Towards Detecting Gravitational Waves from the Crab Pulsar

by

Sharon Rapoport

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Bachelor of Science in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2009

Author . . . .
Department of Physics
May 8, 2009

Certified by . . . .
Nergis Mavalvala
Associate Professor
Thesis Supervisor

Accepted by . . . .
Professor David E. Pritchard
Senior Thesis Coordinator, Department of Physics
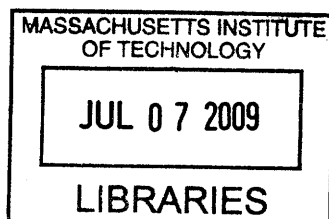
Towards Detecting Gravitational Waves from the Crab Pulsar

by

Sharon Rapoport

Submitted to the Department of Physics
on May 8, 2009, in partial fulfillment of the
requirements for the degree of
Bachelor of Science in Physics

Abstract

A variety of fundamental and technical noise sources impact the ability of the Laser Interferometer Gravitational-Wave Observatory (LIGO) to directly detect gravitational radiation. Noteworthy examples include Newtonian gravity gradient, seismic, acoustic, thermal and photon shot noise. These are the obstacles that must be confronted by the planned upgrade to the LIGO detectors, Advanced LIGO. To achieve improved sensitivity, significant improvements of LIGO's hardware must be paralleled by equivalent advances in the digital realm. Using adaptive filtering techniques, it is possible to cancel noise from known sources. We present results showing successful suppression of power line noise at 60Hz by a factor of 500 using commonly available sensors and standard FIR filters. Attenuation of 60Hz noise is particularly important to LIGO due to the 59.7 Hz radiation frequency of the Crab pulsar. Application of the techniques developed in this work to the LIGO detectors promises to improve the signal to noise ratio at the Crab frequency and thus pave the way toward direct detection of gravitational radiation from a known source.

Thesis Supervisor: Nergis Mavalvala
Title: Associate Professor

Acknowledgments

My time as a UROP student has been such a remarkable experience thanks to the support, kindness and guidance of everyone at LIGO-MIT. I would like to express my deepest gratitude to Rich Mittleman, my supervisor during this period. Since Freshman year he always exhibited great knowledge and patience in answering my questions (physics related or otherwise), spending long hours at his board until he was sure I understood. Thank you. Special thanks to Nergis Mavalvala who shaped this thesis and encouraged me to fulfill my potential. I'd also like to highlight the contribution of Matt Evans, whose comprehensive knowledge of adaptive filtering inspired this project; Brett and Fabrice who provided me with the control theory background necessary for this experiment and John, who was there to preserve my mental health whilst I was losing the battles against Matlab, LaTeXand the English language. Last, but by no means least, I would like to thank my family who supported me through the process of moving to a new continent against their parental instinct to keep offspring close to the nest. They crossed the world to help me with my luggage before Freshman year and will soon cross it again to take my picture as I receive my diploma. I'll love you no matter where I am.

Contents

List of Figures

Chapter 1

Introduction

Adaptive filtering is a technique to reduce measurable noise source from a system with an unknown plant. It is commonly used for systems where the plant is subject to constant changes and can be implemented in different ways to compensate for or ignore instantaneous extremes. In this thesis I apply adaptive filtering methods to cancel AC power line noise at 60 Hz in a data stream that represents the output signal of a gravitational wave (GW) detector. I will discuss the simulation tools employed in this experiment and compare our experimental results to theory.

1.1   Gravitational Waves

Gravitational waves are described as ripples in space-time. They are emitted by accelerating quadrupolar mass distributions, such as the binary orbit of two black holes or two neutron stars. Einstein's theory of general relativity predicts that GWs cause periodic distortion of spacetime transverse to their direction of propagation, they travel at the speed of light, and they occur in two polarizations. When a gravitational wave with the + polarization hits test particles which were initially arrange in a perfect circle, the test particles exhibit a periodic motion as in figure 1-1. A gravitational wave with the × polarization, will have the same impact as the + one, rotated by 45 degrees.
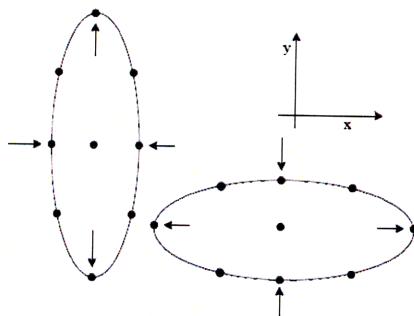
Figure 1-1: A schematic of displacement of test particles influenced by a gravitational wave with a + polarization. The states are separated by 180^0 phase [1].

1.2   Gravitational Wave Detector - LIGO

The Laser Interferometer Gravitational-Wave Observatory (LIGO) is based on a Michelson interferometer with Fabry-Perot cavities in its two perpendicular arms. In addition LIGO also employs power recycling to resonantly enhance the input laser power [3].
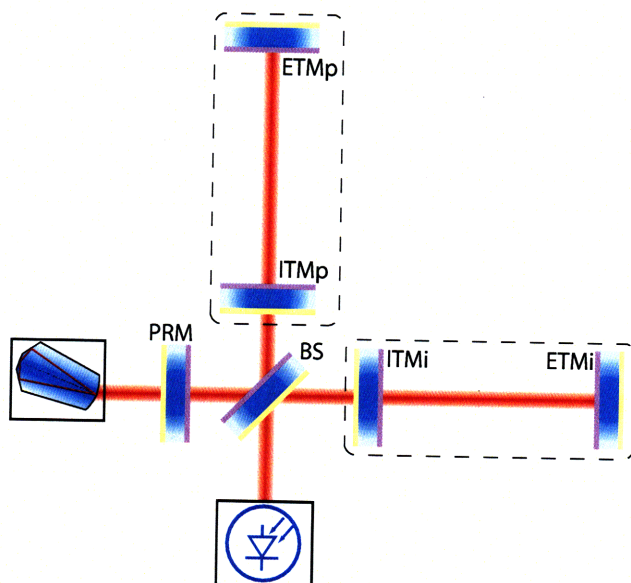

Figure 1-2: A schematic of LIGO, a power recycled Fabry-Perot Michelson interferometer, showing input laser, power recycling mirror (PRM), beam splitter (BS), Fabry-Perot arm cavities (encircled by dashed line) and readout photodiode.

Figure 1-2 shows a schematic of the initial LIGO detector. Laser light enters the interferometer through the partially transmissive power recycling mirror and is subsequently split between the two arms at the beam splitter. Each arm comprises a Fabry-Perot optical resonator, the purpose of which is to enhance the phase response of the detector to differential arm motion. All optics are hung as pendulums mounted on vibration isolation stacks to attenuate seismic disturbances. The reflectivities of the mirrors of the Fabry-Perot cavities are chosen so that most of the incident laser light is returned to the beam splitter where it recombines. The Michelson is nominally held at a dark fringe (to reduce laser noise couplings) so the vast majority of light is transmitted to the PRM which forms a resonator with the arms to effectively increase the input laser power. Any differential arm motion (e.g. a gravitational wave) causes light to leak to the readout photodiode where the signal is detected.

1.3 Suppression of Power Line Noise at 60Hz

LIGO is designed to detect gravitational waves (GW). However, the interferometer is only sensitive to GW when it is in a particular state. The LIGO control system is tasked with maintaining the interferometer mirrors at a specified set point. At present there are over 100 individual control loops which together hold the interferometer at its operating point. Optimizing the control loops will result in a higher signal to noise ratio, increasing the probability of detection.

Adaptive filtering shall become increasingly useful as future upgrades to LIGO increase sensitivity and uncover hitherto unseen noise sources. Amongst other phenomena, the next-generation detector, Advanced LIGO, will search for binary inspirals of high mass black holes. Such sources emit the majority of their GW at low frequencies, with higher mass binaries emitting at lower frequencies. Below $\sim$10 Hz any detectable signal is swamped by fundamental Newtonian gravity gradient noise and seismic disturbances (see figure 1-3). Using adaptive filters, it is possible to increase sensitivity in this low frequency band which will permit LIGO to make confident detections at redshifts of up to $z = 5$ [5].
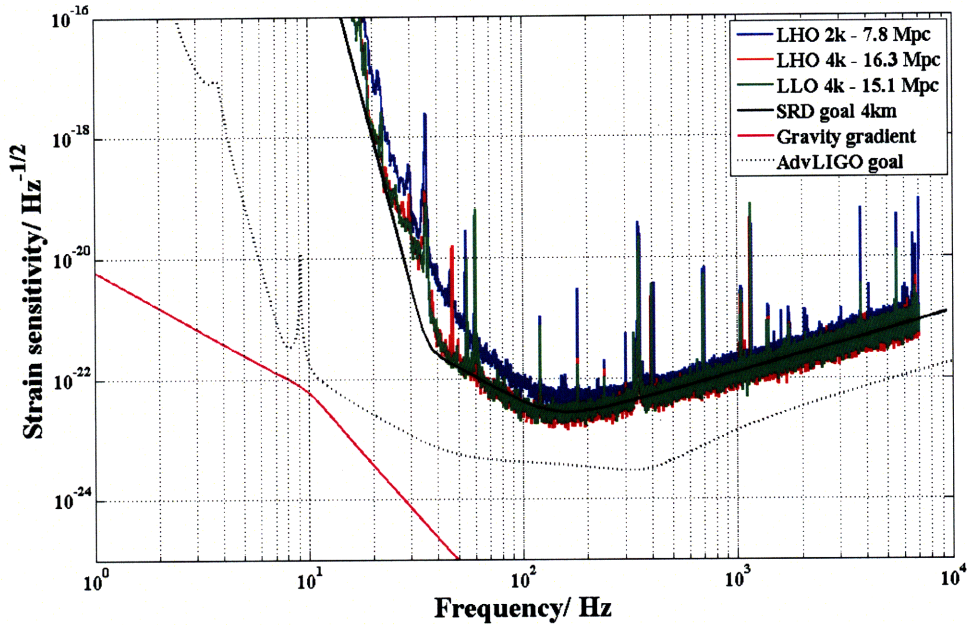
Figure 1-3: LIGO interferometer sensitivity curves from the S5 run (May 2007). Here we see the sensitivity limit at low frequencies ($<\sim$ 30 Hz) due to seismic noise. The ranges are set for a 1.4/1.4 $M_\odot$. In addition to seismic noise, Advanced LIGO will also be affected by Newtonian gravity gradients at the low end of its frequency band. http://www.ligo.caltech.edu/~jzweizig/distribution/LSC_Data/.

Adaptive filtering may be fruitfully applied in the suppression of 60 Hz noise arising from the AC electrical supply. During initial LIGO no attempt was made to actively suppress this noise; rather the collaboration relied on good practice from its engineers to isolate the interferometer's output from the 60Hz line signal. Attenuation of the 60Hz noise is particularly important due to the 59.7 Hz radiation frequency of the Crab pulsar. Figure 1-4 illustrates how the detection of the Crab Pulsar is impossible due to the interference with the 60Hz noise. During the recent science run LIGO was able to beat the classical spin-down [2]. Any mitigation of 60 Hz noise would increase this value and reduce the Crab limit still further, which could result in directly detecting GWs from the Crab pulsar.

Due to the concerns it raises for LIGO and the ubiquity of power line noise in other systems, I have chosen to make suppression of 60 Hz using adaptive filtering the keystone of my senior thesis [2].

14
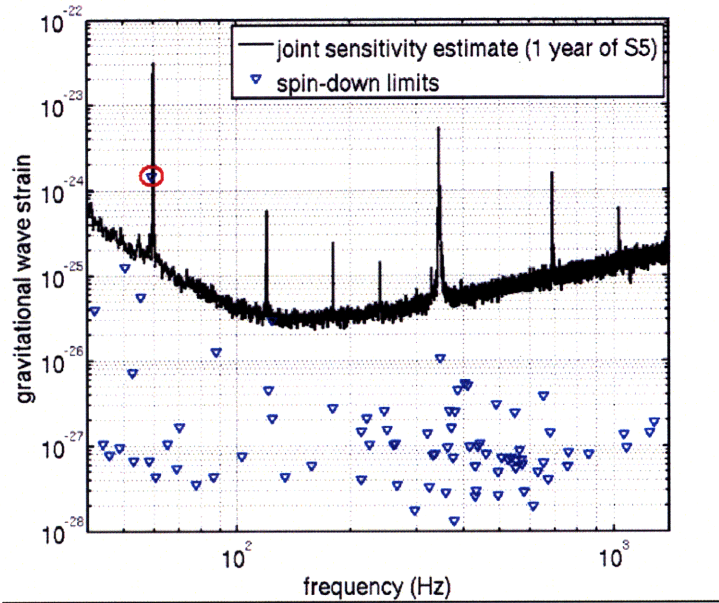
Figure 1-4: Search sensitivity during S5 run and the spin-down limits for pulsars in LIGO's sensitivity band [5]

In this paper I will present recent advances in the application of adaptive filter theory to noise cancellation in LIGO. Our goal is to remove the effects of the 60Hz noise sources from the gravitational wave channel. I will focus my discussion on the development of a multiple-inputs-single-output (MISO) system which has been shown to mitigate the effect of the 60Hz line on our detectors.

Chapter 2

Theory of Filtering Methods

In this chapter I will introduce the basics of digital filtering, needed for the discussion of adaptive filtering techniques that follows.

2.1 Finite Impulse Response (FIR) - Non-Adaptive

Finite impulse response (FIR) filters are commonly used in digital signal processing. This type of filter is referred to as finite since it does not include a feedback path; every signal impulse incident upon the filter will eventually result in an output of zero. The transfer function of a filter can be represented with a finite number of constant poles and zeros[7], where the poles are located at the origin of the Z-plane[1]. The following equation summarizes how an FIR filter with filter coefficients b_i is applied on an input signal $x[n]$ to result in an output signal $y[n]$. For a filter with $N + 1$ taps, the output is simply the convolution of the input signal and the filter (both vectors with size $N + 1$). As the equation demonstrates, there is no feedback (which would result in the dependence of the output in previous output signal), and therefore, once the signal passes and the input is simply zero, the output zeros as well.

$$y[n] = \sum_{i=0}^{N} b_i x[n - i] \tag{2.1}$$

---

[1] The Z plane represents real or complex time-domain inputs, in the complex frequency-domain using a generalized discrete-time Laplace transform.

A different method to represent an FIR filter is in the z-domain. Transferring to the z-domain and denoting the transfer function of the FIR filter as $F(z)$ we obtain

$$F(z) = \sum_{n=0}^{N} b_n z^{-n} \tag{2.2}$$

and the roots of the polynomial $F(z)$ define the zeros of the filter.

Another characteristic property of FIR filters is the number of taps (N) they employ. The number of taps corresponds to the number of poles and zeros. A longer filter is more computationally expensive but results in more aggressive filtering, facilitating sharper cutoffs and better defined bandpass filters [4]. Another important property of FIR filters is that they are unconditionally stable. As mentioned earlier, all incoming signals will result in an output of zero after a well-defined time has elapsed. As I will elaborate on when discussing IIR filters, for a filter with no feedback, all the poles are located at the origin of the Z - plane, which prevent the filter from going unstable.

2.1.1 FIR - Adaptive

An adaptive FIR filter is an FIR filter with time varying coefficients (b_i). The adaptive algorithm aims to minimize the error between some target signal (from which we wish to remove noise) and noise signals (which provide a valid indicator of the disturbance to be removed), taking into consideration their constant online changes. The noise signals are filtered by the ever-changing adaptive FIR filter, such that they emulate the desired output signal. Subtracting the desired known signal from the imitated one provides an error signal which is then used to modify the FIR coefficients. There are several different methods for constructing adaptive filters, however, they are all based on the same general principles of a least-mean-square (LMS) algorithm. A LMS algorithm used in an adaptive filter works to produce the filter's coefficients such that they minimize the least mean squares of the error signal. Since we are looking into a filtering method which is based on the time varying error signal, we decided to use the LMS method. Looking at the final representation of an adaptive Least mean

square (LMS) filter, it is quite simple to understand the significance of the variables that characterize a filter:

$$W(n+1) = W(n) + \mu \epsilon X(n) \qquad (2.3)$$

In equation 2.3, W represents the set of FIR coefficients, ϵ is the input from the error signal. X is the input noise signal vector, which has the same length as the number of taps. μ, the step size, is also referred to as the adaptation rate. Since ϵ, W and X cannot be deterministically controlled, the key variable in the equation is the adaptation rate. This step size has to be chosen carefully. By setting μ large, we can neglect instantaneous changes and the filter will converge more quickly; however, at the same time, this can cause the filter to diverge if the step is too large to reach the optimal minimum. In the search for minimizing the error, the system will be overdamped. This will cause the error to oscillate with high amplitude. If μ is too small, the adaptation time will take longer and the filter will converge slowly (underdamped). Finding the balance between a too small μ and a too large one will maximize the efficiency of the adaptive filter. Figure 2-1 shows an example of a SISO (single-input signal-output) adaptive filter. The noise signal and the error signal are used to modify the coefficients of the FIR filter.
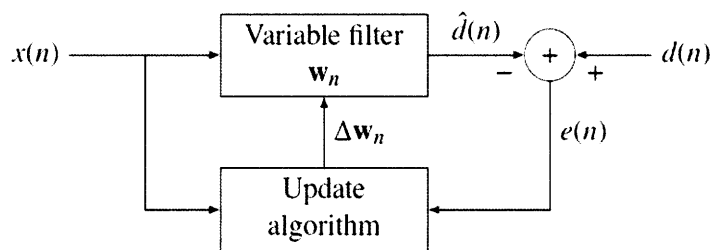

Figure 2-1: The adaptive filter shown can be implemented with a Least-Mean-Square (LMS) algorithm to predict the target signal will the incoming noise signal [6]

2.2 Infinite Impulse Response (IIR)

Unlike FIR filters, an IIR filter will not necessarily yield zero output after an infinite amount of time. This is due to their having poles away from the origin of the Z-plane. The IIR filter has both feedforward and feedback components. The feedback element leads to a dependence of the output on the previous output results. As in the FIR equation 2.1, we will continue denoting the output as $y[n]$, input signal $x[n]$, feedforward coefficients b_i with order P and feedback coefficients a_i with order Q. An IIR filter can then be written in the following way:

$$y[n] = \frac{1}{a_0} \left(\sum_{i=0}^{P} b_i x[n-i] - \sum_{j=1}^{Q} a_j y[n-j] \right) \tag{2.4}$$

and after taking the Z-transform and considering that most IIR filters are designed with $a_0 = 1$, the transfer function of the filter $F(z)$ is:

$$F(z) = \frac{\sum_{i=0}^{P} b_i z^{-i}}{1 + \sum_{j=1}^{Q} a_j z^{-j}} \tag{2.5}$$

Looking at equation 2.4, we can notice the feedback path of the filter which is represented as the dependency of the output $y[n]$ in the previous output $y[n-j]$. This part distinguishes an IIR from an FIR. While an FIR has only a feedforward component and will therefore result in a final output of zero, the IIR has this feedback path which allows an infinite output for a finite input. The poles of the filter are the values of z which will make the denominator of $F(z)$ equal to 0. Now it is clear why an FIR filter has all its poles located at the origin of the z-plane. However, for an IIR, where $a_j \neq 0$, the poles are not limited to the origin. Since a pole implies a z value for which $F(z) \rightarrow \inf$, we learn that IIR filters are not inherently stable. Although they are not guaranteed to be stable, IIR filters are sometimes preferable over FIR filters. Due to the feedback component, an IIR filter results in a much better bandpass cutoff than an FIR filter of the same number of taps.

2.3 Wiener Filter

A Wiener filter is often use to assess the performance of a linear adaptive filter. Although in our experiment we will be using real-valued time series, the Wiener filter theory is formulated for the more general complex valued case. Following the derivation from pp. 96 – 104 of Haykin [6], we define the variables we are going to use to construct the Wiener filter. $y(n)$ will be the output of our linear discrete-time filter which provides an estimate of the desired response $d(n)$, for a given set of inputs $u(0), u(1), u(2), ...,$. The mean-square value of the estimated error is defined as $e^2(n) = (d(n) - y(n))^2$ and is minimized. For an input with $M - 1$ one-sample delays we will denote $\mathbf{R}$ as the $M \times M$ *correlation matrix* of the tap inputs $u(n), u(n - 1), ..., u(n - M + 1)$ such that

$$\mathbf{R} = E[\mathbf{u}(n)\mathbf{u}^H(n)] \tag{2.6}$$

where E denotes the statistical expectation operator and

$$\mathbf{u}(n) = [u(n), u(n - 1), ..., u(n - M + 1)]^T. \tag{2.7}$$

We will further denote $\mathbf{p}$ as the $M \times 1$ *cross-correlation vector* between the tap input and $d(n)$

$$\mathbf{p} = E[\mathbf{u}(n)d^*(n)] \tag{2.8}$$

For $\mathbf{w_0}$ denoting the $M \times 1$ *optimal tap-weight vector* in the mean-square-error sense we find

$$\mathbf{Rw_0} = \mathbf{p} \tag{2.9}$$

which can be solved for the optimal tap-weight

$$\mathbf{w_0} = \mathbf{R}^{-1}\mathbf{p} \tag{2.10}$$

We will use the Wiener filter to predict the optimal noise suppression we can achieve and find the maximum step-size (μ).

2.4 Step-Size Optimization

Choosing the correct step-size for an adaptive filter is crucial. In this experiment we will use the method of steepest descent for choosing our optimal step-size parameter μ. I will briefly describe the method and refer to Haykin p. 203-227 [6] for those interested in further readings. When the method of steepest descent is applied to the Wiener filter it provides an algorithm for keeping track of the time variations without having to resolve the Wiener - Hopf equation 2.9 each time the signal changes. The important point to note is that the method of steepest descent adjusts the successive weight vectors $\mathbf{w}$ in the direction opposite to the gradient vector of the cost function. The cost function $J(\mathbf{w})$ is a continuously differentiable function which maps the elements of $\mathbf{w}$ into real numbers such that for the optimal solution $\mathbf{w_0}$,

$$J(\mathbf{w_0}) \leq J(\mathbf{w_0}) \forall \mathbf{w}. \tag{2.11}$$

Following Haykin's math, one finally reaches the conclusion that to verify the convergence and the stability of the steepest-descent method, the step-size parameter must satisfy

$$0 < \mu < \frac{2}{\lambda_{max}} \tag{2.12}$$

where λ_{max} is the largest eigenvalue of the correlation matrix $\mathbf{R}$. When the step-size is too small, the transient behavior of the steepest descent algorithm is overdamped, and for a large step-size, it is underdamped. Therefore, it is important to run a Wiener simulation for our data and make sure we are choosing a step-size within this range. Solving for the Wiener filter, we can learn what is the optimal result that we can expect from an adaptive filter. In addition, by measuring the correlation matrix R, we can extract λ_{max} and calculate which step-size we should use. The Wiener simulation requires knowing the sensors' signals and the witness signal. Since these

are the signals we can measure, we can use a Matlab code (appendix A) to solve for the correlation matrix and the filter.

Chapter 3

Experiment

In this chapter I describe the apparatus which was built to experimentally study adaptive filtering methods at the LIGO Laboratory, Massachusetts Institute of Technology. I discuss how this was designed to reach optimal results in subtracting the 60Hz AC noise from a chosen witness, and how the methods we used can be implemented in LIGO's sites, minimizing the coupling of the magnetic field caused by the 60Hz AC current to the magnets which are an essential part of the test masses' control system.

3.1 Apparatus

A schematic of the experiment is shown in figure 3-1. The witness was placed in the center of a 25 inch diameter Helmholtz coil (which I will further refer to as the driving coil). The sensors, as the witness, were almost identical coils of wire with a resistance of 2Ω and an inductance of 30mH. The transfer function between the sensors and the driving coil is presented in figure 3-4. The four sensors were placed next to high powered instruments, which we used as our background noise magnetic field generators, around the witness. All the signals of the sensors and the witness were connected to our computers through a low-pass filter with a gain of 110 (see the preamplifier circuit shown in Fig. 3-2). This preamp was modified for our needs with a 1kHz low pass and a 100Hz low pass on the gain's feedback, providing a filter with
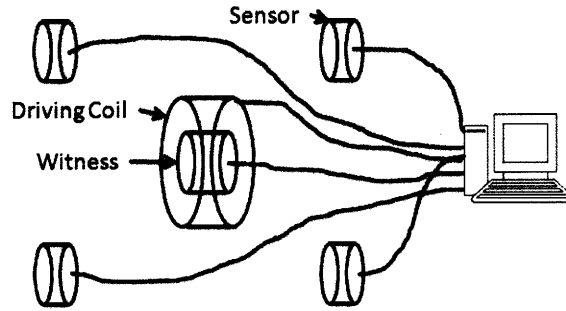
Figure 3-1: This is a schematic of the apparatus that was used in the experiment. The driving coil with the witness at its center were placed between for detecting sensors. The sensors, driving coil and witness were all connected to the ADC through a preamplifier. The driving coil was also connected to an output DAC port.

two poles at 100 and 1,000 Hz (preamp transfer function is shown in figure 3-3). The purpose of the 1kHz low pass was to eliminate high frequency noise. Once we noticed the large contribution of the aliasing of higher harmonics of 60Hz into our band pass, we added the 100Hz low pass on the gain to moderate this effect.
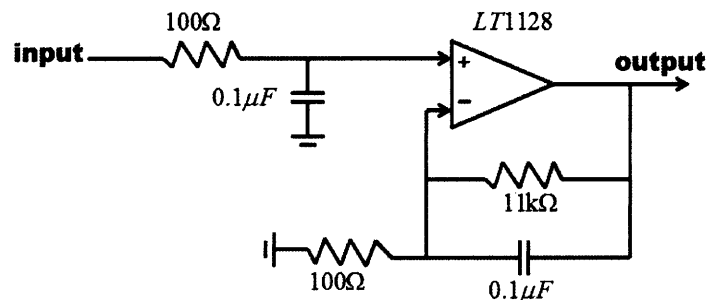

Figure 3-2: A scheme of the preamp used for all sensors, witness and the driving coil. Notice the 1kHz low pass on the input, 100Hz low pass on the gain feedback and a total gain of 110

3.2 Adaptive Filter

Let us start by going through the block diagram in figure 3-5 that represents the apparatus mentioned above. The sensors, as the witness, detect the change in the magnetic field in their specific location. The transfer functions between the sensors'
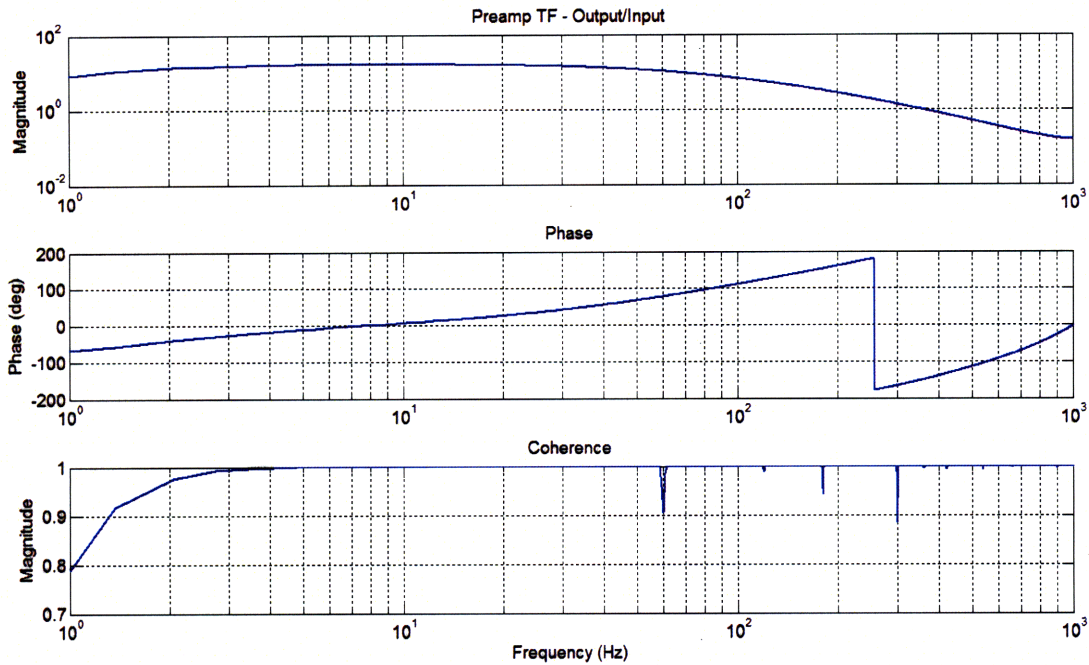
26

Figure 3-3: The transfer function between the output and input of the preamp. It is possible to notice the pole at 100Hz. Since the Nyquist frequency in the measurement is 1024Hz, it is harder to notice the 1kHz pole. This transfer function demonstrated the 110 gain of the preamp.

signals to the witness signal is the filter we are about to simulate using the adaptive filter. The signal of the sensors is also used for driving the driving coil. We measured and found that the driving of the driving coil has a negligible effect on the magnetic field around the lab in comparison to its effect on the witness (see figure 3-6), and therefore, is detected by the sensors, we account for the signal of the sensors as the sum of the magnetic field from the room and the driving coil together. As we are trying to cancel the 60Hz line detected by the witness, and assuming the driving coil is driven to accomplish this mission, we can claim that $(y(n) + d(n)) \times$ plant $= e(n)$ when $e(n)$ is the witness's signal we are trying to minimize and the plant is the transfer function between the sensors' readout to the witness's readout. As discussed in section 2.1.1, the error signal and the sensors signal is used to adapt the FIR filter. However, in our case the math has slightly changed as we will now explain. For sensors signal S, unknown plant P, adaptive filter H, plant model M, error signal e,
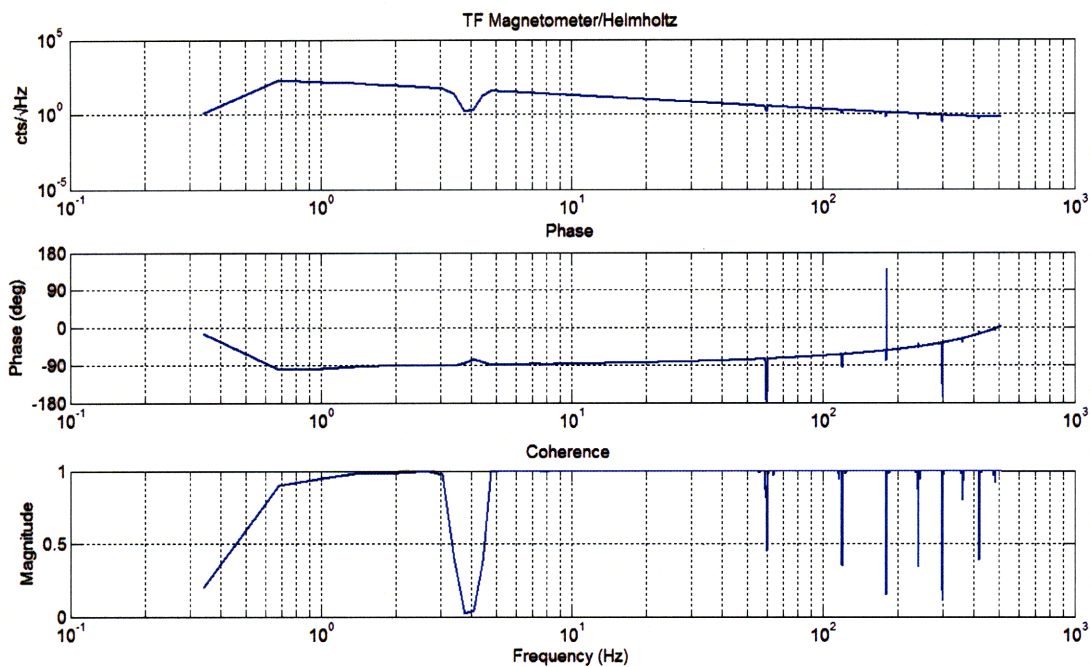
Figure 3-4: The transfer function between the sensor and the driving coil. All 4 sensors had an almost identical TF as this one. As we will discuss in the results section, the 4Hz signal is due to aliasing of higher harmonics of the 60Hz line. Therefore, the coherence at 4Hz drops sharply.

output filter O and error filter G as described in figure 3-5 we can write:

$$d(n) = S \times P$$

$$y(n) = S \times H \times O$$

$$\Rightarrow e(n) = (SP - SHO)M \tag{3.1}$$

Now that we represented $e(n)$ with the parts of our model, we should take a break to understand the maths behind the adaptive part of the filter. Since we are interested in minimizing the absolute error within our bandpass, let us first define this signal and represent it as followed:

$$\epsilon = (Ge)^2 \tag{3.2}$$

and rewrite equation 3.1 as

$$\epsilon^2 = [GM(SP - SHO)]^2. \tag{3.3}$$

A common way to visualize the mean-squared-error signal (ϵ), is that it follows a bowled-shape surface with a unique minimum. The process of adaptation works to seek that minimum using the optimization technique of steepest decent. To compute the gradient vector $\Delta(n)$ we differentiate both sides of equation 3.3 with respect to the filter H and obtain

$$\frac{\partial \epsilon}{\partial H} = 2G^2 e \frac{\partial e}{\partial H} = -2G^2 eSOM \tag{3.4}$$

and we therefore derived

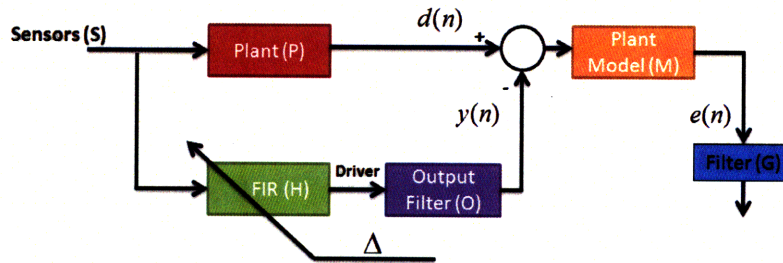$$W(n+1) = W(n) + 2\mu G^2 eMOS \tag{3.5}$$

Figure 3-5: A block diagram of the apparatus. The sensors' signals pass through a plant (which in our case will be the bandpass around 60Hz) and is denoted $d(n)$. At the same time, the unmodified signal passes through the adaptive FIR filter to drive the coil and obtain an output signal $y(n)$. The driving signal is then subtracted from $d(n)$ and the combined signal then passes through the plant model which is the transfer function between the driving coil and the witness. This transfer function simulates how the witness signal reacts to a magnetic field. Following the maths of equation 3.5, $\Delta = 2\mu G^2 eMOS$ and is added to the filter H in each computational cycle.

In the next section we will demonstrate how we implemented the control system
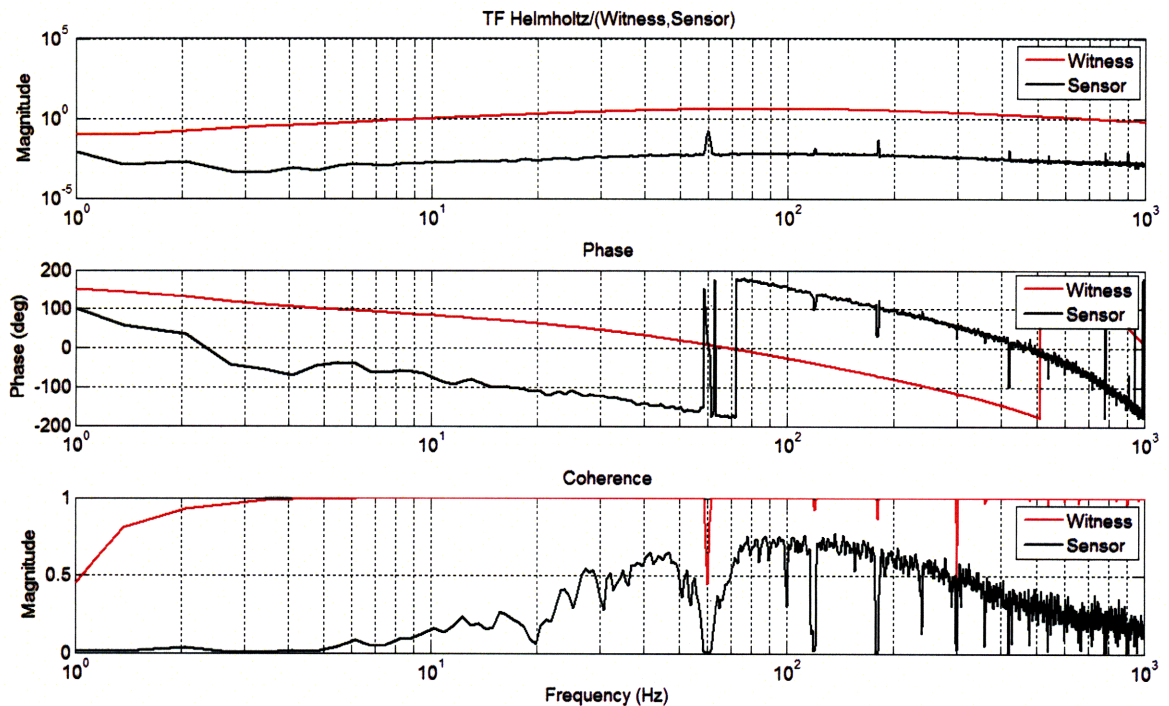
Figure 3-6: In red is the transfer function between the driving coil and the witness. The coherence between the driving coil and the witness is generally 1 in comparison with the effect on the magnetic field detected by a sensor in the lab. The effect on the witness is also approximately 5 orders of magnitude above the effect on the sensor. Therefore, we did not add the change in magnetic field in lab caused by the driving coil to our equations.

in Simulink.

3.3   Software Design

Simulink is an integral application of Matlab which is simple to work with and is compatible with the the hardware used for the experiment. DSpace, which is the hardware and software we used, is a real time program loaded into a front end supplied by the manufacturer dSpace. It has an ADC and a DAC which was used for driving the coil. The software was designed to match the theory from the previous section. Starting from the input signals of the sensors in figure 3-7, we placed an IIR bandpass to limit the optimization of the filter to the desired 60Hz frequency. Since we want

to correlate the witness signal to the filter we obtain, it is of course important to look at the witness signal after it passes the same IIR filter as the sensors one. This agrees with the theory in the previous section if we consider the bandpass filter to be a part of the filter G which appears in the delta as G^2. To find the plant model we took the transfer function between the driving coil and the witness, simply by driving the driving coil with white noise.
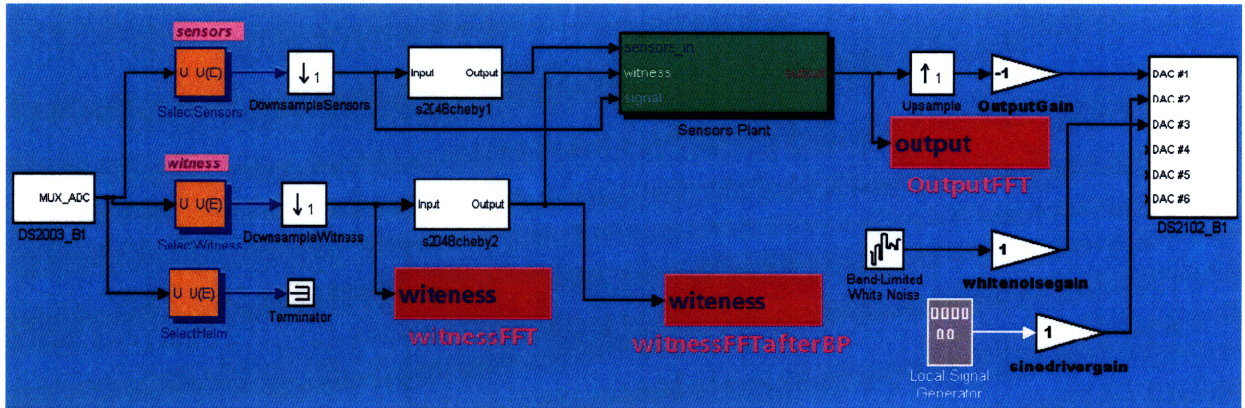

Figure 3-7: The simulink screen used. The input channels in the MUX_ADC include the sensors and the witness signal. The sensors signals pass through an IIR bandpass and then enter the sensor plant box which can be seen in figure 3-8 where they are used to calculate the FIR adaptive filter. The output of the sensors' signals times the FIR filter are then used to drive the driving coil which is connected to DAC#1.

The system described is a multiple-inputs-single-output (MISO) which requires us to choose a method for weighting the inputs and combining them to one output. We examined two methods: 1) Combining the signals from the sensors before calculating the filter. 2) Calculating the filter for each sensor individually, and combining the outputs for a single drive. As can be expected, while the first method costs less calculation time, the second is more effective. The first method requires the computer to make one set of extensive calculations to find the adjustment for the filter. However, it also gives each sensor the same constant weight and does not take into consideration the time dependence of the transfer function between each sensor and the witness. The transfer function between each sensor and the witness depends on the distance between them and the location of the magnetic field source which is not always known

and changes with time. On the other hand, while the second method requires $m \times$ calculations (where m stands for the number of sensors), it accounts for the differences between the transfer functions of each sensor and the witness at any point in time. In case the computer is incapable of calculating m filters, a possible solution is to use a weighting system which is based on the coherence between each sensor and the witness. As our CPU is fast enough to perform these calculations, we chose the second method for the experiment.
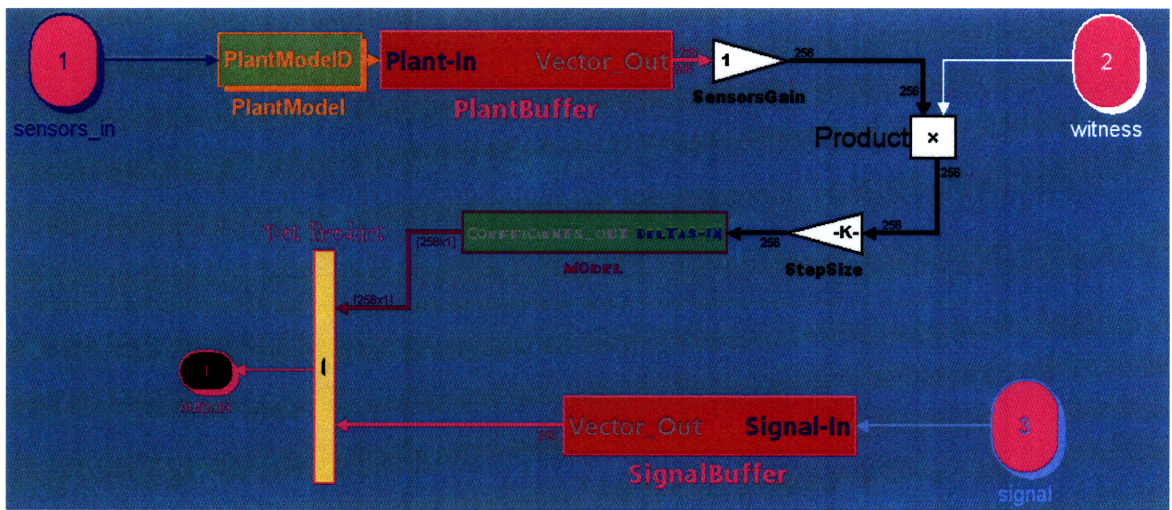

Figure 3-8: This simulink model was used to calculate the Δ of the FIR filter, update the FIR and pass the sensors' signals through it to find the output for the driving coil. The plant model was measured between the driving coil and the witness to find the affect of a changing magnetic field on the witness output. In the MODEL box we updated the filter by adding the new Δ, and the sensors' signals were multiplied by the FIR to calculate the output for the driving coil.

3.4   Procedure

We first used a Matlab code to find the Wiener filter for our data. This was done by taking approximately 10 minutes of data from all our sensors and the witness, and using a Wiener filter algorithm (appendix A) to find the optimal filter W. As described in section 2.4, the maximum step-size was calculated from the correlation matrix R. We initially chose a sampling rate of 1,024Hz and an FIR length of 512

which implies a Nyquist frequency of 512Hz. The desired step-size was entered into the control interface and we measured the performance of the filter over time. We followed this process with sampling rates of 512 and 2,048 Hz and compared the results of each sampling rate for different filter lengths (128, 256, 512). Moreover, to verify the system responds to changes in the magnetic field, we moved the magnetic field generators and followed the changes in the filter. In the next section we will present the results and conclude which setting minimizes the 60Hz noise.

Chapter 4

Results

4.1 Choosing a Sampling Rate

A sampling rate has two major effects: 1) Computer memory consumption 2) Frequency resolution. Since we are using a strong computer and RAM does not restrict us, we will try to find the optimal frequency resolution. In order to do so, we took three sets of 10 minutes data. Each set had a different sampling rate: 512Hz, 1024Hz and 2048Hz. The results we present in figure 4-1, were obtained using the Wiener filter simulation with a filter length N=256. A quick glance at the figure reveals an interesting phenomena - a peak at 4Hz, 4Hz side bands on the 60Hz line, and in the 512Hz sampling rate case, some other peaks with 4Hz gaps between them. Although our first intuition (with a sampling rate of 1024Hz) was to relate the 4Hz peak to a mechanical resonance of our apparatus, we quickly learned by applying changes to it that this was not caused by a mechanical property. The 512Hz data revealed the real cause for these peaks. As the Nyquist frequency for a sampling rate of 512Hz is 256Hz, the unknown peaks are indeed the aliasing of the higher harmonics of the 60Hz line over 256Hz. For example, the peak we see at 32Hz, is nothing else but a "'mirror"' image of the 8th harmonic (at 480Hz) folding over the 256Hz line. The 28Hz peak, is due to the folding of the 9th harmonic at 540Hz. With the 1024Hz sampling rate, the effect is moderated by the preamp's poles at 100Hz and 1kHz. The filter decreases the magnitude of the higher harmonics and they are, therefore, less

noticeable once folded along the Nyquist frequency of 512Hz. At a sampling rate of 2048Hz with a folding frequency of 1024Hz, the effect is no longer detectable. Therefore, in order to continue the experiment with our current hardware, we decided to test our control system using a 2048Hz sampling frequency.
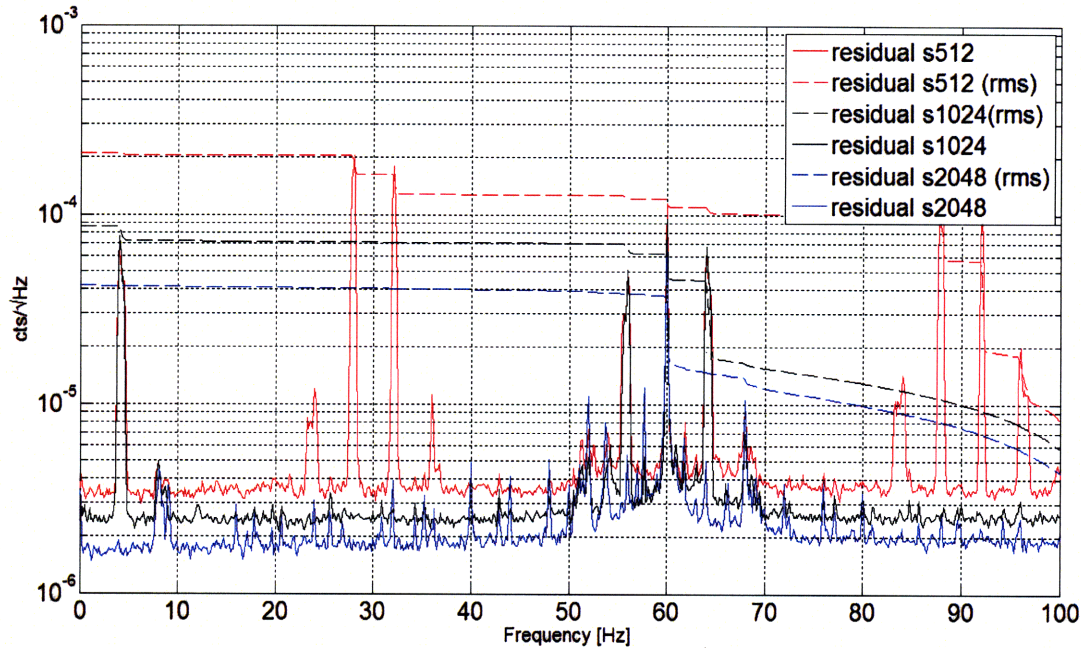

Figure 4-1: Results of a measurement done to determine optimal sampling frequency. The results are of a Wiener filter simulation, and are plotted for different sampling rate frequencies of 512, 1024 and 2048 Hz. The doted curves are the cumulative RMS of the data to allow easier comparison between the curves. The affect of the "'mirror"' imaging along the folding frequency (Nyquest frequency) is mostly detectable at 512Hz (red curve) sampling frequency. Due to the 100Hz and 1kHz poles in the preamp design, the affect is moderated when using a 1024Hz sampling rate (black curve), and is unnoticeable when sampling at 2048Hz. From this measurement we learn that for our setup we can obtain the best results using a 2048Hz sampling rate.

4.2 Choosing a Filter Length

Now that we chose our sampling frequency, we need to test which filter length can maximize the isolation at 60Hz. Again ignoring the heavy calculations a long filter may require, we tested for the differences between a 128, 256 and 512 taps filters. In figure 4-2 we are plotting the results we obtained. A closer look reveals that there

is no significant difference between 256 and 512 taps, however, 128 taps magnify the 4 and 8 Hz side bands. Therefore, we decided to continue our experiment using 256 taps filter.
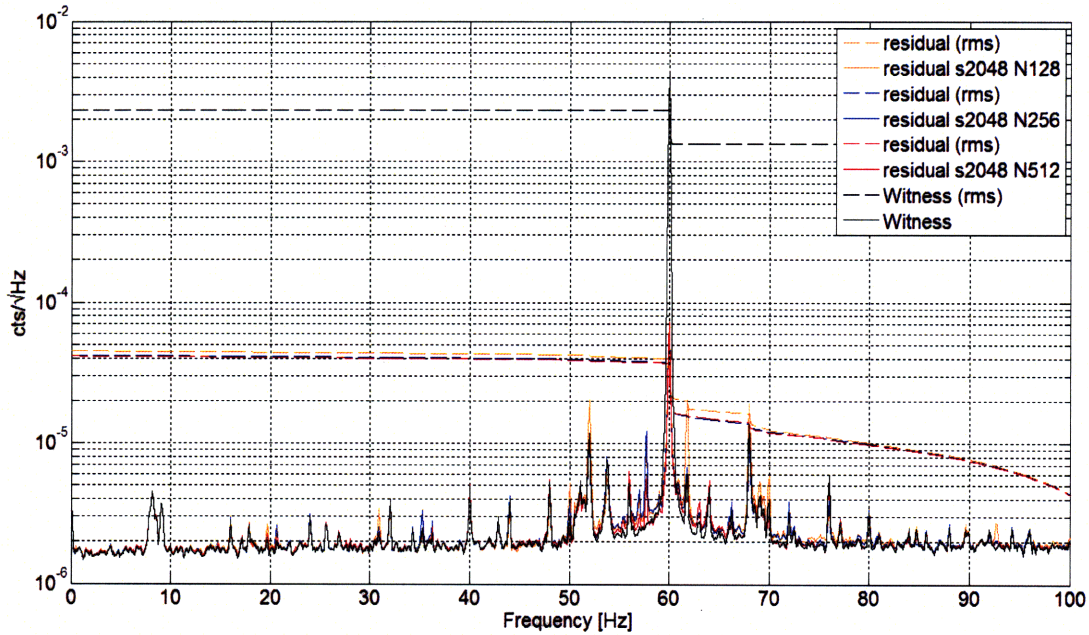

Figure 4-2: A measurement done in order to understand the difference between the lengths of the FIR filter (128, 256 and 512 taps). The black curve is the magnitude of a 10 minutes data set with a sampling rate of 2048Hz. The filtered results were taken from the Wiener filter simulation using the different tap number. With 128 taps (orange curve) we seem to under-model the system with enlarged peaks at 52 and 68 Hz. At 256 and 512 taps there is no noticeable difference which implies we should use 256 taps and decrease our calculation time.

4.3 Isolation

Before finally using the measured parameters, we placed our sensors next to our magnetic field generators and our witness in the center of the driving coil. After collecting data from different times and calculating no significant change in the optimal step size we obtain from the correlation matrix, we used 7×10^{-6} as our step size for all sensors. As this is a small step-size, we expected a long adaptation time. As the purpose of this thesis is to build an adaptive filter, let us now demonstrate that
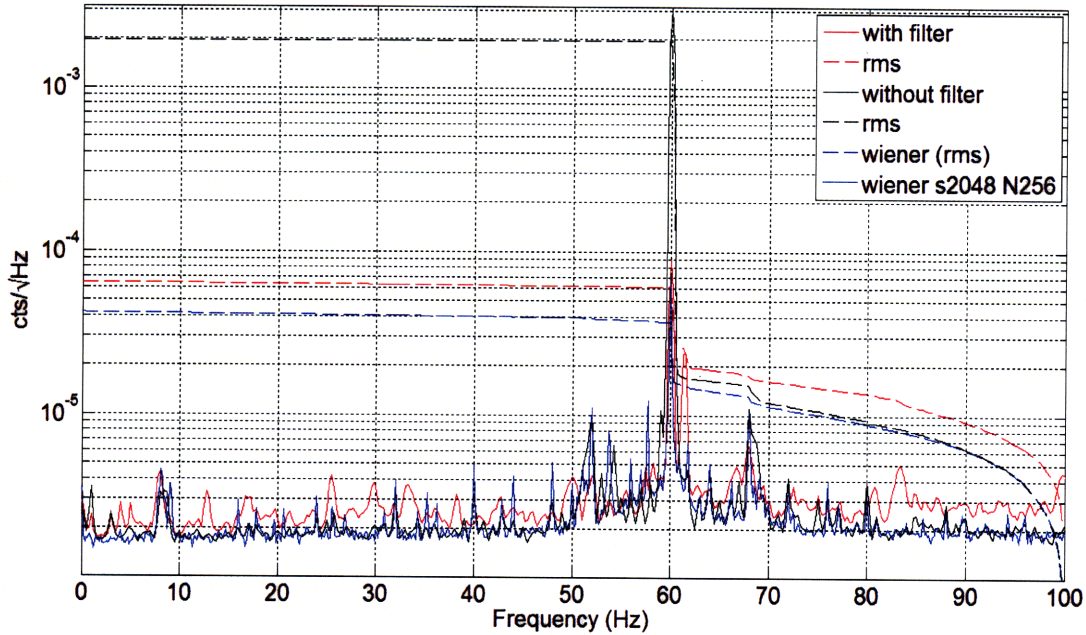
Figure 4-3: Here we compare the witness signal before applying the control system (in black), when using the Wiener filter simulation (in blue) and after adaptation of the filter (in red). This was taken using a sampling frequency of 2048Hz and using an FIR filter with 256 taps. The step size was 7×10^{-6}. Comparing the RMS we can notice how the filter reduced the 60Hz noise from approximately 2×10^{-3} to 5×10^{-5} which is only twice more than the expected Wiener Filter.

other than adapting to a specific filter, the system indeed adapts itself to changes in the magnetic field. In order to verify this, we waited for the filter to adapt as was done in figure 4-3 and then changed the orientation of several of our magnetic field sources. At first, the signal was significantly magnified. In figure 4-4 we track the witness signal with time, and find that the root-mean-square (RMS) decreases from approximately 3.5×10^{-3} to 9×10^{-5}. If we define the magnitude of the RMS (Z) dependence in time as

$$Z = Ae^{-\frac{t}{\tau}} \tag{4.1}$$

where A is a constant and τ is out time constant, we find τ to be 10 minutes.
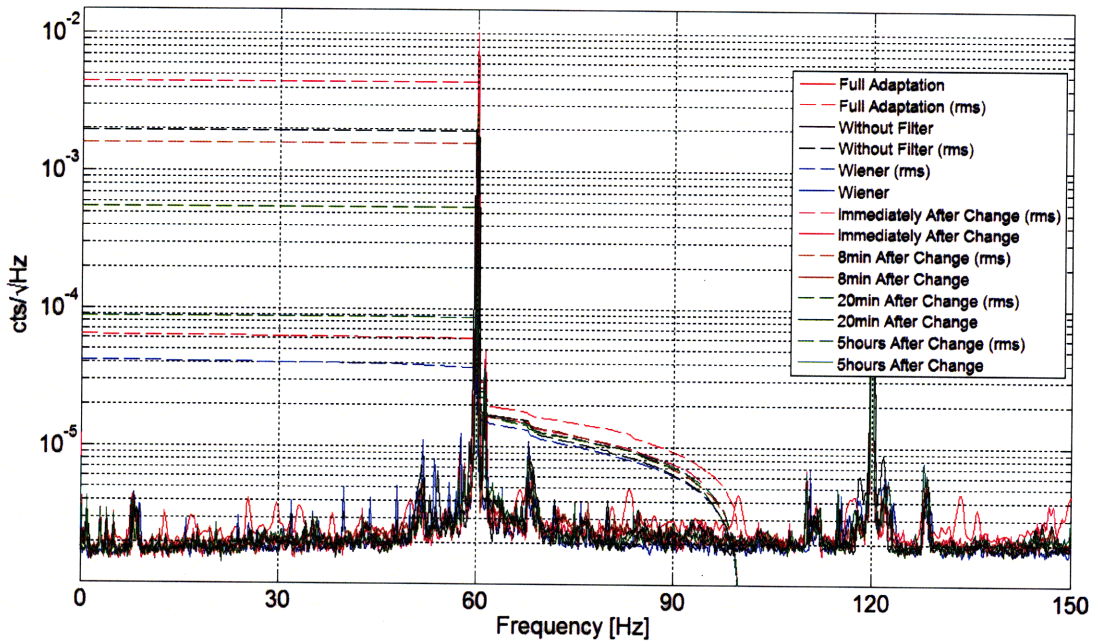
Figure 4-4: Starting from the unfiltered witness signal (in black), we ran the adaptive code until it fully adapted (in red) and compared it to the Witness filter (in blue). We then applied a change to the orientation of the magnetic field generators and followed the witness signal immediately after the change (in pink). We continued following the adaptation after 8 minutes (in brown), 20 minutes (in dark green) and 5 hours (in light green). The adaptation time constant for this adaptation is approximately 10 minutes. We can explain the fact that the signal did not return to its starting point (in red) since the correlation between the sensors to the witness, and therefore the Wiener filter found earlier is irrelevant.

4.4 Filter

After following the development of the witness signal, we would now shift our attention to the filter itself. Looking at the frequency response and the phase of the filter, we could conclude the filter is indeed working to minimize the 60Hz noise. In figure 4-5 we demonstrate how our filter imitated the expected results from the Wiener simulation. The filters drive the driving coil at 60Hz against the magnetic field in the lab to achieve maximum cancellation. Although the 4 and 8 Hz coupling are within our band pass, they do not affect the filter. This is because out filter is a linear filter, and the 4 and 8 Hz coupling are due to non linear aliasing around the Nyquist frequency. When comparing the phase of the filters in figure 4-6, we can see

it also agrees starting from zero phase until 60Hz. At 60Hz the filter gains phase as the output is expected to have a different phase than the input. That is, in order to cancel the 60Hz noise, the driving coil should drive at 60Hz. However there will exist a phase delay between the input signal and the drive.
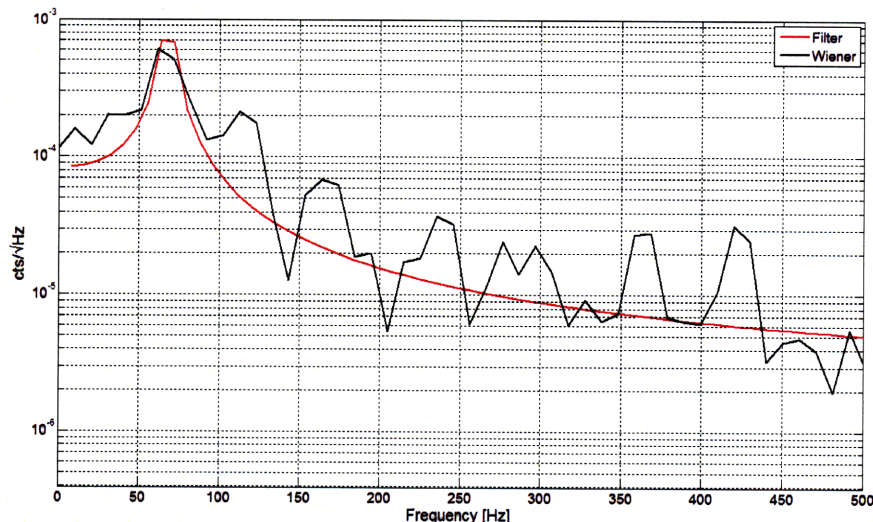

Figure 4-5: This plot is a frequency domain representation of one of the filters. In red, is the actual filter after running the adaptive code. In black, is the filter calculated using the Wiener filter simulation. Notice that the two filters peak at 60Hz and follow the same decay for higher frequencies. From this figure, together with the phase comparison of the two filters in figure 4-6, we can conclude that our filter indeed follows the expected Wiener result.

4.5   Conclusions

In this experiment we were able to suppress the power line noise at 60Hz in a factor of 500. We applied a LMS adaptive filtering method with our apparatus, which resulted in a 10 percent deviation from the optimal Wiener simulation. Therefore, we can conclude that an adaptive filtering method can be a useful instrument to LIGO and a part of the elaborated noise cancellation mechanism it holds. In figure 4-7 we inject a Crab pulsar signal with a 59.7Hz frequency and an amplitude of 10^{-4}. The pulsar's signal is reduced by only 0.29% while the 60Hz signal is reduced by 96.61%. We demonstrated that the adaptive filter is capable of minimizing the 60Hz
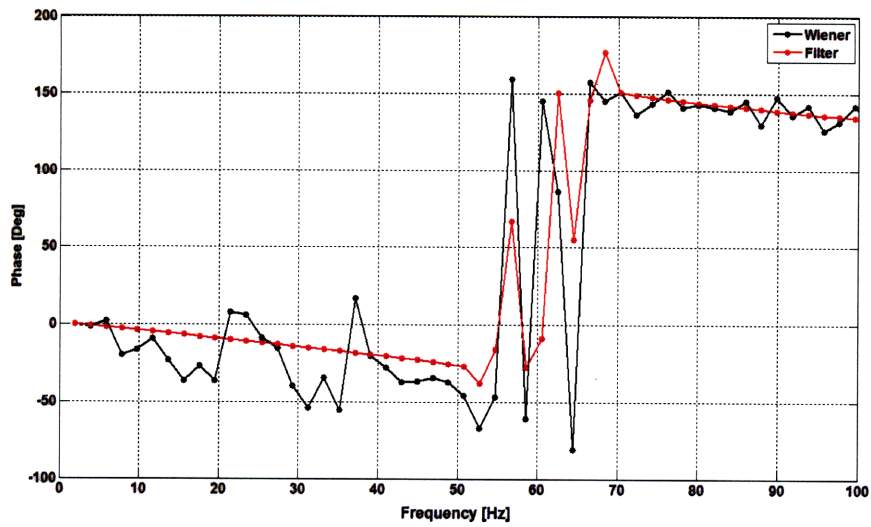
Figure 4-6: This plot compares the phases of the filters of the Wiener calculation (in black) and the real filter (in red). As in figure 4-5, the theory and experiment seem to agree.

noise, leaving the Crab pulsar's signal hardly effected, reducing the SNR for possible detection.
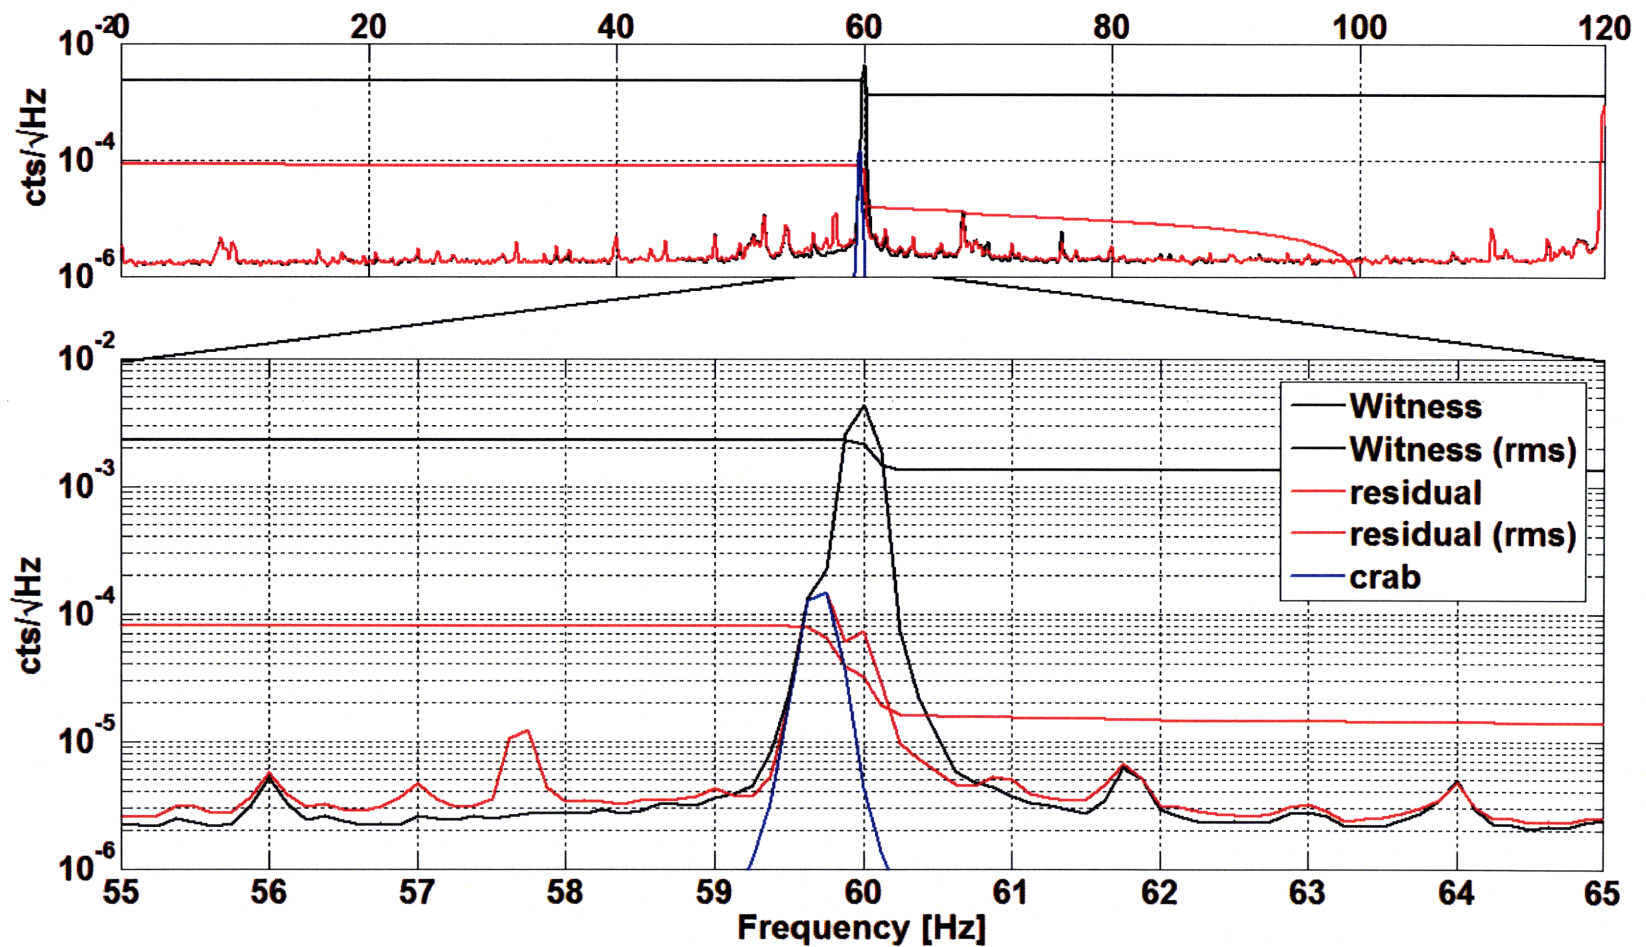
Figure 4-7: We injected a signal from the Crab pulsar at 59.7Hz with an amplitude of 10^{-4} in the signal of the witness (in blue). The Wiener works to minimize the correlated signal between the sensors and the witness, and therefore does not try to cancel the Crab signal. The filter reduced the signal from the pulsar by 0.29%, while reducing the 60Hz signal by 96.61%.

Chapter 5

Future Work

The experiment was limited by several factors that should be considered when implementing the adaptive filtering method in enhanced LIGO or in adaptive noise cancellation experiment in general. The lab where the experiment took place, bordered a machine shop and a different lab with a high power generator from the top. Therefore, although the sensors where placed next to what we believed to generate the strongest magnetic field, this was probably not the case. It is important to recognize the noise sources one is trying to eliminate and select the location and orientation of the sensors accordingly. Moreover, in order to decrease the signal to noise ratio, it is recommended to use a higher gain amplifier to increase the input and output resolution in the desired range. Since we detected the effect of aliasing along the Nyquist frequency, we learned it is important to have a steeper filter at a frequency low enough to avoid the contamination from higher frequencies into the observed frequency band.

5.1   Implementation in LIGO

As the witness signal, the GW signal can be used. The test mess in the chamber is driven by an Optical Shadow sensor and Electro-Magnetic actuator (OSEM) which will replace the driving coil. The sensors will remain in the form of a coil or a magnetometer placed in the vicinity of the test masses (note fig. 1-2). My work in this thesis has shown that if implemented in the LIGO detectors, an adaptive filter

will be useful in suppressing the power line noise at 60Hz and join the numerous methods in improving the interferometer's sensitivity.

Appendix A

Matlab Wiener Filter Code

```
function [W,R,P] = miso_firwiener(N,X,y)
%MISO_FIRWIENER Optimal FIR Wiener filter for multiple inputs.
%   MISO_FIRWIENER(N,X,Y) computes the optimal FIR Wiener filter of order
%   N, given any number of (stationary) random input signals as the columns
%   of matrix X, and one output signal in column vector Y.
%   Author: Keenan Pepper
%   Last modified: 2007/08/02
%   References:
%     [1] Y. Huang, J. Benesty, and J. Chen, Acoustic MIMO Signal
%         Processing, Springer-Verlag, 2006, page 48
% Number of input channels.
M = size(X,2);
% Input covariance matrix.
R = zeros(M*(N+1),M*(N+1));
for m = 1:M
    for i = m:M
        rmi = xcorr(X(:,m),X(:,i),N);
        Rmi = toeplitz(flipud(rmi(1:N+1)),rmi(N+1:2*N+1));
        top = (m-1)*(N+1)+1;
        bottom = m*(N+1);
```

```
        left = (i-1)*(N+1)+1;

        right = i*(N+1);

        R(top:bottom,left:right) = Rmi;

        if i ~= m

            R(left:right,top:bottom) = Rmi';   % Take advantage of hermiticity.

        end

    end

end


% Cross-correlation vector.

P = zeros(1,M*(N+1));

for i = 1:M

    top = (i-1)*(N+1)+1;

    bottom = i*(N+1);

    p = xcorr(y,X(:,i),N);

    P(top:bottom) = p(N+1:2*N+1)';

end


% The following step is very inefficient because it fails to exploit the

% block Toeplitz structure of R. It's done the same way in the built-in

% function "firwiener".

W = P/R;}
```

Bibliography

[1] Gravitational Waves: An Introduction, Indrajit Chakrabarty, arXiv:physics/9908041 v1, Aug 21, 1999.

$http://arxiv.org/PS_cache/physics/pdf/9908/9908041v1.pdf$

[2] B. Abbott el al, Beating the Spin-Down Limit on Gravitational Wave Emission from the Crab Pulsar, The Astrophysical Journal, 683:L45-L49, August 10, 2008

[3] An Overview of Recycling in Laser Interferometric Gravitational Wave Detectors, D.E. McCelland, *Australian Journal of Physics*, Vol.48, 1995, P.953

[4] Digital signal processing (3rd ed.): principles, algorithms, and applications, J. Proakis, D. Manilas, Prentice Hall, 1996, ISBN:0-13-373762-4

[5] M.A. Bizouard el al, 5th Annual Report on GWA-WG2 Joint operation of antennas and network data analysis $www.ego - gw.it/ILIAS - GW/documents/N5AnnualReport2009/NA5 - GWA - 5th - AnnualReport - WG2 - v4.doc$

[6] Simon Haykin, Adaptive Filter Theory, Prentice Hall, 2002, ISBN 0-13-048434-2

[7] Dietrich Schilchtharle, Digital Filters: basics and designs, Springer, 2000, ISBN 3540668411