Development of a Compact, Configurable, Real-time Range Imaging System

by

Adrian Peter Paul Jongenelen

A thesis submitted to the Victoria University of Wellington in fulfilment of the requirements for the degree of Doctor of Philosophy in

Electronics and Computer Systems Engineering.



2011

© 2011 Adrian Peter Paul Jongenelen

Abstract

This thesis documents the development of a time-of-flight (ToF) camera suitable for autonomous mobile robotics applications. By measuring the round trip time of emitted light to and from objects in the scene, the system is capable of simultaneous full-field range imaging. This is achieved by projecting amplitude modulated continuous wave (AMCW) light onto the scene, and recording the reflection using an image sensor array with a high-speed shutter amplitude modulated at the same frequency (of the order of tens of MHz). The effect is to encode the phase delay of the reflected light as a change in pixel intensity, which is then interpreted as distance.

A full field range imaging system has been constructed based on the PMD Technologies PMD19k image sensor, where the high-speed shuttering mechanism is builtin to the integrated circuit. This produces a system that is considerably more compact and power efficient than previous iterations that employed an image intensifier to provide sensor modulation. The new system has comparable performance to commercially available systems in terms of distance measurement precision and accuracy, but is much more flexible with regards to its operating parameters. All of the operating parameters, including the image integration time, sensor modulation phase offset and modulation frequency can be changed in real-time either manually or automatically through software.

This highly configurable system serves as an excellent platform for research into novel range imaging techniques. One promising technique is the utilisation of measurements using multiple modulation frequencies in order to maximise precision over an extended operating range. Each measurement gives an independent estimate of the distance with limited range depending on the modulation frequency. These are combined to give a measurement with extended maximum range using a novel algorithm based on the New Chinese Remainder Theorem. A theoretical model for the measurement precision and accuracy of the new algorithm is presented and verified with experimental results.

All distance image processing is performed on a per-pixel basis in real-time using a Field Programmable Gate Array (FPGA). An efficient hardware implementation of the phase determination algorithm for calculating distance is investigated. The limiting resource for such an implementation is random access memory (RAM), and a detailed analysis of the trade-off between this resource and measurement precision is also presented.

Acknowledgements

This work would not have been possible without the support of many people. First of all, my primary supervisor Professor Dale Carnegie, who encouraged me to undertake this research and was especially motivating when things were at their hardest. My secondary supervisor, Dr. Adrian Dorrington, who is a fountain of knowledge. Also, Dr. Andrew Payne, who has been a real inspiration. All three of you have been immensely helpful, and have provided excellent feedback for all my work, and I cannot thank you enough.

I am extremely grateful for the financial support from the New Zealand Tertiary Education Commission in the form of a Bright Futures Top Achiever Doctoral Scholarship. I also appreciate the contributions in the early stages from the William Georgetti Scholarship administered by the New Zealand Vice-Chancellors' Committee and the Victoria PhD Scholarship from Victoria University of Wellington. As well as providing me with an excellent workspace and all the resources I needed, Victoria University also assisted at the end in the form of a PhD Submission Scholarship, which is greatly appreciated.

Thank you to all of my family who have always shown great support and faith in my work, especially my beautiful wife, Claire. Thank you for loving me and taking great care of me through these years.

To everyone else who has contributed, but whose names I have not mentioned here, thank you very much.

Table of Contents

Abst	ract	iii
Ackn	owledgements	V
Table	e of Contents	.vi
List o	of Tables	X
List (of Figures	.xi
Chap	oter 1 Introduction	1
1.1	Objectives	2
1.2	Outline	4
Chap	oter 2 Full-Field Range Imaging	7
2.1	Alternative Range Imaging Technologies	8
2.1.	1 Passive Stereo and Photogrammetry	8
2.1. 2.1.	 Active Triangulation and Structured Light Interferometry 	10 10
2.2	Direct Time-of-flight Measurement	11
2.3	Indirect Time-of-Flight	13
2.3.	1 Shuttered Light Pulse	13
2.3.	2 Range Gating Segmentation	14
2.3.	3 AMCW Phase Stepped Homodyne	16
2.3.	4 AMCW Heterodyne	20
2.4	CMOS ToF Sensors and Camera Systems	20
2.5	AMCW Measurement Precision	22
2.5.	1 Phase Measurement Precision	22
2.5.	2 Distance Measurement Precision	24
2.6	AMCW Measurement Accuracy	25
2.6.	1 Geometric Distortions	25
2.6.	2 Fixed Pattern Noise and Latency Error	26
2.6.	5 Temperature Kelated Error	20
2.0. 2.6	 EITOIS QUE LO HAIIHOIIICS	21 20
2.0.	6 Distance Ambiguity	20 29
2.7	Summary	30

Chap	oter 3 Range Imaging System Development	.31
3.1	The University of Waikato Range Imager	31
3.2	Improving the Waikato Ranger	34
3.2.	1 VGA Monitor Driver	36
3.2.	2 Camera Link Interface	37
3.3	The Victoria University Range Image System	38
3.3.	2 PMD Daughterboard	43
3.3.	3 PMD Mainboard	45
3.3.	4 VGA DAC and Ethernet PHY/MAC	48
3.3.	5 Stratix III Development Kit	48
3.4	Victoria Ranger Software	55
5.4. 3.4.	2 PC Ethernet Listener	55
3.4.	3 MATLAB Code	57
3.5	Summary	58
Char	ntar A System Characterisation	50
Chap		. 39
4.1	Laser Illumination Waveform Shape	59
4.2	Sensor Modulation Waveform Shape	61
4.3	Waveform shape versus modulation frequency	64
4.4	Noise floor versus modulation frequency	65
4.5	Pixel amplitude versus modulation frequency	66
4.6	Phase versus modulation frequency	70
4.7	Phase and Distance Precision	71
4.7.	1 Precision versus modulation frequency	71
4.7.	 Precision versus measurement Rate	72
4.8	Systematic Errors due to Harmonics	76
49	Power Consumption	
4 10	Summary	83
	Summing	
Char	oter 5 Multiple Modulation Frequency Range Imaging.	.85
5.1	Calculating distance from two phase measurements	87
5.2	Calculation based on New Chinese Remainder Theorem	89
5.3 5.3.	Effect of phase precision on distance precision1Simulated Distance Precision versus Phase Precision	 90 92
5.4	Errors in distance measurement	95
5.4.	1 Simulated Error Rate versus Phase Precision	96
5.5 5.5.	Acquisition of Dual Frequency Measurements	 98 98

5.5.2	Simultaneous Dual Frequency Modulation	
5.5.3 5.5.4	Simulated Distance Precision versus Integration Ratio	103 106
5.5.4 E (T	Simulated Error Rate versus integration Ratio	100
5.0 I	Jual Frequency Experimental Results	106 107
5.6.2	Experimental Distance Precision versus Phase Precision	107
5.6.3	Experimental Error Rate versus Phase Precision	
5.6.4	Experimental Distance Precision versus Integration Ratio	113
5.6.5	Experimental Error Rate versus Integration Ratio	115
5.7 S	Summary	116
Chart	m (Deel dime Distance Determination	117
	er o Real-time Distance Determination	/ 11 118
62 I	Inwranning the Distance Calculation	110 119
	The A commutation Distance Culculation	101
6.3	DAM Dequirements for Arbitrary N	121
632	Output Buffers	122
6.3.3	RAM Requirements for $N = 4$	125
6.4 A	Arc Tangent and Amplitude Calculation	
65 I	mnlamontation Results	128
651	Test Methodology	120 128
6.5.2	Hardware Processing Accuracy	
6.5.3	FPGA Resource Requirements	132
6.6 S	Summary	133
Chart	The Freekration and Constraint	125
Cnapte	er / Evaluation and Conclusion	135
7.1 F	Range imaging examples	
/.1.1	Static Scene	135
1.1.2	Dynamic Scene	
7.2]	Thesis Summary	
7.2.1	Development of the Victoria University range imaging syste	m 141
7.2.2	Dual-frequency Modulation and Real-time Processing	143
7.5 T		
7.4 ŀ	Publications Arising from this Thesis	146
7.5 (Conclusion	148
Refere	nces	151
Appen	dix A PCB Schematics	
A.1 F	PMD Daughterboard	159
A.2 F	PMD Mainboard 1 – Modulation and ADC	160
A.3 F	PMD Mainboard 2 – Connectors and Interface	

A.4	Illumination Board – Top	
A.5	Illumination Board – Laser Diode Module	
A.6	Ethernet Interface Board	
A.7	VGA Interface Board	
A.8	Camera Link Interface Board	166
Арр	endix B DVD Contents	167

List of Tables

Table 2.1: Selection of commercially available compact 3D Cameras......22

List of Figures

Figure 1.1: Marvin, the flagship robot of the Victoria University of Wellington Mechatronics Group
Figure 2.1: Stereoscopic distance measurement
Figure 2.2: Time-of-Flight range measurement. A pulse of light is emitted and reflected from the scene, and the time taken for the light to travel is measured
Figure 2.3: Indirect Time-of-Flight range measurement using a high-speed shutter at the receiver. Only a portion of the returned light is passed through the shutter, encoding range as amplitude, i.e., pixel intensity
Figure 2.4: Timing diagram of shutter and reflected light where the shutter signal is in phase with emitted light. For the first frame, nearer objects return a higher intensity than farther objects. The second frame collects light from the whole pulse, and is used to normalise the measurement
Figure 2.5: Timing diagram of Range Gating Segmentation method. <i>N</i> images are captured with stepped shutter offset. Range is determined by the location of the peak intensity in the image sequence
Figure 2.6: AMCW range imaging
Figure 2.7: Timing diagram for phase stepped AMCW homodyne range imaging. The reflected modulated light is multiplied by the shutter signal and integrated by the camera. The phase of the shutter signal is stepped between frames, encoding range into the phase of the sequence of intensity images
Figure 2.8: Fitting a sinusoid to the intensity data, with labelled phase, φ , amplitude, A, and DC offset, B
Figure 3.1: Block diagram of the Waikato Ranger before improvement
Figure 3.2: Waikato Ranger and Altera Stratix II Development Kit interconnections
Figure 3.3: Top level layout of the Victoria University range imaging system. The system is based around a PMD19k image sensor and an Altera Stratix III development kit. The FPGA drives the illumination and image sensor modulation, retrieves frame date from the sensor, processes it in real-time, and displays it on a VGA monitor or transfers the data to a PC 40
Figure 3.4: Laser diode driver circuit based on the iC-HK Laser Switch
Figure 3.5: Plot of contrast versus modulation high voltage taken from the PMD19k-2 datasheet.44
Figure 3.6: Plot of contrast versus modulation frequency taken from the PMD19k-2 datasheet 44
Figure 3.7: PMD Mainboard sensor modulation driver circuit
Figure 3.8: Nios II processor configuration in Altera SOPC Builder
Figure 3.9: Flow diagram for Nios II application. Task is to process user commands from JTAG UART, and send image frames to PC via Ethernet
Figure 4.1: Laser diode illumination waveform with modulation frequency of 8 MHz as measured by a photodiode
Figure 4.2: Laser diode illumination waveform with modulation frequency of 40 MHz as measured by a photodiode
Figure 4.3: PMD19k sensor modulation waveform with frequency of 8 MHz. Purple and blue traces show the voltage of the A and B modulation inputs respectively. Red trace shows the FFT of their difference

Figure 4.4: PMD19k sensor modulation waveform with frequency of 40 MHz. Purple and blue traces show the voltage of the A and B modulation inputs respectively. Red trace shows the FFT of their difference
Figure 4.5: RMS voltage for illumination measured by photodiode (mV) and PMD19k image sensor modulation input (V) versus modulation frequency. Also shown is an estimation of their correlation product and this product multiplied by frequency
Figure 4.6: Standard deviation of pixel channel outputs for capture with lens cap on and very short integration time
Figure 4.7: Simple static test scene used for characterising tests. In the centre of the field of view is a single white rectangular target, positioned 2.63 m from the range imaging system67
Figure 4.8: Range image of test scene, where distance is represented by colour
Figure 4.9: Amplitude image of test scene. The rectangle in the centre indicates the 40 × 40 group of pixels analysed for various statistics
Figure 4.10: Mean pixel output for central target region versus sensor modulation phase offset for a range of modulation frequencies
Figure 4.11: Semi-logarithmic plots of amplitude of fundamental frequency component versus modulation frequency (blue line, left axis) and the inverse of amplitude and frequency product (green line, right axis)
Figure 4.12: Measured and expected phase versus modulation frequency for a flat target 2.63 m from the imaging system. Also shown is the error in the phase measurement, with a cubic polynomial fit based on frequencies 2 to 44 MHz70
Figure 4.13: Modulation frequency versus Phase standard deviation, σ_{φ} , (blue line, left axis) and distance measurement standard deviation, σ_d (green line, right axis)
Figure 4.14: Logarithmic plot of phase measurement standard deviation versus measurement rate for a range of modulation frequencies73
Figure 4.15: Correlation waveform for $f = 1$ MHz and integration time per frame of 30, 40, 50 and 60 ms73
Figure 4.16: Semi-logarithmic plot of phase standard deviation versus frames per measurement for a number of modulation frequencies where measurement rate is fixed at 16.67 Hz
Figure 4.17: Semi-logarithmic plot of phase standard deviation versus frames per measurement for a number of modulation frequencies where measurement rate is fixed at 4.17 Hz75
Figure 4.18: Semi-logarithmic plot of harmonic amplitude relative to fundamental amplitude for harmonic numbers 2 to 13 and modulation frequencies 1, 10 and 40 MHz76
Figure 4.19: Variation in phase measurement error versus sensor phase offset for modulation frequencies of 1, 10 and 40 MHz, and phase steps per measurement of 3, 4, 5, 6, 8 and 12.77
Figure 4.20: Semi-logarithmic plot of RMS phase error versus modulation frequency for $N = 3, 4, 5, 6, 8$ and 12
Figure 4.21: Power consumption versus modulation frequency for the sensor (CCD camera, image intensifier, DDS) and illumination circuits of the University of Waikato range imaging system
Figure 4.22: Power consumption versus modulation frequency for the sensor and illumination circuits of the Victoria University range imaging system
Figure 4.23: Total power consumption versus modulation frequency for the Waikato University and Victoria University range imaging systems
Figure 5.1: Due to phase wrapping of the returned illumination signal, objects beyond the maximum unambiguous range, d_u , appear to be closer than they really are
Figure 5.2: A target measured twice using different modulation frequencies. Each measurement gives a set of possible object locations, and the true location is determined by where the two measurements are in agreement

Figure 5.3: Experimental σ_{φ} versus modulation frequency for the Victoria University range imaging system. 93
Figure 5.4: Logarithmic plot of simulated and theoretical σ_d versus σ_p for modulation frequency combinations of 40 and 32 MHz, 16 and 8 MHz, 40 and 8 MHz and 8 MHz only
Figure 5.5: Distance measurement error for a number of samples with $\sigma_p = 0.2$ radians. The large spikes are due to errors in the estimate for the integer number of times the phase has wrapped around
Figure 5.6: Logarithmic plot of simulated and theoretical error rate versus σ_p for frequency combinations of 40 and 32 MHz, 16 and 8 MHz, and 40 and 8 MHz
Figure 5.7: The time taken for a single modulation frequency measurement is made up of N integration periods, each of length T/N , and N frame readout periods of length T_R
Figure 5.8: Dual-frequency modulation requires more frames. To keep the measurement rate the same, the total integration time for each frequency, T_A and T_B , is reduced
Figure 5.9: Dual-frequency measurement where measuring each frequency simultaneously reduces the number of readout periods, allowing more time for integration
Figure 5.10: Simulated σ_p versus integration ratio where $f_A = 40$ MHz and $f_B = 32$ and 8 MHz. 105
Figure 5.11: Simulated σ_d versus integration ratio for frequency combinations of 40 and 32 MHz, 40 and 8 MHz, and 8 MHz only. For each multiple frequency measurement a comparison between a fixed weighting scheme of 50% and a variable weighting based on Equation (5.14) is also shown
Figure 5.12: Simulated error rate versus integration ratio for frequency combinations of 40 and 32 MHz, and 40 and 8 MHz
Figure 5.13: Diagram of scene for testing multiple frequency range imaging. The scene consists of 6 boxes positioned at known distances from the imaging system, with one white and one green target on each box
Figure 5.14: Image of test scene where grey scale represents pixel amplitude for a capture with $f = 8$ MHz, $N = 240$ and measurement rate = 0.083 Hz. The numbers indicate where the test targets described in Figure 5.13 can be found in the image
Figure 5.15: Image of test scene where colour represents distance. Pixels with amplitude less than 1.8 mV have been fixed at 12.5 m (further than the distance to the far wall) in order for the image to make better use of the full colour scale
Figure 5.16: Images of test scene where colour represents distance captured using different modulation frequencies: (a) $f = 40$ MHz, (b) $f = 32$ MHz, (c) combination of $f = 40$ MHz and $f = 32$ MHz. Colour scale is the same for all images
Figure 5.17: Plot of experimental σ_p for frequencies of 8, 32 and 40 MHz versus σ_p for 8 MHz. σ_p for 8 MHz scaled by 1.30 and 1.64 is also shown, which is used as the equivalent σ_p for these frequencies so that meaningful comparisons across frequencies can be made
Figure 5.18: Logarithmic plot of experimental σ_d versus σ_p for test pixels measured using frequency combinations of 40 and 32 MHz, 40 and 8 MHz, and 8 MHz only
Figure 5.19: Logarithmic plot of experimental and theoretical error rate versus σ_p for frequency combinations of 40 and 32 MHz, and 40 and 8 MHz
Figure 5.20: Experimental σ_p versus integration ratio where $f_A = 40$ MHz and $f_B = 32$ and 8 MHz. Errorbars indicate the variation in σp across the 16 × 12 group of pixels
Figure 5.21: Experimental σ_d versus integration ratio for frequency combinations of 40 and 32 MHz, 40 and 8 MHz, and 8 MHz only for pixels located at test target 8. Errorbars indicate the variation in σ_d across the 16 × 12 group of pixels
Figure 5.22: Experimental error rate versus integration ratio for frequency combinations of 40 and 32 MHz, and 40 and 8 MHz for pixels located at test target 8
Figure 6.1: Simplified block diagram for calculating phase, p , and amplitude, A from pixel intensity inputs IA_i and IB_i and sine and cosine values C_i and S_i

Chapter 1 Introduction

The Mechatronics Group at Victoria University of Wellington have a number of mobile robot platforms of various shapes and sizes. Our flagship robot is Marvin (Mobile Autonomous Robotic Vehicle for Indoor Navigation) the security robot (pictured in Figure 1.1), whose main role by night is to patrol the corridors and identify intruders, and by day to be a helpful guide to visitors and lost students (Carnegie et al., 2004). For an autonomous mobile robot like Marvin, it is imperative to have a group of depth sensors which can be used to perceive the environment – enabling the robot to avoid obstacles and to interact with objects in its environment, for example, people. Currently, Marvin relies on a combination of laser scanning and infra-red position sensitive devices (Carnegie et al., 2005) and it generally takes a considerable amount of time and dedicated processing to get a complete 3D model of the scene.

A more desirable solution is to scan the entire field-of-view simultaneously – this is full-field range imaging. A number of well known techniques exist based on triangulation methods, such as stereo vision, photogrammetry and structured light; each have their own advantages and disadvantages. The past decade has seen the emergence of a new breed of depth sensor based on the time-of-flight (ToF) principle. These systems utilise active illumination of the scene and, using an array of sensors, measure the time taken for the light to be reflected back to the imaging system.



Figure 1.1: Marvin, the flagship robot of the Victoria University of Wellington Mechatronics Group.

The last four years has seen the development of commercial ToF cameras designed to meet the needs of many applications including automotive safety (Hsu et al., 2006; Ringbeck et al., 2007), entertainment (Yahav et al., 2007), and mobile robotics (Hegde and Ye, 2008; Hussman et al., 2009). While these systems may be appropriate for certain applications, they are not easily reconfigurable outside a prescribed range of operating modes. What we desire is a system that is not only effective at range imaging, but also highly configurable with regards to all its operating parameters. With more direct control over the various aspects of the 3D image acquisition process, we aim to implement and evaluate new techniques for ToF range imaging, where the focus is more on how to utilise these sensors in a novel way.

1.1 Objectives

The objective of this research is to develop a new ToF full-field range imaging system that is appropriate for use in autonomous mobile robotics applications. The system must have the following features:

- Compact The system must be of a physical size small enough to be mounted on a mobile robot, for example, within the torso of Marvin;
- Low-power The system must consume as little power as possible in order to operate for extended periods on a portable power supply;
- Real-time The system must image the scene and determine distance in real-time so the robot can respond immediately to changes in the environment;
- Embedded processing Distance must be calculated by the imaging system itself to reduce the load on the robot's main central processing unit (CPU);
- High precision Distance measurement precision must be comparable to commercial ToF camera systems;
- Applicable to indoor environments The system must operate to a range of at least 6 m under artificial lighting conditions.

Another objective of the system is to be highly configurable so that it can act as a test platform for implementing and evaluating new range imaging methodologies. Configurable parameters include:

- The integration time of the sensor and measurement frame rate;
- Independent control of the illumination and image sensor modulation signals;
- Adjustment of modulation frequency from 1 to 40 MHz;
- Adjustment of modulation phase offset with at least $\pi/6$ phase resolution;
- All parameters must be adjustable in real-time, with the ability to change parameters between successive frames as well as during the sensor integration period.

Once the system has been designed and constructed it must be characterised to measure its performance in terms of distance measurement precision and accuracy. The configurable nature of the system will be especially useful for establishing optimal operating parameters, and in some cases this may indicate a deviation from the typical options used by commercial ToF cameras.

One promising technique for improving maximum range and measurement precision is the use of multiple modulation frequencies. The newly constructed system can be configured to implement a dual-frequency technique, and evaluate the feasibility of such a method.

A further objective is to analyse the requirements of the real-time distance processing. By identifying the critical resource of a hardware implementation, the processor can be scaled appropriately to the task. This can lead to a design that uses a minimum of processing resources, and a more cost effective solution.

1.2 Outline

This thesis is divided into 7 chapters.

Chapter 2 presents an overview of range imaging techniques. The full-field indirect ToF technique is identified as best meeting the project objectives. Understanding the factors affecting precision and accuracy of ToF range imaging are essential in the design of a high performance system.

Chapter 3 describes the hardware development of an easily configurable range imaging system. The design evolves from a modification of an existing range imaging system to a completely new system constructed from a number of subcomponents. This new design has proven invaluable in implementing the ideas presented in the later chapters of this thesis, so that quantitative results can be produced.

Chapter 4 presents results of a series of characterising tests performed on the system described in Chapter 3. Particular emphasis is placed on distance measurement precision, and the effect of operating parameters such as modulation frequency, integration time and phase steps per measurement.

Chapter 5 discusses the concept of distance measurement ambiguity for continuous wave range imaging systems, particularly methods for extending the maximum unambiguous range through the use of multiple modulation frequency

measurements. A theoretical investigation into the trade-off between precision and maximum range is presented, as well as the probability of errors arising from the disambiguation process. The concepts have been tested using the system described in Chapter 3, and experimental results are presented.

Chapter 6 addresses the challenges of processing the raw images produced by the hardware into depth and intensity information as required by a given application. In many situations it is desirable for this processing to occur in real-time. Therefore a particular emphasis is placed on implementing the depth and intensity determination in hardware, and the various advantages and disadvantages of such a method are explored.

The thesis is concluded in **Chapter 7** with a selection of range imaging examples demonstrating the performance of the new system in high-precision and high frame-rate operating modes. The thesis is briefly summarised and a number of suggestions are given for future work within this field. A list of publications arising from this research project, including a book chapter and two journal articles is presented, before concluding with a final summary of the novel contributions of this thesis.

Chapter 2 Full-Field Range Imaging

The goal of full-field range imaging is to create a three dimensional model of an environment. This can be achieved through active or passive means using technology based on a variety of acoustic or optical sensors. Our main application of interest is for use with autonomous mobile robotics, so that a robot may have a more detailed awareness of its environment. Desirable features are a fast acquisition rate, small system size, low power, and high x-y-z resolution.

The fundamental methodologies used by acoustic (ultrasonic) and optical (visible light, infra-red, microwave) ranging systems are essentially the same. However, the wavelength of light compared to that of sound is much smaller, thereby allowing for much higher depth resolution. This thesis discusses only range imaging through optical means.

Optical full-field range imaging technology can be broadly grouped into four different methodologies, each with their own inherent advantages and disadvantages. These are:

- passive stereo and photogrammetry;
- active triangulation and structured light;
- interferometry; and
- time-of-flight measurement.

ToF measurement systems can be further divided into direct and indirect methods, and full-field simultaneous versus point scanning systems. This chapter aims to describe these technologies, with particular emphasis on the full-field, simultaneous indirect ToF technique. In Section 2.1 a brief introduction into alternative range imaging techniques that do not use ToF methods are presented for comparison purposes. Section 2.2 introduces the direct ToF technique where the time taken for light to travel to and from a scene is measured directly and converted to distance. Indirect ToF methods aim to address some of the deficiencies of the direct method with respect to shorter range (metres) distance resolution at video frame rates, and are presented in Section 2.3. Recently, a number of commercial ToF systems have become available for purchase based on custom Complementary Metal Oxide Semiconductor (CMOS) image sensor arrays using the Amplitude Modulated Continuous Wave (AMCW) ToF technique. An investigation into the latest trend towards CMOS ToF sensors and fully integrated camera systems is presented in Section 2.4. In Section 2.5 we discuss the issues relating to distance measurement precision, and in Section 2.6 distance measurement accuracy for AMCW ToF range imaging systems. Finally, a summary of the chapter is presented in Section 2.7.

2.1 Alternative Range Imaging Technologies

In this section we briefly introduce a number of alternative technologies used for full-field range imaging applications. Traditionally, range imaging has been achieved through stereo and active triangulation; however these methods have some limitations as will be discussed below.

2.1.1 Passive Stereo and Photogrammetry

Passive stereo systems operate in a similar manner to eyes in humans and animals. They are termed passive because they do not output light to illuminate the scene, rather relying on ambient lighting. An image is captured from two sensors separated by a known distance (called the base line), so objects in the scene will appear in both images, but viewed from a slightly different angle. The method essentially uses reverse ray tracing and simple trigonometry to find the distance based on the angle of objects in each of the stereo images, as shown in Figure 2.1. The major requirement is in identifying features in each of the two images that represent the same object. While the hardware requirements are small, requiring



Figure 2.1: Stereoscopic distance measurement.

only two standard cameras, the software is complex; though real-time stereo systems have been possible for some time (Kimura et al., 1995).

One limitation is that scenes with few contrasting features have reduced performance of the registration algorithm, and hence for some scenes a range measurement is not possible for all targets. Note that as objects must appear in both images to be measured, the measurement volume is defined by the intersection of the cameras' field of view. Another problem is occlusion, where a range measurement cannot be made if the view from either of the sensors is obstructed. The distance resolution of the method is limited by the resolution of the angle measurements, particularly as the distance becomes much larger than the base line, and increasing the base line is contrary to producing a system that is compact.

Photogrammetry is a similar technology, but uses a larger number of image sensors (three or more) placed around the scene. With image sensors carefully placed from different angles, the problem of occlusion can be diminished as extra images can fill in the missing information. In general, more images leads to a more accurate model of the scene, including statistical information on the measurement quality, but processing becomes even more complex (Atkinson, 2001). Photogrammetry is not appropriate for our application due to the inability

of a mobile robot to simultaneously view a scene from a number of different angles, and the long time required to process the 3D information.

2.1.2 Active Triangulation and Structured Light

Active triangulation works in a similar way to passive stereo but projects a line or pattern of light onto the scene. The image sensor, typically a standard Charged Coupled Device (CCD) or CMOS camera views the projected light from a known offset angle, and depth can be measured by the observed deviation of the projected light (Sato, 1996). A line scanning system makes use of a mechanism (for example, a rotating mirror) for moving the line across the imaging area in order to build up a full-field image of the scene (Hirose, 2001). Full-field systems make use of a two dimensional pattern of light, where distortions to the pattern can be recognised by the imaging system and interpreted as distance (Salvi et al., 2004).

Limitations of this technology with respect to a mobile robotics application are similar to that for stereo vision and photogrammetry – there must be a distance offset between the image sensor and the pattern projector which is related to the distance resolution, as well as measurement volume and occlusion problems if either the illumination or the camera view are obstructed. The use of active illumination also has a disadvantage in that as the distance is increased the projected illumination becomes dispersed. This reduces the resolution of the projected features, and also limits the system's performance under ambient lighting conditions (Hebert, 2000).

2.1.3 Interferometry

Interferometry studies the patterns of interference when two or more waves of light are superimposed onto each other (Hariharan, 1985). In the case with two light sources, the relative distance travelled between them introduces a phase difference, which then affects the intensity of the light when combined. When used with visible light this can provide depth measurement down to sub-micrometre resolution, depending on the wavelength (colour) of the light used. However, the maximum range of such a system is also typically very short – of

the order of centimetres. This short range makes this technology unsuitable for mobile robotics applications.

2.2 Direct Time-of-flight Measurement

Direct time-of-flight technologies, also known as Light Detection and Ranging (LIDAR) measures the time-of-flight taken for active illumination to travel to and from the scene. Typically, systems emit a brief light pulse, and measure the time for the pulse to return as illustrated in Figure 2.2. Distance, *d*, can be calculated by measuring time taken for the light to return to and from the scene by the relationship

$$d = \frac{ct}{2},\tag{2.1}$$

where c is the speed of light, and t is the time between the emission and detection of the light pulse. The distance resolution of such a system is dependent on the resolution of the timer, for example, for 1 cm distance resolution the timer must be able to resolve the returned pulse to within 66 ps. This requires an illumination system that can emit a high power light pulse in a very short amount of time, and a detector system with very high bandwidth, for example, to resolve a 66 ps signal using a standard counter would require a clock operating at 15 GHz.



Figure 2.2: Time-of-Flight range measurement. A pulse of light is emitted and reflected from the scene, and the time taken for the light to travel is measured.

ToF systems can be point scanning, line scanning or full-field scanning depending on the complexity and scalability of the sensor. Point scanning systems are electronically the simplest as they require only a single point source for the illumination and a single sensor element (such as a photodiode). The illumination may be focussed to produce a narrow beam and consequently a high amplitude return, making the technology suitable for long range applications (over 100 metres). A line scanning system can be realised by dispersing the illumination along one dimension and recording the reflection using a one dimensional array of sensor elements. The advantage over a point scanner is that it gives range to all points on the line simultaneously. However, this comes at the cost of increased sensor complexity and reduced amplitude of the returned illumination pulse.

A point or line scanning sensor can be used to construct a full-field range image of a scene by taking multiple measurements moving across the scene in two dimensions (Besl, 1998). The scanning movement can be achieved either by moving the sensor unit itself, redirecting the light pulse using mirrors, or by moving the objects in the scene (for example, using a 2D translation stage). Typically, the accuracy of the system is constrained not by the accuracy of the sensor but by the accuracy of the mechanism used to provide the movement. In addition, performing a large number of measurements in a serial manner does not lend itself well to real-time full-field operation.

A promising technology in the area of simultaneous full-field measurement is the use of two-dimensional arrays of Avalanche Photodiodes (APDs) to directly measure the time-of-flight of a laser pulse (Niclass et al, 2009). The technology is still in its infancy compared to phase demodulation methods (see next section), and sensor arrays typically have fewer pixels due to the relatively large area required for the APD and supporting circuitry.

2.3 Indirect Time-of-Flight

To overcome the high bandwidth requirements of the direct ToF method, a number of techniques exist that seek to encode the distance information indirectly, for example, through amplitude or phase.

2.3.1 Shuttered Light Pulse

In Christie et al. (1995), the authors devise a method where a light pulse is emitted and an image sensor array is used to measure the intensity of the reflected light pulse. However, before reaching the image sensor the light is gated by a highspeed shutter, and consequently only a portion of the reflected light reaches the sensor as shown in Figure 2.3, with an indication of pulse timing shown in Figure 2.4. Reflected pulses from nearer objects have a larger proportion of the pulse integrated by the sensor resulting in a higher intensity pixel, while more distant objects have lower intensity. One problem with this method is that from a single image it is impossible to distinguish between the intensity of a near object of low reflectance and a far object of high reflectance. This can be overcome by taking a second measurement with the shutter entirely open allowing the entire pulse to be integrated, and then using this intensity to normalise the measurement. The distance, *d*, of each pixel can be calculated by

$$d = \frac{cT(1 - I_1/I_2)}{2},$$
 (2.2)

where T is the width of the light pulse and I_1 and I_2 are the pixel intensity measurements from the shuttered and non-shuttered pulses respectively.

The greatest challenge of such a system is in the design of the imaging sensor, which must be shuttered at a very high frequency. Previous systems have used an image intensifier coupled with a standard CCD camera, where the gain of the photocathode can be electronically controlled to give very short integration periods (Christie et al, 1995; Dorrington et al, 2007a). This can give excellent contrast and high spatial resolution, as the spatial resolution is constrained by the limiting resolution of the image intensifier (measured in linepairs/mm) which is typically less than the size of the pixel array for a standard CCD camera. The



pulsed light source

Figure 2.3: Indirect Time-of-Flight range measurement using a high-speed shutter at the receiver. Only a portion of the returned light is passed through the shutter, encoding range as amplitude, i.e., pixel intensity.

Emitted light	<u></u>	
Shutter frame 1		
Reflected light (near object)		
Reflected light (far object)		
Shutter frame 2		
	$\downarrow \qquad \qquad$	

Figure 2.4: Timing diagram of shutter and reflected light where the shutter signal is in phase with emitted light. For the first frame, nearer objects return a higher intensity than farther objects. The second frame collects light from the whole pulse, and is used to normalise the measurement.

main limitation of the method is that distance measurement precision is related to the intensity of the pixels in the image. For objects which give a low intensity reflection (either due to distance or poor object reflectivity), the precision is constrained by the dynamic range of the image sensor.

2.3.2 Range Gating Segmentation

The distance measurement precision of the shuttered light pulse method is dependent on the dynamic range of the image sensor since range is directly proportional to measured amplitude. An alternative proposed in (Busck and Heiselberg, 2004) is to take many frames where the shutter phase relative to the reflected light pulse is stepped as shown in Figure 2.5. The resultant intensity profile of the N captured frames reaches a peak when the shutter and reflected



Figure 2.5: Timing diagram of Range Gating Segmentation method. *N* images are captured with stepped shutter offset. Range is determined by the location of the peak intensity in the image sequence.

light are in phase. The average pixel intensity, *A*, can be calculated by the sum of all the intensity measurements:

$$A = \sum_{i=1}^{N} I_i , \qquad (2.3)$$

where I_i is the pixel intensity value for frame number *i* from 1 to *N*. The distance is then computed by the weighted average

$$d = \frac{c}{2A} \sum_{i=1}^{N} I_i t_i, \qquad (2.4)$$

where t_i is the time delay of the shutter for frame number *i*. As the range measurement is not directly related to the pixel intensity, the system precision is no longer constrained by the dynamic range of the image sensor. However, distance precision is dependent on the number of frames. Therefore such a system requires a large number of frames resulting in a relatively long (of the order of seconds) acquisition time, making the method inappropriate for a real-time mobile robotics application.

2.3.3 AMCW Phase Stepped Homodyne

This technology illuminates the scene using a continuous wave of amplitude modulated light as illustrated in Figure 2.6. The light is reflected back from objects in the scene and integrated over a time period by an image sensor with a shutter amplitude modulated at the same frequency. Due to the time taken for the light to travel to and from the scene, the modulation phase envelope of the returned light is shifted relative to the image sensor modulation (Monson et al., 1999). This phase shift, φ , is related to the distance by

$$d = \frac{c\varphi}{4\pi f},\tag{2.5}$$

where *f* is the modulation frequency.

For simplicity, we will initially assume the illumination and sensor modulation waveforms are sinusoidal. The reflected illumination signal can be expressed as (Luan, 2001)

$$r(t) = R\cos(2\pi f t - \varphi) + R_0 \tag{2.6}$$

where *R* is the average received illumination intensity and accounts for the average transmitted intensity, object reflectivity and the $1/d^2$ relationship between intensity and distance. R_0 is a DC offset caused by background lighting. The



modulated light source

Figure 2.6: AMCW range imaging.

sensor modulation can be expressed as

$$s(t) = S\cos(2\pi f t - \theta) + S_0, \qquad (2.7)$$

where S is the average sensor gain, S_0 a DC offset, and θ is an artificial phase offset introduced between frames. The resultant action of the shuttered image sensor is to multiply and integrate the two signals over a time period, T, yielding a pixel intensity function, $I(\varphi, \theta)$, of

$$I(\varphi,\theta) = \int_0^T r(t) \cdot s(t)dt$$

= $\int_0^T [R\cos(2\pi f t - \varphi) + R_0)] \cdot [S\cos(2\pi f t - \theta) + S_0]dt$
= $\int_0^T \left[\frac{RS}{2}\cos(\theta - \varphi) + \frac{RS}{8\pi f}\cos(4\pi f t - \theta - \varphi) + .$ (2.8)
+ $R\cos(2\pi f t - \theta) + S\cos(2\pi f t - \varphi) + R_0S_0T]dt$

The integration period, *T*, is typically several orders of magnitude larger than the modulation period, 1/f. Consequently after integration, the 2^{nd} , 3^{rd} , 4^{th} and 5^{th} terms can all be expressed as a single constant offset, *B*. This leads to the simplified expression

$$I(\varphi, \theta) = A\cos(\theta - \varphi) + B, \qquad (2.9)$$

where *A* is the pixel intensity signal amplitude and *B* is the DC offset component from all sources, including background illumination.

Between captured image samples the phase offset of the image sensor modulation, θ , is stepped by $2\pi/N$ radians, where N is the number of phase steps per range measurement. This is illustrated in Figure 2.7 for N equal to four, with an example pixel intensity sequence shown in Figure 2.8. Square wave modulation is shown to more easily illustrate the shuttering effect of the sensor. After N samples have been captured, φ , A, and B can be calculated from the Fourier series using



Figure 2.7: Timing diagram for phase stepped AMCW homodyne range imaging. The reflected modulated light is multiplied by the shutter signal and integrated by the camera. The phase of the shutter signal is stepped between frames, encoding range into the phase of the sequence of intensity images.



Figure 2.8: Fitting a sinusoid to the intensity data, with labelled phase, φ , amplitude, A, and DC offset, B.

$$\varphi = \tan^{-1} \left(\frac{\sum_{i=0}^{N-1} I_i \sin(2\pi i/N)}{\sum_{i=0}^{N-1} I_i \cos(2\pi i/N)} \right), \qquad (2.10)$$

$$A = \frac{2}{N} \sqrt{\left[\sum_{i=0}^{N-1} (I_i \cos(2\pi i/N))\right]^2 + \left[\sum_{i=0}^{N-1} (I_i \sin(2\pi i/N))\right]^2}, \qquad (2.11)$$

$$B = \frac{1}{N} \sum_{i=0}^{N-1} I_i, \qquad (2.12)$$

where I_i are the pixel intensity values read from the image sensor for frame number *i*. Most common systems (Lange and Seitz, 2000; Hsu et al, 2006, Ringbeck et al, 2007) have *N* fixed at four, as this simplifies the sine and cosine values to the set {0, 1, -1}, and the calculation of φ and *A* simplify to

$$\varphi = \tan^{-1} \left(\frac{I_1 - I_3}{I_0 - I_2} \right), \tag{2.13}$$

$$A = \frac{1}{2}\sqrt{(I_1 - I_3)^2 + (I_0 - I_2)^2}.$$
 (2.14)

A more accurate calculation of the distance must also take into account the maximum unambiguous range (explained below) of the system, d_u , by

$$d = d_u \left(\frac{\varphi}{2\pi} + k\right),\tag{2.15}$$

where k is an integer and d_u is based on the modulation frequency by

$$d_u = \frac{c}{2f}.$$
 (2.16)

For most applications where the maximum extent of the scene is known, f can be chosen appropriately and k assumed to be 0.

The concept of the maximum unambiguous range, d_u , accounts for the possibility of the light reflected from more distant objects to be delayed by integer multiples of the modulation period. This causes an ambiguity in the measurement where these more distant objects are incorrectly measured as being much closer than they actually are. The integer k in (2.15) is used to account for the extra number of modulation periods that the reflected light has been delayed. This concept is further explained in Chapter 5 along with methods for extending d_u .

Digital switching circuits are a common method for generating the illumination and image sensor modulation, which produce square waveforms that have inherently large amplitude odd-ordered harmonics. A disadvantage to a fixed value for N of 4 is that odd-ordered harmonics of the modulation signal are aliased onto the fundamental frequency bin and contribute to errors in the phase measurement. This effect is further discussed in Section 2.6.4.

2.3.4 AMCW Heterodyne

In Dorrington et al. (2007a), the authors make use of a system similar to phase stepped homodyne but the signal modulating the shutter (the photocathode of an image intensifier tube) has a slightly different frequency to the illumination modulation. If we let the illumination frequency be equal to f, then the frequency of the sensor modulation is $f + f_d$, where f_d is the frequency difference. While the system is in continuous operation the phase of the sensor modulation relative to the illumination is continuously shifting. The difference d is also referred to as the beat frequency, because with the continuously shifting phase offset the intensity of objects in the scene appear to beat on and off at a frequency of f_d . By sampling a full cycle of the beat waveform where the integration of the first frame is synchronised to when the modulation signals are in phase, the phase and amplitude can be calculated in a similar manner to (2.10) and (2.11)

$$\varphi = \tan^{-1} \left(\frac{\sum_{i=0}^{N-1} I_i \sin\left(\frac{2\pi i f_d}{f_s}\right)}{\sum_{i=0}^{N-1} I_i \cos\left(\frac{2\pi i f_d}{f_s}\right)} \right), \qquad (2.17)$$

$$A = \frac{2}{N} \sqrt{\left[\sum_{i=0}^{N-1} \left(I_i \cos\left(\frac{2\pi i f_d}{f_s}\right)\right)\right]^2 + \left[\sum_{i=0}^{N-1} \left(I_i \sin\left(\frac{2\pi i f_d}{f_s}\right)\right)\right]^2}, \quad (2.18)$$

where f_s is the sampling rate, that is, the frame rate of the sensor.

The primary advantage of such a system is that the phase offset is continuously shifting during the sensor integration period, reducing the amplitude of harmonics in the phase measurement (Dorrington et al., 2008), as is further discussed in Section 2.6.4.

2.4 CMOS ToF Sensors and Camera Systems

The past decade has seen a large development in AMCW ToF sensors produced using CMOS technology where the gain modulation of the imaging pixels is controlled on-chip (Lange and Seitz, 2001; Gulden et al., 2002; Gokturk et al., 2004). This technology is highly desirable for many ToF range imaging applications due to their compact size, low power and low cost compared to the use of image intensifier tubes: all features which make them ideal for autonomous mobile robotics applications. The pixel array sizes are currently relatively small compared to image intensifier based systems, however with ongoing improvements to the CMOS pixel design and scalability this can be expected to increase over the next decade.

Most current commercial ToF camera systems make use of a CMOS image sensor with an integrated illumination source. As at July 2010, commercial products are available from

- Fotonic (Stockholm, Sweden), utilising a sensor from Canesta Inc. (Sunnydale, CA, USA);
- Mesa Imaging AG (Zurich, Switzerland);
- PMD Technologies GmbH (Seigen, Germany);
- Optrima NV (Brussels, Belgium);
- IEE SA (Luxembourg City, Luxembourg);
- Panasonic Electric Works Ltd (New Providence, NJ, USA).

Table 2.1 summarises a selection of parameters of interest for products from this group of manufacturers. The quoted depth precisions for the different products are not necessarily directly comparable. Many factors such as ambient lighting conditions, illumination power, integration time, modulation frequency (where applicable) and target reflectivity all affect the depth precision, and these parameters are not necessarily consistent for all camera characterising tests. The range imaging system constructed as part of this research project must have specifications similar to these commercial ToF systems.

Company	Product Name	Pixels	Field-of-	Frame	Depth Precision	Dimensions
		(H×V)	View	Rate		(WxHxD)
			(H×V)	(Hz)		(mm)
Fotonic	Fotonic B70	160×120	70°×50°	50	0.5% @ 3 m	90×90×120
Mesa Imaging AG	SR4000	176×144	43.6°×34.6°	54	0.25 @ 2 m ^A	65×65×68
PMD Technologies	PMD[vision]®	200×200	40°×40°	40	0.075% @ 4 m ^B	194×60×60 ^C
GmbH	CamCube 3.0					
Optrima NV	OptriCam™	120×90	50°×37.5°	50	0.86% @ 3.5 m ^D	74×86×76
	DS10k-A					
IEE SA	3D MLI	61×56	130°×100°	10	1.33% @ 1.5 m	54×144×104
Panasonic Electric	EKL3104	160×120	60°×44°	20 ^E	1.5% @ 2 m ^F	170×54×49
Works Ltd						

Table 2.1: Selection of commercially available compact 3D Cameras.

A - 50% reflective target.

B - Typical value, 75% reflective target, central sensor area.

C - Basic camera configuration with two illumination modules attached.

D-At frame rate of 25 Hz, 60% reflective target, no ambient lighting.

E - Fixed frame rate.

F - 90% reflective target, central sensor area, no ambient lighting.

2.5 AMCW Measurement Precision

The concept of precision used throughout this thesis refers to the repeatability of the measurement, that is, the variation of individual pixels over successive samples. It is calculated by the 1-sigma standard deviation of the samples. This is considered separately to the measurement accuracy of the system, which is a measure of the difference between the mean of the measured samples and the true expected distance measurement.

2.5.1 Phase Measurement Precision

For an AMCW range finding system, the standard deviation of the phase measurement, σ_{φ} , is inversely proportional to the square root of the signal-to-noise (SNR) ratio of the measured pixel intensity values by the formula (Jelalian, 1992)

$$\sigma_{\varphi} = \frac{1}{2m\sqrt{SNR}}, \qquad (2.19)$$
where *m* is the modulation index, describing the relationship between the maximum and minimum levels of the modulation waveform. For a solid state image sensor with the pixel modulation designed into the chip we refer to a similar measure: the demodulation contrast, c_d , calculated as (Lange, 2000b)

$$c_d = \frac{A}{B},\tag{2.20}$$

where A is the measured demodulated pixel amplitude and B is the DC offset. A number of factors affect the demodulation contrast including modulation frequency, f, illumination wavelength and the construction of the image sensor (Lange, 2000b). As an indication of values for c_d , the datasheet for the PMD Technologies 19k image sensor (PMD, 2008) specifies a value of 0.40 with 1 MHz modulation and 650 nm red light, and 0.28 with 10 MHz modulation and 860 nm near infra red light.

The SNR of the pixel intensity is influenced by many factors which are independent of the amplitude of the received light, and are all encompassed within the dark noise measurement of the sensor. These include thermal noise, electronic shot noise, flicker noise, reset noise, dark current and quantisation error of the pixel voltage (Light, 2008). However, photon shot noise is the limiting factor to achieving maximum precision (Lange, 2001). Photon shot noise describes the random nature of the conversion of incidental photons into charge carrying electron-hole pairs. It can be statistically modelled as a Poisson-distributed process, which for a large number of photons (as is the case with reasonable integration time) can be well approximated using a Gaussian distribution (Frank et al., 2009).

In summary, phase standard deviation, σ_{φ} , can be estimated based on the signal amplitude, *A*, and DC offset, *B*, by the formula (Buttgen et al., 2005)

$$\sigma_{\varphi} = \frac{\sqrt{B}}{\sqrt{2}c_d A}.$$
(2.21)

In order to improve phase precision (decrease σ_{φ}) it is necessary to increase the demodulation contrast, or reduce the ratio \sqrt{B}/A . In most cases as an end-user of an image sensor Integrated Circuit (IC), options for modifying the demodulation contrast are limited. However, care should be taken when designing the modulation signal driver circuits to maximise bandwidth, hence reducing the effect of high modulation frequency reducing c_d .

The second approach of reducing the ambient light relative to the illumination signal is commonly achieved through the use of an optical filter placed in front of the image sensor, which is matched to the wavelength of the illumination source. Another option is to increase the power of the illumination, however in some applications this may not be appropriate if eye safety is a concern.

It can also be seen that if the total integration period is extended by a factor of E, then both A and B will be increased equally. This can result in a net reduction in σ_{φ} of \sqrt{E} . This approach is limited due to the effect of saturation, where the storage capacitors of the pixels can collect only a finite amount of charge. The last remaining option is to collect a greater number of measurements, and reduce σ_{φ} through averaging. Improving phase measurement precision through longer integration time or extra measurements is not always appropriate for the application, particularly for real-time measurements of a moving scene.

2.5.2 Distance Measurement Precision

Distance measurement standard deviation, σ_d , is derived from phase measurement standard deviation, σ_{φ} , by the formula

$$\sigma_d = \frac{c\sigma_{\varphi}}{4\pi f} = \frac{d_u \sigma_{\varphi}}{2\pi} \,. \tag{2.22}$$

Assuming steps have been taken to minimise σ_{φ} , the last remaining mechanism for improving distance measurement precision is to increase the modulation frequency, *f*. If the system limits the maximum range to d_u (assuming k = 0) then a trade-off must be made between precision and maximum range. Options for extending the maximum range while still maintaining high precision based on the use of multiple modulation frequencies have been presented in (Jongenelen et al., 2010a; Jongenelen et al., 2011), and these are further discussed in Chapter 5.

Ultimately, the maximum frequency is limited by the bandwidth of the electronics. As the frequency is increased past a certain (system dependent) point, f_c , the demodulation contrast rapidly begins to degrade the phase precision, negating any gain from the increased frequency.

2.6 AMCW Measurement Accuracy

In this section we will discuss the issues relating to measurement accuracy. These are factors that cause a systematic error in the phase or distance measurement which can not be reduced by averaging. However, if they are well understood then steps can be taken to minimise their effects either through appropriately modifying system operating parameters or post processing of the data.

2.6.1 Geometric Distortions

A ToF camera measures the distance of a straight line from objects in the scene to the corresponding pixel that has imaged that object. Therefore, radial data acquired in the ToF camera reference frame must be converted to real world coordinates through a mapping function based on the size of the image sensor and the focal length of the lens. A simple perspective projection approach is possible where a pinhole model for the camera can be used (Rapp, 2007) where an assumption is made that the emitted and received illumination are a coaxial point source. However, in practice the illumination typically originates from a point, or a number of points beside or around the image sensor. For distant objects this approaches an ideal point source, but for short distances systematic errors are likely. In addition, imperfections in the optics can also cause distortions in the distance measurement. Fortunately, these effects are constant during normal system operation and their effects can be compensated for by calibration and post processing of the data (Dorrington et al., 2010; Fuchs and Hirzinger, 2008).

2.6.2 Fixed Pattern Noise and Latency Error

The construction of transistors within a CMOS integrated circuit can not be guaranteed to behave identically, despite the designer's best wishes (Light, 2008). These variations manifest themselves as fixed pattern noise, both within individual pixels and also in the readout circuitry, where entire rows or columns of pixels may be affected differently from other rows or columns.

CMOS ToF sensors are constructed as a two dimensional array of pixels where the global modulation signal is applied to all pixels simultaneously. Due to the finite speed at which signals can propagate through the CMOS device, it is possible for the modulation signal to activate individual pixels at different times. This difference in latency is perceived as a phase shift across a group of pixels.

In addition, some sensors (such as the PMD Technologies 19k) divide the pixel array into a number of blocks, each modulated by a separate external signal (PMD19k datasheet). This is done in order to reduce the load capacitance of each channel, simplifying the requirements of the modulation driver circuit. However, this can lead to a variation in latency between blocks (Payne et al., 2009a).

All of these effects can be characterised and their influence reduced through calibration, for example by way of calculating a 'Fixed Pattern Noise Matrix' which specifies an offset to be applied to each pixel (Kahlman et al, 2006; Lindner and Kolb, 2007).

2.6.3 Temperature Related Error

A well known problem with semiconductor technology is the effect of temperature on charge carrier efficiency. In the case of CMOS image sensors an increased temperature causes a larger number of thermally generated electrons interfering with the collection of photon generated electrons (Kahlmann et al., 2006). In addition to increased random noise on the measurement, increased temperature also introduces an offset to the phase measurement. When performing system calibration it is important to operate the system in a continuous mode for some minutes to allow it to reach a constant internal temperature before

taking measurements. The accuracy will still be affected by changes in the ambient temperature, but the error is systematic and thus removable by incorporating the ambient temperature in the calibration routine (Kahlmann et al., 2006).

2.6.4 Errors due to Harmonics

The ideal form of modulation contains only a single frequency component at the modulation frequency, *f*. Imperfect modulation contains harmonic components at integer multiples of the modulation frequency with their phase and amplitude affected differently due to nonlinearities in the illumination driver and sensor electronics.

After multiplication and integration by the sensor, the pixel intensity can be expressed as the Fourier series summation (Frank et al., 2009)

$$I_i = \sum_{k=1}^{k} A_k \cos(k\theta_i - k\varphi + \alpha_k) + B, \qquad (2.23)$$

where A_k is the correlated amplitude and α_k is the phase for the *k*th harmonic. In the non-ideal case, A_k and α_k are not equal to zero for k > 2 (the harmonic content) and according to the Nyquist-Shannon theorem these will be aliased onto the fundamental component if the sampling frequency is insufficient. For the integer number of phase steps per measurement, N, the harmonic number of the aliased frequencies are given by

$$k = nN \pm 1, \tag{2.24}$$

where *n* is an integer.

In the case where both illumination and sensor modulation waveforms are square, the correlation waveform will be triangular. This situation leads to relatively large odd-ordered harmonics. With N selected as 4, all of these harmonics are aliased onto the signal of interest.

The erroneous phase offset caused by harmonics leads to a systematic phase error that is dependent on the actual object distance, the sampling method and the bandwidth of the electronics. The effect can be characterised and removed through calibration (Fuchs and Hirzinger, 2008; Boehm and Pattinson, 2010), however, recalibration must be performed whenever certain operating parameters are changed, for example, modulation frequency. The use of heterodyne AMCW modulation (see Section 2.3.4) has been demonstrated to reduce the amplitude of aliased harmonics, albeit at a cost of reduced fundamental amplitude, hence reduced precision (Dorrington et al., 2008). Alternatively, a technique exists to cancel a predetermined number of odd-ordered harmonics by shifting the sensor phase offset during the integration period (Payne et al., 2008). Similarly to heterodyne, the method comes at a cost of reduced fundamental amplitude, however the technique is able to cancel odd-ordered harmonics completely.

2.6.5 Scattering and Mixed Pixel Effects

Scattering is where incident light is reflected internally between the lens and the image sensor and mixing into other pixels, causing errors in the distance measurement across the entire image (Mure-Dubois and Hugli, 2007). The effect is particularly evident in scenes with objects returning different intensity reflections, where the object of lower intensity is generally more affected, sometimes to the extent that the direct path return is completely lost (Godbaz et al, 2009a). Because the effect is scene dependant, it is not possible to calibrate the system to alleviate the error. Techniques are under development to compensate for the error, however they require substantially more processing time than the usual per-pixel phase determination, increasing measurement latency and limiting real-time ToF camera frame rates (Mure-Dubois and Hugli, 2007; Godbaz et al., 2009a; Godbaz et al., 2009b).

Another scene dependent issue is that of mixed pixels at object edges. Due to the fixed size of an individual pixel's imaging area, the light returned from object edges may be reflected from both the object and the background. The resultant phase measurement may lie somewhere between the expected measurement for either side of the edge, or if the difference is greater than 2π radians a completely

implausible phase is measured (Godbaz et al., 2007). With imperfect focus or a moving scene, the problem may be spread across multiple pixels. A common technique to deal with this issue is to filter out edges using standard image processing algorithms.

2.6.6 Distance Ambiguity

The use of continuous wave modulation leads to an ambiguity problem if the phase shift of the reflected illumination exceeds 2π radians. This effect has been briefly discussed in Section 2.3.3, where Equation (2.16) gives the maximum unambiguous range, d_u , when taking measurements with a single modulation frequency. For example, with a modulation frequency of 40 MHz, the maximum unambiguous range is 3.75 m. Objects beyond this distance are incorrectly measured to appear closer to the camera by integer multiples of d_u .

An alternative to continuous wave modulation is to use a digitally generated pseudo-random sequence (Buttgen and Seitz, 2008). The primary advantage of this method is that multiple range imaging cameras can operate within the same environment without interference. A secondary feature is that the reflected illumination from objects beyond d_u are no longer correlated with the sensor modulation waveform, and can therefore be identified and removed from the data.

The simplest method to extend d_u is to reduce f, however this often comes at a cost of reduced distance measurement precision. In some applications the extent of the scene to be measured is known and f can be selected appropriately. For the more general case it is desirable to extend d_u whilst maintaining measurement precision. The use of multiple measurements using different modulation frequencies has been demonstrated to achieve this (Payne et al., 2009a; Jongenelen et al., 2010a). This concept and related issues are discussed in more detail in Chapter 5.

2.7 Summary

In this chapter we have presented a number of methods for simultaneous, fullfield range imaging. Time-of-flight methods are of particular interest due to their high-speed acquisition rate and relatively high measurement precision. This chapter has paid particular attention to the AMCW Phase Stepped Homodyne method of indirect ToF range imaging. The technology is amenable to the production of image sensors on a large scale using a CMOS process, reducing overall system cost and power requirements.

A number of factors affect distance measurement precision and accuracy for AMCW ToF range imaging systems. The predominant factor affecting precision is photon shot noise, and methods to improve precision are based on increasing demodulation contrast, increasing signal amplitude and increasing modulation frequency. A number of factors impacting measurement accuracy can be minimised through system characterisation, calibration and post-processing of the data, such as radial lens and geometric distortions, fixed pattern noise and the influence of harmonics in the modulation waveforms. Issues such as scattering and mixed pixel effects are scene dependent, and compensating for these effects is an active area of research.

Chapter 3 Range Imaging System Development

The range imaging system created at the University of Waikato had a number of desirable features at the time this research began. In particular it had a one megapixel camera, which is relatively high for this type of system; high modulation contrast; sub-millimetre depth precision and a large range of operating modes (Dorrington et al., 2007). The major disadvantages, particularly for mobile and low-cost applications were the relatively large size of the system, its high power consumption and slow offline processing.

This chapter begins by describing the Waikato ranger, and the improvements to address some of the disadvantages mentioned above, in particular the incorporation of real-time depth image processing in the embedded hardware (Jongenelen et al., 2009a). To reduce the cost, size and power requirements, an entirely new system has been designed and constructed as a significant part of this research (Jongenelen et al., 2009b). The development of the new Victoria University range imaging system is described, which is significantly more compact, uses less power, and has a similar versatility in terms of operating modes.

3.1 The University of Waikato Range Imager

The initial hardware setup of the Waikato University range imaging system is shown in Figure 3.1. The system illuminates the scene using four 80 mW laser diodes amplitude modulated at a programmable base frequency, *f*. The camera used is a DALSA Pantera 1M60 digital video camera (DALSA Corp., 2004) capable of frame rates up to 57 Hz with a resolution of 1024×1024 pixels, or 220 Hz with a resolution of 128×128 pixels (8×8 binning). To achieve the high frequency shuttering, the camera is coupled to a Photek MCP125 Image



Figure 3.1: Block diagram of the Waikato Ranger before improvement.

Intensifier (Photek Ltd., 2001) with the modulation signal applied to its photocathode.

The system is configured for heterodyne operation where the modulation signals driving the shuttering of the image intensifier and the light source are set to have a small (1 - 10 Hz) frequency difference. In this setup it is imperative that these modulation signals are frequency locked with approximately nine orders of magnitude stability. Furthermore, in order to calculate absolute range measurements it is necessary that the phase difference is known at any point in time during a frame capturing sequence.

To meet these requirements, the system features a circuit board containing three Analog Devices AD9952 Direct Digital Synthesiser (DDS) ICs (Payne et al., 2005). An advantage of these ICs is their ability to be synchronised to a single clock, ensuring minimal phase error between devices. Each IC has an internal 400 MHz system clock that is generated by multiplying a 20 MHz reference from a temperature-compensated crystal oscillator. To generate a desired output frequency, the DDS is programmed with a 32-bit frequency tuning word, *FTW*, using the Serial Peripheral Interface (SPI) of an Atmel 89LS8252 microcontroller connected through a RS232 serial link to a PC. The frequency tuning word is calculated as

$$FTW = \frac{2^{32} f}{f_{SYS}},$$
 (3.1)

where f is the desired frequency and f_{SYS} is the 400 MHz DDS system clock frequency.

The three sinusoidal frequencies generated by the DDS are used for:

- 1) modulating the transmitted light, f,
- 2) modulating the shutter of the camera sensor, $f + f_B$ and
- 3) triggering the frame rate of the camera, f_s .

Each channel has a resolution of 0.093 Hz as calculated by Equation (3.1) with a *FTW* of one. The maximum frequency is limited to 160 MHz by a low-pass filter on the output. An additional CMOS output is provided by the DDS to indicate when all three output signals are in phase. This is necessary to calculate absolute range data as opposed to only relative range (Dorrington et al., 2007).

The frequency provided by the DDS board for the camera trigger input is multiplied by a constant, giving $m.f_s$. This reference signal is passed into a digital counter to divide the frequency back down to the sampling rate, reducing jitter arising from converting the DDS sine wave output to the digital CMOS input required for the camera. A Xilinx Spartan 2 Field Programmable Gate Array (FPGA) is used to divide the frequency of the 3rd channel back down to f_s . In conjunction with the microcontroller on the DDS board and a pair of Digital to Analogue Converter (DAC) ICs, the FPGA also controls the analogue gain of the laser diodes and image intensifier. In the case of the image intensifier, this is used to ensure that the pixels of the camera do not saturate. For the laser diodes it allows them to be initially turned on over a period of several minutes, protecting them from damage caused by overheating.

The PC provides control of the modulation frequency and beat frequency through the RS232 connection with the DDS board. Frames captured by the camera are transferred to the PC via the industry standard Camera Link interface (PULNiX America Inc., 2000), and subsequently processed offline.

3.2 Improving the Waikato Ranger

One improvement to be added to the existing Waikato Ranger is to enable realtime range image processing, the results of which were published in (Jongenelen et al., 2008). The requirement was to construct a hardware design to retrieve the raw image frames from the camera and calculate the distance information in realtime. Processed frames can then be displayed in real-time on a standard Video Graphics Array (VGA) monitor, transferred to a higher order processor (for example the main CPU of a mobile robot) or transferred to a PC for further analysis.

An FPGA has been chosen for this design in preference to a serial processor as it can perform image processing in addition to many other general-purpose digital logic functions within the same package. This feature is particularly useful for interfacing the FPGA to a wide variety of devices each requiring a unique mode of communication, for example the Camera Link interface of the DALSA camera, the 24-bit colour interface to a VGA monitor and the Ethernet port of a PC. The incorporation of an FPGA in the design also simplified the transition to all-digital generation of the modulation signals, as is described in Section 3.3.

The target FPGA for this design is an Altera EP2S60 Stratix II (Altera Corp., 2007) mounted on an Altera Nios II Stratix II (RoHS) Development Kit. Key features that this board offers are:

- Stratix II EP2S60 FPGA
 - 144 18-bit × 18-bit hardware multipliers,
 - o 12 Phase-locked-loops (PLLs),
 - o 2,485 kbits block Random Access Memory (RAM),
 - o 60,440 equivalent logic elements;
- Nios II soft processor core (Altera Corp., 2008a) with 32 MB Double Data Rate Synchronous Dynamic RAM (DDR SDRAM);
- 10/100/1000 Ethernet port;
- Joint Test Action Group (JTAG) programming and debugging port;
- RS232 serial port;
- Two general purpose 41-pin expansion headers.

The Pantera 1M60 camera has a resolution of 1024×1024 pixels, however for this design the camera is operated in 8×8 binning mode reducing the resolution to 128×128 pixels. The first effect of this is that the sensor gain is increased, albeit at the expense of spatial resolution. A secondary consequence is that the maximum frame rate is 220 Hz in this mode, which is an excellent feature for real-time range imaging. Thirdly, the reduced resolution reduces the RAM required for the real-time processing algorithm, allowing the entire design to be implemented using the on-chip RAM resources of the Stratix II FPGA. For further details of the hardware processing algorithm refer Chapter 6.

Figure 3.2 gives an overview of the improved system including connections to the DALSA camera, the PC and a standard VGA monitor for real-time display. The PC is no longer involved in the processing of the depth images and is included purely as a means for long term storage and analysis. All of the hardware instantiated within the FPGA except for the Nios II CPU has been written using the Very-High-Speed Integrated Circuit Hardware Description Language (VHDL). Additional boards have been produced and plugged in to the expansion headers to drive the VGA monitor and interface to the camera using the Camera Link interface. Further detail on these boards and their integration into the system are given later in this section.

Frames are captured by the DALSA camera and transferred directly to the FPGA through the Cameralink interface. These are processed in real-time with resultant frames temporarily stored on the FPGA. These frames are output to the VGA monitor and retrieved by the Nios II CPU to be transferred out to the PC via a Transmission Control Protocol / Internet Protocol (TCP/IP) connection on the Ethernet port. Control registers for the ranging process are configured through the Nios II CPU via either the JTAG debugging interface or the TCP/IP connection. Also shown in Figure 3.2 is the logical block for controlling the gain of the modulation signals to the image intensifier and laser diodes.



Figure 3.2: Waikato Ranger and Altera Stratix II Development Kit interconnections.

3.2.1 VGA Monitor Driver

The requirements for the VGA monitor driver circuit are to convert parallel digital data into the analogue red, green and blue channels required by a standard VGA Additional digital signals are required to drive the horizontal and monitor. vertical synchronisation signals which determine the refresh rate of the monitor. To meet these requirements, a circuit based on the Analog Devices ADV7123 DAC (Analog Devices, 1998) has been constructed. This DAC converts three parallel 10-bit digital inputs into three analogue channels for the red, green and blue signals respectively. The circuit is capable of driving a display with up to 1600×1200 pixels with a 100 MHz pixel clock and is designed to plug into one of the 41-pin expansion headers of the Stratix II Development Kit. All digital inputs to the DAC IC and the vertical and horizontal synchronisation signals are driven directly by the FPGA, where a VHDL module has been written to provide correct timing for a display resolution of 640×480 pixels. There are many resources available on the internet regarding VGA signal timing - one resource in particular that has been helpful for this design is the Altera DE2 Development Kit User Manual (Altera Corporation, 2008b). The schematic for the VGA Monitor Driver can be found in the appendix.

3.2.2 Camera Link Interface

The Camera Link Standard (PULNiX America Inc., 2000) is a specification for camera manufacturers to implement, with the goal to reduce the costs of interfacing high-speed digital camera products to frame grabbers. The standard defines a number of features including the use of Low Voltage Differential Signalling (LVDS) and the pinout for the 3M MDR26 26-pin connector (3M, 2005). In order to retrieve image frames directly from the DALSA camera, a circuit has been constructed to implement the receiver end of the industry standard Camera Link interface.

The DALSA implementation of the Camera Link Standard (DALSA Corp., 2000) used by the DALSA Pantera 1M60 camera utilises the base configuration defined in the standard, requiring a single connector and single National Semiconductor Channel Link IC (National Semiconductor, 2004). The role of the National Semiconductor DS90CR286 Channel Link receiver IC is to convert four serial LVDS channels operating at 280 MHz each into 28 single ended CMOS channels operating at 40 MHz. This accommodates all of the signals required for streaming pixel data from the camera to the frame grabber, or in this case, to the FPGA. The standard also accommodates communication in the direction from the frame grabber to the camera through the use of four predefined LVDS camera control (CC) signals, CC1 to CC4. The 1M60 camera uses only one control signal, CC1, which is used as the external trigger to control the camera integration period. Two-way Universal Asynchronous Receiver/Transmitter (UART) serial communication is also accommodated so the user or frame grabber has access to various control parameters within the Pantera 1M60 camera.

For this project, a printed circuit board has been constructed to establish communications between the DALSA camera and the Stratix II FPGA. The board hosts a National Semiconductor 28-bit Channel Link IC for receiving image frames. This IC has been chosen as it is the serial-to-parallel converter recommended by the Camera Link specification. A National Semiconductor DS90LV047 LVDS Driver IC (National Semiconductor, 2003) is used for transmitting the camera control signal and a Texas Instruments SN75LVDS179 (Texas Instruments, 2003) enables two-way LVDS communication for the serial

UART. The board has two BNC connectors: one for interfacing the FPGA to the DDS chip to synchronise the timing of the camera integration period, the second is spare and is used for displaying debugging signals from the FPGA. Four additional spare connections are routed to 2-pin connectors and a fifth configured as a push button input – also for debugging. The schematic for the Camera Link Interface circuit is provided in the appendix.

3.3 The Victoria University Range Image System

As previously mentioned, the main limitations of the Waikato Ranger were its relatively large size, high power consumption and slow processing speed. A new system is desired that is more compact and has power requirements that are low enough for the system to be powered by the batteries of a mobile robot. It is intended that the system will have comparable measurement performance to other commercial compact ToF camera systems in terms of precision and accuracy, however, ideally it will be much more versatile in its operating modes.

The existing image intensifier tube and associated power supply units are the primary reason for the large physical dimensions. New image sensor ICs have recently become available where the high-speed shuttering mechanism is built-in to the chip. By utilising a sensor of this type the image intensifier is no longer required. In Section 2.4 a number of complete ToF camera systems were presented that are commercially available as of July 2010. At the beginning of this design project in June 2008 many of these systems were not yet produced, or the sensor IC was not available for purchase. One sensor that was available for purchase in low quantities suitable for a prototype was the Photonics PMD19k-2 from PMD Technologies (PMD Technologies GmbH., 2008). This sensor is the core component of the new design. More details are provided in Section 3.3.2.

It is also desirable for the new system to perform distance calculations in realtime, and to incorporate a mechanism for changing operating parameters by software. The ability to finely control the phase and frequency of the modulation signals in real-time is highly desirable for the system to act as a test platform for complex modulation schemes, for example dual-frequency modulation (Jongenelen et al., 2010a) and harmonic cancellation techniques (Payne et al., 2008b). As discussed earlier, an FPGA has been chosen as a suitable device to meet the requirements for real-time distance processing. For this new system, it is also required that the FPGA can drive the modulation signals, thereby removing the requirement for the DDS ICs. A number of FPGA devices were considered from the two leading FPGA manufacturers: Xilinx Corp. and Altera Corp. Previous work by (Payne et al., 2008b) demonstrated the usefulness of the advanced features of the Stratix II family of FPGAs produced by Altera Corp., specifically that the phase-locked-loop resources have the ability to reconfigure the phase and frequency of their outputs in real-time. This feature is excellent for satisfying the requirement to produce complex modulation patterns. The more advanced Altera Stratix III FPGA (Altera Corp., 2008c) mounted on the Altera Stratix III Development Kit has been chosen for this design, as it is expected this will also meet the needs of future design iterations.

The newly constructed Victoria University range imaging system comprises a number of interconnected components as shown in Figure 3.3. Modulation signals are shown as narrow black arrows, control signals as wide black arrows and the path of image frame data shown as white arrows. Custom designed boards include:

- Illumination board for mounting and driving laser illumination;
- PMD Daughterboard for hosting the PMD19k image sensor;
- PMD Mainboard for hosting the PMD Daughterboard and support electronics;
- VGA and Ethernet Physical and Media Access Controller (PHY/MAC) board for real-time image display and data transfer;
- Altera Stratix III Development Kit to provide modulation signals, realtime distance processing, and high-level timing and control.

The following sections describe each of these components in more detail.



Figure 3.3: Top level layout of the Victoria University range imaging system. The system is based around a PMD19k image sensor and an Altera Stratix III development kit. The FPGA drives the illumination and image sensor modulation, retrieves frame data from the sensor, processes it in real-time, and displays it on a VGA monitor or transfers the data to a PC.

3.3.1 Illumination Board

The requirements for the illumination board are to host an illumination source capable of modulating a high intensity output at frequencies up to 80 MHz. Earlier systems, for example (Carnegie et al., 2005), used a bank of light emitting diodes (LEDs) to provide illumination. However, the response time of the LEDs limited the maximum frequency to approximately 20 MHz. Laser diodes were determined to be an excellent alternative to LEDs, offering bandwidth in excess of 100 MHz, and higher illumination intensity (Payne, 2008).

This system uses a bank of 16 laser diodes mounted in a circular pattern around the image sensor lens. The purpose of the circular arrangement is to allow for an approximation that the illumination is a point source originating from the centre of the image sensor, thereby reducing geometric distortions in the distance measurement. There are two independent banks of eight diodes each: one set of red 658 nm wavelength the other with near infra-red (NIR) 808 nm wavelength. These are independently driven in a controlled current configuration giving a total continuous wave optical output power of 800 mW. It was initially anticipated that the use of multiple illumination wavelengths would make the system more robust when imaging scenes with objects of varying reflectivity. However, this project did not fully explore the merits of this feature, and in practice all 16 laser diodes are driven equally. The modulation signals provided by the Stratix III FPGA are provided through the PMD Mainboard via a pair of LVDS buffer ICs, discussed in Section 3.3.3.

Modulation signals connect to the board via a Serial Advanced Technology Attachment (SATA) connector, and are converted from LVDS to single ended CMOS voltage levels with a pair of Texas Instruments SN65LVDT2 Receiver ICs (Texas Instruments, 2002). The 3.3 V output is level shifted to 5 V by a group of ON Semiconductor MC74VHC1GT14 ICs (ON Semiconductor, 2004) to act as the enable signal for the diode driver circuit (discussed below). The illumination board also features a Microchip Technology PIC16F684 microcontroller (Microchip Technology Inc., 2004) whose purpose, in combination with a pair of Analog Devices AD5311 digital-to-analogue converters (Analog Devices, 2007), is to control the current through the laser diodes. Passive low-pass filters on each of the modulation signal inputs (red and NIR) are used to create an analogue signal that is related to the duty cycle of the modulation signal, which can be sensed by the internal Analogue to Digital Converters (ADCs) of the microcontroller. To prevent damage to the diodes, the microcontroller first checks that the duty cycle of the modulation signal is not too high and then begins a warm-up period where the current ramps up in a slow, predefined manner up to the maximum operating current. If the modulation stops, detected by a significant drop in duty cycle, the microcontroller automatically shuts off the current and returns to the warm-up phase.

Figure 3.4 shows the driver circuit diagram for one pair of laser diodes. The IC Haus iC-HK Laser Switch ICs (iC-Haus, 2004) drive the laser diodes in a currentcontrolled configuration. With the EN2 pin fixed at 5 V, the modulation signal which controls EN1 sets the current to a predetermined low value when 0, and a high value when 1. These high and low values are determined by a combination of the CI voltage input and resistors RK1 connected to AGND1 and RK2 connected to AGND2. The red laser diode used has a 658 nm wavelength, maximum continuous operating current of 200 mA, and threshold current of 80 mA. With the CI voltage set by the AD5311 DAC to a maximum of 2.7 V (after the warm-up period), the low current is set to the threshold current with *RK2* chosen as 10 Ω . To achieve the additional 120 mA needed to reach the high current of 200 mA, RK1 is chosen as 5 Ω . The infra-red laser diode used has a 808 nm wavelength, maximum continuous operating current of 240 mA and threshold current of 70 mA. Again using a fixed CI voltage of 2.75 V, RK1 and *RK2* are chosen as 10 Ω and 3 Ω . More detailed schematics for the illumination board can be found in the appendix.



Figure 3.4: Laser diode driver circuit based on the iC-HK Laser Switch.

3.3.2 PMD Daughterboard

The photonic mixer device (PMD) is an image sensor with an array of "smart pixel" elements (Ringbeck et al., 2007). The basic principle is that each pixel has two charge storage capacitors, A and B. The movement of generated charge carriers from the CMOS photodetector is controlled by the modulation signal toward either A or B where charge is accumulated during the integration period, producing measurable voltages, V_A and V_B , across each capacitor. With constant illumination and a 50% duty cycle modulation signal, equal charge will accumulate in each capacitor. If the illumination is also modulated at the same frequency there will be a noticeable difference between V_A and V_B dependent on the phase difference between the received illumination and the sensor modulation signal.

As previously mentioned, a PMD19k-2 version image sensor is used. It features a 160×120 array of "smart pixels" grouped into four independently modulated blocks of 40×120 pixels, with the modulation signal to each block representing a capacitive load of 250 pF. Figure 3.5 shows a plot of sensor demodulation contrast versus the high voltage level of the modulation signal. It is desirable to maximise demodulation contrast, and this is achieved using a high voltage level between 2.0 and 2.5 V.

Figure 3.6 shows a plot of sensor demodulation contrast versus modulation frequency. Maximum contrast is obtained at low frequencies, however higher frequencies are desirable for maximising distance measurement precision as was discussed in Chapter 2. The relationship between frequency and contrast is relatively linear up to 40 MHz, beyond which contrast begins to degrade more rapidly.

A schematic diagram of the PMD Daughterboard is provided in the appendix. This board hosts the image sensor and contains only a bare minimum of passive components required to set various voltage reference levels and filter the power supply. All modulation, analogue video (V_A and V_B) outputs and control signals are routed via a connector to the PMD Mainboard. This allows design changes to be made to the PMD Mainboard without directly affecting the image sensor.

Mounted on the PMD Daughterboard is a cage plate designed to hold a C-Mount lens. For this project a Goyo Optical GM31614MCN 16 mm focal length lens with manual focus and iris is fitted (Goyo Optical Inc., 2003). This gives a fieldof-view (horizontal × vertical) of $22.2^{\circ} \times 16.5^{\circ}$. The lens has a 27 mm filter thread to which a variety of optical filters can be attached. A LP645 filter from Midwest Optical Systems is used which passes over 90% of the reflected light of the red and NIR laser diodes while excluding shorter wavelength visible background light that the image sensor would otherwise detect.

3.3.3 PMD Mainboard

The PMD Mainboard hosts all the electronics for interfacing the FPGA to the image sensor and illumination board. This includes:

- High-speed drive of the high-capacitance sensor modulation inputs;
- Conversion between LVDS signalling with the FPGA and single ended CMOS signalling with the image sensor;
- LVDS buffering for relaying the modulation signal to the illumination board.
- Analogue-to-digital conversion of the image sensor video output streams;



Figure 3.7: PMD Mainboard sensor modulation driver circuit.

The greatest challenge with the PMD Mainboard is to modulate the large capacitive load of the image sensor modulation pins at up to 40 MHz. The Intersil EL7156 ultra-high current pin driver was recommended by PMD Technologies for use with their range of sensors. However, the EL7158 pin driver IC (Intersil Americas Inc., 2007) was chosen for this design as it features a shorter rise and fall time for a 2000 pF load of 12 ns compared to 15 ns for the EL7156. Figure 3.7 shows the schematic for driving one set of the image sensor modulation pins using an EL7158 pin driver IC. The OUT pin of the EL7158 is internally connected to the VH or VL pins depending on the status of the IN pin, which is driven by the FPGA. The isolation of the pin driver output from its power supplies allows VH and VL to be set independently. VL is connected to GND (0 V), and VH is provided by the output of an Analog Devices AD8531 operational amplifier. The op-amp is configured in a voltage follower arrangement, with the output voltage set by the voltage divider network comprising resistors R1 and R2. The purpose of the op-amp is to provide high current at the specified voltage, and three capacitors of 10 µF, 100 nF and 10 pF keep the level stable while the pin driver switches at high frequencies. The values of R1 and R2 are chosen to be 1 k Ω and 680 Ω respectively in order to give a voltage level for VH of 2.02 V, which provides good demodulation contrast as per Figure 3.5. The power supply voltage for the EL7158 has been chosen as 12 V as this gives the highest frequency performance for the pin driver, as per its datasheet.

Analogue-to-digital conversion of the image sensor video outputs is performed using an Analog Devices AD9826 16-bit imaging signal processor (Analog Devices, 2001). This ADC has a number of control registers for setting various operating modes that are configured through a three wire SPI interface. In order to read control register data back from the ADC using unidirectional LVDS, the bidirectional SPI data line must be split into two separate read only and write only lines. This is achieved using a Texas Instruments SN74LVC125 buffer gate with 3-state outputs (Texas Instruments, 2007a), where the direction is controlled by a separate write enable signal. Further information about the configuration of the ADC control registers can be found in the PMD Interface description of Section 3.3.5.

All communication with the FPGA uses LVDS signalling as it reduces radiated electromagnetic noise and is more robust at high speeds than single ended CMOS. Conversion between single ended CMOS to LVDS, required for the datapath from the image sensor outputs and the ADC to the FPGA, is implemented by a pair of Texas Instruments SN65LVDT388 high-speed differential line receivers (Texas Instruments, 2007b). Conversion from LVDS to single ended CMOS, required for the control and timing signals sent from the FPGA, is performed by a pair of Texas Instruments SN65LVDS389 high-speed differential line drivers (Texas Instruments, 2001). The modulation signals are also transmitted using LVDS, and the conversion for the image sensor is performed by a Texas Instruments SN65LVDT9637 receiver IC (Texas Instruments, 2007b). The PMD Mainboard buffers the LVDS modulation signals to the illumination board using a pair of Texas Instruments SN65LVDT100 differential line repeaters (Texas Instruments, The PMD Mainboard also features a 160 pin high-speed mezzanine 2004). connector for high-bandwidth communication with the Stratix III Development Kit. A schematic diagram of the PMD Mainboard can be found in the appendix.

3.3.4 VGA DAC and Ethernet PHY/MAC

Raw and processed image frames are temporarily stored in RAM on the Stratix III Development Kit. Longer term storage is possible through a 10/100 Ethernet link to a PC where the image frames can be saved to a file. This is handled by a DAVICOM DM9000A Ethernet Controller (DAVICOM Semiconductor Inc., 2006). This controller has been selected primarily for its ease of use, as it incorporates both the physical and MAC layers into a single device. The schematic for the Ethernet controller circuit can be found in the appendix. All digital I/O pins of the DM9000A IC are routed via a connector to the FPGA. The circuit also has a group of spare connections to the FPGA for debugging or incorporating additional hardware at a later date.

In addition to transferring frames to a PC, they are also displayed in real-time by a standard 640×480 resolution VGA monitor. This circuit, based on the Analog Devices ADV7123 DAC, is identical to that used for the Waikato University range imaging system in Section 3.2.1.

The VGA DAC and Ethernet controller reside together on a custom built Printed Circuit Board (PCB), and are connected to the Stratix III Development Kit through a 160 pin high-speed mezzanine connector.

3.3.5 Stratix III Development Kit

The Stratix III Development Kit hosts an Altera Stratix III EP3SL150F1152 FPGA (Altera Corp. 2008c). This has been selected primarily for the easily reconfigurable phase-locked-loop (PLL) resources. Features of this FPGA include:

- 384 18-bit × 18-bit hardware multipliers;
- 8 PLLs;
- 5,499 kbits block RAM;
- 142,500 equivalent logic elements.

Each PLL can provide up to 10 phase synchronous output signals of frequencies up to 700 MHz. The FPGA is also responsible for interfacing to the PMD19k image sensor, calculating the distance data in real-time, and communicating with a PC for user control of the system. Image data, either raw or processed is displayed on a VGA monitor and transferred to a PC connected via Ethernet.

All blocks internal to the FPGA are programmed using VHDL, except the Nios II processor where an intellectual property core supplied by Altera is used. VHDL simulation and verification is performed using Riviera software by Aldec Inc. Synthesis of the VHDL for the target FPGA is performed using Quartus software by Altera Corp.

Reconfigurable Phase Locked Loop Blocks

The main advantage of these PLLs is the ability to reconfigure the phase, frequency and duty cycle of each output channel in real-time. This is an especially useful feature for testing the system under a large number of operating conditions. The ability to change the modulation signal both between frames as well as during the sensor integration time provides an excellent platform for further research into various range imaging modulation schemes.

The PLL has a voltage controlled oscillator (VCO) with a frequency, f_{VCO} set by

$$f_{VCO} = \frac{f_{IN}c_m}{c_n},\tag{3.2}$$

where f_{IN} is an input frequency of 50 MHz from the onboard crystal oscillator, and c_m and c_n are configurable integers. The PLL output frequency, f_{OUT} is determined by

$$f_{OUT} = \frac{f_{VCO}}{c_{HI} + c_{LO}},\tag{3.3}$$

where c_{HI} and c_{LO} are configurable integers. The duty cycle can therefore be controlled by the ratio between c_{HI} and c_{LO} . The output phase is set by incrementing the phase in discrete steps, with each step giving a phase shift in radians of

$$\phi_{step} = \frac{2\pi}{8(c_{HI} + c_{LO})}.$$
(3.4)

For example, an output frequency of 40 MHz can be achieved using $c_m = 24$, $c_n = 1$, $c_{HI} = 15$ and $c_{LO} = 15$, giving 240 discrete and equally spaced phase steps before repeating.

The design as shown in Figure 3.3 incorporates two reconfigurable PLL blocks, and each block has two outputs: one for the illumination modulation and the other for the image sensor modulation, which has the same frequency but variable phase offset. The purpose of incorporating two PLL blocks is so the design can generate two modulation frequencies which can be phase stepped independently, as is required to enable dual frequency modulation. Dual frequency modulation is discussed further in Chapter 5. The design can be scaled to incorporate up to six reconfigurable PLLs for testing multiple frequency modulation with up to six frequencies.

Modulation Control Block

The role of this block is to select which of the modulation signals generated by the reconfigurable PLLs reach the illumination board and image sensor.

In single modulation frequency operation only PLL 1 is used. The illumination modulation signal is constantly modulated as this is necessary to keep the laser diodes at a reasonably constant temperature. The image sensor modulation signal is only routed through the multiplexer while the image sensor is integrating; else it is fixed at zero.

In dual modulation frequency operation both PLLs are used. Again, the illumination modulation signal is constantly modulated at one of the two frequencies to keep the laser diodes at a consistent operating temperature. The difference is that the modulation control block also incorporates a mechanism for changing the output frequency during the integration time based on a specified ratio. The same mechanism is used to set which output frequency drives the image sensor.

The modulation control block can also turn off either or both of the NIR or red sets of laser diodes, which is a suitable feature for testing the two sets independently.

PMD Interface

This block is responsible for interfacing to the PMD19k image sensor. It has three main tasks:

- 1. Configure the ADC on the PMD Mainboard;
- Drive all the control signals for the PMD sensor to set timing for integration and frame readout;
- 3. Register pixel data from the image sensor, and assign an appropriate address to each received data word.

The PMD Mainboard ADC is configured through a SPI interface. On system reset, the ADC is configured for 2-channel Sample-Hold Amplifier (SHA) mode. During operation, the user can also reconfigure the analogue gain and offsets for each channel. With an analogue input range of 0 to 4 V, the digital value associated with an input voltage, V_{in} , is given by

$$D_{out} = \left(Vin + 0.3\frac{O(-1)^{S}}{255}\right) \cdot \frac{6}{1 + 5\left[\frac{63 - G}{63}\right]} \cdot \frac{2^{16}}{4},$$
(3.5)

where O is the 8-bit offset value, S is the sign of the offset and G is the 5-bit gain value. Any result less than 0 is returned as 0, and any result larger than 1023 is

returned as 1023. The minimum analogue voltage provided by the PMD19k image sensor is 0.95 V, and the maximum is 3.15 V. To make full use of the ADC output range, the offset, $O(-1)^{s}$ is set to -255 and the gain, G, is set to 20. With these settings, the measured analogue voltage is given from D_{out} by

$$V_{meas} \approx \frac{D_{out}}{22,277} + 0.3$$
. (3.6)

The PMD interface provides all the timing and control signals for the PMD image sensor. This includes the pixel reset and hold lines to specify the sensor integration time and all the clocking signals for pixel data readout. The pixel readout clock has a maximum frequency of 8 MHz, producing a frame readout period, T_R , of approximately 2.5 ms.

As pixel data are converted by the ADC, the PMD interface samples the data at the pixel clock rate, and assigns a row and column address. The system clock is at least twice as fast as the pixel clock, so the PMD interface also provides a pixel valid pulse to signal when pixel data and addresses are valid.

Range Processor

The purpose of this block is to calculate distance and magnitude data from the raw pixel samples. The calculation is performed on a per-pixel basis in real-time at the rate of the pixel clock. The details of performing the distance calculation in hardware are discussed in Chapter 6.

Frame Output Buffer

The outputs from the range processor are subsequently stored in the frame output buffer, which has been configured from block RAM resources internal to the FPGA. It is possible to store up to four sets of data per pixel in the frame output buffer, which are typically configured to be the two raw pixel channels and the phase and amplitude data. The range processor is the only block that writes into the buffer, whereas there are two blocks that read from the buffer: the VGA display driver and the Nios II processor. To implement this, the read port is time multiplexed giving each process equal opportunity to read from the buffer, albeit with an extra clock cycle of latency and a halving of the bandwidth.

VGA Driver

The VGA display driver is responsible for transmitting timing signals to a standard VGA monitor along with image data read from the frame output buffer. The timing is configured for 640×480 display resolution with a pixel clock of 25 MHz and refresh rate of 60 Hz.

All four sets of data in the frame output buffer are displayed, with each RAM location mapped to a 2×2 pixel block on the monitor.

Nios II Processor and DDR2 SDRAM

The Nios II processor core (Altera Corp., 2008a) is a 32-bit embedded processor specifically designed to operate using the on-chip resources of Altera FPGAs. In this design it is primarily responsible for the communication between the range imaging system and an attached PC. The primary advantage of utilising this software defined processor core is that it can be customised to include a range of peripherals specific to this application using Altera's System on Programmable Chip (SOPC) Builder software.

Figure 3.8 shows a screenshot taken from the SOPC builder when creating the Nios II processor for this application. Memory available to the processor is a small 8 kB block of on-chip RAM to store reset and exception code, and two DDR2 SDRAM devices of 16 MB each. These DDR2 devices reside on the Stratix III Development Kit.

A 1 ms interval timer is included to simplify timed delay functions in software. PC communication is enabled through the DM9000A Ethernet controller and a

Conne			Module Name	Description	IRQ
C	\rightarrow	Ŧ	onchip_mem	On-Chip Memory (RAM or ROM)	
		Ξ	сри	Nios II Processor	
\vdash	\vdash		instruction_master	Avalon Memory Mapped Master	
	\succ		data_master	Avalon Memory Mapped Master	← ∖
11-	┝→		jtag_debug_module	Avalon Memory Mapped Slave	
11-	\rightarrow	Ð	ddr2_deva	DDR2 SDRAM High Performance Controller	
	\rightarrow	÷	ddr2_devb	DDR2 SDRAM High Performance Controller	
	\rightarrow	Ð	frame_received	PIO (Parallel I/O)	≻—0
	\searrow	÷	timer_1ms	Interval Timer	1—ל
	\rightarrow	Ð	jtag_uart	JTAG UART	≻-2
	\searrow	÷	dm9000a	DM9000A	≻-3
	\rightarrow	Ð	control	nios_slave_16	
	\rightarrow	÷	TL	nios_slave_16	
	\rightarrow	Ŧ	TR	nios_slave_16	
	\rightarrow	÷	BL	nios_slave_16	
	\hookrightarrow	Ŧ	BR	nios_slave_16	

Figure 3.8: Nios II processor configuration in Altera SOPC Builder.

JTAG UART. The Ethernet connection deals only with bulk data transfers to the PC. The JTAG UART is connected to a text input terminal on the PC to listen and respond to user commands.

There are five "nios_slave_16" modules, which are custom designed peripherals to interface to the frame output buffer (divided into TL for "Top Left", BR for "Bottom Right", etc. referring to their location on the VGA monitor) and control registers.

A one bit input called "frame_received" acts as an active high level triggered interrupt to indicate that a new image frame has been received and is available in the frame output buffer. More information about this interrupt can be found in Section 3.4.1.

Control Registers

The block of control registers function as a group of special function registers addressable by the Nios II processor. In conjunction with the JTAG UART of the Nios II processor, they provide a mechanism for the user to change a number of system parameters as well as automatic control of the system by the Nios II processor. Parameters associated with the various other blocks of the system that can be controlled are:

- Reconfigurable PLL Blocks;
 - VCO frequency;
 - Modulation frequency;
 - Modulation duty cycle;
 - Phase steps per cycle;
 - Initial phase offset/correction;
- Modulation Control Block;
 - Integration Ratio between frequency 1 and frequency 2;
 - NIR and red laser diodes on/off control;
- PMD Interface;
 - ADC configuration via SPI interface;
 - Total integration period;
 - Total frame period, hence measurement rate;
 - Pause / restart capturing;
- Range Processor;
 - Various parameters associated with the distance calculation, for example input gain, input offset, etc.;
 - Frame received interrupt and interrupt acknowledged.

3.4 Victoria Ranger Software

Software written for this project has targeted the embedded Nios II processor, along with a PC application to receive image frames, and MATLAB code to process and present the image frames and associated statistics.

3.4.1 Nios II Code

The Nios II processor is programmed using Altera's Nios II Integrated Development Environment (IDE) software. The language used is C. There are two tasks required of the Nios II processor: allow a user access to the control registers to configure operating parameters, and to enable bulk data transfer of image frames from the ranger to a PC. Figure 3.9 shows the basic structure of the Nios II application. The main loop of the program polls the JTAG UART interface and processes text commands from the user. The user provides these

commands through the JTAG UART terminal provided as part of the Nios II IDE, and are largely based on simple read and write commands of the control registers.

When a "frame_received" interrupt is generated, the interrupt handler routine transfers one or more of the data sets from the frame output buffer into the heap memory of the Nios II processor. It also increments a counter and acknowledges the interrupt. Within the main loop, in addition to checking the JTAG UART the



Figure 3.9: Flow diagram for Nios II application. Task is to process user commands from JTAG UART, and send image frames to PC via Ethernet.

frames received counter is compared to a counter of frames transmitted out of the Ethernet port. Any frames not yet sent are transmitted via the Ethernet using a simple User Datagram Protocol (UDP). Each packet sent, in addition to the frame data includes a row and frame number so that the listening application can organise the data correctly.

One specific text command from the user is to print the current operating parameters. In addition to printing these to the terminal, a specially flagged Ethernet packet is sent to the PC, so that the PC application can associate the incoming image data with the operating parameters used. The packet is flagged by setting the row and frame number to -1, as this will never occur amongst the normal data packets.

3.4.2 PC Ethernet Listener

The PC software to listen for Ethernet packets is programmed in Java using Netbeans IDE. This application listens to the Ethernet port of the PC, and filters out packets that do not relate to a pre-ordained port number and IP address. To enable this, the software makes use of a packet capture library for Java called JPCap. Data from the correct packets are dumped to a file for later processing, primarily in MATLAB.

An additional task for this application is to note the specially flagged Ethernet packets containing information about the system operating parameters. This information is written to a separate file, which is also used in subsequent MATLAB processing.

3.4.3 MATLAB Code

While range images can be processed and displayed on the VGA monitor in realtime, for many applications, particularly for characterisation and research it is desirable to process and analyse the data offline. MATLAB is used extensively for this project in a number of ways:

• Prototyping of range imaging algorithms

- Simulation of hardware implementation
- Processing and analysing range image data to characterise system parameters
- Compiling and presenting experimental results.

3.5 Summary

In this Chapter the hardware and software developments associated with this project have been described.

Initial progress was made towards improving the Waikato University range imaging system by enabling real-time range image processing. This allows the system to be used in a number of applications in which it would have otherwise been too slow to be of any use. Furthermore, it acted as a proof of concept for how range image processing can be done in hardware, paving the way for the development of a more practical range camera suitable for mobile robotics applications.

The most significant work however, has been in the design and construction of the new Victoria University range imaging system. The goal has been to make the system much more compact, while still retaining the configurability of the Waikato Ranger in terms of operating parameters such as modulation frequency, integration time and phase steps per cycle.

This has been achieved using a PMD19k image sensor from PMD Technologies GmbH, which satisfies the requirement for a compact system by performing onchip pixel modulation and removing the bulky image intensifier. A number of custom PCBs have been designed and constructed to host the PMD19k image sensor, and interface it to an Altera Stratix III FPGA. The advanced features of the reconfigurable PLLs of the FPGA make it an ideal platform for further research into time-of-flight range imaging, in particular, multiple modulation frequency range imaging (see Chapter 5) and hardware algorithms for range image processing (see Chapter 6).

Chapter 4 System Characterisation

One of the objectives of this research is to construct a system that is highly flexible, while exhibiting similar performance to commercially available ToF range imaging systems. This chapter investigates the performance of the new system in terms of distance measurement precision and accuracy with a number of software controlled operating parameters.

4.1 Laser Illumination Waveform Shape

In order to test the laser illumination output, a Thorlabs DET10A/M silicon based photodiode (Thorlabs Inc., 2006) has been fixed into place 1 mm in front of one of the 16 laser diodes. The output voltage of the photodiode is observed using a Tektronix DPO4054 digital phosphor oscilloscope (Tektronix Inc., 2005). Due to the imprecise position of the photodiode in front of the laser diode, absolute illumination intensity measurements are not practical with this configuration. However, due to the fixed position of the photodiode meaningful relative measurements of the waveform shape can be made, in addition to useful comparisons across a range of modulation frequencies.

Figure 4.1 shows a screen shot of the oscilloscope with a laser modulation frequency of 8 MHz. The green trace shows the voltage of the enable signal of the laser current driver IC, the yellow trace shows the photodiode output voltage, and the red trace shows a Fast Fourier Transform (FFT) of the data. Numerous statistics are also shown in the bottom left of the display, indicating the frequency as measured by the oscilloscope, amplitude, Root Mean Square (RMS) voltage and positive duty cycle. The most notable feature of the waveform shape is that it closely approximates a square wave, albeit with some overshoot on each of the rising and falling transitions.


Figure 4.1: Laser diode illumination waveform with modulation frequency of 8 MHz as measured by a photodiode.



Figure 4.2: Laser diode illumination waveform with modulation frequency of 40 MHz as measured by a photodiode.

As expected for a square wave, and as can be seen with the FFT, the odd order harmonics have relatively large magnitude. If the image sensor modulation also has these harmonics present, then the correlation waveform may be distorted, causing systematic errors in the distance measurements (Payne et al., 2010).

There is a delay of approximately 10 ns between the transition of the enable input signal and the transition of the illumination output. This equates to an error offset of 1.5 m in the distance measurement, though this error can be removed through calibration. It is also apparent that the steady state low illumination output does not fully reach zero. This can result in a DC offset in the correlation waveform, which can introduce errors in the distance measurements. A DC offset also reduces modulation contrast, leading to reduced measurement precision. The amplitude, that is the difference between the steady state low level and the steady state high level, is approximately 59 mV, and the RMS voltage is approximately 47 mV.

Figure 4.2 shows another screenshot from the oscilloscope, this time for a modulation frequency of 40 MHz. Due to the limited bandwidth of the laser diodes and the driver electronics, the output waveform appears more sinusoidal. There are second and third order harmonics of relatively significant magnitudes still present, as evident from the FFT, although these are much smaller than those present with 8 MHz modulation. The amplitude is also larger for the 40 MHz modulation frequency, as the laser diodes successfully reach down to zero volts.

4.2 Sensor Modulation Waveform Shape

Figure 4.3 shows the A and B modulation signals at the input to the PMD19k image sensor measured on the oscilloscope, with frequency set to 8 MHz. These signals drive all four banks of the image sensor. The purple trace shows the modulation A input, and the blue trace shows the modulation B input. When it comes time to process the images and calculate distance, it is the difference between the A and B channels that is of interest, therefore the FFT of the difference is shown by the red trace. The statistics shown at the bottom left of the screen are for the A modulation input; the statistics for the B modulation input are

not shown as they are very similar. Similarly to the illumination waveform, the shape is basically square with relatively significant odd order harmonic components. When correlated with the illumination waveform, it can be expected that the correlation waveform will look triangular. Depending on the number of samples used to determine the phase of the correlation waveform, these harmonics may be aliased onto the signal of interest, leading to systematic errors (Payne et al., 2008). The amplitude of 2 V shows that the signal is correctly reaching the power rail provided by the AD8532 op-amp as per Section 3.3. The duty cycle is close to 50%, suggesting the DC offset after the subtraction of the A and B pixel values will be reduced.

Figure 4.4 shows the same set of traces as Figure 4.3, this time with a modulation frequency of 40 MHz. The waveform shape is noticeably different, appearing more like a distorted sinusoid than a square wave. Third and fifth order harmonics are still present with relatively large magnitude, and second and fourth order harmonics are also present at a level that is more significant than for the 8 MHz waveform. These even order harmonics contribute a noticeable reduction in the duty cycle of the waveform, which can lead to asymmetric pixel integration and introduce a larger DC component after subtraction. The amplitude is also reduced because the signal does not fully reach 0 V or 2 V. This reduction in amplitude may lead to a reduced amplitude correlation waveform, thereby reducing distance measurement precision.



Figure 4.3: PMD19k sensor modulation waveform with frequency of 8 MHz. Purple and blue traces show the voltage of the A and B modulation inputs respectively. Red trace shows the FFT of their difference.



Figure 4.4: PMD19k sensor modulation waveform with frequency of 40 MHz. Purple and blue traces show the voltage of the A and B modulation inputs respectively. Red trace shows the FFT of their difference.

4.3 Waveform shape versus modulation frequency

Measurements of the illumination waveform have been repeated, this time with the photodiode positioned 200 mm in front of the image sensor on the optical axis of the lens. At this distance it is possible to observe the illumination power contributed by all 16 of the laser diodes.

Figure 4.5 shows the magnitude of the fundamental frequency for the photodiode and the sensor modulation input for frequencies from 1 MHz to 60 MHz. Note that the same y-axis is used for both, but the sensor modulation is measured in volts, and the photodiode output measured in millivolts. The magnitude of the illumination does not vary greatly across this frequency range, peaking at 1.3 mV for a frequency of 28 MHz and then steadily dropping off at approximately 0.018 mV for every 1 MHz. The most marked change from 1.25 MHz and 60 MHz modulation for the illumination output is the shape of the waveform, progressively changing from a square wave to a sinusoid.

The amplitude of the fundamental frequency component for the image sensor modulation is approximately constant at 1.85 V up to 28 MHz, at which point it begins to decrease. The action of the image sensor is to electronically mix the illumination and sensor waveforms so that an FFT of the mixed (or correlated)



Figure 4.5: RMS voltage for illumination measured by photodiode (mV) and PMD19k image sensor modulation input (V) versus modulation frequency. Also shown is an estimation of their correlation product and this product multiplied by frequency.

signal can be used to calculate the phase difference between them. Figure 4.5 also shows the product of the illumination and sensor modulation waveforms by the blue line (arbitrary units) indicating what can be expected of the amplitude of the correlation waveform, with peak amplitude at 24 MHz. Distance measurement standard deviation is proportional to the inverse of the correlation amplitude and the inverse of the modulation frequency. By multiplying the correlation amplitude by the modulation frequency and taking the inverse, an indication of the distance measurement standard deviation can be obtained. This is shown in the plot as the green line (arbitrary units) with a minimum (indicating best precision) at a modulation frequency of approximately 32 MHz.

4.4 Noise floor versus modulation frequency

In the following sections we refer to the precision of the system as a measure of the variation in the measurement over time, that is, the repeatability. This is distinct from the accuracy of the system, which is affected by systematic errors that can often be removed through calibration. Precision is measured as the 1-sigma standard deviation of each pixel over time, and the values given are the mean of the standard deviations for a 40×40 group of pixels.

One factor contributing to overall system precision is the imager noise floor. This is the noise of the pixel outputs under dark conditions with no modulation, where the main sources of noise are related only to the temperature, electrical conditions of the sensor, and the quantisation error of the ADC.

To measure the noise floor, a number of captures have been taken at various modulation frequencies with the lens cap on so that the illumination is still modulating but no light is reaching the sensor. The number of phase steps per measurement has been set to four, and the measurement rate set to 25 Hz with a total of 180 measurements taken. Figure 4.6 shows the standard deviation over the 180 measurements for the raw pixel intensity values output by the sensor on each of the A and B output channels. There appears to be a difference in the noise level between the two channels of approximately 2.5% irrespective of modulation frequency. Aside from the larger noise level for the 1.25 MHz modulation



Figure 4.6: Standard deviation of pixel channel outputs for capture with lens cap on and very short integration time.

frequency, there appears to be no significant relationship between the noise floor and the modulation frequency.

4.5 **Pixel amplitude versus modulation frequency**

As mentioned in the previous section, the function of the image sensor is to electronically mix the illumination waveform with the image sensor modulation. To observe this a static scene is set up with a single white coloured box in the centre of the imaging area. This provides a flat target area 560 mm high, 470 mm wide and 2.63 m from the image sensor lens. The true object distance has been measured using a Leica DISTO distance meter with a manufacturer specified accuracy of ± 1.0 mm (Leica Geosystems, 2010). The scene has been imaged using modulation frequencies from 1 MHz to 60 MHz, with 120 phase steps per measurement and a measurement rate of 0.208 Hz (40 ms per frame). Modulation duty cycle has been fixed at 50% unless otherwise stated.

A 40×40 region of pixels at the centre of the image have been used to provide statistics for what is reasonably assumed to be a flat target with a constant distance from the imaging system. Both pixel output channels of the PMD19k sensor have been captured and processed offline in MATLAB. When referring to the raw pixel data, it is in fact the difference between the two channels that is being referred to.

Figure 4.7 shows a colour photo of the test scene, and Figure 4.8 shows a false colour image of the test scene where colour represents distance. Signal amplitude

is shown in Figure 4.9, where the rectangular region indicates the group of pixels used to provide statistics. Also present in the image is an office chair, some shelving, and the far wall, none of which are used for any measurements.



Figure 4.7: Simple static test scene used for characterising tests. In the centre of the field of view is a single white rectangular target, positioned 2.63 m from the range imaging system.



Figure 4.8: Range image of test scene, where distance is represented by colour.



Figure 4.9: Amplitude image of test scene. The rectangle in the centre indicates the 40 × 40 group of pixels analysed for various statistics.

Figure 4.10 shows, for a number of frequencies, the mean pixel value for the 40×40 region of pixels at the centre of the image across 120 phase steps of the image sensor. This shows the correlation between the returned illumination and the sensor modulation waveforms. It can be seen that as the frequency is increased the waveform progresses from a triangular shape to a sinusoid, and the signal amplitude decreases.

Figure 4.11 shows a semi-logarithmic plot of the amplitude of the fundamental



Figure 4.10: Mean pixel output for central target region versus sensor modulation phase offset for a range of modulation frequencies.



Figure 4.11: Semi-logarithmic plots of amplitude of fundamental frequency component versus modulation frequency (blue line, left axis) and the inverse of amplitude and frequency product (green line, right axis).

frequency component (blue line, left axis) versus the modulation frequency. The pattern is similar to that for the estimated correlation amplitude shown in Figure 4.5, although at the lower frequency end of the spectrum the amplitude is the highest. As an indication of expected distance measurement standard deviation, the inverse of the product of the correlation amplitude and the frequency is also shown (green line, right axis). This shows a predicted minimum standard deviation, hence best precision, at 36 MHz.

4.6 Phase versus modulation frequency

Using the data produced above, the phase of the correlation waveforms has been calculated, and compared to the expected phase for a 2.63 m target object. Figure 4.12 shows the measured phase for the range of tested modulation frequencies compared to the expected phase, as well as the phase error at each frequency. For the operating range of 2 to 42 MHz modulation, a cubic polynomial can be used to model the error with a standard deviation of 43 milliradians for the residuals.



Figure 4.12: Measured and expected phase versus modulation frequency for a flat target 2.63 m from the imaging system. Also shown is the error in the phase measurement, with a cubic polynomial fit based on frequencies 2 to 44 MHz.

4.7 Phase and Distance Precision

The largest noise source in ToF range imaging is photon shot noise (Lange and Seitz, 2000). The process of photons arriving at the sensor and being converted to electrons follows a Poisson statistical process. For a large number of photons where illumination power is sufficiently high and integration time is sufficiently long, this distribution can be modelled as a Gaussian distribution, where the standard deviation indicates the spread. Software controlled operating parameters of the system that influence the precision are the modulation frequency, f, the integration time, T_I , and the number of frames per measurement, N.

4.7.1 Precision versus modulation frequency

As suggested in Section 4.5, the modulation frequency affects the amplitude of the correlation function of the sensor. New captures of the scene used in Section 4.5 were acquired with the same set of modulation frequencies, N of four and a measurement rate of 25 Hz (10 ms per frame). A total of 360 measurements are taken over a 28.8 s period to provide a measure of the precision. Figure 4.13 shows semi-logarithmic plots of the modulation frequency versus the phase and distance measurement standard deviations, σ_{φ} and σ_d respectively. The overall trend is similar to that of Figure 4.11 albeit inverted, with a noticeable anomaly between 21 and 23 MHz. If we initially ignore the anomaly, the amplitude is generally decreasing as the modulation frequency is increased from 1 MHz to 36 MHz, the phase standard deviation is increasing relatively linearly at a rate of 0.082 milliradians per MHz. Throughout this range, the linear increase in phase standard deviation is offset by the increase in frequency, resulting in a general trend of decreasing distance measurement standard deviation, with a minimum of 4.08 mm at 36 MHz. This corresponds to 0.155% of the actual target distance. Beyond 40 MHz, phase standard deviation begins to increase dramatically as the image sensor modulation driver can no longer drive the high capacitive load above this frequency, reducing the amplitude of the correlation waveform and distorting the shape.



Figure 4.13: Modulation frequency versus Phase standard deviation, σ_{φ} , (blue line, left axis) and distance measurement standard deviation, σ_d (green line, right axis).

The dramatic degradation in precision between 21 and 23 MHz is not due to a reduction in amplitude, as this would have been evident in Figure 4.11. It must therefore be due to other factors; however, due to time constraints this has not yet been fully investigated.

4.7.2 Precision versus Measurement Rate

The measurement rate is inversely related to the integration time, which is directly related to the amplitude of the correlation waveform. To measure the effect this has on measurement precision, captures of the static scene have been taken with a number of frequencies, four frames per measurement and a range of measurement rates. Once again 360 measurements have been taken at each setting, and standard deviation calculated as above.

Figure 4.14 shows a logarithmic plot of phase standard deviation versus measurement rate for a group of seven selected frequencies. The general trend is that phase measurement standard deviation increases as the measurement rate is increased. This can be attributed to the reduction in signal amplitude as the integration time is reduced. At the slower measurement rate end of the scale the phase standard deviation shows a sharp increase. This is due to saturation of the pixels, significantly reducing the modulation contrast of the sensor and distorting the correlation waveform. The captures with 40 MHz modulation do not saturate

with these measurement rates because the amplitude is sufficiently reduced due to the high frequency, but saturation is still an issue for slower measurement rates (not shown).



Figure 4.14: Logarithmic plot of phase measurement standard deviation versus measurement rate for a range of modulation frequencies.



Figure 4.15: Correlation waveform for f = 1 MHz and integration time per frame of 30, 40, 50 and 60 ms.

The saturation effect is further illustrated in Figure 4.15, which shows the correlation waveform with f = 1 MHz and integration time per frame of 30, 40, 50 and 60 ms. The expected behaviour is evident between 30 and 40 ms where the amplitude is increasing, however for 50 ms the amplitude is reduced and the waveform shape distorted. At 60 ms the modulation contrast is low enough that the pixel voltages of each of the A and B output channels are so similar (both saturated) that their difference is negligible; hence correlation amplitude is very small.

4.7.3 Precision versus phase steps per measurement

The number of phase steps per measurement, *N*, defines the lower limit for the number of frames required per measurement. In general, using more frames per measurement has the same effect as increasing the total integration time, thereby improving measurement precision. However, for a real-time application where a fast measurement rate is desirable, adding more frames whilst maintaining a constant measurement rate comes at a cost of reduced integration time per frame, therefore the measurement precision is not necessarily improved. In the case of very fast measurement rates, the extra frames may lead to a reduction in precision, as with more frames also comes more frame readout periods. With more time taken up performing frame readout, there is equally less time integrating. In the case of very slow measurement rates, more frames can help mitigate saturation, as each frame collects less light compared to fewer frames of the same total integration time.

In this section we aim to discover what the effect of increasing N while maintaining the same measurement rate has on precision. New captures of the scene have been taken with a range of modulation frequencies and frames per measurement of 3, 4, 6, 8 and 12. Measurement rate has been kept constant at 16.67 Hz, that is, for N = 4 each raw frame takes 15 ms. Figure 4.16 shows a semi-logarithmic plot of the phase measurement standard deviation for these captures. For all the modulation frequencies used, there is a progressive increase in phase standard deviation as the number of frames per measurement is increased. The effect is greatest for the higher frequency of 40 MHz, where the



Figure 4.16: Semi-logarithmic plot of phase standard deviation versus frames per measurement for a number of modulation frequencies where measurement rate is fixed at 16.67 Hz.



Figure 4.17: Semi-logarithmic plot of phase standard deviation versus frames per measurement for a number of modulation frequencies where measurement rate is fixed at 4.17 Hz.

standard deviation for N = 12 is larger than for N = 3 by a factor of approximately 8, compared to a factor of approximately 5.6 for the 1 MHz frequency.

The effect is also dependent on the measurement rate for two reasons: 1) for a faster measurement rate, the effect is larger due to the frame readout time being relatively longer compared to the integration time; 2) for very slow measurement rates, when N is small there is a higher chance of bright or near objects in the scene saturating the sensor. These two outcomes are illustrated in Figure 4.17, where captures with the same set of f and N are shown, but with the measurement

rate fixed at 4.17 Hz. The scale of the y-axis is the same as for Figure 4.16, highlighting how minimal the effect of N is between 6 and 12. However, for N of 3 and 4, the phase standard deviation exhibits a marked degradation in precision due to saturation of the sensor.

4.8 Systematic Errors due to Harmonics

In Chapter 2 the issue of contamination due to harmonics in the modulation signals was discussed. The problem arises due to the method of modulating using digital switching circuits, which in theory produce square modulation waveforms that have large amplitude odd-order harmonics. In practice, the limited bandwidth of the illumination driver, image sensor driver and image sensor itself tend to suppress higher-order harmonics, but depending on the modulation frequency there may still be significant harmonic content, as illustrated in earlier sections of this chapter.

Figure 4.18 shows a semi-logarithmic plot of harmonic amplitude relative to the fundamental amplitude for harmonic numbers 2 to 13 for modulation frequencies of 1, 10 and 40 MHz. For the 1 and 10 MHz modulation frequencies, the data shows odd-ordered harmonics are typically larger than even-ordered harmonics right out to harmonic number 9. The remainder are less than 0.22% of the



Figure 4.18: Semi-logarithmic plot of harmonic amplitude relative to fundamental amplitude for harmonic numbers 2 to 13 and modulation frequencies 1, 10 and 40 MHz.

fundamental amplitude. For the 40 MHz modulation only the 2^{nd} and 3^{rd} order harmonics have relative amplitude greater than this level.

The magnitude of the error caused by the harmonic is related to its phase as well as its amplitude, and is cyclic across the range of phase offsets from 0 to 2π . Figure 4.19 shows the variation in the phase measurement error for the three modulation frequencies across the full phase sweep from 0 to 2π , where the number of phase steps per measurement, *N*, is set to 3, 4, 5, 6, 8 and 12. The scale



Figure 4.19: Variation in phase measurement error versus sensor phase offset for modulation frequencies of 1, 10 and 40 MHz, and phase steps per measurement of 3, 4, 5, 6, 8 and 12.

of the y-axis has deliberately been kept constant so that differences between the different plots can be easily observed. For all modulation frequencies, the largest magnitude errors are evident when sampling with 4 phase steps per measurement. This is because this number of samples causes all of the odd-ordered harmonics to be aliased, and these harmonics typically have the largest amplitude.

The number of periods within the error function, n_p , between 0 and 2π is indicative of the dominant harmonic number, n_d , by the relationship

$$n_d = n_p - 1. \tag{4.1}$$

It is evident that for N = 3, the 5th harmonic is dominant for the 1 and 10 MHz modulation (6 periods) and the 2nd harmonic is dominant for the 40 MHz modulation (3 periods).

In order to establish a figure or merit for each *N*, the RMS phase error across the full phase range is taken. This is presented in Figure 4.20 for a large range of modulation frequencies from 1 to 46 MHz (22 MHz has been omitted due to unusually high noise) and N = 3, 4, 5, 6, 8 and 12. Phase steps per measurement of four gives consistently high RMS phase error. N = 3 and N = 6 behave very similarly for frequencies below 26 MHz as they both alias the same odd-ordered harmonics, and even-ordered harmonics are very small for this frequency range. From approximately 26 MHz the 2nd harmonic begins to increase in relative amplitude, increasing the RMS phase error for N = 3.

It is unfortunate that most commercial systems choose to fix the number of phase steps per measurement at four, when in terms of aliasing odd-order harmonics it is the worst choice. The RMS phase error for N = 4 is between 15 and 50 milliradians, which is larger than the phase standard deviation of a high frame rate operating mode with measurement rate of 33 Hz (shown in Figure 4.14). For this system, N = 5 is a good compromise between reducing systematic phase measurement error and maintaining high measurement rate with relatively high precision.



Figure 4.20: Semi-logarithmic plot of RMS phase error versus modulation frequency for N = 3, 4, 5, 6, 8 and 12.

4.9 Power Consumption

Power consumption is an important consideration for mobile robotics applications, where many systems operate from a battery power supply. In this section we compare the power requirements of the Waikato University image intensifier based system with the new Victoria University PMD-19k based system.

The Waikato University system is comprised of a number of sub-components: the laser illumination, image intensifier and associated power supplies, DDS circuit for modulation signal generation and the DALSA Pantera 1M60 CCD camera. The power consumption for each component has been measured independently across a range of modulation frequencies from 1 to 80 MHz, with camera frame rates of 10 and 100 Hz. The camera and DDS circuit were powered from a Topward TPS4000 power supply set to an output voltage of 12.1 V DC. Current was measured using a Fluke 177 True RMS Multimeter. The laser illumination was powered from a Topward power supply with output voltage of 8.1 V DC and current measured using the multimeter. Power consumption in Watts is the product of the voltage reading from the power supply and the current reading of the multimeter. Measuring the power requirements of the image intensifier is more complex, as there are three DC voltage rails generated by a custom power

supply unit: 7000 V, 700 V and 80 V. Total power consumption of the image intensifier has been measured in Watts using a Cixi Yidong EMA-1 Wall Meter.

In presenting the results, the CCD camera, image intensifier and DDS signal generation circuit have been combined to give a sub-total for the power requirements for the sensor aspect of the system. The laser illumination requirements are kept separate, as this is typically an aspect that may be scaled appropriately for an application, for example, reduced output intensity in order to satisfy eye safety requirements. Figure 4.21 shows the measured power consumption versus the modulation frequency for the illumination and sensor operating at 10 and 100 Hz frame rate (note - frame rate, not measurement rate). The sensor is the largest consumer of power in this system, and the general trend is for the power consumption to increase with modulation frequency. There is a noticeable difference in sensor power consumption when operating at different camera frame rates. The reason for this difference is because the image intensifier, the largest consumer of power in the sensor group, is shut down during the readout period of the camera. Higher frame rates increase the relative time spent performing readout, hence lower power consumption overall. The illumination power consumption is unaffected by the camera frame rate because the laser diodes are continuously in operation, in order to maintain a constant



Figure 4.21: Power consumption versus modulation frequency for the sensor (CCD camera, image intensifier, DDS) and illumination circuits of the University of Waikato range imaging system.

operating temperature.

Power consumption measurements have also been taken for the new Victoria University range imaging system. The power requirements for the system are much more simplified as there are only two power sources needed: a 10 V DC source for the laser illumination circuit, and a 16 V DC source for the Stratix III Development Kit, which powers the PMD-19k sensor and drives the modulation signals. This situation is more desirable than that for the Waikato University system, as there is no requirement for a custom power supply to provide voltages up to 7000 V.

Figure 4.22 shows the power consumption versus modulation frequency for the sensor and illumination circuits for the PMD19k based system. The power used by the sensor is between 30% and 46% that of the University of Waikato sensor subsystem. The general trend for the sensor power is to increase up to a plateau beginning at 30 MHz. At frequencies higher than this level the extra power due to increased frequency is offset by the reduced amplitude of the sensor modulation due to the limited bandwidth of the driver electronics. Low sensor frame rates consume relatively more power than high frame rates for the same reason as the University of Waikato system, that is, the sensor modulation is disabled during the frame readout period.

The power used for the illumination is between 304% and 334% that of the University of Waikato system. This increase is largely attributed to the upgrade from 4 to 16 laser diodes, although the increase is not 400% as expected because the individual lasers used are not as high powered. The infra-red lasers are also slightly more efficient. Figure 4.23 shows the total power consumption for both sensor and illumination subsystems of the Waikato University and Victoria University range imaging systems, at a sensor frame rate of 100 Hz, which is more applicable to real-time operation. The useful frequency range of the Victoria University system is from 1 to approximately 40 MHz, beyond which precision begins to sharply decline. Over this range, the total power consumption is between 63% and 69% that of the Waikato University System, which is a

significant reduction. If the same illumination circuit were used for both designs, the total power consumption could be expected to be between 42% and 48%.



Figure 4.22: Power consumption versus modulation frequency for the sensor and illumination circuits of the Victoria University range imaging system.



Figure 4.23: Total power consumption versus modulation frequency for the Waikato University and Victoria University range imaging systems.

4.10 Summary

In this chapter the performance of the new Victoria University range imaging system based on the PMD19k image sensor has been investigated. The system is capable of changing operating parameters through software, and the parameters that have the greatest impact on measurement precision are the modulation frequency, f, the integration time, T_I , and the number of phase steps per measurement, N.

There is an inverse relationship between distance measurement standard deviation, σ_d , and f for a system where phase measurement standard deviation, σ_{φ} , is assumed to be constant. For this system, σ_{φ} increases at a linear rate between 1 and 40 MHz, beyond which bandwidth limitations of the sensor driver electronics cause a significant reduction in correlation waveform amplitude, greatly degrading phase precision. Peak precision is obtained at f = 36 MHz. For the simple test target positioned 2.63 m from the image sensor, operating in a real-time measurement frame rate of 25 Hz, σ_d was measured as 4.06 mm, corresponding to 0.155% of the target distance. This compares very well to the commercial systems presented in Chapter 2, where depth precision ranged from 0.075% to 1.5%.

Increasing integration time improves measurement precision by collecting and averaging over more photons, enhancing the signal-to-noise ratio. However, if the integration time is too long, pixels imaging nearer or more highly reflective objects will saturate, drastically reducing performance. An alternative to increase total integration time with a reduced risk of pixel saturation is to use more raw frames per measurement. However, if measurement rate is to be held constant, this approach may give reduced precision due to a reduction in integration time equal to the extra frame readout periods required.

The method of producing the modulation signals for this system is to use a digital switching circuit. The resultant modulation waveforms are square, particularly at lower frequencies, which have relatively large amplitude odd-ordered harmonics. Depending on the sampling rate of the correlation waveform, that is, phase steps per measurement, N, these harmonics may be aliased onto the signal of interest

causing systematic errors in the measurement. The results presented here suggest that N = 5 is a good compromise between measurement precision and accuracy for high frame rate operation. For applications requiring high quality measurements of static scenes, where frame rate is of no consideration, a better choice is to use a large value for N, such as 30, and take as many frames as possible with integration adjusted to be as high as possible without causing saturation. Increasing N reduces systematic errors due to harmonics, and taking many frames increases precision by averaging over a very long total integration time.

Finally, the power consumption of the new system was compared to the image intensifier based University of Waikato system. A significant improvement is removing the requirement for the 7000, 700 and 80 V supply lines. The new system uses only two thirds of the total power despite increasing the laser illumination power. If the same laser illumination circuit were used, the total power requirements would be less than half.

Chapter 5 Multiple Modulation Frequency Range Imaging

Recall from Chapter 2 the equations for the maximum unambiguous distance, d_u , and measured distance, d,

$$d_u = \frac{c}{2f},\tag{2.16}$$

$$d = d_u \left(\frac{\varphi}{2\pi} + k\right),\tag{2.15}$$

where *c* is the speed of light, *f* is the modulation frequency, φ is the measured phase offset and the integer *k* accounts for the possible wrapping of phase. This problem of phase wrapping is illustrated in Figure 5.1, where the measured object distance is incorrect by some multiple of the maximum unambiguous distance for the modulation frequency used. The distance precision, σ_d is based on the phase precision, σ_{φ} by the relationship

$$\sigma_d = \frac{c\sigma_{\varphi}}{4\pi f} = \frac{d_u \sigma_{\varphi}}{2\pi}.$$
(2.22)

It can be seen that one method for improving the distance measurement precision is to increase the modulation frequency; however, this comes with a corresponding decrease in the maximum unambiguous range.

An alternative is to utilise phase measurements acquired using two different modulation frequencies, f_A and f_B (Jongenelen, 2011). As illustrated in Figure 5.2, each measurement gives a set of possible object locations, and the true location is



Figure 5.1: Due to phase wrapping of the returned illumination signal, objects beyond the maximum unambiguous range, d_u , appear to be closer than they really are.



Figure 5.2: A target measured twice using different modulation frequencies. Each measurement gives a set of possible object locations, and the true location is determined by where the two measurements are in agreement.

determined by where the two measurements are in agreement. The technique is similar to that used in two-wavelength interferometry (Hariharan, 1986), where the combination of two wavelengths produce a maximum range as if a single effective wavelength, λ_E , were used. For this application we will refer to an effective modulation frequency, f_E , which is based on the greatest common divisor of the two frequencies by

$$f_E = \gcd(f_A, f_B). \tag{5.1}$$

The effective maximum unambiguous range, d_E , from this effective frequency is

$$d_E = \frac{c}{2f_E}.$$
(5.2)

In this chapter we investigate numerical methods for calculating distance with extended maximum unambiguous range based on two phase measurements obtained with different frequency. Initially a naïve method is presented in Section 5.1, and an improved method presented in Section 5.2 based on a

variation of the New Chinese Remainder Theorem (Wang, 2000). An analysis of the effect on distance measurement precision due to noise in the phase values is presented in Section 5.3. Dual-frequency measurements may cause large measurement errors if the integer k is estimated incorrectly, and the probability of this occurrence is analysed in Section 5.4. In Section 5.5 we discuss options for obtaining the required two phase measurements, either independently or simultaneously, and the advantages and disadvantages of each option.

5.1 Calculating distance from two phase measurements

This approach requires two phase measurements, φ_A and φ_B , measured using modulation frequency f_A and f_B respectively, with associated maximum unambiguous distances of d_{uA} and d_{uB} . The ratio $f_A:f_B$ can be expressed by the two co-prime integers $M_A:M_B$, found by

$$M_A = \frac{f_A}{f_E} = \frac{d_E}{d_{uA}},\tag{5.3}$$

$$M_{B} = \frac{f_{B}}{f_{E}} = \frac{d_{E}}{d_{uB}}.$$
 (5.4)

Throughout this chapter it will be assumed that $M_A > M_B$. For convenience, the phases can be scaled to lie within the range [0:1) by working with $p_A = \varphi_A/2\pi$ and $p_B = \varphi_B/2\pi$. In the absence of noise, both of these values represent the distance by

$$d = d_{uA}(k_A + p_A) = d_{uB}(k_B + p_B).$$
(5.5)

The challenge is to establish the integer values of k_A and k_B such that the calculated distances from each are in agreement. Due to errors in the phase measurements, it is unlikely that the calculated distance from each phase will ever be exactly equal, therefore the aim is to find k_A and k_B such that the difference between the two measurements is minimised. This function can be expressed as

$$y(k_A, k_B) = |M_B(k_A + p_A) - M_A(k_B + p_B)|$$
(5.6)

A naïve approach is to evaluate all possible combinations of k_A and k_B and select the options that give the smallest difference. As the value of k_A can range from 0 to $M_A - 1$, and k_B from 0 to $M_B - 1$, the function must be evaluated for a total of $M_A M_B$ iterations.

With values established for k_A and k_B , distance can be computed by the weighted average of the two measurements by

$$d = \frac{c}{2} \left(\frac{w(k_A + p_A)}{f_A} + \frac{(1 - w)(k_B + p_B)}{f_B} \right)$$

= $\frac{d_E}{M_A M_B} \left(w M_B (k_A + p_A) + (1 - w) M_A (k_B + p_B) \right)$, (5.7)

where w is a weighting factor between 0 and 1, chosen to minimise the variance in the output distance estimate. This weighting factor will be further discussed in Section 5.3.

As an example, consider a target at a distance of 11.1 m, measured with two modulation frequencies, $f_A = 50$ MHz, and $f_B = 40$ MHz. Assuming a noiseless scenario, let $p_A = 0.70$ and $p_B = 0.96$. Evaluating Equation (5.6) for $k_A = [0:M_A)$ and $k_B = [0:M_B)$ produces a minimum difference of 0.0 where $k_A = 3$ and $k_B = 2$. Applying Equation (5.7) with a weighting, w, of 0.5 gives a final distance measurement, d = 11.1 m.

A more direct method exists given the constraint that $M_B = M_A - 1$, and is the method used in two-wavelength interferometry. In this case, the effective frequency, f_E is equal to $|f_A - f_B|$, and the effective phase offset from this frequency, p_E , is mod $(p_A - p_B, 1)$. The integer k is assumed to be zero, and the distance calculated by

$$d = \frac{c}{2|f_A - f_B|} \operatorname{mod}(p_A - p_B, 1) = d_E p_E.$$
(5.8)

Using the example given earlier, $p_E = 0.74$, and hence d = 11.1 m. However, in addition to the added constraint for $M_A:M_B$, this method results in a larger standard deviation because it is scaled by d_E , as will be further discussed in Section 5.3. To account for this, the distance estimate can be used to find k_A and k_B by dividing by the individual unambiguous distances and taking the integer quotient

$$k_{A} = \left\lfloor \frac{d}{d_{uA}} \right\rfloor$$

$$k_{B} = \left\lfloor \frac{d}{d_{uB}} \right\rfloor$$
(5.9)

These can then be substituted back into Equation (5.7) with the original phase values, and a higher precision measurement obtained.

5.2 Calculation based on New Chinese Remainder Theorem

An alternative method that works for all co-prime combinations of $M_A:M_B$ involves a modification of a residue-to-binary converter. The distance measurement is calculated in two steps. In the first step we find an estimate for the integer k_B using the modified New Chinese Remainder Theorem (Wang, 2000);

$$e = p_A M_B - p_B M_A$$

$$k_B = \text{mod}[k_0 \text{ round}(e), M_B],$$
(5.10)

where k_0 is the modular multiplicative inverse such that $mod(k_0M_A,M_B) = 1$. This calculation works equally well to find k_A by substituting the relevant subscripts, though the value of k_0 is typically larger. The second step calculates the final distance measurement by

$$d = \frac{d_E}{M_A M_B} (M_A k_B + M_A p_B + w(e - \text{round}(e))).$$
(5.11)

This is a simplification of Equation (5.7) where the calculation of k_A is not required. The third term in the equation effectively weights the two estimates by reintroducing the weighted error term.

Using the example given in the previous section where $f_A = 50$ MHz, $f_B = 40$ MHz, $p_A = 0.70$ and $p_B = 0.96$. The value for k_0 evaluates to 1 in this example, giving a k_B of 2. In the absence of noise, e - round(e) equates to 0, and the final distance measurement is 11.1 m.

5.3 Effect of phase precision on distance precision

For sufficiently large amplitude signals, the phase offset measured using each modulation frequency follows a Gaussian distribution (Frank et al., 2009). Let σ_{pA} and σ_{pB} be the standard deviations of the scaled phase measurements p_A and p_B respectively. The standard deviation of the distance measurement, σ_d computed by either Equations (5.7) or (5.11) is

$$\sigma_{d} = \frac{d_{E}}{M_{A}M_{B}} \sqrt{\left[w\sigma_{pA}M_{B}\right]^{2} + \left[(1-w)\sigma_{pB}M_{A}\right]^{2}}, \qquad (5.12)$$

whereas the standard deviation for Equation (5.8) is

$$\sigma_d = d_E \sqrt{\sigma_{pA}^2 + \sigma_{pB}^2} \,. \tag{5.13}$$

For all combinations of M_A , M_B and w the standard deviation for Equation (5.8) is larger than that for Equations (5.7) or (5.11), in some cases by an order of magnitude.

To maximise precision, it is important to use an appropriate value for w. This value should be chosen in order to place a higher weighting on the phase measurement with higher precision. To find the optimum value, the differential of σ_d^2 with respect to w is taken, and equated to zero, yielding

$$w = \frac{\sigma_{pB}^2}{\sigma_{pB}^2 + \sigma_{pA}^2 \left(\frac{M_B}{M_A}\right)^2}.$$
(5.14)

The difficulty with this equation is that the standard deviation of a given measurement is not necessarily known at the time when d is to be calculated. However, as has been shown in previous sections, the signal amplitude has an inverse relationship with phase standard deviation, and can act as an indicator for precision. Equation (5.14) can therefore be rearranged to make use of the amplitude information by

$$w = \frac{A_A^2}{A_A^2 + A_B^2 \left(\frac{M_B}{M_A}\right)^2},$$
 (5.15)

where A_A and A_B represent the signal amplitudes for the measurements using modulation frequencies f_A and f_B respectively.

By substituting w from Equation (5.14) back into Equation (5.12), the distance standard deviation calculation can be simplified to

$$\sigma_{d} = \frac{d_{E}}{\sqrt{\left(\frac{M_{A}}{\sigma_{pA}}\right)^{2} + \left(\frac{M_{B}}{\sigma_{pB}}\right)^{2}}}.$$
(5.16)

We now compare this to Equation (2.22) where a measurement with a single modulation frequency, $f_S = f_E$, gives a scaled phase standard deviation of σ_{pS} . When using a single frequency it can be expected that twice the integration time is available, halving the phase standard deviation, hence $\sigma_{pS} = \sigma_{pA}/2 = \sigma_{pB}/2$. Nevertheless, the use of two modulation frequencies can be expected to yield an improvement in distance measurement precision by a factor of $\frac{1}{2}\sqrt{M_A^2 + M_B^2}$.

In summary, not only the choice of the individual modulation frequencies used, but also the ratio between them has an affect on the distance measurement precision. In general, for a given effective frequency, f_E , and constant phase precision, higher integer values of M_A and M_B will produce more precise results. Also, for a given ratio of $M_A:M_B$, there is scope to maximise distance precision by modifying system parameters, for example, integration time, to place more emphasis on the higher frequency measurement.

5.3.1 Simulated Distance Precision versus Phase Precision

To verify the expected distance precision for a given phase precision, and M_A and M_B , a simulation has been produced in MATLAB. For a more realistic simulation it can not be assumed that phase precision is constant for all frequencies. Experimental results presented in Chapter 4 show that for the Victoria University range imaging system, σ_p increases at an approximately linear rate over the modulation frequency range from 1 to 40 MHz, as shown in Figure 5.3. For this investigation an effective frequency, f_E , of 8 MHz has been chosen as this gives an integer multiple, M_A , of up to 5 where the frequency is still within the linear range.

For the simulation, σ_p at each higher frequency has been adjusted relative to σ_p at 8 MHz by assuming it increases linearly by 2% of the 8 MHz standard deviation for each 1 MHz frequency increase. Mathematically, σ_p for each frequency greater than 8 MHz has been simulated by

$$\sigma_p(f) = \sigma_p(8MHz) \left(1 + \frac{2(f - 8.10^6)}{100.10^6} \right).$$
(5.17)

For example, if $\sigma_p(8 \text{ MHz}) = 10.0 \text{ milliradians}$, then $\sigma_p(32 \text{ MHz}) = 14.8 \text{ milliradians}$ and $\sigma_p(40 \text{ MHz}) = 16.4 \text{ milliradians}$.



Figure 5.3: Experimental σ_{φ} versus modulation frequency for the Victoria University range imaging system.

The maximum unambiguous distance, d_E is based on an effective frequency of 8 MHz as per Equation (5.2). The simulated true distance has been taken from a uniform random distribution between 0 and d_E . Simulated phase measurements for each frequency have been established using a rearranged version of Equation (5.5) with normally distributed noise added by the formula

$$p_{A} = \operatorname{mod}\left(\frac{d}{d_{uA}} + \frac{\Delta_{A}}{2\pi}, 1\right)$$

$$p_{B} = \operatorname{mod}\left(\frac{d}{d_{uB}} + \frac{\Delta_{B}}{2\pi}, 1\right),$$
(5.18)

where Δ_A and Δ_B are random values taken from independent normal distributions with means of zero and standard deviations of $\sigma_p(f_A)$ and $\sigma_p(f_B)$ respectively.

Simulated modulation frequencies for multiple frequency measurements are the combinations of 16 and 8 MHz ($M_A:M_B = 2:1$), 40 and 8 MHz ($M_A:M_B = 5:1$) and 40 and 32 MHz ($M_A:M_B = 5:4$). All combinations result in an effective frequency, f_E , of 8 MHz. Given the noisy phase measurements p_A and p_B , the distance has been calculated using Equations (5.10) and (5.11), with the weighting factor, w



Figure 5.4: Logarithmic plot of simulated and theoretical σ_d versus σ_p for modulation frequency combinations of 40 and 32 MHz, 16 and 8 MHz, 40 and 8 MHz and 8 MHz only.

chosen by Equation (5.14) using the known phase standard deviations. A total of 10,000 measurements have been simulated for each set of parameters.

Figure 5.4 shows a logarithmic plot of simulated distance standard deviation versus phase standard deviation (adjusted for frequency by Equation (5.17) for the frequency combinations noted above. Solid lines in the plot indicate the theoretical σ_d based on Equation (2.22) for the single frequency 8 MHz operation or results from Equation (5.16) for the multiple modulation frequency measurements. For σ_p less than 100 milliradians the theory demonstrates a very good approximation of the simulated results in this case. The distance standard deviation when $f_A: f_B = 40:32$ shows an improvement in precision by a factor of approximately 3.2 compared with the single 8 MHz modulation frequency. This is not as high as $\sqrt{M_A^2 + M_B^2}$, and the difference can be accounted for by the fact that σ_p has not been held constant across all modulation frequencies. For σ_n greater than 100 milliradians the calculation occasionally incorrectly estimates k_B resulting in a larger than expected σ_d . This effect is discussed in greater detail in Section 5.4.

5.4 Errors in distance measurement

In the event that the integer k_B is incorrectly estimated, the measurement of the distance can be incorrect by some integer multiple of the maximum unambiguous distance, d_{uB} (and vice versa if estimating k_A). Although this is not always the case as it is possible to have an incorrect estimate for k_B , but still calculate an acceptable distance measurement if either p_A or p_B are near the extremes of 0 or 1. Consider the example given above where the true distance is 11.1 m. Expected phase measurements for this distance are $p_A = 0.70$ and $p_B = 0.96$. Suppose noise on p_A reduces the measurement to 0.65 and noise on p_B causes it to wrap around to 0.04. By Equation (5.10), k_B is incorrectly calculated to be 3, yet by Equation (5.11), d still gives the correct estimate of 11.1 m.

An event where *d* is incorrectly estimated by a margin of half the maximum unambiguous range of the higher frequency measurement is deemed to be an error, and the percentage number of errors within a given number of samples is the error rate. An error is typically disproportionate to the level of noise on the phase measurements, depending on the ratio $M_A:M_B$. Mathematically, the occurrence of an error can be defined as

$$error = \begin{cases} 0, & |d - \hat{d}| < \frac{d_{uA}}{2} \\ 0, & |d - \hat{d}| \ge d_E - \frac{d_{uA}}{2} \\ 1, & otherwise \end{cases}$$
(5.19)

where \hat{d} is the true object distance. The purpose of the two options for *error* = 0 is to allow for the possibility that the distance measurement can wrap around the new maximum unambiguous distance, d_E . In terms of the phase measurements, an error can be defined as

$$error = \left| p_B M_A - p_A M_B - \hat{i} \right| > \frac{1}{2},$$
 (5.20)

where \hat{i} is the integer which will give the correct distance estimate. Without prior knowledge of the true object distance, or other constraints relating to object
surface properties, it is impossible to identify an error in an individual pixel. Nevertheless, Equation (5.20) is useful in that it can be used to estimate the error rate given the variables of M_A , M_B and an indication of the phase standard deviations, σ_{pA} and σ_{pB} . The standard deviation of the phase difference, σ_e , can be calculated by

$$\sigma_e = \sqrt{\sigma_{pB}^2 M_A^2 + \sigma_{pA}^2 M_B^2} . \qquad (5.21)$$

The probability that this normally distributed variable is half of a standard deviation from the mean (zero), and hence the probability of an error, P_{error} , can be calculated by

$$P_{error} = 1 - \operatorname{erf}\left(\frac{1}{2\sqrt{2}\sigma_e}\right),\tag{5.22}$$

where erf(x) is the error function, defined as

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$
 (5.23)

In summary, while higher values of M_A and M_B can increase distance measurement precision, they also lead to an increase in the error rate.

5.4.1 Simulated Error Rate versus Phase Precision

To verify the expected error rate for a given σ_p , the simulation performed in Section 5.3.1 was extended to calculate the error rate for frequency combinations of 40 and 32 MHz, 16 and 8 MHz, and 40 and 8 MHz. To clarify the concept of an error as opposed to the usual variation in distance measurement due to noise, Figure 5.5 shows a number of distance measurements with σ_p of 0.2 radians. The large spikes are caused by incorrect estimates of the integer k_B , and result in a distance measurement error over an order of magnitude larger than the distance measurement standard deviation, σ_d .



Figure 5.5: Distance measurement error for a number of samples with $\sigma_p = 0.2$ radians. The large spikes are due to errors in the estimate for the integer number of times the phase has wrapped around.



Figure 5.6: Logarithmic plot of simulated and theoretical error rate versus σ_p for frequency combinations of 40 and 32 MHz, 16 and 8 MHz, and 40 and 8 MHz.

The results of the simulated error rate are shown in Figure 5.6, with the solid lines indicating the theoretical error rate as per Equations (5.21) and (5.22). The simulated behaviour agrees with the theory in this instance. As the error rate drop below 1% the simulation starts to lose statistical significance and differs slightly from the theory, because with the relatively small sample size (10,000) the number of errors counted falls below the minimum level required to maintain a reasonable confidence interval.

5.5 Acquisition of Dual Frequency Measurements

In this section the practicalities of acquiring two measurements of the same scene using two modulation frequencies is addressed.

5.5.1 Independent capture

The simplest method is to take two independent captures of the scene, one at each of the two modulation frequencies. This is most appropriate for static scenes where the positions of the objects will not change between the two measurements. There is no constraint for each capture in terms of the selected system parameters of phase steps per measurement, N, or integration time, T, provided the frequencies used have a known relationship through the ratio of co-prime integers M_A and M_B .

In a real-time application, the measurement frame rate must also be considered. To obtain two phase measurements, a minimum of six frames are required if capturing independently, that is three frames for each modulation frequency. In order to maintain a desired measurement frame rate, this necessitates a reduced integration time for each of the six frames, compared to the case using only a single low frequency (with unambiguous distance appropriate for the scene).



Figure 5.7: The time taken for a single modulation frequency measurement is made up of N integration periods, each of length T/N, and N frame readout periods of length T_R .



Figure 5.8: Dual-frequency modulation requires more frames. To keep the measurement rate the same, the total integration time for each frequency, T_A and T_B , is reduced.

Furthermore, as more frames are required there is proportionally more 'dead time' where the image sensor is performing readout, and not integrating. This is illustrated in Figure 5.7 for a single modulation frequency and Figure 5.8 for dual-frequency modulation. The effect becomes more pronounced as the raw frame rate increases, where the readout period is proportionally larger compared to the integration time.

Effect of total integration time on distance precision

The total integration time, T, is directly proportional to the signal amplitude, A, which is inversely proportional to the phase standard deviation σ_p (Frank et al., 2009). Let T_A and T_B be the total integration time for each of the measurements using two modulation frequencies, and T_S be the total integration time for a single modulation frequency measurement with $f = f_E$, and T_R be the readout time. If $T_A = T_B$, and $T_S + NT_R = T_A + T_B + 2NT_R$, then the ratio between the phase standard deviation of a single frequency measurement, σ_{pS} , to σ_{pA} or σ_{pB} is

$$\frac{\sigma_{pS}}{\sigma_{pA}} = \frac{T_A}{T_S} = \frac{T_A}{2T_A + NT_R}.$$
(5.24)

Consequently, the ratio between the distance measurement standard deviations, σ_{dS} and σ_d is

$$\frac{\sigma_{dS}}{\sigma_d} = \frac{T_A \sqrt{M_A^2 + M_B^2}}{2T_A + NT_R} = \frac{\sqrt{M_A^2 + M_B^2}}{2 + N T_R / T_A}.$$
(5.25)

This ratio will be larger than one for all realistic applications where N = 4, $M_A > 2$, $M_B > 2$ and $T_R/T_A < 0.5$. Therefore, for a given effective frequency (hence effective unambiguous distance, d_E), it is theoretically viable to use two modulation frequencies in order to increase precision for an extended maximum unambiguous range, while still maintaining the same measurement frame rate.

Effect of relative integration time on distance precision

Within a given measurement period, there is the opportunity to spend relatively more time integrating at one of the modulation frequencies, to selectively improve the precision of one phase measurement over the other. Let *R* be the fraction of time spent at the higher modulation frequency, f_{A_2} to the total integration time, in other words, $R = T_A/(T_A+T_B)$. Again, we will assume that σ_{pA} and σ_{pB} are inversely proportional to T_A and T_B , therefore $\sigma_{pA} \propto 1/R$ and $\sigma_{pB} \propto 1/(1-R)$. Substituting into Equation (5.16) gives the relationship

$$\sigma_d \propto \frac{1}{\sqrt{(M_B(1-R))^2 + (M_A R)^2}}$$
 (5.26)

The minimum values for this function are at the extremes of R = 0 and R = 1, although while these might yield the best precision, the error rate will be high.

Effect of total integration time on error rate

When measuring distance using a single measurement with modulation frequency $f_S = f_E$, the integer value for k is always assumed to be zero. Consequently, the possibility of an error as defined by Equation (5.19) is also zero. Therefore, in terms of comparing a single modulation frequency measurement to a two modulation frequency measurement, using two modulation frequencies will always be inferior in terms of the error rate. So, while the precision may be increased, a larger number of results are not useful. This is the trade-off required in order to increase measurement precision.

Effect of relative integration time on error rate

Performing the same substitution for R into Equation (5.21) gives the relationship

$$\sigma_e \propto \sqrt{\left(\frac{M_A}{1-R}\right)^2 + \left(\frac{M_B}{R}\right)^2}$$
 (5.27)

In the interest of reducing the error rate, it is desirable to reduce σ_e . The minimum value for σ_e can be found by differentiating the part within the square root, and solving for 0. The differentiation gives

$$\left(\frac{d\sigma_e}{dR}\right)^2 = M_A^2 \left(\frac{2}{(1-R)^3}\right) - M_B^2 \frac{2}{R^3}.$$
(5.28)

When equated to 0, a unique solution exists for all positive values of M_A and M_B at

$$R = \frac{\sqrt[3]{M_A^4 M_B^2} + \sqrt[3]{M_A^2 M_B^4} + M_B^2}{M_A^2 + M_B^2}.$$
 (5.29)

This result is useful in order to configure the integration ratio appropriately in order to reduce the error rate for a given combination of $M_A:M_B$.

5.5.2 Simultaneous Dual Frequency Modulation

The previous sections have described methods using multiple independent captures at different modulation frequencies. Previous work by (Payne et al., 2009b) has demonstrated that it is also possible to encode both of these measurements simultaneously within a single capture, whereby the modulation frequency is alternated within the integration period. This reduces the required number of readout periods, as illustrated in Figure 5.9. However, Payne et al. did not investigate the error rate for low precision phase measurements nor did the experimental results encompass a range of modulation frequency combinations.

In order to distinguish between the two measurements, the phase step applied to the image sensor of the second modulation frequency is doubled such that during acquisition of N samples, the phase steps from 0 to 4π radians. In effect, the first modulation frequency is sampled over one cycle (from 0 to 2π), while the second modulation frequency is sampled over two cycles (from 0 to 4π), allowing each phase and magnitude value to be resolved from the combined measurement.



Figure 5.9: Dual-frequency measurement where measuring each frequency simultaneously reduces the number of readout periods, allowing more time for integration.

Effect on distance precision

As there are five unknowns (the DC offset and the phase and magnitude of each modulation frequency component) a minimum of five samples are required to uniquely determine these variables; one less than the case where the two captures are performed independently. The removal of one frame readout period, T_R , increases the total integration time available for a given measurement rate, potentially improving the range measurement precision.

Effect on distance accuracy

In the case where captures are performed independently, the number of phase steps per measurement, N, can be appropriately selected in order to reduce the effect of harmonics interfering with the phase measurement, thereby reducing the systematic error caused by this. However, when capturing two frequencies simultaneously, greater care must be taken with selecting N, as it is not only the harmonics of one frequency interfering with itself, but also with the other frequency that must be considered (Payne et al., 2009b).

For a single frequency, candidate harmonics that are aliased onto the fundamental frequency can be found by the formula given in Chapter 2,

$$k = nN \pm 1, \tag{2.24}$$

where k is the harmonic number and n an integer (0, 1, 2, 3...). For example, if N = 5, the candidate harmonics that may interfere with the phase measurement are k = (4, 6, 9, 11, 14, 16...), and as typically only the odd-ordered harmonics have significant magnitude, N = 5 is considered a good choice.

For two frequencies captured simultaneously, candidate harmonics are aliased by the relationship

$$k_{a\to a} = nN \pm 1, \quad k_{b\to a} = n\frac{N}{2} \pm \frac{1}{2}, \quad k_{b\to b} = n\frac{N}{2} \pm 1, \quad k_{a\to b} = nN \pm 2, \quad (5.30)$$

where the subscripted references $k_{a\rightarrow b}$ etc. refer to harmonics of frequency *a* aliasing onto the fundamental component of frequency *b* for example. For the example given above with N = 5, $k_{a\rightarrow a}$ is the same set as earlier, $k_{b\rightarrow b} = (4, 6, 19...), k_{a\rightarrow b} = (3, 7, 8, 12, 13, 17...)$ and $k_{b\rightarrow a} = (2, 3, 7, 8...)$. Odd-ordered harmonics do feature in these lists, particularly the typically large magnitude 3^{rd} harmonic. Therefore, while N = 5 may improve precision by the removal of one frame readout period, it may cause larger systematic errors.

Selecting higher values for N, for example 6, may re-order the values somewhat, but harmonics will still be present. Increasing N further is effectively just trading off reduced errors due to harmonics for lower precision due to reduced integration time as a result of the extra frame readout periods.

In summary, the value of simultaneous dual frequency modulation is highly dependent on the specific properties of the range imaging system. If the modulation suffers from large amplitude harmonics (for example, high bandwidth square wave modulation) the benefits of increased precision due to increased integration time may be negated by the influence of harmonics. However, if these issues can be addressed, either through harmonic cancellation techniques (Payne et al, 2008b) or robust calibration, then simultaneous dual frequency modulation may give improved measurement precision compared to independent capture.

5.5.3 Simulated Distance Precision versus Integration Ratio

To verify the effect of the integration ratio on the distance measurement standard deviation, a simulation has been performed in MATLAB. Modulation frequency combinations investigated are 40 and 32 MHz ($M_A:M_B = 5:4$), and 40 and 8 MHz ($M_A:M_B = 5:1$). To begin, a reasonable estimate for σ_p for a 100% ratio of 8 MHz

modulation frequency (hence 0% at 40 MHz) has been chosen as 0.06 radians. The integration ratio is directly proportional to the correlation function amplitude, which is inversely proportional to σ_p . Therefore, σ_{p8MHz} is simulated across the range of 0 to 100% by the formula

$$\sigma_{p8MHz} = 0.06 \frac{100}{(0.001 + 100 - R)},$$
(5.31)

where R is the ratio of time spent integrating the higher 40 MHz frequency between 0 and 100%. The purpose of the addition of 0.001 is to ensure that at 100% the phase standard deviation is not infinite.

Phase standard deviation for the 32 and 8 MHz frequencies have been chosen in a similar manner. To improve the realism of the simulation, σ_p for 40 and 32 MHz have been scaled by factors based on Equation (5.17), that is

$$\sigma_{p40MHz} = 1.64 \cdot 0.06 \cdot \frac{1}{(0.001 + R)}$$

$$\sigma_{p32MHz} = 1.48 \cdot 0.06 \cdot \frac{1}{(0.001 + 100 - R)}$$
(5.32)

This process can be summarised by the plot shown in Figure 5.10.

With appropriate values for σ_p established, 100,000 measurements have been simulated for each of 21 equally spaced values for the integration ratio from 0 to 100% and the frequency combinations given above. Distance has been calculated using Equations (5.10) and (5.11), with the weighting, w, determined both as a fixed weighting of 50% and a variable weighting based on Equation (5.14). Note that in this case outliers due to disambiguation errors have been removed before calculating σ_d in order to obtain a better estimate of the underlying distribution.



Figure 5.11: Simulated σ_d versus integration ratio for frequency combinations of 40 and 32 MHz, 40 and 8 MHz, and 8 MHz only. For each multiple frequency measurement a comparison between a fixed weighting scheme of 50% and a variable weighting based on Equation (5.14) is also shown.

The results of the simulated σ_d versus integration ratio are shown in Figure 5.11. It can be seen for the multiple frequency measurements the use of a variable weighting based on the frequency ratio and amplitude improves precision compared to the fixed weighting of 50%. The curve for $M_A:M_B = 5:1$ clearly shows the trade-off in terms of time split between the high and low frequencies. Increasing the period of the high frequency improves the accuracy only until a certain critical proportion when using a fixed weighting. When using the variable weighting scheme the standard deviation decreases to a minimum where 100% of the result is based on the high frequency measurement. With two similar frequencies, there is a wide band over which precise measurements can be made. Although the precision is improved with only a single frequency high frequency (at both ends of the integration ratio scale), the error rate is such that many of the results would be meaningless. By comparison, for a single 8 MHz modulation frequency the integration ratio is not applicable. For all ratios, the dual high frequency measurements give superior precision, whereas for the high and low frequency combination the precision is more dependent on the integration ratio.

5.5.4 Simulated Error Rate versus Integration Ratio

Figure 5.12 shows the simulated error rate versus the ratio of integration times. Using two relatively high frequencies has the advantage of potentially improving precision, but also has an associated cost in terms of a higher error rate compared to the high and low frequency case. The combination of high and low frequency $(M_B = 1)$ only produces errors at one end of the scale, where the relative time spent on the low frequency is less. This is because in this case the high frequency measurement does not contribute to resolving the ambiguity. The theoretical error rate is indicated by the solid lines, and the simulated data points match up very closely.



Figure 5.12: Simulated error rate versus integration ratio for frequency combinations of 40 and 32 MHz, and 40 and 8 MHz.

5.6 Dual Frequency Experimental Results

In order to test the hypothesis that the use of multiple modulation frequencies improves distance measurement precision, a number of captures have been taken using the range imaging based on the PMD19k image sensor described in Chapter 3. In this Section we describe the experimental setup and the various parameters of interest, and present results in terms of distance measurement standard deviation, σ_d , and error rate.

5.6.1 Experimental Setup

For these captures a scene has been set up using six boxes positioned at 2.50, 3.75, 5.00, 6.25, 7.50 and 8.75 m from the imaging system. One white and one green target has been attached to the boxes at each distance so that a range of amplitude values can be returned for the same distance. When imaged by the system, an 8×8 group of pixels located at each target are analysed, with the distance to their centres measured using a Leica DISTO distance meter (manufacturer specified accuracy of ±1.0 mm). This size group of pixels has been deemed small enough that geometric distortions across the group will be considered negligible. A white coloured target on the far wall is the most distant target, with the centre of an 8×8 group of pixels measured to be 11.30 m. Figure 5.13 shows a diagram of the test scene with each target numbered from 1 to 13.

The scene is long enough that in order to correctly measure the distance to all targets in the scene without ambiguity would require a modulation frequency less than 13.2 MHz. Measurements taken with 40 MHz are particularly interesting, as ambiguity causes some targets to be measured at almost identical distances.

The scene has been initially measured using a modulation frequency, f, of 8 MHz. This frequency gives maximum unambiguous range, d_u , of 18.75 m, which is large enough to measure all targets in the test scene. To reduce the influence of harmonics the number of phase steps per measurement, N, has been set to 240, and to improve precision the measurement frame rate was reduced to 0.083 Hz.



Figure 5.14 shows a grey scale image of the scene where brightness represents the

Figure 5.13: Diagram of scene for testing multiple frequency range imaging. The scene consists of 6 boxes positioned at known distances from the imaging system, with one white and one green target on each box.



Figure 5.14: Image of test scene where grey scale represents pixel amplitude for a capture with f = 8 MHz, N = 240 and measurement rate = 0.083 Hz. The numbers indicate where the test targets described in Figure 5.13 can be found in the image.



Figure 5.15: Image of test scene where colour represents distance. Pixels with amplitude less than 1.8 mV have been fixed at 12.5 m (further than the distance to the far wall) in order for the image to make better use of the full colour scale.

amplitude of the correlation function for each pixel (on a logarithmic scale). The numbers indicate where the test targets described in Figure 5.13 can be found in the image. Brighter coloured targets, such as the white target for each pair give a larger amplitude pixel, as do nearer targets. The brightest target has amplitude of 650 mV peak-to-peak. Most of the far wall (except test target 13) and various other objects in the scene are black or very dark in colour and give a very small amplitude return, down to as little as 0.1 mV peak-to-peak.

Figure 5.15 shows a false colour image of the scene where the colour represents the measured distance for each pixel. With these operating parameters, all of the labelled test targets give a large enough amplitude return to give a reasonable measure of the targets' distance, which correspond well to their distance as measured using the DISTO distance meter. Objects in the scene that give less than 0.18 mV amplitude have had their distance fixed to 12.5 m. Truncating the scene in this manner serves to make better use of the full colour scale for the targets of interest.

The measurements have been repeated with the same measurement rate and N, this time using modulation frequencies of 40 and 32 MHz. Figure 5.16 (a) and Figure 5.16 (b) show images of the scene using 40 and 32 MHz respectively. The colour scale extends to 12.5 m, but each of these frequencies only have maximum unambiguous ranges of 3.75 m (40 MHz) and 4.69 m (32 MHz), so only the nearer portion of the colour scale is used. Objects in the scene beyond these ranges are incorrectly measured to be closer to the imaging system than they actually are. Large step changes in distance are especially recognisable in the 40 MHz image at the location of target 4, where due to only a small amount of noise the distance changes from 0 to 3.75 m across adjacent pixels.

Figure 5.16 (c) shows the combination of the 40 and 32 MHz images using Equations (5.10) and (5.11). The distance to most of the targets are now correctly established similarly to the 8 MHz image of Figure 5.15, however some pixels which gave a very small amplitude return (for example, parts of the chair and the far wall) are incorrectly measured by multiples of either d_{uA} or d_{uB} . Other notable differences between this image and Figure 5.15 are that object edges appear more clearly defined, and the variation in distance due to different target reflectance (for example, between targets 3 and 4) is reduced.



(a)





(c)

Figure 5.16: Images of test scene where colour represents distance captured using different modulation frequencies: (a) f = 40 MHz, (b) f = 32 MHz,
(c) combination of f = 40 MHz and f = 32 MHz. Colour scale is the same for all images.

5.6.2 Experimental Distance Precision versus Phase Precision

To investigate distance measurement precision a number of captures have been taken using simultaneous modulation frequencies with $f_A = 40$ MHz and $f_B = 32$ MHz as well as $f_A = 40$ MHz and $f_B = 8$ MHz. The integration ratio between them is fixed at 50%, phase steps per measurement fixed at 5, and measurement rate set to 10 Hz. 200 measurements have been captured at each setting. Phase, amplitude and distance measurements have been calculated for 8×8 groups of pixels at all 13 target locations. The resultant $8 \times 8 \times 13$ group of test pixels give a large range of actual distance values as well a wide range of amplitudes. σ_p and σ_d are calculated for all the test pixels to indicate their variation in phase and distance over the 200 measurements.

Recall that for the simulations presented in Section 5.3.1 the phase standard deviation, σ_p , for frequencies of 40 MHz and 32 MHz was assumed to be equal to the equivalent 8 MHz measurement scaled by factors of 1.64 and 1.48 respectively as per Equation (5.17). In order to make meaningful comparisons where σ_p from different frequencies is concerned, an equivalent σ_p has been used based on a similar scaling method. Figure 5.17 shows σ_p for 8, 32 and 40 MHz for a range of test pixels. In general, σ_p for 40 MHz is 1.64 times larger than that for 8 MHz and 32 MHz is 1.30 times larger. This precision for the 32 MHz measurement is slightly better than anticipated.



Figure 5.17: Plot of experimental σ_p for frequencies of 8, 32 and 40 MHz versus σ_p for 8 MHz. σ_p for 8 MHz scaled by 1.30 and 1.64 is also shown, which is used as the equivalent σ_p for these frequencies so that meaningful comparisons across frequencies can be made.



Figure 5.18: Logarithmic plot of experimental σ_d versus σ_p for test pixels measured using frequency combinations of 40 and 32 MHz, 40 and 8 MHz, and 8 MHz only.

Figure 5.18 shows the measured σ_d for all test pixels versus the equivalent (scaled) σ_p . The circle, triangle and diamond markers indicate experimental data, and the solid lines indicate the expected σ_d as per Equations (2.22) and (5.16). The 8 MHz only result follows the theoretical relationship exactly because Equation (2.22) simply scales σ_p to obtain σ_d . For the multiple frequency measurements the relationship is more complicated as σ_d is a function of the noise on two phase measurements, yet the x-axis of the plot only shows the standard deviation of the lower frequency. However, in general the results show a good correlation with the theory up to 100 milliradians. Beyond 100 milliradians, as was also the case with the simulated results presented in Figure 5.4, disambiguation errors in the distance measurement cause larger than expected σ_d .

5.6.3 Experimental Error Rate versus Phase Precision

The measured distance for the $8 \times 8 \times 13$ group of test pixels given above have been compared to the actual distance to each target as measured by the DISTO distance meter. Any distance measurement error of greater than half d_{uA} , that is, 1.875 m has been deemed an error, and the error rate is determined by the proportion of error measurements across the 200 measurements for each pixel.

Figure 5.19 shows a plot of the experimental and theoretical error rate versus σ_p (scaled as per the previous discussion) for frequency combinations of 40 and



Figure 5.19: Logarithmic plot of experimental and theoretical error rate versus σ_p for frequency combinations of 40 and 32 MHz, and 40 and 8 MHz.

32 MHz, and 40 and 8 MHz. The combination of two relatively high frequencies gives a higher error rate at lower σ_p as expected by Equations (5.21) and (5.22). The experimental results for this case agree well with the theoretical model, although the data disperses somewhat as the error rate decreases due to improved phase precision. This can be attributed to the low sample size, since for only 200 measurements the error rate resolution in this experiment is 0.5%.

For the high and low frequency combination the experimental results are in the general vicinity of the theoretical model, but do not align with it as well as for the two high frequency combination. A possible reason for this is that the scaling of the σ_p values is not as robust for very large σ_p , as for low amplitude signals the distribution becomes more uniform than Gaussian (Frank et al., 2009).

5.6.4 Experimental Distance Precision versus Integration Ratio

Using the test scene described above a number of measurements have been taken with the same set of frequency combinations, phase steps per measurement and measurement rate. However, for this experiment the ratio of time spent at each of the two frequencies has been varied from 0 to 100%. A 16×12 group of pixels located at target 8 have been analysed as this target has two interesting properties: 1) it is beyond the maximum unambiguous ranges of both 40 and 32 MHz modulation frequencies; 2) under these conditions the mean phase standard



Figure 5.20: Experimental σ_p versus integration ratio where $f_A = 40$ MHz and $f_B = 32$ and 8 MHz. Errorbars indicate the variation in σ_p across the 16 × 12 group of pixels.

deviation for the group of pixels is low enough that disambiguation errors arise reasonably often.

Figure 5.20 shows how the ratio between the time spent integrating for each frequency affects the standard deviation of the measured phase. The relationship is similar to that for the simulation presented earlier in Figure 5.10, where the best phase precision for 8 MHz modulation was 0.06 radians. In general, the lower 8 MHz frequency (min $\sigma_p = 0.052$ radians) has a smaller phase standard deviation than the 32 MHz frequency (min $\sigma_p = 0.070$ radians), which in turn is smaller than the 40 MHz (min $\sigma_p = 0.089$ radians). This can be attributed to a reduction in the demodulation contrast of the image sensor as the frequency is increased.

Similarly to the simulation presented in Section 5.5.3, when calculating distance standard deviation, outliers due to disambiguation errors have been removed from the data to give a better indication of the underlying distribution. Figure 5.21 shows the effect of the integration ratio on the distance standard deviation using a fixed weighting of 50% (dotted lines) and a variable weighting based on relative frequency and amplitude (solid lines). Error bars are also shown to indicate the variation in the standard deviation across the 16×12 group of pixels. Irrespective of the weighting scheme used, the combination of two relatively high frequencies gives a lower standard deviation. The experimental data relates very closely to the simulated behaviour presented in Figure 5.11.



Figure 5.21: Experimental σ_d versus integration ratio for frequency combinations of 40 and 32 MHz, 40 and 8 MHz, and 8 MHz only for pixels located at test target 8. Errorbars indicate the variation in σ_d across the 16 × 12 group of pixels.

5.6.5 Experimental Error Rate versus Integration Ratio

Figure 5.22 shows the effect of the ratio between the integration time for each frequency and the error rate for a group of pixels located at test point 8. The behaviour is very much as expected, following the simulated results shown in Figure 5.12 closely. The fewest errors when using two relatively high frequencies are obtained at approximately 50%, although this error rate is typically higher than the case of using a combination of high and low frequencies. For this sample set of 16×12 pixels over 200 measurements, no errors occurred for the high and low frequency combination at an integration ratio of less than 65%. This is likely due to the small sample set, and not an indicator that errors do not occur below this integration ratio, although the expected error rate is low.



Figure 5.22: Experimental error rate versus integration ratio for frequency combinations of 40 and 32 MHz, and 40 and 8 MHz for pixels located at test target 8.

5.7 Summary

This chapter has investigated the concept of errors arising from measuring distance using two modulation frequencies in time-of-flight range imaging systems. We have presented an efficient algorithm for combining two independent distance measurements to give a new measurement with increased precision and an extended maximum unambiguous range. The main drawback to this method is the possibility that if either of the individual measurements has sufficient noise, the calculated unambiguous distance may contain a large error, disproportionate to the level of noise.

The rate at which these errors occur can be statistically modelled based on the phase standard deviation. Furthermore, the amplitude of the signal can be used as an indicator of the measurement standard deviation, a fact that proves useful in applying a weighting between the two individual samples to increase overall system precision.

These techniques have been applied to test data captured using a custom built range imaging system based on the PMD19k image sensor. The experimental results compare well with the theoretical models, and demonstrate the viability of the proposed techniques.

Chapter 6 Real-time Distance Determination

The processing required to extract range data from the raw pixel intensity information requires a number of image frames to be stored and temporally processed. For real-time applications such as mobile robotics (Carnegie et al., 2005, Hussmann et. al., 2009) or gaming (Yahav et. al., 2007) it is desirable for the imaging system itself to process these frames in hardware, thereby reducing the load on the higher level processor or removing the requirement for one.

In this chapter we discuss the algorithm for calculating distance and intensity information from raw pixel intensity, and in particular how this algorithm is implemented in hardware, specifically a FPGA. Regarding FPGA resource usage, this design requires an amount of RAM proportional to the number of pixels of the image sensor. For the PMD19k system where an Altera Stratix III EP3SL150 FPGA is targeted, it is possible to use only on-chip block RAM. The advantage of on-chip RAM is that it can accommodate very large bus widths, and has very high bandwidth. The main limiting factor is the volume of RAM on the chip, which is typically reflected in the cost of the FPGA.

However, with the general trend toward image sensors with more pixels it will likely be uneconomical to move to a larger FPGA purely for more block RAM, instead the use of external RAM resources will be necessary. In this case, it is RAM bandwidth, not volume, that will need to be considered. For the analysis of RAM volume and bandwidth throughout this chapter let P be the number of pixels to be processed and f_{pixel} is the pixel clock frequency.

6.1 FPGA versus CPU

FPGAs have firmly established themselves as an excellent tool for high speed signal and image processing applications (Underwood, 2004). A number of comparisons have been made in the literature between inherently parallel FPGA processing to serial CPUs, where for many applications, FPGAs can perform the same calculations at a higher throughput (Guo et al., 2004).

The reasons for the higher throughput obtainable by FPGAs can be summarised by

- Parallelism Depending on the application algorithm, multiple operations can be performed in parallel within an FPGA for the same data set.
- Operations per Clock Cycle An FPGA can be configured to perform multiple operations within the same clock cycle. Indeed, some operations, such as a fixed bit shift do not require any additional time or hardware.
- Instruction Efficiency A well designed FPGA datapath uses a minimum number of operations, whereas a CPU inevitably requires some support instructions, for example fetching/storing data from/to memory, updating program counter, etc.

If clock rate is held constant, an FPGA can increase throughput by some orders of magnitude compared to a serial processor. Of course, clock rates are not constant between devices, but while many serial processors can operate with higher clock rates than FPGAs of a similar cost, the total throughput remains lower. Consider a simple example: the calculation of a complex multiplication

$$x_{real} + jx_{imag} = (a_{real} + ja_{imag})(b_{real} + jb_{imag})$$

$$x_{real} = a_{real}b_{real} - a_{imag}b_{imag}$$

$$x_{imag} = a_{real}b_{imag} + a_{imag}b_{real}$$
(6.1)

The operation requires 4 multiplications and 2 additions/subtractions. An FPGA implementation can produce the two results in two clock cycles: on the first clock

cycle the four multiplications are performed in parallel, and the addition and subtraction performed in parallel on the second clock cycle. As an example processor, the Altera Nios II/s core (Altera Corp., 2010) requires 16 instructions: 6 for the multiplication/addition/subtraction operations and 10 for loading from and storing to memory. Assuming a hardware multiplier (3 clock cycles per multiplication) the full operation requires at least 20 clock cycles; 10 times as many as the direct FPGA hardware implementation.

The latest trend in CPU design has been to incorporate multiple processor cores within the same device. A similar approach is taken by Graphics Processing Units (GPUs), which make use of system level parallelism through multiple processing pipelines. These technologies are further blurring the gap between FPGAs and serial processors, and which device is best for the task is becoming increasingly application specific. Algorithms requiring high memory bandwidth, for example 2D convolutions in image processing, typically achieve a higher throughput using FPGAs (Cope et al., 2005).

6.2 Unwrapping the Distance Calculation

Recall from Chapter 2 the equations for determining phase, and hence distance, from the raw pixel intensity data;

$$d = \frac{c}{2f} \frac{\varphi}{2\pi},\tag{2.5}$$

$$\varphi = \tan^{-1} \left(\frac{\sum_{i=0}^{N-1} I_i \sin(2\pi i/N)}{\sum_{i=0}^{N-1} I_i \cos(2\pi i/N)} \right), \qquad (2.10)$$

where I_i are the raw pixel intensity values from I_0 to I_{N-I} , f is the modulation frequency and c is the speed of light.

The phase calculation of Equation (2.10) can be approximated in fixed-point arithmetic as

$$p = ARCTAN(Real, Imag) = ARCTAN(\sum I_i C_i, \sum I_i S_i),$$
(6.2)

where I_i is the unsigned pixel intensity input with bit width, B, and a maximum value of 2^B-1 . C_i and S_i are cosine and sine approximations respectively, scaled to be signed integers in the range -2^{Bs} to $2^{Bs}-1$ with bit width Bs. The resultant phase, p, is also scaled to be an unsigned integer with bit width Br and a maximum value of $2^{Br}-1$. The range calculation must account for this factor by

$$d = \frac{2}{2f} \frac{p}{2^{Br}},$$
 (6.3)

although this final calculation need not necessarily be performed in the hardware, instead returning only *p*. The phase resolution is determined as $2\pi/2^{Br}$, therefore a 10-bit phase output has a resolution of 6.14 milliradians.

Also discussed in Chapter 2 is the signal amplitude which can be calculated by

$$A = \frac{2}{N} \sqrt{\left[\sum_{i=0}^{N-1} (I_i \cos(2\pi i/N))\right]^2 + \left[\sum_{i=0}^{N-1} (I_i \sin(2\pi i/N))\right]^2} .$$
(2.11)

Again, this can be approximated in fixed-point arithmetic as

$$A = \frac{2}{N}\sqrt{Real^{2} + Imag^{2}} = \frac{2}{N}\sqrt{\left[\sum I_{i}C_{i}\right]^{2} + \left[\sum I_{i}S_{i}\right]^{2}}.$$
 (6.4)

As for the phase calculation above, the multiplication by a constant need not be performed in hardware, or indeed at any stage, since most applications are concerned with relative amplitude as opposed to the amplitude in correctly scaled units.



Figure 6.1: Simplified block diagram for calculating phase, p, and amplitude, A from pixel intensity inputs IA_i and IB_i and sine and cosine values C_i and S_i .

For the system based on the PMD19k image sensor, an extra subtraction stage is required in the calculation. Each pixel of the sensor has two output channels which we label IA_i and IB_i . The active pixel intensity, I_i is equal to $IA_i - IB_i$. Figure 6.1 shows a simplified block diagram of the entire hardware calculation. The subscript *i* refers to the frame number, and the subscript *x* refers to the pixel number, although in most cases *x* is not shown as the calculations are equivalent for all pixels. Pixel data are streamed serially from the image sensor, and processing is performed within the datapath. This reduces the requirement for temporary storage of intermediate results.

6.3 The Accumulator Block

Figure 6.2 shows a block diagram of the elements involved with the accumulator function. To find the real and imaginary accumulator values, *Real* and *Imag* of Equation (6.2), three addition and two multiplication operations are required in the processing pipeline for every pixel loaded from the sensor. The sine and cosine values, C_i and S_i , are constant for each acquired image frame. They must be either calculated on the fly before the start of each frame (timing is not critical) or precalculated and stored in a look-up-table. The number of values required is equal to 2N, and since this is typically less than 20 and does not change during normal operation it is most efficient to precalculate and store them using FPGA on-chip distributed RAM (Jongenelen et al., 2010b).



Figure 6.2: Block diagram for accumulator function.

The use of RAM for temporary data storage is unavoidable in this design as each new frame must use the values for *Real* and *Imag* of the previous frame, noted above as $Real_{i-1}$ and $Imag_{i-1}$. The amount of RAM required in bits is proportional to twice the number of pixels multiplied by the bit width of each of the accumulators.

The values for *Real*_{*i*-1} and *Imag*_{*i*-1} are either the previous result which was stored in RAM or a zero value if *i* is equal to zero, indicating the first frame of the sequence. This ensures each new phase measurement is independent of previous measurements, and if the scene changes the resultant phase and amplitude will respond appropriately. It also enables an upper limit to be set on the *Real* and *Imag* values, allowing the bit width of the RAM to be correctly sized. For applications involving static scenes, this function may be disabled thereby incorporating an averaging function at no extra hardware cost, provided the RAM is large enough, or other measures are taken to prevent overflow, for example using floating point arithmetic.

6.3.1 RAM Requirements for Arbitrary *N*

A decision must be made as to the bit width of the sine and cosine values, *Bs*, as there will be some inherent quantization that occurs. As *N* frames are accumulated the quantisation errors of the sine and cosine values are also accumulated, as



Figure 6.3: The accumulation of worst-case errors in sine and cosine values must stay below the precision of the input in order to introduce no extra errors.

illustrated in Figure 6.3. The worst case error for a sine or cosine value is $1/2^{Bs}$, therefore, *Bs* must be selected such that the accumulation of *N* worst case errors does not exceed $1/2^{B}$, the precision of the raw pixel input. This leads to the relationship

$$Bs > B + \log_2(N)$$
. (6.5)

Similarly, the bit width of the accumulator must also be at least *Bs*. However, more width is required as with the accumulation of *N* frames there must be some guard bits at the most significant end to prevent overflow. The worst case pixel input is a square wave of maximum amplitude and 0 phase offset, that is, 2^{B} -1 for the first *N*/2 samples, 0 thereafter. The accumulated imaginary value after *N*/2 samples will reach a maximum dependent on *N* by

$$Imag_{max} = \sum_{i=0}^{N/2} (2^B - 1)\sin(2\pi i/N) = (2^B - 1)\frac{N}{\pi}.$$
 (6.6)

The number of extra bits required to prevent overflow is $\log_2(N/\pi)$. The final bit width of the accumulator, *Ba*, to give negligible quantisation error relative to the input signal is given by the ceiling function

$$Ba = \left\lceil B + \log_2\left(\frac{N^2}{\pi}\right) \right\rceil.$$
(6.7)

The total required RAM in Bits for both the real and imaginary accumulators is

$$RAM = 2BaP. (6.8)$$

The RAM is accessed four times per pixel, two reads and two writes, so the required bandwidth in Bits/s is equal to

$$BW = 4Baf_{pixel}.$$
 (6.9)

When using a combination of multiple modulation frequencies, a single-bin FFT is required for each frequency. Therefore the accumulators are independent, multiplying the above RAM volume and bandwidth requirements by the number of modulation frequencies used.

Alternatively, a different approach can be taken where instead of storing the accumulated result, only the raw frames from the sensor are stored. In this case the required bit width is *B* and total RAM is BP(N-1). For each newly received sample, the previous *N*-1 samples must be read from RAM, therefore total RAM bandwidth is $Bf_{pixel}N$ where there are *N*-1 reads and 1 write. The number of multiplications required per pixel clock cycle is increased from 2 to 2*N*.

At first glance storing the raw frames seems impractical as the required amount of RAM is significantly higher, however there are two notable advantages to such a system

- 1. Since there is always a full cycle of samples stored, a sliding window approach can be taken and hence a new result can be calculated at the same rate as the sampling rate. While this will not affect precision it may prove useful in applications which require a fast response to sharp changes in the scene. For instance, actual measurement values may not stabilise until the end of a full measurement cycle, but changes will be noticed immediately.
- 2. When multiple modulation frequencies are encoded simultaneously it may prove more efficient to store the raw frames instead of instantiating individual accumulators to process each frequency. For example, with two modulation frequencies and N = 5, the required RAM using two sets of

accumulators (as per Equation (6.8) and then doubled) would be 4(B+3)P, whereas storing the raw frames requires only 4BP.

6.3.2 Output Buffers

The output buffers are memory allocated for temporarily storing the calculated phase and amplitude to be read out by a subsequent process, for example a processor handling communications to transfer results to a PC. Given an input bit width *B*, the quantisation error on the raw pixel intensity will be at most $1/2^B$ due to quantization. If we assume that the error on the input is small compared to the amplitude of the signal, then the phase error in radians is approximately equal to the input error as shown in Figure 6.4. Due to the effect of averaging over *N* normally distributed intensity samples, the noise due to quantisation is reduced by $\sqrt{N}/2$, requiring an additional $\log_2(\sqrt{N}/2)$ bits to represent. To maintain this precision over a full 2π radians requires a further $\log_2(2\pi)$ bits. The final bit width for the phase output buffer to account only for quantisation error can be calculated as the ceiling function

$$Br = \left[B + \log_2 \left(\frac{2\pi\sqrt{N}}{\sqrt{2}} \right) \right]. \tag{6.10}$$

It is likely that a DC offset will be present on the input signal, limiting the dynamic range to a figure less than 2^{B} . In this situation it is more convenient to work in terms of the input signal-to-noise ratio, SNR_{I} . Br can therefore be calculated by the alternative formula

$$Br = \left\lceil \log_2 \left(\frac{2\pi SNR_I \sqrt{N}}{\sqrt{2}} \right) \right\rceil.$$
(6.11)

The RAM required for the output buffer to store both phase and amplitude results is therefore 2BrP. The bandwidth is dependent on the process which is reading the output buffers, although in most cases it will read it at at least the same rate as



Figure 6.4: Phase error is approximately equal to input error if the error is small relative to the amplitude.

the processing algorithm is writing into it. Therefore the required bandwidth for the output buffer is $4Brf_{pixel}$.

6.3.3 **RAM Requirements for** N = 4

A special case exists where N = 4 as the sine and cosine values reduce to $\{-1, 0, 1\}$ and Equation (6.2) simplifies to

$$p = ARCTAN(I_1 - I_3, I_0 - I_2).$$
(6.12)

There are no multiplication operations required. The RAM required to implement this method is 2(B+1)P, as the sine and cosine values can be exactly represented using only 2 bits, and the result of the subtraction will never need more than 1 extra bit to prevent overflow. The required RAM bandwidth is reduced to $2(B+1)f_{pixel}$ as, for each new sample, only one of the real or imaginary accumulators is accessed at a time.

This is the most efficient implementation in terms of RAM usage, but it is inflexible as N cannot be changed. Research has shown that range accuracy might be impaired in systems when using 4 samples per measurement, as odd order harmonics are aliased onto the fundamental (Lange and Seitz, 2000). However, there are ways to mitigate this error including Heterodyning (Dorrington et al., 2007) and Harmonic Rejection techniques (Payne et al., 2008b), or by selecting alternative values of N to reduce the aliased odd-order harmonic content.

6.4 Arc Tangent and Amplitude Calculation

A number of options exist for calculating the arc tangent function efficiently in hardware (Bellis and Marnane, 2000). The simplest method is to use a look-up-table indexed using the result of the division of the real by the imaginary part of the equation (Jongenelen et al., 2008). This technique is costly in terms of RAM for the look-up-table as well as general purpose FPGA logic resources required to implement the division, although efficient division implementations do exist (Bailey, 2006).

A more efficient implementation is to use a Coordinate Rotation Digital Computer (CORDIC) algorithm (Walther, 1971). The CORDIC arc tangent algorithm works by rotating the phase of a given input vector to reduce the imaginary component to zero. If x_0 and y_0 are the real and imaginary components of the input vector respectively and the accumulated phase, z_0 , is initialised to 0, then after *i* iterations subsequent values for *x*, *y* and *z* are computed as (Andraka, 1998)

$$x_{i+1} = x_i - y_i \cdot d_i \cdot 2^i$$

$$y_{i+1} = y_i + x_i \cdot d_i \cdot 2^i$$

$$z_{i+1} = z_i - d_i \cdot \tan^{-1}(2^{-i})$$

(6 13)

where the decision variable, d_i , is 1 if y_i is less than zero, otherwise –1. Figure 6.5 illustrates an example arc tangent calculation using CORDIC.

A CORDIC based implementation is desirable as it can be realised using only shift and addition operations, and without the use of any RAM resources. In addition, the amplitude result from Equation (6.4) is obtained as a by-product of determining the phase, requiring no additional computation.

The angles, $\tan^{-1}(2^{-i})$, can be stored in a small lookup table when implementing CORDIC arc tangent. They can also be scaled by an arbitrary factor, enabling the range to be calculated directly without any additional multiplications. The accuracy of a fixed-point CORDIC arc tangent implementation is determined by the bit width of the phase look-up-table and intermediate results, and the number



Figure 6.5: Given an input vector expressed in rectangular form, $x_0 + iy_0$, CORDIC arc tangent function finds the sum of phase rotations required to reduce the imaginary part to zero.

of CORDIC iterations. For small angles, the value of $atan(\theta) \approx \theta$, so provided there are enough bits to work with, each iteration increases phase accuracy by two (Walther, 1971).

6.5 Implementation Results

The phase determination algorithm described in Section 6.2 has been implemented in VHDL on the full-field range imaging described in Chapter 3. An investigation into the trade-offs of FPGA and RAM resources versus measurement precision have been conducted, and the results presented here.

6.5.1 Test Methodology

Captures of a scene have been taken under conditions as close to ideal as possible; all practicable steps have been taken to reduce noise, particularly from ambient light such as sunlight and indoor lighting. This provides a measure of the best precision achievable for this particular system, and therefore allows an efficient processing method to be tailored to this system. Scene setup is identical to that described in Chapter 4, with a single white target positioned 2.63 m from the sensor. The actual value of N for the captures is 20, modulation frequency of

10 MHz and the total number of cycles processed is 25. Measurement rate has been varied from 1 to 3 Hz to give a variation in phase precision to be processed.

6.5.2 Hardware Processing Accuracy

Figure 6.6 shows a plot of phase standard deviation, σ_{φ} , versus amplitude, *A*, for all pixels of the imaged scene as processed using MATLAB. Even in these ideal conditions, the amplitude is never larger than approximately one eighth of the full scale 16-bit input, which indicates that providing $\log_2(N)$ extra bits to prevent overflow is being overly pessimistic. For the pixels with high amplitude, the standard deviation is approximately 1 milliradian. This corresponds to a dynamic range for the phase measurement of approximately 76 dB and can be immediately useful in scaling the output buffer width, *Br*, to 13 bits. By working backwards using Equation (6.11) we can estimate the SNR of the input as 45.2 dB.

Figure 6.7 shows the percentage of all pixels with phase standard deviation of less than the x-axis. The relevance of this plot is that it acts as a benchmark for analysing the processing of the pixels in hardware. When processing in hardware it is important that any phase error introduced is less than σ_{φ} .



Figure 6.6: Phase standard deviation, σ_o , versus amplitude, A, for all pixels as calculated by MATLAB.



Figure 6.7: Percentage of all pixels with phase standard deviation less than the x-axis.



Figure 6.8: Phase error for all pixels processed in hardware (with respect to MATLAB processing) versus accumulator bit widths (*Ba*) of 10 and 16.

The same data has then been processed in the hardware with varying bit widths used for the accumulators, Ba, – the largest factor determining RAM requirements. Figure 6.8 shows plots of the introduced phase error, Δ_{φ} , from the hardware processing (compared to MATLAB processing) versus amplitude. The bit width of the accumulator, Ba, is shown for 10 and 16 bits. It can be seen that for low amplitude signals, an accumulator width of 10 bits gives a very large

error, sometimes as high as π radians, because there are not enough significant bits for the CORDIC algorithm to make a meaningful approximation of the phase. Indeed, for the 10 bit accumulator width the phase error is generally larger for all pixels, and often higher than the corresponding standard deviation expected for that amplitude as was shown in Figure 6.6. For an accumulator bit width of 16, the error introduced from the hardware processing is less than the standard deviation of the measurement.

Figure 6.9 shows the percentage of pixels processed in hardware with phase error (compared to MATLAB processing) less than the x-axis. Each line represents processing using a different accumulator bit width, from 10 to 16. The dashed line in the graph shows the standard deviation from Figure 6.6, which represents the target accuracy to ensure that errors introduced by the processing do not have a large influence on overall precision. It can be seen that for bit widths greater than or equal to 15, 100% of the pixels are processed to an accuracy better than the phase standard deviation, hence. This reinforces the estimated required bit width given in Equation (6.7), where N is 20 and B is 8. Although a bit width of 14 bits does not always meet the benchmark, it is probably sufficient because it exceeds the requirements for over 95% of the pixels in the image.



Figure 6.9: Percentage of all pixels processed in hardware with phase error (with respect to MATLAB processing) less than the x-axis for varying accumulator bit widths (*Ba*).
6.5.3 **FPGA Resource Requirements**

Resource requirements in FPGAs are commonly specified as the number of basic combinatorial logic elements required for a design. In the Stratix III FPGA the basic logic element is the Adaptive Logic Module, ALM. Figure 6.10 shows the FPGA resources used to implement the phase determination algorithm (not including output buffer) for an arbitrary value of N up to 25 (before possible accumulator overflow). The RAM requirement closely follows Equation (6.8) where the number of pixels, P, is $120 \times 160 = 19,200$ for the PMD19k image sensor. It differs in that actual storage has been implemented for 128×160 pixels, as 128 rows are much simpler to address than 120, and simpler to construct from smaller elements. The FPGA used in this system has enough on-chip Block RAM that external RAM is not required, however it is desirable to reduce this requirement where possible as it can be expected that new image sensors will have more pixels, and upgrading the FPGA purely to gain more Block RAM is not cost effective.

The CORDIC function accounts for approximately 90% of the ALM resources used for the algorithm. For accumulator widths less than or equal to 17 the number of CORDIC iterations is equal to the accumulator width. The number of CORDIC iterations in this design has been capped at 17, hence beyond this accumulator width the ALM requirement plateaus. The phase accuracy of CORDIC beyond 17 iterations (less than 0.1 milliradians) is far better than will ever be necessary for this system.

From the phase error results shown in Figure 6.9 it can be seen that an accumulator bit width of 15 bits yields acceptable results for this particular system. With two output buffers 13 bits wide, total required RAM is 1120 kBits, which equates to 20% of total on-chip RAM available on the FPGA we have used. The 611 ALMs required equates to only 1% of FPGA logic resources available.



Figure 6.10: FPGA resource requirements versus accumulator bit width.

6.6 Summary

In this chapter an efficient real-time FPGA algorithm for determining phase, and hence distance, for pixels in a ToF range imaging system has been presented. The scarcest resource in such a system is RAM, and an analysis has been presented to maximise the efficiency of this resource whilst maintaining acceptable processing accuracy.

The algorithm has been implemented on a range imaging system based on the PMD19k-2 image sensor, with range processing performed in real-time by a Stratix III EP3SL150 FPGA. Particular attention has been paid to the trade-off of the bit widths of the accumulators and the processing accuracy. The results have shown that for this particular system an accumulator width of 15 bits gives a phase error from the processing less than the system standard deviation as calculated by MATLAB. With a bit width of 13 for each of the two output buffers this equates to 1120 kBits of RAM, which is 20% of the RAM available on the FPGA we have used, whereas only 1% of general purpose logic elements are used.

In this case, the Stratix III FPGA, while suitable for testing as a prototype, is much larger than necessary. A more cost effective solution could be obtained using a lower-end FPGA with sufficiently high bandwidth external RAM. Such a solution will also scale better as the number of pixels in state-of-the-art image sensors increases.

Chapter 7 Evaluation and Conclusion

This chapter summarises the work presented in this thesis. A number of range imaging examples are presented to show high precision and high frame rate operating modes. Then, a summary of the thesis and some ideas for the future direction of this project are given. A number of international publications have arisen from the work presented in this thesis, and a list of these is presented in this chapter. Finally, the thesis is concluded with a statement detailing the original contributions of this work.

7.1 Range imaging examples

In order to demonstrate the flexibility of the Victoria University range imaging system, two scenes have been captured. In the first example a static scene is captured with the system operating in a high precision mode, in the second example objects are moving within the scene and the system is operating in a high frame rate mode.

7.1.1 Static Scene

For this capture a static scene has been set up with one object dominating the frame: 'Rebecca' the mannequin head positioned approximately one metre from the ranging system. A colour photograph of the scene is shown in Figure 7.1. System operating parameters have been selected for maximum precision: modulation frequency of 36 MHz, phase steps per measurement of 30 and integration time per frame of 20 ms giving a range image measurement rate of 1.67 Hz. Because the object is so close to the camera, the integration time per frame is relatively short for a high precision operating mode in order to prevent saturation. A total of 100 measurements have been taken over one minute and averaged to improve precision.



Figure 7.1: Colour photograph of a static scene with 'Rebecca' the mannequin head.

Figure 7.2 shows false colour images of the measured distance (left) and amplitude (right) for the mannequin's face. Objects beyond 1100 mm have been truncated in order to make better use of the colour scale. Large facial features such as the nose and neck can be determined, but the image is very noisy due to fixed pattern noise. To correct for this a very simple calibration has been performed by capturing an image of a flat surface (right image of Figure 7.3) and subtracting the data from the noisy distance data of the mannequin. The resultant distance image shown in the left of Figure 7.3 is noticeably smoother with improved definition of the cheek, forehead and chin.

Figure 7.4 shows a 3D reconstruction of the mannequin's face, with intensity data overlaid. Darker parts of the image, for example the eyes, eyebrows and hair, exhibit a systematic distance error tending toward the camera (similar distance as the nose). This can be attributed to the mixed pixel effect described in Chapter 2. No compensation for this has been attempted in this case.



Figure 7.2: False colour images of the mannequin's face depicting distance (left) and amplitude (right)



Figure 7.3: False colour distance image of the mannequin's face (left) after calibration by subtracting measured distance of a flat surface (right).



Figure 7.4: 3D Reconstruction of the mannequin's face, with intensity data overlaid.

7.1.2 Dynamic Scene

In this example the system is operating in a high frame rate mode at the expense of precision. A video sequence has been taken of a person moving within the scene, and can be found on the companion CD for this thesis. Figure 7.5 shows a colour image of the scene and person. Phase steps per measurement are set to the minimum value possible of three. Integration time per frame is 20 ms, giving a measurement frame rate of 16.7 Hz. Modulation frequency has been chosen as 24 MHz, which gives a maximum unambiguous range with a single frequency of 6.25 m – larger than the extent of the room.

Figure 7.6 shows a sequence of images where the person is walking toward the camera. Even in this high frame rate mode of operation, large objects such as the person, the box, the chair and the far wall give sufficient information to establish a coarse measure of their distance. This feature is particularly useful for obstacle avoidance in a mobile robotics application. One notable consequence of objects moving within the scene is the blurring of edges when moving laterally to the camera.



Figure 7.5: Colour photograph of a person standing in a moving scene.

Due to the temporal nature of the distance measurement it is possible for an object to be imaged by a pixel in one frame and then removed for subsequent frames, where a background object is imaged instead. This can lead to a situation where the measured distance for that pixel is neither relevant for the object nor the background – sometimes lying somewhere between, others times protruding towards the camera. Particularly for a robust obstacle detection application, these errors will need to be dealt with. One possible solution is post processing of the data to filter out sharp changes in measured distance.

The last image in Figure 7.6 includes two grids indicating locations on the white box and the cupboards on the back wall where precision has been measured. The standard deviation throughout the video sequence for the pixels imaging the white box is 18.3 ± 3.9 mm, which is approximately 0.6% of the measured distance. For the pixels imaging the cupboard, the standard deviation is 200.5 ± 32.5 mm, which is approximately 4.8% of the measured distance. These results compare reasonably well to the commercial systems presented in Chapter 2 where depth precision ranged from 0.075% to 1.5%. However, care should be taken when comparing these results as few experimental conditions can be considered constant.



Figure 7.6: Sequence of distance images showing a person walking toward the camera. The boxed region is used to find precision statistics.

7.2 Thesis Summary

A new full-field range imaging system has been developed. The work in this thesis can be divided into two parts: the design, construction and characterisation of the new system, and the use of the system as a tool for testing range imaging concepts such as dual-frequency modulation and real-time distance processing.

7.2.1 Development of the Victoria University range imaging system

A significant output of this thesis is the development of the Victoria University range imaging system. The main objectives for this system were for it to be compact, portable and provide full-field distance measurement in real-time.

The real-time distance processing aspect was addressed first. New printed circuit boards were designed and constructed to interface an Altera Stratix II FPGA to the DALSA Pantera 1M60 camera used by the University of Waikato Range Imaging System. This required a detailed understanding of the serial LVDS CameraLink interface in order to correctly communicate with the DALSA camera and transfer image frames directly to the FPGA. From these raw intensity frames, distance is calculated in real-time within the FPGA and displayed on a VGA monitor or transmitted to a PC via an Ethernet connection. The major achievement of this design was to demonstrate real-time range imaging – however an entirely new range imaging system was needed to make the system more compact and portable.

The most technologically challenging aspect of ToF range imaging is to produce a very high-speed (order of MHz) shutter mechanism on the image sensor. With the University of Waikato system this was realised by coupling the camera to an image intensifier tube, where gain was controlled by modulating the voltage of the photocathode. While this system gave excellent measurement precision (< 1 mm), it was very large and required a relatively large amount of power, including two high voltage supplies of 700 and 7000 V – unsuitable for a mobile robotics application. Therefore, the more compact and power economical

Victoria University range imaging system was designed and built based on the PMD Technologies PMD19k image sensor.

The primary advantage of the PMD19k image sensor is that gain modulation of the receiver is controlled on-chip, removing the requirement for the large and power hungry image intensifier. The final design consists of a number of subsystems:

- the illumination board to host 16 laser diodes for high-power AMCW illumination;
- the PMD Daughterboard and PMD Mainboard to host the PMD19k image sensor, convert the analogue video stream to digital data and interface to the FPGA;
- the VGA DAC and Ethernet I/O board for connecting to a VGA monitor for real-time image display and PC communication;
- the Altera Stratix III Development Kit for distance measurement processing, modulation signal generation and high level control of the system.

The Altera Stratix III FPGA was chosen as it includes a number of PLLs where the frequency and phase of the output can be reconfigured in real-time. This is a very important feature that enables the all-important precise control of the frequency and phase of the image sensor and illumination modulation signals. The system is flexible, allowing both automated and manual user control of operating parameters such as modulation frequency, duty cycle, phase steps per measurement and integration time. More complex modulation techniques are also possible, for example, changing the phase or frequency within the integration period to produce simultaneous dual-frequency modulation.

This work has also described the characterisation of the Victoria University range imaging system. The main outcome of the characterisation has been to establish the precision and accuracy of the system in a variety of operating conditions. Based on the amplitude of the correlation waveform, peak precision is obtained at a modulation frequency of 36 MHz, which gives a maximum unambiguous range of 4.17 m when using a single modulation frequency. Actual obtainable precision

is dependent on a number of factors, many of which are application and scene dependent. In general, brighter and more highly reflective objects give increased precision provided the sensor does not saturate, and integrating over a longer time period or a larger number of frames also increases precision. As an example, a white target positioned 2.63 m from the sensor measured using 36 MHz modulation frequency, four phase steps per measurement and measurement frame rate of 25 Hz gave a 1-sigma standard deviation of 4.08 mm, or 0.156% of target distance. However, it was also shown that four phase steps per measurement produced a systematic error due to the aliasing of harmonics, with RMS error of 9.13 mm under the settings given above. A good compromise was established using five phase steps per measurement, where the RMS error due to harmonics is an order of magnitude lower while still maintaining a high frame-rate.

7.2.2 Dual-frequency Modulation and Real-time Processing

The easily reconfigurable design of the Victoria University range imaging system has been essential for implementing and evaluating a number of more complex ToF distance measurements methodologies. One such methodology is the use of multiple modulation frequencies in order to extend maximum unambiguous range whilst maintaining high measurement precision. In this thesis we have discussed the advantages and disadvantages of a dual-frequency system in terms of distance measurement precision, and the expected rate of 'disambiguation errors', where due to insufficient precision of either of the phase measurements the resultant distance is incorrectly estimated by some metres.

A novel algorithm to obtain unambiguous distance from two independent phase measurements has been presented based on a modification of the New Chinese Remainder Theorem. This algorithm has been used as the basis for determining the expected standard deviation and error rate for resultant distance measurements based on the precision of the phase measurements. In addition, the signal amplitude is used to act as an indicator of the phase precision, allowing the algorithm to appropriately weight each input and provide a higher precision output. The dual-frequency modulation technique has been experimentally tested using the system described earlier. While there were some discrepancies, mostly due to the assumption made about the constant nature of phase precision versus modulation frequency, the results compared well to the theoretical model and demonstrated the viability of the dual-frequency technique.

One of the key objectives of this research has been to implement range image processing in real-time. This has been achieved using a fixed-point arithmetic implementation based on an Altera Stratix III FPGA. Additionally, an analysis has been presented regarding the trade-off between introduced quantisation errors and resource usage, particularly the typically scarce resource of RAM. The main result of this analysis has been to provide a more meaningful measure of the FPGA requirements as determined by the desired measurement precision. It is also noted that a more affordable system utilising a lower specification FPGA is possible, if sufficient external RAM resources are available.

7.3 Future Work

The work presented here has exceeded the requirements for the original objectives. However, a number of additional enhancements can be made to further improve the system.

For static scenes, measurement precision can be increased by averaging over a large number of frames. Nevertheless, there comes a point where systematic errors exceed the random errors, thereby limiting the maximum depth resolution of the system. This system would benefit from a much more robust calibration routine to account for errors due to geometric distortions, latency and fixed pattern noise. Errors related to mixed pixel effects (also known as scattering) are scene specific, and therefore impossible to account for through calibration. The next step in real-time distance processing is to include static corrections from data gained through calibration, as well as investigate incorporating processing techniques for real-time scattering compensation.

For dynamic scenes, the ability to measure velocity would be an excellent feature. The effect of movement can be divided into two cases: for lateral movement a sharp change in phase will be evident as the edge of an object moves across the boundary of a pixel; for axial movement it can be likened to a Doppler-effect, where the modulation frequency is subtly altered. This will require the development of new or modified algorithms for phase determination.

The system hardware itself can benefit from a number of improvements. The first possibility is to source a newer model image sensor IC with an increased number of pixels, and hence higher spatial resolution. This is an ongoing trend with image sensor IC manufacturers, for example, PMD Technologies currently have a new model available with 204×204 pixels. It is also expected that the precision of the current system can be improved by further increasing the modulation frequency. The main limitation of the current design is the ability of the Intersil EL7158 modulation driver ICs to drive the two 1000 pF loads of the image sensor at high frequency. The next generation design will divide the image sensor into four independent blocks of 250 pF, each with independent driver ICs. The downside to this is a possible loss of synchronisation causing errors in the distance measurement between blocks. However, if the effect is appropriately characterised it can be compensated for through calibration. Another solution is to source or design a more advanced driver IC that can drive the high capacitive load at higher frequency.

The Stratix III Development Kit has been a key component for this system as the FPGA was one of the first to include real-time reconfigurable PLLs. The disadvantage of this system is that a number of components included with the development kit are unnecessary for the design, and consequently the whole system is larger than is necessary. A next generation design can make use of new FPGAs which include reconfigurable PLLs in the lower end product families, for example, the Altera Cyclone III. By migrating to a dedicated printed circuit board with a bare minimum of components the cost and physical dimensions of the system can be reduced.

The investigation into dual-frequency modulation can also be extended to an arbitrary number of modulation frequencies. It is expected that with *n* modulation frequencies related by the co-prime factors M_1 , M_2 , ..., M_n , the maximum unambiguous range will be related to $M_1 \times M_2 \times \ldots \times M_n$. At currently achievable modulation frequencies this may lead to extremely large unambiguous ranges. Another interesting possibility is that with an increase in the number of frequencies some redundancy can be included in the measurement, thereby potentially identifying or correcting disambiguation errors.

By utilising the flexibility of this system, a number of more complex operating schemes can be studied. One such possibility is to implement an automatically adjusting integration time mechanism such that integration time is maximised and no pixels are saturated (or some minimum threshold of pixels). In terms of the real-time hardware processing, a DC offset on the correlated signal limits the amplitude of the signal that can be processed before the accumulators overflow. A mechanism for subtracting the DC offset, either on a per pixel or entire frame basis, could reduce the likelihood of accumulator overflow and potentially reduce accumulator RAM resource usage.

7.4 Publications Arising from this Thesis

The work presented in this thesis has resulted in the primary authorship of papers accepted for publication in the IEEE Transactions on Instrumentation and Measurement and the Journal of Real-time Image Processing. It has also resulted in a book chapter contribution to a special issue of the Lecture Notes on Electrical Engineering, and papers published in the proceedings of three conferences. This work has also contributed towards four other publications where I am listed as a co-author.

Jongenelen, A. P. P., Bailey, D. G., Payne, A. D., Dorrington, A. A. and Carnegie, D. A. (2011). "Analysis of Errors in ToF Range Imaging with Dual-Frequency Modulation," *IEEE Transactions on Instrumentation and Measurement*, vol. 60 (5), pp 1861 – 1868.

Jongenelen, A. P. P., Bailey, D. G., Payne, A. D., Carnegie, D. A. and Dorrington, A. A. (2010). "Efficient FPGA Implementation of Homodyne-Based Time of Flight Range Imaging," *Journal of Real-time Image Processing*, Springer-Verlag, published online.

Jongenelen, A. P. P., Payne, A. D., Carnegie, D. A., Dorrington, A. A. and Cree, M. J. (2009). "Development of a Real-Time Full-Field Range Imaging System," in *Recent Advances in Sensing Technology, Lecture Notes on Electrical Engineering*, pp. 113 – 129, Springer-Verlag.

Jongenelen, A. P. P., Carnegie, D. A., Payne, A. D. and Dorrington, A. A. (2010). "Maximizing Precision Over Extended Unambiguous Range for TOF Range Imaging Systems," *Proc. International Instrumentation and Measurement Technology (I2MTC'10)*, pp. 1575 – 1580.

Jongenelen, A. P. P., Carnegie, D. A., Payne, A. D. and Dorrington, A. A. (2009). "Development and Characterisation of an Easily Configurable Range Imaging System," *Proc. Image and Vision Computing New Zealand (IVCNZ'09)*, pp. 79– 84.

Jongenelen, A. P. P., Payne, A. D., Carnegie, D. A. and Dorrington, A. A. (2008) "Heterodyne Range Imaging in Real-time," *Proc. International Conference on Sensing Technology (ICST'08)*, pp. 57–62.

Carnegie, D.A., McClymont, J.R.K., Jongenelen, A.P.P., Drayton, B., Dorrington, A.A., and Payne, A.D. (2010). "Design and Construction of a Configurable Full-Field Range Imaging System for Mobile Robotic Applications," *New Developments in Applications in Sensing Technology*, Springer-Verlag, publication pending.

Carnegie, D. A., McClymont, J., Jongenelen, A. P. P., Dorrington, A. A., Payne,A. D. (2010) "Development of a Full-Field Range Imaging Sensor," *Proc. International Conference on Sensing Technology (ICST'10)*, publication pending. Payne, A. D., Jongenelen, A. P. P., Dorrington, A. A., Cree, M. J. and Carnegie, D. A. (2009). "Multiple Frequency Range Imaging to Remove Measurement Ambiguity," *Proc. Optical 3-D Measurement Techniques*, pp. 139–148.

Dorrington, A. A., Cree, M. J., Carnegie, D. A., Payne, A. D., Conroy, R. M., Godbaz, J. P. and Jongenelen, A. P. P. (2008). "Video-rate or high-precision: A flexible range imaging camera," *Proc. SPIE-IS&T Electronic Imaging*, vol. 6813.

7.5 Conclusion

Original contributions made by this thesis include:

- The design and construction of a full-field range imaging system that is adaptable to a range of applications, from high-precision to high frame-rate operating modes;
- The integration of a real-time reconfigurable phase-locked-loop device, which has enabled operating parameters of the design to be updated in real-time. This can be either automated, or controlled directly by the user;
- The ability to reconfigure operating parameters in real-time allows for a much broader range of modulation schemes to be implemented and evaluated;
- The development of a novel algorithm for calculating distance from two phase measurements using the modified New Chinese Remainder Theorem;
- A detailed examination of the improved precision obtainable by a dualfrequency system, and the associated cost in terms of expected error rate;
- An in depth analysis of FPGA resources used by the design and the trade-off between resource requirements and best obtainable measurement precision, allowing a processing unit to be scaled appropriately for a given application.

The main motivation of this research has been to design a full-field range imaging system that is suitable for mobile robotic applications with the added flexibility necessary to evaluate novel modulation and real-time distance processing techniques. The constructed prototype design meets the objectives for a mobile robot application as follows:

- Configurable measurement rate up to 50 Hz;
 - High-speed mode with precision of tens of centimetres to allow the robot to respond quickly to changes in its environment;
 - High-precision mode with sub-centimetre precision to analyse objects in the environment, for example, a person's face;
- Low voltage and power requirement, which allows the system to operate for extended periods from a battery power supply;
- Compact size reduces strain on mechanical elements and increases usability in confined spaces;
- Embedded real-time processing reduces the load on the robot's main CPU, if it has one.

The final result is a new 3D sensor for enabling machine vision. With the ever expanding study of autonomous robotic systems, this technology will be a key element in constructing machines capable of operating in even the most complex environments.

References

- 3M, (2005). "Mini D Ribbon (MDR) Connectors," Technical Datasheet, www.3m.com
- Altera Corporation, (2007). "Stratix II Device Handbook, vol. 1," www.altera.com
- Altera Corporation, (2008a). "Nios II Software developers Handbook, vol. 1," www.altera.com
- Altera Corporation, (2008b). "DE2 Development and Education Board," www.altera.com
- Altera Corporation, (2008c). "Stratix III Device Handbook, vol. 1," www.altera.com
- Altera Corporation, (2010). "Nios II Processor Reference Handbook," www.altera.com
- Analog Devices, (1998). "ADV7123 CMOS, 240 MHz Triple 10-Bit High Speed Video DAC," Technical Datasheet, www.analog.com
- Analog Devices, (2001). "AD9826 Complete 16-bit Imaging Signal Processor," Technical Datasheet, www.analog.com
- Analog Devices, (2007). "AD5301/AD5311/AD5321 2.5 V to 5.5 V, 120 μA, 2-Wire Interface, Voltage-Output 8-/10-/12-Bit DACs," Technical Datasheet, www.analog.com
- Andraka, R., (1998). "A survey of CORDIC algorithms for FPGA based computers," *Proc. ACM/SIGDA Sixth International Symposium on Field Programmable Gate Arrays*, pp. 191-200.
- Atkinson, K. B., (2001). *Close Range Photogrammetry and Machine Vision*, Whittles Publishing, Scotland, UK.
- Bailey, D. G., (2006). "Space Efficient Division on FPGAs," Proc. Electronics New Zealand Conf., pp. 206-211.
- Bellis, S. J. and Marnane, W., P., (2000). "A CORDIC Arctangent FPGA Implementation for a High-Speed 3D-Camera System," *Field-Programmable Logic and Applications: The Roadmap to Reconfigurable Computing, Lecture Notes in Computer Science*, vol. 1896, pp. 485-494.

- Besl, P. J., (1988). "Active, optical range imaging sensors," *Machine Vision and Applications*, vol. 1, pp. 127-152.
- Boehm, J. and Pattinson, T., (2010). "Accuracy of Exterior Orientation for a Range Camera," *Proc. ISPRS Commission V Mid-Term Symposium: Close Range Image Measurement Techniques*, vol. 38 (5), pp. 103-108.
- Busck, J. and Heiselberg, H., (2004). "Gated viewing and high-accuracy threedimensional laser radar," *Applied Optics*, vol. 43 (24), pp. 4705-4710.
- Buttgen, B., Oggier, T., Lehmann, M., Kaufmann, R. and Lustenberger, F., (2005). "CCD/CMOS Lock-In Pixel for Range Imaging: Challenges, Limitations and State-of-the-Art," *Proc. Range Imaging Research Day*, pp. 21–32.
- Buttgen, B. and Seitz, P., (2008). "Robust Optical Time-of-Flight Range Imaging Based on Smart Pixel Structures," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 55 (6), pp. 1512-1525.
- Carnegie, D. A., Loughnane D. L. and Hurd S. A., (2004). "The design of a mobile autonomous robot for indoor security applications," Proc. Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture, vol. 218 (5), pp. 533-543.
- Carnegie, D. A., Cree, M. J., Dorrington, A. A. and Payne, A. D., (2005). "A High Resolution Full-Field Range Imaging System for Robotic Devices," *Proc. IASTED International Conf. on Robotics and Applications*, pp. 306-311.
- Carnegie, D. A., Cree, M. J. and Dorrington, A. A., (2005). "A high-resolution full-field range imaging system," *Review of Scientific Instruments*, vol. 76 (8), pp. 083702.
- Christie, S., Hill, S. L., Bury, B., Gray, J. O. and Booth, K. M., (1995). "Design and Development of a Multi-Detecting 2-Dimensional Ranging Sensor," *Measurement Science & Technology*, vol. 6 (9), pp. 1301-1308.
- Cope, B., Cheung, P. Y. K., Luk, W. and Witt, S., (2005). "Have GPUs made FPGAs redundant in the field of Video Processing?," *Proc. IEEE International Conf. on Field-Programmable Technology*, pp. 111-118.
- DALSA Corp., (2000), "Camera Link, DALSA Implementation Road Map," www.dalsa.com
- DALSA Corp., (2004). "Pantera TF 1M60 and 1M30 User's Manual and Reference," www.dalsa.com
- DAVICOM Semiconductor Inc., (2006). "DM9000A Ethernet Controller with General Purpose Interface," Technical Datasheet, www.davicom.com.tw

- Dorrington, A. A., Cree, M. J., Payne, A. D., Conroy, R. M. and Carnegie, D. A., (2007a). "Achieving sub-millimetre precision with a solid-state full-field heterodyning range imaging camera," *Measurement Science & Technology*, vol. 18, pp. 2809-2816.
- Dorrington, A. A., Cree, M. J., Carnegie, D. A., Payne, A. D., Conroy, R. M., Godbaz, J. P. and Jongenelen, A. P. P., (2008). "Video-rate or highprecision: a flexible range imaging camera," *Proc. SPIE Image Processing: Machine Vision Applications*, vol. 6813, pp. 681307.
- Dorrington, A. A., Payne, A. D. and Cree, M. J., (2010). "An Evaluation of Timeof-Flight Range Cameras for Close Range Metrology Applications," *Proc. ISPRS Commission V Mid-term Symposium: Close Range Image Measurement Techniques*, pp. 201-206.
- Frank, M., Plaue, M., Rapp, H., Kothe, U., Jahne, B. and Hamprecht, F. A., (2009). "Theoretical and experimental error analysis of continuous-wave time-of-flight range cameras," *Optical Engineering*, vol. 48 (1), pp. 013602.
- Fuchs, S. and Hirzinger, G., (2008). "Extrinsic and Depth Calibration of ToFcameras," *IEEE Conf. on Computer Vision and Pattern Recognition*, pp. 1-6.
- Godbaz, J. P., Cree, M. J., Dorrington, A. A., and Kunnemeyer, R., (2007). "Defocus restoration for a full-field heterodyne ranger via multiple return separation," *Proc. Image and Vision Computing New Zealand*, pp. 52-57.
- Godbaz, J. P., Cree, M. J. and Dorrington, A. A., (2009a). "Undue Influence: Mitigating Range-Intensity Coupling in AMCW 'Flash' Lidar using Scene Texture," Proc. 24th International Conf. Image and Vision Computing New Zealand, pp. 304-309.
- Godbaz, J. P., Cree, M. J. and Dorrington, A. A., (2009b). "Multiple Return Separation for a Full-Field Ranger Via Continuous Waveform Modelling," *Image Processing: Machine Vision Applications II*, vol. 7251, pp 72510.
- Gokturk, S. B., Yalcin, H. and Bamji, C., (2004). "A Time-Of-Flight Depth Sensor – System Description, Issues and Solutions," Proc. 2004 IEEE Conf. Computer Vision and Pattern Recognition Workshop, vol. 3, pp. 35.
- Goyo Optical Inc., (2003). "GM31614MCN Manual Industrial Lens," Technical Datasheet, www.goyooptical.com
- Gulden, P., Vossiek, M., Heide, P. and Schwarte, R., (2002). "Novel opportunities for optical level gauging and 3-D-imaging with the photoelectronic mixing device," *IEEE Transactions on Instrumentation and Measurement*, vol. 51 (4), pp. 679-684.

- Guo, Z., Najjar, W., Vahid, F. and Vissers, K., (2004). "A Quantitative Analysis of the Speedup Factors of FPGAs over Processors," *Proc. ACM/SIGDA* 12th International Symposium on Field programmable gate arrays, pp. 162-170.
- Hariharan, P., (1985). Optical Interferometry, Academic Press, Sydney, Australia.
- Hebert, M., (2000). "Active and Passive Range Sensing for Robotics," *Proc. IEEE International Conf. on Robotics and Automation*, vol. 1, pp. 102-110.
- Hirose, M., Furuhashi, H. and Araki, K., (2001). "Automatic Registration of Multi-view Range Images Measured Under Free Condition," Proc. Seventh International Conf. on Virtual Systems and Multimedia, 738-746.
- Hsu, S., Acharya, S., Rafii, A. and New, R., (2006). "Performance of a Time-of-Flight Range Camera for Intelligent Vehicle Safety Applications," *Advanced Microsystems for Automotive Applications*, pp. 205-214.
- Hussmann, S., Schauer, D. and MacDonald, B., (2009). "Integration of a 3D-TOF camera into an autonomous, mobile robot system," *IEEE Instrumentation and Measurement Technology Conf.*, pp. 547-552.
- iC-Haus, (2004). "iC-HK, iC-HKB 155 MHz Laser Switch," Technical Datasheet, www.ichaus.com
- Intersil Americas Inc., (2007). "EL7158 Ultra-High Current Pin Driver," Technical Datasheet, www.intersil.com
- Jelalian, A. V., (1992). Laser Radar Systems, Artech House, Boston, USA.
- Jongenelen, A. P. P., Payne, A. D., Carnegie, D. A. and Dorrington, A. A., (2008). "Heterodyne Range Imaging in Real-time," *Proc. International Conf. on Sensing Technology*, pp. 57-62.
- Jongenelen, A. P. P., Payne, A. D., Carnegie, D. A., Dorrington, A. A. and Cree, M. J., (2009a). "Development of a Real-Time Full-Field Range Imaging System," *Recent Advances in Sensing Technology, Lecture Notes on Electrical Engineering*, pp. 113-129.
- Jongenelen, A. P. P., Carnegie, D. A., Payne, A. D. and Dorrington, A. A., (2009b). "Development and Characterisation of an Easily Configurable Range Imaging System," *Proc. Image and Vision Computing New Zealand*, pp. 79-84.
- Jongenelen, A. P. P., Carnegie, D. A., Payne, A. D. and Dorrington, A. A., (2010a). "Maximizing Precision Over Extended Unambiguous Range for TOF Range Imaging Systems," *Proc. International Instrumentation and Measurement Technology*, pp. 1575-1580.

- Jongenelen, A. P. P., Bailey, D. G., Payne, A. D., Carnegie, D. A. and Dorrington, A. A., (2010b). "Efficient FPGA Implementation of Homodyne-Based Time of Flight Range Imaging," *Journal of Real-time Image Processing*, Springer-Verlag, published online.
- Jongenelen, A. P. P., Bailey, D. G., Payne, A. D., Dorrington, A. A. and Carnegie, D. A., (2011). "Analysis of Errors in ToF Range Imaging with Dual-Frequency Modulation," *IEEE Transactions on Instrumentation and Measurement*, vol. 60 (5), pp. 1861 – 1868.
- Kahlmann, T., Remondino, F. and Ingensand, H., (2006). "Calibration for Increased Accuracy of the Range Imaging Camera SwissRanger," Proc. ISPRS Commission V Symposium - Image Engineering and Vision Metrology, pp. 136-141.
- Kimura, S., Kano, H., Kanade, T., Yoshida, A., Kawamura, E. and Oda, K., (1995). "CMU Video-Rate Stereo Machine," *Mobile Mapping Symposium ASPRS*, pp. 9-18.
- Lange, R. and Seitz, P., (2000). "Seeing distances a fast time-of-flight 3D camera," *Sensor Review*, vol. 20 (3), pp. 212-217.
- Lange, R. and Seitz, P., (2001). "Solid-state time-of-flight range camera," *IEEE Journal of Quantum Electronics*, vol. 37 (3), pp. 390-397.
- Lange, R., Seitz, P., Biber, A. and Lauxtermann, S., (2000). "Demodulation Pixels in CCD and CMOS Technologies for Time-of-Flight Ranging," Proc. SPIE - Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications, vol. 3965, pp. 177-189.
- Lange, R., (2000). "3D Time-of-flight distance measurement with custom solidstate image sensors in CMOS/CCD-technology," *Ph.D. dissertation, University of Siegen.*
- Leica Geosystems, (2007). "Leica DISTO D3 User Manual English," www.leica.com
- Lidner M. and Kolb, A., (2007). "Calibration of the Intensity-Related Distance Error of the PMD TOF-Camera," *SPIE: Intelligent Robots and Computer Vision* XXV, vol. 6764, pp. 676435.
- Light, R., (2008). "Design and Implementation of an Integrating Modulated Light Camera," *Ph.D. thesis, University of Nottingham.*
- Luan, X., (2001). "Experimental Investigation of Photonic Mixer Device and Development of TOF 3D Ranging Systems Based on PMD Technology," Doctor of Technical Sciences dissertation, University of Siegen.

- Microchip Technology Inc., (2004). "PIC16F684 14-Pin Flash-Based, 8-Bit CMOS Microcontrollers with nanoWatt Technology," Technical Datasheet, www.microchip.com
- Monson, T. C., Grantham, J. W., Childress, S. W., Sackos, J. T., Nellums, R. O. and Lebien, S. M., (1999). "Characterization of scannerless ladar," *Proc. SPIE Laser Radar Technology and Applications IV*, vol. 3707, pp. 409-420.
- Mure-Dubois, J. and Hugli, H., (2007). "Time-of-Flight Imaging of Indoor Scenes with Scattering Compensation," *Proc. Conf. on Optical 3-D Measurement Techniques*, pp. 117-123.
- National Semiconductor, (2003). "DS90LV047A 3V LVDS Quad CMOS Differential Line Driver," Technical Datasheet, www.national.com
- National Semiconductor, (2004). "DS90CR286A/DS90CR216A +3.3V Rising Edge Data Strobe LVDS Receiver 28-Bit Channel Link—66 MHz," Technical Datasheet, www.national.com
- Niclass, C., Favi, C., Kluter, T., Monnier, F. and Charbon, E, (2009). "Single-Photon Synchronous Detection," *IEEE Journal of Solid-State Circuits*, vol. 44 (7), pp.1977-1989.
- ON Semiconductor, (2004). "MC74VHC1GT14 Schmitt-Trigger Inverter / CMOS Logic Level Shifter," Technical Datasheet, onsemi.com
- Payne, A. D., Carnegie, D. A., Dorrington, A. A. and Cree, M. J., (2005). "A Synchronised Direct Digital Synthesiser," *Proc. First International Conf. on Sensing Technology*, pp. 174-179.
- Payne, A. D., Dorrington, A. A., Cree, M. J. and Carnegie, D. A., (2008). "Improved Linearity Using Harmonic Error Rejection in a Full-Field Range Imaging System," *Proc. SPIE - 3D Image Capture and Applications VII*, vol. 6805, pp. 68050D.
- Payne, A. D., Dorrington, A. A., Cree, M. J. and Carnegie, D. A., (2009a). "Characterization of modulated time-of-flight range image sensors," *Proc. SPIE-IS&T Electronic Imaging*, vol. 7239, pp. 723904.
- Payne, A. D., Jongenelen, A. P. P., Dorrington, A. A., Cree, M. J. and Carnegie, D. A., (2009b). "Multiple Frequency Range Imaging to Remove Measurement Ambiguity," Proc. 9th Conf. on Optical 3-D Measurement Techniques, pp. 139-148.
- Payne, A. D., Dorrington, A. A., Cree, M. J. and Carnegie, D. A (2010). "Improved measurement linearity and precision for AMCW time-offlight range imaging cameras," *Applied Optics*, vol. 49 (23), pp. 4392-4403.

- Photek Ltd., (2001). "MCP Image Intensifiers," Technical Datasheet, www.photek.com
- PMD Technologies, (2008). "PMD19k datasheet," Technical Datesheet, www.pmdtec.com
- PULNiX America Inc., (2000). "Specifications of the Camera Link Interface Standard for Digital Cameras and Frame Grabbers," www.imagelabs.com
- Rapp, H., (2007). "Experimental and Theoretical Investigation of Correlating TOF-Camera Systems," *Diploma Thesis in Physics, University of Heidelberg.*
- Ringbeck, T., Moller, T. and Hagebeuker, B., (2007). "Multidimensional measurement by using 3-D PMD sensors," *Advances in Radio Science*, vol. 5, pp. 135-146.
- Salvi, J., Pages, J. and Batlle, J., (2004). "Pattern codification strategies in structured light systems," *Pattern Recognition*, vol. 37 (4), pp. 827-849.
- Sato, K., (1996). "Range imaging based on moving pattern light and spatiotemporal matched filter," *International Conf. on Image Processing*, vol. 1, pp. 33-36.
- Tektronix Inc., (2005). "Digital Phosphor Oscilloscopes DPO4000 Series," Technical Datasheet, www.tektronix.com/oscilloscopes
- Texas Instruments, (2001). "SN65LVDS387, SN75LVDS387, SN65LVDS389 SN75LVDS389, SN65LVDS391, SN75LVDS391 High-speed Differential Line Drivers," Technical Datasheet, www.ti.com
- Texas Instruments, (2002). "SN65LVDS1, SN65LVDS2, SN65LVDT2 Highspeed Differential Line Drivers/Receivers," Technical Datasheet, www.ti.com
- Texas Instruments, (2003). "DS90LV047A 3V LVDS Quad CMOS Differential Line Driver," Technical Datasheet, www.ti.com
- Texas Instruments, (2004). "SN65LVDS100, SN65LVDT100, SN65LVDS101, SN65LVDT101 High-speed Differential Line Drivers/Receivers," Technical Datasheet, www.ti.com
- Texas Instruments, (2007a). "SN74ALVC125 Quadruple Bus Buffer Gate with 3-State Outputs," Technical Datasheet, www.ti.com
- Texas Instruments, (2007b). "SN65LVDS386/ 388A/ 390, SN65LVDT386/ 388A/ 390, SN75LVDS386/ 388A/ 390, SN75LVDT386/ 388A/ 390 High-speed Differential Line Receivers," Technical Datasheet, www.ti.com

- Thorlabs Inc., (2006). "DET10A Operating Manual High Speed Silicon Detector," Technical Datasheet, www.thorlabs.com
- Underwood, K., (2004). "FPGAs vs. CPUs: Trends in Peak Floating-Point Performance," *Proc. ACM/SIGDA* 12th International symposium on Field programmable gate arrays, pp 171-180.
- Wang, Y., (2000). "Residue-to-binary converters based on New Chinese Remainder Theorems," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 47 (3), pp. 197-205.
- Walther, J. S., (1971). "A unified algorithm for elementary functions," *AFIPS Conf. Proc.*, vol. 38, pp. 379-385.
- Yahav, G., Iddan, G. J. and Mandelboum, D, (2007). "3D Imaging Camera for Gaming Application," *International Conf. on Consumer Electronics Digest of Technical Papers*, pp. 1-2.

Appendix A PCB Schematics

A.1 PMD Daughterboard









A.3 PMD Mainboard 2 – Connectors and Interface

A.4 Illumination Board – Top





A.5 Illumination Board – Laser Diode Module

Note: The AGND resistors depend on the type of diode being used, as they set the on/off currents through the diodes.

ML101J27 red laser diode - AGND1 = 0, AGND2 = 15

L980P100 IR laser diode - AGND1 = 1.5, AGND2 = 8.25

A.6 Ethernet Interface Board





A.7 VGA Interface Board



A.8 Camera Link Interface Board

Appendix B DVD Contents

DVD 1 – Range Imaging Data

Contains raw and processed range imaging data used for the results in this thesis. Each subfolder has data relating to one capturing session, with '.set' storing operating parameters for each capture. Subfolders include:

- Co250_Feb10_overnight dual-frequency results;
- Lb310_Jul22_cap_on pixel noise floor results;
- Lb310_Jul23_box25_fullscreen displaying box scene;
- Lb310_Jul28_box25_overnight most characterisation results;
- Lb310_Jul30_Box25 some more characterisation results;
- rebecca15 images of the mannequin head;
- videoRate range video of person.
DVD 2 – Source code and Documentation

Contains C, Java, MATLAB and VHDL source code as well documentation in the form of publications and presentations. Also has various other results, for example the video produced for Chapter 7, and oscilloscope screenshots for Chapter 4. Subfolders include:

- Altera project files for Stratix III FPGA;
- Documents publications resulting from this research;
 - Thesis contains this thesis and all images;
- Java PC Ethernet listener applications;
- Illumination Test oscilloscope screenshots;
- MATLAB scripts for algorithm prototyping and data analysis;
 - Thesis scripts for generating thesis results;
- Multifreq compiled data for Chapter 5 results;
- Nios C code for Nios II application;
- Presentations powerpoint slides;
- VHDL FPGA programming source code.

The source code is also freely available online at:

http://github.com/adrianj/RangeImaging