LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

August 3, 2009

Wiener Filter Optimization for Seismic and Acoustic Noise Cancellation

Clara Bennett
Mentors: Rana Adhikari and Jenne Driggers

Distribution of this document:

LIGO Scientific Collaboration

California Institute of Technology
LIGO Project, MS 18-34
Pasadena, CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project, Room NW17-161
Cambridge, MA 02139
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

LIGO Hanford Observatory
Route 10, Mile Marker 2
Richland, WA 99352
Phone (509) 372-8106
Fax (509) 372-8137
E-mail: info@ligo.caltech.edu

LIGO Livingston Observatory
19100 LIGO Lane
Livingston, LA 70754
Phone (225) 686-3100
Fax (225) 686-7189
E-mail: info@ligo.caltech.edu

http://www.ligo.caltech.edu/

1 Introduction

The purpose of this project is to improve the ability to subtract seismic and acoustic noise from the interferometer signal, taking steps toward the eventual goal of an adaptive filtering system. Currently, there is a static Wiener filtering program in place, but little had been done previously to optimize the filtration with a limited number of noise-witness-measurement instruments.

2 Wiener Filtering [7]

Wiener filters comprise a special class of linear optimum discrete-time filters. As the name suggests, these filters are linear, which makes the mathematics far simpler, and operate in discrete time, allowing for implementation in digital devices. The details of the filter also depend on whether the impulse response is infinite or finite, and the statistical criterion used to optimize the filter.

Infinite response (IIR) filters have feedback loops that feed back to the input, while finite response (FIR) filters utilize only forward paths. The advantage of using FIR filters is their inherent stability—the feedback loops of IIR filters can unwanted oscillations if they are improperly designed. The stability issue for IIR filters is manageable, but if the filter is also required to be *adaptive*—able to track a non-stationary minimum on the error performance surface—the compounded stability problems of adaptive *and* IIR filtration become far too difficult to be practical. For these reasons, the FIR filters (which are, in fact, a special case of IIR) are preferred.

In terms of potential statistical criteria, the choices include minimizing the mean-square of the error estimate, the expectation value of the error, or the expectation third or greater powers of the error. Of these, the first choice is best, as it yields a cost function, J, which is second-order in dependence on the filter coefficients. Thus, J has a distinct and unique minimum, which conveniently provides a single point of optimum filtration [1].

3 Mode Cleaner

Before entering into the interferometer, the laser beam is run through a triangular set of mirrors, known as a *mode cleaner*. The optics are suspended on pendula which are supported by a series of mass-spring seismic isolation stacks, just like the test masses in the main body of the interferometer. The mode cleaner's purpose is to suppress the non-fundamental modes in the laser beam, as well as to stabilize the frequency of the laser and reduce jitter.

[1] For a mathematical basis for Wiener filtering, see Section B in the Appendices.

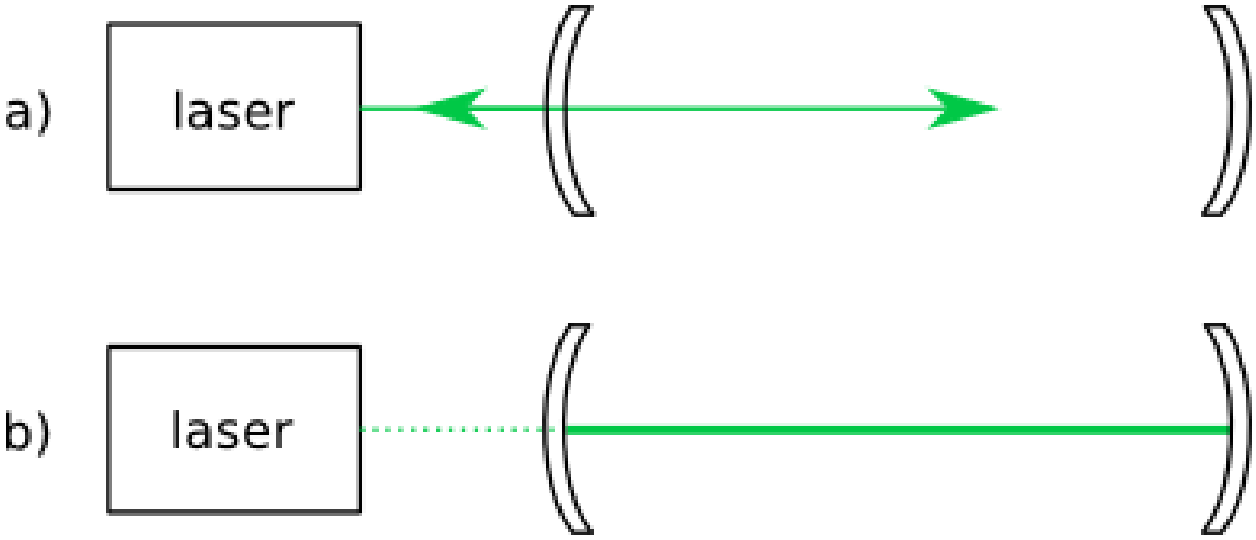

Figure 1 *– A Fabry-Perot Cavity*
a) The laser beam entering the cavity is partially reflected and partially transmitted.
b) If the cavity is the proper length, light will be stored inside, interfering constructively and building up power that is far greater than that of the incoming beam. At equilibrium, the light being transmitted at the first mirror will exactly cancel (destructively interfere) with the incoming beam.

3.1 Fabry-Perot Cavities

A basic Fabry-Perot cavity acts as a ”bottle” for laser light (see Figure 1), provided that the cavity is exactly an integer number of wavelengths of the light being stored. It consists of two highly reflective mirrors. A laser beam is pumped in through the first mirror and bounces back off the far end. If the two beams are in phase, they will constructively interfere. Eventually, the trapped light reaches an equilibrium such that the incoming power is equivalent to the power escaping the cavity from scattering, etc.

The LIGO mode cleaner has a triangular structure (see Figure 2), but the basic principle is the same. The cavity is designed to allow only the desired modes of light to resonate, suppressing all others. The advantage of the three-mirror design over the simpler two-mirror cavity is that the triangular setup prevents reflection of the light back into the pre-stabilized laser (PSL). There can still be some reflection due to scattering on the mirrors, which have finite roughness, especially if the triangle has a very narrow angle, but three mirrors still have a clear advantage over two in this respect.

3.2 Mode Cleaner Feedback

The mode cleaner is held at a fixed perimeter length using the Pound-Drever-Hall (PDH) locking scheme. The basic idea of the PDH technique is to lock the laser to an isolated Fabry-Perot reference cavity: the cavity, which is highly frequency sensitive, provides a measurement of the laser’s

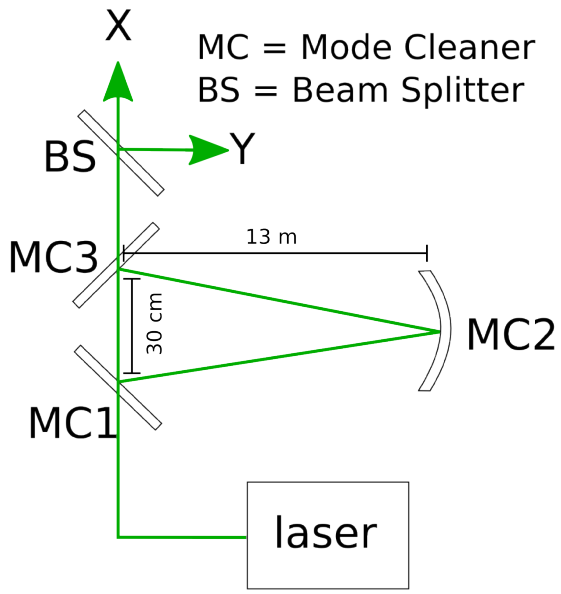

Figure 2 *– Diagram of the mode cleaner. The laser light continues on from the beam splitter into the X and Y arms of the interferometer. Two sets of three accelerometers each are located at the MC1 and MC2 optical chambers.*

frequency, which is then fed back into the laser to servo stabilize the frequency [6] [4] [2]. This scheme is used as the first stage of stabilization for the LIGO laser [8].

The mode cleaner is then locked to the laser beam, actuated by the Optical Shadow Sensor and Magnetic Actuators (OSEMs—see Figure 3). As the name indicates, the OSEMs drive the mirrors magnetically, to avoid transferring mechanical vibration. They can adjust the cavity length as necessary to keep the mode cleaner in lock and resonating only the desired modes. Information about how the mirrors *would* have moved without the actuators is taken from the force applied to the mirror by the OSEMs to maintain stability. Thus the mode cleaner length signal (MC_L) does not reflect actual length changes in the mode cleaner: rather, it records the length changes that were prevented by the OSEMs.

The OSEMs also have a shadow sensor component: the mirrors themselves are mounted with permanent magnets which swing between the arms of the OSEM when the mirror moves. On one arm, an LED with wavelength $\lambda = 880nm$ is focused into a narrow beam to a sensor on the other arm. The photodiode is covered by a thin, coated glass filter that only accepts light from a small range of angles and rejects longer wavelengths, so that no light from the main laser ($\lambda = 1064nm$) can trigger the photodiode. The shadow sensor acts to prevent amplification of mirror motion at high frequencies and to begin the locking process. If the mirrors oscillate quickly enough, the feedback loop cannot respond quickly enough to adjust the mirror position before it reverses direction and thus the actuators end up pushing in the wrong direction. The OSEM detects these erratic motions and shuts down the feedback loop. The OSEMs are also used to get the mirrors

[2] For more on the Pound-Drever-Hall scheme, see Section A.2 in the Appendices.

close to the proper lock position, as the PDH technique will fail if the mode cleaner begins too far from resonance.

3.3 Isolation Stacks

The interferometer is highly sensitive to changes in cavity length, which is excellent for trying to detect gravitational waves, but it also means that noise sources need to be well suppressed in order to take advantage of the sensitivity. Unfiltered normal ground motion would completely obliterate any gravity wave signal, so the instrument needs to be isolated from seismic motions.

The basic concept of vibration isolation is to connect the mass we wish to isolate to the ground through something with spring-like properties. Ignoring dissipative forces for the moment, we can describe a simple, one-dimensional mass-spring-ground system as

$$m\ddot{x}_m = -k(x_m - x_g), \tag{1a}$$

where x_m is the position of the mass and x_g is the position of the ground in an inertial reference frame (so that the quantity $x_m - x_g$ is the change in length of the spring). Recall that a mass-spring system has a resonant frequency $\omega_0 = \sqrt{\frac{k}{m}}$. Then, transforming Eq. 1a into frequency space, we get

$$\omega^2 x_m = -\frac{k}{m}(x_m - x_g) = \omega_0^2(x_m - x_g). \tag{1b}$$

We can then express the transfer function from ground motion, x_g, to mass motion, x_m, as

$$\frac{x_m}{x_g} = \frac{\omega_0^2}{\omega_0^2 - \omega^2}. \tag{2}$$

If we consider end behaviors, we see that the spring acts as a low-pass filter:

$$\frac{x_m}{x_g} \approx \frac{\omega_0^2}{\omega^2}, \qquad \omega \gg \omega_0, \tag{3a}$$

while,

$$\frac{x_m}{x_g} \approx 1, \qquad \omega \ll \omega_0. \tag{3b}$$

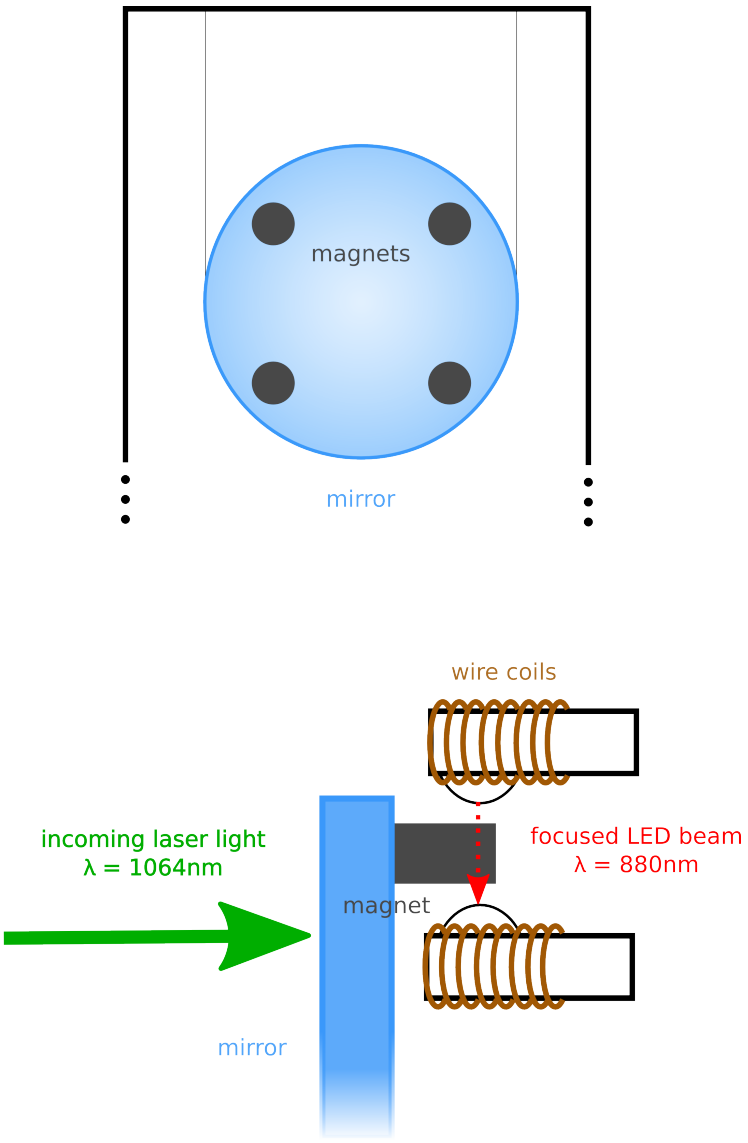

Figure 3 *– The mode cleaner mirror suspension, indicating the positioning of the OSEM magnets (top). A side view of the mirror and the OSEM. The magnet attached to the mirror interrupts the LED beam if the mirror moves, which the OSEM senses. It then readjusts the mirror position with the induced magnetic field created when current is run through the wire coils (bottom).*

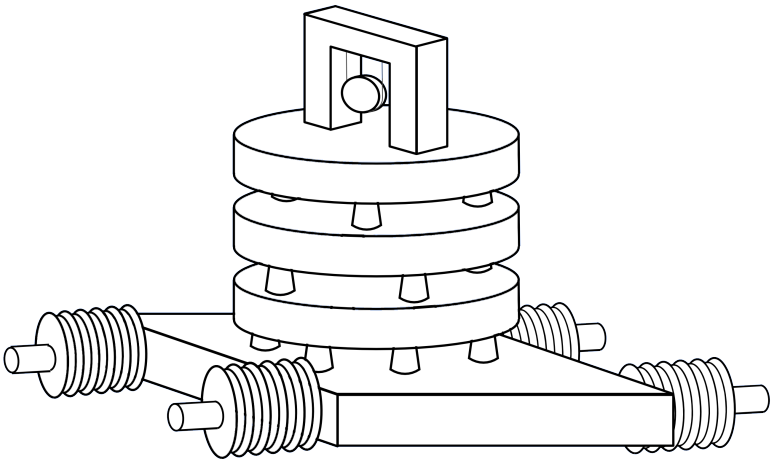

Figure 4 *– A simplified rendering of the mode cleaner isolation stacks. A single column of masses (the main interferometer test masses are mounted on three-legged isolation stacks) stacked with springs supports the top table, on which the test mass mirror pendulum suspension is mounted. The base table is attached to the ground through metal bellows, which further isolate the system from translational motions.*

Essentially, when the mass-spring configuration is driven at frequencies that are large compared to the fundamental frequency of the system, the inertia of the mass prevents it from being translated appreciably before the driving force reverses direction. At the other extreme, the driving force moves so slowly that the spring is never appreciably deformed and thus acts as a rigid mounting structure [6].

All of the test masses, including the mirrors in the mode cleaner, are mounted on complex passive isolation systems (see Figure 4). As previously mentioned, the test masses are all hung as pendula, largely eliminating the transmission of high-frequency motion to the mirrors. The pendulum suspensions are then mounted on seismic isolation stacks, which consist of four layers of steel masses separated by lossy springs [3]. Each layer in the stack stifles the motion of the layer below it. The seismic stacks do not rest directly on the ground—they are mounted through metal bellows, rather than a rigid structure, which further isolates the chamber from translational motions.

The complexity of the isolation stacks means that the coupling of ground motion to test mass motion can be quite complicated. The purpose of the Wiener filter is to account for these relationships and subtract out the noise, but, having a limited number of instrument signals from which to calculate the filter, the mechanics of the isolation stack cannot be ignored. Identifying which of the six degrees of freedom to measure is key and can be aided by even a simple model of the isolation structure, which is currently being worked on.

[3] The springs are made of materials that dissipate the vibrational energy viscoelastically, rather than viscously (in proportion to the velocity), as the former suppresses the rigid-body modes well while also outperforming the latter in high-frequency suppression through the stack levels [5]. The transfer function for ideal viscous damping is $\frac{x_m}{x_g} = \frac{\omega_0^2}{\omega_0^2 - \omega^2 i\omega b/m}$, while viscoelastic damping is described by $\frac{x_m}{x_g} = \frac{\omega_0^2(1+i\phi)}{\omega_0^2 - \omega^2 + i\omega_0^2\phi}$, where b and ϕ are constants that characterize the strength of the damping [6].

4 Instrument Position Optimization

Currently, there are six accelerometers available, four mounted on individual plates, and two mounted on a single plate at right angles to one another. There are also two three-dimensional Guralp seismometers and a single-axis Ranger seismometer measuring in the y-direction at MC2. We are trying to optimize the positioning of these instruments in order to get the best Wiener filtering results. At this time, only the accelerometers' placements have been adjusted.

Initially, I calculated new Wiener filters for each accelerometer configuration to determine which provided the best noise subtraction for the mode cleaner length (MC_L) signal. However, the noise in the signal varies greatly throughout the day, getting more noisy during the day and quieter at night. It is not so easy to compare the quality of the Wiener filters when the raw signal is so variable (see Figure 5), and it would be rather inefficient to try to take the data at the same time of day for each setup, especially with no guarantee that noise levels would be the same day-to-day.

Another way to evaluate the effectiveness of each instrument position is to look at the *coherence* between the instrument signal and the MC_L signal. The coherence is a measurement of the relatedness of two signals, similar in function to the r^2 correlation parameter that characterizes goodness-of-fit for linear regression lines. Also, instrument signals with good cohesion at some fixed phase difference will still be reflected in the coherence function—an important feature for the project, as our sensing instruments detect propagating motion and will have some nonzero physical separation. The coherence is given by [4]

$$C_{xy} = \frac{P_{xy}}{P_{xx}P_{yy}}, \tag{4}$$

where P_{xy} denotes the *cross-spectral density* between signals x and y, and P_{xx} is the *power spectral density* of x with itself (also referred to as the *autospectral density*). The coherence of the instruments with the MC_L signal does not seem to vary with any significance throughout the day, as we see in Figure 6. Combined with the fact that shorter time scales are required to calculate the coherence reliably than the Wiener filters, coherence measurement appears to be the better tool for optimizing the witness instrument positions.

So far, we have only used the accelerometers and seismometers to measure translational motions, but there are three additional rotational modes to consider in terms of ground motion transfer. So, the next step in optimizing instrument locations will be to determine which of the rotational directions are the most important and how to arrange the instruments to best record the motions. Rotation is more tricky than translation, as two instruments are required to record rotational motion in only one direction, and the magnitudes are much smaller, so we need to be much more sensitive to be able to measure it. Hopefully, the two three-axis Guralp seismometers will be operational soon so that low-frequency rotational motions can be investigated as well.

[4] For a more detailed explanation of the coherence, related quantities, and measurement techniques, please see Section C in the Appendices.

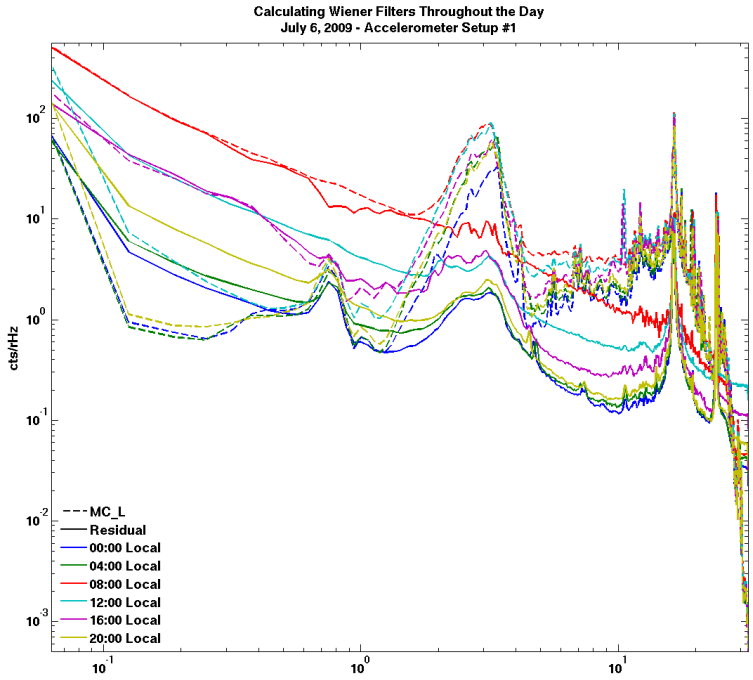

Figure 5 *– Wiener filtered MC_L signal over a full day. The quality of filters produced from data that are far noisier is clearly difficult to compare to the quieter signals at night. Thus, calculating new Wiener filters does not seem to be the best method to determine good instrument placement.*

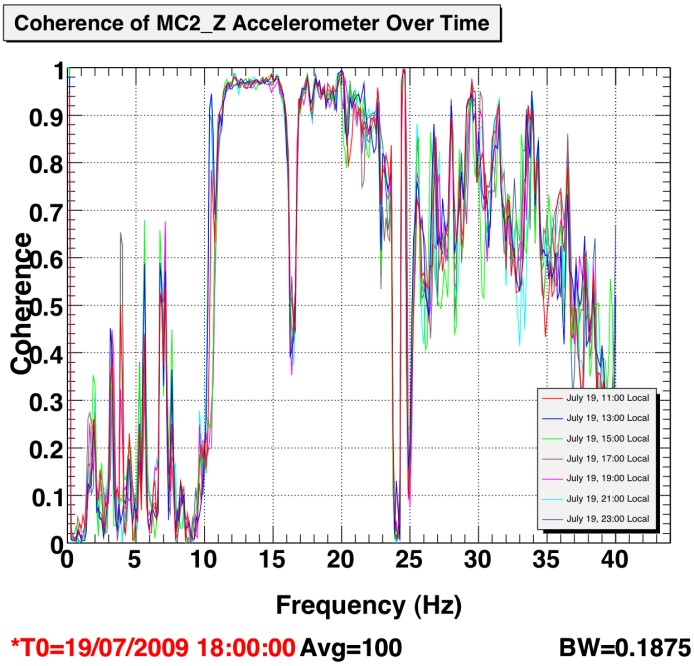

Figure 6 *– Coherence of the MC2_Z accelerometer over a 12-hour period—a full day's worth of coherence data could not be obtained, as the mode cleaner has been falling in and out of lock. The coherence is relatively unchanged throughout, despite there being far more ground motion during daylight hours.*

5 Accelerometers

The 40m lab is equipped with six low noise, low frequency Wilcoxon seismic accelerometers [9]. Originally, they had been mounted in two tri-axis sets, with one placed at the MC1 chamber and one at the MC2. More recently, they have been separated so that four of the accelerometers are independently mounted and two are mounted together at a right angle. Because the seismometers and microphones required other work to get up and running, the placement variations so far have been done only with the accelerometers.

5.1 Problems Encountered

After separating the seismometers so that four were independently mounted and two were mounted together at a right angle, I ran the cables so that all six could be arranged around the MC2 optical chamber. My plan was to look at the coherence between the accelerometers and the MC_L signal at several locations and find where to best detect different modes of motion. I was intending on doing the same thing at the MC1 chamber, but began to run into problems with the mode cleaner constantly falling out of lock. It seems as though the chamber is more sensitive to external forces than we thought, and the clamping of several accelerometers all around the chamber may have caused the stack to shift. So, the coherence optimization has been put on hold for the time being, and the accelerometers may have to be restricted to being mounted on the ground.

5.2 Further Work

Once the mode cleaner lock issues have been solved, work will continue with finding the best places for the accelerometers. We will also need to look at the coherence of rotational motions with the MC_L signal, which can be accomplished by subtracting the signals from two spatially separated instruments measuring motion in the same direction. The simple model of the isolation stacks that is being worked on will be helpful for determining which translational and rotational motions will be the most important to measure.

5.3 Note on the Setups

The accelerometer setups are numbered by the positions of each accelerometer. There was not a good way to come up with a simple and descriptive method of naming each configuration based on the actual location of the accelerometers, so they are instead numbered in the order that each configuration was tried. The first several sets of measurements were taken while the accelerometers were still mounted in three-axis sets, so these are named as ”Setup N-M,” where N is the ordered position of the MC1 set and M is of the MC2 set.

Once the accelerometers were separated, it no longer made sense to refer to them as ”MC1_X,”

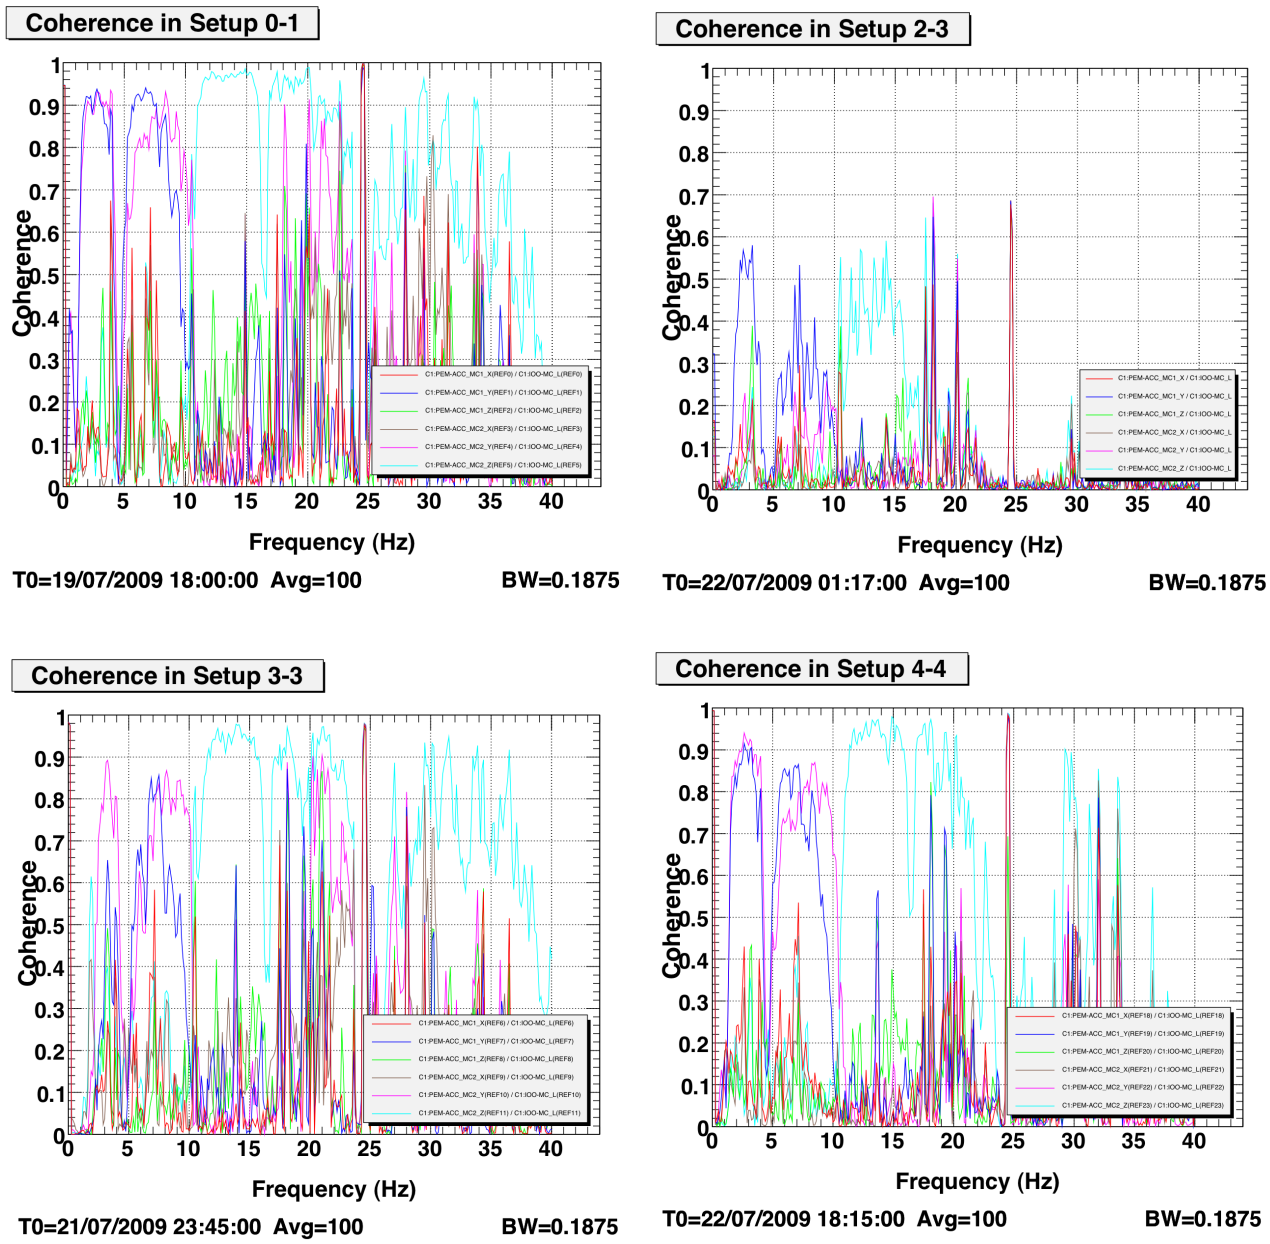

Figure 7 *– Coherences of all six accelerometer channels with the MC_L signal in four setup positions.*

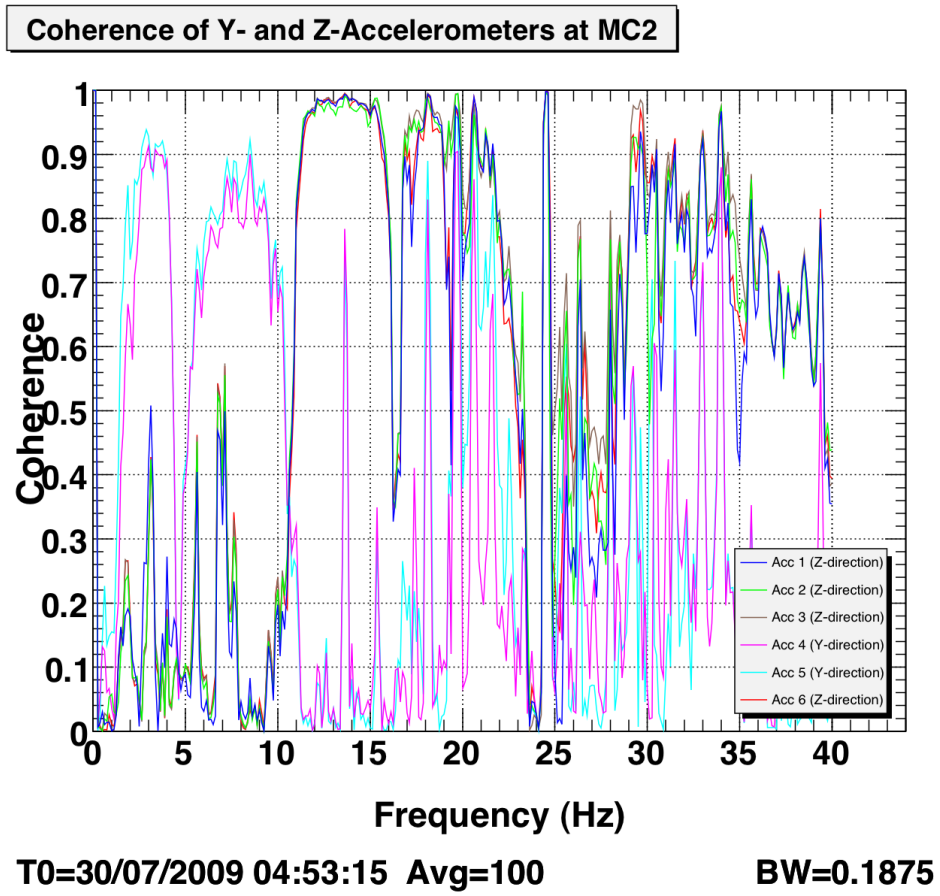

Figure 8 *– Coherence of accelerometers with the MC_L signal in several different positions near the MC2 optical chamber, pointing in the y- and z-directions.*

etc, so they are simply numbered 1-6. Their positions will be numbered with a directional indication, since they are no longer constrained to two sets of x-y-z. Because of the mode cleaner lock problems described in 5.1, there have not been enough measurements taken to necessitate implementation of a post-breakup numbering system.

6 Seismometers

6.1 Guralp Breakout Box

Currently, the 40m lab is in possession of two three-axis Guralp seismometers, and the breakout box is equipped to handle three. The breakout box was originally stuffed with normal thick-film resistors, which tend to be noisy at low frequencies . In order to avoid adding noise to the seismometer signal as much as possible, the resistors needed to be replaced with extra-low-noise Vishay Dale thin-metal-film resistors, which get much closer to the Johnson noise limit. The three channels for the first Guralp seismometer had already been redone, and I restuffed the second set of channels for the second seismometer. Several of the resistances and capacitances were adjusted as well, as shown in Figure 11. Also, capacitors were added across resistors R4 and R13 (drawn in blue in the diagram), which were carefully chosen to have the exact same capacitance, within the limits of our measuring devices.

6.2 AD620 Chip and Output Noise

Before installing both Guralp seismometers, we wanted to measure the noise in the breakout box channels, as well as in the AD620 chip (circuit element U1 in Figure 11). Previous measurements had seemed unreasonably high, and it is useful to know how noisy the channels are with and without the circuit changes. The chip noise in the restuffed channels seems to agree with measurements independently done on the AD620 by Ben Abbott. However, the noise in the original channels is looking problematic.

Initially, the noise measured in the AD620 chip and across the output of the Guralp breakout box channels were far higher than expected in the restuffed channels. Several attempts were made to better isolate the box, using highly insulating foam slabs and moving the measurement setup to a quieter area of the lab, without much improvement. I decided to go ahead and restuff the second set of channels, since the third set had the original configuration of resistors and capacitors, so they could be used for comparison noise measurements if necessary. While doing so, I noticed that the resistors in the first three channels did not all match the diagram: all of the resistors labeled as 100Ω in the circuit diagram had $100k\Omega$ resistors in the actual circuitboard. After verifying that the resistances noted on the diagram were indeed the correct ones, I replaced the appropriate resistors and proceeded to restuff the second channel set. Unsurprisingly, the changes brought both the chip and the output noises down into a credible range. The second vertical channel (VERT2) exhibited slightly higher noise levels than the other restuffed channels that were measured, but

after inspecting all the various circuit elements for damage and repairing a few questionable solder joints, the anomaly disappeared.

Once the circuit board was corrected, I re-measured the noise in channels from each of the three sets. The original circuit left out the R2 resistor, which was soldered into all of the third set of channels after taking a few measurements without it. While the AD620 noise in VERT1 and VERT2 matches the independent measurements pretty well, the noise across the chip in both the VERT3 and N/S 3 channels were about a factor of two lower before R2 was soldered in. Afterwards, all the noise measurements of the AD620 looked about the same (see Figure 9 [5]), so it seems that the noise difference was due to the presence or absence of the R2 resistor.

The Guralp R2 resistor is meant to set the gain on the AD620 chip (see Figure 10a), using the equation:

$$G = \frac{R1 + R2}{R_G} + 1. \tag{5}$$

The internal chip resistors, R1 and R2, are set to be exactly 24.7 kΩ each, so

$$G = \frac{49.4k\Omega}{R_G} + 1. \tag{6}$$

Thus, with no resistor in the Guralp R2 slot, $R_G = \infty$ and G = 1. Using the 5.49 kΩ resistor gives G = 10. From Figure 10b, which shows the noise across the inputs of the AD620, we might expect that the inclusion of the 5.49 kΩ resistor will bring the noise down by about an order of magnitude, but that does not seem to have happened.

6.3 Further Work

Now that the Guralp breakout box is complete and has had careful noise measurements taken, we can test the noise in the seismometers themselves. This will be done by placing both Guralp seismometers and the Ranger together in a quiet area of the interferometer lab, isolating them from any non-seismic motion sources with highly insulative foam. Theoretically, the seismometers ought to measure the same ground motion, since their spatial separation will be small, so we should be able to learn something about the noise in the instruments themselves by comparing the signals.

Additionally, the seismometers will need to be optimized in terms of position in the same way that the accelerometers are being done.

[5] VERT3 is shown *with* R2 and N/S 3 *without*. The measurements of both channels were in agreement for both settings, so only one of each was plotted.

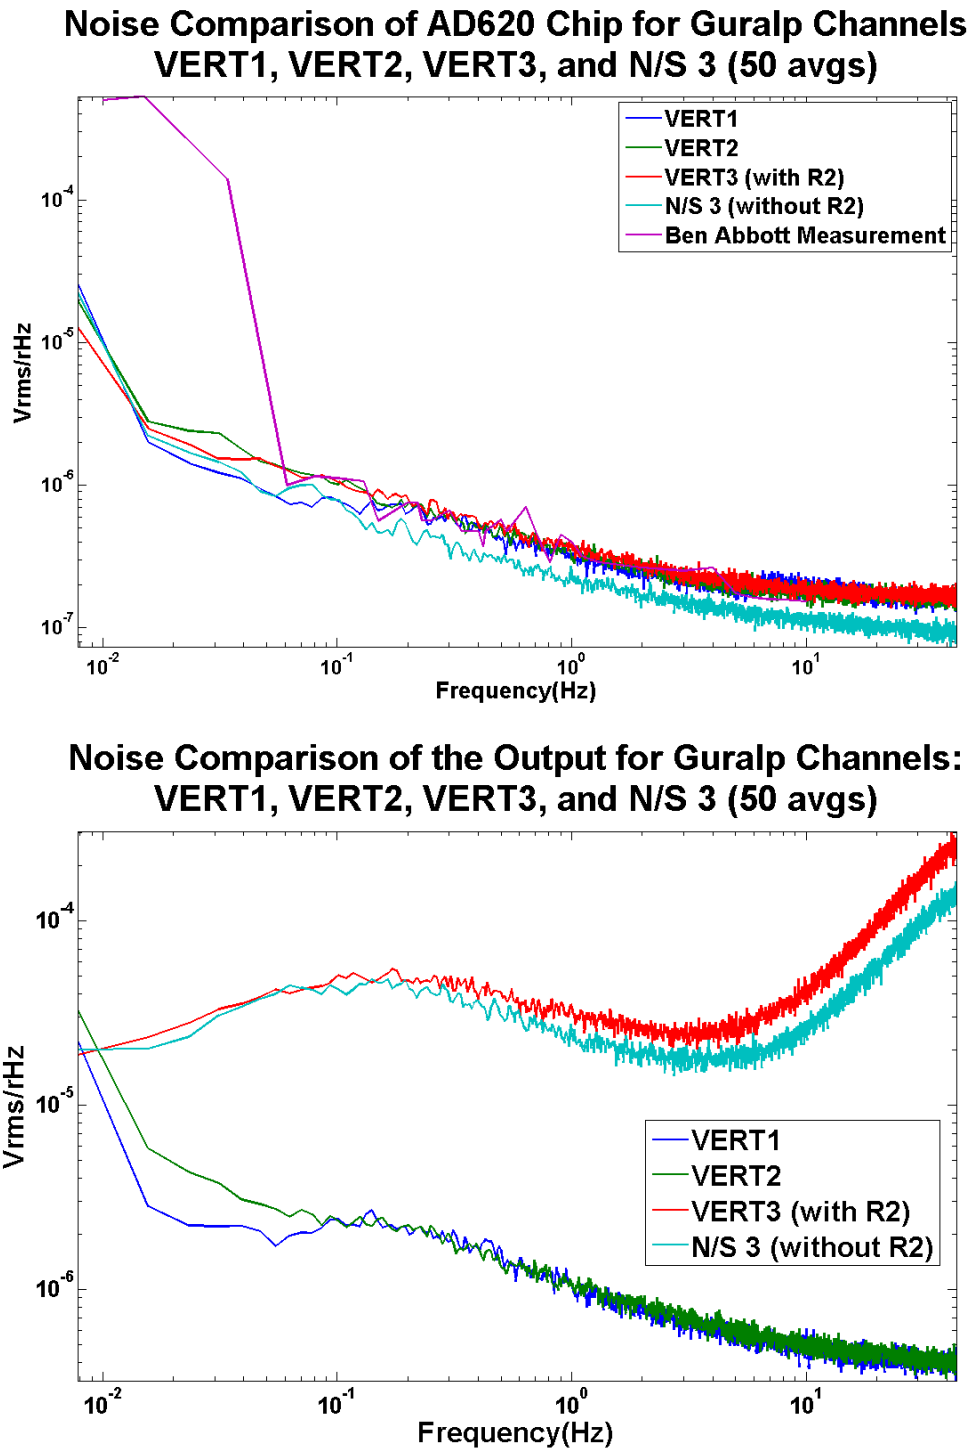

Figure 9 *– Noise across AD620 and output for the three vertical and the third north/south channel in the Guralp breakout box. VERT1 and VERT2 have been restuffed with the low-noise thin-film resistors, and VERT3 and N/S 3 are unchanged from the original circuitry, except that VERT3 has had a 5.49 kΩ resistor placed in the R2 slot. The AD620 chip noise is also compared with measurements taken by Ben Abbott at LIGO Wilson House [3].*

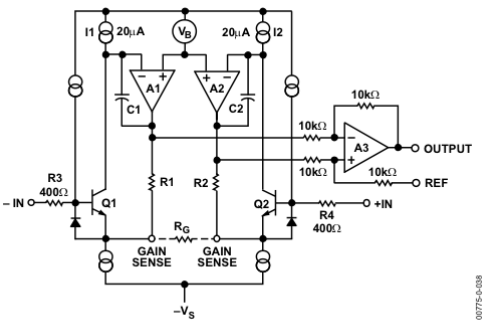

Figure 38. Simplified Schematic of AD620

a)

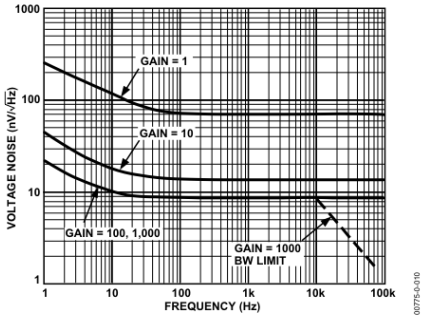

Figure 10. Voltage Noise Spectral Density vs. Frequency (G = 1–1000)

b)

Figure 10 *–*
a) Simplified diagram of the AD620 chip. It is designed so that the gain may be programmed by the single external resistor, R_G. With no resistor in place, G = 1; with a 5.49 kΩ resistor, G = 10.
b) Manufacturer's noise specifications for various gain settings. [1]

7 Microphones

While the 40m laboratory does currently have accelerometers and seismometers set up to measure seismic noise sources, it has only a single, low-quality microphone that is too noisy to be useful. Two condenser microphones were purchased to provide better measurements of acoustic noise sources. Condenser microphones are popular in laboratory applications due to the comparatively small mass that must be moved by the incoming sound waves, requiring the sound waves to do less work. The particular model purchased covers the 20Hz - 20kHz frequency range.

Due to the size of the lab, the required cables were longer than could be found ready-made. So, two 80-ft XLR cables were made for that purpose. Also, the Analog-to-Digital Converter (ADC) can only handle about ±2V. Crude sound-testing of the microphone - preamplifier setup showed that, even on the lowest gain setting, the voltage output could reach ±6V. However, due to the length of the cables, it would be advantageous to be able to operate the preamplifier with a high gain setting, with which we were able to create voltage outputs of about ±70V. Voltage dividers were designed so that, even with an input of ±100V, the ADC would not be overloaded (see Figure 12).

7.1 Problems Encountered

After installing the microphones and testing to make sure that the data acquisition channels were working properly, I began looking for the best places to install the microphones. I started with one in the prestabilized laser (PSL) chamber and one hanging over the asymmetrical (AS) table, where the pre-mode cleaner (PMC) optics are located. Calculating the Wiener filter using only the microphone signals had essentially no effect on the MC_L signal. No matter how the microphones were adjusted, the coherence with the MC_L signal always seemed to drop off sharply around 1 kHz, so the microphones' response to a generated sine signal sweeping over a frequency range of several kHz was investigated (see Figure 13).

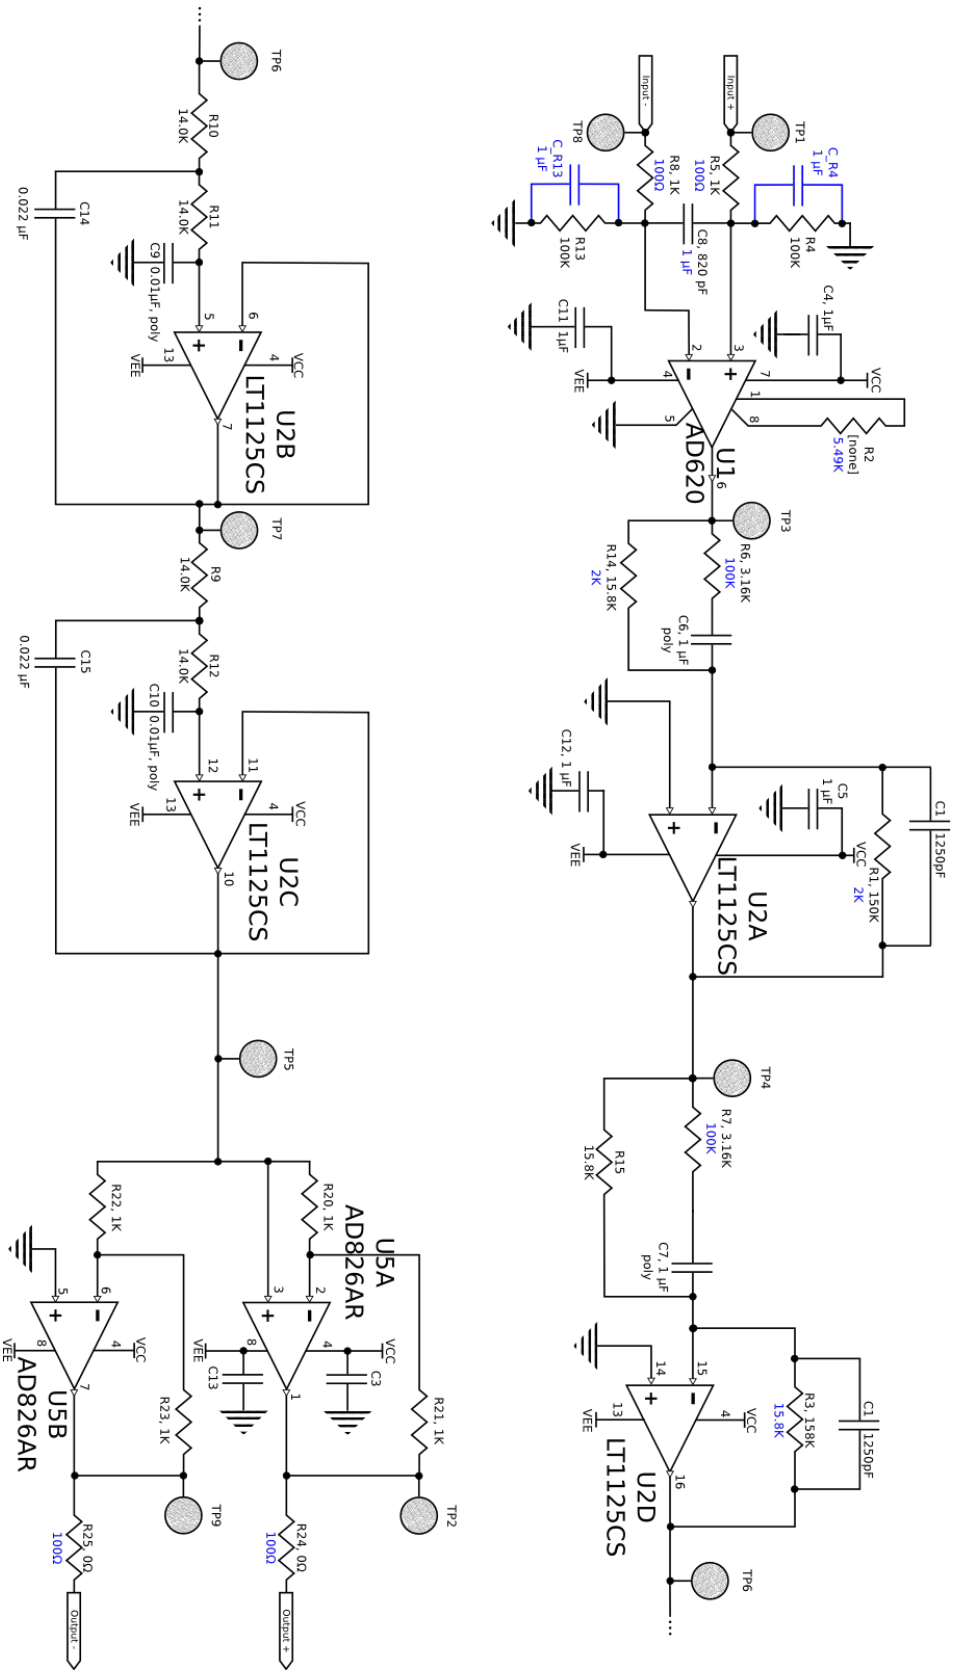

Figure 11 *– Circuit diagram of the Guralp box channels. New resistor and capacitor values are shown in blue, the originals in black. Test point 6 is shown twice to indicate how the two parts fit together.*

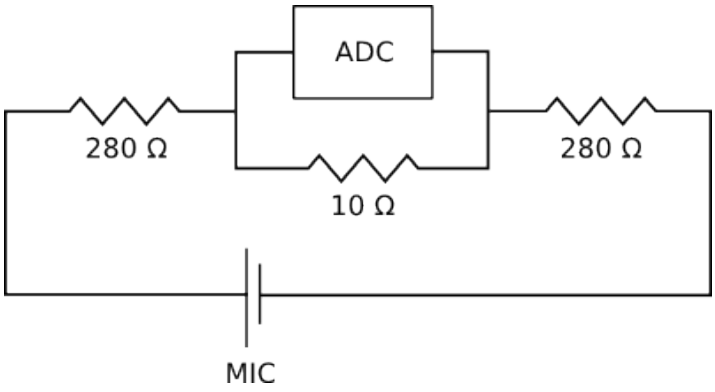

Figure 12 *– Voltage divider design.*

The frequency response measurements yielded some strange results, including very slow beating at some drive frequencies on the order of 10 - 100 mHz. Attempts to better isolate the test setup and produce more reliable input signals did not resolve the beating, but it did not appear for the same input frequencies every time. I did notice clear spatial interference patterns while taking measurements and found that I was able to consistently suppress any periodic voltage modulations by using physical barriers to better enclose the setup. The fact that these adjustments served to stabilize the signal, rather than simply amplify it, producing a fairly consistent voltage response throughout the frequency range, indicates that wave reflection and interference were causing the anomalous measurements.

To test this theory, four power spectra were taken of the response to the swept sine input, each with a different microphone orientation. The purpose was to attempt to average out the effects of sound wave reflection from surfaces in the laboratory. Figure 14 shows the results of the testing, along with the manufacturer's specifications for the microphones' frequency response. The power spectra do not follow the manufacturer's values exactly, but the general shape is in agreement and there are no frequency bands for which the microphone response was consistently low. So, it seems safe to eliminate the microphones and pre-amplifiers as the source of the high-frequency signal suppression.

The ADC that is being used by the microphones has anti-aliasing filters which may be the cause of the sharp frequency drop-off, as it certainly does not seem to be a problem between the microphone and the ADC [6]. As there are not any free channels without the anti-aliasing filters that can be used to test this hypothesis, the filters will simply need to be removed from the ADC to see what difference is made in the signal that is transmitted to the computers.

Appendices

[6] All of the described measurements were taken with and without the voltage divider in place, without producing any significant differences, aside from the obvious magnitude difference, so it seems clear that the voltage divider is not to blame.

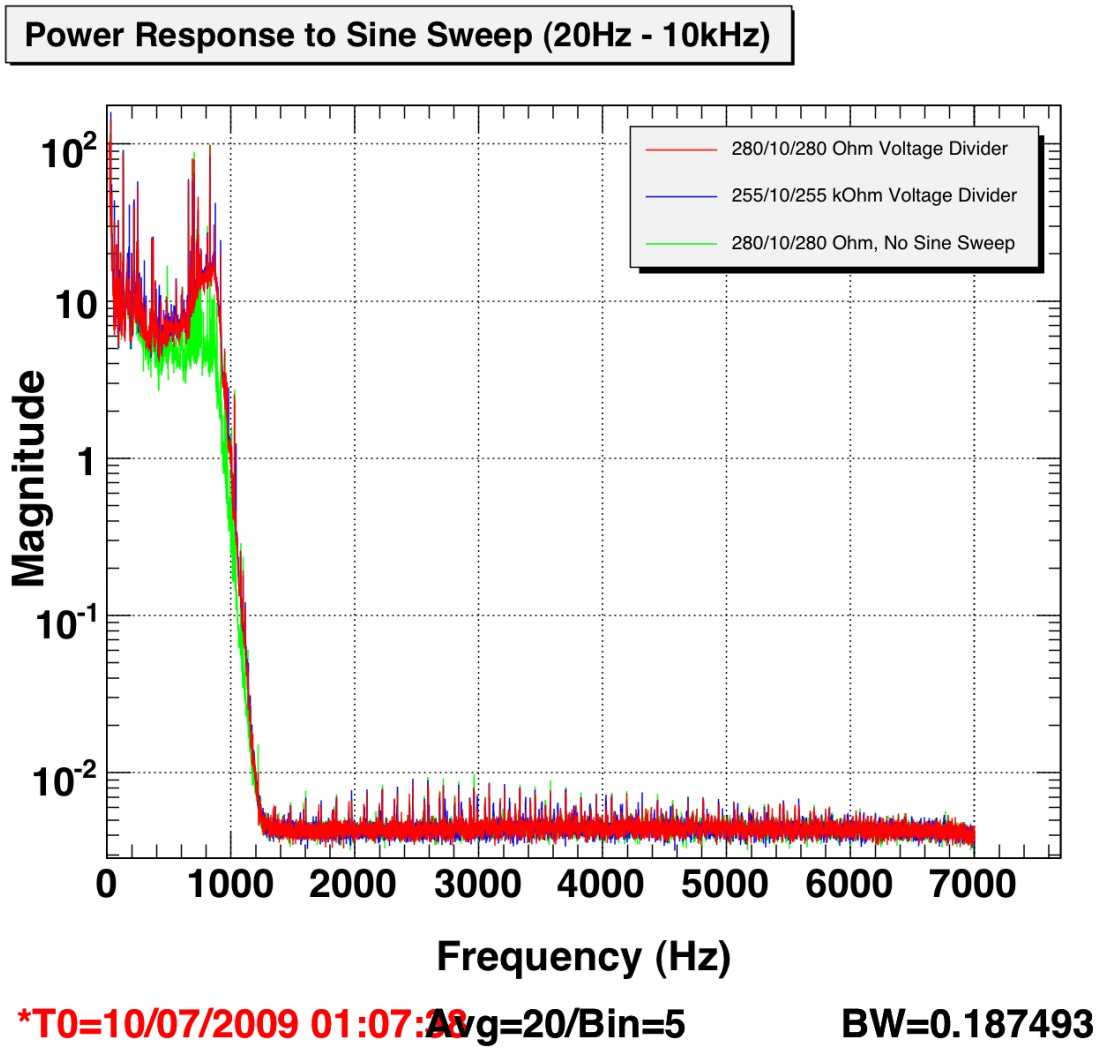

Figure 13 *– Power spectra of microphones with and without a swept sine input and with high and low resistance voltage dividers. It is apparent that something is preventing data transmission from the microphone to the data acquisition system above ~1 kHz.*

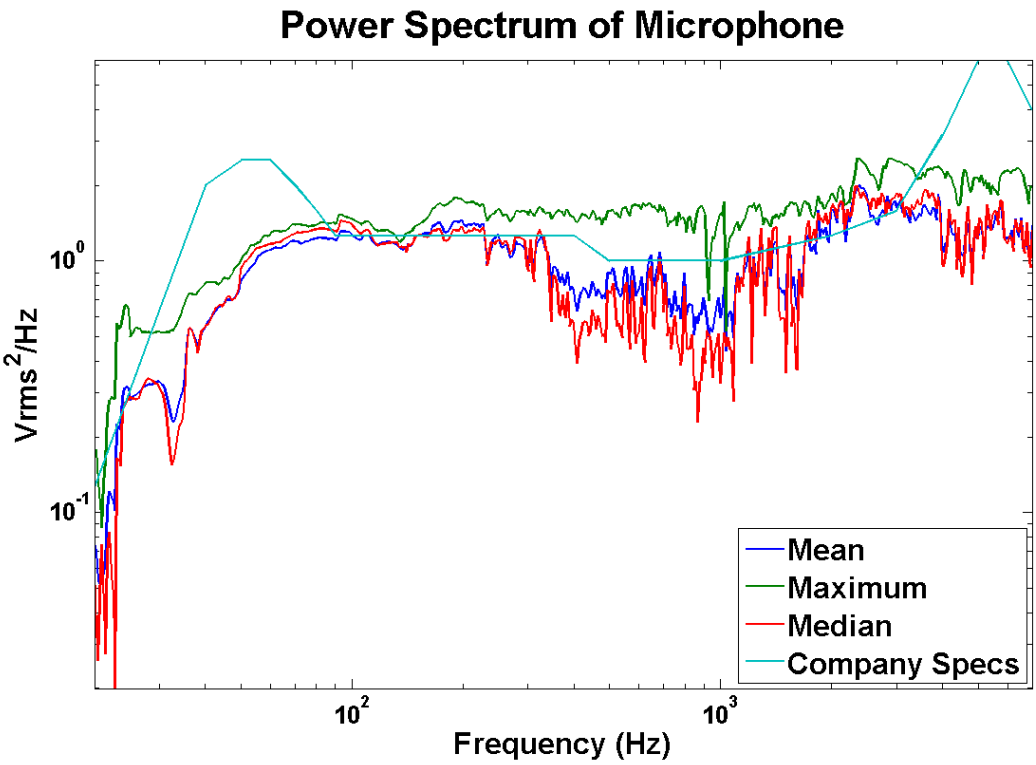

Figure 14 *– Curves taken from four different sine-sweep power spectra for the Bluebird microphones. The microphone was placed in a different position each time in order to try to account for interference patterns created by testing in a less-than-ideal acoustic chamber. The manufacturer specifications are also plotted for comparison purposes [2].*

A More On the Mode Cleaner

A.1 Hermite-Gaussian Modes [6]

There are many transverse amplitude distributions that maintain their forms functionally as they travel. These are known as the Hermite-Gaussian modes, and, for the mode cleaner, these are the eigenmodes of the laser cavity. The modes are composed of Gaussians multiplied by Hermite polynomials in the two directions transverse to propagation, defined as

$$\psi_m(\xi) \equiv H_m(\xi) e^{-\xi^2/2} \tag{7}$$

Recall that the Hermite polynomials represent the quantum mechanical wave function modes for a simple harmonic oscillator potential. Indexing the two polynomials (one for each transverse dimension) as m and n, we obtain the general Hermite-Gaussian mode

$$u_{mn}(x,y,z) = \frac{C_m n}{\sqrt{1+\frac{x^2\lambda^2}{\pi^2\omega_0^4}}} e^{i(m+n+1)\phi} \xi_m\left(\frac{\sqrt{2}y}{\omega}\right) \xi_n\left(\frac{\sqrt{2}z}{\omega}\right) e^{-ik(y^2+z^2)/2R} \tag{8}$$

Notice that for $m = n = 0$, or the 00 mode, the formula reduces to a pure Gaussian. This is the desired fundamental mode. Also, because the mn mode has a phase factor of $e^{i(m+n+1)\phi}$, modes will have different phases through propagation. So, different modes will resonate for various wavelengths of light, or for different cavity lengths at a fixed wavelength. For a cavity of length L, the resonant wavelengths are

$$\lambda_{mnp} = 2L\left[\frac{p+2(m+n+1)}{\pi} tan^{-1}\left(\frac{L\lambda}{2\pi\omega_0^2}\right)\right]^{-1}, \tag{9}$$

where the p index indicates the number of wavelengths along the length of the resonating cavity.

By taking advantage of these phase differences, it is possible to set up a Fabry-Perot cavity that will resonate for only the desired mode(s) and suppress all others, thus "cleaning" the mode of laser beam. For the purposes of the LIGO interferometer, the mode cleaner is set up to isolate the purely Gaussian mode (00).

A.2 Pound-Drever-Hall Locking [4]

The success of the interferometric gravitational wave detector is dependent on its ability to provide a very stable laser frequency with which to measure length changes in the interferometer. One can

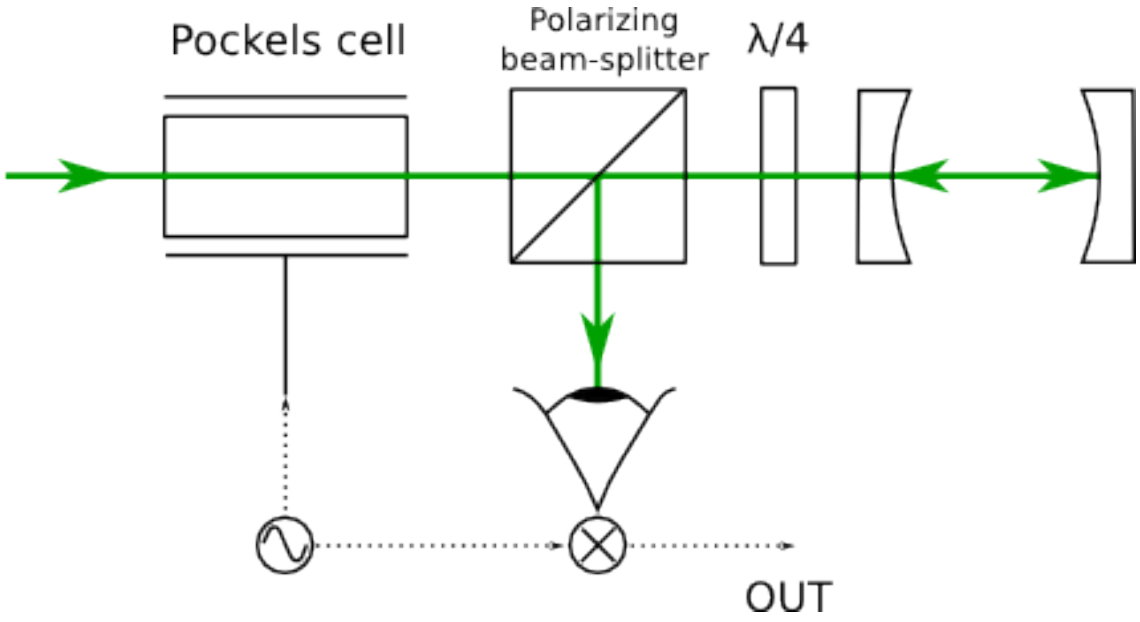

Figure 15 *– Schematic of the components used for the Power-Drever-Hall reflection cavity locking scheme, adapted from Figure 12.6 in Saulson [6].*

very accurately measure the frequency of laser light by passing it through a Fabry-Perot cavity, which only transmits wavelengths that are some integer division of twice the physical length of the cavity. The transmission function is shown in Figure 16a.

If one chose to operate just next to the resonant frequency, where the transmission function is steep and approximately linear, then small frequency changes would beget proportional intensity changes. So, theoretically, maintaining a constant intensity would result in a stable laser frequency. One of the flaws of this method, which was used prior to the development of Pound-Drever-Hall, is that it cannot distinguish between frequency and intensity fluctuations in the laser. A better method would be to hold the *reflected* intensity at zero (which would occur for resonant wavelengths), but, as the reflected intensity curve is symmetric about the minimum (see Figure 16b), one cannot tell in which direction the frequency needs to be adjusted. Fortunately, the derivative of the reflected light curve is anti-symmetric and its sign is easily determined by varying the laser frequency slightly and observing the response.

From here, the information needs to be fed back into the laser to hold resonance. The Pound-Drever-Hall method achieves this by sending the Pockels-cell-modulated laser beam through an optical isolator [7] to a photodetector, where it can be compared to the oscillatior that drives the Pockels cell. The two signals are mixed, essentially producing an output that is a product of the inputs. This product will consist of a dc (low-frequency) signal and a signal at $2f_{mod}$, where f_{mod} is the modulation frequency. A low-pass filter removes the higher frequency to isolate the dc signal, which is then sent to the servos to lock the laser.

[7]The isolator consists of a polarized beam-splitter and a quarter-wave plate (see Figure 15). The light entering through the Fabry-Perot cavity is of a single linear polarization which is converted to circular polarization by the quarter-wave plate. On the way back out, the quarter-wave plate converts the light to the opposite linear polarization so that it is entirely reflected into a photodetector rather than transmitted and sent back toward the laser [6].

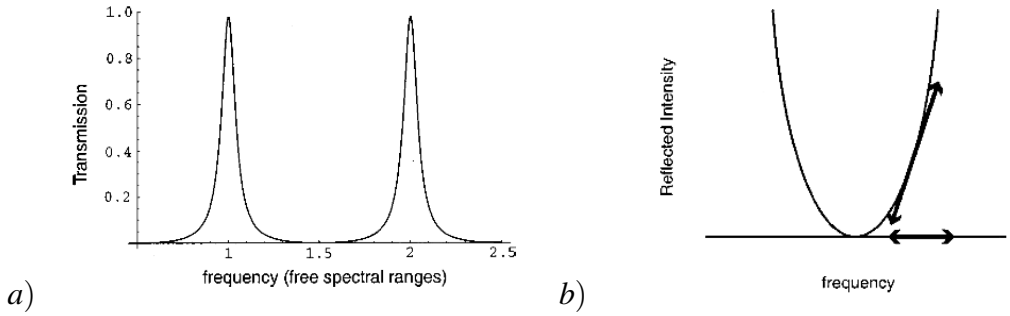

a) *b)*

Figure 16 *– Laser light a) transmission through and b) reflection from a Fabry-Perot cavity as a function of the frequency [4].*

B Wiener Filtering: A Brief Motivation [8] [7]

To begin, consider a filter input sampled on a discrete time series $u(0), u(1), u(2), \ldots$, and a corresponding filter inpulse response $\omega_0, \omega_1, \omega_2, \ldots$ (assumed to be complex-valued and infinite). The filter output, $y(n)$, at discrete time n, is an estimation of the desired output $d(n)$ [9]. It is given by the linear convolution sum

$$y(n) = \sum_{k=0}^{\infty} \omega_k^* u(n-k), \quad n = 0, 1, 2, \ldots, \tag{10a}$$

with the error being defined as

$$e(n) = d(n) - y(n). \tag{10b}$$

We choose to optimize the filter by minimizing the mean-square of $e(n)$, so

$$J = E[e(n)e^*(n)] = E[|e(n)|^2], \tag{10c}$$

where E is the expectation operator.

Without going into the full details, by assuming (reasonably) that the filter coefficients, ω_k, are complex and requiring that all terms of the gradient of the cost function, ∇J, be zero, one arrives at the following optimization condition:

$$E[u(n-k)e_0^*(n)] = 0, \quad k = 0, 1, 2, \ldots. \tag{11}$$

[8] For a full derivation, please see Chapter 2 of Haykin [7]

[9] Note that the filter input and desired response are both assumed to have a zero mean. In the case of a nonzero mean, it is simply subtracted from the proper quantity before filtering.

The minimum of the cost function, naturally, occurs where $\nabla J_k = 0$ for all k. The condition in Eq.(11) implies that, in order for J to be at its minimum value, the corresponding estimation error, $e_0(n)$ must be orthogonal to the input for all times n.

We may substitute Eqs.(10a) and (10b) into Eq.(11) to reconfigure the necessary and sufficient condition of optimization, yielding

$$\sum_{i=0}^{\infty} \omega_{oi} E[u(n-k)u^*(n-i)] = E[u(n-k)d^*(n)], \quad k = 0,1,2,\ldots, \tag{12a}$$

where ω_{oi} is the ith optimum filter coefficient. The two expectations in Eq. (12a) are redefined as

$$r(i-k) = E[u(n-k)u^*(n-i)], \tag{12b}$$

$$p(-k) = E[u(n-k)d^*(n)], \tag{12c}$$

where $r(i-k)$ represents the autocorrelation function of the filter output and the cross-correlation of the filter input, $u(n-k)$, and the desired response, $d(n)$, is given by $p(-k)$. Substituting, this brings us to

$$\sum_{i=0}^{\infty} \omega_{oi} r(i-k) = p(-k), \quad k = 0,1,2,\ldots, \tag{13}$$

an infinite set known as the Wiener-Hopf equations—the fundamental equation set for Wiener filters.

Up until this point, it was assumed that the impulse response was infinite (IIR). It turns out that, aside from being more practical, the special case of FIR filters simplifies the solution to the Wiener-Hopf equations tremendously. For a filter of length M, the impulse response of the filter is given by the set of tap weights, $\omega_0, \omega_1, \ldots, \omega_{M-1}$ (the coefficients of the filter at each sample, or tap point). Thus, we may reduce the Wiener-Hopf equations to a finite set of M equations

$$\sum_{i=0}^{M-1} \omega_{oi} r(i-k) = p(-k), \quad k = 0,1,2,\ldots,M-1, \tag{14}$$

where the o subscript indicates the optimum value.

The Wiener-Hopf equations may be described as matrices. Consider $\mathbf{R}$ to be an $M \times M$ matrix

representing the correlation of the tap inputs of our FIR filter:

$$\mathbf{R} = E[\mathbf{u}(n)\mathbf{u}^H(n)], \tag{15a}$$

where

$$\mathbf{u}(n) = [u(n), u(n-1), \ldots, u(n-M+1)]^T. \tag{15b}$$

Then, $\mathbf{R}$ is expanded into

$$\mathbf{R} = \begin{bmatrix} r(0) & r(1) & \cdots & r(M-1) \\ r^*(1) & r(0) & \cdots & r(M-2) \\ \vdots & \vdots & \ddots & \vdots \\ r^*(M-1) & r^*(M-2) & \cdots & r(0) \end{bmatrix}. \tag{15c}$$

We may also define a cross-correlation vector between the tap inputs and $d(n)$ as follows:

$$\mathbf{p} = E[\mathbf{u}(n)d^*(n)], \tag{16a}$$

which expands to

$$\mathbf{p} = [p(0), p(-1), \ldots, p(1-M)]^T. \tag{16b}$$

This brings us to the matrix formulation of the Wiener-Hopf equations

$$\mathbf{R}\mathbf{w}_o = \mathbf{p} \tag{17}$$

where $\mathbf{w}_o$ is the $M \times 1$ tap-weight vector:

$$\mathbf{w_o} = [\omega_{o0}, \omega_{o1}, \ldots, \omega_{o,M-1}]^T. \tag{18}$$

Rearranging, and assuming the matrix $\mathbf{R}$ to be nonsingular, we find the formula for the optimum filter tap weights to be

$$\mathbf{w}_o = \mathbf{R}^{-1}\mathbf{p}.$$

C Coherence [10]

The coherence of two signals is an indication of the strength of their relationship to each other, similar to the correlation coefficient, r^2, in a linear regression model. We give here a brief derivation and discussion of the coherence.

C.1 Autocorrelation and Cross-correlation

For a time series, x_t, the recall that the *variance* of the signal is estimated by

$$\sigma^2 = \frac{1}{N-1}\sum_{n=1}^{N}(x_n - \mu)^2, \tag{20a}$$

where μ is the mean value of the time series. Similarly, we can calculate the *covariance* of two signals x and y as

$$\sigma_{xy} = \frac{1}{N-1}\sum_{n=1}^{N}(x_i - \mu_x)(y_i - \mu_y). \tag{20b}$$

Notice that for $x = y$, the above reduces to the variance for a single variable—that is, $\sigma_{xx} = \sigma_x^2$—and that $\sigma xy = \sigma yx$.

Now, suppose we have some time-dependent discrete signal x_t and we delay it by some T number of samples. Then, we can calculate the covariance between x_t and x_{t-T}, known as the *autocovariance* function:

$$\sigma_{xx}(T) = \frac{1}{N-1}\sum_{t=1}^{N}(x_{t-T} - \mu_x)(x_t - \mu_x). \tag{21a}$$

Normalizing, we get the *autocorrelation* function

$$r_{xx}(T) = \frac{\sigma_{xx}(T)}{\sigma_{xx}(0)}, \tag{21b}$$

where $\sigma_{xx}(0) = \sigma_x^2$, as pointed out above. We see that $r_{xx}(0) = 1$, which is natural, as the unshifted signal is perfectly correlated with itself. We can also calculate the *cross-correlation* between two

signals x_t and y_t by normalizing the *cross-coveriance*,

$$\sigma_{xy}(T) = \frac{1}{N-1}\sum_{t=1}^{N}(x_{t-T} - \mu_x)(y_t = \mu_y), \tag{22a}$$

which reduces to the autocovariance function for $x = y$. The cross-correlation is then

$$r_{xy} = \frac{\sigma_{xy}(T)}{\sqrt{\sigma_{xx}(0)\sigma_{yy}(0)}}, \tag{22b}$$

where, again, $\sigma_{xx}(0) = \sigma_x^2$ and $\sigma_{yy} = \sigma_y^2$ are the variances in x_t and y_t, respectively. Evaluated at $T = 0$,

$$r_{xy}(0) = \frac{\sigma_{xy}}{\sigma_x\sigma_y}, \tag{22c}$$

the correlation between x_t and y_t.

C.2 Spectral Densities

C.2.1 Power Spectral Density

For a discretely sampled signal with N sampling points, the power is given by

$$P_x = \sum_{n=1}^{N}|x(n)|^2. \tag{23}$$

We would like to find an expression for P_x in frequency k-space, so let us re-express $|x(n)|$ as a Fourier series in its complex conjugated form

$$x^*(n) = \frac{1}{N}\sum_{k=1}^{N}X^*(k)e^{-i2\pi(k-1)n/N}, \tag{24}$$

where X^* is the complex conjugate of the Fourier transform of x. So,

$$P_x = \sum_{n=1}^{N}x^*(n)x(n) = \sum_{n=1}^{N}x(n)\sum_{k=1}^{N}X^*(k)e^{-i2\pi(k-1)n/N}, \tag{25a}$$

which can be rewritten by changing the order of summation as

$$P_x = \sum_{k=1}^{N} X^*(k) \sum_{n=1}^{N} x(n) e^{-i2]pi(k-1)n/N}. \tag{25b}$$

The inner sum we recognize as the Fourier transform of $x(n)$, $X(k)$. Thus,

$$P_x = \frac{1}{N} \sum_{k=1}^{N} X^*(k) X(k) = \frac{1}{N} \sum_{k=1}^{N} |X(k)|^2, \tag{26}$$

giving a relationship between the signal power in n-space and k-space. The summation term,

$$P_x(k) = |X(k)|^2 \tag{27}$$

is called the *power spectral density*, or PSD.

C.2.2 Cross Spectral Density

The Cross Spectral Density (CSD) is defined as the Fourier transform of the cross-covariance:

$$P_{xy} = \sum_{n=-\infty}^{\infty} \sigma_{xy}(n) e^{-i\omega n}. \tag{28}$$

Generalizing the cross-covariance as

$$\sigma_{xy} = \sum_{l=-\infty}^{\infty} x(l) y(l-n) \tag{29}$$

and substituting, we arrive at

$$P_{xy}\omega = \sum_{n=-\infty}^{\infty} \sum_{l=-\infty}^{\infty} x(l) y(l-n) e^{-i\omega n}. \tag{30}$$

Let $k = l - n$, so that $e^{-i\omega n} = e^{i\omega l} e^{i\omega k}$. Then the CSD may be re-expressed as

$$P_{xy} = X(\omega) Y(-\omega) \tag{31}$$

where X and Y are given by

$$X(\omega) = \sum_{l=-\infty}^{\infty} x(l)e^{-i\omega l} \qquad Y(-\omega) = \sum_{k=-\infty}^{\infty} y(k)e^{+i\omega k} \tag{32}$$

If the signals are real, then $Y^*(\omega) = Y(-\omega)$, where Y^* is the complex conjugate of Y. So, the cross-spectral density can be written as

$$P_{xy}(\omega) = X(\omega)Y*(\omega). \tag{33}$$

So, by Eqs. 30 and 33, we see that the CSD can be calculated either by Fourier transforming the cross-covariance between two time series x and y, or by Fourier transforming the signals individually and multiplying one by the complex conjugate of the other.

C.3 The Coherence and Its Interpretation

For a several-variable time-series of m variables, the frequency domain relationships between series can be expressed as the Hermitian *power spectral density matrix* [10]:

$$\mathbf{P}(f) = \begin{pmatrix} P_{11}(f) & P_{12}(f) & \cdots & P_{1m}(f) \\ P_{12}(f) & P_{22}(f) & \cdots & P_{2m}(f) \\ \vdots & \vdots & \ddots & \vdots \\ P_{1m}(f) & P_{2m}(f) & \cdots & P_{mm}(f) \end{pmatrix}. \tag{34}$$

The diagonal elements are the PSDs for each time-series variable x_i, and the non-diagonal elements are the CSDs between variables x_i and x_k, $i \neq k$.

For any pair of variables whose relationships are described by $\mathbf{P}(f)$, the *complex coherence function* is

$$r_{ij}(f) = \frac{P_{ij}}{\sqrt{P_{ii}(f)P_{jj}(f)}}, \tag{35}$$

where P_{ij} is the function in the ith row and jth column of $\mathbf{P}(f)$. The *mean squared coherence* (MSC), or simply the *coherence*, between two members of a multivariate time-series is simply the squared magnitude $|r_{ij}(f)|^2$.

[10] Recall that the cross-covariance, $\sigma_{xy} = \sigma_{yx}$, is not order-specific. It follows then that $P_{xy} = P_{yx}$, so the power spectral density matrix is symmetric.

The coherence function is a measurement of the linear correlation between two variables at each sampled frequency and thus is closely tied to *linear filters*—of which Wiener filters are a subset—where one signal is interpreted as a weighted (noisy) version of the other. Linear filters, then, can be viewed as a set of linear regressions for each frequency, in which the best regression coefficient or filter is

$$H(f) = \frac{P_{xy}(f)}{P_{xx}(f)},$$

which analogous to the regression coefficient $a = \left(\frac{\sigma_{xy}}{\sigma_{xx}}\right)$. Taking Eqs. 35 and **??**, we see that the coherence relates to the best filter as

$$MSC = r_{xy}^2 = |H(f)|^2 \frac{P_{xx}(f)}{P)yy(f)}, \tag{37}$$

which, again, is analogous to $r^2 = a^2 \left(\frac{\sigma_{xx}}{\sigma_{yy}}\right)$. Because the coherence is so closely related to the mechanism of linear filters, it is perfect for assessing the quality of Wiener filtration that is likely to come from any instrument position.

C.3.1 Phase

The complex coherence function can also be used to calculate the *phase spectrum* between two signals:

$$\theta_{ij}(f) = tan^{-1}\left[\frac{Im(r_{ij}(f))}{Re(r_{ij}(f))}\right]. \tag{38}$$

For any given frequency, if one signal remains in a fixed phase relationship with the other, the coherence will be high, even if the signals are completely out-of-phase. It follows that, if the phase of one channel is changing with respect to the other, the coherence will not be high at that particular frequency. The *coherence time* refers to the time interval over which the phase relationship remains constant. In practical terms, this means that, for two instruments measuring a propagating signal at some finite displacement from each other, the coherence function will reflect signal cohesion, even though the signal may reach the instruments at slightly different times.

References

[1] Analog Devices, Inc., Norwood, MA. *AD620: Low Cost, Low Power Instrumentation Amplifier*, Revision G edition, 2004.

[2] Baltic Latvian Universal Electronics (BLUE). *Bluebird Operating Manual*, 2008. URL {http://www.bluemic.com/bluebird/}.

[3] Ben Abbott. Output Noise of Opamps with Inputs Grounded. Independent noise testing of the AD620 and other opamps.

[4] Eric D. Black. An introduction to Pound-Drever-Hall laser frequency stabilization. *American Journal of Physics*, 69(1), Jan 2001.

[5] Joseph Giaime and Partha Saha and David Shoemaker and Lisa Sievers. A passive vibration isolation stack for LIGO: Design, modeling, and testing. *Review of Scientific Instruments*, 67 (1), Jan 1996.

[6] Peter R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*, pages 90–100, 127–140, 191–214. World Scientific Publishing, River Edge, New Jersey, 1994.

[7] Simon Haykin. *Adaptive Filter Theory*, pages 1–14, 35–50, 69–75, 94–126. Prentice Hall, New Jersey, Fourth edition, 2002.

[8] The LIGO Scientific Collaboration. LIGO: The Laser Interferometer Gravitational-Wave Observatory. *Reports on Progress in Physics*, 2008.

[9] Wilcoxon Research Inc, Gaithersburg, MD. *Model 731 A: Ultra-quiet, ultra low frequency, seismic accelerometer*, Revision C.2 edition, May 2004. URL {http://www.wilcoxon.com/}.

[10] William D. Penny. Signal Processing Course, Apr 2000. URL {http://www.fil.ion.ucl.ac.uk/~wpenny/course/course.html}.