University of Louisville

ThinkIR: The University of Louisville's Institutional Repository

Electronic Theses and Dissertations

5-2011

High dynamic range imaging for the detection of motion.

Jeffrey Robert Hay
University of Louisville

Follow this and additional works at: https://ir.library.louisville.edu/etd

Part of the Physics Commons

HIGH DYNAMIC RANGE IMAGING FOR THE DETECTION OF MOTION

By

Jeffrey Robert Hay
B.S., University of Louisville, 2005
M.S., University of Louisville, 2007

A Dissertation
Submitted to the Faculty of the
College of Arts and Sciences of the University of Louisville
in Partial Fulfillment of the Requirements
for the Degree of

Doctor of Philosophy

Department of Physics and Astronomy
University of Louisville
Louisville, Kentucky

May 2011

HIGH DYNAMIC RANGE IMAGING FOR THE DETECTION

OF MOTION

By

Jeffrey Robert Hay
B.S., University of Louisville, 2005
M.S., University of Louisville, 2007

A Dissertation Approved on

April 8, 2011

by the following Dissertation Committee:

_____
John Kielkopf

_____
Gamini Sumanasekera

_____
James Lauroesch

_____
John Naber

DEDICATION

To my wife and children.

ACKNOWLEDGMENTS

I would like to thank all the un-named friends and family who helped me get to this point. There are so many people that play a role throughout your life and I am fortunate and grateful to have so many supportive and caring people around me.

I would like to thank Dr. Adam Willitsford for being a great friend and providing valuable feedback and a great collaboration. Having someone to bounce ideas off of was something much needed. I could always turn to him for great advice.

Thanks to my friends, Eston, Corey, Kyle and Matt, who I could always count on, even to this day. They helped create an environment that valued the importance of knowledge and helped shaped who I am today.

Thank you to Dru and Bill Kemp for being so supportive. They were always there when needed and for that I am very grateful. To Dru I am thankful for all the excellent advice along the way. It always seemed to come at a much needed time.

I would like to thank my parents Pam and Terry Earehart who have supported me every step of the way, offering me guidance and encouragement throughout my entire educational experience. My mother especially who always took the necessary time to provide enrichment in my life. My father Robert Hay,

ABSTRACT

HIGH DYNAMIC RANGE IMAGING FOR THE DETECTION OF MOTION

Jeffrey Robert Hay

April 8, 2011

High dynamic range imaging involves imaging at a bit depth higher than the typical 8-12 bits offered by standard video equipment. We propose a method of imaging a scene at high dynamic range, 14+ bits, to detect motion correlated with changes in the measured optical signal. Features within a scene, namely edges, can be tracked through a time sequence and produce a modulation in light levels associated with the edge moving across a region being sampled by the detector. The modulation in the signal is analyzed and a model is proposed that allows for an absolute measurement of the displacement of an edge. In addition, turbulence present in the received optical path produces a modulation in the received signal that can be directly related to the various turbulent eddy sizes. These features, present in the low frequency portion of the spectrum, are correlated to specific values for a relative measurement of the turbulence intensity. In some cases a single element sensor is used for a measurement at a single point. Video technology is also utilized to produce simultaneous measurements across the entire scene. Several applications are explored and

the results discussed. Key applications include: the use of this technique to analyze the motions of bridges for the assessment of structural health, non-contact methods of measuring the blood pulse waveform and respiration rate of an individual(s), and the imaging of turbulence, including clear air turbulence, for relative values of intensity. Resonant frequencies of bridges can be measured with this technique as well as eddies formed from turbulent flow.

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER I:

INTRODUCTION

Innovations in electronics have made it possible to make more precise measurements with sensor technology. In optical systems these measurements are at times only limited by the Poisson noise inherent in the photons being measured. Remote sensing is a field that has benefited immensely from these innovations leading to equipment with lower noise, higher sensitivity and more capability. These new technologies have advanced the field allowing new measurements to be made leading to more insight into the system they are probing.

The amount of information contained within the photons that travel from one object to another is vast. Take for example a photon emitted shortly after the beginning of our Universe or even billions of years ago. It has traveled those billions of years reaching us here on Earth and contained with it is information that reveals to us the nature of the Big Bang, the shape of the Universe, as well as it's age. [1-4] Although grand as it may sound, this happens every day to everyone; light travels from a object containing information about that object's surface color and material, reaches one's eye and that data is processed in such a way that makes that information meaningful. Take that a step further and add the temporal dimension to the process and a continual stream of photons are

arriving at the eye revealing properties such as velocity and acceleration of the source.

1.1 Goals and Current Work

The focus of this work is to develop a model that outlines the principles behind time varying intensities of light and its relationship to physical parameters of the source structure. The work involved details the development of a method to detect subtle motion of a remote object by correlating the emitted radiation of the object to its motion. Features within the structure are directly related to the amount of light that the object transmits, for example the reflection coefficient will determine the amount of light an object's surface will reflect, or the density of a gas will determine the amount of light that is transmitted through the gas. These features can be analyzed in the temporal domain revealing details as to the motion of the object.

This work involves the ability to make measurements that require a high dynamic range, that is, the signal being detected may be a very small component on a very large background. These signals are often as small as one part in a million.(2^{20} or 20 bits) This high dynamic range imaging allows for subtle motions to be detected that are correlated with small variations in the signal. Displacements smaller than the wavelength of light (<500 nm) have been detected in this way and will be described in this dissertation. These displacements are detected using modulation of the light from the surface features of the object. Other methodologies will be described that utilize similar techniques to see light modulation from the minor fluctuation in turbulent air.

Small variations in the density of the air result in fluctuations of light intensities. These variations are associated with the motion of the turbulence and allow the described methodology to be applied for the detection of even the small density fluctuations from clear air turbulence.

From the developed model several areas of application will be outlined. Data has been collected and processed on several systems from mechanical machine health to human biological functions where the methodology developed in the described work can be applied. Each application will be described in detail and the data presented.

1.2 Hypothesis and Goals

The following research utilizes high dynamic range imaging to detect very small variations in light to describe the motion of remote systems from where the light originates and is encoded. The information in the time varying light signal reveals information about the system allowing the exploration of the following hypotheses:

1) The time varying intensities within a pixel can be correlated to the time varying displacement of the image of edges within a single pixel/detector element.

2) High dynamic range imaging and data processing can offer the necessary system performance required to make measurements sensitive enough to remotely detect the motions of objects.

3

3) The information encoded in the time varying signal reveals information about the remote structure affording information that can characterize that system.

The primary goals of the research are to:

1) Correlate the time varying signal of a surface feature of a remote object to a quantitative description of the motion.

2) Develop methods of data acquisition and processing that are suitable for optical, non-contact measurement and analysis of remote systems.

3) Apply the technique to systems where a remote measure of motion and dynamics would be of use.

Throughout this dissertation, the term pixel will refer to a single sensor element of an imaging device. It can be singular in the case of a single element sensor or found in multiples as is the case with a video imaging device.

CHAPTER II:

BACKGROUND

In some sense remote sensing traces its origin back to thousands of years ago when humans looked to the heavens to record information about the night sky. Whether it be to characterize the objects they saw or record the motions of stars as a method of keeping time, humans have always been collecting information about remote objects. The eye has been a very powerful tool to perform such task, revealing information such as the wavelength of light reflected from an object, its intensity, and its place or origin.

Instrumentation has helped advance the field of remote sensing improving the information we are able to receive. For example, in the 1600's the telescope was invented, vastly improving the ability to see the details by eye of remote objects through magnification and increased light gathering. In the 1800's methods were developed allowing for the recording of focused light as images, later to become what we know as photography. That was followed by the electronic detection and digital recording in the 1900's, in which the charges produced by photons striking a surface are captured. measured, and recorded. Although these are a only a few examples in the progression of imaging technologies, numerous other achievements have made it possible for the remote sensing field to advance to where it is today.

2.1 Technological Landscape

The methodologies used in this research involve the detection and precision measurement of motion. For qualitative comparison, an outline of the current technological landscape of other motion sensing devices is described. Their advantages and disadvantages are reviewed to determine where the proposed technology fits within the existing field.

2.1.1 Accelerometers

Probably the most common sensor technology used for the detection of motion is the accelerometer. Although the accelerometer is not strictly a remote sensing device in the sense of non-contact detection, wireless technology allows the data to be streamed off the device to a remote location.

Figure 1. A three axis accelerometer records motion from all three spatial axes simultaneously. The above accelerometer is a 3-axis Crossbow CXL02TG3 with a sensitivity of 833 mV/g, typical of those used in our lab. The fourth connector shown is the temperature output from an integrated thermometer.

Accelerometers detect motion by measuring the acceleration of an object through various inertial sensors. They output a voltage that is usually proportionate to the acceleration. From acceleration, the temporal signal can be integrated once to determine the velocity and twice to determine position. In that fashion, an object's relative position can be tracked in space over time. Accelerometers often measure all three axes in space.

Accelerometers are highly sensitive devices that are today ubiquitous in motion sensing. Less sensitive microfabricated devices that cost only a few

dollars to produce are commonly found in household items such as game remotes and cell phones and are robust in terms of environmental conditions in which they operate reliably. Even the most sensitive devices used for inertial measurement and seismology cost only a few thousand dollars to manufacture.

The major drawback to accelerometers is the fact that they are required to be in contact with an object to measure its motion. The problem with this is twofold. For one, this requires physical access to the object to be measured, and in many cases it is not cost effective or safe to do so. For example, objects that are in motion may not be a suitable environment for an electronic device such as an accelerometer to be placed. In other cases, for example monitoring large structures such as bridges, access may be difficult to achieve. If it is possible to gain access for accelerometer placement, connections for power and telemetry are still required. When the power requirements are more severe data are transmitted by means of wireless communication, the power requirements are more severe. [5]

The second problem with accelerometer technology is that it requires physical contact with the target. In many cases measurements are necessary that a disturbance such as a contacting accelerometer would corrupt. The mass from an accelerometer would affect the dynamics of the system so unless $M_{acc} << M_{sys}$ the observer alters the measurement. In the case of the optical measurement M_{acc} would now be the effective photon mass, so that its recoil from photon momentum is all that limits the accuracy. Even the load from the

weight of the accelerometer may be enough to fundamentally change the quantity of what is being measured.

2.1.2 Doppler Laser Vibrometers

A commonly used optical technology to sense motion is the laser vibrometer. The vibrometer works by measuring the beat frequency of a modulated laser with the Doppler shifted backscattered light. [6,7] The beat frequency is proportional to $\frac{v}{c}$. Consequently, the laser vibrometer measures velocity then does a Fourier transform to get frequencies. The displacement is the integral of velocity and acceleration is its derivative and vibrometer gives precisely calibrated information on the dynamics of the target.
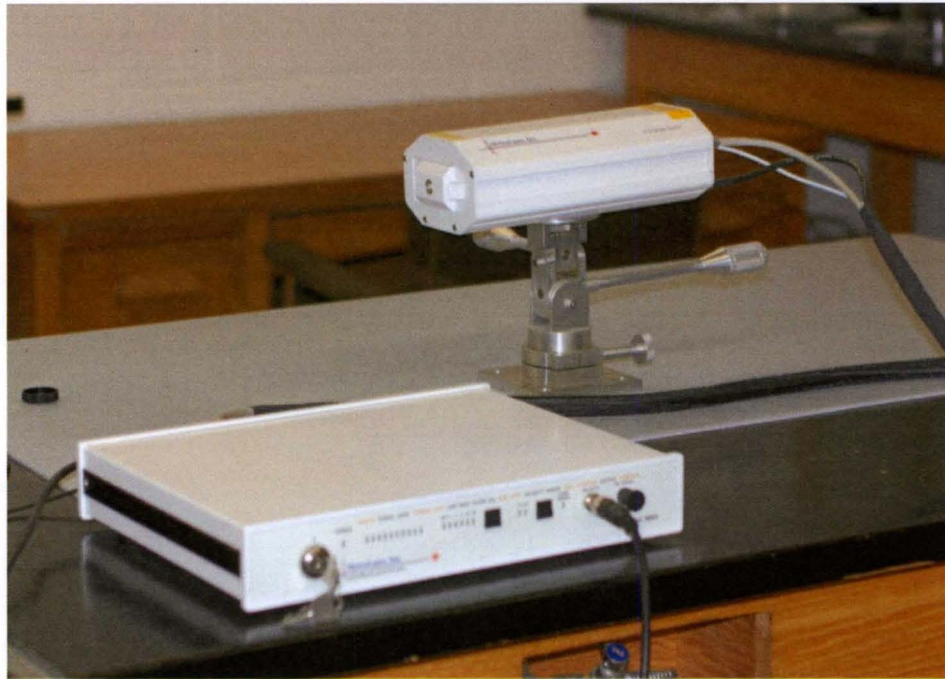
Figure 2. A laser vibrometer measures velocity by means of Doppler shift. The one shown is a Metro Laser Vibromet V500 which is used in our lab. It is capable of measuring displacements ranging from 0.1 nm to 10 mm and frequencies from DC to 20 kHz with 3% error.

Commercial laser vibrometers are highly sensitive and can measure motions as small as 10 nanometers. [8,9] They are non-contacting and are able to measure motions of very small targets in areas that may be inaccessible or unsafe to access. For example an object may be too hot to affix a sensor directly on the surface or may be in an area too hot to even access. Other objects may be too small to affix a sensor to, such as cells or small structures.

The laser vibrometer technology has practical limitations however. First, it is an active device that requires illumination of the target and deflection of weak

backscatter by a laser. There may be situation in which illuminating a surface with obvious laser light may be undesirable. Second, because of the detection requirement the laser vibrometer is a short range sensor, and objects more than a few hundred meters away cannot be measured with "eye safe" laser powers. [10] Third, since it is a pointed device, it can measure only a single point on a surface at a time and cannot be used to detect relative motions with a single object. Finally, laser vibrometers are costly when compared to accelerometers. The least expensive commercial units at this time are about $20,000.

The laser vibrometer is state of the art sensor for several applications for which a passive remote high dynamic range senor is suitable. One such use is the application of the vibrometer to the measurement of the carotid blood pulse. [11-14] The laser vibrometer can be used to integrate the velocity measurement to determine an absolute value of displacement for the blood pulse waveform, as will be shown with the carotid artery. This offers the advantage of the non-contact nature of the vibrometer and provides a highly accurate baseline measurement. However the disadvantage of cost and single point sensing are significant in clinical applications.

2.1.3 Strain Gauges

Another common technology for measuring relative motion are strain gauges. Strain gauges are small devices attached firmly to the surface of an object. They work by measuring the resistive changes associated with the tension and compression that results from small relative displacements in the point of contact of the strain gauges with the surface.
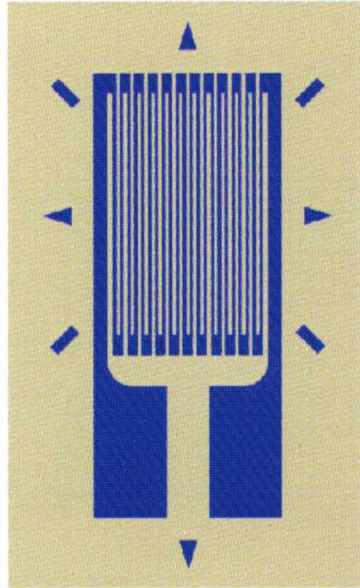
Figure 3. A strain gauge is a common device used to measure surface motion associated with tension and compression. When the material in a strain gauge is stretched or compressed the resistivity changes allowing the displacement to be measured. Source: Creative Commons.

An obvious disadvantage to the strain gauge is the requirement that it is attached to the surface. Often the strain gauges are permanently affixed to the surface with glue and require surface preparation. This could restrict the surfaces that are available for use with strain gauges due to this requirement. This also means that access to the surface would necessary ruling out inaccessible areas or surfaces that are in hazardous areas. Finally a voltage is required to be applied to the strain gauges requiring some sort of power supply and maintenance. The data would need to be transferred from the site as well, requiring either a wired solution or powered wireless device. In addition because

strain gauges measure very small changes in resistance, they are affected by temperature fluctuations.

2.1.4 Extensometers

Extensometers are a common device used to measure small displacements in some direct or mechanical way. They come in a variety of forms some of which are more directly related to the optical methods of this dissertation. Extensometers generally come in two classes, contacting and non-contacting.

Early contacting extensometers use a lever arm effect where two points are fixed to an object and as the two points separate one arm rotates. The measurement of displacement is made at the end of the arm where the distance moved is greater for the given angle. [15] More common was the dial extensometer which use a sensitive dial gauge as a means of measuring the displacement. [16]
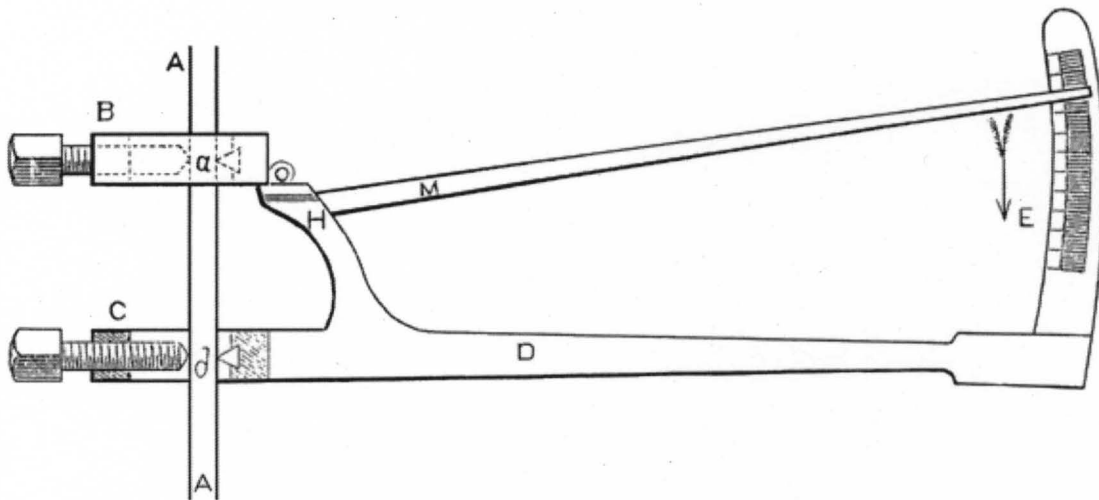
Figure 4. A contact extensometer is shown that measures the displacement by translating a small motion of the sample into a larger one through a lever arm. More sophisticated extensometers use laser speckling and video technology to make highly accurate quantitative measurements of displacement. Source: Creative Commons.

Non-contacting extensometers are more aligned with the technology presented in this dissertation. They come in two primary forms, laser and video. Laser extensometers work by illuminating the surface with a laser and recording the illumination with a CCD camera. As the surface is displaced the CCD captures changes in that illumination and through signal processing the displacement is determined. No contact is necessary and surface preparation is not required. However this techniques requires an active laser illumination and therefore will be limited in range. It will also have the limitation of only being useful in situations where active illumination is acceptable.

The second form of extensometer is the video extensometer. They work by preparing the surface with markings that are recorded with a video camera. When a displacement occurs, these markings are tracked with the video system and the displacement is determined through the geometry of the optical system and correlation functions. [17] This system has the advantage of being non-contacting. However, surface preparation requires access to the surface and is prone to error. [18] Furthermore, these systems are highly dependent of the illumination of the scene with a steady light source and shading from outside illumination. One form of video extensometry, laser speckle extensometers, works by illuminating the scene with a laser to produce a speckle interference from the surface. These patterns are recorded by one or more cameras to track the patterns produced by the laser through data processing. This technique requires no surface preparation but still suffers from the drawbacks associated with laser illumination and the difficulties of analysis of the speckle.

2.1.5 Tonometry

Tonometry is another technique that, while measuring pressure and not motion, is applicable to a biomedical application discussed in this dissertation. It works by applying a force to a surface and using the deformation of that surface as a measure of the pressure inside. [19] Applanation tonometry is the common forms of this technique. The drawback to this process is the requirement of making contact to the surface and the necessity of applying pressure to the measurement site. The application of force to the measurement site can disrupt the actual measurement being made altering the surface and potentially

15

changing subtle variations in the measurement. A non-contact form of this measurement is done by applying puffs of air to the surface and measuring the deformation but the same disadvantages apply since a force is still being applied to the measurement site. [20,21]

2.1.6 Manometers

The sphygmomanometer, like the tonometer measures a pressure at which specific response occurs. This device is also known as the common blood pressure cuff and has widespread reliable use. [22] However it requires contact and suffers from the drawback that it alters the measurement site and potentially affects subtle changes of pressure.
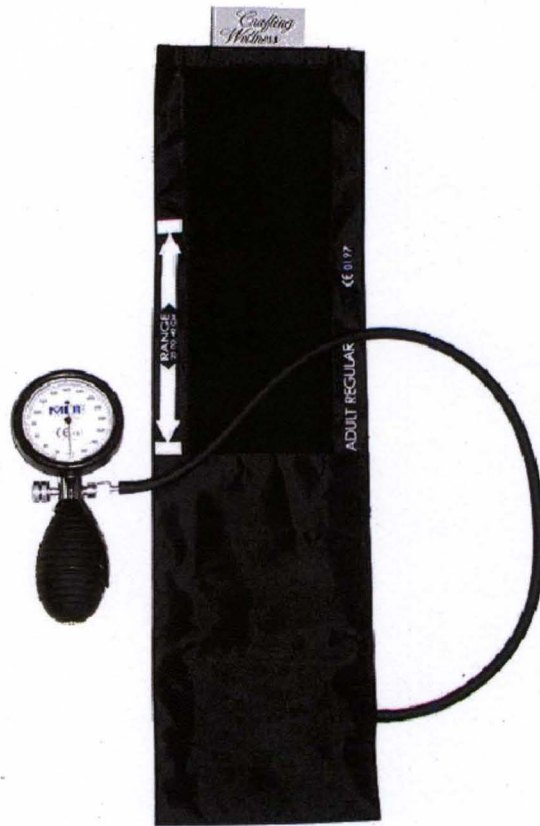
Figure 5. A sphygmomanometer measures blood pressure by listening, in conjunction with a stethoscope, for the return of blood flow after the artery has been occluded at a measured cuff pressure. More advanced sphygmomanometers derive the blood pressure from readings of the surface pressure of the artery. Source: Creative Commons.

CHAPTER III:

METHOD OF OPERATION

3.1 Physical Principles

The principle behind high dynamic range imaging for the detection of motion is to use the dynamic range of an optical sensor to determine a physical displacement of an object in the scene. For example, a single element photodiode sensor can have a saturation level of 6 V, with a typical noise floor of 1 μV rms, a range of 6×10^6. That gives a dynamic range between 22 and 23 bits since 2^{23} is about 10^7. The goal would be to transform that exceptional range of signal into a comparable measurement of displacement by mapping a feature spatially across a single detector element.

2^{16}

Typical Bit Depth
Intensity
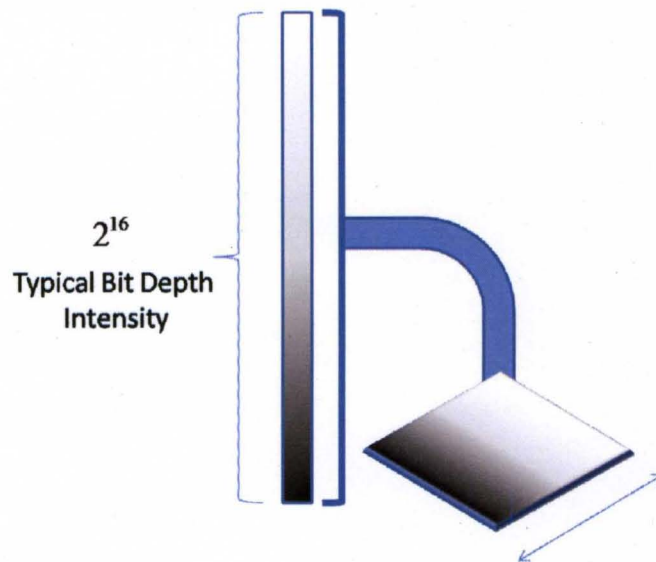
Figure 6. Mapping the signal to a pixel. Here we see how the bit depth of the camera can be mapped to the spatial dimension of the pixel. A typical video pixel may have a bit depth of 2^{16} differing levels of signal. If each possible level of intensity was correlated to a spatial location of a feature within the pixel the ability to spatial resolve that feature can be greatly improve. Normally a feature can simply be determined to be in the pixel or not.

In the case of a dynamic range of 22 bits, that represents approximately 4 million discrete values the sensor can achieve in signal. If those values can represent a physical displacement in the scene, we now have approximately 4 million discrete units of displacement mapped across the detector as opposed to ~65,000 units, as is the case of a 14 bit camera, or even 1 unit when the actual pixel is the spatial measurement!

19

The system works by measuring the modulation of light level imaged on the detector. This can be done with a single element sensor, as is the case with a photodiode, or an array of sensors like a high speed camera. In the case of a camera, each pixel acts as an independent sensor. The modulation in light level is attributed to the fact that, in some way, the amount of light reaching the sensor element is changing.

The simplest and cleanest example would be a black and white edge. Imagine the target of interest was something similar to what is shown in Figure 7.

Figure 7. An imaged black and white edge is shown, one that is used in the lab to provide a clean and simple method of testing our sensor. Here we see a high contrast black and white edge affixed to a speaker cone.

Figure 7 shows a simple black and white edge drawn on a piece of paper attached to the cone of a speaker. For simplicity, let us assume that the black portion of the target absorbs 100% of the light incident upon it and the white

portion of the target reflects 100% of the light. If the sensor were to image the scene with the black and white edge in the center of the sensor detection element, the signal of the scene would be 50% of the value if the entire scene were white. Let us assume that the detector operates just below saturation levels, that is a fully white scene, where the entire target is completely white, gives a level of nearly 6 V on the detector. Now if the black and white edge is imaged on the detector, again with the edge falling in the center of the scene, the sensor would give a value of 3 V, assuming linearity in the sensor. This is the DC value of the signal which represents the ambient illumination of the static target. In this scenario we are also assuming there are no fluctuations in the light level, that is, that the system is constant both spatially and temporally.

Let us now assume that the speaker vibrates at a known low frequency, say 10 Hz. The characteristics of the signal modulation can be determined by performing a Fast Fourier Transform of the temporal signal, as sampled by the detector through a series of discrete measurements, equally spaced in time, of the image of the black and white boundary. Looking at a single element in the scene, we would see the signal of that point vary as the edge moves in and out of the field of view or across the area image on the single pixel. This would occur due to the fact that the ratio of the black to white area seen by the detector would change. The signal at a single element at a given time is ultimately determined by the amount of light collected from the surface. As previously mentioned, a completely black surface would ideally result in an signal of 0 while a completely white surface would result in a nearly saturated pixel. Let's assume we are

working with a 16 bit sensor so that the fully saturated element is 65,536. If the speaker were vibrating such that the edge vibrated completely across the pixel's field of view (moved from one side of the imaged area to the other), the amplitude of vibration would be equal to the resolution of a single pixel. So, if that pixel were aligned so that the turning points of the vibration were on the very edge of the pixel's field of view, we would see the pixel's values change from 0 to 65536 at a frequency of 10 Hz. The corresponding modulation associated with this movement would be 65536 for 100% modulation. If, instead, the speaker vibrated with a smaller amplitude such that the turning points of vibrations fell within the pixel's field of view, we would not expect the signal to be 0 or 65536 as there would always be some part of the black and some part of the white surface in the field of view. Let us assume that the edge instead moves across 80% of the field of view, that is the amplitude of the vibration is such that the turning points are 10% of the pixels field of view from either edge. At the maximum and minimum values of the vibration, 10% of the white edge and 10% of the black edge would remain in the field of view. This would result in values of approximately 6554 and 58982, respectively, for a modulation of 52428 or 80% of 65536.
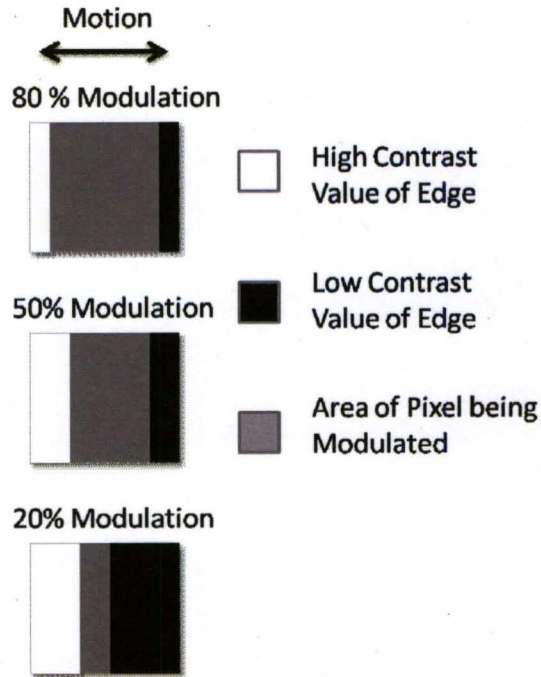
Figure 8. Various levels of modulation are shown that could illuminate a single pixel. The white portion on the left stays illuminated by the brighter side of the edge the entire time. The gray area in the middle represents the area of the pixel where the edge is seen to move back and forth. It is periodically covered by the edge and undergoes modulation. The black portion on the right stays illuminated by the darker side of the edge the entire time. Three examples of varying levels of modulation a pixel may undergo are depicted.

Figure 8 shows various level of modulation as the image of the edge moves across a sensor. Various levels of modulation are proportional to the percentage of the pixel that the edge moves across when imaged on the detector.

The percentage that the edge moves across the field of view can then be correlated to an actual physical displacement. Figure 9 shows the geometry associated with the distance subtended as measured on the detector's surface. The actual displacement of the edge, which is the distance from the two turning points of the vibration, is related to the distance across the sensor in which that edge move by: 1) the distance to the target, 2) the focal length of the lens and 3) the angle subtended by that distance.
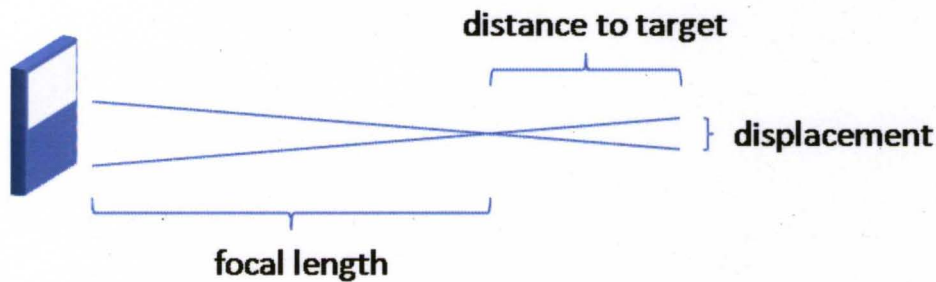
Figure 9. Geometry of the displacement to a pixel shows how the physical parameters relate to the modulation on the sensor. A pixel is shown being modulated due to a displacement of the target.

For a very small displacement, the angle is expected to be very small as well. In that case, we expect the small angle approximation would be appropriate leaving only the focal length of the lens and distance to the target to determine the displacement.

So far we have shown that the percentage of modulation is related to the edge's movement across the surface. If we assume that the edge stretches across the entire pixel, as shown in Figure 9 then we can assume that the modulation is directly proportionate to the sensor width. Again we are assuming

that the edge is black and white and goes from full saturation on the white edge to 0 on the black edge. If the sensor is experiencing 50% modulation, the image of the edge must be moving across 50% of the pixel's width. In other words the edge must be moving across half the pixel during this modulation. If we assume a square pixel, that would mean that the modulation tells us how much of the sensor is being covered by the image of the edge.

Now we can derive an actual physical displacement of the vibration from the modulation levels of the sensor. According to Figure 9, the fractional modulation of the sensor is then determined by the following relationship

$$m = \frac{f}{r} \frac{dx}{p}$$

where the pixel is mapped back on the target width $w = \frac{pr}{f}$ and:

f is the focal length of the lens

r is the distance to the target

p is the pixel width

dx is the displacement

This tells us that the fractional modulation of a pixel detecting a black and white edge in a scene can give us the physical displacement of the edge if we know

25

the focal length of the lens, distance to the target, and the width of the pixel. Of course, if the edge motion has principal components at distinct frequencies, then a Fourier analysis of the signal will reveal those components.

There are two significant problems with this simple example. One is the fact that a signal level of a pixel can never achieve a value of 0; noise due to Poisson statistics in the photon flux will always be present at the very least. Second is the fact that the signal values across the edge will rarely if ever be the full dynamic range of the sensor. Let us deal with the problem of noise first.

Photon detection inherently has a noise associated with it. Poisson statistics apply to the measurement of the number of in a defined time interval, and the uncertainty limits the accuracy of a photonic signal. Assuming that all others sources of system noise can be eliminated, Poisson noise remains and it is unavoidable. The uncertainty in a count of N photons is $\sqrt{N}$. Since the number of photons is proportional to the signal level for the pixel, the statistics tell us that ultimately them most favorably signal-to-noise ratio is $\frac{N}{\sqrt{N}}$ or $\sqrt{N}$. Therefore the smallest modulation we can detect will be when that modulation is of the same fractional order of magnitude as the noise to signal, or

$$\frac{\sqrt{N}}{N} \text{ or } N^{-\frac{1}{2}}$$

Similarly, this establishes a fundamental limit on the smallest detectable displacement of the system, since

$$m = \frac{f}{r}\frac{\delta x}{p}$$

where δx is the displacement for a modulation m. We solve for the displacement corresponding to modulation m

$$\delta x = \frac{rp}{f}m$$

Since we are limited by Poisson Noise at the smallest possible fractional modulation,

$$\delta x = \frac{rp}{f}N^{-\frac{1}{2}}$$

This tells us that the smallest detectable signal is determined by the distance to the target, the focal length of the lens, the width of the pixel, and now the number of photons detected.

Again we are still dealing with a case where the modulation level is based on a situation with an ideal black and white edge. Obviously this will not always be the case so let us look at the second problem described above, a situation where this is no longer the case.
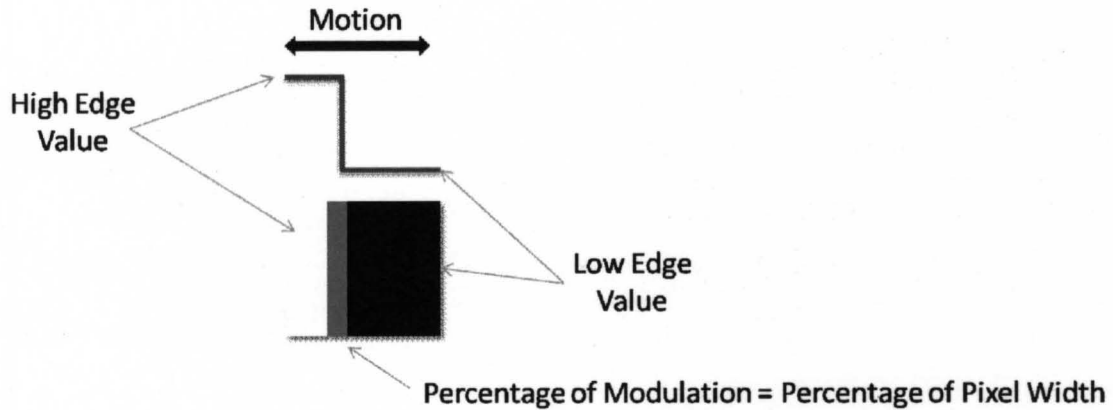
Figure 10. Modulation across an edge is represented. Here we see how a fraction of total modulation is related to a fraction of the size of a pixel. As a black and white edge is imaged on a pixel, the physical area it covers moves across the sensor element and modulates the signal. If a black and white edge periodically transverses a certain percentage of the area of the pixel when it is imaged on the detector, the signal shows that same percentage modulation.

In the event that the edge is no longer purely black and white, that is 100% absorbing and reflective, respectively, there will be some albedo A_1 defining the brighter side, and some albedo A_0 defining the dimmer side. For example, a region with an albedo A_1 would have an albedo 1, scattering incident light with 100% efficiency, and a region with an albedo A_0 would have an albedo 0 absorbing 100% of the incident light.

A position of the edge can be defined by x such that at x=0 the dark region with albedo A_0 fills the pixel, and with x=w the pixel is filled with the region of

albedo A_1. We allow for the possibility of a background, B, contributing to the signal regardless of where the position of the edge, x, is located. A maximum signal is defined as S_{max} that includes the background and is measured for a perfect surface with albedo 1. The signal can then be expressed in terms of the position of the edge and is given by,

$$S = \left(\left(\frac{x}{w} \right) A_1 + \left(\frac{(w-x)}{w} \right) A_0 \right) (S_{max} - B) + B$$

Now consider a small motion in the target that displaces the edge. The resulting change in the signal will be given by the derivative of this expression which is,

$$dS = \left(\frac{dx}{w} \right) (A_1 - A_0)(S_{max} - B)$$

We see the signal change is still proportional to the displacement of the edge, but now also the difference in the albedo of the two sides and the maximum signal above background

We measure the signal S_1 when only the brighter side, A_1, fills the pixel by setting x=w and signal S_0 when only the darker side, A_0, fills the pixel by setting x=0 to obtain two reference signals

$$S_1 = A_1(S_{max} - B) + B$$

$$S_0 = A_0(S_{max} - B) + B$$

The difference in these two values is

$$S_1 - S_0 = (A_1 - A_0)(S_{max} - B)$$

The signal change can be expressed in terms of this difference by

$$\frac{dS}{(S_1 - S_0)} = \frac{dx}{w}$$

$$\frac{dS}{(S_1 - S_0)} = \frac{f}{r}\frac{dx}{p}$$

again where $w = \frac{pr}{f}$.

We can identify this relative change in the signal as the fractional modulation m_f of the signal as it represents the signal change with respect to the difference in the two signal on either side of the edge so that

$$\frac{dS}{(S_1 - S_0)} = m_f$$

This gives is an equation of the form

$$m_f = \frac{f}{rp}dx$$

that relates the modulation to the measurable parameter of the lens, the detector, and the distance to the target. With this definition, a measurement of m_f can be connected to a physical displacement dx if the signals for the two sides of the edge can be measured individually.

Comparing this to the originally derived equation

$$m = \frac{f}{r}\frac{dx}{p}$$

30

we see that m_f has replaced m. Again recall that

$$m_f = \frac{dS}{S_1 - S_0}$$

This tells us that the measured change in the signal dS is scaled by the difference in the signals on either side of the edge. Physically this makes sense as the detector is limited in measuring a modulation that must fall within the range of the those signals.

Recalling

$$\frac{dS}{(S_1 - S_0)} = \frac{f}{r}\frac{dx}{p}$$

we can express this as

$$dS = \frac{f}{r}\frac{dx}{p}(S_1 - S_0)$$

so that the displacement becomes

$$dx = \frac{rp}{f}\frac{dS}{(S_1 - S_0)}$$

Now we see that the only additional quantities necessary to make an absolute measurement, regardless of the surface albedos and background are the signals on either side of the edge, S_1 and S_0.

Again doing the same substitution for the smallest displacement we see

$$\delta x = \frac{rpN^{-\frac{1}{2}}}{f(S_1 - S_0)}$$

where N is an average signal measured in photon count. So now we see that the smallest displacement is determined by the following:

1. The distance to the target   r

2. The width of the pixel.   p

3. The average number of photons.   N
   per pixel per time element.

4. The focal length of the lens.   f

5. The signal difference across the edge. (S_1-S_0)

To optimize the system to detect the smallest displacement we would need to do the following:

Maximize	**Minimize**
Focal Length	Distance to the target
Signal Difference across the edge	Width of the Pixel
Photon Count	

Intuitively this makes sense as getting closer to the target, increasing the focal length, and decreasing the width of the pixel makes the spatial resolution more favorable on the sensor. Increasing the photon count is favorable as it increases the signal-to-noise ratio. However it is the signal difference across the edge that is unique to this system. Quite simply, the higher the difference across

32

the edge the larger the modulated signal can be, allowing for a higher signal-to-noise. It is a linear effect. The largest factors are distance and total photon count.

More importantly, though, the signal difference across the edge gives us terms of the absolute physical measurement. As it turns out, the system becomes completely self calibrating. The sensor is capable of measuring an absolute displacement as long as the signals across the edge is known. It does not matter if the target or illumination changes (assuming it does not change during the integration time), as long as the signals across the edge is known. In the situation of an image array or camera, that value can be achieved simply by measuring the neighboring pixels. A single element sensor would need to sample areas adjacent to the edge to ensure that the sensor measures the values neighboring either side.

CHAPTER IV:

DATA COLLECTION

4.1 Instrumentation

The premise behind the high dynamic range sensing and its ability to see small motion lies in the ability to see minute fluctuations in light levels and correlate those measurements to physical motions. Two primary types of instrumentation are used to do this; a) a highly optimized single element sensor that has a dynamic range on the order of 20 bits resolving light level changes as small as 1 part in a million, and b) video camera technology that has been optimized for high dynamic range imaging. This optimization is done in either data post-processing, and or through the selection of a camera with a suitable bit depth.

4.1.1 Single Element Sensor

The single element sensor in our experimental work consists of an Tamron 23FM16SP, 16 mm focal length f/1.4 C-mount lens, coupled to an un-cooled Thorlabs FGA04 InGaAs photodiode by a 10 meter Thorlabs AFS105/125Y fiber optic cable. The signal is amplified by a Analog Devices AD549LHZ operational amplifier with a dark current noise of $0.1 \mu V$. The circuitry is enclosed in an RF shielded enclosure which also contains the powers supply

34

and Tadian TL5903 3.6 V lithium ion batteries.   The photodiode has a 100 μm diameter cathode, the same size as the fiber optic core, and the sensitive area is aligned to the fiber by an FCA fiber optic connector. A schematic of the system is shown in Figure 11.
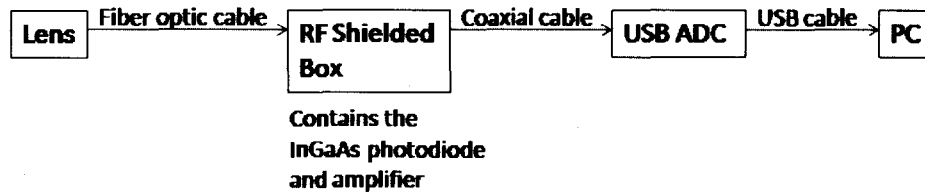
Figure 11. A schematic showing the single element sensor.

The single element sensor has a load resistor that is 100 x 10^6 Ω. In terms of the noise level,

$$V = 1\mu V = 10^{-6}V$$

$$I = \frac{V}{R} = \frac{10^{-6}V}{10^8\Omega} = 10^{-14}A$$

e = 1.6 x 10^{-19} C so, I = 6 x 10^4 e at 1μV and 36 x 10^{10} e at 6V. Now,

$$\frac{N}{\sqrt{N}} = 6x10^5$$

So at the maximum signal, the noise, of the order of 1ppm, is shot noise due to the number of electrons detected.
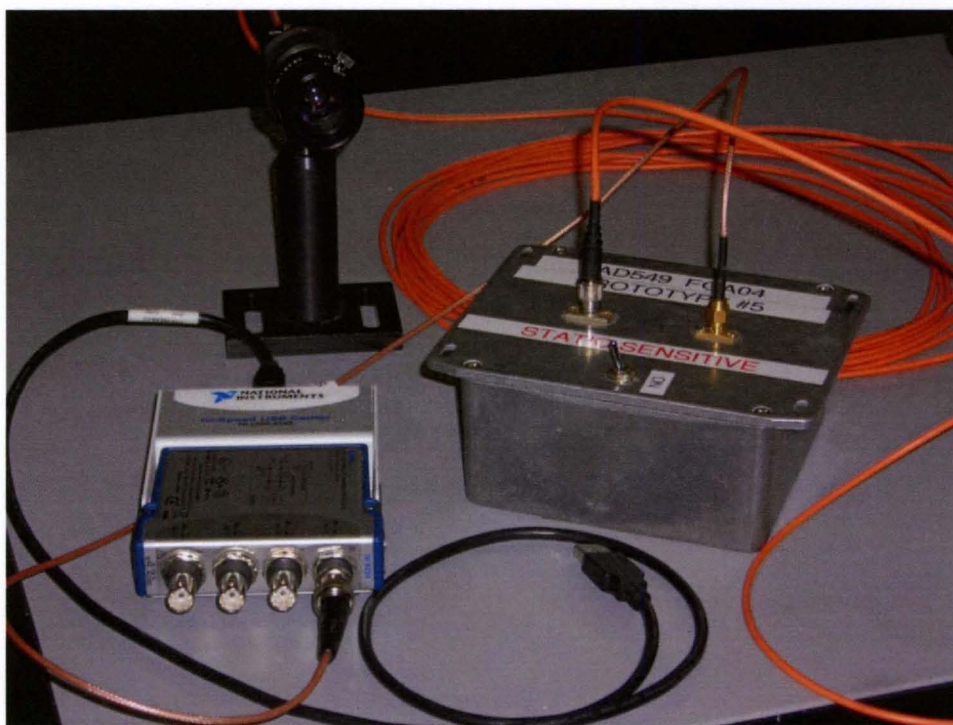
Figure 12. The single element photodiode system is shown. The basic setup for the single element sensor comprises a lens coupled to a fiber optic cable. This transmits the signal to an RF shielded box containing a photodiode and amplifier. Once amplified the signal is then transmitted to an analog to digital converter.

In order to receive as much as light as possible the largest lens aperture is generally chosen. In addition to the Tamron 16mm lens, a Celestron ONYX 80ED 500 mm focal length telescope is used in situations where higher magnification is more suitable. The telescope has an f ratio of f/6.25. The telescope is a much slower optical system then the 16mm lens. However, the focal length of the telescope is much greater and for distant targets it affords proportionately greater spatial accuracy. Recalling the equation on page 32

outlining the smallest displacement, the telescope's longer focal length allows for a much smaller detectable displacement because the image is spread over a larger area in the focal plane. This is especially advantageous when we have high light levels as happens in outdoor sunlit scenes.

The signal is digitized in a National Instruments NI-9234 24 bit analog to digital converter that is AC or DC coupled and measures a ±5 V range. For in-lab experiments, the signal is digitized with an equivalent 24-bit NI PCI-4472 card on a PCI bus in a desktop computer that is AC or DC coupled and measures a ±10 V range. The signal is then read by LabVIEW data acquisition software. Typical cards used are capable of acquiring at a rate of up to 124,000 samples per second but data is usually acquired at rates near 2048 samples per second, more appropriate for low frequency mechanical oscillators. This allows for a maximum frequency measurement of 1024 Hz. Typical frequencies of interest in applications covered in this dissertation are within this range. Software written in LabVIEW is used to process the data as well.

4.1.2 Video Camera Acquisition

The single element sensor has the advantage of extremely high dynamic range and avoids saturation in even the brightest sunlit environments, but there are many advantages to using video camera technology in replace of the single element sensor. When used in a similar fashion to the single element sensor, that is for collecting light from a particular location and looking at the temporal features, the camera acts as an array of single element sensors, albeit with diminished dynamic range.

The best bit depth afforded by a commercially available camera that we use is the Allied Vision Pike FO32B offering a 14 bit A/D converter. The Pike F032B interfaces with the computer through a IEEE 1394b Firewire connection capable of transmitting data at rates up to 800 Mb/s. These connections can be daisy chained to allow more than one connection at a time. The Pike contains a Kodak KAI-0340 CCD Type 1/3 inch sensor with square 7.4 μm pixel size in a 640x480 array. The Firewire card we use is a SIIG Firewire 800 2-Port ExpressCard which supplies 12 V power to the camera. The software used to capture video in Windows 7 is StreamPix 5 written by a third party company NorPix. The software is optimized for high speed cameras and offers the ability to allocate RAM prior to acquisitions for streaming directly to RAM.

Figure 13. The Pike F-032 has a native dynamic range of 14 bits and capable of acquiring a 640x480 image at 107 frames per second. It's small size and Firewire interface make it a good choice for data acquisition in the field.

Although cameras offer the advantage of a larger array of sampled points on a target compared to a single sensor, their lower bit depth is problematic.

However, we have developed ways to overcome this limitation. Oversampling or co-adding frames allows the effective bit depth of the camera to be expanded as needed. To do this even for low frequencies, a high speed camera is necessary in order to oversample in time and retain the spatial resolution. Here, high speed will refer to a camera that is faster than a standard 60 frames per second video rate. Essentially the idea is to acquire more frames each second than necessary to resolve the frequencies or motion of interest. Frames are co-added in post processing to increase the highest possible photon count. This technique allows the camera technology to near the 18-19 bit range. In fact, the Pike FO32B has onboard processing that allows the addition of frames before they are exported allowing for higher dynamic range directly from the camera. Even lower bit depth cameras are useful if they have very high speed. For example, a second camera used in this work is the Photron APX-RS which has a native 10-bit depth. With oversampling and sub-imaging the 17-18 bit range can be achieved. The post processing of the images will be discussed more in the next section.

The Photron APX-RS has a CMOS type camera sensor that is much faster than the CCD used in the Pike camera. The sensor has a 17 μm square pixel with a full resolution of 1024 x 1024. The camera interfaces through gigabit ethernet and uses Photron's FASTCAM Viewer software. Through a gigabit TCP/IP connection the camera can communicate directly to the host computer or be placed on a local network. It operates at maximum frame rate of 3,000 fps at full 1024x1024 resolution. The fastest achievable speed with sub-imaging, in which only part of a frame is read, is 250,000 fps, much faster than the Pike

FO32B camera offers. Again, temporal oversampling can be used to overcome the 10-bit digitization. For example, windowing at a resolution of 128x128 gives a frame rate of 60,000 fps. If we are only interested in a 100 Hz signal, we can sum 300 frames for an effective frame rate of 200 fps. These additional 300 frames give a little over 8 more bits of dynamic range or an 18-19 bit image, coming close to the dynamic range of the single element sensor. Another factor limiting the Photron Fastcam is its use of onboard storage. During acquisition at high speeds images are streamed to the 8 GB of onboard memory. This limits the duration of acquisition, as opposed to the Pike camera, which exports data in real time via Firewire. An advantage to storage however, is the stability in the timing of the individual frames acquired by the Photon APX-RS, which is highly accurate. With the Pike, the system performance of the PC can unpredictably slow the acquisition leading to an inconsistent interval between frames. An optimized system with tuned software capable of handling the data in real-time is required.

4.2 Calibration

A robust method of operation has been described with the above instrumentation. A test of the system in terms of its operation with a well defined target is necessary. The simplest test would be to see the response of the system to the motion of a black and white edge across a pixel. To do this in the laboratory we use the single element sensor with a Tamron 16mm lens coupled to an optical fiber that feeds into the low noise amplifier. The signal is digitized at 24 bits with a NI 4472 PCI analog to digital converter card and read into the

LabVIEW data acquisition software. From there the signal is processed through a Fast Fourier Transform and the frequency spectrum is determined. The modulation of the signal is determined by integrating under the frequency peak of interest. To calibrate the signal a MetroLaser Inc. VibroMet 500V laser vibrometer is used to determine the target motion. The laser vibrometer measures the velocity of the target by Doppler shift of the modulated laser and is considered a reliable standard. The velocity is then integrated in the temporal domain to determine displacement for comparison to our image data.

The setup consist of a Pioneer TS-G1642R speaker with a nominal power of 30 watts and a 4Ω impedance. The speaker is driven by an Applied Research and Technology SLA-1 linear power amplifier with an input signal from a Model 3003 BK Precision 10 MHz sine and square wave handheld signal generator. A low mass plastic tube is affixed to the cone to allow a black and white edge target to be coupled to the speaker vibration while also exposing a surface perpendicular to the direction of vibration. The optical sensor and laser vibrometer measure the target at right angles to one another as seen in Figure 14. Both points of measurement are coupled to the same rigid body.
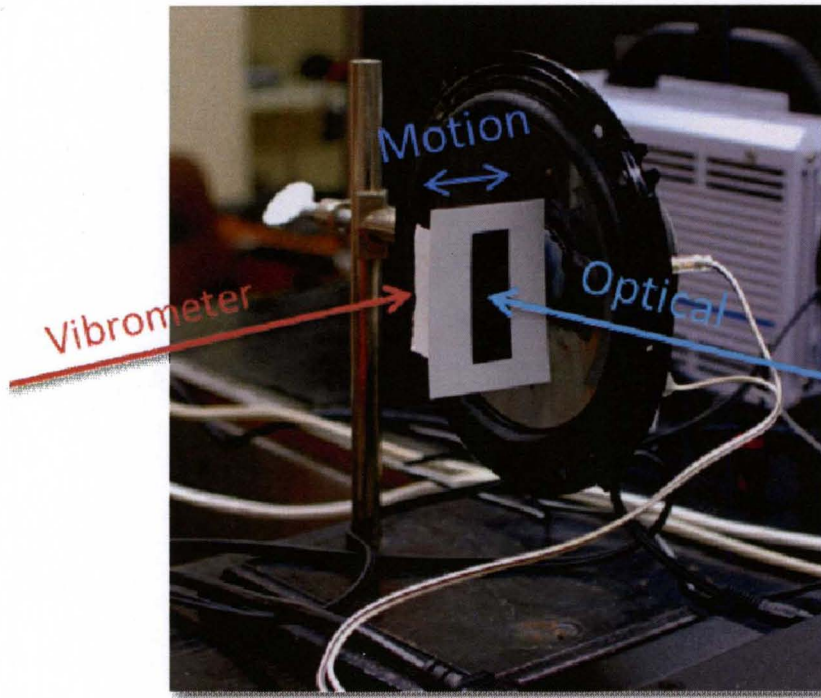
Figure 14. The optical calibration setup consists of a speaker oscillating while both the optical and vibrometer measurements are made simultaneously. A black and white edge is affixed to the speaker by means of a plastic cylinder glued directly to the speaker cone. This allows for a flat surface on which the laser vibrometer can measure as well as a surface to attach the black and white edge for the optical sensor. The surface that the laser vibrometer measures is perpendicular to the direction of motion while the surface of the black and white edge is parallel to the direction of motion.

The system was characterized to evaluate the motion of the speaker. The speaker was driven at a single frequency of 625 Hz throughout the test to ensure

42

the speaker's response was consistent between measurements. This frequency was chosen because the speaker showed well behaved motion in this region of the audio spectrum. The vibrometer was set to its most sensitive setting and the built-in low pass filter, set at 1 KHz, was applied. The vibrometer measured the velocity of the speaker during vibration from the face of the cylinder attached to the speaker as seen in Figure 14. This temporal signal was then integrated in LabVIEW to determine displacement. The optical measurement was made on the black and white edge attached to the side of the cylinder from a direction perpendicular to the motion of the speaker. To set the target field precisely, the fiber optic cable was illuminated on the photodiode end of the system and the light sent back though the fiber was projected on the paper. The lens was then aligned so that the projected spot was placed on the paper with edge in the middle. Data were collected with the optical system and vibrometer simultaneously.

Speaker Vibration

Optical signal relative to derived displacement
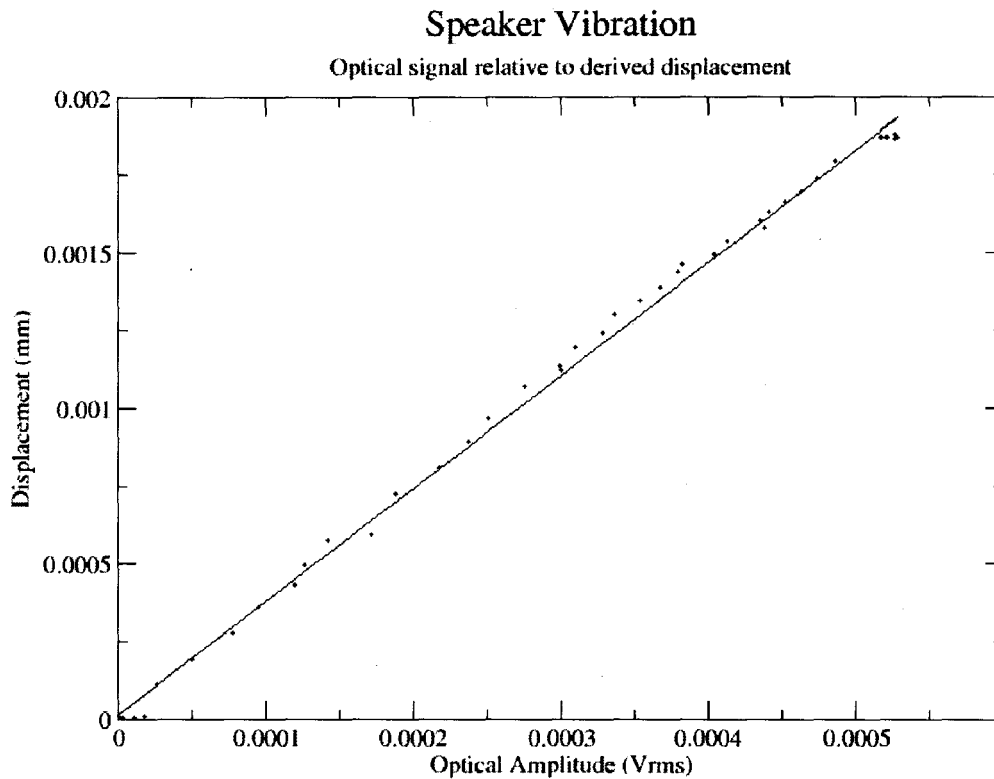
Figure 15. Calibration of the optical system to absolute displacement shows a strong linear relationship. The optical system measured the modulation of light associated with the periodic motion of a black and white edge while the vibrometer measures the displacement of a surface being displaced equally with the edge. The result shows that the optical system's modulation signal is linear with the measurement of displacement determined by integrating the laser vibrometer's measured value of velocity.

The displacement as measured by the vibrometer was then correlated with the optical measurement as shown in Figure 15. The optical system shows a strong linear relationship that supports our proposed model of the system. The fit is characterized as:

$$y = 1.4206 \times 10^{-05} \times (3.629 \pm 0.032)x$$

with a correlation coefficient of 0.9984. Note the value of 1.42×10^{-05} represents a value of approximately 14 nm, a small variation from where the graph should be intersecting the origin. It is interesting to note the two regions on either end of the plot shown in Figure 15. On both the upper and lower measurements of displacement, the optical system continued to track the measurement accurately while the vibrometer failed to continue to accurately measure the displacement as seen by the data flattening out on the y-dimension. This indicates that the optical system has a higher dynamic range in a position measurement than the vibrometer under these conditions and settings.

4.3 Confirming an Absolute Measurement

Section 3.1 outlines a method for determining the absolute value of displacement based on the measurement and system used. The equation on page 31 determines the absolute displacement that is measured by the system which is,

45

$$dx = \frac{rp}{f} \frac{dS}{(S_1 - S_0)}$$

To test this we measure the displacement of a black and white edge that is affixed to a vibrating speaker vibrating at 20 Hz, the plane of the paper being parallel to the motion of the speaker, the same setup used for calibration in Section 4.2 . The single element sensor is chosen for purposes of simplicity and positioned 980 mm. from the black and white edge. A 16mm lens is coupled to the fiber which transmits the light to the photodiode and its low noise amplifier. The signal is digitized to 24 bits with the NI-9234 USB ADC and is then analyzed with programs written in the LabVIEW environment. A white piece of paper is placed on a perpendicular surface to the black and white edge to provide a surface for the vibrometer's laser to illuminate.
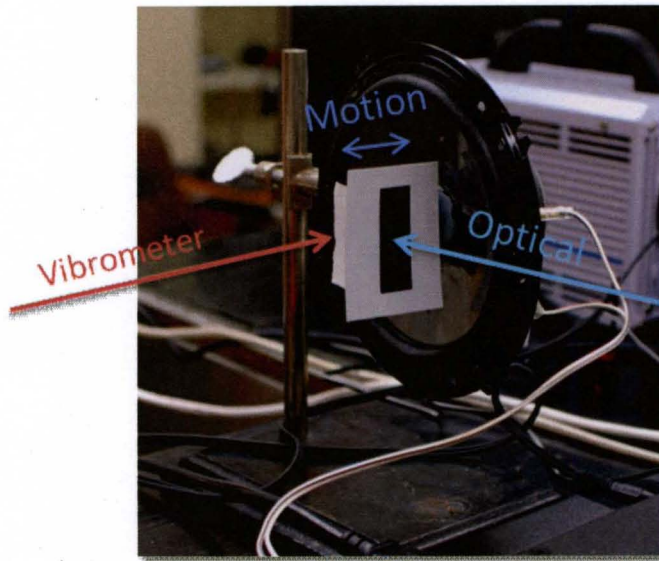
Figure 16. Confirming the absolute measurement uses the same setup as the he optical calibration, consisting of a speaker oscillating while both the optical and vibrometer measurements can be made from the same displacement of the speaker. A black and white edge is affixed to the speaker by means of a plastic cylinder glued directly to the speaker cone. This allows for a flat surface on which the laser vibrometer can measure as well as a surface to attach the black and white edge. The surface the laser vibrometer measures is perpendicular to the direction of motion while the surface of the black and white edge is parallel to the direction of motion.

The vibrometer is set to the most sensitive Lo setting with a 1Khz low pass filter applied. Its velocity signal is also digitized to 24-bits with the NI-9234 and read by the LabVIEW software. The vibrometer signal is integrated from the native

velocity measurement to displacement for a confirming measure to compare to the optical signal.

Looking at the equation

$$dx = \frac{rp}{f} \frac{dS}{(S_1 - S_0)}$$

we see that the necessary information need to determine the displacement form the optical signal is the following:

1. distance from the lens to the edge

2. width of the detector element

3. focal length of the lens

4. amount of detected modulation

5. signal difference between the high and low side of edge

The signal difference is found by focusing the optical detector entirely on a black portion of the end then a white portion of the edge and measuring the DC value. The DC values are measured at 90.36mV and 527.79 mV respectively for a signal difference of 437.43 mV. The distance to the lens is 980 mm, the width of the fiber element is 0.1 mm, the focal length of the lens is 16 mm and the measured modulated signal, dS, from the motion of the edge is 15.3 mV. The modulation for both the optical and vibrometer signal is determined by integrating over 0.2 Hz bins centered around the 20 Hz frequency of interest. The calculated modulation from the optical signal can now be determined,

$$dx = \frac{rp}{f} \frac{dS}{(S_1 - S_0)}$$

$$dx = \left(\frac{980 \; mm \times 0.1 \; mm}{16 \; mm}\right) \frac{15.3 \; mV}{437.432 \; mV} = 0.2142 \; mm = 214.2 \; \mu m$$

This compares to a reference value of 221.6 μm as measured by the integrated vibrometer signal. The laser vibrometer has a specified error of ±3% which result in a measured value of 221.6 μm ± 6.65 μm. The percentage error from the reference measurement is determined to be

$$\frac{221.6 \; \mu m - 214.2 \; \mu m}{221.6 \; \mu m} = 3.34 \; \%$$

The optical detector is able to measure the displacement of the edge to within 7.4 μm which is within 3.34% of the reference vibrometer measurement. A significant portion of this inaccuracy may be due to the laser vibrometer error.

4.4 Comparison of Instrument Signal Quality

It is important to understand the key differences between the different types of sensors. Various parameters lend themselves to being more suitable for different types of measurements. One particularly important comparison is that of signal quality. For that, we measure the performance of each sensor under the same conditions for a comparison of noise levels and signal-to-noise ratio.

To compare the sensors we have them image the same scene, namely the speaker setup with a black and white edge described in Section 4.2 Each sensor was placed so the sensor array was perpendicular to the direction of

49

motion as illustrated by the optical path in Figure 14. The single element photodiode system was placed 45 cm from the black and white high contrast edge, the Fastcam 80 cm and the Pike 61 cm. The focal length of the lenses were 16 mm, 50 mm and 16 mm respectively.

The goal was to compare the signal of each type under optimal conditions, trying to achieve the highest signal-to-noise ratio. To do this, each camera operated at its fastest frame-rate at a resolution of 256x256. The images were co-added to the highest value that would still allow an effective frame rate of 40 frames per second. This rate was selected since it was the lowest effective frame rate after co-adding that both cameras' true frame rates could have in common. This allows for a low frequency range of 0-20 Hz to be measured, a common range for many of the applications discussed in this dissertation.

The black and white edge was illuminated by an incandescent light bulb at a level of 0.5 V as measured by the single element photodiode. 1.0 V is typical of a white surface illuminated by direct sunlight so a 0.5 V measurement of a equally divided black and white surface would approximate a measurement made outdoors.

The Fastcam was set to a resolution of 256x256 with the fastest frame rate of 28,000 frames per second. To keep an effective frame rate of 40 frames per second, 700 Fastcam frames were co-added after each one was dark subtracted. The Pike was also set to a resolution of 256x256 with the fastest frame rate of 360 frames per second. This allowed for an effective frame rate of

40 frames per second with 9 frames co-added. In terms of bits, this processing added 3 bits to the Pike and 9 bits to the Fastcam for effective bit depths of 17 and 19 respectively. It is evident that the fast frame rate of the Fastcam gives it an advantage in bit depth due to the number of frames that can be added in the measurement time interval. The single element sensor acquired at a rate of 2048 samples per second and has an effective bit depth of 19 to 20 bits at 1.0 V.

The Metro Laser Vibromet V500 vibrometer was also used contemporaneously to determine the absolute displacement of the speaker cone motion during the measurement. The vibrometer operated at 2048 samples per second on the most sensitive setting of Lo with a 1.0 kHz low-pass filter applied. The velocity of the cone was determined by measuring the laser vibrometer signal in LabVIEW and the amplitude of displacement was calculated through integration. The velocity is related to displacement by the following:

$$x = x_0 \sin \omega t$$

$$v = x_0 \omega \cos \omega t$$

Since we are only interested in the modulation of the signal, the amplitude is unaffected by a change in phase so that, with a 90° phase change , we have

$$v = x_0 \omega \sin \omega t$$

Now,

$$\frac{x}{v} = \frac{x_0 \sin \omega t}{x_0 \omega \sin \omega t}$$

$$x = \frac{v}{\omega}$$

where $\omega = 2\pi f$ and f = frequency of the cone, 10 Hz for this measurement. The result as measured by the vibrometer is a displacement amplitude of 14.1 μm for this comparative test.

In addition to co-adding frames, all the pixels along the edge are added together to increase the signal-to-noise ratio which will explained in detail in Section 5.2.2 For the Pike we add 288 pixels over 2 columns and for the Fastcam 160 pixels over 1 column. The number of columns was chosen based on the highest signal-to-noise ratio and the fact that the modulation can be blurred over more than one column from the optics.   Each sensor acquired data for a trial time of one second. For the cameras, the pixels were added in the temporal domain to make a single virtual pixel before a Fast Fourier Transform was applied.

The signal-to-noise ratio was determined by integrating under the peak frequency of interest, 10 Hz, in the frequency domain. That result was divided by an integration in the spectrum of equal length where no signal was seen. The results show that the single element photodiode sensor had a signal-to-noise ratio of 35.5, the Pike had a value of 35.1 and the Fastcam 73.0. The relative spot sizes for the camera were 2.81 mm, 0.28 mm and 0.27mm respectively. These results show that the Fastcam has the highest signal-to-noise ratio of the three sensors. The most likely factors responsible for this performance is the

high speed of the camera allowing a large number of frames to be co-added as well as the use of the 50 mm lens.

The various contrasting qualities of the two technologies, the single element sensor and video cameras have been indicated in the following table, Figure 17.

Single Element Sensor	Video Camera Array
Targets only one location	Simultaneously samples multiple locations
20 bits per sample	10-14 bits per sample
Sample rates up to 102,400 S/s	Full frame sample rate of 108 fps (Pike) and 3000 fps (Fastcam)
USB interface	Firewire and Ethernet interface
Acquisition Length limited by Computer Disk Space	Acquisition limited by Disk Space (Pike) Acquisition limited by onboard 8 GB Storage (Fastcam)
Low System Usage	High System Usage (Pike) Low System Usage (Fastcam)

Figure 17. Comparison of the single element to the video array. A table shows the complementary qualities of the two technologies used for measurement, the single element sensor and the video sensor array.

4.5 Problems and Mitigation

4.5.1 Sensor Jitter

One of the largest problems associated with this technology is sensor jitter, that is, any signal originating from motion of the sensor and not the object

being measured. The result of this type of motion can be modulation in the signal that is not from the source but instead from something close to the sensor, for example ground shake coupled to it through its support. These signals can be deceiving and often appear no different than the signal originating from a distant target, especially when urban seismic noise is present.

Lutz Hall Ventilation
Vibration of Ventilation Unit through Edge Detection
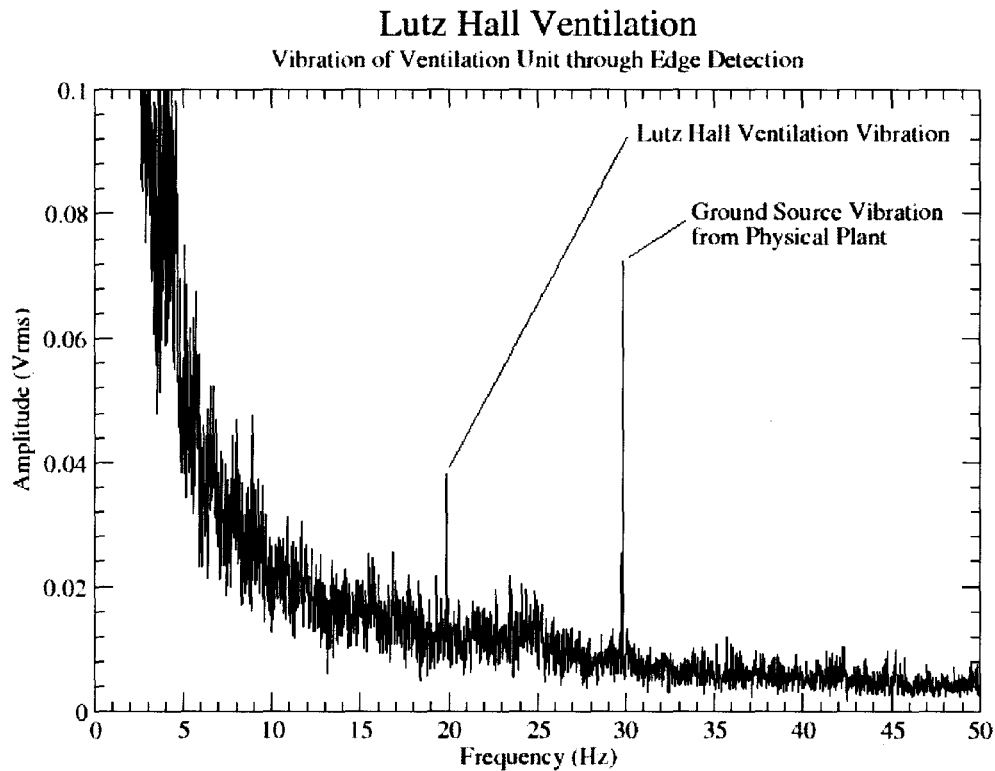
Figure 18. Ground source jitter originating from a nearby heating and cooling facility shows a clear signal in the frequency spectrum of a single channel sensor. If the source of a signal cannot be determined, signals originating from an imaged distant object may be indistinguishable from jitter of the sensor. In this case one component of the signal is originating from a ventilation unit being imaged on the sensor while another is originating from the sensor itself shaking because of vibrations coupled to the ground. Without this prior knowledge of background features it is difficult to attribute the source in a single measurement of a new target.

Figure 18 shows a measurement taken on a ventilation unit on top of Lutz Hall from a location in front of the University's Physical Plant, which houses the University's heating and cooling facility. The spectrum clearly shows the vibration due to the ventilation unit at 20 Hz, or 1200 rpm. However an even larger signal at 30 Hz is evident. The 30 Hz signal was verified by a Physical Plant employee to be the operating rate of the pumps housed in the Physical Plant.

Attributing various signals with the video sensor technology can be done by looking for differential signals, that is for frequencies that appear to be unique to certain areas of the image rather than throughout the entire scene. It would be expected that if a signal arises from sensor jitter then the entire scene would be modulated at levels based on the signal difference across the individual edges in the scene. If such a common signal arises, it can be rejected as originating from local effects, while the unique frequencies in the different sections of the scene can be attributed to those distant objects being imaged.

The single element sensor cannot differentiate between multiple points in the scene simultaneously, although it can be used for sequential measurements at precisely defined points in the scene to the same effect. It is also possible to use more than one single element sensor to sample simultaneously various points within the scene and look for differential motion. A second sensor can also be positioned to image a target that is known to be static. In this way, any motion seen from that sensor can be attributed to sensor jitter. In this scenario it is important that each sensor be physically attached to the same platform so any motion of the sensors is the same.

Beyond analyzing the signal it is possible to reduce or characterize sensor jitter in other ways. Vibration reduction technology can be used to reduce unwanted signals from reaching the sensor, such as passive air cushions or tripod vibration reduction pads, and active stabilization platforms. Accelerometers can also be placed near the sensor location to determine the frequencies present in the ground or the platform motion so they can be rejected.

CHAPTER V:

DATA PROCESSING

5.1 Co-adding Frames for Increased Dynamic Range

As we have seen in the comparative lab tests one of the most effective methods of increasing the dynamic range of the video sensor is done through the co-addition of frames. Each pixel sampling the same spatial location is added to the next successive pixel on the temporal axis. This oversampling, is done by increasing the frame rate of the camera and adding successive frames in memory to double the on-chip analog bit depth for each factor of two in frames added. For example, adding two frames increases the bit depth by one, and adding another set of two beyond that for a total of four frames added frames gives a second bit and so on. Some cameras like the Allied Vision Pike FO32B will do this on the camera up to 16 bits. This has two advantages, 1) the added dynamic range in the sensor, and 2) the reduction of bandwidth coming from the camera as the pixel values begin to fill the already allocated space of 16 bits for each frame. The normal mode allows for 14 bits of data in a 16 bit image, so the camera is already streaming a 16 bit image. Adding four frames allows for 4, 14 bit images to be compressed into one.

A key factor in this technique is the necessity to be able to nearly saturate the frames before they are co-added, assuming the detector stays non-linear. If not, the levels would need to stay within the non-linear regime. The Pike camera uses a Kodak KAI-0340 CCD, an interline transfer sensor with a read noise of 14 electrons, 7.4 μm square pixels, and a well depth of 40,000 electrons. Thus at full well the Poisson shot noise is 200 electrons, much bigger than the read noise. A 14-bit ADC saturates at half the well depth. Without doing this there would be no advantage in increasing the bit depth with the co-addition of frames. Basically each frame would be reduced to half of the pixel value because the integration time of each frame is reduced. When the pixel is saturated, halving the integration time leads to the pixel values being reduced by some value less than half so there is a net gain in photons. The goal is to increase the total number of photons collected and thereby reduce the shot noise in each pixel.

Co-addition can be done to the limit at which the frequency of interest can still be resolved. For example, if a 20 Hz signal is being measured, the Nyquist limit requires at a minimum a 40 samples/second acquisition rate, so for a camera operating at 320 frames per second, at least 4 frames can be co-added for a reduction in frame rate to 80 samples per second and a bit depth increase of 2. [23] In practice a frame rate above 40 samples per second would be ideal to fully resolve a 20 Hz signal.

5.2 Binning Pixels to Increase Sensitivity

Another technique used to increase the sensitivity of the system is to spatially bin pixels. This increases the photon count, producing a virtual pixel

59

larger than the physical one and still providing multiple samples from an extended target. This can be especially advantageous when the vibrating edge spans multiple pixels in a scene. There are two ways in which pixels can be binned: parallel and perpendicular to the direction of target motion. We look at both cases.

5.2.1 Binning Parallel to Motion

The smallest target motion detectable in the system is

$$\delta x_o = \left(\frac{r}{f}\right) w N_e^{-\frac{1}{2}}$$

where f is the focal length of the optical system, r is the distance to the target, p is the width of a signal pixel, and N_e is the number of photoelectrons in the system.
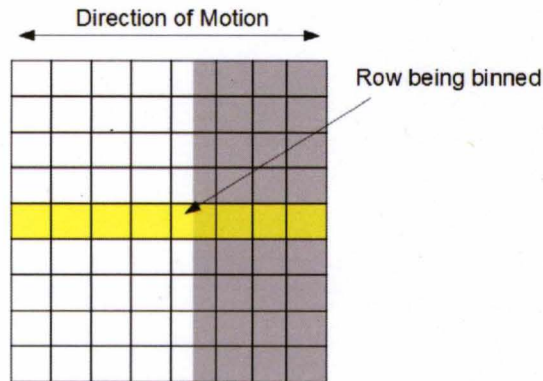
Figure 19. Binning pixels parallel to the direction of motion is one way to increase the signal count by creating a virtual pixel. Pixels parallel to the direction of motion offers no gain in signal-to-noise however.

60

We now look at binning pixels in the direction parallel to the motion of the target as imaged on the sensor to see the effect on the system's sensitivity to motion. This would increase the effective well depth thus increasing the photon count and in turn increase the signal-to-noise ratio. As we bin pixels $N_n \to N_b N_e$ where N_b is the total number of pixels being binned. Likewise, the width of the "binned pixel" becomes $p \to N_b p$. Now we have

$$\delta x_b = \left(\frac{r}{f}\right)(N_b p)(N_b N_e)^{-\frac{1}{2}}$$

Taking the ratio of the two we have

$$\frac{\delta x_b = \left(\frac{r}{f}\right)(N_b p)(N_b N_e)^{-\frac{1}{2}}}{\delta x_o = \left(\frac{r}{f}\right)p N_e^{-\frac{1}{2}}}$$

$$\frac{\delta x_b}{\delta x_o} = N_b b^{-\frac{1}{2}} = N_b^{\frac{1}{2}}$$

By binning additional pixels we see a decrease in the sensitivity to the smallest detectable target motion by a factor of $N_b^{\frac{1}{2}}$.

Although the increase in the total photon count by adding pixels does increase signal-to-noise, we are not interested in the overall signal for the new effective pixel, but instead interested in the modulation. It is unaltered by binning pixels in the direction parallel to the target motion. In the frequency domain, we can view this as increasing the signal in the DC bin while the signal level of the modulation remains the same. The overall effect of binning pixels in this case is

an actual increase in noise by an amount proportional to the number of real pixels binned. Since this process is governed by photon noise we see the standard factor of $N_b^{\frac{1}{2}}$ which is responsible for decreasing the sensitivity of the system.

5.2.2 Binning Perpendicular to Motion

Now we look at binning pixels in the direction perpendicular to the target motion. Note that this is only of interest in the event that the motion is being detected from an edge imaged parallel to the pixels being binned such that the motion across that edge is correlated. As before, the smallest target motion detectable in the system is

$$\delta x_o = \left(\frac{r}{f}\right) p N_e^{-\frac{1}{2}}$$

In the situation where we bin perpendicular to the target motion, again $N_n \rightarrow N_b N_e$. However w does not change as this is the pixel width parallel to the motion. Now looking at the ratio of the two we have

$$\frac{\delta x_b = \left(\frac{r}{f}\right) p (N_b N_e)^{-\frac{1}{2}}}{\delta x_o = \left(\frac{r}{f}\right) p N_e^{-\frac{1}{2}}}$$

$$\frac{\delta x_b}{\delta x_o} = N_b^{-\frac{1}{2}}$$

So we see that binning in the direction perpendicular to motion increases the sensitivity to detect target motion by a factor of $N_b^{\frac{1}{2}}$.
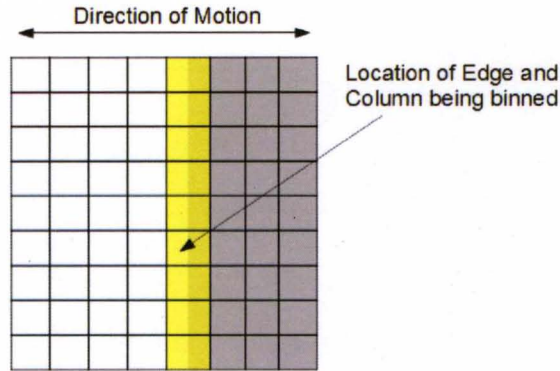
Figure 20. Binning pixels perpendicular to the direction of motion can also increase the signal. Binning pixels perpendicular to the direction of motion offers an increase in signal-to-noise. This differs from the previous case in that we are binning over pixels that see a modulation and effectively increase the desired signal that results from that modulation. In the frequency domain, this amounts to increasing the signal resulting from modulation by a factor of N_b, while the noise is increased by a factor of $N_b^{\frac{1}{2}}$ as governed by photon statistics. This, in turn, gives us an effective factor of $N_b^{\frac{1}{2}}$ increase in signal-to-noise in the signal of interest.

Looking at the combination of the two binning circumstances we see that they effectively cancel each other in terms of the effect of the system's sensitivity to detect motion. This would be the same as acquiring over a single larger pixel of proportionately larger size.

5.3 Handling Large Data Sets

A single element sensor alone can give us information about the target, such as local motion, vibration, and displacement. However to gain spatial

resolution a multi-element array is necessary if data are to be acquired simultaneously at various locations on the target.

As we acquire imaged data in real time, it can easily be seen that the size of the data sets become quite large. When using a standard high speed camera, even at relatively slow frame rates of a few hundred frames per second, at a high resolution one can see that some method would be necessary to process and inspect the data for information of importance under such high bandwidth.

One technique that can be utilized is transforming the time domain into the frequency domain and replacing the temporal dimension with the frequency domain. For example, suppose we are acquiring data with a camera that has a resolution of 640 x 480 pixels. We are interested in frequencies below 100 Hz so to meet the Nyquist criteria we sample at a frame rate of 200 frames per second. [23] In this scenario we collect 200 frames every second, each with a resolution of 640 x 480 and requiring 16 bits of storage. If you can imagine stacking these frames one after another, we build a three dimensional data cube that consists of spatial dimensions in the x and y dimension, and time in the z direction giving us a data set that has the dimensions 640 x 480 x 200. Each element of the array can be thought of as its own sensor so that we have 307200 sensors sampling at a rate of 200 samples per second. If we were to process all of these sensor with a Fourier Transform we would then recover the frequency domain for that element. That new frequency domain data set could then replace the time domain data set for each pixel giving us a new data cube that now has two spatial dimensions in the x and y direction and frequency in the z direction

64

with dimensions 640 x 480 x 100. This data cube will be discussed in more detail in the following section. Recall that the Nyquist sampling criterion limits the highest frequency to half the frame rate. [23]

5.4 Creating a 3-D Data Cube

In the case of the single element sensor, a series of data, consisting of a 1 x N array would be acquired and processed with a Fourier transform. The result is an $\frac{N}{2}$ set of data in the frequency domain where N is the number of samples in the time domain. This a 1 x $\frac{N}{2}$ data array that details the frequency spectrum at a particular spatial location in the scene, wherever the single element sensor is imaging. However, when data are acquired with a multi-element sensor, this allows for multiple points to be monitored simultaneously. The series of "images" represents the spatial field and the combination of fields results in an $X \, x \, Y \, x \, N_i$ data array where X represents the spatial dimension across the width of the image, Y represents the spatial dimension across the height of the image, and N_i represent the number of frames. Since each frame contains multiple samples, each frame N_i will contain $X \, x \, Y$ samples.
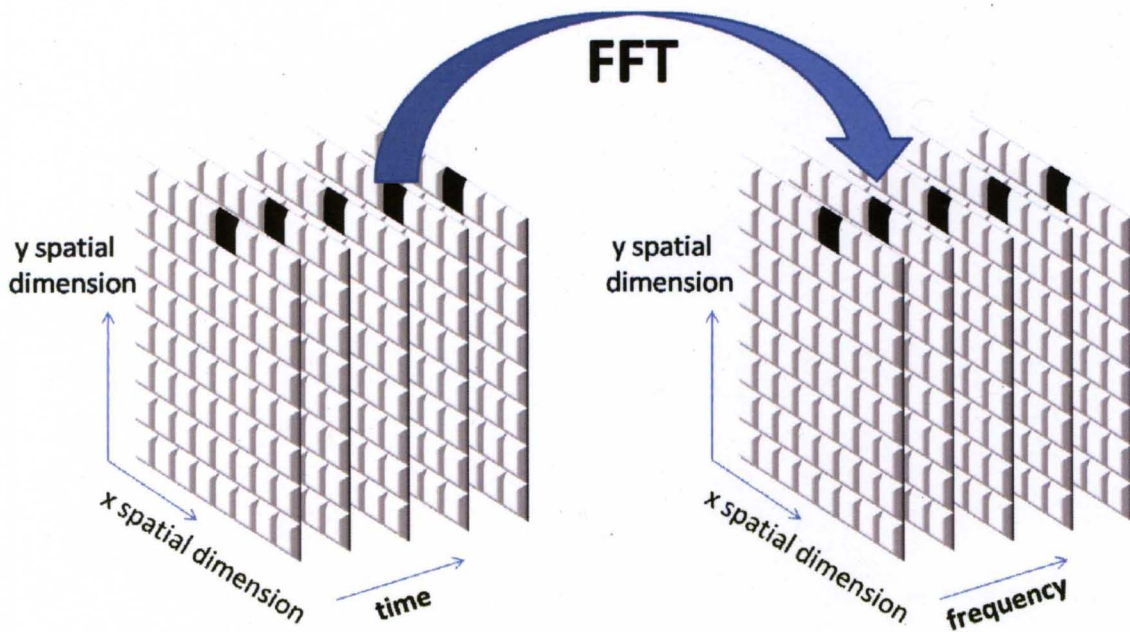
Figure 21. Transforming a time series data cube into a frequency series data cube is done by performing a Fourier transform on each series of pixels in the time domain. Depicted is a representation of how a time series is transformed into a frequency series which then replaces the temporal axis with the frequency axis to create a frequency cube.

Now we treat each pixel in the image just like a single element sensor, applying a Fourier Transform to the series of measurements in time. The result is that the time domain series is now replaced by a frequency domain series. If there were N_i frames acquired with each frame representing the scene at a particular point in time, we now have $\frac{N_i}{2}$ frames each representing a particular frequency. As with any Fourier Transform, we are limited by the Nyquist theorem so that the initial frame of the sequence depicts the 0 Hz component and the last element depicts

$\frac{R}{2}$, where R is the sampling rate. [23] In whole, the transform of each pixel results in a $X \ x \ Y \ x \ \frac{N_i}{2}$ data volume enclosing the X dimension of target width, the Y dimension of target height, and the $\frac{N_i}{2}$ dimension of target frequency.
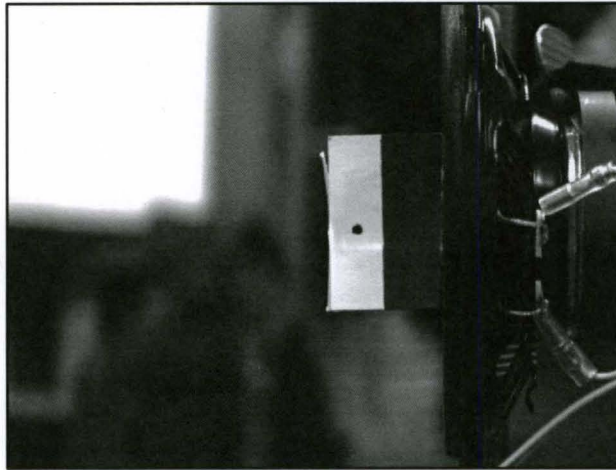
Figure 22. A black and white edge to be imaged in 3D is shown. This edge is used to capture a series of time domain images for transformation into a frequency domain data cube. The edge was attached to the speaker and oscillated at 10 Hz.

The advantages of acquiring data with a series of images is that multiple points in the scene are acquired simultaneously. This means that the resulting data cube gives the frequencies found within the scene as well as the location at which those frequencies occur spatially.
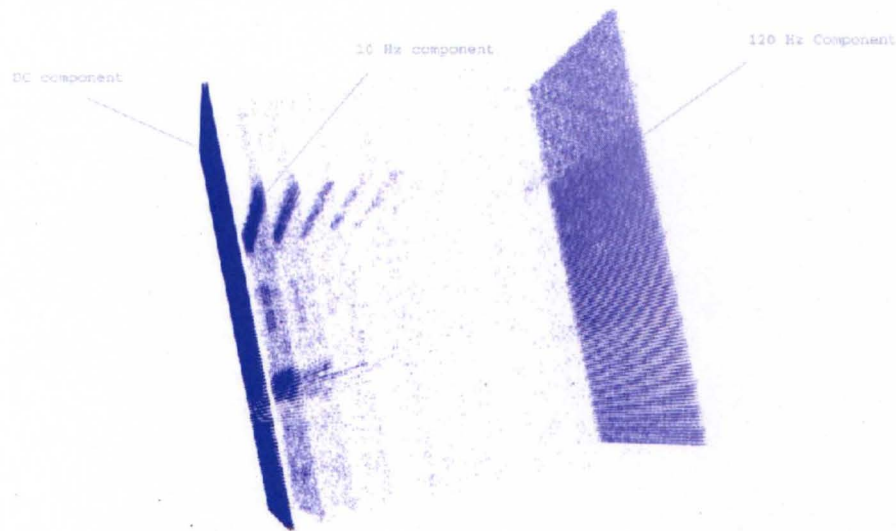
Figure 23. A frequency domain data cube is shown. The figure is in inverse grayscale with the darker features indicating more signal. The data cube was created by performing a Fourier transform on each series of pixels in the time domain and replacing those values with the frequency domain values. The result is a cube with frequency replacing the time axis. The left side of the image shows a signal slice representing the bright DC portion of the image. 10 units to the right is the slice representing the 10 Hz component. The edges are clearly seen due to their high modulation. Various harmonics are seen in slices to the right. The 120 Hz slice is represented on the right which is due to the room illumination from AC lighting. In that plane the white target strip is the data and the dark spot on the target is the white spot in the data. In the principle 10 Hz component only the edge features that strongly modulate the signal are apparent.

When we transform the time domain signal to the frequency domain we lose the temporal information in lieu of the frequency information. However, we can regain some temporal data by acquiring over a given time step, processing that data with a Fourier Transform to create a data cube, and then repeat that process over the next time step. The result is a series of time steps that can be viewed sequentially as a series of data cubes, that is, a 3D movie. This gives us the temporal resolution of the length of one time step equal to the length of an acquisition interval in the temporal domain. For example consider acquiring at a rate of 100 frames per second continuously. If you process each set of 100 frames you end of with a data cube showing 0-50 Hz representing 1 second of acquisition time. If you continuously do this you create a series of data cubes each representing 1 second of data, giving you a sequence of 3D snapshots at 1 second of temporal resolution. Supplemental Video 1 shows a time series of data cubes. The data was collected on a black and white edge attached to a speaker as shown in Figure 22. The speaker's frequency of vibration was increased in 10 Hz intervals over time starting at 10 Hz. The video shows the data cube being rotated to show the spatial and frequency axes. Black indicates more intense modulation while white is less. The black and white edge is clearly evident in the data cube and can be seen to increase along the frequency axis indicating a faster vibration rate.

More finely grained 3D temporal sampling can be achieved if instead of shifting the window in which you process the Fourier Transform by a multiple number of frames, you only shift it by one sample or frame. In this way you

change the starting element of the Fourier transform by one point such that each data cube overlaps with the previous cube in terms of the number of frames processed by N-1 where N is the number of frames that are used in the Fourier Transform. The means that the cube is only altered each iteration by a single sample. The temporal sampling that can now be achieved is

$$\Delta t = \frac{R}{N_i}$$

where,

N_i is the number of frames processed in the Fourier Transform

R is the sample rate for frame rate

Δt is the temporal resolution

Keep in mind that the each data cube still represents 1 second worth of data in this example or an arbitrary amount of time equal to the acquisition time. This means that the data cube contains information that occurs over the entire acquisition so any motion in the scene will be still appeared "smeared" as it has occurred over the entire acquisition window.

Figure 23 shows a frequency data cube of the vibrating speaker that is in Figure 22. Attached to the speaker is a piece of paper that shows a high contrast edge. The Fourier Transform process has changed the time dimension of the series of frames into a frequency dimension. The data cube shows the various "slices" depicting these frequencies. The speaker was vibrated at 10 Hz and illuminated in 60 Hz fluorescent room lighting (120 Hz lighting modulation) with

an incandescent bulb using the equipment and setup described in section 4.2 The data cube is labeled to show prominent features in the frequency space. On the left is the DC component, that is the component of the cube that has no oscillation in illumination. The first prominent feature to the right of DC would be the 10 Hz component from the speaker oscillation itself. The edges are clearly evident, one from the black and white boundary and another from the white end of the paper in contrast to the dark background. The black dot provide another high contrast edge. Continuing to the right we see a series of harmonics at 20 and 30 Hz and so on as they diminish in intensity. Finally we see the 120 Hz component from the room lights. Although the AC lights operate at 60 Hz, the detector see the power associated with the lights which is expressed as a cos^2 term. This mean that the 60 Hz voltage oscillation appears as a 120 Hz illumination oscillation to the detector. Essentially both the positive and negative amplitude of the voltage oscillation in the lamps produce a brightening.

It is clear from the data cube in Figure 23 that each frequency is clearly resolved. The frequencies present in the scene, as well as their location, can easily be determine from the data cube. Also, it is clear that unwanted frequencies can be rejected easily as the 120 Hz oscillation from the room lights is clearly separated from other frequencies in the scene. Furthermore, bright DC backgrounds are easily rejected as this component is only present in the lowest frequency slice.

The software used to render the data cube in Figure 23 as well as in Supplemental Video 1, was Partiview, open source code developed for 3 dimensional visualization. [24]

CHAPTER VI:

GPU COMPUTING

Significant amounts of data are being streamed from the camera when using video technology as a source for data collection. We have seen for example, that a 640x480 pixel image (0.3 megapixel) contains 307,200 elements. For post-processing this is not much of a problem. Data of that size can be easily handled, but if the data of interest would need to be processed in real time this can pose a significant problem. In this case the processing necessary to handle this sort of data acquisition in real time would involve applying a Fourier Transform to the temporal axis of each set of pixel values in the same time it took for that set of frames to be acquired. Typically for the Pike F032B, this would amount to 307,200 Fourier Transforms, each of length 107 frames, in one second. This would represent a full frame acquisition at the fastest frame rate with the operation being performed on one second worth of data. Modern day CPU's are still not capable of handling the processing of data sets of such size in real time. In fact, the CPU is generally occupied with streaming the image files to disk during this time. However, the Graphics Processing Units, or GPU's are highly optimized for type of application. [25,26] The reason is because the Fourier Transforms are highly parallelizable, that is each series of pixels values are independent of one another and can be processed in parallel. [25] NVIDIA

73

GPU's are being built with numerous cores, as many as 512 per card at the time of this writing. In fact, applications are being targeted by GPU manufactures and high performance computing is an available option for the desktop because of the GPU. The most recent computer to take the title as the fastest supercomputer in the world was built around GPUs. [27] The NVIDIA technology in particular has a robust API and compiler under the Compute Unified Device Architecture (CUDA), developed by NVIDIA. The language is written natively in C and has thorough examples and documentation freely available for download to gain entry into GPU computing. [26]

Figure 24. NVIDIA Tesla Card is shown. The NVIDIA Tesla 1060 seen here has 240 cores per card, each operating at 1.3 Ghz, all capable of running in parallel. Credit: Mahogny, Creative Commons.

A GPU based processing system could acquire the data from the system, streaming directly to disc. Once a determined number of frames were acquired, the GPU unit would proceed to perform a Fast Fourier Transform on each set of pixels in parallel. In the case of the Pike FO32B, acquiring at a frame rate of 107 frames per second at a resolution of 640x480, this amounts to doing 307,200

Fourier Transforms, each of 107 in length in one second. Using a GPU with 448 cores, this would mean that 686 Fourier Transforms would need to be made per core in one second for this to be achievable.

6.1 The GPU Test System

Our system consists of two GPU units, an NVIDIA Tesla 1060 and a Quadro FX5800. Both have 240 CUDA cores operating at 1.3 Ghz for a total of 480 cores available for processing. Each card has 4 GB of available RAM allowing large series of images to be transferred and stored in GPU memory making the process faster and more efficient. The machine being used to process the data is a Dell Precision T7500 Workstation, with 4 quadcore W5590 Xeon processors, each operating at 3.33 Ghz. The system has 24 GB of system memory. Although this is not typical of a standard workstation, the system's purpose here is to provide a proof of concept and show feasibility that the necessary computation can be performed in real time. The envisioned machine would more likely be a laptop capable of performing similar operations. Today the highest end laptops rival the described system. Systems are available with mobile platforms that operate with six core Xeon processors operating at 3.33 GHZ, 24 GB of system memory, and dual GeForce GTX 460M GPU units offering a total of 704 CUDA cores with 4 GB of GPU memory. Although these GPU units are designed for gaming they are fully capable of offering the computing capabilities at much lower cost than the Dell Workstation.

A benchmark was performed using our test machine to check the feasibility of performing the necessary Fourier transforms in real time to process

the camera data. The code was written using the CUDA platform which handles parallelizing the processing so that the FFTs are performed in parallel on the GPU. The benchmark code was based on code written by Kashif Rasul. [28] The code implemented only utilized one of the two GPUs on the system, the Tesla 1060. Obviously the speeds would significantly improve if both cards were being used to process the data.

GPU Benchmark of FFTs
NVIDIA Tesla 1060
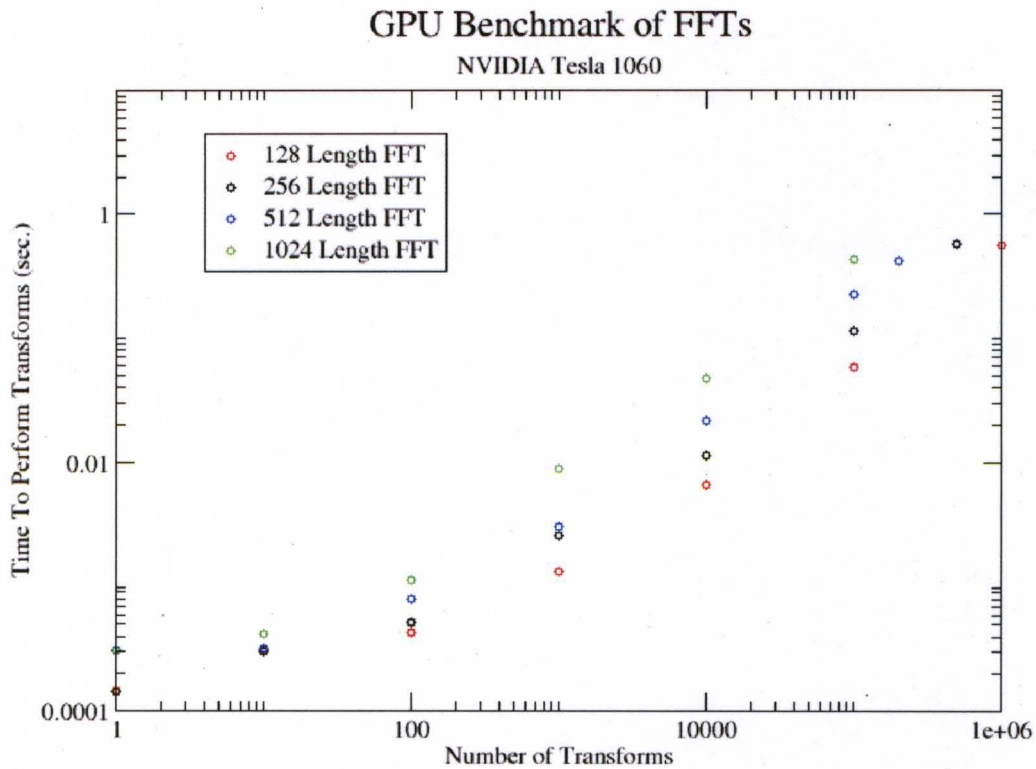
Figure 25. Benchmarking the graphics processor unit shows that the necessary Fourier transforms can be performed to process 307,200 pixels with 107 frames. Three different length FFT's were performed, 128, 256, 512, and 1024. With the Pike operating at 107 frames per second, the 128 length FFT approximates the amount of data streamed from the camera. Up to 1 million FFTs of length 128 were performed in under 1 second in this test.

Figure 25 shows the result of benchmarking the GPU. Multiple tests were performed of various length FFTs to determine the time it takes to process multiple elements. The data clearly shows that, as you would expect, the more

77

FFTs performed, the longer it takes; and likewise the larger the FFT, the longer it takes. What is more important however is that not only can 307,200 FFTs of size 128 be performed under 1 second but a total of 1,000,000 can be performed in less than a second. This clearly show that a data cube consisting of images 640x480 in size with 107 sample length (1 second worth of data from the Pike camera running at 107 frames per second) could be transformed into the frequency domain in real time with time to spare. In fact a sample length of 256 can be done in under a second with the frame size of the Pike. The would allow a camera with a 640x480 sensor size to run at a speed of 256 frames per second and be processed in real time.

CHAPTER VII:

APPLICATIONS

7.1 Structural Health

7.1.1 Bridges

It is well documented that the current US infrastructure is aging and much of it is in need of repair. Railroad and highway bridges are critical components of that infrastructure. Recent studies have shown that 12.1% of the bridges in the United States are considered structurally deficient with another 14.8% categorized as functionally obsolete. [29] Bridges were generally built to last 50 years with the average bridge now 43 years old. [29] Current methods of analyzing bridges to test for defects are outdated, often subjective, and at times based on invalid information. [30] A system that could remotely and quickly assess a bridge's structural health at low cost would be of interest and value. [31]

It has been shown that variations in the natural resonances in bridges can be useful as an indicator of the structural health. [32] In addition, a system capable of a remote measurement of these resonances could be useful in helping to indicate a bridge's structural health, especially in aftermath of a large scale catastrophic event, i.e an earthquake, when numerous bridges are in need

of assessment. A remote system could make multiple measurements of different bridges quickly and with relatively low cost and resources.

The system discussed has proven capable of measuring very small motions optically and has shown it is capable of revealing resonant properties of bridges that may be used to ascertain its structural health. To demonstrate this, two bridges have been analyzed using our system for the detection of motion and analysis of structural dynamics: the Eastern Parkway Bridge carrying Eastern Parkway passing over the CSX railroad at Floyd and Brook Streets, and the Beulah Church Road bridge for I-265 passing over Beulah Church Road .

The first, the Eastern Parkway Bridge, is a concrete bridge that spans two roads and a railway as seen in Figure 26.

Figure 26.   The Eastern Parkway Bridge while under re-
construction in 2010 is shown. The bridge is located in Louisville

Kentucky and spans Brook St., Floyd St., and a CSX railway

crossing.   The bridge underwent repairs beginning in June 2009

due to degradation of its reinforced concrete.   The majority of the

piers were demolished and repoured, new railings were added, and

the sidewalk was replaced.

The Eastern Parkway Bridge was chosen because of its close proximity to our

laboratory and its state of decay.   The bridge was scheduled for renovation

starting in August of 2009 which allowed us to make measurements on the

bridge before, during and after construction. In addition, the construction allowed

us to have access to the underside of the bridge deck to outfit an *in situ*

accelerometer due to the fact that the construction crew already had the

equipment in place and assisted in affixing the sensor.

An accelerometer was placed on the underside of the bridge deck at a location determined by three primary factors: a) in the middle of a span to maximize motion, b) in close proximity to a building that would house the data collection system, and c) in a location that showed no signs of water damage to the bridge to prevent the equipment from getting wet.

Figure 27. The accelerometer affixed to the Eastern Parkway Bridge is located mid-span between piers 5 and 6, numbered from the abutment on the western side of the bridge. The accelerometer is a Crossbow CXL02TG3 capable of measuring all 3 axes of motion. The data from the accelerometer were recorded continuously and archived over the 2 year project time period.

Figure 28. A close-up of the accelerometer is shown in more detail.

Figure 28 shows the accelerometer in detail. The cabling can be seen extending out from the housing containing the accelerometer and down to a nearby building. The building contains a computer that collects data from the bridge continuously. The accelerometer is a 3-axis Crossbow CXL02TG3 with a sensitivity of 833 mV/g. The data acquisition system comprises of a computer running OpenSUSE 11.1 with LabVIEW software. The signal is digitized with a NI-4472 PCI ADC card. A LabVIEW program acquires data from the bridge at a rate of 2048 samples per second. Each measurement file contains 2048 samples for a frequency resolution 1Hz bins up to 1024 Hz although files can be appended for a longer time base. We initially acquired 1024 samples at 1024 samples per second but later decided it would be valuable to see the frequency response in the 512-1024 Hz range. The data collection runs constantly, 24 hours a day 7 days a week, storing all data to external drives for easy access by exchanging drives.
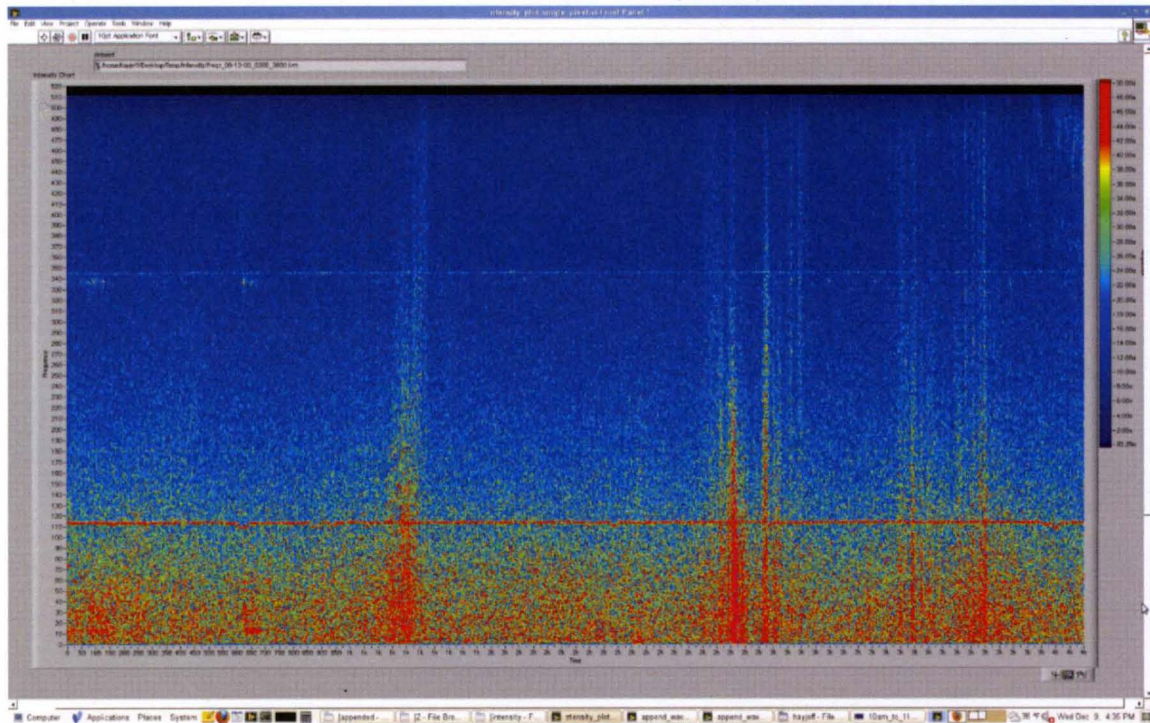
Figure 29. Continuous data collection on the Eastern Parkway Bridge shows a large amount of motion detected by the attached accelerometer. Here we see the LabVIEW analysis of the attached accelerometer data. The x-axis is time showing 1am to 2am while the y-axis is frequency from 0 to 512 Hz. Individual events are clearly scene in the data as sharp rises in the intensity of frequencies. We hypothesis that most of these events are attributed to trains on the CSX line running under the bridge.

The attached accelerometer allows us to have a sensor affixed to the bridge at all times that can be used to calibrate our remote standoff measurement at any given time. In addition, it gives us quantitative information as to the types and amount of motion present.

One of the fundamental properties of bridge is the resonant frequency of the structure. The frequencies arise from natural vibrations in the bridge that are determined by the structural properties of the bridge itself, such as its length of span, material, and support contacts. An accelerometer placed on the bridge can easily detect such motion especially when an event occurs that drives these resonant modes. In fact these vibrations are easily felt when you stand on most bridges, especially with traffic or strong winds exciting their motion.

Prior to the placement of the affixed accelerometer to the underside of the decking, tests were performed on the bridge to measure the acceleration by simply placing an accelerometer on the sidewalk while traffic was flowing. These measurements were performed prior to the closure of the bridge for repairs.

Figure 30. Measuring acceleration of the surface of the Eastern Parkway Bridge was done with a portable system consisting of a laptop, USB ADC, and accelerometer. Data were collected with a Crossbow CLX01LF3 3-axis accelerometer and digitized with a 24 bit an NI-9162 analog to digital converter. Vibrations from passing traffic were easily felt and measured.

Data were collected with a Crossbow CLX01LF3 accelerometer and digitized with an NI-9234 USB ADC. The accelerometer signals were then read by the LabVIEW software on a Panasonic CF-30 Toughbook with Opensuse 11.1 installed.

Accelerations were clearly present on the bridge and could be felt as vehicles passed by. Figure 31 shows the frequency response of the bridge in all three axis as measured from the acceleration. Strong vibrations are present in the low frequency, <100 Hz, portion of the spectrum.

Eastern Parkway Bridge 8-3-09

15 sample average
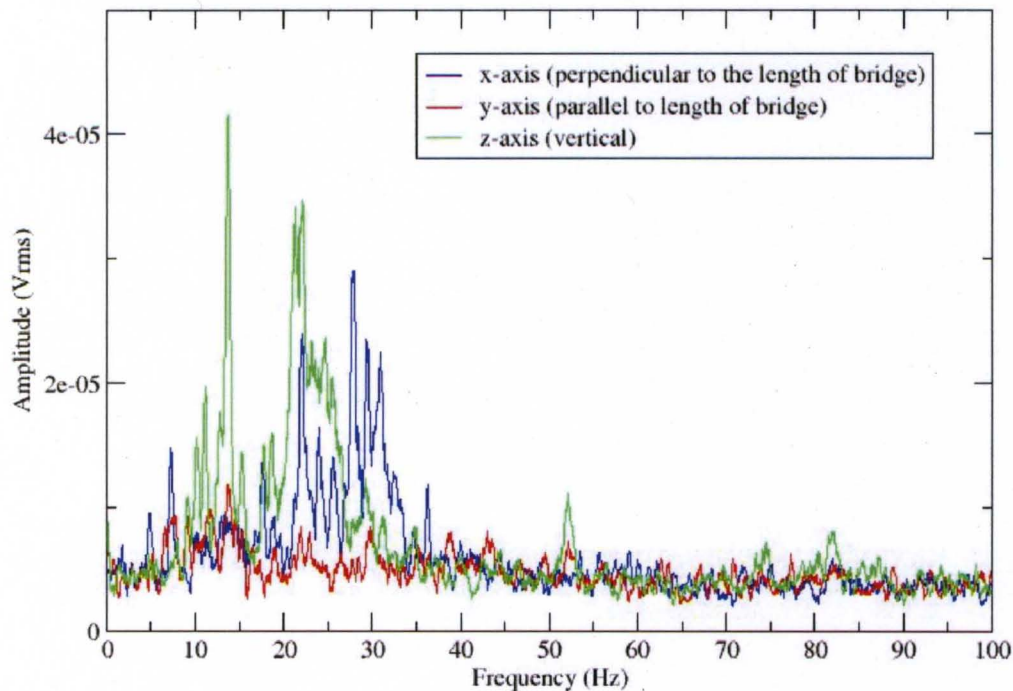
Figure 31. Acceleration of the Eastern Parkway Bridge is seen in all three directions. As expected the vertical component is very strong. The horizontal direction, perpendicular to the length of the bridge is also quite strong as well. The larger strength contributions are on the lower portion of the structural frequency spectrum below 50 Hz.

The ability to see these motions with our optical system would allow these measurements to be made remotely without the need to access the bridge. Measurements were made on the Eastern Parkway Bridge to validate the optical system's ability to measure frequencies found *in situ* sensors on the bridge. The

permanently attached accelerometer offers a validation that the measurements made through the optical system are real and occurring on the bridge. Optical data were collected from the bridge with the with our portable data acquisition system. This consisted of a the photodiode detector coupled through the fiber optical cable to an ONYX 80ED Celestron telescope with a focal length of 500mm at f/6.25. The data were then digitized through the NI 9234 USB analog to digital converter and read into the Toughbook laptop through the LabVIEW data acquisition software. Recalling the formula that determines the smallest detectable displacement on page 32, the longer focal length of the small telescope offers a substantial increase in the ability of the system to see the motion of the bridge.

Figure 32. Data collection on the Eastern Parkway Bridge with the portable optical system. Data is collected with a ONYX 80ED Celestron telescope with a 500 mm focal length. The increase in focal length for this use allows a smaller region to be imaged on the sensor, increasing the spatial resolution. The railing was being removed from the bridge during this data collection process. Credit: John Kielkopf

Figure 32 shows the data collection on the bridge during construction. At this particular time, the construction on the bridge involved removing the last section

of railing over the railroad. In this particular setup the detector was approximately 26.5 meters from the section of bridge being imaged. With a 500 mm focal length lens and 100 μm sensor size the size of the area being imaged on the sensor is 5.3 mm. Data were collected with the optical system by focusing on regions with sufficiently high contrast. Concrete offers many suitable locations of measurement as the variations in the texture of the various constituents, aggregate, cement, etc., offer different levels of reflectivity, often with sharp edges between transitions. Data was collected at a rate of 2048 samples per second with 2048 samples per acquisition. Data was post-processed to append data files to allow for a longer sampling period.

Eastern Parkway Bridge
Sidewalk Removal (4-19-10)
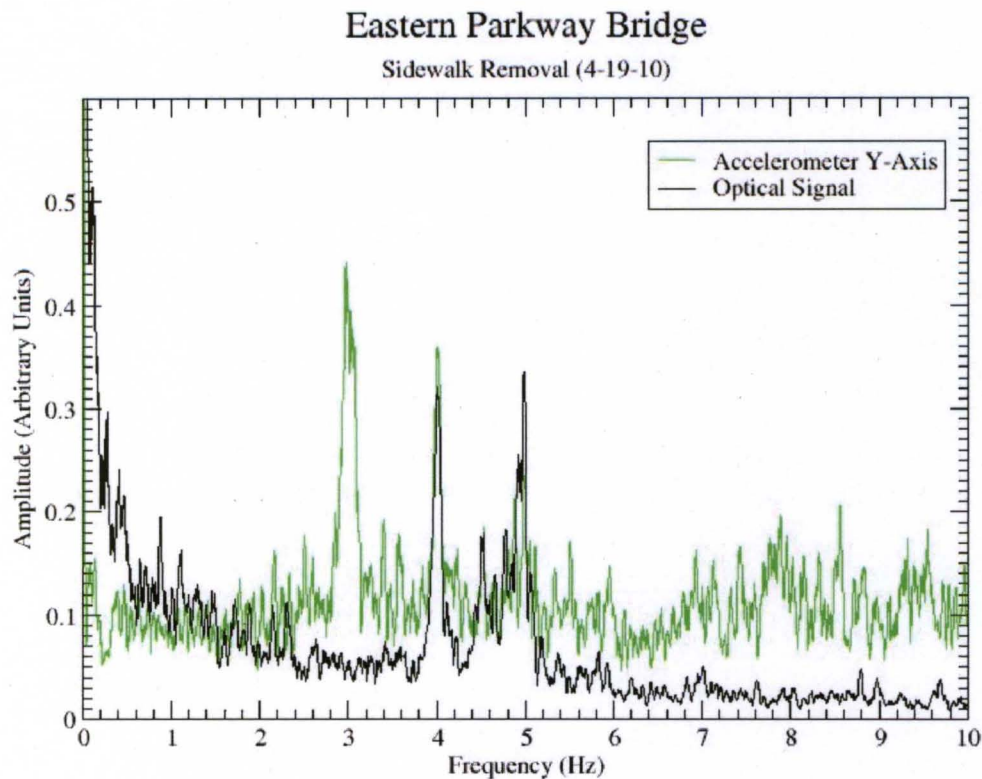
Figure 33. Validation of the optical system with an *in situ* accelerometer is seen. Simultaneous measurements were made on the Eastern Parkway Bridge with the *in situ* accelerometer and optical system. The data show both systems are capable of detecting the same frequency indicating the optical system has made a positive detection of a bridge resonance.

Figure 33 shows comparison of the data taken simultaneously with the optical system and the attached accelerometer under the bridge. The green line shows the frequency spectrum of the bridge as measured by the accelerometer while the black line shows the measurement of the optical system measuring

displacement. A 4 Hz signal is clearly seen in both measurements with a signal-to-noise ratio of 3.10 for the accelerometer and 11.33 for the optical measurement. The accelerometer shows another clear peak at 3 Hz while the optical signal shows another peak at 5 Hz. The 5 Hz signal may be detected by the accelerometer but is just above the noise. These detections may be the result of the locations of the placement of the individual sensors.
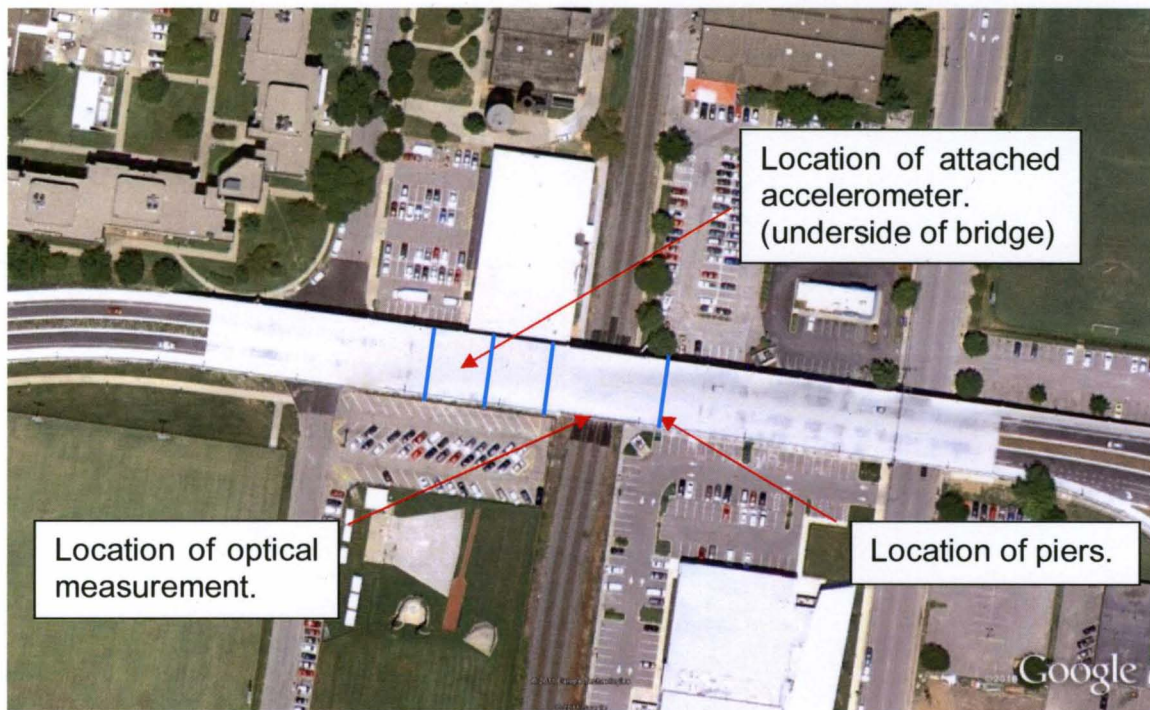
Figure 34. The Eastern Parkway Bridge as seen from above. The location of the individual points of measurement with the optical system and in-situ accelerometer are indicated. The two measurements were made on separate spans with one span separating them. © 2011 Google, © Europa Technologies.

The optical and accelerometer measurements were made a two different locations on the bridge most notably differentiated by the separate spans. They are also separated by a third span. It is possible and most likely that each of the spans exhibits its own mode of vibration and unique frequencies. The accelerometer is much more sensitive to this motion than the optical measurements so the vibration from one span may propagate to others and still be measured by the accelerometer, explaining why the accelerometer is seeing multiple frequencies. The optical signal will be much weaker so the frequencies associated with other spans may not be as strong in detection. In addition, the construction that was occurring at the time was located over the span where the optical measurement was made so it would be likely that that span was being driven to have a higher amplitude of motion. The span on which the accelerometer was placed was likely to have no unnatural driving forces at that time.

Once the bridge was reopened, we were able to access multiple locations on bridge as it has a pedestrian walk path on both sides. A collaboration with the Civil Engineering Department at the University of Louisville also led to a Finite Element Analysis (FEA) being performed on the Eastern Parkway Bridge. The span over the CSX railway was determined by the FEA to be the location that would undergo the maximum deflection, primarily due to it being the longest span. This location was chosen to measure the vibration of the bridge. In addition, the FEA modal revealed the fundamental vibration across the 2 to 3 pier span to be 6.193 Hz. This site was chosen for an optical measurement.

Accelerometers were placed on the sidewalk of the Eastern Parkway Bridge on the CSX span and data were collected. Data were transmitted from the bridge to the laptop using a National Instruments WLS-9163. This is a WiFi cradle that connects to the NI-9234 A/D that transmits a wireless signal over Wireless G directly to the laptop through an Ad-Hoc connection. The device is powered by a Celestron Power Tank 17 portable external DC power supply. An FFT was performed to determine the frequency spectrum at that location.
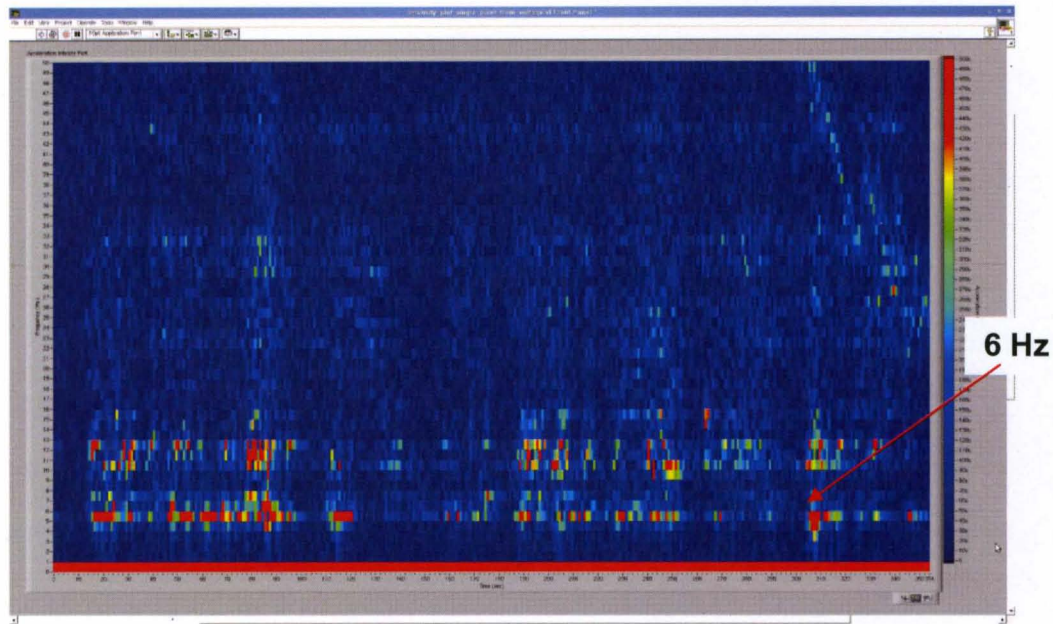
Figure 35. The time evolution of the frequency spectrum of an accelerometer over the CSX railway line on the Eastern Parkway Bridge. Strong amplitude are seen in the 6 Hz region of the spectrum. The x axis represents time in seconds. The y axis represents frequency. Here we see the 0-50 Hz region of the spectrum for the vertical axis of the bridge. The intensity of the spectrum is represented in false color with red indicating the strongest modulation and blue the weakest. The bridge was open to traffic during data collection on 11-18-10 and the short periods of strong modulation likely correspond to traffic passing over the bridge.

Figure 35 shows a series of frequency spectrum determined from the collected data by the accelerometer. Features around the 6 Hz range are clearly

95

seen. Looking at a time interval in more detail that shows more activity, we can clearly see features around the 6 Hz region of the spectrum as seen in Figure 36.

Acceleration
CSX Span
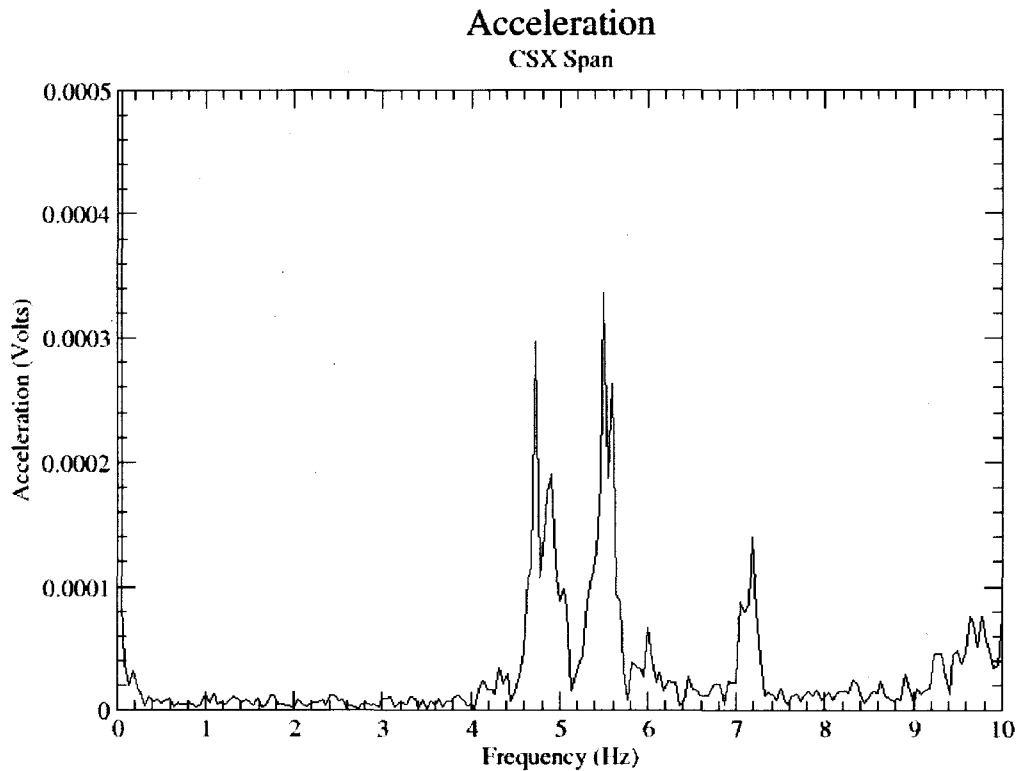
Figure 36. The frequency spectrum of the over the CSX span on the Eastern Parkway Bridge shows strong components near the 6 Hz region of the spectrum. Several modes are seen at 4.7, 4.9, 5.5, 6.0 and 7.2 Hz. The data represent the vertical axis of the bridge and was collected on 11-18-10.

The optical data collected over the 2 to 3 pier span also shows the 6 Hz feature as well as seen in Figure 37 . The optical data was collected with the single element photodiode system along with the Celestron ONYX 80ED 500 mm focal

length telescope. The detector was approximately 26.5 meters from the pier 2 to 3 section of bridge being imaged. With a 500 mm focal length lens and 100 μm sensor size the size of the area being imaged on the sensor is 5.3 mm.

Optical Signal of Motion on Eastern Parkway Bridge
High contrast edge on span between piers 2 and 3
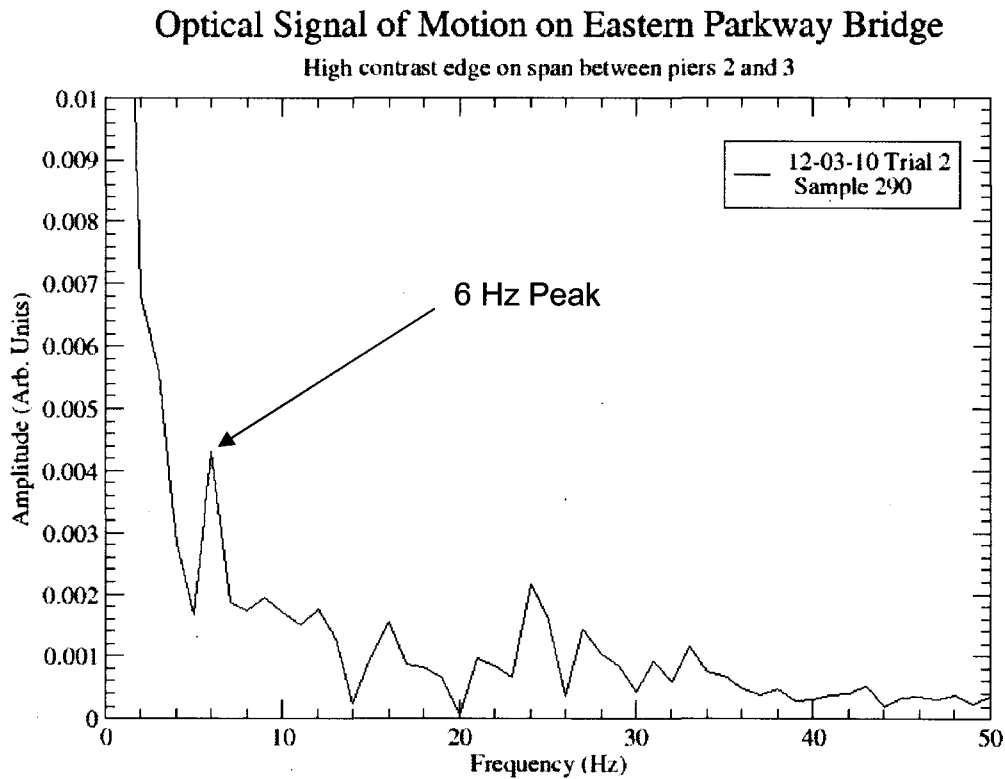
Figure 37. Optical measurement between piers 2 and 3 on the Eastern Parkway Bridge shows a 6 Hz peak and predicted by the Finite Element Analysis(FEA) model.

This offers a validation that the optical measurement can reveal the fundamental frequency in the bridge.

7.2 Machine Health

7.2.1 Mechanical Systems

Mechanical systems exhibit motion that is related to physical parameters of the system. In the simplest form that motion may be a vibration resulting from a fan rotating at a given frequency or a piston oscillating at a given rate. Because these internal functions often produce large amounts of energy, that energy is able to propagate to the surface while still remaining coherent. From there we are able to make measurements on the system and relate to processes at their source.

An example of such a system would be the air handler inside our laboratory. Measurements were made on the air hander to determine the frequency spectrum it produces from internal motions present in the unit. The unit was measured on the edge of the outside panel of the unit while it was running, with no internal features visible, as seen in Figure 38.

> Location at which the measurement was made.

Figure 38. The air handler unit inside our laboratory exhibits motions that reveal information about the internal mechanical system. Data were collected on the air handler unit with the single element optical system. A 50 mm lens was chosen and the right edge of the system was imaged. The results yielded information about the hidden mechanical rotation, internal gearing, shaft rotations, and belt slippage. Credit: John Kielkopf

Data were collected on the system from the single element sensor system from approximately 1.8 meters away on the edge of the enclosure. A 50 mm lens used for collection resulted in a spot size 4.5 mm in diameter on the edge of the enclosure. The measurement location was confirmed by sending light

through the fiber optic cable from the photodiode end and projecting the spot onto the air handler unit. The lens was then positioned such that the edge of the air handler unit was located in the middle of the projected spot. The single element sensor system signal was processed through the desktop PCI NI-4472 A/D card. The data were collected at a rate of 1024 samples/second for 10 seconds and transformed into the frequency domain with LabVIEW generating the spectrum of mechanical vibrations in the air handler unit.
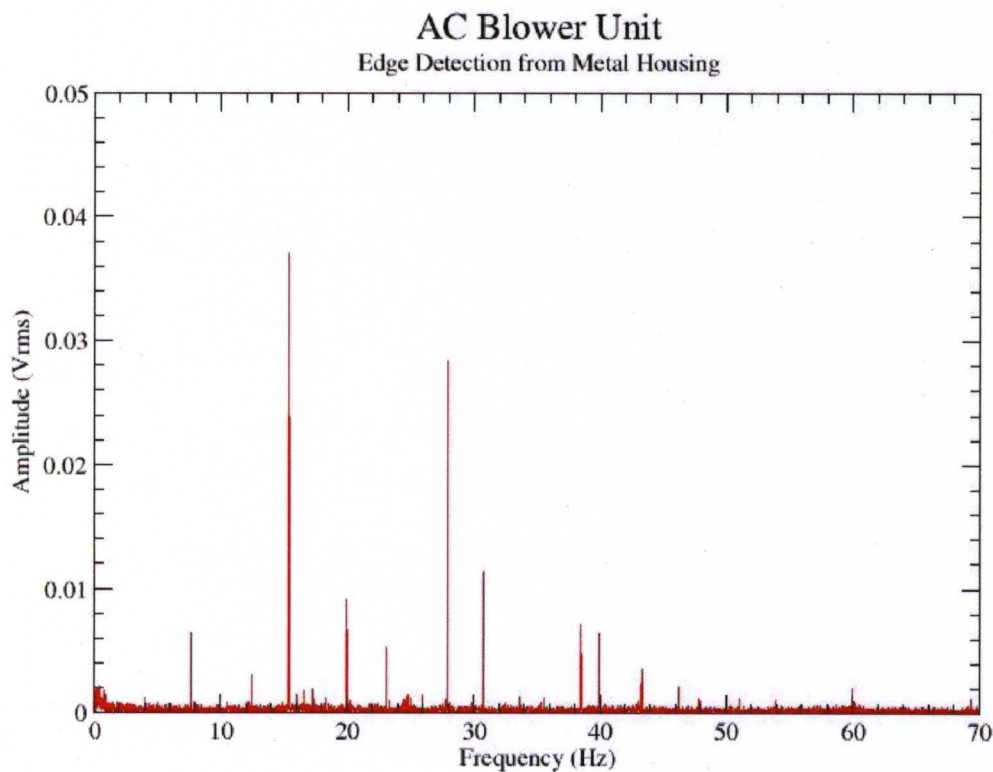
Figure 39. The air hander unit's spectrum reveal several distinct frequencies. Most notably are 15.5 Hz and 27.9 Hz peaks correspond to the rotation rates of the two internal pulleys.

Upon analyzing the internal mechanical system it is possible to relate the acquired frequency spectrum to the physical parameters of the system. Inside the air handler unit we can see that the unit is driven by a motor that is specified to operate at 1725 rpm which corresponds to a measured frequency of 27.9 Hz or 1674 rpm detected with a signal-to-noise ratio of 52.8. This shows that the unit is operating within 3% of the specification, most likely resulting from slippage of the unit, but still within tolerance.
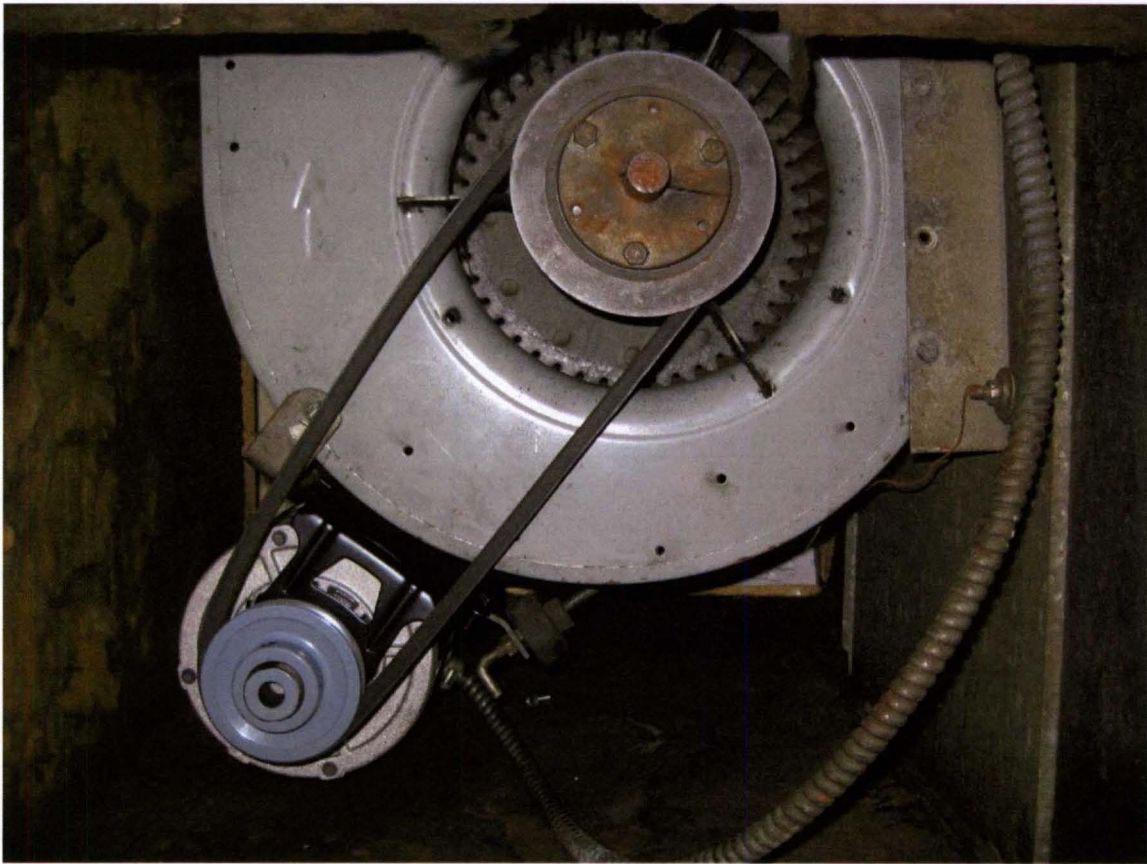
Figure 40. The internal view of the air handler unit shows the motor and belt driven system. The two wheels, which give rise to the two prominent peaks in the frequency spectrum are seen. Credit: John Kielkopf

101

The diameter of the motor's rotating pulley is 5.75 cm. This wheel then drives a second pulley, connected through a single belt. The second larger pulley has a diameter of 10.5 cm. This represents a factor of 1.8 in reduction in the speed of rotation. The resulting rotation rate of the second wheel would then be 930 rpm or 15.5 Hz, which is clearly visible in the spectrum depicted in Figure 39. This peak, which is the rotation of the blower itself, is prominent at a signal-to-noise ratio of 62.4. So from measuring the outside surface with our standoff system, it is possible to determine the rotational rates of the enclosed pulleys and whether or not the system is operating at the specified rates within tolerance.

7.2.2 Rooftop HVAC Units

Certain mechanical systems exhibit vibrations that may be useful to measure from a remote sensor. Situations may make remote sensing a more suitable and safer solution in the event, for example, that the machinery is inaccessible, in a dangerous location or in motion. An example of such machinery is an HVAC system where the machinery is relied upon to continually exchange the air in building to provide a safe environment. HVAC ventilation systems are often in a location not easily accessible but can easily be viewed from a remote location on the ground. A line of sight measurement can be made that measures the vibration resulting from the mechanical systems. These vibrations are indicators that the machinery is operating normally.

Location of single element sensor measurement.

Figure 41. The Chemistry Building on Belknap Campus has single HVAC stacks on top of the building. Measurements taken on one of these units reveal to operators the mechanical systems in the building. Such measurements can be easily made from the ground hundreds of meters from the building.

Figure 41 shows a typical HVAC system on top of the chemistry building on the University of Louisville Belknap campus. The unit would be a forced air system driven by a pump and ultimately responsible for air circulation in the building for evacuation of gases from fume hoods. Their proper functioning is a primary safety concern. Data were collected on the edge shown in Figure 41. The sides of the system offer us high contrast edges to make high signal-to-noise measurements. Data were sampled at a rate of 2048 sample/second for a length 17 seconds. A 500 mm Celestron ONYX 80ED refractor telescope was coupled

to a fiber that terminated at the InGaAs photodiode which was then amplified and digitized with a 24-bit NI-9234 National Instruments analog-to-digital converter. The data was collected at a distance of approximately 61.5 m resulting in a imaged diameter of 1.2 mm. The temporal data were processed with a Fourier Transform in Grace, an open source 2-D plotting tool, producing a frequency spectrum. [33]

Chemistry Building Ventilation Unit
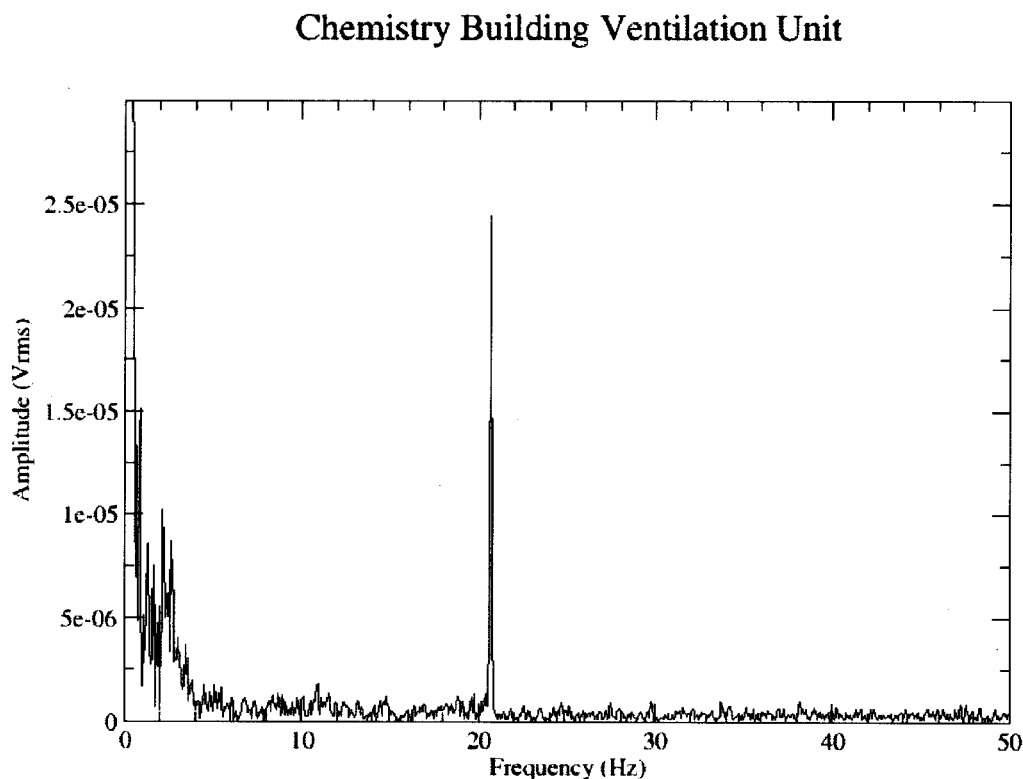
Figure 42. The frequency spectrum from the HVAC system on top of the Chemistry Building is shown for a single HVAC unit. A prominent feature at 20.68 Hz is clearly seen and corresponds to the rotational rate of the system's blower motor.

Figure 42 shows the resulting frequency spectrum from a 17s measurement of the chemistry building ventilation unit. A clear 20.68 Hz peak is seen from the vibration resulting from the mechanical pumping systems of the HVAC unit with a signal-to-noise ratio of 24.6. This corresponds to 1241 revolutions per minute. Most motors work efficiently without overheating if the rate is within the manufacturer's tolerance. This vibration frequency reveals the revolutions per minute at which the motor is operating, and shows that it is running normally. This assessment can be done remotely and quickly on a number of HVAC units all from the same test location. It is possible to survey several buildings from the same test point. This capability could be especially useful in the event that the units are inaccessible, in a location that is laborious to access, or during an emergency that restricts access.

Figure 43. Multiple HVAC unit on the Chemistry Building are visible from a single vantage point. All of these units could easily be measured from the same location on the ground.

7.3 Surface Motion

It is possible to analyze the surface motion of an object if sufficient modulation of the reflected surface light occurs. Usually small variations which result in edges, although not always straight, are present on surfaces and sufficient in contrast to make a measure of the modulation. It is important to note that if the edges on the surface are not straight across a pixel or otherwise unknown before the measurement, a calibrated quantitative measurement is not possible, although an un-calibrated quantitative detection is.

A surface that is generally well-behaved and characterized in terms of its surface motion while undergoing a vibration is a drumhead. When struck in a

particular way, a drum undergoes well known modes of vibration. Optical imaging has been performed on drum vibration with the use of projected fringe patterns recorded at different angles with high speed cameras to reconstruct the surface displacement. Two modes of vibration were seen in the drum head after it was struck in the center, (1,0) and (1,1). [34]
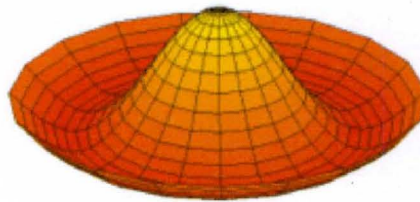
Figure 44. The second mode of vibration of a drum head ,(0,2), shows the pattern resulting from the drum resonating. A single point of maximum amplitude is characteristic of this mode of vibration. Credit: Oleg Alexandrov, Creative Commons.

It is possible to analyze this behavior with a drum head by measuring the modulation of the reflected surface light while it is resonating. Upon close inspection of a white powder coated drum head surface, it is possible to see the various levels of sharp contrast that can give rise to modulated light when the surface moves.
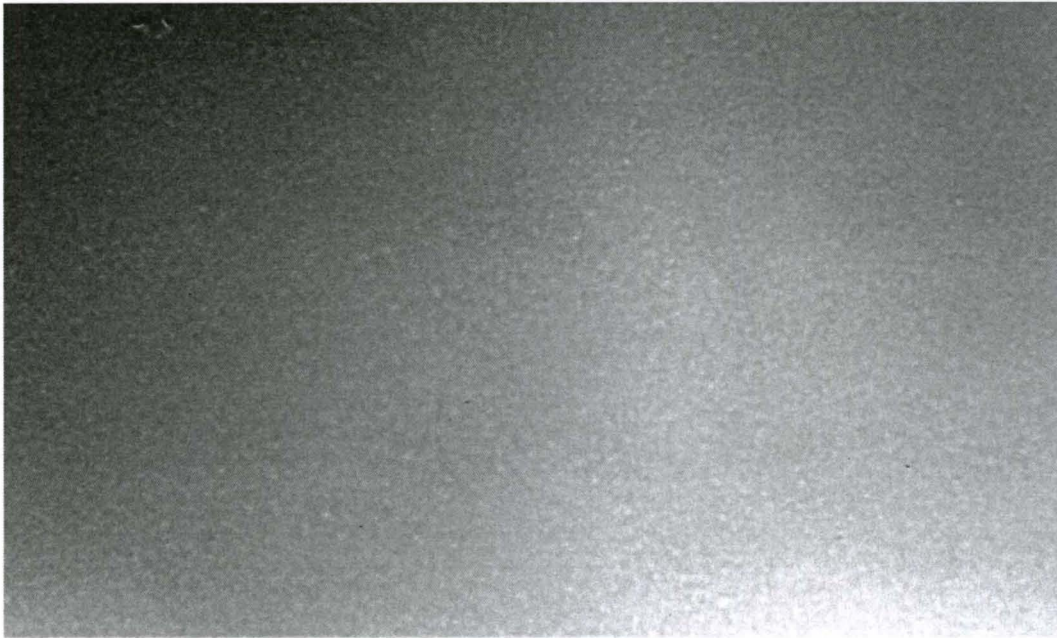
Figure 45. The surface of the drum head is made up of many edges of high contrast. These surface features allow for measurement of modulated light as the drum resonates.

Data were collected on a 18 inch Ludwig floor tom drum head with the Fastcam APX-RS at a frame rate of 500 fps after the opposite batter head was struck. The drum was also struck separately and measured with a Vibromet 500V Laser Doppler vibrometer for confirmation of the resonant frequency. Both measurements confirm the drum was resonating at 86 Hz. The laser Doppler vibrometer sampled at a rate of 1000 samples per second for 1 second for a total of 1000 samples.

Single Drum Strike

Laser Vibrometer Measurement
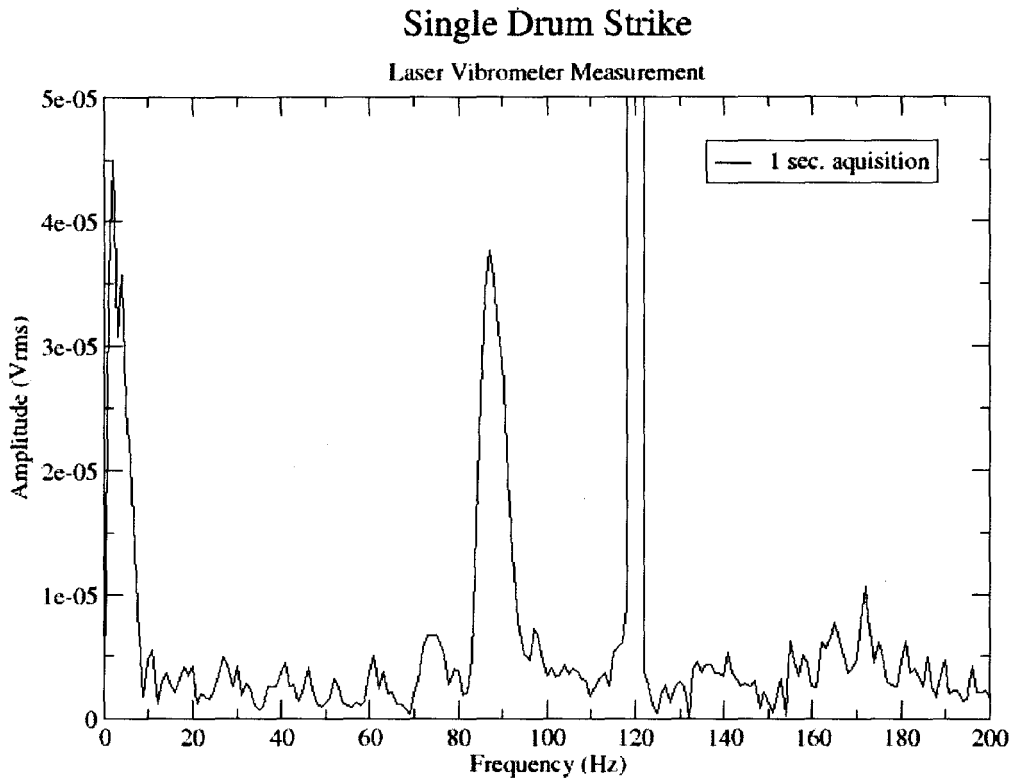
Figure 46. The spectrum of a drum strike shows a clear resonance

at 86 Hz as measured by a laser Doppler vibrometer. The 120 Hz

signal is from the AC coupled modulated room lights.

The data was transformed from the time domain to the frequency domain through

a Fast Fourier Transform. Figure 46 shows the spectrum taken with the laser

Doppler vibrometer. The 86 Hz resonance is clearly visible in the spectrum with

a signal-to-noise ratio of 8.35. Also seen is the 120 Hz peak from the room

lighting.

The video images where processed similarly for each pixel as a time

series. Three eigenimages were created representing the modulation of light at

those particular frequencies. Each frame represents the integrated values centered around the particular frequency of interest with 1 Hz bins.
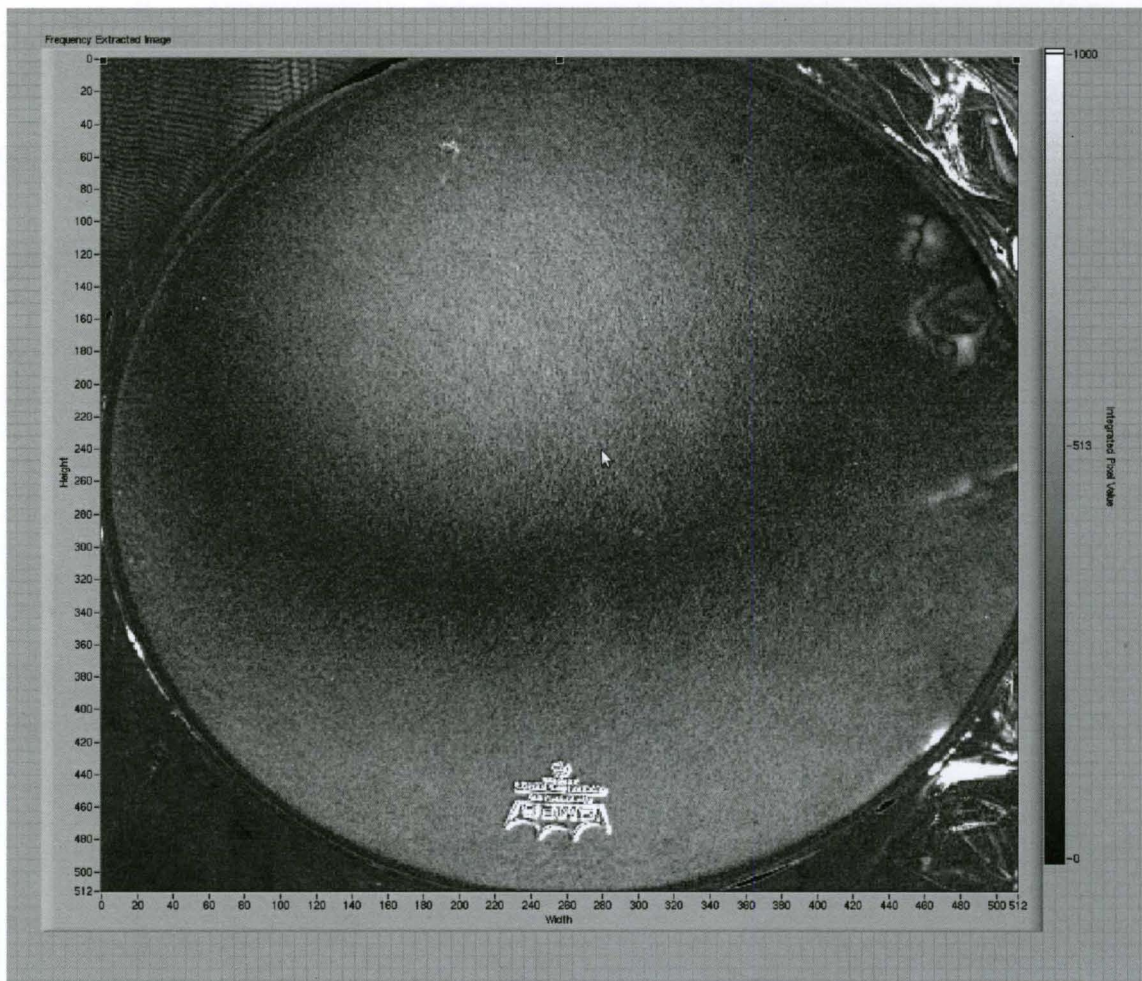
Figure 47. The 86 Hz response of a drum head clearly shows a vibrational mode. A single point of strong modulation is evident in the upper portion of the drum as shown here. That point is surrounded by an area of weaker modulation much like the second mode of vibration of a circular membrane.

Figure 47 shows the integrated values centered around the 86 Hz peak. The vibrational pattern of a drum head is clearly seen in the image. The brighter

110

portions of the image represent greater modulation. Portions showing strong modulation on the drum head have a signal-to-noise ration of 3.96.

To ensure the pattern seen in Figure 47 was not due to uneven illumination in the scene, two other images were created that would show the illumination pattern, a DC image representing the portion of the scene illuminated by a steady light source, and an image depicting the 120 Hz portion of the spectrum, representing illumination from AC room lighting. Neither image revealed a vibrational pattern on the drum head, indicating the pattern seen in the 86 Hz eigenimage is in fact due to the 86 Hz vibration.
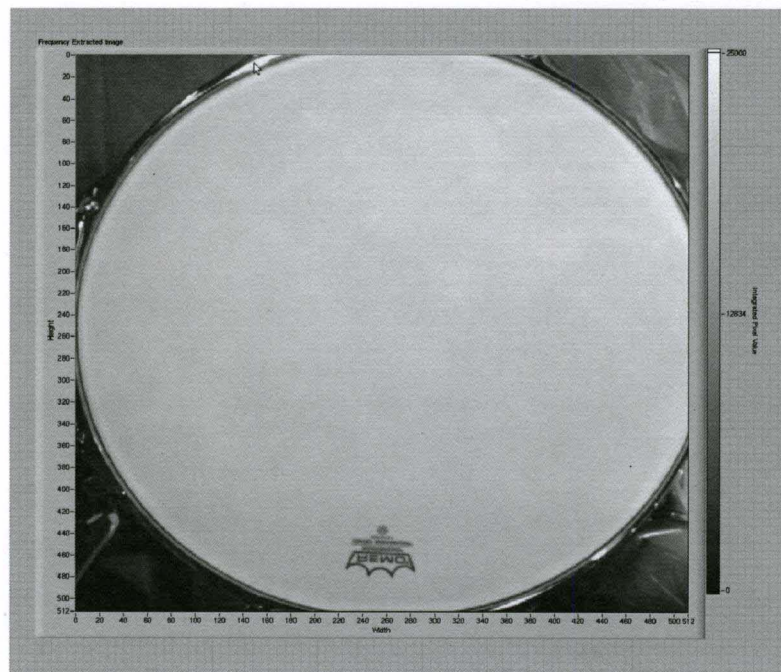
Figure 48. The DC portion of the drum illumination show no unevenness in the drum head. This indicates that vibrational pattern is not the result of uneven DC illumination.
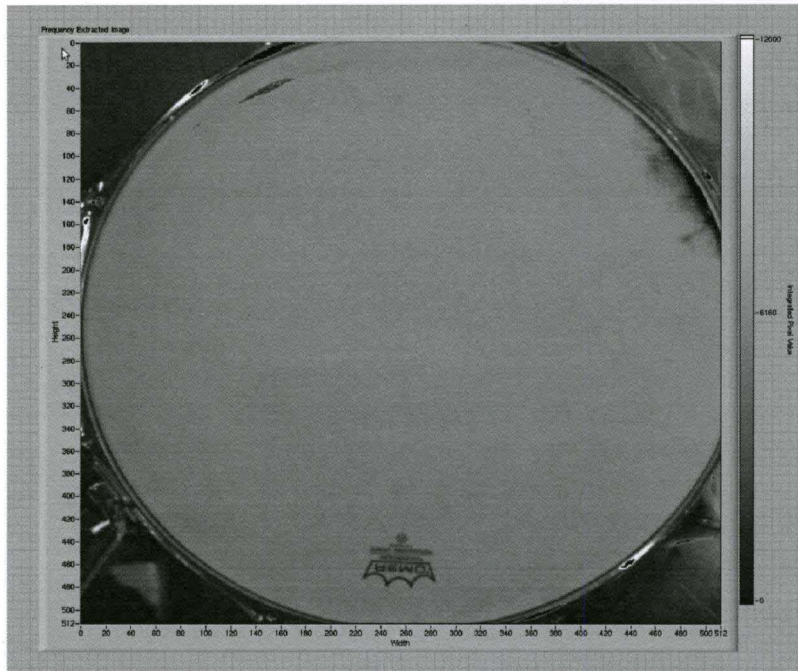
Figure 49. The 120 Hz portion of the drum head illumination shows no unevenness. This indicates that vibrational pattern is not the result of uneven incandescent lamp illumination.

7.4 Deflection of a Beam

With the ability to measure displacement it is possible to utilize the camera sensor to create a spatial measurement of a deflection across the field of view. A very simplistic approach would be to measure the deflection of a straight beam which allows for a consistent edge. To do this we affixed a common laboratory meter stick to clamps on either end. We then attached the center of the meter stick to a vibrating speaker as shown in Figure 50. We used the Pioneer TS-G1642R speaker along with the Applied Research and Technology SLA-1 linear power amplifier and BK Precision Model 3003 signal generator with a similar setup as described in Section 4.2 except the speaker was oriented vertically.

This created a situation where the end points were fixed and the center vibrated at a maximum, driving a fundamental mode with nodes at each end.
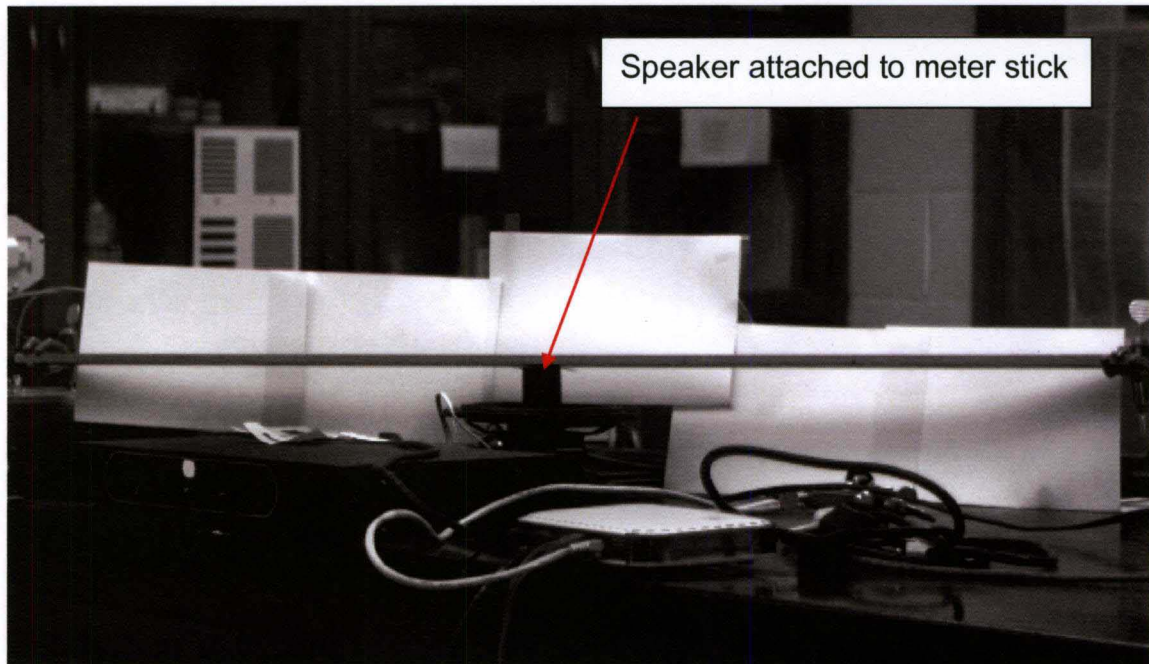
Figure 50. The deflection of a beam was measured using a vibrating meter stick clamped on both ends. The center of the beam was affixed to a plastic cylinder that was attached to a speaker cone. That speaker was vibrated at 20 Hz causing a deflection in the beam as it oscillated. The center of the stick showed maximum deflection as expected with the fundamental mode of vibration.

The background was composed of white paper to increase the contrast of the meter stick edge with the background and to improve the smallest detectable motion. The speaker was vibrated at 20 Hz while a sequence of images were recorded with the Pike F032B camera operating at a resolution of 640 x 480 and

107 frames per second. Data were acquired for 20 seconds resulting in 2140 frames. An eigenimage shown on the lower portion of Figure 51 was produced depicting the 20 Hz motion in the scene which correlates to the absolute displacement. The pixel values were determined by performing a Fast Fourier Transform on each series of pixels along the temporal axis and integrating the 20 Hz signal.
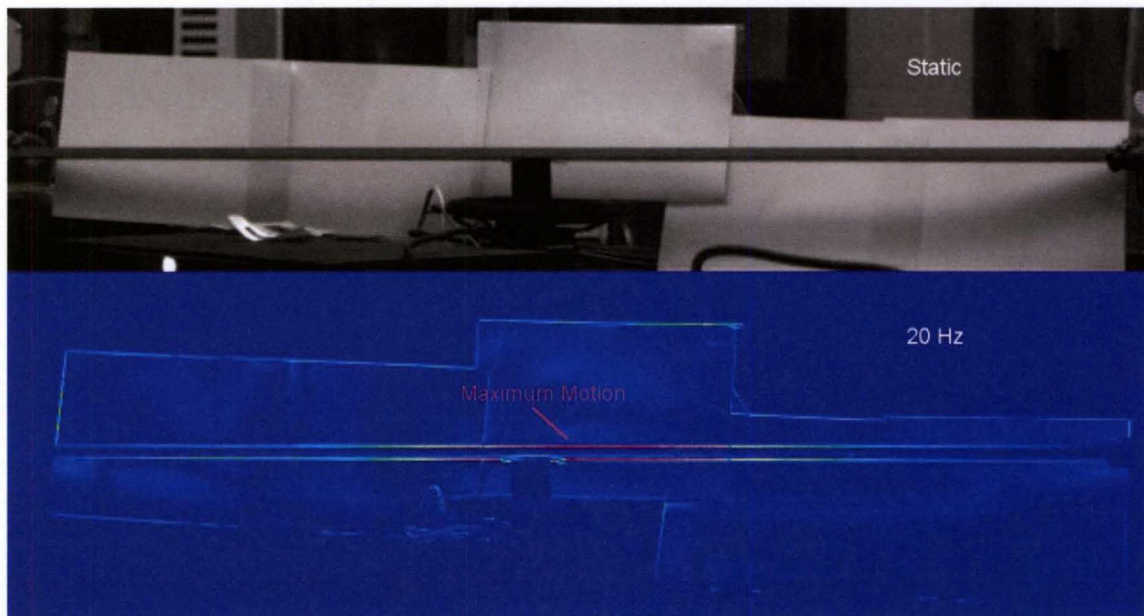
Figure 51. A comparison of the static and displaced beam shows the displacement to be greatest in the center of the element. The upper portion of the image is a single frame in a time series. The lower portion of the image is the eigenimage displaying the 20 Hz component of the image. Red denotes regions of strongest modulation while the blue denotes regions of the weakest modulation. The strong modulation in the center corresponds to the region with the greatest displacement.

Figure 51 shows both the static image and the eigenimage displaying the contribution of motion at 20 Hz driven by the speaker. It is clear that the most intense modulation in the signal is occurring in the center of the element which corresponds to the point of greatest deflection. The signal-to-noise ratio of an area of strong modulation in the eigenimage is 250.0. A signal cross section of the deflection was determined along the edge of the upper portion of the meter stick by summing 3 vertical pixels to account for some spatial smearing of the modulation across multiple pixels as well as the meter stick not being perfectly level. These values were then plotted as seen in Figure 52.

Displacement of a Horizontal Element
Vibrating at 20 Hz
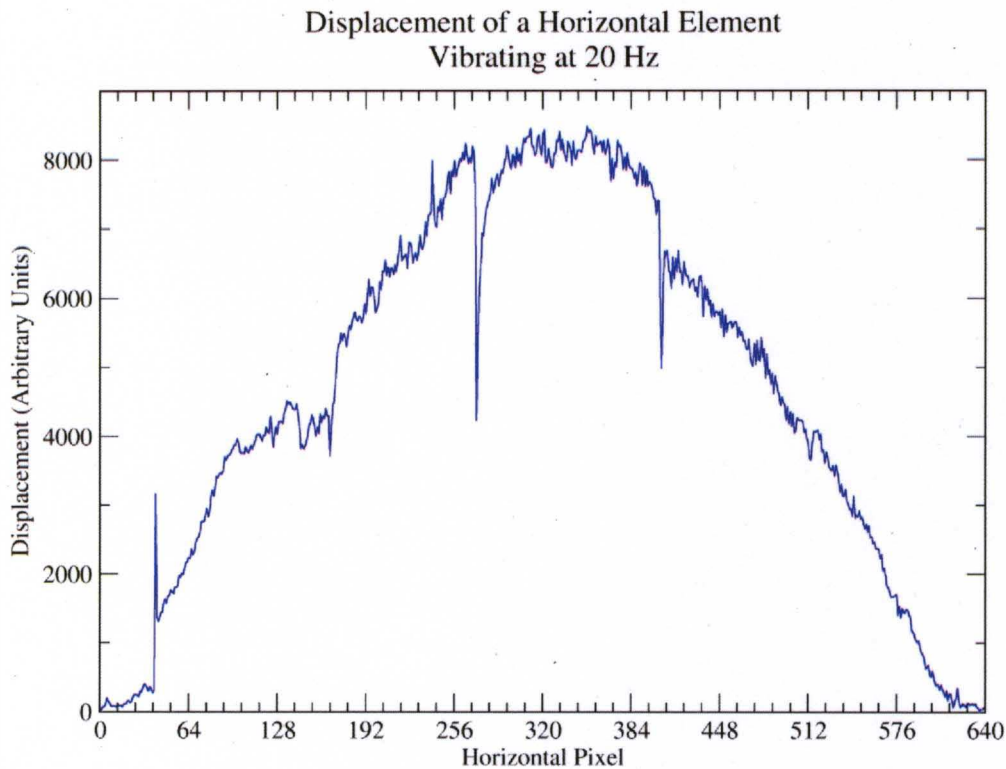
Figure 52. The plotted relative displacement of a beam shows the expected behavior of a vibrating element clamped on both ends. This data show the signal at the 20 Hz modulation along a horizontal slice of pixels on the edge of the meter stick. 3 vertical pixels where summed because the modulation smeared the 20 Hz signal across multiple pixels. Since the modulation is linearly dependent with displacement, the resulting plots is proportionate to the actual displacement of the meter stick.

The increase in displacement toward the center of the meter stick is clearly evident in the plot. It exhibits the behavior one would expect from a

vibrating element clamped at both ends in a fundamental mode. The dips in the plot are due to the fact that the contrast diminishes at points in the background where the papers are joined thus leading to a drop in the signal proportionate to the contrast level, as expected.

7.5 Biomedical

An area in which remote non-contact sensing is of interest is in the field of biomedical monitoring. Removing the tether of wires has the benefit of creating a less restrictive environment for the patient.

7.5.1 Blood Pulse

The sensitivity of the detector allows for applications where precision measurements are necessary. Although the technology lends itself to remote sensing of structures, the fact that it is a simple stand-off measurement of such high precision has merit for other applications as well. One such application is in the biomedical field and involves the sensing of the blood pulse.

Cardiovascular disease is the single largest cause of death in the United States, responsible for 1 in every 2.9 deaths in 2007. [35] Coronary heart disease, which is caused by arthrosclerosis claimed 406,531 deaths in 2007. The cost of cardiovascular disease alone to the US in 2007 was $286.6 billion. [35] Research shows that the prevalence of cardiovascular disease is rising at a rate of 1% to 2% per year as the age of the population within the US is expected to rise. [36] A substantial need for increased cardiovascular care is expected in the next four decades. [36] Arteries at childhood are compliant and

become progressively less compliant or stiffer with age. New and affordable technologies that quantitatively measure arterial stiffness through simple measurements, and possibly track arterial stiffness through aging, would be clinically relevant. [37-39] Studies by Millasseau et al have shown that the blood pulse waveform can offer valuable information about cardiovascular stiffness including an accurate assessment of cardiovascular age. [40] In addition, models put forward by Pauca, for example, have shown that the radial pulse profile can be used to assess the aortic blood pressure, leading to the possibility of a low cost, non-contact, blood pressure monitor. [41]

There are multiple spots on the body where an artery is close enough to the surface of the skin that the displacement of the artery wall during pulsation from blood flow is evident. A few locations are the wrist, with the brachial artery, the neck, with the carotid artery, and the thigh with the femoral artery. [42] This results in multiple locations in which the arterial blood pulse can be sensed outside the body. The Pulse Wave Velocity (PWV) is a widely accepted method of determining arterial stiffness that is determined by measuring the transit time of a blood pulse wave from one location in the body to another. [43-47] With multiple locations being present that allow for a blood pulse measurement to be made optically from the surface of the skin, the proposed detection would be capable of measuring the PWV. In addition, if the orientation of the wrist or neck is oriented correctly the pulsation can be seen as a moving edge from the detector's point of view. It is in this fashion that we are able to make a

measurement of the blood pulse that fits within the model outlined for our detection method.

The wrist is the easiest of the various locations to access so we will primarily focus on this measurement. The setup involves using a similar system to other laboratory tests, a Tamron 16mm lens coupled to the fiber that feeds the photodiode sensor sensor system. The signal is then processed with a 24-bit NI 9234 analog-to-digital converter read by LabVIEW. The lens is placed relatively close to the wrist, ~10cm to increase sensitivity to the motion resulting in an image sample area with a diameter of 0.63 mm.

Figure 53. A single element sensor is placed above the wrist to make the blood pulse measurement. The wrist is then viewed edge on so the profile of the compliance of the skin to the blood pulse can be determined.

Figure 54. Illumination of the wrist for measurement of the blood pulse is done with a small DC powered incandescent lamp flashlight. The wrist is viewed from an angle that enables a measurement of the blood pulse as an edge moving across the scene seen by the detector. Credit: John Kielkopf

The wrist is illuminated by a small DC source, in this case a flashlight, to increase the signal-to-noise ratio. The sensor is positioned above the wrist in such a way that the profile of the wrist, where the pulsation is evident, is seen by the detector. An appropriate background, in this case a gray foam, is positioned below the wrist to increase the contrast of the edge formed by the profile of the wrist against the background.

Data acquired at a rate of 1000 samples/second give a temporal resolution of a millisecond. On each pulse, the artery expands, pushing outwards on the skin which is then seen by the detector to move across the field. Because we look at an edge of the wrist much in the same way we do a vibration, the model correlating the measurement to a physical displacement,

previously put forward, is still valid to describe this motion. The detector sees a greater illumination as the artery pulses outwards, forcing more of the illuminated wrist into the field of view and likewise sees a lesser illumination when the artery contracts allowing the detector to see more of the background. The signal is modulated in proportion to the displacement of the skin. Figure 55 show a measurement taken from the wrist.
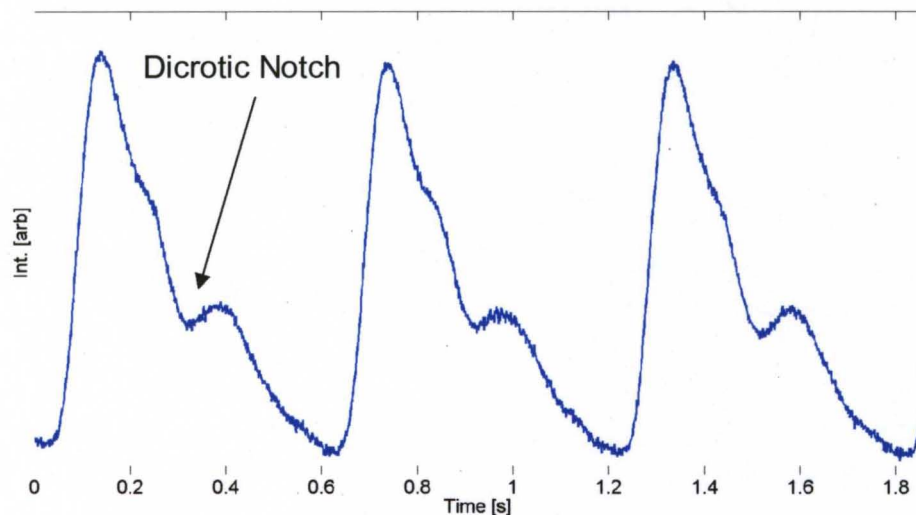
Figure 55. The blood bulse profile seen in edge mode. Shown is the time series of the blood pulse as measured by the single element sensor. The periodic blood pulse profile is clearly measured with features such as the dicrotic notch well defined. Rendering: Adam Willitsford.

The blood pulse profile is clearly seen in Figure 55. Not only is the repetitive process detected, but finer structure, such as the dichotic notch, can be seen within the pulse leading to the possibility that subtle changes in the arterial

121

wall can be identified and measured. The dicrotic notch is the result of the closure of the aortic valves and is used to indicate the end of the systole ejection period. [48] Studies have shown that the shape of the dicrotic notch can be an indicator of cardiovascular health. [49,50] Lax *et al.* referred to the dicrotic notch as the single most important diagnostic feature of the blood pulse waveform. [51] Figure 56 shows a plot taken from Laurent *et. al* showing the conventionally measured blood pulse measurement exhibiting similar shape. Figure 57 also taken from Laurent et. al shows how the measurement of the blood pulse is related to useful information such as blood pressure and distensibility. [52]
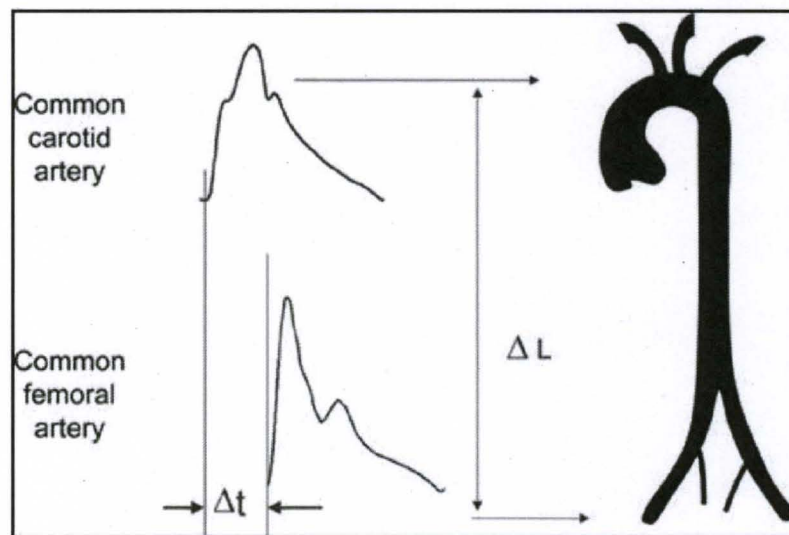
Figure 56. Graphic of the blood pulse transit time from Laurent et al. illustrates the concept of the blood pulse transit time. [52] The waveform depicted closely resembles the waveform measured with the single element photodiode system. Credit: Laurent et al.
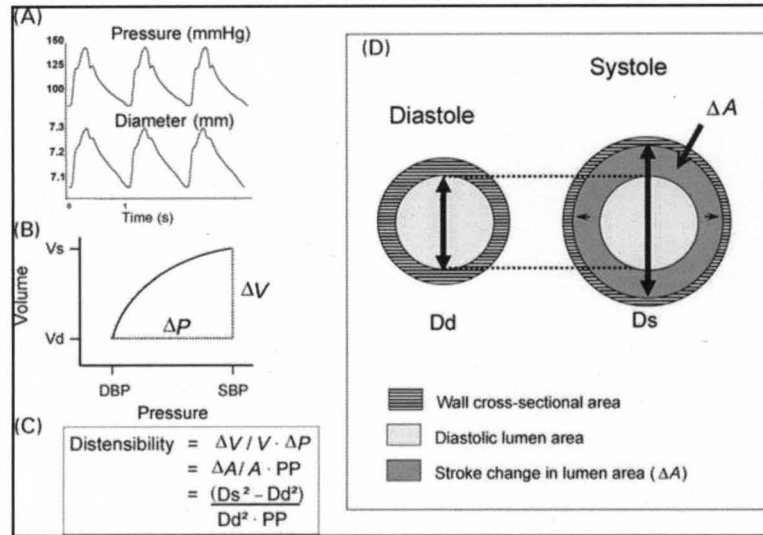
Figure 57. Graphic of the distensibility from Laurent et al. illustrates the concept of distensibility and its relation to vessel diameter. [52] Credit: Laurent et al.

This methodology offers a high level of precision for the measurement of the arterial wall displacement. As previously detailed the smallest detectable change in the displacement is given by the following equation, assuming the contrast across the edge is maximized, which we are free to do in a controlled environment,

$$\delta x = \frac{rw}{f} N^{-\frac{1}{2}}$$

where again,

> *f* is the focal length of the lens
>
> *r* is the distance to the target

w is the pixel width

N is the number of photons

Here we effectively have a 20 bit sensor so N $=2^{20}$, while f =1.6 cm, r = 10 cm and w = 100 x 10^{-4} cm from the 100 micron fiber. Using these values we have,

$$\delta x = \frac{0.1m \times 100 \times 10^{-6}m}{.016m \times \sqrt{2^{20}}}$$

$$\delta x = .61 \text{ μm}$$

We see that the system has a very precise spatial resolution with the above figures, ones that would be typical under which this measurement is made.

The blood pressure was tracked with a Vibromet 500V Laser Vibrometer for comparison of displacement and confirmation that the shape that we were measuring was in fact the blood pulse. The laser vibrometer was focused on the exact location where the blood pulse was causing motion on the wrist. Figure 58 shows a superimposed plot of the optical data as well as the vibrometer data. These data were not acquired simultaneously and were scaled for comparison.

Blood Pulse Profile
Optical and Velocimeter Comparison (not aquired simultaneously)
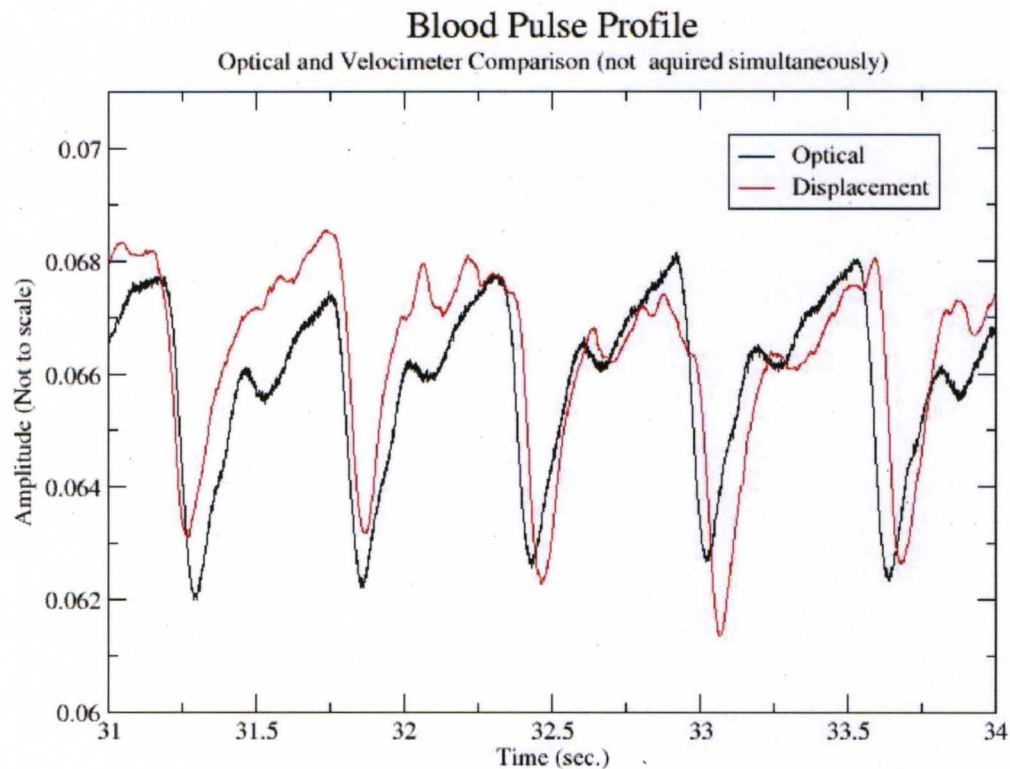
Figure 58. Optical and vibrometer comparison of blood pulse as measured by the single element photodiode system and the laser Doppler vibrometer show a correlation in the rise and fall of the blood pulse waveform. The laser Dopper vibrometer was used to confirm the measurement of the blood pulse waveform on the wrist.

We see that the overall shape and pulsation tracks well with the vibrometer measurement, indicating that we are in fact measuring the pulsation on the surface of the skin from the blood pulse. Notice how much more cleanly the dicrotic notch is distinguished in the optical measurement.

7.5.2 Respiration

Applications such as structural monitoring and analysis of mechanical systems are not the only areas in which sensing a vibration remotely with a stand-off device is useful. Essentially the concept will be useful for any application in which a range of frequencies defines a phenomenon. Human respiration is another biomedical area where the activity produces a pronounced and well defined motion. Respiration is detectable by looking for the motion of the chest cavity related to the inhalation and exhalation of the air. Respiration rates that are typical are in the range of 6 to 31 breathes per minute. [53]

Respiration rates and patterns have been shown to be an early indicator of diseases, common in both children and adults. [54] Conventional respiration monitoring devices usually require contact, such as affixing a sensor directly to the chest cavity to measure physical motion. This methodology is often problematic for the long term monitoring of patients, for example in overnight sleep apnea studies. [55,56] Additionally, these contact sensors often require wires, restricting patient movement and can become detached. The same is true for monitors that focus on infants. [56,57] This underscores the relevance of a non-contact respiration monitor.

Figure 59. A common button down shirt for measurement of respiration shows a large number of edges to detect for measurement. The high contrast edges make a high quality measurement possible.

Figure 60. A common white t-shirt for measurement of respiration shows a number of edges from the stitching from which a modulation can be detected.

127

Because the chest wall has such a pronounced motion it makes a very good target in analyzing the frequency response of the subject's respiration. As discussed before, contrasting edges are needed to make a measurement with good signal-to-noise value. As it turns out, edges are abundant when it comes to measuring a subject's chest cavity motion. Even a uniform white t-shirt contains numerous features such as the stitching within the shirt is easily resolved with a modest resolution camera. The same would be true of a subject that is sleeping underneath a sheet or cover, the material itself has stitching that produce a large number of edges. The various folds that are bound to be present in the material become edges that are useful in making the measurement. Figure 59 and Figure 60 show two shirts on the subject being measured. Figure 59 shows a typical dress shirt. It is clear that the image contains a large number of edges that have high contrast values. On close inspection even the stitching produces edges on the pixel level, all that is needed to provide a signal. Figure 60 shows a common white t-shirt, in which the stitching is clearly visible. The high dynamic range of the sensor makes it possible for such fine detail to produce a respiration signal from such an image.

Data were collected with an Allied Vision Pike FO32B camera that has a 14-bit depth. The camera operated at 30 frames per second and continuously acquired images. A sequence of 20 seconds of data, 600 frames, were processed into the frequency spectrum. Each pixel in the temporal series was transformed to the frequency domain resulting in a 300 frame frequency data

cube ranging from 0-15 Hz. This enabled us to study frequencies as high as 15 Hz with a frequency resolution of .05 Hz.

Whole Frame Respiration Detection
Whole frame average of all edges in scene
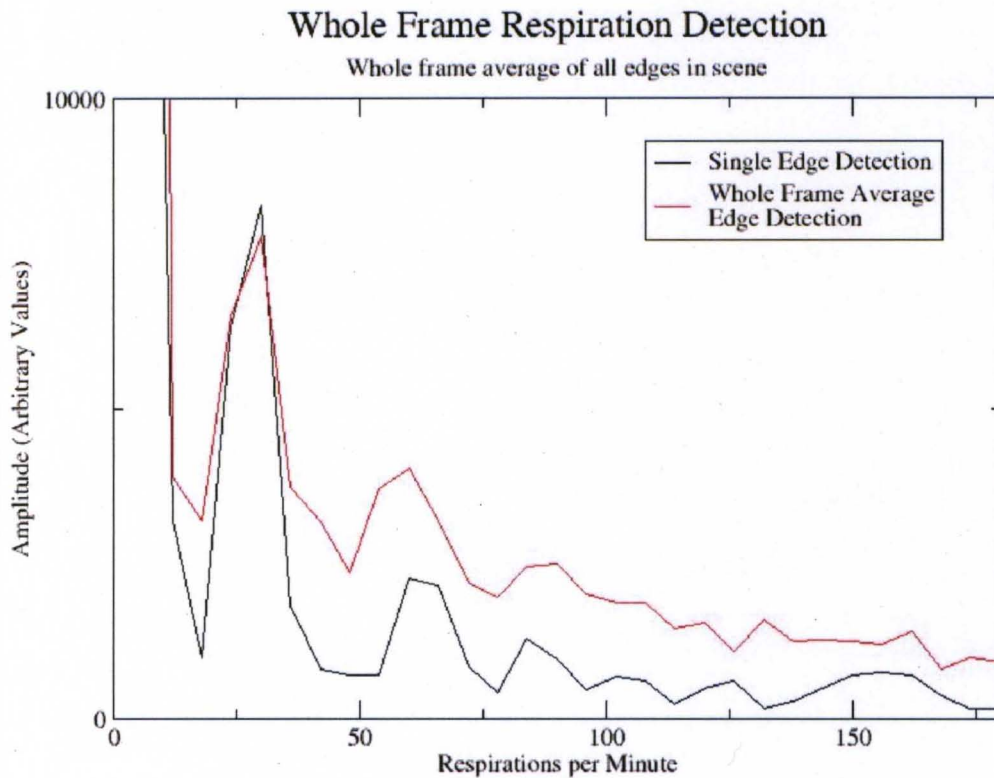
Figure 61. The whole frame average of respiration can be used to determine respiration for a subject without spatial resolution, regardless of the fact that the entire frame is averaged. This shows that the respiration rate of an individual can be determined when their location in the scene is unknown.

Figure 61 shows the frequencies evident from the individual in the white shirt. A clear peak at 26 respirations per minute is clearly evident, showing that the respiration has been detected. The red line, with a signal-to-noise ratio of 32.68,

129

shows the frequency response from the average values of the frequencies over the entire frame. What this shows is that an exact or best point of measurement is not needed to determine the respiration but instead an entire area can be averaged and the respiration rate can still be determined. The whole frame measurement shows a major feature at 20 respirations per minute, with a signal to-noise-ratio of nearly 50. A single point measurement has an even higher signal, and this show that there are individual points that are better than others for determining the respiration rate. A cross-correlation technique can be implemented so that multiple pixels can be used to make the measurement but only those sharing common frequencies will be used. In this manner a larger area can be searched for the respiration while still maintaining a high signal-to-noise ratio.

Respiration Rate
Edge Detection from Uniform White Shirt
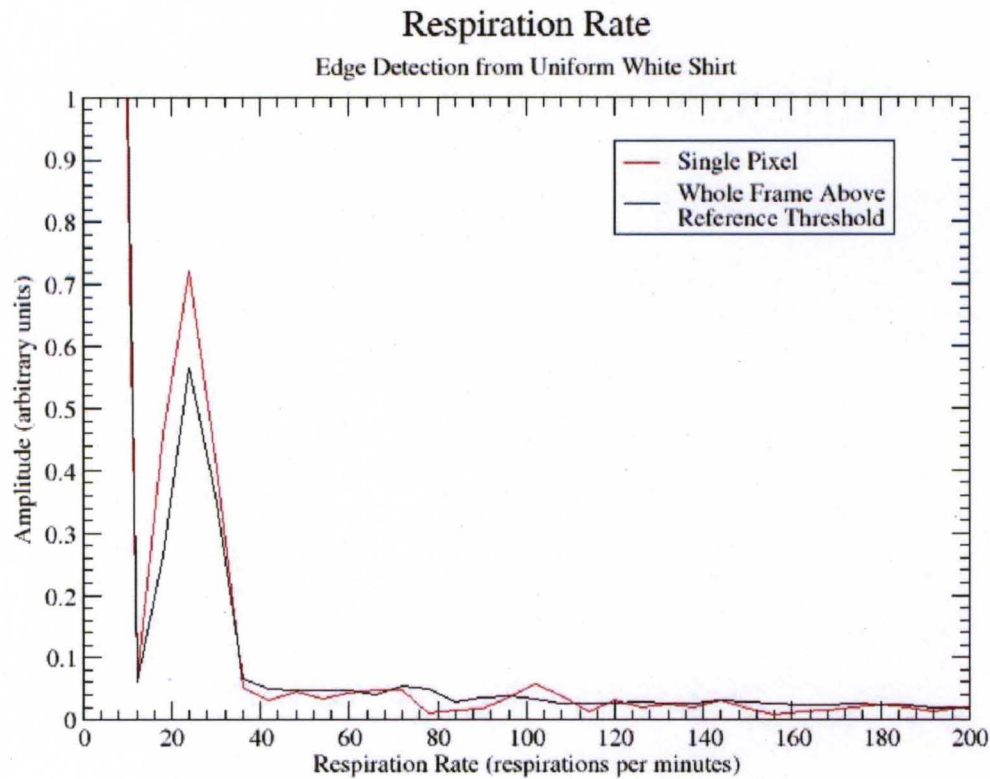
Figure 62. Comparison of single pixel to whole frame respiration levels shows the relative signal values of a single pixel of high contrast to an averaging of the entire frame. As expected, the single pixels offers a higher signal values of respiration than the average pixel value over the entire image. In addition, it offers lower noise.

Figure 62 shows a comparison of a single pixel that shows a strong feature at the respiration frequency and a whole frame average of only pixels above a threshold amplitude of 4000 in pixel value. This allows only those pixels showing an amplitude of a specific frequency above a fixed threshold to be

included in the average. The threshold amplitude was set such that only pixels with frequencies detected at a suitable level above the noise were used in the average. This method rejects the pixels that do not exhibit the signals related to the respiration rate, effectively rejecting a large source of noise and the irrelevant part of the frame. Figure 62 shows the single pixel amplitude in red and the average pixels above a reference threshold in black. The single pixel measurement has a signal-to-noise ratio of 62.0 while the threshold sampling has a signal-to-noise ratio of 48.5. We can see that the noise level has been greatly diminished from the previous plot in Figure 61 where all the pixels where included in the average.

To test the ability to resolve other frequencies, the individual being measured breathed at different rates, categorized as regular, accelerated, slowed, and no breathing. The results of those measurements are shown in Figure 63.

Respiration Rates
Detection from Chest Respiration
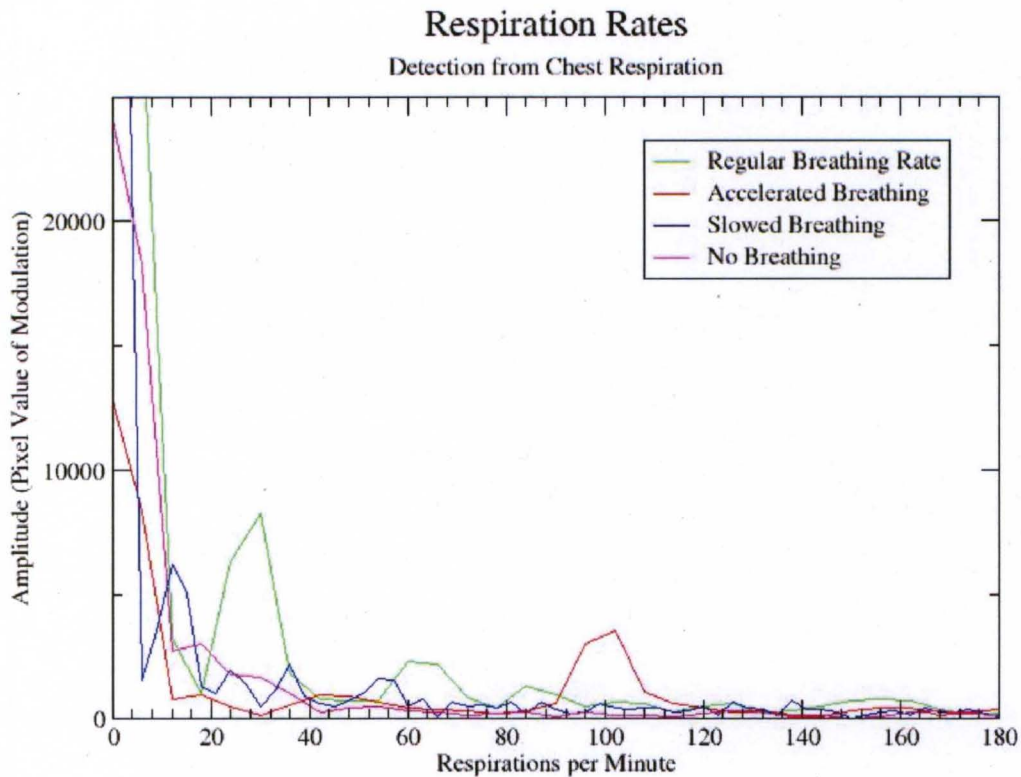
Figure 63. Detection of controlled respiration rates has been demonstrated. Rates at various frequencies are measured from the same individual while purposefully controlling breathing. Even the lack of respiration is detected by the absence of a discernable peak.

The green plot in Figure 63 shows the respiration rate for normal breathing, that is the individual continued to breathe at a rate in which they were comfortable. The normal breathing plot at 30 breaths per minute has a signal-to-noise ratio of 138.3. The individual then intentionally breathed faster, and the measurement at 100 breaths per minute is shown in red. We can see that the faster measured

respiration has a signal-to-noise ratio of 13.3. The individual then slowed their breathing and this is shown in the blue plot. The respiration rate is determined to be ~10 per minute with a signal-to-noise ratio of 204.0. Finally the individual stopped breathing for the entire 20 s acquisition, with the result shown in pink. It is clear that no peak is evident in the plot and that a lack of breathing has been detected.

Individual slices, or eigenimages, in the frequency data cube can be analyzed for specific detected frequencies to see where the respiration is taking place and what pixels exhibit the strongest signal.
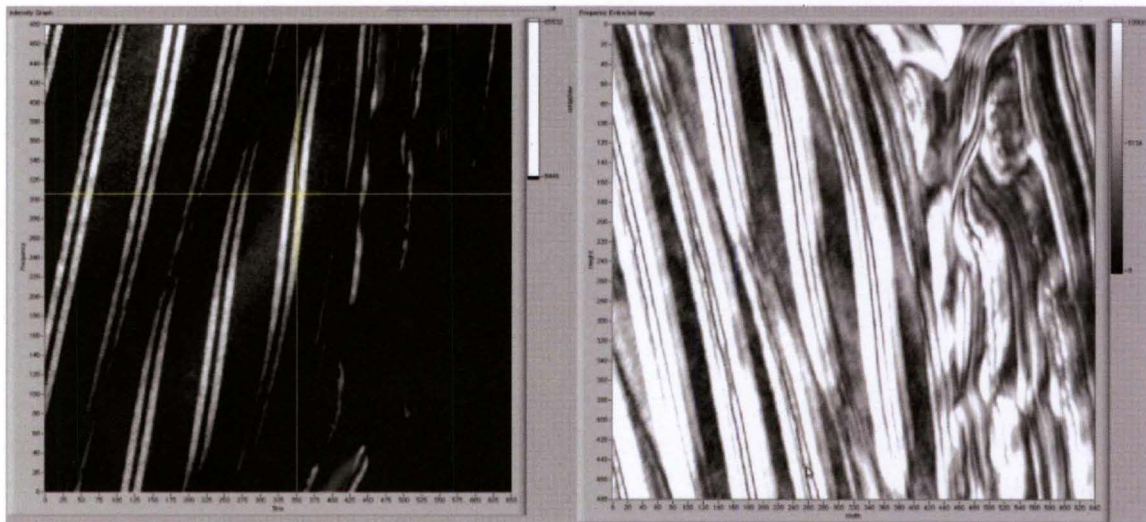
Figure 64. An eigenimage of respiration levels for a common shirt is shown. Comparison of a time domain image and frequency domain image for a button-down shirt shows a single frame from the time series on the left and only the amplitude of the frequency corresponding to the respiration rate on the right.

Figure 64 shows the result of an analysis of only the frequencies associated with the respiration. Data were taken on the chest of an individual wearing a striped shirt. The left side of the image shows a single frame from the time series. The right side show the signal at each pixel for a specific frequency, in this case the peak respiration rate. We see that the stripes in the fabric produce high contrast edges that help to increase the modulated signal. The highlighted pixel in yellow on the left side of the image shows how one pixel, in this case one that lies on a high contrast edge, can be picked out to look at the frequency response. This is done because the amplitude of the signal is expected to be higher for that pixel. The right side shows how strongly the respiration rate is detected. The entire chest is moving and because the field of view is small the respiration rate is seen across the whole image.
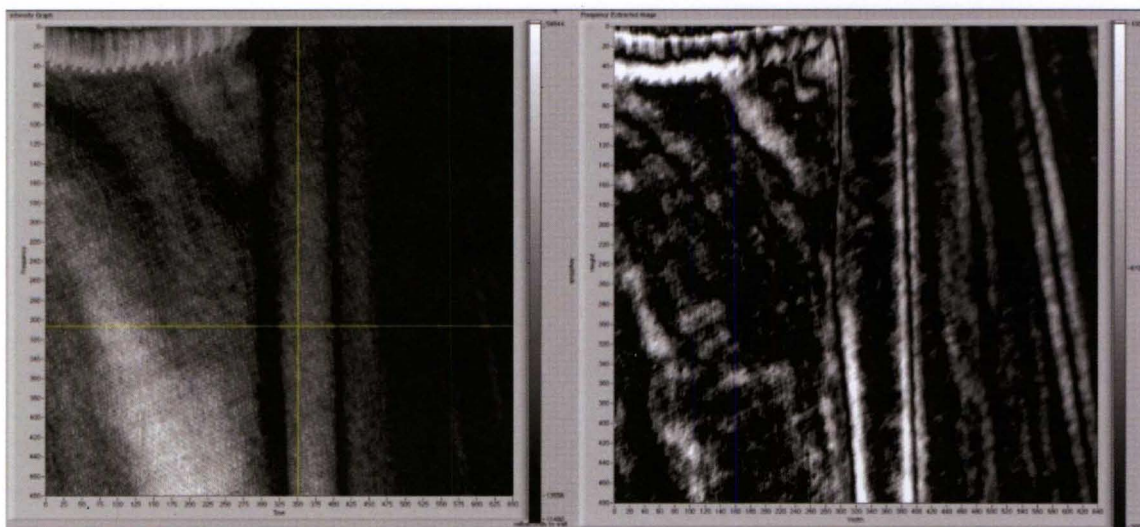
Figure 65. Eigenimage of respiration levels for a common T-shirt is shown. Comparison of a time domain image and frequency domain images for a T-shirt show a single frame from the time series on the left and only the amplitude of the frequency corresponding to the respiration rate on the right.

Figure 65 has that same types of images, except here we show the result from looking at a white T-shirt. Again, the left image is a single frame from the original time series while the right is the frequency-processed image showing only the signals association with the respiration rate. Even with the white T-shirt it is clear that there are many pixels that show a high amplitude in modulation for the respiration rate, well above the background which is the black pixels. It is interesting to have the ability to detect modulation from features such as stitching and folds in the fabric. Both are able to give us information about the respiration rate as they have a high enough contrasting edge to produce suitable modulation in the signal at the respiration rate. This shows us that even from what appears

to be a completely uniform surface, there are enough edges to produce a signal strong enough for the detection of the respiration rate. Even of more importance is the fact that these surfaces are typical in situations that would be present when you would be monitoring a respiration rate, at home or in a clinical setting i.e. when the subject is resting with their clothes on or sleeping in a bed covered by a sheet.

It could be useful if respiration rates from several individuals could be detected in the same field simultaneously and their individual rates spatially resolved and correlated with the subjects. To do this eigenimages of separate frequencies can be created to spatially isolate respiration rates and correlate them to individual subjects. Data were collected on two individuals in the same scene with the Pike FO32B camera at a frame rate of 107 frames per second for 30 seconds. This gives the ability to resolve .033 Hz, or 2 breathes per minute, well below the typical human rate of 6 to 31 breathes per minute. [53] A Fast Fourier Transform was performed on each pixel in a time series creating a frequency data cube. The cube was analyzed for individual eigenimages showing high levels of modulation with correlated spatial properties.
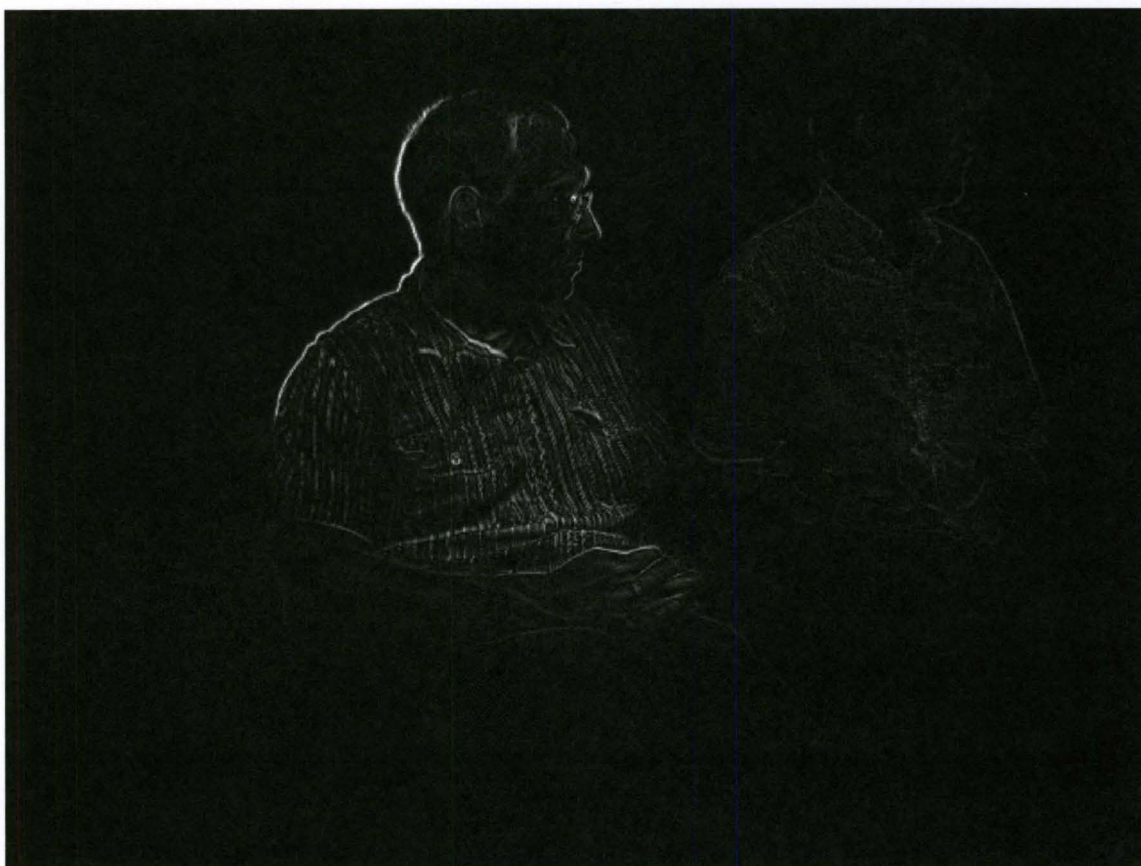
Figure 66. A respiration rate of 8 breathes per minute is measured with the high speed camera and its eigenimage is shown. Clearly the respiration rate of the individual to the left is selected. The fact that the respiration rates can be measured spatially allows this particular rate to be attributed to the one particular individual in the scene breathing at that rate. A much smaller amplitude is still shown for the individual to the right. That the other subject is detected is due to the fact that the respiration rates are close in the frequency domain and the wings of the peaks reach the frequency bin of the other rate.

Two distinct eigenimages have high modulation levels, at 8 and 12 breathes per minute. In addition, they show strong spatial correlation to the location of the separate individuals in the frame. Figure 66 shows an eigenimage at a respiration rate of 8 breaths per minute. The strong spatial correlation to the individual on the right suggests this is the rate at which that person is breathing. In addition, the spatial location of the modulation is located at the chest cavity, where the strongest motions associated with respiration would occur. A single point of strong modulation in that area shows a signal-to-noise ratio of 696.2.
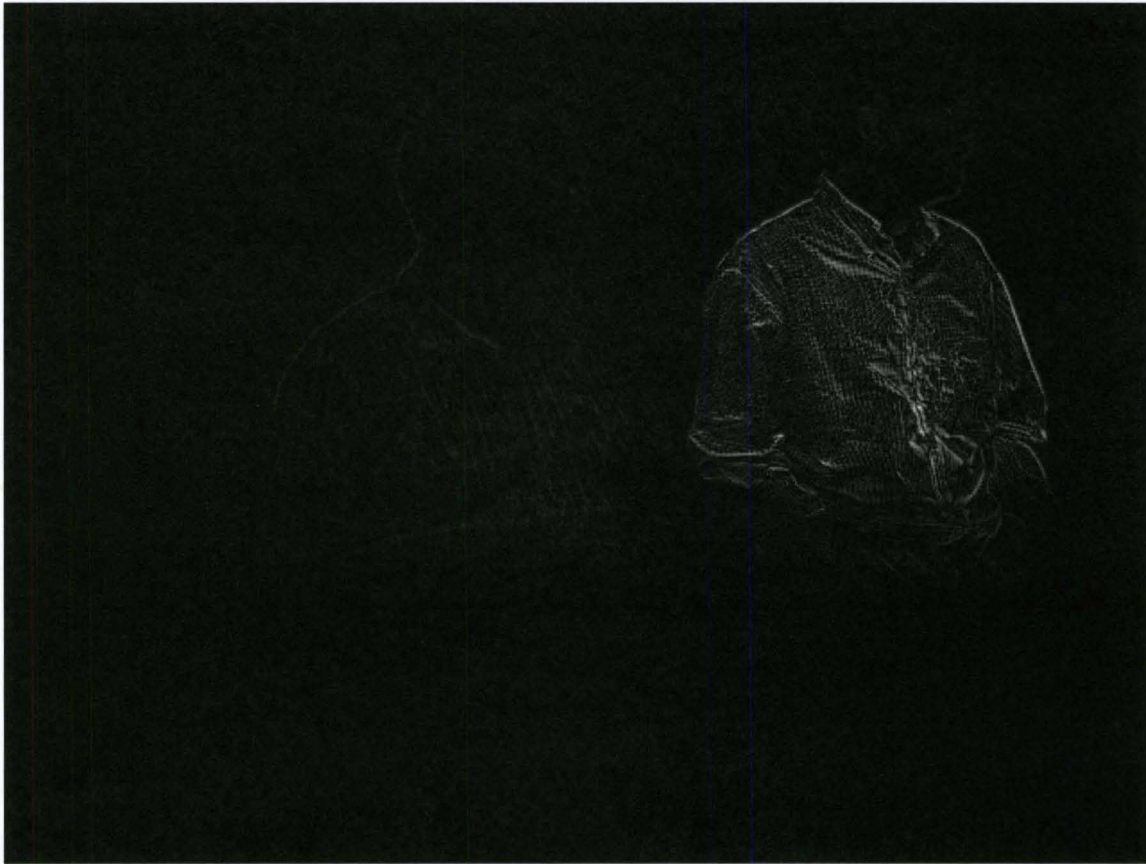
Figure 67. A respiration rate of 12 breathes per minute is measured with the high speed camera and its eigenimage is shown. Clearly the respiration rate of the individual to the right is the subject breathing at that rate. The fact that the respiration rates can be isolated spatially allows particular rates to be attributed to particular individuals in the scene. In this case, a much smaller amplitude is still shown for the individual to the left who is breathing at a different rate.

Figure 67 shows the eigenimage depicting a respiration rate of 0.2 Hz or 12 breathes per minute. This image was created using the same data set as Figure 66. A point of strong modulation in this image shows a signal-to-noise ratio of

929.7.   This demonstrates how multiple respiration rates can be isolated simultaneously in a field of view.

These fields can be used to construct the respiration waveform for the isolated subjects.   A threshold for the modulation level can be determined indicating a positive determination of a respiration occurring at that rate and those pixels can then be summed in signal over each frame in the time series. Essentially we look at the eigenimage and determine where the strongest modulation is occurring within a frequency range.   Then we sum those pixels in each frame and plot that sum over time or frame number.   The result is a single virtual pixel tracking the chest motion over time by recording the optical signal.

Respiration Waveform
Integrated Values of Entire Chest
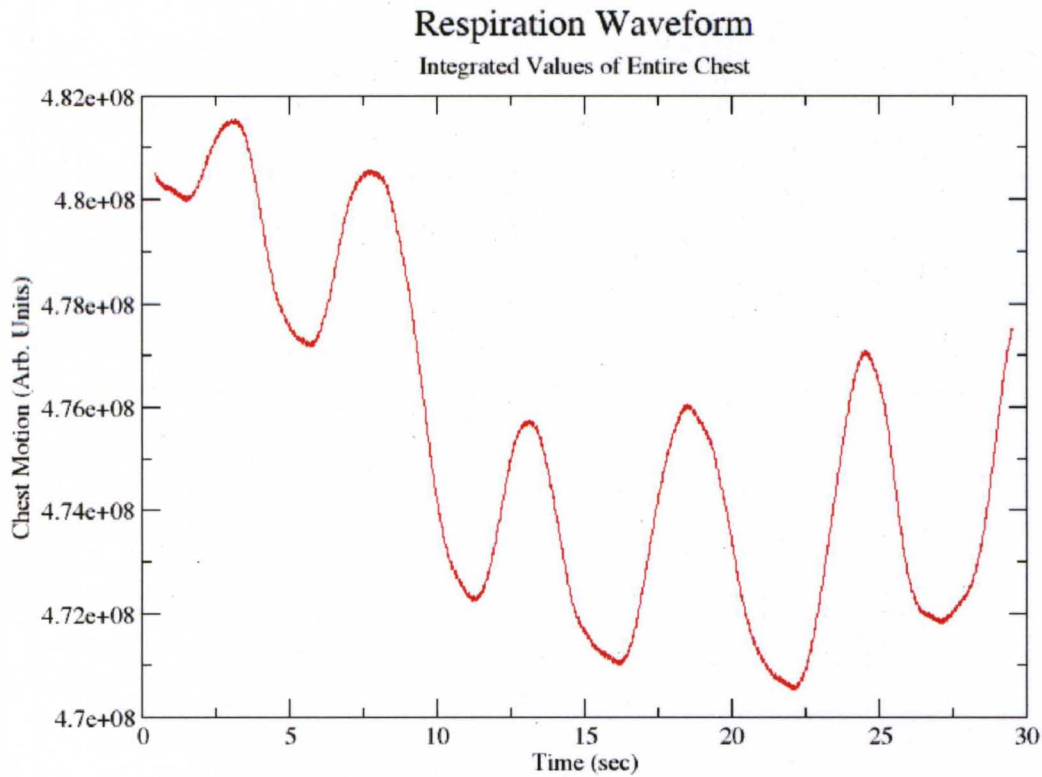
Figure 68. The respiration waveform can be extracted from the individual time series frames. To extract the waveform, the pixels with the highest modulation in the eigenimage are determined by setting a minimum threshold. Those pixels are then summed, and the results for successive temporal images are plotted against time.

Figure 68 shows the result of summing pixels of high modulation in the time domain. A threshold pixel value of 500 in the frequency domain was chosen to be well above the background level which is measured to have a typical average pixel value of 6.45. All pixels showing a signal of 500 or greater in the eigenimage for the identified respiration rates were summed, and that summed

142

signal was plotted against time. The result is a distinctive waveform depicting the respiration over the 30 second acquisition time. This allows respiration rates to be determined for several subjects within a scene and their individual respiration waveforms to be monitored.

7.6 Measuring Flow Through a Pipe

It is possible to detect and measure fluid flow in a pipe by studying turbulence originating from the flow itself. The turbulence does not have to be imaged to accomplish this, since its signature may be transferred to the surface displacement on the pipe. We can detect this with either a single element sensor or camera array by measuring a feature on the surface with sufficient contrast. The edge of the pipe should offer the highest contrast but we are not limited to it, since marks on the pipe could serve as well. The motion of that feature is followed by sampling the signal from sampling pixels over time. The time domain data are transformed to the frequency domain where turbulence gives a unique structure in the low frequency portion of the spectrum at acoustic frequencies. The frequency spectrum is then analyzed to find the frequency that corresponds to the upper scale length of the turbulence. Models are used to determine actual physical size of that characteristic scale length based on the diameter of the pipe.

It is known that the diameter, D, of a pipe determines the large scale eddy size. [58] Knowing the physical size of the scale length L and the frequency at which it occurs f, which is determined by the point where the frequency spectrum plateaus in the direction of the DC component, the velocity of the flow can be determined. [59] Simply, these parameter are related by

143

$$L = \frac{v}{f}$$

where,

$L = Upper\ scale\ length$

$v = velocity\ of\ the\ fluid\ flow$

$f = frequency\ of\ the\ upper\ scale\ length$

Now we have,

$$v = Lf$$

and the fluid flow rate can be determined from the calculated scale length L and the plateau frequency f.

The volumetric flux, q is given by,

$$q = v \times A$$

$$q = Lf\frac{\pi D^2}{4}$$

where,

$A = \frac{\pi D^2}{4} = cross\ sectional\ area\ of\ the\ pipe$

Since the diameter of the pipe determines the upper scale length, the volumetric flux also can also be determined from a turbulence measurement. If the edge of the pipe is used for detection, its absolute displacement can be determined as well. Sampling rates would be determined by the frequencies of interest.

144

CHAPTER VIII:

IMAGING CLEAR AIR TURBULENCE

8.1 Introduction

In theory any phenomenon that produces a modulation in light levels could be measured by the proposed system. Features such as clearly defined edges on a surface are not the only phenomenon that produce a modulation in light levels and one such example is light transmitted through or scattered from a liquid or gas in turbulent flows.

The direction and intensity of light passing through a turbulent medium fluctuates with changes in the index of refraction along the optical path. These fluctuations are due to density or temperature inhomogeneities within turbulent eddies in the fluid. There are significant problems in detecting these fluctuations in a wide field of view of a daytime scene without special sensors and signal processing. For one, the variations in amplitudes are much too small to detect directly with commercially available video cameras and for another, the fluctuations in signal occur in too short of a period of time to be noticed by eye, or in some cases to be measureable at the standard video frame rate.

Turbulence produces variations in air density. In moving air these fluctuations are transported across the line of sight so that variations in intensity of transmitted light are correlated with the scale length of the packets and the flow speed. Imaging data treated in the temporal frequency domain quantify the turbulent fluctuations and reject the significantly stronger constant background. Fourier processing can filter the steady background, and only the fluctuations from turbulence are seen. We use these variations in light levels temporally at each point, and spatially across a field of view as a method to detect turbulence. In order to see these small variations in a large background, we need a system that offers a large dynamic range and low noise. Furthermore in order to map the turbulence in the field of view we need a detection scheme that spatially images temporal fluctuations.

Here we outline a technique for passive detection of turbulence using high dynamic range imaging (14 - 20 bits) to follow these subtle changes in transmitted ambient light intensity. In some cases the turbulence will be clear air turbulence, CAT, turbulence that is not visible to the naked eye. The ambient light fluctuations are used to create a temporal frequency spectrum for turbulent and non-turbulent cases by Fast Fourier Transform (FFT) processing. Turbulent fields show stronger low frequency (2-10 Hz) components over non-turbulent conditions, thus enabling a simple and straight forward means of threshold detection of turbulence. [60]

Four cases are presented in which we detect and image turbulence using two different techniques. Each technique relies on a large dynamic range and

sensitivity in order to detect changes in light intensity both temporally and spatially that are imperceptible to the casual observer.

The first imaging technique was implemented using the single InGaAs photodiode with the high gain low noise amplifier yielding a dynamic range around 20 bits, while maintaining a relatively low cost. The signal was digitized using an NI-9234 24-bit ADC from National Instruments and processed using LabVIEW data acquisition software. The large dynamic range allows us to detect variations in signal as small as one part in a million at a single sample point while remaining unsaturated from bright sunlit backgrounds.

The second technique utilizes the high speed cameras, a Photron Fastcam (APX RS) and an Allied Vision Pike. The required large dynamic range was achieved by oversampling at rates as fast as 60,000 fps and then co-adding images to increase the effective bit depth of the camera. The high speed camera has the advantage of spatially resolving the turbulence by simultaneously acquiring data points across the field of view.

The inherent noise of the single element photodiode system was measured to be on the order of 600 nV. The amplifier saturates at 6 V giving the photodiode system a possible dynamic range of 70 dB, where a factor of 10 corresponds to 10 dB.   However 1 V is more typical of a bright sunlit backgrounds which gives the detector a dynamic range of 60 dB in real world conditions. The camera was operated at near saturation levels. The detector noise will be governed by Poisson Statistics which gives us a noise level of $\sqrt{N}$, where N is the photon count. The noise in the low frequency portion of the

147

spectrum, the area of interest to us, is not governed by detector noise but instead dominated by background noise, some of which is background turbulence. This makes the overall noise of the system much larger than detector noise.

First we looked at a "sheet" of water vapor originating from a cold pipe. We clearly saw the turbulence in this case and used this as a validation tool in the detection of turbulence. We measured the changes in light signal and imaged the different scale lengths associated with turbulence to ensure it corresponds to the well-known modeled and expected behavior.

The second case involved imaging conditions where we fully expected there would be strong turbulence created from the temperature gradient above a hotplate. This allowed us to image a case of CAT where we can validate the results with a reasonable assumption of where the turbulence will be located.

Finally we looked at a case of CAT inside a cardboard tube used to shield the optical path from external disturbances of the imaged air. We injected a flow of gas from a compressed nitrogen tank into the tube and imaged the resulting turbulence. This case of CAT was much weaker in terms of modulated light than the others and imperceptible by eye.

8.2 Water Vapor Turbulence Frequency Spectrum

We started by observing a well-known laboratory phenomenon where turbulence is clearly discernible because of entrained water vapor. Downward airflow over a super-cooled pipe was measured in order to identify specific features and regions in the turbulent flow and associate them with their respective frequency spectra. Atmospheric water vapor condensed on the pipe,

and then sublimated as a dense aerosol tracer into the turbulent flow driven by gravity.

We measured light signal as a function of time for one sampled spatial element of this scene using a single photodiode sensor and a lens focused on a spot approximately 5 mm in diameter. The field of view of the sensor was moved in 1 cm steps from the bottom of the pipe in order to determine the temporal signal in regions with different spatial turbulent scale lengths. Data were collected for approximately 1 minute at each sampled element at 2048 samples/s to a maximum distance of 20 cm below the pipe. The first few centimeters were dominated by a laminar flow where no turbulence is discernible. Farther from the pipe we saw the beginnings of the classical turbulence where larger eddies formed first, followed by smaller eddies until viscosity dispersed the energy and the turbulence was no longer visible. Figure 69 shows frequency spectra of the signal 9 cm below the pipe, where turbulence appeared to be at a maximum, and 20 cm below the pipe where the turbulence was no longer visible. The low frequency rise in amplitude characteristic of turbulence is clearly evident in the measurement at 9 cm and is much greater than at 20 cm below the pipe. We measured a 9 db signal-to-noise ratio for the region integrated from 2 to 10 Hz. The 20 cm measurement represents the background signal, which for our purposes can be regarded as system noise, although we think it contains a characteristic room air turbulence background as well. This low frequency rise from 2 to 10 Hz gives us a very simple method of detecting turbulent conditions.

Turbulent Flow Below Pipe

Single Element Sensor
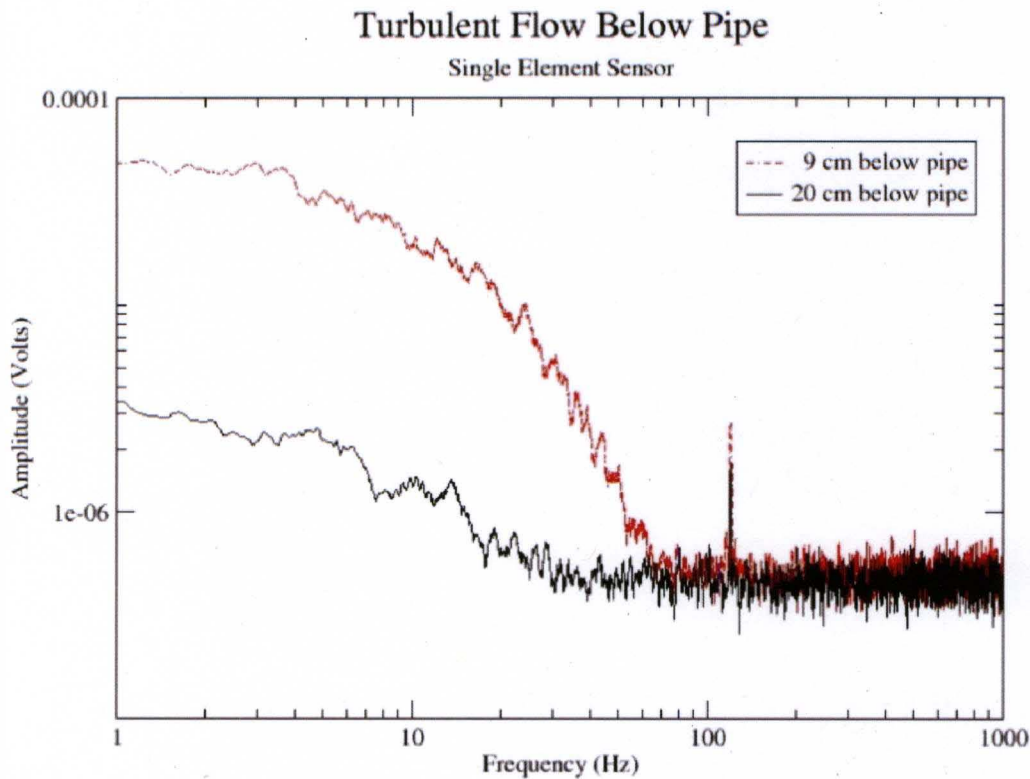
Figure 69. Frequency spectrum of water vapor turbulence. Water vapor turbulence imaged 9 cm (red) and 20 cm (black) below the pipe (see Figure 70) with the single element sensor. The low frequency rise in the 9 cm plot shows a simple and clear method to detect turbulence. The 20 cm plot is indicative of the background noise. The integrated signals from 2-10 Hz are measured and compared to determine the relative signal above that background. The strong peak at 120 Hz is due to the room lighting.

Next we imaged the same scene of water vapor in a turbulent sheet using a Photron Fastcam high speed camera. Image data were collected at a rate of

500 frames per second with the resolution of the camera set to its full frame of 1024x1024 pixels. We collected a total 6144 frames (limited by on-camera storage) yielding approximately 12 seconds of image data. We then co-added 4 frames (reducing the effective frame rate to 125 fps) in order to increase the dynamic range as well as reduce the size of the data set for faster processing.

The co-added frames were then used to produce a frequency spectrum for each pixel to produce a 1024x1024x1536 data space with two spatial axes and one temporal axis. Taking a pixel-by-pixel Fourier transform (not a spatial transform) along the time axis resulted in a transformed data space in which the third axis is frequency rather than time. Separate sample planes in the transformed space are images of the scene in distinct frequency ranges. An integration of the 2-10 Hz components displayed in false color reveals the spatial distribution of the elements of the scene with signals in the 2-10 Hz region characteristic of turbulence. To eliminate bias due to static non-uniformities (that is, variations in illumination or imager response) across the scene, a corresponding 0-2 Hz image was created, representing the constant "flat field" across the image. This reference image was used to normalize the intensities for the images of turbulence at higher frequencies. On the left side of Figure 70 is a single gray-scale frame from the high speed visible camera. At 9 cm below the pipe the region is highly turbulent, and at 20 cm below the pipe the turbulence is nearly indistinguishable from the background. The normalized integrated frequency intensity is shown in false color on the right side of Figure 70. The false color image makes it possible to clearly see the relative intensities of the

turbulence as well as the spatial distribution of the turbulence. The signal-to-noise ratio for a pixel showing strong turbulence is in this image is 14 dB. Notice that laminar (non-turbulent) flow directly below the pipe has been rejected in the frequency image as it is not seen in the integrated 2-10 Hz component of the spectrum.
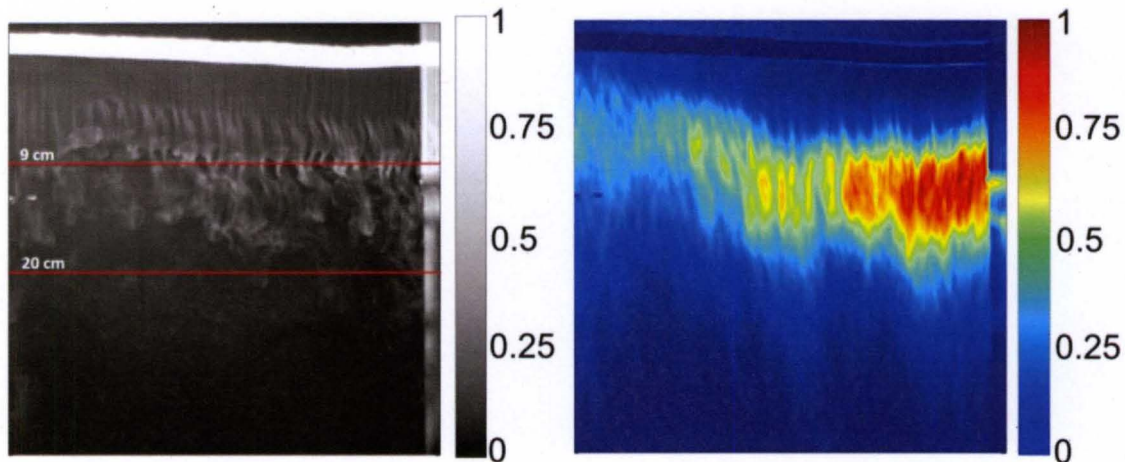
Figure 70. Frequency spectrum image of water vapor turbulence. (Left) Gray scale image of image of water vapor in turbulent air flowing downward over a liquid nitrogen pipe as seen by the high speed camera. (Right) False color image of the normalized integrated signal from 2-10 Hz. Dark blue denotes no turbulence and red is associated with the largest turbulence. The most turbulent region stands out in red in the filtered spectral image on the right. Rendering: Adam Willitsford

Figure 71 shows the spectral images for individual frequencies ranging from 2 to 24 Hz. Each image is the integrated signal centered on the labeled integer values with a width of 1 Hz, e.g. the 6 Hz image is the integrated intensity

of 5.5 to 6.5 Hz. As expected, the amplitudes diminish as the frequency gets higher. In addition, the scale lengths diminish as the frequency gets larger, which clearly shows that longer scale lengths are associated with lower frequencies while shorter scale lengths are associated with higher frequencies. [61]

2 Hz   4 Hz   6 Hz   8 Hz

10 Hz   12 Hz   14 Hz   16 Hz

18 Hz   20 Hz   22 Hz   24 Hz
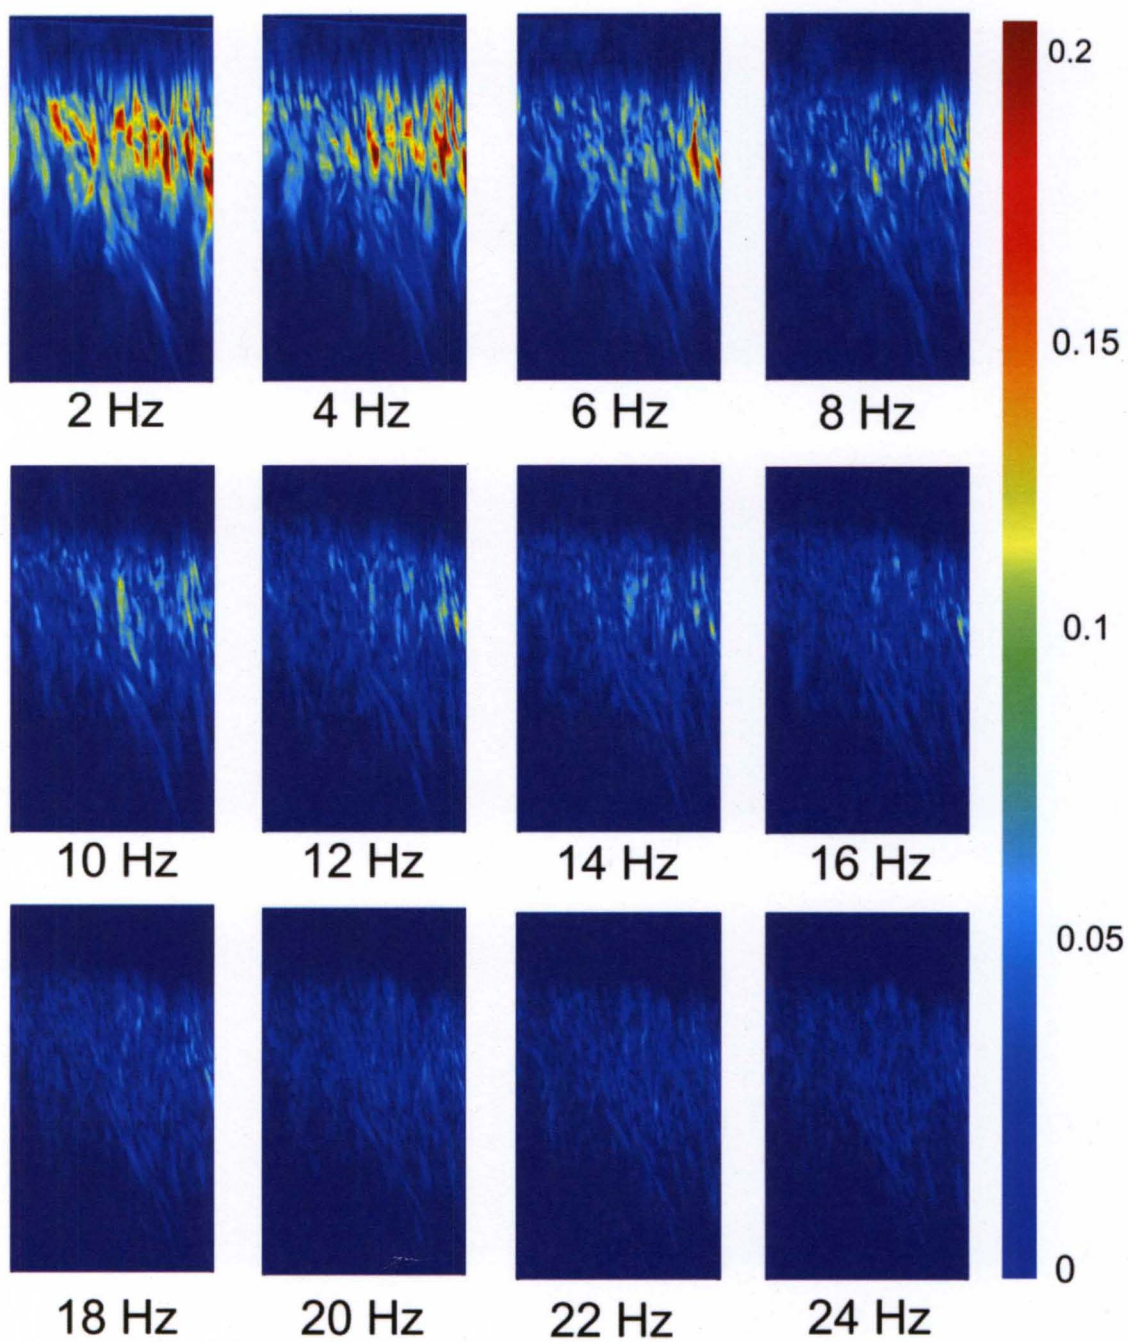
Figure 71. Frequency Resolved Images of Water Vapor Turbulence. Each frame is the integrated signal centered on the labeled frequency with a width of 1 Hz. For example, the 6 Hz image is the integrated signal of 5.5 to 6.5 Hz. Dark blue denotes no turbulence and red is associated with the largest turbulence as indicated by the color bar on the right. The relative strength of the turbulence diminishes as the frequencies increases with the strongest turbulence associated with the lowest frequencies. The scale lengths decrease as the frequencies increase with the larger scale lengths being seen at the lower frequencies. Rendering: Adam Willitsford

8.3 Clear Air Turbulence - Temperature Gradient from a Hotplate

In the previous example the turbulent flow was visible to the camera and to the eye because water vapor aerosols traced the gas flow. To image a scene that has known CAT, but is far less obvious, we consider a scene in strongly turbulent air but absent an aerosol tracer. Using the methodology of extracting images in the Fourier domain, we also sense fluctuations in the index of refraction due to air density changes as these disturbances subtly alter the detected signal. We consider first a case where these effects are very pronounced.

We disturbed the observed optical path with a hotplate that was placed in front of a small-scale grid target seen in the left side of Figure 72. The black and white pattern was placed behind the turbulence to maximize the modulation of

155

the signal resulting from when the refraction in the air shifts image elements on and off the black and white boundaries as seen at the detector. To image the grid pattern, the high speed Pike camera was again operated at 107 frames/s. The right side of Figure 72 shows the false color map of air above a hotplate with integrated 2-10 Hz components normalized against a constant flat background The signal-to-noise ratio in this spectral image was measured to be 8 dB for a strongly turbulent region. Turbulence is clearly seen, formed from the thermally driven inhomogeneities in the index of refraction of turbulent air in the large temperature gradient above the hotplate. A movie sequence, Supplemental Video 2, was created showing the real-time variation of the turbulence above the hotplate. Each frame is the result of a 107 image FFT, representing 1 s worth of data. To time resolve the sequence, each new frame created from the FFT was staggered by 3 frames to create a 0.03 s interval in each frame of the movie. The left portion of the image shows every 3 frames of the original video sequence while the right shows the frequency extracted images depicting the turbulence.
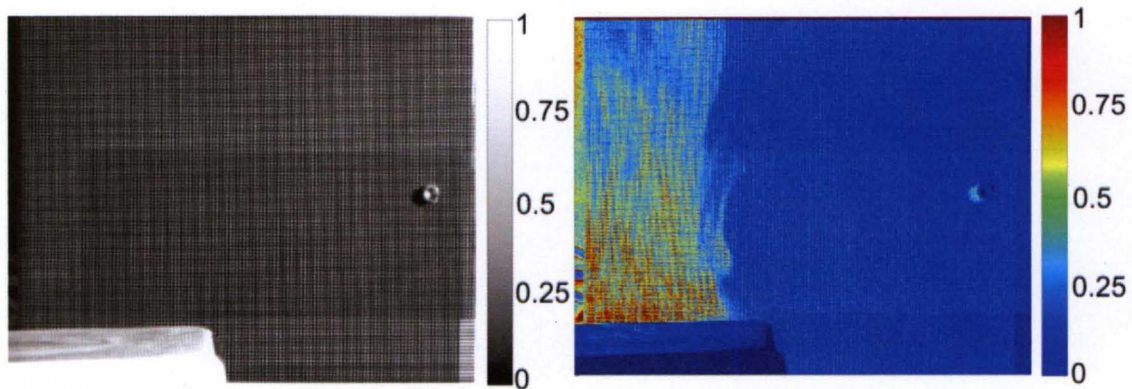
Figure 72. Thermal turbulence of a hotplate. (Left) Gray scale image of the hotplate and the black and white grid pattern. The black and white grid pattern was used to increase the modulation levels of the light. (Right) The integrated 2-10 Hz frequency intensity is shown for turbulence created above the hotplate. Blue denotes no turbulence and red is associated with the largest turbulence. Rendering: Adam Willitsford

8.4 Clear Air Turbulence - Temperature Gradient from a Candle

Similar measurements can be made from turbulence resulting from the heated air above a candle.
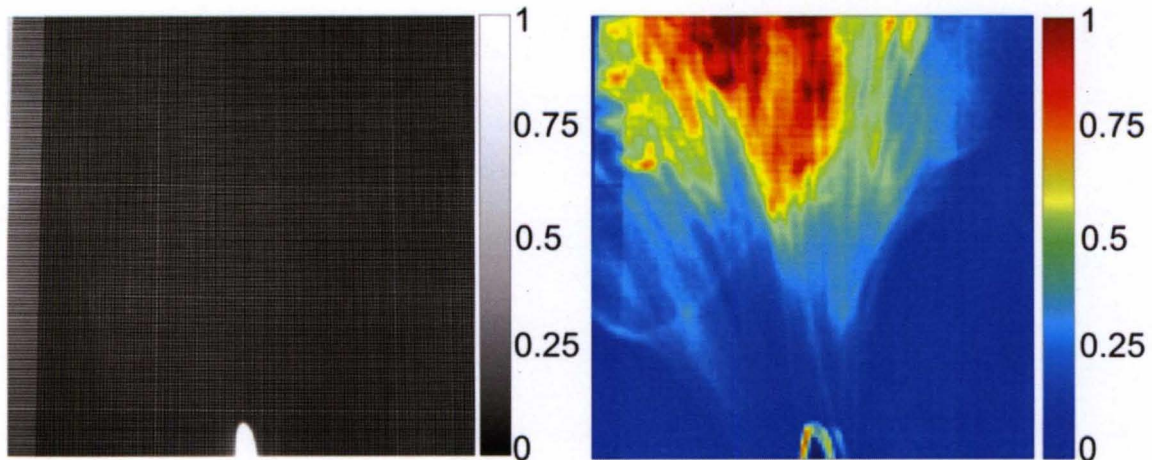
Figure 73. Thermal turbulence of a candle. (Left) Gray scale image of the candle and the black and white grid pattern. The black and white grid pattern was used to increase the modulation levels of the light. (Right) The normalized integrated 2-10 Hz frequency intensity is shown for turbulence created above a burning candle. Dark blue denotes no turbulence and red is associated with the largest turbulence. Rendering: Adam Willitsford

We disturbed the observed optical path with a candle that was placed in front of a small-scale grid target seen in the left side of Figure 73. To image the grid pattern, the high speed Photron camera was again operated at 500 frames/s and post-processed to 125 frames/s as previously described. The right side of Figure 73 shows the false color map of air above a burning candle with integrated 2-10 Hz components normalized against a constant flat background The signal-to-

noise ratio in this spectral image was measured to be 14 dB for a strongly turbulent region. The color image shown is reduced in size by spatially binning 8x8 areas to remove the grid pattern for purposes of clarity. Turbulence is clearly seen, formed from the thermally driven inhomogeneities in the index of refraction of turbulent air in the strong temperature gradient above the candle. A movie sequence, Supplemental Video 3, was created showing the real-time variation of the turbulence above the candle. Each frame is the result of a 125 image FFT, representing 1 s worth of data. To time resolve the sequence, each new frame created from the FFT was staggered by 1 frame to create a 0.008 s interval in each frame of the movie. The video was rendered at 60 frames per second so the resulting playback is reduced to approximately half speed.

8.5 Clear Air Turbulence - Invisible Compressed Gas Flow

The example of the hotplate is a clear demonstration of CAT imaging. However small turbulent fluctuations are still discernible in the individual camera frames, and they are sometimes perceivable to the eye. Much more subtle effects can be extracted that are visually imperceptible without processing. This is suggested by the low frequency rise in the background shown in Figure 69. To detect CAT that is both indiscernible in single video frames and to the eye, we used a flow of dry nitrogen in quiescent room air as a target.

To eliminate any outside effects on our measurements, namely room air currents, we placed a cardboard tube, approximately 1.6 m in length, in the optical path. This shielded any outside flow from entering the optical path. We drilled a small hole at the midpoint of the tube to allow the injection of

159

compressed dry nitrogen to create turbulence along the line of sight Figure 74

shows an increased amplitude in the low frequency spectrum through the tube

when the compressed gas is injected (red line labeled "With Flow in Tube") as

compared to the background measurement without the flow (black line labeled

"No Flow in Tube"). The signal-to-noise ratio for the 2-10 Hz region is 8 dB. The

peaks below 90 Hz seen in the background signal in Figure 74 coincide with

frequencies measured by an accelerometer placed on the tubing. The 97 Hz

signal may represent an organ pipe resonance in the tube which is open at both

ends. We calculated the pipe resonance to be approximately 106 Hz only 8.5%

difference from the measured value. Although not displayed in the graph, it is

interesting to note that these measurements were repeated without the tubing in

place. The removal of the tube gave us a measure of the room's inherent air

turbulence. The frequency background spectra were comparable to the

turbulence excited in a quiet tube. The turbulent flow inherent to the room was

tracked to an air-handling register in the ceiling above the workspace.

Gas Flow in Tubing
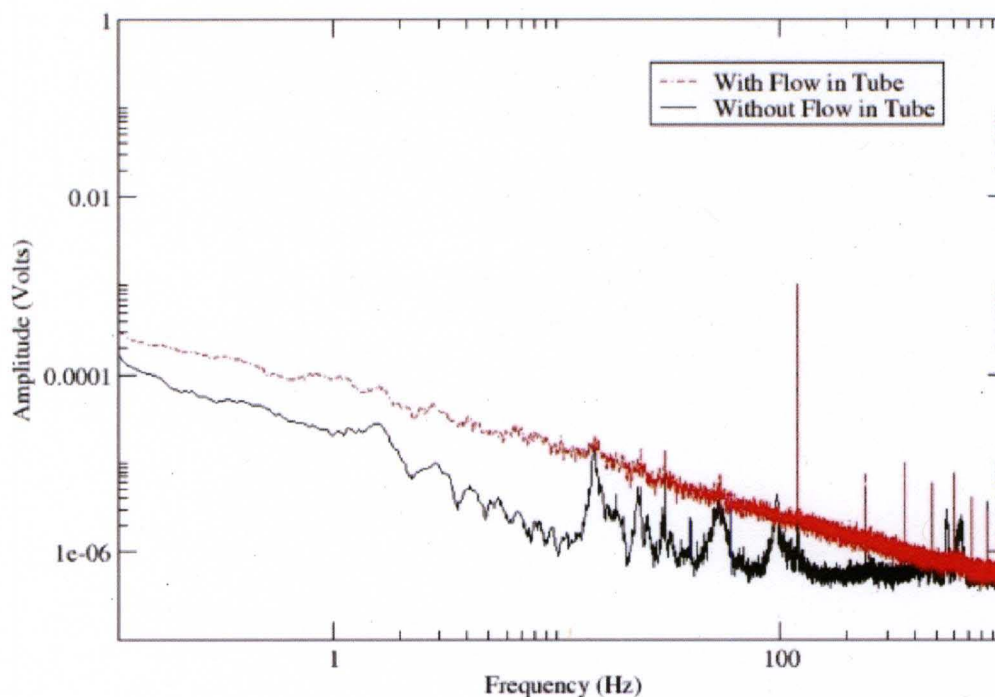
Figure 74. CAT from a compressed flow of dry nitrogen. The black line denotes the non turbulent case with the tubing to reduce the effects of room air currents. The red line shows the turbulent case with nitrogen injected into the tube.

The successful detection of the CAT indicates that modulations in light levels were present upon the introduction of gas into the tube. The modulation level was lower than previous examples since individual frames did not show obvious variations on playback, in contrast to the case with the hotplate. To spatially resolve the turbulence in the scene and see the evolution of the turbulence through time, the high speed camera was optimized to ensure the highest sensitivity possible.

The high speed Photron Fastcam was set up to image through the tube to our grid pattern at the far end of the tube. We recorded images at a rate of 60,000 fps. At this speed the camera is limited to a maximum region of 128x128 pixels. We recorded for a total of ~6.5 s, limited by the cameras internal storage, yielding approximately 390,000 frames. Approximately halfway through acquisition (t ≈3) we injected compressed air into the tube creating turbulent flow in clear air. To the eye there was no distinguishable change in the image, true clear air turbulence. To image this flow we post-processed the 390,000 frames in order to pull the clear air turbulence frequency spectra out of the noise.

We have co-added 300 frames which increases the dynamic range of the camera from 10 bits to a little over 18 bits allowing close to the sensitivity of the single element sensor. This co-addition also reduces the effective frame rate to 200 fps. As with the water vapor images, FFTs were performed on each time sequence for each pixel creating a single frequency spectral image representing the intensity of the 2-10 Hz region of the spectrum which was then normalized with the 0-2 Hz DC signal as previously described.

We processed the data using 200 frames per FFT giving us a spectral bin width of 1 Hz and a temporal size of 1 s for each frequency spectral image. To see the evolution of the turbulence over time we staggered the initial frame of the 200 frame length FFT by 10 frames which allows us to create a new frequency spectrum image every ten frames (0.1 second step between frames) for a total of 111 frequency images over the entire acquisition time. We see the spatial

evolution of the turbulence in the scene associated with the 2-10 Hz frequency range when played back sequentially as shown in Supplemental Video 4.

By injecting the turbulent flow halfway through a single acquisition we benefit by comparing cases of flow and no flow during the same image sequence. The evolution of the turbulence was clearly seen with our imaging and post-processing technique using the normalized integrated intensities associated with the 2-10 Hz components of the frequency spectrum. A measured signal-to-noise ratio for strong turbulence was 11 dB. Figure 75 shows 9 images representing the 2-10 Hz normalized integrated intensity centered around the time indicated. The sharp increase in signal marks the initial formation of the turbulence at ~3s. As the turbulence evolves, large scale structure is first seen which then breaks into smaller structure as theorized and modeled by Kolmogorov. [62]

t = 0.05 s   t = 2.85 s   t = 3.25 s

t = 3.6 s   t = 3.85 s   t = 4.45 s

t = 4.85 s   t = 5.25 s   t = 5.45 s
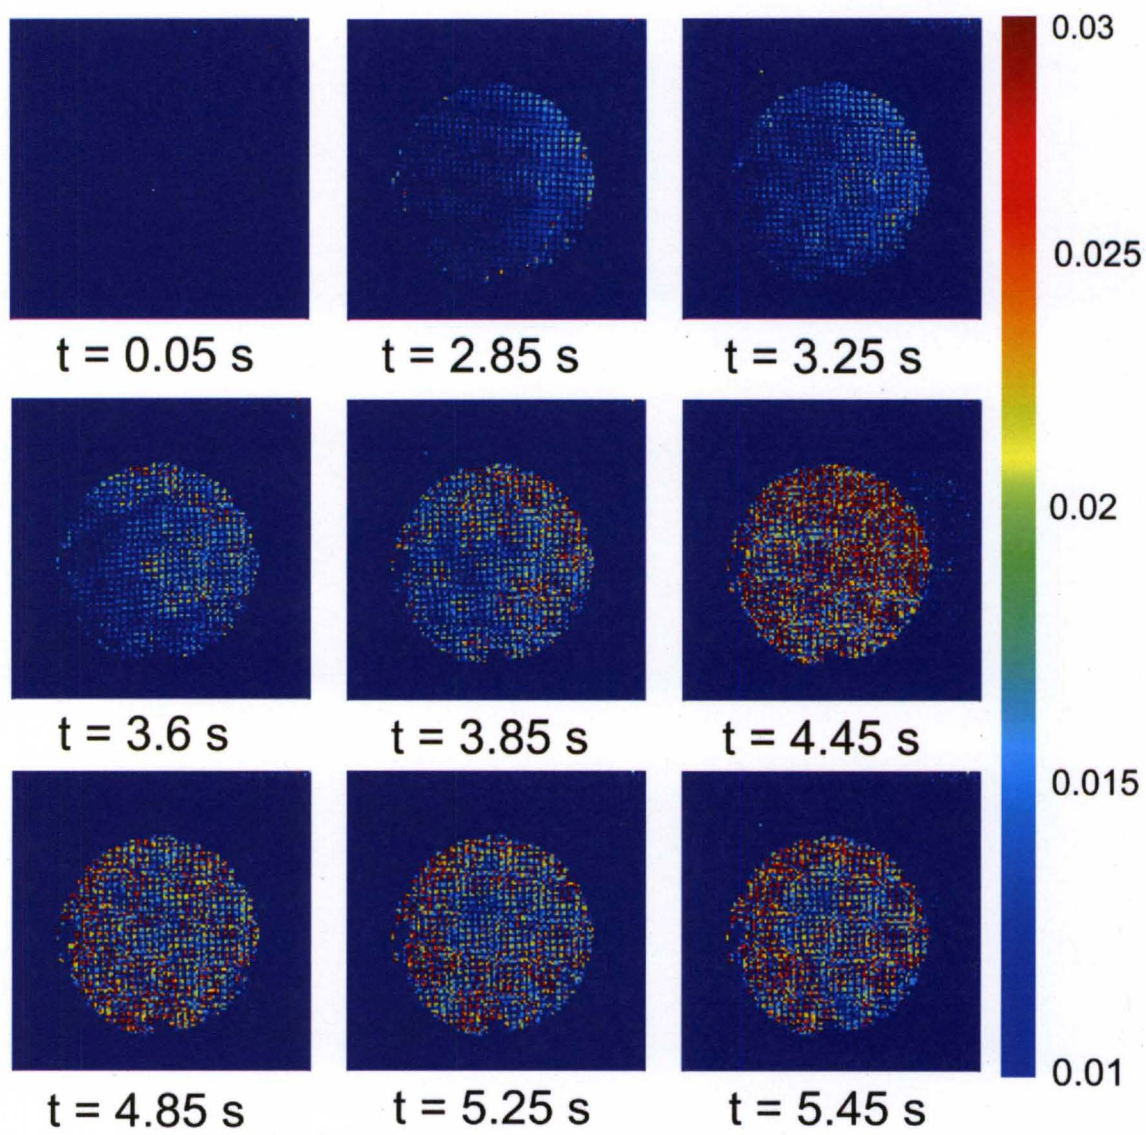
Figure 75. Frequency Spectrum Images of CAT from a Compressed Flow of Nitrogen. Evolution of turbulence when air is injected into a tube at t ≈ 3 s. Each image represents the normalized integrated intensity of 2-10 Hz, using a sliding time window of 1 second centered around each indicated time. Dark blue denotes no turbulence and red is associated with the largest turbulence. Turbulence first forms in the upper right section for the image where the nitrogen was injected. Eddies form quickly thereafter forming structures with smaller scale. Rendering: Adam Willitsford

8.6 Conclusions

We have demonstrated the ability to image CAT using high speed cameras and data processing methods that extend their dynamic range and sensitivity and utilize the frequency domain as a method to filter the turbulent signal. In addition, we have demonstrated that we are able to spatially resolve turbulence and place values on the relative intensities of turbulence throughout the scene. Our spatially resolved images in the frequency domain show the distinct characteristics of turbulence, namely the changes in signal and scale length. Our measurement is not limited to detecting the turbulence of an inhomogeneous gas, but can also measure CAT from a homogenous gas. This technique has the potential to be extended to applications including the passive detection of turbulence for aircraft safety on the runaway during take-off, and for aircraft in flight. This could help to increase the safety of aircraft by offering

165

additional early detection of turbulence. This method also has the potential to be used for detecting the turbulence associated with pipeline leaks that result in turbulent flow in the surrounding medium.

CHAPTER IX:

SUMMARY, CONCLUSIONS AND FUTURE WORK

9.1 Summary

The work of this dissertation involved developing a new technique for the remote measurement of motion. The concept of analyzing the fluctuations of light intensities was explored. A model that describes this technique was developed and tested to show the relationship of the measured signal to physical parameter. Finally, the technique was tested with multiple applications to show feasibility of the technique to be used as a valid instrument for detecting and measuring motion.

Many devices currently exist that aim to measure motion either through direct contact or by remote sensing. However each device suffers from its own limitations that hinder its ability to make a precise non-contacting measure of motion remotely, and at low cost. The applied research described here have been developed to do this and have been tested on several practical applications.

Previous work has been done that shows that features of objects, or varying levels of contrast of objects, can be used to determine position with sub-pixel accuracy. We have extended this concept and applied it to the

measurement of motion. There are distinct advantages of our new approach over others currently available, those being: it is passive, offers high dynamic range, it is suitable in dynamic lighting situations, and it is extremely sensitive to small displacements.

We have developed this technique around two primary modes of data acquisition. One is a single element sensor that consists of imaging lens coupled to a fiber optical cable that directs the focused light to a photodiode. Its photo current is then amplified and digitized to 24-bit accuracy. The second system is composed of high dynamic range video technology that achieves the necessary bit depth by a combination of the co-addition of frames, regioning, and pixel binning, when necessary.

The video technology was used to record the time varying images processed to produce eigenimages of the field of view. These images are able to show the principal frequency components of the scene over two spatial dimensions. Furthermore data cube sequences were produced that create a series of spatial images representing a time series of all the unique frequencies in the scene along with their spatial dimension, limited only by the optical system, sampling rate, and total data acquisition time. Work was done to show the feasibility of collecting and data processing these measurements in real time with the use of the GPU.

Finally the technique was applied to show the ability to make measurements on real systems in real-world conditions and return useful and accurate information.

9.2 Conclusions

The high dynamic range imaging methodology was shown to yield signals that are linearly proportional to displacement when measuring the motion of an edge. A comparison with a laser Doppler vibrometer yielded a calibration for operation with high dynamic range passively under room lighting conditions.

It is possible to extend this measurement of displacement to a measurement of periodic motion of a vibrating source. Frequencies relating to motion have been identified and measured relating to the motion of the object being measured. In the case of the single element sensor, this has been done in real time to show the frequency spectrum. In the case of video technology, this depends on post processed data to reveal the frequency spectrum of each pixel and to create a three dimensional data cube where two dimensions are the spatial dimensions and the third is frequency. Each successive image plane in the cube shows the frequency response at that particular frequency simultaneously across all points in the scene. The resulting data volume reveals dynamical information about each point in the scene.

Although the video data was post-processed for this research, work was done to show the feasibility of processing in real time. An NVIDIA GPU was used to process thousands of Fast Fourier Tranforms (FFT's) in parallel.

Benchmark tests showed that the necessary number of FFT's of sufficient length could be done within the data acquisition window to allow real time processing of the video signals into frequency data cubes. This hardware would allow for data to be collected and processed in real time, a condition that would be useful for many applications.

The technique was applied to the motion of the structure of bridges and shown to yield that the fundamental resonant frequencies at specific locations. In addition the displacement of bridge components under real traffic load can be determined and monitored in real-time. Similar measurements were performed on mechanical systems providing vibration analysis through non-contact optical sensing. Such measurements allow a remote determination of whether mechanical devices are operating normally. The technique was applied to biomedical problems and used to determine the blood pressure pulse profile simultaneously ant many multiple points on the body with a high degree of accuracy, rivaling current state of the art measurements that require contact or invasive sensors. Respiration rates of individuals were measured across a scene and it was shown that the individual subjects could be recognized and monitored simultaneously. Finally, the technique was applied to measure and characterize turbulence. The inertial sub-region of turbulence in a gas flow was measured revealing information about the structure of the turbulence, the energy it contains, and the rate of flow. Clear air turbulence was successfully detected and imaged as well. It was shown that these measurements show the dynamic evolution of turbulence.

9.3 Future Work

More work will be done with the technology discussed in this dissertation to extend the scope of applications as well as refine the technology. Currently the system is known to make a measurement of displacement as accurate as that from an industry standard Doppler Laser Vibrometer. A next step will be to develop algorithms to make the system self calibrating based on features within the scene. More work will be performed with the GPU to integrate it into the workflow of data acquisition, processing and display so that the data is acquired and displayed seamlessly and usefully in real time. A system is envisioned where time series video is taken and displayed in 3 dimensions. The user could then select part or the entire image and have it transformed to show a specific frequency. Data would be acquired and processed in the GPU, in real time, within one software package. This system would be useful for field work to allow real time feedback in the data acquisition process, especially to structural health monitoring.

Work to extend the scope beyond the applications described here will follow, since the ones described are by no means all inclusive. For example, data such as that collected from the drum head imaging could be of used in manufacturing for quality control.

More work is presently planned with analyzing the structural health of bridges with both the single element sensor and video camera technology. A synchronized system of *in situ* strain gauges and accelerometers is being developed to further correlate the optical measurements with structural defects

and physical properties of the bridge. The system under development will allow for simultaneous measurements of strain, acceleration and displacement at many locations on a bridge.

REFERENCES

[1] P. de Bernardis, P. A. R. Ade, J. J. Bock, J. R. Bond, J. Borrill, A. Boscaleri, K. Coble, B. P. Crill, G. De Gasperis, P. C. Farese, P. G. Ferreira, K. Ganga, M. Giacometti, E. Hivon, V. V. Hristov, A. Iacoangeli, A. H. Jaffe, A. E. Lange, L. Martinis, S. Masi, P. V. Mason, P. D. Mauskopf, A. Melchiorri, L. Miglio, T. Montroy, C. B. Netterfield, E. Pascale, F. Piacentini, D. Pogosyan, S. Prunet, S. Rao, G. Romeo, J. E. Ruhl, F. Scaramuzzi, D. Sforna, and N. Vittorio, Nature **404**, 955-959 (2000).

[2] S. Perlmutter, G. Aldering, M. Della Valle, S. Deustua, R. S. Ellis, S. Fabbro, A. Fruchter, G. Goldhaber, A. Goobar, D. E. Groom, I. M. Hook, A. G. Kim, M. Y. Kim, R. A. Knop, C. Lidman, R. G. McMahon, P. Nugent, R. Pain, N. Panagia, C. R. Pennypacker, P. Ruiz-Lapuente, B. Schaefer, and N. Walton, Astro-Ph/9712212 (1997).

[3] A. G. Riess, A. V. Filippenko, P. Challis, A. Clocchiatti, A. Diercks, P. M. Garnavich, R. L. Gilliland, C. J. Hogan, S. Jha, R. P. Kirshner, B. Leibundgut, M. M. Phillips, D. Reiss, B. P. Schmidt, R. A. Schommer, R. C. Smith, J. Spyromilio, C. Stubbs, N. B. Suntzeff, and J. Tonry, The Astronomical Journal **116**, 1009-1038 (1998).

[4] R. A. Sunyaev and Y. B. Zel'dovich, Annu. Rev. Astro. Astrophys. **18**, 537-560 (1980).

[5] J. P. Lynch and K. J. LOH, The Shock and Vibration Digest **38**, 91-128 (2006).

[6] J. Oldengarm, A. H. van Krieken, and H. J. Raterink, Optics & Laser Technology **5**, 249-252 (1973).

[7] N. Halliwell, Journal of Sound and Vibration **62**, 312-315 (1979).

[8] P. Castellini, G. Revel, and E. Tomasini, The Shock and Vibration Digest **30**, 443-456 (1998).

[9] Http://Www.metrolaserinc.com.

[10] P. Castellini, M. Martarelli, and E. Tomasini, Mechanical Systems and Signal Processing **20**, 1265-1285 (2006).

[11] C. L. Desjardins, L. T. Antonelli, and E. Soares, in *Proceedings of SPIE* (San Jose, CA, USA, 2007), pp. 64301C-64301C-9.

[12] Mei Chen, J. O'Sullivan, N. Singla, E. Sirevaag, S. Kristjansson, Po-Hsiang Lai, A. Kaplan, and J. Rohrbaugh, Information Forensics and Security, IEEE Transactions On **5**, 449-460 (2010).

[13] K. Meigas, H. Hinrikus, R. Kattai, and J. Lass, J. Biomed. Opt. **8**, 152 (2003).

[14] L. T. Antonelli, J. F. Lomba, and W. J. Ohley, U.S. Patent No. 7128714 (31 October

2006).

[15] C. Huston, Journal of the Franklin Institute **107**, 41-44 (1879).

[16] J. R. Davis, *Tensile Testing*, 2nd ed. (ASM International, 2004).

[17] G. F. Lucas, P. C. McKeighan, and J. S. Ransom, *Nontraditional Methods of Sensing Stress, Strain, and Damage in Materials and Structures* (ASTM International, 2001).

[18] W. E. Luecke and J. D. French, Journal of the American Ceramic Society **79**, 1617-1626 (1996).

[19] G. L. Pressman and P. M. Newgard, Bio-Medical Electronics, IEEE Transactions On **10**, 73-81 (1963).

[20] B. Grolman, Am J Optom Arch Am Acad Optom **49**, 646-660 (1972).

[21] M. Shields, Survey of Ophthalmology **24**, 211-219.

[22] J. Booth, Proc R Soc Med **70**, 793-799 (1977).

[23] H. Nyquist, American Institute of Electrical Engineers, Transactions of The **47**, 617-644 (1928).

[24] Http://Virdir.ncsa.illinois.edu/Partiview/.

[25] K. Moreland and E. Angel, in *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hardware* (Eurographics Association, Aire-La-Ville, Switzerland, Switzerland, 2003), pp. 112–119.

[26] J. Nickolls, I. Buck, M. Garland, and K. Skadron, Queue **6**, 40-53 (2008).

[27] R. Stone and H. Xin, Science **330**, 746 -747 (2010).

[28] K. Rasul, Https://Github.com/Kashif/Mathematica_cuda/Tree/Master/Src/cuFourier.

[29] *Infrastructure Fact Sheet* (American Society of Civil Engineers, 2009).

[30] A. E. Aktan, A. J. Helmicki, and V. J. Hunt, Smart Materials and Structures **7**, 674 (1998).

[31] J. Ko and Y. Ni, Engineering Structures **27**, 1715-1725 (2005).

[32] O. S. Salawu, Engineering Structures **19**, 718–723 (1997).

[33] Http://Plasma-Gate.weizmann.ac.il/Grace/.

[34] Q. Zhang and X. Su, Opt. Express **13**, 3110-3116 (2005).

[35] V. L. Roger, A. S. Go, D. M. Lloyd-Jones, R. J. Adams, J. D. Berry, T. M. Brown, M. R. Carnethon, S. Dai, G. de Simone, E. S. Ford, C. S. Fox, H. J. Fullerton, C. Gillespie, K. J. Greenlund, S. M. Hailpern, J. A. Heit, P. M. Ho, V. J. Howard, B. M. Kissela, S. J. Kittner, D. T. Lackland, J. H. Lichtman, L. D. Lisabeth, D. M. Makuc, G. M. Marcus, A. Marelli, D. B. Matchar, M. M. McDermott, J. B. Meigs, C. S. Moy, D. Mozaffarian, M. E. Mussolino, G. Nichol, N. P. Paynter, W. D. Rosamond, P. D. Sorlie, R. S. Stafford, T. N. Turan, M. B. Turner, N. D. Wong, J. Wylie-Rosett, and on behalf of the American Heart Association Statistics Committee and Stroke Statistics Subcommittee, Circulation (2010).

[36] W. Rosamond, K. Flegal, K. Furie, A. Go, K. Greenlund, N. Haase, S. M. Hailpern, M. Ho, V. Howard, B. Kissela, S. Kittner, D. Lloyd-Jones, M. McDermott, J. Meigs, C. Moy, G. Nichol, C. O'Donnell, V. Roger, P. Sorlie, J. Steinberger, T. Thom, M. Wilson, Y. Hong, and for the American Heart Association Statistics Committee and Stroke Statistics Subcommittee, Circulation CIRCULATIONAHA.107.187998 (2007).

[37] B. P. Imholz, W. Wieling, G. A. van Montfrans, and K. H. Wesseling, Cardiovasc Res **38**, 605-616 (1998).

174

[38] S. Laurent, P. Boutouyrie, R. Asmar, I. Gautier, B. Laloux, L. Guize, P. Ducimetiere, and A. Benetos, Hypertension **37**, 1236-1241 (2001).

[39] S. Laurent, S. Katsahian, C. Fassot, A. Tropeano, I. Gautier, B. Laloux, and P. Boutouyrie, Stroke **34**, 1203-1206 (2003).

[40] S. C. Millasseau, R. P. Kelly, J. M. Ritter, and P. J. Chowienczyk, Clin. Sci **103**, 371-377 (2002).

[41] A. L. Pauca, M. F. O'Rourke, and N. D. Kon, Hypertension **38**, 932-937 (2001).

[42] R. Kelly, C. Hayward, A. Avolio, and M. O'Rourke, Circulation **80**, 1652-1659 (1989).

[43] J. Blacher, R. Asmar, S. Djane, G. M. London, and M. E. Safar, Hypertension **33**, 1111-1117 (1999).

[44] R. Asmar, A. Benetos, J. Topouchian, P. Laurent, B. Pannier, A. Brisac, R. Target, and B. I. Levy, Hypertension **26**, 485-490 (1995).

[45] A. Yamashina, H. Tomiyama, K. Takeda, H. Tsuda, T. Arai, K. Hirose, Y. Koji, S. Hori, and Y. Yamamoto, Hypertension Research **25**, 359–364 (2002).

[46] G. L. WOOLAM, P. L. SCHNUR, C. VALLBONA, and H. E. HOFF, Circulation **25**, 533-539 (1962).

[47] J. Sugawara, K. Hayashi, T. Yokoi, M. Y. Cortez-Cooper, A. E. DeVan, M. A. Anton, and H. Tanaka, J Hum Hypertens **19**, 401-406 (2005).

[48] S. A. Hoeksel, J. R. Jansen, J. A. Blom, and J. J. Schreuder, J Clin Monit **13**, 309-316 (1997).

[49] W. B. Kannel, P. A. Wolf, D. L. McGee, T. R. Dawber, P. McNamara, and W. P. Castelli, JAMA: The Journal of the American Medical Association **245**, 1225 -1229 (1981).

[50] A. W. FEINBERG, H. LAX, and W. Urban, Circulation **18**, 1125-1130 (1958).

[51] H. Lax, A. W. Feinberg, and B. M. Cohen, Journal of Chronic Diseases **3**, 618-631 (1956).

[52] S. Laurent, J. Cockcroft, L. Van Bortel, P. Boutouyrie, C. Giannattasio, D. Hayoz, B. Pannier, C. Vlachopoulos, I. Wilkinson, and H. Struijker-Boudier, European Heart Journal **27**, 2588-2605 (2006).

[53] *A Treatise on Man and the Development of His Faculties : Quetelet, Adolphe, 1796-1874 : Free Download & Streaming : Internet Archive.*

[54] J. McBride, D. Knight, J. Piper, and G. B. Smith, Resuscitation **65**, 41-44 (2005).

[55] W. P. Glezen, S. B. Greenberg, R. L. Atmar, P. A. Piedra, and R. B. Couch, Jama **283**, 499-505 (2000).

[56] *State of Lung Disease in Diverse Communities* (American Lung Association, 2010), pp. 69-72.

[57] *Sleep Apnea* (National Heart Lung and Blood Insitute, n.d.).

[58] G. Gioia, P. Chakraborty, and F. A. Bombardelli, Phys. Fluids **18**, 038107 (2006).

[59] N. Short, W. Fitelson, and C. H. Townes, ApJ **599**, 1469-1477 (2003).

[60] T. Von Karman, Proceedings of the National Academy of Sciences of the United States of America **34**, 530 (1948).

[61] A. Wheelon, N. Short, and C. H. Townes, The Astrophysical Journal Supplement Series **172**, 720–731 (2007).

[62] A. N. Kolmogorov, Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences **434**, 9-13 (1991).

CURRICULUM VITAE

NAME: Jeffrey Robert Hay

ADDRESS: Department of Physics and Astronomy
                     102 Natural Science Building
                     University of Louisville
                     Louisville, KY 40220

EDUCATION: B.S., Physics
                     University of Louisville
                     2005

                     M.S., Physics
                     University of Louisville
                     2007

                     Ph.D., Physics
                     University of Louisville
                     2011

PROFESSIONAL EXPERIENCE :

                     Graduate Research Assistant
                     2007- Present

                     GEMS Fellow
                     (Groundwork Education in Mathematics and Science)
                     2006-2008

                     Graduate Teaching Assistant
                     2005-2007

                     Supervisor United Parcel Service
                     2001-2003

PROFESSIONAL ACTIVITIES, AWARDS AND HONORS:

Member, Society of Physics Students
2002-2008

Member, Sigma Pi Sigma
(Physics National Honors Society)
2005

Bullitt Scholarship Recipient
2003-2004

Global Venture Challenge Team Member
Finalists (Oak Ridge National Labs)
2010

SELECTED RELEVANT PUBLICATIONS/PATENTS:

Hay, J. Willitsford, A., and Kielkopf, J.   "Detection and Imaging of Turbulence in a Fluid". U.S. Patent Application Serial No. 12/911,278, filed October 25, 2010, University of Louisville Research Foundation, Inc. Ref. No. 11014. J. Hay, Willitsford, A., and J. Kielkopf, University Of Louisville.

Hay, J and Kielkopf, J.   "System and Method for Precision Measurement of Position Motion and Resonances" US Patent Application Serial No. 61/447,846, filed March 1, 2011, University of Louisville Research Foundation, Inc. Ref. No. 11047. J. Hay, J. Kielkopf, University Of Louisville.