

Spring 2016

Well, It's About Time

Daniel David Gagné
Bard College

Follow this and additional works at: https://digitalcommons.bard.edu/senproj_s2016



Part of the [Physics Commons](#)



This work is licensed under a [Creative Commons Attribution-Noncommercial-No Derivative Works 4.0 License](#).

Recommended Citation

Gagné, Daniel David, "Well, It's About Time" (2016). *Senior Projects Spring 2016*. 338.
https://digitalcommons.bard.edu/senproj_s2016/338

This Open Access work is protected by copyright and/or related rights. It has been provided to you by Bard College's Stevenson Library with permission from the rights-holder(s). You are free to use this work in any way that is permitted by the copyright and related rights. For other uses you need to obtain permission from the rights-holder(s) directly, unless additional rights are indicated by a Creative Commons license in the record and/or on the work itself. For more information, please contact digitalcommons@bard.edu.

Well, It's About Time

A Senior Project submitted to
The Division of Science, Mathematics, and Computing
of
Bard College

by
Daniel Gagné

Annandale-on-Hudson, New York
May, 2016

Abstract

A clock can be thought of as anything that oscillates with a known and stable frequency, and is able to count its oscillations. Multivibrators are examples of such a device. Multivibrators resonate with a calculable resonant frequency, and this frequency is stable enough to be a standard by which time is kept. Though we assume clocks to be perfect timekeepers, however, this is not the case. Clocks not only have hardware imperfections, they are also susceptible to environmental noises—the combined effects of which leads to variations in timekeeping. This project will aim to explore and characterize the impact of environmental noises on the timekeeping ability of the multivibrator.

Contents

Abstract	1
Dedication	4
Acknowledgments	5
1 Prologue	6
2 Introduction	10
2.1 Clocks	11
3 Statistical Methods	13
3.1 Allan Variance	13
3.2 N-Cornered Hat Method	15
3.3 Frequency Counter	16
3.4 Signal Measurement	17
3.4.1 Nyquist Frequency	18
4 The Multivibrator	20
4.1 Operational Amplifiers	21
4.2 How Our Clock Works: A Derivation	22
4.2.1 Capacitors	23
4.2.2 Putting it all Together	24
4.3 The Frequency of the Multivibrator	25
5 Experiment Outline	27
5.1 Vibrations	28
5.2 Electromagnetic Waves	31

<i>Contents</i>	3
6 Results	33
6.1 Computer Reference Tests Without Added Environmental Noise	33
6.2 Vibration Experiment	37
6.3 Magnetic Field Test	40
6.4 Our clocks Compared to one Another	43
7 Conclusion	49
8 Appendix	51
8.1 Appendix A: The Temperature Controller	51
8.1.1 Building the Controller	52
8.2 Appendix B: The RLC Circuit	55
8.3 Appendix C: Miscellaneous Relativity	57
Bibliography	60

Dedication

I'd like to dedicate my project to my family. Without their support and love I wouldn't be the man I am today. I am forever grateful.

Acknowledgments

I was lucky enough to have an amazing support system throughout the process of this project. First and foremost, I'd like to thank my two project advisors Harold Haggard and Joshua Cooperman—I feel blessed to have been able to work alongside them. I'd like to thank my advisor Matthew Deady for his guidance and help throughout my four years here. I am forever in debt to my friends—it is thanks to them that these past four years have been the best of my life. I want to thank my family for being an amazing support system, and last, but certainly not least, I'd like to thank my girlfriend Carla Sanchez for her love and support throughout this process and these past two years.

1

Prologue

The beginning phase of this project was quite the rollercoaster. I knew from the beginning I wanted to do something astrophysics-y, I just wasn't sure what. My project advisor was supposed to be Professor Peter Skiff, but he ended up going on sabbatical in the fall of 2015. My advisor, Professor Matthew Deady, brought to my attention that a recently-hired postdoc (Joshua Cooperman) and one of the new professors in the department (Professor Harold Haggard) had caught wind of my interests and were interested in working with me. I felt relieved and excited, and so we met later that week to discuss topics.

We threw around ideas about looking into exoplanets, we read papers on black holes, and we pondered the existence of the universe. So many exciting avenues of exploration! Eventually crunch time was upon us, so it was time for me to think hard, and decide what I wanted to do my project on. More thought eventually kicked out exoplanets, the universe was going to be around for a while so I could tackle that problem later on (grad school maybe?)—but it was the paper on black holes which really captivated me.

You can't talk about black holes without talking about General Relativity, and you certainly can't talk about General Relativity without talking about spacetime curvature.

I was instantly taken with this, and after our discussion we had decided that my project would be to try to measure relativistic time dilation. I remember being gripped with excitement as I left Hal's office.

My excitement; however, quickly changed to skepticism. How on Earth were we supposed to measure these effects? The more I thought about this the more daunting the task seemed to be.

In the words of every child's favorite scientist, consider the following:

Special Relativistic time dilation is described by

$$t_a = t_b \sqrt{1 - \frac{v^2}{c^2}} \quad (1.0.1)$$

here v is the velocity of clock a with respect to clock b , and c is the speed of light. In this scenario, observer a is moving with velocity v relative to the ground and measures the passing of time as t_a , while observer b is stationary relative to the ground, and measures the passing of time as t_b .

Gravitational time dilation is described by

$$t_a = t_b \sqrt{1 - \frac{2GM}{(r+h)c^2}} \quad (1.0.2)$$

here G is the gravitational constant, M (for our purposes) is the mass of the Earth, r is the radius of the Earth, and h is height measured from the surface of the Earth. In this scenario, t_b is the passing of time of an observer at height h relative to observer a who measures t_a on the ground. You may have already noticed, but let me make it absolutely clear. Especially for every-day activity, these effects are *tiny*. For special relativistic time dilation to start to become relevant at all, you'd have to start moving at significant fractions of the speed of light (the speed of light being $3 \times 10^8 m/s$), and for gravitational time dilation to start to become relevant, you'd have to reach heights extending into outer-space. On average, I spend most of my time at a pretty steady altitude nowhere near outer-space, and I don't

normally reach speeds of more than even 100km/h. How on Earth were we to measure these effects?

Hal emailed me one day with some reading about an experiment carried out by Joseph C. Hafele (a physicist) and Richard E. Keating (an astronomer). They flew four cesium atomic clocks on a plane around the world (one trip eastward and one trip westward). They calculated beforehand the expected time difference the clocks should measure (between the clocks on the plane and a reference clock on the ground), and their results were pretty much spot on, thanks to the incredible precision of atomic clocks. What a brilliant way to measure these effects! Planes fly at a decent height, and go pretty fast, so with a long-enough trip and a precise-enough clock, these effects, though still small, should be measurable!

I mentioned I was flying to Europe for winter break, so one idea was to try to replicate their experiment. Some googling brought us to a webpage for a Microsemi chip scale atomic clock. The following meetings were spent thinking about Special and General Relativity—working through the theory and developing a basic understanding for how things worked. The purchase of three clocks was approved by the department, so it was on to figure out a way to bring the clocks on a commercial flight (something I was a bit nervous about), and to try to get a hold of a salesman from Microsemi. After tireless efforts by my advisors to try to get in touch with a sales rep, we found out that the clocks are manufactured per order request and take around six months to make and ship. This, on top of the fact that our calculations showed an excruciatingly tiny time-difference measurement (which wasn't large enough to rule out simple error), led that idea 'cross the river Styx.

Though we weren't able to stick with measuring relativistic effects, lots of interesting questions popped up in our discussions on relativity. A particular interesting topic of discussion was questioning whether a body on a rotating disc experienced general or special relativistic time dilation. The work we did and the questions we asked during the relativity

phase of this project can be found in Appendix C. Throughout all this discussion about time dilation and clocks keeping track of time at different rates, I found myself wondering if relativistic effects were the only effects that could cause clocks to keep incorrect time. This led us to question how clocks actually work, and we wondered how the "goodness" of clocks could be quantized. I use "goodness" as a purposefully vague word, because we weren't even sure if there was a quantitative way to measure how "good" a clock is—not to mention, we weren't quite sure what a clock being "good" even entailed.

We looked into how every-day clocks worked, and found out that quartz oscillators are usually used in clocks and watches—utilizing an operational amplifier feedback system. In my electronics class we had worked with operational amplifiers, and had built a circuit using one which vibrates at a resonant frequency. We connected the dots, and decided to carry out our timekeeping experiments on this "multivibrator" circuit.

2

Introduction

The notions of "past, present, and future" allow us to coordinate events, and the invention of clocks and timekeeping has let us coordinate events with great accuracy and precision. But how accurate and how precise? What's the best we can do when it comes to keeping track of time?

This may seem like a silly question. All one needs to do to find the time is to look at a clock, but the time shown to us isn't the whole story. Clocks usually only display the time down to the second, but this certainly isn't the smallest unit of time; rather, while the average wall clock reads the minute-accurate time for the full minute, there are seconds, milliseconds, microseconds, etc, that are ticking by which aren't shown to you. Imagine time as a tangible object you can keep cutting in half and splitting up into smaller and smaller pieces endlessly. Pieces so small it becomes difficult to comprehend what they even mean. So when I ask how accurate or precise we can be, I mean to ask what the smallest unit of time we can accurately measure and keep track of is.

Physicists in particular have been perplexed by this question ever since Albert Einstein developed his theories of relativity. The effects Einstein predicted were tiny when not

moving near relativistic speeds (near the speed of light), as we showed in equations 1.0.1 and 1.0.2, so the quest was on to build a clock that could make precise and accurate-enough measurements to measure even the tiniest of these effects encountered by every-day activity.

The most accurate clocks developed thus far are atomic clocks. There is a Strontium atomic clock in Boulder, Colorado that only loses or gains one second in 15 billion years—a huge improvement from the cesium atomic clocks (used to calculate UTC) that lose or gain a second in hundreds of millions of years. The strontium atomic clock is so accurate it's able to "detect" an altitude change of 2cm. This kind of accuracy is rare, though. Everyday clocks aren't able to "notice" or keep track of time dilation encountered by everyday activities. What, then, makes these atomic clocks so special? To answer these questions, we'll have to define what a clock really is, and figure out how they work.

2.1 Clocks

Ever wondered how clocks keep time? How they know how long a minute is? Or a second? The answer lies within the watch's circuitry. In most wristwatches, quartz crystals are used as the frequency standard—making use of the piezoelectric effect. Materials like quartz, when squeezed, will generate an electric signal. The reverse is also true, and this is what clocks utilize. If you send electricity through these piezoelectric materials, they will "squeeze" themselves. What's interesting about quartz, is that it will vibrate at the same "resonant" frequency no matter (to an extent) how you send electricity through it. A good analogy is to think of a piano. Pressing the middle C key will play a middle C no matter how hard or soft you hit the key, or how long you hold it for. The string the hammer hits will always vibrate at the same frequency. You can then, in essence, think of clocks as counters. The clock "knows" how many times the quartz is supposed to squeeze

itself per second, so it counts up to that number and then registers a second as having gone by.

It seems, then, that a clock is just anything that resonates at a somewhat-stable resonant frequency. As long as you know what that frequency is, it's easy to keep track of time—all you need to do is build a counter. The only reason some clocks are better than others is the choice of frequency standard. Most wristwatches, as I already said, use quartz, whereas atomic clocks use energy-level-jumping electrons as their standard (which happens to be much more stable, and therefore more precise). This is why that strontium clock is so amazing. It's frequency standard is so well-researched and stable, that even tiny variations in height make a noticeable difference in the clock's frequency standard.

3

Statistical Methods

3.1 Allan Variance

We’ve talked a good amount about frequency standards being stable and precise, but what does it mean for a frequency to be stable? What tool can we use to measure the stability of a signal? This is where Allan variance comes in. Allan variance is an important part of this experiment and will be the metric by which we measure the stability of our clocks’ frequencies. Allan variance is a measure of the frequency stability of an oscillator over a time interval τ . The equation for Allan variance is

$$\sigma_y^2(\tau) = \frac{1}{2} \overline{(\Delta y)^2}. \quad (3.1.1)$$

To measure the Allan variance of a clock, many frequency measurements have to be made with the same time interval τ between measurements (neglecting the amount of time it takes to make a measurement). In the equation, y represents the value of each of those frequency readings made. Δy represents the difference of all calculated y ’s—for instance, $\Delta y_1 = y_2 - y_1$, $\Delta y_2 = y_3 - y_2$, and so on (all the while being wary of the sign of each individual Δy , as some will be positive and some will be negative). To calculate Allan variance, we compute and square each Δy and then take their average—the bar over the

right hand side of 3.1.1 denotes this average. The average is then multiplied by a factor of a half, and we have our Allan variance. In Essence, Allan variance is a measure of the variance of all the measured y 's.

Since we are using an external counter to take frequency measurements, we are using an external reference clock. The counter we're going to use to perform these measurements uses the computer clock as a reference clock, and calculates the frequency of our circuit from there. So, when we measure the Allan variance with this method, we are comparing our circuits to the computer clock. We can, however, compare each of our circuits to one another by the clever equation

$$y_{ij} = \frac{\nu_{iR}}{\nu_{jR}} - 1. \quad (3.1.2)$$

This relation is saying that if we want to take clock i and compare it to clock j , we take the computer-measured frequencies of i (the computer y 's) and divide each y by the computer-measured frequency of clock j at that τ , then subtract 1. This gives us a new set of y 's per τ , and we can find the Allan variance of these y 's to get the variance of i with respect to j . Remember, we're assuming in our experiment that the three clocks we're using are exactly similar, but this will be a good way to actually figure out how comparable our clocks are. As the Allan variance of one clock with respect to another approaches zero, the clocks become more and more comparable—with an Allan variance of 0 denoting "perfectly" comparable clocks—so we'll be looking to see very small numbers when computing Allan variances by this method. This method is important to understand for another reason too. When scientists engineer clocks—trying to make a better clock than already exists—they have to figure out a way to quantify just how good the newly-developed clock is. Because, however, this newly-developed clock is assumed to be better than all other existing clocks, the scientist can't compute the Allan variance of the new clock to any existing one—that quantity wouldn't make sense! Instead, copies of this "better" clock have to be made, and

then compared to one another. So, by taking the variances of our clocks with respect to one another, we can explore this method and, in essence, figure out just how good the clocks are as if they were assumed to be the best in existence.

For our experiment, we will choose τ 's equal to (measured in seconds) $\{.5, 1, 2, 5, 8, 10, 15, 20, 30, 40, 50, 60, 80, 90, 120 \text{ and } 240\}$. After we have calculated the Allan variances for each τ , we will then make a plot of Allan variance vs τ to see how the Allan variance changes as τ increases. We expect to see the Allan variance increase as τ increases, because a longer sampling time means a longer experiment—thus giving the clocks more time to vary and drift.

3.2 N-Cornered Hat Method

The N-Cornered Hat Method is a method by which one can extract the Allan variance of a clock after having measured its Allan Variance against several other similar references. This will be a necessary tool, because as explained in the previous section, though we are going to measure the Allan variance for the clocks with respect to the computer, we're also going to do that measurement for each clock with respect to one another. The N-Cornered Hat Method is a generalization of the Three-Cornered Hat Method, in which only three clocks are compared to one another. The way to extract a clock's Allan Variance using the N-Cornered Hat Method is

$$\sigma_i^2(\tau) = \frac{1}{N-2} \left[\sum_{j=1}^N \sigma_{ij}^2(\tau) - B \right] \quad (3.2.1)$$

where N is the number of hats, and σ_{ij}^2 is the Allan Variance of clock i with respect to clock j , and

$$B = \frac{1}{2(N-1)} \sum_{k=1}^N \left[\sum_{j=1}^N \sigma_{kj}^2 \right] \quad (3.2.2)$$

can be thought of as a correction term. It represents subtracting the measured Allan Variances of all the other clocks (not your test clock you're interested in) compared to

each other. This project will be using three clocks, so in 3.2.1 and 3.2.2 we can let $N = 3$. The equations then boil down to the following

$$\begin{aligned}\sigma_A^2(\tau) &= \frac{1}{2}(\sigma_{AC}^2(\tau) + \sigma_{AB}^2(\tau) - \sigma_{BC}^2(\tau)) \\ \sigma_B^2(\tau) &= \frac{1}{2}(\sigma_{BA}^2(\tau) + \sigma_{BC}^2(\tau) - \sigma_{AC}^2(\tau)) \\ \sigma_C^2(\tau) &= \frac{1}{2}(\sigma_{CA}^2(\tau) + \sigma_{CB}^2(\tau) - \sigma_{AB}^2(\tau))\end{aligned}$$

where A , B , and C are the names of the three clocks we're using. Up to here I've made sure to be explicit about what is acting as a reference for our clocks, because in order to measure Allan variance you need to have a reference to compare your test clock to. You may have noticed, however, that on the left-hand side of the above equation, the reference clock isn't labeled. We have extracted each clock's Allan variance from being tied down to one clock in particular, but that doesn't mean this Allan variance is without reference; rather, it means that the clock is now using the *ensemble* of clocks as a reference. The frequency of this "ensemble reference" has to be somewhat guessed—and a good measure of this, because our clocks are so comparable, is to simply take the average frequency at which all the clocks are functioning. The reason this is so important will be made clear in Chapter 6.

3.3 Frequency Counter

In order to measure the frequency of the multivibrators, we're going to use National Instruments' data acquisition (DAQ) board and National Instruments' software *LabView*. Using *LabView*, I built a program which takes the output signal generated by the multivibrator as an input. This signal is then sent through the *LabView* program, which is set up to take a frequency measurement after the elapse of each time interval τ (the same from the Allan variance section), and to output a frequency reading.

This seems simple enough, but what’s really going on inside the DAQ board is quite interesting. Just how does it measure frequency? Obviously it has to somehow count the peaks per second that run through it, but how does this happen? It turns out, that the counter isn’t continuously measuring the frequency when it’s triggered; rather it measures the signal voltage at a certain *sampling* frequency. To make sense of this, imagine you’re standing at the beach watching waves crash ashore. If you wanted to count how many waves per time interval came to shore, you could just stand there and watch and count—this would be continuous sampling (which the counter doesn’t do). Instead, if you were to count waves like the DAQ board, you wouldn’t keep your eyes open to watch the waves throughout the entire time interval, you’d be continuously opening and closing your eyes every so often—and only counting the waves you see. As you can imagine, depending on how long of an interval you take between closing and opening your eyes, (this is the sampling frequency) you may miss some waves in between, and your count would be off. This is why it’s important to specify a large enough sampling frequency (relative to the input signal coming from the oscillator) for the DAQ board to use. We don’t want the counter to miss any waves. We can make sure we’ve specified a large-enough sampling frequency by following the guidelines in the following section.

3.4 Signal Measurement

According to Fourier’s theorem, any periodic and reasonably continuous function can be written as a sum of sinusoids. Mathematically, this function can be written as

$$f(x) = \frac{A_0}{2} + \sum_{n=1}^N A_n \sin\left(\frac{2\pi nx}{P} + \Phi_n\right)$$

for integer N ’s with $N \in [0, \infty)$. Here, P is the period of $f(x)$, A_0 is a real number, and Φ_n is a phase shift, which may or may not be present (note $0 \leq \Phi_n \leq 2\pi$). As explained earlier, our multivibrator is going to be outputting a periodic signal, and the counter is

going to sample the frequency at a sampling frequency—which can be assumed to be sinusoidal. The caveat is this: If there is an infinite sum of signals that can be equal to the one signal we’re looking at, how can we be sure the counter is digitally rendering (and measuring) the correct signal?

3.4.1 Nyquist Frequency

If the counter samples at equally-spaced times (with an equal amount of time between samples) then the sampling frequency is

$$f_s = \frac{1}{\Delta t} \quad (3.4.1)$$

with Δt being the time between samples. Because our counter samples like this (and not continuously), there is then a maximum frequency the counter can detect for whatever sampling frequency it’s sampling at. This is due to the fact that when sampling, we need to, at the very least: first sample the signal’s positive peak, then the negative peak at the second sample, the positive peak again at the third sample, and so on. Therefore, when you’re sampling with $f_s = \frac{1}{\Delta t}$, the maximum frequency the counter can measure has a period of $T = 2\Delta t$. This limitation is known as the Nyquist Frequency with

$$f_{\text{nyquist}} = \frac{1}{2\Delta t} = \frac{f_s}{2} \quad (3.4.2)$$

You might be wondering what would happen if we sampled a signal too infrequently or frequently. Well, you can never sample too frequently (to an extent)—the higher f_s is compared to f the more accurate the reading will be (folding will appear if your sampling frequency $f_s \gg f$ but we don’t have to worry about that). A problem does arise; however, when we sample too infrequently. When the counter measures a frequency, it takes samples of the signal, and uses these samples to digitally recreate the signal. Then, the counter measures the frequency of the digitally rendered signal not the original signal itself. When sampling a frequency f at sampling frequency f_s , the resulting digital wave-

form (constructed by the counter) will only have frequency f if $f \leq f_{nyquist}$ or $f \leq \frac{f_s}{2}$. If, instead, the frequency of the signal $f > \frac{f_s}{2}$ then the resulting digitized waveform will have frequency $f_{alias} \neq f$ given by

$$f_{alias} = |f - nf_s| \quad \text{with} \quad n = 1, 2, 3, \dots \quad (3.4.3)$$

here f_{alias} is within $0 \leq f_{alias} \leq f_{nyquist}$. This is called the alias frequency because it can be thought of as a false identification—in the same way a person under a fake name is using an alias.

4

The Multivibrator

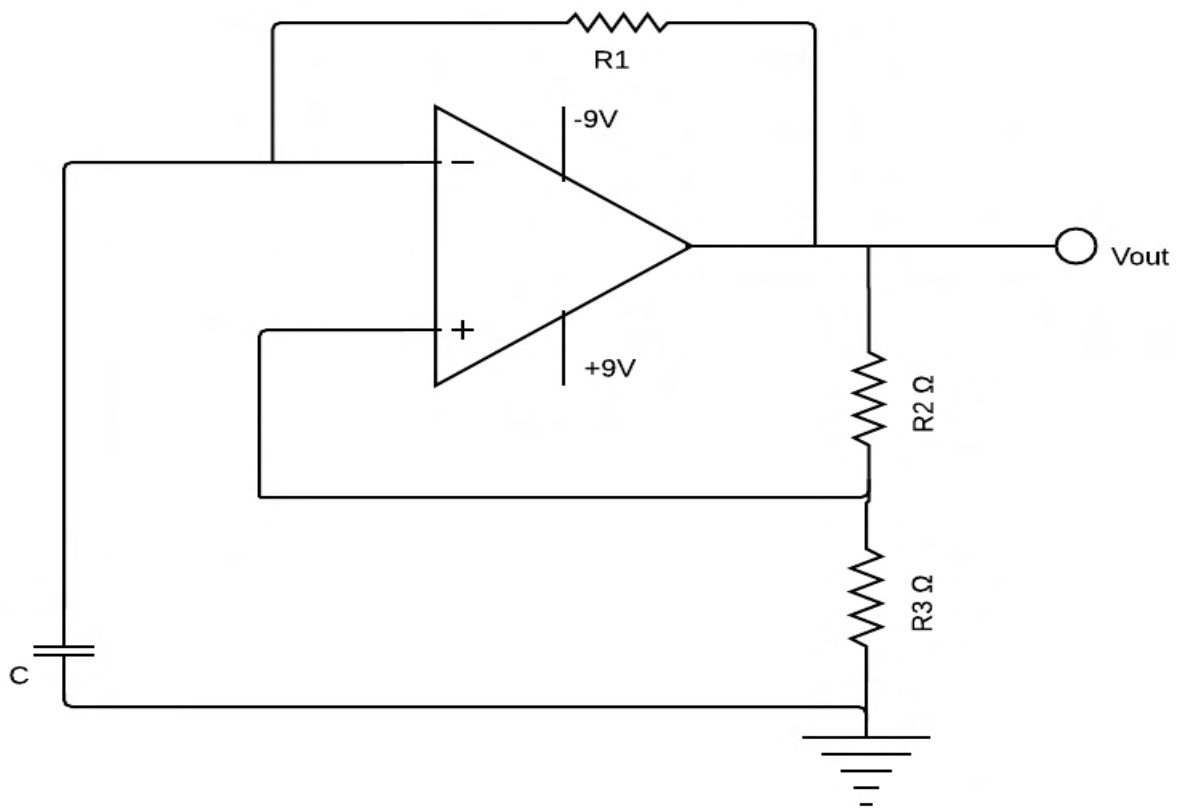


Figure 4.0.1. Multivibrator circuit diagram

The clock I'm going to use for my project will be an operational amplifier feedback circuit. The circuit diagram for my circuit is on the previous page. The circuit will be comprised of three resistors: $R1 = 100\Omega$, and $R2 = R3 = 1K\Omega$, a capacitor: $C = 100\mu F$, and an operational amplifier. A major component in understanding why this circuit behaves the way it does will be to understand operational amplifiers.

4.1 Operational Amplifiers

Operational amplifiers (op amps) are high-gain voltage amplifiers with differential input. Op amps are one of, if not the most, widely-used circuit elements. They make changing around and playing with output voltages easy. For instance, op amps can be used to take the integral and derivative of an input signal. The figure to the right shows how an op amp is represented on a circuit drawing. The inputs are V_+ and V_- (non-inverting and inverting respectively), V_{out} is the output voltage, and V_{S+} and V_{S-} are the "powering voltages"

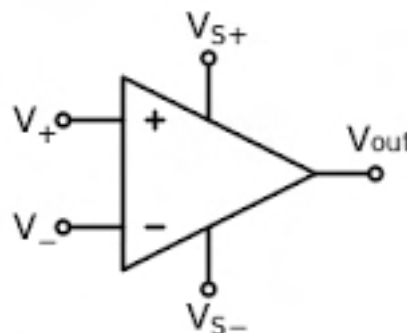


Figure 4.1.1. op amp circuit diagram [8]

of the amplifier (they're just DC voltages which power the inner-workings of the op amp).

There are two "golden rules" which are essential to understanding how op amps function in a circuit:

1. The input impedance of the $+/-$ inputs is infinite. The output impedance is 0.
2. In a circuit with negative feedback, the op amp will try to adjust its voltage output so that $V_+ = V_-$.

Rule 1 means that the $+/-$ inputs draw no current, and the output is a perfect source of current and voltage (there is no impedance or hindrance in outputting voltage or current).

Rule 2 is pretty self-explanatory. The output will act however it needs to make $V_+ = V_-$.

This is what makes it easy to manipulate the output voltage. If you know how the amplifier will act to make $V_+ = V_-$, then you can "create" whatever output voltage you desire by setting up the circuit in different ways.

4.2 How Our Clock Works: A Derivation

The easiest way to figure out how any circuit works, is usually to start where the input voltage is, and then traverse the circuit making note of potential drops. Our circuit, however, does not have an external alternating V_{in} ! Somehow this circuit doesn't need an external alternating input voltage to produce a periodic output voltage. Without being able to start at an input voltage, figuring out how this circuit works is not so straightforward. We can start, however, by noticing two important characteristics of this circuit:

1. The output voltage feeds back into the "+" and "-" inputs, and
2. The non-inverting input (+) is held by the $R_1 R_2$ voltage divider at a fraction of V_{out} .

The importance of (1), is that V_{out} will then have an influence on the voltages at the + and - ports of the op amp, given by $V_{out} = A(V_+ - V_-)$, where A is called the *gain* of the op amp, and is an incredibly large number (on the order of 1×10^6 sometimes). (2) shows us a direct relation between V_{out} and V_+ . A voltage divider acts to divide voltage between two points—proportional to the strength of the resistors. In this configuration, the voltage divider acts to establish the relation

$$V_+ = V_{out} \frac{R_1}{R_1 + R_2}. \quad (4.2.1)$$

From here on out we will refer to this proportionality factor as β , that is, $\beta = \frac{R_1}{R_1 + R_2}$. The multivibrator we used has resistors $R_1 = R_2 = 1K\Omega$, and therefore a β proportionality factor of $\beta = \frac{1}{2}$. This means that V_+ will be held at $\frac{V_{out}}{2}$.

These two observations, however, aren't enough to figure out how this circuit ticks. The second key to understanding why this circuit does what it does is to understand the role of the capacitor.

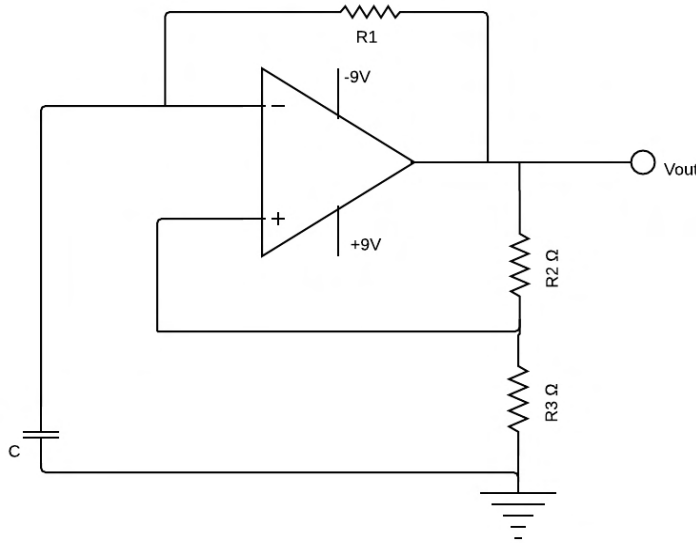


Figure 4.2.1. Our Clock

4.2.1 Capacitors

As the circuit diagram to the left shows, the capacitor is, through R_1 , connected to the output of the op amp—and because of this, it experiences a voltage drop across its plates. The capacitor we're using is a parallel plate capacitor, and when a voltage is applied across it, charges will start to accumulate. Then, because all

the negative charges go to one plate and the positive charges to the other, as more and more charges accumulate, it becomes harder for subsequent charges to reach the plate, because the plate will act to repel more charges (like charges repel). Due to this phenomena, a capacitor can only store a certain amount of charge given by the relation

$$CV = Q \quad (4.2.2)$$

where C is the *capacitance* of the capacitor, V the voltage across the plates, and Q the amount of charge stored.

Because the force between charges is not linear, the capacitor does not charge up linearly. This means, that as more charge accumulates, the capacitor charges at a slower compared to what it was just a second ago. Capacitors discharge in a similar manner as they charge, so if it were possible to get the capacitor to charge and discharge, there actually *would* be an AC voltage located in the circuit. It turns out that this circuit is able to cause this charging and discharging on its own, and I will explain how in a moment. The voltage across a charging and discharging capacitor is shown in Figure 4.2.2, along with a picture demonstrating the accumulation of charge on its plates (note that "dielectric" simply refers to material stored between the plates to increase the capacitance).

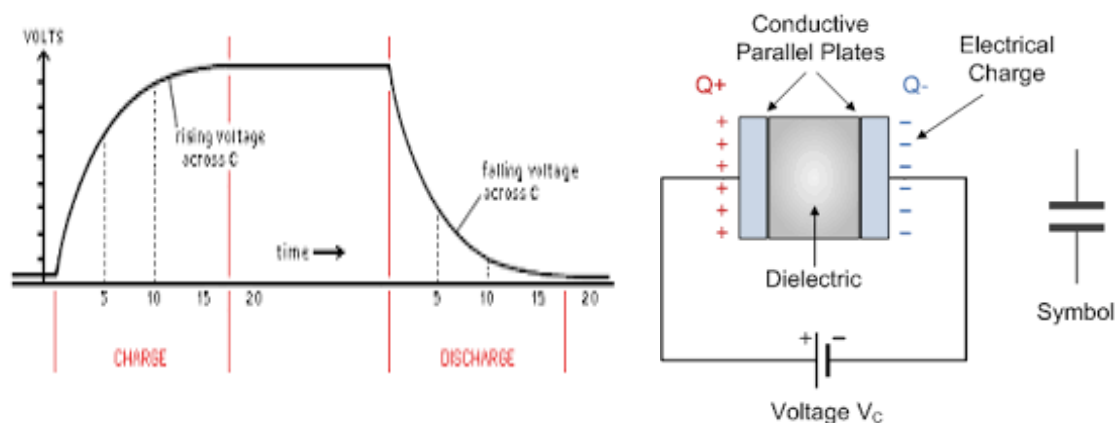


Figure 4.2.2. Capacitor charging and discharging (left) [6] and parallel plate capacitor (right) [7]

4.2.2 Putting it all Together

As you can tell from Figure 4.2.2, the charging and discharging of the capacitor doesn't resemble a square wave at all, so how and why does the op amp generate one? Let's imagine a scenario in our circuit where the capacitor is fully discharged. When the capacitor is fully discharged, the op amp is saturated—meaning V_{out} is very close to V_{S+} (the powering voltage). When this occurs, the capacitor will start to charge up, due to the output voltage through resistor R_1 . The capacitor will asymptotically approach attaining V_{S+} between

its plates—at a rate determined by the RC time constant, τ_C , given by

$$\tau_C = R_1 C. \quad (4.2.3)$$

The capacitor, then, in a way, decides what voltage is at the V_- port. So, as the capacitor charges, there will reach a point where V_- is equal to or greater than V_+ . When this happens, we recall our relation $V_{out} = A(V_+ - V_-)$, and note that V_{out} switches sign! When V_{out} switches sign, there is suddenly a negative voltage across the capacitor. When this sudden voltage change occurs, the capacitor will discharge at a rate again determined by the time constant τ_C . When the capacitor has discharged enough, we will reach a point where $V_+ > V_-$, V_{out} switches sign, and the process repeats. This is where the square wave comes from. The op amp output is really just oscillating between V_{S+} and V_{S-} .

Now, you might be wondering whether the output is really a square wave. Indeed, $V_{out} = A(V_+ - V_-)$ and is only dependent on V_+ and V_- , and you'd be right in thinking V_{out} might change slightly due to the changing quantity of $(V_+ - V_-)$. This, however, is where the gain comes in. As mentioned before, the gain A of an op amp is so large, that even though $(V_+ - V_-)$ is changing a bit, this difference is, in the end, completely drowned out. The large gain makes the op amp ignore the changing quantity of $(V_+ - V_-)$ and just output the largest voltage it can—which as said before, is right around $\pm V_S$.

4.3 The Frequency of the Multivibrator

Pretending our capacitor is the AC input signal, we can start at the capacitor and traverse the circuit using Kirchoff's Loop Laws. We traverse in such a way that:

$$V_C(t) = -V_+ + (V_{Sat} + V_+)(1 - e^{\frac{-t}{\tau}})$$

where V_{Sat} is the saturation voltage of the op amp—which in turn is just the magnitude of the output voltage (note that though the op amp wants to output the same magnitude

square wave as the voltages powering the op amp, $\pm V_S$, due to imperfections in the op amp, it's only able to get close to $\pm V_S$, and this voltage is called V_{Sat}). Let's let $V_C(t) = 0$ when $t = 0$. The time, then, that it takes to reach $+V_+$ is

$$+V_+ = -V_+ + (V_{Sat} + V_+)(1 - e^{\frac{-t}{\tau}})$$

then adding V_+ and dividing through by $(V_{Sat} + V_+)$

$$\frac{2V_+}{V_{Sat} + V_+} = (1 - e^{\frac{-t}{\tau}}).$$

Next recall that the two resistors R_2 and R_3 act as a voltage divider with factor β —and allows us to rewrite the left-hand side of the equation as

$$\frac{2\beta}{1 + \beta} = (1 - e^{\frac{-t}{\tau}})$$

which leads to

$$e^{\frac{-t}{\tau}} = \frac{1 - \beta}{1 + \beta}.$$

We can then take the natural log and solve for t

$$t = \tau \ln\left[\frac{1 + \beta}{1 - \beta}\right].$$

This gives the time for the capacitor to charge up, so if we want the whole period, we have to multiply by a factor of 2 and we get (also substituting in RC for τ)

$$T = 2RC \ln\left[\frac{1 + \beta}{1 - \beta}\right].$$

I noted earlier that our factor of β was a half, so we can simplify this down even more to

$$T = 2RC \ln[3]. \tag{4.3.1}$$

Plugging in numbers and taking note of the relation $f = \frac{1}{T}$, we see our clock should have a frequency of about 40Hz.

5

Experiment Outline

This experiment will investigate the Allan variances of three "similar" clocks with reference to a computer clock. We will do this in an uncontrolled setting (just in the lab), and we will also run tests with different environmental simulations—i.e. we're going to test the effect vibrations and electromagnetic waves have on the stability of our multivibrators (we were also going to try and measure temperature dependence, but we ran into difficulties—explained in Appendix A). On the following page I've pictured multivibrators B and C on the breadboard in the setting with no intentionally-added environmental noises. The breadboard is able to generate a +15V and -15V DC signal, and this is what was used to power the op amps. The setup is pictured at the top of the following page.

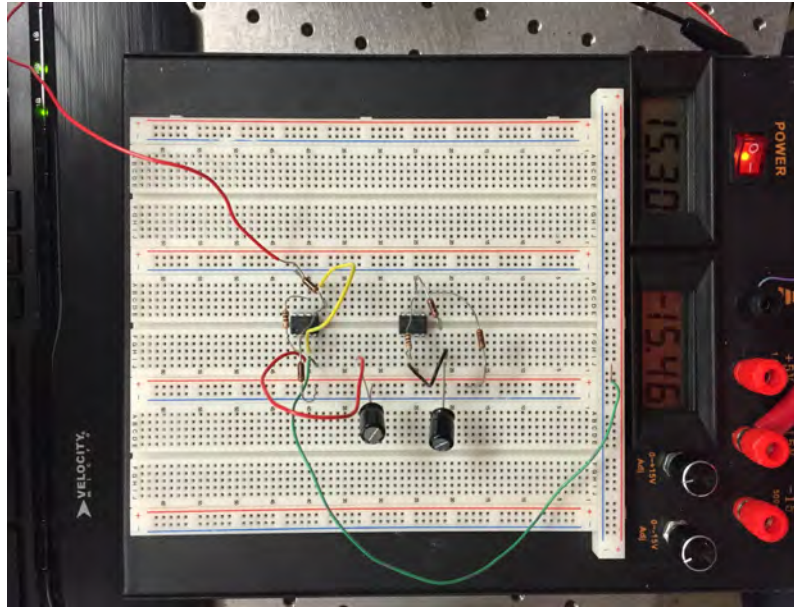


Figure 5.0.1. Clocks B and C in the uncontrolled setting

The reasoning behind our choice for environmental factors comes from what this experiment was originally going to be. As explained earlier, I was going to try and use atomic clocks on a plane to measure time dilation. Being on a plane, the clocks would encounter large accelerations and vibrations due to the plane's movement, they would be susceptible to any electromagnetic fields generated by the plane or technology onboard, and because the cabin is temperature controlled, the clocks would experience a somewhat-stable temperature environment. Though the multivibrators we're using might respond to these environmental factors differently, we are setting up these tests to experiment with ways to test these environmental factors, and to see which factors may contribute the most noise.

5.1 Vibrations

External motion, such as vibrations or walking with the oscillator may cause slightly erratic frequency of oscillation. To test how our oscillators are affected by external motion, we're

going to mount an oscillator on a speaker. The speaker will be powered by a function generator—allowing us to choose the frequency of vibration. We will take one of our multivibrators and subject it to different frequencies of vibration, and compare the Allan variances of the different vibration frequencies to discern whether or not vibrations have a notable impact on the clock’s frequency stability. Figure 5.1.1 shows the set-up for this part of our experiment. The breadboard is taped down on a metal sheet (to prevent it from flying off when the speaker is turned on), the green wires supply power to the op amp, and the red alligator clip brings the output voltage to the counter to make frequency measurements.

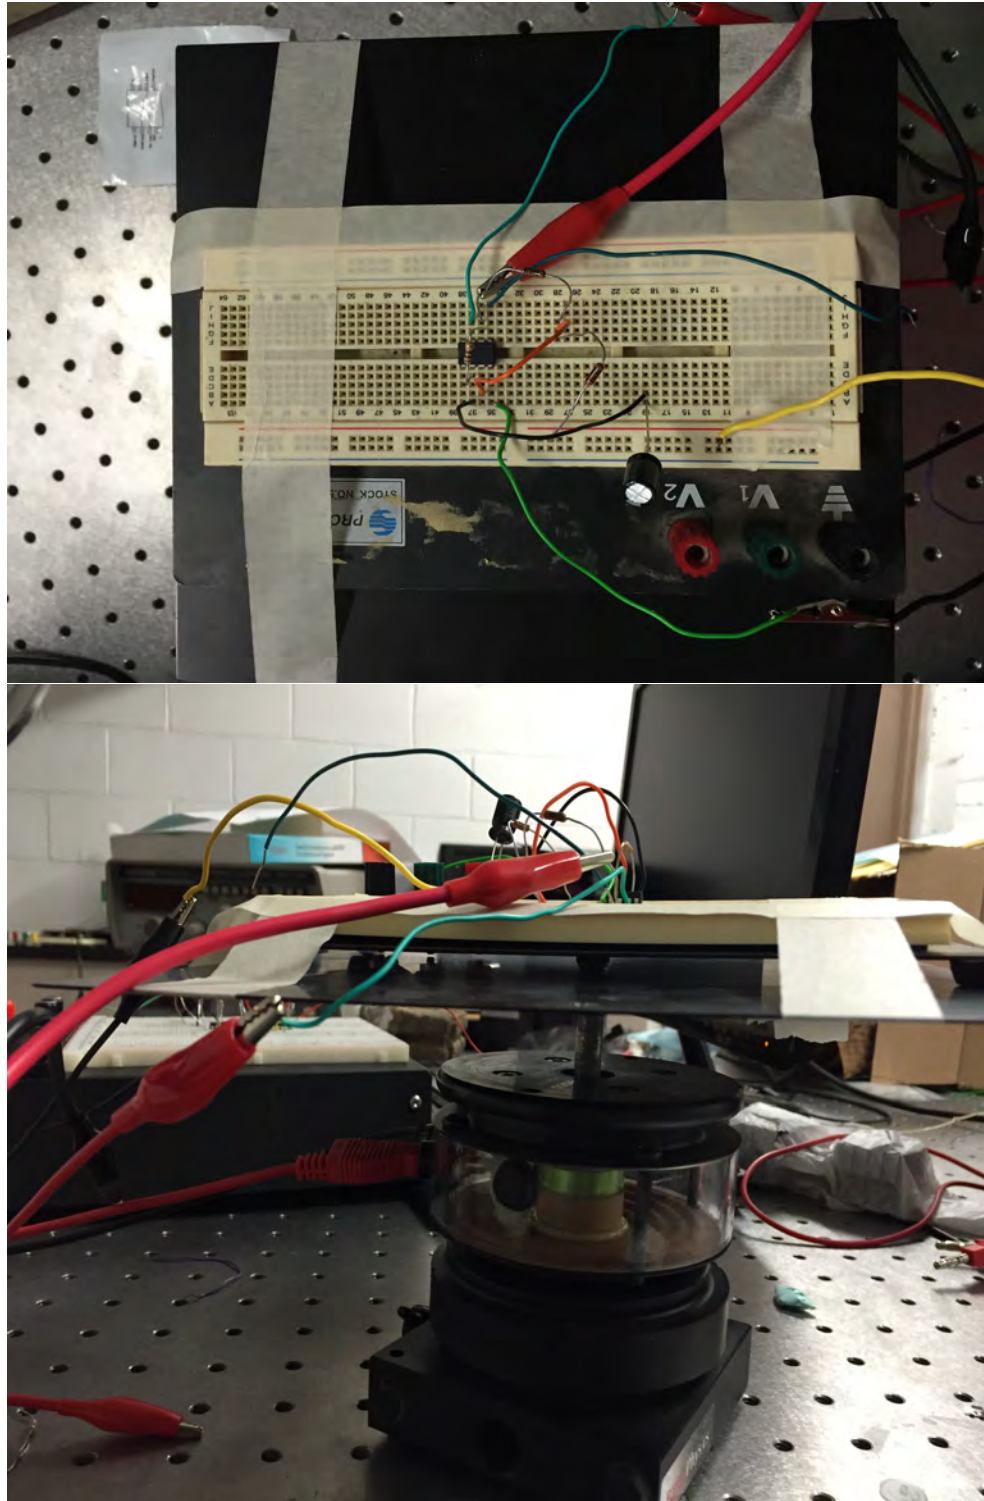


Figure 5.1.1. Vibration Test Setup. Top view (top) and side view (bottom)

5.2 Electromagnetic Waves

Since the output signal of the multivibrator is an electrical current, it stands to reason that a strong enough electromagnetic field could disrupt this signal and mess with the frequency the multivibrator outputs. To try and measure this, we will generate an electromagnetic field, put one of our multivibrators in it, and see how it functions compared to when it wasn't in any generated field.

To generate a magnetic field I created an electromagnet using a large solenoid, a large magnet, and a function generator. I connected leads from the function generator and plugged them into the solenoid. I then placed the large magnet inside the solenoid. When the function generator is turned on, an AC current is created through the solenoid, and current flowing in this type of geometry creates an electromagnetic field through the solenoid's center. The large magnet in the middle simply acts to strengthen the field. The strength of electromagnetic fields dies off at a rate of $\frac{1}{r^2}$ (where r is radial distance from the source), so placed the electromagnet as close to the multivibrator as I could to have it immersed in a strong field. The figure directly below shows the magnetic field due to a solenoid, and on the next page I've pictured the experimental setup.

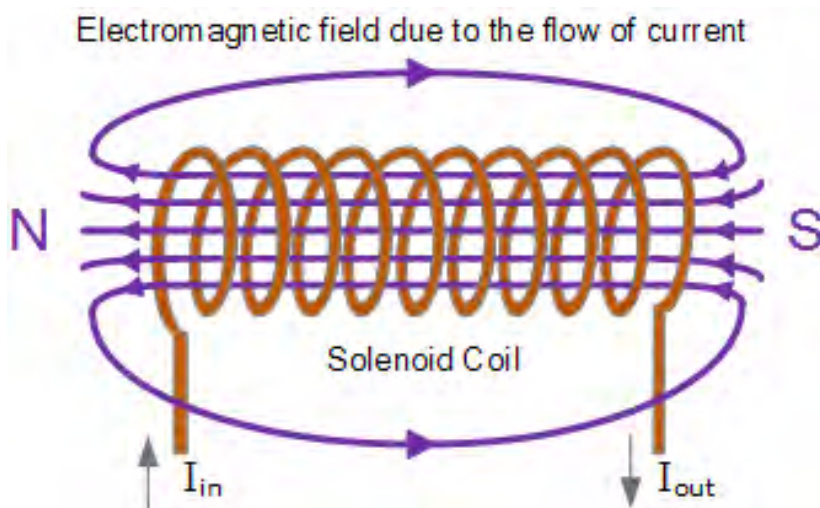


Figure 5.2.1. Magnetic field due to Solenoid [4]

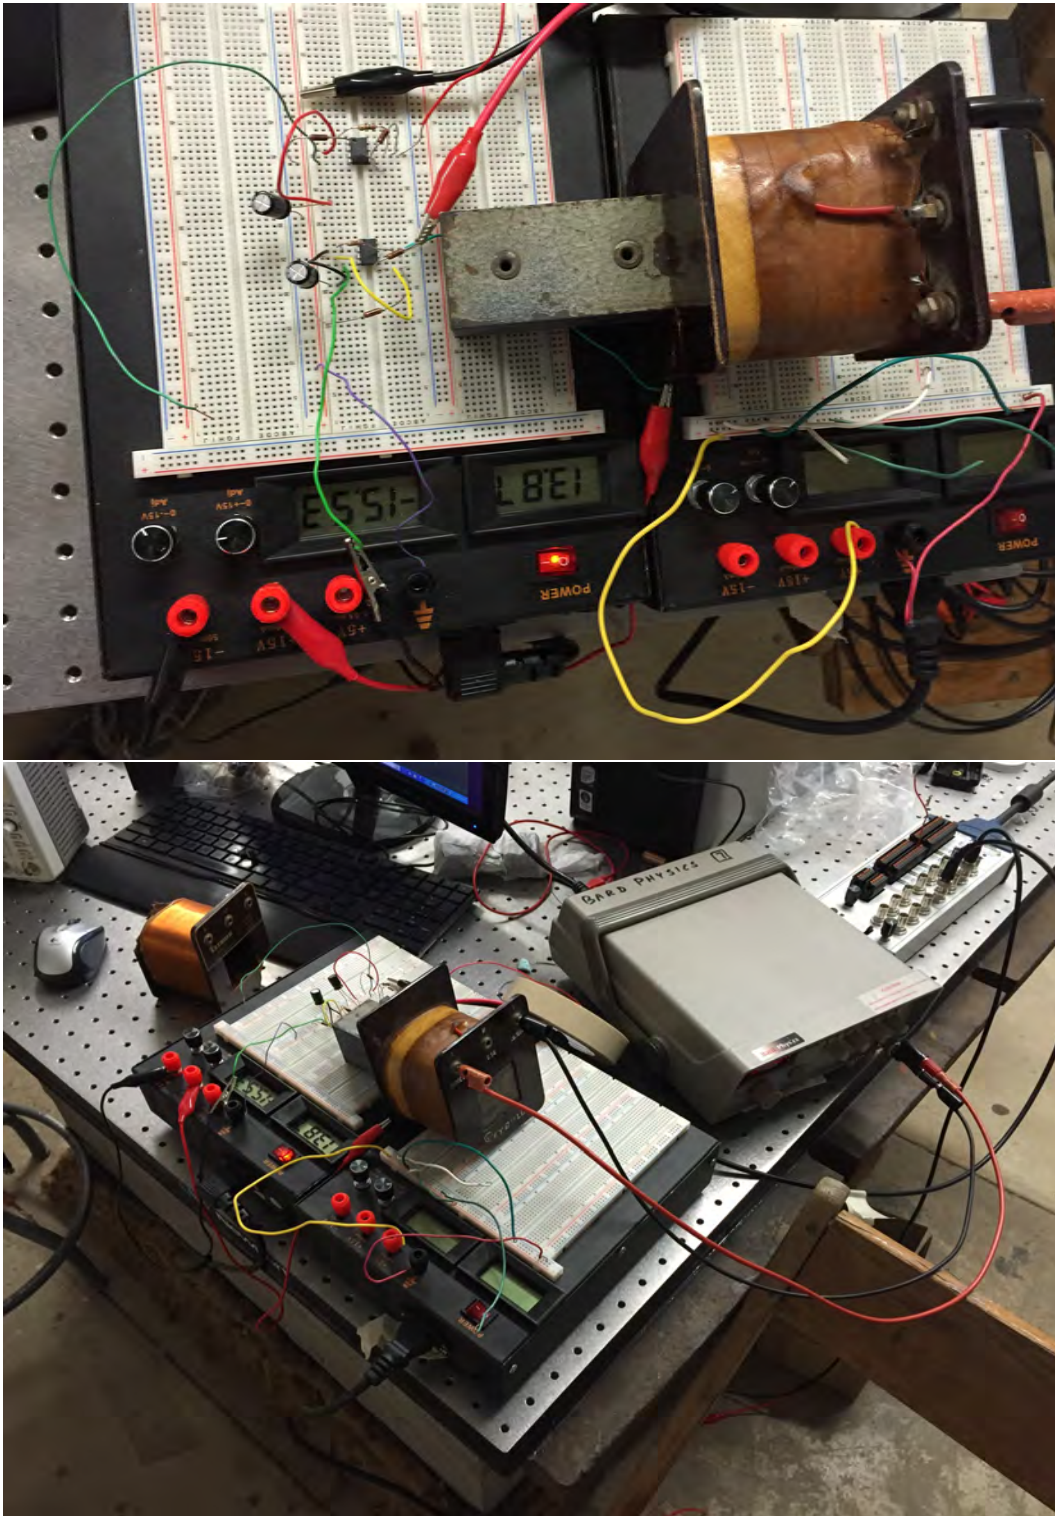


Figure 5.2.2. magnetic field test setup

6

Results

6.1 Computer Reference Tests Without Added Environmental Noise

As stated earlier, the τ 's we measured were (in seconds) $\{.5, 1, 2, 5, 8, 10, 15, 20, 30, 40, 50, 60, 80, 90, 120 \text{ and } 240\}$. We took 34 frequency measurements per τ for each clock, and with mathematica, calculated the Allan variances of the clocks with reference to the computer clock.

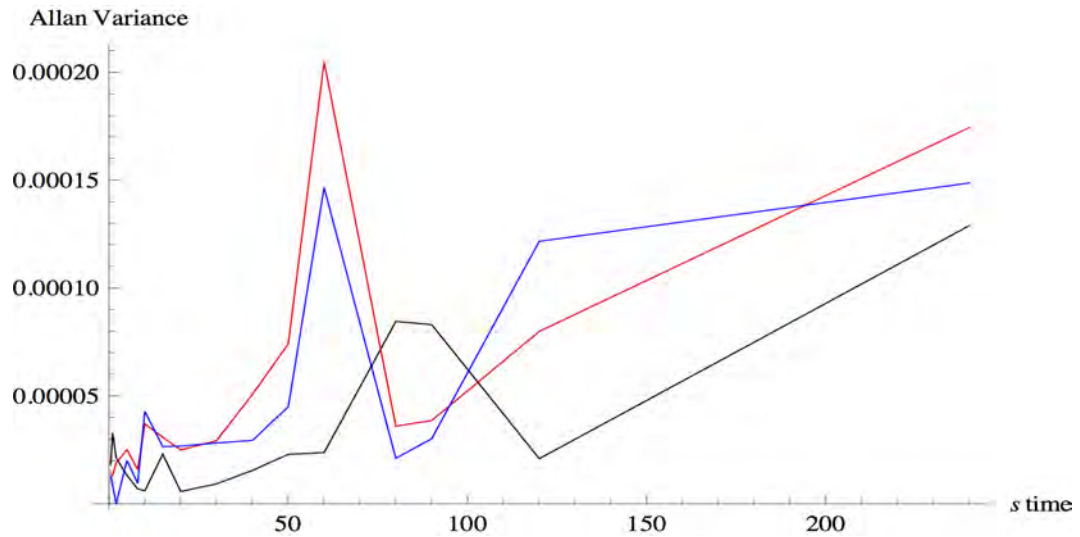


Figure 6.1.1. Allan variance of Clocks A (red), B (blue), and C (black)

It's quite noticeable how similar Clocks A and B are. Their two lines follow each other rather closely. Clock C differs slightly from A and B, however. Up until $\tau = 80$, Clock C is consistently below A and B—meaning it's performing better than those clocks at those smaller τ 's. Clock C jumps above A and B at 80s and 90s but then sharply falls back below them for the rest of the plot. Though Clock C varies a bit from A and B, it's quite safe to say that the clocks are similar to one another. These slight deviations from one clock to another are most likely due to imperfections in circuit elements or perhaps other noises that we (due to lack of time) were not able to characterize.

The largest spike we see comes from Clock A, which reaches upwards of 2×10^{-4} . With exception of $\tau = 240$, none of the other variances really come close to reaching this number. Though I can't say for sure, judging by the sharp increase of the lines coming from 120s to 240s, it appears as if for longer τ 's the variance will continue to trend towards increasing. I can say this with some small confidence because this change in τ is the largest of any of the other ones—so it's easy to imagine this increasing nature might continue, and, though the lines jump around a lot, they certainly trend towards increasing Allan variance with increasing τ .

We only had enough time to run each clock once without an environmental control. Because, however, we're essentially using exactly similar clocks, we can be clever. We can treat running these three clocks once as running one of our clocks three times. This allows us to, then, take our analysis one step further—using fractional differences. We can describe a fractional difference as the following

$$(\text{Fractional Difference})_{ij} = \frac{\sigma_j^2 - \sigma_i^2}{\sigma_i^2},$$

where $\text{Fractional Difference}_{ij}$ is the fractional difference of clock i with respect to clock j . This type of number will allow us to have a quantitative sense of how clocks A, B, and C differ from one another.

Clocks A and B are the only two clocks we subjected to environmental noise (we were going to use Clock C for temperature), so if we compute the fractional differences of A vs. B, and A vs. C, along with B vs. A, and B vs. C, we can get a sense of how the fractional differences of three comparable clocks should look. When we, then, compute fractional differences later on with respect to the environmental noises, we can compare those results to the fractional differences *without* added noises, and see whether or not the fractional differences remained similar or different. If the plots look wildly different, we may be able to conclude that the environmental noise contributed to this difference. Below are the two plots of fractional differences for A with references B and C, and for B with references A and C. In each case, I also plotted the point-average. It's also important to note that the graphs are really the graphs of the absolute value of the fractional differences. The main point of looking at these is to see what types of magnitudes we can expect—so I used the absolute value to prevent being distracted by the sign of the fractional difference.

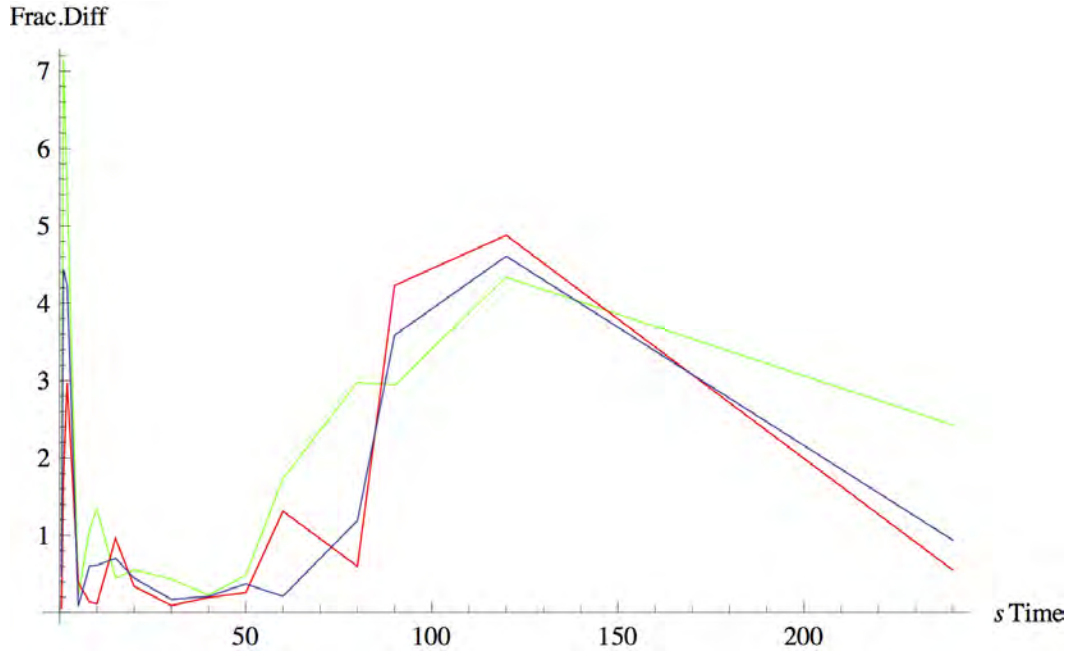


Figure 6.1.2. Fractional difference plot of BA (green) BC (red) and the point-average (blue)

This plot is a little crazy. There's a huge spike in the BA line at around 1s, then from about 8s to 80s things are pretty steadily at or around 1 to 2 in magnitude, but at 90s there's another spike in amplitude, followed by a steep drop at 240s.

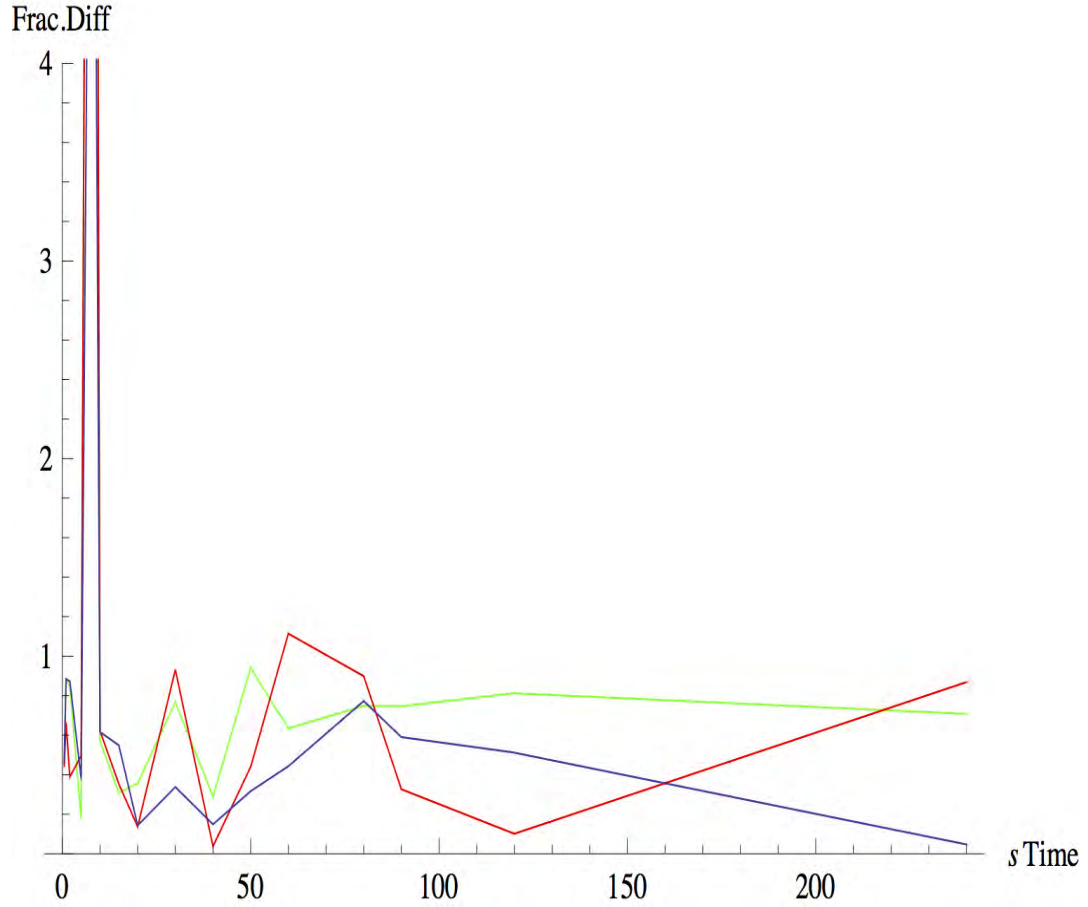


Figure 6.1.3. Fractional difference plot of AB (green), AC (red) and the point-average (blue)

This plot is a little easier to make sense of. There's one huge spike in the plot (which is cut-off now to get a better sense what's happening at the lower amplitudes), but the rest of the graph remains below 1.25 in magnitude. So, when looking at the fractional differences for the environmental noise clock A experienced, we should expect to see relatively small values. Below is the same plot but without the spike at 8s cut off.

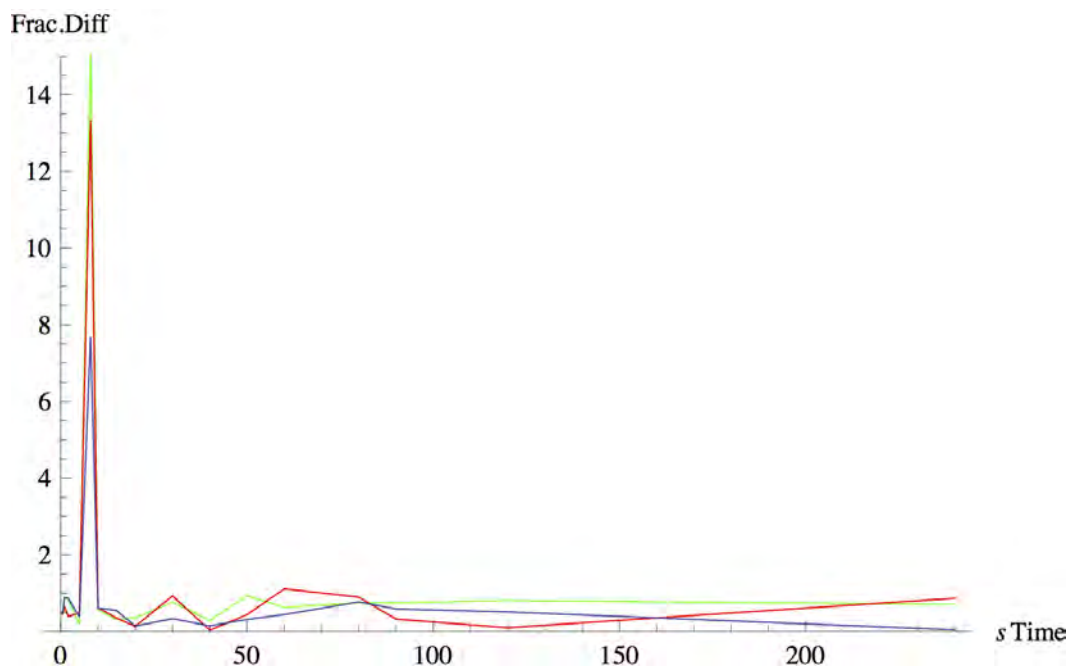


Figure 6.1.4. Fractional difference plot of A with spike shown

6.2 Vibration Experiment

For this experiment, we took clock A and subjected it to vibrations generated by a speaker. The vibrations were square waves, and were of frequencies: 17.8Hz, 29.8Hz, and 130.8Hz. Our reasoning behind choosing these frequencies really has to do with what frequency the vibrator is vibrating at. We wanted to see what happened when we jostled the vibrator right on its resonant frequency (29.8Hz), and to cover all our bases, we vibrated it at a frequency somewhat close to resonance (17.8Hz) and at a frequency far from resonance (130.8Hz). At the outset of this experiment, I didn't expect the vibrations to have much of an effect on the signal generated by the multivibrator, and if there were any effect, I assumed it'd lead to larger Allan variances—denigrating the stability of the clock. Our results showed that, to an extent.

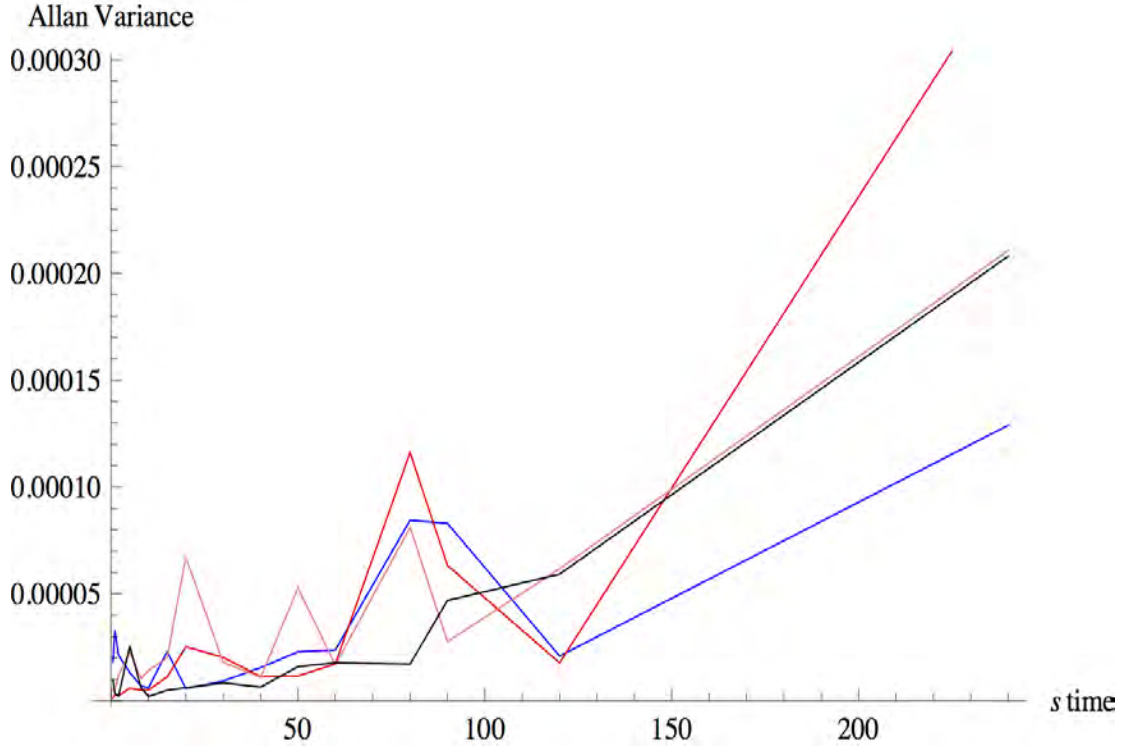


Figure 6.2.1. Allan variances of the vibration test

Shown above is the graph of clock A's Allan variance with respect to the computer clock without vibrations (this'll be called the control clock), and with all three vibration tests. The blue line is the control clock, the red line is Clock A being vibrated at 29.8Hz, the pink one is clock A being vibrated at 17.8Hz, and the black line is Clock A being vibrated at 130.8Hz. Near small τ 's, it's quite clear that vibrations don't really have any effect on the clock—the three lines bounce erratically in a similar manner to the control. The major difference can be seen at $\tau = 240s$. There is a huge spike in all three vibration tests, which leaves a huge gap between them and the control (similar to the spikes we saw in Figure 6.1.1). Unfortunately, 240s τ is the highest τ value we went up to, so, like with the previous graph, we can't make any definitive statements about what would happen at larger τ 's—but there certainly appears to be an upward trend in the data—suggesting increasing Allan variance with increasing τ . To get a sense of how much of an impact

these tests had, I've made a fractional difference plot below which we can compare to figure 6.1.3.

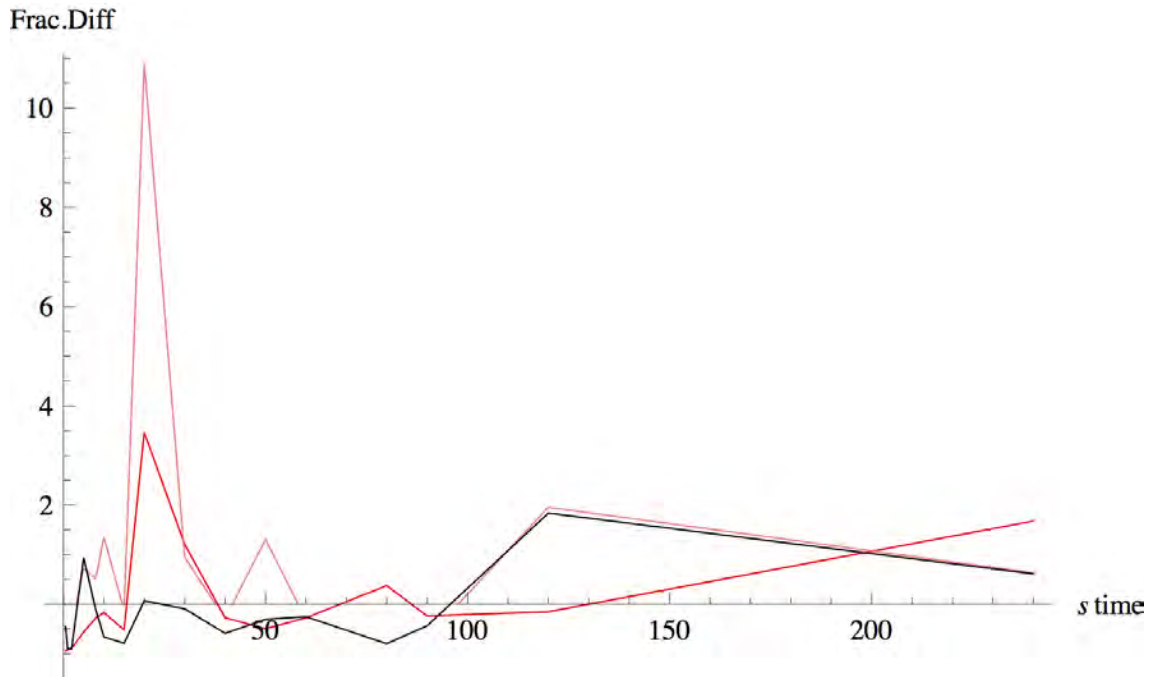


Figure 6.2.2. Fractional differences for vibration test: 29.8Hz (red), 17.8Hz (pink) 130.8Hz (black)

Here (as with Figure 6.2.1) the red line is the fractional difference of A with respect to A being vibrated at 29.8Hz, the pink line refers to the 17.8Hz test, and the black line is the 130.8Hz test.

The way I set up our fractional difference equation means that when one of the lines dips into being negative, it has a smaller Allan variance than the control clock (it performs better than the control), and when it is positive, it has a larger Allan variance (it doesn't perform as well as the control). It's also important to note that a value of 0 would denote two measurements of the same value. It is clear from the plots, that the 29.8Hz and 17.8Hz, are more frequently positive, and have large spikes in the positive range—which, when compared to Figure 6.1.3, might point towards vibrations at these frequencies as

having a detrimental effect on A's frequency stability. The 130.8Hz line starts around small positive values, then jumps to sticking to small negative values, but then there's a huge spike to around 2 for the 120s τ . Though there is one large spike in this line, it's very similar to the fractional difference plot with no added noise, so it appears as if vibrations at this frequency didn't really have any effect on the clock. This could be due to the fact that at higher and higher frequencies, the tray being vibrated would approach appearing as if it wasn't even moving at all. Choosing a high frequency, then, experimentally resembles A not being vibrated, and it appears as if this may have shown in the data.

6.3 Magnetic Field Test

Through the use of the function generator we subjected our clock to a 34.5Hz oscillating field. In the results pictured below, the red line is Clock B with no added electromagnetic field, and the blue line is Clock B inside the added field. The lines follow one another pretty closely, except for a large spike in the blue line at $\tau = 120s$. Other than this spike, the graphs follow one another—with the electromagnetic field line pretty consistently being a bit above Clock B with no field.

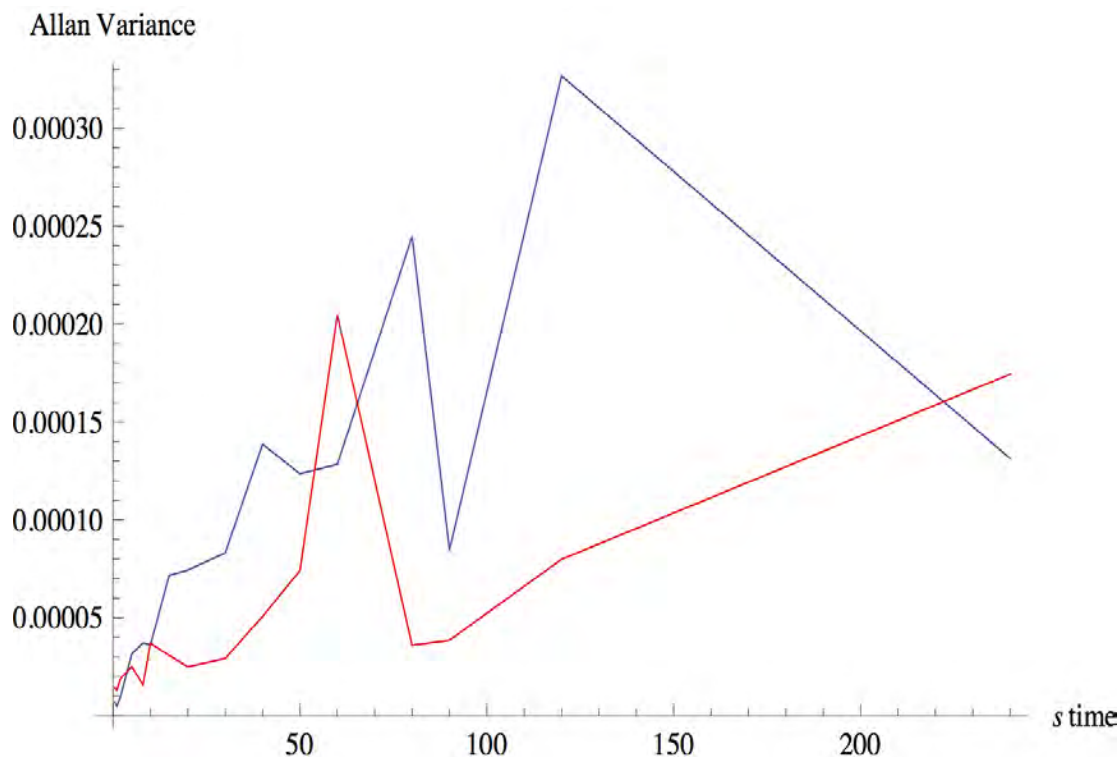


Figure 6.3.1. Magnetic field run (blue) compared to B with no added field

I believe a reason we don't really see much of an effect could be because the field was weak. The magnetic field is created by the accelerating current, and because we wanted to oscillate the field near the frequency of our clock (a slow frequency), the current wasn't accelerating very fast—leading to a field of decreased strength. I noticed that the field had a maximum strength when the frequency of the current through the solenoid was around 7KHz, so, out of curiosity, I immersed Clock B in this field too. The reason I didn't originally want to immerse the clock in a field oscillating this quickly is because due to the fact that a 7KHz frequency is about 234 times faster than the clock ticks, it would appear as if the field is "standing still" relative to the clock, and I hypothesized that a static field won't have much of an effect. The results of this test are pictured below on the following page.

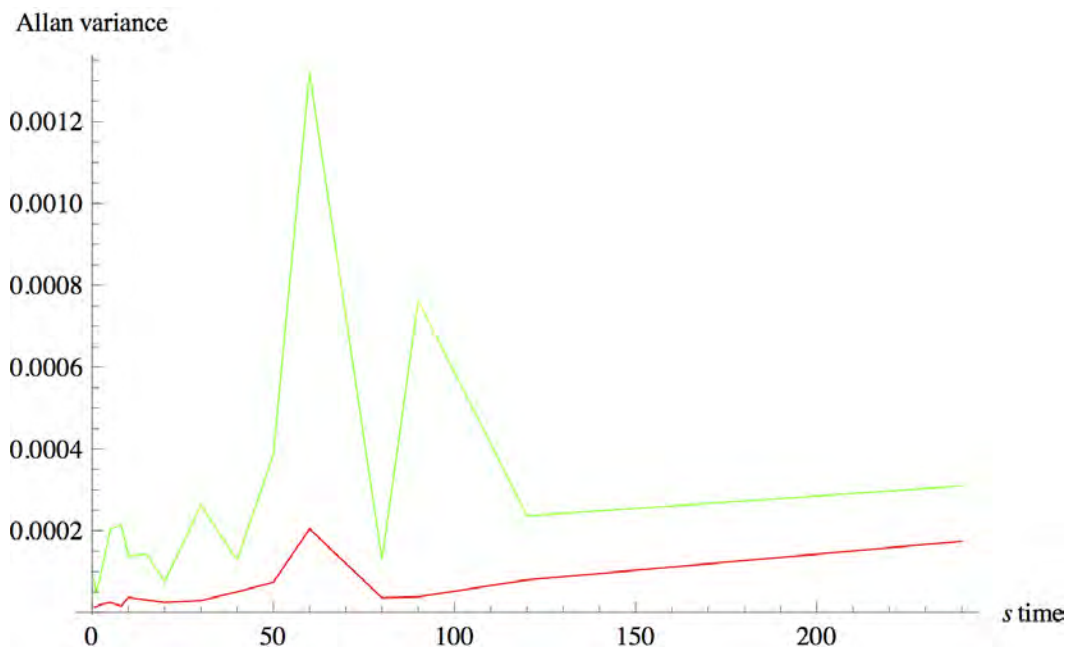


Figure 6.3.2. strong field test (green) compared to B with no added field

Here the red line once again denotes Clock B with no added field, and the green line is Clock B in the strong magnetic field. As is obviously depicted by the graph, the strong magnetic field test has a larger Allan variance than B at all measured τ 's. There are two large spikes at 60s and 90s, but other than that the strong field line is somewhat close to B. To more easily see how large these differences are, I made another fractional difference plot, which is shown on the following page.

The orange line is the fractional difference compared to the strong field, and the blue line is the fractional difference compared to the weak field with frequency near our multivibrator's frequency. The strong field line has large fractional differences, and is only in the positive range. The weak field line dips down to be negative 3 times (at .5s, 60s, and 240s) but for the most part has small amplitudes—the largest amplitude appearing at 80s. Due to the fact that the blue line has mostly small amplitudes, it can be pretty safe to assume this field didn't really have any effect on B's frequency stability, and that the spike at 80s is most likely a statistical anomaly. The strong field line, however, tends

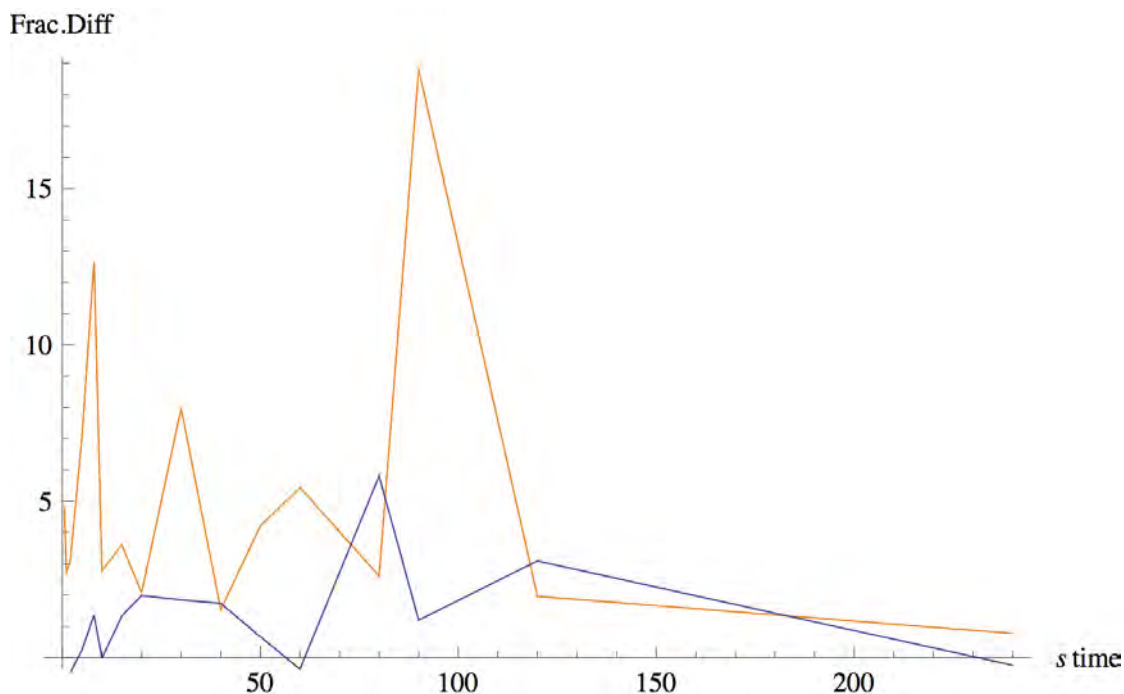


Figure 6.3.3. Fractional differences for magnetic field test. Orange is the large field run, blue is the weak field run

towards larger amplitudes when compared to Figure 6.1.2, so it appears as if this strong field caused loss in stability of B's frequency standard.

6.4 Our clocks Compared to one Another

I stated earlier that along with comparing our clocks to the computer, we'd want to compare them with each other, and then extract the variance of each clock from that data. Below I have plotted C using A as a reference (CA), A using B as a reference (AB) and B using C as a reference (BC).

All three clocks start out with the same relative shape, up until the 50s mark, then BC breaks away and dips down, while AB and CA stay pretty close throughout the entire plot.

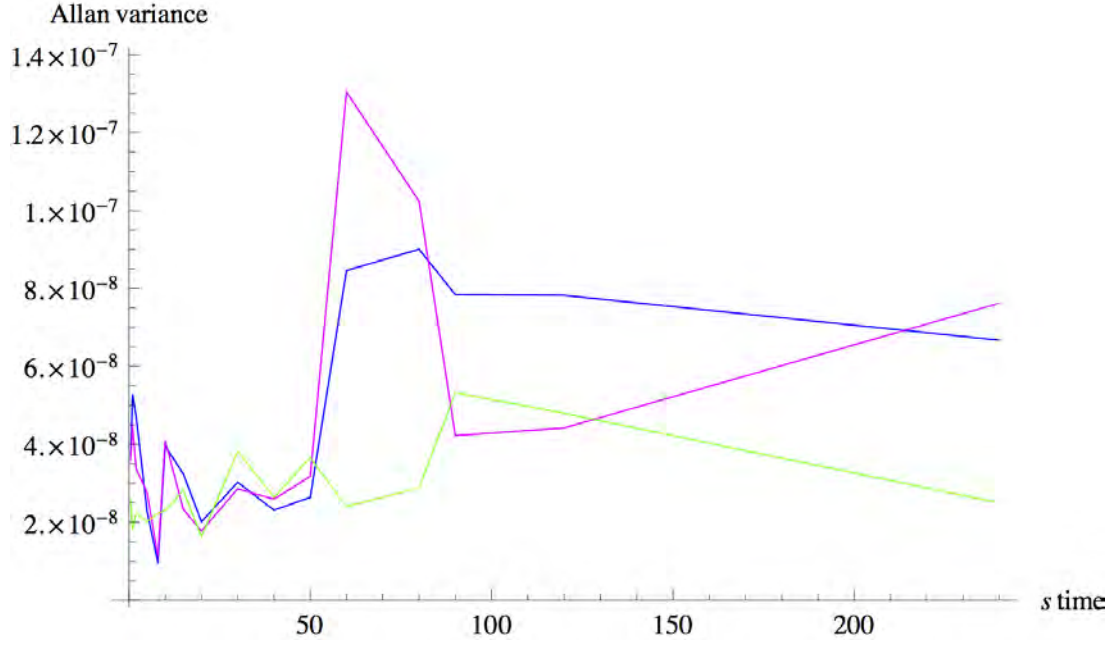


Figure 6.4.1. Allan variances of CA (blue), AB (purple), and BC (green)

We then used the three-cornered hat method to extract the Allan variance for each clock at each τ . The results are plotted below.

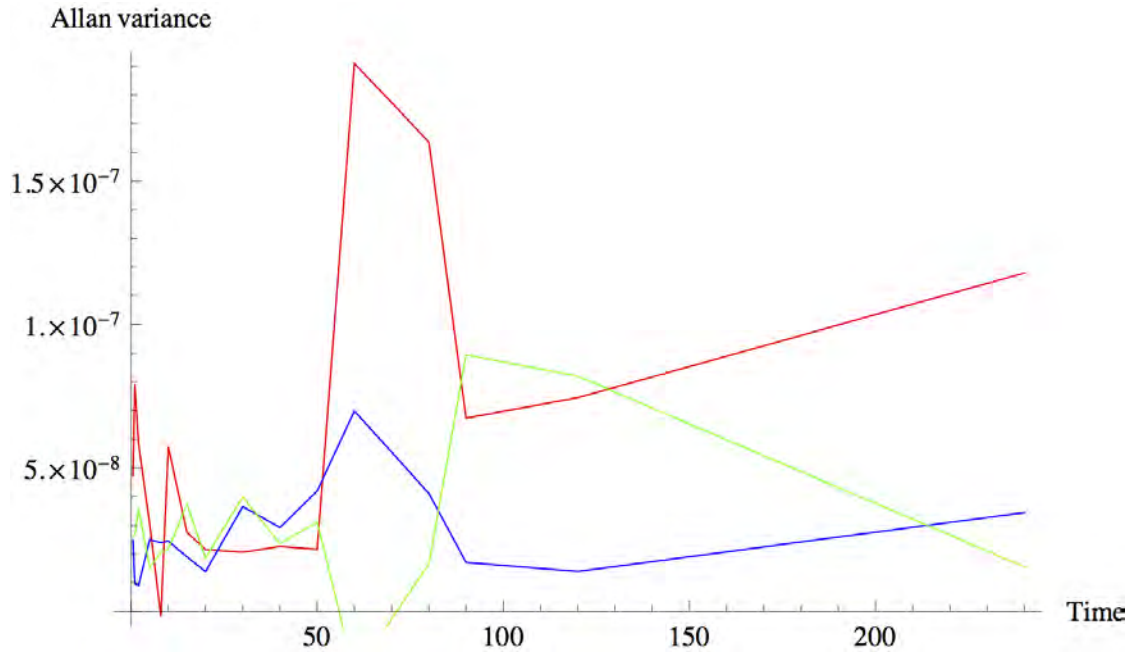


Figure 6.4.2. Extracted Allan variances of each clock

The most notable part of this graph appears at 8s and 60s. At 8s Clock A barely dips below zero, but at 60s Clock C becomes significantly negative. Obviously a negative Allan variance doesn't make any sense, so there must be something wrong with our method. The first thing that comes to mind has to do with the following relation from the N-cornered hat method

$$\sigma_{ij}^2 = \sigma_{ji}^2.$$

This states that the variance of clock i vs clock j is the same as j vs i . Our data, however, does not follow this relation, and it's most likely due to how we set up our experiment. The N-cornered hat method assumes you actually use the clocks as references—it doesn't assume you do what we did, which was use a computer clock and then try to cleverly express our frequencies in terms of one another. If we use A to measure B, for instance, we would count how many times B oscillated per x-number of A oscillations. We would then get a proportion of

$$\frac{N_B}{N_A}.$$

Remember then, that because we're using A as our reference, N_A (the number of oscillations of A) is held constant, and N_B (the number of oscillations of B) is allowed to vary. You would then measure differences in this proportion as your Δy 's. If you then wanted to measure A vs B, you would really, in essence, just be inverting this proportion, and measuring the difference of that as your Δy 's, and clearly, these *differences* should be the same—regardless of the order of the proportion. This isn't the case with the method we're doing, however. We're using the relation

$$\nu(ij) = \frac{\nu_{iR}}{\nu_{jR}} - 1.$$

If $\nu_{ij} = \nu_{ji}$, then

$$\nu_{ij} - \nu_{ji} = 0$$

which leads to

$$\left(\frac{\nu_{iR}}{\nu_{jR}} - 1\right) - \left(\frac{\nu_{jR}}{\nu_{iR}} - 1\right) = 0.$$

The ones nicely cancel, and we're left with

$$\frac{\nu_{iR}}{\nu_{jR}} - \frac{\nu_{jR}}{\nu_{iR}} = 0.$$

Finding a common denominator we see

$$\frac{\nu_{iR}^2 - \nu_{jR}^2}{\nu_{iR}\nu_{jR}} = 0,$$

and here we see our problem. Because i and j are two independent clocks, there's absolutely no reason ν_{iR} *must* equal to ν_{jR} —which leads to $\sigma_{ij}^2 \neq \sigma_{ji}^2$.

This lack of symmetry comes from the approximation $\nu_{ij} = \frac{\nu_{iR}}{\nu_{jR}} - 1$. This is what allows for negative Allan variances, so we shouldn't be too surprised to see that pop up in our data—nor should we worry too much. The whole point of this approximation was to be able to get a ballpark sense of how good our clock was—how much time would it gain or lose for a certain length of measurement. The way we can use this data to do this is via the following relation

$$\sigma_i = \overline{\% \text{Change}}.$$

This might seem surprising, but remember, that the Allan variance of one of our clocks with reference to another is defined as

$$\sigma_{ij}^2 = \frac{1}{2} \overline{\left(\left(\frac{\nu_{iR}}{\nu_{jR}} \right)_{n+1} - \left(\frac{\nu_{iR}}{\nu_{jR}} \right)_n \right)^2}.$$

You'll notice, that when the square root is taken, this is really just an average percent change in frequency with respect to some reference frequency. We can then generalize this reference frequency as the *ensemble* frequency (remember this from chapter 3 section 2), and see that

$$\sigma_i \nu_{\text{ensemble}} = \Delta \nu_i.$$

This is how we can use the variances calculated via the 3-cornered hat method to figure out how good our clocks are. We can calculate the expected $\Delta\nu$, to get an upper and lower bound of what we can expect our frequency to be, and then calculate the time lost or gained. For instance, the average Allan variance for our clocks at a τ of 240s is about 5.5965×10^{-8} (I calculated this average from the raw data). If we take the square root of this number and insert it into our relation above, and multiply it by $\nu_{\text{ensemble}} = 29.4648$ (just the average measured frequency of all the clocks over all τ 's), we see

$$\Delta\nu = .00697.$$

From this, we see that over a length of 240s, our clocks should be vibrating with a frequency around

$$f = 29.4638\text{Hz} \pm .00697.$$

If we compare the percent difference between the upper/lower bound and the average, we can use this percent change to calculate the time our clock may have lost or gained any allotted time. For instance, using the numbers above, the percent change between the average and the upper bound is $\frac{.00697}{29.4638} \approx .0237\%$.

To get a sense of how stable this is over longer periods (like hours) we could pretend that our τ at 240s stayed constant for all higher τ 's—this might seem silly, but this'll give us a ballpark idea of what kind of a magnitude of instability our clock may have. The original experiment was going to use atomic clocks to try to measure changes around 40ns over a 40 hour timescale. In order for them to do this, they would have to be stable up to tens of nanoseconds over this time period. To get a sense of how our clocks would do, we can multiply this .0237% change by 144,000 seconds, and the resulting number will be the amount of time we can expect our clock to have gained or lost. Doing the math, in 40 hours, our clock would have lost or gained about 36s—meaning it's nowhere near as good as the atomic clocks which would have been used for the plane experiment. The above

method demonstrates how you can take three comparable clocks, compare them only to one another, and from that, discern how good they are relative to any other clock.

7

Conclusion

After running our tests and looking at the data, it becomes quite obvious we're lacking statistics. Each full-length τ run took about 7.5 hours, so we weren't able to run each test multiple times (which would have been the right way to do something like this)—making any definitive statements about our data difficult. We were able to, however, achieve some pretty nice results. Overall, our Allan variance plots for the clocks with a computer reference have upward-trending behaviors, which is what we'd expect to see, and though we weren't able to be quite definitive about the effect from the vibrational noise, it did appear as if the strong field had a detrimental effect on Clock B's frequency stability. We were also able to use the three-cornered hat method to discern—simply from using our clocks as references to one another—that our clocks are nowhere near as sensitive as atomic clocks (surprise surprise).

It turns out that we were lacking the very thing we were trying to measure. Had we more time, we would have not only run each test multiple times, but hopefully we would have been able to figure out the PID situation, and then have clocks experience multiple noises at once. It's important to remember, however, that the goal at the onset of this

experiment with the multivibrators was to be able to come up with ways to create and measure the effects of environmental noise on clocks, and to develop a way to extrapolate the amount of time we may expect our clock to gain or lose in a given time period—hoping then to be able to use these methods on better clocks. With these specific goals in mind, I would consider this experiment a success. I’m confident that (with perhaps a little refinement here and there) these methods could be used on more sensitive clocks to develop an understanding of their behavior when subjected to different types of noises.

8

Appendix

8.1 Appendix A: The Temperature Controller

As mentioned in chapter 6, we were going to try and run a temperature control test using *Labview* and a homemade temperature controller with directions from John Essick's book "Hands-On Introduction to *LabView* for Scientists and Engineers". Our logic behind this test was: when molecules absorb heat, their kinetic energy increases, and they move faster and more erratically—and conversely, when molecules release heat, they move more sluggishly—so this change in molecular motion, we thought, could have some effect on the stability of the oscillator's resonant frequency.

In order to test this effect, we were going to make a PID controller in *LabView*. A PID, or "proportional-integral-derivative controller", is a control loop feedback system. The PID controller constantly calculates an error value between a setpoint (chosen by the experimenter) and a variable (in this case temperature). The PID attempts to minimize the error between the setpoint and variable over time. In this case, the PID Controller would attempt to minimize the difference in temperature between the setpoint (the temperature we want) and the temperature the clock is actually at.

As already stated, the PID controller will be created with the help of *LabView*—a programming interface. The output from the program was going to be connected to the circuit with the oscillator and thermistor. A thermistor is a type of resistor that has a predictable relation between its temperature and its resistance. The PID controller will be set to maintain a stable voltage drop across the thermistor. As the room heats up or cools down the thermistor, the thermistor’s resistance will change—thus changing the potential drop across it. When the PID detects this change in voltage, it will act to add or subtract voltage across the thermistor, so as to keep the thermistor generating the amount of heat necessary to keep the oscillator at the desired set-point temperature.

In the end; however, we weren’t able to run this test because the data acquisition board (this is what *LabView* uses to measure the voltage) we have is a counter... and a counter only. To get the software PID to work, we’d need a data acquisition board which can take voltage readings, and we simply didn’t have one.

We tried to build a hardware PID, but that proved to be a difficult task, and with time running out and difficulties occurring in other parts of the experiment, the temperature experiment had to get cut. Before we officially removed the temperature test; however, we had finished building the temperature controller—described in the section below.

8.1.1 *Building the Controller*

Building the temperature controller with Richard Murphy was a fun experience. To build the controller, we had to buy: thermistors, a heat sink, a 4” cooling fan, thermal grease, and a thermal electric cooler.

First, using a milling machine, we cut down an aluminum block into 2”x1.5”x5/16” dimensions. We then used an electric drill press to drill four holes into the block. When looking at the block from the top, there will be two 8-32 Clearance Holes drilled into the block 1.450” apart (measured along the 2” edge and positioned along the centerline).

Then, a 6-32 Set Screw was drilled in 1/2" into the center line of the block measured from the 2" edge. This hole is for a set screw which will hold the thermistor in place. The fourth hole, in which the thermistor will be inserted, is located 1" in along the 2" long side, right in the middle. This hole and the 6-32 Set Screw hole will intersect, allowing the set screw to press down on the thermistor.

Once the aluminum block was finished, we needed to figure out how to fasten the cooling fan to the bottom of the heat sink. We measured the dimensions of the cooling fan, and then marked the cooling fan dimensions directly on the heat sink. We needed to position the fan on the heat sink near the center, but we happened upon a slight difficulty. On the underside of the heat sink there are spiky fins. When the cooling fan was placed perfectly in the middle (on the smooth side), we noticed that if we were to try and drill with the fan laid out in such a manner, we'd drill into fins. So, we had to move the fan slightly off center, and rotate it about 30° or so. We then marked the fan's position, and with the electronic drill press we drilled 8-32 size holes clean through the flat side of the heat sink.

Once the holes were drilled in the heat sink, the next task was to figure out how to mount the cooling fan in such a way as to be attached to the heat sink, but also lifted above the ground—so as to be able to generate air flow. We took a thin aluminum pipe and cut it into four 3" pieces to be used as legs. We then took each leg and drilled an 8-32 hole about an inch or so deep into one of the sides. Then, we tapped this hole so it would be able to be screwed onto an 8-32 screw. Using a chuck, we smoothed out the other side of each leg. Each corner of the cooling fan has a hole for screws.

We put the fan face down on the table and, one at a time, took four 8-32 screws, put on a locking washer, and slid each screw into a corner hole with the head facing down, and then screwed in each leg—clamping the leg and screw to the fan, thus holding it in place. Pictures of the setup are on the following page.

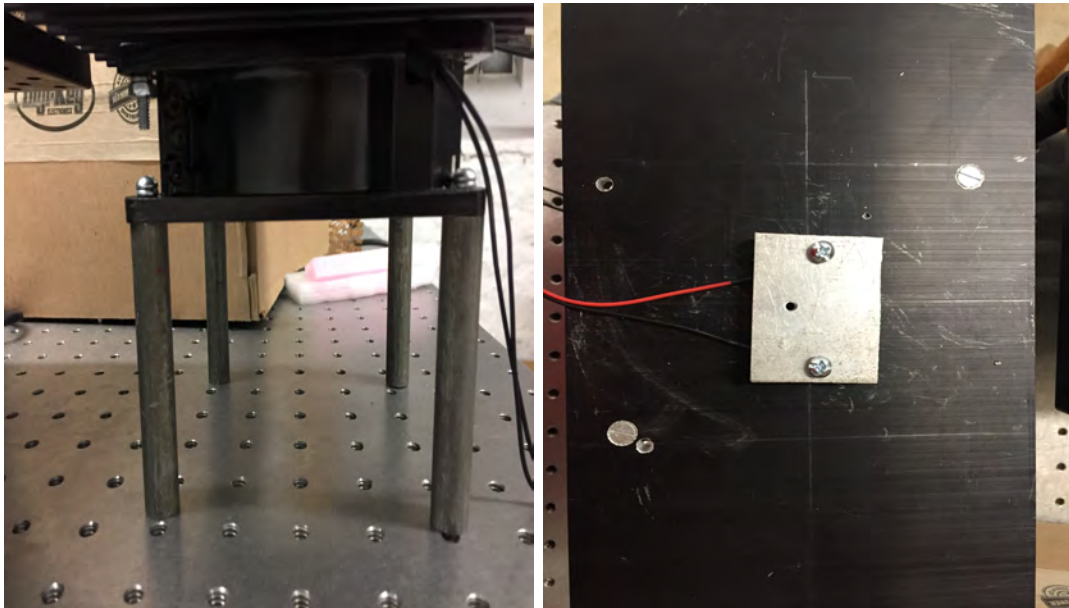
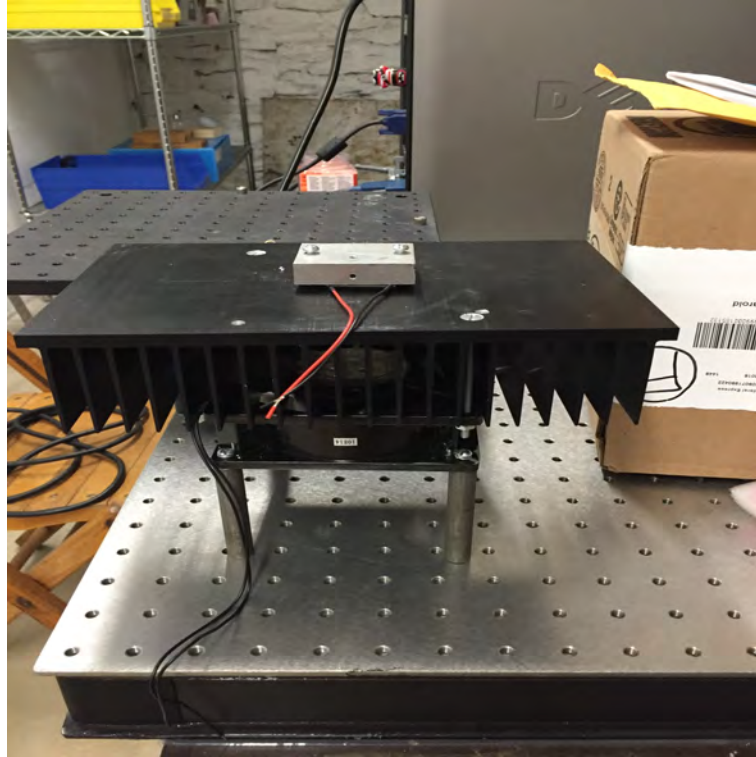


Figure 8.1.1. Temperature controller

8.2 Appendix B: The RLC Circuit

In the prologue I state that after we realized we couldn't use the atomic clocks for our experiments, we jumped right to the multivibrator. This isn't quite the case. Before we decided to work on the multivibrator circuit, we were trying to use another resonant circuit: the RLC (resistor-inductor-capacitor) circuit. We were going to use an RLC initially because the RLC circuit is the poster child of resonant circuits—they're so good at resonating at a tuned frequency that a series RLC is the equivalent circuit buildup for quartz crystal oscillators. To get the RLC to resonate, we needed to input a square-wave. When the square-wave hits the RLC, it generates a sinusoidal waveform that oscillates at the RLC resonant frequency. This all worked perfectly, but the reason we had to make the switch from RLC to multivibrator has to do with the damping factor and how the counter makes frequency measurements. The resistor in an RLC circuit dampens the oscillations, which leads to decreases in amplitude until the generated signal dies out. So the RLC output would start at a maximum amplitude (right when the square wave started) and wouldn't reach that maximum amplitude again until it was once again struck by the square-wave. Unfortunately, the counter measures frequency by looking at maximum voltages. It looks for some maximum voltage, and then counts the amount of time the system takes to reach that voltage again, and calls this time the period of the signal. Because the RLC signal would only reach its maximum amplitude right after it was hit by a square-wave, the counter only ended up measuring the frequency of the square-wave being sent in. So, with time running out and data needing to be collected, instead of trying to figure out how to get around this, we made the switch to the multivibrator circuits. A picture of the RLC predicament is shown on the following page.

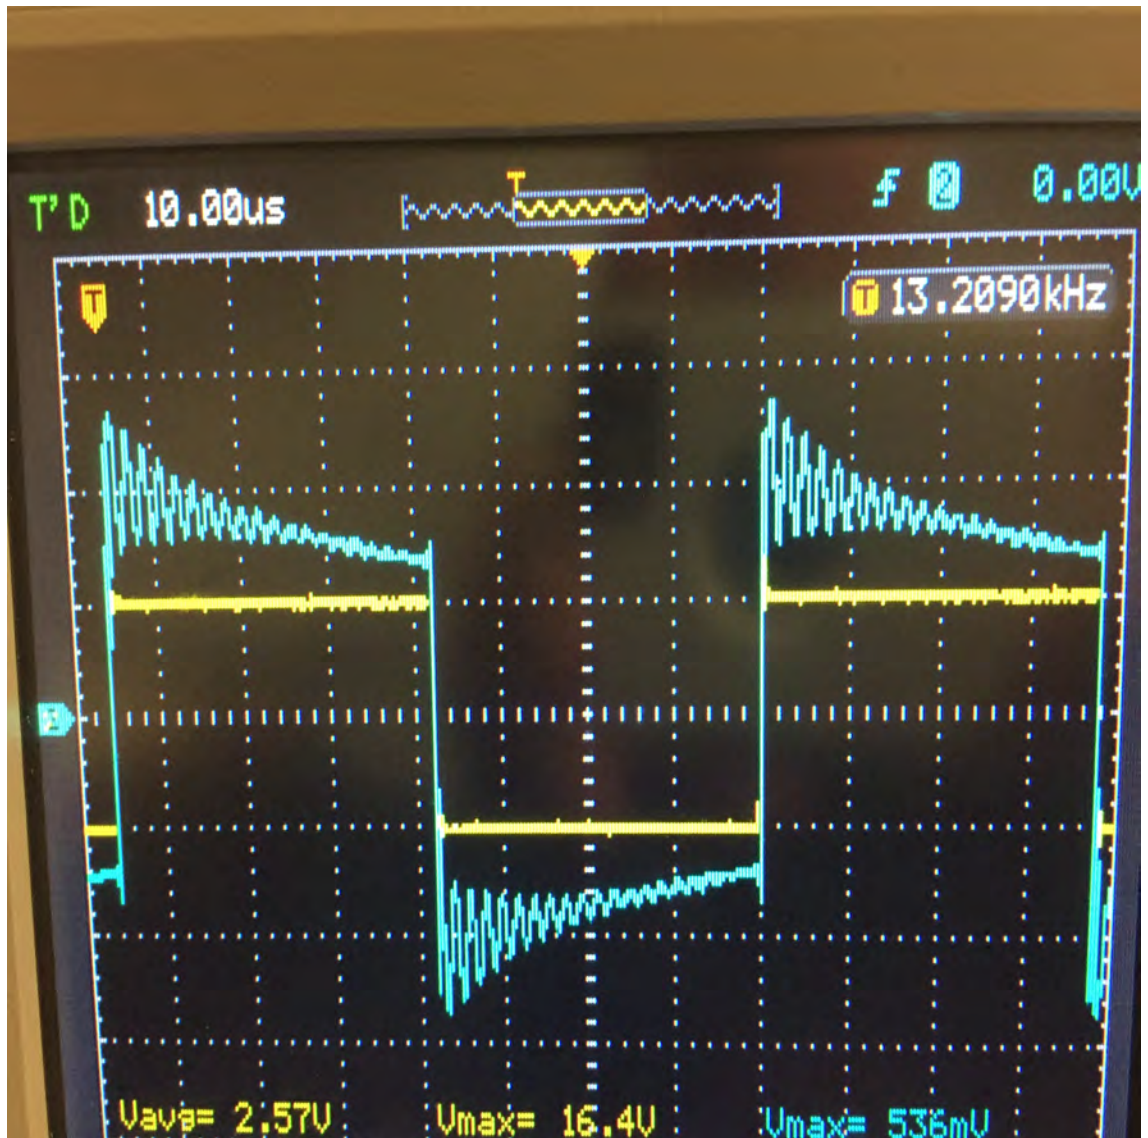


Figure 8.2.1. RLC output (blue) and square-wave input (yellow). Note the decreasing amplitude of the RLC output

8.3 Appendix C: Miscellaneous Relativity

As stated in the prologue, initially this experiment was going to try to measure relativistic time dilation via a plane flight to Europe. Before we realized we weren't going to be able to get the clocks in time for my trip, we had calculated that the net time difference the clocks would see should be around 40 or so nanoseconds (a rough figure), and we determined that this wouldn't be a large enough difference to confidently say we've measured time dilation—a difference that small could be chalked up to clock error or oscillatory drift. When we realized this, we didn't immediately give up on trying to measure time dilation.

We thought about putting a clock at higher elevation in the Catskills (which would measure General relativistic time dilation), and leaving it there for some time, and we thought of ways to make one move rapidly in a controlled environment (to measure Special relativistic time dilation). One idea we came up with to measure special relativistic time dilation was to mount the clock on a motor at some fixed radius, and have it rotated around really fast. Before we really tried to nail down the logistics of this, we found ourselves wondering whether or not the clock would actually experience Special relativistic time dilation—or General relativistic time dilation. The clock is moving—so one might immediately think special relativity is what causing the time dilation, but special relativity only holds in inertial frames, and circular motion is not an inertial frame. To try and get a hold on this problem we did some calculations.

We first looked at the problem with as if it were a special relativity problem. As stated in 1.0.1, the time measured by the atomic clock (we'll call this t_{atomic}) would measure

$$t_{\text{atomic}} = t_{\text{ground}} \sqrt{1 - \frac{v^2}{c^2}}$$

where t_{ground} is the time measured in the frame which is stationary relative to the Earth. From kinematics we know that in a rotated frame, $v = \omega R$, where R is the radius of the

circular path traveled, so we can plug this in and see

$$t_{\text{atomic}} = t_{\text{ground}} \sqrt{1 - \frac{(\omega R)^2}{c^2}}. \quad (8.3.1)$$

This is the full result if it were a special relativistic effect.

We then decided to go about the same process but assume the result was a general relativistic effect. From equation 1.0.2, we have

$$t_{\text{atomic}} = t_{\text{ground}} \sqrt{1 - \frac{2GM}{rc^2}}.$$

This, however, is the time dilation due to gravity. We can notice, however, that $\frac{GM}{r}$ is just a potential energy term. This is the gravitational potential energy you experience due to the earth. We can, then, come up with an analogy to this term for our situation. We can arrive at this analogy by starting with the force the clock would be feeling in this scenario. $F = ma = \frac{mv^2}{r}$ (rotational kinematics), and, because we know the force will always be pulling towards the center of the rotational motion, we can define this direction as the negative radial direction. Then $F = -\frac{mv^2}{r}$. Another kinematic relation we know is

$$U(r) = - \int F(r) dr$$

where $U(r)$ is the radial potential energy and $F(r)$ is the radial force. From our relation just above we have $F = -\frac{v^2}{r}$, so plugging this in gives us

$$U(r) = \int \frac{v^2}{r} dr.$$

We remember that we can substitute $v = \omega r$ in for v and arrive at

$$U(r) = \int \omega^2 r dr.$$

Our integral will be from $r = 0$ to $r = R$, so plugging this in and integrating gives us

$$U(r) = \int \frac{\omega^2 R^2}{2}.$$

We can then plug this into 1.0.2, and see

$$t_{\text{atomic}} = t_{\text{ground}} \sqrt{1 - \frac{\omega^2 R^2}{c^2}}$$

which is precisely the expression we arrived at earlier in 7.3.1! It's quite interesting that you get the same result regardless of what type of an effect you model it as. I was quite surprised when we arrived at this conclusion—I had expected that we'd come to two different expressions, and then have to analyze them and try to determine which one fits the situation better.

Bibliography

- [1] *Multivibrator*, <https://en.wikipedia.org/wiki/Multivibrator>.
- [2] *Op-Amp Multivibrator*, <http://www.electronics-tutorials.ws/waveforms/astable.html>.
- [3] *Application of the 3-Cornered Hat Method to the Analysis of Frequency Stability*, <http://www.wiley.com/3-CornHat.htm>.
- [4] *Solenoid Magnetic Field Image*, http://www.electronics-tutorials.ws/io/io_6.html.
- [5] *RLC Circuit*, https://en.wikipedia.org/wiki/RLC_circuit.
- [6] *Capacitor Charging and Discharging Image*, <http://i.stack.imgur.com/l3iBb.gif>.
- [7] *Parallel Plate Capacitor Image*, http://www.electronics-tutorials.ws/capacitor/cap_1.html.
- [8] *Op Amp Circuit Diagram Image*, https://en.wikipedia.org/wiki/Operational_amplifier.
- [9] John Essick, *Introduction to Labview for Scientists and Engineers*, Oxford University Press, New York, 2009.
- [10] Joshua Cooperman, *On the Allan variance of comparable clocks* (April 19, 2016), 1-2.