Proposed Pixel for Custom Laser Doppler Vibrometry Camera

Stephen Jackson

Thesis submitted to the University of Nottingham for the degree of Doctor of Philosophy

September 2011

Abstract

Non-destructive testing is an important field of research in many areas of industry and science. This field covers methods that test some property of an object while not impacting its future usefulness. This project is specifically interested in a method of NDE called Laser Doppler Vibrometry (LDV), which uses light to probe the motion of an object's surface.

This thesis presents research into the possibility of developing a full-field LDV camera capable of measuring the vibration at a number of points across an object simultaneously. The approach chosen was to develop a single processing element which contained a light detector as well as the processing required to produce an output signal proportional to the vibration of the surface imaged by the element. This processing element, or pixel, could then be used to produce a full-field LDV camera capable of imaging an object's surface to provide vibration amplitude and frequency for a number of points across the whole surface simultaneously

Some early work was carried out measuring the vibration of a piezoelectric target with a CMOS camera and the results were compared with theory. The method of processing chosen is called zero-crossing demodulation and a circuit was designed and simulated using this method. This circuit was then laid-out and a CMOS chip was fabricated with the pixel on it. The circuit was designed such that the different parts of it could be tested separately using electrical signals, and a number of electrical tests were performed to check how well the operation of the fabricated pixel compared with simulation. Using a laser interferometer the pixel was finally tested with real vibrating objects and compared with a commercial LDV.

Definition of Terms

AOM Acousto-optic Modulator

CMOS Complementary Metal-Oxide Semiconductor

DSP Digital Signal Processor

ESPI Electronic Speckle Pattern Interferometry

FET Field Effect Transistor

FPGA Field Programmable Gate Array

Full-Field A method of measuring the properties, in this case vibration, of a surface concurrently - all points on the surface at the same time

HDA Hysteretic Differentiator Amplifier

LED Light emitting diode

LDV Laser Doppler Vibrometry

MI Michelson Interferometer

MLV Matrix Laser Vibrometer

MZI Mach-Zehnder Interferometer

NDE Non-destructive Evaluation

PCB Printed Circuit Board

Acknowledgments

Firstly I would like to thank my supervisors Prof. Stephen Morgan and Prof. Barrie Hayes-Gill for their encouragement, support and unbounded patience during this project.

I would also like to thank Dr. John Himsworth for his help during the design and layout of the processing circuit for this project, teaching me some of the software tools while also trying to do his own design and layout, and Dr. Diwei He for his help with the many questions I sent his way about his CMOS chip - several components from which were incorporated in my design.

This project was supported by the Engineering and Physical Sciences Research Council (EPSRC) and the School of Electrical and Electronic Engineering at the University.

Finally I would like to thank my housemates for the duration of my final year for putting up with me during my write-up, and my family for their support during the more difficult times.

Contents

1 Introduction and Literature Review

1.1 Aim of the Project

This thesis reports on a PhD project at The University of Nottingham to investigate a system for measuring the vibration of a region of the surface of an object in a non-contact manner. The vibration of the entire region needed to be measured concurrently requiring what is called a full-field measurement method. Full-field measurement preserves the relationship between separate points of the object and therefore can provide analysis of fast transient events. Using a non-contact method prevents the measurement from changing the vibration of the object through loading. This is especially important with light objects such as membranes. The method chosen was Laser Doppler vibrometry (LDV) and this is a well known method of vibration measurement, however it is normally a single point technique and this point must be scanned across an object to image an area of it. This method takes time and is best suited to repetitive measurements, however LDV does have the advantage of high bandwidth and good sensitivity.

The main task of this project was to miniaturise the LDV sensor and digital processing onto a small area of silicon, forming an individual, discrete processing and detector element. This element will be called the 'pixel' throughout this thesis as it performs similarly to the pixel in a camera. Light falling onto the pixel is processed in some way to provide electrical signals which can be sampled electronically. This processing pixel could then be repeated across a grid to form a LDV camera. If the object's surface was illuminated with diffuse laser light and appropriate optical processing was included, then the pixels would provide all the vibration information for the area of the object they image - a full-field LDV camera.

1.2 Rationale for the Project

The main problem with full-field, high-frequency vibration measurement is the need to capture a number of images in a very short time. If this system needs to produce real-time images of that vibration then each image must be processed in less than the image time. If the images were 32 x 32 pixels then there would be 1024 samples for each image. To image a vibration at $1\,kHz$ and with an amplitude of $10\,\mu m$ using a Doppler based system the light from the object would need to be imaged at at least twice the highest frequency present in reflected light. As will be shown in Chapter 2 the bandwidth would be $400\,kHz$ requiring images to be taken at $800\,kHz$. For the example image size of 32 x 32 this would result in just over 8×10^8 samples per second. This is a great deal of data and would require a fast, and expensive, processing system. The problems with this example system are that the images must be moved from the sensor to the processing at the image rate, and the image data must be processed to provide the vibration information.

One solution is to put the processing on the same chip as the camera. Now the images are transferred from the camera region to the processing on the same die. With this method the high bandwidth data transfer happens on the same silicon die where distances and capacitances are small and the transfer of large quantities of data is more efficient. The images could then be processed on the chip and the vibration data then transferred off for viewing/ storage. Given that the example vibration was at $1\,kHz$, the processed data need only be transferred off the chip at $2\,kHz$, four hundred times slower than the original image rate.

Silicon chips are priced on the area of silicon used so adding the processing onto the camera will make it more expensive, and a system capable of processing that much data would be complex. A solution might be to split the image in half and have two processing units working on the different sides of the image. Extending

this philosophy, if the processing could be made simple enough (small enough) it might be added to each pixel. Now each pixel is doing its own processing. The pixel samples the light falling on it and from that (high frequency) data produces an output related to the vibration signal.

The aim of this project is therefore to develop a pixel which could be used in a full-field LDV camera. The pixel must be small to make it a viable option for a large camera. If the pixel is too large it would be too costly to manufacture a camera with a useful resolution. Although as a prototype pixel, the one designed during this project will not apply stringent space requirements. It was considered more important to produce a working proof-of-principle pixel at this stage, than a very small one which can often cause unwanted interference and noise on a CMOS chip. The pixel must also provide a signal proportional to the vibration which can be sampled at an rate appropriate to the vibration, rather than the Doppler signal.

1.3 Overview of Thesis Structure

This chapter will provide an introduction to the project, the aims of the work and a review of previous research in the area. Applications of the camera will also be given to show the breadth of possible uses for a full-field vibrometry camera.

The broad field under which this work falls is Non-destructive Evaluation (NDE) which will be introduced in Section 1.4 before focusing on those methods most relevant to this work. There are a number of optical methods which could be used to fulfill the requirements of this project and the one chosen was Laser Doppler Vibrometry (LDV). This choice will be discussed Section 1.5.7. This section will also describe both single point and full-field types of LDV and compare it with other competing technologies. As part of the introduction to this work the Doppler effect will be described and this chapter will provide groundwork for the

processing method chosen.

Following on from this literature review, Chapter 2 will provide a theoretical description of the method chosen as well as some results demonstrating the theoretical operation of the processing. Chapter 2 will also describe the whole experimental setup which the pixel forms part of. This experimental setup takes the form of an optical interferometer. Chapter 2 forms the basis of the design of the pixel which is covered in Chapter 3. A circuit for the proposed pixel was designed and laid-out for fabrication on a $0.35\,\mu m$ CMOS silicon chip. Chapter 3 also describes the simulation of the pixel circuit as a whole as well as the different processing stages, showing how their individual operation affects the operation of the whole pixel. While the chip was being fabricated a board was designed to allow testing of the chip. This is covered in Chapter 4. The testing started with electrical tests to check the operation of the different parts of the chip individually. After this a modulated light source was used to check the operation of the pixel as a whole. Chapter 5 describes how the pixel and board were placed into the optical interferometer setup and used to measure real vibrating objects. The results from these test were compared with a commercial LDV. Finally Chapter 6 concludes the thesis with some suggestions for further work in this field, with the possibility of a complete camera based on the pixel described in this thesis.

1.4 Non-destructive Evaluation

1.4.1 Introduction

Non-destructive Evaluation (NDE) covers a large range of techniques and is widely used in manufacturing, construction and testing. NDE is the examination of an object, material or system with technology that does not affect the object's future usefulness. We naturally use NDE concepts in our everyday lives. For example when shopping for fruit, say a banana, we make a control decision (do I buy it?)

based on whether the outside is yellow, or has started to turn brown. We may additionally hold the banana and through touch decide whether it is still firm or has started to degrade.

Non-destructive Evaluation is a very important area of research in today's world. There are many times when an item or device must be checked without impacting its future operation, for example checking if a kiwifruit is ripe where the ripeness can be tested using a non-contact laser technique [1].

With increased focus on health and safety the applications of NDE are likely to increase. As with many areas of technology the systems used for NDE are constantly advancing. The general concept of NDE is simple: run a given test on the object in question, compare the resulting data with a given criteria, and finally produce a control decision for that object. The control decision is often about the quality or integrity of the object. Although the testing is often at the start of an object's useful lifetime, to check manufacturing, NDE tests can be used at any time during its lifetime. For example, a significant portion of the U.S. commercial and military air fleet has remained in operation well beyond their design lifetimes through retirement for cause testing [2]. This means that the structure is only retired if it is found to be defective. Although there will be a cost involved in the testing of the aircraft, the financial advantage of keeping them in the air, earning money, is obvious.

As well as the method chosen, NDE also considers the testing procedure. Instead of testing all welds on a structure, for example an aluminium frame building, the probability of a weld failing can be weighed against the consequence of that weld failing. This risk assessment can then be used to decide how often and how thoroughly the weld needs to be tested. The development of NDE on the technological as well as business front has resulted from the continual drive to provide higher productivity with tighter tolerances, while still maintaining the safety of

the part.

There are many different methods of NDE but these can be split into contact and non-contact methods. Each method has different advantage and related applications. This section will give an overview of NDE by listing different methods and some of the applications. The methods are split into contact and non-contact, with Ultrasonics and Vibrometry having both contact and non-contact methodologies.

1.4.2 Contact Methods

As the name suggests contact methods are those in which physical contact must be made with the object to test it. These tend to be the oldest methods and most easily related to our own methods of assessing objects. The first two methods of contact NDE are called Liquid Penetrant (LP) [3, 4] and Magnetic Particle (MP) [5, 6] testing. Both these methods allow small surface cracks to be more easily visible. LP uses a liquid dye which seeps into any surface defects through capillary action. After allowing some time for the dye to enter the cracks the excess is removed and a developer added to help show where the cracks are. An every day example of liquid penetrant testing is when a china tea or coffee cup gets a small crack. Over time some of the drink in question collects in the crack and it is shown up as a brown line on the china. There are a number of different dyes and developers and the correct choice can affect the sensitivity of the method. MP uses a similar approach except the object is magnetised then covered with finely divided ferromagnetic particles. Any defects at or near the surface will affect the magnetic field and attract the magnetic particles their edges. The particles are coloured (or fluorescent) for better identification of defects.

The next method considered is called Acoustic Emission [7] which uses a number of acoustic sensors across an object to detect small internal sounds produced

by some material under stress. These sounds are produced by the rapid redistribution of stress within the material, and are often the result of a permanent change in the object (for example a crack). Although the object is changed during this method of testing, the method is included in NDE because the damage is usually very small, such as a microscopic crack, compared with failure of the object. AE offers some level of internal information although the signals can be hard to analyse due to object shape and background noise. AE arguably dates back to about 6500 BC when pottery makers would listen to pots cooling after firing. With experience they learnt that pots that produced more cracking sounds were structurally defective and might fail prematurely. A method related to AE is Ultrasonics [8, 9]. This method is one of the most widely used methods of NDE today and works by probing the properties of an object using high frequency sound waves (ultrasonic waves). These waves are generated using a voltage pulse into a transducer, often coupled to the object using a gel or water. These waves interact with the geometric and mechanical properties of the object and are detected, either with the original transducer (pulse-echo) or another transducer (pitch-catch). The received signal is transformed back into a voltage by the detector for analysis. Several transmitters and receivers can be used, and can be scanned to give a 2D, or even 3D, image of the object. Ultrasonics can measure object material properties such as density and the elastic constant, the geometry of a flaw, delamination of composite layers and mechanical properties and micro-structure.

Non-contacting ultrasonic testing is one of the most recent advances in the method. This method uses very high efficiency transducers using a compressed fibre layer for acoustic impedance matching with air [10]. Applications of this non-contact method include the food, pharmaceutical and biotechnical industries.

Another non-contact utrasound method is laser ultrasound. This method uses a laser to generate the ultrasonic waves in the material and either uses a conven-

tional sensor to detect the waves, or another laser. This laser can either cause a local explosion at the point of impact, or with lower power or low-absorbing materials, localised rapid heating. These two methods of ultrasonic excitation differ in the transfer energy into the different wave type within the material, and also the associated far field directivity pattern. There are two methods of detecting the waves with a laser. One uses the change in angular defection of the beam as the wave passes under it while the other uses an optical interferometer to measure perpendicular surface motion [11]. This method can be used from a distance of the order of meters away.

The waves can also be detected through contact sensors such as electromagnetic acoustic transducers [12] which are fixed in place while the generation laser is scanned. This method has been shown to allow detection of small cracks in metal.

Recently new techniques have been developed which use piezoelectric films deposited on the test object as a receiver for laser generated ultrasound [13]. These contact sensors allow high temperature measurement, curved surface and multiple wave detection. These are intended to reduce the cost of the laser ultrasound equipment. It was demonstrated that the equipment was capable of $14.9\,\mu m$ thickness accuracy for a $12.7\,mm$ thick steel plate.

An important method of measuring how an object deforms is with a strain gauge. These measure the change in shape of a section of the object due to a static or dynamic stress. There are a number of methods of measuring strain, perhaps the best known from classrooms being the electrical strain gauge. This is a device which is bonded to the object and measures small changes in the resistance of a wire due to its dimensions changing with the object surface. They are relatively simple and can be amplified and digitised easily. Multiple gauges can be used to measure more than one direction of movement. More recently

optical strain gauges have been developed using optic fibres. The advantage of these is that they can be built into composite structures and offer local or global strain measurements. Fibre-optic strain gauges come in two basic varieties:

- Point Sensors - There are two basic types of single point strain gauges. One is based on the Fabry-Perot principle and uses two fibres held end-to-end within a small tube, with a small gap between them. When light is launched down one fibre it is reflected off the end of that fibre and the start of the other. These two light signals then interfere and so the distance between the two fibres can be measured. This method can typically measure deformation of the order of the $100\,\mu m$. Better sensitivity can be achieved by using a fibre Bragg grating. This is achieved by periodic changes in the refractive index of the fibre over a small region. Only light with the same wavelength as the grating pitch will be reflected back down the fibre. As the grating pitch is strain and temperature dependent this fibre can be used to measure both. If just one is required than a separate reference fibre Bragg grating must be used to measure temperature alone. Deformations in the micrometer range can be measured with fibre Bragg gratings

- Long Gauge Sensors - This system consists of two fibres running parallel in the direction the strain is to be measured. One is free while the other is attached to the structure at both ends and prestressed. An optical interferometer is used to measure the length difference between the two fibres with a resolution in the micrometer range.

There are also distributed fibre sensors which use Brillouin- and Raman-scattering to measure temperature or strain profiles along a fibre.

Accelerometers[4, 14] are another popular device for measuring the response of a single point of an object. They are electronic devices which convert accel-

eration into a signal which can be digitally measured. These are normally use either piezoelectric crystals to generate small amounts of charge or through capacitive effects of micro scale structures on silicon chips. The latter are called Microelectromechanical Systems (MEMS) and are miniaturized mechanical and electromechanical elements which can be fabricated out of silicon using similar batch techniques to integrated circuits. These accelerometers are much smaller and more power efficient [15]. Accelerometers typically operate below their internal resonant frequency of about $30kHz$. They also have a good dynamic range with the complete system giving up to $120\,dB$ with a frequency response in the region of $10 \rightarrow 20\,kHz$ [14].

A variation on the accelerometer principle is the velocity sensor which consists of a coil and permanent magnet. The magnet moves with the object and so generates a voltage across the coil. Similarly to accelerometers they need to be mounted on the object and can only measure motion in one direction. Multiple velocity sensors or accelerometers can be used to gain information on a number of different directions. They are also susceptible to magnetic fields and have a relatively low maximum frequency. A velocity transducer might have a dynamic range of about $60\,dB$ over the range of $10 \rightarrow 1 - 2\,kHz$ [14].

1.4.3 Non-contact Methods

As the title suggests this section considers those methods of NDE which do not require contact with the object. It will introduce a number of non-contact methods and give examples of their use as an introduction to the specific branch of NDE of interest to this project. There are a number of reasons why non-contact methods may be required. In some situations contact may change the response of the object, for example a thin membrane [16]. These methods can also be used on moving parts like rotating shafts or milling tools[17]. Non-contact can be good

for hostile environments which might damage or interfere with the equipment, for example jet engine exhausts. Another advantage of some non-contact methods is that they do not require that time be spent attaching a sensor to the object. Eddy Current [18] (EC) testing is an example of this and consist of two coils: one, called the excitation coil, is connected to an oscillating voltage source and the other, called the pickup coil, is connected to a voltmeter. The pickup coil detects the magnetic field of the excitation coil and produces a voltage measure by the voltmeter. When a conductive material is brought close to the coils, continuous, circular currents are induced in it. These are called eddy currents. They in turn produce a magnetic field which opposes the field by the excitation coil. Any defects in the material, or changes in its conductivity will change the eddy currents and therefore the magnitude of the opposing magnetic field. The pickup coil then measures the decrease in total magnetic field. EC is used extensively in testing of rails where the sensor is mounted on a moving vehicle and can test the tracks at a rate of between 75 and 100km/h [19]. Another similar approach uses proximity probes to measure the distance to a surface. These are based on capacitive or magnetic properties, with the most popular using the change in the inductance of the system brought about by a change in the dimension of the gap. As such the surface must be electrically conducting and the gap must have a high dielectric value. The probe uses a high frequency signal which is amplitude modulated by the gap, and as such amplitude demodulation techniques are used [14]. This method has a dynamic range of between $20\,dB$ and $40\,dB$, which is limited by the linearity of the system. This method finds use in measuring vibration of rotating machines, such as shafts, and is also affected by out-of-roundness and residual magnetic properties of the shaft in question.

One of the most popular approaches to non-contact NDE is to use light, or more generally electromagnetic radiation. This forms the basis for the remaining

method of NDE considered here. The first considered uses the radiation produced by objects due to their temperature called Thermography [20, 21]. There are two types of thermography, active and passive [21]. Passive thermography is used in predictive maintenance where operating machinery and electronics is imaged to find hot areas resulting from operating problems. Active thermography uses an external heat source and thermal imager to assess subsurface structure within an object. There are a number of different types of active thermography but all have four major components. They work by external heating of an objects surface and this heat is then transported into the object. The resulting thermal distribution is imaged with an infrared camera and through the spatial and temporal distribution of temperature across the object information such as material and structure can be found.

There are other methods of measuring the temperature of an object including beam deflection and temperature sensitive coatings, but only infra-red cameras offer full-field, high speed and high sensitivity imaging. Figure 1 shows three images of a set of holes in a graphite-epoxy composite, showing how the heat distribution changes with time.

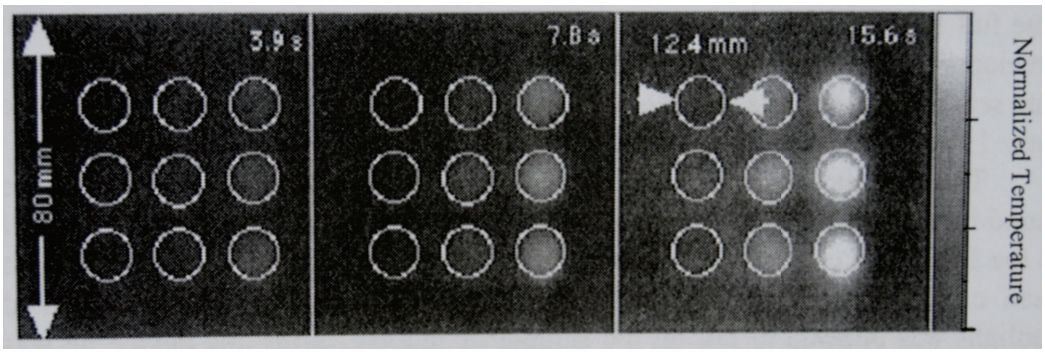

Figure 1: Infrared image of flat bottomed holes in graphite-epoxy composite at different heating times. The holes are 12.4 mm in diameter at differing depths from 0.8 mm to 2.8 mm. The holes are not visible from the top and are marked with the white circles from [20].

Although this is not as widely used as other NDE techniques, microwave [22]

NDE is used for its non-contact nature and its ability to penetrate dielectrics. Microwaves are electromagnetic radiation with wavelengths between about $100cm$ and $1mm$ with the wavelength chosen for the particular object or measurement property. Either the scattered or transmitted waves are detected after interaction with the object and can detect such defects as water content in dielectric materials, delamination in rocket casings, defects in rocket propellants and detecting inclusions and porosity in ceramics and molded rubber. The final category of non-contact NDE considered is Radiology [23]. There are a number of different techniques for gathering information about the internal structure of objects which fall into this category. They are all concerned with the use of high-energy radiation and particles to gather information about the internal structure of objects. All radiation has its source in the atomic structure of the atom, either from the electrons orbiting the atom changing to more stable energy states or from interaction between subatomic particles. X-ray and Gamma ray techniques use high-energy radiation with the difference between the two being the source of the radiation. Those using particles are classified by the type of particle, these being α, β, neutrons and protons.

X-rays are one of the most common methods of NDE. They are absorbed by heavy, dense materials but not by lighter substances. It is one of the few methods that can examine the interior of an object and one of the few methods that can image any object.

Neutron methods use neutrons which are uncharged subatomic particles with a mass slightly larger than a proton. They are attenuated differently from X-rays in that they are absorbed by lighter rare earth elements especially hydrogen. This allows neutrons to image the hydrocarbon foams in petrochemical reactors with minimal absorption from the surrounding steel vessel. They are also used for materials, such as nuclear weapons, which have a high opacity to X-rays.

Although not as common, β particles and protons are also used for NDE. β particles are energetic electrons or positrons. They have a continuous energy spectrum and slow down gradually in materials through multiple collisions. They are therefore well suited for area density measurements with the object to be imaged placed between the source and the detector.

Protons are basically hydrogen nuclei generated by an accelerator and as such interact with the electrons within a substance. The rate the proton slows within a material is dependent only on the electron density along that path rather than the atomic number. Relatively fewer protons are required for accurate information about the object than X-rays, however imaging depths are quite small and the method works with a single thin beam which is scanned across the object.

1.4.4 Conclusion

This section has given an overview of some of the different methods of NDE, and these are summarised in Table 1. All the methods measure a given parameter or parameters of the test object and use them to make a decision about the health of that object. This shows the great variety of NDE methods available, however it has omitted optical methods of NDE as the work of this thesis falls into this category and so they will be considered in the next section.

1.5 Non-Contact, Optical Vibration Analysis

This thesis is interested in measuring vibration across a whole surface using a non-contact methodology. There are a range of optical NDE methods which fulfill these requirement and this section will introduce those relevant to this project. Optical NDE methods are those which use electromagnetic radiation in the optical range of the spectrum to probe the properties of an object. They can provide very precise, high spatial density measurements across the surface of the object

Table 1: Summary of Different Methods of Non-destructive Testing

NDE Method	Measurement Distance	Types of defect	Advantages	Disadvantages
Liquid Penetrant	contact	surface cracks	Reliable, economical, can be automated	Surface must be clean, only surface cracks
Magnetic Particle	contact	surface/ near surface cracks	Reliable, economical, can detect through coatings	Only ferromagnetic objects, large currents required, may require demagnetisation
Acoustic Emission	contact	internal defects	Large volume tested in situe, predictive maintenance, tested while in operation	Susceptible to noise, interpretation can be complex, only defects which produce emission found
Accelerometers	contact	movement	Cheap, reliable, can measure changes in motion/ vibration of object, mature technology	Point measurement, must be attached to the object, can load the test object
Strain Gauge	contact	change in shape	Cheap, reliable, can measure changes in strain/ shape of object, mature technology	Point measurement, must be attached to the object, can load the test object
Ultrasonics	contact/non-contact	internal defects	Low health and environmental risk, wide range of geometries and materials, mature technology	Experienced technician, mainly contact method, only flaws parallel to wave travel, point based
Eddy Current	non-contact	surface/ near-surface	Any conductive material, fast, no surface preparation, portable, many parameters	Only conductive materials, interpretations of results complex, only parameters which affect eddy currents
Thermography	non-contact	internal defects	Full-field, video frame rate method, can detect cracks, voids and delaminations	Only properties which affect thermal characteristics, may need high power heat source
Microwave	non-contact	internal or surface	Can be scanned over larger areas, can be low power, low environmental impact	Internal - dielectric materials, surface-conducting materials
Radiology	non-contact	internal	Works on almost all materials, fast area inspection, different parameters with different methods	Health risks from radiation and high voltages, expensive, only open cracks/ voids
Proximity Probes	non-contact	movement	Cheap, robust, does not require contact with the object	Low dynamic range, only electrically conductive parts

[24]. In addition to this they normally operate without needing contact with the object by shining light onto its surface and then analysing the reflected light. The optical methods of importance to this project are those which are used to measure vibration. These measure the motion of the surface and can be used to gather information of the material properties of the object as well as damage or inconsistency.

Optical NDE techniques can be split into single point measurements and full-field techniques. Full-field techniques operate like a conventional camera (often using one), and image an area of the object concurrently, or via rapid electronic scanning. Single point measurement on the other hand can only image one point on the object at a time. This point can then be moved around to obtain information on an area of the object. Before considering the techniques themselves two important optical effects will be introduced: the Doppler effect and the speckle effect.

1.5.1 Advantages of Optical Vibration Analysis

Accelerometers can be used to measure the vibration of an object's surface but the additional mass of the accelerometer may affect the vibration of the object and so these are best suited to massive objects such as machines and engines where the added mass of the accelerometer has negligible effect on the vibration. Proximity probes can also measure in a non-contact way, but they are limited to conductive surfaces and have limited spatial resolution. Optical vibration analysis offers a solution to this as it is non-contact and can give good spatial resolution across the surface. There are a number of applications where non-contact vibration analysis is required:

Model Validation This application seeks to determine the degree to which an updated mathematical model provides an accurate representation of the behavior of the structure being modeled [24]. This is an important stage in the development of a finite element model of an object. The first stage is to design a model and appropriate validation tests. The second stage is to compare the model with sampled values to check the correlation between them and optimise the model.

Membrane Vibration An important application of vibration analysis is for the vibration of membranes. Membranes and membrane analysis is becoming more and more common in many areas of engineering and science. In paper manufacturing it is important to reduce the out-of-plane vibration of the paper as these reduce the operation speed of the paper machines[25]. Recent interest in membrane structures for use in space has been motivated by the potential for reduced launch mass and stowed volume [26].

Historically it was difficult to measure the time-history of a membrane's vibration as the thin and light nature of it meant that attaching any probe would change the vibration characteristics. With the advances in optical methods of measuring vibrating surfaces it is now possible to image the membrane at appropriate frame rates and over its whole surface.

Damage Detection One of the most common application of vibration analysis is to detecting damage to an object. This is achieved by vibrating the object and measuring how it responds to the vibration. Damaged areas of the object exhibit an altered response to those vibrations and so can be identified. Measuring the vibration of a milling tool during operation can give information on misalignment and out-of-roundness [17] of the tools.

1.5.2 The Doppler Effect

Before continuing, two physical effects which are important in a number of optical NDE methods will be described. They are the Doppler effect and the speckle effect. The Doppler effect is defined as the apparent change in time between received events due to the relative velocity of the transmitter and receiver, and the finite velocity of transmission of the events [27]. An everyday example is how the sound of a car passing at speed sounds higher pitched as it approaches (the speed of the car Doppler shifts the frequency of the noise it produces) and then lower pitched as it recedes (the Doppler effect now shifting the frequency of the cars noise lower as it is now moving away from the observer). The first application of the Doppler effect was in astronomy, and before Doppler and Fizeau published their work on the subject. In 1676 when Romer used the apparent variation of the periods of Jupiter's satellites to deduce the velocity of light he was, in effect, using a Doppler method. Doppler himself first proposed his principle in 1842 in relation first to sound and then to light when trying to explain the difference in colour of stars. Fizeau then pointed out that this was not the correct explanation and he was the first to suggest the use of sharp spectral lines to observe Doppler shift. Since the second world war the applications of the Doppler effect have been steadily growing. It is now used in navigation, missile and satellite tracking, radar, etc. It has more recently been used to monitor small sections of objects, not moving as a whole, but vibrating.

1.5.3 The Speckle Effect

The speckle effect occurs when a rough surface is illuminated by coherent light [28, 29]. By rough, we mean that the height of the variations is comparable to the wavelength of the laser light. Figure 2 shows an example speckle pattern. As can be seen, the pattern is a composed of bright and dark spots in a random

configuration.

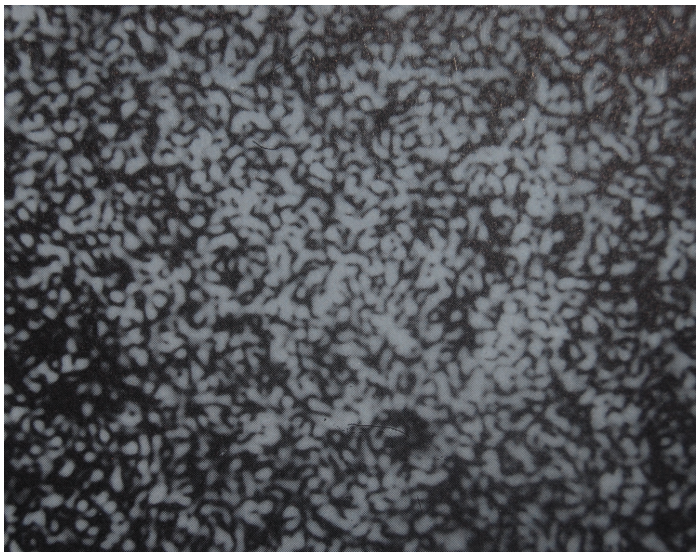

Figure 2: Magnified Laser speckle image, from [30].

The different levels of the surface each introduce a phase shift in the light reflected from them and these different wave fronts interfere at the detector. If the random phase shifts result in constructive interference then there will be a bright area on the detector. If they interfere destructively there will be a dark area.

The distribution of speckles is random and defined by a Gaussian probability distribution function provided that there are a large number of scatterers, as is the case for most real surfaces. This also requires that the surface variations are great enough that the phase differences of the different point sources is equally distributed over the whole 2π range. Effectively no value is more likely than any other.

If the detector images the light from the object directly (without any lenses) then the speckles are called Objective Speckle and the size of each speckle is dependent on the size of the illuminated area, the distance from the object to the detector and the wavelength of the laser light. This is shown in Figure 3 and in

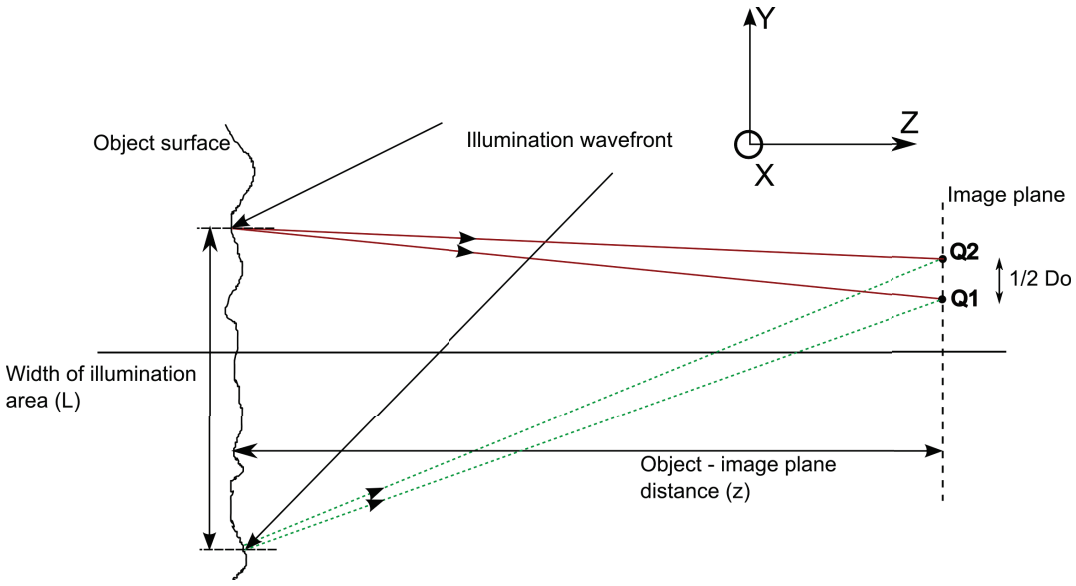

Figure 3: Creation of an objective speckle pattern. An optically rough surface is illuminated by coherent light and the objective speckle pattern is viewer at an image plane.

this case the average diameter (D_O) of the speckles is given by equation 1 where z is the distance between the object and viewing plane, L is the width of the illumination area and λ is wavelength of the laser light [31].

$$D_O = \frac{z}{L}\lambda \tag{1}$$

Most NDE systems where speckle is present use an optical imaging system of some sort. When an optical system is used to image the reflected light the resulting speckle pattern is called a Subjective Speckle Pattern and the size of the speckles is set by the lensing system. This is because the lensing system has a finite resolution limit and so a given speckle is formed from light reflected from an area of the object. Considering the setup in Figure 4, the light from point P1 forms a diffraction pattern on the image plane centered at Q1. This light will have a random phase due to the height of the point P1. The light arriving at Q1 will also have contributions from the points around P1. P2 is a point far enough from

P1 that its diffraction pattern, centered at Q2) on the image plane has its first minimum at Q1. This means that the light from P2 makes no contribution to the illumination of Q1, regardless of its phase. Points further from P1 than P2 will make only a small contribution to the light at Q1 since the secondary maxima of a diffraction pattern are much smaller than the primary. Therefore the intensity of the the light at Q1 is made up from contributions from an area of the object twice the distance from P1 to P2 in size (D_O).

In this situation the diameter of the speckles (D_S) is given by 2, where v is the distance from the lens to the image plane, a is the diameter of the lens and λ is the wavelength of the laser light [31].

$$D_S = 2.4\frac{v}{a}\lambda \tag{2}$$

Using simple geometry the size of the area on the object which produces this speckle is D_O and will be given by $2.4\frac{u}{a}\lambda$ where u is the distance from the object to the lens. This is the same as the resolution limit of the system.

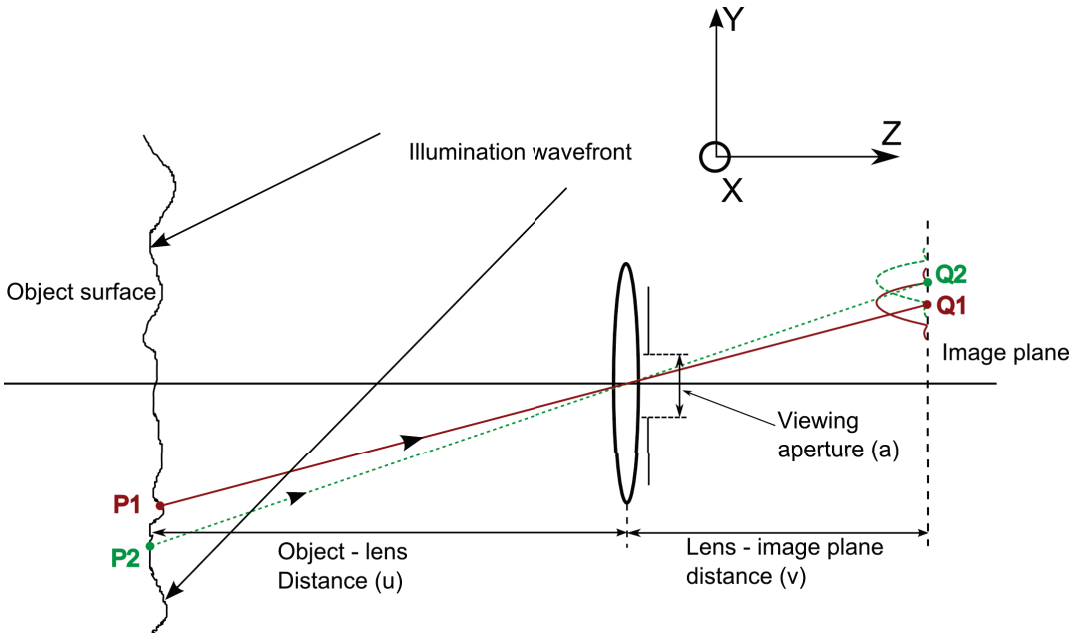

Figure 4: Creation of a subjective speckle pattern. An optically rough surface is illuminated by coherent light. Point P2 is far enough from P1 that the light from its diffraction pattern (Q2) makes no contribution to Q1.

Speckles are often thought of as 2-dimensional spots on the image plane, but they are in fact 3-dimensional volumes of illumination. Leushacke and Manfred [32] have derived the three-dimensional correlation coefficients of a speckle pattern and have shown that the depth of the speckle ($D_{S\,depth}$) is given by Equation 3.

$$D_{S\,depth} = 8\lambda \left(\frac{v}{a}\right)^2 \tag{3}$$

They also showed that the shape of the speckle in the transverse direction is "cigar shaped", and that they point in the direction of the optical axis.

1.5.4 Moiré Analysis

Moiré is an optical effect created by the superimposition of two repetitive structures. There are many examples of this is everyday life such as when two fabrics are overlaid or a pin stripped shirt is viewed on television. The pixels of the

television forming the second structure. If the structures are slightly different, when they are overlapped some of the gaps in one structure will be covered by the other, resulting in a pattern of dark and light regions called a Moiré Pattern. The structures used are typically grids of lines. Figure 5 shows four simulated examples of Moiré showing the patterns that result from rotation and change in spacing of one of the grids (strain).

The main property of a Moiré pattern which makes it useful for NDE is that a relatively small change in the underlying grids results in a large change in the Moiré pattern, effectively magnifying the movement. Another way of thinking about this is from a spatial frequency point of view. The grids being two high frequency signals which, through superposition, are mixed down to base-band. Any different in frequency will then become a spatial beat signal.

Geometric Moiré [28, 33, 34] is a form of Moiré pattern which does not rely on optical interference and in its basic form it is the least sensitive type of Moiré. Its sensitivity is set by the spatial frequency of the grids used. It can measure in-plane and out-of-plane components of displacement, and do so across the whole field simultaneously. Another set of Moiré methods are grouped under Moiré Interferometry. These methods use the same relationships as geometric Moiré, and in fact all Moiré methods can be shown to be cases of optical interference [33].

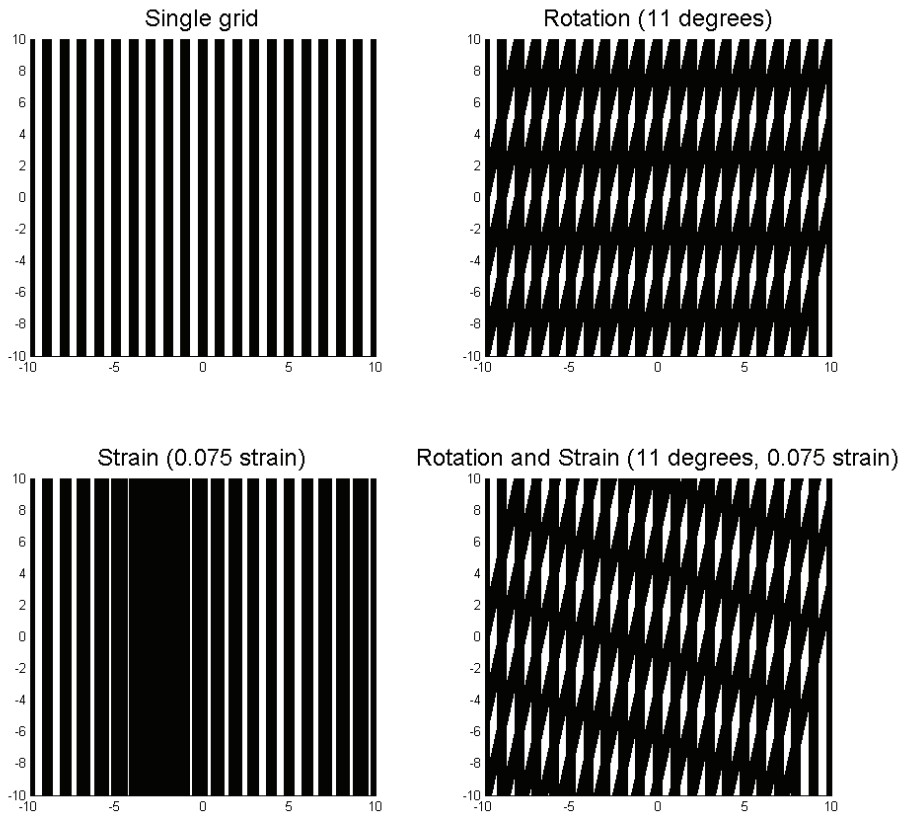

Figure 5: Examples of geometric Moiré Patterns, generated using Matlab. Top left: single grid. Top right: second grid rotated by 11 degrees. Bottom left: second grid spacing changed by 7.5%. Bottom right: combination of rotation and change in spacing.

The methods of applying Moiré to in-plane motion are also greatly varied and are summarised below:

- Direct superposition of a master grid over a grid attached to the object. The recording device need only be able to resolve the Moiré pattern, rather than the actual grid lines which are typically much finer.
- A single grid is attached to the object and two photographic exposures are taken, one before stress or movement and the other after. This method has the advantage that optical processing can be applied to increase sensitivity.

The photograph must have sufficient resolution to record the grid lines.

- If an exposure is taken of the grid as for the previous method, but this exposure is developed and then replaced. A second exposure can then be taken of the first exposure superimposed over the fringes. This also requires that the exposure be able to record the grid lines but allows a high degree of improvement in sensitivity.

- Finally a master grid can be optically superimposed onto the image of the object. This allows real-time imaging of the fringes but optical distortions become a problem as well as less control over sensitivity.

There are similarly a number of Moiré methods which apply to out-of-plane motion:

- The first way places the grid between the object and camera, using off-axis light to cast the shadow of the grating onto the object, forming the second grating which when viewed through the grating together form the moire pattern.

- A variation on this uses the surface of the object to reflect an image of the grating back to the camera. The surface of the object needs to be polished, to provide the reflection.

- The last projects the shadow of the grid onto the object and takes two exposures of the shadows, one before the loading and one after, and then superimposes the images afterwards to produce the Moiré pattern.

The highest sensitivity method of Moiré Analysis is called Moiré Interferometry [30, 33, 34] and can measure movements of the order of one wavelength ($< 1\,um$). The advantage over the mid sensitive Moiré systems is that it does not need to optically image the gratings. The basic operation of a Moiré Interferometer is

to coat the object with a phase grating (which alters the phase rather than the amplitude of the reflected light) and then illuminate it with the fringe pattern created by two slightly off-normal, coherent beams of light.

The grid is illuminated by two beams of coherent light which strike the grating at a slight angle from the surface normal, they interfere to produce a set of fringes which act as the second grating. It utilises the property of a grating that it can refract incident light to reflect back part of the illuminating beams. The angle of the refracted beams is controlled by the grating spacing. The setup can be configured such that the diffracted beams are collinear (the angle of diffraction matches the angle of the incident beams).

If the gating attached to the object changes spacing this will change the spatial frequency and so the angle of the diffracted light. The diffracted beams will no longer be collinear and so create a fringe pattern. The theoretical limit of this method is a measurement of $\lambda/2$ per fringe order[33] where λ is the wavelength of the illuminating light.

Using Moiré Interferometry it is possible to detect full-field strain in on very small objects, even in harsh environments. An example is the interfaces of micro-electronic devices after exposure to temperatures of 100'C [35]. Digital Moiré is a new approach to deformation analysis [36]. The second grating being formed either by subsampling the imaging system or within software. The use of computer generated gratings also adds to the versatility of this branch of Moiré analysis. These gratings can be projected onto the object or used in the form of a shadow mask, and allow contouring through computer controlled phase shifting of the grating. As the grating is computer controlled it can be easily altered depending on the application.

Moiré has been applied to non-contact vibration measurement by projecting lines onto the surface and demonstrated to be able to measure vibration ampli-

tudes of $50\,\mu m \rightarrow 10\,mm$ at $1\,kHz$ [37]. This system used two optic fibres to deliver laser light to the surface of the object. This light interfered to form the projection grating. Moiré has also been applied to rotating objects with amplitudes of $75\,\mu m \rightarrow 500\,\mu m$ [38].

1.5.5 Holography

This method was invented by Dennis Gabor in 1948 as a method for improving the usefulness of electron microscopy images. Due to the restrictions of the technology at the time not much was made of his work until the invention of the laser in 1960. Off-axis holography was then developed by Leith and Upatneiks in the early 1960's which solved the previous problems with holography [28].

The idea behind holography is to capture both the amplitude and phase of the light reflected from an object. No current capturing method can directly capture optical phase so a reference beam is interfered with the light from the object thus converting the phase data into amplitude data.

Considering each point on the object to be a source of plane waves, these will all be traveling in slightly different directions. If these waves are incident on a photographic film at the same time as a reference beam (plane wave) they will interfere and producing an individual grating. For a complex object this will result in a complex set of gratings. Once developed, shining light through the film in the direction of the reference beam results in reconstruction of the original wavefronts and therefore provides a theoretically perfect image of the object.

Holographic Interferometry [39] Holographic Interferometry (HI) is the branch of NDE which uses the property of a hologram that it is a sufficiently good image of the object that it can be compared interferometrically with another hologram or the object. This has been applied to fields of vibrometry and deformation

analysis and can give a quick indication of the location of stress concentrations.

HI can be used as a real-time technique by developing a holographic image of the undistorted object then replacing it in the experimental setup. For the observer this gives the effect of a holographic image superimposed on the object. When the object is deformed the interference of the object's reflected light and the hologram produce correlation fringes which show the degree of relative movement between the original object and the current object. Due to the sensitivity of HI the exposed holographic image must be placed precisely in the experimental setup.

Another method is called Double Exposure HI and works by exposing the photographic negative twice, once for the undistorted object and once for the distorted object. This gives similar results to a single real-time image but does not require the removal, developing and replacing of the holographic image. For periodically vibrating objects the exposure can be increased as is done in time-averaged HI. This shows where there are nodes and gives an indication of the amplitude of vibrations. This is because the surface of the object spends most of its time in the maxima and minima of the vibration and this essentially gives two exposures of the averages at these points.

Advances have been made in digital holography. In the work of Schnars [40] the object and reference beam was set-up in such a way that the interference fringes of the hologram were resolvable on a CCD. The stored holograms were then processed to extract the deformation field of the imaged surface. This method suffer from the relatively low resolving capabilities of the CCD.

1.5.6 Laser Speckle Methods

With the invention of the laser it was hoped that this source of high-intensity, collimated, monochromatic, coherent light would provide the answer to a great many illumination problems. Unfortunately it was found that images of objects

illuminated by a laser showed a grainy pattern of light and dark areas, known as a laser speckle pattern. This effect soon formed the basis for a number of techniques loosely grouped under speckle methods. These methods offer full-field, non-contact analysis of objects and are widely used in NDE. This section will start by providing some background on the speckle effect and then look at these different methods in more detail.

Speckle Pattern Photography [28] Speckle Pattern Photography is a method for measuring surface displacement or deformation. The idea is to use individual speckles on the surface of the object to map how it moves. It was shown in Section 1.5.3 that a single subjective speckle is the result of light reflected from a region of the object, set by the geometry of the illumination and viewing optical systems. If that region moves (through translation or object deformation - within certain limits and keeping the optical system the same) then the speckle moves too. The speckle is also not changed by the movement.

Taking two images of the speckle pattern, one before and the other after the deformation allows the local displacement of the surface to be measured through movement of the relevant speckles.

During its initial development, this method had several different names given by concurrent but independent researchers. These include speckle interferometry and speckle moire interferometry. Stetson [41] offered a simple solution to the naming problem and this also shows the distinction between speckle photography and speckle pattern interferometry. If the method uses incoherent recording of two speckle fields, in other words a reference beam is not used, then it is speckle photography and is basically a moire method. If a reference beam is present then it is speckle interferometry. Even if the fields are still being recorded incoherently the reference beam changes the speckle pattern.

Two speckle photographs are taken on the same photographic plate resulting in a complex diffraction grating as adjacent speckles overlap with each other. The light distribution of the speckles is a sinc function due to the diffraction effects of the lens system and this means that each speckle has the basic form of a central peak (the speckle) with the brightness falling to zero before oscillating in a dampened way as one moves further from the speckle. These oscillations of speckle pairs overlap to form a grating and for smoothly varying displacements the appearance is that of a poor moire grating.

Extraction of the displacement information from the developed photograph could be done through comparing the displacement of individual speckles under a microscope, but this is somewhat impractical. There are two other methods of analysing the photographs.

- The simplest is to shine a slender laser beam through the developed plate and observe the light diffracted by the grating structure of the speckles. This diffraction gives the direction and amplitude of smooth displacements from a small region on the photograph. As has already been stated this can be mapped to a corresponding region on the object. This is a point by point method.

- The second is a full-field method where the developed plate is imaged onto a camera using a special dual lens optical configuration with an optical spatial filter. The spatial filter takes the form of an opaque card with a small hole in it. This hole can be moved to select a given direction of displacement to be chosen as well as the sensitivity of the system. The camera then images the speckle pattern, superimposed on the image of the object, with diffraction fringes over the whole image which indicate the amplitude of the displacement in the chosen direction.

Speckle Correlation Interferometry [28] This form of speckle metrology is an interferometric approach which uses the phase information in the speckle field to measure object displacement. There are a number of techniques and they are typically grouped under Speckle Correlation Interferometry or Speckle Interferometry (SI). They all use a coherent reference source which mixes with the speckle pattern at the detector in an interferometric manner. The reference can either be a beam or another speckle pattern and is used to modify the speckle pattern interferometrically before it is recorded. The recorded speckles for two states of the test object are then superimposed and interfered to create fringes which show the displacement the surface of the object had undergone.

Consider Figure 6 which shows a configuration similar to speckle pattern photography. The object is illuminated with coherent light, but no reference is used. The speckle pattern is formed when the light reflected from the object is coherently interfered on the sensor. However by moving the image plane it is possible to map out the depth (z-axis) of the speckles. The speckles are elongated in z-axis direction resulting in little change for out-of-plane movement of the object. The speckles are not visible until they strike the sensor and then each speckle is created by the constructive interference of the light from the object. The absolute phase of the light in the speckle is not important, or recoverable after imaging.

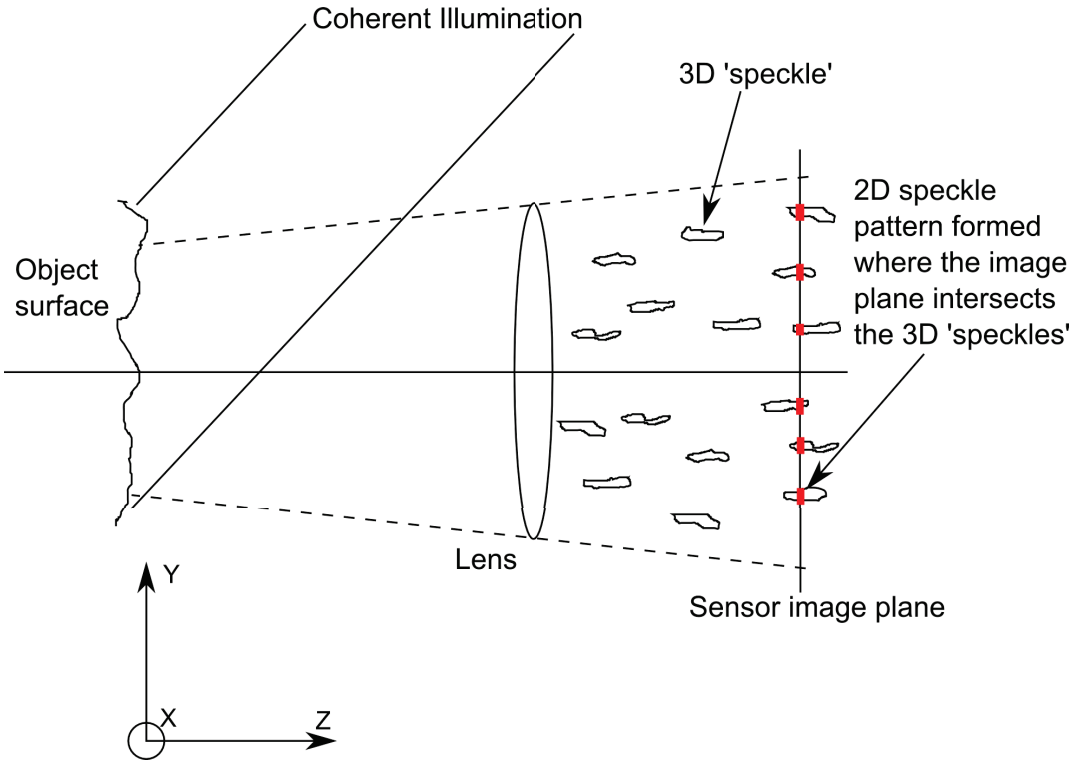

Figure 6: Three-dimensional conceptualization of speckles

If a reference is used then it changes the speckle pattern allowing some of the phase information to be extracted. This reference can either directly illuminate the sensor, as shown in Figure 7, or it can illuminate the object. Considering Figure 7, the reference is a beam incident with angle 'A' to the z-axis. The value of 'A' is not actually important when the reference directly illuminates on the sensor as its relative phase does not change with object motion. Assuming that 'A' is zero, as is the case for most setups, the result is that the three-dimensional speckle field is dependent on the phase difference between reference and object. This makes the speckles short and periodic as opposed to long and random as is the case for speckle photography. In this case, if the average phase of a speckle changes (due to motion of the object) by half a wavelength then it will change from light to dark.

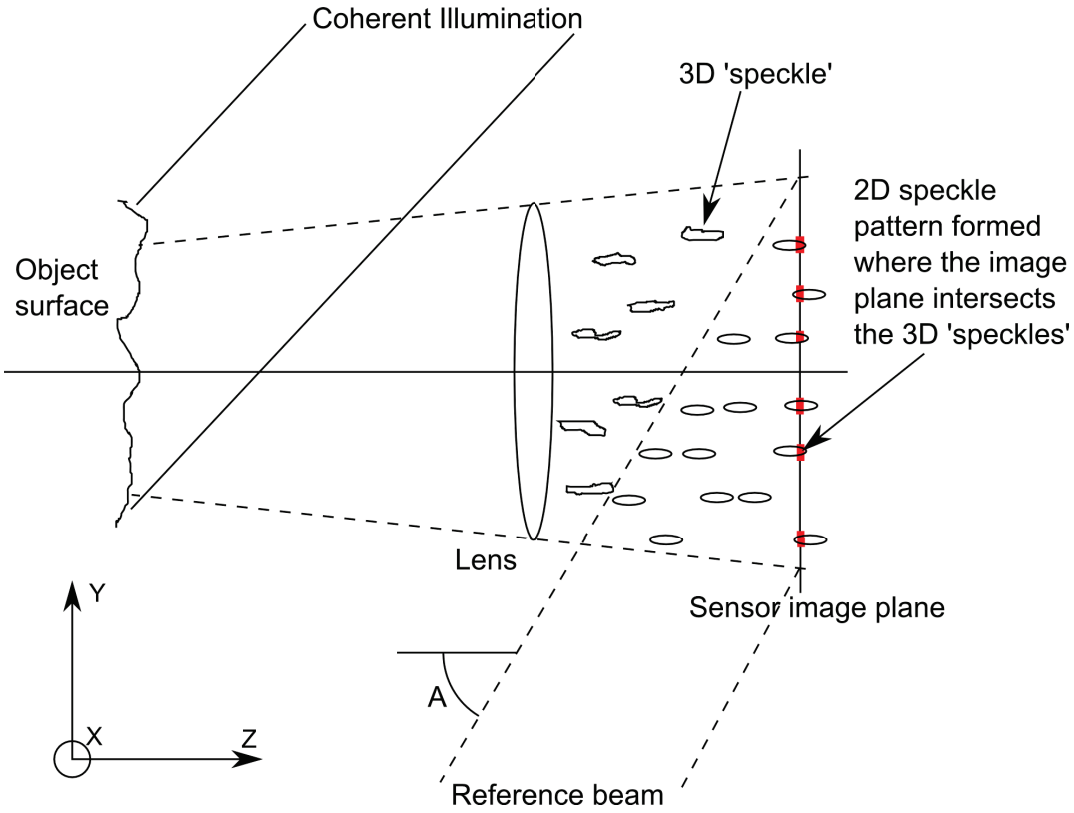

Figure 7: Three-dimensional conceptualization of speckles when a reference beam is present.

This makes SI very sensitive to out-of-plane motion but less so to in-plane movement.

Correlation In order to generate the fringes indicating displacement between two images in SI they need to be correlated. For both in and out-of-plane movement the requirement for correlation is set by two criteria:

1. If the phase change across the element on the object is greater or equal to 2π the corresponding speckle will become decorrelated in phase. This is called phase decorrelation.

2. When the speckle moves by its own diameter then it cannot be compared and so becomes decorrelated. This is also the case when the pixel moves out

of the imaging element (pixel in the case of electronic imaging and resolution cell or grain in photography). This is called memory loss. The speckle size is often set to be the same size as the imaging element.

As a rule of thumb, if the phase change or speckle movement is less than one tenth of the values set by 1 or 2 then the speckle is not decorrelated. Roughly speaking the criteria above define the absolute maximum movement that SI can work with, but they set the minimum which is needed for Speckle Photography as for this the speckles must have a small distance between them. Speckle photography also works best when the speckles are small, allowing smaller movements to separate the speckles, while SI works best with large speckles as they require more movement to decorrelated.

For out-of-plane movement memory loss is found to be the limiting factor for decorrelation but typical systems have a large capability for correlation or memory effect in the z-direction. The ability of the system to track movement is also best in the center of the field of view. With in-plane movement memory loss is again the limiting factor and is the same across the whole field of view.

Measuring Out-of-plane Displacement Two configurations for out-of-plane displacement are shown in Figure 8. With the red one the object is illuminated with a beam at angle 'B' and the reflected light is imaged on the sensor. The reference beam is focused coaxially through the use of a beam splitter. The purpose of the reference beam is to provide a known phase to mix with the speckles so that changes in the longitudinal position of the surface (which change the relative phase of the reflected light) of the object can be measured [28]. If the object is moved as a whole then the phase change of each speckle will be the same. If the object distorts then some speckles will move more than others.

For the green setup the reference beam illuminates the surface of the object

with an angle to the normal of 'C'. When 'B' and 'C' are almost the same this configuration has a sensitivity of several millimeters per fringe, although the maximum displacement is still controlled by the decorrelation effects mentioned earlier.

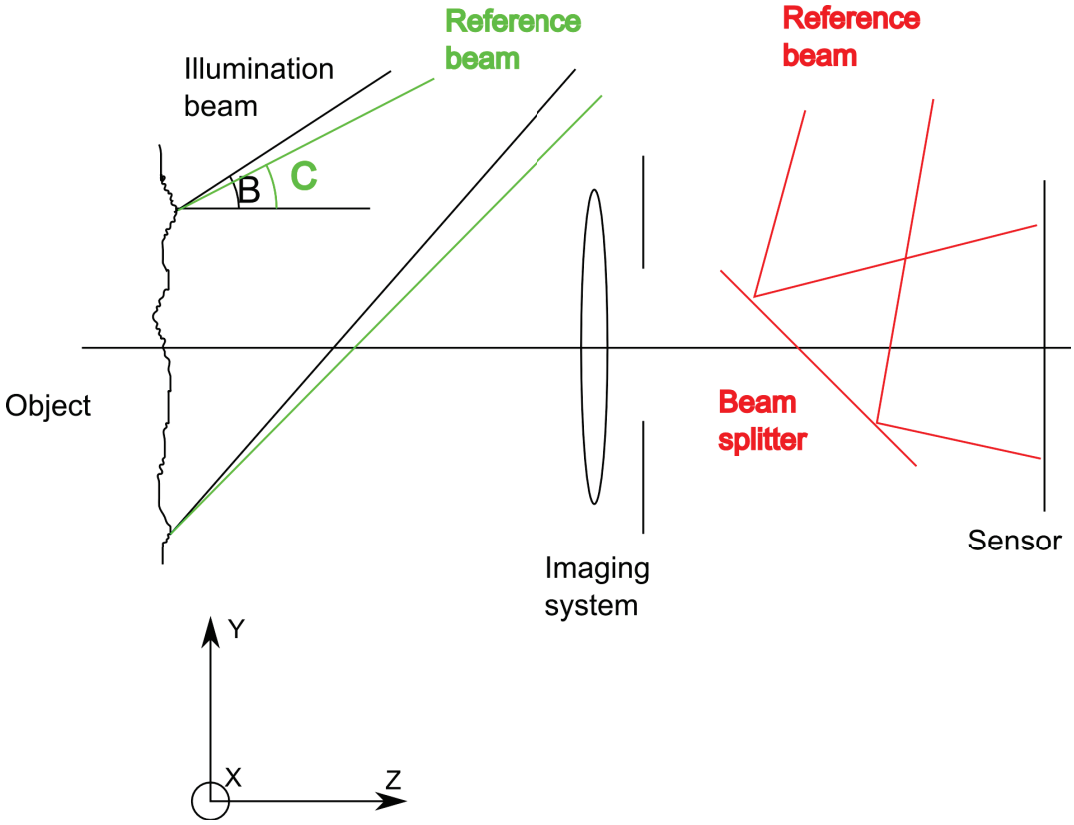

Figure 8: Experimental setup for out-of-plane measurement.

Measuring In-plane displacement The configuration for in-plane measurements is shown in Figure 9. In this configuration the relative phase of the illuminating beams is constant along planes parallel to the surface of the object so if the object moves in the z-direction there is no change in the phase of the reflected light. This means that only displacement in the xy-plane is measured.

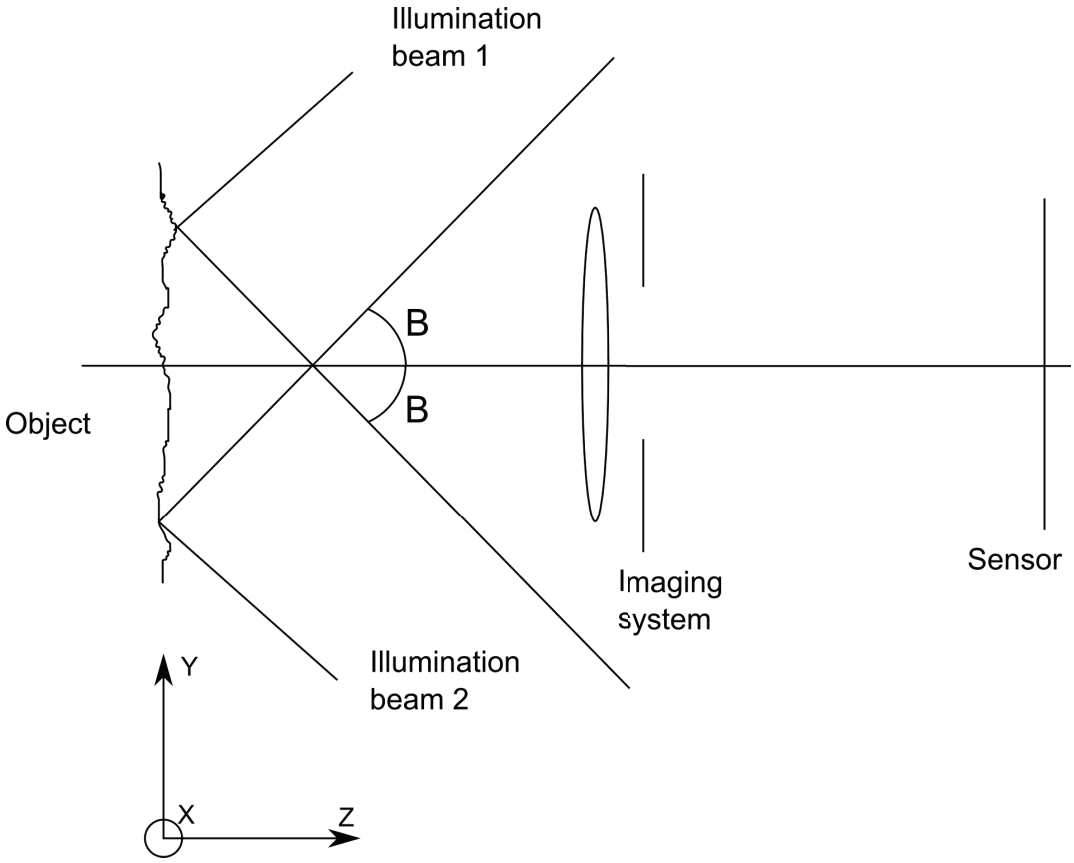

Figure 9: Experimental setup for in-plane measurement.

The processing of the images follows that of holography. However if the speckles are large enough to image using digital methods then this makes the system much easier and forms a branch of optical NDE called Electronic Speckle Pattern Interferometry (ESPI).

Electronic Speckle Pattern Interferometry This form of speckle pattern imagery has become popular for NDE measurements. It is based around imaging the speckle pattern with a digital camera, such as a CCD, and processing the images such that fringes show how the object surface changes. These systems can measure in-plane displacement, out-of-plane displacement as has already been discussed previously. The experimental setup can be changed by shifting the image slightly and this gives out-of-plane displacement gradient. This is mainly

used for testing of materials like composites to observe crack initiation and impact damage [24].

The methods so far discussed are basically static methods as they are two images, with a change occurring between them. Dynamic measurements are possible using time-averaged ESPI [42], stroboscopic ESPI [43] and pulsed ESPI [44]. These methods are basically aimed at cyclical signals, but can be used for transient signals with an appropriate trigger.

Time-averaged ESPI This method records a single image with an exposure long enough to capture many cycles. The interference fringes are averaged over time, losing any phase information. As such only the amplitude of any vibrations is captured.

Stroboscopic ESPI This is similar to time-averaged ESPI, but uses double exposures, synchronised with the vibration. These exposures can be moved to gain phase information as well as amplitude information.

Pulsed ESPI This slight variation on Stroboscopic ESPI used a pulsed source to provide a sequence of images with short periods between them. This can capture transient events. The pulses are normally in the region of $1\,us$ to $1000\,us$ and the light source is commonly a ruby laser.

Finally, it is possible to gain three-dimensional motion measurements by using three separate cameras or three separate illumination sources.

1.5.7 Laser Doppler Vibrometry

The final method of non-contact optical vibrometry is Laser Doppler Vibrometry (LDV). This is an interferometric method using laser illumination and based on the change in the optical frequency through the Doppler effect. The Doppler shift of the laser optical frequency is due to its motion of the surface of the object and

so allows that motion to be measured. This section will give a brief historical look at the use of the Doppler effect in laser metrology starting with the first uses of the effect with lasers after their invention some sixty years ago. The first use was in the form of a flow meter however it was soon applied to moving structures and laser Doppler vibrometry was born.

Summary of LDV systems

Single point The most basic LDV system uses a single laser beam to illuminate the target. This single point method was the first developed and has the advantage of continuous monitoring of that point on the object as well as the ability to operate at very high bandwidths. The obvious disadvantage is that just one point on the object is monitored. The laser beam can, however, be moved to scan the measuring point over the object. This is called Scanning LDV.

Scanning LDV In Scanning LDV (SLDV) the beam is reflected off of two mirrors such that rotation of one mirror shifts the direction of the beam in the horizontal direction, while rotation of the other shifts it the vertical direction. If this system is combined with a CCD or CMOS camera then the position of the beam on the object can be monitored. The measurement point is moved from point to point, remaining at each location for long enough to collect enough data for a measurement before moving to the next. This allows an area of the object to be imaged, but does not provide information on the time relation between events. Any transient events occurring across the object are not tracked by this method. It does however work well for repeated events such as measuring nodes in a vibrating membrane. A slightly different type of SLDV is Continuous Scanning LDV.

Continuous Scanning LDV Continuous Scanning LDV (CSLDV) is similar to SLDV in that the measurement point is moved across the object. The different is that the point is not held at each point for a period of time. Instead it is moved continuously across the object. This is achieved by driving the position mirrors with sinusoidal voltages and the received signal is a combination of the sinusoidal mirror signals combined with the signal from the vibrating surface, requiring processing to extract the surface motion.

Tracking LDV Tracking LDV is a special type of single point measurement ideally suited to measuring moving objects. The measurement point is moved using the same system as SLDV such that it remains targeted at point on a moving object. An example of this is the side wall of a tyre. During rotation the any point on the tyre follows an almost circular path. To gather information about how that point vibrates with each rotation an LDV point is set to track the rotation of the tyre. This type of LDV has the added advantage that speckle noise is limited as the laser is held on a single point on the object. The disadvantage is the setup is more complicated, both in terms of equipment and alignment.

Full-field LDV The are a number of systems which provide full-field laser Doppler vibrometry. These can be split into those based around a single point system and those based around a high speed camera. Those based on single point systems place a number of discrete points across the surface of the object and measure the Doppler shift of each beam at the same time. This can be achieved by adding a small, constant frequency shift to each of the beams. A single detector system is then used and the signals from the different points separated during processing. Those based on a full-field camera shine diffuse laser light on the target and use a high frame rate camera to sample the light reflected. Processing is then applied to the samples for each pixel to extract the vibration information.

History Although the Doppler effect had been known for many years its use in optical applications had been limited by the light sources available. With the invention of the laser in the late 1950's a source of coherent (all the electromagnetic waves being in phase), spectrally pure (only one colour, or very nearly) light was made available. One of the first pieces of work with lasers was by Rabinowitz et. al. in 1962 and demonstrated an optical interferometer [45]. Lasers provide a good source of coherent, monochromatic, well collimated, bright light. Within the restrictions of coherence this light can be interfered and mixed, and an interferometer provides a good optical setup to demonstrate and utilise the effects of optical interference. The one demonstrated by Rabinowitz et. al. is called a Michelson Interferometer a diagram of which is shown in Figure 10. It is composed of a laser source producing a collimated beam. This beam may be expanded. The beam from the laser passes through a beam splitter which splits it into two beams, or rather splits a fraction (typically 50%) of the light off while letting the rest pass through. These two beams then reflect off mirrors before passing back to the beam splitter. In the simple setup shown (and used by Rabinowitz et. al.) there will be another round of splitting each of the reflected beams, but a fraction of the light from each mirror will pass coaxially to the sensor. A theoretical consideration of the Michelson Interferometer will be given in Chapter 2.

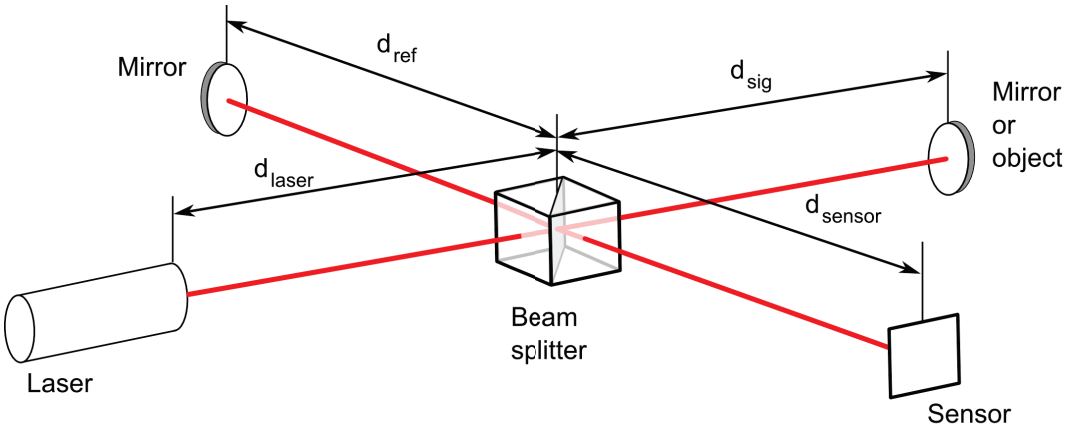

Figure 10: Michelson Interferometer.

Another early piece of work with the new lasers was to confirm the theory published by Raman and Nathe in 1935 and 1936 that ultrasonic waves traveling through a medium would diffract incident light into several orders, see Figure 11. This figure shows a laser beam incident on a water tank with an ultrasonic transducer generating ultrasonic compression waves through the water, almost perpendicular to the beam's travel. The beam is actually at a slight angle to the waves called the Bragg angle.

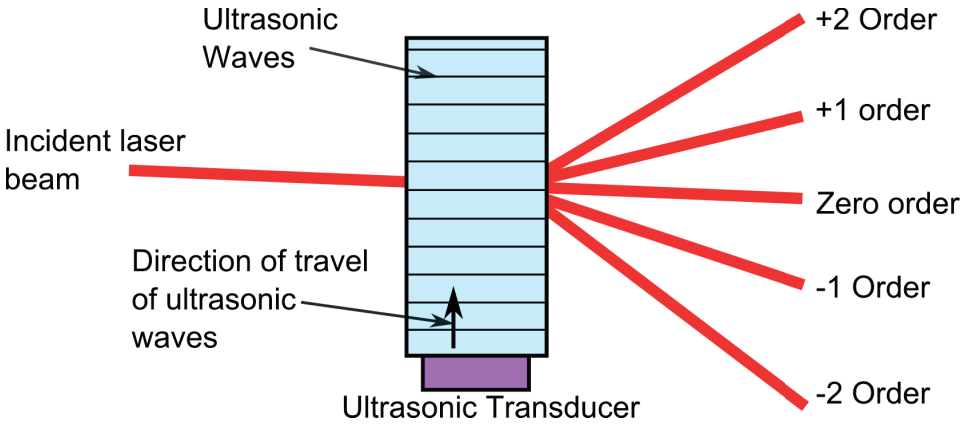

Figure 11: Ultrasonic diffraction of a laser beam.

The theory predicted that as well as diffracting the beam into number of orders, the ultrasonic waves would also add a frequency shift to the light. However, given

that the frequency of light is very high ($\sim 4.7 \times 10^{14}\,Hz$) and the change in frequency was predicted to be the multiples of the ultrasonic driving frequency (initially $25\,MHz$ or less, but now days easily up to $80\,MHz$) there were no spectrometers sensitive enough to measure this change. In 1963 Cummins et. al. [46] demonstrated this shift in frequency quantitatively by optically mixing the light from one diffracted order from the water cell (nowadays these come in the form of crystals with attached ultrasonic transducers and are called Bragg Cells) with the original laser light split off the main beam before it entered the water tank. This is shown in Figure 12. The He-Ne maser at the top of this figure is what would now be called a laser, the beam of which is split into two beams using a half-silvered mirror (M3). The beam reflected off M4 passes through a water tank which has an ultrasonic transducer attached to it. M5 and the slit are used to select a particular diffracted order and M7 is another half-silvered mirror which is used to combine the diffracted beam with the original beam. This combined beam then interferes optically at the photomultiplier.

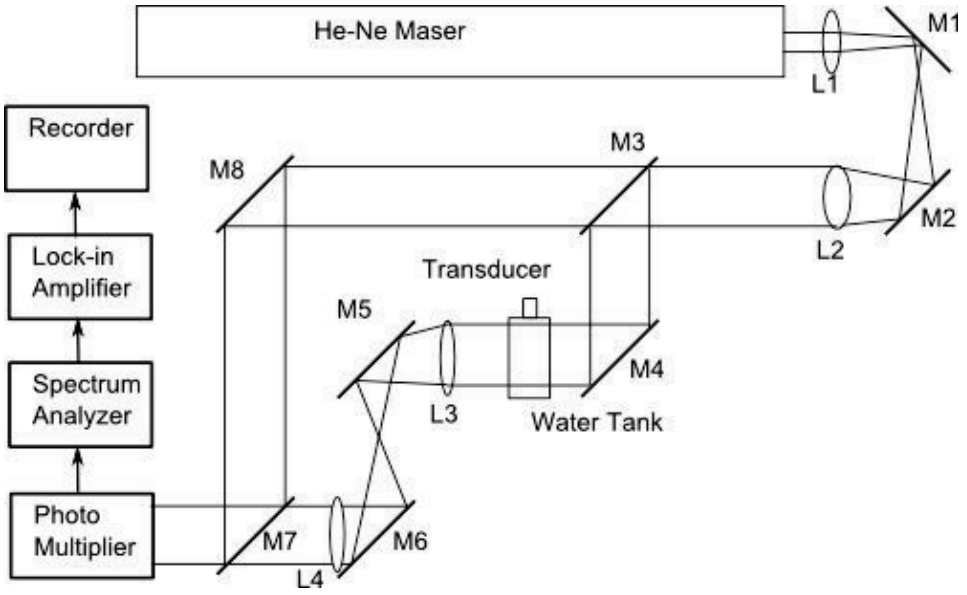

Figure 12: Experimental apparatus used by Cummins et. al. [46].

This experiment demonstrates some useful principles. Firstly the optical setup

used is what is now called a Mach-Zehnder Interferometer (MZI). This is a amplitude splitting device, like the Michelson Interferometer, and has found many applications from measuring the density of gas flows to electron interference fringes [47]. The MZI will also be considered in more detail in Chapter 2. Something else this experiment shows is that Bragg Cells can be used to change the optical frequency of laser light in a controlled way.

As the Doppler effect changes optical frequency, and as this frequency is too high to be directly detected, some method of processing the signal optically is needed. This can be done with a MI, but there are certain disadvantages with this type of system. Given that the laser light is Doppler shifted then mixed with itself, the two options (Doppler shifted to a higher frequency and shifted to a lower frequency) will both give the same results. This is called a homodyne setup, see Chapter 2. This is because the MI basically produces a signal at the different in frequency of the two beams, without consideration on whether one is higher or lower than the other. A better method is to use a MZI setup with a Bragg Cell to add a constant frequency shift to one of the beams. As long as this reference shift is greater than any expected Doppler shift then the difference in frequency of the beams will be centered around the reference shift, and so direction can be measured. This is called a heterodyne setup and more details will be given in Chapter 2.

Cummins and Knable demonstrated this later that year [48] with a Mach-Zehnder Interferometer setup specifically designed to demonstrate this frequency shifting, heterodyne optical setup.

This is not the only way of introducing a frequency shift into laser light. Stevenson, in 1970 [49], demonstrated a rotating diffraction grating to give multiple diffracted orders, each with a frequency shift. This frequency shift being proportional to the rate of spin of the disk. This method was applied to a flow ex-

periment, as will be discussed shortly. There are some problems with the use of a spinning disk; a broad spectrum, low optical efficiency and mechanical vibrations. The main advantages, certainly at the time, were that a disk is relatively cheap and can give small frequency shifts. At roughly $6000\,rpm$ the frequency shift was just under $500\,kHz$ compared with a minimum frequency of about $10\,MHz$.

The next year, 1964, Cummins et. al. [50] showed a way around this, albeit with two modulators and the resultant losses, complexity and cost. This used a similar setup to theirs of the previous year, but heralded as a spectrometer. This shows a slightly different approach to being able to measure, albeit indirectly, the optical frequency of light. In their experiment, shown in Figure 13, they used Bragg Cells (SSB modulators) to add a $12M\,Hz$ shift in optical frequency between a reference and signal beam of a MZI. This setup shows another important use of Bragg Cells: although Bragg Cells will work at a center frequency which is often in the MHz, it is possible to shift both beams by different amounts such that the final frequency difference between the beams is less which may be better for the electronics, where MHz signals cannot be tolerated. One of the beams was then scattered off a solution with suspended microscopic particles. As before the two beams were then mixed at a detector. In this setup any change in colour (frequency) of the light scattered from the solution will shift the difference in optical frequency between the two beams and change the beat signal at the detector. A very sensitive spectrometer.

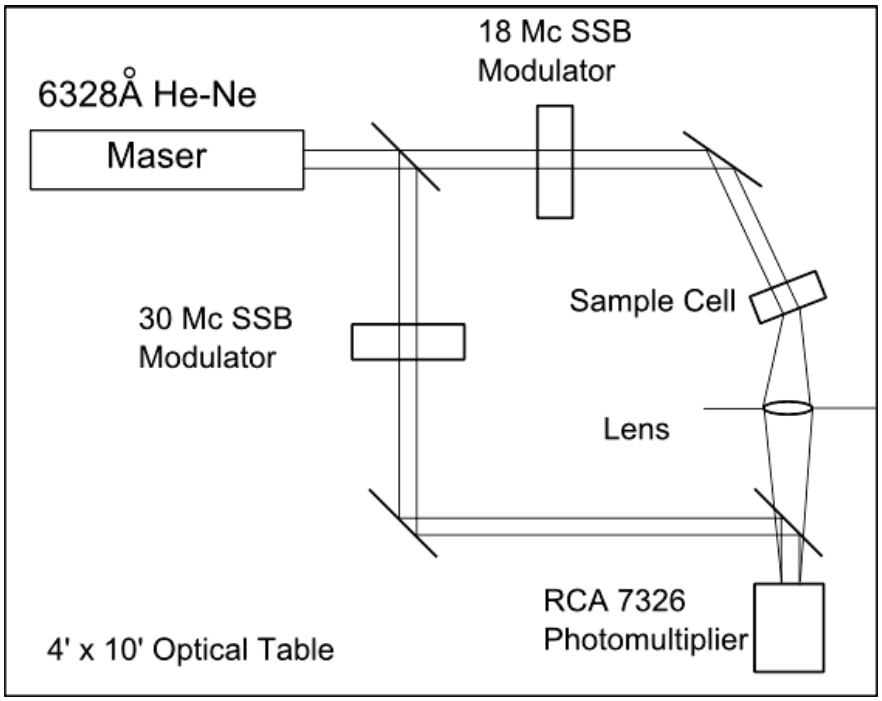

Figure 13: Experimental setup for measurement of Rayleigh scattered light used by Cummins et. al. [50].

That same year Yeh and Cummins reported local fluid flow measurements made using the laser Doppler effect [51]. This demonstrates the optical method of NDE called Laser Doppler Anemometry. Their experimental setup was a relatively simple heterodyne MZI with one arm shifted by a Bragg Cell and the other shifted by a fluid flow. The flow was seeded with polystyrene spheres, and it is actually these which reflect the light out of the flow tube, and give it the Doppler shift.

Laser Doppler Anemometry Laser Doppler Anemometry (LDA) was one of the first applications of the laser Doppler effect and measures the flow of a gas or liquid through the Doppler Shift of the light reflected off of small particles suspended in the flow. This method was first applied to fluids and then applied to gases by Foreman et. al. [52] in 1965 with accurate measurements of particle speeds of up to $\sim 14\,m/s$. The experimental setup now had both the reference beams and signal beams passing through the flow, as shown in Figure 14. This

makes for a simpler optical setup, but does not change the theoretical operation of the system.

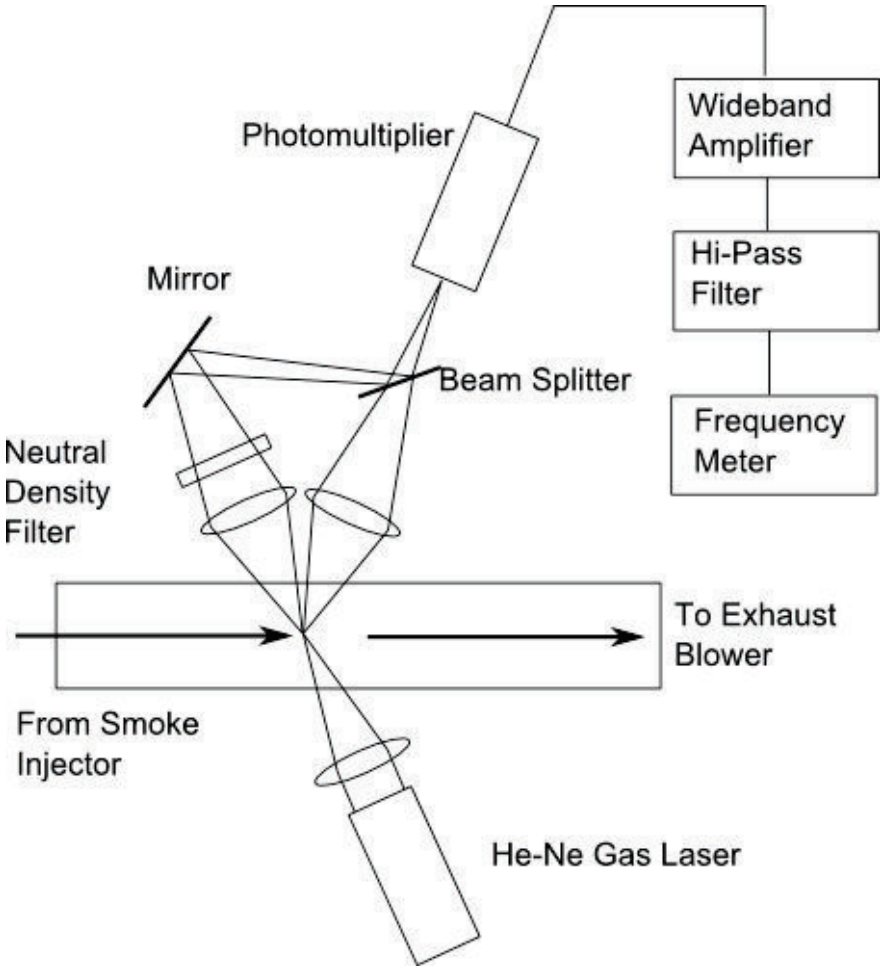

Figure 14: Experimental setup used to measure flow by Foreman et. al.[52].

Foreman continued his work the next year with a similar setup which could measure flows of up to 6000 ft/s ($\sim 1800\, m/s$)[53]. In this paper he also described the operation of the LDA by using the vector description of the light beams, and reflecting one beam off the particle, adding the correct Doppler shift, while the other just passed to the sensor. In 1969 Rudd [54] proposed a different mathematical model for the LDAs' operation. This was based on the idea that the two beams interfered in the region they overlapped, in the flow. This interference produced a grating of interference fringes which the particle then passed over. As

the particle was moving it would reflect back the light falling on it, making a sinusoidal signal. The faster the particle, the higher the sinusoidal frequency. This shows an interesting conceptual difference in considering the use of light. As the light equation can be considered in terms of both frequency and phase, there are often two ways of looking at the Doppler effect, one based on the simple frequency shift and the other based on a phase comparison. This comparative method of looking at the Doppler effect will be covered in Chapter 2.

In 1970, Greated [55] used an LDA to measure turbulence and by 1975 LDAs were portable and designed to traverse $0.5\,m$ flows[56] while still being able to measure velocities from a few cm/s to $50\,m/s$. This apparatus again used a Mach-Zehnder interferometer, but with a spinning disk to introduce the controlled frequency shift of one arm of the interferometer, rather than a Bragg Cell. Work has also started to focus on the electronic processing of the LDA signals. Wilmhurst and Rizzo [57] demonstrated in 1974 a processing scheme that helped compensate for the loss of signal which occurred when two particles were reflecting light.

There are other experimental setups available for LDA, such as that used by Brayton et. al. [58]. This setup used multiple beams and allowed the simultaneous measurement of two orthogonal velocity components of a fluid, at a given point. Work also started to look at the what else the LDA could do. In a paper on particle scattering characteristics of LDAs et. al. describe how the signal can be used for particle sizing [59]. People also started to look at measuring the motion of a solid, rather than a fluid - laser Doppler vibrometry.

There has also been research into a method of full-field flow measurement. Mayers and Lee published a paper describing the process which is based on the absorption of light by Iodine. The molecular properties of Iodine mean that it has a number of sharp absorption lines. In this method of flow measurement the light from a laser is tuned to fall midway along the edge of one of these absorption

lines. Any Doppler shift in the reflected light will shift it on the absorption line and so change how much light is absorbed by the Iodine [60].

Laser Doppler Vibrometry During the early 1970's a number of papers were written on the use of laser Doppler interferometers for vibrometry. In 1977 Ilyasov et. al. [61] demonstrated a vibrometer based on the Michelson Interferometer with increased noise immunity through the use of a vibrating mirror which modulates the reference beam to shift the signal away from DC, an alternative to spinning disk or Bragg Cell. This keeps the signal away from low frequency sources of interference. The setup also used a modified beam splitter which passed most of the power into the signal beam. The reason for this is that the best signal amplitude is detected when the returning signal beam is the same intensity as the reference beam. The signal beam is reflected off the object, often with great losses, and therefore needs more power to start with. Passing most of the power into the signal beam and relatively little into the reference compensates for the losses during reflection from the object without the need to attenuate the reference beam and so gives better signal to noise ratio. This method of measuring vibration has now become known as a Laser Doppler Vibrometer (or, occasionally Velocimeter) LDV.

In 1981 Buunen and Vlaming [62] used the laser Doppler effect in the form of a vibrometer to measure the movement of the small bone in the middle ear of a cat which connects to the eardrum. This shows the sensitivity of interferometric measurements. Instead of a Bragg Cell a spinning disk with a diffraction grating on it was used to give the frequency shift in one of the beams. The electronic processing of their device used a lock-in amplifier to improve the sensitivity of the system. The vibrometer results were compared with a microphone and using a loud speaker were calibrated. This showed a minimum vibration amplitude of

$0.01\,nm$ at $1\,kHz$, although the results from the experiment were normally in the range of $1 \rightarrow 10\,nm$. The frequency range tested was $\sim 100\,Hz \rightarrow \sim 10\,kHz$.

Also in 1981, Dandliker and Willemin [63] used a modified Mach-Zehnder interferometer setup for LDV measurements of both in-plane and out-of-plane vibrations. Their experimental setup is shown in Figure 15. In this figure, M1 and M2 are Bragg Cells giving and D1 is a photodetector. The target was a diffusely scattering surface, as would be expected on a real-life object.

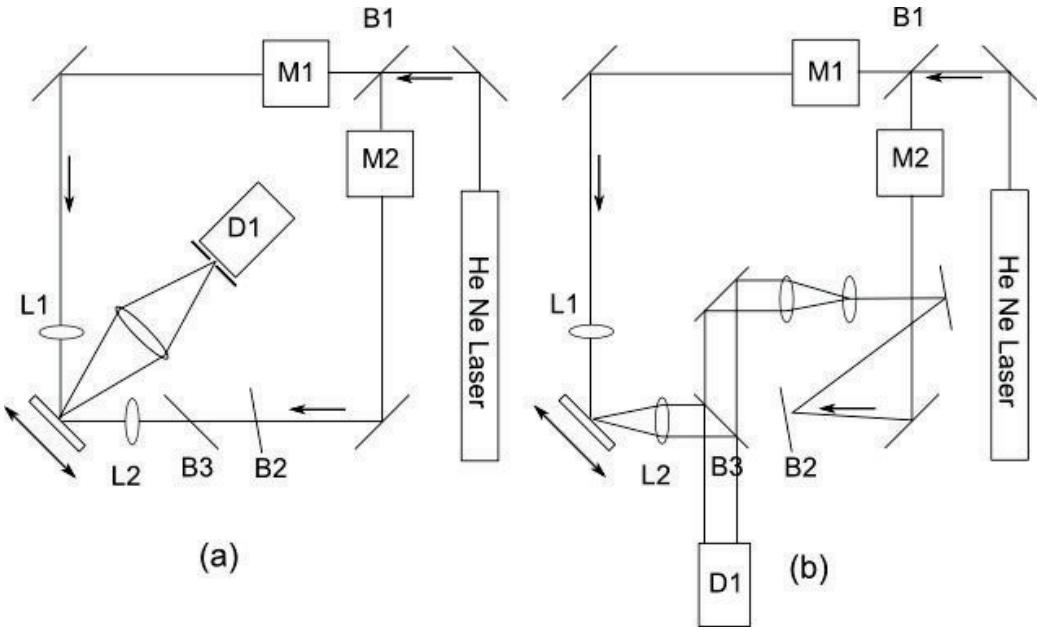

Figure 15: Experimental setup for measuring both in-plane and out-of-plane vibrations, from [63]. (a) is for in-plane vibrations and uses two illumination beams, (b) is for out-of-plane vibrations and just uses one illumination beam.

Interestingly the imaging aperture was set to image at least 100 speckles onto the detector. This was done to compensate for signal loss if as speckles are known to disappear if the object moves too far. As expected the results were in the form of a frequency modulated signal, where the modulation contains the movement information and the center frequency is the difference between the Bragg Cell frequencies. This setup is also very similar to those for ESPI, except that a number of speckles are imaged at once, and add incoherently to give the resultant

signal. This setup gave a measurement spot size of $35\,um$ and for a frequency range of $5\,kHz \rightarrow 5\,MHz$ the minimum vibration amplitude for both in-plane and out-of-plane vibrations was $0.3\,nm$.

Two years later, Willemin and Dandiker [64] showed the same optical setup with a more advanced electronic processing system which provided real-time measurement of both phase and amplitude of in-plane and out-of-plane vibrations. The electronic processing was based on mixing the detected optical signal with a local oscillator (derived from that signal) to give the phase modulation encoded onto the optical signal by the Doppler effect.

As with many new measurement methods, once the basic principle had been proven people started to look at how to add value: the LDV was extended to other applications. Halliwell et. al. [65, 66] showed in 1984 and 1985 the use of a torsional vibrometer by Doppler shifting the light from both arms of an interferometer with the object. This device was designed to be handheld and operate at frequencies below $30\,kHz$ with a maximum dynamic range of $80\,dB$ for torsional vibrations when the frequency was varied.

In 1984 Maynart built a vibrometer based on an LDA [67]. The surface was speckled and used a standard LDA integrated unit, as had now become available [68]. In a change from previous speckle interferometers, only one speckle was imaged onto the sensor. Although that speckle might disappear as the surface moved, resulting in dropout times, it does work well for high contrast low-level signals. Collecting a large number of speckles gives stronger signals (more light collected) but with poorer contrast as the speckles are non-correlated.

Considering the signal received from a single speckle. Given that the experimental setup used the, now familiar, Bragg Cells to add a frequency difference to the beams, the speckle would be modulated at the difference frequency (Bragg Cells frequency) giving an AC signal. This is not the speckle disappearing due to

motion of the object surface, but rather the result of the Bragg Cells shifting of one of the beams. If the surface did move, this AC signal would be modulated by the Doppler effect - increasing or decreasing in frequency depending on the direction of motion. Should that speckle disappear then the AC signal would disappear also. Over a large vibration amplitude this would result in an amplitude modulated AC signal (occasionally going to zero), as well as frequency modulated (by the Doppler effect).

Three types of electronic processing were tried. The first was a frequency counter from an LDA setup, the second used a Schmitt trigger to convert the AC signal into a square wave whose duty cycle was proportional to the instantaneous frequency, and which could be low-pass filtered to provide the demodulated output. The final method suggested was to use a completely digital, bidirectional counter controlled by the optical signal and the Bragg Cell control frequency (effectively a local oscillator). In this last system the accumulated count would be proportional to the instantaneous displacement. This would give $\lambda/2$ accuracy, where λ is the wavelength of the laser light.

Work now started to be done on making the LDV more practical. In 1986 Pickering and Halliwell published work on a portable LDV based around a Michelson Interferometer, and using a spinning disk for frequency shift of one of the beams. Effectively the reference arm reflected off a spinning disk rather than a mirror. The spinning disk does add noise to the optical signal in the form of a moving speckle pattern, but, as pointed out by the author, in many real-life situations the target surface rotates and so will introduce a shift in the speckle pattern anyway.

The final instrument was a tripod mounted and shoebox sized, and used a $1\,mW$ laser output for safety. The spinning disk was spun to give a $1.5\,MHz$ reference frequency shift, giving a practical operational range of $\pm 0.3\,m/s$. Retroreflective tape or paint was used to give a good reflected signal and overall the

system has a noise floor of 10^{-4} to $10^{-5}\,m/s$ over the range $0 \rightarrow 20\,kHz$.

Further work in miniaturising the optical setup of the LDV was done by Suchoski et. al. [69] who designed a multifunction integrated optic circuit containing all the optic signal processing functions of beam splitting (and combining), coupling and frequency shifting a single chip. The laser was a diode laser with a very short coherence length and this was compensated by an optic fibre delay line. The system was shown to work with object reflectivities down to $0.5\,\%$.

In 1994 an LDV was designed by Bokelberg et. al. [70, 71] which could measure vibration and rotation in three orthogonal directions. This setup used three laser sources, reflected off a tetrahedral target to three detectors. The laser sources and detectors arranged around concentric circles. The signals from the sensors were sampled and the processing done by a computer. This system could measure down to $20\,um$ translation and $0.01'$ rotation.

The systems described so far have measured vibration at a single point, with a single detector. By 1999 Stanbridge and Ewins [72] had developed a scanning vibrometer which continuously scanned the beam over a vibrating structure. The vibration modes are then modulated by the motion of the beam over the surface. With appropriate demodulation the mode shape of the vibrating surface can be extracted.

German company called Polytec began to develop and manufacture innovative laser-based test and measurement systems in the 1970's, they are now a major producer of LDV systems. Bauer et. al. published a paper in 2002 [73] which described a MZI heterodyne optical system with digital decoding. This is the basic setup which is used by most commercial LDV systems. As these are single point systems they have large bandwidths ($2\,MHz$) so can measure large or fast vibrations, and also have very good noise floors (less than $1 \times 10^{-12}\,m$).

The applications of LDV systems started to become more varied, the value

added component of the systems started to become processing of the vibrometry data. Castellini and Montanini [74] published a paper in 2002 on using a tracking LDV to measure vibration sources in cars, as part of ongoing research into making cars quieter. Three sources of vibration were measured: the tyre side wall, the timing belt and the rubber blade of the windscreen wiper. Tracking LDV is like single point LDV expect that it is applied to a moving target and the laser beam is moved through the use of controllable mirrors to track that target, therefore keeping the laser spot in the same place on the target. For use on the tyre the tracking system gave a better than $1\,mm$ resolution and the time history of the vibration of that spot showed how the tyre wall vibrated as it touched the road.

Work has also been done on using standard CMOS cameras to image full-field flows. The work of Serov, Steinacher and Lasser in 2005 [75] demonstrated the use of a CMOS camera with built-in processing for 2D imaging of microvascular flows. This method has been used successfully in a number of clinical investigations ranging from skin irritants and burns to organ transplants.

More NDE oriented applications of an LDV is shown by the paper of Prazenica et. al. in 2007 [76]. They used a scanning LDV to study mapping the vibrations of the frescoes at the US Capitol building in order to find structural faults. A survey needed to be done on around $700\,m^2$ of frescoes and in some places these were difficult to reach - showing an advantage of the scanning LDV system. This system is similar to the one used by Castellini and Montanini except that the laser point is held at each one of a grid of locations for a measurement before moving on. In this case around 400 points per meter were used and either a speaker or shaker was used to excite the structures with a chirp signal with a bandwidth ranging from under $100\,Hz$ to $1\,kHz$. It was a Polytec scanning system that was used for this study. The results showed that speckle and surface irregularities added significant noise to the signals and this could affect the processing. Filtering was,

however, used and the paper showed that this is a viable method for detection of faults.

Bucaro, Romano and Abraham [77] also showed in 2003 that from surface vibrations of a plate with an inclusion it would be possible to locate that inclusion. They continued this work in 2007 [78] by detecting and locating detached ribs under thin metal and composite plates from scanning LDV data. The experimental setup is shown in Figure 16. Again an actuator was used to vibrate the object and a commercial vibrometer was used with the novelty lying in the processing of the full-field vibration data. The plate was vibrated with a single frequency and from the amplitude of the vibration it was possible to detect where the ribs were attached, and hence where the detachment was, as shown in Figure 17. For these results the actual vibrations were of the order of $120\,nm$ amplitude. The plate size was $30\,cm \times 30\,cm$ and the scanning LDV was set to 245x245 points, sampled at $64\,kHz$ with 512 samples a piece. 16 averages for each point were used.

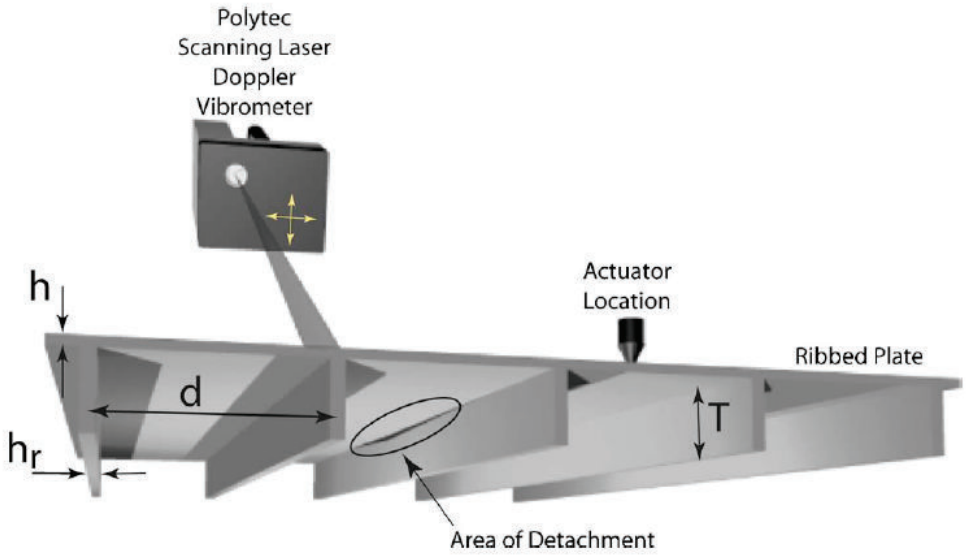

Figure 16: Experimental setup to find the location of rib detachment used by [78]. This system used a commercial scanning LDV.

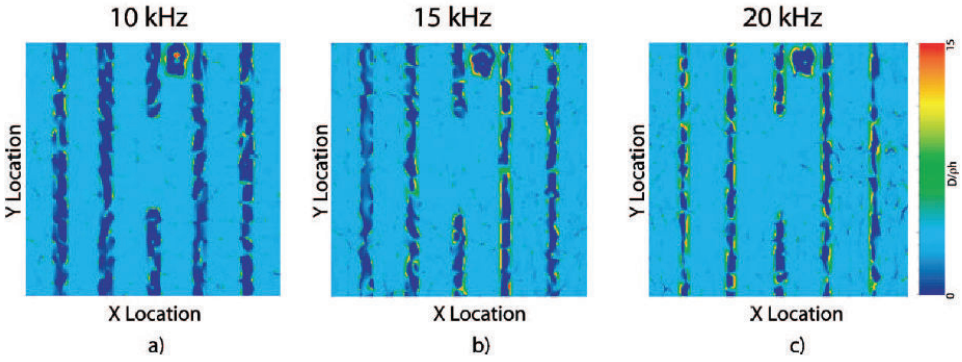

Figure 17: Results of rib detachment location experiment from [78].

Full-Field LDV There are a number of systems which provide full field vibrometry measurements based on the LDV principle. These can be split into those which use a single point LDV system and those which use a full field camera with processing of each pixel's signal over time. The former split the beam from an LDV into a number of separate beams which each measure the vibration of a point on the surface of the object. This splitting typically adds a constant frequency shift to each of the beams to allow the different signals to be distinguished. A single detector is used to measure the reflected light and so the bandwidth of the LDV is effectively divided by the number of points. The methods based on a camera use diffuse laser illumination of the surface of the object and have a number of detectors (pixels) to record the light coming back. The signal from each detector is demodulated to extract the vibration signal. This uses similar processing to a single point LDV and the pixel bandwidth is limited by how fast images can be captured and processed.

In 2002 Aguanno et. al. used a commercial Complementary Metal-Oxide Semiconductor (CMOS) camera with an incorporated Digital Signal Processor (DSP) for quasi real time analysis of full-field vibration analysis [79]. The CMOS-DSP camera used the same carrier based demodulation scheme as some single point systems. Each pixel produces a signal related to the light irradiance falling

on it and this was processed by the DSP to extract the vibration information. The advantages of the CMOS system are what pixels can be accessed randomly and so a small area of the image could be imaged very rapidly. The system still used a single processing device (the DSP) even though a full-field imaging system is used (the CMOS camera). This means that the image size can be decreased resulting in increased frame rate. A single pixel could be sampled at $500\,kHz$ while a window of 64 x 64 pixels could be sampled at $200\,Hz$.

Aguanno et. al. have further developed the CMOS-DSP system to give quasi real-time measurements of a vibrating membrane [80, 81]. This system used a Mach-Zehnder optical interferometer with Bragg Cells to add a small optical frequency shift to one of the beams. This system is the same as the one chosen by this project and is discussed in Section 2.2.4 of Chapter 2. Figure 18 shows the experimental setup used by Aguanno et. al. The light from the laser is split into two beams by BS-1 and both beam are frequency shifted by Bragg Cells (AOMs), except that one is shifted by $102\,Hz$ more than the other. One of these beams is reflected off the vibrating target (the membrane) before being imaged by the CMOS-DSP camera. The other beam is directed to the camera and combined with the beam from the speaker using beam splitter BS-2. After recombination at the camera there is a beat signal of $102\,Hz$ due to the difference in frequency of the two beams. With this configuration the light that is reflected off the speaker is Doppler shifted by the motion of the membrane. This changes the relative frequency between the two beams and so the beat frequency changes. Processing this beat signal is the basis of the processing method used, and contains the vibration frequency and amplitude.

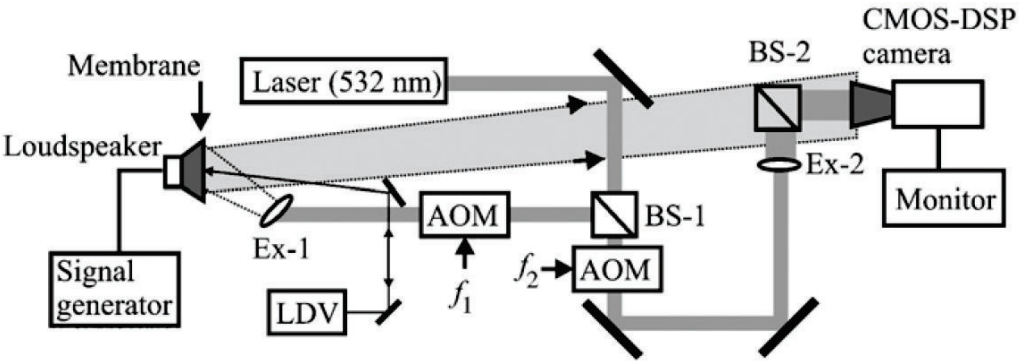

Figure 18: Heterodyne speckle interferometer experimental setup used by Aguanno et. al. [81]

With this new experimental setup it is possible to differentiate of the direction of the motion as the beat signal is either decreased or increased by the Doppler shift of the one beam. The CMOS camera in this system was capable of sampling a 32 x 32 image at $1.8\,kHz$. The system performed quasi real-time measurements on a vibrating membrane and the results again compared with a commercial single point LDV system [81]. The difference between the velocity amplitudes measured by the two systems was of the order of 4% of the average velocity amplitude at a given point.

In 2004 Di Sante and Scalise demonstrated a fibre optic probe based laser vibrometer [82]. The use of fibre optic probes offers the possibility of a smaller, lighter, cheaper and more flexible system due to the use of optic fibres to direct the light to the desired position on the structure. The system also allowed the direction of the motion to be found through the use of a polarising beam splitter to direct the light onto two detectors. The phase shift between the two signals giving a method of finding the direction of the motion. A number of such configurations could be used to give a multi-point LDV, but at the expense of a relatively complex system. The system was compared with a commercial LDV and showed a maximum deviation of the measured data of about 4%.

Other attempts to create a multi-point LDV have concentrated on creating a linear array of points. Li et. al. used Bragg Cells to split a single laser beam into three beams, each with a slight Doppler shift due to the Bragg Cells [83]. The light reflected back from the target object (a spinning disk) was then imaged by a single detector and the additional Doppler shift added to each beam by the velocity of the object removed using an Fast Fourier Transform (FFT) algorithm. Cupido et. al. used a multiple beamsplitter to generate a linear array of sixteen points which were imaged onto a test object. The reflected light was then imaged onto two detector arrays using a similar processing scheme to Di Sante and Scalise [82]. Using a different beamsplitter this system could be configured to given any pattern of points.

Optic fibre has been used extensively to allow a more flexible system architecture for laser measurement systems. Connecting the measurement probe to the rest of the system allows a more flexible and easy to use system. Kilpatrick and Markov have developed a matrix laser vibrometer (MLV) which provides an array of 16x16 points of measurement [84, 85]. The system uses a mobile sensor head connected by a five meter fibre optic umbilical to the rest of the system allowing easy positioning during testing. The system has aspects of both LDV and ESPI. The beam from a laser is split into a grid of 16x16 beams using a diffractive optical element and these beams are focused on the test object. Figure 19 shows the layout of the system with the main electronic and optical equipment in a single unit and the MLV probe separated by a fibre optic umbilical. The light reflected back is focused onto the end of a fibre bundle and the other end of this bundle is imaged using an array of detectors and the data processed to give concurrent measurements from all 16x16 points. The system was demonstrated to be able to measure a $1\,nm$ vibration at $2\,kHz$.

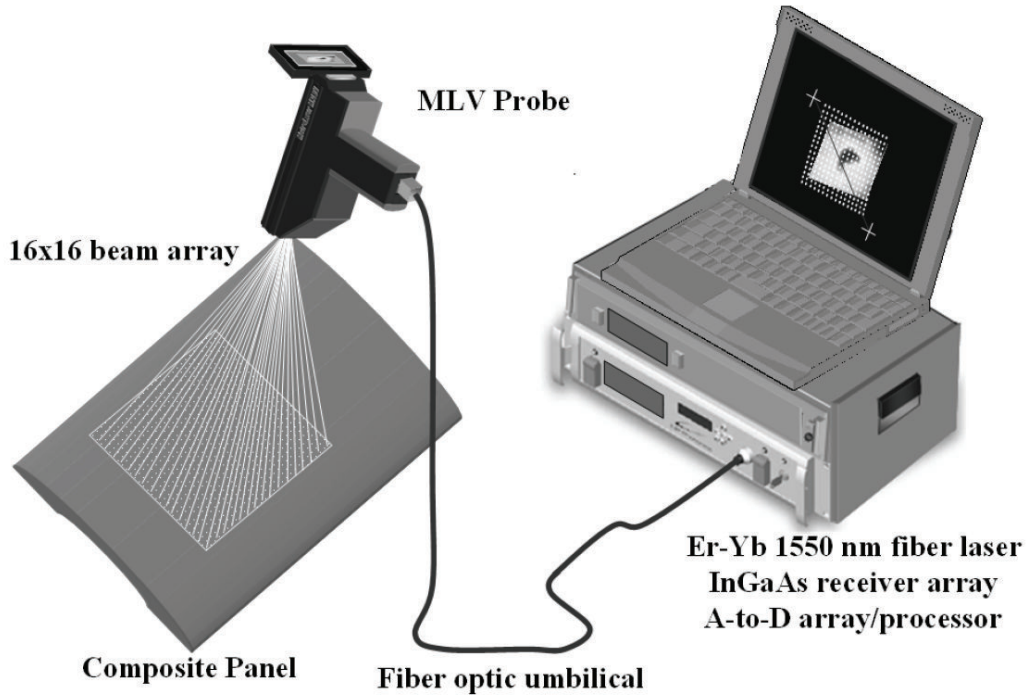

Figure 19: Configuration of the matrix laser vibrometer of Kilpatrick and Markov [85].

1.5.8 Conclusion

This section has covered a number of optical methods of measuring surface vibration using non-contact methods. Table 2 summarises the different methods of optical, non-contact vibration analysis considered above. The most sensitive methods are those which use optical interferometry which allow sensitivities of less than the wavelength of light. Two important effects were also discussed, namely the speckle effect and the Doppler effect.

1.6 Conclusion

This chapter has introduced the project and its aims. These are to investigate a full-field LDV camera based on a single pixel which implements LDV processing. It also gave an overview of the field of Non-destructive Evaluation and briefly

Table 2: Summary of the different methods of optical non-contact vibration analysis.

Method	Advantages	Disadvantages
Moire	Several methods available. Can use digital projection for automation.	The grid used must have high spatial frequency for sensitive measurements
Holography	Provides a great deal of information on the object. Sensitive to movement of the order of the wavelength of light.	Very precise setup required. Photographic process adds significant processing delay.
Speckle Photography	Wide range of displacement sensitivities. Can be setup to measure only one displacement direction at a time.	Sensitivity is set by the resolution of the media, however photographic plates require long development times.
ESPI	Digital imaging provides easier storage and more flexible processing. High frame rates can capture fast transient events.	Fast cameras are expensive. Measures motion indirectly through changes in the speckle pattern, so the surface must produce one.
LDV	High bandwidth and good sensitivity. Relatively simple configurations and fibre optic probes allow easy positioning. Includes full-field methods.	Full-field methods have limited resolution. Normally a single point method. Scanning systems are more complex and require mechanical scanning components.

discussed the different methods, including their advantages, disadvantages and basic operation. A quick history of laser based Doppler systems was given, starting with the laser Doppler anemometer then moving on to laser Doppler vibrometer systems which were described from their early days in the 1970's and 1980's to the present day, including other full-field vibrometry methods. A summary of the different LDV systems was then given. Finally this led to a description of the device designed during this thesis as well as the benefits of full-field vibrometry.

Many of the LDV systems mentioned during the this review were used to measure diffuse surfaces and as such they worked in the presence of speckle patterns, whether they used them or not. Dandliker and Willemin [63] in 1981 showed a system which effectively captured a number of speckles and processed them all together. Maynart [67] three years later also worked with speckles, although this paper used both a single speckle and multiple speckles. The speckle patterns were not used to make the actual measurements, like speckle photography, but rather

the light within the speckle was mixed to give the required beat signal - the idea behind LDV. It can be shown that ESPI, in the form of time varying speckle, and LDV are effectively the same [86]. When considering LDV the speckles are effectively processed in depth, as LDVs are inherently single point and the light in a given speckle is Doppler shifted when the object surface vibrates. For ESPI the speckles are also processed in depth, but this is done by mixing the speckles with a reference and then taking multiple images. Given the reference has a constant phase, for every 2π change in the phase of the light in a given speckle, it will flash once. This is the same as LDV, except from an image vs. time point of view rather than a sample vs. time one.

To create a camera to image this changing speckle pattern, pixels must be made which will process the light falling on them. As this light is contained in speckles it makes sense to work on the level of speckles. This means considering a single speckle falling on a single pixel, independent of the rest of the camera. This is the operation of an LDV, and so the LDV philosophy will form the starting point for the design of a pixel for processing Doppler vibration signals. The final proposed camera could be though of as a grid of very small, single point laser Doppler vibrometers, all working independently. And it is this concept which the pixel for this project will be designed to fulfill.

The field of optical NDE is an important area of research and this chapter has shown the range of systems available for measuring object vibration without needing contact with the object. Full-field laser vibrometry is the measurement of the vibration of a surface illuminated by a laser. For it to be full-field the whole surface must be imaged concurrently, like a conventional camera takes images.

2 Ground Work and Theory Interferometers and the Proposed Pixel Architecture

2.1 Introduction

So far Chapter 1 has introduced the NDE method called Laser Doppler Vibrometry (LDV) with some of its applications and general theory. This chapter will now provide a more detailed look at interferometers, the speckle effect and LDV. This will start by describing the two types of optical interferometer which are normally used for LDV systems. These interferometers will be described in terms of the optical setup and theory of operation, and simulated results will be given to demonstrate the expected signals. After this, the speckle effect will be described in more detail with some simulations of the result of non-uniform illumination as it might be applied to this work. This will then lead on to a summary of methods of LDV processing based on the signals from the interferometer, from which a candidate for a full-field LDV will be chosen. The design of a single pixel using the chosen processing method forms the core of this work and will be covered in Chapter 3.

2.2 Interferometers

An Interferometer is a device which directs radiation over multiple paths, before recombining the separate beams resulting in optical superposition and mixing [87]. This results in intensity changes due to the different path lengths. Two interferometers shall be considered, namely the Michelson Interferometer and the Mach-Zehnder Interferometer [87, 88]. These are the most popular interferometers used for LDV [54, 56, 67, 89]. Before continuing it is important to consider that

the interferometers will be used for NDE. For this reason the two beams will be used differently. One will be used to illuminate the test object and will be called the signal beam. The other will be used as an optical reference in the form of coherent EM waves at a given frequency (often shifted slightly from the original laser). This will be called the reference beam. The sections of the interferometer with the different beams are often called arms as they split off the source beam, leading to the terms signal arm and the reference arm.

The light source used for the interferometers will be a laser as this gives a good source of collimated, monochromatic light. This light gives a good approximation to a plane, harmonic, coherent electromagnetic wave. The coherent property of the light is important as this allows light from the laser to be split (for example by an interferometer) and then recombined such that it interferes. An ideal laser would produce a continuous train of sinusoidally varying electromagnetic radiation of constant amplitude and frequency. This would give a single spectral line - all the energy in a single optical frequency.

This is not the case in reality as the electron transitions responsible for the generation of light have a finite duration. This results in a number of finite length wave trains. The result of this finite nature is that the spectral line is spread over a range of frequencies. Since these wave trains come from atoms which are in random thermal motion there is also a slight Doppler shift to the light. Collisions between atoms interrupt wave trains adding further spectral spread. These wave trains mean that for a given beam of light, as used in this thesis, the electromagnetic waves are nicely sinusoidal for only short periods of time, and the phase can only be reliably predicted during these times. Comparing the light at a given time with the light from too long in the past would show a random change in phase - the electromagnetic waves would be uncorrelated (incoherent). The light from these two points in time could not be interfered.

This time can be converted into a distance for convenience by multiplying it by the speed of the electromagnetic waves - the speed of light. This distance is called the coherence length of the laser and can vary from mm's to m's. When considering an interferometer: if the lengths of the reference and signal arms differ by more than the coherence length of the laser they will no longer interfere. The coherence length for the laser used in this project is in the cm's region.

The equations which describe EM radiation are collectively called Maxwell's equations after James Clerk Maxwell who summarised and extended the previous empirical knowledge on the topic [47]. Maxwell's equations have a number of solutions and Equations 4 (electric field) and 5 (magnetic field) are solutions of these equations for a plane, harmonic wave. In Equation 4, $\overrightarrow{E_0}$ represents the vector amplitude to take into account the polarization of the wave, however this can be ignored (written as a scalar) as the laser light is polarized and this polarization will be fixed during the work. ω_c is the angular frequency of the wave, defined as $2.\pi.f_c$, where f_c is the optical frequency. For the laser used (a He-Ne gas laser) the optical frequency is about $4.7 \times 10^{14}\,Hz$. k is the wave number and is defined as $2\pi/\lambda$ where λ is the wavelength of the light, for the laser this is $632.8\,nm$. $\overrightarrow{r}$ is the vector describing the direction of travel of the light. For the experimental setups used for this thesis the direction will be well defined and constrained such that $\overrightarrow{r}$ also becomes a scalar quantity r, simplifying the mathematics.

$$\overrightarrow{E} = \overrightarrow{E_0}.cos\,(\omega_c t - k.\overrightarrow{r}) \tag{4}$$

$$\overrightarrow{H} = \overrightarrow{H_0}.cos\,(\omega_c t - k.\overrightarrow{r}) \tag{5}$$

In both equations $\omega_c t$ describes the wave as being sinusoidally varying with

time as it moves past a fixed point in space, $\overrightarrow{r}$ fixed. If t is fixed, then $k.\overrightarrow{r}$ describes how the wave is also sinusoidal through space at a given point in time. Given the symmetry of the equations used to derive these it is only necessary to consider either the E or H fields. The E field only will be considered in this project as it is considerably more effective at doing work on charges and will be called the optical field.

Using Euler's Formula Equation 4 can be written as by Equation 6 where E_0 is the amplitude of the electric field oscillation.

$$E\left(t,r\right) = E_0.cos\left(\omega.t - k.r\right) = Re\left\{e^{i.(\omega.t-k.r)}\right\} \tag{6}$$

It should be understood that it will always be the real part of the exponential which is of interest and so the $Re\{\}$ will be suppressed. Equation 6 can also be split into two terms: one with the optical frequency and the other with the displacement, as shown in Equation 7. This is useful as it shows that the equation can be split into one part which varies at the optical frequency and another which is dependent on spatial position (or distance the wave has traveled). This can be considered a phase delay of $\frac{2\pi}{\lambda}.r$ where λ is the wavelength of the light in that medium and r is the distance in a given coordinate system.

$$E\left(t,r\right) = E_0.e^{i.\omega.t}e^{-i.k.r} \tag{7}$$

2.2.1 Michelson Interferometer

The Michelson Interferometer (MI) is the simplest of the two interferometers, both in construction and for alignment [31, 67, 88]. It is composed of a single beam splitter which splits the laser source beam into two beams, each of these is reflected off a mirror back to the beam splitter. These two returning beams undergo another round of splitting with a fraction of the radiation from each

mirror going to the sensor. This light is then interfered at the sensor and due to the sensor being sensitive to the irradiance they are mixed. The setup is shown in Figure 20. The distance from the source laser, mirrors and sensor to the beam splitter are shown. The beam may be expanded between the source and the beam splitter if full-field imaging is required.

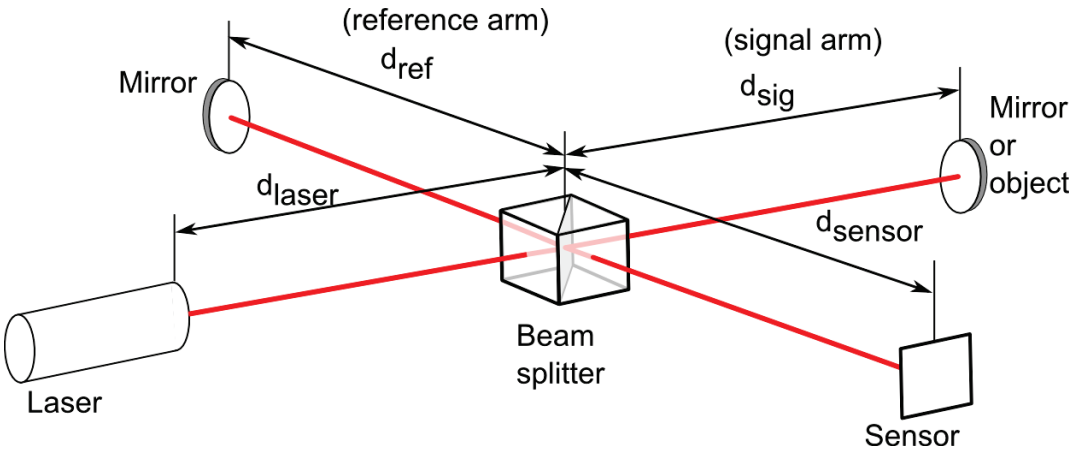

Figure 20: Michelson Interferometer experimental setup.

The mirror on the signal arm may be removed and replaced with an object. Then if the surface is displaced, d_{sig} will be varied. This change in the signal arm length is what the interferometer measures. This is how the MI is used for LDV, by recovering this change in distance the vibration of the object surface can be calculated. Given that the mirror on the signal arm is replaced by the object during operation, the mirror being used just for alignment, this will be called the object mirror. Appendix A gives a summary of the alignment procedure.

Theory Looking at Figure 20, the laser is split into two beams. Both beams must cover the d_{laser} and d_{sensor}, however, the signal beam overs d_{sig} and the reference beam covers d_{ref}. If the reference and signal arms are different lengths ($d_{ref} \neq d_{sig}$) then the beams will have a relative phase shift by the time they reach the sensor. If d_{sig} changes with time, in other words the object mirror/ object

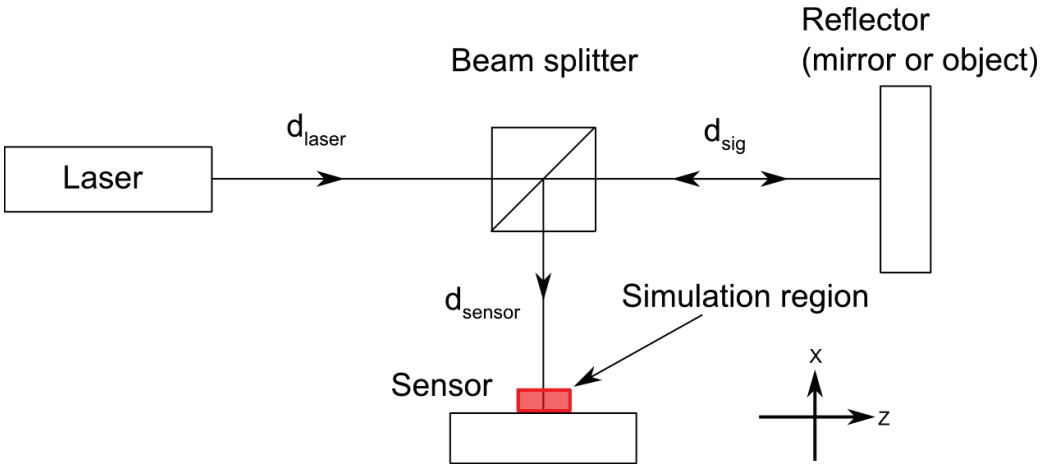

Figure 21: The path of the signal beam of the MI optical setup after removing the reference arm.

surface moves, then the phase difference between the signal and reference arms will also change with time. The MI setup can be simplified to the setup shown in Figure 21 by ignoring the reference arm for now. This simplified setup has a laser source with a simple path to a reflector. The beam splitter is included simply to show that the light to and from the reflector is coincident, while still having the sensor and laser in different places.

If the distance between the laser and the sensor via the reflector (while it's stationary) is $d_0 = d_{laser} + d_{sensor} + 2d_{sig}$ and the reflector is moving with a sinusoidal motion (vibrating), then the distance from the laser to the sensor $d(t)$ can be written as Equation 8, where A_{obj} is the amplitude of the vibration and ω_{obj} is its angular frequency. $A_{obj}.cos(\omega_{obj}t)$ is doubled as the light must cover the distance the reflector has moved twice, once on the way toward it and once on the way back.

$$d(t) = d_0 + 2.A_{obj}.cos(\omega_{obj}t) \tag{8}$$

Signal Beam Now that we have an equation for the distance the signal beam has traversed we can write an equation for the light falling on the sensor from the

object mirror [31, 87]. The mathematical approach is the same as that used by Mansor for deriving the Doppler equations for ultrasound waves [90] and starts by considering the light equation from a purely time based point of view. There is a time delay between the laser source and the sensor due to the distance the light has to cover to travel between them. This time delay can be included in the light equation through the variable $t \rightarrow t - \delta t$ as shown in Equation 9.

$$E_{sig}(t) = E_{sig0} cos\left(\omega_c . \left(t - \frac{d(t)}{c}\right)\right) \tag{9}$$

Considering that the time taken is a function of the distance and speed: $\delta t = d/c$, where d is the distance and c is the speed of light in air.

$$E_{sig}(t) = E_{sig0} cos\left(\omega_c . t - \frac{2\pi f_c d(t)}{c}\right) \tag{10}$$

Expanding this equation.

$$E_{sig}(t) = E_{sig0} cos\left(\omega_c . t - \frac{2\pi f_c 2 . A_{obj} . cos(\omega_{obj} t)}{c} - \frac{2\pi f_c d_0}{c}\right) \tag{11}$$

Using the $k = 2\pi/\lambda$ and $\lambda = c/f_c$.

$$E_{sig}(t) = E_{sig0} cos\left[\omega_c . t - k2A_{obj} cos(\omega_{obj} t) - kd_0\right] \tag{12}$$

Equation 12 therefore shows how the signal beam is affected by the movement of the object. There are three terms in this equation: $\omega_c . t$ which represents the oscillation at the optical frequency, $k2A_{obj} . cos(\omega_{obj} t)$ which represents how the movement of the object mirror affects the signal and kd_0 which is a constant phase shift due to the distance from the beam splitter to the object mirror when its at rest.

Doppler Effect The Doppler Effect describes the change in frequency of a signal by the relative motion of the transmitter and receiver[27], and forms the basis of LDV. The frequency received by a receiver (R) from a transmitter (T) is derived in Gill [27]. A generalized diagram of T and R is shown in Figure 22. The velocity of T is u and the velocity of R is v.

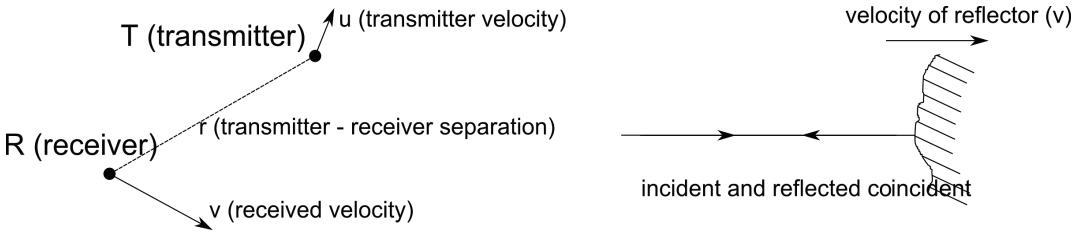

Figure 22: The Doppler effect derivation. Diagram showing transmitter and receiver, and their respective velocities. In the general case on the left and in the specific case of the interferometers used in this thesis on the right.

These velocities are, in a general case, free to have any direction, however in the case of the systems in this project the directions of movement and the light paths will be constrained. For the MI the signal beam will be aligned perpendicular to the object's surface. Only motion in the direction of the beam will change its frequency and therefore register. Also, the laser (T) and sensor (R) will always be stationary and only the object mirror moving. This means that there will be two Doppler shifts, one as the object surface 'receives' the light and a second when it re-transmits it.

In this case the Doppler Shift can be written as Equation 13 [27, 91], where v is the velocity of interest (that of the object mirror), f' is the resultant frequency (after shift), f_0 is the original frequency and c is the speed of the signal (the speed of light in this case). Δf is the Doppler shift.

$$f' = f_0 \left(1 - \frac{v}{c}\right) \rightarrow \Delta f = -\frac{v}{c} \tag{13}$$

The Doppler shift is applied twice as the mirror is first the receiver, and then

the transmitter, resulting in two Doppler shifts. This gives us an equation for the frequency change due to the velocity of the reflecting surface. The velocity is found by differentiating the displacement of the surface. The displacement is given by Equation 8. This velocity is shown in Equation 14.

$$v\left(t\right) = \frac{d}{dt}\left\{A_{obj} cos\left(\omega_{obj} t\right)\right\} = -\omega_{obj}.A_{obj}.sin\left(\omega_{obj}.t\right) \tag{14}$$

The frequency of the reflected signal can now be found by inserting Equation14 in Equation13 to give Equation 15.

$$f\left(t\right) = f_c + \frac{2.f_c.\omega_{obj}.A_{obj}}{c} sin\left(\omega_{obj}.t\right) \tag{15}$$

If the angle (θ) of the sine or cosine $(cos\left(\theta\right))$ changes with time the rate of change is called the angular frequency, and is defined as by Equation 16 [92].

$$\omega = \frac{d\,\theta}{dt} \tag{16}$$

If the angular frequency is constant then it is related to the frequency of the sinusoid by Equation 17.

$$\omega = 2\pi f \tag{17}$$

If the angular frequency changes with time then the integration shown in Equation 18 must be used.

$$\theta\left(t\right) = \int \omega\left(t\right).dt \tag{18}$$

Using this results the phase of the electromagnetic wave can be found from the frequency as shown in Equation 19. The ω_{obj} cancels and $\lambda = c/f_c$ can be used to simplify the Doppler part of the signal. Finally $k = 2\pi/\lambda$ is used resulting in the

phase given in Equation12, $\omega_c.t - k2A_{obj}cos\left(\omega_{obj}t\right) - kd_0$, where ϕ_0 is a constant phase term which practically results from the distance the beam has traveled, d_0.

$$\phi\left(t\right) = 2\pi \int \left\{f\left(t\right)\right\} dt = \omega_c t - \frac{2.2\pi.f_c.\omega_{obj}.A_{obj}}{c.\omega_{obj}} cos\left(\omega_{obj}.t\right) - \phi_0 \tag{19}$$

This derivation is quite a bit more complicated than the delay method first used, but it does highlight that much of the work with optical equations shows two ways of looking at the theory: from a frequency point of view or from a phase point of view. These are in-fact the same as frequency is rate of change of phase - the two are inextricably linked.

Reference Beam The reference beam is simpler and is written in Equation 20. For this equation the light beam is phase shifted by its travel from laser to the sensor via the reference mirror with total distance $d_1 = d_{laser} + 2.d_{ref} + d_{sensor}$.

$$E_{ref}\left(t\right) = E_{ref0}.cos\left[\omega_c.t - kd_1\right] \tag{20}$$

Irradiance The two waves will interfere at the sensor where the resultant electric field will be the sum of the two waves [93]: Equation 21

$$E_{tot}\left(t\right) = E_{sig}\left(t\right) + E_{ref}\left(t\right) \tag{21}$$

Irradiance is the average rate of flow of energy per unit time per unit area (W/m^2) and is also sometimes called the intensity [2, 47, 94]. The vector expressing the irradiance (I) is time average of the cross product of the electric and magnetic fields and is called the Poynting Vector. The irradiance itself is shown in Equation 22 when considering only the electric field. The average is done over time T which must be considerably longer than the period of the EM wave. In the case of light all current methods of detecting its irradiance do so over a sufficiently

long time - they cannot respond to the optical frequency oscillations and therefore average them out. In Equation 22 c is the speed of light and ϵ_0 is the electric permittivity of free space. $c = 2.99792458 \times 10^8\, m/s$, $\epsilon_0 = 8.8542 \times 10^{-12}\, C^2/N.m^2$.

$$I = \left\langle c^2 \epsilon_0 . \overrightarrow{E} \times \overrightarrow{H} \right\rangle_T = \frac{c.\epsilon_0}{2} E_0^2 = c.\epsilon_0 \left\langle E^2 \right\rangle_T \tag{22}$$

When working with complex numbers then the irradiance can be written as shown in Equation 23. In this equation $\widetilde{E}$ is the complex function describing the EM wave, $\widetilde{E}^*$ is its complex conjugate, $\langle\rangle$ is the time average over a long time compared with the period of the function, and $Re\{\}$ represents just the real part of a function.

$$I = c.\epsilon_0 \left\langle \left(Re\left\{ \widetilde{E} \right\} \right)^2 \right\rangle = c.\epsilon_0 . \frac{1}{2} \widetilde{E} . \widetilde{E}^* \tag{23}$$

From here on it is more convenient to use the complex forms of the light equation along with the equation for the irradiance, $I = c.\epsilon_0 . \frac{1}{2} \widetilde{E} . \widetilde{E}^*$. First the equation for the total electric field is rewritten as Equation 24.

$$E_{tot}(t) = E_{sig0} . e^{i(\omega_c . t)} . e^{-ik\left(2A_{obj} cos\left(\omega_{obj} t\right) + d_0\right)} + E_{ref0} . e^{i(\omega_c t)} . e^{-ikd_1} \tag{24}$$

The sensor is sensitive to the irradiance of the light which is proportional to the time average of the square of the electric field. This can be calculated using the steps shown in Equation 25, where $I_s(t)$ is the irradiance due to the light (the constants in the equation for the irradiance have been ignored for now as the actual intensity value is not important at the moment, only how it changes).

$$I_s(t) = \frac{1}{2} E_{tot} . E_{tot}^* \tag{25}$$

$$I_s(t) = \frac{1}{2}\left[E_{sig0}.e^{i(\omega_c.t)}.e^{-ik\left(2A_{obj}cos\left(\omega_{obj}t\right)+d_0\right)} + E_{ref0}.e^{i(\omega_c t)}.e^{-ikd_1}\right] \quad (26)$$
$$.\left[E_{sig0}.e^{-i(\omega_c.t)}.e^{ik\left(2A_{obj}cos\left(\omega_{obj}t\right)+d_0\right)} + E_{ref0}.e^{-i(\omega_c t)}.e^{ikd_1}\right]$$

The terms with ω_c are at the optical frequency and are therefore averaged out by the sensor and are removed for simplicity.

$$I_s(t) = \frac{1}{2}\left[E_{sig0}^2 + E_{ref0}^2 + E_{sig0}E_{ref0}.e^{-ik\left(2A_{obj}cos\left(w_{obj}t\right)+d_0\right)+ikd_1}\right] + \quad (27)$$
$$\frac{1}{2}\left[E_{sig0}E_{ref0}.e^{ik\left(2A_{obj}cos\left(w_{obj}t\right)+d_0\right)-ikd_1}\right] \quad (28)$$

Using the complex form of cosine gives Equation 29.

$$I_s(t) = \frac{E_{sig0}^2 + E_{ref0}^2}{2} + E_{sig0}E_{ref0}cos\left[k\left(d_1 - d_0 - 2A_{obj}cos\left(\omega_{obj}t\right)\right)\right] \quad (29)$$

Simulation Introduction Now that we have an equation for the irradiance it is possible to simulate the sensor values. Again, looking at Figure 21, there is a region marked just in front of the sensor. This simulation will consider this region, viewed from the above, extending in space across the sensor and back toward the beam splitter. The first thing to consider is that the optical frequency components of the beam have been removed. This was because the sensor is unable to measure these and so will only register the average irradiance due to the beams. Any phase difference, however, will register to the sensor. Figure 23 shows an example, with arbitrary units, of how the phase difference between two signals will affect the resultant after mixing. This figure shows two different frequency signals (blue and green) of the form $cos(\omega t)$. The red signal shows the superposition of the other

two. In the case of electric fields, the red would be the resultant electric field. The green's phase changes more rapidly than that of the blue so the phase difference is always increasing, but given that $cos\,(0) = cos\,(2\pi)$, at certain points they have the same phase or go completely out of phase (phase difference of π).

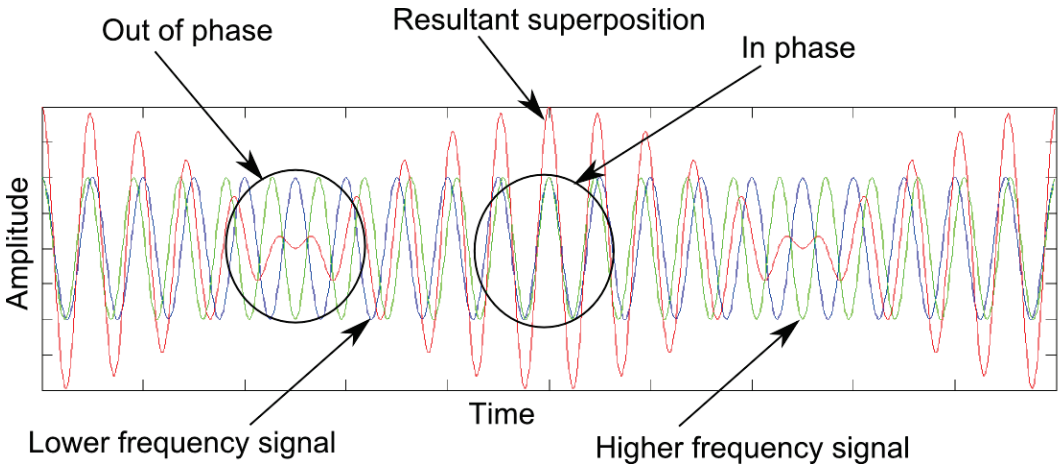

Figure 23: Superposition of waves showing two signals of slightly different frequency being mixed (blue and green). The resultant of this superposition is shown in red. When the two signals are in phase there is a large resultant, and when they are out of phase there is a small resultant.

If they were both the same frequency, but with a constant phase difference of π then the resultant would be zero for all time - they would cancel/ be out of phase. Given that the signals in the MI have the same frequency when the object mirror is stationary they will be superimposed at the sensor and only their relative phase will control the signal - giving a constant irradiance value. If they are in phase then there will be a big signal, the sensor will average out the actual optical oscillations, and the irradiance will be high. If they are out of phase then they will cancel and the irradiance will be small. Equation 29 shows this relationship.

As the reference and signal beams travel in time (and space), their relative phase is constant, when the mirrors are stationary, and if the mirrors move then the phase difference will be controlled by this motion. The optical frequency component of the beams is therefore unimportant as firstly, it cannot be detected,

and secondly the irradiance is dependent on the phase.

The movement of the waves is however important when there is a time dependent variation in the phase difference between the beams (as happens when one of the mirrors moves). By removing the optical frequency components from the wave equations the beams are effectively fixed in time. Figure 24 shows the format of the results from the simulations. A grey-scale image on the left which is the irradiance which would be experienced by a surface at that point in front of the sensor. This is a measure of the field strength and shows what would register if sensor was swept over this region, either holding time fixed or infinitely fast. Alternatively this simulation shows what would happen over time as the light strikes the sensor. For a given point on the sensor, this would result in the trace on the right of the figure which is the signal that would be detected by an infinitely small sensor at position 'A'. Practically this is the signal that would be detected as time went by with the real interferometer setup.

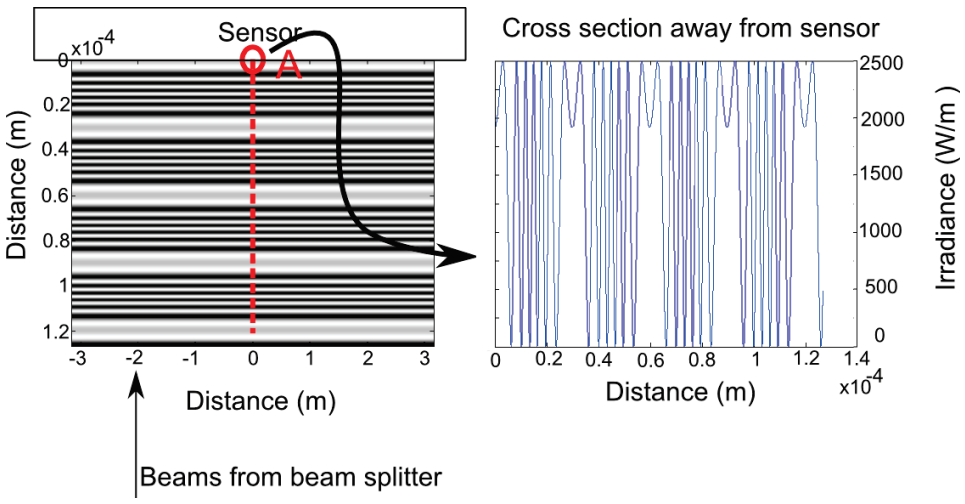

Figure 24: Simulation of MI. The irradiance in the region just in front of the sensor. The image on the left is the irradiance in the simulation region. The sensor is at the top and the beam splitter is off the bottom. The plot on the right is a section through the simulated region, away from the sensor.

Perfectly Aligned without Object Movement If the beam amplitudes are the same ($E_{sig0} = E_{ref0} = E_0$) and the object mirror is not moving ($A_{obj} = 0$) then Equation 29 is simplified. If the difference in arm distances are the same ($d_0 = d_1$) then the result will be $I_s(t) = E_0^2 + E_0^2 cos(0) = 2E_0^2$, ie the sum of the irradiances of the two beams. If the difference in arm distances is half a wavelength ($\lambda/2 \rightarrow \pi$) then the result will be $I_s(t) = E_0^2 + E_0^2 cos(\pi) = 0$, and there will be no irradiance at the sensor.

Perfectly Aligned With Object Movement Figure 25 shows the simulation results for the MI. The top shows what would happen if the object mirror was now vibrated. The phase difference between the two beams changes with time. As the simulation is done with time fixed, but the beams are in motion, this change in phase takes the form of phase difference at different distance from the sensor. For this set of results the traces on the right shows how the irradiance changes, with the vibration of the object superimposed.

It should be pointed out that the frequency of the vibration of the object has been exaggerated to allow several of the waves to be observable in the space allowed. The speed of light means that for any reasonable oscillation, the whole of the reflected signal beam will effectively be at the same phase, as a change in phase due to the motion of the object mirror will almost instantly propagate through the whole interferometer. The simulation is intended to give a better understanding of the system. The horizontal axis is, however, correctly scaled.

It can be seen that the resultant signal has almost stationary points with higher frequency oscillations in between. The stationary points are when the object mirror is almost stationary, albeit only momentarily, during its oscillations. The times in between are when the reflector is moving, either towards or away from the beam splitter. From a phase point of view the relative phase between the

reflected beam and the reference beam is changing with the motion of the object an so they go in and out of phase several times. Once for each $\lambda/2$ the surface has moved.

Out of Alignment Without Object Movement It is very hard to get the beams perfectly aligned. If they are not perfectly coincident then this will result in fringes across the sensor. The second result in Figure 25 shows a 1 degree misalignment of each beam from the perpendicular (2 degrees in total) with the object is stationary. This shows one problem with misaligned beams - the width of the fringes. These fringes are constant in the 'z' direction (in the direction in which the beams should be traveling), but they affect the irradiance in the 'x' direction. If the sensing element is much bigger than this, or put differently if the misalignment is too great, then the fringes will average out over the sensor and changes in the irradiance due to the motion of the object mirror will tend to zero.

Out of Alignment With Object Movement Once the object is moving these fringes effectively oscillates as the phase of the signal beam either leads or trails that of the reference. It should be noted that if the sensor is smaller than one fringe then it will get almost the same signal as the perfectly aligned case, as shown by the bottom trace in Figure 25. The traces on the right show another problem with misalignment - change in the effective wavelength of the light. The wavelength which affects the signal is now the one which is resolved onto the 'z' axis. As there is an angle between the beam and the axis, this 'z' axis wavelength is now smaller than the original wavelength of the beam. This shows in the trace as the beginnings of another cycles of the higher frequency oscillations doe to the phase change of the signal beams. The top and bottom results shown in Figure 25 are with the same actual vibration amplitude.

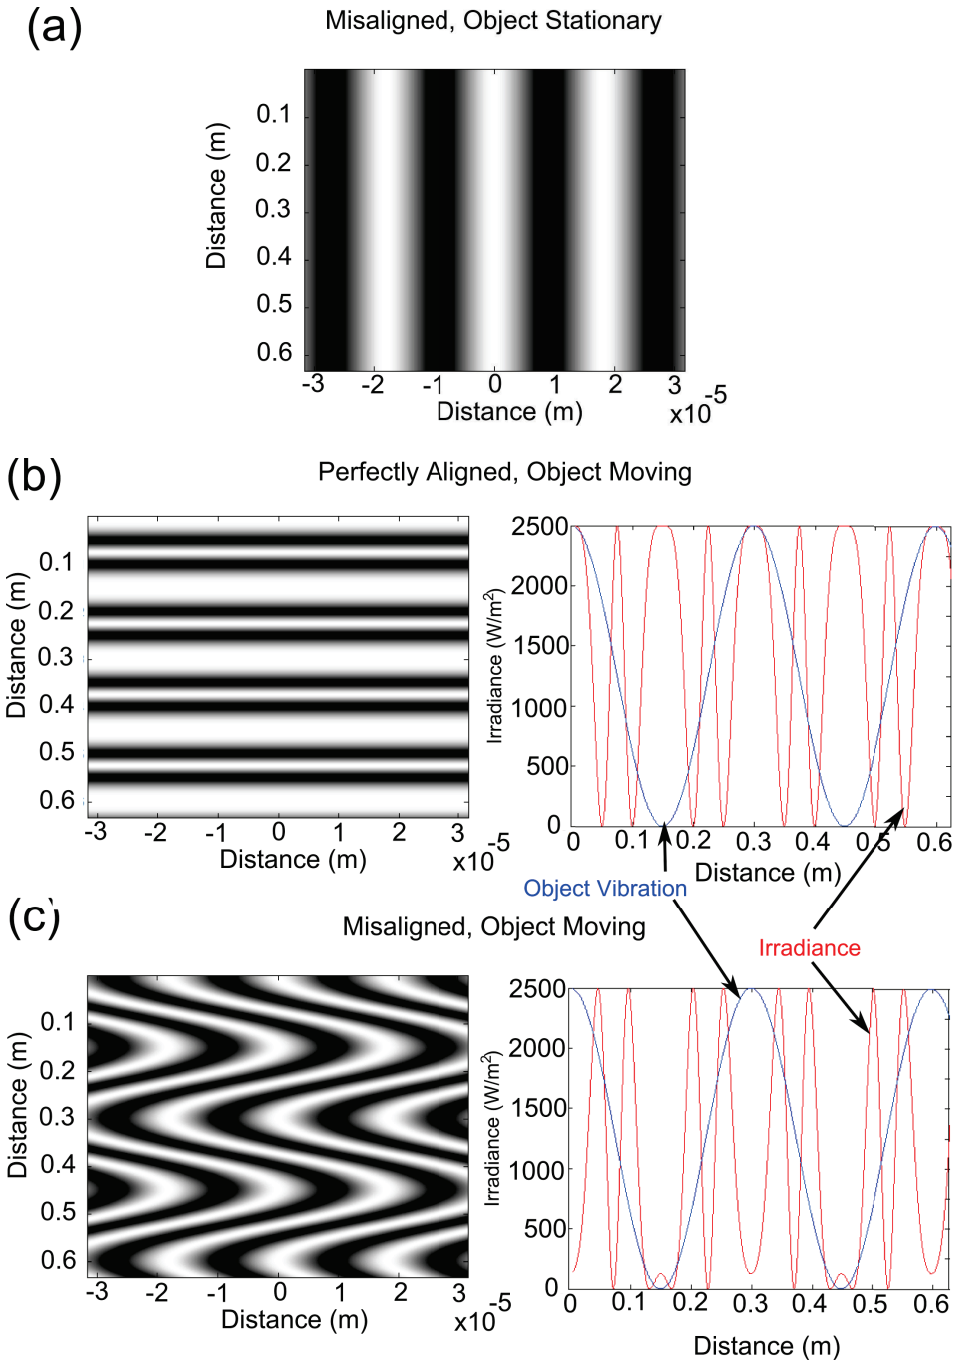

Figure 25: Simulations of the MI signal for different combinations of object motion (in blue) and beam alignment (in red). (a) shows the irradiance distribution when the interferometer is aligned but the object is stationary. The detector will measure a constant irradiance, varying between the maximum and zero depending on where the detector is on the fringe pattern. (b) shows the irradiance distribution (left) and detector signal (right) for an aligned interferometer with a vibrating object. The blue trace shows how the object is vibrating and the phase relationship between the irradiance signal and the vibration signal. (c) is the same setup as (b), but with the interferometer misaligned.

2.2.2 Standard CMOS Camera Results

Introduction A custom CMOS camera was used to measure the irradiance from small metal fingers called bimorphs (Part number: BM15015-06HC, stock number on RS Components Ltd: 285-784). Figure 26 shows a bimorph. These are small piezoelectric strips which are fixed at their base. The lower end of the strip is fixed and attached to electrodes. The top end then vibrates due to the oscillating voltage applied to the bottom. The signal came from a signal generator, with a transformer wired to increase the voltage to a level that would drive the bimorphs. When exposed to an oscillating voltage the strip vibrates at the same frequency.

The first experiment run was a vibrometry experiment based on previous work Dr. Diwei He [95]. This used a CMOS imaging chip to measure the frequency and amplitude of a vibrating bimorph. This experiment was intended to show the VHDL and computer software working in a laboratory setting and give an initial set of results which could be used as a starting point for comparison of future experiments and algorithms.

Experimental Setup This experiment was designed around a Michelson Interferometer (MI) as shown in Figure 27. An Field Programmable Gate Array (FPGA) board was used to control the experiment, with a computer running the processing. This is shown digramatically in Figure 28.

- For the MI, a nominally $5\,mW$ He-Ne laser (Melles Griot - 25-LHP-121-230, $3.8mW$ measured) at $632.8\,nm$, $1\,mm \oslash$ beam.
- MI setup also used a X10 beam expander (Thorlabs BE10M) in the laser arm to illuminate an area rather than a point.
- 32 x 32 pixel log CMOS camera. Each pixel had a photosensitive region of $25\,\mu m \times 25\,\mu m$. This was an in-house custom CMOS camera used for

measuring blood flow [96]. Each pixel was composed of a light sensitive part (a photodiode) with circuitry which produced a voltage which is logarithmically related to the irradiance. There was also some additional circuitry which performed the flow processing. This was however disabled, effectively resulting in a high-speed CMOS camera.

This camera had 1024 pixels, any one of which could be connected to the ADC through the de-multiplexers on the chip itself and the camera board.

- 1024 x 768 CCD camera (Thorlabs DCU223M) to check focus and positioning of the bimorph.

- Field Programmable Gate Array (FPGA) board with an Analog to Digital Converter (ADC) board, designed within the department. The FPGA was programmed to control the CMOS camera. It produced a digital output which, through the de-multiplexers, could address any pixel on the chip. Once a pixel had been addressed the output was connected to the ADC card and sampled at 625 $ksps$. The FPGA synchronized all these operations and then transferred the data to a computer using a USB connection.

- Software on the computer then processed the data to extract the vibration information. The software used was a mixture of C++ to gather the packets of information from the USB, and MATLAB to the actual processing.

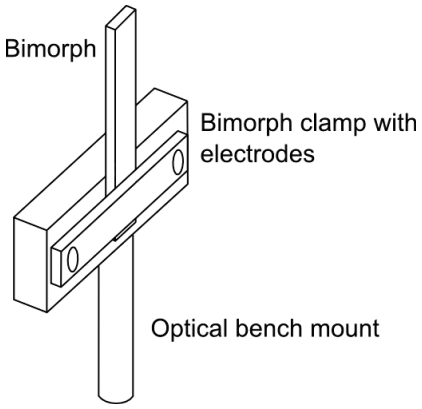

Figure 26: Bimorph and mount. This was the vibrating target used in the experiment.

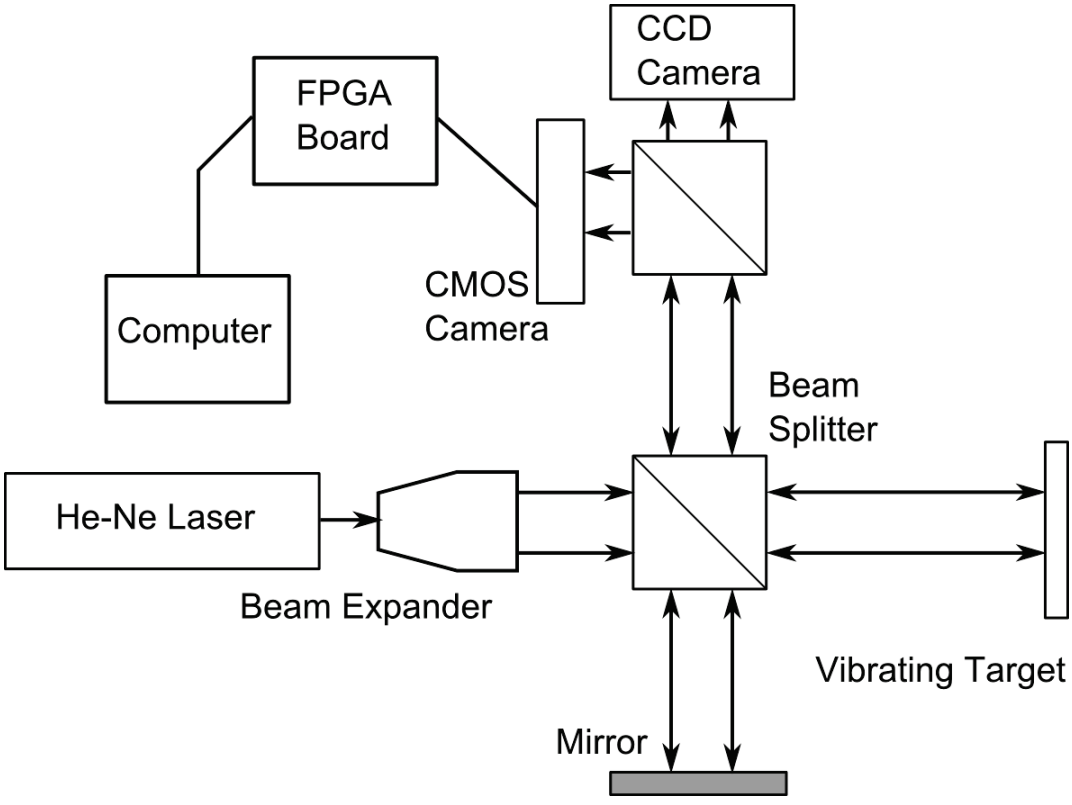

Figure 27: Diagram of the optics bench experimental setup. A He-Ne laser is expanded and used to illuminate a vibrating taget and a reference mirror through the use of a beamsplitter. The returning light is recombined and imaged by a commercial CCD camera and a custom CMOS camera. The CMOS camera was used to capture the optical interference signals and this data was processed to calculate the vibration amplitude and frequency.

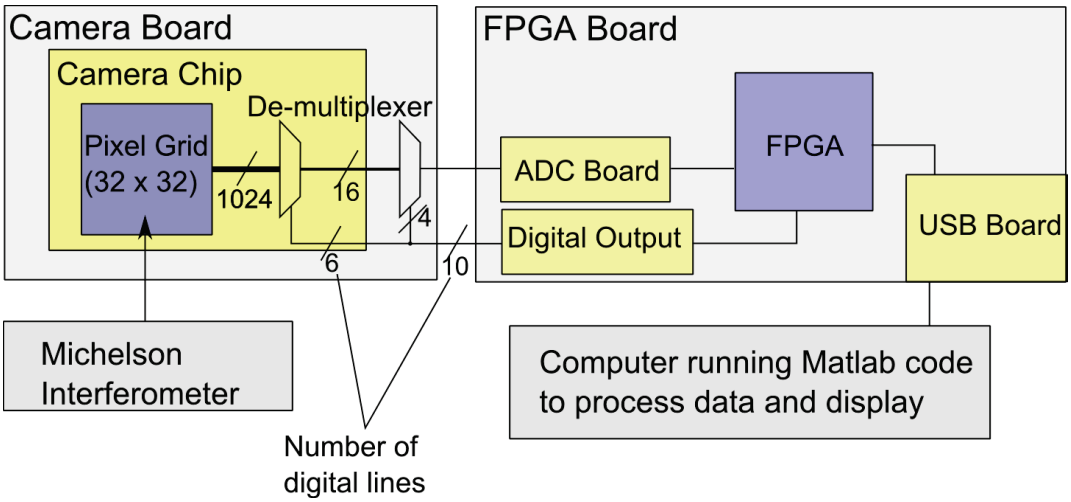

Figure 28: Schematic of the operation of the CMOS camera and FPGA board. The Michelson interferometer provided optical signals to the camera which was operated by the FPGA board. This board sampled the voltages from the different pixels and transfered the data to a computer for processing.

The operation of the system was as follows:

1. A pixel was selected by controlling the de-multiplexers.

2. Once selected, 1024 samples were taken of the voltage generated by that pixel. 1024 samples was chosen as a balance between the time taken to sample each pixel and the time data available for each pixel. The intention was to provide a real-time display and 1024 samples for each of the 32 x 32 pixels gave $\sim 1.7\,s$ for each complete frame. The time take for 1024 samples also gave a complete cycle of $610\,HZ$ and an small enough FFT bin size to resolve the vibration frequencies.

3. These samples were transmitted to the computer where they were buffered.

4. The next pixel was then selected and 1024 samples captured from it.

5. This procedure is repeated until all the pixels have been sampled.

6. There is now information for the whole camera buffered at the computer, where it is processed while the next set of camera information is buffered. The communications side of the software was all written in C++ and once the buffering was complete the data was then shifted to the Matlab engine for processing and display.

Results Figure 29 shows a picture of the bimorph after positioning. It shows the top of the bimorph and was taken with white light to remove the speckle effect. The processing method used is the same as that of Kongsavatsak et. at. [97]. The processing stages were to differentiate the signal from each pixel, then take the envelope using a Hilbert Transform. A section of the data captured by the FPGA board is shown in the top trace of Figure 30. This is the data for six pixels, as transmitted to the computer and shows one of the main drawbacks with CMOS camera chips: each pixel has a slightly different conversion between light irradiance and voltage. This is called fixed pattern noise and is constant for a given camera chip, set by slight variations during manufacture. These variations are unavoidable, but can be minimized through careful layout or, more commonly, calibrated out.

The top trace of Figure 30 shows that each pixel has a trace which oscillates, this results from the motion of the region of the object covered by that pixel and takes the same form as the results shown in Figure 25. The next trace in Figure 30 is after the DC fixed-pattern noise of each pixel has been removed, and the signal has been differentiated. The envelope of the signal now has the vibration information which is extracted by a Hilbert transform as shown in the bottom trace.

Finally a Fourier transform is taken of the bottom trace to find the vibration frequency and the amplitude of the vibration is found by counting the number of

zero crossings in the original pixel data and dividing this by the number of cycles of the vibration frequency.

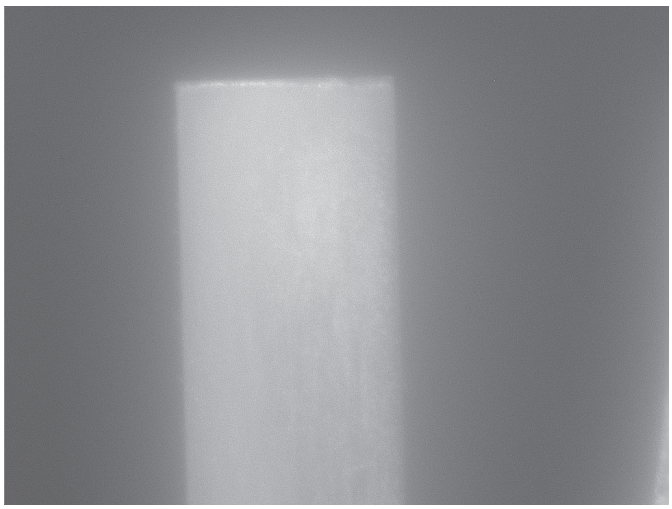

Figure 29: Off-the-shelf CCD image of the bimorph tip. This was taken in white light to check the position and focus of the bimorph.

Figure 31 shows the results from the computer processing. The bimorph was being vibrated at $1\,kHz$ with a $500\,mV_{pk-pk}$ sinusoidal voltage. Figure 31 shows the processed frequency and amplitude. This processed image also shows the effects of speckle on the processed images. The blank areas within the bounds of the bimorphs are where no speckles were present. The resolution of the image was set by the number of pixels of the CMOS camera, 32 x 32.

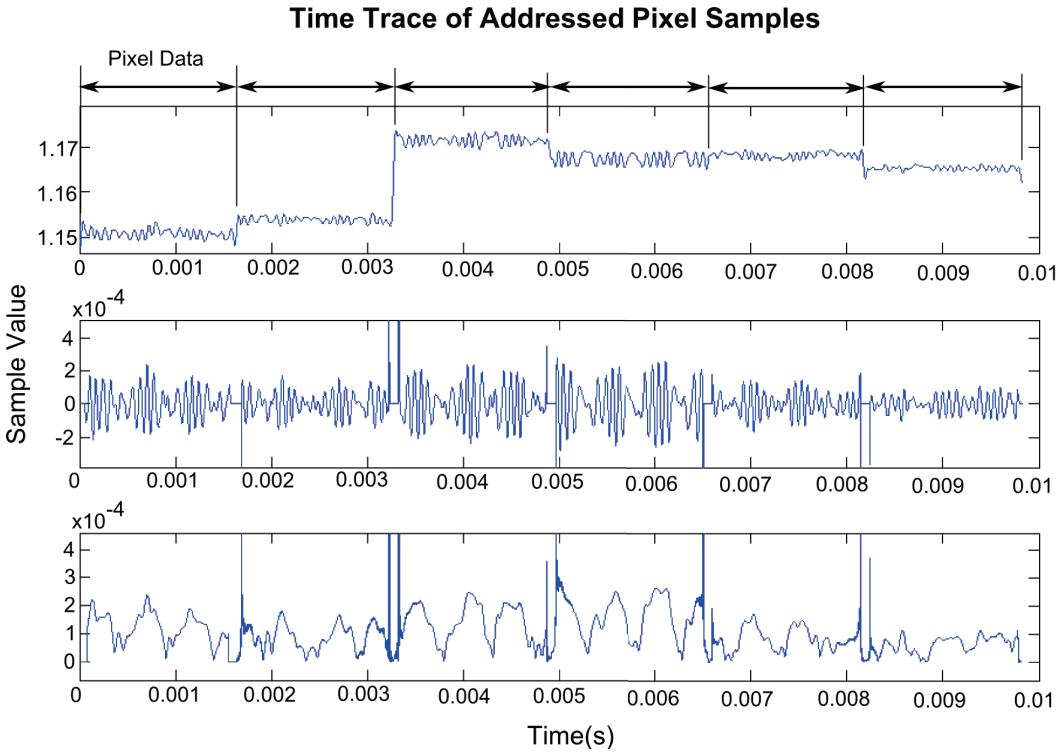

Figure 30: Section of the data captured by the FPGA board. The top trace shows a time trace of the samples from the CMOS camera. It shows a sequence of six pixels, each sampled for 1024 sampes. The middle trace shows the pixel data after filtering. The bottom trace shows the processed vibration signals detected by each pixel.

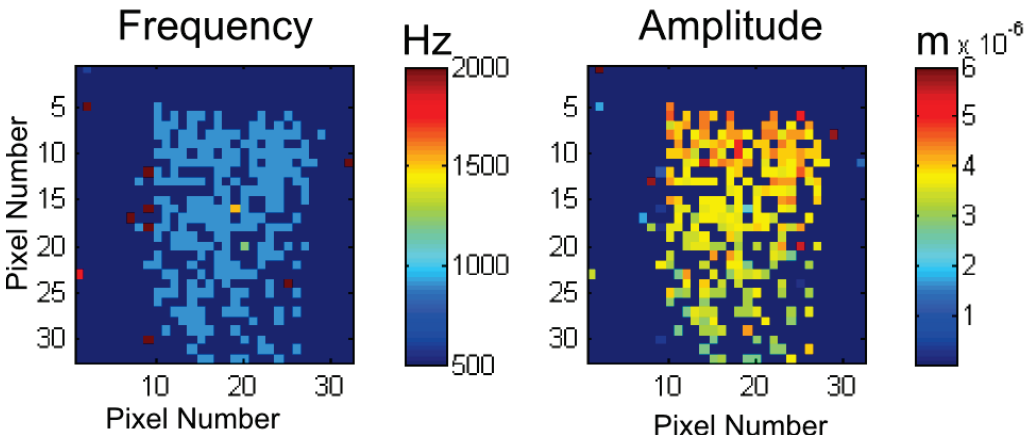

Figure 31: Results from bimorph experiment. These figures show the processed vibration frequency and amplitudes from the 32x32 pixels of the CMOS camera.

2.2.3 Polarizing Michelson Interferometer

As was mentioned in passing, the beams returning from the mirrors in the MI go through a second round of splitting before reaching the sensor. This results in half of each returning beam going to the laser, rather than the sensor. If a polarizing beam splitter is used in the MI then the beams going out to each mirror will be orthogonally polarized, see Figure 32. If this polarization is rotated before returning to the beam splitter, then all of the returning light will go to the sensor (discounting any losses).

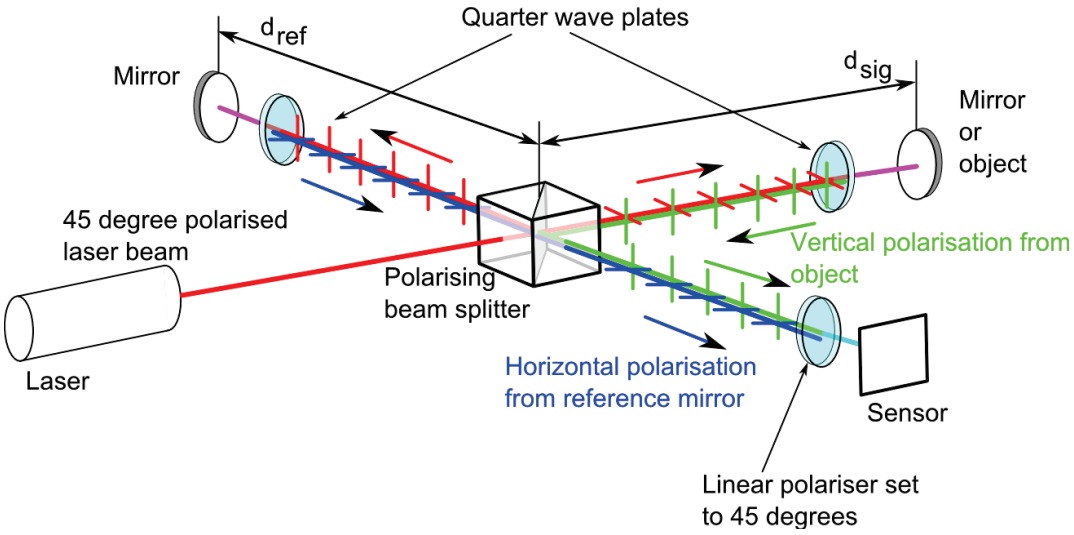

Figure 32: Michelson Interferometer using polarization.

This happens because a polarizing beam splitter reflects vertical polarization and passes horizontal. If this polarization is reversed then the previously reflected beam will change to a pass-through beam and go straight to the sensor. The previously passed beam will be change to a reflect-beam and will also go to the sensor. The polarization is rotated by 90 degrees through the use of a quarter wave plate and a mirror. The quarter wave plate changes the polarization from linear to rotational. The mirror effectively reverses the direction of rotation when it reflects the beam. On returning through the quarter wave plate the polarization is

changed back to linear, but now shifted by ninety degrees. The initial polarization of the laser will control how much light goes into each arm or the interferometer. The figure shows a 45 degree polarized beam which will give each arm the same amount of light. Finally, in order for the beams to interfere they must have the same polarization. They are therefore passed through a linear polarizer set to 45 degrees. Both horizontal and vertical polarized beams have a component at 45 degrees so they will interfere at the sensor as before.

Although half the light is blocked by the linear polarizer, resulting in the same losses as the normal MI, the principle can be extended to more complex interferometers where there are multiple beam splitters and larger beam-splitting losses. Although not of benefit for the Michaelson Interferometer the polarizing principle is introduced here in the simpler system before extending it to the more complex Mach-Zehnder Interferometer, discussed in the next section.

2.2.4 Mach-Zehnder Interferometer

The Mach-Zehnder Interferometer (MZI) is a more complex device than the MI as it passes the two beams over entirely different paths [31, 87, 88]. The device is shown in Figure 33. It is composed of three beam splitters. The input laser beam is first split into two beams, one of which is reflected off a mirror to the final beam splitter. This forms the reference arm. The other beam is passed to a middle beam splitter which passes it to the object mirror. The light reflected back from the object is then reflected by this middle beam splitter toward the final beam splitter where it is combined with the reference arm and interferes at the sensor. The arm with the middle beam splitter is called the signal arm. The advantages of this setup are that the two beams can easily be treated differently, for example the reference beam can be frequency shifted through the use of a Bragg Cell as introduced in Chapter 1, Section 1.5.7. Appendix B gives a summary of the

alignment procedure for the MZI.

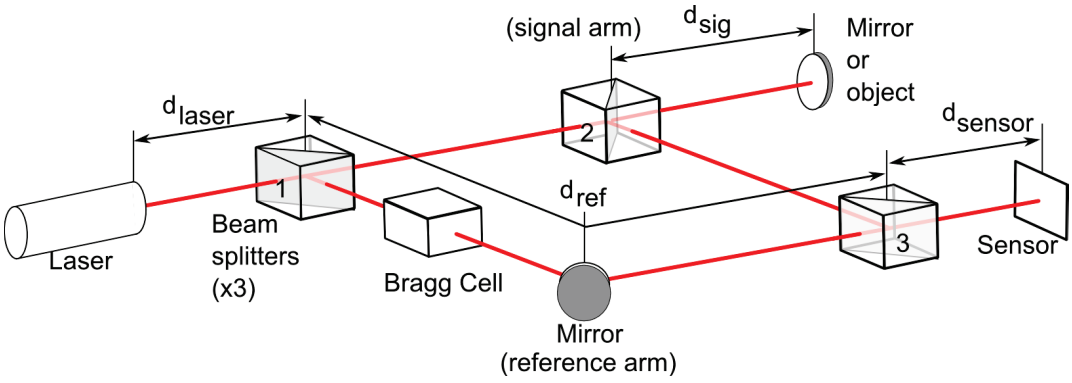

Figure 33: Mach-Zehnder Interferometer experimental setup.

Theory The theoretical derivation of the signal for the MZI is basically the same as the MI as they both have two beams which cover different paths to the sensor. The main difference is that in this case a Bragg Cell is inserted into the reference arm to frequency shift the beam. This has many advantages. Firstly it moves the irradiance signal away from low frequency noise [61], it gives a signal when the object is stationary and allows the direction of motion to be measured [98].

If the Bragg Cell frequency is f_B such that the resultant angular frequency is $\omega_B = 2\pi f_B$, then the irradiance Equation 29 becomes Equation 30.

$$I_s(t) = \frac{E_{sig0}^2 + E_{ref0}^2}{2} + E_{sig0}E_{ref0}cos\left[\omega_B.t + k\left(d_1 - d_0 - 2A_{obj}cos\left(\omega_{obj}t\right)\right)\right] \quad (30)$$

This can now be compared with a Frequency Modulated (FM) signal [92] which takes the form of Equation 31, where $K"$ is a constant with the units of radians per volt-second, A_m is the amplitude of the message signal $(sin\left(\omega_m.t\right))$, ω_B is the carrier frequency, and ω_m is the angular frequency of the message.

$$x_{FM}(t) = A.cos\left(\omega_B.t + \frac{A_m K^{"}}{\omega_m} sin(\omega_m.t)\right) \tag{31}$$

Comparing Equation 31 with 30 it can be seen that they are basically the same. For now d_1 can be set the same as d_0 without loss of generality as this difference in distance only adds a constant phase shift to the signal. In this case the Bragg Cell frequency is now the carrier frequency and this terminology will be used from here on. The cosine and sine functions only differ due to the original message signal. Equation 31 represents the FM equation which originated with a cosinusoidal message signal. This signal is integrated through the definition of angular frequency from phase, to become a sine. In the case of Equation 30 the original signal is also a cosine however this is displacement and, to match the derivation of the FM equation this would need to be differentiated to get velocity (now a sine). This velocity would then need to be changed to a frequency using the Doppler effect equations and then inserted into the carrier equation. This sine would become a cosine after integration as in the case of the FM derivation.

This only leaves the $2.k.A_{obj}$ and $\frac{A_m K^{"}}{\omega_m}$. Equation 19 showed how the ω_{obj} has been canceled given that $\frac{A_m K^{"}}{\omega_m} = \beta$ in the general form of the FM equation this substitution can still be made such that $2.k.A_{obj} = \beta$ where β is called the modulation index [92]. The bandwidth of the signal is therefore given by Equation 32.

$$BW = 2\Delta f + 2f_{obj} = 2f_{obj}(1 + \beta) \tag{32}$$

This shows that the bandwidth is related to both the vibration frequency (f_{obj}), and the vibration amplitude which is contained in β. The carrier frequency needs to be greater than half the bandwidth to keep all the Doppler signal in positive frequencies.

Simulation The simulation of the MZI follows the same format as that of the MI. The results are shown in Figure 34 for three cases: perfectly aligned with no movement, perfectly aligned with object vibration and miss-aligned with object vibration. If there is no change in the beam frequencies then the results would be the same as that for the MI, these simulations therefore show the results for when a Bragg Cell in inserted in the reference beam. This effectively adds a constant frequency difference between the two beams and give a similar result to the simulation shown in Figure 23. The beams are constantly going in and out of phase. The top simulation in Figure 34 therefore shows a sinusoidal variation in the irradiance. Once the object motion is added in, as shown in the middle plots of this figure. The frequency of this variation changes with the movement of the object. The Bragg Cell frequency can be considered a carrier which is frequency modulated by the frequency shift from the vibration of the object mirror. Finally, the bottom plots of Figure 34 shows what happens if the beams are miss-aligned - the grid made by the Bragg Cell variations changes angle to the sensor plane. Again this is a problem if the sensor is too small.

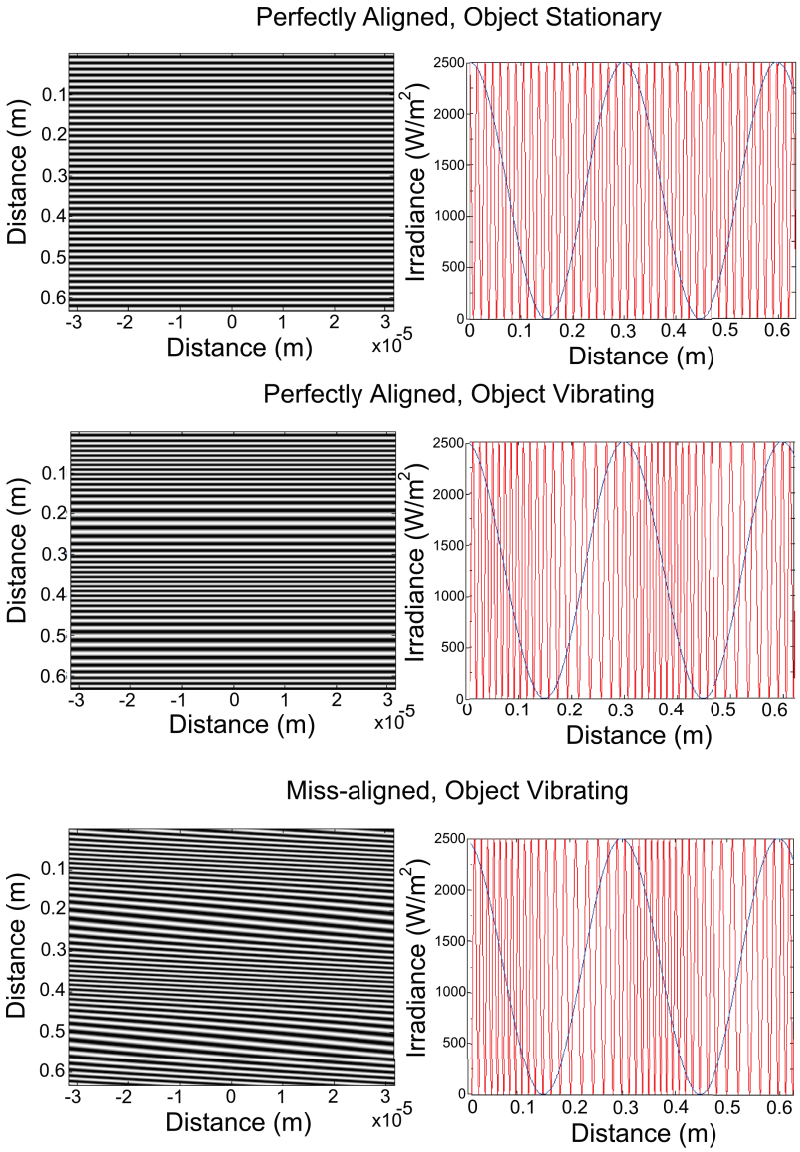

Figure 34: MZI simulation results for three cases of alignment and object motion.

2.2.5 Polarizing Mach-Zehnder Interferometer

The use of polarizing elements was introduced for the MI but is of more benefit to the MZI. Using polarization it is possible to direct light through beamsplitters without losses due to unwanted splitting of the beam. With the MZI the number of beam splitters means that losing half the light power at each quickly reduces the signal strength. Figure 32 shows a polarizing MZI. Beam splitter 3 can be

a non polarizing type which will result in two images. This can be useful if one is used for the main sensor and another for a CCD camera to image the target object normally.

The system of controlling the polarization is similar to the polarizing MI. A half wave plate is used to rotate the polarization of the reference beam and the signal beam again has its polarization flipped by a quarter wave plate. A linear polarizer is still needed to allow interference of the two beams at the sensor.

For full field imaging, lenses were inserted in the gap between beam splitter 2 and 3 to image the reflected light from the object onto the sensor. Beam expanders were inserted in the beams between beam splitter 1 and 2 and between the mirror and beam splitter 3.

It should be noted that the MZI, when used in this case, does not have matched path lengths. The distance d_{sig} is only covered by the signal beam. If this distance is longer than the coherence length of the laser used then the beams will no longer interfere. For the work done in this thesis the coherence length of the laser was of the cm's order of magnitude and care was taken to keep this distance as short as possible.

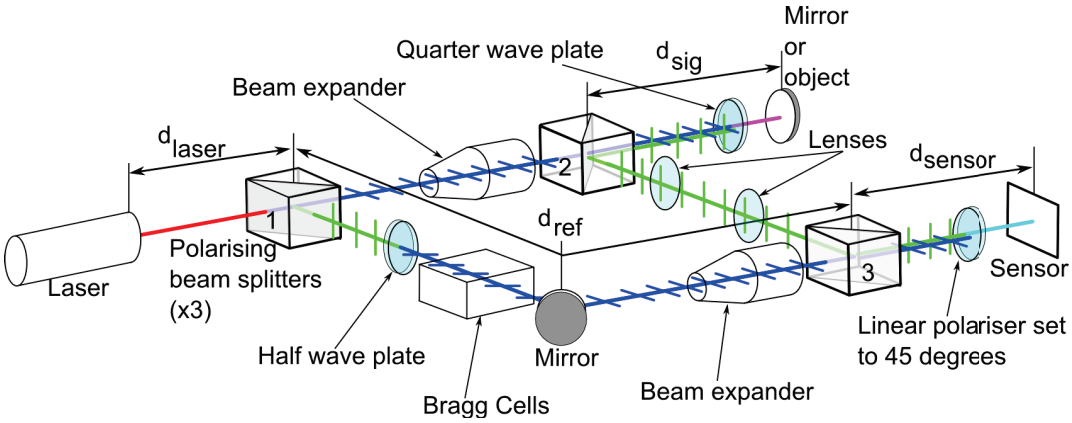

Figure 35: Mach-Zehnder Interferometer using polarization.

2.3 Simulating the Speckle Effect

2.3.1 Introduction

The speckle effect is an unavoidable effect of coherent light being reflected off a surface which is rough on the order of the wavelength of the light. Even if this effect is not exploited by a system, it will be present and can cause problems. The simulation was written in Matlab. Given the nature of the problem it was possible to divide it down into separate threads of execution which could be run concurrently. This favors current processors which are almost all multi-core. The Matlab code was written to split the field of simulations in two such that they could be processed simultaneously.

2.3.2 Simulation Setup

The simulation will take the form of a dual lens optical system and the simulation setup is shown in Figure 36. In this figure, A is the basic setup with two lenses with three planes, the object plane is the object being imaged, the Fourier plane is an intermediary plane which forms the input to the second lens and finally the sensor plane where the image of the object is produced. B is the first stage of testing the simulation model: a single point source. This should give the Point Spread Function (PSF) of the system on the sensor plane. C is the second stage of testing and is when a plane wave illuminates the system. This should give a plane wave as the output.

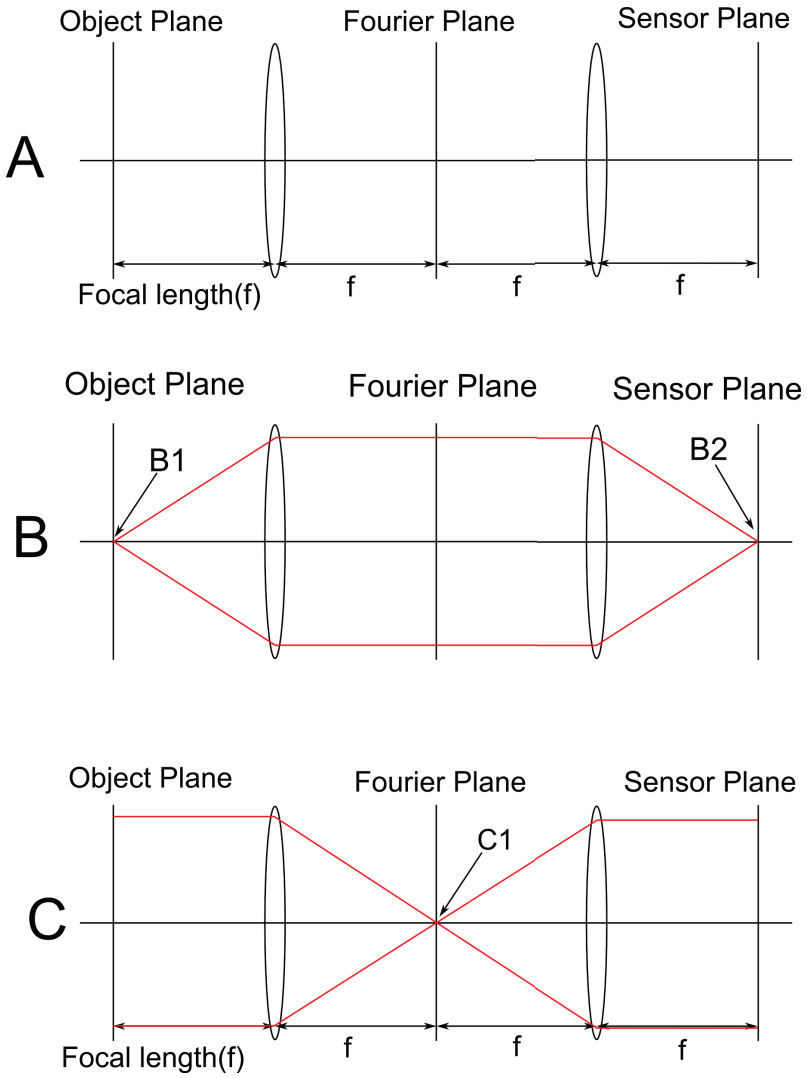

Figure 36: Speckle simulation setup. A - basic setup, B - testing the simulation with a point source, and C - testing the simulation with a plane wave.

2.3.3 Simulation Theory

The simulation is based on the Fourier model of a lens. It is split into two stages. The first creates the Fourier plane as the Fourier transform of the object plane, and the second creates the sensor plane as the Fourier transform of the Fourier plane. It can be shown that for an object plane a certain distance (d) in front of a lens, the image created one focal length behind the lens is given by Equation 33 [99].

$$U_f(u,v) = \frac{exp\left[i.\frac{k}{2.f}\left(1-\frac{d}{f}\right)(u^2+v^2)\right]}{i.\lambda.f}\iint_{-\infty}^{+\infty} U_i(\xi,\eta)\, exp\left[-i\frac{2.\pi}{\lambda.f}(\xi.u+\eta.v)\right] d\xi.d\eta \tag{33}$$

Where:

- k is the wavenumber $2\pi/\lambda$ where λ is the wavelength of the light used.
- f is the focal length of the lens
- u, v are the coordinates on the object plane
- ξ, η are the coordinates on the Fourier plane
- U_i is the optical distribution at the object plane. This is the illumination of the object and the real part is the distribution of brightness while the imaginary part represents the distribution of phase across the object.
- U_f is the optical distribution at the Fourier plane.

This equation does not take into account the physical aperture of the lens, which can be modeled as a round hole in an infinite opaque plane and adds a pupil function to the integral. The need for a pupil function is shown in Figure 37. Using a geometric optics approach, and working backwards from the Fourier plane, a single spot on the axis would produce the bounding rays shown in green. Effectively the lens would shape the spherical wavefront from the point into a plane wave traveling along the horizontal axis. The light from this point completely covers the object and so, working from object plane to Fourier plane, all points on the object are needed for this point on the Fourier plane.

Now, considering a point off-axis as shown by the red rays. The lens still transforms the spherical wavefront from the point into a plane wave, but now it is traveling at an angle to the horizontal axis. The light no longer covers the object

so, working in reverse again, not all points on the object are needed for this second point.

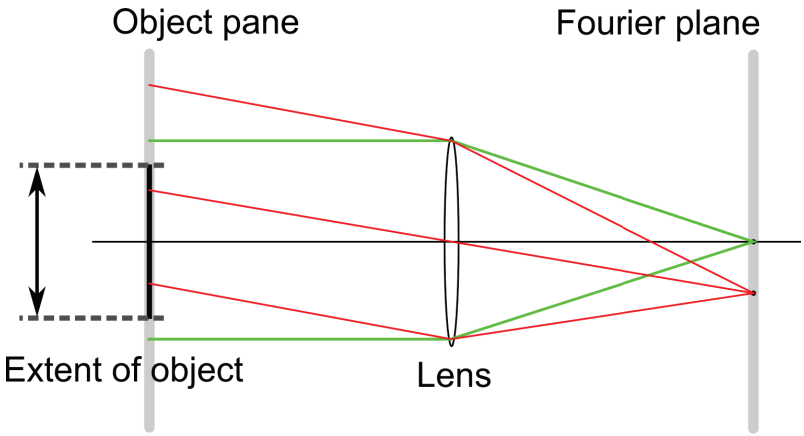

Figure 37: Demonstrating the effects of the physical aperture of the lens.

Figure 38 shows this in three-dimensions. The basic idea behind the pupil function is that the point on the Fourier plane is projected onto the object plane, and then a theoretical circle is placed around this point on the object plane. If a point on the object plane is within this circle then it will be used for the calculation of the point on the Fourier plane. The pupil function takes the form of Equation 34 where r is the radius of the lens, d is the distance from the object plane to the lens and f is the distance from the lens to the Fourier plane. u, v, ξ and η are as shown in Figure 38. If Equation 34 evaluates true (1) then the point at (u, v) should be used in the calculation of the optical distribution at (ξ, η), and false (0) if not.

$$\left(u - \frac{d}{f}.\xi\right)^2 + \left(v - \frac{d}{f}.\eta\right)^2 < r^2 \tag{34}$$

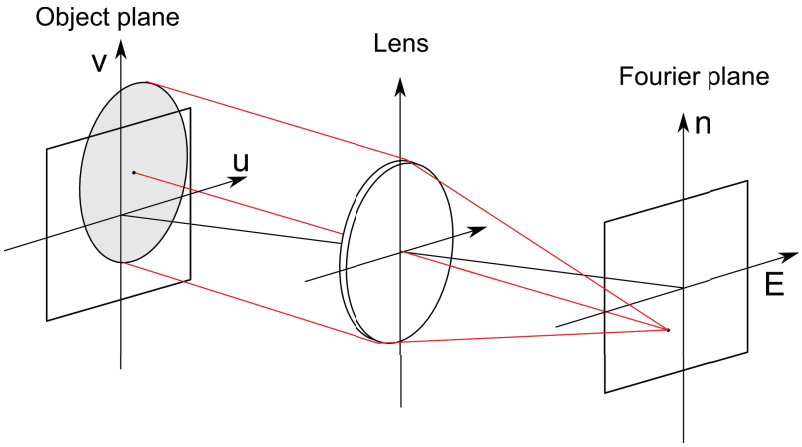

Figure 38: Demonstration of the effects of the physical aperture of the lens in three-dimensions.

When $d = f$ this simplifies the phase factor preceding the double integral in Equation 33, and including the pupil function gives Equation 35.

$$U_f(u,v) = \frac{1}{i.\lambda.f} \iint_{-\infty}^{+\infty} U_i(\xi,\eta) \,.P(u-\xi, v-\eta) \,.exp\left[-i\frac{2.\pi}{\lambda.f}(\xi.u+\eta.v)\right] d\xi.d\eta \tag{35}$$

This integral was processed by two nested 'for' loops and initially creates the Fourier plane from the object plane. This entire process is then repeated with the Fourier plane being the input plane and the output plane becoming the sensor plane, thus calculating the effects of the second lens.

It is also important to mention how the input optical distribution was changed to incorporate distance variations. This was used to simulate the optically rough surface. Looking at the equation for light, but ignoring the variation at the optical frequency, it becomes Equation 36. Here, E_0 is the real part of the distribution - how bright the object is, and the exponential is the imaginary part - the slight change in how far the light must travel from a given point (r).

$$E(r) = E_0.exp[i.k.r] \tag{36}$$

Taking this into account the input optical distribution becomes Equation 37, assuming that the brightness is uniform and where $r(u,v)$ is the height distribution across the object.

$$U_i(u,v) = E_0.exp[i.k.r(u,v)] \tag{37}$$

2.3.4 Simulation Testing

Before the simulation could be used to model the speckle pattern, it needed to be validated to check that it was responding correctly. This testing was done with some well defined problems. For each simulation a number of plots were produced. There were three plots for each image plane: the real distribution at that plane, the imaginary distribution at that plane and the irradiance at that plane. The simulations took approximately one hour for small regions on a single imaging plane up to six hours for full optical fields using two imaging planes.

Point Spread Function The first test was to check that the simulation produced the expected PSF. As well as demonstrating a well known theoretical result, the PSF demonstrates the minimum feature the lens system can resolve. This helps set the pitch for the different planes.

The object plane only needed to be a single point source. For this the object plane was set to a 3 x 3 grid with the central point set to 1 and the others set to 0. The point source needed to be smaller than the PSF (calculated to be $2.2\,um$) for the system to give the correct result so the width of the grid was set to half this value.

The size of the Fourier plane would be set by the object plane and the lens system. There are two ways of looking at this. From geometric optics, a single on-axis point source will be converted to a plane wave by the lens. Figure 39

demonstrates how a lens slows light and focuses light. Working from left to right a plane wave is focused to a point on the other side of the lens. Working back from that point the lens will shape the light into a plane wave, illuminating an area the size of the lens aperture. As such the Fourier plane would have a plane wave incident on it, shaped by the circular aperture of the lens - a disk of illumination the same size as the lens.

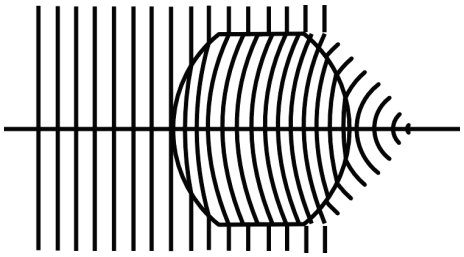

Figure 39: Lens delaying plane wave front.

The other way of thinking about this problem is as a Fourier transform. A single point theoretically contains all frequencies, and so would illuminate the whole of the Fourier plane. The shape of the lens would place some limits on this illumination (adding a low pass filter) and it would take the form of a disk of illumination the same size as the lens.

The sensor plane would need to image the PSF. For this the individual elements would need to be smaller than the PSF, and the grid would need to be big enough to contain the concentric circles which make up the function. With some trial and error runs it was set to $12\,um$ square with 70 points per side.

The simulation results are shown in Figure 40 which shows the sensor plane irradiance on the left and a slice through this on the right compared with the theoretical results. This shows that the simulation correctly reproduces the PSF of the lens system. Figure 41 shows a three-dimensional mesh of the PSF, showing the concentric rings around it.

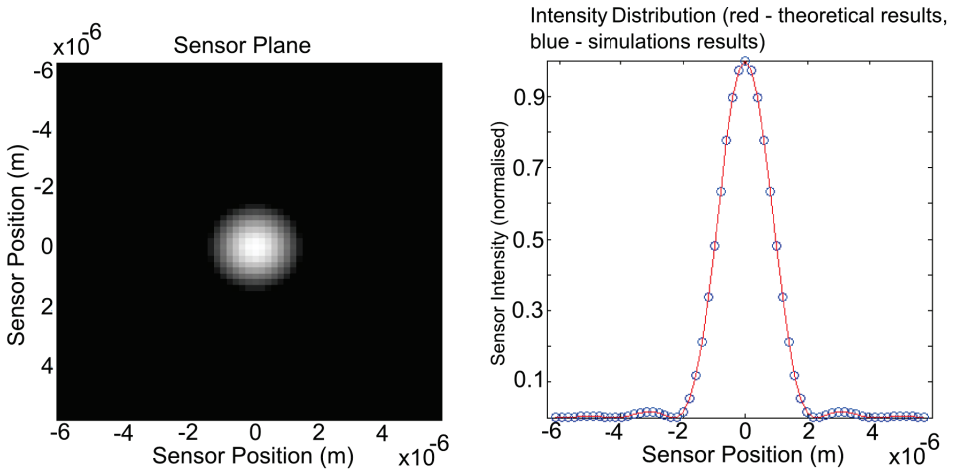

Figure 40: Intensity distribution at the sensor plane compared with theoretical results.

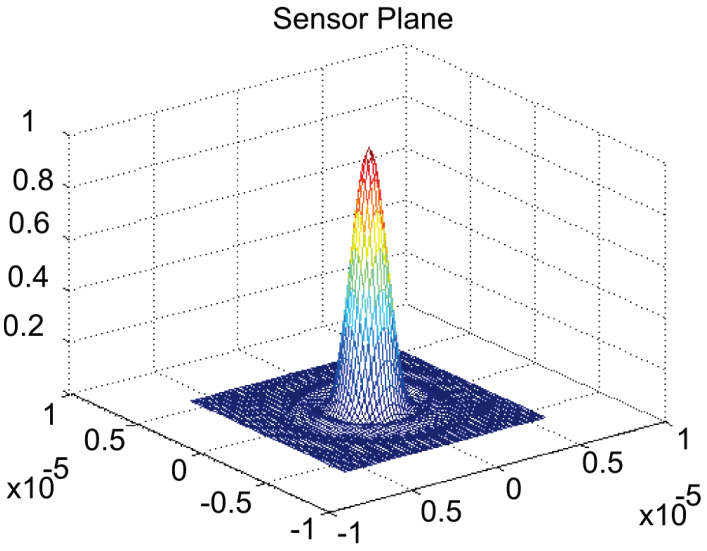

Figure 41: Mesh view of Irradiance at the sensor plane due to the PSF of the lens system.

Plane Wave This tested how the system would respond to a plane wave input. This would represent a DC input from a Fourier point of view and as such produce a single point in the middle of the Fourier plane, which would in turn produce an illuminated disk on the sensor plane. Figure 42 shows the results of this simulation. On the left is the Fourier plane which shows a PSF are the incoming plane wave

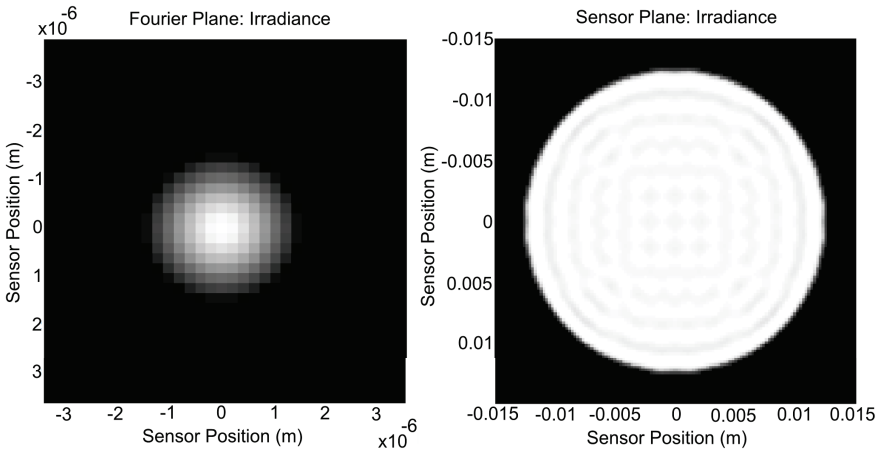

Figure 42: Fourier plane showing a PSF on the left and the sensor plane on the right showing an almost flat illumination.

is focused by the first lens. On the right is the sensor plane which shows an almost flat illumination controlled by the radius of the lens. The irregularities of the sensor plane image are due to some loss of higher frequency information by clipping the Fourier plane. This is a result of the simulation. If infinitely small points were used and the Fourier plane made large enough to contain all possible light from the lens then the sensor plane image would tend to a perfectly flat circle.

Diffraction Grating This test showed the system processing a grid input, also called a diffraction grating. This is a single step more than a plane wave and demonstrates how the Fourier plane shows the Fourier transform of the input. Figure 43 shows the results of this simulation. The Fourier plane was clipped as a horizontal grating only produces information in the horizontal direction, in this case three PSFs indicating the DC component of the grating in the middle with a spot on each side indicating the positive and negative frequency components of the grating. There are some other things which can be done to improve the speed of the simulations, as well as clipping the Fourier plane. None of the planes need have point sizes smaller than the PSF for the system. If the object is smaller than

the lens then it will set the size of the PSF. This can be described qualitatively by considering a plane illumination larger than the lens, say, a very big square. In this case, as already demonstrated the Fourier plane shows a PSF of the correct size, set by the lens. If that square is reduced the PSF will get larger, until at the limit the square becomes a point and the Fourier plane shows a plane wave, limited only by the lens aperture.

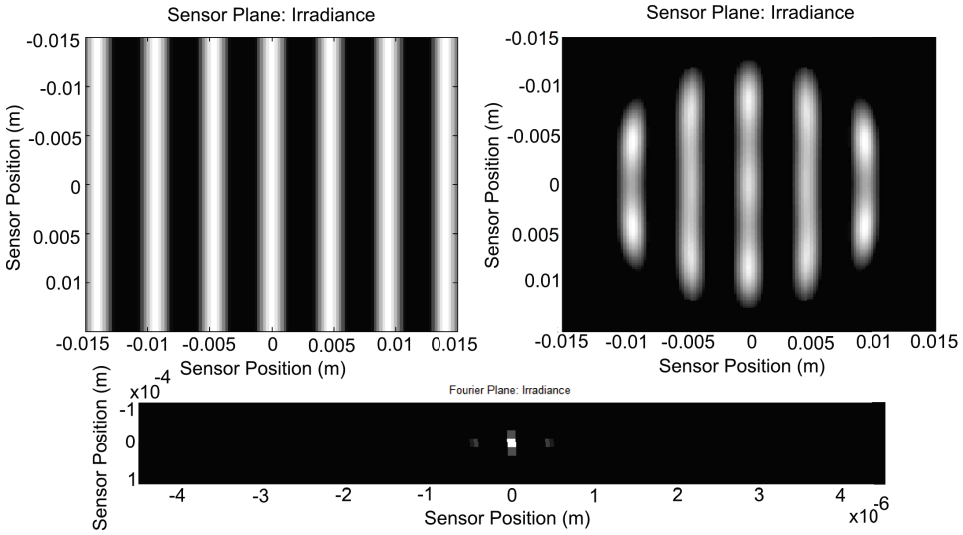

Figure 43: Object plane on the top left and is transformed by the first lens to the image on the bottom, the Fourier plane. This is then in turn transformed by the second lens to the sensor plane image shown at the top right.

Using a Fast Fourier Transform A much more efficient algorithm for simulating the effects of the lenses is a Fast Fourier Transform (FFT). This is an optimised form of the Discrete Fourier Transform which is itself a discrete form of the Fourier transform. It works hundreds of times quicker and is provided by Matlab. Equation 35 is effectively a Fourier transform if the plane are set at the focal lengths of the lens [99], and as such simulation of the distribution created by a lens can be done using the Fourier transform [100]. This can be demonstrated by comparing the FFT of a grating with that from the simulation setup used so far, as shown in Figure 44. Both the simulations used so far and the FFT give similar

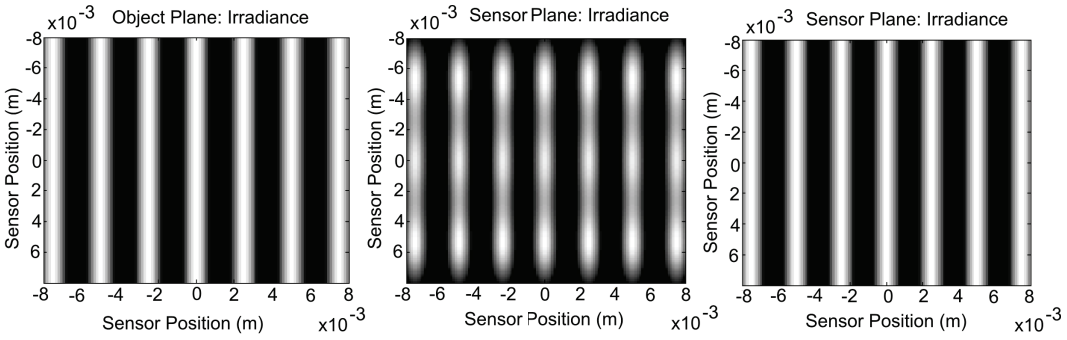

Figure 44: Comparison of the FFT of an image with the simulation used so far. Left is the original object plane grating. The sensor plane image is shown in the middle for the simulation and on the right for the FFT.

results. The FFT does not have the slight darkening around the edges which results from the finite aperture of the lens, however if the object is smaller than the lens aperture, as will be the case for the objects considered when simulating the speckle effect, this will have little effect [99].

2.3.5 Speckle Pattern

Now that the system had been validated with simple optical situations a simulated rough object was simulated. The object used for this simulation was stored for future use to allow comparison. It was generated using a flat probability function to set the height variations across it of maximum amplitude $2\,um$. The object was $10\,mm$ square and an FFT was run twice. First to transform the object plane to the Fourier plane, and then again to transform the Fourier plane to the sensor plane. The object and Fourier planes were padded with zeros to increase their size five fold. This caused the resultant Fourier transforms to have better resolution by effectively decreasing the input frequency [99, 101]. The results are shown in Figure 45. These results have the statistical relationship that the standard deviation of the irradiance is equal to its mean[86, 100], which one of the properties of a speckle field.

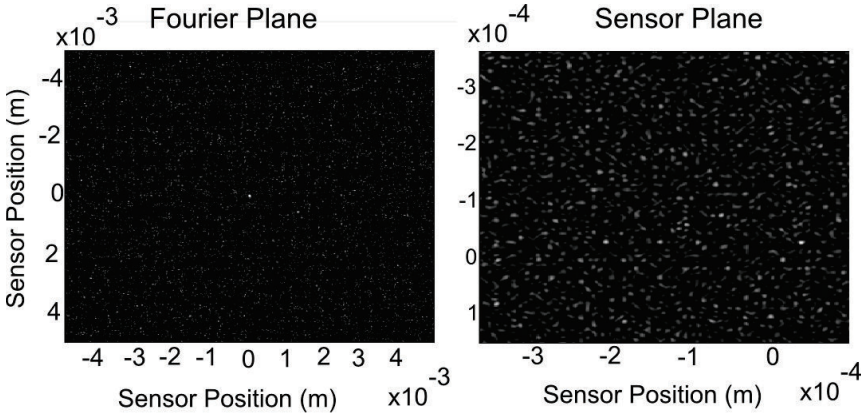

Figure 45: Speckle simulation showing the Fourier and sensor planes.

2.3.6 Conclusion

This section has shown the development of a simulation for the speckle effect. It started with demonstrating the PSF of a lens and how lenses can be used to perform a Fourier transform of the object plane light. It was then demonstrated that an FFT could be used to simulate the effects of a lens system and some results were shown for a simple grating. This simulation was then used to simulate an optically rough surface.

This model will be used to simulate the speckle effect in when the object is illumination with point sources in Chapter 6.

2.4 Proposed Pixel Architecture

2.4.1 Introduction

So far this chapter has described the interferometric setups and given simulations and experimental results for their operation. Chapter 1 showed that ESPI and LDV were basically when operating with speckles.

The difference is that LDV has always considered a single point to gather time dependent irradiance values and process these to extract the vibration informa-

tion. ESPI works by comparing full images and as such the processing is based around having these complete images. LDV processing is therefore better suited to full-field method of vibration analysis utilizing pixel level processing as the pixels can then run independently of each other, each processing the light falling on it. This independence is based on miniaturizing the LDV processing and applying it to each pixel of a camera and this seems the best solution to the full-field LDV problem. As this chapter has shown that the signal received from an interferometric setup with a Brag Cell in one arm will be a Frequency Modulated (FM) signal with the vibration as the modulating signal.

Each pixel will now need to process this FM signal, with circuitry that is as small as possible, and extract the vibration information. The processing needs to be small for two reasons. The first is so that there is the maximum fraction of the area of each pixel sensitive to light. The second is so that many pixels can be placed on the smallest possible silicon chip to keep the final camera economical. This section will now look at a number of different processing methods for FM signals.

Frequency Modulation (FM) is more commonly used in radio communications and there are many different methods of creating the FM signal which encode a given message signal, and then processing that signal to get back the original signal. The process of taking an FM signal and generating the message signal (in this case the motion of the reflecting surface) is called demodulation. This effectively requires a circuit which is sensitive to the frequency of the incoming signal. The ideal output being a linear function of input frequency.

2.4.2 Slope Detection

Perhaps the simplest method is to use a circuit which performs this task directly, using a specially configured filter response. This is called slope detection [92]

and uses a linear section of the frequency response of a bandpass filter. Using a balanced discriminator it is possible to double the linear length by centering one side of the circuit above the carrier frequency and the other side below. In practice this form of circuit does not have a wide enough linear range to be useful and is difficult to tune [92]. Additionally filters can be hard to make on-chip due to the size of capacitors and inductors.

2.4.3 Phase Locked Loop

The basic idea of the Phase Locked Loop (PLL) is to demodulate the FM signal through mixing with a source generated from the incoming signal. It is effectively a very narrow filter composed of a multiplier, a linear filter/ amplifier and a Voltage Controlled Amplifier (VCO). The control signal to the VCO is proportional to the message signal. A limiter circuit is needed before the PLL circuit to avoid distortion and the design of the filter must be such that it rejects sum frequencies from the mixer but passes all the significant parts of the message signal [92]. There are a number of commercial PLL circuits which can be bought nowadays however they are fairly complicated circuits and are again require filters.

Work has been done making completely digital demodulators based on PLLs [102]. The circuit designed by Pasternack and Whalin in 1968 demonstrated a fully integrated PLL with good lock, although using a relatively complicated circuit with multiple clocks. Hoeft and Zingmond [103] demonstrated an integrated demodulator in 1975 which also contained a PLL. This system, however, was composed of a single chip which needs external capacitors, inductors and resistors for operation. This allows these components, which take up the most space on a chip, to be left outside. It also allows some calibration of the chips operation.

Although some of these systems were designed for multiply channels, for example in the case of stereo FM signals [104], any circuit for a FF LDV camera

needs to be small and simple enough to be added to each pixel, which is precluded by the complexity of the PLL [105] as well as need for analog filters.

2.4.4 Phase Shift Discriminator

This form of demodulation used to be popular for FM demodulation, but has now been mostly replaced by the Quadrature Detector which is better suited to integrated circuits [92]. It is based on a doubly tuned circuit which generates a difference voltage which is zero for an unmodulated signal. The output is not necessarily linear with frequency and the circuit parameters need careful choice.

2.4.5 Quadrature Detection

Quadrature Detectors are based on mixing the original FM signal with a 90 degree phase-shifted version of the incoming signal. A high-reactance capacitor is used to initially give the signal the 90 degree phase shift. This signal is then applied to an LC-tuned circuit which is designed to resonate at the carrier frequency. A shift in the frequency of this signal then adds an additional leading or trailing phase to this signal. These two signals are mixed, and the output filtered with a low-pass filter, to give the original signal [92].

These circuits give good linearity and can be partly incorporated onto a chip, however the resonant circuit is normally left outside the chip, which would not be possible for an array of demodulators.

2.4.6 Digital Demodulation

There are a number of digital demodulation methods which work on sampled data. This means that the signal must be sampled first, however, once this has been done these types of demodulation can provide very low noise and very flexible means of demodulating FM signals [106, 107]. Unfortunately converting a signal

to a digital form within each pixel is not a practical approach do to the size of analog to digital converters on a chip.

2.4.7 Zero-Crossing Detector

This form of demodulator works by measuring the time between zero-crossings of the incoming signal. The information contained in an FM signal is held completely in the frequency of the signal and therefore the zero-crossings give a measure of this frequency [92]. A zero crossing demodulator can either count the number of zero-crossings over a long time compared with the period of the carrier, or measure the time for each zero-crossing.

Several papers have been written on the theory of Zero-Crossing (ZC) demodulation. In 1972 Voelcker [108] wrote a paper proof of how zero-crossings hold all the information required to demodulate an angle modulated signal, such as FM. He split the different methods of demodulation into three families. The first contained those that take the FM signal and first convert it to an AM signal before using an AM demodulation method to complete the demodulation. Slope detection is an example of this. The second family was those that use a form feedback and compare the incoming signal with a reference, often generated from the original signal. PLL systems are in this category. The final family was those that used pulse-conversion methods, ZC methods falling into this family. Voelcker showed that the a ZC demodulation scheme allows demodulation of signals where the carrier frequency is only just greater than half the message bandwidth. This means that the system can, theoretically, demodulate signal with a very wide frequency deviation. This wide-bandwidth FM signal is problematic for conventional demodulation schemes [109].

In 1977, Wiley et. al. [109] extended the work of Voelcker and showed a type of ZC demodulator which utilized an iterative approach to reduce the dis-

tortion introduced by a single ZC demodulation. This was, however, an entirely computational method.

Ray demonstrated a system based on a zero-crossing (ZC) detector [110] in 1980, although this system used the ZC circuit to control when the main signal was sampled. The theory was that if a ZC was detected, and a fixed time after that the main signal was sampled, then these samples would give a measure of how much the frequency of the signal had changed during that fixed time. This is shown in Figure 46. (a) shows how a sample at $2T$ is zero for an unmodulated carrier. (b) is a simple message signal. (c) is the modulated signal and shows how a sample from a fixed time after a rising-edge ZC is positive if the frequency has increased, zero if the frequency is unchanged, and negative if the frequency is lower than the original carrier.

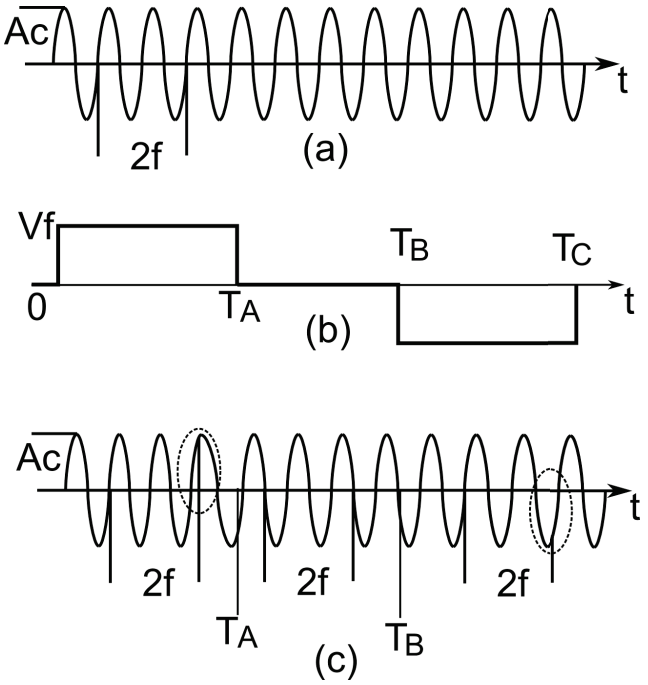

Figure 46: Figure taken from [110] showing how the demodulation system works. (a) is an unmodulated signal, (b) is a simple message signal, and (c) is the modulated signal.

In 1981 Wiley [111] published a mathematical approach to the use of ZC

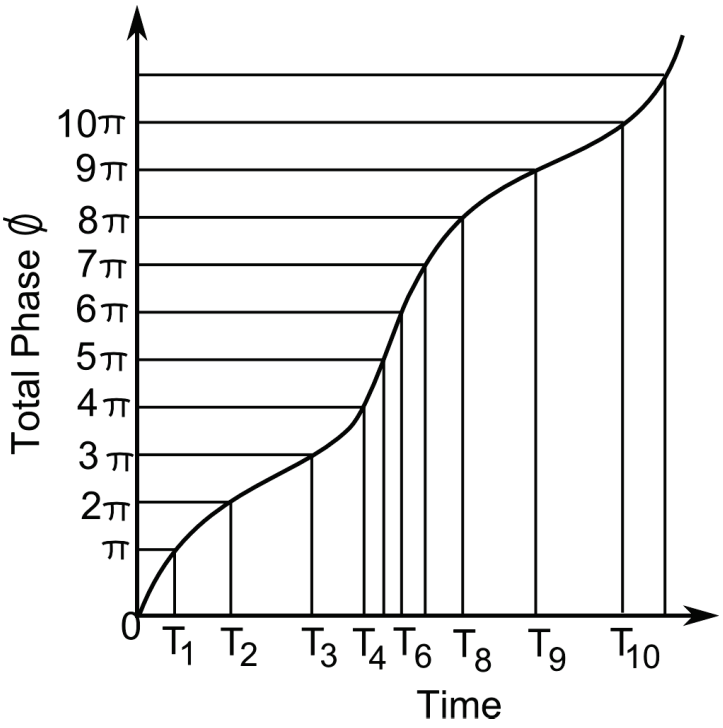

Figure 47: Total phase versus time for a frequency modulated signal, from [111].

demodulation for FM signals and showed how higher order interpretations could be used to improve the demodulation. Figure 47 shows how the zero-crossings of the frequency modulated signal occur at varying times. In this figure the vertical axis is the phase of the modulated signal, and the horizontal axis shows when the phase is a multiple of π. If the phase is controlled by a function such as $\omega_c.t - \beta.cos\,(\omega_m.t)$, and provided that ω_c is larger than the maximum change in frequency which can be introduced by $\beta.cos\,(\omega_m.t)$, then it can be seen that the phase is a linear function of time with a time dependent variation controlled by the message. This theoretical derivation required at least three ZCs and used a relatively complex computational scheme run on a computer.

Three years later in 1984, Ray used the ZC principle to produce a wide-bandwidth demodulator [112] based on an Intel microprocessor to allow a real-time ZC demodulator. This paper showed that a simplified real-time ZC demodulation method could be designed using a only integer mathematics and provided an

output which did not require the level of filtering of pulse-count ZC demodulators.

In 1998 the ZC method was applied to a Sagnac interferometer as the processing method [113]. This form of interferometer is used to measure rotation and the ZC circuit was demonstrated as a low-cost and low component count demodulation scheme. In more recent years research in ZC demodulation has turned to Bluetooth short-range communications systems with papers by Schotland and Jung [114, 115] suggesting the ZC demodulation could lead to lower implementation complexity. Schotland et. al. followed this up in 2006 with a hardware implementation of a ZC system based on a Texas Instruments DSP [116].

In 2009 Zhang et. al. proposed a method of demodulating wideband signals which used two ZC units to control the rest of the demodulation circuitry [117]. The demodulation scheme is shown in Figure 48. The top of this figure shows the incoming sinusoidal signal and the bottom is a block diagram of the demodulation circuit. The ZC units are used to trigger the integration of the incoming signal over specific times. This allowed the removal of Gaussian noise.

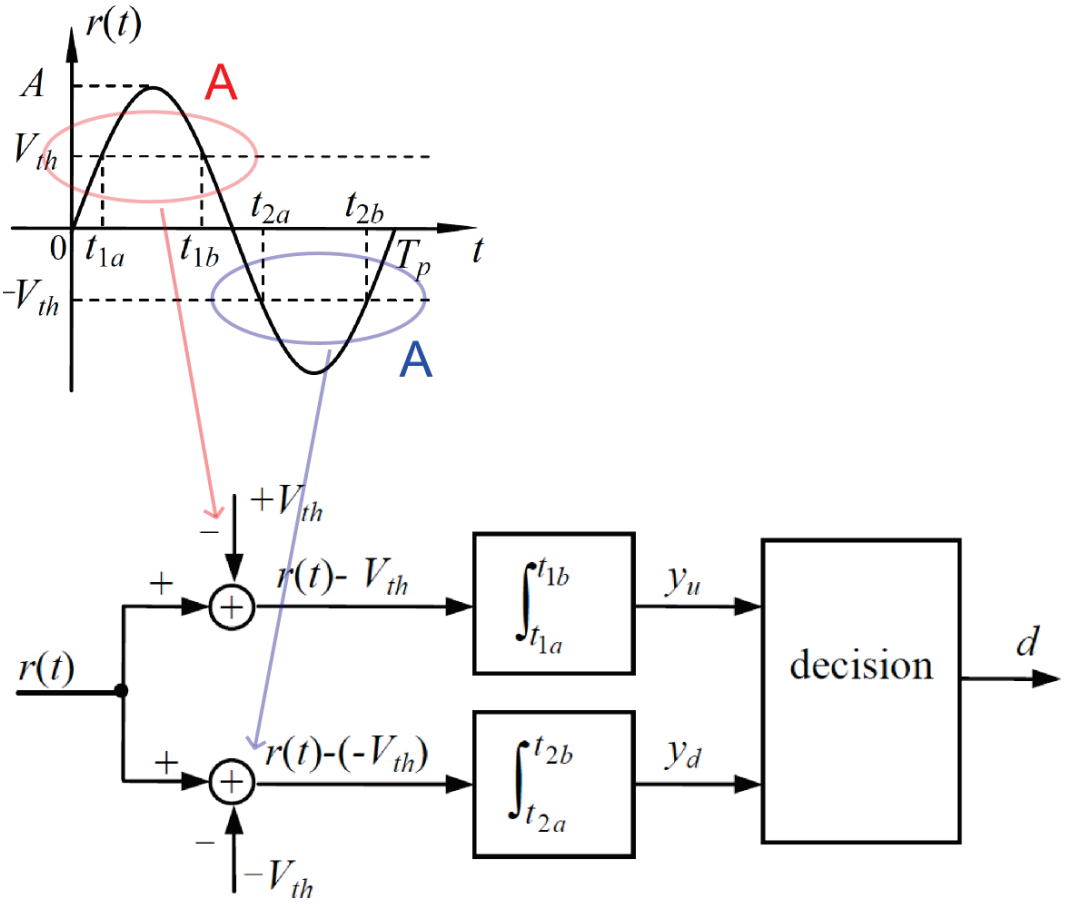

Figure 48: Use of two ZC units to control an ultra wideband demodulator, taken from [117]. A shows the points on the incoming signal which trigger the first ZC unit and B shows the points which trigger the second ZC unit.

ZC demodulation has the advantage that it is a relatively simple circuit which should be easy to miniaturized for use as a pixel level processing method. A ZC detector also does not require any outside control during operation and this means that the proposed camera would not need to be set or synchronized for a given vibration. Finally this method can handle high bandwidth message signals, compared with the bandwidth of the carrier, which allows a lower Bragg Cell frequency to be used for a given Doppler signal bandwidth and therefore reduces the requirement for a high speed photodetector for each pixel. There are some disadvantages with this method as well. The first is that the output is proportional

to the period of the signal rather than the frequency. As the period is $1/frequency$ and its the frequency which contains the information, this results in a non-linear conversion. The other disadvantage is that the signal is measured at fixed points, twice per cycle, and so effectively averages any frequency changes during the time between each zero-crossing. These will be addressed more in Chapter 3.

2.5 Conclusion

This chapter has given an overview of the LDV which is the focus of this thesis. Following on from Chapter 1 which gave a review of the field of LDV, this chapter has shown the different interferometer experimental setups and simulated the signals which would be produced when using them. The experimental setup used in Chapter 5 and 6 will be an MZI. This also helps to define the operating conditions for the pixel which will be important for the design of the pixel in Chapter 3. From the interferometer simulation results the type of signal which the LDV will need to process was derived, namely an FM signal. A number of different types of demodulation methods were then summarised including the zero-crossing method, which seems the most appropriate for use in at a pixel level to process the optical signals. This demodulation method will be used to design the pixel in Chapter 3.

3 Design of a Laser Doppler Vibrometry Pixel Implementation of Theory

3.1 Introduction

The previous chapters have introduced Laser Doppler Vibrometry(LDV) as a form of optical, non-contact, Non-destructive Evaluation(NDE). This method of NDE is a popular method for vibration analysis and has been chosen as the method investigated by this thesis.

In Chapter 2, the experimental setups for LDV were introduced including the Michelson Interferometer (MI) and the Mach-Zehnder Interferometer(MZI). The MZI was shown to overcome some of the limitations of the MI setup. A theoretical derivation of the signal produced by this interferometer was given and this showed that the signal produced by the interferometer could be considered a Frequency Modulated(FM) signal. The carrier frequency of this signal being provided by the use of Bragg Cells within the interferometer, and this carrier frequency modulated by the motion of the reflector. Given the form of this signal, a processing method: zero-crossing (ZC) demodulation, was then proposed for use in a custom CMOS chip.

Although the circuit developed for this project is referred to as "the chip", the chip itself was composed of two separate circuits designed at the University of Nottingham, of which the LDV pixel was one. This multi-project philosophy shares the cost of manufacturing CMOS chips. This chapter will now describe the design and simulation of a single pixel, implementing this processing method, and added to a chip called BVIPS2. BVIPS2 was designed within the Applied Optics Group at the University of Nottingham, primarily as a laser Doppler blood flow imaging chip, and the ZC pixel was placed on some spare room on the chip.

This chapter will start by describing the ZC processing method, deriving the

equation for the processing and simulating them. This section will then show the non-linear properties of the method, discuss when they could take effect and how they could be of benefit for real world applications. Next the design of the actual circuit will be covered showing how the theoretical demodulator was implemented on a CMOS chip. Finally the resulting processing circuit will be simulated to check the operation and range of operation of the various components which make up the circuit. The software program used to design the circuit is called Cadence. It also provides a number of tools to simulate the pixel's processing using electrical models of the components. A specific simulation which was used in this thesis was the Monte Carlo simulation. This was also provided by Cadence and designed to simulation the random process variation expected when the chip went through for fabrication. Small changes in the parameters of the components on a chip are unavoidable and will change the response of those components. The Monte Carlo runs a given simulation a number of times, changing the component parameters in a random way to match the expected variation due to fabrication. This simulation can be used to check approximately how many chips may fail due to the component variation. Another piece of software used to simulate the various parts of the pixel is Matlab. This is a generic mathematical simulation program based on matrix arithmetic. It was used to provide ideal simulations of various stages of the pixel's processing, excluding any electrical non-idealities of the actual circuit.

3.2 Processing Theory

Chapter 2 has derived Equation 30 for the irradiance of the light from the interferometer falling on the processing chip. The light sensitive part of the pixel is a photodiode and the information about the movement of the object is contained in the frequency of that photodiode's signal. This equation was shown to represent a frequency modulated signal with the carrier centered at a frequency set by the

Bragg Cells within the interferometer part of the LDV system. ZC demodulation is a form of FM demodulation (see Section 2.4) and provides a period to voltage converter that processes the Doppler signal to produce a signal with the same frequency as the vibration signal (and hence the same bandwidth) and an amplitude which is proportional to the amplitude of the vibration.

Equation 38 shows the ideal equation for a period to voltage converter. In this equation the output voltage(V_{out}) is proportional to the the time varying period(T) of the input signal through k_{pv}, which is the period to voltage conversion constant with units of V/s. As will be explained in Section 3.3 the actual relationship between the output voltage and the period is given in Equation 39 where V_S is the supply voltage.

$$V_{out}(t) = k_{pv}T(t) \tag{38}$$

$$V_{out}(t) = V_S - k_{pv}T(t) \tag{39}$$

Equation 30 Section 2.2.4 in Chapter 2 shows the expected input light irradiance falling on the chip. This signal is converted to a current by the photodiode and then to a voltage by a transimpedance amplifier. The photodiode provides a current(called the photocurrent, [118]) proportional to the light intensity with a constant of proportionality called the responsivity(R), units A/W. If I_{in} is the current through the photodiode it can be described by the equation $I_{in}(t) = I_0 + R \cdot P_{opt}(t)$ where I_0 is the leakage current through the photodiode, called the dark current and P_{opt} is the optical power of the light falling on the photodiode [119, 120].

The processing circuit will be covered in more detail in Section 3.3.1. This section is interested in the frequency to voltage conversion and is therefore only

considers the AC component of the input signal as this contains the Doppler information.

Equation 40 shows the part of the photocurrent due to the Doppler signal. E_{sig0} and E_{ref0} are the amplitudes of the signal and reference beams within the interferometer, ω_B is angular frequency due to the shift from the Bragg Cells (The Bragg Cell frequency will be called the carrier from here on for simplicity and in keeping with the FM nature of the signal), A_{obj} and ω_{obj} are the amplitude and frequency of the vibration, and $k = 2\pi/\lambda$ is the wavenumber for light of wavelength λ. $d_1 - d_0$ is the different in path lengths between the reference and signal beams when the object is stationary. This can be included in a single constant $\phi_0 = k\,(d_1 - d_0)$ and therefore simply adds a phase shift to the modulated signal. This can be set to zero as the Doppler signal modulates the phase (frequency) irrespective of the starting phase.

The zero-crossing method works using the times the input signal crosses its mean, rather than $0\,V$ as the name suggests. If the time varying part of the input signal is considered then Equation 41 shows how its phase($\phi\,(t)$) varies with time. The k has been expanded. Considering two adjacent falling-edge zero-crossings at t_1 and t_2. The phase of the input signal would have to have changed by 2π as shown in Equation 42. If the bandwidth of the Doppler signal ($\frac{2\pi 2A_{osc}}{\lambda} cos(w_{osc}t)$) is much less than that of the carrier then it can be considered not to have changed much between t_1 and t_2 and this gives Equation 43. In this equation, δf is the Bragg Cell frequency, and t is a time value between t_1 and t_2, and δf is the carrier frequency. Given that the function does not vary much between t_1 and t_2 the exact value of t between these two values has little effect.

$$I_s\,(t) = \frac{E_{sig0}^2 + E_{ref0}^2}{2} + E_{sig0}E_{ref0}cos\left[\omega_B.t + k\,(d_1 - d_0 - 2A_{obj}cos\,(\omega_{obj}t))\right] \quad (40)$$

$$\phi(t) = \omega_B.t + \frac{2\pi 2A_{osc}}{\lambda} cos(w_{osc}t) \tag{41}$$

$$\phi(t_2) - \phi(t_1) = 2\pi \tag{42}$$

$$2\pi = \delta f\,(t_2 - t_1) + \frac{2\pi 2A_{osc}}{\lambda} cos\,(\omega_{osc}t) \tag{43}$$

The period of the input signal would be the time between the zero-crossings and therefore given by Equation 44. Substituting Equation 46 into Equation 43 gives Equation 45. Rearranging to show how the period of the input signal changes with time gives Equation 46 and this would be the period used in Equation 39. T_c in Equation 46 is the period of the carrier. This equation shows that the period of the input signal varies sinusoidally about the period of the carrier.

$$T = t_2 - t_1 \tag{44}$$

$$1 = T.\delta f + \frac{2A_{osc}}{\lambda} cos\,(\omega_{osc}t) \tag{45}$$

$$T(t) = T_c \left(1 - \frac{2A_{osc}}{\lambda} cos\,(\omega_{osc}t)\right) \tag{46}$$

Substituting Equation 46 into Equation 39 gives Equation 47 and this shows the output voltage is a sinusoid with the same frequency as the vibration. The amplitude is proportional to the amplitude of the vibration, but is also affected by the k_{pv} and the carrier.

$$V_{out}(t) = V_S - k_{pv}T_c + k_{pv}T_c \frac{2A_{osc}}{\lambda} cos\,(\omega_{osc}t) \tag{47}$$

In practice k_{pv} would be adjusted such that a given constant frequency (and

therefore constant period) would result in a given output voltage. Provided the input frequency stayed at this value the output would be a DC voltage. Using this idea of an input frequency producing an output voltage to set the value of k_{pv} it is possible to show the operation of the algorithm as a plot of input frequency vs. output voltage as shown in Figure 49. The three curves were produced by setting k_{pv} such that a certain frequency($10k\,Hz$, $20k\,Hz$ and $40k\,Hz$) would give an output of $1.65\,V$. This figure shows that the processing has a bandpass characteristic and changes to k_{pv} affect the minimum and maximum frequency accepted by the processing as well as the gradient of the curve.

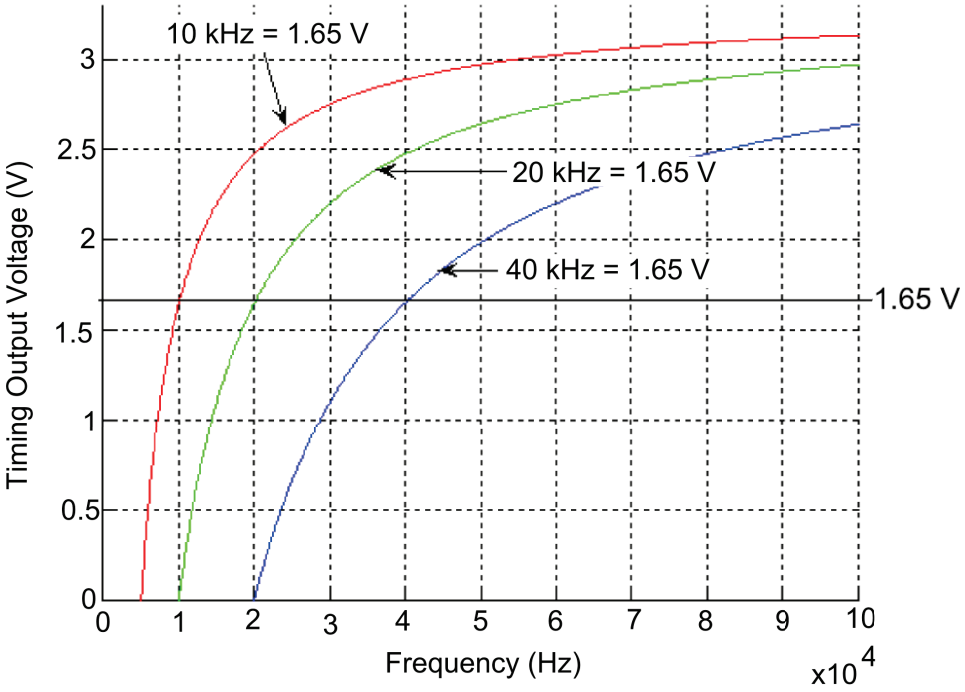

Figure 49: Output voltage vs. input frequency for various values of k_{pv}. $V_S = 3.3\,V$ with the input frequency swept from $0\,Hz$ to $100k\,Hz$

The gradient changes how much distortion there is in the output signal. This distortion is due to the circuit implementing a period to voltage converter not a frequency to voltage converter so there is a $1/f$ relationship. If the changes in frequency are small, as assumed in Equation 43, then the non-linear gradient

will have little effect. For larger signal bandwidths the response of the pixel processing will tend to compress those frequencies higher than the carrier into a smaller voltage range and expand lower frequencies.

3.2.1 Discussion

The previous sections have shown the theory behind the operation of this LDV system. The optical part of the system is an MZI with one of the beams frequency shifted. This provided an input signal to the chip processing which is then demodulated to produce a voltage related to the vibration of the object by outputting a voltage which is proportional to the input signals period. The non linear response of this demodulation method was then shown as the information in the input signal is encoded in the signal's frequency, not period.

This response could be corrected with appropriate additional processing. There is still an advantage over standard camera chips as the signal now has only the vibration bandwidth rather than the wider bandwidth signal contained in the Doppler shifted light. The external processing would also be relatively simple once the characteristics of the circuit's response had been calibrated. A different method of processing could also be used, however zero-crossing demodulation uses mostly digital components and lends itself to the size constraints of pixel level processing as covered in Chapter 2.

For some applications this characteristic of the processing would either have no effect or would be beneficial. If the processing only needed to find the frequency of a vibrating object, and this object vibrated at a single frequency, then the non-linearity would have no effect. Also, in practical applications of LDV the desire is often to find the amplitude of vibration for a given frequency. An object would be vibrated at that frequency and the camera used to map the amplitudes of the vibration across the object [77, 78]. Under these circumstances the response of the

processing provides a better dynamic range than a linear system. These systems tend to use a sinusoidal vibration. If the processing was set such that the output voltage was $V_S/2$ for an input at the carrier frequency, say $40k\,Hz$, then during vibration the voltage would swing either side of this value.

The situation for small vibration amplitudes is shown in Figure 50. This figure shows the frequency to voltage conversion characteristic of the processing. Superimposed on this characteristic plot is a possible input, sinusoidally modulated, frequency modulated signal. This input would be created by an object vibrating with a sinusoidal motion. This object would reflect light with a Doppler shifted frequency proportional to the velocity of the object. The input signal is therefore a sinusoidal signal on the frequency axis showing how the frequency of the input varies due to the modulation(motion of the object). The processing would then produce an output voltage related to this input frequency by the conversion characteristic of the processing. Figure 50 shows the input signal as well as lines mapping the mean input value (the carrier frequency) and the peaks onto voltages. This shows the range of output voltages from the processing for the given input. As can be seen the upper peak, above $V_S/2$, is smaller than the one below $V_S/2$. Here the lower peak could be used and would provide good sensitivity.

Figure 51 shows a similar case to Figure 50, but for a larger object movement, and hence greater velocity for a fixed object vibration frequency. This greater velocity would introduce a greater frequency shift into the reflected light and thus a greater shift in input frequency away from the carrier. Again there are lines mapping the mean, higher and lower frequency values for the input. In this case the lower peak would clip as it is limited by the lower supply rail($0\,V$), however the peak above $V_S/2$ would respond less, and provide good dynamic range.

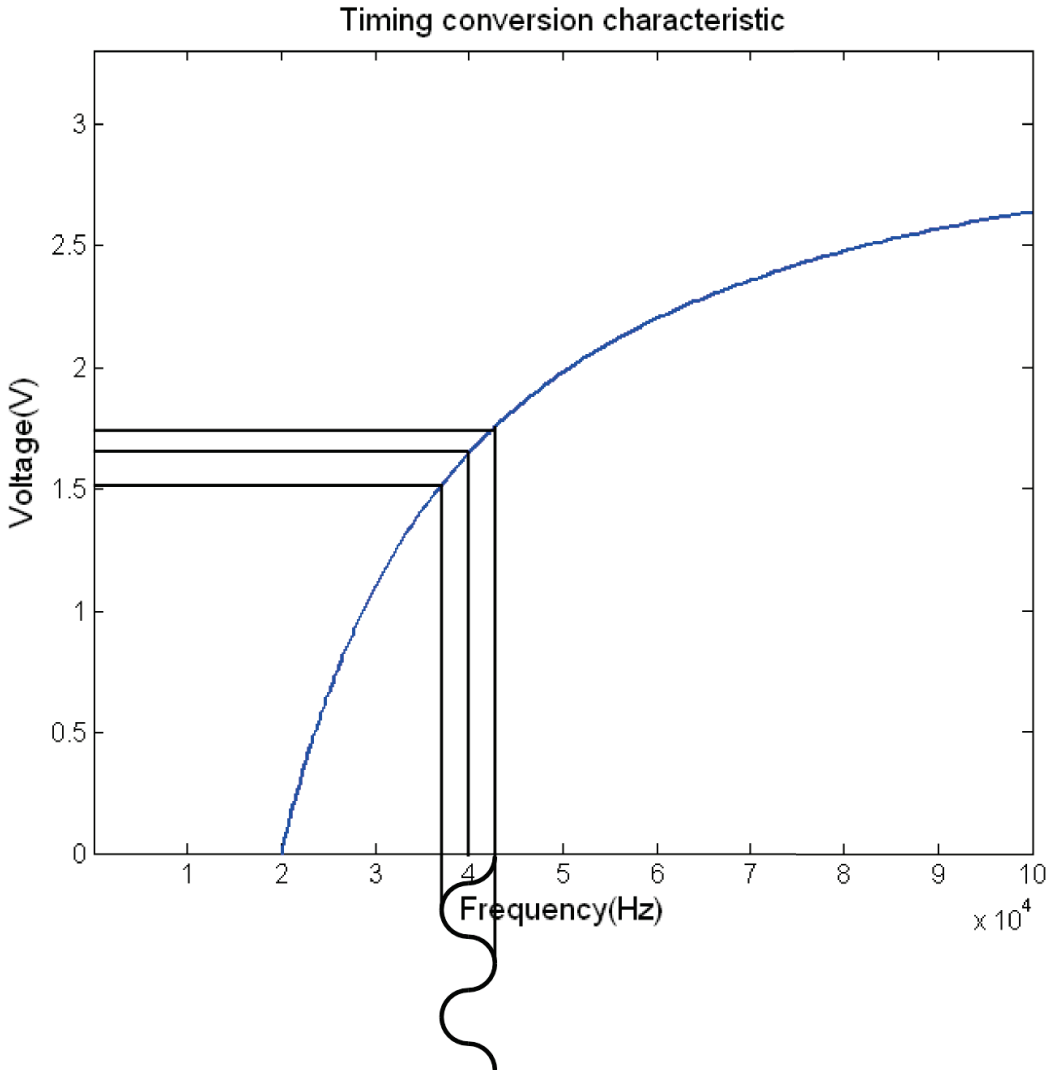

Figure 50: Processing response for small vibration amplitudes

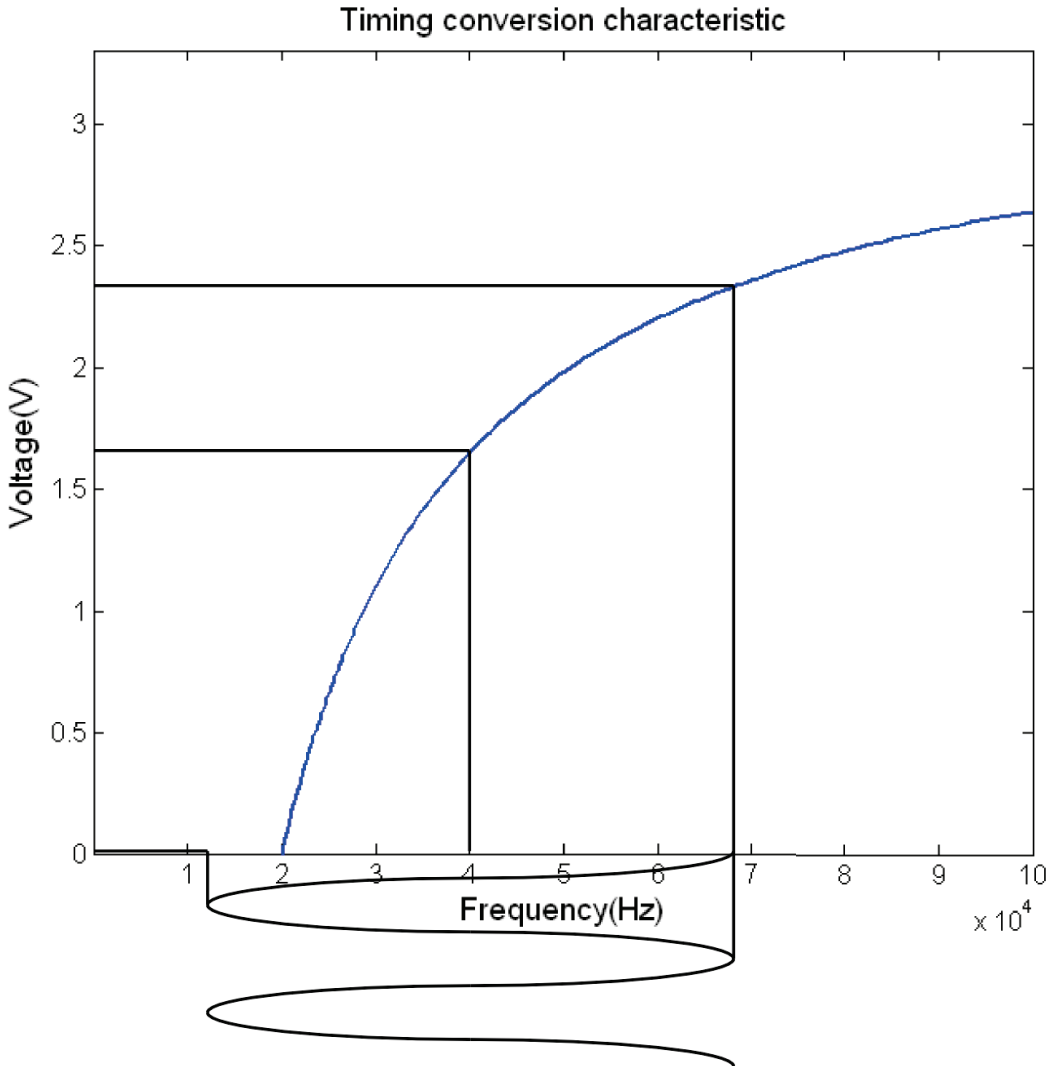

Figure 51: Processing response for large vibration amplitudes

3.3 Design of Circuit

The processing circuit described in this chapter was designed using Cadence as mentioned in the introduction. It was designed using a pipeline approach such that the input signal moved through several stages, changing from analog to digital and finally back to an analog signal. The use of digital components potential allows the zero-crossing method to be rearranged and scaled down and makes it a good option for pixel level processing.

Section 3.2 has shown how a zero-crossing demodulator can be used to process an FM signal producing a signal proportional to the original modulating signal. The optical processing provided by the interferometer produces a Doppler signal at a frequency which the camera can detect. The actual mixing is done at the photodiode due to it being sensitive to the incident light power. The photocurrent in the photodiode forms the input to the processing and resembles Figure 40. The output of the processing circuit is a varying, continuous voltage signal driving one of the pins of the chip.

The circuit was designed using a mixture of available components from previous chips and newly designed components. This allowed relatively quick prototyping of the processing by reusing proven components at the expense of poor fill factor. This was considered an acceptable trade-off as this is the first iteration of the proposed processing and forms a proof of principle.

Section 3.3.1 will cover the design of the circuit at a basic level showing how the input signal is processed. Section 3.3.2 will then provide a more detailed description of the processing and some considerations during the design process. The timing aspect of the circuit will then be covered in Section 3.3.3 as this forms the main part of the processing. Finally the aspects of this proposed circuit which affect its use as a full field LDV camera will be considered.

3.3.1 Basic Operation

The basic operation of the circuit is shown in Figure 52 and can be summarised as follows. The current in the photodiode is converted to a voltage using a current to voltage converter. The AC component of this voltage is amplified while leaving the DC almost unchanged using a Hysteretic Differentiator Amplifier (HDA) [95]. These two components are grouped into the Analog Front End (AFE). The output from the AFE is then clipped using a comparator. On each falling edge of the

comparator output two sequential, non-overlapping pulses are generated and used to control a timing circuit. One pulse is active low and the other active high. The timing circuit provides a voltage which is dependent on the time between the sets of control pulses (and hence the input signal's falling edges). This voltage is buffered and then drives the output pin of the circuit.

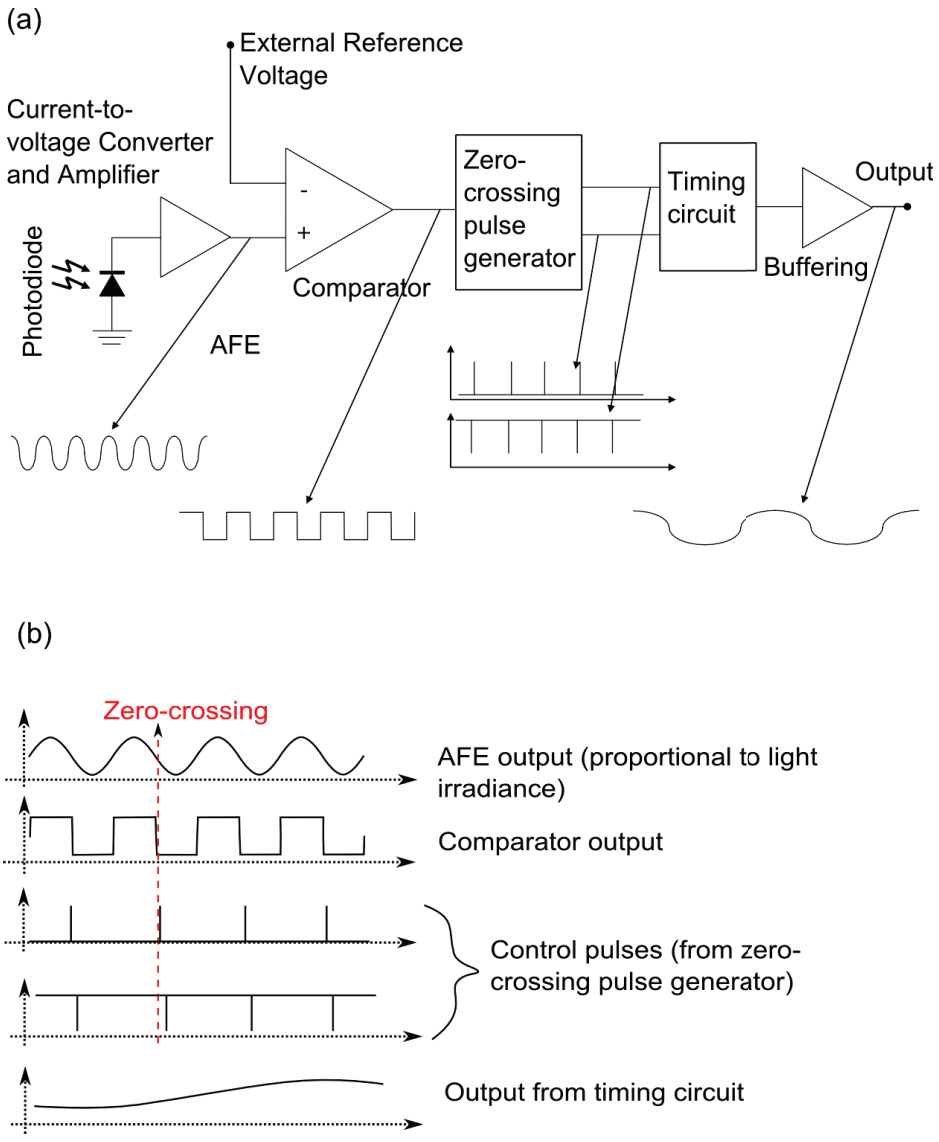

Figure 52: Overview of LDV processing circuit operation. (a) is a block diagram showing the main components of the circuit. (b) is a timing diagram showing how the different signals are related.

3.3.2 Detailed Operation

The design of custom CMOS chips allows efficient, dedicated solutions to a given problem. One drawback is that the final chip cannot be changed after it has been fabricated. For this reason as much controllability and observability is designed into the circuit. The controllability can take the form of digital signals which enable or disable parts of the processing. The observability includes signals from different points in the circuit being routed to pins to allow the internal operation of the circuit to be checked. This processing circuit used several forms of both controllability and observability, limited by the available spare space and pins on the BVIPS2 chip. Figure 53 shows a diagram of the different observability, controllability and noise reduction measures added to the processing circuit. These will be covered separately in the next sections.

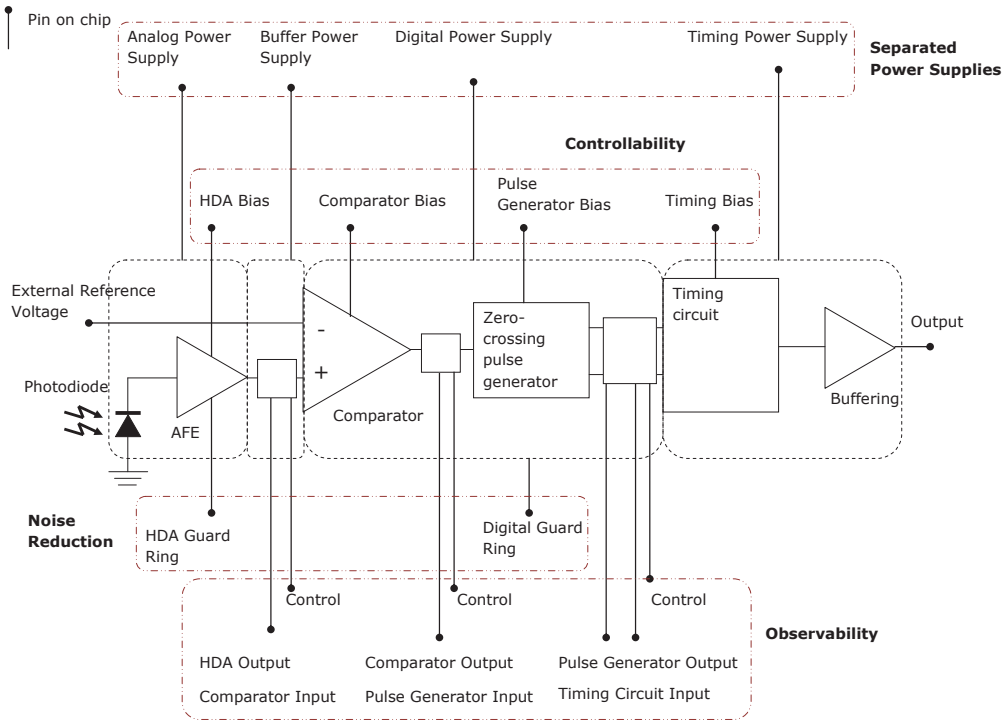

Figure 53: Detailed diagram of processing

Observability Observability is the addition of extra pins to allow access to signals from different parts of a circuit. This allows some recoverability during prototyping if a certain part of the chip does not work as expected. The observability within this circuit consisted of several signals from different stages in the processing being routed to pins on the chip. Each of these signals, with the exception of the final output, has a digital control line which also allowed signals to be inputted to that stage of the processing bypassing the signal from the previous processing stages. This gave the option of checking components in isolation. The observability points were:

AFE to comparator This was the signal between the AFE and the comparator. As an output it showed the amplified voltage signal from the photocurrent, which would be processed by the circuit. This allowed some verification of the light signal the camera was receiving. As an input it allowed the optical input signal to be bypassed and the comparator tested with an external voltage signal.

Comparator to Pulse Generator As an output this signal allowed the output of the comparator to be checked for switching. As an input it allowed the comparator to be bypassed and an external square wave applied to the pulse generator.

Pulse Generator to Timing As an output this set of pins showed the two control pulses for the timing circuit. As an input they allowed the timing circuit to be checked separately from the rest of the circuit.

Final Output This was an output only pin and gave the buffered output from the timing circuit.

Controllability Controllability covers all external signals, whether analog or digital, which influence how a circuit operates. These are useful for both prototyping and the final chip. The controllability was provided through the use of external bias resistors. Several components within the circuit depended on a current source for their operation. This current source, through the use of a current mirror setup, could be controlled by an external resistor. An example of this is the timing part of the circuit. Its operation depends on a capacitor discharging through a current source. If the current in that source was increased, using a smaller external resistor, then the capacitor would discharge faster. This would result in the output voltage from the timing part of the circuit being lower for a given input frequency. The available biases were:

Comparator Controlled the hysteresis level of the comparator. This controlled the minimum amplitude an input signal to the comparator would need to make the comparator switch. This formed a noise rejection/decision component in the processing to help decide whether a signal was genuine or noise.

Pulse Generator Controlled the width of the pulses that the generator produced. This controlled how long the timing capacitor had to charge as well as the error introduced due to the time taken to do the sample/hold and reset.

Timing Controlled the rate the timing capacitor discharged through a current source and thus controlled the period to voltage conversion.

Power Supplies and Noise Reduction Another key consideration with chip design is minimizing noise coupled between different parts of the circuit. This was especially important with this circuit as it had both analog and digital parts.

An example is the rapidly switching output from the comparator. This produced current spikes on the supply to this component. If the comparator shared its supply with the AFE then these spikes could be carried to the sensitive analog parts and introduce noise to the signal.

For this reason the power supplies for the different stages of processing were split to minimize noise coupling from one stage to another. There were four different power supplies for the processing:

Analog The pixel, I/V and HDA

Buffer The buffering between the HDA and the comparator

Digital The comparator and pulse generator

Timing The timing and it's buffer

Buffers were also placed between different components to help isolate them. Again, this was most important when the signal moved between analog and digital parts of the circuit.

The analog and buffer power supplies shared a single analog ground pin and the digital and timing power supplies shared a digital ground pin. It is worth mentioning that the grounds were connected within the chip through the substrate. The pins on the chip were grouped into digital and analog ones. Each group needed its own ground pin and having an additional ground contact would also share the load.

One other addition to the circuit was guard rings. These are rings of n-type silicon wells or p-type substrate contacts to help minimize the movement of carriers through the silicon substrate. It was anticipated that this would limit the coupling of noise from high speed, high transient current digital components to the neighboring analog components. The HDA had guard rings built in as part of

the design of the component as it was especially susceptible to free carriers within the silicon. The digital parts of the circuit also had some guard rings around the comparator and pulse generators as these were expected to have the highest transient currents.

3.3.3 Operation of The Timing Circuit

The timing component formed the main processing stage within the circuit and as such will be covered in more detail. Section 3.2 and Section 3.3.1 have covered the theory behind the zero-crossing demodulator and given a basic idea of how the processing operates. The purpose of the components preceding the timing component were to take a sinusoidal input optical intensity waveform and generate two digital signals on its falling edge. These two signals are the control signals in Figure 52 and are called the sample and reset signals. The sample signal was generated immediately after the falling edge of the comparator output and the reset signal was generated immediately after the sample signal as shown in Figure 54. The timing circuit would take these signals as an input and produce a voltage waveform based on the time between the sets of pulses. Figure 55 shows a schematic for the timing circuit. It basically consists of a timing capacitor connected to a set of switches to control the voltage across it. The timing circuit operates in one of two phases, or modes. The first is the timing phase. During this period the Timing Capacitor is discharged through a current source. The second phase is the sampling/ charging phase. During this phase the capacitor is sampled and then charged to the supply voltage.

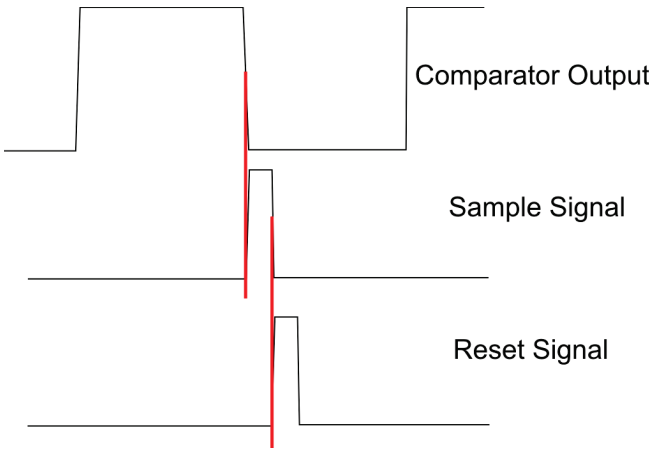

Figure 54: Timing diagram showing the relationship between the comparator falling zero-crossing and the sample and reset signals.

This sample signal is active high. It is used to drive transmission gates which require an active high signal and an active low signal so the sample signal is inverted within the circuit. The red and blue lines in Figure 55 show the inverted and non-inverted control signals respectively. To simplify the diagram the actual connections between the sample signal and the transmission gates have been left out, but the coloured lines on each transmission gate shown the relevant connections.

The circuit has three inputs and one output. These are the sample and control signals, and the bias voltage. The timing capacitor is the core component and can be connected in one of three ways, depending on the control signals: the power supply, the current source or sample/ hold buffer which is buffered and drives the output of the chip. The reset signal controls a transistor switch which can connect the capacitor to the supply voltage. This occurs when the reset signal goes low. The bias voltage controls the current through a transistor which provides the current source for the timing part of the circuit's operation. The bias controls how much current can flow through the current source, and therefore how quickly the timing capacitor discharge. The sample signal controls several

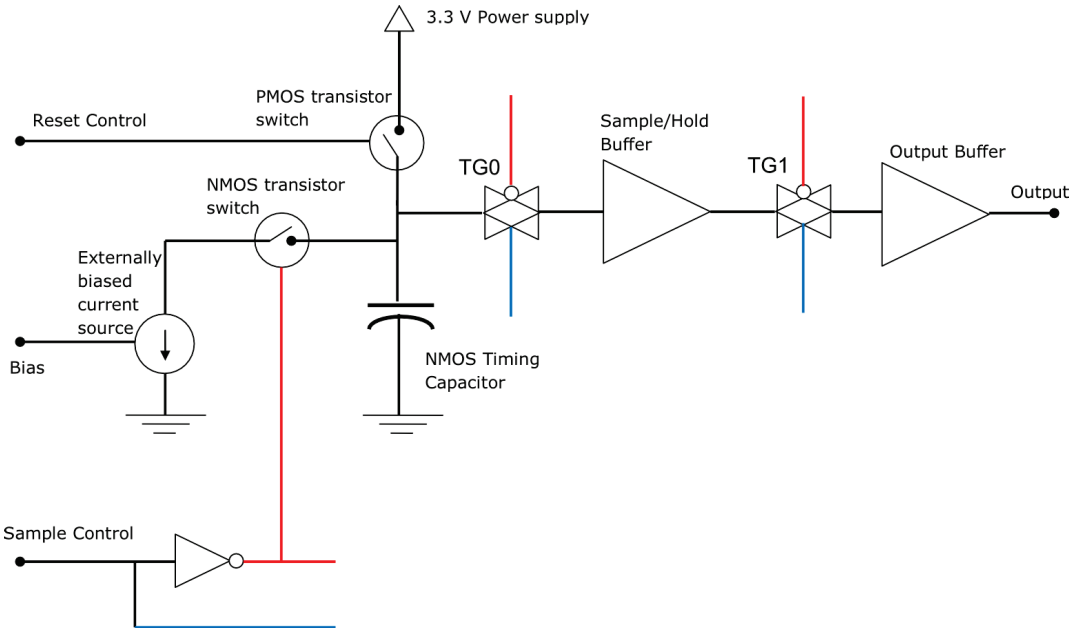

Figure 55: Schematic for timing circuit

transistor switches or transmission gates within the circuit to the sample the voltage across the capacitor at the end of each timing phase.

Figure 56 shows a timing diagram of how the reset and sample signals are related to each other and to the timing capacitor voltage. It shows the control signals in black and the voltage across the capacitor in red.

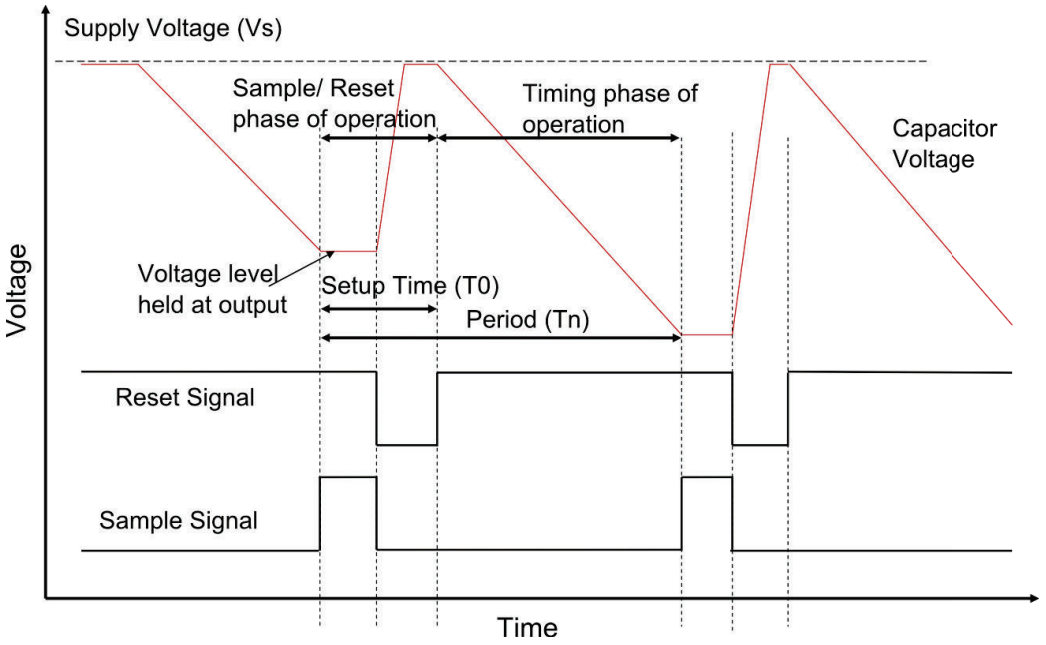

Figure 56: Plots of sample and reset signals for timing circuit as well as the capacitor voltage.

Timing Phase Considering the timing phase first. Both control signals are inactive and the Timing Capacitor is discharging through the current source. Figure 57 shows the position of the switches during this phase. During this phase the output voltage is approximately equal to the input voltage at the input to the output buffer. Due to the switch connected to this input being open this voltage is due to the charge stored on the input capacitance of the buffer. This voltage is labeled "Stored Voltage" on Figure 57. This works as the output buffer has a relatively large input capacitance and very high input impedance. The charge stored at the buffer's input discharges slowly and therefore can be used as temporary storage.

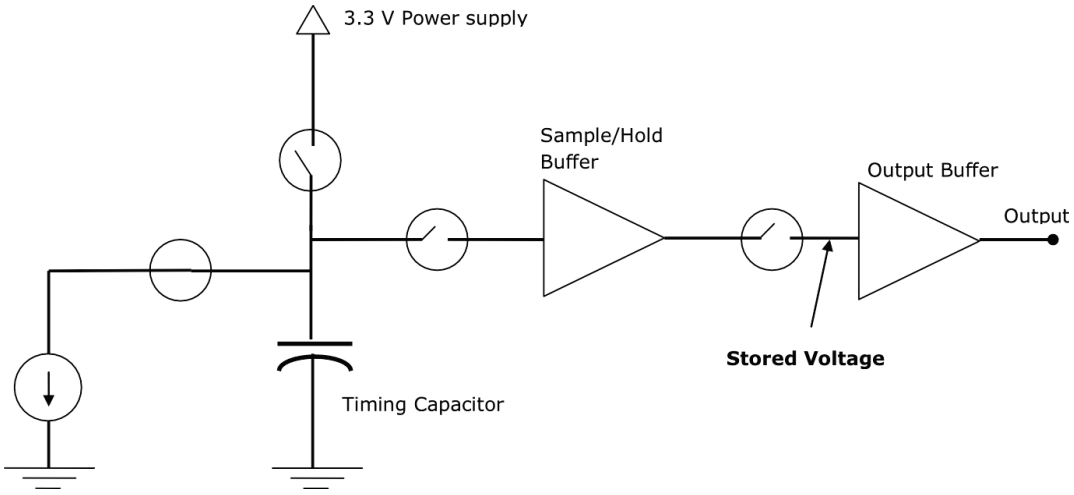

Figure 57: Position of switches during reset pulse and timing capacitor discharge

Sample/ Reset Phase As shown in Figure 56 there are two stages to the sample/ reset phase. The first is triggered by the sample signal going active and this stops the Timing Capacitor discharging and samples the voltage it had reached during the previous timing phase. Figure 58 shows the positions of the switches during this operation. The input capacitance (labeled "Stored Voltage") of the output buffer is charged to the voltage the timing capacitor had reached during the previous cycle. The sample/hold buffer is used to limit the loading of the Timing Capacitor during the sample phase of the circuit's operation. It has a much smaller input capacitance and keeps the effects of charge sharing between the Timing Capacitor and the buffer to a minimum. Also the sample/ hold buffer provides better isolation between the stored voltage and the capacitor voltage during the timing phase of the circuit operation than the switches would alone due to their finite off conductivity.

The second part of the sample/ reset phase is to reset the Timing Capacitor by charging it to the supply voltage. The sample signal goes inactive and the Timing capacitor is disconnected from the sample/ hold buffer and connected to the current source. The reset signal then goes active and charges the capacitor.

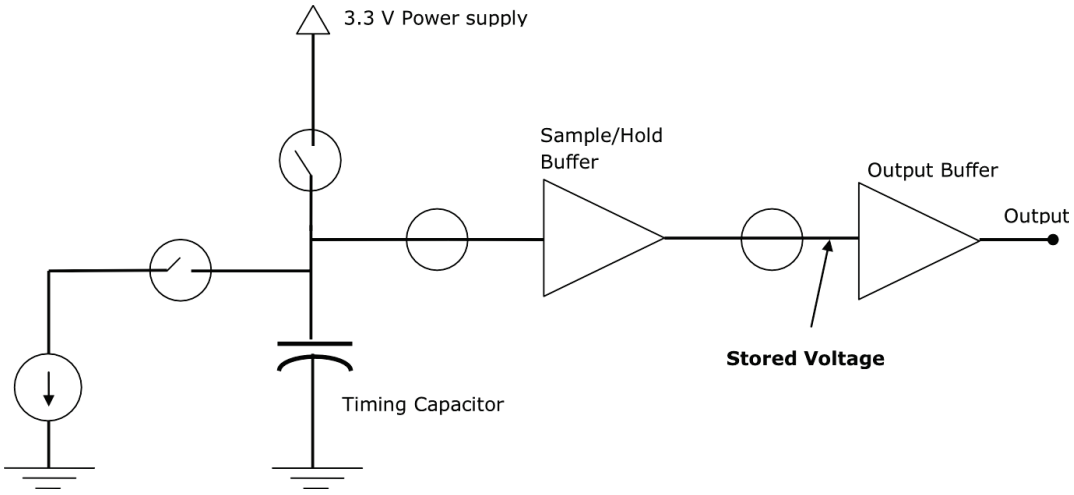

Figure 58: Position of switches during sample pulse

Once the reset signal has gone inactive too the circuit returns to the timing phase of its operation.

The control signals (sample then reset) are generated at the falling edges of the input photocurrent signal (from the irradiance). As this signal is a sinusoid with varying frequency, and therefore varying period, a number of control signals sets are produced and this results in the waveforms shown in Figure 56. At each falling edge of the comparator output the capacitor is first sampled and then recharged.

Figure 59 shows a simulation of the pixel in operation generated using Cadence. It demonstrates how the final voltage of the capacitor follows the frequency of the photodiode signal and how the output voltage follows the final voltage of the capacitor. The flat sections of the output signal show where the hold part of the timing cycle is active.

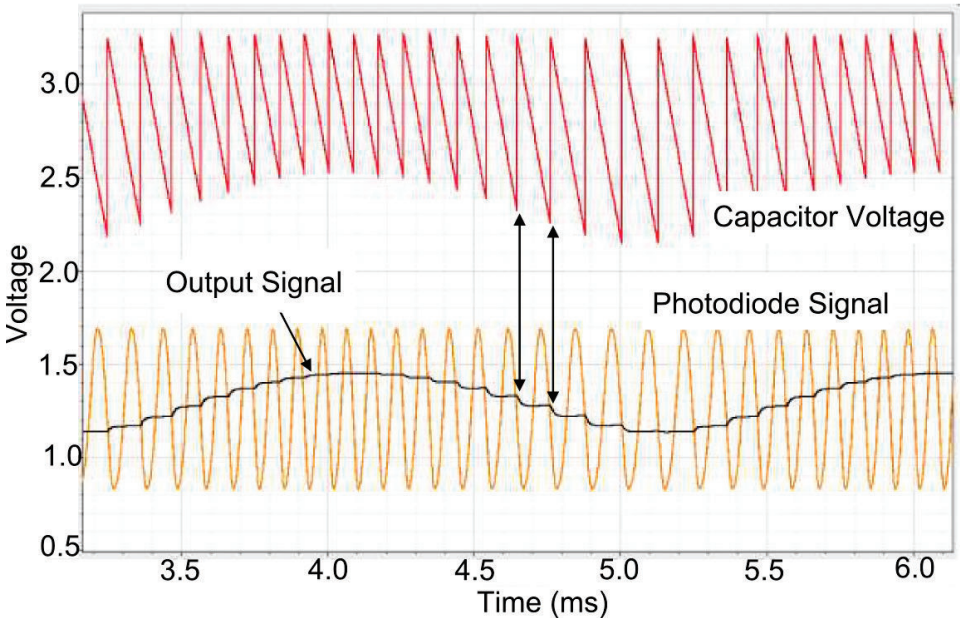

Figure 59: Simulation of zero-crossing processing.

3.3.4 Conclusion

This section has described the design of the processing circuit. It has covered the processing method used, showing how the input signal from the photodiode is demodulated to produce the output waveform. A more detailed diagram of the processing circuit has shown how controllability and observability have been added to the design and some of the measures used to try and minimise noise within the processing.

The timing part of the circuit was then covered in more detail, and its operation described. The description of the whole circuit has shown the order in which components appear in the processing and the architecture of the pipeline.

3.4 Simulation of Circuit

3.4.1 Introduction

The processing circuit covered by this report was designed to provide a zero-crossing demodulator for use as a single pixel of a full field LDV camera. The theory behind the processing has been given in Chapter 2 and the design of the processing has been covered in Section 3.3. This section will present simulations of the zero-crossing demodulation circuit and its individual components.

This processing circuit was intended as a prototype to provide a proof of principal of a pixel-level, zero-crossing demodulator or LDV. The simulations within this section therefore concentrate on the range of operation of the pixel with reference to use in a Mach-Zehnder interferometer. This includes the range of light levels and modulation depths which will produce an output signal. A number of other input conditions will also be considered and together these show the range of input signals the processing will work with. The range of input variables considered is:

- Light DC level
- Light signal modulation depth - This will be used instead of signal amplitude due to the nature of the signal from the interferometer. That signal consists of a DC term, related to the intensities of the signal and reference beams, and a sinusoidal term, the amplitude of which is also related to the intensities of the two beams. Using $I = \frac{c.\epsilon_0}{2}E_0^2$ which relates the irradiance (or light power per unit area), Equation 30 can be written in the form $I(t) = I_{sig} + I_{ref} + E_{sig}E_{ref}cos(\phi(t)) = I_{sig} + I_{ref} + 2\sqrt{I_{sig}I_{ref}}.cos(\phi(t))$, where I_{sig} is the irradiance due to the signal beam alone. If the irradiances of the two beams are identical then the input signal will have 100% modulation depth. If the intensities are mismatched then the modulation depth will decrease.

- Vibration signal maximum frequency and amplitude. As shown in Chapter 2 the bandwidth of the photocurrent signal, the input to the processing pipeline is dependent on both the vibration frequency and the vibration amplitude. As such, for a given maximum bandwidth, the vibration frequency could be increased if the amplitude was decreased.

- Maximum noise on the input signal. This is related to the range of signals the comparator will switch for, and therefore propagate through the circuit, while rejecting noise on that signal. This is related to the comparator hysteresis range.

Another aspect of the pixel which will be simulated is the expected process variation in the finished chip. This is important as variation in certain component parameters during fabrication is unavoidable. An example of a parameter which might vary is the width of a transistor. Random variation between fabricated chips of the width of the transistors will cause the width to length ratio to change and therefore the operation of the transistor. This random variation can be minimised but never removed. One common method for minimising process variation is the common centroid technique [121]. Anther method is using larger transistors so that the fractional variation is smaller. This pixel is intended to form part of a camera chip and given that variation across the chip will change the parameters of each pixel, it is important to check how this will affect the range of inputs the camera can accept with most of the pixels still operating correctly.

A Monte Carlo simulation provided by Cadence was used to simulate process variation and produces a probability distribution of expected values for a given simulation result by repeating the simulation a large number of times with the expected parameter variations.

3.4.2 Procedure

The simulation of the individual parts of the processing circuit was carried out using Cadence, in roughly the order of the processing pipeline. To simplify the simulation some of the components were grouped into logical blocks. Each block was tested in isolation and a section is given over to each of these. These sections start with an introduction to the components giving their schematics. This is followed by any specific information for the commands required to run the simulation. The main section then covers the actual simulation and provides plots and results.

The order of simulations is as follows:

1. Comparator - This would act as the control of which signals would be processed and which would not. Effectively rejecting noise signals.

2. AFE - The AFE contained the I/V and HDA. These could be grouped as only the biasing of the HDA was accessible externally, and these components could not be separated during testing of the finished chip. The comparator was also added for some of the tests because the behavior of the comparator was partially affected by the input signal.

3. Pulse Generator - This contained the two multi-vibrators used to generate the control pulses for the timing circuit.

4. Timing - The main part of the processing. This included the timing capacitor, associated control switches and output buffers.

5. Whole circuit - The final simulation was a Monte-Carlo simulation to check how much variation in the final pixel could be expected. This is required to check if the pixel is suitable for incorporation into a camera without losing a significant number of pixels.

3.5 Comparator

3.5.1 Introduction

The comparator provided a key component in the design as it provided an interface between the exclusively analog AFE, and the pulse generators which are basically digital. The purpose of the comparator was to clip the input and as such the output was expected to be a frequency modulated square wave. The comparator acted as a decision component. Only signals larger than the hysteresis of the comparator would cause an output. Those smaller, including noise, would be rejected by the comparator. Simulations were carried out to determine the range of operation of the comparator and how the hysteresis level could be varied by the external biasing.

The circuit for the comparator was taken from a modulated light camera designed by Light [121], and used without change.

3.5.2 Simulation

Four sets of simulations were carried out on the comparator circuit. Transient simulation was used to show the operation of the comparator circuit over time. This was followed by simulation of the hysteresis function of the circuit and finally the frequency response. Due to the initial state of the comparator simulation balancing the currents in the two sides of the differential pair setup, the comparator would act as an amplifier until it had switched once. To get around this a piecewise-linear voltage source(vpwl) was used to force the comparator into an output high state initially. It was added to the non-inverting input in series with the sine wave source and set to give $3.3\,V$ for $50\,\mu s$ and then drop to 0 V after $60\,\mu s$. The simulation was set to discard the first $100\,\mu s$ of output.

Transient Operation The purpose of the comparator was to produce a clipped output based on an input sinusoid. The inverting input would be externally set to a voltage equal to the mean of the input signal. The incoming signal would be passed to the non-inverting input. Figure 60 shows a transient plot of an input waveform larger than the hysteresis level for the comparator. This plot shows how the comparator takes a small input sinusoid and produces a square wave output. The rising edge of the square wave matches the time the input voltage (at the non-inverting input) crosses the reference voltage (at the inverting input) which is set to $1.65\,V$. The falling edge of the output then matches the point the input voltage drops below the reference voltage. The falling edge of the comparator took $3\,ns$ to complete the transition. This value would be useful for the pulse generator simulation.

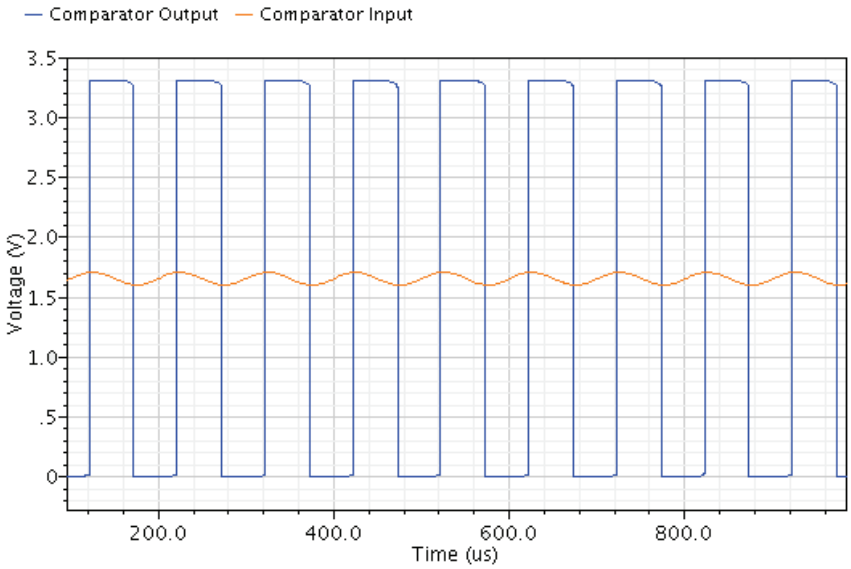

Figure 60: Comparator output for and input sinusoid

Hysteresis Hysteresis is the property of some switching circuits where the input voltage level at which the circuit output switches, changes depending on the current output state. During preliminary testing it was found that the bias of

the comparator affected the hysteresis level. This would be useful in setting the noise rejection of the comparator. To test the hysteresis a number of transient simulations were done with different input AC amplitudes. For those below the hysteresis level the comparator would not switch repeatedly. Those signals of larger amplitude would cause it to switch.

Figure 61 shows the hysteresis level in operation. It was generated by running several transient simulations using a sinusoidal input voltage to the non-inverting input. The amplitude of the input signal was swept from $40m\,V$ to $55m\,V$ while the bias resistance was kept fixed. The DC level of the signal was kept at half the supply($1.65\,V$) and the inverting input was also set to $1.65\,V$. The amplitude of the output signal is plotted against the input signal amplitude. There is a sudden change from not switching(less than $30m\,V$) to switching ($1.65\,V$) when the input signal amplitude increases above the hysteresis level.

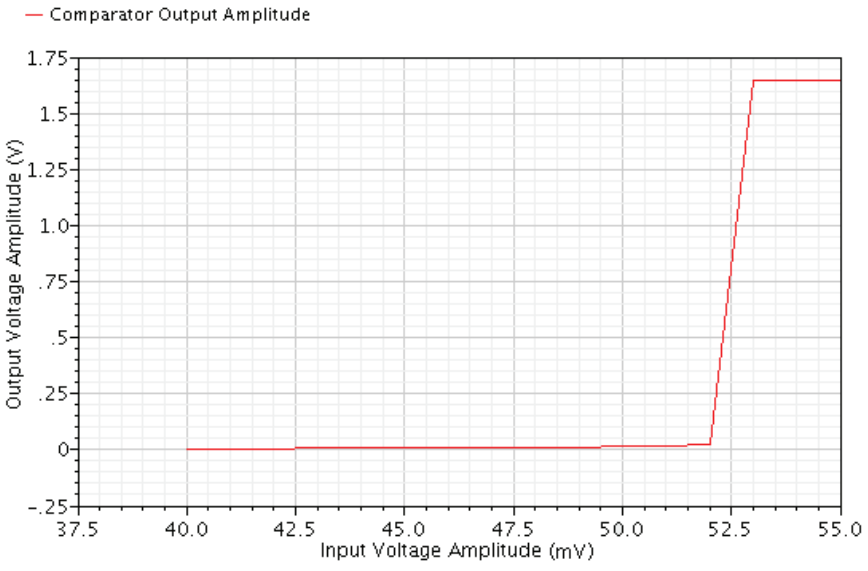

Figure 61: Output signal amplitude vs. sweep of comparator input signal amplitude from 40m V to 55m V. Bias resistor set to $16k\,\Omega$. Switching level and input DC level set to 1.65 V

The next simulation was to find how this hysteresis level would be affected by changes in the bias resistance. These results are shown in Figure 62. This shows

that the largest hysteresis would occur at $\sim 16.5k\,\Omega$ with a value of $53m\,V$. The DC level of the signal was kept at half the supply$(1.65\,V)$.

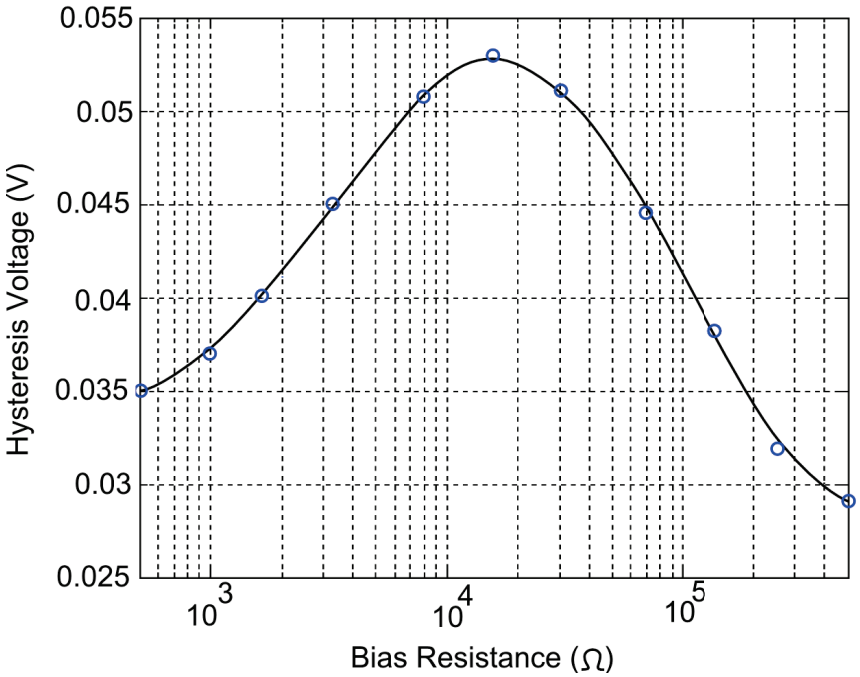

Figure 62: Change in hysteresis voltage with bias resistance

The hysteresis level could be reduced by setting the bias level higher or lower than this value. At $500k\,\Omega$ the hysteresis level was $29m\,V$.

Given the process variation which occurs during fabrication each pixel's comparator would be slightly different and the hysteresis level would vary. These effects were simulated by running a 100 iteration Monte-Carlo simulation for each of a range of input AC amplitudes, and checking what fraction of the iterations worked for that amplitude. Figure 63 shows the results of these simulations. As can be seen a few comparators would switch for input amplitudes about $15m\,V$ below the expected value while a few would require the input amplitude to be about $15m\,V$ above the expected value before they switched.

Simulation also showed that a further variable, the DC level, affected the hysteresis level. Figure 64 show simulation results for the hysteresis level with a

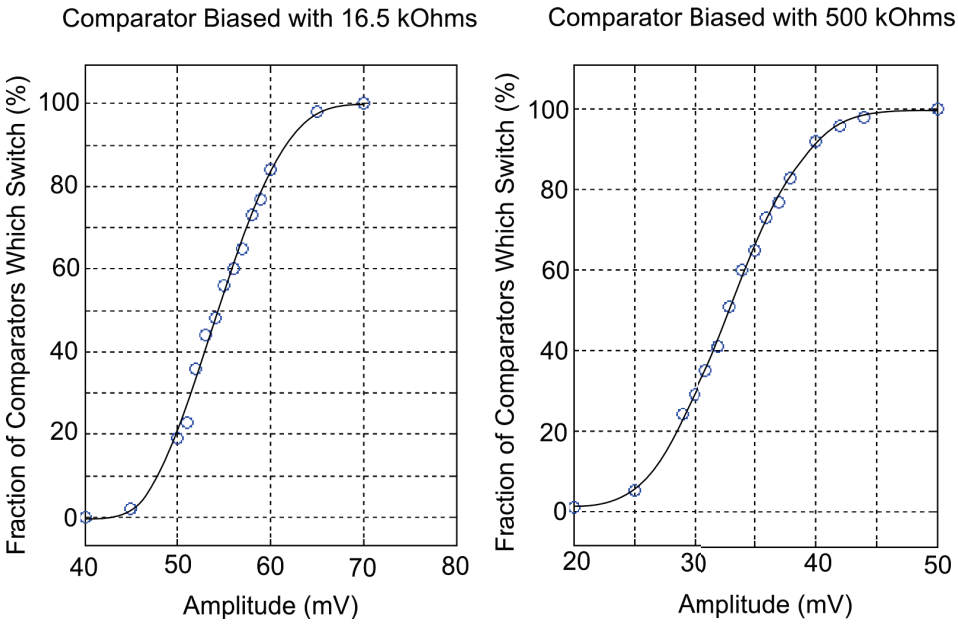

Figure 63: Monte-Carlo simulation of the fraction of finished chips that would switch for a given amplitude. Two resistor biases were used: $16.5k\,\Omega$ and $500k\,\Omega$.

sweep of the input DC voltage (of both the input signal and the reference voltage) while holding the bias resistance constant at $16.5k\,\Omega$ and the AC amplitude at 53m V.

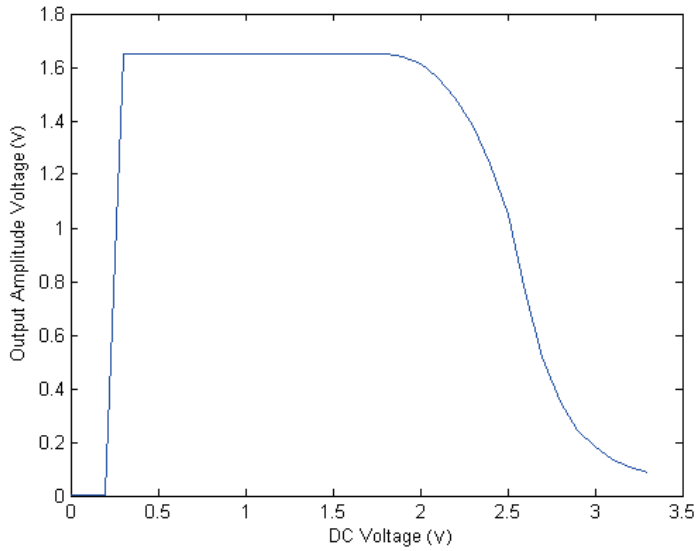

Figure 64: Hysteresis level variation with DC input level.

This plot shows that there is a lower cutoff for a given hysteresis level and at

upper DC levels the comparator starts to act as an amplifier forming an upper cutoff. Figure 65 shows the transient results for some of the points above 1.8 V. As the DC level increases the output amplitude decreases. The comparator no longer clipping the input.

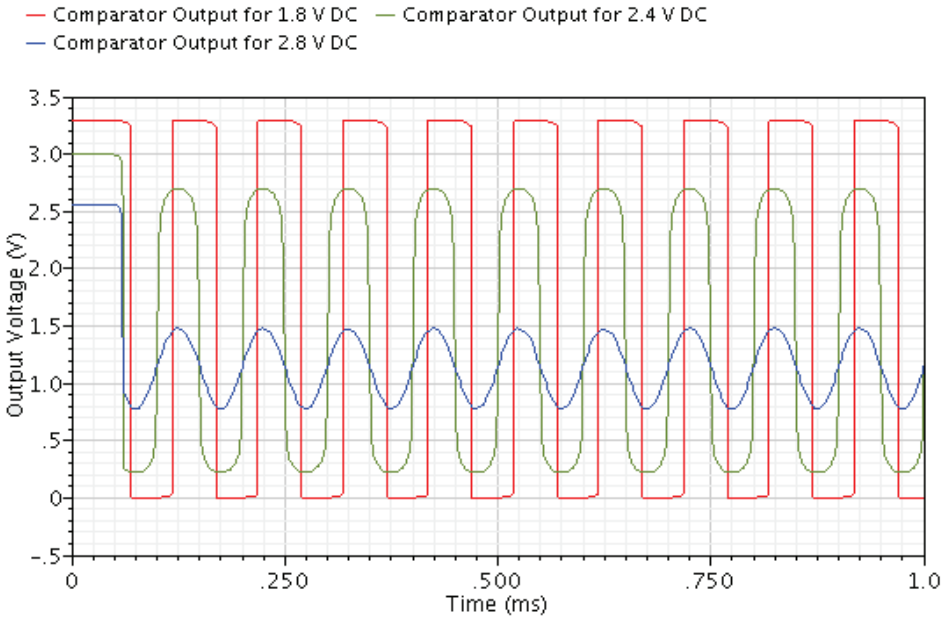

Figure 65: Comparator output vs input DC level transient waveforms

Frequency Response Simulation of the comparator with different frequency inputs show that for $53\,mV$ input signal the output switches correctly up to $1.78M\,Hz$. This upper cutoff frequency can, however, be increased by increasing the amplitude of the input signal.

3.5.3 Conclusion

The comparator was affected by a number of factors including DC level, AC amplitude and bias resistance. It could be controlled using the bias resistance to give a range of hysteresis levels from $29m\,V$ to $53m\,V$. Given the variables that affect the comparator output it was decided to group the comparator with the AFE to simplify the simulation. This would allow the output of the comparator

(whether it switched) to be related to the external variables, such as DC light level, rather than internal ones like the AFE output DC level. If a specific set of input criteria were needed the circuit could still be simulated in more detail to find which component was limiting the processing for that value.

These preliminary tests on the comparator still provide a groundwork for the operation of the whole circuit.

3.6 Analog Front End and Comparator

3.6.1 Introduction

The Analog Front End (AFE) and comparator controlled how the input light signal would be prepared for the digital part of the processing. The AFE takes the incident irradiance and produces a proportional output voltage. The comparator then clips that voltage into a square wave. The square wave provides a binary input to the rest of the circuit (high or low). The processing can therefore be split into the AFE and comparator, which interface with the external world, and the pulse generator and timing, which provide the processing.

Both the AFE and comparator were designed for previous chips and used without change. The AFE was taken from a blood flow camera designed by Diwei He [96] and contained the photodetector, current-to-voltage converter and amplifier as a single unit. The photodetector used by the AFE is a pn photodiode, which generates a current proportional to the incident irradiance. The photodiode was fabricated using a $0.35\,\mu m$ process and is a "well-substrate" type with a responsivity of $0.3\,A/W$ at $667\,nm$. This type of detector was chosen due to its fast time response, medium current gain and dark current, large dynamic range, and simple layout compared with other CMOS detectors. The circuit used to convert the current (called the photocurrent) to a voltage is called a transimpedance amplifier and the one chosen for the AFE has a logarithmic response. It was chosen due to

its small layout size and different responses to AC and DC signals. The reasons this detector was chosen for Diwei He's work are valid for this work.

- The small layout size is vital for a practical camera chip.
- The AC and DC differentiation results in amplification of AC which keeping the DC the same. The DC component of the irradiance does not contain any vibration information so preferential amplification of AC is very useful.
- The large dynamic range means that the detector will work well with very large signals and very small signals without saturating. This allows the detector to work with a wide range of laser powers and object reflectivities which greatly simplifies practical application of the pixel.
- The medium current gain affects how powerful a laser is needed. If needed, more powerful lasers are available, and the dynamic range of the sensor means that the more powerful laser can be used even for highly reflective objects where it is not needed.
- Fast time response is a very important property of the detector as the on-pixel processing is of the most benefit for high bandwidth vibrations (either large amplitude or large frequency) where other methods start to struggle with the data rate.

Figure 66 shows a schematic of the AFE. The I/V on the left contains the photodiode (n-well, $25\,\mu m$ x $25\,\mu m$, $0.3\,A/w$ responsivity) and converts the current through it to a voltage. This voltage is then buffered and drives to the HDA which is effectively a differential input amplifier with the inverting input connected to a low pass filtered version of the output. Effectively biasing out the mean of the input signal. The result of this is that the HDA provides reasonable AC gain (~ 40) but very little DC gain.

The specification for the AFE and comparator would then control the external light signals the processing could accept. The variables for the input signal were:

- Frequency
- DC level
- Modulation depth - used to calculate the AC amplitude.

One other variable was the biasing resistor for the comparator. This controlled the hysteresis level of the comparator and as such would select how large a signal had to be to make the comparator switch. For these tests it was decided to set the bias to give the largest hysteresis. Section 3.5.2 has shown that there is a single value for maximum hysteresis and this gives a worst case scenario.

There were two other biases shown in Figure 66: Bias SF and Bias I/V. These were internally biased and as such are fixed.

This section will start by describing the AFE and providing some results to demonstrate its limitations. Then the AFE and Comparator will be simulated together to check what range of input values that cause the comparator to give an output. In other words the operational range of the AFE and Comparator. For this section the circuit (AFE and comparator) will be considered to operate for the given input conditions if the comparator provides a switching, clipped output due to the input.

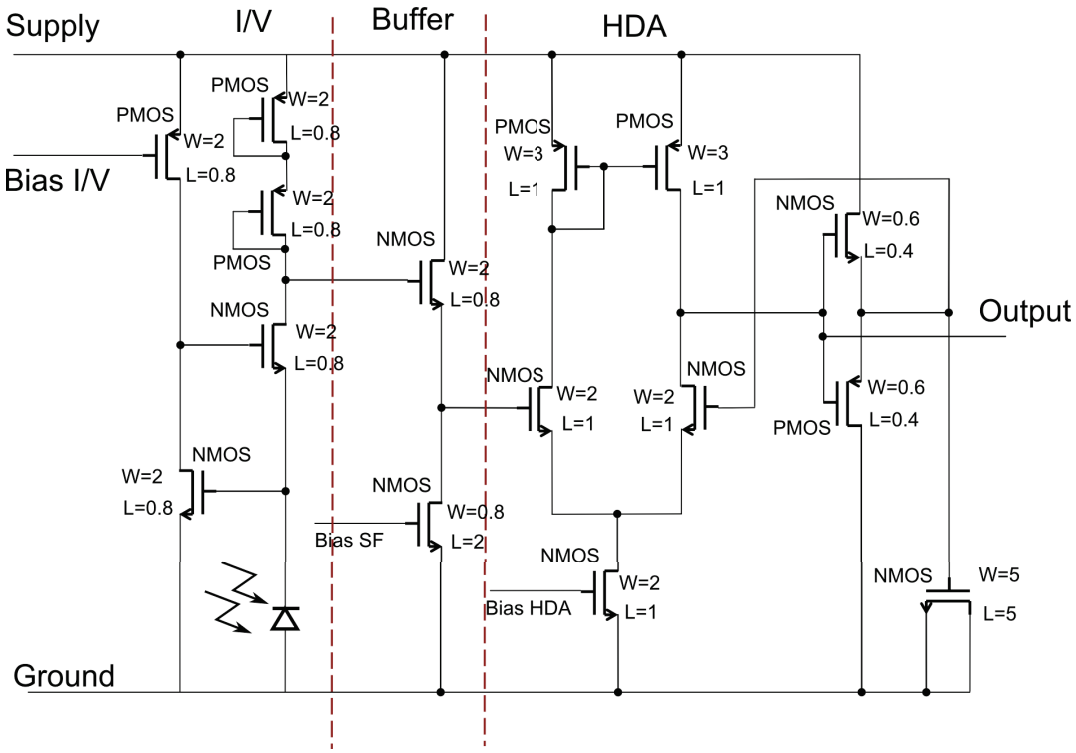

Figure 66: Circuit schematic for AFE.

3.6.2 Simulation

Several simulations were carried out on the AFE to check how it would perform under the expected operating conditions. For accurate simulation of the HDA in the AFE the default setting of the simulation parameter gmin needed to be changed to 1×10^{-24}. This is the conductance added across nonlinear devices to prevent nodes from floating if a device is completely off.

The operation of the AFE-Comparator components and the AFE itself was tested over the frequency range from $10\,Hz$ to $10\,MHz$ and the DC photocurrent range $100\,fA$ to $10\,\mu A$. This introduced some complexity into the simulation procedure. The wide range of frequencies resulted in a wide range of minimum simulation times. To improve efficiency the simulations were split into frequency groups with shorter simulations for higher frequencies to keep them to about 100 cycles. For photocurrents below $10\,pA$, time was needed at the start of the

simulation to allow the output's DC voltage level to settle.

Dark Current As already mentioned, the detector used has a logarithmic response. This is due to the photocurrent being relatively small. The means that the transistors this current flows through are in a region of operation called "subthreshold". This this operational region the current to voltage characteristic they exhibit is logarithmic. The output voltage would start at a certain voltage with no illumination and decrease logarithmically with increasing light level. The voltage level for no illumination is set by the current leaking through the photodiode as well as the buffer and HDA circuitry. The process parameters calculated the leakage current (also called the dark current) as $2.51\,fA$ and a Cadence simulation gave a value of $2.509\,fA$ with an output voltage of $2.064\,V$.

DC Sweep The next step was to simulate the AFE with a sweep of DC illuminations. This is shown in Figure 67 and shows the logarithmic response of the photodiode configuration. The output from the AFE also shows how the HDA introduces a slight DC shift into the signal. Above about $0.8\,V$ the output voltage is higher than the input voltage, and the logarithmic response of the pixel is shown. Below this voltage the response of the pixel starts to break down. It was not expected that this would be a problem as this would require a photocurrent greater than $100\,nA$ which would require direct exposure to the laser intended for the interferometer. The light level at the sensor could also easily be reduced.

The upper cutoff for photocurrents in the femto-amp region was an inherent limitation due to the dark current of the photodiode.

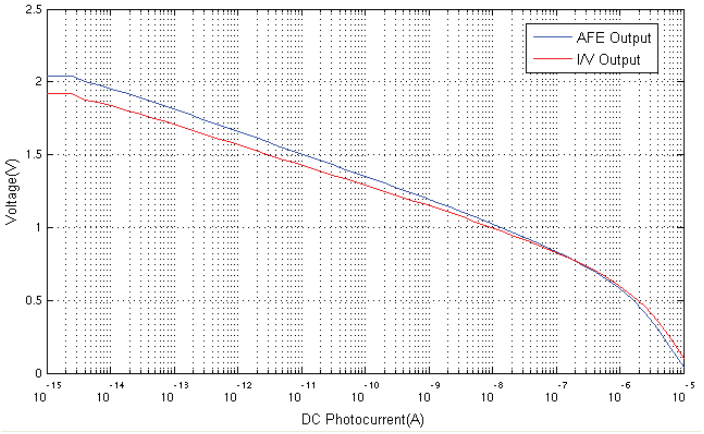

Figure 67: AFE output with sweep if DC illumination

Frequency Response Preliminary simulation showed that the frequency response of the AFE would likely be the limiting factor in the response of the entire processing circuit. These simulations also showed that the frequency response was dependent on the DC photocurrent. Figure 68 shows the frequency response of the AFE for $100pA$ DC current. This is the result of an AC simulation and shows the expected curve. However, this is generated by a simulation using a linearized model of the transistors. A transient simulation is more accurate and this is demonstrated in Figure 69. This shows the frequency response of the AFE falling off at lower frequencies and a slight change in the $3dB$ point. This is due to distortion of the output and results from the relatively large modulation depth as shown in Figure 70.

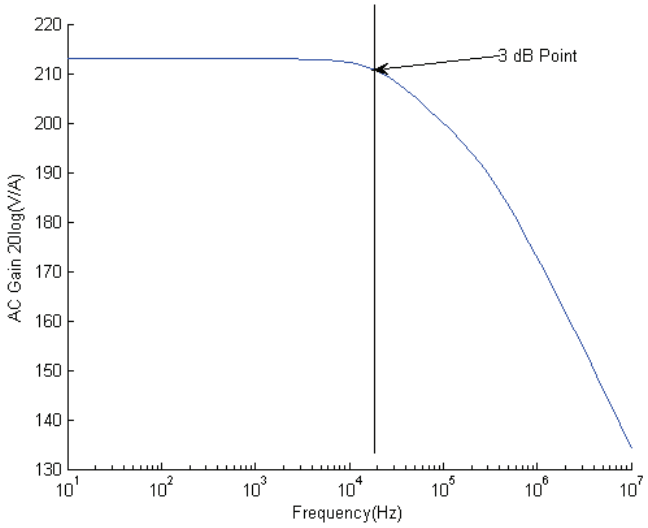

Figure 68: AFE frequency response $100\,pA$ (AC simulation)

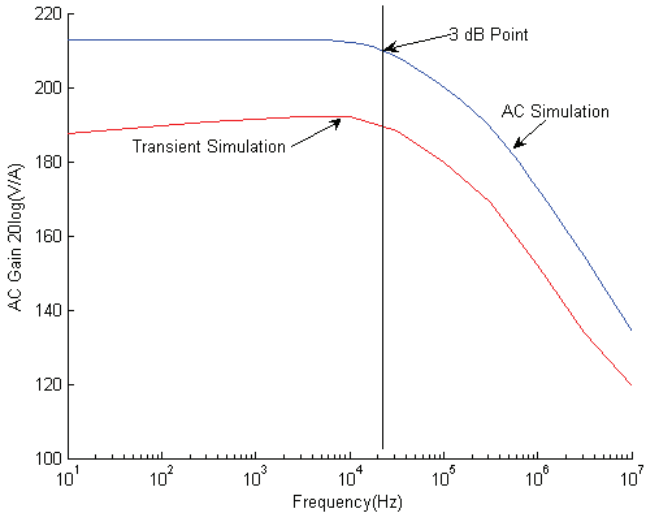

Figure 69: AFE frequency response $100\,pA$ (transient simulation)

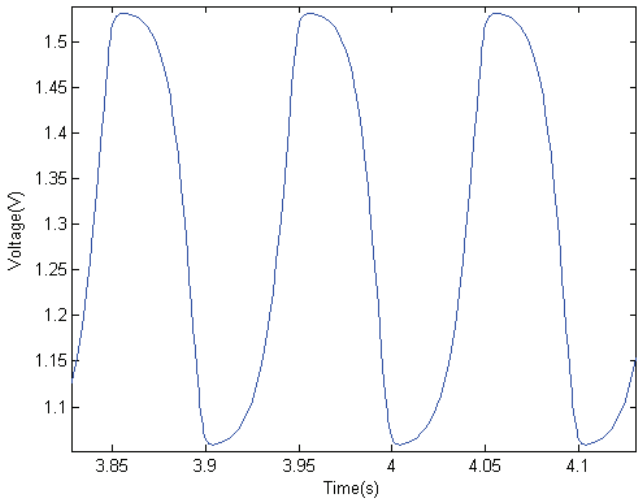

Figure 70: AFE output showing distortion. $100\,pA$ DC current, 10 Hz signal

10% Modulation Depth Transient simulations will be used for the remainder of this section. The following were done with a modulation depth of 10%. This modulation depth was chosen as initial tests had shown that this was the approximate modulation depth which could be expected from the experimental setup. Further simulation of the AC response with different DC photocurrents show a change in the 3 dB point and the gain. This is shown in Figure 71. The two plots are for $10p\,A$ and $100p\,A$. The $3\,dB$ point for the $10p\,A$ plot is at $\sim 208\,dB\Omega$ and $\sim 3k\,Hz$ while for $100\,pA$ the value are $180\,dB\Omega$ and $20k\,Hz$. These simulation results show that the frequency response of the AFE is in part controlled by the DC photocurrent of the input. Since a single frequency response plot does not describe the operation of the AFE under all conditions, a set of simulations were done, varying the DC photocurrent, and used to create a single 2-dimensional plot. This is shown in Figure 72. This plot was created by taking plots like those found in Figure 71 and recording only the frequency of the $3\,dB\Omega$ and $6\,dB\Omega$ points (both upper and lower, where appropriate). The 6 dB lines are shown to give an indication of roll off for the different photocurrents. On Figure 72 the x-axis shows the DC photocurrent and the y-axis shows the frequency. The red

line then shows how the 3dB pass band moves as the DC photocurrent is varied. This shows that at low levels of DC photocurrent(the left of the plot) the AFE acts as a low pass filter with a cutoff at less than $100\,Hz$(the lower left corner). As the DC photocurrent is increased the upper cutoff frequency increases and at about $10p\,A$ the frequency response of the AFE starts to show a bandpass characteristic with both an upper cutoff and a lower cutoff. At about $10n\,A$ to $100n\,A$ the frequency response reaches its greatest bandwidth at $\sim 1M\,Hz$. Above $1u\,A$ the bandwidth falls off rapidly and the AFE effectively stops working. This could be due to the output DC voltage being too low for the buffers in the AFE.

The relationship between bandwidth and DC photocurrent is covered in more detail in [95] and is due to the impedance of the diode connected transistor which provides the resistive element of the current to voltage part of the circuit.

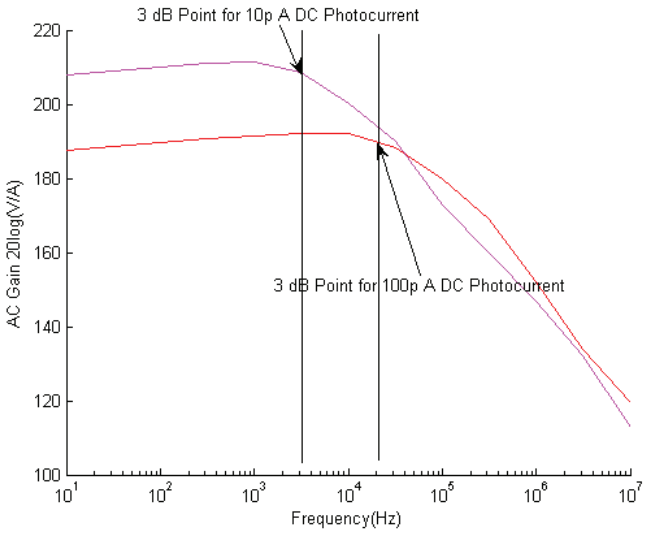

Figure 71: AFE frequency response 10p A(transient simulation)

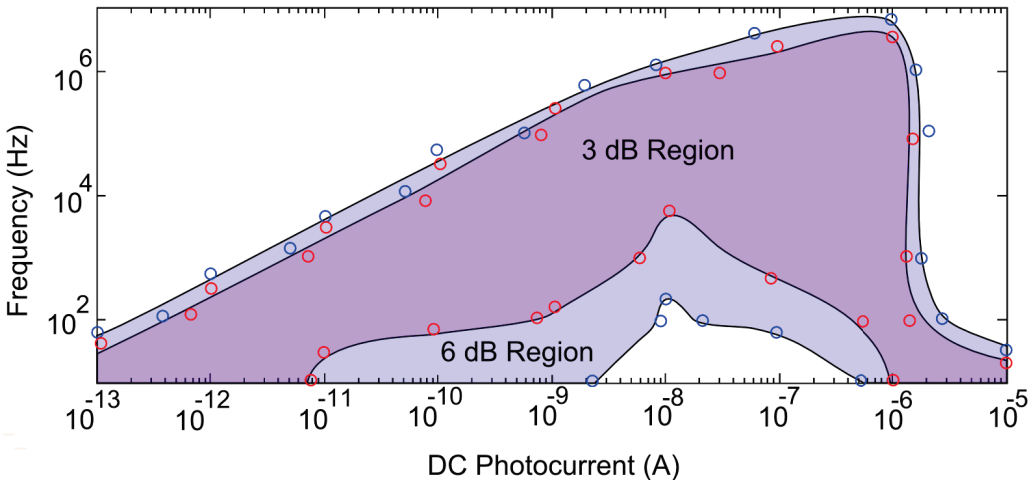

Figure 72: AFE frequency response for the range $10\,Hz$ to $10M\,Hz$ and $100f\,A$ to $10\mu\,A$ with $10\,\%$ modulation depth.

Given the complexity of the AFE and comparator (as shown in Section 3.5.2) it was more practical to simulate the AFE in conjunction with the comparator. This allowed a relatively easy measure of operation: if the comparator provided a clipped output for a given set of input criteria, then it could be considered that the processing circuit would operate under those criteria. This was achieved in a similar way to the plot in Figure 72. The same grid was used except the switching output of the comparator was given instead of $3\,dB\Omega$ or $6\,dB\Omega$ lines. The results of this are shown in Figure 73. All area inside the green line gave an output according to the simulation.

1% Modulation Depth The preceding procedure was then repeated for modulation depth of 1% and gives Figure 74 and Figure 75. These figures now show a simple low-pass filter characteristic, rather than bandpass, where the cut-off frequency increases with increased DC photocurrent.

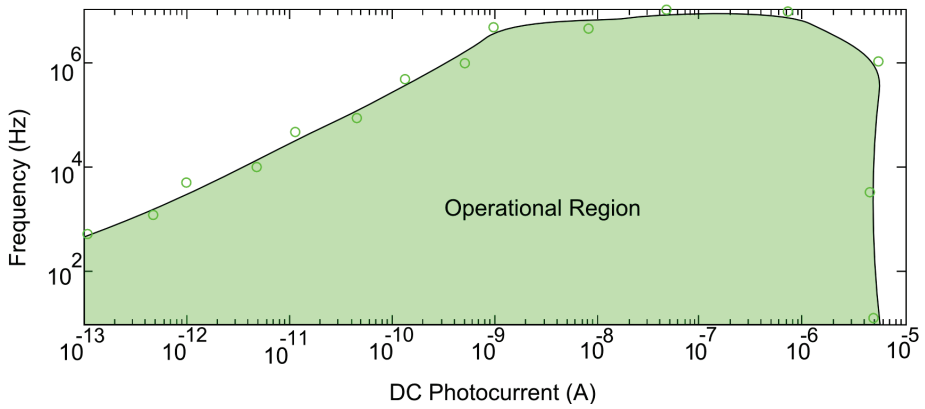

Figure 73: AFE and comparator operational range(below the green line) at 10 % modulation depth

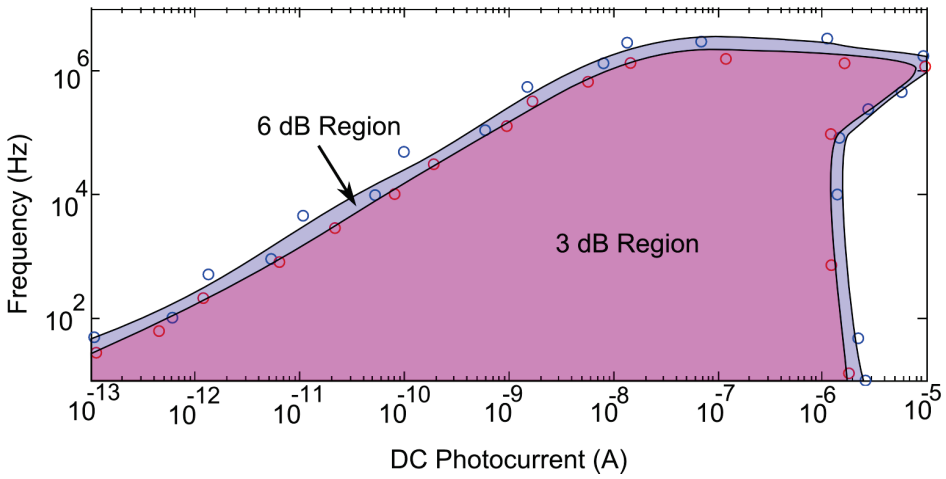

Figure 74: AFE frequency response for the range 10 Hz to 10M Hz and 100 fA to 10 μA with 1 % modulation depth.

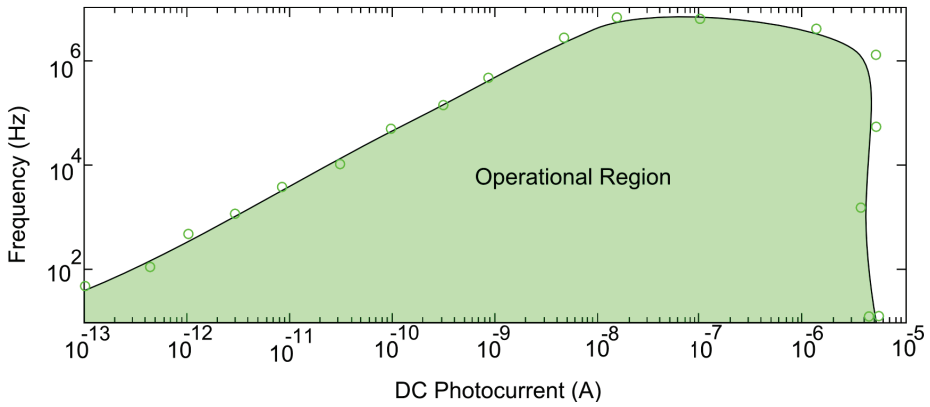

Figure 75: AFE and comparator operational range at 1 % modulation depth

Noise Any noise added to the input signal by the AFE would affect the switching of the comparator. The hysteresis of the comparator would prevent multiple switching of the output for a given input edge, however it would not prevent phase noise on the timing of the switching. This variation in the exact time of the square wave rising or falling edges, as compared with the incoming signals rising and falling edges, would either artificially increase or decrease the apparent frequency of the signal as seen by the timing component.

The left plot of figure 76 shows a simulation of the Power Spectral Density(PSD) of the noise at the output of the AFE with $10\,pA$ DC photocurrent. As with the frequency response of the AFE, the noise PSD varies with DC photocurrent. The right plot of Figure 76 shows the rms noise for photocurrents of $1\,pA$ to $200\,pA$. This shows that the noise decreases with increased DC illumination. The major noise source in the circuit is the two diode connected PMOS transistors in the I/V converter.

Monte Carlo The AFE was simulated using a 500 iteration Monte-Carlo analysis for nine input scenarios: $I_{DC} = \{1p\,A,\ 100p\,A,\ 10n\,A\}\,; Frequency = \{1k\,Hz,\ 10k\,Hz,\ 100k\,H$ to see the effect of IC process variation on circuit performance. The results are

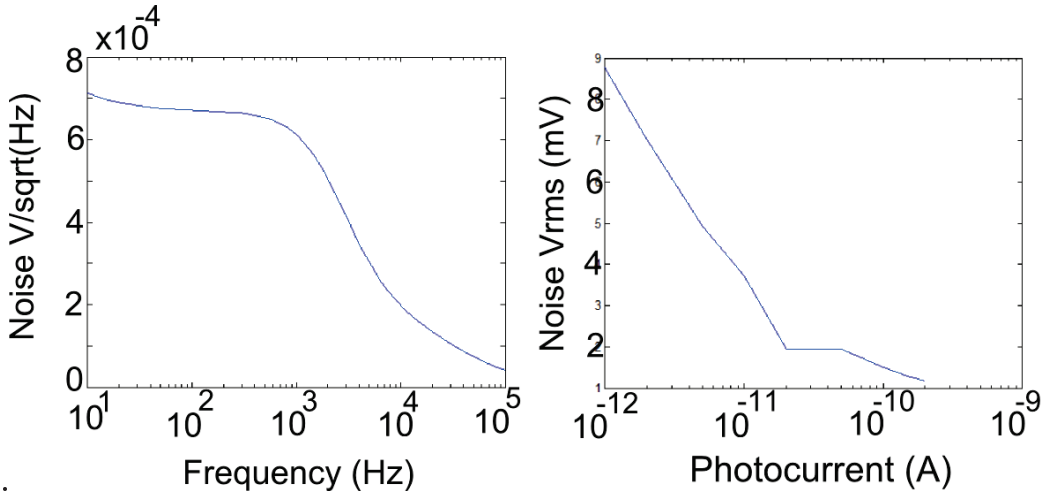

Figure 76: Noise PSD for the AFE on the left and the rms output noise of the AFE vs. DC photocurrent.

given in Tables 3(DC output voltage level from the AFE), 4(transimpedance gain of the AFE) and 5(output rms noise).

Table 3: Variation in output DC voltage level(units V)

	$1p\,A$	$100p\,A$	$10n\,A$
$1k\,Hz$ - mean	1.676	1.2509	1.01
σ	146.8m	184m	150.3m
$10k\,Hz$ - mean	1.676	1.334	998.4m
σ	146.8m	141m	147.4m
$100k\,Hz$ - mean	1.676	1.367	0.990
σ	146.8m	155.8m	151.6m

Table 4: Variation in AFE transimpedance(units V/A)

	$1p\,A$	$100p\,A$	$10n\,A$
$1k\,Hz$ - mean	1.256T	35.14G	367.4M
σ	63.15G	2.437G	34.79M
$10k\,Hz$ - mean	96.33G	40.70G	429.8M
σ	7.843G	1.862G	29.96M
$100k\,Hz$ - mean	44.64G	12.98G	473.3M
σ	7.180G	653.5M	26.74M

Table 6: Monte Carlo simulation results for combined AFE and comparator. Number of pixels where the comparator output switched out of 100 simulated with different process variations.

	$1p\,A$	$100p\,A$	$10n\,A$
$1k\,Hz$	3	99	100
$10k\,Hz$	0	100	100
$100k\,Hz$	0	1	0

Table 5: Variation in AFE output noise(units V_{rms})

	$1p\,A$	$100p\,A$	$10n\,A$
mean	8.77m	1.50m	637u
σ	0.39m	0.05m	28u

Given that the operation of the AFE and comparator would be dependent on the operation of both, they were simulated together and give the results in Table 6.

3.6.3 Discussion

It has been shown that the frequency response of the AFE is affected by the DC light level as shown by He [96]. Increasing the DC light illumination increases the bandwidth of the processing. The DC photocurrent is set by the laser illumination, the object reflective properties and the optical setup. This means that there is some control over this level and may allow processing of high frequency vibrations if a high enough illumination power, or more efficient optics, can be used.

The Monte Carlo simulation has shown the expected variation in the outputs from the circuit. This is an important indication of the suitability of the circuit for full field use. It is possible for normal variation within the process to stop a pixel from operating. With a single pixel it could be discarded if it failed, however when in a camera of, say, 32x32 pixels there would be 1024 pixels and there would be some failed pixels which could not be replaced. Also, with a single pixel the controls, for example the biases, can be adjusted for that pixel, however with an

array they would have to be adjusted to accommodate the average pixel and the outliers might stop working.

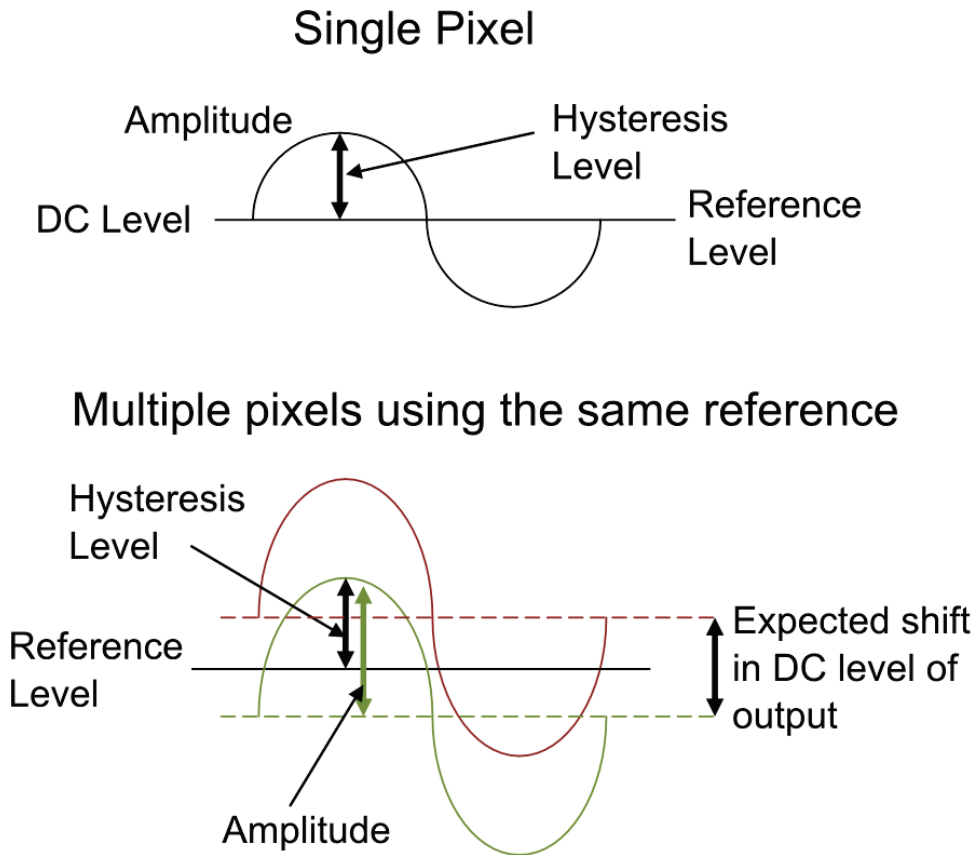

Figure 77: Sketches of the output voltage of the AFE with time.

Figure 77 shows an example of a single pixel's AFE output with the smallest amplitude which causes the comparator to switch (the smallest AC amplitude which is above the hysteresis level of the comparator). The Monte Carlo simulation has shown how the DC level, Transimpedance and noise of the AFE output would vary from pixel to pixel in a camera. Figure 77 shows an example of two pixel AFE outputs with DC levels at the highest and the lowest expected levels due to the process variation (two standard deviations above and below the mean). Given that the inverting input of the comparator is externally set, all pixels would share the same voltage and this is the reference level in the figure. As shown it would

be set to the middle of the DC level variation range. Therefore each AFE output amplitude would have to be large enough to cross this reference level by more than the comparator hysteresis.

To quantify the Monte Carlo results: the amplitude of the AFE output(V_{AC}) must exceed the worst case shift in the AFE DC level shift (2 standard deviations from the mean level, $2V_{STD}$) by the hysteresis(V_{HYS}) for the circuit to operate. See Equation 48. Given that V_{AC} is given by the AC transimpedance of the AFE (T_{IMP}) multiplied by the AC photocurrent this produces Equation 49.

$$V_{AC} \geqslant 2V_{STD} + V_{HYS} \tag{48}$$

$$T_{IMP} \times I_{AC} \geqslant 2V_{STD} + V_{HYS} \tag{49}$$

An example for a signal with a frequency of $10k\,Hz$, $100p\,A$ DC photocurrent and 10% modulation depth is as follows. The mean transimpedance would be $40.7G\,V/A$, the standard deviation in the DC level would be $141m\,V$, therefore Equation 50 shows the maximum value the hysteresis could have. This shows that most of the comparators would switch regardless of their bias and even with the worst case of process variation.

$$407m\,V \geqslant 282m\,V + V_{HYS} \Rightarrow V_{HYS} \leqslant 125m\,V \tag{50}$$

Using a transimpedance value two standard deviations below the mean ($36.98G\,V/A$) would give $V_{HYS} \leqslant 88m\,V$. This is still above the worst case comparator hysteresis level. If the frequency was increased to $100k\,Hz$ then the maximum hysteresis would become $V_{HYS} \leqslant -182m\,V$ with the mean transimpedance. This indicates that those AFEs on the lower edge of the distribution(those two standard deviations from the mean DC level) would not provide an output regardless of the

comparator. For those with exactly the mean DC level the equation would reduce to the amplitude needing to just exceed the comparator hysteresis. The amplitude of the signal with $100pA$ DC photocurrent, 10% modulation depth and $100kHz$ would be $130mV$ for the mean transimpedance which is greater than the worst case comparator hysteresis. Therefore at $100kHz$ a fraction of the pixels on a camera would work, although most likely not all. This is confirmed by Table 6.

The results show that the largest variation in the AFE output is the DC level and for a large camera this would cause some of the pixels to fail for certain inputs. A method of correcting this problem, in a full field camera, would be to generate the comparator reference voltage from within the pixel, rather than setting a global one. This might be done by passing the AFE output through a low pass filter to the comparator inverting input.

3.6.4 Conclusion

This section has simulated the AFE and the AFE combined with the comparator. It has shown the operation of the AFE and how the input signal's DC, AC and frequency affect the output. The AFE was then simulated with the comparator to show the range of input levels which would cause the comparator to switch and therefore be processed by the circuit. Simulations were done to check the expected noise on the output of the AFE and the frequency distribution of that noise. Finally Monte Carlo simulations were done to show the expected variation in the properties of the AFE and comparator if the pixel was used in a full field LDV camera.

3.7 Pulse Generator

3.7.1 Introduction

The pulse generator was one of the components designed specifically for this chip. It was designed to generate two pulses on the falling edge of the input signal. These pulses were needed to control the timing circuit and as such they needed to be non-overlapping, sequential pulses.

The pulse circuit used was a mono-stable, multi-vibrator circuit designed to give a variable width pulse on the falling edge of the input. It was based on a multi-vibrator design from [122]. The original design was for a mono-stable multi-vibrator using an AND gate followed by a series capacitor and pull-up resistor, and finally an inverter with feedback to the other input of the AND gate. This design was simplified to use only inverters and the pull-up resistor replaced with a NMOS transistor connected to ground as a current source to better suit placement on-chip. The feedback was removed as it did not affect the operation of the circuit. The final circuit diagram is given in Figure 78, and Figure 79 is a simulation to show the operation of the circuit. The circuit consisted of a chain of inverters. The first one (INV0) and the last (INV2) were minimum sized inverters. For this process this means both PMOS and NMOS transistors had widths of $0.4\,\mu m$ and lengths of $0.35\,\mu m$. The timing of the pulse was controlled by the second inverter (INV1), the poly-poly capacitor and the current source.

The input to the circuit would be the output from the comparator. As such it would be approximately a square wave.

1. Most of the time the input to INV0 would be a DC voltage, due to the nature of a square wave. The input to INV1 would therefore be kept low ($\sim 0\,V$) due to the current source and the blocking capacitor, therefore the output from the circuit would be low.

2. If the input to INV0 rose (changed from low to high ($\sim 3.3\,V$) the output from it would fall and this transient voltage would propagate through the capacitor and appear at the input to INV1, however the output from it would still remain high, and therefore the output from the whole circuit would remain low.

3. When the circuit input to INV0 falls (changes from high to low) the capacitor will be pulled high($\sim 3.3\,V$), this will in turn take INV1 input high and the circuit output will go high. The middle plot in Figure 79 shows the input to INV0 falling, and then holding, the output from INV0 rising and also holding, and the input to INV1 rising.

4. This change in the input to INV1 causes its output to go low, causing the output from the circuit to go high, as shown in the top plot of Figure 79.

5. The input to INV1 is held high by the charge on the capacitor. This charge will flow through the current source, and when the switching level of INV1 is reached the output will again go high, and the circuit output will again fall low. The relation between INV1 input and output is shown in the bottom plot of Figure 79.

INV0 provided some separation between the circuit which drives the pulse generator and the timing aspect of the circuit, sharpening the falling edge if needed. INV2 cleaned up the output from INV1 to give sharper rising and falling edges. The capacitor size, discharge rate and INV1 switching voltage all control the pulse width.

The multi-vibrator circuit produced a single pulse. To generate the two pulses required by the timing circuit two pulse generator circuits were run in series. The first would trigger on the falling edge of the comparator output. The output from the generator would be used as the first pulse and also the input to the second

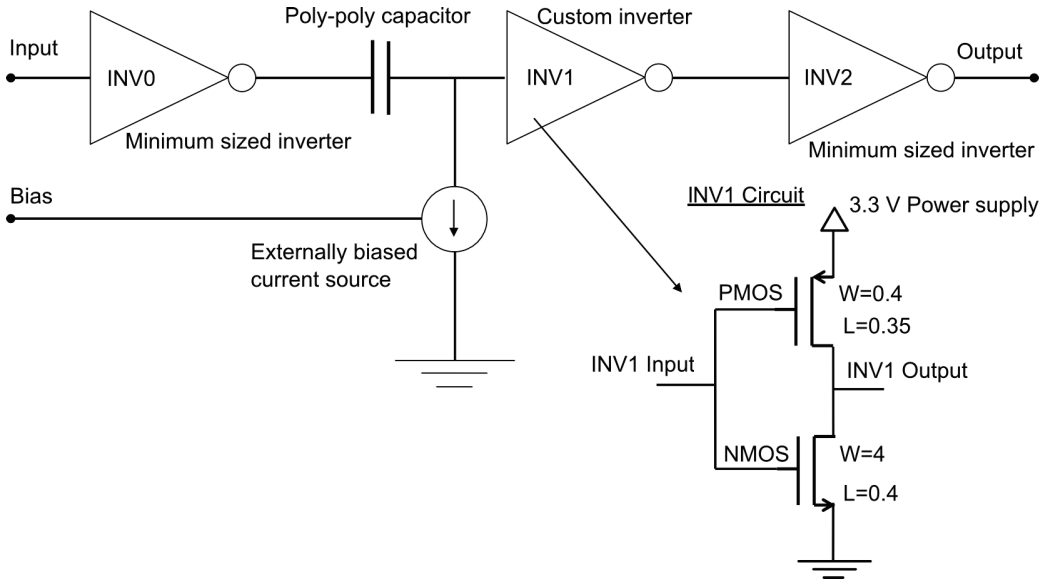

Figure 78: Circuit diagram of pulse generator

pulse generator. This generator would trigger on the falling edge of the first pulse and generate the second pulse.

During the time the pulses were active the timing circuit would not be discharging the capacitor and therefore not timing the period of the incoming signal. This introduced a systematic error into the processing circuit which could be minimised by using a short pulse compared with the smallest period of the incoming signal. The pulse would, however, need to be long enough to properly transfer the voltage levels through the timing circuit and also recharge the capacitor. Section 3.8 covers the requirements of the pulse width. It was decided that having the pulse period controllable would give an opportunity to compare the effects of different pulse widths. The current source was therefore biased using an external resistor in a current mirror configuration.

This section will simulate the pulse width with different bias resistances and show the relationship between these.

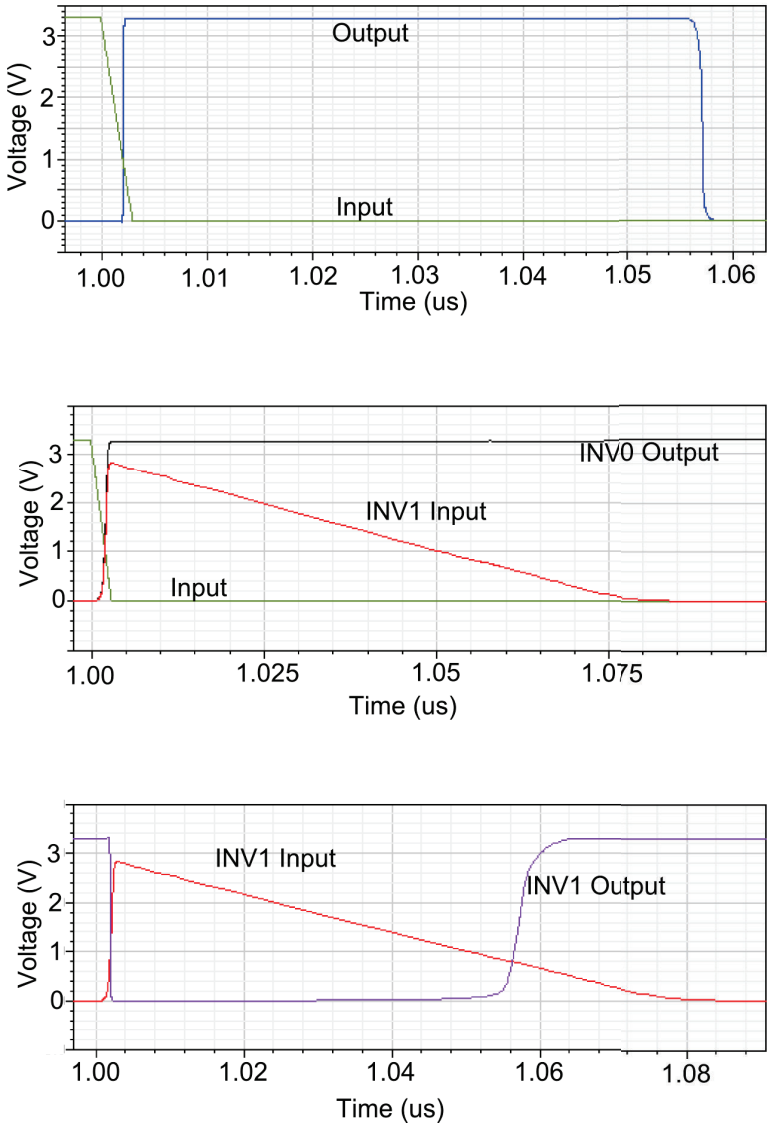

Figure 79: Transient simulation showing the operation of the pulse generator for a falling input (changing from high to low) with bias resistor of $10k\,\Omega$

3.7.2 Simulation

The simulation of the

Switching level for inverters INV0 and INV2 were minimum sized inverters to save space and they had a switching level of $1.192\,V$. INV1 was designed to have

a lower switching level to give the capacitor more time to discharge. The level was kept high enough that the low voltage non-linearities of the current source did not affect the discharge curve. These non-linearities occurred when the drain-source voltage across the current source transistor dropped below the threshold voltage at about $550mV$. INV1 had a switching level above this of $814\,mV$.

Pulse Width The pulse generator was based on the discharge of a capacitor through a current source, driving an inverter. There were a number of considerations with this design. The main property of the circuit was the relationship between the bias resistance and the pulse width. The biasing of the current source would control the width of the pulses the generator produced. This would allow different pulse widths to be tested with the circuit to see how much effect they had on the processing. The relationship between the bias and pulse width could be measured by varying the resistance while performing a number of transient simulations measuring the width of the pulses. The width of the pulse was defined as the time difference between when the rising and falling edge crossed 1.65 V. This value was chosen as the edges of the pulse were relatively rapid compared with it's width($3n\,s$ rise and fall times). The pulses also drove either single transistors, transmission gates or other inverters all of which would have different switch-on voltages.

Figure 80 shows the resulting plot of the pulse width vs. bias resistance. There is a linear relationship between bias resistance and pulse width (doubling the resistance will double the pulse width). For ease of use the plot has logarithmic axes resulting in the curve. The minimum resistance is set by the anti-static input protection of the pads of the chip which have an inherent $500\,\Omega$ series resistance.

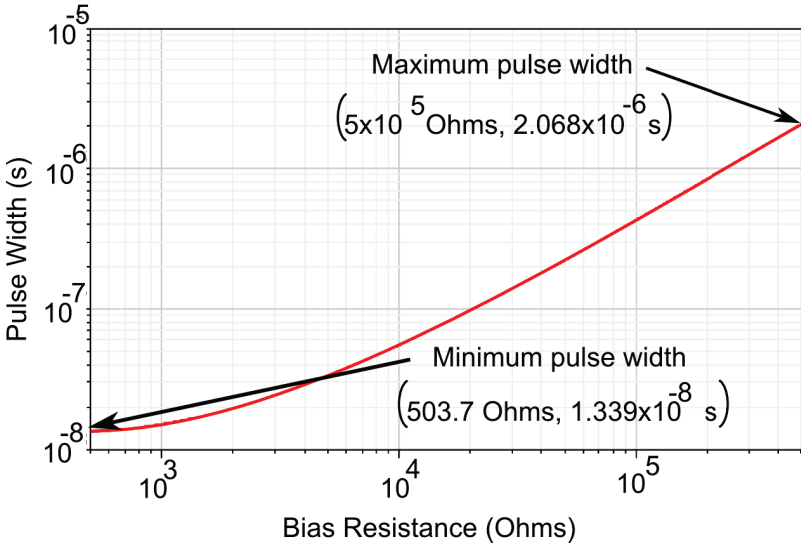

Figure 80: Pulse width of a single pulse generator circuit vs. bias resistance for the current source

Monte Carlo The next simulation was a Monte Carlo simulation of the pulse width with 500 iterations. The results from the simulation is given in Tables 7. These show that the pulse width is very susceptible to variation with between 22% and 50% variation covering two standard deviations of the results.

Table 7: Monte Carlo simulation of the variation of the pulse width of the pulse generator circuit vs. current source bias resistance.

Bias(Ω)	Mean Reset Pulse Width(s)	Reset Pulse Width standard deviation(s)
500	14.29n	1.591n
30k	142.1n	19.72n
500k	2.177u	553.8n

3.7.3 Conclusion

Simulation of the pulse generator has shown that it produces pulses of variable width controlled by an external bias resistor. The pulse width shows a linear relationship with the bias resistance. The minimum pulse width is limited by the resistance of the input protection pad circuitry on the chip and is $500\,\Omega$. This would give a pulse width of $13.4n\,s$. The pulse width then increases to $2.07u\,s$ at

$500k\,\Omega$. It was also found during transient simulation that the rising and falling edges of the pulse occur within $3n\,s$.

The Monte Carlo variation would become important when considering the effects of the pulse width on the operation of the timing circuit. The timing circuit will be covered in Section 3.8 and it will be shown that there is a minimum value for the control pulses such that the timing circuit can respond. The pulse width is also globally controlled so it would have to be set large enough such that the process variation did not cause some of the timing circuits to get pulses with too short widths. The pulse width also introduces an error into the timing circuit and therefore need to be kept as small as possible, requiring a trade between operation of most of the pixels and error introduced into those which do work.

3.8 Timing

3.8.1 Introduction

The timing circuit formed the main processing component in the circuit. It was designed to produce a voltage dependent on the time between a set of control pulses. The component was specifically designed for this chip and the circuit was kept as simple as possible for prototyping purposes. The operation of the circuit is given in Section 3.3.3 and the circuit diagram is repeated here for convenience in Figure 81. There were several aspects of the timing circuit which would affect the final pixel.

- Controllability of the period to voltage conversion. This is the range of discharge currents that could be set using the external bias resistor. This range would determine the range of period to voltage conversion constants (k_{pv}) the processing could operate with, see Section 3.2.

- Minimum time needed to charge the timing capacitor. This was the time

required for the timing capacitor to charge from the voltage it had reached during the previous cycle to the supply voltage. The time given to charging would be set by the width of the control pulses. If they were too short the initial voltage of the capacitor might vary depending on the last voltage it had discharged to. If the capacitor had discharged by only $0.5\,V$ then it might easily charge during a short reset pulse, however if it had almost reached $0\,V$ then it would need more time to charge. Long pulses would also introduce a long wait during each period when the capacitor was not discharging and thus distort the output.

- Minimum time needed to transfer the timing capacitor voltage through the sample/ hold architecture. This consisted of the sample/ hold buffer and the output buffers. This was the time it took for the various switches to set to their appropriate positions and then the voltage level to propagate through the buffers. If the sample pulse was not long enough to accommodate these changes the output voltage would not match the capacitor voltage.

- Leakage current through switches and transmission gates. Ideally the switches within the circuit would keep the sampled voltage separate from the capacitor voltage during the timing phase the circuits operation. Any leakage current through the switches would affect the output accuracy, especially when timing long period input signals.

- Buffer offset voltages. Due to their design the buffers within the circuit displayed a voltage difference between the input and output voltages.

3.8.2 Simulation

Figure 82 shows how the linearity of the discharge curve varies with the voltage of the capacitor. The upper trace shows the voltage across the timing capacitor and

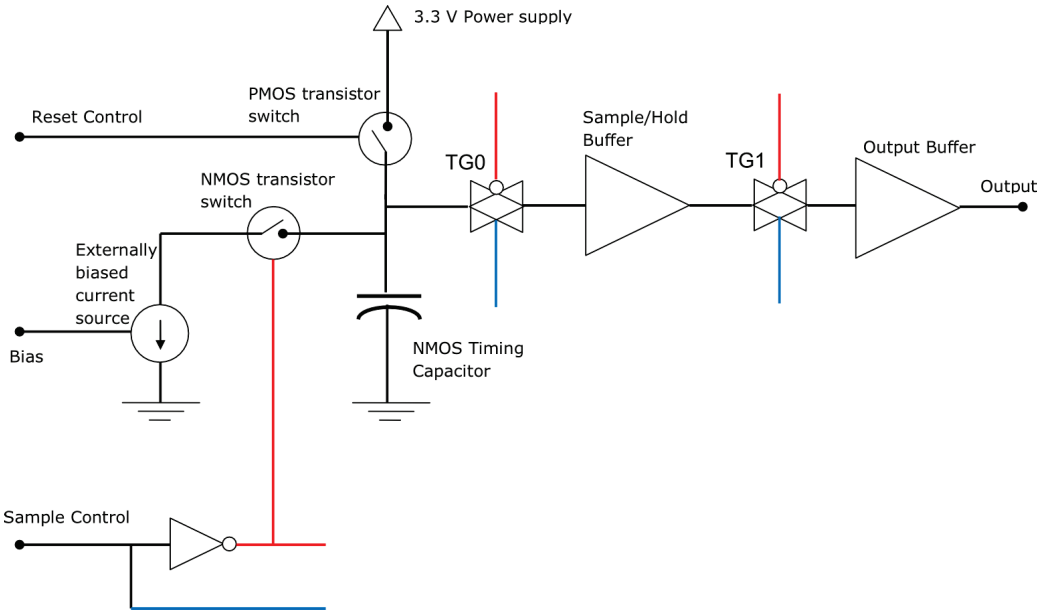

Figure 81: Schematic for timing circuit

the lower trace shows the derivative of this voltage. The sharp rise at the start of the curve is the result of the reset signal connecting the capacitor to the supply to charge it. Figure 82 shows that the gradient of the curve is almost linear until $12.43\,\mu s$ where it deviates from the initial value by 6%. This occurs at $539\,mV$ and was chosen as this was the smallest deviation which contained all of the previous curve. This gives a quick figure for the linearity of the curve and is not intended to be rigorous. Preliminary simulation had shown that several components would affect the non-linearity of the timing circuit conversion and the final operation would be affected by them all. The value of $539\,mV$ gives an approximate level at which the non-linearities of the discharge curve start to take affect.

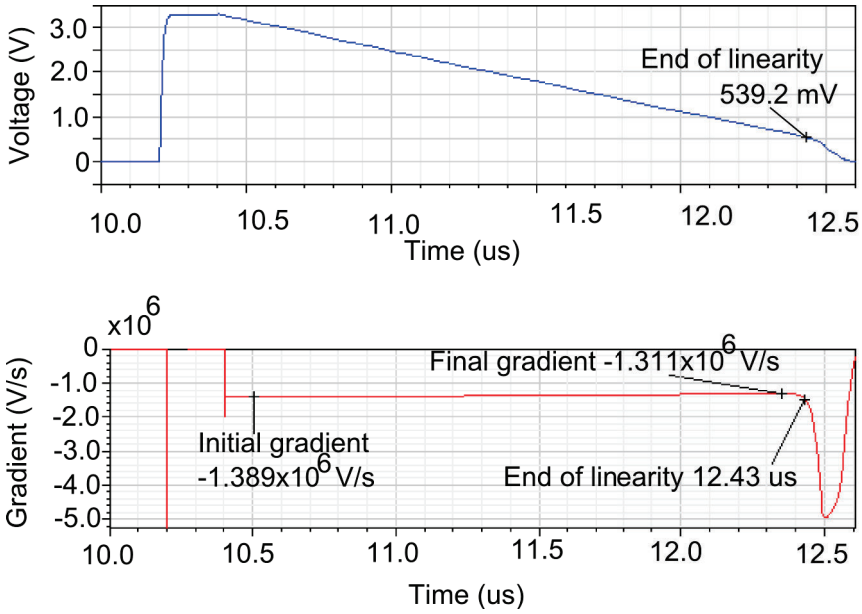

Figure 82: Capacitor voltage linearity range of the timing capacitor in the timing circuit.

Systematic Errors During the discharge time the circuit is relatively quiescent and the errors introduced to the output are entirely associated with the discharge gradient. During the sample and reset pulses, however, are the times when very rapid voltage changes occur within the timing circuit and where a number of other errors could affect the output. These errors would distort the period to voltage conversion response of the circuit.

The purpose of this simulation section is to investigate the operation of the timing circuit. Section 3.2 gave a derivation of the ideal equation for a period to voltage converter.

This section will concentrate on showing that the timing circuit exhibits the required period to voltage conversion characteristic and on the sources of error within the circuit and the method. It is hoped that this will allow future work to address specific areas of the timing circuit to fulfill a given application.

Capacitor Discharge Rate The first aspect of the timing circuit simulated was the discharge characteristics of the timing capacitor. Figure 83 shows a simulation of the discharge curve of the timing capacitor with the control pulses for reference. The timing capacitor voltage has a finite rise time which corresponds to the time needed to charge. It has an almost linear discharge curve until the end where the current source transistor no longer has a large enough drain-source voltage to provide a constant current.

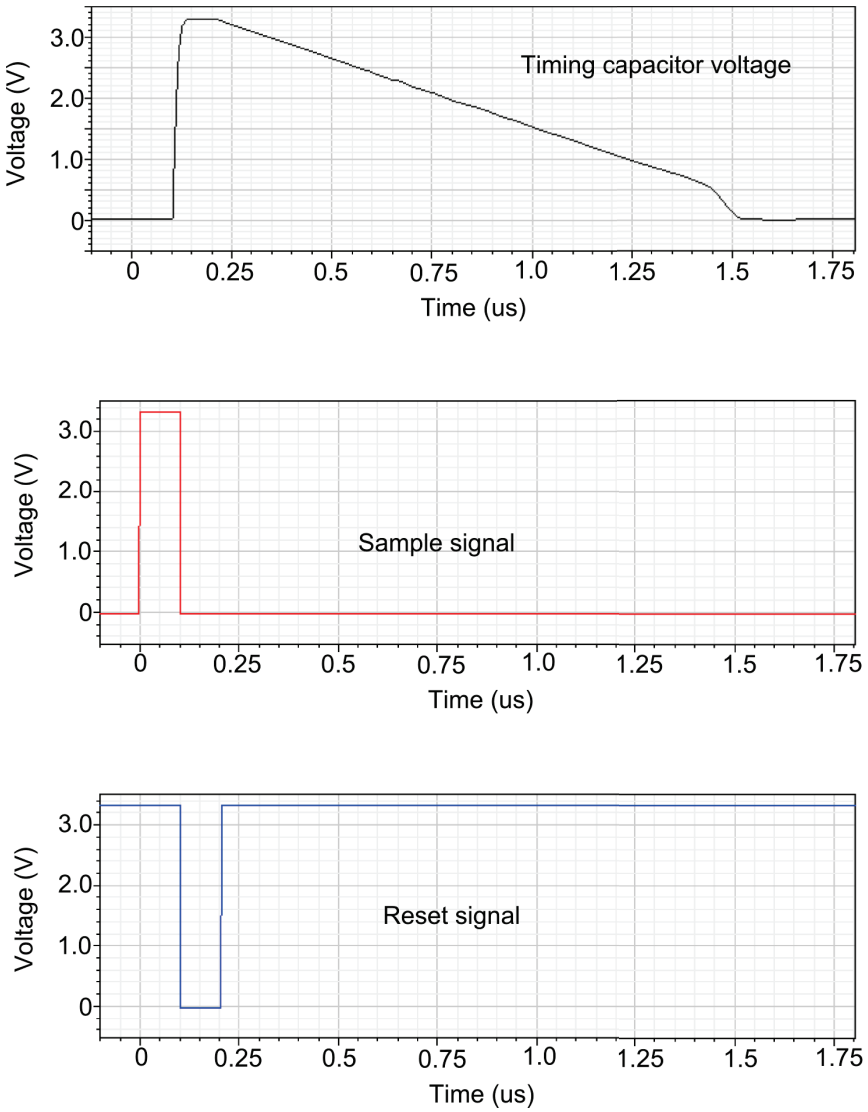

Figure 83: Simulation of relation between timing capacitor and control signals for the timing circuit.

The period to voltage conversion is dependent on the gradient of the timing capacitor voltage and this gradient corresponds to k_{pv} from Section 3.2. The range of values of k_{pv} the timing circuit could implement was found by sweeping the bias resistance from $500\,\Omega$ to $500k\,\Omega$ and plotting this against the resulting value of k_{pv}. The gradient was found as the average value between 3.05 V and 0.55V and is shown in Figure 84. The gradient is always negative as it is the discharge of

the timing capacitor. This plot shows that the minimum discharge time for the capacitor is $1.36\,\mu s$ and the maximum time is $1.1\,ms$.

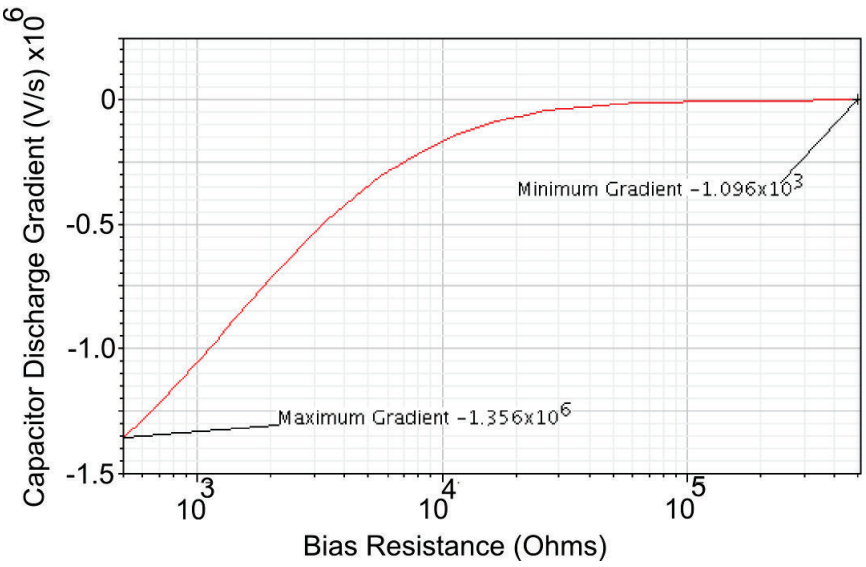

Figure 84: Timing capacitor discharge gradient range for the timing circuit.

For a more illustrative plot these results can be converted to a plot relating input frequency to output voltage. For each value of k_{pv} the input frequency required to give half the supply voltage can be plotted against the bias resistance needed for that value of k_{pv}. This is shown in Figure 85. For example a bias resistance of $10k\,\Omega$ would set the timing capacitor discharge such that an input frequency of $80k\,Hz$ would result in an output voltage of $1.65\,V$. The plot also shows upper and lower frequencies. These correspond to the frequencies that would give an output voltage of $3\,V$(a drop of $0.3\,V$ from supply) and $0.3\,V$(drop of $3\,V$) respectively. They are included to show that the bandwidth available also depends on the value of k_{pv}.

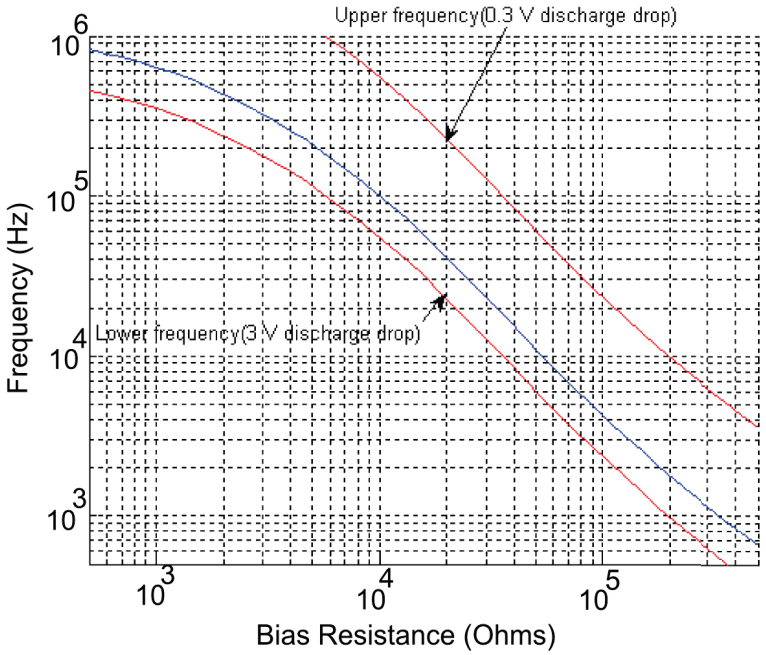

Figure 85: Relationship between the input frequency which corresponds to a given output voltage and the bias resistance of the timing circuit. The blue line represents an output of 1.65 V, the two red lines represent output of 3 V(top) and 0.3 V(bottom)

Minimum Reset Pulse Width Ideally the control pulses to the timing circuit should be as short as possible to introduce the least error and allow higher frequencies to be timed. There is however a minimum duration for each pulse to allow all parts of the timing circuit to respond. The minimum reset pulse width is set by the time it takes for the timing capacitor to charge to the supply voltage. This was found by simulating the time required for the timing capacitor to charge from $0\,V$ to within 1% of $3.3\,V$. It was found to be $53.2n\,s$.

Minimum Sample Pulse Width Similar to the reset pulse the sample pulse also had several factors which affected its minimum width. These factors are all associated with the propagation of the last timing capacitor voltage, before the sample pulse, through the sample/ hold buffer chain and switches to the output. There were five components that formed the sample hold architecture.

Their order was Transmission Gate (TG), sample/hold buffer (OTA Buffer), a second TG, output buffer and OPAMP. Although not shown in Figure 81 as it is not part of the pixel processing, the OPAMP is needed to drive large loads off-chip. For a camera only one OPAMP would be needed, and connected to whichever pixel was being sampled.

The timing buffer was an OTA buffer, modified to have a lower frequency roll-off at $16k\,Hz$. Although the sampled voltage would have to propagate through the whole chain, only the two TGs and the sample/ hold Buffer would have to respond within the sample pulse, the others would have the time between sample pulses to track the input voltage (the input signal frequency).

There are two cases for the voltage change during a sample pulse: either the new timing capacitor voltage is higher than the previous voltage, causing a rising edge to propagate through the buffer chain to the output, or the capacitor voltage is lower that the previous voltage, causing a falling edge to propagate through the buffer chain. These will be looked at separately for the worst case: a change of $3.3\,V$, the supply voltage.

Rising Edge To measure the maximum rise times for the buffer chain a sudden change from $0\,V$ to $3.3\,V$ was simulated and the time taken for each stage in the chain measured. The TG, sample/ hold buffer and output buffer all showed reasonably fast rise times, measured as the time for the output voltage to get within 1% of the final voltage, see Table 8. This table also shows the maximum voltage the buffer could output. These values would limit the useful range of voltages that could be outputted and therefore the range of frequencies that could be demodulated. The OPAMP has a relatively linear rise-time until about $2.7\,V$ which matched the output buffer, however the rise time was much slower above that. The input changed very rapidly from $0\,V$ to $3.3\,V$. The output

from the OPAMP tracks its input (the output from the output buffer) until it reaches $\sim 2.7\,V$ then the rate of change of the output from the OPAMP slows significantly. This limits the voltage range to less than $2.7\,V$ for fast voltage changes.

Table 8: Rise times for the timing circuit buffers and the transmission gate

Component	Time to 1% Final Voltage	Final Voltage
Transmission Gate	$4n\,s$	$3.3\,V$
sample/ hold buffer	$19.7n\,s$	$3.203\,V$
Timing Buffer	$447u\,s$	$3.201\,V$

Falling Edge The fall times were much more uniform and are summarised in Table 9. Again there were some delays in the output reaching the final voltage. Previous tests had shown that the timing capacitors discharge curve is increasingly non-linear below $539m\,V$ so the time for the voltage levels to reach $500m\,V$ is shown as this covers the likely useful range of the timing circuit.

Table 9: Fall time for timing buffers, transmission gate and OPAMP

Component	Time to 1% Final Voltage	Final Voltage	Time to $500m\,V$
Transmission Gate	$1.06n\,s$	$0\,V$	-
sample/ hold buffer	$19.3m\,s$	$22.57m\,V$	$51.2n\,s$
output buffer	$36.0m\,s$	$36.28m\,V$	$464u\,s$
OPAMP	$12.9m\,s$	$73.59m\,V$	$464u\,s$

Charge Transfer to OTA Buffer One final consideration for the sample pulse width is the time taken for a small part of the charge on the timing capacitor to flow to the input capacitance of the sample/ hold buffer. This time was found to be less than $20n\,s$ and therefore approximately the same length as the minimum pulse width and unlikely to be a limiting factor unless faster buffers are used.

Conclusion These simulations demonstrate the variable slew rates of the various buffers within the timing circuit. In the middle range of their operation the

buffers all change voltage relatively quickly with the slowest rise rate of $7.2kV/s$ and the slowest fall rate of $5.6kV/s$. These slew-rates are applicable within the range $\sim 500mV \rightarrow\sim 3.1V$. Outside this range the buffers may still operate, but with increasing delay times for voltage propagation, and therefore errors in the demodulation process. The sample/ hold buffer and TGs can respond at a minimum fall rate of $52MV/s$ and rise rate of $162MV/s$. Given that the minimum pulse width is $13.4ns$ these components can cover 20% of the supply range within that time.

It should be noted that the supply range voltage change used in the simulations would be unlikely to occur under normal circumstances as it would require that the vibration frequency be close to the carrier frequency. For the Doppler bandwidth to not exceed the bandwidth available, set by the Bragg Cells, the maximum amplitude of the vibration would be exceedingly small. Given that the Doppler bandwidth is affected by the frequency and amplitude of the vibration, to gain a useful amplitude the carrier frequency will need to be larger than the vibration frequency. As such there will be several cycles of the carrier for each vibration cycle. This effectively samples the vibration waveform several times and as such the voltage difference between each sample would be a fraction of the supply range.

Charge Sharing and Charge Injection One source of error during the sample signal time is the transfer of the voltage from the timing capacitor to the input of the TG and OTA Buffer. Figure 86 shows a magnified region of the timing capacitor voltage and sample buffer's input and output. The point labeled Timing capacitor voltage before sample shows the time of the sample pulse start. The horizontal line shows the voltage the timing capacitor had reached. This voltage drops initially, matching the input to the OTA Buffer, due to a small portion of the charge on the timing capacitor flowing to the input of the buffer to equalise their

voltages. This voltage drop is proportional to the initial voltage (the final voltage the timing capacitor had reached), and can be reduced by having a larger ratio between the timing capacitance and the input capacitance of the buffer. Making the input capacitance of the buffer smaller can be difficult while maintaining the operation of the buffer, while increasing the size of the timing capacitor would require more space and decrease the fill factor of the pixel. For this circuit the timing capacitor was a $10\,\mu m$ x $10\,\mu m$ NMOS capacitor.

Related to the charge sharing is charge injection from the closing of the TG. This is also shown in Figure 86 as the input voltage of the sample/ hold buffer is higher than the voltage during the sample pulse and the initial voltage. This charge injection comes from the channels within the transistors of the TG closing. The charge injected to the input of the buffer is small, however the capacitance of the buffer input is also small so a relatively large voltage change results.

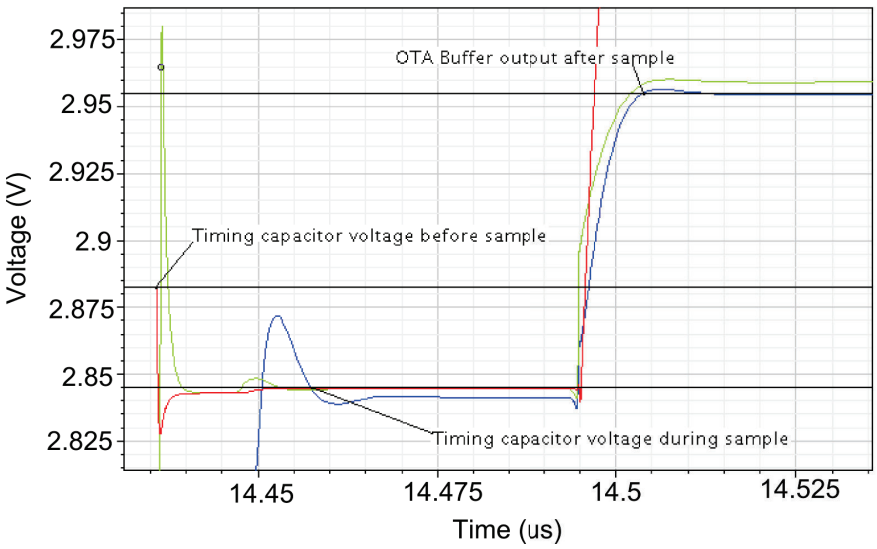

Figure 86: Charge sharing between timing capacitor and buffer input for the timing circuit

Buffer DC Shift and Switch Off Leakage Current Two more errors can be seen over the whole duration of the operation of the timing circuit. They are

the DC voltage shift between the input and output of the buffers and the small current which leaks through the TGs when they are closed. These are shown in Figure 87. This figure was generated using a simulation with a bias resistance of $10k\,\Omega$ and an input frequency of $80k\,Hz$. The sample and reset signals are added for reference. During the sample pulse the switches are open and have minimal DC shift across them. This is shown by TG1's input and output being the same voltage. The difference between the voltages labeled is due to the buffers. The sample/hold buffer adding $4m\,V$, the output buffer removing $9m\,V$ and the OPAMP adding a very small voltage. These voltage shifts are always in effect and should be considered when looking at the available output voltage range of the circuit.

After the sample pulse has finished it can then be seen that TG0's output starts to increase. This is due to a very small current($<600f\,A$) leaking through the TG from the timing capacitor voltage(now approximately $3.3\,V$) to the input capacitance of the sample/hold buffer. This causes the output to gradually track the capacitor voltage.

Control Pulse Width Variation It has been shown that the width of the control pulses has an effect on the timing circuit's operation. Aside from the operation of the timing circuit electronics, the pulse width has an effect on the period to voltage conversion. During the pulse's active time the capacitor is not discharging and therefore not timing the input signal. This results in the voltage the capacitor reaches during each cycle being slightly higher than predicted from the value of k_{pv}, See Equation 51. In this equation T_0 is the time required for both control pulses to operate.

$$V_{out}(t) = 3.3 - k_{pv}(T(t) - T_0) \tag{51}$$

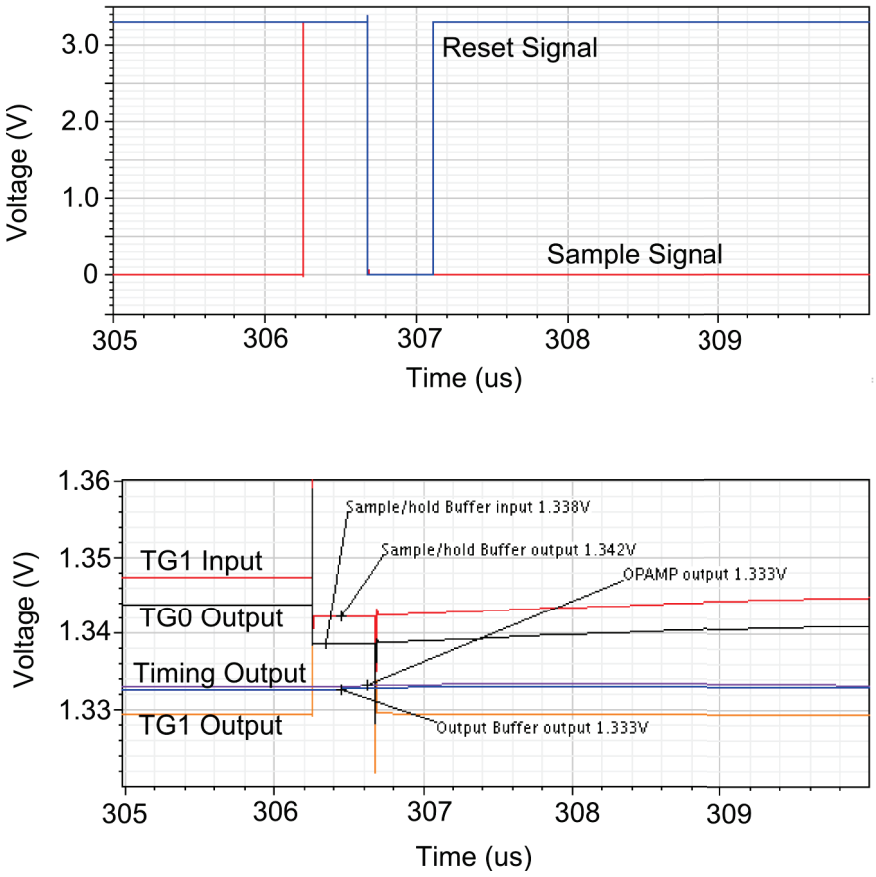

Figure 87: DC shift through buffers and switch leakage current

As mentioned in Section 3.2 the value of k_{pv} would practically be set such that a given constant frequency would produce a certain output voltage. Thus for that frequency the value of the pulse width would be compensated for by a slightly larger value of k_{pv}than would otherwise be required, giving the required output voltage. Figure 88 shows a Matlab simulation of the period to voltage conversion characteristic for the timing circuit if the pulse width were taken into consideration. The red plot shows an ideal conversion(zero pulse width) to give $1.65\,V$ output for $25k\,Hz$ input. The green plot shows the conversion characteristic with the same calibration point but with a pulse width of $15u\,s$. The value of k_{pv} has been set to give the correct value for $25k\,Hz$ but the conversion starts to diverge from the ideal one for frequencies above and below this value.

If the pulse width were then to vary, as shown by the blue and black plots, the effect is to move the calibrated value. An example of variation of the pulse width would be due to process variation in a camera chip. With many pixels all sharing the same timing bias, and therefore the same value of k_{pv}, any variation in the pulse generators would change the pulse width for each pixel and therefore the conversion characteristic.

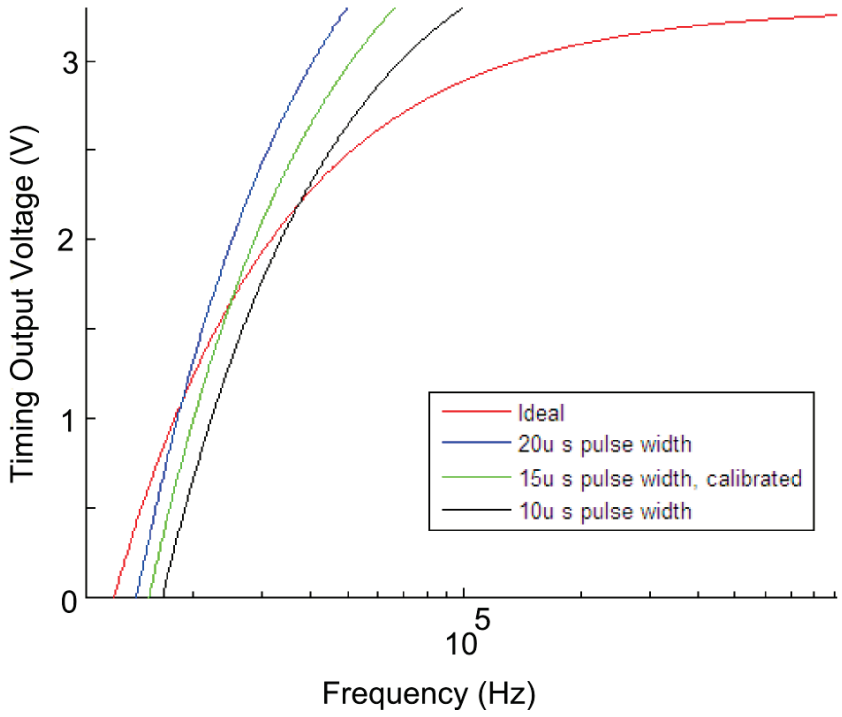

Figure 88: Simulation of timing period to voltage characteristic with variation in control pulse width.

Processing Systematic Errors The processing has been shown to exhibit several sources of error. These will affect the voltage the timing circuit will produce for a given input frequency, or input frequency modulation, and therefore distort the period to voltage conversion response of the circuit to some degree. To quantify these errors a simulation of the period to voltage conversion of the timing circuit was done and compared with a theoretical ideal one(see Section 3.2). A plot of the ideal and simulated input frequency to voltage curves is shown in Figure

89 for three values of the conversion constant(k_{pv}) chosen from Figure 85. For each value of k_{pv} the required bias resistance was therefore known. The figure shows close agreement between ideal and simulated for frequencies which produce voltages close to the middle of the output voltage range ($> 0.5\,V$ and $< 2.7\,V$). Outside this range, either above or below the frequency to voltage conversion rolls off. This means that changes in frequency do not significantly change the output voltage. The upper roll off is a result of the buffers which make up the buffer chain working best in the middle of the supply voltage range. As previous simulations have shown the OPAMP output does not exceed $2.7\,V$ for fast transients which would be expected as when the voltage output is close to $3.3\,V$ the input frequency is high. The lower roll off is due to the discharge curve of the timing capacitor. As can be seen in Figure 83 the discharge curve falls off rapidly under about $0.5\,V$.

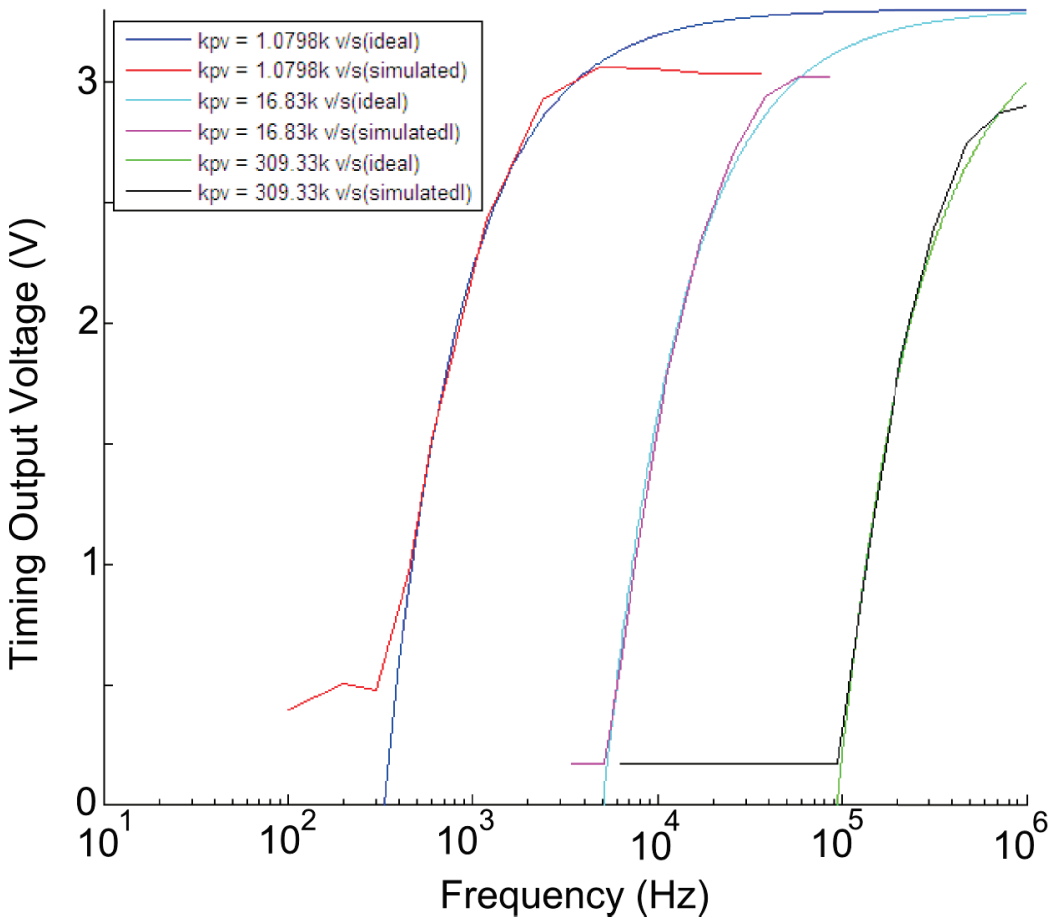

Figure 89: Output voltage vs. frequency showing non-dualities of timing circuit

Transient Operation The next set of tests were transient simulations to check the timing circuit in operation. A theoretical, ideal timing circuit was designed in Matlab to provide a comparison for the simulation results. This code was designed to use the same voltage range and method as the timing circuit, but without the non-idealities of the various components.

The simulations were run with synthetic control pulses of width $110n\,s$, without using the pulse generators to test only the timing circuit. For these simulations k_{pv} was set to $45.75k\,V/s$ or a timing bias resistance of $26.54k\,\Omega$ and several simulations(both Matlab and Cadence) were run with slightly different input signals. The input was created for a theoretical vibrating object with frequency of $500\,Hz$

and a range of amplitudes. The other variable that could be set was the carrier frequency(δf). This would set both the voltage output for a stationary object and the maximum Doppler signal bandwidth available before the output would distort.

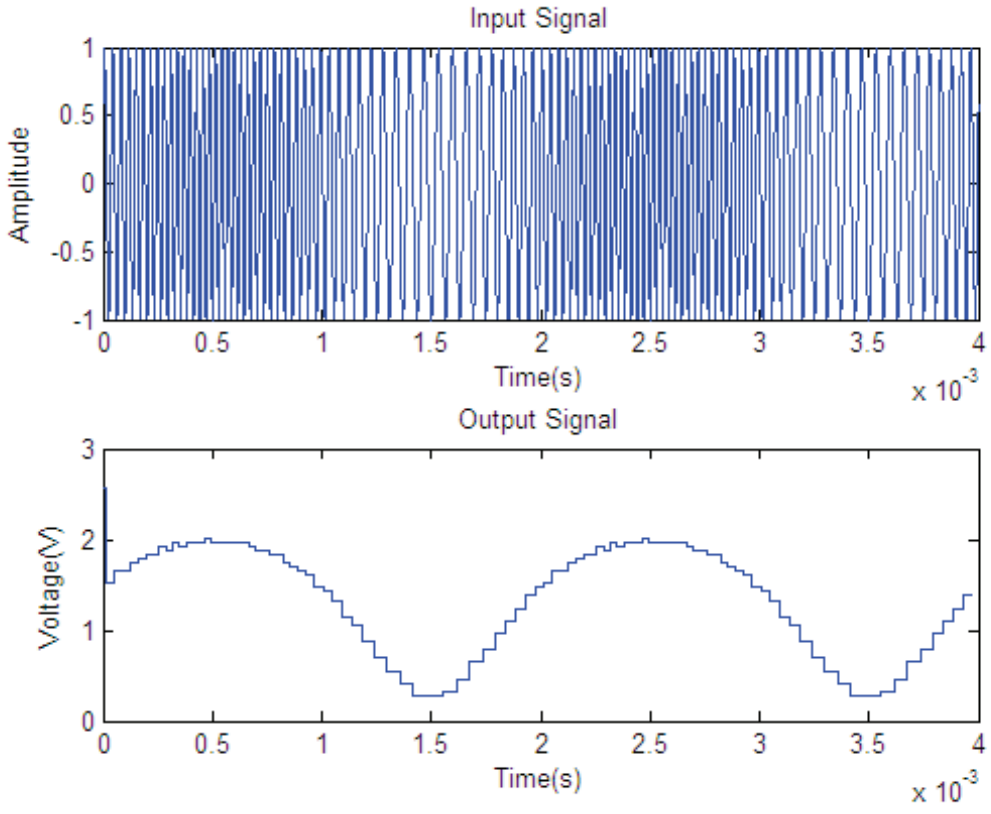

Figure 90: Transient Matlab simulation of timing demodulation output and the input signal. Vibration of frequency of $500\,Hz$ and amplitude of $500n\,m$ with $\delta f = 25k\,Hz$ and $k_{pv} = 16.83k\,V/s$

Figure 90 shows a Matlab simulation of an input and output signal for the timing circuit. The input is a frequency modulated signal generated to simulate a vibrating object vibrating with $500\,Hz$ with an amplitude of $500\,nm$. The output is a series of steps resulting from the zero crossings of the input signal. It can be seen that high frequencies produce shorter steps of higher voltages while at low frequency there are larger steps. The frequency of the output signal is equal to the

frequency of the vibration signal, and the voltage level at either peak represents the amplitude. A number of Matlab and Cadence simulations were done with a range of vibration amplitudes from $500nm$ to $2um$. The resulting outputs are shown in Figure 91.

The $500nm$ plot shows a signal with a smaller bandwidth than δf, and as such the period to voltage response does not distort the output much. The $1.5um$ and $2um$ plots show what happens when the bandwidth of the Doppler signal approaches $2\delta f$. There are three effects which come into play at lower frequencies.

Timing Capacitor Discharge Curve The first and most obvious is that at low frequencies (and low voltages) the output clips to about $250mV$. This is due to the discharge curve of the timing capacitor as shown in Figure 83 and the low voltage response of the OPAMP. The timing capacitor discharges much quicker in this region, however the OPAMP responds very slowly to voltage changes below about $74mV$ (see Table 9) and the combination of this is that the output clips at low frequencies.

Period to Voltage Conversion Characteristic The trace in Figure 90 shows a Matlab simulation which is not affected by buffer output characteristics. It shows that the change between levels is larger at low frequencies. This is because the timing circuit gives a linear response to period and therefore shows a $1/f$ response to frequency.

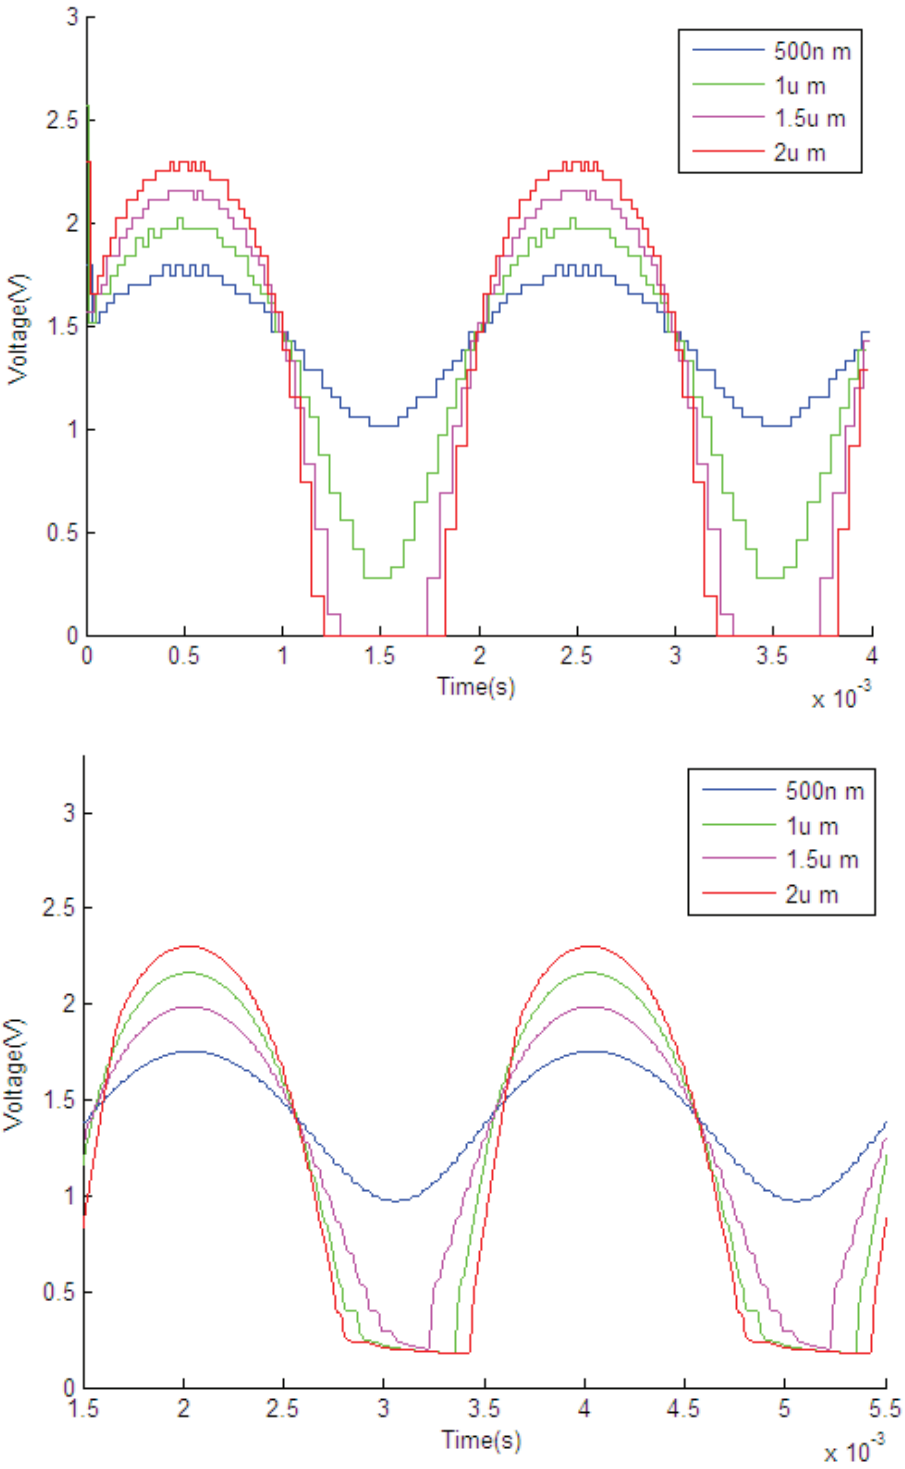

Figure 91: Transient Matlab simulation of timing demodulation output. Vibration of frequency of $1k\,Hz$ and amplitudes of $500n\,m$, $1u\,m$, $1.5u\,m$ and $2u\,m$ with $\delta f = 25k\,Hz$ and $k_{pv} = 45.75k\,V/s$

Zero-Crossing Sampling Errors The frequency swing below δf approaches $0\,Hz$(half the bandwidth is below δf and half the bandwidth is above). This can be seen in the time trace of the input signal in Figure 92. Comparing the $1.5u\,m$ signals and the $2u\,m$ signals in this figure, the $2u\,m$ trace shows sections of very low frequency at $1.5m\,s$ and $3.5m\,s$. At these points there are fewer zero crossings and therefore larger gaps in both time and voltage between levels in the output. These are the sections of Figure 91 where the voltage approaches $0\,V$.

The period is only calculated every cycle of the frequency modulated carrier, every zero crossing. When the zero crossings are far apart the phase of the carrier is varying more slowly, however the rate of change of phase of the vibration signal is not related to this. As such the zero crossings start to average more and more of the vibration signal. Another way of thinking about this is that the system implements a variable sample rate, each zero-crossing producing a sample, except that each sample integrates the phase variation of the vibration since the last sample. As the samples become further apart compared with the period of the vibration signal more of it is integrated and the value starts to differ significantly from a spot value as would be produced by a normal sampling method(an Analog to Digital Converter). The extreme example of this is that if the samples were of the order of one period of the vibration signal apart they would effectively average to zero the vibrations phase contribution to the phase of the carrier. Figure 91 also shows that voltages below about $500m\,V$ start to clip due to the buffers, but the upper peaks remain undistorted showing the high dynamic range this timing method gives. The buffers also have finite bandwidth so cannot reproduce the steps shown in the Matlab simulation. This provides some filtering to the output and thereby reducing the noise on the output signal.

The errors mentioned above, related to having a low frequency input, are less important when considering those frequencies higher than δf. The synthetic

sample-rate is high compared with the frequency of the signal resulting in the samples more closely representing a spot value for the vibration signal. The time and voltage difference between samples is small and required slew rate of the sample/ hold part of the architecture is correspondingly less stringent.

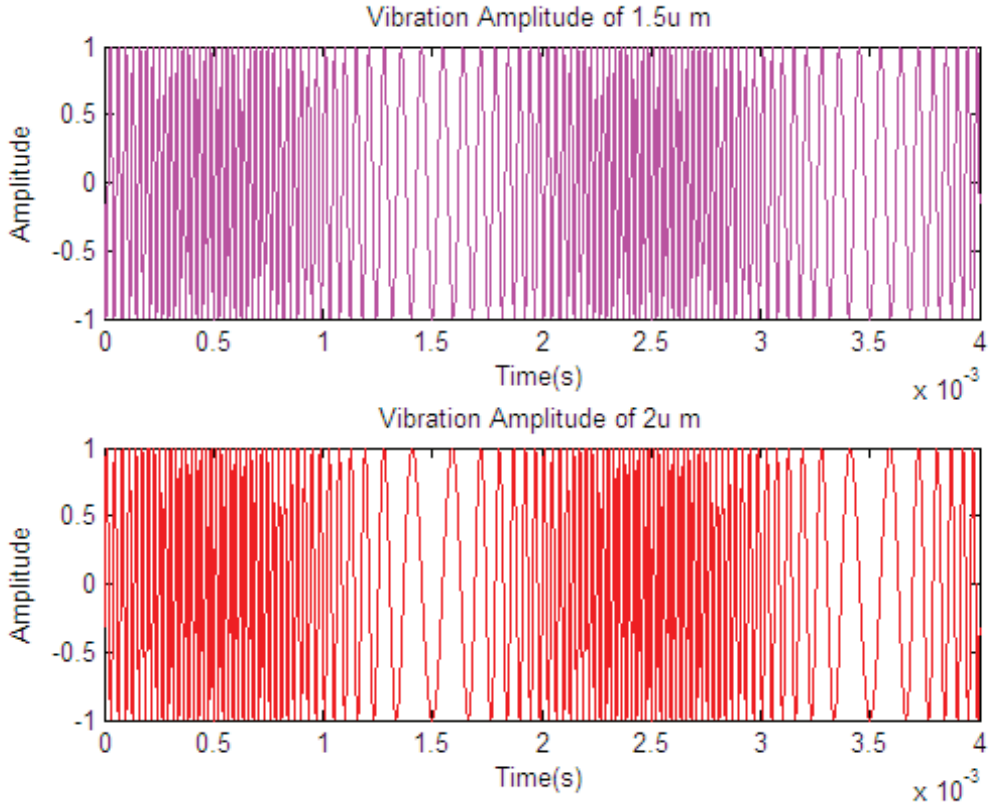

Figure 92: Transient Matlab simulation of timing circuit input. Vibration of frequency of $500\,Hz$ and amplitudes of $1.5u\,m$ and $2u\,m$ with $\delta f = 25k\,Hz$ and $k_{pv} = 45.75k\,V/s$

Scaling The Output to Give The Vibration Amplitude The relationship between an output voltage and the amplitude of the vibration can be found using the equation for the AC component of the input to the timing circuit as shown in Equation 52. This time varying part of the input which has been clipped by the comparator into a square wave.

$$3.3 \times square\left(\omega_c t - \frac{4\pi A_{osc}}{\lambda} cos\left(\omega_{osc} t\right)\right) \tag{52}$$

Considering the phase of this equation as $\phi\left(t\right) = \omega_c t - \frac{4\pi A_{osc}}{\lambda} cos\left(\omega_{osc} t\right)$, the frequency can be found from $f\left(t\right) = d.\phi\left(t\right)/dt$ giving Equation 53. The timing circuit will effectively average the frequency between zero-crossings, however just looking at the highest peak the instantaneous frequency would be given by Equation 54. Substituting this into Equation 39, using $T = 1/f$ to give the period, and then rearranging to give the amplitude gives Equation 55. This equation relates the maximum voltage ($V_{out_{max}}$) to the amplitude of the displacement of the vibrating surface (A_{osc}).

$$f\left(t\right) = \omega_c + \frac{\omega_{osc} 4\pi A_{osc}}{\lambda} sin\left(\omega_{osc} t\right) \tag{53}$$

$$f_{max} = f_c + \frac{f_{osc} 4\pi A_{osc}}{\lambda} \tag{54}$$

$$A_{osc} = \frac{\lambda}{4\pi}\left[\left(\frac{\frac{k_{pv}}{V_s - V_{out_{max}}} - \delta f}{f_{osc}}\right) - 1\right] \tag{55}$$

The output signals in Figure 91 can be scaled to compensate for the nonlinear frequency to voltage conversion of the timing circuit. This uses Equation 55 and is shown in Figure 93 for the Cadence simulation results. The output from the timing circuit is related to the input vibration signal. The frequency of the carrier is modulated by the Doppler effect which is related to the velocity of the reflecting surface. The velocity is the derivative of the displacement of the reflecting object and this introduces an unavoidable phase shift of $\pi/2$ between the vibration displacement and the output from the timing circuit. For comparison Figure 93 was generated with the vibration displacement shifted to compensate

for the phase shift of the output. The plots show a close match for the upper peak and a close match for the lower peak while the signal remains within the bandwidth of the carrier.

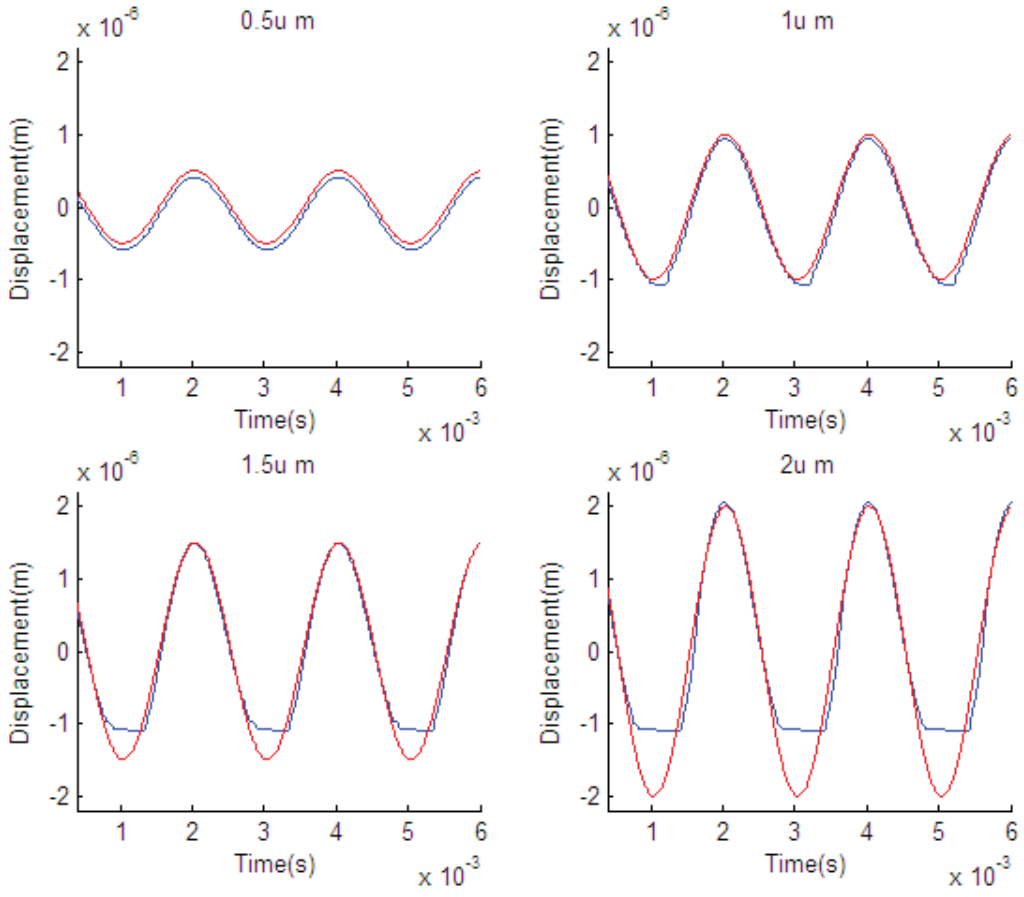

Figure 93: Matlab simulation of scaled outputs from the timing circuit shown with the original vibration signal overlain. The vibration signal has been phase shifted by $\pi/2$ so that the agreement between the scaled output and the vibration can be see.

Monte Carlo The final simulation for the timing circuit checked how much variation in the output voltage would be expected due to process variation during manufacture. This Monte Carlo simulation (using Cadence) used 500 iterations and checked the mean value of the output voltage for a fixed input frequency. Each iteration has the same vibration frequency and amplitude, and the same values for

k_{pv}, f_c, etc. The results are shown in Figure 94. This test is important to check how well the circuit is suited to use in a camera. It checks how much variation there would be across the camera due to process variation in each pixel. If a camera with 500 pixels were made and tested on an object of the same vibration frequency and amplitude then the outputs from the different pixels would be expected to show the same spread as Figure 94.

The large peak below $0.5\,V$ shows signals which have clipped to the minimum output voltage of the buffers and would not provide a correct output. The spread shows that each pixel has a different value of k_{pv} for a given bias resistance, which was kept constant for the simulation. Further simulation showed that the problem was associated with the current flowing through the current source transistor. The different outputs were due to the slope of the discharge curve of the timing capacitor. The bias voltage to the current source did vary, however it did not correspond to the variations in the current in the current source. More simulations were run and it was found that the major source of variation was the timing capacitor and the current source transistor.

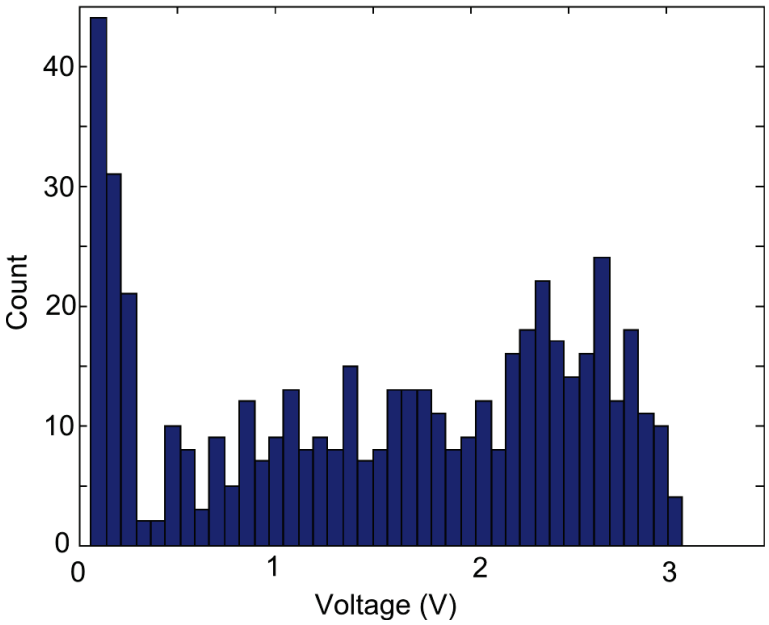

Figure 94: Monte Carlo simulation of variation in output voltage for a frequency of $25k\,Hz$ and with a control pulse width of $110n\,s$.

The Monte Carlo simulation was then run again with the entire pixel circuit. Figure 95 shows the DC output from the circuit for a fixed input frequency with $100p\,A_{DC}$ and 10% modulation depth. The nominal output voltage for these settings would be $1.47\,V$. This shows similar variation to that of the timing circuit alone indicating that the variation is mainly due to the timing circuit.

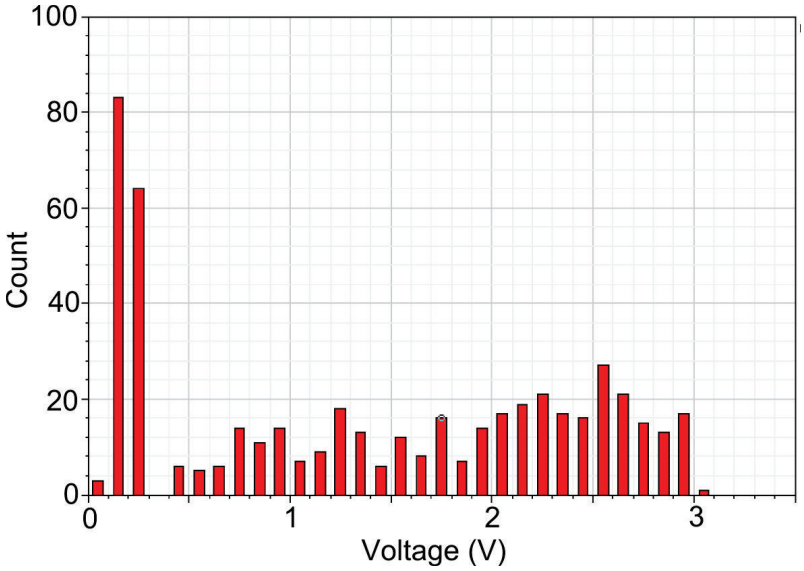

Figure 95: Monte Carlo simulation of pixel.

3.8.3 Conclusion

This section has simulated the timing circuit showing its operation as well as simulating the individual components. The different causes of error within the timing circuit have been shown and discussed. The buffers were considered as their operation introduces errors within the output not related with the method itself. Once these errors had been shown the operation of the timing circuit was compared with the theory and this showed that the errors were most prevalent close to the rails($0\,V$ or $3.3\,V$). The zero crossing method was demonstrated with an idealised Matlab simulation and this compared well with the Cadence simulations of the actual circuit.

Finally a Monte Carlo simulation was carried out to check process variation of the circuit. It was found that there was significant variation in the conversion constant with different chips. This would effectively give each pixel of a camera a different DC value for a given frequency/ amplitude. This is a form of fixed pattern noise and is not uncommon in CMOS imaging arrays and can be removed

through calibration of the value of k_{pv} for each pixel.

3.9 Layout

The final stage in designing a chip is to layout the components from the design schematic. Each component was laid out during design to make a single self-contained block. This made the final layout stage easier as it added a hierarchy, grouping components into functional blocks. These blocks were then laid out to form the final circuit. Given that this was a multi-project chip there was a limit to the space available for this design as well as the shape. There are also circuits provided by the manufacturer which are bonded to the pins of the chip, and are a required part of the design. Figure 96 shows the organisation of the chip. On the left is the layout of the entire chip with pads around the edge and circuitry in the middle. The light blue is an array of photodetectors for another project on the chip. The next is the layout of the circui for this project including all pads and buffers. The third is a layout of the pixel and the OPAMP buffers and the final is the pixel on its own. This shows how much space OPAMP buffers take up.

Figure 97 shows the layout of the pixel with the different operational parts labelled and also shows the amount of free space added to keep components isolated from each other. This space could be reduced in future designs.

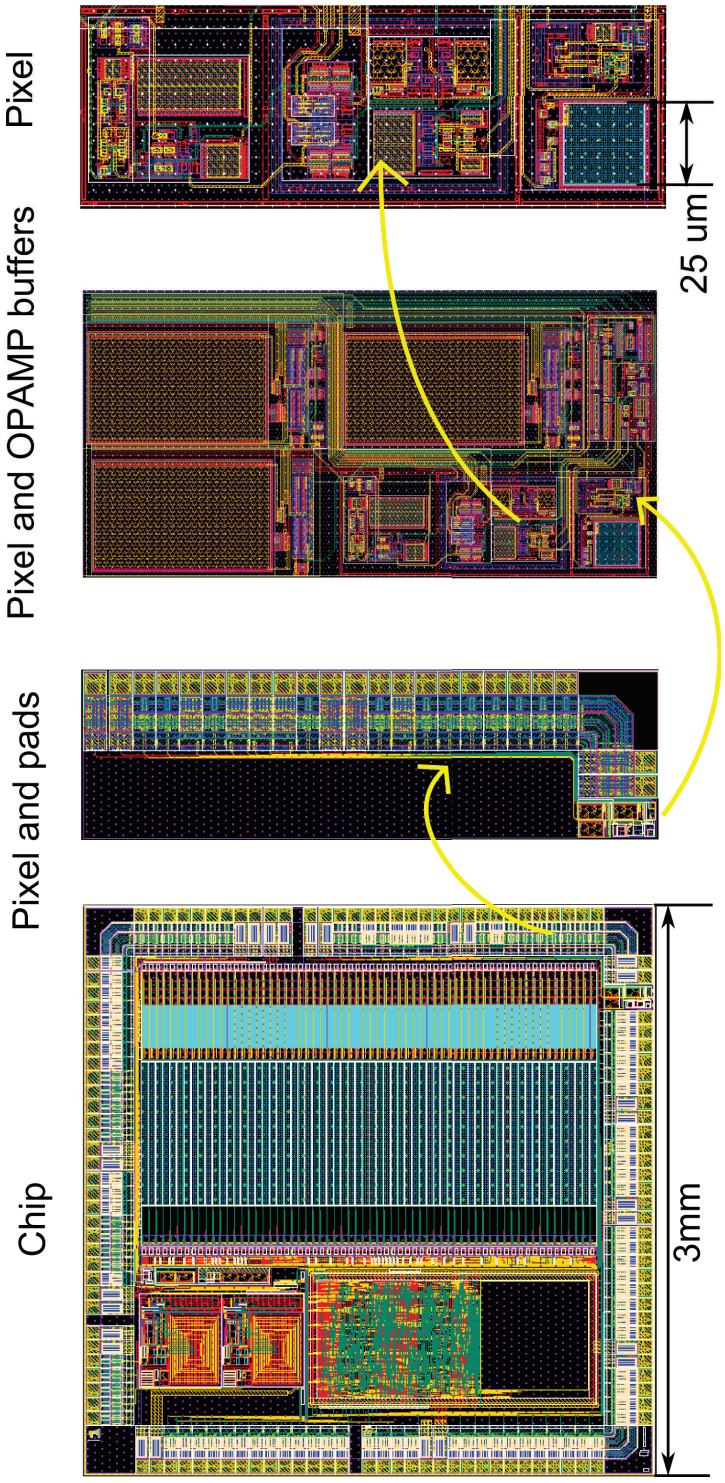

Figure 96: Layout of the chip, OPAMP buffers and pads, and pixel.

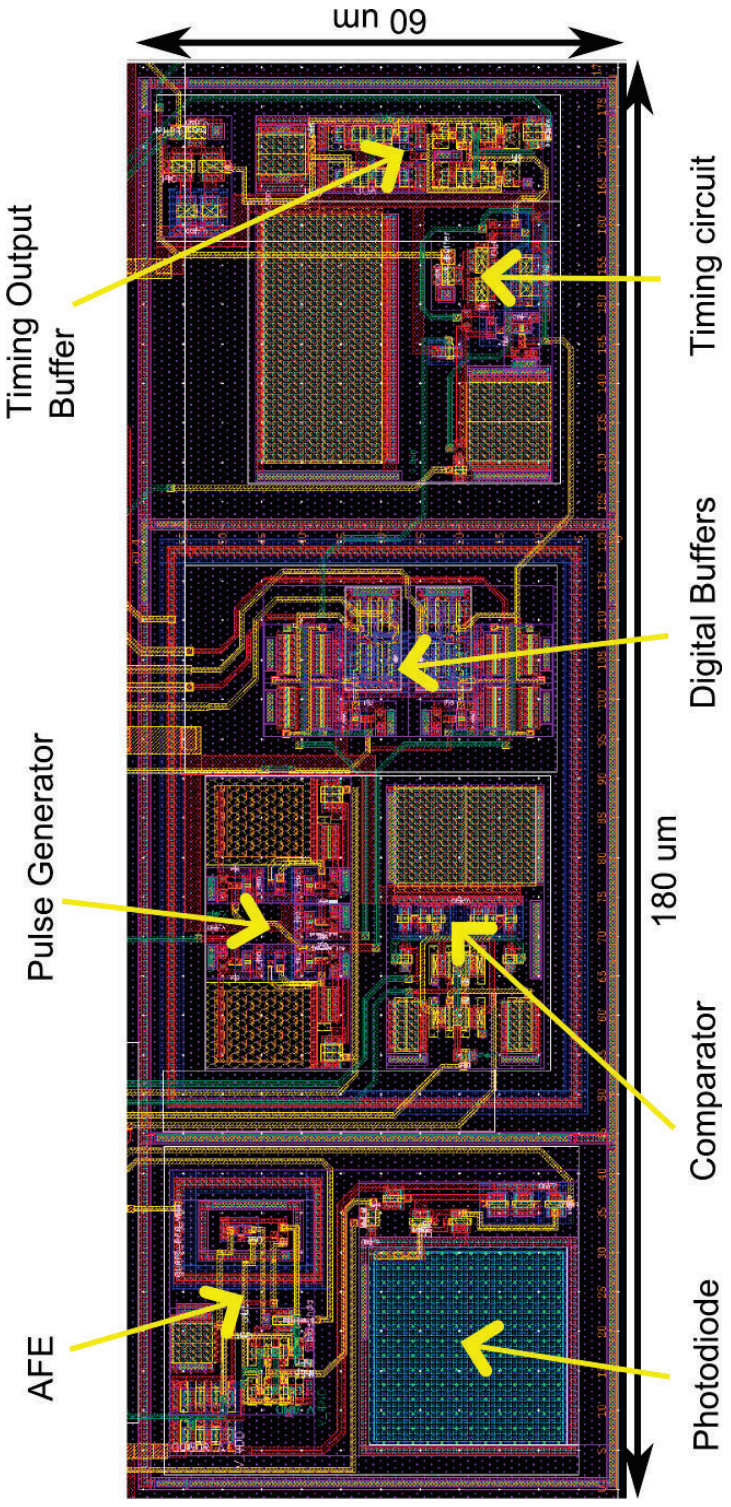

Figure 97: Layout for the final pixel, excluding the output buffers.

3.10 Conclusion

This chapter has followed the design and simulation of the proposed pixel circuit from the equation for the input light intensity, measured by a photodiode, to the output voltage. It has given a schematic of the processing circuit with the different controls and observability which were added during the design and layout of the CMOS circuit. The different circuits within the timing component have been simulated to show how they affect the operation of the pixel as a whole. This simulation has shown that the AFE is a limiting factor when considering the range of light levels and frequencies which the pixel can process. The comparator and AFE were simulated together and showed how process variation could result in some of the pixels of a large camera failing for certain input signals.

Finally the timing circuit was considered in detail showing all the sources of error and how these affected the conversion response of it. This component was shown to give a large range of outputs due to process variation. Although the AFE and comparator showed process variation, the output of the comparator gave a binary state. It either switched or not, so and process variation preceding the comparator would affect whether a pixel worked, but not affect the output it the pixel did operate. The pulse generator also showed large process variation which affected the output conversion characteristic.

The process variation could however be characterised and removed with relatively simple external processing if needed. There are also some design changes which could be done for further chips to minimise the variation. This variation would not prevent operation of a single pixel as appropriate settings of the timing discharge rate the process variation could be calibrated out. This process variation was only noticed after the chip had been fabricated however this pixel still forms a proof of principal for this method of pixel level processing. Further work would need to be carried out to alter the pixel for use in a camera and the process

variation would need to be combated at the same time. A possible solution would be to implement a common centroid design for the current source and a larger timing capacitor.

Chapters 4 and 5 will describe the tests carried out on the finished chip showing it in operation. These tests checked the operation of the different parts of the circuit with electrical signals, the operation of the entire pixel with a modulated light source, and finally the operation of the pixel in an interferometer with a vibrating target.

4 Initial Testing of Laser Doppler Vibrometry Pixel

4.1 Introduction

The final stage in the design of a custom VLSI chip is testing that all the individual parts of the chip operated as expected, as well as the chip as a whole. This testing was performed by applying various inputs signals to the input pins of the chip and checking the output signals against simulation and theory. The previous chapters of this thesis have covered the development of a proposed pixel for a full-field Laser Doppler Vibrometry (LDV) camera, as well as the theory behind that method. Chapter 3 has described the design and simulation of that pixel circuit. This chapter will now cover the testing of the chip containing that pixel.

The testing results gained in this chapter will be considered against two criteria. The first is how they compare with simulation. This validates the simulation process as well as testing that the chip does work as it was designed. The second criteria, which is important where the results vary from simulation, is whether that variation would impact on the overall operation of the chip. An example is a DC shift between simulation and experimental results in the output of the AFE. This DC shift does not affect the basic operation of the AFE and does not affect the next stage of processing (comparator) and therefore does not show a problem with the chip.

The whole design was submitted for manufacture, a process which typically takes about three months. A number of chips were manufactured and packaged with glass lids to allow light to reach the photosensitive parts of the chip. During the manufacturing a Printed Circuit Board (PCB) was designed to allow testing of the chip. This was a two layer PCB which provided power to the circuit and allowed access to the observability pins covered in Section 3.3.2.

This chapter covers the testing of the chip to ascertain how well its operation

matched the simulated results found in Chapter 3. Three different types of test were applied to this chip. The first used the observability pins to drive the different parts of the pixel circuit with voltage signals. The output from whichever part of the circuit was under test could then be monitored to check its operation. This set of tests is found in Section 4.3. The next set of tests performed on the chip were to check the operation of the whole circuit when running with a modulated light source input. This light source allowed the operation of the light sensitive part of the circuit as well the the operation of the whole demodulation pipeline to be tested under controlled conditions. These are shown in Section 4.4. The final set of tests, using an interferometer and a controlled vibrating object are covered in Chapter 5. This final tests showed the operation of the pixel in a real world setup measuring a vibrating object.

4.2 Test PCB

In order to test the chip a Printed Circuit Board (PCB) was designed, manufactured and populated. This board was needed to mount and power the chip as well as provide access to its pins. The test PCB provided jumpers and headers to allow access to the signals to and from the different pins on the chip. See Section 3.3.1 for the various pins on the chip.

A picture of the front the board is shown in Figure 98. This shows how the actual silicon chip is held behind a glass slide across the top of the chip package. Figure 99 shows the back of the test board with the various interfaces and controls. This side of the board held all interfaces to the chip so that they could be easily reached without interfering with the light falling on the chip. Figure 99 is referred to below in listing the features of the PCB.

- Four separate voltage regulators with bypass capacitors, one for each of the chip power supplies. There was a jumper on each power supply to allow

measurement of the current. In the figure they are the red jumpers shown in the orange squares.

- The SMA coaxial sockets are shown in green and were used to get outputs from the chip to the oscilloscope, or signals from the signal generator to the inputs of the chip.

- All the bias pins on the chip were controlled by nine turn variable resistors (shown in blue squares on the figure). There was also the footprint for an inductor in series with each bias line. These would allow some noise rejection by keeping the current in the bias line constant. The timing bias also had an in-line surface mount resistor to allow a wider range of bias resistances.

- The output pin was buffered twice, with an anti-aliasing filter placed between these buffers. This anti-aliasing filter could easily be removed and replaced with another to keep the testing board flexible. The initial filters were simple low-pass filters with cutoffs of $100\,kHz$. The buffered output was passed to one of the SMA connectors. The anti-aliasing filters shown in red were designed as small boards which could be unplugged and swapped to get different frequency responses.

- The digital reset and sample pins were also buffered with a tri-state buffer and had pins for input and output. The tri-state buffers could either add an extra stage of buffering when the pins were driving external loads, or set the outputs to a high-resistance setting when the sample and reset pins were being used as inputs.

- Those observability pins for the AFE and comparator could be set to either output, with similar buffering to the output pin, or input. When set to input the buffering was bypassed (using the green jumpers shown in orange

rectangles in the figure) and a second SMA connector was used to input a signal directly to the pin on the chip.

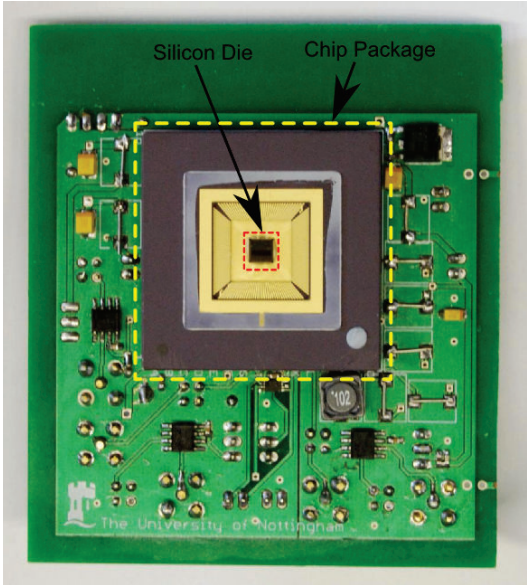

Figure 98: Test PCB front showing the silicon die within the chip.

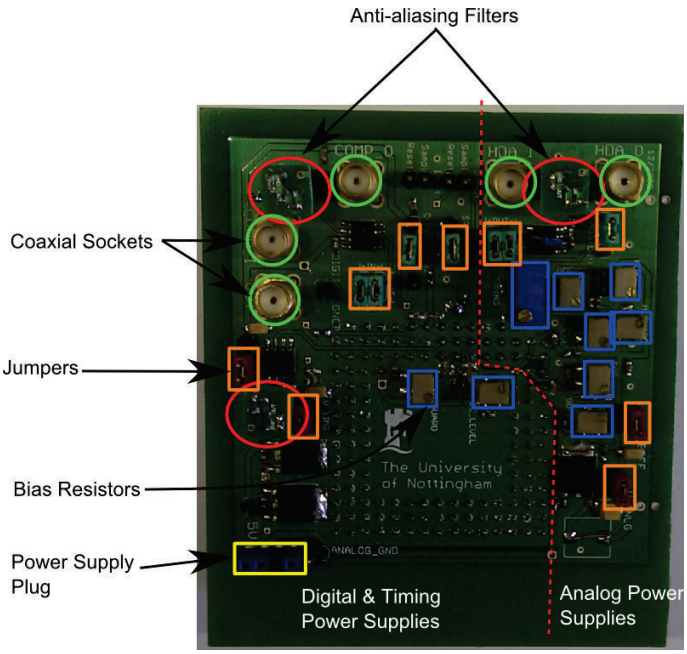

Figure 99: Test PCB board back showing the various interfaces and bias resistors.

4.3 Electrical Tests

4.3.1 Introduction

This set of tests was performed to check how the different parts of the circuit operated with electrical inputs. This allowed the most direct testing of components on the chip. This basic set of tests is also the most controlled but also the furthest removed from practical use of the chip. The sections of the chip which both their inputs and output accessible were the comparator and the pulse generator, while the timing circuit just had its output accessible. They were tested in this order and the results are shown in this section.

4.3.2 Test Setup

The test setup for the electrical tests is shown in Figure 100. In this figure the signal generator (Agilent 33220A - 20 MHz arbitrary signal generator) was used to provide the electrical input signals to the chip. This was connected to the board using coaxial cables to reduce external noise coupled onto the input. The cables also connected together the grounds of the board, signal generator and oscilloscope. The oscilloscope (Tektronix MSO 4054 - Mixed Signal Oscilloscope) was used to view and record the resulting waveforms from the board. This oscilloscope could capture up to 4 channels with up to 2.5 $Gsps$. Each channel could capture up to ten million points of data. The board was powered using a desktop power supply (Manson EP-613) set to 5 V, regulated to 3.3 V on the board.

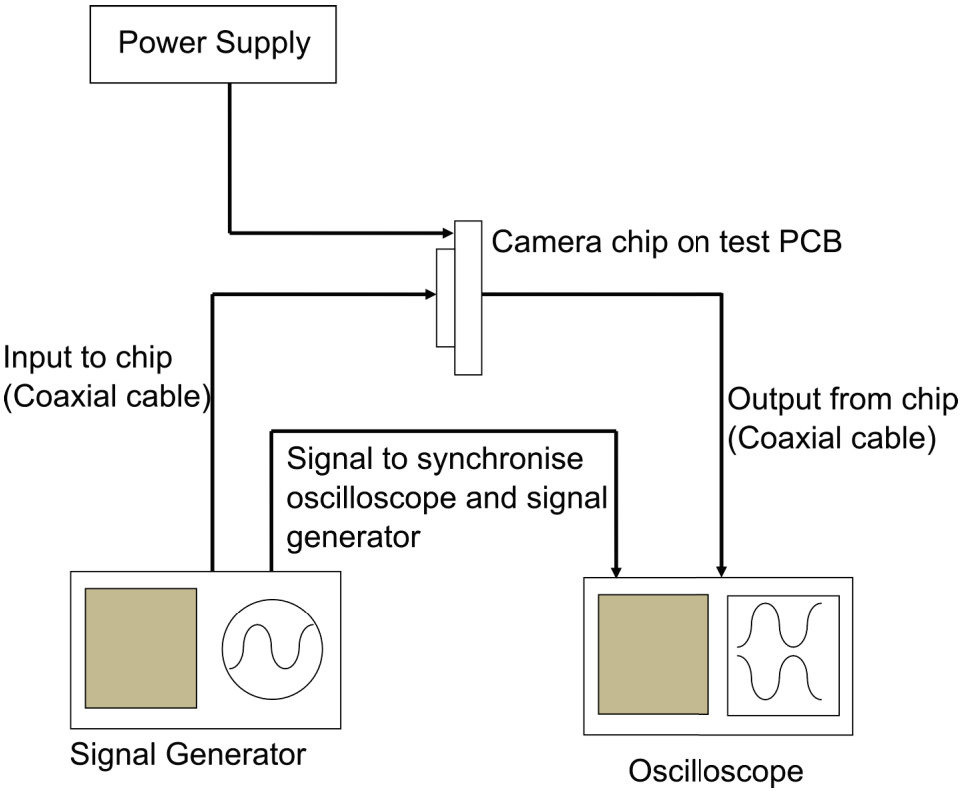

Figure 100: Experimental setup for electrical tests.

4.3.3 Comparator

The comparator was tested first as it was the interface between the analog part of the pixel circuit and the digital time measuring part of the chip. A schematic showing its input and outputs is given in Figure 101. With reference to this figure the comparator circuit was wired with its input going to the non-inverting input. This input signal would be the voltage representation of the optical irradiance falling on the chip. This voltage signal would be generated by the Analog Front End (AFE) and buffered before being passed to the comparator. The signal passed to the inverting input was a reference DC voltage provided externally to the chip, by the test PCB. The AFE input should be clipped around this reference voltage producing a clipped representation of the optical input irradiance, if the external reference voltage was set to the mean value of the AFE output.

The input to the non-inverting input is therefore the one of interest for the processing and will be called the input to the comparator from now on. The inverting input will be called the reference as it is a DC voltage provided externally to the chip.

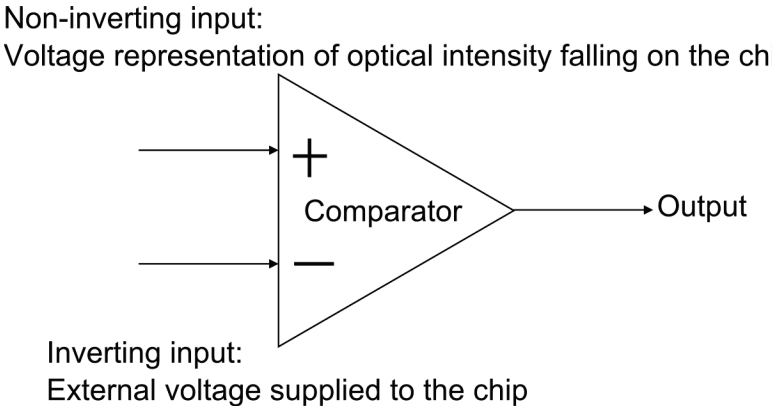

Figure 101: The inputs and output from the comparator.

Method The testing of the comparator consisted of driving the input with a range input signals while setting the reference to the mean of the input signal. The input signals were varied in amplitude (testing the hysteresis response of the comparator), DC level (testing the range of input DC voltages) and the frequency. Additionally the hysteresis of the comparator could be set by and external resistor, see Section 3.5, so the value of the hysteresis was set to the worst case, minimum value.

Basic Operation The first test was to check the basic operation of the comparator. A plot of the oscilloscope traces is given in Figure 102. This figure shows an input sine wave with a DC level of $1.65\,V$, mid way in the supply voltage range ($3.3\,V$), with a frequency of $25k\,Hz$ and the largest possible amplitude ($1.65\,V$). The maximum voltage of the clipped comparator output shown in Figure 102 ($3.368\,V$) is due to the power supply for the chip.

The minimum voltage ($0.3047\,V$) is due to the comparator output circuit. The

high and low levels of the comparator output will not affect the pulse generator's operation as this component just needs the voltage to switch from high to low, crossing $1.2\,V$. The inverter at the input of the pulse generator will clip the voltage waveform better. This plot also shows how the comparator output has a finite slew rate, however the falling edge is faster than the rising edge as shown during simulation. This asymmetry in the rise and fall times of the signal is due to the output stage of the comparator. This stage is designed to give a digital output, either high or low. To do this it has a PMOS and an NMOS transistor connected in the same manner as an inverter. The PMOS pulls the output high while the NMOS pulls it low. Both transistors have the same width to length ration ($1\,\mu m \times 1\,\mu m$) however due to the increased mobility of electrons versus holes the NMOS gives a faster falling edge than the PMOS transistor's rising edge.

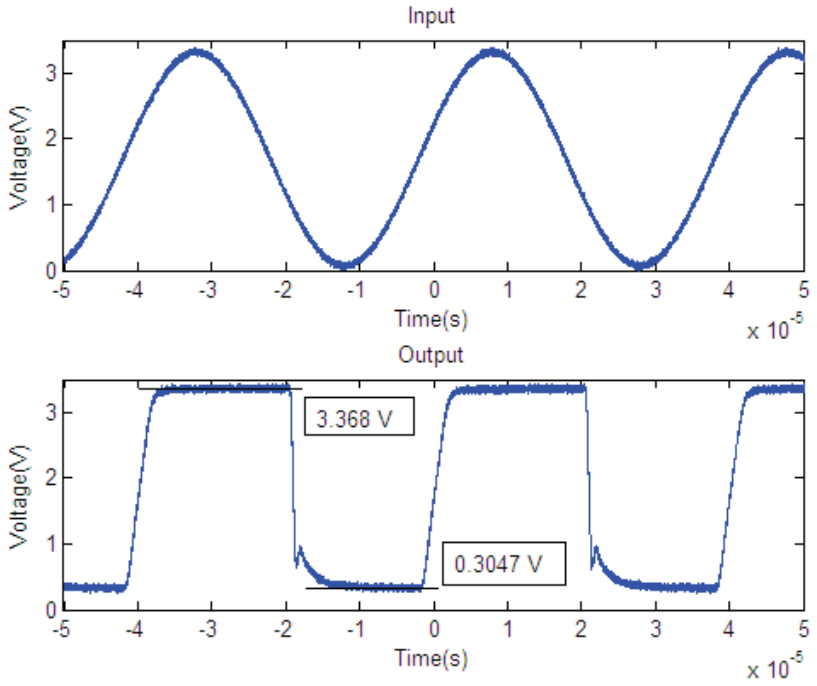

Figure 102: Comparator operation for an input of DC voltage $1.65\,V_{DC}$, AC amplitude $3.3\,V_{PP}$, and frequency $25k\,Hz$

Frequency Response The comparator showed a maximum input frequency of $121\,kHz$ when at $502\,k\Omega$ bias resistance and with the minimum amplitude from the tests on the hysteresis ($22\,mV$). With lower bias resistance the maximum frequency was above this, and the output buffer started to become the limiting factor. These values are considerably higher than the operational bandwidth of the AFE.

Hysteresis The hysteresis of the comparator is the property of the circuit that it rejects signals smaller than a certain value. The hysteresis of the comparator can be tested by finding the minimum amplitude of the input signal which still causes the output to switch. As the hysteresis was controlled by the bias resistance this was swept through the possible values ($500\,k\Omega \rightarrow 500\,\Omega$) by adjusting the variable bias resistor, and at each bias resistance the amplitude of the input signal was changed to give the smallest signal which still caused the comparator to switch. The output signal from the comparator was checked on the oscilloscope.

For this test the DC value of the input signal as kept at $1.65\,V$ and the frequency of the input sine wave was set to $25\,kHz$. The results are shown in Figure 103. This figure shows how the hysteresis peaks at $16.8\,k\Omega$ and then falls for values either side of the maximum value. The simulated values for the hysteresis from Chapter 3 are also shown for comparison. The curves have the maximum values at the same bias, and the same shape. The simulated results are above the experimental results, probably due to process variation in the chip. Unfortunately, it was not possible to run a swept variable simulation for the process variation with the version of Cadence used to design the chip. The swept variable simulation was the type of simulation needed to find the hysteresis point and could not be run as Monte Carlo simulation.

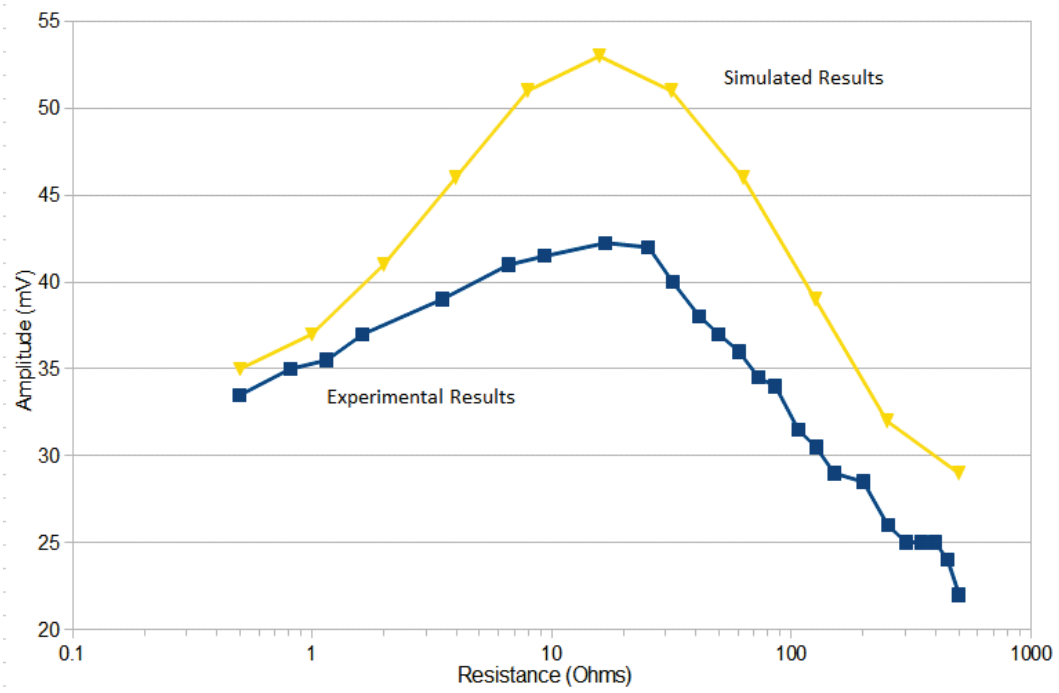

Figure 103: Comparator hysteresis results. Amplitude of the minimum signal which caused the comparator to switch against the bias resistance.

Input DC Range The next set of tests performed was to set the bias resistance to the largest hysteresis value ($16.8\,k\Omega$) and then sweeping the DC value of the input signal from $0\,V$ to $3.3\,V$ and checking the amplitude of the output signal from the comparator. This is the same as the simulation in 3.5 and shows how the comparator is sensitive to the DC level of the input signal. The results are shown in Figure 104. The experimental and simulated match for the operational range of the comparator. The particular fabricated device showed a slightly better high-voltage acceptance but a slightly worse low-voltage acceptance below $0.3\,V$. These variations can again be put down to process variation during fabrication of the chip.

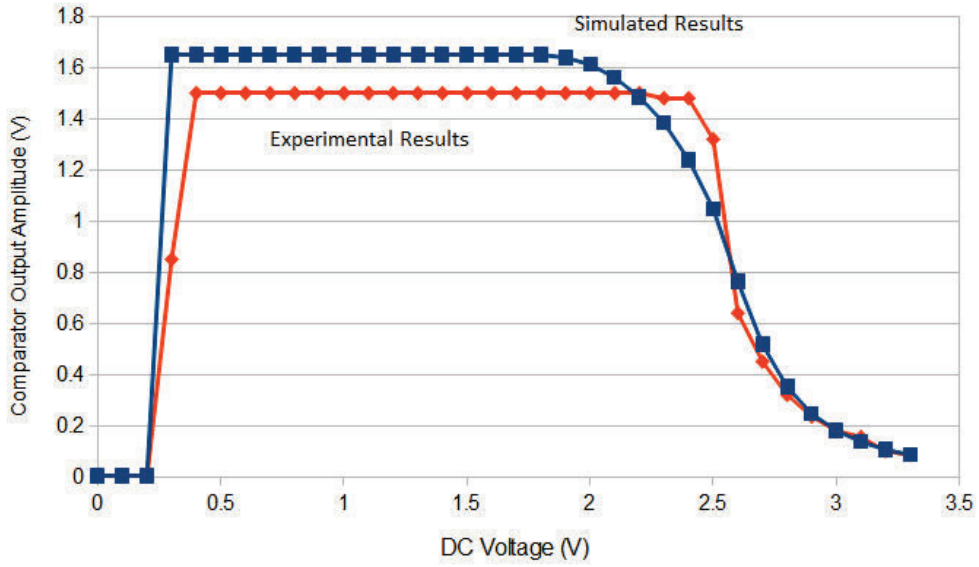

Figure 104: Comparison of simulated and experimental results for the DC sensitivity of the comparator.

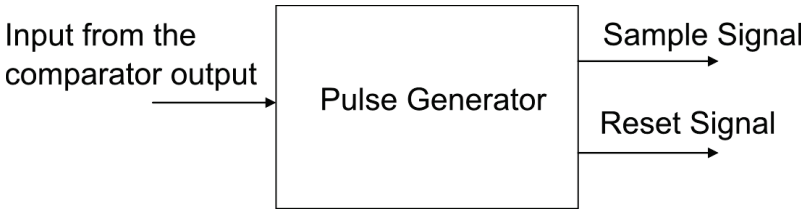

Figure 105: Pulse generator input and outputs.

4.3.4 Pulse Generator

The pulse generator was designed to provide two non-overlapping pulses of controllable width on the falling edge of its input, see Section 3.7 in Chapter 3. A block diagram of it is shown in Figure 105. On this figure the input to the pulse generator would be provided by the output from the comparator. The pulse generator was one of the components designed specifically for this chip.

Method The pulse generator was tested by providing a square wave input from the signal generator and checking the pulses generated on the oscilloscope. Some care was needed with the probes for the oscilloscope due to the fast nature of

the pulses. For this a double coaxial cable was built to carry the signals to the oscilloscope. The pulse widths could be varied by changing the bias resistance of the pulse generator. The bias resistance would be varied and the corresponding pulse width compared with theory.

Results The pulse generator worked as expected when set to short pulse widths. Figure 106 shows a plot of the oscilloscope traces of the pulse generator in operation. In this figure the top plot shows the input to the pulse generator, the output from the comparator. The lower plot shows the sample and reset signals. The smoothing of these signals is probably due to the output buffers used on these lines, rather than internal to the chip, and therefore will not affect the timing circuit. The internal buffers on these lines are much smaller and drive much smaller capacitance than the output buffers. There is some coupling between the sample and reset signals, especially evident from $150\,ns$ onwards. There is also some other internal coupling, possibly to the comparator output as shown by 'A' on Figure 106. Here the reset and sample signals move in opposition indicating some other coupled signal. The work correctly up to $1 M\,Hz$ and above, however at higher frequencies the pulses start to overlap.

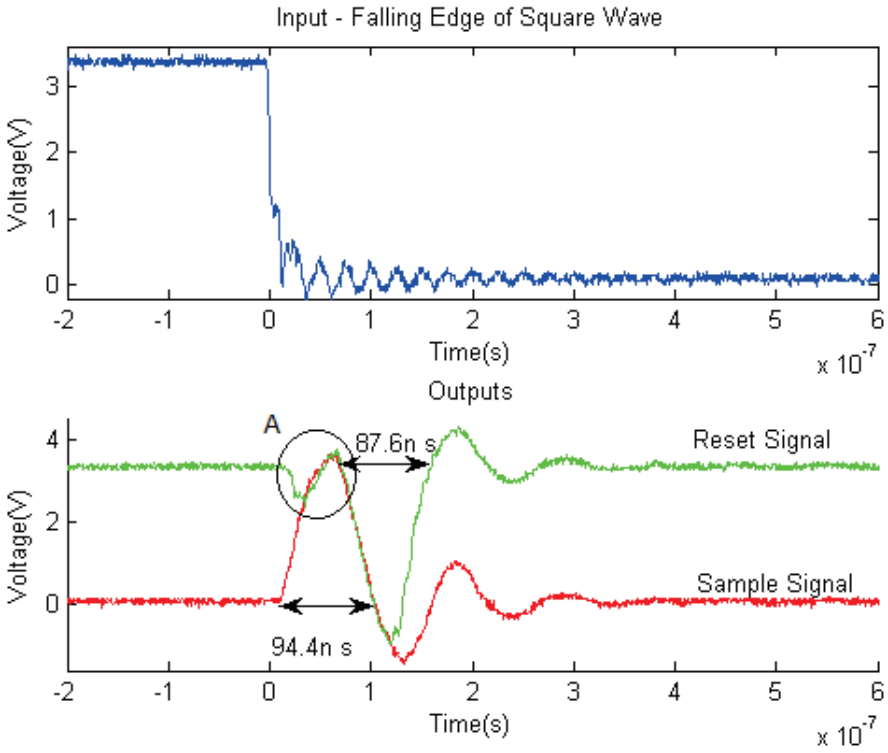

Figure 106: Pulse generator operation for bias resistance of $12.33k\,\Omega$. Input: $1.65\,V_{DC}$, $3.3\,V_{PP}$, $1k\,Hz$

The pulse generator started to fail at longer pulse widths. When the pulse width was increased the reset pulse stopped completing its pulse cycle correctly. This is shown in Figures 107 and 108 for different pulse widths. Figure 107 shows a slightly longer pulse width than Figure 106, while Figure 108 is a considerably longer pulse width.

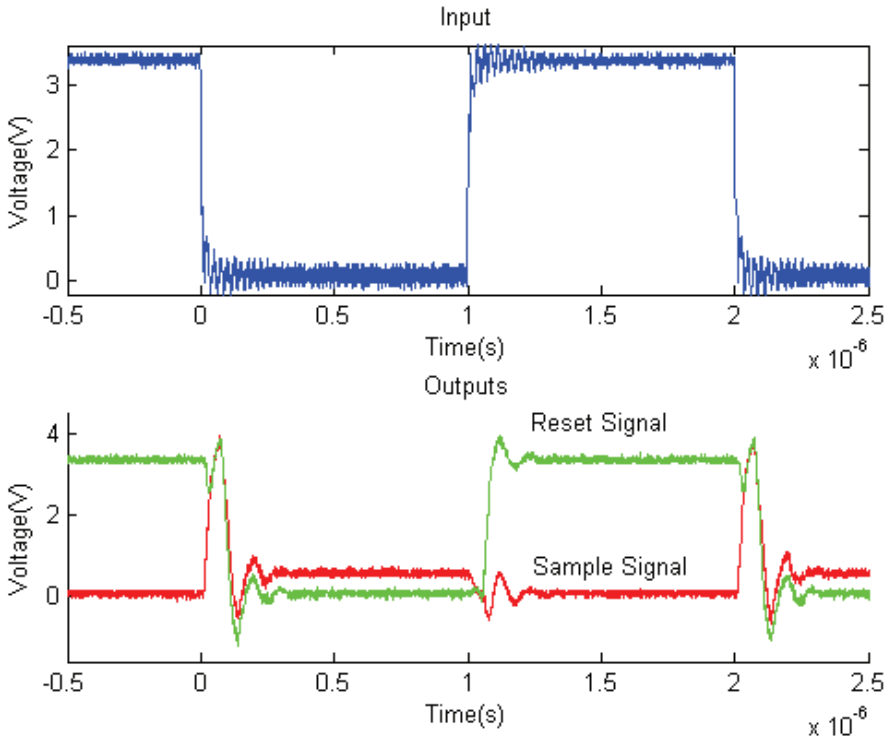

Figure 107: Pulse generator failing for bias resistance of $15.56k\,\Omega$. Input: $1.65\,V_{DC}$, $3.3\,V_{PP}$, $500k\,Hz$

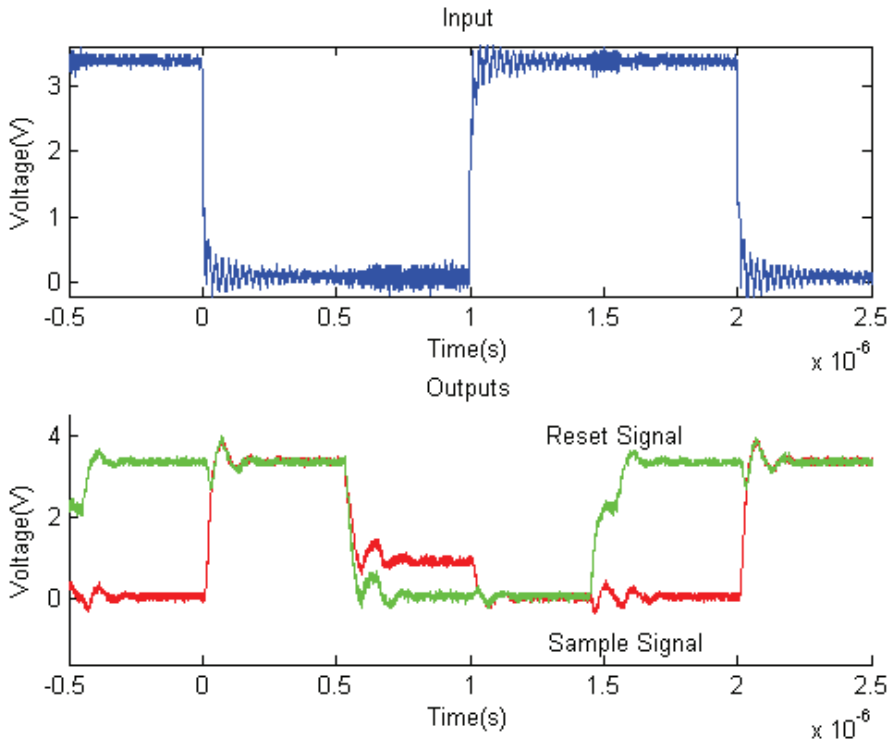

Figure 108: Pulse generator failing for bias resistance of $103.9k\,\Omega$. Input: $1.65\,V_{DC}$, $3.3\,V_{PP}$, $500k\,Hz$

The reason for this appears to be due to an effect called latch-up. This is the undesirable triggering of parasitic structures in a silicon chip. CMOS chips are made of two types of silicon: n-type and p-type. These are pure silicon (called intrinsic silicon) which has had specific impurity atoms implanted in it. These impurities change the characteristics of the silicon and allow it to conduct current. n-type has impurities which bond with the silicon atomic lattice but have an extra electron, which can now move around the lattice. p-type silicon has impurities which have fewer electrons that the silicon outer shell. As such they can accept an electron from a neighboring atom, and so allow current to flow. Conceptually this is often thought of as a positive lack-of-an-electron moving in the opposite direction (called a hole) [118].

When laying-out a CMOS circuit the transistors are composed of n-type and

p-type silicon, placed as required. However in reality these are just areas of p-type and n-type silicon. If an area has several structures near one another then unwanted, or parasitic, transistors can be formed by the joining of two intended transistors. This can lead to the unwanted effect called latch-up. Figure 109 shows an inverter composed of two CMOS transistors, and how these can form a pair of parasitic bipolar junction transistors (BJTs). In this figure (a) is the schematic of an inverter. This inverter can be laid out as shown in (b). n+ and p+ show heavily doped n-type and p-type silicon respectively. The inner pair of n+ and p+ are called substrate connections. The dotted line shows the path through the different n and p regions which forms the parasitic transistor construct, as shown in (c) and (d).

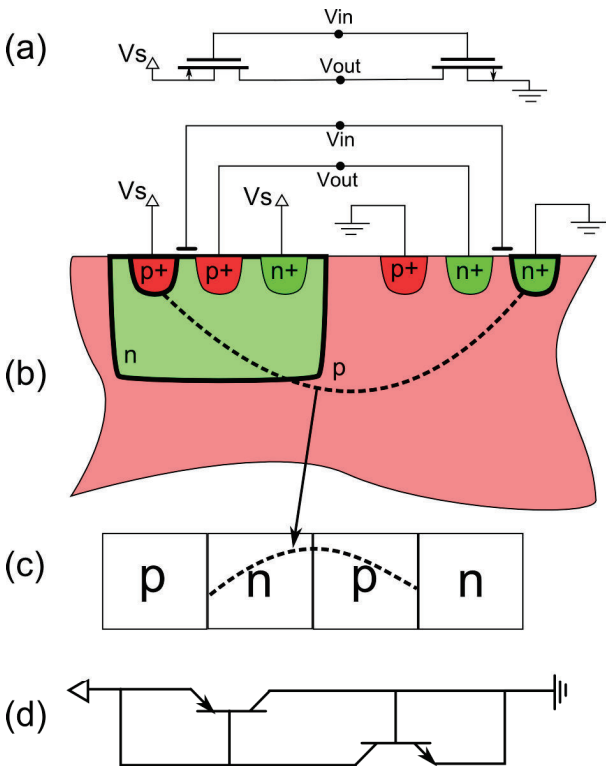

Figure 109: An inverter schematic (a), the silicon implementation of this device (b), and how four of the regions of the inverter layout can form a parasitic BJT pair (c and d).

Latch-up is caused by carrier injection into the silicon [123]. If these charges

(either electrons or holes) reach other structures then may cause those structures to go into latch-up. Figure 109 shows the parasitic transistors which are switched on during latch-up, and may be self-sustaining. Once in latch-up a connection is made between the supply and ground.

There are two types of latch-up: internal, or external [124]. For internal latch-up the parasitic transistors are triggered by perturbations of the inputs of the intended transistors of which they are made. External latch-up is where the perturbations are applied elsewhere in the circuit and the movement of carriers through the silicon causes the parasitic transistors to turn on.

Latch-up can be triggered by electro-static discharge (ESD) events or rapid fluctuations on one of the pins of a transistor, such as an n-well being driven below ground [125]. Inverters are susceptible to latch-up and this appears to be the fault with the pulse generator. Once latch-up has occurred the output of the inverter will not change with the input, and may be stuck in a logic state [126].

Given that latch-up connects the supply to ground, there is an increased current draw by the chip. The test PCB had jumpers to allow measurement of the current being supplied to the chip. A $2.2\,\Omega$ resistor was placed in series between the regulator and the chip's digital supply pin. The voltage across this resistor was measured and converted to a current and is shown in Figure 110 for both pulse working and pulse not working. The current drawn when the pulse was not working is approximately six times higher and exhibits high frequency oscillations each time the reset is incorrectly being held low as shown in Figure 111.

Although the circuit simulations used were not designed to detect latch-up, they do show where high currents or large transient voltages occur in the chip. Simulations were run on the pulse circuit to check for high transient currents and voltages. Figure 112 is included to show the components of the pulse generator. The results are shown in Figure 113. The top trace shows the input waveform from

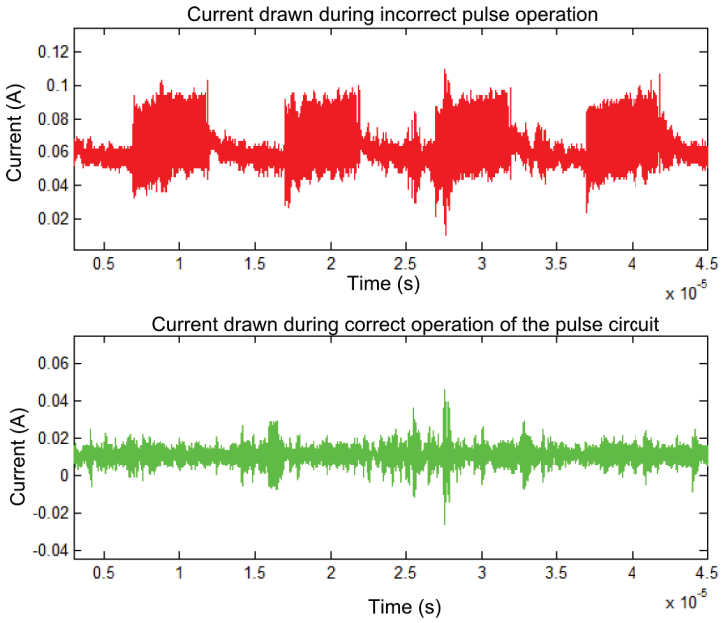

Figure 110: Change in current draw when the pulse generator stops working. Top trace shows the current drawn by the digital side of the circuit when the pulse is not working. The lower trace shows the same current when the pulse is working.

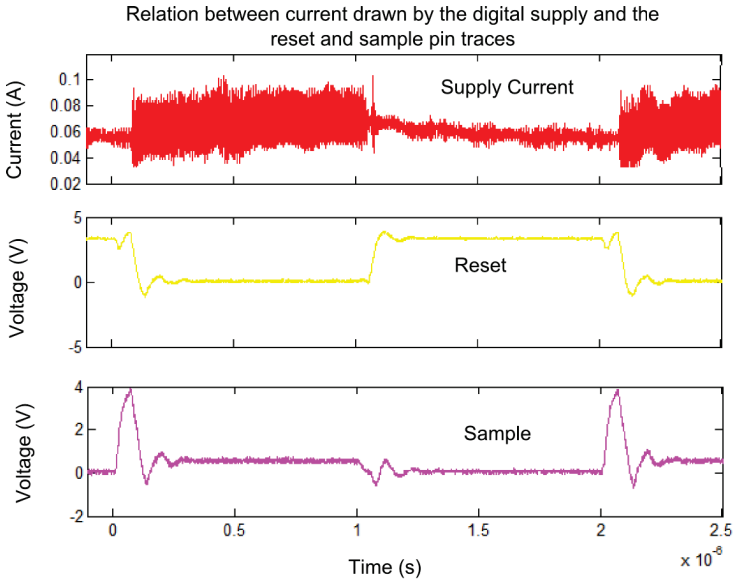

Figure 111: Comparison between the supply current to the digital side of the circuit with the reset and sample signals.

the comparator. This is the falling edge and the one most of the circuit responds to. The next pair of plots show the current traces for the NMOS transistor which forms part of INV0. The output of this inverter drives the poly-poly capacitor and so is exposed to unrestrained voltages. The current flows through the INV0 NMOS transistor show that the reset side of the circuit has very high transient currents, especially into the bulk pin of the transistor, compared with the sample's one. These high currents might inject minority carriers into the silicon, although their short duration will limit their effect. Another possibility is the voltage on the input to INV1. Traces of this voltage are shown in the third plots down in Figure 113. This voltage shows a transient, negative spike on the reset circuits INV input. This is not present on the sample side during the falling edge of the input. It is present during the rising edge (fourth set of plots of Figure 113) but it is a less rapid spike, so if this is the problem it might not affect the sample pulse.

It is impossible to tell what the actual fault is, but these show that latch-up is a possible culprit. Future designs should include guard rings and try to space the sample and reset circuit out as much as possible as this also helps reduce external latch-up.

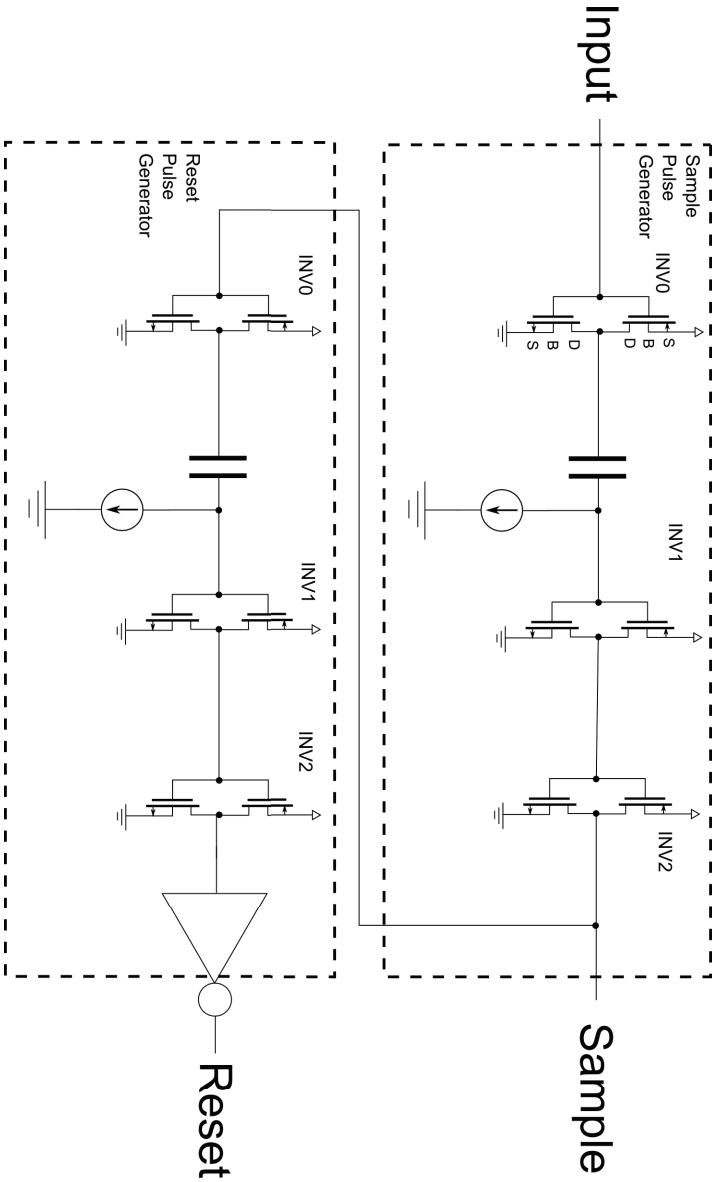

Figure 112: Schematic of pulse generator showing the separate circuits for sample and reset signals. These circuits are composed of inverters with poly-poly capacitors.

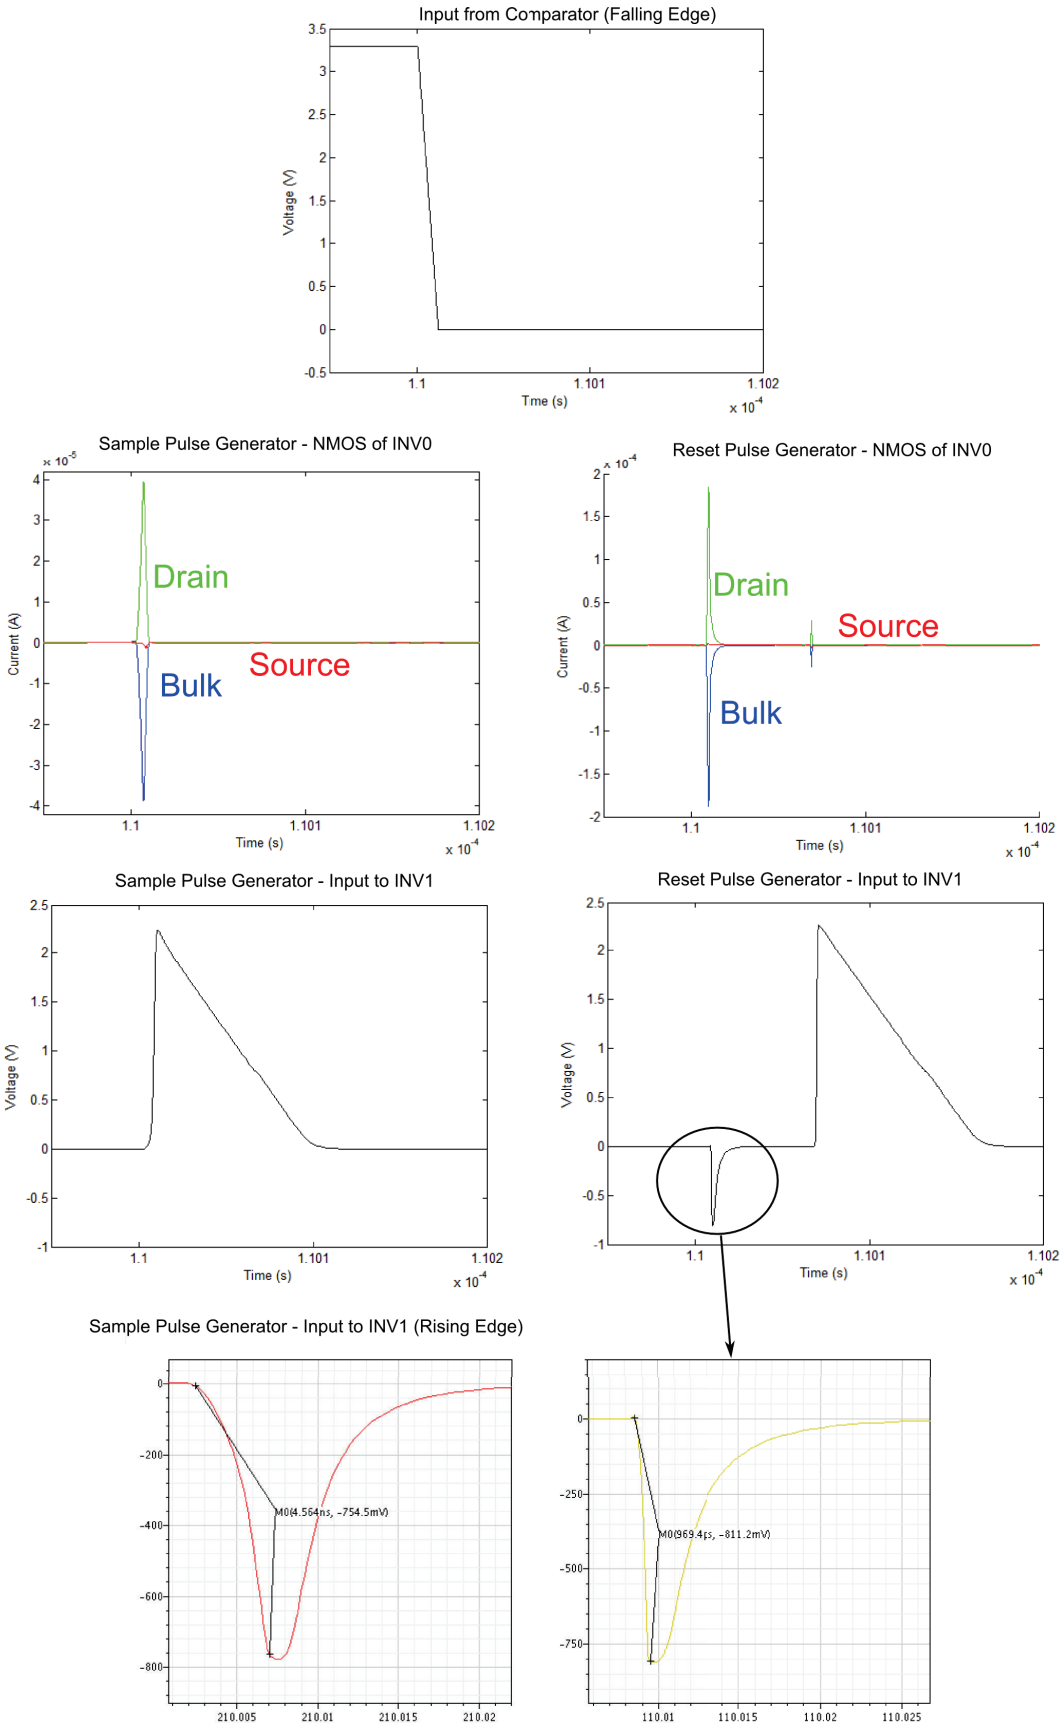

Figure 113: Results from detailed simulation of the pulse generator circuit.

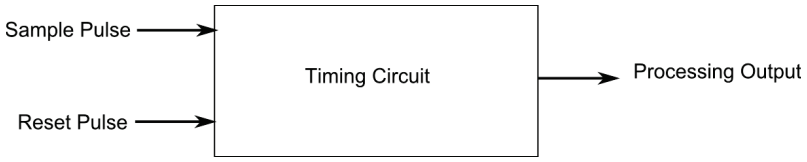

Figure 114: Timing circuit block diagram.

Latch-up is often combated through the use of guard-rings. There are regions of p- or n-type silicon added around the intended transistors to absorb any injected charges. p-wells are connected to ground and n-wells are connected to the supply [123]. They can also be left floating, although this is less effective.

As the pulse generator worked for shorter pulses this did not prevent testing of the timing circuit, and the chip as a whole.

4.3.5 Timing

The timing circuit was the other circuit designed specifically for this chip. It took digital inputs and produce an analog output. A block diagram is shown in Figure 114. The output was buffered within the circuit before being buffered by an OPAMP on the test PCB. The inputs were two digital control signals, and an analog bias voltage. The control signals came from the two pulse generators run from the comparator output. The bias voltage was controlled by an external resistor, and controlled the timing circuits operation.

Method The timing component was tested in two stages. First the signal generator was set to produce the two control pulses. This allowed testing of the circuit in isolation and given the trouble with the pulse generator it allowed different pulse widths to be tested with the circuit. The second set of tests checked how the timing circuit responded to a frequency modulated input. The signal generator could not produce frequency modulated control pulses so for this stage of testing the on-chip pulse generator had to be used. It was set to give the longest

pulse width while still operating correctly.

Results The results from the first tests are plotted in Figure 115 along with a Matlab simulation using the same k_{pv} as the timing circuit. The timing output for each input frequency was averaged to remove noise and this gives the experimental results. This shows the timing component working as predicted. The Matlab simulation is an ideal period to voltage conversion and therefore varies from $0\,V$ to $2.7\,V$. The timing outputs follow the simulated results over the range $15k\,Hz$ to $55k\,Hz$. The experimental results for $5k\,Hz$ and $10k\,Hz$ show the effects of the output voltage range of the buffers used in the timing circuit. This would constrain the voltage range the timing circuit could work in.

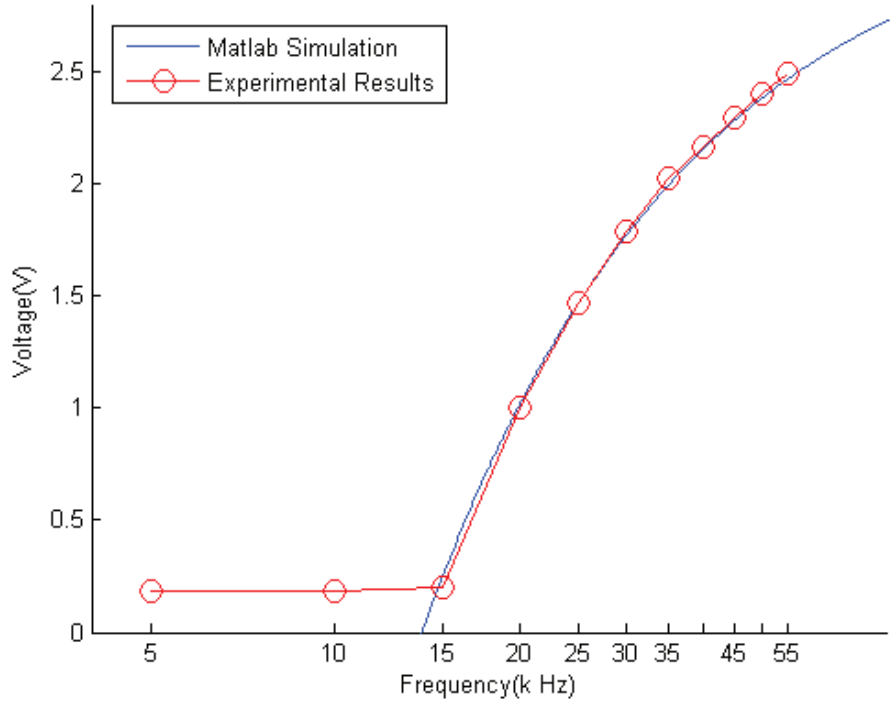

Figure 115: Timing output voltage vs. input frequency for electrical input with $k_{pv} = 45.75k\,V/s$

The next test was to check how the timing component responded to a modulated input. This was done by using a frequency modulated output from the

signal generator driving the input to the pulse generator. The results are given in Figure 116. This can be compared with Figure 91 in Section 3.8.2 of Chapter 3.

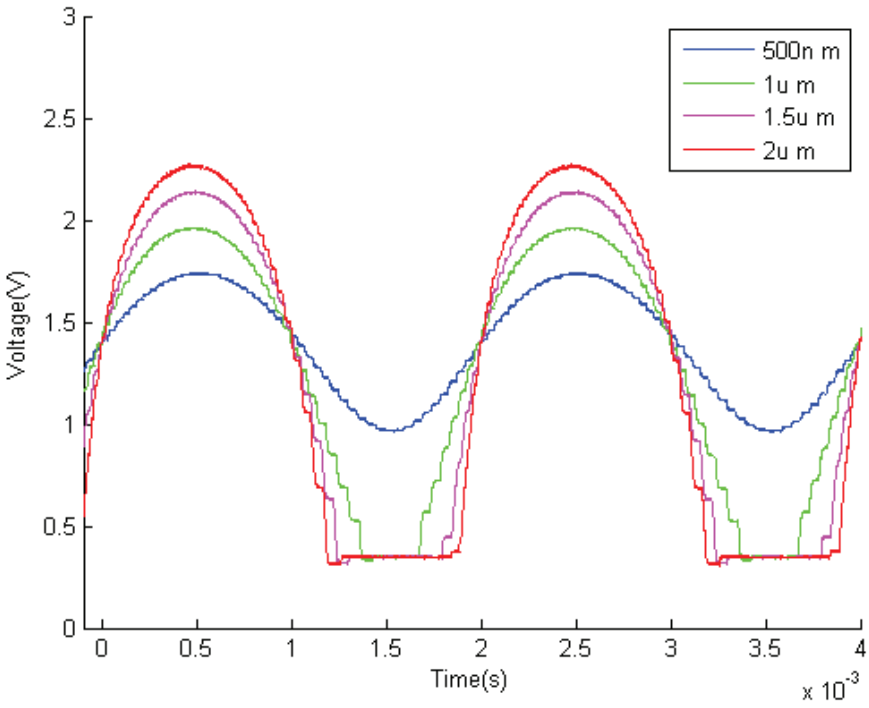

Figure 116: Timing output for different modulated inputs. Using the pulse generator to generate the control pulses. $25k\,Hz$ carrier frequency, $k_{pv} = 45.75k\,v/s$, $500\,Hz$ vibration frequency.

Using Equation 55 in the same section the peaks can be scaled into amplitude. The results are $455n\,m$, $991n\,m$, $1.53u\,m$ and $2.06u\,m$. These are all within 9% of the correct values. Figure 117 shows a comparison between the Matlab and Cadence simulations and the experimental timing output for $500n\,m$, which shows good agreement.

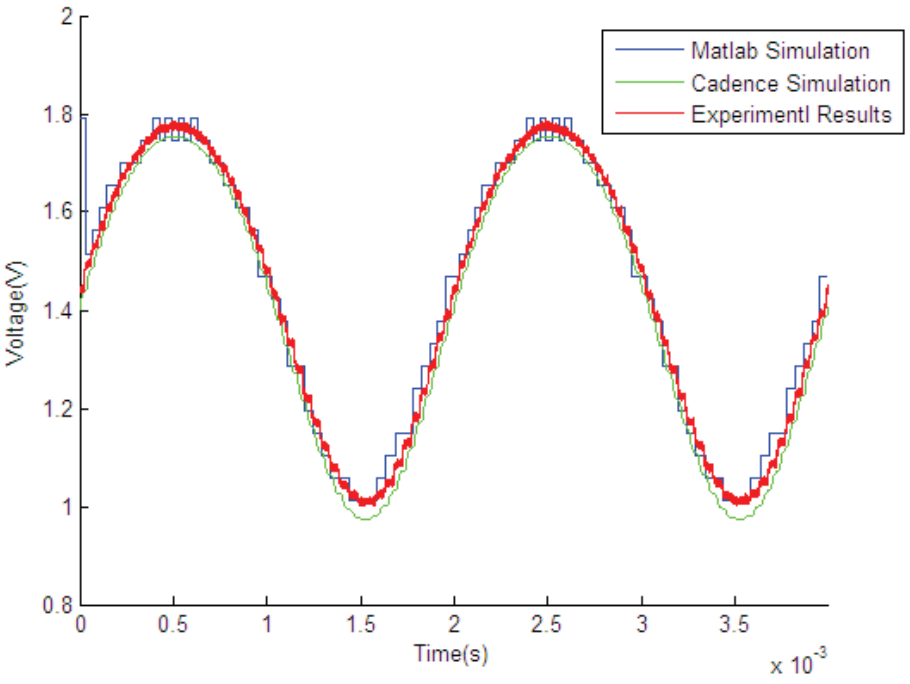

Figure 117: Comparison of Matlab and Cadence simulation with experimental results using the input to the comparator to bypass the AFE

4.3.6 Conclusion

The comparator was shown to operate as expected. Although the results for the hysteresis section showed a change in the magnitude of the hysteresis level when varying the bias resistance, the shape of the graph still agreed well with simulation. The pulse generator showed a fault with pulse widths longer than $\sim 100\,ns$ and this was attributed to latch-up. The pulse generator worked for shorter pulse widths and allowed testing of the timing circuit.

The different parts of the pixel circuit were tested and showed good agreement with theory, with the exception of the pulse generator. The testing of the timing circuit used the comparator, pulse generators and timing circuit and therefore shows the majority of the parts of the pixel operating as expected with a modulated input. The final part of the circuit is the AFE which could only be tested using an optical input.

4.4 Modulated Light Tests

4.4.1 Introduction

The second set of tests performed on the chip were to test the light sensitive circuitry and used modulated light source. This formed a middle stage between electrical tests and full interferometer testing while also checking the operation of the Analog Front End (AFE) which had not been tested yet. These tests all work with changes (modulation) in light power and do not require coherent light sources, although a laser was used, simply as a high power source source of light. This makes the experimental setup easier as the precision requirements of optical mixing are not relevant.

These tests started by checking the AFE's operation before moving onto checking the chip operation as a whole in response to a light input signal. The experimental setup will be described first as well as tests which were done on the surface properties of different objects. These tests give an idea of the reflected light level which can be expected and are important to make sure that the light levels used in these tests are actually physically possible and practical.

4.4.2 Test Setup

The second tests used a similar setup to the electrical tests in Section 4.3 just with the addition of a modulated light source to drive the light sensitive AFE. The setup is shown in Figure 118. The light source was an Light Emitting Diode (LED) array with a lens tube to keep the distance between the array and the chip constant as well as block out ambient light. The light sources were calibrated using a Thorlabs PDA36A-EC reference photodiode and amplifier, sampled by the oscilloscope. This was a $3.6\,mm$ square photodiode attached to a transimpedance amplifier, similar to that of the AFE, but with a linear response to light irradiance. The

device was built into a small, optics bench-mountable box with gain adjustable from $0\,dB$ to $70\,dB$. The output voltage could then be converted to an incident light power.

Coaxial cables were used to connect the SMA connectors on the test board to the SMB connectors on the oscilloscope to minimise noise coupling. The oscilloscope had four channels which allowed several signals to be captured at once. The signal generator had two channels which allowed the sample and reset signals to be generated for the timing component testing.

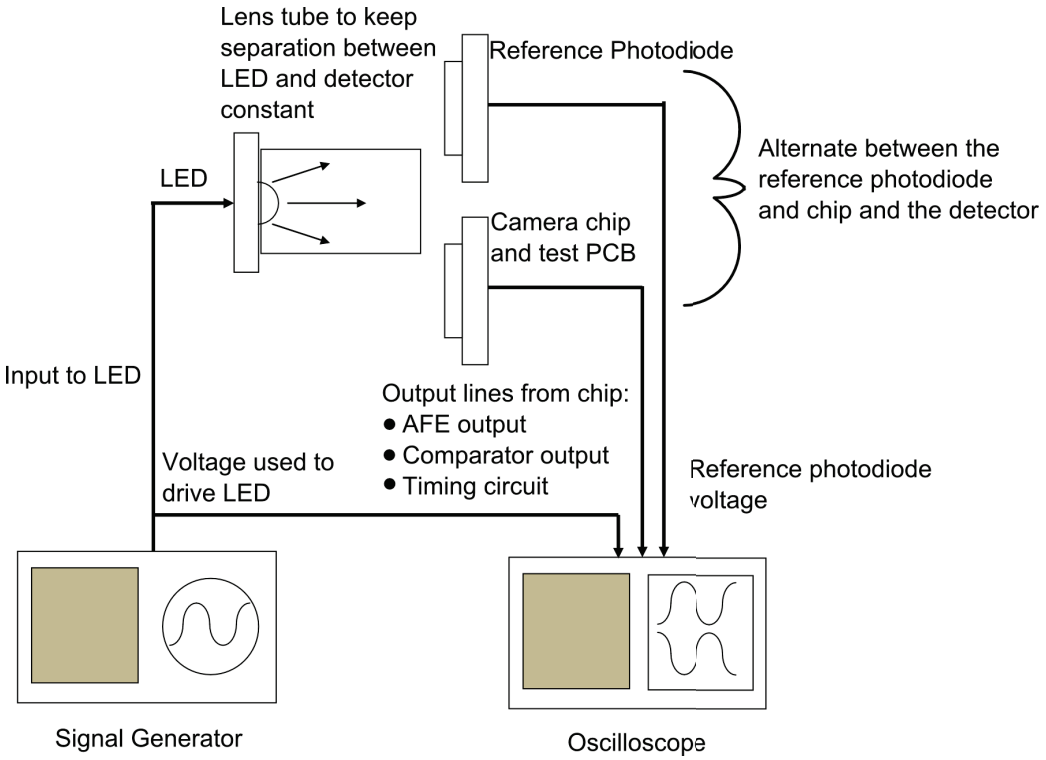

Figure 118: Experimental setup for the optical part of the testing.

4.4.3 Light Sources

Two light sources were used for this stage of testing. The first has already been mentioned, the LED array. The other was a laser, such as would be used for the interferometer, as a bright light source. Unfortunately the laser could not

be modulated so was only used for the DC tests of the AFE. It could, however, provide a much brighter light source than the LEDs.

The LED configuration was a grid of 4 LEDs packed as close as possible with the holes on 0.1" Vero Board. Neutral Density (ND) filters were used to decrease the light power in a controlled, linear way. For the laser this was the only way to decrease the light power. For the LEDs the voltage could be decreased, but the light from an LED is dependent on the current through the device, and the current has a non-linear relationship to the voltage across the diode. This means that the light power at $4\,V$ will not necessarily by twice the light power of $2\,V$. ND filters will always decrease the light power by a fixed amount.

Both light sources were calibrated at maximum power (without filtering) and then ND filters were used to decrease this power. The added advantage of this approach is that the modulation depth of the light signal from the LED array when it is being modulated will not be affected by decreasing the overall light power. As the modulation depth is more relevant to the final interferometer setup than the amplitude, this make comparison easier.

The reference photodiode was used to measure the light power at a certain distance from the source. LEDs provide a good source of light which can be easily modulated by a signal generator. Unfortunately the light distribution from a single LED is not uniform hence a block of LEDs was used to provide a more uniform distribution. The distance between the LEDs and the detector also had to be kept constant so that the light intensity experienced by the reference photodiode and the chip was the same. Increasing the distance flattened the intensity distribution at the cost of reduced intensity. A CCD was used to check the flatness of the LED intensity before the experiment and the distance was balanced with the light intensity such that the experiment could use light levels which might be expected from a real life setup.

The beam profile of the LED light source was checked with a CCD camera. This was to confirm that the light intensity was approximately flat over the central illuminated region. This check was needed as the reference photodiode used to check the light intensity was $3.6\,mm$ square while the processing circuit light sensitive area was $25\,um$ square. This separation was set to $30\,mm$ and the CCD image showed that intensity was flat within 10% and the difference between the average light intensity captured by the processing detector and the reference photodiode was 1.1%.

This LED array was capable of producing a maximum light level which would result in a log pixel photocurrent of $200\,pA$. The laser was capable of producing a light level which would correspond to a photocurrent of $10\,\mu A$.

4.4.4 Light Level

The first stage in testing the chip with light was to find the expected light levels. This was particularly important as the DC level of the light would partly set the bandwidth of the pixel's front end. The DC light level would also set where the signal fell on the log curve of the AFE's light response which would control the gain of the AFE. Finally the output of the AFE would need to fall in the input range of the comparator.

To check the operation of the AFE under as close to real conditions as possible a partial interferometer was setup to gain some idea of the light levels that could be expected. One arm of the interferometer was constructed to allow the signal beam intensity to be estimated. This is shown in Figure 119. The important parts of this setup are that a beam splitter is incorporated to check how efficient these would be as well as the lens system. The lenses would control how much light was collected from the object. A number of samples were used to check different surface properties. When an object reflects light its surface properties effect how

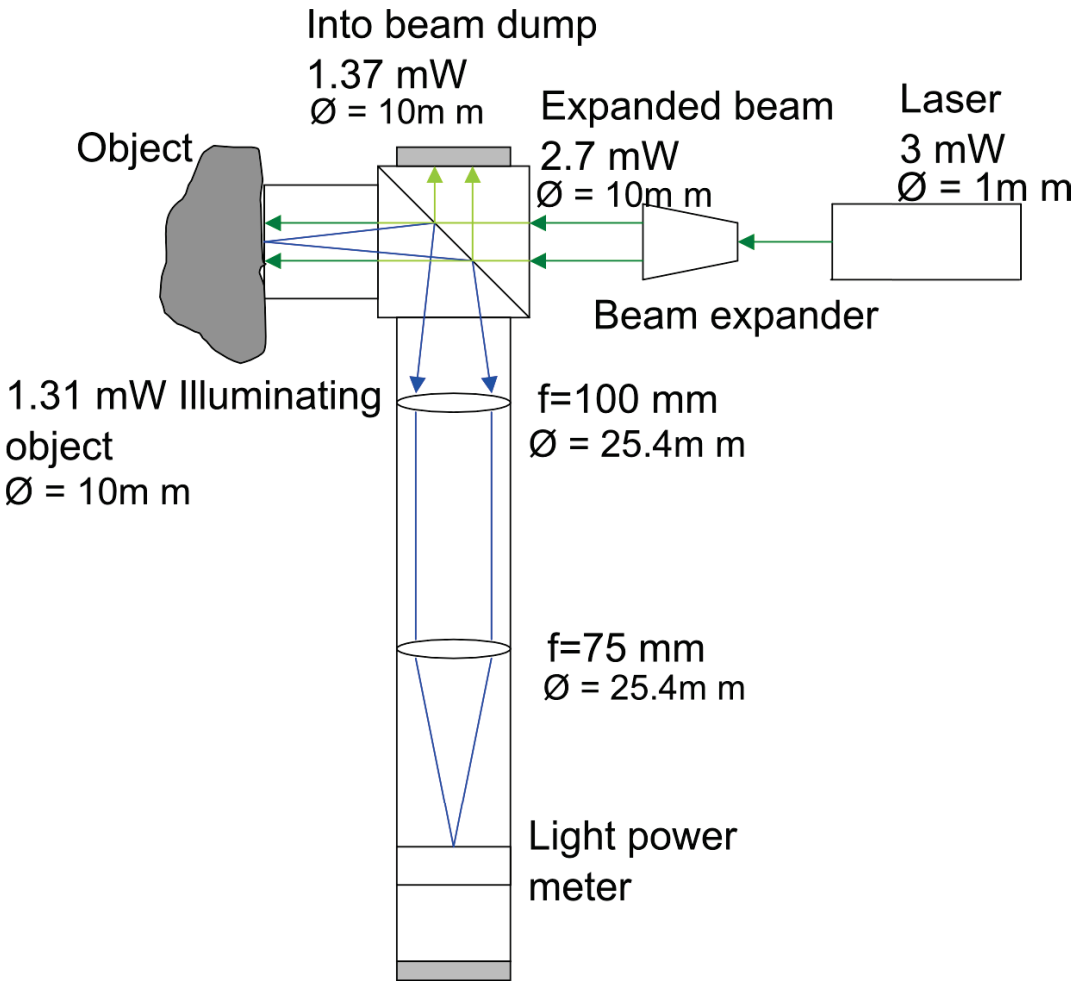

Figure 119: Signal arm of the interferometer

much light is reflected. A mirror will reflect almost all the light with a well defined angle of reflection. As the surface becomes more diffuse the light reflected starts to form a cone rather than a single reflection direction. A larger lens will collect more light from this cone and give better results. However when using the cage system the size of the lens was fixed. The results from the experiment are given in Table 10.

Using this information it was possible to calculate a theoretical light level for the finished interferometer. In an ideal setup the signal and reference beams would have the same intensity. Assuming some pre-treating(cleaning/ painting) the light arriving at the sensor from the signal beam might be between $185m\,W/m^2$ and

Table 10: Signal light levels from different materials

Information	Light Power(W)($\phi = 10\,mm$)	Light Irradiance(W/m^2)
Laser Source	$3m(\phi = 1\,mm)$	$3.8\,k$
Expanded beam	$2.7\,m$	34
Reflected beam	$1.37\,m$	17.4
Transmitted beam	$1.31\,m$	16.7
Matt black reflector	$2.26\,u$	$28.8\,m$
White paper	$8.02\,u$	$102\,m$
Glossy painted steel sheet	$14.5\,u$	$185\,m$
Pitted plated steel bar	$29.7\,u$	$378\,m$
Flat rough aluminium bar	$53.4\,u$	$680\,m$
Mirror	$558\,u$	7.1

$680m\,W/m^2$. These figures would give DC photocurrents of between $116p\,A$ and $425p\,A$ with a $3m\,W$ laser. It would be relatively easy to increase the power of the laser, for example to $300m\,W$, and the could be used for larger illumination areas or larger bandwidths.

4.4.5 Analog Front End and Log Pixel

The AFE was the only circuit which had not been tested up to this point. This was because it needed an optical input, rather than electrical. The AFE formed the first stage of the processing and was the only entirely analogue component. The tests performed were to check that the output of the AFE behaved as expected and would provide an appropriate response to light. The most important parameters would be the AC response of the AFE. The signal from the interferometer would be frequency modulated by the vibration of the object and therefore only the frequency of the signal detected by the chip would be important. The AFE would need to have the required bandwidth and amplification to pass this frequency onto the comparator faithfully.

The second consideration would be the DC output of the AFE. This would need to fit within the input range of the comparator. As has already been shown

the comparator has a range of $0.4\,V$ to $2.5\,V$ over which the input will cause a clipped output.

DC Response The first test was to check that DC response. The results are shown in Figure 120 as a plot of photocurrent against output voltage. The photocurrent could not be measured but it could be calculated from the incident light power. The experimental values show the the expected shape from the simulations, however with a DC shift. This was also noted in the work of Diwei He[96] when testing the chip for which the AFE was originally designed. This type of variation is expected due to process variation, and while it shifts the DC output of the AFE, this is not an important parameter. As already stated none of the vibration information is held in the DC irradiance and therefore the actual DC level coming out of the AFE is only important if it prevents operation of the next stage in the processing: the comparator. The output DC voltages of the AFE fit well within the DC input voltage range of the comparator as shown in Figure 104 in Section 4.3.3.

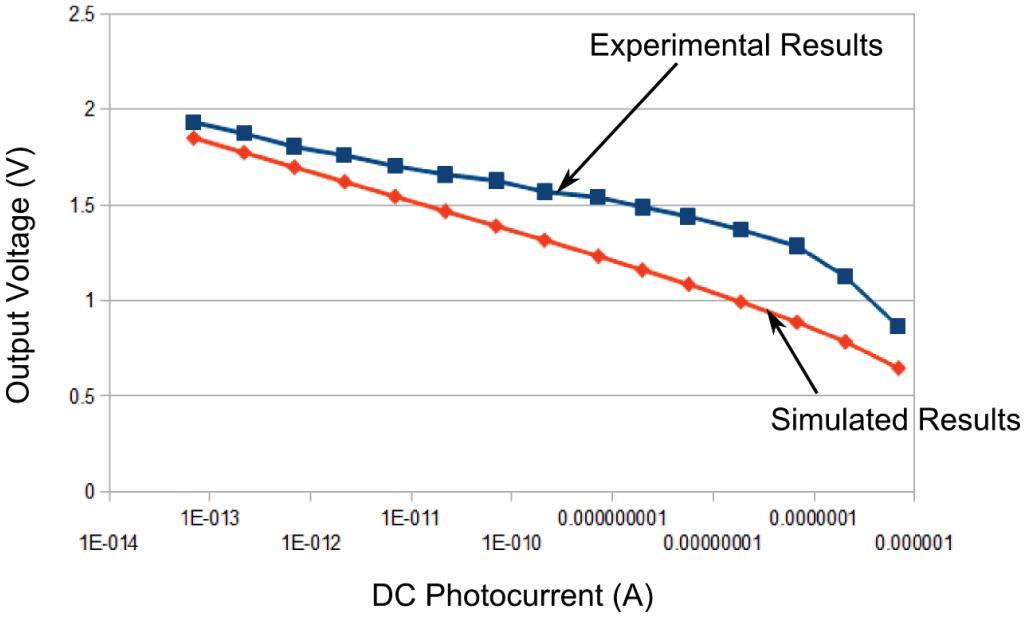

Figure 120: DC response of AFE for photocurrents from $100f\,A$ to $1\,\mu A$.

AC Response The second check was for the AC response of the AFE. This was an important parameter as the AFE needed to provide a voltage which had been amplified sufficiently to switch the comparator. The light source for the modulated signal was the LED array. This had a limited maximum light power and could provide DC photocurrents between $1\,pA$ and $200\,pA$. The LED array was set to give a DC illumination which gave the required DC photocurrent and calibrated to give a 10% modulation depth signal. This calibration was done with the reference photodiode. This signal was swept across a range of frequencies and gave the results in Figure 121. This figure also has the appropriate simulations for comparison. The lines with the points marked are the experimental values and the continuous lines are the simulation values. These show almost matched gains for the AFE compared with simulated values.

Noise The noise was calculated on a range of DC photocurrents and is shown in Figure 122. This also shows the simulated values from Chapter 3 for comparison. It can be seen that the AFE exhibits slightly higher noise compared with simulation, which is probably due to the PCB and interconnects rather than the chip as the shift is not affected much by photocurrent.

Some time plots of the output traces from the AFE are given in Figure 123. They show a decrease in amplitude with increased frequency, however there is also a variation in DC level of the signal as the frequency increases. This low-frequency noise causes problems for the comparator as the reference voltage is fixed. Changes in the DC level would need to be taken into account when setting the hysteresis of the comparator.

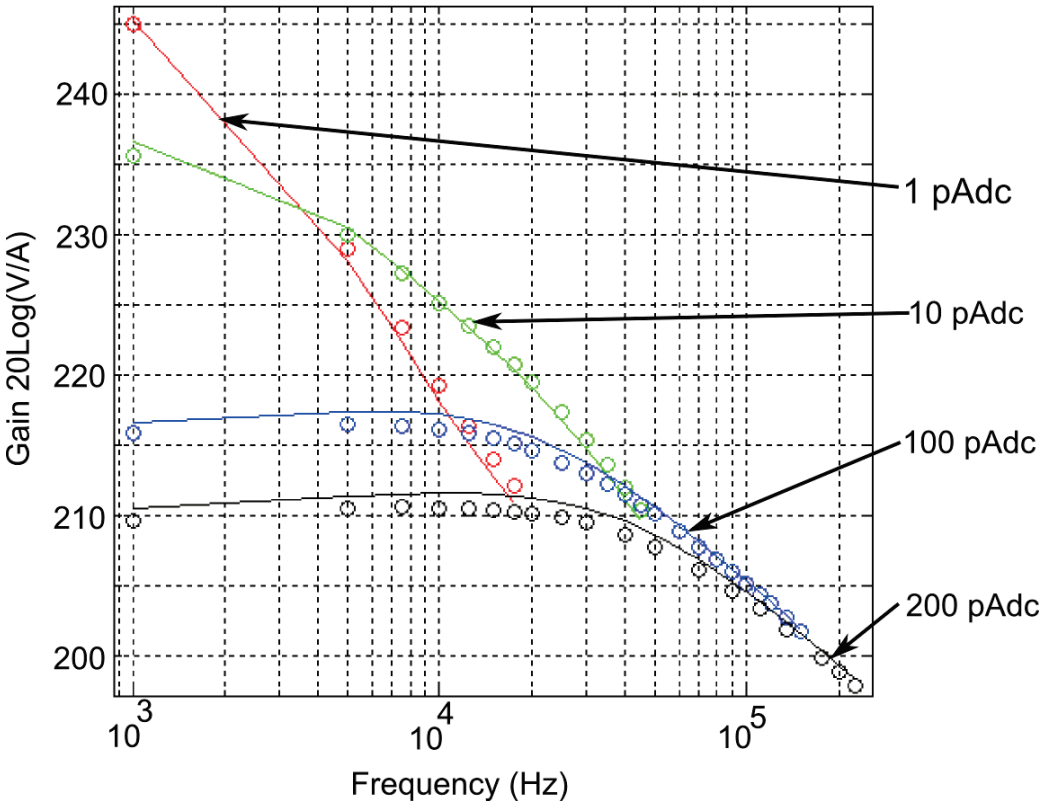

Figure 121: Comparison between simulated and experimental transimpedance gain of the AFE.

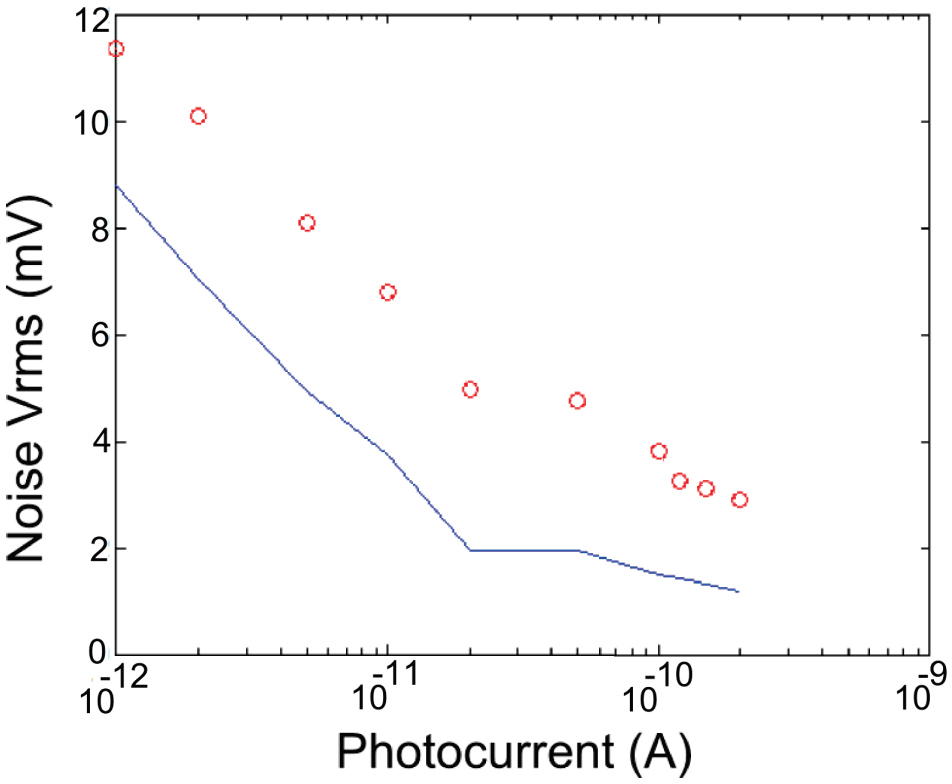

Figure 122: Output noise for the AFE vs. DC photocurrent.

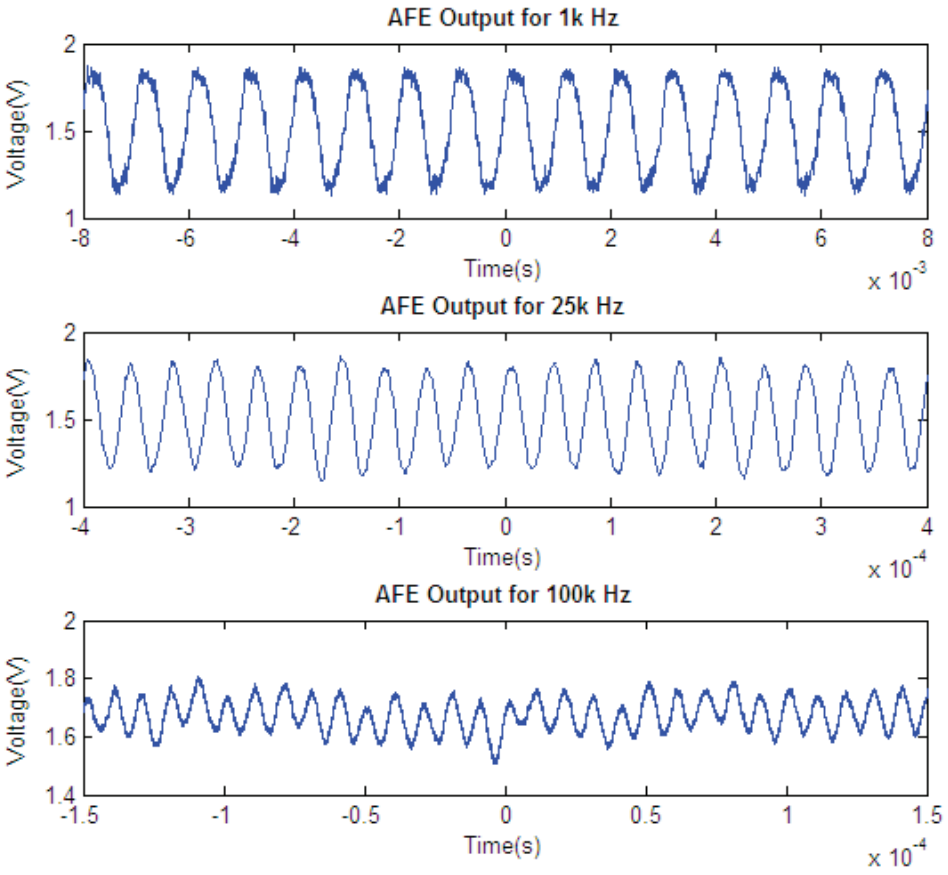

Figure 123: Plots of the output voltage from the AFE for different frequency optical inputs.

4.4.6 Whole Circuit Operation

The final test was to check how the whole circuit would perform using the LED array to provide an optical signal. For comparison purpose the same theoretical vibration signals as used in Section 3.8.2 were recreated using the signal generator. This was an object vibrating at $500\,Hz$ with amplitudes of $500n\,m$, $1u\,m$, $1.5u\,m$ and $2u\,m$. The Brag Cell frequency was set to $25k\,Hz$.

It was found that the noise on the AFE output caused phase shifts in the comparator output which in turn were interpreted as frequency changes by the timing circuit. This resulted in a great deal of noise on the output of the timing circuit for low modulation depth signals. Therefore, although the simulations were

for 10% modulation depth a level of 35% was used as this provided a more stable signal.

Section 3.6 has shown that the bandwidth is related to the DC light level. This required a DC illumination of $100pA$ to provide the bandwidth for the intended signal. This was achieved and Figure 124 shows the oscilloscope traces for the three outputs: comparator, HDA and final timing output.

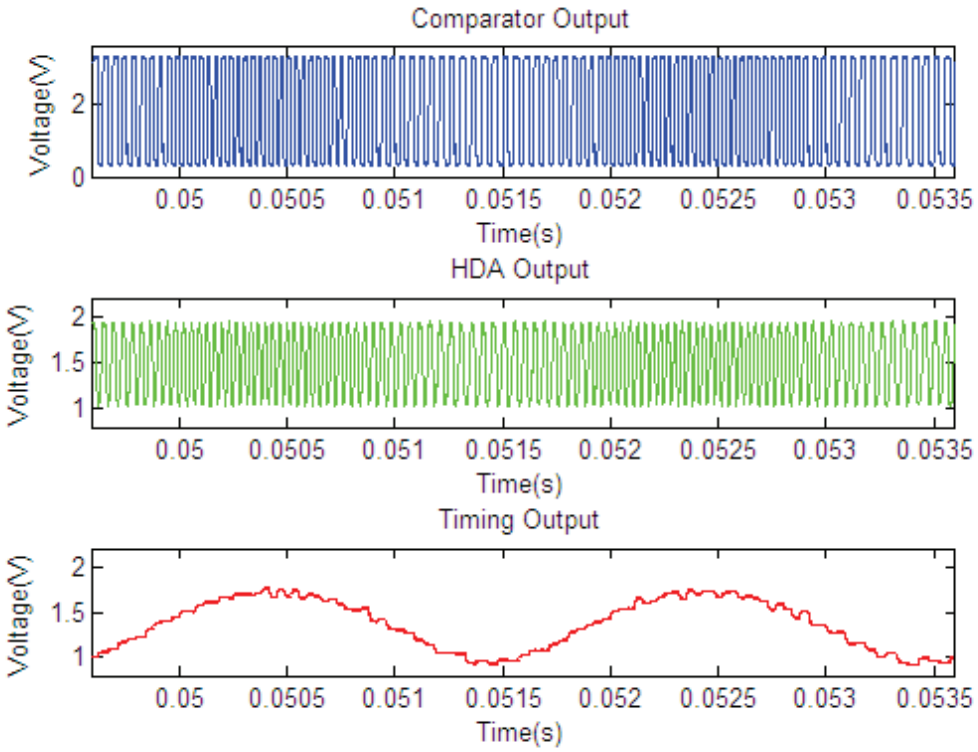

Figure 124: Oscilloscope traces showing the chip in operation

This output can be compared with those simulated in Section3.8.2 and this is shown in Figure 125. This figure shows a close agreement with frequency of the traces and their amplitude. There is also a slight DC shift between the traces. The periodic variation on the experimental results is in keeping with the Matlab simulations, however the Cadence simulations show more filtering of the output signal. These variations were due to the voltage steps at each falling zero-crossing

as the timing capacitor voltage was transferred to the output.

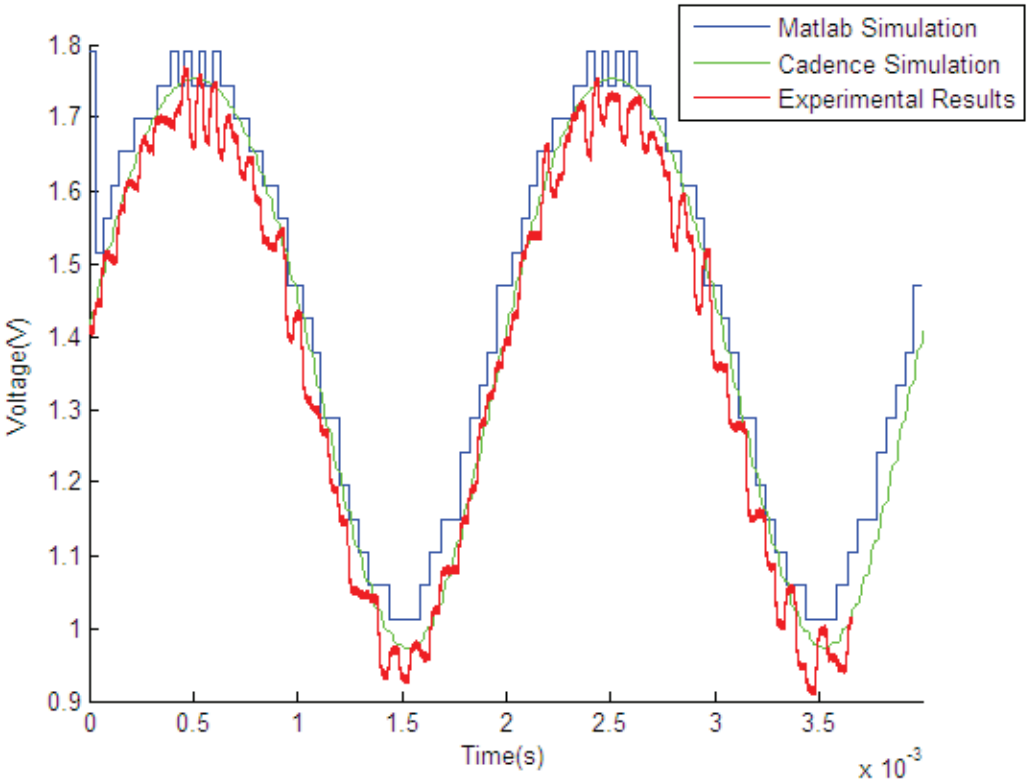

Figure 125: Comparison of Matlab simulation, Cadence simulation and experimental results for $500\,Hz$ vibration with $500n\,m$ amplitude using a $25k\,Hz$ Brag Cell frequency and a modulation depth of 34% for the experimental results and 10% for the simulation.

However when the optical signal modulation depth was reduced to 14% the experimental signal shows much more noise. This is shown in Figure 126.

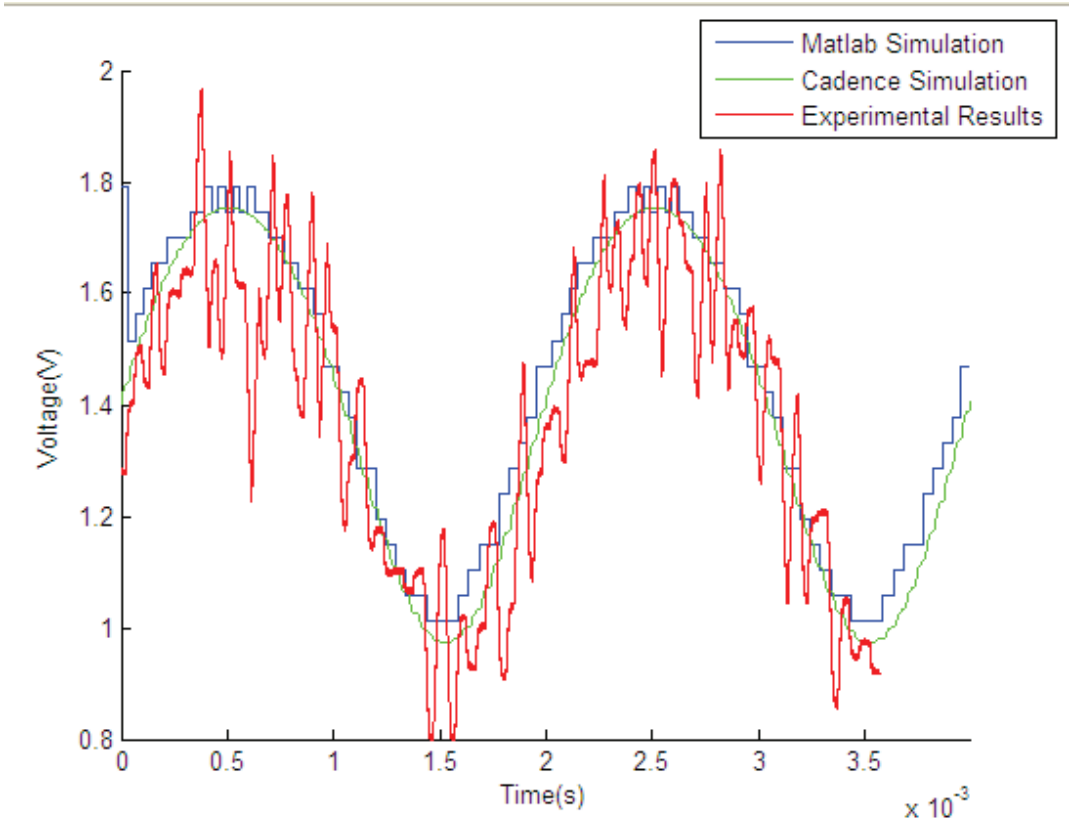

Figure 126: Comparison of Matlab simulation, Cadence simulation and experimental results for $500\,Hz$ vibration with $500nm$ amplitude using a $25k\,Hz$ Brag Cell frequency and a modulation depth of 14% for the experimental results and 10% for the simulation.

This noise is now much greater than that predicted by with either Maltab or Cadence. To check this was due to the AFE output. The AFE output from the chip was recorded and passed through the Matlab simulator to check whether the noise was still present, or whether the comparator, pulse generator or timing circuit were the source. This is shown in Figure 127.

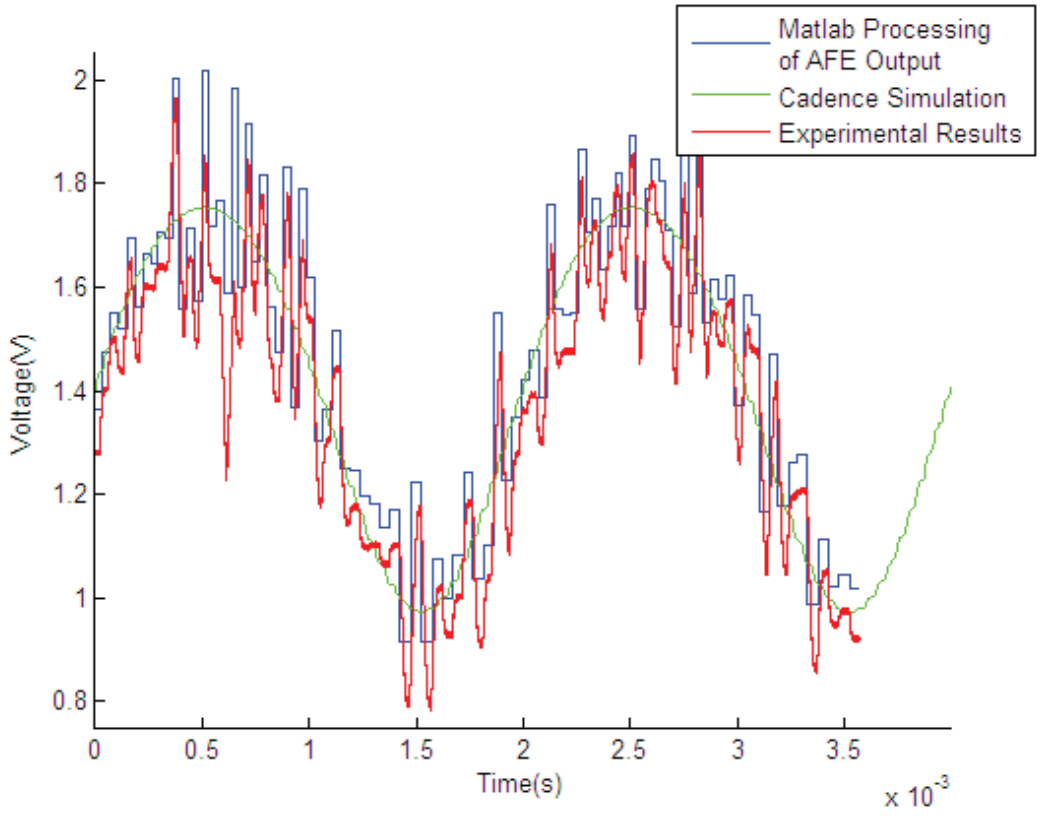

Figure 127: Comparison of Matlab processing on experimental comparator output, Cadence simulation and experimental results for $500\,Hz$ vibration with $500n\,m$ amplitude using a $25k\,Hz$ Brag Cell frequency and a modulation depth of 14% for the experimental results and 10% for the simulation.

This noise on the output of the circuit is still present with Matlab processing therefore it originated within the AFE. Figure 128 shows a comparison between the AFE output, the comparator output and the timing output from the experimental results used to generate Figures 126 and 127. This shows that transient variations in the phase of the AFE output causes the comparator pulse width to change. This in turn affects the discharge level for the timing circuit for a single cycle. This can cause a significant change in voltage of the timing output.

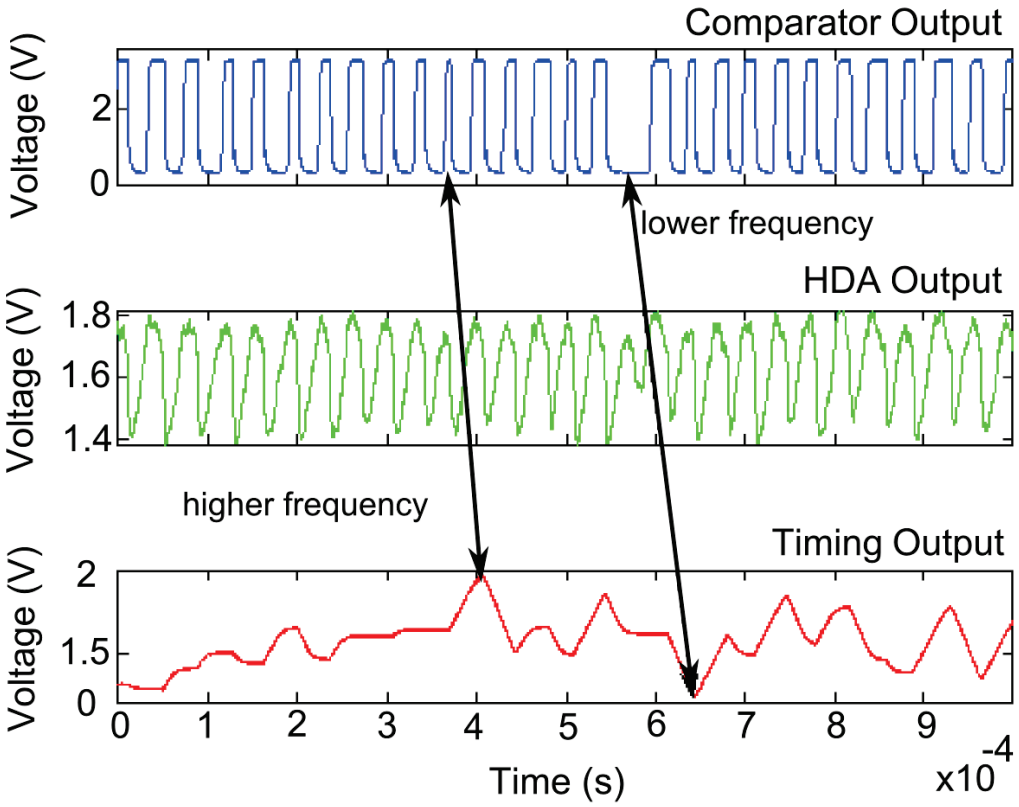

Figure 128: Oscilloscope traces showing how variation in the AFE output produces noise in the timing output

4.5 Conclusion

This section has shown the operation of the chip with both optical and electrical input signals. The electrical tests were used to check the operation of the parts of the circuit which did the processing, and they showed the circuit operating as expected from simulation, with the exception of the pulse generator. The light sensitive AFE was then tested with a modulated light source. This showed that the AFE worked as predicted, although with a slight DC shift in its output voltage. This did not, however affect the rest of the circuit.

The tests did, however, show that noise on the AFE output was a major source of demodulation errors on the output. This noise required that the modulation

depth to the optical signal be larger to overcome the phase changes the noise added to the comparator output. The tests also showed that the pulse generator failed when set to long pulse widths. This was attributed to latch-up within the circuit, and guard rings were suggested as a solution for future work.

5 Testing of the Laser Doppler Vibrometry Pixel with Vibrating Objects Comparison with Current Technology

5.1 Introduction

So far this thesis has covered the design of full-field, laser Doppler vibrometry pixel, intended for a FF LDV camera. Chapter 1 has shown where the work falls within the field of non-destructive evaluation and described the LDV method at a high level. Chapter 2 gave more detail into the theory and apparatus used for this method. Chapter 3 covered the design of the pixel and Chapter 4 has shown the initial testing of the finished pixel.

This chapter will not show some real-world applications of the finished system. This will show the pixel in operation on a controlled vibrating object (the piezoelectric bimorph from Chapter 2) as well as a vibrating membrane. The membrane tests show a case where the vibrating object is less well controlled, as well as an ideal application for a FF LDV camera, of which the circuit developed here could form a single pixel. A comparison with a commercial single point vibrometer will be made.

5.2 Interferometer Tests

5.2.1 Introduction

The final testing stage was to incorporate the chip described in Chapter 3 into an Mach-Zehnder interferometer like that described in Chapter 2. The chip forms the electrical detection and processing part of the laser Doppler system, sampling the irradiance from the interferometer and demodulating this into a vibration signal.

As such the chip requires an interferometer to be setup to provide the input.

This testing stage would test both the chip and the interferometer and show the whole system in operation. The first part was to setup and align the interferometer. A bimorph (Part number: BM15015-06HC, stock number on RS Components Ltd: 285-784) was used as the vibrating object and it was vibrated at different vibration amplitudes and frequencies and the results compared with a commercial single point vibrometer.

The test required vibrating the bimorph and measuring the frequency and amplitude of the demodulated output signal from the chip. The frequency should be the same as the input signal to the bimorph and the scaled amplitude of the output should be the same as the amplitude of the vibration, as measured by the commercial LDV.

The demodulated output from the chip was scaled to give a signal measured in nanometers. This is covered in Section 3.8.2 in Chapter 3. The amplitude of the vibration was found using a commercial, single point vibrometer (Polytec OFV 534/ OFV 2570).

5.2.2 Test Setup

The interferometer was set-up with the following components:

- $2.8mW$ He Ne laser (Melles Griot - 25-LHP-121-230) producing $632.8\,nm$ light. This laser produces collimated, polarized light $1\,mm$ in diameter.
- Polarizing 50:50 beam splitters
- A half wave plate
- Two quarter wave plates
- Lenses

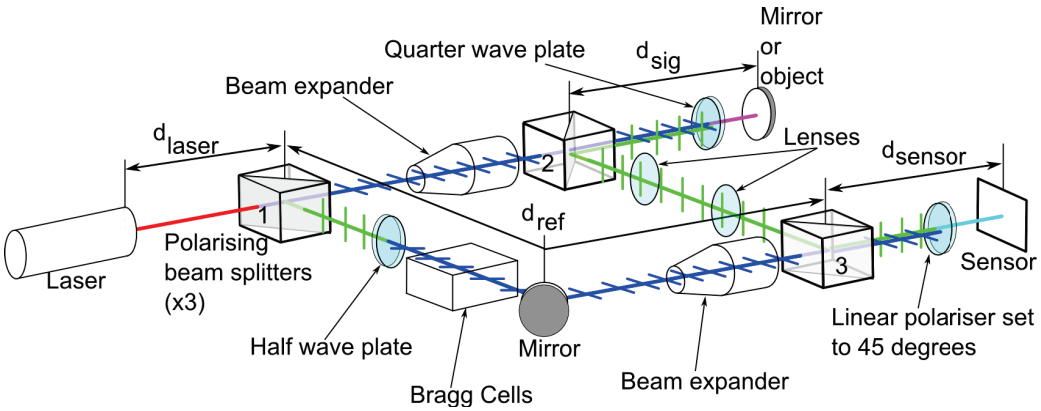

Figure 129: Diagram of the MZI interferometer.

- Optical cage system
- Mini optics bench

The system diagram of the MZI interferometer from Chapter 2 is repeated in Figure 129 and a photograph of the set-up is shown in Figure 130

The vibrating object was a piezoelectric finger called a bimorph and was placed close to the interferometer for coherence reasons, see Section 2.2. The bimorph was driven by the signal generator using a step-up transformer to increase the voltage. The surface of the bimorph gave reasonable reflectivity although optical speckle was evident.

The interferometer was setup on an optics bench to provide a vibration damped surface. The commercial LDV was then aligned such that its measurement beam was directed at the back of the bimorph in question. The guaranteed that the vibration amplitude of the bimorph would be the same for the MZI LDV and the commercial LDV. This allowed real-time comparison and made the testing more efficient.

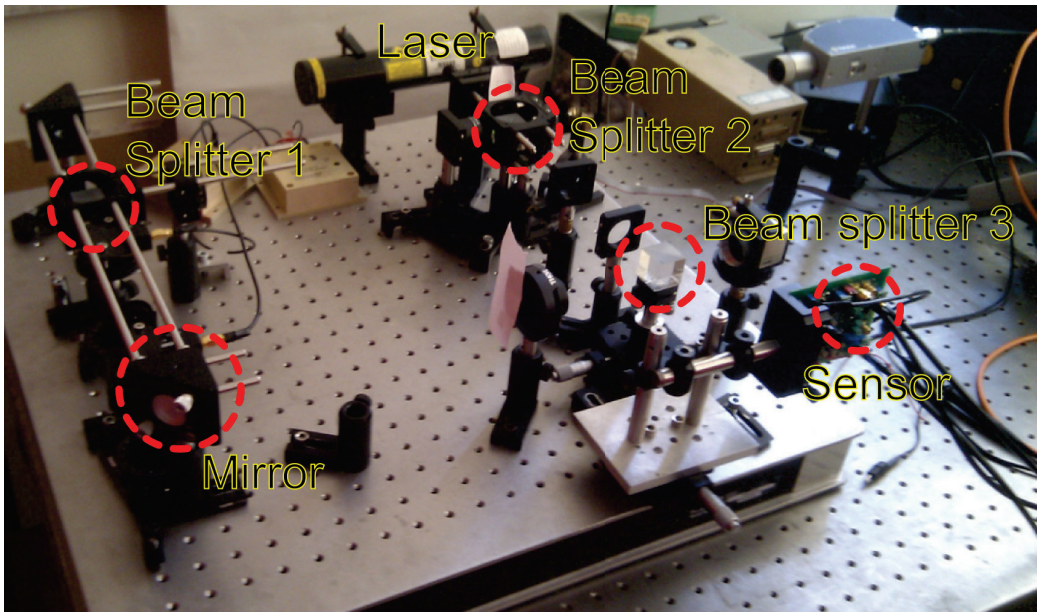

Figure 130: Photograph of the experimental set-up for testing of the ZC pixel.

5.2.3 Method and Results

The interferometer was constructed and aligned, see Appendix B. The bimorph was then positioned and imaged onto the chip using white light. As the chip had a single pixel it was important to carefully position the light sensitive part of the pixel in a speckle. Once the pixel was in position the Bragg Cells were switched on. The Bragg Cells were designed to operate at $\sim 80\,MHz$ which is considerably higher than the frequency response of the pixel. In order to provide a smaller frequency shift two Bragg Cells were used in series - one feeding into the other as shown in Figure 131.

The drive frequency of the first Bragg Cell is slightly higher than the drive frequency of the second. If first the +1 order and then the -1 order are chosen then the final beam is shifted by the difference between the two drive frequencies. For simplicity this frequency difference is referred to as the Bragg Cell frequency. This signal showed good stability with $< 0.25\%$ deviation from the desired frequency and $< 1.1\%$ variation over ten minutes. The exiting beam is parallel to the

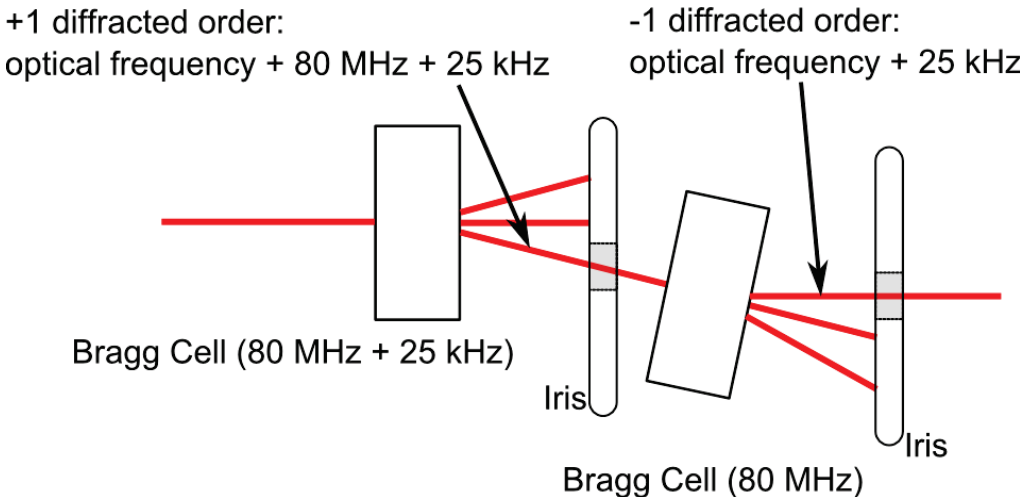

Figure 131: Using two Bragg Cells to provide a smaller frequency shift than the cell drive frequency.

entering beam, but it is shifted laterally. This was corrected by translating the beam splitter 3 on Figure 129.

The outputs from the pixel were checked on the oscilloscope. The Bragg Cell frequency was set to $25\,kHz$ through the control software. The expected output from the pixel was the detection of the carrier frequency ($25\,kHz$) and a DC output from the processing indicating that there was a constant frequency input. The results are shown in Figure 132.

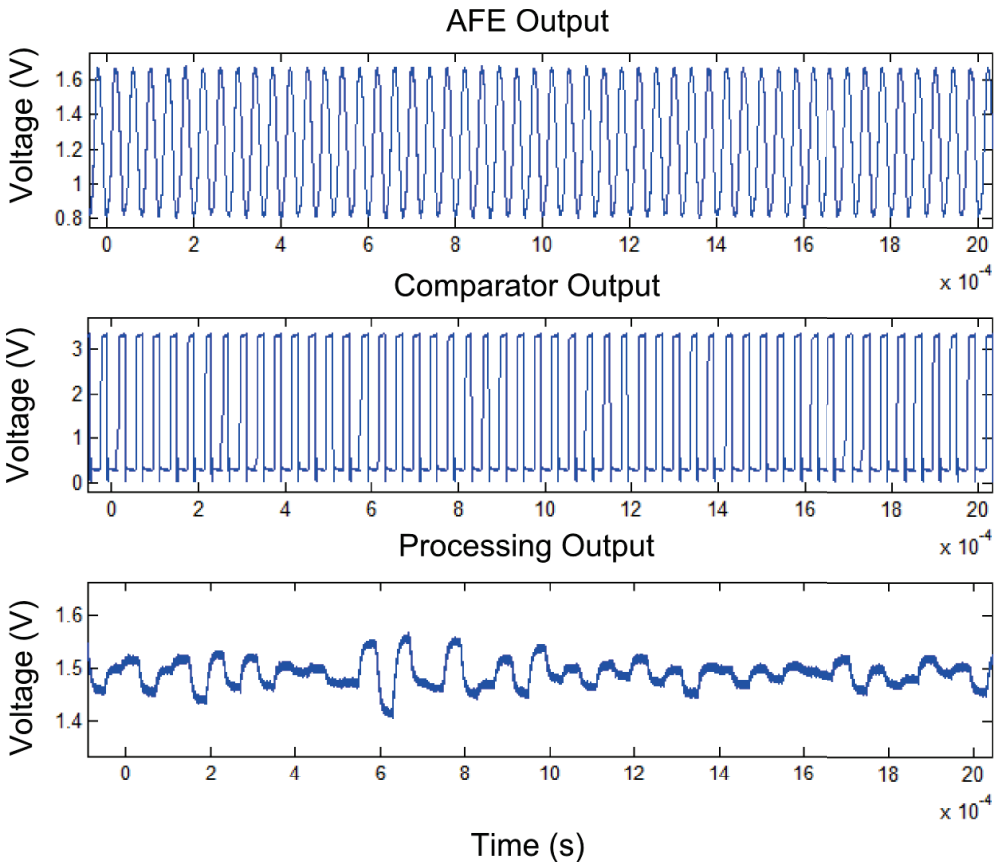

Figure 132: Outputs from the chip during testing. The Bragg Cells running at $25\,kHz$ with the object stationary.

The operation of the timing part of the pixel was described in Section 3.8 and tested in Section 4.3.5. For these tests the bias resistance for the timing part of the circuit was set to $26\,k\Omega$ and the Bragg Cell frequency was shifted while measuring the demodulated voltage from the pixel. The results are given in Table 11 and Figure 133 and show the processing operating as expected. This also allowed calibration of the actual value of k_{pv} and actual supply voltage. The calibrated value for the supply voltage was $3.28\,V$ and the timing conversion constant was $44.58\,kV/s$.

Table 11: Table of calibration values for the output voltage against the input frequency.

Input Frequency (kHz)	Output Voltage (V)
18	0.81
22	1.24
25	1.49
27	1.63
30	1.8

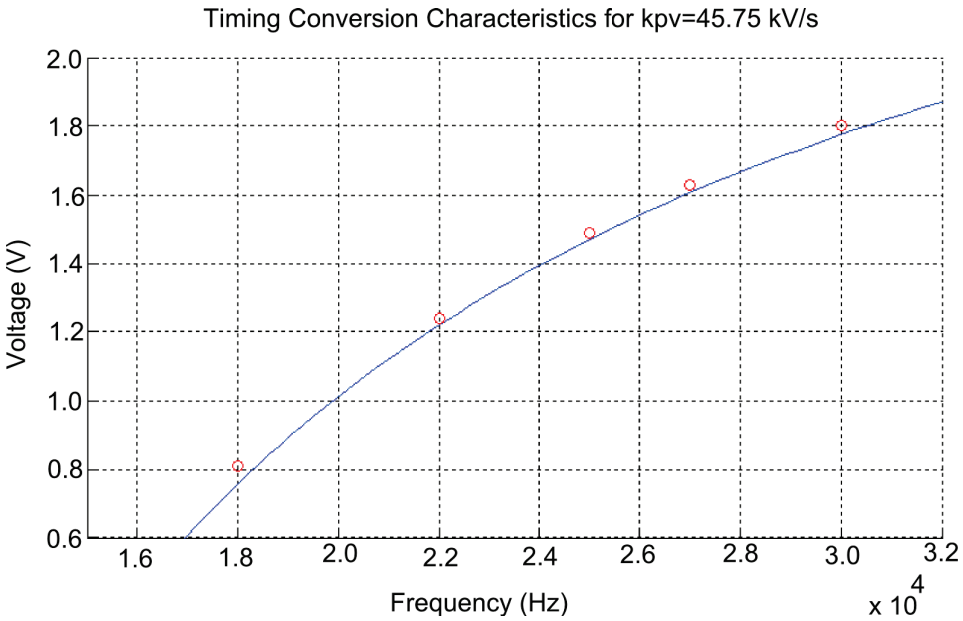

Figure 133: Plot of calibration results for the frequency to voltage conversion characteristic of the timing circuit compared with the theoretical value for $3.3\,V$ supply voltage and $45.75\,kV/s$ k_{pv}.

Once this setup stage had been finished the Bragg Cell frequency was set back to $25\,kHz$ and the bimorph vibrated at $1\,kHz$ with a range of different amplitudes from $200\,mV$ to $1.1\,V$.

Figure 134 shows some results comparing the measurements from the commercial LDV and the pixel. The commercial signal needed digital filtering to remove wide-band noise so the two signals could be compared. The filter used was a 64 pole, FIR, low-pass digital filter with a cutoff frequency of $10\,kHz$. The data from the pixel and the commercial LDV was captured by a digital oscilloscope

and processed in Matlab. In Figure134 both sets of data were scaled to give the vibration amplitude and overlaid to allow comparison. Three different drive voltage amplitudes are shown, $200\,mVpp$, $500\,mVpp$ and $1.1\,Vpp$, all at $1\,kHz$ drive frequency. At $500\,mVpp$ both signals are comparable but the pixel's signal has more low frequency noise which appears as periodic fluctuations. These are the result of the timing circuit operation - each time the timing capacitor is sampled and reset, the output voltage is updated. This update process introduces step in the output voltage. At $200\,mV$ the commercial LDV's signal to noise ratio starts to be decrease, while the pixel is still exhibiting a similar noise profile. At $1.1\,V$ the signal bandwidth has increased enough that the negative peaks starts to clip, however the positive peak still gives a measure of the vibration amplitude. This is a result of the timing conversion and shows how this non-linear conversion gives increased dynamic range.

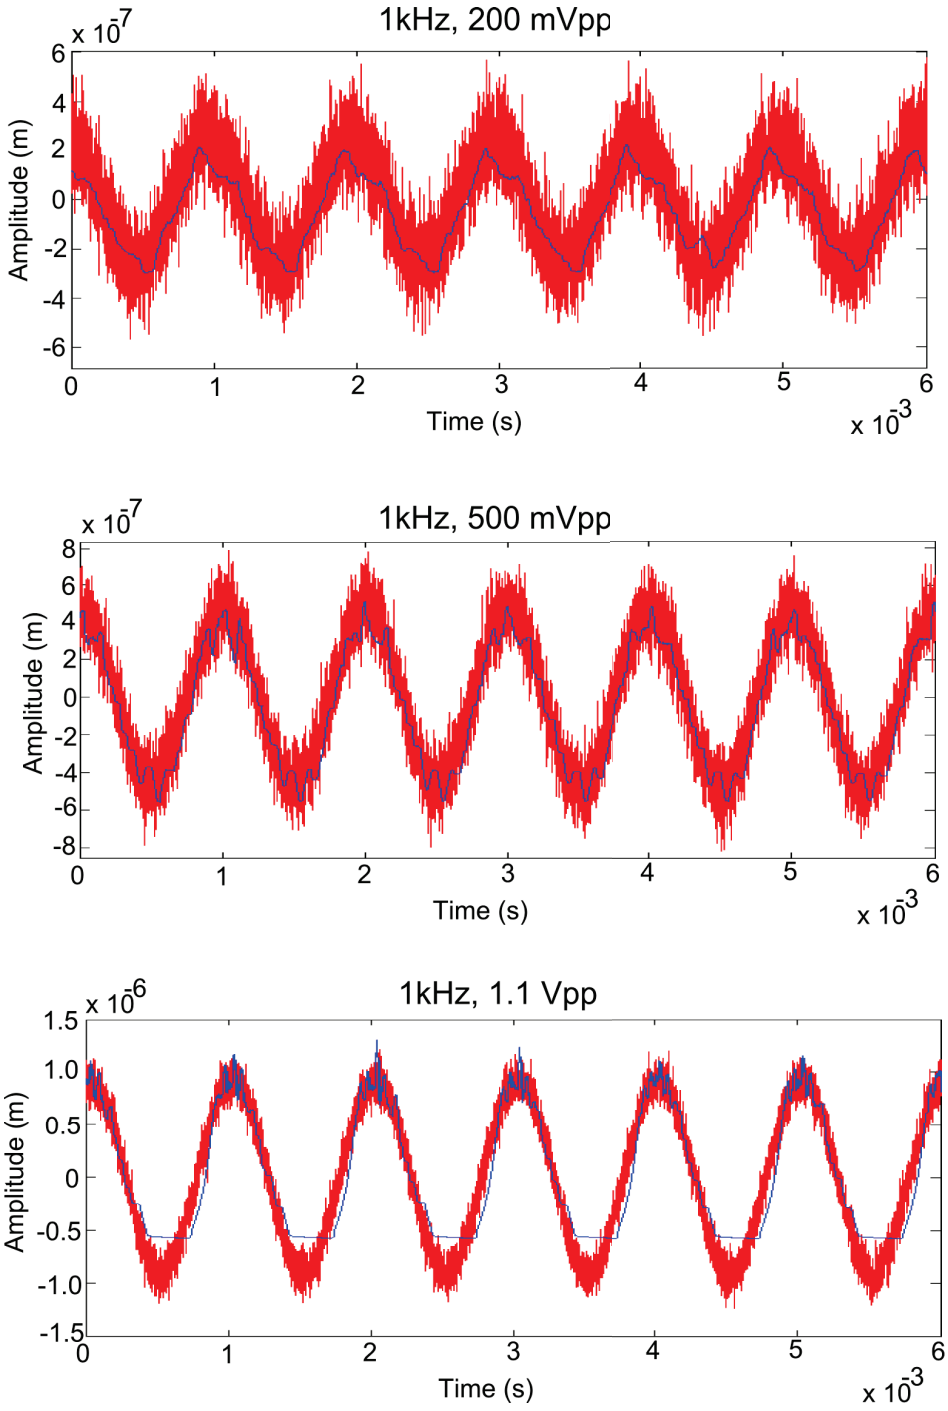

Figure 134: Results comparing the measurements from the commercial LDV with the pixel's processing for different drive amplitudes.

The amplitudes measured by the two devices are listed in Table 12. This table also shows the percentage error in the pixel's values. It should be pointed out that

the commercial output has been filtered to reduce the wide-band noise, while the output from the pixel has been left unchanged.

Table 12: Comparison between the amplitudes of the vibrations measured by the pixel and the commercial LDV.

Voltage Amplitude (mV)	Commercial LDV (nm)	Pixel (nm)	Error(%)
200	197	188	±4.8
228	255	209	±18
500	461	440	±4.6
600	556	529	±4.9
1000	915	890	±2.7
1100	1001	942	±6.3

5.2.4 Conclusion

This stage of testing has shown the MZI LDV in operation with a real-life vibrating object. The output has been compared with a commercial single point system and shown to be in good agreement with $< 5\%$ amplitude error in most cases, and $< 0.1\%$ error for frequency. This forms the final stage of the proof of principle of this LDV system. There is still narrow-band noise on the output from the pixel due to the operation of the timing circuit. This noise can be reduced by increasing the Bragg Cell frequency compared with the signal bandwidth. This will reduce the phase difference between each zero-crossing and therefore also reduce the voltage step at each crossing. The definite steps that appear on the output are due to timing circuit resetting due to the zero-crossings. This could be reduced if the Bragg Cell frequency was increased. This would, however require an AFE with a wider frequency response.

5.3 Full-Field Application - Membrane Vibrations

5.3.1 Introduction

A good application of full-field vibration imaging is a vibrating membrane. A circular membrane is used and demonstrates a situation where non-contact vibrometry is vital [16], and full-field vibrometry is an advantage. An ideal membrane should have no stiffness therefore it's oscillations depend only on the restoring force supplied by the tension its under [127]. If the membrane is vibrated at an appropriate frequency then waves traveling across the membrane can form standing waves. The places where the vibration is zero are called nodes and can be used to describe the standing wave pattern which is called a mode.

There are many different modes in which a membrane can vibrate. Some of these are shown in Figure 135. The values in brackets are the mode designations for an ideal membrane in the form (m,n). They list the standing waves which make up the vibration nodes in radial coordinates. n shows the number of standing waves in the radial direction and m shows the number of standing waves in circular direction.

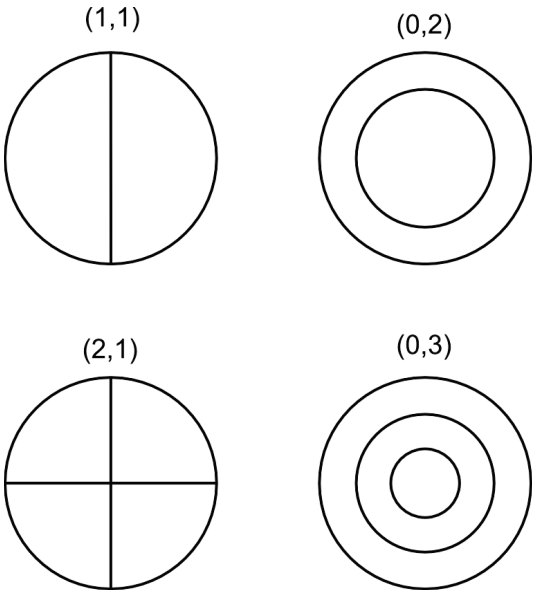

Figure 135: Some of the vibration modes of a circular membrane, taken from [127].

The field of membrane vibration research has been around for at least three hundred years [16] and has been used in such fields as musical instruments, such as violins and drums, and has applications in architectural and civil structures. Recent interest in membrane structures for use in space due to the possibility of reduced stowed volume and launch mass [26]. Examples of such structures are solar sails, concentrators and shields. Important considerations for these type of structures include surface smoothness, profile and slope error. The reasons for vibration analysis in these cases includes reducing vibrations for sensitive devices, and active shape control for which use would be made of an on-board LDV.

This makes vibrating membranes a good example of the applications of a FF LDV system. This section will describe the design, construction and use of a circular membrane experiment using the LDV system described so far. The experimental setup is similar to that of Jenkins et. al.[16] using a mounted membrane with a speaker as the vibration source.

5.3.2 Experimental Setup

In the case of this experiment the membrane was a semi-reflective plastic sheet mounted on the end of a cylindrical support, like a drum. A frame was needed to hold the membrane in place, tension it and support the speaker which vibrated the membrane. Two photographs of the membrane support are given in Figure 136. This structure was made in two parts as shown in Figure 137. The first part, the membrane mount, simply clamped the membrane material in place around a circular hole. When a metal cylinder was pushed through the hole in the membrane mount it would uniformly tension the membrane. The second part was the rest of the mount. This was based around a wooden support onto which was bolted the steel cylinder. A hole was cut in the center of this support to allow sound from the speaker mounted at the back to travel through into the steel cylinder, and against the tensioned membrane.

The final membrane structure is shown in Figure 138.

- The membrane diameter was $9.4\,cm$ formed by a Teflon coated, round cake baking form. This was a simple solution to the problem of a rigid, smooth and round form.

- The wood for the mount and frame was $15\,mm$ compressed chip board which showed excellent stiffness and vibration resistance.

- The speaker was a $75\,mm$ base driver which has sufficient range to cover the frequencies of interest ($< 2\,kHz$).

- The membrane was a plastic/ silver foil. Although unusual this material was very light, smooth and resistant to stretching. When pulled tight it gave a mirror surface, and did not show appreciable change in properties, even after prolonged mounting.

During the design and operation of this mounting system several changes were made to improve the results. The first was that it was found that the speaker cage (the metal frame around the vibrating cone of the speaker) also vibrated when the cone was energised. When the cage was rigidly mounted against the rest of the frame, these vibrations would pass through to the metal cylinder and this had the effect of vibrating the membrane at its edges as well as the middle and distorted the vibration modes. To combat this the speaker was disconnected from the rest of the mount and then pulled against foam pads using rubber bands. This change meant that any vibrations of the speaker cage would not be transferred to the membrane frame while still keeping the speaker in the same place.

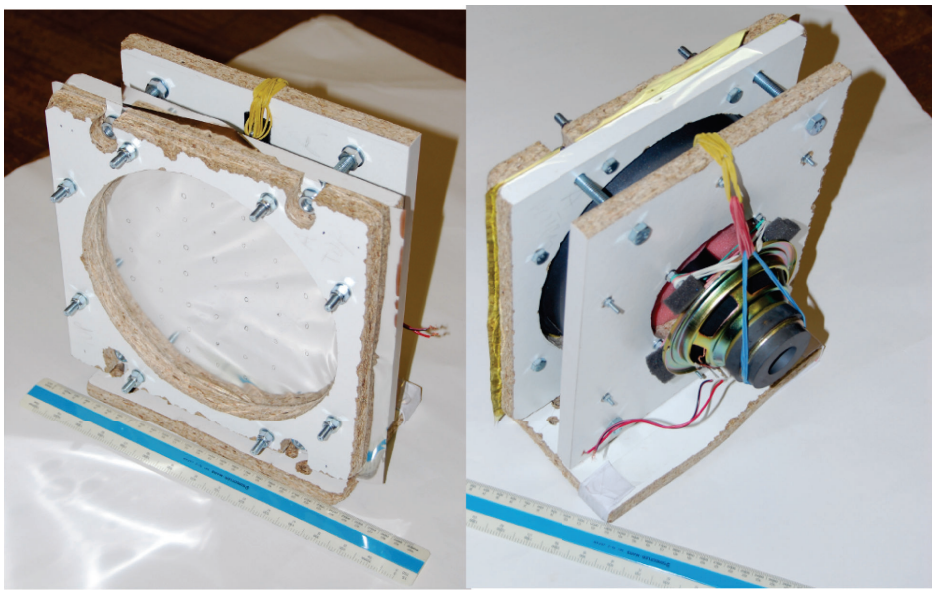

Figure 136: Photograph of the membrane and support frame.

The second change was to limit the sound excitation from the speaker to a 1" hole behind the center of the membrane. This did two things. Firstly it limited the sound source which excited the membrane to a source in the middle, rather than the whole size of the speaker. Secondly it decoupled the air behind the membrane (within the metal cylinder) from the air around the speaker. Before the speaker had been sealed against the back of the metal cylinder. This resulted

in the whole system operating a bit like a balloon. When the speaker moved forward the membrane was pushed out, and when the speaker moved back the membrane was pulled back. This restrained the free vibration of the membrane.

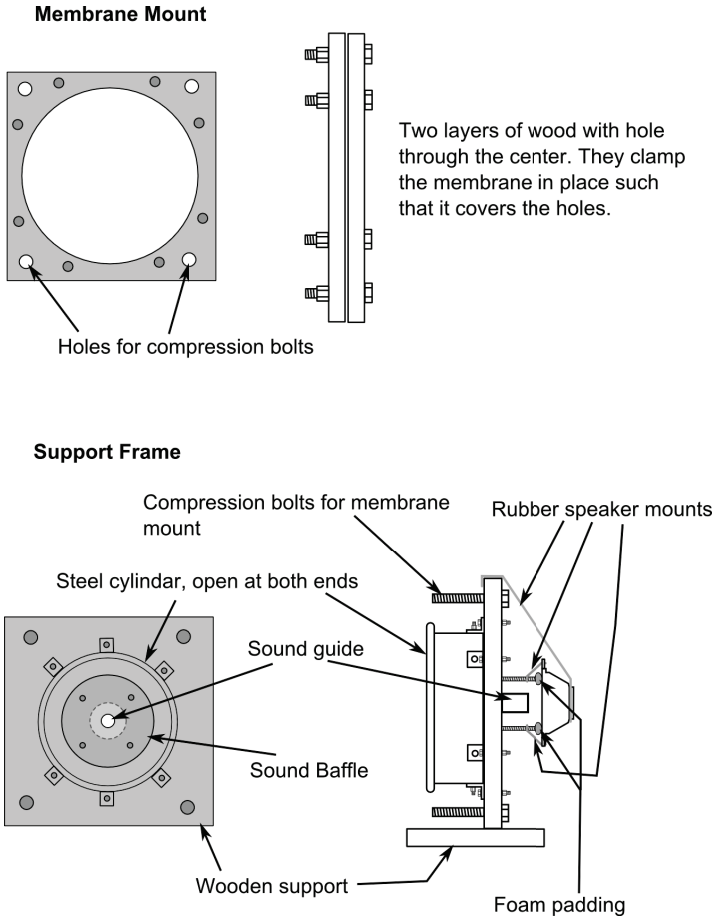

Figure 137: Diagram of membrane mount and the frame that supported the speaker and membrane.

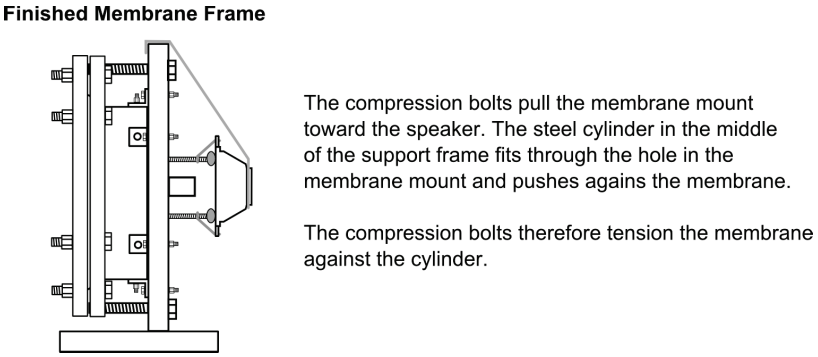

Figure 138: Finished membrane frame.

5.3.3 Experimental Method

A grid of points (with $2\,mm$ separation) was drawn on it using permanent marker which allowed measurement of where the laser measuring point of the MZI was on the surface. The membrane mount was pulled against the support frame using the compression bolts. Turning each bolt the same number of turns ensured that the tension was uniform. The membrane was held vertically as shown in Figure 138 and the LDV directed against its surface. A signal generator (Agilent 33220A) was used to drive the speaker and an oscilloscope was used to capture the results (Tektronix MSO 4054). The following procedure was carried out:

1. The LDV was directed at the center of the membrane.
2. The speaker was driven with a signal of $20\,Hz$.
3. The speaker signal frequency was increased until the vibration reached a maximum. The membrane frame was then moved to scan the LDV measurement spot across it.
4. The location of any zero-vibration nodes were noted.
5. The frequency was then increased and steps 3 and 4 repeated.

Once the first two mode of vibration of the membrane had been found, to illustrate the sensitivity of a membrane system, the top left compression bolt was loosened by half a turn and the node maps remeasured.

5.3.4 Results

The results of this experiment are shown in Figure 139. The top left and bottom left plots show the 0,2 and 0,3 modes of vibration of the membrane with equal tension on the compression bolts. The circle on all the plots of Figure139 shows

the circumference of the membrane. These modes are the most likely for this setup due to the vibration source being in the middle of the membrane.

The top right and bottom right plots show the results after change of the tension at the top left compression bolt. The frequencies at which the modes occurred changed and the modes themselves were distorted. The 0,2 mode was mainly unchanged, but 0,3 showed a shape similar to that of a square membrane. Also the nodes for this mode were badly defined which suggests that the standing waves were not perfect. They also shifted slightly over time, resulting in the sharp distortions at the bottom of the 0,3 plot.

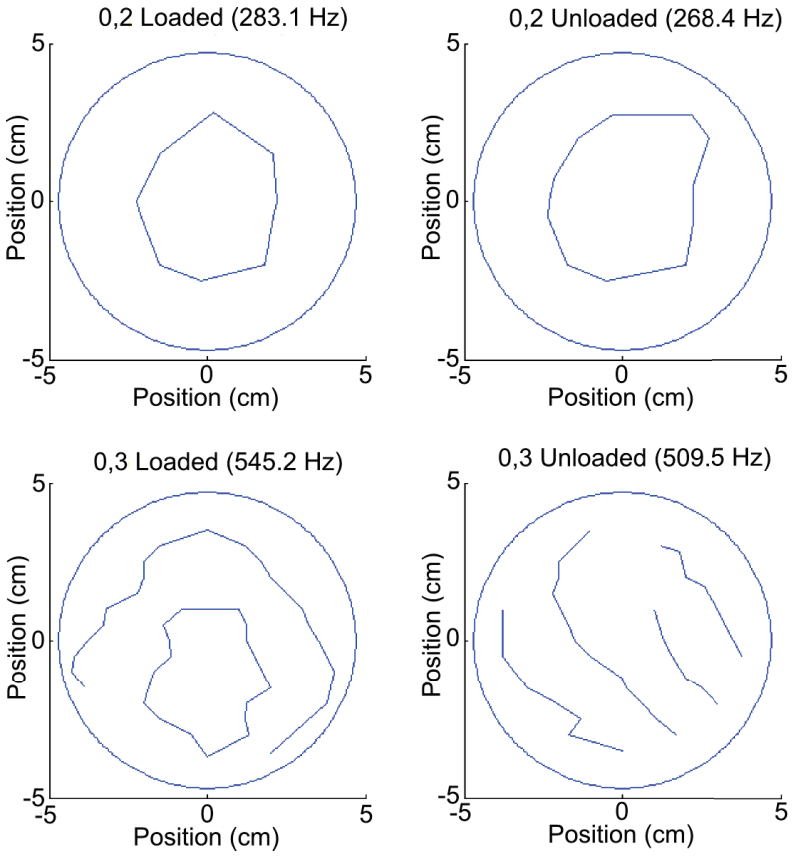

Figure 139: Results from two vibration modes with and without tension on the top left compression bolt.

5.3.5 Discussion

This experiment shows how sensitive a membrane is to changes in its mounting. Referring to Figure 139 the 0,2 mode is shifted down from its ideal position. This could be due to a slight mismatch in the tension of the mounting. The results show how sensitive the system is to change in the tension [128]. This prevents the use of any form of contact vibration measurement. The unloaded 0,3 mode also shows where a full-field measurement would be of benefit. The nodes shifted in an unpredictable way during the experiment and if the whole surface of the membrane was being imaged then a more complete idea of how the nodes were changing would be possible.

5.4 Conclusion

This section has shown some experimental results for the LDV developed during this project. The first experiment was designed to compare the LDV pixel developed in Chapter 3 with a commercial single point LDV. This comparison allowed validation of the pixel and its demodulation method. The vibrating test object was a bimorph like that used in Chapter 2. The pixel showed that it could provide a demodulated output which gave an accurate measure of the vibration frequency with $<0.1\%$ error. The demodulated amplitude had a error of $<5\%$ in most cases. The error was mainly due to the timing circuit operation introducing steps into the output waveform. This error would be reduced by increasing the carrier frequency of the system. The pixel amplitudes were always below that of the commercial LDV, most likely due to the timing processing missing the peaks of the sinusoids.

The second experiment was to measure the vibration modes of a circular membrane. This demonstrated an ideal application of a full-field LDV. The results showed the sensitivity of the membrane to tension by changing the tension ap-

plied to one corner of the membrane mount. This application also shows the benefits of being able to image the whole surface at once as semi-stable modes of vibration are hard to characterise.

6 Conclusion and Future Work

The aim of this project was to investigate the possibility of designing a system for measuring the vibration of a region of the surface of an object in a non-contact manner. It was decided to develop a small circuit to provide the processing required to perform laser Doppler vibrometry. This small circuit was contained a light detector and a number of processing steps to provide a voltage proportional to the vibration of the surface. This single unit or 'pixel' could then be repeated in an array to produce a Full-field Laser Doppler Vibrometry camera (FF-LDV).

A literature review was carried out giving an overview of the applications of such a camera in the field of Non-destructive Evaluation (NDE) as well as other methods of providing full-field vibration measurements. The pixel was based on a LDV as these are a mature single point technology and the individual pixels of a FF-LDV need to perform the same task - measure the vibration at a given point on a surface. Zero-crossing demodulation was chosen as the processing method for the pixel and a circuit was designed to implement this processing. Several simulations were carried out on the circuit and finally it was laid-out as part of a CMOS chip and sent for fabrication.

The finished CMOS chip was tested both electrically and as part of a complete LDV system on both a vibrating bimorph and a vibrating membrane. The results from the bimorph were compared with a commercial single point vibrometer and showed good agreement with $< 5\%$ amplitude error in most cases, and $< 0.1\%$ error for frequency. The applications of such a FF-LDV camera would be in the field of non-destructive evaluation, specifically vibrometry. Some specific applications where full-field LDV might be of benefit are:

- Understanding vibration effects on the vibration resistance of hard disk drives in laptops [129]. This paper measured the vibration using accelerome-

ters and compared the results with finite element simulation. Application of a full-field vibration measurement method to this problem would allow better updating of the simulation model [24] to take into account the complexities of laptop construction (PCB thickness, component loading, different supports).

- The analysis of sound fields to identify sources in reverberant environments [130]. In this case the cabin of an aircraft. The complexity of the environment makes it difficult to differentiate between sources and their reflections. The use of a full-field LDV camera would remove the need to time realign the sampled data, as in the case of a scanning vibrometer [131].

The benefits of full field measurements of objects are that the whole object is imaged concurrently. This concurrent imaging is important to give the phase relation between different points. Scanning systems are capable of measuring vibrations over a region of an object but care must be taken to synchronise the data collected from each point to get an accurate map of the vibrations over the object. This makes transient motion harder to capture. Imaging vibrations over the whole object also has the advantage of faster imaging.

6.1 Future Work

The layout of the final pixel is shown in Figure 97 in Chapter 3. This pixel was not optimised for space, and multiple guard rings were added to separate the digital and analog components. The digital buffers are also not needed for the operation of the pixel, but rather to allow signals from a given pixel to drive off-chip circuits typically requires a great deal more power. They could be moved outside the camera array and connected to whichever pixel's output was needed. The current pixel size is $180\,\mu m$ x $60\,\mu m$. If a chip a similar size to the 32 x 32 pixel CMOS

chip used in Chapter 2 was made it would have 20 x 20 pixels. An obvious goal of any future work with this pixel would be to reduce the size. There are several places where savings could be made.

- Incorporate the AFE and comparator. Currently the AFE and comparator are completely separate. This does have the advantage of noise immunity. A solution would be to place a second stage of amplification in the AFE using a similar technique to the HDA in the current AFE. This is an amplifier where the feedback is driven with a low-pass filtered signal from the output. This would remove the need for a reference voltage.

  The pulse generator does not need a very rapid falling edge as the input is clipped with an inverter. This would remove the need for generating fast signals in the new AFE. The output from this new AFE would also have to have a DC value at the switching voltage of the inverter.

- The timing circuit could then be moved much closer to the pulse generator and it could be scaled down as it contains mainly digital parts which are are not as constrained as analog to specific width to length ratios.

A smaller camera size (say 32 x 32) could still be scanned over the object to gain better resolution, while still maintaining the advantages over current scanning techniques. While this is not concurrent, the small camera could provide sub-images and a number of pixels on each edge could be set to overlap between each scanned location. This would allow better stitching of the different sub-images and still remove the need to provide a secondary reference (for example an accelerometer) to time realign the data.

A specific fault with the pixel designed within this project was the problem with the pulse generator (see Chapter 4), attributed to latch-up within the circuit. This fault prevented pulses longer than about $100\,ns$ from being generated due to

the reset pulse latching 'high' rather than falling after the required time. While this did not prevent the chip from operating, or prevent full testing, it should be addressed in future chip for two reasons:

- Latch-up connects the chip supply to ground and can draw very high currents. If chip variation in a camera based on the pixel (say 32 x 32 pixels - 1024 pixels in total) were to change the threshold where the latch-up occurred such that 1% of the pixels failed regardless of the pulse width, then approximately ten pixels would fail on a given chip. From the results in Chapter 4 this might draw $0.5\,A$ more current than expected causing parts of the chip to overheat and fail.

- The presence of latch-up shows that carriers are being injected into the silicon and these can be an important source of noise for analog components.

Once solution to the latch-up problem is guard rings, and it was suggested that these be placed around the separate pulse generators. This will become especially important if the pixel is scaled down.

6.2 Combating the Speckle Effect

It was shown that the speckle effect is an important concern with this system. This is the interference effect that occurs when illuminating an surface which has height variations greater than the wavelength of the light. The light reflected back is not uniform, but composed of bright and dark areas called speckles. Each pixel of the camera will only be able to process the vibration data if a bright speckle falls on the pixel. This is the same problem as ESPI has and is an unavoidable result of real-world surfaces. There are however ways to minimise the effects of the speckles.

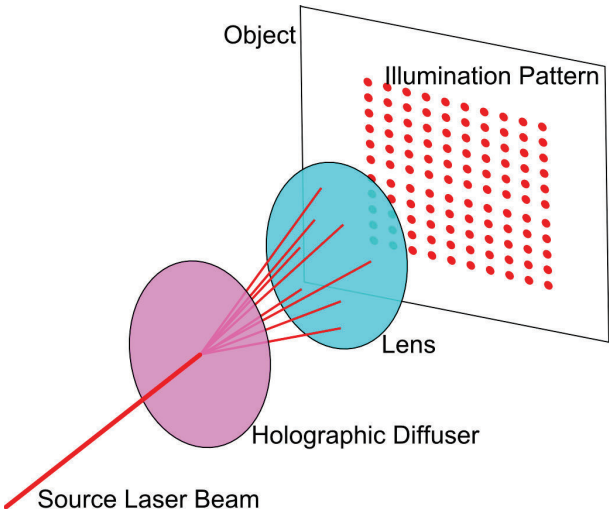

Figure 140: Use of a holographic diffuser to split a laser into a number of diverging beams, which can be focused onto the object.

There is a possible solution to this problem: illuminate the surface of the object with a grid of points, rather than a uniform illumination. This is possible using a holographic diffuser as shown in Figure 140. The use of engineered diffusers to create arbitrary beam shape has been demonstrated by Morris et. al. of RPC Photonics [132].

It has been shown that the size of a speckle in the longitudinal direction (the direction of travel of the light) is larger than the transverse direction [32]. Using a focused laser beam it can be shown that the size of the speckles is related to the size of the illumination spot. The beam was being focused by a 1" diameter lens a distance of $48\,cm$ onto an optically rough surface. Figure 141 shows a laser being focused to an small area. The speckles are $\sim 2.2\,mm$ in transverse size. If this spot is reduced in size then the speckle pattern changes to that in Figure 142 with speckles $\sim 5.8\,mm$ in transverse size. The speckles have increased in transverse size and so will also be larger in longitudinal size.

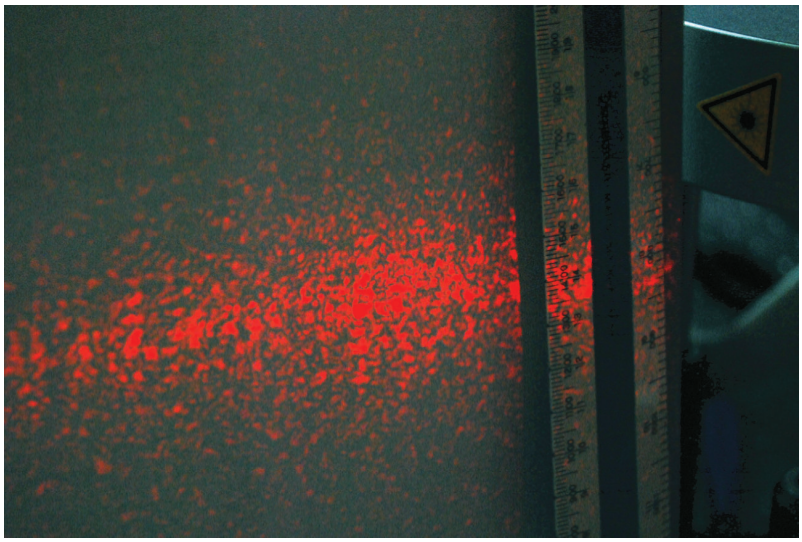

Figure 141: Speckle size for large illumination area.

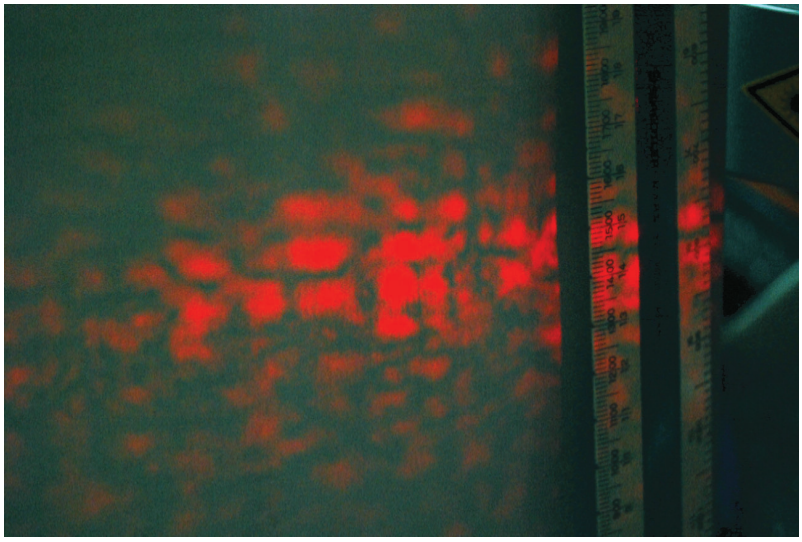

Figure 142: Speckle size for small illumination area.

Chapter 2 developed a simulation of the speckle effect. If the object is illuminated by a set of spots, rather than a uniform illumination then the speckle pattern will be superimposed on this illumination pattern as shown in Figure 143. The image is inverted due to the imaging system used in the simulation.

If just the central section of this image is considered as shown in image (a) of Figure 144 then the simulation can be rerun with a small change in displacement

of the scatterers under the central illumination spot.

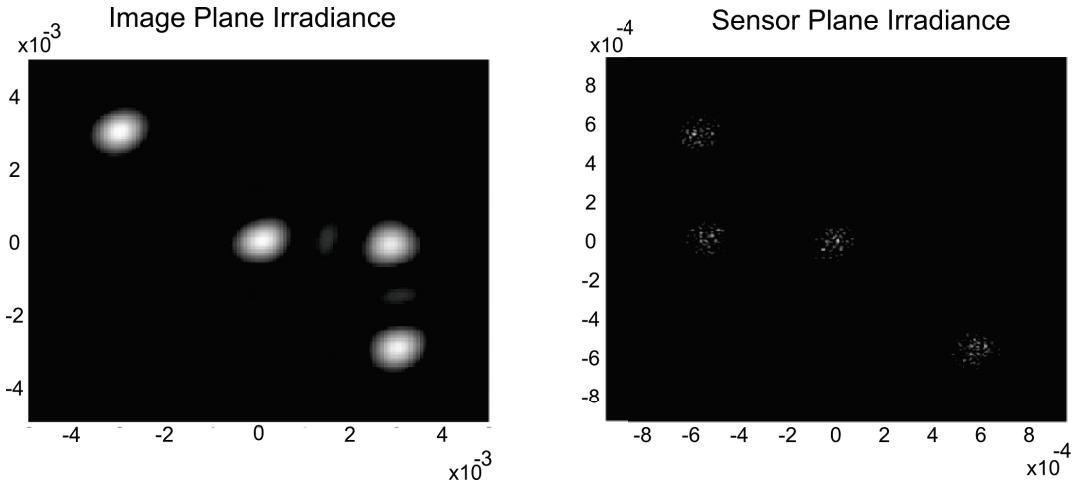

Figure 143: Simulation showing how the speckles are superimposed on an image.

The central region was moved by $\lambda/8$ while keeping the rest of the object the same. Figure 144 (a) shows that the speckle pattern irradiance is unchanged. This is expected as the movement is not enough to decorrelated the speckles, see Section 1.5.6 of Chapter 1. As the basic speckle pattern has not changed the pixel under this region of the pattern will be able to process the vibration data in the light. There is however a change in the real and imaginary components of the speckle field as shown by (b) and (c) of Figure 144.

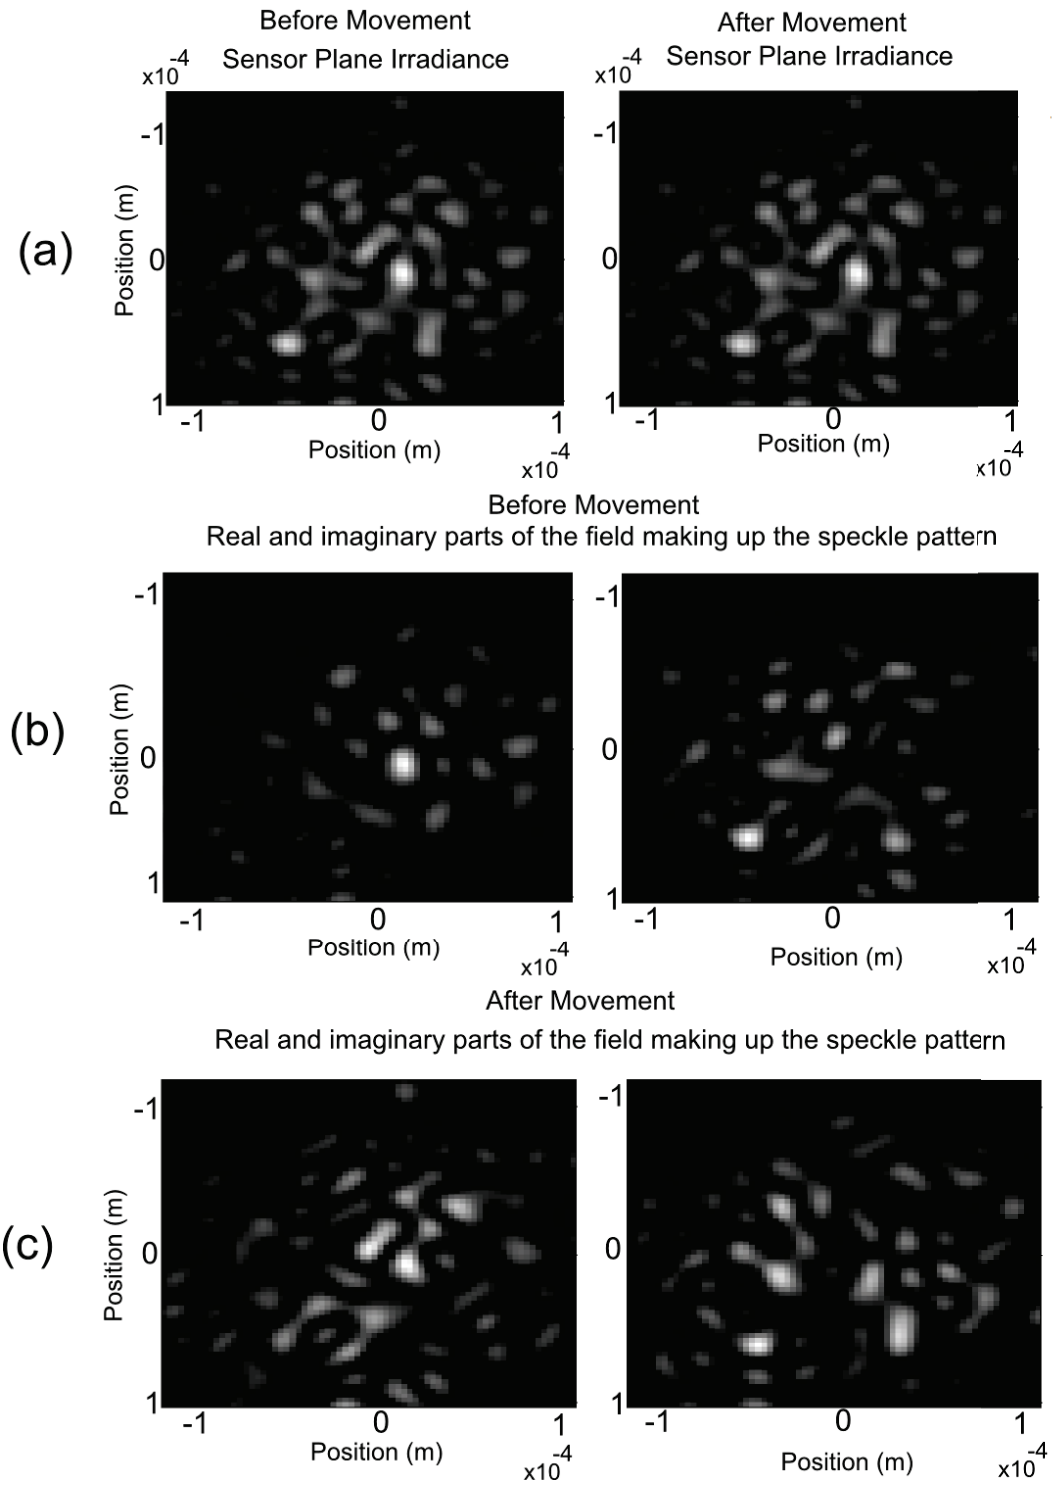

Figure 144: Simulation of a single region of the image. (a) shows the speckle irradiance before and after movement. These are the same as the movement is small enough to preserve each of the speckles. (b) shows the real and imaginary components of the speckle field before movement. (c) shows the real and imaginary values after movement.

The real and imaginary parts change as the phase of the light from the central region of the object has changed. The light is described in a complex form such as Equation 56, where E_0 is the electric field amplitude, ω is the angular frequency, k is wavenumber ($k = 2\pi/\lambda$), r is the displacement and t is time. As the irradiance is calculated by averaging the square of the field variations over a long time compared with the frequency, the irradiance is controlled by the $e^{i\{k.r\}}$ term.

$$E(t) = E_0.e^{i\{\omega.t\}}e^{i\{k.r\}} \tag{56}$$

The change in distance of the central section effectively multiplies the speckle field by $e^{i\{k.D\}}$ where D is the distance moved. As e^{ϕ} can be written as $cos(\phi) + i.sin(\phi)$, ie. a real part and an imaginary part, these will both change when the object surface moves, but the overall magnitude of the light will stay the same - the irradiance pattern will remain the same.

Although qualitative, the speckle size vs. illumination size and speckle simulation, suggest that a grid of illuminations would be a partial solution to the speckle problem.

6.3 Conclusion

This thesis has shown the design, construction and testing of a proposed pixel for a full-field LDV camera. The work has shown how the processing for an LDV can be scaled down to fit into a small area on a CMOS chip, while still demodulating the Doppler signal from an interferometer to give vibration information.

The chip processing follows a pipeline methodology using discrete components for each stage. The main stage of the processing implemented a zero-crossing demodulator which already occupies only $50\,\mu m$ x $30\,\mu m$ and as much of it is digital it should be possible to scale it down. The non-linearity of the processing

method was shown to give better dynamic range allowing both small and large vibrations to be detected. There was a relatively large variation in the timing conversion constant for a given pixel due to manufacturing variation and this would require careful calibration of a camera based on this pixel, but future work on minimising this variation would be worthwhile. It was suggested that a larger timing capacitor would solve the problem.

The chip was tested in several stages and it was found that the processing operated as expected, and compared well with a commercial LDV. A problem was found with the pulse generator used to provide control signals to the processing part of the chip, and latch-up was suggested as a probably reason for the problem. Future work should provide gaurd rings to prevent this problem from affecting future pixels.

Finally the chip was tested with a real-world vibrating object, a membrane, to demonstrate an application where full-field, non-contact, vibration measurement is an advantage. This project has demonstrated that it is feasible to design a pixel for a full-field LDV camera, and this would be a worthwhile area for further work.

Bibliography

References

[1] S. Terasaki, N. Wada, N. Sakurai, N. Muramatsu, R. Yamamoto, and D. J. Nevins, "Nondestructive measurement of kiwifruit ripeness using a laser doppler vibrometer," *Transactions of the Asae*, vol. 44, pp. 81–87, Jan. 2001.

[2] P. J. Shull, *Nondestructive Evaluation*. No. 1, Marcel Dekker, 1st ed., 2002.

[3] F. A. Iddings and P. J. Shull, *Nondestructive Evaluation*. No. 2, Marcel Dekker, 1st ed., 2002.

[4] P. E. Mix, *Introduction to Nondestructive Testing*. No. 6, New Jersey: Wiley, 2nd ed., 2005.

[5] A. Lindgren, P. J. Shull, K. Joseph, and D. Hagemaier, *Nondestructive Evaluation*. No. 4, Marcel Dekker, 1st ed., 2002.

[6] P. E. Mix, *Introduction to Nondestructive Testing*. No. 7, New Jersey: Wiley, 2nd ed., 2005.

[7] P. E. Mix, *Introduction to Nondestructive Testing*. No. 2, New Jersey: Wiley, 2nd ed., 2005.

[8] P. E. Mix, *Introduction to Nondestructive Testing*. No. 11, New Jersey: Wiley, 2nd ed., 2005.

[9] P. J. Shull and B. R. Tittmann, *Nondestructive Evaluation*. No. 3, Marcel Dekker, 1st ed., 2002.

[10] M. Bhardwaj and A. Bhardwarj, "High efficiency non-contact transducers and a very high coupling piezoelectric composite," in *16th World Conference on NDT*, 2004.

[11] I. H. Liu and C.-H. Yang, "A novel procedure employing laser ultrasound technique and simplex algorism for the characterization of mechanical and geometrical properties in zircaloy tubes with different levels of hydrogen charging," *Journal of Nuclear Materials*, vol. 408, pp. 96–101, 2011.

[12] S. Dixon, S. E. Burrows, B. Dutton, and Y. Fan, "Detection of cracks in metal sheets using pulsed laser generated ultrasound and emat detection," *Ultrasonics*, vol. 51, pp. 7–16, 2011.

[13] K. T. Wu, C. K. Jen, M. Kobayashi, and A. Blouin, "Integrated piezoelectric ultrasonic receivers for laser ultrasound in non-destructive testing of metals," *Journal of Nondestructive Evaluation*, vol. 30, pp. 1–8, 2010.

[14] R. B. Randall, *Vibration-based Condition Monitoring.* Wiley, 2011.

[15] S. Kalicinski, T. Bieniek, P. Janus, and P. Grabiec, "Determination of electrical and mechanical parameters in capacitive mems accelerometers using electrical measurements," *Microelectronics Reliability*, vol. 51, no. 7, pp. 1192–, 2011.

[16] C. H. M. Jenkins and U. A. Korde, "Membrane vibration experiments: An historical review and recent results," *Journal of Sound and Vibration*, vol. 295, pp. 602–613, 2006.

[17] K. Tatar and P. Gren, "Measurement of milling tool vibrations during cutting using a laser vibrometer," *Machine Tools & Manufacture*, vol. 48, pp. 380–387, 2008.

[18] P. E. Mix, *Introduction to Nondestructive Testing*. No. 3, New Jersey: Wiley, 2nd ed., 2005.

[19] M. P. Papaelias, C. Roberts, and C. L. Davis, "A review of non-destructive evaluation of rails: state-of-the-art and future development," *Proc. IMechE*, vol. 222, pp. 367–384, 2008.

[20] J. M. Spicer and R. Osiander, *Nondestructive Evaluation*. No. 8, Marcel Dekker, 1st ed., 2002.

[21] Y. Y. Hung, Y. S. Chen, S. P. Ng, L. Lui, Y. H. Huang, B. L. Luk, C. M. L. Wu, and P. S. Chung, "Review and comparison of shearography and active thermography for nondestructive evaluation," *Materials Science and Engineering Reports*, vol. 64, pp. 73–112, 2008.

[22] A. J. Bahr, R. Zoughi, and N. Quaddoumi, *Nondestructive Evaluation*. No. 9, Marcel Dekker, 1st ed., 2002.

[23] H. E. Martz, C. M. Logan, and P. J. Schull, *Nondestructive Evaluation*. No. 7, Marcel Dekker, 1st ed., 2002.

[24] C. Zang, G. Chen, D. J. Ewins, and I. A. Sever, "Review of current status of full-field measurement equipment hardware and software, test planning and integration tools, suggestions for future development," tech. rep., VIVACE, 2005.

[25] C. Shin, J. Chung, and W. Kim, "Dynamic charactersitics of the out-of plane vibration for an axially moving membrane," *Journal of Sound and Vibration*, vol. 286, pp. 1019–1031, 2005.

[26] C. H. Jenkins, V. D. Kalanovic, K. Padmanabhan, and S. M. Faisal, "Intelligent shape control for precision membrane antennae and reflectors in space," *Smart Mater. Struct.*, vol. 8, pp. 857–867, 1999.

[27] T. P. Gill, *The Doppler Effect*. Logos Press LTD, 1965.

[28] G. L. Cloud, *Optical Methods of Engineering Analysis*. Cambridge University Press, 1st ed., 1995.

[29] M. Francon, *Laser Speckle and Applications in Optics*. Academic Press, 1979.

[30] G. L. Cloud and P. H. Franco, *Optical Methods of Engineering Analysis*. Cambridge University Press, 1st ed., 1995.

[31] R. Jones and C. Wykes, *Holographic and Speckle Interferometry*. Cambridge University Press, 1989.

[32] L. Leushacke and M. Kirchner, "Three-dimensional correlation coefficient of speckle intensity for rectangular and circular apertures," *Journal of the Optical Society of America*, vol. 7, p. 827, 1990.

[33] D. Post, B. Han, and P. Ifju, *High Sensitivity Moire*. Springer-Verlag, 1994.

[34] J. McKelvie, "MoireŽ strain analysis: an introduction, review and critique, including related techniques and future potential," *The Journal of Strain Analysis for Engineering Design*, vol. 33, pp. 137–151, 1998.

[35] H. Liu, C. Basaran, A. N. Cartwright, and W. Casey, "Application of moiré interferometry to determine strain fields and debonding of solder joints in bga packages heng," *IEEE Transactionas on Components and Packaging Technologies*, vol. 27, pp. 217–223, 2004.

[36] A. Asundi, *Trends in Optical Non-Destructive Testing and Inspection*. Elsevier, 2000.

[37] G. S. Spagnolo, D. Ambrosini, and D. Paoletti, "Measurement of vibration amplitude by an optical fibre-based moire interferometer," *Optics And Lasers In Engineering*, vol. 30, no. 2, pp. 213–223, 1998.

[38] K. G. Harding and J. S. Harris, "Projection moire interferometer for vibration analysis," *Applied Optics*, vol. 22, no. 6, pp. 856–861, 1983.

[39] T. M. Kreis, *Trends in Optical Non-Destructive Testing and Inspection*. Elsevier, 2000.

[40] U. Schnars, "Direct phase determination in hologram interferometry with use of digitally recorded holograms," *Journal of the Optical Society of America*, vol. 11, pp. 2011–2015, 1994.

[41] K. A. Stetson, "Review of speckle photography and interferometry," *Optical Engineering*, vol. 14, no. 5, pp. 482–489, 1975.

[42] T. R. Moore and J. J. Skubal, "Time-averaged electronic speckle pattern interferometry in the presence of ambient motion," *Applied Optics*, vol. 47, pp. 4640–4648, 2008.

[43] H. Van der Auweraer, H. Steinbichler, S. Vanlanduit, C. Haberstok, R. Freymann, D. Storer, and V. Linet, "Application of stroboscopic and pulsed-laser electronic speckle pattern interferometry (espi) to modal analysis problems," *Measurement Science and Technology*, vol. 13, pp. 451–463, 2002.

[44] A. Davila, D. Kerra, and G. Kaufmann, "Fast electro-optical system for pulsedespi carrier fringe generation," *Optics Communications*, vol. 123, pp. 457–464, 1996.

[45] R. Rabinowitz, S. Jacobs, R. Targ, and G. Gould, "Homodyne detection of phase modulated light," *Proceedings of the IRE*, vol. 50, no. 11, p. 2365, 1962.

[46] H. Cummins, N. Knable, L. Gampel, and Y. Yeh, "Frequency shifts in light diffracted by ultrasonic waves in liquid media," *Applied Physics Letters*, vol. 2, no. 3, pp. 62–64, 1963.

[47] E. Hecht, *Optics*. Addison Wesley, fourth edition ed., 2002.

[48] H. Z. Cummins and N. Knable, "Single sideband modulation of coherent light by bragg reflection from acoustical waves," *Proceedings of the IEEE*, vol. 51, no. 9, p. 1246, 1963.

[49] W. H. Stevenson, "Optical frequency shifting by means of a rotating diffraction grating," *Applied Optics*, vol. 9, no. 3, pp. 649–652, 1970.

[50] H. Z. Cummins, N. Knable, and Y. Yeh, "Observation of diffusion broadening of rayleigh scattered light," *Physical Review Letters*, vol. 12, no. 6, pp. 150–153, 1964.

[51] Y. Yeh and H. Z. Cummins, "Localized fluid flow measurements with an he-ne laser spectrometer," *Applied Physics Letters*, vol. 4, no. 10, pp. 176–178, 1964.

[52] J. W. Foreman, E. W. George, and R. D. Lewis, "Measurement of localized flow velocities in gases with a laser doppler flowmeter," *Applied Physics Letters*, vol. 7, no. 4, pp. 77–18, 1965.

[53] J. W. Foreman, J. E. W. George, J. L. Jetton, R. D. Lewis, J. R. Thornton, and H. J. Watson, "Fluid flow measurements with a laser doppler velocime-

ter," *IEEE Journal of Quantum Electronics*, vol. QE-2, no. 8, pp. 260–266, 1966.

[54] M. J. Rudd, "A new theoretical model for the laser dopplermeter," *Journal of Scientific Instruments*, vol. 2, no. 2, pp. 55–58, 1969.

[55] C. Greated, "Measurement of turbulence statistics with a laser velocimeter," *Journal of E: Scientific Instruments*, vol. 3, pp. 158–160, 1970.

[56] J. E. Rizzo, "A laser doppler interferometer," *Journal of Physics E: Scientific Instruments*, vol. 8, pp. 47–52, 1975.

[57] T. H. Wilmhurst and J. E. Rizzo, "An autodyne frequency tracker for laser doppler anemometry," *Journal of Physics E: Scientific Instruments*, vol. 7, pp. 924–930, 1974.

[58] D. B. Brayton, H. T. Kalb, and F. L. Crosswy, "Two-component dual-scatter laser doppler velocimeter with frequency burst signal readout," *Applied Optics*, vol. 12, no. 6, pp. 1145–1156, 1973.

[59] D. M. Robinson and P. W. Chu, "Diffraction analysis of doppler signal characteristics for a cross-beam laser doppler velocimeter," *Applied Optics*, vol. 14, no. 9, pp. 2177–2183, 1975.

[60] J. F. Meyers, J. W. Lee, and A. A. Cavone, "Signal processing schemes for doppler global velocimetry," in *International Congress On Instrumentation In Aerospace Simulation Facilities*, 1991.

[61] R. S. ILYASOV, Y. N. MALININ, and Y. E. POLSKII, "Laser vibrometer with increased accuracy and interference immunity," *Measurement Techniques*, vol. 20, no. 10, pp. 1470–1472, 1977.

[62] T. J. F. Buunen and S. M. G. Vlaming, "Laser-doppler velocity meter aplied to tympanic membrane vibrations in cat," *Acoustical Society of America*, vol. 69, no. 3, pp. 744–750, 1981.

[63] R. Dandliker and J. F. Willemin, "Measuring microvibrations by heterodyne speckle interferometry," *Optics Letters*, vol. 6, no. 4, p. 165, 1981.

[64] J. F. Willemin and R. Dandiker, "Measuring amplitude and plase of microvibrations by heterodyne speckle interferometry," *Optics Letters*, vol. 8, pp. 102–104, 1983.

[65] N. A. Halliwell, C. J. D. Pickering, and P. G. Eastwood, "The laser torsional vibrometer: A new instrument," *Journal of Sound and Vibration*, vol. 93, no. 4, pp. 588–592, 1984.

[66] N. A. Halliwell and P. G. Eastwood, "The laser torsional vibrometer," *Journal of Sound and Vibration*, vol. 101, no. 3, pp. 446–449, 1985.

[67] R. MEYNART, "A simple interferometric vibrometer based on a laser doppler-velocimeter," *Journal of Physics E-scientific Instruments*, vol. 17, no. 11, pp. 984–987, 1984.

[68] F. Durst and J. H. Whitlaw, "Integrated optical units for laser anemometry," *Journal of Physics E-scientific Instruments*, vol. 4, p. 804, 1971.

[69] P. G. Suchoski, J. P. Walters, and M. R. Fernald, "Miniature laser veibrometer system with multifunctional integrated optic circuit," *IEEE Photonics Technology Letters*, vol. 2, no. 1, p. 81, 1990.

[70] E. H. Bokelberg, H. J. Sommer, and M. W. Trethewey, "A six-degree-of-freedom laser vibrometer, part i: Theoretical development," *Journal of Sound and Vibration*, vol. 178, no. 5, pp. 643–654, 1994.

[71] E. H. Bokelberg, H. J. Sommer, and M. W. Trethewey, "A six-degree-of-freedom laser vibrometer, part ii: Experimental validation," *Journal of Sound and Vibration*, vol. 178, no. 5, pp. 655–667, 1994.

[72] A. B. Stanbridge and D. J. Ewins, "Modal testing using a scanning laser doppler vibrometer," *Mechanical Systems and Signal Processing*, vol. 13, no. 2, pp. 255–270, 1999.

[73] M. Bauer, F. Ritter, and G. Siegmund, "High-precision laser vibrometers based on digital doppler-signal processing," in *Fifth International Conference on Vibration Measurements by Laser Techniques*, vol. 4827, Proceedings of SPIE, 2002.

[74] P. Castellini and R. Montanini, "Automotive components vibration measurements by tracking laser doppler vibrometry: Advances in signal processing," *Measurement Science and Technology*, vol. 13, pp. 1266–1279, 2002.

[75] A. Serov, B. Steinacher, and T. Lasser, "Full-field laser doppler perfusion imaging and monitoring with an intelligent cmos camera," *Optics Express*, vol. 13, no. 10, pp. 3681–3689, 2005.

[76] R. J. Prazenica, A. J. Kurdila, and J. F. Vignola, "Spatial filtering and proper orthogonal decomposition of scanning laser doppler vibrometry data for the nondestructive evaluation of frescoes," *Journal of Sound and Vibration*, vol. 304, pp. 735–751, 2007.

[77] J. A. Bucaro, A. J. Romano, and P. Abraham, "Detection and localizatio of inclusions in plates using inversion of point actuated surface displacements," *Journal of the Acoustical Society of America*, vol. 115, no. 1, pp. 201–206, 2004.

[78] A. J. Romano, J. A. Bucaro, J. F. Vignola, and P. Abraham, "Detection and localization of rib detechment in thin metal and composite polates by inversion of laser doppler vibrometry scans," *Journal of the Acoustical Society of America*, vol. 121, no. 5, p. 2667, 2007.

[79] M. V. Aguanno, M. J. Connelly, and M. P. Whelan, "Full field laser vibrometry employing a novel cmos-dsp camera," in *Proceedings Of The Society Of Photo-optical Instrumentation Engineers (spie)*, vol. 4827 of *Fifth International Conference On Vibration Measurements By Laser Techniques: Advances And Applications*, (Ancona, Italy, Jun 18-21), 2002.

[80] M. V. Aguanno, F. Lakestani, M. P. Whelan, and M. J. Connelly, "Single pixel carrier based approach for full field laser interferometry using a cmos-dsp camera," *Detectors and Associated Signal Processing*, vol. 5251, pp. 304–312, 2004.

[81] M. V. Aguanno, F. Lakestani, M. P. Whelan, and M. J. Connelly, "Heterodyne speckle interferometer for full-field velocity profile measurements of a vibrating membrane by electronic scanning," *Optics and Lasers in Engineering*, vol. 45, pp. 677–683, 2007.

[82] R. Di Sante and L. Scalise, "A novel fiber optic sensor for multiple and simultaneous measurement of vibration velocity," *Review of Scientific Instruments*, vol. 75, pp. 1952–1958, 2004.

[83] E. Li, J. Xi, J. Chicharo, J. Yao, and D. Yu, "Multi-point laser doppler velocimeter," *Optics Communications*, vol. 245, pp. 309–313, 2005.

[84] J. M. Kilpatrick and V. Markov, "Matrix laser vibrometer for transient modal imaging and rapid nondestructive testing," in *Proceedings of the*

SPIE, vol. 7098 of *Eighth International Conference On Vibration Measurements By Laser Techniques: Advances And Applications*, (Ancona, Italy, Jun 18-20), 2008.

[85] J. M. Kilpatrick and V. B. Markov, "Full-field laser vibrometer for instantaneous vibration measurement and non-destructive inspection," *Measurement Technology and Intelligent Instruments*, vol. 437, pp. 407–411, 2010.

[86] J. D. Briers, "Laser doppler and time-varying speckle: A reconciliation," *Journal of the Optical Society of America*, vol. 13, no. 2, p. 345, 1996.

[87] W. H. Steel, *Interferometry*. Cambridge University Press, second ed., 1983.

[88] E. Hecht, *Optics*. Addison Wesley, fourth edition ed., 2002.

[89] P. Castellini, M. Martarelli, and E. P. Tomasini, "Laser doppler vibrometry: Development of advanced solutions answering to technology's needs," *Mechanical Systems and Signal Processing*, vol. 20, pp. 1265–1285, 2006.

[90] W. Mansor, *Fetal systolic timing intervals using Doppler ultrasound*. PhD thesis, The University of Nottingham, 2005.

[91] B. M. Watrasiewicz and M. J. Rudd, *Laser Doppler Measurements*. Butterworths, 1976.

[92] H. Stark, F. B. Tuteur, and J. B. Anderson, *Modern Electrical Communications*. Prentice Hall, second ed., 1979.

[93] E. Hecht, *Optics*. Addison Wesley, fourth edition ed., 2002.

[94] G. W. Stroke, *An Introduction to Coherent Optics and Holography*. New York: Academic Press, 1969.

[95] D. He, *Full Field Laser Doppler Blood Flow Sensor*. PhD thesis, The University of Nottingham, 2007.

[96] D. He, C. Kongsavatsak, B. Hayes-Gill, J. Crowe, and S. Morgan, "32x32 pixel array complementary metal-oxide semiconductor imaging sensor for laser doppler blood-flow measurement," *Optical Engineering*, vol. 50, no. 5, 2011.

[97] C. Kongsavatsak, D. He, B. Hayes-Gill, J. Crowe, and S. Morgan, "Complementary metal-oxide-semiconductor imaging array with laser doppler blood flow processing," *Optical Engineering*, vol. 47, no. 10, 2008.

[98] H. H. Nassif, M. Gindy, and J. Davis, "Comparison of laser doppler vibrometer with contact sensor for monitoring bridge deflection and vibration," *NDT & E International*, vol. 38, pp. 213–218, 2005.

[99] J. W. Goodman, *Introduction to Fourier Optics*. McGraw-Hill Book Company, Inc, second ed., 1996.

[100] J. W. Goodman, *Speckle Phenomena in Optics*. Roberts and Company, 2007.

[101] R. R. Bracewell, *The Fourier Transform and Its Applications*. McGraw-Hill Book Company, Inc, 1986.

[102] G. Pasternack and R. L. Whalin, "Analysis and synthesis of a digital phase-locked loop for fm demodulation," *The Bell System Technical Journal*, vol. 47, no. 10, pp. 2207–2237, 1968.

[103] W. Hoeft and I. Zingmond, "A monolithic if subsystem for demosulation of fm signals," *IEEE Transactions on Consumer Electronics*, vol. 21, no. 4, pp. 348–355, 1975.

[104] L. I. Suckle, "Fm stereo multiplex demodulation with a non-symmetrical waveform," *IEEE Transactiona on Broadcast and Television Receivers*, vol. BT18, pp. 43–48, 1972.

[105] W. Kiranon and P. Wardkein, "A novel fm demodulation scheme," *IEEE transactions on consumer electronics*, vol. 41, no. 4, pp. 1103–1107, 1995.

[106] J. Epps, E. Ambikairajah, and T. Thiruvaran, "Robust fm demodulation of discrete-time signals using least squares differential ratio," *Electronic Letters*, vol. 43, no. 13, pp. 727–729, 2007.

[107] S. L. Reddy, B. Santhanam, and M. Hayat, "Cochannel fm demodulation via the multi angle-centered discrete fractional fourier transform," in *IEEE 13th digital signal processing workshop & 5th ieee processing education workshop, vols 1 and 2, proceedings*, pp. 535–539, 2009.

[108] H. B. Voelcker, "Zero-crossing properties of angle-modulated signals," *IEEE Transactions on Communications*, vol. COM-20, no. 3, pp. 307–315, 1972.

[109] R. G. Wiley, H. Schwarzlander, and D. D. Weiner, "Demodulation procedure for very wide-band fm," *IEEE Transactions on Communications*, vol. COM-25, no. 3, pp. 318–327, 1977.

[110] S. K. Ray, "A new method for the demodulation fo fm signals," *IEEE Transactions on Communications*, vol. COM-28, pp. 142–144, 1980.

[111] R. G. Wiley, "Approximate fm demodulation using zero crossings," *IEEE Transactions on Communications*, vol. COM-29, no. 7, pp. 1061–1065, 1981.

[112] S. K. Ray, "Zero-crossing-based approximate demodulation of wide-deviation fm," *IEE Proceedings-F Radar and Signal Processing*, vol. 131, no. 1, pp. 47–51, 1984.

[113] A. Tselikov, J. U. de Arruda, and J. Blake, "Zero-crossing demodulation for open-loop sagnca interferometers," *Journal of Lightwave Technology*, vol. 16, no. 9, p. 1613, 1998.

[114] T. Schotland and P. Jung, "Bluetooth receiver based on zero-crossing demodulation," *Electronics Letters*, vol. 39, p. 397, 2003.

[115] T. Schotland and P. Jung, "Bluetooth receiver with zero-crossing zrero-forcing demodulation," *Electronics Letters*, vol. 39, p. 1275, 2003.

[116] T. Schotland, C. Spiegel, A. Waadt, A. Burnic, and P. Jung, "Applying zero-crossing demodulation for the bluetooth enhanced data rate mode," in *Proceedings of the Sixth IASTED International Multi-Conference on Wireless and Optical Communications*, pp. 20–26, 2006.

[117] L. Zhang, S. Wang, T. P. Deng, and Z. Qin, "A novel demodulation method based on zero-crossing detection for uwb signal reception," in *5th international conference on wireless communications, networking and mobile computing*, pp. 1730–1733, 2009.

[118] R. R. Spencer and M. S. Ghausi, *Introduction to Electronic Circuit Deign*. Prentice Hall, 2003.

[119] S. M. Sze, *Physics of Semiconductor Devices*. New York: John Wiley & Sons, 1981.

[120] B. G. Streetman, *Solid State Electronic Devices*. Englewood Cliffs, N. J.: Prentice Hall, 1995.

[121] R. A. Light, *Design and Implementation of an Integrating Modulated Light Camera*. PhD thesis, University of Nottingham, 2008.

[122] R. Jacob Baker, *CMOS Circuit Design, Layout, and Simulation*. New York: IEEE Press, 1998.

[123] J. QUINCKE, "Novel test structures for the investigation of the efficiency of guard rings used for i/o-latch-up prevention," in *Proceedings of the 1990 International Conference On Microelectronics Test Structures*, 1990 International Conf On Microelectronic Test Structures, (San Diego), pp. 35–39, Mar 05-07, 1990.

[124] F. Farbiz and E. Rosenbaum, "Understanding transient latchup hazards and the impact of guard rings," *2010 International Reliability Physics Symposium*, pp. 466–473, 2010.

[125] S. Gupta, J. C. Beckman, and S. L. Kosier, "Improved latch-up immunity in junction-isolated smart power ics with unbiased guard ring," *Ieee Electron Device Letters*, vol. 22, pp. 600–602, Dec. 2001.

[126] J. B. Baker, *CMOS: Circuit Design, Layout, and Simulation*. Wiley-IEEE Press, 2010.

[127] T. D. Rossing and N. H. Fletcher, *Principles of Vibration and Sound*. Springer, second ed., 2004.

[128] C. Shin, J. Chung, and W. Kim, "Dynamic characteristics of the out-of-plane vibration for an axially moving membrane," *Journal of Sound and Vibration*, vol. 286, pp. 1019–1031, 2005.

[129] Y. Y. Hu, S. Yoshida, S. Nakamura, K. Watanabe, W. Z. Lin, E. T. Ong, and J. Q. Mou, "Analysis of built-in speaker-induced structural-acoustic vibration of hard disk drives in notebook pcs," *IEEE Transactions on Magnetics*, vol. 45, pp. 4950–4955, 2009.

[130] G. M. Revel, M. Martarelli, and P. Chiariotte, "A new laser vibrometry-based 2d selective intensity method for source identification in reverberant fields: part i. development of the technique and preliminary validation," *Measurement Science and Technology*, vol. 21, pp. 1–8, 2010.

[131] G. M. Revel, M. Martarelli, and P. Chiariotte, "A new laser vibrometry-based 2d selective intensity method for source identification in reverberant fields: part ii. application to an aircraft cabin," *Measurement Science and Technology*, vol. 21, pp. 1–10, 2010.

[132] G. M. Morris, T. R. M. Sales, S. Chakmakjian, and D. J. Schertler, "Engineered diffusers for display and illumination systems: Design, fabrication, and applications," tech. rep., RPC Photonics, Inc., 330 Clay Road, Rochester, NY 14623.

Appendix A - MI Alignment

Of practical note is the method of aligning the interferometer. There are three issues which need to be considered. The first is the coherence length of the laser. The two arms of the interferometer should not vary by more than the coherence length and this is set by the laser used. In the case of the He-Ne laser used here, its coherence length is about $20\,cm$. The next thing to avoid is long light paths, whether going from the laser to the beam splitter or from there to the mirrors. Any changes in the refractive index along the beam will introduce noise. Given that air's refractive index changes with temperature, any changes in temperature along the beam will affect it. The final point is that the reference and signal beams should be as co-axial as possible when they reach the sensor, otherwise the interference effects will be too small to detect (spatially).

- Working with an optics bench, the holes provide a reference grid. First setup the laser. Run the beam parallel with the holes in the optics bench and parallel with its surface. The lateral alignment is not as important as the vertical alignment as moving optical components around the bench, and rotating them is much easier than changing their height. Once the beam is pretty much parallel with the grid (lateral), set an iris to about the same size as the beam and align it in front of the laser - this forms a height marker. Now move the iris along the beam and check that the vertical height is well aligned. This iris forms a reference for the height of any beam in the experiment.

- The laser now forms the alignment reference for the experiment. All beams with either be parallel with this beam or perpendicular to it. Place the beam splitter as close to the laser as possible, taking into account the other optical components. There are always some losses from a beam splitter and

these have their uses. It should be possible to see where the laser beam enters the beam splitter as red disks on the air-glass interface. It should also be possible to see a disk on the internal splitting surface. Use these to place the beam splitter such that the laser disks are in the center of the three cube sides (entry, pass-through and reflected), as well as the center of the splitting interference. This should leave the reflected beam relatively perpendicular to the laser source beam.

- Place the object mirror such that the spot on the beam splitter from the returning beam coincides with the source beam. This now sets the mirror normal to the reference beam, and given that the beam will have traveled d_{sig} twice, and we can align the disks relatively accurately by eye (the laser used for the experiment was a He-Ne with a $1\,mm$ diameter beam - it should be possible to align this to better than half a millimeter), this mirror should be pretty well aligned.

- If the beam splitter is slightly rotated the beam reflected by it will travel at an angle to the perpendicular, see Figure 145. Rotation of the beam splitter affects the passed beam much less than the reflected beam so the signal arm already aligned should be fine. If the reference arm mirror can be aligned such that the outward and return beam disks are aligned well then this rotation simply rotates a whole section of the interferometer with no major problems. The second round of splitting will however show you that the beam splitter is misaligned.

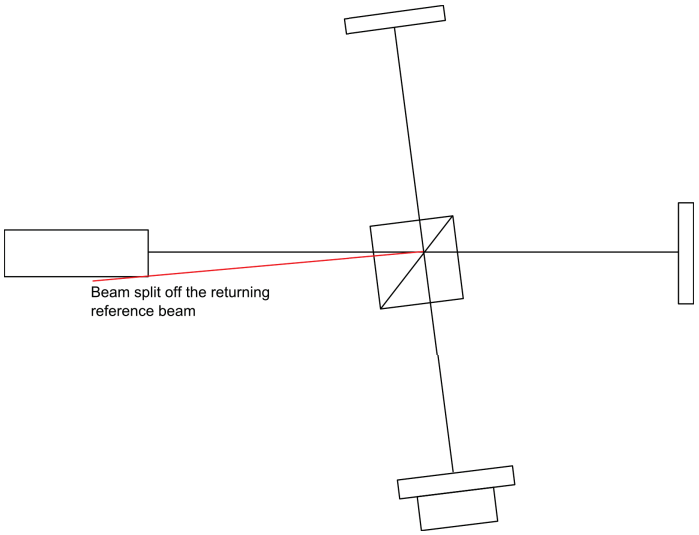

Figure 145: Misalignment of a Michelson Interferometer.

Appendix B - MZI Alignment

Alignment of the MZI is more difficult than the MI due to the beams traveling over completely different paths. As with the MZI the laser should be aligned level, however the alignment was simplified through the use of an optical cage system. This is a system consisting of metal rods which can be used to build a rigid frame to hold lenses and help align the beam splitters. The procedure for alignment was therefore as below:

- Once the laser is level and fairly closely aligned to the optics bench, mark where the beam reaches a fixed surface - such as a wall. The main consideration for interference is that the beams reach the sensor co-axially (and having covered a similar distance). Therefore the interferometer needs to be aligned to a fixed coordinate system. The original laser beam is one axis of this system (from here on called the main axis). The first beam splitter can now be placed and its reflected beam forms the other axis (from here on called the reflected axis). Placing a mirror at the other end of this split

off beam allows the angle of the reflection to be set perpendicular. This is like aligning an MI.

- The arm which splits off the first beam splitter and the mirror was fixed in an optics cage as well as the parallel section with the two other beam splitters. These must be kept these the same length. The one around the first beam splitter and mirror was mounted and using alignment plates it was set to match the reflected axis. The alignment plates were black plated which could be dropped into the cage system and had a $1\,mm$ hole in them as well as concentric circles to aid alignment.

- The beam from the mirror was then checked for height by comparing it with the main beam. Long sections of cage system were then used to make sure that the distance from the mirror to beam splitter 3 was the same as the distance to beam splitter 2. The second section of cage system was also mounted and now the beam from the mirror needed to be aligned to the beam from beam splitter 2, arriving at beam splitter 3.

- This is partly trial and error but marking beam positions on the walls around the optics bench does help. If beam splitter 1 is left unchanged it forms the reflected axis. Beam splitters 2 and 3 could then be leveled to match this system.