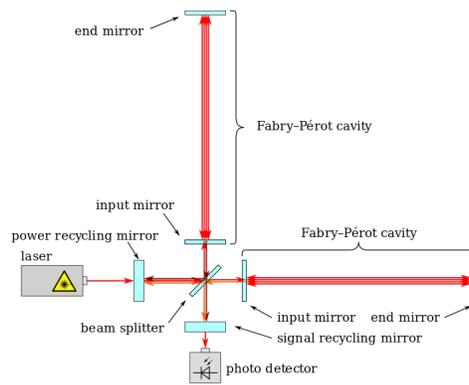


Figure 3.4: LIGO includes two Fabry-Perot cavities, one in each of its arms, to both stabilize the frequency and boost the power in its laser beam. In addition, it has a system to return light exiting the Michelson system that would otherwise travel back toward the input laser. These systems are shown schematically in this figure. Others are left off for clarity. Image Credit: Caltech/MIT/LIGO Lab

this noise. LIGO uses several techniques to manage these complications, which we now describe.

We should mention before going on that all of the optical elements mentioned from here on are contained within a high vacuum system and are on seismically isolated supports. The reasons for this will become clear in the next section.

Take a look at Figures 3.4 and 3.6. They show several of the modifications that LIGO employs to increase its sensitivity. The first of these is an *Input Mode Cleaner* (IMC - shown in Fig 3.6). The system consists of optical elements (mirrors) and electronics that remove unwanted light frequencies from the input laser, tuning it to the desired $1.064\ \mu\text{m}$ wavelength. This device is called a *resonant cavity*, and it is one of several that LIGO employs. The laser light bounces between two mirrors repeatedly, and the mirrors are adjusted so as to select the desired laser frequency and phase, and deselect others. Cleaning the input laser light is necessary because the laser does not emit perfectly coherent light at a single phase. It is close, at least by some criteria, but not close enough for this application. The IMC narrows the frequency fluctuation and phase spread of the light before it enters the interferometer.

On exiting the IMC, the laser passes through the *power recycling mirror* (PRM) system. This system does just what the name says. It boosts the light entering it, both from the input laser and light that exits the interferometer back toward the input laser, and collects it in another resonant cavity, similar to the one in the IMC. The light bounces back and forth in the cavity, building up the intensity of the laser. The PRM boosts the input laser light from 125 W to 5.7 kW before sending it into the interferometer.

The cavities used for the IMC and PRM are quite similar. A schematic for both is shown in Figure 3.5

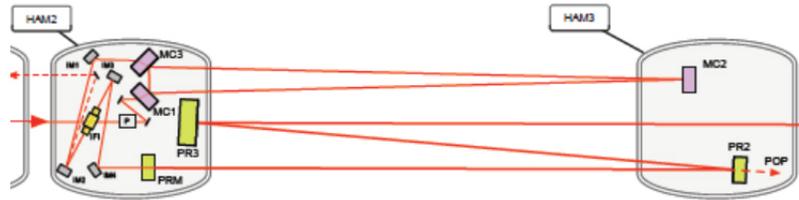


Figure 3.5: LIGO uses a Power Recycling Mirror (PRM) system to return lost light to the interferometer, thus boosting its power from 125 W as it leaves the Input Mode Cleaner (MC) to 5.7 kW when it leaves this system. The optics for these systems are enclosed in seismic isolation units called Horizontal Access Module, or HAMs. Image Credit: Caltech/MIT/LIGO Lab

After the light exits the PRM, it enters the Michelson interferometer, the first element of which is the beam splitter. In Figure 3.4, the light enters the interferometer from the left. The beam splitter then sends half the light up into the vertical arm in the figure, with the other half continuing in the original direction into the horizontal arm. Both beams encounter yet another resonant cavity in their respective arms. They are called *Fabry-Perot* cavities (FP). Each of the FP cavities are another kind of interferometer; it has a partially reflecting mirror at one end, where the light enters, and a fully reflecting one at the far end. You can probably see that, with the correct spacing between the mirrors, they can also be used as resonance cavities. This is how LIGO uses them. Each of the laser beams effectively traverses its cavity 450 times before exiting - back along the direction from which it originally came in this case. This has two functions. First, it boosts the power of the light at a chosen frequency, just as happened in the PRM. The second effect is to increase the effective length of each arm, and therefore also increase the sensitivity of the interferometer to changes in length of the arms. We have prepared some exercises in the homework for this section to help you understand how this happens.

The previous two exercises should have made it clear why LIGO uses a modified Michelson interferometer, with Fabry-Perot cavities along each arm. The LIGO arms are already quite long, at 4 km, but adding the FP cavities makes the effect length of each arm more like 1800 km long. That is a huge increase in sensitivity, while keeping the size of the experiment still relatively compact.

We mentioned earlier that the FP cavities boost the power in the laser. The beam enters the interferometer at 5.7 kW and is bumped up to 815 kW by the FP cavities. The

cavities also spread the beam out to a diameter of 12 cm from its initial size of ~ 1 mm when it exits the ICM. It is important that the beam not be too small while inside the FP cavities for a couple of reasons. First, a high intensity beam will heat the surfaces of the mirrors, and that will cause big problems. We discuss these in the next section. The other reason that a large beam is desirable is that it interacts with a larger part of the mirror surface, and thus with a larger number of atoms. More atoms helps to average out Brownian motion of the mirror surfaces, to which LIGO is sensitive. After all, the amount of displacement in each arm length by a passing gravitational wave is smaller than the diameter of a proton by a factor of a thousand.

Finally, after the two laser beams are combined back into one and exit the interferometer, they pass through one final resonant cavity. This cavity again boosts the strength of the beam, increasing the signal-to-noise before the light is sent to the photodetector.

3.4 Overview: Advanced LIGO

LIGO has a lot of parts. The diagram in Figure 3.6 gives a succinct overview of the various systems we have described. You might find it helpful in keeping track of the material you have just read. You should also look at the video on the class website. You will find it here: <https://vimeo.com/170085135>. The password is “gwavesarefun”.

The labels in Figure 3.6 are provided in the table below.

- *laser* - the input laser system
- ϕ_m - phase modulator. Measures the lengths of the various cavities and helps hold them on resonance.
- IMC - Input Mode Cleaner. A resonant cavity to narrow the phase and frequency of the laser light before insertion into the interferometer
- FI - Faraday Isolator. This is a sort of optical diode. It lets light pass in one direction, but not the other. This one is used to prevent light reflected by the interferometer from getting back into the laser.
- PRM - Power Recycling Mirrors. This resonant cavity boosts the power of the laser from its input value of 125 W up to 5.7 kW. The part of the returning beams that does not exit the beam splitter toward the photodetector is recycled by the PRM and sent back into the interferometer.
- BS - Beam splitter. This mirror is the first element of the Michelson interferometer. It sends half the light into one arm, and half the light into the other. It also combines the returning beams into an output beam and sends it out of the interferometer.

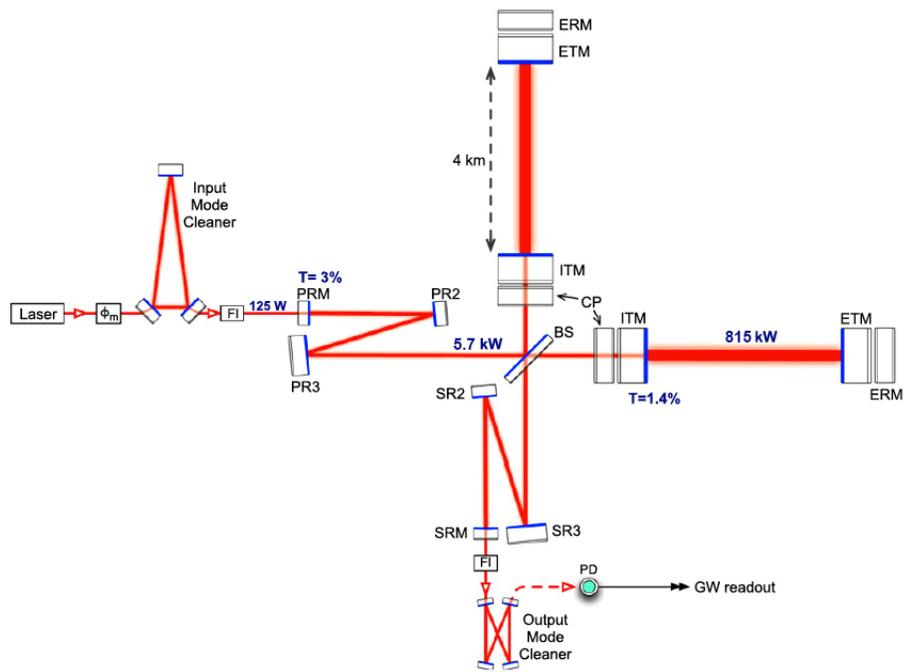


Figure 3.6: See bullet list for description of each labeled part in this diagram. Image taken from [LSCollaboration \(2015\)](#)

- CP - compensation plate. This is part of the Fabry-Perot (FP) cavity. It compensates for distortions to the shape of the mirrors in the FP that are caused by thermal distortions when the laser heats the reflective coatings on the test masses.
- ITM - Input Test Mass. This is the input side of the Fabry-Perot cavity. It has a partially reflecting surface, transmitting only 1.4% of the light.
- ETM - End Test Mass. Far end of the FP cavity. This is the test mass that moves away from the central part of the interferometer when a gravitational wave passes through. The ETM has a reflecting surface and is suspended on a four-stage pendulum, discussed in the next section.
- ERM - End Reaction Mass. This is a fused silica block that is used to help position and align the test mass. It contains conductive traces that allow it to be actuated electrostatically.
- SRM - Signal Recycling Mirrors. After exiting the interferometer, the light enters the SRM, where the now the signal is resonated, like the laser is by the PRM, to help establish the frequency response of the detector to gravitational waves. The beam diameter is reduced to ~ 8 mm from the ~ 12 cm diameter in each FP.
- FI - Faraday Isolator. This second FI prevents light from the OMC from being reflected back into the interferometer.
- OMC - Output Mode Cleaner. Removes the sidebands and cleans unwanted frequencies from the beam before sending it to the optical sensor.
- GWR - Gravitational Wave Readout. This is an optical detector that measures the intensity of the laser beam and converts it to an electrical signal. As the relative phase of the beams in the two FP cavities changes, the intensity of the combined beam at the GWR changes in response.

Aside from the input laser and phase modulator, all the components described above are housed in a high-vacuum environment and are seismically isolated. The systems are housed in HAM units: Horizontal Access Modules. These, as you might guess, allow access to the insides of the vacuum system for installation and maintenance of hardware. The image in Figure 3.7 shows how this is done. Note the vacuum tubes on either side of the HAM that connect it to the rest of the experiment. The towers holding the optics (beside the bunny-suited scientist) sit upon a In-vacuum Seismic Isolator (ISI), the need for which will be discussed below.

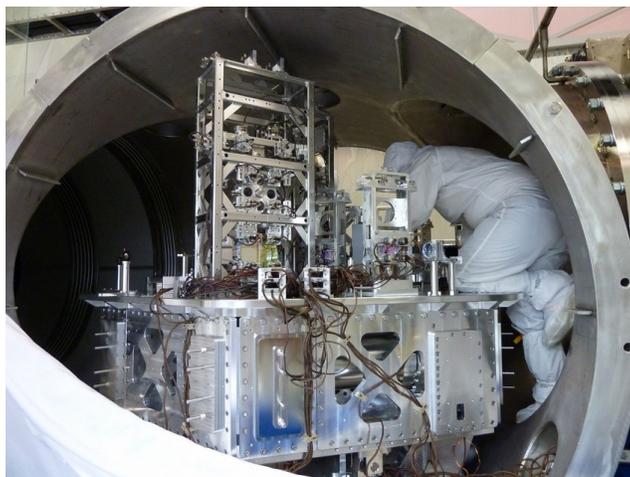


Figure 3.7: A look into one of the instrument vacuum chambers, called HAMs, that hold instruments inside the LIGO interferometer vacuum system. Side covers have been removed to allow installation or adjustment of components in this photo - while the machine is shut down, of course. Image Credit: Caltech/MIT/LIGO Lab

3.5 A Last Small Point About Interferometers

In our discussion of the basic-model Michelson interferometer, the kind without all the resonant cavities and Fabry-Perot arms, we said that when the arms have the same length, the two beams will interfere destructively upon exiting the interferometer. This is true, but it is a bit misleading in the case of LIGO.

We are looking for LIGO's arms to be altered by a passing gravitational wave. The way this happens is that the wave temporarily squeezes and stretches the space along the arms as it passes. This stretching and squeezing of space does not just affect the distance between the mirrors in the FP cavities, it stretches and squeezes the light in those cavities, and by exactly the same amount. How, then, does LIGO measure the space distortions?

This seems a bit puzzling if you have in your mind that the number of wavelengths between the ends of the FP are what matter. The number of wavelengths does not change, but that's okay, because it isn't actually what matters. Instead, consider the light travel time along each arm as the space gets squeezed and stretched. The speed of light is not changing along the light path, only the wavelength of the light changes. That isn't important though, as all that matters is the light travel time along each arm. If the beams start out in phase at exactly the same time, and then one returns before the other, they will no longer be in phase. This phase shift will destroy the exact cancellation in the outgoing beam, and so the photo detector will record some light. The effect is discussed in detail in [Saulson \(1997\)](#).

That's it, at least in overview. The version of the LIGO detector we have described is its second iteration, called Advanced LIGO, or just aLIGO. Earlier versions did not detect gravitational waves, but they performed necessary shakeouts and tests of systems, technology and analysis software. They paved the way for the success of the current observatory. The development continues even now, with improvements planned that will make future generations of LIGO detectors even more sensitive.

Even with all the additional elements that stabilize the laser beam and boost its power, there are sources of noise that can mask a gravitational wave signal. Most of the upgrades to the LIGO systems are made to diminish or eliminate this noise. The techniques LIGO uses to surmount these challenges, either in current or planned incarnations of the detector, will be the subject of the next section.