
A Discussion of Techniques to Mitigate Non-linearity
in Scale Factor within a Linear Accelerometer Under
Capacitive Drive

SENIOR HONORS THESIS

WRITTEN BY

KRISHNA SONI

TUFTS UNIVERSITY

CHARLES STARK DRAPER LABORATORY

2013-2014

Abstract

The development of highly accurate, strategic-grade sensors is integral for the development of Inertial Navigation Systems (IMUs). Currently, sensors which make up IMUs are limited by bandwidth, size, and stability. The Zero Force Accelerometer (ZFA) aims to exceed these specifications as the next generation of high precision sensor. It is necessary then to mitigate any source of error in order to increase the integrity of the device. Small displacement error, temperature drift and variable analog inputs call for an in-situ scale factor calculation, but at a cost. In particular, a g-dependent term was observed, propagating to 61,599 ppm/g. A new algorithm was proposed exploiting changing parameters on the fly, and calculating and combining two scale factor measurements to eliminate the g-dependency. The results showed two orders of magnitude improvement of scale factor error (311.43 ppm/g), but still does not meet strategic-grade specifications, so further study is desired.

IMPORTANT NOTE

Due to contractual obligations with Charles Stark Draper Laboratory, parts of this thesis have been redacted to satisfy the needs of all parties involved. In particular, data taken from the Zero-Force Accelerometer project have been removed. The original, authoritative copy of this thesis has been reviewed by all members of the thesis committee, and has been submitted to the University archives but will not be available to the public. Aside from any redacted material, the original content of this thesis has not been changed.

REDACTED COPY

Acknowledgments

I am incredibly grateful to the Electro-Optics group at Draper Laboratory for allowing me to work with them over the past year on this project. In particular I would like to thank Tim McCarthy, Mark Bowser, Michael Tomaino-Iannucci, and Steve Burgess, all of whom I worked closely with and helped me make this thesis a reality. It is incredible to think how much I learned in the lab, and how much I took away from it to become a better scientist. I am also very grateful for David Butts and Krish Kotru, who both helped mentor me during my time at the lab, as well as gave me invaluable advice in moving forward with my scientific career. I would like to thank my thesis advisor, Professor William Oliver who was incredibly patient with me and my work throughout this whole process. I would like to express my extreme gratitude to Professor Sliwa, who volunteered to be included on the committee very last minute, but managed to provide a great amount of input on my project. I would also like to thank my friends who have been supportive to me throughout this whole process especially my friend Sarah Nasser, who convinced me that I was capable doing a project of this scale in the first place. Finally, I would like to thank my parents, both of whom have not faltered from wanting to see me succeed in this endeavor and all of my future pursuits.

Contents

1	Introduction	9
1.1	Background and Discussion	9
1.2	Structure of Accelerometers	11
1.3	Sources of Error In Linear Accelerometers	14
1.4	Current State of Accelerometer Errors	16
2	Analysis of Linear Accelerometer Mechanics	18
2.1	Effects of Natural Frequency and Dampening	20
2.2	Vibrational Effects	23
2.3	Non-Linearity Due to Resonance	23
2.4	Conclusions	28
3	Classical Analysis of the ZFA	30
3.1	Overview of the ZFA	30
3.2	PID Control	32
3.3	Attaining Closed Loop Operation	34
3.4	Analysis of ZFA Measurement Characteristics	37
3.5	Correspondence to Input Angle	38
4	The Scale Factor	41
4.1	Definition and Motivation	41

4.2	Statically Calculated Scale Factor	42
4.3	In-Situ Scale Factor Algorithm	44
5	Contributions: Sequential Two-Dithered Approach	54
5.1	Motivation	54
5.2	Using Two Dither Measurements	55
5.3	Choosing the Two Frequencies	56
5.4	Implementing the Algorithm	57
5.5	Discussion of Results	60
6	Conclusion	61
6.1	Reflections on Work	61
6.2	Known Problems and Future Studies	62

REDACTED COPY

List of Figures

1.1	Basics of Stable Platform Algorithm	11
1.2	Basics of Strap-down Algorithm	11
1.3	A simple diagram of a single degree of freedom damped harmonic oscillator[11]	12
2.1	A basic example of a linear accelerometer. [11]	19
2.2	Figures showing effects of resonance. The graph on the left shows the effects of displacement amplitude as the driving frequency approaches the natural frequency. The plot on the right shows the effects of varying frequency about the natural frequency with phase[4].	26
2.3	A damped harmonic oscillator with undergoing an impulse in direction y relative to the lab frame	27
3.1	A detailed structural diagram of the Zero-Force Accelerometer	31
3.2	Diagram of Generic PID Control Loop	33
3.3	Diagram of Mechanical Loop in ZFA	36
3.4	Block Diagram of ZFA Response Loop	36
3.5	The ZFA mounted on the index head with 0g incident on the proof-mass	39
3.6	The ZFA mounted on the index head with 1g Incident on the proof-mass	39
3.7	Plot showing linear relationship between the sine of input angle θ and squared voltage over the capacitive plates V_{Cap}	40

4.1 Output acceleration versus known input acceleration. The blue line shows the measured acceleration while the red line shows the ideal acceleration output. An ideal acceleration output shows one-to-one correspondence to the known input acceleration. 43

4.2 Rectified Force Error (ppm) versus Input Acceleration (g) using Single Dithered Scale Factor. This plot shows the coherence of the output acceleration from the idealized output of the sensor (i.e. 1-to-1 ratio between input and output acceleration). The error propagation here constitutes 21.5%/g error in scale factor over a range of 1g. . 43

4.3 A trivial example of a dithered system. A sine wave is integrated over one cycle. The blue line represents the pure dither signal, while the pink line represents the motion of a proof-mass in a perturb system. Clearly, the pink line does not integrate to zero after one period, so there is information to be gained (and compensated for) from the perturbed system. 46

4.4 Diagram of Mechanical Loop in ZFA with the dither 47

4.5 Scale Factor deviation (ppm) versus input acceleration (g) with associated error. Scale Factor Data was averaged for each point and compared with the initial scale factor reading to show propagated error. The trend shows an error propagation of 61,529 ppm/g. This error is less than that of the statically applied scale factor but still well above the approved stability for a strategic grade sensor. 52

4.7 Rectified Force Error (ppm) versus Input Acceleration (g) using Single Dithered Scale Factor. The error propagation here constitutes 15%/g error in Scale Factor over a range of 1g. 52

4.6 Input Acceleration (g) versus output acceleration (g) taking into account applied single dither scale factor (pink line), and statically calculated scale factor (blue line). Voltage and Scale factor were measured at each point, and then applied according to Equation 4.1 to find output acceleration. As expected, the acceleration curve shows large deviations from the idealized output of the sensor (red line). 53

5.1 PSD Plot of capacitive plate voltage response (mV) for a frequency range of 0 to 400HZ 57

5.2 Plot showing scale factor calculated at 300Hz (blue line) and 175Hz (pink line) and the combined scale factor (red line). Clearly, by applying equation 5.5 to the calculated scale factors using a single dither, we can eliminate the g-dependency. 58

5.3 Sequential Two-Dither Scale Factor Error (ppm) vs. Input Acceleration (g). This plot shows as little to no trending in scale factor, but has non-linear outputs that could be caused by other factors. The scale factor error is greatly improved, and equates to about ± 311.43 ppm/g. 59

5.4 Correspondence between Input Acceleration (g) and Output Acceleration (g). Compared to the static and single dithered scale factors, the output acceleration calculated with the sequential scale factor shows strong correspondence to the input acceleration. 59

5.5 As highlighted in Figure 5.3, the scale factor algorithm shows strong correspondence to the idealized output acceleration measurement. The propagated error (aside from the outliers at 0.31g) gives 7.31%/g error in scale factor up to 1g. 60

List of Tables

1.1 Characteristic Performance Metrics in Highest Precision Navigation Systems 17

4.1 Measured Structural Components of ZFA [5] 42

5.1 Measured Resonant Frequencies of ZFA[2] 57

REDACTED COPY

Chapter 1

Introduction

The following section will briefly go over the basics of accelerometers, and the motivation for constructing such devices at strategic-grade precision.

1.1 Background and Discussion

Inertial sensing devices play a pivotal role in the advancement of modern navigation systems and scientific studies. Advances in inertial sensing allow the use of self-contained, highly precise navigation systems, which may be utilized in a variety of ways, from guidance systems in spacecraft and aircraft to applications in measuring gravitational fields of planets. New advances in inertial sensing aim to greatly improve engineering and scientific applications.

Inertial sensors with high stability and precision allow for precise navigation applications when external references such as Global Position Systems (GPS) are either unreliable or become unavailable. Inertial navigation has become integral in applications such as space and air flight, long duration naval navigation, guided missiles, precision satellite pointing, and autonomous navigation. These applications have sparked the need for devices that precisely measure linear acceleration and rotation, by virtue of accelerometers and gyroscopes, in order to provide accurate depictions of

position and velocity and thus reducing the need for external referencing.

In practice, inertial navigation is structured upon self-contained navigation provided by a known orientation, velocity and position. Inertial measurement units (IMUs) typically contain three orthogonally oriented gyroscopes and accelerometers. Given a known initial position, and velocity an IMU integrates small differences in both acceleration and angular acceleration over time to provide the true position and velocity. For example, an aircraft with an IMU moving east at 300 mph, will deduce that it 300 miles east of its last known position after 1 hour of flight time.

IMUs may fall under one of two categories: Stable Platform Systems or Strap-Down Systems. In stable platform systems, inertial sensors are gimbaled in so that they can be isolated from any external rotational motion, maintaining orientation in the global frame. The gimbals (frames) which allow free movement in all three degrees of rotational space. The mounted gyroscopes detect any change in platform rotations. These signals are fed back to torque motors which control the gimbal angle, and keep the device aligned with the global frame. To track position, the IMU performs double integration of data from mounted accelerometers (Note that due to this method, there is an offset in the vertical orientation of the acceleration due to gravity that must be accounted for).

Strap-Down systems entail sensors that are rigidly mounted into the device, and as a result make measurements from the body frame rather than the global frame. To keep track of orientation, signals from gyroscopes are integrated. Similarly, to track position, the signals from the three orthogonally oriented accelerometers are resolved using the global orientation due to the signals from the gyros, and subsequently integrated to find the position relative to that measurement (as described in the below figure). Strap-Down systems, because of their inherent design have a largely decreased mechanical complexity, with a trade off of a more computer intensive algorithm to find position and orientation. Such trade-offs result in physically smaller systems with reduced mechanical complexity. Such systems are now the dominant type of Inertial Navigation Systems (INS) due to decreased costs in computing power. The below figures show basic flow charts for calculations

of stable platform and strap down calculations.[11]




Figure 1.1: Basics of Stable Platform Algorithm

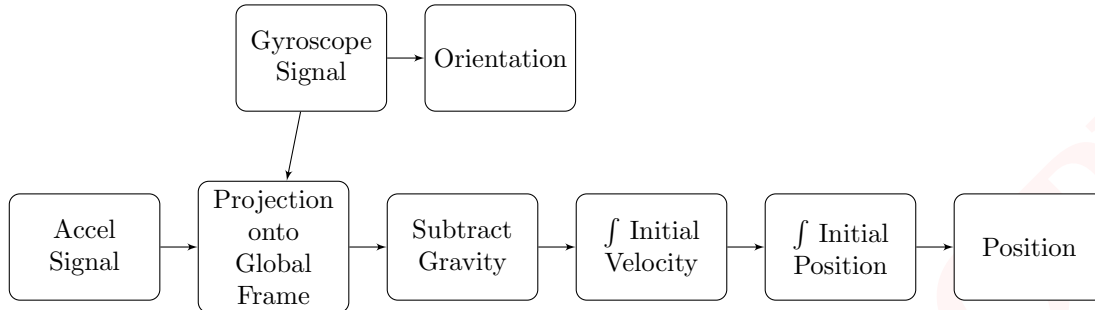


Figure 1.2: Basics of Strap-down Algorithm

1.2 Structure of Accelerometers

In any IMU, the quality of sensors integrated into the system will determine the overall accuracy of the device. For our discussion, we will be concentrating on the attributes of linear accelerometers, specifically mechanically driven, solid state devices such as the Zero Force Accelerometer. This section will go over the basic structure of linear accelerometers.

A linear accelerometer can be characterized as an accelerometer that only measures acceleration in one direction. The simplest example of this is the single degree of freedom harmonic oscillator, consisting of a mass-spring system with a viscous dampener. Here, if all aspects of the system are well known, the acceleration can be found easily as a function of the displacement pick-off (A more formal analysis of how this is done is encompassed in Chapter 2). While this example is trivial, it forms the basis of how modern accelerometers are constructed[4].




Figure 1.3: A simple diagram of a single degree of freedom damped harmonic oscillator[11]

1.2.1 Closed Loop vs. Open Loop

An important aspect of accelerometer operation is the distinction between open loop and closed loop operation. In open loop operation, a proof mass is let to freely move. The device's measurement is purely based on the measurement of the displacement of the proof mass. In closed loop operation, the proof mass is maintained in a fixed position and the force necessary to maintain that position is proportional the acceleration on the device. For strategic application, closed-loop operation is preferred. In open-loop operation, devices tend to experience cross coupling error, caused by the proof mass experiencing acceleration that is outside the axis being measured. In closed-loop systems, since proof masses are generally held in the same position, they are generally unaffected by pickoff nonlinearity, as they are only measuring very small distances (where in open-loop systems this is crucial). Further, open-loop systems are subject to pickoff nonlinearity and lagging effects from their mechanical springs. Due to the nature of the systems, closed loop systems may have mechanical springs that are orders of magnitude weaker in spring rate. Furthermore, since the proof-mass moves so little, the spring rate profiles are well defined and maintainable. This thereby reduces the effects of hysteresis from mechanical springs.

There are many ways to mechanize accelerometers, some examples include pendulous accelerometers, piezoelectric accelerometers, MEMS, and capacitive accelerometers. Each of these accelerometer builds fit within different specifications and types.

1.2.2 Mechanical

Mechanical accelerometers work by measuring the displacement of a mass on a spring using a displacement pick-off measured directly or by virtue of an optical source. The signal measured is proportional to the force acting along the input axis. Mechanical accelerometers, while simplistic in design, tend to have large measurement errors due to many moving parts and reliance on the accuracy and linearity of the displacement measurement.

1.2.3 Solid State

Solid State accelerometers vary in both utility and type. Different types of accelerometers used today include quartz, acoustic and silicone devices. Such devices are based on the principles of mechanical accelerometers but prove to be more accurate due to minimization and durability, making them ideal for strategic applications.

1.2.4 Capacitive Accelerometers and The ZFA

One type of accelerometer is a capacitive accelerometer. In practice a proof mass is held between two plates, forming two capacitors; each with a single fixed plate and each sharing the surface of the proof mass as a movable plate. In closed loop operation, the proof-mass is held in its null position by instigating a voltage across either stationary plate in accordance to a displacement measurement of the proof-mass, usually obtained by an optical pick-off. When the device experiences acceleration with respect to the lab frame, the proof-mass's motion is nulled by virtue of instigating a variable voltage across the plate that it is accelerating towards. Thus, the change in voltage (relative to its initial voltage) measured across the excited capacitor is proportional to the square root of the

acceleration of the device. The ZFA is a type of capacitive accelerometer that will be used as the primary example in this thesis. It is of note however, that any capacitive accelerometer will experience error regardless of design[9].

1.3 Sources of Error In Linear Accelerometers

Inertial navigation systems use a combination of measurements from accelerometers and gyroscopes that are integrated in order to determine a mechanized dead reckoning position. However, with any kind of measurement, there is associated error. In that way, it is readily known that systematic errors quickly integrate to large position errors over time.

As discussed, there are numerous sources of error in any kind of calculation. In the construction of a highly accurate sensor, all sources of error must be accounted for in order to attain the highest level of performance possible. For our discussion I will focus on the errors propagated from linear accelerometers. It is important to note that these errors arise from integrating twice in order to track position (where gyroscopes are only integrated once to track position).

1.3.1 Constant Bias

The bias of an accelerometer is the offset of its output signal from its true value. Because this is simply an offset, the constant bias error ϵ , when doubly integrated grows quadratically with time. Thus, the accumulated position error (s) may be described as function of time by,

$$s(t) = \epsilon \frac{t^2}{2} \quad (1.1)$$

Bias stability can be estimated by measuring the average long term output of an accelerometer when it is undergoing acceleration. However, it is of note that this measurement may be corrupted by gravity. As a result, it is necessary to know the precise orientation of the device with respect to a gravitational field in order to correctly measure the bias.

1.3.2 Hysteresis and Cross-Coupling Error

In accelerometers which operate in open-loop, there is a possibility that the proof-mass may deviate from the input axis of applied acceleration, and as a result develop off-axis errors. These errors are generally caused by manufacturing imperfections which give rise to non-orthogonality relative to the sensor axis. Such errors, defined as cross-coupling errors, cause large imperfections in acceleration measurements. Hysteresis is defined by the system not returning to its unperturbed state after a previous perturbation. An example of this is the presence of deflection or strain within a mechanical accelerometers spring after acceleration has been applied and removed. In the presence of hysteresis, the sensor will not be able to return to its null position resulting in bias instability. Open-loop systems, which rely heavily on measuring large displacement pick-offs, suffer most from these types of errors. Since the proof-mass moves over larger distances the spring rate must be higher to increase durability. However, this in turn increases effects of hysteresis as stiffer spring rates may lead to inconsistent behavior when the system is perturbed. Further, since the proof-mass moves freely on the input axis, there is increased susceptibility of cross-coupling error. As noted before, this type of error is minimal in closed loop systems. Here, the proof-mass moves very little and is maintained in almost constant position. This leads to lower spring rates, decreasing effects of hysteresis, and minimized contributions of cross coupling error since the position of the proof-mass is maintained.

1.3.3 Random Noise

All sensors experience some kind of random noise that may come from a number of sources. For example, electrical noise limits the resolution of inertial sensors where the output signal is very weak. Other sensors, like pendulous accelerometers, may experience noise due to mechanical instabilities. Further, accelerometers may also experience noise from the environment (usually vibrational noise). Vibration can be characterized by oscillatory motion at low frequencies. This type of motion can interact with scale factor and cross-coupling errors to produce oscillating sensor errors. Over time, this will average to zero but asymmetry or non-linearity of cross-coupling errors will result in a component of vibrational error that does not cancel over time, behaving like a bias. These type of

errors may be alleviated by low-pass filtering and increasing the operating frequency of the device[9].

1.3.4 Calibration Error

Calibration error refers to errors in scale factor, alignments, and linearities of the accelerometer. Scale factor error is characterized by a deviation of the input-output gradient of the sensor from unity following a unit conversion by the device. Accelerometer scale factor is proportional to the true force along the input measurement axis. Scale factor errors can occur from variable analog inputs, temperature sensitivities, mechanical drift, as well as optical drift. In an extremely dynamic environment, many of these factors will readily vary and cause large non-linearities and drift error. As a result, to achieve strategic performance, there is a need to dynamically account for scale factor error. This thesis will specifically concentrate on techniques used to compensate for and mitigate scale factor non-linearities in high precision capacitive accelerometers like the ZFA [11].

1.4 Current State of Accelerometer Errors

At present, the highest performance grade inertial navigation systems are rated accurate at a level of $\sim 100m/hr$. The current state of the art mechanical and optical sensors meet these performance levels. The Pendulous Integrating Gyroscopic Accelerometer (PIGA) has demonstrated high sensitivity and linearity over a large dynamic range. The device detects changes in acceleration by measuring displacements of a pendulous rotating mass and rebalancing the motion with gyroscopic forces exerted by a large rotating mass. These sensors are costly to manufacture, and have limited bandwidth. Long term stability of mechanical sensors is limited by the precision of machining, which can be highly expensive and difficult to reproduce. Over time, mechanical wear will also subsequently decrease the sensor lifetime. Optical gyroscopes, such as the IFOG (Interferometric Fiber Optic Gyroscopes) are now replacing their mechanical counterparts in many precision INS systems. The benefits of optical gyroscopes include compact design and high sensitivity over a dynamic range, making them extremely suitable for strap-down systems. However, their production calls for complicated mechanical design and machining, as well as largely increased computational

complexity.

For example, a given accelerometer may have some sensitivity σ_a (quoted velocity of random walk, in g/\sqrt{Hz}), and this leads to a position error in scaling which tends towards, $\Delta_x \sim \sigma_a T^{3/2}$. Similarly, a constant acceleration bias a leads to a position error that tends towards $\Delta_x \sim \frac{1}{2}aT^2$. Similarly, there are large errors from noise in gyroscope (specified by random walks in units of deg/\sqrt{Hz}) that rapidly accumulate over time. Other forms of error include temperature effects from the device, calibration inaccuracies, and bias instabilities (modeled by a second order random walk), and inaccurate compensation of local gravity (which accelerometers may not measure directly). The below table highlights these sources of error with their respective analytical error from integration[3]:

Table 1.1: Characteristic Performance Metrics in Highest Precision Navigation Systems

Sensor Metric	100 m/hr	5 m/hr (Expected in Future)
Accelerometer		
Sensitivity ($\mu g/\sqrt{Hz}$)	50	2
Bias Stability (μg)	1	< 0.1
Scale Factor Stability (ppm)	1-10	< 0.1
Size, Weight, and Power		
Volume (cc)	100-600	Similar
Power (W)	5-35	Similar
Weight (g)	300-2200	Similar

In the future, high accuracy navigation necessitates that INS achieves accuracy much lower than the current performance, especially in applications such as spacecraft design. The current goal of performance accuracy dictates that devices now achieve performance levels of $< 5m/hr$. For applications which last less than 1 hour, this would be on the same order of accuracy of GPS which is accurate up to $\sim 2m$. At current state, such systems (like those using atomic interferometry) are both costly and will not be readily available due to current technological constraints. Newer high performance sensors hope to help meet these needs by providing a low cost and stable device that is usable over large dynamic range, and meets these qualifications.

Chapter 2

Analysis of Linear Accelerometer

Mechanics

To further analyze a complex system like the ZFA, it is ideal to break down our system into the simplified model. As discussed, at its core the ZFA can be interpreted as a mass-spring system with capacitive re-balancing. The following chapter will encompass a formal analysis of a basic linear accelerometer, and derive different error propagations that the ZFA seeks to address.

The below figure shows a simplified version of a linear accelerometer (for now, dampening has been excluded for clarity). Here, the proof-mass is held by a spring, and a displacement pick-off determines its position. For simplicity in this analysis input acceleration from gravity has been excluded[4].




Figure 2.1: A basic example of a linear accelerometer. [11]

Newton's law states that if a mass m undergoes acceleration a , there must be a force acting on the mass in the form of $F = ma$. Similarly, if a spring, with a spring constant k , is extended by distance away from its equilibrium x_o at some distance Δx there must be a restoring force acting on the spring in the form $F = k\Delta x$. In the Figure 2.1 above, the spring is relaxed, and exerts no force on the proof-mass when it is at equilibrium. However, when the spring is extended the proof-mass experiences acceleration due to the restoring force of the spring of the form,

$$ma = -k\Delta x \quad (2.1)$$

Where m is the mass of the proof-mass in kg, a is the acceleration in m/s^2 , k is the spring constant in N/m and Δx is the extension of the spring. Indeed, solving for the acceleration, we find that this measurement is reduced to a measurement of spring extension, where

$$a = -\frac{k}{m}\Delta x \quad (2.2)$$

It is of note, that if the acceleration is reversed, this argument still applies, as the spring is compressed instead of extended. Equation 2.2 still reflects the relationship between displacement and acceleration. This argument, while trivial, forms the basis for many common accelerometer designs, including the ZFA. In this way any acceleration measurement can be construed as a measurement of linear displacement relative to a equilibrium or null position. While many designs differ in how they make this measurement, the same principle still applies.

2.1 Effects of Natural Frequency and Dampening

The crude example shown in the previous section does not consider effects of dampening, which further complicates this analysis. We already know that when a spring is extended beyond its equilibrium position and released, it will overshoot the equilibrium position and continuously oscillate back and forth until its motion is dampened by some other force and brings it to rest[4]. In particular, the mass spring system experiences oscillations of a characteristic natural frequency, defined by:

$$\omega_o = \sqrt{\frac{k}{m}} \quad (2.3)$$

We will introduce a linear force that dampens the oscillations of the proof-mass, and brings the proof mass back to equilibrium parameterized by a dampening coefficient μ (defined in units of s^{-1}) of the form $-\mu\dot{x}$. Our force equation then becomes,

$$F = -kx - \mu\dot{x} \quad (2.4)$$

The time evolution equation for this system then becomes,

$$\ddot{x} + \frac{\mu}{m}\dot{x} + \frac{k}{m}x = 0 \quad (2.5)$$

This result is the classic example of the damped harmonic oscillator. Here, we find a general solution of the form $x(t) = x_o e^{\lambda t}$, where λ and x_o are both constants. Applying our generalized solution

and factoring we find,

$$\lambda^2 + \frac{\mu}{m} + \frac{k}{m} = 0 \quad (2.6)$$

Solving for λ using the quadratic formula,

$$\lambda = \frac{-\frac{\mu}{m} \pm \sqrt{\frac{\mu^2}{m^2} - \frac{4k}{m}}}{2} \quad (2.7)$$

Factoring for a more descriptive solution we find,

$$\lambda = \frac{\mu}{2m} \pm i\omega_o \sqrt{1 - \frac{\mu^2}{4km}} \quad (2.8)$$

Leaving us with the solution,

$$x_{\pm}(t) = X_{o\pm} e^{-\frac{\mu}{2m}t} e^{\pm i\omega' t} \quad (2.9)$$

Where X_o is the initial amplitude of the system, and $\omega' = \omega_o \sqrt{1 - \frac{\mu^2}{4km}}$. Notice that due to the second exponential term (particularly the value of ω') the system may behave differently depending on its initial conditions. We will examine three cases in which we vary the parameters of ω' .

2.1.1 Critically Damped

In a critically damped system, $\omega' = 0$ or $\mu^2 = 4km$. Here, the solution takes the form,

$$x_{\pm}(t) = X_{o\pm} e^{-\frac{\mu}{2m}t} \quad (2.10)$$

So there is no oscillation, and the system will go to equilibrium in the shortest possible time (depending on the values of μ and m).

2.1.2 Over damped

The system is over-damped when ω' is imaginary, or when $4km < \mu^2$. The solution then becomes,

$$x_{\pm}(t) = X_{o\pm} e^{-\frac{\mu}{2m} \mp \omega' t} \quad (2.11)$$

Here, the system damps very slowly (depending on the magnitude of ω') and does not reach equilibrium (due to the first term in the exponential expression).

2.1.3 Under damped

Here, ω' is real ($4km > \mu^2$), so the solution becomes,

$$x_{\pm}(t) = X_{o\pm} e^{-\frac{\mu}{2m} t} e^{\pm i\omega' t} \quad (2.12)$$

Here, the system will experience true oscillations that will exponentially decrease until it approaches zero.

A robust accelerometer must be able to handle dynamic environments and attain accurate readings that do not reflect errors that are inherent within the system. Here, the type of dampening that a system experiences is extremely important. For an over damped system, any kind of force (gravity or otherwise) will drive the proof mass to exponentially approach (but not reach) its equilibrium at an exponentially decaying rate. This is not ideal as there is no basis for measurement because the decay time is so high. Similarly for under damped systems, the system exhibits oscillations about the equilibrium point, but never truly reaches its equilibrium. In an accelerometer, this can translate into displacement errors which would inhibit a proper acceleration measurement. Therefore, a critically damped system, one which reaches its equilibrium as fast as possible is ideal for accelerometer measurements, as any force imposed on the system would be nulled quickly and minimize displacement error. However, this does not describe the whole picture, as external forces in a dynamic environment may cause non-linearities in measurement.

2.2 Vibrational Effects

Vibrational effects, as well as effects from natural frequency, are best described by applied vibration.[4] For example, if we set our mass-spring system on a base which oscillates periodically to and from some position with amplitude χ_o , the resultant acceleration due to this motion is,

$$a(t) = -\omega^2 \chi_o \sin(\omega t) \quad (2.13)$$

where ω is defined in terms of the applied frequency f_v , and $\omega = 2\pi f_v$. From using equation 2.1, the motion of the proof mass can be defined in terms of Δx ,

$$\Delta x = \frac{m\chi_o}{k} \omega^2 \sin(\omega t) \quad (2.14)$$

From our model, our mass-spring accelerometer has been fixed to a base that is vibrating. Here, $\frac{m\chi_o}{k} \omega^2$ is the amplitude of the proof-mass within the accelerometer. From equation, it can easily be seen that the peak amplitude due to vibration varies as function of the vibration frequency squared, and linearly with the amplitude of displacement due to vibration. Note however, this result is obtained without considering the natural frequency of the spring-mass system[9]. When this is taken into account, we see that we will need further analysis in order to take into account the natural frequency of the system.

2.3 Non-Linearity Due to Resonance

Until now in this example we have ignored effects of both the natural frequency of the mass-spring system and dampening. For further analysis, we will assume that our dampening coefficient, μ is non-zero in addition to our applied period vibration.[4] As a result, force equation becomes,

$$F = -k(x - \chi) - m\mu\dot{x} \quad (2.15)$$

Where χ is the displacement of the table under vibration, and x is the displacement of the proof-mass. Knowing this, our time evolution equation becomes ¹,

$$\ddot{x} + \mu\dot{x} + \omega_o^2 x = \omega_o^2 \chi_o \cos(\omega t - \phi) \quad (2.16)$$

Here, ϕ refers to the phase lag of the oscillation. The above equation forms the well-known model for a driven, damped harmonic oscillator. The general form of this equation is,

$$m\ddot{x} + \mu\dot{x} + kx = F(t) \quad (2.17)$$

Where $F(t)$ represents the driving force imposed on the oscillator. The general solution of the driven damped harmonic oscillator equation is of two parts, which encompass the transient and steady state solution.

$$x(t) = x_{Transient} + x_{Steady-State} \quad (2.18)$$

The initial behavior of a damped, driven harmonic oscillator is quite complex. The transient solution describes the initial behavior of the system while the steady state solution describes the system after a long time. For simplicity, we will first look at the steady state solution of this system.

2.3.1 The Steady-State Solution

The steady-state form of oscillation, is where the spring-mass system settles down to a constant amplitude x_o , and the energy lost due to dampening per cycle is matched by the work done due to the driving force of vibration[6]. This suggests that we are to find a solution of the form,

$$x(t) = x_o \cos(\omega t - \phi) \quad (2.19)$$

¹This result is obtained by dividing our time evolution equation through by m and substituting for $\omega = \sqrt{\frac{k}{m}}$.

The commonly known solutions for both x_o and ϕ are defined by,

$$x_o = \frac{\omega_o^2 \chi_o}{\sqrt{(\omega_o^2 - \omega^2)^2 + \mu^2 \omega^2}}$$
$$\phi = \arctan\left(\frac{\mu \omega}{\omega_o^2 - \omega^2}\right)$$

It is important to look at the dependence of both amplitude and phase lag relative to the system's natural frequency. The best way to do this is graphically. The below figure shows ω/ω_o vs. x_o/χ_o and ω_o/ω vs. ϕ/π . The functions of x_o/χ_o and ϕ are plotted as functions of ω with different varying dampening coefficients. We vary μ/ω_o from 1, 1/2, 1/4, 1/8, and 1/6. The amount of damping in the system is decreased the amplitude of the response becomes progressively more prevalent at the natural frequency of oscillation of the system, a feature commonly known as resonance, and the natural frequency ω_o is known as the resonant frequency. Thus, a weakly damped oscillator where $\mu \ll \omega_o$ may be driven to a large amplitude by a relatively small external force which oscillates close to the resonant frequency. Note that the response of the oscillator is in phase ($\phi \sim 0$) when the external drive frequency is well below the resonant frequency. When at the resonant frequency, it is in phase quadrature ($\phi \sim \pi/2$) and in anti-phase for frequencies well above the resonant frequency. Here a common metric to describe the effect of amplification due to resonance is called the quality factor, Q_f , defined by ω_o/μ [8].




Figure 2.2: Figures showing effects of resonance. The graph on the left shows the effects of displacement amplitude as the driving frequency approaches the natural frequency. The plot on the right shows the effects of varying frequency about the natural frequency with phase[4].

Figure 2.2 shows that frequency responses due to resonance drastically effect the performance of an accelerometer readout as it takes very little force to assert a large frequency response under those conditions. As a result, This becomes non-reflective of the acceleration that the device is experiencing, and proves to be troublesome for accurate acceleration readings. In practice, accelerometers have "stops" built in to their design acting as safe guards to restrict proof-mass motion within the usable range of the device. In general, an applied frequency f that is less than the natural frequency f_N has little effect on the mass-spring response of the system. Similarly for an applied frequency much larger than the natural frequency any response is nearly independent of the applied frequency. For this reason, accelerometers do not operate near the natural frequency due to large non-linearities in output.

2.3.2 The Transient Solution

From our analysis of the steady-state solution, we both explained the behavior of our system and explained causes of non-linearities due to resonance. However, the steady state solution does not consider responses to forces that occur over very short time intervals. To model this effect, we

will first consider our damped harmonic oscillator, with a one dimensional degree of freedom and motion in the frame of the device.[6]




Figure 2.3: A damped harmonic oscillator with undergoing an impulse in direction y relative to the lab frame

Here, y represents the displacement of the frame, x represents the displacement of the internal mass, and z represents the relative displacement of the internal mass with respect to the frame. It can be easily seen that $x = z - y$ in this example. The relative motion then is taken as a measure of acceleration, \ddot{y} , of the frame. Knowing this, we can define our equation of motion as follows,

$$m\ddot{z} + \mu(\dot{z} - \dot{y}) + k(z - y) = 0 \quad (2.20)$$

Substituting $z = x + y$ and factoring we find,

$$\ddot{x} + \frac{\mu}{m}\dot{x} + \frac{k}{m}x = \ddot{y} \quad (2.21)$$

Instead of solving this solution directly, we will look at this equation under dimensionless parameters to give our analysis a wider range of applicability[6]. Defining the following:

$$\begin{aligned}
\gamma &= \ddot{y}/\ddot{y}_{max} \\
\tau &= t/T \\
\xi &= -kx/m\ddot{y}_{max} \\
D &= \mu/\sqrt{mk} = \mu/\mu_c \\
R &= 2\pi\sqrt{m/k}/T
\end{aligned} \tag{2.22}$$

where μ_c is the critical value of the dampening coefficient, T is the period of the acceleration pulse to be measured, $2\pi\sqrt{m/k}$ is the undamped period of the oscillator (reflective of the inverse of the oscillator's operating frequency), and \ddot{y} is the peak value of acceleration. Substituting these expressions we find our equation of motion to be,

$$\left(\frac{R}{2\pi}\right)^2 \ddot{\xi} + \frac{RD}{\pi} \dot{\xi} + \xi = \gamma \tag{2.23}$$

By inspection, for a high frequency accelerometer R is very small and the first two terms of Equation 2.23 drop out. Thus, the response ξ is linearly proportional to the dimensionless acceleration γ . This is ideal, as the acceleration is simply a function of the proof-mass's displacement with no added effects of oscillation. As R becomes larger, the first and second terms have more effect. The primary effect of the second term is introducing a time lag in between the response ξ and the acceleration γ . The first term tends to make the response ξ oscillate in value above and below the value of acceleration γ .

2.4 Conclusions

From our analysis it can be shown that a robust, strategic-grade accelerometer must meet specifications that minimize errors in output measurement. An ideal accelerometer will exhibit critical

dampening to react readily to dynamic input. It must also operate at high frequency, well above natural frequencies of dynamic inputs, such as vibration, to mitigate non-linearities in measurement in addition to alleviating oscillatory error due to transient responses. Modern accelerometer builds such the ZFA aim to achieve these specifications, not only to minimize measurement error but also to increase reliability and performance.

REDACTED COPY

Chapter 3

Classical Analysis of the ZFA

This chapter will go over how the Zero Force Accelerometer takes measurements, as well as describe a classical analysis of the system in question to prove the integrity of the device.

3.1 Overview of the ZFA

The ZFA is a strategic grade, low noise sensor currently in development at Charles Stark Draper Laboratory. At its core, the ZFA can be thought of as a mass on a spring system, kept in closed loop operation with capacitive re balancing. Here, the proof mass is suspended between the two capacitive plates by three springs at 120° intervals on the side of the proof mass. The system is kept under relatively high vacuum (about $1 * 10^{-5} Torr$) so there is relatively low dampening. The springs have very high spring rates in the direction parallel to the capacitive plates, but extremely low spring rates in the direction normal to the proof mass, allowing for free motion along the input axis, and decreasing the tendency towards cross coupling error. The movement of the proof-mass is tracked using two optical broad band ASE sources on either side of the apparatus. One optical source measures the absolute displacement of the proof-mass relative to a reference at the capacitive drive plate while the other measures the displacement of the proof-mass relative to its null¹.

¹The second optical readout is used to "lock-in" the movement of the proof-mass by virtue of a two-stage servo loop described in the preceding sections. The null reference position may be changed readily by the user of the device




Figure 3.1: A detailed structural diagram of the Zero-Force Accelerometer

3.2 PID Control

The ZFA measures acceleration by determining the amount of electrical force needed to bring the proof-mass back to its null. The proof mass's motion is nulled by virtue of a two-stage feedback loop system.

The ZFA utilizes various control systems in order to perform at its full capacity. It is then useful to go over the basics of control systems, mainly a proportional-integral-derivative controller, and their significance in maintaining closed loop operation in the ZFA.

PID controllers have been widely used in industry, science, and other control system applications. A PID controller calculates deviation from a measured output, or process variable, from a desired set point. The system attempts to compensate for this deviation by adjusting the control system outputs. The PID controller algorithm involves three separate time dependent parameters: the proportional, integral, and derivative control.

The PID control scheme involves three parameters, K_p , K_i , and K_d representing the proportional, integral and derivative gains. These three parameters are constants which compose a function of time that reflects a response to adjust a system process to a given set point. The adjusted process within the control loop is defined as a manipulated variable. In the case of the ZFA, the manipulated variable of this system is the displacement of the proof-mass relative to the null point. The control loop looks to minimize the error or drift of the manipulated variable from its set-point. Here our error $e(t)$ is the difference of the manipulated variable ($q(t)$) and the set point SP (i.e. $e(t) = SP - q(t)$). The proportional, integral, and derivative terms are summed to produce the output of the controller and act as a function of $e(t)$ [7]. Defining $q(t)$ as the controller response, we can define the PID control algorithm by,

$$q(t) = K_p e(t) + K_i \int_0^t e(\tau) d\tau + K_d \frac{d}{dt} e(t) \quad (3.1)$$

The below figure shows a block diagram for the PID Control loop.

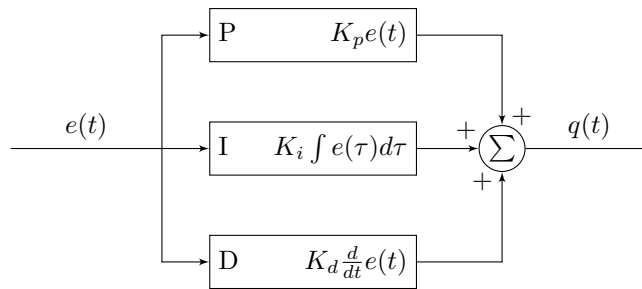


Figure 3.2: Diagram of Generic PID Control Loop

3.2.1 Proportional Term

The proportional term produces a value which is proportional to the current error value, shown in the equation below

$$P_{out} = K_p e(t) \quad (3.2)$$

Where P_{out} is the output parameter, K_p is the parameter gain, and $e(t)$ is the calculated error at a given time. Large changes in gain correspond to large changes in the control loop output. As a result, a large output gain may make the control system unstable. Similarly, a small change in gain corresponds to small changes in control loop output, which would cause little change to correct for the calculated error. As a result, the proportional gain provides the bulk contribution to the output.

3.2.2 Integral Term

With a standalone proportional controller, there is an operational steady-state error, since it is required that there must be a non-zero error, to drive it. The integral gain provides a correction for the accumulation of past errors over time. This parameter produces an accumulated offset which corrects for error that wasn't corrected for previously. In other words, this term corrects for residual steady-state error from proportional control. The integrated error is multiplied by an integral gain,

K_i , and added to the controller output.

$$I_{out} = K_i \int_0^t e(\tau) d\tau \quad (3.3)$$

3.2.3 Derivative Term

The derivative term is used to anticipate future system behavior by determining the change in error with time. The magnitude of the derivative term is determined by the derivative gain, K_d multiplied by the derivative of the error.

$$D_{out} = K_d \frac{d}{dt} e(t) \quad (3.4)$$

Notice that since we are taking the derivative of the error term, at high derivative gains the output is highly sensitive to signal noise. This effect may be mitigated by virtue of a noise reduction filter, such as a low pass filter.

3.3 Attaining Closed Loop Operation

The ZFA uses multiple iterations of control loop systems in order to attain closed loop operation. The follow section will describe briefly how this is done, and how this is used to calculate our scale factor.

3.3.1 The Optical Loop

Due to contractual obligations, the details of the optical loop will not be specified, however the basics will be explained. Here, the optical loop "locks" on the null point of the proof-mass (we will assume that this is the control system's set point). The optical loop then gives a displacement readout relative to the null of the form

$$\delta_{err}(t) = \delta_{SP} - \delta(t) \quad (3.5)$$

Here, $\delta(t)$ reflects displacement readout of the proof-mass as a function of time, δ_{SP} is the null point of the proof-mass, and δ_{err} reflects how far the proof-mass has deviated from its null position (interpreted as the error signal in our control loop system). For the purposes of this thesis, we will assume that the optical loop produces an extremely precise displacement readout with high resolution. This means that extremely small movements of the proof-mass² are accounted for.

3.3.2 The Mechanical Loop

The optical loop sends the displacement error signal to the mechanical loop, which provides physical feedback proportional depending on how much the proof mass has deviated from its null point. The signal is sent to an FPGA (Field programmable gate array). The FPGA interprets the error signal from the optical loop, and calculates a response in accordance to the PID control algorithm. A digital response signal is then sent to a DAC (digital to analog converter) which converts the digital signal (in bits) to a real voltage. The voltage signal is finally sent to a high voltage amplifier, which in turn instigates a voltage across the ZFA's capacitive plates. This physically moves the proof-mass, bringing it to null and the ZFA into closed loop operation. Here, define the response signal α as a function of the error signal from the optical loop using Equation 3.1³

$$\alpha(t) = K_p \delta_{error}(t) + K_i \int_0^t \delta_{error}(\tau) d\tau + K_d \frac{d}{dt} \delta_{error}(t) \quad (3.6)$$

The below figure shows a logic diagram of the ZFA in closed loop operation and the mechanical loop.

²While the specifics of how the optical loop works cannot be disclosed due to ITAR restrictions, the resolution of the optical loop response can be assumed to be on the order of nanometers

³The constants K_p, K_i and K_d are set arbitrarily, in accordance to the desired operating frequency, or loop bandwidth, of the system. The mechanical loop bandwidth is reflective of the time scale that a perturbation of the system goes to zero. Here, a higher loop bandwidth correlates to a higher rate of dampening due to perturbations, which is desirable for making measurements in steady-state. For the experiments highlighted in this thesis, the loop bandwidth for the ZFA was 120Hz.

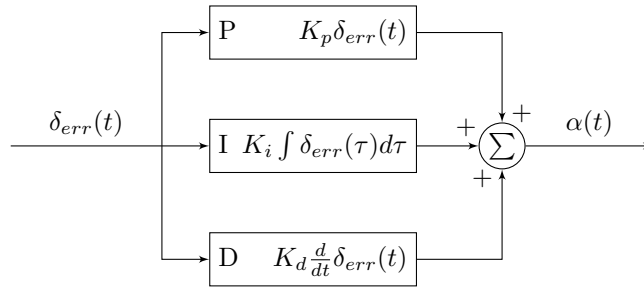


Figure 3.3: Diagram of Mechanical Loop in ZFA

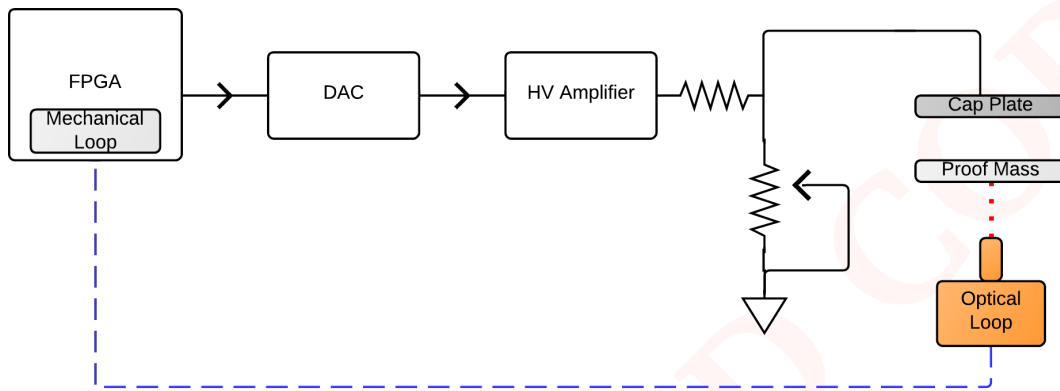


Figure 3.4: Block Diagram of ZFA Response Loop

It is of note that when the ZFA is in closed loop operation, the proof-mass moves very little due to the resolution of the optical loop. This will have a large impact on our analysis in proving the ZFA's measurement capabilities.

3.4 Analysis of ZFA Measurement Characteristics

The following section will go over the mechanics of the ZFA under closed loop operation, as well as explain the measurement characteristics of the device.

It is of note that the ZFA takes measurements of one of two plate displacements depending on what the direction of the proof-mass acceleration. For simplicity we will only consider acceleration in the direction of positive gravitational force for our measurements relative to our input axis. For now, we will assume the system is in steady-state, and has no transient effects from perturbations of small time-scales. From our force equation we find in the absence of gravity

$$F = F_{Spring} + F_{Cap} + F_{Damp} \quad (3.7)$$

Here, F_{spring} is the spring force, F_{Cap} is the capacitive drive force and F_{Damp} is the dampening force. It is of note that since the system is held in vacuum, the effects of dampening are very minimal, and can be ignored⁴. This leaves us with the time evolution equation,

$$m\ddot{x} = k\delta + \frac{\epsilon_o AV^2}{2d^2} \quad (3.8)$$

Here, the k is the spring rate, A is the area of the capacitive plates, V is the voltage across the stationary plate and the proof-mass, δ is the displacement of the proof-mass relative to the stationary capacitive plate, d is the distance between the stationary capacitive plate and the null point⁵. In steady-state operation, there are extremely small changes in the proof-mass position ($< \sim 1nm$) so any contribution by the displacement and mechanical loop can be discarded as well. This leaves us with

$$m\ddot{x} = \frac{\epsilon_o AV^2}{2d^2} \quad (3.9)$$

⁴Here, there is also dampening effects from the springs, which are very small and were ignored for clarity in our argument

⁵From the most current ZFA build specifications, $k = 35N/m$, $A = 2.76cm^2$, $d = 75\mu m$, and $m = 0.0016kg$ [5].

When there is zero acceleration of the proof-mass, there is still a voltage output which balances the gravitational force to hold the proof-mass at its null position. This can be defined as a reference voltage, V_{Ref} . Thus, the acceleration measurement is not related to the total voltage measured across the plates, but rather to the difference between the measured voltage and the reference voltage. For clarity in this thesis, V will refer to the output voltage defined by $V = V_{Measured} - V_{Ref}$. For our test setup, we will assume $V_{Ref} = 0$.

3.5 Correspondence to Input Angle

It is ideal to show that the voltage measured by the ZFA shows correspondence to the acceleration of the proof-mass. For our experimental setup, the ZFA is mounted within a vacuum chamber on a stationary index head which can rotate from 0 to 180 degrees. At the start of the test, the ZFA is oriented such that the proof-mass is orthogonal to the downward gravitational acceleration of the earth (at 0g). Measurements are taken over time, and averaged at subsequent angles relative to the initial angle. Figures 3.4 and 3.5 show a crude depiction of the experimental setup[9].




Figure 3.5: The ZFA mounted on the index head with 0g incident on the proof-mass




Figure 3.6: The ZFA mounted on the index head with 1g Incident on the proof-mass

Given this setup, there is a known input acceleration incident on the proof-mass of the form $\ddot{x} = g \sin \theta$, where g is the gravitational acceleration on earth and θ is measured input angle relative to the starting angle at 0g. Given this relation, and plugging in equation 3.9, we see that $V^2 \sim \sin \theta$. The below plot shows the correspondence between V^2 and $\sin \theta$.




Figure 3.7: Plot showing linear relationship between the sine of input angle θ and squared voltage over the capacitive plates V_{Cap}

The above plot shows a strong linear correspondence to the relationship described in Equation 3.9. This corroborates our assumption that V_{Cap}^2 is linearly proportional to the input acceleration. While it is useful to know that ZFA can make acceleration measurements, it is necessary to be able to convert the output of the sensor into relevant, units of acceleration.

Chapter 4

The Scale Factor

This chapter will go over the basics of a scale factor, and how it is calculated within the context of a capacitive accelerometer. It will then motivate the idea of an in-situ scale factor calculation, and go over the algorithm of how this is done within the context of a capacitive accelerometer like the ZFA. Finally, this chapter will address the problems found within the algorithm, and possible solutions.

4.1 Definition and Motivation

A scale factor, by definition, is the ratio of a sensor's output to its mechanical input.[10] In the context of a capacitive accelerometer in closed loop, this is the ratio of the measured voltage squared by the sensor and gravitational acceleration of the form, g/V^2 . The scale factor is reflective of both the structural make-up of the sensor. In the case of the capacitive accelerometer, the scale factor, denoted as SF , may be defined in terms of the force derived in Equation 3.9. The output acceleration is therefore defined by,

$$\ddot{x}_{out} = SF * V^2 \quad (4.1)$$

Here, the scale factor directly translates the output voltage (which, when measured is in the form of digital LSB Counts) into units of acceleration and is of the form,

$$SF = \frac{\epsilon_o A G_{DAC}^2}{2md^2} \quad (4.2)$$

Here, G_{DAC} is an arbitrary constant set by the DAC that translates a digital LSB unit reading into a voltage, of the form V/LSB_{Counts} ¹. It can be easily seen that inaccuracies in the scale factor will result in large propagated errors for the measured output acceleration. Thus, it is extremely important to measure the scale factor at the highest precision in order maintain the integrity and sensitivity of the sensor.

4.2 Statically Calculated Scale Factor

In Equation 4.1 we defined a relationship to translate the output voltage of the ZFA into acceleration. Here, the scale factor can be easily calculated given the structural make-up of the device. Table 4.1 lists the relevant structural quantities of the ZFA that was used for experimentation in this thesis[5].²

Table 4.1: Measured Structural Components of ZFA [5]

Quantity	Value
Plate to Proof-Mass Gap (d)	70.51 μm
Capacitive Plate Area (A)	2.76 cm^2
Proof-Mass Mass (m)	0.0016 kg

Knowing these values, the scale factor can be easily calculated. The integrity of the scale factor can be measured by observing the correspondence between a well known input acceleration and the measured output acceleration (i.e. applying the scale factor to the voltage output). Figure 4.1 shows the results of the plotted acceleration output versus a known input acceleration. Figure 4.2 shows the propagated error (ppm) relative to the ideal output acceleration.

¹In practice the scale factor is applied digitally via the FPGA. The reason for this is to provide the utility to improve resolution using bit handling, which is beyond the scope of this thesis.

²These values are based off of both build notes from the ZFA as well as physical measurements taken of the device



Figure 4.1: Output acceleration versus known input acceleration. The blue line shows the measured acceleration while the red line shows the ideal acceleration output. An ideal acceleration output shows one-to-one correspondence to the known input acceleration.




Figure 4.2: Rectified Force Error (ppm) versus Input Acceleration (g) using Single Dithered Scale Factor. This plot shows the coherence of the output acceleration from the idealized output of the sensor (i.e. 1-to-1 ratio between input and output acceleration). The error propagation here constitutes 21.5%/g error in scale factor over a range of 1g.

4.2.1 Discussion of Results

The above graphs show clear deviation of the acceleration output which was calculated to be³ 73,075 ppm/g. While this error profile does not reflect only scale factor error (bias instability, manufacturing failures, etc.) this clearly does not meet the needed stability requirements for a strategic grade sensor highlighted in Table 1.1. In this example, the scale factor calculation ignores many factors which produce non-linear and incorrect results. Temperature defects, varying analog inputs (e.g. vibration), mechanical drift, change in gain in analog electronics, and optical drift can change the state of the device and produce outputs that are inconsistent to the calculated scale factor. In order to improve upon this, the scale factor must be continually calculated and accounted for.

4.3 In-Situ Scale Factor Algorithm

This section will go over the motivation for an in-situ scale factor and the method in which it is calculated. It will also encompass a formal analysis of the ZFA when the algorithm is implemented, and provide insight into any problems found within the algorithm.

4.3.1 Motivation

As highlighted in the previous section, a static scale factor calculation is ineffective in generating sensor data with very high accuracy. Taking a step back, we note that the scale factor is defined by,

$$SF = \frac{\epsilon_o A}{2md^2} \quad (4.3)$$

The scale factor is inversely dependent on d^2 . This means that for a small change in d , the amount of voltage needed to null the proof-mass may change, and as a result the scale factor may change. Many factors may lead to this result. For example, temperature changes, may cause the system to slightly

³The result highlighted is an average error propagation over five tests of the same nature. Further, this error profile was limited as the prototype used for testing is only stable up to 0.35g. Further iterations and testing are needed to produce results that reflect up to $\pm 1g$ of measurements

contract or expand and result in a different plate to proof-mass gap distance than was originally calculated. Mechanical drift (caused by misalignments or instrument wear), varying analog inputs, and optical drift error may also contribute to this. In dynamic environments, these factors become extremely prevalent. As a result, there is a need to dynamically calibrate and calculate the scale factor to mitigate error in sensor outputs[?].

4.3.2 Dithering

In order to understand the idea of a dynamically calculated scale factor, we must first introduce the idea of a dither. Dithering is a commonly used practice in engineering to calibrate sensors. The process involves forcing the proof-mass to move at a high frequency, usually sinusoidally, slightly above and below its null position. The distance is tracked and integrated over one or a multiple of the period of the dither signal. Since a pure sine function integrates to zero over one cycle, a non-zero integration denotes a perturbation in the system. Thus, by applying a dither information can be gained through analyzing the amplitudes of the motion of the dithered proof-mass over time.[1]




Figure 4.3: A trivial example of a dithered system. A sine wave is integrated over one cycle. The blue line represents the pure dither signal, while the pink line represents the motion of a proof-mass in a perturb system. Clearly, the pink line does not integrate to zero after one period, so there is information to be gained (and compensated for) from the perturbed system.

In the case of the ZFA, the dither signal is applied in the mechanical loop, and is summed with the displacement error signal gained from the optical loop. Because the mechanical loop is continuously trying to null the proof-mass, the applied dither signal forces the proof-mass to move sinusoidally about its null position. Thus, by integrating both the displacement and the output force response of the mechanical loop, information can be gained on how the system has changed. This leads to being able to calculate the scale factor dynamically, and perform an in-situ calibration as prescribed[1].

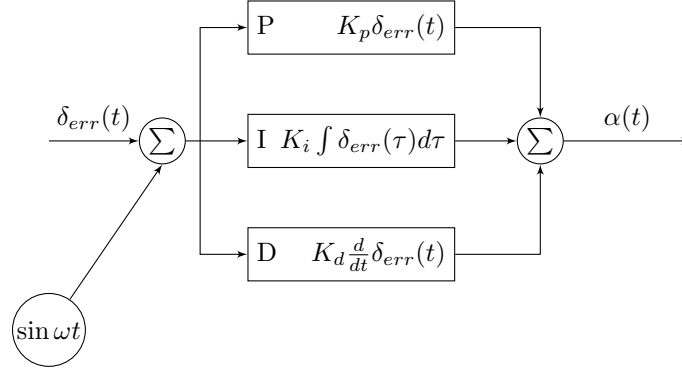


Figure 4.4: Diagram of Mechanical Loop in ZFA with the dither

4.3.3 Lock-In Integration

In order to calculate the values of displacement x and the response force due to the mechanical loop α , we must perform a lock-in integration. Assuming the dither is engaged at a frequency ω , our equations for x and α are as follows[10],

$$x(t) = \delta \sin(\omega t + \phi) \quad (4.4)$$

$$\alpha(t) = \alpha \sin(\omega t + \phi) \quad (4.5)$$

Looking at the equation for displacement, we multiply Equation 4.4 by a sine wave to mix the information at the dither frequency and DC, leaving,

$$x(t) \sin(\omega t) = \delta \sin(\omega t + \phi) \sin \omega t \quad (4.6)$$

Using the trigonometric identity, $\sin(a) \sin(b) = \frac{1}{2}[\cos(a - b) - \cos(a + b)]$ we can expand this expression and simplify to find,

$$x(t) \sin(\omega t) = \frac{\delta}{2} (\cos \phi - \cos(2\omega t + \phi)) \quad (4.7)$$

From this expression, integrating over an arbitrary number of dither periods (or cycles) will reject the in-phase component of the content at the other frequency.

$$\begin{aligned} I_o &= \int_{t=0}^{N/f_o} x(t) \sin(2\pi f_o t) dt = \int_{t=0}^{N/f_o} \frac{\delta}{2} (\cos \phi - \cos(4\pi f_o t + \phi)) dt \\ &= \frac{\delta}{2} \left(\frac{N}{f_o} \cos(\phi - 1/(4\pi f_o)) [-\sin(4\pi f_o \frac{N}{f_o} + \phi_o) + \sin(\phi)] \right) \end{aligned}$$

However, since N is an integer, $\sin(4\pi f_o \frac{N}{f_o} + \phi_o) = \sin \phi$, thus leaving us with

$$I_x = \frac{\delta N}{2f_o} \cos \phi = \int_{t=0}^{N/f_o} x(t) \sin(2\pi f_o t) dt \quad (4.8)$$

Similarly, the quadrature component can be defined as,

$$Q_x = \frac{\delta N}{2f_o} \sin \phi = \int_{t=0}^{N/f_o} x(t) \cos(2\pi f_o t) dt \quad (4.9)$$

Taking the magnitude, we find that,

$$\sqrt{I_x^2 + Q_x^2} = \frac{\delta N}{2f_o} \quad (4.10)$$

Rearranging, we further find that,

$$\delta = \frac{2f_o}{N} \sqrt{\left[\int_{t=0}^{N/f_o} x(t) \cos(2\pi f_o t) dt \right]^2 + \left[\int_{t=0}^{N/f_o} x(t) \sin(2\pi f_o t) dt \right]^2} \quad (4.11)$$

Similarly,

$$\alpha = \frac{2f_o}{N} \sqrt{\left[\int_{t=0}^{N/f_o} \alpha(t) \cos(2\pi f_o t) dt \right]^2 + \left[\int_{t=0}^{N/f_o} \alpha(t) \sin(2\pi f_o t) dt \right]^2} \quad (4.12)$$

Thus, by following this prescription we are able to get the amplitudes of the dither signal[10]. It is important to note that this method is phase independent. For the sake of clarity, I will exclude the phase, ϕ in our derivations of scale factor.

4.3.4 Analysis of Dithered Scale Factor

Given the basic idea of the scale factor, it is important to perform a classical analysis in order to derive the scale factor and expose any sensitivities the system may have.[2] We start again with our force equation, and assume the proof-mass is dithered at a high frequency well above its natural frequency (in other words, ignoring the spring rate term).

$$F = \frac{\epsilon_o A V^2}{2d^2} \quad (4.13)$$

Now given that the proof-mass is dithered with a sinusoidal signal, the force measured reflects small movements of the proof-mass as well as the response of the mechanical loop due to the dither. This results in the following expression,

$$F(t) = \epsilon_o A G_{DAC}^2 \frac{(\sqrt{V_{DC}^2} + \alpha \sin \omega t)^2}{2(d - \delta \sin \omega t)^2} \quad (4.14)$$

Here, $d - \delta$ reflects the displacement of the proof-mass relative to the gap d (where δ is the displacement) and G_{DAC} (of the form (V/LSB)) is an arbitrary constant within the DAC that converts LSB counts into voltage units. The output voltage V has been linearized over the DC voltage measured over the plate V_{DC} and the mechanical loop response due to the dither, α .⁴ Since the dither is sinusoidal, both δ and α vary with the dither frequency ω . Here, V_{DC} can be interpreted as reflective of the "true" acceleration of the proof-mass, the value of $\alpha \sin \omega t$ represents corrections made to the voltage output.⁵ Now, factoring our expression in 4.5 we find,

$$F(t) = \frac{\epsilon_o A G_{DAC}^2 (V_{DC}^2 + \alpha \sin \omega t)}{2d^2 (1 - \frac{\delta}{d} \sin \omega t)^2} \quad (4.15)$$

⁴Here, α is a function of the displacement error (defined in Equation 3.5) relative to the proof-mass's null position (denoted in Equation 3.6).

⁵In steady-state, the term $\alpha \sin \omega t$ averages to zero over time, so any non-zero value of α is reflective of compensation by the algorithm

From Equation 4.3 we know, $SF = \frac{\epsilon_o AG_{DAC}^2}{2md^2}$. We also know that $\frac{\delta}{d} \ll 1$. Making that substitution, and using binomial expansion about $\frac{\delta}{d}$, we find that:

$$\ddot{x}_{Total} = SF(V_{DC}^2 + \alpha \sin \omega t)(1 + 2(\frac{\delta}{d}) \sin \omega t + 3(\frac{\delta}{d})^2 \sin^2 \omega t + \dots) \quad (4.16)$$

Here, \ddot{x}_{Total} reflects the total acceleration experienced by the proof-mass. Expanding and grouping terms,

$$\begin{aligned} \ddot{x}_{Total} = SF & [V_{DC}^2(1 + \frac{3}{2}(\frac{\delta}{d})^2 - \frac{\alpha}{2}(\frac{\delta}{d})^2) + [V_{DC}^2 * 2(\frac{\delta}{d}) + \alpha[1 + \frac{9}{4}(\frac{\delta}{d})^2]] \sin \omega t \\ & - [V_{DC}^2(\frac{3}{2})(\frac{\delta}{d})^2 - \frac{\alpha}{2}(\frac{\delta}{d})] \cos 2\omega t - [(\frac{3}{4})\alpha(\frac{\delta}{d})^2 \sin 3\omega t] \end{aligned} \quad (4.17)$$

After simplifying terms, we can perform our lock-in integration as explained in 4.3.2 and demodulate about ω leaving us with the expression,

$$\ddot{x}_{Total} = SF[\alpha[1 + \frac{9}{4}(\frac{\delta}{d})^2] - V_{DC}^2 * 2\frac{\delta}{d}] \quad (4.18)$$

Looking at the left side of the equation, we can define \ddot{x}_{Total} as a summation of the input gravitational acceleration \ddot{x}_{Input} that the proof-mass experiences, and the acceleration due to the dither, \ddot{x}_{Dither} where,

$$\ddot{x}_{Total} = \ddot{x}_{Input} + \ddot{x}_{Dither} \quad (4.19)$$

Since \ddot{x}_{Dither} reflects the acceleration of the proof-mass due to the dither, we may substitute in the equation for displacement of the proof-mass (shown in the denominator of Equation 4.5) of the form,

$$x_{Dither}(t) = \delta \sin \omega t \quad (4.20)$$

Differentiating with respect to time, and substituting Equation 4.11 into Equation 4.10 we find the equation for total acceleration of the proof-mass to be,

$$\ddot{x}_{total} = \ddot{x}_{Input} - \delta \omega^2 \sin \omega t \quad (4.21)$$

For this analysis, we will assume that there is no input acceleration imposed on the proof-mass (i.e. $\ddot{x}_{Input} \rightarrow 0$). Thus, equating 4.13 with 4.9 we get the expression,

$$-\delta\omega^2 \sin \omega t = SF[\alpha[1 + \frac{9}{4}(\frac{\delta}{d})^2] - V_{DC}^2 * 2\frac{\delta}{d}] \quad (4.22)$$

Rearranging the equation and solving for SF , we find

$$SF = \frac{-\omega^2}{\frac{\alpha}{\delta}(1 + \frac{9}{4}(\frac{\delta}{d})^2) - \frac{2}{d}V_{DC}^2} \quad (4.23)$$

Equation 4.14 shows the equation for scale factor considering the small deviations in proof-mass displacement due to the dither. The most important thing to note about this expression is that there is a dependency on V_{DC} . As a result, the Scale Factor, as calculated by this algorithm, is g-dependent, and will change as the input acceleration changes.

The below plots show the scale factor deviation (in parts-per-million) as a function of input acceleration, as well as the correspondence between the input and output acceleration of the ZFA when applying the algorithm as it stands.




Figure 4.5: Scale Factor deviation (ppm) versus input acceleration (g) with associated error. Scale Factor Data was averaged for each point and compared with the initial scale factor reading to show propagated error. The trend shows an error propagation of 61,529 ppm/g. This error is less than that of the statically applied scale factor but still well above the approved stability for a strategic grade sensor.




Figure 4.7: Rectified Force Error (ppm) versus Input Acceleration (g) using Single Dithered Scale Factor. The error propagation here constitutes 15%/g error in Scale Factor over a range of 1g.




Figure 4.6: Input Acceleration (g) versus output acceleration (g) taking into account applied single dither scale factor (pink line), and statically calculated scale factor (blue line). Voltage and Scale factor were measured at each point, and then applied according to Equation 4.1 to find output acceleration. As expected, the acceleration curve shows large deviations from the idealized output of the sensor (red line).

4.3.5 Discussion of Results

The results in the above section clearly corroborate the analysis made in 4.3.3 that concluded the dithered scale factor is g -dependent. This is not ideal, as there would be a different scale factor for varying input g -levels. This would cause large, g -dependent error propagation in scale factor producing inaccurate acceleration measurements. While our analysis dictates that the scale factor must be calculated dynamically, because of this large g -dependent error the scale factor algorithm must be changed.

Chapter 5

Contributions: Sequential Two-Dithered Approach

This chapter will highlight the contributions I helped make to solve the issue of a g-dependent scale factor by making measurements using two dither signals and combining them to produce a corrected scale factor. The chapter will first motivate the idea and derive an alternative to the g-dependent scale factor highlighted in Chapter 6, and then discuss the results obtained using the newly adapted algorithm. Finally it will discuss the conclusions made through this analysis, and future prospects for research.

5.1 Motivation

As highlighted in Chapter 4, the in-situ calculation using a single dither produces a g-dependent Scale Factor. As a result, it is necessary to rectify the current algorithm in order to take into account this large source of error. Given our expression for scale factor in Equation 4.14, it can be seen that for any given input acceleration the quantity V_{DC} , which reflects unperturbed input acceleration of the proof-mass, is unchanged[2]. Further, V_{DC} is independent of the other parameters in the expression. So for any number of measurements, V_{DC} will stay the same for any measurements

taken at the same input acceleration. The FPGA gives us the utility to be able to change certain parameters, in particular the dither frequency, on the fly. Knowing this, we are able to exploit this fact by making successive measurements of scale factor, and combining these measurements in order to compensate for g-dependency.

5.2 Using Two Dither Measurements

Starting with Equation 4.14, we can form the following expression,

$$\frac{\alpha}{\delta} = -\frac{\omega^2}{SF} \left(1 - \frac{9}{4} \left(\frac{\delta}{d} \right)^2 \right) + \frac{2}{d} V_{DC}^2 \quad (5.1)$$

By observation, δ and α (a function of δ) are both well known and measurable. Further V_{DC} is independent of these quantities. Since it is possible to change the dither frequency, ω arbitrarily, we can make two measurements of scale factor with same input acceleration at different dither frequencies, ω_1 and ω_2 .

$$\frac{\alpha_1}{\delta_1} = -\frac{\omega_1^2}{SF} \left(1 - \frac{9}{4} \left(\frac{\delta_1}{d} \right)^2 \right) + \frac{2}{d} V_{DC}^2 \quad (5.2)$$

$$\frac{\alpha_2}{\delta_2} = -\frac{\omega_2^2}{SF} \left(1 - \frac{9}{4} \left(\frac{\delta_2}{d} \right)^2 \right) + \frac{2}{d} V_{DC}^2 \quad (5.3)$$

At different dither frequencies, the displacement measurement is different due to the fact that the proof-mass is oscillating at a different rate (which will also result in a different mechanical loop response). Here, α_1 and δ_1 represent the phase and mechanical loop response at the first frequency and α_2 and δ_2 represent those same quantities at the second frequency measured [2]. By subtracting 5.3 and 5.4, we can easily eliminate the V_{DC} term, yielding the expression,

$$\frac{\alpha_1}{\delta_1} - \frac{\alpha_2}{\delta_2} = \frac{1}{SF} \left(\omega_2^2 - \omega_1^2 - \frac{9}{4} \left(\frac{\delta_1^2}{d^2} - \frac{\delta_2^2}{d^2} \right) \right) \quad (5.4)$$

Here, $\frac{\delta_1^2}{d^2} - \frac{\delta_2^2}{d^2} \ll \omega_2^2 - \omega_1^2$ so the rightmost term in the parenthesis may be discarded. Solving for SF , we yield the expression for the corrected Scale Factor,

$$SF_{corrected} = \frac{\omega_2^2 - \omega_1^2}{\frac{\alpha_1}{\delta_1} - \frac{\alpha_2}{\delta_2}} \quad (5.5)$$

Equation 5.5 shows that by calculating the scale factor at different times and different dither frequencies, and then combining the results we are able to negate the dependence on input acceleration. Further, it is of note that this expression is a simply a function of the displacement of the proof-mass. As a result, it is easily implemented in software, and is relatively lightweight in terms of CPU usage and space.

5.3 Choosing the Two Frequencies

In order to make an accurate scale factor calculation using the two-dithered approach, it is important to choose two frequencies that are not close to any resonant frequency of the ZFA, or harmonics of each other. To make sure of this, a Power Spectral Density measurement (shown in Figure 5.1) was taken of the device under closed loop operation. Table 5.1 further lists notable resonant frequencies and their causes under closed loop operation.



Figure 5.1: PSD Plot of capacitive plate voltage response (mV) for a frequency range of 0 to 400Hz

Table 5.1: Measured Resonant Frequencies of ZFA[2]

Resonant Frequency (Hz)	Cause
23	Proof-Mass Natural Frequency
60	Utility Frequency
80-100	External Vibrations from Environment
120	Mechanical Loop Operation

For this thesis, the dither frequencies used in testing the the two-dithered scale factor algorithm were 400Hz and 250Hz. This pair does not have any shared harmonics¹ and lie well above natural resonant frequencies of the ZFA.²

5.4 Implementing the Algorithm

In order to make a continuous measurement of the scale factor, the measurements at different frequencies must be made sequentially. The algorithm is as follows. The first measurement is made

¹There is a shared harmonic at 2000Hz, but this is well above our operating capacity, and has very little effect on the measurements taken by the device

²In practice, both the operating frequencies and the dither frequencies still lie somewhat close to the other pronounced resonances. The prototype used in the testing for this thesis was limited to operating frequencies \leq 450Hz. However later iterations will be capable of achieving $>$ 1KHz operation.

at the first specified frequency for a given time period, τ , then a second one is measured immediately after at the second frequency over the same time period. At the end of the second time period, the measurements are combined in the form of Equation 5.5 and outputted by the FPGA.³ To show confidence in this method, measurements were taken at two different dither frequencies with the single dithered algorithm, and the corrected scale factor was calculated via post-processing.⁴




Figure 5.2: Plot showing scale factor calculated at 300Hz (blue line) and 175Hz (pink line) and the combined scale factor (red line). Clearly, by applying equation 5.5 to the calculated scale factors using a single dither, we can eliminate the g-dependency.

From the above figure, we can clearly see that the algorithm compensates for the g-dependency in scale factor for a single dithered frequency by virtue of applying the corrected scale factor equation on two different runs. It is ideal then to look at this algorithm when it is fully integrated with the FPGA. The following charts show the scale factor error (ppm) over g-level, and the correspondence of input to output acceleration using the fully integrated sequential scale factor algorithm.

³This idea may be expanded upon in further studies as exact way of how this works in considering how the LSB bits are used to calculate the scale factor is beyond the scope of this thesis. For now, we'll just assume that the algorithm simply outputs a corrected scale factor after 2τ .

⁴Figure 5.2 was mainly a qualitative test to show that the algorithm does indeed correlate to the theory. While the proceeding tests show the integrity of the algorithm numerically, it is important to note that the method works via manual post-processing.



Figure 5.3: Sequential Two-Dither Scale Factor Error (ppm) vs. Input Acceleration (g). This plot shows as little to no trending in scale factor, but has non-linear outputs that could be caused by other factors. The scale factor error is greatly improved, and equates to about ± 311.43 ppm/g.




Figure 5.4: Correspondence between Input Acceleration (g) and Output Acceleration (g). Compared to the static and single dithered scale factors, the output acceleration calculated with the sequential scale factor shows strong correspondence to the input acceleration.



Figure 5.5: As highlighted in Figure 5.3, the scale factor algorithm shows strong correspondence to the idealized output acceleration measurement. The propagated error (aside from the outliers at 0.31g) gives 7.31%/g error in scale factor up to 1g.

5.5 Discussion of Results

From the above plots, it can be clearly seen that the scale factor error due to g-dependency has been greatly reduced. Particularly, Figure 5.2 shows little to know trending of Scale Factor over increasing input acceleration. This equates to almost over an order of magnitude improvement over both the statically calculated scale factor, and the single dithered scale factor. Further, the acceleration output using the sequential scale factor, in comparison with the other methods used in this thesis, shows a strong correspondence to the input acceleration, and equates to about half as much scale factor error over a range of 1g.

Chapter 6

Conclusion

6.1 Reflections on Work

Though the ZFA seems like a very simplistic inertial sensor, there are many factors in both its design and measurement methods which complicate its ability to make high accuracy measurements. From our analysis, we have successfully concluded that the current algorithm using a single dither to dynamically calculate scale factor results in large g-dependent error. Using two measurements at different frequencies, I was able to mitigate the effects of g-dependent error in the scale factor calculation and improve the scale factor error by almost an order of magnitude.

While the analysis of the scale factor was relatively simple, the implementation and testing of the device proved to be fairly difficult. Since the ZFA is a prototype accelerometer, there were many difficulties that impeded testing that included mechanical failures, various software bugs, and tuning and sensitivity issues. This resulted in multiple iterations of noisy and inconclusive data. Thankfully, the data shown above is valid, corroborates my hypothesis, and my proposed change to the scale factor algorithm.

6.2 Known Problems and Future Studies

The biggest problem with this algorithm as it stands is that it takes very long to make a calibration measurement. The sequential two-dithered algorithm requires that two measurements be made over a given time interval. While this mitigates effects from g-sensitive measurements, in dynamic environments where conditions of the system are constantly changing this can prove to be very problematic. In the future, a more efficient, time-effective methods will be explored and implemented.¹

Furthermore, there were many issues in testing that could be improved in the future. For example, the device I did my testing on was only stable up to a given voltage, so all results shown in this thesis were only reflective from an input acceleration range of 0g to 0.34g. Further, because of mechanical restrictions, the device used was not able to operate above 450Hz. We had used the upper frequency ranges for our dither, and held the mechanical loop at 120Hz. However, this is still very close to different vibrational and other analog inputs which may have frequency ranges from 70Hz to 110Hz. As a result, the data taken may reflect unwanted inputs from the outside environment. In future studies, more robust builds of the ZFA will be made that are capable of operating in frequencies greater than 1 Khz, and capable of operating up to and past 1g input acceleration.

There were many factors that contributed to measurement error that were not considered in this thesis. I did not take into account proof-mass movement or misalignments. I have assumed that the two-stage control loop system is fairly robust, however small proof-mass movements due to bias or other defects may cause non-linearities in both scale factor and sensor output. This was readily observed on all of the results taken in this thesis, as none of the data displayed showed a strong linear trend but rather was generally sporadic. Further, temperature effects, bias instability, as well as optical and mechanical drift error were not considered in this research project. Further

¹A proposed method would be to impose two dither signals on the proof-mass, and demodulate over each respective one to obtain two separate measurements. However, there are many other complications (that are beyond the scope of this thesis) that arise from this, so further study is needed.

studies could encompass looking at temperature profiles of the ZFA, mitigating sources of bias, and signal-noise analysis to the improve resolution and accuracy of the device.

While I and the rest of the Draper research team were fairly successful in mitigating effects of a g-dependent scale factor, this is only a small piece in a larger problem of creating a sensor that will be able to perform in dynamic environments at strategic grade capacity.

REDACTED COPY

Bibliography

- [1] Paritosh Banerjee Amitava Bose, Somnath Puri. *Wilhelmina D. Kroll Modern Inertial Sensors and Systems*, volume 2nd. PHI Learning Private Limited, 2008.
- [2] Mark Bowser. Scale factor transfer function. CSDL Internal Document, 2013.
- [3] David Butts. *Light pulse atom interferometry at short interrogation times for inertial navigation*. PhD thesis, Massachusetts Institute of Technology, 2012.
- [4] Richard Fitzpatrick. *Oscillations and waves*.
- [5] Donald Fyler. Don fyler zfa build notes. CSDL Internal Document, 2013.
- [6] Samuel Levy. Response of accelerometers to transient accelerations. *Journal of Research of the National Bureau of Standards*.
- [7] David Sellers. An overview of proportional plus integral plus derivative control and suggestions for its successful implementation. Technical report, Portland Energy Conservation Inc., 2001.
- [8] John R. Taylor. *Classical Mechanics*. University Science Books, 2005.
- [9] David Titterton. *Strapdown Inertial Navigation Technology*. American Institute of Aeronautics, Institution of Engineering and Technology, 2004.
- [10] Michael Tomaino-Iannucci. Zfa in-situ scale factor calibration. CSDL Internal Document, February 2013.

- [11] Oliver Woodman. An introduction to inertial navigation. Technical report, University of Cambridge Computer Laboratory, 2007.

REDACTED COPY